



Complete Software Guide for Junos OS for QFX10000 Switches, Release 15.1X53-D30



Modified: 2015-11-20

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Complete Software Guide for Junos OS for QFX10000 Switches, Release 15.1X53-D30
Copyright © 2016, Juniper Networks, Inc.
All rights reserved.

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at <http://www.juniper.net/support/eula.html>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

About the Documentation	cxxxix
Documentation and Release Notes	cxxxix
Using the Examples in This Manual	cxxxix
Merging a Full Example	cxl
Merging a Snippet	cxl
Documentation Conventions	cxli
Documentation Feedback	cxliii
Requesting Technical Support	cxliii
Self-Help Online Tools and Resources	cxliii
Opening a Case with JTAC	cxliv

Guide 1	Junos OS Basics Feature Guide for QFX10000 Switches
----------------	--

Part 1	Autoinstallation
Chapter 1	Understanding Autoinstallation 3
	Understanding Autoinstallation of Configuration Files 3
	Typical Uses for Autoinstallation 3
	Autoinstallation Configuration Files and IP Addresses 4
	Typical Autoinstallation Process on a New Switch 4
	Configuring Autoinstallation of Configuration Files (CLI Procedure) 5
	Upgrading Software by Using Automatic Software Download 7
	Verifying Autoinstallation Status 8
	Verifying That Automatic Software Download Is Working Correctly 9
Part 2	Basic System Management
Chapter 2	Understanding Basic System Management 13
	Compressing the Current Configuration File 14
	Configuring a DNS Name Server for Resolving a Hostname into Addresses 15
	Configuring Console and Auxiliary Port Properties 15
	Configuring the Hostname of a Router or Switch by Using a Configuration Group 16
	Configuring the Junos OS to Determine Conditions That Trigger Alarms on Different Interface Types 18
	Configuring Junos OS to Disable Protocol Redirect Messages on the Router or Switch 18
	Configuring Junos OS to Disable the Reporting of IP Address and Timestamps in Ping Responses 19
	Configuring Junos OS to Extend the Default Port Address Range 19

Configuring the Junos OS to Select a Fixed Source Address for Locally Generated TCP/IP Packets	20
Configuring System Alarms to Appear Automatically Upon Login	20
Configuring Time-Based User Access	20
Configuring the Timeout Value for Idle Login Sessions	22
Including the Year or Millisecond in Timestamps	22
Mapping the Hostname of the Switch to IP Addresses	23
Modifying the Default Time Zone for a Router or Switch Running Junos OS	24
Rebooting and Halting a Device	25
Reaching a Domain Name System Server	26
Reverting to the Default Factory Configuration by Using the request system zeroize Command	28
Saving Core Files Generated by Junos OS Processes	29
Specifying the Physical Location of the Switch	29
Specifying Access Privileges for Junos OS Operational Mode Commands	30
Updating the IANA Time Zone Database on Junos OS Devices	32
Importing and Installing Time Zone Files	32
Configuring a Custom Time Zone	33
Viewing Core Files from Junos OS Processes	34
Example: Configuring the Name of the Switch, IP Address, and System ID	34

Part 3

Chapter 3

Command Line Interface (CLI)

Understanding the CLI	37
CLI User Interface Overview	37
CLI Overview	37
CLI Key Features	38
CLI Command Modes	38
Configuring Login Tips	40
Format for Specifying Filenames and URLs in Junos OS CLI Commands	40
Getting Started with Enhanced Layer 2 Software	41
Understanding Enhanced Layer 2 Software Support	42
Using the ELS Translator	42
Configuring a VLAN	43
Configuring the Native VLAN Identifier	44
Configuring Layer 2 Interfaces	44
Configuring Layer 3 Interfaces	45
Configuring an IRB Interface	45
Configuring an Aggregated Ethernet Interface and Configuring LACP on That Interface	46
Enhanced Layer 2 CLI Configuration Statement and Command Changes	47
Changes to the ethernet-switching-options Hierarchy Level	47
Changes to the Port Mirroring Hierarchy Level	49
Changes to the Layer 2 Control Protocol Hierarchy Level	49
Changes to the dot1q-tunneling Statement	50
Changes to the L2 Learning Protocol	50
Changes to Nonstop Bridging	50
Changes to Port Security and DHCP Snooping	51
Changes to Configuring VLANs	52

	Changes to Storm Control Profiles	55
	Changes to the Interfaces Hierarchy	56
	Changes to IGMP Snooping	57
	Junos OS Operational Mode Commands That Combine Other Commands	57
	Overview of Junos OS CLI Operational Mode Commands	58
	CLI Command Categories	58
	Commonly Used Operational Mode Commands	59
	Overview of Navigating the CLI	61
	CLI Command Hierarchy	61
	CLI Configuration Statements	62
	Moving Among Hierarchy Levels	62
	Understanding the Brief, Detail, Extensive, and Terse Options of Junos OS	
	Operational Commands	62
	Understanding Junos OS CLI Configuration Mode	63
	Configuration Mode Commands	65
	Configuration Statements and Identifiers	66
	Configuration Statement Hierarchy	68
Part 4	Dynamic Host Control Protocol (DHCP)	
Chapter 4	Understanding DHCP	73
	Understanding DHCP Services for Switches	73
	DHCP Client/Server Model	73
	Using DHCP	74
	DHCP Relay Servers and DHCP Servers	74
	Legacy DHCP and Extended DHCP for Server Versions	75
	Configuring DHCP on a Switch	76
	How DHCP Works	76
	Configuring a DHCP Client (CLI Procedure)	77
	Configuring a DHCP Server on Switches (CLI Procedure)	78
	Configuring an Extended DHCP Server on a Switch	79
	Configuring a Legacy DHCP Server on a Switch (CLI Procedure)	79
Part 5	Internet Control Message Protocol (ICMP)	
Chapter 5	Understanding ICMP	85
	Configuring Junos OS to Disable the Reporting of IP Address and Timestamps	
	in Ping Responses	85
	Configuring Junos OS to Disable the Routing Engine Response to Multicast Ping	
	Packets	86
	Configuring the Junos OS ICMPv4 Rate Limit for ICMPv4 Routing Engine	
	Messages	86
Part 6	Junos OS With Upgraded FreeBSD	
Chapter 6	Understanding Junos OS With Upgraded FreeBSD	89
	Understanding Junos OS with Upgraded FreeBSD	89
	Understanding Junos OS with Upgraded FreeBSD Package Names	92
	Understanding Junos OS with Upgraded FreeBSD Disk Volumes	94

	Understanding Junos OS with Upgraded FreeBSD Snapshots	95
	Upgrading Junos OS with Upgraded FreeBSD	96
	To Install Junos OS with Upgraded FreeBSD Over a Plain Junos OS	98
	To Install Junos OS with Upgraded FreeBSD Over Junos OS with Upgraded FreeBSD of an Earlier Release	101
	To Install Junos OS with Upgraded FreeBSD Over Junos OS with Upgraded FreeBSD of a Later Release	102
	Downgrading Junos OS from Upgraded FreeBSD	102
Part 7	Licenses	
Chapter 7	Using Licenses	107
	Junos OS Feature Licenses	107
	Software Features That Require Licenses on the QFX Series	108
	Junos OS Feature License Keys	112
	Release-Tied License Keys and Upgrade Licenses on MX Series Routers	112
	Licensable Ports on MX5, MX10, and MX40 Routers	113
	Port Activation on MX104 Routers	114
	Generating License Keys	115
	Adding New Licenses (CLI Procedure)	117
	Deleting a License (CLI Procedure)	118
	Saving License Keys	119
	Verifying Junos OS License Installation	120
	Displaying Installed Licenses	120
	Displaying License Usage	121
Part 8	Login Classes	
Chapter 8	Understanding Login Classes	125
	Configuring Junos OS User Accounts	125
	Configuring the Junos OS to Display a System Login Announcement	126
	Configuring Junos OS to Display a System Login Message	126
Part 9	Network Time Protocol (NTP)	
Chapter 9	Understanding NTP	131
	Understanding NTP Time Servers	131
	Configuring NTP Authentication Keys	132
	Configuring the NTP Time Server and Time Services	133
	Configuring the Switch to Operate in Client Mode	133
	Configuring the Router or Switch to Operate in Symmetric Active Mode	134
	Configuring the Router or Switch to Operate in Broadcast Mode	134
	Configuring the Router or Switch to Operate in Server Mode	135
	Configuring the Switch to Listen for Broadcast Messages Using NTP	135
	Configuring the Switch to Listen for Multicast Messages Using NTP	136
	Setting the Date and Time	136
	Synchronizing and Coordinating Time Distribution Using NTP	137
	Configuring NTP	137
	Configuring the NTP Boot Server	137

	Specifying a Source Address for an NTP Server	138
	Example: Configuring NTP	139
	Example: Configuring NTP as a Single Time Source for Router and Switch Clock Synchronization	142
Part 10	Passwords	
Chapter 10	Understanding Passwords	145
	Configuring the Root Password	145
	Recovering the Root Password	147
	Example: Changing the Requirements for Junos OS Plain-Text Passwords	148
Chapter 11	Precision Time Protocol (PTP)	151
	Understanding Transparent Clocks in Precision Time Protocol	152
	Configuring Transparent Clock Mode for Precision Time Protocol	152
Part 11	Recovery Installation	
Chapter 12	Understanding Recovery Installation	157
	Creating an Emergency Boot Device	157
	Performing a Recovery Installation Using an Emergency Boot Device	159
	Recovering from a Failed Software Installation	160
Part 12	Routine Monitoring	
Chapter 13	Understanding Routine Monitoring	165
	Monitoring System Process Information	165
	Monitoring System Properties	166
	Monitoring Interface Status and Traffic	167
	Other Tools to Configure and Monitor Devices Running Junos OS	168
	Fabric Management Overview	168
Part 13	Standard Software Installation and Upgrade	
Chapter 14	Understanding Standard Software Installation and Upgrade	173
	Junos OS Package Names	173
	Software Installation Overview	174
	Creating an Emergency Boot Device	175
	Installing and Recovering Software Using the Open Network Install Environment (ONIE)	176
	Understanding the Open Network Install Environment	178
	Downloading Software Files with a Browser	178
	Connecting to the Console Port	179
	Backing Up the Current Configuration Files	179
	Uninstalling the Existing Version of Junos OS	180
	Installing a Junos OS Software Package That Resides on a Webserver or DHCP Server with DHCP Options Configured	180
	Installing Junos OS Software Using Secure Copy Protocol (SCP)	181
	Installing Junos OS Software Using FTP or TFTP Without a Webserver	182

	Installing Junos OS Software Using DHCP Server with No DHCP Options Configured	182
	Installing Junos OS Software Using Webserver Without DHCP Configured	184
	Installing Junos OS Software Using USB Media	185
	Verifying Software Installation	185
	Troubleshooting Unsupported Network Operating System	187
	Troubleshooting ONIE ISO Files	187
	Creating an Emergency Boot Device for ONIE Software and Network Operating System	187
	Performing a Recovery Installation	188
	Upgrading Software	189
	Downloading Software Files with a Browser	190
	Accessing Software Downloaded to a Remote Location	191
	Connecting to the Console Port	191
	Backing Up the Current Configuration Files	191
	Installing a Standard Software Package	191
	Installing a Standard Software Package on a QFX10008 Switch	192
Part 14	Zero Touch Provisioning (ZTP)	
Chapter 15	Understanding ZTP	197
	Understanding Zero Touch Provisioning	197
	Understanding Zero Touch Provisioning	197
	Zero Touch Provisioning Process	199
	Zero Touch Provisioning Restart Process Triggers	202
	Configuring Zero Touch Provisioning	204
	Monitoring Zero Touch Provisioning	209
	Using the Console to Monitor Zero Touch Provisioning	209
	Using System Log Alerts to Monitor Zero Touch Provisioning	210
	Using Error Messages to Monitor Zero Touch Provisioning	210
	Using System Log Files to Monitor Zero Touch Provisioning	210
	Using the show dhcp client binding Command	211
	Using the show dhcp client statistics Command	211
Part 15	Configuration Statements and Operational Commands	
Chapter 16	Basic System Management Configuration Statements	215
	arp (System)	217
	auxiliary	218
	checksum	219
	compress-configuration-files (System)	220
	console (Physical Port)	221
	default-address-selection	222
	domain-name	223
	domain-search	223
	ethernet (Alarm)	224
	host-name	224
	icmpv4-rate-limit	225

	internet-options	226
	link-down	227
	location	228
	login-alarms	229
	management-ethernet (Alarm)	230
	max-configurations-on-flash	231
	name-server	231
	no-ping-record-route	232
	no-ping-time-stamp	232
	no-redirects (IPv4 Traffic)	233
	optional	234
	ports	234
	retry	235
	saved-core-context	236
	saved-core-files	236
	source-port (Port Addresses)	237
	static-host-mapping	238
	system	239
	time-format	244
	time-zone	245
	traceoptions (Layer 2 Learning)	247
	uid	249
	use-imported-time-zones	249
Chapter 17	DHCP Configuration Statements	251
	client-identifier (DHCP Client)	251
	lease-time	252
	retransmission-attempt	253
	retransmission-interval	254
	server-address	255
	update-server	256
	vendor-option	257
Chapter 18	ICMP Configuration Statement	259
	no-multicast-echo	259
Chapter 19	Junos OS With Upgraded FreeBSD Operational Commands	261
	request system reboot (Junos OS with Upgraded FreeBSD)	262
	request system snapshot (Junos OS with Upgraded FreeBSD)	265
	show system snapshot (Junos OS with Upgraded FreeBSD)	267
Chapter 20	Login Classes Configuration Statements	269
	access-end	270
	access-start	270
	allow-commands	271
	allow-configuration	272
	allowed-days	272
	announcement	273
	authentication (Login)	274
	class (Defining Login Classes)	275

	class (Assigning a Class to an Individual User)	276
	deny-commands	276
	deny-configuration	277
	idle-timeout	278
	load-key-file	279
	login	280
	login-tip	281
	message	281
	user (Access)	282
Chapter 21	NTP Configuration Statements	283
	authentication-key	284
	boot-server (NTP)	285
	broadcast	286
	broadcast-client	287
	multicast-client	287
	ntp	288
	peer	289
	server (NTP)	290
	source-address (NTP, RADIUS, System Logging, or TACACS+)	291
	trusted-key	292
Chapter 22	Password Configuration Statements	293
	authentication (Login)	294
	change-type	295
	format	296
	maximum-length	297
	minimum-changes	298
	minimum-length	299
	minimum-lower-cases	300
	minimum-numeric	301
	minimum-punctuations	302
	minimum-upper-cases	303
	password (Login)	304
	permissions	304
	retry-options	305
	root-authentication	307
	ssh-dsa	308
	ssh-rsa	308
	user (Access)	309
Chapter 23	PTP Configuration Statement	311
	e2e-transparent	311
Chapter 24	RADIUS and TACACS+ Accounting Configuration Statements	313
	accounting-port (RADIUS Server)	314
	destination (Accounting)	315
	events	316
	port (TACACS+ Server)	316
	radius (System)	317

	secret	318
	server (RADIUS Accounting)	319
	server (TACACS+ Accounting)	319
	single-connection	320
	tacplus	320
	timeout	321
Chapter 25	Autoinstallation Operational Commands	323
	show system autoinstallation status	324
Chapter 26	Basic System Management Operational Commands	327
	commit	329
	clear log	332
	clear chassis display message	333
	clear system commit	336
	clear system reboot	337
	file	341
	file archive	342
	file checksum md5	344
	file checksum sha1	345
	file checksum sha-256	346
	file compare	347
	file delete	351
	file list	352
	file rename	354
	file show	356
	load	358
	ping	360
	request chassis beacon	364
	request chassis fpc	366
	request chassis pic	371
	request chassis routing-engine master	376
	request message	381
	request system configuration rescue delete	382
	request system configuration rescue save	383
	request system halt	384
	request system logout	390
	request system power-off	391
	request system reboot	396
	request system storage cleanup	399
	request system zeroize	409
	restart	414
	save	425
Chapter 27	CLI Operational Commands	427
	show cli	428
	show cli authorization	430
	show cli directory	434
	show cli history	435

Chapter 28	Licensing Operational Commands	437
	request system license add	438
	request system license delete	439
	request system license save	440
	show system license	441
Chapter 29	NTP Operational Commands	449
	show ntp associations	450
	show ntp status	452
Chapter 30	PTP Operational Command	455
	show ptp global-information	456
Chapter 31	Routine Monitoring Operational Commands	457
	show chassis alarms	459
	show chassis beacon	474
	show chassis environment	476
	show chassis environment fpc	540
	show chassis environment pem	566
	show chassis environment routing-engine	575
	show chassis fabric errors	580
	show chassis fabric fpcs	585
	show chassis fabric plane-location	624
	show chassis fabric sibs	629
	show chassis fabric summary	641
	show chassis fabric topology	647
	show chassis fan	678
	show chassis firmware	692
	show chassis fpc	703
	show chassis hardware	743
	show chassis lcd	921
	show chassis led	935
	show chassis location	945
	show chassis mac-addresses	949
	show chassis pic	954
	show chassis routing-engine	970
	show chassis zones	996
	show host	1002
	show log	1003
	show subscribers	1007
	show system alarms	1027
	show system audit	1030
	show system boot-messages	1038
	show system buffers	1045
	show system certificate	1052
	show system commit	1054
	show system configuration archival	1057
	show system configuration rescue	1058
	show system connections	1060
	show system core-dumps	1079

	show system directory-usage	1093
	show system processes	1097
	show system reboot	1126
	show system resource-cleanup processes	1129
	show system services service-deployment	1131
	show system software	1132
	show system statistics	1140
	show system storage	1177
	show system uptime	1185
	show system users	1190
	show system virtual-memory	1195
	show version	1253
	start shell	1266
	test configuration	1268
	traceroute	1269
	traceroute monitor	1273
Chapter 32	Standard Software Installation and Upgrade Operational Commands . .	1275
	request system software add	1276
	request system software delete	1286
	request system software download	1290
	request system software rollback	1292
	request system software validate	1297
	rollback	1301
	show system rollback	1302
Guide 2	Configuration and File Management Feature Guide for QFX10000 Switches	
Part 16	Configuration and File Management Overview	
Chapter 33	Understanding Configuration Files and the configure Command	1309
	Configuration File Terms	1309
	Understanding Configuration Files	1310
	Forms of the configure Command	1310
Chapter 34	Understanding How Configuration Files are Committed and Stored	1313
	Junos OS Commit Model for Router or Switch Configuration	1313
	Understanding How the Junos OS Configuration Is Stored	1314
Part 17	Managing Configuration Files	
Chapter 35	Using the Default Configuration	1319
	QFX3500 Switch Default Configuration	1319
Chapter 36	Loading, Uploading, Transferring, and Saving Configuration Files	1325
	Loading a Configuration from a File	1325
	Examples: Loading a Configuration from a File	1328
	Loading a Previous Configuration File	1330
	Saving a Configuration to a File	1331

	Uploading a Configuration File	1332
	Using Junos OS to Configure a Router or Switch to Transfer Its Configuration to an Archive Site	1333
	Configuring the Transfer of the Currently Active Configuration to an Archive Site	1333
	Configuring the Periodic Transfer of the Active Configuration to an Archive Site	1334
	Configuring the Transfer of the Currently Active Configuration When a Configuration Is Committed	1334
	Configuring Archive Sites for the Transfer of Active Configuration Files	1334
Chapter 37	Reverting a Configuration or Rolling Back to the Factory Default	1337
	Rolling Back Junos OS Configuration Changes	1337
	Returning to the Most Recently Committed Junos OS Configuration	1338
	Returning to a Previously Committed Junos OS Configuration	1339
	Returning to a Configuration Prior to the One Most Recently Committed	1339
	Displaying Previous Configurations	1339
	Comparing Configuration Changes with a Prior Version	1340
	Creating and Returning to a Rescue Configuration	1342
	Saving a Configuration to a File	1343
	Reverting to the Default Factory Configuration	1344
Chapter 38	Using Rescue Configurations	1345
	Setting or Deleting the Rescue Configuration	1345
	Creating and Returning to a Rescue Configuration	1345
	Reverting to the Rescue Configuration	1346
Chapter 39	Comparing or Compressing Configuration Files	1349
	Comparing Configuration Changes with a Prior Version	1349
	Compressing the Current Configuration File	1351
Chapter 40	Troubleshooting Procedures	1353
	Loading a Previous Configuration File	1353
	Reverting to the Default Factory Configuration	1354
	Reverting to the Rescue Configuration	1354
	Cleaning Up the System File Storage Space	1355
Part 18	Configuration Statements and Operational Commands	
Chapter 41	Configuration Statements	1359
	archival	1360
	archive-sites (Configuration File)	1361
	configuration	1363
	transfer-interval (Configuration)	1364
	transfer-on-commit	1365
Chapter 42	Operational Commands	1367
	clear log	1368
	clear system commit	1369
	file archive	1370
	file checksum md5	1372

file checksum sha1	1373
file checksum sha-256	1374
file compare	1375
file delete	1379
file list	1380
file rename	1382
file show	1384
request system configuration rescue delete	1386
request system configuration rescue save	1387
show system commit	1388
show system configuration archival	1391
show system configuration rescue	1392
show system rollback	1394
test configuration	1396

Guide 3 Network Management and Monitoring Feature Guide for QFX10000 Switches

Part 19	Overview	
Chapter 43	Understanding Network Management	1401
	Understanding Device and Network Management Features	1401
Chapter 44	Junos Space	1405
	Understanding Junos Space Support	1405
Part 20	Configuration	
Chapter 45	Configuring Network Management	1409
	Understanding Tracing and Logging Operations	1409
	Configuring Console and Auxiliary Port Properties	1411
	Configuring SSH Service for Remote Access to the Router or Switch	1412
	Configuring the Root Login Through SSH	1413
	Configuring the SSH Protocol Version	1413
	Configuring the Client Alive Mechanism	1413
	Configuring Telnet Service for Remote Access to a Switch	1414
	Pinging Hosts	1414
	Monitoring Traffic Through the Router or Switch	1415
	Displaying Real-Time Statistics About All Interfaces on the Router or Switch	1416
	Displaying Real-Time Statistics About an Interface on the Router or Switch	1416
Chapter 46	Configuring Automation	1419
	Overview of Junos Automation Enhancements	1419
	Features of the Junos Automation Enhancements	1419
	Overview of Python with Junos Automation Enhancements	1420
	Understanding Automation Scripts Support	1423

	How Commit Scripts Work	1424
	Commit Script Input	1425
	Commit Script Output	1425
	Commit Scripts and the Junos OS Commit Model	1426
	Standard Commit Model	1426
	Commit Model with Commit Scripts	1427
	Avoiding Potential Conflicts When Using Multiple Commit Scripts	1429
	Overview of Generating Persistent or Transient Configuration Changes	1430
	Differences Between Persistent and Transient Changes	1430
	Interaction of Configuration Changes and Configuration Groups	1433
	Tag Elements and Templates for Generating Changes	1434
	Required Boilerplate for Commit Scripts	1434
	How Op Scripts Work	1436
	Required Boilerplate for Op Scripts	1436
	Installing Junos OS Software with Junos Automation Enhancements	1438
	Invoking the Python Interpreter	1442
	Controlling the Execution of Commit Scripts	1443
	Enabling Commit Scripts to Execute	1443
	Removing Commit Scripts from the Configuration	1444
	Deactivating Commit Scripts	1445
	Activating Inactive Commit Scripts	1445
	Displaying Commit Script Output	1445
Chapter 47	Configuring sFlow Technology	1449
	Understanding How to Use sFlow Technology for Network Monitoring on a Switch	1449
	Sampling Mechanism and Architecture of sFlow Technology on Switches	1449
	Adaptive Sampling	1451
	sFlow Agent Address Assignment	1452
	sFlow Limitations on Switches	1452
	Example: Monitoring Network Traffic Using sFlow Technology	1453
	Configuring sFlow Technology	1457
Chapter 48	Configuring SNMP	1459
	Utility MIB	1460
	SNMPv3 Overview	1460
	Minimum SNMPv3 Configuration on a Device Running Junos OS	1461
	Understanding RMON	1463
	RMON Overview	1463
	Alarm Thresholds and Events	1464
	RMON MIB Event, Alarm, Log, and History Control Tables	1464
	Understanding Health Monitoring	1467
	SNMP MIBs Support	1468
	MIBs Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis	1468
	MIBs Supported on QFabric Systems	1478

SNMP Traps Support	1484
SNMP Traps Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis	1484
SNMPv1 Traps	1484
SNMPv2 Traps	1489
SNMP Traps Supported on QFabric Systems	1493
MIB Objects for the QFX Series	1497
QFX Series Standalone Switches	1497
QFabric Systems	1497
QFabric System QFX3100 Director Device	1498
QFabric System QFX3008-I Interconnect Device	1498
QFabric System QFX3600-I Interconnect Device	1498
QFabric System Node Devices	1499
Configuring SNMP	1499
Configuring the SNMP Community String	1503
Configuring SNMP Trap Groups	1504
Adding a Group of Clients to an SNMP Community	1505
Configuring the Interfaces on Which SNMP Requests Can Be Accepted	1506
Configuring MIB Views	1506
Configuring RMON Alarms and Events	1507
Configuring SNMP	1508
Configuring an Event	1508
Configuring an Alarm	1509
Configuring Health Monitoring	1509
Creating SNMPv3 Users	1510
Configuring Access Privileges for a Group	1511
Assigning a Security Name to a Group	1513
Configuring SNMPv3 Traps on a Device Running Junos OS	1514
Configuring SNMP Informs	1515
Monitoring RMON MIB Tables	1516
Monitoring SNMP	1516
Tracing SNMP Activity on a Device Running Junos OS	1518
Configuring the Number and Size of SNMP Log Files	1519
Configuring Access to the Log File	1519
Configuring a Regular Expression for Lines to Be Logged	1520
Configuring the Trace Operations	1520
Using the Enterprise-Specific Utility MIB to Enhance SNMP Coverage	1521
Example: Configuring SNMP	1523
Chapter 49	
Configuring System Logging	1527
Overview of Junos OS System Log Messages	1528
Overview of Single-Chassis System Logging Configuration	1528
Junos OS Minimum System Logging Configuration	1530
Junos OS System Log Configuration Statements	1530
Adding a Text String to System Log Messages	1531
Directing System Log Messages to a Log File	1532
Directing System Log Messages to a Remote Machine	1533
Directing System Log Messages to a User Terminal	1533
Directing System Log Messages to the Console	1534

	Disabling the System Logging of a Facility	1534
	Displaying a Log File from a Single-Chassis System	1535
	Including Priority Information in System Log Messages	1536
	Including the Year or Millisecond in Timestamps	1537
	Logging Messages in Structured-Data Format	1538
	Interpreting Messages Generated in Structured-Data Format	1539
	Interpreting Messages Generated in Standard Format	1542
	Specifying Log File Size, Number, and Archiving Properties	1543
	Specifying the Facility and Severity of Messages to Include in the Log	1544
	Junos OS System Logging Facilities and Message Severity Levels	1546
	Default Facilities for System Log Messages Directed to a Remote Destination	1548
	Junos OS System Log Alternate Facilities for Remote Logging	1548
	Changing the Alternative Facility Name for System Log Messages Directed to a Remote Destination	1549
	Using Regular Expressions to Refine the Set of Logged Messages	1551
	Managing Host OS System Log and Core Files	1553
	Viewing Log Files On the Host OS System	1554
	Copying Log Files From the Host System To the Switch	1554
	Viewing Core Files On the Host OS System	1554
	Copying Core Files From the Host System To the Switch	1554
	Cleaning Up Temporary Files on the Host OS	1555
	Monitoring System Log Messages	1555
	Examples: Configuring System Logging	1556
	Examples: Assigning an Alternative Facility	1558
	Example: Configuring System Log Messages	1559
Part 21	Configuration Statements	
Chapter 50	Network Management	1565
	connection-limit	1566
	destination-override	1567
	no-remote-trace	1567
	protocol-version	1568
	rate-limit	1569
	ssh	1570
	telnet	1571
	tracing	1572
Chapter 51	Automation	1573
	allow-transients	1574
	apply-macro	1575
	checksum	1576
	command	1577
	commit	1578
	description	1579
	direct-access	1579
	file (Commit Scripts)	1580
	file (Op Scripts)	1581
	no-allow-url	1582

	op	1583
	optional	1584
	refresh (Commit Scripts)	1584
	refresh (Op Scripts)	1585
	refresh-from (Commit Scripts)	1585
	refresh-from (Op Scripts)	1586
	scripts	1587
	source (Commit Scripts)	1589
	source (Op Scripts)	1589
Chapter 52	sFlow Technology	1591
	agent-id	1591
	collector (sFlow Technology)	1592
	interfaces (sFlow)	1593
	polling-interval	1594
	sample-rate	1595
	sflow	1596
	source-ip	1597
	traceoptions (sFlow Technology)	1598
	udp-port	1599
Chapter 53	SNMP	1601
	access (SNMP)	1604
	address (SNMP)	1604
	address-mask	1605
	agent-address	1605
	alarm (SNMP RMON)	1606
	authentication-md5	1607
	authentication-none	1608
	authentication-password	1609
	authentication-sha	1610
	authorization	1611
	bucket-size	1612
	categories	1613
	client-list	1613
	client-list-name	1614
	clients	1614
	commit-delay	1615
	community (SNMP)	1616
	community (RMON)	1617
	community-name	1618
	contact	1619
	description	1619
	description (RMON)	1620
	destination-port (SNMP)	1620
	engine-id	1621
	event	1622
	falling-event-index (RMON)	1623
	falling-threshold (Health Monitor)	1624
	falling-threshold (RMON)	1625

falling-threshold-interval	1626
filter-duplicates	1626
filter-interfaces	1627
group (Defining Access Privileges for an SNMPv3 Group)	1627
group (Configuring Group Name)	1628
health-monitor	1629
history	1630
interface (SNMP)	1631
interface (RMON)	1632
interval (Health Monitor)	1633
interval (RMON)	1634
local-engine	1635
location	1636
message-processing-model	1636
name	1637
nonvolatile	1637
notify	1638
notify-filter (Applying to the Management Target)	1639
notify-filter (Configuring the Profile Name)	1639
notify-view	1640
oid	1640
oid (SNMPv3)	1641
owner	1642
parameters	1642
port	1643
privacy-3des	1644
privacy-aes128	1645
privacy-des	1646
privacy-none	1646
privacy-password	1647
read-view	1648
remote-engine	1649
request-type	1650
retry-count	1651
rising-event-index	1652
rising-threshold (Health Monitor)	1653
rising-threshold (RMON)	1654
rmon	1655
sample-type	1656
security-level (Defining Access Privileges)	1657
security-level (Generating SNMP Notifications)	1658
security-model (Access Privileges)	1659
security-model (Group)	1660
security-model (SNMP Notifications)	1660
security-name (Community String)	1661
security-name (Security Group)	1662
security-name (SNMP Notifications)	1663
security-to-group	1664
snmp	1665

	snmp-community	1669
	source-address (SNMP)	1669
	startup-alarm	1670
	syslog-subtag	1671
	tag (Configuring Notification Targets)	1671
	tag (Configuring the SNMP Community)	1672
	tag-list	1672
	target-address	1673
	target-parameters	1674
	targets	1675
	timeout	1675
	traceoptions (SNMP)	1676
	trap-group	1678
	trap-options	1679
	type (RMON Notification)	1680
	type	1681
	user	1681
	usm	1682
	v3	1684
	vacm	1686
	variable	1687
	version	1688
	view (Configuring a MIB View)	1689
	view (Associating MIB View with a Community)	1690
	write-view	1690
Chapter 54	System Logging	1691
	archive (All System Log Files)	1692
	archive (Individual System Log File)	1694
	console (System Logging)	1695
	explicit-priority	1696
	facility-override	1696
	file (System Logging)	1697
	files	1698
	host	1699
	log-prefix (System)	1701
	match	1701
	size (System)	1702
	structured-data	1703
	syslog (System)	1704
	time-format	1706
	user (System Logging)	1707
Part 22	Operational Commands	
Chapter 55	Network Management	1711
	monitor traffic	1712
	ping	1724

Chapter 56	sFlow Technology	1729
	clear sflow collector statistics	1730
	show sflow	1731
	show sflow collector	1733
	show sflow interface	1734
Chapter 57	SNMP	1737
	clear snmp history	1738
	clear snmp statistics	1739
	request snmp spoof-trap	1741
	request snmp utility-mib clear instance	1747
	request snmp utility-mib set instance	1748
	show snmp health-monitor	1749
	show snmp inform-statistics	1754
	show snmp mib	1756
	show snmp rmon	1759
	show snmp rmon history	1763
	show snmp statistics	1764
	show snmp v3	1772
Chapter 58	System Logging	1775
	request app-engine file-copy (crash log) from-jhost to-vjunos	1776
	request app-engine cleanup	1778
	show app-engine crash	1780
	show app-engine logs	1782
	show log	1784
Guide 4	User and Access Management Feature Guide for QFX10000 Switches	

Part 23	User and Access Management Overview	
Chapter 59	Understanding the Software	1793
	Understanding Junos OS Infrastructure and Processes	1793
	Routing Engine and Packet Forwarding Engine	1793
	Junos OS Processes	1794
	Understanding LLDP	1795
	Monitoring SNMP	1796
Chapter 60	Understanding Access and Authentication Methods	1799
	Understanding Junos OS Access Privilege Levels	1799
	Junos OS Login Class Permission Flags	1799
	Allowing or Denying Individual Commands for Junos OS Login Classes	1802
	Junos OS User Authentication Methods	1804
	Understanding Login Authentication	1804
	MAC RADIUS Authentication	1805

Part 24	Configuring Access	
Chapter 61	Configuring and Managing Root Users	1809
	Configuring Management Access	1809
	Configuring Access Privilege Levels	1809
	Configuring Login Tips	1810
	Recovering the Root Password	1810
	Example: Configuring a Plain-Text Password for Root Logins	1812
	Example: Configuring SSH Authentication for Root Logins	1814
	Understanding Troubleshooting Resources	1814
	Troubleshooting Overview	1816
	Recovering the Root Password	1818
Chapter 62	Configuring and Managing User Accounts	1821
	Junos OS User Accounts Overview	1821
	Junos OS Login Classes Overview	1823
	Special Requirements for Junos OS Plain-Text Passwords	1824
	Regular Expressions for Allowing and Denying Junos OS Configuration Mode Hierarchies	1826
	Regular Expressions for Allowing and Denying Junos OS Operational Mode Commands	1827
	Examples of Defining Access Privileges Using allow/deny-configuration Statements	1829
	Example: Configuring User Accounts	1830
	Example: Configuring User Permissions with Access Privilege Levels	1831
	Example: Configuring User Permissions with Access Privileges for Operational Mode Commands	1832
	Example: Changing the Requirements for Junos OS Plain-Text Passwords	1834
	Example: Limiting the Number of Login Attempts for SSH and Telnet Sessions to Prevent Unauthorized Access	1835
	Understanding Troubleshooting Resources	1836
	Troubleshooting Overview	1838
	Recovering the Root Password	1840
Part 25	Configuring Authentication	
Chapter 63	Configuring and Managing Local Password Authentication	1845
	Junos OS User Accounts Overview	1845
	Junos OS User Authentication Methods	1847
	Junos OS Login Classes Overview	1848
	Regular Expressions for Allowing and Denying Junos OS Configuration Mode Hierarchies	1848
	Regular Expressions for Allowing and Denying Junos OS Operational Mode Commands	1849
	Special Requirements for Junos OS Plain-Text Passwords	1850
	Configuring Junos OS User Accounts	1853
	Configuring a Local Administrator Account	1853
	Example: Creating Login Classes with Specific Privileges	1854

	Example: Configuring System Authentication for RADIUS, TACACS+, and Password Authentication	1855
	Example: Changing the Requirements for Junos OS Plain-Text Passwords . . .	1857
Chapter 64	Configuring and Managing TACACS+ Authentication	1859
	Junos OS Authentication Order for RADIUS, TACACS+, and Password Authentication	1859
	Using RADIUS or TACACS+ Authentication	1859
	Using Local Password Authentication	1860
	Order of Authentication Attempts	1860
	Juniper Networks Vendor-Specific TACACS+ Attributes	1863
	Example: Configuring System Authentication for RADIUS, TACACS+, and Password Authentication	1865
Chapter 65	Configuring and Managing RADIUS Authentication	1869
	Junos OS Authentication Order for RADIUS, TACACS+, and Password Authentication	1869
	Using RADIUS or TACACS+ Authentication	1869
	Using Local Password Authentication	1870
	Order of Authentication Attempts	1870
	Configuring RADIUS Authentication (QFX Series or OCX Series)	1873
	Configuring RADIUS Server Details	1874
	Configuring MS-CHAPv2 for Password-Change Support	1875
	Specifying a Source Address for the Junos OS to Access External RADIUS Servers	1876
	Using Regular Expressions on a RADIUS or TACACS+ Server to Allow or Deny Access to Commands	1876
	Example: Configuring RADIUS Authentication	1878
	Example: Configuring RADIUS Template Accounts	1879
	Configuring a Local Administrator Account	1879
	Example: Configuring System Authentication for RADIUS, TACACS+, and Password Authentication	1880
Chapter 66	Configuring and Managing RADIUS Accounting	1883
	Understanding RADIUS Accounting	1883
	Junos OS Authentication Order for RADIUS, TACACS+, and Password Authentication	1884
	Using RADIUS or TACACS+ Authentication	1884
	Using Local Password Authentication	1885
	Order of Authentication Attempts	1885
	Juniper Networks Vendor-Specific RADIUS Attributes	1888
	Configuring RADIUS System Accounting	1891
	Configuring Auditing of User Events on a RADIUS Server	1891
	Specifying RADIUS Server Accounting and Auditing Events	1891
	Configuring RADIUS Server Accounting	1892
	Configuring RADIUS Authentication (QFX Series or OCX Series)	1893
	Configuring RADIUS Server Details	1894
	Configuring MS-CHAPv2 for Password-Change Support	1895
	Specifying a Source Address for the Junos OS to Access External RADIUS Servers	1896

	Using Regular Expressions on a RADIUS or TACACS+ Server to Allow or Deny Access to Commands	1896
	Example: Configuring RADIUS System Accounting	1898
Chapter 67	Configuring and Managing RADIUS Template Accounts	1899
	Overview of Template Accounts for RADIUS and TACACS+ Authentication . .	1899
	Example: Configuring RADIUS Template Accounts	1899
Chapter 68	Configuring and Managing VSAs for RADIUS and TACACS+	1901
	Understanding Vendor-Specific Attributes (VSAs)	1901
	Juniper-Switching-Filter VSA Match Conditions and Actions	1902
	Juniper Networks Vendor-Specific TACACS+ Attributes	1904
Part 26	Configuration Statements and Operational Commands	
Chapter 69	Configuration Statements	1909
	access	1911
	accounting (Access Profile)	1912
	accounting-options	1913
	accounting-server	1915
	accounting-stop-on-access-deny	1916
	accounting-stop-on-failure	1917
	advertisement-interval	1918
	agent-address	1919
	archival	1920
	archive-sites (Configuration File)	1921
	authentication-order	1922
	authentication-server	1923
	authorization	1924
	categories	1925
	client-list	1925
	client-list-name	1926
	clients	1926
	commit-delay	1927
	community (SNMP)	1928
	configuration	1929
	connection-limit	1930
	contact	1931
	disable (LLDP)	1931
	falling-threshold (Health Monitor)	1932
	filter-duplicates	1932
	full-name	1933
	health-monitor	1933
	hold-multiplier	1934
	idle-timeout (Access)	1935
	interface (LLDP)	1936
	interval (Health Monitor)	1937
	lldp	1938
	lldp-configuration-notification-interval	1940
	location	1940

	management-address	1941
	name	1942
	nas-ip-address	1942
	nonvolatile	1943
	oid	1943
	order	1944
	port (RADIUS Server)	1945
	profile	1946
	protocols	1947
	protocol-version	1960
	ptopo-configuration-maximum-hold-time	1961
	ptopo-configuration-trap-interval	1961
	radius	1962
	radius-options (edit system)	1963
	radius-server	1964
	rate-limit	1965
	remote-debug-permission	1966
	retry	1967
	rising-threshold (Health Monitor)	1968
	root-login	1969
	services (Switches)	1970
	snmp	1971
	ssh	1975
	system	1976
	tacplus-options	1982
	targets	1983
	traceoptions (LLDP)	1984
	transfer-interval (Configuration)	1986
	transfer-on-commit	1987
	trap-group	1988
	trap-options	1989
	user (Access)	1990
	version	1991
Chapter 70	Operational Commands	1993
	clear lldp neighbors	1994
	clear lldp statistics	1995
	request component login	1996
	show ethernet-switching interfaces	1998
	show lldp	2002
	show lldp local-information	2007
	show lldp neighbors	2009
	show lldp statistics	2014
	show route instance	2016
	show snmp statistics	2020
	ssh	2028

Guide 5 Ethernet Switching Feature Guide for QFX10000 Switches

Part 27	ELS CLI
Chapter 71	Enhanced Layer 2 Software (ELS) CLI 2035
	Getting Started with Enhanced Layer 2 Software 2035
	Understanding Enhanced Layer 2 Software Support 2035
	Using the ELS Translator 2036
	Configuring a VLAN 2037
	Configuring the Native VLAN Identifier 2037
	Configuring Layer 2 Interfaces 2038
	Configuring Layer 3 Interfaces 2038
	Configuring an IRB Interface 2039
	Configuring an Aggregated Ethernet Interface and Configuring LACP on That Interface 2039
	Enhanced Layer 2 CLI Configuration Statement and Command Changes . . 2040
	Changes to the ethernet-switching-options Hierarchy Level 2041
	Changes to the Port Mirroring Hierarchy Level 2042
	Changes to the Layer 2 Control Protocol Hierarchy Level 2043
	Changes to the dot1q-tunneling Statement 2043
	Changes to the L2 Learning Protocol 2044
	Changes to Nonstop Bridging 2044
	Changes to Port Security and DHCP Snooping 2044
	Changes to Configuring VLANs 2046
	Changes to Storm Control Profiles 2049
	Changes to the Interfaces Hierarchy 2050
	Changes to IGMP Snooping 2051
Part 28	Ethernet Ring Protection
Chapter 72	Understanding Ethernet Ring Protection 2055
	Ethernet Ring Protection Switching Overview 2055
	Understanding Ethernet Ring Protection Switching Functionality 2056
	Acronyms 2056
	Ring Nodes 2057
	Ring Node States 2057
	Default Logging of Basic State Transitions on EX Series Switches 2057
	Failure Detection 2058
	Logical Ring 2058
	FDB Flush 2058
	Traffic Blocking and Forwarding 2058
	RPL Neighbor Node 2059
	RAPS Message Blocking and Forwarding 2059
	Dedicated Signaling Control Channel 2060
	RAPS Message Termination 2060
	Revertive and Non-revertive Modes 2061
	Multiple Rings 2061
	Node ID 2061

Ring ID	2061
Bridge Domains with the Ring Port (MX Series Routers Only)	2061
Wait-to-Block Timer	2061
Adding and Removing a Node	2062
Configuring Ethernet Ring Protection Switching (CLI Procedure)	2063
Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS	2066

Part 29

Chapter 73

Bridging and VLANs

Using Bridging and VLANs	2085
Overview of Layer 2 Networking	2085
Understanding Layer 2 Broadcasting	2087
Layer 2 Learning and Forwarding for VLANs Overview	2088
Understanding Bridging and VLANs	2089
History of VLANs	2089
How Bridging of VLAN Traffic Works	2090
Packets Are Either Tagged or Untagged	2091
Switch Interface Modes—Access, Trunk, or Tagged Access	2091
Access Mode	2092
Trunk Mode	2092
Trunk Mode and Native VLAN	2092
Tagged-Access Mode	2093
Additional Advantages of Using VLANs	2093
Maximum VLANs and VLAN Members Per Switch	2094
A Default VLAN Is Configured on Most Switches	2095
Assigning Traffic to VLANs	2095
Assign VLAN Traffic According to the Interface Port Source	2095
Assign VLAN Traffic According to the Source MAC Address	2096
Forwarding VLAN Traffic	2096
VLANs Communicate with Integrated Routing and Bridging Interfaces or Routed VLAN Interfaces	2096
Example: Setting Up Basic Bridging and a VLAN on the QFX Series	2096
Example: Setting Up Bridging with Multiple VLANs	2114
Configuring VLANs	2120
Configuring the Native VLAN Identifier (CLI Procedure)	2121
Creating a Series of Tagged VLANs	2123
Understanding Integrated Routing and Bridging	2124
Example: Configuring Routing Between VLANs on One Switch	2126
Configuring IRB Interfaces	2131
Adding a Static MAC Address Entry to the Ethernet Switching Table (CLI Procedure)	2133
Configuring Static ARP Entries	2133
Troubleshooting Ethernet Switching	2134

Part 30	MAC Addresses	
Chapter 74	Using MAC Addresses	2137
	Introduction to the Media Access Control (MAC) Layer 2 Sublayer	2137
	Understanding MAC Learning	2138
	Disabling MAC Learning	2138
	Example: Disabling MAC Learning	2139
	Configuring MAC Notification (CLI Procedure)	2140
	Enabling MAC Notification	2140
	Disabling MAC Notification	2141
	Setting the MAC Notification Interval	2141
	Verifying That MAC Notification Is Working Properly	2141
	Configuring MAC Limiting (CLI Procedure)	2142
	Limiting the Number of MAC Addresses Learned by an Interface	2143
	Limiting the Number of MAC Addresses Learned by a VLAN	2143
	Configuring MAC Table Aging	2144
Part 31	Spanning Trees	
Chapter 75	Using Spanning Trees	2147
	Overview of Spanning-Tree Protocols	2148
	Understanding Spanning Tree Protocols on a QFabric System	2148
	Understanding MSTP	2149
	Example: Configuring Network Regions for VLANs with MSTP	2149
	Configuring MSTP	2170
	Understanding RSTP	2172
	Example: Configuring Faster Convergence and Improved Network Stability with RSTP	2173
	Configuring RSTP (CLI Procedure)	2190
	Understanding VSTP	2191
	Example: Configuring VSTP on QFX Series Switches and EX4600 Switches	2191
	Configuring VSTP (CLI Procedure)	2197
	Understanding BPDU Protection for STP, RSTP, and MSTP	2199
	Example: Configuring BPDU Protection on Edge Interfaces to Prevent STP Miscalculations	2200
	Configuring BPDU Protection on Spanning Tree Interfaces	2205
	Understanding Loop Protection for STP, RSTP, VSTP, and MSTP	2206
	Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree	2207
	Understanding Root Protection for STP, RSTP, VSTP, and MSTP	2211
	Example: Configuring Root Protection to Enforce Root Bridge Placement in Spanning Trees	2212
	Unblocking an Interface That Receives BPDUs in Error (CLI Procedure)	2217
Part 32	Q-in-Q Tunneling	
Chapter 76	Using Q-in-Q Tunneling	2221
	Understanding Q-in-Q Tunneling	2221
	How Q-in-Q Tunneling Works	2222
	How VLAN Translation Works	2222

	Sending and Receiving Untagged Packets	2222
	Disabling MAC Address Learning	2223
	Mapping C-VLANs to S-VLANs	2223
	All-in-One Bundling	2223
	Many-to-Many Bundling	2224
	Mapping a Specific Interface	2224
	Constraints for Q-in-Q Tunneling and VLAN Translation	2224
	Configuring Q-in-Q Tunneling	2226
	Using the Different Mapping Methods	2226
	Configuring All-in-One Bundling	2227
	Configuring Many-to-Many Bundling	2228
	Configuring a Specific Interface Mapping with VLAN ID Translation Option	2231
	Configuring All-in-One Bundling	2233
	Configuring Many-to-Many Bundling	2235
	Configuring a Specific Interface Mapping with VLAN ID Translation Option	2237
Part 33	Proxy ARP	
Chapter 77	Using Proxy ARP	2243
	Understanding Proxy ARP	2243
	What Is ARP?	2243
	Proxy ARP Overview	2243
	Best Practices for Proxy ARP	2244
	Configuring Proxy ARP (CLI Procedure)	2245
	Verifying That Proxy ARP Is Working Correctly	2245
Part 34	Reflective Relay	
Chapter 78	Using Reflective Relay	2249
	Understanding Reflective Relay for Use with VEPA Technology	2249
	VEPA	2249
	Reflective Relay	2249
	Configuring Reflective Relay	2250
	Example: Configuring Reflective Relay for Use with VEPA Technology	2251
Part 35	Configuration Statements and Operational Commands	
Chapter 79	Ethernet Ring Protection Configuration Statements	2259
	control-channel	2260
	control-vlan	2261
	data-channel	2262
	east-interface	2263
	ethernet-ring	2264
	guard-interval	2265
	hold-interval (Protection Group)	2266
	protection-group	2267
	restore-interval	2268
	ring-protection-link-end	2269
	ring-protection-link-owner	2269

	west-interface	2270
Chapter 80	VLAN Configuration Statements	2271
	[edit vlans] Configuration Statement Hierarchy on the QFX Series	2272
	Supported Statements in the [edit vlans] Hierarchy Level	2272
	Unsupported Statements in the [edit vlans] Hierarchy Level	2274
	description (VLAN)	2275
	dhcp-relay	2276
	filter (VLANs)	2281
	forwarding-options	2282
	interface (VLANs)	2287
	interface-mac-limit	2288
	interface-mode	2290
	irb (Interfaces)	2292
	l3-interface (VLAN)	2295
	mac (Static MAC-Based VLANs)	2296
	members	2297
	native-vlan-id	2298
	packet-action	2299
	port-mode	2302
	service-id	2303
	switch-options	2304
	static (Static MAC-Based VLANs)	2305
	static-mac	2305
	vlan-id (VLANs)	2306
	vlan-id-list	2307
	vlan-rewrite	2308
	vlan-tagging	2309
	vlan-tags	2310
	vlans	2311
Chapter 81	MAC Address Configuration Statements	2315
	global-mac-table-aging-time	2315
	mac-limit	2316
	mac-notification	2317
	mac-statistics	2318
	mac-table-size	2320
	notification-interval	2321
Chapter 82	STP Configuration Statements	2323
	bpdu-block	2324
	bpdu-block-on-edge	2325
	bpdu-timeout-action	2326
	bridge-priority (Spanning Trees)	2327
	configuration-name	2328
	cost	2329
	disable (Spanning Trees)	2330
	bpdu-block	2331
	bpdu-block-on-edge	2332
	bpdu-timeout-action	2333

	bridge-priority (Spanning Trees)	2334
	configuration-name	2335
	cost	2336
	disable (Spanning Trees)	2337
	disable-timeout (Spanning Trees)	2338
	edge	2339
	force-version (IEEE 802.1D STP)	2340
	forward-delay	2341
	hello-time	2342
	interface (BPDU)	2343
	interface (Spanning Tree)	2344
	max-age	2345
	max-hops	2346
	mode	2347
	msti	2348
	mstp	2349
	no-root-port	2351
	priority (Protocols STP)	2352
	protocol	2353
	protocols (STP Type)	2354
	revision-level	2355
	rstp	2356
	traceoptions (Spanning Tree)	2357
	vlan (MSTP)	2360
	vlan (VSTP)	2361
	vlan (VSTP)	2362
	vlan-group	2363
	vstp	2364
Chapter 83	Q-in-Q Configuration Statements	2367
	flexible-vlan-tagging	2368
	input-vlan-map	2369
	native-vlan-id	2370
	output-vlan-map (Gigabit Ethernet IQ and 10-Gigabit Ethernet with SFPP) . .	2371
	pop	2372
	push	2373
	swap	2374
	vlan-id-list	2375
Chapter 84	Reflective Relay Configuration Statement	2377
	reflective-relay	2377
Chapter 85	Bridging and VLANs Monitoring Commands	2379
	clear ethernet-switching table	2380
	show ethernet-switching interfaces	2382
	show ethernet-switching table	2386
	show system statistics arp	2394
	show vlans	2395

Chapter 86	MAC Address Operational Commands	2405
	show ethernet-switching mac-learning-log	2406
	show ethernet-switching mac-notification	2408
	show ethernet-switching statistics aging	2410
	show ethernet-switching statistics mac-learning	2412
Chapter 87	Spanning Tree Monitoring Commands	2417
	clear error bpdu interface	2418
	clear spanning-tree statistics	2419
	show spanning-tree bridge	2420
	show spanning-tree interface	2425
	show spanning-tree mstp configuration	2431
	show spanning-tree statistics	2433
Guide 6	OVSDB and VXLAN Feature Guide for QFX10000 Switches (VMware NSX)	
Part 36	Overview	
Chapter 88	OVSDB and VXLAN Overview	2439
	Understanding the Junos OS Implementation of OVSDB and VXLAN in a VMware NSX for Multi-Hypervisor Environment	2439
	Understanding VXLANs	2442
	VXLAN Benefits	2442
	How Does VXLAN Work?	2443
	VXLAN Implementation Methods	2443
	Using an MX Series Router or EX9200 Switch as a VTEP	2444
	Manual VXLANs Require PIM	2444
	Understanding the OVSDB Protocol Running on Juniper Networks Devices	2445
	OVSDDB Support on Juniper Networks Devices	2446
	OVSDDB Schema for Physical Devices	2447
	Understanding How Layer 2 BUM and Layer 3 Routed Multicast Traffic Are Handled with OVSDB	2450

Part 37**Configuring OVSDB and VXLAN****Chapter 89****Configuring OVSDB-Managed VXLANs with an SDN Controller 2455**

OVSDB and VXLAN Configuration Workflows for VMware NSX Environment . .	2455
OVSDB and VXLAN Configuration Workflow for QFX Series Switches . . .	2455
OVSDB and VXLAN Configuration Workflow for MX Series Routers and EX9200 Switches	2457
Understanding How to Set Up OVSDB Connections on a Juniper Networks Device	2458
Creating and Installing an SSL Key and Certificate on a Juniper Networks Device for Connection with SDN Controllers	2459
Setting Up OVSDB on Juniper Networks Devices That Support the Dynamic Configuration of VXLANs	2461
Understanding Dynamically Configured VXLANs in an OVSDB Environment . .	2462
Performing Tasks Before and After the Dynamic Configuration of OVSDB-Managed VXLANs	2462
What the Juniper Networks Switch Actually Creates Dynamically	2467
Dynamic Association of a Trunk Interface Supporting Untagged Packets to a Dynamically Created VXLAN	2467
Dynamic Association of a Trunk Interface Supporting Tagged Packets to a Dynamically Created VXLAN	2468
VMware NSX Configuration for Juniper Networks Devices Functioning as Virtual Tunnel Endpoints	2469
Creating a Gateway	2470
Creating a Gateway Service	2470
Creating a Logical Switch Port	2471
Example: Setting Up a VXLAN Layer 2 Gateway and OVSDB Connections in a VMware NSX Environment (Trunk Interfaces Supporting Untagged Packets)	2472
Example: Setting Up a VXLAN Layer 2 Gateway and OVSDB Connections in a VMware NSX Environment (Trunk Interfaces Supporting Tagged Packets)	2480
Verifying That a Logical Switch and Corresponding Junos OS OVSDB-Managed VXLAN Are Working Properly	2488

Chapter 90**Configuring VXLANs Without an SDN Controller 2491**

VXLAN Constraints on QFX Series Switches	2491
Manually Configuring VXLANs on a QFX5100 Switch	2492
Configuring a Source IP Address	2493
Configuring PIM for VXLANs	2493
Configuring VXLANs	2493
Examples: Manually Configuring VXLANs on QFX Series Switches	2494
Example: Configuring a VXLAN Transit Switch	2494
Example: Configuring a VXLAN Layer 2 Gateway	2496
Verifying That a Local VXLAN VTEP Is Configured Correctly	2502
Verifying MAC Learning from a Remote VTEP	2502

Part 38	Troubleshooting	
Chapter 91	Troubleshooting Tasks	2507
	Troubleshooting a Nonoperational Logical Switch and Corresponding Junos OS	
	OVSDB-Managed VXLAN	2507
	Verifying VXLAN Reachability	2509
	Monitoring a Remote VTEP Interface	2509
Part 39	Configuration Statements and Operational Commands	
Chapter 92	OVSDB Configuration Statements	2513
	controller (OVSDB)	2514
	inactivity-probe-duration	2515
	interfaces (OVSDB)	2516
	maximum-backoff-duration	2517
	ovsdb	2518
	ovsdb-managed	2519
	port (OVSDB)	2520
	protocol (OVSDB)	2521
	traceoptions (OVSDB)	2522
Chapter 93	VXLAN Configuration Statements	2525
	decapsulate-accept-inner-vlan	2525
	encapsulate-inner-vlan	2526
	multicast-group	2526
	ovsdb-managed	2527
	unreachable-vtep-aging-timer	2528
	vni	2528
	vtep-source-interface	2529
	vxlan	2529
Chapter 94	OVSDB Operational Commands	2531
	clear ovsdb commit failures	2532
	show ovsdb commit failures	2534
	show ovsdb controller	2536
	show ovsdb interface	2538
	show ovsdb logical-switch	2540
	show ovsdb mac	2543
	show ovsdb statistics interface	2547
	show ovsdb virtual-tunnel-end-point	2549
	show vpls mac-table	2551
Chapter 95	VXLAN Operational Commands	2557
	show bridge mac-table	2558
	show vpls mac-table	2563

Guide 7 High Availability Feature Guide for QFX10000 Switches

Part 40	Adaptive Load Balancing	
Chapter 96	Configuring Load Balancing	2573
	Understanding Aggregated Ethernet Load Balancing	2573
	Configuring Adaptive Load Balancing	2575
Part 41	Graceful Restart	
Chapter 97	Configuring Graceful Restart	2579
	Graceful Restart Concepts	2579
	Configuring Routing Protocols Graceful Restart	2580
	Enabling Graceful Restart	2580
	Configuring Graceful Restart Options for BGP	2581
	Configuring Graceful Restart Options for ES-IS	2582
	Configuring Graceful Restart Options for IS-IS	2582
	Configuring Graceful Restart Options for OSPF and OSPFv3	2583
	Configuring Graceful Restart Options for RIP and RIPng	2585
	Configuring Graceful Restart Options for PIM Sparse Mode	2585
	Tracking Graceful Restart Events	2586
Part 42	Nonstop Bridging	
Chapter 98	Configuring Nonstop Bridging	2589
	Nonstop Bridging Concepts	2589
	Nonstop Bridging System Requirements	2591
	Platform Support	2591
	Protocol Support	2592
	Configuring Nonstop Bridging on Switches (CLI Procedure)	2593
Part 43	Nonstop Active Routing	
Chapter 99	Configuring Nonstop Active Routing	2597
	Nonstop Active Routing Concepts	2597
	Nonstop Active Routing System Requirements	2600
	Nonstop Active Routing Platform and Switching Platform Support	2600
	Nonstop Active Routing Protocol and Feature Support	2601
	Nonstop Active Routing BFD Support	2604
	Nonstop Active Routing BGP Support	2605
	Nonstop Active Routing Layer 2 Circuit and VPLS Support	2606
	Nonstop Active Routing PIM Support	2606
	Nonstop Active Routing MSDP Support	2609
	Nonstop Active Routing Support for RSVP-TE LSPs	2609
	Example: Configuring Nonstop Active Routing on Switches	2611

Part 44	Graceful Routing Engine Switchover	
Chapter 100	Configuring Graceful Routing Engine Switchover	2617
	Understanding Graceful Routing Engine Switchover	2617
	Graceful Routing Engine Switchover Concepts	2617
	Effects of a Routing Engine Switchover	2622
	Graceful Routing Engine Switchover System Requirements	2623
	Graceful Routing Engine Switchover Platform Support	2623
	Graceful Routing Engine Switchover Feature Support	2624
	Graceful Routing Engine Switchover DPC Support	2626
	Graceful Routing Engine Switchover and Subscriber Access	2626
	Graceful Routing Engine Switchover PIC Support	2626
	Configuring Graceful Routing Engine Switchover	2627
	Enabling Graceful Routing Engine Switchover	2627
	Configuring Graceful Routing Engine Switchover with Graceful Restart	2627
	Synchronizing the Routing Engine Configuration	2627
	Verifying Graceful Routing Engine Switchover Operation	2629
	Configuring Graceful Routing Engine Switchover in a Virtual Chassis (CLI Procedure)	2629
	Resetting Local Statistics	2630
Part 45	Virtual Router Redundancy Protocol	
Chapter 101	Configuring Virtual Router Redundancy Protocol	2633
	Understanding VRRP	2633
	Overview of VRRP	2633
	Sample VRRP Topology	2634
	Example: Configuring VRRP for Load Sharing	2636
	Configuring Basic VRRP Support for QFX	2641
	Configuring VRRP Authentication (IPv4 Only)	2642
	Configuring the Startup Period for VRRP Operations	2643
	Configuring the Advertisement Interval for the VRRP Master	2643
	Modifying the Advertisement Interval in Seconds	2644
	Modifying the Advertisement Interval in Milliseconds	2644
	Configuring VRRP Preemption and Hold Time	2644
	Configuring VRRP Preemption	2645
	Configuring the Preemption Hold Time	2645
	Overriding the Hold Time	2645
	Configuring a Route to Be Tracked	2646
	Configuring a Logical Interface to Be Tracked	2646
	Configuring a Backup to Accept Packets Destined for the Virtual IP Address	2648
	Configuring Passive ARP Learning for VRRP Backups	2649
	Configuring the Silent Period	2649
	Configuring Inheritance for a VRRP Group	2650
	Troubleshooting VRRP	2651
Part 46	Configuration Statements and Operational Commands	
Chapter 102	Configuration Statements (Adaptive Load Balancing)	2655
	adaptive	2656

Chapter 103	Configuration Statements (Graceful Restart)	2657
	disable	2658
	disable (BGP Graceful Restart)	2659
	graceful-restart (Enabling Globally)	2660
	graceful-restart (Protocols BGP)	2662
	graceful-restart	2664
	helper-disable (OSPF)	2666
	no-strict-lsa-checking	2667
	notify-duration	2668
	redundancy (Graceful Switchover)	2669
	restart-duration	2670
	restart-time (BGP Graceful Restart)	2671
	stale-routes-time	2672
Chapter 104	Configuration Statements (Graceful Switchover)	2673
	graceful-switchover	2673
	redundancy (Graceful Switchover)	2674
Chapter 105	Configuration Statements (Nonstop Bridging and Routing)	2675
	nonstop-bridging	2675
	nonstop-routing	2676
	synchronize	2677
	traceoptions	2679
Chapter 106	Configuration Statements (VRRP)	2683
	accept-data	2684
	advertise-interval	2685
	asymmetric-hold-time	2686
	authentication-key	2687
	authentication-type	2688
	bandwidth-threshold	2689
	failover-delay	2690
	fast-interval	2691
	hold-time (VRRP)	2692
	interface	2693
	preempt (VRRP)	2694
	priority (Protocols VRRP)	2695
	priority-cost (VRRP)	2696
	priority-hold-time	2697
	route (Interfaces)	2698
	startup-silent-period	2699
	traceoptions	2700
	track (VRRP)	2702
	virtual-address	2703
	vrrp-group	2704
Chapter 107	Operational Mode Commands (Graceful Restart)	2707
	Verifying Graceful Restart Operation	2707
	Graceful Restart Operational Mode Commands	2707
	Verifying BGP Graceful Restart	2708

	Verifying IS-IS and OSPF Graceful Restart	2708
	Verifying CCC and TCC Graceful Restart	2709
	show bgp neighbor	2710
	show log	2727
	show (ospf ospf3) overview	2731
Chapter 108	Operational Mode Commands (Graceful Switchover)	2737
	show system switchover	2738
	show task replication	2744
Chapter 109	Operational Mode Command (Nonstop Routing)	2747
	show task replication	2748
Chapter 110	Operational Mode Commands (VRRP)	2751
	show vrrp	2752

Guide 8 Interfaces Feature Guide for QFX10000 Switches

Part 47	Ethernet OAM Link Fault Management	
Chapter 111	Understanding Ethernet OAM Link Fault Management	2767
	Understanding Ethernet OAM Link Fault Management for an EX Series	
	Switch	2767
	Configuring Ethernet OAM Link Fault Management (CLI Procedure)	2768
	Example: Configuring Ethernet OAM Link Fault Management on EX Series	
	Switches	2770
Part 48	Generic Routing Encapsulation (GRE)	
Chapter 112	Understanding GRE	2777
	Understanding Generic Routing Encapsulation	2777
	Overview of GRE	2777
	GRE Tunneling	2778
	Encapsulation and De-Encapsulation on the Switch	2778
	Number of Source and Destination Tunnels Allowed on a Switch	2778
	Class of Service on GRE Tunnels	2779
	Applying Firewall Filters to GRE Traffic	2779
	Using a Firewall Filter to De-encapsulate GRE Traffic on a QFX5100 and	
	OCX Series Switches	2779
	Configuration Limitations	2780
	Configuring Generic Routing Encapsulation Tunneling	2780
	Configuring a GRE Tunnel	2781
	Verifying That Generic Routing Encapsulation Tunneling Is Working Correctly	2781

Part 49	Interfaces	
Chapter 113	Understanding Interfaces	2785
	Interfaces Overview	2785
	Network Interfaces	2785
	Special Interfaces	2786
	Understanding Interface Naming Conventions	2787
	Physical Part of an Interface Name	2788
	Logical Part of an Interface Name on a Switch Running QFabric Software Package	2794
	Logical Part of a Channelized Interface Name on a Switch Running Enhanced Layer 2 Software	2795
	Wildcard Characters in Interface Names	2795
	Understanding Interface Ranges	2795
	Understanding Management Interfaces	2796
	Understanding Port Ranges and System Modes	2798
	Port Ranges for Different Media Types	2798
	Supported System Modes	2822
	Configuring the Interface Address	2824
	Configuring Gigabit and 10-Gigabit Ethernet Interfaces	2826
	Configuring Port Mode on QFX5100-48S, QFX5100-48T, QFX5100-24Q, and EX4600 Switches	2827
	Configuring the Link Settings for Gigabit Ethernet Interfaces on QFX5100-48S, QFX5100-24Q Switches, and EX4600 Switches	2828
	Configuring Gigabit Ethernet Interfaces on QFX5100-48T Switches	2828
	Configuring the Link Settings for 10-Gigabit Ethernet Interfaces on QFX5100-48S, QFX5100-24Q Switches, and EX4600 Switches	2830
	Configuring the Link Settings for 10-Gigabit Ethernet Interfaces on QFX5100-48T Switches	2830
	Configuring the IP Options on QFX5100-48S, QFX5100-48T, QFX5100-24Q, and EX4600 Switches	2830
	Configuring Ethernet Loopback Capability	2831
	Configuring an LPM Table With Junos OS Release 13.2X51-D10	2832
	Channelizing Interfaces	2835
	Monitoring Interface Status and Traffic	2844
	Troubleshooting Network Interfaces	2845
	The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP or SFP+ module is down	2845
Part 50	IP Directed Broadcast	
Chapter 114	Understanding IP Directed Broadcast	2849
	Understanding IP Directed Broadcast	2849
	IP Directed Broadcast Overview	2849
	IP Directed Broadcast Implementation	2850
	When to Enable IP Directed Broadcast	2850
	When Not to Enable IP Directed Broadcast	2850
	Configuring IP Directed Broadcast (CLI Procedure)	2851
	Example: Configuring IP Directed Broadcast	2852

Part 51	Layer 3 Logical Interfaces	
Chapter 115	Understanding Layer 3 Logical Interfaces	2859
	Understanding Layer 3 Logical Interfaces	2859
	Configuring a Layer 3 Logical Interface	2860
	Verifying That Layer 3 Logical Interfaces Are Working	2860
Part 52	Link Aggregation Groups (LAGs) and Link Aggregation Control Protocol (LACP)	
Chapter 116	Understanding LAGs and LACP	2865
	Understanding Aggregated Ethernet Interfaces and LACP	2865
	Link Aggregation Group	2866
	Link Aggregation Control Protocol (LACP)	2867
	Configuring Aggregated Ethernet LACP	2868
	Configuring Link Aggregation	2869
	Creating an Aggregated Ethernet Interface	2870
	Configuring the VLAN Name and VLAN ID Number	2870
	Configuring Aggregated Ethernet LACP	2871
	Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure)	2872
	Configuring LACP Link Protection for a Single Link at the Global Level	2873
	Configuring LACP Link Protection for a Single Link at the Aggregated Interface Level	2873
	Configuring Subgroup Bundles to Provide LACP Link Protection to Multiple Links in an Aggregated Ethernet Interface	2874
	Configuring Periodic Rebalancing of Subscribers in an Aggregated Ethernet Interface	2875
	Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch	2876
	Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch	2880
	Verifying the Status of a LAG Interface	2885
	Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets	2885
	Verifying the LACP Setup	2885
	Verifying That LACP Packets Are Being Exchanged	2886
	Troubleshooting an Aggregated Ethernet Interface	2886
Part 53	Load Balancing	
Chapter 117	Understanding Load Balancing	2891
	Configuring Load Balancing Based on MAC Addresses	2891
Part 54	Local Link Bias	
Chapter 118	Understanding Local Link Bias	2895
	Understanding Local Link Bias	2895
	Configuring Local Link Bias (CLI Procedure)	2897

Part 55	Redundant Trunk Groups	
Chapter 119	Understanding Redundant Trunk Groups	2901
	Understanding Redundant Trunk Links	2901
	Example: Configuring Redundant Trunk Links for Faster Recovery	2903
Part 56	Resilient Hashing	
Chapter 120	Understanding Resilient Hashing	2911
	Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop	
	ECMP Traffic	2911
	Understanding the Hashing Algorithm	2911
	IP (IPv4 and IPv6)	2913
	MPLS	2914
	MAC-in-MAC Packet Hashing	2915
	Layer 2 Header Hashing	2916
	Understanding the Use of Resilient Hashing to Minimize Flow Remapping in	
	Trunk/ECMP Groups	2917
	Why You Might Want to Use Resilient Hashing and How It Works with Static Hashing	2917
	Limitations and Caveats for Resilient Hashing	2918
	Resilient Hashing on LAGs	2918
	Resilient Hashing on ECMP	2919
	Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop	
	ECMP Traffic (QFX 10002 and QFX 10008 Switches)	2919
	Understanding the Hashing Algorithm	2919
	IP (IPv4 and IPv6)	2920
	MPLS	2921
	Layer 2 Header Hashing on the QFX10002 and QFX 10008 Switches	2922
	Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP	
	Traffic (CLI Procedure)	2923
	Configuring the Hashing Algorithm to Use Fields in the Layer 2 Header for Hashing	2924
	Configuring the Hashing Algorithm to Use Fields in the IP Payload for Hashing	2924
	Configuring the Hashing Algorithm to Use Fields in the IPv6 Payload for Hashing	2924
	Configuring Resilient Hashing for Trunk/ECMP Groups	2925
	Configuring Resilient Hashing on LAGs	2925
	Configuring Resilient Hashing on ECMP Groups	2926
Part 57	Uplink Failure Detection	
Chapter 121	Understanding Uplink Failure Detection	2929
	Overview of Uplink Failure Detection	2929
	Uplink Failure Detection Configuration	2929
	Failure Detection Pair	2930
	Configuring Interfaces for Uplink Failure Detection	2931
	Example: Configuring Interfaces for Uplink Failure Detection	2932

Part 58	Configuration Statements and Operational Commands	
Chapter 122	Ethernet OAM Link Fault Management Configuration Statements . . .	2939
	action (OAM LFM)	2940
	action-profile	2941
	allow-remote-loopback	2942
	ethernet (OAM LFM)	2943
	event-thresholds	2945
	event (OAM LFM)	2946
	frame-error	2946
	frame-period	2947
	frame-period-summary	2947
	oam	2948
	interface (OAM LFM)	2950
	link-adjacency-loss	2951
	link-discovery	2951
	link-down	2952
	link-event-rate	2952
	link-fault-management	2953
	negotiation-options	2954
	no-allow-link-events	2954
	pdu-interval	2955
	pdu-threshold	2955
	remote-loopback	2956
	symbol-period	2956
	syslog (OAM LFM)	2957
	traceoptions (OAM LFM)	2958
Chapter 123	GRE Configuration Statements	2961
	allow-fragmentation	2962
	copy-tos-to-outer-ip-header	2962
	destination (Tunnels)	2963
	do-not-fragment	2964
	family	2965
	routing-instance	2966
	source	2966
	tunnel	2967
	tunnel	2968
	unit (Interfaces)	2969
	tunnel-port	2970
Chapter 124	Interfaces Configuration Statements	2971
	address	2973
	alarm	2977
	auto-negotiation	2978
	channel-speed	2979
	configured-flow-control	2980
	craft-lockout	2981
	description	2982
	ethernet (Alarm)	2983

	ethernet-switch-profile	2984
	ethernet-switching	2986
	ether-options	2987
	eui-64	2988
	family	2989
	fibre-channel (Alarm)	2993
	filter	2994
	flow-control	2996
	fpc	2997
	gratuitous-arp-reply	2998
	hold-time (Physical Interface)	2999
	irb (Interfaces)	3001
	inet (interfaces)	3004
	inet6 (interfaces)	3005
	interface-mode	3006
	interface-range	3008
	interfaces	3010
	link-down	3017
	link-mode	3018
	link-speed	3019
	loopback (Aggregated Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet)	3020
	mac	3020
	management-ethernet (Alarm)	3021
	member	3022
	member-range	3023
	mtu	3024
	no-gratuitous-arp-request	3025
	pic	3026
	rx-buffers	3027
	source	3028
	speed	3029
	tag-protocol-id (TPIDs Expected to Be Sent or Received)	3030
	targeted-broadcast	3031
	traceoptions (Individual Interfaces)	3032
	tx-buffers	3033
	unit	3035
	vlan-id	3036
	vlan-tagging	3037
Chapter 125	IP Directed Broadcast Configuration Statement	3039
	targeted-broadcast	3039
Chapter 126	LAGs and LACP Configuration Statements	3041
	aggregated-devices	3042
	aggregated-ether-options	3043
	chassis	3045
	802.3ad	3046
	device-count	3047
	ethernet	3047
	force-up	3048

	lacp (802.3ad)	3049
	lacp (Aggregated Ethernet)	3050
	link-protection	3051
	periodic	3052
Chapter 127	Load Balancing Configuration Statements	3053
	hash-key (Forwarding Options)	3054
	multiservice	3056
Chapter 128	Local Link Bias Configuration Statement	3057
	local-bias	3057
Chapter 129	Redundant Trunk Groups Configuration Statements	3059
	group (Redundant Trunk Groups)	3060
	interface (Redundant Trunk Groups)	3061
	preempt-cutover-timer	3062
	redundant-trunk-group	3063
Chapter 130	Resilient Hashing Configuration Statements	3065
	ecmp-resilient-hash	3065
	enhanced-hash-key	3066
	hash-mode	3069
	hash-key (Forwarding Options)	3071
	hash-seed	3073
	inet (enhanced-hash-key)	3074
	inet6 (enhanced-hash-key)	3076
	ipv6-flow-label	3078
	resilient-hash	3078
Chapter 131	Uplink Failure Detection Configuration Statements	3079
	group	3079
	link-to-disable	3080
	link-to-monitor	3080
	uplink-failure-detection	3081
Chapter 132	Ethernet OAM Link Fault Management Operational Command	3083
	show oam ethernet link-fault-management	3084
Chapter 133	Interfaces Operational Commands	3089
	monitor interface	3090
	show interfaces diagnostics optics	3102
	show interfaces ge	3116
	show interfaces (GRE)	3131
	show interfaces irb	3140
	show interfaces queue	3146
	show interfaces xe	3185
Chapter 134	LAGs and LACP Operational Commands	3203
	show lacp interfaces	3204
	show lacp statistics interfaces (View)	3209

Chapter 135	Redundant Trunk Group Operational Command	3211
	show redundant-trunk-group	3212
Chapter 136	Resilient Hashing Operational Command	3215
	show forwarding-options enhanced-hash-key	3216
Chapter 137	Uplink Failure Detection Operational Command	3221
	show uplink-failure-detection	3222

Guide 9 Routing Options Feature Guide for QFX10000 Switches

Part 59	Configuring Routing Options	
	Overview of Routing Options	3227
	Configuring Static Routing	3228
	Configuring Per-Packet Load Balancing	3228
	Examples: Configuring Per-Packet Load Balancing	3230
	Using a QFX10000 Switch with 6PE	3231
	Understanding Virtual Router Routing Instances	3232
	Configuring Virtual Router Routing Instances	3232
	Understanding Distributed Periodic Packet Management	3233
	Configuring Distributed Periodic Packet Management	3234
	Disabling or Enabling Distributed Periodic Packet Management Globally	3234
	Disabling or Enabling Distributed Periodic Packet Management for LACP Packets	3235
	Understanding Bidirectional Forwarding Detection (BFD)	3236
Chapter 138	Examples: Configuring BFD for Static Routes	3237
	Understanding BFD for Static Routes for Faster Network Failure Detection	3237
	Example: Configuring BFD for Static Routes for Faster Network Failure Detection	3241
	Example: Enabling BFD on Qualified Next Hops in Static Routes for Route Selection	3246
Chapter 139	Example: Configuring BFD Authentication for Static Routes	3253
	Understanding BFD Authentication for Static Route Security	3253
	BFD Authentication Algorithms	3254
	Security Authentication Keychains	3254
	Strict Versus Loose Authentication	3255
	Example: Configuring BFD Authentication for Securing Static Routes	3255
	Monitoring Routing Information	?
	Verifying That Virtual Router Routing Instances Are Working	?
Chapter 140	Troubleshooting Virtual Routing Instances	3263
	Direct Routes Not Leaked Between Routing Instances	3263

Part 60

Chapter 141

Configuration Statements and Operational Commands

Configuration Statements (Routing Options)	3267
active	3270
aggregate (Routing)	3272
as-path (Routing Options)	3274
autonomous-system	3276
backup-pe-group	3278
backups	3279
bandwidth (Multicast Flow Map)	3280
bfd-liveness-detection (Routing Options Static Route)	3281
bgp-orf-cisco-mode	3285
bmp	3287
brief	3289
centralized	3290
community (Routing Options)	3291
confederation	3293
disable (Routing Options)	3294
description (Routing Instances)	3294
discard	3295
export	3296
export-rib	3297
fate-sharing	3300
flow	3301
flow-map	3302
forwarding-cache (Flow Maps)	3303
forwarding-cache (Multicast)	3304
forwarding-options (chassis)	3306
forwarding-table	3308
generate	3309
graceful-restart (Enabling Globally)	3311
import	3312
import-policy	3313
import-rib	3314
indirect-next-hop	3315
install (Routing Options)	3316
instance-export	3317
instance-import	3317
instance-type	3318
interface (Multicast Static Routes)	3319
interface (Routing Instances)	3320
interface (Routing Options)	3321
interface-routes	3322
local-address (Routing Options)	3323
martians	3324
maximum-bandwidth (Routing Options)	3325
maximum-paths	3326
maximum-prefixes	3328
med-igp-update-interval	3329

metric (Aggregate, Generated, or Static Route)	3330
multicast (Routing Options)	3331
no-qos-adjust	3332
options (Routing Options)	3333
pim-to-igmp-proxy	3334
pim-to-mld-proxy	3335
policy (Aggregate and Generated Routes)	3336
policy (Flow Maps)	3337
policy-options	3338
policy-statement	3339
ppm	3343
ppm (Ethernet Switching)	3344
preference (Routing Options)	3345
prefix	3346
protocols	3347
qualified-next-hop (Static Routes)	3350
readvertise	3352
redundant-sources	3353
resolution	3354
resolution-ribs	3355
resolve	3356
restart-duration (Routing Options)	3357
retain	3358
reverse-oif-mapping	3359
rpf-check-policy (Routing Options RPF)	3360
rib (General)	3361
rib (Route Resolution)	3363
rib-group (Routing Options)	3364
rib-groups	3365
route-distinguisher-id	3367
route-record	3368
router-id	3369
routing-instances	3370
routing-options	3371
scope	3371
scope-policy	3372
source (Source-Specific Multicast)	3373
source-routing	3374
ssm-groups	3375
ssm-map (Routing Options Multicast)	3376
static (Routes)	3377
subscriber-leave-timer	3379
tag (Routing Options)	3380
threshold (Multicast Forwarding Cache)	3381
timeout (Flow Maps)	3382
timeout (Multicast)	3383
traceoptions	3384
upstream-interface	3387

Chapter 142	Operational Commands (Routing Options)	3389
	clear ipv6 neighbors	3391
	show as-path	3392
	show as-path domain	3396
	show as-path summary	3399
	show ipv6 neighbors	3401
	show ipv6 router-advertisement	3403
	show route	3406
	show route active-path	3412
	show route all	3417
	show route aspath-regex	3419
	show route best	3421
	show route brief	3424
	show route community	3426
	show route community-name	3428
	show route damping	3430
	show route detail	3435
	show route exact	3453
	show route export	3455
	show route extensive	3458
	show route flow validation	3475
	show route forwarding-table	3477
	show route inactive-path	3491
	show route inactive-prefix	3494
	show route instance	3496
	show route label	3503
	show route label-switched-path	3506
	show route martians	3508
	show route next-hop	3510
	show route no-community	3516
	show route protocol	3519
	show route range	3529
	show route receive-protocol	3533
	show route resolution	3542
	show route snooping	3545
	show route source-gateway	3548
	show route summary	3554
	show route table	3558
	show route terse	3586

Guide 10 BGP Feature Guide for QFX10000 Switches

Part 61	Overview	
Chapter 143	BGP Overview	3593
	Understanding BGP	3593
	Autonomous Systems	3594
	AS Paths and Attributes	3594
	External and Internal BGP	3594

	Multiple Instances of BGP	3595
	BGP Routes Overview	3596
	BGP Messages Overview	3597
	Open Messages	3597
	Update Messages	3597
	Keepalive Messages	3598
	Notification Messages	3598
	Route-Refresh Messages	3598
Part 62	Configuring BGP	
Chapter 144	Basic BGP Configuration	3601
	Examples: Configuring External BGP Peering	3601
	Understanding External BGP Peering Sessions	3601
	Example: Configuring External BGP Point-to-Point Peer Sessions	3602
	Example: Configuring External BGP on Logical Systems with IPv6 Interfaces	3609
	Examples: Configuring Internal BGP Peering	3624
	Understanding Internal BGP Peering Sessions	3624
	Example: Configuring Internal BGP Peer Sessions	3625
	Example: Configuring Internal BGP Peering Sessions on Logical Systems	3636
	Configuring BGP Monitoring Protocol Version 3	3647
Chapter 145	BGP Path Attribute Configuration	3651
	Understanding the Advertisement of Multiple Paths to a Single Destination in BGP	3651
	Example: Configuring BGP Local Preference	3652
	Understanding the Local Preference Metric for Internal BGP Routes	3652
	Example: Configuring the Local Preference Value for BGP Routes	3653
	Examples: Configuring BGP MED	3665
	Understanding the MED Attribute That Determines the Exit Point in an AS	3665
	Example: Configuring the MED Attribute That Determines the Exit Point in an AS	3668
	Example: Configuring the MED Using Route Filters	3680
	Example: Configuring the MED Using Communities	3693
	Example: Associating the MED Path Attribute with the IGP Metric and Delaying MED Updates	3694
	Examples: Configuring BGP Local AS	3704
	Understanding the BGP Local AS Attribute	3704
	Example: Configuring a Local AS for EBGp Sessions	3709
	Example: Configuring a Private Local AS for EBGp Sessions	3719
	Example: Configuring the Accumulated IGP Attribute for BGP	3724
	Understanding the Accumulated IGP Attribute for BGP	3725
	Example: Configuring the Accumulated IGP Attribute for BGP	3725

Chapter 146	BGP Policy Configuration	3765
	Example: Configuring BGP Interactions with IGPs	3765
	Understanding Routing Policies	3765
	Example: Injecting OSPF Routes into the BGP Routing Table	3766
	Example: Configuring BGP Route Advertisement	3769
	Configuring Routing Policies to Control BGP Route Advertisements	3769
	Applying Routing Policy	3769
	Setting BGP to Advertise Inactive Routes	3770
	Configuring BGP to Advertise the Best External Route to Internal Peers	3770
	Configuring How Often BGP Exchanges Routes with the Routing Table	3772
	Disabling Suppression of Route Advertisements	3773
	Example: Configuring BGP Prefix-Based Outbound Route Filtering	3773
	Example: Configuring EBGp Multihop	3777
	Understanding BGP Multihop for Peering	3777
	Example: Configuring EBGp Multihop Sessions	3777
	Example: Configuring BGP Route Preference (Administrative Distance)	3786
	Understanding Route Preference Values (Administrative Distance)	3786
	Example: Configuring the Preference Value for BGP Routes	3787
	Example: Configuring BGP Path Selection	3793
	Understanding BGP Path Selection	3793
	Routing Table Path Selection	3795
	Effects of Advertising Multiple Paths to a Destination	3796
	Example: Ignoring the AS Path Attribute When Selecting the Best Path	3796
	Example: Removing Private AS Numbers	3803
	Understanding Private AS Number Removal from AS Paths	3804
	Example: Removing Private AS Numbers from AS Paths	3805
Chapter 147	BGP BFD Configuration	3811
	Example: Configuring BFD for BGP	3811
	Understanding BFD for BGP	3811
	Example: Configuring BFD on Internal BGP Peer Sessions	3812
	Example: Configuring BFD Authentication for BGP	3820
	Understanding BFD Authentication for BGP	3820
	BFD Authentication Algorithms	3821
	Security Authentication Keychains	3822
	Strict Versus Loose Authentication	3822
	Example: Configuring BFD Authentication for BGP	3822
	Configuring BFD Authentication Parameters	3822
	Viewing Authentication Information for BFD Sessions	3824

Chapter 148	BGP Load Balancing Configuration	3827
	Examples: Configuring BGP Multipath	3827
	Understanding BGP Multipath	3827
	Example: Load Balancing BGP Traffic	3828
	Example: Configuring Single-Hop EBGP Peers to Accept Remote Next Hops	3833
	Example: Advertising Multiple Paths in BGP	3844
	Configuring ECMP Next Hops for RSVP and LDP LSPs for Load Balancing . . .	3870
Chapter 149	IBGP Scaling Configuration	3873
	Example: Configuring BGP Route Reflectors	3873
	Understanding BGP Route Reflectors	3873
	Example: Configuring a Route Reflector	3876
	Example: Configuring BGP Confederations	3890
	Understanding BGP Confederations	3890
	Example: Configuring BGP Confederations	3891
Chapter 150	BGP Security Configuration	3897
	Example: Configuring BGP Route Authentication	3897
	Understanding Route Authentication	3897
	Example: Configuring Route Authentication for BGP	3898
	Examples: Configuring TCP and BGP Security	3904
	Understanding Security Options for BGP with TCP	3904
	Example: Configuring a Filter to Block TCP Access to a Port Except from Specified BGP Peers	3904
	Example: Configuring a Filter to Limit TCP Access to a Port Based On a Prefix List	3909
	Example: Limiting TCP Segment Size for BGP	3912
Chapter 151	BGP Flap Configuration	3919
	Example: Preventing BGP Session Resets	3919
	Understanding BGP Session Resets	3919
	Example: Preventing BGP Session Flaps When VPN Families Are Configured	3919
	Examples: Configuring BGP Flap Damping	3926
	Understanding Damping Parameters	3926
	Example: Configuring BGP Route Flap Damping Parameters	3927
	Example: Configuring BGP Route Flap Damping Based on the MBGP MVPN Address Family	3936
Part 63	Monitoring and Troubleshooting	
Chapter 152	BGP Monitoring Configuration	3949
	Example: Configuring BGP Trace Operations	3949
	Understanding Trace Operations for BGP Protocol Traffic	3949
	Example: Viewing BGP Trace Files on Logical Systems	3951
	Tracing BMP Operations	3956
Chapter 153	Routine Monitoring	3959
	Monitoring BGP Routing Information	3959

Part 64	Configuration Statement and Operational Commands	
Chapter 154	Configuration Statements	3963
	accept-remote-nexthop	3966
	advertise-external	3967
	advertise-inactive	3969
	advertise-peer-as	3970
	algorithm (BGP BFD Authentication)	3971
	apply-groups	3973
	apply-groups-except	3973
	authentication (BGP BFD Liveness Detection)	3974
	authentication-algorithm	3976
	authentication-key (Protocols BGP and BMP)	3979
	authentication-key-chain (Protocols BGP and BMP)	3980
	bfd-liveness-detection (Protocols BGP)	3981
	bgp	3985
	bgp-orf-cisco-mode	3986
	cluster	3988
	connection-mode	3989
	damping (Protocols BGP)	3990
	description (Protocols BGP)	3992
	detection-time (BFD Liveness Detection)	3993
	disable (Protocols BGP)	3994
	disable (BGP Graceful Restart)	3995
	export (Protocols BGP)	3996
	family	3997
	graceful-restart (Protocols BGP)	4001
	group (Protocols BGP)	4003
	hold-down	4006
	hold-down-interval (BGP BFD Liveness Detection)	4008
	hold-time (Protocols BGP)	4010
	import	4012
	include-mp-next-hop	4014
	initiation-message	4015
	keep	4016
	key-chain (BGP BFD Authentication)	4018
	local-address (Protocols BGP)	4020
	local-address (Protocols BMP)	4022
	local-as	4023
	local-port	4025
	local-preference	4026
	log-updown	4027
	loops (BGP Address Family)	4028
	loose-check (BGP BFD Authentication)	4030
	maximum-ecmp	4031
	metric-out	4032
	minimum-interval (BFD Liveness Detection)	4034
	minimum-interval (transmit-interval)	4036
	minimum-receive-interval (BFD Liveness Detection)	4038

monitor (Protocols BMP)	4039
mtu-discovery	4040
multihop	4042
multiplier (BFD Liveness Detection)	4044
neighbor (Protocols BGP)	4046
no-adaptation (BFD Liveness Detection)	4050
no advertise-peer-as	4051
no-aggregator-id	4052
no-client-reflect	4053
out-delay	4054
outbound-route-filter	4056
passive (Protocols BGP)	4057
path-selection	4058
peer-as (Protocols BGP)	4060
post-policy	4061
pre-policy	4062
preference (Protocols BGP)	4063
priority (Protocols BMP)	4064
remove-private	4065
restart-time (BGP Graceful Restart)	4067
route-monitoring	4068
session-mode	4070
stale-routes-time	4071
station	4072
station-address	4073
station-port	4074
statistics-timeout	4075
tcp-mss (Protocols BGP)	4076
threshold (detection-time)	4077
threshold (transmit-interval)	4079
traceoptions (Protocols BGP)	4081
traceoptions (Protocols BMP)	4084
transmit-interval (BFD Liveness Detection)	4086
version (BFD Liveness Detection)	4088
Chapter 155	
Operational Commands	4091
clear bgp damping	4092
clear bgp neighbor	4093
clear bgp table	4096
show bgp bmp	4098
show bgp group	4100
show bgp neighbor	4107
show bgp summary	4124
show policy damping	4130
show route damping	4132
show route detail	4137

Guide 11	IS-IS Feature Guide for QFX10000 Switches	
Chapter 156	Overview	4157
	IS-IS Overview	4157
	IS-IS Terminology	4158
	ISO Network Addresses	4158
	IS-IS Packets	4160
	Persistent Route Reachability	4161
	IS-IS Support for Multipoint Network Clouds	4161
	Installing a Default Route to the Nearest Routing Device That Operates at Both IS-IS Levels	4161
Part 65	Configuring IS-IS Features	
Chapter 157	Configuring IS-IS	4165
	Understanding Loop-Free Alternate Routes for IS-IS	4165
	Configuring Link Protection for IS-IS	4167
	Configuring Node-Link Protection for IS-IS	4168
	Excluding an IS-IS Interface as a Backup for Protected Interfaces	4168
	Configuring RSVP Label-Switched Paths as Backup Paths for IS-IS	4168
	Using Operational Mode Commands to Monitor Protected IS-IS Routes	4169
	Example: Configuring IS-IS	4169
	Example: Configuring a Multi-Level IS-IS Topology to Control Interarea Flooding	4175
	Configuring IS-IS Authentication	4183
	Configuring IS-IS Authentication Without Network-Wide Deployment	4184
	Example: Redistributing OSPF Routes into IS-IS	4185
	Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies	4193
	Understanding IS-IS IPv4 and IPv6 Unicast Topologies	4193
	Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies	4193
	Example: Configuring IS-IS Multicast Topology	4201
	IS-IS Multicast Topologies Overview	4202
	Example: Configuring IS-IS Multicast Topology	4203
	Example: Configuring Link and Node Protection for IS-IS Routes	4217
	Understanding Loop-Free Alternate Routes for IS-IS	4218
	Configuring Link Protection for IS-IS	4219
	Configuring Node-Link Protection for IS-IS	4220
	Excluding an IS-IS Interface as a Backup for Protected Interfaces	4220
	Configuring RSVP Label-Switched Paths as Backup Paths for IS-IS	4221
	Using Operational Mode Commands to Monitor Protected IS-IS Routes	4221
	Example: Configuring Node-Link Protection for IS-IS Routes in a Layer 3 VPN	4221
	Example: Configuring an IS-IS Default Route Policy on Logical Systems	4232
	Example: Configuring IS-IS for CLNS	4237
	Understanding IS-IS for CLNS	4237
	Example: Configuring IS-IS for CLNS	4237
	Example: Enabling Packet Checksums on IS-IS Interfaces for Error Checking	4239

Chapter 158	Configuring BFD Authentication for IS-IS	4243
	Understanding BFD Authentication for IS-IS	4243
	BFD Authentication Algorithms	4243
	Security Authentication Keychains	4244
	Strict Versus Loose Authentication	4244
	Example: Configuring BFD for IS-IS	4245
	Example: Configuring BFD Authentication for IS-IS	4250
Chapter 159	Configuring Hitless Authentication Key Rollover for IS-IS	4255
	Understanding Hitless Authentication Key Rollover for IS-IS	4255
	Example: Configuring Hitless Authentication Key Rollover for IS-IS	4256
Chapter 160	Configuration Statements	4261
	authentication-key (Protocols IS-IS)	4263
	authentication-key-chain (Protocols IS-IS)	4264
	authentication-type (Protocols IS-IS)	4265
	bfd-liveness-detection (Protocols IS-IS)	4266
	checksum (Protocols IS-IS)	4268
	csnp-interval	4269
	disable (Protocols IS-IS)	4270
	export	4271
	external-preference (Protocols IS-IS)	4272
	family (Protocols IS-IS)	4273
	graceful-restart (Protocols IS-IS)	4274
	hello-authentication-key	4275
	hello-authentication-key-chain	4276
	hello-authentication-type	4277
	hello-interval (Protocols IS-IS)	4278
	hello-padding	4280
	hold-time (Protocols IS-IS)	4282
	ignore-attached-bit	4283
	interface	4284
	ipv4-multicast	4286
	ipv4-multicast-metric	4287
	ipv6-multicast	4288
	ipv6-multicast-metric	4288
	ipv6-unicast	4289
	ipv6-unicast-metric	4290
	isis	4291
	level (Global IS-IS)	4292
	link-protection (Protocols IS-IS)	4293
	loose-authentication-check	4293
	lsp-interval	4294
	lsp-lifetime	4295
	max-areas	4296
	mesh-group (Protocols IS-IS)	4297
	metric (Protocols IS-IS)	4298
	no-adjacency-holddown	4299
	no-authentication-check	4300

no-csnp-authentication	4300
no-eligible-backup (Protocols IS-IS)	4301
no-hello-authentication	4301
no-ipv4-multicast	4302
no-ipv4-routing	4303
no-ipv6-multicast	4304
no-ipv6-routing	4305
no-ipv6-unicast	4306
no-psnp-authentication	4306
no-unicast-topology	4307
node-link-protection (Protocols IS-IS)	4308
overload (Protocols IS-IS)	4309
passive (Protocols IS-IS)	4312
point-to-point	4313
preference (Protocols IS-IS)	4314
prefix-export-limit (Protocols IS-IS)	4315
priority (Protocols IS-IS)	4316
reference-bandwidth (Protocols IS-IS)	4317
rib-group (Protocols IS-IS)	4318
spf-options (Protocols IS-IS)	4319
topologies (Protocols IS-IS)	4320
traceoptions (Protocols IS-IS)	4321
traffic-engineering (Protocols IS-IS)	4324
wide-metrics-only	4327
Chapter 161	
Operational Commands	4329
clear isis adjacency	4330
clear isis database	4332
clear isis overload	4334
clear isis statistics	4336
show isis adjacency	4338
show isis authentication	4342
show isis backup coverage	4344
show isis backup label-switched-path	4346
show isis backup spf results	4348
show isis database	4352
show isis hostname	4359
show isis interface	4361
show isis overview	4365
show isis route	4369
show isis spf	4373
show isis statistics	4378

Guide 12 OSPF Feature Guide for QFX10000 Switches

Part 66**Overview****Chapter 162****OSPF Overview 4385**

OSPF Overview	4386
OSPF Default Route Preference Values	4388
OSPF Routing Algorithm	4388
OSPF Three-Way Handshake	4389
OSPF Version 3	4390
OSPF Packets Overview	4391
OSPF Packet Header	4391
Hello Packets	4391
Database Description Packets	4392
Link-State Request Packets	4392
Link-State Update Packets	4392
Link-State Acknowledgment Packets	4392
Link-State Advertisement Packet Types	4393
Understanding OSPF External Metrics	4393

Part 67**Configuring OSPF****Chapter 163****Basic OSPF Area Configuration 4397**

Understanding OSPF Areas	4397
Areas	4397
Area Border Routers	4398
Backbone Areas	4398
AS Boundary Routers	4398
Backbone Router	4398
Internal Router	4398
Stub Areas	4398
Not-So-Stubby Areas	4399
Transit Areas	4399
Examples: Configuring OSPF Designated Routers	4399
OSPF Designated Router Overview	4400
Example: Configuring an OSPF Router Identifier	4400
Example: Controlling OSPF Designated Router Election	4402
Examples: Configuring OSPF Areas	4404
Understanding OSPF Areas and Backbone Areas	4404
Example: Configuring a Single-Area OSPF Network	4406
Example: Configuring a Multiarea OSPF Network	4408

Chapter 164	Advanced OSPF Area Configuration	4413
	Examples: Configuring OSPF Stub and Not-So-Stubby Areas	4413
	Understanding OSPF Stub Areas, Totally Stubby Areas, and Not-So-Stubby Areas	4413
	Example: Configuring OSPF Stub and Totally Stubby Areas	4414
	Example: Configuring OSPF Not-So-Stubby Areas	4418
	Example: Configuring OSPF Multiarea Adjacency	4423
	Multiarea Adjacency for OSPF	4423
	Example: Configuring Multiarea Adjacency for OSPF	4424
	Example: Disabling OSPFv2 Compatibility with RFC 1583	4428
	OSPFv2 Compatibility with RFC 1583 Overview	4428
	Example: Disabling OSPFv2 Compatibility with RFC 1583	4428
Chapter 165	OSPF Interface Configuration	4431
	Examples: Configuring OSPF Interfaces	4431
	About OSPF Interfaces	4431
	Example: Configuring an Interface on a Broadcast or Point-to-Point Network	4432
	Example: Configuring an OSPFv2 Interface on a Nonbroadcast Multiaccess Network	4435
	Example: Configuring an OSPFv2 Interface on a Point-to-Multipoint Network	4438
	Example: Configuring OSPF Demand Circuits	4439
	Example: Configuring a Passive OSPF Interface	4442
	Example: Configuring OSPFv2 Peer interfaces	4444
	Example: Configuring Multiple Address Families for OSPFv3	4446
	Understanding Multiple Address Families for OSPFv3	4446
	Example: Configuring Multiple Address Families for OSPFv3	4446
Chapter 166	OSPF Route Control Configuration	4451
	Examples: Configuring OSPF Route Summarization	4451
	Understanding OSPF Route Summarization	4451
	Example: Summarizing Ranges of Routes in OSPF Link-State Advertisements	4452
	Example: Limiting the Number of Prefixes Exported to OSPF	4457
	Configuring OSPF Refresh and Flooding Reduction in Stable Topologies	4459
	Examples: Configuring OSPF Traffic Control	4460
	Understanding OSPF Traffic Control	4461
	Controlling the Cost of Individual OSPF Network Segments	4461
	Dynamically Adjusting OSPF Interface Metrics Based on Bandwidth	4462
	Controlling OSPF Route Preferences	4462
	Example: Controlling the Cost of Individual OSPF Network Segments	4462
	Example: Dynamically Adjusting OSPF Interface Metrics Based on Bandwidth	4466
	Example: Controlling OSPF Route Preferences	4468
	Example: Configuring OSPF Overload Mode	4470
	OSPF Overload Function Overview	4471
	Example: Configuring OSPF to Make Routing Devices Appear Overloaded	4472

Chapter 167	OSPF Fault Detection Configuration	4475
	Example: Configuring OSPF Timers	4475
	OSPF Timers Overview	4475
	Example: Configuring OSPF Timers	4476
	Example: Configuring BFD for OSPF	4482
	Understanding BFD for OSPF	4482
	Example: Configuring BFD for OSPF	4484
	Example: Configuring BFD Authentication for OSPF	4488
	Understanding BFD Authentication for OSPF	4488
	BFD Authentication Algorithms	4489
	Security Authentication Keychains	4490
	Strict Versus Loose Authentication	4490
	Configuring BFD Authentication for OSPF	4490
	Configuring BFD Authentication Parameters	4490
	Viewing Authentication Information for BFD Sessions	4492
Chapter 168	OSPF Redundancy Features Configuration	4495
	Examples: Configuring Graceful Restart for OSPF	4495
	Graceful Restart for OSPF Overview	4495
	Helper Mode for Graceful Restart	4496
	Planned and Unplanned Graceful Restart	4496
	Example: Configuring Graceful Restart for OSPF	4497
	Example: Configuring the Helper Capability Mode for OSPFv2 Graceful Restart	4501
	Example: Configuring the Helper Capability Mode for OSPFv3 Graceful Restart	4504
	Example: Disabling Strict LSA Checking for OSPF Graceful Restart	4507
Chapter 169	OSPF Traffic Engineering Configuration	4511
	Examples: Configuring OSPF Traffic Engineering	4511
	OSPF Support for Traffic Engineering	4511
	Example: Enabling OSPF Traffic Engineering Support	4513
	Example: Configuring the Traffic Engineering Metric for a Specific OSPF Interface	4518
	Example: Configuring OSPF Passive Traffic Engineering Mode	4519
	OSPF Passive Traffic Engineering Mode	4519
	Example: Configuring OSPF Passive Traffic Engineering Mode	4520
Chapter 170	OSPF Database Protection Configuration	4523
	Example: Configuring OSPF Database Protection	4523
	OSPF Database Protection Overview	4523
	Configuring OSPF Database Protection	4524
Chapter 171	OSPF Policy Configuration	4527
	Examples: Configuring OSPF Routing Policy	4527
	Understanding OSPF Routing Policy	4527
	Routing Policy Terms	4528
	Routing Policy Match Conditions	4528

	Routing Policy Actions	4529
	Example: Injecting OSPF Routes into the BGP Routing Table	4529
	Example: Redistributing Static Routes into OSPF	4532
	Example: Configuring an OSPF Import Policy	4535
	Example: Configuring a Route Filter Policy to Specify Priority for Prefixes Learned Through OSPF	4539
	Examples: Configuring Routing Policy for Network Summaries	4543
	Import and Export Policies for Network Summaries Overview	4543
	Example: Configuring an OSPF Export Policy for Network Summaries	4544
	Example: Configuring an OSPF Import Policy for Network Summaries	4554
Part 68	Monitoring and Troubleshooting	
Chapter 172	Routine Monitoring	4567
	Monitoring OSPF Routing Information	4567
Chapter 173	OSPF Monitoring Configuration	4569
	Example: Configuring OSPF Trace Options	4569
	Tracing OSPF Protocol Traffic	4569
	Example: Tracing OSPF Protocol Traffic	4570
Part 69	Configuration Statements and Operational Commands	
Chapter 174	OSPF Configuration Statements	4579
	area	4580
	area-range	4582
	authentication	4584
	context-identifier (Protocols OSPF)	4585
	bfd-liveness-detection	4586
	database-protection	4590
	disable (OSPF)	4592
	export	4594
	external-preference	4595
	graceful-restart	4596
	import	4598
	interface (Protocols OSPF)	4599
	no-nssa-abr	4601
	no-rfc-1583	4602
	ospf	4603
	ospf3	4603
	overload (Protocols OSPF)	4604
	preference (Protocols OSPF)	4606
	prefix-export-limit (Protocols OSPF)	4607
	reference-bandwidth (Protocols OSPF)	4608
	rib-group (Protocols OSPF)	4609
	topology (OSPF)	4610
	traceoptions (Protocols OSPF)	4611
	traffic-engineering (OSPF)	4614

Chapter 175	OSPF Operational Commands	4617
	clear (ospf ospf3) database	4618
	clear (ospf ospf3) database-protection	4621
	clear (ospf ospf3) io-statistics	4622
	clear (ospf ospf3) neighbor	4623
	clear (ospf ospf3) statistics	4625
	clear (ospf ospf3) overload	4627
	show (ospf ospf3) backup coverage	4628
	show (ospf ospf3) backup neighbor	4631
	show ospf context-identifier	4633
	show ospf database	4635
	show (ospf ospf3) interface	4643
	show (ospf ospf3) io-statistics	4649
	show (ospf ospf3) log	4651
	show (ospf ospf3) neighbor	4654
	show (ospf ospf3) overview	4660
	show (ospf ospf3) route	4665
	show (ospf ospf3) statistics	4671

Guide 13 RIP Feature Guide for QFX10000 Switches

Chapter 176	Overview	4677
	RIP Overview	4677
	Distance-Vector Routing Protocols	4677
	RIP Protocol Overview	4678
	RIP Packets	4679
	Maximizing Hop Count	4680
	Split Horizon and Poison Reverse Efficiency Techniques	4680
	Limitations of Unidirectional Connectivity	4681
Part 70	Configuring RIP	
Chapter 177	Example: Configuring RIP	4685
	Understanding Basic RIP Routing	4685
	Example: Configuring a Basic RIP Network	4685
Chapter 178	Example: Configuring Authentication for RIP Routes	4693
	Understanding RIP Authentication	4693
	Example: Configuring Route Authentication for RIP	4693
	Enabling Authentication with Plain-Text Passwords (CLI Procedure)	4698
	Enabling Authentication with MD5 Authentication (CLI Procedure)	4699
Chapter 179	Example: Configuring BFD for RIP	4701
	Understanding BFD for RIP	4701
	Example: Configuring BFD for RIP	4702
Chapter 180	Example: Configuring BFD Authentication for RIP	4709
	Understanding BFD Authentication for RIP	4709
	BFD Authentication Algorithms	4709
	Security Authentication Keychains	4710

	Strict Versus Loose Authentication	4710
	Example: Configuring BFD Authentication for RIP	4711
Chapter 181	Example: Applying Policies to RIP Routes Imported from Neighbors . . .	4719
	Understanding RIP Import Policy	4719
	Example: Applying Policies to RIP Routes Imported from Neighbors	4719
Chapter 182	Examples: Controlling Traffic with Metrics in a RIP Network	4725
	Understanding Traffic Control with Metrics in a RIP Network	4725
	Example: Controlling Traffic in a RIP Network with an Incoming Metric	4726
	Example: Controlling Traffic in a RIP Network with an Outgoing Metric	4728
	Example: Configuring the Metric Value Added to Imported RIP Routes	4729
Chapter 183	Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets	4735
	Understanding the Sending and Receiving of RIPv1 and RIPv2 Packets	4735
	Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets . .	4735
Chapter 184	Example: Redistributing Routes Among RIP Instances	4741
	Understanding Route Redistribution Among RIP instances	4741
	Example: Redistributing Routes Between Two RIP Instances	4742
Chapter 185	Example: Configuring RIP Timers	4749
	Understanding RIP Timers	4749
	Example: Configuring RIP Timers	4750
Chapter 186	Example: Tracing RIP Protocol Traffic	4757
	Understanding RIP Trace Operations	4757
	Example: Tracing RIP Protocol Traffic	4758
	Monitoring RIP Routing Information	?
Chapter 187	Configuration Statements	4763
	any-sender	4764
	authentication-key	4765
	authentication-type (Protocols RIP)	4766
	bfd-liveness-detection	4767
	check-zero	4769
	export	4770
	group (Protocols RIP)	4771
	holddown (Protocols RIP)	4773
	import (Protocols RIP)	4774
	message-size	4775
	metric-in (Protocols RIP)	4776
	metric-out	4777
	neighbor	4778
	preference (Protocols RIP)	4779
	receive (Protocols RIP)	4780
	rib-group (Protocols RIP)	4781
	rip	4781
	route-timeout (Protocols RIP)	4782
	send (Protocols RIP)	4783
	traceoptions (Protocols RIP)	4784

	update-interval (Protocols RIP)	4787
Chapter 188	Operational Commands	4789
	clear rip general-statistics	4790
	clear rip statistics	4791
	show rip general-statistics	4792
	show rip neighbor	4794
	show rip statistics	4796

Guide 14 MPLS Feature Guide for QFX10000 Switches

Part 71	LDP	
Chapter 189	Using LDP	4803
	LDP Introduction	4804
	Junos OS LDP Protocol Implementation	4804
	LDP Operation	4804
	Tunneling LDP LSPs in RSVP LSPs	4805
	Tunneling LDP LSPs in RSVP LSPs Overview	4805
	Label Operations	4806
	LDP Message Types	4807
	Discovery Messages	4807
	Session Messages	4808
	Advertisement Messages	4808
	Notification Messages	4808
	LDP Session Protection	4808
	LDP Graceful Restart	4809
	Minimum LDP Configuration	4810
	Enabling and Disabling LDP	4810
	Enabling Strict Targeted Hello Messages for LDP	4810
	Filtering Inbound LDP Label Bindings	4811
	Examples: Filtering Inbound LDP Label Bindings	4812
	Filtering Outbound LDP Label Bindings	4813
	Examples: Filtering Outbound LDP Label Bindings	4814
	Specifying the Transport Address Used by LDP	4815
	Collecting LDP Statistics	4815
	LDP Statistics Output	4816
	Disabling LDP Statistics on the Penultimate-Hop Router	4816
	LDP Statistics Limitations	4817
	Tracing LDP Protocol Traffic	4817
	Tracing LDP Protocol Traffic at the Protocol and Routing Instance Levels	4818
	Tracing LDP Protocol Traffic Within FECs	4819
	Examples: Tracing LDP Protocol Traffic	4819
	Example: Configuring LDP Downstream on Demand	4820
	Configuring the LDP Timer for Hello Messages	4825
	Configuring the LDP Timer for Link Hello Messages	4826
	Configuring the LDP Timer for Targeted Hello Messages	4826

Configuring the Delay Before LDP Neighbors Are Considered Down	4826
Configuring the LDP Hold Time for Link Hello Messages	4827
Configuring the LDP Hold Time for Targeted Hello Messages	4827
Configuring the Interval for LDP Keepalive Messages	4827
Configuring the LDP Keepalive Timeout	4828
Configuring LDP Route Preferences	4828
Configuring LDP Graceful Restart	4828
Enabling Graceful Restart	4829
Disabling LDP Graceful Restart or Helper Mode	4829
Configuring Reconnect Time	4830
Configuring Recovery Time and Maximum Recovery Time	4830
Configuring the Prefixes Advertised into LDP from the Routing Table	4831
Example: Configuring the Prefixes Advertised into LDP	4831
Configuring LDP LSP Traceroute	4832
Configuring Miscellaneous LDP Properties	4833
Configuring LDP to Use the IGP Route Metric	4833
Preventing Addition of Ingress Routes to the inet.0 Routing Table	4833
Multiple-Instance LDP and Carrier-of-Carriers VPNs	4834
Configuring MPLS and LDP to Pop the Label on the Ultimate-Hop Router	4834
Enabling LDP over RSVP-Established LSPs	4834
Enabling LDP over RSVP-Established LSPs in Heterogeneous Networks	4835
Configuring the TCP MD5 Signature for LDP Sessions	4835
Configuring LDP Session Protection	4836
Disabling SNMP Traps for LDP	4837
Configuring LDP Synchronization with the IGP on LDP Links	4837
Configuring LDP Synchronization with the IGP on the Router	4838
Configuring the Label Withdrawal Timer	4838
Ignoring the LDP Subnet Check	4838
Chapter 190 Configuration Statements and Monitoring Commands for LDP	4841
Configuration Statements for LDP	4841
allow-subnet-mismatch	4842
authentication-algorithm	4843
authentication-key (Protocols LDP)	4845
authentication-key-chain (Protocols LDP)	4846
deaggregate	4847
disable (Protocols LDP)	4848
dod-request-policy	4849
downstream-on-demand	4849
ecmp	4850
egress-policy	4850
explicit-null (Protocols LDP)	4851
export (Protocols LDP)	4851
fec	4852
graceful-restart (Protocols LDP)	4853
hello-interval (Protocols LDP)	4854
helper-disable (LDP)	4855
hold-time (Protocols LDP)	4856

ignore-lsp-metrics	4857
igp-synchronization	4857
import (Protocols LDP)	4858
interface (Protocols LDP)	4859
keepalive-interval	4860
keepalive-timeout	4861
l2-smart-policy	4861
label-withdrawal-delay	4862
ldp	4863
ldp-synchronization	4866
ldp-tunneling	4866
log-updown (Protocols LDP)	4867
maximum-neighbor-recovery-time	4868
no-forwarding	4869
policing (Protocols LDP)	4870
preference (Protocols LDP)	4871
reconnect-time	4872
recovery-time	4873
session (ldp)	4874
session-protection	4875
strict-targeted-hellos	4875
targeted-hello	4876
traceoptions (Protocols LDP)	4877
track-igp-metric	4879
traffic-statistics (Protocols LDP)	4880
transport-address	4882
Monitoring Commands for LDP	4882
clear ldp neighbor	4884
clear ldp session	4885
clear ldp statistics	4886
ping mpls ldp	4887
show ldp database	4890
show ldp fec-filters	4899
show ldp interface	4900
show ldp neighbor	4902
show ldp path	4904
show ldp route	4906
show ldp session	4910
show ldp statistics	4916
show ldp traffic-statistics	4920
traceroute mpls ldp	4924

Part 72

Chapter 191

MPLS

Using MPLS 4929

MPLS Overview For QFX Series and EX4600 Switches 4930

Why Use MPLS? 4931

Why Not Use MPLS? 4931

How Do I Configure MPLS?	4931
Configure the MPLS LER (Ingress) Switch and the Egress Switch . . .	4932
Configure LSRs for MPLS	4932
What Does the MPLS Protocol Do?	4932
How Does MPLS Interface to Other Protocols?	4933
If I Have Used Cisco MPLS, What Do I Need to Know?	4933
MPLS Feature Support on QFX Series and EX4600 Switches	4934
MPLS Commands Supported by QFX Series and EX4600 Switches . . .	4935
MPLS Features Supported by QFX Series and EX4600 Switches	4935
MPLS Limitations on QFX Series and EX4600 Switches	4943
MPLS Limitations on QFX3500 Switches	4943
MPLS Limitations on QFX5100 and EX4600 Switches	4944
MPLS Limitations on QFX5100 Virtual Chassis and Virtual Chassis Fabric	4945
MPLS Limitations on QFX10000 Switches	4946
Understanding MPLS Components for QFX Series and EX4600 Switches . .	4946
Provider Edge Switches	4946
MPLS Protocol and Label-Switched Paths	4947
IP Over MPLS for Customer Edge Interfaces	4947
BGP Layer 3 VPN Configuration	4947
Routing Instances for Layer 3 VPN	4947
Routing Instances for Layer 2 VPN and Layer 3 VPN	4948
Ethernet Encapsulation for Layer 2 VPN	4948
Provider Switch	4948
Components Required for All Switches in the MPLS Network	4948
Interior Gateway Protocol	4948
Traffic Engineering	4949
MPLS Protocol	4949
RSVP	4949
Family mpls	4950
Understanding MPLS Label Operations	4951
MPLS Label-Switched Paths and MPLS Labels	4951
Reserved Labels	4952
MPLS Label Operations	4952
Penultimate-Hop Popping and Ultimate-Hop Popping	4954
Understanding BGP	4955
Autonomous Systems	4955
AS Paths and Attributes	4955
External and Internal BGP	4956
Multiple Instances of BGP	4956
IPv6 Layer 3 VPNs	4957
Ethernet Pseudowire Overview	4958
Understanding CoS MPLS EXP Classifiers and Rewrite Rules	4959
EXP Classifiers	4960
EXP Rewrite Rules	4961
Schedulers	4962
Understanding Ethernet-over-MPLS (L2 Circuit)	4962
Ethernet-over-MPLS in Data Centers	4962

Understanding Using MPLS-Based Layer 3 VPNs on Switches	4963
MPLS-Based Layer 3 VPNs	4963
Carrier-of-Carriers VPNs	4965
Internet Service Provider as the Customer	4966
VPN Service Provider as the Customer	4966
Interprovider and Carrier-of-Carriers VPNs	4966
Fast Reroute Overview	4967
Graceful Restart and MPLS-Related Protocols	4969
LDP	4970
RSVP	4970
CCC and TCC	4970
Types of LSPs	4971
Configuring CoS Bits for an MPLS Network	4972
Configuring Ethernet over MPLS (L2 Circuit)	4972
Configuring the Local PE Switch for Port-Based Layer 2 Circuit (Pseudo-wire)	4973
Configuring the Remote PE Switch for Port-Based Layer 2 Circuit (Pseudo-wire)	4974
Configuring the Local PE Switch for VLAN-Based Layer 2 Circuit	4975
Configuring the Remote PE Switch for VLAN-Based Layer 2 Circuit	4975
Configuring a Global MPLS EXP Classifier	4977
Configuring MPLS Firewall Filters and Policers	4977
Configuring an MPLS Firewall Filter	4979
Applying an MPLS Firewall Filter to an MPLS Interface	4979
Configuring Policers for LSPs	4980
Configuring MPLS to Gather Statistics	4980
Configuring MPLS on Provider Edge Switches	4981
Configuring the Ingress PE Switch	4982
Configuring the Egress PE Switch	4983
Configuring MPLS on Provider Switches	4985
Configuring Reporting of Automatic Bandwidth Allocation Statistics for LSPs	4986
Configuring Static Label Switched Paths for MPLS	4989
Configuring the Ingress PE Switch	4990
Configuring the Provider and the Egress PE Switch	4991
Configuring Rewrite Rules for MPLS EXP Classifiers	4992
Example: Configuring MPLS-Based Layer 3 VPNs	4993
Example: Tunneling IPv6 Traffic over MPLS IPv4 Networks	5002
Verifying That MPLS Is Working Correctly	5010
Verifying the Physical Layer on the Switches	5011
Verifying the Routing Protocol	5011
Verifying the Core Interfaces Being Used for the MPLS Traffic	5011
Verifying RSVP	5012
MPLS Configuration Guidelines	5012
Supported MPLS Scaling Values	5013
MPLS Stitching For Virtual Machine Connection	5014
When Would I Use Stitching?	5015
How Does MPLS Stitching Work?	5015
How Do I Configure Stitching?	5016

	Which Switches Support Stitching?	5016
	Q&A	5016
Chapter 192	Configuration Statements and Command Statements for MPLS	5017
	Configuration Statements for MPLS	5017
	[edit protocols mpls] Hierarchy Level	5019
	adaptive	5024
	adjust-interval	5025
	adjust-threshold	5025
	adjust-threshold-overflow-limit	5026
	adjust-threshold-underflow-limit	5026
	admin-down	5027
	advertisement-hold-time	5027
	associate-backup-pe-groups	5028
	auto-bandwidth (MPLS Tunnel)	5029
	backup-pe-group	5030
	bandwidth (Fast Reroute, Signaled, and Multiclass LSPs)	5031
	bandwidth-model	5032
	bypass (Static LSP)	5033
	class-of-service (Protocols MPLS)	5034
	corouted-bidirectional	5035
	corouted-bidirectional-passive	5035
	description (Protocols MPLS)	5036
	diffserv-te	5037
	disable (Protocols MPLS)	5038
	exclude (for Fast Reroute)	5039
	exclude-srlg	5040
	explicit-null (Protocols MPLS)	5041
	fast-reroute (Protocols MPLS)	5042
	forwarding-table	5043
	from (Protocols MPLS)	5043
	gp-id	5044
	hop-limit	5045
	include-all (for Fast Reroute)	5046
	include-any (for Fast Reroute)	5046
	ingress (LSP)	5047
	install (Protocols MPLS)	5048
	interface (Protocols MPLS)	5049
	ipv6-tunneling	5050
	label-switched-path (Protocols MPLS)	5051
	ldp-tunneling	5053
	link-protection (Static LSPs)	5054
	log-updown (Protocols MPLS)	5055
	lsp-attributes	5056
	maximum-bandwidth (Protocols MPLS)	5057
	metric (Protocols MPLS)	5057
	minimum-bandwidth	5058
	monitor-bandwidth	5058
	mtu-signaling	5059

no-cspf	5060
no-decrement-ttl	5061
no-install-to-address	5062
no-propagate-ttl	5063
record	5064
no-trap	5065
node-link-protection (Protocols MPLS)	5066
oam (Protocols MPLS)	5067
optimize-aggressive	5068
optimize-hold-dead-delay	5069
optimize-switchover-delay	5070
optimize-timer (Protocols MPLS)	5071
p2mp (Protocols MPLS)	5072
path (Protocols MPLS)	5073
path-mtu	5074
policing (Protocols MPLS)	5075
policy-statement	5076
pop	5080
preference (Protocols MPLS)	5081
primary (Protocols MPLS)	5082
push	5083
record	5084
retry-limit	5085
revert-timer	5086
rsvp-error-hold-time	5087
secondary (Protocols MPLS)	5088
select	5089
signal-bandwidth	5089
smart-optimize-timer	5090
standby	5091
static-label-switched-path	5092
statistics (Protocols MPLS)	5094
swap	5095
switching-type	5096
te-class-matrix	5097
traffic-engineering (Protocols MPLS)	5098
transit-lsp-association	5099
te-class-matrix	5100
to	5101
traceoptions (Protocols MPLS)	5102
traffic-engineering (Protocols MPLS)	5104
transit-lsp-association	5105
Monitoring Commands for MPLS	5105
clear mpls lsp	5107
monitor label-switched-path	5109
ping mpls bgp	5112
ping mpls l2circuit	5114
ping mpls l3vpn	5117
ping mpls lsp-end-point	5120

request mpls lsp adjust-autobandwidth	5122
show security keychain	5124
show link-management	5127
show link-management peer	5131
show link-management routing	5133
show link-management statistics	5136
show link-management te-link	5138
show mpls call-admission-control	5140
show mpls cspf	5142
show mpls diffserv-te	5144
show route forwarding-table	5146
show mpls interface	5154
show link-management statistics	5156
show link-management te-link	5158
show mpls call-admission-control	5160
show mpls cspf	5162
show mpls diffserv-te	5164
show route forwarding-table	5166
show mpls interface	5174
show mpls lsp	5176
show mpls lsp autobandwidth	5194
show mpls path	5197
show route table	5199
show route forwarding-table	5227
show mpls static-lsp	5241
show ted database	5244
show ted link	5252
show ted protocol	5255

Part 73

Chapter 193

RSVP

Using RSVP	5259
Understanding MPLS Components for QFX Series and EX4600 Switches . . .	5259
Provider Edge Switches	5259
MPLS Protocol and Label-Switched Paths	5260
IP Over MPLS for Customer Edge Interfaces	5260
BGP Layer 3 VPN Configuration	5260
Routing Instances for Layer 3 VPN	5260
Routing Instances for Layer 2 VPN and Layer 3 VPN	5260
Ethernet Encapsulation for Layer 2 VPN	5261
Provider Switch	5261
Components Required for All Switches in the MPLS Network	5261
Interior Gateway Protocol	5261
Traffic Engineering	5262
MPLS Protocol	5262
RSVP	5262
Family mpls	5262
RSVP Overview	5263
MTU Signaling in RSVP	5264

	Tunneling LDP LSPs in RSVP LSPs	5264
	Tunneling LDP LSPs in RSVP LSPs Overview	5265
	Configuring MPLS on Provider Edge Switches	5265
	Configuring the Ingress PE Switch	5266
	Configuring the Egress PE Switch	5267
	Configuring MPLS on Provider Switches	5269
	Verifying That MPLS Is Working Correctly	5270
	Verifying the Physical Layer on the Switches	5270
	Verifying the Routing Protocol	5271
	Verifying the Core Interfaces Being Used for the MPLS Traffic	5271
	Verifying RSVP	5271
	Dynamic Bandwidth Management Using Container LSP Overview	5272
	Understanding RSVP Multipath Extensions	5272
	Junos OS RSVP Multipath Implementation	5273
	Current Traffic Engineering Challenges	5274
	Using Container LSP as a Solution	5277
	Accommodating the New Demand X	5277
	Creating New LSPs to Meet Demand X	5278
	Assigning Bandwidth to the New LSPs	5278
	Controlling the LSP Paths	5278
	Junos OS Container LSP Implementation	5279
	Container LSP Terminology	5279
	LSP Splitting	5280
	LSP Merging	5282
	Node and Link Protection	5284
	Naming Convention	5284
	Normalization	5285
	Constraint-Based Routing Path Computation	5290
	Sampling	5291
	Support for NSR, IPG-FA, and Static Routes	5291
	Configuration Statements Supported for Container LSPs	5294
	Impact of Configuring Container LSPs on Network Performance	5298
	Supported and Unsupported Features	5298
Chapter 194	Configuration Statements and Command Statements for RSVP	5301
	Configuration Statements for RSVP	5301
	[edit protocols rsvp] Hierarchy Level	5302
	admin-group	5304
	authentication-key (Protocols RSVP)	5305
	aggregate (Protocols RSVP)	5306
	bandwidth (Protocols RSVP)	5307
	bypass (Signaled LSP)	5308
	class-of-service (Protocols RSVP)	5309
	container-label-switched-path	5310
	disable (Protocols RSVP)	5311
	fast-reroute (Protocols RSVP)	5312
	graceful-deletion-timeout	5312
	graceful-restart (Enabling Globally)	5313
	hello-acknowledgements	5314

hello-interval (Protocols RSVP)	5315
helper-disable (Multiple Protocols)	5316
hop-limit	5317
interface (Protocols RSVP)	5318
keep-multiplier	5319
link-protection (RSVP)	5320
load-balance (Protocols RSVP)	5321
max-bypasses	5322
maximum-helper-recovery-time	5323
maximum-helper-restart-time (RSVP)	5324
no-cspf (Protocols RSVP)	5325
no-interface-hello	5325
no-local-reversion	5326
no-node-id-subobject	5327
no-p2mp-sublsp	5327
node-hello	5328
optimize-timer (Protocols RSVP)	5328
path (Protocols RSVP)	5329
preemption	5330
priority (Protocols RSVP)	5331
refresh-time	5332
reliable	5332
setup-protection	5333
subscription	5334
soft-preemption (Protocols RSVP)	5335
splitting-merging	5336
traceoptions (Protocols RSVP)	5338
tunnel-services (RSVP)	5340
update-threshold	5340
Monitoring Commands for RSVP	5341
clear mpls container-lsp	5342
clear rsvp session	5344
clear rsvp statistics	5346
ping mpls rsvp	5347
request mpls container-lsp	5352
clear mpls container-lsp	5353
show rsvp interface	5355
show rsvp neighbor	5360
show rsvp session	5365
show rsvp statistics	5375
show rsvp version	5379
traceroute mpls rsvp	5382

Guide 15 Multicast Protocols Feature Guide for QFX10000 Switches

Chapter 195	Overview	5389
	Multicast Overview	5389
	Comparing Multicast to Unicast	5389
	IP Multicast Uses	5391
	IP Multicast Terminology	5392
	Reverse-Path Forwarding for Loop Prevention	5393
	Shortest-Path Tree for Loop Prevention	5393
	Administrative Scoping for Loop Prevention	5394
	Multicast Leaf and Branch Terminology	5394
	IP Multicast Addressing	5394
	Multicast Addresses	5395
	Layer 2 Frames and IPv4 Multicast Addresses	5395
	Multicast Interface Lists	5397
	Multicast Routing Protocols	5398
	T Series Router Multicast Performance	5401
Part 74	Managing Group Membership	
Chapter 196	Using IGMP	5405
	Understanding Group Membership Protocols	5405
	Understanding IGMP	5406
	Configuring IGMP	5408
	Enabling IGMP	5409
	Changing the IGMP Version	5410
	Modifying the IGMP Host-Query Message Interval	5411
	Modifying the IGMP Last-Member Query Interval	5412
	Specifying Immediate-Leave Host Removal for IGMP	5413
	Filtering Unwanted IGMP Reports at the IGMP Interface Level	5414
	Accepting IGMP Messages from Remote Subnetworks	5415
	Modifying the IGMP Query Response Interval	5416
	Modifying the IGMP Robustness Variable	5417
	Limiting the Maximum IGMP Message Rate	5418
	Enabling IGMP Static Group Membership	5418
	Recording IGMP Join and Leave Events	5425
	Limiting the Number of IGMP Multicast Group Joins on Logical Interfaces	5426
	Tracing IGMP Protocol Traffic	5427
	Disabling IGMP	5429
Chapter 197	Using IGMP Snooping	5431
	IGMP Snooping Overview	5431
	How IGMP Snooping Works	5431
	How IGMP Snooping Works with Routed VLAN Interfaces	5432
	How Hosts Join and Leave Multicast Groups	5432
	IGMP Snooping and Forwarding Interfaces	5433

	General Forwarding Rules	5433
	Configuring IGMP Snooping	5434
	Configuring VLAN-Specific IGMP Snooping Parameters	5435
	Example: Configuring IGMP Snooping	5436
	Monitoring IGMP Snooping	5438
	Verifying the IGMP Snooping Group Timeout Value	5439
Chapter 198	Using MLD	5441
	Understanding MLD	5441
	Examples: Configuring MLD	5444
	Understanding MLD	5444
	Configuring MLD	5447
	Enabling MLD	5448
	Modifying the MLD Version	5449
	Modifying the MLD Host-Query Message Interval	5449
	Modifying the MLD Query Response Interval	5450
	Modifying the MLD Last-Member Query Interval	5450
	Specifying Immediate-Leave Host Removal for MLD	5451
	Filtering Unwanted MLD Reports at the MLD Interface Level	5452
	Example: Modifying the MLD Robustness Variable	5453
	Limiting the Maximum MLD Message Rate	5454
	Enabling MLD Static Group Membership	5454
	Create a MLD Static Group Member	5455
	Automatically create static groups	5456
	Automatically increment group addresses	5457
	Specify multicast source address (in SSM mode)	5458
	Automatically specify multicast sources	5459
	Automatically increment source addresses	5460
	Exclude multicast source addresses (in SSM mode)	5461
	Example: Recording MLD Join and Leave Events	5462
	Configuring the Number of MLD Multicast Group Joins on Logical Interfaces	5464
	Disabling MLD	5466
Chapter 199	Using MLD Snooping	5467
	Understanding MLD Snooping	5467
	How MLD Snooping Works	5468
	MLD Message Types	5469
	How Hosts Join and Leave Multicast Groups	5469
	Support for MLDv2 Multicast Sources	5470
	MLD Snooping and Forwarding Interfaces	5470
	General Forwarding Rules	5471
	Examples of MLD Snooping Multicast Forwarding	5471
	Scenario 1: Switch Forwarding Multicast Traffic to a Multicast Router and Hosts	5471
	Scenario 2: Switch Forwarding Multicast Traffic to Another Switch	5472
	Scenario 3: Switch Connected to Hosts Only (No MLD Querier)	5473

	Scenario 4: Layer 2/Layer 3 Switch Forwarding Multicast Traffic	
	Between VLANs	5474
	Configuring MLD Snooping on a VLAN (CLI Procedure)	5475
	Enabling or Disabling MLD Snooping on VLANs	5476
	Configuring the MLD Version	5477
	Enabling Immediate Leave	5477
	Configuring an Interface as a Multicast-Router Interface	5478
	Configuring Static Group Membership on an Interface	5479
	Changing the Timer and Counter Values	5480
	Example: Configuring MLD Snooping	5481
	Verifying MLD Snooping	5484
	Verifying MLD Snooping Memberships	5485
	Verifying MLD Snooping Interfaces	5485
	Viewing MLD Snooping Statistics	5486
	Viewing MLD Snooping Routing Information	5487
Part 75	Configuring PIM	
Chapter 200	Using PIM Basic Features	5491
	PIM Overview	5491
	Basic PIM Network Components	5493
	PIM on Aggregated Interfaces	5494
	Changing the PIM Version	5494
	Modifying the PIM Hello Interval	5494
	Preserving Multicast Performance by Disabling Response to the ping Utility	5495
	Configuring PIM Trace Options	5496
	Configuring Interface Priority for PIM Designated Router Selection	5498
	Configuring PIM Designated Router Election on Point-to-Point Links	5499
	Configuring BFD for PIM	5500
	Configuring BFD Authentication for PIM	5502
	Configuring BFD Authentication Parameters	5502
	Viewing Authentication Information for BFD Sessions	5503
	Disabling PIM	5505
	Disabling the PIM Protocol	5505
	Disabling PIM on an Interface	5506
	Disabling PIM for a Family	5506
	Disabling PIM for a Rendezvous Point	5507
Chapter 201	Using PIM Sparse Mode	5509
	Understanding PIM Sparse Mode	5509
	Rendezvous Point	5511
	RP Mapping Options	5511
	Designated Router	5512
	Enabling PIM Sparse Mode	5512
	Configuring PIM Join Load Balancing	5513
	Modifying the Join State Timeout	5517
	Example: Enabling Join Suppression	5517

Chapter 202	Using PIM Dense Mode and PIM Sparse-Dense Mode	5523
	Understanding PIM Dense Mode	5523
	Understanding PIM Sparse-Dense Mode	5525
	Mixing PIM Sparse and Dense Modes	5525
	Configuring PIM Dense Mode Properties	5526
	Configuring PIM Sparse-Dense Mode Properties	5527
Chapter 203	Using Source-Specific Multicast	5529
	Source-Specific Multicast Groups Overview	5529
	Understanding PIM SSM	5530
	PIM SSM	5531
	Example: Configuring PIM SSM on a Network	5533
	Example: Configuring an SSM-Only Domain	5534
	Example: Configuring SSM Mapping	5535
	Example: Configuring Source-Specific Multicast Groups with Any-Source Override	5537
	Example: Configuring SSM Maps for Different Groups to Different Sources	5541
	Multiple SSM Maps and Groups for Interfaces	5541
	Example: Configuring Multiple SSM Maps Per Interface	5541
Chapter 204	Using Static RP	5545
	Understanding Static RP	5545
	Configuring Local PIM RPs	5545
	Configuring the Static PIM RP Address on the Non-RP Routing Device	5547
Chapter 205	Using Anycast RP	5551
	Understanding RP Mapping with Anycast RP	5551
	Example: Configuring PIM Anycast With or Without MSDP	5552
	Configuring a PIM Anycast RP Router with MSDP	5555
	Configuring a PIM Anycast RP Router Using Only PIM	5556
	Configuring All PIM Anycast Non-RP Routers	5557
	Example: Configuring Multiple RPs in a Domain with Anycast RP	5558
Chapter 206	Using Auto-RP	5561
	Understanding PIM Auto-RP	5561
	Configuring PIM Auto-RP	5561
Chapter 207	Using PIM Bootstrap Router	5567
	Understanding the PIM Bootstrap Router	5567
	Configuring PIM Bootstrap Properties for IPv4 or IPv6	5567
	Example: Rejecting PIM Bootstrap Messages at the Boundary of a PIM Domain	5569
	Example: Configuring PIM BSR Filters	5569
Chapter 208	Using PIM Filtering	5571
	Understanding Multicast Message Filters	5571
	Filtering MAC Addresses	5572
	Filtering RP and DR Register Messages	5572
	Configuring Interface-Level PIM Neighbor Policies	5573
	Filtering Outgoing PIM Join Messages	5574
	Filtering Incoming PIM Join Messages	5575

	Configuring Register Message Filters on a PIM RP and DR	5577
Chapter 209	Using PIM RPT and SPT Cutover	5579
	Understanding Multicast Rendezvous Points, Shared Trees, and Rendezvous-Point Trees	5579
	Building an RPT Between the RP and Receivers	5580
	PIM Sparse Mode Source Registration	5581
	Multicast Shortest-Path Tree	5584
	SPT Cutover	5585
	SPT Cutover Control	5588
	Example: Configuring the PIM Assert Timeout	5588
	Example: Configuring the PIM SPT Threshold Policy	5590
Part 76	Configuring MSDP	
Chapter 210	Using MSDP	5597
	Understanding MSDP	5597
	Configuring MSDP	5598
	Filtering MSDP SA Messages	5600
	Tracing MSDP Protocol Traffic	5600
	Configuring the Interface to Accept Traffic from a Remote Source	5602
	Example: Configuring MSDP	5603
	Example: Configuring MSDP with Active Source Limits and Mesh Groups	5604
	Example: Configuring PIM Anycast With or Without MSDP	5610
	Configuring a PIM Anycast RP Router with MSDP	5613
Part 77	Configuration Statements and Operational Commands	
Chapter 211	Configuration Statements (IGMP)	5617
	accounting (Protocols IGMP)	5618
	accounting (Protocols IGMP Interface)	5618
	asm-override-ssm	5619
	disable (Protocols IGMP)	5619
	exclude (Protocols IGMP)	5620
	group (Protocols IGMP)	5621
	group-count	5622
	group-increment (Protocols IGMP)	5622
	group-limit (Protocols IGMP)	5623
	group-policy (Protocols IGMP)	5623
	igmp	5624
	immediate-leave (Protocols IGMP)	5626
	interface (Protocols IGMP)	5627
	maximum-transmit-rate (Protocols IGMP)	5628
	oif-map (IGMP Interface)	5628
	passive (IGMP)	5629
	promiscuous-mode (Protocols IGMP)	5630
	query-interval (Protocols IGMP)	5631
	query-last-member-interval (Protocols IGMP)	5632
	query-response-interval (Protocols IGMP)	5633
	robust-count (Protocols IGMP)	5634

	source (Protocols IGMP)	5635
	source-count (Protocols IGMP)	5636
	source-increment (Protocols IGMP)	5637
	static (Protocols IGMP)	5638
	traceoptions (Protocols IGMP)	5639
	version (Protocols IGMP)	5641
Chapter 212	Configuration Statements (IGMP Snooping)	5643
	all	5644
	data-forwarding	5644
	disable (IGMP Snooping)	5645
	group (IGMP Snooping)	5645
	group-limit (IGMP and MLD Snooping)	5646
	host-only-interface	5647
	igmp-querier	5647
	igmp-snooping	5648
	immediate-leave (Bridge Domains)	5649
	interface (Bridge Domains)	5650
	interface (IGMP Snooping)	5651
	l2-querier	5651
	multicast-router-interface (IGMP Snooping)	5652
	query-interval (Bridge Domains)	5653
	query-last-member-interval (Bridge Domains)	5654
	query-response-interval (Bridge Domains)	5655
	receiver	5656
	robust-count (IGMP Snooping)	5656
	source-address	5657
	src-address (IGMP Querier)	5658
	source-vlans	5658
	static (IGMP Snooping)	5659
	traceoptions (IGMP Snooping)	5660
	version (IGMP Snooping)	5662
	vlan (IGMP Snooping)	5663
Chapter 213	Configuration Statements (MLD Snooping)	5665
	all	5666
	data-forwarding	5666
	disable (IGMP Snooping)	5667
	group (IGMP Snooping)	5667
	group-limit (IGMP and MLD Snooping)	5668
	host-only-interface	5669
	igmp-querier	5669
	igmp-snooping	5670
	immediate-leave (Bridge Domains)	5671
	interface (Bridge Domains)	5672
	interface (IGMP Snooping)	5673
	l2-querier	5673
	mld-snooping	5674
	multicast-router-interface (IGMP Snooping)	5675
	query-interval (Bridge Domains)	5676

	query-last-member-interval (Bridge Domains)	5677
	query-response-interval (Bridge Domains)	5678
	receiver	5679
	robust-count (IGMP Snooping)	5679
	source-address	5680
	src-address (IGMP Querier)	5681
	source-vlans	5681
	static (IGMP Snooping)	5682
	traceoptions (IGMP Snooping)	5683
	version (IGMP Snooping)	5685
	vlan (IGMP Snooping)	5686
Chapter 214	Configuration Statements (PIM)	5687
	address (Anycast RPs)	5689
	address (Local RPs)	5690
	address (Static RPs)	5691
	algorithm	5692
	anycast-pim	5693
	assert-timeout	5694
	authentication	5695
	auto-rp	5696
	bfd-liveness-detection	5697
	bootstrap	5698
	bootstrap-export	5699
	bootstrap-import	5700
	bootstrap-priority	5701
	dense-groups	5702
	detection-time (BFD for PIM)	5703
	disable (PIM)	5704
	dr-election-on-p2p	5705
	dr-register-policy	5705
	embedded-rp	5706
	export (Bootstrap)	5707
	export (Protocols PIM)	5708
	family (Bootstrap)	5709
	family (Protocols PIM)	5710
	family (Local RP)	5711
	group (RPF Selection)	5712
	group-ranges	5713
	hello-interval	5714
	hold-time (Protocols PIM)	5715
	import (Protocols PIM Bootstrap)	5716
	import (Protocols PIM)	5717
	infinity	5718
	interface	5719
	join-load-balance	5720
	join-prune-timeout	5721
	key-chain	5721
	local	5722

	local-address (Protocols PIM)	5723
	loose-check	5724
	mapping-agent-election	5725
	maximum-rps	5726
	minimum-interval (PIM BFD Liveness Detection)	5727
	minimum-interval (PIM BFD Transmit Interval)	5728
	minimum-receive-interval	5729
	mode (Protocols PIM)	5729
	multiplier	5730
	neighbor-policy	5730
	next-hop (PIM RPF Selection)	5731
	no-adaptation (PIM BFD Liveness Detection)	5731
	override-interval	5732
	pim	5733
	prefix-list (PIM RPF Selection)	5736
	priority (Bootstrap)	5737
	priority (PIM Interfaces)	5738
	priority (PIM RPs)	5739
	propagation-delay	5740
	register-probe-time	5741
	reset-tracking-bit	5742
	rib-group (Protocols PIM)	5743
	rp	5744
	rp-register-policy	5746
	rp-set	5747
	rpf-selection	5748
	source (PIM RPF Selection)	5749
	spt-threshold	5750
	static (Protocols PIM)	5751
	threshold (PIM BFD Detection Time)	5752
	threshold (PIM BFD Transmit Interval)	5753
	transmit-interval (PIM BFD Liveness Detection)	5754
	traceoptions (Protocols PIM)	5755
	version (BFD)	5758
	version (PIM)	5759
	wildcard-source (PIM RPF Selection)	5760
Chapter 215	Configuration Statements (Source-Specific Multicast)	5761
	asm-override-ssm	5761
	policy (SSM Maps)	5762
	ssm-groups	5763
	ssm-map (Protocols IGMP)	5764
	ssm-map (Routing Options Multicast)	5765
	ssm-map-policy (IGMP)	5766
Chapter 216	Configuration Statements (MSDP)	5767
	active-source-limit	5768
	authentication-key	5769
	data-encapsulation	5770
	default-peer	5771

	disable (Protocols MSDP)	5772
	export (Protocols MSDP)	5773
	group (Protocols MSDP)	5774
	import (Protocols MSDP)	5775
	local-address (Protocols MSDP)	5776
	maximum (MSDP Active Source Messages)	5777
	mode (Protocols MSDP)	5778
	msdp	5779
	peer (Protocols MSDP)	5781
	rib-group (Protocols MSDP)	5782
	source (Protocols MSDP)	5783
	threshold (MSDP Active Source Messages)	5784
	traceoptions (Protocols MSDP)	5785
Chapter 217	Operational Commands (IGMP)	5789
	clear igmp membership	5790
	clear igmp statistics	5793
	show igmp group	5795
	show configuration protocols igmp	5799
	show igmp interface	5801
	show igmp statistics	5805
	show system statistics igmp	5808
Chapter 218	Operational Commands (IGMP Snooping)	5813
	clear igmp-snooping membership	5814
	clear igmp-snooping statistics	5815
	show igmp-snooping membership	5816
	show igmp-snooping route	5819
	show igmp-snooping statistics	5821
	show igmp-snooping vlans	5823
Chapter 219	Operational Commands (PIM)	5825
	clear multicast bandwidth-admission	5827
	clear multicast scope	5829
	clear multicast sessions	5830
	clear multicast statistics	5831
	clear pim join	5832
	clear pim register	5834
	clear pim statistics	5836
	mtrace	5839
	mtrace from-source	5842
	mtrace monitor	5845
	mtrace to-gateway	5847
	show multicast flow-map	5850
	show multicast interface	5852
	show multicast mrinfo	5854
	show multicast next-hops	5856
	show multicast pim-to-igmp-proxy	5859
	show multicast pim-to-mld-proxy	5861
	show multicast route	5863

	show multicast rpf	5869
	show multicast scope	5873
	show multicast sessions	5875
	show multicast usage	5878
	show pim bootstrap	5881
	show pim interfaces	5883
	show pim join	5886
	show pim neighbors	5907
	show pim rps	5911
	show pim source	5918
	show pim statistics	5921
Chapter 220	Operational Commands (MSDP)	5931
	clear msdp cache	5932
	clear msdp statistics	5933
	show msdp	5934
	show msdp source	5936
	show msdp source-active	5938
	show msdp statistics	5941
	test msdp	5945

Guide 16 Security Feature Guide for QFX10000 Switches

Part 78

Chapter 221

Firewall Filters

Configuring Firewall Filters	5951
Overview of Firewall Filters	5951
Firewall Filter Types	5952
Firewall Filter Components	5953
Firewall Filter Processing	5953
How Many Filters Are Supported?	5953
Understanding How Firewall Filters Are Evaluated	5954
Understanding How Firewall Filters Control Packet Flows	5956
Understanding Firewall Filter Match Conditions	5957
Filter Match Conditions	5957
Numeric Filter Match Conditions	5957
Interface Filter Match Conditions	5958
IP Address Filter Match Conditions	5958
MAC Address Filter Match Conditions	5959
Bit-Field Filter Match Conditions	5959
Firewall Filter Match Conditions and Actions for QFX10000 Switches	5960
Understanding How a Firewall Filter Tests a Protocol	5971
Understanding Firewall Filter Planning	5972
Planning the Number of Firewall Filters to Create	5973
Understanding How Many Firewall Filters Are Supported	5973
Egress Filters	5975
Avoid Configuring too Many Filters	5975
Configuring TCAM Error Messages	5976
Policers can Limit Egress Filters	5976
Planning for Filter-Specific Policers	5977

Planning for Filter-Based Forwarding	5977
Understanding Firewall Filter Processing Points for Bridged and Routed	
Packets	5978
Configuring Firewall Filters	5978
Configuring a Firewall Filter	5979
Applying a Firewall Filter to a Port	5980
Applying a Firewall Filter to a VLAN	5981
Applying a Firewall Filter to a Layer 3 (Routed) Interface	5981
Applying Firewall Filters to Interfaces	5982
Understanding Filter-Based Forwarding	5983
Example: Using Filter-Based Forwarding to Route Application Traffic to a Security	
Device	5983
Monitoring Firewall Filter Traffic	5987
Monitoring Traffic for All Firewall Filters and Policers That Are	
Configured	5987
Monitoring Traffic for a Specific Firewall Filter	5988
Monitoring Traffic for a Specific Policer	5988
Verifying That Firewall Filters Are Operational	5988
Troubleshooting Firewall Filters	5989
Troubleshooting QFX10000 Switches	5989
Do Not Combine Match Conditions for Different Layers	5989
Layer 2 Packets Cannot be Discarded with Firewall Filters	5989
Protect-RE (loopback) Firewall Filter Does Not Filter Packets Applied	
to EMO Interfaces	5990
Troubleshooting Other Switches	5990
Firewall Filter Configuration Returns a No Space Available in TCAM	
Message	5990
Filter Counts Previously Dropped Packet	5992
Matching Packets Not Counted	5993
Counter Reset When Editing Filter	5993
Cannot Include loss-priority and policer Actions in Same Term	5993
Cannot Egress Filter Certain Traffic Originating on QFX Switch	5994
Firewall Filter Match Condition Not Working with Q-in-Q Tunneling	5994
Egress Firewall Filters with Private VLANs	5994
Egress Filtering of L2PT Traffic Not Supported	5995
Cannot Drop BGP Packets in Certain Circumstances	5995
Invalid Statistics for Policer	5995
Policers can Limit Egress Filters	5995

Part 79

Chapter 222

Policers

Configuring Policers	5999
Overview of Policers	5999
Policer Overview	6000
Policer Types	6000
Policer Actions	6001
Policer Colors	6002
Filter-Specific Policers	6002
Suggested Naming Convention for Policers	6003

Policer Counters	6003
Policer Algorithms	6003
How Many Policers Are Supported?	6004
Policers Can Limit Egress Firewall Filters	6004
Understanding Policers with Link Aggregation Groups	6005
Understanding Color-Blind Mode for Single-Rate Tricolor Marking	6005
Understanding Color-Aware Mode for Single-Rate Tricolor Marking	6006
Summary of PLP Changes	6006
Effect on Green Packets (Low PLP)	6007
Effect on Yellow Packets (Medium PLP)	6007
Effect on Red Packets (High PLP)	6007
Understanding Color-Blind Mode for Two-Rate Tricolor Marking	6008
Understanding Color-Aware Mode for Two-Rate Tricolor Marking	6008
Summary of PLP Changes	6008
Effect on Green Packets (Low PLP)	6009
Effect on Yellow Packets (Medium PLP)	6009
Effect on Red Packets (High PLP)	6010
Example: Using Two-Color Policers and Prefix Lists	6010
Example: Using Policers to Manage Oversubscription	6013
Assigning Forwarding Classes and Loss Priority	6015
Configuring Color-Blind Egress Policers for Medium-Low PLP	6016
Configuring Two-Color and Three-Color Policers to Control Traffic Rates	6017
Configuring Two-Color Policers	6017
Configuring Three-Color Policers	6018
Specifying Policers in a Firewall Filter Configuration	6018
Applying a Firewall Filter That Includes a Policer	6019
Verifying That Two-Color Policers Are Operational	6019
Verifying That Three-Color Policers Are Operational	6020
Troubleshooting Policer Configuration	6020
Incomplete Count of Packet Drops	6020
Counter Reset When Editing Filter	6021
Invalid Statistics for Policer	6021
Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured	6021
Filter-Specific Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured	6022
Policers Can Limit Egress Filters	6023

Part 80

Chapter 223

Configuring Port Security

Port Security	6027
Overview of Access Port Protection	6027
Mitigation of Ethernet Switching Table Overflow Attacks	6027
Mitigation of Rogue DHCP Server Attacks	6028
Protection Against DHCP Starvation Attacks	6028
Understanding MAC Limiting and MAC Move Limiting for Port Security	6029
MAC Limiting	6029
MAC Move Limiting	6030
Actions for MAC Limiting	6030

MAC Addresses That Exceed the MAC Limit or MAC Move Limit	6030
Verifying That MAC Limiting Is Working Correctly	6031
Verifying That MAC Limiting for Dynamic MAC Addresses Is Working Correctly	6031
Verifying That Allowed MAC Addresses Are Working Correctly	6032
Verifying That Interfaces Are Shut Down	6032
Customizing the Ethernet Switching Table Display to View Information for a Specific Interface	6033
Verifying That MAC Move Limiting Is Working Correctly	6034
Verifying That the Port Error Disable Setting Is Working Correctly	6034
Configuring Persistent MAC Learning (CLI Procedure)	6035
Understanding Trusted and Untrusted Ports	6037
Understanding Trusted DHCP Servers for Port Security	6037
Verifying That a Trusted DHCP Server Is Working Correctly	6037
Understanding DHCP Option 82 for Port Security	6039
DHCP Option 82 Processing	6039
Suboption Components of Option 82	6040
Configurations That Support Option 82	6040
Understanding Static ARP Entries	6041

Part 81

Chapter 224

Configuring Device Security

Device Security	6045
Understanding Storm Control	6045
Example: Configuring Storm Control to Prevent Network Outages	6046
Verifying That the Port Error Disable Setting Is Working Correctly	6049
Understanding Unicast RPF	6050
Unicast RPF for Switches Overview	6050
Unicast RPF Implementation	6051
Unicast RPF Packet Filtering	6051
Bootstrap Protocol (BOOTP) and DHCP Requests	6051
Default Route Handling	6051
When to Enable Unicast RPF	6051
When Not to Enable Unicast RPF	6052
Limitations of the Unicast RPF Implementation on EX3200, EX4200, and EX4300 Switches	6053
Configuring Unicast RPF (CLI Procedure)	6054
Disabling Unicast RPF (CLI Procedure)	6055
Verifying Unicast RPF Status	6056
Understanding Unknown Unicast Forwarding	6059
Configuring Unknown Unicast Forwarding (CLI Procedure)	6059
Configuring Unknown Unicast Forwarding on EX4300 Switches	6060
Configuring Unknown Unicast Forwarding on EX9200 Switches	6060

Part 82	DDoS Protection	
Chapter 225	Overview of DDoS Protection	6065
	Understanding Distributed Denial-of-Service Protection on QFX Series	
	Switches	6065
	Policer Types and Packet Priorities	6066
	Example of Policer Priority Behavior	6066
	Policer Enforcement Points on QFX Series Switches	6067
Chapter 226	Configuring DDoS Protection	6069
	Configuring Protection Against DDoS Attacks	6069
	Example: Configuring DDoS Protection on QFX Series Switches	6070
	Disabling DDoS Protection Policers and Logging Globally	6074
	Configuring DDoS Protection Policers on QFX Series Switches	6074
	Configuring the Aggregate Policer for a Protocol Group	6075
	Configuring Packet-Type Policers for a Protocol Group	6076
	Configuring Policers on Individual Line Cards	6077
	Disabling Policers and Policer Logging	6077
	Tracing DDoS Protection Operations	6078
	Configuring the DDoS Protection Trace Log Filename	6079
	Configuring the Number and Size of DDoS Protection Log Files	6079
	Configuring Access to the DDoS Protection Log File	6080
	Configuring a Regular Expression for DDoS Protection Messages to Be	
	Logged	6080
	Configuring the DDoS Protection Tracing Flags	6080
	Configuring the Severity Level to Filter Which DDoS Protection Messages	
	Are Logged	6081
Chapter 227	Monitoring DDoS Protection	6083
	Verifying and Managing DDoS Protection	6083
Part 83	Configuration Statements and Operational Commands	
Chapter 228	Configuration Statements (Firewall Filters)	6087
	family	6088
	filter	6089
	filter (Layer 2 and Layer 3 Interfaces)	6090
	filter (VLANs)	6091
	firewall	6092
	from	6093
	input (Forwarding Table)	6094
	interface-specific	6094
	output (Forwarding Table)	6095
	term	6096
	then (Filters)	6097
Chapter 229	Configuration Statements (Policers)	6099
	action	6100
	bandwidth-limit	6100
	burst-size-limit	6101
	color-aware	6102

	color-blind	6103
	committed-burst-size	6104
	committed-information-rate	6105
	excess-burst-size	6106
	filter-specific	6107
	firewall	6108
	if-exceeding	6109
	loss-priority high then discard (Three-Color Policer)	6110
	peak-burst-size	6111
	peak-information-rate	6112
	policer	6113
	single-rate	6114
	then (Policers)	6115
	three-color-policer	6116
	two-rate	6117
Chapter 230	Configuration Statements (Device Security)	6119
	action-shutdown	6120
	bandwidth-level	6121
	bandwidth-percentage	6122
	interface (Unknown Unicast Forwarding)	6123
	no-broadcast	6124
	no-multicast	6125
	no-unknown-unicast	6126
	rpf-check	6127
	storm-control	6128
	storm-control-profiles	6129
	unknown-unicast-forwarding	6130
Chapter 231	Configuration Statements (Port Security)	6131
	circuit-id	6132
	fc-map	6134
	fcoe-trusted	6136
	mac-move-limit	6137
	no-allowed-mac-log	6138
	no-gratuitous-arp-request	6139
	persistent-learning	6139
	port-error-disable	6140
	vendor-id	6142
	write-interval	6143
Chapter 232	Configuration Statements (DDoS Protection)	6145
	bandwidth (DDoS)	6146
	bandwidth-scale (DDoS)	6147
	burst (DDoS)	6148
	burst-scale (DDoS)	6149
	bypass-aggregate (DDoS)	6150
	ddos-protection (DDoS)	6151
	disable-fpc (DDoS)	6152
	disable-logging (DDoS)	6153

	fpc (DDoS)	6154
	global (DDoS)	6155
	priority (DDoS)	6156
	protocols (DDoS)	6157
	traceoptions (DDoS)	6162
Chapter 233	Operational Commands (Firewall Filters)	6165
	clear firewall	6166
	show firewall	6167
	show firewall policer	6171
	show interfaces filters	6173
	show pfe filter hw summary	6175
Chapter 234	Operational Commands (Port Security)	6177
	clear ethernet-switching port-error	6178
Chapter 235	Operational Commands (DDoS Protection)	6179
	clear ddos-protection protocols	6180
	show ddos-protection protocols	6182
	show ddos-protection protocols parameters	6201
	show ddos-protection protocols statistics	6208
	show ddos-protection statistics	6218
	show ddos-protection version	6220
	show ddos-protection protocols violations	6221

Guide 17 **System Services Feature Guide for QFX10000 Switches**

Part 84	Port Mirroring
Chapter 236	Configuring Port Mirroring
	6227
	Understanding Port Mirroring
	6227
	Port Mirroring Overview
	6227
	Port Mirroring Instance Types
	6228
	Port-Mirroring Terminology
	6228
	Port Mirroring and STP
	6230
	Port Mirroring Constraints and Limitations
	6230
	Local and Remote Port Mirroring
	6230
	Remote Port Mirroring Only
	6232
	Configuring Port Mirroring
	6232
	Configuring Port Mirroring for Local Analysis
	6233
	Configuring Port Mirroring for Remote Analysis
	6233
	Filtering the Traffic Entering an Analyzer
	6234
	Examples: Configuring Port Mirroring for Local Analysis
	6235
	Example: Mirroring Employee Web Traffic with a Firewall Filter
	6237
	Example: Configuring Port Mirroring for Remote Analysis
	6240

	Example: Mirroring Employee Web Traffic with a Firewall Filter	6245
	Troubleshooting Port Mirroring	6248
	Port Mirroring Constraints and Limitations	6248
	Local and Remote Port Mirroring	6248
	Remote Port Mirroring Only	6250
	Egress Port Mirroring with VLAN Translation	6250
	Egress Port Mirroring with Private VLANs	6250
Part 85	DHCP and DHCP Relay	
Chapter 237	Using DHCP and DHCP Relay	6255
	DHCP and BOOTP Relay Overview	6255
	Configuring DHCP and BOOTP	6255
	Verifying and Managing DHCP Local Server Configuration	6256
	Verifying and Managing DHCPv6 Local Server Configuration	6257
	Verifying and Managing DHCP Relay Configuration	6257
	Verifying and Managing DHCPv6 Relay Configuration	6257
	Understanding Layer 3 Logical Interfaces	6258
Part 86	Configuration Statements and Operational Commands	
Chapter 238	Configuration Statements (Port Mirroring)	6261
	analyzer	6262
	egress	6263
	ethernet-switching (Port Mirroring)	6264
	family (Port Mirroring)	6265
	inet (Port Mirroring)	6266
	ingress (Port Mirroring)	6267
	input	6268
	instance (Port Mirroring)	6269
	interface (Port Mirroring)	6270
	ip-address (Port Mirroring)	6271
	no-tag	6271
	output	6272
	port-mirroring	6273
	routing-instance (Port Mirroring)	6274
	vlan (Port Mirroring)	6275
Chapter 239	Configuration Statements (DHCP and DHCP Relay)	6277
	dhcp-local-server	6278
	dhcp-relay	6284
	dhcp-service	6289
	forwarding-options	6290
	forward-only (DHCP Relay Agent)	6295
Chapter 240	Configuration Statements (Encryption)	6297
	authentication-key-chains	6298
	cache-size	6299
	cache-timeout-negative	6300
	ca-name	6301

	certificates	6302
	certification-authority	6303
	crl (Encryption Interface)	6304
	encoding	6304
	enrollment-retry	6305
	enrollment-url	6305
	file	6306
	key (Authentication Keychain)	6307
	key-chain (Security)	6308
	ldap-url	6309
	local	6310
	maximum-certificates	6311
	path-length	6311
	secret	6312
	security	6313
	ssh-known-hosts	6314
	start-time (Authentication Key Transmission)	6315
	traceoptions	6317
Chapter 241	Operational Command (Port Mirroring)	6319
	show analyzer	6320
Chapter 242	Operational Commands (DHCP Local Server)	6323
	clear dhcp server binding	6324
	clear dhcp server statistics	6327
	clear dhcpv6 server binding	6329
	clear dhcpv6 server statistics	6331
	request dhcp server reconfigure	6332
	request dhcpv6 server reconfigure	6334
	request system reboot	6336
	show dhcp server binding	6341
	show dhcp server statistics	6347
	show dhcpv6 server binding	6350
	show dhcpv6 server statistics	6356
Chapter 243	Operational Commands (DHCP Relay Agent)	6359
	clear dhcp relay binding	6360
	clear dhcp relay statistics	6362
	clear dhcpv6 relay binding	6365
	clear dhcpv6 relay statistics	6368
	show dhcp relay binding	6370
	show dhcp relay statistics	6375
	show dhcpv6 relay binding	6378
	show dhcpv6 relay statistics	6384
	show route extensive	6387
	show route protocol	6404

Guide 18	Storage Feature Guide for QFX10000 Switches
Chapter 244	Storage Overview 6417
	Overview of FIP 6417
	Understanding Fibre Channel Terminology 6417
Part 87	Configuring FCoE on a Transit Switch
Chapter 245	Using FCoE on a Transit Switch 6431
	Understanding FCoE Transit Switch Functionality 6431
	Understanding FCoE 6436
	FCoE Devices 6437
	FCoE Frames 6438
	Virtual Links 6439
	FCoE VLANs 6439
	Configuring VLANs for FCoE Traffic on an FCoE Transit Switch 6442
	Troubleshooting Dropped FCoE Traffic 6447
	Understanding MC-LAGs on an FCoE Transit Switch 6451
	Supported Topology 6451
	Transit Switches (Server Access) 6453
	MC-LAG Switches (FCoE Aggregation) 6453
	FIP Snooping and FCoE Trusted Ports 6453
	CoS and Data Center Bridging (DCB) 6454
	Example: Configuring CoS for FCoE Transit Switch Traffic Across an MC-LAG . . 6454
	Understanding FCoE and FIP Session High Availability 6480
	High Availability for Fibre Channel Process Termination (FCoE-FC Gateway Mode, QFX3500 Only) 6480
	High Availability for FIP Snooping 6480
	Nonstop Software Upgrade (QFabric Systems) 6481
	Troubleshooting Dropped FIP Traffic 6482
	Understanding OxID Hash Control for FCoE Traffic Load Balancing on Standalone Switches 6484
	Enabling and Disabling CoS OxID Hash Control on Standalone Switches . . . 6485
Part 88	Configuring DCBX and PFC
Chapter 246	Using DCBX and PFC 6489
	Understanding DCB Features and Requirements 6490
	Lossless Transport 6490
	PFC 6491
	Buffer Management 6491
	Physical Interfaces 6491
	ETS 6491
	DCBX 6492
	Understanding CoS Flow Control (Ethernet PAUSE and PFC) 6493
	General Information about Ethernet PAUSE and PFC and When to Use Them 6493
	Ethernet PAUSE 6494
	Symmetric Flow Control 6495
	Asymmetric Flow Control 6496

PFC	6499
Lossless Transport Support Summary	6502
Example: Configuring CoS PFC for FCoE Traffic	6504
Understanding DCBX	6514
DCBX Basics	6514
DCBX Modes and Support	6515
DCBX Modes (Versions)	6515
Autonegotiation	6517
CNA Support for DCBX Modes	6518
Interface Support for DCBX	6518
DCBX Attribute Types	6518
Asymmetric Attributes	6518
Symmetric Attributes	6519
DCBX Application Protocol TLV Exchange	6519
Application Protocol TLV Exchange	6519
FCoE Application Protocol TLV Exchange	6519
Disabling Application Protocol TLV Exchange	6520
DCBX and PFC	6520
DCBX and ETS	6521
Default DCBX ETS Advertisement	6521
ETS Advertisement and Peer Configuration	6521
ETS Recommendation TLV	6522
Configuring the DCBX Mode	6523
Configuring DCBX Autonegotiation	6524
Disabling the ETS Recommendation TLV	6527
Understanding DCBX Application Protocol TLV Exchange	6528
Applications	6528
Application Maps	6529
Classifying and Prioritizing Application Traffic	6530
Enabling Interfaces to Exchange Application Protocol Information	6531
Disabling DCBX Application Protocol Exchange	6531
Defining an Application for DCBX Application Protocol TLV Exchange	6532
Configuring an Application Map for DCBX Application Protocol TLV Exchange	6533
Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange	6534
Example: Configuring DCBX Application Protocol TLV Exchange	6535
Chapter 247 Learn About Technology	6547
Data Center Technology Overview Videos	6547
Learn About Video: Why Do We Need an IP Fabric?	6547
Learn About Video: What is the Best Control Plane Protocol to Use in a Data Center IP Fabric?	6547
Learn About Video: Why Use an Overlay Network in a Data Center?	6547
Conceptual Documents That Contain Technology Overview Videos	6548

Part 89	Configuration Statements and Operational Commands	
Chapter 248	Configuration Statements (FCoE on a Transit Switch)	6551
	family fcoe	6552
	oxid	6553
Chapter 249	Configuration Statements (DCBX and PFC)	6555
	application (Application Maps)	6556
	application (Applications)	6557
	application-map	6558
	application-maps	6559
	applications (Applications)	6560
	applications (DCBX)	6561
	code-points (Application Maps)	6562
	dcbx	6563
	dcbx-version	6564
	destination-port (Applications)	6565
	disable (DCBX)	6566
	enhanced-transmission-selection	6567
	ether-type	6568
	interface (DCBX)	6569
	no-recommendation-tlv	6570
	policy-options	6571
	priority-flow-control	6572
	protocol (Applications)	6573
	recommendation-tlv	6574
Chapter 250	Operational Commands (FCoE on a Transit Switch)	6575
	clear fip vlan-discovery statistics	6576
	show fip vlan-discovery	6577
Chapter 251	Operational Commands (DCBX and PFC)	6579
	show dcbx	6580
	show dcbx neighbors	6581

Guide 19 Traffic Management Feature Guide for QFX Switches

Part 90	CoS Overview	
Chapter 252	Basic Concepts	6607
	Overview of Junos OS CoS	6608
	CoS Standards	6608
	How Junos OS CoS Works	6609
	Default CoS Behavior	6610
	CoS Support on QFX Series Switches, EX4600 Switches, and QFabric Systems	6610
	Configuring CoS	6618
	Understanding Junos CoS Components	6621
	Code-Point Aliases	6622
	Policers	6622

Classifiers	6622
Forwarding Classes	6623
Forwarding Class Sets	6624
Flow Control (Ethernet PAUSE, PFC, and ECN)	6624
WRED Profiles and Tail Drop	6625
Schedulers	6626
Rewrite Rules	6626
Assigning CoS Components to Interfaces	6627
Understanding CoS Packet Flow	6628
Understanding Default CoS Settings	6630
Default Forwarding Classes and Queue Mapping	6631
Default Forwarding Class Sets (Priority Groups)	6631
Default Code-Point Aliases	6632
Default Classifiers	6633
Default Rewrite Rules	6636
Default Drop Profile	6636
Default Schedulers	6636
Default Scheduler Maps	6638
CoS Inputs and Outputs Overview	6639
Overview of Policers	6639
Policer Overview	6640
Policer Types	6640
Policer Actions	6641
Policer Colors	6642
Filter-Specific Policers	6642
Suggested Naming Convention for Policers	6642
Policer Counters	6643
Policer Algorithms	6643
How Many Policers Are Supported?	6643
Policers Can Limit Egress Firewall Filters	6644
Part 91	Classifying and Rewriting Traffic
Chapter 253	Using Classifiers, Forwarding Classes, and Rewrite Rules
Understanding CoS Classifiers	6650
Interfaces and Output Queues	6650
Classifier Support by Type	6651
Behavior Aggregate Classifiers	6651
Default Behavior Aggregate Classification	6652
Importing a Classifier	6652
PFC Priorities	6653
Fixed Classifiers on Ethernet Interfaces	6654
Multifield Classifiers	6654
MPLS EXP Classifiers	6655

Packet Classification for IRB Interfaces	6655
Defining CoS BA Classifiers (DSCP, DSCP IPv6, IEEE 802.1p)	6656
Example: Configuring Classifiers	6657
Understanding CoS MPLS EXP Classifiers and Rewrite Rules	6660
EXP Classifiers	6661
EXP Rewrite Rules	6662
Schedulers	6663
Configuring a Global MPLS EXP Classifier	6663
Configuring Rewrite Rules for MPLS EXP Classifiers	6664
Understanding Host Inbound Traffic Classification	6665
Understanding Default CoS Scheduling and Classification	6666
Default Classification	6666
Default Scheduling	6669
Default DCBX Advertisement	6671
Default Scheduling and Classification Summary	6671
Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces	6672
Supported Classifier and Rewrite Rule Types	6672
Ethernet Interfaces Supported for Classifier and Rewrite Rule Configuration	6674
Interface Types That Support Classifier and Rewrite Rule Configuration	6674
Classifier and Rewrite Rule Physical and Logical Ethernet Interface Support	6675
Routed VLAN Interfaces (RVIs) and Integrated Routing and Bridging (IRB) Interfaces	6677
Default Classifiers	6677
Default Rewrite Rules	6678
Classifier Precedence	6678
Classifier Precedence on Physical Ethernet Interfaces (QFX5200, QFX5100, EX4600, QFX3500, and QFX3600 Switches, and QFabric Systems)	6678
Classifier Precedence on Logical Ethernet Interfaces (All Switches)	6679
Classifier Behavior and Limitations	6680
Rewrite Rule Precedence and Behavior	6681
Classifier and Rewrite Rule Configuration Interaction with Ethernet Interface Configuration	6681
QFX5100, QFX5200, EX4600, QFX3500, and QFX3600 Switch Scenarios	6683
Understanding CoS Code-Point Aliases	6685
Defining CoS Code-Point Aliases	6687
Understanding CoS Forwarding Classes	6688
Default Forwarding Classes	6688
Forwarding Class Configuration Rules	6689
Queue Assignment Rules	6689
Scheduling Rules	6690
Rewrite Rules	6691
Lossless Transport Support	6691
Defining CoS Forwarding Classes	6691
Example: Configuring Forwarding Classes	6692

	Understanding CoS Forwarding Class Sets (Priority Groups)	6695
	Defining CoS Forwarding Class Sets	6696
	Example: Configuring Forwarding Class Sets	6697
	Understanding CoS Rewrite Rules	6701
	Defining CoS Rewrite Rules	6704
	Troubleshooting an Unexpected Rewrite Value	6705
Part 92	Scheduling Traffic	
Chapter 254	Using Schedulers	6711
	Understanding Default CoS Scheduling and Classification	6712
	Default Classification	6712
	Default Scheduling	6715
	Default DCBX Advertisement	6717
	Default Scheduling and Classification Summary	6717
	Understanding CoS Scheduling Behavior and Configuration Considerations . .	6718
	Understanding CoS Virtual Output Queues (VOQs) on QFX10000 Switches . .	6724
	VOQ Architecture	6724
	Round-Trip Time Buffering	6725
	Requesting and Granting Egress Port Bandwidth	6725
	VOQ Advantages	6726
	Eliminate Head-of-Line Blocking	6726
	Increase Fabric Efficiency and Utilization	6728
	Understanding CoS Port Schedulers on QFX Switches	6729
	Queue Scheduling Components	6730
	Default Schedulers	6732
	Scheduling Priority	6733
	Bandwidth Scheduling	6735
	Minimum Guaranteed Bandwidth	6735
	Maximum Bandwidth (Rate Shaping on Low Priority Queues and LAGs)	6736
	Limiting Bandwidth Consumed by Strict-High Priority Queues	6737
	Sharing Extra Bandwidth (Excess Rate on Low Priority Queues)	6738
	Scheduler Drop-Profile Maps	6739
	Buffer Size	6739
	Explicit Congestion Notification	6741
	Scheduler Maps	6741
	Defining CoS Queue Schedulers for Port Scheduling	6742
	Example: Configuring Queue Schedulers for Port Scheduling	6744
	Troubleshooting Egress Bandwidth That Exceeds the Configured Minimum Bandwidth	6749
	Troubleshooting Egress Bandwidth That Exceeds the Configured Maximum Bandwidth	6751
	Defining CoS Queue Scheduling Priority	6752
	Example: Configuring Queue Scheduling Priority	6753
	Understanding CoS Traffic Control Profiles	6757
	Understanding CoS Priority Group Scheduling	6758
	Priority Group Scheduling Components	6759
	Default Traffic Control Profile	6759

Guaranteed Rate (Minimum Guaranteed Bandwidth)	6759
Sharing Extra Bandwidth	6760
Shaping Rate (Maximum Bandwidth)	6760
Scheduler Maps	6760
Defining CoS Traffic Control Profiles (Priority Group Scheduling)	6761
Example: Configuring Traffic Control Profiles (Priority Group Scheduling)	6762
Understanding CoS Hierarchical Port Scheduling (ETS)	6765
Hierarchical Scheduling Tiers	6765
Hierarchical Scheduling and ETS	6766
ETS Advertisement in DCBX	6768
Hierarchical Scheduling Process	6768
Strict-High Priority Queues and Hierarchical Scheduling	6769
Default Hierarchical Scheduling	6770
Example: Configuring CoS Hierarchical Port Scheduling (ETS)	6771
Understanding CoS Priority Group and Queue Guaranteed Minimum	
Bandwidth	6795
Guaranteeing Bandwidth Using Hierarchical Scheduling	6796
Priority Group Guaranteed Rate (Guaranteed Minimum Bandwidth)	6797
Queue Transmit Rate (Guaranteed Minimum Bandwidth)	6797
Example: Configuring Minimum Guaranteed Output Bandwidth	6798
Understanding CoS Priority Group Shaping and Queue Shaping (Maximum	
Bandwidth)	6803
Priority Group Shaping	6803
Queue Shaping	6803
Shaping Maximum Bandwidth Using Hierarchical Scheduling	6804
Example: Configuring Maximum Output Bandwidth	6805
Understanding CoS WRED Drop Profiles	6809
Drop Profile Parameters	6809
Default Drop Profile	6810
Packet Drop Method	6811
Drop Profile Maps	6811
Congestion Prevention	6811
Configuring a WRED Drop Profile and Applying it to an Output Queue	6811
Drop Profiles on Explicit Congestion Notification Enabled Queues	6812
Configuring CoS WRED Drop Profiles	6813
Example: Configuring WRED Drop Profiles	6814
Configuring CoS Drop Profile Maps	6817
Example: Configuring Drop Profile Maps	6817
Understanding CoS Explicit Congestion Notification	6820
How ECN Works	6820
ECN Bits in the DiffServ Field	6821
End-to-End ECN Behavior	6822
ECN Compared to PFC and Ethernet PAUSE	6824
WRED Drop Profile Control of ECN Thresholds	6825
Support, Limitations, and Notes	6828
Example: Configuring ECN	6828

Part 93	Data Center Bridging (PFC, DCBX) and Flow Control	
Chapter 255	Using Data Center Bridging and Flow Control	6837
	Understanding DCB Features and Requirements	6838
	Lossless Transport	6839
	PFC	6839
	Buffer Management	6839
	Physical Interfaces	6839
	ETS	6840
	DCBX	6840
	Understanding CoS Hierarchical Port Scheduling (ETS)	6841
	Hierarchical Scheduling Tiers	6841
	Hierarchical Scheduling and ETS	6842
	ETS Advertisement in DCBX	6844
	Hierarchical Scheduling Process	6844
	Strict-High Priority Queues and Hierarchical Scheduling	6845
	Default Hierarchical Scheduling	6846
	Example: Configuring CoS Hierarchical Port Scheduling (ETS)	6847
	Disabling the ETS Recommendation TLV	6872
	Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows	6873
	Default Lossless Priority Configuration	6873
	Configuring Lossless Priorities	6876
	Configuring Lossless Forwarding Classes (Packet Drop Attribute)	6876
	Congestion Notification Profiles (PFC Configuration)	6878
	Configuring DCBX (Application Protocol TLV Exchange)	6884
	Fate Sharing Among Traffic Classes	6885
	Transit Switch Configuration Versus FCoE-FC Gateway	
	Configuration	6886
	Configuration Results and Commit Checks	6887
	Configuration Rules and Recommendations	6889
	Lossless Transport Features Introduced in Junos OS Release 12.3 (Legacy	
	Non-ELS CLI)	6890
	Backward Compatibility with Junos OS Releases Earlier Than Release 12.3	
	(Legacy Non-ELS CLI)	6891
	Understanding CoS Flow Control (Ethernet PAUSE and PFC)	6893
	General Information about Ethernet PAUSE and PFC and When to Use	
	Them	6893
	Ethernet PAUSE	6894
	Symmetric Flow Control	6895
	Asymmetric Flow Control	6895
	PFC	6899
	Lossless Transport Support Summary	6902
	Configuring CoS PFC (Congestion Notification Profiles)	6905
	Example: Configuring CoS PFC for FCoE Traffic	6908
	Troubleshooting Dropped FCoE Traffic	6917
	Understanding PFC Functionality Across Layer 3 Interfaces	6921
	Example: Configuring PFC Across Layer 3 Interfaces	6923
	Example: Configuring CoS for FCoE Transit Switch Traffic Across an MC-LAG	6943

Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch)	6969
Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface	6977
Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces	6986
Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI)	7000
Understanding DCBX	7016
DCBX Basics	7016
DCBX Modes and Support	7018
DCBX Modes (Versions)	7018
Autonegotiation	7020
CNA Support for DCBX Modes	7020
Interface Support for DCBX	7020
DCBX Attribute Types	7020
Asymmetric Attributes	7021
Symmetric Attributes	7021
DCBX Application Protocol TLV Exchange	7022
Application Protocol TLV Exchange	7022
FCoE Application Protocol TLV Exchange	7022
Disabling Application Protocol TLV Exchange	7023
DCBX and PFC	7023
DCBX and ETS	7023
Default DCBX ETS Advertisement	7023
ETS Advertisement and Peer Configuration	7024
ETS Recommendation TLV	7024
Configuring the DCBX Mode	7026
Configuring DCBX Autonegotiation	7027
Understanding DCBX Application Protocol TLV Exchange	7030
Applications	7030
Application Maps	7031
Classifying and Prioritizing Application Traffic	7032
Enabling Interfaces to Exchange Application Protocol Information	7033
Disabling DCBX Application Protocol Exchange	7033
Defining an Application for DCBX Application Protocol TLV Exchange	7034
Configuring an Application Map for DCBX Application Protocol TLV Exchange	7035
Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange	7036
Example: Configuring DCBX Application Protocol TLV Exchange	7037
Chapter 256 Learn About Technology	7049
Data Center Technology Overview Videos	7049
Learn About Video: Why Do We Need an IP Fabric?	7049
Learn About Video: What is the Best Control Plane Protocol to Use in a Data Center IP Fabric?	7049
Learn About Video: Why Use an Overlay Network in a Data Center?	7049

	Conceptual Documents That Contain Technology Overview Videos	7050
Part 94	Configuration Statements and Operational Commands	
Chapter 257	Configuration Statements (Basic Concepts)	7053
	class-of-service	7054
	traceoptions (Class of Service)	7058
Chapter 258	Configuration Statements (Classifiers and Rewrite Rules)	7061
	class (Forwarding Classes)	7062
	class (Forwarding Class Sets)	7063
	classifiers	7064
	code-point (Rewrite Rules)	7066
	code-point-aliases	7067
	code-points (CoS)	7068
	dscp	7069
	dscp-ipv6	7071
	exp	7073
	forwarding-class	7075
	forwarding-class-sets	7076
	forwarding-classes	7077
	ieee-802.1	7079
	import	7081
	interfaces (Class of Service)	7082
	loss-priority (Classifiers)	7084
	loss-priority (Rewrite Rules)	7085
	queue-num	7086
	rewrite-rules	7088
	unit	7089
Chapter 259	Configuration Statements (Scheduling)	7091
	buffer-size	7092
	drop-probability	7097
	drop-profile	7098
	drop-profile-map	7098
	drop-profiles	7099
	excess-rate	7100
	explicit-congestion-notification	7101
	fill-level	7102
	forwarding-class	7104
	forwarding-class-set	7105
	guaranteed-rate	7106
	interpolate	7107
	loss-priority (Drop Profiles)	7108
	output-traffic-control-profile	7108
	priority (Schedulers)	7109
	protocol (Drop Profile Map)	7110
	scheduler	7111
	scheduler-map	7111
	scheduler-maps	7112

	schedulers	7113
	shaping-rate	7114
	traffic-control-profiles	7116
	transmit-rate	7117
Chapter 260	Configuration Statements (Data Center Bridging and Flow Control)	7121
	application (Application Maps)	7122
	application (Applications)	7123
	application-map	7124
	application-maps	7125
	applications (Applications)	7126
	applications (DCBX)	7127
	cable-length (Congestion Notification)	7128
	code-point (Input Congestion Notification)	7129
	code-point (Output Congestion Notification)	7130
	code-points (Application Maps)	7131
	congestion-notification-profile	7132
	dcbx	7134
	dcbx-version	7135
	destination-port (Applications)	7136
	disable (DCBX)	7137
	enhanced-transmission-selection	7138
	ether-type	7139
	flow-control	7140
	flow-control-queue (Output Congestion Notification)	7142
	ieee-802.1 (Input Congestion Notification)	7143
	ieee-802.1 (Output Congestion Notification)	7144
	input (Congestion Notification)	7145
	interface (DCBX)	7146
	interfaces (Class of Service)	7147
	mrui	7149
	output (Congestion Notification)	7150
	pfc (Input Congestion Notification)	7151
	policy-options	7152
	priority-flow-control	7153
	protocol (Applications)	7154
	recommendation-tlv	7155
	rx-buffers	7156
	tx-buffers	7158
Chapter 261	Operational Commands (Basic Concepts)	7161
	Monitoring Interfaces That Have CoS Components	7161
	show class-of-service	7163
	show class-of-service interface	7167
	show class-of-service shared-buffer	7196
	show pfe filter hw summary	7199
	show pfe next-hop	7201
	show pfe route	7206
	show pfe terse	7215
	show pfe version	7217

	show interfaces voq	7218
Chapter 262	Operational Commands (Classifiers and Rewrite Rules)	7231
	Monitoring CoS Classifiers	7231
	Monitoring CoS Forwarding Classes	7232
	Monitoring CoS Rewrite Rules	7235
	Monitoring CoS Code-Point Value Aliases	7236
	show class-of-service classifier	7238
	show class-of-service code-point-aliases	7241
	show class-of-service forwarding-class	7243
	show class-of-service forwarding-class-set	7245
	show class-of-service forwarding-table	7247
	show class-of-service forwarding-table classifier	7251
	show class-of-service forwarding-table classifier mapping	7253
	show class-of-service forwarding-table rewrite-rule	7255
	show class-of-service forwarding-table rewrite-rule mapping	7257
	show class-of-service interface	7258
	show class-of-service multi-destination	7287
	show class-of-service rewrite-rule	7288
Chapter 263	Operational Commands (Scheduling)	7291
	Monitoring CoS Scheduler Maps	7291
	show class-of-service drop-profile	7293
	show class-of-service forwarding-table	7296
	show class-of-service forwarding-table drop-profile	7300
	show class-of-service forwarding-table scheduler-map	7302
	show class-of-service interface	7304
	show class-of-service scheduler-map	7333
	show class-of-service traffic-control-profile	7336
	show interfaces queue	7340
	show interfaces voq	7379
Chapter 264	Operational Commands (Data Center Bridging and Flow Control)	7393
	show class-of-service congestion-notification	7394
	show dcbx	7397
	show dcbx neighbors	7398
Part 95	Index	
	Index	7423

List of Figures

Part 3	Command Line Interface (CLI)	
Chapter 3	Understanding the CLI	37
	Figure 1: Commands That Combine Other Commands	58
	Figure 2: CLI Command Hierarchy	61
	Figure 3: Command Output Options	63
	Figure 4: Configuration Mode Hierarchy of Statements	68
Part 4	Dynamic Host Control Protocol (DHCP)	
Chapter 4	Understanding DHCP	73
	Figure 5: DHCP Client/Server Model	74
	Figure 6: DHCP Four-Step Transfer	76
Part 17	Managing Configuration Files	
Chapter 36	Loading, Uploading, Transferring, and Saving Configuration Files	1325
	Figure 7: Overriding the Current Configuration	1328
	Figure 8: Using the replace Option	1328
	Figure 9: Using the merge Option	1329
	Figure 10: Using a Patch File	1329
	Figure 11: Using the set Option	1330
Part 20	Configuration	
Chapter 46	Configuring Automation	1419
	Figure 12: Commit Script Input and Output	1424
	Figure 13: Standard Commit Model	1427
	Figure 14: Commit Model with Commit Scripts Added	1427
	Figure 15: Configuration Evaluation by Multiple Commit Scripts	1429
	Figure 16: Op Script Input and Output	1436
Chapter 47	Configuring sFlow Technology	1449
	Figure 17: sFlow Technology Monitoring System	1454
Chapter 48	Configuring SNMP	1459
	Figure 18: Setting Thresholds	1464
	Figure 19: Inform Request and Response	1515
Part 28	Ethernet Ring Protection	
Chapter 72	Understanding Ethernet Ring Protection	2055
	Figure 20: Protocol Packets from the Network to the Router	2059

	Figure 21: Protocol Packets from the Router or Switch to the Network	2059
	Figure 22: Ethernet Ring Protection Switching Example	2067
Part 29	Bridging and VLANs	
Chapter 73	Using Bridging and VLANs	2085
	Figure 23: IRB with One Switch	2126
Part 31	Spanning Trees	
Chapter 75	Using Spanning Trees	2147
	Figure 24: Network Topology for MSTP	2151
	Figure 25: Network Topology for RSTP	2174
	Figure 26: BPDU Protection Topology	2201
	Figure 27: Network Topology for Loop Protection	2208
	Figure 28: Network Topology for Root Protection	2213
Part 34	Reflective Relay	
Chapter 78	Using Reflective Relay	2249
	Figure 29: Reflective Relay Topology	2253
Part 36	Overview	
Chapter 88	OVSDB and VXLAN Overview	2439
	Figure 30: High-Level View of NSX for Multi-Hypervisor Architecture	2440
	Figure 31: Integration of Juniper Networks Device into NSX for Multi-Hypervisor Environment	2441
	Figure 32: VXLAN Packet Format	2443
Part 37	Configuring OVSDB and VXLAN	
Chapter 89	Configuring OVSDB-Managed VXLANs with an SDN Controller	2455
	Figure 33: VXLAN-OVSDB Layer 2 Gateway Topology	2474
	Figure 34: VXLAN/OVSDB Layer 2 Gateway Topology	2481
Chapter 90	Configuring VXLANs Without an SDN Controller	2491
	Figure 35: QFX5100 Acting as a VXLAN Transit Switch	2495
	Figure 36: QFX5100 Acting as a VTEP	2497
Part 42	Nonstop Bridging	
Chapter 98	Configuring Nonstop Bridging	2589
	Figure 37: Nonstop Bridging Switchover Preparation Process	2590
	Figure 38: Nonstop Bridging During a Switchover	2591
Part 43	Nonstop Active Routing	
Chapter 99	Configuring Nonstop Active Routing	2597
	Figure 39: Nonstop Active Routing Switchover Preparation Process	2598
	Figure 40: Nonstop Active Routing During a Switchover	2599

Part 44	Graceful Routing Engine Switchover	
Chapter 100	Configuring Graceful Routing Engine Switchover	2617
	Figure 41: Preparing for a Graceful Routing Engine Switchover	2620
	Figure 42: Graceful Routing Engine Switchover Process	2621
Part 45	Virtual Router Redundancy Protocol	
Chapter 101	Configuring Virtual Router Redundancy Protocol	2633
	Figure 43: Basic VRRP Topology	2635
	Figure 44: VRRP Load-Sharing Configuration	2637
Part 50	IP Directed Broadcast	
Chapter 114	Understanding IP Directed Broadcast	2849
	Figure 45: Topology for IP Directed Broadcast	2853
Part 54	Local Link Bias	
Chapter 118	Understanding Local Link Bias	2895
	Figure 46: Egress Traffic Flow with Local Link Bias	2895
	Figure 47: Egress Traffic Flow without Local Link Bias	2896
Part 55	Redundant Trunk Groups	
Chapter 119	Understanding Redundant Trunk Groups	2901
	Figure 48: Redundant Trunk Group, Link 1 Active	2902
	Figure 49: Redundant Trunk Group, Link 2 Active	2902
	Figure 50: Topology for Configuring the Redundant Trunk Links	2905
Part 57	Uplink Failure Detection	
Chapter 121	Understanding Uplink Failure Detection	2929
	Figure 51: Uplink Failure Detection Configuration on Switches	2930
	Figure 52: Uplink Failure Detection Configuration on Switches	2933
Part 59	Configuring Routing Options	
Chapter 138	Examples: Configuring BFD for Static Routes	3237
	Figure 53: Customer Routes Connected to a Service Provider	3242
	Figure 54: BFD Enabled on Qualified Next Hops	3247
Chapter 139	Example: Configuring BFD Authentication for Static Routes	3253
	Figure 55: Customer Routes Connected to a Service Provider	3256
Part 61	Overview	
Chapter 143	BGP Overview	3593
	Figure 56: ASs, EBGp, and IBGP	3594

Part 62	Configuring BGP	
Chapter 144	Basic BGP Configuration	3601
	Figure 57: BGP Peering Session	3601
	Figure 58: Typical Network with BGP Peer Sessions	3603
	Figure 59: Typical Network with BGP Peer Sessions	3610
	Figure 60: Internal and External BGP	3624
	Figure 61: Typical Network with IBGP Sessions	3627
	Figure 62: Typical Network with IBGP Sessions	3637
Chapter 145	BGP Path Attribute Configuration	3651
	Figure 63: Typical Network with IBGP Sessions and Multiple Exit Points	3653
	Figure 64: Default MED Example	3666
	Figure 65: Typical Network with IBGP Sessions and Multiple Exit Points	3669
	Figure 66: Typical Network with IBGP Sessions and Multiple Exit Points	3681
	Figure 67: Topology for Delaying the MED Update	3696
	Figure 68: Local AS Configuration	3707
	Figure 69: Topology for Configuring the Local AS	3710
	Figure 70: Topology for Configuring a Private Local AS	3720
	Figure 71: Advertisement of Multiple Paths in BGP	3727
Chapter 146	BGP Policy Configuration	3765
	Figure 72: BGP Prefix-Based Outbound Route Filtering	3774
	Figure 73: Typical Network with EBGP Multihop Sessions	3778
	Figure 74: BGP Preference Value Topology	3789
	Figure 75: Topology for Ignoring the AS-Path Length	3798
	Figure 76: Topology for Removing a Private AS from the Advertised AS Path	3805
Chapter 147	BGP BFD Configuration	3811
	Figure 77: Typical Network with IBGP Sessions	3813
Chapter 148	BGP Load Balancing Configuration	3827
	Figure 78: BGP Load Balancing	3829
	Figure 79: Topology for Accepting a Remote Next Hop	3834
	Figure 80: Advertisement of Multiple Paths in BGP	3845
Chapter 149	IBGP Scaling Configuration	3873
	Figure 81: Simple Route Reflector Topology (One Cluster)	3874
	Figure 82: Basic Route Reflection (Multiple Clusters)	3875
	Figure 83: Hierarchical Route Reflection (Clusters of Clusters)	3875
	Figure 84: IBGP Network Using a Route Reflector	3877
	Figure 85: BGP Confederations	3891
	Figure 86: Typical Network Using BGP Confederations	3892
Chapter 150	BGP Security Configuration	3897
	Figure 87: Authentication for BGP	3900
	Figure 88: Typical Network with BGP Peer Sessions	3905
	Figure 89: TCP Maximum Segment Size for BGP	3913
Chapter 151	BGP Flap Configuration	3919
	Figure 90: Topology for the EBGP Case	3922
	Figure 91: Topology for the RR Case	3922

	Figure 92: BGP Flap Damping Topology	3928
	Figure 93: MBGP MVPN with BGP Route Flap Damping	3936
Chapter 156	Overview	4157
	Figure 94: Install Default Route to Nearest Routing Device That Operates at Both Level 1 and Level 2	4161
Part 65	Configuring IS-IS Features	
Chapter 157	Configuring IS-IS	4165
	Figure 95: Link Protection and Node-Link Protection Comparison for IS-IS Routes	4167
	Figure 96: Simple IS-IS Topology	4170
	Figure 97: IS-IS Multi-Level Topology	4175
	Figure 98: IS-IS Route Redistribution Topology	4186
	Figure 99: IS-IS IPv4 and IPv6 Unicast Topologies	4195
	Figure 100: Configuring IS-IS Multicast Topology	4204
	Figure 101: Link Protection and Node-Link Protection Comparison for IS-IS Routes	4219
	Figure 102: IS-IS Node-Link Protection Topology	4222
	Figure 103: IS-IS with a Default Route to an ISP	4232
	Figure 104: IS-IS Checksum Topology	4240
Chapter 158	Configuring BFD Authentication for IS-IS	4243
	Figure 105: Configuring BFD for IS-IS	4245
	Figure 106: IS-IS BFD Authentication Topology	4251
Chapter 159	Configuring Hitless Authentication Key Rollover for IS-IS	4255
	Figure 107: Hitless Authentication Key Rollover for IS-IS	4257
Part 66	Overview	
Chapter 162	OSPF Overview	4385
	Figure 108: OSPF Three-Way Handshake	4389
Part 67	Configuring OSPF	
Chapter 163	Basic OSPF Area Configuration	4397
	Figure 109: Multiarea OSPF Topology	4405
	Figure 110: OSPF Topology with a Virtual Link	4406
	Figure 111: Typical Single-Area OSPF Network Topology	4407
	Figure 112: Typical Multiarea OSPF Network Topology	4409
Chapter 164	Advanced OSPF Area Configuration	4413
	Figure 113: OSPF AS Network with Stub Areas and NSSAs	4413
	Figure 114: OSPF Network Topology with Stub Areas and NSSAs	4416
	Figure 115: OSPF Network Topology with Stub Areas and NSSAs	4420
Chapter 165	OSPF Interface Configuration	4431
	Figure 116: IPv4 Unicast Realm	4448
Chapter 166	OSPF Route Control Configuration	4451

	Figure 117: Summarizing Ranges of Routes in OSPF	4453
	Figure 118: OSPF Metric Configuration	4464
Chapter 171	OSPF Policy Configuration	4527
	Figure 119: Sample Topology Used for an OSPF Export Network Summary Policy	4545
	Figure 120: Sample Topology Used for an OSPF Import Network Summary Policy	4556
Chapter 176	Overview	4677
	Figure 121: Distance-Vector Protocol	4678
	Figure 122: Split Horizon Example	4680
	Figure 123: Poison Reverse Example	4681
	Figure 124: Limitations of Unidirectional Connectivity	4681
Part 70	Configuring RIP	
Chapter 177	Example: Configuring RIP	4685
	Figure 125: Sample RIP Network Topology	4686
Chapter 178	Example: Configuring Authentication for RIP Routes	4693
	Figure 126: RIP Authentication Network Topology	4694
Chapter 179	Example: Configuring BFD for RIP	4701
	Figure 127: RIP BFD Network Topology	4704
Chapter 180	Example: Configuring BFD Authentication for RIP	4709
	Figure 128: RIP BFD Authentication Network Topology	4711
Chapter 181	Example: Applying Policies to RIP Routes Imported from Neighbors	4719
	Figure 129: RIP Import Policy Network Topology	4720
Chapter 182	Examples: Controlling Traffic with Metrics in a RIP Network	4725
	Figure 130: Controlling Traffic in a RIP Network with the Incoming Metric	4727
	Figure 131: Controlling Traffic in a RIP Network with the Outgoing Metric	4728
	Figure 132: RIP Incoming Metrics Network Topology	4730
Chapter 183	Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets	4735
	Figure 133: Sending and Receiving RIPv1 and RIPv2 Packets Network Topology	4736
Chapter 184	Example: Redistributing Routes Among RIP Instances	4741
	Figure 134: Redistributing Routes Between RIP Instances Network Topology	4743
Chapter 185	Example: Configuring RIP Timers	4749
	Figure 135: RIP Timers Network Topology	4751
Chapter 186	Example: Tracing RIP Protocol Traffic	4757
	Figure 136: RIP Trace Operations Network Topology	4759
Part 71	LDP	
Chapter 189	Using LDP	4803

	Figure 137: Swap and Push When LDP LSPs Are Tunneled Through RSVP LSPs	4806
	Figure 138: Double Push When LDP LSPs Are Tunneled Through RSVP LSPs	4807
Part 72	MPLS	
Chapter 191	Using MPLS	4929
	Figure 139: Label Encoding	4952
	Figure 140: MPLS Label Swapping	4953
	Figure 141: ASs, EBGp, and IBGP	4956
	Figure 142: Carrier-of-Carriers VPN Architecture	4965
	Figure 143: Detours Established for an LSP Using Fast Reroute	4968
	Figure 144: Detour After the Link from Router B to Router C Fails	4968
	Figure 145: Detours Merging into Other Detours	4969
	Figure 146: Ethernet over MPLS Layer 2 Circuit	4973
	Figure 147: MPLS-Based Layer 3 VPN	4994
	Figure 148: IPv6 Networks Linked by MPLS IPv4 Tunnels	5003
	Figure 149: Virtual Machines on Either Side of Routers	5015
Part 73	RSVP	
Chapter 193	Using RSVP	5259
	Figure 150: RSVP Reservation Request and Data Flow	5263
	Figure 151: Sample Topology	5274
Chapter 195	Overview	5389
	Figure 152: Multicast Terminology in an IP Network	5393
	Figure 153: Converting MAC Addresses to Multicast Addresses	5397
Part 74	Managing Group Membership	
Chapter 198	Using MLD	5441
	Figure 154: Routing Devices Start Up on a Subnet	5442
	Figure 155: Querier Routing Device Is Determined	5442
	Figure 156: General Query Message Is Issued	5442
	Figure 157: Reports Are Received by the Querier Routing Device	5443
	Figure 158: Host Has No Interested Receivers and Sends a Done Message to Routing Device	5443
	Figure 159: Host Address Timer Expires and Address Is Removed from Multicast Address List	5443
	Figure 160: Routing Devices Start Up on a Subnet	5445
	Figure 161: Querier Routing Device Is Determined	5445
	Figure 162: General Query Message Is Issued	5446
	Figure 163: Reports Are Received by the Querier Routing Device	5446
	Figure 164: Host Has No Interested Receivers and Sends a Done Message to Routing Device	5446
	Figure 165: Host Address Timer Expires and Address Is Removed from Multicast Address List	5447
Chapter 199	Using MLD Snooping	5467
	Figure 166: Multicast Traffic Flow with MLD Snooping Enabled	5468

	Figure 167: Scenario 1: Switch Forwarding Multicast Traffic to a Multicast Router and Hosts	5472
	Figure 168: Scenario 2: Switch Forwarding Multicast Traffic to Another Switch	5473
	Figure 169: Scenario 3: Switch Connected to Hosts Only (No MLD Querier) . . .	5474
	Figure 170: Scenario 4: Layer 2/Layer 3 Switch Forwarding Multicast Traffic Between VLANs	5475
	Figure 171: MLD Snooping Topology Example	5482
Part 75	Configuring PIM	
Chapter 201	Using PIM Sparse Mode	5509
	Figure 172: Rendezvous Point As Part of the RPT and SPT	5511
	Figure 173: Join Suppression	5519
Chapter 202	Using PIM Dense Mode and PIM Sparse-Dense Mode	5523
	Figure 174: Multicast Traffic Flooded from the Source Using PIM Dense Mode . .	5524
	Figure 175: Prune Messages Sent Back to the Source to Stop Unwanted Multicast Traffic	5525
Chapter 203	Using Source-Specific Multicast	5529
	Figure 176: Receiver Announces Desire to Join Group G and Source S	5532
	Figure 177: Router 3 (Last-Hop Router) Joins the Source Tree	5532
	Figure 178: (S,G) State Is Built Between the Source and the Receiver	5532
	Figure 179: Network on Which to Configure PIM SSM	5533
	Figure 180: Receiver Sends Messages to Join Group G and Source S	5538
	Figure 181: Router 3 (Last-Hop Router) Joins the Source Tree	5538
	Figure 182: (S,G) State Is Built Between the Source and the Receiver	5539
	Figure 183: Simple RPF Topology	5539
Chapter 209	Using PIM RPT and SPT Cutover	5579
	Figure 184: Building an RPT Between the RP and the Receiver	5581
	Figure 185: PIM Register Message and PIM Join Message Exchanged	5582
	Figure 186: Traffic Sent from the Source to the RP Router	5583
	Figure 187: Traffic Sent from the RP Router Toward the Receiver	5583
	Figure 188: Receiver DR Sends a PIM Join Message to the Source	5585
	Figure 189: PIM Prune Message Is Sent from the Receiver's DR Toward the RP Router	5586
	Figure 190: RP Router Receives PIM Prune Message	5586
	Figure 191: RP Router Sends a PIM Prune Message to the Source DR	5587
	Figure 192: Source's DR Stops Sending Duplicate Multicast Packets Toward the RP Router	5587
	Figure 193: PIM Assert Topology	5589
Part 76	Configuring MSDP	
Chapter 210	Using MSDP	5597
	Figure 194: Accepting Multicast Traffic from a Remote Source	5602
	Figure 195: Source-Active Message Flooding	5607

Part 78	Firewall Filters	
Chapter 221	Configuring Firewall Filters	5951
	Figure 196: Evaluation of Terms Within a Firewall Filter	5955
	Figure 197: Application of Firewall Filters to Control Packet Flow	5956
Part 79	Policers	
Chapter 222	Configuring Policers	5999
	Figure 198: Flow of Tricolor Marking Policer Operation	6000
Part 80	Configuring Port Security	
Chapter 223	Port Security	6027
	Figure 199: Switch Relays DHCP Requests to Server	6041
Part 81	Configuring Device Security	
Chapter 224	Device Security	6045
	Figure 200: Symmetrically Routed Interfaces	6052
	Figure 201: Asymmetrically Routed Interfaces	6053
Part 84	Port Mirroring	
Chapter 236	Configuring Port Mirroring	6227
	Figure 202: Network Topology for Local Port Mirroring Example	6236
Part 87	Configuring FCoE on a Transit Switch	
Chapter 245	Using FCoE on a Transit Switch	6431
	Figure 203: FCoE Transit Switch Connecting FCoE Devices to an FC Switch	6434
	Figure 204: ENode Components	6438
	Figure 205: Supported Topology for an MC-LAG on an FCoE Transit Switch	6452
	Figure 206: Supported Topology for an MC-LAG on an FCoE Transit Switch	6457
Part 88	Configuring DCBX and PFC	
Chapter 246	Using DCBX and PFC	6489
	Figure 207: PFC for FCoE Traffic Configuration Components Block Diagram	6507
Part 90	CoS Overview	
Chapter 252	Basic Concepts	6607
	Figure 208: Packet Flow Across the Network	6610
	Figure 209: CoS Classifier, Queues, and Scheduler	6629
	Figure 210: Packet Flow Through Configurable CoS Components	6630
	Figure 211: Flow of Tricolor Marking Policer Operation	6640
Part 92	Scheduling Traffic	
Chapter 254	Using Schedulers	6711

Figure 212: Congestion Occurs on EP-2	6726
Figure 213: EP-2 Backpressures IP-1	6726
Figure 214: Backpressure from EP-2 Causes IP-1 to Buffer Traffic Instead of Sending Traffic, Affecting EP-3	6727
Figure 215: Congestion on EP-2 Clears, Allowing IP-1 to Resume Sending Traffic to Both Egress Ports	6727
Figure 216: Each Egress Port Has a Separate Virtual Output Queue on IP-1	6727
Figure 217: Congestion on EP-2 Does Not Affect Uncongested Port EP-3	6728
Figure 218: Hierarchical Scheduling Tiers	6767
Figure 219: Hierarchical Scheduling Packet Flow	6769
Figure 220: Hierarchical Port Scheduling Components Block Diagram	6776
Figure 221: Hierarchical Port Scheduling Packet Flow Block Diagram	6776
Figure 222: Allocating Guaranteed Bandwidth Using Hierarchical Scheduling	6796
Figure 223: Setting Maximum Bandwidth Using Hierarchical Scheduling	6804
Figure 224: Explicit Congestion Notification	6822

Part 93

Chapter 255

Data Center Bridging (PFC, DCBX) and Flow Control

Using Data Center Bridging and Flow Control	6837
Figure 225: Hierarchical Scheduling Tiers	6843
Figure 226: Hierarchical Scheduling Packet Flow	6845
Figure 227: Hierarchical Port Scheduling Components Block Diagram	6852
Figure 228: Hierarchical Port Scheduling Packet Flow Block Diagram	6852
Figure 229: PFC for FCoE Traffic Configuration Components Block Diagram	6910
Figure 230: Enabling PFC Across Layer 3 Interface Hops	6922
Figure 231: Enabling PFC Across Layer 3 Interface Hops	6925
Figure 232: Supported Topology for an MC-LAG on an FCoE Transit Switch	6946
Figure 233: Topology of the Two Lossless FCoE Priorities Example	6988
Figure 234: Topology of the Lossless FCoE and iSCSI Priorities Example	7002

List of Tables

	About the Documentation	cxxxix
	Table 1: Notice Icons	cxli
	Table 2: Text and Syntax Conventions	cxli
Part 3	Command Line Interface (CLI)	
Chapter 3	Understanding the CLI	37
	Table 3: ELS Support	42
	Table 4: Renaming the ethernet-switching-options hierarchy	47
	Table 5: RTG Statements	48
	Table 6: Deleted Statements	48
	Table 7: Port Mirroring hierarchy	49
	Table 8: Layer 2 Control Protocol	49
	Table 9: dot1q-tunneling	50
	Table 10: mac-table-aging-time statement	50
	Table 11: Nonstop Bridging statement	50
	Table 12: Port Security statements	51
	Table 13: DHCP Snooping Statements	52
	Table 14: VLAN hierarchy	52
	Table 15: Statements Moved to a Different Hierarchy	54
	Table 16: Changes to the Storm Control Profile hierarchy level	55
	Table 17: Changes to the Interfaces hierarchy	56
	Table 18: IGMP Snooping hierarchy	57
	Table 19: Commonly Used Operational Mode Commands	60
	Table 20: Summary of Configuration Mode Commands	65
	Table 21: Configuration Mode Top-Level Statements	67
Part 4	Dynamic Host Control Protocol (DHCP)	
Chapter 4	Understanding DHCP	73
	Table 22: Legacy DHCP and Extended DHCP Server Hierarchy Levels	75
	Table 23: DHCP Client Settings	77
Part 6	Junos OS With Upgraded FreeBSD	
Chapter 6	Understanding Junos OS With Upgraded FreeBSD	89
	Table 24: Upgraded FreeBSD Kernel Support by Hardware Platform	90
	Table 25: New and Changed Commands and Statements for Junos OS with Upgraded FreeBSD	91
	Table 26: Deprecated Commands and Statements for Junos OS with Upgraded FreeBSD	91

	Table 27: Upgrade Path to Junos OS with the Upgraded FreeBSD	97
Part 7	Licenses	
Chapter 7	Using Licenses	107
	Table 28: Disaggregated Junos OS Feature Licenses and Model Numbers for QFX Series Devices	109
	Table 29: Standard Junos OS Feature Licenses and Model Numbers for QFX Series Devices	109
	Table 30: Upgrade Licenses for Enhancing Port Capacity	114
	Table 31: Port Activation License Model for MX104 Routers	115
Part 12	Routine Monitoring	
Chapter 13	Understanding Routine Monitoring	165
	Table 32: Summary of System Process Information Output Fields	165
	Table 33: Summary of Key System Properties Output Fields	166
Part 15	Configuration Statements and Operational Commands	
Chapter 25	Autoinstallation Operational Commands	323
	Table 34: show system autoinstallation status Output Fields	324
Chapter 26	Basic System Management Operational Commands	327
	Table 35: request system storage cleanup Output Fields	401
Chapter 27	CLI Operational Commands	427
	Table 36: show cli Output Fields	428
	Table 37: show cli authorization Output Fields	430
	Table 38: show cli directory Output Fields	434
	Table 39: show cli history Output Fields	435
Chapter 28	Licensing Operational Commands	437
	Table 40: show system license Output Fields	441
Chapter 29	NTP Operational Commands	449
	Table 41: show ntp associations Output Fields	450
	Table 42: show ntp status Output Fields	452
Chapter 30	PTP Operational Command	455
	Table 43: show ptp global-information Output Fields	456
Chapter 31	Routine Monitoring Operational Commands	457
	Table 44: show chassis alarms Output Fields	466
	Table 45: show chassis led Output Fields	474
	Table 46: show chassis environment Output Fields	483
	Table 47: show chassis environment fpc Output Fields	543
	Table 48: show chassis environment pem Output Fields	568
	Table 49: show chassis environment routing-engine Output Fields	577
	Table 50: show chassis fabric errors Output Fields	582
	Table 51: show chassis fabric fpcs Output Fields	588
	Table 52: show chassis fabric plane-location Output Fields	625

	Table 53: show chassis fabric sibs Output Fields	630
	Table 54: show chassis fabric summary Output Fields	641
	Table 55: show chassis fabric topology Output Fields	649
	Table 56: show chassis fan Output Fields	680
	Table 57: show chassis firmware Output Fields	695
	Table 58: show chassis fpc Output Fields	711
	Table 59: Routing Engines Displaying DIMM Information	746
	Table 60: show chassis hardware Output Fields	751
	Table 61: show chassis lcd Output Fields	923
	Table 62: show chassis led Output Fields	936
	Table 63: show chassis location Output Fields	947
	Table 64: show chassis mac-addresses Output Fields	951
	Table 65: show chassis pic Output Fields	958
	Table 66: show chassis routing-engine Output Fields	973
	Table 67: show chassis zones Output Fields	997
	Table 68: show subscribers Output Fields	1010
	Table 69: show system alarms Output Fields	1028
	Table 70: show system buffers Output Fields	1048
	Table 71: show system certificate Output Fields	1052
	Table 72: show system commit Output Fields	1054
	Table 73: show system connections Output Fields	1063
	Table 74: show system core-dumps Output Fields	1082
	Table 75: show system directory-usage Output Fields	1095
	Table 76: show system processes Output Fields	1105
	Table 77: show system resource-cleanup processes Output Fields	1129
	Table 78: show system services service-deployment Output Fields	1131
	Table 79: show system storage Output Fields	1180
	Table 80: show system uptime Output Fields	1187
	Table 81: show system users Output Fields	1192
	Table 82: show system virtual-memory Output Fields	1198
	Table 83: traceroute Output Fields	1271
	Table 84: traceroute monitor Output Fields	1274
Part 16	Configuration and File Management Overview	
Chapter 33	Understanding Configuration Files and the configure Command	1309
	Table 85: Configuration File Terms	1309
	Table 86: Forms of the configure Command	1311
Part 17	Managing Configuration Files	
Chapter 36	Loading, Uploading, Transferring, and Saving Configuration Files	1325
	Table 87: Options for the load Command	1332
Part 18	Configuration Statements and Operational Commands	
Chapter 42	Operational Commands	1367
	Table 88: show system commit Output Fields	1388

Part 19	Overview	
Chapter 43	Understanding Network Management	1401
	Table 89: Device and Network Management Features on the QFX Series, OCX Series, and EX4600	1401
Part 20	Configuration	
Chapter 45	Configuring Network Management	1409
	Table 90: Output Control Keys for the monitor interface Command	1417
Chapter 46	Configuring Automation	1419
	Table 91: Differences Between Persistent and Transient Changes	1431
	Table 92: Commit Script Configuration and Operational Mode Commands	1445
Chapter 48	Configuring SNMP	1459
	Table 93: RMON Event Table	1465
	Table 94: RMON Alarm Table	1465
	Table 95: jnxRmon Alarm Table	1466
	Table 96: RMON History Control Table	1466
	Table 97: Monitored Object Instances	1467
	Table 98: Standard MIBs Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis	1468
	Table 99: Juniper Networks Enterprise-Specific MIBs Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis	1474
	Table 100: Standard MIBs Supported on QFabric Systems	1478
	Table 101: Juniper Networks Enterprise-Specific MIBs Supported on QFabric Systems	1482
	Table 102: Standard SNMP Version 1 Traps Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis	1485
	Table 103: Enterprise-Specific SNMPv1 Traps Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis	1487
	Table 104: Standard SNMPv2 Traps Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis	1489
	Table 105: Enterprise-Specific SNMPv2 Traps Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis	1491
	Table 106: Standard SNMPv2 Traps Supported on QFabric Systems	1494
	Table 107: Enterprise-Specific SNMPv2 Traps Supported on QFabric Systems	1495
	Table 108: SNMP Tracing Flags	1520
Chapter 49	Configuring System Logging	1527
	Table 109: Minimum Configuration Statements for System Logging	1530
	Table 110: Fields in Structured-Data Messages	1539
	Table 111: Facility and Severity Codes in the priority-code Field	1541
	Table 112: Fields in Standard-Format Messages	1542
	Table 113: Junos OS System Logging Facilities	1545
	Table 114: System Log Message Severity Levels	1546
	Table 115: Junos OS System Logging Facilities	1546
	Table 116: System Log Message Severity Levels	1547
	Table 117: Default Facilities for Messages Directed to a Remote Destination	1548

	Table 118: Facilities for the facility-override Statement	1548
	Table 119: Regular Expression Operators for the match Statement	1552
Part 22	Operational Commands	
Chapter 55	Network Management	1711
	Table 120: Match Conditions for the monitor traffic Command	1714
	Table 121: Logical Operators for the monitor traffic Command	1716
	Table 122: Arithmetic and Relational Operators for the monitor traffic Command	1717
Chapter 56	sFlow Technology	1729
	Table 123: show sflow Output Fields	1731
	Table 124: show sflow collector Output Fields	1733
	Table 125: show sflow interface Output Fields	1734
Chapter 57	SNMP	1737
	Table 126: show snmp health-monitor Output Fields	1749
	Table 127: show snmp inform-statistics Output Fields	1754
	Table 128: show snmp mib Output Fields	1757
	Table 129: show snmp rmon Output Fields	1759
	Table 130: show snmp statistics Output Fields	1765
	Table 131: show snmp statistics subagents Output Fields	1768
	Table 132: show snmp v3 Output Fields	1773
Chapter 58	System Logging	1775
	Table 133: request app-engine cleanup Output Fields	1778
	Table 134: show app-engine crash Output Fields	1780
	Table 135: show app-engine logs Output Fields	1782
Part 23	User and Access Management Overview	
Chapter 59	Understanding the Software	1793
	Table 136: Junos OS Processes	1794
Chapter 60	Understanding Access and Authentication Methods	1799
	Table 137: Login Class Permission Flags	1800
Part 24	Configuring Access	
Chapter 61	Configuring and Managing Root Users	1809
	Table 138: Troubleshooting Resources on the QFX and OCX Series	1814
	Table 139: Troubleshooting on the QFX Series	1816
Chapter 62	Configuring and Managing User Accounts	1821
	Table 140: Predefined System Login Classes	1823
	Table 141: Special Requirements for Plain-Text Passwords	1824
	Table 142: Configuration Mode Hierarchies—Common Regular Expression Operators	1827
	Table 143: Common Regular Expression Operators to Allow or Deny Operational Mode Commands	1828
	Table 144: Troubleshooting Resources on the QFX and OCX Series	1836

	Table 145: Troubleshooting on the QFX Series	1838
Part 25	Configuring Authentication	
Chapter 63	Configuring and Managing Local Password Authentication	1845
	Table 146: Predefined System Login Classes	1848
	Table 147: Configuration Mode Hierarchies—Common Regular Expression Operators	1849
	Table 148: Common Regular Expression Operators to Allow or Deny Operational Mode Commands	1850
	Table 149: Special Requirements for Plain-Text Passwords	1850
Chapter 64	Configuring and Managing TACACS+ Authentication	1859
	Table 150: Order of Authentication Attempts	1861
	Table 151: Juniper Networks Vendor-Specific TACACS+ Attributes	1864
Chapter 65	Configuring and Managing RADIUS Authentication	1869
	Table 152: Order of Authentication Attempts	1871
Chapter 66	Configuring and Managing RADIUS Accounting	1883
	Table 153: Order of Authentication Attempts	1886
	Table 154: Juniper Networks Vendor-Specific RADIUS Attributes	1889
Chapter 68	Configuring and Managing VSAs for RADIUS and TACACS+	1901
	Table 155: Match Conditions	1902
	Table 156: Actions for VSAs	1903
	Table 157: Juniper Networks Vendor-Specific TACACS+ Attributes	1904
Part 26	Configuration Statements and Operational Commands	
Chapter 70	Operational Commands	1993
	Table 158: show ethernet-switching interfaces Output Fields	1998
	Table 159: show lldp Output Fields	2002
	Table 160: show lldp local-information Output Fields	2007
	Table 161: show lldp neighbors Output Fields	2009
	Table 162: show lldp statistics Output Fields	2014
	Table 163: show route instance Output Fields	2016
	Table 164: show snmp statistics Output Fields	2021
	Table 165: show snmp statistics subagents Output Fields	2024
Part 27	ELS CLI	
Chapter 71	Enhanced Layer 2 Software (ELS) CLI	2035
	Table 166: ELS Support	2035
	Table 167: Renaming the ethernet-switching-options hierarchy	2041
	Table 168: RTG Statements	2042
	Table 169: Deleted Statements	2042
	Table 170: Port Mirroring hierarchy	2043
	Table 171: Layer 2 Control Protocol	2043
	Table 172: dot1q-tunneling	2043
	Table 173: mac-table-aging-time statement	2044

	Table 174: Nonstop Bridging statement	2044
	Table 175: Port Security statements	2045
	Table 176: DHCP Snooping Statements	2045
	Table 177: VLAN hierarchy	2046
	Table 178: Statements Moved to a Different Hierarchy	2048
	Table 179: Changes to the Storm Control Profile hierarchy level	2049
	Table 180: Changes to the Interfaces hierarchy	2050
	Table 181: IGMP Snooping hierarchy	2051
Part 28	Ethernet Ring Protection	
Chapter 72	Understanding Ethernet Ring Protection	2055
	Table 182: Components to Configure for This Example	2068
Part 29	Bridging and VLANs	
Chapter 73	Using Bridging and VLANs	2085
	Table 183: Components of the Basic Bridging Configuration Topology	2097
	Table 184: Components of the Multiple VLAN Topology	2115
	Table 185: Sample IRB Values	2125
	Table 186: Number of Supported IRBs/RVIs by Platform	2125
	Table 187: Components of the Multiple VLAN Topology	2127
Part 31	Spanning Trees	
Chapter 75	Using Spanning Trees	2147
	Table 188: Components of the Topology for Configuring MSTP on EX Series Switches	2151
	Table 189: Components of the Topology for Configuring RSTP	2175
	Table 190: Interfaces of the Topology for Configuring VSTP	2193
	Table 191: Components of the Topology for Configuring BPDU Protection on EX Series Switches	2201
	Table 192: Topology for Configuring Loop Protection on the QFX Series	2208
	Table 193: Topology for Configuring Root Protection on the QFX Series	2214
Part 34	Reflective Relay	
Chapter 78	Using Reflective Relay	2249
	Table 194: Components of the Topology for Configuring Reflective Relay	2253
Part 35	Configuration Statements and Operational Commands	
Chapter 80	VLAN Configuration Statements	2271
	Table 195: Unsupported [edit vlans] Configuration Statements on EX Series Switches	2274
Chapter 85	Bridging and VLANs Monitoring Commands	2379
	Table 196: show ethernet-switching interfaces Output Fields	2382
	Table 197: show ethernet-switching table Output Fields	2386
	Table 198: show vlans Output Fields	2396

Chapter 86	MAC Address Operational Commands	2405
	Table 199: show ethernet-switching mac-learning-log Output Fields	2406
	Table 200: show ethernet-switching mac-notification Output Fields	2408
	Table 201: show ethernet-switching statistics aging Output Fields	2410
	Table 202: show ethernet-switching statistics mac-learning Output Fields . . .	2413
Chapter 87	Spanning Tree Monitoring Commands	2417
	Table 203: show spanning-tree bridge Output Fields	2420
	Table 204: show spanning-tree Interface Output Fields	2425
	Table 205: show spanning-tree mstp configuration Output Fields	2431
	Table 206: show spanning-tree statistics Output Fields	2433
Part 36	Overview	
Chapter 88	OVSDB and VXLAN Overview	2439
	Table 207: NSX for Multi-Hypervisor Components and Related Products . . .	2439
	Table 208: OVSDB Support on Juniper Networks Devices	2446
	Table 209: OVSDB Schema Tables	2448
Part 37	Configuring OVSDB and VXLAN	
Chapter 89	Configuring OVSDB-Managed VXLANs with an SDN Controller	2455
	Table 210: OVSDB and VXLAN Configuration Workflow for QFX Series Switches	2456
	Table 211: OVSDB and VXLAN Configuration Workflow for MX Series Routers and EX9200 Switches	2457
	Table 212: Workflow of Tasks and Events for the Dynamic Configuration of OVSDB-Managed VXLANs in an NSX Environment	2463
	Table 213: Workflow of Tasks and Events for the Dynamic Configuration of OVSDB-Managed VXLANs in a Contrail Environment	2465
	Table 214: Key Configurations to Create a Gateway in NSX Manager	2470
	Table 215: Key Configurations to Create a Gateway Service in NSX Manager . .	2471
	Table 216: Key Configurations to Create a Logical Switch Port in NSX Manager	2471
	Table 217: NSX Manager and Junos OS Entities That Must Be Configured . . .	2474
	Table 218: Components of the Topology for Setting Up a VXLAN Layer 2 Gateway and OVSDB Connections	2475
	Table 219: NSX Manager and Junos OS Entities That Must Be Configured . . .	2482
	Table 220: NSX Manager Configurations and Dynamic Configurations by Juniper Networks Switch	2483
	Table 221: Components for Two VXLAN Topologies Configured on a Juniper Networks Switch that Functions as a Hardware VTEP	2483
Part 39	Configuration Statements and Operational Commands	
Chapter 94	OVSDB Operational Commands	2531
	Table 222: show ovssdb commit failures Output Fields	2535
	Table 223: show ovssdb controller Output Fields	2536
	Table 224: show ovssdb interface Output Fields	2538
	Table 225: show ovssdb logical-switch Output Fields	2541

	Table 226: show ovssdb mac Output Fields	2544
	Table 227: show ovssdb statistics interface Output Fields	2547
	Table 228: show ovssdb virtual-tunnel-end-point Output Fields	2549
	Table 229: show vpls mac-table Output fields	2552
Chapter 95	VXLAN Operational Commands	2557
	Table 230: show bridge mac-table Output Fields	2559
	Table 231: show vpls mac-table Output fields	2564
Part 43	Nonstop Active Routing	
Chapter 99	Configuring Nonstop Active Routing	2597
	Table 232: Nonstop Active Routing Platform Support	2600
	Table 233: Nonstop Active Routing Protocol and Feature Support	2602
Part 44	Graceful Routing Engine Switchover	
Chapter 100	Configuring Graceful Routing Engine Switchover	2617
	Table 234: Effects of a Routing Engine Switchover	2622
	Table 235: Graceful Routing Engine Switchover Feature Support	2624
Part 45	Virtual Router Redundancy Protocol	
Chapter 101	Configuring Virtual Router Redundancy Protocol	2633
	Table 236: Settings for VRRP Load-Sharing Example	2637
	Table 237: Interface State and Priority Cost Usage	2648
Part 46	Configuration Statements and Operational Commands	
Chapter 107	Operational Mode Commands (Graceful Restart)	2707
	Table 238: show bgp neighbor Output Fields	2711
	Table 239: show ospf overview Output Fields	2732
Chapter 108	Operational Mode Commands (Graceful Switchover)	2737
	Table 240: show system switchover Output Fields	2740
	Table 241: show task replication Output Fields	2744
Chapter 109	Operational Mode Command (Nonstop Routing)	2747
	Table 242: show task replication Output Fields	2748
Chapter 110	Operational Mode Commands (VRRP)	2751
	Table 243: show vrrp Output Fields	2752
Part 48	Generic Routing Encapsulation (GRE)	
Chapter 112	Understanding GRE	2777
	Table 244: Firewall Filter Application Points for Tunneled Packets	2779
	Table 245: Features Not Supported with GRE	2780
Part 49	Interfaces	
Chapter 113	Understanding Interfaces	2785

	Table 246: Network Interface Types and Purposes	2785
	Table 247: Special Interface Types and Purposes	2787
	Table 248: Valid Port Ranges on QFX3500 Switches Running QFabric Software Package	2800
	Table 249: Valid Port Ranges on QFX3500 Switches Running Enhanced Layer 2 Software	2804
	Table 250: Valid Port Ranges on QFX3600 Switches Running QFabric Software Package	2807
	Table 251: Valid Port Ranges on QFX3600 Switches Running Enhanced Layer 2 Software	2810
	Table 252: Valid Port Ranges on QFX3600 Node Devices Running QFabric Software Package	2813
	Table 253: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running Enhanced Layer 2 Software	2815
	Table 254: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running QFabric Software Package	2819
	Table 255: System Modes Supported on QFX5100 Switches Running Enhanced Layer 2 Software	2822
	Table 256: Example LPM Table Combinations Using l2-profile-one With Junos OS 13.2X51-D10	2833
	Table 257: Example LPM Table Combinations Using lpm-profile With Junos OS 13.2X51-D10	2834
	Table 258: QFX10002-36Q Switch and QFX10000-36Q Line Card Port Mappings	2837
	Table 259: QFX10002-72Q Switch Port Mappings	2839
Part 50	IP Directed Broadcast	
Chapter 114	Understanding IP Directed Broadcast	2849
	Table 260: Components of the IP Directed Broadcast Topology	2853
Part 52	Link Aggregation Groups (LAGs) and Link Aggregation Control Protocol (LACP)	
Chapter 116	Understanding LAGs and LACP	2865
	Table 261: Components of the Topology for Configuring a LAG Between a QFX3500 Switch and Aggregation Switch	2877
Part 55	Redundant Trunk Groups	
Chapter 119	Understanding Redundant Trunk Groups	2901
	Table 262: Components of the Redundant Trunk Link Topology	2906
Part 56	Resilient Hashing	
Chapter 120	Understanding Resilient Hashing	2911
	Table 263: IPv4 and IPv6 Hashing Fields	2913
	Table 264: MPLS Hashing Fields	2914
	Table 265: MAC-in-MAC Hashing Fields	2915
	Table 266: Layer 2 Header Hashing Fields	2916

	Table 267: Destination Path Results for Static Hashing and for Resilient Hashing When Members Are Added to or Deleted from Trunk Groups	2918
	Table 268: IPv4 and IPv6 Hashing Fields for the QFX10002 Switch and QFX 10008 Switch	2920
	Table 269: MPLS Hashing Fields for the QFX10002 Switch and QFX 10008 Switch	2921
	Table 270: Layer 2 Header Hashing Fields for the QFX10002 Switch	2922
Part 57	Uplink Failure Detection	
Chapter 121	Understanding Uplink Failure Detection	2929
	Table 271: Settings for Uplink Failure Protection Example	2933
Part 58	Configuration Statements and Operational Commands	
Chapter 132	Ethernet OAM Link Fault Management Operational Command	3083
	Table 272: show oam ethernet link-fault-management Output Fields	3084
Chapter 133	Interfaces Operational Commands	3089
	Table 273: Output Control Keys for the monitor interface interface-name Command	3090
	Table 274: Output Control Keys for the monitor interface traffic Command . . .	3091
	Table 275: monitor interface Output Fields	3092
	Table 276: show interfaces diagnostics optics Output Fields	3102
	Table 277: show interfaces ge Output Fields	3117
	Table 278: GRE show interfaces Output Fields	3132
	Table 279: show interfaces irb Output Fields	3140
	Table 280: Layer 2 Overhead and Transmitted Packets or Byte Counts	3147
	Table 281: show interfaces queue Output Fields	3150
	Table 282: Byte Count by Interface Hardware	3153
	Table 283: show interfaces xe Output Fields	3186
Chapter 134	LAGs and LACP Operational Commands	3203
	Table 284: show lacp interfaces Output Fields	3205
	Table 285: show lacp statistics interfaces Output Fields	3209
Chapter 135	Redundant Trunk Group Operational Command	3211
	Table 286: show redundant-trunk-group Output Fields	3212
Chapter 136	Resilient Hashing Operational Command	3215
	Table 287: show forwarding-options enhanced-hash-key Output Fields	3216
Chapter 137	Uplink Failure Detection Operational Command	3221
	Table 288: show uplink-failure-detection Output Fields	3222
Part 59	Configuring Routing Options	
	Table 289: Filtering Route Messages	?
	Table 290: Summary of Key Routing Information Output Fields	?
Part 60	Configuration Statements and Operational Commands	
Chapter 141	Configuration Statements (Routing Options)	3267

	Table 291: Unified Forwarding Table Profiles	3307
Chapter 142	Operational Commands (Routing Options)	3389
	Table 292: show as-path Output Fields	3393
	Table 293: show as-path domain Output Fields	3396
	Table 294: show as-path summary Output Fields	3399
	Table 295: show ipv6 neighbors Output Fields	3401
	Table 296: show ipv6 router-advertisement Output Fields	3403
	Table 297: show route Output Fields	3407
	Table 298: show route damping Output Fields	3431
	Table 299: show route detail Output Fields	3435
	Table 300: Next-Hop Types Output Field Values	3440
	Table 301: State Output Field Values	3442
	Table 302: Communities Output Field Values	3444
	Table 303: show route export Output Fields	3455
	Table 304: show route extensive Output Fields	3458
	Table 305: show route flow validation Output Fields	3475
	Table 306: show route forwarding-table Output Fields	3480
	Table 307: show route instance Output Fields	3497
	Table 308: show route martians Output Fields	3508
	Table 309: show route receive-protocol Output Fields	3534
	Table 310: show route resolution Output Fields	3543
	Table 311: show route summary Output Fields	3554
	Table 312: show route table Output Fields	3559
	Table 313: Next-hop Types Output Field Values	3564
	Table 314: State Output Field Values	3566
	Table 315: Communities Output Field Values	3568
	Table 316: show route terse Output Fields	3586
Part 62	Configuring BGP	
Chapter 145	BGP Path Attribute Configuration	3651
	Table 317: MED Options for Routing Table Path Selection	3667
Chapter 146	BGP Policy Configuration	3765
	Table 318: Default Route Preference Values	3786
Chapter 151	BGP Flap Configuration	3919
	Table 319: Damping Parameters	3927
Part 64	Configuration Statement and Operational Commands	
Chapter 155	Operational Commands	4091
	Table 320: show bgp bmp Output Fields	4098
	Table 321: show bgp group Output Fields	4101
	Table 322: show bgp neighbor Output Fields	4108
	Table 323: show bgp summary Output Fields	4125
	Table 324: show policy damping Output Fields	4130
	Table 325: show route damping Output Fields	4133
	Table 326: show route detail Output Fields	4137
	Table 327: Next-Hop Types Output Field Values	4142

	Table 328: State Output Field Values	4144
	Table 329: Communities Output Field Values	4146
Part 65	Configuring IS-IS Features	
Chapter 157	Configuring IS-IS	4165
	Table 330: IPv4 Statements	4202
	Table 331: IPv6 Statements	4202
Chapter 160	Configuration Statements	4261
	Table 332: Default Metric Values for Routes Exported into IS-IS	4298
Chapter 161	Operational Commands	4329
	Table 333: show isis adjacency Output Fields	4339
	Table 334: show isis authentication Output Fields	4342
	Table 335: show isis backup coverage Output Fields	4344
	Table 336: show isis backup label-switched-path Output Fields	4346
	Table 337: show isis backup spf results Output Fields	4349
	Table 338: show isis database Output Fields	4353
	Table 339: show isis hostname Output Fields	4359
	Table 340: show isis interface Output Fields	4362
	Table 341: show isis overview Output Fields	4365
	Table 342: show isis route Output Fields	4370
	Table 343: show isis spf Output Fields	4373
	Table 344: show isis statistics Output Fields	4379
Part 66	Overview	
Chapter 162	OSPF Overview	4385
	Table 345: Default Route Preference Values for OSPF	4388
Part 69	Configuration Statements and Operational Commands	
Chapter 175	OSPF Operational Commands	4617
	Table 346: show (ospf ospf3) backup coverage Output Fields	4628
	Table 347: show (ospf ospf3) backup neighbor Output Fields	4631
	Table 348: show ospf context-identifier Output Fields	4634
	Table 349: show ospf database Output Fields	4636
	Table 350: show (ospf ospf3) interface Output Fields	4644
	Table 351: show (ospf ospf3) io-statistics Output Fields	4649
	Table 352: show (ospf ospf3) log Output Fields	4651
	Table 353: show (ospf ospf3) neighbor Output Fields	4655
	Table 354: show ospf overview Output Fields	4661
	Table 355: show (ospf ospf3) route Output Fields	4666
	Table 356: show (ospf ospf3) statistics Output Fields	4671
Part 70	Configuring RIP	
Chapter 178	Example: Configuring Authentication for RIP Routes	4693
	Table 357: Configuring Simple RIP Authentication	4698
	Table 358: Configuring MD5 RIP Authentication	4699

Chapter 188	Operational Commands	4789
	Table 359: show rip general-statistics Output Fields	4792
	Table 360: show rip neighbor Output Fields	4795
	Table 361: show rip statistics Output Fields	4797
Part 71	LDP	
Chapter 189	Using LDP	4803
	Table 362: from Operators That Apply to LDP Received-Label Filtering	4811
	Table 363: to Operators for LDP Outbound-Label Filtering	4813
Chapter 190	Configuration Statements and Monitoring Commands for LDP	4841
	Table 364: show ldp database Output Fields	4891
	Table 365: show ldp fec-filters Output Fields	4899
	Table 366: show ldp interface Output Fields	4900
	Table 367: show ldp neighbor Output Fields	4902
	Table 368: show ldp path Output Fields	4904
	Table 369: show ldp route Output Fields	4906
	Table 370: show ldp session Output Fields	4910
	Table 371: show ldp statistics Output Fields	4916
	Table 372: show ldp traffic-statistics Output Fields	4921
	Table 373: traceroute mpls ldp Output Fields	4925
Part 72	MPLS	
Chapter 191	Using MPLS	4929
	Table 374: QFX10000 Switch MPLS Features with Junos OS Release Support	4935
	Table 375: QFX3500, QFX5100, and QFX5200 MPLS Features with Junos OS Release Support	4937
	Table 376: EX4600 Switch MPLS Features with Junos OS Release Support	4941
	Table 377: Comparison of Interprovider and Carrier-of-Carriers VPNs	4966
	Table 378: Supported Match Conditions for MPLS Firewall Filters	4978
	Table 379: Supported Actions for MPLS Firewall Filters	4978
	Table 380: Local CE Switch in the MPLS-Based Layer 3 VPN Topology	4995
	Table 381: Remote CE Switch in the MPLS-Based Layer 3 VPN Topology	4995
	Table 382: Layer 3 VPN Components of the Local PE Switch	4995
	Table 383: Layer 3 VPN Components of the Remote PE Switch	4996
	Table 384: MPLS Scaling Values	5013
Chapter 192	Configuration Statements and Command Statements for MPLS	5017
	Table 385: Output Control Keys for the monitor label-switched-path Command	5109
	Table 386: monitor label-switched-path Output Fields	5110
	Table 387: show security keychain Output Fields	5124
	Table 388: show link-management Output Fields	5127
	Table 389: show link-management peer Output Fields	5131
	Table 390: show link-management routing Output Fields	5133
	Table 391: show link-management statistics Output Fields	5136
	Table 392: show link-management te-link Output Fields	5138

	Table 393: show mpls call-admission-control Output Fields	5140
	Table 394: show mpls cspf Output Fields	5142
	Table 395: show mpls diffserv-te Output Fields	5144
	Table 396: show route forwarding-table Output Fields	5147
	Table 397: show mpls interface Output Fields	5154
	Table 398: show link-management statistics Output Fields	5156
	Table 399: show link-management te-link Output Fields	5158
	Table 400: show mpls call-admission-control Output Fields	5160
	Table 401: show mpls cspf Output Fields	5162
	Table 402: show mpls diffserv-te Output Fields	5164
	Table 403: show route forwarding-table Output Fields	5167
	Table 404: show mpls interface Output Fields	5174
	Table 405: show mpls lsp Output Fields	5178
	Table 406: show mpls lsp autobandwidth Output Fields	5194
	Table 407: show mpls path Output Fields	5197
	Table 408: show route table Output Fields	5200
	Table 409: Next-hop Types Output Field Values	5205
	Table 410: State Output Field Values	5207
	Table 411: Communities Output Field Values	5209
	Table 412: show route forwarding-table Output Fields	5230
	Table 413: show mpls static-lsp Output Fields	5242
	Table 414: show ted database Output Fields	5244
	Table 415: show ted link Output Fields	5252
	Table 416: show ted protocol Output Fields	5255
Part 73	RSVP	
Chapter 193	Using RSVP	5259
	Table 417: LSP Sequence Order for Bin Packing	5275
	Table 418: LSP Sequence Order for Deadlock	5275
	Table 419: LSP Sequence Order for Predictability	5277
	Table 420: LSP Sequence Order for Predictability	5277
	Table 421: Normalization with Per-LSP Autobandwidth Adjustment Changes	5286
	Table 422: Normalization with Traffic Growth	5288
	Table 423: Applicability of RSVP LSPs Configuration to a Container LSP	5294
Chapter 194	Configuration Statements and Command Statements for RSVP	5301
	Table 424: show rsvp interface Output Fields	5356
	Table 425: show rsvp neighbor Output Fields	5360
	Table 426: show rsvp session Output Fields	5367
	Table 427: show rsvp statistics Output Fields	5375
	Table 428: show rsvp version Output Fields	5379
	Table 429: traceroute mpls rsvp Output Fields	5383
Chapter 195	Overview	5389
	Table 430: Multicast Routing Protocols Compared	5400
Part 74	Managing Group Membership	
Chapter 196	Using IGMP	5405

	Table 431: IGMP Event Messages	5425
Chapter 197	Using IGMP Snooping	5431
	Table 432: Components of the IGMP Snooping Topology	5436
	Table 433: Summary of IGMP Snooping Output Fields	5438
Chapter 198	Using MLD	5441
	Table 434: MLD Event Messages	5462
Part 75	Configuring PIM	
Chapter 203	Using Source-Specific Multicast	5529
	Table 435: ASM and SSM Terminology	5531
Chapter 206	Using Auto-RP	5561
	Table 436: Local RP and Auto-RP Message Types	5562
Chapter 208	Using PIM Filtering	5571
	Table 437: PIM Join Filter Match Conditions	5576
Part 76	Configuring MSDP	
Chapter 210	Using MSDP	5597
	Table 438: Source-Active Message Flooding Explanation	5606
Part 77	Configuration Statements and Operational Commands	
Chapter 217	Operational Commands (IGMP)	5789
	Table 439: show igmp group Output Fields	5795
	Table 440: show igmp group Output Fields	5799
	Table 441: show igmp interface Output Fields	5801
	Table 442: show igmp statistics Output Fields	5805
Chapter 218	Operational Commands (IGMP Snooping)	5813
	Table 443: show igmp-snooping membership Output Fields	5816
	Table 444: show igmp-snooping route Output Fields	5819
	Table 445: show igmp-snooping statistics Output Fields	5821
	Table 446: show igmp-snooping vlans Output Fields	5823
Chapter 219	Operational Commands (PIM)	5825
	Table 447: mtrace Output Fields	5839
	Table 448: mtrace from-source Output Fields	5843
	Table 449: mtrace monitor Output Fields	5845
	Table 450: mtrace to-gateway Output Fields	5848
	Table 451: show multicast flow-map Output Fields	5850
	Table 452: show multicast interface Output Fields	5852
	Table 453: show multicast minfo Output Fields	5854
	Table 454: show multicast next-hops Output Fields	5857
	Table 455: show multicast pim-to-igmp-proxy Output Fields	5859
	Table 456: show multicast pim-to-mln-proxy Output Fields	5861
	Table 457: show multicast route Output Fields	5864
	Table 458: show multicast rpf Output Fields	5870

	Table 459: show multicast scope Output Fields	5873
	Table 460: show multicast sessions Output Fields	5875
	Table 461: show multicast usage Output Fields	5879
	Table 462: show pim bootstrap Output Fields	5881
	Table 463: show pim interfaces Output Fields	5883
	Table 464: show pim join Output Fields	5888
	Table 465: show pim neighbors Output Fields	5908
	Table 466: show pim rps Output Fields	5912
	Table 467: show pim source Output Fields	5919
	Table 468: show pim statistics Output Fields	5922
Chapter 220	Operational Commands (MSDP)	5931
	Table 469: show msdp Output Fields	5934
	Table 470: show msdp source Output Fields	5937
	Table 471: show msdp source-active Output Fields	5939
	Table 472: show msdp statistics Output Fields	5941
Part 78	Firewall Filters	
Chapter 221	Configuring Firewall Filters	5951
	Table 473: Supported Firewall Filter Numbers for Specific Switches	5953
	Table 474: Actions for Firewall Filters	5960
	Table 475: Supported Match Conditions for Firewall Filters on QFX10000 Switches	5961
	Table 476: Actions for Firewall Filters on QFX10000 Switches	5969
	Table 477: Action Modifiers for Firewall Filters on QFX10000 Switches	5969
	Table 478: Supported Firewall Filter Numbers	5974
Part 79	Policers	
Chapter 222	Configuring Policers	5999
	Table 479: Policer Actions	6001
	Table 480: Color-Blind Mode TCM Color-to-PLP Mapping	6006
	Table 481: Color-Aware Mode Single-Rate PLP Mapping	6006
	Table 482: Color-Blind Mode TCM Color-to-PLP Mapping	6008
	Table 483: Color-Aware Mode Two-Rate PLP Mapping	6008
	Table 484: Servers Connected to Switch	6013
	Table 485: Unicast Forwarding Classes	6015
Part 83	Configuration Statements and Operational Commands	
Chapter 232	Configuration Statements (DDoS Protection)	6145
	Table 486: Packet Types Supported by DDoS Protection on QFX Switches . . .	6157
	Table 487: Protocol Groups Supported by DDoS Protection on QFX Switches	6159
Chapter 233	Operational Commands (Firewall Filters)	6165
	Table 488: show firewall Output Fields	6167
	Table 489: show firewall policer Output Fields	6171
	Table 490: show interfaces filters Output Fields	6173
	Table 491: show pfe filter hw summary Output Fields	6175

Chapter 235	Operational Commands (DDos Protection)	6179
	Table 492: Supported Protocol Groups	6186
	Table 493: show ddos-protection protocols Output Fields	6192
	Table 494: show ddos-protection protocols parameters Output Fields	6202
	Table 495: show ddos-protection protocols statistics Output Fields	6209
	Table 496: show ddos-protection statistics Output Fields	6218
	Table 497: show ddos-protection version Output Fields	6220
	Table 498: show ddos-protection protocols violations Output Fields	6221
Part 84	Port Mirroring	
Chapter 236	Configuring Port Mirroring	6227
	Table 499: Port Mirroring Terms and Definitions	6228
Part 86	Configuration Statements and Operational Commands	
Chapter 241	Operational Command (Port Mirroring)	6319
	Table 500: show analyzer Output Fields	6320
Chapter 242	Operational Commands (DHCP Local Server)	6323
	Table 501: show dhcp server binding Output Fields	6342
	Table 502: show dhcp server statistics Output Fields	6348
	Table 503: show dhcpv6 server binding Output Fields	6351
	Table 504: show dhcpv6 server statistics Output Fields	6357
Chapter 243	Operational Commands (DHCP Relay Agent)	6359
	Table 505: clear dhcp relay statistics Output Fields	6363
	Table 506: show dhcp relay binding Output Fields	6371
	Table 507: show dhcp relay statistics Output Fields	6376
	Table 508: show dhcpv6 relay binding Output Fields	6379
	Table 509: show dhcpv6 relay statistics Output Fields	6384
	Table 510: show route extensive Output Fields	6387
Chapter 244	Storage Overview	6417
	Table 511: Fibre Channel Terms	6418
Part 87	Configuring FCoE on a Transit Switch	
Chapter 245	Using FCoE on a Transit Switch	6431
	Table 512: Components of the CoS for FCoE Traffic Across an MC-LAG Configuration Topology	6457
Part 88	Configuring DCBX and PFC	
Chapter 246	Using DCBX and PFC	6489
	Table 513: Asymmetric Ethernet PAUSE Flow Control Configuration	6496
	Table 514: Flow Control State Advertised to the Connected Peer (Autonegotiation)	6497
	Table 515: Asymmetric Ethernet PAUSE Behavior on Local and Peer Interfaces	6498
	Table 516: Default PFC Priority to Queue and Forwarding Class Mapping	6500

	Table 517: Components of the PFC for FCoE Traffic Configuration Topology . .	6506
	Table 518: Summary of Differences Between IEEE DCBX and DCBX Version 1.01	6516
	Table 519: Default IEEE 802.1 Classifiers for Trunk Ports and Tagged-Access Ports (Default Trusted Classifier)	6536
	Table 520: Default IEEE 802.1 Unicast Classifiers for Access Ports (Default Untrusted Classifier)	6537
	Table 521: Components of DCBX Application Protocol Exchange Configuration Topology	6538
Part 89	Configuration Statements and Operational Commands	
Chapter 250	Operational Commands (FCoE on a Transit Switch)	6575
	Table 522: show fip vlan-discovery Output Fields	6577
Chapter 251	Operational Commands (DCBX and PFC)	6579
	Table 523: show dcbx output fields	6580
	Table 524: show dcbx neighbors Output Fields	6581
Part 90	CoS Overview	
Chapter 252	Basic Concepts	6607
	Table 525: QFX10000, QFX5100, QFX5200, and EX4600 Switch CoS Features (As of Software Release 15.1X53-D30)	6611
	Table 526: QFX3500 and QFX3600 Switch, and QFabric System CoS Features (As of Software Release 15.1X53-D30)	6612
	Table 527: Ethernet Interface Support for Classifier and Rewrite Rule Configuration (QFX10000 Switches)	6614
	Table 528: Ethernet Interface Support for Classifier and Rewrite Rule Configuration (QFX5100, QFX5200, EX4600, QFX3500, and QFX3600 Switches, and QFabric Systems)	6614
	Table 529: CoS Operational Comparison Between QFX5100 and QFX5200 . .	6615
	Table 530: Classifier Support by Classifier Type	6618
	Table 531: Rewrite Rule Support by Rewrite Rule Type	6618
	Table 532: Default Forwarding Classes and Queue Mapping	6631
	Table 533: Default IEEE 802.1 Code-Point Aliases	6632
	Table 534: Default DSCP and DCSP IPv6 Code-Point Aliases	6632
	Table 535: Default IEEE 802.1 Classifiers for Ports in Trunk Mode (Trusted Classifier)	6633
	Table 536: Default IEEE 802.1 Classifiers for Ports in Access Mode (Untrusted Classifier)	6634
	Table 537: Default DSCP IP and IPv6 Classifiers	6634
	Table 538: Default EXP Classifiers	6635
	Table 539: Default Drop Profile	6636
	Table 540: Default Schedulers	6636
	Table 541: Default Scheduler Maps	6638
	Table 542: CoS Mappings—Inputs and Outputs	6639
	Table 543: Policer Actions	6641

Part 91**Classifying and Rewriting Traffic****Chapter 253****Using Classifiers, Forwarding Classes, and Rewrite Rules 6649**

Table 544: Classifier Support by Classifier Type	6651
Table 545: Default BA Classification	6652
Table 546: Default IEEE 802.1p Code Point to PFC Priority, Output Queue, and Forwarding Class Mapping	6653
Table 547: ba-classifier Loss Priority Assignments	6658
Table 548: Default IEEE 802.1 Classifiers for Ports in Trunk Mode (Trusted Classifier)	6667
Table 549: Default IEEE 802.1 Classifiers for Ports in Access Mode (Untrusted Classifier)	6667
Table 550: Default DSCP IP and IPv6 Classifiers	6668
Table 551: Default EXP Classifiers	6669
Table 552: Default Scheduler Configuration	6669
Table 553: Supported Classifiers and Rewrite Rules	6672
Table 554: Ethernet Interface Support for Classifier and Rewrite Rule Configuration (QFX5100, QFX5200, EX4600, QFX3500, and QFX3600 Switches, and QFabric Systems)	6676
Table 555: Ethernet Interface Support for Classifier and Rewrite Rule Configuration (QFX10000 Switches)	6676
Table 556: Default IEEE 802.1 Code-Point Aliases	6685
Table 557: Default DSCP and DSCP IPv6 Code-Point Aliases	6686
Table 558: Default Forwarding Classes	6688
Table 559: Forwarding-Class-to-Queue Example Configuration	6694
Table 560: Components of the Forwarding Class Sets Configuration Example	6698
Table 561: Configuring Rewrite Rules	6702

Part 92**Scheduling Traffic****Chapter 254****Using Schedulers 6711**

Table 562: Default IEEE 802.1 Classifiers for Ports in Trunk Mode (Trusted Classifier)	6713
Table 563: Default IEEE 802.1 Classifiers for Ports in Access Mode (Untrusted Classifier)	6713
Table 564: Default DSCP IP and IPv6 Classifiers	6714
Table 565: Default EXP Classifiers	6715
Table 566: Default Scheduler Configuration	6715
Table 567: Output Queue Scheduler Components	6730
Table 568: Related Scheduling Components	6731
Table 569: Default Scheduler Configuration	6732
Table 570: Components of the Port Output Queue Scheduler Configuration Example	6747
Table 571: Components of the Queue Scheduler Priority Configuration Example	6754
Table 572: Priority Group Scheduler Components	6759
Table 573: Other Scheduling Components	6759
Table 574: Hierarchical Scheduling Tiers	6766

	Table 575: Components of the Hierarchical Port Scheduling (ETS) Configuration Topology	6773
	Table 576: Components of the Minimum Guaranteed Output Bandwidth Configuration Example	6800
	Table 577: Components of the Maximum Output Bandwidth Configuration Example	6806
	Table 578: ECN Bit Codes	6821
	Table 579: Traffic Behavior on ECN-Enabled Queues	6824
	Table 580: Components of the ECN Configuration Example	6830
Part 93	Data Center Bridging (PFC, DCBX) and Flow Control	
Chapter 255	Using Data Center Bridging and Flow Control	6837
	Table 581: Hierarchical Scheduling Tiers	6842
	Table 582: Components of the Hierarchical Port Scheduling (ETS) Configuration Topology	6849
	Table 583: Mapping of Default Forwarding Class to Queue, IEEE 802.1p Priority, and Drop Attribute	6875
	Table 584: FCoE and No-Loss Forwarding Class Configuration in Junos OS Release 12.3	6878
	Table 585: Default Output Flow Control Profile	6883
	Table 586: User-Configured Output Flow Control Profile	6884
	Table 587: Results of Lossless Priority Configuration	6888
	Table 588: Asymmetric Ethernet PAUSE Flow Control Configuration	6896
	Table 589: Flow Control State Advertised to the Connected Peer (Autonegotiation)	6897
	Table 590: Asymmetric Ethernet PAUSE Behavior on Local and Peer Interfaces	6898
	Table 591: Default PFC Priority to Queue and Forwarding Class Mapping	6900
	Table 592: Components of the PFC for FCoE Traffic Configuration Topology	6909
	Table 593: Components of the PFC Across Layer 3 Interfaces Topology	6925
	Table 594: Components of the CoS for FCoE Traffic Across an MC-LAG Configuration Topology	6946
	Table 595: Components of the Configuration Topology for FCoE Traffic That Does Not Use Priority 3	6970
	Table 596: Components of the Two Lossless FCoE Priorities on an Interface Configuration Topology	6978
	Table 597: Components of the Two Lossless FCoE Priorities Configuration Topology	6988
	Table 598: Components of the Lossless FCoE and iSCSI Priorities Configuration Topology	7002
	Table 599: Summary of Differences Between IEEE DCBX and DCBX Version 1.01	7018
	Table 600: Default IEEE 802.1 Classifiers for Trunk Ports and Tagged-Access Ports (Default Trusted Classifier)	7038
	Table 601: Default IEEE 802.1 Unicast Classifiers for Access Ports (Default Untrusted Classifier)	7039
	Table 602: Components of DCBX Application Protocol Exchange Configuration Topology	7040

Part 94	Configuration Statements and Operational Commands	
Chapter 259	Configuration Statements (Scheduling)	7091
	Table 603: Default Output Queue Buffer Sizes (QFX10000 Switches)	7095
	Table 604: Default Output Queue Buffer Sizes (QFX5100, EX4600, QFX3500, and QFX3600 Switches, and QFabric Systems)	7096
	Table 605: Default Transmit Rates for QFX5100, EX4600, QFX3500, and QFX3600 Switches, and QFabric Systems	7119
	Table 606: Default Transmit Rates for QFX10000 Switches	7120
Chapter 261	Operational Commands (Basic Concepts)	7161
	Table 607: Summary of Key CoS Interfaces Output Fields	7161
	Table 608: show class-of-service Output Fields	7163
	Table 609: show class-of-service interface Output Fields	7168
	Table 610: show class-of-service shared-buffer Output Fields	7196
	Table 611: show pfe filter hw summary Output Fields	7199
	Table 612: show pfe next-hop Output Fields	7203
	Table 613: show pfe route Output Fields	7208
	Table 614: QFX Series, EX4600 switches, and OCX Series show pfe route Hardware Table Output Fields	7209
	Table 615: show interfaces voq Output Fields	7219
Chapter 262	Operational Commands (Classifiers and Rewrite Rules)	7231
	Table 616: Summary of Key CoS Classifier Output Fields	7232
	Table 617: Summary of Key CoS Forwarding Class Output Fields on Switches that Separate Unicast and Multidestination Traffic	7233
	Table 618: Summary of Key CoS Forwarding Class Output Fields on Switches That Do Not Separate Unicast and Multidestination Traffic	7234
	Table 619: Summary of Key CoS Rewrite Rule Output Fields	7236
	Table 620: Summary of Key CoS Value Alias Output Fields	7237
	Table 621: show class-of-service classifier Output Fields	7238
	Table 622: show class-of-service code-point-aliases Output Fields	7241
	Table 623: show class-of-service forwarding-class Output Fields	7243
	Table 624: show class-of-service forwarding-class-set Output Fields	7245
	Table 625: show class-of-service forwarding-table classifier Output Fields . . .	7251
	Table 626: show class-of-service forwarding-table classifier mapping Output Fields	7253
	Table 627: show class-of-service forwarding-table rewrite-rule Output Fields	7255
	Table 628: show class-of-service forwarding-table rewrite-rule mapping Output Fields	7257
	Table 629: show class-of-service interface Output Fields	7259
	Table 630: show class-of-service multi-destination Output Fields	7287
	Table 631: show class-of-service rewrite-rule Output Fields	7288
Chapter 263	Operational Commands (Scheduling)	7291
	Table 632: Summary of Key CoS Scheduler Maps Output Fields	7291
	Table 633: show class-of-service drop-profile Output Fields	7293
	Table 634: show class-of-service forwarding-table drop-profile Output Fields	7300

	Table 635: show class-of-service forwarding-table scheduler-map Output Fields	7302
	Table 636: show class-of-service interface Output Fields	7305
	Table 637: show class-of-service scheduler-map Output Fields	7333
	Table 638: show class-of-service traffic-control-profile Output Fields	7336
	Table 639: Layer 2 Overhead and Transmitted Packets or Byte Counts	7341
	Table 640: show interfaces queue Output Fields	7344
	Table 641: Byte Count by Interface Hardware	7347
	Table 642: show interfaces voq Output Fields	7380
Chapter 264	Operational Commands (Data Center Bridging and Flow Control)	7393
	Table 643: show class-of-service congestion-notification Output Fields	7394
	Table 644: show dcbx output fields	7397
	Table 645: show dcbx neighbors Output Fields	7398

About the Documentation

- Documentation and Release Notes on page cxxxix
- Using the Examples in This Manual on page cxxxix
- Documentation Conventions on page cxli
- Documentation Feedback on page cxliii
- Requesting Technical Support on page cxliii

Documentation and Release Notes

To obtain the most current version of all Juniper Networks® technical documentation, see the product documentation page on the Juniper Networks website at <http://www.juniper.net/techpubs/>.

If the information in the latest release notes differs from the information in the documentation, follow the product Release Notes.

Juniper Networks Books publishes books by Juniper Networks engineers and subject matter experts. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration. The current list can be viewed at <http://www.juniper.net/books>.

Using the Examples in This Manual

If you want to use the examples in this manual, you can use the **load merge** or the **load merge relative** command. These commands cause the software to merge the incoming configuration into the current candidate configuration. The example does not become active until you commit the candidate configuration.

If the example configuration contains the top level of the hierarchy (or multiple hierarchies), the example is a *full example*. In this case, use the **load merge** command.

If the example configuration does not start at the top level of the hierarchy, the example is a *snippet*. In this case, use the **load merge relative** command. These procedures are described in the following sections.

Merging a Full Example

To merge a full example, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration example into a text file, save the file with a name, and copy the file to a directory on your routing platform.

For example, copy the following configuration to a file and name the file **ex-script.conf**. Copy the **ex-script.conf** file to the **/var/tmp** directory on your routing platform.

```
system {
  scripts {
    commit {
      file ex-script.xml;
    }
  }
}
interfaces {
  fxp0 {
    disable;
    unit 0 {
      family inet {
        address 10.0.0.1/24;
      }
    }
  }
}
```

2. Merge the contents of the file into your routing platform configuration by issuing the **load merge** configuration mode command:

```
[edit]
user@host# load merge /var/tmp/ex-script.conf
load complete
```

Merging a Snippet

To merge a snippet, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration snippet into a text file, save the file with a name, and copy the file to a directory on your routing platform.

For example, copy the following snippet to a file and name the file **ex-script-snippet.conf**. Copy the **ex-script-snippet.conf** file to the **/var/tmp** directory on your routing platform.

```
commit {
  file ex-script-snippet.xml; }
```

2. Move to the hierarchy level that is relevant for this snippet by issuing the following configuration mode command:

```
[edit]
user@host# edit system scripts
[edit system scripts]
```

3. Merge the contents of the file into your routing platform configuration by issuing the **load merge relative** configuration mode command:

```
[edit system scripts]
user@host# load merge relative /var/tmp/ex-script-snippet.conf
load complete
```

For more information about the **load** command, see the *CLI User Guide*.

Documentation Conventions

Table 1 defines notice icons used in this guide.

Table 1: Notice Icons







Icon	Meaning	Description
	Informational note	Indicates important features or instructions.
	Caution	Indicates a situation that might result in loss of data or hardware damage.
	Warning	Alerts you to the risk of personal injury or death.
	Laser warning	Alerts you to the risk of personal injury from a laser.
	Tip	Indicates helpful information.
	Best practice	Alerts you to a recommended use or implementation.

Table 2 defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the configure command: user@host> configure
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> show chassis alarms No alarms currently active

Table 2: Text and Syntax Conventions (*continued*)

Convention	Description	Examples
<i>Italic text like this</i>	<ul style="list-style-type: none">Introduces or emphasizes important new terms.Identifies guide names.Identifies RFC and Internet draft titles.	<ul style="list-style-type: none">A policy <i>term</i> is a named structure that defines match conditions and actions.<i>Junos OS CLI User Guide</i>RFC 1997, <i>BGP Communities Attribute</i>
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# set system domain-name <i>domain-name</i>
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none">To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level.The console port is labeled CONSOLE.
< > (angle brackets)	Encloses optional keywords or variables.	stub <default-metric <i>metric</i> >;
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	broadcast multicast (<i>string1</i> <i>string2</i> <i>string3</i>)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only
[] (square brackets)	Encloses a variable for which you can substitute one or more values.	community name members [<i>community-ids</i>]
Indentation and braces ({ })	Identifies a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } }
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
GUI Conventions		
Bold text like this	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none">In the Logical Interfaces box, select All Interfaces.To cancel the configuration, click Cancel.
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select Protocols>Ospf .

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can provide feedback by using either of the following methods:

- Online feedback rating system—On any page of the Juniper Networks TechLibrary site at <http://www.juniper.net/techpubs/index.html>, simply click the stars to rate the content, and use the pop-up form to provide us with information about your experience. Alternately, you can use the online feedback form at <http://www.juniper.net/techpubs/feedback/>.
- E-mail—Send your comments to techpubs-comments@juniper.net. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or Partner Support Service support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at <http://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <http://www.juniper.net/support/warranty/>.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: <http://www.juniper.net/customers/support/>
- Search for known bugs: <http://www2.juniper.net/kb/>
- Find product documentation: <http://www.juniper.net/techpubs/>
- Find solutions and answer questions using our Knowledge Base: <http://kb.juniper.net/>
- Download the latest versions of software and review release notes: <http://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <http://kb.juniper.net/InfoCenter/>

- Join and participate in the Juniper Networks Community Forum:
<http://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <http://www.juniper.net/cm/>

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: <https://tools.juniper.net/SerialNumberEntitlementSearch/>

Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at <http://www.juniper.net/cm/>.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see <http://www.juniper.net/support/requesting-support.html>.

Junos OS Basics Feature Guide for QFX10000 Switches

PART 1

Autoinstallation

- [Understanding Autoinstallation on page 3](#)

CHAPTER 1

Understanding Autoinstallation

- [Understanding Autoinstallation of Configuration Files on page 3](#)
- [Configuring Autoinstallation of Configuration Files \(CLI Procedure\) on page 5](#)
- [Upgrading Software by Using Automatic Software Download on page 7](#)
- [Verifying Autoinstallation Status on page 8](#)
- [Verifying That Automatic Software Download Is Working Correctly on page 9](#)

Understanding Autoinstallation of Configuration Files

Autoinstallation is the automatic configuration of a device over the network from a preexisting configuration file that you create and store on a configuration server—typically a Trivial File Transfer Protocol (TFTP) server. You can use autoinstallation to configure new devices automatically and to deploy multiple devices from a central location in the network.

You enable autoinstallation so that the switches in your network implement autoinstallation when they are powered on. To configure autoinstallation, you specify a configuration server, an autoinstallation interface, and a protocol for IP address acquisition.



NOTE: The QFX5200 switches only work with HTTP for autoinstallation. TFTP and FTP protocols are not supported.

This topic describes:

- [Typical Uses for Autoinstallation on page 3](#)
- [Autoinstallation Configuration Files and IP Addresses on page 4](#)
- [Typical Autoinstallation Process on a New Switch on page 4](#)

Typical Uses for Autoinstallation

Typical uses for autoinstallation of the software include:

- To deploy and update multiple devices from a central location in the network.
- To update a device—Autoinstallation occurs when a device that has been manually configured for autoinstallation is powered on.

Autoinstallation Configuration Files and IP Addresses

For the autoinstallation process to work, you must store one or more host-specific or default configuration files on a configuration server in the network and have a service available—typically Dynamic Host Configuration Protocol (DHCP)—to assign an IP address to the switch.

You can set up the following configuration files for autoinstallation on the switch:

- **network.conf**—Default configuration file for autoinstallation, in which you specify IP addresses and associated hostnames for devices on the network.
- **switch.conf**—Default configuration file for autoinstallation with a minimum configuration sufficient for you to telnet to the device and configure it manually.
- **hostname.conf**—Host-specific configuration file for autoinstallation on a device that contains all the configuration information necessary for the switch. In the filename, **hostname** is replaced with the hostname assigned to the switch.

If the server with the autoinstallation configuration file is not on the same LAN segment as the new device, or if a specific device is required by the network, you must configure an intermediate device directly attached to the new switch, through which the new switch can send TFTP, Boot Protocol (BOOTP), and Domain Name System (DNS) requests. In this case, you specify the IP address of the intermediate device as the location to receive TFTP requests for autoinstallation.

Typical Autoinstallation Process on a New Switch

When the switch configured for autoinstallation is powered on, it performs the following autoinstallation tasks:

1. The switch sends out DHCP or BOOTP requests on each connected interface simultaneously to obtain an IP address.

If a DHCP server responds to these requests, it provides the switch with some or all of the following information:

- An IP address and subnet mask for the autoinstallation interface.
- The location of the (typically) TFTP server, Hypertext Transfer Protocol (HTTP) server, or FTP server on which the configuration file is stored.
- The name of the configuration file to be requested from the TFTP server.
- The IP address or hostname of the TFTP server.

If the DHCP server provides the server's hostname, a DNS server must be available on the network to resolve the name to an IP address.

- The IP address of an intermediate device if the configuration server is on a different LAN segment from the switch.

2. After the switch acquires an IP address, the autoinstallation process on the switch attempts to download a configuration file in the following ways:
 - a. If the DHCP server specifies the host-specific configuration file **hostname.conf**, the switch uses that filename in the TFTP server request. The autoinstallation process on the new switch makes three unicast TFTP requests for **hostname.conf**. If these attempts fail, the switch broadcasts three requests to any available TFTP server for the file.
 - b. If the switch does not locate a **hostname.conf** file, the autoinstallation process sends three unicast TFTP requests for a **network.conf** file that contains the switch's hostname-to-IP-address mapping information. If these attempts fail, the switch broadcasts three requests to any available TFTP server for the file.
 - c. If the switch fails to find a **network.conf** file that contains a hostname entry for the switch, the autoinstallation process sends out a DNS request and attempts to resolve the switch's IP address to a hostname.
 - d. If the switch determines its hostname, it sends a TFTP request for the **hostname.conf** file.
 - e. If the switch is unable to map its IP address to a hostname, it sends TFTP requests for the default configuration file **switch.conf**. The TFTP request procedure is the same as for the **network.conf** file.
3. After the switch locates a configuration file on a TFTP server, the autoinstallation process downloads the file, installs the file on the switch, and commits the configuration.

**Related
Documentation**

- [Configuring Autoinstallation of Configuration Files \(CLI Procedure\) on page 5](#)
- [Connecting and Configuring an EX Series Switch \(CLI Procedure\)](#)
- [Connecting and Configuring an EX Series Switch \(J-Web Procedure\)](#)
- [Configuration Files Terms](#)

Configuring Autoinstallation of Configuration Files (CLI Procedure)

Autoinstallation is the automatic configuration of a device over the network from a pre-existing configuration file that you create and store on a configuration server—typically a Trivial File Transfer Protocol (TFTP) server. You can use autoinstallation to automatically deploy multiple devices from a central location in the network.

To specify autoinstallation to run when you power on a switch already installed in your network, you can enable it by specifying one or more interfaces, protocols, and configuration servers to be used for autoinstallation.

Before you explicitly enable and configure autoinstallation on the switch, perform these tasks as needed for your network's configuration:

- Have a service available—typically Dynamic Host Configuration Protocol (DHCP)—to assign an IP address to the switch
- Configure a DHCP server on your network to meet your network requirements. You can configure a switch to operate as a DHCP server. For more information, see [“Configuring a DHCP Server on Switches \(CLI Procedure\)” on page 78](#).
- Create one of the following configuration files, and store it on a TFTP server (or HTTP server or FTP server) in the network:
 - A host-specific file with the name **hostname.conf** for each switch undergoing autoinstallation. Replace **hostname** with the name of a switch. The **hostname.conf** file typically contains all the configuration information necessary for the switch with this hostname.
 - A default configuration file named **switch.conf** with the minimum configuration necessary to enable you to telnet into the new switch for further configuration.
- Physically attach the switch to the network using a Gigabit Ethernet port.
- If you configure the DHCP server to provide only the TFTP server hostname, add an IP address-to-hostname mapping entry for the TFTP server to the DNS database file on the Domain Name System (DNS) server in the network.
- If the switch is not on the same network segment as the DHCP server (or other device providing IP address resolution), configure an existing device as an intermediate device to receive TFTP and DNS requests and forward them to the TFTP server and the DNS server. You must configure the LAN or serial interface on the intermediate device with the IP addresses of the hosts providing TFTP and DNS services. Connect this interface to the switch.
- If you are using **hostname.conf** files for autoinstallation, you must also complete the following tasks:
 - Configure the DHCP server to provide a **hostname.conf** filename to each switch. Each switch uses its **hostname.conf** filename to request a configuration file from the TFTP server. Copy the necessary **hostname.conf** configuration files to the TFTP server.
 - Create a default configuration file named **network.conf**, and copy it to the TFTP server. This file contains IP-address-to-hostname mapping entries. If the DHCP server does not send a **hostname.conf** filename to a new switch, the switch uses **network.conf** to resolve its hostname based on its IP address.

Alternatively, you can add the IP-address-to-hostname mapping entry for the switch to a DNS database file.

The switch uses the hostname to request a **hostname.conf** file from the TFTP server.

To configure autoinstallation:

1. Specify the URL address of one or more servers from which to obtain configuration files.

```
[edit system]
user@switch# set autoinstallation configuration-servers tftp://tftpconfig.sp.com
```



NOTE: You can also use an FTP address, for example, `ftp://user:password@sftpconfig.sp.com`.

2. Configure one or more Ethernet interfaces to perform autoinstallation and one or two procurement protocols for each interface. The switch uses the protocols to send a request for an IP address for the interface:

```
[edit system]
user@switch# set autoinstallation interfaces ge-0/0/0 bootp
```

**Related
Documentation**

- [Verifying Autoinstallation Status on page 8](#)
- [Understanding Autoinstallation of Configuration Files on page 3](#)
- [Understanding DHCP Services for Switches on page 73](#)

Upgrading Software by Using Automatic Software Download

The automatic software download feature uses the Dynamic Host Configuration Protocol (DHCP) message exchange process to download and install software packages. You configure the automatic software download feature on switches that act as DHCP clients. You must enable automatic software download on a switch before the software upgrade can occur.

You configure a path to a software package file on the DHCP server. The server communicates the path to the software package file through DHCP server messages.

If you enable automatic software download, the DHCP client switch compares the software package name in the DHCP server message with the name of the software package that booted the switch. If the software packages are different, the DHCP client switch downloads and installs the software package specified in the DHCP server message.

Before you upgrade software by using automatic software download, ensure that you have configured DHCP services for the switch, including configuring a path to a boot server and a boot file.

To configure a path to a boot server and a boot file:

1. Configure the name of the boot server advertised to DHCP clients. The client uses a boot file located on the boot server to complete DHCP setup. This configuration is equivalent to DHCP Option 66:

```
[edit system services dhcp]
user@switch# set boot-server (address | hostname)
```

2. Set the boot file advertised to DHCP clients. After the client receives an IP address and the boot file location from the DHCP server, the client uses the boot image stored in the boot file to complete the DHCP setup. This configuration is equivalent to DHCP Option 67:

```
[edit system services dhcp]
```

```
user@switch# set boot-file filename
```

To enable automatic software download on a switch that acts as a DHCP client:

```
[edit chassis]  
user@switch# set auto-image-upgrade
```

After automatic software download is enabled on your DHCP client switch and after DHCP services are enabled on your network, an automatic software download can occur at any time as part of the DHCP message exchange process.

If an automatic software download occurs, you see the following message on the switch:

```
Auto-image upgrade started  
On successful installation system will reboot automatically
```

The switch reboots automatically to complete the upgrade.

**Related
Documentation**

- [Verifying That Automatic Software Download Is Working Correctly on page 9](#)
- [Understanding Software Installation on EX Series Switches](#)
- [Configuring a DHCP Server on Switches \(CLI Procedure\) on page 78](#)
- [Configuring DHCP Services \(J-Web Procedure\)](#)

Verifying Autoinstallation Status

Purpose Display the status of the autoinstallation feature.

Action From the CLI, enter the **show system autoinstallation status** command.

Sample Output

```
user@switch> show system autoinstallation status  
Autoinstallation status:  
Master state: Active  
Last committed file: None  
Configuration server of last committed file: 10.25.100.1  
Interface:  
  Name: ge-0/0/0  
  State: Configuration Acquisition  
  Acquired:  
    Address: 192.168.124.75  
    Hostname: host-ge-000  
    Hostname source: DNS  
    Configuration filename: switch-ge-000.conf  
    Configuration filename server: 10.25.100.3  
  Address acquisition:  
    Protocol: DHCP Client  
    Acquired address: None  
    Protocol: RARP Client  
    Acquired address: None  
Interface:  
  Name: ge-0/0/1  
  State: None  
  Address acquisition:  
    Protocol: DHCP Client  
    Acquired address: None
```


Protocol: RARP Client
Acquired address: None

Meaning The output shows the settings configured for autoinstallation. Verify that the values displayed are correct for the switch when it is deployed on the network.

Related Documentation • [Configuring Autoinstallation of Configuration Files \(CLI Procedure\) on page 5](#)

Verifying That Automatic Software Download Is Working Correctly

Purpose Verify that the automatic software download feature is working correctly.

Action Use the `show system services dhcp client interface-name` command to verify that the automatic software download feature has been used to install a software package.

```
user@switch> show system services dhcp client ge-0/0/1.0
Logical Interface Name      ge-0/0/1.0
Hardware address           00:0a:12:00:12:12
Client Status              bound
Vendor Identifier          ether
Server Address             10.1.1.1
Address obtained           10.1.1.89
Lease Obtained at          2009-08-20 18:13:04 PST
Lease Expires at           2009-08-22 18:13:04 PST

DHCP Options :
Name: name-server, Value: [ 10.209.194.131, 2.2.2.2, 3.3.3.3 ]
Name: server-identifier, Value: 10.1.1.1
Name: router, Value: [ 10.1.1.80 ]
Name: boot-image,
Value: jinstall-ex-4200-9.6R1.5-domestic-signed.tgz
Name: boot-image-location,
Value: 10.1.1.25:/bootfiles/
```

Meaning The output from this command shows the name and location of the software package under DHCP options when automatic software download was last used to install a software package. The sample output in DHCP options shows that the last DHCP server message to arrive on the DHCP client had a boot server address of 192.168.1.165 and a boot file named jinstall-ex-4200-9.6R1.5-domestic-signed.tgz. If automatic software download was enabled on this client switch during the last DHCP message exchange, these values were used by the switch to upgrade the software.

Related Documentation • [Upgrading Software by Using Automatic Software Download on page 7](#)

PART 2

Basic System Management

- [Understanding Basic System Management on page 13](#)

CHAPTER 2

Understanding Basic System Management

- [Compressing the Current Configuration File on page 14](#)
- [Configuring a DNS Name Server for Resolving a Hostname into Addresses on page 15](#)
- [Configuring Console and Auxiliary Port Properties on page 15](#)
- [Configuring the Hostname of a Router or Switch by Using a Configuration Group on page 16](#)
- [Configuring the Junos OS to Determine Conditions That Trigger Alarms on Different Interface Types on page 18](#)
- [Configuring Junos OS to Disable Protocol Redirect Messages on the Router or Switch on page 18](#)
- [Configuring Junos OS to Disable the Reporting of IP Address and Timestamps in Ping Responses on page 19](#)
- [Configuring Junos OS to Extend the Default Port Address Range on page 19](#)
- [Configuring the Junos OS to Select a Fixed Source Address for Locally Generated TCP/IP Packets on page 20](#)
- [Configuring System Alarms to Appear Automatically Upon Login on page 20](#)
- [Configuring Time-Based User Access on page 20](#)
- [Configuring the Timeout Value for Idle Login Sessions on page 22](#)
- [Including the Year or Millisecond in Timestamps on page 22](#)
- [Mapping the Hostname of the Switch to IP Addresses on page 23](#)
- [Modifying the Default Time Zone for a Router or Switch Running Junos OS on page 24](#)
- [Rebooting and Halting a Device on page 25](#)
- [Reaching a Domain Name System Server on page 26](#)
- [Reverting to the Default Factory Configuration by Using the request system zeroize Command on page 28](#)
- [Saving Core Files Generated by Junos OS Processes on page 29](#)
- [Specifying the Physical Location of the Switch on page 29](#)
- [Specifying Access Privileges for Junos OS Operational Mode Commands on page 30](#)
- [Updating the IANA Time Zone Database on Junos OS Devices on page 32](#)

- [Viewing Core Files from Junos OS Processes on page 34](#)
- [Example: Configuring the Name of the Switch, IP Address, and System ID on page 34](#)

Compressing the Current Configuration File

By default, the current operational configuration file is compressed and is stored in the file **juniper.conf.gz** the **/config** file system, along with the last three committed versions of the configuration. If you have large networks, the current configuration file might exceed the available space in the **/config** file system. Compressing the current configuration file enables the file to fit in the file system, typically reducing the size of the file by 90 percent. You might want to compress your current operation configuration files when they reach 3 megabytes (MB) in size.

When you compress the current configuration file, the names of the configuration files change. To determine the size of the files in the **/config** file system, issue the **file list /config detail** command.



NOTE: We recommend that you compress the configuration files (this is the default) to minimize the amount of disk space that they require.

- If you want to compress the current configuration file, include the **compress-configuration-files** statement at the **[edit system]** hierarchy level:

```
[edit system]
compress-configuration-files;
```
- Commit the current configuration file to include the **compression-configuration-files** statement. Commit the configuration again to compress the current configuration file:

```
[edit system]
user@host# set compress-configuration-files
user@host# commit
commit complete
user@host# commit
commit complete
```
- If you do not want to compress the current operational configuration file, include the **no-compress-configuration-files** statement at the **[edit system]** hierarchy level:

```
[edit system]
no-compression-configuration-files;
```
- Commit the current configuration file to include the **no-compress-configuration-files** statement. Commit the configuration again to uncompress the current configuration file:

```
[edit system]
user@host# commit
commit complete
user@host# commit
commit complete
```

- Related Documentation**
- [Junos OS Commit Model for Router or Switch Configuration on page 1313](#)
 - [compress-configuration-files on page 220](#)

Configuring a DNS Name Server for Resolving a Hostname into Addresses

To have the router or switch resolve hostnames into addresses, you must configure one or more Domain Name System (DNS) name servers by including the **name-server** statement at the **[edit system]** hierarchy level:

```
[edit system]
name-server {
    address;
}
```

The following example shows how to configure two DNS name servers:

```
[edit]
user@switch# set system name-server 192.168.1.253
[edit]
user@switch# set system name-server 192.168.1.254
[edit]
user@switch# show
system {
    name server {
        192.168.1.253;
        192.168.1.254;
    }
}
```

- Related Documentation**
- [name-server on page 231](#)

Configuring Console and Auxiliary Port Properties

The console port and auxiliary port on a switch provide out-of-band remote access to the switch. You can configure the console and auxiliary ports so that an external data terminal may be connected to the switch. The console port is enabled by default. The console port speed is 9600 baud, except on OCX Series devices, on which it is 115200 baud. The auxiliary port is disabled by default.

By default, terminal connections to the console and auxiliary ports are secure. When you configure the console and auxiliary ports as insecure, root logins are not allowed to establish terminal connections, and superusers and anyone with a user identifier (UID) of 0 are not allowed to establish terminal connections in multiuser mode.

To configure the console and auxiliary port properties on the switch:

1. To specify that the console port session should terminate if the connection to the data carrier is lost:

```
[edit system ports]
user@switch# set console log-out-on-disconnect
```

2. To specify the auxiliary port terminal type:

```
[edit system ports]
user@switch# set auxiliary type (ansi | small-xterm | vt100 | xterm)
```

For example, to specify the auxiliary port terminal type of **xterm** with a display of 80 columns by 65 rows:

```
[edit system ports]
user@switch# set auxiliary type xterm
```

3. To check the configuration:

```
[edit system ports]
user@switch# show
console log-out-on-disconnect;
auxiliary type xterm;
```

- Related Documentation
- [auxiliary on page 218](#)
 - [console \(Physical Port\) on page 221](#)
 - [ports on page 234](#)

Configuring the Hostname of a Router or Switch by Using a Configuration Group

The hostname of a device is its identification. A router or switch must have its identity established to be accessible on the network to other devices. That is perhaps the most important reason to have a hostname, but a hostname has other purposes: Junos OS uses the configured hostname as part of the command prompt, to prepend log files and other accounting information, as well as in other places where knowing the device identity is useful. We recommend that the hostname be descriptive and memorable.

You can configure the hostname at the **[edit system]** hierarchy level, a procedure shown in *Example: Configuring the Unique Identity of a Router for Making it Accessible on the Network*. Optionally, instead of configuring the hostname at the **[edit system]** hierarchy level, you can use a configuration group, as shown in this procedure. This is a recommended best practice for configuring the hostname, especially if the device has dual Routing Engines. This procedure uses groups called **re0** and **re1** as an example.



NOTE: If you configure hostnames that are longer than the CLI screen width, regardless of the terminal screen width setting, the commit operation occurs successfully, starting with Junos OS release 13.2R3. Even if the terminal screen width is less than the hostname length, commit is successful. In Junos OS releases earlier than Release 13.2R3, if you configured such hostnames by using the `host-name hostname` statement at the `[edit system]` hierarchy level and the terminal screen width was less than the length of the hostname by using the `set cli screen-width` statement, a foreign file propagation (ffp) failure error message is displayed when you attempt to commit the configuration. In such a case, because of the ffp failure, the commit operation does not complete and you cannot recover the router unless you make the modification in the backend in the `juniper.conf.gz` file and commit the change from the shell prompt.

To set the hostname using a configuration group:

1. Include the **host-name** statement in the configuration at the `[edit groups group-name system]` hierarchy level.

The name value must be less than 256 characters.

```
[edit groups group-name system]
host-name hostname;
```

For example:

```
[edit groups re0 system]
root@# set host-name san-jose-router

[edit groups re1 system]
root@# set host-name san-jose-router1
```

2. If you used one or more configuration groups, apply the configuration groups, substituting the appropriate group names.

For example:

```
[edit]
user@host# set apply-groups [re0 re1]
```

3. Commit the changes.

```
[edit]
root@# commit
```

The hostname subsequently appears in the device CLI prompt.

```
san-jose-router@#
```

Related Documentation

- [Understanding Hostnames](#)

Configuring the Junos OS to Determine Conditions That Trigger Alarms on Different Interface Types

For the different types of PICs, you can configure which conditions trigger alarms and whether they trigger a red or yellow alarm. Red alarm conditions light the **RED ALARM** LED and trigger an audible alarm if one is connected. Yellow alarm conditions light the **YELLOW ALARM** LED and trigger an audible alarm if one is connected.



NOTE: By default, any failure condition on the integrated-services interface (Adaptive Services PIC) triggers a red alarm.

To configure conditions that trigger alarms and that can occur on any interface of the specified type, include the **alarm** statement at the **[edit chassis]** hierarchy level.

```
[edit chassis]
alarm {
  interface-type {
    alarm-name (red | yellow | ignore);
  }
}
```

alarm-name is the name of an alarm.

Related Documentation

- *System-Wide Alarms and Alarms for Each Interface Type*
- *Chassis Conditions That Trigger Alarms*
- *Silencing External Devices Connected to Alarm Relay Contacts*

Configuring Junos OS to Disable Protocol Redirect Messages on the Router or Switch

By default, the router or switch sends protocol redirect messages. To disable the sending of redirect messages by the router or switch, include the **no-redirects** statement at the **[edit system]** hierarchy level:

```
[edit system]
no-redirects;
```

To reenable the sending of redirect messages on the router or switch, delete the **no-redirects** statement from the configuration.

To disable the sending of redirect messages on a per-interface basis, include the **no-redirects** statement at the **[edit interfaces interface-name unit logical-unit-number family family]** hierarchy level.

Related Documentation

- *Configuring Junos OS to Ignore ICMP Source Quench Messages*
- *Configuring Junos OS to Select a Fixed Source Address for Locally Generated TCP/IP Packets*
- *Junos OS Network Interfaces Library for Routing Devices*

Configuring Junos OS to Disable the Reporting of IP Address and Timestamps in Ping Responses

When you issue the **ping** command with the **record-route** option, the Routing Engine displays the path of the ICMP echo request packets and timestamps in the ICMP echo responses by default.

You can configure the Routing Engine to disable the setting of the **record-route** option in the IP header of the ping request packets. Disabling the **record-route** option prevents the Routing Engine from recording and displaying the path of the ICMP echo request packets in the response.

- To configure the Routing Engine to disable the setting of the **record-route** option, include the **no-ping-record-route** statement at the **[edit system]** hierarchy level:

```
[edit system]
no-ping-record-route;
```

- To disable the reporting of timestamps in the ICMP echo responses, include the **no-ping-time-stamp** option at the **[edit system]** hierarchy level:

```
[edit system]
no-ping-time-stamp;
```

By configuring the **no-ping-record-route** and **no-ping-timestamp** options, you can prevent unauthorized persons from discovering information about the provider edge (PE) router or switch and its loopback address.

Related Documentation

- [Configuring Junos OS to Disable the Routing Engine Response to Multicast Ping Packets on page 86](#)

Configuring Junos OS to Extend the Default Port Address Range

By default, the upper range of a port address is 5000. You can increase the range from which the port number can be selected to decrease the probability that someone can determine your port number.

- To configure the Junos OS to extend the default port address range, include the **source-port** statement at the **[edit system internet-options]** hierarchy level:

```
[edit system internet-options]
source-port upper-limit upper-limit;
```

upper-limit *upper-limit* is the upper limit of a source port address and can be a value from 5000 through 65,355.

Related Documentation

- [Configuring Junos OS to Disable TCP RFC 1323 Extensions](#)
- [Configuring Junos OS ARP Learning and Aging Options for Mapping IPv4 Network Addresses to MAC Addresses](#)
- [source-port on page 237](#)

Configuring the Junos OS to Select a Fixed Source Address for Locally Generated TCP/IP Packets

By default, the source address included in locally generated Transmission Control Protocol/IP (TCP/IP) packets, such as FTP traffic, and in User Datagram Protocol (UDP) and IP packets, such as Network Time Protocol (NTP) requests, is chosen as the local address for the interface on which the traffic is transmitted. This means that the local address chosen for packets to a particular destination might change from connection to connection based on the interface that the routing protocol has chosen to reach the destination when the connection is established. If multiple equal-cost next hops are present for a destination, locally generated packets use the **lo0** address as a source.

- To configure the software to select a fixed address to use as the source for locally generated IP packets, include the **default-address-selection** statement at the **[edit system]** hierarchy level:

```
[edit system]  
default-address-selection;
```

If you include the **default-address-selection** statement in the configuration, the Junos OS chooses the system default address as the source for most locally generated IP packets. The default address is usually an address configured on the **lo0** loopback interface. For example, if you specified that SSH and telnet use a particular address, but you also have **default-address selection** configured, the system default address is used.

Related Documentation

- [Configuring Junos OS to Disable Protocol Redirect Messages on the Router or Switch on page 18](#)
- [default-address-selection on page 222](#)

Configuring System Alarms to Appear Automatically Upon Login

You can configure Juniper Networks routers and switches to run the **show system alarms** command whenever a user with the login class **admin** logs in to the router or switch. To do so, include the **login-alarms** statement at the **[edit system login class admin]** hierarchy level.

```
[edit system login class admin]  
login-alarms;
```

For more information on the **show system alarms** command, see the [CLI Explorer](#).

Related Documentation

- [show system alarms on page 1027](#)

Configuring Time-Based User Access

The Junos OS enables you to configure time-based restrictions for user access to log in to a device. This is useful for restricting the time and duration of user logins for all users

belonging to a login class. You can specify the days of the week when users can log in, the access start time, and the access end time.

- To configure user access on specific days of the week, without any restrictions on the duration of login, include the **allowed-days** statement only.

```
[edit system]
login {
  class class-name {
    allowed-days [ days-of-the-week ];
  }
}
```

- To configure user access on all the days of the week for a specific duration, include the **access-start** and **access-end** statements only.

```
[edit system]
login {
  class class-name {
    access-start HH:MM;
    access-end HH:MM;
  }
}
```

- To configure user access on specific days of the week for a specified duration, include the **allowed-days**, **access-start**, and **access-end** statements.

```
[edit system]
login {
  class class-name {
    allowed-days [ days-of-the-week ];
    access-start HH:MM;
    access-end HH:MM;
  }
}
```

Specify the start time and end time in **HH:MM** (24-hour) format, where **HH** represents the hours and **MM** represents the minutes.



NOTE: Access start time and end time that spans across 12:00 AM on a specified day results in the user having access until the next day, even if the access day is not explicitly configured. For instance, the following configuration results in the user having access until 6:00 AM on Tuesday and Thursday, although the **allowed-days** statement specifies access only on Monday and Wednesday:

```
[edit system]
login {
  class operator-night-shift {
    allowed-days [ monday wednesday ];
    access-start 2000;
    access-end 0600;
  }
}
```

- Related Documentation**
- *Examples: Configuring Time-Based User Access*
 - *Defining Junos OS Login Classes*
 - *access-end*
 - *access-start*
 - *allowed-days*
 - [access-end on page 270](#)

Configuring the Timeout Value for Idle Login Sessions

An idle login session is one in which the CLI operational mode prompt is displayed but there is no input from the keyboard. By default, a login session remains established until a user logs out of the router or switch, even if that session is idle. To close idle sessions automatically, you must configure a time limit for each login class. If a session established by a user in that class remains idle for the configured time limit, the session automatically closes.

To define the timeout value for idle login sessions, include the **idle-timeout** statement at the **[edit system login class *class-name*]** hierarchy level:

```
[edit system login class class-name]  
idle-timeout minutes;
```

Specify the number of minutes that a session can be idle before it is automatically closed.

If you have configured a timeout value, the CLI displays messages similar to the following when timing out an idle user. It starts displaying these messages 5 minutes before timing out the user.

```
user@host# Session will be closed in 5 minutes if there is no activity.  
Warning: session will be closed in 1 minute if there is no activity  
Warning: session will be closed in 10 seconds if there is no activity  
Idle timeout exceeded: closing session
```

If you configure a timeout value, the session closes after the specified time has elapsed, unless the user is running telnet or monitoring interfaces using the **monitor interface** or **monitor traffic** command.

- Related Documentation**
- *Defining Junos OS Login Classes*
 - *idle-timeout (System-Login)*

Including the Year or Millisecond in Timestamps

By default, the timestamp recorded in a standard-format system log message specifies the month, date, hour, minute, and second when the message was logged, as in the following example:

```
Aug 21 12:36:30
```

To include the year, the millisecond, or both in the timestamp, include the **time-format** statement at the **[edit system syslog]** hierarchy level:

```
[edit system syslog]
time-format (year | millisecond | year millisecond);
```

However, the timestamp for traceoption messages is specified in milliseconds by default, and is independent of the **[edit system syslog time-format]** statement.

The modified timestamp is used in messages directed to each destination configured by a **file**, **console**, or **user** statement at the **[edit system syslog]** hierarchy level, but not to destinations configured by a **host** statement.

The following example illustrates the format for a timestamp that includes both the millisecond (401) and the year (2006):

```
Aug 21 12:36:30.401 2006
```



NOTE: Messages logged in structured-data format include the year and millisecond by default. If you include the structured-data statement at the **[edit system syslog file filename]** hierarchy level along with the time-format statement, the time-format statement is ignored and messages are logged in structured-data format.

For information about the structured-data statement, see *Logging Messages in Structured-Data Format*. For information about the contents of a structured-data message, see the *Junos OS System Log Reference for Security Devices*.

- Related Documentation**
- *Single-Chassis System Logging Configuration Overview*
 - *Examples: Configuring System Logging*

Mapping the Hostname of the Switch to IP Addresses

To map a hostname of a switch to one or more IP addresses, include the **inet** statement at the **[edit system static-host-mapping hostname]** hierarchy level:

```
[edit system]
static-host-mapping {
  hostname {
    inet [ addresses ];
    alias [ aliases ];
  }
}
```

hostname is the name specified by the **host-name** statement at the **[edit system]** hierarchy level.

For each host, you can specify one or more aliases.

- Related Documentation**
- [Reaching a Domain Name System Server on page 26](#)
 - [Example: Configuring the Unique Identity of a Router for Making it Accessible on the Network](#)
 - [static-host-mapping on page 238](#)

Modifying the Default Time Zone for a Router or Switch Running Junos OS

The default local time zone on the router or switch is UTC (Coordinated Universal Time, formerly known as Greenwich Mean Time, or GMT).

- To modify the local time zone, include the **time-zone** statement at the **[edit system]** hierarchy level:

```
[edit system]
time-zone (GMT hour-offset | time-zone);
```

You can use the **GMT hour-offset** option to set the time zone relative to UTC (GMT) time. By default, **hour-offset** is 0. You can configure this to be a value from -14 to +12.

You can also specify the **time-zone** value as a string such as PDT (Pacific Daylight Time) or WET (Western European Time), or specify the continent and major city.



NOTE: Junos OS complies with the POSIX time-zone standard, which is counter-intuitive to the way time zones are generally indicated relative to UTC. A time zone ahead of UTC (east of the Greenwich meridian) is commonly indicated as GMT +*n*; for example, the Central European Time (CET) zone is indicated as GMT +1. However, this is not true for POSIX time zone designations. POSIX indicates CET as GMT-1. If you include the `set system time-zone GMT+1` statement for a router in the CET zone, your router time will be set to one hour behind GMT, or two hours behind the actual CET time. For this reason, you might find it easier to use the POSIX time-zone strings, which you can list by entering `set system time-zone ?`.

For the time zone change to take effect for all processes running on the router or switch, you must reboot the router or switch.

The following example shows how to change the current time zone to **America/New_York**:

```
[edit]
user@host# set system time-zone America/New_York
[edit]
user@host# show
system {
    time-zone America/New_York;
}
```

- Related Documentation**
- [Understanding NTP Time Servers on page 131](#)
 - [Updating the IANA Time Zone Database on Junos OS Devices on page 32](#)

Rebooting and Halting a Device

To reboot the switch, issue the **request system reboot** command.

```
user@switch> request system reboot ?
Possible completions:
<[Enter]>      Execute this command
all-members    Reboot all virtual chassis members
at             Time at which to perform the operation
both-routing-engines Reboot both the Routing Engines
fast-boot      Enable fast reboot
hypervisor     Reboot Junos OS, host OS, and Hypervisor
in            Number of minutes to delay before operation
local         Reboot local virtual chassis member
member        Reboot specific virtual chassis member (0..9)
message       Message to display to all users
other-routing-engine Reboot the other Routing Engine
|            Pipe through a command
{master:0}
```

```
user@switch> request system reboot
Reboot the system ? [yes,no] (no) yes
Rebooting switch
```



NOTE: Not all options shown in the preceding command output are available on all QFX Series, OCX Series, and EX4600 switches. See the documentation for the **request system reboot** command for details about options.



NOTE: When you issue the **request system reboot hypervisor** command on QFX10000 switches, the reboot takes longer than a standard Junos OS reboot.

Similarly, to halt the switch, issue the **request system halt** command.



CAUTION: Before entering this command, you must have access to the switch's console port in order to bring up the Routing Engine.

```
user@switch> request system halt ?
Possible completions:
<[Enter]>      Execute this command
all-members    Halt all virtual chassis members
at             Time at which to perform the operation
backup-routing-engine Halt backup Routing Engine
both-routing-engines Halt both Routing Engines
in            Number of minutes to delay before operation
local         Halt local virtual chassis member
member        Halt specific virtual chassis member (0..9)
message       Message to display to all users
other-routing-engine Halt other Routing Engine
|            Pipe through a command
```



NOTE: When you issue this command on an individual component in a QFabric system, you will receive a warning that says “Hardware-based members will halt, Virtual Junos Routing Engines will reboot.” If you want to halt only one member, use the `member` option. You cannot issue this command from the QFabric CLI.

Issuing the **request system halt** command on the switch halts the Routing Engine. To reboot a Routing Engine that has been halted, you must connect through the console.

**Related
Documentation**

- [clear system reboot on page 337](#)
- [request system reboot on page 396](#)
- [request system halt on page 384](#)
- [request system power-off on page 391](#)
- *Connecting a QFX Series Device to a Management Console*

Reaching a Domain Name System Server

Domain name system (DNS) servers are used for resolving hostnames to IP addresses.

For redundancy, it is a best practice to configure access to multiple DNS servers. You can configure a maximum of three DNS servers. The approach is similar to the way Web browsers resolve the names of a Web site to its network address. Additionally, Junos OS enables you configure one or more domain names, which it uses to resolve hostnames that are not fully qualified (in other words, the domain name is missing). This is convenient because you can use a hostname in configuring and operating Junos OS without the need to reference the full domain name. After adding DNS server addresses and domain names to your Junos OS configuration, you can use DNS resolvable hostnames in your configuration and commands instead of IP addresses.

Optionally, instead of configuring the name server at the **[edit system]** hierarchy level, you can use a configuration group, as shown in this procedure. This is a recommended best practice for configuring the name server. This procedure uses a group called **global** as an example.

Before you begin, configure your DNS servers with the hostname and an IP address for your Junos OS device. It does not matter which IP address you assign as the address of your Junos OS device in the DNS server, as long it is an address that reaches your device. Normally, you would use the management interface IP address, but you can choose the loopback interface IP address, or a network interface IP address, or even configure multiple addresses on the DNS server.

To configure the router or switch to resolve hostnames into addresses:

1. Reference the IP addresses of your DNS servers.

```
[edit groups group-name system]
name-server {
```

```

    address;
}

```

The following example shows how to reference two DNS servers:

```

[edit groups global system]
user@host# set name-server 192.168.1.253
user@host# set name-server 192.168.1.254

```

```

user@host# show
name server {
    192.168.1.253;
    192.168.1.254;
}

```

2. (Optional) Configure the name of the domain in which the device itself is located.

This is a good practice. Junos OS then uses this configured domain name as the default domain name to append to hostnames that are not fully qualified.

```

[edit system]
domain-name domain-name;

```

The following example shows how to configure the domain name:

```

[edit groups global system]
user@host# set domain-name company.net

```

```

user@host# show
domain-name company.net;

```

3. (Optional) Configure a list of domains to be searched.

If your device can reach several different domains, you can configure these as a list of domains to be searched. Junos OS then uses this list to set an order in which it appends domain names when searching for the IP address of a host.

```

[edit groups global system]
domain-search [ domain-list ];

```

The domain list can contain up to six domain names, with a total of up to 256 characters.

The following example shows how to configure two domains to be searched. This example configures Junos OS to search the company.net domain and then the domainone.net domain and then the domainonealternate.com domain when attempting to resolve unqualified hosts.

```

[edit groups global system]
domain-search [ company.net domainone.net domainonealternate.com ]

```

4. If you used a configuration group, apply the configuration group, substituting **global** with the appropriate group name.

```

[edit]
user@host# set apply-groups global

```

5. Commit the configuration.

```
user@host# commit
```

6. Verify the configuration.

If you have configured your DNS server with the hostname and an IP address for your Junos OS device, you can issue the following commands to confirm that DNS is working and reachable. You can either use the configured hostname to confirm resolution to the IP address or use the IP address of your device to confirm resolution to the configured hostname.

```
user@host> show host host-name
user@host> show host host-ip-address
```

For example:

```
user@host> show host san-jose-router1
san-jose-router1.company.net
san-jose-router1.company.net has address 192.168.187.1

user@host> show host 192.168.187.1
1.187.168.192.in-addr.arpa domain name pointer san-jose-router1.company.net.
```

Related Documentation • [DNS Overview](#)

Reverting to the Default Factory Configuration by Using the `request system zeroize` Command

The **`request system zeroize`** command is a standard Junos OS operational mode command that removes all configuration information and resets all key values. The operation unlinks all user-created data files, including customized configuration and log files, from their directories. The switch then reboots and reverts to the factory-default configuration.

To completely erase user-created data so that it is unrecoverable, use the **`request system zeroize media`** command.



CAUTION: Before issuing **`request system zeroize`**, use the **`request system snapshot`** command to back up the files currently used to run the switch to a secondary device.

To revert to the factory-default configuration by using the **`request system zeroize`** command:

1.

```
user@switch> request system zeroize
```

```
warning: System will be rebooted and may not boot without configuration
```

```
Erase all data, including configuration and log files? [yes,no] (yes)
```
2. Type **yes** to remove configuration and log files and revert to the factory default configuration.
3. Complete the initial configuration of the switch.

Related Documentation • [request system zeroize on page 409](#)

Saving Core Files Generated by Junos OS Processes

By default, when an internal Junos OS process generates a core file, the file and associated context information are saved for debugging purposes in a compressed tar file named `/var/tmp/process-name.core.core-number.tgz`. The contextual information includes the configuration and system log message files.

- To disable the saving of core files and associated context information:

```
[edit system]
no-saved-core-context;
```

- To save the core files only:

```
[edit system]
saved-core-files number;
```

Where *number* is the number of core files to save and can be a value from 1 through 10.

- To save the core files along with the contextual information:

```
[edit system]
saved-core-context;
```

Related Documentation

- [Viewing Core Files from Junos OS Processes on page 34](#)

Specifying the Physical Location of the Switch

To specify the physical location of the switch, specify the following options for the **location** statement at the **[edit system]** hierarchy level:

- **altitude *feet***—Number of feet above sea level.
- **building *name***—Name of the building, 1 to 28 characters in length. If the string contains spaces, enclose it in quotation marks (" ").
- **country-code *code***—Two-letter country code.
- **floor *number***—Floor in the building.
- **hcoord *horizontal-coordinate***—Bellcore Horizontal Coordinate.
- **lata *service-area***—Long-distance service area.
- **latitude *degrees***—Latitude in degree format.
- **longitude *degrees***—Longitude in degree format.
- **npa-nxx *number***—First six digits of the phone number (area code and exchange).
- **postal-code *postal-code***—Postal code.
- **rack *number***—Rack number.
- **vcoord *vertical-coordinate***—Bellcore Vertical Coordinate.

The following example shows how to specify the physical location of the switch:

```
[edit system]
location {
  altitude feet;
  building name;
  country-code code;
  floor number;
  hcoord horizontal-coordinate;
  lata service-area;
  latitude degrees;
  longitude degrees;
  npa-nxx number;
  postal-code postal-code;
  rack number;
  vcoord vertical-coordinate;
}
```

Related Documentation • [Example: Configuring the Name of the Switch, IP Address, and System ID on page 34](#)

Specifying Access Privileges for Junos OS Operational Mode Commands

You can specify extended regular expressions by using the **allow-commands** and **deny-commands** statements to define a user's access privileges to individual operational mode commands. Doing so takes precedence over a login class permissions bit set for a user. You can include one **deny-commands** and one **allow-commands** statement in each login class.

To explicitly provide use of an individual operational mode command that would otherwise be denied, include the **allow-commands** statement at the **[edit system login class *class-name*]** hierarchy level:

```
[edit system login class class-name]
allow-commands "regular-expression";
```

To explicitly deny access to an individual operational mode command that would otherwise be supported, include the **deny-commands** statement at the **[edit system login class *class-name*]** hierarchy level:

```
[edit system login class class-name]
deny-commands "regular-expression";
```



NOTE: The regular expression to allow/deny commands for any login class is supported at the commands level but not at the arg level. For example, you can completely block ping, but not ping *arg1*.

If the regular expression contains any spaces, operators, or wildcard characters, enclose the expression in quotation marks. Regular expressions are not case-sensitive.

```
allow-commands "show interfaces";
```



NOTE: Modifiers are not supported within the regular expression string to be matched. If a modifier is used, then nothing is matched.

For example, the `deny command set protocols` does not match anything, whereas `protocols` matches *protocols*.

Explicitly providing access to operational mode commands using the **allow-commands** statement adds to the regular permissions set using the **permissions** statement. Likewise, explicitly denying access to operational mode commands using the **deny-commands** statement removes permissions for the specified commands from the default permissions provided by the **permissions** statement.

For example, if a login class has the permission **view** and the **allow-commands** statement includes the **request system software add** command, the specified login class user can install software, in addition to the permissions specified by the **view** permissions flag. Likewise, if a login class has the permission **all** and the **deny-commands** statement includes the **request system software add** command, the specified login class user can perform all operations allowed by the **all** permissions flag, except installing software using the **request system software add** command.

If you allow and deny the same commands, the **allow-commands** permissions take precedence over the permissions specified by **deny-commands**. For example, if you include **allow-commands "request system software add"** and **deny-commands "request system software add"**, the login class user is allowed to install software using the **request system software add** command.

If you specify a regular expression for **allow-commands** and **deny-commands** with two different variants of a command, the longest match is always executed.

For example, if you specify a regular expression for **allow-commands** with the **commit-synchronize** command and a regular expression for **deny-commands** with the **commit** command, users assigned to such a login class would be able to issue the **commit synchronize** command, but not the **commit** command. This is because **commit-synchronize** is the longest match between **commit** and **commit-synchronize**, and it is specified for **allow-commands**.

Likewise, if you specify a regular expression for **allow-commands** with the **commit** command and a regular expression for **deny-commands** with the **commit-synchronize** command, users assigned to such a login class would be able to issue the **commit** command, but not the **commit-synchronize** command. This is because **commit-synchronize** is the longest match between **commit** and **commit-synchronize**, and it is specified for **deny-commands**.

Anchors are required when specifying complex regular expressions with **allow-commands** or **deny-commands** statements. For example, when specifying multiple commands using the pipe (|) symbol for **allow-commands**, the following syntax is incorrect:

allow-commands = "(monitor.*)"|(ping.*)"|(show.*)"|(exit)". Instead, you must specify the expression using the following syntax: **allow-commands = "^(^monitor) | (^ping) | (^show) | (^exit)"** OR **allow-commands = "^(monitor | ping | show | exit)"**

Related Documentation

- [Example: Configuring User Permissions with Access Privileges for Operational Mode Commands on page 1832](#)
- [Regular Expressions for Allowing and Denying Junos OS Operational Mode Commands on page 1827](#)
- *allow-commands*
- *deny-commands*

Updating the IANA Time Zone Database on Junos OS Devices

Junos OS devices use the tz database, also known as the IANA Time Zone Database to manage time zones. This database is periodically updated by IANA to reflect political and time changes. As such, you may need from time to time to update this file to ensure the Junos devices continue to accurately reflect worldwide time zones and daylight savings time intervals.

To update the IANA Time Zone Database, perform the following steps:

1. [Importing and Installing Time Zone Files on page 32](#)
2. [Configuring a Custom Time Zone on page 33](#)

Importing and Installing Time Zone Files

The IANA Time Zone Database is maintained by the Internet Assigned Numbers Authority (IANA), which is a department of the Internet Corporation for Assigned Names and Numbers (ICANN). You can download the latest IANA Time Zone Database file from the following URL: <http://www.iana.org/time-zones>.

The following steps will guide you through one method of installing the file to your device. However, depending on your network access and other preferences, you may need to modify these steps.

1. Log into the Junos device.
2. If you are in the CLI interface, open the shell interface.

```
device@user# start shell
```
3. Create a **tz** directory in the **/var/tmp** and navigate to that directory.

```
# mkdir /var/tmp/tz  
# cd /var/tmp/tz
```
4. Using FTP, download the time zone files archive.



NOTE: FTP must be enabled on your device before you can use FTP. FTP is enabled by adding the **ftp** statement into the **[edit system services]** hierarchy.

```
# ftp ftp.iana.org/tz  
# bin
```



```
# get tzdata-latest.tar.gz
```



NOTE: If needed, you can edit the above untarred files to create or modify the time zones.

5. Select the names of time zone files to compile and feed them to the following script. For example, to generate **northamerica** and **asia** tz files:

```
# /usr/libexec/ui/compile-tz northamerica asia
```

6. Enable the use of the generated tz files using the CLI:

```
[edit]
# set system use-imported-time-zones
[edit]
# set system time-zone ?
```

This should show the newly generated tz files in **/var/db/zoneinfo/**.

7. Set the time zone and commit the configuration:

```
[edit]
# set system time-zone <your-time-zone>
# commit
```

8. Verify that the time zone change has taken effect:

```
[edit]
# run show system uptime
```

Configuring a Custom Time Zone

To use a custom time zone, follow these steps:

1. Download a time zones archive (from a known or designated source) to the router or switch. Compile the time zone archive using the **zic** time zone compiler, which generates **tz** files.
2. Using the CLI, configure the router or switch to enable the use of the generated tz files as follows:

```
[edit]
user@host# set system use-imported-time-zones
```

3. Display the imported time zones (saved in the directory **/var/db/zoneinfo/**):

```
[edit]
user@host# set system time-zone ?
```

If you do not configure the router to use imported time zones, the Junos OS default time zones are shown (saved in the directory **/usr/share/zoneinfo/**).

Related Documentation

- *Modifying the Default Time Zone for a Router or Switch Running Junos OS*
- *NTP Overview*
- [Understanding NTP Time Servers on page 131](#)

- [Example: Configuring NTP as a Single Time Source for Router and Switch Clock Synchronization on page 142](#)
- *use-imported-time-zones*

Viewing Core Files from Junos OS Processes

When an internal Junos OS process generates a core file, the output found at `/var/crash/` and `/var/tmp/` can now be viewed. This provides a quick method of finding core issues across large networks.

Use the CLI command `show system core-dumps` to view core files.

```
root@host> show system core-dumps
-rw-----  1 root  wheel  268369920 Jun 18 17:59 /var/crash/vmcore.0
-rw-rw----  1 root  field   3371008 Jun 18 17:53 /var/tmp/rpd.core.0
-rw-r--r--  1 root  wheel   27775914 Jun 18 17:59 /var/crash/kernel.0
```

Related Documentation • [Saving Core Files from Junos OS Processes](#)

Example: Configuring the Name of the Switch, IP Address, and System ID

The following example shows how to configure the switch name, map the name to an IP address and alias, and configure a system identifier:

```
[edit]
user@switch# set system host-names switch-sj1
[edit]
user@switch# set system static-host-mapping switch-sj1 inet 192.168.1.77
[edit]
user@switch# set system static-host-mapping switch-sj1 alias sj1
[edit]
user@switch# set system static-host-mapping switch-sj1 sysid 1921.6800.1077
[edit]
user@switch# show
system {
    host-name switch-sj1;
    static-host-mapping {
        switch-sj1 {
            inet 192.168.1.77;
            alias sj1;
            sysid 1921.6800.1077;
        }
    }
}
```

Related Documentation • [Getting Started Guide for Routing Devices](#)

PART 3

Command Line Interface (CLI)

- [Understanding the CLI on page 37](#)

CHAPTER 3

Understanding the CLI

- [CLI User Interface Overview on page 37](#)
- [Configuring Login Tips on page 40](#)
- [Format for Specifying Filenames and URLs in Junos OS CLI Commands on page 40](#)
- [Getting Started with Enhanced Layer 2 Software on page 41](#)
- [Junos OS Operational Mode Commands That Combine Other Commands on page 57](#)
- [Overview of Junos OS CLI Operational Mode Commands on page 58](#)
- [Overview of Navigating the CLI on page 61](#)
- [Understanding the Brief, Detail, Extensive, and Terse Options of Junos OS Operational Commands on page 62](#)
- [Understanding Junos OS CLI Configuration Mode on page 63](#)

CLI User Interface Overview

- [CLI Overview on page 37](#)
- [CLI Key Features on page 38](#)
- [CLI Command Modes on page 38](#)

CLI Overview

The command-line interface (CLI) is the software interface you use to access, monitor, configure, troubleshoot, and manage a device running Junos OS. You can access the CLI either from the console or through a network connection. The CLI is a Juniper Networks-specific command shell that runs on top of a FreeBSD UNIX-based operating system kernel.

The CLI provides a variety of UNIX utilities, such as Emacs-style keyboard sequences, which allows you to perform the following actions:

- Move around on a command line and scroll through recently executed commands.
- Match regular expressions to locate and replace values and identifiers in a configuration.
- Filter command output.
- Log file entries.
- Store and archive device files on a UNIX-based file system.

You can exit the CLI environment and create a UNIX C shell or Bourne shell to navigate the file system, manage processes, and perform other tasks.

CLI Key Features

The CLI commands and statements follow a hierarchical organization and have consistent syntax. The CLI provides the following features for ease of use:

- Consistent command names—Commands that provide the same type of function have the same name, regardless of the portion of the software on which they are operating. For example, all **show** commands display software information and statistics, and all **clear** commands erase various types of system information.
- Lists and short descriptions of available commands—Information about available commands is provided at each level of the CLI command hierarchy. If you type a question mark (?) at any level, you see a list of the available commands along with a short description of each command. This means that if you already are familiar with the Junos OS, you can use many of the CLI commands without referring to the documentation.
- Command completion—Command completion for command names (keywords) and for command options is available at each level of the hierarchy. To complete a command or option that you have partially typed, press Tab or the Spacebar. If the partially typed letters begin a string that uniquely identifies a command, the complete command name appears. Otherwise, a beep indicates that you have entered an ambiguous command, and the possible completions are displayed. Completion also applies to other strings, such as filenames, interface names, usernames, and configuration statements.

CLI Command Modes

The CLI has two modes, operational mode and configuration mode.

- Operational mode—This mode displays the current status of the device. In operational mode, you enter commands to monitor and troubleshoot Junos OS and devices and network connectivity. Operational mode is indicated by the > prompt—for example, **user@switch> clear**
- Configuration mode—A Junos OS device configuration is stored as a hierarchy of statements. In configuration mode, you can define all properties of the Juniper Networks Junos OS, including interfaces, VLANs, Virtual Chassis information, user access, and several system hardware properties. To enter configuration mode, enter the **configure** command. Configuration mode is indicated by the # prompt and includes the current location in the configuration hierarchy—for example:

```
[edit interfaces ge-0/0/12]  
user@switch#
```

In configuration mode, you are actually viewing and changing the candidate configuration file. The candidate configuration allows you to make configuration changes without causing operational changes to the current operating configuration, called the active configuration. When you commit the changes you added to the candidate configuration,

the system updates the active configuration. Candidate configurations enable you to alter your configuration without causing potential damage to your current network operations.

To activate your configuration changes, enter the **commit** command.

When you commit the candidate configuration, you can require an explicit confirmation for the commit to become permanent by using the **commit confirmed** command. This is useful for verifying that a configuration change works correctly and does not prevent management access to the switch. After you issue the **commit confirmed** command, you must issue another **commit** command within the defined period of time (10 minutes by default), or the system reverts to the previous configuration.

You can also activate your configuration changes and exit configuration mode with a single command, **commit and-quit**. This command succeeds only if there are no mistakes or syntax errors in the configuration.

To return to operational mode, go to the top of the configuration hierarchy and then quit—for example:

```
[edit interfaces ge-0/0/12]
user@switch# top
[edit]
user@switch# exit
```

When you monitor and configure a device running Junos OS, you may need to switch between operational mode and configuration mode. When you change to configuration mode, the command prompt also changes. The operational mode prompt is a right angle bracket (>) and the configuration mode prompt is a pound sign (#).

When you log in to the switch and type the **cli** command, you are automatically in operational mode. To switch to configuration mode, type the **configure** command or the **edit** command.

The CLI prompt changes from **user@switch>** to **user@switch#**, and a banner appears to indicate the hierarchy level.

To return to operational mode as well as commit your changes, enter **command and-quit**. To return to operational mode without committing any of your changes, enter **exit**.

To display the output of an operational mode command, such as **show**, while in configuration mode, issue the **run** configuration mode command and then specify the operational mode command.

Related Documentation

- [Configuring Login Tips on page 40](#)
- [Overview of Navigating the CLI on page 61](#)
- *CLI User Guide*
- [Other Tools to Configure and Monitor Devices Running Junos OS on page 168](#)

Configuring Login Tips

The Junos OS CLI provides the option of configuring login tips for the user. By default, the **tip** command is not enabled when a user logs in.

- To enable tips, include the **login-tip** statement at the **[edit system login class *class-name*]** hierarchy level:

```
[edit system login class class-name]  
login-tip;
```

Adding this statement enables the **tip** command for the class specified, provided the user logs in using the CLI.

Related Documentation

- *Defining Junos OS Login Classes*

Format for Specifying Filenames and URLs in Junos OS CLI Commands

In some command-line interface (CLI) commands and configuration statements—including **file copy**, **file archive**, **load**, **save**, **set system login user *username* authentication load-key-file**, and **request system software add**—you can include a filename. On a routing matrix, you can include chassis information (for example, **lcc0**, **lcc0-re0**, or **lcc0-re1**) as part of the filename.

A *routing matrix* is a multichassis architecture composed of either one TX Matrix router and from one to four T640 routers connected to the TX Matrix router, or one TX Matrix Plus router and from one to four T1600 routers connected to the TX Matrix Plus router. From the perspective of the user interface, the routing matrix appears as a single router. On a routing matrix composed of the TX Matrix router and T640 routers, the TX Matrix router controls all the T640 routers. On a routing matrix composed of a TX Matrix Plus router and T1600 routers, the TX Matrix Plus router controls all the T1600 routers.

You can specify a filename or URL in one of the following ways:

- **filename**—File in the user's current directory on the local CompactFlash card (not applicable on the QFX Series). You can use wildcards to specify multiple source files or a single destination file. Wildcards are not supported in Hypertext Transfer Protocol (HTTP) or FTP.



NOTE: Wildcards are supported only by the **file (compare | copy | delete | list | rename | show)** commands. When you issue the **file show** command with a wildcard, it must resolve to one filename.

- **path/filename**—File on the local flash disk.
- **/var/filename** or **/var/path/filename**—File on the local hard disk. You can also specify a file on a local Routing Engine for a specific T640 router or a T1600 router in a routing matrix:


```
user@host> file delete lcc0-re0:/var/tmp/junk
```

- **a:filename** or **a:path/filename**—File on the local removable media. The default path is / (the root-level directory). The removable media can be in MS-DOS or UNIX (UFS) format.
- **hostname:/path/filename**, **hostname:filename**, **hostname:path/filename**, or **"scp://hostname/path/filename"**—File on an scp/ssh client. This form is not available in the worldwide version of the Junos OS. The default path is the user's home directory on the remote system. You can also specify **hostname** as **username@hostname**.
- **ftp://hostname/path/filename**—File on an FTP server. You can also specify **hostname** as **username@hostname** or **username:password@hostname**. The default path is the user's home directory. To specify an absolute path, the path must start with %2F; for example, **ftp://hostname/%2Fpath/filename**. To have the system prompt you for the password, specify **prompt** in place of the password. If a password is required and you do not specify the password or **prompt**, an error message is displayed:

```
user@host> file copy ftp://username@ftp.hostname.net/filename
file copy ftp.hostname.net: Not logged in.
```

```
user@host> file copy ftp://username:prompt@ftp.hostname.net/filename
Password for username@ftp.hostname.net:
```

- **re0:/path/filename** or **re1:/path/filename**—File on a local Routing Engine. You can also specify a file on a local Routing Engine for a specific T640 router or a T1600 router in a routing matrix:

```
user@host> show log lcc0-re1:chassisd
```



NOTE: You cannot specify a URL for a file on a Hypertext Transfer Protocol (HTTP) server, because HTTP URLs are not writable.

Related Documentation

- *Format for Specifying IP Addresses, Network Masks, and Prefixes in Junos OS Configuration Statements*
- *Default Directories for Junos OS File Storage on the Router or Switch*

Getting Started with Enhanced Layer 2 Software

- [Understanding Enhanced Layer 2 Software Support on page 42](#)
- [Using the ELS Translator on page 42](#)
- [Configuring a VLAN on page 43](#)
- [Configuring the Native VLAN Identifier on page 44](#)
- [Configuring Layer 2 Interfaces on page 44](#)
- [Configuring Layer 3 Interfaces on page 45](#)
- [Configuring an IRB Interface on page 45](#)

- [Configuring an Aggregated Ethernet Interface and Configuring LACP on That Interface on page 46](#)
- [Enhanced Layer 2 CLI Configuration Statement and Command Changes on page 47](#)

Understanding Enhanced Layer 2 Software Support

Enhanced Layer 2 Software (ELS) is automatically supported if your device is running a Junos OS release that supports it. You do not need to take any action to enable ELS, and you cannot disable ELS.

Table 1 lists the EX Series and QFX Series switches that support ELS and the initial ELS release they support..

Table 3: ELS Support

Device	Initial ELS Release
EX4300 switches	13.2X50-D10
EX4600 switches	13.2X51-D25
EX9200 switches	12.3R2
QFX3500 switches	13.2X50-D15
QFX3600 switches	13.2X50-D15
QFX5100 switches	13.2X51-D10
QFX10000 switches	15.1X53-D10

ELS is supported on the EX4300, EX4600, and EX9200 switches for all Junos OS releases, starting with the initial releases shown in [Table 3](#).

ELS support was introduced on QFX3500 and QFX3600 switches in Junos OS Release 13.2X50-D15. ELS is supported only on the software package that supports Virtual Chassis (the **jinstall-qfx-3-*** software package) for QFX3500 and QFX3600 switches.

For QFX5100 switches, ELS support was introduced in Junos OS Release 13.2X51-D10 and is supported on the **jinstall-qfx-5-*** software package.



NOTE: ELS is not supported on software packages that can be installed in a QFabric system.

Using the ELS Translator

ELS Translator is a Web-based tool that converts Junos OS Layer 2 configurations to Enhanced Layer 2 Software (ELS) configurations. This conversion tool supports all Juniper Networks EX Series, MX Series, and QFX Series devices with ELS installed. ELS Translator is hosted on the Juniper Networks Customer Support website for EX Series switches, MX

Series routers, and QFX Series switches and is available to registered users, internal users, partners, and premium service contract customers. You need to log in using your Juniper Networks username and password to access the ELS Translator.

[Click here](#) to access the ELS translator.

If you are upgrading from a version of Junos OS that does not support ELS to a version of Junos OS that supports ELS, we recommend that you update your configuration with the ELS Translator using the following procedure:

1. Log in to your device by using the console port.



NOTE: Perform this procedure only from the console port. You will lose connectivity to your device if you perform this procedure from a management port or any other interface.

2. Copy the entire existing configuration to another file. Save the file in a remote location. See [“Saving a Configuration to a File” on page 1331](#).
3. Retain the portion of your existing configuration related to management network connectivity (such as **[edit system]** hierarchy level). Delete all other top-level configuration hierarchy levels (such as the **[edit interfaces]**, **[edit protocols]**, and **[edit vlans]**). Issue the **commit** command to remove the deleted configuration hierarchy levels.
4. Perform the software upgrade. Reboot your device to complete the upgrade. See [“Software Installation Overview” on page 174](#).



NOTE: Ensure that the console port connection is up during the reboot.

5. [Click here](#) to access the ELS Translator in a web browser. Follow the instructions on the page to update your configuration.
6. Return to your console port connection. When the switch has rebooted to complete the software upgrade, copy the configuration from the ELS Translator to your switch. See [“Uploading a Configuration File” on page 1332](#).
7. Commit the new configuration.



NOTE: It is possible that scripts do not translate correctly. Therefore, review translated scripts carefully before loading the converted configuration on your switch or other device.

Configuring a VLAN

You can configure one or more VLANs to perform Layer 2 bridging. The Layer 2 bridging functions include integrated routing and bridging (IRB) for support for Layer 2 bridging and Layer 3 IP routing on the same interface. EX Series and QFX Series switches can

function as Layer 2 switches, each with multiple bridging, or broadcast, domains that participate in the same Layer 2 network. You can also configure Layer 3 routing support for a VLAN.

To configure a VLAN:

1. Create the VLAN by setting a unique VLAN name and configuring the VLAN ID:

```
[edit]
user@host# set vlans vlan-name vlan-id vlan-id-number
```

2. Assign at least one interface to the VLAN:

```
[edit]
user@host# set interface interface-name family ethernet-switching vlan members vlan-name
```

Configuring the Native VLAN Identifier

EX Series and QFX Series switches support receiving and forwarding routed or bridged Ethernet frames with 802.1Q VLAN tags. Typically, trunk ports, which connect switches to each other, accept untagged control packets, but do not accept untagged data packets. You can enable a trunk port to accept untagged data packets by configuring a native VLAN ID on the interface on which you want the untagged data packets to be received.

To configure the native VLAN ID:

1. On the interface on which you want untagged data packets to be received, set the interface mode to **trunk**, which specifies that the interface is in multiple VLANs and can multiplex traffic between different VLANs.

```
[edit interfaces]
user@host# set interface-name unit logical-unit-number family ethernet-switching
interface-mode trunk
```

2. Configure the native VLAN ID:

```
[edit interfaces]
user@host# set interface-name native-vlan-id number
```

3. Assign the interface to the native VLAN ID:

```
[edit interfaces]
user@host# set interface-name unit logical-unit-number family ethernet-switching vlan
members native-vlan-id-number
```

Configuring Layer 2 Interfaces

To ensure that your high-traffic network is tuned for optimal performance, explicitly configure some settings on the switch's network interfaces.

To configure a Gigabit Ethernet interface or a 10-Gigabit Ethernet interface as a **trunk** interface:

```
[edit]
user@host# set interfaces interface-name unit logical-unit-number family ethernet-switching
interface-mode trunk
```

To configure a Gigabit Ethernet interface or a 10-Gigabit Ethernet interface as a **access** interface:

```
[edit]
```

```
user@host# set interfaces interface-name unit logical-unit-number family ethernet-switching
interface-mode access
```

Configuring Layer 3 Interfaces

To configure a Layer 3 interface, you must assign an IP address to the interface. You assign an address to an interface by specifying the address when you configure the protocol family. For the **inet** or **inet6** family, configure the interface IP address.

You can configure interfaces with a 32-bit IP version 4 (IPv4) address and optionally with a destination prefix, sometimes called a subnet mask. An IPv4 address utilizes a 4-octet dotted decimal address syntax (for example, 192.16.1.1). An IPv4 address with destination prefix utilizes a 4-octet dotted decimal address syntax with a destination prefix appended (for example, 192.16.1.1/30).

To specify an IP4 address for the logical unit:

```
[edit]
user@host# set interfaces interface-name unit logical-unit-number family inet address ip-address
```

You represent IP version 6 (IPv6) addresses in hexadecimal notation by using a colon-separated list of 16-bit values. You assign a 128-bit IPv6 address to an interface.

To specify an IP6 address for the logical unit:

```
[edit]
user@host# set interfaces interface-name unit logical-unit-number family inet6 address ip-address
```

Configuring an IRB Interface

Integrated routing and bridging (IRB) provides support for Layer 2 bridging and Layer 3 IP routing on the same interface. IRB enables you to route packets to another routed interface or to another VLAN that has a Layer 3 protocol configured. IRB interfaces enable the device to recognize packets that are being sent to local addresses so that they are bridged (switched) whenever possible and are routed only when necessary. Whenever packets can be switched instead of routed, several layers of processing are eliminated. An interface named **irb** functions as a logical router on which you can configure a Layer 3 logical interface for VLAN. For redundancy, you can combine an IRB interface with implementations of the Virtual Router Redundancy Protocol (VRRP) in both bridging and virtual private LAN service (VPLS) environments.

To configure an IRB interface:

1. Create a Layer 2 VLAN by assigning it a name and a VLAN ID:

```
[edit]
user@host# set vlans vlan-name vlan-id vlan-id
```

2. Create an IRB logical interface:

```
[edit]
user@host# set interface irb unit logical-unit-number family inet address ip-address
```

3. Associate the IRB interface with the VLAN:

```
[edit]
user@host# set vlans vlan-name l3-interface irb.logical-unit-number
```

Configuring an Aggregated Ethernet Interface and Configuring LACP on That Interface

Use the link aggregation feature to aggregate one or more links to form a virtual link or link aggregation group (LAG). The MAC client can treat this virtual link as if it were a single link to increase bandwidth, provide graceful degradation as failure occurs, and increase availability.

To configure an aggregated Ethernet interface:

1. Specify the number of aggregated Ethernet interfaces to be created:

```
[edit chassis]
user@host# set aggregated-devices ethernet device-count number
```

2. Specify the name of the link aggregation group interface:

```
[edit interfaces]
user@host# set interfaces aex
```

3. Specify the minimum number of links for the aggregated Ethernet interface (*aex*)—that is, the defined bundle— to be labeled *up*:

```
[edit interfaces]
user@host# set aex aggregated-ether-options minimum-links number
```

4. Specify the link speed for the aggregated Ethernet bundle:

```
[edit interfaces]
user@host# set aex aggregated-ether-options link-speed link-speed
```

5. Specify the members to be included within the aggregated Ethernet bundle:

```
[edit interfaces]
user@host# set interface-name ether-options 802.3ad aex
user@host# set interface-name ether-options 802.3ad aex
```

6. Specify an interface family for the aggregated Ethernet bundle:

```
[edit interfaces]
user@host# set aex unit 0 family inet address ip-address
```

For aggregated Ethernet interfaces on the device, you can configure the Link Aggregation Control Protocol (LACP). LACP bundles several physical interfaces to form one logical interface. You can configure aggregated Ethernet with or without LACP enabled.

When LACP is enabled, the local and remote sides of the aggregated Ethernet links exchange protocol data units (PDUs), containing information about the state of the link. You can configure Ethernet links to actively transmit PDUs, or you can configure the links to passively transmit them, sending out LACP PDUs only when they receive them from another link. One side of the link must be configured as active for the link to be up.

To configure LACP:

1. Enable one side of the aggregated Ethernet link as active:

```
[edit interfaces]
user@host# set aex aggregated-ether-options lacp active
```

2. Specify the interval at which the interfaces send LACP packets:

```
[edit interfaces]
user@host# set aex aggregated-ether-options lacp periodic interval
```

Enhanced Layer 2 CLI Configuration Statement and Command Changes

The enhanced Layer 2 CLI feature is introduced in Junos OS Release 12.3R2. The enhanced Layer 2 CLI feature changes the CLI for some of the Layer 2 features on EX Series switches. This enhanced CLI will be used to configure Layer 2 features on future EX Series hardware platforms, and also to configure Layer 2 features on other Juniper Networks products.

The following sections provide a list of existing commands that were moved to new hierarchy levels or changed on EX Series switches as part of this CLI enhancement effort. These sections are provided as a high-level reference only. For detailed information about these commands, use the links to the configuration statements provided or see the technical documentation.

- [Changes to the ethernet-switching-options Hierarchy Level on page 47](#)
- [Changes to the Port Mirroring Hierarchy Level on page 49](#)
- [Changes to the Layer 2 Control Protocol Hierarchy Level on page 49](#)
- [Changes to the dot1q-tunneling Statement on page 50](#)
- [Changes to the L2 Learning Protocol on page 50](#)
- [Changes to Nonstop Bridging on page 50](#)
- [Changes to Port Security and DHCP Snooping on page 51](#)
- [Changes to Configuring VLANs on page 52](#)
- [Changes to Storm Control Profiles on page 55](#)
- [Changes to the Interfaces Hierarchy on page 56](#)
- [Changes to IGMP Snooping on page 57](#)

Changes to the ethernet-switching-options Hierarchy Level

This section outlines the changes to the **ethernet-switching-options** hierarchy level.



NOTE: The **ethernet-switching-options** hierarchy level has been renamed as **switch-options**.

Table 4: Renaming the ethernet-switching-options hierarchy

Original Hierarchy	Changed Hierarchy
<pre>ethernet-switching-options { authentication-whitelist { ... } }</pre>	<pre>switch-options { ... authentication-whitelist { ... } }</pre>

Table 4: Renaming the ethernet-switching-options hierarchy (*continued*)

Original Hierarchy	Changed Hierarchy
<pre> ethernet-switching-options { interfaces <i>interface-name</i> { no-mac-learning; ... } } </pre>	<pre> switch-options { interfaces <i>interface-name</i> { no-mac-learning; ... } } </pre>
<pre> ethernet-switching-options { unknown-unicast-forwarding { (...) } } </pre>	<pre> switch-options { unknown-unicast-forwarding { (...) } } </pre>
<pre> ethernet-switching-options { voip { interface (all [<i>interface-name</i> access-ports]) { forwarding-class (assured-forwarding best-effort expedited-forwarding network-control); vlan <i>vlan-name</i>; ... } } } </pre>	<pre> switch-options { voip { interface (all [<i>interface-name</i> access-ports]) { forwarding-class (assured-forwarding best-effort expedited-forwarding network-control); vlan <i>vlan-name</i>; ... } } } </pre>

Table 5: RTG Statements

Original Hierarchy	Changed Hierarchy
<pre> ethernet-switching-options { redundant-trunk-group { group <i>name</i> { description; interface <i>interface-name</i> { primary; } preempt-cutover-timer <i>seconds</i>; ... } } } </pre>	<pre> switch-options { redundant-trunk-group { group <i>name</i> { description; interface <i>interface-name</i> { primary; } preempt-cutover-timer <i>seconds</i>; ... } } } </pre>

Table 6: Deleted Statements

Original Hierarchy	Changed Hierarchy
<pre> ethernet-switching-options { mac-notification { notification-interval <i>seconds</i>; ... } } </pre>	<p>The statements have been removed from the switch-options hierarchy.</p>

Table 6: Deleted Statements (*continued*)

Original Hierarchy	Changed Hierarchy
<pre> ethernet-switching-options { traceoptions { file filename <files number> <no-stamp> <replace> <size size> <world-readable no-world-readable>; flag flag <disable>; ... } } </pre>	<p>The statements have been removed from the switch-options hierarchy.</p>
<pre> ethernet-switching-options { port-error-disable { disable-timeout timeout; ... } } </pre>	<p>NOTE: The port-error-disable statement has been replaced with a new statement.</p> <pre> interfaces interface-name family ethernet-switching { recovery-timeout seconds; } </pre>

Changes to the Port Mirroring Hierarchy Level



NOTE: Statements have moved from the **ethernet-switching-options** hierarchy level to the **forwarding-options** hierarchy level.

Table 7: Port Mirroring hierarchy

Original Hierarchy	Changed Hierarchy
<pre> ethernet-switching-options { analyzer (Port Mirroring) { name { ... } } } </pre>	<pre> forwarding-options { analyzer (Port Mirroring) { name { ... } } } </pre>

Changes to the Layer 2 Control Protocol Hierarchy Level

The Layer 2 control protocol statements have moved from the **ethernet-switching-options** hierarchy to the **protocols** hierarchy.

Table 8: Layer 2 Control Protocol

Original Hierarchy	Changed Hierarchy
<pre> ethernet-switching-options { bpdu-block { ... } } </pre>	<pre> protocols { layer2-control { bpdu-block { ... } } } </pre>

Changes to the dot1q-tunneling Statement

The **dot1q-tunneling** statement has been replaced with a new statement and moved to a different hierarchy level.

Table 9: dot1q-tunneling

Original Hierarchy	Changed Hierarchy
<pre> ethernet-switching-options { dot1q-tunneling { ether-type (0x8100 0x88a8 0x9100); ... } }</pre>	<pre> interfaces interface-name { ether-options { ethernet-switch-profile { tag-protocol-id [tpids]; } } } interfaces interface-name { aggregated-ether-options { ethernet-switch-profile { tag-protocol-id [tpids]; } } }</pre>

Changes to the L2 Learning Protocol

The **mac-table-aging-time** statement has been replaced with a new statement and moved to a different hierarchy level.

Table 10: mac-table-aging-time statement

Original Hierarchy	Changed Hierarchy
<pre> ethernet-switching-options { mac-table-aging-time seconds; ... }</pre>	<pre> protocols { l2-learning { global-mac-table-aging-time seconds; ... } }</pre>

Changes to Nonstop Bridging

The **nonstop-bridging** statement has moved to a different hierarchy level.

Table 11: Nonstop Bridging statement

Original Hierarchy	Changed Hierarchy
<pre> ethernet-switching-options { nonstop-bridging; }</pre>	<pre> protocols { layer2-control { nonstop-bridging { } } }</pre>

Changes to Port Security and DHCP Snooping

Port security and DHCP snooping statements have moved to different hierarchy levels.



NOTE: The statement `examine-dhcp` does not exist in the changed hierarchy. DHCP snooping is now enabled automatically when other DHCP security features are enabled on a VLAN. See *Configuring Port Security Features* for additional information.

Table 12: Port Security statements

Original Hierarchy	Changed Hierarchy
<pre> ethernet-switching-options { secure-access-port { interface (all interface-name) { (dhcp-trusted no-dhcp-trusted); static-ip ip-address { mac mac-address; vlan vlan-name; } } } vlan (all vlan-name) { (arp-inspection no-arp-inspection); dhcp-option82 { disable; circuit-id { prefix hostname; use-interface-description; use-vlan-id; } remote-id { prefix (hostname mac none); use-interface-description; use-string string; } vendor-id [string]; } (examine-dhcp no-examine-dhcp); } (ip-source-guard no-ip-source-guard); } </pre>	<pre> vlans vlan-name forwarding-options{ dhcp-security { arp-inspection; group group-name { interface interface-name { static-ip ip-address { mac mac-address; } } } overrides { no-option82; trusted; } } ip-source-guard; no-dhcp-snooping; option-82 { circuit-id { prefix { host-name; routing-instance-name; } use-interface-description (device logical); use-vlan-id; } } remote-id { host-name; use-interface-description (device logical); use-string string; } vendor-id { use-string string; } } </pre>



NOTE: DHCP snooping statements have moved to a different hierarchy level.

Table 13: DHCP Snooping Statements

Original Hierarchy	Changed Hierarchy
<pre> ethernet-switching-options { secure-access-port { dhcp-snooping-file { location <i>local_pathname</i> <i>remote_URL</i>; timeout <i>seconds</i>; write-interval <i>seconds</i>; } } </pre>	<pre> system [processes [dhcp-service dhcp-snooping-file <i>local_pathname</i> <i>remote_URL</i>; write-interval <i>interval</i>; }] } </pre>

Changes to Configuring VLANs

The statements for configuring VLANs have moved to a different hierarchy level.



NOTE: When configuring xSTP on EX4300 and EX4600 switches, you must add all the interfaces in the applied VLANs in configurations. For MSTP, configure all interfaces in all VLANs at the [edit protocols mstp interface] hierarchy level.

Table 14: VLAN hierarchy

Original Hierarchy	Changed Hierarchy
<pre> ethernet-switching-options { secure-access-port vlan (all <i>vlan-name</i>) { mac-move-limit } } </pre>	<pre> vlangs <i>vlan-name</i> switch-options { mac-move-limit } </pre>
<pre> ethernet-switching-options { static { vlan <i>vlan-id</i> { mac <i>mac-address</i> next-hop <i>interface-name</i>; ... } } } </pre>	<p>NOTE: Statement is replaced with a new statement and has moved to a different hierarchy level.</p> <pre> vlangs { <i>vlan-name</i> { switch-options { interface <i>interface-name</i> { static-mac <i>mac-address</i>; ... } } } } </pre>
<pre> vlangs { <i>vlan-name</i> { interface <i>interface-name</i> { egress; ingress; mapping (native (push swap) policy tag (push swap)); pvlan-trunk; ... } } } </pre>	<p>These statements have been removed. You can assign interfaces to a VLAN using the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family ethernet-switching vlan members <i>vlan-name</i>] hierarchy.</p>

Table 14: VLAN hierarchy (*continued*)

Original Hierarchy	Changed Hierarchy
<pre>vlan { vlan-name { isolation-id id-number; ... } }</pre>	Statements have been removed.
<pre>vlan { vlan-name { l3-interface vlan.logical-interface-number; ... } }</pre>	<p>NOTE: Syntax is changed.</p> <pre>vlan { vlan-name { l3-interface irb.logical-interface-number; ... } }</pre>
<pre>vlan { vlan-name { l3-interface-ingress-counting layer-3-interface-name; ... } }</pre>	Statement is removed. Ingress traffic is automatically tracked.
<pre>vlan { vlan-name { no-local-switching; ... } }</pre>	Statement is removed.
<pre>vlan { vlan-name { no-mac-learning; ... } }</pre>	<p>Statement has been moved to different hierarchy.</p> <pre>vlan { vlan-name { switch-options { no-mac-learning limit ... } } }</pre>
<pre>vlan { vlan-name { primary-vlan vlan-name; ... } }</pre>	Statement has been removed.
<pre>vlan { vlan-name { vlan-prune; ... } }</pre>	Statement is removed.

Table 14: VLAN hierarchy (*continued*)

Original Hierarchy	Changed Hierarchy
<pre>vlan { vlan-name { vlan-range <i>vlan-id-low-vlan-id-high</i>; ... } }</pre>	<p>NOTE: Statement has been replaced with a new statement.</p> <pre>vlan { vlan-name { <i>vlan-id-list</i> [<i>vlan-id-numbers</i>]; ... } }</pre>
<pre>vlan { vlan-name { l3-interface <i>vlan.logical-interface-number</i>; ... } }</pre>	<p>NOTE: Syntax is changed.</p> <pre>vlan { vlan-name { l3-interface <i>irb.logical-interface-number</i>; ... } }</pre>

Table 15: Statements Moved to a Different Hierarchy

Original Hierarchy	Changed Hierarchy
<pre>vlan { vlan-name { dot1q-tunneling { customer-vlans (<i>id</i> <i>native</i> <i>range</i>); layer2-protocol-tunneling all <i>protocol-name</i> { drop-threshold <i>number</i>; shutdown-threshold <i>number</i>; ... } } } }</pre>	<pre>interface <i>interface-name</i> { encapsulation extended-vlan-bridge; <i>flexible-vlan-tagging</i>; <i>native-vlan-id number</i>; unit <i>logical-unit-number</i> { <i>input-vlan-map action</i>; <i>output-vlan-map action</i>; <i>vlan-id number</i>; <i>vlan-id-list</i> [<i>vlan-id vlan-id-vlan-id</i>]; } }</pre>
<pre>vlan { vlan-name { filter { input <i>filter-name</i> output <i>filter-name</i>; ... } } }</pre>	<pre>vlan { vlan-name { forwarding-options { filter { input <i>filter-name</i> output <i>filter-name</i>; ... } } } }</pre>

Table 15: Statements Moved to a Different Hierarchy (*continued*)

Original Hierarchy	Changed Hierarchy
<pre> vlangs { vlan-name { mac-limit limit action action; ... } } </pre>	<pre> vlangs { vlan-name { switch-options { interface-mac-limit limit { packet-action action; ... } } } } vlangs { vlan-name { switch-options { interface interface-name { interface-mac-limit limit { packet-action action; ... } } } } } </pre>
<pre> vlangs { vlan-name { mac-table-aging-time seconds; ... } } </pre>	<pre> protocols { l2-learning { global-mac-table-aging-time seconds; ... } } </pre>

Changes to Storm Control Profiles

Storm control is configured in two steps. The first step is to create a storm control profile at the **[edit forwarding-options]** hierarchy level, and the second step is to bind the profile to a logical interface at the **[edit interfaces]** hierarchy level. See *Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches* for the changed procedure.

Table 16: Changes to the Storm Control Profile hierarchy level

Original Hierarchy	Changed Hierarchy
<pre> ethernet-switching-options { storm-control { (...) } } </pre>	<pre> forwarding-options { storm-control-profiles profile-name { (...) } } interfaces interface-name unit number family ethernet-switching { storm-control storm-control-profile; } </pre>

Changes to the Interfaces Hierarchy



NOTE: Statements have been moved to a different hierarchy.

Table 17: Changes to the Interfaces hierarchy

Original Hierarchy	Changed Hierarchy
<pre> interfaces <i>interface-name</i> { ether-options { link-mode <i>mode</i>; speed (auto-negotiation <i>speed</i>) } } </pre>	<pre> interfaces <i>interface-name</i> { link-mode <i>mode</i>; speed <i>speed</i>) } </pre>
<pre> interfaces <i>interface-name</i> { unit <i>logical-unit-number</i> { family ethernet-switching { native-vlan-id <i>vlan-id</i> } } } </pre>	<pre> interfaces <i>interface-name</i> { native-vlan-id <i>vlan-id</i> } </pre>
<pre> interfaces <i>interface-name</i> { unit <i>logical-unit-number</i> { family ethernet-switching { port-mode <i>mode</i> } } } </pre>	<p>NOTE: Statement has been replaced with a new statement.</p> <pre> interfaces <i>interface-name</i> { unit <i>logical-unit-number</i> { family ethernet-switching { interface-mode <i>mode</i> } } } </pre>
<pre> interfaces vlan </pre>	<p>NOTE: Statement has been replaced with a new statement.</p> <pre> interfaces irb </pre>

Changes to IGMP Snooping

Table 18: IGMP Snooping hierarchy

Original Hierarchy	Changed Hierarchy
<pre> protocols { igmp-snooping { traceoptions { file filename <files number> <no-stamp> <replace> <size maximum-file-size> <world-readable no-world-readable>; flag flag <flag-modifier> <disable>; } vlan (all vlan-identifier) { disable; data-forwarding { receiver { install; source-vlans vlan-name; } source { groups ip-address; } } immediate-leave; interface (all interface-name) { multicast-router-interface; static { group multicast-ip-address; } } proxy { source-address ip-address; } robust-count number; } } } </pre>	<pre> protocols { igmp-snooping { vlan vlan-name { immediate-leave; interface interface-name { group-limit <1..65535> host-only-interface multicast-router-interface; immediate-leave; static { group multicast-ip-address { source <> } } } } l2-querier { source-address ip-address; } proxy { source-address ip-address; } query-interval number; query-last-member-interval number; query-response-interval number; robust-count number; traceoptions { file filename <files number> <no-stamp> <replace> <size maximum-file-size> <world-readable no-world-readable>; flag flag <flag-modifier>; } } } </pre>

Junos OS Operational Mode Commands That Combine Other Commands

In some cases, some Junos OS operational commands are created from a combination of other operational commands. These commands can be useful shortcuts for collecting information about the device, as shown in [Figure 1](#).

Figure 1: Commands That Combine Other Commands

The **request support information** command provides output from a combination of other operational commands.

```

user@host> request support information

root@host> show system uptime

Current time: 2007-02-16 13:10:08 PST
System booted: 2007-02-02 09:21:50 PST (2w0d 03:48 ago)
Protocols started: 2007-02-02 09:24:42 PST (2w0d 03:45 ago)
Last configured: 2007-02-16 03:04:58 PST (10:05:10 ago) by root
1:10PM up 14 days, 3:48, 2 users, load averages: 0.01, 0.02, 0.00

root@host> show version detail

Hostname: host
Model: m320
JUNOS Base OS boot [8.3-R1.1]

root@host> show system core-dumps

/var/tmp/*core*: No such file or directory
/var/crash/kernel.*: No such file or directory

/var/crash/cores:
total 9780
-rw-r--r-- 1 root wheel 4990976 Feb 9 15:39
core-FPC2.core.0.060209.1539

root@host> show chassis hardware detail

Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis
Backplane     REV 07   710-001517   AW44 31       M20 Backplane
Power Supply B REV 09   740-001466   0042 33       DC Power Supply

```

Related Documentation

- [Overview of Junos OS CLI Operational Mode Commands on page 58](#)
- [Understanding the Brief, Detail, Extensive, and Terse Options of Junos OS Operational Commands on page 62](#)

Overview of Junos OS CLI Operational Mode Commands

This topic provides an overview of Junos OS CLI operational mode commands and contains the following sections:

- [CLI Command Categories on page 58](#)
- [Commonly Used Operational Mode Commands on page 59](#)

CLI Command Categories

When you log in to a device running Junos OS and the CLI starts, there are several broad groups of CLI commands:

- Commands for controlling the CLI environment—Some set commands in the **set** hierarchy configure the CLI display screen. For information about these commands, see *Understanding the Junos OS CLI Modes, Commands, and Statement Hierarchies*.
- Commands for monitoring and troubleshooting—The following commands display information and statistics about the software and test network connectivity. Detailed command descriptions are provided in the *Junos OS Interfaces Command Reference*.

- **clear**—Clear statistics and protocol database information.
- **mtrace**—Trace mtrace packets from source to receiver.
- **monitor**—Perform real-time debugging of various software components, including the routing protocols and interfaces.
- **ping**—Determine the reachability of a remote network host.
- **show**—Display the current configuration and information about interfaces, routing protocols, routing tables, routing policy filters, system alarms, and the chassis.
- **test**—Test the configuration and application of policy filters and autonomous system (AS) path regular expressions.
- **traceroute**—Trace the route to a remote network host.
- Commands for connecting to other network systems—The **ssh** command opens Secure Shell connections, and the **telnet** command opens telnet sessions to other hosts on the network. For information about these commands, see the [CLI Explorer](#).
- Commands for copying files—The **copy** command copies files from one location on the router or switch to another, from the router or switch to a remote system, or from a remote system to the router or switch. For information about these commands, see the [CLI Explorer](#).
- Commands for restarting software processes—The commands in the **restart** hierarchy restart the various Junos OS processes, including the routing protocol, interface, and SNMP. For information about these commands, see the [CLI Explorer](#).
- A command—**request**—for performing system-level operations, including stopping and rebooting the router or switch and loading Junos OS images. For information about this command, see the [CLI Explorer](#).
- A command—**start**—to exit the CLI and start a UNIX shell. For information about this command, see the [CLI Explorer](#).
- A command—**configure**—for entering configuration mode, which provides a series of commands that configure Junos OS, including the routing protocols, interfaces, network management, and user access. For information about the CLI configuration commands, see “[Understanding Junos OS CLI Configuration Mode](#)” on page 63.
- A command—**quit**—to exit the CLI. For information about this command, see the [CLI Explorer](#).
- For more information about the CLI operational mode commands, see the [CLI Explorer](#).

Commonly Used Operational Mode Commands

[Table 19](#) lists some operational commands you may find useful for monitoring router or switch operation. For a complete description of operational commands, see the Junos OS command references.



NOTE: The QFX3500 switch does not support the IS-IS, OSPF, BGP, MPLS, and RSVP protocols.

Table 19: Commonly Used Operational Mode Commands

Items to Check	Description	Command
Software version	Versions of software running on the router or switch	show version
Log files	Contents of the log files	monitor
	Log files and their contents and recent user logins	show log
Remote systems	Host reachability and network connectivity	ping
	Route to a network system	traceroute
Configuration	Current system configuration	show configuration
Manipulate files	List of files and directories on the router or switch	file list
	Contents of a file	file show
Interface information	Detailed information about interfaces	show interfaces
Chassis	Chassis alarm status	show chassis alarms
	Information currently on craft display	show chassis craft-interface
	Router or switch environment information	show chassis environment
	Hardware inventory	show chassis hardware
Routing table information	Information about entries in the routing tables	show route
Forwarding table information	Information about data in the kernel's forwarding table	show route forwarding-table
IS-IS	Adjacent routers or switches	show isis adjacency
OSPF	Display standard information about OSPF neighbors	show ospf neighbor
BGP	Display information about BGP neighbors	show bgp neighbor

Table 19: Commonly Used Operational Mode Commands (*continued*)

Items to Check	Description	Command
MPLS	Status of interfaces on which MPLS is running	show mpls interface
	Configured LSPs on the router or switch, as well as all ingress, transit, and egress LSPs	show mpls lsp
	Routes that form a label-switched path	show route label-switched-path
RSVP	Status of interfaces on which RSVP is running	show rsvp interface
	Currently active RSVP sessions	show rsvp session
	RSVP packet and error counters	show rsvp statistics

Related Documentation

- [Junos OS Operational Mode Commands That Combine Other Commands on page 57](#)
- [Understanding the Brief, Detail, Extensive, and Terse Options of Junos OS Operational Commands on page 62](#)

Overview of Navigating the CLI

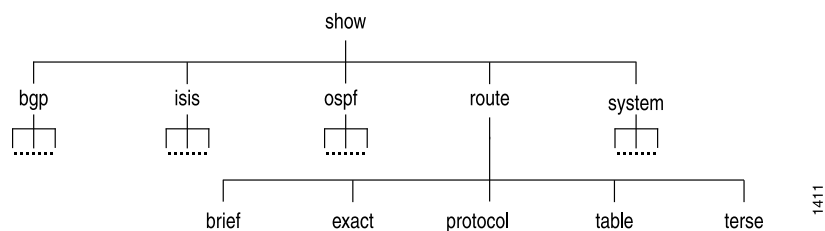
This topic describes how to navigate the CLI.

- [CLI Command Hierarchy on page 61](#)
- [CLI Configuration Statements on page 62](#)
- [Moving Among Hierarchy Levels on page 62](#)

CLI Command Hierarchy

CLI commands are organized in a hierarchy. Commands that perform a similar function are grouped together under the same level of the hierarchy. For example, all commands that display information about the system and the system software are grouped under the **show system** command, and all commands that display information about the routing table are grouped under the **show route** command. [Figure 2](#) illustrates a portion of the **show** command hierarchy.

Figure 2: CLI Command Hierarchy



To execute a command, you enter the full command name, starting at the top level of the hierarchy. For example, to display a brief view of your Ethernet switching options for your interfaces, use the command **show ethernet-switching-options interfaces**.

CLI Configuration Statements

The configuration statement hierarchy has two types of statements: *container statements*, which are statements that contain other statements, and *leaf statements*, which do not contain other statements. All of the container and leaf statements together form the *configuration hierarchy*.

The **protocols** statement is a top-level statement at the trunk of the configuration tree. The **ospf**, **area**, and **interface** statements are all subordinate container statements of a higher statement (they are branches of the hierarchy tree), and the **hello-interval** statement is a leaf on the tree.

Moving Among Hierarchy Levels

You can use the CLI commands to navigate the levels of the configuration statement hierarchy:

- **edit**— Moves to an existing configuration statement hierarchy or creates a hierarchy and moves to that level.
- **exit**— Moves up the hierarchy to the previous level where you were working. This command is, in effect, the opposite of the **edit** command. Alternatively, you can use the **quit** command. The **exit** and **quit** commands are interchangeable.
- **up**— Moves up the hierarchy one level at a time.
- **top**— Moves directly to the top level of the hierarchy.

Related Documentation

- [CLI User Interface Overview on page 37](#)
- [CLI User Guide](#)

Understanding the Brief, Detail, Extensive, and Terse Options of Junos OS Operational Commands

The Junos OS operational mode commands can include **brief**, **detail**, **extensive**, or **terse** options. You can use these options to control the amount of information you want to view.

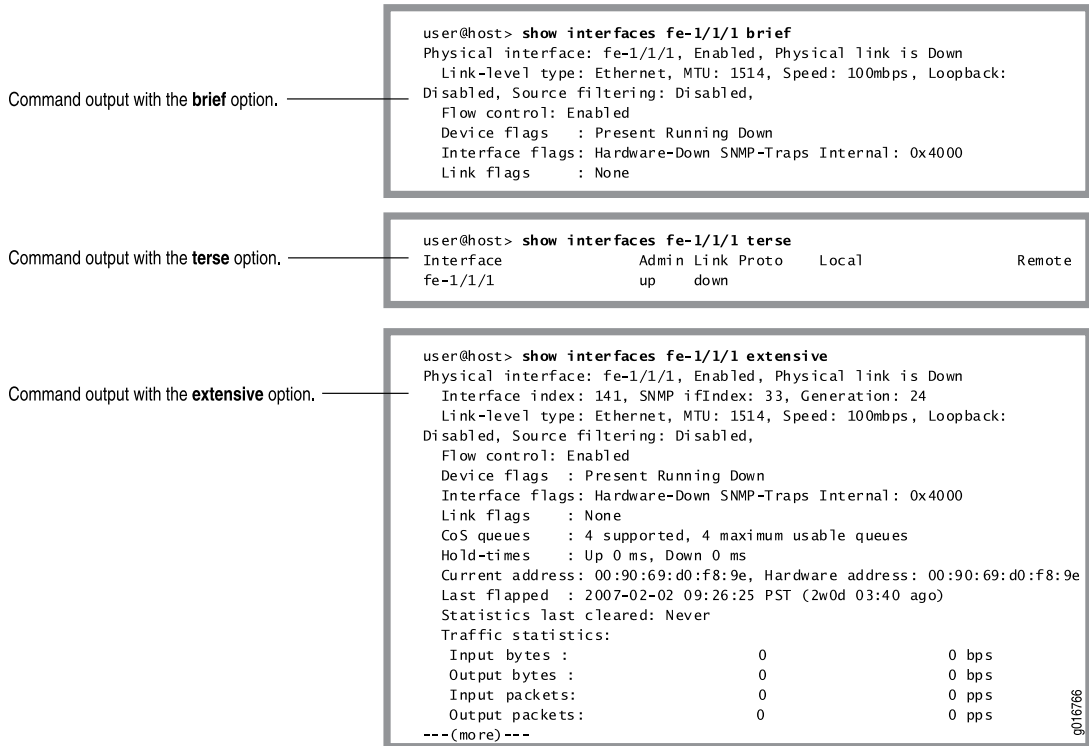
1. Use the **?** prompt to list options available for the command. For example:

```
user@host> show interfaces fe-1/1/1 ?
Possible completions:
<[Enter]>          Execute this command
brief              Display brief output
descriptions       Display interface description strings
detail             Display detailed output
extensive          Display extensive output
media              Display media information
snmp-index         SNMP index of interface
```

statistics	Display statistics and detailed output
terse	Display terse output
	Pipe through a command

2. Choose the option you wish to use with the command. (See [Figure 3.](#))

Figure 3: Command Output Options



Related Documentation

- [Overview of Junos OS CLI Operational Mode Commands on page 58](#)
- [Controlling the Scope of an Operational Mode Command](#)

Understanding Junos OS CLI Configuration Mode

You can configure all properties of Junos OS, including interfaces, general routing information, routing protocols, and user access, as well as several system hardware properties.

As described in *Understanding the Junos OS CLI Modes, Commands, and Statement Hierarchies*, a router configuration is stored as a hierarchy of statements. In configuration mode, you create the specific hierarchy of configuration statements that you want to use. When you have finished entering the configuration statements, you commit them, which activates the configuration on the router.

You can create the hierarchy interactively or you can create an ASCII text file that is loaded onto the router or switch and then committed.

This topic covers:

- [Configuration Mode Commands on page 65](#)
- [Configuration Statements and Identifiers on page 66](#)
- [Configuration Statement Hierarchy on page 68](#)

Configuration Mode Commands

Table 20 summarizes each CLI configuration mode command. The commands are organized alphabetically.

Table 20: Summary of Configuration Mode Commands

Command	Description
activate	Remove the inactive: tag from a statement, effectively reading the statement or identifier to the configuration. Statements or identifiers that have been activated take effect when you next issue the commit command.
annotate	Add comments to a configuration. You can add comments only at the current hierarchy level.
commit	Commit the set of changes to the database and cause the changes to take operational effect.
copy	Make a copy of an existing statement in the configuration.
deactivate	Add the inactive: tag to a statement, effectively commenting out the statement or identifier from the configuration. Statements or identifiers marked as inactive do not take effect when you issue the commit command.
delete	Delete a statement or identifier. All subordinate statements and identifiers contained within the specified statement path are deleted with it.
edit	Move inside the specified statement hierarchy. If the statement does not exist, it is created.
exit	Exit the current level of the statement hierarchy, returning to the level prior to the last edit command, or exit from configuration mode. The quit and exit commands are synonyms.
extension	Manage configurations that are contributed by SDK application packages. Either display or delete user-defined configuration contributed by the named SDK application package. A configuration defined in any native Junos OS package is never deleted by the extension command.
help	Display help about available configuration statements.
insert	Insert an identifier into an existing hierarchy.
load	Load a configuration from an ASCII configuration file or from terminal input. Your current location in the configuration hierarchy is ignored when the load operation occurs.

Table 20: Summary of Configuration Mode Commands (*continued*)

Command	Description
quit	Exit the current level of the statement hierarchy, returning to the level prior to the last edit command, or exit from configuration mode. The quit and exit commands are synonyms.
rename	Rename an existing configuration statement or identifier.
replace	Replace identifiers or values in a configuration.
rollback	Return to a previously committed configuration. The software saves the last 10 committed configurations, including the rollback number, date, time, and name of the user who issued the commit configuration command.
run	Run a top-level CLI command without exiting from configuration mode.
save	Save the configuration to an ASCII file. The contents of the current level of the statement hierarchy (and below) are saved, along with the statement hierarchy containing it. This allows a section of the configuration to be saved, while fully specifying the statement hierarchy.
set	Create a statement hierarchy and set identifier values. This is similar to edit except that your current level in the hierarchy does not change.
show	Display the current configuration.
status	Display the users currently editing the configuration.
top	Return to the top level of configuration command mode, which is indicated by the [edit] banner.
up	Move up one level in the statement hierarchy.
update	Update a private database.
wildcard	Delete a statement or identifier. All subordinate statements and identifiers contained within the specified statement path are deleted with it. You can use regular expressions to specify a pattern. Based on this pattern, you search for items that contain these patterns and delete them.

Configuration Statements and Identifiers

You can configure router or switch properties by including the corresponding statements in the configuration. Typically, a statement consists of a keyword, which is fixed text, and, optionally, an identifier. An identifier is an identifying name that you can define, such as

the name of an interface or a username, which enables you and the CLI to differentiate among a collection of statements.

Table 21 describes top-level CLI configuration mode statements.



NOTE: The QFX3500 switch does not support the IS-IS, OSPF, BGP, LDP, MPLS, and RSVP protocols.

Table 21: Configuration Mode Top-Level Statements

Statement	Description
access	Configure the Challenge Handshake Authentication Protocol (CHAP). For information about the statements in this hierarchy, see the <i>Junos OS Administration Library for Routing Devices</i> .
accounting-options	Configure accounting statistics data collection for interfaces and firewall filters. For information about the statements in this hierarchy, see the <i>Network Management Administration Guide for Routing Devices</i> .
chassis	Configure properties of the router chassis, including conditions that activate alarms and SONET/SDH framing and concatenation properties. For information about the statements in this hierarchy, see the <i>Junos OS Administration Library for Routing Devices</i> .
class-of-service	Configure class-of-service parameters. For information about the statements in this hierarchy, see the <i>Class of Service Feature Guide for Routing Devices</i> .
firewall	Define filters that select packets based on their contents. For information about the statements in this hierarchy, see the <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i> .
forwarding-options	Define forwarding options, including traffic sampling options. For information about the statements in this hierarchy, see the <i>Junos OS Network Interfaces Library for Routing Devices</i> .
groups	Configure configuration groups. For information about statements in this hierarchy, see the <i>Junos OS Administration Library for Routing Devices</i> .
interfaces	Configure interface information, such as encapsulation, interfaces, virtual channel identifiers (VCIs), and data-link connection identifiers (DLCIs). For information about the statements in this hierarchy, see the <i>Junos OS Network Interfaces Library for Routing Devices</i> .
policy-options	Define routing policies, which allow you to filter and set properties in incoming and outgoing routes. For information about the statements in this hierarchy, see the <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i> .
protocols	Configure routing protocols, including BGP, IS-IS, LDP, MPLS, OSPF, RIP, and RSVP. For information about the statements in this hierarchy, see the chapters that discuss how to configure the individual routing protocols in the <i>Junos OS Routing Protocols Library for Routing Devices</i> and the <i>MPLS Applications Feature Guide for Routing Devices</i> .

Table 21: Configuration Mode Top-Level Statements (*continued*)


Statement	Description
routing-instances	Configure multiple routing instances. For information about the statements in this hierarchy, see the <i>Junos OS Routing Protocols Library for Routing Devices</i> .
routing-options	Configure protocol-independent routing options, such as static routes, autonomous system numbers, confederation members, and global tracing (debugging) operations to log. For information about the statements in this hierarchy, see the <i>Junos OS Routing Protocols Library for Routing Devices</i> .
security	Configure IP Security (IPsec) services. For information about the statements in this hierarchy see the <i>Junos OS Administration Library for Routing Devices</i> .
snmp	Configure SNMP community strings, interfaces, traps, and notifications. For information about the statements in this hierarchy, see the <i>Network Management Administration Guide for Routing Devices</i> .
system	Configure systemwide properties, including the hostname, domain name, Domain Name System (DNS) server, user logins and permissions, mappings between hostnames and addresses, and software processes. For information about the statements in this hierarchy, see the <i>Junos OS Administration Library for Routing Devices</i> .

For specific information on configuration statements, see the Junos OS configuration guides.

Configuration Statement Hierarchy

The Junos OS configuration consists of a hierarchy of *statements*. There are two types of statements: *container statements*, which are statements that contain other statements, and *leaf statements*, which do not contain other statements (see [Figure 4](#)). All of the container and leaf statements together form the *configuration hierarchy*.

Figure 4: Configuration Mode Hierarchy of Statements



Each statement at the top level of the configuration hierarchy resides at the trunk (or root level) of a hierarchy tree. The top-level statements are container statements, containing other statements that form the tree branches. The leaf statements are the leaves of the hierarchy tree. An individual hierarchy of statements, which starts at the trunk of the hierarchy tree, is called a *statement path*. Figure 4 illustrates the hierarchy tree, showing a statement path for the portion of the protocol configuration hierarchy that configures the hello interval on an interface in an OSPF area.

The **protocols** statement is a top-level statement at the trunk of the configuration tree. The **ospf**, **area**, and **interface** statements are all subordinate container statements of a higher statement (they are branches of the hierarchy tree); and the **hello-interval** statement is a leaf on the tree which in this case contains a data value: the length of the hello interval, in seconds.

The CLI represents the statement path shown in Figure 4 as **[edit protocols ospf area *area-number* interface *interface-name*]** and displays the configuration as follows:

```
protocols {
  ospf {
    area 0.0.0.0 {
      interface so-0/0/0 {
        hello-interval 5;
      }
      interface so-0/0/1 {
        hello-interval 5;
      }
    }
  }
}
```

The CLI indents each level in the hierarchy to indicate each statement's relative position in the hierarchy and generally sets off each level with braces, using an open brace at the beginning of each hierarchy level and a closing brace at the end. If the statement at a hierarchy level is empty, the braces are not printed.

Each leaf statement ends with a semicolon. If the hierarchy does not extend as far as a leaf statement, the last statement in the hierarchy ends with a semicolon.

The configuration hierarchy can also contain “oneliners” at the last level in the hierarchy. Oneliners remove one level of braces in the syntax and display the container statement, its identifiers, the child or leaf statement and its attributes all on one line. For example, in the following sample configuration hierarchy, the line **level 1 metric 10** is a oneliner because the **level** container statement with identifier **1**, its child statement **metric**, and its corresponding attribute **10** all appear on a single line in the hierarchy:

```
[edit protocols]
isis {
  interface ge-0/0/0.0 {
    level 1 metric 10;
  }
}
```

Likewise, in the following example, **dynamic-profile *dynamic-profile-name* aggregate-clients;** is a oneliner because the **dynamic-profile** statement, its identifier ***dynamic-profile-name***, and leaf statement **aggregate-clients** all appear on one line when you run the **show** command in the configuration mode:

```
[edit forwarding-options]
user@host# show
dhcp-relay {
  dynamic-profile dynamic-profile-name aggregate-clients;
}
```

**Related
Documentation**

- *Entering and Exiting the Junos OS CLI Configuration Mode*

PART 4

Dynamic Host Control Protocol (DHCP)

- [Understanding DHCP on page 73](#)

CHAPTER 4

Understanding DHCP

- [Understanding DHCP Services for Switches on page 73](#)
- [Configuring a DHCP Client \(CLI Procedure\) on page 77](#)
- [Configuring a DHCP Server on Switches \(CLI Procedure\) on page 78](#)

Understanding DHCP Services for Switches

A Dynamic Host Configuration Protocol (DHCP) server on a switch can provide many valuable TCP/IP network services. For example, DHCP can dynamically allocate the four required IP parameters to each computer on the LAN: IP address, network mask, switch address, and name server address. Additionally, DHCP on the switch can automatically upgrade software on client systems.


This topic describes:

- [DHCP Client/Server Model on page 73](#)
- [Using DHCP on page 74](#)
- [DHCP Relay Servers and DHCP Servers on page 74](#)
- [Legacy DHCP and Extended DHCP for Server Versions on page 75](#)
- [Configuring DHCP on a Switch on page 76](#)
- [How DHCP Works on page 76](#)

DHCP Client/Server Model

DHCP IP address allocation works on a client/server model in which the server, in this case a switch, assigns the client reusable IP information from an address pool. A DHCP client might receive offer messages from multiple DHCP servers and can accept any one of the offers; however, the client usually accepts the first offer it receives. See [Figure 5](#).

Figure 5: DHCP Client/Server Model



Using DHCP

DHCP automates network-parameter assignment to network devices. Even in small networks, DHCP is useful because it makes it easy to add new machines to the network.

DHCP access service minimizes the overhead required to add clients to the network by providing a centralized, server-based setup, which means that you do not have to manually create and maintain IP address assignments for clients. In addition, when you use DHCP to manage a pool of IP addresses among hosts, you reduce the number of IP addresses needed on the network. DHCP does this by leasing an IP address to a host for a limited period of time, allowing the DHCP server to share a limited number of IP addresses. DHCP also provides a central database of devices that are connected to the network and eliminates duplicate resource assignments. In addition to IP addresses for clients, DHCP provides other configuration information, particularly the IP addresses of local caching Domain Name System (DNS) resolvers, network boot servers, or other service hosts.

Another valuable DHCP feature is automatic software download for installation of software packages on switches. DHCP clients configured for automatic software download receive messages as part of the DHCP message exchange process—when the software package name in the DHCP server message is different from that of the software package that booted the DHCP client switch, the new software is downloaded and installed. See [“Upgrading Software by Using Automatic Software Download” on page 7](#).

DHCP Relay Servers and DHCP Servers

You can configure a switch either as a DHCP server or as a DHCP relay server, but not both. Whereas a DHCP server replies to a client with an IP address, a DHCP relay server relays DHCP messages to and from the configured DHCP server, even if the client and server are on different IP networks.

Configure a switch to be a DHCP relay agent if you have locally attached hosts and a remote DHCP server. For directions on configuring a DHCP relay server, see *DHCP/BOOTP Relay for Switches Overview*.

Legacy DHCP and Extended DHCP for Server Versions

Two versions of both DHCP server and DHCP relay agent are available on EX Series, QFX Series, and OCX Series switches. The original legacy DHCP server and legacy DHCP relay agent can be used in the same network as the extended DHCP servers and extended DHCP relay agent—extended DHCP is also referred to as virtual router (VR) aware DHCP.

You cannot configure legacy DHCP and extended DHCP versions on the same switch. Because the newer extended DHCP server version has more features, we recommend that you configure the extended DHCP server if it is supported by the switch. See *EX Series Switch Software Features Overview* for a list of switches that support the extended DHCP server.

The extended DHCP server version has the following added features:

- Graceful Routing Engine switchover (GRES), which provides mirroring support for clients. For details, see *High Availability Features for EX Series Switches Overview*.
- Virtual routing and forwarding (VRF), which allows multiple instances of a routing table to simultaneously coexist on the same switch. For details, see *Understanding Virtual Routing Instances on EX Series Switches*.



NOTE: Legacy DHCP supports the circuit ID and the remote ID fields for the relay agent option (option 82). Extended DHCP for the relay agent option supports only circuit ID. See *EX Series Switch Software Features Overview* for a list of switches that support extended DHCP (VR-aware DHCP).

Legacy DHCP and extended DHCP servers can be configured at the hierarchy levels shown in [Table 22](#):

Table 22: Legacy DHCP and Extended DHCP Server Hierarchy Levels

DHCP Service	Hierarchy
Extended DHCP server	<code>edit system services dhcp-local-server</code>
Extended DHCP address pool	<code>edit access address-assignment pool</code>
Legacy DHCP server	<code>edit system services dhcp</code>
Legacy DHCP relay	<code>edit forwarding-options helpers bootp</code>
Extended DHCP relay	<code>edit forwarding-options dhcp-relay</code>
Legacy DHCP address pool	<code>edit system services dhcp pool</code>

DHCP clients on a switch are always configured at the hierarchy level `[edit interfaces interface-name family dhcp]`.

Configuring DHCP on a Switch

A DHCP configuration consists of two parts: the configuration for a DHCP server and the configuration for DHCP clients. The DHCP server configuration is simple if you accept the default configurations.

When you configure a legacy DHCP server, you only need to define the DHCP server name and the interface on the switch. You can use the default configuration for the rest of the settings. When you configure an extended DHCP server, you need to only define a DHCP pool, indicate IP addresses for the pool, and create a server group. You can use the default configuration for the rest of the settings.


For directions for configuring either a legacy DHCP server or an extended DHCP server, see [“Configuring a DHCP Server on Switches \(CLI Procedure\)” on page 78](#).

To configure a DHCP client, set the client's DHCP interface address in the **[edit interfaces interface-name unit 0 family inet dhcp]** hierarchy. For directions for configuring a DHCP client on a switch, see [“Configuring a DHCP Client \(CLI Procedure\)” on page 77](#).

How DHCP Works

DHCP consists of a four-step transfer process beginning with a broadcast DHCP discovery message from the client. As the second step, the client receives a DHCP offer message from the server. This message includes the IP address and mask, and some other specific parameters. The client then sends a DHCP request message to accept the IP address and other parameters that it received from the server in the previous step. The DHCP server sends a DHCP response message and removes the now-allocated address from the DHCP address pool. See [Figure 6](#).

Figure 6: DHCP Four-Step Transfer



NOTE: Because the DHCP discovery message from the client is a broadcast message and because broadcast messages cross other segments only when they are explicitly routed, you might have to configure a DHCP relay agent on the switch interface so that all DHCP discovery messages from the clients are forwarded to one DHCP server.

Related Documentation

- [Configuring a DHCP Client \(CLI Procedure\) on page 77](#)
- [Configuring a DHCP Server on Switches \(CLI Procedure\) on page 78](#)

- [Configuring an Extended DHCP Relay Server on EX Series Switches \(CLI Procedure\)](#)
- [Configuring a DHCP SIP Server \(CLI Procedure\)](#)
- [Upgrading Software by Using Automatic Software Download on page 7](#)
- [Monitoring DHCP Services](#)

Configuring a DHCP Client (CLI Procedure)

A Dynamic Host Configuration Protocol (DHCP) server can provide many valuable TCP/IP network services. DHCP can dynamically allocate IP parameters, such as an IP address, to clients, and it can also deliver software upgrades to clients.

DHCP configuration consists of two components, configuration of DHCP clients and configuration of a DHCP server. Client configuration determines how clients send a message requesting an IP address, whereas a DHCP server configuration enables the server to send an IP address configuration back to the client. This topic describes configuring a DHCP client. For directions for configuring a DHCP server, see [“Configuring a DHCP Server on Switches \(CLI Procedure\)” on page 78](#) or [Configuring a DHCP Server on Switches \(CLI Procedure\)](#).

You can change DHCP client configurations from the switch, using client identifiers to indicate which clients you want to configure.

To configure a DHCP client, you configure an interface to belong to the DHCP family and specify additional attributes, as desired:

```
[edit]
user@switch# set interfaces interface-name unit number family inet dhcp
configuration-statement
```

The options that you can configure are listed in [Table 23](#). Replace the variable *configuration-statement* with one or more of the statements listed in this table. If you do not explicitly configure these options, the switch uses default values for them.

Table 23: DHCP Client Settings

Configuration Statement	Description
client-identifier	Unique client ID—By default this consists of the hardware type (01 for Ethernet) and the MAC address (a.b.c.d). For this example, the value would be 01abcd.
lease-time	Time in seconds that a client holds the lease for an IP address assigned by a DHCP server. If a client does not request a specific lease time, then the server sends the default lease time. The default lease time on a Junos OS DHCP server is 1 day.
retransmission-attempt	Number of times the client attempts to retransmit a DHCP packet.
retransmission-interval	Time between transmission attempts.
server-address	IP address of the server that the client queries for an IP address.

Table 23: DHCP Client Settings (*continued*)

Configuration Statement	Description
update-server	TCP/IP settings learned from an external DHCP server to the DHCP server running on the switch are propagated.
vendor-option	Vendor class ID (CPU's manufacturer ID string) for the DHCP client.

**Related
Documentation**

- [Configuring a DHCP Server on Switches \(CLI Procedure\) on page 78](#)
- [Understanding DHCP Services for Switches on page 73](#)

Configuring a DHCP Server on Switches (CLI Procedure)



NOTE: This task uses Junos OS for EX Series switches that does not support the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that supports ELS, see *Configuring a DHCP Server on Switches (CLI Procedure)*. For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on page 41.

A Dynamic Host Configuration Protocol (DHCP) server can provide two valuable TCP/IP network services. DHCP can dynamically allocate IP parameters, such as an IP address, to clients and it can also deliver software upgrades to clients.

A DHCP configuration consists of two components—an optional reconfiguration of default settings on DHCP clients and the configuration of a DHCP server. This topic covers configuration of the DHCP server. For information about reconfiguring a DHCP client, see “[Configuring a DHCP Client \(CLI Procedure\)](#)” on page 77.

You can configure either of two versions of a DHCP server on a switch— the extended server version or the legacy server version. We recommend that you configure the extended server unless you need to keep your DHCP server configuration backward-compatible with the legacy server version.

This topic includes the following tasks:

1. [Configuring an Extended DHCP Server on a Switch on page 79](#)
2. [Configuring a Legacy DHCP Server on a Switch \(CLI Procedure\) on page 79](#)

Configuring an Extended DHCP Server on a Switch

To configure an extended DHCP server, you must configure a DHCP pool, indicate IP addresses for the pool, and create a server group. Additional configurations are optional.

Do not assign addresses that are already in use in the network to address pools. The extended DHCP server does not check whether addresses are already in use before it assigns them to clients.

1. Create an address pool for DHCP IP addresses:

```
[edit]
user@switch# set access address-pool address-pool
```

2. Configure an address-assignment pool that can be used by different client applications for DHCP dynamic assignment:

```
[edit access address-assignment]
user@switch# set pool address-pool-name
```

3. Create a server group on the switch, providing a group name and an interface name for DHCP:

```
[edit system services dhcp-local-server]
user@switch# set group group-name interface interface-name
```

4. (Optional) Process the information protocol data units (PDUs):

```
[edit system services dhcp-local-server]
user@switch# set overrides process-inform
```

5. (Optional) Redefine the order of attribute matching for pool selection:

```
[edit system services dhcp-local-server]
user@switch# set pool-match-order ip-address-first
```

6. (Optional) Enable dynamic reconfiguration triggered by the DHCP extended server for all DHCP clients or only for the DHCP clients serviced by the specified group of interfaces:

```
[edit system services dhcp-local-server]
user@switch# set reconfigure

[edit system services dhcp-local-server group group-name]
user@switch# set reconfigure
```

Configuring a Legacy DHCP Server on a Switch (CLI Procedure)

To configure a legacy DHCP server, you must configure a pool of IP addresses for dynamic assignment. You only need to supply a series of network addresses. Additional configurations are optional.

1. Configure a pool of IP addresses for dynamic assignment:

```
[edit system services dhcp]
user@switch# set pool network-range
```



NOTE: Step 2 through Step 15 are for assigning global values at the `[edit system services dhcp]` hierarchy level. You can also assign the same values to a specific pool by using those same commands at the `[edit system services dhcp pool network-range]` hierarchy level.

2. (Optional) Change the domain search list used to resolve hostnames:

```
[edit system services dhcp]
user@switch# set domain-search [ domain-list ]
```

3. (Optional) Change the domain name server (DNS) name that the DHCP server advertises to clients:

```
[edit system services dhcp]
user@switch# set name-server address
```

4. (Optional) Change the DHCP options:

```
[edit system services dhcp]
user@switch# set option id-number
```

5. (Optional) Change the devices advertised to clients:

```
[edit system services dhcp]
user@switch# set router address
```

6. (Optional) Configure the name of the boot server advertised to DHCP clients. The client uses a boot file located on the boot server to complete the DHCP setup. This configuration step is equivalent to DHCP Option 66:

```
[edit system services dhcp]
user@switch# set boot-server (address | hostname)
```

7. (Optional) Set the boot file advertised to DHCP clients. After the client receives an IP address and the boot file location from the DHCP server, the client uses the boot image stored in the boot file to complete DHCP setup. This configuration step is equivalent to DHCP Option 67:

```
[edit system services dhcp]
user@switch# set boot-file filename
```

8. (Optional) Change the SIP server:

```
[edit system services dhcp]
user@switch# set sip-server addresses-or-names
```

For more information, see *Configuring a DHCP SIP Server (CLI Procedure)*.

9. (Optional) Change the DHCP client's hardware address:

```
[edit system services dhcp]
user@switch# set static-binding mac-address
```

10. (Optional) Change the NetBIOS name server:

```
[edit system services dhcp]
user@switch# set wins-server address
```


- Related Documentation**
- [Configuring a DHCP Client \(CLI Procedure\) on page 77](#)
 - [Configuring a DHCP SIP Server \(CLI Procedure\)](#)
 - [Understanding DHCP Services for Switches on page 73](#)

PART 5

Internet Control Message Protocol (ICMP)

- [Understanding ICMP on page 85](#)

CHAPTER 5

Understanding ICMP

- [Configuring Junos OS to Disable the Reporting of IP Address and Timestamps in Ping Responses on page 85](#)
- [Configuring Junos OS to Disable the Routing Engine Response to Multicast Ping Packets on page 86](#)
- [Configuring the Junos OS ICMPv4 Rate Limit for ICMPv4 Routing Engine Messages on page 86](#)

Configuring Junos OS to Disable the Reporting of IP Address and Timestamps in Ping Responses

When you issue the **ping** command with the **record-route** option, the Routing Engine displays the path of the ICMP echo request packets and timestamps in the ICMP echo responses by default.

You can configure the Routing Engine to disable the setting of the **record-route** option in the IP header of the ping request packets. Disabling the **record-route** option prevents the Routing Engine from recording and displaying the path of the ICMP echo request packets in the response.

- To configure the Routing Engine to disable the setting of the **record route** option, include the **no-ping-record-route** statement at the **[edit system]** hierarchy level:

```
[edit system]  
no-ping-record-route;
```

- To disable the reporting of timestamps in the ICMP echo responses, include the **no-ping-time-stamp** option at the **[edit system]** hierarchy level:

```
[edit system]  
no-ping-time-stamp;
```

By configuring the **no-ping-record-route** and **no-ping-timestamp** options, you can prevent unauthorized persons from discovering information about the provider edge (PE) router or switch and its loopback address.

Related Documentation

- [Configuring Junos OS to Disable the Routing Engine Response to Multicast Ping Packets on page 86](#)

Configuring Junos OS to Disable the Routing Engine Response to Multicast Ping Packets

By default, the Routing Engine responds to Internet Control Message Protocol (ICMP) echo requests sent to multicast group addresses. To disable the Routing Engine from responding to ICMP echo requests sent to multicast group addresses, include the **no-multicast-echo** statement at the **[edit system]** hierarchy level:

```
[edit system]
no-multicast-echo;
```

By configuring the Routing Engine to ignore multicast ping packets, you can prevent unauthorized persons from discovering the list of provider edge (PE) routers or switches in the network.

- Related Documentation**
- [Configuring Junos OS to Disable the Reporting of IP Address and Timestamps in Ping Responses on page 19](#)

Configuring the Junos OS ICMPv4 Rate Limit for ICMPv4 Routing Engine Messages

To limit the rate at which ICMPv4 messages can be generated and received by the Routing Engine, include the **icmpv4-rate-limit** statement at the **[edit system internet-options]** hierarchy level:

```
icmpv4-rate-limit bucket-size bucket-size packet-rate packet-rate;
```

The bucket size is the number of seconds in the rate-limiting bucket. The packet rate is the rate-limiting packets earned per second. Specify a **bucket-size** from 0 through 4294967295 seconds. The default value is 5 seconds. Specify a **packet-rate** from 0 through 4,294,967,295. The default value is 1000.

- Related Documentation**
- [icmpv4-rate-limit on page 225](#)

PART 6

Junos OS With Upgraded FreeBSD

- [Understanding Junos OS With Upgraded FreeBSD on page 89](#)

CHAPTER 6

Understanding Junos OS With Upgraded FreeBSD

- [Understanding Junos OS with Upgraded FreeBSD on page 89](#)
- [Understanding Junos OS with Upgraded FreeBSD Package Names on page 92](#)
- [Understanding Junos OS with Upgraded FreeBSD Disk Volumes on page 94](#)
- [Understanding Junos OS with Upgraded FreeBSD Snapshots on page 95](#)
- [Upgrading Junos OS with Upgraded FreeBSD on page 96](#)
- [Downgrading Junos OS from Upgraded FreeBSD on page 102](#)

Understanding Junos OS with Upgraded FreeBSD

Starting with Junos OS Release 15.1, certain hardware platforms run a Junos OS based on an upgraded FreeBSD kernel instead of older versions of FreeBSD. Basing Junos OS on the newer kernel provides Junos OS with sophisticated processing, efficiency, and security features which do not then have to be reproduced in Junos OS.



NOTE: Upgrading to Junos OS Release 15.1 reformats the file system. Only specific files and directories are preserved unless precautions are taken. For details, see [“Upgrading Junos OS with Upgraded FreeBSD” on page 96](#).

Junos OS with an upgraded FreeBSD kernel provides a clean-slate implementation of Junos OS on top of a pristine (minimally modified) and current version of the FreeBSD OS.



NOTE: In Junos OS releases earlier than 15.1, the partition swap pages were counted as part of the memory file system partition. Using this method leaves 4 GB of memory as the maximum that is theoretically accessible when you are using a 32-bit image. However, when Junos OS with upgraded FreeBSD is run, the system only counts the actual partition size, which leaves around 3.4 GB of available physical address space, or only 3 GB of usable RAM. Therefore, we recommend you use a 64-bit image with Junos OS with upgraded FreeBSD.

The platforms currently running Junos OS with upgraded FreeBSD are listed in [Table 24](#).

Table 24: Upgraded FreeBSD Kernel Support by Hardware Platform

Platforms	CPU Type	Release Introduced
MX240, MX460, MX960, MX2010, MX2020	Intel	15.1
EX9200	Intel	15.1

The major processing changes are as follows:

- Interactions between Junos OS and the upgraded FreeBSD kernel use well-established interfaces because Junos OS is now layered on a minimally modified and current version of FreeBSD.
- Symmetric multiprocessing (SMP) is enabled by default.
- FreeBSD provides a consistent runtime environment for all Junos OS platforms.

There are also major changes in file structures and software packages. These changes are as follows:

- New packages use XML description files instead of scripts.
- Hybrid packages are used to install legacy or replacement build images in the general form **junos-upgrade-x.tgz** where x is a variable such as **mx-x86-64-15.1-20150114** (the whole package name is **junos-upgrade-mx-x86-64-15.1-20150114.tgz**).
- Multiple package sets (a collection of installed packages) are stored on the router at the same time. Sets can be either active (the currently used set), pending (the set that should be used at the next reboot), or previous (a formerly active set). Non-recovery snapshots (but not recoverable image snapshots) are available for the package sets to preserve package content lists.

There is now a separate Operations, Administration, and Maintenance (OAM) volume (**oam**) distinct from the Junos OS volume (**junos**). This provides support for downgrades from replacement build images (that is, those using the upgraded FreeBSD kernel) to the legacy Junos OS with a different kernel. The OAM volume allows you to recover the Junos OS volume using recovery snapshots.

One major change is the distinction between recovery snapshots and non-recovery snapshots.

The major characteristics of the recovery snapshots are as follows:

- Recovery snapshots are full copies of the packages and configuration taken at the time the snapshot command is issued.
- Recovery snapshots reside on the OAM volume or USB medium.

The major characteristics of the non-recovery snapshots are as follows:

- Non-recovery snapshots are snapshots residing on the Junos OS volume that refer to the current running set of packages and a copy of the configuration at the time the snapshot command is issued.
- Non-recovery snapshots do not need to copy the whole Junos OS installation and so are very fast.
- Non-recovery snapshots can be requested as the boot image for the next reboot.

The upgraded FreeBSD kernel requires changes to several commands and statements and their related parameters. The new and changed actions are summarized in [Table 25](#). For details on the changes, see the topics covering the specific command or statement.

Table 25: New and Changed Commands and Statements for Junos OS with Upgraded FreeBSD

Command or Statement	Release Introduced	Change
<code>request system snapshot delete <i>snapshot</i></code>	15.1	New action
<code>request system snapshot recovery</code>	15.1	New action
<code>request system snapshot load <i>snapshot</i></code>	15.1	New action
<code>request system recover <i>volume</i></code>	15.1	New action: <i>volume</i> is either <code>/junos-volume</code> or <code>/oam-volume</code>
<code>request system snapshot</code>	15.1	Changed action
<code>show system snapshot</code>	15.1	Changed action
<code>request system reboot <i>media</i></code>	15.1	Changed action with new media options

The new FreeBSD kernel also requires that several commands and statements are now deprecated. In some cases, these commands and statements generate an error, and, in other cases, the result is appropriate for the new kernel. The deprecated commands and statements are summarized in [Table 26](#). For details, see the topics covering the specific command or statement.

Table 26: Deprecated Commands and Statements for Junos OS with Upgraded FreeBSD

Deprecated Command or Configuration Statement	Release Deprecated
Deprecated Command	
<code>request system partition abort</code>	15.1
<code>request system partition compact-flash</code>	15.1
<code>request system partition hard-disk</code>	15.1

Table 26: Deprecated Commands and Statements for Junos OS with Upgraded FreeBSD (*continued*)

Deprecated Command or Configuration Statement	Release Deprecated
<code>request system snapshot <config-partition></code>	15.1
<code>request system snapshot <root-partition></code>	15.1
<code>request system snapshot <slice></code>	15.1
<code>request system software delete-backup</code>	15.1
<code>request system software rollback <force></code>	15.1
<code>show system processes providers</code>	15.1
<code>show system snapshot <slice></code>	15.1
Deprecated Configuration Statement	
<code>set system mirror-flash-on-disk</code>	15.1

Related Documentation

- [Upgrading Junos OS with Upgraded FreeBSD on page 96](#)
- [Downgrading Junos OS from Upgraded FreeBSD on page 102](#)
- [request system snapshot \(Junos OS with Upgraded FreeBSD\) on page 265](#)
- [show system snapshot \(Junos OS with Upgraded FreeBSD\) on page 267](#)
- [request system reboot \(Junos OS with Upgraded FreeBSD\) on page 262](#)

Understanding Junos OS with Upgraded FreeBSD Package Names

Starting with Junos OS Release 15.1, certain hardware platforms run a Junos OS based on an upgraded FreeBSD kernel (hereafter called Junos OS with upgraded FreeBSD). In releases earlier than Junos OS Release 15.1, software packages came in several major software package categories, such as domestic, world-wide, or Federal Information Processing Standard (FIPS). However, Junos OS with upgraded FreeBSD has a new naming convention: There is only one category, and FIPS, instead of being a separate category, is an option you select on installation. This topic describes the simplified naming convention for Junos OS with upgraded FreeBSD.

If your hardware platform is listed in the table in [“Understanding Junos OS with Upgraded FreeBSD” on page 89](#), then you must use the new package names for download and installation.

The components of the new package naming conventions are as follows:

- **Prefix**—This is normally **junos-install**. This prefix takes the place of what was called a bundle name, but we do use the term *bundle* in the new package naming convention.
- **Media keyword**—Added to the prefix, a media keyword is only used when the image is not for use with the **request system software add** command. Values for the **media** keyword include **usb** for images installed from a USB drive or **net** for images installed over a network, for example, the entire prefix of your package might be **junos-install-usb-**.
- **Platform**—This field indicates the major product group, such as **mx** or **qfx**.
- **Architecture**—This field indicates the CPU architecture of the platforms. Values include **x86** for Intel and **arm** for Advanced RISC Machines CPUs.
- **Application Binary Interface (ABI)**—This field indicates the “word length” of the CPU architecture. Values include **32** for 32-bit architectures and **64** for 64-bit architectures.
- **Release**—This field indicates the release number, such as **15.1R1.9**.
- **Edition**—The edition field is null (empty) for the standard (domestic) images. For the Russian Trade Zone (RTZ) images, this field is set to **limited**.

As before, all images are in tarred and gzipped (**.tgz**) format.



NOTE: There are no longer “export” World-Wide images or separate FIPS images. The keyword “signed” no longer appears because all Junos OS images are signed for validation.

Examples of valid Junos OS software package names include the following:

- **junos-install-mx-x86-32-15.1R1.9.tgz**—an image for a supported MX platform outside the RTZ.
- **junos-install-mx-x86-32-15.1R1.9-limited.tgz**—an image for a supported MX platform used in the RTZ.
- **junos-install-usb-mx-x86-32-15.1R1.9.tgz**—an image stored on and installed from a USB drive for a supported MX platform outside the RTZ.

Because an upgrade to Junos OS with upgraded FreeBSD from a release earlier than Junos OS 15.1 restructures the disk file system, you can lose many configuration and log files that you might want to keep. Items that are essential can be preserved by moving or copying them to the **/var/preserve** directory.

Related Documentation

- [Understanding Junos OS with Upgraded FreeBSD on page 89](#)
- [Upgrading Junos OS with Upgraded FreeBSD on page 96](#)

Understanding Junos OS with Upgraded FreeBSD Disk Volumes

Starting with Junos OS Release 15.1, certain hardware platforms have a new disk naming convention. These platforms run a Junos OS based on an upgraded FreeBSD kernel instead of older versions of FreeBSD.

The hardware platforms listed in the table in [“Understanding Junos OS with Upgraded FreeBSD” on page 89](#) have two volumes. The main device is the **/junos** volume and contains all of the software and files needed for the day-to-day running of the device. The compact flash drive is the **/oam** volume and stores recovery snapshot backup information. In case of failure of the main drive (that is, the **/junos** volume), the **/oam** volume can be used to boot the system.

Because the **/junos** and **/oam** volumes have very different purposes, their content is different. Technically, these volumes are **dev/gpt/oam** and **dev/gpt/junos**, but the short forms (**/junos** and **/oam**) are used in this topic. Essentially, the **/junos** volume is used for the running device software and holds configuration information and logs, whereas the **/oam** volume is used for backup copies of everything needed in the event that the **/junos** volume fails.

The **/junos** volume contains a directory named **/packages/db** that has all the components present on the device, such as **os-kernel-123**, **os-kernel-456**, and so on. A sibling directory named **/package-sets** is also present. Package sets are an important concept in Junos OS with upgraded FreeBSD.

The **/package-sets** directory contains a package listing that gathers all the components of the running Junos OS into an XML format in the **/active** subdirectory. So **os-kernel-123** could be a component in the **/package-sets/active** subdirectory, but then **os-kernel-456** could not be in the same XML package. Package sets do not contain the kernel software itself (for example), but tell the device where to find the kernel component needed for the software package. The same kernel can be present in several package listings, but only one package can be active and running on the device at any given time.

There are several directories on the **/junos** volume where a particular software package listing can be found:

- **/previous**—The package set in this directory contains the list of all the components that ran on the device before the last upgrade.
- **/active**—The package set in this directory contains the list of all the software components currently running on the device.
- **/pending**—The package set in this directory contains the list of all the software components on the device that will run after the next reboot.



NOTE: After a successful reboot, the package set in the **/pending** directory becomes the active package set, and the package set in the **/active** directory becomes the previous set.

The **/junos** volume also contains non-recovery snapshots taken with the **request system snapshot** command. These types of snapshots are new to Junos OS with upgraded FreeBSD and cannot be used for recovery of a failed system. Non-recovery snapshots are a special type of package set that includes a copy of the configuration. There can be many non-recovery snapshots on the device, and the files can be renamed. Multiple non-recovery snapshots, essentially lists of software components and configuration files, can be helpful when major software or configuration changes are occurring and establishment of a known stable system baseline is required.

On the other hand, a recovery snapshot, created with the **request system snapshot recovery** command, is stored on the **/oam** volume and is always replaced when a new recovery snapshot is taken.

The **/oam** volume should contain all the information needed to reboot the system if there is a failure of the **/junos** volume and restore the system to the state running at the time of the failure. In order to perform this reboot, the **/oam** volume needs to have all of the information required to provide the system with a running configuration. This information is provided by the recovery snapshot, created with the **request system snapshot recovery** command. Although it can take a while to perform, the recovery snapshot establishes an **.izo** or **.iso** image of the running Junos OS.

In the case of a total failure of the **/junos** volume, the system can be booted from the **/oam** volume. The recovery snapshot can then restore the repaired system.

**Related
Documentation**

- [Understanding Junos OS with Upgraded FreeBSD on page 89](#)
- [Upgrading Junos OS with Upgraded FreeBSD on page 96](#)

Understanding Junos OS with Upgraded FreeBSD Snapshots

Starting with Junos OS Release 15.1, certain hardware platforms have two types of snapshots. These platforms run a Junos OS based on an upgraded FreeBSD kernel instead of older versions of FreeBSD. The two types of snapshots have different content, locations, and purposes, so it is important that they are created and maintained properly. One major change is the distinction between recovery snapshots and non-recovery snapshots. The hardware platforms listed in the table in “[Understanding Junos OS with Upgraded FreeBSD](#)” on page 89 have these two different types of snapshots.

Recovery snapshots are full copies of the packages and configuration taken at the time the snapshot command is issued. Recovery snapshots reside on the OAM volume or USB medium. Recovery snapshots take some time to complete because of the level of detail captured. Recovery snapshots can be used to recover the Junos OS volume. There is only ever one recovery snapshot on the system.

On the other hand, non-recovery snapshots are snapshots residing on the Junos OS volume that refer to the current running set of packages and a copy of the configuration at the time the snapshot command is issued. Non-recovery snapshots do not need to copy the whole Junos OS installation and so are very fast. They also consume little space, except for the **config.tgz** file. Non-recovery snapshots can be requested as the boot image for the next reboot. You can rename non-recovery snapshots and retain more than one.

You rename the non-recovery snapshots with the same procedure used to rename any other file on the system.



NOTE: We recommend that you generate both a non-recovery and a recovery snapshot after you successfully upgrade to Junos OS with upgraded FreeBSD. These snapshots should be refreshed periodically.

Package sets relate to non-recovery and recovery snapshots. The **/active**, **/pending**, and **/previous** sets are all package sets. A non-recovery snapshot is also a package set in a sense, with the addition of a copy of the configuration at the time that the non-recovery snapshot is taken.

Packages that are no longer referenced by any package set or non-recovery snapshot are automatically deleted. We recommend deleting any old non-recovery snapshots after an upgrade so that old packages can be deleted and space recovered.

Some helpful commands for non-recovery snapshots are:

- **request system snapshot**—Use this command to create a non-recovery snapshot.
- **show system snapshot**—Use this command to list all the available non-recovery snapshots.
- **request system snapshot delete**—Use this command to delete a non-recovery snapshot.
- **request system snapshot recovery**—Use this command to create a recovery snapshot. You can use other parameters to determine the details of the recovery snapshot created. There is only ever one recovery snapshot on the system.

**Related
Documentation**

- [request system snapshot \(Junos OS with Upgraded FreeBSD\) on page 265](#)
- [show system snapshot \(Junos OS with Upgraded FreeBSD\) on page 267](#)
- [request system reboot \(Junos OS with Upgraded FreeBSD\) on page 262](#)
- [request system software validate on](#)
- [Understanding Junos OS with Upgraded FreeBSD Package Names on page 92](#)
- [Understanding Junos OS with Upgraded FreeBSD Disk Volumes on page 94](#)

Upgrading Junos OS with Upgraded FreeBSD

Starting with Junos OS Release 15.1, certain hardware platforms run an upgraded FreeBSD kernel instead of older versions of FreeBSD.

Before you begin:

1. Verify that the upgrade applies to your router or switch model, as listed in [“Understanding Junos OS with Upgraded FreeBSD” on page 89](#).
2. Download the Junos OS package.

The current Junos OS release determines the upgrade path to Junos OS with upgraded FreeBSD, as shown in [Table 27](#). Other upgrade paths might work, but they are not supported.

Table 27: Upgrade Path to Junos OS with the Upgraded FreeBSD

Current Router's Junos OS Release	Upgrade Path
12.3 or earlier	Upgrade to 13.3. or 14.2 first, then upgrade to 15.1 or later (multiple steps).
13.3 or later	Use upgrade package to upgrade from the current Junos OS release to Junos OS with upgraded FreeBSD (single step).
15.1 or later	Use upgrade package to upgrade from the current Junos OS release to Junos OS with upgraded FreeBSD (single step).



NOTE: You can also downgrade from Junos OS Release 15.1 to an earlier release of Junos OS, as long as the path complies with the Junos OS policy of skipping at most two releases earlier.

The upgrade process only preserves the following directories:

- `/config`
- `/etc/localtime`
- `/var/db`
- `/var/etc/master.passwd`
- `/var/etc/inetd.conf`
- `/var/etc/pam.conf`
- `/var/etc/resolv.conf`
- `/var/etc/syslog.conf`
- `/var/etc/localtime`
- `/var/etc/exports`
- `/var/etc/extensions.allow`
- `/var/preserve`
- `/var/tmp/baseline-config.conf`
- `/var/tmp/preinstall_boot_loader.conf`

For specific installation procedures, see the following:

- [To Install Junos OS with Upgraded FreeBSD Over a Plain Junos OS on page 98](#)
- [To Install Junos OS with Upgraded FreeBSD Over Junos OS with Upgraded FreeBSD of an Earlier Release on page 101](#)
- [To Install Junos OS with Upgraded FreeBSD Over Junos OS with Upgraded FreeBSD of a Later Release on page 102](#)

To Install Junos OS with Upgraded FreeBSD Over a Plain Junos OS



NOTE: If you have important files in other directories, copy them from the router or switch to a secure location before upgrading the router or switch.



NOTE: The following procedure refers to routers, but it also applies to switches.

To install Junos OS with upgraded FreeBSD over a plain Junos OS:

1. Enter the **request system software add *package-name* no-validate** command from the operational mode in the CLI:



NOTE: The **no-copy** option is enabled by default.

Use the **no-validate** option with the **request system software add** command. If you leave out the **no-validate** option, the command uses the **validate** option by default, and direct validation of running configuration does not work for upgrading to Junos OS with upgraded FreeBSD from Junos OS based on older versions of the FreeBSD kernel.



NOTE: You can also use **reboot** option along with **request system software add** command, but it is not recommended to do this in a single step while upgrading from a FreeBSD 6.1 based Junos OS to FreeBSD 10 based Junos OS.



NOTE: To validate current configuration on an upgrade to Junos OS with upgraded FreeBSD from Junos OS, use the **request system software validate on** command.

```
user@host>request system software add
/var/tmp/junos-install-mx-x86-32-15.1R1.9.tgz no-validate
Installing package '/var/tmp/junos-install-mx-x86-32-15.1R1.9.tgz' ...
Verified manifest signed by PackageProductionEc_2015
Verified manifest signed by PackageProductionRSA_2015
```

```

Verified contents.iso
Verified issu-indb.tgz
Verified junos-x86-32.tgz
Verified kernel
Verified metatags
Verified package.xml
Verified pkgtools.tgz
camcontrol: not found
camcontrol: not found
Verified manifest signed by PackageProductionEc_2015
Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Saving package file in
/var/sw/pkg/junos-install-x86-32-domestic-20150618.043753_builder_junos_151_r1.tgz
...
Saving state for rollback ...

```

The new Junos OS image is installed on the router.

2. Reboot the device to start the new software using the **request system reboot** command:

```

user@host> request system reboot
Reboot the system? [yes, no] (no) yes

```



NOTE: You must reboot the device to load the newly installed version of Junos OS on the device.

To abort the installation, do not reboot the device. Instead, finish the installation and then issue the **request system software delete *package-name*** command where package is, for example, **junos-install-mx-x86-32-15.1R1.9.tgz**. This is your last chance to stop the installation.

The software is loaded when you reboot the system. Installation can take between 5 and 10 minutes. The device then reboots from the boot device on which the software was just installed. When the reboot is complete, the device displays the login prompt.

While the software is being upgraded, the Routing Engine on which you are performing the installation does not route traffic.

3. Log in and issue the **show version** command to verify the version of the software installed.

```

user@host> show version
Hostname: spice
Model: mx240
Junos: 15.1R1.9
JUNOS OS Kernel 32-bit [20150617.306001_builder_stable_10]
JUNOS OS runtime [20150617.306001_builder_stable_10]
JUNOS OS time zone information [20150617.306001_builder_stable_10]
JUNOS py base [20150618.043753_builder_junos_151_r1]
JUNOS OS crypto [20150617.306001_builder_stable_10]
JUNOS network stack and utilities [20150618.043753_builder_junos_151_r1]
JUNOS libs [20150618.043753_builder_junos_151_r1]

```

```
JUNOS runtime [20150618.043753_builder_junos_151_r1]
JUNOS platform support [20150618.043753_builder_junos_151_r1]
JUNOS modules [20150618.043753_builder_junos_151_r1]
JUNOS daemons [20150618.043753_builder_junos_151_r1]
JUNOS Voice Services Container package [20150618.043753_builder_junos_151_r1]
JUNOS Services SSL [20150618.043753_builder_junos_151_r1]
JUNOS Services Stateful Firewall [20150618.043753_builder_junos_151_r1]
JUNOS Services RPM [20150618.043753_builder_junos_151_r1]
JUNOS Services PTSP Container package [20150618.043753_builder_junos_151_r1]
JUNOS Services NAT [20150618.043753_builder_junos_151_r1]
JUNOS Services Mobile Subscriber Service Container package
[20150618.043753_builder_junos_151_r1]
JUNOS Services MobileNext Software package
[20150618.043753_builder_junos_151_r1]
JUNOS Services LL-PDF Container package [20150618.043753_builder_junos_151_r1]
JUNOS Services Jflow Container package [20150618.043753_builder_junos_151_r1]
JUNOS Services IPSec [20150618.043753_builder_junos_151_r1]
JUNOS IDP Services [20150618.043753_builder_junos_151_r1]
JUNOS Services HTTP Content Management package
[20150618.043753_builder_junos_151_r1]
JUNOS Services Crypto [20150618.043753_builder_junos_151_r1]
JUNOS Services Captive Portal and Content Delivery Container package
[20150618.043753_builder_junos_151_r1]
JUNOS Border Gateway Function package [20150618.043753_builder_junos_151_r1]
JUNOS Appld Services [20150618.043753_builder_junos_151_r1]
JUNOS Services Application Level Gateways [20150618.043753_builder_junos_151_r1]
JUNOS Services AACL Container package [20150618.043753_builder_junos_151_r1]
JUNOS Packet Forwarding Engine Support (MX/EX92XX Common)
[20150618.043753_builder_junos_151_r1]
JUNOS Packet Forwarding Engine Support (M/T Common)
[20150618.043753_builder_junos_151_r1]
JUNOS Online Documentation [20150618.043753_builder_junos_151_r1]
JUNOS FIPS mode utilities [20150618.043753_builder_junos_151_r1]
```



NOTE: The output shows the OS kernel, OS runtime, and other packages installed on the router.

To Install Junos OS with Upgraded FreeBSD Over Junos OS with Upgraded FreeBSD of an Earlier Release



NOTE: If you have important files in other directories, copy them from the router or switch to a secure location before upgrading the router or switch.



NOTE: The following procedure refers to routers, but it also applies to switches.

To install Junos OS with upgraded FreeBSD over Junos OS with upgraded FreeBSD of an earlier release:

1. Enter the **request system software add *package-name* validate reboot** command from the operational mode in the CLI:



NOTE: The **no-copy** option is enabled by default.

Use the **validate** and **reboot** options with the **request system software add** command. The command uses the **validate** option by default. We encourage users to validate using the **validate** option when upgrading from Junos OS to Junos OS or from Junos OS with upgraded FreeBSD to Junos OS with upgraded FreeBSD.

If you leave out the **reboot** option, you can take care of that in a separate reboot step.

The new Junos OS image is installed on the router.

2. Verify the installation of Junos OS with upgraded FreeBSD.

```
user@host> show version
```



NOTE: The output shows the OS kernel, OS runtime, and other packages installed on the router.

To Install Junos OS with Upgraded FreeBSD Over Junos OS with Upgraded FreeBSD of a Later Release



NOTE: If you have important files in other directories, copy them from the router or switch to a secure location before upgrading the router or switch.



NOTE: The following procedure refers to routers, but it also applies to switches.

To install Junos OS with upgraded FreeBSD over Junos OS with upgraded FreeBSD of a later release:

1. Enter the **request system software add *package-name* validate reboot** command from the operational mode in the CLI:



NOTE: The **no-copy** option is enabled by default.

Use the **validate** and **reboot** options with the **request system software add** command. The command uses the **validate** option by default. We encourage users to validate using the **validate** option when upgrading from Junos OS to Junos OS or from Junos OS with upgraded FreeBSD to Junos OS with upgraded FreeBSD.

If you leave out the **reboot** option, you can take care of that in a separate reboot step.

The new Junos OS image is installed on the router.

2. Verify the installation of Junos OS with upgraded FreeBSD.

```
user@host> show version
```



NOTE: The output shows the OS kernel, OS runtime, and other packages installed on the router.

Related Documentation

- [Downgrading Junos OS from Upgraded FreeBSD on page 102](#)
- [Understanding Junos OS with Upgraded FreeBSD on page 89](#)
- [request system snapshot \(Junos OS with Upgraded FreeBSD\) on page 265](#)
- [request system reboot \(Junos OS with Upgraded FreeBSD\) on page 262](#)

Downgrading Junos OS from Upgraded FreeBSD

Starting with Junos OS Release 15.1, certain hardware platforms run a Junos OS based on an upgraded FreeBSD kernel instead of older versions of FreeBSD. If you have previously upgraded to Junos OS with upgraded FreeBSD, you can downgrade to earlier versions

versions of Junos OS, as long as the downgrade conforms to the Junos OS policy of skipping at most two earlier releases.

Before you begin:

1. Verify that you have previously upgraded to Junos OS with the upgraded FreeBSD kernel, as described in [“Upgrading Junos OS with Upgraded FreeBSD” on page 96](#).
2. Download the Junos OS package.

This example uses the package `/var/tmp/jinstall-13.3R2.7-domestic-signed.tgz` to install Junos OS with a pre-upgraded FreeBSD kernel on the master Routing Engine (**re0**).



NOTE: The following procedure refers to routers, but it also applies to switches.

To downgrade from Junos OS with upgraded FreeBSD to Junos OS:

1. Enter the **request system software add package-name no-validate reboot** command from the operational mode in the CLI:

Use the **no-validate** and **reboot** options with the **request system software add** command. If you leave out the **no-validate** option, the command uses the **validate** option by default, and direct validation of running configuration does not work for downgrading to Junos OS from Junos OS with upgraded FreeBSD.



NOTE: To validate current configuration on a downgrade to Junos OS from Junos OS with upgraded FreeBSD, use the **request system software validate on** command.

If you leave out the **reboot** option, you can take care of that in a separate reboot step.

The following example uses the **re0** option.

```
user@host>request system software add
/var/tmp/jinstall-13.3R2.7-domestic-signed.tgz re0 no-validate reboot
THIS IS A SIGNED PACKAGE Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install Rebooting. Please wait ...
shutdown: [pid 11001] Shutdown NOW! *** FINAL System shutdown message
from root@host *** System going down IMMEDIATELY Shutdown NOW! System
shutdown time has arrived\x07\x07 regress@host> Connection to
host.juniper.net closed by remote host. Connection to
host.juniper.net closed. ... user@router> show version
Hostname: host
Model: mx240
Junos: 13.3R2.7
JUNOS Base OS boot [13.3R2.7]
JUNOS Base OS Software Suite [13.3R2.7]
JUNOS Kernel Software Suite [13.3R2.7]
JUNOS Crypto Software Suite [13.3R2.7]
JUNOS Packet Forwarding Engine Support (M/T/EX Common) [13.3R2.7]
```

JUNOS Packet Forwarding Engine Support (MX Common) [13.3R2.7]
JUNOS Online Documentation [13.3R2.7]
JUNOS Services AACL Container package [13.3R2.7]
...

2. Verify the downgrade of the software package.

```
user@host> show version
```



NOTE: The output shows the OS kernel, OS runtime, and other packages installed on the router.

This example uses the package `/var/tmp/jinstall-13.3R2.7-domestic-signed.tgz` to install Junos OS with a pre-upgraded FreeBSD kernel on the master Routing Engine (re0).



NOTE: The following procedure refers to routers, but it also applies to switches.

To downgrade from Junos OS with upgraded FreeBSD to an earlier release of Junos OS with upgraded FreeBSD:

1. Enter the **request system software add package-name validate reboot** command from the operational mode in the CLI:

Use the **validate** and **reboot** options with the **request system software add** command. The command uses the **validate** option by default.

If you leave out the **reboot** option, you can take care of that in a separate reboot step.

2. Verify the downgrade of the software package.

```
user@host> show version
```



NOTE: The output shows the OS kernel, OS runtime, and other packages installed on the router.

Related Documentation

- [Upgrading Junos OS with Upgraded FreeBSD on page 96](#)
- [Understanding Junos OS with Upgraded FreeBSD on page 89](#)
- [request system snapshot \(Junos OS with Upgraded FreeBSD\) on page 265](#)
- [request system reboot \(Junos OS with Upgraded FreeBSD\) on page 262](#)

PART 7

Licenses

- [Using Licenses on page 107](#)

CHAPTER 7

Using Licenses

- [Junos OS Feature Licenses on page 107](#)
- [Software Features That Require Licenses on the QFX Series on page 108](#)
- [Junos OS Feature License Keys on page 112](#)
- [Generating License Keys on page 115](#)
- [Adding New Licenses \(CLI Procedure\) on page 117](#)
- [Deleting a License \(CLI Procedure\) on page 118](#)
- [Saving License Keys on page 119](#)
- [Verifying Junos OS License Installation on page 120](#)

Junos OS Feature Licenses

Some Junos OS software features require a license to activate the feature. To enable a licensed feature, you need to purchase, install, manage, and verify a license key that corresponds to each licensed feature. To conform to Junos OS feature licensing requirements, you must purchase one license per feature per device. The presence of the appropriate software license key on your device determines whether you are eligible to configure and use the licensed feature.

To speed deployment of licensed features, Junos OS software implements an honor-based licensing structure and provides you with a 30-day grace period to use a licensed feature without a license key installed. The grace period begins when you configure the feature and your device uses the licensed feature for the first time, but not necessarily when you install the license. After the grace period expires, the system generates system log messages saying that the feature requires a license. To clear the error message and use the licensed feature properly, you must install and verify the required license.

For information about how to purchase software licenses, contact your Juniper Networks sales representative.

Related Documentation

- *License Enforcement*
- *Working with License Keys for SRX Series Devices*
- *Junos OS Feature License Model Number for SRX Series Services Gateways*

Software Features That Require Licenses on the QFX Series



NOTE: If you try to configure a feature that is not licensed, you will receive syslog messages saying that you are using a feature that is licensable and that you do not possess a license for the feature. If you try to commit configuration changes for a feature that is not licensed, you will receive a commit warning saying that you have exceeded the allowed license limit for the feature.



NOTE: Virtual Extensible Local Area Network (VXLAN) is not supported on QFX3500 and QFX3600 devices. When you issue the `show licenses` command, you will see VXLAN in the CLI output, but the feature is not enabled.



NOTE: There is no separate license for Virtual Chassis like there is for Virtual Chassis Fabric.

We provide licenses for both standard Junos OS features and disaggregated Junos OS features.

Disaggregated Junos OS is offered starting with the QFX5200 switch. The QFX5200 hardware model number is hardware only and does not include the right to use the basic Junos OS functions. You can purchase feature licenses from three different tiers for disaggregated Junos OS. The disaggregated Junos OS feature licenses are available on a perpetual basis.

The following disaggregated Junos OS feature licenses are available:

- Junos Base Services (JBS)

The JBS license includes basic layer 2 switching, basic layer 3 routing, multicast, automation, programmability, Zero Touch Provisioning (ZTP) and basic monitoring.

- Junos Advanced Services (JAS)

The JAS license includes features supported in the JBS license and Border Gateway Protocol (BGP), Intermediate System-to-Intermediate System (IS-IS), and Virtual Extensible Local Area Network (VXLAN).

- Junos Premium Service (JPS)

The JPS license includes features supported in the JAS license and the Multi-protocol Label Switching (MPLS) feature set.

[Table 28](#) lists the disaggregated Junos OS feature licenses and supported QFX Series devices, and [Table 29](#) lists the standard Junos OS features licenses and supported QFX Series devices.

For information about how to purchase a software license, contact your Juniper Networks sales representative.

Table 28: Disaggregated Junos OS Feature Licenses and Model Numbers for QFX Series Devices

Licensed Software Feature	Supported Devices	Number of Licenses Required	Model Number
Junos base services (JBS) license includes basic layer 2 switching, basic layer 3 routing, multicast, automation, programmability, Zero Touch Provisioning (ZTP) and basic monitoring.	QFX5200-32C switch	One per switch.	QFX5000-35-JBS
Junos advanced services (JAS) license includes features supported in JBS license and Border Gateway Protocol (BGP), Intermediate System-to-Intermediate System (IS-IS), and Virtual Extensible Local Area Network (VXLAN)	QFX5200-32C switch	One per switch.	QFX5000-35-JAS
Junos premium services (JPS) license includes features supported in JAS license and Multi-protocol Label Switching (MPLS) feature set.	QFX5200-32C switch	One per switch.	QFX5000-35-JPS

Table 29: Standard Junos OS Feature Licenses and Model Numbers for QFX Series Devices

Licensed Software Feature	Supported Devices	Number of Licenses Required	Model Number
QFX Series premium feature license for Border Gateway Protocol (BGP), Intermediate System-to-Intermediate System (IS-IS), and Virtual Extensible Local Area Network (VXLAN), and Open vSwitch Database (OVSDB)	QFX10002-36Q switch	One per switch	QFX10002-36Q-PFL
	QFX10002-72Q switch		QFX10002-72Q-PFL
	QFX10008 switch		QFX10008-PFL

Table 29: Standard Junos OS Feature Licenses and Model Numbers for QFX Series Devices (*continued*)

Licensed Software Feature	Supported Devices	Number of Licenses Required	Model Number
QFX Series advanced feature license for Border Gateway Protocol (BGP), Intermediate System-to-Intermediate System (IS-IS), Multi-protocol Label Switching (MPLS), and Virtual Extensible Local Area Network (VXLAN), and Open vSwitch Database (OVSDB)	QFX10002-36Q switch	One per switch	QFX10002-36Q-AFL
	QFX10002-72Q switch		QFX10002-72Q-AFL
	QFX10008 switch		QFX10008-AFL
QFX Series advanced feature license for Border Gateway Protocol (BGP), Intermediate System-to-Intermediate System (IS-IS), Multi-protocol Label Switching (MPLS), and Virtual Extensible Local Area Network (VXLAN), and Open vSwitch Database (OVSDB)	QFX3500, QFX3600, QFX5100-48S, and QFX5100-48T switches	One per switch, two per Virtual Chassis, and two per Virtual Chassis Fabric	QFX-JSL-EDGE-ADV1
QFX Series advanced feature license for Border Gateway Protocol (BGP), Intermediate System-to-Intermediate System (IS-IS), Multi-protocol Label Switching (MPLS), and Virtual Extensible Local Area Network (VXLAN) and Open vSwitch Database (OVSDB)	QFX5100-24Q and QFX5100-96S switches	One per switch, two per Virtual Chassis, and two per Virtual Chassis Fabric	QFX5100-HDNSE-LIC
QFX Series advanced feature license for Border Gateway Protocol (BGP)	QFX3100 Director device	One per Node device in a network Node group	QFX-JSL-DRCTR-ADV1
QFX Series advanced feature license for Fibre Channel	QFX3500 switch	One per switch on which fibre channel ports are configured	QFX-JSL-EDGE-FC

Table 29: Standard Junos OS Feature Licenses and Model Numbers for QFX Series Devices (*continued*)

Licensed Software Feature	Supported Devices	Number of Licenses Required	Model Number
QFX Series advanced feature license for Fibre Channel	QFX3100 Director device	One per QFX3500 Node device on which fibre channel ports are configured	QFX-JSL-DRCTR-FC
QFX Series advanced feature license for Fibre Channel - Capacity 16	QFX3100 Director device	One for up to 16 QFX3500 Node devices on which fibre channel ports are configured	QFX-JSL-DRCTR-FC-C16
QFX Series feature license for enabling fabric mode	QFX3500 and QFX3600 device	One per device	QFX3000-JSL-EDGE-FAB
QFX Series feature license for base software for QFX3000-G QFabric system	QFX3100 Director device	One per QFX3000-G QFabric system	QFX3008-JSL-DRCTR-FAB
QFX Series feature license for base software for QFX3000-M QFabric system	QFX3100 Director device	One per QFX3000-M QFabric system	QFX3000M-JSL-DRCTR-FAB
QFX and EX Series feature license for enabling Media Access Control security (MACsec)	QFX switches that support MACsec. See <i>Understanding Media Access Control Security (MACsec)</i> .	One per switch, two per Virtual Chassis,	EX-QFX-MACSEC-AGG
Virtual Chassis Fabric (VCF)	All member devices in a Virtual Chassis Fabric (VCF)	Two per Virtual Chassis Fabric (VCF)	QFX-VCF-LIC

Related Documentation

- [Junos OS Feature Licenses on page 107](#)
- [Junos OS Feature License Keys on page 112](#)
- [Generating License Keys on page 115](#)
- [Adding New Licenses \(CLI Procedure\) on page 117](#)
- [Deleting a License \(CLI Procedure\) on page 118](#)
- [Saving License Keys on page 119](#)
- [Verifying Junos OS License Installation on page 120](#)

Junos OS Feature License Keys

Some Junos OS software features require a license to be activated. To enable each licensed feature, you must purchase, install, manage, and verify a license key that corresponds to the licensed feature.

Release-Tied License Keys and Upgrade Licenses on MX Series Routers

The Junos OS licensing infrastructure currently associates a license feature with attributes such as date, platform, and validity. In addition to these attributes, for MX Series routers running Junos OS Release 12.2 and later, a licensed feature can be associated with a release number at the time of generating the license key. This type of release-tied license key is used to validate a particular licensed feature while attempting a software upgrade. The upgrade process aborts if the release number in the license key is earlier than the Junos OS release number to which the system is being upgraded.

Additionally, an upgrade license key can be generated for a release-tied licensed feature. An upgrade license key is used for carrying forward a capacity license to the upgrade release. Although an upgrade license might be an acceptable license on the current release, it does not add to the existing capacity limit. The capacity added in the upgrade license key is valid for the upgrade software release only.

The release number embedded in the license key indicates the maximum release number up to which Junos OS can be upgraded.

As an example, assume that your system is running Junos OS Release 12.2 and is using the **scale-subscriber** licensed feature with a later release-tied upgrade license key installed. If you request a software upgrade to the later release of Junos OS, the software upgrade operation fails and the following error message is displayed:

```
mgd: error: No valid upgrade license found for feature 'scale-subscriber'.  
Aborting Software upgrade.  
Validation failed
```

In this example, to successfully upgrade to the later release of Junos OS, the release number included in the upgrade license key should be greater than or equal to the later release number. Also, you can perform software upgrades up to the previous release without any additional license keys to retain the existing scale limit.

**NOTE:**

When you install a release-tied license, the following apply:

- You can purchase an upgrade capacity license only if a base capacity license for the same scale-tier has already been generated or purchased.
- You cannot install an upgrade license if the capacity does not match any of the existing base capacity licenses on the system.
- The license installation fails when you install a lower release number license key on a higher software release number.
- A release-tied license can be installed on a Junos OS release number that is lower than or equal to the release number included in the license key. For example, a 12.2 license key is valid on Junos OS Release 12.1.
- An upgrade license is valid only on the target release number specified in the license key, but can be installed on an earlier Junos OS release. For example, a 4 K scale-tier upgrade license for Junos OS Release 12.2 can be installed on an earlier release, and the installed count of licenses remains unaltered.
- Release-tied licenses of the previous release are not deleted on upgrading Junos OS to a newer release version.

Licensable Ports on MX5, MX10, and MX40 Routers

Starting with Junos OS Release 12.2, license keys are available to enhance the port capacity on MX5, MX10, and MX40 routers up to the port capacity of an MX80 router. The MX5, MX10, and MX40 routers are derived from the modular MX80 chassis with similar slot and port assignments, and provide all functionality available on an MX80 router, but at a lower capacity. Restricting port capacity is achieved by making a set of MIC slots and ports licensable. MICs without a license are locked, and are unlocked or made usable by installing appropriate upgrade licenses.

The base capacity of a router is identified by the Ideeprom assembly ID (I2C ID), which defines the board type. However, the Junos OS licensing infrastructure allows the use of restricted ports without a license for a grace period of 30 days. After the grace period expires, the router reverts back to the base capacity if no upgrade license is purchased and installed for the locked ports. The I2C ID along with an upgrade license determine the final capacity of an MX5, MX10, or MX40 router.

The MX5, MX10, MX40, and MX80 routers support the following types of MICs:

- A built-in 10-Gigabit Ethernet MIC with four 10-Gigabit Ethernet ports
- Two front-pluggable MICs

A feature ID is assigned to every license upgrade for enhancing port capacity. [Table 30](#) displays the chassis types and their associated port capacity, I2C ID, base capacity, feature ID, feature name, and the final capacity after a license upgrade.

Table 30: Upgrade Licenses for Enhancing Port Capacity

Chassis Type	Port Capacity	I2C ID	Base Capacity	Feature ID and Feature Name	Upgrade Capacity
MX5	20G	0x556	Slot 1 • 1/MIC0	f1—MX5 to MX10 upgrade	Slot 1 and 2 • 1/MIC0 • 1/MIC1
MX10	40G	0x555	Slot 1 and 2 • 1/MIC0 • 1/MIC1	f2—MX10 to MX40 upgrade	Slot 2 and first 2 ports on Slot 0 • 1/MIC1 • First 2 ports on 0/MIC0
MX40	60G	0x554	Slot 1, Slot 2 and first 2 ports on Slot 0 • 1/MIC0 • 1/MIC1 • First 2 ports on 0/MIC0	f3—MX40 to MX80 upgrade	Slot 2 and all ports on Slot 0 • 1/MIC1 • All 4 ports on 0/MIC0

When installing an upgrade license for enhancing port capacity on MX5, MX10 and MX40 routers, consider the following:

- To upgrade an MX5 router to MX80 router capacity, licenses for all three features (f1, f2, f3) must be installed. All three features can be provided in a single license key.
- To upgrade an MX10 router to MX40 router capacity, installing a license key with f2 feature is sufficient.
- Non-applicable feature IDs in a license key reject the upgrade license. For example:
 - An f1 feature ID on an MX10 upgrade license key rejects the license.
 - Feature IDs f1 and f2 on an MX40 upgrade license key reject the entire license.

Port Activation on MX104 Routers

Starting with Junos OS Release 13.3, license keys are available to activate the ports on the MX104 router. MX104 routers have four built-in ports. By default, in the absence of valid licenses, all four built-in ports are deactivated. By installing licenses, you can activate any two of the four or all of the four built-in ports. For instance, you can install a license to activate the first two built-in ports (xe-2/0/0 and xe-2/0/1) or you can install a license to activate the next two built-in ports (xe-2/0/2 and xe-2/0/3). You can also install a license to activate all four built-in ports (xe-2/0/0, xe-2/0/1, xe-2/0/2, and xe-2/0/3). If you have already activated two of the built-in ports, you can install an additional license to activate the other two built-in ports on the MX104 router.

A feature ID is assigned to every license for activating the built-in ports on the MX104 router. The port license model with the feature ID is described in [Table 31](#).

Table 31: Port Activation License Model for MX104 Routers

Feature ID	Feature Name	Functionality
F1	MX104 2X10G Port Activate (0 and 1)	Ability to activate first two built-in ports (xe-2/0/0 and xe-2/0/1)
F2	MX104 2X10G Port Activate (2 and 3)	Ability to activate next two built-in ports (xe-2/0/2 and xe-2/0/3)

Both the features are also provided in a single license key for ease of use. To activate all four ports, you must either install the licenses for both the features listed in [Table 31](#) or the single license key for both features. If you install the single license key when feature IDs F1 and F2 are already installed, the license does not get rejected. Also, MX104 routers do not support the graceful license expiry policy. A graceful license expiry policy allows the use of a feature for a certain period of time (usually a grace period of 30 days), and reverts if the license for that feature is not installed after the grace period.

Related Documentation

- [Junos OS Feature Licenses on page 107](#)
- [License Enforcement](#)
- [Software Feature Licenses](#)
- [Verifying Junos OS License Installation on page 120](#)
- [show system license on page 441](#)

Generating License Keys

When you purchase a Junos OS software feature license for a device, you receive an e-mail containing an authorization code for the feature license from Juniper Networks. You can use the authorization code to generate a unique license key (a combination of the authorization code and the device's serial number) for the device, and then add the license key on the device.

Before generating the license keys for a device:

- Purchase the required licenses for the device. See [“Software Features That Require Licenses on the QFX Series” on page 108](#).
- Note down the authorization code in the e-mail you received from Juniper Networks when you purchased the license.
- Determine the serial number of the device. For instructions, see *Locating the Serial Number on a QFX3500 Device or Component*.

To generate the license keys for a device:



NOTE: This procedure shows you how to generate license keys on a QFX Series device, but you can follow the same procedure for any device.

1. In a browser, log in to the Juniper Networks License Management System at <https://www.juniper.net/lcrs/license.do>.

The Manage Product Licenses page appears.



NOTE: To access the licensing site, you must have a service contract with Juniper Networks and an access account. If you need help obtaining an account, complete the registration form at the Juniper Networks website <https://www.juniper.net/registration/Register.jsp>.

2. On the Generate Licenses tab, select **QFX Series Product** from the drop-down list, and click **Go**.

The Generate Licenses - QFX Series Product page appears.

3. Select the **QFX Series Product Device** option button, and click **Continue**.

The Generate Licenses - QFX Series Product Devices page appears.

4. In the **Device Serial Number** field, enter the serial number for the device.
5. In the **Authorization Code** field, enter the authorization code in the e-mail you received from Juniper Networks when you purchased the license.
6. (Optional) If you want to enter another authorization code for the same device, click **Enter More Authorization Codes** to display a new authorization code field. Enter the authorization code in this field.
7. Click **Confirm**.

The Confirm License Information page appears, displaying a summary of the information you submitted to the License Management System.

8. Review the information to ensure everything is correct and then click **Generate License**.

The Generate Licenses - QFX Series Product Devices page appears, displaying a summary of your license keys, including a link that displays the details of your new license keys.

9. Select the file format in which you want to obtain your new license keys.
10. Select the delivery method you want to use to obtain your new license keys.

To download the license keys:

- Select the **Download to this computer** option button, and click **OK**.

To e-mail the license keys:

- Select the **Send e-mail to e-mail ID** option button, and click **OK**.

Related Documentation

- [Software Features That Require Licenses on the QFX Series on page 108](#)
- [Adding New Licenses \(CLI Procedure\) on page 117](#)
- [Locating the Serial Number on a QFX3500 Device or Component](#)

Adding New Licenses (CLI Procedure)

Before adding new licenses, complete the following tasks:

- Purchase the required licenses.
- Establish basic network connectivity with the router or switch. For instructions on establishing basic connectivity, see the *Getting Started Guide* or *Quick Start Guide* for your device.



NOTE: On QFabric systems, install your licenses in the default partition of the QFabric system and not on the individual components (Node devices and Interconnect devices).

To add a new license key to the device using the CLI:

1. From the CLI operational mode, enter one of the following CLI commands:

- To add a license key from a file or URL, enter the following command, specifying the filename or the URL where the key is located:

```
user@host> request system license add filename | url
```

- To add a license key from the terminal, enter the following command:

```
user@host> request system license add terminal
```

2. When prompted, enter the license key, separating multiple license keys with a blank line.

If the license key you enter is invalid, an error appears in the CLI output when you press Ctrl+d to exit license entry mode.

3. Go on to [“Verifying Junos OS License Installation” on page 120](#).

On routers that have graceful Routing Engine switchover (GRES) enabled, after successfully adding the new license on the master Routing Engine, the license keys are automatically synchronized on the backup Routing Engine as well. However, in case GRES is not enabled, the new license is added on each Routing Engine separately. This ensures

that the license key is enabled on the backup Routing Engine during changeover of mastership between the Routing Engines.

To add a new license key to a router with dual Routing Engines without GRES:

1. After adding the new license key on the master Routing Engine, use the **request chassis routing-engine master switch** command to have the backup Routing Engine become the master Routing Engine.
2. Log in to the active Routing Engine and add the new license key, repeat the same step.



NOTE: Adding a license key to the router or switch might be delayed if a kernel resynchronization operation is in progress at that time. The following message is displayed on the CLI when the license-adding operation is about to be delayed:

A kernel re-sync operation is in progress. License update may take several minutes to complete.

**Related
Documentation**

- [Deleting a License \(CLI Procedure\) on page 118](#)
- [Junos OS Feature Licenses on page 107](#)
- [Verifying Junos OS License Installation on page 120](#)
- [request system license add on page 438](#)

Deleting a License (CLI Procedure)

Before deleting a license, establish basic network connectivity with the router or switch. For instructions on establishing basic connectivity, see the *Getting Started Guide* or *Quick Start Guide* for your router or switch.

You have the options to delete a single license, delete all licenses, or delete a list of licenses enclosed in brackets.

1. Display the licenses available to be deleted.

```
user@host> request system license delete license-identifier-list ?
```

```
Possible completions:
```

E00468XXX4	License key identifier
JUNOS10XXX1	License key identifier
JUNOS10XXX2	License key identifier
JUNOS10XXX3	License key identifier
JUNOS10XXX4	License key identifier
[Open a set of values

2. To delete a license key or keys from a device using the CLI operational mode, select one of the following methods:

- Delete a single license by specifying the license ID. Using this option, you can delete only one license at a time.

```
user@host> request system license delete license-identifier
```

- Delete all license keys from the current device.

```
user@host> request system license delete all
```

- Delete multiple license keys from the current device. Specify the license identifier for each key and enclose the list of identifiers in brackets.

```
user@host> request system license delete license-identifier-list [JUNOS10XXX1
JUNOS10XXX3 JUNOS10XXX4 ...]
```

```
Delete license(s) ?
[yes,no] (no) yes
```

3. Go on to [“Verifying Junos OS License Installation” on page 120](#).



NOTE: Deleting a license key from the router or switch might be delayed if a kernel resynchronization operation is in progress at that time. The following message is displayed on the CLI when the license-deleting operation is about to be delayed:

A kernel re-sync operation is in progress. License update may take several minutes to complete.

Related Documentation

- [Adding New Licenses \(CLI Procedure\) on page 117](#)
- [Saving License Keys on page 119](#)
- [Junos OS Feature Licenses on page 107](#)
- [Verifying Junos OS License Installation on page 120](#)
- [request system license delete on page 439](#)

Saving License Keys

Before saving a license, establish basic network connectivity with the router or switch. For instructions on establishing basic connectivity, see the *Getting Started Guide* or *Quick Start Guide* for your router or switch.

To save the licenses installed on a device to a file using the CLI:

1. From the CLI operational mode, enter one of the following CLI commands:

- To save the installed license keys to a file or URL, enter the following command:

```
user@host> request system license save filename | url
```

For example, the following command saves the installed license keys to a file named `license.config`:

- To save a license key from the terminal, enter the following command:

```
user@host> request system license save ftp://user@host/license.config
```

2. Go on to [“Verifying Junos OS License Installation” on page 120](#).

- Related Documentation**
- [Adding New Licenses \(CLI Procedure\) on page 117](#)
 - [Deleting a License \(CLI Procedure\) on page 118](#)
 - [Junos OS Feature Licenses on page 107](#)
 - [Verifying Junos OS License Installation on page 120](#)

Verifying Junos OS License Installation

To verify Junos OS license management, perform the following tasks:

- [Displaying Installed Licenses on page 120](#)
- [Displaying License Usage on page 121](#)

Displaying Installed Licenses

Purpose Verify that the expected licenses are installed and active on the router or switch.

Action From the CLI, enter the **show system license** command.

Sample Output

```
user@host> show system license
License usage:

```

Feature name	Licenses used	Licenses installed	Licenses needed	Expiry
subscriber-acct	0	1	0	permanent
subscriber-auth	0	1	0	permanent
subscriber-addr	0	1	0	permanent
subscriber-vlan	0	1	0	permanent
subscriber-ip	0	1	0	permanent
scale-subscriber	0	1000	0	permanent
scale-l2tp	0	1000	0	permanent
scale-mobile-ip	0	1000	0	permanent

```

Licenses installed:
License identifier: E000185416
License version: 2
Features:
subscriber-acct - Per Subscriber Radius Accounting
permanent
subscriber-auth - Per Subscriber Radius Authentication
permanent
subscriber-addr - Address Pool Assignment
permanent
subscriber-vlan - Dynamic Auto-sensed Vlan
permanent
subscriber-ip - Dynamic and Static IP
permanent
```

Meaning The output shows a list of the license usage and a list of the licenses installed on the router or switch. Verify the following information:

- Each license is present. Licenses are listed in ascending alphanumeric order by license ID.
- The state of each license is **permanent**.



NOTE: A state of invalid indicates that the license key is not a valid license key. Either it was entered incorrectly or it is not valid for the specific device.

- The feature for each license is the expected feature. The features enabled are listed by license. An all-inclusive license has all features listed.
- All configured features have the required licenses installed. The Licenses needed column must show that no licenses are required.

Displaying License Usage

Purpose Verify that the licenses fully cover the feature configuration on the router or switch.

Action From the CLI, enter the **show system license usage** command.

Sample Output

```
user@host> show system license usage
```

	Licenses used	Licenses installed	Licenses needed	Expiry
Feature name				
subscriber-addr	1	0	1	29 days
scale-subscriber	0	1000	0	permanent
scale-l2tp	0	1000	0	permanent
scale-mobile-ip	0	1000	0	permanent

Meaning The output shows any licenses installed on the router or switch and how they are used. Verify the following information:

- Any configured licenses appear in the output. The output lists features in ascending alphabetical order by license name. The number of licenses appears in the third column. Verify that you have installed the appropriate number of licenses.
- The number of licenses used matches the number of configured features. If a licensed feature is configured, the feature is considered used. The sample output shows that the subscriber address pooling feature is configured.
- A license is installed on the router or switch for each configured feature. For every feature configured that does not have a license, one license is needed.

For example, the sample output shows that the subscriber address feature is configured but that the license for the feature has not yet been installed. The license must be installed within the remaining grace period to be in compliance.

PART 8

Login Classes

- [Understanding Login Classes on page 125](#)

CHAPTER 8

Understanding Login Classes

- [Configuring Junos OS User Accounts on page 125](#)
- [Configuring the Junos OS to Display a System Login Announcement on page 126](#)
- [Configuring Junos OS to Display a System Login Message on page 126](#)

Configuring Junos OS User Accounts

User accounts provide one way for users to access the router or switch. For each account, you define the login name for the user and, optionally, information that identifies the user. After you have created an account, the software creates a home directory for the user.

To create user accounts, include the **user** statement at the **[edit system login]** hierarchy level:

```
[edit system login]
user username {
  class class-name;
  class {
    (encrypted-password "password" | plain-text-password);
    ssh-rsa "public-key";
    ssh-dsa "public-key";
  }
  full-name complete-name;
  uid uid-value;
  class class-name;
}
```

Related Documentation

- [Example: Configuring User Accounts on page 1830](#)
- [Configuring a Local Administrator Account on page 1853](#)
- [Junos OS User Accounts Overview on page 1821](#)
- [*Limiting the Number of User Login Attempts for SSH and Telnet Sessions*](#)

Configuring the Junos OS to Display a System Login Announcement

By default, no login announcement is displayed. To configure a system login announcement, include the **announcement** statement at the **[edit system login]** hierarchy level:

```
[edit system login]
announcement text;
```

If the announcement text contains any spaces, enclose the text in quotation marks.

A system login *announcement* appears after the user logs in. A system login *message* appears before the user logs in.



TIP: You can use the same special characters described to format your system login announcement.

Related Documentation

- *Defining Junos OS Login Classes*
- *Configuring the Junos OS to Display a System Login Message*

Configuring Junos OS to Display a System Login Message

By default, no login message is displayed on the router or switch. To configure a system login message, include the **message** statement at the **[edit system login]** hierarchy level:

```
[edit system login]
message text;
```

If the message text contains any spaces, enclose it in quotation marks.

You can format the message using the following special characters:

- \n—New line
- \t—Horizontal tab
- \'—Single quotation mark
- \"—Double quotation mark
- \\—Backslash

The following is a sample login message configuration:

```
[edit]
system {
  login {
    message "\n\n\n\tUNAUTHORIZED USE OF THIS SYSTEM\n\tIS STRICTLY PROHIBITED!\n\n\tPlease contact\n\t'company-noc@company.com\t' to gain\n\taccess to this equipment if you need authorization.\n\n\n";
  }
}
```

```
    }  
}
```

The preceding login message configuration example produces a login message similar to the following:

```
server% telnet router1  
Trying 1.1.1.1...  
Connected to router1.  
Escape character is '^['.
```

```
UNAUTHORIZED USE OF THIS SYSTEM  
IS STRICTLY PROHIBITED!
```

```
Please contact 'company-noc@company.com' to gain  
access to this equipment if you need authorization.
```

```
router1 (ttyp0)
```

```
login:
```

A system login message appears before the user logs in. A system login announcement appears after the user logs in.

- Related Documentation**
- *Defining Junos OS Login Classes*
 - [message on page 281](#)

PART 9

Network Time Protocol (NTP)

- [Understanding NTP on page 131](#)

CHAPTER 9

Understanding NTP

- [Understanding NTP Time Servers on page 131](#)
- [Configuring NTP Authentication Keys on page 132](#)
- [Configuring the NTP Time Server and Time Services on page 133](#)
- [Configuring the Switch to Listen for Broadcast Messages Using NTP on page 135](#)
- [Configuring the Switch to Listen for Multicast Messages Using NTP on page 136](#)
- [Setting the Date and Time on page 136](#)
- [Synchronizing and Coordinating Time Distribution Using NTP on page 137](#)
- [Example: Configuring NTP on page 139](#)
- [Example: Configuring NTP as a Single Time Source for Router and Switch Clock Synchronization on page 142](#)

Understanding NTP Time Servers

The IETF defined the Network Time Protocol (NTP) to synchronize the clocks of computer systems connected to each other over a network. Most large networks have an NTP server that ensures that time on all devices is synchronized, regardless of the device location. If you use one or more NTP servers on your network, ensure you include the NTS server addresses in your Junos OS configuration.

When configuring the NTP, you can specify which system on the network is the authoritative time source, or time server, and how time is synchronized between systems on the network. To do this, you configure the router, switch, or security device to operate in one of the following modes:

- **Client mode**—In this mode, the local router or switch can be synchronized with the remote system, but the remote system can never be synchronized with the local router or switch.
- **Symmetric active mode**—In this mode, the local router or switch and the remote system can synchronize with each other. You use this mode in a network in which either the local router or switch or the remote system might be a better source of time.



NOTE: Symmetric active mode can be initiated by either the local or the remote system. Only one system needs to be configured to do so. This means that the local system can synchronize with any system that offers symmetric active mode without any configuration whatsoever. However, we strongly encourage you to configure authentication to ensure that the local system synchronizes only with known time servers.

- Broadcast mode—In this mode, the local router or switch sends periodic broadcast messages to a client population at the specified broadcast or multicast address. Normally, you include this statement only when the local router or switch is operating as a transmitter.
- Server mode—In this mode, the local router or switch operates as an NTP server.



NOTE: In NTP server mode, the Junos OS supports authentication as follows:

- If the NTP request from the client comes with an authentication key (such as a key ID and message digest sent with the packet), the request is processed and answered based on the authentication key match.
- If the NTP request from the client comes without any authentication key, the request is processed and answered without authentication.

**Related
Documentation**

- [Example: Configuring NTP as a Single Time Source for Router and Switch Clock Synchronization on page 142](#)

Configuring NTP Authentication Keys

Time synchronization can be authenticated to ensure that the switch obtains its time services only from known sources. By default, network time synchronization is unauthenticated. The switch will synchronize to whatever system appears to have the most accurate time. We strongly encourage you to configure authentication of network time services.

To authenticate other time servers, include the **trusted-key** statement at the **[edit system ntp]** hierarchy level. Only time servers that transmit network time packets containing one of the specified key numbers are eligible to be synchronized. Additionally, the key needs to match the value configured for that key number. Other systems can synchronize to the local switch without being authenticated.

```
[edit system ntp]  
trusted-key [ key-numbers ];
```

Each key can be any 32-bit unsigned integer except 0. Include the **key** option in the **peer**, **server**, or **broadcast** statements to transmit the specified authentication key when

transmitting packets. The key is necessary if the remote system has authentication enabled so that it can synchronize to the local system.

To define the authentication keys, include the **authentication-key** statement at the **[edit system ntp]** hierarchy level:

```
[edit system ntp]
authentication-key key-number type type value password;
```

number is the key number, **type** is the authentication type (only Message Digest 5 [MD5] is supported), and **password** is the password for this key. The key number, type, and password must match on all systems using that particular key for authentication.

Related Documentation

- [Understanding NTP Time Servers on page 131](#)
- [Example: Configuring NTP as a Single Time Source for Router and Switch Clock Synchronization on page 142](#)
- [trusted-key on page 292](#)
- *authentication-key*

Configuring the NTP Time Server and Time Services

When you use NTP, configure the switch to operate in one of the following modes:

- Client mode
- Symmetric active mode
- Broadcast mode
- Server mode

The following topics describe how to configure these modes of operation:

1. [Configuring the Switch to Operate in Client Mode on page 133](#)
2. [Configuring the Router or Switch to Operate in Symmetric Active Mode on page 134](#)
3. [Configuring the Router or Switch to Operate in Broadcast Mode on page 134](#)
4. [Configuring the Router or Switch to Operate in Server Mode on page 135](#)

Configuring the Switch to Operate in Client Mode

To configure the local router or switch to operate in client mode, include the **server** statement and other optional statements at the **[edit system ntp]** hierarchy level:

```
[edit system ntp]
server address <key key-number> <version value> <prefer>;
authentication-key key-number type type value password;
boot-server address;
trusted-key [ key-numbers ];
```

Specify the address of the system acting as the time server. You must specify an address, not a hostname.

To include an authentication key in all messages sent to the time server, include the **key** option. The key corresponds to the key number you specify in the **authentication-key** statement, as described in .

By default, the router or switch sends NTP version 4 packets to the time server. To set the NTP version level to 1, 2, or 3, include the **version** option.

If you configure more than one time server, you can mark one server preferred by including the **prefer** option.

The following example shows how to configure the router or switch to operate in client mode:

```
[edit system ntp]
authentication-key 1 type md5 value "$9$EgfcvX7VY4ZEcwgoHjkP5Q3CuREyv87";
boot-server 10.1.1.1;
server 10.1.1.1 key 1 prefer;
trusted-key 1;
```

Configuring the Router or Switch to Operate in Symmetric Active Mode

To configure the local router or switch to operate in symmetric active mode, include the **peer** statement at the **[edit system ntp]** hierarchy level:

```
[edit system ntp]
peer address <key key-number> <version value> <prefer>;
```

Specify the address of the remote system. You must specify an address, not a hostname.

To include an authentication key in all messages sent to the remote system, include the **key** option. The key corresponds to the key number you specify in the **authentication-key** statement.

By default, the router or switch sends NTP version 4 packets to the remote system. To set the NTP version level to 1, 2 or 3, include the **version** option.

If you configure more than one remote system, you can mark one system preferred by including the **prefer** option:

```
peer address <key key-number> <version value> prefer;
```

Configuring the Router or Switch to Operate in Broadcast Mode

To configure the local router or switch to operate in broadcast mode, include the **broadcast** statement at the **[edit system ntp]** hierarchy level:

```
[edit system ntp]
broadcast address <key key-number> <version value> <ttl value>;
```

Specify the broadcast address on one of the local networks or a multicast address assigned to NTP. You must specify an address, not a hostname. If the multicast address is used, it must be **224.0.1.1**.

To include an authentication key in all messages sent to the remote system, include the **key** option. The key corresponds to the key number you specify in the **authentication-key** statement.

By default, the router or switch sends NTP version 4 packets to the remote system. To set the NTP version level to 1, 2, or 3, include the **version** option.

Configuring the Router or Switch to Operate in Server Mode

In server mode, the router or switch acts as an NTP server for clients when the clients are configured appropriately. The only prerequisite for “server mode” is that the router or switch must be receiving time from another NTP peer or server. No other configuration is necessary on the router or switch.

To configure the local router or switch to operate as an NTP server, include the following statements at the **[edit system ntp]** hierarchy level:

```
[edit system ntp]
authentication-key key-number type type value password;
server address <key key-number> <version value> <prefer>;
trusted-key [ key-numbers ];
```

Specify the address of the system acting as the time server. You must specify an address, not a hostname.

To include an authentication key in all messages sent to the time server, include the **key** option. The key corresponds to the key number you specify in the **authentication-key** statement.

By default, the router or switch sends NTP version 4 packets to the time server. To set the NTP version level to 1, or 2, or 3, include the **version** option.

If you configure more than one time server, you can mark one server preferred by including the **prefer** option.

The following example shows how to configure the router or switch to operate in server mode:

```
[edit system ntp]
authentication-key 1 type md5 value "$9$txEruBEreWx-wtuLNdboaUjH.T3AtOESe";
server 172.17.27.46 prefer;
trusted-key 1;
```

Related Documentation

- [Understanding NTP Time Servers on page 131](#)
- [Example: Configuring NTP as a Single Time Source for Router and Switch Clock Synchronization on page 142](#)

Configuring the Switch to Listen for Broadcast Messages Using NTP

When you are using NTP, you can configure the local switch to listen for broadcast messages on the local network to discover other servers on the same subnet by including the **broadcast-client** statement at the **[edit system ntp]** hierarchy level:

```
[edit system ntp]  
broadcast-client;
```

When the switch detects a broadcast message for the first time, it measures the nominal network delay using a brief client-server exchange with the remote server. It then enters *broadcast client* mode, in which it listens for, and synchronizes to, succeeding broadcast messages.

To avoid accidental or malicious disruption in this mode, both the local and remote systems must use authentication and the same trusted key and key identifier.

**Related
Documentation**

- [Configuring the Switch to Listen for Multicast Messages Using NTP on page 136](#)
- [Configuring the NTP Time Server and Time Services on page 133](#)
- [Example: Configuring NTP as a Single Time Source for Router and Switch Clock Synchronization on page 142](#)

Configuring the Switch to Listen for Multicast Messages Using NTP

When you are using NTP, you can configure the local switch to listen for multicast messages on the local network to discover other servers on the same subnet by including the **multicast-client** statement at the **[edit system ntp]** hierarchy level:

```
[edit system ntp]  
multicast-client <address>;
```

When the switch receives a multicast message for the first time, it measures the nominal network delay using a brief client-server exchange with the remote server. It then enters *multicast client* mode, in which it listens for, and synchronizes to, succeeding multicast messages.

You can specify one or more IP addresses. (You must specify an address, not a hostname.) If you do, the router or switch joins those multicast groups. If you do not specify any addresses, the software uses **224.0.0.1**.

To avoid accidental or malicious disruption in this mode, both the local and remote systems must use authentication and the same trusted key and key identifier.

**Related
Documentation**

- [Configuring the Switch to Listen for Broadcast Messages Using NTP on page 135](#)
- [Configuring the NTP Time Server and Time Services on page 133](#)
- [Example: Configuring NTP as a Single Time Source for Router and Switch Clock Synchronization on page 142](#)

Setting the Date and Time

1. Enter operational mode in the CLI.
2. Enter the following command:

```
user@switch> set date YYYYMMDDHHMM.ss source-address
```


For example, the following command sets the date and time.

```
user@switch# set date 201102151010.55
```

3. To set the date and time from an NTP server, enter the following command:

```
user@switch# set date ntp servers
```

For example, the following command sets the date and time from an NTP server:

```
user@switch# set date ntp 200.40.40.1
```

4. To set the date and time from more than one NTP server, enter the same command:

```
user@switch# set date ntp servers
```

For example, the following command sets the date and time from more than one NTP server:

```
user@switch# set date ntp 200.40.40.1 200.40.40.2
```

Related Documentation

- [set date](#)

Synchronizing and Coordinating Time Distribution Using NTP

Using NTP to synchronize and coordinate time distribution in a large network involves these tasks:

1. [Configuring NTP on page 137](#)
2. [Configuring the NTP Boot Server on page 137](#)
3. [Specifying a Source Address for an NTP Server on page 138](#)

Configuring NTP

- To configure NTP on the switch, include the **ntp** statement at the **[edit system]** hierarchy level:

```
[edit system]
ntp {
  authentication-key number type type value password;
  boot-server (address | hostname);
  broadcast <address> <key key-number> <version value> <ttl value>;
  broadcast-client;
  multicast-client <address>;
  peer address <key key-number> <version value> <prefer>;
  server address <key key-number> <version value> <prefer>;
  source-address source-address;
  trusted-key [ key-numbers ];
}
```

Configuring the NTP Boot Server

When you boot the switch, it issues an **ntpdate** request, which polls a network server to determine the local date and time. You need to configure a server that the switch uses to determine the time when the switch boots. Otherwise, NTP will not be able to

synchronize to a time server if the server's time appears to be very far off of the local switch's time.

- To configure the NTP boot server, include the **boot-server** statement at the **[edit system ntp]** hierarchy level:

```
[edit system ntp]
boot-server (address | hostname);
```

Specify either the IP address or the hostname of the network server.

Specifying a Source Address for an NTP Server

For IP version 4 (IPv4), you can specify that if the NTP server configured at the **[edit system ntp]** hierarchy level is contacted on one of the loopback interface addresses, the reply always uses a specific source address. This is useful for controlling which source address NTP uses to access your network when it is either responding to or sending an NTP client request from your network.

To configure the specific source address that the reply will always use, and the source address that requests initiated by NTP server will use, include the **source-address** statement at the **[edit system ntp]** hierarchy level:

```
[edit system ntp]
source-address source-address;
```

source-address is a valid IP address configured on one of the switch interfaces.



NOTE: If a firewall filter is applied on the loopback interface, ensure that the source address specified for the NTP server at the **[edit system ntp]** hierarchy level is explicitly included as one of the match criteria in the firewall filter. This enables the Junos OS to accept traffic on the loopback interface from the specified source address.

The following example shows a firewall filter with the source address 10.0.10.100 specified in the **from** statement included at the **[edit firewall filter firewall-filter-name]** hierarchy:

```
[edit firewall filter Loopback-Interface-Firewall-Filter]
term Allow-NTP {
  from {
    source-address {
      172.17.27.46/32; // IP address of the NTP server
      10.0.10.100/32; // Source address specified for the NTP server
    }
  }
  then accept;
}
```

If no source address is configured for the NTP server, include the primary address of the loopback interface in the firewall filter.

- Related Documentation**
- [Understanding NTP Time Servers on page 131](#)
 - [Example: Configuring NTP as a Single Time Source for Router and Switch Clock Synchronization on page 142](#)

Example: Configuring NTP

The Network Time Protocol (NTP) provides the mechanisms to synchronize time and coordinate time distribution in a large, diverse network. NTP uses a returnable-time design in which a distributed subnet of time servers operating in a self-organizing, hierarchical primary-secondary configuration synchronizes local clocks within the subnet and to national time standards by means of wire or radio. The servers also can redistribute reference time using local routing algorithms and time daemons.

This example shows how to configure NTP:

- [Requirements on page 139](#)
- [Overview on page 139](#)
- [Configuration on page 139](#)
- [Verification on page 140](#)

Requirements

This example uses the following software and hardware components:

- Junos OS Release 11.1 or later
- A switch connected to a network on which an NTP boot server and NTP server reside

Overview

Debugging and troubleshooting are much easier when the timestamps in the log files of all switches are synchronized, because events that span a network can be correlated with synchronous entries in multiple logs. We recommend using the Network Time Protocol (NTP) to synchronize the system clocks of your switch and other network equipment.

In this example, an administrator wants to synchronize the time in a switch to a single time source. We recommend using authentication to make sure that the NTP peer is trusted. The **boot-server** statement identifies the server from which the initial time of day and date are obtained when the switch boots. The **server** statement identifies the NTP server used for periodic time synchronization. The **authentication-key** statement specifies that an HMAC-Message Digest 5 (MD5) scheme is used to hash the key value for authentication, which prevents the switch from synchronizing with an attacker's host that is posing as the time server.

Configuration

To configure NTP:

CLI Quick Configuration	<p>To quickly configure NTP, copy the following commands and paste them into the switch's terminal window:</p> <pre>[edit system] set ntp boot-server 10.1.4.1 set ntp server 10.1.4.2 set ntp authentication-key 2 type md5 value "\$9\$aHlj8"</pre>
Step-by-Step Procedure	<p>To configure NTP :</p> <ol style="list-style-type: none">1. Specify the boot server: <pre>[edit system] user@switch# set ntp boot-server 10.1.4.1</pre>2. Specify the NTP server: <pre>[edit system] user@switch# set ntp server 10.1.4.2</pre>3. Specify the key number, authentication type (MD5), and key for authentication: <pre>[edit system] user@switch# set ntp authentication-key 2 type md5 value "\$9\$aHlj8"</pre>
Results	<p>Check the results:</p> <pre>[edit system] user@switch# show ntp { boot-server 10.1.4.1; authentication-key 2 type md5 value "\$9\$aHlj8"; ## SECRET-DATA server 10.1.4.2; }</pre>

Verification

To confirm that the configuration is correct, perform these tasks:

- [Checking the Time on page 140](#)
- [Displaying the NTP Peers on page 141](#)
- [Displaying the NTP Status on page 141](#)

Checking the Time

Purpose	Check the time that has been set on the switch.
Action	<p>Enter the show system uptime operational mode command to display the time.</p> <pre>user@switch> show system uptime fpc0: ----- Current time: 2009-06-12 12:49:03 PDT System booted: 2009-05-15 06:24:43 PDT (4w0d 06:24 ago) Protocols started: 2009-05-15 06:27:08 PDT (4w0d 06:21 ago) Last configured: 2009-05-27 14:57:03 PDT (2w1d 21:52 ago) by admin1 12:49PM up 28 days, 6:24, 1 user, load averages: 0.05, 0.06, 0.01</pre>

Meaning The output shows that the current date and time are June 12, 2009 and 12:49:03 PDT. The switch booted 4 weeks, 6 hours, and 24 minutes ago, and its protocols were started approximately 3 minutes before it booted. The switch was last configured by user **admin1** on May 27, 2009, and there is currently one user logged in to the switch.

The output also shows that the load average is 0.05 seconds for the last minute, 0.06 seconds for the last 5 minutes, and 0.01 seconds for the last 15 minutes.

Displaying the NTP Peers

Purpose Verify that the time has been obtained from an NTP server.

Action Enter the **show ntp associations** operational mode command to display the NTP server from switch obtained its time.

```
user@switch> show ntp associations
      remote      refid      st t when poll reach  delay  offset  jitter
=====
*ntp5.domain1.ne .GPS.          1 u  414 1024  377   3.435   4.002   0.765
```

Meaning The asterisk (*) in front of the NTP server name, or peer, indicates that the time is synchronized and obtained from this server. The delay, offset, and jitter are displayed in milliseconds.

Displaying the NTP Status

Purpose View the configuration of the NTP server and the status of the system.

Action Enter the **show ntp status** operational mode command to view the status of the NTP.

```
user@switch> show ntp status
status=0644 leap_none, sync_ntp, 4 events, event_peer/strat_chg,
version="ntpd 4.2.0-a Mon Apr 13 19:09:05 UTC 2009 (1)",
processor="powerpc", system="JUNOS9.5R1.8", leap=00, stratum=2,
precision=-18, rootdelay=2.805, rootdispersion=42.018, peer=48172,
refid=172.17.28.5,
reftime=cddd397a.60e6d7bf Fri, Jun 12 2009 13:30:50.378, poll=10,
clock=cddd3b1b.ec5a2bb4 Fri, Jun 12 2009 13:37:47.923, state=4,
offset=3.706, frequency=-23.018, jitter=1.818, stability=0.303
```

Meaning The output shows status information about the switch and the NTP.

- Related Documentation**
- [Understanding NTP Time Servers on page 131](#)
 - [ntp on page 288](#)
 - [Configuring the NTP Time Server and Time Services on page 133](#)
 - [CLI Explorer](#)
 - [Junos OS Baseline Network Operations Guide](#)

Example: Configuring NTP as a Single Time Source for Router and Switch Clock Synchronization

Debugging and troubleshooting are much easier when the timestamps in the log files of all the routers or switches are synchronized, because events that span the network can be correlated with synchronous entries in multiple logs. We strongly recommend using the Network Time Protocol (NTP) to synchronize the system clocks of routers, switches, and other network equipment.

By default, NTP operates in an entirely unauthenticated manner. If a malicious attempt to influence the accuracy of a router or switch's clock succeeds, it could have negative effects on system logging, make troubleshooting and intrusion detection more difficult, and impede other management functions.

The following sample configuration synchronizes all the routers or switches in the network to a single time source. We recommend using authentication to make sure that the NTP peer is trusted. The **boot-server** statement identifies the server from which the initial time of day and date is obtained when the router boots. The **server** statement identifies the NTP server used for periodic time synchronization. The **authentication-key** statement specifies that an HMAC-Message Digest 5 (MD5) scheme should be used to hash the key value for authentication, which prevents the router or switch from synchronizing with an attacker's host posing as the time server.

```
[edit]
system {
  ntp {
    authentication-key 2 type md5 value "$9$aHlj8gqQ1gJyJghgJgiiii"; # SECRET-DATA
    boot-server 10.1.4.1;
    server 10.1.4.2;
  }
}
```

Related Documentation

- [NTP Overview](#)
- [Understanding NTP Time Servers on page 131](#)
- [show ntp associations on page 450](#)
- [show ntp status on page 452](#)

PART 10

Passwords

- [Understanding Passwords on page 145](#)

CHAPTER 10

Understanding Passwords

- [Configuring the Root Password on page 145](#)
- [Recovering the Root Password on page 147](#)
- [Example: Changing the Requirements for Junos OS Plain-Text Passwords on page 148](#)

Configuring the Root Password

The Junos OS is preinstalled on the router or switch. When the router or switch is powered on, it is ready to be configured. Initially, you log in as the user **root** with no password. The root directory of a UNIX device is the entry point to all other folders and files on that device. As a result, access to the root directory is restricted by default to a predefined user account known as the *root user*. The root user (also referred to as *superuser*) has unrestricted access and full permissions within the system. The expression “log in as root” is commonly used when an action requires the user to log into the device as the root user.



NOTE: If you configure a blank password using the **encrypted-password** statement at the **[edit system root-authentication]** hierarchy level for root authentication, you can commit a configuration but you *cannot* log in as the root user and gain root level access to the router or switch.

After you log in, you should configure the root (superuser) password by including the **root-authentication** statement at the **[edit system]** hierarchy level and configuring one of the password options:

```
[edit system]
root-authentication {
  (encrypted-password "password"| plain-text-password);
  load-key-file URL filename;
  ssh-dsa "public-key" <from hostname>;
  ssh-ecdsa "public-key" <from hostname>;
  ssh-rsa "public-key" <from hostname>;
}
```

If you configure the **plain-text-password** option, you are prompted to enter and confirm the password:

```
[edit system]
```

```
user@host# set root-authentication plain-text-password
New password: type password here
Retype new password: retype password here
```

The default requirements for plain-text passwords are:

- The password must be between 6 and 128 characters long
 - You can include most character classes in a password (uppercase letters, lowercase letters, numbers, punctuation marks, and other special characters). Control characters are not recommended.
 - Valid passwords must contain at least one change of case or character class.

You can use the **load-key-file** *URL filename* statement to load an SSH key file that was previously generated using **ssh-keygen**. The *URL filename* is the path to the file's location and name. When using this option, the contents of the key file are copied into the configuration immediately after entering the **load-key-file** *URL* statement. This command loads RSA (SSH version 1 and SSH version 2) and DSA (SSH version 2) public keys.

Optionally, you can use the **ssh-dsa**, **ssh-ecdsa**, or **ssh-rsa** statements to directly configure SSH RSA, DSA, or ECDSA keys to authenticate root logins. You can configure more than one public key for SSH authentication of root logins as well as for user accounts. When a user logs in as root, the public keys are referenced to determine whether the private key matches any of them.

To view the SSH keys entries, use the configuration mode **show** command. For example:

```
[edit system]
user@host# set root-authentication load-key-file my-host::ssh/id_dsa.pub
.file.19692 | 0 KB | 0.3 kB/s | ETA: 00:00:00 | 100%
[edit system]
user@host# show
root-authentication {
  ssh-rsa "1024 35 9727638204084251055468226757249864241630322
20740496252839038203869014158453496417001961060835872296
15634757491827360336127644187426594689320773910834481012
68312595772262546166799927831612350043866091586628382248
97467326056611921489539813965561563786211940327687806538
16960202749164163735913269396344008443 boojum@juniper.net"; #
  SECRET-DATA
}
```

Junos-FIPS software has special password requirements. FIPS passwords must be between 10 and 20 characters in length. Passwords must use at least three of the five defined character sets (uppercase letters, lowercase letters, digits, punctuation marks, and other special characters). If Junos-FIPS is installed on the router or switch, you cannot configure passwords unless they meet this standard. If you use the **encrypted-password** option, then a null-password (empty) is not permitted.

You cannot configure a blank password for **encrypted-password** using blank quotation marks (" "). You must configure a password whose number of characters range from 1 through 128 characters and enclose the password in quotation marks.

Related Documentation

- *Protecting Network Security by Configuring the Root Password*
- [Example: Configuring a Plain-Text Password for Root Logins on page 1812](#)
- [Example: Configuring SSH Authentication for Root Logins on page 1814](#)
- [Example: Changing the Requirements for Junos OS Plain-Text Passwords on page 148](#)
- *Recovering the Root Password*

Recovering the Root Password

If you forget the root password, you can use the password recovery procedure to reset the root password.



NOTE: The root password cannot be recovered on a QFabric system.



NOTE: You need console access to the switch to recover the root password.

To recover the root password:

1. Power off the switch by switching off the AC power outlet of the device or, if necessary, by pulling the power cords out of the device's power supplies.
2. Turn off the power to the management device, such as a PC or laptop computer, that you want to use to access the CLI.
3. Plug one end of the Ethernet rollover cable supplied with the device into the RJ-45-to-DB-9 serial port adapter supplied with the device.
4. Plug the RJ-45-to-DB-9 serial port adapter into the serial port on the management device.
5. Connect the other end of the Ethernet rollover cable to the console port on the device.
6. Turn on the power to the management device.
7. On the management device, start your asynchronous terminal emulation application (such as Microsoft Windows Hyperterminal) and select the appropriate **COM** port to use (for example, **COM1**).
8. Configure the port settings as follows:
 - Bits per second: 9600
 - Data bits: 8
 - Parity: None
 - Stop bits: 1
 - Flow control: None

9. Power on the device by (if necessary) plugging the power cords into the device's power supply, or turning on the power to the device by switching on the AC power outlet the device is plugged into.

The terminal emulation screen on your management device displays the device's boot sequence.

10. When the following prompt appears, press the Spacebar to access the device's bootstrap loader command prompt:

```
Hit [Enter] to boot immediately, or space bar for command prompt.  
Booting [kernel] in 9 seconds...
```

11. At the following prompt, enter **boot -s** to start up the system in single-user mode.

```
ok boot -s
```

12. At the following prompt, enter **recovery** to start the root password recovery procedure.

```
Enter full pathname of shell or 'recovery' for root password recovery or RETURN  
for /bin/sh: recovery
```

13. Enter configuration mode in the CLI.

14. Set the root password. For example:

```
user@switch# set system root-authentication plain-text-password
```

15. At the following prompt, enter the new root password. For example:

```
New password: juniper1  
Retype new password:
```

16. At the second prompt, reenter the new root password.

17. After you have finished configuring the password, commit the configuration.

```
root@host# commit  
commit complete
```

18. Exit configuration mode in the CLI.

19. Exit operational mode in the CLI.

20. At the prompt, enter **y** to reboot the device.

```
Reboot the system? [y/n] y
```

Related Documentation

- *Configuring the Root Password*

Example: Changing the Requirements for Junos OS Plain-Text Passwords

This example shows how to set various maximum and minimum requirements for plain-text passwords to increase password strength.

- [Requirements on page 149](#)
- [Overview on page 149](#)
- [Configuration on page 149](#)

Requirements

This example requires a device running Junos 12.2 or greater. The **minimum-length** and **maximum-length** password requirements statements are available in earlier releases, however, you must have Junos OS Release 12.2 or greater to configure **minimum-lower-cases**, **minimum-numeric**s, **minimum-punctuations**, or **minimum-upper-cases**.

Overview

You can use a variety of requirements to strengthen plain-text passwords for greater security. Junos OS provides a number of possible configurations at the **[edit system login password]** hierarchy level that allow you to require users to create plain-text passwords that conform to a particular set of requirements that may include such things as length, number of changes, type of characters, numbers, or letter case.

Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set system login password minimum-length 12
set system login password maximum-length 22
set system login password minimum-numeric 1
set system login password minimum-upper-cases 1
set system login password minimum-lower-cases 1
set system login password minimum-punctuations 1
```

Configuring Requirements for Plain-Text Passwords

Step-by-Step Procedure This example configures password requirements that require the user to create a password that has a minimum length of 12 characters, a maximum length of 22 characters, and that includes at least one lower-case letter, at least one upper-case letter, at least one punctuation character, and at least one numeric character.

1. Navigate to configuration mode in the **[system login password]** hierarchy level.


```
user@host> edit
[edit]
user@host# edit system login password
```
2. Set a minimum length requirement of 12 characters and a maximum length requirement of 22 characters for user passwords.


```
[edit system login password]
user@host# set minimum-length 12
[edit system login password]
user@host# set maximum-length 22
```
3. Require users to set a password that has at least one lower-case letter and at least one upper-case letter.


```
[edit system login password]
```

```
user@host# set minimum-lower-cases 1
[edit system login password]
user@host# set minimum-upper-cases 1
```

4. Require users to set a password that has at least one punctuation-class character and at least one number.

```
[edit system login password]
user@host# set minimum-punctuations 1
[edit system login password]
user@host# set minimum-numeric 1
```

Results

From configuration mode, confirm your configuration by entering the show command at the edit system login password hierarchy level. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit system login password]
user@host# show
minimum-length 12;
maximum-length 22;
minimum-numeric 1;
minimum-upper-cases 1;
minimum-lower-cases 1;
```

Related Documentation

- [Special Requirements for Junos OS Plain-Text Passwords on page 1824](#)
- *password (Login)*

CHAPTER 11

Precision Time Protocol (PTP)

- [Understanding Transparent Clocks in Precision Time Protocol on page 152](#)
- [Configuring Transparent Clock Mode for Precision Time Protocol on page 152](#)

Understanding Transparent Clocks in Precision Time Protocol

The Precision Time Protocol (PTP) standardized by IEEE 1588 improves the current methods of synchronization used within a distributed network. You can use PTP across packet-based networks including, but not limited to, Ethernet networks. Queuing and buffering delays in the switch can cause variable delay to packets, which affect path delay measurements. Queuing delays vary based on the network load and also depend on the architecture of the switch.

Transparent clocks measure and adjust for packet delay. The transparent clock computes the variable delay as the PTP packets pass through the switch. The switches act as transparent clocks only and operate between the master and slave clocks in a distributed network. Transparent clocks improve synchronization between the master and slave clocks and ensure that the master and slave clocks are not impacted by the affects of packet delay variation.

The transparent clock measures the residence time (the time that the packet spends passing through the switch), adds the residence time into the correction field of the PTP packet. The slave clock accounts for the packet delay by using both the origin timestamp and the information in the correction field.

End-to-end transparent clocks are supported. With an end-to-end transparent clock, only the residence time is included in the correction field of the PTP packets. The residence time can be sent in a one-step process, which means that the timestamps are sent in one packet. In a two-step process, estimated timestamps are sent in one packet, and additional packets contain updated timestamps.

You can enable or disable a transparent clock globally for the switch. With a global configuration, the same configuration is applied to each interface. If the transparent clock is disabled, PTP packet correction fields are not updated. If the transparent clock is enabled, the PTP packet correct fields are updated.

PTP over Ethernet, UDP over IPv4 and IPv6, and unicast and multicast transparent clocks are also supported.

Related Documentation

- [Configuring Transparent Clock Mode for Precision Time Protocol on page 152](#)
- [e2e-transparent on page 311](#)

Configuring Transparent Clock Mode for Precision Time Protocol

In a distributed network, you can configure Precision Time Protocol (PTP) transparent clocks to help synchronize the timing across the network.

To configure the transparent clock mode for Precision Time Protocol (PTP):

1. In configuration mode, go to the **[edit protocols ptp]** hierarchy level.

```
[edit]  
user@host# edit protocols ptp
```


2. Specify transparent clock mode:

```
[edit protocols ptp]  
user@host# set e2e-transparent
```

**Related
Documentation**

- [Understanding Transparent Clocks in Precision Time Protocol on page 152](#)
- [e2e-transparent on page 311](#)
- [\[edit protocols ptp\] Hierarchy Level](#)
- [show ptp global-information on page 456](#)

PART 11

Recovery Installation

- [Understanding Recovery Installation on page 157](#)

Understanding Recovery Installation

- [Creating an Emergency Boot Device on page 157](#)
- [Performing a Recovery Installation Using an Emergency Boot Device on page 159](#)
- [Recovering from a Failed Software Installation on page 160](#)

Creating an Emergency Boot Device

If Junos OS on the device is damaged in some way that prevents the software from loading properly, you can use an emergency boot device to repartition the primary disk and load a fresh installation of Junos OS. Use the following procedure to create an emergency boot device.

Before you begin, you need to download the installation media image for your device and Junos OS release from <http://www.juniper.net/customers/support/>.



NOTE: You can create the emergency boot device on another Juniper Networks switch or router, or any PC or laptop that supports Linux. The steps you take to create the emergency boot device vary, depending on the device.

To create an emergency boot device:

1. Use FTP to copy the installation media image into the **/var/tmp** directory on the device.
2. Insert a USB device into the USB port.
3. From the Junos OS command-line interface (CLI), start the shell:

```
user@device> start shell
%
```

4. Switch to the root account using the **su** command:

```
% su
Password: password
```



NOTE: The password is the root password for the device. If you logged in to the device as root, you do not need to perform this step.

5. Enter the following command on the device:

```
root@device% dd if=/var/tmp/filename of=/dev/da1 bs=16k
```

The device writes the installation media image to the USB device:

```
root@device% dd if=install-media-qfx-5e-15.1X53-D30.5-domestic.img of=/dev/da0
bs=1m
1399+0 records in
1399+0 records out
1466957824 bytes transferred in 394.081902 secs (3722469 bytes/sec)
```

6. Enter the following command:

```
root@device% dd if=/var/tmp/filename of=/dev/da0 bs=1048576
```

The device writes the installation media image to the USB device:

```
root@device% dd if=/var/tmp/jinstall-vjunos-usb-13.2.img of=/dev/da0 bs=1048576
11006+1 records in
11006+1 records out
180332544 bytes transferred in 71.764266 secs (2512846 bytes/sec)
```



NOTE: The device automatically create a recovery Junos OS image.

The “Select a recovery image” menu appears on the console when one of these switches is booted and unable to load a version of Junos OS. You can follow the instructions in the “Select a recovery image” menu to load the Junos OS image for one of these switches.

7. Log out of the shell:

```
root@device% exit
% exit
user@device>
```

**Related
Documentation**

- *USB Port Specifications for the QFX Series*
- *Performing a Recovery Installation*
- *Performing a QFabric System Recovery Installation on the Director Group*
- [Performing a Recovery Installation Using an Emergency Boot Device on page 159](#)

Performing a Recovery Installation Using an Emergency Boot Device

If Junos OS on your device is damaged in some way that prevents the software from loading correctly, you may need to perform a recovery installation using an emergency boot device (for example, a USB flash drive) to restore the default factory installation. Once you have recovered the software, you need to restore the device configuration. You can either create a new configuration as you did when the device was shipped from the factory, or if you saved the previous configuration, you can simply restore that file to the device.

If at all possible, you should try to perform the following steps before you perform the recovery installation:

1. Ensure that you have an emergency boot device to use during the installation. See [“Creating an Emergency Boot Device” on page 157](#) for information on how to create an emergency boot device.
2. Copy the existing configuration in the file `/config/juniper.conf.gz` from the device to a remote system, such as a server, or to an emergency boot device. For extra safety, you can also copy the backup configurations (the files named `/config/juniper.conf.n`, where *n* is a number from 0 through 9) to a remote system or to an emergency boot device.

You can use the system snapshot feature to complete this step. The system snapshot feature takes a “snapshot” of the files currently used to run the QFX Series switch—the complete contents of the `/config` and `/var` directories, which include the running Juniper Networks Junos OS, the active configuration, and the rescue configuration—and copies all of these files into a memory source. See *Creating a Snapshot and Using It to Boot a Device*.



NOTE: System snapshot is not supported on the QFX10000 and QFX5200 switches.



WARNING: The recovery installation process completely overwrites the entire contents of the internal flash storage.

3. Copy any other stored files to a remote system as desired.

To reinstall Junos OS:

1. Insert the emergency boot device into the device.
2. Power cycle the device.

The emergency boot device is detected. At this time, you can load the Junos OS from the emergency boot device onto the internal flash storage.

3. The software prompts you with the following option if you have a snapshot saved on the emergency boot device:

```
Junos Snapshot Installer - (c) Juniper Networks 2013
Reboot
Install Junos Snapshot [14.1X53-D11_vjunos.61]
Boot to host shell [debug]
```

Select **Install Junos Snapshot** install the snapshot.

The software prompts you with the following option if you have Junos OS software from the factory installed on the emergency boot device.

```
Juniper Linux Installer - (c) Juniper Networks 2014
Reboot
Install Juniper Linux Platform
Boot to host shell [debug]
```

Select **Install Juniper Linux Platform** to install the Junos OS software from the emergency boot device.

4. The device copies the software from the emergency boot device, occasionally displaying status messages. Copying the software can take up to 12 minutes.

When the software is finished being copied from the emergency device to the device, the device reboots from the internal flash storage on which the software was just installed. When the reboot is complete, the device displays the Junos OS login prompt:

```
root@switch#
```

5. Create a new configuration as you did when the device was shipped from the factory, or restore the previously saved configuration file to the device.
6. Remove the emergency boot device.

**Related
Documentation**

- [Creating an Emergency Boot Device on page 157](#)

Recovering from a Failed Software Installation

Problem	Description: If the Junos OS appears to have been installed but the CLI does not work, or if the switch has no software installed, you can use this recovery installation procedure to install the Junos OS.
Solution	If a Junos OS image already exists on the switch, you can either install the new Junos OS package in a separate partition, in which case both Junos OS images remain on the switch, or you can remove the existing Junos OS image before you start the new installation process.



NOTE: QFX5100, EX4600, and OCX Series switches do not have a separate partition to reinstall a Junos OS image.

A recovery image is created automatically on these switches. If a previously-running switch is powered on and unable to boot using a Junos OS image, you can boot the switch using the recovery Junos OS image by selecting an option in the “Select a recovery image” menu.

We suggest creating a system snapshot on your switch onto the external USB flash drive, and using the snapshot for recovery purposes. The system snapshot feature takes a “snapshot” of the files currently used to run the device—the complete contents of the /config directories, which include the running Juniper Networks Junos OS, the active configuration, and the rescue configuration, as well as the host OS—and copies all of these files into an external USB flash drive. See *Creating a Snapshot and Using It to Boot a QFX3500 and QFX3600 Switch* or *Creating a Snapshot and Using It to Boot a Device*.

System snapshot is not supported on QFX10002 switches.

To perform a recovery installation:

1. Power on the switch. The loader script starts.
2. After the message **Loading /boot/defaults/loader.conf** appears, you are prompted with the following message:

Hit [Enter] to boot immediately, or space bar for command prompt.

Press the Spacebar to enter the manual loader. The **loader>** prompt appears.



NOTE: The loader prompt does not appear on QFX5100, QFX5200, EX4600, and OCX Series switches.

On QFX5100, QFX5200, EX4600, and OCX Series switches only, a recovery image is automatically saved if a previously-running switch is powered on and unable to boot using a Junos OS image.

The “Select a recovery image” menu appears on the console when one of these switches is booted and unable to load a version of Junos OS. Follow the instructions in the “Select a recovery image” menu to load the recovery version of Junos OS for one of these switches.

You can ignore the remainder of this procedure if you are using a QFX5100, QFX5200, EX4600, or OCX Series switch.

3. Enter the following command:

```
loader> install [- --format] [- --external] source
```

where:

- **format**—Enables you to erase the installation media before installing the installation package. If you do not include this option, the system installs the new Junos OS in a different partition from that of the most recently installed Junos OS.
- **external**—Installs the installation package onto external media (a USB stick, for example).
- **source**—Represents the name and location of the Junos OS package, either on a server on the network or as a file on an external media, as shown in the following two examples:
 - Network address of the server and the path on the server; for example,
tftp://192.17.1.28/junos/jinstall-qfx-5e-flex-15.1X53-D30.5-domestic-signed.tgz
 - Junos OS package on a USB device (commonly stored in the root drive as the only file), for example,
file:///jinstall-qfx-5e-flex-15.1X53-D30.5-domestic-signed.tgz.

The installation now proceeds normally and ends with a login prompt.

- | | |
|------------------------------|---|
| Related Documentation | <ul style="list-style-type: none">• <i>Creating a Snapshot and Using It to Boot a QFX3500 and QFX3600 Switch</i>• <i>Creating a Snapshot and Using It to Boot a Device</i> |
|------------------------------|---|

PART 12

Routine Monitoring

- [Understanding Routine Monitoring on page 165](#)

CHAPTER 13

Understanding Routine Monitoring

- [Monitoring System Process Information on page 165](#)
- [Monitoring System Properties on page 166](#)
- [Monitoring Interface Status and Traffic on page 167](#)
- [Other Tools to Configure and Monitor Devices Running Junos OS on page 168](#)
- [Fabric Management Overview on page 168](#)

Monitoring System Process Information

Purpose View the processes running on the device.

Action To view the software processes running on the device:
[edit system]

user@switch> [show system processes](#)

Meaning [Table 32](#) summarizes the output fields in the system process information display.

The display includes the total CPU load and total memory utilization.

Table 32: Summary of System Process Information Output Fields

Field	Values
PID	Identifier of the process.
Name	Owner of the process.
State	Current state of the process.
CPU Load	Percentage of the CPU that is being used by the process.
Memory Utilization	Amount of memory that is being used by the process.
Start Time	Time of day when the process started.

- Related Documentation**
- [Monitoring System Properties on page 166](#)
 - [show system uptime on page 1185](#)

Monitoring System Properties

Purpose View system properties such as the name, IP address, and resource usage.

Action To monitor system properties in the CLI, enter the following commands:

- [show system uptime](#)
- [show system users](#)
- [show system storage](#)

Meaning [Table 33](#) summarizes key output fields in the system properties display.

Table 33: Summary of Key System Properties Output Fields

Field	Values	Additional Information
General Information		
Serial Number	Serial number of device.	
Junos OS Version	Version of Junos OS active on the switch, including whether the software is for domestic or export use.	Export software is for use outside the USA and Canada.
Hostname	Name of the device.	
IP Address	IP address of the device.	
Loopback Address	Loopback address.	
Domain Name Server	Address of the domain name server.	
Time Zone	Time zone on the device.	
Time		
Current Time	Current system time, in Coordinated Universal Time (UTC).	
System Booted Time	Date and time when the device was last booted and how long it has been running.	
Protocol Started Time	Date and time when the protocols were last started and how long they have been running.	

Table 33: Summary of Key System Properties Output Fields (*continued*)

Field	Values	Additional Information
Last Configured Time	Date and time when a configuration was last committed. This field also shows the name of the user who issued the last commit command.	
Load Average	CPU load average for 1, 5, and 15 minutes.	
Storage Media		
Internal Flash Memory	Usage details of internal flash memory.	
External Flash Memory	Usage details of external USB flash memory.	
Logged in Users Details		
User	Username of any user logged in to the switch.	
Terminal	Terminal through which the user is logged in.	
From	System from which the user has logged in. A hyphen indicates that the user is logged in through the console.	
Login Time	Time when the user logged in.	This is the user@switch field in show system users command output.
Idle Time	How long the user has been idle.	

- Related Documentation**
- [Monitoring System Process Information on page 165](#)
 - [show system processes on page 1097](#)

Monitoring Interface Status and Traffic

Purpose View interface status to monitor interface bandwidth utilization and traffic statistics.

- Action**
- To view interface status for all the interfaces, enter **show interfaces xe**.
 - To view status and statistics for a specific interface, enter **show interfaces xe interface-name**.
 - To view status and traffic statistics for all interfaces, enter either **show interfaces xe detail** or **show interfaces xe extensive**.

Meaning For details about output from the CLI commands, see [show interfaces xe](#).

Other Tools to Configure and Monitor Devices Running Junos OS

Apart from the command-line interface, Junos OS also supports the following applications, scripts, and utilities that enable you to configure and monitor devices running Junos OS:

- Junos XML Management Protocol Application Programming Interface (API)—Application programmers can use the Junos XML Management Protocol API to monitor and configure Juniper Networks devices. Juniper Networks provides a Perl module with the API to help you more quickly and easily develop custom Perl scripts for configuring and monitoring the devices.
- NETCONF Application Programming Interface (API)—Application programmers can also use the NETCONF API to monitor and configure Juniper Networks devices.
- Junos OS commit scripts—You can define scripts to enforce custom configuration tasks, enforce consistency, prevent common mistakes, and more. Every time you commit a new candidate configuration, the active commit scripts are called to inspect the new candidate configuration. If a configuration violates your custom rules, the script can instruct the Junos OS to perform various actions, including making changes to the configuration and generating custom, warning, and system log messages.
- Junos OS Op scripts—You can add your own commands to the operation-mode CLI. You can use these scripts to automate troubleshooting of known network problems and correct them.
- Junos OS event scripts—You can use event scripts to diagnose and fix issues, monitor the overall status of the system, and examine errors periodically. Event scripts are similar to op scripts except that certain events on the switch will trigger these scripts.
- CHEF—You can use CHEF automate the provisioning and management of compute, networking, and storage resources. Chef for Junos OS provides support for Chef on selected Junos OS devices, allowing you to automate common switching network configurations.
- Puppet—You can use PUPPET for configuration management. Puppet provides an efficient and scalable solution for managing the configurations of large numbers of devices. System administrators take advantage of Puppet to manage compute resources such as physical and virtual servers.

Related Documentation

- [CLI User Interface Overview on page 37](#)
- [*NETCONF XML Management Protocol Developer Guide*](#)
- [Understanding Device and Network Management Features on page 1401](#)

Fabric Management Overview

Fabric management is responsible for setting up and managing the connections between Packet Forwarding Engines (PFEs) in a switching fabric. The switching fabric consists of Switch Interface Boards (SIBs) which provide the interconnection between the PFEs across Flexible PIC Concentrators (FPCs).

Fabric management is event-driven. When the system is in a steady state, fabric management only collects fabric statistics periodically, monitors the health of the hardware, and responds to CLI queries. Other fabric management activity is triggered as a response to managed events and faults to minimize packet loss.

Fabric management handles the following managed events:

- Online and planned offline of field replaceable units (FRUs)—FPCs and SIBs are two classes of FRUs.
- Software upgrade
- Firmware upgrade

Fabric management handles the following faults:

- Hardware faults—Includes hardware errors at the link level, board level, or chassis level. These faults occur during initialization or during runtime. Temperature and voltage triggered faults, card level power brick failure, and control board failure are some examples of hardware faults.
- Operator faults—These faults can occur during run-time or install time. Accidental removal of a functional FRU, bent pins at the mid-plane due to improper insertion of a line card, and power loss of a chassis or a FRU are some examples of operator faults.
- Software faults—The most likely software faults that affect fabric management are process crash and restart, kernel crash and reboot, and process starvation, causing loss of control plane keep-alives.

The following commands are used to monitor fabric management:

- **show chassis fabric fpcs**—Indicates whether any fabric links are in error on the FPCs' side.
- **show chassis fabric plane location**—Displays the fabric plane location of each SIB.
- **show chassis fabric sibs**—Indicates whether any fabric links are in error on the SIBs' side.
- **show chassis fabric summary**—Displays the state of all fabric planes.
- **show chassis fabric topology**—Displays the input-output link topology.

Related Documentation

- [show chassis fabric fpcs on page 585](#)
- [show chassis fabric plane-location on page 624](#)
- [show chassis fabric sibs on page 629](#)
- [show chassis fabric summary on page 641](#)
- [show chassis fabric topology on page 647](#)

PART 13

Standard Software Installation and Upgrade

- [Understanding Standard Software Installation and Upgrade on page 173](#)

CHAPTER 14

Understanding Standard Software Installation and Upgrade

- [Junos OS Package Names on page 173](#)
- [Software Installation Overview on page 174](#)
- [Creating an Emergency Boot Device on page 175](#)
- [Installing and Recovering Software Using the Open Network Install Environment \(ONIE\) on page 176](#)
- [Upgrading Software on page 189](#)

Junos OS Package Names

You upgrade the Juniper Networks Junos OS on the switch by copying a software package to your switch or another system on your local network and then installing the new software package on the switch.

A software package name is in the following format:



NOTE: A signed domestic package is used as an example only. Other types of software packages might be available in future releases.

package-name-m.nZx.y-domestic-signed.tgz

where:

- ***package-name*** is the name of the package—for example, ***jinstall-qfx***.
- ***m.n*** is the software release, with ***m*** representing the major release number and ***n*** representing the minor release number—for example, ***15.1***.
- ***Z*** indicates the type of software release, where ***R*** indicates released software and ***B*** indicates beta-level software.
- ***x.y*** represents the maintenance software release, with ***x*** representing the maintenance software release number and ***y*** representing the maintenance software spin number—for example, ***1.5***.

A sample switch software package name is:

```
jinstall-qfx-5e-flex-15.1X53-D30.5-domestic-signed.tgz
```



NOTE: The image in the .tgz file is not automatically cleaned up after the image is successfully installed. If necessary, you can free up the storage space by using the `request system storage cleanup` command.

**Related
Documentation**

- [Upgrading Software on page 189](#)
- [Installing and Recovering Software Using the Open Network Install Environment \(ONIE\) on page 176](#)
- [Upgrading Software on a QFabric System](#)
- [Software Installation Overview on page 174](#)

Software Installation Overview

A device is delivered with the Junos OS preinstalled. As new features and software fixes become available, you can upgrade your software to use them.

When you power on the switch, it starts (boots) using the installed software.

You upgrade the Junos OS on a switch by copying a software package to a switch or other system on your local network and then using the CLI to install the new software on the switch. You then reboot the switch, which boots from the upgraded software. After a successful upgrade, you should back up the new current configuration to a secondary device.

During a successful upgrade, the installation package removes all files from the `/var/tmp` directory of the switch and completely reinstalls the existing software. It retains configuration files, and similar information, such as secure shell and host keys, from the previous version. The previous software package is preserved in a separate disk partition, and you can manually revert to it if necessary. If the software installation fails for any reason, such as loss of power during the installation process, the system returns to the originally active installation when you reboot.

If you encounter any difficulties during software installation or an upgrade, you can use the recovery installation procedure to install the Junos OS on the switch.

**Related
Documentation**

- [Upgrading Software on page 189](#)
- [Recovering from a Failed Software Installation on page 160](#)
- [Performing a Recovery Installation](#)

Creating an Emergency Boot Device

If Junos OS on the device is damaged in some way that prevents the software from loading properly, you can use an emergency boot device to repartition the primary disk and load a fresh installation of Junos OS. Use the following procedure to create an emergency boot device.

Before you begin, you need to download the installation media image for your device and Junos OS release from <http://www.juniper.net/customers/support/>.



NOTE: You can create the emergency boot device on another Juniper Networks switch or router, or any PC or laptop that supports Linux. The steps you take to create the emergency boot device vary, depending on the device.

To create an emergency boot device:

1. Use FTP to copy the installation media image into the `/var/tmp` directory on the device.
2. Insert a USB device into the USB port.
3. From the Junos OS command-line interface (CLI), start the shell:

```
user@device> start shell
%
```

4. Switch to the root account using the `su` command:

```
% su
Password: password
```



NOTE: The password is the root password for the device. If you logged in to the device as root, you do not need to perform this step.

5. Enter the following command on the device:

```
root@device% dd if=/var/tmp/filename of=/dev/dal bs=16k
```

The device writes the installation media image to the USB device:

```
root@device% dd if=install-media-qfx-5e-15.1X53-D30.5-domestic.img of=/dev/da0
bs=1m
1399+0 records in
1399+0 records out
1466957824 bytes transferred in 394.081902 secs (3722469 bytes/sec)
```

6. Enter the following command:

```
root@device% dd if=/var/tmp/filename of=/dev/da0 bs=1048576
```

The device writes the installation media image to the USB device:

```
root@device% dd if=/var/tmp/jinstall-vjunos-usb-13.2.img of=/dev/da0 bs=1048576
11006+1 records in
11006+1 records out
180332544 bytes transferred in 71.764266 secs (2512846 bytes/sec)
```



NOTE: The device automatically create a recovery Junos OS image.

The “Select a recovery image” menu appears on the console when one of these switches is booted and unable to load a version of Junos OS. You can follow the instructions in the “Select a recovery image” menu to load the Junos OS image for one of these switches.

7. Log out of the shell:

```
root@device% exit
% exit
user@device>
```

**Related
Documentation**

- [USB Port Specifications for the QFX Series](#)
- [Performing a Recovery Installation](#)
- [Performing a QFabric System Recovery Installation on the Director Group](#)
- [Performing a Recovery Installation Using an Emergency Boot Device on page 159](#)

Installing and Recovering Software Using the Open Network Install Environment (ONIE)



NOTE: ONIE is not supported on QFX10002 and QFX10008 switches in Junos OS Release 15.1X53-D30.

ONIE, the open network install environment from Cumulus Networks, is a network OS installer that installs Junos OS and third party applications on a switch. Juniper Network switches come preinstalled with ONIE. When you turn on a switch, the ONIE discovery and execution (ODE) application locates the management Ethernet interface and the Junos OS software package, which can be found either locally on the switch or on the network using HTTP, FTP, or TFTP. After the switch discovers and downloads the Junos OS software package, the switch installs the Junos OS software, reboots, and then boots from Junos OS. Junos OS then becomes the default software image .



NOTE: If you want to use the Junos OS CLI to install software, see [“Upgrading Software” on page 189](#).

Upgrading involves these tasks:

- [Understanding the Open Network Install Environment on page 178](#)
- [Downloading Software Files with a Browser on page 178](#)
- [Connecting to the Console Port on page 179](#)
- [Backing Up the Current Configuration Files on page 179](#)
- [Uninstalling the Existing Version of Junos OS on page 180](#)

- [Installing a Junos OS Software Package That Resides on a Webserver or DHCP Server with DHCP Options Configured on page 180](#)
- [Installing Junos OS Software Using Secure Copy Protocol \(SCP\) on page 181](#)
- [Installing Junos OS Software Using FTP or TFTP Without a Webserver on page 182](#)
- [Installing Junos OS Software Using DHCP Server with No DHCP Options Configured on page 182](#)
- [Installing Junos OS Software Using Webserver Without DHCP Configured on page 184](#)
- [Installing Junos OS Software Using USB Media on page 185](#)
- [Verifying Software Installation on page 185](#)
- [Troubleshooting Unsupported Network Operating System on page 187](#)
- [Troubleshooting ONIE ISO Files on page 187](#)
- [Creating an Emergency Boot Device for ONIE Software and Network Operating System on page 187](#)
- [Performing a Recovery Installation on page 188](#)

Understanding the Open Network Install Environment

Juniper Network switches come preinstalled with Junos OS, which is designed to be used with the Junos OS CLI. The Junos OS CLI uses .tgz files; an example filename is `jinstall-qfx-5e-flex-15.1X53-D30.5-domestic-signed.tgz`. However, if you want to use ONIE to install Junos OS, you need to uninstall the existing Junos OS and reinstall the Junos OS image that has a .bin extension—for example, `jnr-qfx-5e-flex-nos-updater-15.1X53-D30.5.bin` file. When you uninstall the existing Junos OS image (.tgz file) and reinstall with the image that ONIE requires (.bin file), any existing configuration information is destroyed. ONIE only supports the ability to install a fresh version of Junos OS or reinstall the Junos OS software. ONIE does not support upgrading to a different version of Junos OS.

When you log into the switch with ONIE, you see the install boot menu:

- Juniper Linux (This is a default menu option.)
- Juniper Linux Debug
- Juniper Linux Recovery
- Go to ONIE Loader
 - ONIE: Install OS (This is a default menu option.)
 - ONIE: Rescue
 - ONIE: Uninstall OS
 - ONIE: Update ONIE
 - ONIE: Embed ONIE

You can use the following commands to install and uninstall Junos OS and start and stop the ONIE ODE application:

- **onie-nos-install**

Installs Junos OS from any URL, such as `http://`, `ftp://`, and `file://`.

- **onie-uninstaller**

Uninstalls Junos OS.

- **onie-discovery-start**

The discovery process starts automatically. However, if you stop the discovery process by issuing the **onie-discovery-stop** command, you can restart the discovery process by issuing the **onie-discovery-start** command.

- **onie-discovery-stop**

Stops the discovery process. To restart the discovery process, issue the **onie-discovery-start** command.

Downloading Software Files with a Browser

To download the software package from the Juniper Networks Support website at <http://www.juniper.net/support/>.



NOTE: To access the download site, you must have a service contract with Juniper Networks and an access account. If you need help obtaining an account, complete the registration form at the Juniper Networks website <https://www.juniper.net/registration/Register.jsp>.

This procedure shows you how to upgrade software on a QFX Series device, but you can follow the same procedure for any device unless otherwise specified.

To download a software package:

1. Using a Web browser, navigate to the <http://www.juniper.net/support>.
2. Click **Download Software**.
3. In the By Technology box, click **Switching > QFXSeries > QFX5200**.
4. In the QFXSeries section, click the name of the platform for which you want to download software.
5. Click the **Software** tab and select the install package from the Install Package box.
A login screen appears.
6. Enter your name and password and press Enter.
7. Read the End User License Agreement, click the **I agree** radio button, and then click **Proceed**.
8. Save the `jnpr-qfx-5e-flex-nos-updater-15.1X53-D30.5.bin` file to your computer.
9. Open or save the installation package either to the local system in the `var/tmp` directory or to a remote location. If you are copying the installation package to a remote system, make sure that you can access it using HTTP, TFTP, FTP, or SCP.

Connecting to the Console Port

We recommend that you connect to the console port while installing the installation package so you can respond to any required user input and detect any errors that might occur.

Backing Up the Current Configuration Files

Before you install the new installation package, we strongly recommend that you back up your current configuration files because the upgrade process removes all of the stored files on the switch.

To back up your current configuration files, enter the **save** command:

```
user@switch# save filename
```

Executing this command saves a copy of your configuration files to a remote location such as an external USB device.

Uninstalling the Existing Version of Junos OS

The switch comes preinstalled with a version of Junos OS that is to be used with the Junos OS CLI. However, if you want to use ONIE to install Junos OS, you need to uninstall the existing Junos OS and reinstall the Junos OS image that has a .bin extension—for example, `jnpr-qfx-5e-flex-nos-updater-15.1X53-D30.5.bin` file.

To uninstall your existing Junos OS version:

1. Select **Go to ONIE Loader** from the GNU GRUB menu.
2. Select **ONIE: Uninstall OS** from the GNU GRUB menu.

The Junos OS is uninstalled, and the switch reboots.

By default, the ONIE discovery and execution (ODE) application attempts to discover and fetch an image from a configured DHCP or webserver and the management IP address of the switch and the IP address of the default gateway. If you want to manually configure static addressing for the management IP address of the switch, issue **onie-discovery-stop** command at the ONIE prompt, and then manually configure the management IP address and IP address of the default gateway.

For example:

```
ONIE:/ # onie-discovery-stop
ONIE:/ # ifconfig eth0 10.204.32.96 netmask 255.255.254.0
ONIE:/ # route add default gw 10.204.47.254
```

To restart the ONIE discovery and execution (ODE) application, issue the **onie-discovery-start** command.

For example:

```
ONIE:/ # onie-discovery-start
```

Installing a Junos OS Software Package That Resides on a Webserver or DHCP Server with DHCP Options Configured

To install a Junos OS software package residing on a webserver or DHCP server:

1. Copy the software image with the filename **onie-installer** to the `var/www/html` directory of the webserver or DHCP server.
2. Configure the DHCP option 114 in the DHCP server to redirect to the webserver to fetch the Junos OS software image.
3. Uninstall the preinstalled Junos OS version.
 - Select **Go to ONIE Loader** from the GNU GRUB menu.
 - Select **ONIE: Uninstall OS** from the GNU GRUB menu.

The Junos OS is uninstalled, and the switch reboots.

4. Configure DHCP option 114 and other DHCP options as necessary.

Here is a sample Windows Open DHCP server configuration with DHCP option 114 configured.

```
#Following are range-specific DHCP options.
#You can copy more option names from [GLOBAL_OPTIONS]
IP=10.204.42.250
SubnetMask=255.255.240.0
Router=10.204.47.254
114="http://10.207.66.147/onie-installer"
```

Here is a sample boot initialization log, showing the options you just configured:

```
Info: Trying DHCPv4 on interface: eth0
ONIE: Using DHCPv4 addr: eth0: 10.204.42.250 / 255.255.240.0
ONIE: Starting ONIE Service Discovery
Info: Fetching http://10.207.66.147/onie-installer ...
ONIE: Executing installer: http://10.207.66.147/onie-installer  <-----
automatically redirects to web sever to fetch Junos OS image.
Verifying image checksum ... OK.
Preparing image archive ... OK.
Installing Juniper NOS...
```

The log shows that the installation process has fetched the Junos OS software image from the DHCP server and is installing the Junos OS software.

The switch reboots and the GNU GRUB menu is displayed.

Installing Junos OS Software Using Secure Copy Protocol (SCP)

To install Junos OS software using SCP:

1. Uninstall the preinstalled Junos OS version.
 - Select **Go to ONIE Loader** from the GNU GRUB menu.
 - Select **ONIE: Uninstall OS** from the GNU GRUB menu.

The Junos OS is uninstalled, and the switch reboots.

By default, the ONIE discovery and execution (ODE) application attempts to discover and fetch an image from a configured webserver. If you do not have DHCP configured, you will need to stop the ONIE discovery and execution (ODE) application and manually configure static addressing for the management IP address of the switch,

For example:

```
ONIE:/ # onie-discovery-stop
ONIE:/ # ifconfig eth0 10.204.32.96 netmask 255.255.254.0
ONIE:/ # route add default gw 10.204.47.254
```

2. Use SCP to copy the Junos OS image from a server or other location to the `/var/tmp` directory on the switch.

For example:

```
user@server scp jnpr-qfx-5e-flex-nos-updater-15.1X53-D30.5.bin
root@10.204.32.196:/var/tmp/
```

3. Issue the `onie-nos-install` command in the `/var/tmp` directory to install Junos OS software.

```
ONIE:/var/tmp # onie-nos-install file:///var/tmp/
jnpr-qfx-5e-flex-nos-updater-15.1X53-D30.5.bin
```

The switch reboots and displays the GNU GRUB menu.

Installing Junos OS Software Using FTP or TFTP Without a Webserver

To install Junos OS software using FTP or TFTP:

1. Uninstall the preinstalled Junos OS version.
 - Select **Go to ONIE Loader** from the GNU GRUB menu.
 - Select **ONIE: Uninstall OS** from the GNU GRUB menu.

The Junos OS is uninstalled, and the switch reboots.

By default, the ONIE discovery and execution (ODE) application attempts to discover and fetch an image from a configured webserver. If you do not have DHCP configured, you will need to stop the ONIE discovery and execution (ODE) application and manually configure static addressing for the management IP address of the switch,

For example:

```
ONIE:/ # onie-discovery-stop
ONIE:/ # ifconfig eth0 10.204.32.96 netmask 255.255.254.0
ONIE:/ # route add default gw 10.204.47.254
```

2. Copy the Junos OS image to an FTP or TFTP directory.
3. Issue the **onie-nos-install** command at the ONIE prompt to install the Junos OS software.

If you are using FTP:

```
ONIE:/ # onie-nos-install ftp://<username>:<password>@10.209.152.22/
jnpr-qfx-5e-flex-nos-updater-15.1X53-D30.5.bin
```

If you are using TFTP:



NOTE: The software image should be located in the `/tftp/boot` directory.

```
ONIE:/ # onie-nos-install tftp://10.207.66.147/
jnpr-qfx-5e-flex-nos-updater-15.1X53-D30.5.bin
```

The switch reboots and displays the GNU GRUB menu.

Installing Junos OS Software Using DHCP Server with No DHCP Options Configured

Use this installation method if you cannot modify or set the DHCP options on your DHCP server.

To install the Junos OS software using a DHCP server with no DHCP options configured:

1. Copy the software image with the filename **jnpr-qfx-5e-flex-nos-updater-15.1X53-D30.5.bin** to the **var/www/html** directory of the webserver or DHCP server.
2. Uninstall the preinstalled Junos OS version.
 - Select **Go to ONIE Loader** from the GNU GRUB menu.

- Select **ONIE: Uninstall OS** from the GNU GRUB menu.

The Junos OS is uninstalled, and the switch reboots.

3. Issue the **onie-nos-install** command at the ONIE prompt to install the Junos OS software.

For example:

```
ONIE:/ # onie-nos-install http://10.207.66.147/
jnpr-qfx-5e-flex-nos-updater-15.1X53-D30.5.bin
```

Here is sample log with the options you just configured:

```
ONIE:/ # ifconfig
eth0      Link encap:Ethernet  HWaddr 94:DE:80:AA:F2:E1
          inet addr:10.204.42.250  Bcast:10.204.47.255  Mask:255.255.240.0
<<<---- --> Received IP address from DHCP server, but auto redirected to web
server. Installation will not happen because DHCP option (114) is not
configured.
```

```
inet6 addr: fe80::96de:80ff:feaa:f2e1/64 Scope:Link
UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
RX packets:444 errors:0 dropped:0 overruns:0 frame:0
TX packets:17 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:48170 (47.0 KiB)  TX bytes:2678 (2.6 KiB)
Memory:80180000-801a0000
```

This log shows that the installation process has fetched the Junos OS software image from the webserver and is installing the Junos OS software.

```
Stopping: discover... done.
Info: Fetching
http://10.207.66.147/jnpr-qfx-5e-flex-nos-updater-15.1X53-D30.5.bin
Connecting to 10.207.66.147 (10.207.66.147:80)
installer      100% |*****| 464M 0:00:00
ETA
ONIE: Executing installer:
http://10.207.66.147/jnpr-ocx-11-flex-nos-updater-14.1X53-D20.1.bin
Verifying image checksum ... OK.
Preparing image archive ... OK.
Installing Juniper NOS...
```

The switch reboots and the GNU GRUB menu is displayed.

Installing Junos OS Software Using Webserver Without DHCP Configured

Use this installation method if you do not have a DHCP server.

To install the Junos OS software using a webserver without DHCP configured:

1. Because the switch comes preinstalled with the Junos OS to be used with the Junos OS CLI, you need to uninstall this version of software before you can install the Junos OS image to be used with ONIE.

- Select **Go to ONIE Loader** from the GNU GRUB menu.
- Select **ONIE: Uninstall OS** from the GNU GRUB menu.

The Junos OS is uninstalled, and the switch reboots.

2. Select **ONIE: Uninstall OS** from the GNU GRUB menu.

The Junos OS is uninstalled, and the switch reboots.

By default, the ONIE discovery and execution (ODE) application attempts to discover and fetch an image from a configured webserver. Because you do not have DHCP configured, you will need to stop the ONIE discovery and execution (ODE) application and manually configure static addressing for the management IP address of the switch.

For example:

```
ONIE:/ # onie-discovery-stop
ONIE:/ # ifconfig eth0 10.204.32.96 netmask 255.255.254.0
ONIE:/ # route add default gw 10.204.47.254
```

3. Copy the software image to the **var/www/html** directory of the webserver.
4. Issue the **onie-nos-install** command at the ONIE prompt to install the Junos OS software.

For example:

```
ONIE:/ # onie-nos-install http://10.204.35.100/
jnpr-qfx-5e-flex-nos-updater-15.1X53-D30.5.bin
```

Here is sample log:

```
Stopping: discover... done.
Info: Fetching http://10.204.35.100/
jnpr-qfx-5e-flex-nos-updater-15.1X53-D30.5.bin ...
Connecting to 10.204.35.100 (10.204.35.100:80)
installer          100% | ***** | 464M 0:00:00
ETA
ONIE: Executing installer: http://10.204.35.100/
jnpr-qfx-5e-flex-nos-updater-15.1X53-D30.5.bin
Verifying image checksum ... OK.
Preparing image archive ... OK.
Installing Juniper NOS...
```

The log shows that the installation process has fetched the Junos OS software image from the webserver and is installing the Junos OS software.

The switch reboots and the GNU GRUB menu is displayed.

Installing Junos OS Software Using USB Media



NOTE: This installation method is not supported on the QFX5200 switches.

To install the Junos OS software using USB media:

1. Make sure the USB stick is formatted as VFAT or ext2.
2. Copy the software image with the filename **onie-installer** to the USB stick and insert the USB stick into the front USB port of the switch.
3. Uninstall the preinstalled Junos OS version.
 - Select **Go to ONIE Loader** from the GNU GRUB menu.
 - Select **ONIE: Uninstall OS** from the GNU GRUB menu.

The Junos OS is uninstalled, the switch reboots, and the software is installed from the USB stick.

Here is a sample log showing the USB installation process:

```

Version      : master-201412161452.0.1
Build Date: 2014-12-16T14:56+0800
Info: Mounting kernel filesystems... done.
Info: Mounting LABEL=ONIE-BOOT on /mnt/onie-boot ...
Info: Using eth0 MAC address: 94:de:80:aa:f2:e1
Info: eth0: Checking link... scsi 6:0:0:0: Direct-Access      SanDisk Ultra
        2.01 P6
sd 6:0:0:0: [sdb] 31266816 512-byte logical blocks: (16.0 GB/14.9 GiB)
sd 6:0:0:0: [sdb] Write Protect is off
sd 6:0:0:0: [sdb] Write cache: disabled, read cache: enabled, doesn't support
        DPO or FUA
sd 6:0:0:0: [sdb] Attached SCSI disk
up.
Info: Trying DHCPv4 on interface: eth0
ONIE: Using DHCPv4 addr: eth0: 10.204.42.250 / 255.255.240.0
Starting: dropbear ssh daemon... done.
Starting: telnetd... done.
discover: installer mode detected. Running installer.
Starting: discover... done.

Please press Enter to activate this console. Info: eth0: Checking link... up.
Info: Trying DHCPv4 on interface: eth0
ONIE: Using DHCPv4 addr: eth0: 10.204.42.250 / 255.255.240.0
ONIE: Starting ONIE Service Discovery
ONIE: Executing installer: file://dev/sdb/onie-installer <<<<--- Installing
from external(dev/sdb) attached USB.
Verifying image checksum ... OK.
Preparing image archive ... OK.
Installing Juniper NOS..

```

Verifying Software Installation

Purpose Verify that the software was installed successfully on the switch.

Action To verify that the software was properly installed, issue the **show version** command.

```
user@switch > show version
re0:
```

```
-----

Model: qfx5200
Junos: 15.1X53-D30.5
JUNOS OS Kernel 32-bit FLEX [20151030.110850_fbsd-builder_stable_10]
JUNOS OS libs [20151030.110850_fbsd-builder_stable_10]
JUNOS OS runtime [20151030.110850_fbsd-builder_stable_10]
JUNOS OS time zone information [20151030.110850_fbsd-builder_stable_10]
JUNOS py extensions [20151221.192301_builder_junos_151_x53_d30]
JUNOS py base [20151221.192301_builder_junos_151_x53_d30]
JUNOS OS vmguest [20151030.110850_fbsd-builder_stable_10]
JUNOS OS crypto [20151030.110850_fbsd-builder_stable_10]
JUNOS network stack and utilities [20151221.192301_builder_junos_151_x53_d30]
JUNOS libs [20151221.192301_builder_junos_151_x53_d30]
JUNOS runtime [20151221.192301_builder_junos_151_x53_d30]
JUNOS qfx runtime [20151221.192301_builder_junos_151_x53_d30]
JUNOS qfx platform support [20151221.192301_builder_junos_151_x53_d30]
JUNOS modules [20151221.192301_builder_junos_151_x53_d30]
JUNOS qfx modules [20151221.192301_builder_junos_151_x53_d30]
JUNOS Data Plane Crypto Support [20151221.192301_builder_junos_151_x53_d30]
JUNOS daemons [20151221.192301_builder_junos_151_x53_d30]
JUNOS qfx daemons [20151221.192301_builder_junos_151_x53_d30]
JUNOS Voice Services Container package [20151221.192301_builder_junos_151_x53_d30]
JUNOS Services SSL [20151221.192301_builder_junos_151_x53_d30]
JUNOS Services Stateful Firewall [20151221.192301_builder_junos_151_x53_d30]
JUNOS Services RPM [20151221.192301_builder_junos_151_x53_d30]
JUNOS Services PTSP Container package [20151221.192301_builder_junos_151_x53_d30]
JUNOS Services NAT [20151221.192301_builder_junos_151_x53_d30]
JUNOS Services Mobile Subscriber Service Container package
[20151221.192301_builder_junos_151_x53_d30]
JUNOS Services MobileNext Software package
[20151221.192301_builder_junos_151_x53_d30]
JUNOS Services LL-PDF Container package [20151221.192301_builder_junos_151_x53_d30]
JUNOS Services Jflow Container package [20151221.192301_builder_junos_151_x53_d30]
JUNOS Services IPSec [20151221.192301_builder_junos_151_x53_d30]
JUNOS IDP Services [20151221.192301_builder_junos_151_x53_d30]
JUNOS Services HTTP Content Management package
[20151221.192301_builder_junos_151_x53_d30]
JUNOS Services Crypto [20151221.192301_builder_junos_151_x53_d30]
JUNOS Services Captive Portal and Content Delivery Container package
[20151221.192301_builder_junos_151_x53_d30]
JUNOS Border Gateway Function package [20151221.192301_builder_junos_151_x53_d30]
JUNOS AppId Services [20151221.192301_builder_junos_151_x53_d30]
JUNOS Services Application Level Gateways
[20151221.192301_builder_junos_151_x53_d30]
JUNOS Services ACL Container package [20151221.192301_builder_junos_151_x53_d30]
JUNOS Packet Forwarding Engine Support (DC-PFE)
[20151221.192301_builder_junos_151_x53_d30]
JUNOS Online Documentation [20151221.192301_builder_junos_151_x53_d30]
JUNOS Host Software [3.10.62-ltsi-WR6.0.0.21_standard:2.4.0]
JUNOS Host qfx-5e platform package [15.1X53-D30.5]
JUNOS Host qfx-5e control-plane flex package [15.1X53-D30.5]
JUNOS Host qfx-5e base package [15.1X53-D30.5]
JUNOS Host qfx-5e data-plane package [15.1X53-D30.5]
```

In this example, the CLI output shows that Junos OS Release Software 15.1X53-D30.5 was installed on the switch.

Troubleshooting Unsupported Network Operating System

Problem **Description:** ONIE has detected an unsupported network operating system.

Solution You can only run the Junos OS network operating system with the .bin file extension on the switch. If you receive a message that says, “**NOS Updater has detected incompatible NOS. Please uninstall it manually before proceeding. Exiting...**,” you must uninstall the unsupported network operating system from the switch. The following procedure explains how to uninstall the unsupported network operating system.

1. Select **Go to ONIE Loader** from the GNU GRUB menu.
2. Select **ONIE: Uninstall OS** from the GNU GRUB menu.

The network operating system is uninstalled, and the switch reboots.

Troubleshooting ONIE ISO Files

Problem **Description:** In the event that you cannot get ONIE ISO file sources from Alpha Networks, you can perform this procedure to get ISO files sources from an open source.

Solution 1. Download the ONIE ISO files by issuing the **git clone** <https://github.com/opencomputeproject/onie> command.

For example:

```
$ git clone https://github.com/opencomputeproject/onie
```

2. To compile ONIE, make sure you are in the **build-config** directory and then issue the **make -j4 MACHINEROOT=../machine/alpha_networks MACHINE=alpha_networks_snx60a0_486f all** command.

For example:

```
$ cd build-config
$ make -j4 MACHINEROOT=../machine/alpha_networks
MACHINE=alpha_networks_snx60a0_486f all
```

3. Download the built image located in the **build/images/onie-recovery-x86_64-alpha_networks_snx60a0_486f-r0.iso** directory.

Creating an Emergency Boot Device for ONIE Software and Network Operating System

If the Open Network Install Environment (ONIE) software is damaged or corrupted in some way, or the switch went into rescue mode, you can use an emergency boot device to repartition the primary disk and load a fresh installation of ONIE. Use the following procedure to create an emergency boot device.

Before you begin, you need to have the 201502181747 version of ONIE software.



NOTE: In the following procedure, we assume that you are creating the emergency boot device on a switch. You can create the emergency boot device on any PC or laptop that supports Linux. The emergency boot device will contain both ONIE software and the Network Operating Software (NOS).

To create an emergency boot device:

1. Insert the USB stick into the front USB port of the switch.
2. Issue the following command from the directory on the switch in which the ISO file is located:

```
ONIE:/ # dd if=onie-recovery-x86_64-alpha_networks_snx60a0_486f-r0.iso  
of=<usb-detected-drive> bs=1M count=20
```

You can also issue the **dd** command using the full path to where the ISO file is located.

For example, if the ISO file is located in the **/var/tmp/** directory:

```
root@device%dd if=jnpr-qfx-5e-flex-onie-usb-15.1X53-D30.5.iso of=<usb-detected-drive>  
/dev/da0 bs=1m
```

The switch writes the installation media image to the USB device:

3. Remove the USB stick from the USB port of the switch.

Performing a Recovery Installation

In the event that the Open Network Install Environment (ONIE) is corrupted, the switch goes into rescue mode, or you need to reinstall ONIE software for any reason, you need to perform a recovery installation.



NOTE: All Junos OS partitions are destroyed during a recovery installation.



NOTE: Before you can perform a recovery installation, make sure you have an emergency boot device loaded with ONIE software.

1. Insert the emergency boot device into the device.
2. Power cycle the device.
3. Press the **ESC** button to go into the Boot Manager menu.
4. Select **Boot Manager**, and then press **Enter**.
5. Select **Unigen PQS1000** under **Legacy USB**, and then press **Enter**.
6. Select **ONIE: Embed ONIE** from the **ONIE Installer** menu, and then press **Enter**.

The recovery installation proceeds using the emergency boot device.

7. Remove the emergency boot device.
8. Verify that the ONIE software was installed by looking at the installation log file.

For example:

```
Info: Found static url: file:///lib/onie/onie-updater
ONIE: Executing installer: file:///lib/onie/onie-updater
Verifying image checksum ... OK.
Preparing image archive ... OK.
ONIE: Version          : master-201502181747
```

Installation log files are displayed automatically during the installation process, but if you want to verify installation log files at a different time, you can find them in the `/var/log/` directory. To view an installation log file, issue the **tail -f /var/log/onie.log** command.

9. Issue the **var/tmp # parted /dev/sda print** command to verify that the ONIE partitions have been created.

For example:

```
ONIE:/ var/tmp # parted /dev/sda print
Model: ATA TS32GMTS400 (scsi)
Disk /dev/sda: 32.0GB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:
```

Number	Start	End	Size	File system	Name	Flags
1	1049kB	3146kB	2097kB		GRUB-B00T	hidden, bios_grub
2	3146kB	137MB	134MB	ext4	ONIE-B00T	hidden

Related Documentation

- [Upgrading Software on page 189](#)
- [Upgrading Software by Using Automatic Software Download on page 7](#)
- [Configuring a DHCP Server on Switches \(CLI Procedure\) on page 78](#)
- For more information on how to use ONIE, [Open Network Install Environment Documentation](#)

Upgrading Software

To upgrade Junos OS, you need to install the appropriate upgrade package on the device. Upgrading involves these tasks:



NOTE: If you want to use the Open Network Install Environment (ONIE) to install software, see “Installing and Recovering Software Using the Open Network Install Environment (ONIE)” on page 176.

1. [Downloading Software Files with a Browser on page 190](#)
2. [Accessing Software Downloaded to a Remote Location on page 191](#)
3. [Connecting to the Console Port on page 191](#)
4. [Backing Up the Current Configuration Files on page 191](#)

5. [Installing a Standard Software Package on page 191](#)
6. [Installing a Standard Software Package on a QFX10008 Switch on page 192](#)

Downloading Software Files with a Browser

To download the software package from the Juniper Networks Support website, go to <http://www.juniper.net/support/>.



NOTE: To access the download site, you must have a service contract with Juniper Networks and an access account. If you need help obtaining an account, complete the registration form at the Juniper Networks website <https://www.juniper.net/registration/Register.jsp>.

This procedure shows you how to upgrade software on a QFX Series device, but you can follow the same procedure for any device unless otherwise specified.

1. Using a Web browser, navigate to the <http://www.juniper.net/support>.
2. Click **Download Software**.
3. In the **Switching** box, click **Junos OS Platforms**.
4. In the **QFX Series** section, click the name of the platform for which you want to download software.
5. Click the **Software** tab and select the release number from the **Release** drop-down list.
6. In the **Install Package** section of the **Software** tab, select the **Install Package** for the release.
A login screen appears.
7. Enter your name and password and press **Enter**.
8. Read the End User License Agreement, click the **I agree** radio button, and then click **Proceed**.
9. Save the **jinstall-qfx-<version>-domestic-signed.tgz** file on your computer.
10. Open or save the installation package either to the local system in the **var/tmp** directory or to a remote location. If you are saving the installation package to a remote system, make sure that you can access it using HTTP, TFTP, FTP, or scp.

Accessing Software Downloaded to a Remote Location

To access the installation package if you downloaded it to a remote location (for example, any system other than the switch):

1. From the command line, make sure you are in the **/var/tmp** directory of the switch.
2. Start the shell interface:
`user@switch> start shell`
3. Initiate an FTP, TFTP, or scp session.

In this example, FTP is used.

`>ftp`

4. Use FTP to access the remote location where the installation package resides.

`ftp ftp://<hostname>/<pathname>/<package-name-m.mZx-distribution>.tgz.`

where `<package-name-m.mZx-distribution>.tgz` is

`jinstall-qfx-5e-flex-15.1X53-D30.5-domestic-signed.tgz`

5. When prompted, enter your username and password.
6. Use the **get** command to transfer the installation package from the remote location to your **/var/tmp** directory on your switch.

`get <package-name-m.mZx-distribution>.tgz`

7. Close the FTP session:

`bye`

Connecting to the Console Port

We recommend that you connect to the console port while installing the installation package so you can respond to any required user input and detect any errors that may occur.

Backing Up the Current Configuration Files

Before you install the new installation package, we strongly recommend that you back up your current configuration files because the upgrade process removes all of the stored files on the switch.

To back up your current configuration files, enter the **save** command:

`user@switch# save filename`

Executing this command saves a copy of your configuration files to a remote location such as an external USB device.

Installing a Standard Software Package



NOTE: Before you install the software, back up any critical files in **/var/home**. For more information regarding how to back up critical files, contact Customer Support at <http://www.juniper.net/support>.

Install the software in one of two ways:



NOTE: On the switch, use the **force-host** option to force installing the latest version of the Host OS. However, by default, if the Host OS version is different from the one that is already installed on the switch, the latest version is installed without using the **force-host** option.

If the installation package resides locally on the switch, execute the **request system software add <pathname> <source> reboot** command.

For example:

```
user@switch> request system software add  
/var/tmp/jinstall-qfx-5e-flex-15.1X53-D30-domestic.tgz reboot
```

If the Install Package resides remotely, execute the **request system software add <pathname> <source> reboot** command.

For example:

```
user@switch> request system software add  
ftp://ftpsrvr/directory/jinstall-qfx-5e-flex-15.1X53-D30-domestic.tgz reboot
```

After the reboot has finished, verify that the new version of software has been properly installed by executing the **show version** command.

```
user@switch> show version
```

Installing a Standard Software Package on a QFX10008 Switch



NOTE: Before you install the software, back up any critical files in `/var/home`. For more information regarding how to back up critical files, contact Customer Support at <http://www.juniper.net/support>.



NOTE: On the switch, use the `force-host` option to force-install the latest version of the Host OS. However, by default, if the Host OS version is different from the one that is already installed on the switch, the latest version is installed without using the `force-host` option.

The switch contains two routing engines, so you will need to install the software on each routing engine (re0 and re1).

If the installation package resides locally on the switch, execute the **request system software add <pathname><source> (re0 | re1)** command.

To install the software on re0:

```
user@switch> request system software add
/var/tmp/jinstall-qfx-10-m-flex-15.1X53-D30-secure-domestic-signed.tgz re0
```

If the Install Package resides remotely from the switch, execute the **request system software add <pathname><source> re0** command.

For example:

```
user@switch> request system software add
ftp://ftpsrvr/directory/jinstall-qfx-10-m-flex-15.1X53-D30-secure-domestic-signed.tgz re0
```

To install the software on re1:

```
user@switch> request system software add
/var/tmp/jinstall-qfx-10-m-flex-15.1X53-D30-secure-domestic-signed.tgz re1
```

If the Install Package resides remotely from the switch, execute the **request system software add <pathname><source> re1** command.

For example:

```
user@switch> request system software add
ftp://ftpsrvr/directory/jinstall-qfx-10-m-flex-15.1X53-D30-secure-domestic-signed.tgz re1
```

Reboot both routing engines.

For example:

```
user@switch> request system reboot both-routing-engines
```

After the reboot has finished, verify that the new version of software has been properly installed by executing the **show version** command.

```
user@switch> show version
```

Related Documentation

- [Installing and Recovering Software Using the Open Network Install Environment \(ONIE\) on page 176](#)
- [Software Installation Overview on page 174](#)

- [Recovering from a Failed Software Installation on page 160](#)
- *Upgrading Jloader Software on QFX Series Devices*
- [request system software add on page 1276](#)
- *Installation and Upgrade Guide*

PART 14

Zero Touch Provisioning (ZTP)

- [Understanding ZTP on page 197](#)

CHAPTER 15

Understanding ZTP

- Understanding Zero Touch Provisioning on page 197
- Configuring Zero Touch Provisioning on page 204
- Monitoring Zero Touch Provisioning on page 209

Understanding Zero Touch Provisioning

- Understanding Zero Touch Provisioning on page 197
- Zero Touch Provisioning Process on page 199
- Zero Touch Provisioning Restart Process Triggers on page 202

Understanding Zero Touch Provisioning



NOTE: To see which platforms support Zero Touch Provisioning (ZTP) in a browser, go to [Feature Explorer](#). In the Explore Features section of the Feature Explorer page, select All Features. In the Features Grouped by Feature Family box, select Zero Touch Provisioning. You can also type the name of the feature in the Search for Features edit box. In previous Junos OS releases on EX Series switches, ZTP was called EZ Touchless Provisioning.



NOTE: QFX5200 switches only work with HTTP in 15.1X53-D30. FTP and TFTP protocols are not supported.

ZTP allows you to provision new Juniper Networks switches in your network automatically, without manual intervention. When you physically connect a switch to the network and boot it with a default factory configuration, it attempts to upgrade the Junos OS software automatically and autoinstall a configuration file from the network.

The switch uses information that you configure on a Dynamic Host Configuration Protocol (DHCP) server to locate the necessary software image and configuration files on the network. If the DHCP server does not respond or provide the software image and configuration files, the switch boots with the preinstalled software and default factory configuration. On switches running Enhanced Layer 2 Software, Junos Extended Dynamic Host Configuration Protocol (JDHCP) is used instead of legacy DHCP. JDHCP supports

the same functionality as DHCP, and all configuration options remain the same. JDHCP is an enhanced version of legacy DHCP software.



NOTE: For detailed information regarding the DHCP and DHCP options, refer to RFC2131 (<http://www.ietf.org/rfc/rfc2131.txt>) and RFC2132 (www.ietf.org/rfc/rfc2132.txt). Also, this document refers to Internet Systems Consortium (ISC) DHCP version 4.2. For more information regarding this version, refer to <http://www.isc.org/software/dhcp/documentation>.

The Zero Touch Provisioning process will either upgrade or downgrade the Junos OS version. During a downgrade:

- On an EX Series switch, If you downgrade to a software version earlier than Junos OS Release 12.2, in which ZTP is not supported, the configuration file autoinstall phase of the Zero Touch Provisioning process does not happen.
- On an EX Series switch, to downgrade to a software version that does not support resilient dual-root partitions (Junos OS Release 10.4R2 or earlier), you must perform some manual work on the switch. For more information, see *Understanding Resilient Dual-Root Partitions on Switches*.



NOTE: On QFX3500 and QFX3600 switches running the original CLI, you cannot use ZTP to upgrade from Junos OS Release 12.2 or later to Junos OS Release 13.2X51-D15 or later.

Zero Touch Provisioning Process

When you boot a switch with the default factory configuration, the following process happens:



NOTE: If you are performing Zero Touch Provisioning (ZTP) with a Junos OS image that contains enhanced automation for the QFX5100 switch, configure root authentication, and the provider name, license type, and deployment scope for Chef and Puppet at the [edit system] hierarchy in the configuration file that is fetched from the server:

```
{master:0}
root# set root-authentication (encrypted-password password | plain-text-password
password | ssh-dsa public-key | ssh-rsa public-key)
root# set extensions providers juniper license-type customer deployment-scope
commercial
root# set extensions providers chef license-type customer deployment-scope
commercial
```

1. If DHCP option 43, suboption 00 (the name of the software image file on the FTP, HTTP, or TFTP server) is configured, the switch compares the version of the provided software image to the version of the software installed on the switch.



NOTE: When the DHCP server cannot use suboption 00, configure the image file using suboption 04. If both suboption 00 and suboption 4 are defined, suboption 04 is ignored.

2. If DHCP option 43, suboption 02 (a symbolic link to the software image file on the FTP, HTTP, or TFTP server), the switch compares the version of the provided software image to the version of the software installed on the switch.
 - If the Junos OS versions are different, the switch downloads the software image from the FTP, HTTP, or TFTP server, installs the Junos OS, and reboots using the default factory configuration.
 - If the software versions are the same, the switch does not upgrade the software.
3. If DHCP option 43, suboption 01 (the name of the configuration file on the FTP, TFTP, or HTTP server) is configured:

If DHCP option 43 suboption 01 is not specified, the switch uses the default factory configuration.

If both DHCP option 43 suboption 01 and suboption 2 are specified, suboption 01 is processed before suboption 02. The Junos OS is upgraded, and then the configuration file is applied.



NOTE: On EX4300 and QFX5100 switches running Enhanced Layer 2 Software, and QFX5100 switches running a Junos OS image that contains enhanced automation, you can specify the name of a script file or a configuration file in suboption 01. ZTP determines if the file is a script file based on the first line that is included in the file. If the first line contains

#! characters followed by an interpreter path—for example, `#!/usr/libexec/ui/cscript`—ZTP determines that the file is a script file, and executes the script file with the specified interpreter path. If the script returns an error, ZTP will fetch the script file and execute the script file until the script executes successfully. If the file does not contain special characters or an interpreter path, ZTP determines that the file is a configuration file.

4. If DHCP option 43, suboption 03 (the transfer mode setting) is configured, the switch accesses the FTP, HTTP, or TFTP server using the specified transfer mode setting—for example, FTP.

If DHCP option 43, suboption 03, is not configured, TFTP becomes the transfer mode automatically.

5. If DHCP option 43, suboption 04 (the name of the software image file on the FTP, HTTP, or TFTP server) is configured, the switch compares the version of the provided software image to the version of the software installed on the switch.



NOTE: When the DHCP server cannot use suboption 00, configure the image file using suboption 04. If both suboption 00 and suboption 4 are defined, suboption 04 is ignored.



NOTE: DHCP option 43 suboptions 05 through 255 are reserved.

6. If DHCP option 150 or option 66 is specified, the IP address of the FTP, HTTP, or TFTP server is configured.



NOTE: You must configure either option 150 or option 66. If you configure both option 150 and option 66, option 150 takes precedence, and option 66 is ignored. Also, make sure you specify an IP address, not a hostname, because name resolution is not supported.

7. (Optional) If DHCP option 7 is specified, you can configure one or more system log (syslog) servers.

8. (Optional) If DHCP option 42 is specified, you can configure one or more NTP servers.
9. (Optional) If DHCP option 12 is specified, you can configure the hostname of the switch.

Zero Touch Provisioning Restart Process Triggers

ZTP restarts when any of the following events occur:

- Request for configuration file, script file, or image file fails.
- Configuration file is incorrect, and commit fails.
- No configuration file and no image file is available.
- Image file is corrupted, and installation fails.
- No file server information is available.
- DHCP client does not have valid ZTP parameters configured.
- When none of the DHCP client interfaces goes to a bound state.
- ZTP transaction fails after six attempts to fetch configuration file or image file.

When any of these events occur, ZTP resets the DHCP client state machine on all of the DHCP client-configured interfaces (management and network) and then restarts the state machine. Restarting the state machine enables the DHCP client to get the latest DHCP server-configured parameters.

Before ZTP restarts, approximately 15 to 30 seconds must elapse to allow enough time to build a list of bound and unbound DHCP client interfaces.

The list of bound and unbound DHCP client interfaces can contain:

- No entries.
- Multiple DHCP client interfaces.

Priority is given to the DHCP client interfaces that have received all ZTP parameters (software image file, configuration file, and file server information) from the DHCP server.

After the lists of bound and unbound client interfaces are created, and a DHCP client gets selected for ZTP activity, then any existing default route is deleted, and the DHCP client interface that was selected adds a new default route. In order to add a new default route, only one ZTP instance can be active.

After ZTP restarts, the DHCP client attempts fetching files from the DHCP server for up to six times, with ten to fifteen seconds elapsing between attempts. Every attempt, whether successful or not, is logged and can be seen on the console.

If there is a failure, or the number of attempts exceeds the limit, ZTP stops. ZTP then clears the DHCP client bindings and restarts state machine on the DHCP-configured interfaces.

The ZTP restart process continues until there is either a successful software upgrade, or an operator manually commits a user configuration and deletes the ZTP configuration.

**Related
Documentation**

- [Configuring Zero Touch Provisioning on page 204](#)
- [Monitoring Zero Touch Provisioning on page 209](#)

Configuring Zero Touch Provisioning



NOTE: To see which platforms support Zero Touch Provisioning (ZTP), in a browser, go to [Feature Explorer](#). In the Explore Features section of the Feature Explorer page, select All Features. In the Features Grouped by Feature Family box, select Zero Touch Provisioning. You can also type the name of the feature in the Search for Features edit box. In previous Junos OS releases on EX Series switches, Zero Touch Provisioning was called EZ Touchless Provisioning.

Zero Touch Provisioning allows you to provision new devices in your network automatically, without manual intervention. When you physically connect a device to the network and boot it with a default configuration, it attempts to upgrade the Junos OS software automatically and autoinstall a configuration file from the network.

The device uses information that you configure on a Dynamic Host Configuration Protocol (DHCP) server to locate the necessary software image and configuration files on the network. If the DHCP server does not respond or provide the software image and configuration files, the device continues using the preinstalled Junos OS software and default factory configuration. On switches running Enhanced Layer 2 Software, Junos Extended Dynamic Host Configuration Protocol (JDHCP) is used instead of legacy DHCP. JDHCP supports the same functionality as DHCP, and all configuration options remain the same. JDHCP is an enhanced version of legacy DHCP software. If you are performing Zero Touch Provisioning with a Junos OS image that contains enhanced automation for the QFX5100 switch, you can use DHCP option 43 suboption 01 to run script files, not just load configuration files. Using scripts, you can create device-specific configuration files, and perform HTTP request operations to web servers to download specific configuration files or Junos OS software.



NOTE: If the ZTP configuration is enabled, the switch broadcasts DHCP DISCOVER packets on its interfaces. If the DHCP server on the network responds with DHCP vendor options set with the necessary values to initiate ZTP, then ZTP proceeds. To disable broadcasting the DHCP DISCOVER packets without performing the ZTP process, manually delete the `auto-image-upgrade` statement located at the `[edit chassis]` hierarchy. If ZTP completes without errors, the `auto-image-upgrade` statement is automatically deleted.

Before you begin, ensure that the switch has access to the following network resources:

- A DHCP server to lease IP addresses and information on software images and configuration files on the network.

Refer to your DHCP server documentation for configuration instructions.

- The File Transfer Protocol (anonymous FTP), Hypertext Transfer Protocol (HTTP), Trivial File Transfer Protocol (TFTP) server on which the software image and configuration files are stored



NOTE: Although TFTP is supported, we recommend that you use FTP or HTTP instead, because these transport protocols are more reliable.

- A Domain Name System (DNS) server to perform reverse DNS lookup
- (Optional) An NTP server to perform time synchronization on the network
- (Optional) A system log (syslog) server to manage system log messages and alerts



CAUTION: We recommend that you do not commit a user configuration while the device is performing ZTP activity—for example, updating the software image or applying a configuration file.

Perform the following steps to configure ZTP:

1. Boot the device.

The device continues to use the preinstalled Junos OS software and default factory configuration.

2. Issue the **request system zeroize** command on the device.
3. Download the software image file and the configuration file to the FTP, HTTP, TFTP server that the device will download these files from.

You can download either one or both of these files.



NOTE: If you are performing Zero Touch Provisioning with a Junos OS image that contains enhanced automation for the QFX5100 device, configure root authentication, and the provider name, license type, and deployment scope for Chef and Puppet at the [edit system] hierarchy in the configuration file that is fetched from the server:

```
{master:0}
root# set root-authentication (encrypted-password password |
plain-text-password password | ssh-dsa public-key | ssh-rsa public-key)
root# set extensions providers juniper license-type customer deployment-scope
commercial
root# set extensions providers chef license-type customer deployment-scope
commercial
```

4. Configure the DHCP server to provide the necessary information to the device.

Configure IP address assignment.

You can configure dynamic or static IP address assignment for the device's management address. To determine the device's management MAC address for static IP address mapping, add 1 to the last byte of the device's MAC address, which you noted before you began this procedure.

5. Define the format of the vendor-specific information for DHCP option 43 in the dhcpd.conf file.

Here is an example of an ISC DHCP 4.2 server `dhcpd.conf` file:

```
option space NEW_OP; option;
option NEW_OP.image-file-name code 0 = text;
option NEW_OP.config-file-name code 1 = text;
option NEW_OP.image-file-type code 2 = text;
option NEW_OP.transfer-mode code 3 = text;
option NEW_OP.alt-image-file-name code 4 = text;
option NEW_OP-encapsulation code 43 = encapsulate NEW_OP;
```

6. Configure the following DHCP option 43 suboptions:

- Suboption 00: The name of the software image file to install



NOTE: When the DHCP server cannot use suboption 00, configure the image file using suboption 04. If both suboption 00 and suboption 4 are defined, suboption 04 is ignored.

```
option NEW_OP.image-file-name
"/dist/images/jinstall-ex-4300-13.2R1.1-domestic-signed.tgz";
```

- Suboption 01: The name of the configuration file to install



NOTE: On EX4300 and QFX5100 devices running Enhanced Layer 2 Software, and QFX5100 devices running a Junos OS image that contains enhanced automation, you can specify the name of a script file or a configuration file. ZTP determines if the file is a script file based on the first line that is included in the file. If the first line contains `#!` characters followed by an interpreter path, ZTP determines that the file is a script file, and executes the script file with the specified interpreter path. In order for a script to execute, the script file must provide the ability to fetch and load a valid configuration file on the device during the ZTP process.

The following list provides the types of scripts and their associated interpreter paths:

- Shell script interpreter path: `#!/bin/sh`
- SLAX script interpreter path: `#!/usr/libexec/ui/cscript`
- Python script interpreter path: `#!/usr/bin/python`

Unsigned Python scripts are only supported on limited platforms, such as the QFX5100 device. If you try to execute unsigned Python scripts on devices that do not provide support, error messages will be issued.

If the file does not contain special characters (`#!`), ZTP determines that the file is a configuration file and loads the configuration file.

```
option NEW_OP.config-file-name "/dist/config/jn-switch35.config";
```

- Suboption 02: The symbolic link to the software image file to install

```
option NEW_OP.image-file-type "symlink";
```



NOTE: If you do not specify suboption 2, the Zero Touch Provisioning process handles the software image as a filename, not a symbolic link.

- Suboption 03: The transfer mode that the device uses to access the TFTP/FTP/HTTP server

```
option NEW_OP.transfer-mode "ftp";
```



NOTE: If suboption 03 is not configured, TFTP becomes the transfer mode by default.

- Suboption 04: The name of the software image file to install



NOTE: When the DHCP server cannot use suboption 00, configure the image file using suboption 04. If both suboption 00 and suboption 4 are defined, suboption 04 is ignored.



NOTE: DHCP option 43 suboptions 05 through 255 are reserved.

```
option NEW_OP.alt-image-file-name
"/dist/images/jinstall-ex-4300-13.2R1.1-domestic-signed.tgz";
```

7.



NOTE: You must configure either option 150 or option 66. If you configure both option 150 and option 66, option 150 takes precedence, and option 66 is ignored. Also, make sure you specify an IP address, not a hostname, because name resolution is not supported.

Configure DHCP option 150 to specify the IP address of the FTP, HTTP, or TFTP server.

```
option option-150 code 150 "10.100.31.71";
```

8. Configure DHCP option 66 to specify the IP address of the FTP, HTTP, or TFTP server.

```
option tftp-server-name "10.100.31.71";
```

9. (Optional) Configure DHCP option 7 to specify one or more system log (syslog) servers.

```
option log-servers 10.100.31.72;
```

10. (Optional) Configure DHCP option 42 to specify one or more NTP servers.

```
option ntp-servers 10.100.31.73;
```

11. (Optional) Configure DHCP option 12 to specify the hostname of the device.

```
option hostname "jn-switch35";
```

The following sample configuration shows the DHCP options you just configured:

```
host jn-switch35 {
  hardware ethernet ac:4b:c8:29:5d:02;
  fixed-address 10.100.31.36;
  option tftp-server-name "10.100.31.71";
  option host-name "jn-switch35";
  option log-servers 10.100.31.72;
  option ntp-servers 10.100.31.73;
  option NEW_OP.image-file-name
    "/dist/images/jinstall-ex-4300-13.2R1.1-domestic-signed.tgz";
  option NEW_OP.transfer-mode "ftp";
  option NEW_OP.config-file-name "/dist/config/jn-switch35.config";
}
```

Based on the DHCP options you just configured, the following statements are appended to the Junos OS configuration file (for example, `jn-switch35.config`):

```
system {
  host-name jn-switch35;
  syslog {
    host 10.100.31.72 {
      any any;
    }
  }
  ntp {
    server 10.100.31.73;
  }
}
```

12. Connect the device to the network that includes the DHCP server and the FTP, HTTP, or TFTP server.
13. Boot the device with the default configuration.
14. Monitor the ZTP process by looking at the following log files.



NOTE: When SLAX (live operating system based on Linux) scripts are issued, the `op-script.log` and `event-script.log` files are produced.

- `/var/log/dhcp_logfile`
- `/var/log/event-script.log`
- `/var/log/image_load_log`
- `/var/log/messages`
- `/var/log/op-script.log`
- `/var/log/script_output`

You can also monitor the ZTP process by looking at error messages and issuing operational commands. See [“Monitoring Zero Touch Provisioning” on page 209](#) for more information.

- Related Documentation**
- [Understanding Zero Touch Provisioning on page 197](#)
 - [Understanding NTP Time Servers on page 131](#)
 - [Op Script Overview](#)
 - [Monitoring Zero Touch Provisioning on page 209](#)
 - [Understanding DHCP Services for Switches on page 73](#)
 - [Reverting to the Default Factory Configuration by Using the request system zeroize Command on page 28](#)

Monitoring Zero Touch Provisioning

You can use the console and operational commands to monitor Zero Touch Provisioning.

1. [Using the Console to Monitor Zero Touch Provisioning on page 209](#)
2. [Using System Log Alerts to Monitor Zero Touch Provisioning on page 210](#)
3. [Using Error Messages to Monitor Zero Touch Provisioning on page 210](#)
4. [Using System Log Files to Monitor Zero Touch Provisioning on page 210](#)
5. [Using the show dhcp client binding Command on page 211](#)
6. [Using the show dhcp client statistics Command on page 211](#)

Using the Console to Monitor Zero Touch Provisioning

The following Zero Touch Provisioning (ZTP) activities are displayed on the console during the ZTP process:

- Starting and ending times of ZTP process.
- Lists of bound and unbound DHCP client interfaces.
- DHCP options that DHCP servers send to DHCP clients.
- Logs indicating which interfaces are used for ZTP.
- ZTP parameters that DHCP clients obtain from DHCP servers.
- File names of configuration and image files, names of file servers, protocols used to fetch files, and times when DHCP servers fetch configuration and image files.
- Failure states caused by files not being on servers, or unreachable servers, and time outs.
- Number of attempts made, and number of attempts remaining, for retry in current ZTP cycle.
- Completion of file transfers.
- Installation, reboot, and state of ZTP process.
- Internal state errors and termination of ZTP process.
- Logs for when default routes were added or deleted.

Using System Log Alerts to Monitor Zero Touch Provisioning

Purpose In this example, the system log alert alerts you that the auto-image upgrade will start.

Action Use the following system log alert to monitor the auto-image upgrade process.

```
"ALERT:Auto-image upgrade will start. This can terminate config CLI session(s).  
Modified configuration will be lost. To stop Auto-image, in CLI do the  
following: 'edit; delete chassis auto-image-upgrade; commit'."
```

```
"Checking whether image upgrade is already invoked"
```

Meaning This system log alert indicates that the auto-image upgrade will start, and provides information on how to stop the auto-image upgrade process.

Using Error Messages to Monitor Zero Touch Provisioning

Purpose Error messages provide information on which DHCP options are not configured.

Action Use the information in the following error message to find out which DHCP options are not configured.

```
"DHCP Log Server Option"  
"DHCP Host Name Option"  
"DHCP NTP Server Option"
```

Meaning The error message indicates that the DHCP log server, hostname, and NTP server options are not configured.

Using System Log Files to Monitor Zero Touch Provisioning

Purpose System log files provide information on the state of the auto-upgrade process, lists of bound and unbound DHCP client interfaces, IP addresses of file servers, names and locations of image and configuration files, and successful and failed attempts at fetching configuration and image files.

Action Use the information in the following system log files to monitor the auto-upgrade process.

```
Auto Image Upgrade: Start fetching config-file file from server 1.1.1.1 through  
irb using ftp
```

```
Auto Image Upgrade: Tried [2] attempts to fetch config-file file from server  
1.1.1.1 through irb. Summary: "Retrieving /config-file  
:: Failed to open file.". To retry [4] times.
```

```
Auto Image Upgrade: Tried [4] attempts to fetch config-file file from server  
1.1.1.1 through irb. Summary: "Retrieving /config-fileconfig-file  
:: Failed to open file.". To retry [2] times.
```

```
Auto Image Upgrade: Tried [6] attempts to fetch config-file file from server  
1.1.1.1 through irb. Summary: "Retrieving /config-file  
:: Failed to open file.". To retry [0] times.
```

Auto Image Upgrade: All [6] attempts to fetch config-file file from server 1.1.1.1 through irb FAILED. Start retry again in few minutes.

Meaning These system log files indicate that there were six failed attempts to fetch the configuration file from the file server, the IP address of the file server, the DHCP client interface name, and the number of times the retry process occurred.

Using the show dhcp client binding Command

Purpose Issue the **show dhcp client binding** command to display DHCP client binding information

Action Issue the **show dhcp client binding** command to display the IP address of the DHCP client, the hardware address of the DHCP client, number of seconds in which the DHCP client's IP address lease expires, state of the DHCP client IP address in the binding table, and the name of the interface that has active client bindings.

show dhcp client binding

```
user@switch# show dhcp client binding
IP address      Hardware address Expires   State      Interface
0.0.0.0         00:22:83:2a:db:dc 0         SELECTING  irb.0
6.6.6.13        00:22:83:2a:db:dd 49201     BOUND      vme.0
0.0.0.0         00:22:83:2a:db:df 0         SELECTING  xe-0/0/0.0
0.0.0.0         00:22:83:2a:db:e0 0         SELECTING  xe-0/0/1.0
```

Meaning The output of this command shows that there is one client interface that is bound, and that there are three interfaces that are receiving DHCP offers from the DHCP server.

Using the show dhcp client statistics Command

Purpose Issue the **show dhcp client statistics** command to display DHCP client statistics.

Action Issue the **show dhcp client statistics** command to display DHCP client statistics, such as the number of packets dropped, and the number DHCP and BOOTP messages sent and received.

show dhcp client statistics

```
user@switch# show dhcp client statistics
Packets dropped:
  Total          14
  Send error     14
Messages received:
  BOOTREPLY      5
  DHCPOFFER      1
  DHCPACK        4
  DHCPNAK        0
  DHCPFORCERENEW 0
Messages sent:
  BOOTREQUEST    6751
  DHCPDECLINE    0
  DHCPDISCOVER   6747
  DHCPREQUEST    4
  DHCPINFORM     0
  DHCPRELEASE    0
```

DHCPRENEW	0
DHCPREBIND	0

Meaning The output of this command displays how many packets were dropped with errors, the number of BOOTREPLY and DHCPOFFER messages that were received, and the number of BOOTREQUEST and DHCPREQUEST messages that were sent.

Related Documentation

- [Understanding Zero Touch Provisioning on page 197](#)
- [Configuring Zero Touch Provisioning on page 204](#)

PART 15

Configuration Statements and Operational Commands

- [Basic System Management Configuration Statements on page 215](#)
- [DHCP Configuration Statements on page 251](#)
- [ICMP Configuration Statement on page 259](#)
- [Junos OS With Upgraded FreeBSD Operational Commands on page 261](#)
- [Login Classes Configuration Statements on page 269](#)
- [NTP Configuration Statements on page 283](#)
- [Password Configuration Statements on page 293](#)
- [PTP Configuration Statement on page 311](#)
- [RADIUS and TACACS+ Accounting Configuration Statements on page 313](#)
- [Autoinstallation Operational Commands on page 323](#)
- [Basic System Management Operational Commands on page 327](#)
- [CLI Operational Commands on page 427](#)
- [Licensing Operational Commands on page 437](#)
- [NTP Operational Commands on page 449](#)
- [PTP Operational Command on page 455](#)
- [Routine Monitoring Operational Commands on page 457](#)
- [Standard Software Installation and Upgrade Operational Commands on page 1275](#)

CHAPTER 16

Basic System Management Configuration Statements

- [arp \(System\) on page 217](#)
- [auxiliary on page 218](#)
- [checksum on page 219](#)
- [compress-configuration-files \(System\) on page 220](#)
- [console \(Physical Port\) on page 221](#)
- [default-address-selection on page 222](#)
- [domain-name on page 223](#)
- [domain-search on page 223](#)
- [ethernet \(Alarm\) on page 224](#)
- [host-name on page 224](#)
- [icmpv4-rate-limit on page 225](#)
- [internet-options on page 226](#)
- [link-down on page 227](#)
- [location on page 228](#)
- [login-alarms on page 229](#)
- [management-ethernet \(Alarm\) on page 230](#)
- [max-configurations-on-flash on page 231](#)
- [name-server on page 231](#)
- [no-ping-record-route on page 232](#)
- [no-ping-time-stamp on page 232](#)
- [no-redirects \(IPv4 Traffic\) on page 233](#)
- [optional on page 234](#)
- [ports on page 234](#)
- [retry on page 235](#)
- [saved-core-context on page 236](#)
- [saved-core-files on page 236](#)

- [source-port \(Port Addresses\) on page 237](#)
- [static-host-mapping on page 238](#)
- [system on page 239](#)
- [time-format on page 244](#)
- [time-zone on page 245](#)
- [traceoptions \(Layer 2 Learning\) on page 247](#)
- [uid on page 249](#)
- [use-imported-time-zones on page 249](#)

arp (System)

Syntax	<pre>arp { aging-timer <i>minutes</i>; gratuitous-arp-delay <i>seconds</i>; gratuitous-arp-on-ifup; interfaces { <i>interface-name</i> { aging-timer <i>minutes</i>; } } passive-learning; purging; }</pre> <p>For EX-Series switches:</p> <pre>arp { aging-timer <i>minutes</i>; }</pre>
Hierarchy Level	[edit system]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Specify ARP options. You can enable backup VRRP routers to learn ARP requests for VRRP-IP to VRRP-MAC address translation. You can also set the time interval between ARP updates.</p> <p>For EX-Series switches, set only the time interval between ARP updates.</p>
Options	<p>aging-timer—Time interval in minutes between ARP updates. In environments where the number of ARP entries to update is high (for example, on routers only, metro Ethernet environments), increasing the time between updates can improve system performance.</p> <p>passive-learning (QFX-Series only)—Configure switches to learn the ARP mappings (IP-to-MAC address) for hosts sending the requests.</p> <p>Default: 20 minutes</p> <p>Range: 1 to 240 minutes</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Junos OS ARP Learning and Aging Options for Mapping IPv4 Network Addresses to MAC Addresses</i> • <i>Junos OS Network Interfaces Library for Routing Devices</i>

- [Junos OS System Basics Configuration Guide](#) .


auxiliary

Syntax	<pre>auxiliary { disable; insecure; type <i>terminal-type</i>; }</pre>
Hierarchy Level	[edit system ports]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the characteristics of the auxiliary port.
Default	The auxiliary port is disabled.
Options	<p>disable—Disable the port.</p> <p>insecure—Disable superuser access or root logins to establish a terminal connection.</p> <p>type <i>terminal-type</i>—Type of terminal that is connected to the port.</p> <p>Range: <code>ansi</code>, <code>vt100</code>, <code>small-xterm</code>, <code>xterm</code></p> <p>Default: The terminal type is unknown, and the user is prompted for the terminal type.</p>
Required Privilege Level	<p><code>system</code>—To view this statement in the configuration.</p> <p><code>system-control</code>—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Console and Auxiliary Port Properties on page 15

checksum

Syntax	<code>checksum (md5 sha-256 sha1) <i>hash</i>;</code>
Hierarchy Level	[edit event-options event-script file <i>filename</i>], [edit system scripts commit file <i>filename</i>],
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	For Junos commit scripts and op scripts, specify the MD5, SHA-1, or SHA-256 checksum hash. When it executes a local event or commit script, the Junos OS verifies the authenticity of the script by using the configured checksum hash.
Options	<p>md5 <i>hash</i>—MD5 checksum of this script.</p> <p>sha-256 <i>hash</i>—SHA-256 checksum of this script.</p> <p>sha1 <i>hash</i>—SHA-1 checksum of this script.</p>
Required Privilege Level	<p>maintenance—To view this statement in the configuration.</p> <p>maintenance-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Checksum Hashes for a Commit Script</i> • <i>Configuring Checksum Hashes for an Event Script</i> • <i>Configuring Checksum Hashes for an Op Script</i> • file checksum md5 on page 344 • file checksum sha-256 on page 346 • file checksum sha1 on page 345

compress-configuration-files (System)

Syntax	(compress-configuration-files no-compress-configuration-files);
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Compress the current operational configuration file. The file is stored in the file juniper.conf , in the /config file system, along with the last three committed versions of the configuration. However, with large networks, the current configuration file might exceed the available space in the /config file system. Compressing the current configuration file allows the file to fit in the file system, typically reducing the size of the file by 90 percent. The current configuration file is compressed on the second commit of the configuration after the first commit is made to include the compress-configuration-files statement.
<div> NOTE: We recommend that you enable compression of the configuration files to minimize the amount of disk space that they require.</div>	
Default	The current operational configuration file is uncompressed.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Compressing the Current Configuration File on page 14

console (Physical Port)

Syntax	<pre>console { disable; insecure; log-out-on-disconnect; type <i>terminal-type</i>; }</pre>
Hierarchy Level	[edit system ports]
Release Information	<p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Configure the characteristics of the console port.
Default	The console port is enabled and its speed is 9600 baud.
Options	<p>disable—Disable console login connections.</p> <p>insecure—Disable root login connections to the console and auxiliary ports. Configuring the console port as insecure also prevents superusers and anyone with a user identifier (UID) of 0 from establishing terminal connections in multiuser mode. This option can be used to prevent a user from attempting password recovery by booting into single-user mode, if the user does not know the root password.</p> <p>log-out-on-disconnect—Log out the session when the data carrier on the console port is lost.</p> <p>type <i>terminal-type</i>—Type of terminal that is connected to the port: ansi, vt100, small-xterm, or xterm.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Console and Auxiliary Port Properties on page 15

default-address-selection

Syntax	default-address-selection;
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Use the loopback interface, lo0 , as the source address for all locally generated IP packets when the packet is sent through a routed interface, but not when the packet is sent through a local interface such as fxp0 . The lo0 interface is the interface to the switch's Routing Engine.
Default	<p>The default address is used as the source address for all locally generated IP packets on outgoing interfaces that are unnumbered. If an outgoing interface is numbered, the default address is chosen using the following sequence:</p> <ul style="list-style-type: none">• The primary address on the loopback interface lo0 that is <i>not</i> 127.0.0.1 is used.• The primary address for the primary interface or the preferred address (if configured) for the primary interface is used. <p>By default, the primary address on an interface is selected as the numerically lowest local address configured on the interface.</p> <p>An interface's <i>primary address</i> is used by default as the local address for broadcast and multicast packets sourced locally and sent out through the interface. An interface's <i>preferred address</i> is the default local address used for packets sourced by the local switch to destinations on the subnet. By default, the numerically lowest local address configured for the interface is chosen as the preferred address on the subnet.</p> <p>To configure a different primary address or preferred address, include the primary or preferred statement at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i>] or [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i>] hierarchy levels.</p>
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Junos OS to Select a Fixed Source Address for Locally Generated TCP/IP Packets on page 20• <i>Junos OS Network Interfaces Library for Routing Devices</i>

domain-name

Syntax	<code>domain-name <i>domain-name</i>;</code>
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the name of the domain in which the switch is located. This is the default domain name that is appended to hostnames that are not fully qualified.
Options	<i>domain-name</i> —Name of the domain.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Reaching a Domain Name System Server on page 26

domain-search

Syntax	<code>domain-search <i>domain-list</i>;</code>
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure a list of domains to be searched.
Options	<i>domain-list</i> —List of domain names to search. The list can contain up to 6 domain names, with a total of up to 256 characters.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Reaching a Domain Name System Server on page 26

ethernet (Alarm)

Syntax	ethernet { link-down (red yellow ignore); }
Hierarchy Level	[edit chassis alarm], [edit chassis interconnect-device <i>name</i> alarm], [edit chassis node-group <i>name</i> alarm]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure alarms for an Ethernet interface.
Options	The remaining statement is explained separately.—
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Understanding Alarms</i>• <i>Interface Alarm Messages</i>

host-name

Syntax	host-name <i>hostname</i> ;
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Set the hostname of the switch.
Options	<i>hostname</i> —Name of the switch.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Hostname of a Router or Switch by Using a Configuration Group on page 16

icmpv4-rate-limit

Syntax	icmpv4-rate-limit { bucket-size <i>seconds</i> ; packet-rate <i>pps</i> ; }
Hierarchy Level	[edit system internet-options]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure rate-limiting parameters for ICMPv4 messages sent.
Options	<p>bucket-size <i>seconds</i>—Number of seconds in the rate-limiting bucket. Range: 0 through 4294967295 seconds Default: 5</p> <p>packet-rate <i>pps</i>—Rate-limiting packets earned per second. Range: 0 through 4294967295 pps Default: 1000</p>
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Junos OS ICMPv4 Rate Limit for ICMPv4 Routing Engine Messages</i> • <i>Configuring Junos OS ICMPv6 Rate Limit for ICMPv6 Routing Engine Messages</i>

internet-options

Syntax	<pre>internet-options { icmpv4-rate-limit bucket-size <i>bucket-size</i> packet-rate <i>packet-rate</i>; source-port upper-limit <i>upper-limit</i>; }</pre>
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Configure system IP options to protect against certain types of denial-of-service (DoS) attacks.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Junos OS ICMPv4 Rate Limit for ICMPv4 Routing Engine Messages on page 86• Configuring Junos OS to Extend the Default Port Address Range on page 19

link-down

Syntax	link-down (red yellow ignore);
Hierarchy Level	[edit chassis alarm ethernet], [edit chassis alarm fibre-channel], [edit chassis interconnect-device <i>name</i> alarm ethernet], [edit chassis node-group <i>name</i> alarm fibre-channel]
Release Information	Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify either red, yellow, or ignore to display when the link is down.
Options	<p>red—Indicates that one or more hardware components have failed or exceeded temperature thresholds, or an alarm condition configured on an interface has triggered a critical warning.</p> <p>yellow—Indicates a noncritical condition on the device that, if left unchecked, might cause an interruption in service or degradation in performance. A yellow alarm condition requires monitoring or maintenance.</p> <p>ignore—Suppresses or ignores the alarm.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

location

Syntax	<pre>location { altitude <i>feet</i>; building <i>name</i>; country-code <i>code</i>; floor <i>number</i>; hcoord <i>horizontal-coordinate</i>; lata <i>service-area</i>; latitude <i>degrees</i>; longitude <i>degrees</i>; npa-nxx <i>number</i>; postal-code <i>postal-code</i>; rack <i>number</i>; vcoord <i>vertical-coordinate</i>; }</pre>
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the system location.
Options	<p>altitude <i>feet</i>—Number of feet above sea level.</p> <p>building <i>name</i>—Name of the building. The name of the building can be 1 to 28 characters in length. If the string contains spaces, enclose it in quotation marks (" ").</p> <p>country-code <i>code</i>—Two-letter country code.</p> <p>floor <i>number</i>—Floor in the building.</p> <p>hcoord <i>horizontal-coordinate</i>—Bellcore Horizontal Coordinate.</p> <p>lata <i>service-area</i>—Long-distance service area.</p> <p>latitude <i>degrees</i>—Latitude in degree format.</p> <p>longitude <i>degrees</i>—Longitude in degree format.</p> <p>npa-nxx <i>number</i>—First six digits of the phone number (area code and exchange).</p> <p>postal-code <i>postal-code</i>—Postal code.</p> <p>rack <i>number</i>—Rack number.</p> <p>vcoord <i>vertical-coordinate</i>—Bellcore Vertical Coordinate.</p>
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.

- Related Documentation**
- [Specifying the Physical Location of the Switch on page 29](#)

login-alarms

Syntax	login-alarms;
Hierarchy Level	[edit system login class <i>class-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Show system alarms automatically when an admin user logs in to the router or switch.
Options	<i>class-name</i> —Login class name.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring System Alarms to Appear Automatically Upon Login on page 20

management-ethernet (Alarm)

Syntax	management-ethernet { link-down (red yellow ignore); }
Hierarchy Level	[edit chassis alarm], [edit chassis interconnect-device <i>name</i> alarm], [edit chassis node-group <i>name</i> alarm]
Release Information	Statement introduced in Junos OS Release 12.2 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure alarms for a management Ethernet interface.



NOTE: If you configure a yellow alarm on the Interconnect device, it will be handled as a red alarm.

Options	The remaining statement is explained separately.—
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Understanding Alarms</i>• <i>Interface Alarm Messages</i>

max-configurations-on-flash

Syntax	<code>max-configurations-on-flash <i>number</i>;</code>
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify the number of configurations stored on the internal fixed media storage (for example, USB device).
Options	<i>number</i> —The number of configurations stored on the CompactFlash card. Range: 0 through 49. The most recently saved configuration is number 0, and the oldest saved configuration is number 49.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Saving a Configuration to a File on page 1331• Setting or Deleting the Rescue Configuration on page 1345• Uploading a Configuration File on page 1332• Uploading a Configuration File

name-server

Syntax	<code>name-server { <i>address</i>; }</code>
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure one or more Domain Name System (DNS) name servers.
Options	<i>address</i> —Address of the name server. To configure multiple name servers, include multiple <i>address</i> options.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a DNS Name Server for Resolving a Hostname into Addresses on page 15

no-ping-record-route

Syntax	no-ping-record-route;
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 9.4. Statement introduced in Junos OS Release 9.4 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the Junos OS to disable the reporting of the IP address in ping responses.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Junos OS to Disable the Reporting of IP Address and Timestamps in Ping Responses on page 19

no-ping-time-stamp

Syntax	no-ping-time-stamp;
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 9.4. Statement introduced in Junos OS Release 9.4 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the Junos OS to disable the recording of timestamps in ping responses.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Junos OS to Disable the Reporting of IP Address and Timestamps in Ping Responses on page 19

no-redirects (IPv4 Traffic)

Syntax	no-redirects;
Hierarchy Level	[edit system], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 12.3 for EX Series switches. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Stop protocol redirect messages for IPv4 traffic from being sent on the entire switch or on an interface on the router or switch.</p> <p>To disable the sending of protocol redirect messages for the entire router or switch, include the no-redirects statement at the [edit system] hierarchy level.</p> <p>To disable the sending of protocol redirect messages on a specific interface, include the no-redirects statement at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>] hierarchy level.</p>
Default	The router or switch sends redirect messages.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Junos OS to Disable Protocol Redirect Messages on the Router or Switch on page 18 • <i>Understanding the Protocol Redirect Mechanism on EX Series Switches</i> • <i>Configuring Junos OS to Disable Sending Protocol Redirect Messages on EX Series Switches (CLI Procedure)</i> • <i>Junos OS Network Interfaces Library for Routing Devices</i>

optional

Syntax	optional;
Hierarchy Level	[edit system scripts commit file <i>filename</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	For Junos OS commit scripts, allow a commit operation to succeed even if the script specified in the file statement is missing from the /var/db/scripts/commit directory on the router.
Required Privilege Level	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Controlling Execution of Commit Scripts During Commit Operations</i>

ports

Syntax	<pre>ports { auxiliary { disable; insecure; type <i>terminal-type</i>; } console { disable; insecure; log-out-on-disconnect; type <i>terminal-type</i>; } }</pre>
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the properties of the console and auxiliary ports. The ports are located on the craft interface. See the switch hardware documentation for port locations. The remaining statements are explained separately.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Console and Auxiliary Port Properties on page 15

retry

Syntax	<code>retry number;</code>
Hierarchy Level	[edit system radius server <i>server-address</i>], [edit system accounting destination radius server <i>server-address</i>]
Release Information	Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Number of times the router or switch is allowed to try to contact a RADIUS authentication or accounting server.
Options	<i>number</i> —Number of retries allowed for contacting a RADIUS server. Range: 1 through 10 Default: 3



NOTE: The [edit system accounting] hierarchy is not available on QFabric systems.

Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring RADIUS Authentication (QFX Series or OCX Series) on page 1873 • Configuring RADIUS Accounting • timeout on page 321

saved-core-context

Syntax	(saved-core-context no-saved-core-context);
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure whether the switch saves core files generated by internal Junos OS processes, along with contextual information (system log files and a copy of the current configuration): <ul style="list-style-type: none">• saved-core-context—The switch saves each core file and its associated context in a compressed tar file named <code>/var/tmp/process-name.core.core-number.tgz</code>.• no-saved-core-context—The switch does not save core files and their associated context.
Default	The switch saves core files.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Saving Core Files from Junos OS Processes</i>• saved-core-files on page 236

saved-core-files

Syntax	saved-core-files <i>number</i> ;
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Save core files generated by internal Junos OS processes, but not the associated contextual information (configuration and system log files).
Options	<i>number</i> —Maximum number of core files to save. Range: 1 through 10
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Saving Core Files from Junos OS Processes</i>• saved-core-context on page 236

source-port (Port Addresses)

Syntax	source-port upper-limit <upper-limit>;
Hierarchy Level	[edit system internet-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the range of port addresses.
Options	upper-limit <i>upper-limit</i> —(Optional) The range of port addresses and can be a value from 5000 through 65,355.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Junos OS to Extend the Default Port Address Range on page 19

static-host-mapping

Syntax	<pre>static-host-mapping { hostname { alias [<i>alias</i>]; inet [<i>address</i>]; sysid <i>system-identifier</i>; } }</pre>
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Map a hostname to one or more IP addresses and aliases, and configure an International Organization for Standardization (ISO) system identifier (system ID).
Options	<p>alias <i>alias</i>—Alias for the hostname.</p> <p>hostname—Fully qualified hostname.</p> <p>inet <i>address</i>—IP address. You can specify one or more IP addresses for the host.</p> <p>sysid <i>system-identifier</i>—ISO system identifier (system ID). This is the 6-byte portion of the Intermediate System-to-Intermediate System (IS-IS) network service access point (NSAP). We recommend that you use the host's IP address represented in binary-coded decimal (BCD) format. For example, the IP address 208.197.169.18 is 2081.9716.9018 in BCD.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring the Hostname of a Router or Switch by Using a Configuration Group on page 16

system

```
Syntax  system {
    accounting {
        events [ login change-log interactive-commands ];
        destination {
            radius {
                server {
                    server-address {
                        accounting-port port-number;
                        retry number;
                        secret password;
                        source-address address;
                        timeout seconds;
                    }
                }
            }
        }
        tacplus {
            server {
                server-address {
                    port port-number;
                    secret password;
                    single-connection;
                    timeout seconds;
                }
            }
        }
    }
    archival {
        configuration {
            archive-sites {
                ftp://<username>:<password>@<host>:<port>/<url-path>;
                ftp://<username>:<password>@<host>:<port>/<url-path>;
            }
            transfer-interval interval;
            transfer-on-commit;
        }
    }
    arp {
        aging-timer minutes;
        interfaces;
    }
    authentication-order [ authentication-methods ];
    (compress-configuration-files | no-compress-configuration-files);
    default-address-selection;
    domain-name domain-name;
    domain-search [ domain-list ];
    host-name hostname;
    internet-options {
        icmpv4-rate-limit bucket-size bucket-size packet-rate packet-rate;
        source-port upper-limit <upper-limit>;
    }
    location {
```

```
altitude feet;  
building name;  
country-code code;  
floor number;  
hcoord horizontal-coordinate;  
lata service-area;  
latitude degrees;  
longitude degrees;  
npa-nxx number;  
postal-code postal-code;  
rack number;  
vcoord vertical-coordinate;  
}  
login {  
  announcement text;  
  class class-name {  
    access-end;  
    access-start;  
    allow-configuration "regular-expression";  
    allowed-days "regular-expression";  
    deny-commands "regular-expression";  
    deny-configuration "regular-expression";  
    idle-timeout minutes;  
    login-tip;  
    permissions [ permissions ];  
  }  
  message text;  
  password {  
    change-type (set-transitions | character-set);  
    format (md5 | sha1 | des);  
    maximum-length length;  
    minimum-changes number;  
    minimum-length length;  
  }  
  retry-options {  
    backoff-factor seconds;  
    backoff-threshold number;  
    minimum-time seconds;  
    tries-before-disconnect number;  
  }  
  user username {  
    authentication {  
      (encrypted-password "password" | plain-text-password);  
      load-key-file URL;  
      remote-debug-permission (qfabric-admin | qfabric-operator | qfabric-user);  
      ssh-rsa "public-key";  
      ssh-dsa "public-key";  
    }  
    uid uid-value;  
    class class-name;  
    full-name complete-name;  
  }  
}  
name-server {  
  address;  
}
```



```

no-multicast-echo;
no-redirects;
no-ping-record-route;
no-ping-time-stamp;
ntp {
    authentication-key number type type value password;
    serveraddress <key key-number> <version value> <prefer>;
}
ports {
    auxiliary {
        disable;
        insecure;
        type terminal-type;
    }
    console {
        disable;
        insecure;
        log-out-on-disconnect;
        type terminal-type;
    }
}
radius-server server-address {
    accounting-port port-number;
    port number;
    retry number;
    secret password;
    source-address source-address;
    timeout seconds;
}
radius-options {
    password-protocol mschap-v2;
}
attributes {
    nas-ip-address ip-address;
}
root-authentication {
    (encrypted-password "password" | plain-text-password);
    ssh-rsa "public-key";
    ssh-dsa "public-key";
}
(saved-core-context | no-saved-core-context);
saved-core-files saved-core-files;
services {
    finger {
        connection-limit limit;
        rate-limit limit;
    }
    flow-tap-dtcp {
        ssh {
            connection-limit limit;
            rate-limit limit;
        }
    }
}
ftp {
    connection-limit limit;
    rate-limit limit;
}

```

```
}
service-deployment {
  servers server-address {
    port port-number;
  }
  source-address source-address;
}
ssh {
  root-login (allow | deny | deny-password);
  protocol-version [v1 v2];
  connection-limit limit;
  rate-limit limit;
}
telnet {
  connection-limit limit;
  rate-limit limit;
}
web-management {
  http {
    interfaces [ interface-names ];
    port port;
  }
  https {
    interfaces [ interface-names ];
    local-certificate name;
    port port;
  }
  session {
    idle-timeout [ minutes ];
    session-limit [ session-limit ];
  }
}
xnm-clear-text {
  connection-limit limit;
  rate-limit limit;
}
xnm-ssl {
  connection-limit limit;
  local-certificate name;
  rate-limit limit;
}
}
static-host-mapping {
  hostname {
    alias [ alias ];
    inet [ address ];
    sysid system-identifier;
  }
}
syslog {
  archive {
    files number;
    size maximum-file-size;
    start-time "YYYY-MM-DD.hh:mm";
    transfer-interval minutes;
    (world-readable | no-world-readable);
  }
}
```

```

}
console {
    facility severity;
}
file filename {
    archive {
        files number;
        size maximum-file-size;
        start-time "YYYY-MM-DD.hh:mm";
        transfer-interval minutes;
        (world-readable | no-world-readable);
    }
    explicit-priority;
    facility severity;
    match "regular-expression";
    structured-data {
        brief;
    }
}
host (hostname | other-routing-engine | scc-master) {
    explicit-priority;
    facility-override facility;
    facility severity;
    log-prefix string;
    match "regular-expression";
}
source-address source-address;
time-format (millisecond | year | year millisecond);
user (username | *) {
    facility severity;
    match "regular-expression";
}
}
tacplus-options {
    service-name service-name;
    (no-cmd-attribute-value | exclude-cmd-attribute);
}
tacplus-server server-address {
    port
    secret password;
    single-connection;
    source-address source-address;
    timeout seconds;
}
time-zone (GMThour-offset | time-zone);
}
tracing {
    destination-override {
        syslog host;
    }
}
use-imported-time-zones;
}

```

Hierarchy Level [edit]

Release Information Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure system management properties.



NOTE: The `radius-server source-address` and `radius-options` statements are not available on the QFabric system.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

time-format

Syntax `time-format (year | millisecond | year millisecond);`

Hierarchy Level [edit system syslog]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Include the year, the millisecond, or both, in the timestamp on every standard-format system log message. The additional information is included for messages directed to each destination configured by a `file`, `console`, or `user` statement at the [edit system syslog] hierarchy level, but not to destinations configured by a `host` statement.

Default The timestamp specifies the month, date, hour, minute, and second when the message was logged—for example, **Aug 21 12:36:30**.



NOTE: When the `structured-data` statement is included at the [edit system syslog file *filename*] hierarchy level, this statement is ignored for the file.

Options `millisecond`—Include the millisecond in the timestamp.

`year`—Include the year in the timestamp.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- [Including the Year or Millisecond in Timestamps on page 22](#)

time-zone

Syntax	<code>time-zone (GMT <i>hour-offset</i> <i>time-zone</i>);</code>
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Set the local time zone. To have the time zone change take effect for all processes running on the switch, you must reboot the switch.
Default	UTC
Options	<p>GMT <i>hour-offset</i>—Set the time zone relative to UTC time.</p> <p>Range: -14 through +12</p> <p>Default: 0</p> <p><i>time-zone</i>—Specify the time zone as UTC, which is the default time zone, or as a string such as PDT (Pacific Daylight Time), or use one of the following continents and major cities:</p> <p>Africa/Abidjan, Africa/Accra, Africa/Addis_Ababa, Africa/Algiers, Africa/Asmera, Africa/Bamako, Africa/Bangui, Africa/Banjul, Africa/Bissau, Africa/Blantyre, Africa/Brazzaville, Africa/Bujumbura, Africa/Cairo, Africa/Casablanca, Africa/Ceuta, Africa/Conakry, Africa/Dakar, Africa/Dar_es_Salaam, Africa/Djibouti, Africa/Douala, Africa/El_Aaiun, Africa/Freetown, Africa/Gaborone, Africa/Harare, Africa/Johannesburg, Africa/Kampala, Africa/Khartoum, Africa/Kigali, Africa/Kinshasa, Africa/Lagos, Africa/Libreville, Africa/Lome, Africa/Luanda, Africa/Lubumbashi, Africa/Lusaka, Africa/Malabo, Africa/Maputo, Africa/Maseru, Africa/Mbabane, Africa/Mogadishu, Africa/Monrovia, Africa/Nairobi, Africa/Ndjamena, Africa/Niamey, Africa/Nouakchott, Africa/Ouagadougou, Africa/Porto-Novo, Africa/Sao_Tome, Africa/Timbuktu, Africa/Tripoli, Africa/Tunis, Africa/Windhoek</p> <p>America/Adak, America/Anchorage, America/Anguilla, America/Antigua, America/Aruba, America/Asuncion, America/Barbados, America/Belize, America/Bogota, America/Boise, America/Buenos_Aires, America/Caracas, America/Catamarca, America/Cayenne, America/Cayman, America/Chicago, America/Cordoba, America/Costa_Rica, America/Cuiaba, America/Curacao, America/Dawson, America/Dawson_Creek, America/Denver, America/Detroit, America/Dominica, America/Edmonton, America/El_Salvador, America/Ensenada, America/Fortaleza, America/Glace_Bay, America/Godthab, America/Goose_Bay, America/Grand_Turk, America/Grenada, America/Guadeloupe, America/Guatemala, America/Guayaquil, America/Guyana, America/Halifax, America/Havana, America/Indiana/Knox, America/Indiana/Marengo, America/Indiana/Vevay, America/Indianapolis, America/Inuvik, America/Iqaluit, America/Jamaica, America/Jujuy, America/Juneau, America/La_Paz, America/Lima, America/Los_Angeles, America/Louisville, America/Maceio, America/Managua, America/Manaus, America/Martinique, America/Mazatlan, America/Mendoza, America/Menominee, America/Mexico_City, America/Miquelon, America/Montevideo, America/Montreal, America/Montserrat, America/Nassau, America/New_York, America/Nipigon, America/Nome, America/Noronha, America/Panama, America/Pangnirtung, America/Paramaribo, America/Phoenix, America/Port-au-Prince, America/Port_of_Spain, America/Porto_Acre, America/Puerto_Rico, America/Rainy_River,</p>

America/Rankin_Inlet, America/Regina, America/Rosario, America/Santiago, America/Santo_Domingo, America/Sao_Paulo, America/Scoresbysund, America/Shiprock, America/St_Johns, America/St_Kitts, America/St_Lucia, America/St_Thomas, America/St_Vincent, America/Swift_Current, America/Tegucigalpa, America/Thule, America/Thunder_Bay, America/Tijuana, America/Tortola, America/Vancouver, America/Whitehorse, America/Winnipeg, America/Yakutat, America/Yellowknife

Antarctica/Casey, Antarctica/DumontDURville, Antarctica/Mawson, Antarctica/McMurdo, Antarctica/Palmer, Antarctica/South_Pole

Arctic/Longyearbyen

Asia/Aden, Asia/Alma-Ata, Asia/Amman, Asia/Anadyr, Asia/Aqttau, Asia/Aqtobe, Asia/Ashkhabad, Asia/Baghdad, Asia/Bahrain, Asia/Baku, Asia/Bangkok, Asia/Beirut, Asia/Bishkek, Asia/Brunei, Asia/Calcutta, Asia/Chungking, Asia/Colombo, Asia/Dacca, Asia/Damascus, Asia/Dubai, Asia/Dushanbe, Asia/Gaza, Asia/Harbin, Asia/Hong_Kong, Asia/Irkutsk, Asia/Ishigaki, Asia/Jakarta, Asia/Jayapura, Asia/Jerusalem, Asia/Kabul, Asia/Kamchatka, Asia/Karachi, Asia/Kashgar, Asia/Katmandu, Asia/Krasnoyarsk, Asia/Kuala_Lumpur, Asia/Kuching, Asia/Kuwait, Asia/Macao, Asia/Magadan, Asia/Manila, Asia/Muscat, Asia/Nicosia, Asia/Novosibirsk, Asia/Omsk, Asia/Phnom_Penh, Asia/Pyongyang, Asia/Qatar, Asia/Rangoon, Asia/Riyadh, Asia/Saigon, Asia/Seoul, Asia/Shanghai, Asia/Singapore, Asia/Taipei, Asia/Tashkent, Asia/Tbilisi, Asia/Tehran, Asia/Thimbu, Asia/Tokyo, Asia/Ujung_Pandang, Asia/Ulan_Bator, Asia/Urumqi, Asia/Vientiane, Asia/Vladivostok, Asia/Yakutsk, Asia/Yekaterinburg, Asia/Yerevan

Atlantic/Azores, Atlantic/Bermuda, Atlantic/Canary, Atlantic/Cape_Verde, Atlantic/Faeroe, Atlantic/Jan_Mayen, Atlantic/Madeira, Atlantic/Reykjavik, Atlantic/South_Georgia, Atlantic/St_Helena, Atlantic/Stanley

Australia/Adelaide, Australia/Brisbane, Australia/Broken_Hill, Australia/Darwin, Australia/Hobart, Australia/Lindeman, Australia/Lord_Howe, Australia/Melbourne, Australia/Perth, Australia/Sydney

Europe/Amsterdam, Europe/Andorra, Europe/Athens, Europe/Belfast, Europe/Belgrade, Europe/Berlin, Europe/Bratislava, Europe/Brussels, Europe/Bucharest, Europe/Budapest, Europe/Chisinau, Europe/Copenhagen, Europe/Dublin, Europe/Gibraltar, Europe/Helsinki, Europe/Istanbul, Europe/Kaliningrad, Europe/Kiev, Europe/Lisbon, Europe/Ljubljana, Europe/London, Europe/Luxembourg, Europe/Madrid, Europe/Malta, Europe/Minsk, Europe/Monaco, Europe/Moscow, Europe/Oslo, Europe/Paris, Europe/Prague, Europe/Riga, Europe/Rome, Europe/Samara, Europe/San_Marino, Europe/Sarajevo, Europe/Simferopol, Europe/Skopje, Europe/Sofia, Europe/Stockholm, Europe/Tallinn, Europe/Tirane, Europe/Vaduz, Europe/Vatican, Europe/Vienna, Europe/Vilnius, Europe/Warsaw, Europe/Zagreb, Europe/Zurich

Indian/Antananarivo, Indian/Chagos, Indian/Christmas, Indian/Cocos, Indian/Comoro, Indian/Kerguelen, Indian/Mahe, Indian/Maldives, Indian/Mauritius, Indian/Mayotte, Indian/Reunion

Pacific/Apia, Pacific/Auckland, Pacific/Chatham, Pacific/Easter, Pacific/Efate, Pacific/Enderbury, Pacific/Fakaofu, Pacific/Fiji, Pacific/Funafuti, Pacific/Galapagos, Pacific/Gambier, Pacific/Guadalcanal, Pacific/Guam, Pacific/Honolulu, Pacific/Johnston, Pacific/Kiritimati, Pacific/Kosrae, Pacific/Kwajalein, Pacific/Majuro, Pacific/Marquesas, Pacific/Midway, Pacific/Nauru, Pacific/Niue, Pacific/Norfolk, Pacific/Noumea, Pacific/Pago_Pago, Pacific/Palau, Pacific/Pitcairn, Pacific/Ponape, Pacific/Port_Moresby, Pacific/Rarotonga, Pacific/Saipan, Pacific/Tahiti, Pacific/Tarawa, Pacific/Tongatapu, Pacific/Truk, Pacific/Wake, Pacific/Wallis, Pacific/Yap

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation • [Modifying the Default Time Zone for a Router or Switch Running Junos OS on page 24](#)

traceoptions (Layer 2 Learning)

Syntax	<pre> traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> (detail disable receive send); in-memory-debug; level; no-remote-trace; } </pre>
Hierarchy Level	[edit protocols l2-learning]
Release Information	Statement introduced in Junos OS Release 13.2 for the QFX Series.
Description	Define tracing operations for Layer 2 learning.
Default	The traceoptions feature is disabled by default.
Options	<p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log.</p> <p>You can specify the following options:</p> <ul style="list-style-type: none"> • no-world-readable—(Optional) Restrict file access to the user who created the file. • size <i>size</i> —(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the files option. Use xk to specify KB, xm to specify MB, or xg to specify gigabytes. • world-readable—(Optional) Enable unrestricted file access. <p>flag <i>flag</i> —Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:</p> <ul style="list-style-type: none"> • all—All tracing operations. • bmac-next-hop—Trace backbone MAC next hop operations. • bridge-bmac-next-hop—Trace backbone MAC next hop bridge operations. • bridging-interface—Trace interface bridge operations. • bridging-domain—Trace bridging domain operations. • configuration—Trace configuration operations. • flood-next-hop—Trace flood next hop operations. • initialization—Trace initialization operations. • interface-device—Trace interface device operations. • interface-family—Trace interface family operations.

- **interface-logical**—Trace logical interface operations.
- **ipc**—Trace inter-process communications operations.
- **irb**—Trace integrated routing and bridging operations.
- **isid**—Trace i-tagged service ID operations.
- **kack**—Trace kernel-acknowledgment.
- **learning-domain**—Trace learning domain operations.
- **logical-system**—Trace logical system operations.
- **mac-learning**—Trace MAC address learning.
- **mc-ae**—Trace multichassis aggregated Ethernet interface operations.
- **redundant-trunk-group**—Trace redundant trunk group operations.
- **routing-instance**—Trace routing instance operations.
- **routing-socket**—Trace routing socket operations.
- **storm-control**—Trace storm control operations.
- **unknown-unicast-forwarding**—Trace unknown unicast forwarding events.
- **vpls-ping**—Trace Virtual Private VLAN Service (VPLS) ping operations.

in-memory-debug—Enable trace parameters in the memory.

level—Specify level of debugging output.

no-remote-trace—Disable remote tracing.

Required Privilege Level	routing —To view this statement in the configuration.
	routing-control —To add this statement to the configuration.

uid

Syntax	<code>uid <i>uid-value</i>;</code>
Hierarchy Level	[edit system login user]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure a user identifier for a login account.
Options	<i>uid-value</i> —Number associated with the login account. This value must be unique on the router or switch. Range: 100 through 64000
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Junos OS User Accounts on page 125

use-imported-time-zones

Syntax	<code>use-imported-time-zones;</code>
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure a custom time zone from a locally generated time zone database.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Updating the IANA Time Zone Database on Junos OS Devices on page 32

CHAPTER 17

DHCP Configuration Statements

- [client-identifier \(DHCP Client\) on page 251](#)
- [lease-time on page 252](#)
- [retransmission-attempt on page 253](#)
- [retransmission-interval on page 254](#)
- [server-address on page 255](#)
- [update-server on page 256](#)
- [vendor-option on page 257](#)

client-identifier (DHCP Client)

Syntax	client-identifier (ascii <i>ascii</i> hexadecimal <i>hexadecimal</i>);
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet dhcp]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify an ASCII or hexadecimal identifier for the Dynamic Host Configuration Protocol (DHCP) client. The DHCP server identifies a client by a client-identifier value, which must be unique for each client.
Default	If you do not include client-identifier in the configuration, the DHCP server uses the client hardware type and MAC address to identify the client.
Options	ascii <i>ascii</i> —Identifier consisting of ASCII characters, such as a fully qualified domain name. hexadecimal <i>hexadecimal</i> —Identifier consisting of hexadecimal numbers (0-9, a-f, A-F). Do not use colons.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	• Configuring a DHCP Client (CLI Procedure) on page 77

lease-time

Syntax	lease-time (<i>length</i> infinite);
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet dhcp]
Release Information	Statement introduced in Junos OS Release 9.2.
Description	Request a specific lease time for the IP address.
Default	If no lease time is requested by client, then the server sends the lease time. The default lease time on a Junos OS DHCP server is one day.
Options	seconds —Request a lease time of a specific duration. Range: 60 through 2147483647 seconds infinite —Request that the lease never expire.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Understanding Interfaces</i>• <i>unit</i>• <i>family</i>

retransmission-attempt

Syntax	<code>retransmission-attempt <i>number</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet dhcp]
Release Information	Statement introduced in Junos OS Release 8.5 for J Series devices. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 9.2 for SRX Series devices. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify the number of times the device retransmits a Dynamic Host Control Protocol (DHCP) packet if a DHCP server fails to respond. After the specified number of attempts, no further attempts at reaching a server are made.
Options	<i>number</i> —Number of retransmit attempts.. Range: 0 through 6 Default: 4
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring a DHCP Client (CLI Procedure) on page 77 • <i>Example: Configuring the Device as a DHCP Client</i> • <i>interfaces</i> • <i>unit</i> • <i>family</i>

retransmission-interval

Syntax	retransmission-interval <i>seconds</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet dhcp]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify the time between successive retransmissions of the client DHCP request if a DHCP server fails to respond.
Options	<i>seconds</i> —Number of seconds between successive retransmissions. Range: 4 through 64 seconds Default: 4 seconds
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a DHCP Client (CLI Procedure) on page 77


server-address

Syntax	<code>server-address <i>ip-address</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet dhcp]
Release Information	Statement introduced in Junos OS Release 8.5 for J Series devices. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 9.2 for SRX Series devices. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify the address of the DHCP server that the client should accept DHCP offers from. If this option is included in the DHCP configuration, the client accepts offers only from this server and ignores all other offers.
Default	The client accepts the first offer it receives from any DHCP server.
Options	<i>ip-address</i> —DHCP server address.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring a DHCP Client (CLI Procedure) on page 77 • <i>Example: Configuring the Device as a DHCP Client</i> • <i>interfaces</i> • <i>unit</i> • <i>family</i>

update-server

Syntax	update-server;
Hierarchy Level	[edit Interfaces <i>interface-name</i> unit <i>logical-unit-number</i> inet dhcp]
Release Information	Statement introduced in Junos OS Release 8.5 for J Series devices. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 9.2 for SRX Series devices. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Propagate TCP/IP settings learned from an external DHCP server to the DHCP server running on the switch, router, or device.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a DHCP Client (CLI Procedure) on page 77• <i>Example: Configuring the Device as a DHCP Client</i>• <i>interfaces</i>• <i>unit</i>• <i>family</i>

vendor-option

Syntax	<pre> vendor-option { default-local-server-group <i>local-server-group-name</i> default-relay-server-group <i>server-group-name</i> drop; equals starts-with } </pre>
Hierarchy Level	[edit forwarding-options dhcp-relay relay-option-60]
Release Information	Statement introduced before Junos OS Release 12.1 for EX Series switches. Statement deprecated in Junos OS Release 12.3 for EX Series switches.
Description	Configure the match criteria when you use the DHCP vendor class identifier option (option 60) in DHCP client packets to forward client traffic to specific DHCP servers. The extended DHCP relay agent compares the option 60 vendor-specific strings received in DHCP client packets against the match criteria that you specify. If there is a match, you can define certain actions for the associated DHCP client packets.
<div style="display: flex; align-items: center;">  <div> <p>NOTE: The <code>vendor-option</code> statement has been deprecated and might be removed from future product releases. We recommend that you phase out its use. See <i>option-number</i>.</p> </div> </div>	
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring an Extended DHCP Relay Server on EX Series Switches (CLI Procedure)</i> • <i>Understanding the Extended DHCP Relay Agent for EX Series Switches</i>

CHAPTER 18

ICMP Configuration Statement

- [no-multicast-echo on page 259](#)

no-multicast-echo

Syntax	no-multicast-echo
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 8.1.
Description	Disable the Routing Engine from responding to ICMP echo requests sent to multicast group addresses.
Default	The Routing Engine responds to ICMP echo requests sent to multicast group addresses.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Junos OS to Disable the Routing Engine Response to Multicast Ping Packets on page 86

CHAPTER 19

Junos OS With Upgraded FreeBSD Operational Commands

- request system reboot (Junos OS with Upgraded FreeBSD)
- request system snapshot (Junos OS with Upgraded FreeBSD)
- show system snapshot (Junos OS with Upgraded FreeBSD)

request system reboot (Junos OS with Upgraded FreeBSD)

Syntax	<pre>request system reboot <all-members> <at <i>time</i>> <both-routing-engines> <in <i>minutes</i>> <local> <media (oam junos network usb)> <member <i>member-id</i>> <message "<i>text</i>"> <other-routing-engine></pre>
Release Information	<p>Command introduced in Junos OS Release 15.1 for MX240, MX480, MX960, MX2010, and MX2020 routers and EX9200 switches.</p> <p>Command introduced in Junos OS Release 15.1X53-D30 for QFX5200 switches.</p>
Description	Reboot the software.
Options	<p>none—Reboot the software immediately.</p> <p>all-members—(Optional) Reboot the software on all members of the Virtual Chassis configuration.</p> <p>at <i>time</i>—(Optional) Time at which to reboot the software, specified in one of the following ways:</p> <ul style="list-style-type: none">• now—Stop or reboot the software immediately. This is the default.• +<i>minutes</i>—Number of minutes from now to reboot the software.• <i>yymmddhhmm</i>—Absolute time at which to reboot the software, specified as year, month, day, hour, and minute. Omitting a value will default to the current date for that value.• <i>hh:mm</i>—Absolute time on the current day at which to stop the software, specified in 24-hour time. <p>both-routing-engines—(Optional) Reboot both Routing Engines at the same time.</p> <p>in <i>minutes</i>—(Optional) Number of minutes from now to reboot the software. This option is an alias for the at +<i>minutes</i> option.</p> <p>local—(Optional) Reboot the software on the local Virtual Chassis member.</p> <p>media (oam junos network usb)—(Optional) Reboot the boot media:</p> <ul style="list-style-type: none">• oam—Reboot from the oam volume.• junos—Reboot from the junos volume.• network—Reboot from the network.• usb—Reboot from the USB device.

member *member-id*—(Optional) Reboot the software on the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

message "*text*"—(Optional) Message to display to all system users before stopping or rebooting the software.

other-routing-engine—(Optional) Reboot the other Routing Engine from which the command is issued. For example, if you issue the command from the master Routing Engine, the backup Routing Engine is rebooted. Similarly, if you issue the command from the backup Routing Engine, the master Routing Engine is rebooted.

Additional Information Reboot requests are recorded in the system log files, which you can view with the **show log** command (see [show log](#)). Also, the names of any running processes that are scheduled to be shut down are changed. You can view the process names with the **show system processes** command (see [show system processes](#)).



NOTE: To reboot a router or switch that has two Routing Engines, reboot the backup Routing Engine (if you have upgraded it) first, and then reboot the master Routing Engine.

Required Privilege Level maintenance

Related Documentation

- [request system snapshot \(Junos OS with Upgraded FreeBSD\) on page 265](#)
- [show system snapshot \(Junos OS with Upgraded FreeBSD\) on page 267](#)
- [clear system reboot on page 337](#)
- [request system halt on page 384](#)
- [Understanding Junos OS with Upgraded FreeBSD on page 89](#)

List of Sample Output

[request system reboot on page 263](#)
[request system reboot \(at 2300\) on page 263](#)
[request system reboot \(in 2 Hours\) on page 264](#)
[request system reboot \(Immediately\) on page 264](#)
[request system reboot \(at 1:20 AM\) on page 264](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system reboot

```
user@host> request system reboot
Reboot the system ? [yes,no] (no)
```

request system reboot (at 2300)

```
user@host> request system reboot at 2300 message "Maintenance time!"
```

```
Reboot the system ? [yes,no] (no) yes
```

```
shutdown: [pid 186]
```

```
*** System shutdown message from root@berry.network.net ***
```

```
System going down at 23:00
```

request system reboot (in 2 Hours)

The following example, which assumes that the time is 5 PM (17:00), illustrates three different ways to request the system to reboot in two hours:

```
user@host> request system reboot at +120
```

```
user@host> request system reboot in 120
```

```
user@host> request system reboot at 19:00
```

request system reboot (Immediately)

```
user@host> request system reboot at now
```

request system reboot (at 1:20 AM)


To reboot the system at 1:20 AM, enter the following command. Because 1:20 AM is the next day, you must specify the absolute time.

```
user@host> request system reboot at 06060120
```

```
request system reboot at 120
```

```
Reboot the system at 120? [yes,no] (no) yes
```


request system snapshot (Junos OS with Upgraded FreeBSD)

Syntax	request system snapshot <delete <i>snapshot-name</i> > <load <i>snapshot-name</i> > <media <i>type</i> > <recovery>
Release Information	Command introduced in Junos OS Release 15.1 for MX240, MX480, MX960, MX2010, and MX2020 routers and EX9200 switches. Command introduced in Junos OS Release 15.1X53-D30 for QFX5200 switches.
Description	On the router or switch, back up the currently running and active file system partitions to standby partitions that are not running. Non-recovery snapshots are named snap.date.time and stored in the /packages/sets directory.
<div>  <p>CAUTION: After you run the request system snapshot command, you cannot return to the previous version of the software, because the running and backup copies of the software are identical.</p> </div>	
Options	<p>none—On the router or switch, back up the currently running and active file system partitions to standby partitions that are not running. Specifically, this creates a non-recovery snapshot named snap.<date>.<time> which is stored in /packages/sets.</p> <p>delete <i>snapshot-name</i>—(Optional) Delete a specific non-recovery snapshot from /packages/sets. Wildcards are supported, so request system snapshot delete snap* deletes all snapshots.</p> <p>load <i>snapshot-name</i>—(Optional) Load a specific snapshot from /packages/sets.</p> <p>media <i>type</i>—(Optional) Specify the boot device the software is copied to:</p> <ul style="list-style-type: none"> usb—(MX960 routers only) Copy software to the device connected to the USB port. <p>recovery—Create a recovery snapshot and store it in the /oam volume.</p>
Additional Information	Before upgrading the software on the router or switch, when you have a known stable system, issue the request system snapshot command to back up the software, including the configuration, to the /packages/sets file systems. After you have upgraded the software on the router or switch and are satisfied that the new packages are successfully installed and running, issue the request system snapshot command again to back up the new software to the /packages/sets file systems.
Required Privilege Level	maintenance

- Related Documentation**
- [request system reboot \(Junos OS with Upgraded FreeBSD\) on page 262](#)
 - [show system snapshot \(Junos OS with Upgraded FreeBSD\) on page 267](#)
 - [Understanding Junos OS with Upgraded FreeBSD on page 89](#)

List of Sample Output

- [request system snapshot recovery on page 266](#)
- [request system snapshot delete on page 266](#)
- [request system snapshot on page 266](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

[request system snapshot recovery](#)

```
user@host> request system snapshot recovery
Creating image ...
Compressing image . . .
Image size is 777MB

Recovery snapshot created successfully
```


[request system snapshot delete](#)

```
user@host> request system snapshot delete snap.20150112.122106
NOTICE: Snapshot 'snap.20150112.122106' deleted successfully
```

[request system snapshot](#)

```
user@host> request system snapshot
NOTICE: Snapshot snap.20150119.122106 created successfully
```

show system snapshot (Junos OS with Upgraded FreeBSD)

Syntax	show system snapshot
Release Information	<p>Command introduced in Junos OS Release 15.1 for MX240, MX480, MX960, MX2010, and MX2020 routers and EX9200 switches.</p> <p>Command introduced in Junos OS Release 15.1X53-D30 for QFX5200 switches.</p>
Description	Display information about the non-recovery backup software, which is located in the junos file system on the hard disk drive or solid-state drive (SSD).
<div>  NOTE: To back up software, use the <code>request system snapshot</code> command. </div>	
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • request system snapshot (Junos OS with Upgraded FreeBSD) on page 265 • request system reboot (Junos OS with Upgraded FreeBSD) on page 262 • Understanding Junos OS with Upgraded FreeBSD on page 89
List of Sample Output	show system snapshot on page 267
Output Fields	When you enter this command, you are provided feedback on the status of your request. If there are no snapshots available, the command returns null output.

Sample Output

show system snapshot

```

user@host> show system snapshot
Snapshot snap.20141219.122106:
Location: /packages/sets/snap.20141219.122106
Creation date: Dec 19 12:21:06 2014
Junos version: 15.1-20141216_ib_15_1_psd.0

```


CHAPTER 20

Login Classes Configuration Statements

- [access-end on page 270](#)
- [access-start on page 270](#)
- [allow-commands on page 271](#)
- [allow-configuration on page 272](#)
- [allowed-days on page 272](#)
- [announcement on page 273](#)
- [authentication \(Login\) on page 274](#)
- [class \(Defining Login Classes\) on page 275](#)
- [class \(Assigning a Class to an Individual User\) on page 276](#)
- [deny-commands on page 276](#)
- [deny-configuration on page 277](#)
- [idle-timeout on page 278](#)
- [load-key-file on page 279](#)
- [login on page 280](#)
- [login-tip on page 281](#)
- [message on page 281](#)
- [user \(Access\) on page 282](#)

access-end

Syntax	access-end <i>HH:MM</i> ;
Hierarchy Level	[edit system login class]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the end time for login access.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Time-Based User Access on page 20

access-start

Syntax	access-start <i>HH:MM</i> ;
Hierarchy Level	[edit system login class]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the start time for login access.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Time-Based User Access on page 20

allow-commands

Syntax	<code>allow-commands "regular-expression";</code>
Hierarchy Level	[edit system login class <i>class-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify the operational mode commands that members of a login class can use.
Default	If you omit this statement and the deny-commands statement, users can issue only those commands for which they have access privileges through the permissions statement.
Options	regular-expression —Extended (modern) regular expression as defined in POSIX 1003.2. If the regular expression contains any spaces, operators, or wildcard characters, enclose it in quotation marks.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Specifying Access Privileges for Junos OS Operational Mode Commands on page 30 • deny-commands on page 276 • user on page 282

allow-configuration

Syntax	<code>allow-configuration "regular-expression";</code>
Hierarchy Level	[edit system login class <i>class-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Explicitly allow configuration access to the specified levels in the hierarchy even if the permissions set with the permissions statement do not grant such access by default.
Default	If you omit this statement and the deny-configuration statement, users can edit only those commands for which they have access privileges through the permissions statement.
Options	regular-expression —Extended (modern) regular expression as defined in POSIX 1003.2. If the regular expression contains any spaces, operators, or wildcard characters, enclose it in quotation marks.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Specifying Access Privileges for Junos OS Configuration Mode Hierarchies</i>• Regular Expressions for Allowing and Denying Junos OS Configuration Mode Hierarchies on page 1826• deny-configuration on page 277• user on page 282

allowed-days

Syntax	<code>allowed-days [<i>days-of-the-week</i>];</code>
Hierarchy Level	[edit system login class <i>class-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify the days of the week when users can log in.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Time-Based User Access on page 20

announcement

Syntax	announcement text;
Hierarchy Level	[edit system login]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure a system login announcement. This announcement appears after a user logs in.
Options	text —Text of the announcement. If the text contains any spaces, enclose it in quotation marks.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Junos OS to Display a System Login Announcement</i>• Configuring Junos OS to Display a System Login Message on page 126• message on page 281

authentication (Login)

Syntax	<pre>authentication { encrypted-password <i>password</i>; load-key-file <i>URL</i>; plain-text-password <i>password</i>; remote-debug-permission (qfabric-admin qfabric-operator qfabric-user); ssh-dsa "<i>public-key</i>"; ssh-rsa "<i>public-key</i>"; }</pre>
Hierarchy Level	[edit system login user <i>username</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Authentication methods that a user can use to log in to the switch. You can assign multiple authentication methods to a single user.
Options	<p>encrypted-password "<i>password</i>"—Message Digest 5 (MD5) or other encrypted authentication. Specify the MD5 or other password. You can specify only one encrypted password for each user.</p> <p>You cannot configure a blank password for encrypted-password using blank quotation marks (" "). You must configure a password of 1 through 128 characters and enclose the password in quotation marks.</p> <p>load-key-file—Load RSA (SSH version 1 and SSH version 2) and DSA (SSH version 2) public keys from a file. The file is a URL containing one or more SSH keys.</p> <p>plain-text-password—Plain-text password. The command-line interface (CLI) prompts you for the password and then encrypts it.</p> <p>remote-debug-permission (QFabric systems only)—QFabric component authentication. Specifies permission levels for users to access individual components in a QFabric system.</p> <p>ssh-dsa "<i>public-key</i>"—SSH version 2 authentication. Specify the SSH public key. You can specify one or more public keys for each user.</p> <p>ssh-rsa "<i>public-key</i>"—SSH version 1 and SSH version 2 authentication. Specify the SSH public key. You can specify one or more public keys for each user.</p>
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Junos OS User Accounts on page 125• root-authentication on page 307

class (Defining Login Classes)

Syntax	<pre> class <i>class-name</i> { access-end; access-start; allow-commands "<i>regular-expression</i>"; allow-configuration "<i>regular-expression</i>"; deny-commands "<i>regular-expression</i>"; deny-configuration "<i>regular-expression</i>"; idle-timeout <i>minutes</i>; login-tip; permissions [<i>permissions</i>]; } </pre>
Hierarchy Level	[edit system login]
Release Information	<p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Define a login class.
Options	<p><i>class-name</i>—A name you choose for the login class.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Defining Junos OS Login Classes user on page 282

class (Assigning a Class to an Individual User)

Syntax	<pre>class <i>class-name</i> { operator; read-only; super-user; unauthorized; }</pre>
Hierarchy Level	[edit system login user <i>username</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure a user's login class. You must configure one class for each user.
Options	<i>class-name</i> —One of the classes defined at the [edit system login class] hierarchy level.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Junos OS User Accounts on page 125

deny-commands

Syntax	<pre>deny-commands "<i>regular-expression</i>";</pre>
Hierarchy Level	[edit system login class]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify the operational mode commands that the user is denied permission to issue even though the permissions set with the permissions statement would allow their use.
Default	If you omit this statement and the allow-commands statement, users can issue only those commands for which they have access privileges through the permissions statement.
Options	<i>regular-expression</i> —Extended (modern) regular expression as defined in POSIX 1003.2. If the regular expression contains any spaces, operators, or wildcard characters, enclose it in quotation marks.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Specifying Access Privileges for Junos OS Operational Mode Commands on page 30• allow-commands on page 271• user on page 282


deny-configuration

Syntax	<code>deny-configuration "regular-expression";</code>
Hierarchy Level	[edit system login class]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Explicitly deny configuration access to the specified levels in the hierarchy even if the permissions set with the permissions statement grant such access by default.
Default	If you omit this statement and the allow-configuration statement, users can edit those levels in the configuration hierarchy for which they have access privileges through the permissions statement.
Options	regular-expression —Extended (modern) regular expression as defined in POSIX 1003.2. If the regular expression contains any spaces, operators, or wildcard characters, enclose it in quotation marks.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Specifying Access Privileges Using allow/deny-configuration Statements</i> • allow-configuration on page 272 • user on page 282

idle-timeout

Syntax	<code>idle-timeout <i>minutes</i>;</code>
Hierarchy Level	<code>[edit system login class <i>class-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	For a login class, configure the maximum time that a session can be idle before the user is logged off the switch. The session times out after remaining at the CLI operational mode prompt for the specified time.
Default	If you omit this statement, a user is never forced off the system after extended idle times.
Options	<i>minutes</i> —Maximum idle time. Range: 0 through 4294967295 minutes
Required Privilege Level	admin —To view this statement in the configuration. admin-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Timeout Value for Idle Login Sessions on page 22

load-key-file

Syntax	<code>load-key-file URL filename;</code>
Hierarchy Level	[edit system root-authentication], [edit system login user <i>username</i> authentication]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	 NOTE: ECDSA is not supported on the QFabric system. Load RSA (SSH version 1 and SSH version 2) and DSA or ECDSA (SSH version 2) public keys from a previously-generated named file at a specified URL location or local path. The file contains one or more SSH keys that are copied into the configuration when the command is issued.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Root Password on page 145 • <i>Configuring the Root Password</i> • <i>Configuring Junos OS User Accounts by Using a Configuration Group</i>

login

```
Syntax login {
    announcement text;
    class class-name {
        access-end "regular-expression";
        access-start "regular-expression";
        allow-commands "regular-expression";
        allow-configuration "regular-expression";
        deny-commands "regular-expression";
        deny-configuration "regular-expression";
        idle-timeout minutes;
        login-tip;
        permissions [ permissions ];
    }
    message text;
    password {
        change-type (set-transitions | character-set);
        format (md5 | sha1 | des);
        maximum-length length;
        minimum-changes number;
        minimum-length length;
    }
    retry-options {
        backoff-factor seconds;
        backoff-threshold number;
        minimum-time seconds;
        tries-before-disconnect number;
    }
    user username {
        authentication authentication;
        (encrypted-password "password" | plain-text-password);
        load-key-file URL;
        remote-debug-permission (qfabric-admin | qfabric-operator | qfabric-user);
        ssh-dsa "public-key";
        ssh-rsa "public-key";
    }
    class class-name;
    full-name complete-name;
    uid uid-value;
}
```

Hierarchy Level [edit system]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure user access to the switch.

Options The remaining statements are explained separately.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- [Defining Junos OS Login Classes](#)

login-tip

Syntax login-tip;

Hierarchy Level [edit system login class *class-name*]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Enable CLI tips at login.

Default Disabled.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- [Configuring Login Tips on page 40](#)

message

Syntax message *text*;

Hierarchy Level [edit system login]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure a system login message. This message appears before a user logs in.

Options *text*—Text of the message.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration

Related Documentation

- [Configuring Junos OS to Display a System Login Message on page 126](#)
- [announcement on page 273](#)

user (Access)

Syntax	<pre>user username { authentication { (encrypted-password "password" plain-text-password); load-key-file URL; remote-debug-permission (qfabric-admin qfabric-operator qfabric-user); ssh-dsa "public-key" <from hostname>; ssh-rsa "public-key" <from hostname>; } class class-name; full-name "complete-name"; uid uid-value; }</pre>
Hierarchy Level	[edit system login]
Release Information	Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure access permission for individual users.
Options	The remaining statements are explained separately.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Junos OS User Accounts on page 125• class on page 275

CHAPTER 21

NTP Configuration Statements

- [authentication-key on page 284](#)
- [boot-server \(NTP\) on page 285](#)
- [broadcast on page 286](#)
- [broadcast-client on page 287](#)
- [multicast-client on page 287](#)
- [ntp on page 288](#)
- [peer on page 289](#)
- [server \(NTP\) on page 290](#)
- [source-address \(NTP, RADIUS, System Logging, or TACACS+\) on page 291](#)
- [trusted-key on page 292](#)

authentication-key

Syntax	<code>authentication-key <i>key-number</i> type <i>type</i> value <i>password</i>;</code>
Hierarchy Level	[edit system ntp]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Configure Network Time Protocol (NTP) authentication keys so that the router or switch can send authenticated packets. If you configure the router or switch to operate in authenticated mode, you must configure a key.</p> <p>Both the keys and the authentication scheme (MD5) must be identical between a set of peers sharing the same key number.</p>
Options	<p><i>key-number</i>—An integer in the range of 1 to 65533.</p> <p><i>type type</i>—Authentication type. It can only be md5.</p> <p><i>value password</i>—Key itself, consisting of 1 through 8 ASCII characters. If the key contains spaces, enclose it in quotation marks.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Understanding NTP Time Servers on page 131• Configuring NTP Authentication Keys on page 132

boot-server (NTP)

Syntax	<code>boot-server (address hostname);</code>
Hierarchy Level	[edit system ntp]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Configure the server that NTP queries when the router or switch boots to determine the local date and time.</p> <p>When you boot the router or switch, it issues an ntpdate request, which polls a network server to determine the local date and time. You need to configure a server that the router or switch uses to determine the time when the router or switch boots. Otherwise, NTP cannot synchronize to a time server if the server time significantly differs from the local router's or switch's time. You can configure either an IP address or a hostname for the boot server. If you configure a hostname instead of an IP address, the ntpdate request resolves the hostname to an IP address when the router or switch boots up.</p>
Options	<ul style="list-style-type: none"> • address—IP address of an NTP boot server. • hostname—Hostname of an NTP boot server.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding NTP Time Servers on page 131 • Configuring NTP Authentication Keys on page 132 • Synchronizing and Coordinating Time Distribution Using NTP on page 137

broadcast

Syntax	<code>broadcast address <key key-number> <version value> <tll value>;</code>
Hierarchy Level	[edit system ntp]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the local router or switch to operate in broadcast mode with the remote system at the specified address to send periodic broadcast messages to a client population. Normally, you include this statement only when the local router or switch is operating as a transmitter.
Options	<p>address—Broadcast address on one of the local networks or a multicast address assigned to NTP. You must specify an address, not a hostname. If the multicast address is used, it must be 224.0.1.1.</p> <p>key key-number—(Optional) All packets sent to the address include authentication fields that are encrypted using the specified key number (any unsigned 32-bit integer).</p> <p>tll value—(Optional) Time-to-live (TTL) value to use. Range: 1 through 255 Default: 1</p> <p>version value—(Optional) Specify the version number to be used in outgoing NTP packets. Range: 1 through 4 Default: 4</p>
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding NTP Time Servers on page 131• Configuring NTP Authentication Keys on page 132• Configuring the NTP Time Server and Time Services on page 133

broadcast-client

Syntax	<code>broadcast-client;</code>
Hierarchy Level	<code>[edit system ntp]</code>
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the local switch to listen for broadcast messages on the local network to discover other servers on the same subnet.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding NTP Time Servers on page 131 • Configuring NTP Authentication Keys on page 132 • Configuring the Switch to Listen for Broadcast Messages Using NTP on page 135

multicast-client

Syntax	<code>multicast-client <address>;</code>
Hierarchy Level	<code>[edit system ntp]</code>
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	For Network Time Protocol (NTP), configure the local router or switch to listen for multicast messages on the local network to discover other servers on the same subnet.
Options	address —(Optional) One or more IP addresses. If you specify addresses, the router or switch joins those multicast groups. Default: 224.0.1.1.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding NTP Time Servers on page 131 • Configuring NTP Authentication Keys on page 132 • Configuring the Switch to Listen for Multicast Messages Using NTP on page 136

ntp

Syntax	<pre>ntp { authentication-key <i>number</i> type <i>type</i> value <i>password</i>; boot-server <i>address</i>; broadcast <<i>address</i>> <<i>key key-number</i>> <<i>version value</i>> <<i>ttl value</i>>; broadcast-client; multicast-client <<i>address</i>>; peer <i>address</i> <<i>key key-number</i>> <<i>version value</i>> <<i>prefer</i>>; server <i>address</i> <<i>key key-number</i>> <<i>version value</i>> <<i>prefer</i>>; source-address <i>source-address</i>; trusted-key [<i>key-numbers</i>]; }</pre>
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure Network Time Protocol (NTP) on the switch. The remaining statements are explained separately.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding NTP Time Servers on page 131• Configuring NTP Authentication Keys on page 132• Synchronizing and Coordinating Time Distribution Using NTP on page 137

peer

Syntax	<code>peer address <key key-number> <version value> <prefer>;</code>
Hierarchy Level	[edit system ntp]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	For NTP, configure the local router or switch to operate in symmetric active mode with the remote system at the specified address. In this mode, the local router or switch and the remote system can synchronize with each other. This configuration is useful in a network in which either the local router or switch or the remote system might be a better source of time.
Options	<p>address—Address of the remote system. You must specify an address, not a hostname.</p> <p>key key-number—(Optional) All packets sent to the address include authentication fields that are encrypted using the specified key number.</p> <p>Range: Any unsigned 32-bit integer</p> <p>prefer—(Optional) Mark the remote system as the preferred host, which means that if all other factors are equal, this remote system is chosen for synchronization among a set of correctly operating systems.</p> <p>version value—(Optional) Specify the NTP version number to be used in outgoing NTP packets.</p> <p>Range: 1 through 4</p> <p>Default: 4</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Understanding NTP Time Servers on page 131 • Configuring NTP Authentication Keys on page 132 • Configuring the NTP Time Server and Time Services on page 133

server (NTP)

Syntax	<code>server address <key key-number> <version value> <prefer>;</code>
Hierarchy Level	[edit system ntp]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	For NTP, configure the switch to operate in client mode with the remote system at the specified server address. In this mode, the local switch can be synchronized with the remote system, but the remote system can never be synchronized with the local switch.
Options	<p>address—Address of the remote system. You must specify an address, not a hostname.</p> <p>key key-number—(Optional) Use the specified key number to encrypt authentication fields in all packets sent to the specified address.</p> <p>Range: Any unsigned 32-bit integer</p> <p>prefer—(Optional) Mark the remote system as preferred host, which means that if all other things are equal, this remote system is chosen for synchronization among a set of correctly operating systems.</p> <p>version value—(Optional) Specify the version number to be used in outgoing NTP packets.</p> <p>Range: 1 through 4</p> <p>Default: 4</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Understanding NTP Time Servers on page 131• Configuring NTP Authentication Keys on page 132

source-address (NTP, RADIUS, System Logging, or TACACS+)

Syntax	<code>source-address <i>source-address</i>;</code>
Hierarchy Level	[edit system accounting destination radius server <i>server-address</i>], [edit system accounting destination tacplus server <i>server-address</i>], [edit system ntp], [edit system radius-server <i>server-address</i>], [edit system syslog], [edit system tacplus-server <i>server-address</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify a source address for each configured TACACS+ server, RADIUS server, NTP server, or the source address to record in system log messages that are directed to a remote machine.
Options	<i>source-address</i> —Valid IP address configured on one of the switch interfaces. For system logging, the address is recorded as the message source in messages sent to the remote machines specified in all host <i>hostname</i> statements at the [edit system syslog] hierarchy level.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring RADIUS Authentication (QFX Series or OCX Series) on page 1873 • Synchronizing and Coordinating Time Distribution Using NTP on page 137 • Specifying an Alternative Source Address for System Log Messages Directed to a Remote Destination

trusted-key

Syntax	<code>trusted-key [<i>key-numbers</i>];</code>
Hierarchy Level	[edit system ntp]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	For NTP, configure the keys to use when you configure the switch to synchronize its time with other systems on the network.
Options	<i>key-numbers</i> —One or more key numbers. Each key can be any 32-bit unsigned integer except 0.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring NTP Authentication Keys on page 132• <i>authentication-key</i>• server on page 290

CHAPTER 22

Password Configuration Statements

- [authentication \(Login\) on page 294](#)
- [change-type on page 295](#)
- [format on page 296](#)
- [maximum-length on page 297](#)
- [minimum-changes on page 298](#)
- [minimum-length on page 299](#)
- [minimum-lower-cases on page 300](#)
- [minimum-numeric on page 301](#)
- [minimum-punctuations on page 302](#)
- [minimum-upper-cases on page 303](#)
- [password \(Login\) on page 304](#)
- [permissions on page 304](#)
- [retry-options on page 305](#)
- [root-authentication on page 307](#)
- [ssh-dsa on page 308](#)
- [ssh-rsa on page 308](#)
- [user \(Access\) on page 309](#)

authentication (Login)

Syntax	<pre>authentication { encrypted-password <i>password</i>; load-key-file <i>URL</i>; plain-text-password <i>password</i>; remote-debug-permission (qfabric-admin qfabric-operator qfabric-user); ssh-dsa "<i>public-key</i>"; ssh-rsa "<i>public-key</i>"; }</pre>
Hierarchy Level	[edit system login user <i>username</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Authentication methods that a user can use to log in to the switch. You can assign multiple authentication methods to a single user.
Options	<p>encrypted-password "<i>password</i>"—Message Digest 5 (MD5) or other encrypted authentication. Specify the MD5 or other password. You can specify only one encrypted password for each user.</p> <p>You cannot configure a blank password for encrypted-password using blank quotation marks (" "). You must configure a password of 1 through 128 characters and enclose the password in quotation marks.</p> <p>load-key-file—Load RSA (SSH version 1 and SSH version 2) and DSA (SSH version 2) public keys from a file. The file is a URL containing one or more SSH keys.</p> <p>plain-text-password—Plain-text password. The command-line interface (CLI) prompts you for the password and then encrypts it.</p> <p>remote-debug-permission (QFabric systems only)—QFabric component authentication. Specifies permission levels for users to access individual components in a QFabric system.</p> <p>ssh-dsa "<i>public-key</i>"—SSH version 2 authentication. Specify the SSH public key. You can specify one or more public keys for each user.</p> <p>ssh-rsa "<i>public-key</i>"—SSH version 1 and SSH version 2 authentication. Specify the SSH public key. You can specify one or more public keys for each user.</p>
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Junos OS User Accounts on page 125• root-authentication on page 307

change-type

Syntax	change-type (character-sets set-transitions);
Hierarchy Level	[edit system login password]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Set requirements for using character sets in plain-text passwords. When you combine this statement with the minimum-changes statement, you can check for the total number of character sets included in the password or for the total number of character-set changes in the password. Newly created passwords must meet these requirements.
Options	Specify one of the following: <ul style="list-style-type: none">• character-sets—Number of character sets in the password. Valid character sets include uppercase letters, lowercase letters, numbers, punctuation, and other special characters.• set-transitions—Number of transitions between character sets.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Special Requirements for Junos OS Plain-Text Passwords on page 1824• minimum-changes on page 298

format

Syntax	format (des md5 sha1);
Hierarchy Level	[edit system login password]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the authentication algorithm for plain-text passwords.
Default	For Junos OS, the default encryption format is md5 . For Junos OS-FIPS software, the default encryption format is sha1 .
Options	The hash algorithm that authenticates the password can be one of three algorithms: <ul style="list-style-type: none">• des—Has a block size of 8 bytes; its key size is 48 bits long.• md5—Produces a 128-bit digest.• sha1—Produces a 160-bit digest.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Special Requirements for Junos OS Plain-Text Passwords on page 1824

maximum-length

Syntax	maximum-length <i>length</i> ;
Hierarchy Level	[edit system login passwords]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify the maximum number of characters allowed in plain-text passwords. Newly created passwords must meet this requirement.
Default	For Junos OS-FIPS software, the maximum number of characters for plain-text passwords is 20 . For Junos OS, no maximum is set.
Options	length —Maximum number of characters the password can include. Range: 1 to 64 characters
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Special Requirements for Junos OS Plain-Text Passwords on page 1824• minimum-length on page 299

minimum-changes

Syntax	<code>minimum-changes <i>number</i>;</code>
Hierarchy Level	[edit system login passwords]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Specify the minimum number of character sets (or character set changes) required in plain-text passwords. Newly created passwords must meet this requirement.</p> <p>This statement is used in combination with the change-type statement. If the change type is character-sets, then the number of character sets included in the password is checked against the specified minimum. If the change type is set-transitions, then the number of character set changes in the password is checked against the specified minimum.</p>
Default	For Junos OS, the minimum number of changes is 1. For Junos-FIPS software, the minimum number of changes is 3.
Options	<i>number</i> —Minimum number of character sets (or character set changes) required for the password.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Special Requirements for Junos OS Plain-Text Passwords on page 1824• change-type on page 295

minimum-length

Syntax	minimum-length <i>length</i> ;
Hierarchy Level	[edit system login password]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify the minimum number of characters required in plain-text passwords. Newly created passwords must meet this requirement.
Default	For Junos OS, the minimum number of characters for plain-text passwords is six. For Junos-FIPS software, the minimum number of characters for plain-text passwords is 10.
Options	length —Minimum number of characters the password must include. Range: 6 to 20 characters
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Special Requirements for Junos OS Plain-Text Passwords on page 1824 • maximum-length on page 297

minimum-lower-cases

Syntax	<code>minimum-lower-cases <i>number</i>;</code>
Hierarchy Level	[edit system login password]
Release Information	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Specify the minimum number of lower-case letters required in plain-text passwords. Newly created passwords must meet this requirement.</p> <p>This statement can be used in combination with all of the other requirement options for plain-text passwords, such as minimum-length, minimum-punctuations, minimum-upper-cases, and so on.</p> <p>Using several password minimum requirement options will cause the minimum-length to be reset if the total sum of the required minimums exceeds the minimum-length setting.</p>
Options	<i>number</i> —The minimum number of lower-case letters required for the password.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Special Requirements for Junos OS Plain-Text Passwords on page 1824• Example: Changing the Requirements for Junos OS Plain-Text Passwords on page 148• <i>password (Login)</i>

minimum-numeric

Syntax	<code>minimum-numeric <i>number</i>;</code>
Hierarchy Level	[edit system login password]
Release Information	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Specify the minimum number of numeric class characters required in plain-text passwords. Newly created passwords must meet this requirement.</p> <p>This statement can be used in combination with all of the other requirement options for plain-text passwords, such as minimum-length, minimum-punctuations, minimum-lower-cases, and so on.</p> <p>Using several password minimum requirement options will cause the minimum-length to be reset if the total sum of the required minimums exceeds the minimum-length setting.</p>
Options	<i>number</i> —The minimum number of numeric class characters required for the password.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Special Requirements for Junos OS Plain-Text Passwords on page 1824 • Example: Changing the Requirements for Junos OS Plain-Text Passwords on page 148 • <i>password (Login)</i>

minimum-punctuations

Syntax	<code>minimum-punctuations <i>number</i>;</code>
Hierarchy Level	[edit system login password]
Release Information	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Specify the minimum number of punctuation class characters required in plain-text passwords. Newly created passwords must meet this requirement.</p> <p>This statement can be used in combination with all of the other requirement options for plain-text passwords, such as minimum-length, minimum-upper-cases, minimum-lower-cases, and so on.</p> <p>Using several password minimum requirement options will cause the minimum-length to be reset if the total sum of the required minimums exceeds the minimum-length setting.</p>
Options	<i>number</i> —The minimum number of punctuation class characters required for the password.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Special Requirements for Junos OS Plain-Text Passwords on page 1824• Example: Changing the Requirements for Junos OS Plain-Text Passwords on page 148• <i>password (Login)</i>

minimum-upper-cases

Syntax	<code>minimum-upper-cases <i>number</i>;</code>
Hierarchy Level	[edit system login password]
Release Information	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Specify the minimum number of upper-case letters required in plain-text passwords. Newly created passwords must meet this requirement.</p> <p>This statement can be used in combination with all of the other requirement options for plain-text passwords, such as minimum-length, minimum-punctuations, minimum-lower-cases, and so on.</p> <p>Using several password minimum requirement options will cause the minimum-length to be reset if the total sum of the required minimums exceeds the minimum-length setting.</p>
Options	<i>number</i> —The minimum number of upper-case letters required for the password.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Special Requirements for Junos OS Plain-Text Passwords on page 1824 • Example: Changing the Requirements for Junos OS Plain-Text Passwords on page 148 • <i>password (Login)</i>

password (Login)

Syntax	<pre>password { change-type (set-transitions character-set); format (md5 sha1 des); maximum-length length; minimum-changes number; minimum-length length; }</pre>
Hierarchy Level	[edit system login]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure special requirements such as character length and encryption format for plain-text passwords. Newly created passwords must meet these requirements. The remaining statements are explained separately.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Special Requirements for Junos OS Plain-Text Passwords on page 1824

permissions


Syntax	<pre>permissions { storage; storage-control; }</pre>
Hierarchy Level	[edit system login class]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the login access privileges to be provided on the switch.
Options	<i>permissions</i> —Privilege type.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Access Privilege Levels on page 1809• Table 137• user on page 282

retry-options

Syntax	<pre> retry-options { backoff-threshold <i>number</i>; backoff-factor <i>seconds</i>; maximum-time <i>seconds</i>; minimum-time <i>seconds</i>; tries-before-disconnect <i>number</i>; } </pre>
Hierarchy Level	[edit system login]
Release Information	<p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Maximum number of times a user can attempt to enter a password while logging in through SSH or Telnet before being disconnected.
Options	<p>backoff-threshold <i>number</i>—Threshold for the number of failed login attempts before the user experiences a delay when attempting to reenter a password. Use the backoff-factor option to specify the length of delay, in seconds.</p> <p>Range: 1 through 3</p> <p>Default: 2</p> <p>backoff-factor <i>seconds</i>—Length of delay after each failed login attempt. The length of delay increases by this value for each subsequent login attempt after the value specified in the backoff-threshold option.</p> <p>Range: 5 through 10</p> <p>Default: 5</p> <p>maximum-time <i>seconds</i>—Maximum length of time that the connection remains open for the user to enter a username and password to log in. If the user remains idle and does not enter a username and password within the configured maximum-time, the connection is closed.</p> <p>Range: 20 through 300</p> <p>Default: 120</p> <p>minimum-time <i>seconds</i>—Minimum length of time that the connection remains open while the user is attempting to enter a password to log in.</p> <p>Range: 20 through 60</p> <p>Default: 20</p> <p>tries-before-disconnect <i>number</i>—Maximum number of times a user is allowed to attempt to enter a password to log in through SSH or Telnet.</p> <p>Range: 1 through 10</p> <p>Default: 10</p>

Required Privilege	admin—To view this statement in the configuration.
Level	admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Limiting the Number of User Login Attempts for SSH and Telnet Sessions</i>

root-authentication

Syntax	<pre>root-authentication { (encrypted-password "password" load-key-password URL plain-text-password); ssh-dsa "public-key"; ssh-rsa "public-key"; }</pre>
Hierarchy Level	[edit system]
Release Information	<p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Configure the authentication methods for the root-level user, whose username is root .
Options	<p>encrypted-password "password"— Specify the MD5 or other encrypted authentication password. You can specify only one encrypted password.</p> <p>You cannot configure a blank password for the encrypted-password option using blank quotation marks (" "). You must configure a password of 1 through 128 characters and enclose the password in quotation marks.</p>
	<div>  <p>CAUTION: Do not use the encrypted-password option unless the password is <i>already</i> encrypted, and you are entering the encrypted version of the password. If you commit the encrypted-password option with a plain-text password or with blank quotation marks (" "), you will not be able to log in to the device as root, and you will need to use the password recovery process.</p> </div>
	<p>plain-text-password—Plain-text password. The CLI prompts you for the password and then encrypts it. The CLI displays the encrypted version, and the software places the encrypted version in its user database. You can specify only one plain-text password.</p>
	<p>ssh-dsa "public-key"—SSH version 2 authentication. Specify the DSA (SSH version 2) public key. You can specify one or more public keys.</p>
	<p>ssh-rsa "public-key"—SSH version 1 authentication. Specify the RSA (SSH version 1 and SSH version 2) public key. You can specify one or more public keys.</p>
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the Root Password</i> • <i>Recovering the Root Password</i> • authentication on page 274

ssh-dsa

Syntax	<code>ssh-dsa "public-key";</code>
Hierarchy Level	[edit system root-authentication] [edit system login user authentication]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify the DSA (SSH version 2) public key. You can specify one or more public keys.
Options	<code>ssh-dsa "public-key"</code> —SSH version 2 authentication.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Root Password on page 145• authentication on page 274• <i>root-authentication</i>

ssh-rsa

Syntax	<code>ssh-rsa "public-key";</code>
Hierarchy Level	[edit system root-authentication] [edit system login user authentication]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify the RSA (SSH version 1) public key. You can specify one or more public keys.
Options	<code>ssh-rsa "public-key"</code> —SSH version 1 authentication. Specify the RSA (SSH version 1 and SSH version 2) public key. You can specify one or more public keys.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Root Password</i>• authentication on page 274• <i>root-authentication</i>

user (Access)

Syntax	<pre> user username { authentication { (encrypted-password "password" plain-text-password); load-key-file URL; remote-debug-permission (qfabric-admin qfabric-operator qfabric-user); ssh-dsa "public-key" <from hostname>; ssh-rsa "public-key" <from hostname>; } class class-name; full-name "complete-name"; uid uid-value; } </pre>
Hierarchy Level	[edit system login]
Release Information	Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure access permission for individual users.
Options	The remaining statements are explained separately.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Junos OS User Accounts on page 125 • class on page 275

CHAPTER 23

PTP Configuration Statement

- [e2e-transparent on page 311](#)

e2e-transparent

Syntax	e2e-transparent;
Hierarchy Level	[edit protocols ptp]
Release Information	Statement introduced in Junos OS Release 14.1X53-D25 for the QFX Series.
Description	Configure the transparent clock for Precision Time Protocol (PTP). The end to end (E2E) transparent clock mode is supported. With an end-to-end transparent clock, only the residence time is included in the timestamp in the packet. Transparent clock functionality is supported on both PTP over IP and PTP over Ethernet. With PTP over Ethernet, one or two VLANs are supported. Transparent clock functionality is enabled globally and might be required in scenarios in which the interface on which packets are received and transmitted is unknown.
Options	There are no options.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• [edit protocols ptp] Hierarchy Level• Configuring Transparent Clock Mode for Precision Time Protocol on page 152

CHAPTER 24

RADIUS and TACACS+ Accounting Configuration Statements

- [accounting-port \(RADIUS Server\) on page 314](#)
- [destination \(Accounting\) on page 315](#)
- [events on page 316](#)
- [port \(TACACS+ Server\) on page 316](#)
- [radius \(System\) on page 317](#)
- [secret on page 318](#)
- [server \(RADIUS Accounting\) on page 319](#)
- [server \(TACACS+ Accounting\) on page 319](#)
- [single-connection on page 320](#)
- [tacplus on page 320](#)
- [timeout on page 321](#)

accounting-port (RADIUS Server)

Syntax	<code>accounting-port <i>port-number</i>;</code>
Hierarchy Level	[edit system accounting destination radius server <i>server-address</i>], [edit system radius-server <i>server-address</i>]
Release Information	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the accounting port number on which to contact the RADIUS server.
Options	<i>number</i> —Port number on which to contact the RADIUS server. Default: 1813
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring RADIUS Server Authentication• Configuring RADIUS System Accounting on page 1891

destination (Accounting)

```
Syntax destination {
    radius {
        server {
            server-address {
                accounting-port port-number;
                retry number;
                secret password;
                source-address address;
                source-address-inet6 IPv6-source-address;
                timeout seconds;
            }
        }
    }
    tacplus {
        server {
            server-address {
                port port-number;
                secret password;
                single-connection;
                timeout seconds;
            }
        }
    }
}
```

Hierarchy Level [edit system accounting]

Release Information Statement introduced before Junos OS Release 7.4.
radius statement added in Junos OS Release 7.4.
 Statement introduced in Junos OS Release 9.0 for EX Series switches.
 Statement introduced in Junos OS Release 11.1 for the QFX Series.
 Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure the authentication server.

Options **source-address-inet6 IPv6-source-address**—A valid IPv6 address configured on one of the routers or switch interfaces.

The remaining statements are explained separately.

Required Privilege Level system—To view this statement in the configuration.
 system-control—To add this statement to the configuration.

Related Documentation

- [Configuring RADIUS System Accounting on page 1891](#)
- [Configuring TACACS+ System Accounting](#)

events

Syntax	<code>events [<i>events</i>];</code>
Hierarchy Level	[edit system accounting]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the types of events to track and log.
Options	<i>events</i> —Event types; can be one or more of the following: <ul style="list-style-type: none">• change-log—Audit configuration changes.• interactive-commands—Audit interactive commands (any command-line input).• login—Audit logins.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring TACACS+ System Accounting</i>

port (TACACS+ Server)

Syntax	<code>port <i>port-number</i>;</code>
Hierarchy Level	[edit system accounting destination tacplus server <i>server-address</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the port number on which to contact the TACACS+ server.
Options	<i>number</i> —Port number on which to contact the TACACS+ server. Default: 49
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring TACACS+ System Accounting</i>

radius (System)

Syntax	<pre>radius { server { server-address { accounting-port port-number; secret password; source-address address; retry number; timeout seconds; } } }</pre>
Hierarchy Level	[edit system accounting destination]
Release Information	<p>Statement introduced in Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Configure the RADIUS accounting server.
Options	<p>server-address—Address of the RADIUS accounting server.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring RADIUS System Accounting on page 1891

secret

Syntax	<code>secret <i>password</i>;</code>
Hierarchy Level	[edit system accounting destination radius server <i>server-address</i>], [edit system accounting destination tacplus server <i>server-address</i>], [edit system radius-server <i>server-address</i>], [edit system tacplus-server <i>server-address</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the password to use with the RADIUS or TACACS+ server. The secret password used by the local switch must match that used by the server.
Options	<i>password</i> —Password to use; can include spaces included in quotation marks.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring RADIUS Accounting</i>• Configuring RADIUS Authentication (QFX Series or OCX Series) on page 1873• <i>Configuring TACACS+ System Accounting</i>

server (RADIUS Accounting)

Syntax	<pre>server { server-address { accounting-port <i>port-number</i>; retry <i>number</i> secret <i>password</i>; source-address <i>address</i>; timeout <i>seconds</i>; } }</pre>
Hierarchy Level	[edit system accounting destination radius]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure RADIUS logging. The remaining statements are explained separately.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring RADIUS System Accounting on page 1891

server (TACACS+ Accounting)

Syntax	<pre>server { server-address { port <i>port-number</i>; secret <i>password</i>; single-connection; timeout <i>seconds</i>; } }</pre>
Hierarchy Level	[edit system accounting destination tacplus]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure TACACS+ logging. The remaining statements are explained separately.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring TACACS+ System Accounting

single-connection

Syntax	single-connection;
Hierarchy Level	[edit system accounting destination tacplus server <i>server-address</i>], [edit system tacplus <i>server-address</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Optimize attempts to connect to a TACACS+ server. The software maintains one open TCP connection to the server for multiple requests rather than opening a connection for each connection attempt.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring TACACS+ Authentication (QFX Series)</i>• <i>Configuring TACACS+ System Accounting</i>

tacplus

Syntax	<pre>tacplus { server { server-address { port <i>port-number</i>; secret <i>password</i>; single-connection; timeout <i>seconds</i>; } } }</pre>
Hierarchy Level	[edit system accounting destination]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure TACACS+.
Options	<i>server-address</i> —Address of the TACACS+ authentication server. The remaining statements are explained separately.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring TACACS+ System Accounting</i>

timeout

Syntax	<code>timeout seconds;</code>
Hierarchy Level	[edit system radius-server <i>server-address</i>], [edit system tacplus-server <i>server-address</i>], [edit system accounting destination radius server <i>server-address</i>], [edit system accounting destination tacplus server <i>server-address</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the length of time that the local router or switch waits to receive a response from a RADIUS or TACACS+ server.
Options	seconds —Length of time to wait. Range: 1 through 90 seconds Default: 3 seconds
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring RADIUS Accounting</i>• <i>Configuring TACACS+ System Accounting</i>• retry on page 235

CHAPTER 25

Autoinstallation Operational Commands

- `show system autoinstallation status`

show system autoinstallation status

Syntax	show system autoinstallation status
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command supported in Junos OS Release 12.2 for ACX Series Universal Access Routers.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	(ACX Series routers, and EX Series switches only) Display autoinstallation status information.
Options	This command has no options.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>ACX Series Autoinstallation Overview</i> • <i>Before You Begin Autoinstallation on an ACX Series Universal Access Router</i> • <i>Autoinstallation Configuration of ACX Series Universal Access Routers</i> • <i>USB Autoinstallation on ACX Series Routers</i> • <i>Autoinstalling a Configuration File from a Disk-on-Key USB Memory Stick onto an EX2200 or EX3300 Switch</i> • <i>Verifying Autoinstallation on ACX Series Universal Access Routers</i> • <i>autoinstallation</i>
List of Sample Output	show system autoinstallation status on page 325
Output Fields	Table 34 describes the output fields for the show system autoinstallation status command. Output fields are listed in the approximate order in which they appear.

Table 34: show system autoinstallation status Output Fields

Field Name	Field Description
Autoinstallation status	<p>Display autoinstallation status information:</p> <ul style="list-style-type: none"> • Last committed file—File last committed for autoinstallation configuration. • Configuration server of last committed file—IP address or URL of the server configured to retrieve configuration information for the last committed configuration file. • Interface—Interface configured for autoinstallation. <ul style="list-style-type: none"> • Name—Name of the interface. • State—Interface state. • Address acquisition—Display IP address acquired and protocol used for acquisition upon startup. <ul style="list-style-type: none"> • Protocol—Protocol used for acquisition: BOOTP/DHCP or RARP. • Acquired address—IP address acquired from the DHCP server.

Sample Output

show system autoinstallation status

```
user@host> show system autoinstallation status
Autoinstallation status:
Master state: Active
Last committed file: None
Configuration server of last committed file: 0.0.0.0
Interface:
  Name: ge-0/0/1
  State: None
  Address acquisition:
    Protocol: DHCP Client
    Acquired address: None
    Protocol: RARP Client
    Acquired address: None
```


CHAPTER 26

Basic System Management Operational Commands

- commit
- clear log
- clear chassis display message
- clear system commit
- clear system reboot
- file
- file archive
- file checksum md5
- file checksum sha1
- file checksum sha-256
- file compare
- file delete
- file list
- file rename
- file show
- load
- ping
- request chassis beacon
- request chassis fpc
- request chassis pic
- request chassis routing-engine master
- request message
- request system configuration rescue delete
- request system configuration rescue save
- request system halt
- request system logout

- `request system power-off`
- `request system reboot`
- `request system storage cleanup`
- `request system zeroize`
- `restart`
- `save`

commit

Syntax `commit <<at <"string">> <and-quit> <check> <comment <"comment-string">>
<confirmed> <display detail> <minutes> <synchronize><force>>`

Release Information Command introduced before Junos OS Release 7.4.

Description Commit the set of changes to the database and cause the changes to take operational effect.

Options `at <"string">`—(Optional) Save software configuration changes and activate the configuration at a future time, or upon reboot.

string is **reboot** or the future time to activate the configuration changes. Enclose the **string** value (including **reboot**) in quotation marks (" "). You can specify time in two formats:

- A time value in the form **hh:mm[:ss]** (hours, minutes, and optionally seconds)—Commit the configuration at the specified time, which must be in the future but before 11:59:59 PM on the day the **commit at** configuration command is issued. Use 24-hour time for the **hh** value; for example, **04:30:00** is 4:30:00 AM, and **20:00** is 8:00 PM. The time is interpreted with respect to the clock and time zone settings on the router.
- A date and time value in the form **yyyy-mm-dd hh:mm[:ss]** (year, month, date, hours, minutes, and, optionally, seconds)—Commit the configuration at the specified day and time, which must be after the **commit at** command is issued. Use 24-hour time for the **hh** value. For example, **2003-08-21 12:30:00** is 12:30 PM on August 21, 2003. The time is interpreted with respect to the clock and time zone settings on the router.

For example, **commit at "18:00:00"**. For date and time, include both values in the same set of quotation marks. For example, **commit at "2005-03-10 14:00:00"**.

A *commit check* is performed when you issue the **commit at** configuration mode command. If the result of the check is successful, then the current user is logged out of configuration mode, and the configuration data is left in a read-only state. No other commit can be performed until the scheduled commit is completed.



NOTE: If Junos OS fails before the configuration changes become active, all configuration changes are lost.

You cannot enter the **commit at** configuration command when there is a pending reboot.

You cannot enter the **request system reboot** command once you schedule a commit operation for a specific time in the future.

You cannot commit a configuration when a scheduled commit is pending. For information about how to use the **clear** command to cancel a scheduled configuration, see [CLI Explorer](#).

and-quit—(Optional) Commit the configuration and, if the configuration contains no errors and the commit succeeds, exit from configuration mode.

check—(Optional) Verify the syntax of the configuration, but do not activate it.

comment <"*comment-string*">—(Optional) Add a comment that describes the committed configuration. The comment can be as long as 512 bytes and must be typed on a single line. You cannot include a comment with the **commit check** command. Enclose *comment-string* in quotation marks (" "). For example, **commit comment "Includes changes recommended by SW Lab"**.

confirmed <*minutes*>—(Optional) Require that the commit be confirmed within the specified amount of time. To confirm a commit, enter either a **commit** or **commit check** command. If the commit is not confirmed within the time limit, the configuration rolls back automatically to the precommit configuration and a broadcast message is sent to all logged-in users. To show when a rollback is scheduled, enter the **show system commit** command. The allowed range is 1 through 65,535 minutes, and the default is 10 minutes.

In Junos OS Release 11.4 and later, you can also use the **commit confirmed** command in the **[edit private]** configuration mode.

display detail—(Optional) Monitors the commit process.



NOTE: In Junos OS Release 10.4 and later, if the number of commit details or messages exceeds a page when used with the **| display detail** pipe option, the **more** pagination option on the screen is no longer available. Instead, the messages roll up on the screen by default, just like using the **commit** command with the **| no more** pipe option.

synchronize <*force*>—(Optional) If your router has two Routing Engines, you can manually direct one Routing Engine to synchronize its configuration with the other by issuing the **commit synchronize** command. The Routing Engine on which you execute this command (request Routing Engine) copies and loads its candidate configuration to the other (responding Routing Engine). Both Routing Engines then perform a syntax check on the candidate configuration file being committed. If no errors are found, the configuration is activated and becomes the current operational configuration on both Routing Engines. The **commit synchronize** command does not work if the responding Routing Engine has uncommitted configuration changes. However, you can enforce commit synchronization on the Routing Engines by using the **force** option. When you issue the **commit synchronize** command with the **force** option from one Routing Engine, the configuration sessions on the other Routing Engine is terminated and its configuration synchronized with that on the Routing Engine from which you issued the command.



NOTE: When you issue the `commit synchronize` command, you must use the `apply-groups re0` and `re1` commands. For information about how to use groups, see *Disabling Inheritance of a Junos OS Configuration Group*.

The responding Routing Engine must use Junos OS Release 5.0 or later.

Required Privilege Level

`configure`—To enter configuration mode.



NOTE: If you are using Junos OS in a Common Criteria environment, system log messages are created whenever a secret attribute is changed (for example, password changes or changes to the RADIUS shared secret). These changes are logged during the following configuration load operations:

```
load merge
load replace
load override
load update
```

For more information, see the [Secure Configuration Guide for Common Criteria and Junos-FIPS](#).



NOTE: Do not use the `load override` or `load replace` command instead of `set` command in the management software Junos Space or NSM documentation.

Related Documentation

- *Verifying a Junos OS Configuration, Committing a Junos OS Configuration*
- *Scheduling a Junos OS Commit Operation*
- *Deactivating and Reactivating Statements and Identifiers in a Junos OS Configuration*
- *Monitoring the Junos OS Commit Process*
- *Adding a Comment to Describe the Committed Configuration*

clear log

Syntax	<code>clear log <i>filename</i></code> <code><all></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Description	Remove contents of a log file.
Options	<i>filename</i> —Name of the specific log file to delete. all —(Optional) Delete the specified log file and all archived versions of it.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show log on page 1003
List of Sample Output	clear log on page 332
Output Fields	See file list for an explanation of output fields.

Sample Output

clear log

The following sample commands list log file information, clear the contents of a log file, and then display the updated log file information:

```
user@host> file list lcc0-re0:/var/log/sampled detail
lcc0-re0:
-----
-rw-r-----  1 root  wheel          26450 Jun 23 18:47 /var/log/sampled
total 1

user@host> clear log lcc0-re0:sampled
lcc0-re0:
-----

user@host> file list lcc0-re0:/var/log/sampled detail
lcc0-re0:
-----
-rw-r-----  1 root  wheel           57 Sep 15 03:44 /var/log/sampled
total 1
```

clear chassis display message

List of Syntax	Syntax on page 333 Syntax (TX Matrix Router) on page 333 Syntax (TX Matrix Plus Router) on page 333 Syntax (QFabric Systems) on page 333
Syntax	clear chassis display message
Syntax (TX Matrix Router)	clear chassis display message <lcc <i>number</i> scc>
Syntax (TX Matrix Plus Router)	clear chassis display message <lcc <i>number</i> sfc <i>number</i> >
Syntax (QFabric Systems)	clear chassis display message <node-device <i>name</i> interconnect-device <i>name</i> >
Release Information	<p>Command introduced in Junos OS Release 7.5.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>sfc option for the TX Matrix Plus routers introduced in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>(M40e, M160, M320, T Series routers, EX Series, and QFabric systems only) Clear or stop a text message on the craft interface display, which is on the front of the router or switch or on the LCD panel display on the router or switch. The craft interface alternates the display of text messages with standard craft interface messages, switching between messages every 2 seconds. By default, on both the router and the switch, the text message is displayed for 5 minutes. The craft interface display has four 20-character lines. The LCD panel display has two 16-character lines, and text messages appear only on the second line.</p>
Options	<p>none—Clear or stop a text message on the craft interface display.</p> <p>interconnect-device <i>name</i>—(QFabric systems only) (Optional) On a QFabric system, clear or stop a text message on the LCD panel display on the specified Interconnect device.</p> <p>lcc <i>number</i>—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.</p> <p>Replace <i>number</i> with the following values depending on the LCC configuration:</p> <ul style="list-style-type: none"> 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix. 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.

- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

node-device *name*—(QFabric systems only) (Optional) On a QFabric system, clear or stop a text message on the LCD panel display on the specified Node device in a Node group.

scc—(TX Matrix routers only) (Optional) Clear or stop a text message on the craft interface on the TX Matrix router (switch-card chassis).

sfc *number*—(TX Matrix Plus routers only) (Optional) Clear or stop a text message on the craft interface on the TX Matrix Plus router (or switch-fabric chassis). Replace *number* with 0.

Required Privilege Level clear

Related Documentation

- *Configuring the LCD Panel on EX Series Switches (CLI Procedure)*
- *set chassis display message*
- *show chassis craft-interface*

List of Sample Output [clear chassis display message on page 334](#)

Output Fields See *show chassis craft-interface* for an explanation of output fields.

Sample Output

clear chassis display message

The following example displays and then clears the text message on the craft interface display:

```
user@host> show chassis craft-interface
Red alarm:      LED off, relay off
Yellow alarm:   LED off, relay off
Host OK LED:    On
Host fail LED:  Off
FPCs           0  1  2  3  4  5  6  7
-----
Green  ..  *..  *  *.
Red    .....
LCD screen:
      +-----+
      |NOC contact Dusty |
      |(888) 526-1234    |
      +-----+

user@host> clear chassis display message

user@host> show chassis craft-interface
Red alarm:      LED off, relay off
Yellow alarm:   LED off, relay off
```

```
Host OK LED:  On
Host fail LED: Off
FPCs      0  1  2  3  4  5  6  7
-----
Green  ..  *..  *  *.
Red    .....
LCD screen:
+-----+
|host    |
|Up: 0+17:05:47|
|        |
|Temperature OK|
+-----+
```

clear system commit

Syntax	clear system commit
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Clear any pending commit operation.
Options	This command has no options.
Required Privilege Level	maintenance (or the actual user who scheduled the commit)
Related Documentation	<ul style="list-style-type: none">• show system commit on page 1054
List of Sample Output	clear system commit on page 336 clear system commit (None Pending) on page 336 clear system commit (User Does Not Have Required Privilege Level) on page 336
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear system commit

```
user@host> clear system commit
Pending commit cleared.
```

clear system commit (None Pending)

```
user@host> clear system commit
No commit scheduled.
```

clear system commit (User Does Not Have Required Privilege Level)

```
user@host> clear system commit
error: Permission denied
```


clear system reboot

List of Syntax	Syntax on page 337 Syntax (EX Series Switches) on page 337 Syntax (TX Matrix Router) on page 337 Syntax (TX Matrix Plus Router) on page 337 Syntax (QFX Series) on page 337
Syntax	clear system reboot <both-routing-engines>
Syntax (EX Series Switches)	clear system reboot <all-members> <both-routing-engines> <local> <member <i>member-id</i> >
Syntax (TX Matrix Router)	clear system reboot <both-routing-engines> <all-chassis all-lcc lcc <i>number</i> scc>
Syntax (TX Matrix Plus Router)	clear system reboot <both-routing-engines> <all-chassis all-lcc lcc <i>number</i> sfc <i>number</i> >
Syntax (QFX Series)	clear system reboot <infrastructure <i>name</i> > <interconnect-device <i>name</i> > <node-group <i>name</i> >
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Clear any pending system software reboots or halts. When issued on a TX Matrix router without any options, the default behavior clears all pending system software reboots or halts on all T640 routers connected to the TX Matrix router. When issued on a TX Matrix Plus router without any options, the default behavior clears all pending system software reboots or halts on all T1600 or T4000 routers connected to the TX Matrix Plus router.
Options	<p>none—Clear all pending system software reboots or halts.</p> <p>all-chassis—(TX Matrix routers and TX Matrix Plus routers only) (Optional) Clear all halt or reboot requests for all the Routing Engines in the chassis.</p> <p>all-lcc—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, clear all halt or reboot requests for all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, clear all halt or reboot requests on the l connected T1600 or T4000 LCCs.</p>

all-members—(EX4200 switches only) (Optional) Clear all halt or reboot requests on all members of the Virtual Chassis configuration.

both-routing-engines—(Systems with multiple Routing Engines) (Optional) Clear all halt or reboot requests on both Routing Engines. On a TX Matrix router, clear both Routing Engines on all chassis connected to the TX Matrix router. Likewise, on a TX Matrix Plus router, clear both Routing Engines on all chassis connected to the TX Matrix Plus router.

infrastructure *name*—(QFabric systems) (Optional) Clear all halt or reboot requests on the fabric control Routing Engines or fabric manager Routing Engines.

interconnect-device *name*—(QFabric systems) (Optional) Clear all halt or reboot requests on the Interconnect device.

lcc *number*—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, clear all halt or reboot requests for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, clear all halt or reboot requests for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(EX4200 switches only) (Optional) Clear all halt or reboot requests on the local Virtual Chassis member.

member *member-id*—(EX4200 switches only) (Optional) Clear all halt or reboot requests on the specified member of the Virtual Chassis configuration. Replace *member-id* with a value from 0 through 9.

node-group *name*—(QFabric systems) (Optional) Clear all halt or reboot requests on the Node group.

scc—(TX Matrix routers only) (Optional) Clear all halt or reboot requests for the TX Matrix router (or switch-card chassis).

sfc *number*—(TX Matrix Plus routers only) (Optional) Clear all halt or reboot requests for the TX Matrix Plus router. Replace *number* with 0.

Required Privilege Level maintenance

Related Documentation	<ul style="list-style-type: none">• request system reboot on page 6336• Routing Matrix with a TX Matrix Plus Router Solutions Page
List of Sample Output	clear system reboot on page 340 clear system reboot (TX Matrix Router) on page 340 clear system reboot (QFX Series) on page 340
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear system reboot

```
user@host> clear system reboot
reboot requested by root at Sat Dec 12 19:37:34 1998
[process id 17855]
Terminating...
```

clear system reboot (TX Matrix Router)

```
user@host> clear system reboot
scc-re0:
-----
No shutdown/reboot scheduled.
lcc0-re0:
-----
No shutdown/reboot scheduled.
lcc2-re0:
-----
No shutdown/reboot scheduled.
```

clear system reboot (QFX Series)

```
user@switch> clear system reboot node-group node1
No shutdown/reboot scheduled.
```

file

Syntax	file <archive checksum compare copy delete list rename show source address>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Archive files from the device, copy files to and from the router or switch, calculate the file checksum, compare files, delete a file from the device, list files on the device, rename a file, show file contents, or show the local address to initiate a connection.
Options	<p>archive (Optional)—Archive, and optionally compress, one or multiple local system files as a single file, locally or at a remote location.</p> <p>checksum (Optional)—Calculate the Message Digest 5 (MD5) checksum of a file.</p> <p>compare (Optional)—Compare two local files and describe the differences between them in default, context, or unified output styles.</p> <p>copy (Optional)—Copy files from one place to another on the local switch or between the local switch and a remote system.</p> <p>delete (Optional)—Delete a file on the local switch.</p> <p>list (Optional)—Display a list of files on the local switch.</p> <p>rename (Optional)—Rename a file on the local switch.</p> <p>show (Optional)—Display the contents of a file.</p> <p>source address (Optional)—Specify the source address of the local file.</p>
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none"> • <i>Viewing Files and Directories on a Device Running Junos OS</i>

file archive

Syntax	<code>file archive destination <i>destination</i> source <i>source</i> <compress></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Description	Archive, and optionally compress, one or multiple local system files as a single file, locally or at a remote location.
Options	<p>destination <i>destination</i>—Destination of the archived file or files. Specify the destination as a URL or filename. The Junos OS adds one of the following suffixes if the destination filename does not already have it:</p> <ul style="list-style-type: none">• For archived files—The suffix .tar• For archived and compressed files—The suffix .tgz <p>source <i>source</i>—Source of the original file or files. Specify the source as a URL or filename.</p> <p>compress—(Optional) Compress the archived file with the GNU zip (gzip) compression utility. The compressed files have the suffix .tgz.</p>
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none">• Format for Specifying Filenames and URLs in Junos OS CLI Commands on page 40
List of Sample Output	file archive (Multiple Files) on page 342 file archive (Single File) on page 342 file archive (with Compression) on page 343
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

file archive (Multiple Files)

The following sample command archives all message files in the local directory `/var/log/messages` as the single file `messages-archive.tar`.

```
user@host> file archive source /var/log/messages* destination /var/log/messages-archive.tar
/usr/bin/tar: Removing leading / from absolute path names in the archive.
user@host>
```

file archive (Single File)

The following sample command archives one message file in the local directory `/var/log/messages` as the single file `messages-archive.tar`.

```
user@host> file archive source /var/log/messages destination /var/log/messages-archive.tar
/usr/bin/tar: Removing leading / from absolute path names in the archive.
user@host
```

file archive (with Compression)

The following sample command archives and compresses all message files in the local directory **/var/log/messages** as the single file **messages-archive.tgz**.

```
user@host> file archive compress source /var/log/messages* destination
/var/log/messages-archive.tgz
/usr/bin/tar: Removing leading / from absolute path names in the archive.
```

file checksum md5

Syntax	<code>file checksum md5 <pathname> filename</code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Calculate the Message Digest 5 (MD5) checksum of a file.
Options	pathname —(Optional) Path to a filename. filename —Name of a local file for which to calculate the MD5 checksum.
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Checksum Hashes for a Commit Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i>• <i>Configuring Checksum Hashes for an Event Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i>• <i>Configuring Checksum Hashes for an Op Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i>• <i>Executing an Op Script from a Remote Site</i> in the <i>Junos OS Configuration and Operations Automation Guide</i>• file checksum sha-256 on page 346• file checksum sha1 on page 345• <i>op</i>
List of Sample Output	file checksum md5 on page 344
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

file checksum md5

```
user@host> file checksum md5 jbundle-5.3R2.4-export-signed.tgz
MD5 (jbundle-5.3R2.4-export-signed.tgz) = 2a3b69e43f9bd4893729cc16f505a0f5
```


file checksum sha1

Syntax	<code>file checksum sha1 <pathname> filename</code>
Release Information	<p>Command introduced in Junos OS Release 9.5.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	Calculate the Secure Hash Algorithm (SHA-1) checksum of a file.
Options	<p>pathname—(Optional) Path to a filename.</p> <p>filename—Name of a local file for which to calculate the SHA-1 checksum.</p>
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Checksum Hashes for a Commit Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i> • <i>Configuring Checksum Hashes for an Event Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i> • <i>Configuring Checksum Hashes for an Op Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i> • <i>Executing an Op Script from a Remote Site</i> in the <i>Junos OS Configuration and Operations Automation Guide</i> • file checksum md5 on page 344 • file checksum sha-256 on page 346 • <i>op</i>
List of Sample Output	file checksum sha1 on page 345
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

file checksum sha1

```
user@host> file checksum sha1 /var/db/scripts/opscript.slax
```

```
SHA1 (/var/db/scripts/commitscript.slax) = ba9e47120c7ce55cff29afd73eacd370e162c676
```

file checksum sha-256

Syntax	<code>file checksum sha-256 <pathname> filename</code>
Release Information	Command introduced in Junos OS Release 9.5. Command introduced in Junos OS Release 9.5 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Calculate the Secure Hash Algorithm 2 family (SHA-256) checksum of a file.
Options	pathname —(Optional) Path to a filename. filename —Name of a local file for which to calculate the SHA-256 checksum.
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Checksum Hashes for a Commit Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i>• <i>Configuring Checksum Hashes for an Event Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i>• <i>Configuring Checksum Hashes for an Op Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i>• <i>Executing an Op Script from a Remote Site</i> in the <i>Junos OS Configuration and Operations Automation Guide</i>• file checksum md5 on page 344• file checksum sha1 on page 345• <i>op</i>
List of Sample Output	file checksum sha-256 on page 346
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

file checksum sha-256

```
user@host> file checksum sha-256 /var/db/scripts/commitscript.slax

SHA256 (/var/db/scripts/commitscript.slax) =
94c2b061fb55399e15babd2529453815601a602b5c98e5c12ed929c9d343dd71
```

file compare

Syntax	<pre>file compare (files <i>filename filename</i>) <context unified> <ignore-white-space></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.</p>
Description	<p>Compare two local files and describe the differences between them in default, context, or unified output styles:</p> <ul style="list-style-type: none"> • Default—In the first line of output, c means lines were changed between the two files, d means lines were deleted between the two files, and a means lines were added between the two files. The numbers preceding this alphabetical marker represent the first file, and the lines after the alphabetical marker represent the second file. A left angle bracket (<) in front of output lines refers to the first file. A right angle bracket (>) in front of output lines refers to the second file. • Context—The display is divided into two parts. The first part is the first file; the second part is the second file. Output lines preceded by an exclamation point (!) have changed. Additions are marked with a plus sign (+), and deletions are marked with a minus sign (-). • Unified—The display is preceded by the line number from the first and the second file (xx,xxx,x). Before the line number, additions to the file are marked with a plus sign (+), and deletions to the file are marked with a minus sign (-). The body of the output contains the affected lines. Changes are viewed as additions plus deletions.
Options	<p>files <i>filename</i>—Names of two local files to compare.</p> <p>context—(Optional) Display output in context format.</p> <p>ignore-white-space—(Optional) Ignore changes in the amount of white space.</p> <p>unified—(Optional) Display output in unified format.</p>
Required Privilege Level	none
Related Documentation	<ul style="list-style-type: none"> • Format for Specifying Filenames and URLs in Junos OS CLI Commands on page 40 • Viewing Core Files from Junos OS Processes on page 34
List of Sample Output	<p>file compare files on page 349</p> <p>file compare files context on page 349</p> <p>file compare files unified on page 349</p> <p>file compare files unified ignore-white-space on page 349</p>

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

file compare files

```
user@host> file compare files /tmp/one /tmp/two
100c100
<          full-name "File 1";
---
>          full-name "File 2";
102c102
<          class foo; # 'foo' is not defined
---
>          class super-user;
```

file compare files context

```
user@host> file compare files /tmp/one /tmp/two context
*** /tmp/one   Wed Dec  3 17:12:50 2003
--- /tmp/two   Wed Dec  3 09:13:14 2003
*****
*** 97,104 ****
        }
    }
    user bill {
!         full-name "Bill Smith";
!         class foo; # 'foo' is not defined
        authentication {
            encrypted-password SECRET;
        }
--- 97,105 ----
    }
    user bill {
!         full-name "Bill Smith";
!         uid 1089;
!         class super-user;
        authentication {
            encrypted-password SECRET;
        }
    }
```

file compare files unified

```
user@host> file compare files /tmp/one /tmp/two unified
--- /tmp/one   Wed Dec  3 17:12:50 2003
+++ /tmp/two   Wed Dec  3 09:13:14 2003
@@ -97,8 +97,9 @@
    }
}
user bill {
-     full-name "Bill Smith";
-     class foo; # 'foo' is not defined
+     full-name "Bill Smith";
+     uid 1089;
+     class super-user;
    authentication {
        encrypted-passwordSECRET;
    }
}
```

file compare files unified ignore-white-space

```
user@host> file compare files /tmp/one /tmp/two unified ignore-white-space
```

```
--- /tmp/one    Wed Dec  3 09:13:10 2003
+++ /tmp/two    Wed Dec  3 09:13:14 2003
@@ -99,7 +99,7 @@
     user bill {
         full-name "Bill Smith";
         uid 1089;
-        class foo; # 'foo' is not defined
+        class super-user;
         authentication {
             encrypted-password <SECRET>; # SECRET-DATA
         }
     }
```

file delete

Syntax	<code>file delete <i>filename</i></code> <code><purge></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Description	Delete a file on the local router or switch.
Options	<i>filename</i> —Name of the file to delete. For a routing matrix, include chassis information in the filename if the file to be deleted is not local to the Routing Engine from which the command is issued. <i>purge</i> —(Optional) Overwrite regular files before deleting them.
Required Privilege Level	maintenance
List of Sample Output	file delete on page 351 file delete (Routing Matrix) on page 351
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

file delete

```
user@host> file list /var/tmp
dcd.core
rpd.core
snmpd.core

user@host> file delete /var/tmp/snmpd.core
user@host> file list /var/tmp
dcd.core
rpd.core
```

file delete (Routing Matrix)

```
user@host> file list lcc0-re0:/var/tmp
dcd.core
rpd.core
snmpd.core

user@host> file delete lcc0-re0:/var/tmp/snmpd.core
user@host> file list /var/tmp
dcd.core
rpd.core
```

file list

Syntax	<code>file list</code> <code><detail recursive></code> <code><filename></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display a list of files on the local router or switch.
Options	none —Display a list of all files for the current directory. detail recursive —(Optional) Display detailed output or descend recursively through the directory hierarchy, respectively. filename —(Optional) Display a list of files. For a routing matrix, the filename must include the chassis information.
Additional Information	The default directory is the home directory of the user logged in to the router or switch. To view available directories, enter a space and then a backslash (/) after the file list command. To view files within a specific directory, include a backslash followed by the directory and, optionally, subdirectory name after the file list command.
Required Privilege Level	maintenance
List of Sample Output	file list on page 352 file list (Routing Matrix) on page 352
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

file list

```
user@host> file list /var/tmp
dcd.core
rpd.core
snmpd.core
```

file list (Routing Matrix)

```
user@host> file list lcc0-re0:var/tmp
lcc0-re0:
-----
/var/tmp/:
.gdbinit
.pccardd
Test/
chassisd*
chassisd.nathan*
check_time*
```



```
cores/  
diagTestPrep*  
diagtest*  
diagtest.regress*  
do_switchovers*  
dump_test*  
err.manoj.log  
esw_clearstats*  
esw_counter*  
esw_debug*  
esw_debug_ge*  
esw_filt_test*  
esw_filter_tnp_addr*  
esw_getstats*  
esw_phy*  
esw_stats*
```

file rename

Syntax	<code>file rename <i>source destination</i></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Description	Rename a file on the local router or switch.
Options	<i>destination</i> —New name for the file. <i>source</i> —Original name of the file. For a routing matrix, the filename must include the chassis information.
Required Privilege Level	maintenance
List of Sample Output	file rename on page 354 file rename (Routing Matrix) on page 354
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

file rename

The following example lists the files in `/var/tmp`, renames one of the files, and then displays the list of files again to reveal the newly named file.

```
user@host> file list /var/tmp
dcd.core
rpd.core
snmpd.core

user@host> file rename /var/tmp/dcd.core /var/tmp/dcd.core.990413
user@host> file list /var/tmp
dcd.core.990413
rpd.core
snmpd.core
```

file rename (Routing Matrix)

The following example lists the files in `/var/tmp`, renames one of the files, and then displays the list of files again to reveal the newly named file.

```
user@host> file list lcc0-re1:/var/tmp
lcc0-re1:
-----

/var/tmp:
.pccardd
sartre.conf
snmpd
syslogd.core-tarball.0.tgz
```

```
user@host> file rename lcc0-re0:/var/tmp/snmpd /var/tmp/snmpd.rr
user@host> file list lcc0-re1:/var/tmp
lcc0-re1:
```

```
-----

/var/tmp:
.pccardd
sartre.conf
snmpd.rr
syslogd.core-tarball.0.tgz
```

file show

Syntax	<code>file show <i>filename</i></code> <code><encoding (base64 raw)></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display the contents of a file.
Options	<i>filename</i> —Name of a file. For a routing matrix, the filename must include the chassis information. <code>encoding (base64 raw)</code> —(Optional) Encode file contents with base64 encoding or show raw text.
Required Privilege Level	maintenance
List of Sample Output	file show on page 356 file show (Routing Matrix) on page 356
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

file show

```
user@host> file show /var/log/messages
Apr 13 21:00:08 romney /kernel: so-1/1/2: loopback suspected; going to standby.
Apr 13 21:00:40 romney /kernel: so-1/1/2: loopback suspected; going to standby.
Apr 13 21:02:48 romney last message repeated 4 times
Apr 13 21:07:04 romney last message repeated 8 times
Apr 13 21:07:13 romney /kernel: so-1/1/0: Clearing SONET alarm(s) RDI-P
Apr 13 21:07:29 romney /kernel: so-1/1/0: Asserting SONET alarm(s) RDI-P
...
```

file show (Routing Matrix)

```
user@host> file show lcc0-re0:/var/tmp/.gdbinit
lcc0-re0:
-----
#####
# Settings
#####

set print pretty

#####
# Basic stuff
#####

define msgbuf
    printf "%s", msgbufp->msg_ptr
end
```

```
# hex dump of a block of memory
# usage: dump address length
define dump
  p $arg0, $arg1
  set $ch = $arg0
  set $j = 0
  set $n = $arg1
  while ($j < $n)
    #printf "%x %x ",&$ch[$j],$ch[$j]
    printf "%x ",$ch[$j]
    set $j = $j + 1
    if (!($j % 16))
      printf "\n"
    end
  end
end
end
```

load

Syntax	<code>load (factory-default merge override patch replace set update) load (<i>filename</i> terminal) <relative></code>
QFX Series	<code>load (dhcp-snooping <i>filename</i>)</code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Load a configuration from an ASCII configuration file, from terminal input, or from the factory default. Your current location in the configuration hierarchy is ignored when the load operation occurs.



NOTE: `load` can be run from configuration mode only.

Options **dhcp-snooping**—(QFX Series switches) Loads DHCP snooping entries.

factory-default—Loads the factory configuration. The factory configuration contains the manufacturer's suggested configuration settings. The factory configuration is the router or switch's first configuration and is loaded when the router or switch is first installed and powered on.



NOTE: To load the factory default configuration, you must first *unprotect* any protected hierarchies in the configuration.

filename—Name of the file to load. For information about specifying the filename, see *Viewing Files and Directories on a Device Running Junos OS*.

merge—Combine the configuration that is currently shown in the CLI with the configuration.

override—Discard the entire configuration that is currently shown in the CLI and load the entire configuration. Marks every object as changed.

patch—Change part of the configuration and mark only those parts as changed.

replace—Look for a **replace** tag in *filename*, delete the existing statement of the same name, and replace it with the configuration.

set—Merge a set of commands with an existing configuration. This option executes the configuration instructions line by line as they are stored in a file or from a terminal. The instructions can contain any configuration mode command, such as **set**, **edit**, **exit**, and **top**.

relative—(Optional) Execute set of commands until the current edit point.

terminal—Use the text you type at the terminal as input to the configuration. Type Ctrl+d to end terminal input.

update—Discard the entire configuration that is currently shown in the CLI, and load the entire configuration. Marks changed objects only.



NOTE: If you are using Junos OS in a Common Criteria environment, system log messages are created whenever a secret attribute is changed (for example, password changes or changes to the RADIUS shared secret). These changes are logged during the following configuration load operations:

```
load merge
load replace
load override
load update
```

For more information, see the *Secure Configuration Guide for Common Criteria and Junos-FIPS*.

Required Privilege Level configure—To enter configuration mode, but other required privilege levels depend on where the statement is located in the configuration hierarchy.

Related Documentation • [Loading a Configuration from a File on page 1325](#)


ping

List of Syntax [Syntax on page 360](#)
 [Syntax \(QFX Series\) on page 360](#)

Syntax `ping host`
 `<bypass-routing>`
 `<count requests>`
 `<detail>`
 `<do-not-fragment>`
 `<inet | inet6>`
 `<interface source-interface>`
 `<interval seconds>`
 `<logical-system logical-system-name>`
 `<loose-source value>`
 `<mac-address mac-address>`
 `<no-resolve>`
 `<pattern string>`
 `<rapid>`
 `<record-route>`
 `<routing-instance routing-instance-name>`
 `<size bytes>`
 `<source source-address>`
 `<strict >`
 `<strict-source value.>`
 `<tos type-of-service>`
 `<ttl value>`
 `<verbose>`
 `<vpls instance-name>`
 `<wait seconds>`

Syntax (QFX Series) `ping host`
 `<bypass-routing>`
 `<count requests>`
 `<detail>`
 `<do-not-fragment>`
 `<inet>`
 `<interface source-interface>`
 `<interval seconds>`
 `<logical-system logical-system-name>`
 `<loose-source value>`
 `<mac-address mac-address>`
 `<no-resolve>`
 `<pattern string>`
 `<rapid>`
 `<record-route>`
 `<routing-instance routing-instance-name>`
 `<size bytes>`
 `<source source-address>`
 `<strict>`
 `< strict-source value>`
 `<tos type-of-service>`
 `<ttl value>`
 `<verbose>`

<wait *seconds*>

Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Check host reachability and network connectivity. The ping command sends Internet Control Message Protocol (ICMP) ECHO_REQUEST messages to elicit ICMP ECHO_RESPONSE messages from the specified host. Press Ctrl+c to interrupt a ping command.</p>
Options	<p>host—IP address or hostname of the remote system to ping.</p> <p>bypass-routing—(Optional) Bypass the normal routing tables and send ping requests directly to a system on an attached network. If the system is not on a directly attached network, an error is returned. Use this option to ping a local system through an interface that has no route through it.</p> <p>count requests—(Optional) Number of ping requests to send. The range of values is 1 through 2,000,000,000. The default value is an unlimited number of requests.</p> <p>detail—(Optional) Include in the output the interface on which the ping reply was received.</p> <p>do-not-fragment—(Optional) Set the do-not-fragment (DF) flag in the IP header of the ping packets. For IPv6 packets, this option disables fragmentation.</p> <div style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;"> <p> NOTE: In Junos OS Release 11.1 and later, when issuing the ping command for an IPv6 route with the do-not-fragment option, the maximum ping packet size is calculated by subtracting 48 bytes (40 bytes for the IPV6 header and 8 bytes for the ICMP header) from the MTU. Therefore, if the ping packet size (including the 48-byte header) is greater than the MTU, the ping operation might fail.</p> </div> <p>inet—(Optional) Ping Packet Forwarding Engine IPv4 routes.</p> <p>inet6—(Optional) Ping Packet Forwarding Engine IPv6 routes.</p> <p>interface source-interface—(Optional) Interface to use to send the ping requests.</p> <p>interval seconds—(Optional) How often to send ping requests. The range of values, in seconds, is 1 through infinity. The default value is 1.</p> <p>logical-system logical-system-name—(Optional) Name of logical system from which to send the ping requests.</p> <p>Alternatively, enter the set cli logical-system logical-system-name command and then run the ping command. To return to the main router or switch, enter the clear cli logical-system command.</p>

loose-source value—(Optional) Intermediate loose source route entry (IPv4). Open a set of values.

mac-address mac-address—(Optional) Ping the physical or hardware address of the remote system you are trying to reach.

no-resolve—(Optional) Do not attempt to determine the hostname that corresponds to the IP address.

pattern string—(Optional) Specify a hexadecimal fill pattern to include in the ping packet.

rapid—(Optional) Send ping requests rapidly. The results are reported in a single message, not in individual messages for each ping request. By default, five ping requests are sent before the results are reported. To change the number of requests, include the **count** option.

record-route—(Optional) Record and report the packet's path (IPv4).

routing-instance routing-instance-name—(Optional) Name of the routing instance for the ping attempt.

size bytes—(Optional) Size of ping request packets. The range of values, in bytes, is 0 through 65,468. The default value is 56, which is effectively 64 bytes because 8 bytes of ICMP header data are added to the packet.

source source-address—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (lo.0).

strict—(Optional) Use the strict source route option (IPv4).

strict-source value—(Optional) Intermediate strict source route entry (IPv4). Open a set of values.

tos type-of-service—(Optional) Set the type-of-service (ToS) field in the IP header of the ping packets. The range of values is 0 through 255.

If the device configuration includes the **dscp-code-point value** statement at the **[edit class-of-service host-outbound-traffic]** hierarchy level, the configured DSCP value overrides the value specified in this command option. In this case, the ToS field of ICMP echo request packets sent on behalf of this command carries the DSCP value specified in the **dscp-code-point** configuration statement instead of the value you specify in this command option.

ttl value—(Optional) Time-to-live (TTL) value to include in the ping request (IPv6). The range of values is 0 through 255.

verbose—(Optional) Display detailed output.

vpls instance-name—(Optional) Ping the instance to which this VPLS belongs.

wait seconds—(Optional) Maximum wait time, in seconds, after the final packet is sent. If this option is not specified, the default delay is 10 seconds. If this option is used without the count option, a default count of 5 packets is used.

Required Privilege Level	network
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Junos OS ICMPv4 Rate Limit for ICMPv4 Routing Engine Messages</i>
List of Sample Output	ping hostname on page 363 ping hostname rapid on page 363 ping hostname size count on page 363
Output Fields	<p>When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. These packets are not counted in the received packets count. They are accounted for separately.</p>

Sample Output

ping hostname

```
user@host> ping skye
PING skye.net (192.168.169.254): 56 data bytes
64 bytes from 192.168.169.254: icmp_seq=0 ttl=253 time=1.028 ms
64 bytes from 192.168.169.254: icmp_seq=1 ttl=253 time=1.053 ms
64 bytes from 192.168.169.254: icmp_seq=2 ttl=253 time=1.025 ms
64 bytes from 192.168.169.254: icmp_seq=3 ttl=253 time=1.098 ms
64 bytes from 192.168.169.254: icmp_seq=4 ttl=253 time=1.032 ms
64 bytes from 192.168.169.254: icmp_seq=5 ttl=253 time=1.044 ms
^C [abort]
```

ping hostname rapid

```
user@host> ping skye rapid
PING skye.net (192.168.169.254): 56 data bytes
!!!!
--- skye.net ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.956/0.974/1.025/0.026 ms
```

ping hostname size count

```
user@host> ping skye size 200 count 5
PING skye.net (192.168.169.254): 200 data bytes
208 bytes from 192.168.169.254: icmp_seq=0 ttl=253 time=1.759 ms
208 bytes from 192.168.169.254: icmp_seq=1 ttl=253 time=2.075 ms
208 bytes from 192.168.169.254: icmp_seq=2 ttl=253 time=1.843 ms
208 bytes from 192.168.169.254: icmp_seq=3 ttl=253 time=1.803 ms
208 bytes from 192.168.169.254: icmp_seq=4 ttl=253 time=17.898 ms

--- skye.net ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max = 1.759/5.075/17.898 ms
```

request chassis beacon

Syntax (QFX Series)	<code>request chassis beacon</code> <code><all (off on)></code> <code><fpc slot-number (off on)></code> <code><interconnect-device name (cb slot-number fpc slot-number (off on)></code> <code><node-device name (off on)></code>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	(QFX Series only) Enable or disable the beacon LED on a QFX Series device.
Options	<p>all—Turn the beacon LED either on or off on all QFabric system Interconnect and Node devices.</p> <p>cb slot-number—Turn the beacon LED either on or off on the Control Board of the QFX3008-I Interconnect device.</p> <p>fpc slot-number—Turn the beacon LED either on or off on the Flexible PIC Concentrator on the standalone QFX3500 switch or the Interconnect device.</p> <p>interconnect-device name—Turn the beacon LED either on or off on the Interconnect device.</p> <p>node-device name—Turn the beacon LED either on or off on the Node device.</p> <p>off—Turn the beacon LED off.</p> <p>on—Turn the beacon LED on.</p>
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none">• show chassis beacon on page 474
List of Sample Output	request chassis beacon fpc 0 on (QFX Series) on page 364 request chassis beacon node-device (QFabric System) on page 364 request chassis beacon on interconnect-device fpc (QFabric System) on page 365
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis beacon fpc 0 on (QFX Series)

```
user@switch> request chassis beacon fpc 0 on  
  
Beacon set to ON
```

request chassis beacon node-device (QFabric System)

```
user@switch> request chassis beacon node-device node1 on
```

node1 ON

request chassis beacon on interconnect-device fpc (QFabric System)

user@switch> request chassis beacon on interconnect-device fpc 2

FPC 2 ON

request chassis fpc

List of Syntax	Syntax on page 366 Syntax (TX Matrix and TX Matrix Plus Routers) on page 366 Syntax (MX Series Routers) on page 366 Syntax (MX2020 3D Universal Edge Routers) on page 366 Syntax (MX2010 3D Universal Edge Routers) on page 366 Syntax (QFabric System) on page 366 Syntax (PTX Series Packet Transport Routers) on page 366
Syntax	<code>request chassis fpc (offline online restart) slot <i>slot-number</i></code>
Syntax (TX Matrix and TX Matrix Plus Routers)	<code>request chassis fpc (offline online restart) slot <i>slot-number</i> <lcc <i>number</i>></code>
Syntax (MX Series Routers)	<code>request chassis fpc (offline online restart) slot <i>slot-number</i> <all-members> <local> <member <i>member-id</i>></code>
Syntax (MX2020 3D Universal Edge Routers)	<code>request chassis fpc (offline online restart) slot <i>slot-number</i></code>
Syntax (MX2010 3D Universal Edge Routers)	<code>request chassis fpc (offline online restart) slot <i>slot-number</i></code>
Syntax (QFabric System)	<code>request chassis fpc <interconnect-device <i>name</i> slot <i>slot-number</i> (offline online)> <(offline online) interconnect-device <i>name</i> slot <i>slot-number</i>> <slot <i>slot-number</i> interconnect-device <i>name</i> (offline online)></code>
Syntax (PTX Series Packet Transport Routers)	<code>request chassis fpc (offline online restart) slot <i>slot-number</i></code>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS 11.3 for QFX Series.</p> <p>Command introduced in Junos OS Release 12.1x48 for PTX Series Packet Transport Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	(M20, M40, M40e, M120, M160, M320, MX Series, and T Series routers, QFabric systems, EX Series switches, and PTX Series Packet Transport Routers only) Control the operation of the Flexible PIC Concentrator (FPC). For information about the meaning of “FPCs” on the switches, see <i>EX Series Switches Hardware and CLI Terminology Mapping</i> .



NOTE: Beginning in Junos OS Release 12.3, it is possible that FPCs brought offline using the `request chassis fpc slot fpc-slot offline` operational-mode CLI command can come online during a configuration commit or power-supply replacement procedure. As an alternative, use the `set fpc fpc-slot power off` configuration-mode command at the `[edit chassis]` hierarchy level to ensure that the FPCs remain offline.



NOTE: If a CLI-based firmware upgrade is in progress, it prevents the specified FPC from restarting. Starting in Junos OS Release 15.1, the following message is displayed:

```
user@host> request chassis fpc slot 0 restart
FPC 0 Firmware update in progress. Wait!!!
```



NOTE: The command `request chassis fpc (offline | online | restart) slot slot-number` is not supported on PTX1000 router.

Options **offline**—Take the FPC offline.

online—Bring the FPC online.

interconnect-device *name*—(QFabric systems only) Bring the Flexible Port Concentrator (FPC) on the QFX3008-I Interconnect device either offline or online:

- (QFabric System) On a QFabric system, specify the name of the QFX3008-I Interconnect device containing the Flexible Port Concentrator (FPC) you want to bring either offline or online.

restart—Restart the FPC.

slot *slot-number*—FPC slot number:

- M20 router—0 through 3.
- M120 router—0 through 5.
- MX240 router—0 through 2. On the MX240 router, slot-number corresponds to the Dense Port Concentrator (DPC) slot number. If an MPC is installed, slot-number corresponds to the MPC slot number.
- MX480 router—0 through 5. On the MX480 router, slot-number corresponds to the Dense Port Concentrator (DPC) slot number. If an MPC is installed, slot-number corresponds to the MPC slot number.

- MX960 router—0 through 11. On the MX960 router, slot-number corresponds to the Dense Port Concentrator (DPC) slot number. If an MPC is installed, slot-number corresponds to the MPC slot number.
- MX2020 router—0 through 19.
- MX2010 router—0 through 9.
- TX Matrix and TX Matrix Plus routers only—On the TX Matrix router, if you specify the number of the T640 router by using the **lcc number** option (the recommended method), replace **slot-number** with a value from 0 through 7. Otherwise, replace **slot-number** with a value from 0 through 31.

Likewise, on a TX Matrix Plus router, if you specify the number of the T1600 or T4000 router by using the **lcc number** option (the recommended method), replace **slot-number** with a value from 0 through 7. Otherwise, replace **slot-number** with a value from 0 through 31. In case of TX Matrix Plus router with 3D SIBs, replace **slot-number** with a value from 0 through 63. For example, the following commands have the same result:

```
user@host> request chassis fpc lcc 1 slot 1 offline
user@host> request chassis fpc slot 9 offline
```

- Other routers—0 through 7.
- QFabric System—Replace **slot-number** with a value from 0 through 2.
- EX Series switches:
 - EX4200 switches in a Virtual Chassis configuration—Replace **slot-number** with a value from 0 through 9.
 - EX6210 switches—Replace **slot-number** with a value from 0 through 9.



NOTE: These commands are not supported for slots 4 and 5 when a Switch Fabric and Routing Engine (SRE) module is installed in those slots. These commands are supported for slots 4 and 5 only if a line card is installed in them.

- EX8208 switches—Replace **slot-number** with a value from 0 through 7.
- EX8216 switches—Replace **slot-number** with a value from 0 through 15.
- PTX5000 Packet Transport Router—Replace **slot-number** with a value from 0 through 7.

all-members—(MX Series routers only) (Optional) Change FPC status of all members of the Virtual Chassis configuration.

local—(MX Series routers only) (Optional) Change FPC status of the local Virtual Chassis member.

member *member-id*—(MX Series routers only) (Optional) Change FPC status of the specified member of the Virtual Chassis configuration. Replace ***member-id*** with a value of 0 or 1.

lcc *number*—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

Required Privilege Level maintenance

Related Documentation

- [show chassis fpc on page 703](#)
- *show chassis fpc-feb-connectivity*
- [show chassis fabric fpcs on page 585](#)
- *Configuring the Junos OS to Make a Flexible PIC Concentrator Stay Offline*
- *Configuring the Junos OS to Resynchronize FPC Sequence Numbers with Active FPCs when an FPC Comes Online*
- *MX960 Flexible PIC Concentrator Description*

List of Sample Output

- [request chassis fpc on page 369](#)
- [request chassis fpc \(MX Series Routers with Media Services Blade \[MSB\]\) on page 369](#)
- [request chassis fpc \(MX2020 Router\) on page 370](#)
- [request chassis fpc \(MX2010 Router\) on page 370](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis fpc

```
user@host> request chassis fpc online slot 0
FPC 0 already online
```

request chassis fpc (MX Series Routers with Media Services Blade [MSB])

```
user@host> request chassis fpc slot 0
Possible completions:
  offline           Take FPC offline
  online            Bring FPC online
```

restart Restart FPC

request chassis fpc (MX2020 Router)

```
user@host >request chassis fpc online slot 2
FPC 2 already online
```

request chassis fpc (MX2010 Router)

```
user@host >request chassis fpc offline slot 5
Offline initiated, use "show chassis fpc" to verify
```

request chassis pic

List of Syntax	Syntax on page 371 Syntax (ACX4000 Series Routers) on page 371 Syntax (MX Series Routers) on page 371 Syntax (TX Matrix and TX Matrix Plus Routers) on page 371
Syntax	<code>request chassis pic (offline online) fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i></code>
Syntax (ACX4000 Series Routers)	<code>request chassis pic (offline online) fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i></code>
Syntax (MX Series Routers)	<code>request chassis pic (offline online) fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i> <member <i>member-id</i>></code>
Syntax (TX Matrix and TX Matrix Plus Routers)	<code>request chassis pic (offline online) fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i> <lcc <i>number</i>></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.3 for ACX4000 Routers. Command introduced in Junos OS Release 13.2 for the QFX Series. Option member introduced in Junos OS Release 14.2 for MX Series routers.
Description	Control the operation of the PIC.



NOTE: The `request chassis pic (offline | online) fpc-slot slot number pic-slot slot-number` command is not supported for built-in PICs on MX Series routers.

To view a list of built-in PICs on the router or switch chassis, use the `show chassis hardware` command.



NOTE: This command is not supported on MX960 and MX2020 routers with MPC5EQ.



NOTE: T1600 routers and TX Matrix Plus routers with 100-Gigabit Ethernet PICs require two adjacent PIC slots, 0 and 1, for each PIC. Therefore, only online and offline command options to PIC slot 0 are allowed. Use of the online and offline command options for PIC slot 1 with the described router and PIC combination is not allowed.



NOTE: In T Series routers, when the PIC state is set from offline to online or vice-versa before the processing is complete for the previous command, you are provided feedback on the status of your request. The following sample messages are displayed if you try to set a PIC offline or online:

```
user@switch> request chassis pic fpc-slot 1 pic-slot 0 online
fpc 1 pic 0 online initiated, use "show chassis fpc pic-status" to verify
```

```
user@switch> request chassis pic fpc-slot 1 pic-slot 0 online
FPC 1 PIC 0 already transitioning to online
```

When the same PIC is set to a different state while the transition is in progress, you are provided feedback on the status of your request.

```
user@switch> request chassis pic fpc-slot 1 pic-slot 0 offline
FPC 1, PIC 0 already transitioning to online. Please retry later.
```



NOTE: If a CLI-based firmware upgrade is in progress, it prevents the specified PIC from restarting. Starting in Junos OS Release 15.1, the following message is displayed:

```
user@host> request chassis pic fpc-slot 0 pic-slot 1 offline
PIC's Firmware update in progress. Wait!!!
```



NOTE: The command `request chassis pic (offline | online) fpc-slot slot-number pic-slot slot-number` is not supported on PTX1000 routers.

Options **offline**—Take the PIC offline.

online—Bring the PIC online.

fpc-slot *slot-number*—Flexible PIC Concentrator (FPC) slot number. Replace *slot-number* with a value appropriate for your router or switch:

- ACX4000 routers—1 or 2.
- EX Series switches:
 - EX3200 switches and EX4200 standalone switches—0.
 - EX4200 switches in a Virtual Chassis configuration—0 through 9 (switch's member ID).
 - EX8208 switches—0 through 7 (line card).
 - EX8216 switches—0 through 15 (line card).
- M5, M7i, M10, and M10i routers—0 or 1.
- M20 routers—0 through 3.

- M40 and M40e routers—0 through 7.
- M120 routers—0 through 5.
- M160 routers—0 through 7.
- M320 routers—0 through 7.
- MX 5, MX10, and MX40 routers—0 or 1.
- MX80 routers—0 or 1.
- MX240 routers—0 through 2
- MX480 routers—0 through 5
- MX2020 routers—0 through 19.
- MX2010 routers—0 through 9.
- MX960 routers—0 through 11.
- PTX5000 routers—0 or 1.
- T Series routers—0 through 7.
- TX Matrix and TX Matrix Plus routers only—On a TX Matrix router, if you specify the number of the T640 router by using the **lcc number** option (the recommended method), replace **slot-number** with a value from 0 through 7. Otherwise, replace **slot-number** with a value from 0 through 31.

Likewise, on a TX Matrix Plus router, if you specify the **number** of the T1600 or T4000 router by using the lcc number option (the recommended method), replace **slot-number** with a value from 0 through 7. Otherwise, for the FPC slot number, replace **slot-number** with a value from 0 through 31. On a TX Matrix Plus router with 3D SIBs to assign the FPC slot number, replace **slot-number** with a value from 0 through 63. For example, the following commands have the same result:

```
user@host> request chassis pic fpc-slot 1 lcc 1 pic-slot 0 offline
user@host> request chassis pic fpc-slot 9 pic-slot 0 offline
```

- QFX5100 standalone switches—0.

lcc number—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace **number** with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

member *member-id*—(MX Series routers only) (Optional) Change the PIC status on the specified member of the Virtual Chassis configuration. Replace *member-id* with the value that is assigned to the specified member.

offline—Take the PIC offline.

online—Bring the PIC online.

pic-slot *slot-number*—PIC slot number.

- EX3200 and EX4200 switches—0 for built-in network interfaces and 1 for interfaces on uplink modules.
- EX8208 and EX8216 switches—0.
- M Series routers—0, 1, 2, or 3
- MX960 router—***slot-number*** corresponds to the slot number of the Packet Forwarding Engine.
- PTX5000 routers—0 or 1.
- T320 router—0 or 1.
- T640 router—0, 1, 2, or 3.
- T1600 router —0, 1, 2, or 3.
- T4000 router—0, 1, 2, or 3.
- QFX5100 standalone switches—0, 1, or 2. PIC 0 is used for all interfaces that are not configured on expansion modules, and PIC 1 and PIC 2 are used for interfaces configured on expansion modules.

Required Privilege Level maintenance

Related Documentation

- [show chassis hardware on page 743](#)
- [show chassis pic on page 954](#)

List of Sample Output [request chassis pic on page 374](#)
[request chassis pic online member \(MX Series Routers\) on page 374](#)
[request chassis pic offline member \(MX Series Routers\) on page 375](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis pic

```
user@host> request chassis pic pic-slot 0 online fpc-slot 0
FPC 0, PIC 0 is already online
```

request chassis pic online member (MX Series Routers)

```
user@host> request chassis pic online member 1 fpc-slot 11 pic-slot 3
```

```
fpc 11 pic 3 online initiated
```

request chassis pic offline member (MX Series Routers)

```
user@host> request chassis pic offline member 1 fpc-slot 11 pic-slot 3  
fpc 11 pic 3 offline initiated
```

request chassis routing-engine master

List of Syntax	Syntax on page 376 Syntax (M Series, MX Series, T Series Routers) on page 376 Syntax (TX Matrix Routers) on page 376 Syntax (TX Matrix Plus Routers) on page 376 Syntax (MX Series Virtual Chassis) on page 376 Syntax (QFX Series) on page 376
Syntax	request chassis routing-engine master (acquire release switch) <no-confirm>
Syntax (M Series, MX Series, T Series Routers)	request chassis routing-engine master (acquire release switch) <no-confirm> <check>
Syntax (TX Matrix Routers)	request chassis routing-engine master (acquire release switch) (lcc <i>number</i> scc all-chassis) <no-confirm>
Syntax (TX Matrix Plus Routers)	request chassis routing-engine master (acquire release switch) (lcc <i>number</i> sfc all-chassis all-lcc) <no-confirm>
Syntax (MX Series Virtual Chassis)	request chassis routing-engine master (acquire release switch) <all-members> <check> <local> <member <i>member-id</i> > <no-confirm>
Syntax (QFX Series)	request chassis routing-engine master (release switch) <check> <interconnect-device <i>name</i> > <node-group <i>name</i> > <no-confirm>
Release Information	Command introduced before Junos OS Release 7.4. all-chassis option added in Junos OS Release 8.0. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.3 for QFX Series. Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers. Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers. Command introduced in Junos OS Release 13.2 for MX104 3D Universal Edge Routers. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	For routers or switches with multiple Routing Engines, control which Routing Engine is the master.



CAUTION: (Routing matrix based on the TX Matrix or TX Matrix Plus routers only) Within the routing matrix, we recommend that all Routing Engines run the same Junos OS Release. If you run different releases on the Routing Engines and a change in mastership occurs on any backup Routing Engine in the routing matrix, one or all routers (in a routing matrix based on the TX Matrix router or in a routing matrix based on a TX Matrix Plus router) might become logically disconnected from the TX Matrix router and cause data loss. For more information, see the [TX Matrix Router Hardware Guide](#) or the *Junos OS High Availability Library for Routing Devices*.



NOTE: Successive graceful Routing Engine switchover events must be a minimum of 240 seconds (4 minutes) apart after both Routing Engines have come up.

If the router or switch displays a warning message similar to “Standby Routing Engine is not ready for graceful switchover. Packet Forwarding Engines that are not ready for graceful switchover might be reset,” do not attempt switchover. If you choose to proceed with switchover, only the Packet Forwarding Engines that were not ready for graceful switchover are reset. None of the Flexible PIC concentrators (FPCs) should spontaneously restart. We recommend that you wait until the warning no longer appears and then proceed with the switchover.

You will receive an error message stating “Command aborted. Not ready for mastership switch, try after n seconds” when this command is re-entered before 240 seconds have elapsed on EX Series switches.



NOTE: On a QFabric system, to avoid traffic loss on the network Node group, switch mastership of the routing engine to the backup routing engine, and then reboot.

Options **acquire**—Attempt to become the master Routing Engine.

release—Request that the other Routing Engine become the master.

switch—Toggle mastership between Routing Engines.



NOTE: The **acquire** option should be used with caution because acquiring a Routing Engine may result in a corrupted database. If possible, use the **switch** option instead.

The **acquire**, **release**, and **switch** options have the following suboptions:

all-chassis—(TX Matrix and TX Matrix Plus routers only) On a routing matrix composed of a TX Matrix router and the attached T640 routers, switch mastership on all the Routing Engines in the routing matrix. Likewise, on a routing matrix composed of a TX Matrix Plus router and the attached T1600 or T4000 routers, switch mastership on all the Routing Engines in the routing matrix.

all-lcc—(TX Matrix Plus routers only) Request to acquire mastership for all line-card chassis (LCC).

all-members—(MX Series routers only) (Optional) Control Routing Engine mastership on the Routing Engines in all member routers of the Virtual Chassis configuration.

check—(QFabric systems, MX104, MX480, MX960, MX2010, and MX2020 routers, and PTX5000 routers only) (Optional) Available with the **switch**, **release**, and **acquire** options. Check graceful switchover status of the standby Routing Engine before toggling mastership between Routing Engines.

interconnect-device *name*—(QFabric systems only) (Optional) Control Routing Engine mastership on the Routing Engines on an Interconnect device.

lcc *number*—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(MX Series routers only) (Optional) Control Routing Engine mastership on the Routing Engines in the local Virtual Chassis member.

member *member-id*—(MX Series routers only) (Optional) Control Routing Engine mastership on the Routing Engines of the specified member in the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

no-confirm—(Optional) Do not request confirmation for the switch.

node-group *name*—(QFabric systems only) (Optional) Control Routing Engine mastership on the Routing Engines on a Node group.

scc—(TX Matrix routers only) TX Matrix (switch-card chassis).

sfc—(TX Matrix Plus routers only) TX Matrix Plus router (or switch-fabric chassis).

Additional Information Because both Routing Engines are always running, the transition from one to the other as the master Routing Engine is immediate. However, the changeover interrupts communication to the System and Switch Board (SSB). The SSB takes several seconds to reinitialize the Flexible PIC Concentrators (FPCs) and restart the PICs. Interior gateway protocol (IGP) and BGP convergence times depend on the specific network environment.

By default, the Routing Engine in slot 0 (**RE0**) is the master and the Routing Engine in slot 1 (**RE1**) is the backup. To change the default master Routing Engine, include the **routing-engine** statement at the **[edit chassis redundancy]** hierarchy level in the configuration. For more information, see the *Junos OS Administration Library for Routing Devices*.

To have the backup Routing Engine become the master Routing Engine, use the **request chassis routing-engine master switch** command. If you use this command to change the master and then restart the chassis software for any reason, the master reverts to the default setting.



NOTE: Although the configurations on the two Routing Engines do not have to be the same and are not automatically synchronized, we recommend making both configurations the same.

Required Privilege Level maintenance

Related Documentation

- [show chassis routing-engine on page 970](#)
- [Configuring Routing Engine Redundancy](#)
- [Switching the Global Master and Backup Roles in a Virtual Chassis Configuration](#)

List of Sample Output

- [request chassis routing-engine master acquire on page 379](#)
- [request chassis routing-engine master switch on page 379](#)
- [request chassis routing-engine master switch check on page 380](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

[request chassis routing-engine master acquire](#)

```
user@host> request chassis routing-engine master acquire

warning: Traffic will be interrupted while the PFE is re-initialized

warning: The other routing engine's file system could be corrupted

Reset other routing engine and become master ? [yes,no] (no)
```

[request chassis routing-engine master switch](#)

```
user@host> request chassis routing-engine master switch
```

```
warning: Traffic will be interrupted while the PFE is re-initialized
Toggle mastership between Routing Engines ? [yes,no] (no) yes
```

```
Resolving mastership...
Complete. The other Routing Engine becomes the master.
```

Switch mastership back to the local Routing Engine:

```
user@host> request chassis routing-engine master switch
```

```
warning: Traffic will be interrupted while the PFE is re-initialized
Toggle mastership between routing engines ? [yes,no] (no) yes
```

```
Resolving mastership...
Complete. The local routing engine becomes the master.
```

request chassis routing-engine master switch check

Usage shown for M Series, MX Series, and T Series routers.

```
{master}[edit]
```

```
user@host> request chassis routing-engine master switch check
```

```
warning: Standby Routing Engine is not ready for graceful switchover.
```

```
{master}[edit]
```

```
user@host> request chassis routing-engine master switch check
Switchover Ready
```

You can similarly check the backup Routing Engine.

request message

Syntax	<code>request message all message "text"</code> <code>request message message "text" (terminal <i>terminal-name</i> user <i>user-name</i>)</code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Display a message on the screens of all users who are logged in to the router or switch or on specific screens.
Options	all —Display a message on the terminal of all users who are currently logged in. message "text" —Message to display. terminal <i>terminal-name</i> —Name of the terminal on which to display the message. user <i>user-name</i> —Name of the user to whom to direct the message.
Required Privilege Level	maintenance
List of Sample Output	request message message on page 381
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

request message message

```
user@host> request message message "Maintenance window in 10 minutes" user maria
Message from user@host on tty0 at 20:27 ...
Maintenance window in 10 minutes
EOF
```

request system configuration rescue delete

Syntax request system configuration rescue delete

Release Information Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.

Description Delete an existing rescue configuration.



NOTE: The [edit system configuration] hierarchy is not available on QFabric systems.

Options This command has no options.

Required Privilege Level maintenance

Related Documentation

- [request system configuration rescue save on page 383](#)
- [request system software rollback on page 1292](#)
- [show system commit on page 1054](#)

List of Sample Output [request system configuration rescue delete on page 382](#)


Output Fields This command produces no output.

Sample Output

request system configuration rescue delete

```
user@host> request system configuration rescue delete
```

request system configuration rescue save

Syntax	request system configuration rescue save
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.</p>
Description	Save the most recently committed configuration as the rescue configuration so that you can return to it at any time by using the rollback command.
<div>  <p>NOTE: The [edit system configuration] hierarchy is not available on QFabric systems.</p> </div>	
Options	This command has no options.
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none"> • request system software delete on page 1286 • request system software rollback on page 1292 • show system commit on page 1054
List of Sample Output	request system configuration rescue save on page 383
Output Fields	This command produces no output.

Sample Output

request system configuration rescue save

```
user@host> request system configuration rescue save
```

request system halt

List of Syntax	Syntax on page 384 Syntax (EX Series Switches) on page 384 Syntax (PTX Series) on page 384 Syntax (TX Matrix Router) on page 384 Syntax (TX Matrix Plus Router) on page 384 Syntax (MX Series Router) on page 385 Syntax (QFX Series) on page 385
Syntax	<code>request system halt</code> <code><at <i>time</i>></code> <code><backup-routing-engine></code> <code><both-routing-engines></code> <code><other-routing-engine></code> <code><in <i>minutes</i>></code> <code><media (compact-flash disk removable-compact-flash usb)></code> <code><message "<i>text</i>"></code>
Syntax (EX Series Switches)	<code>request system halt</code> <code><all-members></code> <code><at <i>time</i>></code> <code><backup-routing-engine></code> <code><both-routing-engines></code> <code><in <i>minutes</i>></code> <code><local></code> <code><media (external internal)></code> <code><member <i>member-id</i>></code> <code><message "<i>text</i>"></code> <code><other-routing-engine></code> <code><slice <i>slice</i>></code>
Syntax (PTX Series)	<code>request system halt</code> <code><at <i>time</i>></code> <code><backup-routing-engine></code> <code><both-routing-engines></code> <code><other-routing-engine></code> <code><in <i>minutes</i>></code> <code><media (compact-flash disk)></code> <code><message "<i>text</i>"></code>
Syntax (TX Matrix Router)	<code>request system halt</code> <code><all-lcc lcc <i>number</i> scc></code> <code><at <i>time</i>></code> <code><backup-routing-engine></code> <code><both-routing-engines></code> <code><other-routing-engine></code> <code><in <i>minutes</i>></code> <code><media (compact-flash disk)></code> <code><message "<i>text</i>"></code>
Syntax (TX Matrix Plus Router)	<code>request system halt</code> <code><all-chassis all-lcc lcc <i>number</i> sfc <i>number</i>></code>

	<pre> <at <i>time</i>> <backup-routing-engine> <both-routing-engines> <other-routing-engine> <in <i>minutes</i>> <media (compact-flash disk)> <message "<i>text</i>"> </pre>
Syntax (MX Series Router)	<pre> request system halt <all-members> <at <i>time</i>> <backup-routing-engine> <both-routing-engines> <in <i>minutes</i>> <local> <media (external internal)> <member <i>member-id</i>> <message "<i>text</i>"> <other-routing-engine> </pre>
Syntax (QFX Series)	<pre> request system halt <all-members> <at <i>time</i>> <both-routing-engines> <director-device <i>director-device-id</i>> <in <i>minutes</i>> <local> <media > <member <i>member-id</i>> <message "<i>text</i>"> <other-routing-engine> <slice <i>slice</i>> </pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>other-routing-engine option introduced in Junos OS Release 8.0.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>director-device option introduced for QFabric systems in Junos OS Release 12.2.</p> <p>backup-routing-engine option introduced in Junos OS Release 13.1.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Stop the router or switch software.



NOTE: When you issue this command on an individual component—for example, a Node device—in a QFabric system, you will receive a warning that says “Hardware-based members will halt, Virtual Junos Routing Engines will reboot.” If you want to halt only one member of a Node group, issue this command with the **member** option on the Node device CLI, because you cannot issue this command from the QFabric CLI. Also, issuing this command might cause traffic loss on an individual component.

When you issue this command on a QFX5100 switch, you are not prompted to reboot. You must power cycle the switch to reboot.

Options **none**—Stop the router or switch software immediately.

all-chassis—(TX Matrix and TX Matrix Plus routers only) (Optional) Halt all chassis.

all-lcc—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, halt all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, halt all T1600 routers (or line-card chassis) connected to the TX Matrix Plus router.

all-members—(EX4200 switches and MX Series routers only) (Optional) Halt all members of the Virtual Chassis configuration.

at time —(Optional) Time at which to stop the software, specified in one of the following ways:

- **now**—Stop the software immediately. This is the default.
- **+minutes**—Number of minutes from now to stop the software.
- **yymmddhhmm**—Absolute time at which to stop the software, specified as year, month, day, hour, and minute.
- **hh:mm**—Absolute time on the current day at which to stop the software.

backup-routing-engine—(Optional) Halt the backup Routing Engine. This command halts the backup Routing Engine, regardless from which Routing Engine the command is executed. For example, if you issue the command from the master Routing Engine, the backup Routing Engine is halted. If you issue the command from the backup Routing Engine, the backup Routing Engine is halted.

both-routing-engines—(Optional) Halt both Routing Engines at the same time.

director-device *director-device-id*—(QFabric systems only) Halt a specific Director device.

lcc number—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, halt a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, halt a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(EX4200 switches and MX Series routers only) (Optional) Halt the local Virtual Chassis member.

in *minutes*—(Optional) Number of minutes from now to stop the software. This option is an alias for the at *+minutes* option.

media (compact-flash | disk)—(Optional) Boot medium for the next boot.

media (external | internal)—(EX Series and QFX Series switches and MX Series routers only) (Optional) Halt the boot media:

- **external**—Halt the external mass storage device.
- **internal**—Halt the internal flash device.

member *member-id*—(EX4200 switches and MX Series routers only) (Optional) Halt the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

message "*text*"—(Optional) Message to display to all system users before stopping the software.

other-routing-engine—(Optional) Halt the other Routing Engine from which the command is issued. For example, if you issue the command from the master Routing Engine, the backup Routing Engine is halted. Similarly, if you issue the command from the backup Routing Engine, the master Routing Engine is halted.

scc—(TX Matrix routers only) (Optional) Halt the TX Matrix router (or switch-card chassis).

sfc *number*—(TX Matrix Plus routers only) (Optional) Halt the TX Matrix Plus router (or switch-fabric chassis). Replace *number* with 0.

slice *slice*—(EX Series and QFX Series switches only) (Optional) Halt a partition on the boot media. This option has the following suboptions:

- 1—Halt partition 1.
- 2—Halt partition 2.
- **alternate**—Reboot from the alternate partition.

Additional Information On the M7i router, the **request system halt** command does not immediately power down the Packet Forwarding Engine. The power-down process can take as long as 5 minutes.

On a TX Matrix router and TX Matrix Plus router if you issue the **request system halt** command on the master Routing Engine, all the master Routing Engines connected to the routing matrix are halted. If you issue this command on the backup Routing Engine, all the backup Routing Engines connected to the routing matrix are halted.



NOTE: If you have a router or switch with two Routing Engines and you want to shut the power off to the router or switch or remove a Routing Engine, you must first halt the backup Routing Engine (if it has been upgraded), and then halt the master Routing Engine. To halt a Routing Engine, issue the **request system halt** command. You can also halt both Routing Engines at the same time by issuing the **request system halt both-routing-engines** command.

Required Privilege Level maintenance

Related Documentation

- [clear system reboot on page 337](#)
- [request system power-off on page 391](#)
- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

List of Sample Output

- [request system halt on page 389](#)
- [request system halt \(In 2 Hours\) on page 389](#)
- [request system halt \(Immediately\) on page 389](#)
- [request system halt \(At 1:20 AM\) on page 389](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system halt

```
user@host> request system halt
Halt the system ? [yes,no] (no) yes

*** FINAL System shutdown message from root@section2 ***
System going down IMMEDIATELY
Terminated
...
syncing disks... 11 8 done
The operating system has halted.
Please press any key to reboot.
```

request system halt (In 2 Hours)

The following example, which assumes that the time is 5 PM (1700), illustrates three different ways to request that the system stop 2 hours from now:

```
user@host> request system halt at +120
user@host> request system halt in 120
user@host> request system halt at 19:00
```

request system halt (Immediately)

```
user@host> request system halt at now
```

request system halt (At 1:20 AM)

To stop the system at 1:20 AM, enter the following command. Because 1:20 AM is the next day, you must specify the absolute time.

```
user@host> request system halt at yymdd120
request system halt at 120
Halt the system at 120? [yes,no] (no) yes
```

request system logout

Syntax	<code>request system logout (pid <i>pid</i> terminal <i>terminal</i> user <i>username</i>) <all></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Log out users from the router or switch and the configuration database. If a user held the configure exclusive lock, this command clears the exclusive lock.
Options	all —(Optional) Log out all sessions owned by a particular PID, terminal session, or user. (On a TX Matrix or TX Matrix Plus router, this command is broadcast to all chassis.) pid <i>pid</i> —Log out the user session using the specified management process identifier (PID). The PID type must be management process. terminal <i>terminal</i> —Log out the user for the specified terminal session. user <i>username</i> —Log out the specified user.
Required Privilege Level	configure
Related Documentation	<ul style="list-style-type: none">• <i>Log a User Out of the Router</i>
List of Sample Output	request system logout on page 390
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system logout

```
user@host> request system logout user tammy all
Connection closed by foreign host.
```

request system power-off

List of Syntax	Syntax on page 391 Syntax (EX Series Switches) on page 391 Syntax (TX Matrix Router) on page 391 Syntax (TX Matrix Plus Router) on page 391 Syntax (MX Series Router) on page 391 Syntax (QFX Series) on page 392
Syntax	<pre>request system power-off <both-routing-engines> <other-routing-engine> <at <i>time</i>> <in <i>minutes</i>> <media (compact-flash disk removable-compact-flash usb)> <message "<i>text</i>"></pre>
Syntax (EX Series Switches)	<pre>request system power-off <all-members> <at <i>time</i>> <both-routing-engines> <in <i>minutes</i>> <local> <media (external internal)> <member <i>member-id</i>> <message "<i>text</i>"> <other-routing-engine> <slice <i>slice</i>></pre>
Syntax (TX Matrix Router)	<pre>request system power-off <all-chassis all-lcc lcc <i>number</i> scc> <both-routing-engines> <other-routing-engine> <at <i>time</i>> <in <i>minutes</i>> <media (compact-flash disk)> <message "<i>text</i>"></pre>
Syntax (TX Matrix Plus Router)	<pre>request system power-off <all-chassis all-lcc lcc <i>number</i> sfc <i>number</i>> <both-routing-engines> <other-routing-engine> <at <i>time</i>> <in <i>minutes</i>> <media (compact-flash disk)> <message "<i>text</i>"></pre>
Syntax (MX Series Router)	<pre>request system power-off <all-members> <at <i>time</i>> <both-routing-engines> <in <i>minutes</i>> <local></pre>

```
<media (external | internal)>
<member member-id>
<message "text">
<other-routing-engine>
```

Syntax (QFX Series) `request system power-off`
`<at time>`
`<in minutes>`
`<media (external | internal)>`
`<message "text">`
`<slice slice>`

Release Information Command introduced in Junos OS Release 8.0.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Power off the software.



NOTE: When you issue this command on an individual component in a QFabric system, you will receive a warning that says “Hardware-based members will halt, Virtual Junos Routing Engines will reboot.” If you want to halt only one member, use the `member` option. You cannot issue this command from the QFabric CLI.



NOTE: For a standalone chassis (such as MX Series, PTX Series, and T Series routers), the request to power off the system is applicable only to the Routing Engines. When you request to power off both Routing Engines, all the FPCs in the chassis shut down after approximately 10 minutes and the chassis fans run at full speed. The FPCs shut down because they no longer have communication with the Routing Engines and an Inter-Integrated Circuit (I2C) timeout occurred.

Options `none`—Power off the router or switch software immediately.

all-chassis—(Optional) (TX Matrix and TX Matrix Plus router only) Power off all Routing Engines in the chassis.

all-lcc—(Optional) (TX Matrix and TX Matrix Plus router only) On a TX Matrix router, power off all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, power off all T1600 routers (or line-card chassis) connected to the TX Matrix Plus router.

all-members—(EX4200 switches and MX Series routers only) (Optional) Power off all members of the Virtual Chassis configuration.

at *time*—(Optional) Time at which to power off the software, specified in one of the following ways:

- **now**—Power off the software immediately. This is the default.
- **+*minutes***—Number of minutes from now to power off the software.
- ***yymmddhhmm***—Absolute time at which to power off the software, specified as year, month, day, hour, and minute.
- ***hh:mm***—Absolute time on the current day at which to power off the software.

both-routing-engines—(Optional) Power off both Routing Engines at the same time.

in *minutes*—(Optional) Number of minutes from now to power off the software. This option is an alias for the **at +*minutes*** option.

lcc *number*—(Optional) (TX Matrix and TX Matrix Plus router only) On a TX Matrix router, power off a T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, power off a specific router that is connected to the TX Matrix Plus router. Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(EX4200 switches and MX Series routers only) (Optional) Power off the local Virtual Chassis member.

media (compact-flash | disk)—(Optional) Boot medium for the next boot.

media (external | internal)—(EX Series and QFX Series switches and MX Series routers only) (Optional) Power off the boot media:

- **external**—Power off the external mass storage device.
- **internal**—Power off the internal flash device.

member *member-id*—(EX4200 switches and MX Series routers only) (Optional) Power off the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

message "*text*"—(Optional) Message to display to all system users before powering off the software.

other-routing-engine—(Optional) Power off the other Routing Engine from which the command is issued. For example, if you issue the command from the master Routing Engine, the backup Routing Engine is halted. Similarly, if you issue the command from the backup Routing Engine, the master Routing Engine is halted.

scc—(Optional) (TX Matrix router only) Power off only the master Routing Engine or the backup Routing Engine on the TX Matrix router (or switch-card chassis). If you issue the command from the master Routing Engine, the master SCC is powered off. If you issue the command from the backup Routing Engine, the backup SCC is powered off.

sfc number—(Optional) (TX Matrix Plus router only) Power off only the master Routing Engine or the backup Routing Engine on the TX Matrix Plus router (or switch-fabric chassis). If you issue the command from the master Routing Engine, the master SFC is powered off. If you issue the command from the backup Routing Engine, the backup SFC is powered off. Replace *number* with zero.

slice slice—(EX Series and QFX Series switches only) (Optional) Power off a partition on the boot media. This option has the following suboptions:

- **1**—Power off partition 1.
- **2**—Power off partition 2.
- **alternate**—Reboot from the alternate partition.

Additional Information On a routing matrix composed of a TX Matrix router and T640 routers, if you issue the **request system power-off** command on the TX Matrix master Routing Engine, all the master Routing Engines connected to the routing matrix are powered off. If you issue this command on the backup Routing Engine, all the backup Routing Engines connected to the routing matrix are powered off.

Likewise, on a routing matrix composed of a TX Matrix Plus router and T1600 routers, if you issue the **request system power-off** command on the TX Matrix Plus master Routing Engine, all the master Routing Engines connected to the routing matrix are powered off. If you issue this command on the backup Routing Engine, all the backup Routing Engines connected to the routing matrix are powered off.

If you issue the **request system power-off both-routing-engines** command on the TX Matrix or TX Matrix Plus router, all the Routing Engines on the routing matrix are powered off.

Required Privilege Level maintenance

List of Sample Output [request system power-off on page 395](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system power-off

```
user@host> request system power-off message "This router will be powered off in 30 minutes.  
Please save your data and log out immediately."  
warning: This command will not halt the other routing-engine.  
If planning to switch off power, use the both-routing-engines option.  
Power Off the system ? [yes,no] (no) yes
```

```
*** FINAL System shutdown message from remote@nutmeg ***  
System going down IMMEDIATELY
```

```
This router will be powered off in 30 minutes. Please save your data and log out  
immediately.
```

```
Shutdown NOW!  
[pid 5177]
```

request system reboot

Syntax (QFX Series) request system reboot
 <all <graceful>>
 <at time>
 <director-device *name*>
 <director-group <graceful>>
 <fabric <graceful>>
 <fast-boot>
 <in minutes>
 <hypervisor>
 <media >
 <message "text">
 <node-group *name*>
 <slice (1 | 2 | alternate)>

Release Information Command introduced in Junos OS Release 11.1 for the QFX Series.

Description Reboot the Junos OS.



NOTE: On a QFabric system, to avoid traffic loss on the network Node group, switch mastership of the Routing Engine to the backup Routing Engine, and then reboot.

Reboot requests are recorded in the system log files, which you can view with the **show log messages** command. You can view the process names with the **show system processes** command.

Options **none**—Reboots the software immediately.

all—(QFabric systems only) (Optional) Reboots the software on the Director group, fabric control Routing Engines, fabric manager Routing Engines, Interconnect devices, and network and server Node groups.

at time—(Optional) Time at which to reboot the software, specified in one of the following ways:

- **+minutes**—Number of minutes from now to reboot the software.
- **hh:mm**—Absolute time on the current day at which to reboot the software, specified in 24-hour time.
- **now**—Stop or reboot the software immediately. This is the default.
- **yymmddhhmm**—Absolute time at which to reboot the software, specified as year, month, day, hour, and minute.

director-device *name*—(QFabric systems only) (Optional) Reboots the software on the Director device and the default partition (QFabric CLI).

director-group—(QFabric systems only) (Optional) Reboots the software on the Director group and the default partition (QFabric CLI).

fabric—(QFabric systems only) (Optional) Reboots the fabric control Routing Engines and the Interconnect devices.

fast-boot—(QFX5100 switches only) (Optional) Enable fast reboot.

graceful—(QFabric systems only) (Optional) Allows the QFabric component to reboot with minimal impact to network traffic. This option is only available for the **all**, **fabric**, and **director-group** options.

in minutes—(Optional) Number of minutes from now to reboot the software. This option is an alias for the **at +minutes** option.

hypervisor—(Optional) Reboot Junos OS, host OS, and any installed guest VMs.

media (external | internal)—(Optional) Boot medium for the next boot. The external option reboots the switch using a software package stored on an external boot source, such as a USB flash drive. The internal option reboots the switch using a software package stored in an internal memory source.

message "text"—(Optional) Message to display to all system users before rebooting the software.

node-group name—(QFabric systems only) (Optional) Reboots the software on a server Node group or a network Node group.

routing-engine—(Optional) Reboot the Routing Engine.

slice (1 | 2 | alternate)—(Optional) Reboot using the specified partition on the boot media. This option has the following suboptions:



NOTE: The slice option is not supported on the QFX5100 switch, because there is no alternate slice when Junos OS boots as a Virtual Machine (VM). To switch to previous version of Junos OS, issue the **request system software rollback** command.

- **1**—Reboot from partition 1.
- **2**—Reboot from partition 2.
- **alternate**—Reboot from the alternate partition, which is the partition that did not boot the switch at the last bootup.

Required Privilege Level maintenance

Related Documentation

- [clear system reboot on page 337](#)
- [Rebooting and Halting a Device on page 25](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system reboot

```
user@switch> request system reboot
Reboot the system ? [yes,no] (no)
```

request system reboot (At 2300)

```
user@switch> request system reboot at 2300 message ?Maintenance time!?
Reboot the system ? [yes,no] (no) yes
```

```
shutdown: [pid 186]
*** System shutdown message from root@berry.network.net ***
System going down at 23:00
```

request system reboot (In 2 Hours)

The following example, which assumes that the time is 5 PM (1700), illustrates three different ways to request the system to reboot in 2 hours:

```
user@switch> request system reboot at +120
user@switch> request system reboot in 120
user@switch> request system reboot at 19:00
```

request system reboot (Immediately)

```
user@switch> request system reboot at now
```

request system reboot (At 1:20 AM)

To reboot the system at 1:20 AM, enter the following command. Because 1:20 AM is the next day, you must specify the absolute time.

```
user@switch> request system reboot at 06060120
request system reboot at 120
Reboot the system at 120? [yes,no] (no) yes
```

request system reboot director-device

```
user@switch> request system reboot director-device Node1
Issuing this command may interrupt traffic forwarding.
Continue? [yes,no] (no)
```

request system reboot director-group

```
user@switch> request system reboot director-group
Issuing this command may interrupt traffic forwarding.
Continue? [yes,no] (no)
```

request system reboot director-group graceful

```
user@switch> request system reboot director-group graceful
Issuing this command may interrupt traffic forwarding.
Continue? [yes,no] (no)
```

request system storage cleanup

List of Syntax	Syntax on page 399 Syntax (EX Series Switches) on page 399 Syntax (MX Series Router) on page 399 Syntax (QFX Series) on page 399 Syntax (SRX Series) on page 399
Syntax	request system storage cleanup <dry-run>
Syntax (EX Series Switches)	request system storage cleanup <all-members> <dry-run> <local> <member <i>member-id</i> > <satellite [slot-id <i>slot-id</i> device-alias <i>alias-name</i>]>
Syntax (MX Series Router)	request system storage cleanup <all-members> <dry-run> <local> <member <i>member-id</i> > <satellite [slot-id <i>slot-id</i> device-alias <i>alias-name</i>]>
Syntax (QFX Series)	request system storage cleanup <component (<i>serial number</i> <i>UUID</i> all)> <director-group <i>name</i> > <dry-run> <infrastructure <i>name</i> > <interconnect-device <i>name</i> > <name-tag <i>name-tag</i> > <node-group <i>name</i> > <prune> <qfabric (component <i>name</i>) dry-run name-tag repository> <repository (core log)>
Syntax (SRX Series)	request system storage cleanup <dry-run>
Release Information	Command introduced in Junos OS Release 7.4. dry-run option introduced in Junos OS Release 7.6. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 9.2 for SRX Series. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series. satellite option introduced in Junos OS Release 14.2R3.
Description	Free storage space on the router or switch by rotating log files and proposing a list of files for deletion. User input is required for file deletion. On a QFabric system, you can delete debug files located on individual devices or on the entire QFabric system.

Options **all-members**—(EX4200 switches and MX Series routers only) (Optional) Delete files on the Virtual Chassis master Routing Engine only.



NOTE: To delete files on the other members of the Virtual Chassis configuration, log in to each backup Routing Engine and delete the files using the **request system storage cleanup local** command.

component (*UUID | serial number | all*)—(QFabric systems only) (Optional) Delete files located on individual QFabric system devices or on the entire QFabric system.

director-group *name*—(QFabric systems only) (Optional) Delete files on the Director group.

dry-run—(Optional) List files proposed for deletion (without deleting them).

infrastructure *name*—(QFabric systems only) (Optional) Delete files on the fabric control Routing Engine and fabric manager Routing Engine.

interconnect-device *name*—(QFabric systems only) (Optional) Delete files on the Interconnect device.

local—(EX4200 switches and MX Series routers only) (Optional) Delete files on the local Virtual Chassis member.

member *member-id*—(EX4200 switches and MX Series routers only) (Optional) Delete files on the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

name-tag *name-tag*—(QFabric systems only) (Optional) Delete debug files that match a specific regular expression.

node-group *name*—(QFabric systems only) (Optional) Delete files on the Node group.

prune—(QFabric systems only) (Optional) Delete debug files located in either the core or log debug repositories of a QFabric system device.

qfabric component *name*—(QFabric systems only) (Optional) Delete debug files located in the debug repositories of a QFabric system device.

repository (*core | log*)—(QFabric systems only) (Optional) Specify the repository on the QFabric system device for which you want to delete debug files.

satellite [*slot-id slot-id | device-alias alias-name*]—(Junos Fusion only) (Optional) Specify the satellite device in the Junos Fusion by FPC ID or device alias name for which you want to delete debug files.

Additional Information If logging is configured and being used, the **dry-run** option rotates the log files. In that case, the output displays the message “Currently rotating log files, please wait.” If no logging is currently under way, the output displays only a list of files to delete.

Required Privilege Level maintenance

List of Sample Output [request system storage cleanup dry-run on page 402](#)
[request system storage cleanup on page 402](#)
[request system storage cleanup director-group \(QFabric Systems\) on page 402](#)
[request system storage cleanup infrastructure device-name \(QFabric Systems\) on page 404](#)
[request system storage cleanup interconnect-device device-name \(QFabric Systems\) on page 405](#)
[request system storage cleanup node-group group-name \(QFabric Systems\) on page 406](#)
[request system storage cleanup qfabric component device-name \(QFabric Systems\) on page 407](#)
[request system storage cleanup qfabric component device-name repository core \(QFabric Systems\) on page 407](#)
[request system storage cleanup qfabric component all \(QFabric Systems\) on page 408](#)

Output Fields [Table 35](#) describes the output fields for the **request system storage cleanup** command. Output fields are listed in the approximate order in which they appear.

Table 35: request system storage cleanup Output Fields

Field Name	Field Description
List of files to delete:	Shows list of files available for deletion.
Size	Size of the core-dump file.
Date	Last core-dump file modification date and time.
Name	Name of the core-dump file.
Directory to delete:	Shows list of directories available for deletion.
Repository scope:	Repository where core-dump files and log files are stored. The core-dump files are located in the core repository, and the log files are located in the log repository. The default Repository scope is shared since both the core and log repositories are shared by all of the QFabric system devices.
Repository head:	Name of the top-level repository location.
Repository name:	Name of the repository: core or log .
Creating list of debug artifacts to be removed under:	Shows location of files available for deletion.
List of debug artifacts to be removed under:	Shows list of files available for deletion.

Sample Output

request system storage cleanup dry-run

```
user@host> request system storage cleanup dry-run
Currently rotating log files, please wait.
This operation can take up to a minute.
```

List of files to delete:

	Size	Date	Name
	11.4K	Mar 8 15:00	/var/log/messages.1.gz
	7245B	Feb 5 15:00	/var/log/messages.3.gz
	11.8K	Feb 22 13:00	/var/log/messages.2.gz
	3926B	Mar 16 13:57	/var/log/messages.0.gz
	3962B	Feb 22 12:47	/var/log/sampled.1.gz
	4146B	Mar 8 12:20	/var/log/sampled.0.gz
	4708B	Dec 21 11:39	/var/log/sampled.2.gz
	7068B	Jan 16 18:00	/var/log/messages.4.gz
	13.7K	Dec 27 22:00	/var/log/messages.5.gz
	890B	Feb 22 17:22	/var/tmp/sampled.pkts
	65.8M	Oct 26 09:10	/var/sw/pkg/jinstall-7.4R1.7-export-signed.tgz
	63.1M	Oct 26 09:13	/var/sw/pkg/jbundle-7.4R1.7.tgz

request system storage cleanup

```
user@host> request system storage cleanup
Currently rotating log files, please wait.
This operation can take up to a minute.
```

List of files to delete:

	Size	Date	Name
	11.4K	Mar 8 15:00	/var/log/messages.1.gz
	7245B	Feb 5 15:00	/var/log/messages.3.gz
	11.8K	Feb 22 13:00	/var/log/messages.2.gz
	3926B	Mar 16 13:57	/var/log/messages.0.gz
	11.6K	Mar 8 15:00	/var/log/messages.5.gz
	7254B	Feb 5 15:00	/var/log/messages.6.gz
	12.9K	Feb 22 13:00	/var/log/messages.8.gz
	3726B	Mar 16 13:57	/var/log/messages.7.gz
	3962B	Feb 22 12:47	/var/log/sampled.1.gz
	4146B	Mar 8 12:20	/var/log/sampled.0.gz
	4708B	Dec 21 11:39	/var/log/sampled.2.gz
	7068B	Jan 16 18:00	/var/log/messages.4.gz
	13.7K	Dec 27 22:00	/var/log/messages.5.gz
	890B	Feb 22 17:22	/var/tmp/sampled.pkts
	65.8M	Oct 26 09:10	/var/sw/pkg/jinstall-7.4R1.7-export-signed.tgz
	63.1M	Oct 26 09:13	/var/sw/pkg/jbundle-7.4R1.7.tgz

Delete these files ? [yes,no] (yes)

request system storage cleanup director-group (QFabric Systems)

```
user@switch> request system storage cleanup director-group
List of files to delete:
```

	Size	Date	Name
	4.0K	2011-11-07 05:16:29	/tmp/2064.sfcauth
	4.0K	2011-11-07 05:07:34	/tmp/30804.sfcauth
	4.0K	2011-11-07 04:13:41	/tmp/26792.sfcauth

```

4.0K  2011-11-07 04:13:39 /tmp/26432.sfcauth
0      2011-11-07 07:45:40 /tmp/cluster_cleanup.log
1.3M  2011-11-07 07:39:11 /tmp/cn_monitor.20111107-052401.log
4.0K  2011-11-07 07:36:29 /tmp/clustat.28019.log
4.0K  2011-11-07 07:36:29 /tmp/clustat_x.28019.log
9.6M  2011-11-07 05:30:24 /tmp/sfc.2.log
4.0K  2011-11-07 05:28:11 /tmp/mgd-init.1320672491.log
248K  2011-11-07 05:19:24 /tmp/cn_monitor.20111107-045111.log
4.0K  2011-11-07 05:17:18 /tmp/clustat.3401.log
4.0K  2011-11-07 05:17:18 /tmp/clustat_x.3401.log
8.0K  2011-11-07 04:58:25 /tmp/mgd-init.1320670633.log
0      2011-11-07 04:54:01 /tmp/mysql_db_install_5.1.37.log
4.0K  2011-11-07 04:52:08 /tmp/cn_send.log
0      2011-11-07 04:52:00 /tmp/init_eth0.log
4.0K  2011-11-07 04:49:35 /tmp/install_interfaces.sh.log
4.0K  2011-11-07 04:48:15 /tmp/bootstrap.sh.log
160K  2011-11-07 04:47:43 /tmp/bootstrap_cleanup.log
38M   2011-11-07 04:42:42 /tmp/cn_monitor.20111104-110308.log
4.0K  2011-11-07 04:38:47 /tmp/clustat.30913.log
4.0K  2011-11-07 04:38:47 /tmp/clustat_x.30913.log
4.0K  2011-11-07 04:38:03 /tmp/dcf_upgrade.sh.remove.log
4.0K  2011-11-07 04:38:03 /tmp/peer_update.log
4.0K  2011-11-07 04:38:02 /tmp/dcf_upgrade.log
4.0K  2011-11-07 04:38:02 /tmp/perl_mark_upgrade.log
8.0K  2011-11-07 04:13:42 /tmp/install_dcf_rpm.log
4.0K  2011-11-07 04:13:06 /tmp/00_cleanup.sh.1320667986.log
0      2011-11-07 04:13:06 /tmp/ccif_patch_4410_4450.sh.1320667986.log
4.0K  2011-11-07 04:13:06 /tmp/dcf-tools.sh.1320667986.log
0      2011-11-07 04:13:06 /tmp/initial.sh.1320667986.log
0      2011-11-07 04:13:06 /tmp/inventory.sh.1320667986.log
4.0K  2011-11-07 04:13:06 /tmp/qf-db.sh.1320667986.log
4.0K  2011-11-07 04:13:06 /tmp/sfc.sh.1320667986.log
8.0K  2011-11-07 04:13:05 /tmp/jinstall-qfabric.log
8.0K  2011-11-04 11:10:24 /tmp/mgd-init.1320430192.log
4.0K  2011-11-04 11:07:03 /tmp/mysql_dcf_db_install.log
8.0K  2011-11-04 10:55:07 /tmp/ccif_patch_4410_4450.sh.1320429307.log
8.0K  2011-11-04 10:55:07 /tmp/initial.sh.1320429307.log
4.0K  2011-11-04 10:55:07 /tmp/inventory.sh.1320429307.log
8.0K  2011-11-04 10:55:07 /tmp/sfc.sh.1320429307.log
4.0K  2011-11-04 10:54:09 /tmp/ks-script-Ax0tz5.log
4.0K  2011-11-07 04:13:06 /tmp//sfc.sh.1320667986.log
8.0K  2011-11-04 10:55:07 /tmp//sfc.sh.1320429307.log

```

Directory to delete:

```

45M   2011-11-08 10:57:43 /tmp/sfc-captures

```

List of files to delete:

	Size	Date	Name
4.0K	2011-11-08	05:47:47	/tmp/5713.sfcauth
4.0K	2011-11-08	05:14:32	/tmp/14494.sfcauth
4.0K	2011-11-08	05:11:47	/tmp/9978.sfcauth
4.0K	2011-11-08	05:09:37	/tmp/6128.sfcauth
4.0K	2011-11-08	05:04:28	/tmp/29703.sfcauth
4.0K	2011-11-07	11:59:10	/tmp/7811.sfcauth
4.0K	2011-11-07	11:36:08	/tmp/32415.sfcauth
4.0K	2011-11-07	11:30:30	/tmp/22406.sfcauth
4.0K	2011-11-07	11:24:37	/tmp/12131.sfcauth
4.0K	2011-11-07	10:48:42	/tmp/12687.sfcauth
4.0K	2011-11-07	09:27:20	/tmp/31082.sfcauth
4.0K	2011-11-07	07:33:58	/tmp/14633.sfcauth

```

4.0K  2011-11-07 05:08:25 /tmp/15447.sfcauth
4.0K  2011-11-07 04:12:29 /tmp/26874.sfcauth
4.0K  2011-11-07 04:12:27 /tmp/26713.sfcauth
4.0K  2011-11-07 03:49:17 /tmp/17691.sfcauth
4.0K  2011-11-05 01:32:23 /tmp/5716.sfcauth
4.0K  2011-11-07 08:00:17 /tmp/sfcsnmpd.log
4.0K  2011-11-07 07:57:50 /tmp/cluster_cleanup.log
824K  2011-11-07 07:38:37 /tmp/cn_monitor.20111107-053643.log
4.0K  2011-11-07 07:36:30 /tmp/clustat.18399.log
4.0K  2011-11-07 07:36:30 /tmp/clustat_x.18399.log
4.0K  2011-11-07 07:35:47 /tmp/command_lock.log
4.0K  2011-11-07 05:39:54 /tmp/mgd-init.1320673194.log
92K   2011-11-07 05:19:25 /tmp/cn_monitor.20111107-050412.log
4.0K  2011-11-07 05:17:20 /tmp/clustat.30115.log
4.0K  2011-11-07 05:17:20 /tmp/clustat_x.30115.log
8.0K  2011-11-07 05:08:07 /tmp/mgd-init.1320671241.log
4.0K  2011-11-07 05:04:57 /tmp/cn_send.log
0     2011-11-07 05:04:52 /tmp/init_eth0.log
4.0K  2011-11-07 05:02:38 /tmp/install_interfaces.sh.log
4.0K  2011-11-07 05:01:19 /tmp/bootstrap.sh.log
160K  2011-11-07 05:00:47 /tmp/bootstrap_cleanup.log
28M   2011-11-07 04:42:27 /tmp/cn_monitor.20111104-112954.log
4.0K  2011-11-07 04:38:49 /tmp/clustat.6780.log
4.0K  2011-11-07 04:38:49 /tmp/clustat_x.6780.log
4.0K  2011-11-07 04:38:05 /tmp/issue_event.log
4.0K  2011-11-07 04:38:05 /tmp/peer_upgrade_reboot.log
12K   2011-11-07 04:38:05 /tmp/primary_update.log
4.0K  2011-11-07 04:38:04 /tmp/dcf_upgrade.sh.remove.log
4.0K  2011-11-07 04:38:04 /tmp/peer_rexec_upgrade.log
4.0K  2011-11-07 04:13:42 /tmp/peer_install_dcf_rpm.log
4.0K  2011-11-07 04:11:57 /tmp/dcf-tools.sh.1320667917.log
0     2011-11-07 04:11:57 /tmp/initial.sh.1320667917.log
0     2011-11-07 04:11:57 /tmp/inventory.sh.1320667917.log
4.0K  2011-11-07 04:11:57 /tmp/qf-db.sh.1320667917.log
4.0K  2011-11-07 04:11:57 /tmp/sfc.sh.1320667917.log
4.0K  2011-11-07 04:11:56 /tmp/00_cleanup.sh.1320667916.log
0     2011-11-07 04:11:56 /tmp/ccif_patch_4410_4450.sh.1320667916.log
8.0K  2011-11-07 04:11:56 /tmp/jinstall-qfabric.log
4.0K  2011-11-07 04:11:33 /tmp/dcf_upgrade.log
8.0K  2011-11-04 11:53:12 /tmp/mgd-init.1320432782.log
8.0K  2011-11-04 11:06:17 /tmp/ccif_patch_4410_4450.sh.1320429977.log
8.0K  2011-11-04 11:06:17 /tmp/initial.sh.1320429977.log
4.0K  2011-11-04 11:06:17 /tmp/inventory.sh.1320429977.log
8.0K  2011-11-04 11:06:17 /tmp/sfc.sh.1320429977.log
4.0K  2011-11-04 11:05:19 /tmp/ks-script-_tnWeb.log
4.0K  2011-11-07 04:11:57 /tmp//sfc.sh.1320667917.log
8.0K  2011-11-04 11:06:17 /tmp//sfc.sh.1320429977.log

```

Directory to delete:

```
49M   2011-11-08 10:45:20 /tmp/sfc-captures
```

request system storage cleanup infrastructure device-name (QFabric Systems)

```
user@switch> request system storage cleanup infrastructure FC-0
re0:
```

```
-----
```

List of files to delete:

Size	Date	Name
139B	Nov 8 19:03	/var/log/default-log-messages.0.gz

```

5602B Nov  8 19:03 /var/log/messages.0.gz
28.4K Nov  8 10:15 /var/log/messages.1.gz
35.2K Nov  7 13:45 /var/log/messages.2.gz
207B Nov  7 16:02 /var/log/wtmp.0.gz
27B Nov  7 12:14 /var/log/wtmp.1.gz
184.4M Nov  7 12:16
/var/sw/pkg/jinstall-dc-re-11.3I20111104_1216_dc-builder-domestic-signed.tgz
124.0K Nov  7 15:59 /var/tmp/gres-tp/env.dat
0B Nov  7 12:57 /var/tmp/gres-tp/lock
155B Nov  7 16:02 /var/tmp/krt_gencfg_filter.txt
0B Nov  7 12:35 /var/tmp/last_ccif_update
1217B Nov  7 12:15 /var/tmp/loader.conf.preinstall
184.4M Nov  6 07:11 /var/tmp/mchassis-install.tgz
10.8M Nov  7 12:16
/var/tmp/preinstall/bootstrap-install-11.3I20111104_1216_dc-builder.tar
57.4K Nov  7 12:16 /var/tmp/preinstall/configs-11.3I20111104_1216_dc-builder.tgz

259B Nov  7 12:16 /var/tmp/preinstall/install.conf
734.3K Nov  4 13:46
/var/tmp/preinstall/jboot-dc-re-11.3I20111104_1216_dc-builder.tgz
177.8M Nov  7 12:16
/var/tmp/preinstall/jbundle-dc-re-11.3I20111104_1216_dc-builder-domestic.tgz
124B Nov  7 12:15 /var/tmp/preinstall/metatags
1217B Nov  7 12:16 /var/tmp/preinstall_boot_loader.conf
0B Nov  7 16:02 /var/tmp/rtsdb/if-rtsdb

```

request system storage cleanup interconnect-device device-name (QFabric Systems)

```

user@switch> request system storage cleanup interconnect IC-WS001
re1:

```

List of files to delete:

Size	Date	Name
11B	Nov 7 15:55	/var/jail/tmp/alarmd.ts
128B	Nov 8 19:06	/var/log/default-log-messages.0.gz
9965B	Nov 8 19:06	/var/log/messages.0.gz
15.8K	Nov 8 12:30	/var/log/messages.1.gz
15.8K	Nov 8 11:00	/var/log/messages.2.gz
15.7K	Nov 8 07:30	/var/log/messages.3.gz
15.8K	Nov 8 04:00	/var/log/messages.4.gz
15.7K	Nov 8 00:30	/var/log/messages.5.gz
18.7K	Nov 7 21:00	/var/log/messages.6.gz
17.6K	Nov 7 19:00	/var/log/messages.7.gz
58.3K	Nov 7 16:00	/var/log/messages.8.gz
20.3K	Nov 7 15:15	/var/log/messages.9.gz
90B	Nov 7 15:41	/var/log/wtmp.0.gz
57B	Nov 7 12:41	/var/log/wtmp.1.gz
124.0K	Nov 7 15:42	/var/tmp/gres-tp/env.dat
0B	Nov 7 12:40	/var/tmp/gres-tp/lock
0B	Nov 7 12:41	/var/tmp/if-rtsdb/env.lck
12.0K	Nov 7 15:41	/var/tmp/if-rtsdb/env.mem
132.0K	Nov 7 15:55	/var/tmp/if-rtsdb/shm_usr1.mem
2688.0K	Nov 7 15:41	/var/tmp/if-rtsdb/shm_usr2.mem
2048.0K	Nov 7 15:41	/var/tmp/if-rtsdb/trace.mem
730B	Nov 7 19:57	/var/tmp/juniper.conf+.gz
155B	Nov 7 15:53	/var/tmp/krt_gencfg_filter.txt
0B	Nov 7 15:41	/var/tmp/rtsdb/if-rtsdb

re0:

List of files to delete:

	Size	Date	Name
	11B	Nov 7 15:55	/var/jail/tmp/alarmd.ts
	121B	Nov 8 19:06	/var/log/default-log-messages.0.gz
	16.7K	Nov 8 19:06	/var/log/messages.0.gz
	22.2K	Nov 8 17:45	/var/log/messages.1.gz
	18.4K	Nov 8 17:00	/var/log/messages.2.gz
	21.6K	Nov 8 16:00	/var/log/messages.3.gz
	17.9K	Nov 8 14:30	/var/log/messages.4.gz
	19.4K	Nov 8 13:30	/var/log/messages.5.gz
	18.2K	Nov 8 12:30	/var/log/messages.6.gz
	20.4K	Nov 8 11:30	/var/log/messages.7.gz
	21.4K	Nov 8 10:15	/var/log/messages.8.gz
	21.0K	Nov 8 09:00	/var/log/messages.9.gz
	19.9K	Nov 8 08:13	/var/log/snmp-traps.0.gz
	203B	Nov 8 15:36	/var/log/wtmp.0.gz
	57B	Nov 7 12:41	/var/log/wtmp.1.gz
	124.0K	Nov 7 15:42	/var/tmp/gres-tp/env.dat
	0B	Nov 7 12:40	/var/tmp/gres-tp/lock
	0B	Nov 7 12:41	/var/tmp/if-rtbdb/env.lck
	12.0K	Nov 7 15:41	/var/tmp/if-rtbdb/env.mem
	132.0K	Nov 7 15:55	/var/tmp/if-rtbdb/shm_usr1.mem
	2688.0K	Nov 7 15:41	/var/tmp/if-rtbdb/shm_usr2.mem
	2048.0K	Nov 7 15:41	/var/tmp/if-rtbdb/trace.mem
	727B	Nov 7 15:54	/var/tmp/juniper.conf+.gz
	155B	Nov 7 15:55	/var/tmp/krt_gencfg_filter.txt
	0B	Nov 7 15:41	/var/tmp/rtbdb/if-rtbdb

request system storage cleanup node-group group-name (QFabric Systems)

```
user@switch> request system storage cleanup node-group NW-NG-0
BBAK0372:
```

List of files to delete:

	Size	Date	Name
	126B	Nov 8 19:07	/var/log/default-log-messages.0.gz
	179B	Nov 7 13:32	/var/log/install.0.gz
	22.9K	Nov 8 19:07	/var/log/messages.0.gz
	26.5K	Nov 8 17:30	/var/log/messages.1.gz
	20.5K	Nov 8 13:15	/var/log/messages.2.gz
	33.2K	Nov 7 17:45	/var/log/messages.3.gz
	35.5K	Nov 7 15:45	/var/log/messages.4.gz
	339B	Nov 8 17:10	/var/log/wtmp.0.gz
	58B	Nov 7 12:40	/var/log/wtmp.1.gz
	124.0K	Nov 8 17:08	/var/tmp/gres-tp/env.dat
	0B	Nov 7 12:39	/var/tmp/gres-tp/lock
	0B	Nov 7 12:59	/var/tmp/if-rtbdb/env.lck
	12.0K	Nov 8 17:09	/var/tmp/if-rtbdb/env.mem
	2688.0K	Nov 8 17:09	/var/tmp/if-rtbdb/shm_usr1.mem
	132.0K	Nov 8 17:09	/var/tmp/if-rtbdb/shm_usr2.mem
	2048.0K	Nov 8 17:09	/var/tmp/if-rtbdb/trace.mem
	1082B	Nov 8 17:09	/var/tmp/juniper.conf+.gz
	155B	Nov 7 17:39	/var/tmp/krt_gencfg_filter.txt
	0B	Nov 8 17:09	/var/tmp/rtbdb/if-rtbdb

EE3093:

List of files to delete:

	Size	Date	Name
	11B	Nov 8 17:33	/var/jail/tmp/alarmd.ts
	119B	Nov 8 19:08	/var/log/default-log-messages.0.gz
	180B	Nov 7 17:41	/var/log/install.0.gz
	178B	Nov 7 13:32	/var/log/install.1.gz
	2739B	Nov 8 19:08	/var/log/messages.0.gz
	29.8K	Nov 8 18:45	/var/log/messages.1.gz
	31.8K	Nov 8 17:15	/var/log/messages.2.gz
	20.6K	Nov 8 16:00	/var/log/messages.3.gz
	15.4K	Nov 8 10:15	/var/log/messages.4.gz
	15.4K	Nov 8 02:15	/var/log/messages.5.gz
	25.5K	Nov 7 20:45	/var/log/messages.6.gz
	48.0K	Nov 7 17:45	/var/log/messages.7.gz
	32.8K	Nov 7 13:45	/var/log/messages.8.gz
	684B	Nov 8 17:02	/var/log/wtmp.0.gz
	58B	Nov 7 12:40	/var/log/wtmp.1.gz
	124.0K	Nov 7 17:34	/var/tmp/gres-tp/env.dat
	0B	Nov 7 12:40	/var/tmp/gres-tp/lock
	0B	Nov 7 12:59	/var/tmp/if-rtbdb/env.lck
	12.0K	Nov 7 17:39	/var/tmp/if-rtbdb/env.mem
	2688.0K	Nov 7 17:39	/var/tmp/if-rtbdb/shm_usr1.mem
	132.0K	Nov 7 17:40	/var/tmp/if-rtbdb/shm_usr2.mem
	2048.0K	Nov 7 17:39	/var/tmp/if-rtbdb/trace.mem
	155B	Nov 7 17:40	/var/tmp/krt_gencfg_filter.txt
	0B	Nov 7 17:39	/var/tmp/rtbdb/if-rtbdb

request system storage cleanup qfabric component device-name (QFabric Systems)

```

user@switch> request system storage cleanup qfabric component A0001/YA0197
Repository type: regular
Repository head: /pbstorage
Creating list of debug artifacts to be removed under:
/pbstorage/rdumps/A0001/YA0197
Removing debug artifacts ... (press control C to abort)
Removing /pbstorage/rdumps/A0001/YA0197/cosd.core.0.0.05162011123308.gz ... done
Removing /pbstorage/rdumps/A0001/YA0197/cosd.core.1.0.05162011123614.gz ... done
Removing /pbstorage/rdumps/A0001/YA0197/cosd.core.2.0.05162011123920.gz ... done
Removing /pbstorage/rdumps/A0001/YA0197/livecore.05132011163930.gz ... done
Removing /pbstorage/rdumps/A0001/YA0197/tnetd.core.0.1057.05162011124500.gz ...
done
Removing /pbstorage/rdumps/A0001/YA0197/vmcore.05132011120528.gz ... done
Removing /pbstorage/rdumps/A0001/YA0197/vmcore.kz ... done
Creating list of debug artifacts to be removed under: /pbstorage/rlogs/A0001/YA0197
Removing debug artifacts ... (press control C to abort)
Removing /pbstorage/rlogs/A0001/YA0197/kdumpinfo.05132011120528 ... done
Removing /pbstorage/rlogs/A0001/YA0197/kernel.tarball.0.1039.05122011234415.tgz
... done
Removing /pbstorage/rlogs/A0001/YA0197/kernel.tarball.1.1039.05132011175544.tgz
... done
Removing /pbstorage/rlogs/A0001/YA0197/tnetd.tarball.0.1057.05162011175453.tgz
... done

```

request system storage cleanup qfabric component device-name repository core (QFabric Systems)

```

user@switch> request system storage cleanup qfabric component EE3093 repository core
Repository scope: shared
Repository head: /pbdata/export

```

```
Repository name: core
Creating list of debug artifacts to be removed under: /pbdata/export/r dumps/EE3093
NOTE: core repository under /pbdata/export/r dumps/EE3093 empty
```

request system storage cleanup qfabric component all (QFabric Systems)

```
user@switch> request system storage cleanup qfabric component all
Repository scope: shared
Repository head: /pbdata/export
Creating list of debug artifacts to be removed under: /pbdata/export/r dumps
NOTE: core repository under /pbdata/export/r dumps/all empty
Creating list of debug artifacts to be removed under: /pbdata/export/r logs
List of debug artifacts to clean up ... (press control C to abort)
/pbdata/export/r logs/73747cd8-0710-11e1-b6a4-00e081c5297e/install-11072011125819.log
/pbdata/export/r logs/77116f18-0710-11e1-a2a0-00e081c5297e/install-11072011125819.log
/pbdata/export/r logs/BBAK0372/install-11072011121538.log
/pbdata/export/r logs/BBAK0394/install-11072011121532.log
/pbdata/export/r logs/EE3093/install-11072011121536.log
/pbdata/export/r logs/WS001/YN5999/install-11072011121644.log
/pbdata/export/r logs/WS001/YW3803/install-11072011122429.log
/pbdata/export/r logs/cd78871a-0710-11e1-878e-00e081c5297e/install-11072011125932.log
/pbdata/export/r logs/d0afda1e-0710-11e1-a1d0-00e081c5297e/install-11072011125930.log
/pbdata/export/r logs/d0afda1e-0710-11e1-a1d0-00e081c5297e/install-11072011133211.log
/pbdata/export/r logs/d0afda1e-0710-11e1-a1d0-00e081c5297e/install-11072011155302.log
/pbdata/export/r logs/d31ab7a6-0710-11e1-ad1b-00e081c5297e/install-11072011125931.log
/pbdata/export/r logs/d4d0f254-0710-11e1-90c3-00e081c5297e/install-11072011125932.log
```


request system zeroize

Syntax request system zeroize
 <media>
 <local>

Release Information Command introduced before Junos OS Release 9.0.
 Command introduced in Junos OS Release 11.2 for EX Series switches.
 Option **media** added in Junos OS Release 11.4 for EX Series switches.
 Command introduced in Junos OS Release 12.2 for MX Series routers.
 Command introduced in Junos OS Release 12.3 for the QFX Series.
 Option **local** added in Junos OS Release 14.1.
 Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description



NOTE: The **media** option is not available on the QFX Series.

Remove all configuration information on the Routing Engines and reset all key values. If the device has dual Routing Engines, the command is broadcast to all Routing Engines on the device. The command removes all data files, including customized configuration and log files, by unlinking the files from their directories. The command removes all user-created files from the system including all plain-text passwords, secrets, and private keys for SSH, local encryption, local authentication, IPsec, RADIUS, TACACS+, and SNMP.

This command reboots the device and sets it to the factory default configuration. After the reboot, you cannot access the device through the management Ethernet interface. Log in through the console as **root** and start the Junos OS CLI by typing **cli** at the prompt.



NOTE: If you configure the **commit synchronize** statement at the **[edit system]** hierarchy level and issue a **commit** in the master Routing Engine, the master configuration is automatically synchronized with the backup. However, if the backup Routing Engine is down when you issue the **commit**, the Junos OS displays a warning and commits the candidate configuration in the master Routing Engine. When the backup Routing Engine comes up, its configuration will automatically be synchronized with the master. A newly inserted backup Routing Engine automatically synchronizes its configuration with the master Routing Engine configuration.

To completely erase user-created data so that it is unrecoverable, use the **media** option.

Options **media**—(Optional) In addition to removing all configuration and log files, causes memory and the media to be scrubbed, removing all traces of any user-created files. Every storage device attached to the system is scrubbed, including disks, flash drives, removable USBs, and so on. The duration of the scrubbing process is dependent on the size of the media being erased. As a result, the **request system zeroize media**

operation can take considerably more time than the **request system zeroize** operation. However, the critical security parameters are all removed at the beginning of the process.

local—(Optional) Remove all the configuration information and restore all the key values on the active Routing Engine.

Required Privilege Level

maintenance

Related Documentation

- *request system snapshot*
- *Reverting to the Default Factory Configuration for the EX Series Switch*
- *Reverting to the Rescue Configuration for the EX Series Switch*
- [Reverting to the Default Factory Configuration on page 1344](#)
- [Reverting to the Rescue Configuration on page 1346](#)
- [Reverting to the Default Factory Configuration by Using the request system zeroize Command on page 28](#)

List of Sample Output [request system zeroize on page 410](#)
[request system zeroize media on page 411](#)

Sample Output

request system zeroize

```
user@host> request system zeroize
warning: System will be rebooted and may not boot without configuration
Erase all data, including configuration and log files? [yes,no] (no) yes

0 1 1 0 0 0 done

syncing disks... All buffers synced.
Uptime: 5d19h20m26s
recorded reboot as normal shutdown
Rebooting...

U-Boot 1.1.6 (Mar 11 2011 - 04:39:06)

Board: EX4200-24T 2.11
EPLD: Version 6.0 (0x85)
DRAM: Initializing (1024 MB)
FLASH: 8 MB

Firmware Version: --- 01.00.00 ---
USB: scanning bus for devices... 2 USB Device(s) found
      scanning bus for storage devices... 1 Storage Device(s) found

ELF file is 32 bit
Consoles: U-Boot console

FreeBSD/PowerPC U-Boot bootstrap loader, Revision 2.4
(user@juniper.net, Fri Mar 11 03:03:36 UTC 2011)
Memory: 1024MB
bootsequencing is enabled
```

```

bootsuccess is set
new boot device = disk0s1:
Loading /boot/defaults/loader.conf
/kernel data=0x915c84+0xa1260 syms=[0x4+0x7cbd0+0x4+0xb1c19]

Hit [Enter] to boot immediately, or space bar for command prompt.
Booting [/kernel]...
Kernel entry at 0x800000e0 ...
GDB: no debug ports present
KDB: debugger backends: ddb
KDB: current backend: ddb
Copyright (c) 1996-2011, Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1992-2006 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
    The Regents of the University of California. All rights reserved.
JUNOS 11.1R1.8 #0: 2011-03-09 20:14:25 UTC

user@juniper.net:/volume/build/junos/11.1/release/11.1R1.8/obj-powerpc/bsd/kernels/
JUNIPER-EX/kernel
Timecounter "decrementer" frequency 50000000 Hz quality 0
cpu0: Freescale e500v2 core revision 2.2
cpu0: HID0 80004080
...

```

request system zeroize media

```

user@host> request system zeroize media
warning: System will be rebooted and may not boot without configuration
Erase all data, including configuration and log files? [yes,no] (no) yes

warning: ipsec-key-management subsystem not running - not needed by configuration.
warning: zeroizing fpc0

{master:0}
root> Waiting (max 60 seconds) for system process `vnlr' to stop...done
...
Syncing disks, vnodes remaining...2 4 2 4 3 2 1 1 0 0 0 done

syncing disks... All buffers synced.
Uptime: 14m50s
recorded reboot as normal shutdown
Rebooting...

U-Boot 1.1.6 (Apr 21 2011 - 13:58:42)

Board: EX4200-48PX 1.1
EPLD: Version 8.0 (0x82)
DRAM: Initializing (512 MB)
FLASH: 8 MB
NAND: No NAND device found!!!
0 MiB

Firmware Version: --- 01.00.00 ---
USB: scanning bus for devices... 2 USB Device(s) found
      scanning bus for storage devices... 1 Storage Device(s) found

ELF file is 32 bit
Consoles: U-Boot console

```

```

FreeBSD/PowerPC U-Boot bootstrap loader, Revision 2.2
(vtseng@svl-junos-pool27.juniper.net, Fri Feb 26 17:48:51 PST 2010)
Memory: 512MB
Loading /boot/defaults/loader.conf
/kernel data=0x9abfdc+0xb06e4 syms=[0x4+0x83b30+0x4+0xbd7c6]

Hit [Enter] to boot immediately, or space bar for command prompt.
Booting [/kernel] in 1 second... Booting [/kernel]...
Kernel entry at 0x800000e0 ...
GDB: no debug ports present
KDB: debugger backends: ddb
KDB: current backend: ddb
Copyright (c) 1996-2011, Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1992-2006 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
The Regents of the University of California. All rights reserved.
JUNOS 11.4R1.2 #0: 2011-10-27 18:05:39 UTC
user@juniper.net:/volume/build/junos/11.4/release/11.4R1.2/obj-powerpc/
bsd/kernels/JUNIPER-EX/kernel
can't re-use a leaf (all_slot_serialid)!
Timecounter "decrementer" frequency 50000000 Hz quality 0
cpu0: Freescale e500v2 core revision 2.2
cpu0: HID0 80004080<EMCP,TBEN,EN_MAS7_UPDATE>
real memory = 511705088 (488 MB)
avail memory = 500260864 (477 MB)
ETHERNET SOCKET BRIDGE initialising
Initializing EXSERIES platform properties ...
. . .
Automatic reboot in progress...
Media check on da0 on ex platforms
** /dev/da0s2a
FILE SYSTEM CLEAN; SKIPPING CHECKS
clean, 20055 free (31 frags, 2503 blocks, 0.0% fragmentation)
zeroizing /dev/da0s1a ...
. . .
zeroizing /dev/da0s3d ...
. . .
zeroizing /dev/da0s3e ...
. . .
zeroizing /dev/da0s4d ...
. . .
zeroizing /dev/da0s4e ...
. . .

syncing disks... All buffers synced.
Uptime: 3m40s
Rebooting...

U-Boot 1.1.6 (Apr 21 2011 - 13:58:42)

Board: EX4200-48PX 1.1
EPLD: Version 8.0 (0x82)
DRAM: Initializing (512 MB)
FLASH: 8 MB
NAND: No NAND device found!!!
0 MiB

Firmware Version: --- 01.00.00 ---
USB: scanning bus for devices... 2 USB Device(s) found
      scanning bus for storage devices... 1 Storage Device(s) found

```

```

ELF file is 32 bit
Consoles: U-Boot console

FreeBSD/PowerPC U-Boot bootstrap loader, Revision 2.2
(vtseng@svl-junos-pool27.juniper.net, Fri Feb 26 17:48:51 PST 2010)
Memory: 512MB
Loading /boot/defaults/loader.conf
/kernel data=0x9abfdc+0xb06e4 syms=[0x4+0x83b30+0x4+0xbd7c6]

Hit [Enter] to boot immediately, or space bar for command prompt.
Booting [/kernel] in 1 second... Booting [/kernel]...
Kernel entry at 0x800000e0 ...
GDB: no debug ports present
KDB: debugger backends: ddb
KDB: current backend: ddb
Copyright (c) 1996-2011, Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1992-2006 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
The Regents of the University of California. All rights reserved.
JUNOS 11.4R1.2 #0: 2011-10-27 18:05:39 UTC
user@juniper.net:/volume/build/junos/11.4/release/11.4R1.2/obj-powerpc/
bsd/kernels/JUNIPER-EX/kernel
can't re-use a leaf (all_slot_serialid)!
Timecounter "decrementer" frequency 50000000 Hz quality 0
cpu0: Freescale e500v2 core revision 2.2
cpu0: HID0 80004080 <EMCP,TBEN,EN_MAS7_UPDATE>
real memory = 511705088 (488 MB)
avail memory = 500260864 (477 MB)
ETHERNET SOCKET BRIDGE initialising
Initializing EXSERIES platform properties ...
...
Automatic reboot in progress...
Media check on da0 on ex platforms
** /dev/da0s1a
FILE SYSTEM CLEAN; SKIPPING CHECKS
clean, 20064 free (48 frags, 2502 blocks, 0.1% fragmentation)
zeroizing /dev/da0s2a ...
...
Creating initial configuration...mgd: error: Cannot open configuration file:
/config/juniper.conf
mgd: warning: activating factory configuration
mgd: commit complete
mgd: -----
mgd: Please login as 'root'. No password is required.
mgd: To start Initial Setup, type 'ezsetup' at the JUNOS prompt.
mgd: To start JUNOS CLI, type 'cli' at the JUNOS prompt.
mgd: -----
Setting initial options: debugger_on_panic=NO debugger_on_break=NO.
Starting optional daemons: .
Doing initial network setup:
...

Amnesiac (ttyu0)

```

restart

List of Syntax [Syntax on page 414](#)

[Syntax \(ACX Series Routers\) on page 414](#)
[Syntax \(EX Series Switches\) on page 414](#)
[Syntax \(MX Series Routers\) on page 415](#)
[Syntax \(QFX Series\) on page 415](#)
[Syntax \(Routing Matrix\) on page 415](#)
[Syntax \(TX Matrix Routers\) on page 415](#)
[Syntax \(TX Matrix Plus Routers\) on page 416](#)
[Syntax \(MX Series Routers\) on page 416](#)
[Syntax \(QFX Series\) on page 416](#)

Syntax restart

```

<adaptive-services | ancpd-service | application-identification | audit-process |
  auto-configuration | captive-portal-content-delivery | ce-l2tp-service | chassis-control |
  class-of-service | clksyncd-service | database-replication | datapath-trace-service
  | dhcp-service | diameter-service | disk-monitoring | dynamic-flow-capture |
  ecc-error-logging | ethernet-connectivity-fault-management
  | ethernet-link-fault-management | event-processing | firewall
  | general-authentication-service | gracefully | iccp-service | idp-policy | immediately
  | interface-control | ipsec-key-management | kernel-replication | l2-learning | l2cpd-service
  | l2tp-service | l2tp-universal-edge | lacp | license-service | link-management
  | local-policy-decision-function | mac-validation | mib-process | mounstd-service
  | mpls-traceroute | mspd | multicast-snooping | named-service | nfsd-service |
  packet-triggered-subscribers | peer-selection-service | pgm | pic-services-logging | pki-service
  | ppp | ppp-service | pppoe | protected-system-domain-service |
  redundancy-interface-process | remote-operations | root-system-domain-service | routing
  <logical-system logical-system-name> | sampling | sbc-configuration-process | sdk-service
  | service-deployment | services | snmp | soft | static-subscribers | statistics-service |
  subscriber-management | subscriber-management-helper | tunnel-oamd | usb-control |
  vrrp | web-management>
<gracefully | immediately | soft>

```

Syntax (ACX Series Routers)

```

restart
<adaptive-services | audit-process | auto-configuration | autoinstallation | chassis-control |
  class-of-service | clksyncd-service | database-replication | dhcp-service | diameter-service
  | disk-monitoring | dynamic-flow-capture | ethernet-connectivity-fault-management
  | ethernet-link-fault-management | event-processing | firewall
  | general-authentication-service | gracefully | immediately | interface-control |
  ipsec-key-management | l2-learning | lacp | link-management | mib-process | mounstd-service
  | mpls-traceroute | mspd | named-service | nfsd-service | pgm | pki-service | ppp | pppoe |
  redundancy-interface-process | remote-operations | routing | sampling | sdk-service
  | secure-neighbor-discovery | service-deployment | services | snmp | soft | statistics-service |
  subscriber-management | subscriber-management-helper | tunnel-oamd | vrrp>

```

Syntax (EX Series Switches)

```

restart
<autoinstallation | chassis-control | class-of-service | database-replication | dhcp |
  dhcp-service | diameter-service | dot1x-protocol | ethernet-link-fault-management |
  ethernet-switching | event-processing | firewall | general-authentication-service |
  interface-control | kernel-replication | l2-learning | lacp | license-service | link-management
  | lldpd-service | mib-process | mounstd-service | multicast-snooping | pgm |

```

	<p>redundancy-interface-process remote-operations routing secure-neighbor-discovery service-deployment sflow-service snmp vrrp web-management></p>
Syntax (MX Series Routers)	<p>restart</p> <p><adaptive-services ancpd-service application-identification audit-process auto-configuration captive-portal-content-delivery ce-l2tp-service chassis-control class-of-service clksyncd-service database-replication datapath-trace-service dhcp-service diameter-service disk-monitoring dynamic-flow-capture ecc-error-logging ethernet-connectivity-fault-management ethernet-link-fault-management event-processing firewall general-authentication-service gracefully iccp-service idp-policy immediately interface-control ipsec-key-management kernel-replication l2-learning l2cpd-service l2tp-service l2tp-universal-edge lacp license-service link-management local-policy-decision-function mac-validation mib-process mountd-service mpls-traceroute mspd multicast-snooping named-service nfsd-service packet-triggered-subscribers peer-selection-service pgm pic-services-logging pki-service ppp ppp-service pppoe protected-system-domain-service redundancy-interface-process remote-operations root-system-domain-service routing routing <logical-system <i>logical-system-name</i>> sampling sbc-configuration-process sdk-service service-deployment services snmp soft static-subscribers statistics-service subscriber-management subscriber-management-helper tunnel-oamd usb-control vrrp web-management></p> <p><all-members></p> <p><gracefully immediately soft></p> <p><local></p> <p><member <i>member-id</i>></p>
Syntax (QFX Series)	<p>restart</p> <p><adaptive-services audit-process chassis-control class-of-service dialer-services diameter-service dlsd ethernet-connectivity event-processing fibre-channel firewall general-authentication-service igmp-host-services interface-control ipsec-key-management isdn-signaling l2ald l2-learning l2tp-service mib-process named-service network-access-service nstrace-process pgm ppp pppoe redundancy-interface-process remote-operations <i>logical-system-name</i>> routing sampling secure-neighbor-discovery service-deployment snmp usb-control web-management></p> <p><gracefully immediately soft></p>
Syntax (Routing Matrix)	<p>restart</p> <p><adaptive-services audit-process chassis-control class-of-service disk-monitoring dynamic-flow-capture ecc-error-logging event-processing firewall interface-control ipsec-key-management kernel-replication l2-learning l2tp-service lacp link-management mib-process pgm pic-services-logging ppp pppoe redundancy-interface-process remote-operations routing <logical-system <i>logical-system-name</i>> sampling service-deployment snmp></p> <p><all all-lcc lcc <i>number</i>></p> <p><gracefully immediately soft></p>
Syntax (TX Matrix Routers)	<p>restart</p> <p><adaptive-services audit-process chassis-control class-of-service dhcp-service diameter-service disk-monitoring dynamic-flow-capture ecc-error-logging event-processing firewall interface-control ipsec-key-management kernel-replication l2-learning l2tp-service lacp link-management mib-process pgm pic-services-logging ppp pppoe redundancy-interface-process remote-operations routing <logical-system <i>logical-system-name</i>> sampling service-deployment snmp statistics-service></p>

	<p><all-chassis all-lcc lcc <i>number</i> scc></p> <p><gracefully immediately soft></p>
Syntax (TX Matrix Plus Routers)	<p>restart</p> <p><adaptive-services audit-process chassis-control class-of-service dhcp-service diameter-service disk-monitoring dynamic-flow-capture ecc-error-logging event-processing firewall interface-control ipsec-key-management kernel-replication l2-learning l2tp-service lacp link-management mib-process pgm pic-services-logging ppp pppoe redundancy-interface-process remote-operations routing <logical-system <i>logical-system-name</i>> sampling service-deployment snmp statistics-service></p> <p><all-chassis all-lcc all-sfc lcc <i>number</i> sfc <i>number</i>></p> <p><gracefully immediately soft></p>
Syntax (MX Series Routers)	<p>restart</p> <p><adaptive-services ancpd-service application-identification audit-process auto-configuration captive-portal-content-delivery ce-l2tp-service chassis-control class-of-service clksyncd-service database-replication datapath-trace-service dhcp-service diameter-service disk-monitoring dynamic-flow-capture ecc-error-logging ethernet-connectivity-fault-management ethernet-link-fault-management event-processing firewall general-authentication-service gracefully iccp-service idp-policy immediately interface-control ipsec-key-management kernel-replication l2-learning l2cpd-service l2tp-service l2tp-universal-edge lacp license-service link-management local-policy-decision-function mac-validation mib-process mobile-ip mounstd-service mpls-traceroute mspd multicast-snooping named-service nfsd-service packet-triggered-subscribers peer-selection-service pgcp-service pgm pic-services-logging pki-service ppp ppp-service pppoe protected-system-domain-service redundancy-interface-process remote-operations root-system-domain-service routing routing <logical-system <i>logical-system-name</i>> sampling sbc-configuration-process sdk-service service-deployment services services pgcp gateway <i>gateway-name</i> snmp soft static-subscribers statistics-service subscriber-management subscriber-management-helper tunnel-oamd usb-control vrrp web-management></p> <p><all-members></p> <p><gracefully immediately soft></p> <p><local></p> <p><member <i>member-id</i>></p>
Syntax (QFX Series)	<p>restart</p> <p><adaptive-services audit-process chassis-control class-of-service dialer-services diameter-service dlsd ethernet-connectivity event-processing fibre-channel firewall general-authentication-service igmp-host-services interface-control ipsec-key-management isdn-signaling l2ald l2-learning l2tp-service mib-process named-service network-access-service nstrace-process pgm ppp pppoe redundancy-interface-process remote-operations <i>logical-system-name</i>> routing sampling secure-neighbor-discovery service-deployment snmp usb-control web-management></p> <p><gracefully immediately soft></p>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 12.2 for ACX Series routers.</p>

Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Options added:

- **dynamic-flow-capture** in Junos OS Release 7.4.
- **dls** in Junos OS Release 7.5.
- **event-processing** in Junos OS Release 7.5.
- **ppp** in Junos OS Release 7.5.
- **l2ald** in Junos OS Release 8.0.
- **link-management** in Release 8.0.
- **pgcp-service** in Junos OS Release 8.4.
- **sbc-configuration-process** in Junos OS Release 9.5.
- **services pgcp gateway** in Junos OS Release 9.6.
- **sfc** and **all-sfc** for the TX Matrix Router in Junos OS Release 9.6.

Description Restart a Junos OS process.



CAUTION: Never restart a software process unless instructed to do so by a customer support engineer. A restart might cause the router or switch to drop calls and interrupt transmission, resulting in possible loss of data.

Options **none**—Same as **gracefully**.

adaptive-services—(Optional) Restart the configuration management process that manages the configuration for stateful firewall, Network Address Translation (NAT), intrusion detection services (IDS), and IP Security (IPsec) services on the Adaptive Services PIC.

all-chassis—(TX Matrix and TX Matrix Plus routers only) (Optional) Restart the software process on all chassis.

all-lcc—(TX Matrix and TX Matrix Plus routers only) (Optional) For a TX Matrix router, restart the software process on all T640 routers connected to the TX Matrix router. For a TX Matrix Plus router, restart the software process on all T1600 routers connected to the TX Matrix Plus router.

all-members—(MX Series routers only) (Optional) Restart the software process for all members of the Virtual Chassis configuration.

all-sfc—(TX Matrix Plus routers only) (Optional) For a TX Matrix Plus router, restart the software processes for the TX Matrix Plus router (or switch-fabric chassis).

ancpd-service—(Optional) Restart the Access Node Control Protocol (ANCP) process, which works with a special Internet Group Management Protocol (IGMP) session to collect outgoing interface mapping events in a scalable manner.

application-identification—(Optional) Restart the process that identifies an application using intrusion detection and prevention (IDP) to allow or deny traffic based on applications running on standard or nonstandard ports.

audit-process—(Optional) Restart the RADIUS accounting process that gathers statistical data that can be used for general network monitoring, analyzing, and tracking usage patterns, for billing a user based on the amount of time or type of services accessed.

auto-configuration—(Optional) Restart the Interface Auto-Configuration process.

autoinstallation—(EX Series switches only) (Optional) Restart the autoinstallation process.

captive-portal-content-delivery—(Optional) Restart the HTTP redirect service by specifying the location to which a subscriber's initial Web browser session is redirected, enabling initial provisioning and service selection for the subscriber.

ce-l2tp-service—(M10, M10i, M7i, and MX Series routers only) (Optional) Restart the Universal Edge Layer 2 Tunneling Protocol (L2TP) process, which establishes L2TP tunnels and Point-to-Point Protocol (PPP) sessions through L2TP tunnels.

chassis-control—(Optional) Restart the chassis management process.

class-of-service—(Optional) Restart the class-of-service (CoS) process, which controls the router's or switch's CoS configuration.

clksyncd-service—(Optional) Restart the external clock synchronization process, which uses synchronous Ethernet (SyncE).

database-replication—(EX Series switches and MX Series routers only) (Optional) Restart the database replication process.

datapath-trace-service—(Optional) Restart the packet path tracing process.

dhcp—(EX Series switches only) (Optional) Restart the software process for a Dynamic Host Configuration Protocol (DHCP) server. A DHCP server allocates network IP addresses and delivers configuration settings to client hosts without user intervention.

dhcp-service—(Optional) Restart the Dynamic Host Configuration Protocol process.

dialer-services—(EX Series switches only) (Optional) Restart the ISDN dial-out process.

diameter-service—(Optional) Restart the diameter process.

disk-monitoring—(Optional) Restart disk monitoring, which checks the health of the hard disk drive on the Routing Engine.

dls—(QFX Series only) (Optional) Restart the data link switching (DLSw) service.

dot1x-protocol—(EX Series switches only) (Optional) Restart the port-based network access control process.

dynamic-flow-capture—(Optional) Restart the dynamic flow capture (DFC) process, which controls DFC configurations on Monitoring Services III PICs.

ecc-error-logging—(Optional) Restart the error checking and correction (ECC) process, which logs ECC parity errors in memory on the Routing Engine.

ethernet-connectivity-fault-management—(Optional) Restart the process that provides IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management (CFM) database information for CFM maintenance association end points (MEPs) in a CFM session.

ethernet-link-fault-management—(EX Series switches and MX Series routers only)
(Optional) Restart the process that provides the OAM link fault management (LFM) information for Ethernet interfaces.

ethernet-switching—(EX Series switches only) (Optional) Restart the Ethernet switching process.

event-processing—(Optional) Restart the event process (eventd).

fibre-channel—(QFX Series only) (Optional) Restart the Fibre Channel process.

firewall—(Optional) Restart the firewall management process, which manages the firewall configuration and enables accepting or rejecting packets that are transiting an interface on a router or switch.

general-authentication-service—(EX Series switches and MX Series routers only)
(Optional) Restart the general authentication process.

gracefully—(Optional) Restart the software process.

iccp-service—(Optional) Restart the Inter-Chassis Communication Protocol (ICCP) process.

idp-policy—(Optional) Restart the intrusion detection and prevention (IDP) protocol process.

immediately—(Optional) Immediately restart the software process.

interface-control—(Optional) Restart the interface process, which controls the router's or switch's physical interface devices and logical interfaces.

ipsec-key-management—(Optional) Restart the IPsec key management process.

isdn-signaling—(QFX Series only) (Optional) Restart the ISDN signaling process, which initiates ISDN connections.

kernel-replication—(Optional) Restart the kernel replication process, which replicates the state of the backup Routing Engine when graceful Routing Engine switchover (GRES) is configured.

l2-learning—(Optional) Restart the Layer 2 address flooding and learning process.

l2cpd-service—(Optional) Restart the Layer 2 Control Protocol process, which enables features such as Layer 2 protocol tunneling and nonstop bridging.

l2tp-service— (M10, M10i, M7i, and MX Series routers only) (Optional) Restart the Layer 2 Tunneling Protocol (L2TP) process, which sets up client services for establishing Point-to-Point Protocol (PPP) tunnels across a network and negotiating Multilink PPP if it is implemented.

l2tp-universal-edge— (MX Series routers only) (Optional) Restart the L2TP process, which establishes L2TP tunnels and PPP sessions through L2TP tunnels.

lACP— (Optional) Restart the Link Aggregation Control Protocol (LACP) process. LACP provides a standardized means for exchanging information between partner systems on a link to allow their link aggregation control instances to reach agreement on the identity of the LAG to which the link belongs, and then to move the link to that LAG, and to enable the transmission and reception processes for the link to function in an orderly manner.

lcc *number*— (TX Matrix and TX Matrix Plus routers only) (Optional) For a TX Matrix router, restart the software process for a specific T640 router that is connected to the TX Matrix router. For a TX Matrix Plus router, restart the software process for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

license-service— (EX Series switches only) (Optional) Restart the feature license management process.

link-management— (TX Matrix and TX Matrix Plus routers and EX Series switches only) (Optional) Restart the Link Management Protocol (LMP) process, which establishes and maintains LMP control channels.

lldpd-service— (EX Series switches only) (Optional) Restart the Link Layer Discovery Protocol (LLDP) process.

local— (MX Series routers only) (Optional) Restart the software process for the local Virtual Chassis member.

local-policy-decision-function— (Optional) Restart the process for the Local Policy Decision Function, which regulates collection of statistics related to applications and application groups and tracking of information about dynamic subscribers and static interfaces.

mac-validation— (Optional) Restart the Media Access Control (MAC) validation process, which configures MAC address validation for subscriber interfaces created on demux interfaces in dynamic profiles on MX Series routers.

member *member-id*— (MX Series routers only) (Optional) Restart the software process for a specific member of the Virtual Chassis configuration. Replace ***member-id*** with a value of 0 or 1.

mib-process— (Optional) Restart the Management Information Base (MIB) version II process, which provides the router's MIB II agent.

mobile-ip— (Optional) Restart the Mobile IP process, which configures Junos OS Mobile IP features.

moundd-service— (EX Series switches and MX Series routers only) (Optional) Restart the service for NFS mount requests.

mpls-traceroute— (Optional) Restart the MPLS Periodic Traceroute process.

mspd— (Optional) Restart the Multiservice process.

multicast-snooping— (EX Series switches and MX Series routers only) (Optional) Restart the multicast snooping process, which makes Layer 2 devices, such as VLAN switches, aware of Layer 3 information, such as the media access control (MAC) addresses of members of a multicast group.

named-service— (Optional) Restart the DNS Server process, which is used by a router or a switch to resolve hostnames into addresses.

network-access-service— (QFX Series only) (Optional) Restart the network access process, which provides the router's Challenge Handshake Authentication Protocol (CHAP) authentication service.

nfsd-service— (Optional) Restart the Remote NFS Server process, which provides remote file access for applications that need NFS-based transport.

packet-triggered-subscribers— (Optional) Restart the packet-triggered subscribers and policy control (PTSP) process, which allows the application of policies to dynamic subscribers that are controlled by a subscriber termination device.

peer-selection-service— (Optional) Restart the Peer Selection Service process.

pgcp-service— (Optional) Restart the pgcpd service process running on the Routing Engine. This option does not restart pgcpd processes running on mobile station PICs. To restart pgcpd processes running on mobile station PICs, use the **services pgcp gateway** option.

pgm— (Optional) Restart the process that implements the Pragmatic General Multicast (PGM) protocol for assisting in the reliable delivery of multicast packets.

pic-services-logging— (Optional) Restart the logging process for some PICs. With this process, also known as fsad (the file system access daemon), PICs send special logging information to the Routing Engine for archiving on the hard disk.

pki-service—(Optional) Restart the PKI Service process.

ppp—(Optional) Restart the Point-to-Point Protocol (PPP) process, which is the encapsulation protocol process for transporting IP traffic across point-to-point links.

ppp-service—(Optional) Restart the Universal edge PPP process, which is the encapsulation protocol process for transporting IP traffic across universal edge routers.

pppoe—(Optional) Restart the Point-to-Point Protocol over Ethernet (PPPoE) process, which combines PPP that typically runs over broadband connections with the Ethernet link-layer protocol that allows users to connect to a network of hosts over a bridge or access concentrator.

protected-system-domain-service—(Optional) Restart the Protected System Domain (PSD) process.

redundancy-interface-process—(Optional) Restart the ASP redundancy process.

remote-operations—(Optional) Restart the remote operations process, which provides the ping and traceroute MIBs.

root-system-domain-service—(Optional) Restart the Root System Domain (RSD) service.

routing—(ACX Series routers, QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the routing protocol process.

routing <logical-system *logical-system-name*>—(Optional) Restart the routing protocol process, which controls the routing protocols that run on the router or switch and maintains the routing tables. Optionally, restart the routing protocol process for the specified logical system only.

sampling—(Optional) Restart the sampling process, which performs packet sampling based on particular input interfaces and various fields in the packet header.

sbc-configuration-process—(Optional) Restart the session border controller (SBC) process of the border signaling gateway (BSG).

scc—(TX Matrix routers only) (Optional) Restart the software process on the TX Matrix router (or switch-card chassis).

sdk-service—(Optional) Restart the SDK Service process, which runs on the Routing Engine and is responsible for communications between the SDK application and Junos OS. Although the SDK Service process is present on the router, it is turned off by default.

secure-neighbor-discovery—(QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the secure Neighbor Discovery Protocol (NDP) process, which provides support for protecting NDP messages.

sfc *number*—(TX Matrix Plus routers only) (Optional) Restart the software process on the TX Matrix Plus router (or switch-fabric chassis). Replace *number* with 0.

service-deployment—(Optional) Restart the service deployment process, which enables Junos OS to work with the Session and Resource Control (SRC) software.

services—(Optional) Restart a service.

services pgcp gateway gateway-name—(Optional) Restart the pgcpd process for a specific border gateway function (BGF) running on an MS-PIC. This option does not restart the pgcpd process running on the Routing Engine. To restart the pgcpd process on the Routing Engine, use the **pgcp-service** option.

sflow-service—(EX Series switches only) (Optional) Restart the flow sampling (sFlow technology) process.

snmp—(Optional) Restart the SNMP process, which enables the monitoring of network devices from a central location and provides the router's or switch's SNMP master agent.

soft—(Optional) Reread and reactivate the configuration without completely restarting the software processes. For example, BGP peers stay up and the routing table stays constant. Omitting this option results in a graceful restart of the software process.

static-subscribers—(Optional) Restart the static subscribers process, which associates subscribers with statically configured interfaces and provides dynamic service activation and activation for these subscribers.

statistics-service—(Optional) Restart the process that manages the Packet Forwarding Engine statistics.

subscriber-management—(Optional) Restart the Subscriber Management process.

subscriber-management-helper—(Optional) Restart the Subscriber Management Helper process.

tunnel-oamd—(Optional) Restart the Tunnel OAM process, which enables the Operations, Administration, and Maintenance of Layer 2 tunneled networks. Layer 2 protocol tunneling (L2PT) allows service providers to send Layer 2 protocol data units (PDUs) across the provider's cloud and deliver them to Juniper Networks EX Series Ethernet Switches that are not part of the local broadcast domain.

usb-control—(MX Series routers) (Optional) Restart the USB control process.

vrrp—(ACX Series routers, EX Series switches, and MX Series routers only) (Optional) Restart the Virtual Router Redundancy Protocol (VRRP) process, which enables hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts.

web-management—(QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the Web management process.

Required Privilege Level reset

**Related
Documentation**

- [Overview of Junos OS CLI Operational Mode Commands on page 58](#)

List of Sample Output [restart interfaces on page 424](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

restart interfaces

```
user@host> restart interfaces
interfaces process terminated
interfaces process restarted
```


save

Syntax	<code>save <i>filename</i></code>
QFX Series	<code>save (dhcp-snooping <i>filename</i>)</code>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Save the configuration to an ASCII file. The contents of the current level of the statement hierarchy (and below) are saved, along with the statement hierarchy containing it. This allows a section of the configuration to be saved, while fully specifying the statement hierarchy.</p> <p>When saving a file to a remote system, the software uses the scp/ssh protocol.</p>
Options	<p><i>filename</i>—Name of the saved file. You can specify a filename in one of the following ways:</p> <ul style="list-style-type: none"> • <i>filename</i>—File in the user's home directory (the current directory) on the local flash drive. • <i>path/filename</i>—File on the local flash drive. • <i>/var/filename</i> or <i>/var/path/filename</i>—File on the local hard disk. • <i>a:filename</i> or <i>a:path/filename</i>—File on the local drive. The default path is / (the root-level directory). The removable media can be in MS-DOS or UNIX (UFS) format. • <i>hostname:/path/filename</i>, <i>hostname:filename</i>, <i>hostname:path/filename</i>, or <i>scp://hostname/path/filename</i>—File on an scp/ssh client. This form is not available in the worldwide version of Junos OS. The default path is the user's home directory on the remote system. You can also specify <i>hostname</i> as <i>username@hostname</i>. • <i>ftp://hostname/path/filename</i>—File on an FTP server. You can also specify <i>hostname</i> as <i>username @hostname</i> or <i>username:password @hostname</i>. The default path is the user's home directory. To specify an absolute path, the path must start with the string %2F; for example, <i>ftp://hostname/%2Fpath/filename</i>. To have the system prompt you for the password, specify <i>prompt</i> in place of the password. If a password is required, and you do not specify the password or <i>prompt</i>, an error message is displayed: <pre> user@host> file copy ftp://username@ftp.hostname.net//filename file copy ftp.hostname.net: Not logged in. user@host> file copy ftp://username:prompt@ftphostname.net//filename </pre> <p>Password for <i>username@ftp.hostname.net</i>:</p> • <i>http://hostname/path/filename</i>—File on a Hypertext Transfer Protocol (HTTP) server. You can also specify <i>hostname</i> as <i>username@hostname</i> or <i>username:password@hostname</i>. If a password is required and you omit it, you are prompted for it. • <i>re0:/path/filename</i> or <i>re1:/path/filename</i>—File on a local Routing Engine.

Required Privilege Level configure—To enter configuration mode.

Related Documentation • *Deactivating and Reactivating Statements and Identifiers in a Junos OS Configuration*

CHAPTER 27

CLI Operational Commands

- `show cli`
- `show cli authorization`
- `show cli directory`
- `show cli history`

show cli

Syntax	show cli
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display configured CLI settings.
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show cli on page 428
Output Fields	Table 36 lists the output fields for the show cli command. Output fields are listed in the approximate order in which they appear.

Table 36: show cli Output Fields

Field Name	Field Description
CLI complete-on-space	Capability to complete a partial command entry when you type a space or a tab: on or off .
CLI idle-timeout	Maximum time that an individual session can be idle before the user is logged out from the router or switch. When this feature is enabled, the number of minutes is displayed. Otherwise, the state is disabled .
CLI restart-on-upgrade	CLI is set to prompt you to restart the router or switch after upgrading the software: on or off .
CLI screen-length	Number of lines of text that the terminal screen displays.
CLI screen-width	Number of characters in a line on the terminal screen.
CLI terminal	Terminal type.
CLI is operating in	Mode: enhanced .
CLI timestamp	Date and time format for the timestamp. If the timestamp is not set, the state is disabled .
CLI working directory	Pathname of the working directory.

Sample Output

show cli

```
user@host> show cli
CLI complete-on-space set to on
CLI idle-timeout disabled
CLI restart-on-upgrade set to on
CLI screen-length set to 47
CLI screen-width set to 132
```

```
CLI terminal is 'vt100'  
CLI is operating in enhanced mode  
CLI timestamp disabled  
CLI working directory is '/var/home/regress'
```

show cli authorization

Syntax	show cli authorization
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Display the permissions for the current user.
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show cli authorization on page 432
Output Fields	Table 37 lists the output fields for the show cli authorization command. In the table, all possible permissions are displayed and output fields are listed in alphabetical order.

Table 37: show cli authorization Output Fields

Field Name	Field Description
access	Can view access configuration information.
access-control	Can modify access configuration.
admin	Can view user account information.
admin-control	Can modify user account information.
clear	Can clear learned network information.
configure	Can enter configuration mode.
control	Can modify any configuration.
edit	Can edit configuration files.
field	Reserved for field (debugging) support.
firewall	Can view firewall configuration information.
firewall-control	Can modify firewall configuration information.
floppy	Can read from and write to removable media.
flow-tap	Can view flow-tap configuration information.

Table 37: show cli authorization Output Fields (*continued*)

Field Name	Field Description
flow-tap-control	Can configure flow-tap configuration information.
idp-profiler-operation	Can configure Profiler data.
interface	Can view interface configuration information.
interface-control	Can modify interface configuration information.
maintenance	Can perform system maintenance.
network	Can access the network by entering the ping , ssh , telnet , and traceroute commands.
pgcp-session-mirroring	Can view Packet Gateway Control Protocol session mirroring configuration.
pgcp-session-mirroring-control	Can modify Packet Gateway Control Protocol session mirroring configuration all-control.
reset	Can reset or restart interfaces and system processes.
rollback	Can roll back to previous configurations.
routing	Can view routing configuration information.
routing-control	Can modify routing configuration information.
secret	Can view passwords and authentication keys in the configuration.
secret-control	Can modify passwords and authentication keys in the configuration.
security	Can view security configuration information.
security-control	Can modify security configuration information.
shell	Can start a local shell.
snmp	Can view SNMP configuration information.
snmp-control	Can modify SNMP configuration information.
system	Can view system configuration information.
system-control	Can modify system configuration information.
trace	Can view trace file settings information.

Table 37: show cli authorization Output Fields (*continued*)

Field Name	Field Description
trace-control	Can modify trace file settings information.
view	Can view current values and statistics.
view-configuration	Can view all configuration information (not including secrets).

Sample Output

show cli authorization

```

user@host> show cli authorization
Current user: 'remote' login: 'user' class ''
Permissions:
  admin      -- Can view user accounts
  admin-control-- Can modify user accounts
  clear      -- Can clear learned network information
  configure  -- Can enter configuration mode
  control    -- Can modify any configuration
  edit       -- Can edit full files
  field      -- Special for field (debug) support
  floppy     -- Can read and write from the floppy
  interface  -- Can view interface configuration
  interface-control-- Can modify interface configuration
  network    -- Can access the network
  reset      -- Can reset/restart interfaces and daemons
  routing    -- Can view routing configuration
  routing-control-- Can modify routing configuration
  shell      -- Can start a local shell
  snmp       -- Can view SNMP configuration
  snmp-control-- Can modify SNMP configuration
  system     -- Can view system configuration
  system-control-- Can modify system configuration
  trace      -- Can view trace file settings
  trace-control-- Can modify trace file settings
  view       -- Can view current values and statistics
  maintenance -- Can become the super-user
  firewall   -- Can view firewall configuration
  firewall-control-- Can modify firewall configuration
  secret     -- Can view secret configuration
  secret-control-- Can modify secret configuration
  rollback   -- Can rollback to previous configurations
  security   -- Can view security configuration
  security-control-- Can modify security configuration
  access     -- Can view access configuration
  access-control-- Can modify access configuration
  view-configuration-- Can view all configuration (not including secrets)
  flow-tap   -- Can view flow-tap configuration
  flow-tap-control-- Can configure flow-tap service
Individual command authorization:
  Allow regular expression: none
  Deny regular expression: none
  Allow configuration regular expression: none
  Deny configuration regular expression: none

```


show cli directory

Syntax show cli directory

Release Information Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Display the current working directory.

Options This command has no options.

Required Privilege Level view

List of Sample Output [show cli directory on page 434](#)

Output Fields [Table 38](#) lists the output fields for the **show cli directory** command. Output fields are listed in the approximate order in which they appear.

Table 38: show cli directory Output Fields

Field Name	Field Description
Current directory	Pathname of the current working directory.

Sample Output

show cli directory

```
user@host> show cli directory
Current directory: /var/home/regress
```

show cli history

Syntax	<code>show cli history</code> <code><count></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Display a list of previous CLI commands.
Options	none —Display all previous CLI commands. count —(Optional) Maximum number of commands to display.
Required Privilege Level	view
List of Sample Output	show cli history on page 435
Output Fields	Table 39 lists the output fields for the show cli history command. Output fields are listed in the approximate order in which they appear.

Table 39: show cli history Output Fields

Field Name	Field Description
<i>timestamp</i>	Time at which the command was entered.
<i>command-syntax</i>	Command that was entered.

Sample Output

show cli history

```
user@host> show cli history
11:14:14 -- show arp
11:22:10 -- show cli authorization
11:27:12 -- show cli history
```


CHAPTER 28

Licensing Operational Commands

- request system license add
- request system license delete
- request system license save
- show system license

request system license add

Syntax	<code>request system license add (<i>filename</i> terminal)</code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 9.5 for SRX Series devices. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Add a license key.
Options	<i>filename</i> —License key from a file or URL. Specify the filename or the URL where the key is located. <i>terminal</i> —License key from the terminal.
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none">• Adding New Licenses (CLI Procedure) on page 117
List of Sample Output	request system license add on page 438
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system license add

```
user@host> request system license add terminal
E408408918 aeaqib qcsbj a okbuq rcmxnq vjocwf uxfsta
          z5ufjb kdrmt6 57bimv 2f3ddp qttcdn 627q4a
          jx4s5x hiri
E408408918: successfully added
add license complete (no errors)
```

request system license delete

Syntax	<code>request system license delete (<i>license-identifier</i> license-identifier-list [<i>licenseid001</i> <i>licenseid002</i> <i>licenseid003</i>] all)</code>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Option license-identifier-list introduced in Junos OS Release 13.1.</p>
Description	Delete a license key. You can choose to delete one license at a time, all licenses at once, or a list of license identifiers enclosed in brackets.
Options	<p>license-identifier—Text string that uniquely identifies a license key.</p> <p>license-identifier-list [<i>licenseid001</i> <i>licenseid002</i> <i>licenseid003</i>....]—Delete multiple license identifiers as a list enclosed in brackets.</p> <p>all—Delete all licenses on the device.</p>
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none"> • Deleting a License (CLI Procedure) on page 118

request system license save

Syntax	<code>request system license save (<i>filename</i> terminal)</code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 9.5 for SRX Series devices.
Description	Save installed license keys to a file or URL.
Options	<i>filename</i> —License key from a file or URL. Specify the filename or the URL where the key is located. <i>terminal</i> —License key from the terminal.
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none">• Saving License Keys on page 119
List of Sample Output	request system license save on page 440
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system license save

```
user@host> request system license save ftp://user@host/license.conf
```


show system license

Syntax	show system license <installed keys usage>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 13.3 for the MX104 3D Universal Edge Routers.
Description	Display licenses and information about how they are used.
Options	<p>none—Display all license information.</p> <p>installed—(Optional) Display installed licenses only.</p> <p>keys—(Optional) Display a list of license keys. Use this information to verify that each expected license key is present.</p> <p>usage—(Optional) Display the state of licensed features.</p>
Required Privilege Level	maintenance
List of Sample Output	show system license on page 442 show system license installed on page 443 show system license keys on page 443 show system license usage on page 443 show system license (MX104 Routers) on page 443 show system license installed (MX104 Routers) on page 444 show system license keys (MX104 Routers) on page 444 show system license usage (MX104 Routers) on page 444 show system license (MX104 Routers) on page 444 show system license installed (MX104 Routers) on page 445 show system license keys (MX104 Routers) on page 445 show system license usage (MX104 Routers) on page 445 show system license (MX104 Routers) on page 446 show system license installed (MX104 Routers) on page 446 show system license keys (MX104 Routers) on page 446 show system license usage (MX104 Routers) on page 447 show system license (QFX Series) on page 447
Output Fields	Table 40 lists the output fields for the show system license command. Output fields are listed in the approximate order in which they appear.

Table 40: show system license Output Fields

Field Name	Field Description
Feature name	Name assigned to the configured feature. You use this information to verify that all the features for which you installed licenses are present.

Table 40: show system license Output Fields (*continued*)

Field Name	Field Description
Licenses used	<p>Number of licenses used by a router or switch. You use this information to verify that the number of licenses used matches the number configured. If a licensed feature is configured, the feature is considered used.</p> <p>NOTE: In Junos OS Release 10.1 and later, the Licenses used column displays the actual usage count based on the number of active sessions or connections as reported by the corresponding feature daemons. This is applicable for scalable license-based features such as Subscriber Access (scale-subscriber), L2TP (scale-l2tp), Mobile IP (scale-mobile-ip), and so on.</p>
Licenses installed	<p>Information about the installed license key:</p> <ul style="list-style-type: none"> • License identifier—Identifier associated with a license key. • State—State of the license key: valid or invalid. An invalid state indicates that the key was entered incorrectly or is not valid for the specific device. • License version—Version of a license. The version indicates how the license is validated, the type of signature, and the signer of the license key. • Valid for device—Device that can use a license key. • Group defined—Group membership of a device. • Features—Feature associated with a license, such as data link switching (DLSw).
Licenses needed	Number of licenses required for features being used but not yet properly licensed.
Expiry	Amount of time left within the grace period before a license is required for a feature being used.

Sample Output

show system license

```
user@host> show system license
```

```
License usage:
```

Feature name	Licenses used	Licenses installed	Licenses needed	Expiry
subscriber-accounting	2	2	0	permanent
subscriber-authentication	1	2	0	permanent
subscriber-address-assignment	2	2	0	permanent
subscriber-vlan	2	2	0	permanent
subscriber-ip	0	2	0	permanent
scale-subscriber	2	3	0	permanent
scale-l2tp	4	5	0	permanent
scale-mobile-ip	1	2	0	permanent

```
Licenses installed:
```

```
License identifier: XXXXXXXXXX
```

```
License version: 2
```

```
Features:
```

```
subscriber-accounting - Per Subscriber Radius Accounting
permanent
subscriber-authentication - Per Subscriber Radius Authentication
permanent
subscriber-address-assignment - Radius/SRC Address Pool Assignment
permanent
subscriber-vlan - Dynamic Auto-sensed Vlan
```

```

    permanent
subscriber-ip    - Dynamic and Static IP
    permanent

```

show system license installed

```

user@host> show system license installed
License identifier: XXXXXXXXXX
License version: 2
Features:
  subscriber-accounting - Per Subscriber Radius Accounting
    permanent
  subscriber-authentication - Per Subscriber Radius Authentication
    permanent
  subscriber-address-assignment - Radius/SRC Address Pool Assignment
    permanent
  subscriber-vlan - Dynamic Auto-sensed Vlan
    permanent
  subscriber-ip - Dynamic and Static IP
    permanent

```

show system license keys

```

user@host> show system license keys
XXXXXXXXXX xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx
          xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx
          xxxxxx xxxxxx xxx

```

show system license usage

```

user@host> show system license usage
License usage:

```

Feature name	Licenses used	Licenses installed	Licenses needed	Expiry
subscriber-accounting	2	2	0	permanent
subscriber-authentication	1	2	0	permanent
subscriber-address-assignment	2	2	0	permanent
subscriber-vlan	2	2	0	permanent
subscriber-ip	0	2	0	permanent
scale-subscriber	2	3	0	permanent
scale-l2tp	4	5	0	permanent
scale-mobile-ip	1	2	0	permanent

show system license (MX104 Routers)

In the following output, ports 0 and 1 are activated by installing the license to activate the first two built-in ports.

```

user@host> show system license
License usage:

```

Feature name	Licenses used	Licenses installed	Licenses needed	Expiry
scale-subscriber	0	1000	0	permanent
scale-l2tp	0	1000	0	permanent
scale-mobile-ip	0	1000	0	permanent
MX104-2x10Gig-port-0-1	0	1	0	permanent

```

Licenses installed:
License identifier: XXXXXXXXXX
License version: 2
Features:

```

```

MX104-2x10Gig-port-0-1 - MX104 2X10Gig Builtin Port(xe-2/0/0 & xe-2/0/1)
upgrade
    permanent

```

show system license installed (MX104 Routers)

In the following output, ports 0 and 1 are activated by installing the license to activate the first two built-in ports.

```

user@host > show system license installed
License identifier: XXXXXXXXXX
License version: 2
Features:
MX104-2x10Gig-port-0-1 - MX104 2X10Gig Builtin Port(xe-2/0/0 & xe-2/0/1)
upgrade
    permanent

```

show system license keys (MX104 Routers)

In the following output, ports 0 and 1 are activated by installing the license to activate the first two built-in ports.

```

user@host > show system license keys

XXXXXXXXXX xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx
          xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx
          xxxxxx xxxx

```

show system license usage (MX104 Routers)

In the following output, ports 0 and 1 are activated by installing the license to activate the first two built-in ports.

```

user@host > show system license usage

```

Feature name	Licenses used	Licenses installed	Expiry needed	
scale-subscriber	0	1000	0	permanent
scale-l2tp	0	1000	0	permanent
scale-mobile-ip	0	1000	0	permanent
MX104-2x10Gig-port-0-1	0	1	0	permanent

show system license (MX104 Routers)

In the following output, ports 2 and 3 are activated by installing the license to activate the next two built-in ports after installing the license to activate the first two built-in ports.

```

user@host > show system license
License usage:

```

Feature name	Licenses used	Licenses installed	Licenses needed	Expiry
scale-subscriber	0	1000	0	permanent
scale-l2tp	0	1000	0	permanent
scale-mobile-ip	0	1000	0	permanent
MX104-2x10Gig-port-0-1	0	1	0	permanent
MX104-2x10Gig-port-2-3	0	1	0	permanent

```

Licenses installed:
License identifier: XXXXXXXXXX
License version: 2

```

```

Features:
MX104-2x10Gig-port-0-1 - MX104 2X10Gig Builtin Port(xe-2/0/0 & xe-2/0/1)
upgrade
  permanent

License identifier: XXXXXXXXXX
License version: 2
Features:
MX104-2x10Gig-port-2-3 - MX104 2X10Gig Builtin Port(xe-2/0/2 & xe-2/0/3)
upgrade
  permanent

```

show system license installed (MX104 Routers)

In the following output, ports 2 and 3 are activated by installing the license to activate the next two built-in ports after installing the license to activate the first two built-in ports.

```

user@host > show system license installed
License identifier: XXXXXXXXXX
License version: 2
Features:
MX104-2x10Gig-port-0-1 - MX104 2X10Gig Builtin Port(xe-2/0/0 & xe-2/0/1)
upgrade
  permanent

License identifier: XXXXXXXXXX
License version: 2
Features:
MX104-2x10Gig-port-2-3 - MX104 2X10Gig Builtin Port(xe-2/0/2 & xe-2/0/3)
upgrade
  permanent

```

show system license keys (MX104 Routers)

In the following output, ports 2 and 3 are activated by installing the license to activate the next two built-in ports after installing the license to activate the first two built-in ports.

```

user@host > show system license keys

XXXXXXXXXX xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx
          xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx
          xxxxxx xxxx

XXXXXXXXXX xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx
          xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx
          xxxxxx xxxx

```

show system license usage (MX104 Routers)

In the following output, ports 2 and 3 are activated by installing the license to activate the next two built-in ports after installing the license to activate the first two built-in ports.

```

user@host > show system license usage

```

Feature name	Licenses used	Licenses installed	Expiry needed	
scale-subscriber	0	1000	0	permanent
scale-l2tp	0	1000	0	permanent

scale-mobile-ip	0	1000	0	permanent
MX104-2x10Gig-port-0-1	0	1	0	permanent
MX104-2x10Gig-port-2-3	0	1	0	permanent

show system license (MX104 Routers)

In the following output, ports 0,1,2, and 3 are activated by installing a single license key to activate all four built-in ports.

```
user@host > show system license
```

License usage:

Feature name	Licenses used	Licenses installed	Licenses needed	Expiry
scale-subscriber	0	1000	0	permanent
scale-l2tp	0	1000	0	permanent
scale-mobile-ip	0	1000	0	permanent
MX104-2x10Gig-port-0-1	0	1	0	permanent
MX104-2x10Gig-port-2-3	0	1	0	permanent

Licenses installed:

License identifier: XXXXXXXXXX

License version: 2

Features:

MX104-2x10Gig-port-0-1 - MX104 2X10Gig Builtin Port(xe-2/0/0 & xe-2/0/1)

upgrade

permanent

MX104-2x10Gig-port-2-3 - MX104 2X10Gig Builtin Port(xe-2/0/2 & xe-2/0/3)

upgrade

permanent

show system license installed (MX104 Routers)

In the following output, ports 0,1,2, and 3 are activated by installing a single license key to activate all four built-in ports.

```
user@host > show system license installed
```

License identifier: XXXXXXXXXX

License version: 2

Features:

MX104-2x10Gig-port-0-1 - MX104 2X10Gig Builtin Port(xe-2/0/0 & xe-2/0/1)

upgrade

permanent

MX104-2x10Gig-port-2-3 - MX104 2X10Gig Builtin Port(xe-2/0/2 & xe-2/0/3)

upgrade

permanent

show system license keys (MX104 Routers)

In the following output, ports 0,1,2, and 3 are activated by installing a single license key to activate all four built-in ports.

```
user@host > show system license keys
```

```
XXXXXXXX XXXXXX XXXXXX XXXXXX XXXXXX
XXXXXXXX XXXXXX XXXXXX XXXXXX XXXXXX
XXXXXXXX XXXXXX X
```

show system license usage (MX104 Routers)

In the following output, ports 0,1,2, and 3 are activated by installing a single license key to activate all four built-in ports.

```
user@host > show system license usage
```

Feature name	Licenses used	Licenses installed	Licenses needed	Expiry
scale-subscriber	0	1000	0	permanent
scale-l2tp	0	1000	0	permanent
scale-mobile-ip	0	1000	0	permanent
MX104-2x10Gig-port-0-1	0	1	0	permanent
MX104-2x10Gig-port-2-3	0	1	0	permanent

show system license (QFX Series)

```
user@switch> show system license
```

License usage:

Feature name	Licenses used	Licenses installed	Licenses needed	Expiry
qfx-edge-fab	1	1	1	permanent

Licenses installed:
 License identifier: JUNOS417988
 License version: 1
 Features:
 qfx-edge-fab - QFX3000 Series QF/Node feature license
 permanent

CHAPTER 29

NTP Operational Commands

- `show ntp associations`
- `show ntp status`

show ntp associations

Syntax	<code>show ntp associations</code> <code><no-resolve></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Display Network Time Protocol (NTP) peers and their state.
Options	none —Display NTP peers and their state. no-resolve —(Optional) Suppress symbolic addressing.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> show ntp status on page 452
List of Sample Output	show ntp associations on page 451
Output Fields	Table 41 describes the output fields for the show ntp associations command. Output fields are listed in the approximate order in which they appear.

Table 41: show ntp associations Output Fields

Field Name	Field Description
remote	Address or name of the remote NTP peer.
refid	Reference identifier of the remote peer. If the reference identifier is not known, this field shows a value of 0.0.0.0 .
st	Stratum of the remote peer.
t	Type of peer: b (broadcast), l (local), m (multicast), or u (unicast).
when	When the last packet from the peer was received.
poll	Polling interval, in seconds.
reach	Reachability register, in octal.
delay	Current estimated delay of the peer, in milliseconds.
offset	Current estimated offset of the peer, in milliseconds.
disp	Current estimated dispersion of the peer, in milliseconds.

Table 41: show ntp associations Output Fields (*continued*)

Field Name	Field Description
<i>peer-name</i>	<p>Peer name and status of the peer in the clock selection process:</p> <ul style="list-style-type: none"> • space—Discarded because of a high stratum value or failed sanity checks. • x—Designated "falseticker" by the intersection algorithm. • .—Culled from the end of the candidate list. • — —Discarded by the clustering algorithm. • +—Included in the final selection set. • #—Selected for synchronization, but the distance exceeds the maximum. • *—Selected for synchronization. • o—Selected for synchronization, but the packets-per-second (pps) signal is in use.

Sample Output

show ntp associations

```

user@host> show ntp associations
      remote          refid      st t when poll reach  delay  offset  disp
=====
*wolfe-gw.junipe tick.ucla.edu    2 u  43   64  377    1.86    0.319    0.08

```

show ntp status

Syntax	<code>show ntp status</code> <code><no-resolve></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Display the values of internal variables returned by Network Time Protocol (NTP) peers.
Options	none —Display the values of internal variables returned by NTP peers. no-resolve —(Optional) Suppress symbolic addressing.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> show ntp associations on page 450
List of Sample Output	show ntp status on page 453
Output Fields	Table 42 describes the output fields for the show ntp status command. Output fields are listed in the approximate order in which they appear.

Table 42: show ntp status Output Fields

Field Name	Field Description
status	System status word, a code representing the status items listed.
leap_none	Indicates a normal synchronized state with no leap seconds imminent. Other options could be leap_add_sec , leap_del_sec , or leap_alarm , indicating a leap second will be added, deleted, or a leap second requirement is upcoming.
sync_ntp	Indicates the current synchronization source, in this case, an NTP server. Other options include sync_alarm and sync_unspec , both indicating that the router has not been synched.
x events	Indicates the number of events that have occurred since that last code change. An event is often the receipt of an NTP polling message.
event_peer/strat_chg	Describes the most recent event, in this case, the stratum of the peer server changed.
version	A detailed description of the version of NTP being used.
processor	Indicates the current hardware platform and version of the processor.
system	Detailed description of the name and version of the operating system in use.
leap	The number of leap seconds in use.

Table 42: show ntp status Output Fields (*continued*)

Field Name	Field Description
stratum	The stratum of the peer server. Anything greater than 1 is a secondary reference source, and the number roughly represents the number of hops away from the stratum 1 server.. Stratum 1 is a primary reference, such as an atomic clock.
precision	The precision of the peer clock, how precisely the frequency and time can be maintained with this particular timekeeping system.
rootdelay	The total roundtrip delay to the primary reference source, in seconds.
rootdispersion	The maximum error relative to the primary reference source, in seconds.
peer	An identification number of the peer in use.
refid	Reference identifier of the remote peer. If the reference identifier is not known, this field shows a value of 0.0.0.0.
reftime	The local time, in timestamp format, when the local clock was last updated. If the local clock has never been synchronized, the value is zero.
poll	The NTP broadcast message polling interval, in seconds.
clock	The current time on the local router clock.
state	The current mode of NTP operation, where 1 is symmetric active, 2 is symmetric passive, 3 is client, 4 is server, and 5 is broadcast.
offset	Current estimated offset of the peer, in milliseconds. Indicates the time difference between the reference clock and the local clock.
frequency	The frequency of the clock.
jitter	Indicates the magnitude of jitter, in milliseconds, between several time queries.
stability	A measure of how well this clock can maintain a constant frequency.

Sample Output

show ntp status

```

user@host> show ntp status
assID=0 status=0544 leap_none, sync_local_proto, 4 events, event_peer/strat_chg,
version="ntpd 4.2.2p1@1.1570-o Tue May 19 13:57:55 UTC 2009 (1)",
processor="x86_64", system="Linux/2.6.18-164.el5", leap=00, stratum=4,
precision=-10, rootdelay=0.000, rootdispersion=11.974, peer=59475,
refid=LOCAL(0),
reftime=d495c32c.0e71eaf2 Mon, Jan 7 2013 13:57:00.056, poll=10,
clock=d495c32c.cebd43bd Mon, Jan 7 2013 13:57:00.807, state=4,
offset=0.000, frequency=0.000, jitter=0.977, noise=0.977,
stability=0.000, tai=0

```


CHAPTER 30

PTP Operational Command

- `show ptp global-information`

show ptp global-information

Syntax	show ptp global-information
Release Information	Command introduced in Junos OS Release 14.1X53-D25 for the QFX Series.
Description	Show Precision Time Protocol (PTP)—related global information.
Options	This command has no options.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Understanding Transparent Clocks in Precision Time Protocol on page 152 • Configuring Transparent Clock Mode for Precision Time Protocol on page 152
List of Sample Output	show ptp global-information on page 456
Output Fields	Table 43 lists the output fields for the show ptp global-information command. Output fields are listed in the approximate order in which they appear.

Table 43: show ptp global-information Output Fields

Field Name	Field Description
PTP Global Configuration	Displays if PTP is configured globally.
Transparent-clock-config	Displays if transparent clock mode is enabled or disabled.
Transparent-clock-status	Displays if transparent clock mode is active or inactive. If the mode is active, the switch is ready to update the correction field in the PTP packets. If the mode is inactive, the switch is not ready to update the correction field in the PTP packets.

Sample Output

show ptp global-information

```

user@switch> show ptp global-information
PTP Global Configuration:
Transparent-clock-config : ENABLED
Transparent-clock-status  : ACTIVE

```


CHAPTER 31

Routine Monitoring Operational Commands

- `show chassis alarms`
- `show chassis beacon`
- `show chassis environment`
- `show chassis environment fpc`
- `show chassis environment pem`
- `show chassis environment routing-engine`
- `show chassis fabric errors`
- `show chassis fabric fpcs`
- `show chassis fabric plane-location`
- `show chassis fabric sibs`
- `show chassis fabric summary`
- `show chassis fabric topology`
- `show chassis fan`
- `show chassis firmware`
- `show chassis fpc`
- `show chassis hardware`
- `show chassis lcd`
- `show chassis led`
- `show chassis location`
- `show chassis mac-addresses`
- `show chassis pic`
- `show chassis routing-engine`
- `show chassis zones`
- `show host`
- `show log`
- `show subscribers`

- `show system alarms`
- `show system audit`
- `show system boot-messages`
- `show system buffers`
- `show system certificate`
- `show system commit`
- `show system configuration archival`
- `show system configuration rescue`
- `show system connections`
- `show system core-dumps`
- `show system directory-usage`
- `show system processes`
- `show system reboot`
- `show system resource-cleanup processes`
- `show system services service-deployment`
- `show system software`
- `show system statistics`
- `show system storage`
- `show system uptime`
- `show system users`
- `show system virtual-memory`
- `show version`
- `start shell`
- `test configuration`
- `traceroute`
- `traceroute monitor`

show chassis alarms

List of Syntax	Syntax on page 459 Syntax (TX Matrix Routers) on page 459 Syntax (TX Matrix Plus Routers) on page 459 Syntax (MX Series Routers) on page 459 Syntax (MX104, MX2010, and MX2020 3D Universal Edge Routers) on page 459 Syntax (QFX Series) on page 459 Syntax (OCX Series) on page 459 Syntax (PTX Series Packet Transport Routers) on page 459 Syntax (ACX Series Universal Access Routers) on page 459
Syntax	show chassis alarms
Syntax (TX Matrix Routers)	show chassis alarms <lcc <i>number</i> scc>
Syntax (TX Matrix Plus Routers)	show chassis alarms <lcc <i>number</i> sfc <i>number</i> >
Syntax (MX Series Routers)	show chassis alarms <all-members> <local> <member <i>member-id</i> >
Syntax (MX104, MX2010, and MX2020 3D Universal Edge Routers)	show chassis alarms <satellite [slot-id <i>slot-id</i>]>
Syntax (QFX Series)	show chassis alarms <interconnect-device <i>name</i> > <node-device <i>name</i> >
Syntax (OCX Series)	show chassis alarms
Syntax (PTX Series Packet Transport Routers)	show chassis alarms
Syntax (ACX Series Universal Access Routers)	show chassis alarms
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc option for the TX Matrix Plus router introduced in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 12.1 for the PTX Series Packet Transport Routers.

Command introduced in Junos OS Release 12.2 for the ACX Series Universal Access Routers.

Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.

Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.

Command introduced in Junos OS Release 13.2 for MX104 3D Universal Edge Routers.

Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

satellite option introduced in Junos OS Release 14.2R3 for Junos Fusion.

Description Display information about the conditions that have been configured to trigger alarms.

Options **none**—Display information about the conditions that have been configured to trigger alarms.

all-members—(MX Series routers only) (Optional) Display information about alarm conditions for all the member routers of the Virtual Chassis configuration.

interconnect-device *name*—(QFabric systems only) (Optional) Display information about alarm conditions for the Interconnect device.

lcc *number*—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(MX Series routers only) (Optional) Display information about alarm conditions for the local Virtual Chassis member.

member *member-id*—(MX Series routers only) (Optional) Display information about alarm conditions for the specified member of the Virtual Chassis configuration. Replace *member-id* variable with a value of 0 or 1.

node-device *name*—(QFabric systems only) (Optional) Display information about alarm conditions for the Node device.

satellite [*slot-id slot-id*]—(Junos Fusion only) (Optional) Display information about alarm conditions for the specified satellite device in a Junos Fusion, or for all satellite devices in the Junos Fusion if no satellite devices are specified.

scc—(TX Matrix router only) (Optional) Show information about the TX Matrix router (switch-card chassis).

sfc number—(TX Matrix Plus router only) (Optional) Show information about the respective TX Matrix Plus router, which is the switch-fabric chassis. Replace *number* variable with 0.

Additional Information You cannot clear the alarms for chassis components. Instead, you must remedy the cause of the alarm. When a chassis alarm LED is lit, it indicates that you are running the router or switch in a manner that we do not recommend.

On routers, you can manually silence external devices connected to the alarm relay contacts by pressing the alarm cutoff button, located on the craft interface. Silencing the device does not remove the alarm messages from the display (if present on the router) or extinguish the alarm LEDs. In addition, new alarms that occur after you silence an external device reactivate the external device.

In Junos OS release 11.1 and later, alarms for fans also show the slot number of the fans in the CLI output.

In Junos OS Release 11.2 and later, the command output on EX8200 switches shows the detailed location (**Plane/FPC/PFE**) for link errors in the chassis.

In Junos OS Release 10.2 and later, an alarm is shown on T Series routers for a standby sonic clock generator (SCG) that is offline or absent.

You may often see the following error messages, in which only the error code is shown and no other information is provided:

```
Apr 12 08:04:10 send: red alarm set, device FPC 6, reason FPC 6 Major Errors - Error code: 257
Apr 12 08:04:19 send: red alarm set, device FPC 1, reason FPC 1 Major Errors - Error code: 559
```

To understand what CM_ALARM error codes mean, you need to first identify the structure of the CM Alarm codes. A CM_ALARM code has the following structure:

Bits:	Error type:
1-31	Major (1)
0	Minor (0)

According to the table above, the LSB (bit 0) identifies the **Error Type** (major alarm, if the bit is set and minor alarm if the bit is unset). The rest of the bits (1 - 31) identify the actual error code.

Take an example of the following error code, which was logged on a T1600:

```
Apr 12 08:04:10 send: red alarm set, device FPC 1, reason FPC 1 Major Errors - Error code: 559
```

First, you have to convert 559 to binary; that is **1000101111**. The LSB in this case is 1, which means that this is a major alarm. After removing the LSB, you are left with **100010111**,

which is equal to 279 in decimal. This is the actual error code, its meaning can be found from the following list:

Chip Type: L Chip	Code
CMALARM_LCHIP_LOUT_DESRD_PARITY_ERR	1
CMALARM_LCHIP_LOUT_DESRD_UNINIT_ERR	2
CMALARM_LCHIP_LOUT_DESRD_ILLEGALLINK_ERR	3
CMALARM_LCHIP_LOUT_DESRD_ILLEGALSIZE_ERR	4
CMALARM_LCHIP_LOUT_HDRF_TOERR_ERR	5
CMALARM_LCHIP_LOUT_HDRF_PARITY_ERR	6
CMALARM_LCHIP_LOUT_HDRF_UCERR_ERR	7
CMALARM_LCHIP_LOUT_NLIF_CRCDROP_ERR	8
CMALARM_LCHIP_LOUT_NLIF_CRCERR_ERR	9
CMALARM_LCHIP_UCODE_TIMEOUT_ERR	10
CMALARM_LCHIP_LIN_SRCTL_ACCT_DROP_ERR	11
CMALARM_LCHIP_LIN_SRCTL_ACCT_ADDR_SIZE_ERR	12
CMALARM_LCHIP_SRAM_PARITY_ERR	13
CMALARM_LCHIP_UCODE_OVFLW_ERR	14
CMALARM_LCHIP_LOUT_HDRF_MTU_ERR	15
Chip Type: M Chip	Code
CMALARM_MCHIP_ECC_UNCORRECT_ERR	128
Chip Type: N Chip	Code
CMALARM_NCHIP_RDDMA_JBUS_TIMEOUT_ERR	256
CMALARM_NCHIP_RDDMA_FIFO_OVFLW_ERR	257
CMALARM_NCHIP_RDDMA_FIFO_UNFLW_ERR	258
CMALARM_NCHIP_RDDMA_SIZE_ERR	259

CMALARM_NCHIP_RDDMA_JBUS_CRC_ERR	260
CMALARM_NCHIP_WRDMA_PKTR_ERR	261
CMALARM_NCHIP_WRDMA_PKT_CRC_ERR	262
CMALARM_NCHIP_WRDMA_JBUS_TIMEOUT_ERR	263
CMALARM_NCHIP_WRDMA_FIFO_OVFLW_ERR	264
CMALARM_NCHIP_WRDMA_FIFO_UNFLW_ERR	265
CMALARM_NCHIP_WRDMA_PKT_LEN_ERR	266
CMALARM_NCHIP_WRDMA_JBUS_CRC_ERR	267
CMALARM_NCHIP_PKTR_DMA_AGE_ERR	268
CMALARM_NCHIP_PKTR_ICELLSIG_ERR	269
CMALARM_NCHIP_PKTR_FTTL_ERR	270
CMALARM_NCHIP_RODR_OFFSET_OVFLW_ERR	271
CMALARM_NCHIP_PKTR_TMO_CELL_ERR	272
CMALARM_NCHIP_PKTR_TMO_OUTRANGE_ERR	273
CMALARM_NCHIP_PKTR_MD_REQUEST_Q_OVFLW_ERR	274
CMALARM_NCHIP_PKTR_DMA_BUFFER_OVFLW_ERR	275
CMALARM_NCHIP_PKTR_GRT_OVFLW_ERR	276
CMALARM_NCHIP_FRQ_ERR	277
CMALARM_NCHIP_RODR_IN_Q_OVFLW_ERR	278
CMALARM_NCHIP_DBUF_CRC_ERR	279
<hr/>	
Chip Type: R Chip	Code
CMALARM_RCHIP_SRAM_PARITY_ERR	512
<hr/>	
Chip Type: R Chip	Code
CMALARM_ICHIP_WO_DESRD_ID_ERR	601
CMALARM_ICHIP_WO_DESRD_DATA_ERR	602

CMALARM_ICHIP_WO_DESRD_OFLOW_ERR	603
CMALARM_ICHIP_WO_HDRF_UCERR_ERR	604
CMALARM_ICHIP_WO_HDRF_MTUERR_ERR	605
CMALARM_ICHIP_WO_HDRF_PARITY_ERR	606
CMALARM_ICHIP_WO_HDRF_TOERR_ERR	607
CMALARM_ICHIP_WO_IP_CRC_ERR	608
CMALARM_ICHIP_WO_IP_INTER_ERR	609
CMALARM_ICHIP_WI_WAN_TIMEOUT_ERR	625
CMALARM_ICHIP_WI_FAB_TIMEOUT_ERR	626
CMALARM_ICHIP_RLDRAM_BIST_ERR	630
CMALARM_ICHIP_SDRAM_BIST_ERR	631
CMALARM_ICHIP_RLDRAM_PARITY_ERR	632
CMALARM_ICHIP_SDRAM_UNCORRECT_ERR	633
CMALARM_ICHIP_SDRAM_CORRECT_ERR	634
CMALARM_ICHIP_FUSE_DONE_ERR	635

According to the table above, the **279** error code corresponds to **CMALARM_NCHIP_DBUF_CRC_ERR**; this means that new CRC errors were seen on the NCHIP of this particular FPC, which is FPC as per the logs.

If you do not want to convert decimal to binary and vice versa, you may use the following shortcut:

For major alarms, the **Actual Error Code = (Error Code - 1)/2**, where **Error Code** is the code that you get in the log message. For example, if you get the following log:

Apr 12 08:04:10 send: red alarm set, device FPC 6, reason FPC 6 Major Errors - Error code: 257

Actual Error Code = $(257-1)/2 = 128$. Similarly, for minor alarms, Actual Error Code = $(\text{Error Code})/2$

Required Privilege Level view

- Related Documentation**
- *Configuring an RMON Alarm Entry and Its Attributes*
 - *Chassis Conditions That Trigger Alarms*

- List of Sample Output**
- [show chassis alarms \(Alarms Active\) on page 466](#)
 - [show chassis alarms \(No Alarms Active\) on page 466](#)
 - [show chassis alarms \(Fan Tray\) on page 466](#)
 - [show chassis alarms \(MX104 Router\) on page 466](#)
 - [show chassis alarms \(MX2010 Router\) on page 466](#)
 - [show chassis alarms \(MX2020 Router\) on page 466](#)
 - [show chassis alarms \(MX960, MX480, and MX240 Routers showing Major CB Failure\) on page 467](#)
 - [show chassis alarms \(T4000 Router\) on page 467](#)
 - [show chassis alarms \(Unreachable Destinations Present on a T Series Router\) on page 467](#)
 - [show chassis alarms \(FPC Offline Due to Unreachable Destinations on a T Series Router\) on page 467](#)
 - [show chassis alarms \(SCG Absent on a T Series Router\) on page 468](#)
 - [show chassis alarms \(Alarms Active on a TX Matrix Router\) on page 468](#)
 - [show chassis alarms \(TX Matrix Plus router with 3D SIBs\) on page 468](#)
 - [show chassis alarms \(Alarms on a T4000 Router After the enhanced-mode Statement is Enabled\) on page 470](#)
 - [show chassis alarms \(Backup Routing Engine\) on page 470](#)
 - [show chassis alarms \(EX Series Switch\) on page 471](#)
 - [show chassis alarms \(Alarms Active on the QFX Series and OCX Series Switches\) on page 471](#)
 - [show chassis alarms node-device \(Alarms Active on the QFabric System\) on page 471](#)
 - [show chassis alarms \(Alarms Active on the QFabric System\) on page 471](#)
 - [show chassis alarms \(Alarms Active on an EX8200 Switch\) on page 471](#)
 - [show chassis alarms \(Alarms Active on a PTX5000 Packet Transport Router\) on page 472](#)
 - [show chassis alarms \(Mix of PDUs Alarm on a PTX5000 Packet Transport Router with FPC2-PTX-PIA\) on page 472](#)
 - [show chassis alarms \(PDU Converter Failed Alarm on a PTX5000 Packet Transport Router with FPC2-PTX-PIA\) on page 472](#)
 - [show chassis alarms \(No Power for System Alarm on a PTX5000 Packet Transport Router with FPC2-PTX-PIA\) on page 473](#)
 - [show chassis alarms \(Alarms Active on an ACX2000 Universal Access Router\) on page 473](#)
 - [show chassis alarms \(Active Alarm to Indicate Status of the Bad SCB Clock on MX Series\) on page 473](#)
 - [show chassis alarms \(Active Alarms on PTX5000, MX240, MX480, MX960, MX2010, and MX2020 Routers with Smart Disk Error\) on page 473](#)
- Output Fields** [Table 44](#) lists the output fields for the **show chassis alarms** command. Output fields are listed in the approximate order in which they appear.

Table 44: show chassis alarms Output Fields

Field Name	Field Description
Alarm time	Date and time the alarm was first recorded.
Class	Severity class for this alarm: Minor or Major .
Description	Information about the alarm.

Sample Output

show chassis alarms (Alarms Active)

```

user@host> show chassis alarms
3 alarms are currently active
Alarm time      Class  Description
2000-02-07 10:12:22 UTC Major fxp0: ethernet link down
2000-02-07 10:11:54 UTC Minor YELLOW ALARM - PEM 1 Removed
2000-02-07 10:11:03 UTC Minor YELLOW ALARM - Lower Fan Tray Removed

```

show chassis alarms (No Alarms Active)

```

user@host> show chassis alarms
No alarms are currently active

```

show chassis alarms (Fan Tray)

```

user@host> show chassis alarms
4 alarms currently active
Alarm time      Class  Description
2010-11-11 20:27:38 UTC Major Side Fan Tray 7 Failure
2010-11-11 20:27:13 UTC Minor Side Fan Tray 7 Overspeed
2010-11-11 20:27:13 UTC Major Side Fan Tray 5 Failure
2010-11-11 20:27:13 UTC Major Side Fan Tray 0 Failure

```

show chassis alarms (MX104 Router)

```

user@host >show chassis alarms
1 alarms currently active
Alarm time      Class  Description
2013-06-05 14:43:31 IST Minor Backup RE Active

```

show chassis alarms (MX2010 Router)

```

user@host> show chassis alarms
7 alarms currently active
Alarm time      Class  Description
2012-08-07 00:46:06 PDT Major Fan Tray 2 Failure
2012-08-06 18:24:36 PDT Minor Redundant feed missing for PSM 6
2012-08-06 07:41:04 PDT Minor Redundant feed missing for PSM 8
2012-08-04 02:42:06 PDT Minor Redundant feed missing for PSM 5
2012-08-03 21:14:24 PDT Minor Loss of communication with Backup RE
2012-08-03 12:26:03 PDT Minor Redundant feed missing for PSM 4
2012-08-03 10:40:18 PDT Minor Redundant feed missing for PSM 7

```

show chassis alarms (MX2020 Router)

```

user@host> show chassis alarms

```

```

1 alarms currently active
Alarm time Class Description
2012-10-03 12:14:59 PDT Minor Plane 0 not online

```

show chassis alarms (MX960, MX480, and MX240 Routers showing Major CB Failure)

A major CB 0 failure alarm occurs in the event of a bad CB (unknown or mismatched CBs do not trigger this alarm in Junos Release OS 12.3R9 and later). Following GRES or recovery, if the hardware issue persists, the traffic moves to the good CB and continues. If the alarm was triggered by something transient like a power zone budget on GRES, bringing the CB back online can clear the alarm. Otherwise, replace the bad CB. Note that fabric link speed is not impacted by an offline SCB. The alarm might be raised on CB0, CB1, and CB2.

```

user@host> show chassis alarms
6 alarms currently active
Alarm time          Class Description
2014-10-31 16:49:41 EDT Major PEM 3 Not OK
2014-10-31 16:49:41 EDT Major PEM 2 Not OK
2014-10-31 16:49:31 EDT Major CB 0 Failure
2014-10-31 16:49:31 EDT Minor CB 0 Fabric Chip 0 Not Online
2014-10-31 16:49:31 EDT Minor CB 0 Fabric Chip 1 Not Online
2014-10-31 16:49:31 EDT Minor Backup RE Active

```

show chassis alarms (T4000 Router)

```

user@host> show chassis alarms
9 alarms currently active
Alarm time          Class Description
2007-06-02 01:41:10 UTC Minor RE 0 Not Supported
2007-06-02 01:41:10 UTC Minor CB 0 Not Supported
2007-06-02 01:41:10 UTC Minor Mixed Master and Backup RE types
2007-05-30 19:37:33 UTC Major SPMB 1 not online
2007-05-30 19:37:29 UTC Minor Front Bottom Fan Tray Absent
2007-05-30 19:37:13 UTC Major PEM 1 Input Failure
2007-05-30 19:37:13 UTC Major PEM 0 Not OK
2007-05-30 19:37:03 UTC Major PEM 0 Improper for Platform
2007-05-30 19:37:03 UTC Minor Backup RE Active

```

show chassis alarms (Unreachable Destinations Present on a T Series Router)

```

user@host> show chassis alarms
10 alarms currently active
Alarm time          Class Description
2011-08-30 18:43:53 PDT Major FPC 7 has unreachable destinations
2011-08-30 18:43:53 PDT Major FPC 5 has unreachable destinations
2011-08-30 18:43:52 PDT Major FPC 3 has unreachable destinations
2011-08-30 18:43:52 PDT Major FPC 2 has unreachable destinations
2011-08-30 18:43:52 PDT Minor SIB 0 Not Online
2011-08-30 18:43:33 PDT Minor SIB 4 Not Online
2011-08-30 18:43:28 PDT Minor SIB 3 Not Online
2011-08-30 18:43:05 PDT Minor SIB 2 Not Online
2011-08-30 18:43:28 PDT Minor SIB 1 Not Online
2011-08-30 18:43:05 PDT Major PEM 1 Not Ok

```

show chassis alarms (FPC Offline Due to Unreachable Destinations on a T Series Router)

```

user@host> show chassis alarms
10 alarms currently active
Alarm time          Class Description

```

```

2011-08-30 18:43:53 PDT Major FPC 7 offline due to unreachable destinations
2011-08-30 18:43:53 PDT Major FPC 5 offline due to unreachable destinations
2011-08-30 18:43:52 PDT Major FPC 3 offline due to unreachable destinations
2011-08-30 18:43:52 PDT Major FPC 2 offline due to unreachable destinations
2011-08-30 18:43:52 PDT Minor SIB 0 Not Online
2011-08-30 18:43:33 PDT Minor SIB 4 Not Online
2011-08-30 18:43:28 PDT Minor SIB 3 Not Online
2011-08-30 18:43:05 PDT Minor SIB 2 Not Online
2011-08-30 18:43:28 PDT Minor SIB 1 Not Online
2011-08-30 18:43:05 PDT Major PEM 1 Not Ok

```

show chassis alarms (SCG Absent on a T Series Router)

```

user@host> show chassis alarms
4 alarms currently active
Alarm time          Class Description
2011-01-23 21:42:46 PST Major SCG 0 NO EXT CLK MEAS-BKUP SCG ABS

```

show chassis alarms (Alarms Active on a TX Matrix Router)

```

user@host> show chassis alarms
scc-re0:
-----
8 alarms currently active
Alarm time          Class Description
2004-08-05 18:43:53 PDT Minor LCC 0 Minor Errors
2004-08-05 18:43:53 PDT Minor SIB 3 Not Online
2004-08-05 18:43:52 PDT Major SIB 2 Absent
2004-08-05 18:43:52 PDT Major SIB 1 Absent
2004-08-05 18:43:52 PDT Major SIB 0 Absent
2004-08-05 18:43:33 PDT Major LCC 2 Major Errors
2004-08-05 18:43:28 PDT Major LCC 0 Major Errors
2004-08-05 18:43:05 PDT Minor LCC 2 Minor Errors
lcc0-re0:
-----
5 alarms currently active
Alarm time          Class Description
2004-08-05 18:43:53 PDT Minor SIB 3 Not Online
2004-08-05 18:43:49 PDT Major SIB 2 Absent
2004-08-05 18:43:49 PDT Major SIB 1 Absent
2004-08-05 18:43:49 PDT Major SIB 0 Absent
2004-08-05 18:43:28 PDT Major PEM 0 Not OK
lcc2-re0:
-----
5 alarms currently active
Alarm time          Class Description
2004-08-05 18:43:35 PDT Minor SIB 3 Not Online
2004-08-05 18:43:33 PDT Major SIB 2 Absent
2004-08-05 18:43:33 PDT Major SIB 1 Absent
2004-08-05 18:43:33 PDT Major SIB 0 Absent
2004-08-05 18:43:05 PDT Minor PEM 1 Absent

```

show chassis alarms (TX Matrix Plus router with 3D SIBs)

```

user@host> show chassis alarms
sfc0-re0:
-----
Alarm time          Class Description
2014-04-08 14:35:13 IST Minor FPM 0 SFC Config Size Changed
2014-04-08 14:32:58 IST Major Fan Tray Failure
2014-04-08 14:31:53 IST Major SIB F13 6 Fault

```

```

2014-04-08 14:31:43 IST Major SIB F13 11 Fault
2014-04-08 14:31:08 IST Minor Check SIB F13 12 CXP 14 Fbr Cbl
2014-04-08 14:31:08 IST Minor Check SIB F13 12 CXP 8 Fbr Cbl
2014-04-08 14:31:08 IST Minor Check SIB F13 12 CXP 3 Fbr Cbl
2014-04-08 14:31:08 IST Major SIB F13 12 CXP 15 fault
2014-04-08 14:31:08 IST Minor SIB F13 12 CXP 14 LOL
2014-04-08 14:31:08 IST Minor Check SIB F13 12 CXP 14
2014-04-08 14:31:08 IST Major SIB F13 12 CXP 10 fault
2014-04-08 14:31:08 IST Minor SIB F13 12 CXP 8 LOL
2014-04-08 14:31:08 IST Minor Check SIB F13 12 CXP 8
2014-04-08 14:31:08 IST Major SIB F13 12 CXP 7 fault
2014-04-08 14:31:08 IST Major SIB F13 12 CXP 4 fault
2014-04-08 14:31:08 IST Minor SIB F13 12 CXP 3 LOL
2014-04-08 14:31:08 IST Minor Check SIB F13 12 CXP 3
2014-04-08 14:31:08 IST Minor Check SIB F13 6 CXP 14 Fbr Cbl
2014-04-08 14:31:08 IST Minor Check SIB F13 6 CXP 12 Fbr Cbl
2014-04-08 14:31:08 IST Minor Check SIB F13 6 CXP 8 Fbr Cbl
2014-04-08 14:31:08 IST Minor Check SIB F13 6 CXP 6 Fbr Cbl
2014-04-08 14:31:08 IST Minor Check SIB F13 6 CXP 4 Fbr Cbl
2014-04-08 14:31:08 IST Minor Check SIB F13 6 CXP 2 Fbr Cbl
2014-04-08 14:31:08 IST Minor Check SIB F13 6 CXP 0 Fbr Cbl
2014-04-08 14:31:08 IST Minor SIB F13 6 CXP 14 LOL
2014-04-08 14:31:08 IST Minor Check SIB F13 6 CXP 14
2014-04-08 14:31:08 IST Minor SIB F13 6 CXP 12 LOL
2014-04-08 14:31:08 IST Minor Check SIB F13 6 CXP 12
2014-04-08 14:31:08 IST Major SIB F13 6 CXP 10 fault
2014-04-08 14:31:08 IST Minor SIB F13 6 CXP 8 LOL
2014-04-08 14:31:08 IST Minor Check SIB F13 6 CXP 8
2014-04-08 14:31:08 IST Minor SIB F13 6 CXP 6 LOL
2014-04-08 14:31:08 IST Minor Check SIB F13 6 CXP 6
2014-04-08 14:31:08 IST Minor SIB F13 6 CXP 4 LOL
2014-04-08 14:31:08 IST Minor Check SIB F13 6 CXP 4
2014-04-08 14:31:08 IST Minor SIB F13 6 CXP 2 LOL
2014-04-08 14:31:08 IST Minor Check SIB F13 6 CXP 2
2014-04-08 14:31:08 IST Minor SIB F13 6 CXP 0 LOL
2014-04-08 14:31:08 IST Minor Check SIB F13 6 CXP 0
2014-04-08 14:31:08 IST Minor SIB F13 12 CXP 14 XC HSL Link Error
2014-04-08 14:29:27 IST Minor LCC 0 Minor Errors
2014-04-08 14:28:37 IST Major LCC 0 Major Errors
2014-04-08 14:28:37 IST Major LCC 2 Major Errors
2014-04-08 14:28:37 IST Minor LCC 2 Minor Errors
2014-04-08 14:28:24 IST Major SIB F2S 4/6 Absent
2014-04-08 14:28:24 IST Major SIB F2S 4/4 Absent
2014-04-08 14:28:24 IST Major SIB F2S 4/2 Absent
2014-04-08 14:28:24 IST Major SIB F2S 4/0 Absent
2014-04-08 14:28:24 IST Major SIB F2S 3/6 Absent
2014-04-08 14:28:24 IST Major SIB F2S 3/4 Absent
2014-04-08 14:28:24 IST Major SIB F2S 3/2 Absent
2014-04-08 14:28:24 IST Major SIB F2S 3/0 Absent
2014-04-08 14:28:24 IST Major SIB F13 9 Absent
2014-04-08 14:28:24 IST Major SIB F13 8 Absent
2014-04-08 14:28:24 IST Major SIB F13 7 Absent
2014-04-08 14:28:24 IST Major SIB F13 4 Absent
2014-04-08 14:28:24 IST Major SIB F13 1 Absent
2014-04-08 14:28:22 IST Major PEM 0 Input Failure
2014-04-08 14:28:22 IST Major PEM 0 Not OK

```

```
lcc0-re0:
```

```
-----
12 alarms currently active
```

```
Alarm time          Class  Description
```

```

2014-04-08 14:36:08 IST Minor CB 1 M/S Switch Changed
2014-04-08 14:36:08 IST Minor CB 1 CHASSIS ID Changed
2014-04-08 14:35:43 IST Minor CB 0 M/S Switch Changed
2014-04-08 14:35:43 IST Minor CB 0 CHASSIS ID Changed
2014-04-08 14:29:30 IST Minor SIB 4 Not Online
2014-04-08 14:29:30 IST Minor SIB 3 Not Online
2014-04-08 14:29:30 IST Minor SIB 2 Not Online
2014-04-08 14:29:24 IST Major Rear Fan Tray Failure
2014-04-08 14:29:24 IST Major Front Bottom Fan Tray Improper for Platform
2014-04-08 14:29:24 IST Major Front Top Fan Tray Improper for Platform
2014-04-08 14:28:37 IST Major SIB 4 Absent
2014-04-08 14:28:37 IST Major SIB 3 Absent

```

```
lcc2-re0:
```

```

-----
12 alarms currently active
Alarm time      Class Description
2014-04-08 14:36:02 IST Minor CB 1 M/S Switch Changed
2014-04-08 14:36:02 IST Minor CB 1 CHASSIS ID Changed
2014-04-08 14:35:42 IST Minor CB 0 M/S Switch Changed
2014-04-08 14:34:42 IST Minor CB 0 CHASSIS ID Changed
2014-04-08 14:29:29 IST Minor SIB 0 CXP 7 Unsupported Optics
2014-04-08 14:29:27 IST Major Front Bottom Fan Tray Improper for Platform
2014-04-08 14:29:27 IST Major Front Top Fan Tray Improper for Platform
2014-04-08 14:29:25 IST Minor SIB 4 Not Online
2014-04-08 14:29:25 IST Minor SIB 3 Not Online
2014-04-08 14:28:47 IST Major PEM 0 Not OK
2014-04-08 14:28:36 IST Major SIB 2 Absent
2014-04-08 14:28:36 IST Minor Host 0 Boot from alternate media

```

```
lcc6-re0:
```

```

-----
2 alarms currently active
Alarm time      Class Description
2013-11-06 04:03:56 PST Minor SIB 1 CXP 0 XC HSL Link Error
2013-11-06 03:49:32 PST Major PEM 1 Not OK

```

show chassis alarms (Alarms on a T4000 Router After the enhanced-mode Statement is Enabled)

To enable improved virtual private LAN service (VPLS) MAC address learning on T4000 routers, you must include the **enhanced-mode** statement at the **[edit chassis network-services]** hierarchy level and reboot the router. When router reboots, only the T4000 Type 5 FPCs are required to be present on the router. If there are any other FPCs (apart from T4000 Type 5 FPCs) on the T4000 router, such FPCs become offline, and FPC misconfiguration alarms are generated. The **show chassis alarm** command output displays FPC misconfiguration (**FPC *fpc-slot* misconfig**) as the reason for the generation of the alarms.

```

user@host> show chassis alarms
2 alarms currently active
Alarm time      Class Description
2011-10-22 10:10:47 PDT Major FPC 1 misconfig
2011-10-22 10:10:46 PDT Major FPC 0 misconfig

```

show chassis alarms (Backup Routing Engine)

```

user@host> show chassis alarms
2 alarms are currently active
Alarm time      Class Description

```

```

2005-04-07 10:12:22 PDT Minor Host 1 Boot from alternate media
2005-04-07 10:11:54 PDT Major Host 1 compact-flash missing in Boot List

```

show chassis alarms (EX Series Switch)

```

user@switch> show chassis alarms
4 alarms currently active
Alarm time      Class Description
2014-03-12 15:36:09 UTC Minor Require a Fan Tray upgrade
2014-03-12 15:00:02 UTC Major PEM 0 Input Failure
2014-03-12 15:00:02 UTC Major PEM 0 Not OK
2014-03-12 14:59:51 UTC Minor Host 1 Boot from alternate media

```

show chassis alarms (Alarms Active on the QFX Series and OCX Series Switches)

```

user@switch> show chassis alarms
1 alarms currently active
Alarm time      Class Description
2012-03-05 2:10:24 UTC Major FPC 0 PEM 0 Airflow not matching Chassis Airflow

```

show chassis alarms node-device (Alarms Active on the QFabric System)

```

user@switch> show chassis alarms node-device ED3691
node-device ED3694
3 alarms currently active
Alarm time      Class Description
2011-08-24 16:04:15 UTC Major ED3694:fte-0/1/2: Link down
2011-08-24 16:04:14 UTC Major ED3694:fte-0/1/0: Link down
2011-08-24 14:21:14 UTC Major ED3694 PEM 0 is not supported/powered

```

show chassis alarms (Alarms Active on the QFabric System)

```

user@switch> show chassis alarms
IC-A0001:
-----
1 alarms currently active
Alarm time      Class Description
2011-08-24 16:04:15 UTC Minor Backup RE Active

ED3694:
-----
3 alarms currently active
Alarm time      Class Description
2011-08-24 16:04:15 UTC Major ED3694:fte-0/1/2: Link down
2011-08-24 16:04:14 UTC Major ED3694:fte-0/1/0: Link down
2011-08-24 14:21:14 UTC Major ED3694 PEM 0 is not supported/powered

SNG-0:
-----

NW-NG-0:
-----
1 alarms currently active
Alarm time      Class Description
2011-08-24 15:49:27 UTC Major ED3691 PEM 0 is not supported/powered

```

show chassis alarms (Alarms Active on an EX8200 Switch)

```

user@switch> show chassis alarms

6 alarms currently active

```

Alarm time	Class	Description
2010-12-02 19:15:22 UTC	Major	Fan Tray Failure
2010-12-02 19:15:22 UTC	Major	Fan Tray Failure
2010-12-02 19:15:14 UTC	Minor	Check CB 0 Fabric Chip 1 on Plane/FPC/PFE: 1/5/0, 1/5/1, 1/5/2, 1/5/3, 1/7/0, 1/7/1, 1/7/2, 1/7/3, 2/5/0, 2/5/1, ...
2010-12-02 19:15:14 UTC	Minor	Check CB 0 Fabric Chip 0 on Plane/FPC/PFE: 1/5/0, 1/5/1, 1/5/2, 1/5/3, 1/7/0, 1/7/1, 1/7/2, 1/7/3, 2/5/0, 2/5/1, ...
2010-12-02 19:14:18 UTC	Major	PSU 1 Output Failure
2010-12-02 19:14:18 UTC	Minor	Loss of communication with Backup RE

show chassis alarms (Alarms Active on a PTX5000 Packet Transport Router)

```
user@host> show chassis alarms
```

```
23 alarms currently active
Alarm time      Class  Description
2011-07-12 16:22:05 PDT  Minor  No Redundant Power for Rear Chassis
2011-07-12 16:22:05 PDT  Major  PDU 0 PSM 1 Not OK
2011-07-12 16:21:57 PDT  Minor  No Redundant Power for Fan 0-2
2011-07-12 16:21:57 PDT  Major  PDU 0 PSM 0 Not OK
2011-07-12 15:56:06 PDT  Major  PDU 1 PSM 2 Not OK
2011-07-12 15:56:06 PDT  Minor  No Redundant Power for FPC 0-7
2011-07-12 15:56:06 PDT  Major  PDU 0 PSM 3 Not OK
2011-07-12 15:28:20 PDT  Major  PDU 0 PSM 2 Not OK
2011-07-12 15:19:14 PDT  Minor  Backup RE Active
```

show chassis alarms (Mix of PDUs Alarm on a PTX5000 Packet Transport Router with FPC2-PTX-P1A)

All PDUs installed on a PTX5000 router must be of the same type. The **Mix of PDUs** or **Power Manager Non Operational** alarm is raised when different types of PDUs are installed on a PTX5000 router.

```
user@host> show chassis alarms
15 alarms currently active
Alarm time      Class  Description
2013-03-19 23:03:53 PDT  Minor  No Redundant Power
2013-03-19 23:03:48 PDT  Minor  Mix of PDUs
2013-03-19 23:03:47 PDT  Minor  PDU 1 PSM 3 Absent
2013-03-19 23:03:47 PDT  Minor  PDU 1 PSM 2 Absent
2013-03-19 23:03:47 PDT  Minor  PDU 1 PSM 1 Absent
2013-03-19 23:03:47 PDT  Minor  PDU 1 PSM 0 Absent
2013-03-19 23:03:46 PDT  Major  No CG Online
```

show chassis alarms (PDU Converter Failed Alarm on a PTX5000 Packet Transport Router with FPC2-PTX-P1A)

The **PDU Converter Failed** alarm is raised when one or more 36 V booster converter of a DC PDU fails. If two or more 36 V booster converter fails, fan trays fail and the router might get over heated. Therefore, when this alarm is raised, check the PDU and replace it, if required.

```
user@host> show chassis alarms
11 alarms currently active
Alarm time      Class  Description
2013-12-11 22:14:13 PST  Minor  No Redundant Power for System
2013-12-11 22:14:10 PST  Major  PDU 0 PSM 7 Not OK
2013-12-11 22:14:10 PST  Major  PDU 0 PSM 6 Not OK
2013-12-11 22:14:10 PST  Major  PDU 0 PSM 5 Not OK
2013-12-11 22:14:10 PST  Major  PDU 0 PSM 4 Not OK
2013-12-11 22:14:10 PST  Major  PDU 0 PSM 3 Not OK
2013-12-11 22:14:10 PST  Major  PDU 0 PSM 2 Not OK
```



```

2013-12-11 22:14:10 PST Major PDU 0 PSM 1 Not OK
2013-12-11 22:14:10 PST Major PDU 0 PSM 0 Not OK
2013-12-11 22:14:10 PST Major PDU 0 Not OK
2013-12-11 22:14:01 PST Major PDU 0 Converter Failed

```

show chassis alarms (No Power for System Alarm on a PTX5000 Packet Transport Router with FPC2-PTX-P1A)

```

user@host> show chassis alarms
8 alarms currently active
Alarm time          Class Description
2013-11-19 01:58:41 PST Major No Power for System
2013-11-19 01:58:37 PST Major PDU 0 PSM 1 Not OK
2013-11-19 01:56:46 PST Major PDU 0 PSM 2 Not OK
2013-11-19 01:54:26 PST Major PDU 0 PSM 3 Not OK
2013-11-19 01:53:30 PST Major PDU 1 PSM 3 Not OK
2013-11-19 01:53:29 PST Major PDU 1 PSM 2 Not OK
2013-11-19 01:53:29 PST Major PDU 1 PSM 1 Not OK
2013-11-19 01:53:29 PST Major PDU 1 PSM 0 Not OK

```

show chassis alarms (Alarms Active on an ACX2000 Universal Access Router)

```

user@host> show chassis alarms
7 alarms currently active
Alarm time          Class Description
2012-05-22 11:19:09 UTC Major xe-0/3/1: Link down
2012-05-22 11:19:09 UTC Major xe-0/3/0: Link down
2012-05-22 11:19:09 UTC Major ge-0/1/7: Link down
2012-05-22 11:19:09 UTC Major ge-0/1/6: Link down
2012-05-22 11:19:09 UTC Major ge-0/1/3: Link down
2012-05-22 11:19:09 UTC Major ge-0/1/2: Link down
2012-05-22 11:19:09 UTC Major ge-0/1/1: Link down

```

show chassis alarms (Active Alarm to Indicate Status of the Bad SCB Clock on MX Series)

```

user@host> show chassis alarms
1 alarm currently active
Alarm time          Class Description
2013-08-06 07:48:35 PDT Major CB 0 19.44 MHz clock failure

```

show chassis alarms (Active Alarms on PTX5000, MX240, MX480, MX960, MX2010, and MX2020 Routers with Smart Disk Error)

```

user@host> show chassis alarms
4 alarms currently active
Alarm time          Class Description
2016-01-11 16:02:10 UTC MINOR Host 0 disk drive 2 smart error
2016-01-11 16:02:10 UTC MINOR Host 0 disk drive 1 smart error
2016-01-11 16:02:05 UTC MINOR Host 1 disk drive 2 smart error
2016-01-11 16:02:05 UTC MINOR Host 1 disk drive 1 smart error

```

show chassis beacon

show chassis beacon (QFX Series)	<pre>show chassis beacon <cb slot-number> <fpc slot-number> <interconnect-device name (cb slot-number fpc slot-number)> <node-device name></pre>
Release Information	<p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Display the beacon LED status on a QFX3500, QFX3600, QFX5100, EX4600, OCX Series standalone switch, Node device, and an Interconnect device. You can also display the beacon LED status of the Control Boards and Flexible PIC Concentrators on the Interconnect device.
Options	<p>cb slot-number— (QFabric systems only) (Optional) Display the status of the beacon LEDs for the Control Board on the Interconnect device.</p> <p>fpc slot-number— (QFabric systems only) (Optional) Display the status of the beacon LEDs for the Flexible PIC Concentrator (FPC) on the Interconnect device. (Optional) Display the status of the beacon LEDs for the Flexible PIC Concentrator on the standalone switch.</p> <p>interconnect-device name— (QFabric systems only) (Optional) Display the status of the beacon LEDs for the Interconnect device.</p> <p>node-device name— (QFabric systems only) (Optional) Display the status of the beacon LEDs for the Node device.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> request chassis beacon on page 364
List of Sample Output	<p>show chassis beacon (QFX Series and OCX Series) on page 475</p> <p>show chassis beacon interconnect-device (QFabric System) on page 475</p> <p>show chassis beacon interconnect-device fpc (QFabric System) on page 475</p> <p>show chassis beacon node-device (QFabric System) on page 475</p> <p>show chassis beacon node-device fpc (QFabric System) on page 475</p>
Output Fields	Table 45 lists the output fields for the show chassis beacon command. Output fields are listed in the approximate order in which they appear.

Table 45: show chassis led Output Fields

Field Name	Field Description
Slot	FPC slot number of the device whose content is being displayed. On QFX3500 standalone switches, the number is always 0.

Table 45: show chassis led Output Fields (*continued*)

Field Name	Field Description
Beacon State	Status of the beacon state: <ul style="list-style-type: none"> • Off—The beacon is OFF. • On—The beacon is ON.

Sample Output

show chassis beacon (QFX Series and OCX Series)

```

user@switch> show chassis beacon
Slot          Beacon State
FPC           0          OFF

```

show chassis beacon interconnect-device (QFabric System)

```

user@switch> show chassis beacon interconnect-device interconnect1
Chassis              OFF
CB 0                  OFF
CB 1                  OFF
FC 0 FPC 0           OFF
FC 1 FPC 1           OFF
RC 0 FPC 8           OFF
RC 1 FPC 9           OFF

```

show chassis beacon interconnect-device fpc (QFabric System)

```

user@switch> show chassis beacon interconnect-device interconnect1 fpc 0
FPC 0                ON

```

show chassis beacon node-device (QFabric System)

```

user@switch> show chassis beacon node-device node1
node1                ON

```

show chassis beacon node-device fpc (QFabric System)

```

user@switch> show chassis beacon node-device node1 fpc 0
FPC 0                ON

```

show chassis environment

- List of Syntax**
- Syntax on page 476
 - Syntax (T320, T640, T1600, and T4000 Routers) on page 476
 - Syntax (TX Matrix Routers) on page 476
 - Syntax (TX Matrix Plus Routers) on page 476
 - Syntax (MX Series Routers) on page 476
 - Syntax (MX104 3D Universal Edge Routers) on page 476
 - Syntax (MX2010 and MX2020 3D Universal Edge Routers) on page 477
 - Syntax (EX8200 Switches) on page 477
 - Syntax (EX Series Switches except EX8200) on page 477
 - Syntax (QFX Series) on page 477
 - Syntax (OCX Series) on page 477
 - Syntax (PTX Series Packet Transport Routers) on page 477
 - Syntax (ACX Series Universal Access Routers) on page 478

Syntax show chassis environment

Syntax (T320, T640, T1600, and T4000 Routers)

```
show chassis environment
<cb cb-slot-number>
<fpc fpc-slot-number>
<fpm>
<pem pem-slot-number>
<routing-engine re-slot-number>
<scg scg-slot-number>
<sib sib-slot-number>
```

Syntax (TX Matrix Routers)

```
show chassis environment
<lcc number | scc>
```

Syntax (TX Matrix Plus Routers)

```
show chassis environment
<cb cb-slot-number>
<cip cip-slot-number>
<fpc fpc-slot-number>
<fpm>
<lcc number>
<pem pem-slot-number>
<routing-engine re-slot-number>
<scg scg-slot-number>
<sfc number>
<sib sib-slot-number>
```

Syntax (MX Series Routers)

```
show chassis environment
<all-members>
<local>
<member member-id>
```

Syntax (MX104 3D Universal Edge Routers)

```
show chassis environment
<cb>
<pem pem-slot-number>
<routing-engine re-slot-number>
```

Syntax (MX2010 and MX2020 3D Universal Edge Routers)	<pre> show chassis environment <adc <i>adc-slot-number</i>> <all-members> <cb <i>cb-slot-number</i>> <fan <i>fantray-slot-number</i>> <fpc <i>fpc-slot-number</i>> <fpm> <local> <member <i>member-id</i>> <monitored> <psm <i>psm-slot-number</i>> <routing-engine <i>re-slot-number</i>> <sfb <i>sfb-slot-number</i>> <satellite [<i>slot-id slot-id</i> device-alias <i>alias-name</i>]> </pre>
Syntax (EX8200 Switches)	<pre> show chassis environment <all-members> <cb <i>cb-slot-number</i>> <fpc <i>fpc-slot-number</i>> <local> <member <i>member-id</i>> <psu <i>psu-slot-number</i>> <routing-engine <i>re-slot-number</i>> </pre>
Syntax (EX Series Switches except EX8200)	<pre> show chassis environment <all-members> <fpc <i>fpc-slot-number</i>> <local> <member <i>member-id</i>> <power-supply-unit> <routing-engine> <satellite [<i>slot-id slot-id</i> device-alias <i>alias-name</i>]> </pre>
Syntax (QFX Series)	<pre> show chassis environment <cb <i>slot-number</i> <interconnect-device <i>name</i>>> <fpc <i>slot-number</i> <interconnect-device <i>name</i>>> <interconnect-device <i>name</i> <slot-number> <node-device <i>name</i>> <pem <i>slot-number</i> (interconnect-device <i>name slot-number</i>) (node-device <i>name</i>)> <routing-engine <i>name</i> <interconnect-device <i>name slot-number</i>>> </pre>
Syntax (OCX Series)	<pre> show chassis environment </pre>
Syntax (PTX Series Packet Transport Routers)	<pre> show chassis environment <cb <i>cb-slot-number</i>> <ccg <i>ccg-slot-number</i>> <fpc <i>fpc-slot-number</i>> <fpm> <monitored> <pdu <i>pdu-slot-number</i>> <routing-engine <i>re-slot-number</i>> <sib <i>sib-slot-number</i>> </pre>

Syntax (ACX Series Universal Access Routers) `show chassis environment`
`<cb cb-slot-number>`
`<pem pem-slot-number>`
`<routing-engine re-slot-number>`

Release Information Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.
Command introduced in Junos OS Release 11.1 for QFX Series.
Command introduced in Junos OS Release 12.1x48 for PTX Series Packet Transport Routers.
monitored option added in Junos OS Release 12.1x48 for PTX Series Packet Transport Routers.
Command introduced in Junos OS Release 12.1 for T4000 Core Routers.
Command introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.
Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.
Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.
pem option introduced in Junos OS Release 12.3 for ACX4000 Universal Access Routers.
Command introduced in Junos OS Release 13.2 for MX104 3D Universal Edge Routers.
all-members, **local**, and **member *member-id*** options introduced in Junos OS Release 15.1 for MX2020 and MX2010 routers.
satellite option introduced in Junos OS Release 14.2R3.

Description Display environmental information about the router or switch chassis, including the temperature and information about the fans, power supplies, and Routing Engine.

In addition, on ACX4000 routers, display temperature information about the different channels of a Modular Interface Card (MIC). The number of channels displayed depends on the type of MIC installed.

Starting with Junos OS Release 14.1, the **show chassis environment *cb cb-slot-number* | *ccg ccg-slot-number* | *fpc fpc-slot-number* | *fpm* | *monitored* | *pdu pdu-slot-number* | *routing-engine re-slot-number* | *sib sib-slot-number*** operational mode command output displays environmental information for the new DC power supply module (PSM) and power distribution unit (PDU) that are added to provide power to the high-density FPC (FPC2-PTX-PIA) and other components in a PTX5000 Packet Transport Router.

Options **none**—Display environmental information about the router or switch chassis. On a TX Matrix router, display environmental information about the TX Matrix router and its attached T640 routers. On a TX Matrix Plus router, display environmental information about the TX Matrix Plus router and its attached routers.

all-members—(MX Series routers and EX Series switches only) (Optional) Display chassis environmental information for all the members of the Virtual Chassis configuration.

adc *adc-slot-number*—(MX2020 and MX2010 routers only) (Optional) Display chassis environmental information for the adapter cards. For MX2020 routers, replace ***adc-slot-number*** with a value from 0 through 19. For MX2010 routers, replace ***adc-slot-number*** with a value from 0 through 9.

cb *cb-slot-number*—(ACX Series Universal Access Routers, EX Series switches, M120, M320, and M40e routers, MX Series routers, MX2020 routers, MX2010 routers, PTX Series Packet Transport Routers, QFX Series, and T Series routers, and TX Matrix Plus routers only) (Optional) Display chassis environmental information for the Control Board. On devices other than EX Series switches, replace ***cb-slot*** with **0** or **1**. For the EX Series switches, see *EX Series Switches Hardware and CLI Terminology Mapping* for information on CB slot numbering.

cip *cip-slot-number*—(TX Matrix Plus routers only) (Optional) Display chassis environmental information for the Connection Interface Panel (CIP). Replace the ***cip-slot-number*** variable with a value of **0** or **1**.

cb *interconnect-device name*—(QFabric systems only) (Optional) Display chassis environmental information for the Control Board on an Interconnect device.

ccg *ccg-slot-number*—(PTX Series only) (Optional) Display chassis environmental information for the Centralized Clock Generator. Replace ***cb-slot*** with a value of **0** or **1**.

fan *fantray-slot-number*—(MX2020 and MX2010 routers only) (Optional) Display chassis environmental information for the fan trays. Replace ***fantray-slot-number*** with a value from **0** through **3**.

fpc *fpc-slot*—(EX Series switches, M120, M320, and M40e routers, MX Series routers, MX2010 routers, MX2020 routers, PTX Series Packet Transport Routers, QFX Series, QFX3500 switches, QFabric systems, T Series routers, and TX Matrix Plus routers) (Optional) Display chassis environmental information for a specified Flexible PIC Concentrator. For MX2010 routers, replace ***fpc-slot*** with a value from **0** through **9**. For MX2020 routers, replace ***fpc-slot*** with a value from **0** through **19**. For information about FPC numbering, see [show chassis environment fpc](#). On a QFabric system, display chassis environmental information for a specified Flexible PIC Concentrator on an Interconnect device. On an EX Series switch, display chassis environmental information for a specified Flexible PIC Concentrator; see *EX Series Switches Hardware and CLI Terminology Mapping* for information on FPC numbering. On a TX Matrix Plus router with 3D SIBs replace ***fpc-slot*** with a value from **0** through **63**.

fpm—(M120, M320, and M40e routers, MX2010 routers, MX2020 routers, PTX Series, Packet Transport Routers, T Series routers, and TX Matrix Plus routers only) (Optional) Display chassis environmental information for the craft interface (FPM).

interconnect-device *name*—(QFabric systems only) (Optional) Display chassis environmental information for the Interconnect device.

lcc *number*—(TX Matrix routers and TX Matrix Plus routers only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.

- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(MX Series routers and EX Series switches only) (Optional) Display chassis environmental information for the local Virtual Chassis member.

member *member-id*—(MX Series routers and EX Series switches only) (Optional) Display chassis environmental information for the specified member of the Virtual Chassis configuration. On MX Series routers, replace *member-id* with a value of 0 or 1. For EX Series switches, see *member* for member ID values.

monitored—(MX2020 routers and PTX Series Packet Transport Routers only) (Optional) Display chassis environmental information for monitored temperatures only. Temperatures that are not included in temperature alarm computations are not displayed.

node-device *name*—(QFabric systems only) (Optional) Display chassis environmental information for the Node device.

pdu *pdu-slot-number*—(PTX Series only) (Optional) Display chassis environmental information for the specified power distribution unit.

pem—(QFX3500 switches and QFabric systems only) (Optional) Display chassis environmental information for the Power Entry Module on the specified Interconnect device or Node device.

pem *pem-slot-number*—(ACX Series Universal Access Routers, M120, M320, and M40e routers, MX Series routers, MX104 routers, QFX Series, and T Series routers only) (Optional) Display chassis environmental information for the Power Entry Module on the specified Power Entry Module. For information about the options, see [show chassis environment pem](#).

psm *psm-slot-number*—(MX2020 and MX2010 routers only) (Optional) Display chassis environmental information for the power supply module. For MX2020 routers, replace *psm-slot-number* with a value from 0 through 17. For MX2010 routers, replace *psm-slot-number* with a value from 0 through 8.

psu *psu-slot-number*—(EX Series switches only) (Optional) Display chassis environmental information for a specified power supply. See *EX Series Switches Hardware and CLI Terminology Mapping* for detailed information.

routing-engine—(QFX3500 switches and QFabric systems only) (Optional) Display chassis environmental information for the Routing Engine on the specified Interconnect device.

routing-engine *re-slot-number*—(Optional) Display chassis environmental information for the specified Routing Engine. For information about the options, see [show chassis environment routing-engine](#).

satellite [**slot-id** *slot-id* | **device-alias** *alias-name*]**—**(Junos Fusion only) (Optional) Display chassis environmental information for the specified satellite device in a Junos Fusion, or for all satellite devices in the Junos Fusion if no satellite devices are specified.

scg**—**(T Series routers only) (Optional) Display chassis environmental information about the SONET Clock Generator.

scc**—**(TX Matrix routers only) (Optional) Display chassis environmental information about the TX Matrix router (switch-card chassis).

sfb **sfb-slot-number****—**(MX2020 and MX2010 routers only) (Optional) Display chassis environmental information for the power supply module. Replace **sfb-slot-number** with a value from 0 through 7.

sfc **number****—**(TX Matrix Plus routers only) (Optional) Display chassis environmental information about the respective TX Matrix Plus router (switch-fabric chassis). Replace **number** variable with 0.

sib **sib-slot-number****—**(M320 routers, PTX Series Packet Transport Routers, and T Series routers only) (Optional) Display chassis environmental information about the specified switch interface board. For information about the options, see *show chassis environment sib*.

Required Privilege Level

view

Related Documentation

- *show chassis environment adc*
- *show chassis environment cb*
- *show chassis environment ccg*
- *show chassis environment cip*
- [show chassis environment fpc on page 540](#)
- *show chassis environment fpm*
- *show chassis environment lcc*
- *show chassis environment mcs*
- *show chassis environment monitored*
- *show chassis environment pcg*
- *show chassis environment pdu*
- [show chassis environment pem on page 566](#)
- *show chassis environment psm*
- *show chassis environment psu*
- [show chassis environment routing-engine on page 575](#)
- *show chassis environment scg*

- *show chassis environment sfb*
- *show chassis environment sib*
- *show chassis environment sfc*

List of Sample Output	show chassis environment (M5 Router) on page 484
	show chassis environment (M7i Router) on page 484
	show chassis environment (M10 Router) on page 485
	show chassis environment (M10i Router) on page 485
	show chassis environment (M20 Router) on page 485
	show chassis environment (M40 Router) on page 486
	show chassis environment (M40e Router) on page 486
	show chassis environment (M120 Router) on page 487
	show chassis environment (M160 Router) on page 488
	show chassis environment (M320 Router) on page 488
	show chassis environment (MX104 Router) on page 489
	show chassis environment (MX240 Router) on page 489
	show chassis environment (MX240 Router with SCBE) on page 490
	show chassis environment (MX480 Router) on page 491
	show chassis environment (MX480 Router with SCBE) on page 492
	show chassis environment (MX960 Router) on page 493
	show chassis environment (MX960 Router with SCBE) on page 494
	show chassis environment (MX960 Router with MPC5EQ) on page 496
	show chassis environment (MX2020 Router) on page 501
	show chassis environment (MX2020 Router with MPC5EQ and MPC6E) on page 510
	show chassis environment (MX2010 Router) on page 514
	show chassis environment (T320 Router) on page 519
	show chassis environment (T640 Router) on page 520
	show chassis environment (T4000 Router) on page 521
	show chassis environment (TX Matrix Router) on page 522
	show chassis environment (T1600 Router) on page 524
	show chassis environment (TX Matrix Plus Router) on page 525
	show chassis environment (TX Matrix Plus router with 3D SIBs) on page 527
	show chassis environment (EX4200 Standalone Switch) on page 530
	show chassis environment (EX8216 Switch) on page 530
	show chassis environment (EX9200 Switch) on page 531
	show chassis environment (QFX Series and OCX Series) on page 532
	show chassis environment interconnect-device (QFabric System) on page 532
	show chassis environment node-device (QFabric System) on page 534
	show chassis environment pem node-device (QFabric System) on page 534
	show chassis environment (PTX5000 Packet Transport Router) on page 534
	show chassis environment (PTX5000 Packet Transport Router with FPC2-PTX-P1A) on page 537
	show chassis environment (PTX1000 Packet Transport Router) on page 537
	show chassis environment (ACX2000 Universal Access Router) on page 538
	show chassis environment (ACX4000 Universal Access Router) on page 538

Output Fields [Table 46](#) lists the output fields for the **show chassis environment** command. Output fields are listed in the approximate order in which they appear.

Table 46: show chassis environment Output Fields

Field Name	Field Description
Class	<p>Information about the category or class of chassis component:</p> <ul style="list-style-type: none"> • Power: Power information: <ul style="list-style-type: none"> • (M5, M10, M20, and M40 routers and EX Series switches only) Power supply status: OK, Testing, (during initial power-on), Failed, or Absent. • (M7i, M10i, M40e, M120, M160, M320, and T Series routers and EX Series switches only) Power Entry Modules status: OK, Testing, (during initial power-on), Check, Failed, or Absent. • (PTX Series only) Power information is reported in PDU or PSM combinations. The status is: OK, Testing, (during initial power-on), Check, Failed, or Absent. • Temp: Temperature of air flowing through the chassis in degrees Celsius (C) and Fahrenheit (F). <ul style="list-style-type: none"> • On PTX Series Packet Transport Routers and MX2010 and MX2020 Routers, multiple cooling zones are supported. FRU temperatures in each zone are coordinated with the fan speed of fan trays in those zones. • EX2200 switches have a side-to-rear cooling system. The Local Intake temperature is measured by the sensor on the right side of the chassis, and the Remote Intake temperature is measured by the sensor on the left side of the chassis. • Pic: On ACX4000 Routers, multiple temperature channels on a MIC. The status is: OK and the Measurement is in degrees Celsius (C) and Fahrenheit (F). • Fan: Fan status: OK, Testing (during initial power-on), Failed, or Absent. On PTX Series Packet Transport Routers and MX2010 and MX2020 Routers, multiple fan trays are supported. Fan status is reported in Fan Tray or Fan combinations. Measurement indicates actual fan RPM (PTX and MX2010 and MX2020 Routers only). • Misc: Information about other components of the chassis. <ul style="list-style-type: none"> • On some routers, this field indicates the status of one or more additional components. • On the M40e, M160, and M320 router, Misc includes CIP (Connector Interface Panel). OK indicates that the CIP is present. Absent indicates that the CIP is not present. • On T Series routers, Misc includes CIP and SPMB (Switch Processor Mezzanine Board). OK indicates that the CIP or SPMB is present. Absent indicates that the CIP or SPMB is not present. • On PTX Series Packet Transport Routers, Misc includes the SPMB (Switch Processor Mezzanine Board). The SPMB is located on the control boards. OK indicates that the control board is present. Absent indicates that the control board is not present.
Item	<p>(MX2010 and MX2020 Routers) Information about the chassis component: Routing Engines, Controls Boards (CBs), Switch Fabric Boards (SFBs), PICs, Flexible PIC Concentrators (FPCs), and Adapter Cards (ADCs).</p> <p>(MX104 Routers) Information about the chassis components: Routing Engines, Control Board (CB), Power Entry Module (PEM), and Compact Forwarding Engine Board (AFEB).</p> <p>(QFabric Systems) Information about the chassis component: Control Boards, Routing Engines, Flexible PIC Concentrators (FPCs), and Power Entry Modules (PEMs), Node Devices, and Interconnect Devices.</p> <p>(QFX Series) Information about the chassis component: Flexible PIC Concentrators (FPCs), and Power Entry Modules (PEMs).</p>

Table 46: show chassis environment Output Fields (*continued*)

Field Name	Field Description
Status	<p>(MX104, MX2010, and MX2020 Routers) Status of the specified chassis component. For example, if the Class is Fan, the fan status can be:</p> <ul style="list-style-type: none"> • OK: The fans are operational. • Testing: The fans are being tested during initial power-on. • Failed: The fans have failed or the fans are not spinning. • Absent: The fan tray is not installed. <p>If the Class is Power, the power supply status can be:</p> <ul style="list-style-type: none"> • OK: The power component is operational. • Testing: The power component is being tested during initial power-on. • Check: There is insufficient power---that is, fewer than the minimum required feeds are connected. • Failed: The inputs leads have failed. • Absent: The power component is not installed.
Measurement	<p>(MX104, MX2010, and MX2020 Routers) Dependant on the Class. For example, if the Class is Temp, indicates the temperature in degree Celsius and degrees Fahrenheit. If the Class is Fan, indicates actual fan RPM.</p>

Sample Output

show chassis environment (M5 Router)

```

user@host> show chassis environment
Class Item           Status Measurement
Power Power Supply A   OK
        Power Supply B Absent
Temp  FPC 0            OK      30 degrees C / 86 degrees F
        FEB           OK      33 degrees C / 91 degrees F
        PS Intake      OK      27 degrees C / 80 degrees F
        PS Exhaust     OK      27 degrees C / 80 degrees F
        Routing Engine OK      34 degrees C / 93 degrees F
Fans  Left Fan 1       OK      Spinning at normal speed
        Left Fan 2       OK      Spinning at normal speed
        Left Fan 3       OK      Spinning at normal speed
        Left Fan 4       OK      Spinning at normal speed
Misc  Craft Interface  OK

```

show chassis environment (M7i Router)

```

user@host> show chassis environment
Class Item           Status Measurement
Power Power Supply 0   OK
        Power Supply 1 Absent
Temp  Intake           OK      22 degrees C / 71 degrees F
        FPC 0          OK      23 degrees C / 73 degrees F
        Power Supplies OK      23 degrees C / 73 degrees F
        CFEB Intake    OK      24 degrees C / 75 degrees F
        CFEB Exhaust   OK      29 degrees C / 84 degrees F
        Routing Engine OK      26 degrees C / 78 degrees F
Fans  Fan 1            OK      Spinning at normal speed
        Fan 2          OK      Spinning at normal speed

```

Fan 3	OK	Spinning at normal speed
Fan 4	OK	Spinning at normal speed

show chassis environment (M10 Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Power	Power Supply A	OK	
	Power Supply B	Failed	
Temp	FPC 0	OK	36 degrees C / 96 degrees F
	FPC 1	OK	35 degrees C / 95 degrees F
	FEB	OK	34 degrees C / 93 degrees F
	PS Intake	OK	31 degrees C / 87 degrees F
	PS Exhaust	OK	34 degrees C / 93 degrees F
	Routing Engine	OK	35 degrees C / 95 degrees F
Fans	Left Fan 1	OK	Spinning at normal speed
	Left Fan 2	OK	Spinning at normal speed
	Left Fan 3	OK	Spinning at normal speed
	Left Fan 4	OK	Spinning at normal speed
Misc	Craft Interface	OK	

show chassis environment (M10i Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Power	Power Supply 0	OK	
	Power Supply 1	OK	
	Power Supply 2	Absent	
	Power Supply 3	Absent	
Temp	Intake	OK	26 degrees C / 78 degrees F
	FPC 0	OK	27 degrees C / 80 degrees F
	FPC 1	OK	28 degrees C / 82 degrees F
	Lower Power Supplies	OK	29 degrees C / 84 degrees F
	Upper Power Supplies	OK	28 degrees C / 82 degrees F
	CFEB Intake	OK	27 degrees C / 80 degrees F
	CFEB Exhaust	OK	36 degrees C / 96 degrees F
	Routing Engine 0	OK	31 degrees C / 87 degrees F
	Routing Engine 1	OK	27 degrees C / 80 degrees F
Fans	Fan Tray 0 Fan 1	OK	Spinning at normal speed
	Fan Tray 0 Fan 2	OK	Spinning at normal speed
	Fan Tray 0 Fan 3	OK	Spinning at normal speed
	Fan Tray 0 Fan 4	OK	Spinning at normal speed
	Fan Tray 0 Fan 5	OK	Spinning at normal speed
	Fan Tray 0 Fan 6	OK	Spinning at normal speed
	Fan Tray 0 Fan 7	OK	Spinning at normal speed
	Fan Tray 0 Fan 8	OK	Spinning at normal speed
	Fan Tray 1 Fan 1	Absent	
	Fan Tray 1 Fan 2	Absent	
	Fan Tray 1 Fan 3	Absent	
	Fan Tray 1 Fan 4	Absent	
	Fan Tray 1 Fan 5	Absent	
	Fan Tray 1 Fan 6	Absent	
	Fan Tray 1 Fan 7	Absent	
	Fan Tray 1 Fan 8	Absent	

show chassis environment (M20 Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Power	Power Supply A	OK	
	Power Supply B	Absent	
Temp	FPC 0	OK	28 degrees C / 82 degrees F
	FPC 1	OK	27 degrees C / 80 degrees F
	Power Supply A	OK	22 degrees C / 71 degrees F
	Power Supply B	Absent	
	SSB 0	OK	30 degrees C / 86 degrees F
	Backplane	OK	22 degrees C / 71 degrees F
Fans	Routing Engine 0	OK	26 degrees C / 78 degrees F
	Routing Engine 1	Testing	
	Rear Fan	OK	Spinning at normal speed
	Front Upper Fan	OK	Spinning at normal speed
	Front Middle Fan	OK	Spinning at normal speed
	Front Bottom Fan	OK	Spinning at normal speed
Misc	Craft Interface	OK	

show chassis environment (M40 Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Power	Power Supply A	OK	
	Power Supply B	Absent	
Temp	FPC 3	OK	24 degrees C / 75 degrees F
	FPC 6	OK	26 degrees C / 78 degrees F
	SCB	OK	26 degrees C / 78 degrees F
	Backplane @ A1	OK	28 degrees C / 82 degrees F
	Backplane @ A2	OK	23 degrees C / 73 degrees F
	Routing Engine	OK	26 degrees C / 78 degrees F
Fans	Top Impeller	OK	Spinning at normal speed
	Bottom impeller	OK	Spinning at normal speed
	Rear Left Fan	OK	Spinning at normal speed
	Rear Center Fan	OK	Spinning at normal speed
	Rear Right Fan	OK	Spinning at normal speed
Misc	Craft Interface	OK	

show chassis environment (M40e Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Power	PEM 0	OK	
	PEM 1	Absent	
Temp	PCG 0	OK	44 degrees C / 111 degrees F
	PCG 1	OK	47 degrees C / 116 degrees F
	Routing Engine 0	OK	40 degrees C / 104 degrees F
	Routing Engine 1	OK	37 degrees C / 98 degrees F
	MCS 0	OK	45 degrees C / 113 degrees F
	MCS 1	OK	42 degrees C / 107 degrees F
	SFM 0 SPP	OK	40 degrees C / 104 degrees F
	SFM 0 SPR	OK	44 degrees C / 111 degrees F
	SFM 1 SPP	OK	43 degrees C / 109 degrees F
	SFM 1 SPR	OK	45 degrees C / 113 degrees F
	FPC 0	OK	38 degrees C / 100 degrees F
	FPC 1	OK	40 degrees C / 104 degrees F
	FPC 2	OK	38 degrees C / 100 degrees F
	FPC 4	OK	34 degrees C / 93 degrees F
	FPC 5	OK	43 degrees C / 109 degrees F
	FPC 6	OK	41 degrees C / 105 degrees F
	FPC 7	OK	43 degrees C / 109 degrees F

	FPM CMB	OK	28 degrees C / 82 degrees F
	FPM Display	OK	28 degrees C / 82 degrees F
Fans	Rear Bottom Blower	OK	Spinning at normal speed
	Rear Top Blower	OK	Spinning at normal speed
	Front Top Blower	OK	Spinning at normal speed
	Fan Tray Rear Left	OK	Spinning at normal speed
	Fan Tray Rear Right	OK	Spinning at normal speed
	Fan Tray Front Left	OK	Spinning at normal speed
	Fan Tray Front Right	OK	Spinning at normal speed
Misc	CIP	OK	

show chassis environment (M120 Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Temp	PEM 0	OK	
	PEM 1	OK	
	Routing Engine 0	OK	43 degrees C / 109 degrees F
	Routing Engine 1	OK	44 degrees C / 111 degrees F
	CB 0 Intake	OK	33 degrees C / 91 degrees F
	CB 0 Exhaust A	OK	36 degrees C / 96 degrees F
	CB 0 Exhaust B	OK	35 degrees C / 95 degrees F
	CB 1 Intake	OK	34 degrees C / 93 degrees F
	CB 1 Exhaust A	OK	38 degrees C / 100 degrees F
	CB 1 Exhaust B	OK	35 degrees C / 95 degrees F
	FEB 3 Intake	OK	35 degrees C / 95 degrees F
	FEB 3 Exhaust A	OK	37 degrees C / 98 degrees F
	FEB 3 Exhaust B	OK	39 degrees C / 102 degrees F
	FEB 4 Intake	OK	33 degrees C / 91 degrees F
	FEB 4 Exhaust A	OK	39 degrees C / 102 degrees F
	FEB 4 Exhaust B	OK	36 degrees C / 96 degrees F
	FPC 2 Exhaust A	OK	32 degrees C / 89 degrees F
	FPC 2 Exhaust B	OK	31 degrees C / 87 degrees F
	FPC 3 Exhaust A	OK	32 degrees C / 89 degrees F
	FPC 3 Exhaust B	OK	33 degrees C / 91 degrees F
	FPC 4 Exhaust A	OK	32 degrees C / 89 degrees F
	FPC 4 Exhaust B	OK	30 degrees C / 86 degrees F
Fans	Front Top Tray Fan 1	OK	Spinning at normal speed
	Front Top Tray Fan 2	OK	Spinning at normal speed
	Front Top Tray Fan 3	OK	Spinning at normal speed
	Front Top Tray Fan 4	OK	Spinning at normal speed
	Front Top Tray Fan 5	OK	Spinning at normal speed
	Front Top Tray Fan 6	OK	Spinning at normal speed
	Front Top Tray Fan 7	OK	Spinning at normal speed
	Front Top Tray Fan 8	OK	Spinning at normal speed
	Front Bottom Tray Fan 1	OK	Spinning at normal speed
	Front Bottom Tray Fan 2	OK	Spinning at normal speed
	Front Bottom Tray Fan 3	OK	Spinning at normal speed
	Front Bottom Tray Fan 4	OK	Spinning at normal speed
	Front Bottom Tray Fan 5	OK	Spinning at normal speed
	Front Bottom Tray Fan 6	OK	Spinning at normal speed
	Front Bottom Tray Fan 7	OK	Spinning at normal speed
	Front Bottom Tray Fan 8	OK	Spinning at normal speed
	Rear Top Tray Fan 1	OK	Spinning at normal speed
	Rear Top Tray Fan 2	OK	Spinning at normal speed
	Rear Top Tray Fan 3	OK	Spinning at normal speed
	Rear Top Tray Fan 4	OK	Spinning at normal speed
	Rear Top Tray Fan 5	OK	Spinning at normal speed
	Rear Top Tray Fan 6	OK	Spinning at normal speed
	Rear Top Tray Fan 7	OK	Spinning at normal speed

Rear Top Tray Fan 8	OK	Spinning at normal speed
Rear Bottom Tray Fan 1	OK	Spinning at normal speed
Rear Bottom Tray Fan 2	OK	Spinning at normal speed
Rear Bottom Tray Fan 3	OK	Spinning at normal speed
Rear Bottom Tray Fan 4	OK	Spinning at normal speed
Rear Bottom Tray Fan 5	OK	Spinning at normal speed
Rear Bottom Tray Fan 6	OK	Spinning at normal speed
Rear Bottom Tray Fan 7	OK	Spinning at normal speed
Rear Bottom Tray Fan 8	OK	Spinning at normal speed

show chassis environment (M160 Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Power	PEM 0	OK	PEM 1
Temp	PCG 0	OK	45 degrees C / 113 degrees F
	PCG 1	Absent	
	Routing Engine 0	OK	35 degrees C / 95 degrees F
	Routing Engine 1	Absent	
	MCS 0	OK	50 degrees C / 122 degrees F
	SFM 0 SPP	OK	47 degrees C / 116 degrees F
	SFM 0 SPR	OK	49 degrees C / 120 degrees F
	SFM 1 SPP	OK	50 degrees C / 122 degrees F
	SFM 1 SPR	OK	50 degrees C / 122 degrees F
	SFM 2 SPP	OK	51 degrees C / 123 degrees F
	SFM 2 SPR	OK	52 degrees C / 125 degrees F
	SFM 3 SPP	OK	52 degrees C / 125 degrees F
	SFM 3 SPR	OK	48 degrees C / 118 degrees F
	FPC 0	OK	45 degrees C / 113 degrees F
	FPC 6	OK	43 degrees C / 109 degrees F
	FPM CMB	OK	31 degrees C / 87 degrees F
	FPM Display	OK	33 degrees C / 91 degrees F
Fans	Rear Bottom Blower	OK	Spinning at normal speed
	Rear Top Blower	OK	Spinning at normal speed
	Front Top Blower	OK	Spinning at normal speed
	Fan Tray Rear Left	OK	Spinning at normal speed
	Fan Tray Rear Right	OK	Spinning at normal speed
	Fan Tray Front Left	OK	Spinning at normal speed
	Fan Tray Front Right	OK	Spinning at normal speed
Misc	CIP	OK	

show chassis environment (M320 Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Temp	PEM 0	Absent	
	PEM 1	Absent	
	PEM 2	OK	
	PEM 3	OK	
	Routing Engine 0	OK	33 degrees C / 91 degrees F
	Routing Engine 1	OK	32 degrees C / 89 degrees F
	CB 0	OK	36 degrees C / 96 degrees F
	CB 1	OK	36 degrees C / 96 degrees F
	SIB 0	OK	38 degrees C / 100 degrees F
	SIB 1	OK	29 degrees C / 84 degrees F
	SIB 2	OK	38 degrees C / 100 degrees F
	SIB 3	OK	41 degrees C / 105 degrees F
	FPC 0 Intake	OK	28 degrees C / 82 degrees F
	FPC 0 Exhaust	OK	40 degrees C / 104 degrees F
	FPC 1 Intake	OK	29 degrees C / 84 degrees F

	FPC 1 Exhaust	OK	39 degrees C / 102 degrees F
	FPC 2 Intake	OK	28 degrees C / 82 degrees F
	FPC 2 Exhaust	OK	38 degrees C / 100 degrees F
	FPC 3 Intake	OK	28 degrees C / 82 degrees F
	FPC 3 Exhaust	OK	39 degrees C / 102 degrees F
	FPC 6 Intake	OK	27 degrees C / 80 degrees F
	FPC 6 Exhaust	OK	39 degrees C / 102 degrees F
	FPC 7 Intake	OK	27 degrees C / 80 degrees F
	FPC 7 Exhaust	OK	42 degrees C / 107 degrees F
	FPM GBUS	OK	30 degrees C / 86 degrees F
Fan	Top Left Front fan	OK	Spinning at normal speed
	Top Right Rear fan	OK	Spinning at normal speed
	Top Right Front fan	OK	Spinning at normal speed
	Top Left Rear fan	OK	Spinning at normal speed
	Bottom Left Front fan	OK	Spinning at normal speed
	Bottom Right Rear fan	OK	Spinning at normal speed
	Bottom Right Front fan	OK	Spinning at normal speed
	Bottom Left Rear fan	OK	Spinning at normal speed
	Rear Fan 1 (TOP)	OK	Spinning at normal speed
	Rear Fan 2	OK	Spinning at normal speed
	Rear Fan 3	OK	Spinning at normal speed
	Rear Fan 4	OK	Spinning at normal speed
	Rear Fan 5	OK	Spinning at normal speed
	Rear Fan 6	OK	Spinning at normal speed
	Rear Fan 7 (Bottom)	OK	Spinning at normal speed
Misc	CIP	OK	

show chassis environment (MX104 Router)

```

user@host> show chassis environment
Class Item                               Status Measurement
Temp PEM 0                               OK          34 degrees C / 93 degrees F
      PEM 1                               Absent
      ABB 0 Intake                         OK          33 degrees C / 91 degrees F
      ABB 0 Exhaust A                     OK          42 degrees C / 107 degrees F
      ABB 0 Exhaust B                     OK          43 degrees C / 109 degrees F
      ABB 1 Intake                         Absent
      ABB 1 Exhaust A                     Absent
      ABB 1 Exhaust B                     Absent
      Routing Engine 0                     OK          34 degrees C / 93 degrees F
      Routing Engine 0 CPU                 OK          46 degrees C / 114 degrees F
      Routing Engine 1                     Absent
      Routing Engine 1 CPU                 Absent
      AFEB 0 AFEB Processor                OK          33 degrees C / 91 degrees F
Fans  Fan 1                               OK          Spinning at normal speed
      Fan 2                               OK          Spinning at normal speed
      Fan 3                               OK          Spinning at normal speed
      Fan 4                               OK          Spinning at normal speed
      Fan 5                               OK          Spinning at normal speed

```

show chassis environment (MX240 Router)

```

user@host> show chassis environment
Class Item                               Status Measurement
Temp PEM 0                               OK          40 degrees C / 104 degrees F
      PEM 1                               OK          45 degrees C / 113 degrees F
      PEM 2                               Absent
      PEM 3                               Absent
      Routing Engine 0                     OK          39 degrees C / 102 degrees F
      Routing Engine 1                     OK          37 degrees C / 98 degrees F

```

CB 0 Intake	OK	36 degrees C / 96 degrees F
CB 0 Exhaust A	OK	34 degrees C / 93 degrees F
CB 0 Exhaust B	OK	38 degrees C / 100 degrees F
CB 0 ACBC	OK	37 degrees C / 98 degrees F
CB 0 SF A	OK	49 degrees C / 120 degrees F
CB 0 SF B	OK	41 degrees C / 105 degrees F
CB 1 Intake	OK	37 degrees C / 98 degrees F
CB 1 Exhaust A	OK	34 degrees C / 93 degrees F
CB 1 Exhaust B	OK	39 degrees C / 102 degrees F
CB 1 ACBC	OK	38 degrees C / 100 degrees F
CB 1 SF A	OK	47 degrees C / 116 degrees F
CB 1 SF B	OK	41 degrees C / 105 degrees F
FPC 1 Intake	OK	33 degrees C / 91 degrees F
FPC 1 Exhaust A	OK	38 degrees C / 100 degrees F
FPC 1 Exhaust B	OK	53 degrees C / 127 degrees F
FPC 1 I3 0 TSensor	OK	50 degrees C / 122 degrees F
FPC 1 I3 0 Chip	OK	53 degrees C / 127 degrees F
FPC 1 I3 1 TSensor	OK	49 degrees C / 120 degrees F
FPC 1 I3 1 Chip	OK	52 degrees C / 125 degrees F
FPC 1 I3 2 TSensor	OK	47 degrees C / 116 degrees F
FPC 1 I3 2 Chip	OK	49 degrees C / 120 degrees F
FPC 1 I3 3 TSensor	OK	44 degrees C / 111 degrees F
FPC 1 I3 3 Chip	OK	46 degrees C / 114 degrees F
FPC 1 IA 0 TSensor	OK	45 degrees C / 113 degrees F
FPC 1 IA 0 Chip	OK	44 degrees C / 111 degrees F
FPC 1 IA 1 TSensor	OK	44 degrees C / 111 degrees F
FPC 1 IA 1 Chip	OK	48 degrees C / 118 degrees F
FPC 2 Intake	OK	32 degrees C / 89 degrees F
FPC 2 Exhaust A	OK	40 degrees C / 104 degrees F
FPC 2 Exhaust B	OK	52 degrees C / 125 degrees F
FPC 2 I3 0 TSensor	OK	52 degrees C / 125 degrees F
FPC 2 I3 0 Chip	OK	56 degrees C / 132 degrees F
FPC 2 I3 1 TSensor	OK	52 degrees C / 125 degrees F
FPC 2 I3 1 Chip	OK	55 degrees C / 131 degrees F
FPC 2 I3 2 TSensor	OK	49 degrees C / 120 degrees F
FPC 2 I3 2 Chip	OK	52 degrees C / 125 degrees F
FPC 2 I3 3 TSensor	OK	44 degrees C / 111 degrees F
FPC 2 I3 3 Chip	OK	48 degrees C / 118 degrees F
FPC 2 IA 0 TSensor	OK	50 degrees C / 122 degrees F
FPC 2 IA 0 Chip	OK	48 degrees C / 118 degrees F
FPC 2 IA 1 TSensor	OK	47 degrees C / 116 degrees F
FPC 2 IA 1 Chip	OK	53 degrees C / 127 degrees F
Fans Front Fan	OK	Spinning at normal speed
Middle Fan	OK	Spinning at normal speed
Rear Fan	OK	Spinning at normal speed

show chassis environment (MX240 Router with SCBE)

user@host> show chassis environment			
Class	Item	Status	Measurement
Temp	PEM 0	OK	40 degrees C / 104 degrees F
	PEM 1	OK	45 degrees C / 113 degrees F
	PEM 2	Absent	
	PEM 3	Absent	
	Routing Engine 0	OK	39 degrees C / 102 degrees F
	Routing Engine 1	OK	37 degrees C / 98 degrees F
	CB 0 Intake	OK	36 degrees C / 96 degrees F
	CB 0 Exhaust A	OK	34 degrees C / 93 degrees F
	CB 0 Exhaust B	OK	38 degrees C / 100 degrees F
	CB 0 ACBC	OK	37 degrees C / 98 degrees F
	CB 0 XF A	OK	49 degrees C / 120 degrees F

	CB 0 XF B	OK	41 degrees C / 105 degrees F
	CB 1 Intake	OK	37 degrees C / 98 degrees F
	CB 1 Exhaust A	OK	34 degrees C / 93 degrees F
	CB 1 Exhaust B	OK	39 degrees C / 102 degrees F
	CB 1 ACBC	OK	38 degrees C / 100 degrees F
	CB 1 XF A	OK	47 degrees C / 116 degrees F
	CB 1 XF B	OK	41 degrees C / 105 degrees F
	FPC 1 Intake	OK	33 degrees C / 91 degrees F
	FPC 1 Exhaust A	OK	38 degrees C / 100 degrees F
	FPC 1 Exhaust B	OK	53 degrees C / 127 degrees F
	FPC 1 I3 0 TSensor	OK	50 degrees C / 122 degrees F
	FPC 1 I3 0 Chip	OK	53 degrees C / 127 degrees F
	FPC 1 I3 1 TSensor	OK	49 degrees C / 120 degrees F
	FPC 1 I3 1 Chip	OK	52 degrees C / 125 degrees F
	FPC 1 I3 2 TSensor	OK	47 degrees C / 116 degrees F
	FPC 1 I3 2 Chip	OK	49 degrees C / 120 degrees F
	FPC 1 I3 3 TSensor	OK	44 degrees C / 111 degrees F
	FPC 1 I3 3 Chip	OK	46 degrees C / 114 degrees F
	FPC 1 IA 0 TSensor	OK	45 degrees C / 113 degrees F
	FPC 1 IA 0 Chip	OK	44 degrees C / 111 degrees F
	FPC 1 IA 1 TSensor	OK	44 degrees C / 111 degrees F
	FPC 1 IA 1 Chip	OK	48 degrees C / 118 degrees F
	FPC 2 Intake	OK	32 degrees C / 89 degrees F
	FPC 2 Exhaust A	OK	40 degrees C / 104 degrees F
	FPC 2 Exhaust B	OK	52 degrees C / 125 degrees F
	FPC 2 I3 0 TSensor	OK	52 degrees C / 125 degrees F
	FPC 2 I3 0 Chip	OK	56 degrees C / 132 degrees F
	FPC 2 I3 1 TSensor	OK	52 degrees C / 125 degrees F
	FPC 2 I3 1 Chip	OK	55 degrees C / 131 degrees F
	FPC 2 I3 2 TSensor	OK	49 degrees C / 120 degrees F
	FPC 2 I3 2 Chip	OK	52 degrees C / 125 degrees F
	FPC 2 I3 3 TSensor	OK	44 degrees C / 111 degrees F
	FPC 2 I3 3 Chip	OK	48 degrees C / 118 degrees F
	FPC 2 IA 0 TSensor	OK	50 degrees C / 122 degrees F
	FPC 2 IA 0 Chip	OK	48 degrees C / 118 degrees F
	FPC 2 IA 1 TSensor	OK	47 degrees C / 116 degrees F
	FPC 2 IA 1 Chip	OK	53 degrees C / 127 degrees F
Fans	Front Fan	OK	Spinning at normal speed
	Middle Fan	OK	Spinning at normal speed
	Rear Fan	OK	Spinning at normal speed

show chassis environment (MX480 Router)

user@host> show chassis environment			
Class	Item	Status	Measurement
Temp	PEM 0	OK	35 degrees C / 95 degrees F
	PEM 1	OK	40 degrees C / 104 degrees F
	PEM 2	Absent	
	PEM 3	Absent	
	Routing Engine 0	OK	44 degrees C / 111 degrees F
	Routing Engine 1	OK	45 degrees C / 113 degrees F
	CB 0 Intake	OK	36 degrees C / 96 degrees F
	CB 0 Exhaust A	OK	38 degrees C / 100 degrees F
	CB 0 Exhaust B	OK	39 degrees C / 102 degrees F
	CB 0 ACBC	OK	37 degrees C / 98 degrees F
	CB 0 SF A	OK	51 degrees C / 123 degrees F
	CB 0 SF B	OK	44 degrees C / 111 degrees F
	CB 1 Intake	OK	36 degrees C / 96 degrees F
	CB 1 Exhaust A	OK	39 degrees C / 102 degrees F
	CB 1 Exhaust B	OK	40 degrees C / 104 degrees F
	CB 1 ACBC	OK	37 degrees C / 98 degrees F

CB 1 SF A	OK	50 degrees C / 122 degrees F
CB 1 SF B	OK	43 degrees C / 109 degrees F
FPC 0 Intake	OK	36 degrees C / 96 degrees F
FPC 0 Exhaust A	OK	39 degrees C / 102 degrees F
FPC 0 Exhaust B	OK	51 degrees C / 123 degrees F
FPC 0 I3 0 TSensor	OK	49 degrees C / 120 degrees F
FPC 0 I3 0 Chip	OK	56 degrees C / 132 degrees F
FPC 0 I3 1 TSensor	OK	47 degrees C / 116 degrees F
FPC 0 I3 1 Chip	OK	52 degrees C / 125 degrees F
FPC 0 I3 2 TSensor	OK	46 degrees C / 114 degrees F
FPC 0 I3 2 Chip	OK	48 degrees C / 118 degrees F
FPC 0 I3 3 TSensor	OK	42 degrees C / 107 degrees F
FPC 0 I3 3 Chip	OK	45 degrees C / 113 degrees F
FPC 0 IA 0 TSensor	OK	45 degrees C / 113 degrees F
FPC 0 IA 0 Chip	OK	45 degrees C / 113 degrees F
FPC 0 IA 1 TSensor	OK	44 degrees C / 111 degrees F
FPC 0 IA 1 Chip	OK	48 degrees C / 118 degrees F
FPC 1 Intake	OK	37 degrees C / 98 degrees F
FPC 1 Exhaust A	OK	41 degrees C / 105 degrees F
FPC 1 Exhaust B	OK	52 degrees C / 125 degrees F
FPC 1 I3 0 TSensor	OK	51 degrees C / 123 degrees F
FPC 1 I3 0 Chip	OK	57 degrees C / 134 degrees F
FPC 1 I3 1 TSensor	OK	48 degrees C / 118 degrees F
FPC 1 I3 1 Chip	OK	52 degrees C / 125 degrees F
FPC 1 I3 2 TSensor	OK	46 degrees C / 114 degrees F
FPC 1 I3 2 Chip	OK	50 degrees C / 122 degrees F
FPC 1 I3 3 TSensor	OK	42 degrees C / 107 degrees F
FPC 1 I3 3 Chip	OK	46 degrees C / 114 degrees F
FPC 1 IA 0 TSensor	OK	49 degrees C / 120 degrees F
FPC 1 IA 0 Chip	OK	48 degrees C / 118 degrees F
FPC 1 IA 1 TSensor	OK	46 degrees C / 114 degrees F
FPC 1 IA 1 Chip	OK	50 degrees C / 122 degrees F
Fans Top Rear Fan	OK	Spinning at normal speed
Bottom Rear Fan	OK	Spinning at normal speed
Top Middle Fan	OK	Spinning at normal speed
Bottom Middle Fan	OK	Spinning at normal speed
Top Front Fan	OK	Spinning at normal speed
Bottom Front Fan	OK	Spinning at normal speed

show chassis environment (MX480 Router with SCBE)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Temp	PEM 0	OK	35 degrees C / 95 degrees F
	PEM 1	OK	40 degrees C / 104 degrees F
	PEM 2	Absent	
	PEM 3	Absent	
	Routing Engine 0	OK	44 degrees C / 111 degrees F
	Routing Engine 1	OK	45 degrees C / 113 degrees F
	CB 0 Intake	OK	36 degrees C / 96 degrees F
	CB 0 Exhaust A	OK	38 degrees C / 100 degrees F
	CB 0 Exhaust B	OK	39 degrees C / 102 degrees F
	CB 0 ACBC	OK	37 degrees C / 98 degrees F
	CB 0 XF A	OK	51 degrees C / 123 degrees F
	CB 0 XF B	OK	44 degrees C / 111 degrees F
	CB 1 Intake	OK	36 degrees C / 96 degrees F
	CB 1 Exhaust A	OK	39 degrees C / 102 degrees F
	CB 1 Exhaust B	OK	40 degrees C / 104 degrees F
	CB 1 ACBC	OK	37 degrees C / 98 degrees F
	CB 1 XF A	OK	50 degrees C / 122 degrees F
	CB 1 XF B	OK	43 degrees C / 109 degrees F

FPC 0 Intake	OK	36 degrees C / 96 degrees F
FPC 0 Exhaust A	OK	39 degrees C / 102 degrees F
FPC 0 Exhaust B	OK	51 degrees C / 123 degrees F
FPC 0 I3 0 TSensor	OK	49 degrees C / 120 degrees F
FPC 0 I3 0 Chip	OK	56 degrees C / 132 degrees F
FPC 0 I3 1 TSensor	OK	47 degrees C / 116 degrees F
FPC 0 I3 1 Chip	OK	52 degrees C / 125 degrees F
FPC 0 I3 2 TSensor	OK	46 degrees C / 114 degrees F
FPC 0 I3 2 Chip	OK	48 degrees C / 118 degrees F
FPC 0 I3 3 TSensor	OK	42 degrees C / 107 degrees F
FPC 0 I3 3 Chip	OK	45 degrees C / 113 degrees F
FPC 0 IA 0 TSensor	OK	45 degrees C / 113 degrees F
FPC 0 IA 0 Chip	OK	45 degrees C / 113 degrees F
FPC 0 IA 1 TSensor	OK	44 degrees C / 111 degrees F
FPC 0 IA 1 Chip	OK	48 degrees C / 118 degrees F
FPC 1 Intake	OK	37 degrees C / 98 degrees F
FPC 1 Exhaust A	OK	41 degrees C / 105 degrees F
FPC 1 Exhaust B	OK	52 degrees C / 125 degrees F
FPC 1 I3 0 TSensor	OK	51 degrees C / 123 degrees F
FPC 1 I3 0 Chip	OK	57 degrees C / 134 degrees F
FPC 1 I3 1 TSensor	OK	48 degrees C / 118 degrees F
FPC 1 I3 1 Chip	OK	52 degrees C / 125 degrees F
FPC 1 I3 2 TSensor	OK	46 degrees C / 114 degrees F
FPC 1 I3 2 Chip	OK	50 degrees C / 122 degrees F
FPC 1 I3 3 TSensor	OK	42 degrees C / 107 degrees F
FPC 1 I3 3 Chip	OK	46 degrees C / 114 degrees F
FPC 1 IA 0 TSensor	OK	49 degrees C / 120 degrees F
FPC 1 IA 0 Chip	OK	48 degrees C / 118 degrees F
FPC 1 IA 1 TSensor	OK	46 degrees C / 114 degrees F
FPC 1 IA 1 Chip	OK	50 degrees C / 122 degrees F
Fans Top Rear Fan	OK	Spinning at normal speed
Bottom Rear Fan	OK	Spinning at normal speed
Top Middle Fan	OK	Spinning at normal speed
Bottom Middle Fan	OK	Spinning at normal speed
Top Front Fan	OK	Spinning at normal speed
Bottom Front Fan	OK	Spinning at normal speed

show chassis environment (MX960 Router)

```

user@host> show chassis environment

```

Class	Item	Status	Measurement
Temp	PEM 0	Absent	
	PEM 1	Absent	
	PEM 2	Check	
	PEM 3	OK	35 degrees C / 95 degrees F
	Routing Engine 0	OK	37 degrees C / 98 degrees F
	Routing Engine 1	Absent	
	CB 0 Intake	OK	24 degrees C / 75 degrees F
	CB 0 Exhaust A	OK	30 degrees C / 86 degrees F
	CB 0 Exhaust B	OK	27 degrees C / 80 degrees F
	CB 1 Intake	Absent	
	CB 1 Exhaust A	Absent	
	CB 1 Exhaust B	Absent	
	CB 1 ACBC	Absent	
	CB 1 SF A	Absent	
	CB 1 SF B	Absent	
	CB 2 Intake	Absent	
	CB 2 Exhaust A	Absent	
	CB 2 Exhaust B	Absent	
	CB 2 ACBC	Absent	
	CB 2 SF A	Absent	

	CB 2 SF B	Absent	
	FPC 4 Intake	OK	24 degrees C / 75 degrees F
	FPC 4 Exhaust A	OK	36 degrees C / 96 degrees F
	FPC 4 Exhaust B	OK	38 degrees C / 100 degrees F
	FPC 7 Intake	OK	24 degrees C / 75 degrees F
	FPC 7 Exhaust A	OK	36 degrees C / 96 degrees F
	FPC 7 Exhaust B	OK	42 degrees C / 107 degrees F
Fans	Top Fan Tray Temp	Failed	
	Top Tray Fan 1	OK	Spinning at normal speed
	Top Tray Fan 2	OK	Spinning at normal speed
	Top Tray Fan 3	OK	Spinning at normal speed
	Top Tray Fan 4	OK	Spinning at normal speed
	Top Tray Fan 5	OK	Spinning at normal speed
	Top Tray Fan 6	OK	Spinning at normal speed
	Bottom Fan Tray Temp	Failed	
	Bottom Tray Fan 1	OK	Spinning at normal speed
	Bottom Tray Fan 2	OK	Spinning at normal speed
	Bottom Tray Fan 3	OK	Spinning at normal speed
	Bottom Tray Fan 4	OK	Spinning at normal speed
	Bottom Tray Fan 5	OK	Spinning at normal speed
	Bottom Tray Fan 6	OK	Spinning at normal speed

show chassis environment (MX960 Router with SCBE)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Temp	PEM 0	Absent	
	PEM 1	OK	50 degrees C / 122 degrees F
	PEM 2	OK	50 degrees C / 122 degrees F
	PEM 3	OK	50 degrees C / 122 degrees F
	Routing Engine 0	OK	42 degrees C / 107 degrees F
	Routing Engine 0 CPU	OK	51 degrees C / 123 degrees F
	Routing Engine 1	OK	39 degrees C / 102 degrees F
	Routing Engine 1 CPU	OK	44 degrees C / 111 degrees F
	CB 0 Intake	OK	35 degrees C / 95 degrees F
	CB 0 Exhaust A	OK	36 degrees C / 96 degrees F
	CB 0 Exhaust B	OK	43 degrees C / 109 degrees F
	CB 0 ACBC	OK	38 degrees C / 100 degrees F
	CB 0 XF A	OK	53 degrees C / 127 degrees F
	CB 0 XF B	OK	47 degrees C / 116 degrees F
	CB 1 Intake	OK	35 degrees C / 95 degrees F
	CB 1 Exhaust A	OK	35 degrees C / 95 degrees F
	CB 1 Exhaust B	OK	41 degrees C / 105 degrees F
	CB 1 ACBC	OK	38 degrees C / 100 degrees F
	CB 1 XF A	OK	52 degrees C / 125 degrees F
	CB 1 XF B	OK	47 degrees C / 116 degrees F
	CB 2 Intake	OK	32 degrees C / 89 degrees F
	CB 2 Exhaust A	OK	30 degrees C / 86 degrees F
	CB 2 Exhaust B	OK	35 degrees C / 95 degrees F
	CB 2 ACBC	OK	33 degrees C / 91 degrees F
	CB 2 XF A	OK	51 degrees C / 123 degrees F
	CB 2 XF B	OK	50 degrees C / 122 degrees F
	FPC 0 Intake	OK	35 degrees C / 95 degrees F
	FPC 0 Exhaust A	OK	39 degrees C / 102 degrees F
	FPC 0 Exhaust B	OK	50 degrees C / 122 degrees F
	FPC 0 I3 0 TSensor	OK	50 degrees C / 122 degrees F
	FPC 0 I3 0 Chip	OK	56 degrees C / 132 degrees F
	FPC 0 I3 1 TSensor	OK	47 degrees C / 116 degrees F
	FPC 0 I3 1 Chip	OK	50 degrees C / 122 degrees F
	FPC 0 I3 2 TSensor	OK	45 degrees C / 113 degrees F
	FPC 0 I3 2 Chip	OK	48 degrees C / 118 degrees F

FPC 0 I3 3 TSensor	OK	41 degrees C / 105 degrees F
FPC 0 I3 3 Chip	OK	44 degrees C / 111 degrees F
FPC 0 IA 0 TSensor	OK	45 degrees C / 113 degrees F
FPC 0 IA 0 Chip	OK	45 degrees C / 113 degrees F
FPC 0 IA 1 TSensor	OK	44 degrees C / 111 degrees F
FPC 0 IA 1 Chip	OK	48 degrees C / 118 degrees F
FPC 1 Intake	OK	36 degrees C / 96 degrees F
FPC 1 Exhaust A	OK	47 degrees C / 116 degrees F
FPC 1 Exhaust B	OK	43 degrees C / 109 degrees F
FPC 1 LU 0 TCAM TSensor	OK	53 degrees C / 127 degrees F
FPC 1 LU 0 TCAM Chip	OK	57 degrees C / 134 degrees F
FPC 1 LU 0 TSensor	OK	53 degrees C / 127 degrees F
FPC 1 LU 0 Chip	OK	60 degrees C / 140 degrees F
FPC 1 MQ 0 TSensor	OK	53 degrees C / 127 degrees F
FPC 1 MQ 0 Chip	OK	56 degrees C / 132 degrees F
FPC 1 LU 1 TCAM TSensor	OK	51 degrees C / 123 degrees F
FPC 1 LU 1 TCAM Chip	OK	52 degrees C / 125 degrees F
FPC 1 LU 1 TSensor	OK	51 degrees C / 123 degrees F
FPC 1 LU 1 Chip	OK	53 degrees C / 127 degrees F
FPC 1 MQ 1 TSensor	OK	51 degrees C / 123 degrees F
FPC 1 MQ 1 Chip	OK	58 degrees C / 136 degrees F
FPC 2 Intake	OK	35 degrees C / 95 degrees F
FPC 2 Exhaust A	OK	39 degrees C / 102 degrees F
FPC 2 Exhaust B	OK	54 degrees C / 129 degrees F
FPC 2 I3 0 TSensor	OK	52 degrees C / 125 degrees F
FPC 2 I3 0 Chip	OK	59 degrees C / 138 degrees F
FPC 2 I3 1 TSensor	OK	48 degrees C / 118 degrees F
FPC 2 I3 1 Chip	OK	52 degrees C / 125 degrees F
FPC 2 I3 2 TSensor	OK	47 degrees C / 116 degrees F
FPC 2 I3 2 Chip	OK	49 degrees C / 120 degrees F
FPC 2 I3 3 TSensor	OK	41 degrees C / 105 degrees F
FPC 2 I3 3 Chip	OK	44 degrees C / 111 degrees F
FPC 2 IA 0 TSensor	OK	47 degrees C / 116 degrees F
FPC 2 IA 0 Chip	OK	46 degrees C / 114 degrees F
FPC 2 IA 1 TSensor	OK	45 degrees C / 113 degrees F
FPC 2 IA 1 Chip	OK	49 degrees C / 120 degrees F
FPC 3 Intake	OK	34 degrees C / 93 degrees F
FPC 3 Exhaust A	OK	34 degrees C / 93 degrees F
FPC 3 Exhaust B	OK	47 degrees C / 116 degrees F
FPC 3 I3 0 TSensor	OK	48 degrees C / 118 degrees F
FPC 3 I3 0 Chip	OK	52 degrees C / 125 degrees F
FPC 3 I3 1 TSensor	OK	46 degrees C / 114 degrees F
FPC 3 I3 1 Chip	OK	48 degrees C / 118 degrees F
FPC 3 IA 0 TSensor	OK	41 degrees C / 105 degrees F
FPC 3 IA 0 Chip	OK	40 degrees C / 104 degrees F
FPC 5 Intake	OK	42 degrees C / 107 degrees F
FPC 5 Exhaust A	OK	42 degrees C / 107 degrees F
FPC 5 Exhaust B	OK	53 degrees C / 127 degrees F
FPC 5 LU 0 TSensor	OK	53 degrees C / 127 degrees F
FPC 5 LU 0 Chip	OK	54 degrees C / 129 degrees F
FPC 5 LU 1 TSensor	OK	53 degrees C / 127 degrees F
FPC 5 LU 1 Chip	OK	61 degrees C / 141 degrees F
FPC 5 LU 2 TSensor	OK	53 degrees C / 127 degrees F
FPC 5 LU 2 Chip	OK	51 degrees C / 123 degrees F
FPC 5 LU 3 TSensor	OK	53 degrees C / 127 degrees F
FPC 5 LU 3 Chip	OK	53 degrees C / 127 degrees F
FPC 5 MQ 0 TSensor	OK	47 degrees C / 116 degrees F
FPC 5 MQ 0 Chip	OK	52 degrees C / 125 degrees F
FPC 5 MQ 1 TSensor	OK	47 degrees C / 116 degrees F
FPC 5 MQ 1 Chip	OK	52 degrees C / 125 degrees F
FPC 5 MQ 2 TSensor	OK	47 degrees C / 116 degrees F

	FPC 5 MQ 2 Chip	OK	46 degrees C / 114 degrees F
	FPC 5 MQ 3 TSensor	OK	47 degrees C / 116 degrees F
	FPC 5 MQ 3 Chip	OK	45 degrees C / 113 degrees F
	FPC 7 Intake	OK	36 degrees C / 96 degrees F
	FPC 7 Exhaust A	OK	35 degrees C / 95 degrees F
	FPC 7 Exhaust B	OK	33 degrees C / 91 degrees F
	FPC 7 QX 0 TSensor	OK	42 degrees C / 107 degrees F
	FPC 7 QX 0 Chip	OK	47 degrees C / 116 degrees F
	FPC 7 LU 0 TCAM TSensor	OK	42 degrees C / 107 degrees F
	FPC 7 LU 0 TCAM Chip	OK	44 degrees C / 111 degrees F
	FPC 7 LU 0 TSensor	OK	42 degrees C / 107 degrees F
	FPC 7 LU 0 Chip	OK	46 degrees C / 114 degrees F
	FPC 7 MQ 0 TSensor	OK	42 degrees C / 107 degrees F
	FPC 7 MQ 0 Chip	OK	45 degrees C / 113 degrees F
	FPC 8 Intake	OK	33 degrees C / 91 degrees F
	FPC 8 Exhaust A	OK	33 degrees C / 91 degrees F
	FPC 8 Exhaust B	OK	36 degrees C / 96 degrees F
	FPC 8 I3 0 TSensor	OK	38 degrees C / 100 degrees F
	FPC 8 I3 0 Chip	OK	43 degrees C / 109 degrees F
	FPC 8 BDS 0 TSensor	OK	37 degrees C / 98 degrees F
	FPC 8 BDS 0 Chip	OK	36 degrees C / 96 degrees F
	FPC 8 IA 0 TSensor	OK	37 degrees C / 98 degrees F
	FPC 8 IA 0 Chip	OK	37 degrees C / 98 degrees F
	FPC 10 Intake	OK	38 degrees C / 100 degrees F
	FPC 10 Exhaust A	OK	36 degrees C / 96 degrees F
	FPC 10 Exhaust B	OK	41 degrees C / 105 degrees F
	FPC 10 I3 0 TSensor	OK	40 degrees C / 104 degrees F
	FPC 10 I3 0 Chip	OK	42 degrees C / 107 degrees F
	FPC 10 I3 1 TSensor	OK	40 degrees C / 104 degrees F
	FPC 10 I3 1 Chip	OK	44 degrees C / 111 degrees F
	FPC 10 I3 2 TSensor	OK	42 degrees C / 107 degrees F
	FPC 10 I3 2 Chip	OK	43 degrees C / 109 degrees F
	FPC 10 I3 3 TSensor	OK	39 degrees C / 102 degrees F
	FPC 10 I3 3 Chip	OK	44 degrees C / 111 degrees F
	FPC 10 IA 0 TSensor	OK	36 degrees C / 96 degrees F
	FPC 10 IA 0 Chip	OK	36 degrees C / 96 degrees F
	FPC 10 IA 1 TSensor	OK	43 degrees C / 109 degrees F
	FPC 10 IA 1 Chip	OK	42 degrees C / 107 degrees F
Fans	Top Fan Tray Temp	OK	37 degrees C / 98 degrees F
	Top Tray Fan 1	OK	Spinning at normal speed
	Top Tray Fan 2	OK	Spinning at normal speed
	Top Tray Fan 3	OK	Spinning at normal speed
	Top Tray Fan 4	OK	Spinning at normal speed
	Top Tray Fan 5	OK	Spinning at normal speed
	Top Tray Fan 6	OK	Spinning at normal speed
	Bottom Fan Tray Temp	OK	28 degrees C / 82 degrees F
	Bottom Tray Fan 1	OK	Spinning at normal speed
	Bottom Tray Fan 2	OK	Spinning at normal speed
	Bottom Tray Fan 3	OK	Spinning at normal speed
	Bottom Tray Fan 4	OK	Spinning at normal speed
	Bottom Tray Fan 5	OK	Spinning at normal speed
	Bottom Tray Fan 6	OK	Spinning at normal speed

show chassis environment (MX960 Router with MPC5EQ)

user@host> show chassis environment			
Class	Item	Status	Measurement
Temp	PEM 0	OK	50 degrees C / 122 degrees F
	PEM 1	OK	45 degrees C / 113 degrees F
	PEM 2	OK	45 degrees C / 113 degrees F
	PEM 3	Absent	

Routing Engine 0	OK	31 degrees C / 87 degrees F
Routing Engine 0 CPU	OK	30 degrees C / 86 degrees F
Routing Engine 1	Present	
Routing Engine 1 CPU	Present	
CB 0 Intake	OK	29 degrees C / 84 degrees F
CB 0 Exhaust A	OK	29 degrees C / 84 degrees F
CB 0 Exhaust B	OK	34 degrees C / 93 degrees F
CB 0 ACBC	OK	32 degrees C / 89 degrees F
CB 0 XF A	OK	49 degrees C / 120 degrees F
CB 0 XF B	OK	45 degrees C / 113 degrees F
CB 1 Intake	OK	26 degrees C / 78 degrees F
CB 1 Exhaust A	OK	26 degrees C / 78 degrees F
CB 1 Exhaust B	OK	27 degrees C / 80 degrees F
CB 1 ACBC	OK	26 degrees C / 78 degrees F
CB 1 XF A	OK	32 degrees C / 89 degrees F
CB 1 XF B	OK	32 degrees C / 89 degrees F
CB 2 Intake	OK	28 degrees C / 82 degrees F
CB 2 Exhaust A	OK	27 degrees C / 80 degrees F
CB 2 Exhaust B	OK	33 degrees C / 91 degrees F
CB 2 ACBC	OK	30 degrees C / 86 degrees F
CB 2 XF A	OK	48 degrees C / 118 degrees F
CB 2 XF B	OK	46 degrees C / 114 degrees F
FPC 0 Intake	OK	38 degrees C / 100 degrees F
FPC 0 Exhaust A	OK	48 degrees C / 118 degrees F
FPC 0 Exhaust B	OK	49 degrees C / 120 degrees F
FPC 0 XL TSen	OK	48 degrees C / 118 degrees F
FPC 0 XL Chip	OK	50 degrees C / 122 degrees F
FPC 0 XL_XR0 TSen	OK	48 degrees C / 118 degrees F
FPC 0 XL_XR0 Chip	OK	53 degrees C / 127 degrees F
FPC 0 XL_XR1 TSen	OK	48 degrees C / 118 degrees F
FPC 0 XL_XR1 Chip	OK	54 degrees C / 129 degrees F
FPC 0 XQ TSen	OK	48 degrees C / 118 degrees F
FPC 0 XQ Chip	OK	52 degrees C / 125 degrees F
FPC 0 XQ_XR0 TSen	OK	48 degrees C / 118 degrees F
FPC 0 XQ_XR0 Chip	OK	62 degrees C / 143 degrees F
FPC 0 XQ_XR1 TSen	OK	48 degrees C / 118 degrees F
FPC 0 XQ_XR1 Chip	OK	62 degrees C / 143 degrees F
FPC 0 XM 0 TSen	OK	53 degrees C / 127 degrees F
FPC 0 XM 0 Chip	OK	63 degrees C / 145 degrees F
FPC 0 XM 1 TSen	OK	53 degrees C / 127 degrees F
FPC 0 XM 1 Chip	OK	46 degrees C / 114 degrees F
FPC 0 PLX PCIe Switch TSe	OK	53 degrees C / 127 degrees F
FPC 0 PLX PCIe Switch Chi	OK	66 degrees C / 150 degrees F
FPC 1 Intake	OK	31 degrees C / 87 degrees F
FPC 1 Exhaust A	OK	38 degrees C / 100 degrees F
FPC 1 Exhaust B	OK	49 degrees C / 120 degrees F
FPC 1 LU 0 TSen	OK	41 degrees C / 105 degrees F
FPC 1 LU 0 Chip	OK	47 degrees C / 116 degrees F
FPC 1 LU 1 TSen	OK	41 degrees C / 105 degrees F
FPC 1 LU 1 Chip	OK	42 degrees C / 107 degrees F
FPC 1 LU 2 TSen	OK	41 degrees C / 105 degrees F
FPC 1 LU 2 Chip	OK	46 degrees C / 114 degrees F
FPC 1 LU 3 TSen	OK	41 degrees C / 105 degrees F
FPC 1 LU 3 Chip	OK	51 degrees C / 123 degrees F
FPC 1 XM 0 TSen	OK	41 degrees C / 105 degrees F
FPC 1 XM 0 Chip	OK	49 degrees C / 120 degrees F
FPC 1 XF 0 TSen	OK	41 degrees C / 105 degrees F
FPC 1 XF 0 Chip	OK	63 degrees C / 145 degrees F
FPC 1 PLX Switch TSen	OK	41 degrees C / 105 degrees F
FPC 1 PLX Switch Chip	OK	43 degrees C / 109 degrees F
FPC 3 Intake	OK	31 degrees C / 87 degrees F

FPC 3 Exhaust A	OK	37 degrees C / 98 degrees F
FPC 3 Exhaust B	OK	43 degrees C / 109 degrees F
FPC 3 LU 0 TSen	OK	42 degrees C / 107 degrees F
FPC 3 LU 0 Chip	OK	43 degrees C / 109 degrees F
FPC 3 LU 1 TSen	OK	42 degrees C / 107 degrees F
FPC 3 LU 1 Chip	OK	46 degrees C / 114 degrees F
FPC 3 LU 2 TSen	OK	42 degrees C / 107 degrees F
FPC 3 LU 2 Chip	OK	40 degrees C / 104 degrees F
FPC 3 LU 3 TSen	OK	42 degrees C / 107 degrees F
FPC 3 LU 3 Chip	OK	41 degrees C / 105 degrees F
FPC 3 MQ 0 TSen	OK	37 degrees C / 98 degrees F
FPC 3 MQ 0 Chip	OK	37 degrees C / 98 degrees F
FPC 3 MQ 1 TSen	OK	37 degrees C / 98 degrees F
FPC 3 MQ 1 Chip	OK	40 degrees C / 104 degrees F
FPC 3 MQ 2 TSen	OK	37 degrees C / 98 degrees F
FPC 3 MQ 2 Chip	OK	36 degrees C / 96 degrees F
FPC 3 MQ 3 TSen	OK	37 degrees C / 98 degrees F
FPC 3 MQ 3 Chip	OK	38 degrees C / 100 degrees F
FPC 4 Intake	OK	34 degrees C / 93 degrees F
FPC 4 Exhaust A	OK	45 degrees C / 113 degrees F
FPC 4 Exhaust B	OK	47 degrees C / 116 degrees F
FPC 4 XL TSen	OK	44 degrees C / 111 degrees F
FPC 4 XL Chip	OK	47 degrees C / 116 degrees F
FPC 4 XL_XR0 TSen	OK	44 degrees C / 111 degrees F
FPC 4 XL_XR0 Chip	OK	48 degrees C / 118 degrees F
FPC 4 XL_XR1 TSen	OK	44 degrees C / 111 degrees F
FPC 4 XL_XR1 Chip	OK	47 degrees C / 116 degrees F
FPC 4 XQ TSen	OK	44 degrees C / 111 degrees F
FPC 4 XQ Chip	OK	47 degrees C / 116 degrees F
FPC 4 XQ_XR0 TSen	OK	44 degrees C / 111 degrees F
FPC 4 XQ_XR0 Chip	OK	57 degrees C / 134 degrees F
FPC 4 XQ_XR1 TSen	OK	44 degrees C / 111 degrees F
FPC 4 XQ_XR1 Chip	OK	58 degrees C / 136 degrees F
FPC 4 XM 0 TSen	OK	51 degrees C / 123 degrees F
FPC 4 XM 0 Chip	OK	61 degrees C / 141 degrees F
FPC 4 XM 1 TSen	OK	51 degrees C / 123 degrees F
FPC 4 XM 1 Chip	OK	47 degrees C / 116 degrees F
FPC 4 PLX PCIe Switch TSe	OK	51 degrees C / 123 degrees F
FPC 4 PLX PCIe Switch Chi	OK	60 degrees C / 140 degrees F
FPC 5 Intake	OK	34 degrees C / 93 degrees F
FPC 5 Exhaust A	OK	45 degrees C / 113 degrees F
FPC 5 Exhaust B	OK	47 degrees C / 116 degrees F
FPC 5 XL TSen	OK	45 degrees C / 113 degrees F
FPC 5 XL Chip	OK	47 degrees C / 116 degrees F
FPC 5 XL_XR0 TSen	OK	45 degrees C / 113 degrees F
FPC 5 XL_XR0 Chip	OK	49 degrees C / 120 degrees F
FPC 5 XL_XR1 TSen	OK	45 degrees C / 113 degrees F
FPC 5 XL_XR1 Chip	OK	49 degrees C / 120 degrees F
FPC 5 XQ TSen	OK	45 degrees C / 113 degrees F
FPC 5 XQ Chip	OK	48 degrees C / 118 degrees F
FPC 5 XQ_XR0 TSen	OK	45 degrees C / 113 degrees F
FPC 5 XQ_XR0 Chip	OK	60 degrees C / 140 degrees F
FPC 5 XQ_XR1 TSen	OK	45 degrees C / 113 degrees F
FPC 5 XQ_XR1 Chip	OK	58 degrees C / 136 degrees F
FPC 5 XM 0 TSen	OK	50 degrees C / 122 degrees F
FPC 5 XM 0 Chip	OK	48 degrees C / 118 degrees F
FPC 5 XM 1 TSen	OK	50 degrees C / 122 degrees F
FPC 5 XM 1 Chip	OK	47 degrees C / 116 degrees F
FPC 5 PLX PCIe Switch TSe	OK	50 degrees C / 122 degrees F
FPC 5 PLX PCIe Switch Chi	OK	59 degrees C / 138 degrees F
FPC 7 Intake	OK	32 degrees C / 89 degrees F

FPC 7 Exhaust A	OK	32 degrees C / 89 degrees F
FPC 7 Exhaust B	OK	33 degrees C / 91 degrees F
FPC 7 LU 0 TSen	OK	49 degrees C / 120 degrees F
FPC 7 LU 0 Chip	OK	44 degrees C / 111 degrees F
FPC 7 LU 1 TSen	OK	49 degrees C / 120 degrees F
FPC 7 LU 1 Chip	OK	47 degrees C / 116 degrees F
FPC 7 LU 2 TSen	OK	49 degrees C / 120 degrees F
FPC 7 LU 2 Chip	OK	39 degrees C / 102 degrees F
FPC 7 LU 3 TSen	OK	49 degrees C / 120 degrees F
FPC 7 LU 3 Chip	OK	43 degrees C / 109 degrees F
FPC 7 XM 0 TSen	OK	49 degrees C / 120 degrees F
FPC 7 XM 0 Chip	OK	57 degrees C / 134 degrees F
FPC 7 XM 1 TSen	OK	49 degrees C / 120 degrees F
FPC 7 XM 1 Chip	OK	48 degrees C / 118 degrees F
FPC 7 PLX Switch TSen	OK	49 degrees C / 120 degrees F
FPC 7 PLX Switch Chip	OK	45 degrees C / 113 degrees F
FPC 8 Intake	OK	36 degrees C / 96 degrees F
FPC 8 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 8 Exhaust B	OK	46 degrees C / 114 degrees F
FPC 8 XL TSen	OK	46 degrees C / 114 degrees F
FPC 8 XL Chip	OK	47 degrees C / 116 degrees F
FPC 8 XL_XR0 TSen	OK	46 degrees C / 114 degrees F
FPC 8 XL_XR0 Chip	OK	53 degrees C / 127 degrees F
FPC 8 XL_XR1 TSen	OK	46 degrees C / 114 degrees F
FPC 8 XL_XR1 Chip	OK	52 degrees C / 125 degrees F
FPC 8 XQ TSen	OK	46 degrees C / 114 degrees F
FPC 8 XQ Chip	OK	46 degrees C / 114 degrees F
FPC 8 XQ_XR0 TSen	OK	46 degrees C / 114 degrees F
FPC 8 XQ_XR0 Chip	OK	59 degrees C / 138 degrees F
FPC 8 XQ_XR1 TSen	OK	46 degrees C / 114 degrees F
FPC 8 XQ_XR1 Chip	OK	57 degrees C / 134 degrees F
FPC 8 XM 0 TSen	OK	52 degrees C / 125 degrees F
FPC 8 XM 0 Chip	OK	61 degrees C / 141 degrees F
FPC 8 XM 1 TSen	OK	52 degrees C / 125 degrees F
FPC 8 XM 1 Chip	OK	47 degrees C / 116 degrees F
FPC 8 PLX PCIe Switch TSe	OK	52 degrees C / 125 degrees F
FPC 8 PLX PCIe Switch Chi	OK	63 degrees C / 145 degrees F
FPC 9 Intake	OK	31 degrees C / 87 degrees F
FPC 9 Exhaust A	OK	34 degrees C / 93 degrees F
FPC 9 Exhaust B	OK	35 degrees C / 95 degrees F
FPC 9 QX 0 TSen	OK	42 degrees C / 107 degrees F
FPC 9 QX 0 Chip	OK	45 degrees C / 113 degrees F
FPC 9 LU 0 TCAM TSen	OK	42 degrees C / 107 degrees F
FPC 9 LU 0 TCAM Chip	OK	41 degrees C / 105 degrees F
FPC 9 LU 0 TSen	OK	42 degrees C / 107 degrees F
FPC 9 LU 0 Chip	OK	43 degrees C / 109 degrees F
FPC 9 MQ 0 TSen	OK	42 degrees C / 107 degrees F
FPC 9 MQ 0 Chip	OK	43 degrees C / 109 degrees F
FPC 9 QX 1 TSen	OK	38 degrees C / 100 degrees F
FPC 9 QX 1 Chip	OK	40 degrees C / 104 degrees F
FPC 9 LU 1 TCAM TSen	OK	38 degrees C / 100 degrees F
FPC 9 LU 1 TCAM Chip	OK	38 degrees C / 100 degrees F
FPC 9 LU 1 TSen	OK	38 degrees C / 100 degrees F
FPC 9 LU 1 Chip	OK	41 degrees C / 105 degrees F
FPC 9 MQ 1 TSen	OK	38 degrees C / 100 degrees F
FPC 9 MQ 1 Chip	OK	41 degrees C / 105 degrees F
FPC 10 Intake	OK	35 degrees C / 95 degrees F
FPC 10 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 10 Exhaust B	OK	46 degrees C / 114 degrees F
FPC 10 XL TSen	OK	42 degrees C / 107 degrees F
FPC 10 XL Chip	OK	44 degrees C / 111 degrees F

FPC 10 XL_XR0 TSen	OK	42 degrees C / 107 degrees F
FPC 10 XL_XR0 Chip	OK	47 degrees C / 116 degrees F
FPC 10 XL_XR1 TSen	OK	42 degrees C / 107 degrees F
FPC 10 XL_XR1 Chip	OK	48 degrees C / 118 degrees F
FPC 10 XQ TSen	OK	42 degrees C / 107 degrees F
FPC 10 XQ Chip	OK	46 degrees C / 114 degrees F
FPC 10 XQ_XR0 TSen	OK	42 degrees C / 107 degrees F
FPC 10 XQ_XR0 Chip	OK	57 degrees C / 134 degrees F
FPC 10 XQ_XR1 TSen	OK	42 degrees C / 107 degrees F
FPC 10 XQ_XR1 Chip	OK	53 degrees C / 127 degrees F
FPC 10 XM 0 TSen	OK	51 degrees C / 123 degrees F
FPC 10 XM 0 Chip	OK	61 degrees C / 141 degrees F
FPC 10 XM 1 TSen	OK	51 degrees C / 123 degrees F
FPC 10 XM 1 Chip	OK	49 degrees C / 120 degrees F
FPC 10 PLX PCIe Switch TSe	OK	51 degrees C / 123 degrees F
FPC 10 PLX PCIe Switch Chi	OK	61 degrees C / 141 degrees F
FPC 11 Intake	OK	33 degrees C / 91 degrees F
FPC 11 Exhaust A	OK	33 degrees C / 91 degrees F
FPC 11 Exhaust B	OK	34 degrees C / 93 degrees F
FPC 11 LU 0 TSen	OK	50 degrees C / 122 degrees F
FPC 11 LU 0 Chip	OK	48 degrees C / 118 degrees F
FPC 11 LU 1 TSen	OK	50 degrees C / 122 degrees F
FPC 11 LU 1 Chip	OK	50 degrees C / 122 degrees F
FPC 11 LU 2 TSen	OK	50 degrees C / 122 degrees F
FPC 11 LU 2 Chip	OK	41 degrees C / 105 degrees F
FPC 11 LU 3 TSen	OK	50 degrees C / 122 degrees F
FPC 11 LU 3 Chip	OK	48 degrees C / 118 degrees F
FPC 11 XM 0 TSen	OK	50 degrees C / 122 degrees F
FPC 11 XM 0 Chip	OK	57 degrees C / 134 degrees F
FPC 11 XM 1 TSen	OK	50 degrees C / 122 degrees F
FPC 11 XM 1 Chip	OK	52 degrees C / 125 degrees F
FPC 11 PLX Switch TSen	OK	50 degrees C / 122 degrees F
FPC 11 PLX Switch Chip	OK	45 degrees C / 113 degrees F
Fans Top Fan Tray Temp	OK	42 degrees C / 107 degrees F
Top Tray Fan 1	OK	Spinning at high speed
Top Tray Fan 2	OK	Spinning at high speed
Top Tray Fan 3	OK	Spinning at high speed
Top Tray Fan 4	OK	Spinning at high speed
Top Tray Fan 5	OK	Spinning at high speed
Top Tray Fan 6	OK	Spinning at high speed
Top Tray Fan 7	OK	Spinning at high speed
Top Tray Fan 8	OK	Spinning at high speed
Top Tray Fan 9	OK	Spinning at high speed
Top Tray Fan 10	OK	Spinning at high speed
Top Tray Fan 11	OK	Spinning at high speed
Top Tray Fan 12	OK	Spinning at high speed
Bottom Fan Tray Temp	OK	33 degrees C / 91 degrees F
Bottom Tray Fan 1	OK	Spinning at high speed
Bottom Tray Fan 2	OK	Spinning at high speed
Bottom Tray Fan 3	OK	Spinning at high speed
Bottom Tray Fan 4	OK	Spinning at high speed
Bottom Tray Fan 5	OK	Spinning at high speed
Bottom Tray Fan 6	OK	Spinning at high speed
Bottom Tray Fan 7	OK	Spinning at high speed
Bottom Tray Fan 8	OK	Spinning at high speed
Bottom Tray Fan 9	OK	Spinning at high speed
Bottom Tray Fan 10	OK	Spinning at high speed
Bottom Tray Fan 11	OK	Spinning at high speed
Bottom Tray Fan 12	OK	Spinning at high speed

show chassis environment (MX2020 Router)

user@host> show chassis environment

Class	Item	Status	Measurement
Temp	PSM 0	Absent	
	PSM 1	Absent	
	PSM 2	OK	41 degrees C / 105 degrees F
	PSM 3	OK	39 degrees C / 102 degrees F
	PSM 4	OK	39 degrees C / 102 degrees F
	PSM 5	OK	38 degrees C / 100 degrees F
	PSM 6	OK	38 degrees C / 100 degrees F
	PSM 7	OK	38 degrees C / 100 degrees F
	PSM 8	OK	37 degrees C / 98 degrees F
	PSM 9	Absent	
	PSM 10	Absent	
	PSM 11	OK	47 degrees C / 116 degrees F
	PSM 12	OK	45 degrees C / 113 degrees F
	PSM 13	OK	44 degrees C / 111 degrees F
	PSM 14	OK	44 degrees C / 111 degrees F
	PSM 15	OK	43 degrees C / 109 degrees F
	PSM 16	OK	42 degrees C / 107 degrees F
	PSM 17	OK	41 degrees C / 105 degrees F
	PDM 0	OK	
	PDM 1	Absent	
	PDM 2	Absent	
	PDM 3	OK	
	CB 0 IntakeA-Zone0	OK	45 degrees C / 113 degrees F
	CB 0 IntakeB-Zone1	OK	34 degrees C / 93 degrees F
	CB 0 IntakeC-Zone0	OK	48 degrees C / 118 degrees F
	CB 0 ExhaustA-Zone0	OK	45 degrees C / 113 degrees F
	CB 0 ExhaustB-Zone1	OK	37 degrees C / 98 degrees F
	CB 0 TCBC-Zone0	OK	41 degrees C / 105 degrees F
	CB 1 IntakeA-Zone0	OK	46 degrees C / 114 degrees F
	CB 1 IntakeB-Zone1	OK	42 degrees C / 107 degrees F
	CB 1 IntakeC-Zone0	OK	49 degrees C / 120 degrees F
	CB 1 ExhaustA-Zone0	OK	46 degrees C / 114 degrees F
	CB 1 ExhaustB-Zone1	OK	41 degrees C / 105 degrees F
	CB 1 TCBC-Zone0	OK	46 degrees C / 114 degrees F
	SPMB 0 Intake	OK	33 degrees C / 91 degrees F
	SPMB 1 Intake	OK	42 degrees C / 107 degrees F
	Routing Engine 0	OK	35 degrees C / 95 degrees F
	Routing Engine 0 CPU	OK	34 degrees C / 93 degrees F
	Routing Engine 1	OK	44 degrees C / 111 degrees F
	Routing Engine 1 CPU	OK	42 degrees C / 107 degrees F
	SFB 0 Intake-Zone0	OK	55 degrees C / 131 degrees F
	SFB 0 Exhaust-Zone1	OK	48 degrees C / 118 degrees F
	SFB 0 IntakeA-Zone0	OK	50 degrees C / 122 degrees F
	SFB 0 IntakeB-Zone1	OK	40 degrees C / 104 degrees F
	SFB 0 Exhaust-Zone0	OK	52 degrees C / 125 degrees F
	SFB 0 SFB-XF2-Zone1	OK	61 degrees C / 141 degrees F
	SFB 0 SFB-XF1-Zone0	OK	69 degrees C / 156 degrees F
	SFB 0 SFB-XF0-Zone0	OK	68 degrees C / 154 degrees F
	SFB 1 Intake-Zone0	OK	56 degrees C / 132 degrees F
	SFB 1 Exhaust-Zone1	OK	47 degrees C / 116 degrees F
	SFB 1 IntakeA-Zone0	OK	51 degrees C / 123 degrees F
	SFB 1 IntakeB-Zone1	OK	40 degrees C / 104 degrees F
	SFB 1 Exhaust-Zone0	OK	51 degrees C / 123 degrees F
	SFB 1 SFB-XF2-Zone1	OK	62 degrees C / 143 degrees F
	SFB 1 SFB-XF1-Zone0	OK	67 degrees C / 152 degrees F
	SFB 1 SFB-XF0-Zone0	OK	69 degrees C / 156 degrees F
	SFB 2 Intake-Zone0	OK	56 degrees C / 132 degrees F

SFB 2 Exhaust-Zone1	OK	47 degrees C / 116 degrees F
SFB 2 IntakeA-Zone0	OK	51 degrees C / 123 degrees F
SFB 2 IntakeB-Zone1	OK	40 degrees C / 104 degrees F
SFB 2 Exhaust-Zone0	OK	53 degrees C / 127 degrees F
SFB 2 SFB-XF2-Zone1	OK	65 degrees C / 149 degrees F
SFB 2 SFB-XF1-Zone0	OK	69 degrees C / 156 degrees F
SFB 2 SFB-XF0-Zone0	OK	70 degrees C / 158 degrees F
SFB 3 Intake-Zone0	OK	57 degrees C / 134 degrees F
SFB 3 Exhaust-Zone1	OK	48 degrees C / 118 degrees F
SFB 3 IntakeA-Zone0	OK	52 degrees C / 125 degrees F
SFB 3 IntakeB-Zone1	OK	41 degrees C / 105 degrees F
SFB 3 Exhaust-Zone0	OK	53 degrees C / 127 degrees F
SFB 3 SFB-XF2-Zone1	OK	66 degrees C / 150 degrees F
SFB 3 SFB-XF1-Zone0	OK	69 degrees C / 156 degrees F
SFB 3 SFB-XF0-Zone0	OK	71 degrees C / 159 degrees F
SFB 4 Intake-Zone0	OK	58 degrees C / 136 degrees F
SFB 4 Exhaust-Zone1	OK	49 degrees C / 120 degrees F
SFB 4 IntakeA-Zone0	OK	54 degrees C / 129 degrees F
SFB 4 IntakeB-Zone1	OK	42 degrees C / 107 degrees F
SFB 4 Exhaust-Zone0	OK	53 degrees C / 127 degrees F
SFB 4 SFB-XF2-Zone1	OK	64 degrees C / 147 degrees F
SFB 4 SFB-XF1-Zone0	OK	68 degrees C / 154 degrees F
SFB 4 SFB-XF0-Zone0	OK	71 degrees C / 159 degrees F
SFB 5 Intake-Zone0	OK	58 degrees C / 136 degrees F
SFB 5 Exhaust-Zone1	OK	50 degrees C / 122 degrees F
SFB 5 IntakeA-Zone0	OK	53 degrees C / 127 degrees F
SFB 5 IntakeB-Zone1	OK	43 degrees C / 109 degrees F
SFB 5 Exhaust-Zone0	OK	54 degrees C / 129 degrees F
SFB 5 SFB-XF2-Zone1	OK	66 degrees C / 150 degrees F
SFB 5 SFB-XF1-Zone0	OK	69 degrees C / 156 degrees F
SFB 5 SFB-XF0-Zone0	OK	74 degrees C / 165 degrees F
SFB 6 Intake-Zone0	OK	58 degrees C / 136 degrees F
SFB 6 Exhaust-Zone1	OK	49 degrees C / 120 degrees F
SFB 6 IntakeA-Zone0	OK	53 degrees C / 127 degrees F
SFB 6 IntakeB-Zone1	OK	43 degrees C / 109 degrees F
SFB 6 Exhaust-Zone0	OK	53 degrees C / 127 degrees F
SFB 6 SFB-XF2-Zone1	OK	65 degrees C / 149 degrees F
SFB 6 SFB-XF1-Zone0	OK	68 degrees C / 154 degrees F
SFB 6 SFB-XF0-Zone0	OK	72 degrees C / 161 degrees F
SFB 7 Intake-Zone0	OK	57 degrees C / 134 degrees F
SFB 7 Exhaust-Zone1	OK	50 degrees C / 122 degrees F
SFB 7 IntakeA-Zone0	OK	53 degrees C / 127 degrees F
SFB 7 IntakeB-Zone1	OK	43 degrees C / 109 degrees F
SFB 7 Exhaust-Zone0	OK	54 degrees C / 129 degrees F
SFB 7 SFB-XF2-Zone1	OK	68 degrees C / 154 degrees F
SFB 7 SFB-XF1-Zone0	OK	69 degrees C / 156 degrees F
SFB 7 SFB-XF0-Zone0	OK	73 degrees C / 163 degrees F
FPC 0 Intake	OK	41 degrees C / 105 degrees F
FPC 0 Exhaust A	OK	48 degrees C / 118 degrees F
FPC 0 Exhaust B	OK	62 degrees C / 143 degrees F
FPC 0 LU 0 TSen	OK	59 degrees C / 138 degrees F
FPC 0 LU 0 Chip	OK	62 degrees C / 143 degrees F
FPC 0 LU 1 TSen	OK	59 degrees C / 138 degrees F
FPC 0 LU 1 Chip	OK	64 degrees C / 147 degrees F
FPC 0 LU 2 TSen	OK	59 degrees C / 138 degrees F
FPC 0 LU 2 Chip	OK	53 degrees C / 127 degrees F
FPC 0 LU 3 TSen	OK	59 degrees C / 138 degrees F
FPC 0 LU 3 Chip	OK	53 degrees C / 127 degrees F
FPC 0 MQ 0 TSen	OK	47 degrees C / 116 degrees F
FPC 0 MQ 0 Chip	OK	49 degrees C / 120 degrees F
FPC 0 MQ 1 TSen	OK	47 degrees C / 116 degrees F

FPC 0 MQ 1 Chip	OK	51 degrees C / 123 degrees F
FPC 0 MQ 2 TSen	OK	47 degrees C / 116 degrees F
FPC 0 MQ 2 Chip	OK	44 degrees C / 111 degrees F
FPC 0 MQ 3 TSen	OK	47 degrees C / 116 degrees F
FPC 0 MQ 3 Chip	OK	45 degrees C / 113 degrees F
FPC 1 Intake	OK	40 degrees C / 104 degrees F
FPC 1 Exhaust A	OK	49 degrees C / 120 degrees F
FPC 1 Exhaust B	OK	58 degrees C / 136 degrees F
FPC 1 LU 0 TSen	OK	55 degrees C / 131 degrees F
FPC 1 LU 0 Chip	OK	56 degrees C / 132 degrees F
FPC 1 LU 1 TSen	OK	55 degrees C / 131 degrees F
FPC 1 LU 1 Chip	OK	58 degrees C / 136 degrees F
FPC 1 LU 2 TSen	OK	55 degrees C / 131 degrees F
FPC 1 LU 2 Chip	OK	49 degrees C / 120 degrees F
FPC 1 LU 3 TSen	OK	55 degrees C / 131 degrees F
FPC 1 LU 3 Chip	OK	51 degrees C / 123 degrees F
FPC 1 MQ 0 TSen	OK	47 degrees C / 116 degrees F
FPC 1 MQ 0 Chip	OK	48 degrees C / 118 degrees F
FPC 1 MQ 1 TSen	OK	47 degrees C / 116 degrees F
FPC 1 MQ 1 Chip	OK	50 degrees C / 122 degrees F
FPC 1 MQ 2 TSen	OK	47 degrees C / 116 degrees F
FPC 1 MQ 2 Chip	OK	44 degrees C / 111 degrees F
FPC 1 MQ 3 TSen	OK	47 degrees C / 116 degrees F
FPC 1 MQ 3 Chip	OK	44 degrees C / 111 degrees F
FPC 2 Intake	OK	39 degrees C / 102 degrees F
FPC 2 Exhaust A	OK	49 degrees C / 120 degrees F
FPC 2 Exhaust B	OK	61 degrees C / 141 degrees F
FPC 2 LU 0 TSen	OK	58 degrees C / 136 degrees F
FPC 2 LU 0 Chip	OK	60 degrees C / 140 degrees F
FPC 2 LU 1 TSen	OK	58 degrees C / 136 degrees F
FPC 2 LU 1 Chip	OK	65 degrees C / 149 degrees F
FPC 2 LU 2 TSen	OK	58 degrees C / 136 degrees F
FPC 2 LU 2 Chip	OK	51 degrees C / 123 degrees F
FPC 2 LU 3 TSen	OK	58 degrees C / 136 degrees F
FPC 2 LU 3 Chip	OK	53 degrees C / 127 degrees F
FPC 2 MQ 0 TSen	OK	47 degrees C / 116 degrees F
FPC 2 MQ 0 Chip	OK	50 degrees C / 122 degrees F
FPC 2 MQ 1 TSen	OK	47 degrees C / 116 degrees F
FPC 2 MQ 1 Chip	OK	52 degrees C / 125 degrees F
FPC 2 MQ 2 TSen	OK	47 degrees C / 116 degrees F
FPC 2 MQ 2 Chip	OK	45 degrees C / 113 degrees F
FPC 2 MQ 3 TSen	OK	47 degrees C / 116 degrees F
FPC 2 MQ 3 Chip	OK	46 degrees C / 114 degrees F
FPC 3 Intake	OK	40 degrees C / 104 degrees F
FPC 3 Exhaust A	OK	49 degrees C / 120 degrees F
FPC 3 Exhaust B	OK	61 degrees C / 141 degrees F
FPC 3 LU 0 TSen	OK	58 degrees C / 136 degrees F
FPC 3 LU 0 Chip	OK	61 degrees C / 141 degrees F
FPC 3 LU 1 TSen	OK	58 degrees C / 136 degrees F
FPC 3 LU 1 Chip	OK	62 degrees C / 143 degrees F
FPC 3 LU 2 TSen	OK	58 degrees C / 136 degrees F
FPC 3 LU 2 Chip	OK	51 degrees C / 123 degrees F
FPC 3 LU 3 TSen	OK	58 degrees C / 136 degrees F
FPC 3 LU 3 Chip	OK	53 degrees C / 127 degrees F
FPC 3 MQ 0 TSen	OK	48 degrees C / 118 degrees F
FPC 3 MQ 0 Chip	OK	50 degrees C / 122 degrees F
FPC 3 MQ 1 TSen	OK	48 degrees C / 118 degrees F
FPC 3 MQ 1 Chip	OK	54 degrees C / 129 degrees F
FPC 3 MQ 2 TSen	OK	48 degrees C / 118 degrees F
FPC 3 MQ 2 Chip	OK	45 degrees C / 113 degrees F
FPC 3 MQ 3 TSen	OK	48 degrees C / 118 degrees F

FPC 3 MQ 3 Chip	OK	48 degrees C / 118 degrees F
FPC 4 Intake	OK	40 degrees C / 104 degrees F
FPC 4 Exhaust A	OK	49 degrees C / 120 degrees F
FPC 4 Exhaust B	OK	62 degrees C / 143 degrees F
FPC 4 LU 0 TSen	OK	59 degrees C / 138 degrees F
FPC 4 LU 0 Chip	OK	62 degrees C / 143 degrees F
FPC 4 LU 1 TSen	OK	59 degrees C / 138 degrees F
FPC 4 LU 1 Chip	OK	65 degrees C / 149 degrees F
FPC 4 LU 2 TSen	OK	59 degrees C / 138 degrees F
FPC 4 LU 2 Chip	OK	51 degrees C / 123 degrees F
FPC 4 LU 3 TSen	OK	59 degrees C / 138 degrees F
FPC 4 LU 3 Chip	OK	53 degrees C / 127 degrees F
FPC 4 MQ 0 TSen	OK	48 degrees C / 118 degrees F
FPC 4 MQ 0 Chip	OK	52 degrees C / 125 degrees F
FPC 4 MQ 1 TSen	OK	48 degrees C / 118 degrees F
FPC 4 MQ 1 Chip	OK	53 degrees C / 127 degrees F
FPC 4 MQ 2 TSen	OK	48 degrees C / 118 degrees F
FPC 4 MQ 2 Chip	OK	46 degrees C / 114 degrees F
FPC 4 MQ 3 TSen	OK	48 degrees C / 118 degrees F
FPC 4 MQ 3 Chip	OK	47 degrees C / 116 degrees F
FPC 5 Intake	OK	41 degrees C / 105 degrees F
FPC 5 Exhaust A	OK	50 degrees C / 122 degrees F
FPC 5 Exhaust B	OK	63 degrees C / 145 degrees F
FPC 5 LU 0 TSen	OK	60 degrees C / 140 degrees F
FPC 5 LU 0 Chip	OK	63 degrees C / 145 degrees F
FPC 5 LU 1 TSen	OK	60 degrees C / 140 degrees F
FPC 5 LU 1 Chip	OK	66 degrees C / 150 degrees F
FPC 5 LU 2 TSen	OK	60 degrees C / 140 degrees F
FPC 5 LU 2 Chip	OK	56 degrees C / 132 degrees F
FPC 5 LU 3 TSen	OK	60 degrees C / 140 degrees F
FPC 5 LU 3 Chip	OK	54 degrees C / 129 degrees F
FPC 5 MQ 0 TSen	OK	49 degrees C / 120 degrees F
FPC 5 MQ 0 Chip	OK	52 degrees C / 125 degrees F
FPC 5 MQ 1 TSen	OK	49 degrees C / 120 degrees F
FPC 5 MQ 1 Chip	OK	53 degrees C / 127 degrees F
FPC 5 MQ 2 TSen	OK	49 degrees C / 120 degrees F
FPC 5 MQ 2 Chip	OK	48 degrees C / 118 degrees F
FPC 5 MQ 3 TSen	OK	49 degrees C / 120 degrees F
FPC 5 MQ 3 Chip	OK	47 degrees C / 116 degrees F
FPC 6 Intake	OK	42 degrees C / 107 degrees F
FPC 6 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 6 Exhaust B	OK	63 degrees C / 145 degrees F
FPC 6 LU 0 TSen	OK	61 degrees C / 141 degrees F
FPC 6 LU 0 Chip	OK	64 degrees C / 147 degrees F
FPC 6 LU 1 TSen	OK	61 degrees C / 141 degrees F
FPC 6 LU 1 Chip	OK	66 degrees C / 150 degrees F
FPC 6 LU 2 TSen	OK	61 degrees C / 141 degrees F
FPC 6 LU 2 Chip	OK	56 degrees C / 132 degrees F
FPC 6 LU 3 TSen	OK	61 degrees C / 141 degrees F
FPC 6 LU 3 Chip	OK	56 degrees C / 132 degrees F
FPC 6 MQ 0 TSen	OK	50 degrees C / 122 degrees F
FPC 6 MQ 0 Chip	OK	56 degrees C / 132 degrees F
FPC 6 MQ 1 TSen	OK	50 degrees C / 122 degrees F
FPC 6 MQ 1 Chip	OK	59 degrees C / 138 degrees F
FPC 6 MQ 2 TSen	OK	50 degrees C / 122 degrees F
FPC 6 MQ 2 Chip	OK	49 degrees C / 120 degrees F
FPC 6 MQ 3 TSen	OK	50 degrees C / 122 degrees F
FPC 6 MQ 3 Chip	OK	49 degrees C / 120 degrees F
FPC 7 Intake	OK	41 degrees C / 105 degrees F
FPC 7 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 7 Exhaust B	OK	63 degrees C / 145 degrees F

FPC 7 LU 0 TSen	OK	60 degrees C / 140 degrees F
FPC 7 LU 0 Chip	OK	61 degrees C / 141 degrees F
FPC 7 LU 1 TSen	OK	60 degrees C / 140 degrees F
FPC 7 LU 1 Chip	OK	65 degrees C / 149 degrees F
FPC 7 LU 2 TSen	OK	60 degrees C / 140 degrees F
FPC 7 LU 2 Chip	OK	54 degrees C / 129 degrees F
FPC 7 LU 3 TSen	OK	60 degrees C / 140 degrees F
FPC 7 LU 3 Chip	OK	53 degrees C / 127 degrees F
FPC 7 MQ 0 TSen	OK	50 degrees C / 122 degrees F
FPC 7 MQ 0 Chip	OK	53 degrees C / 127 degrees F
FPC 7 MQ 1 TSen	OK	50 degrees C / 122 degrees F
FPC 7 MQ 1 Chip	OK	54 degrees C / 129 degrees F
FPC 7 MQ 2 TSen	OK	50 degrees C / 122 degrees F
FPC 7 MQ 2 Chip	OK	47 degrees C / 116 degrees F
FPC 7 MQ 3 TSen	OK	50 degrees C / 122 degrees F
FPC 7 MQ 3 Chip	OK	47 degrees C / 116 degrees F
FPC 8 Intake	OK	41 degrees C / 105 degrees F
FPC 8 Exhaust A	OK	50 degrees C / 122 degrees F
FPC 8 Exhaust B	OK	62 degrees C / 143 degrees F
FPC 8 LU 0 TSen	OK	59 degrees C / 138 degrees F
FPC 8 LU 0 Chip	OK	62 degrees C / 143 degrees F
FPC 8 LU 1 TSen	OK	59 degrees C / 138 degrees F
FPC 8 LU 1 Chip	OK	64 degrees C / 147 degrees F
FPC 8 LU 2 TSen	OK	59 degrees C / 138 degrees F
FPC 8 LU 2 Chip	OK	55 degrees C / 131 degrees F
FPC 8 LU 3 TSen	OK	59 degrees C / 138 degrees F
FPC 8 LU 3 Chip	OK	54 degrees C / 129 degrees F
FPC 8 MQ 0 TSen	OK	49 degrees C / 120 degrees F
FPC 8 MQ 0 Chip	OK	51 degrees C / 123 degrees F
FPC 8 MQ 1 TSen	OK	49 degrees C / 120 degrees F
FPC 8 MQ 1 Chip	OK	52 degrees C / 125 degrees F
FPC 8 MQ 2 TSen	OK	49 degrees C / 120 degrees F
FPC 8 MQ 2 Chip	OK	46 degrees C / 114 degrees F
FPC 8 MQ 3 TSen	OK	49 degrees C / 120 degrees F
FPC 8 MQ 3 Chip	OK	47 degrees C / 116 degrees F
FPC 9 Intake	OK	42 degrees C / 107 degrees F
FPC 9 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 9 Exhaust B	OK	63 degrees C / 145 degrees F
FPC 9 LU 0 TSen	OK	60 degrees C / 140 degrees F
FPC 9 LU 0 Chip	OK	65 degrees C / 149 degrees F
FPC 9 LU 1 TSen	OK	60 degrees C / 140 degrees F
FPC 9 LU 1 Chip	OK	67 degrees C / 152 degrees F
FPC 9 LU 2 TSen	OK	60 degrees C / 140 degrees F
FPC 9 LU 2 Chip	OK	54 degrees C / 129 degrees F
FPC 9 LU 3 TSen	OK	60 degrees C / 140 degrees F
FPC 9 LU 3 Chip	OK	54 degrees C / 129 degrees F
FPC 9 MQ 0 TSen	OK	51 degrees C / 123 degrees F
FPC 9 MQ 0 Chip	OK	55 degrees C / 131 degrees F
FPC 9 MQ 1 TSen	OK	51 degrees C / 123 degrees F
FPC 9 MQ 1 Chip	OK	59 degrees C / 138 degrees F
FPC 9 MQ 2 TSen	OK	51 degrees C / 123 degrees F
FPC 9 MQ 2 Chip	OK	49 degrees C / 120 degrees F
FPC 9 MQ 3 TSen	OK	51 degrees C / 123 degrees F
FPC 9 MQ 3 Chip	OK	49 degrees C / 120 degrees F
FPC 10 Intake	OK	44 degrees C / 111 degrees F
FPC 10 Exhaust A	OK	49 degrees C / 120 degrees F
FPC 10 Exhaust B	OK	55 degrees C / 131 degrees F
FPC 10 LU 0 TSen	OK	54 degrees C / 129 degrees F
FPC 10 LU 0 Chip	OK	55 degrees C / 131 degrees F
FPC 10 LU 1 TSen	OK	54 degrees C / 129 degrees F
FPC 10 LU 1 Chip	OK	59 degrees C / 138 degrees F

FPC 10 LU 2 TSen	OK	54 degrees C / 129 degrees F
FPC 10 LU 2 Chip	OK	52 degrees C / 125 degrees F
FPC 10 LU 3 TSen	OK	54 degrees C / 129 degrees F
FPC 10 LU 3 Chip	OK	51 degrees C / 123 degrees F
FPC 10 MQ 0 TSen	OK	48 degrees C / 118 degrees F
FPC 10 MQ 0 Chip	OK	49 degrees C / 120 degrees F
FPC 10 MQ 1 TSen	OK	48 degrees C / 118 degrees F
FPC 10 MQ 1 Chip	OK	52 degrees C / 125 degrees F
FPC 10 MQ 2 TSen	OK	48 degrees C / 118 degrees F
FPC 10 MQ 2 Chip	OK	47 degrees C / 116 degrees F
FPC 10 MQ 3 TSen	OK	48 degrees C / 118 degrees F
FPC 10 MQ 3 Chip	OK	47 degrees C / 116 degrees F
FPC 11 Intake	OK	30 degrees C / 86 degrees F
FPC 11 Exhaust A	OK	35 degrees C / 95 degrees F
FPC 11 Exhaust B	OK	30 degrees C / 86 degrees F
FPC 11 LU 0 TSen	OK	57 degrees C / 134 degrees F
FPC 11 LU 0 Chip	OK	58 degrees C / 136 degrees F
FPC 11 LU 1 TSen	OK	57 degrees C / 134 degrees F
FPC 11 LU 1 Chip	OK	62 degrees C / 143 degrees F
FPC 11 LU 2 TSen	OK	57 degrees C / 134 degrees F
FPC 11 LU 2 Chip	OK	53 degrees C / 127 degrees F
FPC 11 LU 3 TSen	OK	57 degrees C / 134 degrees F
FPC 11 LU 3 Chip	OK	54 degrees C / 129 degrees F
FPC 11 MQ 0 TSen	OK	52 degrees C / 125 degrees F
FPC 11 MQ 0 Chip	OK	52 degrees C / 125 degrees F
FPC 11 MQ 1 TSen	OK	52 degrees C / 125 degrees F
FPC 11 MQ 1 Chip	OK	57 degrees C / 134 degrees F
FPC 11 MQ 2 TSen	OK	52 degrees C / 125 degrees F
FPC 11 MQ 2 Chip	OK	48 degrees C / 118 degrees F
FPC 11 MQ 3 TSen	OK	52 degrees C / 125 degrees F
FPC 11 MQ 3 Chip	OK	52 degrees C / 125 degrees F
FPC 12 Intake	OK	40 degrees C / 104 degrees F
FPC 12 Exhaust A	OK	47 degrees C / 116 degrees F
FPC 12 Exhaust B	OK	52 degrees C / 125 degrees F
FPC 12 LU 0 TSen	OK	51 degrees C / 123 degrees F
FPC 12 LU 0 Chip	OK	52 degrees C / 125 degrees F
FPC 12 LU 1 TSen	OK	51 degrees C / 123 degrees F
FPC 12 LU 1 Chip	OK	55 degrees C / 131 degrees F
FPC 12 LU 2 TSen	OK	51 degrees C / 123 degrees F
FPC 12 LU 2 Chip	OK	47 degrees C / 116 degrees F
FPC 12 LU 3 TSen	OK	51 degrees C / 123 degrees F
FPC 12 LU 3 Chip	OK	50 degrees C / 122 degrees F
FPC 12 MQ 0 TSen	OK	46 degrees C / 114 degrees F
FPC 12 MQ 0 Chip	OK	46 degrees C / 114 degrees F
FPC 12 MQ 1 TSen	OK	46 degrees C / 114 degrees F
FPC 12 MQ 1 Chip	OK	50 degrees C / 122 degrees F
FPC 12 MQ 2 TSen	OK	46 degrees C / 114 degrees F
FPC 12 MQ 2 Chip	OK	44 degrees C / 111 degrees F
FPC 12 MQ 3 TSen	OK	46 degrees C / 114 degrees F
FPC 12 MQ 3 Chip	OK	46 degrees C / 114 degrees F
FPC 13 Intake	OK	40 degrees C / 104 degrees F
FPC 13 Exhaust A	OK	48 degrees C / 118 degrees F
FPC 13 Exhaust B	OK	52 degrees C / 125 degrees F
FPC 13 LU 0 TSen	OK	51 degrees C / 123 degrees F
FPC 13 LU 0 Chip	OK	52 degrees C / 125 degrees F
FPC 13 LU 1 TSen	OK	51 degrees C / 123 degrees F
FPC 13 LU 1 Chip	OK	55 degrees C / 131 degrees F
FPC 13 LU 2 TSen	OK	51 degrees C / 123 degrees F
FPC 13 LU 2 Chip	OK	48 degrees C / 118 degrees F
FPC 13 LU 3 TSen	OK	51 degrees C / 123 degrees F
FPC 13 LU 3 Chip	OK	48 degrees C / 118 degrees F

FPC 13 MQ 0 TSen	OK	46 degrees C / 114 degrees F
FPC 13 MQ 0 Chip	OK	46 degrees C / 114 degrees F
FPC 13 MQ 1 TSen	OK	46 degrees C / 114 degrees F
FPC 13 MQ 1 Chip	OK	50 degrees C / 122 degrees F
FPC 13 MQ 2 TSen	OK	46 degrees C / 114 degrees F
FPC 13 MQ 2 Chip	OK	44 degrees C / 111 degrees F
FPC 13 MQ 3 TSen	OK	46 degrees C / 114 degrees F
FPC 13 MQ 3 Chip	OK	46 degrees C / 114 degrees F
FPC 14 Intake	OK	40 degrees C / 104 degrees F
FPC 14 Exhaust A	OK	50 degrees C / 122 degrees F
FPC 14 Exhaust B	OK	51 degrees C / 123 degrees F
FPC 14 LU 0 TSen	OK	50 degrees C / 122 degrees F
FPC 14 LU 0 Chip	OK	50 degrees C / 122 degrees F
FPC 14 LU 1 TSen	OK	50 degrees C / 122 degrees F
FPC 14 LU 1 Chip	OK	54 degrees C / 129 degrees F
FPC 14 LU 2 TSen	OK	50 degrees C / 122 degrees F
FPC 14 LU 2 Chip	OK	47 degrees C / 116 degrees F
FPC 14 LU 3 TSen	OK	50 degrees C / 122 degrees F
FPC 14 LU 3 Chip	OK	49 degrees C / 120 degrees F
FPC 14 MQ 0 TSen	OK	47 degrees C / 116 degrees F
FPC 14 MQ 0 Chip	OK	46 degrees C / 114 degrees F
FPC 14 MQ 1 TSen	OK	47 degrees C / 116 degrees F
FPC 14 MQ 1 Chip	OK	51 degrees C / 123 degrees F
FPC 14 MQ 2 TSen	OK	47 degrees C / 116 degrees F
FPC 14 MQ 2 Chip	OK	45 degrees C / 113 degrees F
FPC 14 MQ 3 TSen	OK	47 degrees C / 116 degrees F
FPC 14 MQ 3 Chip	OK	48 degrees C / 118 degrees F
FPC 15 Intake	OK	44 degrees C / 111 degrees F
FPC 15 Exhaust A	OK	49 degrees C / 120 degrees F
FPC 15 Exhaust B	OK	60 degrees C / 140 degrees F
FPC 15 LU 0 TSen	OK	50 degrees C / 122 degrees F
FPC 15 LU 0 Chip	OK	56 degrees C / 132 degrees F
FPC 15 LU 1 TSen	OK	50 degrees C / 122 degrees F
FPC 15 LU 1 Chip	OK	50 degrees C / 122 degrees F
FPC 15 LU 2 TSen	OK	50 degrees C / 122 degrees F
FPC 15 LU 2 Chip	OK	58 degrees C / 136 degrees F
FPC 15 LU 3 TSen	OK	50 degrees C / 122 degrees F
FPC 15 LU 3 Chip	OK	63 degrees C / 145 degrees F
FPC 15 XM 0 TSen	OK	50 degrees C / 122 degrees F
FPC 15 XM 0 Chip	OK	56 degrees C / 132 degrees F
FPC 15 XF 0 TSen	OK	50 degrees C / 122 degrees F
FPC 15 XF 0 Chip	OK	68 degrees C / 154 degrees F
FPC 15 PLX Switch TSen	OK	50 degrees C / 122 degrees F
FPC 15 PLX Switch Chip	OK	56 degrees C / 132 degrees F
FPC 16 Intake	OK	42 degrees C / 107 degrees F
FPC 16 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 16 Exhaust B	OK	53 degrees C / 127 degrees F
FPC 16 LU 0 TSen	OK	51 degrees C / 123 degrees F
FPC 16 LU 0 Chip	OK	52 degrees C / 125 degrees F
FPC 16 LU 1 TSen	OK	51 degrees C / 123 degrees F
FPC 16 LU 1 Chip	OK	55 degrees C / 131 degrees F
FPC 16 LU 2 TSen	OK	51 degrees C / 123 degrees F
FPC 16 LU 2 Chip	OK	48 degrees C / 118 degrees F
FPC 16 LU 3 TSen	OK	51 degrees C / 123 degrees F
FPC 16 LU 3 Chip	OK	49 degrees C / 120 degrees F
FPC 16 MQ 0 TSen	OK	49 degrees C / 120 degrees F
FPC 16 MQ 0 Chip	OK	48 degrees C / 118 degrees F
FPC 16 MQ 1 TSen	OK	49 degrees C / 120 degrees F
FPC 16 MQ 1 Chip	OK	53 degrees C / 127 degrees F
FPC 16 MQ 2 TSen	OK	49 degrees C / 120 degrees F
FPC 16 MQ 2 Chip	OK	46 degrees C / 114 degrees F

FPC 16 MQ 3 TSen	OK	49 degrees C / 120 degrees F
FPC 16 MQ 3 Chip	OK	49 degrees C / 120 degrees F
FPC 17 Intake	OK	43 degrees C / 109 degrees F
FPC 17 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 17 Exhaust B	OK	55 degrees C / 131 degrees F
FPC 17 LU 0 TSen	OK	54 degrees C / 129 degrees F
FPC 17 LU 0 Chip	OK	57 degrees C / 134 degrees F
FPC 17 LU 1 TSen	OK	54 degrees C / 129 degrees F
FPC 17 LU 1 Chip	OK	60 degrees C / 140 degrees F
FPC 17 LU 2 TSen	OK	54 degrees C / 129 degrees F
FPC 17 LU 2 Chip	OK	53 degrees C / 127 degrees F
FPC 17 LU 3 TSen	OK	54 degrees C / 129 degrees F
FPC 17 LU 3 Chip	OK	53 degrees C / 127 degrees F
FPC 17 MQ 0 TSen	OK	49 degrees C / 120 degrees F
FPC 17 MQ 0 Chip	OK	50 degrees C / 122 degrees F
FPC 17 MQ 1 TSen	OK	49 degrees C / 120 degrees F
FPC 17 MQ 1 Chip	OK	54 degrees C / 129 degrees F
FPC 17 MQ 2 TSen	OK	49 degrees C / 120 degrees F
FPC 17 MQ 2 Chip	OK	47 degrees C / 116 degrees F
FPC 17 MQ 3 TSen	OK	49 degrees C / 120 degrees F
FPC 17 MQ 3 Chip	OK	51 degrees C / 123 degrees F
FPC 18 Intake	OK	44 degrees C / 111 degrees F
FPC 18 Exhaust A	OK	53 degrees C / 127 degrees F
FPC 18 Exhaust B	OK	57 degrees C / 134 degrees F
FPC 18 LU 0 TSen	OK	56 degrees C / 132 degrees F
FPC 18 LU 0 Chip	OK	57 degrees C / 134 degrees F
FPC 18 LU 1 TSen	OK	56 degrees C / 132 degrees F
FPC 18 LU 1 Chip	OK	62 degrees C / 143 degrees F
FPC 18 LU 2 TSen	OK	56 degrees C / 132 degrees F
FPC 18 LU 2 Chip	OK	53 degrees C / 127 degrees F
FPC 18 LU 3 TSen	OK	56 degrees C / 132 degrees F
FPC 18 LU 3 Chip	OK	55 degrees C / 131 degrees F
FPC 18 MQ 0 TSen	OK	51 degrees C / 123 degrees F
FPC 18 MQ 0 Chip	OK	54 degrees C / 129 degrees F
FPC 18 MQ 1 TSen	OK	51 degrees C / 123 degrees F
FPC 18 MQ 1 Chip	OK	58 degrees C / 136 degrees F
FPC 18 MQ 2 TSen	OK	51 degrees C / 123 degrees F
FPC 18 MQ 2 Chip	OK	50 degrees C / 122 degrees F
FPC 18 MQ 3 TSen	OK	51 degrees C / 123 degrees F
FPC 18 MQ 3 Chip	OK	53 degrees C / 127 degrees F
FPC 19 Intake	OK	48 degrees C / 118 degrees F
FPC 19 Exhaust A	OK	56 degrees C / 132 degrees F
FPC 19 Exhaust B	OK	64 degrees C / 147 degrees F
FPC 19 LU 0 TSen	OK	63 degrees C / 145 degrees F
FPC 19 LU 0 Chip	OK	64 degrees C / 147 degrees F
FPC 19 LU 1 TSen	OK	63 degrees C / 145 degrees F
FPC 19 LU 1 Chip	OK	70 degrees C / 158 degrees F
FPC 19 LU 2 TSen	OK	63 degrees C / 145 degrees F
FPC 19 LU 2 Chip	OK	61 degrees C / 141 degrees F
FPC 19 LU 3 TSen	OK	63 degrees C / 145 degrees F
FPC 19 LU 3 Chip	OK	62 degrees C / 143 degrees F
FPC 19 MQ 0 TSen	OK	56 degrees C / 132 degrees F
FPC 19 MQ 0 Chip	OK	60 degrees C / 140 degrees F
FPC 19 MQ 1 TSen	OK	56 degrees C / 132 degrees F
FPC 19 MQ 1 Chip	OK	62 degrees C / 143 degrees F
FPC 19 MQ 2 TSen	OK	56 degrees C / 132 degrees F
FPC 19 MQ 2 Chip	OK	56 degrees C / 132 degrees F
FPC 19 MQ 3 TSen	OK	56 degrees C / 132 degrees F
FPC 19 MQ 3 Chip	OK	57 degrees C / 134 degrees F
ADC 0 Intake	OK	40 degrees C / 104 degrees F
ADC 0 Exhaust	OK	52 degrees C / 125 degrees F

ADC 0 ADC-XF1	OK	59 degrees C / 138 degrees F
ADC 0 ADC-XF0	OK	66 degrees C / 150 degrees F
ADC 1 Intake	OK	38 degrees C / 100 degrees F
ADC 1 Exhaust	OK	50 degrees C / 122 degrees F
ADC 1 ADC-XF1	OK	59 degrees C / 138 degrees F
ADC 1 ADC-XF0	OK	63 degrees C / 145 degrees F
ADC 2 Intake	OK	37 degrees C / 98 degrees F
ADC 2 Exhaust	OK	52 degrees C / 125 degrees F
ADC 2 ADC-XF1	OK	53 degrees C / 127 degrees F
ADC 2 ADC-XF0	OK	61 degrees C / 141 degrees F
ADC 3 Intake	OK	40 degrees C / 104 degrees F
ADC 3 Exhaust	OK	51 degrees C / 123 degrees F
ADC 3 ADC-XF1	OK	61 degrees C / 141 degrees F
ADC 3 ADC-XF0	OK	64 degrees C / 147 degrees F
ADC 4 Intake	OK	39 degrees C / 102 degrees F
ADC 4 Exhaust	OK	51 degrees C / 123 degrees F
ADC 4 ADC-XF1	OK	60 degrees C / 140 degrees F
ADC 4 ADC-XF0	OK	63 degrees C / 145 degrees F
ADC 5 Intake	OK	38 degrees C / 100 degrees F
ADC 5 Exhaust	OK	54 degrees C / 129 degrees F
ADC 5 ADC-XF1	OK	56 degrees C / 132 degrees F
ADC 5 ADC-XF0	OK	67 degrees C / 152 degrees F
ADC 6 Intake	OK	39 degrees C / 102 degrees F
ADC 6 Exhaust	OK	52 degrees C / 125 degrees F
ADC 6 ADC-XF1	OK	59 degrees C / 138 degrees F
ADC 6 ADC-XF0	OK	66 degrees C / 150 degrees F
ADC 7 Intake	OK	39 degrees C / 102 degrees F
ADC 7 Exhaust	OK	54 degrees C / 129 degrees F
ADC 7 ADC-XF1	OK	62 degrees C / 143 degrees F
ADC 7 ADC-XF0	OK	70 degrees C / 158 degrees F
ADC 8 Intake	OK	39 degrees C / 102 degrees F
ADC 8 Exhaust	OK	52 degrees C / 125 degrees F
ADC 8 ADC-XF1	OK	61 degrees C / 141 degrees F
ADC 8 ADC-XF0	OK	65 degrees C / 149 degrees F
ADC 9 Intake	OK	41 degrees C / 105 degrees F
ADC 9 Exhaust	OK	51 degrees C / 123 degrees F
ADC 9 ADC-XF1	OK	63 degrees C / 145 degrees F
ADC 9 ADC-XF0	OK	63 degrees C / 145 degrees F
ADC 10 Intake	OK	48 degrees C / 118 degrees F
ADC 10 Exhaust	OK	53 degrees C / 127 degrees F
ADC 10 ADC-XF1	OK	67 degrees C / 152 degrees F
ADC 10 ADC-XF0	OK	66 degrees C / 150 degrees F
ADC 12 Intake	OK	49 degrees C / 120 degrees F
ADC 12 Exhaust	OK	54 degrees C / 129 degrees F
ADC 12 ADC-XF1	OK	67 degrees C / 152 degrees F
ADC 12 ADC-XF0	OK	67 degrees C / 152 degrees F
ADC 13 Intake	OK	49 degrees C / 120 degrees F
ADC 13 Exhaust	OK	57 degrees C / 134 degrees F
ADC 13 ADC-XF1	OK	66 degrees C / 150 degrees F
ADC 13 ADC-XF0	OK	69 degrees C / 156 degrees F
ADC 14 Intake	OK	51 degrees C / 123 degrees F
ADC 14 Exhaust	OK	59 degrees C / 138 degrees F
ADC 14 ADC-XF1	OK	69 degrees C / 156 degrees F
ADC 14 ADC-XF0	OK	74 degrees C / 165 degrees F
ADC 15 Intake	OK	50 degrees C / 122 degrees F
ADC 15 Exhaust	OK	59 degrees C / 138 degrees F
ADC 15 ADC-XF1	OK	68 degrees C / 154 degrees F
ADC 15 ADC-XF0	OK	69 degrees C / 156 degrees F
ADC 16 Intake	OK	52 degrees C / 125 degrees F
ADC 16 Exhaust	OK	58 degrees C / 136 degrees F
ADC 16 ADC-XF1	OK	68 degrees C / 154 degrees F

	ADC 16 ADC-XF0	OK	70 degrees C / 158 degrees F
	ADC 17 Intake	OK	52 degrees C / 125 degrees F
	ADC 17 Exhaust	OK	59 degrees C / 138 degrees F
	ADC 17 ADC-XF1	OK	69 degrees C / 156 degrees F
	ADC 17 ADC-XF0	OK	71 degrees C / 159 degrees F
	ADC 18 Intake	OK	53 degrees C / 127 degrees F
	ADC 18 Exhaust	OK	59 degrees C / 138 degrees F
	ADC 18 ADC-XF1	OK	68 degrees C / 154 degrees F
	ADC 18 ADC-XF0	OK	73 degrees C / 163 degrees F
	ADC 19 Intake	OK	50 degrees C / 122 degrees F
	ADC 19 Exhaust	OK	59 degrees C / 138 degrees F
	ADC 19 ADC-XF1	OK	68 degrees C / 154 degrees F
	ADC 19 ADC-XF0	OK	72 degrees C / 161 degrees F
Fans	Fan Tray 0 Fan 1	OK	7440 RPM
	Fan Tray 0 Fan 2	OK	7200 RPM
	Fan Tray 0 Fan 3	OK	6960 RPM
	Fan Tray 0 Fan 4	OK	7200 RPM
	Fan Tray 0 Fan 5	OK	7080 RPM
	Fan Tray 0 Fan 6	OK	6840 RPM
	Fan Tray 1 Fan 1	OK	6840 RPM
	Fan Tray 1 Fan 2	OK	6960 RPM
	Fan Tray 1 Fan 3	OK	6960 RPM
	Fan Tray 1 Fan 4	OK	7080 RPM
	Fan Tray 1 Fan 5	OK	6960 RPM
	Fan Tray 1 Fan 6	OK	6960 RPM
	Fan Tray 2 Fan 1	OK	8640 RPM
	Fan Tray 2 Fan 2	OK	8640 RPM
	Fan Tray 2 Fan 3	OK	8760 RPM
	Fan Tray 2 Fan 4	OK	8760 RPM
	Fan Tray 2 Fan 5	OK	8640 RPM
	Fan Tray 2 Fan 6	OK	8640 RPM
	Fan Tray 3 Fan 1	OK	8520 RPM
	Fan Tray 3 Fan 2	OK	8520 RPM
	Fan Tray 3 Fan 3	OK	8640 RPM
	Fan Tray 3 Fan 4	OK	8640 RPM
	Fan Tray 3 Fan 5	OK	8520 RPM
	Fan Tray 3 Fan 6	OK	8520 RPM

show chassis environment (MX2020 Router with MPC5EQ and MPC6E)

Class	Item	Status	Measurement
Temp	PSM 0	OK	32 degrees C / 89 degrees F
	PSM 1	OK	32 degrees C / 89 degrees F
	PSM 2	OK	32 degrees C / 89 degrees F
	PSM 3	OK	32 degrees C / 89 degrees F
	PSM 4	OK	32 degrees C / 89 degrees F
	PSM 5	OK	33 degrees C / 91 degrees F
	PSM 6	OK	32 degrees C / 89 degrees F
	PSM 7	OK	32 degrees C / 89 degrees F
	PSM 8	OK	32 degrees C / 89 degrees F
	PSM 9	Absent	
	PSM 10	Absent	
	PSM 11	Absent	
	PSM 12	OK	33 degrees C / 91 degrees F
	PSM 13	OK	33 degrees C / 91 degrees F
	PSM 14	OK	34 degrees C / 93 degrees F
	PSM 15	OK	34 degrees C / 93 degrees F
	PSM 16	OK	33 degrees C / 91 degrees F
	PSM 17	OK	33 degrees C / 91 degrees F
	PDM 0	OK	
	PDM 1	OK	

PDM 2	OK	
PDM 3	OK	
CB 0 IntakeA-Zone0	OK	34 degrees C / 93 degrees F
CB 0 IntakeB-Zone1	OK	26 degrees C / 78 degrees F
CB 0 IntakeC-Zone0	OK	38 degrees C / 100 degrees F
CB 0 ExhaustA-Zone0	OK	34 degrees C / 93 degrees F
CB 0 ExhaustB-Zone1	OK	27 degrees C / 80 degrees F
CB 0 TCBC-Zone0	OK	32 degrees C / 89 degrees F
CB 1 IntakeA-Zone0	OK	24 degrees C / 75 degrees F
CB 1 IntakeB-Zone1	OK	22 degrees C / 71 degrees F
CB 1 IntakeC-Zone0	OK	34 degrees C / 93 degrees F
CB 1 ExhaustA-Zone0	OK	31 degrees C / 87 degrees F
CB 1 ExhaustB-Zone1	OK	24 degrees C / 75 degrees F
CB 1 TCBC-Zone0	OK	27 degrees C / 80 degrees F
SPMB 0 Intake	OK	25 degrees C / 77 degrees F
SPMB 1 Intake	OK	23 degrees C / 73 degrees F
Routing Engine 0	OK	28 degrees C / 82 degrees F
Routing Engine 0 CPU	OK	25 degrees C / 77 degrees F
Routing Engine 1	OK	25 degrees C / 77 degrees F
Routing Engine 1 CPU	OK	24 degrees C / 75 degrees F
SFB 0 Intake-Zone0	OK	45 degrees C / 113 degrees F
SFB 0 Exhaust-Zone1	OK	34 degrees C / 93 degrees F
SFB 0 IntakeA-Zone0	OK	32 degrees C / 89 degrees F
SFB 0 IntakeB-Zone1	OK	28 degrees C / 82 degrees F
SFB 0 Exhaust-Zone0	OK	36 degrees C / 96 degrees F
SFB 0 SFB-XF2-Zone1	OK	46 degrees C / 114 degrees F
SFB 0 SFB-XF1-Zone0	OK	48 degrees C / 118 degrees F
SFB 0 SFB-XF0-Zone0	OK	60 degrees C / 140 degrees F
SFB 1 Intake-Zone0	OK	44 degrees C / 111 degrees F
SFB 1 Exhaust-Zone1	OK	34 degrees C / 93 degrees F
SFB 1 IntakeA-Zone0	OK	35 degrees C / 95 degrees F
SFB 1 IntakeB-Zone1	OK	27 degrees C / 80 degrees F
SFB 1 Exhaust-Zone0	OK	37 degrees C / 98 degrees F
SFB 1 SFB-XF2-Zone1	OK	47 degrees C / 116 degrees F
SFB 1 SFB-XF1-Zone0	OK	49 degrees C / 120 degrees F
SFB 1 SFB-XF0-Zone0	OK	56 degrees C / 132 degrees F
SFB 2 Intake-Zone0	OK	41 degrees C / 105 degrees F
SFB 2 Exhaust-Zone1	OK	34 degrees C / 93 degrees F
SFB 2 IntakeA-Zone0	OK	35 degrees C / 95 degrees F
SFB 2 IntakeB-Zone1	OK	28 degrees C / 82 degrees F
SFB 2 Exhaust-Zone0	OK	37 degrees C / 98 degrees F
SFB 2 SFB-XF2-Zone1	OK	47 degrees C / 116 degrees F
SFB 2 SFB-XF1-Zone0	OK	55 degrees C / 131 degrees F
SFB 2 SFB-XF0-Zone0	OK	55 degrees C / 131 degrees F
SFB 3 Intake-Zone0	OK	43 degrees C / 109 degrees F
SFB 3 Exhaust-Zone1	OK	33 degrees C / 91 degrees F
SFB 3 IntakeA-Zone0	OK	35 degrees C / 95 degrees F
SFB 3 IntakeB-Zone1	OK	27 degrees C / 80 degrees F
SFB 3 Exhaust-Zone0	OK	36 degrees C / 96 degrees F
SFB 3 SFB-XF2-Zone1	OK	46 degrees C / 114 degrees F
SFB 3 SFB-XF1-Zone0	OK	46 degrees C / 114 degrees F
SFB 3 SFB-XF0-Zone0	OK	57 degrees C / 134 degrees F
SFB 4 Intake-Zone0	OK	36 degrees C / 96 degrees F
SFB 4 Exhaust-Zone1	OK	32 degrees C / 89 degrees F
SFB 4 IntakeA-Zone0	OK	31 degrees C / 87 degrees F
SFB 4 IntakeB-Zone1	OK	26 degrees C / 78 degrees F
SFB 4 Exhaust-Zone0	OK	32 degrees C / 89 degrees F
SFB 4 SFB-XF2-Zone1	OK	44 degrees C / 111 degrees F
SFB 4 SFB-XF1-Zone0	OK	45 degrees C / 113 degrees F
SFB 4 SFB-XF0-Zone0	OK	52 degrees C / 125 degrees F
SFB 5 Intake-Zone0	OK	31 degrees C / 87 degrees F

SFB 5 Exhaust-Zone1	OK	30 degrees C / 86 degrees F
SFB 5 IntakeA-Zone0	OK	26 degrees C / 78 degrees F
SFB 5 IntakeB-Zone1	OK	24 degrees C / 75 degrees F
SFB 5 Exhaust-Zone0	OK	29 degrees C / 84 degrees F
SFB 5 SFB-XF2-Zone1	OK	43 degrees C / 109 degrees F
SFB 5 SFB-XF1-Zone0	OK	47 degrees C / 116 degrees F
SFB 5 SFB-XF0-Zone0	OK	49 degrees C / 120 degrees F
SFB 6 Intake-Zone0	OK	30 degrees C / 86 degrees F
SFB 6 Exhaust-Zone1	OK	29 degrees C / 84 degrees F
SFB 6 IntakeA-Zone0	OK	25 degrees C / 77 degrees F
SFB 6 IntakeB-Zone1	OK	24 degrees C / 75 degrees F
SFB 6 Exhaust-Zone0	OK	29 degrees C / 84 degrees F
SFB 6 SFB-XF2-Zone1	OK	43 degrees C / 109 degrees F
SFB 6 SFB-XF1-Zone0	OK	44 degrees C / 111 degrees F
SFB 6 SFB-XF0-Zone0	OK	45 degrees C / 113 degrees F
SFB 7 Intake-Zone0	OK	31 degrees C / 87 degrees F
SFB 7 Exhaust-Zone1	OK	30 degrees C / 86 degrees F
SFB 7 IntakeA-Zone0	OK	26 degrees C / 78 degrees F
SFB 7 IntakeB-Zone1	OK	24 degrees C / 75 degrees F
SFB 7 Exhaust-Zone0	OK	28 degrees C / 82 degrees F
SFB 7 SFB-XF2-Zone1	OK	50 degrees C / 122 degrees F
SFB 7 SFB-XF1-Zone0	OK	43 degrees C / 109 degrees F
SFB 7 SFB-XF0-Zone0	OK	47 degrees C / 116 degrees F
FPC 0 Intake	OK	31 degrees C / 87 degrees F
FPC 0 Exhaust A	OK	49 degrees C / 120 degrees F
FPC 0 Exhaust B	OK	43 degrees C / 109 degrees F
FPC 0 XL TSen	OK	42 degrees C / 107 degrees F
FPC 0 XL Chip	OK	46 degrees C / 114 degrees F
FPC 0 XL_XR0 TSen	OK	42 degrees C / 107 degrees F
FPC 0 XL_XR0 Chip	OK	48 degrees C / 118 degrees F
FPC 0 XL_XR1 TSen	OK	42 degrees C / 107 degrees F
FPC 0 XL_XR1 Chip	OK	48 degrees C / 118 degrees F
FPC 0 XQ TSen	OK	42 degrees C / 107 degrees F
FPC 0 XQ Chip	OK	44 degrees C / 111 degrees F
FPC 0 XQ_XR0 TSen	OK	42 degrees C / 107 degrees F
FPC 0 XQ_XR0 Chip	OK	57 degrees C / 134 degrees F
FPC 0 XQ_XR1 TSen	OK	42 degrees C / 107 degrees F
FPC 0 XQ_XR1 Chip	OK	55 degrees C / 131 degrees F
FPC 0 XM 0 TSen	OK	48 degrees C / 118 degrees F
FPC 0 XM 0 Chip	OK	62 degrees C / 143 degrees F
FPC 0 XM 1 TSen	OK	48 degrees C / 118 degrees F
FPC 0 XM 1 Chip	OK	44 degrees C / 111 degrees F
FPC 0 PLX PCIe Switch TSe	OK	48 degrees C / 118 degrees F
FPC 0 PLX PCIe Switch Chi	OK	57 degrees C / 134 degrees F
FPC 1 Intake	OK	29 degrees C / 84 degrees F
FPC 1 Exhaust A	OK	36 degrees C / 96 degrees F
FPC 1 Exhaust B	OK	44 degrees C / 111 degrees F
FPC 1 LU 0 TSen	OK	38 degrees C / 100 degrees F
FPC 1 LU 0 Chip	OK	45 degrees C / 113 degrees F
FPC 1 LU 1 TSen	OK	38 degrees C / 100 degrees F
FPC 1 LU 1 Chip	OK	38 degrees C / 100 degrees F
FPC 1 LU 2 TSen	OK	38 degrees C / 100 degrees F
FPC 1 LU 2 Chip	OK	42 degrees C / 107 degrees F
FPC 1 LU 3 TSen	OK	38 degrees C / 100 degrees F
FPC 1 LU 3 Chip	OK	47 degrees C / 116 degrees F
FPC 1 XM 0 TSen	OK	38 degrees C / 100 degrees F
FPC 1 XM 0 Chip	OK	44 degrees C / 111 degrees F
FPC 1 XF 0 TSen	OK	38 degrees C / 100 degrees F
FPC 1 XF 0 Chip	OK	54 degrees C / 129 degrees F
FPC 1 PLX Switch TSen	OK	38 degrees C / 100 degrees F
FPC 1 PLX Switch Chip	OK	41 degrees C / 105 degrees F

FPC 2 Intake	OK	28 degrees C / 82 degrees F
FPC 2 Exhaust A	OK	28 degrees C / 82 degrees F
FPC 2 Exhaust B	OK	28 degrees C / 82 degrees F
FPC 2 LU 0 TSen	OK	40 degrees C / 104 degrees F
FPC 2 LU 0 Chip	OK	40 degrees C / 104 degrees F
FPC 2 LU 1 TSen	OK	40 degrees C / 104 degrees F
FPC 2 LU 1 Chip	OK	41 degrees C / 105 degrees F
FPC 2 LU 2 TSen	OK	40 degrees C / 104 degrees F
FPC 2 LU 2 Chip	OK	34 degrees C / 93 degrees F
FPC 2 LU 3 TSen	OK	40 degrees C / 104 degrees F
FPC 2 LU 3 Chip	OK	38 degrees C / 100 degrees F
FPC 2 XM 0 TSen	OK	40 degrees C / 104 degrees F
FPC 2 XM 0 Chip	OK	47 degrees C / 116 degrees F
FPC 2 XM 1 TSen	OK	40 degrees C / 104 degrees F
FPC 2 XM 1 Chip	OK	42 degrees C / 107 degrees F
FPC 2 PLX Switch TSen	OK	40 degrees C / 104 degrees F
FPC 2 PLX Switch Chip	OK	39 degrees C / 102 degrees F
FPC 3 Intake	OK	27 degrees C / 80 degrees F
FPC 3 Exhaust A	OK	38 degrees C / 100 degrees F
FPC 3 Exhaust B	OK	31 degrees C / 87 degrees F
FPC 3 QX 0 TSen	OK	38 degrees C / 100 degrees F
FPC 3 QX 0 Chip	OK	42 degrees C / 107 degrees F
FPC 3 LU 0 TCAM TSen	OK	38 degrees C / 100 degrees F
FPC 3 LU 0 TCAM Chip	OK	43 degrees C / 109 degrees F
FPC 3 LU 0 TSen	OK	38 degrees C / 100 degrees F
FPC 3 LU 0 Chip	OK	42 degrees C / 107 degrees F
FPC 3 MQ 0 TSen	OK	38 degrees C / 100 degrees F
FPC 3 MQ 0 Chip	OK	39 degrees C / 102 degrees F
FPC 3 QX 1 TSen	OK	32 degrees C / 89 degrees F
FPC 3 QX 1 Chip	OK	36 degrees C / 96 degrees F
FPC 3 LU 1 TCAM TSen	OK	32 degrees C / 89 degrees F
FPC 3 LU 1 TCAM Chip	OK	35 degrees C / 95 degrees F
FPC 3 LU 1 TSen	OK	32 degrees C / 89 degrees F
FPC 3 LU 1 Chip	OK	37 degrees C / 98 degrees F
FPC 3 MQ 1 TSen	OK	32 degrees C / 89 degrees F
FPC 3 MQ 1 Chip	OK	36 degrees C / 96 degrees F
FPC 4 Intake	OK	29 degrees C / 84 degrees F
FPC 4 Exhaust A	OK	36 degrees C / 96 degrees F
FPC 4 Exhaust B	OK	40 degrees C / 104 degrees F
FPC 4 XL TSen	OK	39 degrees C / 102 degrees F
FPC 4 XL Chip	OK	42 degrees C / 107 degrees F
FPC 4 XL_XR0 TSen	OK	39 degrees C / 102 degrees F
FPC 4 XL_XR0 Chip	OK	45 degrees C / 113 degrees F
FPC 4 XL_XR1 TSen	OK	39 degrees C / 102 degrees F
FPC 4 XL_XR1 Chip	OK	46 degrees C / 114 degrees F
FPC 4 XQ TSen	OK	39 degrees C / 102 degrees F
FPC 4 XQ Chip	OK	42 degrees C / 107 degrees F
FPC 4 XQ_XR0 TSen	OK	39 degrees C / 102 degrees F
FPC 4 XQ_XR0 Chip	OK	54 degrees C / 129 degrees F
FPC 4 XQ_XR1 TSen	OK	39 degrees C / 102 degrees F
FPC 4 XQ_XR1 Chip	OK	53 degrees C / 127 degrees F
FPC 4 XM 0 TSen	OK	45 degrees C / 113 degrees F
FPC 4 XM 0 Chip	OK	59 degrees C / 138 degrees F
FPC 4 XM 1 TSen	OK	45 degrees C / 113 degrees F
FPC 4 XM 1 Chip	OK	41 degrees C / 105 degrees F
FPC 4 PLX PCIe Switch TSe	OK	45 degrees C / 113 degrees F
FPC 4 PLX PCIe Switch Chi	OK	58 degrees C / 136 degrees F
FPC 5 Intake	OK	29 degrees C / 84 degrees F
FPC 5 Exhaust A	OK	33 degrees C / 91 degrees F
FPC 5 Exhaust B	OK	39 degrees C / 102 degrees F
FPC 5 LU 0 TSen	OK	40 degrees C / 104 degrees F

FPC 5 LU 0 Chip	OK	40 degrees C / 104 degrees F
FPC 5 LU 1 TSen	OK	40 degrees C / 104 degrees F
FPC 5 LU 1 Chip	OK	45 degrees C / 113 degrees F
FPC 5 LU 2 TSen	OK	40 degrees C / 104 degrees F
FPC 5 LU 2 Chip	OK	40 degrees C / 104 degrees F
FPC 5 LU 3 TSen	OK	40 degrees C / 104 degrees F
FPC 5 LU 3 Chip	OK	46 degrees C / 114 degrees F
FPC 5 MQ 0 TSen	OK	32 degrees C / 89 degrees F
FPC 5 MQ 0 Chip	OK	33 degrees C / 91 degrees F
FPC 5 MQ 1 TSen	OK	32 degrees C / 89 degrees F
FPC 5 MQ 1 Chip	OK	35 degrees C / 95 degrees F
FPC 5 MQ 2 TSen	OK	32 degrees C / 89 degrees F
FPC 5 MQ 2 Chip	OK	32 degrees C / 89 degrees F
FPC 5 MQ 3 TSen	OK	32 degrees C / 89 degrees F
FPC 5 MQ 3 Chip	OK	32 degrees C / 89 degrees F
FPC 9 Intake	OK	25 degrees C / 77 degrees F
FPC 9 Exhaust A	OK	37 degrees C / 98 degrees F
FPC 9 Exhaust B	OK	40 degrees C / 104 degrees F
FPC 9 XL 0 TSen	OK	40 degrees C / 104 degrees F

...

show chassis environment (MX2010 Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Temp	PSM 0	OK	7 degrees C / 44 degrees F
	PSM 1	OK	7 degrees C / 44 degrees F
	PSM 2	OK	7 degrees C / 44 degrees F
	PSM 3	OK	6 degrees C / 42 degrees F
	PSM 4	OK	6 degrees C / 42 degrees F
	PSM 5	OK	6 degrees C / 42 degrees F
	PSM 6	OK	6 degrees C / 42 degrees F
	PSM 7	OK	7 degrees C / 44 degrees F
	PSM 8	OK	7 degrees C / 44 degrees F
	PDM 0	OK	
	PDM 1	Absent	
	CB 0 IntakeA-Zone0	OK	14 degrees C / 57 degrees F
	CB 0 IntakeB-Zone1	OK	7 degrees C / 44 degrees F
	CB 0 IntakeC-Zone0	OK	22 degrees C / 71 degrees F
	CB 0 ExhaustA-Zone0	OK	14 degrees C / 57 degrees F
	CB 0 ExhaustB-Zone1	OK	9 degrees C / 48 degrees F
	CB 0 TCBC-Zone0	OK	11 degrees C / 51 degrees F
	CB 1 IntakeA-Zone0	OK	9 degrees C / 48 degrees F
	CB 1 IntakeB-Zone1	OK	5 degrees C / 41 degrees F
	CB 1 IntakeC-Zone0	OK	20 degrees C / 68 degrees F
	CB 1 ExhaustA-Zone0	OK	12 degrees C / 53 degrees F
	CB 1 ExhaustB-Zone1	OK	7 degrees C / 44 degrees F
	CB 1 TCBC-Zone0	OK	10 degrees C / 50 degrees F
	SPMB 0 Intake	OK	5 degrees C / 41 degrees F
	SPMB 1 Intake	OK	4 degrees C / 39 degrees F
	Routing Engine 0	OK	9 degrees C / 48 degrees F
	Routing Engine 0 CPU	OK	9 degrees C / 48 degrees F
	Routing Engine 1	OK	6 degrees C / 42 degrees F
	Routing Engine 1 CPU	OK	6 degrees C / 42 degrees F
	SFB 0 Intake-Zone0	OK	26 degrees C / 78 degrees F
	SFB 0 Exhaust-Zone1	OK	17 degrees C / 62 degrees F
	SFB 0 IntakeA-Zone0	OK	16 degrees C / 60 degrees F
	SFB 0 IntakeB-Zone1	OK	11 degrees C / 51 degrees F
	SFB 0 Exhaust-Zone0	OK	18 degrees C / 64 degrees F
	SFB 0 SFB-XF2-Zone1	OK	25 degrees C / 77 degrees F
	SFB 0 SFB-XF1-Zone0	OK	23 degrees C / 73 degrees F

SFB 0 SFB-XF0-Zone0	OK	33 degrees C / 91 degrees F
SFB 1 Intake-Zone0	OK	27 degrees C / 80 degrees F
SFB 1 Exhaust-Zone1	OK	15 degrees C / 59 degrees F
SFB 1 IntakeA-Zone0	OK	20 degrees C / 68 degrees F
SFB 1 IntakeB-Zone1	OK	10 degrees C / 50 degrees F
SFB 1 Exhaust-Zone0	OK	19 degrees C / 66 degrees F
SFB 1 SFB-XF2-Zone1	OK	26 degrees C / 78 degrees F
SFB 1 SFB-XF1-Zone0	OK	27 degrees C / 80 degrees F
SFB 1 SFB-XF0-Zone0	OK	32 degrees C / 89 degrees F
SFB 2 Intake-Zone0	OK	21 degrees C / 69 degrees F
SFB 2 Exhaust-Zone1	OK	13 degrees C / 55 degrees F
SFB 2 IntakeA-Zone0	OK	18 degrees C / 64 degrees F
SFB 2 IntakeB-Zone1	OK	9 degrees C / 48 degrees F
SFB 2 Exhaust-Zone0	OK	16 degrees C / 60 degrees F
SFB 2 SFB-XF2-Zone1	OK	24 degrees C / 75 degrees F
SFB 2 SFB-XF1-Zone0	OK	21 degrees C / 69 degrees F
SFB 2 SFB-XF0-Zone0	OK	26 degrees C / 78 degrees F
SFB 4 Intake-Zone0	OK	28 degrees C / 82 degrees F
SFB 4 Exhaust-Zone1	OK	16 degrees C / 60 degrees F
SFB 4 IntakeA-Zone0	OK	18 degrees C / 64 degrees F
SFB 4 IntakeB-Zone1	OK	11 degrees C / 51 degrees F
SFB 4 Exhaust-Zone0	OK	19 degrees C / 66 degrees F
SFB 4 SFB-XF2-Zone1	OK	27 degrees C / 80 degrees F
SFB 4 SFB-XF1-Zone0	OK	27 degrees C / 80 degrees F
SFB 4 SFB-XF0-Zone0	OK	32 degrees C / 89 degrees F
SFB 5 Intake-Zone0	OK	22 degrees C / 71 degrees F
SFB 5 Exhaust-Zone1	OK	14 degrees C / 57 degrees F
SFB 5 IntakeA-Zone0	OK	18 degrees C / 64 degrees F
SFB 5 IntakeB-Zone1	OK	10 degrees C / 50 degrees F
SFB 5 Exhaust-Zone0	OK	17 degrees C / 62 degrees F
SFB 5 SFB-XF2-Zone1	OK	22 degrees C / 71 degrees F
SFB 5 SFB-XF1-Zone0	OK	29 degrees C / 84 degrees F
SFB 5 SFB-XF0-Zone0	OK	27 degrees C / 80 degrees F
SFB 6 Intake-Zone0	OK	27 degrees C / 80 degrees F
SFB 6 Exhaust-Zone1	OK	13 degrees C / 55 degrees F
SFB 6 IntakeA-Zone0	OK	19 degrees C / 66 degrees F
SFB 6 IntakeB-Zone1	OK	10 degrees C / 50 degrees F
SFB 6 Exhaust-Zone0	OK	20 degrees C / 68 degrees F
SFB 6 SFB-XF2-Zone1	OK	24 degrees C / 75 degrees F
SFB 6 SFB-XF1-Zone0	OK	32 degrees C / 89 degrees F
SFB 6 SFB-XF0-Zone0	OK	33 degrees C / 91 degrees F
SFB 7 Intake-Zone0	OK	25 degrees C / 77 degrees F
SFB 7 Exhaust-Zone1	OK	13 degrees C / 55 degrees F
SFB 7 IntakeA-Zone0	OK	14 degrees C / 57 degrees F
SFB 7 IntakeB-Zone1	OK	8 degrees C / 46 degrees F
SFB 7 Exhaust-Zone0	OK	17 degrees C / 62 degrees F
SFB 7 SFB-XF2-Zone1	OK	21 degrees C / 69 degrees F
SFB 7 SFB-XF1-Zone0	OK	21 degrees C / 69 degrees F
SFB 7 SFB-XF0-Zone0	OK	33 degrees C / 91 degrees F
FPC 0 Intake	OK	13 degrees C / 55 degrees F
FPC 0 Exhaust A	OK	13 degrees C / 55 degrees F
FPC 0 Exhaust B	OK	14 degrees C / 57 degrees F
FPC 0 LU 0 TSen	OK	28 degrees C / 82 degrees F
FPC 0 LU 0 Chip	OK	25 degrees C / 77 degrees F
FPC 0 LU 1 TSen	OK	28 degrees C / 82 degrees F
FPC 0 LU 1 Chip	OK	27 degrees C / 80 degrees F
FPC 0 LU 2 TSen	OK	28 degrees C / 82 degrees F
FPC 0 LU 2 Chip	OK	19 degrees C / 66 degrees F
FPC 0 LU 3 TSen	OK	28 degrees C / 82 degrees F
FPC 0 LU 3 Chip	OK	23 degrees C / 73 degrees F
FPC 0 XM 0 TSen	OK	28 degrees C / 82 degrees F

FPC 0 XM 0 Chip	OK	33 degrees C / 91 degrees F
FPC 0 XM 1 TSen	OK	28 degrees C / 82 degrees F
FPC 0 XM 1 Chip	OK	26 degrees C / 78 degrees F
FPC 0 PLX Switch TSen	OK	28 degrees C / 82 degrees F
FPC 0 PLX Switch Chip	OK	26 degrees C / 78 degrees F
FPC 1 Intake	OK	10 degrees C / 50 degrees F
FPC 1 Exhaust A	OK	24 degrees C / 75 degrees F
FPC 1 Exhaust B	OK	28 degrees C / 82 degrees F
FPC 1 LU 0 TSen	OK	22 degrees C / 71 degrees F
FPC 1 LU 0 Chip	OK	31 degrees C / 87 degrees F
FPC 1 LU 1 TSen	OK	22 degrees C / 71 degrees F
FPC 1 LU 1 Chip	OK	21 degrees C / 69 degrees F
FPC 1 LU 2 TSen	OK	22 degrees C / 71 degrees F
FPC 1 LU 2 Chip	OK	25 degrees C / 77 degrees F
FPC 1 LU 3 TSen	OK	22 degrees C / 71 degrees F
FPC 1 LU 3 Chip	OK	33 degrees C / 91 degrees F
FPC 1 XM 0 TSen	OK	22 degrees C / 71 degrees F
FPC 1 XM 0 Chip	OK	30 degrees C / 86 degrees F
FPC 1 XF 0 TSen	OK	22 degrees C / 71 degrees F
FPC 1 XF 0 Chip	OK	37 degrees C / 98 degrees F
FPC 1 PLX Switch TSen	OK	22 degrees C / 71 degrees F
FPC 1 PLX Switch Chip	OK	22 degrees C / 71 degrees F
FPC 2 Intake	OK	9 degrees C / 48 degrees F
FPC 2 Exhaust A	OK	10 degrees C / 50 degrees F
FPC 2 Exhaust B	OK	10 degrees C / 50 degrees F
FPC 2 LU 0 TSen	OK	26 degrees C / 78 degrees F
FPC 2 LU 0 Chip	OK	25 degrees C / 77 degrees F
FPC 2 LU 1 TSen	OK	26 degrees C / 78 degrees F
FPC 2 LU 1 Chip	OK	26 degrees C / 78 degrees F
FPC 2 LU 2 TSen	OK	26 degrees C / 78 degrees F
FPC 2 LU 2 Chip	OK	17 degrees C / 62 degrees F
FPC 2 LU 3 TSen	OK	26 degrees C / 78 degrees F
FPC 2 LU 3 Chip	OK	22 degrees C / 71 degrees F
FPC 2 XM 0 TSen	OK	26 degrees C / 78 degrees F
FPC 2 XM 0 Chip	OK	34 degrees C / 93 degrees F
FPC 2 XM 1 TSen	OK	26 degrees C / 78 degrees F
FPC 2 XM 1 Chip	OK	26 degrees C / 78 degrees F
FPC 2 PLX Switch TSen	OK	26 degrees C / 78 degrees F
FPC 2 PLX Switch Chip	OK	20 degrees C / 68 degrees F
FPC 3 Intake	OK	12 degrees C / 53 degrees F
FPC 3 Exhaust A	OK	16 degrees C / 60 degrees F
FPC 3 Exhaust B	OK	26 degrees C / 78 degrees F
FPC 3 LU 0 TSen	OK	23 degrees C / 73 degrees F
FPC 3 LU 0 Chip	OK	26 degrees C / 78 degrees F
FPC 3 LU 1 TSen	OK	23 degrees C / 73 degrees F
FPC 3 LU 1 Chip	OK	27 degrees C / 80 degrees F
FPC 3 LU 2 TSen	OK	23 degrees C / 73 degrees F
FPC 3 LU 2 Chip	OK	22 degrees C / 71 degrees F
FPC 3 LU 3 TSen	OK	23 degrees C / 73 degrees F
FPC 3 LU 3 Chip	OK	21 degrees C / 69 degrees F
FPC 3 MQ 0 TSen	OK	15 degrees C / 59 degrees F
FPC 3 MQ 0 Chip	OK	18 degrees C / 64 degrees F
FPC 3 MQ 1 TSen	OK	15 degrees C / 59 degrees F
FPC 3 MQ 1 Chip	OK	20 degrees C / 68 degrees F
FPC 3 MQ 2 TSen	OK	15 degrees C / 59 degrees F
FPC 3 MQ 2 Chip	OK	17 degrees C / 62 degrees F
FPC 3 MQ 3 TSen	OK	15 degrees C / 59 degrees F
FPC 3 MQ 3 Chip	OK	16 degrees C / 60 degrees F
FPC 4 Intake	OK	11 degrees C / 51 degrees F
FPC 4 Exhaust A	OK	22 degrees C / 71 degrees F
FPC 4 Exhaust B	OK	28 degrees C / 82 degrees F

FPC 4 LU 0 TSen	OK	22 degrees C / 71 degrees F
FPC 4 LU 0 Chip	OK	33 degrees C / 91 degrees F
FPC 4 LU 1 TSen	OK	22 degrees C / 71 degrees F
FPC 4 LU 1 Chip	OK	21 degrees C / 69 degrees F
FPC 4 LU 2 TSen	OK	22 degrees C / 71 degrees F
FPC 4 LU 2 Chip	OK	26 degrees C / 78 degrees F
FPC 4 LU 3 TSen	OK	22 degrees C / 71 degrees F
FPC 4 LU 3 Chip	OK	33 degrees C / 91 degrees F
FPC 4 XM 0 TSen	OK	22 degrees C / 71 degrees F
FPC 4 XM 0 Chip	OK	30 degrees C / 86 degrees F
FPC 4 XF 0 TSen	OK	22 degrees C / 71 degrees F
FPC 4 XF 0 Chip	OK	37 degrees C / 98 degrees F
FPC 4 PLX Switch TSen	OK	22 degrees C / 71 degrees F
FPC 4 PLX Switch Chip	OK	23 degrees C / 73 degrees F
FPC 5 Intake	OK	12 degrees C / 53 degrees F
FPC 5 Exhaust A	OK	12 degrees C / 53 degrees F
FPC 5 Exhaust B	OK	12 degrees C / 53 degrees F
FPC 5 LU 0 TSen	OK	27 degrees C / 80 degrees F
FPC 5 LU 0 Chip	OK	28 degrees C / 82 degrees F
FPC 5 LU 1 TSen	OK	27 degrees C / 80 degrees F
FPC 5 LU 1 Chip	OK	27 degrees C / 80 degrees F
FPC 5 LU 2 TSen	OK	27 degrees C / 80 degrees F
FPC 5 LU 2 Chip	OK	19 degrees C / 66 degrees F
FPC 5 LU 3 TSen	OK	27 degrees C / 80 degrees F
FPC 5 LU 3 Chip	OK	22 degrees C / 71 degrees F
FPC 5 XM 0 TSen	OK	27 degrees C / 80 degrees F
FPC 5 XM 0 Chip	OK	36 degrees C / 96 degrees F
FPC 5 XM 1 TSen	OK	27 degrees C / 80 degrees F
FPC 5 XM 1 Chip	OK	26 degrees C / 78 degrees F
FPC 5 PLX Switch TSen	OK	27 degrees C / 80 degrees F
FPC 5 PLX Switch Chip	OK	24 degrees C / 75 degrees F
FPC 6 Intake	OK	12 degrees C / 53 degrees F
FPC 6 Exhaust A	OK	17 degrees C / 62 degrees F
FPC 6 Exhaust B	OK	28 degrees C / 82 degrees F
FPC 6 LU 0 TSen	OK	24 degrees C / 75 degrees F
FPC 6 LU 0 Chip	OK	29 degrees C / 84 degrees F
FPC 6 LU 1 TSen	OK	24 degrees C / 75 degrees F
FPC 6 LU 1 Chip	OK	30 degrees C / 86 degrees F
FPC 6 LU 2 TSen	OK	24 degrees C / 75 degrees F
FPC 6 LU 2 Chip	OK	24 degrees C / 75 degrees F
FPC 6 LU 3 TSen	OK	24 degrees C / 75 degrees F
FPC 6 LU 3 Chip	OK	22 degrees C / 71 degrees F
FPC 6 MQ 0 TSen	OK	16 degrees C / 60 degrees F
FPC 6 MQ 0 Chip	OK	19 degrees C / 66 degrees F
FPC 6 MQ 1 TSen	OK	16 degrees C / 60 degrees F
FPC 6 MQ 1 Chip	OK	20 degrees C / 68 degrees F
FPC 6 MQ 2 TSen	OK	16 degrees C / 60 degrees F
FPC 6 MQ 2 Chip	OK	17 degrees C / 62 degrees F
FPC 6 MQ 3 TSen	OK	16 degrees C / 60 degrees F
FPC 6 MQ 3 Chip	OK	16 degrees C / 60 degrees F
FPC 7 Intake	OK	10 degrees C / 50 degrees F
FPC 7 Exhaust A	OK	10 degrees C / 50 degrees F
FPC 7 Exhaust B	OK	11 degrees C / 51 degrees F
FPC 7 LU 0 TSen	OK	26 degrees C / 78 degrees F
FPC 7 LU 0 Chip	OK	26 degrees C / 78 degrees F
FPC 7 LU 1 TSen	OK	26 degrees C / 78 degrees F
FPC 7 LU 1 Chip	OK	29 degrees C / 84 degrees F
FPC 7 LU 2 TSen	OK	26 degrees C / 78 degrees F
FPC 7 LU 2 Chip	OK	19 degrees C / 66 degrees F
FPC 7 LU 3 TSen	OK	26 degrees C / 78 degrees F
FPC 7 LU 3 Chip	OK	24 degrees C / 75 degrees F

FPC 7 XM 0 TSen	OK	26 degrees C / 78 degrees F
FPC 7 XM 0 Chip	OK	34 degrees C / 93 degrees F
FPC 7 XM 1 TSen	OK	26 degrees C / 78 degrees F
FPC 7 XM 1 Chip	OK	32 degrees C / 89 degrees F
FPC 7 PLX Switch TSen	OK	26 degrees C / 78 degrees F
FPC 7 PLX Switch Chip	OK	22 degrees C / 71 degrees F
FPC 8 Intake	OK	10 degrees C / 50 degrees F
FPC 8 Exhaust A	OK	22 degrees C / 71 degrees F
FPC 8 Exhaust B	OK	28 degrees C / 82 degrees F
FPC 8 LU 0 TSen	OK	20 degrees C / 68 degrees F
FPC 8 LU 0 Chip	OK	33 degrees C / 91 degrees F
FPC 8 LU 1 TSen	OK	20 degrees C / 68 degrees F
FPC 8 LU 1 Chip	OK	23 degrees C / 73 degrees F
FPC 8 LU 2 TSen	OK	20 degrees C / 68 degrees F
FPC 8 LU 2 Chip	OK	26 degrees C / 78 degrees F
FPC 8 LU 3 TSen	OK	20 degrees C / 68 degrees F
FPC 8 LU 3 Chip	OK	33 degrees C / 91 degrees F
FPC 8 XM 0 TSen	OK	20 degrees C / 68 degrees F
FPC 8 XM 0 Chip	OK	29 degrees C / 84 degrees F
FPC 8 XF 0 TSen	OK	20 degrees C / 68 degrees F
FPC 8 XF 0 Chip	OK	38 degrees C / 100 degrees F
FPC 8 PLX Switch TSen	OK	20 degrees C / 68 degrees F
FPC 8 PLX Switch Chip	OK	24 degrees C / 75 degrees F
FPC 9 Intake	OK	11 degrees C / 51 degrees F
FPC 9 Exhaust A	OK	11 degrees C / 51 degrees F
FPC 9 Exhaust B	OK	11 degrees C / 51 degrees F
FPC 9 LU 0 TSen	OK	25 degrees C / 77 degrees F
FPC 9 LU 0 Chip	OK	24 degrees C / 75 degrees F
FPC 9 LU 1 TSen	OK	25 degrees C / 77 degrees F
FPC 9 LU 1 Chip	OK	26 degrees C / 78 degrees F
FPC 9 LU 2 TSen	OK	25 degrees C / 77 degrees F
FPC 9 LU 2 Chip	OK	16 degrees C / 60 degrees F
FPC 9 LU 3 TSen	OK	25 degrees C / 77 degrees F
FPC 9 LU 3 Chip	OK	21 degrees C / 69 degrees F
FPC 9 XM 0 TSen	OK	25 degrees C / 77 degrees F
FPC 9 XM 0 Chip	OK	32 degrees C / 89 degrees F
FPC 9 XM 1 TSen	OK	25 degrees C / 77 degrees F
FPC 9 XM 1 Chip	OK	25 degrees C / 77 degrees F
FPC 9 PLX Switch TSen	OK	25 degrees C / 77 degrees F
FPC 9 PLX Switch Chip	OK	21 degrees C / 69 degrees F
ADC 0 Intake	OK	12 degrees C / 53 degrees F
ADC 0 Exhaust	OK	20 degrees C / 68 degrees F
ADC 0 ADC-XF1	OK	26 degrees C / 78 degrees F
ADC 0 ADC-XF0	OK	32 degrees C / 89 degrees F
ADC 1 Intake	OK	11 degrees C / 51 degrees F
ADC 1 Exhaust	OK	21 degrees C / 69 degrees F
ADC 1 ADC-XF1	OK	24 degrees C / 75 degrees F
ADC 1 ADC-XF0	OK	31 degrees C / 87 degrees F
ADC 2 Intake	OK	14 degrees C / 57 degrees F
ADC 2 Exhaust	OK	21 degrees C / 69 degrees F
ADC 2 ADC-XF1	OK	28 degrees C / 82 degrees F
ADC 2 ADC-XF0	OK	34 degrees C / 93 degrees F
ADC 3 Intake	OK	13 degrees C / 55 degrees F
ADC 3 Exhaust	OK	19 degrees C / 66 degrees F
ADC 3 ADC-XF1	OK	24 degrees C / 75 degrees F
ADC 3 ADC-XF0	OK	31 degrees C / 87 degrees F
ADC 4 Intake	OK	9 degrees C / 48 degrees F
ADC 4 Exhaust	OK	22 degrees C / 71 degrees F
ADC 4 ADC-XF1	OK	28 degrees C / 82 degrees F
ADC 4 ADC-XF0	OK	35 degrees C / 95 degrees F
ADC 5 Intake	OK	12 degrees C / 53 degrees F

ADC 5 Exhaust	OK	22 degrees C / 71 degrees F
ADC 5 ADC-XF1	OK	28 degrees C / 82 degrees F
ADC 5 ADC-XF0	OK	34 degrees C / 93 degrees F
ADC 6 Intake	OK	11 degrees C / 51 degrees F
ADC 6 Exhaust	OK	21 degrees C / 69 degrees F
ADC 6 ADC-XF1	OK	26 degrees C / 78 degrees F
ADC 6 ADC-XF0	OK	35 degrees C / 95 degrees F
ADC 7 Intake	OK	14 degrees C / 57 degrees F
ADC 7 Exhaust	OK	22 degrees C / 71 degrees F
ADC 7 ADC-XF1	OK	26 degrees C / 78 degrees F
ADC 7 ADC-XF0	OK	34 degrees C / 93 degrees F
ADC 8 Intake	OK	14 degrees C / 57 degrees F
ADC 8 Exhaust	OK	21 degrees C / 69 degrees F
ADC 8 ADC-XF1	OK	24 degrees C / 75 degrees F
ADC 8 ADC-XF0	OK	31 degrees C / 87 degrees F
ADC 9 Intake	OK	10 degrees C / 50 degrees F
ADC 9 Exhaust	OK	22 degrees C / 71 degrees F
ADC 9 ADC-XF1	OK	28 degrees C / 82 degrees F
ADC 9 ADC-XF0	OK	36 degrees C / 96 degrees F
Fans Fan Tray 0 Fan 1	OK	3480 RPM
Fans Fan Tray 0 Fan 2	OK	3480 RPM
Fans Fan Tray 0 Fan 3	OK	3480 RPM
Fans Fan Tray 0 Fan 4	OK	3360 RPM
Fans Fan Tray 0 Fan 5	OK	3360 RPM
Fans Fan Tray 0 Fan 6	OK	3480 RPM
Fans Fan Tray 1 Fan 1	OK	3360 RPM
Fans Fan Tray 1 Fan 2	OK	3360 RPM
Fans Fan Tray 1 Fan 3	OK	3360 RPM
Fans Fan Tray 1 Fan 4	OK	3480 RPM
Fans Fan Tray 1 Fan 5	OK	3480 RPM
Fans Fan Tray 1 Fan 6	OK	3480 RPM
Fans Fan Tray 2 Fan 1	OK	3360 RPM
Fans Fan Tray 2 Fan 2	OK	3360 RPM
Fans Fan Tray 2 Fan 3	OK	3480 RPM
Fans Fan Tray 2 Fan 4	OK	3480 RPM
Fans Fan Tray 2 Fan 5	OK	3360 RPM
Fans Fan Tray 2 Fan 6	OK	3480 RPM
Fans Fan Tray 3 Fan 1	OK	3360 RPM
Fans Fan Tray 3 Fan 2	OK	3360 RPM
Fans Fan Tray 3 Fan 3	OK	3480 RPM
Fans Fan Tray 3 Fan 4	OK	3480 RPM
Fans Fan Tray 3 Fan 5	OK	3480 RPM
Fans Fan Tray 3 Fan 6	OK	3360 RPM

show chassis environment (T320 Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Power	PEM 0	OK	
	PEM 1	Absent	
Temp	SCG 0	OK	28 degrees C / 82 degrees F
	SCG 1	OK	28 degrees C / 82 degrees F
	Routing Engine 0	OK	31 degrees C / 87 degrees F
	Routing Engine 1	OK	30 degrees C / 86 degrees F
	CB 0	OK	32 degrees C / 89 degrees F
	CB 1	OK	32 degrees C / 89 degrees F
	SIB 0	OK	33 degrees C / 91 degrees F
	SIB 1	OK	33 degrees C / 91 degrees F
	SIB 2	OK	34 degrees C / 93 degrees F
	FPC 0 Top	OK	38 degrees C / 100 degrees F
	FPC 0 Bottom	OK	32 degrees C / 89 degrees F

	FPC 1 Top	OK	38 degrees C / 100 degrees F
	FPC 1 Bottom	OK	33 degrees C / 91 degrees F
	FPC 2 Top	OK	36 degrees C / 96 degrees F
	FPC 2 Bottom	OK	31 degrees C / 87 degrees F
	FPM GBUS	OK	26 degrees C / 78 degrees F
	FPM Display	OK	29 degrees C / 84 degrees F
Fans	Top Left Front fan	OK	Spinning at normal speed
	Top Left Middle fan	OK	Spinning at normal speed
	Top Left Rear fan	OK	Spinning at normal speed
	Top Right Front fan	OK	Spinning at normal speed
	Top Right Middle fan	OK	Spinning at normal speed
	Top Right Rear fan	OK	Spinning at normal speed
	Bottom Left Front fan	OK	Spinning at normal speed
	Bottom Left Middle fan	OK	Spinning at normal speed
	Bottom Left Rear fan	OK	Spinning at normal speed
	Bottom Right Front fan	OK	Spinning at normal speed
	Bottom Right Middle fan	OK	Spinning at normal speed
	Bottom Right Rear fan	OK	Spinning at normal speed
	Rear Tray Top fan	OK	Spinning at normal speed
	Rear Tray Second fan	OK	Spinning at normal speed
	Rear Tray Middle fan	OK	Spinning at normal speed
	Rear Tray Fourth fan	OK	Spinning at normal speed
Misc	Rear Tray Bottom fan	OK	Spinning at normal speed
	CIP	OK	
	SPMB 0	OK	
	SPMB 1	OK	

show chassis environment (T640 Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Temp	PEM 0	Absent	
	PEM 1	OK	22 degrees C / 71 degrees F
	SCG 0	OK	30 degrees C / 86 degrees F
	SCG 1	OK	30 degrees C / 86 degrees F
	Routing Engine 0	Present	
	Routing Engine 1	OK	27 degrees C / 80 degrees F
	CB 0	Present	
	CB 1	OK	33 degrees C / 91 degrees F
	SIB 0	Absent	
	SIB 1	Absent	
	SIB 2	Absent	
	SIB 3	Absent	
	SIB 4	Absent	
	FPC 4 Top	Testing	
	FPC 4 Bottom	Testing	
	FPC 5 Top	Testing	
	FPC 5 Bottom	Testing	
	FPC 6 Top	Testing	
	FPC 6 Bottom	Testing	
	FPM GBUS	OK	23 degrees C / 73 degrees F
	FPM Display	Absent	
Fans	Top Left Front fan	OK	Spinning at normal speed
	Top Left Middle fan	OK	Spinning at normal speed
	Top Left Rear fan	OK	Spinning at normal speed
	Top Right Front fan	OK	Spinning at normal speed
	Top Right Middle fan	OK	Spinning at normal speed
	Top Right Rear fan	OK	Spinning at normal speed
	Bottom Left Front fan	OK	Spinning at normal speed
	Bottom Left Middle fan	OK	Spinning at normal speed

	Bottom Left Rear fan	OK	Spinning at normal speed
	Bottom Right Front fan	OK	Spinning at normal speed
	Bottom Right Middle fan	OK	Spinning at normal speed
	Bottom Right Rear fan	OK	Spinning at normal speed
	Fourth Blower from top	OK	Spinning at normal speed
	Bottom Blower	OK	Spinning at normal speed
	Middle Blower	OK	Spinning at normal speed
	Top Blower	OK	Spinning at normal speed
	Second Blower from top	OK	Spinning at normal speed
Misc	CIP	OK	
	SPMB 0	OK	
	SPMB 1	OK	

show chassis environment (T4000 Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Temp	PEM 0	OK	33 degrees C / 91 degrees F
	PEM 1	Absent	
	SCG 0	OK	33 degrees C / 91 degrees F
	SCG 1	OK	33 degrees C / 91 degrees F
	Routing Engine 0	OK	33 degrees C / 91 degrees F
	Routing Engine 0 CPU	OK	50 degrees C / 122 degrees F
	Routing Engine 1	OK	32 degrees C / 89 degrees F
	Routing Engine 1 CPU	OK	46 degrees C / 114 degrees F
	CB 0	OK	32 degrees C / 89 degrees F
	CB 1	OK	33 degrees C / 91 degrees F
	SIB 0	OK	42 degrees C / 107 degrees F
	SIB 1	OK	42 degrees C / 107 degrees F
	SIB 2	OK	42 degrees C / 107 degrees F
	SIB 3	OK	43 degrees C / 109 degrees F
	SIB 4	OK	45 degrees C / 113 degrees F
	FPC 0 Fan Intake	OK	34 degrees C / 93 degrees F
	FPC 0 Fan Exhaust	OK	48 degrees C / 118 degrees F
	FPC 0 PMB	OK	47 degrees C / 116 degrees F
	FPC 0 LMB0	OK	50 degrees C / 122 degrees F
	FPC 0 LMB1	OK	41 degrees C / 105 degrees F
	FPC 0 LMB2	OK	35 degrees C / 95 degrees F
	FPC 0 PFE1 LU2	OK	46 degrees C / 114 degrees F
	FPC 0 PFE1 LU0	OK	41 degrees C / 105 degrees F
	FPC 0 PFE0 LU0	OK	57 degrees C / 134 degrees F
	FPC 0 XF1	OK	46 degrees C / 114 degrees F
	FPC 0 XF0	OK	52 degrees C / 125 degrees F
	FPC 0 XM1	OK	41 degrees C / 105 degrees F
	FPC 0 XM0	OK	50 degrees C / 122 degrees F
	FPC 0 PFE0 LU1	OK	56 degrees C / 132 degrees F
	FPC 0 PFE0 LU2	OK	45 degrees C / 113 degrees F
	FPC 0 PFE1 LU1	OK	37 degrees C / 98 degrees F
	FPC 3 Fan Intake	OK	36 degrees C / 96 degrees F
	FPC 3 Fan Exhaust	OK	51 degrees C / 123 degrees F
	FPC 3 PMB	OK	43 degrees C / 109 degrees F
	FPC 3 LMB0	OK	57 degrees C / 134 degrees F
	FPC 3 LMB1	OK	54 degrees C / 129 degrees F
	FPC 3 LMB2	OK	38 degrees C / 100 degrees F
	FPC 3 PFE1 LU2	OK	63 degrees C / 145 degrees F
	FPC 3 PFE1 LU0	OK	45 degrees C / 113 degrees F
	FPC 3 PFE0 LU0	OK	69 degrees C / 156 degrees F
	FPC 3 XF1	OK	62 degrees C / 143 degrees F
	FPC 3 XF0	OK	63 degrees C / 145 degrees F
	FPC 3 XM1	OK	43 degrees C / 109 degrees F

	FPC 3 XM0	OK	67 degrees C / 152 degrees F
	FPC 3 PFE0 LU1	OK	63 degrees C / 145 degrees F
	FPC 3 PFE0 LU2	OK	66 degrees C / 150 degrees F
	FPC 3 PFE1 LU1	OK	41 degrees C / 105 degrees F
	FPC 5 Top	OK	39 degrees C / 102 degrees F
	FPC 5 Bottom	OK	38 degrees C / 100 degrees F
	FPC 6 Fan Intake	OK	33 degrees C / 91 degrees F
	FPC 6 Fan Exhaust	OK	49 degrees C / 120 degrees F
	FPC 6 PMB	OK	40 degrees C / 104 degrees F
	FPC 6 LMB0	OK	60 degrees C / 140 degrees F
	FPC 6 LMB1	OK	58 degrees C / 136 degrees F
	FPC 6 LMB2	OK	40 degrees C / 104 degrees F
	FPC 6 PFE1 LU2	OK	69 degrees C / 156 degrees F
	FPC 6 PFE1 LU0	OK	45 degrees C / 113 degrees F
	FPC 6 PFE0 LU0	OK	71 degrees C / 159 degrees F
	FPC 6 XF1	OK	58 degrees C / 136 degrees F
	FPC 6 XF0	OK	65 degrees C / 149 degrees F
	FPC 6 XM1	OK	39 degrees C / 102 degrees F
	FPC 6 XM0	OK	66 degrees C / 150 degrees F
	FPC 6 PFE0 LU1	OK	69 degrees C / 156 degrees F
	FPC 6 PFE0 LU2	OK	69 degrees C / 156 degrees F
	FPC 6 PFE1 LU1	OK	42 degrees C / 107 degrees F
	FPM GBUS	OK	24 degrees C / 75 degrees F
	FPM Display	OK	27 degrees C / 80 degrees F
Fans	Top Left Front fan	OK	Spinning at high speed
	Top Left Middle fan	OK	Spinning at high speed
	Top Left Rear fan	OK	Spinning at high speed
	Top Right Front fan	OK	Spinning at high speed
	Top Right Middle fan	OK	Spinning at high speed
	Top Right Rear fan	OK	Spinning at high speed
	Bottom Left Front fan	OK	Spinning at high speed
	Bottom Left Middle fan	OK	Spinning at high speed
	Bottom Left Rear fan	OK	Spinning at high speed
	Bottom Right Front fan	OK	Spinning at high speed
	Bottom Right Middle fan	OK	Spinning at high speed
	Bottom Right Rear fan	OK	Spinning at high speed
	Rear Tray Top fan	OK	Spinning at high speed
	Rear Tray Second fan	OK	Spinning at high speed
	Rear Tray Third fan	OK	Spinning at high speed
	Rear Tray Fourth fan	OK	Spinning at high speed
	Rear Tray Fifth fan	OK	Spinning at high speed
	Rear Tray Sixth fan	OK	Spinning at high speed
	Rear Tray Seventh fan	OK	Spinning at high speed
Misc	Rear Tray Bottom fan	OK	Spinning at high speed
	CIP	OK	
	SPMB 0	OK	
	SPMB 1	OK	

show chassis environment (TX Matrix Router)

```
user@host> show chassis environment
scc-re0:
```

Class	Item	Status	Measurement
Temp	PEM 0	Absent	
	PEM 1	OK	29 degrees C / 84 degrees F
	Routing Engine 0	OK	34 degrees C / 93 degrees F
	Routing Engine 1	OK	34 degrees C / 93 degrees F
	CB 0	OK	32 degrees C / 89 degrees F
	CB 1	OK	32 degrees C / 89 degrees F
	SIB 0	OK	44 degrees C / 111 degrees F

	SIB 0 (B)	OK	44 degrees C / 111 degrees F
	FPM GBUS	OK	27 degrees C / 80 degrees F
	FPM Display	OK	32 degrees C / 89 degrees F
Fans	Top Left Front fan	OK	Spinning at normal speed
	Top Left Middle fan	OK	Spinning at normal speed
	Top Left Rear fan	OK	Spinning at normal speed
	Top Right Front fan	OK	Spinning at normal speed
	Top Right Middle fan	OK	Spinning at normal speed
	Top Right Rear fan	OK	Spinning at normal speed
	Bottom Left Front fan	OK	Spinning at normal speed
	Bottom Left Middle fan	OK	Spinning at normal speed
	Bottom Left Rear fan	OK	Spinning at normal speed
	Bottom Right Front fan	OK	Spinning at normal speed
	Bottom Right Middle fan	OK	Spinning at normal speed
	Bottom Right Rear fan	OK	Spinning at normal speed
	Rear Tray Top fan	OK	Spinning at normal speed
	Rear Tray Second fan	OK	Spinning at normal speed
	Rear Tray Third fan	OK	Spinning at normal speed
	Rear Tray Fourth fan	OK	Spinning at normal speed
	Rear Tray Fifth fan	OK	Spinning at normal speed
	Rear Tray Sixth fan	OK	Spinning at normal speed
	Rear Tray Seventh fan	OK	Spinning at normal speed
	Rear Tray Bottom fan	OK	Spinning at normal speed
Misc	CIP 0	OK	
	CIP 1	OK	
	SPMB 0	OK	
	SPMB 1	OK	

1cc0-re0:

Class	Item	Status	Measurement
Temp	PEM 0	OK	29 degrees C / 84 degrees F
	PEM 1	Absent	
	SCG 0	OK	35 degrees C / 95 degrees F
	SCG 1	Absent	
	Routing Engine 0	OK	39 degrees C / 102 degrees F
	Routing Engine 1	OK	36 degrees C / 96 degrees F
	CB 0	OK	32 degrees C / 89 degrees F
	CB 1	OK	32 degrees C / 89 degrees F
	SIB 0	OK	40 degrees C / 104 degrees F
	SIB 0 (B)	OK	51 degrees C / 123 degrees F
	FPC 0 Top	OK	45 degrees C / 113 degrees F
	FPC 0 Bottom	OK	31 degrees C / 87 degrees F
	FPC 1 Top	OK	34 degrees C / 93 degrees F
	FPC 1 Bottom	OK	31 degrees C / 87 degrees F
	FPM GBUS	OK	30 degrees C / 86 degrees F
	FPM Display	OK	34 degrees C / 93 degrees F
Fans	Top Left Front fan	OK	Spinning at normal speed
	Top Left Middle fan	OK	Spinning at normal speed
	Top Left Rear fan	OK	Spinning at normal speed
	Top Right Front fan	OK	Spinning at normal speed
	Top Right Middle fan	OK	Spinning at normal speed
	Top Right Rear fan	OK	Spinning at normal speed
	Bottom Left Front fan	OK	Spinning at normal speed
	Bottom Left Middle fan	OK	Spinning at normal speed
	Bottom Left Rear fan	OK	Spinning at normal speed
	Bottom Right Front fan	OK	Spinning at normal speed
	Bottom Right Middle fan	OK	Spinning at normal speed
	Bottom Right Rear fan	OK	Spinning at normal speed
	Rear Tray Top fan	OK	Spinning at normal speed
	Rear Tray Second fan	OK	Spinning at normal speed

	Rear Tray Third fan	OK	Spinning at normal speed
	Rear Tray Fourth fan	OK	Spinning at normal speed
	Rear Tray Fifth fan	OK	Spinning at normal speed
	Rear Tray Sixth fan	OK	Spinning at normal speed
	Rear Tray Seventh fan	OK	Spinning at normal speed
	Rear Tray Bottom fan	OK	Spinning at normal speed
Misc	CIP	OK	
	SPMB 0	OK	
	SPMB 1	OK	

```
lcc2-re0:
```

Class	Item	Status	Measurement
Temp	PEM 0	OK	29 degrees C / 84 degrees F
	PEM 1	Absent	
	SCG 0	OK	32 degrees C / 89 degrees F
	SCG 1	Absent	
	Routing Engine 0	OK	31 degrees C / 87 degrees F
	Routing Engine 1	OK	32 degrees C / 89 degrees F
	CB 0	OK	30 degrees C / 86 degrees F
	SIB 0	OK	38 degrees C / 100 degrees F
	SIB 0 (B)	OK	49 degrees C / 120 degrees F
	FPC 0 Top	OK	45 degrees C / 113 degrees F
	FPC 0 Bottom	OK	33 degrees C / 91 degrees F
	FPC 1 Top	OK	37 degrees C / 98 degrees F
	FPC 1 Bottom	OK	33 degrees C / 91 degrees F
	FPM GBUS	OK	30 degrees C / 86 degrees F
	FPM Display	OK	34 degrees C / 93 degrees F
Fans	Top Left Front fan	OK	Spinning at normal speed
	Top Left Middle fan	OK	Spinning at normal speed
...			

show chassis environment (T1600 Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Temp	PEM 0	OK	27 degrees C / 80 degrees F
	PEM 1	Absent	
	SCG 0	OK	31 degrees C / 87 degrees F
	SCG 1	OK	35 degrees C / 95 degrees F
	Routing Engine 0	OK	30 degrees C / 86 degrees F
	Routing Engine 1	OK	30 degrees C / 86 degrees F
	CB 0	OK	31 degrees C / 87 degrees F
	CB 1	OK	31 degrees C / 87 degrees F
	SIB 0	OK	41 degrees C / 105 degrees F
	SIB 0 (B)	OK	34 degrees C / 93 degrees F
	SIB 1	OK	0 degrees C / 32 degrees F
	SIB 1 (B)	OK	0 degrees C / 32 degrees F
	SIB 2	OK	0 degrees C / 32 degrees F
	SIB 2 (B)	OK	0 degrees C / 32 degrees F
	SIB 3	OK	0 degrees C / 32 degrees F
	SIB 3 (B)	OK	0 degrees C / 32 degrees F
	SIB 4	OK	0 degrees C / 32 degrees F
	SIB 4 (B)	OK	0 degrees C / 32 degrees F
	FPC 0 Top	OK	49 degrees C / 120 degrees F
	FPC 0 Bottom	OK	50 degrees C / 122 degrees F
	FPC 1 Top	OK	48 degrees C / 118 degrees F
	FPC 1 Bottom	OK	49 degrees C / 120 degrees F
	FPM GBUS	OK	27 degrees C / 80 degrees F
	FPM Display	OK	30 degrees C / 86 degrees F
Fans	Top Left Front fan	OK	Spinning at normal speed

Top Left Middle fan	OK	Spinning at normal speed
Top Left Rear fan	OK	Spinning at normal speed
Top Right Front fan	OK	Spinning at normal speed
Top Right Middle fan	OK	Spinning at normal speed
Top Right Rear fan	OK	Spinning at normal speed
Bottom Left Front fan	OK	Spinning at normal speed
Bottom Left Middle fan	OK	Spinning at normal speed
Bottom Left Rear fan	OK	Spinning at normal speed
Bottom Right Front fan	OK	Spinning at normal speed
Bottom Right Middle fan	OK	Spinning at normal speed
Bottom Right Rear fan	OK	Spinning at normal speed
Rear Tray Top fan	OK	Spinning at normal speed
Rear Tray Second fan	OK	Spinning at normal speed
Rear Tray Third fan	OK	Spinning at normal speed
Rear Tray Fourth fan	OK	Spinning at normal speed
Rear Tray Fifth fan	OK	Spinning at normal speed
Rear Tray Sixth fan	OK	Spinning at normal speed
Rear Tray Seventh fan	OK	Spinning at normal speed
Rear Tray Bottom fan	OK	Spinning at normal speed
Misc CIP	OK	
SPMB 0	OK	
SPMB 1	OK	

show chassis environment (TX Matrix Plus Router)

```
user@host> show chassis environment
sfc0-re0:
```

Class	Item	Status	Measurement
Temp	PEM 0	OK	28 degrees C / 82 degrees F
	PEM 1	Absent	
	Routing Engine 0	OK	27 degrees C / 80 degrees F
	Routing Engine 1	OK	29 degrees C / 84 degrees F
	CB 0 Intake	OK	26 degrees C / 78 degrees F
	CB 0 Exhaust A	OK	25 degrees C / 77 degrees F
	CB 0 Exhaust B	OK	25 degrees C / 77 degrees F
	CB 1 Intake	OK	26 degrees C / 78 degrees F
	CB 1 Exhaust A	OK	26 degrees C / 78 degrees F
	CB 1 Exhaust B	OK	26 degrees C / 78 degrees F
	SIB F13 0	OK	47 degrees C / 116 degrees F
	SIB F13 0 (B)	OK	48 degrees C / 118 degrees F
	SIB F13 1	OK	38 degrees C / 100 degrees F
	SIB F13 1 (B)	OK	37 degrees C / 98 degrees F
	SIB F2S 0/0	OK	27 degrees C / 80 degrees F
	SIB F2S 0/2	OK	28 degrees C / 82 degrees F
	SIB F2S 0/4	OK	27 degrees C / 80 degrees F
	SIB F2S 0/6	OK	28 degrees C / 82 degrees F
	SIB F2S 1/0	OK	26 degrees C / 78 degrees F
	SIB F2S 1/2	OK	26 degrees C / 78 degrees F
	SIB F2S 1/4	OK	26 degrees C / 78 degrees F
	SIB F2S 1/6	OK	26 degrees C / 78 degrees F
	SIB F2S 2/0	OK	25 degrees C / 77 degrees F
	SIB F2S 2/2	OK	25 degrees C / 77 degrees F
	SIB F2S 2/4	OK	23 degrees C / 73 degrees F
	CIP 0 Intake	OK	23 degrees C / 73 degrees F
	CIP 0 Exhaust A	OK	24 degrees C / 75 degrees F
	CIP 0 Exhaust B	OK	24 degrees C / 75 degrees F
	CIP 1 Intake	OK	24 degrees C / 75 degrees F
	CIP 1 Exhaust A	OK	25 degrees C / 77 degrees F
	CIP 1 Exhaust B	OK	25 degrees C / 77 degrees F
Fans	Fan Tray 0 Fan 1	OK	Spinning at normal speed

Fan Tray 0 Fan 2	OK	Spinning at normal speed
Fan Tray 0 Fan 3	OK	Spinning at normal speed
Fan Tray 0 Fan 4	OK	Spinning at normal speed
Fan Tray 0 Fan 5	OK	Spinning at normal speed
Fan Tray 0 Fan 6	OK	Spinning at normal speed
Fan Tray 1 Fan 1	OK	Spinning at normal speed
Fan Tray 1 Fan 2	OK	Spinning at normal speed
Fan Tray 1 Fan 3	OK	Spinning at normal speed
Fan Tray 1 Fan 4	OK	Spinning at normal speed
Fan Tray 1 Fan 5	OK	Spinning at normal speed
Fan Tray 1 Fan 6	OK	Spinning at normal speed
Fan Tray 2 Fan 1	OK	Spinning at normal speed
Fan Tray 2 Fan 2	OK	Spinning at normal speed
Fan Tray 2 Fan 3	OK	Spinning at normal speed
Fan Tray 2 Fan 4	OK	Spinning at normal speed
Fan Tray 2 Fan 5	OK	Spinning at normal speed
Fan Tray 2 Fan 6	OK	Spinning at normal speed
Fan Tray 2 Fan 7	OK	Spinning at normal speed
Fan Tray 2 Fan 8	OK	Spinning at normal speed
Fan Tray 2 Fan 9	OK	Spinning at normal speed
Fan Tray 3 Fan 1	OK	Spinning at normal speed
Fan Tray 3 Fan 2	OK	Spinning at normal speed
Fan Tray 3 Fan 3	OK	Spinning at normal speed
Fan Tray 3 Fan 4	OK	Spinning at normal speed
Fan Tray 3 Fan 5	OK	Spinning at normal speed
Fan Tray 3 Fan 6	OK	Spinning at normal speed
Fan Tray 3 Fan 7	OK	Spinning at normal speed
Fan Tray 3 Fan 8	OK	Spinning at normal speed
Fan Tray 3 Fan 9	OK	Spinning at normal speed
Fan Tray 4 Fan 1	OK	Spinning at normal speed
Fan Tray 4 Fan 2	OK	Spinning at normal speed
Fan Tray 4 Fan 3	OK	Spinning at normal speed
Fan Tray 4 Fan 4	OK	Spinning at normal speed
Fan Tray 4 Fan 5	OK	Spinning at normal speed
Fan Tray 4 Fan 6	OK	Spinning at normal speed
Fan Tray 4 Fan 7	OK	Spinning at normal speed
Fan Tray 4 Fan 8	OK	Spinning at normal speed
Fan Tray 4 Fan 9	OK	Spinning at normal speed
Fan Tray 5 Fan 1	OK	Spinning at normal speed
Fan Tray 5 Fan 2	OK	Spinning at normal speed
Fan Tray 5 Fan 3	OK	Spinning at normal speed
Fan Tray 5 Fan 4	OK	Spinning at normal speed
Fan Tray 5 Fan 5	OK	Spinning at normal speed
Fan Tray 5 Fan 6	OK	Spinning at normal speed
Fan Tray 5 Fan 7	OK	Spinning at normal speed
Fan Tray 5 Fan 8	OK	Spinning at normal speed
Fan Tray 5 Fan 9	OK	Spinning at normal speed
Misc SPMB 0	OK	
SPMB 1	OK	

1cc0-re0:

Class	Item	Status	Measurement
Temp	PEM 0	OK	27 degrees C / 80 degrees F
	PEM 1	Absent	
	SCG 0	OK	31 degrees C / 87 degrees F
	SCG 1	OK	35 degrees C / 95 degrees F
	Routing Engine 0	OK	30 degrees C / 86 degrees F
	Routing Engine 1	OK	30 degrees C / 86 degrees F
	CB 0	OK	31 degrees C / 87 degrees F
	CB 1	OK	31 degrees C / 87 degrees F

	SIB 0	OK	41 degrees C / 105 degrees F
	SIB 0 (B)	OK	34 degrees C / 93 degrees F
	SIB 1	OK	0 degrees C / 32 degrees F
	SIB 1 (B)	OK	0 degrees C / 32 degrees F
	SIB 2	OK	0 degrees C / 32 degrees F
	SIB 2 (B)	OK	0 degrees C / 32 degrees F
	SIB 3	OK	0 degrees C / 32 degrees F
	SIB 3 (B)	OK	0 degrees C / 32 degrees F
	SIB 4	OK	0 degrees C / 32 degrees F
	SIB 4 (B)	OK	0 degrees C / 32 degrees F
	FPC 0 Top	OK	49 degrees C / 120 degrees F
	FPC 0 Bottom	OK	50 degrees C / 122 degrees F
	FPC 1 Top	OK	48 degrees C / 118 degrees F
	FPC 1 Bottom	OK	49 degrees C / 120 degrees F
	FPM GBUS	OK	27 degrees C / 80 degrees F
	FPM Display	OK	30 degrees C / 86 degrees F
Fans	Top Left Front fan	OK	Spinning at normal speed
	Top Left Middle fan	OK	Spinning at normal speed
	Top Left Rear fan	OK	Spinning at normal speed
	Top Right Front fan	OK	Spinning at normal speed
	Top Right Middle fan	OK	Spinning at normal speed
	Top Right Rear fan	OK	Spinning at normal speed
	Bottom Left Front fan	OK	Spinning at normal speed
	Bottom Left Middle fan	OK	Spinning at normal speed
	Bottom Left Rear fan	OK	Spinning at normal speed
	Bottom Right Front fan	OK	Spinning at normal speed
	Bottom Right Middle fan	OK	Spinning at normal speed
	Bottom Right Rear fan	OK	Spinning at normal speed
	Rear Tray Top fan	OK	Spinning at normal speed
	Rear Tray Second fan	OK	Spinning at normal speed
	Rear Tray Third fan	OK	Spinning at normal speed
	Rear Tray Fourth fan	OK	Spinning at normal speed
	Rear Tray Fifth fan	OK	Spinning at normal speed
	Rear Tray Sixth fan	OK	Spinning at normal speed
	Rear Tray Seventh fan	OK	Spinning at normal speed
	Rear Tray Bottom fan	OK	Spinning at normal speed
Misc	CIP	OK	
	SPMB 0	OK	
	SPMB 1	OK	

show chassis environment (TX Matrix Plus router with 3D SIBs)

```
user@host> show chassis environment
sfc0-re0:
```

Class	Item	Status	Measurement
Temp	PEM 0	Check	30 degrees C / 86 degrees F
	PEM 1	OK	33 degrees C / 91 degrees F
	Routing Engine 0	OK	28 degrees C / 82 degrees F
	Routing Engine 0 CPU	OK	42 degrees C / 107 degrees F
	Routing Engine 1	OK	29 degrees C / 84 degrees F
	Routing Engine 1 CPU	OK	44 degrees C / 111 degrees F
	CB 0 Intake	OK	30 degrees C / 86 degrees F
	CB 0 Exhaust A	OK	28 degrees C / 82 degrees F
	CB 0 Exhaust B	OK	30 degrees C / 86 degrees F
	CB 1 Intake	OK	31 degrees C / 87 degrees F
	CB 1 Exhaust A	OK	27 degrees C / 80 degrees F
	CB 1 Exhaust B	OK	31 degrees C / 87 degrees F
	SIB F13 0 Board	OK	44 degrees C / 111 degrees F
	SIB F13 0 XF Junction	OK	62 degrees C / 143 degrees F
	SIB F13 3 Board	OK	45 degrees C / 113 degrees F

	SIB F13 3 XF Junction	OK	60 degrees C / 140 degrees F
	SIB F13 6 Board	OK	47 degrees C / 116 degrees F
	SIB F13 6 XF Junction	OK	62 degrees C / 143 degrees F
	SIB F2S 0/0 Board	OK	32 degrees C / 89 degrees F
	SIB F2S 0/0 XF Junction	OK	42 degrees C / 107 degrees F
	SIB F2S 0/2 Board	OK	31 degrees C / 87 degrees F
	SIB F2S 0/2 XF Junction	OK	41 degrees C / 105 degrees F
	SIB F2S 0/4 Board	OK	31 degrees C / 87 degrees F
	SIB F2S 0/4 XF Junction	OK	42 degrees C / 107 degrees F
	SIB F2S 0/6 Board	OK	31 degrees C / 87 degrees F
	SIB F2S 0/6 XF Junction	OK	41 degrees C / 105 degrees F
	SIB F2S 1/0 Board	OK	31 degrees C / 87 degrees F
	SIB F2S 1/0 XF Junction	OK	41 degrees C / 105 degrees F
	SIB F2S 1/2 Board	OK	29 degrees C / 84 degrees F
	SIB F2S 1/2 XF Junction	OK	39 degrees C / 102 degrees F
	SIB F2S 1/4 Board	OK	29 degrees C / 84 degrees F
	SIB F2S 1/4 XF Junction	OK	35 degrees C / 95 degrees F
	SIB F2S 1/6 Board	OK	30 degrees C / 86 degrees F
	SIB F2S 1/6 XF Junction	OK	41 degrees C / 105 degrees F
	SIB F2S 2/0 Board	OK	30 degrees C / 86 degrees F
	SIB F2S 2/0 XF Junction	OK	42 degrees C / 107 degrees F
	SIB F2S 2/2 Board	OK	28 degrees C / 82 degrees F
	SIB F2S 2/2 XF Junction	OK	39 degrees C / 102 degrees F
	SIB F2S 2/4 Board	OK	29 degrees C / 84 degrees F
	SIB F2S 2/4 XF Junction	OK	42 degrees C / 107 degrees F
	SIB F2S 2/6 Board	OK	29 degrees C / 84 degrees F
	SIB F2S 2/6 XF Junction	OK	41 degrees C / 105 degrees F
	CIP 0 Intake	OK	25 degrees C / 77 degrees F
	CIP 0 Exhaust A	OK	26 degrees C / 78 degrees F
	CIP 0 Exhaust B	OK	26 degrees C / 78 degrees F
	CIP 1 Intake	OK	26 degrees C / 78 degrees F
	CIP 1 Exhaust A	OK	27 degrees C / 80 degrees F
	CIP 1 Exhaust B	OK	27 degrees C / 80 degrees F
Fans	Fan Tray 0 Fan 1	OK	Spinning at normal speed
	Fan Tray 0 Fan 2	OK	Spinning at normal speed
	Fan Tray 0 Fan 3	OK	Spinning at normal speed
	Fan Tray 0 Fan 4	OK	Spinning at normal speed
	Fan Tray 0 Fan 5	OK	Spinning at normal speed
	Fan Tray 0 Fan 6	OK	Spinning at normal speed
	Fan Tray 1 Fan 1	OK	Spinning at normal speed
	Fan Tray 1 Fan 2	OK	Spinning at normal speed
	Fan Tray 1 Fan 3	OK	Spinning at normal speed
	Fan Tray 1 Fan 4	OK	Spinning at normal speed
	Fan Tray 1 Fan 5	OK	Spinning at normal speed
	Fan Tray 1 Fan 6	OK	Spinning at normal speed
	Fan Tray 2 Fan 1	OK	Spinning at normal speed
	Fan Tray 2 Fan 2	OK	Spinning at normal speed
	Fan Tray 2 Fan 3	OK	Spinning at normal speed
	Fan Tray 2 Fan 4	OK	Spinning at normal speed
	Fan Tray 2 Fan 5	OK	Spinning at normal speed
	Fan Tray 2 Fan 6	OK	Spinning at normal speed
	Fan Tray 2 Fan 7	OK	Spinning at normal speed
	Fan Tray 2 Fan 8	OK	Spinning at normal speed
	Fan Tray 2 Fan 9	OK	Spinning at normal speed
	Fan Tray 3 Fan 1	OK	Spinning at normal speed
	Fan Tray 3 Fan 2	OK	Spinning at normal speed
	Fan Tray 3 Fan 3	OK	Spinning at normal speed
	Fan Tray 3 Fan 4	OK	Spinning at normal speed
	Fan Tray 3 Fan 5	OK	Spinning at normal speed
	Fan Tray 3 Fan 6	OK	Spinning at normal speed
	Fan Tray 3 Fan 7	OK	Spinning at normal speed

Fan Tray 3 Fan 8	OK	Spinning at normal speed
Fan Tray 3 Fan 9	OK	Spinning at normal speed
Fan Tray 4 Fan 1	OK	Spinning at normal speed
Fan Tray 4 Fan 2	OK	Spinning at normal speed
Fan Tray 4 Fan 3	OK	Spinning at normal speed
Fan Tray 4 Fan 4	OK	Spinning at normal speed
Fan Tray 4 Fan 5	OK	Spinning at normal speed
Fan Tray 4 Fan 6	OK	Spinning at normal speed
Fan Tray 4 Fan 7	OK	Spinning at normal speed
Fan Tray 4 Fan 8	OK	Spinning at normal speed
Fan Tray 4 Fan 9	OK	Spinning at normal speed
Fan Tray 5 Fan 1	OK	Spinning at normal speed
Fan Tray 5 Fan 2	OK	Spinning at normal speed
Fan Tray 5 Fan 3	OK	Spinning at normal speed
Fan Tray 5 Fan 4	OK	Spinning at normal speed
Fan Tray 5 Fan 5	OK	Spinning at normal speed
Fan Tray 5 Fan 6	OK	Spinning at normal speed
Fan Tray 5 Fan 7	OK	Spinning at normal speed
Fan Tray 5 Fan 8	OK	Spinning at normal speed
Fan Tray 5 Fan 9	Check	
Misc SPMB 0	OK	
SPMB 1	OK	

lcc0-re0:

Class	Item	Status	Measurement
Temp	PEM 0	OK	29 degrees C / 84 degrees F
	PEM 1	Check	29 degrees C / 84 degrees F
	SCG 0	OK	32 degrees C / 89 degrees F
	SCG 1	OK	33 degrees C / 91 degrees F
	Routing Engine 0	OK	32 degrees C / 89 degrees F
	Routing Engine 0 CPU	OK	51 degrees C / 123 degrees F
	Routing Engine 1	OK	32 degrees C / 89 degrees F
	Routing Engine 1 CPU	OK	49 degrees C / 120 degrees F
	CB 0	OK	34 degrees C / 93 degrees F
	CB 1	OK	34 degrees C / 93 degrees F
	SIB 0	OK	39 degrees C / 102 degrees F
	SIB 0 (B)	Absent	
	SIB 1	OK	39 degrees C / 102 degrees F
	SIB 1 (B)	Absent	
	SIB 2	OK	39 degrees C / 102 degrees F
	SIB 2 (B)	Absent	
	FPC 4 Top	OK	43 degrees C / 109 degrees F
	FPC 4 Bottom	OK	43 degrees C / 109 degrees F
	FPC 7 Fan Intake	OK	35 degrees C / 95 degrees F
	FPC 7 Fan Exhaust	OK	50 degrees C / 122 degrees F
	FPC 7 PMB	OK	50 degrees C / 122 degrees F
	FPC 7 LMB0	OK	55 degrees C / 131 degrees F
	FPC 7 LMB1	OK	49 degrees C / 120 degrees F
	FPC 7 LMB2	OK	39 degrees C / 102 degrees F
	FPC 7 PFE1 LU2	OK	55 degrees C / 131 degrees F
	FPC 7 PFE1 LU0	OK	45 degrees C / 113 degrees F
	FPC 7 PFE0 LU0	OK	62 degrees C / 143 degrees F
	FPC 7 XF1	OK	52 degrees C / 125 degrees F
	FPC 7 XF0	OK	61 degrees C / 141 degrees F
	FPC 7 XM1	OK	39 degrees C / 102 degrees F
	FPC 7 XM0	OK	56 degrees C / 132 degrees F
	FPC 7 PFE0 LU1	OK	60 degrees C / 140 degrees F
	FPC 7 PFE0 LU2	OK	55 degrees C / 131 degrees F
	FPC 7 PFE1 LU1	OK	41 degrees C / 105 degrees F
	FPM GBUS	OK	24 degrees C / 75 degrees F

	FPM Display	OK	28 degrees C / 82 degrees F
Fans	Top Left Front fan	OK	Spinning at normal speed
	Top Left Middle fan	OK	Spinning at normal speed
	Top Left Rear fan	OK	Spinning at normal speed
	Top Right Front fan	OK	Spinning at normal speed
	Top Right Middle fan	OK	Spinning at normal speed
	Top Right Rear fan	OK	Spinning at normal speed
	Bottom Left Front fan	OK	Spinning at normal speed
	Bottom Left Middle fan	OK	Spinning at normal speed
	Bottom Left Rear fan	OK	Spinning at normal speed
	Bottom Right Front fan	OK	Spinning at normal speed
	Bottom Right Middle fan	OK	Spinning at normal speed
	Bottom Right Rear fan	OK	Spinning at normal speed
	Rear Tray fan 1 (Top)	OK	Spinning at normal speed
	Rear Tray fan 2	OK	Spinning at normal speed
	Rear Tray fan 3	OK	Spinning at normal speed
	Rear Tray fan 4	OK	Spinning at normal speed
Misc	Rear Tray fan 5	OK	Spinning at normal speed
	Rear Tray fan 6	OK	Spinning at normal speed
	Rear Tray fan 7	OK	Spinning at normal speed
	Rear Tray fan 8	OK	Spinning at normal speed
	Rear Tray fan 9	OK	Spinning at normal speed
	Rear Tray fan 10	OK	Spinning at normal speed
	Rear Tray fan 11	OK	Spinning at normal speed
	Rear Tray fan 12	OK	Spinning at normal speed
	Rear Tray fan 13	OK	Spinning at normal speed
	Rear Tray fan 14	OK	Spinning at normal speed
	Rear Tray fan 15	OK	Spinning at normal speed
	Rear Tray fan 16 (Bottom)	OK	Spinning at normal speed
	CIP	OK	
	SPMB 0	OK	
	SPMB 1	OK	

show chassis environment (EX4200 Standalone Switch)

```
user@switch> show chassis environment
```

Class	Item	Status	Measurement
Power	FPC 0 Power Supply 0	OK	
	FPC 0 Power Supply 1	Absent	
Temp	FPC 0 CPU	OK	41 degrees C / 105 degrees F
	FPC 0 EX-PFE1	OK	42 degrees C / 107 degrees F
	FPC 0 EX-PFE2	OK	46 degrees C / 114 degrees F
	FPC 0 GEPHY Front Left	OK	25 degrees C / 77 degrees F
	FPC 0 GEPHY Front Right	OK	27 degrees C / 80 degrees F
	FPC 0 Uplink Conn	OK	29 degrees C / 84 degrees F
Fans	FPC 0 Fan 1	OK	Spinning at normal speed
	FPC 0 Fan 2	OK	Spinning at normal speed
	FPC 0 Fan 3	OK	Spinning at normal speed

show chassis environment (EX8216 Switch)

```
user@switch> show chassis environment
```

Class	Item	Status	Measurement
Power	PSU 0	OK	
	PSU 1	OK	
	PSU 2	OK	
	PSU 3	Check	
	PSU 4	Absent	
	PSU 5	Absent	
Temp	CB 0 Intake	OK	23 degrees C / 73 degrees F
	CB 0 Exhaust	OK	26 degrees C / 78 degrees F

	CB 1 Intake	OK	22 degrees C / 71 degrees F
	CB 1 Exhaust	OK	25 degrees C / 77 degrees F
	FPC 4 Intake	OK	49 degrees C / 120 degrees F
	FPC 4 Exhaust	OK	59 degrees C / 138 degrees F
	SIB 5 Intake	OK	25 degrees C / 77 degrees F
	SIB 5 Exhaust	OK	35 degrees C / 95 degrees F
	SIB 6 Intake	OK	25 degrees C / 77 degrees F
	SIB 6 Exhaust	OK	38 degrees C / 100 degrees F
Fans	Top Fan 1	OK	Spinning at normal speed
	Top Fan 2	OK	Spinning at normal speed
	Top Fan 3	OK	Spinning at normal speed
	Top Fan 4	OK	Spinning at normal speed
	Top Fan 5	OK	Spinning at normal speed
	Top Fan 6	OK	Spinning at normal speed
	Top Fan 7	OK	Spinning at normal speed
	Top Fan 8	OK	Spinning at normal speed
	Top Fan 9	OK	Spinning at normal speed
	Bottom Fan 1	OK	Spinning at normal speed
	Bottom Fan 2	OK	Spinning at normal speed
	Bottom Fan 3	OK	Spinning at normal speed
	Bottom Fan 4	OK	Spinning at normal speed
	Bottom Fan 5	OK	Spinning at normal speed
	Bottom Fan 6	OK	Spinning at normal speed
	Bottom Fan 7	OK	Spinning at normal speed
	Bottom Fan 8	OK	Spinning at normal speed
	Bottom Fan 9	OK	Spinning at normal speed

show chassis environment (EX9200 Switch)

```
user@switch> show chassis environment
```

Class	Item	Status	Measurement
Temp	PEM 0	Check	
	PEM 1	OK	40 degrees C / 104 degrees F
	PEM 2	OK	40 degrees C / 104 degrees F
	PEM 3	Absent	
	Routing Engine 0	OK	35 degrees C / 95 degrees F
	Routing Engine 0 CPU	OK	33 degrees C / 91 degrees F
	Routing Engine 1	OK	38 degrees C / 100 degrees F
	Routing Engine 1 CPU	OK	33 degrees C / 91 degrees F
	CB 0 Intake	OK	35 degrees C / 95 degrees F
	CB 0 Exhaust A	OK	33 degrees C / 91 degrees F
	CB 0 Exhaust B	OK	40 degrees C / 104 degrees F
	CB 0 ACBC	OK	39 degrees C / 102 degrees F
	CB 0 XF A	OK	49 degrees C / 120 degrees F
	CB 0 XF B	OK	46 degrees C / 114 degrees F
	CB 1 Intake	OK	37 degrees C / 98 degrees F
	CB 1 Exhaust A	OK	32 degrees C / 89 degrees F
	CB 1 Exhaust B	OK	39 degrees C / 102 degrees F
	CB 1 ACBC	OK	41 degrees C / 105 degrees F
	CB 1 XF A	OK	49 degrees C / 120 degrees F
	CB 1 XF B	OK	49 degrees C / 120 degrees F
	FPC 2 Intake	OK	37 degrees C / 98 degrees F
	FPC 2 Exhaust A	OK	40 degrees C / 104 degrees F
	FPC 2 Exhaust B	OK	34 degrees C / 93 degrees F
	FPC 2 LU 0 TCAM TSen	OK	44 degrees C / 111 degrees F
	FPC 2 LU 0 TCAM Chip	OK	48 degrees C / 118 degrees F
	FPC 2 LU 0 TSen	OK	44 degrees C / 111 degrees F
	FPC 2 LU 0 Chip	OK	60 degrees C / 140 degrees F
	FPC 2 MQ 0 TSen	OK	44 degrees C / 111 degrees F
	FPC 2 MQ 0 Chip	OK	51 degrees C / 123 degrees F
	FPC 3 Intake	OK	39 degrees C / 102 degrees F

```

FPC 3 Exhaust A          OK          51 degrees C / 123 degrees F

[...Output truncated...]

Fans  Top Rear Fan       OK          Spinning at intermediate-speed
      Bottom Rear Fan    OK          Spinning at intermediate-speed
      Top Middle Fan     OK          Spinning at intermediate-speed
      Bottom Middle Fan  OK          Spinning at intermediate-speed
      Top Front Fan      OK          Spinning at intermediate-speed
      Bottom Front Fan   OK          Spinning at intermediate-speed

```

show chassis environment (QFX Series and OCX Series)

```

user@switch> show chassis environment
Class Item                               Status      Measurement
Power FPC 0 Power Supply 0              OK
      FPC 0 Power Supply 1              OK
Temp  FPC 0 Sensor TopLeft I            OK          26 degrees C / 78 degrees F
      FPC 0 Sensor TopRight I           OK          24 degrees C / 75 degrees F
      FPC 0 Sensor TopLeft E            OK          30 degrees C / 86 degrees F
      FPC 0 Sensor TopRight E           OK          30 degrees C / 86 degrees F
      FPC 0 Sensor TopMiddle I          OK          30 degrees C / 86 degrees F
      FPC 0 Sensor TopMiddle E          OK          38 degrees C / 100 degrees F
      FPC 0 Sensor Bottom I             OK          34 degrees C / 93 degrees F
      FPC 0 Sensor Bottom E             OK          38 degrees C / 100 degrees F
      FPC 0 Sensor Die Temp             OK          38 degrees C / 100 degrees F
      FPC 0 Sensor Mgmt Brd I           OK          24 degrees C / 75 degrees F
      FPC 0 Sensor Switch I             OK          28 degrees C / 82 degrees F
Fans  FPC 0 Fan 1 (left)                 Failed
      FPC 0 Fan 2 (right)               OK          Spinning at normal speed
      FPC 0 Fan 3 (middle)              OK          Spinning at normal speed

```

show chassis environment interconnect-device (QFabric System)

```

user@switch> show chassis environment interconnect-device IC-A0004
Class Item                               Status      Measurement
CB 0
CB 0 L Intake                    OK          30 degrees C / 86 degrees F
CB 0 R Intake                    OK          31 degrees C / 87 degrees F
CB 0 L Exhaust                   OK          32 degrees C / 89 degrees F
CB 0 R Exhaust                   OK          33 degrees C / 91 degrees F
Routing Engine 0 CPU temp        OK          51 degrees C / 123 degrees F
CB 1
CB 1 L Intake                    OK          27 degrees C / 80 degrees F
CB 1 R Intake                    OK          29 degrees C / 84 degrees F
CB 1 L Exhaust                   OK          31 degrees C / 87 degrees F
CB 1 R Exhaust                   OK          32 degrees C / 89 degrees F
Routing Engine 1 CPU temp        OK          40 degrees C / 104 degrees F
FC 0 FPC 0
FPC 0 L Intake                   OK          25 degrees C / 77 degrees F
FPC 0 R Intake                   OK          28 degrees C / 82 degrees F
FPC 0 L Exhaust                  OK          28 degrees C / 82 degrees F
FPC 0 R Exhaust                  OK          29 degrees C / 84 degrees F
FC 7 FPC 7
FPC 7 L Intake                   OK          25 degrees C / 77 degrees F
FPC 7 R Intake                   OK          26 degrees C / 78 degrees F
FPC 7 L Exhaust                  OK          28 degrees C / 82 degrees F
FPC 7 R Exhaust                  OK          29 degrees C / 84 degrees F
RC 0 FPC 8
FPC 8 L Intake                   OK          25 degrees C / 77 degrees F
FPC 8 R Intake                   OK          26 degrees C / 78 degrees F

```

FPC 8 L Exhaust	OK	32 degrees C / 89 degrees F
FPC 8 R Exhaust	OK	30 degrees C / 86 degrees F
RC 7 FPC 15		
FPC 15 L Intake	OK	24 degrees C / 75 degrees F
FPC 15 R Intake	OK	25 degrees C / 77 degrees F
FPC 15 L Exhaust	OK	33 degrees C / 91 degrees F
FPC 15 R Exhaust	OK	31 degrees C / 87 degrees F
Fans TFT 0 Fan 0	OK	Spinning at normal speed
Fans TFT 0 Fan 1	OK	Spinning at normal speed
Fans TFT 0 Fan 2	OK	Spinning at normal speed
Fans TFT 0 Fan 3	OK	Spinning at normal speed
Fans TFT 0 Fan 4	OK	Spinning at normal speed
Fans TFT 0 Fan 5	OK	Spinning at normal speed
Fans BFT 1 Fan 0	OK	Spinning at normal speed
Fans BFT 1 Fan 1	OK	Spinning at normal speed
Fans BFT 1 Fan 2	OK	Spinning at normal speed
Fans BFT 1 Fan 3	Check	
Fans BFT 1 Fan 4	OK	Spinning at normal speed
Fans BFT 1 Fan 5	OK	Spinning at normal speed
Fans SFT 0 Fan 0 Rotor 0	OK	Spinning at normal speed
Fans SFT 0 Fan 0 Rotor 1	OK	Spinning at normal speed
Fans SFT 0 Fan 1 Rotor 0	OK	Spinning at normal speed
Fans SFT 0 Fan 1 Rotor 1	OK	Spinning at normal speed
Fans SFT 0 Fan 2 Rotor 0	OK	Spinning at normal speed
Fans SFT 0 Fan 2 Rotor 1	OK	Spinning at normal speed
Fans SFT 0 Fan 3 Rotor 0	OK	Spinning at normal speed
Fans SFT 0 Fan 3 Rotor 1	OK	Spinning at normal speed
Fans SFT 1 Fan 0 Rotor 0	OK	Spinning at normal speed
Fans SFT 1 Fan 0 Rotor 1	OK	Spinning at normal speed
Fans SFT 1 Fan 1 Rotor 0	OK	Spinning at normal speed
Fans SFT 1 Fan 1 Rotor 1	OK	Spinning at normal speed
Fans SFT 1 Fan 2 Rotor 0	OK	Spinning at normal speed
Fans SFT 1 Fan 2 Rotor 1	OK	Spinning at normal speed
Fans SFT 1 Fan 3 Rotor 0	OK	Spinning at normal speed
Fans SFT 1 Fan 3 Rotor 1	OK	Spinning at normal speed
Fans SFT 2 Fan 0 Rotor 0	OK	Spinning at normal speed
Fans SFT 2 Fan 0 Rotor 1	OK	Spinning at normal speed
Fans SFT 2 Fan 1 Rotor 0	OK	Spinning at normal speed
Fans SFT 2 Fan 1 Rotor 1	OK	Spinning at normal speed
Fans SFT 2 Fan 2 Rotor 0	OK	Spinning at normal speed
Fans SFT 2 Fan 2 Rotor 1	OK	Spinning at normal speed
Fans SFT 2 Fan 3 Rotor 0	OK	Spinning at normal speed
Fans SFT 2 Fan 3 Rotor 1	OK	Spinning at normal speed
Fans SFT 3 Fan 0 Rotor 0	OK	Spinning at normal speed
Fans SFT 3 Fan 0 Rotor 1	OK	Spinning at normal speed
Fans SFT 3 Fan 1 Rotor 0	OK	Spinning at normal speed
Fans SFT 3 Fan 1 Rotor 1	OK	Spinning at normal speed
Fans SFT 3 Fan 2 Rotor 0	OK	Spinning at normal speed
Fans SFT 3 Fan 2 Rotor 1	OK	Spinning at normal speed
Fans SFT 3 Fan 3 Rotor 0	OK	Spinning at normal speed
Fans SFT 3 Fan 3 Rotor 1	OK	Spinning at normal speed
Fans SFT 4 Fan 0 Rotor 0	OK	Spinning at normal speed
Fans SFT 4 Fan 0 Rotor 1	OK	Spinning at normal speed
Fans SFT 4 Fan 1 Rotor 0	OK	Spinning at normal speed
Fans SFT 4 Fan 1 Rotor 1	OK	Spinning at normal speed
Fans SFT 4 Fan 2 Rotor 0	OK	Spinning at normal speed
Fans SFT 4 Fan 2 Rotor 1	OK	Spinning at normal speed
Fans SFT 4 Fan 3 Rotor 0	OK	Spinning at normal speed
Fans SFT 4 Fan 3 Rotor 1	OK	Spinning at normal speed
Fans SFT 5 Fan 0 Rotor 0	OK	Spinning at normal speed
Fans SFT 5 Fan 0 Rotor 1	OK	Spinning at normal speed

Fans	SFT 5	Fan 1	Rotor 0	OK	Spinning at normal speed
Fans	SFT 5	Fan 1	Rotor 1	OK	Spinning at normal speed
Fans	SFT 5	Fan 2	Rotor 0	OK	Spinning at normal speed
Fans	SFT 5	Fan 2	Rotor 1	OK	Spinning at normal speed
Fans	SFT 5	Fan 3	Rotor 0	OK	Spinning at normal speed
Fans	SFT 5	Fan 3	Rotor 1	OK	Spinning at normal speed
Fans	SFT 6	Fan 0	Rotor 0	OK	Spinning at normal speed
Fans	SFT 6	Fan 0	Rotor 1	OK	Spinning at normal speed
Fans	SFT 6	Fan 1	Rotor 0	OK	Spinning at normal speed
Fans	SFT 6	Fan 1	Rotor 1	OK	Spinning at normal speed
Fans	SFT 6	Fan 2	Rotor 0	OK	Spinning at normal speed
Fans	SFT 6	Fan 2	Rotor 1	OK	Spinning at normal speed
Fans	SFT 6	Fan 3	Rotor 0	OK	Spinning at normal speed
Fans	SFT 6	Fan 3	Rotor 1	OK	Spinning at normal speed
Fans	SFT 7	Fan 0	Rotor 0	OK	Spinning at normal speed
Fans	SFT 7	Fan 0	Rotor 1	OK	Spinning at normal speed
Fans	SFT 7	Fan 1	Rotor 0	OK	Spinning at normal speed
Fans	SFT 7	Fan 1	Rotor 1	OK	Spinning at normal speed
Fans	SFT 7	Fan 2	Rotor 0	OK	Spinning at normal speed
Fans	SFT 7	Fan 2	Rotor 1	OK	Spinning at normal speed
Fans	SFT 7	Fan 3	Rotor 0	OK	Spinning at normal speed
Fans	SFT 7	Fan 3	Rotor 1	OK	Spinning at normal speed
Power	PEM 0			OK	30 degrees C / 86 degrees F
Power	PEM 1			OK	30 degrees C / 86 degrees F
Power	PEM 2			OK	30 degrees C / 86 degrees F
Power	PEM 3			Absent	
Power	PEM 4			Absent	
Power	PEM 5			Absent	

show chassis environment node-device (QFabric System)

```

user@switch> show chassis environment node-device node1
Class Item                               Status Measurement
Power node1 Power Supply 0              Absent
      node1 Power Supply 1              Absent
Fans  node1 Fan Tray 0                  Testing
      node1 Fan Tray 1                  Testing
      node1 Fan Tray 2                  Testing

```

show chassis environment pem node-device (QFabric System)

```

user@switch> show chassis environment pem node-device node1
FPC 0 PEM 0 status:
  State           Check
  Airflow         Front to Back
  Temperature     OK
  AC Input:       OK
  DC Output       Voltage(V) Current(A) Power(W) Load(%)
                  12          10        120      18
FPC 0 PEM 1 status:
  State           Online
  Airflow         Back to Front
  Temperature     OK
  AC Input:       OK
  DC Output       Voltage(V) Current(A) Power(W) Load(%)
                  11          10        110      17

```

show chassis environment (PTX5000 Packet Transport Router)

```

user@host> show chassis environment
Class Item                               Status Measurement
Temp PDU 0                               OK

```

PDU 0 PSM 0	OK	36 degrees C / 96 degrees F
PDU 0 PSM 1	OK	38 degrees C / 100 degrees F
PDU 0 PSM 2	OK	38 degrees C / 100 degrees F
PDU 0 PSM 3	OK	37 degrees C / 98 degrees F
PDU 1	Absent	
CCG 0	OK	44 degrees C / 111 degrees F
CCG 1	OK	44 degrees C / 111 degrees F
Routing Engine 0	OK	62 degrees C / 143 degrees F
Routing Engine 0 CPU	OK	75 degrees C / 167 degrees F
Routing Engine 1	OK	51 degrees C / 123 degrees F
Routing Engine 1 CPU	OK	64 degrees C / 147 degrees F
CB 0 Intake	OK	38 degrees C / 100 degrees F
CB 0 Exhaust A	OK	46 degrees C / 114 degrees F
CB 0 Exhaust B	OK	42 degrees C / 107 degrees F
CB 1 Intake	OK	35 degrees C / 95 degrees F
CB 1 Exhaust A	OK	39 degrees C / 102 degrees F
CB 1 Exhaust B	OK	36 degrees C / 96 degrees F
SIB 0 Exhaust	OK	47 degrees C / 116 degrees F
SIB 0 Junction	OK	45 degrees C / 113 degrees F
SIB 1 Exhaust	OK	44 degrees C / 111 degrees F
SIB 1 Junction	OK	43 degrees C / 109 degrees F
SIB 2 Exhaust	OK	47 degrees C / 116 degrees F
SIB 2 Junction	OK	42 degrees C / 107 degrees F
SIB 3 Exhaust	OK	43 degrees C / 109 degrees F
SIB 3 Junction	OK	43 degrees C / 109 degrees F
SIB 4 Exhaust	OK	47 degrees C / 116 degrees F
SIB 4 Junction	OK	42 degrees C / 107 degrees F
SIB 5 Exhaust	OK	42 degrees C / 107 degrees F
SIB 5 Junction	OK	40 degrees C / 104 degrees F
SIB 6 Exhaust	OK	46 degrees C / 114 degrees F
SIB 6 Junction	OK	42 degrees C / 107 degrees F
SIB 7 Exhaust	OK	43 degrees C / 109 degrees F
SIB 7 Junction	OK	39 degrees C / 102 degrees F
SIB 8 Exhaust	OK	44 degrees C / 111 degrees F
SIB 8 Junction	OK	41 degrees C / 105 degrees F
FPC 0 PMB	OK	35 degrees C / 95 degrees F
FPC 0 Intake	OK	33 degrees C / 91 degrees F
FPC 0 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 0 Exhaust B	OK	43 degrees C / 109 degrees F
FPC 0 TL0	OK	48 degrees C / 118 degrees F
FPC 0 TQ0	OK	53 degrees C / 127 degrees F
FPC 0 TL1	OK	56 degrees C / 132 degrees F
FPC 0 TQ1	OK	58 degrees C / 136 degrees F
FPC 0 TL2	OK	55 degrees C / 131 degrees F
FPC 0 TQ2	OK	56 degrees C / 132 degrees F
FPC 0 TL3	OK	59 degrees C / 138 degrees F
FPC 0 TQ3	OK	59 degrees C / 138 degrees F
FPC 2 PMB	OK	35 degrees C / 95 degrees F
FPC 2 Intake	OK	34 degrees C / 93 degrees F
FPC 2 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 2 Exhaust B	OK	52 degrees C / 125 degrees F
FPC 2 TL0	OK	53 degrees C / 127 degrees F
FPC 2 TQ0	OK	53 degrees C / 127 degrees F
FPC 2 TL1	OK	57 degrees C / 134 degrees F
FPC 2 TQ1	OK	58 degrees C / 136 degrees F
FPC 2 TL2	OK	54 degrees C / 129 degrees F
FPC 2 TQ2	OK	59 degrees C / 138 degrees F
FPC 2 TL3	OK	60 degrees C / 140 degrees F
FPC 2 TQ3	OK	64 degrees C / 147 degrees F
PIC 2/0 Ambient	OK	49 degrees C / 120 degrees F
FPC 3 PMB	OK	34 degrees C / 93 degrees F

FPC 3 Intake	OK	35 degrees C / 95 degrees F
FPC 3 Exhaust A	OK	54 degrees C / 129 degrees F
FPC 3 Exhaust B	OK	49 degrees C / 120 degrees F
FPC 3 TL0	OK	49 degrees C / 120 degrees F
FPC 3 TQ0	OK	55 degrees C / 131 degrees F
FPC 3 TL1	OK	56 degrees C / 132 degrees F
FPC 3 TQ1	OK	58 degrees C / 136 degrees F
FPC 3 TL2	OK	56 degrees C / 132 degrees F
FPC 3 TQ2	OK	59 degrees C / 138 degrees F
FPC 3 TL3	OK	62 degrees C / 143 degrees F
FPC 3 TQ3	OK	63 degrees C / 145 degrees F
PIC 3/1	Absent	
FPC 5 PMB	OK	35 degrees C / 95 degrees F
FPC 5 Intake	OK	34 degrees C / 93 degrees F
FPC 5 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 5 Exhaust B	OK	53 degrees C / 127 degrees F
FPC 5 TL0	OK	54 degrees C / 129 degrees F
FPC 5 TQ0	OK	52 degrees C / 125 degrees F
FPC 5 TL1	OK	61 degrees C / 141 degrees F
FPC 5 TQ1	OK	60 degrees C / 140 degrees F
FPC 5 TL2	OK	55 degrees C / 131 degrees F
FPC 5 TQ2	OK	55 degrees C / 131 degrees F
FPC 5 TL3	OK	59 degrees C / 138 degrees F
FPC 5 TQ3	OK	58 degrees C / 136 degrees F
PIC 5/0 Ambient	OK	51 degrees C / 123 degrees F
PIC 5/1 Ambient	OK	34 degrees C / 93 degrees F
PIC 5/1 cfp-5/1/0	OK	34 degrees C / 93 degrees F
PIC 5/1 cfp-5/1/1	OK	36 degrees C / 96 degrees F
FPC 6 PMB	OK	36 degrees C / 96 degrees F
FPC 6 Intake	OK	33 degrees C / 91 degrees F
FPC 6 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 6 Exhaust B	OK	39 degrees C / 102 degrees F
FPC 6 TL0	OK	44 degrees C / 111 degrees F
FPC 6 TQ0	OK	54 degrees C / 129 degrees F
FPC 6 TL1	OK	59 degrees C / 138 degrees F
FPC 6 TQ1	OK	58 degrees C / 136 degrees F
FPC 6 TL2	OK	60 degrees C / 140 degrees F
FPC 6 TQ2	OK	57 degrees C / 134 degrees F
FPC 6 TL3	OK	65 degrees C / 149 degrees F
FPC 6 TQ3	OK	60 degrees C / 140 degrees F
FPC 7 PMB	OK	35 degrees C / 95 degrees F
FPC 7 Intake	OK	33 degrees C / 91 degrees F
FPC 7 Exhaust A	OK	53 degrees C / 127 degrees F
FPC 7 Exhaust B	OK	40 degrees C / 104 degrees F
FPC 7 TL0	OK	46 degrees C / 114 degrees F
FPC 7 TQ0	OK	58 degrees C / 136 degrees F
FPC 7 TL1	OK	53 degrees C / 127 degrees F
FPC 7 TQ1	OK	59 degrees C / 138 degrees F
FPC 7 TL2	OK	56 degrees C / 132 degrees F
FPC 7 TQ2	OK	61 degrees C / 141 degrees F
FPC 7 TL3	OK	63 degrees C / 145 degrees F
FPC 7 TQ3	OK	63 degrees C / 145 degrees F
FPM I2CS	OK	37 degrees C / 98 degrees F
Fans Fan Tray 0 Fan 1	OK	3042 RPM
Fan Tray 0 Fan 2	OK	3042 RPM
Fan Tray 0 Fan 3	OK	3000 RPM
Fan Tray 0 Fan 4	OK	3042 RPM
Fan Tray 0 Fan 5	OK	3000 RPM
Fan Tray 0 Fan 6	OK	3042 RPM
Fan Tray 0 Fan 7	OK	3085 RPM
Fan Tray 0 Fan 8	OK	3042 RPM

Fan Tray 0 Fan 9	OK	3042 RPM
Fan Tray 0 Fan 10	OK	3085 RPM
Fan Tray 0 Fan 11	OK	3085 RPM
Fan Tray 0 Fan 12	OK	3128 RPM
Fan Tray 0 Fan 13	OK	3128 RPM
Fan Tray 0 Fan 14	OK	3042 RPM
Fan Tray 1 Fan 1	OK	2299 RPM
Fan Tray 1 Fan 2	OK	2399 RPM
Fan Tray 1 Fan 3	OK	2299 RPM
Fan Tray 1 Fan 4	OK	2266 RPM
Fan Tray 1 Fan 5	OK	2266 RPM
Fan Tray 1 Fan 6	OK	2366 RPM
Fan Tray 2 Fan 1	OK	2199 RPM
Fan Tray 2 Fan 2	OK	2133 RPM
Fan Tray 2 Fan 3	OK	2366 RPM
Fan Tray 2 Fan 4	OK	2233 RPM
Fan Tray 2 Fan 5	OK	2399 RPM
Fan Tray 2 Fan 6	OK	2233 RPM
Misc SPMB 0 Intake	OK	50 degrees C / 122 degrees F
SPMB 1 Intake	OK	40 degrees C / 104 degrees F

show chassis environment (PTX5000 Packet Transport Router with FPC2-PTX-P1A)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Temp	PDU 0	OK	
	PDU 0 PSM 0	OK	41 degrees C / 105 degrees F
	PDU 0 PSM 1	Absent	
	PDU 0 PSM 2	OK	43 degrees C / 109 degrees F
	PDU 0 PSM 3	Absent	
	PDU 0 PSM 4	OK	44 degrees C / 111 degrees F
	PDU 0 PSM 5	Absent	
	PDU 0 PSM 6	OK	45 degrees C / 113 degrees F
	PDU 0 PSM 7	Absent	
	PDU 1	OK	
	PDU 1 PSM 0	Absent	
	PDU 1 PSM 1	OK	45 degrees C / 113 degrees F
	PDU 1 PSM 2	Absent	
	PDU 1 PSM 3	OK	43 degrees C / 109 degrees F
	PDU 1 PSM 4	Absent	
	PDU 1 PSM 5	OK	46 degrees C / 114 degrees F
	PDU 1 PSM 6	Absent	
	PDU 1 PSM 7	OK	46 degrees C / 114 degrees F
	CCG 0	OK	27 degrees C / 80 degrees F
	CCG 1	OK	29 degrees C / 84 degrees F
	...		

show chassis environment (PTX1000 Packet Transport Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Power	FPC 0 Power Supply 0	Absent	
	FPC 0 Power Supply 1	Absent	
	FPC 0 Power Supply 2	OK	
	FPC 0 Power Supply 3	OK	
Temp	FPC 0 Intake Temp Sensor	OK	25 degrees C / 77 degrees F
	FPC 0 Exhaust Temp Sensor	OK	35 degrees C / 95 degrees F
	FPC 0 Mezz Temp Sensor 0	OK	25 degrees C / 77 degrees F
	FPC 0 Mezz Temp Sensor 1	OK	34 degrees C / 93 degrees F
	FPC 0 PE2 Temp Sensor	OK	34 degrees C / 93 degrees F
	FPC 0 PE1 Temp Sensor	OK	32 degrees C / 89 degrees F

	FPC 0 PF0 Temp Sensor	OK	40 degrees C / 104 degrees F
	FPC 0 PE0 Temp Sensor	OK	33 degrees C / 91 degrees F
	FPC 0 PE5 Temp Sensor	OK	34 degrees C / 93 degrees F
	FPC 0 PE4 Temp Sensor	OK	34 degrees C / 93 degrees F
	FPC 0 PF1 Temp Sensor	OK	41 degrees C / 105 degrees F
	FPC 0 PE3 Temp Sensor	OK	36 degrees C / 96 degrees F
	FPC 0 CPU Die Temp Sensor	OK	40 degrees C / 104 degrees F
	FPC 0 OCX0 Temp Sensor	OK	37 degrees C / 98 degrees F
Fans	FPC 0 Fan Tray 0	OK	Spinning at normal speed
	FPC 0 Fan Tray 1	OK	Spinning at normal speed
	FPC 0 Fan Tray 2	OK	Spinning at normal speed

show chassis environment (ACX2000 Universal Access Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
	PCB Left	OK	44 degrees C / 111 degrees F
	SFP+ Xcvr	OK	50 degrees C / 122 degrees F
	FEB	OK	70 degrees C / 158 degrees F
	PCB Up	OK	63 degrees C / 145 degrees F
	PCB Mid	OK	66 degrees C / 150 degrees F
	Telecom Mod	OK	65 degrees C / 149 degrees F
	Routing Engine	OK	54 degrees C / 129 degrees F
	Heater off		

show chassis environment (ACX4000 Universal Access Router)

On the ACX4000 router, the MIC output of the **show chassis environment** command varies depending on the number of temperature channels present in the installed MIC.

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Temp	PEM 0	OK	33 degrees C / 91 degrees F
	PEM 1	Absent	
	PCB Bottom	OK	30 degrees C / 86 degrees F
	PCB Middle	OK	34 degrees C / 93 degrees F
	BCM56445	OK	33 degrees C / 91 degrees F
	SFP+ Xcvr	OK	32 degrees C / 89 degrees F
	Fan tray inlet	OK	39 degrees C / 102 degrees F
	Exhaust	OK	30 degrees C / 86 degrees F
	Routing Engine	OK	32 degrees C / 89 degrees F
Pic	Heater off		
	PIC 0/0 Channel 0	OK	28 degrees C / 82 degrees F
	PIC 0/0 Channel 1	OK	29 degrees C / 84 degrees F
	PIC 0/0 Channel 2	OK	0 degrees C / 32 degrees F
	PIC 0/0 Channel 3	OK	0 degrees C / 32 degrees F
	PIC 0/0 Channel 4	OK	0 degrees C / 32 degrees F
	PIC 0/0 Channel 5	OK	0 degrees C / 32 degrees F
	PIC 0/0 Channel 6	OK	0 degrees C / 32 degrees F
	PIC 0/0 Channel 7	OK	0 degrees C / 32 degrees F
	PIC 0/0 Channel 8	OK	0 degrees C / 32 degrees F
	PIC 0/0 Channel 9	OK	0 degrees C / 32 degrees F
	PIC 1/0 Channel 0	OK	33 degrees C / 91 degrees F
	PIC 1/0 Channel 1	OK	31 degrees C / 87 degrees F
	PIC 1/0 Channel 2	OK	30 degrees C / 86 degrees F
	PIC 1/0 Channel 3	OK	0 degrees C / 32 degrees F
	PIC 1/0 Channel 4	OK	0 degrees C / 32 degrees F
	PIC 1/0 Channel 5	OK	0 degrees C / 32 degrees F
	PIC 1/0 Channel 6	OK	0 degrees C / 32 degrees F
	PIC 1/0 Channel 7	OK	0 degrees C / 32 degrees F

	PIC 1/0 Channel 8	OK	0 degrees C / 32 degrees F
	PIC 1/1 Channel 0	OK	31 degrees C / 87 degrees F
	PIC 1/1 Channel 1	OK	29 degrees C / 84 degrees F
	PIC 1/1 Channel 2	OK	28 degrees C / 82 degrees F
	PIC 1/1 Channel 3	OK	0 degrees C / 32 degrees F
	PIC 1/1 Channel 4	OK	0 degrees C / 32 degrees F
	PIC 1/1 Channel 5	OK	0 degrees C / 32 degrees F
	PIC 1/1 Channel 6	OK	0 degrees C / 32 degrees F
	PIC 1/1 Channel 7	OK	0 degrees C / 32 degrees F
	PIC 1/1 Channel 8	OK	0 degrees C / 32 degrees F
Fans	Fan 1	OK	Spinning at normal speed
	Fan 2	OK	Spinning at normal speed

show chassis environment fpc

List of Syntax	Syntax on page 540 Syntax (TX Matrix and TX Matrix Plus Routers) on page 540 Syntax (MX Series Routers) on page 540 Syntax (MX2010 3D Universal Edge Routers) on page 540 Syntax (MX2020 3D Universal Edge Routers) on page 540 Syntax (QFX Series) on page 540 Syntax (OCX Series) on page 540
Syntax	show chassis environment fpc <slot>
Syntax (TX Matrix and TX Matrix Plus Routers)	show chassis environment fpc <lcc number> <slot>
Syntax (MX Series Routers)	show chassis environment fpc <slot> <all-members> <local> <member member-id>
Syntax (MX2010 3D Universal Edge Routers)	show chassis environment fpc <slot>
Syntax (MX2020 3D Universal Edge Routers)	show chassis environment fpc <slot> <satellite [slot-id slot-id [device-alias alias-name]]
Syntax (QFX Series)	show chassis environment fpc <fpc-slot> interconnect-device name
Syntax (OCX Series)	show chassis environment fpc <fpc-slot>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for QFX Series. Command introduced in Junos OS Release 12.1x48 for PTX Series Packet Transport Routers. Command introduced in Junos OS Release 12.1 for T4000 Core Routers. Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers. Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series. satellite option introduced in Junos OS Release 14.2R3.
Description	(M40e, M120, M160, M320, MX Series, T Series routers, EX Series, QFX Series, and PTX Series routers only) Display environmental information about Flexible PIC Concentrators (FPCs).

Options **none**—Display environmental information about all FPCs. On a TX Matrix router, display environmental information about all FPCs on the TX Matrix router and its attached T640 routers. On a TX Matrix Plus router, display environmental information about all FPCs on the TX Matrix Plus router and its attached routers.

all-members—(MX Series routers only) (Optional) Display environmental information for the FPCs in all the members of the Virtual Chassis configuration.

interconnect-device *name*—(QFabric systems only) (Optional) Display chassis environmental information for the Interconnect device.

lcc *number*—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(MX Series routers only) (Optional) Display environmental information for the FPCs in the local Virtual Chassis member.

member *member-id*—(MX Series routers only) (Optional) Display environmental information for the FPCs in the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

satellite [*slot-id slot-id* | *device-alias alias-name*]—(Junos Fusion only) (Optional) Display environmental information for the FPCs in the specified satellite device in a Junos Fusion, or for all satellite devices in the Junos Fusion if no satellite devices are specified.

slot* or *fpc-slot—(Optional) Display environmental information about an individual FPC:

- (TX Matrix and TX Matrix Plus routers only) On a TX Matrix router, if you specify the number of the T640 router by using only the **lcc *number*** option (the recommended method), replace ***slot*** with a value from 0 through 7. Similarly, on a TX Matrix Plus router, if you specify the number of the router by using only the **lcc *number*** option (the recommended method), replace ***slot*** with a value from 0 through 7. Otherwise, replace ***slot*** with a value from 0 through 31. For example, the following commands have the same result:

```
user@host> show chassis environment fpc 1 lcc 1
user@host> show chassis environment fpc 9
```

- M120 router—Replace ***slot*** with a value from 0 through 5.

- MX240 router—Replace **slot** with a value from 0 through 2.
- MX480 router—Replace **slot** with a value from 0 through 5.
- MX960 router—Replace **slot** with a value from 0 through 11.
- MX2010 router—Replace **slot** with a value from 0 through 9.
- MX2020 router—Replace **slot** with a value from 0 through 19.
- Other routers—Replace **slot** with a value from 0 through 7.
- EX Series switches:
 - EX3200 switches and EX4200 standalone switches—Replace **slot** with 0.
 - EX4200 switches in a Virtual Chassis configuration—Replace **slot** with a value from 0 through 9 (switch's member ID).
 - EX6210 switches—Replace **slot** with a value from 0 through 3 (line card only), 4 or 5 (line card or Switch Fabric and Rotuing Engine (SRE) module), or 6 through 9 (line card only).
 - EX8208 switches—Replace **slot** with a value from 0 through 7 (line card).
 - EX8216 switches—Replace **slot** with a value from 0 through 15 (line card).
- QFX3500 switches —Replace **fpc-slot** with 0 through 15.
- PTX5000 Packet Transport Router—Replace **fpc-slot** with 0 through 7.

Required Privilege Level view

- Related Documentation**
- [request chassis fpc on page 366](#)
 - [show chassis fpc on page 703](#)
 - *show chassis fpc-feb-connectivity*
 - *Configuring the Junos OS to Resynchronize FPC Sequence Numbers with Active FPCs when an FPC Comes Online*
 - *MX960 Flexible PIC Concentrator Description*

- List of Sample Output**
- [show chassis environment fpc \(M120 Router\) on page 544](#)
 - [show chassis environment fpc \(M160 Router\) on page 545](#)
 - [show chassis environment fpc \(M320 Router\) on page 545](#)
 - [show chassis environment fpc \(MX2020 Router\) on page 546](#)
 - [show chassis environment fpc \(MX2010 Router\) on page 549](#)
 - [show chassis environment fpc \(MX240 Router\) on page 552](#)
 - [show chassis environment fpc \(MX480 Router\) on page 553](#)
 - [show chassis environment fpc \(MX960 Router\) on page 553](#)
 - [show chassis environment fpc \(MX480 Router with 100-Gigabit Ethernet CFP\) on page 554](#)
 - [show chassis environment fpc \(MX240, MX480, MX960 with Application Services Modular Line Card on page 555](#)

[show chassis environment fpc \(T320, T640, and T1600 Routers\) on page 556](#)
[show chassis environment fpc \(T4000 Router\) on page 557](#)
[show chassis environment fpc lcc \(TX Matrix Router\) on page 561](#)
[show chassis environment fpc lcc \(TX Matrix Plus Router\) on page 562](#)
[show chassis environment fpc \(QFX Series and OCX Series\) on page 563](#)
[show chassis environment fpc interconnect-device \(QFabric Systems\) on page 563](#)
[show chassis environment fpc 0 \(PTX5000 Packet Transport Router\) on page 563](#)
[show chassis environment fpc 07 \(PTX5000 Packet Transport Router with FPC2-PTX-P1A\) on page 564](#)
[show chassis environment FPC 1 \(MX Routers with Media Services Blade \[MSB\]\) on page 565](#)

Output Fields Table 47 lists the output fields for the **show chassis environment fpc** command. Output fields are listed in the approximate order in which they appear.

Table 47: show chassis environment fpc Output Fields

Field Name	Field Description
State	<p>Status of the FPC:</p> <ul style="list-style-type: none"> • Unknown—FPC is not detected by the router. • Empty—No FPC is present. • Present—FPC is detected by the chassis daemon but is either not supported by the current version of the Junos OS, or the FPC is coming up but not yet online. • Ready—FPC is in intermediate or transition state. • Announce online—Intermediate state during which the FPC is coming up but not yet online, and the chassis manager acknowledges the chassisd FPC online initiative. • Online—FPC is online and running. • Offline—FPC is powered down. • Diagnostics—FPC is set to operate in diagnostics mode.
Temperature	(M40e and M160 routers and QFX Series only) Temperature of the air flowing past the FPC.
PMB Temperature	<p>(PTX Series only) Temperature of the air flowing past the PMB (bottom of the FPC).</p> <p>The PTX5000 Packet Transport Router with FPC2-PTX-P1A include multiple temperatures for PMB (TEMPO and TEMP1).</p>
PMB CPU Temperature	(PTX5000 Packet Transport Router with FPC2-PTX-P1A only) Temperature of the air flowing past the PMB CPU.
Temperature Intake	(M320 routers, MX2010 routers, MX2020 routers, and PTX Series only) Temperature of the air flowing into the chassis.
Temperature Top	(T Series routers only) Temperature of the air flowing past the top of the FPC.
Temperature Exhaust	<p>(M120 and M320 routers, MX2010 routers, MX2020 routers, and PTX Series only) Temperature of the air flowing out of the chassis.</p> <p>The PTX Series Packet Transport Routers, and the MX2010 and MX2020 routers include exhaust temperatures for multiple zones (Exhaust A and Exhaust B).</p>

Table 47: show chassis environment fpc Output Fields (*continued*)

Field Name	Field Description
Temperature Bottom	(T Series routers only) Temperature of the air flowing past the bottom of the FPC.
TL <i>n</i> Temperature	(PTX Series only) Temperature of the air flowing past the specified TL area of the packet forwarding engine (PFE) on the FPC.
TQ <i>n</i> Temperature	(PTX Series only) Temperature of the air flowing past the specified TQ area of the packet forwarding engine (PFE) on the FPC.
Temperature MMBO	(T640 router only) Temperature of the air flowing past the type 3 FPC.
Temperature MMB1	(M320 and T Series routers only) Temperature of the air flowing past the type 1, type 2, and type 3 FPC.
Power	Information about the voltage supplied to the FPC. The left column displays the required power, in volts. The right column displays the measured power, in millivolts.
CMB Revision or BUS revision	Revision level of the chassis management bus device (M Series router) or bus (T Series routers).

Sample Output

show chassis environment fpc (M120 Router)

```

user@host> show chassis environment fpc
FPC 2 status:
  State                               Online
  Temperature Exhaust A               32 degrees C / 89 degrees F
  Temperature Exhaust B               31 degrees C / 87 degrees F
  Power A-Board
    1.2 V                             1202 mV
    1.5 V                             1508 mV
    1.8 V                             1798 mV
    2.5 V                             2507 mV
    3.3 V                             3351 mV
    5.0 V                             4995 mV
    3.3 V bias                         3296 mV
    1.2 V Rocket IO                   1205 mV
    1.5 V Rocket IO                   1501 mV
  I2C Slave Revision                 12
FPC 3 status:
  State                               Online
  Temperature Exhaust A               31 degrees C / 87 degrees F
  Temperature Exhaust B               33 degrees C / 91 degrees F
  Power A-Board
    1.2 V                             1211 mV
    1.5 V                             1501 mV
    1.8 V                             1798 mV
    2.5 V                             2471 mV
    3.3 V                             3293 mV
    5.0 V                             4930 mV
    3.3 V bias                         3296 mV
    1.2 V Rocket IO                   1205 mV
    1.5 V Rocket IO                   1501 mV

```



```

Power B-Board
  1.2 V          1214 mV
  1.5 V          1501 mV
  2.5 V          2471 mV
  3.3 V          3300 mV
  5.0 V          4943 mV
  3.3 V bias     3296 mV
  1.2 V Rocket IO 1205 mV
  1.5 V Rocket IO 1501 mV
I2C Slave Revision 12
FPC 4 status:
State              Online
Temperature Exhaust A 32 degrees C / 89 degrees F
Temperature Exhaust B 30 degrees C / 86 degrees F
Power A-Board
  1.2 V          1195 mV
  1.5 V          1504 mV
  1.8 V          1801 mV
  2.5 V          2504 mV
  3.3 V          3293 mV
  5.0 V          4917 mV
  3.3 V bias     3296 mV
  1.2 V Rocket IO 1202 mV
  1.5 V Rocket IO 1492 mV
I2C Slave Revision 12

```

show chassis environment fpc (M160 Router)

```

user@host> show chassis environment fpc
FPC 0 status:
State              Online
Temperature         42 degrees C / 107 degrees F
Power:
  1.5 V             1500 mV
  2.5 V             2509 mV
  3.3 V             3308 mV
  5.0 V             4991 mV
  5.0 V bias        4952 mV
  8.0 V bias        8307 mV
CMB Revision        12
FPC 1 status:
State              Online
Temperature         45 degrees C / 113 degrees F
Power:
  1.5 V             1498 mV
  2.5 V             2501 mV
  3.3 V             3319 mV
  5.0 V             5020 mV
  5.0 V bias        5025 mV
  8.0 V bias        8307 mV
CMB Revision        12

```

show chassis environment fpc (M320 Router)

```

user@host> show chassis environment fpc
FPC 0 status:
State              Online
Temperature Intake   27 degrees C / 80 degrees F
Temperature Exhaust  38 degrees C / 100 degrees F
Temperature MMB1     31 degrees C / 87 degrees F
Power:

```

```

1.5 V          1487 mV
1.5 V *        1494 mV
1.8 V          1821 mV
2.5 V          2533 mV
3.3 V          3323 mV
5.0 V          5028 mV
3.3 V bias     3296 mV
5.0 V bias     4984 mV
CMB Revision   16
FPC 1 status:
State          Online
Temperature Intake    27 degrees C / 80 degrees F
Temperature Exhaust   37 degrees C / 98 degrees F
Temperature MMB1      32 degrees C / 89 degrees F
Power:
1.5 V          1504 mV
1.5 V *        1499 mV
1.8 V          1820 mV
2.5 V          2529 mV
3.3 V          3328 mV
5.0 V          5013 mV
3.3 V bias     3294 mV
5.0 V bias     4984 mV
CMB Revision   16
FPC 2 status:
State          Online
Temperature Intake    28 degrees C / 82 degrees F
Temperature Exhaust   38 degrees C / 100 degrees F
Temperature MMB1      32 degrees C / 89 degrees F
Power:
1.5 V          1498 mV
1.5 V *        1487 mV
1.8 V          1816 mV
2.5 V          2531 mV
3.3 V          3324 mV
5.0 V          5025 mV
3.3 V bias     3277 mV
5.0 V bias     5013 mV
CMB Revision   17
FPC 3 status:
...

```

show chassis environment fpc (MX2020 Router)

```

user@host> show chassis environment fpc
FPC 0 status:
State          Online
Temperature Intake    41 degrees C / 105 degrees F
Temperature Exhaust A 48 degrees C / 118 degrees F
Temperature Exhaust B 60 degrees C / 140 degrees F
Temperature LU 0 TSen 56 degrees C / 132 degrees F
Temperature LU 0 Chip 59 degrees C / 138 degrees F
Temperature LU 1 TSen 56 degrees C / 132 degrees F
Temperature LU 1 Chip 61 degrees C / 141 degrees F
Temperature LU 2 TSen 56 degrees C / 132 degrees F
Temperature LU 2 Chip 52 degrees C / 125 degrees F
Temperature LU 3 TSen 56 degrees C / 132 degrees F
Temperature LU 3 Chip 52 degrees C / 125 degrees F
Temperature MQ 0 TSen 49 degrees C / 120 degrees F
Temperature MQ 0 Chip 49 degrees C / 120 degrees F
Temperature MQ 1 TSen 49 degrees C / 120 degrees F

```

```

Temperature MQ 1 Chip      52 degrees C / 125 degrees F
Temperature MQ 2 TSen      49 degrees C / 120 degrees F
Temperature MQ 2 Chip      45 degrees C / 113 degrees F
Temperature MQ 3 TSen      49 degrees C / 120 degrees F
Temperature MQ 3 Chip      46 degrees C / 114 degrees F
Power
AS-BIAS3V3-z12105         3299 mV
AS-VDD1V8-z12006          1807 mV
AS-VDD2V5-z12006          2512 mV
AS-AVDD1V0-z12004         997 mV
AS-PCIE_1V0-z12004        996 mV
AS-VDD3V3-z12004          3294 mV
AS-VDD_1V5A-z12004        1501 mV
AS-VDD_1V5B-z12004        1498 mV
AS-LU0_1V0-z12004         998 mV
AS-LU1_1V0-z12004        1002 mV
AS-MQ0_1V0-z12004         999 mV
AS-MQ1_1V0-z12004         994 mV
AS-LU2_1V0-z12004        1000 mV
AS-LU3_1V0-z12004         998 mV
AS-MQ2_1V0-z12004        1002 mV
AS-MQ3_1V0-z12004         999 mV
AS-PMB_1V1-z12006        1096 mV
I2C Slave Revision        68
FPC 1 status:
State                      Online
Temperature Intake          39 degrees C / 102 degrees F
Temperature Exhaust A       48 degrees C / 118 degrees F
Temperature Exhaust B       55 degrees C / 131 degrees F
Temperature LU 0 TSen       52 degrees C / 125 degrees F
Temperature LU 0 Chip       54 degrees C / 129 degrees F
Temperature LU 1 TSen       52 degrees C / 125 degrees F
Temperature LU 1 Chip       56 degrees C / 132 degrees F
Temperature LU 2 TSen       52 degrees C / 125 degrees F
Temperature LU 2 Chip       49 degrees C / 120 degrees F
Temperature LU 3 TSen       52 degrees C / 125 degrees F
Temperature LU 3 Chip       50 degrees C / 122 degrees F
Temperature MQ 0 TSen       48 degrees C / 118 degrees F
Temperature MQ 0 Chip       48 degrees C / 118 degrees F
Temperature MQ 1 TSen       48 degrees C / 118 degrees F
Temperature MQ 1 Chip       51 degrees C / 123 degrees F
Temperature MQ 2 TSen       48 degrees C / 118 degrees F
Temperature MQ 2 Chip       45 degrees C / 113 degrees F
Temperature MQ 3 TSen       48 degrees C / 118 degrees F
Temperature MQ 3 Chip       45 degrees C / 113 degrees F
Power
AS-BIAS3V3-z12105         3291 mV
AS-VDD1V8-z12006          1786 mV
AS-VDD2V5-z12006          2496 mV
AS-AVDD1V0-z12004        1000 mV
AS-PCIE_1V0-z12004        1000 mV
AS-VDD3V3-z12004          3294 mV
AS-VDD_1V5A-z12004        1500 mV
AS-VDD_1V5B-z12004        1498 mV
AS-LU0_1V0-z12004        1003 mV
AS-LU1_1V0-z12004        1000 mV
AS-MQ0_1V0-z12004        1000 mV
AS-MQ1_1V0-z12004         995 mV
AS-LU2_1V0-z12004        1002 mV
AS-LU3_1V0-z12004         997 mV
AS-MQ2_1V0-z12004        1000 mV

```

```

AS-MQ3_1V0-z12004          998 mV
AS-PMB_1V1-z12006          1096 mV
I2C Slave Revision          68
FPC 2 status:
State                        Online
Temperature Intake           39 degrees C / 102 degrees F
Temperature Exhaust A        48 degrees C / 118 degrees F
Temperature Exhaust B        58 degrees C / 136 degrees F
Temperature LU 0 TSen         55 degrees C / 131 degrees F
Temperature LU 0 Chip         57 degrees C / 134 degrees F
Temperature LU 1 TSen         55 degrees C / 131 degrees F
Temperature LU 1 Chip         63 degrees C / 145 degrees F
Temperature LU 2 TSen         55 degrees C / 131 degrees F
Temperature LU 2 Chip         51 degrees C / 123 degrees F
Temperature LU 3 TSen         55 degrees C / 131 degrees F
Temperature LU 3 Chip         52 degrees C / 125 degrees F
Temperature MQ 0 TSen         48 degrees C / 118 degrees F
Temperature MQ 0 Chip         50 degrees C / 122 degrees F
Temperature MQ 1 TSen         48 degrees C / 118 degrees F
Temperature MQ 1 Chip         52 degrees C / 125 degrees F
Temperature MQ 2 TSen         48 degrees C / 118 degrees F
Temperature MQ 2 Chip         47 degrees C / 116 degrees F
Temperature MQ 3 TSen         48 degrees C / 118 degrees F
Temperature MQ 3 Chip         47 degrees C / 116 degrees F
Power
AS-BIAS3V3-z12105           3299 mV
AS-VDD1V8-z12006            1805 mV
AS-VDD2V5-z12006            2510 mV
AS-AVDD1V0-z12004           999 mV
AS-PCIE_1V0-z12004           998 mV
AS-VDD3V3-z12004            3296 mV
AS-VDD_1V5A-z12004          1492 mV
AS-VDD_1V5B-z12004          1497 mV
AS-LU0_1V0-z12004           997 mV
AS-LU1_1V0-z12004           1000 mV
AS-MQ0_1V0-z12004           998 mV
AS-MQ1_1V0-z12004           1001 mV
AS-LU2_1V0-z12004           996 mV
AS-LU3_1V0-z12004           995 mV
AS-MQ2_1V0-z12004           998 mV
AS-MQ3_1V0-z12004           997 mV
AS-PMB_1V1-z12006           1100 mV
I2C Slave Revision          68
FPC 3 status:
State                        Online
Temperature Intake           41 degrees C / 105 degrees F
Temperature Exhaust A        48 degrees C / 118 degrees F
Temperature Exhaust B        58 degrees C / 136 degrees F
Temperature LU 0 TSen         56 degrees C / 132 degrees F
Temperature LU 0 Chip         59 degrees C / 138 degrees F
Temperature LU 1 TSen         56 degrees C / 132 degrees F
Temperature LU 1 Chip         61 degrees C / 141 degrees F
Temperature LU 2 TSen         56 degrees C / 132 degrees F
Temperature LU 2 Chip         51 degrees C / 123 degrees F
Temperature LU 3 TSen         56 degrees C / 132 degrees F
Temperature LU 3 Chip         53 degrees C / 127 degrees F
Temperature MQ 0 TSen         50 degrees C / 122 degrees F
Temperature MQ 0 Chip         51 degrees C / 123 degrees F
Temperature MQ 1 TSen         50 degrees C / 122 degrees F
Temperature MQ 1 Chip         55 degrees C / 131 degrees F
Temperature MQ 2 TSen         50 degrees C / 122 degrees F

```

```

Temperature MQ 2 Chip      47 degrees C / 116 degrees F
Temperature MQ 3 TSen      50 degrees C / 122 degrees F
Temperature MQ 3 Chip      50 degrees C / 122 degrees F
Power
  AS-BIAS3V3-z12105        3305 mV
  AS-VDD1V8-z12006         1810 mV
  AS-VDD2V5-z12006         2508 mV
  AS-AVDD1V0-z12004         999 mV
  AS-PCIE_1V0-z12004        1001 mV
  AS-VDD3V3-z12004         3294 mV
  AS-VDD_1V5A-z12004        1500 mV
  AS-VDD_1V5B-z12004        1498 mV
  AS-LU0_1V0-z12004         998 mV
  AS-LU1_1V0-z12004         998 mV
  AS-MQ0_1V0-z12004         999 mV
  AS-MQ1_1V0-z12004         998 mV
  AS-LU2_1V0-z12004        1000 mV
  AS-LU3_1V0-z12004        1001 mV
  AS-MQ2_1V0-z12004         996 mV
  AS-MQ3_1V0-z12004         998 mV
  AS-PMB_1V1-z12006        1098 mV
I2C Slave Revision        68
FPC 4 status:
...

```

show chassis environment fpc (MX2010 Router)

```

user@host> show chassis environment fpc
FPC 0 status:
  State      Online
  Temperature Intake      36 degrees C / 96 degrees F
  Temperature Exhaust A   42 degrees C / 107 degrees F
  Temperature Exhaust B   51 degrees C / 123 degrees F
  Temperature LU 0 TSen    49 degrees C / 120 degrees F
  Temperature LU 0 Chip    50 degrees C / 122 degrees F
  Temperature LU 1 TSen    49 degrees C / 120 degrees F
  Temperature LU 1 Chip    54 degrees C / 129 degrees F
  Temperature LU 2 TSen    49 degrees C / 120 degrees F
  Temperature LU 2 Chip    45 degrees C / 113 degrees F
  Temperature LU 3 TSen    49 degrees C / 120 degrees F
  Temperature LU 3 Chip    46 degrees C / 114 degrees F
  Temperature MQ 0 TSen    40 degrees C / 104 degrees F
  Temperature MQ 0 Chip    41 degrees C / 105 degrees F
  Temperature MQ 1 TSen    40 degrees C / 104 degrees F
  Temperature MQ 1 Chip    44 degrees C / 111 degrees F
  Temperature MQ 2 TSen    40 degrees C / 104 degrees F
  Temperature MQ 2 Chip    38 degrees C / 100 degrees F
  Temperature MQ 3 TSen    40 degrees C / 104 degrees F
  Temperature MQ 3 Chip    41 degrees C / 105 degrees F
Power
  AS-BIAS3V3-z12105        3300 mV
  AS-VDD1V8-z12006         1805 mV
  AS-VDD2V5-z12006         2505 mV
  AS-AVDD1V0-z12004         998 mV
  AS-PCIE_1V0-z12004        999 mV
  AS-VDD3V3-z12004         3303 mV
  AS-VDD_1V5A-z12004        1497 mV
  AS-VDD_1V5B-z12004        1497 mV
  AS-LU0_1V0-z12004         998 mV
  AS-LU1_1V0-z12004        1003 mV
  AS-MQ0_1V0-z12004         998 mV

```

```

AS-MQ1_1V0-z12004      998 mV
AS-LU2_1V0-z12004      997 mV
AS-LU3_1V0-z12004      1001 mV
AS-MQ2_1V0-z12004      996 mV
AS-MQ3_1V0-z12004      994 mV
AS-PMB_1V1-z12006      1097 mV
I2C Slave Revision      68
FPC 1 status:
State                    Online
Temperature Intake       34 degrees C / 93 degrees F
Temperature Exhaust A    46 degrees C / 114 degrees F
Temperature Exhaust B    54 degrees C / 129 degrees F
Temperature LU 0 TSen     45 degrees C / 113 degrees F
Temperature LU 0 Chip     55 degrees C / 131 degrees F
Temperature LU 1 TSen     45 degrees C / 113 degrees F
Temperature LU 1 Chip     44 degrees C / 111 degrees F
Temperature LU 2 TSen     45 degrees C / 113 degrees F
Temperature LU 2 Chip     50 degrees C / 122 degrees F
Temperature LU 3 TSen     45 degrees C / 113 degrees F
Temperature LU 3 Chip     58 degrees C / 136 degrees F
Temperature XM 0 TSen     45 degrees C / 113 degrees F
Temperature XM 0 Chip     51 degrees C / 123 degrees F
Temperature XF 0 TSen     45 degrees C / 113 degrees F
Temperature XF 0 Chip     63 degrees C / 145 degrees F
Temperature PLX Switch TSen 45 degrees C / 113 degrees F
Temperature PLX Switch Chip 47 degrees C / 116 degrees F
Power
MPC-BIAS3V3-z12105      3300 mV
MPC-VDD3V3-z16100       3294 mV
MPC-VDD2V5-z16100       2505 mV
MPC-VDD1V8-z12004       1796 mV
MPC-AVDD1V0-z12004      991 mV
MPC-VDD1V2-z16100       1196 mV
MPC-VDD1V5A-z12004      1491 mV
MPC-VDD1V5B-z12004      1492 mV
MPC-XF_0V9-z12004       996 mV
MPC-PCIE_1V0-z16100     1003 mV
MPC-LU0_1V0-z12004      996 mV
MPC-LU1_1V0-z12004      996 mV
MPC-LU2_1V0-z12004      998 mV
MPC-LU3_1V0-z12004      994 mV
MPC-12VA-BMR453         12031 mV
MPC-12VB-BMR453         12003 mV
MPC-PMB_1V1-z12006      1104 mV
MPC-PMB_1V2-z12106      1194 mV
MPC-XM_0V9-vt273m       911 mV
I2C Slave Revision      110
FPC 8 status:
State                    Online
Temperature Intake       32 degrees C / 89 degrees F
Temperature Exhaust A    44 degrees C / 111 degrees F
Temperature Exhaust B    37 degrees C / 98 degrees F
Temperature LU 0 TCAM TSen 41 degrees C / 105 degrees F
Temperature LU 0 TCAM Chip 49 degrees C / 120 degrees F
Temperature LU 0 TSen     41 degrees C / 105 degrees F
Temperature LU 0 Chip     52 degrees C / 125 degrees F
Temperature MQ 0 TSen     41 degrees C / 105 degrees F
Temperature MQ 0 Chip     47 degrees C / 116 degrees F
Temperature LU 1 TCAM TSen 39 degrees C / 102 degrees F
Temperature LU 1 TCAM Chip 42 degrees C / 107 degrees F
Temperature LU 1 TSen     39 degrees C / 102 degrees F

```

```

Temperature LU 1 Chip      46 degrees C / 114 degrees F
Temperature MQ 1 TSen      39 degrees C / 102 degrees F
Temperature MQ 1 Chip      45 degrees C / 113 degrees F
Power
MPC-BIAS3V3-z12105        3296 mV
MPC-VDD3V3-z12006         3298 mV
MPC-VDD2V5-z12006         2505 mV
MPC-TCAM_1V0-z12004        997 mV
MPC-AVDD1V0-z12006        1007 mV
MPC-VDD1V8-z12006         1803 mV
MPC-PCIE_1V0-z12006        1004 mV
MPC-LU0_1V0-z12004         1000 mV
MPC-MQ0_1V0-z12004         999 mV
MPC-VDD_1V5-z12004        1498 mV
MPC-PMB_1V1-z12006        1102 mV
MPC-9VA-BMR453            9009 mV
MPC-9VB-BMR453            8960 mV
MPC-PMB_1V2-z12105        1202 mV
MPC-LU1_1V0-z12004        1005 mV
MPC-MQ1_1V0-z12004        1000 mV
I2C Slave Revision        70
FPC 9 status:
State                      Online
Temperature Intake          34 degrees C / 93 degrees F
Temperature Exhaust A       41 degrees C / 105 degrees F
Temperature Exhaust B       54 degrees C / 129 degrees F
Temperature LU 0 TSen        51 degrees C / 123 degrees F
Temperature LU 0 Chip        52 degrees C / 125 degrees F
Temperature LU 1 TSen        51 degrees C / 123 degrees F
Temperature LU 1 Chip        55 degrees C / 131 degrees F
Temperature LU 2 TSen        51 degrees C / 123 degrees F
Temperature LU 2 Chip        47 degrees C / 116 degrees F
Temperature LU 3 TSen        51 degrees C / 123 degrees F
Temperature LU 3 Chip        47 degrees C / 116 degrees F
Temperature MQ 0 TSen        40 degrees C / 104 degrees F
Temperature MQ 0 Chip        42 degrees C / 107 degrees F
Temperature MQ 1 TSen        40 degrees C / 104 degrees F
Temperature MQ 1 Chip        44 degrees C / 111 degrees F
Temperature MQ 2 TSen        40 degrees C / 104 degrees F
Temperature MQ 2 Chip        38 degrees C / 100 degrees F
Temperature MQ 3 TSen        40 degrees C / 104 degrees F
Temperature MQ 3 Chip        40 degrees C / 104 degrees F
Power
AS-BIAS3V3-z12105          3302 mV
AS-VDD1V8-z12006           1808 mV
AS-VDD2V5-z12006           2513 mV
AS-AVDD1V0-z12004           997 mV
AS-PCIE_1V0-z12004           999 mV
AS-VDD3V3-z12004            3294 mV
AS-VDD_1V5A-z12004          1503 mV
AS-VDD_1V5B-z12004          1502 mV
AS-LU0_1V0-z12004           996 mV
AS-LU1_1V0-z12004           999 mV
AS-MQ0_1V0-z12004           997 mV
AS-MQ1_1V0-z12004           999 mV
AS-LU2_1V0-z12004           997 mV
AS-LU3_1V0-z12004           998 mV
AS-MQ2_1V0-z12004           1000 mV
AS-MQ3_1V0-z12004           1000 mV
AS-PMB_1V1-z12006           1102 mV
I2C Slave Revision        68

```

show chassis environment fpc (MX240 Router)

user@host> show chassis environment fpc

FPC 1 status:

State	Online
Temperature Intake	34 degrees C / 93 degrees F
Temperature Exhaust A	39 degrees C / 102 degrees F
Temperature Exhaust B	53 degrees C / 127 degrees F
Temperature I3 0 TSensor	51 degrees C / 123 degrees F
Temperature I3 0 Chip	54 degrees C / 129 degrees F
Temperature I3 1 TSensor	50 degrees C / 122 degrees F
Temperature I3 1 Chip	53 degrees C / 127 degrees F
Temperature I3 2 TSensor	48 degrees C / 118 degrees F
Temperature I3 2 Chip	51 degrees C / 123 degrees F
Temperature I3 3 TSensor	45 degrees C / 113 degrees F
Temperature I3 3 Chip	48 degrees C / 118 degrees F
Temperature IA 0 TSensor	45 degrees C / 113 degrees F
Temperature IA 0 Chip	45 degrees C / 113 degrees F
Temperature IA 1 TSensor	45 degrees C / 113 degrees F
Temperature IA 1 Chip	49 degrees C / 120 degrees F

Power

1.5 V	1492 mV
2.5 V	2507 mV
3.3 V	3306 mV
1.8 V PFE 0	1801 mV
1.8 V PFE 1	1804 mV
1.8 V PFE 2	1798 mV
1.8 V PFE 3	1798 mV
1.2 V PFE 0	1169 mV
1.2 V PFE 1	1189 mV
1.2 V PFE 2	1182 mV
1.2 V PFE 3	1176 mV

I2C Slave Revision

42

FPC 2 status:

State	Online
Temperature Intake	33 degrees C / 91 degrees F
Temperature Exhaust A	41 degrees C / 105 degrees F
Temperature Exhaust B	53 degrees C / 127 degrees F
Temperature I3 0 TSensor	53 degrees C / 127 degrees F
Temperature I3 0 Chip	58 degrees C / 136 degrees F
Temperature I3 1 TSensor	52 degrees C / 125 degrees F
Temperature I3 1 Chip	56 degrees C / 132 degrees F
Temperature I3 2 TSensor	50 degrees C / 122 degrees F
Temperature I3 2 Chip	52 degrees C / 125 degrees F
Temperature I3 3 TSensor	46 degrees C / 114 degrees F
Temperature I3 3 Chip	49 degrees C / 120 degrees F
Temperature IA 0 TSensor	51 degrees C / 123 degrees F
Temperature IA 0 Chip	49 degrees C / 120 degrees F
Temperature IA 1 TSensor	48 degrees C / 118 degrees F
Temperature IA 1 Chip	53 degrees C / 127 degrees F

Power

1.5 V	1492 mV
2.5 V	2445 mV
3.3 V	3293 mV
1.8 V PFE 0	1827 mV
1.8 V PFE 1	1775 mV
1.8 V PFE 2	1788 mV
1.8 V PFE 3	1798 mV
1.2 V PFE 0	1250 mV
1.2 V PFE 1	1234 mV
1.2 V PFE 2	1231 mV


```

1.2 V PFE 3          1192 mV
I2C Slave Revision   42

```

show chassis environment fpc (MX480 Router)

```

user@host> show chassis environment fpc
FPC 1 status:
State                               Online
Temperature Intake                  36 degrees C / 96 degrees F
Temperature Exhaust A               41 degrees C / 105 degrees F
Temperature Exhaust B               55 degrees C / 131 degrees F
Temperature I3 0 TSensor            55 degrees C / 131 degrees F
Temperature I3 0 Chip               57 degrees C / 134 degrees F
Temperature I3 1 TSensor            53 degrees C / 127 degrees F
Temperature I3 1 Chip               53 degrees C / 127 degrees F
Temperature I3 2 TSensor            52 degrees C / 125 degrees F
Temperature I3 2 Chip               49 degrees C / 120 degrees F
Temperature I3 3 TSensor            47 degrees C / 116 degrees F
Temperature I3 3 Chip               47 degrees C / 116 degrees F
Temperature IA 0 TSensor            54 degrees C / 129 degrees F
Temperature IA 0 Chip               58 degrees C / 136 degrees F
Temperature IA 1 TSensor            48 degrees C / 118 degrees F
Temperature IA 1 Chip               53 degrees C / 127 degrees F
Power
1.5 V                               1479 mV
2.5 V                               2542 mV
3.3 V                               3319 mV
1.8 V PFE 0                         1811 mV
1.8 V PFE 1                         1804 mV
1.8 V PFE 2                         1804 mV
1.8 V PFE 3                         1814 mV
1.2 V PFE 0                         1192 mV
1.2 V PFE 1                         1202 mV
1.2 V PFE 2                         1205 mV
1.2 V PFE 3                         1189 mV
I2C Slave Revision                  40

```

show chassis environment fpc (MX960 Router)

```

user@host> show chassis environment fpc
FPC 5 status:
State                               Online
Temperature Intake                  27 degrees C / 80 degrees F
Temperature Exhaust A               34 degrees C / 93 degrees F
Temperature Exhaust B               40 degrees C / 104 degrees F
Temperature I3 0 TSensor            39 degrees C / 102 degrees F
Temperature I3 0 Chip               41 degrees C / 105 degrees F
Temperature I3 1 TSensor            38 degrees C / 100 degrees F
Temperature I3 1 Chip               37 degrees C / 98 degrees F
Temperature I3 2 TSensor            37 degrees C / 98 degrees F
Temperature I3 2 Chip               34 degrees C / 93 degrees F
Temperature I3 3 TSensor            32 degrees C / 89 degrees F
Temperature I3 3 Chip               33 degrees C / 91 degrees F
Temperature IA 0 TSensor            39 degrees C / 102 degrees F
Temperature IA 0 Chip               44 degrees C / 111 degrees F
Temperature IA 1 TSensor            36 degrees C / 96 degrees F
Temperature IA 1 Chip               44 degrees C / 111 degrees F
Power
1.5 V                               1479 mV
2.5 V                               2523 mV
3.3 V                               3254 mV

```

```

1.8 V PFE 0          1798 mV
1.8 V PFE 1          1798 mV
1.8 V PFE 2          1807 mV
1.8 V PFE 3          1791 mV
1.2 V PFE 0          1173 mV
1.2 V PFE 1          1179 mV
1.2 V PFE 2          1179 mV
1.2 V PFE 3          1185 mV
I2C Slave Revision   6
FPC 6 status:
State                Online
Temperature Intake    25 degrees C / 77 degrees F
Temperature Exhaust A 38 degrees C / 100 degrees F
Temperature Exhaust B 38 degrees C / 100 degrees F
Temperature I3 0 TSensor 40 degrees C / 104 degrees F
Temperature I3 0 Chip  40 degrees C / 104 degrees F
Temperature I3 1 TSensor 40 degrees C / 104 degrees F
Temperature I3 1 Chip  38 degrees C / 100 degrees F
Temperature I3 2 TSensor 37 degrees C / 98 degrees F
Temperature I3 2 Chip  32 degrees C / 89 degrees F
Temperature I3 3 TSensor 34 degrees C / 93 degrees F
Temperature I3 3 Chip  33 degrees C / 91 degrees F
Temperature IA 0 TSensor 45 degrees C / 113 degrees F
Temperature IA 0 Chip  47 degrees C / 116 degrees F
Temperature IA 1 TSensor 37 degrees C / 98 degrees F
Temperature IA 1 Chip  42 degrees C / 107 degrees F
Power
1.5 V                1485 mV
2.5 V                2510 mV
3.3 V                3332 mV
1.8 V PFE 0          1801 mV
1.8 V PFE 1          1814 mV
1.8 V PFE 2          1804 mV
1.8 V PFE 3          1820 mV
1.2 V PFE 0          1192 mV
1.2 V PFE 1          1189 mV
1.2 V PFE 2          1202 mV
1.2 V PFE 3          1156 mV
I2C Slave Revision   40

```

show chassis environment fpc (MX480 Router with 100-Gigabit Ethernet CFP)

```

user@host> show chassis environment fpc
FPC 0 status:
State                Online
Temperature Intake    32 degrees C / 89 degrees F
Temperature Exhaust A 39 degrees C / 102 degrees F
Temperature Exhaust B 37 degrees C / 98 degrees F
Temperature QX 0 TSen 44 degrees C / 111 degrees F
Temperature QX 0 Chip 48 degrees C / 118 degrees F
Temperature LU 0 TCAM TSen 44 degrees C / 111 degrees F
Temperature LU 0 TCAM Chip 47 degrees C / 116 degrees F
Temperature LU 0 TSen 44 degrees C / 111 degrees F
Temperature LU 0 Chip 48 degrees C / 118 degrees F
Temperature MQ 0 TSen 44 degrees C / 111 degrees F
Temperature MQ 0 Chip 47 degrees C / 116 degrees F
Power
MPC-BIAS3V3-z12105    3297 mV
MPC-VDD3V3-z12105     3306 mV
MPC-VDD2V5-z12105     2498 mV
MPC-TCAM_1V0-z12004   999 mV

```

```

MPC-AVDD1V0-z12006      999 mV
MPC-VDD1V8-z12006      1796 mV
MPC-PCIE_1V0-z12006     1002 mV
MPC-LU0_1V0-z12004      997 mV
MPC-MQ0_1V0-z12004      995 mV
MPC-VDD_1V5-z12004     1496 mV
MPC-PMB_1V1-z12006     1094 mV
MPC-9VA-BMR453          9054 mV
MPC-9VB-BMR453          9037 mV
MPC-PMB_1V2-z12106     1191 mV
MPC-QXM0_1V0-z12006    1000 mV
I2C Slave Revision      66
FPC 1 status:
State                    Online
Temperature Intake       35 degrees C / 95 degrees F
Temperature Exhaust A    50 degrees C / 122 degrees F
Temperature Exhaust B    56 degrees C / 132 degrees F
Temperature LU 0 TSen    46 degrees C / 114 degrees F
Temperature LU 0 Chip    59 degrees C / 138 degrees F
Temperature LU 1 TSen    46 degrees C / 114 degrees F
Temperature LU 1 Chip    45 degrees C / 113 degrees F
Temperature LU 2 TSen    46 degrees C / 114 degrees F
Temperature LU 2 Chip    60 degrees C / 140 degrees F
Temperature LU 3 TSen    46 degrees C / 114 degrees F
Temperature LU 3 Chip    71 degrees C / 159 degrees F
Temperature XM 0 TSen    46 degrees C / 114 degrees F
Temperature XM 0 Chip    -18 degrees C / 0 degrees F
Temperature XF 0 TSen    46 degrees C / 114 degrees F
Temperature XF 0 Chip    76 degrees C / 168 degrees F
Power
MPC-BIAS3V3-z12105      3292 mV
MPC-VDD3V3-z16100       3303 mV
MPC-VDD2V5-z16100       2501 mV
MPC-VDD1V8-z12004       1801 mV
MPC-AVDD1V0-z12006       996 mV
MPC-VDD1V2-z16100       1199 mV
MPC-VDD1V5A-z12004      1493 mV
MPC-VDD1V5B-z12004      1498 mV
MPC-XF_0V9-z12006        996 mV
MPC-PCIE_1V0-z16100     1000 mV
MPC-LU0_1V0-z12004       994 mV
MPC-LU1_1V0-z12004       994 mV
MPC-LU2_1V0-z12004       992 mV
MPC-LU3_1V0-z12004       993 mV
MPC-12VA-BMR453         12003 mV
MPC-12VB-BMR453         12043 mV
MPC-PMB_1V1-z12006      1091 mV
MPC-PMB_1V2-z12106      1196 mV
MPC-XM_0V9-vt273m       899 mV
I2C Slave Revision      106

```

show chassis environment fpc (MX240, MX480, MX960 with Application Services Modular Line Card)

```

user@host>show chassis environment fpc 1
FPC 1 status:
State                    Online
Temperature Intake       36 degrees C / 96 degrees F
Temperature Exhaust A    39 degrees C / 102 degrees F
Temperature LU TSen      52 degrees C / 125 degrees F
Temperature LU Chip      54 degrees C / 129 degrees F
Temperature XM TSen      52 degrees C / 125 degrees F

```

Temperature XM Chip	60 degrees C / 140 degrees F
Temperature PCIe TSen	52 degrees C / 125 degrees F
Temperature PCIe Chip	69 degrees C / 156 degrees F
Power	
MPC-BIAS3V3-z12106	3302 mV
MPC-VDD3V3-z16100	3325 mV
MPC-AVDD1V0-z16100	1007 mV
MPC-PCIE_1V0-z16100	904 mV
MPC-LU0_1V0-z12004	996 mV
MPC-VDD_1V5-z12004	1498 mV
MPC-12VA-BMR453	11733 mV
MPC-12VB-BMR453	11728 mV
MPC-XM_0V9-vt273m	900 mV
I2C Slave Revision	81

show chassis environment fpc (T320, T640, and T1600 Routers)

```

user@host> show chassis environment fpc
FPC 0 status:
State                               Online
Temperature Top                     42 degrees C / 107 degrees F
Temperature Bottom                   36 degrees C / 96 degrees F
Temperature MMB1                     39 degrees C / 102 degrees F
Power:
  1.8 V                             1959 mV
  2.5 V                             2495 mV
  3.3 V                             3344 mV
  5.0 V                             5047 mV
  1.8 V bias                         1787 mV
  3.3 V bias                         3291 mV
  5.0 V bias                         4998 mV
  8.0 V bias                         7343 mV
BUS Revision                         40
FPC 1 status:
State                               Online
Temperature Top                     42 degrees C / 107 degrees F
Temperature Bottom                   39 degrees C / 102 degrees F
Temperature MMB1                     40 degrees C / 104 degrees F
Power:
  1.8 V                             1956 mV
  2.5 V                             2498 mV
  3.3 V                             3340 mV
  5.0 V                             5023 mV
  1.8 V bias                         1782 mV
  3.3 V bias                         3277 mV
  5.0 V bias                         4989 mV
  8.0 V bias                         7289 mV
BUS Revision                         40
FPC 2 status:
State                               Online
Temperature Top                     43 degrees C / 109 degrees F
Temperature Bottom                   39 degrees C / 102 degrees F
Temperature MMB1                     41 degrees C / 105 degrees F
Power:
  1.8 V                             1963 mV
  2.5 V                             2503 mV
  3.3 V                             3340 mV
  5.0 V                             5042 mV
  1.8 V bias                         1797 mV
  3.3 V bias                         3311 mV
  5.0 V bias                         5013 mV

```

8.0 V bias	7221 mV
BUS Revision	40

show chassis environment fpc (T4000 Router)

```

user@host> show chassis environment fpc
FPC 0 status:
State                               Online
Fan Intake                         34 degrees C / 93 degrees F
Fan Exhaust                        48 degrees C / 118 degrees F
PMB                                47 degrees C / 116 degrees F
LMB0                               50 degrees C / 122 degrees F
LMB1                               41 degrees C / 105 degrees F
LMB2                               35 degrees C / 95 degrees F
PFE1 LU2                          46 degrees C / 114 degrees F
PFE1 LU0                          41 degrees C / 105 degrees F
PFE0 LU0                          57 degrees C / 134 degrees F
XF1                               47 degrees C / 116 degrees F
XF0                               52 degrees C / 125 degrees F
XM1                               41 degrees C / 105 degrees F
XM0                               50 degrees C / 122 degrees F
PFE0 LU1                          56 degrees C / 132 degrees F
PFE0 LU2                          45 degrees C / 113 degrees F
PFE1 LU1                          37 degrees C / 98 degrees F

Power 1
  1.0 V                            991 mV
  1.2 V bias                      1195 mV
  1.8 V                            1788 mV
  2.5 V                            2483 mV
  3.3 V                            3289 mV
  3.3 V bias                      3299 mV
  12.0 V A                        10608 mV
  12.0 V B                        10637 mV

Power 2
  0.9 V                            881 mV
  0.9 V PFE0                      916 mV
  0.9 V PFE1                      903 mV
  1.0 V PFE0                     1012 mV
  1.0 V PFE1                     1002 mV
  1.1 V                            1095 mV
  1.5 V_0                        1494 mV
  1.5 V_1                        1479 mV

Power 3
  1.0 V PFE0                     1000 mV
  1.0 V PFE1                     1002 mV
  1.0 V PFE0 *                   995 mV
  1.0 V PFE1 *                   995 mV
  1.8 V PFE 0                    1788 mV
  1.8 V PFE 1                    1789 mV
  2.5 V                            2482 mV
  12.0 V                         11614 mV

Power 4
  1.0 V PFE0 LU0                 1003 mV
  1.0 V PFE1 LU0                 1003 mV
  1.0 V PFE1 LU2                 1004 mV
  1.0 V PFE0 LU0 *               995 mV
  1.0 V PFE1 LU0 *               998 mV
  1.0 V PFE1 LU2 *               996 mV
  12.0 V                         11643 mV
  12.0 V C                       11711 mV

Power (Base/PMB/MMB)

```

LMB0 VDD2V5	2488 mV
LMB0 VDD1V8	1788 mV
LMB0 VDD1V5	1496 mV
LMB0 PFE0 LU0 AVDD1V0	1002 mV
LMB0 PFE0 LU0 VDD1V0	1000 mV
LMB0 VDD12V0	10752 mV
LMB1 VDD2V5	2472 mV
LMB1 VDD1V8	1792 mV
LMB1 VDD1V5	1480 mV
LMB1 PFE0 LU2 AVDD1V0	994 mV
LMB1 PFE0 LU2 VDD1V0	1002 mV
LMB1 VDD12V0	10800 mV
LMB2 VDD2V5	2472 mV
LMB2 VDD1V8	1792 mV
LMB2 VDD1V5	1486 mV
LMB2 PFE1 LU1 AVDD1V0	996 mV
LMB2 PFE1 LU1 VDD1V0	998 mV
LMB2 VDD12V0	10704 mV
PMB 1.05v	1049 mV
PMB 1.5v	1500 mV
PMB 2.5v	2500 mV
PMB 3.3v	3299 mV
Bus Revision	113
FPC 3 status:	
State	Online
Fan Intake	37 degrees C / 98 degrees F
Fan Exhaust	51 degrees C / 123 degrees F
PMB	43 degrees C / 109 degrees F
LMB0	57 degrees C / 134 degrees F
LMB1	54 degrees C / 129 degrees F
LMB2	38 degrees C / 100 degrees F
PFE1 LU2	63 degrees C / 145 degrees F
PFE1 LU0	45 degrees C / 113 degrees F
PFE0 LU0	69 degrees C / 156 degrees F
XF1	62 degrees C / 143 degrees F
XF0	63 degrees C / 145 degrees F
XM1	43 degrees C / 109 degrees F
XM0	67 degrees C / 152 degrees F
PFE0 LU1	63 degrees C / 145 degrees F
PFE0 LU2	66 degrees C / 150 degrees F
PFE1 LU1	41 degrees C / 105 degrees F
Power 1	
1.0 V	1002 mV
1.2 V bias	1201 mV
1.8 V	1785 mV
2.5 V	2485 mV
3.3 V	3288 mV
3.3 V bias	3285 mV
12.0 V A	10412 mV
12.0 V B	10515 mV
Power 2	
0.9 V	882 mV
0.9 V PFE0	920 mV
0.9 V PFE1	905 mV
1.0 V PFE0	1015 mV
1.0 V PFE1	1001 mV
1.1 V	1094 mV
1.5 V_0	1495 mV
1.5 V_1	1478 mV
Power 3	
0.92 V PFE1	998 mV

1.0 V PFE0	997 mV
1.0 V PFE0 *	992 mV
1.0 V PFE1 *	991 mV
1.8 V PFE 0	1780 mV
1.8 V PFE 1	1797 mV
2.5 V	2492 mV
12.0 V	11604 mV
Power 4	
1.0 V PFE0 LU0	1003 mV
1.0 V PFE1 LU0	1004 mV
1.0 V PFE1 LU2	1003 mV
1.0 V PFE0 LU0 *	1000 mV
1.0 V PFE1 LU0 *	1001 mV
1.0 V PFE1 LU2 *	1003 mV
12.0 V	11653 mV
12.0 V C	11672 mV
Power (Base/PMB/MMB)	
LMB0 VDD2V5	2512 mV
LMB0 VDD1V8	1790 mV
LMB0 VDD1V5	1500 mV
LMB0 PFE0 LU0 AVDD1V0	1004 mV
LMB0 PFE0 LU0 VDD1V0	1002 mV
LMB0 VDD12V0	10608 mV
LMB1 VDD2V5	2472 mV
LMB1 VDD1V8	1788 mV
LMB1 VDD1V5	1480 mV
LMB1 PFE0 LU2 AVDD1V0	1000 mV
LMB1 PFE0 LU2 VDD1V0	1004 mV
LMB1 VDD12V0	10672 mV
LMB2 VDD2V5	2488 mV
LMB2 VDD1V8	1798 mV
LMB2 VDD1V5	1494 mV
LMB2 PFE1 LU1 AVDD1V0	1000 mV
LMB2 PFE1 LU1 VDD1V0	1004 mV
LMB2 VDD12V0	10528 mV
PMB 1.05v	1050 mV
PMB 1.5v	1500 mV
PMB 2.5v	2499 mV
PMB 3.3v	3299 mV
Bus Revision	113
FPC 5 status:	
State	Online
Temperature Top	39 degrees C / 102 degrees F
Temperature Bottom	38 degrees C / 100 degrees F
Power	
1.8 V	1804 mV
1.8 V bias	1802 mV
3.3 V	3294 mV
3.3 V bias	3277 mV
5.0 V bias	5008 mV
5.0 V TOP	5067 mV
8.0 V bias	6642 mV
Power (Base/PMB/MMB)	
1.2 V	1202 mV
1.5 V	1504 mV
5.0 V BOT	5079 mV
12.0 V TOP Base	11848 mV
12.0 V BOT Base	11780 mV
1.1 V PMB	1111 mV
1.2 V PMB	1189 mV
1.5 V PMB	1494 mV

1.8 V PMB	1819 mV
2.5 V PMB	2503 mV
3.3 V PMB	3294 mV
5.0 V PMB	5035 mV
12.0 V PMB	11788 mV
0.75 MMB TOP	766 mV
1.5 V MMB TOP	1484 mV
1.8 V MMB TOP	1772 mV
2.5 V MMB TOP	2485 mV
1.2 V MMB TOP	1137 mV
5.0 V MMB TOP	4946 mV
12.0 V MMB TOP	11772 mV
3.3 V MMB TOP	3289 mV
0.75 MMB BOT	759 mV
1.5 V MMB BOT	1482 mV
1.8 V MMB BOT	1792 mV
2.5 V MMB BOT	2490 mV
1.2 V MMB BOT	1145 mV
5.0 V MMB BOT	4922 mV
12.0 V MMB BOT	11625 mV
3.3 V MMB BOT	3282 mV
APS 00	2495 mV
APS 01	3308 mV
APS 02	3301 mV
5.0 V PIC 0	4967 mV
APS 10	2512 mV
APS 11	3316 mV
APS 12	3304 mV
5.0 V PIC 1	5081 mV
Bus Revision	49
FPC 6 status:	
State	Online
Fan Intake	34 degrees C / 93 degrees F
Fan Exhaust	49 degrees C / 120 degrees F
PMB	40 degrees C / 104 degrees F
LMB0	60 degrees C / 140 degrees F
LMB1	58 degrees C / 136 degrees F
LMB2	40 degrees C / 104 degrees F
PFE1 LU2	69 degrees C / 156 degrees F
PFE1 LU0	45 degrees C / 113 degrees F
PFE0 LU0	71 degrees C / 159 degrees F
XF1	58 degrees C / 136 degrees F
XF0	65 degrees C / 149 degrees F
XM1	40 degrees C / 104 degrees F
XM0	66 degrees C / 150 degrees F
PFE0 LU1	69 degrees C / 156 degrees F
PFE0 LU2	68 degrees C / 154 degrees F
PFE1 LU1	42 degrees C / 107 degrees F
Power 1	
1.0 V	998 mV
1.2 V bias	1191 mV
1.8 V	1781 mV
2.5 V	2487 mV
3.3 V	3302 mV
3.3 V bias	3300 mV
12.0 V A	10388 mV
12.0 V B	10388 mV
Power 2	
0.9 V	902 mV
0.9 V PFE0	921 mV
0.9 V PFE1	907 mV

1.0 V PFE0	996 mV
1.0 V PFE1	974 mV
1.1 V	1095 mV
1.5 V_0	1495 mV
1.5 V_1	1478 mV
Power 3	
1.0 V PFE0	997 mV
1.0 V PFE1	998 mV
1.0 V PFE0 *	993 mV
1.0 V PFE1 *	991 mV
1.8 V PFE 0	1796 mV
1.8 V PFE 1	1789 mV
2.5 V	2465 mV
12.0 V	11609 mV
Power 4	
1.0 V PFE0 LU0	1003 mV
1.0 V PFE1 LU0	1006 mV
1.0 V PFE1 LU2	1002 mV
1.0 V PFE0 LU0 *	1000 mV
1.0 V PFE1 LU0 *	998 mV
1.0 V PFE1 LU2 *	998 mV
12.0 V	11638 mV
12.0 V C	11702 mV
Power (Base/PMB/MMB)	
LMB0 VDD2V5	2484 mV
LMB0 VDD1V8	1780 mV
LMB0 VDD1V5	1496 mV
LMB0 PFE0 LU0 AVDD1V0	998 mV
LMB0 PFE0 LU0 VDD1V0	1004 mV
LMB0 VDD12V0	10528 mV
LMB1 VDD2V5	2472 mV
LMB1 VDD1V8	1776 mV
LMB1 VDD1V5	1474 mV
LMB1 PFE0 LU2 AVDD1V0	994 mV
LMB1 PFE0 LU2 VDD1V0	1004 mV
LMB1 VDD12V0	10544 mV
LMB2 VDD2V5	2476 mV
LMB2 VDD1V8	1790 mV
LMB2 VDD1V5	1492 mV
LMB2 PFE1 LU1 AVDD1V0	996 mV
LMB2 PFE1 LU1 VDD1V0	1010 mV
LMB2 VDD12V0	10528 mV
PMB 1.05v	1050 mV
PMB 1.5v	1499 mV
PMB 2.5v	2500 mV
PMB 3.3v	3300 mV
Bus Revision	80

show chassis environment fpc lcc (TX Matrix Router)

```
user@host> show chassis environment fpc lcc 0
lcc0-re0:
```

```
-----
FPC 1 status:
```

State	Online
Temperature Top	30 degrees C / 86 degrees F
Temperature Bottom	25 degrees C / 77 degrees F
Temperature MMB0	Absent
Temperature MMB1	27 degrees C / 80 degrees F
Power:	
1.8 V	1813 mV

```

2.5 V                2504 mV
3.3 V                3338 mV
5.0 V                5037 mV
1.8 V bias           1797 mV
3.3 V bias           3301 mV
5.0 V bias           5013 mV
8.0 V bias           7345 mV
BUS Revision         40
FPC 2 status:
State                Online
Temperature Top       37 degrees C / 98 degrees F
Temperature Bottom    26 degrees C / 78 degrees F
Temperature MMB0      32 degrees C / 89 degrees F
Temperature MMB1      27 degrees C / 80 degrees F
Power:
1.8 V                1791 mV
2.5 V                2517 mV
3.3 V                3308 mV
5.0 V                5052 mV
1.8 V bias           1797 mV
3.3 V bias           3289 mV
5.0 V bias           4991 mV
8.0 V bias           7477 mV
BUS Revision         40

```

show chassis environment fpc lcc (TX Matrix Plus Router)

```

user@host> show chassis environment fpc lcc 0
lcc0-re0:

```

```

-----
FPC 1 status:
State                Online
Temperature Top       46 degrees C / 114 degrees F
Temperature Bottom    47 degrees C / 116 degrees F
Power
1.8 V                1788 mV
1.8 V bias           1787 mV
3.3 V                3321 mV
3.3 V bias           3306 mV
5.0 V bias           5018 mV
5.0 V TOP            5037 mV
8.0 V bias           7223 mV
Power (Base/PMB/MMB)
1.2 V                1205 mV
1.5 V                1503 mV
5.0 V BOT            5084 mV
12.0 V TOP Base      11775 mV
12.0 V BOT Base      11794 mV
1.1 V PMB            1108 mV
1.2 V PMB            1196 mV
1.5 V PMB            1499 mV
1.8 V PMB            1811 mV
2.5 V PMB            2515 mV
3.3 V PMB            3318 mV
5.0 V PMB            5030 mV
12.0 V PMB           11832 mV
0.75 MMB TOP         752 mV
1.5 V MMB TOP        1489 mV
1.8 V MMB TOP        1782 mV
2.5 V MMB TOP        2498 mV
1.2 V MMB TOP        1155 mV

```

5.0 V MMB TOP	4902 mV
12.0 V MMB TOP	11721 mV
3.3 V MMB TOP	3316 mV
0.75 MMB BOT	754 mV
1.5 V MMB BOT	1482 mV
1.8 V MMB BOT	1758 mV
2.5 V MMB BOT	2488 mV
1.2 V MMB BOT	1157 mV
5.0 V MMB BOT	4962 mV
12.0 V MMB BOT	11691 mV
3.3 V MMB BOT	3308 mV
APS 00	1484 mV
APS 01	2503 mV
APS 02	3313 mV
5.0 V PIC 0	5025 mV
APS 10	1501 mV
APS 11	2466 mV
APS 12	3311 mV
5.0 V PIC 1	5081 mV
Bus Revision	49

show chassis environment fpc (QFX Series and OCX Series)

```
user@switch> show chassis environment fpc 0
FPC 0 status:
State                Online
Temperature           42 degrees C / 107 degrees F
```

show chassis environment fpc interconnect-device (QFabric Systems)

```
user@switch> show chassis environment fpc interconnect-device interconnect1 0
FC 0 FPC 0 status:
State                Online
Left Intake Temperature 24 degrees C / 75 degrees F
Right Intake Temperature 24 degrees C / 75 degrees F
Left Exhaust Temperature 27 degrees C / 80 degrees F
Right Exhaust Temperature 27 degrees C / 80 degrees F
Power
  BIAS 3V3           3330 mV
  VDD 3V3            3300 mV
  VDD 2V5            2502 mV
  VDD 1V5            1496 mV
  VDD 1V2            1194 mV
  VDD 1V0            1000 mV
  SW0 VDD 1V0        1020 mV
  SW0 CVDD 1V025     1032 mV
  SW1 VDD 1V0        1022 mV
  SW1 CVDD 1V025     1030 mV
  VDD 12V0 DIV3_33   3414 mV
```

show chassis environment fpc 0 (PTX5000 Packet Transport Router)

```
user@host> show chassis environment fpc 0
FPC 0 status:
State                Online
PMB Temperature      35 degrees C / 95 degrees F
Intake Temperature    33 degrees C / 91 degrees F
Exhaust A Temperature 51 degrees C / 123 degrees F
Exhaust B Temperature 43 degrees C / 109 degrees F
TL0 Temperature       48 degrees C / 118 degrees F
TQ0 Temperature       53 degrees C / 127 degrees F
TL1 Temperature       56 degrees C / 132 degrees F
```

TQ1 Temperature	58 degrees C / 136 degrees F
TL2 Temperature	55 degrees C / 131 degrees F
TQ2 Temperature	57 degrees C / 134 degrees F
TL3 Temperature	59 degrees C / 138 degrees F
TQ3 Temperature	59 degrees C / 138 degrees F
Power	
PMB 1.05v	1049 mV
PMB 1.5v	1500 mV
PMB 2.5v	2500 mV
PMB 3.3v	3299 mV
PFE0 1.5v	1500 mV
PFE0 1.0v	999 mV
TQ0 0.9v	900 mV
TL0 0.9v	900 mV
PFE1 1.5v	1499 mV
PFE1 1.0v	999 mV
TQ1 0.9v	899 mV
TL1 0.9v	900 mV
PFE2 1.5v	1500 mV
PFE2 1.0v	1000 mV
TQ2 0.9v	900 mV
TL2 0.9v	900 mV
PFE3 1.5v	1499 mV
PFE3 1.0v	1000 mV
TQ3 0.9v	900 mV
TL3 0.9v	900 mV
Bias 3.3v	3327 mV
FPC 3.3v	3300 mV
FPC 2.5v	2500 mV
SAM 0.9v	900 mV
A 12.0v	2014 mV
B 12.0v	2030 mV

show chassis environment fpc 07 (PTX5000 Packet Transport Router with FPC2-PTX-PIA)

```

user@host> show chassis environment fpc 07
FPC 7 status:
State Online
PMB TEMPO Temperature 32 degrees C / 89 degrees F
PMB TEMP1 Temperature 28 degrees C / 82 degrees F
PMB CPU Temperature 46 degrees C / 114 degrees F
Intake Temperature 35 degrees C / 95 degrees F
Exhaust A Temperature 55 degrees C / 131 degrees F
Exhaust B Temperature 54 degrees C / 129 degrees F
TL5 Temperature 59 degrees C / 138 degrees F
TQ5 Temperature 57 degrees C / 134 degrees F
TL6 Temperature 57 degrees C / 134 degrees F
TQ6 Temperature 51 degrees C / 123 degrees F
TL1 Temperature 76 degrees C / 168 degrees F
TQ1 Temperature 58 degrees C / 136 degrees F
TL2 Temperature 75 degrees C / 167 degrees F
TQ2 Temperature 57 degrees C / 134 degrees F
TL4 Temperature 52 degrees C / 125 degrees F
TQ4 Temperature 66 degrees C / 150 degrees F
TL7 Temperature 52 degrees C / 125 degrees F
TQ7 Temperature 60 degrees C / 140 degrees F
TL0 Temperature 72 degrees C / 161 degrees F
TQ0 Temperature 73 degrees C / 163 degrees F
TL3 Temperature 64 degrees C / 147 degrees F
TQ3 Temperature 70 degrees C / 158 degrees F
Power

```

PMB	1.05v	1049 mV
PMB	3.3v	3299 mV
PMB	1.1v-a	1100 mV
PMB	1.5v	1499 mV
PMB	1.1v-b	1100 mV
Base	3.3v	3300 mV
FPC Base	2.5v	2499 mV
TL1	0.9v	897 mV
TQ1	0.9v	897 mV
PFE1	1.0v	999 mV
PFE1	1.5v	1499 mV
TL2	0.9v	897 mV
TQ2	0.9v	897 mV
PFE2	1.0v	999 mV
PFE2	1.5v	1499 mV
FPC Base	1.0v	1000 mV
FPC Base	1.2v	1199 mV
TL5	0.9v	898 mV
TQ5	0.9v	898 mV
PFE5	1.0v	1000 mV
PFE5	1.5v	1500 mV
TL6	0.9v	897 mV
TQ6	0.9v	897 mV
PFE6	1.0v	1000 mV
PFE6	1.5v	1499 mV
Mezz Base	2.5v	2500 mV
TL0	0.9v	896 mV
TQ0	0.9v	896 mV
PFE0	1.0v	999 mV
PFE0	1.5v	1499 mV

show chassis environment FPC 1 (MX Routers with Media Services Blade [MSB])

```
user@switch> show chassis environment fpc 1
```

```
FPC 1 status:
```

State	Online
Temperature Intake	36 degrees C / 96 degrees F
Temperature Exhaust A	39 degrees C / 102 degrees F
Temperature LU TSen	52 degrees C / 125 degrees F
Temperature LU Chip	54 degrees C / 129 degrees F
Temperature XM TSen	52 degrees C / 125 degrees F
Temperature XM Chip	60 degrees C / 140 degrees F
Temperature PCIe TSen	52 degrees C / 125 degrees F
Temperature PCIe Chip	69 degrees C / 156 degrees F
Power	
MPC-BIAS3V3-z12106	3302 mV
MPC-VDD3V3-z16100	3325 mV
MPC-AVDD1V0-z16100	1007 mV
MPC-PCIE_1V0-z16100	904 mV
MPC-LU0_1V0-z12004	996 mV
MPC-VDD_1V5-z12004	1498 mV
MPC-12VA-BMR453	11733 mV
MPC-12VB-BMR453	11728 mV
MPC-XM_0V9-vt273m	900 mV
I2C Slave Revision	81

show chassis environment pem

List of Syntax	Syntax on page 566 Syntax (ACX4000 Router) on page 566 Syntax (TX Matrix Routers) on page 566 Syntax (TX Matrix Plus Routers) on page 566 Syntax (MX Series Router) on page 566 Syntax (MX104 3D Universal Edge Routers) on page 566 Syntax (QFX Series) on page 566 Syntax (OCX Series) on page 566
Syntax	show chassis environment pem <slot>
Syntax (ACX4000 Router)	show chassis environment pem
Syntax (TX Matrix Routers)	show chassis environment pem <lcc number scc> <slot>
Syntax (TX Matrix Plus Routers)	show chassis environment pem <lcc number sfc number> <slot>
Syntax (MX Series Router)	show chassis environment pem <slot> <all-members> <local> <member member-id>
Syntax (MX104 3D Universal Edge Routers)	show chassis environment pem <slot>
Syntax (QFX Series)	show chassis environment pem <slot (interconnect-device name slot) (node-device name)>
Syntax (OCX Series)	show chassis environment pem <slot>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS 11.3 for the QFX Series. Command introduced in Junos OS 12.3R2 for EX Series. Command introduced in Junos OS Release 13.2 for MX104 3D Universal Edge Routers. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Display Power Entry Module (PEM) environmental status information.



NOTE: The new high-capacity (4100W) enhanced DC PEM on MX960 routers includes a new design that can condition the input voltage. This results in the output voltage differing from the input voltage. The earlier generation of DC PEMs coupled the input power directly to the output, thereby making it safe to assume that the output voltage was equal to the input voltage.

- Options** **none**—Display environmental information about both PEMs. For the TX Matrix router, display environmental information about the PEMs, the TX Matrix router, and its attached T640 routers. For the TX Matrix Plus router, display environmental information about the PEMs, the TX Matrix Plus router, and its attached T1600 routers.
- all-members**—(MX Series routers only) (Optional) Display environmental information about the PEMs in all the member routers of the Virtual Chassis configuration.
- interconnect-device *name***—(QFabric systems only) (Optional) Display chassis environmental information about the PEMs in the Interconnect device.
- lcc *number***—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display environmental information about the PEM in a specified T640 router (or line-card chassis) that is connected to a TX Matrix router. On a TX Matrix Plus router, display environmental information about the PEM in a specified T1600 router (or line-card chassis) that is connected to a TX Matrix Plus router. Replace ***number*** with a value from **0** through **3**.
- local**—(MX Series routers only) (Optional) Display environmental information about the PEM in the local Virtual Chassis member.
- member *member-id***—(MX Series routers only) (Optional) Display environmental information about the PEM in the specified member of the Virtual Chassis configuration. Replace ***member-id*** with a value of **0** or **1**.
- node-device *name***—(QFabric systems only) (Optional) Display chassis environmental information about the PEMs in the Node device.
- scc**—(TX Matrix routers only) (Optional) Display environmental information about the PEM in the TX Matrix router (or switch-card chassis).
- sfc**—(TX Matrix Plus routers only) (Optional) Display environmental information about the PEM in the TX Matrix Plus router (or switch-fabric chassis).
- slot**—(Optional) Display environmental information about an individual PEM. Replace ***slot*** with **0** or **1**.

Required Privilege Level view

Related Documentation • [show chassis hardware on page 743](#)

List of Sample Output

- [show chassis environment pem \(M40e Router\) on page 569](#)
- [show chassis environment pem \(M120 Router\) on page 569](#)
- [show chassis environment pem \(M160 Router\) on page 569](#)
- [show chassis environment pem \(M320 Router\) on page 570](#)
- [show chassis environment pem \(MX104 Router\) on page 570](#)
- [show chassis environment pem \(MX240 Router\) on page 570](#)
- [show chassis environment pem \(MX480 Router\) on page 570](#)
- [show chassis environment pem \(MX960 Router\) on page 571](#)
- [show chassis environment pem \(T320 Router\) on page 571](#)
- [show chassis environment pem \(T640 Router\) on page 571](#)
- [show chassis environment pem \(T4000 Router\) on page 571](#)
- [show chassis environment pem \(T640/T1600/T4000 Routers With Six-Input DC Power Supply\) on page 572](#)
- [show chassis environment pem lcc \(TX Matrix Routing Matrix\) on page 572](#)
- [show chassis environment pem scc \(TX Matrix Routing Matrix\) on page 572](#)
- [show chassis environment pem sfc \(TX Matrix Plus Routing Matrix\) on page 573](#)
- [show chassis environment pem lcc \(TX Matrix Plus Routing Matrix\) on page 573](#)
- [show chassis environment pem node-device \(QFabric System\) on page 574](#)
- [show chassis environment pem \(QFX Series and OCX Series\) on page 574](#)
- [show chassis environment pem interconnect-device \(QFabric System\) on page 574](#)

Output Fields Table 48 lists the output fields for the **show chassis environment pem** command. Output fields are listed in the approximate order in which they appear.

Table 48: show chassis environment pem Output Fields

Field Name	Field Description
PEM <i>slot</i> status	Number of the PEM slot.
State	Status of the PEM.
Temperature	Temperature of the air flowing past the PEM.
AC Input	Status of the AC input for the specified component
AC Output	Status of the AC output for the specified component.
DC input	Status of the DC input for the specified component.
DC output	Status of the DC output for the specified component.
Load	(Not available on M40e or M160 routers) Information about the load on supply, in percentage of rated current being used.
Voltage	(M120, M160, M320, T640, T1600, TX Matrix, and TX Matrix Plus routers only) Information about voltage supplied to the PEM. (MX104 routers only) Information about voltage supplied by the PEM to the system.

Table 48: show chassis environment pem Output Fields (*continued*)

Field Name	Field Description
Current	(T640, T1600, TX Matrix, and TX Matrix Plus routers only) Information about the PEM current.
Power	(T640, T1600, TX Matrix, and TX Matrix Plus routers only) Information about the PEM power.
SCG/CB/SIB	(T640, T1600, TX Matrix, and TX Matrix Plus routers only) SONET Clock Generator/Control Board/Switch Interface Board.
FAN	(T640, T1600, and T4000 routers with six-input DC power supply only) Information about the DC output to the fan.

Sample Output

show chassis environment pem (M40e Router)

```
user@host> show chassis environment pem
PEM 0 status:
  State                Online
  Temperature           OK
  AC input              OK
  DC output             OK
```

show chassis environment pem (M120 Router)

```
user@host> show chassis environment pem
PEM 0 status:
  State                Online
  Temperature           OK
  DC Input:            OK
  DC Output:           OK
  Load                Less than 20 percent
  Voltage:
    48.0 V input       52864 mV
    48.0 V fan supply  41655 mV
    3.3 V              3399 mV
PEM 1 status:
  State                Online
  Temperature           OK
  DC Input:            OK
  DC Output:           OK
  Load                Less than 20 percent
  Voltage:
    48.0 V input       54537 mV
    48.0 V fan supply  42910 mV
    3.3 V              3506 mV
```

show chassis environment pem (M160 Router)

```
user@host> show chassis environment pem
PEM 0 status:
  State                Online
  Temperature           OK
  DC input             OK
  DC output            OK
  Load                Less than 20 percent
```

```
Voltage:
  48.0 V input      54833 mV
  48.0 V fan supply 50549 mV
  8.0 V bias        8239 mV
  5.0 V bias        5006 mV
```

show chassis environment pem (M320 Router)

```
user@host> show chassis environment pem
PEM 2 status:
  State      Online
  Temperature OK
  DC input    OK
  Load       Less than 40 percent
    48.0 V input  51853 mV
    48.0 V fan supply 48877 mV
    8.0 V bias    8449 mV
    5.0 V bias    4998 mV
PEM 3 status:
  State      Online
  Temperature OK
  DC input    OK
  Load       Less than 40 percent
    48.0 V input  51717 mV
    48.0 V fan supply 49076 mV
    8.0 V bias    8442 mV
    5.0 V bias    4998 mV
```

show chassis environment pem (MX104 Router)

```
user@host> show chassis environment pem
PEM 0 status:
  State      Online
  Temperature OK
  DC Output:  OK
  Voltage:
    12.0 V output 12281 mV
    3.3 V output  3353 mV
PEM 1 status:
  State      Empty
```

show chassis environment pem (MX240 Router)

```
user@host> show chassis environment pem
PEM 0 status:
  State      Online
  Temperature OK
  DC Output:  OK
PEM 1 status:
  State      Online
  Temperature OK
  DC Output:  OK
```

show chassis environment pem (MX480 Router)

```
user@host> show chassis environment pem
PEM 0 status:
  State      Online
  Temperature OK
  DC Input:   OK
  DC Output:  OK
  Voltage:
```

```

PEM 1 status:
  State           Online
  Temperature      OK
  DC Input:       OK
  DC Output:      OK
  Voltage:

```

show chassis environment pem (MX960 Router)

```

user@host> show chassis environment pem
PEM 2 status:
  State           Present
PEM 3 status:
  State           Online
  Temperature      OK
  DC Output:      OK

```

show chassis environment pem (T320 Router)

```

user@host> show chassis environment pem
PEM 0 status:
  State           Online
  Temperature      OK
  DC input:       OK

```

show chassis environment pem (T640 Router)

```

user@host> show chassis environment pem
PEM 0 status:
  State           Online
  Temperature      22 degrees C / 71 degrees F
  AC input: OK
  DC output:
    Voltage      Current      Power      Load
    FPC 0        56875        606        34        4
    FPC 1        57016        525        29        3
    FPC 2         0         0         0         0
    FPC 3         0         0         0         0
    FPC 4         0         0         0         0
    FPC 5         0         0         0         0
    FPC 6        57158        1581       90        12
    FPC 7         0         0         0         0
  SCG/CB/SIB     56750        1125       63         5

```

show chassis environment pem (T4000 Router)

```

user@host> show chassis environment pem
PEM 0 status:
  State           Online
  Temperature      33 degrees C / 91 degrees F
  DC Input:       OK
    Voltage(V)    Current(A)    Power(W)    Load(%)
  INPUT 0        54.625       9.812       535        22
  INPUT 1        54.625      10.250      559        23
  INPUT 2        55.125       0.125        6         0
  INPUT 3        54.500      10.062      548        22
  INPUT 4        54.750       9.375       513        21
  INPUT 5        54.750      10.187      557        23
  DC Output      Voltage(V)    Current(A)    Power(W)    Load(%)
  FPC 0         55.750      10.125       564        37
  FPC 1         51.625       0.000        0         0
  FPC 2         52.000       0.000        0         0
  FPC 3         55.062      10.437      574        38

```

FPC 4	52.125	0.000	0	0
FPC 5	55.000	9.375	515	34
FPC 6	55.187	9.687	534	35
FPC 7	51.437	0.000	0	0
SCG/CB/SIB	55.375	15.750	872	35
FAN	54.562	14.750	804	42

show chassis environment pem (T640/T1600/T4000 Routers With Six-Input DC Power Supply)

```

user@host> show chassis environment pem
PEM 1 status:
  State          Online
  Temperature    36 degrees C / 96 degrees F
  DC Input:      OK

```

	Voltage(V)	Current(A)	Power(W)	Load(%)
INPUT 0	0.000	0.000	0	0
INPUT 1	54.875	3.812	209	27
INPUT 2	55.375	3.937	218	29
INPUT 3	54.625	3.750	204	27
INPUT 4	55.125	3.375	186	24
INPUT 5	55.125	3.375	186	24

```

DC Output

```

	Voltage(V)	Current(A)	Power(W)	Load(%)
FPC 0	52.312	0.000	0	0
FPC 1	52.687	0.000	0	0
FPC 2	52.812	0.000	0	0
FPC 3	55.812	7.062	394	52
FPC 4	52.625	0.000	0	0
FPC 5	52.625	0.000	0	0
FPC 6	52.750	0.000	0	0
FPC 7	52.750	0.000	0	0
SCG/CB/SIB	55.937	11.937	667	55
FAN	55.812	4.937	275	36

show chassis environment pem lcc (TX Matrix Routing Matrix)

```

user@host> show chassis environment pem 0 lcc 0
lcc0-re0:
-----
PEM 0 status:
  State          Present
  Temperature    27 degrees C / 80 degrees F
  DC input:      Check
  DC output:

```

	Voltage	Current	Power	Load
FPC 0	0	0	0	0
FPC 1	0	0	0	0
FPC 2	0	0	0	0
FPC 3	0	0	0	0
FPC 4	0	0	0	0
FPC 5	0	0	0	0
FPC 6	0	0	0	0
FPC 7	0	0	0	0
SCG/CB/SIB	0	0	0	0

show chassis environment pem scc (TX Matrix Routing Matrix)

```

user@host> show chassis environment pem scc
scc-re0:
-----
PEM 1 status:
  State          Online
  Temperature    24 degrees C / 75 degrees F
  DC input:      OK

```

DC output:	Voltage	Current	Power	Load
SIB 0	0	0	0	0
SIB 1	0	0	0	0
SIB 2	0	0	0	0
SIB 3	56550	0	0	0
SIB 4	55958	6912	386	51

show chassis environment pem sfc (TX Matrix Plus Routing Matrix)

```
user@host> show chassis environment pem sfc 0
sfc0-re0:
```

```
-----
PEM 0 status:
State                Online
Temperature          35 degrees C / 95 degrees F
DC Input:            OK
DC Output
  Voltage    Current    Power    Load
  Channel 0  53820    14140    761     59
  Channel 1  53550    12720    681     53
  Channel 2  53840    12930    696     54
  Channel 3  53690    14990    804     63
  Channel 4  53620    15070    808     63
  Channel 5  53900    14820    798     62
  Channel 6  54120    5020     271     21
```

show chassis environment pem lcc (TX Matrix Plus Routing Matrix)

```
user@host> show chassis environment lcc 0
```

```
lcc0-re1:
```

```
-----
PEM 0 status:
State                Online
Temperature          38 degrees C / 100 degrees F
DC Input:            OK
DC Output
  Voltage    Current    Power    Load
  FPC 0      0          0         0         0
  FPC 1      0          0         0         0
  FPC 2      0          0         0         0
  FPC 3      0          0         0         0
  FPC 4      56408    7575     427        56
  FPC 5      0          0         0         0
  FPC 6      56266    7956     447        59
  FPC 7      56283    6100     343        45
  SCG/CB/SIB 55916    8950     500        41

PEM 1 status:
State                Present
Temperature          35 degrees C / 95 degrees F
DC Input:            Check
DC Output
  Voltage    Current    Power    Load
  FPC 0      0          0         0         0
  FPC 1      0          0         0         0
  FPC 2      0          0         0         0
  FPC 3      0          0         0         0
  FPC 4      0          0         0         0
  FPC 5      0          0         0         0
  FPC 6      0          0         0         0
  FPC 7      0          0         0         0
  SCG/CB/SIB 0          0         0         0
```

show chassis environment pem node-device (QFabric System)

```
user@switch> show chassis environment pem node-device node1
FPC 0 PEM 0 status:
  State          Check
  Airflow        Front to Back
  Temperature     OK
  AC Input:      OK
  DC Output      Voltage(V) Current(A) Power(W) Load(%)
                  12          10       120     18
FPC 0 PEM 1 status:
  State          Online
  Airflow        Back to Front
  Temperature     OK
  AC Input:      OK
  DC Output      Voltage(V) Current(A) Power(W) Load(%)
                  11          10       110     17
```

show chassis environment pem (QFX Series and OCX Series)

```
user@switch> show chassis environment pem
FPC 0 PEM 1 status:
  State          Online
  Airflow        Front to Back
  Temperature     OK
  AC Input:      OK
  DC Output      Voltage(V) Current(A) Power(W) Load(%)
                  12          17       204     31
```

show chassis environment pem interconnect-device (QFabric System)

```
user@switch> show chassis environment pem interconnect-device IC11
IC1 PEM 1 status:
  State          Online
  Airflow        Front to Back
  Temperature     OK
  AC Input:      OK
  DC Output      Voltage(V) Current(A) Power(W) Load(%)
                  12          18       216     33
```

show chassis environment routing-engine

List of Syntax	Syntax on page 575 Syntax (TX Matrix Routers) on page 575 Syntax (TX Matrix Plus Routers) on page 575 Syntax (MX104, MX2010, and MX2020 3D Universal Edge Routers) on page 575 Syntax (MX Series Routers) on page 575 Syntax (QFX Series) on page 575 Syntax (OCX Series) on page 575
Syntax	show chassis environment routing-engine <slot>
Syntax (TX Matrix Routers)	show chassis environment routing-engine <lcc number scc> <slot>
Syntax (TX Matrix Plus Routers)	show chassis environment routing-engine <lcc number sfc number> <slot>
Syntax (MX104, MX2010, and MX2020 3D Universal Edge Routers)	show chassis environment routing-engine <slot> <satellite [slot-id slot-id device-alias alias-name]>
Syntax (MX Series Routers)	show chassis environment routing-engine <slot> <all-members> <local> <member member-id>
Syntax (QFX Series)	show chassis environment routing-engine interconnect-device name
Syntax (OCX Series)	show chassis environment routing-engine
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 12.1 for the PTX Series Packet Transport Routers.</p> <p>Command introduced in Junos OS Release 12.1 for the T4000 Core Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 13.2 for MX104 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> <p>satellite option introduced in Junos OS Release 14.2R3.</p>
Description	Display Routing Engine environmental status information.

Options **none**—Display environmental information about all Routing Engines. For a TX Matrix router, display environmental information about all Routing Engines on the TX Matrix router and its attached T640 routers. For a TX Matrix Plus router, display environmental information about all Routing Engines on the TX Matrix Plus router and its attached routers.

all-members—(MX Series routers only) (Optional) Display environmental information about the Routing Engines in all member routers in the Virtual Chassis configuration.

interconnect-device *name*—(QFabric systems only) (Optional) Display environmental information about the Routing Engines for the Interconnect device.

lcc *number*—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(MX Series routers only) (Optional) Display environmental information about the Routing Engines in the local Virtual Chassis member.

member *member-id*—(MX Series routers only) (Optional) Display environmental information about the Routing Engines in the specified member in the Virtual Chassis configuration. Replace *member-id* with the value of 0 or 1.

satellite [*slot-id slot-id* | *device-alias alias-name*]—(Junos Fusion only) (Optional) Display environmental information for the specified satellite device in a Junos Fusion, or for all satellite devices in the Junos Fusion if no satellite devices are specified.

scc—(TX Matrix router only) (Optional) Display environmental information about the Routing Engine in the TX Matrix router (switch-card chassis).

sfc—(TX Matrix Plus router only) (Optional) Display environmental information about the Routing Engine in the TX Matrix Plus router (or switch-fabric chassis).

slot—(Optional) Display environmental information about an individual Routing Engine. On M10i, M20, M40e, M120, M160, M320, MX Series, MX104 routers, MX2010 routers, MX2020 routers, and T Series routers, replace *slot* with 0 or 1. On M5, M7i, M10, and M40 routers, replace *slot* with 0. On EX3200 and EX4200 standalone switches, replace *slot* with 0. On EX4200 switches in a Virtual Chassis configuration and on EX8208 and EX8216 switches, replace *slot* with 0 or 1. On the QFX3500 switch, there

is only one Routing Engine, so you do not need to specify the slot number. On PTX Series Packet Transport Routers, replace **slot** with **0** or **1**

Required Privilege Level view

Related Documentation

- [request chassis routing-engine master on page 376](#)
- [show chassis routing-engine on page 970](#)

List of Sample Output

- [show chassis environment routing-engine \(Nonredundant\) on page 577](#)
- [show chassis environment routing-engine \(Redundant\) on page 578](#)
- [show chassis environment routing-engine \(MX104 Router\) on page 578](#)
- [show chassis environment routing-engine \(MX2010 Router\) on page 578](#)
- [show chassis environment routing-engine \(MX2020 Router\) on page 578](#)
- [show chassis environment routing-engine \(TX Matrix Plus Router\) on page 578](#)
- [show chassis environment routing-engine \(T4000 Core Router\) on page 579](#)
- [show chassis environment routing-engine \(QFX Series and OCX Series\) on page 579](#)
- [show chassis environment routing-engine interconnect-device \(QFabric System\) on page 579](#)
- [show chassis environment routing-engine \(PTX5000 Packet Transport Router\) on page 579](#)

Output Fields Table 49 lists the output fields for the **show chassis environment routing-engine** command. Output fields are listed in the approximate order in which they appear.

Table 49: show chassis environment routing-engine Output Fields

Field Name	Field Description
Routing engine <i>slot</i> status	Number of the Routing Engine slot: 0 or 1.
State	Status of the Routing Engine: <ul style="list-style-type: none"> • Online Master—Routing Engine is online, operating as Master. • Online Standby—Routing Engine is online, operating as Standby. • Offline—Routing Engine is offline.
Temperature	Temperature of the air flowing past the Routing Engine.
CPU Temperature	(PTX Series and T4000 Core Routers only) Temperature of the air flowing past the Routing Engine CPU.

Sample Output

show chassis environment routing-engine (Nonredundant)

```
user@host> show chassis environment routing-engine
Routing Engine 0 status:
  State           Online Master
  Temperature     27 degrees C / 80 degrees
```

show chassis environment routing-engine (Redundant)

```
user@host> show chassis environment routing-engine
Route Engine 0 status:
  State:                Online Master
  Temperature:          26 degrees C / 78 degrees F
Route Engine 1 status:
  State:                Online Standby
  Temperature:          26 degrees C / 78 degrees F
```

show chassis environment routing-engine (MX104 Router)

```
user@ host >show chassis environment routing-engine
Routing Engine 0 status:
  State                Online Master
  Temperature          34 degrees C / 93 degrees F
  CPU Temperature      43 degrees C / 109 degrees F
Routing Engine 1 status:
  State                Online Standby
  Temperature          33 degrees C / 91 degrees F
  CPU Temperature      39 degrees C / 102 degrees F
```

show chassis environment routing-engine (MX2010 Router)

```
user@host> show chassis environment routing-engine
Routing Engine 0 status:
  State                Online Master
  Temperature          37 degrees C / 98 degrees F
  CPU Temperature      37 degrees C / 98 degrees F
Routing Engine 1 status:
  State                Online Standby
  Temperature          35 degrees C / 95 degrees F
  CPU Temperature      34 degrees C / 93 degrees F
```

show chassis environment routing-engine (MX2020 Router)

```
user@host> show chassis environment routing-engine
Routing Engine 0 status:
  State                Online Master
  Temperature          35 degrees C / 95 degrees F
  CPU Temperature      34 degrees C / 93 degrees F
Routing Engine 1 status:
  State                Online Standby
  Temperature          44 degrees C / 111 degrees F
  CPU Temperature      43 degrees C / 109 degrees F
```

show chassis environment routing-engine (TX Matrix Plus Router)

```
user@host> show chassis environment routing-engine
sfc0-re0:
-----
Routing Engine 0 status:
  State                Online Master
  Temperature          26 degrees C / 78 degrees F
Routing Engine 1 status:
  State                Online Standby
  Temperature          28 degrees C / 82 degrees F

lcc0-re0:
-----
Routing Engine 0 status:
```

```

State                Online Master
Temperature          30 degrees C / 86 degrees F
Routing Engine 1 status:
State                Online Standby
Temperature          29 degrees C / 84 degrees F

```

show chassis environment routing-engine (T4000 Core Router)

```

user@host> show chassis environment routing-engine
Routing Engine 0 status:
State                Online Master
Temperature          33 degrees C / 91 degrees F
CPU Temperature      50 degrees C / 122 degrees F
Routing Engine 1 status:
State                Online Standby
Temperature          33 degrees C / 91 degrees F
CPU Temperature      46 degrees C / 114 degrees F

```

show chassis environment routing-engine (QFX Series and OCX Series)

```

user@switch> show chassis environment routing-engine
Routing Engine 0 status:
State                Online Master
Temperature          42 degrees C / 107 degrees F

```

show chassis environment routing-engine interconnect-device (QFabric System)

```

user@switch> show chassis environment routing-engine interconnect-device interconnect1
routing-engine interconnect-device interconnect1
Routing Engine 0 status:
State                Online Standby
Temperature          52 degrees C / 125 degrees F
Routing Engine 1 status:
State                Online Master
Temperature          57 degrees C / 134 degrees F

```

show chassis environment routing-engine (PTX5000 Packet Transport Router)

```

user@switch> show chassis environment routing-engine
Routing Engine 0 status:
State                Online Master
Temperature          55 degrees C / 131 degrees F
CPU Temperature      66 degrees C / 150 degrees F
Routing Engine 1 status:
State                Online Standby
Temperature          52 degrees C / 125 degrees F
CPU Temperature      64 degrees C / 147 degrees F

```

show chassis fabric errors

List of Syntax	Syntax on page 580 Syntax (PTX Series Packet Transport Routers) on page 580
Syntax	<pre>show chassis fabric errors <autoheal> <fpc slot-number lcc number> <sib (slot f13 sib-slot f2s sib-slot/sib-f2s-slot-number lcc number)></pre>
Syntax (PTX Series Packet Transport Routers)	<pre>show chassis fabric errors (autoheal fpc slot-number sib sib-slot)</pre>
Release Information	<p>Command introduced in Junos OS Release 10.0.</p> <p>Command introduced in Junos OS Release 12.1X48 for the PTX Series Packet Transport Routers.</p>
Description	Display the first ten and last ten fabric errors for the FPC or Switch Interface Boards (SIBs).



NOTE: This command can only be issued on a master Routing Engine.

- Options**
- autoheal**—(TX Matrix Plus routers and PTX Series Packet Transport Routers only) Show an error log of the first 100 autoheal actions taken on the system.
- fpc slot-number**—Show error log of the first ten and last ten errors for the specified FPC. (PTX5000 Packet Transport Routers only)—Replace **slot-number** with a value from 0 through 7.
- (TX Matrix Plus routers only)—Replace **fpc slot-number** with the following values depending on the LCC configuration:
- On a TX Matrix Plus router with the TXP-T1600 configuration, if you specify the number of a T1600 LCC by using the **lcc number** option (the recommended method), replace **fpc slot-number** with a value from 0 through 7. Otherwise, use a value from 0 through 31.
 - On a TX Matrix Plus router with the TXP-T1600-3D, TXP-T4000-3D, or TXP-Mixed-LCC-3D configuration, if you specify the number of a T1600 or T4000 LCC by using the **lcc number** option (the recommended method), replace **fpc slot-number** with a value from 0 through 7. Otherwise, use a value from 0 through 63.
 - lcc number**—Show error log of the first ten and last ten errors for the specified FPC on a specific network device (line-card chassis) that is part of the routing matrix.

Replace **lcc number** with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

If you specify the number of the network device by using only the **lcc number** option (the recommended method), replace **slot-number** with a value from 0 through 7. Otherwise, replace **slot-number** with a value from 0 through 31. For example, the following commands have the same result:

```
user@host> show chassis fabric errors fpc 1 lcc 1
user@host> show chassis fabric errors fpc 9
```

sib—Show error log of the first ten and last ten errors for the specified SIB. This option has the following suboptions:

- (TX Matrix Plus routers only) **sib-slot**—Specify a value ranging from 0 through 4.
- (PTX Series Packet Transport Routers) **sib-slot**—Specify a value ranging from 0 through 8.
- **f13 sib-slot**—(Optional) Show SIB F13 errors. Specify a valid SIB value number: 0, 1, 3, 4, 6, 7, 8, 9, 11, or 12.
- **f2s sib-slot/sib-f2s-slot-number**—(Optional) Show SIB F2S errors. Replace **sib-slot** with a value from 0 through 4, followed by a **sib-f2s-slot-number** value 0, 2, 4 or 6.
- **lcc number**—(Optional) Show error log of the first ten and last ten SIB errors for the specified network device (line-card chassis).

Replace **number** with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.



NOTE: The `lcc number` suboption is mandatory when using the following format for the command: `show chassis fabric errors sib lcc number sib slot-number`. For instance, issuing `show chassis fabric errors sib lcc 2 3` displays errors detected on LCC 2, SIB 3.

This suboption is not required when the `f13` or `f2s` suboptions are used with the `sib slot-number` option.

Required Privilege Level view

List of Sample Output

- [show chassis fabric errors \(F13 SIB Errors on a TX Matrix Plus Router\) on page 583](#)
- [show chassis fabric errors \(F2S SIB Errors on a TX Matrix Plus Router\) on page 583](#)
- [show chassis fabric errors \(SIB Errors Specific to an LCC Connected to a TX Matrix Plus Router\) on page 583](#)
- [show chassis fabric errors \(FPC Errors Specific to an LCC Connected to a TX Matrix Plus Router\) on page 584](#)
- [show chassis fabric errors \(SIB Errors Specific to an LCC Connected to a TX Matrix Plus Router with 3D SIBs\) on page 584](#)
- [show chassis fabric errors fpc or sib \(PTX Series Packet Transport Routers\) on page 584](#)
- [show chassis fabric errors autoheal \(PTX Series Packet Transport Routers\) on page 584](#)
- [show chassis fabric errors autoheal \(TX Matrix Plus Router with 3D SIBs\) on page 584](#)

Output Fields [Table 50](#) lists the output fields for the `show chassis fabric errors` command. Output fields are listed in the approximate order in which they appear.

Table 50: show chassis fabric errors Output Fields

Field Name	Field Description
Time	Time the error was logged. (TX Matrix Plus routers and PTX Series Packet Transport Routers only) For the <code>autoheal</code> option, shows the timestamp when autoheal was attempted on a SIB that was in fault state.
Error log of first 10 errors	List of the first ten errors.
Error log of last 10 errors	List of the last ten errors.

Table 50: show chassis fabric errors Output Fields (*continued*)

Field Name	Field Description
Error log of first 100 errors	<p>Indicates the autoheal action taken on the SIB. The following actions can occur:</p> <ul style="list-style-type: none"> • Req—A SIB autoheal request was made on a faulty SIB. • Action—Autohealing (taking the SIB offline and then online) is initiated. • Denied—Autohealing (taking the SIB offline and then online) is denied because the SIB went to a fault state before the autoheal configuration period completed. • Set info—Setting information to force skipping autoheal on the SIB so that no further attempts to autoheal the faulty SIB are made. • Clear info—If a user takes a SIB offline and then online, then the autoheal information of the SIB is cleared. If the SIB goes to a fault state, autoheal is attempted on the SIB.
fpc slot number	(PTX5000 Packet Transport Router only)—Range is 0 through 7.
sib slot number	(PTX Series Packet Transport Routers only)—Range is 0 through 8.
lcc number	Not supported on PTX Series Packet Transport Routers.

Sample Output

show chassis fabric errors (F13 SIB Errors on a TX Matrix Plus Router)

```
user@host> show chassis fabric errors sib f13 11
```

```
Time                Error log of first 10 errors
2009-10-06 02:21:17 PDT    LOS on Cable-D(1,0)
```

show chassis fabric errors (F2S SIB Errors on a TX Matrix Plus Router)

```
user@host> show chassis fabric errors sib f2s 0/0
```

```
Time                Error log of first 10 errors
2009-10-06 13:51:42 PDT    Cell drop errors on CLOS F2 SF 0 Port 0 link
```

show chassis fabric errors (SIB Errors Specific to an LCC Connected to a TX Matrix Plus Router)

```
user@host> show chassis fabric errors sib 1 lcc 0
lcc0-re0:
```

```
-----
Time                Error log of first 10 errors
2009-10-06 02:23:16 PDT    Cell drop errors on FPC7_T link
2009-10-06 02:23:16 PDT    Cell drop errors on FPC7_B link
```

show chassis fabric errors (FPC Errors Specific to an LCC Connected to a TX Matrix Plus Router)

```
user@host> show chassis fabric errors fpc 5 lcc 0
lcc0-re0:
```

```
-----
Time                               Error log of first 10 errors
2009-10-06 13:56:59 PDT            PFE_T has link error on plane 1
```

show chassis fabric errors (SIB Errors Specific to an LCC Connected to a TX Matrix Plus Router with 3D SIBs)

```
user@host> show chassis fabric errors sib 1 lcc 0
lcc0-re0:
```

```
-----
Time                               Error log of first 10 errors
2013-02-11 04:46:42 PST            CRC errors on XC link SIB01_XF3#11,0
```

show chassis fabric errors fpc or sib (PTX Series Packet Transport Routers)

```
user@host> show chassis fabric errors fpc 1
```

```
Time                               Error log of first 10 errors
2012-01-06 16:27:03 PST            Link errs on PFE 2, SIB 0, Plane 0
```

```
user@host> show chassis fabric errors sib 1
```

```
Time                               Error log of first 10 errors
2015-01-16 15:34:33 PST            Link errs on PFE 0, FPC 0, Plane 2
2015-01-16 15:44:33 PST            CM set ASIC 1 to FAULT (Fault due to PIO errors)
```

show chassis fabric errors autoheal (PTX Series Packet Transport Routers)

```
user@host> show chassis fabric errors autoheal
```

show chassis fabric errors autoheal (TX Matrix Plus Router with 3D SIBs)

```
user@host> show chassis fabric errors autoheal
```

```
Time                               Error log of first 100 errors
2013-03-25 00:16:10 PDT            Req: Plane 3 F13 8 Cbl 4 (tx) LCC0-SIB3 Cbl 4 (rx)
2013-03-25 00:16:12 PDT            Action: Plane 3 F13 8 Cbl 4 (autohealing)
2013-03-25 00:17:24 PDT            Req: Plane 3 F13 8 Cbl 4 (tx) LCC0-SIB3 Cbl 4 (rx)
2013-03-25 00:17:24 PDT            Denied: Plane 3 F13 8 Cbl 4 (time < configured)
2013-03-25 00:17:24 PDT            Set info: Plane 3 F13 8 Cbl 4 (skip autoheal)
2013-03-25 01:20:17 PDT            Clear info: Plane 3
```


show chassis fabric fpcs

List of Syntax	Syntax on page 585 Syntax (MX Series Routers) on page 585 Syntax (MX2010 and MX2020 3D Universal Edge Routers) on page 585 Syntax (T4000 Core Router) on page 585 Syntax (PTX Series Packet Transport Routers) on page 585 Syntax (TX Matrix Plus Router) on page 585 Syntax (QFX Series Switches) on page 585
Syntax	show chassis fabric fpcs <lcc <i>number</i> >
Syntax (MX Series Routers)	show chassis fabric fpcs <all-members> <local> <member <i>member-id</i> >
Syntax (MX2010 and MX2020 3D Universal Edge Routers)	show chassis fabric fpcs
Syntax (T4000 Core Router)	show chassis fabric fpcs
Syntax (PTX Series Packet Transport Routers)	show chassis fabric fpcs <slot <i>fpc-slot</i> >
Syntax (TX Matrix Plus Router)	show chassis fabric fpcs <lcc <i>number</i> >
Syntax (QFX Series Switches)	show chassis fabric fpcs <slot <i>fpc-slot</i> >
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.4 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1x48 for PTX Series Packet Transport Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 15.1X53-D30 for QFX Series switches.</p>
Description	Display the state of the electrical switch fabric links between the Flexible PIC Concentrators (FPCs) and the Switch Interface Boards (SIBs).
Options	<p>none—Display the switch fabric link state. On a TX Matrix router, display the switching fabric link states for the FPCs in all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, display the switching fabric link states for the FPCs in all routers connected to the TX Matrix Plus router.</p>

all-members—(MX Series routers only) (Optional) Display the switching fabric link states for the FPCs in all members of the Virtual Chassis configuration.

lcc *number*—(TX Matrix router and TX Matrix Plus router only) (Optional) On a TX Matrix router, display the switch fabric link state for the FPCs in the specified T640 router (line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, display the switch fabric link state for the FPCs in the specified router (line-card chassis) that is connected to the TX Matrix Plus router. Replace *number* with a following value depending on the LCC configurations:

- From **0** through **3** on a T640 router on the routing matrix with TX Matrix routers.
- From **0** through **3** on a T1600 router on the routing matrix with TX Matrix Plus routers.
- From **0** through **7** on a T1600 router in a routing matrix with TX Matrix Plus router with 3D SIBs.
- **0, 2, 4, 6** on a T4000 router in a routing matrix with TX Matrix Plus router with 3D SIBs.

local—(MX Series routers only) (Optional) Display the switching fabric link states for the FPCs in the local Virtual Chassis member.

member *member-id*—(MX Series routers only) (Optional) Display the switching fabric link states for the FPCs in the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

slot *fpc-slot*—(PTX Series Packet Transport Routers and QFX Series switches only) (Optional) Display the fabric state of the specified FPC slot. If no value is provided, display the status of all FPCs.

Required Privilege Level

view

Related Documentation

- *request chassis fabric fpc*
- [show chassis fpc on page 703](#)
- *Displaying Information About DPCs or FPCs in an MX Series Router*

List of Sample Output

[show chassis fabric fpcs \(M320 Router\) on page 588](#)
[show chassis fabric fpcs \(MX240 Router\) on page 589](#)
[show chassis fabric fpcs \(MX480 Router\) on page 589](#)
[show chassis fabric fpcs \(MX960 Router\) on page 590](#)
[show chassis fabric fpcs \(MX240 with AS MLC Modular Carrier Card\) on page 592](#)
[show chassis fabric fpcs \(MX480 with AS MLC Modular Carrier Card\) on page 592](#)
[show chassis fabric fpcs \(MX480 Router with MPC4E\) on page 593](#)
[show chassis fabric fpcs \(MX960 with AS MLC Modular Carrier Card on page 594](#)
[show chassis fabric fpcs \(MX2010 Router\) on page 596](#)
[show chassis fabric fpcs \(MX2020 Router\) on page 599](#)
[show chassis fabric fpcs \(MX2020 Router with MPC4E\) on page 602](#)

[show chassis fabric fpcs \(T320 Router\) on page 604](#)
[show chassis fabric fpcs \(T640 Router\) on page 604](#)
[show chassis fabric fpcs \(TX Matrix Router\) on page 604](#)
[show chassis fabric fpcs \(TX Matrix Router with 3D SIBs\) on page 606](#)
[show chassis fabric fpcs lcc \(TX Matrix Router with 3D SIBs\) on page 609](#)
[show chassis fabric fpcs \(T1600 Router\) on page 609](#)
[show chassis fabric fpcs \(T4000 Core Router\) on page 611](#)
[show chassis fabric fpcs \(TX Matrix Plus Router\) on page 612](#)
[show chassis fabric fpcs lcc \(TX Matrix Plus Router\) on page 620](#)
[show chassis fabric fpcs \(EX8200 Switch\) on page 620](#)
[show chassis fabric fpcs \(PTX3000 Router\) on page 621](#)
[show chassis fabric fpcs \(QFX10008 Switch\) on page 622](#)

Output Fields [Table 51](#) lists the output fields for the **show chassis fabric fpcs** command. Output fields are listed in the approximate order in which they appear.

Table 51: show chassis fabric fpcs Output Fields

Field Name	Field Description
Fabric management FPC state	<p>Switching fabric link (link from SIB to FPC) state for each FPC:</p> <ul style="list-style-type: none"> • Unused—FPC is not present. (On MX240 and MX480 routers with AS- MLC modular carrier card or MPC4E only) the fabric plane from the pair that share physical links (1 and 5, and 3 and 7) is inactive. • Destination error on PFEs <i>list of PFE numbers</i>—Destination errors to the listed Packet Forwarding Engines. Indicates that the link is not carrying traffic to the listed Packet Forwarding Engines. NOTE: In Junos OS Release 9.6 and later, the list of Packet Forwarding Engines with destination errors is displayed in the output. In Junos OS Releases before 9.6, the output only indicates that there are destination errors. However, the list of Packet Forwarding Engines with destination errors is not displayed. • Links ok—Link between the spare SIB and FPC is eligible to carry traffic. • Link error—Link between the SIB and FPC has CRC errors. However, the link is still eligible to carry traffic. • Plane disabled—Fabric plane has been disabled for the following reasons: <ul style="list-style-type: none"> • Destination errors have exceeded the thresholds. • Run-time link errors have exceeded the thresholds. • Initialization time link errors detected, and link training was unsuccessful. • Plane Disabled, Links Error (PTX Series Packet Transport Routers and QFX Series switches only)—The plane is disabled because of link errors detected at the FPC RX. • Plane Disabled, Links Down (PTX Series Packet Transport Routers and QFX Series switches only)—The plane is disabled because of link errors detected at the SIB RX. • Plane enabled—Link between the active SIB and FPC is eligible to carry traffic. NOTE: On the Enhanced MX SCB with MPC, a maximum of 4 planes are operational and running. On all the other SCBs with MPC, all the planes are operational and running. • Plane Enabled, Links OK (PTX Series Packet Transport Routers and QFX Series switches only)—The FPC CCL RX link is eligible to carry traffic. • Plane Enabled, Links OK (TX Matrix and TX Matrix Plus routers only)—The FPC HSL RX link is eligible to carry traffic.

Sample Output

show chassis fabric fpcs (M320 Router)

```
user@host> show chassis fabric fpcs
```

```

Fabric management FPC state:
FPC #2
  PFE #1
    SIB #0      Plane enabled
    SIB #1      Plane enabled
    SIB #2      Plane enabled
    SIB #3      Plane enabled

```

show chassis fabric fpcs (MX240 Router)

```

user@host> show chassis fabric fpcs
Fabric management FPC state:
FPC 2
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Links ok
    Plane 5: Links ok
    Plane 6: Links ok
    Plane 7: Links ok
  PFE #1
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Links ok
    Plane 5: Links ok
    Plane 6: Links ok
    Plane 7: Links ok
  PFE #2
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Links ok
    Plane 5: Links ok
    Plane 6: Links ok
    Plane 7: Links ok
  PFE #3
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Links ok
    Plane 5: Links ok
    Plane 6: Links ok
    Plane 7: Links ok

```

show chassis fabric fpcs (MX480 Router)

```

user@host> show chassis fabric fpcs

FPC 0
  PFE #0
    Plane 0: Plane enabled

```

```
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
Plane 6: Links ok
Plane 7: Links ok
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
Plane 6: Links ok
Plane 7: Links ok
PFE #2
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
Plane 6: Links ok
Plane 7: Links ok
PFE #3
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
Plane 6: Links ok
Plane 7: Links ok
FPC 1
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
```

show chassis fabric fpcs (MX960 Router)

```
user@host> show chassis fabric fpcs
FPC 0
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
```

```
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
PFE #2
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
PFE #3
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
FPC 1
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
FPC 2
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
PFE #2
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
```

Plane 4: Links ok
...

show chassis fabric fpcs (MX240 with AS MLC Modular Carrier Card)

In the following output, FPC 1 is the AS MLC modular carrier card (AS MCC).

```
user@host>show chassis fabric fpcs
FPC 1
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Plane enabled
    Plane 5: Unused
    Plane 6: Plane enabled
    Plane 7: Unused
FPC 2
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled
```

show chassis fabric fpcs (MX480 with AS MLC Modular Carrier Card)

In the following output, FPC 5 is the AS MLC modular carrier card (AS MCC).

```
user@host>show chassis fabric fpcs
FPC 2
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled
FPC 4
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Links ok
    Plane 5: Links ok
    Plane 6: Links ok
    Plane 7: Links ok
  PFE #2
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Links ok
    Plane 5: Links ok
```



```

        Plane 6: Links ok
        Plane 7: Links ok
FPC 5
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Plane enabled
    Plane 5: Unused
    Plane 6: Plane enabled
    Plane 7: Unused

```

show chassis fabric fpcs (MX480 Router with MPC4E)

In the following output, **FPC4** is the MPC4E (MPC4E-3D-32XGE-SFPP) card.

```

user@host > show chassis fabric fpcs
Fabric management FPC state:
FPC 0
  PFE #0
    Plane 0: Links ok
    Plane 1: Links ok
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Plane enabled
    Plane 5: Links ok
    Plane 6: Plane enabled
    Plane 7: Links ok
  PFE #1
    Plane 0: Links ok
    Plane 1: Links ok
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Plane enabled
    Plane 5: Links ok
    Plane 6: Plane enabled
    Plane 7: Links ok
FPC 1
  PFE #0
    Plane 0: Links ok
    Plane 1: Links ok
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Plane enabled
    Plane 5: Links ok
    Plane 6: Plane enabled
    Plane 7: Links ok
  PFE #1
    Plane 0: Links ok
    Plane 1: Links ok
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Plane enabled
    Plane 5: Links ok
    Plane 6: Plane enabled
    Plane 7: Links ok
  PFE #2
    Plane 0: Links ok
    Plane 1: Links ok
    Plane 2: Plane enabled

```

```
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Links ok
Plane 6: Plane enabled
Plane 7: Links ok
PFE #3
Plane 0: Links ok
Plane 1: Links ok
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Links ok
Plane 6: Plane enabled

FPC 3
PFE #0
Plane 0: Links ok
Plane 1: Links ok
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Links ok
Plane 6: Plane enabled
Plane 7: Links ok
FPC 4
PFE #0
Plane 0: Links ok
Plane 1: Links ok
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Unused
Plane 6: Plane enabled
Plane 7: Unused
PFE #1
Plane 0: Links ok
Plane 1: Links ok
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Unused
Plane 6: Plane enabled
Plane 7: Unused
```

show chassis fabric fpcs (MX960 with AS MLC Modular Carrier Card)

In the following output, FPC 5 is the AS MLC modular carrier card (AS MCC).

```
user@host>show chassis fabric fpcs
Fabric management FPC state:
FPC 0
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
```

```
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
FPC 1
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
FPC 4
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
PFE #2
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
PFE #3
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
FPC 5
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
FPC 8
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
```

```
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
PFE #2
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
PFE #3
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
```

show chassis fabric fpcs (MX2010 Router)

```
user@host> show chassis fabric fpcs
Fabric management FPC state:
FPC 0
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane disabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane disabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
FPC 1
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane disabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
FPC 2
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane disabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
```

```
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane disabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled

FPC 3
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane disabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled

PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane disabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled

PFE #2
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane disabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled

PFE #3
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane disabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled

FPC 4
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane disabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled

FPC 5
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
```

```
Plane 3: Plane disabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane disabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
FPC 6
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane disabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane disabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #2
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane disabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #3
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane disabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
FPC 7
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane disabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
```

```

PFE #1
  Plane 0: Plane enabled
  Plane 1: Plane enabled
  Plane 2: Plane enabled
  Plane 3: Plane disabled
  Plane 4: Plane enabled
Plane 5: Plane enabled
  Plane 6: Plane enabled
  Plane 7: Plane enabled
FPC 8
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane disabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled
FPC 9
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane disabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled
  PFE #1
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane disabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled

```

show chassis fabric fpcs (MX2020 Router)

```

user@host> show chassis fabric fpcs
Fabric management FPC state:
FPC 0
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled
  PFE #1
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled

```

```
Plane 7: Plane enabled
PFE #2
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #3
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
FPC 1
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #2
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #3
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
FPC 2
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
```



```
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #2
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #3
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
FPC 3
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #2
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #3
```

```
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
FPC 4
...
```

show chassis fabric fpcs (MX2020 Router with MPC4E)

```
user@host > show chassis fabric fpcs
Fabric management FPC state:
FPC 0
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled
  PFE #1
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled
FPC 9
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled
  PFE #1
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled
FPC 10
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
```

```
Plane 6: Plane enabled
Plane 7: Plane enabled
FPC 14
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
FPC 19
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #2
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #3
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
```

show chassis fabric fpcs (T320 Router)

```
user@host> show chassis fabric fpcs
FPC #3
  PFE #1
    SIB #0
      Links ok
    SIB #1
      Plane enabled
    SIB #2
      Plane enabled
FPC #5
  PFE #1
    SIB #0
      Links ok
    SIB #1
      Plane enabled
    SIB #2
      Plane enabled
FPC #7
  PFE #1
    SIB #0
      Links ok
    SIB #1
      Plane enabled
    SIB #2
      Plane enabled
```

show chassis fabric fpcs (T640 Router)

```
user@host> show chassis fabric fpcs
Fabric management FPC state:

FPC #2
  PFE #1
    SIB #0
      Links ok
    SIB #1
      Plane enabled
    SIB #2
      Plane enabled
    SIB #3
      Plane enabled
    SIB #4
      Plane enabled
FPC #3
  PFE #1
    SIB #2
      Plane enabled
    SIB #3
      Link error
      Destination error on PFES
      8   9   10   11   12   13   14   15   16   17   18   19   20   21
      0   1   2   3   4   5   6   7
    SIB #4
      Destination error on PFES
      8   9   10   11   12   13   14   15   16   17   18   19   20   21
      0   1   2   3   4   5   6   7
  ...
```

show chassis fabric fpcs (TX Matrix Router)

```
user@host> show chassis fabric fpcs
```

```
lcc0-re0:
```

```
-----
Fabric management FPC state:
```

```
FPC #0
```

```
  PFE #1
```

```
    SIB #0
```

```
      Links ok
```

```
    SIB #2
```

```
      Links ok
```

```
    SIB #3
```

```
      Links ok
```

```
    SIB #4
```

```
      Links ok
```

```
FPC #2
```

```
  PFE #1
```

```
    SIB #0
```

```
      Links ok
```

```
    SIB #2
```

```
      Links ok
```

```
    SIB #3
```

```
      Links ok
```

```
    SIB #4
```

```
      Links ok FPC #3
```

```
  PFE #1
```

```
    SIB #2
```

```
      Plane enabled
```

```
    SIB #3
```

```
      Link error
```

```
      Destination error on PFES
```

```
      8  9 10 11 12 13 14 15 16 17 18 19 20 21
```

```
    SIB #4
```

```
      Destination error on PFES
```

```
      8  9 10 11 12 13 14 15 16 17 18 19 20 21
```

```
...
```

```
FPC #4
```

```
  PFE #0
```

```
    SIB #4 Links ok
```

```
  PFE #1
```

```
    SIB #4 Links ok
```

```
FPC #5
```

```
  PFE #1
```

```
    SIB #4 Links ok
```

```
FPC #6
```

```
  PFE #1
```

```
    SIB #4 Links ok
```

```
lcc2-re0:
```

```
-----
Fabric management FPC state:
```

```
FPC #0
```

```
  PFE #1
```

```
    SIB #4 Links ok
```

```
FPC #1
```

```
  PFE #1
```

```
    SIB #4 Links ok
```

```
FPC #2
```

```
  PFE #0
```

```
    SIB #4 Links ok
```

```
  PFE #1
```

```
    SIB #4 Links ok
```

```
FPC #4
```

```
PFE #0
  SIB #4 Links ok
PFE #1
  SIB #4 Links ok
FPC #5
  PFE #1
    SIB #4 Links ok
```

show chassis fabric fpcs (TX Matrix Router with 3D SIBs)

```
user@host> show chassis fabric fpcs
lcc0-re0:
```

```
-----
Fabric management FPC state:
```

```
FPC #0
  PFE #0
    SIB #0
      Links ok
    SIB #1
      Links ok
    SIB #2
      Links ok
    SIB #3
      Links ok
    SIB #4
      Links ok
  PFE #1
    SIB #0
      Links ok
    SIB #1
      Links ok
    SIB #2
      Links ok
    SIB #3
      Links ok
    SIB #4
      Links ok
FPC #3
  PFE #0
    SIB #0
      Links ok
    SIB #1
      Links ok
    SIB #2
      Links ok
    SIB #3
      Links ok
    SIB #4
      Links ok
  PFE #1
    SIB #0
      Links ok
    SIB #1
      Links ok
    SIB #2
      Links ok
    SIB #3
      Links ok
    SIB #4
      Links ok
FPC #4
```

```
PFE #0
  SIB #0
    Links ok
  SIB #1
    Links ok
  SIB #2
    Links ok
  SIB #3
    Links ok
  SIB #4
    Links ok
PFE #1
  SIB #0
    Links ok
  SIB #1
    Links ok
  SIB #2
    Links ok
  SIB #3
    Links ok
  SIB #4
    Links ok
FPC #5
  PFE #0
    SIB #0
      Links ok
    SIB #1
      Links ok
    SIB #2
      Links ok
    SIB #3
      Links ok
    SIB #4
      Links ok
  PFE #1
    SIB #0
      Links ok
    SIB #1
      Links ok
    SIB #2
      Links ok
    SIB #3
      Links ok
    SIB #4
      Links ok
FPC #6
  PFE #0
    SIB #0
      Links ok
    SIB #1
      Links ok
    SIB #2
      Links ok
    SIB #3
      Links ok
    SIB #4
      Links ok
  PFE #1
    SIB #0
      Links ok
    SIB #1
```

```

        Links ok
    SIB #2
        Links ok
    SIB #3
        Links ok
    SIB #4
        Links ok

lcc2-re0:
-----

lcc4-re0:
-----
Fabric management FPC state:
FPC #2
  PFE #0
    SIB #0
        Links ok
    SIB #1
        Links ok
    SIB #2
        Links ok
    SIB #3
        Links ok
    SIB #4
        Links ok
  PFE #1
    SIB #0
        Links ok
    SIB #1
        Links ok
    SIB #2
        Links ok
    SIB #3
        Links ok
    SIB #4
        Links ok
FPC #3
  PFE #0
    SIB #0
        Links ok
    SIB #1
        Links ok
    SIB #2
        Links ok
    SIB #3
        Links ok
    SIB #4
        Links ok
  PFE #1
    SIB #0
        Links ok
    SIB #1
        Links ok
    SIB #2
        Links ok
    SIB #3
        Links ok
    SIB #4
        Links ok
```



```
lcc6-re0:
```

show chassis fabric fpcs lcc (TX Matrix Router with 3D SIBs)

```
user@host> show chassis fabric fpcs lcc 4
lcc4-re0:
```

```
Fabric management FPC state:
```

```
FPC #2
```

```
  PFE #0
```

```
    SIB #0
```

```
      Links ok
```

```
    SIB #1
```

```
      Links ok
```

```
    SIB #2
```

```
      Links ok
```

```
    SIB #3
```

```
      Links ok
```

```
    SIB #4
```

```
      Links ok
```

```
  PFE #1
```

```
    SIB #0
```

```
      Links ok
```

```
    SIB #1
```

```
      Links ok
```

```
    SIB #2
```

```
      Links ok
```

```
    SIB #3
```

```
      Links ok
```

```
    SIB #4
```

```
      Links ok
```

```
FPC #3
```

```
  PFE #0
```

```
    SIB #0
```

```
      Links ok
```

```
    SIB #1
```

```
      Links ok
```

```
    SIB #2
```

```
      Links ok
```

```
    SIB #3
```

```
      Links ok
```

```
    SIB #4
```

```
      Links ok
```

```
  PFE #1
```

```
    SIB #0
```

```
      Links ok
```

```
    SIB #1
```

```
      Links ok
```

```
    SIB #2
```

```
      Links ok
```

```
    SIB #3
```

```
      Links ok
```

```
    SIB #4
```

```
      Links ok
```

show chassis fabric fpcs (T1600 Router)

```
user@host> show chassis fabric fpcs
Fabric management FPC state:
FPC #0
```

```
PFE #0
  SIB #0
    Links ok
  SIB #1
    Plane enabled
  SIB #2
    Plane enabled
  SIB #3
    Plane enabled
  SIB #4
    Plane enabled
PFE #1
  SIB #0
    Links ok
  SIB #1
    Plane enabled
  SIB #2
    Plane enabled
  SIB #3
    Plane enabled
  SIB #4
    Plane enabled
FPC #1
  PFE #0
    SIB #0
      Links ok
    SIB #1
      Plane enabled
    SIB #2
      Plane enabled
    SIB #3
      Plane enabled
    SIB #4
      Plane enabled
  PFE #1
    SIB #0
      Links ok
    SIB #1
      Plane enabled
    SIB #2
      Plane enabled
    SIB #3
      Plane enabled
    SIB #4
      Plane enabled
FPC #2
  PFE #0
    SIB #0
      Links ok
    SIB #1
      Plane enabled
    SIB #2
      Plane enabled
    SIB #3
      Plane enabled
    SIB #4
      Plane enabled
FPC #4
  PFE #0
    SIB #0
      Links ok
```

```

SIB #1
    Plane enabled
SIB #2
    Plane enabled
SIB #3
    Plane enabled
SIB #4
    Plane enabled
PFE #1
    SIB #0
        Links ok
    SIB #1
        Plane enabled
    SIB #2
        Plane enabled
    SIB #3
        Plane enabled
    SIB #4
        Plane enabled
FPC #3
    PFE #1
        SIB #2
            Plane enabled
        SIB #3
            Link error
            Destination error on PFES
            8   9   10  11  12  13  14  15  16  17  18  19  20  21
            0   1   2   3   4   5   6   7
SIB #4
            Destination error on PFES
            8   9   10  11  12  13  14  15  16  17  18  19  20  21
            0   1   2   3   4   5   6   7

```

show chassis fabric fpcs (T4000 Core Router)

```

Fabric management FPC state:
FPC #2
    PFE #0
        SIB #0
            Links ok
        SIB #1
            Plane enabled
        SIB #2
            Plane enabled
        SIB #3
            Plane enabled
        SIB #4
            Plane enabled
FPC #3
    PFE #0
        SIB #0
            Links ok
        SIB #1
            Plane enabled
        SIB #2
            Plane enabled
        SIB #3
            Plane enabled
        SIB #4
            Plane enabled
FPC #5
    PFE #0
        SIB #0

```

```

        Links ok
    SIB #1
        Plane enabled
    SIB #2
        Plane enabled
    SIB #3
        Plane enabled
    SIB #4
        Plane enabled
PFE #1
    SIB #0
        Links ok
    SIB #1
        Plane enabled
    SIB #2
        Plane enabled
    SIB #3
        Plane enabled
    SIB #4
        Plane enabled
FPC #6
PFE #0
    SIB #0
        Links ok
    SIB #1
        Plane enabled
    SIB #2
        Plane enabled
    SIB #3
        Plane enabled
    SIB #4
        Plane enabled
PFE #1
    SIB #0
        Links ok
    SIB #1
        Plane enabled
    SIB #2
        Plane enabled
    SIB #3
        Plane enabled
    SIB #4
        Plane enabled
```

show chassis fabric fpcs (TX Matrix Plus Router)

```
user@host> show chassis fabric fpcs
lcc0-re0:
```

```
-----
Fabric management FPC state:
```

```
FPC #0
PFE #1
    SIB #0
        Unused
    SIB #1
        Links ok
    SIB #2
        Links ok
    SIB #3
        Links ok
    SIB #4
```

```

Links ok
FPC #2
  PFE #0
    SIB #0
      Unused
    SIB #1
      Links ok
    SIB #2
      Links ok
    SIB #3
      Links ok
    SIB #4
      Links ok
  PFE #1
    SIB #0
      Unused
    SIB #1
      Links ok
    SIB #2
      Links ok
    SIB #3
      Links ok
    SIB #4
      Links ok
FPC #3
  PFE #1
    SIB #2
      Plane enabled
    SIB #3
      Link error
      Destination error on PFes
      8   9   10  11  12  13  14  15  16  17  18  19  20  21
    SIB #4
      Destination error on PFes
      8   9   10  11  12  13  14  15  16  17  18  19  20  21
FPC #4
  PFE #0
    SIB #0
      Unused
    SIB #1
      Links ok
    SIB #2
      Links ok
    SIB #3
      Links ok
    SIB #4
      Links ok
  PFE #1
    SIB #0
      Unused
    SIB #1
      Links ok
    SIB #2
      Links ok
    SIB #3
      Links ok
    SIB #4
      Links ok
FPC #6
  PFE #0
    SIB #0

```

```

        Unused
    SIB #1
        Links ok
    SIB #2
        Links ok
    SIB #3
        Links ok
    SIB #4
        Links ok
PFE #1
    SIB #0
        Unused
    SIB #1
        Links ok
    SIB #2
        Links ok
    SIB #3
        Links ok
    SIB #4
        Links ok
FPC #7
    PFE #0
        SIB #0
            Unused
        SIB #1
            Links ok
        SIB #2
            Links ok
        SIB #3
            Links ok
        SIB #4
            Links ok
```

lcc1-re0:

Fabric management FPC state:

```
FPC #2
    PFE #0
        SIB #0
            Links ok
        SIB #1
            Links ok
        SIB #2
            Links ok
        SIB #3
            Links ok
        SIB #4
            Links ok
    PFE #1
        SIB #0
            Links ok
        SIB #1
            Links ok
        SIB #2
            Links ok
        SIB #3
            Links ok
        SIB #4
            Links ok
FPC #4
    PFE #0
```

```

SIB #0
    Links ok
SIB #1
    Links ok
SIB #2
    Links ok
SIB #3
    Links ok
SIB #4
    Links ok
PFE #1
    SIB #0
        Links ok
    SIB #1
        Links ok
    SIB #2
        Links ok
    SIB #3
        Destination error on PFES      1      8      9      29      40      65      72      73
                                           93     104
SIB #4
    Links ok
FPC #6
    PFE #0
        SIB #0
            Links ok
        SIB #1
            Links ok
        SIB #2
            Links ok
        SIB #3
            Links ok
        SIB #4
            Links ok
    PFE #1
        SIB #0
            Links ok
        SIB #1
            Links ok
        SIB #2
            Links ok
        SIB #3
            Links ok
        SIB #4
            Links ok
FPC #7
    PFE #0
        SIB #0
            Links ok
        SIB #1
            Links ok
        SIB #2
            Links ok
        SIB #3
            Links ok
        SIB #4
            Links ok

```

```

lcc2-re0:
-----

```

Fabric management FPC state:

FPC #0

PFE #0

SIB #0

Links ok

SIB #1

Links ok

SIB #2

Links ok

SIB #3

Links ok

SIB #4

Links ok

PFE #1

SIB #0

Links ok

SIB #1

Links ok

SIB #2

Links ok

SIB #3

Links ok

SIB #4

Links ok

FPC #2

PFE #0

SIB #0

Links ok

SIB #1

Links ok

SIB #2

Links ok

SIB #3

Links ok

SIB #4

Links ok

PFE #1

SIB #0

Links ok

SIB #1

Links ok

SIB #2

Links ok

SIB #3

Links ok

SIB #4

Links ok

FPC #4

PFE #0

SIB #0

Links ok

SIB #1

Links ok

SIB #2

Links ok

SIB #3

Links ok

SIB #4

Links ok

FPC #5

PFE #0


```

SIB #0
Links ok
SIB #1
Links ok
SIB #2
Links ok
SIB #3
Links ok
SIB #4
Links ok
PFE #1
SIB #0
Links ok
SIB #1
Links ok
SIB #2
Links ok
SIB #3
Links ok
SIB #4
Links ok
FPC #6
PFE #0
SIB #0
Links ok
SIB #1
Links ok
SIB #2
Links ok
SIB #3
Links ok
SIB #4
Links ok
PFE #1
SIB #0
Links ok
SIB #1
Links ok
SIB #2
Links ok
SIB #3
Links ok
SIB #4
Links ok
FPC #7
PFE #0
SIB #0
Links ok
SIB #1
Links ok
SIB #2
Links ok
SIB #3
Links ok
SIB #4
Links ok

```

```

1cc3-re0:

```

```

-----
Fabric management FPC state:
FPC #0

```

```
PFE #0
  SIB #0
    Links ok
  SIB #1
    Links ok
  SIB #2
    Links ok
  SIB #3
    Links ok
  SIB #4
    Links ok
PFE #1
  SIB #0
    Links ok
  SIB #1
    Links ok
  SIB #2
    Links ok
  SIB #3
    Links ok
  SIB #4
    Links ok
FPC #2
  PFE #0
    SIB #0
      Links ok
    SIB #1
      Links ok
    SIB #2
      Links ok
    SIB #3
      Links ok
    SIB #4
      Links ok
  PFE #1
    SIB #0
      Links ok
    SIB #1
      Links ok
    SIB #2
      Links ok
    SIB #3
      Links ok
    SIB #4
      Links ok
FPC #4
  PFE #0
    SIB #0
      Links ok
    SIB #1
      Links ok
    SIB #2
      Links ok
    SIB #3
      Links ok
    SIB #4
      Links ok
  PFE #1
    SIB #0
      Links ok
    SIB #1
```

```

Links ok
SIB #2
Links ok
SIB #3
Links ok
SIB #4
Links ok
FPC #5
PFE #0
SIB #0
Links ok
SIB #1
Links ok
SIB #2
Links ok
SIB #3
Links ok
SIB #4
Links ok
PFE #1
SIB #0
Links ok
SIB #1
Links ok
SIB #2
Links ok
SIB #3
Links ok
SIB #4
Links ok
FPC #6
PFE #0
SIB #0
Links ok
SIB #1
Links ok
SIB #2
Links ok
SIB #3
Links ok
SIB #4
Links ok
PFE #1
SIB #0
Links ok
SIB #1
Links ok
SIB #2
Links ok
SIB #3
Links ok
SIB #4
Links ok
FPC #7
PFE #0
SIB #0
Links ok
SIB #1
Links ok
SIB #2
Links ok
```

```

SIB #3
    Links ok
SIB #4
    Links ok

```

show chassis fabric fpcs lcc (TX Matrix Plus Router)

```

user@host> show chassis fabric fpcs lcc 0
lcc0-re1:
-----
Fabric management FPC state:
FPC #3
  PFE #1
    SIB #2
      Plane enabled
    SIB #3
      Link error
      Destination error on PFes
      8   9   10  11  12  13  14  15  16  17  18  19  20  21
    SIB #4
      Destination error on PFes
      8   9   10  11  12  13  14  15  16  17  18  19  20  21
FPC #4
  PFE #0
    SIB #0 Links ok
    SIB #1 Links ok
    SIB #2 Links ok
    SIB #3 Links ok
    SIB #4 Links ok
  PFE #1
    SIB #0 Links ok
    SIB #1 Links ok
    SIB #2 Links ok
    SIB #3 Links ok
    SIB #4 Links ok
FPC #6
  PFE #0
    SIB #0 Links ok
    SIB #1 Links ok
    SIB #2 Links ok
    SIB #3 Links ok
    SIB #4 Links ok
  PFE #1
    SIB #0 Links ok
    SIB #1 Links ok
    SIB #2 Links ok
    SIB #3 Links ok
    SIB #4 Links ok
FPC #7
  PFE #0
    SIB #0 Links ok
    SIB #1 Links ok
    SIB #2 Links ok
    SIB #3 Links ok
    SIB #4 Links ok

```

show chassis fabric fpcs (EX8200 Switch)

```

user@host> show chassis fabric fpcs
Fabric management FPC state
FPC 6

```

```

PFE #0
  Plane 0: Plane enabled
  Plane 1: Plane enabled
  Plane 2: Plane enabled
  Plane 3: Plane enabled
  Plane 4: Links ok
  Plane 5: Links ok
  Plane 6: Links ok
  Plane 7: Links ok
  Plane 8: Plane enabled
  Plane 9: Plane enabled
  Plane 10: Plane enabled
  Plane 11: Plane enabled
PFE #1
  Plane 0: Plane enabled
  Plane 1: Plane enabled
  Plane 2: Plane enabled
  Plane 3: Plane enabled
  Plane 4: Links ok
  Plane 5: Links ok
  Plane 6: Links ok
  Plane 7: Links ok
  Plane 8: Plane enabled
  Plane 9: Plane enabled
  Plane 10: Plane enabled
  Plane 11: Plane enabled
FPC 7
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Links ok
    Plane 5: Links ok
    Plane 6: Links ok
    Plane 7: Links ok
    Plane 8: Plane enabled
    Plane 9: Plane enabled
    Plane 10: Plane enabled
    Plane 11: Plane enabled
  PFE #1
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Links ok
    Plane 5: Links ok
    Plane 6: Links ok
    Plane 7: Links ok
    Plane 8: Plane enabled
    Plane 9: Plane enabled
    Plane 10: Plane enabled
    Plane 11: Plane enabled

```

show chassis fabric fpcs (PTX3000 Router)

```

user@host> show chassis fabric fpcs slot 8
Fabric management FPC state:
FPC #8
  PFE #0
    SIB0_Fcore0 (plane 0)  Plane Enabled, Links OK

```

```

SIB0_Fcore1 (plane 1) Plane Enabled, Links OK
SIB1_Fcore0 (plane 2) Plane Enabled, Links OK
SIB1_Fcore1 (plane 3) Plane Enabled, Links OK
SIB2_Fcore0 (plane 4) Plane Enabled, Links OK
SIB2_Fcore1 (plane 5) Plane Enabled, Links OK
SIB3_Fcore0 (plane 6) Plane Enabled, Links OK
SIB3_Fcore1 (plane 7) Plane Enabled, Links OK
SIB4_Fcore0 (plane 8) Plane Enabled, Links OK
SIB4_Fcore1 (plane 9) Plane Enabled, Links OK
SIB5_Fcore0 (plane 10) Plane Enabled, Links OK
SIB5_Fcore1 (plane 11) Plane Enabled, Links OK
SIB6_Fcore0 (plane 12) Plane Enabled, Links OK
SIB6_Fcore1 (plane 13) Plane Enabled, Links OK
SIB7_Fcore0 (plane 14) Plane Enabled, Links OK
SIB7_Fcore1 (plane 15) Plane Enabled, Links OK
SIB8_Fcore0 (plane 16) Plane Enabled, Links OK
SIB8_Fcore1 (plane 17) Plane Enabled, Links OK
PFE #1
SIB0_Fcore0 (plane 0) Plane Enabled, Links OK
SIB0_Fcore1 (plane 1) Plane Enabled, Links OK
SIB1_Fcore0 (plane 2) Plane Enabled, Links OK
SIB1_Fcore1 (plane 3) Plane Enabled, Links OK
SIB2_Fcore0 (plane 4) Plane Enabled, Links OK
SIB2_Fcore1 (plane 5) Plane Enabled, Links OK
SIB3_Fcore0 (plane 6) Plane Enabled, Links OK
SIB3_Fcore1 (plane 7) Plane Enabled, Links OK
SIB4_Fcore0 (plane 8) Plane Enabled, Links OK
SIB4_Fcore1 (plane 9) Plane Enabled, Links OK
SIB5_Fcore0 (plane 10) Plane Enabled, Links OK
SIB5_Fcore1 (plane 11) Plane Enabled, Links OK
SIB6_Fcore0 (plane 12) Plane Enabled, Links OK
SIB6_Fcore1 (plane 13) Plane Enabled, Links OK
SIB7_Fcore0 (plane 14) Plane Enabled, Links OK
SIB7_Fcore1 (plane 15) Plane Enabled, Links OK
SIB8_Fcore0 (plane 16) Plane Enabled, Links OK
SIB8_Fcore1 (plane 17) Plane Enabled, Links OK

```

show chassis fabric fpcs (QFX10008 Switch)

```
user@host> show chassis fabric fpcs slot 0
```

```
Fabric management FPC state:
```

```
FPC #0
```

```
PFE #0
```

```

SIB0_PFO (plane 0) Plane Enabled, Links OK
SIB0_PF1 (plane 1) Plane Enabled, Links OK
SIB1_PFO (plane 2) Plane Enabled, Links OK
SIB1_PF1 (plane 3) Plane Enabled, Links OK
SIB2_PFO (plane 4) Plane Enabled, Links OK
SIB2_PF1 (plane 5) Plane Enabled, Links OK
SIB3_PFO (plane 6) Plane Enabled, Links OK
SIB3_PF1 (plane 7) Plane Enabled, Links OK
SIB4_PFO (plane 8) Plane Enabled, Links OK
SIB4_PF1 (plane 9) Plane Enabled, Links OK
SIB5_PFO (plane 10) Plane Enabled, Links OK
SIB5_PF1 (plane 11) Plane Enabled, Links OK

```

```
PFE #1
```

```

SIB0_PFO (plane 0) Plane Enabled, Links OK
SIB0_PF1 (plane 1) Plane Enabled, Links OK
SIB1_PFO (plane 2) Plane Enabled, Links OK
SIB1_PF1 (plane 3) Plane Enabled, Links OK
SIB2_PFO (plane 4) Plane Enabled, Links OK

```

```

SIB2_PF1 (plane 5) Plane Enabled, Links OK
SIB3_PF0 (plane 6) Plane Enabled, Links OK
SIB3_PF1 (plane 7) Plane Enabled, Links OK
SIB4_PF0 (plane 8) Plane Enabled, Links OK
SIB4_PF1 (plane 9) Plane Enabled, Links OK
SIB5_PF0 (plane 10) Plane Enabled, Links OK
SIB5_PF1 (plane 11) Plane Enabled, Links OK
PFE #2
SIB0_PF0 (plane 0) Plane Enabled, Links OK
SIB0_PF1 (plane 1) Plane Enabled, Links OK
SIB1_PF0 (plane 2) Plane Enabled, Links OK
SIB1_PF1 (plane 3) Plane Enabled, Links OK
SIB2_PF0 (plane 4) Plane Enabled, Links OK
SIB2_PF1 (plane 5) Plane Enabled, Links OK
SIB3_PF0 (plane 6) Plane Enabled, Links OK
SIB3_PF1 (plane 7) Plane Enabled, Links OK
SIB4_PF0 (plane 8) Plane Enabled, Links OK
SIB4_PF1 (plane 9) Plane Enabled, Links OK
SIB5_PF0 (plane 10) Plane Enabled, Links OK
SIB5_PF1 (plane 11) Plane Enabled, Links OK
PFE #3
SIB0_PF0 (plane 0) Plane Enabled, Links OK
SIB0_PF1 (plane 1) Plane Enabled, Links OK
SIB1_PF0 (plane 2) Plane Enabled, Links OK
SIB1_PF1 (plane 3) Plane Enabled, Links OK
SIB2_PF0 (plane 4) Plane Enabled, Links OK
SIB2_PF1 (plane 5) Plane Enabled, Links OK
SIB3_PF0 (plane 6) Plane Enabled, Links OK
SIB3_PF1 (plane 7) Plane Enabled, Links OK
SIB4_PF0 (plane 8) Plane Enabled, Links OK
SIB4_PF1 (plane 9) Plane Enabled, Links OK
SIB5_PF0 (plane 10) Plane Enabled, Links OK
SIB5_PF1 (plane 11) Plane Enabled, Links OK

```

show chassis fabric plane-location

List of Syntax	Syntax on page 624 Syntax (MX Series Routers) on page 624 Syntax (MX2010 3D Universal Edge Routers) on page 624 Syntax (MX2020 3D Universal Edge Routers) on page 624 Syntax (TX Matrix Plus Router) on page 624 Syntax (QFX Switches) on page 624
Syntax	show chassis fabric plane-location
Syntax (MX Series Routers)	show chassis fabric plane-location <all-members> <local> <member <i>member-id</i> >
Syntax (MX2010 3D Universal Edge Routers)	show chassis fabric plane-location
Syntax (MX2020 3D Universal Edge Routers)	show chassis fabric plane-location
Syntax (TX Matrix Plus Router)	show chassis fabric plane-location
Syntax (QFX Switches)	show chassis fabric plane-location
Release Information	Command introduced in Junos OS Release 8.0. Command introduced in Junos OS Release 9.4 for EX Series switches. Command introduced in Junos OS Release 12.1x48 for PTX Series Packet Transport Routers. Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers. Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers. Command introduced in Junos OS Release 15.1X53-D30 for QFX Series switches.
Description	(M120, MX Series routers, and EX8200 switches only) Display the Control Board (CB) location of each plane. This command can be used on the master Routing Engine or the backup Routing Engine. For information about the meaning of “CBs” and “fabric plane” on the switches, see <i>EX Series Switches Hardware and CLI Terminology Mapping</i> . (TX Matrix Plus routers only) Display the SIB location of each fabric plane. (PTX Series Packet Transport Routers and QFX Series switches only) Display the fabric plane location of each SIB. (MX2010 and MX2020 Routers only) Display the fabric plane location of each Switch Fabric Board (SFB).

Options **all-members**—(MX Series routers only) (Optional) Display the CB location of each fabric plane on the Routing Engines in all member routers in the Virtual Chassis configuration.

local—(MX Series routers only) (Optional) Display the CB location of each fabric plane on the Routing Engines in the local Virtual Chassis member.

member *member-id*—(MX Series routers only) (Optional) Display the CB location of each fabric plane on the Routing Engines in the specified member in the Virtual Chassis configuration. Replace ***member-id*** with a value of 0 or 1.

Required Privilege Level view

List of Sample Output [show chassis fabric plane-location \(M120 Router\) on page 626](#)
[show chassis fabric plane-location \(MX240 and MX480 Routers\) on page 626](#)
[show chassis fabric plane-location \(MX960 Router\) on page 626](#)
[show chassis fabric plane-location \(MX2010 Router\) on page 626](#)
[show chassis fabric plane-location \(MX2020 Router\) on page 627](#)
[show chassis fabric plane-location \(TX Matrix Plus Router\) on page 627](#)
[show chassis fabric plane-location \(TX Matrix Plus Router with 3D SIBs\) on page 627](#)
[show chassis fabric plane-location \(EX8200 Switch\) on page 627](#)
[show chassis fabric plane-location \(PTX Series Packet Transport Routers\) on page 627](#)
[show chassis fabric plane-location \(QFX 10008 Switch\) on page 628](#)

Output Fields [Table 52](#) lists the output fields for the **show chassis fabric plane-location** command. Output fields are listed in the approximate order in which they appear.

Table 52: show chassis fabric plane-location Output Fields

Field Name	Field Description
Plane <i>n</i>	Plane number. (PTX Series Packet Transport Routers and QFX Series switches) Plane numbers associated with the SIB. (MX2010 and MX2020 Routers only) Plane numbers associated with the SFB.
Control Board <i>n</i>	Control board number.
SFC ABS-SIB-F13	(TX Matrix Plus routers only) Switch Interface Board (SIB) slot number on the F13 SIB.
SFC ABS-SIB-F2S	(TX Matrix Plus routers only) SIB slot number on the F2S SIB.
LCC ST-SIB-L	(TX Matrix Plus routers only) Line-card chassis (LCC) SIB slot number.
SFC SIB F13	(TX Matrix Plus routers with 3D SIBs only) Switch Interface Board (SIB) slot number on the F13 SIB.

Table 52: show chassis fabric plane-location Output Fields (*continued*)

Field Name	Field Description
SFC SIB F2S	(TX Matrix Plus routers with 3D SIBs only) SIB slot number on the F2S SIB.
LCC SIB	(TX Matrix Plus routers with 3D SIBs only) Line-card chassis (LCC) SIB slot number.
SIB	(PTX Series Packet Transport Routers and QFX Series switches) SIB number.
Switch Fabric Board <i>n</i>	(MX2010 and MX2020 Routers only) SFB number.

Sample Output

show chassis fabric plane-location (M120 Router)

```

user@host> show chassis fabric plane-location
-----Fabric Plane Locations-----
Plane 0                      Control Board 0
Plane 1                      Control Board 0
Plane 2                      Control Board 1
Plane 3                      Control Board 1

```

show chassis fabric plane-location (MX240 and MX480 Routers)

```

user@host> show chassis fabric plane-location
-----Fabric Plane Locations-----
Plane 0                      Control Board 0
Plane 1                      Control Board 0
Plane 2                      Control Board 0
Plane 3                      Control Board 0
Plane 4                      Control Board 1
Plane 5                      Control Board 1
Plane 6                      Control Board 1
Plane 7                      Control Board 1

```

show chassis fabric plane-location (MX960 Router)

```

user@host> show chassis fabric plane-location
-----Fabric Plane Locations-----
Plane 0                      Control Board 0
Plane 1                      Control Board 0
Plane 2                      Control Board 1
Plane 3                      Control Board 1
Plane 4                      Control Board 2
Plane 5                      Control Board 2

```

show chassis fabric plane-location (MX2010 Router)

```

user@host> show chassis fabric plane-location
-----Fabric Plane Locations-----
Plane 0                      Switch Fabric Board 0
Plane 1                      Switch Fabric Board 1
Plane 2                      Switch Fabric Board 2
Plane 3                      Switch Fabric Board 3
Plane 4                      Switch Fabric Board 4

```

```

Plane 5          Switch Fabric Board 5
Plane 6          Switch Fabric Board 6
Plane 7          Switch Fabric Board 7

```

show chassis fabric plane-location (MX2020 Router)

```

user@host> show chassis fabric plane-location
-----Fabric Plane Locations-----
Plane 0          Switch Fabric Board 0
Plane 1          Switch Fabric Board 1
Plane 2          Switch Fabric Board 2
Plane 3          Switch Fabric Board 3
Plane 4          Switch Fabric Board 4
Plane 5          Switch Fabric Board 5
Plane 6          Switch Fabric Board 6
Plane 7          Switch Fabric Board 7

```

show chassis fabric plane-location (TX Matrix Plus Router)

```

user@host> show chassis fabric plane-location
Fabric Plane Locations :
Plane      SFC ABS-SIB-F13      SFC ABS-SIB-F2      LCC ST-SIB-L
0           0, 1                0/0, 0/2, 0/4, 0/6      0
1           3, 4                1/0, 1/2, 1/4, 1/6      1
2           6, 7                2/0, 2/2, 2/4, 2/6      2
3           8, 9                3/0, 3/2, 3/4, 3/6      3
4          11, 12              4/0, 4/2, 4/4, 4/6      4

```

show chassis fabric plane-location (TX Matrix Plus Router with 3D SIBs)

```

user@host> show chassis fabric plane-location
sfc0-re0
-----Fabric Plane Locations-----
Plane      SFC SIB F13      SFC SIB F2      LCC SIB
0           0, 1                0/0, 0/2, 0/4, 0/6      0
1           3, 4                1/0, 1/2, 1/4, 1/6      1
2           6, 7                2/0, 2/2, 2/4, 2/6      2
3           8, 9                3/0, 3/2, 3/4, 3/6      3
4          11, 12              4/0, 4/2, 4/4, 4/6      4

```

show chassis fabric plane-location (EX8200 Switch)

```

user@host> show chassis fabric plane-location
-----Fabric Plane Locations-----
Plane 0          Control Board 0
Plane 1          Control Board 0
Plane 2          Control Board 0
Plane 3          Control Board 0
Plane 4          Control Board 1
Plane 5          Control Board 1
Plane 6          Control Board 1
Plane 7          Control Board 1
Plane 8          Control Board 2
Plane 9          Control Board 2
Plane 10         Control Board 2
Plane 11         Control Board 2

```

show chassis fabric plane-location (PTX Series Packet Transport Routers)

```

user@host> show chassis fabric plane-location

```

```
-----Fabric Plane Locations-----
SIB          Planes
0            0    1
1            2    3
2            4    5
3            6    7
4            8    9
5           10   11
6           12   13
7           14   15
8           16   17
```

show chassis fabric plane-location (QFX 10008 Switch)

```
user@host> show chassis fabric plane-location
-----Fabric Plane Locations-----
SIB          Planes
0            0    1
1            2    3
2            4    5
3            6    7
4            8    9
5           10   11
```

show chassis fabric sibs

Syntax	<pre>show chassis fabric sibs <fcc number scc> <slot slot-number></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced on QFX Series switches in Junos OS Release 15.1X53-D30</p>
Description	<p>(TX Matrix routers only) Display the state of the electrical and optical switch fabric link between the SIBs in the TX Matrix router (TX-SIBs) and the SIBs in the T640 routers (T640 LCC SIBs).</p> <p>(M320, T640, T1600, T4000 and PTX Series routers and QFX Series switches) Display the state of the electrical switch fabric link between the SIBs and the FPCs.</p>
Options	<p>none—(TX Matrix routers only) Display the state of the electrical and optical switch fabric link between the SIBs in the TX Matrix router (TX-SIBs) and the SIBs in the T640 routers (T640 LCC SIBs).</p> <p>(M320, T640, T1600, T4000 and PTX Series routers and QFX Series switches) Display the state of the electrical switch fabric link between the SIBs and the FPCs.</p> <p>fcc number—(Optional) Display the switching fabric link state for the T640 SIBs on a specified T640 router (line-card chassis) connected to a TX Matrix router.</p> <p>scc—(Optional) Display the switching fabric link state for the TX-SIBs on the TX Matrix router (switch-card chassis).</p> <p>slot slot-number—(Optional) Display the state of the electrical switch fabric link between the specified SIB slot and the FPCs.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>request chassis sib</i> • <i>show chassis sibs</i> • <i>Monitoring the SIBs</i> • <i>Redundant SIBs Overview</i>
List of Sample Output	<p>show chassis fabric sibs (M320 Router) on page 632</p> <p>show chassis fabric sibs (T640 Router) on page 632</p> <p>show chassis fabric sibs (T1600 Router) on page 633</p> <p>show chassis fabric sibs (T4000 Core Router) on page 635</p> <p>show chassis fabric sibs (TX Matrix Router) on page 636</p> <p>show chassis fabric sibs fcc (TX Matrix Router) on page 638</p> <p>show chassis fabric sibs scc (TX Matrix Router) on page 639</p> <p>show chassis fabric sibs slot (PTX3000 Router) on page 639</p> <p>show chassis fabric sibs (QFX10008 Switch) on page 640</p>

Output Fields Table 53 lists the output fields for the **show chassis fabric sibs** command. Output fields are listed in the approximate order in which they appear.

Table 53: show chassis fabric sibs Output Fields

Field Name	Field Description
Fabric management SIB state	Switching fabric link (link from FPC to SIB) state for each SIB: <ul style="list-style-type: none">• Unused—SIB is not present.• Links ok—Link between the SIB and the FPC is active.• Link error—Link between the SIB and the FPC is not operational.

Table 53: show chassis fabric sibs Output Fields (*continued*)

Field Name	Field Description
Plane state	<p>Possible plane state of the M320 SIB, TX-SIB or T640 SIB:</p> <ul style="list-style-type: none"> • S_ACTIVE—Links on the SIB are operational, and the fabric plane (SIB) is operational and running. • S_SPARE—Links on the SIB are operational and the fabric plane (SIB) is redundant and can be operational if any of the fabric planes in the S_ACTIVE state encounters an error. <p>NOTE: If the plane is unusable by any of the Packet Forwarding Engines, the command output displays an additional string, plane has link errors on # pfes, where, # indicates the total number of links (both from SIB to FPC, and from FPC to SIB) having link errors (detected either during initialization time or runtime) in this particular plane. This does not count links having destination errors.</p> <ul style="list-style-type: none"> • S_EMPTY—No links are present on the SIB, and the fabric plane (SIB) is powered down. • S_ACTIVATING—Links on the SIB are coming online; this is a transitional state. • S_DEACTIVATING—Links on the SIB are going offline; this is a transitional state. • S_FAULTING—Links on the SIB are being marked faulty, and the fabric plane (SIB) is not operational. • S_FAULT—Links on the SIB are in an alarmed state, and the fabric plane (SIB) is not operational for the following reasons: <ul style="list-style-type: none"> • On-board F-chip is not operational. • Fiber optic connector faults. • FPC connector faults. <p>Possible plane state of the QFX Series SIB:</p> <ul style="list-style-type: none"> • Active—Links on the SIB are operational, and the fabric plane (SIB) is operational and running. • Spare—Links on the SIB are operational and the fabric plane (SIB) is redundant and can be operational if any of the fabric planes in the S_ACTIVE state encounters an error. • Empty—No links are present on the SIB, and the fabric plane (SIB) is powered down. • Activating—Links on the SIB are coming online; this is a transitional state. • Deactivating—Links on the SIB are going offline; this is a transitional state. • Faulting—Links on the SIB are being marked faulty, and the fabric plane (SIB) is not operational. • Fault—Links on the SIB are in an alarmed state, and the fabric plane (SIB) is not operational for the following reasons: <ul style="list-style-type: none"> • On-board F-chip is not operational. • Fiber optic connector faults. • FPC connector faults.

Sample Output

show chassis fabric sibs (M320 Router)

```
user@host> show chassis fabric sibs
Fabric management SIB state:
SIB #0
  plane state: S_ACTIVE
  FPC #0
    PFE #1 : Links ok
  FPC #1
    PFE #1 : Links ok
  FPC #2
    PFE #1 : Links ok
  FPC #3
    PFE #1 : Links ok
SIB #1
  plane state: S_ACTIVE
  FPC #0
    PFE #1 : Links ok
  FPC #1
    PFE #1 : Links ok
  FPC #2
    PFE #1 : Links ok
  FPC #3
    PFE #1 : Links ok
SIB #2
  plane state: S_ACTIVE
  FPC #0
    PFE #1 : Links ok
  FPC #1
    PFE #1 : Links ok
  FPC #2
    PFE #1 : Links ok
  FPC #3
    PFE #1 : Links ok
SIB #3
  plane state: S_ACTIVE
  FPC #0
    PFE #1 : Links ok
  FPC #1
    PFE #1 : Links ok
  FPC #2
    PFE #1 : Links ok
  FPC #3
    PFE #1 : Links ok
```

show chassis fabric sibs (T640 Router)

```
user@host> show chassis fabric sibs
Fabric management SIB state:
SIB #0
  plane state: S_SPARE
  FPC #0
    PFE #1 : Links ok
  FPC #2
    PFE #1 : Links ok
  FPC #3
    PFE #0 : Links ok
    PFE #1 : Links ok
SIB #1
```



```

plane state: S_ACTIVE
FPC #0
  PFE #1 : Links ok
FPC #2
  PFE #1 : Links ok
FPC #3
  PFE #0 : Links ok
  PFE #1 : Links ok
SIB #2
plane state: S_ACTIVE
FPC #0
  PFE #1 : Links ok
FPC #2
  PFE #1 : Links ok
FPC #3
  PFE #0 : Links ok
  PFE #1 : Links ok
SIB #3
plane state: S_ACTIVE
FPC #0
  PFE #1 : Links ok
FPC #2
  PFE #1 : Links ok
FPC #3
  PFE #0 : Links ok
  PFE #1 : Links ok
SIB #4
plane state: S_ACTIVE
FPC #0
  PFE #1 : Links ok
FPC #2
  PFE #1 : Links ok
FPC #3
  PFE #0 : Links ok
  PFE #1 : Links ok

```

show chassis fabric sibs (T1600 Router)

```

user@host> show chassis fabric sibs
SIB #0
plane state: S_SPARE
FPC #0
  PFE #0 : Links ok
  PFE #1 : Links ok
FPC #1
  PFE #0 : Links ok
  PFE #1 : Links ok
FPC #2
  PFE #0 : Links ok
FPC #4
  PFE #0 : Links ok
  PFE #1 : Links ok
FPC #5
  PFE #0 : Links ok
FPC #6
  PFE #0 : Links ok
  PFE #1 : Links ok
FPC #7
  PFE #0 : Links ok
  PFE #1 : Links ok
SIB #1

```

```
plane state: S_ACTIVE , plane has link errors on 2 pfes
FPC #0
  PFE #0 : Links ok
  PFE #1 : Links ok
FPC #1
  PFE #0 : Links ok
  PFE #1 : Links ok
FPC #3
  PFE #0 : Links ok
  PFE #1 : Links ok
FPC #4
  PFE #0 : Links ok
  PFE #1 : Links ok
FPC #5
  PFE #0 : Links ok
  PFE #1 : Links ok
FPC #7
  PFE #0 : Links ok
  PFE #1 : Links okSIB #2
plane state: S_ACTIVE
SIB #2
  plane state: S_ACTIVE
FPC #0
  PFE #0 : Links ok
  PFE #1 : Links ok
FPC #1
  PFE #0 : Links ok
  PFE #1 : Links ok
FPC #2
  PFE #0 : Links ok
FPC #4
  PFE #0 : Links ok
  PFE #1 : Links ok
FPC #5
  PFE #0 : Links ok
FPC #6
  PFE #0 : Links ok
  PFE #1 : Links ok
FPC #7
  PFE #0 : Links ok
  PFE #1 : Links ok
SIB #3
plane state: S_ACTIVE
FPC #0
  PFE #0 : Links ok
  PFE #1 : Links ok
FPC #1
  PFE #0 : Links ok
  PFE #1 : Links ok
FPC #2
  PFE #0 : Links ok
FPC #4
  PFE #0 : Links ok
  PFE #1 : Links ok
FPC #5
  PFE #0 : Links ok
FPC #6
  PFE #0 : Links ok
  PFE #1 : Links ok
FPC #7
  PFE #0 : Links ok
```

```

        PFE #1 : Links ok
SIB #4
  plane state: S_ACTIVE
  FPC #0
    PFE #0 : Links ok
    PFE #1 : Links ok
  FPC #1
    PFE #0 : Links ok
    PFE #1 : Links ok
  FPC #2
    PFE #0 : Links ok
  FPC #4
    PFE #0 : Links ok
    PFE #1 : Links ok
  FPC #5
    PFE #0 : Links ok
  FPC #6
    PFE #0 : Links ok
    PFE #1 : Links ok
  FPC #7
    PFE #0 : Links ok
    PFE #1 : Links ok

```

show chassis fabric sibs (T4000 Core Router)

```

user@host> show chassis fabric sibs
Fabric management SIB state:
SIB #0
  plane state: S_SPARE
  FPC #2
    PFE #0 : Links ok
  FPC #3
    PFE #0 : Links ok
  FPC #5
    PFE #0 : Links ok
    PFE #1 : Links ok
  FPC #6
    PFE #0 : Links ok
    PFE #1 : Links ok
SIB #1
  plane state: S_ACTIVE
  FPC #2
    PFE #0 : Links ok
  FPC #3
    PFE #0 : Links ok
  FPC #5
    PFE #0 : Links ok
    PFE #1 : Links ok
  FPC #6
    PFE #0 : Links ok
    PFE #1 : Links ok
SIB #2
  plane state: S_ACTIVE
  FPC #2
    PFE #0 : Links ok
  FPC #3
    PFE #0 : Links ok
  FPC #5
    PFE #0 : Links ok
    PFE #1 : Links ok
  FPC #6

```

```
        PFE #0 : Links ok
        PFE #1 : Links ok
SIB #3
  plane state: S_ACTIVE
  FPC #2
    PFE #0 : Links ok
  FPC #3
    PFE #0 : Links ok
  FPC #5
    PFE #0 : Links ok
    PFE #1 : Links ok
  FPC #6
    PFE #0 : Links ok
    PFE #1 : Links ok
SIB #4
  plane state: S_ACTIVE
  FPC #2
    PFE #0 : Links ok
  FPC #3
    PFE #0 : Links ok
  FPC #5
    PFE #0 : Links ok
    PFE #1 : Links ok
  FPC #6
    PFE #0 : Links ok
    PFE #1 : Links ok
```

show chassis fabric sibs (TX Matrix Router)

```
user@host> show chassis fabric sibs
scc-re0:
-----
Fabric management SIB state:
SIB #1
  plane state: S_ACTIVE , plane has link errors on 2 pfes
  FPC #0
    PFE #0 : Links ok
    PFE #1 : Links ok
  FPC #1
    PFE #0 : Links ok
    PFE #1 : Links ok
  FPC #3
    PFE #0 : Links ok
    PFE #1 : Links ok
  FPC #4
    PFE #0 : Links ok
    PFE #1 : Links ok
  FPC #5
    PFE #0 : Links ok
    PFE #1 : Links ok
  FPC #7
    PFE #0 : Links ok
    PFE #1 : Links ok
SIB #2
  plane state: S_ACTIVE
  LCC #0 : Links ok
  LCC #1 : Links ok
SIB #3
  plane state: S_ACTIVE
  LCC #0 : Links ok
  LCC #1 : Links ok
```

```
SIB #4
  plane state: S_ACTIVE
  LCC #0      : Links ok
  LCC #1      : Links ok
```

```
lcc0-re0:
```

```
-----
Fabric management SIB state:
```

```
SIB #1
  plane state: S_ACTIVE
  FPC #0
    PFE #0 : Links ok
    PFE #1 : Links ok
  FPC #1
    PFE #1 : Links ok
  FPC #2
    PFE #0 : Links ok
    PFE #1 : Links ok
  FPC #3
    PFE #1 : Links ok
  FPC #4
    PFE #1 : Links ok
  FPC #5
    PFE #0 : Links ok
  FPC #6
    PFE #1 : Links ok
  FPC #7
    PFE #1 : Links ok
  SCC      : Links ok
```

```
SIB #2
  plane state: S_ACTIVE
  FPC #0
    PFE #0 : Links ok
    PFE #1 : Links ok
  FPC #1
    PFE #1 : Links ok
  FPC #2
    PFE #0 : Links ok
    PFE #1 : Links ok
  FPC #3
    PFE #1 : Links ok
  FPC #4
    PFE #1 : Links ok
  FPC #5
    PFE #0 : Links ok
  FPC #6
    PFE #1 : Links ok
  FPC #7
    PFE #1 : Links ok
  SCC      : Links ok
```

```
SIB #3
  plane state: S_ACTIVE
  FPC #0
    PFE #0 : Links ok
    PFE #1 : Links ok
  FPC #1
    PFE #1 : Links ok
  FPC #2
    PFE #0 : Links ok
    PFE #1 : Links ok
  FPC #3
```

```
        PFE #1 : Links ok
FPC #4
        PFE #1 : Links ok
FPC #5
        PFE #0 : Links ok
FPC #6
        PFE #1 : Links ok
FPC #7
        PFE #1 : Links ok
SCC      : Links ok
SIB #4
plane state: S_ACTIVE
FPC #0
        PFE #0 : Links ok
        PFE #1 : Links ok
FPC #1
        PFE #1 : Links ok
FPC #2
        PFE #0 : Links ok
        PFE #1 : Links ok
FPC #3
        PFE #1 : Links ok
FPC #4
        PFE #1 : Links ok
FPC #5
        PFE #0 : Links ok
FPC #6
        PFE #1 : Links ok
FPC #7
        PFE #1 : Links ok
SCC      : Links o
```

show chassis fabric sibs lcc (TX Matrix Router)

```
user@host> show chassis fabric sibs lcc 0
lcc1-re0:
```

```
-----
Fabric management SIB state:
```

```
SIB #1
plane state: S_ACTIVE
FPC #0
        PFE #0 : Links ok
FPC #2
        PFE #1 : Links ok
FPC #4
        PFE #0 : Links ok
FPC #5
        PFE #1 : Links ok
FPC #7
        PFE #0 : Links ok
SCC      : Links ok
SIB #2
plane state: S_ACTIVE
FPC #0
        PFE #0 : Links ok
FPC #2
        PFE #1 : Links ok
FPC #4
        PFE #0 : Links ok
FPC #5
        PFE #1 : Links ok
```

```

FPC #7
  PFE #0 : Links ok
  SCC    : Links ok
SIB #3
  plane state: S_ACTIVE
  FPC #0
    PFE #0 : Links ok
  FPC #2
    PFE #1 : Links ok
  FPC #4
    PFE #0 : Links ok
  FPC #5
    PFE #1 : Links ok
  FPC #7
    PFE #0 : Links ok
  SCC    : Links ok
SIB #4
  plane state: S_ACTIVE
  FPC #0
    PFE #0 : Links ok
  FPC #2
    PFE #1 : Links ok
  FPC #4
    PFE #0 : Links ok
  FPC #5
    PFE #1 : Links ok
  FPC #7
    PFE #0 : Links ok
  SCC    : Links ok

```

show chassis fabric sibs scc (TX Matrix Router)

```

user@host> show chassis fabric sibs scc
scc-re0:

```

```

-----
Fabric management SIB state:
SIB #1
  plane state: S_ACTIVE
  LCC #0      : Links ok
  LCC #1      : Links ok
SIB #2
  plane state: S_ACTIVE
  LCC #0      : Links ok
  LCC #1      : Links ok
SIB #3
  plane state: S_ACTIVE
  LCC #0      : Links ok
  LCC #1      : Links ok
SIB #4
  plane state: S_ACTIVE
  LCC #0      : Links ok
  LCC #1      : Links ok

```

show chassis fabric sibs slot (PTX3000 Router)

```

user@host> show chassis fabric sibs slot 0
Fabric management SIB state:
SIB #0 Online
  Fcore #0 (plane 0) Active
    FPC #8
      PFE #0 : OK

```

```
        PFE #1 : OK
    FPC #12
        PFE #0 : OK
        PFE #1 : OK
Fcore #1 (plane 1) Active
    FPC #8
        PFE #0 : OK
        PFE #1 : OK
    FPC #12
        PFE #0 : OK
        PFE #1 : OK
```

show chassis fabric sibs (QFX10008 Switch)

```
user@host> show chassis fabric sibs
Fabric management SIB state:
SIB #0 Online
    FASIC #0 (plane 0) Active
        FPC #0
            PFE #0 : OK
            PFE #1 : OK
            PFE #2 : OK
            PFE #3 : OK
        FPC #1
            PFE #0 : OK
            PFE #1 : OK
    FASIC #1 (plane 1) Active
        FPC #0
            PFE #0 : OK
            PFE #1 : OK
        FPC #12
            PFE #0 : OK
            PFE #1 : OK
SIB #1 Empty
SIB #2 Empty
SIB #3 Empty
SIB #4 Empty
SIB #5 Empty
```


show chassis fabric summary

Syntax	<code>show chassis fabric summary <extended></code>
Release Information	<p>Command introduced in Junos OS Release 8.4.</p> <p>Command introduced in Junos OS Release 9.4 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 15.1X53-D30 for QFX Series switches.</p> <p>extended option added in Junos OS Release 14.1R2.</p>
Description	<p>(MX Series routers and EX8200 switches only) Display the state of all fabric planes and the elapsed uptime.</p> <p>(QFX Series switches) Display the state of all fabric planes.</p>
Options	extended —(Optional) Display the extended summary of fabric planes.
Required Privilege Level	view
List of Sample Output	<p>show chassis fabric summary (MX240 Router) on page 643</p> <p>show chassis fabric summary (MX480 Router) on page 643</p> <p>show chassis fabric summary (MX480 Router with MPC4E) on page 643</p> <p>show chassis fabric summary (MX960 Router) on page 644</p> <p>show chassis fabric summary (MX2010 Router) on page 644</p> <p>show chassis fabric summary (MX2020 Router) on page 644</p> <p>show chassis fabric summary (MX2020 Router with MPC4E) on page 644</p> <p>show chassis fabric summary (EX8200 Switch) on page 645</p> <p>show chassis fabric summary (PTX Series Packet Transport Router) on page 645</p> <p>show chassis fabric summary (QFX 10008 Switch) on page 645</p> <p>show chassis fabric summary extended (MX960 Router) on page 646</p>
Output Fields	Table 54 lists the output fields for the show chassis fabric summary command. Output fields are listed in the approximate order in which they appear.

Table 54: show chassis fabric summary Output Fields

Field Name	Field Description
Plane	(MX Series, MX2020 and MX2010 Routers only) Plane number.

Table 54: show chassis fabric summary Output Fields (*continued*)

Field Name	Field Description
State	<p>(MX Series and QFX Series) State of the SIB or FPC:</p> <ul style="list-style-type: none"> • Online—Switch Interface Board (SIB) is operational and running. <p>NOTE: On the Enhanced MX SCB with Trio MPC, a maximum of 4 planes are operational and running. On all the other SCBs with Trio MPC, all the planes are operational and running.</p> <ul style="list-style-type: none"> • Empty—SIB is powered down. • Check—SIB is in the Check state because of the following reasons: <ul style="list-style-type: none"> • SIB is not inserted properly. • Some destination errors are detected on the SIB. In this case, the Packet Forwarding Engine stops using the SIB to send traffic to the affected destination Packet Forwarding Engine. • Some link errors are detected on the channel between the SIB and a Packet Forwarding Engine. Link errors can be detected at initialization time or runtime: <ul style="list-style-type: none"> • Link errors caused by a link training failure at initialization time—The Packet Forwarding Engine does not use the SIB to send traffic. The show chassis fabric fpcs command shows Plane disabled as status for this link. • Link errors caused by CRC errors detected at runtime—The Packet Forwarding Engine continues to use the SIB to send traffic. The show chassis fabric fpcs command shows Link error as the status for this link. <p>NOTE: The Check state does not apply to PTX Series Packet Transport Routers because there are no SIBs in the Check state.</p> <p>For information about link and destination errors, issue the show chassis fabric fpcs commands.</p> <ul style="list-style-type: none"> • Spare—SIB is redundant and will move to active state if one of the working SIBs fails. <p>NOTE: Spare does not apply to PTX Series Packet Transport Routers because there are no spare SIBs in the device.</p> <p>(MX2010 and MX2020 Routers) State of the SFB.</p> <ul style="list-style-type: none"> • Online—Switch Fabric Board (SFB) is operational and running. • Offline—Switch Fabric Board (SFB) is powered down. • Check—Switch Fabric Board (SFB) is in the check state.
Errors	<p>(PTX Series and QFX Series) Indicates whether there is any error on the SIB.</p> <ul style="list-style-type: none"> • None—No errors • Link Errors—Fabric link errors were found on the SIB RX link. • Cell drops—Fabric cell drops were found on the SIB ASIC. • Link, Cell drops—Both Link errors and cell drops were detected on at least one of the FPC's fabric links.

Table 54: show chassis fabric summary Output Fields (*continued*)

Field Name	Field Description
	<ul style="list-style-type: none"> Asic Errors—A fault affecting one of the ASICs on the SIB is detected. It can be an IO error or an internal error signaled by the ASIC. <p>NOTE: The Errors column is empty only when the FPC or SIB is offline.</p>
Uptime	(MX Series, MX2010 and MX2020 Routers) Elapsed time the plane has been online.
Link Error	Fabric link errors were found on the SIB RX link.
Link TF	Fabric link training failure has occurred.
Destination errors	<ul style="list-style-type: none"> Local—Destination error detected on the FPC or PFE's own self-stream. Remote—Destination error detected on the FPC or PFE's non-self-streams.

Sample Output

show chassis fabric summary (MX240 Router)

```
user@host> show chassis fabric summary
Plane  State  Uptime
0      Online 23 hours, 26 minutes, 54 seconds
1      Online 23 hours, 26 minutes, 54 seconds
2      Check 18 hours, 33 minutes, 42 seconds
3      Online 23 hours, 26 minutes, 54 seconds
4      Spare 23 hours, 26 minutes, 54 seconds
5      Spare 23 hours, 26 minutes, 54 seconds
6      Spare 23 hours, 26 minutes, 54 seconds
7      Spare 23 hours, 26 minutes, 54 seconds
```

show chassis fabric summary (MX480 Router)

```
user@host> show chassis fabric summary
Plane  State  Uptime
0      Online 8 hours, 45 minutes, 29 seconds
1      Online 8 hours, 45 minutes, 28 seconds
2      Online 8 hours, 45 minutes, 28 seconds
3      Online 8 hours, 45 minutes, 28 seconds
4      Spare 8 hours, 45 minutes, 28 seconds
5      Spare 8 hours, 45 minutes, 28 seconds
6      Spare 8 hours, 45 minutes, 28 seconds
7      Check 6 hours, 10 minutes, 12 seconds
```

show chassis fabric summary (MX480 Router with MPC4E)

```
user@host > show chassis fabric summary
Plane  State  Uptime
0      Online 6 hours, 57 minutes, 44 seconds
1      Online 6 hours, 57 minutes, 40 seconds
```

2	Online	6 hours, 57 minutes, 39 seconds
3	Online	6 hours, 57 minutes, 34 seconds
4	Spare	6 hours, 57 minutes, 34 seconds
5	Spare	6 hours, 57 minutes, 29 seconds
6	Spare	6 hours, 57 minutes, 29 seconds
7	Spare	6 hours, 57 minutes, 24 seconds

Note:

For FPC slots with MPC Type 4 or MCC:

Fabric planes 1 and 5, 3 and 7 use shared physical links.

Those slots may run in a reduced bandwidth in case both plane 1 and 5, or both 3 and 7 are active.

show chassis fabric summary (MX960 Router)

```
user@host> show chassis fabric summary
Plane  State      Uptime
0       Online    3 hours, 7 minutes, 9 seconds
1       Online    3 hours, 7 minutes, 4 seconds
2       Online    3 hours, 6 minutes, 59 seconds
3       Online    3 hours, 6 minutes, 54 seconds
4       Empty
5       Empty
```

show chassis fabric summary (MX2010 Router)

```
user@host> show chassis fabric summary
Plane  State      Uptime
0       Online    1 day, 13 hours, 20 minutes, 10 seconds
1       Online    1 day, 13 hours, 19 minutes, 59 seconds
2       Online    1 day, 13 hours, 19 minutes, 49 seconds
3       Offline
4       Online    1 day, 13 hours, 19 minutes, 28 seconds
5       Check    1 day, 13 hours, 19 minutes, 17 seconds
6       Online    1 day, 13 hours, 19 minutes, 6 seconds
7       Online    1 hour, 43 minutes, 5 seconds
```

show chassis fabric summary (MX2020 Router)

```
user@host> show chassis fabric summary
Plane  State      Uptime
0       Online    8 hours, 24 minutes, 1 second
1       Online    8 hours, 47 minutes, 54 seconds
2       Online    8 hours, 47 minutes, 44 seconds
3       Online    8 hours, 47 minutes, 33 seconds
4       Online    8 hours, 47 minutes, 22 seconds
5       Online    8 hours, 47 minutes, 12 seconds
6       Online    8 hours, 47 minutes, 1 second
7       Online    8 hours, 46 minutes, 50 seconds
```

show chassis fabric summary (MX2020 Router with MPC4E)

```
user@host > show chassis fabric summary
Plane  State      Uptime
0       Online    3 days, 6 hours, 58 minutes, 29 seconds
1       Online    3 days, 6 hours, 58 minutes, 18 seconds
2       Online    3 days, 6 hours, 58 minutes, 8 seconds
3       Online    3 days, 6 hours, 57 minutes, 57 seconds
4       Online    3 days, 6 hours, 57 minutes, 46 seconds
5       Online    3 days, 6 hours, 57 minutes, 36 seconds
6       Online    3 days, 6 hours, 57 minutes, 25 seconds
7       Online    3 days, 6 hours, 57 minutes, 14 seconds
```

show chassis fabric summary (EX8200 Switch)

```

user@host> show chassis fabric summary
Plane  State      Uptime
0       Online    12 days, 50 minutes, 54 seconds
1       Online    12 days, 50 minutes, 53 seconds
2       Online    12 days, 50 minutes, 53 seconds
3       Online    12 days, 50 minutes, 52 seconds
4       Spare     12 days, 50 minutes, 49 seconds
5       Spare     12 days, 50 minutes, 47 seconds
6       Spare     12 days, 50 minutes, 47 seconds
7       Spare     12 days, 50 minutes, 46 seconds
8       Online    12 days, 50 minutes, 52 seconds
9       Online    12 days, 50 minutes, 50 seconds
10      Online    12 days, 50 minutes, 50 seconds
11      Online    12 days, 50 minutes, 49 seconds

```

show chassis fabric summary (PTX Series Packet Transport Router)

```

user@host> show chassis fabric summary
FRU           State      Errors
SIB0          Online     Asic Errors
SIB1          Online     Link Errors
SIB2          Online     None
SIB3          Online     Cell drops
SIB4          Offline
SIB5          Online     None
SIB6          Online     Link, Cell drops
SIB7          Online     None
SIB8          Online     Link, Cell drops

FPC0          Online     None
FPC1          Online     Link Errors
FPC2          Online     None
FPC3          Offline
FPC4          Online     None
FPC5          Online     None
FPC6          Empty
FPC7          Empty

```

show chassis fabric summary (QFX 10008 Switch)

```

user@host> show chassis fabric summary
FRU           State      Errors
FPC0          Online     None
FPC1          Online     Link Errors
FPC2          Online     None
FPC3          Offline
FPC4          Online     None
FPC5          Online     None
FPC6          Empty
FPC7          Empty

SIB0          Online     None
SIB1          Online     Link Errors
SIB2          Online     None
SIB3          Online     Cell drops
SIB4          Offline

```

SIB5 Online None

Sample Output

show chassis fabric summary extended (MX960 Router)

```
user@host> show chassis fabric summary extended
Plane  State      Link  Link  Destination errors  Uptime
      Error TF   Local / Remote
0      Online   NO    NO     NO/  NO      7 days, 5 hours, 25 minutes,
20 seconds
1      Online   NO    NO     NO/  NO      7 days, 5 hours, 25 minutes,
11 seconds
2      Online   NO    NO     NO/  NO      7 days, 5 hours, 25 minutes,
5 seconds
3      Online   NO    NO     NO/  NO      7 days, 5 hours, 24 minutes,
59 seconds
4      Spare    NO    NO     NO/  NO      7 days, 5 hours, 24 minutes,
52 seconds
5      Spare    NO    NO     NO/  NO      7 days, 5 hours, 24 minutes,
45 seconds
```

show chassis fabric topology

List of Syntax	Syntax on page 647 Syntax (TX Matrix Router) on page 647 Syntax (TX Matrix Plus Router) on page 647 Syntax (T4000 Core Router) on page 647 Syntax (PTX Series Packet Transport Routers) on page 647 Syntax (QTX Series Switches) on page 647
Syntax	<pre>show chassis fabric topology <lcc number scc> <sib-slot></pre>
Syntax (TX Matrix Router)	<pre>show chassis fabric topology <lcc number scc> <sib-slot></pre>
Syntax (TX Matrix Plus Router)	<pre>show chassis fabric topology <lcc number sfc number> <sib-slot></pre>
Syntax (T4000 Core Router)	<pre>show chassis fabric topology <sib-slot></pre>
Syntax (PTX Series Packet Transport Routers)	<pre>show chassis fabric topology</pre>
Syntax (QTX Series Switches)	<pre>show chassis fabric topology</pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 12.1 for PTX Series Packet Transport Routers.</p> <p>Command introduced in Junos OS Release 15.1X53-D30 for QFX Series switches.</p>
Description	<p>(TX Matrix routers only) Display the state of the switching fabric topology for the Switch Interface Board (SIB) connection between the TX Matrix router and the T640 routers.</p> <p>(TX Matrix Plus routers only) Display the state of the switching fabric topology for the SIB connection between the TX Matrix Plus router and the connected routers.</p> <p>(T320, T640, T1600, and T4000 routers only) Display the state of the switching fabric topology for the connection between the Switch Interface Board (SIB) and the FPCs.</p> <p>(PTX Series Packet Transport Routers and QFX Series switches) Display the input-output link topology.</p>
Options	<p>none—(TX Matrix routers only) Display the state of the switching fabric topology for the Switch Interface Board (SIB) connection between the TX Matrix router and the T640 routers.</p>

(TX Matrix Plus routers only) Display the state of the switching fabric topology for the SIB connection between the TX Matrix Plus router and the connected routers.

(T320, T640, T1600, and T4000 routers only) Display the state of the switching fabric topology for the connection between the Switch Interface Board (SIB) and the FPCs.

(QFX Series switches) Display the input-output link topology.

lcc number—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display the fabric topology state for a specified T640 router (line-card chassis) that is connected to a TX Matrix router. On a TX Matrix Plus router, display the fabric topology state for a specified router (line-card chassis) that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

scc—(TX Matrix routers only) (Optional) Display the fabric topology state for the TX Matrix router (or switch-card chassis).

sfc number—(TX Matrix Plus routers only) (Optional) Display the fabric topology for the switch-fabric chassis. Replace *number* with 0.

sib-slot—(Optional) Display the fabric topology state for a specified SIB slot. Replace *sib-slot* with a value from 0 through 4. On a TX Matrix Plus router, replace *sib-slot* with a value from 0 through 15.

Required Privilege Level view

Related Documentation

- [Layer 2 Wholesale Network Topology Overview](#)

List of Sample Output [show chassis fabric topology scc \(TX Matrix Router\) on page 652](#)
[show chassis fabric topology lcc on page 654](#)
[show chassis fabric topology \(TX Matrix Plus Router\) on page 656](#)
[show chassis fabric topology sfc \(TX Matrix Plus Router\) on page 657](#)
[show chassis fabric topology lcc \(TX Matrix Plus Router\) on page 658](#)
[show chassis fabric topology \(T4000 Core Router\) on page 659](#)
[show chassis fabric topology lcc \(TX Matrix Plus Router with 3D SIBs\) on page 660](#)
[show chassis fabric topology sfc \(TX Matrix Plus Router with 3D SIBs\) on page 662](#)

[show chassis fabric topology \(PTX5000 Router\) on page 666](#)

[show chassis fabric topology \(PTX3000 Router\) on page 669](#)

[show chassis fabric topology \(QFX10008 Switch\) on page 676](#)

Output Fields Table 55 lists the output fields for the **show chassis fabric topology** command. Output fields are listed in the approximate order in which they appear.

Table 55: show chassis fabric topology Output Fields

Field Name	Field Description
in-links	Fabric topology for receive side links.
out-links	Fabric topology for transmit side links.
state	<p>State of the fabric link:</p> <ul style="list-style-type: none"> • RESET—Link between the SIB and the FPC/DPC is powered down on purpose. This is done in all non-dual Packet Forwarding Engine–based boards. • UP—Link between the SIB and the FPC/DCP is up and running. • DOWN—Link between the SIB and the FPC/DCP is powered down. • FAULT—The SIB is in the alarmed state, in which the SIB's plane is not operational for one or more of the following reasons: <ul style="list-style-type: none"> • On-board F-chip is not operational. • Fiber-optic connector faults. • FPC connector faults. • SIB midplane connector faults. <p>NOTE: The following state descriptions are applicable only to PTX Series Packet Transport Routers.</p> <ul style="list-style-type: none"> • OK—The link between the SIB and the FPC is operational. • Down—The link between the SIB and the FPC is powered down. • Error—The CCL link between the SIB and FPC is not operational for one or more of the following reasons: <ul style="list-style-type: none"> • FPC midplane connector failure. • SIB midplane connector failure. • CCL link CRC error.

Table 55: show chassis fabric topology Output Fields (*continued*)

Out-Links: and In-Links (TX Matrix Plus router only)	State of the links from the F13 SIB to the LCC or vice-versa. Out-Links indicate Tx links. In-Links indicate an Rx link. The following additional fields are displayed for each SIB:
	<ul style="list-style-type: none"> • VCSEL Status—Optical (VCSEL channel) link status for the corresponding electrical (HSL2) link. The states include: <ul style="list-style-type: none"> • OK—Optical signal power is good. • Error—Internal error. • LOS—Loss of Signal detected. • High Cur—The Tx Bias-current is higher than threshold on this channel. This is applicable only to Tx Channels. • Low Cur—The Tx Bias-current is lower than threshold on this channel. This is applicable only to Tx Channels. • HSL2 Channel—HSL2 is the electrical link used to connect ASICs to the in-link and out-link. The channel number corresponds to the link and varies based on the ASIC or configuration.

- **HSL2 Status** —The status of the HSL2 Channel. Includes the following states:
 - **Up**—Channel is up.
 - **Down**—Channel is down.
 - **Reset**—Channel has been reset.
 - **Fault**—Channel has faults.

The following is a representation of display output for links originating from the SIBs (LCC or SFC)

SF_[1|3]_port#_FB_[A-D] (VCSEL#, fiber)

- **SF_[1|3]**—Name of the ASIC, with Fabric F1 or F3 mode.
- **port#**—HSL2 port number on the SF ASIC in the LCC.
- **FB_[A-D]**—via fiber bundle A, B, C or D.
- **VCSEL#**—VCSEL module number on SIB.
- **fiber**—Fiber channel number.

The following is a representation of display output for links originating from the SIBs (LCC or SFC)

SF_[1|3]_port#_FB_[A-D] (VCSEL#, fiber)

- **SF_[1|3]**—Name of the ASIC, with Fabric F1 or F3 mode.
- **port#**—HSL2 port number on the SF ASIC in the LCC.
- **FB_[A-D]**—via fiber bundle A, B, C or D.
- **VCSEL#**—VCSEL module number on SIB.
- **fiber**—Fiber channel number.

The following is a sample output with description of the fields displayed in the output for Out-Links:

Out-Links:

=====

SF_30_13_FB_A(21,09) -> FPC7_B_SG(3,3,6)_FB_A(18,09)	OK	203	Up
--	----	-----	----

Table 55: show chassis fabric topology Output Fields (*continued*)

- **SF_30_13**—Name of the ASIC, with Fabric F1 or F3 mode. In this case, 3 is the F3 direction and is used in the Tx path and 0 identifies the serial link on the SF chip (in this case, link goes to sf-3 chip number 0). You can also have F1 mode and Rx path instead.
- **FB_A (21, 09)**—Fiber bundle A, with VCSEL unit number 21 within the SIB, and channel number 9 within the unit number.
- **FPC7_B_SG(3,3,6)**—FPC 7.with bottom Packet Forwarding Engine (T for top PFE and B for bottom PFE), SG ASIC, with number 3 and port number 3, with HSL2 link number with the SIB as 6.
- **FB_A(18, 09)**—Fiber Bundle, with VCSEL unit number 18 within the SIB, and VCSEL channel number 9 within the unit number.

The following is a representation of display output for links originating from the FPCs (In-Links)

FPC#[T|B]_SG(ASIC#, port#, HSL2_bit)_FB_[A-D] (VCSEL#, fiber)

- **FPC#**—FPC number with PFE (0 or 1).
- **T**—Top Packet Forwarding Engine.
- **B**—Bottom Packet Forwarding Engine.
- **SG(ASIC#, port#, HSL2_bit)**—SG ASIC information (ASIC 0-3, port 0-3, HSL2_bit 0-7).
- **FB_[A-D]**—via fiber bundle A, B, C or D.
- **VCSEL#**—VCSEL module number on SIB.
- **fiber**—Fiber channel number.

The following is a representation of display output for links originating from the FPCs (In-Links)

FPC#[T|B]_SG(ASIC#, port#, HSL2_bit)_FB_[A-D] (VCSEL#, fiber)

- **FPC#**—FPC number with PFE (0 or 1).
- **T**—Top Packet Forwarding Engine.
- **B**—Bottom Packet Forwarding Engine.
- **SG(ASIC#, port#, HSL2_bit)**—SG ASIC information (ASIC 0-3, port 0-3, HSL2_bit 0-7).
- **FB_[A-D]**—via fiber bundle A, B, C or D.
- **VCSEL#**—VCSEL module number on SIB.
- **fiber**—Fiber channel number.

The following is a sample output with description of the fields displayed in the output for In-Links:

In-Links :

=====

FPC0_T_SG(0,0,0)_FB_D(04,11) -> SF_10_00_FB_D(01,11) OK 0 Up

- **FPC0**—FPC 0.
- **T**—Top Packet Forwarding Engine.
- **SG (0, 0, 0)**—SG ASIC with port number 0 and link 0.
- **FB_D (04,11)**—Fiber Bundle D with VCSEL 4, channel 11.
- **SF_10**—Indicates F1 mode chip number 0 and Rx path.
- **SF_10_00_FB_D(01,11)** —Indicates F1 mode chip number 0 and Rx path with port 0, fiber bundle D, with VCSEL 1, channel 11.

Table 55: show chassis fabric topology Output Fields (*continued*)

Out-links and In-links (TX Matrix Plus router with 3D SIBs only)	State of the links from the F13 SIB to the SFC/LCC or vice-versa. Out-Links indicate Tx links. In-Links indicate an Rx link. The following additional fields are displayed for each SIB:			
	<ul style="list-style-type: none"> Description of the fields displayed in the output for In-links and Out-links for SFC: 			
	In-links	State	Out-links	State
	CXP0_Evn->F13_SIB0_XF2,04_0	Up	F13_SIB0_XF2,04_0->CXP0_Evn	Up

- CXP0_Evn**—CXP optics with type of port bits such as even or odd. In this case, it indicates CXP optics with even port bit number 0.
- F13_SIB0**—Name of the SFC data plane SIB with the SIB number. In this case, it indicates F13 SIB with number 0.
- XF2,04_0**—Name of the ASIC with port and subchannel number. In this case, it Indicates XF2 chip with port number 4 and subchannel number 0.

- Description of the fields displayed in the output for In-links and Out-links for LCC:

	In-links	State	Out-links	
State				
CXP0_Evn->LCC_SIB0_XF3,10_0	Up	LCC_SIB0_XF3,10_0->CXP0_Evn	Up	

- CXP0_Evn**—CXP optics with the type of port bits such as even or odd. In this case, it indicates CXP optics with even port bit number 0.
- LCC_SIB0**—LCC SIB number. In this case, it indicates LCC SIB with number 0.
- XF3,10_0**—Name of the ASIC with port and subchannel number. In this case, it Indicates XF3 with port number 10 and subchannel number 0.

Sample Output

show chassis fabric topology scc (TX Matrix Router)

```
user@host> show chassis fabric topology scc
scc-re1:
```

```
fchip (mode)
```

```
in-links      state  out-links      state
```

```
Sib #0 :
```

```
SIB0_F0 (F2 ):
```

```
LCC0_SIB-L0_F0,03->SIB-S0_F0,00  UP      SIB-S0_F0,00->LCC0_SIB-L0_F1,00  UP
LCC1_SIB-L0_F0,03->SIB-S0_F0,01  UP      SIB-S0_F0,01->LCC1_SIB-L0_F1,08  UP
LCC2_SIB-L0_F0,03->SIB-S0_F0,02  RESET   SIB-S0_F0,02->LCC2_SIB-L0_F1,08  UP
LCC3_SIB-L0_F0,03->SIB-S0_F0,03  RESET   SIB-S0_F0,03->LCC3_SIB-L0_F1,00  UP
LCC0_SIB-L0_F0,02->SIB-S0_F0,04  UP      SIB-S0_F0,04->LCC0_SIB-L0_F1,01  UP
LCC1_SIB-L0_F0,02->SIB-S0_F0,05  UP      SIB-S0_F0,05->LCC1_SIB-L0_F1,09  UP
LCC2_SIB-L0_F0,02->SIB-S0_F0,06  RESET   SIB-S0_F0,06->LCC2_SIB-L0_F1,09  UP
LCC3_SIB-L0_F0,02->SIB-S0_F0,07  RESET   SIB-S0_F0,07->LCC3_SIB-L0_F1,01  UP
```

LCC0_SIB-L0_F0,07->SIB-S0_F0,08	UP	SIB-S0_F0,08->LCC0_SIB-L0_F1,04	UP
LCC1_SIB-L0_F0,07->SIB-S0_F0,09	UP	SIB-S0_F0,09->LCC1_SIB-L0_F1,12	UP
LCC2_SIB-L0_F0,07->SIB-S0_F0,10	RESET	SIB-S0_F0,10->LCC2_SIB-L0_F1,12	UP
LCC3_SIB-L0_F0,07->SIB-S0_F0,11	RESET	SIB-S0_F0,11->LCC3_SIB-L0_F1,04	UP
LCC0_SIB-L0_F0,06->SIB-S0_F0,12	UP	SIB-S0_F0,12->LCC0_SIB-L0_F1,05	UP
LCC1_SIB-L0_F0,06->SIB-S0_F0,13	UP	SIB-S0_F0,13->LCC1_SIB-L0_F1,13	UP
LCC2_SIB-L0_F0,06->SIB-S0_F0,14	RESET	SIB-S0_F0,14->LCC2_SIB-L0_F1,13	UP
LCC3_SIB-L0_F0,06->SIB-S0_F0,15	RESET	SIB-S0_F0,15->LCC3_SIB-L0_F1,05	UP
SIB0_F1 (F2):			
LCC0_SIB-L0_F0,11->SIB-S0_F1,00	UP	SIB-S0_F1,00->LCC0_SIB-L0_F1,08	UP
LCC1_SIB-L0_F0,11->SIB-S0_F1,01	UP	SIB-S0_F1,01->LCC1_SIB-L0_F1,00	UP
LCC2_SIB-L0_F0,11->SIB-S0_F1,02	RESET	SIB-S0_F1,02->LCC2_SIB-L0_F1,00	UP
LCC3_SIB-L0_F0,11->SIB-S0_F1,03	RESET	SIB-S0_F1,03->LCC3_SIB-L0_F1,08	UP
LCC0_SIB-L0_F0,10->SIB-S0_F1,04	UP	SIB-S0_F1,04->LCC0_SIB-L0_F1,09	UP
LCC1_SIB-L0_F0,10->SIB-S0_F1,05	UP	SIB-S0_F1,05->LCC1_SIB-L0_F1,01	UP
LCC2_SIB-L0_F0,10->SIB-S0_F1,06	RESET	SIB-S0_F1,06->LCC2_SIB-L0_F1,01	UP
LCC3_SIB-L0_F0,10->SIB-S0_F1,07	RESET	SIB-S0_F1,07->LCC3_SIB-L0_F1,09	UP
LCC0_SIB-L0_F0,15->SIB-S0_F1,08	UP	SIB-S0_F1,08->LCC0_SIB-L0_F1,12	UP
LCC1_SIB-L0_F0,15->SIB-S0_F1,09	UP	SIB-S0_F1,09->LCC1_SIB-L0_F1,04	UP
LCC2_SIB-L0_F0,15->SIB-S0_F1,10	RESET	SIB-S0_F1,10->LCC2_SIB-L0_F1,04	UP
LCC3_SIB-L0_F0,15->SIB-S0_F1,11	RESET	SIB-S0_F1,11->LCC3_SIB-L0_F1,12	UP
LCC0_SIB-L0_F0,14->SIB-S0_F1,12	UP	SIB-S0_F1,12->LCC0_SIB-L0_F1,13	UP
LCC1_SIB-L0_F0,14->SIB-S0_F1,13	UP	SIB-S0_F1,13->LCC1_SIB-L0_F1,05	UP
LCC2_SIB-L0_F0,14->SIB-S0_F1,14	RESET	SIB-S0_F1,14->LCC2_SIB-L0_F1,05	
UP			
LCC3_SIB-L0_F0,14->SIB-S0_F1,15	RESET	SIB-S0_F1,15->LCC3_SIB-L0_F1,13	
UP			
SIB0_F2 (F2):			
LCC3_SIB-L0_F0,13->SIB-S0_F2,00	RESET	SIB-S0_F2,00->LCC3_SIB-L0_F1,14	UP
LCC2_SIB-L0_F0,13->SIB-S0_F2,01	RESET	SIB-S0_F2,01->LCC2_SIB-L0_F1,06	
UP			
LCC1_SIB-L0_F0,13->SIB-S0_F2,02	UP	SIB-S0_F2,02->LCC1_SIB-L0_F1,06	UP
LCC0_SIB-L0_F0,13->SIB-S0_F2,03	UP	SIB-S0_F2,03->LCC0_SIB-L0_F1,14	UP
LCC3_SIB-L0_F0,12->SIB-S0_F2,04	RESET	SIB-S0_F2,04->LCC3_SIB-L0_F1,15	
UP			
LCC2_SIB-L0_F0,12->SIB-S0_F2,05	RESET	SIB-S0_F2,05->LCC2_SIB-L0_F1,07	UP
LCC1_SIB-L0_F0,12->SIB-S0_F2,06	UP	SIB-S0_F2,06->LCC1_SIB-L0_F1,07	UP
LCC0_SIB-L0_F0,12->SIB-S0_F2,07	UP	SIB-S0_F2,07->LCC0_SIB-L0_F1,15	UP
LCC3_SIB-L0_F0,09->SIB-S0_F2,08	RESET	SIB-S0_F2,08->LCC3_SIB-L0_F1,10	
UP			
LCC2_SIB-L0_F0,09->SIB-S0_F2,09	RESET	SIB-S0_F2,09->LCC2_SIB-L0_F1,02	
UP			
LCC1_SIB-L0_F0,09->SIB-S0_F2,10	UP	SIB-S0_F2,10->LCC1_SIB-L0_F1,02	UP
LCC0_SIB-L0_F0,09->SIB-S0_F2,11	UP	SIB-S0_F2,11->LCC0_SIB-L0_F1,10	UP
LCC3_SIB-L0_F0,08->SIB-S0_F2,12	RESET	SIB-S0_F2,12->LCC3_SIB-L0_F1,11	
UP			
LCC2_SIB-L0_F0,08->SIB-S0_F2,13	RESET	SIB-S0_F2,13->LCC2_SIB-L0_F1,03	
UP			
LCC1_SIB-L0_F0,08->SIB-S0_F2,14	UP	SIB-S0_F2,14->LCC1_SIB-L0_F1,03	UP
LCC0_SIB-L0_F0,08->SIB-S0_F2,15	UP	SIB-S0_F2,15->LCC0_SIB-L0_F1,11	UP
SIB0_F3 (F2):			
LCC3_SIB-L0_F0,05->SIB-S0_F3,00	RESET	SIB-S0_F3,00->LCC3_SIB-L0_F1,06	
UP			
LCC2_SIB-L0_F0,05->SIB-S0_F3,01	RESET	SIB-S0_F3,01->LCC2_SIB-L0_F1,14	
UP			
LCC1_SIB-L0_F0,05->SIB-S0_F3,02	UP	SIB-S0_F3,02->LCC1_SIB-L0_F1,14	UP
LCC0_SIB-L0_F0,05->SIB-S0_F3,03	UP	SIB-S0_F3,03->LCC0_SIB-L0_F1,06	UP
LCC3_SIB-L0_F0,04->SIB-S0_F3,04	RESET	SIB-S0_F3,04->LCC3_SIB-L0_F1,07	
UP			
LCC2_SIB-L0_F0,04->SIB-S0_F3,05	RESET	SIB-S0_F3,05->LCC2_SIB-L0_F1,15	
UP			

```

LCC1_SIB-L0_F0,04->SIB-S0_F3,06  UP
LCC0_SIB-L0_F0,04->SIB-S0_F3,07  UP
LCC3_SIB-L0_F0,01->SIB-S0_F3,08  RESET
UP
LCC2_SIB-L0_F0,01->SIB-S0_F3,09  RESET
UP
LCC1_SIB-L0_F0,01->SIB-S0_F3,10  UP
LCC0_SIB-L0_F0,01->SIB-S0_F3,11  UP
LCC3_SIB-L0_F0,00->SIB-S0_F3,12  RESET
UP
LCC2_SIB-L0_F0,00->SIB-S0_F3,13  RESET
UP
LCC1_SIB-L0_F0,00->SIB-S0_F3,14  UP
LCC0_SIB-L0_F0,00->SIB-S0_F3,15  UP
Sib #1 :
-----
SIB1_F0 (F2 ):
LCC0_SIB-L1_F0,03->SIB-S1_F0,00  RESET
LCC1_SIB-L1_F0,03->SIB-S1_F0,01  RESET
LCC2_SIB-L1_F0,03->SIB-S1_F0,02  RESET
LCC3_SIB-L1_F0,03->SIB-S1_F0,03  RESET
LCC0_SIB-L1_F0,02->SIB-S1_F0,04  RESET
LCC1_SIB-L1_F0,02->SIB-S1_F0,05  RESET
LCC2_SIB-L1_F0,02->SIB-S1_F0,06  RESET
LCC3_SIB-L1_F0,02->SIB-S1_F0,07  RESET
LCC0_SIB-L1_F0,07->SIB-S1_F0,08  RESET
LCC1_SIB-L1_F0,07->SIB-S1_F0,09  RESET
LCC2_SIB-L1_F0,07->SIB-S1_F0,10  RESET
LCC3_SIB-L1_F0,07->SIB-S1_F0,11  RESET
LCC0_SIB-L1_F0,06->SIB-S1_F0,12  RESET
LCC1_SIB-L1_F0,06->SIB-S1_F0,13  RESET
LCC2_SIB-L1_F0,06->SIB-S1_F0,14  RESET
LCC3_SIB-L1_F0,06->SIB-S1_F0,15  RESET
SIB1_F1 (F2 ):
LCC0_SIB-L1_F0,11->SIB-S1_F1,00  RESET
LCC1_SIB-L1_F0,11->SIB-S1_F1,01  RESET
LCC2_SIB-L1_F0,11->SIB-S1_F1,02  RESET
LCC3_SIB-L1_F0,11->SIB-S1_F1,03  RESET
LCC0_SIB-L1_F0,10->SIB-S1_F1,04  RESET
LCC1_SIB-L1_F0,10->SIB-S1_F1,05  RESET
LCC2_SIB-L1_F0,10->SIB-S1_F1,06  RESET
LCC3_SIB-L1_F0,10->SIB-S1_F1,07  RESET
LCC0_SIB-L1_F0,15->SIB-S1_F1,08  RESET
LCC1_SIB-L1_F0,15->SIB-S1_F1,09  RESET
LCC2_SIB-L1_F0,15->SIB-S1_F1,10  RESET
LCC3_SIB-L1_F0,15->SIB-S1_F1,11  RESET
LCC0_SIB-L1_F0,14->SIB-S1_F1,12  RESET
LCC1_SIB-L1_F0,14->SIB-S1_F1,13  RESET
LCC2_SIB-L1_F0,14->SIB-S1_F1,14  RESET

SIB-S0_F3,06->LCC1_SIB-L0_F1,15  UP
SIB-S0_F3,07->LCC0_SIB-L0_F1,07  UP
SIB-S0_F3,08->LCC3_SIB-L0_F1,02
SIB-S0_F3,09->LCC2_SIB-L0_F1,10
SIB-S0_F3,10->LCC1_SIB-L0_F1,10  UP
SIB-S0_F3,11->LCC0_SIB-L0_F1,02  UP
SIB-S0_F3,12->LCC3_SIB-L0_F1,03
SIB-S0_F3,13->LCC2_SIB-L0_F1,11
SIB-S0_F3,14->LCC1_SIB-L0_F1,11  UP
SIB-S0_F3,15->LCC0_SIB-L0_F1,03  UP

SIB-S1_F0,00->LCC0_SIB-L1_F1,00  UP
SIB-S1_F0,01->LCC1_SIB-L1_F1,08  UP
SIB-S1_F0,02->LCC2_SIB-L1_F1,08  UP
SIB-S1_F0,03->LCC3_SIB-L1_F1,00  UP
SIB-S1_F0,04->LCC0_SIB-L1_F1,01  UP
SIB-S1_F0,05->LCC1_SIB-L1_F1,09  UP
SIB-S1_F0,06->LCC2_SIB-L1_F1,09  UP
SIB-S1_F0,07->LCC3_SIB-L1_F1,01  UP
SIB-S1_F0,08->LCC0_SIB-L1_F1,04  UP
SIB-S1_F0,09->LCC1_SIB-L1_F1,12  UP
SIB-S1_F0,10->LCC2_SIB-L1_F1,12  UP
SIB-S1_F0,11->LCC3_SIB-L1_F1,04  UP
SIB-S1_F0,12->LCC0_SIB-L1_F1,05  UP
SIB-S1_F0,13->LCC1_SIB-L1_F1,13  UP
SIB-S1_F0,14->LCC2_SIB-L1_F1,13  UP
SIB-S1_F0,15->LCC3_SIB-L1_F1,05  UP

SIB-S1_F1,00->LCC0_SIB-L1_F1,08  UP
SIB-S1_F1,01->LCC1_SIB-L1_F1,00  UP
SIB-S1_F1,02->LCC2_SIB-L1_F1,00  UP
SIB-S1_F1,03->LCC3_SIB-L1_F1,08  UP
SIB-S1_F1,04->LCC0_SIB-L1_F1,09  UP
SIB-S1_F1,05->LCC1_SIB-L1_F1,01  UP
SIB-S1_F1,06->LCC2_SIB-L1_F1,01  UP
SIB-S1_F1,07->LCC3_SIB-L1_F1,09  UP
SIB-S1_F1,08->LCC0_SIB-L1_F1,12  UP
SIB-S1_F1,09->LCC1_SIB-L1_F1,04  UP
SIB-S1_F1,10->LCC2_SIB-L1_F1,04  UP
SIB-S1_F1,11->LCC3_SIB-L1_F1,12,05  UP
SIB-S1_F1,12->LCC0_SIB-L1_F1,13  UP
SIB-S1_F1,13->LCC1_SIB-L1_F1,05  UP
SIB-S1_F1,14->LCC2_SIB-L1_F1,05  UP

```

show chassis fabric topology lcc

```

user@host> show chassis fabric topology lcc 0
lcc0-re0:

```

```

-----
      fchip (mode)
in-links          state      out-links          state
-----
Sib #2 :
-----
SIB2_F0 (F1 ):

```

```

FPC0_T->SIB-L2_F0,00    DOWN    SIB-L2_F0,00->SIB-S2_F3,15 DOWN
FPC0_B->SIB-L2_F0,01    UP       SIB-L2_F0,01->SIB-S2_F3,11 DOWN
FPC1_T->SIB-L2_F0,02    DOWN     SIB-L2_F0,02->SIB-S2_F0,04 DOWN
FPC1_B->SIB-L2_F0,03    DOWN     SIB-L2_F0,03->SIB-S2_F0,00 DOWN
FPC2_T->SIB-L2_F0,04    DOWN     SIB-L2_F0,04->SIB-S2_F3,07 DOWN
FPC2_B->SIB-L2_F0,05    DOWN     SIB-L2_F0,05->SIB-S2_F3,03 DOWN
FPC3_T->SIB-L2_F0,06    DOWN     SIB-L2_F0,06->SIB-S2_F0,12 DOWN
FPC3_B->SIB-L2_F0,07    DOWN     SIB-L2_F0,07->SIB-S2_F0,08 DOWN
FPC4_T->SIB-L2_F0,08    DOWN     SIB-L2_F0,08->SIB-S2_F2,15 DOWN
FPC4_B->SIB-L2_F0,09    DOWN     SIB-L2_F0,09->SIB-S2_F2,11 DOWN
FPC5_T->SIB-L2_F0,10    DOWN     SIB-L2_F0,10->SIB-S2_F1,04 DOWN
FPC5_B->SIB-L2_F0,11    DOWN     SIB-L2_F0,11->SIB-S2_F1,00 DOWN
FPC6_T->SIB-L2_F0,12    DOWN     SIB-L2_F0,12->SIB-S2_F2,07 DOWN
FPC6_B->SIB-L2_F0,13    UP       SIB-L2_F0,13->SIB-S2_F2,03 DOWN
FPC7_T->SIB-L2_F0,14    DOWN     SIB-L2_F0,14->SIB-S2_F1,12 DOWN
FPC7_B->SIB-L2_F0,15    DOWN     SIB-L2_F0,15->SIB-S2_F1,08 DOWN

SIB2_F1 (F3 ):
SIB-S2_F0,00->SIB-L2_F1,00 UP    SIB-L2_F1,00->FPC7_B    DOWN
SIB-S2_F0,04->SIB-L2_F1,01 UP    SIB-L2_F1,01->FPC7_T    DOWN
SIB-S2_F3,11->SIB-L2_F1,02 UP    SIB-L2_F1,02->FPC6_B    DOWN
SIB-S2_F3,15->SIB-L2_F1,03 UP    SIB-L2_F1,03->FPC6_T    DOWN
SIB-S2_F0,08->SIB-L2_F1,04 UP    SIB-L2_F1,04->FPC5_B    DOWN
SIB-S2_F0,12->SIB-L2_F1,05 UP    SIB-L2_F1,05->FPC5_T    DOWN
SIB-S2_F3,03->SIB-L2_F1,06 UP    SIB-L2_F1,06->FPC4_B    DOWN
SIB-S2_F3,07->SIB-L2_F1,07 UP    SIB-L2_F1,07->FPC4_T    DOWN
SIB-S2_F1,00->SIB-L2_F1,08 UP    SIB-L2_F1,08->FPC3_B    DOWN
SIB-S2_F1,04->SIB-L2_F1,09 UP    SIB-L2_F1,09->FPC3_T    DOWN
SIB-S2_F2,11->SIB-L2_F1,10 UP    SIB-L2_F1,10->FPC2_B    DOWN
SIB-S2_F2,15->SIB-L2_F1,11 UP    SIB-L2_F1,11->FPC2_T    DOWN
SIB-S2_F1,08->SIB-L2_F1,12 UP    SIB-L2_F1,12->FPC1_B    DOWN
SIB-S2_F1,12->SIB-L2_F1,13 UP    SIB-L2_F1,13->FPC1_T    DOWN
SIB-S2_F2,03->SIB-L2_F1,14 UP    SIB-L2_F1,14->FPC0_B    DOWN
SIB-S2_F2,07->SIB-L2_F1,15 UP    SIB-L2_F1,15->FPC0_T    DOWN

Sib #4 :
-----
SIB4_F0 (F1 ):
FPC0_T->SIB-L4_F0,00    RESET    SIB-L4_F0,00->SIB-S4_F3,15 UP
FPC0_B->SIB-L4_F0,01    UP       SIB-L4_F0,01->SIB-S4_F3,11 UP
FPC1_T->SIB-L4_F0,02    RESET    SIB-L4_F0,02->SIB-S4_F0,04 UP
FPC1_B->SIB-L4_F0,03    RESET    SIB-L4_F0,03->SIB-S4_F0,00 UP
FPC2_T->SIB-L4_F0,04    RESET    SIB-L4_F0,04->SIB-S4_F3,07 UP
FPC2_B->SIB-L4_F0,05    RESET    SIB-L4_F0,05->SIB-S4_F3,03 UP
FPC3_T->SIB-L4_F0,06    RESET    SIB-L4_F0,06->SIB-S4_F0,12 UP
FPC3_B->SIB-L4_F0,07    RESET    SIB-L4_F0,07->SIB-S4_F0,08 UP
FPC4_T->SIB-L4_F0,08    RESET    SIB-L4_F0,08->SIB-S4_F2,15 UP
FPC4_B->SIB-L4_F0,09    RESET    SIB-L4_F0,09->SIB-S4_F2,11 UP
FPC5_T->SIB-L4_F0,10    RESET    SIB-L4_F0,10->SIB-S4_F1,04 UP
FPC5_B->SIB-L4_F0,11    RESET    SIB-L4_F0,11->SIB-S4_F1,00 UP
FPC6_T->SIB-L4_F0,12    RESET    SIB-L4_F0,12->SIB-S4_F2,07 UP
FPC6_B->SIB-L4_F0,13    UP       SIB-L4_F0,13->SIB-S4_F2,03 UP
FPC7_T->SIB-L4_F0,14    RESET    SIB-L4_F0,14->SIB-S4_F1,12 UP
FPC7_B->SIB-L4_F0,15    RESET    SIB-L4_F0,15->SIB-S4_F1,08 UP

SIB4_F1 (F3 ):
SIB-S4_F0,00->SIB-L4_F1,00 UP    SIB-L4_F1,00->FPC7_B    UP
SIB-S4_F0,04->SIB-L4_F1,01 UP    SIB-L4_F1,01->FPC7_T    UP
SIB-S4_F3,11->SIB-L4_F1,02 UP    SIB-L4_F1,02->FPC6_B    UP
SIB-S4_F3,15->SIB-L4_F1,03 UP    SIB-L4_F1,03->FPC6_T    UP
SIB-S4_F0,08->SIB-L4_F1,04 UP    SIB-L4_F1,04->FPC5_B    UP
SIB-S4_F0,12->SIB-L4_F1,05 UP    SIB-L4_F1,05->FPC5_T    UP
SIB-S4_F3,03->SIB-L4_F1,06 UP    SIB-L4_F1,06->FPC4_B    UP
SIB-S4_F3,07->SIB-L4_F1,07 UP    SIB-L4_F1,07->FPC4_T    UP

```

```

SIB-S4_F1,00->SIB-L4_F1,08 UP      SIB-L4_F1,08->FPC3_B    UP
SIB-S4_F1,04->SIB-L4_F1,09 UP      SIB-L4_F1,09->FPC3_T    UP
SIB-S4_F2,11->SIB-L4_F1,10 UP      SIB-L4_F1,10->FPC2_B    UP
SIB-S4_F2,15->SIB-L4_F1,11 UP      SIB-L4_F1,11->FPC2_T    UP
SIB-S4_F1,08->SIB-L4_F1,12 UP      SIB-L4_F1,12->FPC1_B    UP
SIB-S4_F1,12->SIB-L4_F1,13 UP      SIB-L4_F1,13->FPC1_T    UP
SIB-S4_F2,03->SIB-L4_F1,14 UP      SIB-L4_F1,14->FPC0_B    UP
SIB-S4_F2,07->SIB-L4_F1,15 UP      SIB-L4_F1,15->FPC0_T    UP

```

show chassis fabric topology (TX Matrix Plus Router)

```

user@host> show chassis fabric topology
sfc0-re0:

```

```

F13_SIB0
=====

```

```

Out-Links:
=====

```

SFC0_F13_SIB_00	-> LCC00_ST_SIB_L00	VCSEL Status	HSL2 Channel	HSL2 Status
SF_30_00_FB_D(04,11)	-> FPC0_T_SG(0,0,0)_FB_D(01,11)	OK	112	Up
SF_30_00_FB_D(04,10)	-> FPC0_T_SG(0,0,1)_FB_D(01,10)	OK	112	Up
SF_30_00_FB_D(04,09)	-> FPC0_T_SG(0,0,2)_FB_D(01,09)	OK	112	Up
SF_30_00_FB_D(04,08)	-> FPC0_T_SG(0,0,3)_FB_D(01,08)	OK	112	Up
SF_30_00_FB_D(04,07)	-> FPC0_T_SG(0,0,4)_FB_D(01,07)	OK	112	Up
SF_30_00_FB_D(04,06)	-> FPC0_T_SG(0,0,5)_FB_D(01,06)	OK	112	Up
SF_30_00_FB_D(04,05)	-> FPC0_T_SG(0,0,6)_FB_D(01,05)	OK	112	Up
SF_30_00_FB_D(04,04)	-> FPC0_T_SG(0,0,7)_FB_D(01,04)	OK	112	Up
SF_30_01_FB_B(16,11)	-> FPC4_T_SG(2,0,0)_FB_B(13,11)	OK	119	Up
SF_30_01_FB_B(16,10)	-> FPC4_T_SG(2,0,1)_FB_B(13,10)	OK	119	Up
SF_30_01_FB_B(16,09)	-> FPC4_T_SG(2,0,2)_FB_B(13,09)	OK	119	Up
SF_30_01_FB_B(16,08)	-> FPC4_T_SG(2,0,3)_FB_B(13,08)	OK	119	Up
SF_30_01_FB_B(16,07)	-> FPC4_T_SG(2,0,4)_FB_B(13,07)	OK	119	Up
SF_30_01_FB_B(16,06)	-> FPC4_T_SG(2,0,5)_FB_B(13,06)	OK	119	Up
SF_30_01_FB_B(16,05)	-> FPC4_T_SG(2,0,6)_FB_B(13,05)	OK	119	Up
SF_30_01_FB_B(16,04)	-> FPC4_T_SG(2,0,7)_FB_B(13,04)	OK	119	Up
SF_30_02_FB_D(05,08)	-> FPC1_T_SG(0,2,0)_FB_D(02,08)	OK	126	Up
SF_30_02_FB_D(05,07)	-> FPC1_T_SG(0,2,1)_FB_D(02,07)	OK	126	Up
SF_30_02_FB_D(05,06)	-> FPC1_T_SG(0,2,2)_FB_D(02,06)	OK	126	Up
SF_30_02_FB_D(05,05)	-> FPC1_T_SG(0,2,3)_FB_D(02,05)	OK	126	Up
SF_30_02_FB_D(05,03)	-> FPC1_T_SG(0,2,4)_FB_D(02,03)	OK	126	Up
SF_30_02_FB_D(05,02)	-> FPC1_T_SG(0,2,5)_FB_D(02,02)	OK	126	Up
SF_30_02_FB_D(05,01)	-> FPC1_T_SG(0,2,6)_FB_D(02,01)	OK	126	Up
SF_30_02_FB_D(05,00)	-> FPC1_T_SG(0,2,7)_FB_D(02,00)	OK	126	Up
SF_30_03_FB_B(17,08)	-> FPC5_T_SG(2,2,0)_FB_B(14,08)	OK	133	Up
SF_30_03_FB_B(17,07)	-> FPC5_T_SG(2,2,1)_FB_B(14,07)	OK	133	Up
SF_30_03_FB_B(17,06)	-> FPC5_T_SG(2,2,2)_FB_B(14,06)	OK	133	Up
SF_30_03_FB_B(17,05)	-> FPC5_T_SG(2,2,3)_FB_B(14,05)	OK	133	Up
SF_30_03_FB_B(17,03)	-> FPC5_T_SG(2,2,4)_FB_B(14,03)	OK	133	Up
SF_30_03_FB_B(17,02)	-> FPC5_T_SG(2,2,5)_FB_B(14,02)	OK	133	Up
SF_30_03_FB_B(17,01)	-> FPC5_T_SG(2,2,6)_FB_B(14,01)	OK	133	Up
SF_30_03_FB_B(17,00)	-> FPC5_T_SG(2,2,7)_FB_B(14,00)	OK	133	Up
SF_30_04_FB_C(10,11)	-> FPC2_T_SG(1,0,0)_FB_C(07,11)	OK	140	Up
SF_30_04_FB_C(10,10)	-> FPC2_T_SG(1,0,1)_FB_C(07,10)	OK	140	Up
SF_30_04_FB_C(10,09)	-> FPC2_T_SG(1,0,2)_FB_C(07,09)	OK	140	Up
SF_30_04_FB_C(10,08)	-> FPC2_T_SG(1,0,3)_FB_C(07,08)	OK	140	Up
SF_30_04_FB_C(10,07)	-> FPC2_T_SG(1,0,4)_FB_C(07,07)	OK	140	Up
SF_30_04_FB_C(10,06)	-> FPC2_T_SG(1,0,5)_FB_C(07,06)	OK	140	Up


```

SF_30_04_FB_C(10,05) -> FPC2_T_SG(1,0,6)_FB_C(07,05)    OK      140    Up
SF_30_04_FB_C(10,04) -> FPC2_T_SG(1,0,7)_FB_C(07,04)    OK      140    Up
SF_30_05_FB_A(22,11) -> FPC6_T_SG(3,0,0)_FB_A(19,11)    OK      147    Up
SF_30_05_FB_A(22,10) -> FPC6_T_SG(3,0,1)_FB_A(19,10)    OK      147    Up
SF_30_05_FB_A(22,09) -> FPC6_T_SG(3,0,2)_FB_A(19,09)    OK      147    Up
SF_30_05_FB_A(22,08) -> FPC6_T_SG(3,0,3)_FB_A(19,08)    OK      147    Up
SF_30_05_FB_A(22,07) -> FPC6_T_SG(3,0,4)_FB_A(19,07)    OK      147    Up
SF_30_05_FB_A(22,06) -> FPC6_T_SG(3,0,5)_FB_A(19,06)    OK      147    Up
SF_30_05_FB_A(22,05) -> FPC6_T_SG(3,0,6)_FB_A(19,05)    OK      147    Up
SF_30_05_FB_A(22,04) -> FPC6_T_SG(3,0,7)_FB_A(19,04)    OK      147    Up
SF_30_06_FB_C(11,08) -> FPC3_T_SG(1,2,0)_FB_C(08,08)    OK      154    Up
SF_30_06_FB_C(11,07) -> FPC3_T_SG(1,2,1)_FB_C(08,07)    OK      154    Up
SF_30_06_FB_C(11,06) -> FPC3_T_SG(1,2,2)_FB_C(08,06)    OK      154    Up
SF_30_06_FB_C(11,05) -> FPC3_T_SG(1,2,3)_FB_C(08,05)    OK      154    Up
SF_30_06_FB_C(11,03) -> FPC3_T_SG(1,2,4)_FB_C(08,03)    OK      154    Up
SF_30_06_FB_C(11,02) -> FPC3_T_SG(1,2,5)_FB_C(08,02)    OK      154    Up
SF_30_06_FB_C(11,01) -> FPC3_T_SG(1,2,6)_FB_C(08,01)    OK      154    Up
SF_30_06_FB_C(11,00) -> FPC3_T_SG(1,2,7)_FB_C(08,00)    OK      154    Up
...

```

show chassis fabric topology sfc (TX Matrix Plus Router)

```

user@host> show chassis fabric topology sfc 0
sfc0-re0:

```

```

-----
F13_SIB0
=====

```

```

Out-Links:
=====

```

SFC0_F13_SIB_00	-> LCC00_ST_SIB_L00	VCSEL Status	HSL2 Channel	HSL2 Status
=====				
SF_30_00_FB_D(04,11)	-> FPC0_T_SG(0,0,0)_FB_D(01,11)	OK	112	Up
SF_30_00_FB_D(04,10)	-> FPC0_T_SG(0,0,1)_FB_D(01,10)	OK	112	Up
SF_30_00_FB_D(04,09)	-> FPC0_T_SG(0,0,2)_FB_D(01,09)	OK	112	Up
SF_30_00_FB_D(04,08)	-> FPC0_T_SG(0,0,3)_FB_D(01,08)	OK	112	Up
SF_30_00_FB_D(04,07)	-> FPC0_T_SG(0,0,4)_FB_D(01,07)	OK	112	Up
SF_30_00_FB_D(04,06)	-> FPC0_T_SG(0,0,5)_FB_D(01,06)	OK	112	Up
SF_30_00_FB_D(04,05)	-> FPC0_T_SG(0,0,6)_FB_D(01,05)	OK	112	Up
SF_30_00_FB_D(04,04)	-> FPC0_T_SG(0,0,7)_FB_D(01,04)	OK	112	Up
SF_30_01_FB_B(16,11)	-> FPC4_T_SG(2,0,0)_FB_B(13,11)	OK	119	Up
SF_30_01_FB_B(16,10)	-> FPC4_T_SG(2,0,1)_FB_B(13,10)	OK	119	Up
SF_30_01_FB_B(16,09)	-> FPC4_T_SG(2,0,2)_FB_B(13,09)	OK	119	Up
SF_30_01_FB_B(16,08)	-> FPC4_T_SG(2,0,3)_FB_B(13,08)	OK	119	Up
SF_30_01_FB_B(16,07)	-> FPC4_T_SG(2,0,4)_FB_B(13,07)	OK	119	Up
SF_30_01_FB_B(16,06)	-> FPC4_T_SG(2,0,5)_FB_B(13,06)	OK	119	Up
SF_30_01_FB_B(16,05)	-> FPC4_T_SG(2,0,6)_FB_B(13,05)	OK	119	Up
SF_30_01_FB_B(16,04)	-> FPC4_T_SG(2,0,7)_FB_B(13,04)	OK	119	Up
SF_30_02_FB_D(05,08)	-> FPC1_T_SG(0,2,0)_FB_D(02,08)	OK	126	Up
SF_30_02_FB_D(05,07)	-> FPC1_T_SG(0,2,1)_FB_D(02,07)	OK	126	Up
SF_30_02_FB_D(05,06)	-> FPC1_T_SG(0,2,2)_FB_D(02,06)	OK	126	Up
SF_30_02_FB_D(05,05)	-> FPC1_T_SG(0,2,3)_FB_D(02,05)	OK	126	Up
SF_30_02_FB_D(05,03)	-> FPC1_T_SG(0,2,4)_FB_D(02,03)	OK	126	Up
SF_30_02_FB_D(05,02)	-> FPC1_T_SG(0,2,5)_FB_D(02,02)	OK	126	Up
SF_30_02_FB_D(05,01)	-> FPC1_T_SG(0,2,6)_FB_D(02,01)	OK	126	Up
SF_30_02_FB_D(05,00)	-> FPC1_T_SG(0,2,7)_FB_D(02,00)	OK	126	Up
SF_30_03_FB_B(17,08)	-> FPC5_T_SG(2,2,0)_FB_B(14,08)	OK	133	Up
SF_30_03_FB_B(17,07)	-> FPC5_T_SG(2,2,1)_FB_B(14,07)	OK	133	Up
SF_30_03_FB_B(17,06)	-> FPC5_T_SG(2,2,2)_FB_B(14,06)	OK	133	Up

```

SF_30_03_FB_B(17,05) -> FPC5_T_SG(2,2,3)_FB_B(14,05)    OK      133    Up
SF_30_03_FB_B(17,03) -> FPC5_T_SG(2,2,4)_FB_B(14,03)    OK      133    Up
SF_30_03_FB_B(17,02) -> FPC5_T_SG(2,2,5)_FB_B(14,02)    OK      133    Up
SF_30_03_FB_B(17,01) -> FPC5_T_SG(2,2,6)_FB_B(14,01)    OK      133    Up
SF_30_03_FB_B(17,00) -> FPC5_T_SG(2,2,7)_FB_B(14,00)    OK      133    Up
SF_30_04_FB_C(10,11) -> FPC2_T_SG(1,0,0)_FB_C(07,11)    OK      140    Up
SF_30_04_FB_C(10,10) -> FPC2_T_SG(1,0,1)_FB_C(07,10)    OK      140    Up
SF_30_04_FB_C(10,09) -> FPC2_T_SG(1,0,2)_FB_C(07,09)    OK      140    Up
SF_30_04_FB_C(10,08) -> FPC2_T_SG(1,0,3)_FB_C(07,08)    OK      140    Up
SF_30_04_FB_C(10,07) -> FPC2_T_SG(1,0,4)_FB_C(07,07)    OK      140    Up
SF_30_04_FB_C(10,06) -> FPC2_T_SG(1,0,5)_FB_C(07,06)    OK      140    Up
SF_30_04_FB_C(10,05) -> FPC2_T_SG(1,0,6)_FB_C(07,05)    OK      140    Up
SF_30_04_FB_C(10,04) -> FPC2_T_SG(1,0,7)_FB_C(07,04)    OK      140    Up
SF_30_05_FB_A(22,11) -> FPC6_T_SG(3,0,0)_FB_A(19,11)    OK      147    Up
SF_30_05_FB_A(22,10) -> FPC6_T_SG(3,0,1)_FB_A(19,10)    OK      147    Up
SF_30_05_FB_A(22,09) -> FPC6_T_SG(3,0,2)_FB_A(19,09)    OK      147    Up
SF_30_05_FB_A(22,08) -> FPC6_T_SG(3,0,3)_FB_A(19,08)    OK      147    Up
SF_30_05_FB_A(22,07) -> FPC6_T_SG(3,0,4)_FB_A(19,07)    OK      147    Up
SF_30_05_FB_A(22,06) -> FPC6_T_SG(3,0,5)_FB_A(19,06)    OK      147    Up
SF_30_05_FB_A(22,05) -> FPC6_T_SG(3,0,6)_FB_A(19,05)    OK      147    Up
SF_30_05_FB_A(22,04) -> FPC6_T_SG(3,0,7)_FB_A(19,04)    OK      147    Up
SF_30_06_FB_C(11,08) -> FPC3_T_SG(1,2,0)_FB_C(08,08)    OK      154    Up
SF_30_06_FB_C(11,07) -> FPC3_T_SG(1,2,1)_FB_C(08,07)    OK      154    Up
SF_30_06_FB_C(11,06) -> FPC3_T_SG(1,2,2)_FB_C(08,06)    OK      154    Up
SF_30_06_FB_C(11,05) -> FPC3_T_SG(1,2,3)_FB_C(08,05)    OK      154    Up
SF_30_06_FB_C(11,03) -> FPC3_T_SG(1,2,4)_FB_C(08,03)    OK      154    Up
SF_30_06_FB_C(11,02) -> FPC3_T_SG(1,2,5)_FB_C(08,02)    OK      154    Up
SF_30_06_FB_C(11,01) -> FPC3_T_SG(1,2,6)_FB_C(08,01)    OK      154    Up
SF_30_06_FB_C(11,00) -> FPC3_T_SG(1,2,7)_FB_C(08,00)    OK      154    Up
...

```

show chassis fabric topology lcc (TX Matrix Plus Router)

```

user@host> show chassis fabric topology lcc 0
lcc0-re0:

```

```

-----
SIB0

```

```

=====

```

```

Out-Links:

```

```

=====

```

LCC00_ST_SIB_L00	-> SFC0_F13_SIB_00	VCSEL Status	HSL2 Channel	HSL2 Status
=====				
FPC0_T_SG(0,0,0)_FB_D(04,11)	-> SF_10_00_FB_D(01,11)	OK	12	Up
FPC0_T_SG(0,0,1)_FB_D(04,10)	-> SF_10_00_FB_D(01,10)	OK	12	Up
FPC0_T_SG(0,0,2)_FB_D(04,09)	-> SF_10_00_FB_D(01,09)	OK	12	Up
FPC0_T_SG(0,0,3)_FB_D(04,08)	-> SF_10_00_FB_D(01,08)	OK	12	Up
FPC0_T_SG(0,0,4)_FB_D(04,07)	-> SF_10_00_FB_D(01,07)	OK	12	Up
FPC0_T_SG(0,0,5)_FB_D(04,06)	-> SF_10_00_FB_D(01,06)	OK	12	Up
FPC0_T_SG(0,0,6)_FB_D(04,05)	-> SF_10_00_FB_D(01,05)	OK	12	Up
FPC0_T_SG(0,0,7)_FB_D(04,04)	-> SF_10_00_FB_D(01,04)	OK	12	Up
FPC0_B_SG(0,1,0)_FB_D(03,07)	-> SF_10_10_FB_D(00,07)	OK	15	Up
FPC0_B_SG(0,1,1)_FB_D(03,06)	-> SF_10_10_FB_D(00,06)	OK	15	Up
FPC0_B_SG(0,1,2)_FB_D(03,05)	-> SF_10_10_FB_D(00,05)	OK	15	Up
FPC0_B_SG(0,1,3)_FB_D(03,04)	-> SF_10_10_FB_D(00,04)	OK	15	Up
FPC0_B_SG(0,1,4)_FB_D(03,03)	-> SF_10_10_FB_D(00,03)	OK	15	Up
FPC0_B_SG(0,1,5)_FB_D(03,02)	-> SF_10_10_FB_D(00,02)	OK	15	Up
FPC0_B_SG(0,1,6)_FB_D(03,01)	-> SF_10_10_FB_D(00,01)	OK	15	Up
FPC0_B_SG(0,1,7)_FB_D(03,00)	-> SF_10_10_FB_D(00,00)	OK	15	Up

```

FPC1_T_SG(0,2,0)_FB_D(05,08) -> SF_10_02_FB_D(02,08)    OK      18      Up
FPC1_T_SG(0,2,1)_FB_D(05,07) -> SF_10_02_FB_D(02,07)    OK      18      Up
FPC1_T_SG(0,2,2)_FB_D(05,06) -> SF_10_02_FB_D(02,06)    OK      18      Up
FPC1_T_SG(0,2,3)_FB_D(05,05) -> SF_10_02_FB_D(02,05)    OK      18      Up
FPC1_T_SG(0,2,4)_FB_D(05,03) -> SF_10_02_FB_D(02,03)    OK      18      Up
FPC1_T_SG(0,2,5)_FB_D(05,02) -> SF_10_02_FB_D(02,02)    OK      18      Up
FPC1_T_SG(0,2,6)_FB_D(05,01) -> SF_10_02_FB_D(02,01)    OK      18      Up
FPC1_T_SG(0,2,7)_FB_D(05,00) -> SF_10_02_FB_D(02,00)    OK      18      Up
FPC1_B_SG(0,3,0)_FB_D(04,03) -> SF_10_11_FB_D(01,03)    OK      21      Up
FPC1_B_SG(0,3,1)_FB_D(04,02) -> SF_10_11_FB_D(01,02)    OK      21      Up
FPC1_B_SG(0,3,2)_FB_D(04,01) -> SF_10_11_FB_D(01,01)    OK      21      Up
FPC1_B_SG(0,3,3)_FB_D(04,00) -> SF_10_11_FB_D(01,00)    OK      21      Up
FPC1_B_SG(0,3,4)_FB_D(03,11) -> SF_10_11_FB_D(00,11)    OK      21      Up
FPC1_B_SG(0,3,5)_FB_D(03,10) -> SF_10_11_FB_D(00,10)    OK      21      Up
FPC1_B_SG(0,3,6)_FB_D(03,09) -> SF_10_11_FB_D(00,09)    OK      21      Up
FPC1_B_SG(0,3,7)_FB_D(03,08) -> SF_10_11_FB_D(00,08)    OK      21      Up
FPC2_T_SG(1,0,0)_FB_C(10,11) -> SF_10_04_FB_C(07,11)    OK      12      Up
FPC2_T_SG(1,0,1)_FB_C(10,10) -> SF_10_04_FB_C(07,10)    OK      12      Up
FPC2_T_SG(1,0,2)_FB_C(10,09) -> SF_10_04_FB_C(07,09)    OK      12      Up
FPC2_T_SG(1,0,3)_FB_C(10,08) -> SF_10_04_FB_C(07,08)    OK      12      Up
FPC2_T_SG(1,0,4)_FB_C(10,07) -> SF_10_04_FB_C(07,07)    OK      12      Up
FPC2_T_SG(1,0,5)_FB_C(10,06) -> SF_10_04_FB_C(07,06)    OK      12      Up
FPC2_T_SG(1,0,6)_FB_C(10,05) -> SF_10_04_FB_C(07,05)    OK      12      Up
FPC2_T_SG(1,0,7)_FB_C(10,04) -> SF_10_04_FB_C(07,04)    OK      12      Up
FPC2_B_SG(1,1,0)_FB_C(09,07) -> SF_10_14_FB_C(06,07)    OK      15      Up
FPC2_B_SG(1,1,1)_FB_C(09,06) -> SF_10_14_FB_C(06,06)    OK      15      Up
FPC2_B_SG(1,1,2)_FB_C(09,05) -> SF_10_14_FB_C(06,05)    OK      15      Up
FPC2_B_SG(1,1,3)_FB_C(09,04) -> SF_10_14_FB_C(06,04)    OK      15      Up
FPC2_B_SG(1,1,4)_FB_C(09,03) -> SF_10_14_FB_C(06,03)    OK      15      Up
FPC2_B_SG(1,1,5)_FB_C(09,02) -> SF_10_14_FB_C(06,02)    OK      15      Up
FPC2_B_SG(1,1,6)_FB_C(09,01) -> SF_10_14_FB_C(06,01)    OK      15      Up
FPC2_B_SG(1,1,7)_FB_C(09,00) -> SF_10_14_FB_C(06,00)    OK      15      Up
FPC3_T_SG(1,2,0)_FB_C(11,08) -> SF_10_06_FB_C(08,08)    OK      18      Up
FPC3_T_SG(1,2,1)_FB_C(11,07) -> SF_10_06_FB_C(08,07)    OK      18      Up
FPC3_T_SG(1,2,2)_FB_C(11,06) -> SF_10_06_FB_C(08,06)    OK      18      Up
FPC3_T_SG(1,2,3)_FB_C(11,05) -> SF_10_06_FB_C(08,05)    OK      18      Up
FPC3_T_SG(1,2,4)_FB_C(11,03) -> SF_10_06_FB_C(08,03)    OK      18      Up
FPC3_T_SG(1,2,5)_FB_C(11,02) -> SF_10_06_FB_C(08,02)    OK      18      Up
FPC3_T_SG(1,2,6)_FB_C(11,01) -> SF_10_06_FB_C(08,01)    OK      18      Up
...

```

show chassis fabric topology (T4000 Core Router)

```

user@host> show chassis fabric topology 0
fchip (mode)

```

In-links	State	Out-links	State
----------	-------	-----------	-------

```

SIB0 :
-----

```

Onboard Links

SIB0_XF1,14_0->SIB0_XF,00_0	Up	SIB0_XF,00_0->SIB0_XF1,14_0	Up
SIB0_XF,00_0->SIB0_XF1,14_0	Up	SIB0_XF1,14_0->SIB0_XF,00_0	Up
SIB0_XF1,13_0->SIB0_XF,01_0	Up	SIB0_XF,01_0->SIB0_XF1,13_0	Up
SIB0_XF,01_0->SIB0_XF1,13_0	Up	SIB0_XF1,13_0->SIB0_XF,01_0	Up
SIB0_XF1,12_0->SIB0_XF,02_0	Up	SIB0_XF,02_0->SIB0_XF1,12_0	Up
SIB0_XF,02_0->SIB0_XF1,12_0	Up	SIB0_XF1,12_0->SIB0_XF,02_0	Up
SIB0_XF1,11_0->SIB0_XF,03_0	Up	SIB0_XF,03_0->SIB0_XF1,11_0	Up

SIB0_XF,03_0->SIB0_XF1,11_0	Up	SIB0_XF1,11_0->SIB0_XF,03_0	Up
SIB0_XF1,10_0->SIB0_XF,04_0	Up	SIB0_XF,04_0->SIB0_XF1,10_0	Up
SIB0_XF,04_0->SIB0_XF1,10_0	Up	SIB0_XF1,10_0->SIB0_XF,04_0	Up
SIB0_XF1,09_0->SIB0_XF,05_0	Up	SIB0_XF,05_0->SIB0_XF1,09_0	Up
SIB0_XF,05_0->SIB0_XF1,09_0	Up	SIB0_XF1,09_0->SIB0_XF,05_0	Up
SIB0_XF2,14_0->SIB0_XF,06_0	Up	SIB0_XF,06_0->SIB0_XF2,14_0	Up
SIB0_XF,06_0->SIB0_XF2,14_0	Up	SIB0_XF2,14_0->SIB0_XF,06_0	Up
SIB0_XF2,13_0->SIB0_XF,07_0	Up	SIB0_XF,07_0->SIB0_XF2,13_0	Up
SIB0_XF,07_0->SIB0_XF2,13_0	Up	SIB0_XF2,13_0->SIB0_XF,07_0	Up
SIB0_XF2,12_0->SIB0_XF,08_0	Up	SIB0_XF,08_0->SIB0_XF2,12_0	Up
SIB0_XF,08_0->SIB0_XF2,12_0	Up	SIB0_XF2,12_0->SIB0_XF,08_0	Up
SIB0_XF2,11_0->SIB0_XF,09_0	Up	SIB0_XF,09_0->SIB0_XF2,11_0	Up
SIB0_XF,09_0->SIB0_XF2,11_0	Up	SIB0_XF2,11_0->SIB0_XF,09_0	Up
SIB0_XF2,10_0->SIB0_XF,10_0	Up	SIB0_XF,10_0->SIB0_XF2,10_0	Up
SIB0_XF,10_0->SIB0_XF2,10_0	Up	SIB0_XF2,10_0->SIB0_XF,10_0	Up
SIB0_XF2,09_0->SIB0_XF,11_0	Up	SIB0_XF,11_0->SIB0_XF2,09_0	Up
SIB0_XF,11_0->SIB0_XF2,09_0	Up	SIB0_XF2,09_0->SIB0_XF,11_0	Up
SIB0_XF3,13_0->SIB0_XF,12_0	Up	SIB0_XF,12_0->SIB0_XF3,13_0	Up
SIB0_XF,12_0->SIB0_XF3,13_0	Up	SIB0_XF3,13_0->SIB0_XF,12_0	Up
SIB0_XF3,12_0->SIB0_XF,13_0	Up	SIB0_XF,13_0->SIB0_XF3,12_0	Up
SIB0_XF,13_0->SIB0_XF3,12_0	Up	SIB0_XF3,12_0->SIB0_XF,13_0	Up
SIB0_XF3,11_0->SIB0_XF,14_0	Up	SIB0_XF,14_0->SIB0_XF3,11_0	Up
SIB0_XF,14_0->SIB0_XF3,11_0	Up	SIB0_XF3,11_0->SIB0_XF,14_0	Up
SIB0_XF3,10_0->SIB0_XF,15_0	Up	SIB0_XF,15_0->SIB0_XF3,10_0	Up
SIB0_XF,15_0->SIB0_XF3,10_0	Up	SIB0_XF3,10_0->SIB0_XF,15_0	Up

PFE Links

FPC2PFE0->SIB0_XF1,05_0	Up	SIB0_XF1,05_0->FPC2PFE0	Up
FPC3PFE0->SIB0_XF2,15_0	Up	SIB0_XF2,15_0->FPC3PFE0	Up
FPC5PFE0->SIB0_XF2,05_0	Up	SIB0_XF2,05_0->FPC5PFE0	Up
FPC5PFE1->SIB0_XF2,07_0	Up	SIB0_XF2,07_0->FPC5PFE1	Up
FPC6PFE0->SIB0_XF3,01_0	Up	SIB0_XF3,01_0->FPC6PFE0	Up
FPC6PFE0->SIB0_XF3,01_1	Up	SIB0_XF3,01_1->FPC6PFE0	Up
FPC6PFE0->SIB0_XF3,02_0	Up	SIB0_XF3,02_0->FPC6PFE0	Up
FPC6PFE1->SIB0_XF3,03_0	Up	SIB0_XF3,03_0->FPC6PFE1	Up
FPC6PFE1->SIB0_XF3,03_1	Up	SIB0_XF3,03_1->FPC6PFE1	Up
FPC6PFE1->SIB0_XF3,02_1	Up	SIB0_XF3,02_1->FPC6PFE1	Up

show chassis fabric topology lcc (TX Matrix Plus Router with 3D SIBs)

```
user@host> show chassis fabric topology lcc 6
lcc6-re0:
```

fchip (mode)			
In-links	State	Out-links	State
SIB0 :			

CXP0_Evn->LCC_SIB0_XF3,10_0	Up	LCC_SIB0_XF3,10_0->CXP0_Evn	Up
CXP0_Odd->LCC_SIB0_XF3,11_0	Up	LCC_SIB0_XF3,11_0->CXP0_Odd	Up
CXP1_Evn->LCC_SIB0_XF3,12_0	Up	LCC_SIB0_XF3,12_0->CXP1_Evn	Up
CXP1_Odd->LCC_SIB0_XF3,13_0	Up	LCC_SIB0_XF3,13_0->CXP1_Odd	Up
CXP2_Evn->LCC_SIB0_XF2,09_0	Up	LCC_SIB0_XF2,09_0->CXP2_Evn	Up
CXP2_Odd->LCC_SIB0_XF2,10_0	Up	LCC_SIB0_XF2,10_0->CXP2_Odd	Up
CXP3_Evn->LCC_SIB0_XF2,11_0	Up	LCC_SIB0_XF2,11_0->CXP3_Evn	Up
CXP3_Odd->LCC_SIB0_XF2,12_0	Up	LCC_SIB0_XF2,12_0->CXP3_Odd	Up
CXP4_Evn->LCC_SIB0_XF2,13_0	Up	LCC_SIB0_XF2,13_0->CXP4_Evn	Up
CXP4_Odd->LCC_SIB0_XF1,09_0	Up	LCC_SIB0_XF1,09_0->CXP4_Odd	Up
CXP5_Evn->LCC_SIB0_XF2,14_0	Up	LCC_SIB0_XF2,14_0->CXP5_Evn	Up
CXP5_Odd->LCC_SIB0_XF1,10_0	Up	LCC_SIB0_XF1,10_0->CXP5_Odd	Up

```

CXP6_Evn->LCC_SIB0_XF1,11_0      Up      LCC_SIB0_XF1,11_0->CXP6_Evn      Up
CXP6_Odd->LCC_SIB0_XF1,12_0      Up      LCC_SIB0_XF1,12_0->CXP6_Odd      Up
CXP7_Evn->LCC_SIB0_XF1,13_0      Up      LCC_SIB0_XF1,13_0->CXP7_Evn      Up
CXP7_Odd->LCC_SIB0_XF1,14_0      Up      LCC_SIB0_XF1,14_0->CXP7_Odd      Up
SIB1 :
-----
SIB2 :
-----
CXP0_Evn->LCC_SIB2_XF3,10_0      Up      LCC_SIB2_XF3,10_0->CXP0_Evn      Up
CXP0_Odd->LCC_SIB2_XF3,11_0      Up      LCC_SIB2_XF3,11_0->CXP0_Odd      Up
CXP1_Evn->LCC_SIB2_XF3,12_0      Up      LCC_SIB2_XF3,12_0->CXP1_Evn      Up
CXP1_Odd->LCC_SIB2_XF3,13_0      Up      LCC_SIB2_XF3,13_0->CXP1_Odd      Up
CXP2_Evn->LCC_SIB2_XF2,09_0      Up      LCC_SIB2_XF2,09_0->CXP2_Evn      Up
CXP2_Odd->LCC_SIB2_XF2,10_0      Up      LCC_SIB2_XF2,10_0->CXP2_Odd      Up
CXP3_Evn->LCC_SIB2_XF2,11_0      Up      LCC_SIB2_XF2,11_0->CXP3_Evn      Up
CXP3_Odd->LCC_SIB2_XF2,12_0      Up      LCC_SIB2_XF2,12_0->CXP3_Odd      Up
CXP4_Evn->LCC_SIB2_XF2,13_0      Up      LCC_SIB2_XF2,13_0->CXP4_Evn      Up
CXP4_Odd->LCC_SIB2_XF1,09_0      Up      LCC_SIB2_XF1,09_0->CXP4_Odd      Up
CXP5_Evn->LCC_SIB2_XF2,14_0      Up      LCC_SIB2_XF2,14_0->CXP5_Evn      Up
CXP5_Odd->LCC_SIB2_XF1,10_0      Up      LCC_SIB2_XF1,10_0->CXP5_Odd      Up
CXP6_Evn->LCC_SIB2_XF1,11_0      Up      LCC_SIB2_XF1,11_0->CXP6_Evn      Up
CXP6_Odd->LCC_SIB2_XF1,12_0      Up      LCC_SIB2_XF1,12_0->CXP6_Odd      Up
CXP7_Evn->LCC_SIB2_XF1,13_0      Up      LCC_SIB2_XF1,13_0->CXP7_Evn      Up
CXP7_Odd->LCC_SIB2_XF1,14_0      Up      LCC_SIB2_XF1,14_0->CXP7_Odd      Up
SIB3 :
-----
CXP0_Evn->LCC_SIB3_XF3,10_0      Up      LCC_SIB3_XF3,10_0->CXP0_Evn      Up
CXP0_Odd->LCC_SIB3_XF3,11_0      Up      LCC_SIB3_XF3,11_0->CXP0_Odd      Up
CXP1_Evn->LCC_SIB3_XF3,12_0      Up      LCC_SIB3_XF3,12_0->CXP1_Evn      Up
CXP1_Odd->LCC_SIB3_XF3,13_0      Up      LCC_SIB3_XF3,13_0->CXP1_Odd      Up
CXP2_Evn->LCC_SIB3_XF2,09_0      Up      LCC_SIB3_XF2,09_0->CXP2_Evn      Up
CXP2_Odd->LCC_SIB3_XF2,10_0      Up      LCC_SIB3_XF2,10_0->CXP2_Odd      Up
CXP3_Evn->LCC_SIB3_XF2,11_0      Up      LCC_SIB3_XF2,11_0->CXP3_Evn      Up
CXP3_Odd->LCC_SIB3_XF2,12_0      Up      LCC_SIB3_XF2,12_0->CXP3_Odd      Up
CXP4_Evn->LCC_SIB3_XF2,13_0      Up      LCC_SIB3_XF2,13_0->CXP4_Evn      Up
CXP4_Odd->LCC_SIB3_XF1,09_0      Up      LCC_SIB3_XF1,09_0->CXP4_Odd      Up
CXP5_Evn->LCC_SIB3_XF2,14_0      Up      LCC_SIB3_XF2,14_0->CXP5_Evn      Up
CXP5_Odd->LCC_SIB3_XF1,10_0      Up      LCC_SIB3_XF1,10_0->CXP5_Odd      Up
CXP6_Evn->LCC_SIB3_XF1,11_0      Up      LCC_SIB3_XF1,11_0->CXP6_Evn      Up
CXP6_Odd->LCC_SIB3_XF1,12_0      Up      LCC_SIB3_XF1,12_0->CXP6_Odd      Up
CXP7_Evn->LCC_SIB3_XF1,13_0      Up      LCC_SIB3_XF1,13_0->CXP7_Evn      Up
CXP7_Odd->LCC_SIB3_XF1,14_0      Up      LCC_SIB3_XF1,14_0->CXP7_Odd      Up
SIB4 :
-----
CXP0_Evn->LCC_SIB4_XF3,10_0      Up      LCC_SIB4_XF3,10_0->CXP0_Evn      Up
CXP0_Odd->LCC_SIB4_XF3,11_0      Up      LCC_SIB4_XF3,11_0->CXP0_Odd      Up
CXP1_Evn->LCC_SIB4_XF3,12_0      Up      LCC_SIB4_XF3,12_0->CXP1_Evn      Up
CXP1_Odd->LCC_SIB4_XF3,13_0      Up      LCC_SIB4_XF3,13_0->CXP1_Odd      Up
CXP2_Evn->LCC_SIB4_XF2,09_0      Up      LCC_SIB4_XF2,09_0->CXP2_Evn      Up
CXP2_Odd->LCC_SIB4_XF2,10_0      Up      LCC_SIB4_XF2,10_0->CXP2_Odd      Up
CXP3_Evn->LCC_SIB4_XF2,11_0      Up      LCC_SIB4_XF2,11_0->CXP3_Evn      Up
CXP3_Odd->LCC_SIB4_XF2,12_0      Up      LCC_SIB4_XF2,12_0->CXP3_Odd      Up
CXP4_Evn->LCC_SIB4_XF2,13_0      Up      LCC_SIB4_XF2,13_0->CXP4_Evn      Up
CXP4_Odd->LCC_SIB4_XF1,09_0      Up      LCC_SIB4_XF1,09_0->CXP4_Odd      Up
CXP5_Evn->LCC_SIB4_XF2,14_0      Up      LCC_SIB4_XF2,14_0->CXP5_Evn      Up
CXP5_Odd->LCC_SIB4_XF1,10_0      Up      LCC_SIB4_XF1,10_0->CXP5_Odd      Up
CXP6_Evn->LCC_SIB4_XF1,11_0      Up      LCC_SIB4_XF1,11_0->CXP6_Evn      Up
CXP6_Odd->LCC_SIB4_XF1,12_0      Up      LCC_SIB4_XF1,12_0->CXP6_Odd      Up
CXP7_Evn->LCC_SIB4_XF1,13_0      Up      LCC_SIB4_XF1,13_0->CXP7_Evn      Up
CXP7_Odd->LCC_SIB4_XF1,14_0      Up      LCC_SIB4_XF1,14_0->CXP7_Odd      Up

```

show chassis fabric topology sfc (TX Matrix Plus Router with 3D SIBs)

```
user@host> show chassis fabric topology sfc 0
sfc0-re0:
```

fchip (mode)			
In-links	State	Out-links	State
F13_SIB0 :			

CXP0_Evn->F13_SIB0_XF2,04_0	Up	F13_SIB0_XF2,04_0->CXP0_Evn	Up
CXP0_Odd->F13_SIB0_XF2,03_0	Up	F13_SIB0_XF2,03_0->CXP0_Odd	Up
CXP1_Evn->F13_SIB0_XF2,06_0	Up	F13_SIB0_XF2,06_0->CXP1_Evn	Up
CXP1_Odd->F13_SIB0_XF2,05_0	Up	F13_SIB0_XF2,05_0->CXP1_Odd	Up
CXP2_Evn->F13_SIB0_XF2,08_0	Up	F13_SIB0_XF2,08_0->CXP2_Evn	Up
CXP2_Odd->F13_SIB0_XF2,07_0	Up	F13_SIB0_XF2,07_0->CXP2_Odd	Up
CXP3_Evn->F13_SIB0_XF2,10_0	Up	F13_SIB0_XF2,10_0->CXP3_Evn	Up
CXP3_Odd->F13_SIB0_XF2,09_0	Up	F13_SIB0_XF2,09_0->CXP3_Odd	Up
CXP4_Evn->F13_SIB0_XF0,04_0	Up	F13_SIB0_XF0,04_0->CXP4_Evn	Up
CXP4_Odd->F13_SIB0_XF0,03_0	Up	F13_SIB0_XF0,03_0->CXP4_Odd	Up
CXP5_Evn->F13_SIB0_XF0,06_0	Up	F13_SIB0_XF0,06_0->CXP5_Evn	Up
CXP5_Odd->F13_SIB0_XF0,05_0	Up	F13_SIB0_XF0,05_0->CXP5_Odd	Up
CXP6_Evn->F13_SIB0_XF0,08_0	Up	F13_SIB0_XF0,08_0->CXP6_Evn	Up
CXP6_Odd->F13_SIB0_XF0,07_0	Up	F13_SIB0_XF0,07_0->CXP6_Odd	Up
CXP7_Evn->F13_SIB0_XF0,10_0	Up	F13_SIB0_XF0,10_0->CXP7_Evn	Up
CXP7_Odd->F13_SIB0_XF0,09_0	Up	F13_SIB0_XF0,09_0->CXP7_Odd	Up
CXP8_Evn->F13_SIB0_XF3,04_0	Up	F13_SIB0_XF3,04_0->CXP8_Evn	Up
CXP8_Odd->F13_SIB0_XF3,03_0	Up	F13_SIB0_XF3,03_0->CXP8_Odd	Up
CXP9_Evn->F13_SIB0_XF3,06_0	Up	F13_SIB0_XF3,06_0->CXP9_Evn	Up
CXP9_Odd->F13_SIB0_XF3,05_0	Up	F13_SIB0_XF3,05_0->CXP9_Odd	Up
CXP10_Evn->F13_SIB0_XF3,08_0	Up	F13_SIB0_XF3,08_0->CXP10_Evn	Up
CXP10_Odd->F13_SIB0_XF3,07_0	Up	F13_SIB0_XF3,07_0->CXP10_Odd	Up
CXP11_Evn->F13_SIB0_XF3,10_0	Up	F13_SIB0_XF3,10_0->CXP11_Evn	Up
CXP11_Odd->F13_SIB0_XF3,09_0	Up	F13_SIB0_XF3,09_0->CXP11_Odd	Up
CXP12_Evn->F13_SIB0_XF1,04_0	Up	F13_SIB0_XF1,04_0->CXP12_Evn	Up
CXP12_Odd->F13_SIB0_XF1,03_0	Up	F13_SIB0_XF1,03_0->CXP12_Odd	Up
CXP13_Evn->F13_SIB0_XF1,06_0	Up	F13_SIB0_XF1,06_0->CXP13_Evn	Up
CXP13_Odd->F13_SIB0_XF1,05_0	Up	F13_SIB0_XF1,05_0->CXP13_Odd	Up
CXP14_Evn->F13_SIB0_XF1,08_0	Up	F13_SIB0_XF1,08_0->CXP14_Evn	Up
CXP14_Odd->F13_SIB0_XF1,07_0	Up	F13_SIB0_XF1,07_0->CXP14_Odd	Up
CXP15_Evn->F13_SIB0_XF1,10_0	Up	F13_SIB0_XF1,10_0->CXP15_Evn	Up
CXP15_Odd->F13_SIB0_XF1,09_0	Up	F13_SIB0_XF1,09_0->CXP15_Odd	Up
F13_SIB0_XF4,00_0->F13_SIB0_XF2,02_0	Up	F13_SIB0_XF2,02_0->F13_SIB0_XF4,00_0	Up
F13_SIB0_XF4,01_0->F13_SIB0_XF2,01_0	Up	F13_SIB0_XF2,01_0->F13_SIB0_XF4,01_0	Up
F13_SIB0_XF4,02_0->F13_SIB0_XF2,00_0	Up	F13_SIB0_XF2,00_0->F13_SIB0_XF4,02_0	Up
F13_SIB0_XF4,03_0->F13_SIB0_XF2,15_0	Up	F13_SIB0_XF2,15_0->F13_SIB0_XF4,03_0	Up
F13_SIB0_XF4,04_0->F13_SIB0_XF2,14_0	Up	F13_SIB0_XF2,14_0->F13_SIB0_XF4,04_0	Up
F13_SIB0_XF4,05_0->F13_SIB0_XF2,13_0	Up	F13_SIB0_XF2,13_0->F13_SIB0_XF4,05_0	Up
F13_SIB0_XF4,06_0->F13_SIB0_XF2,12_0	Up	F13_SIB0_XF2,12_0->F13_SIB0_XF4,06_0	Up
F13_SIB0_XF4,07_0->F13_SIB0_XF2,11_0	Up	F13_SIB0_XF2,11_0->F13_SIB0_XF4,07_0	Up
F13_SIB0_XF4,08_0->F13_SIB0_XF0,02_0	Up	F13_SIB0_XF0,02_0->F13_SIB0_XF4,08_0	Up
F13_SIB0_XF4,09_0->F13_SIB0_XF0,01_0	Up	F13_SIB0_XF0,01_0->F13_SIB0_XF4,09_0	Up

```

F13_SIB0_XF4,10_0->F13_SIB0_XF0,00_0 Up  F13_SIB0_XF0,00_0->F13_SIB0_XF4,10_0 Up
F13_SIB0_XF4,11_0->F13_SIB0_XF0,15_0 Up  F13_SIB0_XF0,15_0->F13_SIB0_XF4,11_0 Up
F13_SIB0_XF4,12_0->F13_SIB0_XF0,14_0 Up  F13_SIB0_XF0,14_0->F13_SIB0_XF4,12_0 Up
F13_SIB0_XF4,13_0->F13_SIB0_XF0,13_0 Up  F13_SIB0_XF0,13_0->F13_SIB0_XF4,13_0 Up
F13_SIB0_XF4,14_0->F13_SIB0_XF0,12_0 Up  F13_SIB0_XF0,12_0->F13_SIB0_XF4,14_0 Up
F13_SIB0_XF4,15_0->F13_SIB0_XF0,11_0 Up  F13_SIB0_XF0,11_0->F13_SIB0_XF4,15_0 Up
F13_SIB0_XF6,08_0->F13_SIB0_XF3,02_0 Up  F13_SIB0_XF3,02_0->F13_SIB0_XF6,08_0 Up
F13_SIB0_XF6,09_0->F13_SIB0_XF3,01_0 Up  F13_SIB0_XF3,01_0->F13_SIB0_XF6,09_0 Up
F13_SIB0_XF6,10_0->F13_SIB0_XF3,00_0 Up  F13_SIB0_XF3,00_0->F13_SIB0_XF6,10_0 Up
F13_SIB0_XF6,11_0->F13_SIB0_XF3,15_0 Up  F13_SIB0_XF3,15_0->F13_SIB0_XF6,11_0 Up
F13_SIB0_XF6,12_0->F13_SIB0_XF3,14_0 Up  F13_SIB0_XF3,14_0->F13_SIB0_XF6,12_0 Up
F13_SIB0_XF6,13_0->F13_SIB0_XF3,13_0 Up  F13_SIB0_XF3,13_0->F13_SIB0_XF6,13_0 Up
F13_SIB0_XF6,14_0->F13_SIB0_XF3,12_0 Up  F13_SIB0_XF3,12_0->F13_SIB0_XF6,14_0 Up
F13_SIB0_XF6,15_0->F13_SIB0_XF3,11_0 Up  F13_SIB0_XF3,11_0->F13_SIB0_XF6,15_0 Up
F13_SIB0_XF6,00_0->F13_SIB0_XF1,02_0 Up  F13_SIB0_XF1,02_0->F13_SIB0_XF6,00_0 Up
F13_SIB0_XF6,01_0->F13_SIB0_XF1,01_0 Up  F13_SIB0_XF1,01_0->F13_SIB0_XF6,01_0 Up
F13_SIB0_XF6,02_0->F13_SIB0_XF1,00_0 Up  F13_SIB0_XF1,00_0->F13_SIB0_XF6,02_0 Up
F13_SIB0_XF6,03_0->F13_SIB0_XF1,15_0 Up  F13_SIB0_XF1,15_0->F13_SIB0_XF6,03_0 Up
F13_SIB0_XF6,04_0->F13_SIB0_XF1,14_0 Up  F13_SIB0_XF1,14_0->F13_SIB0_XF6,04_0 Up
F13_SIB0_XF6,05_0->F13_SIB0_XF1,13_0 Up  F13_SIB0_XF1,13_0->F13_SIB0_XF6,05_0 Up
F13_SIB0_XF6,06_0->F13_SIB0_XF1,12_0 Up  F13_SIB0_XF1,12_0->F13_SIB0_XF6,06_0 Up
F13_SIB0_XF6,07_0->F13_SIB0_XF1,11_0 Up  F13_SIB0_XF1,11_0->F13_SIB0_XF6,07_0 Up
F13_SIB0_XF2,02_0->F13_SIB0_XF5,00_0 Up  F13_SIB0_XF5,00_0->F13_SIB0_XF2,02_0 Up
F13_SIB0_XF2,01_0->F13_SIB0_XF5,01_0 Up  F13_SIB0_XF5,01_0->F13_SIB0_XF2,01_0 Up
F13_SIB0_XF2,00_0->F13_SIB0_XF5,02_0 Up  F13_SIB0_XF5,02_0->F13_SIB0_XF2,00_0 Up
F13_SIB0_XF2,15_0->F13_SIB0_XF5,03_0 Up  F13_SIB0_XF5,03_0->F13_SIB0_XF2,15_0 Up
F13_SIB0_XF2,14_0->F13_SIB0_XF5,04_0 Up  F13_SIB0_XF5,04_0->F13_SIB0_XF2,14_0 Up
F13_SIB0_XF2,13_0->F13_SIB0_XF5,05_0 Up  F13_SIB0_XF5,05_0->F13_SIB0_XF2,13_0 Up
F13_SIB0_XF2,12_0->F13_SIB0_XF5,06_0 Up  F13_SIB0_XF5,06_0->F13_SIB0_XF2,12_0 Up
F13_SIB0_XF2,11_0->F13_SIB0_XF5,07_0 Up  F13_SIB0_XF5,07_0->F13_SIB0_XF2,11_0 Up

```

F13_SIB0_XF0,02_0->F13_SIB0_XF5,08_0 Up	F13_SIB0_XF5,08_0->F13_SIB0_XF0,02_0 Up
F13_SIB0_XF0,01_0->F13_SIB0_XF5,09_0 Up	F13_SIB0_XF5,09_0->F13_SIB0_XF0,01_0 Up
F13_SIB0_XF0,00_0->F13_SIB0_XF5,10_0 Up	F13_SIB0_XF5,10_0->F13_SIB0_XF0,00_0 Up
F13_SIB0_XF0,15_0->F13_SIB0_XF5,11_0 Up	F13_SIB0_XF5,11_0->F13_SIB0_XF0,15_0 Up
F13_SIB0_XF0,14_0->F13_SIB0_XF5,12_0 Up	F13_SIB0_XF5,12_0->F13_SIB0_XF0,14_0 Up
F13_SIB0_XF0,13_0->F13_SIB0_XF5,13_0 Up	F13_SIB0_XF5,13_0->F13_SIB0_XF0,13_0 Up
F13_SIB0_XF0,12_0->F13_SIB0_XF5,14_0 Up	F13_SIB0_XF5,14_0->F13_SIB0_XF0,12_0 Up
F13_SIB0_XF0,11_0->F13_SIB0_XF5,15_0 Up	F13_SIB0_XF5,15_0->F13_SIB0_XF0,11_0 Up
F13_SIB0_XF3,02_0->F13_SIB0_XF7,08_0 Up	F13_SIB0_XF7,08_0->F13_SIB0_XF3,02_0 Up
F13_SIB0_XF3,01_0->F13_SIB0_XF7,09_0 Up	F13_SIB0_XF7,09_0->F13_SIB0_XF3,01_0 Up
F13_SIB0_XF3,00_0->F13_SIB0_XF7,10_0 Up	F13_SIB0_XF7,10_0->F13_SIB0_XF3,00_0 Up
F13_SIB0_XF3,15_0->F13_SIB0_XF7,11_0 Up	F13_SIB0_XF7,11_0->F13_SIB0_XF3,15_0 Up
F13_SIB0_XF3,14_0->F13_SIB0_XF7,12_0 Up	F13_SIB0_XF7,12_0->F13_SIB0_XF3,14_0 Up
F13_SIB0_XF3,13_0->F13_SIB0_XF7,13_0 Up	F13_SIB0_XF7,13_0->F13_SIB0_XF3,13_0 Up
F13_SIB0_XF3,12_0->F13_SIB0_XF7,14_0 Up	F13_SIB0_XF7,14_0->F13_SIB0_XF3,12_0 Up
F13_SIB0_XF3,11_0->F13_SIB0_XF7,15_0 Up	F13_SIB0_XF7,15_0->F13_SIB0_XF3,11_0 Up
F13_SIB0_XF1,02_0->F13_SIB0_XF7,00_0 Up	F13_SIB0_XF7,00_0->F13_SIB0_XF1,02_0 Up
F13_SIB0_XF1,01_0->F13_SIB0_XF7,01_0 Up	F13_SIB0_XF7,01_0->F13_SIB0_XF1,01_0 Up
F13_SIB0_XF1,00_0->F13_SIB0_XF7,02_0 Up	F13_SIB0_XF7,02_0->F13_SIB0_XF1,00_0 Up
F13_SIB0_XF1,15_0->F13_SIB0_XF7,03_0 Up	F13_SIB0_XF7,03_0->F13_SIB0_XF1,15_0 Up
F13_SIB0_XF1,14_0->F13_SIB0_XF7,04_0 Up	F13_SIB0_XF7,04_0->F13_SIB0_XF1,14_0 Up
F13_SIB0_XF1,13_0->F13_SIB0_XF7,05_0 Up	F13_SIB0_XF7,05_0->F13_SIB0_XF1,13_0 Up
F13_SIB0_XF1,12_0->F13_SIB0_XF7,06_0 Up	F13_SIB0_XF7,06_0->F13_SIB0_XF1,12_0 Up
F13_SIB0_XF1,11_0->F13_SIB0_XF7,07_0 Up	F13_SIB0_XF7,07_0->F13_SIB0_XF1,11_0 Up
F2S_SIB2_XF,12_0->F13_SIB0_XF4,00_0 Up	F13_SIB0_XF4,00_0->F2S_SIB2_XF,12_0 Up
F2S_SIB2_XF,08_0->F13_SIB0_XF4,01_0 Up	F13_SIB0_XF4,01_0->F2S_SIB2_XF,08_0 Up
F2S_SIB2_XF,14_0->F13_SIB0_XF4,02_0 Up	F13_SIB0_XF4,02_0->F2S_SIB2_XF,14_0 Up
F2S_SIB2_XF,10_0->F13_SIB0_XF4,03_0 Up	F13_SIB0_XF4,03_0->F2S_SIB2_XF,10_0 Up
F2S_SIB3_XF,12_0->F13_SIB0_XF4,04_0 Up	F13_SIB0_XF4,04_0->F2S_SIB3_XF,12_0 Up
F2S_SIB3_XF,08_0->F13_SIB0_XF4,05_0 Up	F13_SIB0_XF4,05_0->F2S_SIB3_XF,08_0 Up
F2S_SIB3_XF,14_0->F13_SIB0_XF4,06_0 Up	F13_SIB0_XF4,06_0->F2S_SIB3_XF,14_0 Up

F2S_SIB3_XF,10_0->F13_SIB0_XF4,07_0 Up	F13_SIB0_XF4,07_0->F2S_SIB3_XF,10_0 Up
F2S_SIB0_XF,12_0->F13_SIB0_XF4,08_0 Up	F13_SIB0_XF4,08_0->F2S_SIB0_XF,12_0 Up
F2S_SIB0_XF,08_0->F13_SIB0_XF4,09_0 Up	F13_SIB0_XF4,09_0->F2S_SIB0_XF,08_0 Up
F2S_SIB0_XF,14_0->F13_SIB0_XF4,10_0 Up	F13_SIB0_XF4,10_0->F2S_SIB0_XF,14_0 Up
F2S_SIB0_XF,10_0->F13_SIB0_XF4,11_0 Up	F13_SIB0_XF4,11_0->F2S_SIB0_XF,10_0 Up
F2S_SIB1_XF,12_0->F13_SIB0_XF4,12_0 Up	F13_SIB0_XF4,12_0->F2S_SIB1_XF,12_0 Up
F2S_SIB1_XF,08_0->F13_SIB0_XF4,13_0 Up	F13_SIB0_XF4,13_0->F2S_SIB1_XF,08_0 Up
F2S_SIB1_XF,14_0->F13_SIB0_XF4,14_0 Up	F13_SIB0_XF4,14_0->F2S_SIB1_XF,14_0 Up
F2S_SIB1_XF,10_0->F13_SIB0_XF4,15_0 Up	F13_SIB0_XF4,15_0->F2S_SIB1_XF,10_0 Up
F2S_SIB2_XF,13_0->F13_SIB0_XF6,00_0 Up	F13_SIB0_XF6,00_0->F2S_SIB2_XF,13_0 Up
F2S_SIB2_XF,09_0->F13_SIB0_XF6,01_0 Up	F13_SIB0_XF6,01_0->F2S_SIB2_XF,09_0 Up
F2S_SIB2_XF,15_0->F13_SIB0_XF6,02_0 Up	F13_SIB0_XF6,02_0->F2S_SIB2_XF,15_0 Up
F2S_SIB2_XF,11_0->F13_SIB0_XF6,03_0 Up	F13_SIB0_XF6,03_0->F2S_SIB2_XF,11_0 Up
F2S_SIB3_XF,13_0->F13_SIB0_XF6,04_0 Up	F13_SIB0_XF6,04_0->F2S_SIB3_XF,13_0 Up
F2S_SIB3_XF,09_0->F13_SIB0_XF6,05_0 Up	F13_SIB0_XF6,05_0->F2S_SIB3_XF,09_0 Up
F2S_SIB3_XF,15_0->F13_SIB0_XF6,06_0 Up	F13_SIB0_XF6,06_0->F2S_SIB3_XF,15_0 Up
F2S_SIB3_XF,11_0->F13_SIB0_XF6,07_0 Up	F13_SIB0_XF6,07_0->F2S_SIB3_XF,11_0 Up
F2S_SIB0_XF,13_0->F13_SIB0_XF6,08_0 Up	F13_SIB0_XF6,08_0->F2S_SIB0_XF,13_0 Up
F2S_SIB0_XF,09_0->F13_SIB0_XF6,09_0 Up	F13_SIB0_XF6,09_0->F2S_SIB0_XF,09_0 Up
F2S_SIB0_XF,15_0->F13_SIB0_XF6,10_0 Up	F13_SIB0_XF6,10_0->F2S_SIB0_XF,15_0 Up
F2S_SIB0_XF,11_0->F13_SIB0_XF6,11_0 Up	F13_SIB0_XF6,11_0->F2S_SIB0_XF,11_0 Up
F2S_SIB1_XF,13_0->F13_SIB0_XF6,12_0 Up	F13_SIB0_XF6,12_0->F2S_SIB1_XF,13_0 Up
F2S_SIB1_XF,09_0->F13_SIB0_XF6,13_0 Up	F13_SIB0_XF6,13_0->F2S_SIB1_XF,09_0 Up
F2S_SIB1_XF,15_0->F13_SIB0_XF6,14_0 Up	F13_SIB0_XF6,14_0->F2S_SIB1_XF,15_0 Up
F2S_SIB1_XF,11_0->F13_SIB0_XF6,15_0 Up	F13_SIB0_XF6,15_0->F2S_SIB1_XF,11_0 Up
F13_SIB0_XF5,00_0->F2S_SIB2_XF,12_0 Up	F2S_SIB2_XF,12_0->F13_SIB0_XF5,00_0 Up
F13_SIB0_XF5,01_0->F2S_SIB2_XF,08_0 Up	F2S_SIB2_XF,08_0->F13_SIB0_XF5,01_0 Up
F13_SIB0_XF5,02_0->F2S_SIB2_XF,14_0 Up	F2S_SIB2_XF,14_0->F13_SIB0_XF5,02_0 Up
F13_SIB0_XF5,03_0->F2S_SIB2_XF,10_0 Up	F2S_SIB2_XF,10_0->F13_SIB0_XF5,03_0 Up
F13_SIB0_XF5,04_0->F2S_SIB3_XF,12_0 Up	F2S_SIB3_XF,12_0->F13_SIB0_XF5,04_0 Up

```

F13_SIB0_XF5,05_0->F2S_SIB3_XF,08_0 Up   F2S_SIB3_XF,08_0->F13_SIB0_XF5,05_0 Up
F13_SIB0_XF5,06_0->F2S_SIB3_XF,14_0 Up   F2S_SIB3_XF,14_0->F13_SIB0_XF5,06_0 Up
F13_SIB0_XF5,07_0->F2S_SIB3_XF,10_0 Up   F2S_SIB3_XF,10_0->F13_SIB0_XF5,07_0 Up
F13_SIB0_XF5,08_0->F2S_SIB0_XF,12_0 Up   F2S_SIB0_XF,12_0->F13_SIB0_XF5,08_0 Up
F13_SIB0_XF5,09_0->F2S_SIB0_XF,08_0 Up   F2S_SIB0_XF,08_0->F13_SIB0_XF5,09_0 Up
F13_SIB0_XF5,10_0->F2S_SIB0_XF,14_0 Up   F2S_SIB0_XF,14_0->F13_SIB0_XF5,10_0 Up
F13_SIB0_XF5,11_0->F2S_SIB0_XF,10_0 Up   F2S_SIB0_XF,10_0->F13_SIB0_XF5,11_0 Up
F13_SIB0_XF5,12_0->F2S_SIB1_XF,12_0 Up   F2S_SIB1_XF,12_0->F13_SIB0_XF5,12_0 Up
F13_SIB0_XF5,13_0->F2S_SIB1_XF,08_0 Up   F2S_SIB1_XF,08_0->F13_SIB0_XF5,13_0 Up
F13_SIB0_XF5,14_0->F2S_SIB1_XF,14_0 Up   F2S_SIB1_XF,14_0->F13_SIB0_XF5,14_0 Up
F13_SIB0_XF5,15_0->F2S_SIB1_XF,10_0 Up   F2S_SIB1_XF,10_0->F13_SIB0_XF5,15_0 Up
F13_SIB0_XF7,00_0->F2S_SIB2_XF,13_0 Up   F2S_SIB2_XF,13_0->F13_SIB0_XF7,00_0 Up
F13_SIB0_XF7,01_0->F2S_SIB2_XF,09_0 Up   F2S_SIB2_XF,09_0->F13_SIB0_XF7,01_0 Up
F13_SIB0_XF7,02_0->F2S_SIB2_XF,15_0 Up   F2S_SIB2_XF,15_0->F13_SIB0_XF7,02_0 Up
F13_SIB0_XF7,03_0->F2S_SIB2_XF,11_0 Up   F2S_SIB2_XF,11_0->F13_SIB0_XF7,03_0 Up
F13_SIB0_XF7,04_0->F2S_SIB3_XF,13_0 Up   F2S_SIB3_XF,13_0->F13_SIB0_XF7,04_0 Up
F13_SIB0_XF7,05_0->F2S_SIB3_XF,09_0 Up   F2S_SIB3_XF,09_0->F13_SIB0_XF7,05_0 Up
F13_SIB0_XF7,06_0->F2S_SIB3_XF,15_0 Up   F2S_SIB3_XF,15_0->F13_SIB0_XF7,06_0 Up
F13_SIB0_XF7,07_0->F2S_SIB3_XF,11_0 Up   F2S_SIB3_XF,11_0->F13_SIB0_XF7,07_0 Up
F13_SIB0_XF7,08_0->F2S_SIB0_XF,13_0 Up   F2S_SIB0_XF,13_0->F13_SIB0_XF7,08_0 Up
F13_SIB0_XF7,09_0->F2S_SIB0_XF,09_0 Up   F2S_SIB0_XF,09_0->F13_SIB0_XF7,09_0 Up
F13_SIB0_XF7,10_0->F2S_SIB0_XF,15_0 Up   F2S_SIB0_XF,15_0->F13_SIB0_XF7,10_0 Up
F13_SIB0_XF7,11_0->F2S_SIB0_XF,11_0 Up   F2S_SIB0_XF,11_0->F13_SIB0_XF7,11_0 Up
F13_SIB0_XF7,12_0->F2S_SIB1_XF,13_0 Up   F2S_SIB1_XF,13_0->F13_SIB0_XF7,12_0 Up
F13_SIB0_XF7,13_0->F2S_SIB1_XF,09_0 Up   F2S_SIB1_XF,09_0->F13_SIB0_XF7,13_0 Up
F13_SIB0_XF7,14_0->F2S_SIB1_XF,15_0 Up   F2S_SIB1_XF,15_0->F13_SIB0_XF7,14_0 Up
F13_SIB0_XF7,15_0->F2S_SIB1_XF,11_0 Up   F2S_SIB1_XF,11_0->F13_SIB0_XF7,15_0 Up

...

```

show chassis fabric topology (PTX5000 Router)

```
user@host> show chassis fabric topology
```

In-link : FPC# FE# TQ# (TQ-TX sub-chnl #) ->
 SIB# TF#_FCORE# (TF-RX port#, TF-RX sub-chn#, TF-RX inst#)

Out-link : SIB# TF#_FCORE# (TF-TX port#, TF-TX sub-chn#, TF-TX inst#) ->
 FPC# FE# TQ# (TQ-RX sub-chnl #)

(6, 4, 06) in FPC02FE0TQ0(02)->S01F0_0(6,4,06) will be TF Rx Port 6, TF CCL Rx Sub-Channel 4, TF CCL Rx Instance 6.
 (2, 7, 10) in S01F0_0(2,7,10)->FPC02FE0TQ0(02) will be TF-Tx Port 2, TF CCL Tx Sub-channel 7, TF CCL Tx Instance 10.
 SIB 0 FCHIP 0 FCORE 0 :

In-links	State	Out-links	State
FPC00FE0TQ0(00)->S00F0_0(7,4,07)	OK	S00F0_0(3,7,11)->FPC00FE0TQ0(00)	OK
FPC00FE1TQ1(00)->S00F0_0(7,6,07)	OK	S00F0_0(3,5,11)->FPC00FE1TQ1(00)	OK
FPC00FE2TQ2(00)->S00F0_0(7,5,07)	OK	S00F0_0(3,6,11)->FPC00FE2TQ2(00)	OK
FPC00FE3TQ3(00)->S00F0_0(7,7,07)	OK	S00F0_0(3,4,11)->FPC00FE3TQ3(00)	OK
FPC01FE0TQ0(00)->S00F0_0(7,0,07)	OK	S00F0_0(3,3,11)->FPC01FE0TQ0(00)	OK
FPC01FE1TQ1(00)->S00F0_0(7,1,07)	OK	S00F0_0(3,1,11)->FPC01FE1TQ1(00)	OK
FPC01FE2TQ2(00)->S00F0_0(7,2,07)	OK	S00F0_0(3,2,11)->FPC01FE2TQ2(00)	Error
FPC01FE3TQ3(00)->S00F0_0(7,3,07)	OK	S00F0_0(3,0,11)->FPC01FE3TQ3(00)	OK
FPC02FE0TQ0(00)->S00F0_0(6,4,06)	OK	S00F0_0(2,7,10)->FPC02FE0TQ0(00)	OK
FPC02FE1TQ1(00)->S00F0_0(6,5,06)	OK	S00F0_0(2,5,10)->FPC02FE1TQ1(00)	OK
FPC02FE2TQ2(00)->S00F0_0(6,6,06)	OK	S00F0_0(2,6,10)->FPC02FE2TQ2(00)	OK
FPC02FE3TQ3(00)->S00F0_0(6,7,06)	OK	S00F0_0(2,4,10)->FPC02FE3TQ3(00)	OK
FPC03FE0TQ0(00)->S00F0_0(6,0,06)	Down	S00F0_0(2,3,10)->FPC03FE0TQ0(00)	Down
FPC03FE1TQ1(00)->S00F0_0(6,1,06)	Down	S00F0_0(2,0,10)->FPC03FE1TQ1(00)	Down
FPC03FE2TQ2(00)->S00F0_0(6,2,06)	Down	S00F0_0(2,2,10)->FPC03FE2TQ2(00)	Down
FPC03FE3TQ3(00)->S00F0_0(6,3,06)	Down	S00F0_0(2,1,10)->FPC03FE3TQ3(00)	Down
FPC04FE0TQ0(00)->S00F0_0(5,4,05)	OK	S00F0_0(1,7,09)->FPC04FE0TQ0(00)	OK
FPC04FE1TQ1(00)->S00F0_0(5,5,05)	OK	S00F0_0(1,6,09)->FPC04FE1TQ1(00)	OK
FPC04FE2TQ2(00)->S00F0_0(5,6,05)	OK	S00F0_0(1,4,09)->FPC04FE2TQ2(00)	OK
FPC04FE3TQ3(00)->S00F0_0(5,7,05)	OK	S00F0_0(1,5,09)->FPC04FE3TQ3(00)	OK
FPC05FE0TQ0(00)->S00F0_0(5,0,05)	OK	S00F0_0(1,3,09)->FPC05FE0TQ0(00)	OK
FPC05FE1TQ1(00)->S00F0_0(5,1,05)	OK	S00F0_0(1,0,09)->FPC05FE1TQ1(00)	OK
FPC05FE2TQ2(00)->S00F0_0(5,2,05)	OK	S00F0_0(1,2,09)->FPC05FE2TQ2(00)	OK
FPC05FE3TQ3(00)->S00F0_0(5,3,05)	OK	S00F0_0(1,1,09)->FPC05FE3TQ3(00)	OK
FPC06FE0TQ0(00)->S00F0_0(4,4,04)	Down	S00F0_0(0,7,08)->FPC06FE0TQ0(00)	Down
FPC06FE1TQ1(00)->S00F0_0(4,5,04)	Down	S00F0_0(0,5,08)->FPC06FE1TQ1(00)	Down
FPC06FE2TQ2(00)->S00F0_0(4,6,04)	Down	S00F0_0(0,6,08)->FPC06FE2TQ2(00)	Down
FPC06FE3TQ3(00)->S00F0_0(4,7,04)	Down	S00F0_0(0,4,08)->FPC06FE3TQ3(00)	Down
FPC07FE0TQ0(00)->S00F0_0(4,2,04)	Down	S00F0_0(0,3,08)->FPC07FE0TQ0(00)	Down
FPC07FE1TQ1(00)->S00F0_0(4,0,04)	Down	S00F0_0(0,0,08)->FPC07FE1TQ1(00)	Down
FPC07FE2TQ2(00)->S00F0_0(4,1,04)	Down	S00F0_0(0,1,08)->FPC07FE2TQ2(00)	Down
FPC07FE3TQ3(00)->S00F0_0(4,3,04)	Down	S00F0_0(0,2,08)->FPC07FE3TQ3(00)	Down

SIB 0 FCHIP 0 FCORE 1 :

In-links	State	Out-links	State
FPC00FE0TQ0(01)->S00F0_1(3,4,11)	OK	S00F0_1(7,6,07)->FPC00FE0TQ0(01)	OK
FPC00FE1TQ1(01)->S00F0_1(3,5,11)	OK	S00F0_1(7,4,07)->FPC00FE1TQ1(01)	OK
FPC00FE2TQ2(01)->S00F0_1(3,6,11)	OK	S00F0_1(7,7,07)->FPC00FE2TQ2(01)	OK
FPC00FE3TQ3(01)->S00F0_1(3,7,11)	OK	S00F0_1(7,5,07)->FPC00FE3TQ3(01)	OK
FPC01FE0TQ0(01)->S00F0_1(3,0,11)	OK	S00F0_1(7,2,07)->FPC01FE0TQ0(01)	OK
FPC01FE1TQ1(01)->S00F0_1(3,1,11)	OK	S00F0_1(7,0,07)->FPC01FE1TQ1(01)	OK
FPC01FE2TQ2(01)->S00F0_1(3,2,11)	OK	S00F0_1(7,3,07)->FPC01FE2TQ2(01)	OK
FPC01FE3TQ3(01)->S00F0_1(3,3,11)	OK	S00F0_1(7,1,07)->FPC01FE3TQ3(01)	OK
FPC02FE0TQ0(01)->S00F0_1(2,4,10)	OK	S00F0_1(6,5,06)->FPC02FE0TQ0(01)	OK
FPC02FE1TQ1(01)->S00F0_1(2,5,10)	OK	S00F0_1(6,4,06)->FPC02FE1TQ1(01)	OK
FPC02FE2TQ2(01)->S00F0_1(2,6,10)	OK	S00F0_1(6,7,06)->FPC02FE2TQ2(01)	OK

```

FPC02FE3TQ3(01)->S00F0_1(2,7,10) OK      S00F0_1(6,6,06)->FPC02FE3TQ3(01) OK
FPC03FE0TQ0(01)->S00F0_1(2,0,10) Down    S00F0_1(6,1,06)->FPC03FE0TQ0(01) Down
FPC03FE1TQ1(01)->S00F0_1(2,1,10) Down    S00F0_1(6,0,06)->FPC03FE1TQ1(01) Down
FPC03FE2TQ2(01)->S00F0_1(2,2,10) Down    S00F0_1(6,3,06)->FPC03FE2TQ2(01) Down
FPC03FE3TQ3(01)->S00F0_1(2,3,10) Down    S00F0_1(6,2,06)->FPC03FE3TQ3(01) Down
FPC04FE0TQ0(01)->S00F0_1(1,4,09) OK      S00F0_1(5,5,05)->FPC04FE0TQ0(01) OK
FPC04FE1TQ1(01)->S00F0_1(1,5,09) OK      S00F0_1(5,4,05)->FPC04FE1TQ1(01) OK
FPC04FE2TQ2(01)->S00F0_1(1,6,09) OK      S00F0_1(5,7,05)->FPC04FE2TQ2(01) OK
FPC04FE3TQ3(01)->S00F0_1(1,7,09) OK      S00F0_1(5,6,05)->FPC04FE3TQ3(01) OK
FPC05FE0TQ0(01)->S00F0_1(1,0,09) OK      S00F0_1(5,1,05)->FPC05FE0TQ0(01) OK
FPC05FE1TQ1(01)->S00F0_1(1,1,09) OK      S00F0_1(5,0,05)->FPC05FE1TQ1(01) OK
FPC05FE2TQ2(01)->S00F0_1(1,2,09) OK      S00F0_1(5,3,05)->FPC05FE2TQ2(01) OK
FPC05FE3TQ3(01)->S00F0_1(1,3,09) OK      S00F0_1(5,2,05)->FPC05FE3TQ3(01) OK
FPC06FE0TQ0(01)->S00F0_1(0,4,08) Down    S00F0_1(4,7,04)->FPC06FE0TQ0(01) Down
FPC06FE1TQ1(01)->S00F0_1(0,5,08) Down    S00F0_1(4,0,04)->FPC06FE1TQ1(01) Down
FPC06FE2TQ2(01)->S00F0_1(0,6,08) Down    S00F0_1(4,6,04)->FPC06FE2TQ2(01) Down
FPC06FE3TQ3(01)->S00F0_1(0,7,08) Down    S00F0_1(4,1,04)->FPC06FE3TQ3(01) Down
FPC07FE0TQ0(01)->S00F0_1(0,0,08) Down    S00F0_1(4,3,04)->FPC07FE0TQ0(01) Down
FPC07FE1TQ1(01)->S00F0_1(0,1,08) Down    S00F0_1(4,4,04)->FPC07FE1TQ1(01) Down
FPC07FE2TQ2(01)->S00F0_1(0,2,08) Down    S00F0_1(4,2,04)->FPC07FE2TQ2(01) Down
FPC07FE3TQ3(01)->S00F0_1(0,3,08) Down    S00F0_1(4,5,04)->FPC07FE3TQ3(01) Down

```

SIB 1 FCHIP 0 FCORE 0 :

In-links	State	Out-links	State
FPC00FE0TQ0(02)->S01F0_0(7,4,07)	Error	S01F0_0(3,7,11)->FPC00FE0TQ0(02)	Down
FPC00FE1TQ1(02)->S01F0_0(7,6,07)	OK	S01F0_0(3,5,11)->FPC00FE1TQ1(02)	OK
FPC00FE2TQ2(02)->S01F0_0(7,5,07)	OK	S01F0_0(3,6,11)->FPC00FE2TQ2(02)	OK
FPC00FE3TQ3(02)->S01F0_0(7,7,07)	OK	S01F0_0(3,4,11)->FPC00FE3TQ3(02)	OK
FPC01FE0TQ0(02)->S01F0_0(7,0,07)	OK	S01F0_0(3,3,11)->FPC01FE0TQ0(02)	OK
FPC01FE1TQ1(02)->S01F0_0(7,1,07)	OK	S01F0_0(3,1,11)->FPC01FE1TQ1(02)	OK
FPC01FE2TQ2(02)->S01F0_0(7,2,07)	OK	S01F0_0(3,2,11)->FPC01FE2TQ2(02)	OK
FPC01FE3TQ3(02)->S01F0_0(7,3,07)	OK	S01F0_0(3,0,11)->FPC01FE3TQ3(02)	OK
FPC02FE0TQ0(02)->S01F0_0(6,4,06)	OK	S01F0_0(2,7,10)->FPC02FE0TQ0(02)	OK
FPC02FE1TQ1(02)->S01F0_0(6,5,06)	OK	S01F0_0(2,5,10)->FPC02FE1TQ1(02)	OK
FPC02FE2TQ2(02)->S01F0_0(6,6,06)	OK	S01F0_0(2,6,10)->FPC02FE2TQ2(02)	OK
FPC02FE3TQ3(02)->S01F0_0(6,7,06)	OK	S01F0_0(2,4,10)->FPC02FE3TQ3(02)	OK
FPC03FE0TQ0(02)->S01F0_0(6,0,06)	Down	S01F0_0(2,3,10)->FPC03FE0TQ0(02)	Down
FPC03FE1TQ1(02)->S01F0_0(6,1,06)	Down	S01F0_0(2,0,10)->FPC03FE1TQ1(02)	Down
FPC03FE2TQ2(02)->S01F0_0(6,2,06)	Down	S01F0_0(2,2,10)->FPC03FE2TQ2(02)	Down
FPC03FE3TQ3(02)->S01F0_0(6,3,06)	Down	S01F0_0(2,1,10)->FPC03FE3TQ3(02)	Down
FPC04FE0TQ0(02)->S01F0_0(5,4,05)	OK	S01F0_0(1,7,09)->FPC04FE0TQ0(02)	OK
FPC04FE1TQ1(02)->S01F0_0(5,5,05)	OK	S01F0_0(1,6,09)->FPC04FE1TQ1(02)	OK
FPC04FE2TQ2(02)->S01F0_0(5,6,05)	OK	S01F0_0(1,4,09)->FPC04FE2TQ2(02)	OK
FPC04FE3TQ3(02)->S01F0_0(5,7,05)	OK	S01F0_0(1,5,09)->FPC04FE3TQ3(02)	OK
FPC05FE0TQ0(02)->S01F0_0(5,0,05)	OK	S01F0_0(1,3,09)->FPC05FE0TQ0(02)	OK
FPC05FE1TQ1(02)->S01F0_0(5,1,05)	OK	S01F0_0(1,0,09)->FPC05FE1TQ1(02)	OK
FPC05FE2TQ2(02)->S01F0_0(5,2,05)	OK	S01F0_0(1,2,09)->FPC05FE2TQ2(02)	OK
FPC05FE3TQ3(02)->S01F0_0(5,3,05)	OK	S01F0_0(1,1,09)->FPC05FE3TQ3(02)	OK
FPC06FE0TQ0(02)->S01F0_0(4,4,04)	Down	S01F0_0(0,7,08)->FPC06FE0TQ0(02)	Down
FPC06FE1TQ1(02)->S01F0_0(4,5,04)	Down	S01F0_0(0,5,08)->FPC06FE1TQ1(02)	Down
FPC06FE2TQ2(02)->S01F0_0(4,6,04)	Down	S01F0_0(0,6,08)->FPC06FE2TQ2(02)	Down
FPC06FE3TQ3(02)->S01F0_0(4,7,04)	Down	S01F0_0(0,4,08)->FPC06FE3TQ3(02)	Down
FPC07FE0TQ0(02)->S01F0_0(4,2,04)	Down	S01F0_0(0,3,08)->FPC07FE0TQ0(02)	Down
FPC07FE1TQ1(02)->S01F0_0(4,0,04)	Down	S01F0_0(0,0,08)->FPC07FE1TQ1(02)	Down
FPC07FE2TQ2(02)->S01F0_0(4,1,04)	Down	S01F0_0(0,1,08)->FPC07FE2TQ2(02)	Down
FPC07FE3TQ3(02)->S01F0_0(4,3,04)	Down	S01F0_0(0,2,08)->FPC07FE3TQ3(02)	Down

SIB 1 FCHIP 0 FCORE 1 :

In-links	State	Out-links	State
FPC00FE0TQ0(03)->S01F0_1(3,4,11)	OK	S01F0_1(7,6,07)->FPC00FE0TQ0(03)	OK
FPC00FE1TQ1(03)->S01F0_1(3,5,11)	OK	S01F0_1(7,4,07)->FPC00FE1TQ1(03)	OK
FPC00FE2TQ2(03)->S01F0_1(3,6,11)	OK	S01F0_1(7,7,07)->FPC00FE2TQ2(03)	OK
FPC00FE3TQ3(03)->S01F0_1(3,7,11)	OK	S01F0_1(7,5,07)->FPC00FE3TQ3(03)	OK
FPC01FE0TQ0(03)->S01F0_1(3,0,11)	OK	S01F0_1(7,2,07)->FPC01FE0TQ0(03)	OK
FPC01FE1TQ1(03)->S01F0_1(3,1,11)	OK	S01F0_1(7,0,07)->FPC01FE1TQ1(03)	OK
FPC01FE2TQ2(03)->S01F0_1(3,2,11)	OK	S01F0_1(7,3,07)->FPC01FE2TQ2(03)	OK
FPC01FE3TQ3(03)->S01F0_1(3,3,11)	OK	S01F0_1(7,1,07)->FPC01FE3TQ3(03)	OK
FPC02FE0TQ0(03)->S01F0_1(2,4,10)	OK	S01F0_1(6,5,06)->FPC02FE0TQ0(03)	OK
FPC02FE1TQ1(03)->S01F0_1(2,5,10)	OK	S01F0_1(6,4,06)->FPC02FE1TQ1(03)	OK
FPC02FE2TQ2(03)->S01F0_1(2,6,10)	OK	S01F0_1(6,7,06)->FPC02FE2TQ2(03)	OK
FPC02FE3TQ3(03)->S01F0_1(2,7,10)	OK	S01F0_1(6,6,06)->FPC02FE3TQ3(03)	OK
FPC03FE0TQ0(03)->S01F0_1(2,0,10)	Down	S01F0_1(6,1,06)->FPC03FE0TQ0(03)	Down
FPC03FE1TQ1(03)->S01F0_1(2,1,10)	Down	S01F0_1(6,0,06)->FPC03FE1TQ1(03)	Down
FPC03FE2TQ2(03)->S01F0_1(2,2,10)	Down	S01F0_1(6,3,06)->FPC03FE2TQ2(03)	Down
FPC03FE3TQ3(03)->S01F0_1(2,3,10)	Down	S01F0_1(6,2,06)->FPC03FE3TQ3(03)	Down
FPC04FE0TQ0(03)->S01F0_1(1,4,09)	OK	S01F0_1(5,5,05)->FPC04FE0TQ0(03)	OK
FPC04FE1TQ1(03)->S01F0_1(1,5,09)	OK	S01F0_1(5,4,05)->FPC04FE1TQ1(03)	OK
FPC04FE2TQ2(03)->S01F0_1(1,6,09)	OK	S01F0_1(5,7,05)->FPC04FE2TQ2(03)	OK
FPC04FE3TQ3(03)->S01F0_1(1,7,09)	OK	S01F0_1(5,6,05)->FPC04FE3TQ3(03)	OK
FPC05FE0TQ0(03)->S01F0_1(1,0,09)	OK	S01F0_1(5,1,05)->FPC05FE0TQ0(03)	OK
FPC05FE1TQ1(03)->S01F0_1(1,1,09)	OK	S01F0_1(5,0,05)->FPC05FE1TQ1(03)	OK
FPC05FE2TQ2(03)->S01F0_1(1,2,09)	OK	S01F0_1(5,3,05)->FPC05FE2TQ2(03)	OK
FPC05FE3TQ3(03)->S01F0_1(1,3,09)	OK	S01F0_1(5,2,05)->FPC05FE3TQ3(03)	OK
FPC06FE0TQ0(03)->S01F0_1(0,4,08)	Down	S01F0_1(4,7,04)->FPC06FE0TQ0(03)	Down
FPC06FE1TQ1(03)->S01F0_1(0,5,08)	Down	S01F0_1(4,0,04)->FPC06FE1TQ1(03)	Down
FPC06FE2TQ2(03)->S01F0_1(0,6,08)	Down	S01F0_1(4,6,04)->FPC06FE2TQ2(03)	Down
FPC06FE3TQ3(03)->S01F0_1(0,7,08)	Down	S01F0_1(4,1,04)->FPC06FE3TQ3(03)	Down
FPC07FE0TQ0(03)->S01F0_1(0,0,08)	Down	S01F0_1(4,3,04)->FPC07FE0TQ0(03)	Down
FPC07FE1TQ1(03)->S01F0_1(0,1,08)	Down	S01F0_1(4,4,04)->FPC07FE1TQ1(03)	Down
FPC07FE2TQ2(03)->S01F0_1(0,2,08)	Down	S01F0_1(4,2,04)->FPC07FE2TQ2(03)	Down
FPC07FE3TQ3(03)->S01F0_1(0,3,08)	Down	S01F0_1(4,5,04)->FPC07FE3TQ3(03)	Down

show chassis fabric topology (PTX3000 Router)

```

user@host> show chassis fabric topology
In-link : FPC# FE# TQ# (TQ-TX sub-chnl #) ->
           SIB# TF#_FCORE# (TF-RX port#, TF-RX sub-chn#, TF-RX inst#)

Out-link : SIB# TF#_FCORE# (TF-TX port#, TF-TX sub-chn#, TF-TX inst#) ->
           FPC# FE# TQ# (TQ-RX sub-chnl #)
SIB 0 FCHIP 0 FCORE 0 :

```

In-links	State	Out-links	State
FPC00FE0TQ0(00)->S00F0_0(7,0,07)	Down	S00F0_0(3,0,11)->FPC00FE0TQ0(00)	Down
FPC00FE1TQ1(00)->S00F0_0(7,1,07)	Down	S00F0_0(3,1,11)->FPC00FE1TQ1(00)	Down
FPC02FE0TQ0(00)->S00F0_0(6,0,06)	Down	S00F0_0(2,0,10)->FPC02FE0TQ0(00)	Down
FPC02FE1TQ1(00)->S00F0_0(6,1,06)	Down	S00F0_0(2,1,10)->FPC02FE1TQ1(00)	Down
FPC04FE0TQ0(00)->S00F0_0(5,0,05)	Down	S00F0_0(1,0,09)->FPC04FE0TQ0(00)	Down
FPC04FE1TQ1(00)->S00F0_0(5,1,05)	Down	S00F0_0(1,1,09)->FPC04FE1TQ1(00)	Down
FPC06FE0TQ0(00)->S00F0_0(4,0,04)	Down	S00F0_0(0,0,08)->FPC06FE0TQ0(00)	Down
FPC06FE1TQ1(00)->S00F0_0(4,1,04)	Down	S00F0_0(0,1,08)->FPC06FE1TQ1(00)	Down
FPC08FE0TQ0(00)->S00F0_0(4,2,04)	OK	S00F0_0(0,2,08)->FPC08FE0TQ0(00)	OK
FPC08FE1TQ1(00)->S00F0_0(4,3,04)	OK	S00F0_0(0,3,08)->FPC08FE1TQ1(00)	OK
FPC10FE0TQ0(00)->S00F0_0(5,2,05)	Down	S00F0_0(1,2,09)->FPC10FE0TQ0(00)	Down
FPC10FE1TQ1(00)->S00F0_0(5,3,05)	Down	S00F0_0(1,3,09)->FPC10FE1TQ1(00)	Down
FPC12FE0TQ0(00)->S00F0_0(7,2,07)	OK	S00F0_0(3,2,11)->FPC12FE0TQ0(00)	OK
FPC12FE1TQ1(00)->S00F0_0(7,3,07)	OK	S00F0_0(3,3,11)->FPC12FE1TQ1(00)	OK

```

FPC14FE0TQ0(00)->S00F0_0(7,4,07) Down    S00F0_0(3,4,11)->FPC14FE0TQ0(00) Down
FPC14FE1TQ1(00)->S00F0_0(7,5,07) Down    S00F0_0(3,5,11)->FPC14FE1TQ1(00) Down

```

SIB 0 FCHIP 0 FCORE 1 :

In-links	State	Out-links	State
FPC00FE0TQ0(01)->S00F0_1(3,0,11)	Down	S00F0_1(7,0,07)->FPC00FE0TQ0(01)	Down
FPC00FE1TQ1(01)->S00F0_1(3,1,11)	Down	S00F0_1(7,1,07)->FPC00FE1TQ1(01)	Down
FPC02FE0TQ0(01)->S00F0_1(2,0,10)	Down	S00F0_1(6,0,06)->FPC02FE0TQ0(01)	Down
FPC02FE1TQ1(01)->S00F0_1(2,1,10)	Down	S00F0_1(6,1,06)->FPC02FE1TQ1(01)	Down
FPC04FE0TQ0(01)->S00F0_1(1,0,09)	Down	S00F0_1(4,0,04)->FPC04FE0TQ0(01)	Down
FPC04FE1TQ1(01)->S00F0_1(1,1,09)	Down	S00F0_1(4,1,04)->FPC04FE1TQ1(01)	Down
FPC06FE0TQ0(01)->S00F0_1(0,0,08)	Down	S00F0_1(4,2,04)->FPC06FE0TQ0(01)	Down
FPC06FE1TQ1(01)->S00F0_1(0,1,08)	Down	S00F0_1(4,3,04)->FPC06FE1TQ1(01)	Down
FPC08FE0TQ0(01)->S00F0_1(0,2,08)	OK	S00F0_1(4,4,04)->FPC08FE0TQ0(01)	OK
FPC08FE1TQ1(01)->S00F0_1(0,3,08)	OK	S00F0_1(4,5,04)->FPC08FE1TQ1(01)	OK
FPC10FE0TQ0(01)->S00F0_1(1,2,09)	Down	S00F0_1(5,0,05)->FPC10FE0TQ0(01)	Down
FPC10FE1TQ1(01)->S00F0_1(1,3,09)	Down	S00F0_1(5,1,05)->FPC10FE1TQ1(01)	Down
FPC12FE0TQ0(01)->S00F0_1(2,2,10)	OK	S00F0_1(6,2,06)->FPC12FE0TQ0(01)	OK
FPC12FE1TQ1(01)->S00F0_1(2,3,10)	OK	S00F0_1(6,3,06)->FPC12FE1TQ1(01)	OK
FPC14FE0TQ0(01)->S00F0_1(3,2,11)	Down	S00F0_1(7,2,07)->FPC14FE0TQ0(01)	Down
FPC14FE1TQ1(01)->S00F0_1(3,3,11)	Down	S00F0_1(7,3,07)->FPC14FE1TQ1(01)	Down

SIB 1 FCHIP 0 FCORE 0 :

In-links	State	Out-links	State
FPC00FE0TQ0(02)->S01F0_0(7,0,07)	Down	S01F0_0(3,0,11)->FPC00FE0TQ0(02)	Down
FPC00FE1TQ1(02)->S01F0_0(7,1,07)	Down	S01F0_0(3,1,11)->FPC00FE1TQ1(02)	Down
FPC02FE0TQ0(02)->S01F0_0(6,0,06)	Down	S01F0_0(2,0,10)->FPC02FE0TQ0(02)	Down
FPC02FE1TQ1(02)->S01F0_0(6,1,06)	Down	S01F0_0(2,1,10)->FPC02FE1TQ1(02)	Down

---(more)---[abort]

regress@raki> show chassis fabric topology | no-more

In-link : FPC# FE# TQ# (TQ-TX sub-chnl #) ->
SIB# TF#_FCORE# (TF-RX port#, TF-RX sub-chn#, TF-RX inst#)

Out-link : SIB# TF#_FCORE# (TF-TX port#, TF-TX sub-chn#, TF-TX inst#) ->
FPC# FE# TQ# (TQ-RX sub-chnl #)

SIB 0 FCHIP 0 FCORE 0 :

In-links	State	Out-links	State
FPC00FE0TQ0(00)->S00F0_0(7,0,07)	Down	S00F0_0(3,0,11)->FPC00FE0TQ0(00)	Down
FPC00FE1TQ1(00)->S00F0_0(7,1,07)	Down	S00F0_0(3,1,11)->FPC00FE1TQ1(00)	Down
FPC02FE0TQ0(00)->S00F0_0(6,0,06)	Down	S00F0_0(2,0,10)->FPC02FE0TQ0(00)	Down
FPC02FE1TQ1(00)->S00F0_0(6,1,06)	Down	S00F0_0(2,1,10)->FPC02FE1TQ1(00)	Down
FPC04FE0TQ0(00)->S00F0_0(5,0,05)	Down	S00F0_0(1,0,09)->FPC04FE0TQ0(00)	Down
FPC04FE1TQ1(00)->S00F0_0(5,1,05)	Down	S00F0_0(1,1,09)->FPC04FE1TQ1(00)	Down
FPC06FE0TQ0(00)->S00F0_0(4,0,04)	Down	S00F0_0(0,0,08)->FPC06FE0TQ0(00)	Down
FPC06FE1TQ1(00)->S00F0_0(4,1,04)	Down	S00F0_0(0,1,08)->FPC06FE1TQ1(00)	Down
FPC08FE0TQ0(00)->S00F0_0(4,2,04)	OK	S00F0_0(0,2,08)->FPC08FE0TQ0(00)	OK
FPC08FE1TQ1(00)->S00F0_0(4,3,04)	OK	S00F0_0(0,3,08)->FPC08FE1TQ1(00)	OK
FPC10FE0TQ0(00)->S00F0_0(5,2,05)	Down	S00F0_0(1,2,09)->FPC10FE0TQ0(00)	Down
FPC10FE1TQ1(00)->S00F0_0(5,3,05)	Down	S00F0_0(1,3,09)->FPC10FE1TQ1(00)	Down
FPC12FE0TQ0(00)->S00F0_0(7,2,07)	OK	S00F0_0(3,2,11)->FPC12FE0TQ0(00)	OK
FPC12FE1TQ1(00)->S00F0_0(7,3,07)	OK	S00F0_0(3,3,11)->FPC12FE1TQ1(00)	OK
FPC14FE0TQ0(00)->S00F0_0(7,4,07)	Down	S00F0_0(3,4,11)->FPC14FE0TQ0(00)	Down
FPC14FE1TQ1(00)->S00F0_0(7,5,07)	Down	S00F0_0(3,5,11)->FPC14FE1TQ1(00)	Down

SIB 0 FCHIP 0 FCORE 1 :

In-links	State	Out-links	State
FPC00FE0TQ0(01)->S00F0_1(3,0,11)	Down	S00F0_1(7,0,07)->FPC00FE0TQ0(01)	Down
FPC00FE1TQ1(01)->S00F0_1(3,1,11)	Down	S00F0_1(7,1,07)->FPC00FE1TQ1(01)	Down
FPC02FE0TQ0(01)->S00F0_1(2,0,10)	Down	S00F0_1(6,0,06)->FPC02FE0TQ0(01)	Down
FPC02FE1TQ1(01)->S00F0_1(2,1,10)	Down	S00F0_1(6,1,06)->FPC02FE1TQ1(01)	Down
FPC04FE0TQ0(01)->S00F0_1(1,0,09)	Down	S00F0_1(4,0,04)->FPC04FE0TQ0(01)	Down
FPC04FE1TQ1(01)->S00F0_1(1,1,09)	Down	S00F0_1(4,1,04)->FPC04FE1TQ1(01)	Down
FPC06FE0TQ0(01)->S00F0_1(0,0,08)	Down	S00F0_1(4,2,04)->FPC06FE0TQ0(01)	Down
FPC06FE1TQ1(01)->S00F0_1(0,1,08)	Down	S00F0_1(4,3,04)->FPC06FE1TQ1(01)	Down
FPC08FE0TQ0(01)->S00F0_1(0,2,08)	OK	S00F0_1(4,4,04)->FPC08FE0TQ0(01)	OK
FPC08FE1TQ1(01)->S00F0_1(0,3,08)	OK	S00F0_1(4,5,04)->FPC08FE1TQ1(01)	OK
FPC10FE0TQ0(01)->S00F0_1(1,2,09)	Down	S00F0_1(5,0,05)->FPC10FE0TQ0(01)	Down
FPC10FE1TQ1(01)->S00F0_1(1,3,09)	Down	S00F0_1(5,1,05)->FPC10FE1TQ1(01)	Down
FPC12FE0TQ0(01)->S00F0_1(2,2,10)	OK	S00F0_1(6,2,06)->FPC12FE0TQ0(01)	OK
FPC12FE1TQ1(01)->S00F0_1(2,3,10)	OK	S00F0_1(6,3,06)->FPC12FE1TQ1(01)	OK
FPC14FE0TQ0(01)->S00F0_1(3,2,11)	Down	S00F0_1(7,2,07)->FPC14FE0TQ0(01)	Down
FPC14FE1TQ1(01)->S00F0_1(3,3,11)	Down	S00F0_1(7,3,07)->FPC14FE1TQ1(01)	Down

SIB 1 FCHIP 0 FCORE 0 :

In-links	State	Out-links	State
FPC00FE0TQ0(02)->S01F0_0(7,0,07)	Down	S01F0_0(3,0,11)->FPC00FE0TQ0(02)	Down
FPC00FE1TQ1(02)->S01F0_0(7,1,07)	Down	S01F0_0(3,1,11)->FPC00FE1TQ1(02)	Down
FPC02FE0TQ0(02)->S01F0_0(6,0,06)	Down	S01F0_0(2,0,10)->FPC02FE0TQ0(02)	Down
FPC02FE1TQ1(02)->S01F0_0(6,1,06)	Down	S01F0_0(2,1,10)->FPC02FE1TQ1(02)	Down
FPC04FE0TQ0(02)->S01F0_0(5,0,05)	Down	S01F0_0(1,0,09)->FPC04FE0TQ0(02)	Down
FPC04FE1TQ1(02)->S01F0_0(5,1,05)	Down	S01F0_0(1,1,09)->FPC04FE1TQ1(02)	Down
FPC06FE0TQ0(02)->S01F0_0(4,0,04)	Down	S01F0_0(0,0,08)->FPC06FE0TQ0(02)	Down
FPC06FE1TQ1(02)->S01F0_0(4,1,04)	Down	S01F0_0(0,1,08)->FPC06FE1TQ1(02)	Down
FPC08FE0TQ0(02)->S01F0_0(4,2,04)	OK	S01F0_0(0,2,08)->FPC08FE0TQ0(02)	OK
FPC08FE1TQ1(02)->S01F0_0(4,3,04)	OK	S01F0_0(0,3,08)->FPC08FE1TQ1(02)	OK
FPC10FE0TQ0(02)->S01F0_0(5,2,05)	Down	S01F0_0(1,2,09)->FPC10FE0TQ0(02)	Down
FPC10FE1TQ1(02)->S01F0_0(5,3,05)	Down	S01F0_0(1,3,09)->FPC10FE1TQ1(02)	Down
FPC12FE0TQ0(02)->S01F0_0(7,2,07)	OK	S01F0_0(3,2,11)->FPC12FE0TQ0(02)	OK
FPC12FE1TQ1(02)->S01F0_0(7,3,07)	OK	S01F0_0(3,3,11)->FPC12FE1TQ1(02)	OK
FPC14FE0TQ0(02)->S01F0_0(7,4,07)	Down	S01F0_0(3,4,11)->FPC14FE0TQ0(02)	Down
FPC14FE1TQ1(02)->S01F0_0(7,5,07)	Down	S01F0_0(3,5,11)->FPC14FE1TQ1(02)	Down

SIB 1 FCHIP 0 FCORE 1 :

In-links	State	Out-links	State
FPC00FE0TQ0(03)->S01F0_1(3,0,11)	Down	S01F0_1(7,0,07)->FPC00FE0TQ0(03)	Down
FPC00FE1TQ1(03)->S01F0_1(3,1,11)	Down	S01F0_1(7,1,07)->FPC00FE1TQ1(03)	Down
FPC02FE0TQ0(03)->S01F0_1(2,0,10)	Down	S01F0_1(6,0,06)->FPC02FE0TQ0(03)	Down
FPC02FE1TQ1(03)->S01F0_1(2,1,10)	Down	S01F0_1(6,1,06)->FPC02FE1TQ1(03)	Down
FPC04FE0TQ0(03)->S01F0_1(1,0,09)	Down	S01F0_1(4,0,04)->FPC04FE0TQ0(03)	Down
FPC04FE1TQ1(03)->S01F0_1(1,1,09)	Down	S01F0_1(4,1,04)->FPC04FE1TQ1(03)	Down
FPC06FE0TQ0(03)->S01F0_1(0,0,08)	Down	S01F0_1(4,2,04)->FPC06FE0TQ0(03)	Down
FPC06FE1TQ1(03)->S01F0_1(0,1,08)	Down	S01F0_1(4,3,04)->FPC06FE1TQ1(03)	Down
FPC08FE0TQ0(03)->S01F0_1(0,2,08)	OK	S01F0_1(4,4,04)->FPC08FE0TQ0(03)	OK
FPC08FE1TQ1(03)->S01F0_1(0,3,08)	OK	S01F0_1(4,5,04)->FPC08FE1TQ1(03)	OK
FPC10FE0TQ0(03)->S01F0_1(1,2,09)	Down	S01F0_1(5,0,05)->FPC10FE0TQ0(03)	Down
FPC10FE1TQ1(03)->S01F0_1(1,3,09)	Down	S01F0_1(5,1,05)->FPC10FE1TQ1(03)	Down
FPC12FE0TQ0(03)->S01F0_1(2,2,10)	OK	S01F0_1(6,2,06)->FPC12FE0TQ0(03)	OK
FPC12FE1TQ1(03)->S01F0_1(2,3,10)	OK	S01F0_1(6,3,06)->FPC12FE1TQ1(03)	OK

```

FPC14FE0TQ0(03)->S01F0_1(3,2,11) Down    S01F0_1(7,2,07)->FPC14FE0TQ0(03) Down
FPC14FE1TQ1(03)->S01F0_1(3,3,11) Down    S01F0_1(7,3,07)->FPC14FE1TQ1(03) Down

```

SIB 2 FCHIP 0 FCORE 0 :

In-links	State	Out-links	State
FPC00FE0TQ0(04)->S02F0_0(7,0,07)	Down	S02F0_0(3,0,11)->FPC00FE0TQ0(04)	Down
FPC00FE1TQ1(04)->S02F0_0(7,1,07)	Down	S02F0_0(3,1,11)->FPC00FE1TQ1(04)	Down
FPC02FE0TQ0(04)->S02F0_0(6,0,06)	Down	S02F0_0(2,0,10)->FPC02FE0TQ0(04)	Down
FPC02FE1TQ1(04)->S02F0_0(6,1,06)	Down	S02F0_0(2,1,10)->FPC02FE1TQ1(04)	Down
FPC04FE0TQ0(04)->S02F0_0(5,0,05)	Down	S02F0_0(1,0,09)->FPC04FE0TQ0(04)	Down
FPC04FE1TQ1(04)->S02F0_0(5,1,05)	Down	S02F0_0(1,1,09)->FPC04FE1TQ1(04)	Down
FPC06FE0TQ0(04)->S02F0_0(4,0,04)	Down	S02F0_0(0,0,08)->FPC06FE0TQ0(04)	Down
FPC06FE1TQ1(04)->S02F0_0(4,1,04)	Down	S02F0_0(0,1,08)->FPC06FE1TQ1(04)	Down
FPC08FE0TQ0(04)->S02F0_0(4,2,04)	OK	S02F0_0(0,2,08)->FPC08FE0TQ0(04)	OK
FPC08FE1TQ1(04)->S02F0_0(4,3,04)	OK	S02F0_0(0,3,08)->FPC08FE1TQ1(04)	OK
FPC10FE0TQ0(04)->S02F0_0(5,2,05)	Down	S02F0_0(1,2,09)->FPC10FE0TQ0(04)	Down
FPC10FE1TQ1(04)->S02F0_0(5,3,05)	Down	S02F0_0(1,3,09)->FPC10FE1TQ1(04)	Down
FPC12FE0TQ0(04)->S02F0_0(7,2,07)	OK	S02F0_0(3,2,11)->FPC12FE0TQ0(04)	OK
FPC12FE1TQ1(04)->S02F0_0(7,3,07)	OK	S02F0_0(3,3,11)->FPC12FE1TQ1(04)	OK
FPC14FE0TQ0(04)->S02F0_0(7,4,07)	Down	S02F0_0(3,4,11)->FPC14FE0TQ0(04)	Down
FPC14FE1TQ1(04)->S02F0_0(7,5,07)	Down	S02F0_0(3,5,11)->FPC14FE1TQ1(04)	Down

SIB 2 FCHIP 0 FCORE 1 :

In-links	State	Out-links	State
FPC00FE0TQ0(05)->S02F0_1(3,0,11)	Down	S02F0_1(7,0,07)->FPC00FE0TQ0(05)	Down
FPC00FE1TQ1(05)->S02F0_1(3,1,11)	Down	S02F0_1(7,1,07)->FPC00FE1TQ1(05)	Down
FPC02FE0TQ0(05)->S02F0_1(2,0,10)	Down	S02F0_1(6,0,06)->FPC02FE0TQ0(05)	Down
FPC02FE1TQ1(05)->S02F0_1(2,1,10)	Down	S02F0_1(6,1,06)->FPC02FE1TQ1(05)	Down
FPC04FE0TQ0(05)->S02F0_1(1,0,09)	Down	S02F0_1(4,0,04)->FPC04FE0TQ0(05)	Down
FPC04FE1TQ1(05)->S02F0_1(1,1,09)	Down	S02F0_1(4,1,04)->FPC04FE1TQ1(05)	Down
FPC06FE0TQ0(05)->S02F0_1(0,0,08)	Down	S02F0_1(4,2,04)->FPC06FE0TQ0(05)	Down
FPC06FE1TQ1(05)->S02F0_1(0,1,08)	Down	S02F0_1(4,3,04)->FPC06FE1TQ1(05)	Down
FPC08FE0TQ0(05)->S02F0_1(0,2,08)	OK	S02F0_1(4,4,04)->FPC08FE0TQ0(05)	OK
FPC08FE1TQ1(05)->S02F0_1(0,3,08)	OK	S02F0_1(4,5,04)->FPC08FE1TQ1(05)	OK
FPC10FE0TQ0(05)->S02F0_1(1,2,09)	Down	S02F0_1(5,0,05)->FPC10FE0TQ0(05)	Down
FPC10FE1TQ1(05)->S02F0_1(1,3,09)	Down	S02F0_1(5,1,05)->FPC10FE1TQ1(05)	Down
FPC12FE0TQ0(05)->S02F0_1(2,2,10)	OK	S02F0_1(6,2,06)->FPC12FE0TQ0(05)	OK
FPC12FE1TQ1(05)->S02F0_1(2,3,10)	OK	S02F0_1(6,3,06)->FPC12FE1TQ1(05)	OK
FPC14FE0TQ0(05)->S02F0_1(3,2,11)	Down	S02F0_1(7,2,07)->FPC14FE0TQ0(05)	Down
FPC14FE1TQ1(05)->S02F0_1(3,3,11)	Down	S02F0_1(7,3,07)->FPC14FE1TQ1(05)	Down

SIB 3 FCHIP 0 FCORE 0 :

In-links	State	Out-links	State
FPC00FE0TQ0(06)->S03F0_0(7,0,07)	Down	S03F0_0(3,0,11)->FPC00FE0TQ0(06)	Down
FPC00FE1TQ1(06)->S03F0_0(7,1,07)	Down	S03F0_0(3,1,11)->FPC00FE1TQ1(06)	Down
FPC02FE0TQ0(06)->S03F0_0(6,0,06)	Down	S03F0_0(2,0,10)->FPC02FE0TQ0(06)	Down
FPC02FE1TQ1(06)->S03F0_0(6,1,06)	Down	S03F0_0(2,1,10)->FPC02FE1TQ1(06)	Down
FPC04FE0TQ0(06)->S03F0_0(5,0,05)	Down	S03F0_0(1,0,09)->FPC04FE0TQ0(06)	Down
FPC04FE1TQ1(06)->S03F0_0(5,1,05)	Down	S03F0_0(1,1,09)->FPC04FE1TQ1(06)	Down
FPC06FE0TQ0(06)->S03F0_0(4,0,04)	Down	S03F0_0(0,0,08)->FPC06FE0TQ0(06)	Down
FPC06FE1TQ1(06)->S03F0_0(4,1,04)	Down	S03F0_0(0,1,08)->FPC06FE1TQ1(06)	Down
FPC08FE0TQ0(06)->S03F0_0(4,2,04)	OK	S03F0_0(0,2,08)->FPC08FE0TQ0(06)	OK
FPC08FE1TQ1(06)->S03F0_0(4,3,04)	OK	S03F0_0(0,3,08)->FPC08FE1TQ1(06)	OK
FPC10FE0TQ0(06)->S03F0_0(5,2,05)	Down	S03F0_0(1,2,09)->FPC10FE0TQ0(06)	Down
FPC10FE1TQ1(06)->S03F0_0(5,3,05)	Down	S03F0_0(1,3,09)->FPC10FE1TQ1(06)	Down


```

FPC12FE0TQ0(06)->S03F0_0(7,2,07) OK      S03F0_0(3,2,11)->FPC12FE0TQ0(06) OK
FPC12FE1TQ1(06)->S03F0_0(7,3,07) OK      S03F0_0(3,3,11)->FPC12FE1TQ1(06) OK
FPC14FE0TQ0(06)->S03F0_0(7,4,07) Down    S03F0_0(3,4,11)->FPC14FE0TQ0(06) Down
FPC14FE1TQ1(06)->S03F0_0(7,5,07) Down    S03F0_0(3,5,11)->FPC14FE1TQ1(06) Down

```

SIB 3 FCHIP 0 FCORE 1 :

In-links	State	Out-links	State
FPC00FE0TQ0(07)->S03F0_1(3,0,11)	Down	S03F0_1(7,0,07)->FPC00FE0TQ0(07)	Down
FPC00FE1TQ1(07)->S03F0_1(3,1,11)	Down	S03F0_1(7,1,07)->FPC00FE1TQ1(07)	Down
FPC02FE0TQ0(07)->S03F0_1(2,0,10)	Down	S03F0_1(6,0,06)->FPC02FE0TQ0(07)	Down
FPC02FE1TQ1(07)->S03F0_1(2,1,10)	Down	S03F0_1(6,1,06)->FPC02FE1TQ1(07)	Down
FPC04FE0TQ0(07)->S03F0_1(1,0,09)	Down	S03F0_1(4,0,04)->FPC04FE0TQ0(07)	Down
FPC04FE1TQ1(07)->S03F0_1(1,1,09)	Down	S03F0_1(4,1,04)->FPC04FE1TQ1(07)	Down
FPC06FE0TQ0(07)->S03F0_1(0,0,08)	Down	S03F0_1(4,2,04)->FPC06FE0TQ0(07)	Down
FPC06FE1TQ1(07)->S03F0_1(0,1,08)	Down	S03F0_1(4,3,04)->FPC06FE1TQ1(07)	Down
FPC08FE0TQ0(07)->S03F0_1(0,2,08)	OK	S03F0_1(4,4,04)->FPC08FE0TQ0(07)	OK
FPC08FE1TQ1(07)->S03F0_1(0,3,08)	OK	S03F0_1(4,5,04)->FPC08FE1TQ1(07)	OK
FPC10FE0TQ0(07)->S03F0_1(1,2,09)	Down	S03F0_1(5,0,05)->FPC10FE0TQ0(07)	Down
FPC10FE1TQ1(07)->S03F0_1(1,3,09)	Down	S03F0_1(5,1,05)->FPC10FE1TQ1(07)	Down
FPC12FE0TQ0(07)->S03F0_1(2,2,10)	OK	S03F0_1(6,2,06)->FPC12FE0TQ0(07)	OK
FPC12FE1TQ1(07)->S03F0_1(2,3,10)	OK	S03F0_1(6,3,06)->FPC12FE1TQ1(07)	OK
FPC14FE0TQ0(07)->S03F0_1(3,2,11)	Down	S03F0_1(7,2,07)->FPC14FE0TQ0(07)	Down
FPC14FE1TQ1(07)->S03F0_1(3,3,11)	Down	S03F0_1(7,3,07)->FPC14FE1TQ1(07)	Down

SIB 4 FCHIP 0 FCORE 0 :

In-links	State	Out-links	State
FPC00FE0TQ0(08)->S04F0_0(7,0,07)	Down	S04F0_0(3,0,11)->FPC00FE0TQ0(08)	Down
FPC00FE1TQ1(08)->S04F0_0(7,1,07)	Down	S04F0_0(3,1,11)->FPC00FE1TQ1(08)	Down
FPC02FE0TQ0(08)->S04F0_0(6,0,06)	Down	S04F0_0(2,0,10)->FPC02FE0TQ0(08)	Down
FPC02FE1TQ1(08)->S04F0_0(6,1,06)	Down	S04F0_0(2,1,10)->FPC02FE1TQ1(08)	Down
FPC04FE0TQ0(08)->S04F0_0(5,0,05)	Down	S04F0_0(1,0,09)->FPC04FE0TQ0(08)	Down
FPC04FE1TQ1(08)->S04F0_0(5,1,05)	Down	S04F0_0(1,1,09)->FPC04FE1TQ1(08)	Down
FPC06FE0TQ0(08)->S04F0_0(4,0,04)	Down	S04F0_0(0,0,08)->FPC06FE0TQ0(08)	Down
FPC06FE1TQ1(08)->S04F0_0(4,1,04)	Down	S04F0_0(0,1,08)->FPC06FE1TQ1(08)	Down
FPC08FE0TQ0(08)->S04F0_0(4,2,04)	OK	S04F0_0(0,2,08)->FPC08FE0TQ0(08)	OK
FPC08FE1TQ1(08)->S04F0_0(4,3,04)	OK	S04F0_0(0,3,08)->FPC08FE1TQ1(08)	OK
FPC10FE0TQ0(08)->S04F0_0(5,2,05)	Down	S04F0_0(1,2,09)->FPC10FE0TQ0(08)	Down
FPC10FE1TQ1(08)->S04F0_0(5,3,05)	Down	S04F0_0(1,3,09)->FPC10FE1TQ1(08)	Down
FPC12FE0TQ0(08)->S04F0_0(7,2,07)	OK	S04F0_0(3,2,11)->FPC12FE0TQ0(08)	OK
FPC12FE1TQ1(08)->S04F0_0(7,3,07)	OK	S04F0_0(3,3,11)->FPC12FE1TQ1(08)	OK
FPC14FE0TQ0(08)->S04F0_0(7,4,07)	Down	S04F0_0(3,4,11)->FPC14FE0TQ0(08)	Down
FPC14FE1TQ1(08)->S04F0_0(7,5,07)	Down	S04F0_0(3,5,11)->FPC14FE1TQ1(08)	Down

SIB 4 FCHIP 0 FCORE 1 :

In-links	State	Out-links	State
FPC00FE0TQ0(09)->S04F0_1(3,0,11)	Down	S04F0_1(7,0,07)->FPC00FE0TQ0(09)	Down
FPC00FE1TQ1(09)->S04F0_1(3,1,11)	Down	S04F0_1(7,1,07)->FPC00FE1TQ1(09)	Down
FPC02FE0TQ0(09)->S04F0_1(2,0,10)	Down	S04F0_1(6,0,06)->FPC02FE0TQ0(09)	Down
FPC02FE1TQ1(09)->S04F0_1(2,1,10)	Down	S04F0_1(6,1,06)->FPC02FE1TQ1(09)	Down
FPC04FE0TQ0(09)->S04F0_1(1,0,09)	Down	S04F0_1(4,0,04)->FPC04FE0TQ0(09)	Down
FPC04FE1TQ1(09)->S04F0_1(1,1,09)	Down	S04F0_1(4,1,04)->FPC04FE1TQ1(09)	Down
FPC06FE0TQ0(09)->S04F0_1(0,0,08)	Down	S04F0_1(4,2,04)->FPC06FE0TQ0(09)	Down
FPC06FE1TQ1(09)->S04F0_1(0,1,08)	Down	S04F0_1(4,3,04)->FPC06FE1TQ1(09)	Down
FPC08FE0TQ0(09)->S04F0_1(0,2,08)	OK	S04F0_1(4,4,04)->FPC08FE0TQ0(09)	OK
FPC08FE1TQ1(09)->S04F0_1(0,3,08)	OK	S04F0_1(4,5,04)->FPC08FE1TQ1(09)	OK

```

FPC10FE0TQ0(09)->S04F0_1(1,2,09) Down    S04F0_1(5,0,05)->FPC10FE0TQ0(09) Down
FPC10FE1TQ1(09)->S04F0_1(1,3,09) Down    S04F0_1(5,1,05)->FPC10FE1TQ1(09) Down
FPC12FE0TQ0(09)->S04F0_1(2,2,10) OK      S04F0_1(6,2,06)->FPC12FE0TQ0(09) OK
FPC12FE1TQ1(09)->S04F0_1(2,3,10) OK      S04F0_1(6,3,06)->FPC12FE1TQ1(09) OK
FPC14FE0TQ0(09)->S04F0_1(3,2,11) Down    S04F0_1(7,2,07)->FPC14FE0TQ0(09) Down
FPC14FE1TQ1(09)->S04F0_1(3,3,11) Down    S04F0_1(7,3,07)->FPC14FE1TQ1(09) Down

```

SIB 5 FCHIP 0 FCORE 0 :

In-links	State	Out-links	State
FPC00FE0TQ0(10)->S05F0_0(7,0,07)	Down	S05F0_0(3,0,11)->FPC00FE0TQ0(10)	Down
FPC00FE1TQ1(10)->S05F0_0(7,1,07)	Down	S05F0_0(3,1,11)->FPC00FE1TQ1(10)	Down
FPC02FE0TQ0(10)->S05F0_0(6,0,06)	Down	S05F0_0(2,0,10)->FPC02FE0TQ0(10)	Down
FPC02FE1TQ1(10)->S05F0_0(6,1,06)	Down	S05F0_0(2,1,10)->FPC02FE1TQ1(10)	Down
FPC04FE0TQ0(10)->S05F0_0(5,0,05)	Down	S05F0_0(1,0,09)->FPC04FE0TQ0(10)	Down
FPC04FE1TQ1(10)->S05F0_0(5,1,05)	Down	S05F0_0(1,1,09)->FPC04FE1TQ1(10)	Down
FPC06FE0TQ0(10)->S05F0_0(4,0,04)	Down	S05F0_0(0,0,08)->FPC06FE0TQ0(10)	Down
FPC06FE1TQ1(10)->S05F0_0(4,1,04)	Down	S05F0_0(0,1,08)->FPC06FE1TQ1(10)	Down
FPC08FE0TQ0(10)->S05F0_0(4,2,04)	OK	S05F0_0(0,2,08)->FPC08FE0TQ0(10)	OK
FPC08FE1TQ1(10)->S05F0_0(4,3,04)	OK	S05F0_0(0,3,08)->FPC08FE1TQ1(10)	OK
FPC10FE0TQ0(10)->S05F0_0(5,2,05)	Down	S05F0_0(1,2,09)->FPC10FE0TQ0(10)	Down
FPC10FE1TQ1(10)->S05F0_0(5,3,05)	Down	S05F0_0(1,3,09)->FPC10FE1TQ1(10)	Down
FPC12FE0TQ0(10)->S05F0_0(7,2,07)	OK	S05F0_0(3,2,11)->FPC12FE0TQ0(10)	OK
FPC12FE1TQ1(10)->S05F0_0(7,3,07)	OK	S05F0_0(3,3,11)->FPC12FE1TQ1(10)	OK
FPC14FE0TQ0(10)->S05F0_0(7,4,07)	Down	S05F0_0(3,4,11)->FPC14FE0TQ0(10)	Down
FPC14FE1TQ1(10)->S05F0_0(7,5,07)	Down	S05F0_0(3,5,11)->FPC14FE1TQ1(10)	Down

SIB 5 FCHIP 0 FCORE 1 :

In-links	State	Out-links	State
FPC00FE0TQ0(11)->S05F0_1(3,0,11)	Down	S05F0_1(7,0,07)->FPC00FE0TQ0(11)	Down
FPC00FE1TQ1(11)->S05F0_1(3,1,11)	Down	S05F0_1(7,1,07)->FPC00FE1TQ1(11)	Down
FPC02FE0TQ0(11)->S05F0_1(2,0,10)	Down	S05F0_1(6,0,06)->FPC02FE0TQ0(11)	Down
FPC02FE1TQ1(11)->S05F0_1(2,1,10)	Down	S05F0_1(6,1,06)->FPC02FE1TQ1(11)	Down
FPC04FE0TQ0(11)->S05F0_1(1,0,09)	Down	S05F0_1(4,0,04)->FPC04FE0TQ0(11)	Down
FPC04FE1TQ1(11)->S05F0_1(1,1,09)	Down	S05F0_1(4,1,04)->FPC04FE1TQ1(11)	Down
FPC06FE0TQ0(11)->S05F0_1(0,0,08)	Down	S05F0_1(4,2,04)->FPC06FE0TQ0(11)	Down
FPC06FE1TQ1(11)->S05F0_1(0,1,08)	Down	S05F0_1(4,3,04)->FPC06FE1TQ1(11)	Down
FPC08FE0TQ0(11)->S05F0_1(0,2,08)	OK	S05F0_1(4,4,04)->FPC08FE0TQ0(11)	OK
FPC08FE1TQ1(11)->S05F0_1(0,3,08)	OK	S05F0_1(4,5,04)->FPC08FE1TQ1(11)	OK
FPC10FE0TQ0(11)->S05F0_1(1,2,09)	Down	S05F0_1(5,0,05)->FPC10FE0TQ0(11)	Down
FPC10FE1TQ1(11)->S05F0_1(1,3,09)	Down	S05F0_1(5,1,05)->FPC10FE1TQ1(11)	Down
FPC12FE0TQ0(11)->S05F0_1(2,2,10)	OK	S05F0_1(6,2,06)->FPC12FE0TQ0(11)	OK
FPC12FE1TQ1(11)->S05F0_1(2,3,10)	OK	S05F0_1(6,3,06)->FPC12FE1TQ1(11)	OK
FPC14FE0TQ0(11)->S05F0_1(3,2,11)	Down	S05F0_1(7,2,07)->FPC14FE0TQ0(11)	Down
FPC14FE1TQ1(11)->S05F0_1(3,3,11)	Down	S05F0_1(7,3,07)->FPC14FE1TQ1(11)	Down

SIB 6 FCHIP 0 FCORE 0 :

In-links	State	Out-links	State
FPC00FE0TQ0(12)->S06F0_0(7,0,07)	Down	S06F0_0(3,0,11)->FPC00FE0TQ0(12)	Down
FPC00FE1TQ1(12)->S06F0_0(7,1,07)	Down	S06F0_0(3,1,11)->FPC00FE1TQ1(12)	Down
FPC02FE0TQ0(12)->S06F0_0(6,0,06)	Down	S06F0_0(2,0,10)->FPC02FE0TQ0(12)	Down
FPC02FE1TQ1(12)->S06F0_0(6,1,06)	Down	S06F0_0(2,1,10)->FPC02FE1TQ1(12)	Down
FPC04FE0TQ0(12)->S06F0_0(5,0,05)	Down	S06F0_0(1,0,09)->FPC04FE0TQ0(12)	Down
FPC04FE1TQ1(12)->S06F0_0(5,1,05)	Down	S06F0_0(1,1,09)->FPC04FE1TQ1(12)	Down
FPC06FE0TQ0(12)->S06F0_0(4,0,04)	Down	S06F0_0(0,0,08)->FPC06FE0TQ0(12)	Down
FPC06FE1TQ1(12)->S06F0_0(4,1,04)	Down	S06F0_0(0,1,08)->FPC06FE1TQ1(12)	Down

```

FPC08FE0TQ0(12)->S06F0_0(4,2,04) OK      S06F0_0(0,2,08)->FPC08FE0TQ0(12) OK
FPC08FE1TQ1(12)->S06F0_0(4,3,04) OK      S06F0_0(0,3,08)->FPC08FE1TQ1(12) OK
FPC10FE0TQ0(12)->S06F0_0(5,2,05) Down    S06F0_0(1,2,09)->FPC10FE0TQ0(12) Down
FPC10FE1TQ1(12)->S06F0_0(5,3,05) Down    S06F0_0(1,3,09)->FPC10FE1TQ1(12) Down
FPC12FE0TQ0(12)->S06F0_0(7,2,07) OK      S06F0_0(3,2,11)->FPC12FE0TQ0(12) OK
FPC12FE1TQ1(12)->S06F0_0(7,3,07) OK      S06F0_0(3,3,11)->FPC12FE1TQ1(12) OK
FPC14FE0TQ0(12)->S06F0_0(7,4,07) Down    S06F0_0(3,4,11)->FPC14FE0TQ0(12) Down
FPC14FE1TQ1(12)->S06F0_0(7,5,07) Down    S06F0_0(3,5,11)->FPC14FE1TQ1(12) Down

```

SIB 6 FCHIP 0 FCORE 1 :

In-links	State	Out-links	State
FPC00FE0TQ0(13)->S06F0_1(3,0,11)	Down	S06F0_1(7,0,07)->FPC00FE0TQ0(13)	Down
FPC00FE1TQ1(13)->S06F0_1(3,1,11)	Down	S06F0_1(7,1,07)->FPC00FE1TQ1(13)	Down
FPC02FE0TQ0(13)->S06F0_1(2,0,10)	Down	S06F0_1(6,0,06)->FPC02FE0TQ0(13)	Down
FPC02FE1TQ1(13)->S06F0_1(2,1,10)	Down	S06F0_1(6,1,06)->FPC02FE1TQ1(13)	Down
FPC04FE0TQ0(13)->S06F0_1(1,0,09)	Down	S06F0_1(4,0,04)->FPC04FE0TQ0(13)	Down
FPC04FE1TQ1(13)->S06F0_1(1,1,09)	Down	S06F0_1(4,1,04)->FPC04FE1TQ1(13)	Down
FPC06FE0TQ0(13)->S06F0_1(0,0,08)	Down	S06F0_1(4,2,04)->FPC06FE0TQ0(13)	Down
FPC06FE1TQ1(13)->S06F0_1(0,1,08)	Down	S06F0_1(4,3,04)->FPC06FE1TQ1(13)	Down
FPC08FE0TQ0(13)->S06F0_1(0,2,08)	OK	S06F0_1(4,4,04)->FPC08FE0TQ0(13)	OK
FPC08FE1TQ1(13)->S06F0_1(0,3,08)	OK	S06F0_1(4,5,04)->FPC08FE1TQ1(13)	OK
FPC10FE0TQ0(13)->S06F0_1(1,2,09)	Down	S06F0_1(5,0,05)->FPC10FE0TQ0(13)	Down
FPC10FE1TQ1(13)->S06F0_1(1,3,09)	Down	S06F0_1(5,1,05)->FPC10FE1TQ1(13)	Down
FPC12FE0TQ0(13)->S06F0_1(2,2,10)	OK	S06F0_1(6,2,06)->FPC12FE0TQ0(13)	OK
FPC12FE1TQ1(13)->S06F0_1(2,3,10)	OK	S06F0_1(6,3,06)->FPC12FE1TQ1(13)	OK
FPC14FE0TQ0(13)->S06F0_1(3,2,11)	Down	S06F0_1(7,2,07)->FPC14FE0TQ0(13)	Down
FPC14FE1TQ1(13)->S06F0_1(3,3,11)	Down	S06F0_1(7,3,07)->FPC14FE1TQ1(13)	Down

SIB 7 FCHIP 0 FCORE 0 :

In-links	State	Out-links	State
FPC00FE0TQ0(14)->S07F0_0(7,0,07)	Down	S07F0_0(3,0,11)->FPC00FE0TQ0(14)	Down
FPC00FE1TQ1(14)->S07F0_0(7,1,07)	Down	S07F0_0(3,1,11)->FPC00FE1TQ1(14)	Down
FPC02FE0TQ0(14)->S07F0_0(6,0,06)	Down	S07F0_0(2,0,10)->FPC02FE0TQ0(14)	Down
FPC02FE1TQ1(14)->S07F0_0(6,1,06)	Down	S07F0_0(2,1,10)->FPC02FE1TQ1(14)	Down
FPC04FE0TQ0(14)->S07F0_0(5,0,05)	Down	S07F0_0(1,0,09)->FPC04FE0TQ0(14)	Down
FPC04FE1TQ1(14)->S07F0_0(5,1,05)	Down	S07F0_0(1,1,09)->FPC04FE1TQ1(14)	Down
FPC06FE0TQ0(14)->S07F0_0(4,0,04)	Down	S07F0_0(0,0,08)->FPC06FE0TQ0(14)	Down
FPC06FE1TQ1(14)->S07F0_0(4,1,04)	Down	S07F0_0(0,1,08)->FPC06FE1TQ1(14)	Down
FPC08FE0TQ0(14)->S07F0_0(4,2,04)	OK	S07F0_0(0,2,08)->FPC08FE0TQ0(14)	OK
FPC08FE1TQ1(14)->S07F0_0(4,3,04)	OK	S07F0_0(0,3,08)->FPC08FE1TQ1(14)	OK
FPC10FE0TQ0(14)->S07F0_0(5,2,05)	Down	S07F0_0(1,2,09)->FPC10FE0TQ0(14)	Down
FPC10FE1TQ1(14)->S07F0_0(5,3,05)	Down	S07F0_0(1,3,09)->FPC10FE1TQ1(14)	Down
FPC12FE0TQ0(14)->S07F0_0(7,2,07)	OK	S07F0_0(3,2,11)->FPC12FE0TQ0(14)	OK
FPC12FE1TQ1(14)->S07F0_0(7,3,07)	OK	S07F0_0(3,3,11)->FPC12FE1TQ1(14)	OK
FPC14FE0TQ0(14)->S07F0_0(7,4,07)	Down	S07F0_0(3,4,11)->FPC14FE0TQ0(14)	Down
FPC14FE1TQ1(14)->S07F0_0(7,5,07)	Down	S07F0_0(3,5,11)->FPC14FE1TQ1(14)	Down

SIB 7 FCHIP 0 FCORE 1 :

In-links	State	Out-links	State
FPC00FE0TQ0(15)->S07F0_1(3,0,11)	Down	S07F0_1(7,0,07)->FPC00FE0TQ0(15)	Down
FPC00FE1TQ1(15)->S07F0_1(3,1,11)	Down	S07F0_1(7,1,07)->FPC00FE1TQ1(15)	Down
FPC02FE0TQ0(15)->S07F0_1(2,0,10)	Down	S07F0_1(6,0,06)->FPC02FE0TQ0(15)	Down
FPC02FE1TQ1(15)->S07F0_1(2,1,10)	Down	S07F0_1(6,1,06)->FPC02FE1TQ1(15)	Down
FPC04FE0TQ0(15)->S07F0_1(1,0,09)	Down	S07F0_1(4,0,04)->FPC04FE0TQ0(15)	Down
FPC04FE1TQ1(15)->S07F0_1(1,1,09)	Down	S07F0_1(4,1,04)->FPC04FE1TQ1(15)	Down

```

FPC06FE0TQ0(15)->S07F0_1(0,0,08) Down    S07F0_1(4,2,04)->FPC06FE0TQ0(15) Down
FPC06FE1TQ1(15)->S07F0_1(0,1,08) Down    S07F0_1(4,3,04)->FPC06FE1TQ1(15) Down
FPC08FE0TQ0(15)->S07F0_1(0,2,08) OK       S07F0_1(4,4,04)->FPC08FE0TQ0(15) OK
FPC08FE1TQ1(15)->S07F0_1(0,3,08) OK       S07F0_1(4,5,04)->FPC08FE1TQ1(15) OK
FPC10FE0TQ0(15)->S07F0_1(1,2,09) Down    S07F0_1(5,0,05)->FPC10FE0TQ0(15) Down
FPC10FE1TQ1(15)->S07F0_1(1,3,09) Down    S07F0_1(5,1,05)->FPC10FE1TQ1(15) Down
FPC12FE0TQ0(15)->S07F0_1(2,2,10) OK      S07F0_1(6,2,06)->FPC12FE0TQ0(15) OK
FPC12FE1TQ1(15)->S07F0_1(2,3,10) OK      S07F0_1(6,3,06)->FPC12FE1TQ1(15) OK
FPC14FE0TQ0(15)->S07F0_1(3,2,11) Down    S07F0_1(7,2,07)->FPC14FE0TQ0(15) Down
FPC14FE1TQ1(15)->S07F0_1(3,3,11) Down    S07F0_1(7,3,07)->FPC14FE1TQ1(15) Down

```

SIB 8 FCHIP 0 FCORE 0 :

In-links	State	Out-links	State
FPC00FE0TQ0(16)->S08F0_0(7,0,07)	Down	S08F0_0(3,0,11)->FPC00FE0TQ0(16)	Down
FPC00FE1TQ1(16)->S08F0_0(7,1,07)	Down	S08F0_0(3,1,11)->FPC00FE1TQ1(16)	Down
FPC02FE0TQ0(16)->S08F0_0(6,0,06)	Down	S08F0_0(2,0,10)->FPC02FE0TQ0(16)	Down
FPC02FE1TQ1(16)->S08F0_0(6,1,06)	Down	S08F0_0(2,1,10)->FPC02FE1TQ1(16)	Down
FPC04FE0TQ0(16)->S08F0_0(5,0,05)	Down	S08F0_0(1,0,09)->FPC04FE0TQ0(16)	Down
FPC04FE1TQ1(16)->S08F0_0(5,1,05)	Down	S08F0_0(1,1,09)->FPC04FE1TQ1(16)	Down
FPC06FE0TQ0(16)->S08F0_0(4,0,04)	Down	S08F0_0(0,0,08)->FPC06FE0TQ0(16)	Down
FPC06FE1TQ1(16)->S08F0_0(4,1,04)	Down	S08F0_0(0,1,08)->FPC06FE1TQ1(16)	Down
FPC08FE0TQ0(16)->S08F0_0(4,2,04)	OK	S08F0_0(0,2,08)->FPC08FE0TQ0(16)	OK
FPC08FE1TQ1(16)->S08F0_0(4,3,04)	OK	S08F0_0(0,3,08)->FPC08FE1TQ1(16)	OK
FPC10FE0TQ0(16)->S08F0_0(5,2,05)	Down	S08F0_0(1,2,09)->FPC10FE0TQ0(16)	Down
FPC10FE1TQ1(16)->S08F0_0(5,3,05)	Down	S08F0_0(1,3,09)->FPC10FE1TQ1(16)	Down
FPC12FE0TQ0(16)->S08F0_0(7,2,07)	OK	S08F0_0(3,2,11)->FPC12FE0TQ0(16)	OK
FPC12FE1TQ1(16)->S08F0_0(7,3,07)	OK	S08F0_0(3,3,11)->FPC12FE1TQ1(16)	OK
FPC14FE0TQ0(16)->S08F0_0(7,4,07)	Down	S08F0_0(3,4,11)->FPC14FE0TQ0(16)	Down
FPC14FE1TQ1(16)->S08F0_0(7,5,07)	Down	S08F0_0(3,5,11)->FPC14FE1TQ1(16)	Down

SIB 8 FCHIP 0 FCORE 1 :

In-links	State	Out-links	State
FPC00FE0TQ0(17)->S08F0_1(3,0,11)	Down	S08F0_1(7,0,07)->FPC00FE0TQ0(17)	Down
FPC00FE1TQ1(17)->S08F0_1(3,1,11)	Down	S08F0_1(7,1,07)->FPC00FE1TQ1(17)	Down
FPC02FE0TQ0(17)->S08F0_1(2,0,10)	Down	S08F0_1(6,0,06)->FPC02FE0TQ0(17)	Down
FPC02FE1TQ1(17)->S08F0_1(2,1,10)	Down	S08F0_1(6,1,06)->FPC02FE1TQ1(17)	Down
FPC04FE0TQ0(17)->S08F0_1(1,0,09)	Down	S08F0_1(4,0,04)->FPC04FE0TQ0(17)	Down
FPC04FE1TQ1(17)->S08F0_1(1,1,09)	Down	S08F0_1(4,1,04)->FPC04FE1TQ1(17)	Down
FPC06FE0TQ0(17)->S08F0_1(0,0,08)	Down	S08F0_1(4,2,04)->FPC06FE0TQ0(17)	Down
FPC06FE1TQ1(17)->S08F0_1(0,1,08)	Down	S08F0_1(4,3,04)->FPC06FE1TQ1(17)	Down
FPC08FE0TQ0(17)->S08F0_1(0,2,08)	OK	S08F0_1(4,4,04)->FPC08FE0TQ0(17)	OK
FPC08FE1TQ1(17)->S08F0_1(0,3,08)	OK	S08F0_1(4,5,04)->FPC08FE1TQ1(17)	OK
FPC10FE0TQ0(17)->S08F0_1(1,2,09)	Down	S08F0_1(5,0,05)->FPC10FE0TQ0(17)	Down
FPC10FE1TQ1(17)->S08F0_1(1,3,09)	Down	S08F0_1(5,1,05)->FPC10FE1TQ1(17)	Down
FPC12FE0TQ0(17)->S08F0_1(2,2,10)	OK	S08F0_1(6,2,06)->FPC12FE0TQ0(17)	OK
FPC12FE1TQ1(17)->S08F0_1(2,3,10)	OK	S08F0_1(6,3,06)->FPC12FE1TQ1(17)	OK
FPC14FE0TQ0(17)->S08F0_1(3,2,11)	Down	S08F0_1(7,2,07)->FPC14FE0TQ0(17)	Down
FPC14FE1TQ1(17)->S08F0_1(3,3,11)	Down	S08F0_1(7,3,07)->FPC14FE1TQ1(17)	Down

show chassis fabric topology (QFX10008 Switch)

user@host> show chassis fabric topology

```

In-link : FPC# FE# ASIC# (TX inst#, TX sub-chnl #) ->
          SIB# ASIC#_FCORE# (RX port#, RX sub-chnl #, RX inst#)

```

```

Out-link : SIB# ASIC#_FCORE# (TX port#, TX sub-chnl #, TX inst#) ->
          FPC# FE# ASIC# (RX inst#, RX sub-chnl #)

```

SIB 0 FCHIP 0 FCORE 0 :

In-links	State	Out-links	State
FPC00FE0(1,17)->S00F0_0(01,0,01)	OK	S00F0_0(00,0,00)->FPC00FE0(1,09)	OK
FPC00FE0(1,09)->S00F0_0(02,0,02)	OK	S00F0_0(00,1,00)->FPC00FE0(1,17)	OK
FPC00FE0(1,07)->S00F0_0(02,2,02)	OK	S00F0_0(00,2,00)->FPC00FE0(1,07)	OK
FPC00FE1(1,12)->S00F0_0(01,1,01)	OK	S00F0_0(00,3,00)->FPC00FE1(1,06)	OK
FPC00FE1(1,06)->S00F0_0(01,2,01)	OK	S00F0_0(01,1,01)->FPC00FE1(1,12)	OK
FPC00FE1(1,10)->S00F0_0(01,3,01)	OK	S00F0_0(01,3,01)->FPC00FE1(1,10)	OK
FPC00FE2(1,16)->S00F0_0(00,4,00)	OK	S00F0_0(00,4,00)->FPC00FE2(1,08)	OK
FPC00FE2(1,08)->S00F0_0(01,6,01)	OK	S00F0_0(00,5,00)->FPC00FE2(1,16)	OK
FPC00FE2(1,06)->S00F0_0(01,7,01)	OK	S00F0_0(00,6,00)->FPC00FE2(1,06)	OK

SIB 0 FCHIP 1 FCORE 0 :

In-links	State	Out-links	State
FPC00FE0(1,15)->S00F1_0(15,4,15)	OK	S00F1_0(16,4,16)->FPC00FE0(1,15)	OK
FPC00FE0(1,11)->S00F1_0(17,4,17)	OK	S00F1_0(18,4,18)->FPC00FE0(1,11)	OK
FPC00FE0(1,13)->S00F1_0(17,6,17)	OK	S00F1_0(18,6,18)->FPC00FE0(1,13)	OK
FPC00FE1(1,08)->S00F1_0(15,6,15)	OK	S00F1_0(16,6,16)->FPC00FE1(1,08)	OK
FPC00FE1(1,14)->S00F1_0(17,5,17)	OK	S00F1_0(18,5,18)->FPC00FE1(1,14)	OK
FPC00FE1(1,16)->S00F1_0(17,7,17)	OK	S00F1_0(18,7,18)->FPC00FE1(1,16)	OK
FPC00FE2(1,14)->S00F1_0(16,0,16)	OK	S00F1_0(16,0,16)->FPC00FE2(1,14)	OK
FPC00FE2(1,10)->S00F1_0(18,0,18)	OK	S00F1_0(18,0,18)->FPC00FE2(1,10)	OK
FPC00FE2(1,12)->S00F1_0(18,2,18)	OK	S00F1_0(18,2,18)->FPC00FE2(1,12)	OK

SIB 1

Not Online

SIB 2

Not Online

SIB 3

Not Online

SIB 4

Not Online

SIB 5

Not Online

show chassis fan

List of Syntax	Syntax on page 678 Syntax (ACX4000 Series Router) on page 678 Syntax (MX Series Router) on page 678 Syntax (T Series Routers) on page 678 Syntax (MX104, MX2010, and MX2020 3D Universal Edge Router) on page 678 Syntax (QFX Series) on page 678 Syntax (OCX Series) on page 678 Syntax (TX Matrix Router) on page 678 Syntax (TX Matrix Plus Router) on page 678
Syntax	show chassis fan
Syntax (ACX4000 Series Router)	show chassis fan
Syntax (MX Series Router)	show chassis fan <all-members> <local> <member <i>member-id</i> >
Syntax (T Series Routers)	show chassis fan
Syntax (MX104, MX2010, and MX2020 3D Universal Edge Router)	show chassis fan <satellite [<i>slot-id slot-id</i> <i>device-alias alias-name</i>]>
Syntax (QFX Series)	show chassis fan <interconnect-device <i>name</i> >
Syntax (OCX Series)	show chassis fan
Syntax (TX Matrix Router)	show chassis fan <lcc <i>number</i> <i>scc</i> >
Syntax (TX Matrix Plus Router)	show chassis fan <lcc <i>number</i> <i>sfc number</i> >
Release Information	Command introduced in Junos OS Release 10.0 on MX Series 3D Universal Edge Routers, M120 routers, and M320 routers, T320 routers, T640 routers, T1600 routers, TX Matrix Routers, and TX Matrix Plus routers. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 11.4 for EX Series switches. Command introduced in Junos OS Release 12.3 for PTX5000 Packet Transport Routers. Command introduced in Junos OS Release 12.1 for T4000 routers. Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers. Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers. Command introduced in Junos OS Release 12.3 for ACX Series Routers.

Command introduced in Junos OS Release 13.2 for MX104 3D Universal Edge Routers.
 Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
satellite option introduced in Junos OS Release 14.2R3.

Description (T Series routers, TX Matrix routers, TX Matrix Plus routers, M120 routers, M320 routers, MX104 routers, MX2010 routers, MX2020 routers, MX Series 3D Universal Edge Routers, QFX3008-I Interconnect devices, QFX Series, OCX Series, EX Series switches, and PTX Series Packet Transport Routers only) Show information about the fan tray and fans.

Options **all-members**—(MX Series routers only) (Optional) Display information about the fan tray and fans for all members of the Virtual Chassis configuration.

local—(MX Series routers only) (Optional) Display information about the fan tray and fans for the local Virtual Chassis member.

member *member-id*—(MX Series routers only) (Optional) Display information about the fan tray and fans for the specified member of the Virtual Chassis configuration. For an MX Series Virtual Chassis, replace *member-id* variable with a value 0 or 1.

interconnect-device *name*—(QFX3000-G QFabric systems only) (Optional) Display information about the fan tray and fans for the specified QFX3008-I Interconnect device.

lcc *number*—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display information about the fan tray and fans for the specified T640 router (line-card chassis) that is connected to a TX Matrix router. On a TX Matrix Plus router, display information about the fan tray and fans for the specified router (line-card chassis) that is connected to a TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

satellite [*slot-id slot-id* | *device-alias alias-name*]—(Junos Fusion only) (Optional) Display information about the fan tray and fans for the specified satellite device or devices in a Junos Fusion, or for all satellite devices if no satellite devices are specified.

scc—(TX Matrix routers only) (Optional) Display information about the fan tray and fans for the TX Matrix router (switch-card chassis).

sfc *number*—(TX Matrix Plus routers only) (Optional) Display information about the fan tray and fans for the TX Matrix Plus router (switch-fabric chassis). Replace *number* variable with 0.

Required Privilege Level view

List of Sample Output

- [show chassis fan on page 681](#)
- [show chassis fan \(QFabric Systems\) on page 681](#)
- [show chassis fan \(EX Series Switches\) on page 683](#)
- [show chassis fan \(T320 Router\) on page 683](#)
- [show chassis fan \(T640 Router\) on page 683](#)
- [show chassis fan \(T1600 Router\) on page 684](#)
- [show chassis fan \(T4000 Core Router\) on page 684](#)
- [show chassis fan \(TX Matrix Router\) on page 685](#)
- [show chassis fan \(TX Matrix Plus Router\) on page 685](#)
- [show chassis fan \(TX Matrix Plus Router with 3D SIBs\) on page 687](#)
- [show chassis fan \(PTX5000 Packet Transport Router\) on page 689](#)
- [show chassis fan \(MX104 Router\) on page 689](#)
- [show chassis fan \(MX2010 Router\) on page 689](#)
- [show chassis fan \(MX2020 Router\) on page 690](#)
- [show chassis fan \(ACX4000 Router\) on page 690](#)
- [show chassis fan \(QFX5100 Switch and OCX Series\) on page 690](#)

Output Fields Table 56 lists the output fields for the **show chassis fan** command. Output fields are listed in the approximate order in which they appear.

Table 56: show chassis fan Output Fields

Field Name	Field Description
Item	Fan item identifier.
Status	Status of the fan: <ul style="list-style-type: none"> • OK—Fan is running properly and within the normal range. • Check—Fan is in Check state because of some fault or alarm condition.
RPM	(T Series routers, TX Matrix routers, TX Matrix Plus routers, MX Series 3D Universal Edge Routers, QFX3108 Interconnect devices, and EX Series switches only) Fan speed in revolutions per minute (RPM).
% RPM	(MX2010 routers, MX2020 routers, and PTX Series Packet Transport Routers only) Percentage of the fan speed being used.

Table 56: show chassis fan Output Fields (*continued*)

Field Name	Field Description
Measurement	<p>(T Series routers, TX Matrix routers, TX Matrix Plus routers, MX Series 3D Universal Edge Routers, QFX3108 Interconnect devices, and EX Series switches only) Fan speed status based on different chassis cooling requirements:</p> <ul style="list-style-type: none"> • Spinning at high speed • Spinning at intermediate speed • Spinning at normal speed • Spinning at low speed (except EX Series switches) <p>(MX2010 routers, MX2020 routers, and PTX Series Packet Transport Routers only) Fan speed in revolutions per minute (RPM) for each fan in the fan tray.</p>

Sample Output

show chassis fan

```
user@host> show chassis fan
```

Item	Status	RPM	Measurement
Top Tray Fan 1	OK	3790	Spinning at normal speed
Top Tray Fan 2	OK	3769	Spinning at normal speed
Top Tray Fan 3	OK	3769	Spinning at normal speed
Top Tray Fan 4	OK	3790	Spinning at normal speed
Top Tray Fan 5	OK	3790	Spinning at normal speed
Top Tray Fan 6	OK	3769	Spinning at normal speed
Top Tray Fan 7	OK	3790	Spinning at normal speed
Top Tray Fan 8	OK	3769	Spinning at normal speed
Top Tray Fan 9	OK	3769	Spinning at normal speed
Top Tray Fan 10	OK	3790	Spinning at normal speed
Top Tray Fan 11	OK	3790	Spinning at normal speed
Top Tray Fan 12	OK	3769	Spinning at normal speed
Bottom Tray Fan 1	OK	2880	Spinning at normal speed
Bottom Tray Fan 2	OK	2912	Spinning at normal speed
Bottom Tray Fan 3	OK	2928	Spinning at normal speed
Bottom Tray Fan 4	OK	2896	Spinning at normal speed
Bottom Tray Fan 5	OK	2896	Spinning at normal speed
Bottom Tray Fan 6	OK	2928	Spinning at normal speed

show chassis fan (QFabric Systems)

```
user@host> show chassis fan interconnect-device interconnect1
```

Item	Status	RPM	Measurement
TFT 0 Fan 0	OK	2849	Spinning at normal speed
TFT 0 Fan 1	OK	2821	Spinning at normal speed
TFT 0 Fan 2	OK	2735	Spinning at normal speed
TFT 0 Fan 3	OK	2815	Spinning at normal speed
TFT 0 Fan 4	OK	2828	Spinning at normal speed
TFT 0 Fan 5	OK	2863	Spinning at normal speed
BFT 1 Fan 0	OK	2941	Spinning at normal speed
BFT 1 Fan 1	OK	3008	Spinning at normal speed
BFT 1 Fan 2	OK	3073	Spinning at normal speed
BFT 1 Fan 3	OK	2925	Spinning at normal speed

BFT 1 Fan 4	OK	2863	Spinning at normal speed
BFT 1 Fan 5	OK	2933	Spinning at normal speed
SFT 0 Fan 0 Rotor 0	OK	15472	Spinning at normal speed
SFT 0 Fan 0 Rotor 1	OK	14477	Spinning at normal speed
SFT 0 Fan 1 Rotor 0	OK	15561	Spinning at normal speed
SFT 0 Fan 1 Rotor 1	OK	14210	Spinning at normal speed
SFT 0 Fan 2 Rotor 0	OK	16167	Spinning at normal speed
SFT 0 Fan 2 Rotor 1	OK	14248	Spinning at normal speed
SFT 0 Fan 3 Rotor 0	OK	16463	Spinning at normal speed
SFT 0 Fan 3 Rotor 1	OK	14099	Spinning at normal speed
SFT 1 Fan 0 Rotor 0	OK	15083	Spinning at normal speed
SFT 1 Fan 0 Rotor 1	OK	13533	Spinning at normal speed
SFT 1 Fan 1 Rotor 0	OK	16071	Spinning at normal speed
SFT 1 Fan 1 Rotor 1	OK	14400	Spinning at normal speed
SFT 1 Fan 2 Rotor 0	OK	15517	Spinning at normal speed
SFT 1 Fan 2 Rotor 1	OK	14210	Spinning at normal speed
SFT 1 Fan 3 Rotor 0	OK	16413	Spinning at normal speed
SFT 1 Fan 3 Rotor 1	OK	14400	Spinning at normal speed
SFT 2 Fan 0 Rotor 0	OK	15297	Spinning at normal speed
SFT 2 Fan 0 Rotor 1	OK	14634	Spinning at normal speed
SFT 2 Fan 1 Rotor 0	OK	15561	Spinning at normal speed
SFT 2 Fan 1 Rotor 1	OK	14285	Spinning at normal speed
SFT 2 Fan 2 Rotor 0	OK	15835	Spinning at normal speed
SFT 2 Fan 2 Rotor 1	OK	14400	Spinning at normal speed
SFT 2 Fan 3 Rotor 0	OK	15789	Spinning at normal speed
SFT 2 Fan 3 Rotor 1	OK	14323	Spinning at normal speed
SFT 3 Fan 0 Rotor 0	OK	16314	Spinning at normal speed
SFT 3 Fan 0 Rotor 1	OK	14876	Spinning at normal speed
SFT 3 Fan 1 Rotor 0	OK	15835	Spinning at normal speed
SFT 3 Fan 1 Rotor 1	OK	14323	Spinning at normal speed
SFT 3 Fan 2 Rotor 0	OK	16265	Spinning at normal speed
SFT 3 Fan 2 Rotor 1	OK	14594	Spinning at normal speed
SFT 3 Fan 3 Rotor 0	OK	16071	Spinning at normal speed
SFT 3 Fan 3 Rotor 1	OK	14323	Spinning at normal speed
SFT 4 Fan 0 Rotor 0	OK	15652	Spinning at normal speed
SFT 4 Fan 0 Rotor 1	OK	14438	Spinning at normal speed
SFT 4 Fan 1 Rotor 0	OK	16167	Spinning at normal speed
SFT 4 Fan 1 Rotor 1	OK	14555	Spinning at normal speed
SFT 4 Fan 2 Rotor 0	OK	16023	Spinning at normal speed
SFT 4 Fan 2 Rotor 1	OK	14361	Spinning at normal speed
SFT 4 Fan 3 Rotor 0	OK	16216	Spinning at normal speed
SFT 4 Fan 3 Rotor 1	OK	14438	Spinning at normal speed
SFT 5 Fan 0 Rotor 0	OK	15297	Spinning at normal speed
SFT 5 Fan 0 Rotor 1	OK	14173	Spinning at normal speed
SFT 5 Fan 1 Rotor 0	OK	15472	Spinning at normal speed
SFT 5 Fan 1 Rotor 1	OK	13846	Spinning at normal speed
SFT 5 Fan 2 Rotor 0	OK	15340	Spinning at normal speed
SFT 5 Fan 2 Rotor 1	OK	13917	Spinning at normal speed
SFT 5 Fan 3 Rotor 0	OK	15835	Spinning at normal speed
SFT 5 Fan 3 Rotor 1	OK	13917	Spinning at normal speed
SFT 6 Fan 0 Rotor 0	OK	15743	Spinning at normal speed
SFT 6 Fan 0 Rotor 1	OK	14594	Spinning at normal speed
SFT 6 Fan 1 Rotor 0	OK	16167	Spinning at normal speed
SFT 6 Fan 1 Rotor 1	OK	14634	Spinning at normal speed
SFT 6 Fan 2 Rotor 0	OK	16167	Spinning at normal speed
SFT 6 Fan 2 Rotor 1	OK	14516	Spinning at normal speed
SFT 6 Fan 3 Rotor 0	OK	16666	Spinning at normal speed
SFT 6 Fan 3 Rotor 1	OK	14438	Spinning at normal speed
SFT 7 Fan 0 Rotor 0	OK	15517	Spinning at normal speed
SFT 7 Fan 0 Rotor 1	OK	14438	Spinning at normal speed
SFT 7 Fan 1 Rotor 0	OK	15517	Spinning at normal speed

SFT 7 Fan 1 Rotor 1	OK	14361	Spinning at normal speed
SFT 7 Fan 2 Rotor 0	OK	16167	Spinning at normal speed
SFT 7 Fan 2 Rotor 1	OK	14555	Spinning at normal speed
SFT 7 Fan 3 Rotor 0	OK	15697	Spinning at normal speed
SFT 7 Fan 3 Rotor 1	OK	14361	Spinning at normal speed

show chassis fan (EX Series Switches)

```
user@host> show chassis fan
```

Item	Status	RPM	Measurement
Fan 1	OK	3477	Spinning at normal speed
Fan 2	OK	3477	Spinning at normal speed
Fan 3	OK	3479	Spinning at normal speed
Fan 4	OK	3508	Spinning at normal speed
Fan 5	OK	3517	Spinning at normal speed
Fan 6	OK	3531	Spinning at normal speed
Fan 7	OK	3439	Spinning at normal speed
Fan 8	OK	3424	Spinning at normal speed
Fan 9	OK	3413	Spinning at normal speed
Fan 10	OK	3439	Spinning at normal speed
Fan 11	OK	3446	Spinning at normal speed
Fan 12	OK	3432	Spinning at normal speed

show chassis fan (T320 Router)

```
user@host> show chassis fan
```

Item	Status	RPM	Measurement
Top Left Front fan	OK	2850	Spinning at normal speed
Top Left Middle fan	OK	2820	Spinning at normal speed
Top Left Rear fan	OK	2970	Spinning at normal speed
Top Right Front fan	OK	2790	Spinning at normal speed
Top Right Middle fan	OK	2640	Spinning at normal speed
Top Right Rear fan	OK	2790	Spinning at normal speed
Bottom Left Front fan	OK	2520	Spinning at normal speed
Bottom Left Middle fan	OK	2610	Spinning at normal speed
Bottom Left Rear fan	OK	2550	Spinning at normal speed
Bottom Right Front fan	OK	2610	Spinning at normal speed
Bottom Right Middle fan	OK	2880	Spinning at normal speed
Bottom Right Rear fan	OK	2790	Spinning at normal speed
Rear Tray Top fan	OK	2130	Spinning at normal speed
Rear Tray Second fan	OK	2190	Spinning at normal speed
Rear Tray Middle fan	OK	2250	Spinning at normal speed
Rear Tray Fourth fan	OK	2220	Spinning at normal speed
Rear Tray Bottom fan	OK	2280	Spinning at normal speed

show chassis fan (T640 Router)

```
user@host> show chassis fan
```

Item	Status	RPM	Measurement
Top Left Front fan	OK	3420	Spinning at normal speed
Top Left Middle fan	OK	3420	Spinning at normal speed
Top Left Rear fan	OK	3420	Spinning at normal speed
Top Right Front fan	OK	3420	Spinning at normal speed
Top Right Middle fan	OK	3420	Spinning at normal speed
Top Right Rear fan	OK	3450	Spinning at normal speed
Bottom Left Front fan	OK	3390	Spinning at normal speed
Bottom Left Middle fan	OK	3420	Spinning at normal speed
Bottom Left Rear fan	OK	3390	Spinning at normal speed
Bottom Right Front fan	OK	3390	Spinning at normal speed

Bottom Right Middle fan	OK	3390	Spinning at normal speed
Bottom Right Rear fan	OK	3390	Spinning at normal speed
Rear Tray Top fan	OK	5220	Spinning at normal speed
Rear Tray Second fan	OK	5220	Spinning at normal speed
Rear Tray Third fan	OK	5220	Spinning at normal speed
Rear Tray Fourth fan	OK	5220	Spinning at normal speed
Rear Tray Fifth fan	OK	5220	Spinning at normal speed
Rear Tray Sixth fan	OK	5220	Spinning at normal speed
Rear Tray Seventh fan	OK	5220	Spinning at normal speed
Rear Tray Bottom fan	OK	5220	Spinning at normal speed

show chassis fan (T1600 Router)

```
user@host> show chassis fan
```

Item	Status	RPM	Measurement
Top Left Front fan	OK	3420	Spinning at normal speed
Top Left Middle fan	OK	3420	Spinning at normal speed
Top Left Rear fan	OK	3450	Spinning at normal speed
Top Right Front fan	OK	3420	Spinning at normal speed
Top Right Middle fan	OK	3420	Spinning at normal speed
Top Right Rear fan	OK	3390	Spinning at normal speed
Bottom Left Front fan	OK	3420	Spinning at normal speed
Bottom Left Middle fan	OK	3420	Spinning at normal speed
Bottom Left Rear fan	OK	3390	Spinning at normal speed
Bottom Right Front fan	OK	3390	Spinning at normal speed
Bottom Right Middle fan	OK	3420	Spinning at normal speed
Bottom Right Rear fan	OK	3390	Spinning at normal speed
Rear Tray Top fan	OK	5190	Spinning at normal speed
Rear Tray Second fan	OK	5190	Spinning at normal speed
Rear Tray Third fan	OK	5190	Spinning at normal speed
Rear Tray Fourth fan	OK	5190	Spinning at normal speed
Rear Tray Fifth fan	OK	5190	Spinning at normal speed
Rear Tray Sixth fan	OK	5190	Spinning at normal speed
Rear Tray Seventh fan	OK	5190	Spinning at normal speed
Rear Tray Bottom fan	OK	5190	Spinning at normal speed

show chassis fan (T4000 Core Router)

```
user@host> show chassis fan
```

Item	Status	RPM	Measurement
Top Left Front fan	OK	5190	Spinning at high speed
Top Left Middle fan	OK	5220	Spinning at high speed
Top Left Rear fan	OK	5190	Spinning at high speed
Top Right Front fan	OK	5160	Spinning at high speed
Top Right Middle fan	OK	5190	Spinning at high speed
Top Right Rear fan	OK	5160	Spinning at high speed
Bottom Left Front fan	OK	6030	Spinning at high speed
Bottom Left Middle fan	OK	6090	Spinning at high speed
Bottom Left Rear fan	OK	6090	Spinning at high speed
Bottom Right Front fan	OK	6030	Spinning at high speed
Bottom Right Middle fan	OK	6060	Spinning at high speed
Bottom Right Rear fan	OK	6060	Spinning at high speed
Rear Tray Top fan	OK	10000	Spinning at high speed
Rear Tray Second fan	OK	10000	Spinning at high speed
Rear Tray Third fan	OK	10000	Spinning at high speed
Rear Tray Fourth fan	OK	10000	Spinning at high speed
Rear Tray Fifth fan	OK	10000	Spinning at high speed
Rear Tray Sixth fan	OK	10000	Spinning at high speed

Rear Tray Seventh fan	OK	10000	Spinning at high speed
Rear Tray Bottom fan	OK	10000	Spinning at high speed

show chassis fan (TX Matrix Router)

```
user@host> show chassis fan
scc-re0:
```

Item	Status	RPM	Measurement
Top Left Front fan	OK	3420	Spinning at normal speed
Top Left Middle fan	OK	3390	Spinning at normal speed
Top Left Rear fan	OK	3420	Spinning at normal speed
Top Right Front fan	OK	3390	Spinning at normal speed
Top Right Middle fan	OK	3420	Spinning at normal speed
Top Right Rear fan	OK	3390	Spinning at normal speed
Bottom Left Front fan	OK	3420	Spinning at normal speed
Bottom Left Middle fan	OK	3450	Spinning at normal speed
Bottom Left Rear fan	OK	3420	Spinning at normal speed
Bottom Right Front fan	OK	3420	Spinning at normal speed
Bottom Right Middle fan	OK	3420	Spinning at normal speed
Bottom Right Rear fan	OK	3420	Spinning at normal speed
Rear Tray Top fan	OK	3420	Spinning at normal speed
Rear Tray Second fan	OK	5190	Spinning at normal speed
Rear Tray Third fan	OK	5190	Spinning at normal speed
Rear Tray Fourth fan	OK	5190	Spinning at normal speed
Rear Tray Fifth fan	OK	3420	Spinning at normal speed
Rear Tray Sixth fan	OK	3420	Spinning at normal speed
Rear Tray Seventh fan	OK	3420	Spinning at normal speed
Rear Tray Bottom fan	OK	3420	Spinning at normal speed

```
1cc2-re0:
```

Item	Status	RPM	Measurement
Top Left Front fan	OK	3420	Spinning at normal speed
Top Left Middle fan	OK	3420	Spinning at normal speed
Top Left Rear fan	OK	3450	Spinning at normal speed
Top Right Front fan	OK	3420	Spinning at normal speed
Top Right Middle fan	OK	3450	Spinning at normal speed
Top Right Rear fan	OK	3360	Spinning at normal speed
Bottom Left Front fan	OK	3420	Spinning at normal speed
Bottom Left Middle fan	OK	3480	Spinning at normal speed
Bottom Left Rear fan	OK	3420	Spinning at normal speed
Bottom Right Front fan	OK	3420	Spinning at normal speed
Bottom Right Middle fan	OK	3390	Spinning at normal speed
Bottom Right Rear fan	OK	3420	Spinning at normal speed
Rear Tray Top fan	OK	3420	Spinning at normal speed
Rear Tray Second fan	OK	3420	Spinning at normal speed
Rear Tray Third fan	OK	3420	Spinning at normal speed
Rear Tray Fourth fan	OK	3420	Spinning at normal speed
Rear Tray Fifth fan	OK	3420	Spinning at normal speed
Rear Tray Sixth fan	OK	3420	Spinning at normal speed
Rear Tray Seventh fan	OK	3420	Spinning at normal speed
Rear Tray Bottom fan	OK	3420	Spinning at normal speed

show chassis fan (TX Matrix Plus Router)

```
user@host> show chassis fan
sfc0-re0:
```

Item	Status	RPM	Measurement
Fan Tray 0 Fan 1	OK	4350	Spinning at normal speed

Fan Tray 0 Fan 2	OK	4380	Spinning at normal speed
Fan Tray 0 Fan 3	OK	4410	Spinning at normal speed
Fan Tray 0 Fan 4	OK	4380	Spinning at normal speed
Fan Tray 0 Fan 5	OK	4350	Spinning at normal speed
Fan Tray 0 Fan 6	OK	4380	Spinning at normal speed
Fan Tray 1 Fan 1	OK	4410	Spinning at normal speed
Fan Tray 1 Fan 2	OK	4380	Spinning at normal speed
Fan Tray 1 Fan 3	OK	4410	Spinning at normal speed
Fan Tray 1 Fan 4	OK	4380	Spinning at normal speed
Fan Tray 1 Fan 5	OK	4410	Spinning at normal speed
Fan Tray 1 Fan 6	OK	4410	Spinning at normal speed
Fan Tray 2 Fan 1	OK	4380	Spinning at normal speed
Fan Tray 2 Fan 2	OK	4380	Spinning at normal speed
Fan Tray 2 Fan 3	OK	4380	Spinning at normal speed
Fan Tray 2 Fan 4	OK	4410	Spinning at normal speed
Fan Tray 2 Fan 5	OK	4380	Spinning at normal speed
Fan Tray 2 Fan 6	OK	4410	Spinning at normal speed
Fan Tray 2 Fan 7	OK	4410	Spinning at normal speed
Fan Tray 2 Fan 8	OK	4380	Spinning at normal speed
Fan Tray 2 Fan 9	OK	4380	Spinning at normal speed
Fan Tray 3 Fan 1	OK	4350	Spinning at normal speed
Fan Tray 3 Fan 2	OK	4380	Spinning at normal speed
Fan Tray 3 Fan 3	OK	4410	Spinning at normal speed
Fan Tray 3 Fan 4	OK	4440	Spinning at normal speed
Fan Tray 3 Fan 5	OK	4380	Spinning at normal speed
Fan Tray 3 Fan 6	OK	4410	Spinning at normal speed
Fan Tray 3 Fan 7	OK	4410	Spinning at normal speed
Fan Tray 3 Fan 8	OK	4380	Spinning at normal speed
Fan Tray 3 Fan 9	OK	4410	Spinning at normal speed
Fan Tray 4 Fan 1	OK	4410	Spinning at normal speed
Fan Tray 4 Fan 2	OK	4410	Spinning at normal speed
Fan Tray 4 Fan 3	OK	4380	Spinning at normal speed
Fan Tray 4 Fan 4	OK	4380	Spinning at normal speed
Fan Tray 4 Fan 5	OK	4410	Spinning at normal speed
Fan Tray 4 Fan 6	OK	4410	Spinning at normal speed
Fan Tray 4 Fan 7	OK	4410	Spinning at normal speed
Fan Tray 4 Fan 8	OK	4410	Spinning at normal speed
Fan Tray 4 Fan 9	OK	4410	Spinning at normal speed
Fan Tray 5 Fan 1	OK	4350	Spinning at normal speed
Fan Tray 5 Fan 2	OK	4380	Spinning at normal speed
Fan Tray 5 Fan 3	OK	4380	Spinning at normal speed
Fan Tray 5 Fan 4	OK	4350	Spinning at normal speed
Fan Tray 5 Fan 5	OK	4380	Spinning at normal speed
Fan Tray 5 Fan 6	OK	4410	Spinning at normal speed
Fan Tray 5 Fan 7	OK	4410	Spinning at normal speed
Fan Tray 5 Fan 8	OK	4380	Spinning at normal speed
Fan Tray 5 Fan 9	OK	4410	Spinning at normal speed

1cc0-re0:

Item	Status	RPM	Measurement
Top Left Front fan	OK	3420	Spinning at normal speed
Top Left Middle fan	OK	3420	Spinning at normal speed
Top Left Rear fan	OK	3420	Spinning at normal speed
Top Right Front fan	OK	3450	Spinning at normal speed
Top Right Middle fan	OK	3420	Spinning at normal speed
Top Right Rear fan	OK	3420	Spinning at normal speed
Bottom Left Front fan	OK	3420	Spinning at normal speed
Bottom Left Middle fan	OK	3420	Spinning at normal speed
Bottom Left Rear fan	OK	3390	Spinning at normal speed
Bottom Right Front fan	OK	3420	Spinning at normal speed

Bottom Right Middle fan	OK	3390	Spinning at normal speed
Bottom Right Rear fan	OK	3390	Spinning at normal speed
Rear Tray Top fan	OK	7050	Spinning at normal speed
Rear Tray Second fan	OK	7050	Spinning at normal speed
Rear Tray Third fan	OK	7050	Spinning at normal speed
Rear Tray Fourth fan	OK	7050	Spinning at normal speed
Rear Tray Fifth fan	OK	7050	Spinning at normal speed
Rear Tray Sixth fan	OK	7050	Spinning at normal speed
Rear Tray Seventh fan	OK	7050	Spinning at normal speed
Rear Tray Bottom fan	OK	7050	Spinning at normal speed

show chassis fan (TX Matrix Plus Router with 3D SIBs)

```
user@host> show chassis fan
sfc0-re0:
```

Item	Status	RPM	Measurement
Fan Tray 0 Fan 1	OK	4830	Spinning at normal speed
Fan Tray 0 Fan 2	OK	4860	Spinning at normal speed
Fan Tray 0 Fan 3	OK	4830	Spinning at normal speed
Fan Tray 0 Fan 4	OK	4800	Spinning at normal speed
Fan Tray 0 Fan 5	OK	4830	Spinning at normal speed
Fan Tray 0 Fan 6	OK	4770	Spinning at normal speed
Fan Tray 1 Fan 1	OK	4800	Spinning at normal speed
Fan Tray 1 Fan 2	OK	4770	Spinning at normal speed
Fan Tray 1 Fan 3	OK	4800	Spinning at normal speed
Fan Tray 1 Fan 4	OK	4770	Spinning at normal speed
Fan Tray 1 Fan 5	OK	4770	Spinning at normal speed
Fan Tray 1 Fan 6	OK	4800	Spinning at normal speed
Fan Tray 2 Fan 1	OK	4800	Spinning at normal speed
Fan Tray 2 Fan 2	OK	4800	Spinning at normal speed
Fan Tray 2 Fan 3	OK	4830	Spinning at normal speed
Fan Tray 2 Fan 4	OK	4830	Spinning at normal speed
Fan Tray 2 Fan 5	OK	4830	Spinning at normal speed
Fan Tray 2 Fan 6	OK	4830	Spinning at normal speed
Fan Tray 2 Fan 7	OK	4800	Spinning at normal speed
Fan Tray 2 Fan 8	OK	4830	Spinning at normal speed
Fan Tray 2 Fan 9	OK	4800	Spinning at normal speed
Fan Tray 3 Fan 1	OK	4860	Spinning at normal speed
Fan Tray 3 Fan 2	OK	4860	Spinning at normal speed
Fan Tray 3 Fan 3	OK	4800	Spinning at normal speed
Fan Tray 3 Fan 4	OK	4830	Spinning at normal speed
Fan Tray 3 Fan 5	OK	4830	Spinning at normal speed
Fan Tray 3 Fan 6	OK	4830	Spinning at normal speed
Fan Tray 3 Fan 7	OK	4830	Spinning at normal speed
Fan Tray 3 Fan 8	OK	4800	Spinning at normal speed
Fan Tray 3 Fan 9	OK	4800	Spinning at normal speed
Fan Tray 4 Fan 1	OK	4830	Spinning at normal speed
Fan Tray 4 Fan 2	OK	4830	Spinning at normal speed
Fan Tray 4 Fan 3	OK	4830	Spinning at normal speed
Fan Tray 4 Fan 4	OK	4830	Spinning at normal speed
Fan Tray 4 Fan 5	OK	4830	Spinning at normal speed
Fan Tray 4 Fan 6	OK	4860	Spinning at normal speed
Fan Tray 4 Fan 7	OK	4800	Spinning at normal speed
Fan Tray 4 Fan 8	OK	4860	Spinning at normal speed
Fan Tray 4 Fan 9	OK	4770	Spinning at normal speed
Fan Tray 5 Fan 1	OK	4830	Spinning at normal speed
Fan Tray 5 Fan 2	OK	4830	Spinning at normal speed
Fan Tray 5 Fan 3	OK	4830	Spinning at normal speed
Fan Tray 5 Fan 4	OK	4800	Spinning at normal speed
Fan Tray 5 Fan 5	OK	4800	Spinning at normal speed

Fan Tray 5 Fan 6	OK	4800	Spinning at normal speed
Fan Tray 5 Fan 7	OK	4830	Spinning at normal speed
Fan Tray 5 Fan 8	OK	4830	Spinning at normal speed
Fan Tray 5 Fan 9	Check	2010	

1cc0-re0:

Item	Status	RPM	Measurement
Top Left Front fan	OK	3420	Spinning at normal speed
Top Left Middle fan	OK	3390	Spinning at normal speed
Top Left Rear fan	OK	3390	Spinning at normal speed
Top Right Front fan	OK	3420	Spinning at normal speed
Top Right Middle fan	OK	3420	Spinning at normal speed
Top Right Rear fan	OK	3450	Spinning at normal speed
Bottom Left Front fan	OK	3420	Spinning at normal speed
Bottom Left Middle fan	OK	3390	Spinning at normal speed
Bottom Left Rear fan	OK	3420	Spinning at normal speed
Bottom Right Front fan	OK	3420	Spinning at normal speed
Bottom Right Middle fan	OK	3390	Spinning at normal speed
Bottom Right Rear fan	OK	3420	Spinning at normal speed
Rear Tray fan 1 (Top)	OK	7740	Spinning at normal speed
Rear Tray fan 2	OK	7740	Spinning at normal speed
Rear Tray fan 3	OK	7740	Spinning at normal speed
Rear Tray fan 4	OK	7740	Spinning at normal speed
Rear Tray fan 5	OK	7740	Spinning at normal speed
Rear Tray fan 6	OK	7740	Spinning at normal speed
Rear Tray fan 7	OK	7740	Spinning at normal speed
Rear Tray fan 8	OK	7740	Spinning at normal speed
Rear Tray fan 9	OK	7740	Spinning at normal speed
Rear Tray fan 10	OK	7740	Spinning at normal speed
Rear Tray fan 11	OK	7740	Spinning at normal speed
Rear Tray fan 12	OK	7740	Spinning at normal speed
Rear Tray fan 13	OK	7740	Spinning at normal speed
Rear Tray fan 14	OK	7740	Spinning at normal speed
Rear Tray fan 15	OK	7740	Spinning at normal speed
Rear Tray fan 16 (Bottom)	OK	7740	Spinning at normal speed

1cc2-re0:

Item	Status	RPM	Measurement
Top Left Front fan	OK	3420	Spinning at normal speed
Top Left Middle fan	OK	3390	Spinning at normal speed
Top Left Rear fan	OK	3420	Spinning at normal speed
Top Right Front fan	OK	3420	Spinning at normal speed
Top Right Middle fan	OK	3420	Spinning at normal speed
Top Right Rear fan	OK	3450	Spinning at normal speed
Bottom Left Front fan	OK	3420	Spinning at normal speed
Bottom Left Middle fan	OK	3390	Spinning at normal speed
Bottom Left Rear fan	OK	3420	Spinning at normal speed
Bottom Right Front fan	OK	3420	Spinning at normal speed
Bottom Right Middle fan	OK	3390	Spinning at normal speed
Bottom Right Rear fan	OK	3420	Spinning at normal speed
Rear Tray fan 1 (Top)	OK	7740	Spinning at normal speed
Rear Tray fan 2	OK	7740	Spinning at normal speed
Rear Tray fan 3	OK	7740	Spinning at normal speed
Rear Tray fan 4	OK	7740	Spinning at normal speed
Rear Tray fan 5	OK	7740	Spinning at normal speed
Rear Tray fan 6	OK	7740	Spinning at normal speed
Rear Tray fan 7	OK	7740	Spinning at normal speed
Rear Tray fan 8	OK	7740	Spinning at normal speed
Rear Tray fan 9	OK	7740	Spinning at normal speed

Rear Tray fan 10	OK	7740	Spinning at normal speed
Rear Tray fan 11	OK	7740	Spinning at normal speed
Rear Tray fan 12	OK	7740	Spinning at normal speed
Rear Tray fan 13	OK	7740	Spinning at normal speed
Rear Tray fan 14	OK	7740	Spinning at normal speed
Rear Tray fan 15	OK	7740	Spinning at normal speed
Rear Tray fan 16 (Bottom)	OK	7740	Spinning at normal speed

show chassis fan (PTX5000 Packet Transport Router)

```
user@host> show chassis fan
user@host> show chassis fan
```

Item	Status	% RPM	Measurement
Fan Tray 0 Fan 1	OK	29%	2700 RPM
Fan Tray 0 Fan 2	OK	29%	2700 RPM
Fan Tray 0 Fan 3	OK	29%	2742 RPM
Fan Tray 0 Fan 4	OK	29%	2700 RPM
Fan Tray 0 Fan 5	OK	30%	2828 RPM
Fan Tray 0 Fan 6	OK	30%	2828 RPM
Fan Tray 0 Fan 7	OK	29%	2700 RPM
Fan Tray 0 Fan 8	OK	30%	2785 RPM
Fan Tray 0 Fan 9	OK	30%	2828 RPM
Fan Tray 0 Fan 10	OK	30%	2828 RPM
Fan Tray 0 Fan 11	OK	30%	2785 RPM
Fan Tray 0 Fan 12	OK	30%	2828 RPM
Fan Tray 0 Fan 13	OK	31%	2871 RPM
Fan Tray 0 Fan 14	OK	30%	2828 RPM
Fan Tray 1 Fan 1	OK	42%	3033 RPM
Fan Tray 1 Fan 2	OK	42%	3066 RPM
Fan Tray 1 Fan 3	OK	43%	3099 RPM
Fan Tray 1 Fan 4	OK	43%	3166 RPM
Fan Tray 1 Fan 5	OK	45%	3266 RPM
Fan Tray 1 Fan 6	OK	43%	3133 RPM
Fan Tray 2 Fan 1	OK	29%	2099 RPM
Fan Tray 2 Fan 2	OK	30%	2199 RPM
Fan Tray 2 Fan 3	OK	30%	2166 RPM
Fan Tray 2 Fan 4	OK	33%	2399 RPM
Fan Tray 2 Fan 5	OK	29%	2133 RPM
Fan Tray 2 Fan 6	OK	32%	2366 RPM

show chassis fan (MX104 Router)

```
user@host > show chassis fan
```

Item	Status	RPM	Measurement
Fan 1	OK	5640	Spinning at normal speed
Fan 2	OK	5640	Spinning at normal speed
Fan 3	OK	5760	Spinning at normal speed
Fan 4	OK	5640	Spinning at normal speed
Fan 5	OK	5640	Spinning at normal speed

show chassis fan (MX2010 Router)

```
user@host > show chassis fan
```

Item	Status	% RPM	Measurement
Fan Tray 0 Fan 1	OK	37%	3360 RPM
Fan Tray 0 Fan 2	OK	38%	3480 RPM
Fan Tray 0 Fan 3	OK	37%	3360 RPM
Fan Tray 0 Fan 4	OK	37%	3360 RPM
Fan Tray 0 Fan 5	OK	38%	3480 RPM
Fan Tray 0 Fan 6	OK	37%	3360 RPM
Fan Tray 1 Fan 1	OK	38%	3480 RPM
Fan Tray 1 Fan 2	OK	40%	3600 RPM

Fan Tray 1 Fan 3	OK	38%	3480 RPM
Fan Tray 1 Fan 4	OK	38%	3480 RPM
Fan Tray 1 Fan 5	OK	38%	3480 RPM
Fan Tray 1 Fan 6	OK	38%	3480 RPM
Fan Tray 2 Fan 1	OK	38%	3480 RPM
Fan Tray 2 Fan 2	OK	41%	3720 RPM
Fan Tray 2 Fan 3	OK	38%	3480 RPM
Fan Tray 2 Fan 4	OK	38%	3480 RPM
Fan Tray 2 Fan 5	OK	38%	3480 RPM
Fan Tray 2 Fan 6	OK	38%	3480 RPM
Fan Tray 3 Fan 1	OK	38%	3480 RPM
Fan Tray 3 Fan 2	OK	40%	3600 RPM
Fan Tray 3 Fan 3	OK	40%	3600 RPM
Fan Tray 3 Fan 4	OK	40%	3600 RPM
Fan Tray 3 Fan 5	OK	40%	3600 RPM
Fan Tray 3 Fan 6	OK	38%	3480 RPM

show chassis fan (MX2020 Router)

```
user@host > show chassis fan
```

Item	Status	% RPM	Measurement
Fan Tray 0 Fan 1	OK	37%	3360 RPM
Fan Tray 0 Fan 2	OK	37%	3360 RPM
Fan Tray 0 Fan 3	OK	36%	3240 RPM
Fan Tray 0 Fan 4	OK	37%	3360 RPM
Fan Tray 0 Fan 5	OK	37%	3360 RPM
Fan Tray 0 Fan 6	OK	37%	3360 RPM
Fan Tray 1 Fan 1	OK	37%	3360 RPM
Fan Tray 1 Fan 2	OK	37%	3360 RPM
Fan Tray 1 Fan 3	OK	37%	3360 RPM
Fan Tray 1 Fan 4	OK	37%	3360 RPM
Fan Tray 1 Fan 5	OK	37%	3360 RPM
Fan Tray 1 Fan 6	OK	36%	3240 RPM
Fan Tray 2 Fan 1	OK	37%	3360 RPM
Fan Tray 2 Fan 2	OK	37%	3360 RPM
Fan Tray 2 Fan 3	OK	37%	3360 RPM
Fan Tray 2 Fan 4	OK	37%	3360 RPM
Fan Tray 2 Fan 5	OK	37%	3360 RPM
Fan Tray 2 Fan 6	OK	38%	3480 RPM
Fan Tray 3 Fan 1	OK	38%	3480 RPM
Fan Tray 3 Fan 2	OK	38%	3480 RPM
Fan Tray 3 Fan 3	OK	38%	3480 RPM
Fan Tray 3 Fan 4	OK	37%	3360 RPM
Fan Tray 3 Fan 5	OK	37%	3360 RPM
Fan Tray 3 Fan 6	OK	37%	3360 RPM

show chassis fan (ACX4000 Router)

```
user@host > show chassis fan
```

Item	Status	RPM	Measurement
Fan 1	OK	4140	Spinning at normal speed
Fan 2	OK	4200	Spinning at normal speed

show chassis fan (QFX5100 Switch and OCX Series)

```
user@switch > show chassis fan
```

Item	Status	RPM	Measurement
FPC 0 Tray 0 Fan 0	OK	6428	Spinning at normal speed
FPC 0 Tray 0 Fan 1	OK	5515	Spinning at normal speed
FPC 0 Tray 1 Fan 0	OK	6360	Spinning at normal speed
FPC 0 Tray 1 Fan 1	OK	5532	Spinning at normal speed

show chassis firmware

List of Syntax	Syntax on page 692 Syntax (TX Matrix Routers) on page 692 Syntax (TX Matrix Plus Routers) on page 692 Syntax (MX Series Routers) on page 692 Syntax (MX104, MX2010, and MX2020 3D Universal Edge Routers) on page 692 Syntax (QFX Series) on page 692 Syntax (OCX Series) on page 692 Syntax (ACX Series Universal Access Routers) on page 692 Syntax (EX Series Switches) on page 692
Syntax	show chassis firmware
Syntax (TX Matrix Routers)	show chassis firmware <lcc <i>number</i> scc>
Syntax (TX Matrix Plus Routers)	show chassis firmware <lcc <i>number</i> sfc <i>number</i> >
Syntax (MX Series Routers)	show chassis firmware <all-members> <local> <member <i>member-id</i> >
Syntax (MX104, MX2010, and MX2020 3D Universal Edge Routers)	show chassis firmware <satellite [slot-id <i>slot-id</i> device-alias <i>alias-name</i>]>
Syntax (QFX Series)	show chassis firmware interconnect-device <i>name</i> node-device <i>name</i>
Syntax (OCX Series)	show chassis firmware
Syntax (ACX Series Universal Access Routers)	show chassis firmware
Syntax (EX Series Switches)	show chassis firmware <detail> <satellite [slot-id <i>slot-id</i> device-alias <i>alias-name</i>]>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.4 for EX Series switches. sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced for EX8200 switches in Junos OS Release 10.2 for EX Series switches. Command introduced in Junos OS Release 11.1 for QFX Series. Command introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.

Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.
 Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.
 Command introduced in Junos OS Release 12.3 for ACX4000 Universal Access Routers.
 Command introduced in Junos OS Release 13.2 for MX104 3D Universal Edge Routers.
 Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
satellite option introduced in Junos OS Release 14.2R3.

Description On routers and switches, display the version levels of the firmware running on the System Control Board (SCB), Switching and Forwarding Module (SFM), System and Switch Board (SSB), Forwarding Engine Board (FEB), Flexible PIC Concentrators (FPCs), and Routing Engines. On a TX Matrix Plus router, display the version levels of the firmware running on the FPCs and the Switch Processor Mezzanine Board (SPMBs).

On EX2200, EX3200, EX4200, QFX Series, and OCX Series switches, display the version levels of the firmware running on the switch. On an EX8208 switch, display the version levels of the firmware running on the Switch Fabric and Routing Engine (SRE) modules and on the line cards (shown as FPCs). On an EX8216 switch, display the version levels of the firmware running on the Routing Engine (RE) modules and on the line cards (shown as FPCs).

Options **none**—Display the version levels of the firmware running. For an EX4200 switch that is a member of a Virtual Chassis, display version levels for all members. For a TX Matrix router, display version levels for the firmware on the TX Matrix router and on all the T640 routers connected to the TX Matrix router. For a TX Matrix Plus router, display version levels for the firmware on the TX Matrix Plus router and on all the routers connected to the TX Matrix Plus router.

all-members—(MX Series routers only) (Optional) Display the version levels of the firmware running for all members of the Virtual Chassis configuration.

interconnect-device *name*—(QFabric systems) (Optional) Display the version levels of the firmware running on the Interconnect device.

lcc *number*—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display version levels for the firmware on a specified T640 router (line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, display the version levels for the firmware on a specified router (line-card chassis) that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(MX Series routers only) (Optional) Display the version levels of the firmware running for the local Virtual Chassis member.

member *member-id*—(MX Series routers only) (Optional) Display the version levels of the firmware running for the specified member of the Virtual Chassis configuration. Replace ***member-id*** with a value of 0 or 1.

node-device—(QFabric systems only) (Optional) Display the version levels of the firmware running on the Node device.

satellite [*slot-id slot-id* | *device-alias alias-name*]—(Junos Fusion only) (Optional) Display version levels of the firmware running for the specified satellite device or devices in a Junos Fusion, or for all satellite devices if no satellite devices are specified.

scc—(TX Matrix router only) (Optional) Display version levels for the firmware on the TX Matrix router (switch-card chassis).

sfc *number*—(TX Matrix Plus router only) (Optional) Display version levels for the firmware on the TX Matrix Plus router (or switch-fabric chassis). Replace ***number*** with 0.

detail—(EX3200, EX3300, EX4200, and EX4500 standalone and Virtual Chassis member switches only) (Optional) Display version levels of the firmware running on the switch for its programmable hardware components.

Required Privilege
Level

view

List of Sample Output

[show chassis firmware \(M10 Router\) on page 695](#)
[show chassis firmware \(M20 Router\) on page 695](#)
[show chassis firmware \(M40 Router\) on page 696](#)
[show chassis firmware \(M120 Router\) on page 696](#)
[show chassis firmware \(M160 Router\) on page 696](#)
[show chassis firmware \(MX104 Router\) on page 696](#)
[show chassis firmware \(MX240 Router\) on page 696](#)
[show chassis firmware \(MX480 Router\) on page 697](#)
[show chassis firmware \(MX960 Router\) on page 697](#)
[show chassis firmware \(MX2010 Router\) on page 697](#)
[show chassis firmware \(MX2020 Router\) on page 697](#)
[show chassis firmware \(MX240, MX480, MX960 Router with Application Services Modular Line Card\) on page 698](#)
[show chassis firmware \(EX4200 Switch\) on page 698](#)
[show chassis firmware \(EX8200 Switch\) on page 698](#)
[show chassis firmware \(EX9200 Switch\) on page 699](#)
[show chassis firmware lcc \(TX Matrix Router\) on page 699](#)
[show chassis firmware scc \(TX Matrix Router\) on page 699](#)
[show chassis firmware \(TX Matrix Plus Router\) on page 699](#)
[show chassis firmware lcc \(TX Matrix Plus Router\) on page 701](#)
[show chassis firmware sfc \(TX Matrix Plus Router\) on page 701](#)
[show chassis firmware \(QFX Series and OCX Series\) on page 702](#)
[show chassis firmware \(PTX1000 Packet Transport Routers\) on page 702](#)
[show chassis firmware interconnect-device \(QFabric System\) on page 702](#)

[show chassis firmware \(ACX2000 Universal Access Router\) on page 702](#)
[show chassis firmware detail \(EX3300 Switch\) on page 702](#)
[show chassis firmware \(MX Routers with Media Services Blade \[MSB\]\) on page 702](#)

Output Fields Table 57 lists the output fields for the **show chassis firmware** command. Output fields are listed in the approximate order in which they appear.

Table 57: show chassis firmware Output Fields

Field Name	Field Description
Part	(MX Series, MX2010, and MX2020 routers) Chassis part name.
Type	(MX Series, MX2010, and MX2020 routers) Type of firmware: On routers: ROM or O/S . On switches: uboot or loader .
Version	(MX Series, MX2010, and MX2020 routers) Version of firmware running on the chassis part.
FPC	(<i>detail</i> option only) Number of FPC. For a standalone switch, the value is 0. For a Virtual Chassis configuration, value in the range of 0-9; refers to the member ID assigned to the switch.
AFEB	(MX104 routers) Version of the compact Forwarding Engine Board.
Boot	(<i>detail</i> option only) Version of the SYSPLD.
PoE	(<i>detail</i> option only) Version of the PoE firmware.
PFE-<number>	(<i>detail</i> option only) Version of the PFE used in the switch.
PHY-	(<i>detail</i> option only) Version of the physical layer device (PHY) used in the switch.
microcode	(<i>detail</i> option only) Microcode of the physical layer devices (PHY) used in the switch.
uboot	(<i>detail</i> option only) Version of the u-boot used in the switch.
loader	(<i>detail</i> option only) Version of the loader used in the switch.

Sample Output

show chassis firmware (M10 Router)

```
user@host> show chassis firmware
Part      Type      Version
Forwarding engine board  ROM      Juniper ROM Monitor Version 4.1b2
                                O/S      Version 4.1I1 by tlim on 2000-04-24 11:27
```

show chassis firmware (M20 Router)

```
user@host> show chassis firmware
```

Part	Type	Version
System switch board	ROM	Juniper ROM Monitor Version 3.4b26
	O/S	Version 3.4I16 by smackie on 2000-02-29 2
FPC 1	ROM	Juniper ROM Monitor Version 3.0b1
	O/S	Version 3.4I4 by smackie on 2000-02-25 21
FPC 2	ROM	Juniper ROM Monitor Version 3.0b1
	O/S	Version 3.4I4 by smackie on 2000-02-25 21

show chassis firmware (M40 Router)

```
user@host> show chassis firmware
```

Part	Type	Version
System control board	ROM	Juniper ROM Monitor Version 2.0i126Copyri
	O/S	Version 2.0i1 by root on Thu Jul 23 00:51
FPC 5	ROM	Juniper ROM Monitor Version 2.0i49Copyrig
	O/S	Version 2.0i1 by root on Thu Jul 23 00:59

show chassis firmware (M120 Router)

```
user@host> show chassis firmware
```

FPC 2	ROM	Juniper ROM Monitor Version 8.0b29
	O/S	Version 8.2B1 by builder on 2006-10-18 16:2
FPC 3	ROM	Juniper ROM Monitor Version 8.0b29
	O/S	Version 8.2B1 by builder on 2006-10-18 16:2
FPC 4	ROM	Juniper ROM Monitor Version 8.0b29
	O/S	Version 8.2B1 by builder on 2006-10-18 16:2
FEB 3	ROM	Juniper ROM Monitor Version 8.0b29
	O/S	Version 8.2B1 by builder on 2006-10-18 16:1
FEB 4	ROM	Juniper ROM Monitor Version 8.0b29
	O/S	Version 8.2B1 by builder on 2006-10-18 16:1

show chassis firmware (M160 Router)

```
user@host> show chassis firmware
```

Part	Type	Version
SFM 0	ROM	Juniper ROM Monitor Version 4.0b2
	O/S	Version 4.0I1 by tlim on 2000-02-29 11:50
SFM 1	ROM	Juniper ROM Monitor Version 4.0b2
	O/S	Version 4.0I1 by tlim on 2000-02-29 11:50
FPC 0	ROM	Juniper ROM Monitor Version 4.0b2
	O/S	Version 4.0I1 by tlim on 2000-02-29 11:56
FPC 1	ROM	Juniper ROM Monitor Version 4.0b2
	O/S	Version 4.0I1 by tlim on 2000-02-29 11:56
FPC 2	ROM	Juniper ROM Monitor Version 4.0b3
	O/S	Version 4.0I1 by tlim on 2000-02-29 11:56

show chassis firmware (MX104 Router)

```
user@host > show chassis firmware
```

Part	Type	Version
FPC 0	ROM	Juniper ROM Monitor Version 13.1b24
	O/S	Version 13.2-20130514.1 by builder on 2013-
FPC 1	ROM	Juniper ROM Monitor Version 13.1b24
	O/S	Version 13.2-20130514.1 by builder on 2013-
FPC 2	ROM	Juniper ROM Monitor Version 13.1b24
	O/S	Version 13.2-20130514.1 by builder on 2013-
AFEB	ROM	Juniper ROM Monitor Version 13.1b24
	O/S	Version 13.2-20130514.1 by builder on 2013-

show chassis firmware (MX240 Router)

```
user@host> show chassis firmware
```


Part	Type	Version
FPC 1	ROM	Juniper ROM Monitor Version 8.3b1
	O/S	Version 9.0-20080103.0 by builder on 2008-0
FPC 2	ROM	Juniper ROM Monitor Version 8.3b1
	O/S	Version 9.0-20080103.0 by builder on 2008-0

show chassis firmware (MX480 Router)

```
user@host> show chassis firmware
```

Part	Type	Version
FPC 1	ROM	Juniper ROM Monitor Version 8.3b1
	O/S	Version 9.0-20070916.3 by builder on 2007-0

show chassis firmware (MX960 Router)

```
user@host> show chassis firmware
```

Part	Type	Version
FPC 4	ROM	Juniper ROM Monitor Version 8.0b8
	O/S	Version 8.2I59 by artem on 2006-10-31 19:22
FPC 7	ROM	Juniper ROM Monitor Version 8.2b1
	O/S	Version 8.2-20061026.1 by builder on 2006-1

show chassis firmware (MX2010 Router)

```
user@host> show chassis firmware
```

Part	Type	Version
FPC 0	ROM	Juniper ROM Monitor Version 12.3b1
	O/S	Version 12.3-20121220.0 by builder on 2012-
FPC 1	ROM	Juniper ROM Monitor Version 10.1b3
	O/S	Version 12.3-20121220.0 by builder on 2012-
FPC 2	ROM	Juniper ROM Monitor Version 10.1b3
	O/S	Version 12.3-20121220.0 by builder on 2012-
FPC 3	ROM	Juniper ROM Monitor Version 10.1b3
	O/S	Version 12.3-20121220.0 by builder on 2012-
FPC 4	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20121220.0 by builder on 2012-
FPC 5	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20121220.0 by builder on 2012-
FPC 6	ROM	Juniper ROM Monitor Version 10.4b1
	O/S	Version 12.3-20121220.0 by builder on 2012-
FPC 7	ROM	Juniper ROM Monitor Version 10.1b3
	O/S	Version 12.3-20121220.0 by builder on 2012-
FPC 8	ROM	Juniper ROM Monitor Version 10.4b1
	O/S	Version 12.3-20121220.0 by builder on 2012-
FPC 9	ROM	Juniper ROM Monitor Version 10.4b1
	O/S	Version 12.3-20121220.0 by builder on 2012-
SPMB 0	ROM	Juniper ROM Monitor Version 12.1b1
	O/S	Version 12.3-20121220.0 by builder on 2012-
SPMB 1	ROM	Juniper ROM Monitor Version 12.1b1
	O/S	Version 12.3-20121220.0 by builder on 2012-

show chassis firmware (MX2020 Router)

```
user@host> show chassis firmware
```

Part	Type	Version
FPC 0	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 1	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 2	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 3	ROM	Juniper ROM Monitor Version 10.0b39

	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 4	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 5	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 6	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 7	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 8	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 9	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 10	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 11	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 12	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 13	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 14	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 15	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 16	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 17	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 18	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 19	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
SPMB 0	ROM	Juniper ROM Monitor Version 12.1b1
	O/S	Version 12.3-20130415.0 by builder on 2013-
SPMB 1	ROM	Juniper ROM Monitor Version 12.1b1
	O/S	Version 12.3-20130415.0 by builder on 2013-

show chassis firmware (MX240, MX480, MX960 Router with Application Services Modular Line Card)

```
user@host> show chassis firmware
```

Part	Type	Version
FPC 1	ROM	Juniper ROM Monitor Version 12.1b1
	O/S	Version 12.2I21 by manish on 2012-06-19 17:

show chassis firmware (EX4200 Switch)

```
user@switch> show chassis firmware
```

Part	Type	Version
FPC 0	uboot	U-Boot 1.1.6 (Feb 6 2008 - 11:27:42)
	loader	FreeBSD/PowerPC U-Boot bootstrap loader 2.1
FPC 1	uboot	U-Boot 1.1.6 (Feb 6 2008 - 11:27:42)
	loader	FreeBSD/PowerPC U-Boot bootstrap loader 2.1
FPC 2	uboot	U-Boot 1.1.6 (Feb 6 2008 - 11:27:42)
	loader	FreeBSD/PowerPC U-Boot bootstrap loader 2.1

show chassis firmware (EX8200 Switch)

```
user@switch> show chassis firmware
```

Part	Type	Version
FPC 0	U-Boot loader	U-Boot 1.1.6 (Mar 25 2009 - 06:13:12) 2.4.0 FreeBSD/PowerPC U-Boot bootstrap loader 2.2
FPC 3	U-Boot loader	U-Boot 1.1.6 (Dec 4 2009 - 13:17:34) 3.1.0 FreeBSD/PowerPC U-Boot bootstrap loader 2.2
FPC 5	U-Boot loader	U-Boot 1.1.6 (Mar 25 2009 - 06:13:12) 2.4.0 FreeBSD/PowerPC U-Boot bootstrap loader 2.2
FPC 7	U-Boot loader	U-Boot 1.1.6 (Feb 6 2009 - 05:31:46) 2.4.0 FreeBSD/PowerPC U-Boot bootstrap loader 2.2
Routing Engine 0	U-Boot loader	U-Boot 1.1.6 (Mar 25 2009 - 06:13:12) 2.4.0 FreeBSD/PowerPC U-Boot bootstrap loader 2.2
Routing Engine 1	U-Boot loader	U-Boot 1.1.6 (Mar 25 2009 - 06:13:12) 2.4.0 FreeBSD/PowerPC U-Boot bootstrap loader 2.2

show chassis firmware (EX9200 Switch)

```
user@switch> show chassis firmware
```

Part	Type	Version
FPC 2	ROM	Juniper ROM Monitor Version 11.4b2
	O/S	Version 14.1I20140312_0741_bavig by bavig o
FPC 3	ROM	Juniper ROM Monitor Version 10.4b1
	O/S	Version 14.1I20140312_0741_bavig by bavig o

show chassis firmware lcc (TX Matrix Router)

```
user@host> show chassis firmware lcc 0
lcc0-re0:
```

Part	Type	Version
FPC 1	ROM	Juniper ROM Monitor Version 6.4b18
	O/S	Version 7.0-20040804.0 by builder on 2004-0
FPC 2	ROM	Juniper ROM Monitor Version 6.4b20
	O/S	Version 7.0-20040804.0 by builder on 2004-0
SPMB 0	ROM	Juniper ROM Monitor Version 6.4b18
	O/S	Version 7.0-20040804.0 by builder on 2004-0

show chassis firmware scc (TX Matrix Router)

```
user@host> show chassis firmware scc
scc-re0:
```

Part	Type	Version
SPMB 0	ROM	Juniper ROM Monitor Version 6.4b18
	O/S	Version 7.0-20040804.0 by builder on 2004-0

show chassis firmware (TX Matrix Plus Router)

```
user@host> show chassis firmware
sfc0-re0:
```

Part	Type	Version
Global FPC 4		
Global FPC 6		
Global FPC 7		
Global FPC 12		
Global FPC 14		
Global FPC 15		
Global FPC 20		
Global FPC 21		

```

Global FPC 22
Global FPC 23
Global FPC 24
Global FPC 25
Global FPC 26
Global FPC 28
Global FPC 29
Global FPC 31
SPMB 0          ROM      Juniper ROM Monitor Version 9.5b1
                  O/S      Version 9.6-20090507.0 by builder on 2009-0
SPMB 1          ROM      Juniper ROM Monitor Version 9.5b1
                  O/S      Version 9.6-20090507.0 by builder on 2009-0

```

lcc0-re1:

```

-----
Part          Type      Version
FPC 4          ROM      Juniper ROM Monitor Version 9.0b2
                  O/S      Version 9.6-20090507.0 by builder on 2009-0
FPC 6          ROM      Juniper ROM Monitor Version 9.0b2
                  O/S      Version 9.6-20090507.0 by builder on 2009-0
FPC 7          ROM      Juniper ROM Monitor Version 9.0b2
                  O/S      Version 9.6-20090507.0 by builder on 2009-0
SPMB 0          ROM      Juniper ROM Monitor Version 9.5b1
                  O/S      Version 9.6-20090507.0 by builder on 2009-0
SPMB 1          ROM      Juniper ROM Monitor Version 9.5b1
                  O/S      Version 9.6-20090507.0 by builder on 2009-0

```

lcc1-re1:

```

-----
Part          Type      Version
FPC 4          ROM      Juniper ROM Monitor Version 9.0b2
                  O/S      Version 9.6-20090507.0 by builder on 2009-0
FPC 6          ROM      Juniper ROM Monitor Version 9.0b2
                  O/S      Version 9.6-20090507.0 by builder on 2009-0
FPC 7          ROM      Juniper ROM Monitor Version 9.0b2
                  O/S      Version 9.6-20090507.0 by builder on 2009-0
SPMB 0          ROM      Juniper ROM Monitor Version 9.5b1
                  O/S      Version 9.6-20090507.0 by builder on 2009-0
SPMB 1          ROM      Juniper ROM Monitor Version 9.5b1
                  O/S      Version 9.6-20090507.0 by builder on 2009-0

```

lcc2-re1:

```

-----
Part          Type      Version
FPC 4          ROM      Juniper ROM Monitor Version 9.0b2
                  O/S      Version 9.6-20090507.0 by builder on 2009-0
FPC 5          ROM      Juniper ROM Monitor Version 9.0b2
                  O/S      Version 9.6-20090507.0 by builder on 2009-0
FPC 6          ROM      Juniper ROM Monitor Version 9.0b2
                  O/S      Version 9.6-20090507.0 by builder on 2009-0
FPC 7          ROM      Juniper ROM Monitor Version 7.5b4
                  O/S      Version 9.6-20090507.0 by builder on 2009-0
SPMB 0          ROM      Juniper ROM Monitor Version 9.5b1
                  O/S      Version 9.6-20090507.0 by builder on 2009-0
SPMB 1          ROM      Juniper ROM Monitor Version 9.5b1
                  O/S      Version 9.6-20090507.0 by builder on 2009-0

```

lcc3-re1:

```

-----
Part          Type      Version
FPC 0          ROM      Juniper ROM Monitor Version 9.0b2

```

	O/S	Version 9.6-20090507.0 by builder on 2009-0
FPC 1	ROM	Juniper ROM Monitor Version 9.0b2
	O/S	Version 9.6-20090507.0 by builder on 2009-0
FPC 2	ROM	Juniper ROM Monitor Version 9.0b2
	O/S	Version 9.6-20090507.0 by builder on 2009-0
FPC 4	ROM	Juniper ROM Monitor Version 7.5b4
	O/S	Version 9.6-20090507.0 by builder on 2009-0
FPC 5	ROM	Juniper ROM Monitor Version 9.0b2
	O/S	Version 9.6-20090507.0 by builder on 2009-0
FPC 7	ROM	Juniper ROM Monitor Version 9.0b2
	O/S	Version 9.6-20090507.0 by builder on 2009-0
SPMB 0	ROM	Juniper ROM Monitor Version 9.5b1
	O/S	Version 9.6-20090507.0 by builder on 2009-0
SPMB 1	ROM	Juniper ROM Monitor Version 9.5b1
	O/S	Version 9.6-20090507.0 by builder on 2009-0

show chassis firmware lcc (TX Matrix Plus Router)

```
user@host> show chassis firmware lcc 0
lcc0-re1:
```

Part	Type	Version
FPC 4	ROM	Juniper ROM Monitor Version 9.0b2
	O/S	Version 9.6-20090507.0 by builder on 2009-0
FPC 6	ROM	Juniper ROM Monitor Version 9.0b2
	O/S	Version 9.6-20090507.0 by builder on 2009-0
FPC 7	ROM	Juniper ROM Monitor Version 9.0b2
	O/S	Version 9.6-20090507.0 by builder on 2009-0
SPMB 0	ROM	Juniper ROM Monitor Version 9.5b1
	O/S	Version 9.6-20090507.0 by builder on 2009-0
SPMB 1	ROM	Juniper ROM Monitor Version 9.5b1
	O/S	Version 9.6-20090507.0 by builder on 2009-0

show chassis firmware sfc (TX Matrix Plus Router)

```
user@host> show chassis firmware sfc 0
sfc0-re0:
```

Part	Type	Version
Global FPC 4		
Global FPC 6		
Global FPC 7		
Global FPC 12		
Global FPC 14		
Global FPC 15		
Global FPC 20		
Global FPC 21		
Global FPC 22		
Global FPC 23		
Global FPC 24		
Global FPC 25		
Global FPC 26		
Global FPC 28		
Global FPC 29		
Global FPC 31		
SPMB 0	ROM	Juniper ROM Monitor Version 9.5b1
	O/S	Version 9.6-20090507.0 by builder on 2009-0
SPMB 1	ROM	Juniper ROM Monitor Version 9.5b1
	O/S	Version 9.6-20090507.0 by builder on 2009-0

show chassis firmware (QFX Series and OCX Series)

```
user@switch> show chassis firmware
Part                               Type      Version
FPC 0
Routing Engine 0                  U-Boot    U-Boot 1.1.6 (Sep 15 2010 - 02:11:11) 1.0.5
loader                            FreeBSD/MIPS U-Boot bootstrap loader 0.1
```

show chassis firmware (PTX1000 Packet Transport Routers)

```
user@host> show chassis firmware
Part                               Type      Version
FPC 0
Loader                            FreeBSD/i386 bootstrap loader 1.2
BIOS                              V0018.2U
EC FPGA                           2.0
MAIN_CPLD                         1.f
MEZZ_CPLD                         1.f
RE FPGA                           2.3
```

show chassis firmware interconnect-device (QFabric System)

```
user@switch> show chassis firmware interconnect-device interconnect1
Part                               Type      Version
Routing Engine 0                  U-Boot    U-Boot 1.1.6 (May 10 2011 - 04:52:59) 1.1.1
loader                            FreeBSD/MIPS U-Boot bootstrap loader 0.1
Routing Engine 1                  U-Boot    U-Boot 1.1.6 (May 10 2011 - 04:52:59) 1.1.1
loader                            FreeBSD/MIPS U-Boot bootstrap loader 0.1
```

show chassis firmware (ACX2000 Universal Access Router)

```
user@switch> show chassis firmware
Part                               Type      Version
FPC                               O/S       Version 12.2I13 by jisjoy on 2012-05-29 06:
FEB                               O/S       Version 12.2I13 by jisjoy on 2012-05-29 06:
```

show chassis firmware detail (EX3300 Switch)

```
user@switch> show chassis firmware detail
FPC 0
  Boot SYSPLD                     3
  PoE firmware                    4.1.6
  PFE-0                           3
  PFE-1                           3
  PHY
    microcode                     0x514
  Boot Firmware
    uboot                         U-Boot 1.1.6 (Aug 21 2011 - 01:45:26) 1.0.0
    loader                        FreeBSD/arm U-Boot loader 1.0
```

show chassis firmware (MX Routers with Media Services Blade [MSB])

```
user@switch> show chassis firmware
Part                               Type      Version
FPC 1                             ROM        Juniper ROM Monitor Version 12.1b1
O/S                               Version 12.2I21 by manish on 2012-06-19 17:
```

show chassis fpc

List of Syntax	Syntax on page 703 Syntax (EX Series Switches) on page 703 Syntax (T4000 Routers) on page 703 Syntax (TX Matrix and TX Matrix Plus Routers) on page 703 Syntax (MX Series Router) on page 703 Syntax (MX104, MX2010, and MX2020 3D Universal Edge Routers) on page 703 Syntax (QFX Series) on page 703 Syntax (OCX Series) on page 703 Syntax (PTX Series Packet Transport Routers) on page 703 Syntax (PTX Series Packet Transport Switches) on page 704
Syntax	<pre>show chassis fpc <detail <slot>> <pic-status <slot>></pre>
Syntax (EX Series Switches)	<pre>show chassis fpc <detail <fpc-slot>> <pic-status <fpc-slot>> <fpc-slot></pre>
Syntax (T4000 Routers)	<pre>show chassis fpc <detail <fpc-slot>> <pic-status <fpc-slot>></pre>
Syntax (TX Matrix and TX Matrix Plus Routers)	<pre>show chassis fpc <detail <fpc-slot>> <pic-status <fpc-slot>> <slot></pre>
Syntax (MX Series Router)	<pre>show chassis fpc <detail <slot>> <pic-status <slot>> <all-members> <local> <member member-id></pre>
Syntax (MX104, MX2010, and MX2020 3D Universal Edge Routers)	<pre>show chassis fpc <slot> detail <detail <slot>> <pic-status <slot>> <fpc-slot></pre>
Syntax (QFX Series)	<pre>show chassis fpc <detail> <interconnect-device name <fpc-slot fpc-slot>> <node-device name></pre>
Syntax (OCX Series)	<pre>show chassis fpc <detail></pre>
Syntax (PTX Series Packet Transport Routers)	<pre>show chassis fpc <detail <fpc-slot>> <pic-status <fpc-slot>> <fpc-slot></pre>

Syntax (PTX Series Packet Transport Switches) `show chassis fpc`
`<detail <fpc-slot>> | <pic-status <fpc-slot>>`
`<fpc-slot>`

Release Information Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.1 for QFX Series.
Command introduced in Junos OS Release 12.1x48 for PTX Series Packet Transport Routers.
Command introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.
Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.
Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.
Command introduced in Junos OS Release 13.2 for MX104 3D Universal Edge Routers.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Command introduced in Junos OS Release 15.1 for MX104-40G 3D Universal Edge Routers.

Description Display status information about the installed Flexible PIC Concentrators (FPCs) and PICs.

Options **none**—Display status information for all FPCs. On a TX Matrix router, display status information for all FPCs on the attached T640 routers in the routing matrix. On a TX Matrix Plus router, display status information for all FPCs on the attached T1600 routers in the routing matrix.



NOTE: In EX8200 switches, line cards initialize Packet Forwarding Engine during start up. If an error occurs during hardware initialization, the FPCs with bad hardware parts power down after transferring the debug information to the Routing Engine. The Routing Engine marks the FPC offline, logs the error in system log messages (`/var/log/messages`), and generates an alarm to inform the user.

See the following sample output:

```
user@host> show chassis fpc
```

Utilization (%)	Temp	CPU	Utilization (%)	Memory
Slot State	(C)	Total	Interrupt	DRAM (MB) Heap
Buffer				
0 Empty				
1 Empty				
2 Empty				
3 Empty				
4 Empty				
5 Offline	---	Hard	FPC error---	
6 Empty				
7 Online	26	4	0	1024 0
32				

The following sample output shows the alarm raised for the failed FPCs.

```
user@host > show chassis alarms
```

```
4 alarms currently active
Alarm time      Class  Description
2011-03-24 00:52:51 UTC Major  FPC 5 Hard errors
2011-03-24 00:52:31 UTC Major  Fan Tray Failure
2011-03-24 00:52:31 UTC Major  Fan Tray Failure
2011-03-24 00:51:26 UTC Minor  Loss of communication with Backup
RE
```



NOTE: On T4000 routers, when you include the enhanced-mode statement at the [edit chassis network-services] hierarchy level and reboot the system, only the T4000 Type 5 FPCs present on the router become online while the remaining FPCs are offline, and FPC misconfiguration alarms are generated. The show chassis alarm command output displays FPC misconfiguration (FPC *fpc-slot* misconfig) as the reason for the generation the alarms.

The following sample output shows the FPC status after the enhanced-mode statement is configured on the T4000 router. The T4000 Type 5 FPC present in slot 5 becomes online while the remaining FPCs are offline.

```
user@host> show chassis fpc
```

Utilization (%)	Temp	CPU Utilization (%)		Memory	
Slot State	(C)	Total	Interrupt	DRAM (MB)	Heap
Buffer					
0 offline	---	FPC misconfiguration---			
1 offline	---	FPC misconfiguration---			
2 offline	---	FPC misconfiguration---			
3 Empty					
4 Empty					
5 Online	66	50	0	2816	29
27					

The following sample output shows FPC misconfiguration alarms.

```
user@host > show chassis alarms
```

```
3 alarms currently active
Alarm time      Class Description
2011-03-24 00:52:51 PST Major FPC 1 misconfig
2011-03-24 00:52:31 PST Major FPC 2 misconfig
2011-03-24 00:52:31 PST Major FPC 3 misconfig
```

detail—(Optional) Display detailed status information for all FPCs or for the FPC in the specified slot (see *fpc-slot* or *slot*).

all-members—(MX Series routers and EX Series switches only) (Optional) Display status information for all FPCs on all members of the Virtual Chassis configuration.

interconnect-device *name*—(QFabric switches only) (Optional) Display status information for all FPCs on the Interconnect device.

fpc-slot—(Optional) FPC slot number:

- (TX Matrix and TX Matrix Plus router only)—On a TX Matrix router, if you specify the number of the T640 router (line-card chassis) by using the *lcc number* option (the recommended method), replace *fpc-slot* with a value from 0 through 7. Otherwise, replace *fpc-slot* with a value from 0 through 31. Likewise, on a TX Matrix

Plus router, if you specify the number of the specified router (line-card chassis) by using the **lcc number** option (the recommended method), replace **fpc-slot** with a value from 0 through 7. Otherwise, replace **fpc-slot** with a value from 0 through 31. For example, the following commands have the same result:

```
user@host> show chassis fpc detail 1 lcc 1
user@host> show chassis fpc detail 9
```

- M120 router—Replace **fpc-slot** with a value from 0 through 5.
- MX80 router—Replace **fpc-slot** with a value from 0 through 1.
- MX104 and MX104-40G routers—Replace **fpc-slot** with a value from 0 through 2.
- MX240 router—Replace **fpc-slot** with a value from 0 through 2.
- MX480 router—Replace **fpc-slot** with a value from 0 through 5.
- MX-960 router—Replace **fpc-slot** with a value from 0 through 11.
- MX2010 router—Replace **fpc-slot-number** with a value from 0 through 9.
- MX2020 router—Replace **fpc-slot-number** with a value from 0 through 19.
- Other routers—Replace **fpc-slot** with a value from 0 through 7.
- EX Series switches:
 - EX3200 switches and EX4200 standalone switches—Replace **fpc-slot** with 0.
 - EX4200 switches in a Virtual Chassis configuration—Replace **fpc-slot** with a value from 0 through 9.
 - EX6210 switches—Replace **fpc-slot** with a value from 0 through 9.
 - EX8208 switches—Replace **fpc-slot** with a value from 0 through 7.
 - EX8216 switches—Replace **fpc-slot** with a value from 0 through 15.
- QFX Series:
 - QFX Series and OCX Series switches—Replace **fpc-slot** with 0.
 - QFabric systems—Replace **fpc-slot** with 0 through 31 on the Interconnect device.
- PTX Series Packet Transport Switches:
 - PTX5000 Packet Transport Router—Replace **fpc-slot** with a value from 0 through 7.
- ACX Series Universal Access Routers:
 - ACX1000 and ACX2000 Universal Access Routers—Replace **fpc-slot** with 0.

local—(MX Series routers only) (Optional) Display status information for all FPCs on the local Virtual Chassis member.

member member-id—(MX Series routers and EX Series switches only) (Optional) Display status information for all FPCs on the specified member of the Virtual Chassis configuration. Replace **member-id** with a value of 0 or 1.

node-device *name*—(QFabric switches only) (Optional) Display status information for each Node device. Each Node device is equivalent to an FPC.

pic-status—(Optional) Display status information for all PICs or for the PIC in the specified slot (see *fpc-slot*).



NOTE: On T1600 routers, Type 4 FPCs with ASICs based on the SL2.0 chipset do not support the 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (10x10GE [LAN/WAN] SFPP). If you issue the `show chassis fpc` command with the `pic-status` option, the CLI displays the string “Not Supported” for 10x10GE (LAN/WAN) SFPP PICs installed on such FPCs. The following is a sample output:

```
user@host> show chassis fpc pic-status
```

```
Slot 0  Online      E2-FPC Type 1
  PIC 0  Online      1x G/E SFP, 1000 BASE
  PIC 1  Online      Adaptive Services-II
  PIC 2  Online      1x G/E IQ, 1000 BASE
  PIC 3  Online      1x G/E IQ, 1000 BASE
Slot 1  Online      FPC Type 3-ES
  PIC 0  Present     UNUSED- Not Supported
Slot 2  Online      FPC Type 4-ES
  PIC 0  Offline     4x OC-192 SONET XFP
  PIC 1  Present     10x10GE(LAN/WAN) SFPP- Not Supported
<<<<<<
Slot 4  Offline     FPC Type 1-ES
Slot 5  Offline     FPC Type 2-ES
Slot 6  Online      E2-FPC Type 3
  PIC 0  Online      1x OC-192 SONET XFP
  PIC 1  Online      4x OC-48 SONET
  PIC 2  Online      4x OC-48 SONET
  PIC 3  Online      MultiServices 500
Slot 7  Online      FPC Type 4-ES
  PIC 0  Online      4x 10GE (LAN/WAN) XFP
  PIC 1  Online      4x 10GE (LAN/WAN) XFP
```

In addition, an entry is logged in the system log messages (`/var/log/messages`) that the PIC is not supported. The following is a sample message logged in the system log:

```
Apr  5 08:47:36  router1 chassisd[2770]: CHASSISD_UNSUPPORTED_PIC:
PIC 1 in FPC 2 (type 763, version 257) is not supported
```

If you see this issue, contact Juniper Networks Technical Assistance Center (JTAC) for a possible fix. For more information about this issue and a possible solution, see [PSN-2010-03-696](https://www.juniper.net/psn/2010-03-696).



NOTE: When there is a double-bit ECC error in a network processor's memory, the Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP or Channelized E1/T1 Circuit Emulation MIC is switched to the offline state.

```
user@host> show chassis fpc pic-status
```

```
Slot 1   Online      MPC Type 2 3D Q
PIC 0   Offline      1xC0C12/4xC0C3 CH-CE- ECC error detected
```

lcc number—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

Required Privilege Level view

- Related Documentation**
- [request chassis fpc on page 366](#)
 - [show chassis fpc-feb-connectivity](#)
 - [show chassis fabric fpcs on page 585](#)
 - *Configuring the Junos OS to Resynchronize FPC Sequence Numbers with Active FPCs when an FPC Comes Online*
 - *MX960 Flexible PIC Concentrator Description*
 - *ACX2000 and ACX2100 Routers Hardware and CLI Terminology Mapping*
 - *enhanced-mode*

- List of Sample Output**
- [show chassis fpc \(EX6210 Switch\) on page 713](#)
 - [show chassis fpc \(M10 Router\) on page 714](#)
 - [show chassis fpc \(M20 Router\) on page 714](#)
 - [show chassis fpc detail \(M Series Routers\) on page 714](#)
 - [show chassis fpc detail \(MX80 Router\) on page 714](#)
 - [show chassis fpc \(MX104 Router\) on page 714](#)
 - [show chassis fpc detail \(MX104 Router\) on page 715](#)
 - [show chassis fpc pic-status \(MX104 Router\) on page 715](#)

[show chassis fpc \(MX104-40G Router\) on page 715](#)
[show chassis fpc detail \(MX104-40G Router\) on page 715](#)
[show chassis fpc pic-status \(MX104-40G Router\) on page 716](#)
[show chassis fpc pic-status \(MX104-40G Router with Upgrade License\) on page 716](#)
[show chassis fpc \(MX240 Router\) on page 716](#)
[show chassis fpc \(MX480 Router\) on page 717](#)
[show chassis fpc detail \(EX9200 Switch\) on page 717](#)
[show chassis fpc \(MX480 Router\) on page 717](#)
[show chassis fpc \(MX480 Router with 100-Gigabit Ethernet CFP\) on page 717](#)
[show chassis fpc pic-status \(MX480 Router with 100-Gigabit Ethernet CFP\) on page 718](#)
[show chassis fpc pic-status \(EX Series Switch\) on page 718](#)
[show chassis fpc \(MX480 Router with MPC4E\) on page 718](#)
[show chassis fpc detail \(MX480 Router with MPC4E\) on page 718](#)
[show chassis fpc \(MX480 Router with MPC4E\) on page 719](#)
[show chassis fpc detail \(MX480 Router with MPC4E\) on page 719](#)
[show chassis fpc \(MX960 Router\) on page 719](#)
[show chassis fpc \(MX960 Router with MPC5EQ\) on page 720](#)
[show chassis fpc detail \(MX960 Router with MPC5EQ\) on page 720](#)
[show chassis fpc pic-status \(MX960 Router with MPC5EQ\) on page 722](#)
[show chassis fpc \(MX240, MX480, MX960 Routers with Application Services Modular Line Card\) on page 722](#)
[show chassis fpc \(MX240, MX480, MX960 with Application Services Modular Line Card\) on page 723](#)
[show chassis fpc \(MX240, MX480, MX960, MX2010, and MX2020 3D Universal Edge Routers Configured for Dynamic Power Management\) on page 723](#)
[show chassis fpc \(MX2010 Routers\) on page 723](#)
[show chassis fpc \(MX2020 Routers\) on page 723](#)
[show chassis fpc \(MX2020 Router with MPC4E\) on page 724](#)
[show chassis fpc detail \(MX2020 Router with MPC4E\) on page 724](#)
[show chassis fpc \(MX2020 Router with MPC5EQ and MPC6E\) on page 725](#)
[show chassis fpc detail \(MX2020 Router with MPC5EQ and MPC6E\) on page 725](#)
[show chassis fpc pic-status \(MX2020 Router with MPC5EQ and MPC6E\) on page 727](#)
[show chassis fpc detail \(MX Series Routers\) on page 728](#)
[show chassis fpc detail \(EX Series Switches\) on page 728](#)
[show chassis fpc \(Hardware Not Supported\) on page 728](#)
[show chassis fpc detail \(Hardware Not Supported\) on page 729](#)
[show chassis fpc pic-status on page 729](#)
[show chassis fpc pic-status \(M Series Routers\) on page 729](#)
[show chassis fpc pic-status \(M120 Router\) on page 730](#)
[show chassis fpc pic-status \(MX240, MX480, and MX960 Routers with Application Services Modular Line Card\) on page 730](#)
[show chassis fpc lcc \(TX Matrix Router\) on page 730](#)
[show chassis fpc pic-status \(TX Matrix Router\) on page 730](#)
[show chassis fpc pic-status lcc \(TX Matrix Router\) on page 731](#)
[show chassis fpc \(TX Matrix Plus Router\) on page 731](#)
[show chassis fpc lcc \(TX Matrix Plus Router\) on page 732](#)
[show chassis fpc detail \(TX Matrix Plus Router\) on page 732](#)
[show chassis fpc pic-status \(TX Matrix Plus Router\) on page 734](#)
[show chassis fpc \(T1600 Router\) on page 735](#)

[show chassis fpc detail \(T1600 Router\) on page 736](#)
[show chassis fpc <fpc-slot> \(EX Series Switch\) on page 736](#)
[show chassis fpc slot \(T1600 Router\) on page 736](#)
[show chassis fpc pic-status \(T1600 Router\) on page 736](#)
[show chassis fpc \(T4000 Router\) on page 737](#)
[show chassis fpc detail \(T4000 Router\) on page 737](#)
[show chassis fpc pic-status \(T4000 Router\) on page 738](#)
[show chassis fpc \(QFX Series\) on page 738](#)
[show chassis fpc detail \(QFX3500 Switches\) on page 738](#)
[show chassis fpc pic-status \(QFX3500 Switches\) on page 738](#)
[show chassis fpc interconnect-device \(QFabric Switch\) on page 738](#)
[show chassis fpc interconnect-device \(QFabric Switch\) on page 739](#)
[show chassis fpc interconnect-device detail \(QFabric Switch\) on page 739](#)
[show chassis fpc pic-status interconnect-device \(QFabric Switch\) on page 739](#)
[show chassis fpc pic-status node-device \(QFabric Switch\) on page 740](#)
[show chassis fpc \(PTX5000 Packet Transport Switch\) on page 740](#)
[show chassis fpc detail \(PTX5000 Packet Transport Switch\) on page 740](#)
[show chassis fpc pic-status \(PTX5000 Packet Transport Switch\) on page 741](#)
[show chassis fpc \(ACX2000 Universal Access Router\) on page 741](#)
[show chassis fpc 0 \(ACX2000 Universal Access Router\) on page 741](#)
[show chassis fpc detail \(ACX2000 Universal Access Router\) on page 741](#)
[show chassis fpc pic-status \(ACX2000 Universal Access Router\) on page 742](#)
[show chassis FPC 1 \(MX Routers with Media Services Blade \[MSB\]\) on page 742](#)
[show chassis FPC 1 detail \(MX Routers with Media Services Blade \[MSB\]\) on page 742](#)

Output Fields Table 58 lists the output fields for the **show chassis fpc** command. Output fields are listed in the approximate order in which they appear.

Table 58: show chassis fpc Output Fields

Field Name	Field Description	Level of Output
Slot or Slot State	<p>Slot number and state. The state can be one of the following conditions:</p> <ul style="list-style-type: none"> • Dead—Held in reset because of errors. • Diag—Slot is being ignored while the FPC is running diagnostics. • Dormant—Held in reset. • Empty—No FPC is present. • Offline—(PTX Series Packet Transport Routers only) One of the following two states is displayed: <ul style="list-style-type: none"> • FPC offlined due to unreachable destinations • FPC Offlined due to degraded FPC action • Online—FPC is online and running. • Present—FPC is detected by the chassis daemon but either is not supported by the current version of Junos OS or is inserted in the wrong slot. The output also states either Hardware Not Supported or Hardware Not In Right Slot. The FPC is coming up but not yet online. • Probed—Probe is complete; awaiting restart of the Packet Forwarding Engine. • Probe-wait—Waiting to be probed. 	all levels
Logical slot	Slot number.	all levels

Table 58: show chassis fpc Output Fields (*continued*)

Field Name	Field Description	Level of Output
Temp (C) or Temperature	Temperature of the air passing by the FPC, in degrees Celsius or in both Celsius and Fahrenheit.	all levels all levels
Temperature (PTX Series)	On PTX Series Packet Transport Routers, temperature details are provided in degrees Celsius and Fahrenheit. Output includes: <ul style="list-style-type: none"> • Temperature (PMB)—Temperature of the air passing by the Processor Mezzanine Board (PMB) at the bottom of the FPC. • Temperature (Intake)—Temperature of the air flowing into the chassis. • Temperature (Exhaust)—Exhaust temperatures for multiple zones (Exhaust A and Exhaust B). • Temperature (TLn)—Temperature of the specified Lookup ASIC (TL) of the packet forwarding engine on the FPC. • Temperature (TQn)—Temperature of the specified Queuing and Memory Interface ASIC (TQ) of the packet forwarding engine on the FPC. 	detail
Total CPU Utilization (%)	Total percentage of CPU being used by the FPC's processor.	all levels
Interrupt CPU Utilization (%)	Of the total CPU being used by the FPC's processor, the percentage being used for interrupts.	none specified
1 min CPU utilization (%)	Information about the Routing Engine's CPU utilization in the past 1 minute.	none specified
NOTE: Supported only on MX240, MX480, MX960, MX2010, and MX2020.		
5 min CPU utilization (%)	Information about the Routing Engine's CPU utilization in the past 5 minutes.	none specified
NOTE: Supported only on MX240, MX480, MX960, MX2010, and MX2020.		
15 min CPU utilization (%)	Information about the Routing Engine's CPU utilization in the past 5 minutes.	none specified
NOTE: Supported only on MX240, MX480, MX960, MX2010, and MX2020.		
Memory DRAM (MB)	Total DRAM, in megabytes, available to the FPC's processor.	none specified

Table 58: show chassis fpc Output Fields (*continued*)

Field Name	Field Description	Level of Output
Heap Utilization (%)	Percentage of heap space (dynamic memory) being used by the FPC's processor. If this number exceeds 80 percent, there may be a software problem (memory leak). NOTE: On MX Series routers and EX Series switches in a broadband edge environment, heap utilization levels higher than 70 percent can affect unified ISSU, router stability, or scaling capability.	none specified
Buffer Utilization (%)	Percentage of buffer space being used by the FPC's processor for buffering internal messages.	none specified
Total CPU DRAM	Amount of DRAM available to the FPC's CPU.	detail
Total RLDRAM	Amount of reduced latency dynamic random access memory (RLDRAM) available to the FPC CPU.	detail
Total DDR DRAM	Amount of double data rate dynamic random access memory (DDR DRAM) available to the FPC CPU.	detail
Total SRAM	Amount of static RAM (SRAM) used by the FPC's CPU.	detail
Total SDRAM	Total amount of memory used for storing packets and notifications.	detail
I/O Manager ASICs information	I/O Manager version number, manufacturer, and part number.	detail
Start time	Time when the Routing Engine detected that the FPC was running.	detail
Uptime	How long the Routing Engine has been connected to the FPC and, therefore, how long the FPC has been up and running.	detail
PIC type	(pic-status output only) Type of PIC.	none specified

Sample Output

show chassis fpc (EX6210 Switch)

```

user@switch> show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%)	Memory DRAM (MB)	Utilization (%)
			Total Interrupt		Heap Buffer
0	Empty				
1	Online	7	5 0	1024	0 32
2	Empty				
3	Empty				
4	Online	25	17 2	2048	0 30
5	Online	25	3 0	2048	0 24
6	Online	6	5 0	1024	0 32
7	Empty				
8	Empty				
9	Online	8	7 0	1024	0 32

show chassis fpc (M10 Router)

```

user@host> show chassis fpc
FPC status:

Slot State      Temp
              (C)
0  Online       27
1  Online       28

```

show chassis fpc (M20 Router)

```

user@host> show chassis fpc
FPC status:

Slot State      Temp  CPU Utilization (%)  Memory  Utilization (%)
              (C)  Total  Interrupt  DRAM (MB)  Heap    Buffer
0  Empty         0      0      0          0         0       0
1  Online        38      0      0          8         0       4
2  Online        35      0      0          8         0       3
3  Empty         0      0      0          0         0       0

```

show chassis fpc detail (M Series Routers)

```

user@host> show chassis fpc detail 1
Slot 1 information:
State                               Online
Temperature                         48 degrees C
Total CPU DRAM                     32 MB
Total SRAM                         4 MB
Total SDRAM                        256 MB
I/O Manager ASICs information      Version 2.0, Foundry IBM, Part number 0
I/O Manager ASICs information      Version 2.0, Foundry IBM, Part number 0
Start time                         2000-02-08 02:18:49 UTC
Uptime                             14 hours, 41 minutes, 41 seconds

```

show chassis fpc detail (MX80 Router)

```

user@host> show chassis fpc detail
Slot 0 information:
State                               Online
Temperature                         47 degrees C / 116 degrees F
Total CPU DRAM                     1024 MB
Total SRAM                         331 MB
Total SDRAM                        1280 MB
Start time                         2010-02-08 12:25:33 PST
Uptime                             2 hours, 13 minutes, 19 seconds
Slot 1 information:
State                               Online
Temperature                         47 degrees C / 116 degrees F
Total CPU DRAM                     1024 MB
Total SRAM                         331 MB
Total SDRAM                        1280 MB
Start time                         2010-02-08 12:25:33 PST
Uptime                             2 hours, 13 minutes, 19 seconds

```

show chassis fpc (MX104 Router)

```

user@host> show chassis fpc
Temp CPU Utilization (%)  Memory  Utilization (%)
Slot State      (C)  Total  Interrupt  DRAM (MB)  Heap    Buffer
0  Online        32   15      5         2048     22     13

```

1	Online	32	15	5	2048	22	13
2	Online	32	15	5	2048	22	13

show chassis fpc detail (MX104 Router)

```

user@host> show chassis fpc detail
Slot 0 information:
  State                Online
  Temperature           32 (C)
  Total CPU DRAM        2048 MB
  Total SRAM            403 MB
  Total SDRAM           1316 MB
  Start time            2013-05-23 14:39:18 IST
  Uptime                1 hour, 20 minutes, 22 seconds
Slot 1 information:
  State                Online
  Temperature           32 (C)
  Total CPU DRAM        2048 MB
  Total SRAM            403 MB
  Total SDRAM           1316 MB
  Start time            2013-05-23 14:39:18 IST
  Uptime                1 hour, 20 minutes, 22 seconds
Slot 2 information:
  State                Online
  Temperature           32 (C)
  Total CPU DRAM        2048 MB
  Total SRAM            403 MB
  Total SDRAM           1316 MB
  Start time            2013-05-23 14:39:18 IST
  Uptime                1 hour, 20 minutes, 22 seconds

```

show chassis fpc pic-status (MX104 Router)

```

user@host> show chassis fpc pic-status
Slot 0  Online
Slot 1  Online
  PIC 0  Online    10x 1GE(LAN) -E  SFP
  PIC 1  Online    10x 1GE(LAN) -E  SFP
Slot 2  Online
  PIC 0  Online    4x 10GE(LAN) SFP+

```

show chassis fpc (MX104-40G Router)

```

user@host> show chassis fpc

```

Slot	Temp	State	CPU Utilization (%)	CPU Utilization (%)	Memory	Utilization (%)		
			(C)	Interrupt	1min	5min	15min	DRAM (MB)
0	Online	48	18	6				2048
9		13						
1	Online	48	18	6				2048
9		13						
2	Online	48	18	6				2048
9		13						

show chassis fpc detail (MX104-40G Router)

```

user@host> show chassis fpc detail
Slot 0 information:
  State                Online
  Temperature           48 (C)
  Total CPU DRAM        2048 MB
  Total SRAM            403 MB

```

```

Total SDRAM                1316 MB
Start time                  2015-02-27 03:05:54 PST
Uptime                      2 hours, 38 minutes, 27 seconds
Slot 1 information:
State                       Online
Temperature                  48 (C)
Total CPU DRAM              2048 MB
Total SRAM                  403 MB
Total SDRAM                 1316 MB
Start time                  2015-02-27 03:05:55 PST
Uptime                      2 hours, 38 minutes, 26 seconds
Slot 2 information:
State                       Online
Temperature                  48 (C)
Total CPU DRAM              2048 MB
Total SRAM                  403 MB
Total SDRAM                 1316 MB
Start time                  2015-02-27 03:05:55 PST
Uptime                      2 hours, 38 minutes, 26 seconds

```

show chassis fpc pic-status (MX104-40G Router)

MIC slots 1/0 and 1/1 have been disabled by default on the MX104-40G routers. If you install MICs on those slots, the MIC details are displayed when you run the **show chassis fpc pic-status** command. However, the status of the MIC is displayed as offline.

```

user@host> show chassis fpc pic-status
Slot 0  Online
  PIC 0  Online      MS-MIC-16G
  PIC 2  Online      10x 1GE(LAN) SFP
  PIC 3  Online      10x 1GE(LAN) SFP
Slot 1  Online
  PIC 0  Offline     10x 1GE(LAN) SFP
  PIC 1  Offline     10x 1GE(LAN) SFP
Slot 2  Online
  PIC 0  Online      4x 10GE(LAN) SFP+

```

show chassis fpc pic-status (MX104-40G Router with Upgrade License)

When you install the upgrade license on MX104-40G, MIC slots 1/0 and 1/1 are enabled. If you install MICs on those slots, the MIC details are displayed and the status of the MIC is displayed as online when you run the **show chassis fpc pic-status** command.

```

user@host> show chassis fpc pic-status
Slot 0  Online
  PIC 0  Online      MS-MIC-16G
  PIC 2  Online      10x 1GE(LAN) SFP
  PIC 3  Online      10x 1GE(LAN) SFP
Slot 1  Online
  PIC 0  Online      10x 1GE(LAN) SFP
  PIC 1  Online      10x 1GE(LAN) SFP
Slot 2  Online
  PIC 0  Online      4x 10GE(LAN) SFP+

```

show chassis fpc (MX240 Router)

```

user@host> show chassis fpc

```

Slot	State	Temp	CPU Utilization (%)	Memory	Utilization (%)
		(C)	Total	DRAM (MB)	Heap
0	Empty		Interrupt		Buffer

1	Online	34	6	0	1024	18	30
2	Online	33	9	0	1024	24	30

show chassis fpc (MX480 Router)

```
user@host> show chassis fpc
```

Slot	State	Temp (C)	CPU Utilization (%)		Memory DRAM (MB)	Utilization (%)	
			Total	Interrupt		Heap	Buffer
0	Empty						
1	Online	36	9	0	1024	17	57
2	Empty						
3	Empty						
4	Empty						
5	Empty						

show chassis fpc detail (EX9200 Switch)

```
user@switch> show chassis fpc detail
```

Slot 2 information:

State	Online
Temperature	37
Total CPU DRAM	2048 MB
Total RLD RAM	331 MB
Total DDR DRAM	1536 MB
Start time:	2014-03-12 15:35:28 UTC
Uptime:	1 hour, 4 minutes, 29 seconds
Max Power Consumption	239 Watts

Slot 3 information:

State	Online
Temperature	39
Total CPU DRAM	2048 MB
Total RLD RAM	1036 MB
Total DDR DRAM	6656 MB
Start time:	2014-03-12 15:00:18 UTC
Uptime:	1 hour, 39 minutes, 39 seconds
Max Power Consumption	520 Watts

show chassis fpc (MX480 Router)

```
user@host> show chassis fpc
```

Slot	State	Temp (C)	CPU Utilization (%)		CPU Utilization (%)			Memory DRAM (MB)
			Total	Interrupt	1min	5min	15min	
0	Online		1	0	1	2	3	1024
4		56						
1	Online		1	0	2	2	3	1024
4		56						

show chassis fpc (MX480 Router with 100-Gigabit Ethernet CFP)

```
user@host> show chassis fpc
```

Slot	State	Temp (C)	CPU Utilization (%)		Memory DRAM (MB)	Utilization (%)	
			Total	Interrupt		Heap	Buffer
0	Online	33	4	0	2048	10	13
1	Online	36	7	0	2048	16	13
2	Online	29	6	0	1024	27	29
3	Online	33	0	0	0	0	0
4	Online	36	7	0	2048	19	13
5	Online	34	31	11	2048	14	13

show chassis fpc pic-status (MX480 Router with 100-Gigabit Ethernet CFP)

```

user@host> show chassis fpc pic-status
Slot 1  Online      MPC Type 3
PIC 2   Online      1X100GE CFP
Slot 2  Online      DPCE 40x 1GE R EQ
PIC 0   Online      10x 1GE(LAN) EQ
PIC 1   Online      10x 1GE(LAN) EQ
PIC 2   Online      10x 1GE(LAN) EQ
PIC 3   Online      10x 1GE(LAN) EQ
Slot 3  Online      MPC Type 3
PIC 0   Online      1X100GE CFP
PIC 2   Online      1X100GE CFP
Slot 4  Online      MPC Type 3
PIC 0   Online      1X100GE CFP
PIC 2   Online      1X100GE CFP
Slot 5  Online      MPC Type 2 3D EQ
PIC 0   Online      2x 10GE XFP
PIC 1   Online      2x 10GE XFP
PIC 2   Online      10x 1GE(LAN) SFP
PIC 3   Online      10x 1GE(LAN) SFP

```

show chassis fpc pic-status (EX Series Switch)

```

user@host> show chassis fpc pic-status
Slot 1  Online      EX9200 32x10G SFP
PIC 0   Online      8X10GE SFPP
PIC 1   Online      8X10GE SFPP
PIC 2   Online      8X10GE SFPP
PIC 3   Online      8X10GE SFPP
Slot 2  Online      EX9200 32x10G SFP
PIC 0   Online      8X10GE SFPP
PIC 1   Online      8X10GE SFPP
PIC 2   Online      8X10GE SFPP
PIC 3   Online      8X10GE SFPP

```

show chassis fpc (MX480 Router with MPC4E)

```

user@host> show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%)	Memory Interrupt	Utilization (%)	DRAM (MB)	Heap	Buffer
0	Empty							
1	Empty							
2	Online	38	7	0	2048	19	14	
3	Online	39	8	0	2048	18	14	
4	Online	39	7	0	2048	17	14	
5	Empty							

show chassis fpc detail (MX480 Router with MPC4E)

```

user@host> show chassis fpc detail
Slot 2 information:
State Online
Temperature 38
Total CPU DRAM 2048 MB
Total RLDRAM 1036 MB
Total DDR DRAM 11264 MB
Start time: 2013-02-18 05:06:57 PST
Uptime: 17 hours, 41 minutes, 9 seconds
Max Power Consumption 610 Watts
Slot 3 information:

```

```

State                               Online
Temperature                         38
Total CPU DRAM                     2048 MB
Total RLD RAM                      1036 MB
Total DDR DRAM                     11264 MB
Start time:                        2013-02-18 05:07:00 PST
Uptime:                            17 hours, 41 minutes, 6 seconds
Max Power Consumption              610 Watts
Slot 4 information:
State                               Diagnostics
Temperature                         37
Total CPU DRAM                     0 MB
Total RLD RAM                      0 MB
Total DDR DRAM                     0 MB
Max Power Consumption              520 Watts

```

show chassis fpc (MX480 Router with MPC4E)

```

user@host> show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%) Total	Interrupt	Memory DRAM (MB)	Utilization (%) Heap	Buffer
0	Empty						
1	Empty						
2	Online	38	7	0	2048	19	14
3	Online	39	8	0	2048	18	14
4	Online	39	7	0	2048	17	14
5	Empty						

show chassis fpc detail (MX480 Router with MPC4E)

```

user@host> show chassis fpc detail
Slot 2 information:
State                               Online
Temperature                         38
Total CPU DRAM                     2048 MB
Total RLD RAM                      1036 MB
Total DDR DRAM                     11264 MB
Start time:                        2013-02-18 05:06:57 PST
Uptime:                            17 hours, 41 minutes, 9 seconds
Max Power Consumption              610 Watts
Slot 3 information:
State                               Online
Temperature                         38
Total CPU DRAM                     2048 MB
Total RLD RAM                      1036 MB
Total DDR DRAM                     11264 MB
Start time:                        2013-02-18 05:07:00 PST
Uptime:                            17 hours, 41 minutes, 6 seconds
Max Power Consumption              610 Watts
Slot 4 information:
State                               Diagnostics
Temperature                         37
Total CPU DRAM                     0 MB
Total RLD RAM                      0 MB
Total DDR DRAM                     0 MB
Max Power Consumption              520 Watts

```

show chassis fpc (MX960 Router)

```

user@host> show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%) Total	Interrupt	Memory DRAM (MB)	Utilization (%) Heap	Buffer
------	-------	-------------	------------------------------	-----------	---------------------	-------------------------	--------

```

0 Empty
1 Empty
2 Empty
3 Online      25    19      0    1024    15      57
4 Empty
5 Online      26    27      0    1024    15      57
6 Empty
7 Empty
8 Empty
9 Empty
10 Empty
11 Empty

```

show chassis fpc (MX960 Router with MPC5EQ)

```

user@host> show chassis fpc

```

Slot	State	Temp (C)	CPU Total	Utilization (%) Interrupt	Memory DRAM (MB)	Utilization (%) Heap	Buffer
0	Online	38	16	0	3584	7	13
1	Online	31	15	0	2048	17	13
2	Empty						
3	Online	31	14	0	2048	20	13
4	Online	34	16	0	3584	7	13
5	Online	34	16	0	3584	7	13
6	Empty						
7	Online	32	9	0	2048	18	14
8	Online	36	19	0	3584	7	13
9	Online	31	9	0	2048	13	13
10	Online	35	14	0	3584	7	13
11	Online	33	11	0	2048	18	14

show chassis fpc detail (MX960 Router with MPC5EQ)

```

user@host> show chassis fpc detail

```

Slot 0 information:

State	Online
Temperature	38
Total CPU DRAM	3584 MB
Total XR2	291 MB
Total DDR DRAM	24960 MB
Start time:	2014-04-22 10:01:46 PDT
Uptime:	1 hour, 23 minutes, 40 seconds
Max Power Consumption	607 Watts

Slot 1 information:

State	Online
Temperature	31
Total CPU DRAM	2048 MB
Total RLD RAM	1036 MB
Total DDR DRAM	6656 MB
Start time:	2014-04-22 10:01:50 PDT
Uptime:	1 hour, 23 minutes, 36 seconds
Max Power Consumption	520 Watts

Slot 3 information:

State	Online
Temperature	31
Total CPU DRAM	2048 MB
Total RLD RAM	1324 MB
Total DDR DRAM	5120 MB
Start time:	2014-04-22 10:01:50 PDT
Uptime:	1 hour, 23 minutes, 36 seconds
Max Power Consumption	440 Watts


```

Slot 4 information:
  State                Online
  Temperature          34
  Total CPU DRAM       3584 MB
  Total XR2            291 MB
  Total DDR DRAM       24960 MB
  Start time:          2014-04-22 10:01:54 PDT
  Uptime:              1 hour, 23 minutes, 32 seconds
  Max Power Consumption 607 Watts
Slot 5 information:
  State                Online
  Temperature          34
  Total CPU DRAM       3584 MB
  Total XR2            291 MB
  Total DDR DRAM       24960 MB
  Start time:          2014-04-22 10:01:56 PDT
  Uptime:              1 hour, 23 minutes, 30 seconds
  Max Power Consumption 607 Watts
Slot 7 information:
  State                Online
  Temperature          32
  Total CPU DRAM       2048 MB
  Total RLDRAM         1036 MB
  Total DDR DRAM       11264 MB
  Start time:          2014-04-22 10:02:02 PDT
  Uptime:              1 hour, 23 minutes, 24 seconds
  Max Power Consumption 608 Watts
Slot 8 information:
  State                Online
  Temperature          36
  Total CPU DRAM       3584 MB
  Total XR2            291 MB
  Total DDR DRAM       24960 MB
  Start time:          2014-04-22 10:02:07 PDT
  Uptime:              1 hour, 23 minutes, 19 seconds
  Max Power Consumption 607 Watts
Slot 9 information:
  State                Online
  Temperature          31
  Total CPU DRAM       2048 MB
  Total RLDRAM         734 MB
  Total DDR DRAM       3108 MB
  Start time:          2014-04-22 10:02:05 PDT
  Uptime:              1 hour, 23 minutes, 21 seconds
  Max Power Consumption 368 Watts
Slot 10 information:
  State                Online
  Temperature          35
  Total CPU DRAM       3584 MB
  Total XR2            291 MB
  Total DDR DRAM       24960 MB
  Start time:          2014-04-22 10:02:11 PDT
  Uptime:              1 hour, 23 minutes, 15 seconds
  Max Power Consumption 607 Watts
Slot 11 information:
  State                Online
  Temperature          33
  Total CPU DRAM       2048 MB
  Total RLDRAM         1036 MB
  Total DDR DRAM       11264 MB
  Start time:          2014-04-22 10:02:16 PDT

```

Uptime:	1 hour, 23 minutes, 10 seconds
Max Power Consumption	608 Watts

show chassis fpc pic-status(MX960 Router with MPC5EQ)

```

user@host> show chassis fpc pic-status
Slot 0  Online      MPC5E 3D Q 2CGE+4XGE
        PIC 0  Online      2X10GE SFPP OTN
        PIC 1  Online      1X100GE CFP2 OTN
        PIC 2  Online      2X10GE SFPP OTN
        PIC 3  Online      1X100GE CFP2 OTN
Slot 1  Online      MPCE Type 3 3D
        PIC 0  Online      10X10GE SFPP
        PIC 2  Online      1X100GE CXP
Slot 3  Online      MPC 3D 16x 10GE
        PIC 0  Online      4x 10GE(LAN) SFP+
        PIC 1  Online      4x 10GE(LAN) SFP+
        PIC 2  Online      4x 10GE(LAN) SFP+
        PIC 3  Online      4x 10GE(LAN) SFP+
Slot 4  Online      MPC5E 3D Q 2CGE+4XGE
        PIC 0  Online      2X10GE SFPP OTN
        PIC 1  Online      1X100GE CFP2 OTN
        PIC 2  Online      2X10GE SFPP OTN
        PIC 3  Online      1X100GE CFP2 OTN
Slot 5  Online      MPC5E 3D Q 2CGE+4XGE
        PIC 0  Online      2X10GE SFPP OTN
        PIC 1  Online      1X100GE CFP2 OTN
        PIC 2  Online      2X10GE SFPP OTN
        PIC 3  Online      1X100GE CFP2 OTN
Slot 7  Online      MPC4E 3D 2CGE+8XGE
        PIC 0  Online      4x10GE SFPP
        PIC 1  Online      1X100GE CFP
        PIC 2  Online      4x10GE SFPP
        PIC 3  Online      1X100GE CFP
Slot 8  Online      MPC5E 3D Q 24XGE+6XLGE
        PIC 0  Offline     12X10GE SFPP OTN
        PIC 1  Offline     12X10GE SFPP OTN
        PIC 2  Online      3X40GE QSFPP
        PIC 3  Online      3X40GE QSFPP
Slot 9  Online      MPCE Type 2 3D P
        PIC 0  Online      2x 10GE XFP
        PIC 1  Online      2x 10GE XFP
Slot 10 Online      MPC5E 3D Q 24XGE+6XLGE
        PIC 0  Online      12X10GE SFPP
        PIC 1  Online      12X10GE SFPP
        PIC 2  Offline     3X40GE QSFPP
        PIC 3  Offline     3X40GE QSFPP
Slot 11 Online      MPC4E 3D 2CGE+8XGE
        PIC 0  Online      4x10GE SFPP
        PIC 1  Online      1X100GE CFP
        PIC 2  Online      4x10GE SFPP
        PIC 3  Online      1X100GE CFP

```

show chassis fpc (MX240, MX480, MX960 Routers with Application Services Modular Line Card)

```

user@host> show chassis fpc 1

```

Temp	CPU Utilization (%)	Memory	Utilization (%)			
Slot State	(C)	Total	Interrupt	DRAM (MB)	Heap	Buffer
1 Online	34	5	0	3072	5	13

show chassis fpc (MX240, MX480, MX960 with Application Services Modular Line Card)

```

user@host> show chassis fpc 1 detail
Slot 1 information:
  State                               Online
  Temperature                         34
  Total CPU DRAM                      3072 MB
  Total RLDRAM                       259 MB
  Total DDR DRAM                     4864 MB
  Start time:                        2012-06-19 10:51:43 PDT
  Uptime:                            16 minutes, 48 seconds
  Max Power Consumption               550 Watts

```

show chassis fpc (MX240, MX480, MX960, MX2010, and MX2020 3D Universal Edge Routers Configured for Dynamic Power Management)

```

user@host> show chassis fpc 2 detail

Slot 2 information:
  State                               Online
  Temperature                         37
  Total CPU DRAM                      3584 MB
  Total XR2                          275 MB
  Total DDR DRAM                     20352 MB
  Start time:                        2014-07-18 02:51:23 PDT
  Uptime:                            5 minutes, 19 seconds
  Max MPC Base Power Consumption      485 Watts
  Max MICO Power Consumption          50 Watts
  Max MIC1 Power Consumption          50 Watts
  Max MPC Total Power Consumption     585 Watts

```

show chassis fpc (MX2010 Routers)

```

user@host> show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%)	Memory Interrupt	Utilization (%)	DRAM (MB)	Heap	Buffer
0	Online	34	9	0	2048	18	13	
1	Online	32	9	0	2048	15	13	
2	Empty							
3	Empty							
4	Empty							
5	Empty							
6	Empty							
7	Empty							
8	Online	31	13	0	2048	11	13	
9	Online	33	10	0	2048	18	13	

show chassis fpc (MX2020 Routers)

```

user@host> show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%)	Memory Interrupt	Utilization (%)	DRAM (MB)	Heap	Buffer
0	Online	10	12	0	2048	18	13	
1	Online	8	9	0	2048	18	13	
2	Online	7	9	0	2048	18	13	
3	Online	8	10	0	2048	18	13	
4	Online	9	10	0	2048	18	13	
5	Online	8	9	0	2048	18	13	
6	Online	8	10	0	2048	18	13	
7	Online	9	9	0	2048	18	13	
8	Online	9	10	0	2048	18	13	

9	Online	10	9	0	2048	18	13
10	Online	16	8	0	2048	18	13
11	Online	11	10	0	2048	18	13
12	Online	10	10	0	2048	18	13
13	Online	11	9	0	2048	18	13
14	Online	12	10	0	2048	18	13
15	Online	13	9	0	2048	18	13
16	Online	13	9	0	2048	18	13
17	Online	12	9	0	2048	18	13
18	Online	12	8	0	2048	18	13
19	Online	14	10	0	2048	18	13

show chassis fpc (MX2020 Router with MPC4E)

```
user@host> show chassis fpc
```

Slot	Temp	State	CPU Utilization (%)	Memory	Utilization (%)	Heap	Buffer
			(C) Total	Interrupt	DRAM (MB)		
0		Online	33 12	2	2048	11	13
1		Empty					
2		Empty					
3		Empty					
4		Empty					
5		Empty					
6		Empty					
7		Empty					
8		Empty					
9		Online	31 10	0	2048	11	13
10		Online	32 7	0	2048	14	13
11		Empty					
12		Empty					
13		Empty					
14		Online	28 12	0	2048	15	14
15		Empty					
16		Empty					
17		Empty					
18		Empty					
19		Online	38 8	0	2048	18	13

show chassis fpc detail (MX2020 Router with MPC4E)

```
user@host> show chassis fpc detail
```

Slot 0 information:	
State	Online
Temperature	34
Total CPU DRAM	2048 MB
Total RLDRAM	806 MB
Total DDR DRAM	2632 MB
Start time:	2013-02-17 08:17:35 PST
Uptime:	1 day, 14 hours, 50 minutes, 39 seconds
Max Power Consumption	368 Watts
Slot 9 information:	
State	Online
Temperature	32
Total CPU DRAM	2048 MB
Total RLDRAM	806 MB
Total DDR DRAM	2632 MB
Start time:	2013-02-17 08:17:43 PST
Uptime:	1 day, 14 hours, 50 minutes, 31 seconds
Max Power Consumption	368 Watts
Slot 10 information:	
State	Online

```

Temperature                37
Total CPU DRAM              2048 MB
Total RLD RAM              1036 MB
Total DDR DRAM             6656 MB
Start time:                2013-02-17 08:17:54 PST
Uptime:                    1 day, 14 hours, 50 minutes, 20 seconds
Max Power Consumption      520 Watts
Slot 14 information:
State                      Online
Temperature                32
Total CPU DRAM              2048 MB
Total RLD RAM              1036 MB
Total DDR DRAM             11264 MB
Start time:                2013-02-17 08:18:01 PST
Uptime:                    1 day, 14 hours, 50 minutes, 13 seconds
Max Power Consumption      610 Watts
Slot 19 information:
State                      Online
Temperature                38
Total CPU DRAM              2048 MB
Total RLD RAM              1324 MB
Total DDR DRAM             5120 MB
Start time:                2013-02-17 08:18:08 PST
Uptime:                    1 day, 14 hours, 50 minutes, 6 seconds
Max Power Consumption      440 Watts

```

show chassis fpc (MX2020 Router with MPC5EQ and MPC6E)

```

user@host> show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%)	Memory Utilization (%)	DRAM (MB)	Heap	Buffer
0	Online	31	20	0	3584	7	13
1	Online	28	19	0	2048	17	13
2	Online	27	10	0	2048	18	14
3	Online	26	10	0	2048	13	13
4	Online	29	19	0	3584	7	13
5	Online	28	68	0	2048	20	13
6	Empty						
7	Empty						
8	Empty						
9	Online	36	19	0	3584	10	13
10	Online	37	26	0	3584	10	13
11	Empty						
12	Empty						
13	Empty						
14	Empty						
15	Empty						
16	Empty						
17	Online	28	43	0	3584	10	13
18	Online	29	19	0	3584	7	13
19	Online	31	19	0	3584	7	13

show chassis fpc detail (MX2020 Router with MPCEQ and MPC6E)

```

user@host> show chassis fpc detail
Slot 0 information:
State                      Online
Temperature                31
Total CPU DRAM              3584 MB
Total XR2                   291 MB

```

```

Total DDR DRAM                24960 MB
Start time:                    2014-04-22 23:33:19 PDT
Uptime:                        6 minutes, 24 seconds
Max Power Consumption          607 Watts
Slot 1 information:
State                          Online
Temperature                    28
Total CPU DRAM                 2048 MB
Total RLD RAM                  1036 MB
Total DDR DRAM                 6656 MB
Start time:                    2014-04-22 23:33:24 PDT
Uptime:                        6 minutes, 19 seconds
Max Power Consumption          520 Watts
Slot 2 information:
State                          Online
Temperature                    27
Total CPU DRAM                 2048 MB
Total RLD RAM                  1036 MB
Total DDR DRAM                 11264 MB
Start time:                    2014-04-22 23:33:34 PDT
Uptime:                        6 minutes, 9 seconds
Max Power Consumption          608 Watts
Slot 3 information:
State                          Online
Temperature                    26
Total CPU DRAM                 2048 MB
Total RLD RAM                  734 MB
Total DDR DRAM                 3108 MB
Start time:                    2014-04-22 23:33:39 PDT
Uptime:                        6 minutes, 4 seconds
Max Power Consumption          368 Watts
Slot 4 information:
State                          Online
Temperature                    29
Total CPU DRAM                 3584 MB
Total XR2                      291 MB
Total DDR DRAM                 24960 MB
Start time:                    2014-04-22 23:33:51 PDT
Uptime:                        5 minutes, 52 seconds
Max Power Consumption          607 Watts
Slot 5 information:
State                          Online
Temperature                    28
Total CPU DRAM                 2048 MB
Total RLD RAM                  1324 MB
Total DDR DRAM                 5120 MB
Start time:                    2014-04-22 23:33:57 PDT
Uptime:                        5 minutes, 46 seconds
Max Power Consumption          440 Watts
Slot 9 information:
State                          Online
Temperature                    25
Total CPU DRAM                 3584 MB
Total XR2                      518 MB
Total DDR DRAM                 49920 MB
Start time:                    2014-04-22 23:31:20 PDT
Uptime:                        8 minutes, 23 seconds
Max Power Consumption          1130 Watts
Slot 10 information:
State                          Online
Temperature                    32

```

```

Total CPU DRAM          3584 MB
Total XR2                518 MB
Total DDR DRAM          49920 MB
Start time:              2014-04-22 23:31:25 PDT
Uptime:                  8 minutes, 18 seconds
Max Power Consumption    1130 Watts
Slot 17 information:
State                    Online
Temperature              25
Total CPU DRAM          3584 MB
Total XR2                518 MB
Total DDR DRAM          49920 MB
Start time:              2014-04-22 23:31:29 PDT
Uptime:                  8 minutes, 14 seconds
Max Power Consumption    1130 Watts
Slot 18 information:
State                    Online
Temperature              29
Total CPU DRAM          3584 MB
Total XR2                291 MB
Total DDR DRAM          24960 MB
Start time:              2014-04-22 23:34:11 PDT
Uptime:                  5 minutes, 32 seconds
Max Power Consumption    607 Watts
Slot 19 information:
State                    Online
Temperature              32
Total CPU DRAM          3584 MB
Total XR2                291 MB
Total DDR DRAM          24960 MB
Start time:              2014-04-22 23:34:20 PDT
Uptime:                  5 minutes, 23 seconds
Max Power Consumption    607 Watts

```

show chassis fpc pic-status (MX2020 Router with MPC5EQ and MPC6E)

```

user@host> show chassis fpc pic-status
Slot 0  Online      MPC5E 3D Q 24XGE+6XLGE
PIC 0   Online      12X10GE SFPP OTN
PIC 1   Online      12X10GE SFPP OTN
PIC 2   Offline     3X40GE QSFPP
PIC 3   Offline     3X40GE QSFPP
Slot 1  Online      MPCE Type 3 3D
PIC 0   Online      10X10GE SFPP
PIC 2   Online      1X100GE CXP
Slot 2  Online      MPC4E 3D 2CGE+8XGE
PIC 0   Online      4x10GE SFPP
PIC 1   Online      1X100GE CFP
PIC 2   Online      4x10GE SFPP
PIC 3   Online      1X100GE CFP
Slot 3  Online      MPCE Type 2 3D P
PIC 0   Online      2x 10GE XFP
PIC 1   Online      2x 10GE XFP
Slot 4  Online      MPC5E 3D Q 2CGE+4XGE
PIC 0   Online      2X10GE SFPP OTN
PIC 1   Online      1X100GE CFP2 OTN
PIC 2   Online      2X10GE SFPP OTN
PIC 3   Online      1X100GE CFP2 OTN
Slot 5  Online      MPC 3D 16x 10GE
PIC 0   Online      4x 10GE(LAN) SFP+
PIC 1   Online      4x 10GE(LAN) SFP+

```

```

PIC 2 Online 4x 10GE(LAN) SFP+
PIC 3 Online 4x 10GE(LAN) SFP+
Slot 9 Online MPC6E 3D
PIC 0 Online 2X100GE CFP2 OTN
PIC 1 Online 2X100GE CFP2 OTN
Slot 10 Online MPC6E 3D
PIC 0 Online 24X10GE SFPP OTN
PIC 1 Online 4X100GE CXP
Slot 17 Online MPC6E 3D
PIC 0 Online 24X10GE SFPP
PIC 1 Online 4X100GE CXP
Slot 18 Online MPC5E 3D Q 24XGE+6XLGE
PIC 0 Offline 12X10GE SFPP OTN
PIC 1 Offline 12X10GE SFPP OTN
PIC 2 Online 3X40GE QSFPP
PIC 3 Online 3X40GE QSFPP
Slot 19 Online MPC5E 3D Q 24XGE+6XLGE
PIC 0 Online 12X10GE SFPP OTN
PIC 1 Offline 12X10GE SFPP OTN
PIC 2 Offline 3X40GE QSFPP
PIC 3 Online 3X40GE QSFPP

```

show chassis fpc detail (MX Series Routers)

```

user@host> show chassis fpc detail 2
Slot 0 information:
State Online
Temperature 36 degrees C / 96 degrees F
Total CPU DRAM 1024 MB
Total RLDRAM 256 MB
Total DDR DRAM 4096 MB
Start time: 2009-08-11 21:20:30 PDT
Uptime: 2 hours, 8 minutes, 50 seconds
Max Power Consumption 335 Watts

```

show chassis fpc detail (EX Series Switches)

```

user@host> show chassis fpc detail 2
Slot 1 information:
State Online
Temperature 41
Total CPU DRAM 2048 MB
Total RLDRAM 1036 MB
Total DDR DRAM 11264 MB
Start time: 2013-04-02 00:04:52 PDT
Uptime: 7 days, 9 hours, 47 minutes, 46 seconds
Max Power Consumption 610 Watts
Slot 2 information:
State Online
Temperature 41
Total CPU DRAM 2048 MB
Total RLDRAM 1036 MB
Total DDR DRAM 11264 MB
Start time: 2013-04-02 00:04:56 PDT
Uptime: 7 days, 9 hours, 47 minutes, 42 seconds
Max Power Consumption 610 Watts

```

show chassis fpc (Hardware Not Supported)

```

user@host> show chassis fpc
show chassis fpc
Temp CPU Utilization (%) Memory Utilization (%)

```


Slot	State	(C)	Total	Interrupt	DRAM (MB)	Heap	Buffer
0	Online	-----		CPU less FPC	-----		
1	Present	-----	Hardware	Not In Right Slot	-----		
2	Online		0	0	0	0	0
3	Present	-----	Hardware	Not Supported	-----		
4	Empty						
5	Empty						
6	Online		0	0	0	0	0

show chassis fpc detail (Hardware Not Supported)

```

user@host> show chassis fpc detail
Slot 0 information:
  State                Online
  Total CPU DRAM       ---- CPU less FPC ----
  Start time           2006-07-07 03:21:00 UTC
  Uptime                27 minutes, 51 seconds
Slot 1 information:
  State                Present
  Reason               --- Hardware Not In Right Slot ---
Slot 2 information:
  State                Online
  Total CPU DRAM       32 MB
  Start time           2006-07-07 03:20:59 UTC
  Uptime                27 minutes, 52 seconds
Slot 3 information:
  State                Present
  Reason               --- Hardware Not Supported ---
  Total CPU DRAM       0 MB
Slot 6 information:
  State                Online
  Total CPU DRAM       32 MB
  Start time           2006-07-07 03:21:01 UTC
  Uptime                27 minutes, 50 seconds

```

show chassis fpc pic-status

```

user@host> show chassis fpc pic-status
Slot 0 Online
  PIC 1    1x OC-12 ATM, MM
  PIC 2    1x OC-12 ATM, MM
  PIC 3    1x OC-12 ATM, MM
Slot 1 Online
  PIC 0    1x OC-48 SONET, SMIR
Slot 2 Online
  PIC 0    1x OC-192 SONET, SMSR

```

show chassis fpc pic-status (M Series Routers)

```

user@host> show chassis fpc pic-status
Slot 1 Online      FPC Type 1
  PIC 0 Present    2x OC-3 ATM, MM- Hardware Error
  PIC 1 Online     4x OC-3 SONET, SMIR
Slot 2 Online      E-FPC Type 2
  PIC 0 Online     4x G/E, 1000 BASE-SX
  PIC 1 Online     2x G/E SFP, 1000 BASE
  PIC 3 Online     1x Tunnel
Slot 3 Online      E-FPC Type 1
  PIC 0 Online     1x G/E IQ, 1000 BASE
  PIC 2 Online     1x G/E SFP, 1000 BASE
Slot 4 Online      E-FPC Type 2

```

```

PIC 0 Online      4x G/E SFP, 1000 BASE
PIC 1 Online      4x G/E SFP, 1000 BASE
PIC 2 Online      4x G/E SFP, 1000 BASE
PIC 3 Online      4x G/E SFP, 1000 BASE
Slot 5 Online     FPC Type 2
...

```

show chassis fpc pic-status (M120 Router)

```

user@host> show chassis fpc pic-status
Slot 1 Online      M120 CFPC 10GE
PIC 0 Online       1x 10GE(LAN/WAN) XFP
Slot 3 Online      M120 FPC Type 2 (proto)
PIC 0 Online       2x G/E IQ, 1000 BASE
PIC 1 Online       4x OC-3 SONET, SMIR
PIC 2 Online       2x G/E IQ, 1000 BASE
PIC 3 Online       8x 1GE(LAN), IQ2
Slot 4 Online      M120 FPC Type 3 (proto)
PIC 0 Online       10x 1GE(LAN), 1000 BASE
Slot 5 Online      M120 FPC Type 1 (proto)
PIC 0 Present      1x G/E, 1000 BASE-LX- Not Supported
PIC 1 Online       1x CHOC3 IQ SONET, SMLR
PIC 2 Online       4x CHDS3 IQ
PIC 3 Online       1x G/E SFP, 1000 BASE

```

show chassis fpc pic-status (MX240, MX480, and MX960 Routers with Application Services Modular Line Card)

In the following output **Slot 1 and Slot 5** are the Application Services Modular Carrier Cards (AS MCC), **PIC 0** is the Application Services Modular Storage Card (AS MSC), and **PIC 2** is the Application Services Modular Processing Card (AS MXC).

```

user@host> show chassis fpc pic-status
Slot 2 Online      MPC Type 1 3D Q
Slot 1 Online      AS-MCC
PIC 0 Online       AS-MSC
PIC 2 Online       AS-MXC
Slot 4 Offline     MPC 3D 16x 10GE
Slot 5 Offline     AS-MCC

```

show chassis fpc lcc (TX Matrix Router)

```

user@host> show chassis fpc lcc 0
lcc0-re0:
-----

```

Slot	State	Temp (C)	CPU Total	Utilization (%) Interrupt	Memory Utilization (%) DRAM (MB)	Heap	Buffer
0	Empty						
1	Online	27	2	0	256	8	44
2	Online	27	3	0	256	15	44
3	Empty						
4	Empty						
5	Empty						
6	Empty						
7	Empty						

show chassis fpc pic-status (TX Matrix Router)

```

user@host> show chassis fpc pic-status
lcc0-re0:
-----
Slot 0 Online      FPC Type 3

```

```

PIC 0 Online      1x OC-192 SM SR1
PIC 1 Online      1x OC-192 SM SR2
PIC 2 Online      1x OC-192 SM SR1
PIC 3 Online      1x Tunnel
Slot 1 Online     FPC Type 2
PIC 0 Online      1x OC-48 SONET, SMSR
PIC 1 Online      1x OC-48 SONET, SMSR

```

lcc1-re0:

lcc2-re0:

```

Slot 1 Online     FPC Type 3
PIC 0 Online      1x OC-192 SM SR1
Slot 5 Online     FPC Type 2
PIC 0 Online      1x OC-48 SONET, SMSR
PIC 1 Online      2x G/E, 1000 BASE-LX
PIC 2 Online      2x G/E, 1000 BASE-LX
PIC 3 Online      1x OC-48 SONET, SMSR

```

lcc3-re0:

show chassis fpc pic-status lcc (TX Matrix Router)

```
user@host> show chassis fpc pic-status lcc 0
```

lcc0-re0:

```

Slot 0 Online     FPC Type 3
PIC 0 Online      1x OC-192 SM SR2
Slot 1 Online     FPC Type 2
PIC 0 Online      2x OC-12 ATM2 IQ, MM
PIC 1 Online      1x OC-48 SONET, SMSR
PIC 2 Online      1x OC-48 SONET, SMSR
PIC 3 Online      4x G/E, 1000 BASE-SX

```

show chassis fpc (TX Matrix Plus Router)

```
user@host> show chassis fpc
```

lcc0-re0:

Slot	State	Temp (C)	CPU Utilization (%)	Memory	Utilization (%)
			Total Interrupt	DRAM (MB) Heap	Buffer
0	Empty				
1	Online	38	4 0	2048 3	24
2	Online	43	8 0	2048 6	24
3	Empty				
4	Online	43	6 0	2048 6	24
5	Empty				
6	Online	42	13 0	2048 6	24
7	Online	45	7 0	2048 3	24

lcc2-re0:

Slot	State	Temp (C)	CPU Utilization (%)	Memory	Utilization (%)
			Total Interrupt	DRAM (MB) Heap	Buffer
0	Online	42	10 0	2048 6	24
1	Empty				
2	Online	42	11 0	2048 6	24
3	Online	40	5 0	2048 3	24

4	Online	33	26	0	1024	8	49
5	Empty						
6	Online	43	8	0	2048	6	24
7	Online	46	6	0	2048	3	24

lcc3-re0:

Slot	State	Temp (C)	CPU Utilization (%) Total Interrupt	Memory DRAM (MB)	Utilization (%) Heap Buffer
0	Empty				
1	Empty				
2	Online	39	30	0	2048 7 24
3	Empty				
4	Online	41	8	0	2048 6 24
5	Online	41	12	0	2048 6 24
6	Online	40	8	0	2048 6 24
7	Online	42	4	0	2048 3 24

show chassis fpc lcc (TX Matrix Plus Router)

user@host> show chassis fpc lcc 0

lcc0-re0:

Slot	State	Temp (C)	CPU Utilization (%) Total Interrupt	Memory DRAM (MB)	Utilization (%) Heap Buffer
0	Empty				
1	Online	38	4	0	2048 3 24
2	Online	43	8	0	2048 6 24
3	Empty				
4	Online	43	6	0	2048 6 24
5	Empty				
6	Online	42	14	0	2048 6 24
7	Online	45	6	0	2048 3 24

show chassis fpc detail (TX Matrix Plus Router)

user@host> show chassis fpc details

lcc0-re0:

Slot 1 information:

State	Online
Temperature	38 degrees C / 100 degrees F
Total CPU DRAM	2048 MB
Total SRAM	64 MB
Total SDRAM	1280 MB
Start time	2010-10-04 20:06:22 PDT
Uptime	1 hour, 32 minutes, 51 seconds

Slot 2 information:

State	Online
Temperature	43 degrees C / 109 degrees F
Total CPU DRAM	2048 MB
Total SRAM	128 MB
Total SDRAM	2560 MB
Start time	2010-10-04 20:06:37 PDT
Uptime	1 hour, 32 minutes, 36 seconds

Slot 4 information:

State	Online
Temperature	43 degrees C / 109 degrees F
Total CPU DRAM	2048 MB
Total SRAM	128 MB

```

Total SDRAM                2560 MB
Start time                 2010-10-04 20:06:40 PDT
Uptime                     1 hour, 32 minutes, 33 seconds
Slot 6 information:
State                      Online
Temperature                 42 degrees C / 107 degrees F
Total CPU DRAM             2048 MB
Total SRAM                  128 MB
Total SDRAM                2560 MB
Start time                 2010-10-04 20:06:42 PDT
Uptime                     1 hour, 32 minutes, 31 seconds
Slot 7 information:
State                      Online
Temperature                 45 degrees C / 113 degrees F
Total CPU DRAM             2048 MB
Total SRAM                  64 MB
Total SDRAM                1280 MB
Start time                 2010-10-04 20:06:43 PDT
Uptime                     1 hour, 32 minutes, 30 seconds

```

```
lcc2-re0:
```

```

-----
Slot 0 information:
State                      Online
Temperature                 42 degrees C / 107 degrees F
Total CPU DRAM             2048 MB
Total SRAM                  128 MB
Total SDRAM                2560 MB
Start time                 2010-10-04 20:06:35 PDT
Uptime                     1 hour, 32 minutes, 38 seconds
Slot 2 information:
State                      Online
Temperature                 42 degrees C / 107 degrees F
Total CPU DRAM             2048 MB
Total SRAM                  128 MB
Total SDRAM                2560 MB
Start time                 2010-10-04 20:06:37 PDT
Uptime                     1 hour, 32 minutes, 36 seconds
Slot 3 information:
State                      Online
Temperature                 40 degrees C / 104 degrees F
Total CPU DRAM             2048 MB
Total SRAM                  64 MB
Total SDRAM                1280 MB
Start time                 2010-10-04 20:06:28 PDT
Uptime                     1 hour, 32 minutes, 45 seconds
Slot 4 information:
State                      Online
Temperature                 33 degrees C / 91 degrees F
Total CPU DRAM             1024 MB
Total SRAM                  64 MB
Total SDRAM                1280 MB
Start time                 2010-10-04 20:08:03 PDT
Uptime                     1 hour, 31 minutes, 10 seconds
Slot 6 information:
State                      Online
Temperature                 43 degrees C / 109 degrees F
Total CPU DRAM             2048 MB
Total SRAM                  128 MB
Total SDRAM                2560 MB
Start time                 2010-10-04 20:06:44 PDT

```

```

Uptime                               1 hour, 32 minutes, 29 seconds
Slot 7 information:
State                                Online
Temperature                          46 degrees C / 114 degrees F
Total CPU DRAM                       2048 MB
Total SRAM                           64 MB
Total SDRAM                           1280 MB
Start time                           2010-10-04 20:06:46 PDT
Uptime                               1 hour, 32 minutes, 27 seconds

```

```
lcc3-re0:
```

```

-----
Slot 2 information:
State                                Online
Temperature                          38 degrees C / 100 degrees F
Total CPU DRAM                       2048 MB
Total SRAM                           128 MB
Total SDRAM                           2560 MB
Start time                           2010-10-04 20:17:31 PDT
Uptime                               1 hour, 21 minutes, 42 seconds

```

```

Slot 4 information:
State                                Online
Temperature                          41 degrees C / 105 degrees F
Total CPU DRAM                       2048 MB
Total SRAM                           128 MB
Total SDRAM                           2560 MB
Start time                           2010-10-04 20:17:34 PDT
Uptime                               1 hour, 21 minutes, 39 seconds

```

```

Slot 5 information:
State                                Online
Temperature                          41 degrees C / 105 degrees F
Total CPU DRAM                       2048 MB
Total SRAM                           128 MB
Total SDRAM                           2560 MB
Start time                           2010-10-04 20:17:36 PDT
Uptime                               1 hour, 21 minutes, 37 seconds

```

```

Slot 6 information:
State                                Online
Temperature                          40 degrees C / 104 degrees F
Total CPU DRAM                       2048 MB
Total SRAM                           128 MB
Total SDRAM                           2560 MB
Start time                           2010-10-04 20:17:39 PDT
Uptime                               1 hour, 21 minutes, 34 seconds

```

```

Slot 7 information:
State                                Online
Temperature                          42 degrees C / 107 degrees F
Total CPU DRAM                       2048 MB
Total SRAM                           64 MB
Total SDRAM                           1280 MB
Start time                           2010-10-04 20:17:41 PDT
Uptime                               1 hour, 21 minutes, 32 seconds

```

show chassis fpc pic-status (TX Matrix Plus Router)

```
user@host> show chassis fpc pic-status
```

```
lcc0-re0:
```

```

-----
Slot 1  Online      FPC Type 2-ES
PIC 0   Online      8x 1GE(LAN), IQ2

```

```

Slot 2  Online      FPC Type 4-ES
        PIC 0 Online  4x 10GE (LAN/WAN) XFP
Slot 4  Online      FPC Type 4-ES
        PIC 0 Online  4x 10GE (LAN/WAN) XFP
Slot 6  Online      FPC Type 4-ES
        PIC 0 Online  4x 10GE (LAN/WAN) XFP
        PIC 1 Online  4x 10GE (LAN/WAN) XFP
Slot 7  Online      FPC Type 3-ES
        PIC 0 Online  10x 1GE(LAN), 1000 BASE
        PIC 2 Online  1x OC-192 SM SR2
        PIC 3 Online  10x 1GE(LAN), 1000 BASE

```

1cc2-re0:

```

-----
Slot 0  Online      FPC Type 4-ES
        PIC 0 Online  4x 10GE (LAN/WAN) XFP
Slot 2  Online      FPC Type 4-ES
        PIC 0 Online  4x 10GE (LAN/WAN) XFP
        PIC 1 Online  4x 10GE (LAN/WAN) XFP
Slot 3  Online      FPC Type 2-ES
        PIC 0 Online  8x 1GE(LAN), IQ2
Slot 4  Online      FPC Type 4
        PIC 0 Online  10x10GE(LAN/WAN) SFPP
Slot 6  Online      FPC Type 4-ES
        PIC 0 Online  4x OC-192 SONET XFP
Slot 7  Online      FPC Type 3-ES
        PIC 0 Online  10x 1GE(LAN), 1000 BASE
        PIC 1 Offline 1x 10GE(LAN/WAN) IQ2E
        PIC 2 Online  1x OC-192 SM SR2
        PIC 3 Online  1x Tunnel

```

1cc3-re0:

```

-----
Slot 2  Online      FPC Type 4-ES
        PIC 0 Online  10x10GE(LAN/WAN) SFPP
Slot 4  Online      FPC Type 4-ES
        PIC 0 Online  4x OC-192 SONET XFP
Slot 5  Online      FPC Type 4-ES
        PIC 0 Online  4x OC-192 SONET XFP
        PIC 1 Online  4x 10GE (LAN/WAN) XFP
Slot 6  Online      FPC Type 4-ES
        PIC 1 Online  4x 10GE (LAN/WAN) XFP
Slot 7  Online      FPC Type 3-ES
        PIC 0 Online  10x 1GE(LAN), 1000 BASE
        PIC 1 Online  8x 1GE(TYPE3), IQ2E
        PIC 2 Online  4x OC-48 SONET

```

show chassis fpc (T1600 Router)

user@host> show chassis fpc

Slot	State	Temp (C)	CPU Utilization (%) Total Interrupt	Memory DRAM (MB)	Utilization (%) Heap Buffer
0	Empty				
1	Empty				
2	Online	49	3 0	2048	3 24
3	Online	46	6 0	2048	6 24
4	Empty				
5	Online	46	5 0	2048	3 24
6	Empty				
7	Online	44	8 0	1024	7 49

show chassis fpc detail (T1600 Router)

```

user@host> show chassis fpc detail

show chassis fpc detail
Slot 2 information:
  State      Online
  Temperature 49 degrees C / 120 degrees F
  Total CPU DRAM 2048 MB
  Total SRAM    64 MB
  Total SDRAM   1280 MB
  Start time   2010-10-04 21:12:52 PDT
  Uptime       32 minutes, 9 seconds
Slot 3 information:
  State      Online
  Temperature 47 degrees C / 116 degrees F
  Total CPU DRAM 2048 MB
  Total SRAM    128 MB
  Total SDRAM   2560 MB
  Start time   2010-10-04 21:13:06 PDT
  Uptime       31 minutes, 55 seconds
Slot 5 information:
  State      Online
  Temperature 46 degrees C / 114 degrees F
  Total CPU DRAM 2048 MB
  Total SRAM    64 MB
  Total SDRAM   1280 MB
  Start time   2010-10-04 21:12:56 PDT
  Uptime       32 minutes, 5 seconds
Slot 7 information:
  State      Online
  Temperature 44 degrees C / 111 degrees F
  Total CPU DRAM 1024 MB
  Total SRAM    64 MB
  Total SDRAM   1280 MB
  Start time   2010-10-04 21:14:34 PDT
  Uptime       30 minutes, 27 seconds

```

show chassis fpc <fpc-slot> (EX Series Switch)

```

user@host> show chassis fpc 2

Slot State      Temp  CPU Utilization (%)  Memory  Utilization (%)
      (C) Total  Interrupt  DRAM (MB) Heap  Buffer
2 Online        40    12         0    2048    19    14

```

show chassis fpc slot (T1600 Router)

```

user@host> show chassis fpc slot 2

Slot State      Temp  CPU Utilization (%)  Memory  Utilization (%)
      (C) Total  Interrupt  DRAM (MB) Heap  Buffer
2 Online        49     3         0    2048     3    24

```

show chassis fpc pic-status (T1600 Router)

```

user@host> show chassis fpc pic-status

Slot 2  Online      FPC Type 1-ES
PIC 0   Online      Load Type 1
PIC 1   Online      4x 1GE(LAN), IQ2E

```



```

PIC 3 Online 1x OC-12-3 SFP
Slot 3 Online FPC Type 4-ES
PIC 0 Online 4x 10GE (LAN/WAN) XFP
PIC 1 Online 4x OC-192 SONET XFP
Slot 5 Online FPC Type 2-ES
PIC 0 Online Load Type 2
PIC 1 Online 8x 1GE(LAN), IQ2E
PIC 2 Online 8x 1GE(LAN), IQ2E
PIC 3 Online 1x OC-48-12-3 SFP
Slot 7 Online FPC Type 4
PIC 0 Online 4x 10GE (LAN/WAN) XFP

```

show chassis fpc (T4000 Router)

```
user@host> show chassis fpc
```

```

regress@stymphalian# run show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%)	Memory DRAM (MB)	Utilization (%)
			Total Interrupt	Heap	Buffer
0	Online	48	15 0	2816 21	27
1	Empty				
2	Empty				
3	Online	51	15 0	2816 21	27
4	Empty				
5	Online	39	8 0	2048 6	23
6	Online	49	15 0	2816 21	27
7	Empty				

show chassis fpc detail (T4000 Router)

```
user@host> show chassis fpc detail
```

```
Slot 0 information:
```

```

State Online
Temperature 48 degrees C / 118 degrees F
Total CPU DRAM 2816 MB
Total SRAM 1554 MB
Total SDRAM 10752 MB
Start time 2012-02-09 22:56:25 PST
Uptime 2 hours, 40 minutes, 52 seconds

```

```
Slot 3 information:
```

```

State Online
Temperature 51 degrees C / 123 degrees F
Total CPU DRAM 2816 MB
Total SRAM 1554 MB
Total SDRAM 10752 MB
Start time 2012-02-09 22:56:22 PST
Uptime 2 hours, 40 minutes, 55 seconds

```

```
Slot 5 information:
```

```

State Online
Temperature 39 degrees C / 102 degrees F
Total CPU DRAM 2048 MB
Total SRAM 128 MB
Total SDRAM 2560 MB
Start time 2012-02-09 22:51:27 PST
Uptime 2 hours, 45 minutes, 50 seconds

```

```
Slot 6 information:
```

```

State Online
Temperature 49 degrees C / 120 degrees F
Total CPU DRAM 2816 MB
Total SRAM 1554 MB
Total SDRAM 10752 MB

```

```

Start time          2012-02-09 22:56:29 PST
Uptime              2 hours, 40 minutes, 48 seconds

```

show chassis fpc pic-status (T4000 Router)

```

user@host> show chassis fpc pic-status
Slot 0  Online      FPC Type 5-3D
  PIC 0  Online      12x10GE (LAN/WAN) SFPP
  PIC 1  Online      12x10GE (LAN/WAN) SFPP
Slot 3  Online      FPC Type 5-3D
  PIC 0  Online      1x100GE
  PIC 1  Online      12x10GE (LAN/WAN) SFPP
Slot 5  Online      FPC Type 4-ES
  PIC 0  Online      100GE
  PIC 1  Online      100GE CFP
Slot 6  Online      FPC Type 5-3D
  PIC 0  Online      12x10GE (LAN/WAN) SFPP
  PIC 1  Online      12x10GE (LAN/WAN) SFPP

```

show chassis fpc (QFX Series)

```

user@switch> show chassis fpc
Temp CPU Utilization (%) Memory      Utilization (%)
Slot State              (C) Total Interrupt    DRAM (MB) Heap      Buffer
0 Online                26      2          0        2820      0        49

```

show chassis fpc detail (QFX3500 Switches)

```

user@switch> show chassis fpc detail
Slot 0 information:
  State                      Online
  Temperature                 28 degrees C / 82 degrees F
  Total CPU DRAM              2820 MB
  Total SRAM                  0 MB
  Total SDRAM                 0 MB
  Start time                  2010-09-20 01:34:13 PDT
  Uptime                      3 days, 3 hours, 31 minutes, 48 seconds

```

show chassis fpc pic-status (QFX3500 Switches)

```

user@switch> show chassis fpc pic-status
Slot 0  Online      QFX 48x10G 4x40G Switch
  PIC 0  Online      48x 10G-SFP+
  PIC 1  Online      15x 10G-SFP+

```

show chassis fpc interconnect-device (QFabric Switch)

```

user@switch> show chassis fpc interconnect-device interconnect1
FPC status:
Temp
Slot State      (C)
0 Online        0
1 Online        0
2 Online        0
3 Online        0
4 Online        0
5 Online        0
6 Online        0
7 Online        0
8 Online        0
9 Online        0
10 Online       0

```

11	Online	0
12	Online	0
13	Online	0
14	Online	0
15	Online	0

show chassis fpc interconnect-device (QFabric Switch)

```
user@switch> show chassis fpc interconnect-device interconnect1 3
FPC status:
```

Slot	State	Temp (C)
3	Online	0

show chassis fpc interconnect-device detail (QFabric Switch)

```
user@switch> show chassis fpc interconnect-device interconnect1 3 detail
Slot 3 information:
```

State	Online
Temperature	0 degrees C / 32 degrees F
Start time	2011-08-18 10:45:04 PDT
Uptime	1 minute, 49 seconds

show chassis fpc pic-status interconnect-device (QFabric Switch)

```
user@switch> show chassis fpc pic-status interconnect-device interconnect1
```

Slot 0	Online	QFX 16-port QSFP+ Front Card
PIC 0	Online	16x 40G-QSFP+
PIC 1	Online	16x 40G-GE
Slot 1	Online	QFX 16-port QSFP+ Front Card
PIC 0	Online	16x 40G-QSFP+
PIC 1	Online	16x 40G-GE
Slot 2	Online	QFX 16-port QSFP+ Front Card
PIC 0	Online	16x 40G-QSFP+
PIC 1	Online	16x 40G-GE
Slot 3	Online	QFX 16-port QSFP+ Front Card
PIC 0	Online	16x 40G-QSFP+
PIC 1	Online	16x 40G-GE
Slot 4	Online	QFX 16-port QSFP+ Front Card
PIC 0	Online	16x 40G-QSFP+
PIC 1	Online	16x 40G-GE
Slot 5	Online	QFX 16-port QSFP+ Front Card
PIC 0	Online	16x 40G-QSFP+
PIC 1	Online	16x 40G-GE
Slot 6	Online	QFX 16-port QSFP+ Front Card
PIC 0	Online	16x 40G-QSFP+
PIC 1	Online	16x 40G-GE
Slot 7	Online	QFX 16-port QSFP+ Front Card
PIC 0	Online	16x 40G-QSFP+
PIC 1	Online	16x 40G-GE
Slot 8	Online	QFX Fabric Rear Card
PIC 0	Online	16x 40G-GE
Slot 9	Online	QFX Fabric Rear Card
PIC 0	Online	16x 40G-GE
Slot 10	Online	QFX Fabric Rear Card
PIC 0	Online	16x 40G-GE
Slot 11	Online	QFX Fabric Rear Card
PIC 0	Online	16x 40G-GE
Slot 12	Online	QFX Fabric Rear Card
PIC 0	Online	16x 40G-GE
Slot 13	Online	QFX Fabric Rear Card
PIC 0	Online	16x 40G-GE

```

Slot 14 Online      QFX Fabric Rear Card
PIC 0  Online      16x 40G-GE
Slot 15 Online      QFX Fabric Rear Card
PIC 0  Online      16x 40G-GE

```

show chassis fpc pic-status node-device (QFabric Switch)

```

user@switch> show chassis fpc pic-status node-device node1
Slot node1 Online      QFX 48x10G 4x40G Switch
PIC 0  Online      48x 10G-SFP+
PIC 1  Online      4x 40G-QSFP+

```

show chassis fpc (PTX5000 Packet Transport Switch)

```

user@host> show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%) Total Interrupt	Memory DRAM (MB)	Utilization (%) Heap Buffer
0	Empty				
1	Empty				
2	Online	50	6 0	2816	5 27
3	Empty				
4	Empty				
5	Online	48	9 0	2816	5 27
6	Empty				
7	Online	49	8 0	2816	5 27

show chassis fpc detail (PTX5000 Packet Transport Switch)

```

user@host> show chassis fpc detail
Slot 2 information:
State Online
Temperature 35 degrees C / 95 degrees F (PMB)
Temperature 35 degrees C / 95 degrees F (Intake)
Temperature 50 degrees C / 122 degrees F (Exhaust A)
Temperature 54 degrees C / 129 degrees F (Exhaust B)
Temperature 54 degrees C / 129 degrees F (TL0)
Temperature 52 degrees C / 125 degrees F (TQ0)
Temperature 61 degrees C / 141 degrees F (TL1)
Temperature 58 degrees C / 136 degrees F (TQ1)
Temperature 57 degrees C / 134 degrees F (TL2)
Temperature 58 degrees C / 136 degrees F (TQ2)
Temperature 62 degrees C / 143 degrees F (TL3)
Temperature 61 degrees C / 141 degrees F (TQ3)
Total CPU DRAM 2816 MB
Total SRAM 0 MB
Total SDRAM 0 MB
Start time 2012-01-12 12:05:42 PST
Uptime 3 hours, 14 minutes, 7 seconds
Slot 5 information:
State Online
Temperature 35 degrees C / 95 degrees F (PMB)
Temperature 34 degrees C / 93 degrees F (Intake)
Temperature 48 degrees C / 118 degrees F (Exhaust A)
Temperature 53 degrees C / 127 degrees F (Exhaust B)
Temperature 54 degrees C / 129 degrees F (TL0)
Temperature 52 degrees C / 125 degrees F (TQ0)
Temperature 69 degrees C / 156 degrees F (TL1)
Temperature 56 degrees C / 132 degrees F (TQ1)
Temperature 54 degrees C / 129 degrees F (TL2)
Temperature 56 degrees C / 132 degrees F (TQ2)
Temperature 59 degrees C / 138 degrees F (TL3)

```

```

Temperature                60 degrees C / 140 degrees F (TQ3)
Total CPU DRAM              2816 MB
Total SRAM                  0 MB
Total SDRAM                  0 MB
Start time                  2012-01-12 12:05:43 PST
Uptime                      3 hours, 14 minutes, 6 seconds
Slot 7 information:
State                       Online
Temperature                  35 degrees C / 95 degrees F (PMB)
Temperature                  33 degrees C / 91 degrees F (Intake)
Temperature                  50 degrees C / 122 degrees F (Exhaust A)
Temperature                  55 degrees C / 131 degrees F (Exhaust B)
Temperature                  56 degrees C / 132 degrees F (TL0)
Temperature                  56 degrees C / 132 degrees F (TQ0)
Temperature                  61 degrees C / 141 degrees F (TL1)
Temperature                  57 degrees C / 134 degrees F (TQ1)
Temperature                  55 degrees C / 131 degrees F (TL2)
Temperature                  59 degrees C / 138 degrees F (TQ2)
Temperature                  62 degrees C / 143 degrees F (TL3)
Temperature                  62 degrees C / 143 degrees F (TQ3)
Total CPU DRAM              2816 MB
Total SRAM                  0 MB
Total SDRAM                  0 MB
Start time                  2012-01-12 12:05:44 PST
Uptime                      3 hours, 14 minutes, 5 seconds

```

show chassis fpc pic-status (PTX5000 Packet Transport Switch)

```

user@host> show chassis fpc pic-status
Slot 2  Online      FPC
  PIC 0  Online      24x 10GE(LAN) SFP+
  PIC 1  Online      24x 10GE(LAN) SFP+
Slot 5  Online      FPC
  PIC 0  Online      24x 10GE(LAN) SFP+
  PIC 1  Online      2x 40GE CFP
Slot 7  Online      FPC
  PIC 0  Online      24x 10GE(LAN) SFP+
  PIC 1  Online      2x 40GE CFP

```

show chassis fpc (ACX2000 Universal Access Router)

```

user@host> show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%)	Memory DRAM (MB)	Utilization (%)
			Total Interrupt	Heap	Buffer
0	Online	61	17 6	512 21	37

show chassis fpc 0 (ACX2000 Universal Access Router)

```

user@host> show chassis fpc 0

```

Slot	State	Temp (C)	CPU Utilization (%)	Memory DRAM (MB)	Utilization (%)
			Total Interrupt	Heap	Buffer
0	Online	61	17 6	512 21	37

show chassis fpc detail (ACX2000 Universal Access Router)

```

user@host> show chassis fpc detail
Slot 0 information:
State                       Online
Temperature                  61 degrees C / 141 degrees F
Total CPU DRAM               512 MB
Start time                   2012-05-29 02:52:06 PDT
Uptime                       27 minutes, 17 seconds

```

show chassis fpc pic-status (ACX2000 Universal Access Router)

```
user@host> show chassis fpc pic-status
Slot 0  Online
  PIC 0  Online      16x CHE1T1, RJ48
  PIC 1  Online      8x 1GE(LAN) RJ45
  PIC 2  Online      2x 1GE(LAN) SFP
  PIC 3  Online      2x 10GE(LAN) SFP+
```

show chassis FPC 1 (MX Routers with Media Services Blade [MSB])

```
user@switch> show chassis fpc 1
          Temp  CPU Utilization (%)  Memory  Utilization (%)
Slot State    (C)  Total  Interrupt  DRAM (MB) Heap    Buffer
  1  Online      34     5         0      3072     5      13
```

show chassis FPC 1 detail (MX Routers with Media Services Blade [MSB])

```
user@switch> show chassis fpc 1 detail
Slot 1 information:
  State                               Online
  Temperature                         34
  Total CPU DRAM                      3072 MB
  Total RLDRAM                       259 MB
  Total DDR DRAM                     4864 MB
  Start time:                        2012-06-19 10:51:43 PDT
  Uptime:                            16 minutes, 48 seconds
  Max Power Consumption               550 Watts
```

show chassis hardware

List of Syntax	Syntax on page 743 Syntax (EX Series) on page 743 Syntax (T4000 Router) on page 743 Syntax (TX Matrix Router) on page 743 Syntax (TX Matrix Plus Router) on page 743 Syntax (MX Series Routers) on page 743 Syntax (MX104, MX2010, and MX2020 3D Universal Edge Routers) on page 743 Syntax (QFX Series) on page 744 Syntax (OCX Series) on page 744 Syntax (PTX Series Packet Transport Routers) on page 744 Syntax (ACX Series Universal Access Routers) on page 744
Syntax	show chassis hardware <detail extensive> <clei-models> <models>
Syntax (EX Series)	show chassis hardware <clei-models> <detail extensive> <models> <satellite [slot-id <i>slot-id</i> device-alias <i>alias-name</i>]>
Syntax (T4000 Router)	show chassis hardware <clei-models> <detail extensive> <models>
Syntax (TX Matrix Router)	show chassis hardware <clei-models> <detail extensive> <models> <lcc <i>number</i> scc>
Syntax (TX Matrix Plus Router)	show chassis hardware <clei-models> <detail extensive> <models> <lcc <i>number</i> sfc <i>number</i> >
Syntax (MX Series Routers)	show chassis hardware <detail extensive> <clei-models> <models> <all-members> <local> <member <i>member-id</i> >
Syntax (MX104, MX2010, and MX2020)	show chassis hardware <clei-models>

3D Universal Edge Routers)	<detail extensive> <models> <satellite [slot-id <i>slot-id</i> device-alias <i>alias-name</i>]>
Syntax (QFX Series)	show chassis hardware <detail extensive> <clei-models> <interconnect-device <i>name</i> > <node-device <i>name</i> > <models>
Syntax (OCX Series)	show chassis hardware <detail extensive> <clei-models> <models>
Syntax (PTX Series Packet Transport Routers)	show chassis hardware <detail extensive> <clei-models> <models>
Syntax (ACX Series Universal Access Routers)	show chassis hardware <detail extensive> <clei-models> <models>
Release Information	Command introduced before Junos OS Release 7.4. models option introduced in Junos OS Release 8.2. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for QFX Series. Command introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers. Command introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers. Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers. Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers. Command introduced in Junos OS Release 13.2 for MX104 3D Universal Edge Routers. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series. satellite option introduced in Junos OS Release 14.2R3.
Description	Display a list of all Flexible PIC Concentrators (FPCs) and PICs installed in the router or switch chassis, including the hardware version level and serial number. In the EX Series switch command output, FPC refers to the following: <ul style="list-style-type: none">• On EX2200 switches, EX3200 switches, EX4200 standalone switches, and EX4500 switches—Refers to the switch; FPC <i>number</i> is always 0.• On EX4200 switches in a Virtual Chassis configuration—Refers to the member of a Virtual Chassis; FPC <i>number</i> equals the member ID, from 0 through 9.• On EX8208 and EX8216 switches—Refers to a line card; FPC <i>number</i> equals the slot number for the line card.

On QFX3500, QFX5100, OCX Series standalone switches, and PTX1000 routers, both the FPC and FPC *number* are always 0.

On T4000 Type 5 FPCs, there are no **top temperature sensor** or **bottom temperature sensor** parameters. Instead, **fan intake temperature sensor** and **fan exhaust temperature sensors** parameters are displayed.

Starting from Junos OS Release 11.4, the output of the **show chassis hardware models** operational mode command displays the enhanced midplanes FRU model numbers (CHAS-BP3-MX240-S, CHAS-BP3-MX480-S or CHAS-BP3-MX960-S) based on the router. Prior to release 11.4, the FRU model numbers are left blank when the router has enhanced midplanes. Note that the enhanced midplanes are introduced through the Junos OS Release 13.3, but can be supported on all Junos OS releases.

Starting with Junos OS Release 14.1, the output of the **show chassis hardware detail | extensive | clei-models | models** operational mode command displays the new DC power supply module (PSM) and power distribution unit (PDU) that are added to provide power to the high-density FPC (FPC2-PTX-P1A) and other components in a PTX5000 Packet Transport Router.

Options **none**—Display information about hardware. For a TX Matrix router, display information about the TX Matrix router and its attached T640 routers. For a TX Matrix Plus router, display information about the TX Matrix Plus router and its attached routers.

clei-models—(Optional) Display Common Language Equipment Identifier (CLEI) barcode and model number for orderable field-replaceable units (FRUs).

detail—(Optional) Include RAM and disk information in output.

extensive—(Optional) Display ID EEPROM information.

all-members—(MX Series routers only) (Optional) Display hardware-specific information for all the members of the Virtual Chassis configuration.

interconnect-device name—(QFabric systems only) (Optional) Display hardware-specific information for the Interconnect device.

lcc number—(TX Matrix routers and TX Matrix Plus router only) (Optional) On a TX Matrix router, display hardware information for a specified T640 router (line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, display hardware information for a specified router (line-card chassis) that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.

- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(MX Series routers only) (Optional) Display hardware-specific information for the local Virtual Chassis members.

member *member-id*—(MX Series routers and EX Series switches) (Optional) Display hardware-specific information for the specified member of the Virtual Chassis configuration. Replace *member-id* variable with a value 0 or 1.

models—(Optional) Display model numbers and part numbers for orderable FRUs and, for components that use ID EEPROM format v2, the CLEI code.

node-device *name*—(QFabric systems only) (Optional) Display hardware-specific information for the Node device.

satellite [*slot-id slot-id* | device-alias *alias-name*]—(Junos Fusion only) (Optional) Display hardware information for the specified satellite device in a Junos Fusion, or for all satellite devices in the Junos Fusion if no satellite devices are specified.

scc—(TX Matrix router only) (Optional) Display hardware information for the TX Matrix router (switch-card chassis).

sfc *number*—(TX Matrix Plus router only) (Optional) Display hardware information for the TX Matrix Plus router (switch-fabric chassis). Replace *number* variable with 0.

Additional Information The **show chassis hardware detail** command now displays DIMM information for the following Routing Engines:

Table 59: Routing Engines Displaying DIMM Information

Routing Engines	Routers
RE-S-1800x2 and RE-S-1800x4	MX240, MX480, and MX960 routers
RE-A-1800x2	M120 and M320 routers

In Junos OS Release 11.4 and later, the output for the **show chassis hardware models** operational mode command for MX Series routers display the enhanced midplanes FRU model numbers—CHAS-BP3-MX240-S, CHAS-BP3-MX480-S, or CHAS-BP3-MX960-S—based on the router. In releases before Junos OS Release 11.4, the FRU model numbers are left blank when the router has enhanced midplanes. Note that the enhanced midplanes are introduced through Junos OS Release 13.3, but can be supported on all Junos OS releases.

Required Privilege Level view

Related Documentation	<ul style="list-style-type: none"> • <i>show chassis power</i>
List of Sample Output	<p> show chassis hardware (EX8216 Switch) on page 753 show chassis hardware clei-models (EX8216 Switch) on page 754 show chassis hardware clei-models (T1600 Router) on page 755 show chassis hardware (EX2300-C Switch) on page 756 show chassis hardware (EX2300 Switch) on page 756 show chassis hardware detail (EX4200 Switch) on page 756 show chassis hardware (EX4300 Switch) on page 756 show chassis hardware models (EX4500 Switch) on page 757 show chassis hardware detail (EX9200 Switch) on page 757 show chassis hardware (M7i Router) on page 758 show chassis hardware (M10 Router) on page 758 show chassis hardware models (M10 Router) on page 759 show chassis hardware (M20 Router) on page 759 show chassis hardware models (M20 Router) on page 760 show chassis hardware (M40 Router) on page 760 show chassis hardware (M40e Router) on page 761 show chassis hardware (M120 Router) on page 761 show chassis hardware detail (M120 Router) on page 762 show chassis hardware models (M120 Router) on page 763 show chassis hardware (M160 Router) on page 764 show chassis hardware models (M160 Router) on page 764 show chassis hardware detail (M160 Router) on page 765 show chassis hardware (M320 Router) on page 766 show chassis hardware models (M320 Router) on page 767 show chassis hardware (MX5 Router) on page 768 show chassis hardware (MX10 Router) on page 768 show chassis hardware (MX40 Router) on page 769 show chassis hardware (Fixed MX80 Router) on page 769 show chassis hardware (Modular MX80 Router) on page 770 show chassis hardware (MX104 Router) on page 770 show chassis hardware detail (MX104 Router) on page 771 show chassis hardware extensive (MX104 Router) on page 771 show chassis hardware models (MX104 Router) on page 775 show chassis hardware clei-models (MX104 Router) on page 775 show chassis hardware (MX240 Router) on page 775 show chassis hardware detail (MX 240 Router with Routing Engine Displaying DIMM Information) on page 776 show chassis hardware (MX240 Router with Enhanced MX SCB) on page 776 show chassis hardware (MX480 Router) on page 777 show chassis hardware (MX480 Router with Enhanced MX SCB) on page 777 show chassis hardware (MX480 Routers with MPC5E and Built-in OTN PIC) on page 778 show chassis hardware detail (MX480 Routers with MPC5E and Built-in OTN PIC) on page 779 show chassis hardware extensive (MX480 Routers with MPC5E and Built-in OTN PIC) on page 781 show chassis hardware (MX960 Router) on page 784 </p>

[show chassis hardware \(MX960 Router with Bidirectional Optics\) on page 784](#)
[show chassis hardware \(MX960 Router with Enhanced MX SCB\) on page 785](#)
[show chassis hardware models \(MX960 Router with Enhanced MX SCB\) on page 787](#)
[show chassis hardware \(MX960 Router with MPC5EQ\) on page 787](#)
[show chassis hardware detail \(MX960 Router\) on page 790](#)
[show chassis hardware detail \(MX960 Router with MPC5EQ\) on page 791](#)
[show chassis hardware extensive \(MX960 Router with MPC5EQ\) on page 794](#)
[show chassis hardware models \(MX960 Router with MPC5EQ\) on page 802](#)
[show chassis hardware clei-models \(MX960 Router with MPC5EQ\) on page 803](#)
[show chassis hardware \(MX2010 Router\) on page 803](#)
[show chassis hardware detail \(MX2010 Router\) on page 806](#)
[show chassis hardware extensive \(MX2010 Router\) on page 810](#)
[show chassis hardware models \(MX2010 Router\) on page 816](#)
[show chassis hardware clei-models \(MX2010 Routers\) on page 816](#)
[show chassis hardware \(MX2010 Routers with MPC6E and OTN MIC\) on page 817](#)
[show chassis hardware detail \(MX2010 Routers with MPC6E and OTN MIC\) on page 819](#)
[show chassis hardware extensive \(MX2010 Routers with MPC6E and OTN MIC\) on page 821](#)
[show chassis hardware \(MX2020 Router\) on page 826](#)
[show chassis hardware detail \(MX2020 Router\) on page 834](#)
[show chassis hardware models \(MX2020 Router\) on page 843](#)
[show chassis hardware clei-models \(MX2020 Router\) on page 844](#)
[show chassis hardware \(MX2020 Router with MPC5EQ and MPC6E\) on page 846](#)
[show chassis hardware detail \(MX2020 Router with MPC5EQ and MPC6E\) on page 850](#)
[show chassis hardware extensive \(MX2020 Router with MPC5EQ and MPC6E\) on page 852](#)
[show chassis hardware models \(MX2020 Routers with MPC5EQ and MPC6E\) on page 857](#)
[show chassis hardware clei-models \(MX2020 Router with MPC5EQ and MPC6E\) on page 858](#)
[show chassis hardware \(MX Series routers with ATM MIC\) on page 860](#)
[show chassis hardware \(MX240, MX480, MX960 routers with Application Services Modular Line Card\) on page 860](#)
[show chassis hardware extensive \(MX240, MX480, MX960 Routers with Application Services Modular Line Card\) on page 861](#)
[show chassis hardware \(MX480 Router with MPC4E\) on page 862](#)
[show chassis hardware \(MX2020 Router with MPC4E\) on page 862](#)
[show chassis hardware \(MX5, MX10, MX40, MX80, MX240, MX480, and MX960 Routers with Enhanced 20-Port Gigabit Ethernet MIC\) on page 864](#)
[show chassis hardware models \(MX5, MX10, MX40, MX80, MX240, MX480, and MX960 Routers with Enhanced 20-Port Gigabit Ethernet MIC\) on page 865](#)
[show chassis hardware \(T320 Router\) on page 865](#)
[show chassis hardware \(T640 Router\) on page 866](#)
[show chassis hardware models \(T640 Router\) on page 867](#)
[show chassis hardware extensive \(T640 Router\) on page 867](#)
[show chassis hardware \(T4000 Router\) on page 868](#)
[show chassis hardware \(T4000 Router with 16-GB Line Card Chassis \(LCC\) Routing Engine\) on page 870](#)
[show chassis hardware \(T4000 Router with LSR FPC\) on page 871](#)

[show chassis hardware clei-models \(T4000 Router\) on page 871](#)
[show chassis hardware detail \(T4000 Router\) on page 871](#)
[show chassis hardware models \(T4000 Router\) on page 873](#)
[show chassis hardware lcc \(TX Matrix Router\) on page 874](#)
[show chassis hardware scc \(TX Matrix Router\) on page 875](#)
[show chassis hardware \(T1600 Router\) on page 875](#)
[show chassis hardware \(TX Matrix Plus Router\) on page 877](#)
[show chassis hardware sfc \(TX Matrix Plus Router\) on page 882](#)
[show chassis hardware extensive \(TX Matrix Plus Router\) on page 884](#)
[show chassis hardware clei-models \(TX Matrix Plus Router\) on page 885](#)
[show chassis hardware detail \(TX Matrix Plus Router\) on page 887](#)
[show chassis hardware models \(TX Matrix Plus Router\) on page 889](#)
[show chassis hardware \(TX Matrix Plus Router with 3D SIBs\) on page 892](#)
[show chassis hardware clei-models \(TX Matrix Plus Router with 3D SIBs\) on page 895](#)
[show chassis hardware detail \(TX Matrix Plus Router with 3D SIBs\) on page 899](#)
[show chassis hardware lcc \(TX Matrix Plus Router with 3D SIBs\) on page 902](#)
[show chassis hardware sfc \(TX Matrix Plus Router with 3D SIBs\) on page 903](#)
[show chassis hardware \(16-Port 10-Gigabit Ethernet MPC with SFP+ Optics \[MX Series Routers\]\) on page 905](#)
[show chassis hardware \(MPC3E \[MX Series Routers\]\) on page 905](#)
[show chassis hardware \(QFX3500 Switches\) on page 906](#)
[show chassis hardware detail \(QFX3500 Switches\) on page 907](#)
[show chassis hardware models \(QFX3500 Switches\) on page 908](#)
[show chassis hardware clei-models \(QFX3500 Switches\) on page 908](#)
[show chassis hardware clei-models \(QFX5100 Switches\) on page 908](#)
[show chassis hardware interconnect-device \(QFabric Systems\) on page 908](#)
[show chassis hardware node-device \(QFabric Systems\) on page 908](#)
[show chassis hardware \(PTX5000 Packet Transport Router\) on page 909](#)
[show chassis hardware \(PTX5000 Packet Transport Router with AC PSM and PDU\) on page 910](#)
[show chassis hardware \(PTX5000 Packet Transport Router with FPC2-PTX-P1A\) on page 911](#)
[show chassis hardware clei-models \(PTX5000 Packet Transport Router\) on page 911](#)
[show chassis hardware clei-models \(PTX5000 Packet Transport Router with AC PSM and PDU\) on page 912](#)
[show chassis hardware clei-models \(PTX5000 Packet Transport Router with FPC2-PTX-P1A\) on page 912](#)
[show chassis hardware detail \(PTX5000 Packet Transport Router\) on page 912](#)
[show chassis hardware detail \(PTX5000 Packet Transport Router with AC PSM and PDU\) on page 914](#)
[show chassis hardware detail \(PTX5000 Packet Transport Router with FPC2-PTX-P1A\) on page 914](#)
[show chassis hardware models \(PTX5000 Packet Transport Router\) on page 915](#)
[show chassis hardware models \(PTX5000 Packet Transport Router with AC PSM and PDU\) on page 915](#)
[show chassis hardware models \(PTX5000 Packet Transport Router with FPC2-PTX-P1A\) on page 916](#)
[show chassis hardware extensive \(PTX5000 Packet Transport Router\) on page 916](#)
[show chassis hardware extensive \(PTX1000 Packet Transport Router\) on page 917](#)

[show chassis hardware \(MX Routers with Media Services Blade \[MSB\]\) on page 917](#)

[show chassis hardware extensive \(MX Routers with Media Services Blade \[MSB\]\) on page 918](#)

[show chassis hardware \(QFX3500 Switch running Enhanced Layer 2 Software\) on page 919](#)

[show chassis hardware \(QFX5100 Switch running Enhanced Layer 2 Software\) on page 919](#)

Output Fields [Table 60](#) lists the output fields for the **show chassis hardware** command. Output fields are listed in the approximate order in which they appear.

Table 60: show chassis hardware Output Fields

Field Name	Field Description	Level of Output
Item	<p>Chassis component:</p> <ul style="list-style-type: none"> (EX Series switches)—Information about the chassis, Routing Engine (SRE and Routing Engine modules in EX8200 switches), power supplies, fan trays, and LCD panel. Also displays information about Flexible PIC Concentrators (FPCs) and associated Physical Interface Cards (PICs). Information about the backplane, midplane, and SIBs (SF modules) is displayed for EX8200 switches. See <i>EX Series Switches Hardware and CLI Terminology Mapping</i>. (MX Series routers and EX Series switches)—Information about the backplane, Routing Engine, Power Entry Modules (PEMs), and fan trays. Also displays information about Flexible PIC Concentrators (FPCs) and associated Physical Interface Cards (PICs), Modular Port Concentrators (MPCs) and associated Modular Interface Cards (MICs), or Dense Port Concentrators (DPCs). MX80 routers have a single Routing Engine and a built-in Packet Forwarding Engine that attaches directly to MICs. The Packet Forwarding Engine has two “pseudo” FPCs (FPC 0 and FPC1). MX80 routers also have a Forwarding Engine Board (FEB). MX104 routers have a built-in Packet forwarding Engine and a Forwarding Engine Board (FEB). The Packet Forwarding Engine of the MX104 router has three “pseudo” FPCs (FPC0, FPC1, and FPC2). (M Series routers, except for the M320 router)—Information about the backplane; power supplies; fan trays; Routing Engine; maxicab (the connection between the Routing Engine and the backplane, for the M40 router only); SCB, SSB, SFM, or FEB; MCS and PCG (for the M160 router only); each FPC and PIC; and each fan, blower, and impeller. (M120, M320, and T Series routers)—Information about the backplane, power supplies, fan trays, midplane, FPM (craft interface), CIP, PEM, SCG, CB, FPC, PIC, SFP, SPMB, and SIB. (QFX Series)—Information about the chassis, Pseudo CB, Routing Engine, power supplies, fan trays, Interconnect devices, and Node devices. Also displays information about Flexible PIC Concentrators (FPCs) and associated Physical Interface Cards (PICs). (PTX Series)—Information about the chassis, midplane, craft interface (FPM), power distribution units (PDUs) and Power Supply Modules (PSMs), Centralized Clock Generators (CCGs), Routing Engines, Control Boards (CBs) and Switch Processor Mezzanine Boards (SPMBs), Flexible PIC Concentrators (FPCs), PICs, Switch Interface Boards (SIBs), and fan trays (vertical and horizontal). (MX2010 and MX2020 routers)—Information about the chassis, midplane, craft interface (FPM), power midplane (PMP), Power Supply Modules (PSMs), Power Distribution Modules (PDMs), Routing Engines, Control Boards (CBs) and Switch Processor Mezzanine Boards (SPMBs), Switch Fabric Boards (SFBs), Flexible PIC Concentrators (FPCs), PICs, adapter cards (ADCs) and fan trays. 	All levels
Version	Revision level of the chassis component.	All levels
Part number	Part number of the chassis component.	All levels
Serial number	Serial number of the chassis component. The serial number of the backplane is also the serial number of the router chassis. Use this serial number when you need to contact Juniper Networks Customer Support about the router or switch chassis.	All levels

Table 60: show chassis hardware Output Fields (*continued*)

Field Name	Field Description	Level of Output
Assb ID or Assembly ID	(extensive keyword only) Identification number that describes the FRU hardware.	extensive
Assembly Version	(extensive keyword only) Version number of the FRU hardware.	extensive
Assembly Flags	(extensive keyword only) Flags.	extensive
FRU model number	(clei-models , extensive , and models keyword only) Model number of the FRU hardware component.	none specified
CLEI code	(clei-models and extensive keyword only) Common Language Equipment Identifier code. This value is displayed only for hardware components that use ID EEPROM format v2. This value is not displayed for components that use ID EEPROM format v1.	none specified
EEPROM Version	ID EEPROM version used by the hardware component: 0x00 (version 0), 0x01 (version 1), or 0x02 (version 2).	extensive
Description	<p>Brief description of the hardware item:</p> <ul style="list-style-type: none"> Type of power supply. Type of PIC. If the PIC type is not supported on the current software release, the output states Hardware Not Supported. Type of FPC: FPC Type 1, FPC Type 2, FPC Type 3, FPC Type 4, or FPC TypeOC192. <p>On EX Series switches, a brief description of the FPC.</p> <p>The following list shows the PIM abbreviation in the output and the corresponding PIM name.</p> <ul style="list-style-type: none"> 2x FE—Either two built-in Fast Ethernet interfaces (fixed PIM) or dual-port Fast Ethernet PIM 4x FE—4-port Fast Ethernet ePIM 1x GE Copper—Copper Gigabit Ethernet ePIM (one 10-Mbps, 100-Mbps, or 1000-Mbps port) 1x GE SFP—SFP Gigabit Ethernet ePIM (one fiber port) 2x Serial—Dual-port serial PIM 2x T1—Dual-port T1 PIM 2x E1—Dual-port E1 PIM 2x CT1E1—Dual-port channelized T1/E1 PIM 1x T3—T3 PIM (one port) 1x E3—E3 PIM (one port) 4x BRI S/T—4-port ISDN BRI S/T PIM 4x BRI U—4-port ISDN BRI U PIM 1x ADSL Annex A—ADSL 2/2+ Annex A PIM (one port, for POTS) 1x ADSL Annex B—ADSL 2/2+ Annex B PIM (one port, for ISDN) 2x SHDSL (ATM)—G SHDSL PIM (2-port two-wire module or 1-port four-wire module) 	All levels

Table 60: show chassis hardware Output Fields (*continued*)

Field Name	Field Description	Level of Output
	<ul style="list-style-type: none"> • 1x TGM550—TGM550 Telephony Gateway Module (Avaya VoIP gateway module with one console port, two analog LINE ports, and two analog TRUNK ports) • 1x DS1 TIM510—TIM510 E1/T1 Telephony Interface Module (Avaya VoIP media module with one E1 or T1 trunk termination port and ISDN PRI backup) • 4x FXS, 4x FXO, TIM514—TIM514 Analog Telephony Interface Module (Avaya VoIP media module with four analog LINE ports and four analog TRUNK ports) • 4x BRI TIM521—TIM521 BRI Telephony Interface Module (Avaya VoIP media module with four ISDN BRI ports) • Crypto Accelerator Module—For enhanced performance of cryptographic algorithms used in IP Security (IPsec) services • MPC M16x10GE—16-port 10-Gigabit Module Port Concentrator that supports SFP+ optical transceivers. (Not on EX Series switches.) • For hosts, the Routing Engine type. • For small form-factor pluggable transceiver (SFP) modules, the type of fiber: LX, SX, LH, or T. • LCD description for EX Series switches (except EX2200 switches). • MPC2—1-port MPC2 that supports two separate slots for MICs. • MPC3E—1-port MPC3E that supports two separate slots for MICs (MIC-3D-1X100GE-CFP and MIC-3D-20GE-SFP) on MX960, MX480, and MX240 routers. The MPC3E maps one MIC to one PIC (1 MIC, 1 PIC), which differs from the mapping of legacy MPCs. • 100GBASE-LR4, pluggable CFP optics • Supports the Enhanced MX Switch Control Board with fabric redundancy and existing SCBs without fabric redundancy. • Interoperates with existing MX Series line cards, including Flexible Port Concentrators (FPC), Dense Port Concentrators (DPCs), and Modular Port Concentrators (MPCs). • MPC4E—Fixed configuration MPC4E that is available in two flavors: MPC4E-3D-32XGE-SFP and MPC4E-3D-2CGE-8XGE on MX2020, MX960, MX480, and MX240 routers. • LCD description for MX Series routers 	

Sample Output

show chassis hardware (EX8216 Switch)

```

user@host> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis       REV 06   710-016845   CY0109220035   EX8216
Midplane      REV 06   710-016845   BA0909120112   EX8216-MP
CB 0          REV 22   710-020771   AX0109197723   EX8216-RE320
CB 1          REV 22   710-020771   AX0109197726   EX8216-RE320
  Routing Engine 1    BUILTIN      BUILTIN        RE-EX8216
FPC 3         REV 19   710-020683   BC0109083125   EX8200-48F
  CPU         REV 13   710-020598   BF0109144549   EX8200-CPU
FPC 4         REV 17   710-020683   BC0108500127   EX8200-48F
  CPU         REV 10   710-020598   BF0108460510   EX8200-CPU

```

PIC 0		BUILTIN	BUILTIN	48x 100 Base-QFX/1000
Base-X				
Xcvr 1	REV 01	740-011613	PE70V89	SFP-SX
Xcvr 11	REV 01	740-011613	PE70YCE	SFP-SX
Xcvr 12	REV 01	740-011613	PE70VSH	SFP-SX
Xcvr 13	REV 01	740-011613	E08C02063	SFP-SX
Xcvr 14	REV 01	740-011613	PE70VKU	SFP-SX
Xcvr 15	REV 01	740-011613	E08E03372	SFP-SX
Xcvr 21	REV 01	740-011613	PE70VAD	SFP-SX
Xcvr 22	REV 01	740-011613	E08E01228	SFP-SX
Xcvr 23	REV 01	740-011613	PE70VSL	SFP-SX
Xcvr 24	REV 01	740-011613	E08E03409	SFP-SX
Xcvr 25	REV 01	740-011613	PE70VL4	SFP-SX
Xcvr 26	REV 01	740-011613	PDQ4L2Z	SFP-SX
Xcvr 27	REV 01	740-011613	PE70WFK	SFP-SX
Xcvr 28	REV 01	740-011782	PBD2B5U	SFP-SX
Xcvr 29	REV 01	740-011613	PE70UQX	SFP-SX
Xcvr 30	REV 01	740-011613	PE70VL5	SFP-SX
Xcvr 31	REV 01	740-011613	PE70V0F	SFP-SX
Xcvr 32	REV 01	740-011613	E08C02052	SFP-SX
Xcvr 33	REV 01	740-011613	E08C02197	SFP-SX
Xcvr 34	REV 01	740-011613	PE70V0L	SFP-SX
Xcvr 35	REV 01	740-011613	E08E03390	SFP-SX
Xcvr 36	REV 01	740-011613	PDQ4VL9	SFP-SX
Xcvr 37	REV 01	740-011613	E08E03370	SFP-SX
Xcvr 38	REV 01	740-011613	E08E03362	SFP-SX
Xcvr 39	REV 01	740-011613	E08C02065	SFP-SX
Xcvr 40	REV 01	740-011613	E08E03405	SFP-SX
Xcvr 41	REV 01	740-011613	E08E03411	SFP-SX
Xcvr 43	REV 01	740-011613	E08C02171	SFP-SX
Xcvr 45	REV 01	740-011613	E08E03410	SFP-SX
FPC 13	REV 16	710-016837	BB0109051344	EX8200-8XS
CPU				
SIB 0	REV 10	710-021613	AY0109166244	EX8216-SF320
SIB 1	REV 10	710-021613	AY0109166357	EX8216-SF320
SIB 2	REV 10	710-021613	AY0109166362	EX8216-SF320
SIB 3	REV 10	710-021613	AY0109166338	EX8216-SF320
SIB 4	REV 10	710-021613	AY0109166350	EX8216-SF320
SIB 5	REV 10	710-021613	AY0109166365	EX8216-SF320
SIB 6	REV 10	710-021613	AY0109166361	EX8216-SF320
SIB 7	REV 10	710-021613	AY0109166399	EX8216-SF320
PSU 0	REV 17	740-021466	BG0709170003	EX8200-AC2K
PSU 1	REV 17	740-021466	BG0709170004	EX8200-AC2K
PSU 2	REV 17	740-021466	BG0709170020	EX8200-AC2K
PSU 3	REV 17	740-021466	BG0709170017	EX8200-AC2K
PSU 4	REV 17	740-021466	BG0709170008	EX8200-AC2K
PSU 5	REV 17	740-021466	BG0709170018	EX8200-AC2K
Top Fan Tray				
FTC 0	REV 4	760-022620	CX1209140212	EX8216-FT
FTC 1	REV 4	760-022620	CX1209140212	EX8216-FT
Bottom Fan Tray				
FTC 0	REV 4	760-022620	CX1209140211	EX8216-FT
FTC 1	REV 4	760-022620	CX1209140211	EX8216-FT
LCD 0	REV 04	710-025742	CE0109186919	EX8200 LCD

show chassis hardware clei-models (EX8216 Switch)

```

user@host> show chassis hardware clei-models
Hardware inventory:
Item          Version  Part number  CLEI code  FRU model number
Midplane      REV 08   710-016845

```

```

PSU 0          REV 05  740-023002  COUPAEAEAA  EX8200-PWR-AC3KR
PSU 1          REV 05  740-023002  COUPAEAEAA  EX8200-PWR-AC3KR
PSU 2          REV 05  740-023002  COUPAEAEAA  EX8200-PWR-AC3KR
PSU 3          REV 05  740-023002  COUPAEAEAA  EX8200-PWR-AC3KR
PSU 4          REV 05  740-023002  COUPAEAEAA  EX8200-PWR-AC3KR
PSU 5          REV 05  740-023002  COUPAEAEAA  EX8200-PWR-AC3KR
Top Fan Tray
Bottom Fan Tray

```

show chassis hardware clei-models (T1600 Router)

```

user@host> show chassis hardware clei-models
Hardware inventory:
Item                Version  Part number  CLEI code  FRU model number
Midplane            REV 03  710-005608
FPM Display         REV 05  710-002897
CIP                 REV 06  710-002895
PEM 0               Rev 07  740-017906  IPUPAC7KTA  PWR-T1600-3-80-DC-S
PEM 1               Rev 18  740-002595  PWR-T-DC-S
SCG 0               REV 15  710-003423  SCG-T-S
Routing Engine 0    REV 08  740-014082  RE-A-2000-4096-S
Routing Engine 1    REV 07  740-014082  RE-A-2000-4096-S
CB 0                REV 05  710-007655  CB-T-S
CB 1                REV 03  710-017707  CB-T-S
FPC 0               REV 07  710-013558  T640-FPC2-E2
  PIC 0             REV 01  750-010618  PB-4GE-SFP
  PIC 1             REV 06  750-001900  PB-10C48-SON-SMSR
  PIC 2             REV 14  750-001901  PB-40C12-SON-SMIR
  PIC 3             REV 07  750-001900  PB-10C48-SON-SMSR
FPC 1               REV 06  710-013553  T640-FPC1-E2
  PIC 0             REV 08  750-001072  P-1GE-SX
  PIC 1             REV 10  750-012266  PB-4GE-TYPE1-SFP-IQ2
  PIC 2             REV 22  750-005634  PB-1CHOC12SMIR-QPP
FPC 2
  PIC 0             REV 16  750-007141  PC-10GE-SFP
  PIC 1             REV 06  750-015217  PC-8GE-TYPE3-SFP-IQ2
  PIC 2             REV 05  750-004695  PC-TUNNEL
  PIC 3             REV 17  750-009553  PC-40C48-SON-SFP
FPC 3               REV 01  710-010154  T640-FPC3-E
  PIC 0             REV 07  750-012793  PC-1XGE-TYPE3-XFP-IQ2
  PIC 1             REV 25  750-007141  PC-10GE-SFP
  PIC 2             REV 17  750-009553  PC-40C48-SON-SFP
  PIC 3             REV 32  750-003700  PC-10C192-SON-VSR
FPC 4               REV 16  710-013037  T1600-FPC4-ES
  PIC 1             REV 06  750-034781  PD-1CE-CFP
FPC 5               REV 02  710-013037  T1600-FPC4-ES
  PIC 0             REV 16  750-012518  PD-40C192-SON-XFP
  PIC 1             REV 01  750-010850  PD-10C768-SON-SR
FPC 6               REV 14  710-013037  T1600-FPC4-ES
  PIC 0             REV 11  750-017405  PD-4XGE-XFP
  PIC 1             REV 13  750-017405  PD-4XGE-XFP
FPC 7               REV 09  710-007529  T640-FPC3
  PIC 0             REV 10  750-012793  PC-1XGE-TYPE3-XFP-IQ2
  PIC 1             REV 01  750-015217  PC-8GE-TYPE3-SFP-IQ2
  PIC 2             REV 01  750-015217  PC-8GE-TYPE3-SFP-IQ2
  PIC 3             REV 15  750-009450  PC-10C192-SON-SR2
SIB 0               REV 07  710-013074  SIB-I-T1600-S
SIB 1               REV 07  710-013074  SIB-I-T1600-S
SIB 2               REV 07  710-013074  SIB-I-T1600-S
SIB 3               REV 07  710-013074  SIB-I-T1600-S
SIB 4               REV 07  710-013074  SIB-I-T1600-S

```

Fan Tray 0	FANTRAY-T-S
Fan Tray 1	FANTRAY-T-S
Fan Tray 2	FAN-REAR-TX-T640-S

show chassis hardware (EX2300-C Switch)

```
user@switch> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               HV0215410003  EX2300-C-12P
Pseudo CB 0
Routing Engine 0
FPC 0          REV 04    650-059984   HV0215410003  EX2300-C-12P
  CPU          BUILTIN   BUILTIN      FPC CPU
  PIC 0        REV 04    BUILTIN      BUILTIN      12x10/100/1000 Base-T
  PIC 1        REV 04    650-059984   HV0215410003  2x10G SFP/SFP+
    Xcvr 0     REV 01    740-021309   T09K00695     SFP+-10G-LR
    Xcvr 1     REV 01    740-030658   AD1146A05JT   SFP+-10G-USR
Power Supply 0
JPSU-170W-AC
```

show chassis hardware (EX2300 Switch)

```
user@switch> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               JY0215410033  EX2300-24P
Pseudo CB 0
Routing Engine 0
FPC 0          REV 05    650-059968   JY0215410033  EX2300-24P
  CPU          BUILTIN   BUILTIN      FPC CPU
  PIC 0        REV 05    BUILTIN      BUILTIN      24x10/100/1000 Base-T
  PIC 1        REV 05    650-059968   JY0215410033  4x10G SFP/SFP+
    Xcvr 0     REV 01    740-030658   AD1125A03ES   SFP+-10G-USR
    Xcvr 1     REV 01    740-021308   AJPOTDZ       SFP+-10G-SR
    Xcvr 3     REV 01    740-021309   A9401FL       SFP+-10G-LR
Power Supply 0
JPSU-450W-AC-AFO
Fan Tray 0
(AFO)
Fan Tray 1
(AFO)
Fan Module, Airflow Out
Fan Module, Airflow Out
```

show chassis hardware detail (EX4200 Switch)

```
user@host> show chassis hardware detail
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               BM0208327733  EX4200-24T
Routing Engine 0 REV 11    750-021256   BM0208327733  EX4200-24T, 8 POE
Routing Engine 0
FPC 0          REV 11    750-021256   BM0208327733  EX4200-24T, 8 POE
  CPU          BUILTIN   BUILTIN      FPC CPU
  PIC 0        BUILTIN   BUILTIN      24x 10/100/1000 Base-T
  PIC 1        REV 03B   711-021270   AR0208162285  4x GE SFP
  BRD          REV 08    711-021264   AK0208328289  EX4200-24T, 8 POE
Power Supply 0
REV 03    740-020957   AT0508346354  PS 320W AC
Fan Tray
Fan Tray
```

show chassis hardware (EX4300 Switch)

```
user@host> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
```

```

Chassis                               PD3713160055    EX4300-48P
Routing Engine 0 REV 04 650-044930 PD3713160055    EX4300-48P
FPC 0                                REV 04 650-044930 PD3713160055    EX4300-48P
  CPU                                BUILTIN        BUILTIN        FPC CPU
  PIC 0                             REV 04 BUILTIN        BUILTIN        48x 10/100/1000 Base-T
  PIC 1                             REV 04 BUILTIN        BUILTIN        4x 40GE
Power Supply 0 REV 01 740-046871 1EDA3090026      JPSU-1100-AC-AF0-A
Fan Tray 0 (AF0)                               Fan Module, Airflow Out
Fan Tray 1 (AF0)                               Fan Module, Airflow Out

```

show chassis hardware models (EX4500 Switch)

```

user@host> show chassis hardware models
Hardware inventory:
Item          Version  Part number  Serial number  FRU model number
Routing Engine 0 REV 01 750-035700  GG0210271867  EX4500-40F-FB-C
FPC 0         REV 01 750-035700  GG0210271867  EX4500-40F-FB-C
  PIC 0       BUILTIN        BUILTIN        EX4500-40F-FB-C
Power Supply 1 REV 01 740-029654  H884FS00JC09  EX4500-PWR1-AC-FB

```

show chassis hardware detail (EX9200 Switch)

```

user@switch> show chassis hardware

Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               JN111DA44RFB  EX9208
Midplane      REV 05 710-017414  TS2912        EX9208-BP
FPM Board     REV 02 710-017254  XN1804        Front Panel Display
PEM 0         Rev 01 740-022697  QCS0906C033  PS 1.2-1.7kW; 100-240V
AC in
PEM 1         Rev 01 740-022697  QCS0906C095  PS 1.2-1.7kW; 100-240V
AC in
Routing Engine 0 REV 08 740-031116  9009122883    RE-S-EX9200-1800X4
CB 0          REV 16 750-031391  CAAW4391      EX9200-SCBEF
PC 0          REV 07 750-049612  CABJ9312      EX9200 40x1G Copper
  CPU          REV 04 711-038484  CABH8268      MPCE PMB 2G
  MIC 0        REV 02 750-049607  CABT9623      40x 1GE RJ45
    PIC 0      BUILTIN        BUILTIN        10x 1GE RJ45
    PIC 1      BUILTIN        BUILTIN        10x 1GE RJ45
    PIC 2      BUILTIN        BUILTIN        10x 1GE RJ45
    PIC 3      BUILTIN        BUILTIN        10x 1GE RJ45
FPC 1         REV 10 710-013699  CAAN3529      EX9200-40x1G-SFP
  CPU          REV 04 711-038484  CAAL7608      MPCE PMB 2G
  MIC 0        REV 26 750-028392  CAAS5151      20x 1GE SFP
    PIC 0      BUILTIN        BUILTIN        10x 1GE SFP
    PIC 1      BUILTIN        BUILTIN        10x 1GE SFP
  MIC 1        REV 26 750-028392  CAAC8006      20x 1GE SFP
    PIC 2      BUILTIN        BUILTIN        10x 1GE SFP
      Xcvr 8    REV 01 740-011613  E08L03674    SFP-SX
      Xcvr 9    REV 01 740-011613  E08M00243    SFP-SX
    PIC 3      BUILTIN        BUILTIN        10x 1GE SFP
      FPC 3     REV 10 710-013699  CAAR5261      EX9200-40x1G-SFP
  CPU          REV 04 711-038484  CAAS2118      MPCE PMB 2G
  MIC 0        REV 26 750-028392  CAAS5067      20x 1GE SFP
    PIC 0      BUILTIN        BUILTIN        10x 1GE SFP
      Xcvr 2    REV 01 740-031851  PNA7L8U      SFP-SX
      Xcvr 3    REV 02 740-011613  AM0943SEKGZ  SFP-SX
      Xcvr 4    REV 02 740-011613  AM0943SEJZ9  SFP-SX

```

PIC 1		BUILTIN	BUILTIN	10x 1GE SFP
MIC 1	REV 26	750-028392	CAAS5132	20x 1GE SFP
PIC 2		BUILTIN	BUILTIN	10x 1GE SFP
Xcvr 4	REV 01	740-011613	E08D02625	SFP-SX
Xcvr 9	REV 02	740-011613	PJH4RD9	SFP-SX
PIC 3		BUILTIN	BUILTIN	10x 1GE SFP
Xcvr 0	REV 01	740-011613	AM0813S8YME	SFP-SX
Fan Tray				Left Fan Tray

show chassis hardware (M7i Router)

```
user@host> show chassis hardware
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			31959	M7i
Midplane	REV 02	710-008761	CA0209	M7i Midplane
Power Supply 0	Rev 04	740-008537	PD10272	AC Power Supply
Routing Engine	REV 01	740-008846	1000396803	RE-5.0
CFEB	REV 02	750-009492	CA0166	Internet Processor IIv1
FPC 0				E-FPC
PIC 0	REV 04	750-003163	HJ6416	1x G/E, 1000 BASE-SX
PIC 1	REV 04	750-003163	HJ6423	1x G/E, 1000 BASE-SX
PIC 2	REV 04	750-003163	HJ6421	1x G/E, 1000 BASE-SX
PIC 3	REV 02	750-003163	HJ0425	1x G/E, 1000 BASE-SX
FPC 1				E-FPC
PIC 2	REV 01	750-009487	HM2275	ASP - Integrated
PIC 3	REV 01	750-009098	CA0142	2x F/E, 100 BASE-TX

```
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			B1157	M7i
Midplane	REV 05	710-008761	DM0840	M7i Midplane
Power Supply 0	Rev 08	740-008537	TE53755	AC Power Supply
Routing Engine	REV 07	740-011202	1000736567	RE-850
CFEB	REV 09	750-010463	DK6952	Internet Processor II
FPC 0				E-FPC
PIC 0	REV 12	750-012838	DL7993	4x 1GE(LAN), IQ2
Xcvr 0	REV 01	740-011614	PD94TDJ	SFP-LX10
Xcvr 1	REV 01	740-011615	PAD5EER	UNSUPPORTED
Xcvr 2	REV 01	740-011614	PD94THU	SFP-LX10
Xcvr 3		NON-JNPR	PDC2E7A	SFP-LX10
PIC 1	REV 03	750-023116	JT0203	4x CHSTM1 SDH CE SFP
Xcvr 0	REV 01	740-012434	AGT063832PS	SFP-SR
Xcvr 1	REV 01	740-012434	AGT063832LY	SFP-SR
Xcvr 3	REV 01	740-016064	C06J19018	SFP-LR
PIC 2	REV 15	750-014895	DM5757	MultiServices 100
PIC 3	REV 01	750-025390	JW9448	12x T1/E1 CE
FPC 1				E-FPC
PIC 2		BUILTIN	BUILTIN	1x Tunnel
PIC 3	REV 09	750-009099	DM0899	1x G/E, 1000 BASE
Xcvr 0	REV 01	740-012434	AGT07150HGJ	UNSUPPORTED
Fan Tray				Rear Fan Tray

show chassis hardware (M10 Router)

```
user@host> show chassis hardware
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			1122	M10
Midplane	REV 1.1	710-001950	S/N AC6626	
Power supply A	Rev 01	740-002497	S/N LC36095	AC
Power supply B	Rev 01	740-002497	S/N LC36100	AC

Display	REV 1.2	710-001995	S/N AC6656	
Host			18000005dfb3fb01	teknor
FEB	REV 01	710-001948	S/N AC6632	Internet Processor II
FPC 0				
PIC 0	REV 08	750-001072	S/N AB2485	1x G/E, 1000 BASE-SX
PIC 1	REV 01	750-000613	S/N AA1048	1x OC-12 SONET, SMIR
FPC 1				
Fan Tray 0				FANTRAY-M10I-S
Fan Tray 1				FANTRAY-M10I-S

show chassis hardware models (M10 Router)

```
user@host> show chassis hardware models
```

Hardware inventory:

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 04	710-008920		CHAS-MP-M10i-S
Power Supply 0	Rev 06	740-008537		PWR-M10i-M7i-AC-S
Power Supply 1	Rev 06	740-008537		PWR-M10i-M7i-AC-S
HCM 0	REV 03	710-010580		HCM-M10i-S
HCM 1	REV 03	710-010580		HCM-M10i-S
Routing Engine 0	REV 09	740-009459		RE-400-256-S
CFEB 0	REV 05	750-010465		FEB-M10i-M7i-S
FPC 0				
PIC 0	REV 10	750-002971		PE-40C3-SON-MM
PIC 1	REV 11	750-002992		PE-4FE-TX
PIC 2	REV 03	750-002977		PE-20C3-ATM-MM
PIC 3	REV 08	750-005724		PE-20C3-ATM2-MM
FPC 1				
PIC 2	REV 12	750-008425		PE-AS
PIC 3	REV 13	750-005636		PE-4CHDS3-QPP
Fan Tray 0				FANTRAY-M10I-S
Fan Tray 1				FANTRAY-M10I-S

show chassis hardware (M20 Router)

```
user@host> show chassis hardware
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			20033	M20
Backplane	REV 07	710-001517	S/N AA7940	
Power supply B	Rev 01	740-001465	S/N 000001	AC
Display	REV 02	710-001519	S/N AA9704	
Host 0			98000004f8f27501	teknor
SSB slot 0	REV 01	710-001951	S/N AD5905	Internet Processor II
SSRAM bank 0	REV 01	710-001385	S00480	2 MB
SSRAM bank 1	REV 01	710-001385	S00490	2 MB
SSRAM bank 2	REV 01	710-001385	S001:?	2 MB
SSRAM bank 3	REV 01	710-001385	S00483	2 MB
SSB slot 1	N/A	N/A	N/A	Backup
FPC 1	REV 01	710-001292	S/N AB7528	
SSRAM	REV 01	710-000077	S/N 304209	1 MB
SDRAM bank 0	REV 01	710-000099	S/N 000603	64 MB
SDRAM bank 1	REV 01	710-000099	S/N 000414	64 MB
PIC 0	REV 03	750-000612	S/N AB8433	2x OC-3 ATM, MM
PIC 1	REV 01	750-000616	S/N AA1168	1x OC-12 ATM, MM
PIC 2	REV 01	750-000613	S/N AA1008	1x OC-12 SONET, SMIR
PIC 3	REV 01	750-002501	S/N AD5810	4x E3
FPC 2	REV 01	710-001292	S/N AC0119	
SSRAM	REV 01	710-000077	S/N 503241	1 MB
SDRAM bank 0	REV 01	710-000099	S/N 306835	64 MB
SDRAM bank 1	REV 01	710-000099	S/N 306832	64 MB

Fan Tray 0
Fan Tray 1
Fan Tray 2
Fan Tray 3

Front Upper Fan Tray
Front Middle Fan Tray
Front Bottom Fan Tray
Rear Fan Tray

show chassis hardware models (M20 Router)

user@host> show chassis hardware models

Hardware inventory:

Item	Version	Part number	CLEI code	FRU model number
Backplane	REV 03	710-002334		CHAS-MP-M20-S
Power Supply A	REV 06	740-001465		PWR-M20-AC-S
Display	REV 04	710-001519		CRAFT-M20-S
Routing Engine 0	REV 06	740-003239		RE-333-768-S
Routing Engine 1	REV 06	740-003239		RE-333-768-S
SSB 0	REV 02	710-001951		SSB-E-M20
SSB 1	N/A	N/A		
FPC 0	REV 03	710-003308		FPC-E
PIC 0	REV 08	750-002303		P-4FE-TX
PIC 1	REV 07	750-004745		P-2MCDS3
PIC 2	REV 03	750-002965		PE-4CHDS3
FPC 1	REV 03	710-003308		FPC-E
PIC 0	REV 03	750-002914		P-20C3-ATM-MM
Fan Tray 0				FANTRAY-F-M20-S
Fan Tray 1				FANTRAY-F-M20-S
Fan Tray 2				FANTRAY-F-M20-S
Fan Tray 3				FANTRAY-R-M20-S

show chassis hardware (M40 Router)

user@host> show chassis hardware

Hardware inventory:

Item	Version	Part number	Serial number	Description
Backplane	REV 02	710-000073	S/N AA0053	
Power supply A	Rev 2	740-000235	S/N 000042	DC
Maxicab	REV X1	710-000229	S/N AA0139	
Minicab	REV X1	710-000482	S/N AA0201	
Display	REV 06	710-000150	S/N AA0905	
Host				cpv5000
SCB	REV X1	710-000075	S/N AA0158	Internet Processor I
SSRAM bank 0	REV 02	710-000077	S/N AA2267	1 MB
SSRAM bank 1	REV 02	710-000077	S/N AA2270	1 MB
SSRAM bank 2	REV 02	710-000077	S/N AA2269	1 MB
SSRAM bank 3	REV 02	710-000077	S/N AA2268	1 MB
FPC 0	REV 01	710-000175	S/N AA0048	
SSRAM	REV 01	710-000077	S/N AA2333	1 MB
SDRAM bank 0	REV 01	710-000099	S/N AA2332	64 MB
SDRAM bank 1	REV X1	710-000099	S/N AA2337	64 MB
PIC 0	REV 04	750-000613	S/N aa0343	1x OC-12 SONET, SMIR
PIC 1	REV 04	750-000613	S/N AA0379	1x OC-12 SONET, SMIR
PIC 2	REV 04	750-000613	S/N AA0377	1x OC-12 SONET, SMIR
PIC 3	REV 04	750-000613	S/N AA0378	1x Tunnel
FPC 2	REV 01	710-000175	S/N AA0042	
SSRAM	REV 02	710-000077	S/N AA2288	1 MB
SDRAM bank 0	REV 01	710-000099	S/N AA2331	64 MB
SDRAM bank 1	REV 01	710-000099	S/N AA2330	64 MB
PIC 0	REV X1	750-000603	S/N AA0143	4x OC-3 SONET, SMIR
PIC 1	REV X1	750-000615	S/N AA0149	4x OC-3 SONET, MM
PIC 2	REV X1	750-000611	S/N AA0148	4x OC-3 SONET, MM
PIC 3	REV 04	750-000613	S/N AA0330	1x OC-12 SONET, SMIR
FPC 4	REV 01	710-000175	S/N AA0050	

SSRAM	REV 01	710-000077	S/N AA2327	1 MB
SDRAM bank 0	REV 01	710-000099	S/N AA2329	64 MB
SDRAM bank 1	REV 01	710-000099	S/N AA2328	64 MB
PIC 0	REV 04	750-000613	S/N AA0320	1x OC-12 SONET, SMIR
PIC 2	REV 05	750-000616	S/N AA1341	1x OC-12 ATM, MM
PIC 3	REV 08	750-001072	S/N AB2462	1x G/E, 1000 BASE-SX
FPC 5	REV 10	710-000175	S/N AA7663	
SSRAM	REV 01	710-000077	S/N 501590	1 MB
SDRAM bank 0	REV 01	710-000099	S/N 300949	64 MB
SDRAM bank 1	REV 01	710-000099	S/N 300868	64 MB
PIC 1	REV 01	750-001323	S/N AB1670	1x Tunnel

show chassis hardware (M40e Router)

```
user@host> show chassis hardware
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis				m40e
Midplane	REV 01	710-005071	AX3671	
FPM CMB	REV 03	710-001642	AR9074	
FPM Display	REV 03	710-001647	AR7331	
CIP	REV 04	710-002649	BB4449	
PEM 0	Rev 01	740-003787	MC12364	Power Entry Module
PEM 1	Rev 01	740-003787	MC12383	Power Entry Module
PCG 0	REV 07	710-001568	AG1332	
PCG 1	REV 07	710-001568	AR3789	
Host 0			3e000007c8176601	Present
MCS 0	REV 11	710-001226	AN5813	
SFM 0 SPP	REV 07	710-001228	AG4676	
SFM 0 SPR	REV 05	710-002189	AE4735	Internet Processor II
SFM 1 SPP	REV 07	710-001228	AP1347	
SFM 1 SPR	REV 05	710-002189	BE0063	Internet Processor II
FPC 0	REV 01	710-011725	BE0669	M40e-EP-FPC Type 1
CPU	REV 01	710-004600	BD9504	
PIC 0	REV 03	750-003737	AY3991	4x G/E, 1000 BASE-SX
FPC 1	REV 01	710-005197	BD9842	M40e-FPC Type 2
CPU	REV 01	710-004600	BB4869	
PIC 0	REV 07	750-001900	AR8278	1x OC-48 SONET, SMSR
FPC 2	REV 02	710-005197	BD9824	M40e-FPC Type 2
CPU	REV 01	710-004600	BD9531	
PIC 0	REV 03	750-003737	AY3986	4x G/E, 1000 BASE-SX
FPC 4	REV 02	710-005078	BE0664	M40e-FPC Type 1
CPU	REV 01	710-004600	BD9559	
PIC 0	REV 03	750-001894	AG7963	1x G/E, 1000 BASE-SX
PIC 2	REV 01	750-002575	AF2472	4x OC-3 SONET, SMIR
FPC 6	REV 02	710-005078	BE0652	M40e-FPC Type 1
CPU	REV 01	710-004600	BD9607	
PIC 0	REV 02	750-002911	AN2286	4x F/E, 100 BASE-TX
PIC 2	REV 01	750-002577	AP6345	4x OC-3 SONET, MM

show chassis hardware (M120 Router)

```
user@host> show chassis hardware
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN000054AC	M120
Midplane	REV 01	710-013667	RB4170	M120 Midplane
FPM Board	REV 02	710-011407	CJ9186	M120 FPM Board
FPM Display	REV 02	710-011405	CJ9173	M120 FPM Display
FPM CIP	REV 02	710-011410	CJ9221	M120 FPM CIP
PEM 0	Rev 05	740-011936	RM28320	AC Power Entry Module

PEM 1	Rev 05	740-011936	RM28321	AC Power Entry Module
Routing Engine 0	REV 03	740-014080	1000642883	RE-A-1000
CB 0	REV 03	710-011403	CM8346	M120 Control Board
CB 1	REV 06	710-011403	CP6728	M120 Control Board
FPC 1	REV 02	710-015908	CP6925	M120 CFPC 10GE
PIC 0		BUILTIN	BUILTIN	1x 10GE(LAN/WAN) XFP
Xcvr 0	REV 01	740-014279	62E204N00007	XFP-10G-LR
FPC 3	REV 03	710-011393	CJ9234	M120 FPC Type 2
PIC 0	REV 16	750-008155	NB5229	2x G/E IQ, 1000 BASE
Xcvr 0	REV 01	740-011613	P9F15JB	SFP-SX
Xcvr 1	REV 01	740-007326	P4Q0R9G	SFP-SX
PIC 1	REV 09	750-007745	CG4360	4x OC-3 SONET, SMIR
PIC 2	REV 16	750-008155	ND7787	2x G/E IQ, 1000 BASE
Xcvr 0	REV 01	740-011613	P9F12AS	SFP-SX
Xcvr 1	REV 01	740-011613	P9F1ALU	SFP-SX
PIC 3	REV 07	750-011800	JW1284	8x 1GE(LAN), IQ2
Xcvr 0	REV 01	740-011613	P9F1AM6	SFP-SX
Xcvr 6	REV 01	740-011613	P9F16NN	SFP-SX
Xcvr 7	REV 01	740-011782	P8C29Y7	SFP-SX
Board B	REV 02	710-011395	CN3754	M120 FPC Mezz
FPC 4	REV 02	710-011398	CP6741	M120 FPC Type 3
PIC 0	REV 16	750-007141	NB2855	10x 1GE(LAN), 1000 BASE
Xcvr 0	REV 01	740-011782	P922A1F	SFP-SX
Xcvr 1	REV 01	740-011782	P922A16	SFP-SX
Xcvr 2	REV 01	740-011782	P922A0U	SFP-SX
Xcvr 3	REV 01	740-011782	P9229UZ	SFP-SX
Xcvr 4	REV 01	740-009029	P11JXWP	SFP-LX
Xcvr 6	REV 01	740-011613	P9F1ALW	SFP-SX
FPC 5	REV 01	710-011388	CJ9088	M120 FPC Type 1
PIC 0	*** Hardware Not Supported ***			
PIC 1	REV 05	750-012052	NB0410	1x CHOC3 IQ SONET, SMLR
PIC 2	REV 01	750-013167	CM3824	4x CHDS3 IQ
PIC 3	REV 01	750-010240	CB5366	1x G/E SFP, 1000 BASE
Board B	REV 01	710-011390	CJ9103	M120 FPC Mezz Board
FEB 3	REV 04	710-011663	CP6673	M120 FEB
FEB 4	REV 04	710-011663	CJ9368	M120 FEB
FEB 5	REV 04	710-011663	CJ9386	M120 FEB
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Top Fan Tray
Fan Tray 3				Rear Bottom Fan Tray

show chassis hardware detail (M120 Router)

user@host> show chassis hardware detail

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN000054AC	M120
Midplane	REV 01	710-013667	RB4170	M120 Midplane
FPM Board	REV 02	710-011407	CJ9186	M120 FPM Board
FPM Display	REV 02	710-011405	CJ9173	M120 FPM Display
FPM CIP	REV 02	710-011410	CJ9221	M120 FPM CIP
PEM 0	Rev 05	740-011936	RM28320	AC Power Entry Module
PEM 1	Rev 05	740-011936	RM28321	AC Power Entry Module
Routing Engine 0	REV 03	740-014080	1000642883	RE-A-1000
ad0	248 MB	SILICONSYSTEMS INC	256M 126CT505S0763SC00110	Compact Flash
ad2	38154 MB	HTE541040G9SA00	MPBBTOX2HS2E3M	Hard Disk
CB 0	REV 03	710-011403	CM8346	M120 Control Board
CB 1	REV 06	710-011403	CP6728	M120 Control Board

FPC 1	REV 02	710-015908	CP6925	M120 CFPC 10GE
PIC 0		BUILTIN	BUILTIN	1x 10GE(LAN/WAN) XFP
Xcvr 0	REV 01	740-014279	62E204N00007	XFP-10G-LR
FPC 3	REV 03	710-011393	CJ9234	M120 FPC Type 2
PIC 0	REV 16	750-008155	NB5229	2x G/E IQ, 1000 BASE
Xcvr 0	REV 01	740-011613	P9F15JB	SFP-SX
Xcvr 1	REV 01	740-007326	P4Q0R9G	SFP-SX
PIC 1	REV 09	750-007745	CG4360	4x OC-3 SONET, SMIR
PIC 2	REV 16	750-008155	ND7787	2x G/E IQ, 1000 BASE
Xcvr 0	REV 01	740-011613	P9F12AS	SFP-SX
Xcvr 1	REV 01	740-011613	P9F1ALU	SFP-SX
PIC 3	REV 07	750-011800	JW1284	8x 1GE(LAN), IQ2
Xcvr 0	REV 01	740-011613	P9F1AM6	SFP-SX
Xcvr 6	REV 01	740-011613	P9F16NN	SFP-SX
Xcvr 7	REV 01	740-011782	P8C29Y7	SFP-SX
Board B	REV 02	710-011395	CN3754	M120 FPC Mezz
FPC 4	REV 02	710-011398	CP6741	M120 FPC Type 3
PIC 0	REV 16	750-007141	NB2855	10x 1GE(LAN), 1000 BASE
Xcvr 0	REV 01	740-011782	P922A1F	SFP-SX
Xcvr 1	REV 01	740-011782	P922A16	SFP-SX
Xcvr 2	REV 01	740-011782	P922A0U	SFP-SX
Xcvr 3	REV 01	740-011782	P9229UZ	SFP-SX
Xcvr 4	REV 01	740-009029	P11JXWP	SFP-LX
Xcvr 6	REV 01	740-011613	P9F1ALW	SFP-SX
FPC 5	REV 01	710-011388	CJ9088	M120 FPC Type 1
PIC 0	*** Hardware Not Supported ***			
PIC 1	REV 05	750-012052	NB0410	1x CHOC3 IQ SONET, SMLR
PIC 2	REV 01	750-013167	CM3824	4x CHDS3 IQ
PIC 3	REV 01	750-010240	CB5366	1x G/E SFP, 1000 BASE
Board B	REV 01	710-011390	CJ9103	M120 FPC Mezz Board
FEB 3	REV 04	710-011663	CP6673	M120 FEB
FEB 4	REV 04	710-011663	CJ9368	M120 FEB
FEB 5	REV 04	710-011663	CJ9386	M120 FEB
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Top Fan Tray
Fan Tray 3				Rear Bottom Fan Tray

show chassis hardware models (M120 Router)

```

user@host> show chassis hardware models
Hardware inventory:
Item          Version  Part number  CLEI code  FRU model number
Midplane      REV 01   710-013667
FPM CIP       REV 02   710-011410  CRAFT-M120-S
PEM 0         Rev 05   740-011936  PWR-M120-AC-S
PEM 1         Rev 05   740-011936  PWR-M120-AC-S
Routing Engine 0 REV 03   740-014080  RE-A-1000-2048-S
CB 0          REV 03   710-011403  CB-M120-S
CB 1          REV 06   710-011403  CB-M120-S
FPC 1         REV 02   710-015908  M120-cFPC-1XGE-XFP
FPC 3
PIC 0         REV 16   750-008155  PB-2GE-SFP-QPP
PIC 1         REV 09   750-007745  PC-40C3-SON-SMIR
PIC 2         REV 16   750-008155  PB-2GE-SFP-QPP
PIC 3         REV 07   750-011800  PB-8GE-TYPE2-SFP-IQ2
FPC 4
PIC 0         REV 16   750-007141  PC-10GE-SFP
FPC 5

```

PIC 1	REV 05	750-012052	PB-1CHOC3-SMIR-QPP
PIC 2	REV 01	750-013167	PE-4CHDS3-QPP
PIC 3	REV 01	750-010240	PB-1GE-SFP
Fan Tray 0			FFANTRAY-M120-S
Fan Tray 1			FFANTRAY-M120-S
Fan Tray 2			RFANTRAY-M120-S
Fan Tray 3			RFANTRAY-M120-S

show chassis hardware (M160 Router)

```

user@host> show chassis hardware

```

Item	Version	Part number	Serial number	Description
Chassis			101	M160
Midplane	REV 02	710-001245	S/N AB4107	
FPM CMB	REV 01	710-001642	S/N AA2911	
FPM Display	REV 01	710-001647	S/N AA2999	
CIP	REV 02	710-001593	S/N AA9563	
PEM 0	Rev 01	740-001243	S/N KJ35769	DC
PEM 1	Rev 01	740-001243	S/N KJ35765	DC
PCG 0	REV 01	710-001568	S/N AA9794	
PCG 1	REV 01	710-001568	S/N AA9804	
Host 1			da000004f8d57001	teknor
MCS 1	REV 03	710-001226	S/N AA9777	
SFM 0 SPP	REV 04	710-001228	S/N AA2975	
SFM 0 SPR	REV 02	710-001224	S/N AA9838	Internet Processor I
SFM 1 SPP	REV 04	710-001228	S/N AA2860	
SFM 1 SPR	REV 01	710-001224	S/N AB0139	Internet Processor I
FPC 0	REV 03	710-001255	S/N AA9806	FPC Type 1
CPU	REV 02	710-001217	S/N AA9590	
PIC 1	REV 05	750-000616	S/N AA1527	1x OC-12 ATM, MM
PIC 2	REV 05	750-000616	S/N AA1535	1x OC-12 ATM, MM
PIC 3	REV 01	750-000616	S/N AA1519	1x OC-12 ATM, MM
FPC 1	REV 02	710-001611	S/N AA9523	FPC Type 2
CPU	REV 02	710-001217	S/N AA9571	
PIC 0	REV 03	750-001900	S/N AA9626	1x STM-16 SDH, SMIR
PIC 1	REV 01	710-002381	S/N AD3633	2x G/E, 1000 BASE-SX
FPC 2				FPC Type OC192
CPU	REV 03	710-001217	S/N AB3329	
PIC 0	REV 01			1x OC-192 SM SR-2
Fan Tray 0				Rear Bottom Blower
Fan Tray 1				Rear Top Blower
Fan Tray 2				Front Top Blower
Fan Tray 3				Front Fan Tray

show chassis hardware models (M160 Router)

```

user@host> show chassis hardware models

```

Hardware inventory:				
Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 03	710-009120		CHAS-BP-M320-S
FPM Display	REV 02	710-009351		CRAFT-M320-S
CIP	REV 03	710-005926		CIP-M320-S
PEM 2	Rev X4	740-009148		PWR-M-DC-S
PEM 3	Rev X4	740-009148		PWR-M-DC-S
Routing Engine 0	REV 02	740-008883		RE-1600-2048-S
Routing Engine 1	REV 02	740-008883		RE-1600-2048-S
FPC 0	REV 02	710-010419		M320-FPC1
PIC 0	REV 01	750-001323		P-TUNNEL
PIC 1	REV 02	750-002987		PE-10C12-SON-SMIR
PIC 2	REV 04	750-001894		PB-1GE-SX
PIC 3	REV 04	750-001896		PB-10C12-SON-SMIR

FPC 1	REV 02	710-010419	M320-FPC1
PIC 0	REV 04	750-001894	PB-1GE-SX
PIC 1	REV 04	750-001894	PB-1GE-SX
PIC 3	REV 03	750-001894	PB-1GE-SX
FPC 2	REV 02	710-010419	M320-FPC1
PIC 0	REV 10	750-005634	PB-1CHOC12SMIR-QPP
PIC 1	REV 10	750-005634	PB-1CHOC12SMIR-QPP
PIC 2	REV 07	750-005634	PB-1CHOC12SMIR-QPP
PIC 3	REV 07	750-005634	PB-1CHOC12SMIR-QPP
PIC 1	REV 10	750-005634	PB-1CHOC12SMIR-QPP
PIC 2	REV 07	750-005634	PB-1CHOC12SMIR-QPP
PIC 3	REV 07	750-005634	PB-1CHOC12SMIR-QPP
FPC 3			
PIC 0	REV 03	750-001895	PB-10C12-SON-MM
PIC 1	REV 04	750-001894	PB-1GE-SX
PIC 3	REV 04	750-003141	PB-1GE-SX-B
FPC 4	REV 02	710-010419	M320-FPC1
FPC 5	REV 02	710-010419	M320-FPC1
FPC 6	REV 02	710-010419	M320-FPC1
FPC 7			
PIC 0	REV 15	750-001901	PB-40C12-SON-SMIR
PIC 1	REV 06	750-001900	PB-10C48-SON-SMSR
PIC 2	REV 07	750-001900	PB-10C48-SON-SMSR
PIC 3	REV 05	750-003737	PB-4GE-SX
SIB 0	REV 03	710-009184	SIB-M-S
SIB 1	REV 03	710-009184	SIB-M-S
SIB 2	REV 03	710-009184	SIB-M-S
SIB 3	REV 03	710-009184	SIB-M-S
Fan Tray 0			FFANTRAY-M320-S
Fan Tray 1			FFANTRAY-M320-S
Fan Tray 2			RFANTRAY-M320-S

show chassis hardware detail (M160 Router)

```

user@host> show chassis hardware detail
Hardware inventory:

```

Item	Version	Part number	Serial number	Description
Chassis			101	M160
Midplane	REV 02	710-001245	S/N AB4107	
FPM CMB	REV 01	710-001642	S/N AA2911	
FPM Display	REV 01	710-001647	S/N AA2999	
CIP	REV 02	710-001593	S/N AA9563	
PEM 0	Rev 01	740-001243	S/N KJ35769	DC
PEM 1	Rev 01	740-001243	S/N KJ35765	DC
PCG 0	REV 01	710-001568	S/N AA9794	
PCG 1	REV 01	710-001568	S/N AA9804	
Host 1			da000004f8d57001	teknor
MCS 1	REV 03	710-001226	S/N AA9777	
SFM 0 SPP	REV 04	710-001228	S/N AA2975	
SFM 0 SPR	REV 02	710-001224	S/N AA9838	Internet Processor I
SSRAM bank 0	REV 01	710-000077	S/N 306456	1 MB
SSRAM bank 1	REV 01	710-000077	S/N 306474	1 MB
SSRAM bank 2	REV 01	710-000077	S/N 306388	1 MB
SSRAM bank 3	REV 01	710-000077	S/N 306392	1 MB
SFM 1 SPP	REV 04	710-001228	S/N AA2860	
SFM 1 SPR	REV 01	710-001224	S/N AB0139	Internet Processor I
SSRAM bank 0	REV 01	710-000077	S/N 302917	1 MB
SSRAM bank 1	REV 01	710-000077	S/N 302662	1 MB
SSRAM bank 2	REV 01	710-000077	S/N 302593	1 MB
SSRAM bank 3	REV 01	710-000077	S/N 100160	1 MB
FPC 0	REV 03	710-001255	S/N AA9806	FPC Type 1

CPU	REV 02	710-001217	S/N AA9590	
SSRAM	REV 01	710-000077	S/N 302836	1 MB
SDRAM 0	REV 01	710-001196	S00141	32 MB
SDRAM 1	REV 01	710-001196	S0010;	32 MB
SSRAM	REV 01	710-000077	S/N 302633	1 MB
SDRAM 0	REV 01	710-001196	S00143	32 MB
SDRAM 1	REV 01	710-001196	S00115	32 MB
SSRAM	REV 01	710-000077	S/N 302952	1 MB
SDRAM 0	REV 01	710-001196	S00135	32 MB
SDRAM 1	REV 01	710-001196	S001=3	32 MB
SSRAM	REV 01	710-000077	S/N 302892	1 MB
SDRAM 0	REV 01	710-001196	S00076	32 MB
SDRAM 1	REV 01	710-001196	S001=5	32 MB
PIC 1	REV 05	750-000616	S/N AA1527	1x OC-12 ATM, MM
PIC 2	REV 05	750-000616	S/N AA1535	1x OC-12 ATM, MM
PIC 3	REV 01	750-000616	S/N AA1519	1x OC-12 ATM, MM
FPC 1	REV 02	710-001611	S/N AA9523	FPC Type 2
CPU	REV 02	710-001217	S/N AA9571	
SSRAM	REV 01	710-000077	S/N 306340	1 MB
SDRAM 0	REV 01	710-001196	S00012	32 MB
SDRAM 1	REV 01	710-001196	S0001?	32 MB
SSRAM	REV 01	710-000077	S/N 306454	1 MB
SDRAM 0	REV 01	710-001196	S00028	32 MB
SDRAM 1	REV 01	710-001196	S0002?	32 MB
SSRAM	REV 01	710-000077	S/N 306492	1 MB
SDRAM 0	REV 01	710-001196	S00015	32 MB
SDRAM 1	REV 01	710-001196	S00031	32 MB
SSRAM	REV 01	710-000077	S/N 306363	1 MB
SDRAM 0	REV 01	710-001196	S00013	32 MB
SDRAM 1	REV 01	710-001196	S00032	32 MB
PIC 0	REV 03	750-001900	S/N AA9626	1x STM-16 SDH, SMIR
PIC 1	REV 01	710-002381	S/N AD3633	2x G/E, 1000 BASE-SX
FPC 2				FPC Type OC192
... SSRAM	REV 01	710-000077	S/N 306466	1 MB

show chassis hardware (M320 Router)

user@host> show chassis hardware

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			67245	M320
Midplane	REV 05	710-009120	RB1202	M320 Midplane
FPM GBUS	REV 04	710-005928	HZ5697	M320 Board
FPM Display	REV 05	710-009351	HR1464	M320 FPM Display
CIP	REV 04	710-005926	HT8672	M320 CIP
PEM 0	Rev 05	740-009148	QK34208	DC Power Entry Module
PEM 1	Rev 05	740-009148	QK34262	DC Power Entry Module
PEM 2	Rev 05	740-009148	QF10449	DC Power Entry Module
PEM 3	Rev 05	740-009148	QJ18257	DC Power Entry Module
Routing Engine 0	REV 06	740-008883	P11123901185	RE-4.0
CB 0	REV 07	710-009115	JB2382	M320 Control Board
FPC 0	REV 02	710-005017	CD9926	M320 FPC Type 2
CPU	REV 01	710-011659	CJ6940	M320 PCA SCPU
PIC 0	REV 07	750-001900	AT1594	1x OC-48 SONET, SMSR
PIC 1	REV 03	750-001850	HS2746	1x Tunnel
PIC 2	REV 05	750-010618	JE7117	4x G/E SFP, 1000 BASE
PIC 3	REV 06	750-001900	HE6083	1x OC-48 SONET, SMSR
FPC 2	REV 02	710-005017	CH0319	M320 FPC Type 1
CPU	REV 01	710-011659	CJ6942	M320 PCA SCPU
PIC 0	REV 05	750-003034	BD8705	4x OC-3 SONET, SMIR
FPC 5	REV 02	710-005017	CD9938	M320 FPC Type 2

CPU				
FPC 7	REV 02	710-005017	CD9934	M320 FPC Type 2
CPU				
SIB 0	REV 09	710-009184	JA6540	M320 SIB
SIB 1	REV 09	710-009184	HV9511	M320 SIB
SIB 2	REV 09	710-009184	HW2057	M320 SIB
SIB 3	REV 09	710-009184	JA6687	M320 SIB
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Fan Tray

show chassis hardware models (M320 Router)

```
user@host> show chassis hardware models
Hardware inventory:
```

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 03	710-009120		CHAS-BP-M320-S
FPM Display	REV 02	710-009351		CRAFT-M320-S
CIP	REV 03	710-005926		CIP-M320-S
PEM 2	Rev X4	740-009148		PWR-M-DC-S
PEM 3	Rev X4	740-009148		PWR-M-DC-S
Routing Engine 0	REV 02	740-008883		RE-1600-2048-S
Routing Engine 1	REV 02	740-008883		RE-1600-2048-S
FPC 0	REV 02	710-010419		M320-FPC1
PIC 0	REV 01	750-001323		P-TUNNEL
PIC 1	REV 02	750-002987		PE-10C12-SON-SMIR
PIC 2	REV 04	750-001894		PB-1GE-SX
PIC 3	REV 04	750-001896		PB-10C12-SON-SMIR
FPC 1	REV 02	710-010419		M320-FPC1
PIC 0	REV 04	750-001894		PB-1GE-SX
PIC 1	REV 04	750-001894		PB-1GE-SX
PIC 3	REV 03	750-001894		PB-1GE-SX
FPC 2	REV 02	710-010419		M320-FPC1
PIC 0	REV 10	750-005634		PB-1CHOC12SMIR-QPP
PIC 1	REV 10	750-005634		PB-1CHOC12SMIR-QPP
PIC 2	REV 07	750-005634		PB-1CHOC12SMIR-QPP
PIC 3	REV 07	750-005634		PB-1CHOC12SMIR-QPP
PIC 1	REV 10	750-005634		PB-1CHOC12SMIR-QPP
PIC 2	REV 07	750-005634		PB-1CHOC12SMIR-QPP
PIC 3	REV 07	750-005634		PB-1CHOC12SMIR-QPP
FPC 3				
PIC 0	REV 03	750-001895		PB-10C12-SON-MM
PIC 1	REV 04	750-001894		PB-1GE-SX
PIC 3	REV 04	750-003141		PB-1GE-SX-B
FPC 4	REV 02	710-010419		M320-FPC1
FPC 5	REV 02	710-010419		M320-FPC1
FPC 6	REV 02	710-010419		M320-FPC1
FPC 7				
PIC 0	REV 15	750-001901		PB-40C12-SON-SMIR
PIC 1	REV 06	750-001900		PB-10C48-SON-SMSR
PIC 2	REV 07	750-001900		PB-10C48-SON-SMSR
PIC 3	REV 05	750-003737		PB-4GE-SX
SIB 0	REV 03	710-009184		SIB-M-S
SIB 1	REV 03	710-009184		SIB-M-S
SIB 2	REV 03	710-009184		SIB-M-S
SIB 3	REV 03	710-009184		SIB-M-S
Fan Tray 0				FFANTRAY-M320-S
Fan Tray 1				FFANTRAY-M320-S
Fan Tray 2				RFANTRAY-M320-S

show chassis hardware (MX5 Router)

```

user@host> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               E1368         MX5-T
Midplane      REV 01   711-038215   YF5288        MX5-T
PEM 0         Rev 04   740-028288   VA01215       AC Power Entry Module
PEM 1         Rev 04   740-028288   VA01218       AC Power Entry Module
Routing Engine BUILTIN   BUILTIN      Routing Engine
TFEB 0        BUILTIN   BUILTIN      Forwarding Engine
Processor
  QXM 0       REV 05   711-028408   ZA9136        MPC QXM
  FPC 0       BUILTIN   BUILTIN      MPC BUILTIN
  MIC 0       BUILTIN   BUILTIN      4x 10GE XFP
  PIC 0       BUILTIN   BUILTIN      4x 10GE XFP
  FPC 1       BUILTIN   BUILTIN      MPC BUILTIN
  MIC 0       REV 24   750-028392   YX9820        3D 20x 1GE(LAN) SFP
  PIC 0       BUILTIN   BUILTIN      10x 1GE(LAN) SFP
    Xcvr 0    REV 01   740-031851   AM1045SUAQ3   SFP-SX
    Xcvr 1    REV 01   740-031851   AM1045SUAPA   SFP-SX
    Xcvr 2    REV 01   740-031851   AM1045SUAN7   SFP-SX
    Xcvr 3    REV 01   740-031851   AM1045SU91Q   SFP-SX
    Xcvr 4    REV 01   740-031851   AM1045SUDDR   SFP-SX
    Xcvr 9    REV 01   740-011613   AM0848SB6A1   SFP-SX
  PIC 1       BUILTIN   BUILTIN      10x 1GE(LAN) SFP
    Xcvr 0    REV 01   740-031851   AM1045SUANO   SFP-SX
    Xcvr 1    REV 01   740-011613   AS0812S0719   SFP-SX
    Xcvr 2    REV 01   740-011613   AM0821SA121   SFP-SX
    Xcvr 3    REV 01   740-011613   PF21K21       SFP-SX
    Xcvr 4    REV 01   740-011613   AM0848SB69Z   SFP-SX
    Xcvr 5    REV 01   740-011782   P9P0XV3       SFP-SX
    Xcvr 6    REV 01   740-011613   AM0812S8WJN   SFP-SX
    Xcvr 7    REV 01   740-011613   PAM3G9Q       SFP-SX
    Xcvr 8    REV 01   740-011613   AM0848SB4A6   SFP-SX
    Xcvr 9    REV 01   740-011782   P9MOU37       SFP-SX
  MIC 1       REV 20   750-028380   ZG2657        3D 2x 10GE XFP
  PIC 2       BUILTIN   BUILTIN      1x 10GE XFP
  PIC 3       BUILTIN   BUILTIN      1x 10GE XFP
Fan Tray                               Fan Tray

```

show chassis hardware (MX10 Router)

```

user@host> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               E1372         MX10-T
Midplane      REV 01   711-038211   YF5285        MX10-T
PEM 0         Rev 04   740-028288   VB01678       AC Power Entry Module
Routing Engine BUILTIN   BUILTIN      Routing Engine
TFEB 0        BUILTIN   BUILTIN      Forwarding Engine
Processor
  QXM 0       REV 05   711-028408   ZA9053        MPC QXM
  FPC 0       BUILTIN   BUILTIN      MPC BUILTIN
  MIC 0       BUILTIN   BUILTIN      4x 10GE XFP
  PIC 0       BUILTIN   BUILTIN      4x 10GE XFP
  FPC 1       BUILTIN   BUILTIN      MPC BUILTIN
  MIC 0       REV 24   750-028392   YX9436        3D 20x 1GE(LAN) SFP
  PIC 0       BUILTIN   BUILTIN      10x 1GE(LAN) SFP
    Xcvr 0    REV 01   740-031851   AM1107SUFQW   SFP-SX

```


PIC 1	BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Fan Tray			Fan Tray

show chassis hardware (MX40 Router)

```

user@host> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               E1367         MX40-T
Midplane      REV 01   711-038211   YF5284        MX40-T
PEM 0         Rev 04   740-028288   VB01680       AC Power Entry Module
PEM 1         Rev 04   740-028288   VB01700       AC Power Entry Module
Routing Engine BUILTIN   BUILTIN      Routing Engine
TFEB 0        BUILTIN   BUILTIN      Forwarding Engine
Processor
  QXM 0       REV 05   711-028408   ZA9048        MPC QXM
  FPC 0       BUILTIN   BUILTIN      MPC BUILTIN
    MIC 0     BUILTIN   BUILTIN      4x 10GE XFP
      PIC 0   BUILTIN   BUILTIN      4x 10GE XFP
        Xcvr 0 REV 01   740-014279   M7067UPP      XFP-10G-LR
        Xcvr 1     NON-JNPR K9J02UN      XFP-10G-LR
  FPC 1       BUILTIN   BUILTIN      MPC BUILTIN
    MIC 0     REV 24   750-028392   YX3504        3D 20x 1GE(LAN) SFP
      PIC 0   BUILTIN   BUILTIN      10x 1GE(LAN) SFP
        Xcvr 0 REV 01   740-011613   AM0812S8WTE   SFP-SX
        Xcvr 1 REV 01   740-011613   PFA6KV2       SFP-SX
        Xcvr 2 REV 01   740-031851   AM1045SUDDM   SFP-SX
        Xcvr 3 REV 01   740-011613   PD63C7M       SFP-SX
        Xcvr 4 REV 01   740-011613   PD63DJY       SFP-SX
        Xcvr 5 REV 02   740-011613   AA0950STLL9   SFP-SX
        Xcvr 6 REV 01   740-011782   PAR1YHC       SFP-SX
        Xcvr 7 REV 01   740-011782   P9P0XXL       SFP-SX
        Xcvr 8 REV 01   740-011613   PD63D95       SFP-SX
        Xcvr 9 REV 01   740-031851   AM1045SU9B8   SFP-SX
      PIC 1   BUILTIN   BUILTIN      10x 1GE(LAN) SFP
        Xcvr 0 REV 01   740-011613   PF21L3Z       SFP-SX
        Xcvr 1 REV 01   740-031851   AM1045SU7M9   SFP-SX
        Xcvr 2 REV 01   740-031851   AM1045SUAPT   SFP-SX
        Xcvr 3 REV 01   740-011613   PFF2BZH       SFP-SX
        Xcvr 4 REV 01   740-031851   AM1045SUDDN   SFP-SX
        Xcvr 5 REV 01   740-031851   AM1039S00ZR   SFP-SX
        Xcvr 6 REV 01   740-031851   AM1045SUD6Y   SFP-SX
        Xcvr 8 REV 01   740-011613   PFM1QBS       SFP-SX
        Xcvr 9 REV 01   740-011613   PFF2E25       SFP-SX
    MIC 1     REV 01   750-021130   KG4391        3D 2x 10GE XFP
      PIC 2   BUILTIN   BUILTIN      1x 10GE XFP
        Xcvr 0 REV 01   740-011571   C645XJ04G     XFP-10G-SR
      PIC 3   BUILTIN   BUILTIN      1x 10GE XFP
        Xcvr 0     NON-JNPR CA49BK0AE     XFP-10G-SR
Fan Tray

```

show chassis hardware (Fixed MX80 Router)

```

user@host> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               MX80-48T
Midplane      REV 01   711-031603   KF9250        MX80-48T
Routing Engine BUILTIN   BUILTIN      Routing Engine
FEB 0         BUILTIN   BUILTIN      Forwarding Engine Board
FPC 0         BUILTIN   BUILTIN      MPC BUILTIN

```

MIC 0		BUILTIN	BUILTIN	4x 10GE XFP
PIC 0		BUILTIN	BUILTIN	4x 10GE XFP
Xcvr 0		NON-JNPR	M6439D41	XFP-10G-LR
Xcvr 1	REV 01	740-014279	6XE931N00202	XFP-10G-LR
Xcvr 2	REV 01	740-014289	C715XU05F	XFP-10G-SR
Xcvr 3	REV 01	740-014289	C650XU0EP	XFP-10G-SR
FPC 1		BUILTIN	BUILTIN	MPC BUILTIN
MIC 0	REV 01	711-029399	JR6981	12x 1GE(LAN) RJ45
PIC 0		BUILTIN	BUILTIN	12x 1GE(LAN) RJ45
PIC 1		BUILTIN	BUILTIN	12x 1GE(LAN) RJ45
MIC 1	REV 01	BUILTIN	BUILTIN	12x 1GE(LAN) RJ45
PIC 2		BUILTIN	BUILTIN	12x 1GE(LAN) RJ45
PIC 3		BUILTIN	BUILTIN	12x 1GE(LAN) RJ45
Fan Tray				Fan Tray

show chassis hardware (Modular MX80 Router)

```
user@host> show chassis hardware
```

Hardware inventory:				
Item	Version	Part number	Serial number	Description
Chassis				MX80
Midplane	REV 02	711-031594	JR7084	MX80
PEM 0	Rev 01	740-028288	000018	AC Power Entry Module
Routing Engine		BUILTIN	BUILTIN	Routing Engine
FEB 0		BUILTIN	BUILTIN	Forwarding Engine Board
QXM 0	REV 05	711-028408	JR7041	MPC QXM
FPC 0		BUILTIN	BUILTIN	MPC BUILTIN
MIC 0		BUILTIN	BUILTIN	4x 10GE XFP
PIC 0		BUILTIN	BUILTIN	4x 10GE XFP
FPC 1		BUILTIN	BUILTIN	MPC BUILTIN
MIC 0	REV 02	750-028380	JR6598	3D 2x 10GE XFP
PIC 0		BUILTIN	BUILTIN	1x 10GE XFP
Xcvr 0	REV 01	740-014289	T07M86365	XFP-10G-SR
PIC 1		BUILTIN	BUILTIN	1x 10GE XFP
Xcvr 0	REV 01	740-014289	T07M71094	XFP-10G-SR
MIC 1	REV 02	750-028380	JG8548	3D 2x 10GE XFP
PIC 2		BUILTIN	BUILTIN	1x 10GE XFP
Xcvr 0	REV 02	740-014289	T08L86302	XFP-10G-SR
PIC 3		BUILTIN	BUILTIN	1x 10GE XFP
Xcvr 0	REV 02	740-014289	C810XU0BA	XFP-10G-SR
Fan Tray				Fan Tray

show chassis hardware (MX104 Router)

```
user@host> show chassis hardware
```

Hardware inventory:				
Item	Version	Part number	Serial number	Description
Chassis			G3503	MX104
Midplane	REV 28	750-044219	CAAX5741	MX104
PEM 0	REV 03	740-045933	1H072500016	AC Power Entry Module
PEM 1	REV 03	740-045932	1H073050017	DC Power Entry Module
Routing Engine 0	REV 20	750-044228	CAAY7935	RE-MX-104
Routing Engine 1	REV 13	750-044228	CAAM6380	RE-MX-104
AFEB 0		BUILTIN	BUILTIN	Forwarding Engine
Processor				
FPC 0		BUILTIN	BUILTIN	MPC BUILTIN
FPC 1		BUILTIN	BUILTIN	MPC BUILTIN
MIC 0	REV 15	750-036132	CAAF7948	2xOC12/8xOC3 CC-CE
PIC 0		BUILTIN	BUILTIN	2xOC12/8xOC3 CC-CE

Xcvr 0	REV 01	740-011615	PCQ0U2J	SFP-IR
Xcvr 1	REV 01	740-016068	PJL7A6G	SFP-SR
Xcvr 2	REV 01	740-016068	PJL7A5J	SFP-SR
Xcvr 3	REV 01	740-016065	PJN5HPZ	SFP-SR
Xcvr 4	REV 01	740-029122	PKB38TL	SFP-LR
Xcvr 5	REV 01	740-011787	P6A107G	SFP-LR
Xcvr 6	REV 01	740-029122	PKB38TR	SFP-LR
Xcvr 7	REV 01	740-011787	PBKONK3	SFP-LR
MIC 1				
FPC 2		BUILTIN	BUILTIN	MPC BUILTIN
MIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B10F00465	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B10F00461	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B10G01545	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B10G01385	SFP+-10G-SR
Fan Tray 0	REV 02	711-049570	CAAX6538	Fan Tray

show chassis hardware detail (MX104 Router)

```

user@host> show chassis hardware detail
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               G3503         MX104
Midplane      REV 28   750-044219   CAAX5741      MX104
PEM 0         REV 03   740-045933   1H072500016   AC Power Entry Module
PEM 1         REV 03   740-045932   1H073050017   DC Power Entry Module
Routing Engine 0 REV 20   750-044228   CAAY7935      RE-MX-104
  da0         7836 MB  ATP IG eUSB SSD      Nand Flash 0
  usb0 (addr 1) EHCI root hub 0    Freescale     uhub0
  usb0 (addr 2) USB2513Bi 9491    SMSC          uhub1
  usb0 (addr 3) ATP IG eUSB SSD 44801 ATP Electronics umass0
Routing Engine 1 REV 13   750-044228   CAAM6380      RE-MX-104
  da0         7836 MB  ATP IG eUSB SSD      Nand Flash 0
AFEB 0                               BUILTIN       BUILTIN       Forwarding Engine
Processor
FPC 0                               BUILTIN       BUILTIN       MPC BUILTIN
FPC 1                               BUILTIN       BUILTIN       MPC BUILTIN
MIC 0         REV 15   750-036132   CAAF7948      2x0C12/8x0C3 CC-CE
PIC 0                               BUILTIN       BUILTIN       2x0C12/8x0C3 CC-CE
  Xcvr 0      REV 01   740-011615   PCQ0U2J      SFP-IR
  Xcvr 1      REV 01   740-016068   PJL7A6G      SFP-SR
  Xcvr 2      REV 01   740-016068   PJL7A5J      SFP-SR
  Xcvr 3      REV 01   740-016065   PJN5HPZ      SFP-SR
  Xcvr 4      REV 01   740-029122   PKB38TL      SFP-LR
  Xcvr 5      REV 01   740-011787   P6A107G      SFP-LR
  Xcvr 6      REV 01   740-029122   PKB38TR      SFP-LR
  Xcvr 7      REV 01   740-011787   PBKONK3      SFP-LR
MIC 1
FPC 2                               BUILTIN       BUILTIN       MPC BUILTIN
MIC 0                               BUILTIN       BUILTIN       4x 10GE(LAN) SFP+
PIC 0                               BUILTIN       BUILTIN       4x 10GE(LAN) SFP+
  Xcvr 0      REV 01   740-031980   B10F00465    SFP+-10G-SR
  Xcvr 1      REV 01   740-031980   B10F00461    SFP+-10G-SR
  Xcvr 2      REV 01   740-031980   B10G01545    SFP+-10G-SR
  Xcvr 3      REV 01   740-031980   B10G01385    SFP+-10G-SR
Fan Tray 0    REV 02   711-049570   CAAX6538     Fan Tray

```

show chassis hardware extensive (MX104 Router)

```

user@host> show chassis hardware extensive

```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			G3503	MX104

Jedec Code:	0x7fb0	EEPROM Version:	0x02
		S/N:	G3503
Assembly ID:	0x0560	Assembly Version:	00.00
Date:	00-00-0000	Assembly Flags:	0x00

ID: MX104

Board Information Record:

Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

I2C Hex Data:

Address 0x00: 7f b0 02 ff 05 60 00 00 00 00 00 00 00 00 00 00
 Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x20: 47 33 35 30 33 00 00 00 00 00 00 00 00 00 00 00
 Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

Midplane	REV 28	750-044219	CAAX5741	MX104
----------	--------	------------	----------	-------

Jedec Code:	0x7fb0	EEPROM Version:	0x02
P/N:	750-044219	S/N:	CAAX5741
Assembly ID:	0x0560	Assembly Version:	01.28
Date:	03-27-2013	Assembly Flags:	0x00
Version:	REV 28	CLEI Code:	PROTOXCLEI
ID:	MX104	FRU Model Number:	PROTO-ASSEMBLY

Board Information Record:

Address 0x00: ad 01 08 00 b0 a8 6e a7 f8 00 ff ff ff ff ff ff

I2C Hex Data:

Address 0x00: 7f b0 02 ff 05 60 01 1c 52 45 56 20 32 38 00 00
 Address 0x10: 00 00 00 00 37 35 30 2d 30 34 34 32 31 39 00 00
 Address 0x20: 53 2f 4e 20 43 41 41 58 35 37 34 31 00 1b 03 07
 Address 0x30: dd ff ff ff ad 01 08 00 b0 a8 6e a7 f8 00 ff ff
 Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
 Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
 Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
 Address 0x70: ff ff ff c2 47 33 35 30 33 00 00 00 00 00 00 00

PEM 0	REV 03	740-045933	1H072500016	AC Power Entry Module
-------	--------	------------	-------------	-----------------------

Jedec Code:	0x7fb0	EEPROM Version:	0x02
P/N:	740-045933	S/N:	1H072500016
Assembly ID:	0x0475	Assembly Version:	00.03
Date:	12-14-2012	Assembly Flags:	0x00
Version:	REV 03	CLEI Code:	IPUPAJ9KAA
ID:	AC Power Entry Module	FRU Model Number:	PWR-AMX1100-AC-S

Board Information Record:

Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff 02 02 00 ff

I2C Hex Data:

Address 0x00: 7f b0 02 ff 04 75 00 03 52 45 56 20 30 33 00 00
 Address 0x10: 00 00 00 00 37 34 30 2d 30 34 35 39 33 33 00 00
 Address 0x20: 31 48 30 37 32 35 30 30 30 31 36 00 00 0e 0c 07
 Address 0x30: dc 30 43 ff ff ff ff ff ff ff ff ff ff ff ff ff
 Address 0x40: 02 02 00 ff 01 49 50 55 50 41 4a 39 4b 41 41 50
 Address 0x50: 57 52 2d 41 4d 58 31 31 30 30 2d 41 43 2d 53 00
 Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
 Address 0x70: ff ff ff 70 ff ff ff ff ff ff ff ff ff ff ff ff

PEM 1	REV 03	740-045932	1H073050017	DC Power Entry Module
-------	--------	------------	-------------	-----------------------

Jedec Code:	0x7fb0	EEPROM Version:	0x02
P/N:	740-045932	S/N:	1H073050017
Assembly ID:	0x0476	Assembly Version:	00.03
Date:	01-30-2013	Assembly Flags:	0x00
Version:	REV 03	CLEI Code:	IPUPAJ8KAA

```

ID: DC Power Entry Module          FRU Model Number: PWR-AMX1100-DC-S
Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff 02 02 00 ff
I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 76 00 03 52 45 56 20 30 33 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 34 35 39 33 32 00 00
  Address 0x20: 31 48 30 37 33 30 35 30 30 31 37 00 00 1e 01 07
  Address 0x30: dd 30 44 ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: 02 02 00 ff 01 49 50 55 50 41 4a 38 4b 41 41 50
  Address 0x50: 57 52 2d 41 4d 58 31 31 30 30 2d 44 43 2d 53 00
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 72 ff ff ff ff ff ff ff ff ff ff ff ff
Routing Engine 0 REV 20 750-044228 CAAY7935 RE-MX-104
Jedec Code: 0x7fb0 EEPROM Version: 0x02
P/N: 750-044228 S/N: CAAY7935
Assembly ID: 0x0b81 Assembly Version: 01.20
Date: 03-18-2013 Assembly Flags: 0x00
Version: REV 20 CLEI Code: PROTOXCLEI
ID: RE-MX-104 FRU Model Number: PROTO-ASSEMBLY
Board Information Record:
  Address 0x00: ad 01 00 08 b0 a8 6e a6 fc 10 ff ff ff ff ff ff
I2C Hex Data:
  Address 0x00: 7f b0 02 fe 0b 81 01 14 52 45 56 20 32 30 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 34 34 32 32 38 00 00
  Address 0x20: 53 2f 4e 20 43 41 41 59 37 39 33 35 00 12 03 07
  Address 0x30: dd ff ff ff ad 01 00 08 b0 a8 6e a6 fc 10 ff ff
  Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
  Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff c2 ff ff ff ff ff ff ff ff ff ff ff ff
da0 7836 MB ATP IG eUSB SSD Nand Flash 0
usb0 (addr 1) EHCI root hub 0 Freescale uhub0
usb0 (addr 2) USB2513Bi 9491 SMSC uhub1
usb0 (addr 3) ATP IG eUSB SSD 44801 ATP Electronics umass0
Routing Engine 1 REV 13 750-044228 CAAM6380 RE-MX-104
Jedec Code: 0x7fb0 EEPROM Version: 0x02
P/N: 750-044228 S/N: CAAM6380
Assembly ID: 0x0b81 Assembly Version: 01.13
Date: 09-17-2012 Assembly Flags: 0x00
Version: REV 13 CLEI Code: PROTOXCLEI
ID: RE-MX-104 FRU Model Number: PROTO-ASSEMBLY
Board Information Record:
  Address 0x00: ad 01 00 08 64 87 88 27 08 18 ff ff ff ff ff ff
I2C Hex Data:
  Address 0x00: 7f b0 02 fe 0b 81 01 0d 52 45 56 20 31 33 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 34 34 32 32 38 00 00
  Address 0x20: 53 2f 4e 20 43 41 41 4d 36 33 38 30 00 11 09 07
  Address 0x30: dc ff ff ff ad 01 00 08 64 87 88 27 08 18 ff ff
  Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
  Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff c2 ff ff ff ff ff ff ff ff ff ff ff ff
da0 7836 MB ATP IG eUSB SSD Nand Flash 0
AFEB 0 BUILTIN BUILTIN Forwarding Engine
Processor
FPC 0 BUILTIN BUILTIN MPC BUILTIN
FPC 1 BUILTIN BUILTIN MPC BUILTIN
MIC 0 REV 15 750-036132 CAAF7948 2x0C12/8x0C3 CC-CE
Jedec Code: 0x7fb0 EEPROM Version: 0x02
P/N: 750-036132 S/N: CAAF7948
Assembly ID: 0x0a1a Assembly Version: 01.15

```

```

Date:          07-03-2012      Assembly Flags:    0x00
Version:       REV 15          CLEI Code:       IP9IAM2DAA
ID: 2x0C12/8x0C3 CC-CE        FRU Model Number: MIC-3D-80C3-20C12-ATM
Board Information Record:
Address 0x00: 12 01 05 03 05 ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0a 1a 01 0f 52 45 56 20 31 35 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 33 36 31 33 32 00 00
Address 0x20: 53 2f 4e 20 43 41 41 46 37 39 34 38 00 03 07 07
Address 0x30: dc ff ff ff 12 01 05 03 05 ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 39 49 41 4d 32 44 41 41 4d
Address 0x50: 49 43 2d 33 44 2d 38 4f 43 33 2d 32 4f 43 31 32
Address 0x60: 2d 41 54 4d 00 00 41 00 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff e3 c0 02 a3 9c 00 00 00 00 0a 60 00 00
PIC 0          BUILTIN      BUILTIN      2x0C12/8x0C3 CC-CE
Xcvr 0        REV 01      740-011615    PCQOU2J      SFP-IR
Xcvr 1        REV 01      740-016068    P3L7A6G      SFP-SR
Xcvr 2        REV 01      740-016068    P3L7A5J      SFP-SR
Xcvr 3        REV 01      740-016065    PJN5HPZ      SFP-SR
Xcvr 4        REV 01      740-029122    PKB38TL      SFP-LR
Xcvr 5        REV 01      740-011787    P6A107G      SFP-LR
Xcvr 6        REV 01      740-029122    PKB38TR      SFP-LR
Xcvr 7        REV 01      740-011787    PBKONK3      SFP-LR
MIC 1
FPC 2          BUILTIN      BUILTIN      MPC BUILTIN
MIC 0          BUILTIN      BUILTIN      4x 10GE(LAN) SFP+
Jedec Code:    0x0000      EEPROM Version: 0x00
P/N:          BUILTIN      S/N:          BUILTIN
Assembly ID:    0x0a60      Assembly Version: 00.00
Date:          00-00-0000   Assembly Flags: 0x00
ID: 4x 10GE(LAN) SFP+
Board Information Record:
Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
Address 0x00: 00 00 00 00 0a 60 00 00 00 00 00 00 00 00 00 00
Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 4d 58 43 00
Address 0x20: 42 55 49 4c 54 49 4e 00 4d 58 43 00 00 00 00 00
Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 00 c0 02 a5 04 7f b0 02 ff 0a 1a 01 0f
PIC 0          BUILTIN      BUILTIN      4x 10GE(LAN) SFP+
Xcvr 0        REV 01      740-031980    B10F00465    SFP+-10G-SR
Xcvr 1        REV 01      740-031980    B10F00461    SFP+-10G-SR
Xcvr 2        REV 01      740-031980    B10G01545    SFP+-10G-SR
Xcvr 3        REV 01      740-031980    B10G01385    SFP+-10G-SR
Fan Tray 0    REV 02      711-049570    CAAX6538      Fan Tray
Jedec Code:    0x7fb0      EEPROM Version: 0x02
P/N:          711-049570   S/N:          CAAX6538
Assembly ID:    0x0b82      Assembly Version: 01.02
Date:          03-01-2013   Assembly Flags: 0x00
Version:       REV 02      CLEI Code:    PROTOXCLEI
ID: Fan Tray    FRU Model Number: PROTO-ASSEMBLY
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 82 01 02 52 45 56 20 30 32 00 00
Address 0x10: 00 00 00 00 37 31 31 2d 30 34 39 35 37 30 00 00
Address 0x20: 53 2f 4e 20 43 41 41 58 36 35 33 38 00 01 03 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

```

```

Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff
Address 0x70: ff ff ff c2 ff ff ff ff ff ff ff ff ff ff ff

```

show chassis hardware models (MX104 Router)

```

user@host> show chassis hardware models
Hardware inventory:
Item                Version  Part number  Serial number  FRU model number
Midplane            REV 20    750-044219   CAAS5849       PROTO-ASSEMBLY
PEM 0               REV 01    740-045932   1H072400065
Routing Engine 0    REV 16    750-044228   CAAR5915       PROTO-ASSEMBLY
AFEB 0              BUILTIN   BUILTIN
FPC 0               BUILTIN   BUILTIN
FPC 1               BUILTIN   BUILTIN
  MIC 0             REV 01    750-046905   CAAK7103       MIC-3D-20GE-SFP-EH
FPC 2               BUILTIN   BUILTIN
Fan Tray            REV 02    711-049570   CAAX6538       PROTO-ASSEMBLY

```

show chassis hardware clei-models (MX104 Router)

```

user@host> show chassis hardware clei-models
Hardware inventory:
Item                Version  Part number  CLEI code      FRU model number
Midplane            REV 20    750-044219   PROTOXCLEI     PROTO-ASSEMBLY
PEM 0               REV 01    740-045932
Routing Engine 0    REV 16    750-044228   PROTOXCLEI     PROTO-ASSEMBLY
AFEB 0              BUILTIN
FPC 0               BUILTIN
FPC 1               BUILTIN
  MIC 0             REV 01    750-046905   PROTOXCLEI     MIC-3D-20GE-SFP-EH
FPC 2               BUILTIN
Fan Tray            REV 02    711-049570   CAAX6538       PROTO-ASSEMBLY

```

show chassis hardware (MX240 Router)

```

user@host> show chassis hardware
Hardware inventory:
Item                Version  Part number  Serial number  Description
Chassis                                     JN10C7F7EAFC   MX240
Midplane            REV 01    710-021041   TR1502         MX240 Backplane
FPM Board           REV 01    710-017254   KD4017         Front Panel Display
PEM 0               Rev 02    740-017330   000332         PS 1.2-1.7kW; 100-240V
AC in
PEM 1               Rev 02    740-017330   000226         PS 1.2-1.7kW; 100-240V
AC in
Routing Engine 0    REV 06    740-013063   1000703522     RE-S-2000
Routing Engine 1    REV 06    740-015113   1000687625     RE-S-1300
CB 0                REV 07    710-013385   KC9057         MX SCB
CB 1                REV 05    710-013385   JY4760         MX SCB
FPC 1               REV 01    750-021679   KC7340         DPCE 40x 1GE R
  CPU               REV 06    710-013713   KD4078         DPC PMB
    PIC 0            BUILTIN   BUILTIN        10x 1GE(LAN)
      Xcvr 0          REV 01    740-011613   P9F18ME        SFP-SX
    PIC 1            BUILTIN   BUILTIN        10x 1GE(LAN)
    PIC 2            BUILTIN   BUILTIN        10x 1GE(LAN)
    PIC 3            BUILTIN   BUILTIN        10x 1GE(LAN)
FPC 2               REV 04    710-016669   JS4529         DPCE 40x 1GE R EQ
  CPU               REV 06    710-013713   KB3969         DPC PMB
    PIC 0            BUILTIN   BUILTIN        10x 1GE(LAN) EQ
      Xcvr 0          REV 01    740-011613   PBG3Y79        SFP-SX

```

Xcvr 1	REV 01	740-011613	PBG3XU8	SFP-SX
Xcvr 2	REV 01	740-011613	PBG3YG6	SFP-SX
Xcvr 3	REV 01	740-011613	PBG3XUG	SFP-SX
Xcvr 4	REV 01	740-011613	PBG3XTJ	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 0	REV 01	740-011613	PBG3ZUM	SFP-SX
Xcvr 1	REV 01	740-011613	PBG3Y5H	SFP-SX
Xcvr 2	REV 01	740-011613	PBG3UZT	SFP-SX
Xcvr 3	REV 01	740-011613	PBG3US1	SFP-SX
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 0	REV 01	740-011613	PBG3YG7	SFP-SX
Xcvr 1	REV 01	740-011613	PBG3XZ9	SFP-SX
Xcvr 2	REV 01	740-011613	PBG3XTY	SFP-SX
Xcvr 3	REV 01	740-011613	PBG3UZG	SFP-SX
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 0	REV 01	740-011613	PBG3Y8W	SFP-SX
Xcvr 1	REV 01	740-011613	PBG3YVX	SFP-SX
Xcvr 2	REV 01	740-011613	PBG3YB3	SFP-SX
Xcvr 3	REV 01	740-011613	PBG43VQ	SFP-SX
Fan Tray 0	REV 01	710-021113	JS4642	MX240 Fan Tray

show chassis hardware detail (MX 240 Router with Routing Engine Displaying DIMM Information)

```
user@host> show chassis hardware detail
```

Item	Version	Part number	Serial number	Description
Chassis			JN11279B4AFC	MX240 Backplane
Midplane	REV 07	760-021404	TS2474	MX240 Backplane
FPM Board	REV 03	760-021392	XC2643	Front Panel Display
PEM 0	Rev 03	740-017343	QCS0908A068	DC Power Entry Module
Routing Engine 0	REV 01	740-031117	AARCH00	RE-S-1800x4
ad0 3764 MB	STEC M2+	CF 9.0.2	STIM2Q3209239145303	Removable Compact Flash
ad1 28626 MB	WDC SSD-F0030S-5000		C933Z036237215548S00	Compact Flash
usb0 (addr 1)	EHCI root hub 0		Intel	uhub0
usb0 (addr 2)	product 0x0020 32		vendor 0x8087	uhub1
DIMM 0	VL31B5263E-F8S DIE REV-0 PCB REV-0			MFR ID-ce80
DIMM 1	VL31B5263E-F8S DIE REV-0 PCB REV-0			MFR ID-ce80
DIMM 2	VL31B5263E-F8S DIE REV-0 PCB REV-0			MFR ID-ce80
DIMM 3	SL31B5263E-F8S DIE REV-0 PCB REV-0			MFR ID-ce80
CB 0	REV 03	710-021523	XD7225	MX SCB
Fan Tray 0	REV 01	710-021113	WZ4986	MX240 Fan Tray

show chassis hardware (MX240 Router with Enhanced MX SCB)

```
user@host> show chassis hardware
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN10C7F7EAFC	MX240
Midplane	REV 01	710-021041	TR1502	MX240 Backplane
FPM Board	REV 01	710-017254	KD4017	Front Panel Display
PEM 0	Rev 02	740-017330	000332	PS 1.2-1.7kW; 100-240V
AC in				
PEM 1	Rev 02	740-017330	000226	PS 1.2-1.7kW; 100-240V
AC in				
Routing Engine 0	REV 06	740-013063	1000703522	RE-S-2000
Routing Engine 1	REV 06	740-015113	1000687625	RE-S-1300
CB 0	REV 02	710-031391	YE8494	Enhanced MX SCB
CB 1	REV 05	710-031391	YOP5764	Enhanced MX SCB
FPC 1	REV 01	750-021679	KC7340	DPCE 40x 1GE R
CPU	REV 06	710-013713	KD4078	DPC PMB

PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 0	REV 01	740-011613	P9F18ME	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN)
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN)
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN)
FPC 2	REV 04	710-016669	JS4529	DPCE 40x 1GE R EQ
CPU	REV 06	710-013713	KB3969	DPC PMB
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 0	REV 01	740-011613	PBG3Y79	SFP-SX
Xcvr 1	REV 01	740-011613	PBG3XU8	SFP-SX
Xcvr 2	REV 01	740-011613	PBG3YG6	SFP-SX
Xcvr 3	REV 01	740-011613	PBG3XUG	SFP-SX
Xcvr 4	REV 01	740-011613	PBG3XTJ	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 0	REV 01	740-011613	PBG3ZUM	SFP-SX
Xcvr 1	REV 01	740-011613	PBG3Y5H	SFP-SX
Xcvr 2	REV 01	740-011613	PBG3UZT	SFP-SX
Xcvr 3	REV 01	740-011613	PBG3US1	SFP-SX
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 0	REV 01	740-011613	PBG3YG7	SFP-SX
Xcvr 1	REV 01	740-011613	PBG3XZ9	SFP-SX
Xcvr 2	REV 01	740-011613	PBG3XTY	SFP-SX
Xcvr 3	REV 01	740-011613	PBG3UZG	SFP-SX
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 0	REV 01	740-011613	PBG3Y8W	SFP-SX
Xcvr 1	REV 01	740-011613	PBG3YVX	SFP-SX
Xcvr 2	REV 01	740-011613	PBG3YB3	SFP-SX
Xcvr 3	REV 01	740-011613	PBG43VQ	SFP-SX
Fan Tray 0	REV 01	710-021113	JS4642	MX240 Fan Tray

show chassis hardware (MX480 Router)

```
user@host> show chassis hardware
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN10C7F7FAFB	MX480
Midplane	REV 04	710-017414	TR2071	MX480 Midplane
FPM Board	REV 02	710-017254	KB8459	Front Panel Display
PEM 0	Rev 02	740-017330	QCS07519029	PS 1.2-1.7kW; 100-240V
AC in				
PEM 1	Rev 02	740-017330	QCS07519041	PS 1.2-1.7kW; 100-240V
AC in				
PEM 2	Rev 02	740-017330	QCS07519097	PS 1.2-1.7kW; 100-240V
AC in				
Routing Engine 0	REV 07	740-013063	1000733381	RE-S-2000
Routing Engine 1	REV 07	740-013063	1000733540	RE-S-2000
CB 0	REV 07	710-013385	KA8022	MX SCB
CB 1	REV 07	710-013385	KA8303	MX SCB
FPC 0	REV 09	750-020452	KA8660	DPCE 40x 1GE X EQ
CPU	REV 06	710-013713	KA8185	DPC PMB
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Fan Tray				Left Fan Tray

show chassis hardware (MX480 Router with Enhanced MX SCB)

```
user@host> show chassis hardware
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
------	---------	-------------	---------------	-------------

Chassis			JN10C7F7FAFB	MX480
Midplane	REV 04	710-017414	TR2071	MX480 Midplane
FPM Board	REV 02	710-017254	KB8459	Front Panel Display
PEM 0	Rev 02	740-017330	QCS07519029	PS 1.2-1.7kW; 100-240V
AC in				
PEM 1	Rev 02	740-017330	QCS07519041	PS 1.2-1.7kW; 100-240V
AC in				
PEM 2	Rev 02	740-017330	QCS07519097	PS 1.2-1.7kW; 100-240V
AC in				
Routing Engine 0	REV 07	740-013063	1000733381	RE-S-2000
Routing Engine 1	REV 07	740-013063	1000733540	RE-S-2000
CB 0	REV 07	710-013385	KA8022	Enhanced MX SCB
CB 1	REV 07	710-013385	KA8303	Enhanced MX SCB
FPC 0	REV 09	750-020452	KA8660	DPCE 40x 1GE X EQ
CPU	REV 06	710-013713	KA8185	DPC PMB
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Fan Tray				Left Fan Tray

show chassis hardware (MX480 Routers with MPC5E and Built-in OTN PIC)

```
user@host> show chassis hardware
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN11C0338AFB	MX480
Midplane	REV 05	710-017414	ABAB8430	MX480 Midplane
FPM Board	REV 02	710-017254	ZS8005	Front Panel Display
PEM 0	Rev 05	740-029970	QCS1024U089	PS 1.4-2.52kW; 90-264V
AC in				
PEM 1	Rev 10	740-029970	QCS1314U0FJ	PS 1.4-2.52kW; 90-264V
AC in				
PEM 2	Rev 07	740-029970	QCS1121U076	PS 1.4-2.52kW; 90-264V
AC in				
Routing Engine 0	REV 05	740-031116	9009092471	RE-S-1800x4
Routing Engine 1	REV 05	740-031116	9009097958	RE-S-1800x4
CB 0	REV 16	750-031391	CAAX0789	Enhanced MX SCB
CB 1	REV 16	750-031391	CAAX0856	Enhanced MX SCB
FPC 0	REV 32	750-028467	ABBP1782	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBP5410	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	983152A00038	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11F00211	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AQ72LPB	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AHNOWR5	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11J03627	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11F00300	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQ42WSS	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQ43HGC	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	ANAONDO	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	ANAONGF	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	ANAONG9	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	ANAOMP9	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AQA06CG	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	19T511100493	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	APR040J	SFP+-10G-SR
FPC 1	REV 26	750-046005	CACN1894	MPC5E 3D Q 2CGE+4XGE

CPU	REV 09	711-045719	CACN8698	RMP	CMP
PIC 0		BUILTIN	BUILTIN	2X10GE	SFP OTN
Xcvr 0	REV 01	740-031980	163363A03046	SFP+-10G-SR	
Xcvr 1	REV 01	740-031980	AJ40JS8	SFP+-10G-SR	
PIC 1		BUILTIN	BUILTIN	1X100GE	CFP2 OTN
PIC 2		BUILTIN	BUILTIN	2X10GE	SFP OTN
Xcvr 0	REV 01	740-031980	153363A00593	SFP+-10G-SR	
Xcvr 1	REV 01	740-031980	AJ40JUJ	SFP+-10G-SR	
PIC 3		BUILTIN	BUILTIN	1X100GE	CFP2 OTN
Xcvr 0		NON-JNPR	UQC0B53	CFP2-100G-LR4-D	
FPC 2	REV 26	750-046005	CACN1891	MPC5E 3D Q 2CGE+4XGE	
CPU	REV 09	711-045719	CACN8694	RMP	CMP
PIC 0		BUILTIN	BUILTIN	2X10GE	SFP OTN
Xcvr 0		NON-JNPR	URA012A	SFP+-10G-LR	
PIC 1		BUILTIN	BUILTIN	1X100GE	CFP2 OTN
Xcvr 0		NON-JNPR	J13F47042	CFP2-100G-LR4-D	
PIC 2		BUILTIN	BUILTIN	2X10GE	SFP OTN
Xcvr 0	REV 01	740-031980	AJC0BM3	SFP+-10G-SR	
Xcvr 1	REV 01	740-021308	11T511100917	SFP+-10G-SR	
PIC 3		BUILTIN	BUILTIN	1X100GE	CFP2 OTN
Xcvr 0		NON-JNPR	UQK07SU	CFP2-100G-LR4-D	
FPC 3	REV 03	750-045372	CAAD9425	MPCE Type 3 3D	
CPU	REV 08	711-035209	CAAD9094	HMP	CMP
MIC 0	REV 14	750-033196	CAAW9204	1X100GE	CXP
PIC 0		BUILTIN	BUILTIN	1X100GE	CXP
Xcvr 0	REV 01	740-046563	XD16FC034	CFP2-100G-SR10	
MIC 1	REV 19	750-033199	CAAJ1814	1X100GE	CFP
PIC 2		BUILTIN	BUILTIN	1X100GE	CFP
FPC 4	REV 21.0.11	750-045715	CAAY3568	MPC5E 3D Q 24XGE+6XLGE	
CPU	REV 07	711-045719	CAAW7430	RMP	CMP
PIC 0		BUILTIN	BUILTIN	12X10GE	SFP OTN
Xcvr 0	REV 01	740-031980	AP406NG	SFP+-10G-SR	
Xcvr 1	REV 01	740-021308	AR41NLP	SFP+-10G-SR	
Xcvr 2	REV 01	740-031980	B11D05630	SFP+-10G-SR	
PIC 1		BUILTIN	BUILTIN	12X10GE	SFP OTN
PIC 2		BUILTIN	BUILTIN	3X40GE	QSFP
PIC 3		BUILTIN	BUILTIN	3X40GE	QSFP
WAN MEZZ	REV 12	750-049136	CACM6678	MPC5E 24XGE OTN Mezz	
FPC 5	REV 11	750-045372	CABK7539	MPCE Type 3 3D	
CPU	REV 08	711-035209	CABJ2466	HMP	CMP
MIC 0	REV 19	750-033199	CAAJ9719	1X100GE	CFP
PIC 0		BUILTIN	BUILTIN	1X100GE	CFP
Xcvr 0	REV 01	740-035329	UP1020P	CFP-100G-SR10	
MIC 1	REV 07	750-033196	YZ0797	1X100GE	CXP
PIC 2		BUILTIN	BUILTIN	1X100GE	CXP
Xcvr 0	REV 01	740-046563	XC42FC022	CFP2-100G-SR10	
Fan Tray				Enhanced Left Fan Tray	

show chassis hardware detail (MX480 Routers with MPC5E and Built-in OTN PIC)

```
user@host> show chassis hardware detail
```

Hardware inventory:				
Item	Version	Part number	Serial number	Description
Chassis			JN11C0338AFB	MX480
Midplane	REV 05	710-017414	ABAB8430	MX480 Midplane
FPM Board	REV 02	710-017254	ZS8005	Front Panel Display
PEM 0	Rev 05	740-029970	QCS1024U089	PS 1.4-2.52kW; 90-264V
AC in				
PEM 1	Rev 10	740-029970	QCS1314U0FJ	PS 1.4-2.52kW; 90-264V
AC in				
PEM 2	Rev 07	740-029970	QCS1121U076	PS 1.4-2.52kW; 90-264V

AC in				
Routing Engine 0	REV 05	740-031116	9009092471	RE-S-1800x4
ad0	3896 MB	VRFCF14096DIHK1	VM4096MB 6862	Compact Flash
ad1	30533 MB	UGB94ARF32H0S3-KC	UNIGEN-478612-001127	Disk 1
usb0 (addr 1)		EHCI root hub 0	Intel	uhub0
usb0 (addr 2)		product 0x0020 32	vendor 0x8087	uhub1
DIMM 0		SGU04G72H1BB2SA-BB DIE	REV-52 PCB REV-54	MFR ID-ce80
DIMM 1		SGU04G72H1BB2SA-BB DIE	REV-52 PCB REV-54	MFR ID-ce80
DIMM 2		SGU04G72H1BB2SA-BB DIE	REV-52 PCB REV-54	MFR ID-ce80
DIMM 3		SGU04G72H1BB2SA-BB DIE	REV-52 PCB REV-54	MFR ID-ce80
Routing Engine 1	REV 05	740-031116	9009097958	RE-S-1800x4
ad0	3896 MB	VRFCF14096DIHK1	VM4096MB 6145	Compact Flash
ad1	30533 MB	UGB94ARF32H0S3-KC	UNIGEN-499551-000273	Disk 1
CB 0	REV 16	750-031391	CAAX0789	Enhanced MX SCB
CB 1	REV 16	750-031391	CAAX0856	Enhanced MX SCB
FPC 0	REV 32	750-028467	ABBP1782	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBP5410	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	983152A00038	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11F00211	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AQ72LPB	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AHNRW5	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11J03627	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11F00300	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQ42WSS	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQ43HGC	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	ANAOND0	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	ANAONGF	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	ANAONG9	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	ANAOMP9	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AQA06CG	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	19T511100493	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	APR040J	SFP+-10G-SR
FPC 1	REV 26	750-046005	CACN1894	MPC5E 3D Q 2CGE+4XGE
CPU	REV 09	711-045719	CACN8698	RMPC PMB
PIC 0		BUILTIN	BUILTIN	2X10GE SFPP OTN
Xcvr 0	REV 01	740-031980	163363A03046	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AJ40JS8	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	1X100GE CFP2 OTN
PIC 2		BUILTIN	BUILTIN	2X10GE SFPP OTN
Xcvr 0	REV 01	740-031980	153363A00593	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AJ40JUJ	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	1X100GE CFP2 OTN
Xcvr 0		NON-JNPR	UQC0B53	CFP2-100G-LR4-D
FPC 2	REV 26	750-046005	CACN1891	MPC5E 3D Q 2CGE+4XGE
CPU	REV 09	711-045719	CACN8694	RMPC PMB
PIC 0		BUILTIN	BUILTIN	2X10GE SFPP OTN
Xcvr 0		NON-JNPR	URA012A	SFP+-10G-LR
PIC 1		BUILTIN	BUILTIN	1X100GE CFP2 OTN
Xcvr 0		NON-JNPR	J13F47042	CFP2-100G-LR4-D
PIC 2		BUILTIN	BUILTIN	2X10GE SFPP OTN
Xcvr 0	REV 01	740-031980	AJC0BM3	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	11T511100917	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	1X100GE CFP2 OTN
Xcvr 0		NON-JNPR	UQK07SU	CFP2-100G-LR4-D
FPC 3	REV 03	750-045372	CAAD9425	MPCE Type 3 3D
CPU	REV 08	711-035209	CAAD9094	HMPC PMB 2G
MIC 0	REV 14	750-033196	CAAW9204	1X100GE CXP

PIC 0		BUILTIN	BUILTIN	1X100GE CXP
Xcvr 0	REV 01	740-046563	XD16FC034	CFP2-100G-SR10
MIC 1	REV 19	750-033199	CAAJ1814	1X100GE CFP
PIC 2		BUILTIN	BUILTIN	1X100GE CFP
FPC 4	REV 21.0.11	750-045715	CAAY3568	MPC5E 3D Q 24XGE+6XLGE
CPU	REV 07	711-045719	CAAW7430	RMPC PMB
PIC 0		BUILTIN	BUILTIN	12X10GE SFPP OTN
Xcvr 0	REV 01	740-031980	AP406NG	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AR41NLP	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11D05630	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	12X10GE SFPP OTN
PIC 2		BUILTIN	BUILTIN	3X40GE QSFPP
PIC 3		BUILTIN	BUILTIN	3X40GE QSFPP
WAN MEZZ	REV 12	750-049136	CACM6678	MPC5E 24XGE OTN Mezz
FPC 5	REV 11	750-045372	CABK7539	MPCE Type 3 3D
CPU	REV 08	711-035209	CABJ2466	HMPC PMB 2G
MIC 0	REV 19	750-033199	CAAJ9719	1X100GE CFP
PIC 0		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-035329	UP1020P	CFP-100G-SR10
MIC 1	REV 07	750-033196	YZ0797	1X100GE CXP
PIC 2		BUILTIN	BUILTIN	1X100GE CXP
Xcvr 0	REV 01	740-046563	XC42FC022	CFP2-100G-SR10
Fan Tray				Enhanced Left Fan Tray

show chassis hardware extensive (MX480 Routers with MPC5E and Built-in OTN PIC)

```

user@host> show chassis hardware extensive
Hardware inventory:
Item              Version  Part number  Serial number  Description
Chassis
Jedec Code:       0x7fb0          EEPROM Version: 0x02
S/N:              JN11C0338AFB
Assembly ID:      0x01fe          Assembly Version: 00.00
Date:             00-00-0000      Assembly Flags:  0x02
ID: MX480
Board Information Record:
Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
Address 0x00: 7f b0 02 ff 01 fe 00 00 00 00 00 00 00 00 00 00
Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x20: 4a 4e 31 31 43 30 33 33 38 41 46 42 02 00 00 00
Address 0x30: 00 00 00 ff 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Midplane          REV 05    710-017414  ABAB8430      MX480 Midplane
Jedec Code:       0x7fb0          EEPROM Version: 0x01
P/N:              710-017414      S/N:          ABAB8430
Assembly ID:      0x01fe          Assembly Version: 01.05
Date:             12-13-2011      Assembly Flags: 0x00
Version:          REV 05
ID: MX480 Midplane          FRU Model Number: CHAS-BP-MX480-S
Board Information Record:
Address 0x00: ad 01 08 00 00 23 9c fc 98 00 ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 01 fe 01 05 52 45 56 20 30 35 00 00
Address 0x10: 00 00 00 00 37 31 30 2d 30 31 37 34 31 34 00 00
Address 0x20: 53 2f 4e 20 41 42 41 42 38 34 33 30 00 0d 0c 07
Address 0x30: db ff ff ff ad 01 08 00 00 23 9c fc 98 00 ff ff

```

```

Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 43
Address 0x50: 48 41 53 2d 42 50 2d 4d 58 34 38 30 2d 53 00 00
Address 0x60: 00 00 00 00 00 00 ff ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

FPM Board          REV 02    710-017254    ZS8005          Front Panel Display
Jedec Code:    0x7fb0          EEPROM Version:    0x01
P/N:          710-017254          S/N:          ZS8005
Assembly ID:  0x01ff          Assembly Version:  01.02
Date:         11-21-2011        Assembly Flags:    0x00
Version:      REV 02
ID: Front Panel Display          FRU Model Number:  CRAFT-MX480-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 01 ff 01 02 52 45 56 20 30 32 00 00
Address 0x10: 00 00 00 00 37 31 30 2d 30 31 37 32 35 34 00 00
Address 0x20: 53 2f 4e 20 5a 53 38 30 30 35 00 00 00 15 0b 07
Address 0x30: db ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 43
Address 0x50: 52 41 46 54 2d 4d 58 34 38 30 2d 53 00 00 00 00
Address 0x60: 00 00 00 00 00 00 ff ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

PEM 0              Rev 05    740-029970    QCS1024U089    PS 1.4-2.52kW; 90-264V
AC in
Jedec Code:    0x7fb0          EEPROM Version:    0x01
P/N:          740-029970          S/N:          QCS1024U089
Assembly ID:  0x0432          Assembly Version:  01.05
Date:         06-17-2010        Assembly Flags:    0x00
Version:      Rev 05
ID: PS 1.4-2.52kW; 90-264V AC in FRU Model Number:  PWR-MX480-2520-AC-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff 00 00 00 00
I2C Hex Data:
Address 0x00: 7f b0 01 ff 04 32 01 05 52 65 76 20 30 35 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 32 39 39 37 30 00 00
Address 0x20: 51 43 53 31 30 32 34 55 30 38 39 00 00 11 06 07
Address 0x30: da ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: 00 00 00 00 01 00 00 00 00 00 00 00 00 00 00 50
Address 0x50: 57 52 2d 4d 58 34 38 30 2d 32 35 32 30 2d 41 43
Address 0x60: 2d 53 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

PEM 1              Rev 10    740-029970    QCS1314U0FJ    PS 1.4-2.52kW; 90-264V
AC in
Jedec Code:    0x7fb0          EEPROM Version:    0x01
P/N:          740-029970          S/N:          QCS1314U0FJ
Assembly ID:  0x0432          Assembly Version:  01.10
Date:         04-04-2013        Assembly Flags:    0x00
Version:      Rev 10
ID: PS 1.4-2.52kW; 90-264V AC in FRU Model Number:  PWR-MX480-2520-AC-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff 00 00 00 00
I2C Hex Data:
Address 0x00: 7f b0 01 ff 04 32 01 0a 52 65 76 20 31 30 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 32 39 39 37 30 00 00
Address 0x20: 51 43 53 31 33 31 34 55 30 46 4a 00 00 04 04 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: 00 00 00 00 01 00 00 00 00 00 00 00 00 00 00 50
Address 0x50: 57 52 2d 4d 58 34 38 30 2d 32 35 32 30 2d 41 43
Address 0x60: 2d 53 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

PEM 2              Rev 07    740-029970    QCS1121U076    PS 1.4-2.52kW; 90-264V

```

```

AC in
Jedec Code: 0x7fb0          EEPROM Version: 0x01
P/N: 740-029970          S/N: QCS1121U076
Assembly ID: 0x0432        Assembly Version: 01.07
Date: 05-23-2011          Assembly Flags: 0x00
Version: Rev 07
ID: PS 1.4-2.52kW; 90-264V AC in FRU Model Number: PWR-MX480-2520-AC-S
Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff 00 00 00 00
I2C Hex Data:
  Address 0x00: 7f b0 01 ff 04 32 01 07 52 65 76 20 30 37 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 32 39 39 37 30 00 00
  Address 0x20: 51 43 53 31 31 32 31 55 30 37 36 00 00 17 05 07
  Address 0x30: db ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: 00 00 00 00 01 00 00 00 00 00 00 00 00 00 00 50
  Address 0x50: 57 52 2d 4d 58 34 38 30 2d 32 35 32 30 2d 41 43
  Address 0x60: 2d 53 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Routing Engine 0 REV 05 740-031116 9009092471 RE-S-1800x4
Jedec Code: 0x7fb0          EEPROM Version: 0x02
P/N: 740-031116          S/N: 9009092471
Assembly ID: 0x09c0        Assembly Version: 01.05
Date: 11-01-2011          Assembly Flags: 0x00
Version: REV 05           CLEI Code: COUCALDBAA
ID: RE-S-1800x4           FRU Model Number: RE-S-1800X4-16G-S
Board Information Record:
  Address 0x00: 54 32 30 32 37 43 41 2d 34 32 46 42 23 23 23 00
I2C Hex Data:
  Address 0x00: 7f b0 02 ff 09 c0 01 05 52 45 56 20 30 35 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 33 31 31 31 36 00 00
  Address 0x20: 39 30 30 39 30 39 32 34 37 31 00 00 00 01 0b 07
  Address 0x30: db ff ff ff 54 32 30 32 37 43 41 2d 34 32 46 42
  Address 0x40: 23 23 23 00 01 43 4f 55 43 41 4c 44 42 41 41 52
  Address 0x50: 45 2d 53 2d 31 38 30 30 58 34 2d 31 36 47 2d 53
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 4b ff ff ff ff ff ff ff ff ff ff ff ff
ad0 3896 MB VRFCF14096DIHK1 VM4096MB 6862 Compact Flash
ad1 30533 MB UGB94ARF32H0S3-KC UNIGEN-478612-001127 Disk 1
usb0 (addr 1) EHCI root hub 0 Intel uhub0
usb0 (addr 2) product 0x0020 32 vendor 0x8087 uhub1
DIMM 0 SGU04G72H1BB2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
DIMM 1 SGU04G72H1BB2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
DIMM 2 SGU04G72H1BB2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
DIMM 3 SGU04G72H1BB2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
Routing Engine 1 REV 05 740-031116 9009097958 RE-S-1800x4
Jedec Code: 0x7fb0          EEPROM Version: 0x02
P/N: 740-031116          S/N: 9009097958
Assembly ID: 0x09c0        Assembly Version: 01.05
Date: 02-06-2012          Assembly Flags: 0x00
Version: REV 05           CLEI Code: COUCALDBAA
ID: RE-S-1800x4           FRU Model Number: RE-S-1800X4-16G-S
Board Information Record:
  Address 0x00: 54 32 30 32 37 43 41 2d 34 32 46 42 23 23 23 00
I2C Hex Data:
  Address 0x00: 7f b0 02 ff 09 c0 01 05 52 45 56 20 30 35 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 33 31 31 31 36 00 00
  Address 0x20: 39 30 30 39 30 39 37 39 35 38 00 00 00 06 02 07
  Address 0x30: dc ff ff ff 54 32 30 32 37 43 41 2d 34 32 46 42
  Address 0x40: 23 23 23 00 01 43 4f 55 43 41 4c 44 42 41 41 52
  Address 0x50: 45 2d 53 2d 31 38 30 30 58 34 2d 31 36 47 2d 53
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff

```

```

Address 0x70: ff ff ff 4b ff ff ff ff ff ff ff ff ff ff ff
ad0      3896 MB VRFCF14096DIHK1      VM4096MB 6145      Compact Flash
ad1      30533 MB UGB94ARF32H0S3-KC    UNIGEN-499551-000273 Disk 1

```

```
...
```

show chassis hardware (MX960 Router)

```

user@host> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               MX960
Midplane                               MX960 Midplane
PIM                               Power Inlet Module
PEM 2
PEM 3
Routing Engine 0 REV 00  740-015113  1000617944    RE-S-1300
CB 0                               MX960 Test SCB
FPC 4                               MX960 Test DPC
CPU
PIC 0                               1x 10GE (LAN/WAN)
PIC 1                               10x 1GE
FPC 7                               MX960 Test DPC
CPU
PIC 0                               1x 10GE (LAN/WAN)
Xcvr 0                               XFP-10G-SR
PIC 1                               10x 1GE
Xcvr 1                               SFP-SX
Xcvr 4                               SFP-SX
Xcvr 6                               SFP-SX
Xcvr 9                               SFP-SX
Fan Tray 0
Fan Tray 1

```

show chassis hardware (MX960 Router with Bidirectional Optics)

```

user@host> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               MX960
Midplane                               MX960 Backplane
FPM Board                               Front Panel Display
PDM                               Power Distribution Module
PEM 0                               PS 1.7kW; 200-240VAC in
PEM 1                               PS 1.7kW; 200-240VAC in
PEM 2                               PS 1.7kW; 200-240VAC in
Routing Engine 0 REV 06  740-013063  1000691458    RE-S-2000
CB 0                               MX SCB
CB 1                               MX SCB
FPC 3                               DPCE 40x 1GE R
CPU
FPC 4                               DPCE 40x 1GE R
CPU                               DPC PMB
PIC 0                               10x 1GE (LAN)
Xcvr 1                               SFP-1000BASE-BX40-D
Xcvr 2                               SFP-1000BASE-BX40-D
Xcvr 5                               SFP-1000BASE-BX10-U
Xcvr 6                               SFP-1000BASE-BX40-U
Xcvr 8                               SFP-1000BASE-BX40-U
PIC 1                               10x 1GE (LAN)
Xcvr 0                               SFP-1000BASE-BX10-D

```


Xcvr 1	REV 01	740-020465	75E467X00818	SFP-1000BASE-BX10-D
Xcvr 2	REV 01	740-020465	75E467X00573	SFP-1000BASE-BX10-D
Xcvr 3	REV 01	740-020465	4888227	SFP-1000BASE-BX10-D
Xcvr 4	REV 01	740-020465	4888241	SFP-1000BASE-BX10-D
Xcvr 5	REV 01	740-021340	77E245N00005	SFP-1000BASE-BX10-U
Xcvr 6	REV 01	740-021340	76E245X00487	SFP-1000BASE-BX10-U
Xcvr 7	REV 01	740-021341	5255889	SFP-1000BASE-BX10-U
Xcvr 8	REV 01	740-021341	5255887	SFP-1000BASE-BX10-U
Xcvr 9	REV 01	740-021340	77E245N00004	SFP-1000BASE-BX10-U
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 0	REV 01	740-020424	5007582	SFP-1000BASE-BX10-D
Xcvr 1	REV 01	740-020424	4888187	SFP-1000BASE-BX10-D
Xcvr 2	REV 01	740-020424	4656500	SFP-1000BASE-BX10-D
Xcvr 5	REV 01	740-021341	5255886	SFP-1000BASE-BX10-U
Xcvr 7	REV 01	740-021340	77E245N00003	SFP-1000BASE-BX10-U
Xcvr 8	REV 01	740-021341	5255888	SFP-1000BASE-BX10-U
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 0	REV 01	740-017726	74S184H30341	SFP-EX
Xcvr 1	REV 01	740-017726	4814061	SFP-EX
Xcvr 5	REV 01	740-017726	6ZS184H31108	SFP-EX
Xcvr 9	REV 01	740-021340	76E245X00486	SFP-1000BASE-BX10-U
Fan Tray 0				
Fan Tray 1	REV 03	740-014971	TP0850	Fan Tray

show chassis hardware (MX960 Router with Enhanced MX SCB)

```
user@host> show chassis hardware
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN1096805AFA	MX960
Midplane	REV 03	710-013698	TR0183	MX960 Backplane
Fan Extender	REV 02	710-018051	JY5227	Extended Cable Manager
FPM Board	REV 03	710-014974	JZ6876	Front Panel Display
PDM	Rev 03	740-013110	QCS11035023	Power Distribution Module
PEM 1	Rev 03	740-013682	QCS1109400L	PS 1.7kW; 200-240VAC in
PEM 2	Rev 03	740-013682	QCS11094015	PS 1.7kW; 200-240VAC in
PEM 3	Rev 03	740-013682	QCS11094012	PS 1.7kW; 200-240VAC in
Routing Engine 0	REV 06	740-013063	1000687969	RE-S-2000
Routing Engine 1	REV 06	740-013063	1000687955	RE-S-2000
CB 0	REV 11	750-031391	YZ6072	Enhanced MX SCB
CB 1	REV 11	750-031391	YZ6068	Enhanced MX SCB
CB 2	REV 11	750-031391	YZ6081	Enhanced MX SCB
FPC 0	REV 01	750-018122	KA5576	DPCE 40x 1GE R
CPU	REV 06	710-013713	KB3961	DPC PMB
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 0	REV 01	740-011613	P9F18GF	SFP-SX
Xcvr 2	REV 01	740-011782	P9M0TL9	SFP-SX
Xcvr 7	REV 01	740-011782	P9P0XXH	SFP-SX
Xcvr 9	REV 01	740-011782	P9M0TN1	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 0	REV 01	740-011613	PAJ4UHC	SFP-SX
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 0	REV 01	740-011613	PFF2CD0	SFP-SX
Xcvr 1	REV 01	740-011613	PBG3ZUT	SFP-SX
Xcvr 2	REV 01	740-011613	PFF2DDV	SFP-SX
Xcvr 5	REV 01	740-011613	P8E2SST	SFP-SX
Xcvr 9	REV 01	740-011782	PB8329N	SFP-SX
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 0	REV 01	740-026192	1U0201084503342	SFP-100BASE-BX10-U
Xcvr 1	REV 01	740-026193	1U1201084503313	SFP-100BASE-BX10-D
Xcvr 2	REV 01	740-011613	PAJ4Y5B	SFP-SX

Xcvr 6	REV 01	740-011782	P9M0U3M	SFP-SX
Xcvr 7	REV 01	740-011782	P9M0TLA	SFP-SX
FPC 1	REV 16	750-031089	YL0719	MPC Type 2 3D
CPU	REV 06	711-030884	YL1463	MPC PMB 2G
MIC 0	REV 07	750-028387	JR6500	3D 4x 10GE XFP
PIC 0		BUILTIN	BUILTIN	2x 10GE XFP
Xcvr 0	REV 01	740-014279	733019A00154	XFP-10G-LR
Xcvr 1	REV 02	740-014289	T09F55034	XFP-10G-SR
PIC 1		BUILTIN	BUILTIN	2x 10GE XFP
Xcvr 0	REV 01	740-014279	913019B00791	XFP-10G-LR
Xcvr 1	REV 01	740-014289	98S803A90384	XFP-10G-SR
MIC 1	REV 24	750-028387	YJ3950	3D 4x 10GE XFP
PIC 2		BUILTIN	BUILTIN	2x 10GE XFP
Xcvr 0	REV 02	740-014279	T10B36134	XFP-10G-LR
Xcvr 1	REV 01	740-014289	T07M86354	XFP-10G-SR
PIC 3		BUILTIN	BUILTIN	2x 10GE XFP
FPC 2	REV 08	710-014219	JY9654	DPCE 4x 10GE R
CPU	REV 06	710-013713	JZ6549	DPC PMB
PIC 0		BUILTIN	BUILTIN	1x 10GE(LAN/WAN)
PIC 1		BUILTIN	BUILTIN	1x 10GE(LAN/WAN)
PIC 2		BUILTIN	BUILTIN	1x 10GE(LAN/WAN)
Xcvr 0	REV 03	740-011571	C931BK028	XFP-10G-SR
PIC 3		BUILTIN	BUILTIN	1x 10GE(LAN/WAN)
FPC 3	REV 10	750-024199	XJ6692	MX FPC Type 3
CPU	REV 03	710-022351	XF5182	DPC PMB
PIC 0	REV 17	750-009553	RJ2945	4x OC-48 SONET
Xcvr 1	REV 01	740-011785	PCP3YLL	SFP-SR
Xcvr 3	REV 01	740-011785	PDSOMRY	SFP-SR
PIC 1	REV 32	750-003700	DP2113	1x OC-192 12xMM VSR
FPC 5	REV 25	750-028467	YM8256	MPC 3D 16x 10GE
CPU	REV 10	711-029089	YL3029	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 1	REV 01	740-031980	AHNOX1Z	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
FPC 7	REV 02	750-031092	JR6658	MPC Type 1 3D Q
CPU	REV 01	711-030884	JZ9038	MPC PMB 2G
MIC 0	REV 08	750-028392	JZ8737	3D 20x 1GE(LAN) SFP
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 0	REV 01	740-011782	PBE2C6Y	SFP-SX
Xcvr 2		NON-JNPR	U8105N8	SFP-SX
Xcvr 4	REV 01	740-011613	PFM18EF	SFP-SX
Xcvr 7	REV 01	740-011613	PFF2AM8	SFP-SX
Xcvr 8	REV 01	740-011613	PFF2CT6	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 0	REV 01	740-011782	PB82VHH	SFP-SX
Xcvr 1	REV 01	740-011613	PFF2CSW	SFP-SX
Xcvr 9	REV 01	740-011613	PFF2BY0	SFP-SX
QXM 0	REV 04	711-028408	JR6372	MPC QXM
FPC 8	REV 05	750-024387	JW9754	MX FPC Type 2
CPU	REV 03	710-022351	KF1651	DPC PMB
PIC 0	REV 08	750-014730	DM3664	4x OC-3 1x OC-12 SFP
Xcvr 0	REV 01	740-016065	81S290N00077	SFP-SR
Xcvr 1		NON-JNPR	2191844	SFP-SR
Xcvr 2	REV 01	740-011618	PD81EE5	SFP-IR
PIC 1	REV 08	750-014637	DM3671	4x OC-12-3 SFP
Xcvr 0	REV 01	740-011785	PCK3UNK	SFP-SR
Xcvr 3	REV 01	740-011785	PDSOMPZ	SFP-SR
FPC 10	REV 04	710-013699	JY4654	DPCE 40x 1GE R
CPU	REV 05	710-013713	JS9717	DPC PMB

PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 5	REV 01	740-011782	PAR1L72	SFP-SX
Xcvr 6	REV 01	740-011782	P8N1YQ4	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN)
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 0	REV 01	740-011782	P8Q2AVL	SFP-SX
Xcvr 5	REV 01	740-011782	PAR1L7B	SFP-SX
Xcvr 6	REV 01	740-011782	PAR1L2J	SFP-SX
Xcvr 8	REV 01	740-011782	P8N1YMY	SFP-SX
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN)
Fan Tray 0	REV 03	740-014971	TP0567	Fan Tray
Fan Tray 1	REV 03	740-014971	TP0702	Fan Tray

show chassis hardware models (MX960 Router with Enhanced MX SCB)

```
user@host> show chassis hardware models
Hardware inventory:
```

Item	Version	Part number	Serial number	FRU model number
Midplane	REV 03	710-013698	TR0183	CHAS-BP-MX960-S
Fan Extender	REV 02	710-018051	JY5227	ECM-MX960
FPM Board	REV 03	710-014974	JZ6876	CRAFT-MX960-S
Routing Engine 0	REV 06	740-013063	1000687969	RE-S-2000-4096-S
Routing Engine 1	REV 06	740-013063	1000687955	RE-S-2000-4096-S
CB 0	REV 11	750-031391	YZ6072	SCBE-MX-S
CB 1	REV 11	750-031391	YZ6068	SCBE-MX-S
CB 2	REV 11	750-031391	YZ6081	SCBE-MX-S
FPC 0	REV 01	750-018122	KA5576	DPCE-R-40GE-SFP
FPC 1	REV 16	750-031089	YL0719	MX-MPC2-3D
MIC 0	REV 07	750-028387	JR6500	MIC-3D-4XGE-XFP
MIC 1	REV 24	750-028387	YJ3950	MIC-3D-4XGE-XFP
FPC 2	REV 08	710-014219	JY9654	DPC-R-4XGE-XFP
FPC 3	REV 10	750-024199	XJ6692	MX-FPC3
PIC 0	REV 17	750-009553	RJ2945	PC-40C48-SON-SFP
PIC 1	REV 32	750-003700	DP2113	PC-10C192-SON-VSR
FPC 5	REV 25	750-028467	YM8256	MPC-3D-16XGE-SFP
FPC 7	REV 02	750-031092	JR6658	MX-MPC1-3D-Q
MIC 0	REV 08	750-028392	JZ8737	MIC-3D-20GE-SFP
FPC 8	REV 05	750-024387	JW9754	MX-FPC2
PIC 0	REV 08	750-014730	DM3664	PB-40C3-10C12-SON2-SFP
PIC 1	REV 08	750-014637	DM3671	PB-40C3-40C12-SON-SFP
FPC 10	REV 04	710-013699	JY4654	DPC-R-40GE-SFP
Fan Tray 0	REV 03	740-014971	TP0567	FFANTRAY-MX960-S
Fan Tray 1	REV 03	740-014971	TP0702	FFANTRAY-MX960-S

show chassis hardware (MX960 Router with MPC5EQ)

```
user@host> show chassis hardware
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN1214852AFA	MX960
Midplane	REV 01	710-030012	ACAX3674	MX960 Backplane
FPM Board	REV 03	710-014974	CAAZ9326	Front Panel Display
PDM	Rev 03	740-013110	QCS17025017	Power Distribution Module
PEM 0	Rev 10	740-027760	QCS1702N062	PS 4.1kW; 200-240V AC
in				
PEM 1	Rev 04	740-027760	QCS1422N02C	PS 4.1kW; 200-240V AC
in				
PEM 2	Rev 09	740-027760	QCS1614N01X	PS 4.1kW; 200-240V AC
in				
Routing Engine 0	REV 08	740-031116	9009131803	RE-S-1800x4
Routing Engine 1	REV 08	740-031116	9009124913	RE-S-1800x4

CB 0	REV 18	750-031391	CABF0579	Enhanced MX SCB
CB 1	REV 16	750-031391	CAAZ2471	Enhanced MX SCB
CB 2	REV 16	750-031391	CAAW9595	Enhanced MX SCB
FPC 0	REV 18	750-046005	CACE6574	MPC5E 3D Q 2CGE+4XGE
CPU	REV 09	711-045719	CACG8908	RMPC PMB
PIC 0		BUILTIN	BUILTIN	2X10GE SFPP OTN
Xcvr 0	REV 01	740-021308	AQA0DYT	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQG0MS7	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	1X100GE CFP2 OTN
Xcvr 0	REV 01	740-046563	XD16FC03Z	CFP2-100G-SR10
PIC 2		BUILTIN	BUILTIN	2X10GE SFPP OTN
Xcvr 0	REV 01	740-021308	ANA0NAJ	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQG0MRQ	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	1X100GE CFP2 OTN
Xcvr 0	REV 01	740-049775	J13K72993	CFP2-100G-LR4
FPC 1	REV 11	750-045372	CABK8154	MPCE Type 3 3D
CPU	REV 08	711-035209	CABE7370	HMPC PMB 2G
MIC 0	REV 07	750-033307	CABD5255	10X10GE SFPP
PIC 0		BUILTIN	BUILTIN	10X10GE SFPP
Xcvr 0	REV 01	740-021308	AQ50319	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQ5035V	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQ502XJ	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQ43HHR	SFP+-10G-SR
Xcvr 4	REV 01	740-021308	AQ502YA	SFP+-10G-SR
Xcvr 5	REV 01	740-021308	AQ502EU	SFP+-10G-SR
Xcvr 6	REV 01	740-021308	AQ502HR	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	AQ502A6	SFP+-10G-SR
Xcvr 8	REV 01	740-021308	AQ43H8M	SFP+-10G-SR
MIC 1	REV 14	750-033196	CAAP1398	1X100GE CXP
PIC 2		BUILTIN	BUILTIN	1X100GE CXP
Xcvr 0	REV 01	740-046563	XD16FC064	CFP-100G-SR10
FPC 3	REV 35	750-028467	CAAT9156	MPC 3D 16x 10GE
CPU	REV 11	711-029089	CAAV4645	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AQ43HZ1	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQ43HZC	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQ43HD2	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQ502HN	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AQ43HGF	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQ501RZ	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQ5029V	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQ501X9	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AQ502ZN	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQ43H86	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQ502ZY	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQ502PZ	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AQ503E6	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQ502XN	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11F00213	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQ50336	SFP+-10G-SR
FPC 4	REV 18	750-046005	CACE6568	MPC5E 3D Q 2CGE+4XGE
CPU	REV 09	711-045719	CACG8900	RMPC PMB
PIC 0		BUILTIN	BUILTIN	2X10GE SFPP OTN
Xcvr 0	REV 01	740-021308	AQA095A	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQG0M1E	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	1X100GE CFP2 OTN
Xcvr 0		NON-JNPR	FE13F000F	CFP2-100G-SR10
PIC 2		BUILTIN	BUILTIN	2X10GE SFPP OTN

Xcvr 0	REV 01	740-021308	AQGOLYC	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQGOLYB	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	1X100GE CFP2 OTN
Xcvr 0	REV 01	740-048813	XD32FE00Z	CFP2-100G-SR10
FPC 5	REV 18	750-046005	CACE6577	MPC5E 3D Q 2CGE+4XGE
CPU	REV 09	711-045719	CACG8902	RMPD PMB
PIC 0		BUILTIN	BUILTIN	2X10GE SFPP OTN
Xcvr 0	REV 01	740-021308	AQGOMXE	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQGOLVY	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	1X100GE CFP2 OTN
Xcvr 0	REV 01	740-046563	XD16FC03T	CFP2-100G-SR10
PIC 2		BUILTIN	BUILTIN	2X10GE SFPP OTN
Xcvr 0	REV 01	740-021308	AQGOLW1	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQGOLW3	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	1X100GE CFP2 OTN
Xcvr 0		NON-JNPR	FE13F000J	CFP2-100G-SR10
FPC 7	REV 09	750-037355	CAAF0937	MPC4E 3D 2CGE+8XGE
CPU	REV 08	711-035209	CAAD8004	HMPD PMB 2G
PIC 0		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-021308	ANAOMM3	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-035329	X000C163	CFP-100G-SR10
PIC 2		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-021308	AQGOMS6	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQGOMRX	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQGOM6Y	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQGOLZM	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-035329	X12J00499	CFP-100G-SR10
FPC 8	REV 39	750-045715	CACD1903	MPC5E 3D Q 24XGE+6XLGE
CPU	REV 09	711-045719	CACD1815	RMPD PMB
PIC 0		BUILTIN	BUILTIN	12X10GE SFPP OTN
PIC 1		BUILTIN	BUILTIN	12X10GE SFPP OTN
PIC 2		BUILTIN	BUILTIN	3X40GE QSFPP
Xcvr 0	REV 01	740-046565	QC480289	QSFPP+-40G-SR4
Xcvr 1	REV 01	740-046565	QC480274	QSFPP+-40G-SR4
Xcvr 2	REV 01	740-046565	QD130190	QSFPP+-40G-SR4
PIC 3		BUILTIN	BUILTIN	3X40GE QSFPP
Xcvr 0	REV 01	740-046565	QD130197	QSFPP+-40G-SR4
Xcvr 1	REV 01	740-046565	QD130180	QSFPP+-40G-SR4
Xcvr 2	REV 01	740-046565	QD130199	QSFPP+-40G-SR4
WAN MEZZ	REV 09	750-049136	CABN0415	MPC5E 24XGE OTN Mezz
FPC 9	REV 05	750-044444	CAAY9801	MPCE Type 2 3D P
CPU	REV 04	711-038484	CAAW3673	MPCE PMB 2G
MIC 0	REV 28	750-028387	CAAX1071	3D 4x 10GE XFP
PIC 0		BUILTIN	BUILTIN	2x 10GE XFP
Xcvr 0		NON-JNPR	T12L92342	XFP-10G-SR
Xcvr 1		NON-JNPR	T12L92303	XFP-10G-SR
PIC 1		BUILTIN	BUILTIN	2x 10GE XFP
Xcvr 0		NON-JNPR	CC07BK02X	XFP-10G-SR
QXM 0	REV 06	711-028408	CAAW4883	MPC QXM
QXM 1	REV 06	711-028408	CAAW4603	MPC QXM
FPC 10	REV 21.0.11	750-045715	CAAY3541	MPC5E 3D Q 24XGE+6XLGE
CPU	REV 07	711-045719	CAAW7426	RMPD PMB
PIC 0		BUILTIN	BUILTIN	12X10GE SFPP
Xcvr 0	REV 01	740-031980	AHK01AP	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQ502ZU	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AP41BLS	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQA08YA	SFP+-10G-SR
Xcvr 4	REV 01	740-021308	AQA0K26	SFP+-10G-SR
Xcvr 6	REV 01	740-021308	AQA06S3	SFP+-10G-SR

Xcvr 7	REV 01	740-021308	AQA06AS	SFP+-10G-SR
Xcvr 8	REV 01	740-021308	AQA053N	SFP+-10G-SR
Xcvr 9	REV 01	740-021308	AQA0E97	SFP+-10G-SR
Xcvr 10	REV 01	740-021308	AQA0GS4	SFP+-10G-SR
Xcvr 11	REV 01	740-021308	AQA0JVA	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	12X10GE SFPP
Xcvr 0	REV 01	740-021308	AQA057A	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	ANA0MLS	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQA093A	SFP+-10G-SR
Xcvr 3	REV 01	740-021309	943153A00075	SFP+-10G-LR
Xcvr 4	REV 01	740-021308	AQA077B	SFP+-10G-SR
Xcvr 5	REV 01	740-021308	AQA0JSC	SFP+-10G-SR
Xcvr 6	REV 01	740-021308	AQA0735	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	AQ5028N	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	AP40VN5	SFP+-10G-SR
Xcvr 9	REV 01	740-021308	AQA0K0J	SFP+-10G-SR
Xcvr 10	REV 01	740-021308	AQA07AP	SFP+-10G-SR
Xcvr 11	REV 01	740-021308	AQA08YB	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	3X40GE QSFPP
PIC 3		BUILTIN	BUILTIN	3X40GE QSFPP
WAN MEZZ	REV 07	750-045717	CAAX3123	MPC5E 24XGE Mezz
FPC 11	REV 17	750-037355	CAAT3986	MPC4E 3D 2CGE+8XGE
CPU	REV 08	711-035209	CAAR3972	HMPC PMB 2G
PIC 0		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-021308	AQA0DSE	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQ501Y3	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQ501XU	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQ5036Y	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0		NON-JNPR	X12J00247	CFP-100G-SR10
PIC 2		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-031980	ALQ1DKF	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AJ403YA	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AP40TY0	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	ALQ14G0	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-035329	X12J00095	CFP-100G-SR10
Fan Tray 0	REV 08	740-031521	ACAF4219	Enhanced Fan Tray
Fan Tray 1	REV 08	740-031521	ACAF4225	Enhanced Fan Tray

show chassis hardware detail (MX960 Router)

```
user@host> show chassis hardware detail
```

Hardware inventory:				
Item	Version	Part number	Serial number	Description
Chassis				MX960
Midplane	REV 01	710-013698	AA6082	MX960 Midplane
PIM	Rev 01	740-013110	000008	Power Inlet Module
PEM 2				
PEM 3	Rev 01	740-013682	000038	PS 1.7kW; 200-240VAC in
Routing Engine 0	REV 00	740-015113	1000617944	RE-S-1300
ad0	245 MB	SanDisk SDCFB-256	111419E1805T1141	Compact Flash
ad2	38154 MB	FUJITSU MHT2040BH	NROWT5925N77	Hard Disk
CB 0	REV 05	710-013725	JK6947	MX960 Test SCB
FPC 4	REV 01	710-013305	JM7617	MX960 Test DPC
CPU				
PIC 0		BUILTIN	BUILTIN	1x 10GE(LAN/WAN)
PIC 1		BUILTIN	BUILTIN	10x 1GE
FPC 7	REV 01	710-013305	JL9634	MX960 Test DPC
CPU				
PIC 0		BUILTIN	BUILTIN	1x 10GE(LAN/WAN)

```

Xcvr 0          NON-JNPR      MYBG65I82C      XFP-10G-SR
PIC 1           BUILTIN      BUILTIN         10x 1GE
Xcvr 1          REV 01       740-011782     P7N0368         SFP-SX
Xcvr 4          REV 01       740-011782     P8J1W27         SFP-SX
Xcvr 6          REV 01       740-011782     P8J1VSD         SFP-SX
Xcvr 9          REV 01       740-011782     P8J1W25         SFP-SX
Fan Tray 0
Fan Tray 1

```

show chassis hardware detail (MX960 Router with MPC5EQ)

```
user@host> show chassis hardware detail
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN1214852AFA	MX960
Midplane	REV 01	710-030012	ACAX3674	MX960 Backplane
FPM Board	REV 03	710-014974	CAAZ9326	Front Panel Display
PDM	Rev 03	740-013110	QCS17025017	Power Distribution Module
PEM 0	Rev 10	740-027760	QCS1702N062	PS 4.1kW; 200-240V AC
in				
PEM 1	Rev 04	740-027760	QCS1422N02C	PS 4.1kW; 200-240V AC
in				
PEM 2	Rev 09	740-027760	QCS1614N01X	PS 4.1kW; 200-240V AC
in				
Routing Engine 0	REV 08	740-031116	9009131803	RE-S-1800x4
ad0 3831 MB	UGB30SFA4000T1		SFA4000T1 000016CD	Compact Flash
ad1 30533 MB	UGB94BPH32H0S1-KCI		11000061346	Disk 1
usb0 (addr 1)	EHCI root hub 0		Intel	uhub0
usb0 (addr 2)	product 0x0020 32		vendor 0x8087	uhub1
DIMM 0	VL31B5263F-F8SD DIE	REV-0 PCB REV-0		MFR ID-ce80
DIMM 1	VL31B5263F-F8SD DIE	REV-0 PCB REV-0		MFR ID-ce80
DIMM 2	VL31B5263F-F8SD DIE	REV-0 PCB REV-0		MFR ID-ce80
DIMM 3	VL31B5263F-F8SD DIE	REV-0 PCB REV-0		MFR ID-ce80
Routing Engine 1	REV 08	740-031116	9009124913	RE-S-1800x4
ad0 3831 MB	UGB30SFA4000T1		SFA4000T1 0000106D	Compact Flash
ad1 30533 MB	UGB94BPH32H0S1-KCI		11000052402	Disk 1
CB 0	REV 18	750-031391	CABF0579	Enhanced MX SCB
CB 1	REV 16	750-031391	CAAZ2471	Enhanced MX SCB
CB 2	REV 16	750-031391	CAAW9595	Enhanced MX SCB
FPC 0	REV 18	750-046005	CACE6574	MPC5E 3D Q 2CGE+4XGE
CPU	REV 09	711-045719	CACG8908	RMPC PMB
PIC 0		BUILTIN	BUILTIN	2X10GE SFPP OTN
Xcvr 0	REV 01	740-021308	AQA0DYT	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQGOMS7	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	1X100GE CFP2 OTN
Xcvr 0	REV 01	740-046563	XD16FC03Z	CFP2-100G-SR10
PIC 2		BUILTIN	BUILTIN	2X10GE SFPP OTN
Xcvr 0	REV 01	740-021308	ANA0NAJ	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQGOMRQ	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	1X100GE CFP2 OTN
Xcvr 0	REV 01	740-049775	J13K72993	CFP2-100G-LR4
FPC 1	REV 11	750-045372	CABK8154	MPCE Type 3 3D
CPU	REV 08	711-035209	CABE7370	HMPC PMB 2G
MIC 0	REV 07	750-033307	CABD5255	10X10GE SFPP
PIC 0		BUILTIN	BUILTIN	10X10GE SFPP
Xcvr 0	REV 01	740-021308	AQ50319	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQ5035V	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQ502XJ	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQ43HHR	SFP+-10G-SR
Xcvr 4	REV 01	740-021308	AQ502YA	SFP+-10G-SR
Xcvr 5	REV 01	740-021308	AQ502EU	SFP+-10G-SR

Xcvr 6	REV 01	740-021308	AQ502HR	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	AQ502A6	SFP+-10G-SR
Xcvr 8	REV 01	740-021308	AQ43H8M	SFP+-10G-SR
MIC 1	REV 14	750-033196	CAAP1398	1X100GE CXP
PIC 2		BUILTIN	BUILTIN	1X100GE CXP
Xcvr 0	REV 01	740-046563	XD16FC064	CFP2-100G-SR10
FPC 3	REV 35	750-028467	CAAT9156	MPC 3D 16x 10GE
CPU	REV 11	711-029089	CAAV4645	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AQ43HZ1	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQ43HZC	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQ43HD2	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQ502HN	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AQ43HGF	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQ501RZ	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQ5029V	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQ501X9	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AQ502ZN	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQ43H86	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQ502ZY	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQ502PZ	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AQ503E6	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQ502XN	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11F00213	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQ50336	SFP+-10G-SR
FPC 4	REV 18	750-046005	CACE6568	MPC5E 3D Q 2CGE+4XGE
CPU	REV 09	711-045719	CACG8900	RMPC PMB
PIC 0		BUILTIN	BUILTIN	2X10GE SFPP OTN
Xcvr 0	REV 01	740-021308	AQA095A	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQG0M1E	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	1X100GE CFP2 OTN
Xcvr 0		NON-JNPR	FE13F000F	CFP2-100G-SR10
PIC 2		BUILTIN	BUILTIN	2X10GE SFPP OTN
Xcvr 0	REV 01	740-021308	AQG0LYC	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQG0LYB	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	1X100GE CFP2 OTN
Xcvr 0	REV 01	740-048813	XD32FE00Z	CFP2-100G-SR10
FPC 5	REV 18	750-046005	CACE6577	MPC5E 3D Q 2CGE+4XGE
CPU	REV 09	711-045719	CACG8902	RMPC PMB
PIC 0		BUILTIN	BUILTIN	2X10GE SFPP OTN
Xcvr 0	REV 01	740-021308	AQG0MXE	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQG0LVY	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	1X100GE CFP2 OTN
Xcvr 0	REV 01	740-046563	XD16FC03T	CFP2-100G-SR10
PIC 2		BUILTIN	BUILTIN	2X10GE SFPP OTN
Xcvr 0	REV 01	740-021308	AQG0LW1	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQG0LW3	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	1X100GE CFP2 OTN
Xcvr 0		NON-JNPR	FE13F000J	CFP2-100G-SR10
FPC 7	REV 09	750-037355	CAAF0937	MPC4E 3D 2CGE+8XGE
CPU	REV 08	711-035209	CAAD8004	HMPC PMB 2G
PIC 0		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-021308	ANA0MM3	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-035329	X000C163	CFP-100G-SR10
PIC 2		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-021308	AQG0MS6	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQG0MRX	SFP+-10G-SR

Xcvr 2	REV 01	740-021308	AQG0M6Y	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQG0LZM	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-035329	X12J00499	CFP-100G-SR10
FPC 8	REV 39	750-045715	CACD1903	MPC5E 3D Q 24XGE+6XLGE
CPU	REV 09	711-045719	CACD1815	RMPD PMB
PIC 0		BUILTIN	BUILTIN	12X10GE SFPP OTN
PIC 1		BUILTIN	BUILTIN	12X10GE SFPP OTN
PIC 2		BUILTIN	BUILTIN	3X40GE QSFPP
Xcvr 0	REV 01	740-046565	QC480289	QSFP+-40G-SR4
Xcvr 1	REV 01	740-046565	QC480274	QSFP+-40G-SR4
Xcvr 2	REV 01	740-046565	QD130190	QSFP+-40G-SR4
PIC 3		BUILTIN	BUILTIN	3X40GE QSFPP
Xcvr 0	REV 01	740-046565	QD130197	QSFP+-40G-SR4
Xcvr 1	REV 01	740-046565	QD130180	QSFP+-40G-SR4
Xcvr 2	REV 01	740-046565	QD130199	QSFP+-40G-SR4
WAN MEZZ	REV 09	750-049136	CABN0415	MPC5E 24XGE OTN Mezz
FPC 9	REV 05	750-044444	CAAY9801	MPCE Type 2 3D P
CPU	REV 04	711-038484	CAAW3673	MPCE PMB 2G
MIC 0	REV 28	750-028387	CAAX1071	3D 4x 10GE XFP
PIC 0		BUILTIN	BUILTIN	2x 10GE XFP
Xcvr 0		NON-JNPR	T12L92342	XFP-10G-SR
Xcvr 1		NON-JNPR	T12L92303	XFP-10G-SR
PIC 1		BUILTIN	BUILTIN	2x 10GE XFP
Xcvr 0		NON-JNPR	CC07BK02X	XFP-10G-SR
QXM 0	REV 06	711-028408	CAAW4883	MPC QXM
QXM 1	REV 06	711-028408	CAAW4603	MPC QXM
FPC 10	REV 21.0.11	750-045715	CAAY3541	MPC5E 3D Q 24XGE+6XLGE
CPU	REV 07	711-045719	CAAW7426	RMPD PMB
PIC 0		BUILTIN	BUILTIN	12X10GE SFPP
Xcvr 0	REV 01	740-031980	AHK01AP	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQ502ZU	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AP41BLS	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQA08YA	SFP+-10G-SR
Xcvr 4	REV 01	740-021308	AQA0K26	SFP+-10G-SR
Xcvr 6	REV 01	740-021308	AQA06S3	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	AQA06AS	SFP+-10G-SR
Xcvr 8	REV 01	740-021308	AQA053N	SFP+-10G-SR
Xcvr 9	REV 01	740-021308	AQA0E97	SFP+-10G-SR
Xcvr 10	REV 01	740-021308	AQA0GS4	SFP+-10G-SR
Xcvr 11	REV 01	740-021308	AQA0JVA	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	12X10GE SFPP
Xcvr 0	REV 01	740-021308	AQA057A	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	ANA0MLS	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQA093A	SFP+-10G-SR
Xcvr 3	REV 01	740-021309	943153A00075	SFP+-10G-LR
Xcvr 4	REV 01	740-021308	AQA077B	SFP+-10G-SR
Xcvr 5	REV 01	740-021308	AQA0JSC	SFP+-10G-SR
Xcvr 6	REV 01	740-021308	AQA0735	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	AQ5028N	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	AP40VN5	SFP+-10G-SR
Xcvr 9	REV 01	740-021308	AQA0K0J	SFP+-10G-SR
Xcvr 10	REV 01	740-021308	AQA07AP	SFP+-10G-SR
Xcvr 11	REV 01	740-021308	AQA08YB	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	3X40GE QSFPP
PIC 3		BUILTIN	BUILTIN	3X40GE QSFPP
WAN MEZZ	REV 07	750-045717	CAAX3123	MPC5E 24XGE Mezz
FPC 11	REV 17	750-037355	CAAT3986	MPC4E 3D 2CGE+8XGE
CPU	REV 08	711-035209	CAAR3972	HMPC PMB 2G
PIC 0		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-021308	AQA0DSE	SFP+-10G-SR

Xcvr 1	REV 01	740-021308	AQ501Y3	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQ501XU	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQ5036Y	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0		NON-JNPR	X12J00247	CFP-100G-SR10
PIC 2		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-031980	ALQ1DKF	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AJ403YA	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AP40TY0	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	ALQ14G0	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-035329	X12J00095	CFP-100G-SR10
Fan Tray 0	REV 08	740-031521	ACAF4219	Enhanced Fan Tray
Fan Tray 1	REV 08	740-031521	ACAF4225	Enhanced Fan Tray

show chassis hardware extensive (MX960 Router with MPC5EQ)

```

user@host> show chassis hardware extensive
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis
Jedec Code:   0x7fb0          EEPROM Version: 0x02
S/N:          JN1214852AFA
Assembly ID:  0x0512          Assembly Version: 00.00
Date:         00-00-0000      Assembly Flags:  0x00
ID: MX960
Board Information Record:
Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
Address 0x00: 7f b0 02 ff 05 12 00 00 00 00 00 00 00 00 00 00
Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x20: 4a 4e 31 32 31 34 38 35 32 41 46 41 00 00 00 00
Address 0x30: 00 00 00 ff 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Midplane      REV 01  710-030012  ACAX3674      MX960 Backplane
Jedec Code:   0x7fb0          EEPROM Version: 0x02
P/N:          710-030012      S/N:          ACAX3674
Assembly ID:  0x01df          Assembly Version: 01.01
Date:         01-19-2013      Assembly Flags: 0x00
Version:      REV 01          CLEI Code:    COM8T00CRB
ID: MX960 Backplane          FRU Model Number: CHAS-BP-MX960-S
Board Information Record:
Address 0x00: ad 01 08 00 54 e0 32 bc 68 00 ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 01 df 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 31 30 2d 30 33 30 30 31 32 00 00
Address 0x20: 53 2f 4e 20 41 43 41 58 33 36 37 34 00 13 01 07
Address 0x30: dd ff ff ff ad 01 08 00 54 e0 32 bc 68 00 ff ff
Address 0x40: ff ff ff ff 01 43 4f 4d 38 54 30 30 43 52 42 43
Address 0x50: 48 41 53 2d 42 50 2d 4d 58 39 36 30 2d 53 00 00
Address 0x60: 00 00 00 00 00 00 42 00 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff aa ff ff ff ff ff ff ff ff ff ff ff ff
FPM Board     REV 03  710-014974  CAAZ9326      Front Panel Display
Jedec Code:   0x7fb0          EEPROM Version: 0x01
P/N:          710-014974      S/N:          CAAZ9326
Assembly ID:  0x01e6          Assembly Version: 01.03
Date:         12-31-2012      Assembly Flags: 0x00
Version:      REV 03

```

ID: Front Panel Display FRU Model Number: CRAFT-MX960-S

Board Information Record:

Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

I2C Hex Data:

Address 0x00: 7f b0 01 ff 01 e6 01 03 52 45 56 20 30 33 00 00

Address 0x10: 00 00 00 00 37 31 30 2d 30 31 34 39 37 34 00 00

Address 0x20: 53 2f 4e 20 43 41 41 5a 39 33 32 36 00 1f 0c 07

Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 43

Address 0x50: 52 41 46 54 2d 4d 58 39 36 30 2d 53 00 00 00 00

Address 0x60: 00 00 00 00 00 00 ff ff ff ff ff ff ff ff ff ff

Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

PDM Rev 03 740-013110 QCS17025017 Power Distribution Module

Jedec Code: 0x7fb0 EEPROM Version: 0x01

P/N: 740-013110 S/N: QCS17025017

Assembly ID: 0x0416 Assembly Version: 01.03

Date: 01-10-2013 Assembly Flags: 0x00

Version: Rev 03

ID: Power Distribution Module

Board Information Record:

Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff 00 00 00

I2C Hex Data:

Address 0x00: 7f b0 01 ff 04 16 01 03 52 65 76 20 30 33 00 00

Address 0x10: 00 00 00 00 37 34 30 2d 30 31 33 31 31 30 00 00

Address 0x20: 51 43 53 31 37 30 32 35 30 31 37 00 00 0a 01 07

Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

Address 0x40: ff 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

PEM 0 Rev 10 740-027760 QCS1702N062 PS 4.1kW; 200-240V AC

in

Jedec Code: 0x7fb0 EEPROM Version: 0x01

P/N: 740-027760 S/N: QCS1702N062

Assembly ID: 0x0430 Assembly Version: 01.10

Date: 01-15-2013 Assembly Flags: 0x00

Version: Rev 10

ID: PS 4.1kW; 200-240V AC in FRU Model Number: PWR-MX960-4100-AC-S

Board Information Record:

Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff 00 00 00 00

I2C Hex Data:

Address 0x00: 7f b0 01 ff 04 30 01 0a 52 65 76 20 31 30 00 00

Address 0x10: 00 00 00 00 37 34 30 2d 30 32 37 37 36 30 00 00

Address 0x20: 51 43 53 31 37 30 32 4e 30 36 32 00 00 0f 01 07

Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

Address 0x40: 00 00 00 00 01 00 00 00 00 00 00 00 00 00 00 50

Address 0x50: 57 52 2d 4d 58 39 36 30 2d 34 31 30 30 2d 41 43

Address 0x60: 2d 53 00 00 00 00 00 00 00 00 00 00 00 00 00 00

Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

PEM 1 Rev 04 740-027760 QCS1422N02C PS 4.1kW; 200-240V AC

in

Jedec Code: 0x7fb0 EEPROM Version: 0x01

P/N: 740-027760 S/N: QCS1422N02C

Assembly ID: 0x0430 Assembly Version: 01.04

Date: 06-04-2010 Assembly Flags: 0x00

Version: Rev 04

ID: PS 4.1kW; 200-240V AC in FRU Model Number: PWR-MX960-4100-AC-S

Board Information Record:

Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff 00 00 00 00

I2C Hex Data:

Address 0x00: 7f b0 01 ff 04 30 01 04 52 65 76 20 30 34 00 00

```

Address 0x10: 00 00 00 00 37 34 30 2d 30 32 37 37 36 30 00 00
Address 0x20: 51 43 53 31 34 32 32 4e 30 32 43 00 00 04 06 07
Address 0x30: da ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: 00 00 00 00 01 00 00 00 00 00 00 00 00 00 00 50
Address 0x50: 57 52 2d 4d 58 39 36 30 2d 34 31 30 30 2d 41 43
Address 0x60: 2d 53 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
PEM 2          Rev 09    740-027760    QCS1614N01X    PS 4.1kW; 200-240V AC
in
Jedec Code:    0x7fb0          EEPROM Version:    0x01
P/N:           740-027760      S/N:              QCS1614N01X
Assembly ID:   0x0430          Assembly Version:  01.09
Date:          04-07-2012      Assembly Flags:    0x00
Version:       Rev 09
ID: PS 4.1kW; 200-240V AC in    FRU Model Number:  PWR-MX960-4100-AC-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff 00 00 00 00
I2C Hex Data:
Address 0x00: 7f b0 01 ff 04 30 01 09 52 65 76 20 30 39 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 32 37 37 36 30 00 00
Address 0x20: 51 43 53 31 36 31 34 4e 30 31 58 00 00 07 04 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: 00 00 00 00 01 00 00 00 00 00 00 00 00 00 00 50
Address 0x50: 57 52 2d 4d 58 39 36 30 2d 34 31 30 30 2d 41 43
Address 0x60: 2d 53 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Routing Engine 0 REV 08    740-031116    9009131803    RE-S-1800x4
Jedec Code:    0x7fb0          EEPROM Version:    0x02
P/N:           740-031116      S/N:              9009131803
Assembly ID:   0x09c0          Assembly Version:  01.08
Date:          03-04-2013      Assembly Flags:    0x00
Version:       REV 08          CLEI Code:         COUCASKBAA
ID: RE-S-1800x4          FRU Model Number:  RE-S-1800X4-16G-S
Board Information Record:
Address 0x00: 54 32 30 32 37 44 42 2d 34 34 47 42 23 42 23 00
I2C Hex Data:
Address 0x00: 7f b0 02 ff 09 c0 01 08 52 45 56 20 30 38 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 33 31 31 31 36 00 00
Address 0x20: 39 30 30 39 31 33 31 38 30 33 00 00 00 04 03 07
Address 0x30: dd ff ff ff 54 32 30 32 37 44 42 2d 34 34 47 42
Address 0x40: 23 42 23 00 01 43 4f 55 43 41 53 4b 42 41 41 52
Address 0x50: 45 2d 53 2d 31 38 30 30 58 34 2d 31 36 47 2d 53
Address 0x60: 00 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 59 ff ff ff ff ff ff ff ff ff ff ff ff ff
ad0    3831 MB    UGB30SFA4000T1          SFA4000T1 000016CD Compact Flash
ad1    30533 MB   UGB94BPH32H0S1-KCI    11000061346    Disk 1
usb0 (addr 1) EHCI root hub 0          Intel          uhub0
usb0 (addr 2) product 0x0020 32          vendor 0x8087    uhub1
DIMM 0          VL31B5263F-F8SD DIE REV-0 PCB REV-0    MFR ID-ce80
DIMM 1          VL31B5263F-F8SD DIE REV-0 PCB REV-0    MFR ID-ce80
DIMM 2          VL31B5263F-F8SD DIE REV-0 PCB REV-0    MFR ID-ce80
DIMM 3          VL31B5263F-F8SD DIE REV-0 PCB REV-0    MFR ID-ce80
Routing Engine 1 REV 08    740-031116    9009124913    RE-S-1800x4
Jedec Code:    0x7fb0          EEPROM Version:    0x02
P/N:           740-031116      S/N:              9009124913
Assembly ID:   0x09c0          Assembly Version:  01.08
Date:          01-09-2013      Assembly Flags:    0x00
Version:       REV 08          CLEI Code:         COUCASKBAA
ID: RE-S-1800x4          FRU Model Number:  RE-S-1800X4-16G-S
Board Information Record:
Address 0x00: 54 32 30 32 37 44 42 2d 34 34 47 42 23 42 23 00

```

```

I2C Hex Data:
Address 0x00: 7f b0 02 ff 09 c0 01 08 52 45 56 20 30 38 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 33 31 31 31 36 00 00
Address 0x20: 39 30 30 39 31 32 34 39 31 33 00 00 00 09 01 07
Address 0x30: dd ff ff ff 54 32 30 32 37 44 42 2d 34 34 47 42
Address 0x40: 23 42 23 00 01 43 4f 55 43 41 53 4b 42 41 41 52
Address 0x50: 45 2d 53 2d 31 38 30 30 58 34 2d 31 36 47 2d 53
Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 59 ff ff ff ff ff ff ff ff ff ff ff ff
ad0 3831 MB UGB30SFA4000T1 SFA4000T1 0000106D Compact Flash
ad1 30533 MB UGB94BPH32H0S1-KCI 11000052402 Disk 1
CB 0 REV 18 750-031391 CABF0579 Enhanced MX SCB
Jedec Code: 0x7fb0 EEPROM Version: 0x02
P/N: 750-031391 S/N: CABF0579
Assembly ID: 0x09b0 Assembly Version: 01.18
Date: 04-15-2013 Assembly Flags: 0x00
Version: REV 18 CLEI Code: COUCASRBAA
ID: Enhanced MX SCB FRU Model Number: SCBE-MX-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 09 b0 01 12 52 45 56 20 31 38 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 33 31 33 39 31 00 00
Address 0x20: 53 2f 4e 20 43 41 42 46 30 35 37 39 00 0f 04 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 43 4f 55 43 41 53 52 42 41 41 53
Address 0x50: 43 42 45 2d 4d 58 2d 53 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 43 00 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 7d ff ff ff ff ff ff ff ff ff ff ff ff
CB 1 REV 16 750-031391 CAAZ2471 Enhanced MX SCB
Jedec Code: 0x7fb0 EEPROM Version: 0x02
P/N: 750-031391 S/N: CAAZ2471
Assembly ID: 0x09b0 Assembly Version: 01.16
Date: 03-09-2013 Assembly Flags: 0x00
Version: REV 16 CLEI Code: COUCARCBAB
ID: Enhanced MX SCB FRU Model Number: SCBE-MX-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 09 b0 01 10 52 45 56 20 31 36 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 33 31 33 39 31 00 00
Address 0x20: 53 2f 4e 20 43 41 41 5a 32 34 37 31 00 09 03 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 43 4f 55 43 41 52 43 42 41 42 53
Address 0x50: 43 42 45 2d 4d 58 2d 53 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 42 00 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 6d ff ff ff ff ff ff ff ff ff ff ff ff
CB 2 REV 16 750-031391 CAAW9595 Enhanced MX SCB
Jedec Code: 0x7fb0 EEPROM Version: 0x02
P/N: 750-031391 S/N: CAAW9595
Assembly ID: 0x09b0 Assembly Version: 01.16
Date: 02-01-2013 Assembly Flags: 0x00
Version: REV 16 CLEI Code: COUCARCBAB
ID: Enhanced MX SCB FRU Model Number: SCBE-MX-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 09 b0 01 10 52 45 56 20 31 36 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 33 31 33 39 31 00 00
Address 0x20: 53 2f 4e 20 43 41 41 57 39 35 39 35 00 01 02 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

```

```

Address 0x40: ff ff ff ff 01 43 4f 55 43 41 52 43 42 41 42 53
Address 0x50: 43 42 45 2d 4d 58 2d 53 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 42 00 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 6d ff ff ff ff ff ff ff ff ff ff ff ff
FPC 0          REV 18    750-046005    CACE6574          MPC5E 3D Q 2CGE+4XGE
Jedec Code:    0x7fb0          EEPROM Version:    0x02
P/N:          750-046005          S/N:          CACE6574
Assembly ID:  0x0b8c          Assembly Version:  01.18
Date:         11-20-2013        Assembly Flags:    0x00
Version:      REV 18          CLEI Code:        PROTOXCLEI
ID: MPC5E 3D Q 2CGE+4XGE      FRU Model Number:  PROTO-ASSEMBLY
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 8c 01 12 52 45 56 20 31 38 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 34 36 30 30 35 00 00
Address 0x20: 53 2f 4e 20 43 41 43 45 36 35 37 34 00 14 0b 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
Address 0x70: ff ff ff c2 ff ff ff ff ff ff ff ff ff ff ff ff
CPU          REV 09    711-045719    CACG8908          RMPC PMB
Jedec Code:    0x7fb0          EEPROM Version:    0x02
P/N:          711-045719          S/N:          CACG8908
Assembly ID:  0x0b85          Assembly Version:  01.09
Date:         11-13-2013        Assembly Flags:    0x00
Version:      REV 09
ID: RMPC PMB
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 85 01 09 52 45 56 20 30 39 00 00
Address 0x10: 00 00 00 00 37 31 31 2d 30 34 35 37 31 39 00 00
Address 0x20: 53 2f 4e 20 43 41 43 47 38 39 30 38 00 0d 0b 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 00 50 52 4f 54 4f 58 43 4c 45 49 50
Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
Address 0x70: ff ff ff c2 00 00 00 00 00 00 00 00 00 00 00 00
PIC 0          BUILTIN    BUILTIN          2X10GE SFPP OTN
Jedec Code:    0x0000          EEPROM Version:    0x00
P/N:          BUILTIN          S/N:          BUILTIN
Assembly ID:  0x0a90          Assembly Version:  00.00
Date:         00-00-0000        Assembly Flags:    0x00
ID: 2X10GE SFPP OTN
Board Information Record:
Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
Address 0x00: 00 00 00 00 0a 90 00 00 00 00 00 00 00 00 00 00
Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 25 73 3a 20
Address 0x20: 42 55 49 4c 54 49 4e 00 25 73 3a 20 00 00 00 00
Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 00 c0 02 ae dc 00 00 00 00 0a 6e 00 00
Xcvr 0          REV 01    740-021308    AQA0DYT          SFP+-10G-SR
Xcvr 1          REV 01    740-021308    AQG0MS7          SFP+-10G-SR
PIC 1          BUILTIN    BUILTIN          1X100GE CFP2 OTN
Jedec Code:    0x0000          EEPROM Version:    0x00

```

```

P/N:          BUILTIN          S/N:          BUILTIN
Assembly ID:  0x0a6e          Assembly Version: 00.00
Date:         00-00-0000      Assembly Flags:  0x00
ID: 1X100GE CFP2 OTN
Board Information Record:
Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
Address 0x00: 00 00 00 00 0a 6e 00 00 00 00 00 00 00 00 00 00
Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 25 73 3a 20
Address 0x20: 42 55 49 4c 54 49 4e 00 25 73 3a 20 00 00 00 00
Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 00 c0 03 f3 8c 31 5c e7 80 00 00 00 02
Xcvr 0      REV 01      740-046563      XD16FC03Z      CFP2-100G-SR10
PIC 2      BUILTIN      BUILTIN      2X10GE SFPP OTN
Jedec Code: 0x0000      EEPROM Version: 0x00
P/N:          BUILTIN      S/N:          BUILTIN
Assembly ID:  0x0a90      Assembly Version: 00.00
Date:         00-00-0000      Assembly Flags:  0x00
ID: 2X10GE SFPP OTN
Board Information Record:
Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
Address 0x00: 00 00 00 00 0a 90 00 00 00 00 00 00 00 00 00 00
Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 25 73 3a 20
Address 0x20: 42 55 49 4c 54 49 4e 00 25 73 3a 20 00 00 00 00
Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 00 c0 03 f5 6c 31 5c db 40 00 00 00 02
Xcvr 0      REV 01      740-021308      ANAONAJ      SFP+-10G-SR
Xcvr 1      REV 01      740-021308      AQGOMRQ      SFP+-10G-SR
PIC 3      BUILTIN      BUILTIN      1X100GE CFP2 OTN
Jedec Code: 0x0000      EEPROM Version: 0x00
P/N:          BUILTIN      S/N:          BUILTIN
Assembly ID:  0x0a6e      Assembly Version: 00.00
Date:         00-00-0000      Assembly Flags:  0x00
ID: 1X100GE CFP2 OTN
Board Information Record:
Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
Address 0x00: 00 00 00 00 0a 6e 00 00 00 00 00 00 00 00 00 00
Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 25 73 3a 20
Address 0x20: 42 55 49 4c 54 49 4e 00 25 73 3a 20 00 00 00 00
Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 00 c0 03 ed ec 31 5c e2 e8 00 00 00 02
Xcvr 0      REV 01      740-049775      J13K72993      CFP2-100G-LR4
FPC 1      REV 11      750-045372      CABK8154      MPCE Type 3 3D
Jedec Code: 0x7fb0      EEPROM Version: 0x02
P/N:          750-045372      S/N:          CABK8154
Assembly ID:  0x09db      Assembly Version: 04.11
Date:         05-18-2013      Assembly Flags:  0x00
Version:      REV 11      CLEI Code:      COUIBBNBA
ID: MPCE Type 3 3D      FRU Model Number: MX-MPC3E-3D
Board Information Record:

```

```

Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 09 db 04 0b 52 45 56 20 31 31 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 34 35 33 37 32 00 00
Address 0x20: 53 2f 4e 20 43 41 42 4b 38 31 35 34 00 12 05 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 43 4f 55 49 42 42 4e 42 41 41 4d
Address 0x50: 58 2d 4d 50 43 33 45 2d 33 44 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 44 00 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff cf ff ff ff ff ff ff ff ff ff ff ff ff
CPU          REV 08    711-035209    CABE7370          HMPC PMB 2G
Jedec Code:  0x7fb0          EEPROM Version:  0x01
P/N:         711-035209          S/N:          CABE7370
Assembly ID: 0x0b04          Assembly Version: 01.08
Date:        05-08-2013          Assembly Flags: 0x00
Version:     REV 08
ID: HMPC PMB 2G
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 0b 04 01 08 52 45 56 20 30 38 00 00
Address 0x10: 00 00 00 00 37 31 31 2d 30 33 35 32 30 39 00 00
Address 0x20: 53 2f 4e 20 43 41 42 45 37 33 37 30 00 08 05 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00
MIC 0          REV 07    750-033307    CABD5255          10X10GE SFPP
Jedec Code:  0x7fb0          EEPROM Version:  0x02
P/N:         750-033307          S/N:          CABD5255
Assembly ID: 0x0a2a          Assembly Version: 02.07
Date:        04-25-2013          Assembly Flags: 0x00
Version:     REV 07          CLEI Code:     COUIBBJBAA
ID: 10X10GE SFPP          FRU Model Number: MIC3-3D-10XGE-SFPP
Board Information Record:
Address 0x00: 34 01 03 03 05 ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 fe 0a 2a 02 07 52 45 56 20 30 37 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 33 33 33 30 37 00 00
Address 0x20: 53 2f 4e 20 43 41 42 44 35 32 35 35 00 19 04 07
Address 0x30: dd ff ff ff 34 01 03 03 05 ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 43 4f 55 49 42 42 4a 42 41 41 4d
Address 0x50: 49 43 33 2d 33 44 2d 31 30 58 47 45 2d 53 46 50
Address 0x60: 50 00 00 00 00 00 41 00 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 82 c0 03 f0 bc 57 79 83 80 00 00 00 02
PIC 0          BUILTIN    BUILTIN          10X10GE SFPP
Xcvr 0          REV 01    740-021308    AQ50319          SFP+-10G-SR
    Xcvr 1          REV 01    740-021308    AQ5035V          SFP+-10G-SR
    Xcvr 2          REV 01    740-021308    AQ502XJ          SFP+-10G-SR
    Xcvr 3          REV 01    740-021308    AQ43HHR          SFP+-10G-SR
    Xcvr 4          REV 01    740-021308    AQ502YA          SFP+-10G-SR
    Xcvr 5          REV 01    740-021308    AQ502EU          SFP+-10G-SR
    Xcvr 6          REV 01    740-021308    AQ502HR          SFP+-10G-SR
    Xcvr 7          REV 01    740-021308    AQ502A6          SFP+-10G-SR
    Xcvr 8          REV 01    740-021308    AQ43H8M          SFP+-10G-SR
MIC 1          REV 14    750-033196    CAAP1398          1X100GE CXP
Jedec Code:  0x7fb0          EEPROM Version:  0x02
P/N:         750-033196          S/N:          CAAP1398
Assembly ID: 0x0a29          Assembly Version: 03.14
Date:        10-27-2012          Assembly Flags: 0x00

```



```

Version:      REV 14          CLEI Code:      COUIBBKBAA
ID: 1X100GE CXP          FRU Model Number: MIC3-3D-1X100GE-CXP
Board Information Record:
  Address 0x00: 34 01 07 07 08 ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
  Address 0x00: 7f b0 02 fe 0a 29 03 0e 52 45 56 20 31 34 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 33 33 31 39 36 00 00
  Address 0x20: 53 2f 4e 20 43 41 41 50 31 33 39 38 00 1b 0a 07
  Address 0x30: dc ff ff ff 34 01 07 07 08 ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 43 4f 55 49 42 42 4b 42 41 41 4d
  Address 0x50: 49 43 33 2d 33 44 2d 31 58 31 30 30 47 45 2d 43
  Address 0x60: 58 50 00 00 00 00 41 00 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 96 c0 03 ef cc 57 79 85 08 00 00 00 02
    PIC 2          BUILTIN          BUILTIN          1X100GE CXP
      Xcvr 0      REV 01      740-046563      XD16FC064      CFP2-100G-SR10
FPC 3          REV 35      750-028467      CAAT9156      MPC 3D 16x 10GE
Jedec Code: 0x7fb0          EEPROM Version: 0x01
P/N:      750-028467          S/N:      CAAT9156
Assembly ID: 0x0997          Assembly Version: 01.35
Date:      12-17-2012          Assembly Flags: 0x00
Version:      REV 35
ID: MPC 3D 16x 10GE          FRU Model Number: MPC-3D-16XGE-SFPP
Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
  Address 0x00: 7f b0 01 ff 09 97 01 23 52 45 56 20 33 35 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 32 38 34 36 37 00 00
  Address 0x20: 53 2f 4e 20 43 41 41 54 39 31 35 36 00 11 0c 07
  Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 4d
  Address 0x50: 50 43 2d 33 44 2d 31 36 58 47 45 2d 53 46 50 50
  Address 0x60: 00 00 00 00 00 00 ff ff ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
    CPU          REV 11      711-029089      CAAV4645      AMPC PMB
Jedec Code: 0x7fb0          EEPROM Version: 0x01
P/N:      711-029089          S/N:      CAAV4645
Assembly ID: 0x0998          Assembly Version: 01.11
Date:      12-13-2012          Assembly Flags: 0x00
Version:      REV 11
ID: AMPC PMB
Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
  Address 0x00: 7f b0 01 ff 09 98 01 0b 52 45 56 20 31 31 00 00
  Address 0x10: 00 00 00 00 37 31 31 2d 30 32 39 30 38 39 00 00
  Address 0x20: 53 2f 4e 20 43 41 41 56 34 36 34 35 00 0d 0c 07
  Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
  Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00
    PIC 0          BUILTIN          BUILTIN          4x 10GE(LAN) SFP+
Jedec Code: 0x0000          EEPROM Version: 0x00
P/N:      BUILTIN          S/N:      BUILTIN
Assembly ID: 0x02fe          Assembly Version: 00.00
Date:      00-00-0000          Assembly Flags: 0x00
ID: 4x 10GE(LAN) SFP+
Board Information Record:
  Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
  Address 0x00: 00 00 00 00 02 fe 00 00 00 00 00 00 00 00 00 00

```

```

Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 25 73 3a 20
Address 0x20: 42 55 49 4c 54 49 4e 00 25 73 3a 20 00 00 00 00
Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 00 c0 02 6b 94 00 00 00 00 02 fe 00 00
  Xcvr 0      REV 01  740-021308  AQ43HZ1      SFP+-10G-SR
  Xcvr 1      REV 01  740-021308  AQ43HZC      SFP+-10G-SR
  Xcvr 2      REV 01  740-021308  AQ43HD2      SFP+-10G-SR
  Xcvr 3      REV 01  740-021308  AQ502HN      SFP+-10G-SR
PIC 1        BUILTIN  BUILTIN      4x 10GE(LAN) SFP+
Jedec Code:  0x0000      EEPROM Version: 0x00
P/N:         BUILTIN      S/N:         BUILTIN
Assembly ID: 0x02fe      Assembly Version: 00.00
Date:        00-00-0000   Assembly Flags: 0x00
ID: 4x 10GE(LAN) SFP+
Board Information Record:
Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
Address 0x00: 00 00 00 00 02 fe 00 00 00 00 00 00 00 00 00 00
Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 25 73 3a 20
Address 0x20: 42 55 49 4c 54 49 4e 00 25 73 3a 20 00 00 00 00
Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 00 c0 02 ac 0c 00 00 00 00 02 fe 00 00
  Xcvr 0      REV 01  740-021308  AQ43HGF      SFP+-10G-SR
  Xcvr 1      REV 01  740-021308  AQ501RZ      SFP+-10G-SR
  Xcvr 2      REV 01  740-021308  AQ5029V      SFP+-10G-SR
  Xcvr 3      REV 01  740-021308  AQ501X9      SFP+-10G-SR
PIC 2        BUILTIN  BUILTIN      4x 10GE(LAN) SFP+
Jedec Code:  0x0000      EEPROM Version: 0x00
P/N:         BUILTIN      S/N:         BUILTIN
Assembly ID: 0x02fe      Assembly Version: 00.00
Date:        00-00-0000   Assembly Flags: 0x00
.....

```

show chassis hardware models (MX960 Router with MPC5EQ)

```

user@host> show chassis hardware models
Hardware inventory:
Item          Version  Part number  Serial number  FRU model number
Midplane      REV 01   710-030012  ACAX3674      CHAS-BP-MX960-S
FPM Board     REV 03   710-014974  CAAZ9326      CRAFT-MX960-S
PEM 0         Rev 10   740-027760  QCS1702N062   PWR-MX960-4100-AC-S
PEM 1         Rev 04   740-027760  QCS1422N02C   PWR-MX960-4100-AC-S
PEM 2         Rev 09   740-027760  QCS1614N01X   PWR-MX960-4100-AC-S
Routing Engine 0 REV 08   740-031116  9009131803    RE-S-1800X4-16G-S
Routing Engine 1 REV 08   740-031116  9009124913    RE-S-1800X4-16G-S
CB 0          REV 18   750-031391  CABF0579      SCBE-MX-S
CB 1          REV 16   750-031391  CAAZ2471      SCBE-MX-S
CB 2          REV 16   750-031391  CAAW9595      SCBE-MX-S
FPC 0         REV 18   750-046005  CACE6574      PROTO-ASSEMBLY
FPC 1         REV 11   750-045372  CABK8154      MX-MPC3E-3D
  MIC 0       REV 07   750-033307  CABD5255      MIC3-3D-10XGE-SFPP
  MIC 1       REV 14   750-033196  CAAP1398      MIC3-3D-1X100GE-CXP
FPC 3         REV 35   750-028467  CAAT9156      MPC-3D-16XGE-SFPP
FPC 4         REV 18   750-046005  CACE6568      PROTO-ASSEMBLY
FPC 5         REV 18   750-046005  CACE6577      PROTO-ASSEMBLY

```

FPC 7	REV 09	750-037355	CAAF0937	MPC4E-2CGE-8XGE
FPC 8	REV 39	750-045715	CACD1903	PROTO-ASSEMBLY
FPC 9	REV 05	750-044444	CAAY9801	MX-MPC2E-3D-P
MIC 0	REV 28	750-028387	CAAX1071	MIC-3D-4XGE-XFP
FPC 10	REV 21.0.11	750-045715	CAAY3541	PROTO-ASSEMBLY
FPC 11	REV 17	750-037355	CAAT3986	MPC4E-3D-2CGE-8XGE
Fan Tray 0	REV 08	740-031521	ACAF4219	FFANTRAY-MX960-HC-S
Fan Tray 1	REV 08	740-031521	ACAF4225	FFANTRAY-MX960-HC-S

show chassis hardware clei-models (MX960 Router with MPC5EQ)

```
user@host> show chassis hardware clei-models
```

Hardware inventory:

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 01	710-030012	COM8T00CRB	CHAS-BP-MX960-S
FPM Board	REV 03	710-014974		CRAFT-MX960-S
PEM 0	Rev 10	740-027760		PWR-MX960-4100-AC-S
PEM 1	Rev 04	740-027760		PWR-MX960-4100-AC-S
PEM 2	Rev 09	740-027760		PWR-MX960-4100-AC-S
Routing Engine 0	REV 08	740-031116	COUCASKBAA	RE-S-1800X4-16G-S
Routing Engine 1	REV 08	740-031116	COUCASKBAA	RE-S-1800X4-16G-S
CB 0	REV 18	750-031391	COUCASRBAA	SCBE-MX-S
CB 1	REV 16	750-031391	COUCARCBAB	SCBE-MX-S
CB 2	REV 16	750-031391	COUCARCBAB	SCBE-MX-S
FPC 0	REV 18	750-046005	PROTOXCLEI	PROTO-ASSEMBLY
FPC 1	REV 11	750-045372	COUIBBNBAA	MX-MPC3E-3D
MIC 0	REV 07	750-033307	COUIBBJBAA	MIC3-3D-10XGE-SFPP
MIC 1	REV 14	750-033196	COUIBBKBAA	MIC3-3D-1X100GE-CXP
FPC 3	REV 35	750-028467		MPC-3D-16XGE-SFPP
FPC 4	REV 18	750-046005	PROTOXCLEI	PROTO-ASSEMBLY
FPC 5	REV 18	750-046005	PROTOXCLEI	PROTO-ASSEMBLY
FPC 7	REV 09	750-037355	PROTOXCLEI	MPC4E-2CGE-8XGE
FPC 8	REV 39	750-045715	PROTOXCLEI	PROTO-ASSEMBLY
FPC 9	REV 05	750-044444	COUIBBGBAA	MX-MPC2E-3D-P
MIC 0	REV 28	750-028387	COUIA16BAA	MIC-3D-4XGE-XFP
FPC 10	REV 21.0.11	750-045715	PROTOXCLEI	PROTO-ASSEMBLY
FPC 11	REV 17	750-037355	IPU3A4DHAA	MPC4E-3D-2CGE-8XGE
Fan Tray 0	REV 08	740-031521		FFANTRAY-MX960-HC-S
Fan Tray 1	REV 08	740-031521		FFANTRAY-MX960-HC-S

show chassis hardware (MX2010 Router)

```
user@host > show chassis hardware
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN11E3217AFK	MX2010
Midplane	REV 01	750-044636	ABAB8506	Lower Backplane
Midplane 1	REV 01	711-044557	ZY8296	Upper Backplane
PMP	REV 03	711-032426	ACAJ1388	Power Midplane
FPM Board	REV 06	711-032349	ZX8744	Front Panel Display
PSM 4	REV 0C	740-033727	VK00254	DC 52V Power Supply
Module				
PSM 5	REV 0B	740-033727	VG00015	DC 52V Power Supply
Module				
PSM 6	REV 0B	740-033727	VH00097	DC 52V Power Supply
Module				
PSM 7	REV 0C	740-033727	VJ00151	DC 52V Power Supply
Module				
PSM 8	REV 0C	740-033727	VJ00149	DC 52V Power Supply
Module				
PDM 0	REV 0B	740-038109	WA00008	DC Power Dist Module

PDM 1	REV 0B	740-038109	WA00014	DC Power Dist Module
Routing Engine 0	REV 02	740-041821	9009094134	RE-S-1800x4
Routing Engine 1	REV 02	740-041821	9009094141	RE-S-1800x4
CB 0	REV 08	750-040257	CAAB3491	Control Board
CB 1	REV 08	750-040257	CAAB3489	Control Board
SPMB 0	REV 02	711-041855	CAAA6135	PMB Board
SPMB 1	REV 02	711-041855	CAAA6137	PMB Board
SFB 0	REV 06	711-032385	ZV1828	Switch Fabric Board
SFB 1	REV 07	711-032385	ZZ2568	Switch Fabric Board
SFB 2	REV 07	711-032385	ZZ2563	Switch Fabric Board
SFB 3	REV 07	711-032385	ZZ2564	Switch Fabric Board
SFB 4	REV 07	711-032385	ZZ2580	Switch Fabric Board
SFB 5	REV 07	711-032385	ZZ2579	Switch Fabric Board
SFB 6	REV 07	711-032385	CAAB4882	Switch Fabric Board
SFB 7	REV 07	711-032385	CAAB4898	Switch Fabric Board
FPC 0	REV 33	750-028467	CAAB1919	MPC 3D 16x 10GE
CPU	REV 11	711-029089	CAAB7174	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AMH02RE	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AMH038C	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AMH0390	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AMG0SUA	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AMH0579	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AMG0SGP	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AMH04SV	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AMH04X3	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AMH0135	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AMH02NC	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AMH02XB	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AMH02PN	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AMH057Y	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AMG0JHE	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AMH02HT	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AMH04V4	SFP+-10G-SR
FPC 1	REV 21	750-033205	ZG5027	MPC Type 3
CPU	REV 04	711-035209	YT4780	HMPC PMB 2G
MIC 0	REV 03	750-033307	ZV6299	10X10GE SFPP
PIC 0		BUILTIN	BUILTIN	10X10GE SFPP
Xcvr 0	REV 01	740-031980	083363A00410	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	083363A00334	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	113363A00125	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	083363A00953	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	AHR013D	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	AJ40JUR	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	AJ40JKL	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	AJ30ECK	SFP+-10G-SR
Xcvr 8	REV 01	740-021308	19T511100864	SFP+-10G-SR
Xcvr 9	REV 01	740-021308	19T511100868	SFP+-10G-SR
MIC 1	REV 03	750-033307	ZV6268	10X10GE SFPP
PIC 2		BUILTIN	BUILTIN	10X10GE SFPP
Xcvr 0	REV 01	740-031980	AJC0JML	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AJ403PC	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AJ10N25	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AJ40JF4	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	AJ40JSJ	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	AJ403V7	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	AJ40JN3	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	AJ40JSU	SFP+-10G-SR

Xcvr 8	REV 01	740-021308	19T511100468	SFP+-10G-SR
Xcvr 9	REV 01	740-021308	19T511101363	SFP+-10G-SR
FPC 8	REV 22	750-031089	ZT9746	MPC Type 2 3D
CPU	REV 06	711-030884	ZS1271	MPC PMB 2G
MIC 0	REV 26	750-028392	ABBS1150	3D 20x 1GE(LAN) SFP
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 0	REV 01	740-031851	PLG023C	SFP-SX
Xcvr 1	REV 01	740-031851	PLG09C6	SFP-SX
Xcvr 2	REV 02	740-011613	AM0950SF9L7	SFP-SX
Xcvr 3	REV 02	740-011613	AM1001SFN1H	SFP-SX
Xcvr 4	REV 02	740-011613	AM1001SFM9D	SFP-SX
Xcvr 5	REV 02	740-011613	AM1001SFLTJ	SFP-SX
Xcvr 6	REV 01	740-031851	AC1108S03L9	SFP-SX
Xcvr 7	REV 01	740-031851	AC1102S00NC	SFP-SX
Xcvr 8	REV 01	740-031851	AC1102S00MX	SFP-SX
Xcvr 9	REV 01	740-031851	AC1102S0085	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 0	REV 01	740-031851	AC1102S00KU	SFP-SX
Xcvr 1	REV 01	740-031851	AC1102S00NG	SFP-SX
Xcvr 2	REV 01	740-031851	AC1102S00K3	SFP-SX
Xcvr 3	REV 01	740-031851	AC1102S008R	SFP-SX
Xcvr 4	REV 01	740-031851	AM1107SUFVJ	SFP-SX
Xcvr 5	REV 01	740-031851	AC1108S03LC	SFP-SX
MIC 1	REV 26	750-028387	ABBR9582	3D 4x 10GE XFP
PIC 2		BUILTIN	BUILTIN	2x 10GE XFP
Xcvr 0		NON-JNPR	T10A91703	XFP-10G-SR
Xcvr 1		NON-JNPR	T09L42604	XFP-10G-SR
PIC 3		BUILTIN	BUILTIN	2x 10GE XFP
FPC 9	REV 11	750-036284	ZL3591	MPC 3D 16x 10GE EM
CPU	REV 10	711-029089	ZL0513	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	1YT517101825	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	1YT517101821	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	1YT517101682	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	ALQ13R6	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	1YT517101828	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	1YT517101716	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	1YT517101732	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	ALPOTR1	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	1YT517101741	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	1YT517101829	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	1YT517101669	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	ALQ14E3	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	1YT517101826	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	1YT517101817	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	1YT517101735	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	ALQ159A	SFP+-10G-SR
ADC 0	REV 05	750-043596	CAAC2073	Adapter Card
ADC 1	REV 01	750-043596	ZV4117	Adapter Card
ADC 8	REV 01	750-043596	ZV4107	Adapter Card
ADC 9	REV 02	750-043596	ZW1555	Adapter Card
Fan Tray 0	REV 2A	760-046960	ACAY0015	172mm FanTray - 6 Fans
Fan Tray 1	REV 2A	760-046960	ACAY0019	172mm FanTray - 6 Fans
Fan Tray 2	REV 2A	760-046960	ACAY0020	172mm FanTray - 6 Fans
Fan Tray 3	REV 2A	760-046960	ACAY0021	172mm FanTray - 6 Fans

show chassis hardware detail (MX2010 Router)

```

user@host > show chassis hardware detail
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               JN11E233DAFK  MX2010
Midplane      REV 26   750-044636   ABAB9357      Lower Backplane
Midplane 1    REV 01   711-044557   ABAB8643      Upper Backplane
PMP           REV 04   711-032426   ACAJ1677      Power Midplane
FPM Board     REV 08   760-044634   ABBV9726      Front Panel Display
PSM 0         REV 01   740-045050   1E02224000P   DC 52V Power Supply
Module
PSM 1         REV 01   740-045050   1E02224000M   DC 52V Power Supply
Module
PSM 2         REV 01   740-045050   1E022240010   DC 52V Power Supply
Module
PSM 3         REV 01   740-045050   1E02224000G   DC 52V Power Supply
Module
PSM 4         REV 01   740-045050   1E022240013   DC 52V Power Supply
Module
PSM 5         REV 01   740-045050   1E022240007   DC 52V Power Supply
Module
PSM 6         REV 01   740-045050   1E02224001C   DC 52V Power Supply
Module
PSM 7         REV 01   740-045050   1E02224001D   DC 52V Power Supply
Module
PSM 8         REV 01   740-045050   1E02224001B   DC 52V Power Supply
Module
PDM 0         REV 01   740-045234   1E262250067   DC Power Dist Module
Routing Engine 0 REV 02   740-041821   9009099704    RE-S-1800x4
  ad0  3831 MB  UGB30SFA4000T1  SFA4000T1 00000651 Compact Flash
  ad1  30533 MB UGB94BPH32H0S1-KCI 11000019592 Disk 1
  usb0 (addr 1) EHCI root hub 0 Intel uhub0
  usb0 (addr 2) product 0x0020 32 vendor 0x8087 uhub1
  DIMM 0 SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 1 SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 2 SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 3 SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
Routing Engine 1 REV 02   740-041821   9009099706    RE-S-1800x4
  ad0  3998 MB  Virtium - TuffDrive VCF P1T0200262860208 114 Compact Flash
  ad1  30533 MB UGB94ARF32H0S3-KC UNIGEN-499551-000404 Disk 1
CB 0          REV 13   750-040257   CAAF8436      Control Board
CB 1          REV 13   750-040257   CAAF8434      Control Board
SPMB 0        REV 02   711-041855   ABBV3825      PMB Board
SPMB 1        REV 02   711-041855   ABBV3833      PMB Board
SFB 0         REV 05   711-044466   ABBX5682      Switch Fabric Board
SFB 1         REV 05   711-044466   ABBX5676      Switch Fabric Board
SFB 2         REV 05   711-044466   ABBX5665      Switch Fabric Board
SFB 3         REV 05   711-044466   ABBX5699      Switch Fabric Board
SFB 4         REV 05   711-044466   ABBX5603      Switch Fabric Board
SFB 5         REV 05   711-044466   ABBX5587      Switch Fabric Board
SFB 6         REV 05   711-044466   ABBX5607      Switch Fabric Board
SFB 7         REV 05   711-044466   ABBX5669      Switch Fabric Board
FPC 0         REV 09   750-037355   CAAF0924      MPC Type 4-2
CPU           REV 08   711-035209   CAAB9842      HMPC PMB 2G
PIC 0         BUILTIN  BUILTIN      4x10GE SFPP
  Xcvr 0      REV 01   740-021308   19T511101656 SFP+-10G-SR
  Xcvr 1      REV 01   740-031980   AMA04RU      SFP+-10G-SR
  Xcvr 2      REV 01   740-031980   193363A00558 SFP+-10G-SR
  Xcvr 3      REV 01   740-031980   B10M00202   SFP+-10G-SR
PIC 1         BUILTIN  BUILTIN      1X100GE CFP

```

Xcvr 0		NON-JNPR	X12J00328	CFP-100G-SR10
PIC 2		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-031980	AMA088W	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B10L04211	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	19T511101602	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B10L04151	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0		NON-JNPR	X12J00332	CFP-100G-SR10
FPC 1	REV 18	750-033205	ZE0128	MPC Type 3
CPU	REV 06	711-035209	ZG5431	HMPC PMB 2G
MIC 0	REV 15	750-033199	ZP6435	1X100GE CFP
PIC 0		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-032210	J11E46118	CFP-100G-LR4
MIC 1	REV 15	750-033199	ZP6442	1X100GE CFP
PIC 2		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-032210	UMN03T4	CFP-100G-LR4
FPC 2	REV 16	750-037358	CAAL1001	MPC Type 4-1
CPU	REV 08	711-035209	CAAK7927	HMPC PMB 2G
PIC 0		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	193363A00589	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	973152A00028	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	193363A00376	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	973152A00016	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	193363A00499	SFP+-10G-SR
Xcvr 5	REV 01	740-021308	973152A00039	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	B11E01239	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	973152A00058	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	B10M00075	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	973152A00014	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AMA0638	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	973152A00063	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	AMA0629	SFP+-10G-SR
Xcvr 5	REV 01	740-021308	973152A00053	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	193363A00344	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	973152A00046	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	AMA062M	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	973152A00080	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	193363A00580	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	973152A00064	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	093363A01494	SFP+-10G-SR
Xcvr 5	REV 01	740-021308	973152A00020	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	123363A00047	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	973152A00072	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-021308	03DZ06A01033	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	973152A00022	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	03DZ06A01026	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	973152A00013	SFP+-10G-SR
Xcvr 4	REV 01	740-021308	03DZ06A01028	SFP+-10G-SR
Xcvr 5	REV 01	740-021308	973152A00079	SFP+-10G-SR
Xcvr 6	REV 01	740-021308	03DZ06A01018	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	973152A00025	SFP+-10G-SR
FPC 3	REV 33	750-028467	CAAF5400	MPC 3D 16x 10GE
CPU	REV 11	711-029089	CAAH7626	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	973152A00066	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	973152A00021	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	973152A00062	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	973152A00027	SFP+-10G-SR

PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	973152A00065	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	973152A00069	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	973152A00026	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	973152A00003	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	973152A00035	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	973152A00004	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	973152A00049	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	973152A00055	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	973152A00010	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	973152A00001	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	973152A00073	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	973152A00012	SFP+-10G-SR
FPC 4	REV 21	750-033205	ZG5028	MPC Type 3
CPU	REV 05	711-035209	YX3911	HMPC PMB 2G
MIC 0	REV 03	750-036233	ZL2036	2X40GE QSFP
PIC 0		BUILTIN	BUILTIN	2X40GE QSFP
Xcvr 0	REV 01	740-032986	QB220708	QSFP+-40G-SR4
Xcvr 1	REV 01	740-032986	QB220735	QSFP+-40G-SR4
MIC 1	REV 03	750-036233	ZL2028	2X40GE QSFP
PIC 2		BUILTIN	BUILTIN	2X40GE QSFP
Xcvr 0	REV 01	740-032986	QB220727	QSFP+-40G-SR4
Xcvr 1	REV 01	740-032986	QB220715	QSFP+-40G-SR4
FPC 5	REV 11	750-037358	CAAE2196	MPC Type 4-1
CPU	REV 08	711-035209	CAAD9074	HMPC PMB 2G
PIC 0		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	AMA062S	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AMA062P	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AMA052R	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AMA0632	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	193363A00564	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	193363A00229	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	193363A00363	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	193363A00278	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	AMA04CC	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AD0927A001W	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AMA04N2	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AMA062U	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	193363A00491	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	183363A01511	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	193363A00565	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	193363A00405	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	AMA07QX	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AMA06MS	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	193363A00318	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	193363A00402	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	193363A00174	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	193363A00388	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	193363A00377	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	193363A00234	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	AMA062T	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	193363A00550	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	193363A00364	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AMA0630	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	193363A00509	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	193363A00459	SFP+-10G-SR

Xcvr 6	REV 01	740-031980	113363A00191	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	193363A00352	SFP+-10G-SR
FPC 6	REV 33	750-028467	CAAF5552	MPC 3D 16x 10GE
CPU	REV 11	711-029089	CAAH7601	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AD0927A0036	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AD0927A003M	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AD0927A003G	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AD0927A0031	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	193363A00331	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	193363A00325	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	193363A00417	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	183363A02509	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	T09K75140	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11A04356	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01952	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K01914	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	T09K75157	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	T09K75194	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01926	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K01936	SFP+-10G-SR
FPC 7	REV 16	750-037358	CAAL1012	MPC Type 4-1
CPU	REV 08	711-035209	CAAJ3851	HMPC PMB 2G
PIC 0		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	AMA04NK	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11F00260	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11E02192	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AMA04CP	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	AJ40JJK	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	B11F00238	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	B10M00275	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	193363A00211	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	B11D05577	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11G00586	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AMA08B7	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AMA04Q0	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	B11D05840	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	B11E00467	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	B11E00029	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	19T511101712	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	193363A00568	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B10M00166	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B10M00212	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11D05823	SFP+-10G-SR
Xcvr 4	REV 01	740-021308	03DZ06A01005	SFP+-10G-SR
Xcvr 5	REV 01	740-021308	03DZ06A01003	SFP+-10G-SR
Xcvr 6	REV 01	740-021308	03DZ06A01009	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	03DZ06A01004	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-021308	03DZ06A01017	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	03DZ06A01016	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	03DZ06A01024	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	03DZ06A01008	SFP+-10G-SR
Xcvr 4	REV 01	740-030658	AD0946A02UH	SFP+-10G-USR
Xcvr 5	REV 01	740-021308	T09J67913	SFP+-10G-SR
Xcvr 6	REV 01	740-021308	AD0837E509G	SFP+-10G-SR

Xcvr 7	REV 01	740-021308	03DZ06A01015	SFP+-10G-SR
FPC 8	REV 03	750-045372	CAAD3111	MPC Type 3
CPU	REV 08	711-035209	CAAD8033	HMPC PMB 2G
MIC 0	REV 03	750-036233	ZL2032	2X40GE QSFP
PIC 0		BUILTIN	BUILTIN	2X40GE QSFP
Xcvr 0	REV 01	740-032986	QB230273	QSFP+-40G-SR4
Xcvr 1	REV 01	740-032986	QB230254	QSFP+-40G-SR4
MIC 1	REV 03	750-036233	ZL2021	2X40GE QSFP
PIC 2		BUILTIN	BUILTIN	2X40GE QSFP
Xcvr 0	REV 01	740-032986	QB390962	QSFP+-40G-SR4
Xcvr 1	REV 01	740-032986	QB390960	QSFP+-40G-SR4
FPC 9	REV 09	750-037355	CAAF1531	MPC Type 4-2
CPU	REV 08	711-035209	CAAB9927	HMPC PMB 2G
PIC 0		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-031980	193363A00525	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	193363A00504	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	193363A00368	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AJ40JSS	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	1X100GE CFP
PIC 2		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-031980	123363A00042	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B10M00023	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AJ802EM	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11E02348	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	1X100GE CFP
ADC 0	REV 13	750-043596	ABBX5532	Adapter Card
ADC 1	REV 13	750-043596	ABBX5550	Adapter Card
ADC 2	REV 13	750-043596	ABBX5571	Adapter Card
ADC 3	REV 13	750-043596	ABBX5568	Adapter Card
ADC 4	REV 13	750-043596	ABBX5556	Adapter Card
ADC 5	REV 13	750-043596	ABBX5553	Adapter Card
ADC 6	REV 13	750-043596	ABBX5541	Adapter Card
ADC 7	REV 13	750-043596	ABBX5578	Adapter Card
ADC 8	REV 13	750-043596	ABBX5560	Adapter Card
ADC 9	REV 07	750-043596	ABBV7188	Adapter Card
Fan Tray 0	REV 03	760-046960	ACAY0127	172mm FanTray - 6 Fans
Fan Tray 1	REV 2A	760-046960	ACAY0068	172mm FanTray - 6 Fans
Fan Tray 2	REV 2A	760-046960	ACAY0072	172mm FanTray - 6 Fans
Fan Tray 3	REV 2A	760-046960	ACAY0070	172mm FanTray - 6 Fans

show chassis hardware extensive (MX2010 Router)

```

user@host > show chassis hardware extensive
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               JN11E233DAFK  MX2010
Jedec Code:   0x7fb0                  EEPROM Version: 0x02
S/N:          JN11E233DAFK
Assembly ID:  0x0557                  Assembly Version: 00.00
Date:         00-00-0000              Assembly Flags:  0x00
ID: MX2010
Board Information Record:
Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
Address 0x00: 7f b0 02 ff 05 57 00 00 00 00 00 00 00 00 00 00
Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x20: 4a 4e 31 31 45 32 33 33 44 41 46 4b 00 00 00 00
Address 0x30: 00 00 00 ff 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

```

```

Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Midplane      REV 26    750-044636    ABAB9357      Lower Backplane
Jedec Code:   0x7fb0      EEPROM Version: 0x02
P/N:          750-044636    S/N:          ABAB9357
Assembly ID:  0x0b66      Assembly Version: 01.26
Date:         08-28-2012   Assembly Flags:  0x00
Version:      REV 26      CLEI Code:     PROTOXCLEI
ID: Lower Backplane      FRU Model Number: PROTO-ASSEMBLY
Board Information Record:
Address 0x00: ad 01 08 00 2c 21 72 70 a0 00 ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 66 01 1a 52 45 56 20 32 36 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 34 34 36 33 36 00 00
Address 0x20: 53 2f 4e 20 41 42 41 42 39 33 35 37 00 1c 08 07
Address 0x30: dc ff ff ff ad 01 08 00 2c 21 72 70 a0 00 ff ff
Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
Address 0x70: ff ff ff c2 ff ff ff ff ff ff ff ff ff ff ff ff

Midplane 1    REV 01    711-044557    ABAB8643      Upper Backplane
Jedec Code:   0x7fb0      EEPROM Version: 0x01
P/N:          711-044557    S/N:          ABAB8643
Assembly ID:  0x0b65      Assembly Version: 01.01
Date:         07-27-2012   Assembly Flags:  0x00
Version:      REV 01
ID: Upper Backplane
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 0b 65 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 31 31 2d 30 34 34 35 35 37 00 00
Address 0x20: 53 2f 4e 20 41 42 41 42 38 36 34 33 00 1b 07 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

PMP           REV 04    711-032426    ACAJ1677      Power Midplane
Jedec Code:   0x7fb0      EEPROM Version: 0x01
P/N:          711-032426    S/N:          ACAJ1677
Assembly ID:  0x045d      Assembly Version: 01.04
Date:         07-20-2012   Assembly Flags:  0x00
Version:      REV 04
ID: Power Midplane
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 04 5d 01 04 52 45 56 20 30 34 00 00
Address 0x10: 00 00 00 00 37 31 31 2d 30 33 32 34 32 36 00 00
Address 0x20: 53 2f 4e 20 41 43 41 4a 31 36 37 37 00 14 07 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

FPM Board     REV 08    760-044634    ABBV9726      Front Panel Display
Jedec Code:   0x7fb0      EEPROM Version: 0x02
P/N:          760-044634    S/N:          ABBV9726
Assembly ID:  0x0b64      Assembly Version: 01.08
Date:         09-10-2012   Assembly Flags:  0x00
Version:      REV 08      CLEI Code:     IPMYA4EJRA

```

```

ID: Front Panel Display          FRU Model Number: MX2010-CRAFT-S
Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b 64 01 08 52 45 56 20 30 38 00 00
  Address 0x10: 00 00 00 00 37 36 30 2d 30 34 34 36 33 34 00 00
  Address 0x20: 53 2f 4e 20 41 42 42 56 39 37 32 36 00 0a 09 07
  Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 4d 59 41 34 45 4a 52 41 4d
  Address 0x50: 58 32 30 31 30 2d 43 52 41 46 54 2d 53 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 00 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 93 ff ff ff ff ff ff ff ff ff ff ff ff
PSM 0          REV 01   740-045050   1E02224000P   DC 52V Power Supply
Module
Jedec Code:    0x7fb0          EEPROM Version:    0x02
P/N:           740-045050      S/N:              1E02224000P
Assembly ID:   0x0478          Assembly Version:  01.01
Date:          12-06-2012      Assembly Flags:    0x00
Version:       REV 01          CLEI Code:         XXXXXXXXXX
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-HC-DC-S-A
Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 34 35 30 35 30 00 00
  Address 0x20: 31 45 30 32 32 32 34 30 30 30 50 00 00 06 0c 07
  Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 58 58 58 58 58 58 58 58 58 58 4d
  Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 48 43 2d 44 43 2d
  Address 0x60: 53 2d 41 00 00 00 31 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 4a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 1          REV 01   740-045050   1E02224000M   DC 52V Power Supply
Module
Jedec Code:    0x7fb0          EEPROM Version:    0x02
P/N:           740-045050      S/N:              1E02224000M
Assembly ID:   0x0478          Assembly Version:  01.01
Date:          12-06-2012      Assembly Flags:    0x00
Version:       REV 01          CLEI Code:         XXXXXXXXXX
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-HC-DC-S-A
Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 34 35 30 35 30 00 00
  Address 0x20: 31 45 30 32 32 32 34 30 30 30 4d 00 00 06 0c 07
  Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 58 58 58 58 58 58 58 58 58 58 4d
  Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 48 43 2d 44 43 2d
  Address 0x60: 53 2d 41 00 00 00 31 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 4a 00 00 00 00 00 00 00 00 00 00 00 00
...
PDM 0          REV 01   740-045234   1E262250067   DC Power Dist Module
Jedec Code:    0x7fb0          EEPROM Version:    0x02
P/N:           740-045234      S/N:              1E262250067
Assembly ID:   0x047b          Assembly Version:  01.01
Date:          06-28-2012      Assembly Flags:    0x00
Version:       REV 01          CLEI Code:         IPUPAJSKAA
ID: DC Power Dist Module      FRU Model Number: MX2000-PDM-DC-S-A
Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:

```

```

Address 0x00: 7f b0 02 ff 04 7b 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 34 35 32 33 34 00 00
Address 0x20: 31 45 32 36 32 32 35 30 30 36 37 00 00 1c 06 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 55 50 41 4a 53 4b 41 41 4d
Address 0x50: 58 32 30 30 30 2d 50 44 4d 2d 44 43 2d 53 2d 41
Address 0x60: 00 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff
Address 0x70: ff ff ff 89 00 00 00 00 00 00 00 00 00 00 00 00
Routing Engine 0 REV 02 740-041821 9009099704 RE-S-1800x4
Jedec Code: 0x7fb0 EEPROM Version: 0x02
P/N: 740-041821 S/N: 9009099704
Assembly ID: 0x09c0 Assembly Version: 01.02
Date: 03-15-2012 Assembly Flags: 0x00
Version: REV 02
ID: RE-S-1800x4 FRU Model Number: RE-S-1800X4-16G-S
Board Information Record:
Address 0x00: 54 32 30 32 37 44 41 2d 34 34 47 42 23 41 23 00
I2C Hex Data:
Address 0x00: 7f b0 02 ff 09 c0 01 02 52 45 56 20 30 32 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 34 31 38 32 31 00 00
Address 0x20: 39 30 30 39 30 39 39 37 30 34 00 00 00 0f 03 07
Address 0x30: dc ff ff ff 54 32 30 32 37 44 41 2d 34 34 47 42
Address 0x40: 23 41 23 00 01 00 00 00 00 00 00 00 00 00 00 52
Address 0x50: 45 2d 53 2d 31 38 30 30 58 34 2d 31 36 47 2d 53
Address 0x60: 00 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff
Address 0x70: ff ff ff 8c ff ff ff ff ff ff ff ff ff ff ff ff
ad0 3831 MB UGB30SFA4000T1 SFA4000T1 00000651 Compact Flash
ad1 30533 MB UGB94BPH32H0S1-KCI 11000019592 Disk 1
usb0 (addr 1) EHCI root hub 0 Intel uhub0
usb0 (addr 2) product 0x0020 32 vendor 0x8087 uhub1
DIMM 0 SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
DIMM 1 SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
DIMM 2 SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
DIMM 3 SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
Routing Engine 1 REV 02 740-041821 9009099706 RE-S-1800x4
Jedec Code: 0x7fb0 EEPROM Version: 0x02
P/N: 740-041821 S/N: 9009099706
Assembly ID: 0x09c0 Assembly Version: 01.02
Date: 02-23-2012 Assembly Flags: 0x00
Version: REV 02
ID: RE-S-1800x4 FRU Model Number: RE-S-1800X4-16G-S
Board Information Record:
Address 0x00: 54 32 30 32 37 44 41 2d 34 34 47 42 23 41 23 00
I2C Hex Data:
Address 0x00: 7f b0 02 ff 09 c0 01 02 52 45 56 20 30 32 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 34 31 38 32 31 00 00
Address 0x20: 39 30 30 39 30 39 39 37 30 36 00 00 00 17 02 07
Address 0x30: dc ff ff ff 54 32 30 32 37 44 41 2d 34 34 47 42
Address 0x40: 23 41 23 00 01 00 00 00 00 00 00 00 00 00 00 52
Address 0x50: 45 2d 53 2d 31 38 30 30 58 34 2d 31 36 47 2d 53
Address 0x60: 00 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff
Address 0x70: ff ff ff 8c ff ff ff ff ff ff ff ff ff ff ff ff
ad0 3998 MB Virtium - TuffDrive VCF P1T0200262860208 114 Compact Flash
ad1 30533 MB UGB94ARF32H0S3-KC UNIGEN-499551-000404 Disk 1
CB 0 REV 13 750-040257 CAAF8436 Control Board
Jedec Code: 0x7fb0 EEPROM Version: 0x02
P/N: 750-040257 S/N: CAAF8436
Assembly ID: 0x0b26 Assembly Version: 01.13
Date: 08-29-2012 Assembly Flags: 0x00
Version: REV 13 CLEI Code: PROTOXCLEI
ID: Control Board FRU Model Number: PROTO-ASSEMBLY

```

Board Information Record:

Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
 Address 0x00: 7f b0 02 ff 0b 26 01 0d 52 45 56 20 31 33 00 00
 Address 0x10: 00 00 00 00 37 35 30 2d 30 34 30 32 35 37 00 00
 Address 0x20: 53 2f 4e 20 43 41 41 46 38 34 33 36 00 1d 08 07
 Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
 Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
 Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
 Address 0x70: ff ff ff c2 ff ff ff ff ff ff ff ff ff ff ff ff

...

SPMB 0 REV 02 711-041855 ABBV3825 PMB Board
 Jedec Code: 0x7fb0 EEPROM Version: 0x01
 P/N: 711-041855 S/N: ABBV3825
 Assembly ID: 0x0b29 Assembly Version: 01.02
 Date: 08-14-2012 Assembly Flags: 0x00
 Version: REV 02
 ID: PMB Board

Board Information Record:

Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
 Address 0x00: 7f b0 01 ff 0b 29 01 02 52 45 56 20 30 32 00 00
 Address 0x10: 00 00 00 00 37 31 31 2d 30 34 31 38 35 35 00 00
 Address 0x20: 53 2f 4e 20 41 42 42 56 33 38 32 35 00 0e 08 07
 Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
 Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 Address 0x70: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00

...

SFB 0 REV 05 711-044466 ABBX5682 Switch Fabric Board
 Jedec Code: 0x7fb0 EEPROM Version: 0x02
 P/N: 711-044466 S/N: ABBX5682
 Assembly ID: 0x0b25 Assembly Version: 01.05
 Date: 09-07-2012 Assembly Flags: 0x00
 Version: REV 05 CLEI Code: PROTOXCLEI
 ID: Switch Fabric Board FRU Model Number: PROTO-ASSEMBLY

Board Information Record:

Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
 Address 0x00: 7f b0 02 ff 0b 25 01 05 52 45 56 20 30 35 00 00
 Address 0x10: 00 00 00 00 37 31 31 2d 30 34 34 34 36 36 00 00
 Address 0x20: 53 2f 4e 20 41 42 42 58 35 36 38 32 00 07 09 07
 Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
 Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
 Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
 Address 0x70: ff ff ff c2 00 00 00 01 00 00 00 00 00 00 00 00

...

FPC 0 REV 09 750-037355 CAAF0924 MPC Type 4-2
 Jedec Code: 0x7fb0 EEPROM Version: 0x02
 P/N: 750-037355 S/N: CAAF0924
 Assembly ID: 0x0b4e Assembly Version: 01.09
 Date: 05-21-2012 Assembly Flags: 0x00
 Version: REV 09 CLEI Code: PROTOXCLEI
 ID: MPC Type 4-2 FRU Model Number: MPC4E-2CGE-8XGE

Board Information Record:

Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
 Address 0x00: 7f b0 02 ff 0b 4e 01 09 52 45 56 20 30 39 00 00

```

Address 0x10: 00 00 00 00 37 35 30 2d 30 33 37 33 35 35 00 00
Address 0x20: 53 2f 4e 20 43 41 41 46 30 39 32 34 00 15 05 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 4d
Address 0x50: 50 43 34 45 2d 32 43 47 45 2d 38 58 47 45 00 00
Address 0x60: 00 00 00 00 00 00 30 39 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff c6 ff ff ff ff ff ff ff ff ff ff ff ff
CPU          REV 08    711-035209    CAAB9842          HMPC PMB 2G
Jedec Code:  0x7fb0          EEPROM Version:  0x01
P/N:         711-035209      S/N:         CAAB9842
Assembly ID: 0x0b04          Assembly Version: 01.08
Date:        05-17-2012      Assembly Flags: 0x00
Version:     REV 08
ID: HMPC PMB 2G
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 0b 04 01 08 52 45 56 20 30 38 00 00
Address 0x10: 00 00 00 00 37 31 31 2d 30 33 35 32 30 39 00 00
Address 0x20: 53 2f 4e 20 43 41 41 42 39 38 34 32 00 11 05 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00
PIC 0          BUILTIN      BUILTIN      4x10GE SFP
Jedec Code:  0x0000          EEPROM Version:  0x00
P/N:         BUILTIN        S/N:         BUILTIN
Assembly ID: 0x0a53          Assembly Version: 00.00
Date:        00-00-0000      Assembly Flags: 0x00
ID: 4x10GE SFP
Board Information Record:
Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
Address 0x00: 00 00 00 00 0a 53 00 00 00 00 00 00 00 00 00 00
Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 4d 58 43 00
Address 0x20: 42 55 49 4c 54 49 4e 00 4d 58 43 00 00 00 00 00
Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 00 c0 02 ae 64 00 00 00 00 0a 52 00 00
Xcvr 0      REV 01    740-021308    19T511101656      SFP+-10G-SR
Xcvr 1      REV 01    740-031980    AMA04RU           SFP+-10G-SR
Xcvr 2      REV 01    740-031980    193363A00558      SFP+-10G-SR
Xcvr 3      REV 01    740-031980    B10M00202         SFP+-10G-SR
...
ADC 0      REV 13    750-043596    ABBX5532          Adapter Card
Jedec Code: 0x7fb0          EEPROM Version:  0x02
P/N:        750-043596      S/N:         ABBX5532
Assembly ID: 0x0b3d          Assembly Version: 01.13
Date:       09-12-2012      Assembly Flags: 0x00
Version:    REV 13          CLEI Code:     IPUCBA8CAA
ID: Adapter Card          FRU Model Number: MX2000-LC-ADAPTER
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 3d 01 0d 52 45 56 20 31 33 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 34 33 35 39 36 00 00
Address 0x20: 53 2f 4e 20 41 42 42 58 35 35 33 32 00 0c 09 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff

```

```

Address 0x40: ff ff ff ff 01 49 50 55 43 42 41 38 43 41 41 4d
Address 0x50: 58 32 30 30 30 2d 4c 43 2d 41 44 41 50 54 45 52
Address 0x60: 00 00 00 00 00 00 41 00 00 ff ff ff ff ff ff
Address 0x70: ff ff ff 3a 00 00 00 00 00 00 00 00 00 00 00
...

```

show chassis hardware models (MX2010 Router)

```
user@host > show chassis hardware models
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	FRU model number
FPM Board	REV 06	711-032349	ZX8744	711-032349
PSM 4	REV 0C	740-033727	VK00254	000000000000000000000000
PSM 5	REV 0B	740-033727	VG00015	000000000000000000000000
PSM 6	REV 0B	740-033727	VH00097	000000000000000000000000
PSM 7	REV 0C	740-033727	VJ00151	000000000000000000000000
PSM 8	REV 0C	740-033727	VJ00149	000000000000000000000000
PDM 0	REV 0B	740-038109	WA00008	
PDM 1	REV 0B	740-038109	WA00014	
Routing Engine 0	REV 02	740-041821	9009094134	RE-S-1800X4-16G-S
Routing Engine 1	REV 02	740-041821	9009094141	RE-S-1800X4-16G-S
CB 0	REV 08	750-040257	CAAB3491	750-040257
CB 1	REV 08	750-040257	CAAB3489	750-040257
SFB 0	REV 06	711-032385	ZV1828	711-032385
SFB 1	REV 07	711-032385	ZZ2568	711-032385
SFB 2	REV 07	711-032385	ZZ2563	711-032385
SFB 3	REV 07	711-032385	ZZ2564	711-032385
SFB 4	REV 07	711-032385	ZZ2580	711-032385
SFB 5	REV 07	711-032385	ZZ2579	711-032385
SFB 6	REV 07	711-032385	CAAB4882	711-044170
SFB 7	REV 07	711-032385	CAAB4898	711-044170
FPC 0	REV 33	750-028467	CAAB1919	MPC-3D-16XGE-SFPP
FPC 1	REV 21	750-033205	ZG5027	MX-MPC3-3D
MIC 0	REV 03	750-033307	ZV6299	MIC3-3D-10XGE-SFPP
MIC 1	REV 03	750-033307	ZV6268	MIC3-3D-10XGE-SFPP
FPC 8	REV 22	750-031089	ZT9746	MX-MPC2-3D
MIC 0	REV 26	750-028392	ABBS1150	MIC-3D-20GE-SFP
MIC 1	REV 26	750-028387	ABBR9582	MIC-3D-4XGE-XFP
FPC 9	REV 11	750-036284	ZL3591	MPCE-3D-16XGE-SFPP
ADC 0	REV 05	750-043596	CAAC2073	750-043596
ADC 1	REV 01	750-043596	ZV4117	750-043596
ADC 8	REV 01	750-043596	ZV4107	750-043596
ADC 9	REV 02	750-043596	ZW1555	750-043596
Fan Tray 0	REV 2A	760-046960	ACAY0015	
Fan Tray 1	REV 2A	760-046960	ACAY0019	
Fan Tray 2	REV 2A	760-046960	ACAY0020	
Fan Tray 3	REV 2A	760-046960	ACAY0021	

show chassis hardware clei-models (MX2010 Routers)

```
user@host > show chassis hardware clei-models
```

```
Hardware inventory:
```

Item	Version	Part number	CLEI code	FRU model number
FPM Board	REV 06	711-032349	PROTOXCLEI	711-032349
PSM 4	REV 0C	740-033727	0000000000	000000000000000000000000
PSM 5	REV 0B	740-033727	0000000000	000000000000000000000000
PSM 6	REV 0B	740-033727	0000000000	000000000000000000000000
PSM 7	REV 0C	740-033727	0000000000	000000000000000000000000
PSM 8	REV 0C	740-033727	0000000000	000000000000000000000000
PDM 0	REV 0B	740-038109		
PDM 1	REV 0B	740-038109		

Routing Engine 0	REV 02	740-041821		RE-S-1800X4-16G-S
Routing Engine 1	REV 02	740-041821		RE-S-1800X4-16G-S
CB 0	REV 08	750-040257	PROTOXCLEI	750-040257
CB 1	REV 08	750-040257	PROTOXCLEI	750-040257
SFB 0	REV 06	711-032385	PROTOXCLEI	711-032385
SFB 1	REV 07	711-032385	PROTOXCLEI	711-032385
SFB 2	REV 07	711-032385	PROTOXCLEI	711-032385
SFB 3	REV 07	711-032385	PROTOXCLEI	711-032385
SFB 4	REV 07	711-032385	PROTOXCLEI	711-032385
SFB 5	REV 07	711-032385	PROTOXCLEI	711-0323856
SFB 6	REV 07	711-032385	PROTOXCLEI	711-044170
SFB 7	REV 07	711-032385	PROTOXCLEI	711-044170
FPC 0	REV 33	750-028467		MPC-3D-16XGE-SFPP
FPC 1	REV 21	750-033205		MX-MPC3-3D
MIC 0	REV 03	750-033307	PROTOXCLEI	MIC3-3D-10XGE-SFPP
MIC 1	REV 03	750-033307	PROTOXCLEI	MIC3-3D-10XGE-SFPP
FPC 8	REV 22	750-031089	COUIBAYBAA	MX-MPC2-3D
MIC 0	REV 26	750-028392	COUIA15BAA	MIC-3D-20GE-SFP
MIC 1	REV 26	750-028387	COUIA16BAA	MIC-3D-4XGE-XFP
FPC 9	REV 11	750-036284	CMUIACGBAA	MPCE-3D-16XGE-SFPP
ADC 0	REV 05	750-043596	PROTOXCLEI	750-043596
ADC 1	REV 01	750-043596	PROTOXCLEI	750-043596
ADC 8	REV 01	750-043596	PROTOXCLEI	750-043596
ADC 9	REV 02	750-043596	PROTOXCLEI	750-043596
Fan Tray 0	REV 2A	760-046960		
Fan Tray 1	REV 2A	760-046960		
Fan Tray 2	REV 2A	760-046960		
Fan Tray 3	REV 2A	760-046960		

show chassis hardware (MX2010 Routers with MPC6E and OTN MIC)

```

user@host> show chassis hardware
Hardware inventory:

```

Item	Version	Part number	Serial number	Description
Chassis			JN11C9AFEAFK	MX2010
Midplane	REV 35	750-044636	ABAB9188	Lower Backplane
Midplane 1	REV 02	711-044557	ABAB8729	Upper Backplane
PMP	REV 04	711-032426	ACAJ2432	Power Midplane
FPM Board	REV 09	760-044634	ABCA4314	Front Panel Display
PSM 0	REV 01	740-050037	1EDB321015C	DC 52V Power Supply
Module				
PSM 1	REV 01	740-050037	1EDB321015J	DC 52V Power Supply
Module				
PSM 2	REV 01	740-050037	1EDB32000K8	DC 52V Power Supply
Module				
PSM 3	REV 01	740-050037	1EDB32101JW	DC 52V Power Supply
Module				
PSM 4	REV 01	740-050037	1EDB321015G	DC 52V Power Supply
Module				
PSM 5	REV 01	740-050037	1EDB32101HH	DC 52V Power Supply
Module				
PSM 6	REV 01	740-050037	1EDB32101HD	DC 52V Power Supply
Module				
PSM 7	REV 01	740-050037	1EDB321015F	DC 52V Power Supply
Module				
PSM 8	REV 01	740-050037	1EDB321015B	DC 52V Power Supply
Module				
PDM 0	REV 03	740-045234	1EFA3220433	DC Power Dist Module
PDM 1	REV 03	740-045234	1EFA3220425	DC Power Dist Module
Routing Engine 0	REV 02	740-041821	9009115685	RE-S-1800x4
Routing Engine 1	REV 02	740-041821	9009099711	RE-S-1800x4

CB 0	REV 23	750-040257	CABE8395	Control Board
CB 1	REV 12	750-040257	CAAD9499	Control Board
SPMB 0	REV 02	711-041855	ABCG8426	PMB Board
SPMB 1	REV 02	711-041855	ABBS1481	PMB Board
SFB 0	REV 06	711-044466	ABCD5013	Switch Fabric Board
SFB 1	REV 06	711-044466	ABCD5160	Switch Fabric Board
SFB 2	REV 06	711-044466	ABCD5175	Switch Fabric Board
SFB 3	REV 06	711-044466	ABCD4938	Switch Fabric Board
SFB 4	REV 06	711-044466	ABCD4944	Switch Fabric Board
SFB 5	REV 06	711-044466	ABCD4968	Switch Fabric Board
SFB 6	REV 06	711-044466	ABCD5267	Switch Fabric Board
SFB 7	REV 06	711-044466	ABCD4997	Switch Fabric Board
FPC 0	REV 59	750-044130	ABCT7676	MPC6E 3D
CPU	REV 10	711-045719	ABCK8527	RMPD PMB
XLM 0	REV 13	711-046638	ABCT7810	MPC6E XL
XLM 1	REV 13	711-046638	ABCT7811	MPC6E XL
FPC 2	REV 27	750-033205	ZL6014	MPCE Type 3 3D
CPU	REV 07	711-035209	ZK9068	HMPD PMB 2G
MIC 0	REV 14	750-033196	CAAW9214	1X100GE CXP
PIC 0		BUILTIN	BUILTIN	1X100GE CXP
Xcvt 0	REV 01	740-046563	XC49FC030	CFP2-100G-SR10
MIC 1	REV 18	750-033199	CAAC3231	1X100GE CFP
PIC 2		BUILTIN	BUILTIN	1X100GE CFP
FPC 3	REV 59	750-044130	ABCT7682	MPC6E 3D
CPU	REV 10	711-045719	ABCK8531	RMPD PMB
XLM 0	REV 13	711-046638	ABCT7818	MPC6E XL
XLM 1	REV 13	711-046638	ABCT7819	MPC6E XL
FPC 4	REV 33	750-044130	ABBY9278	MPC6E 3D
CPU	REV 09	711-045719	ABBY8677	RMPD PMB
XLM 0	REV 06.2.00	711-046638	ABBY8844	MPC6E XL
XLM 1	REV 06.2.00	711-046638	ABBY8830	MPC6E XL
FPC 5	REV 59	750-044130	ABCT7675	MPC6E 3D
CPU	REV 10	711-045719	ABCK8526	RMPD PMB
XLM 0	REV 13	711-046638	ABCT7808	MPC6E XL
XLM 1	REV 13	711-046638	ABCT7809	MPC6E XL
FPC 6	REV 30	750-028467	ZM4986	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ZP6541	AMPD PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvt 0	REV 01	740-021308	AQ43GAC	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvt 0	REV 01	740-031980	ALM0A6D	SFP+-10G-SR
Xcvt 1	REV 01	740-031980	AQFORB3	SFP+-10G-SR
Xcvt 2	REV 01	740-031980	153363A00333	SFP+-10G-SR
Xcvt 3	REV 01	740-021308	AN10KYE	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvt 0	REV 01	740-021308	APK04YM	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvt 0	REV 01	740-031980	AQFOH44	SFP+-10G-SR
FPC 8	REV 38	750-031090	CABF7313	MPC Type 2 3D EQ
CPU	REV 08	711-030884	CABE6727	MPC PMB 2G
MIC 0	REV 18	750-028380	YK8253	3D 2x 10GE XFP
PIC 0		BUILTIN	BUILTIN	1x 10GE XFP
Xcvt 0	REV 03	740-014289	AD1148M00TP	XFP-10G-SR
PIC 1		BUILTIN	BUILTIN	1x 10GE XFP
QXM 0	REV 06	711-028408	CABC5614	MPC QXM
QXM 1	REV 06	711-028408	CABC5550	MPC QXM
FPC 9	REV 39	750-044130	ABCK1652	MPC6E 3D
CPU	REV 09	711-045719	ABCK1655	RMPD PMB
MIC 0	REV 09	750-049457	ABCP1230	2X100GE CFP2 OTN
PIC 0		BUILTIN	BUILTIN	2X100GE CFP2 OTN
Xcvt 0		NON-JNPR	37300222WP0002	CFP2-100G-LR4-D

Xcvr 1		NON-JNPR	FD46F001Y	CFP2-100G-SR10
MIC 1	REV 07	750-049457	ABCV6662	2X100GE CFP2 OTN
PIC 1		BUILTIN	BUILTIN	2X100GE CFP2 OTN
Xcvr 0		NON-JNPR	UQD0014	CFP2-100G-LR4-D
Xcvr 1		NON-JNPR	J13J68335	CFP2-100G-LR4-D
XLM 0	REV 07.2.00	711-046638	ABCK5491	MPC6E XL
XLM 1	REV 07.2.00	711-046638	ABCK5475	MPC6E XL
ADC 1	REV 17	750-043596	ABCG9023	Adapter Card
ADC 2	REV 01	750-043596	ZV4079	Adapter Card
ADC 6	REV 17	750-043596	ABCG8866	Adapter Card
ADC 8	REV 17	750-043596	ABCA8993	Adapter Card
Fan Tray 0	REV 06	760-046960	ACAY0354	172mm FanTray - 6 Fans
Fan Tray 1	REV 06	760-046960	ACAY0831	172mm FanTray - 6 Fans
Fan Tray 2	REV 06	760-046960	ACAY0892	172mm FanTray - 6 Fans
Fan Tray 3	REV 06	760-046960	ACAY0839	172mm FanTray - 6 Fans

show chassis hardware detail (MX2010 Routers with MPC6E and OTN MIC)

```

user@host> show chassis hardware detail
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               JN11C9AFEAFK  MX2010
Midplane      REV 35    750-044636   ABAB9188      Lower Backplane
Midplane 1    REV 02    711-044557   ABAB8729      Upper Backplane
PMP            REV 04    711-032426   ACAJ2432      Power Midplane
FPM Board      REV 09    760-044634   ABCA4314      Front Panel Display
PSM 0          REV 01    740-050037   1EDB321015C   DC 52V Power Supply
Module
PSM 1          REV 01    740-050037   1EDB321015J   DC 52V Power Supply
Module
PSM 2          REV 01    740-050037   1EDB32000K8    DC 52V Power Supply
Module
PSM 3          REV 01    740-050037   1EDB32101JW    DC 52V Power Supply
Module
PSM 4          REV 01    740-050037   1EDB321015G    DC 52V Power Supply
Module
PSM 5          REV 01    740-050037   1EDB32101HH    DC 52V Power Supply
Module
PSM 6          REV 01    740-050037   1EDB32101HD    DC 52V Power Supply
Module
PSM 7          REV 01    740-050037   1EDB321015F    DC 52V Power Supply
Module
PSM 8          REV 01    740-050037   1EDB321015B    DC 52V Power Supply
Module
PDM 0          REV 03    740-045234   1EFA3220433    DC Power Dist Module
PDM 1          REV 03    740-045234   1EFA3220425    DC Power Dist Module
Routing Engine 0 REV 02    740-041821   9009115685      RE-S-1800x4
  ad0    3998 MB  Virtium - TuffDrive VCF P1T0200274310822 191 Compact Flash
  ad1    30533 MB UGB94BPH32H0S1-KCI 11000043190      Disk 1
  usb0 (addr 1) EHCI root hub 0      Intel          uhub0
  usb0 (addr 2) product 0x0020 32 vendor 0x8087      uhub1
  DIMM 0      VL31B5263F-F8SD DIE REV-0 PCB REV-0      MFR ID-ce80
  DIMM 1      VL31B5263F-F8SD DIE REV-0 PCB REV-0      MFR ID-ce80
  DIMM 2      VL31B5263F-F8SD DIE REV-0 PCB REV-0      MFR ID-ce80
  DIMM 3      VL31B5263F-F8SD DIE REV-0 PCB REV-0      MFR ID-ce80
Routing Engine 1 REV 02    740-041821   9009099711      RE-S-1800x4
  ad0    3998 MB  Virtium - TuffDrive VCF P1T0200262860208 30 Compact Flash
  ad1    30533 MB UGB94ARF32H0S3-KC  UNIGEN-499551-000146 Disk 1
  CB 0      REV 23    750-040257   CABE8395      Control Board
  CB 1      REV 12    750-040257   CAAD9499      Control Board
  SPMB 0      REV 02    711-041855   ABCG8426      PMB Board

```

SPMB 1	REV 02	711-041855	ABBS1481	PMB Board
SFB 0	REV 06	711-044466	ABCD5013	Switch Fabric Board
SFB 1	REV 06	711-044466	ABCD5160	Switch Fabric Board
SFB 2	REV 06	711-044466	ABCD5175	Switch Fabric Board
SFB 3	REV 06	711-044466	ABCD4938	Switch Fabric Board
SFB 4	REV 06	711-044466	ABCD4944	Switch Fabric Board
SFB 5	REV 06	711-044466	ABCD4968	Switch Fabric Board
SFB 6	REV 06	711-044466	ABCD5267	Switch Fabric Board
SFB 7	REV 06	711-044466	ABCD4997	Switch Fabric Board
FPC 0	REV 59	750-044130	ABCT7676	MPC6E 3D
CPU	REV 10	711-045719	ABCK8527	RMPK PMB
XLM 0	REV 13	711-046638	ABCT7810	MPC6E XL
XLM 1	REV 13	711-046638	ABCT7811	MPC6E XL
FPC 2	REV 27	750-033205	ZL6014	MPCE Type 3 3D
CPU	REV 07	711-035209	ZK9068	HMPK PMB 2G
MIC 0	REV 14	750-033196	CAAW9214	1X100GE CXP
PIC 0		BUILTIN	BUILTIN	1X100GE CXP
Xcvr 0	REV 01	740-046563	XC49FC030	CFP2-100G-SR10
MIC 1	REV 18	750-033199	CAAC3231	1X100GE CFP
PIC 2		BUILTIN	BUILTIN	1X100GE CFP
FPC 3	REV 59	750-044130	ABCT7682	MPC6E 3D
CPU	REV 10	711-045719	ABCK8531	RMPK PMB
XLM 0	REV 13	711-046638	ABCT7818	MPC6E XL
XLM 1	REV 13	711-046638	ABCT7819	MPC6E XL
FPC 4	REV 33	750-044130	ABBY9278	MPC6E 3D
CPU	REV 09	711-045719	ABBY8677	RMPK PMB
XLM 0	REV 06.2.00	711-046638	ABBY8844	MPC6E XL
XLM 1	REV 06.2.00	711-046638	ABBY8830	MPC6E XL
FPC 5	REV 59	750-044130	ABCT7675	MPC6E 3D
CPU	REV 10	711-045719	ABCK8526	RMPK PMB
XLM 0	REV 13	711-046638	ABCT7808	MPC6E XL
XLM 1	REV 13	711-046638	ABCT7809	MPC6E XL
FPC 6	REV 30	750-028467	ZM4986	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ZP6541	AMPK PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AQ43GAC	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	ALMOA6D	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AQFORB3	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	153363A00333	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AN10KYE	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	APK04YM	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AQFOH44	SFP+-10G-SR
FPC 8	REV 38	750-031090	CABF7313	MPC Type 2 3D EQ
CPU	REV 08	711-030884	CABE6727	MPC PMB 2G
MIC 0	REV 18	750-028380	YK8253	3D 2x 10GE XFP
PIC 0		BUILTIN	BUILTIN	1x 10GE XFP
Xcvr 0	REV 03	740-014289	AD1148M00TP	XFP-10G-SR
PIC 1		BUILTIN	BUILTIN	1x 10GE XFP
QXM 0	REV 06	711-028408	CABC5614	MPC QXM
QXM 1	REV 06	711-028408	CABC5550	MPC QXM
FPC 9	REV 39	750-044130	ABCK1652	MPC6E 3D
CPU	REV 09	711-045719	ABCK1655	RMPK PMB
MIC 0	REV 09	750-049457	ABCP1230	2X100GE CFP2 OTN
PIC 0		BUILTIN	BUILTIN	2X100GE CFP2 OTN
Xcvr 0		NON-JNPR	37300222WP0002	CFP2-100G-LR4-D
Xcvr 1		NON-JNPR	FD46F001Y	CFP2-100G-SR10
MIC 1	REV 07	750-049457	ABCV6662	2X100GE CFP2 OTN
PIC 1		BUILTIN	BUILTIN	2X100GE CFP2 OTN

Xcvr 0	NON-JNPR	UQD0014	CFP2-100G-LR4-D
Xcvr 1	NON-JNPR	J13J68335	CFP2-100G-LR4-D
XLM 0	REV 07.2.00	711-046638	ABCK5491
XLM 1	REV 07.2.00	711-046638	ABCK5475
ADC 1	REV 17	750-043596	ABCG9023
ADC 2	REV 01	750-043596	ZV4079
ADC 6	REV 17	750-043596	ABCG8866
ADC 8	REV 17	750-043596	ABCA8993
Fan Tray 0	REV 06	760-046960	ACAY0354
Fan Tray 1	REV 06	760-046960	ACAY0831
Fan Tray 2	REV 06	760-046960	ACAY0892
Fan Tray 3	REV 06	760-046960	ACAY0839

show chassis hardware extensive (MX2010 Routers with MPC6E and OTN MIC)

```

user@host> show chassis hardware extensive
Hardware inventory:
Item              Version  Part number  Serial number  Description
Chassis
Jedec Code:       0x7fb0          EEPROM Version: 0x02
S/N:              JN11C9AFEAFK
Assembly ID:      0x0557          Assembly Version: 00.00
Date:             00-00-0000      Assembly Flags:  0x00
ID: MX2010
Board Information Record:
Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
Address 0x00: 7f b0 02 ff 05 57 00 00 00 00 00 00 00 00 00 00
Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x20: 4a 4e 31 31 43 39 41 46 45 41 46 4b 00 00 00 00
Address 0x30: 00 00 00 ff 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Midplane          REV 35    750-044636  ABAB9188      Lower Backplane
Jedec Code:       0x7fb0          EEPROM Version: 0x02
P/N:              750-044636      S/N:           ABAB9188
Assembly ID:      0x0b66          Assembly Version: 01.35
Date:             06-21-2013      Assembly Flags: 0x00
Version:          REV 35          CLEI Code:     IPMU810ARA
ID: Lower Backplane              FRU Model Number: CHAS-BP-MX2010-S
Board Information Record:
Address 0x00: ad 01 08 00 3c 8a b0 38 68 00 ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 66 01 23 52 45 56 20 33 35 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 34 34 36 33 36 00 00
Address 0x20: 53 2f 4e 20 41 42 41 42 39 31 38 38 00 15 06 07
Address 0x30: dd ff ff ff ad 01 08 00 3c 8a b0 38 68 00 ff ff
Address 0x40: ff ff ff ff 01 49 50 4d 55 38 31 30 41 52 41 43
Address 0x50: 48 41 53 2d 42 50 2d 4d 58 32 30 31 30 2d 53 00
Address 0x60: 00 00 00 00 00 00 30 36 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff f8 ff ff ff ff ff ff ff ff ff ff ff ff
Midplane 1        REV 02    711-044557  ABAB8729      Upper Backplane
Jedec Code:       0x7fb0          EEPROM Version: 0x01
P/N:              711-044557      S/N:           ABAB8729
Assembly ID:      0x0b65          Assembly Version: 01.02
Date:             03-21-2013      Assembly Flags: 0x00
Version:          REV 02
ID: Upper Backplane
Board Information Record:

```

```

Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 0b 65 01 02 52 45 56 20 30 32 00 00
Address 0x10: 00 00 00 00 37 31 31 2d 30 34 34 35 35 37 00 00
Address 0x20: 53 2f 4e 20 41 42 41 42 38 37 32 39 00 15 03 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 00
Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 ff ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

PMP REV 04 711-032426 ACAJ2432 Power Midplane
Jedec Code: 0x7fb0 EEPROM Version: 0x01
P/N: 711-032426 S/N: ACAJ2432
Assembly ID: 0x045d Assembly Version: 01.04
Date: 03-28-2013 Assembly Flags: 0x00
Version: REV 04
ID: Power Midplane
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 04 5d 01 04 52 45 56 20 30 34 00 00
Address 0x10: 00 00 00 00 37 31 31 2d 30 33 32 34 32 36 00 00
Address 0x20: 53 2f 4e 20 41 43 41 4a 32 34 33 32 00 1c 03 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 00
Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 ff ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

FPM Board REV 09 760-044634 ABCA4314 Front Panel Display
Jedec Code: 0x7fb0 EEPROM Version: 0x02
P/N: 760-044634 S/N: ABCA4314
Assembly ID: 0x0b64 Assembly Version: 01.09
Date: 03-28-2013 Assembly Flags: 0x00
Version: REV 09 CLEI Code: IPMYA4EJRA
ID: Front Panel Display FRU Model Number: MX2010-CRAFT-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 64 01 09 52 45 56 20 30 39 00 00
Address 0x10: 00 00 00 00 37 36 30 2d 30 34 34 36 33 34 00 00
Address 0x20: 53 2f 4e 20 41 42 43 41 34 33 31 34 00 1c 03 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 4d 59 41 34 45 4a 52 41 4d
Address 0x50: 58 32 30 31 30 2d 43 52 41 46 54 2d 53 00 00 00
Address 0x60: 00 00 00 00 00 00 41 00 00 ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff 93 ff ff ff ff ff ff ff ff ff ff ff ff

PSM 0 REV 01 740-050037 1EDB321015C DC 52V Power Supply
Module
Jedec Code: 0x7fb0 EEPROM Version: 0x02
P/N: 740-050037 S/N: 1EDB321015C
Assembly ID: 0x0478 Assembly Version: 01.01
Date: 05-28-2013 Assembly Flags: 0x00
Version: REV 01 CLEI Code: IPUPAKRKAA
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
Address 0x20: 31 45 44 42 33 32 31 30 31 35 43 00 00 1c 05 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff

```

```

Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 1          REV 01   740-050037   1EDB321015J   DC 52V Power Supply
Module
Jedec Code:   0x7fb0          EEPROM Version:   0x02
P/N:          740-050037      S/N:             1EDB321015J
Assembly ID:  0x0478          Assembly Version: 01.01
Date:         05-28-2013      Assembly Flags:   0x00
Version:      REV 01          CLEI Code:        IPUPAKRKAA
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
Address 0x20: 31 45 44 42 33 32 31 30 31 35 4a 00 00 1c 05 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 2          REV 01   740-050037   1EDB32000K8   DC 52V Power Supply
Module
Jedec Code:   0x7fb0          EEPROM Version:   0x02
P/N:          740-050037      S/N:             1EDB32000K8
Assembly ID:  0x0478          Assembly Version: 01.01
Date:         05-23-2013      Assembly Flags:   0x00
Version:      REV 01          CLEI Code:        IPUPAKRKAA
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
Address 0x20: 31 45 44 42 33 32 30 30 30 4b 38 00 00 17 05 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 3          REV 01   740-050037   1EDB32101JW   DC 52V Power Supply
Module
Jedec Code:   0x7fb0          EEPROM Version:   0x02
P/N:          740-050037      S/N:             1EDB32101JW
Assembly ID:  0x0478          Assembly Version: 01.01
Date:         05-30-2013      Assembly Flags:   0x00
Version:      REV 01          CLEI Code:        IPUPAKRKAA
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
Address 0x20: 31 45 44 42 33 32 31 30 31 4a 57 00 00 1e 05 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00

```

```

PSM 4          REV 01  740-050037  1EDB321015G      DC 52V Power Supply
Module
Jedec Code:    0x7fb0          EEPROM Version:    0x02
P/N:           740-050037      S/N:              1EDB321015G
Assembly ID:   0x0478          Assembly Version:  01.01
Date:          05-28-2013      Assembly Flags:    0x00
Version:       REV 01          CLEI Code:         IPUPAKRKAA
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
Address 0x20: 31 45 44 42 33 32 31 30 31 35 47 00 00 1c 05 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 5          REV 01  740-050037  1EDB32101HH      DC 52V Power Supply
Module
Jedec Code:    0x7fb0          EEPROM Version:    0x02
P/N:           740-050037      S/N:              1EDB32101HH
Assembly ID:   0x0478          Assembly Version:  01.01
Date:          05-30-2013      Assembly Flags:    0x00
Version:       REV 01          CLEI Code:         IPUPAKRKAA
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
Address 0x20: 31 45 44 42 33 32 31 30 31 48 48 00 00 1e 05 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 6          REV 01  740-050037  1EDB32101HD      DC 52V Power Supply
Module
Jedec Code:    0x7fb0          EEPROM Version:    0x02
P/N:           740-050037      S/N:              1EDB32101HD
Assembly ID:   0x0478          Assembly Version:  01.01
Date:          05-30-2013      Assembly Flags:    0x00
Version:       REV 01          CLEI Code:         IPUPAKRKAA
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
Address 0x20: 31 45 44 42 33 32 31 30 31 48 44 00 00 1e 05 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 7          REV 01  740-050037  1EDB321015F      DC 52V Power Supply
Module
Jedec Code:    0x7fb0          EEPROM Version:    0x02
P/N:           740-050037      S/N:              1EDB321015F

```



```

Assembly ID: 0x0478      Assembly Version: 01.01
Date: 05-28-2013      Assembly Flags: 0x00
Version: REV 01      CLEI Code: IPUPAKRKAA
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S
Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  I2C Hex Data:
    Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
    Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
    Address 0x20: 31 45 44 42 33 32 31 30 31 35 46 00 00 1c 05 07
    Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff
    Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
    Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
    Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
    Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 8      REV 01      740-050037      1EDB321015B      DC 52V Power Supply
Module
Jedec Code: 0x7fb0      EEPROM Version: 0x02
P/N: 740-050037      S/N: 1EDB321015B
Assembly ID: 0x0478      Assembly Version: 01.01
Date: 05-28-2013      Assembly Flags: 0x00
Version: REV 01      CLEI Code: IPUPAKRKAA
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S
Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  I2C Hex Data:
    Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
    Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
    Address 0x20: 31 45 44 42 33 32 31 30 31 35 42 00 00 1c 05 07
    Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff
    Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
    Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
    Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
    Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PDM 0      REV 03      740-045234      1EFA3220433      DC Power Dist Module
Jedec Code: 0x7fb0      EEPROM Version: 0x02
P/N: 740-045234      S/N: 1EFA3220433
Assembly ID: 0x047b      Assembly Version: 01.03
Date: 05-30-2013      Assembly Flags: 0x00
Version: REV 03      CLEI Code: IPUPAJSKAA
ID: DC Power Dist Module FRU Model Number: MX2000-PDM-DC-S
Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  I2C Hex Data:
    Address 0x00: 7f b0 02 ff 04 7b 01 03 52 45 56 20 30 33 00 00
    Address 0x10: 00 00 00 00 37 34 30 2d 30 34 35 32 33 34 00 00
    Address 0x20: 31 45 46 41 33 32 32 30 34 33 33 00 00 1e 05 07
    Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff
    Address 0x40: ff ff ff ff 01 49 50 55 50 41 4a 53 4b 41 41 4d
    Address 0x50: 58 32 30 30 30 2d 50 44 4d 2d 44 43 2d 53 00 00
    Address 0x60: 00 00 00 00 00 00 31 30 33 ff ff ff ff ff ff ff
    Address 0x70: ff ff ff 1d 00 00 00 00 00 00 00 00 00 00 00 00
PDM 1      REV 03      740-045234      1EFA3220425      DC Power Dist Module
Jedec Code: 0x7fb0      EEPROM Version: 0x02
P/N: 740-045234      S/N: 1EFA3220425
Assembly ID: 0x047b      Assembly Version: 01.03
Date: 05-30-2013      Assembly Flags: 0x00
Version: REV 03      CLEI Code: IPUPAJSKAA
ID: DC Power Dist Module FRU Model Number: MX2000-PDM-DC-S
Board Information Record:

```

Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

..

show chassis hardware (MX2020 Router)

user@host > show chassis hardware

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN11E2227AFJ	MX2020
Midplane	REV 27	750-040240	ABAB9384	Lower Power Midplane
Midplane 1	REV 04	711-032386	ABAB9386	Upper Backplane
PMP 1	REV 05	711-032428	ACA11579	Upper Power Midplane
PMP 0	REV 04	711-032426	ACA11524	Lower Power Midplane
FPM Board	REV 06	760-040242	ABBT8837	Front Panel Display
PSM 0	REV 01	740-045050	1E022240056	DC 52V Power Supply
Module				
PSM 1	REV 01	740-045050	1E022240054	DC 52V Power Supply
Module				
PSM 2	REV 01	740-045050	1E02224005H	DC 52V Power Supply
Module				
PSM 3	REV 01	740-045050	1E022240053	DC 52V Power Supply
Module				
PSM 4	REV 01	740-045050	1E02224004K	DC 52V Power Supply
Module				
PSM 7	REV 01	740-045050	1E02224006W	DC 52V Power Supply
Module				
PSM 8	REV 01	740-045050	1E022240062	DC 52V Power Supply
Module				
PSM 9	REV 01	740-045050	1E02224005B	DC 52V Power Supply
Module				
PSM 10	REV 01	740-045050	1E02224005A	DC 52V Power Supply
Module				
PSM 11	REV 01	740-045050	1E022240052	DC 52V Power Supply
Module				
PSM 12	REV 01	740-045050	1E022240051	DC 52V Power Supply
Module				
PSM 13	REV 01	740-045050	1E022240058	DC 52V Power Supply
Module				
PSM 14	REV 01	740-045050	1E02224004L	DC 52V Power Supply
Module				
PSM 15	REV 01	740-045050	1E02224005M	DC 52V Power Supply
Module				
PSM 16	REV 01	740-045050	1E02224006S	DC 52V Power Supply
Module				
PSM 17	REV 01	740-045050	1E02224005Z	DC 52V Power Supply
Module				
PDM 0	REV 01	740-045234	1E012150033	DC Power Dist Module
PDM 1	REV 01	740-045234	1E012150027	DC Power Dist Module
PDM 2	REV 01	740-045234	1E012150028	DC Power Dist Module
PDM 3	REV 01	740-045234	1E012150045	DC Power Dist Module
Routing Engine 0	REV 02	740-041821	9009089704	RE-S-1800x4
Routing Engine 1	REV 02	740-041821	9009094138	RE-S-1800x4
CB 0	REV 14	750-040257	CAAF8430	Control Board
CB 1	REV 08	750-040257	CAAB3482	Control Board
SPMB 0	REV 01	711-041855	ZS2290	PMB Board
SPMB 1	REV 02	711-041855	CAAA6141	PMB Board
SFB 0	REV 03	711-044466	ABBV6789	Switch Fabric Board
SFB 1	REV 05	711-044466	ABBX5666	Switch Fabric Board
SFB 2	REV 05	711-044466	ABBX5678	Switch Fabric Board
SFB 3	REV 05	711-044466	ABBX5687	Switch Fabric Board
SFB 4	REV 05	711-044466	ABBX5609	Switch Fabric Board

SFB 5	REV 05	711-044466	ABBX5675	Switch Fabric Board
SFB 6	REV 03	711-044466	ABBV6805	Switch Fabric Board
SFB 7	REV 05	711-044466	ABBX5701	Switch Fabric Board
FPC 0	REV 30	750-028467	ABBN0284	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN0507	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E00990	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E04357	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01327	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E04375	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02760	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02904	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E03963	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E00756	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04418	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01077	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01128	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01253	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E01140	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11F01626	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01075	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E01177	SFP+-10G-USR
FPC 1	REV 30	750-028467	ABBN0208	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBJ1084	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04745	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11F01570	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E04388	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01439	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04739	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11F01869	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01675	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01901	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01346	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11F01288	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01824	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E04312	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02811	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E03847	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01495	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01265	SFP+-10G-USR
FPC 2	REV 30	750-028467	ZM5111	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ZP6607	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LJA	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MFZ	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NKL	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80KF4	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80FBJ	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MM2	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LJV	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NXV	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80N1H	SFP+-10G-SR

Xcvr 1	REV 01	740-031980	AK80NLS	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80FL5	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NL9	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NG2	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80KDU	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80MG1	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80MM0	SFP+-10G-SR
FPC 3	REV 30	750-028467	ABBN0302	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN0495	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01581	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01176	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01251	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E02752	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E00786	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01020	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01023	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E02819	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02812	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11D04437	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01279	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01333	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E00978	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01018	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01784	SFP+-10G-USR
Xcvr 3	REV 01	740-031980	AK80NKP	SFP+-10G-SR
FPC 4	REV 30	750-028467	ABBN0308	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBJ1095	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04305	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01147	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01195	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01743	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01892	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02880	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E00725	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E01057	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02816	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11C04501	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E02764	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E00789	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01250	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02847	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E00787	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E03803	SFP+-10G-USR
FPC 5	REV 30	750-028467	ABBN0316	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBJ1082	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00523	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K01848	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01865	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K00540	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00422	SFP+-10G-SR

Xcvr 1	REV 01	740-031980	B11K00428	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K00423	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K01855	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K01847	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K00526	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K00529	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K00525	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00425	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K00530	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01851	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K00528	SFP+-10G-SR
FPC 6	REV 32	750-028467	ABBN6832	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBK6534	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80MB4	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80FQ6	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80N1F	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NLQ	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80KDR	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80FGJ	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80N5G	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80KD8	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LET	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80N1X	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NRF	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NL2	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80N3D	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MRB	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LEQ	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LER	SFP+-10G-SR
FPC 7	REV 32	750-028467	ABBN6811	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN7288	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NK8	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80LJG	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LBU	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80N21	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LEU	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NLM	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NL6	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LES	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LEN	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80ME0	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LMG	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80MM1	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80MG7	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80KF9	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NRQ	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NLE	SFP+-10G-SR
FPC 8	REV 23	750-028467	YN2977	MPC 3D 16x 10GE
CPU	REV 10	711-029089	YP1856	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	183363A00875	SFP+-10G-SR

Xcvr 1	REV 01	740-031980	183363A00851	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	183363A00772	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	183363A00882	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	183363A00735	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	183363A00169	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	183363A00726	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	183363A00077	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	183363A00168	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	183363A00676	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	183363A00732	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	183363A00091	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	183363A00725	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	183363A00642	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	183363A00871	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	183363A00853	SFP+-10G-SR
FPC 9	REV 32	750-028467	ABBN6798	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBK6556	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	9ZDZ06A00055	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	183363A00239	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AD0915E003K	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AD0915E003A	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80MRC	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NL5	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NKN	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80N3U	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80N1T	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AJ808DJ	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NG4	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80FND	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80FKQ	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NLT	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NKR	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LKM	SFP+-10G-SR
FPC 10	REV 32	750-028467	ABBN6813	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBK6542	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NA3	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NLF	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80MRH	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80KE4	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	973152A00030	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80L9H	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80ME8	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NLR	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NG1	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MCA	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LFC	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LEM	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80N9X	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80LAC	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LFC2	SFP+-10G-SR

Xcvr 3	REV 01	740-031980	AK80N8T	SFP+-10G-SR
FPC 11	REV 30	750-028467	ABBN0281	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN0526	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01326	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E03973	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E00950	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E00674	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E00775	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E04461	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01074	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E02821	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04501	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E00757	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01623	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E01022	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04359	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02751	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E02736	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E01178	SFP+-10G-USR
FPC 12	REV 32	750-028467	ABBN6796	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN7259	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K01856	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K01853	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01863	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02863	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02668	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02881	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A01671	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02627	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02725	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02692	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02730	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03081	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02736	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02568	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02747	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02579	SFP+-10G-SR
FPC 13	REV 30	750-028467	ABBN0270	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBJ0966	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NL1	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NXW	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80KD2	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80FMD	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NKQ	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MGH	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80N38	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NL7	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LEL	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NKD	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80KCY	SFP+-10G-SR

Xcvr 3	REV 01	740-031980	AK80LHK	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80M5J	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MBE	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NLG	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LFH	SFP+-10G-SR
FPC 14	REV 32	750-028467	ABBN6790	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBK6515	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LZM	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MCC	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80KCM	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80KE0	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021310	C10F99155	SFP+-10G-LRM
Xcvr 1	REV 01	740-021310	C10F99049	SFP+-10G-LRM
Xcvr 2	REV 01	740-021310	C10F99128	SFP+-10G-LRM
Xcvr 3	REV 01	740-021310	C10F99169	SFP+-10G-LRM
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LF3	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02597	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A03060	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03057	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LEX	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80FEU	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80FNM	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AJQQ05G	SFP+-10G-SR
FPC 15	REV 32	750-028467	ABBN6791	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN7289	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00424	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K01849	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01862	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K01852	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00427	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K00430	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01854	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K00426	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00429	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K01864	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01850	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K00522	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E01144	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E00985	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E00796	SFP+-10G-USR
Xcvr 3	REV 01	740-031980	B11K01866	SFP+-10G-SR
FPC 16	REV 30	750-028467	ABBM4592	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN0465	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01435	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01052	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01328	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01254	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02738	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02881	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01624	SFP+-10G-USR

Xcvr 3	REV 01	740-030658	B11E00889	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02883	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E00681	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E04306	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E02813	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01801	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02753	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01156	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E04324	SFP+-10G-USR
FPC 17	REV 32	750-028467	ABBN6810	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN7237	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02638	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02082	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A01674	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03058	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A03048	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02729	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02566	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02567	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02878	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02739	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A01959	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02660	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02731	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02588	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02673	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02654	SFP+-10G-SR
FPC 18	REV 30	750-028467	ABBM4739	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN0487	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02569	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02886	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A03082	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	133363A00297	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02726	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A03050	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02884	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03076	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02581	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02873	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02582	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03083	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031981	UL70BU6	SFP+-10G-LR
Xcvr 1	REV 01	740-031981	UL50QC6	SFP+-10G-LR
Xcvr 2	REV 01	740-031981	UL708N6	SFP+-10G-LR
Xcvr 3	REV 01	740-031981	UL603KK	SFP+-10G-LR
FPC 19	REV 32	750-028467	ABBN6827	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN6508	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A01688	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A01724	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A01773	SFP+-10G-SR

Xcvr 3	REV 01	740-031980	163363A02593	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A03061	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A03056	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02669	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03070	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02572	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02697	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02585	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03052	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02591	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02649	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02577	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02698	SFP+-10G-SR
ADC 0	REV 13	750-043596	ABBX5561	Adapter Card
ADC 1	REV 13	750-043596	ABBX5546	Adapter Card
ADC 2	REV 13	750-043596	ABBX5535	Adapter Card
ADC 3	REV 13	750-043596	ABBX5552	Adapter Card
ADC 4	REV 13	750-043596	ABBX5581	Adapter Card
ADC 5	REV 13	750-043596	ABBX5545	Adapter Card
ADC 6	REV 13	750-043596	ABBX5554	Adapter Card
ADC 7	REV 07	750-043596	ABBV7194	Adapter Card
ADC 8	REV 07	750-043596	ABBV7251	Adapter Card
ADC 9	REV 07	750-043596	ABBV7202	Adapter Card
ADC 10	REV 13	750-043596	ABBX5538	Adapter Card
ADC 11	REV 13	750-043596	ABBX5566	Adapter Card
ADC 12	REV 13	750-043596	ABBX5542	Adapter Card
ADC 13	REV 13	750-043596	ABBX5539	Adapter Card
ADC 14	REV 13	750-043596	ABBX5555	Adapter Card
ADC 15	REV 13	750-043596	ABBX5557	Adapter Card
ADC 16	REV 13	750-043596	ABBX5536	Adapter Card
ADC 17	REV 13	750-043596	ABBX5559	Adapter Card
ADC 18	REV 13	750-043596	ABBX5537	Adapter Card
ADC 19	REV 11	750-043596	ABBW5685	Adapter Card
Fan Tray 0	REV 2A	760-046960	ACAY0030	172mm FanTray - 6 Fans
Fan Tray 1	REV 2A	760-046960	ACAY0039	172mm FanTray - 6 Fans
Fan Tray 2	REV 2A	760-046960	ACAY0033	172mm FanTray - 6 Fans
Fan Tray 3	REV 2A	760-046960	ACAY0062	172mm FanTray - 6 Fans

show chassis hardware detail (MX2020 Router)

```
user@host> show chassis hardware detail
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN11E2227AFJ	MX2020
Midplane	REV 27	750-040240	ABAB9384	Lower Power Midplane
Midplane 1	REV 04	711-032386	ABAB9386	Upper Backplane
PMP 1	REV 05	711-032428	ACA11821	Upper Power Midplane
PMP 0	REV 04	711-032426	ACA11524	Lower Power Midplane
FPM Board	REV 06	760-040242	ABBT8837	Front Panel Display
PSM 0	REV 01	740-045050	1E02224006G	DC 52V Power Supply
Module				
PSM 1	REV 01	740-045050	1E022240053	DC 52V Power Supply
Module				
PSM 2	REV 01	740-045050	1E02224004K	DC 52V Power Supply
Module				
PSM 3	REV 01	740-045050	1E022240056	DC 52V Power Supply
Module				
PSM 4	REV 01	740-045050	1E022240054	DC 52V Power Supply

Module				
PSM 5	REV 01	740-045050	1E02224005H	DC 52V Power Supply
Module				
PSM 6	REV 01	740-045050	1E02224006S	DC 52V Power Supply
Module				
PSM 7	REV 01	740-045050	1E02224005M	DC 52V Power Supply
Module				
PSM 8	REV 01	740-045050	1E022240062	DC 52V Power Supply
Module				
PSM 9	REV 03	740-045050	1EDB2350095	DC 52V Power Supply
Module				
PSM 10	REV 03	740-045050	1EDB235009L	DC 52V Power Supply
Module				
PSM 11	REV 03	740-045050	1EDB2350092	DC 52V Power Supply
Module				
PSM 12	REV 03	740-045050	1EDB23500AT	DC 52V Power Supply
Module				
PSM 13	REV 03	740-045050	1EDB2350094	DC 52V Power Supply
Module				
PSM 15	REV 03	740-045050	1EDB235008X	DC 52V Power Supply
Module				
PDM 0	REV 01	740-045234	1E012150033	DC Power Dist Module
PDM 1	REV 01	740-045234	1E012150027	DC Power Dist Module
PDM 2	REV 01	740-045234	1E262250072	DC Power Dist Module
Routing Engine 0	REV 02	740-041821	9009094138	RE-S-1800x4
ad0	3998 MB	Virtium - TuffDisk	VCF3 20110825A021D0000064	Compact Flash
ad1	30533 MB	UGB94ARF32H0S3-KC	UNIGEN-499551-000347	Disk 1
usb0 (addr 1)		EHCI root hub 0	Intel	uhub0
usb0 (addr 2)		product 0x0020 32	vendor 0x8087	uhub1
DIMM 0		SGU04G72H1BD2SA-BB DIE	REV-52 PCB REV-54	MFR ID-ce80
DIMM 1		SGU04G72H1BD2SA-BB DIE	REV-52 PCB REV-54	MFR ID-ce80
DIMM 2		SGU04G72H1BD2SA-BB DIE	REV-52 PCB REV-54	MFR ID-ce80
DIMM 3		SGU04G72H1BD2SA-BB DIE	REV-52 PCB REV-54	MFR ID-ce80
Routing Engine 1	REV 02	740-041821	9009089709	RE-S-1800x4
ad0	3831 MB	UGB30SFA4000T1	SFA4000T1 00000113	Compact Flash
ad1	30533 MB	UGB94ARF32H0S3-KC	UNIGEN-478612-001044	Disk 1
CB 0	REV 08	750-040257	CAAB3482	Control Board
CB 1	REV 04	750-040257	ZT2864	Control Board
SPMB 0	REV 02	711-041855	CAAA6141	PMB Board
SPMB 1	REV 01	711-041855	ZS2275	PMB Board
SFB 0	REV 05	711-044466	ABBT2161	Switch Fabric Board
SFB 1	REV 05	711-044466	ABBT2159	Switch Fabric Board
SFB 2	REV 05	711-044466	ABBX3718	Switch Fabric Board
SFB 3	REV 05	711-044466	ABBT2152	Switch Fabric Board
SFB 4	REV 05	711-044466	ABBT2160	Switch Fabric Board
SFB 5	REV 05	711-044466	ABBT2145	Switch Fabric Board
SFB 6	REV 05	711-044466	ABBT2150	Switch Fabric Board
SFB 7	REV 05	711-044466	ABBT2163	Switch Fabric Board
FPC 0	REV 30	750-028467	ABBN0284	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN0507	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E00990	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E04357	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01327	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E04375	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02760	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02904	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E03963	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E00756	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+

Xcvr 0	REV 01	740-030658	B11E04418	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01077	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01128	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01253	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E01140	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11F01626	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01075	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E01177	SFP+-10G-USR
FPC 1	REV 30	750-028467	ABBN0308	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBJ1095	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04305	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01147	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01195	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01743	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01892	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02880	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E00725	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E01057	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02816	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11C04501	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E02764	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E00789	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01250	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02847	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E00787	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E03803	SFP+-10G-USR
FPC 2	REV 30	750-028467	ABBN0316	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBJ1082	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00523	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K01848	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01865	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K00540	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00422	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K00428	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K00423	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K01855	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K01847	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K00526	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K00529	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K00525	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00425	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K00530	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01851	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K00528	SFP+-10G-SR
FPC 3	REV 32	750-028467	ABBN6832	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBK6534	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80MB4	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80FQ6	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80N1F	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NLQ	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+

Xcvr 0	REV 01	740-031980	AK80KDR	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80FGJ	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80N5G	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80KD8	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LET	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80N1X	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NRF	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NL2	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80N3D	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MRB	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LEQ	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LER	SFP+-10G-SR
FPC 4	REV 32	750-028467	ABB6811	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABB7288	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NK8	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80LJG	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LBU	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80N21	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LEU	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NLM	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NL6	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LES	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LEN	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80ME0	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LMG	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80MM1	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80MG7	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80KF9	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NRQ	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NLE	SFP+-10G-SR
FPC 5	REV 32	750-028467	ABB6791	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABB7289	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00424	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K01849	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01862	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K01852	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP
Xcvr 0	REV 01	740-031980	B11K00427	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K00430	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01854	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K00426	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00429	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K01864	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01850	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K00522	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E01144	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E00985	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E00796	SFP+-10G-USR
Xcvr 3	REV 01	740-031980	B11K01866	SFP+-10G-SR
FPC 6	REV 30	750-028467	ABBM4592	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABB0465	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+

Xcvr 0	REV 01	740-030658	B11F01435	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01052	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01328	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01254	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02738	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02881	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01624	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E00889	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02883	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E00681	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E04306	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E02813	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01801	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02753	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01156	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E04324	SFP+-10G-USR
FPC 7	REV 32	750-028467	ABBN6810	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN7237	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A03058	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02082	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A01674	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02638	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A03048	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02729	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02566	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02567	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02878	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02739	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A01959	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02660	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02731	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02588	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02673	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02654	SFP+-10G-SR
FPC 8	REV 30	750-028467	ABBM4739	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN0487	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02569	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02886	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A03082	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	133363A00297	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02726	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A03050	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02884	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03076	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02581	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02873	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02582	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03083	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031981	UL70BU6	SFP+-10G-LR
Xcvr 1	REV 01	740-031981	UL50QC6	SFP+-10G-LR

Xcvr 2	REV 01	740-031981	UL708N6	SFP+-10G-LR
Xcvr 3	REV 01	740-031981	UL603KK	SFP+-10G-LR
FPC 9	REV 32	750-028467	ABBN6827	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBK6508	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A01688	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A01724	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A01773	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02593	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A03061	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A03056	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02669	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03070	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02572	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02697	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02585	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03052	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02591	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02649	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02577	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02698	SFP+-10G-SR
FPC 10	REV 30	750-028467	ABBN0302	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN0495	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01581	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01176	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01251	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E02752	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E00786	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01020	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01023	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E02819	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02812	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11D04437	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01279	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01333	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E00978	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01018	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01784	SFP+-10G-USR
Xcvr 3	REV 01	740-031980	AK80NKP	SFP+-10G-SR
FPC 11	REV 32	750-028467	ABBN6790	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBK6515	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LZM	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MCC	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80KCM	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80KE0	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021310	C10F99155	SFP+-10G-LRM
Xcvr 1	REV 01	740-021310	C10F99049	SFP+-10G-LRM
Xcvr 2	REV 01	740-021310	C10F99128	SFP+-10G-LRM
Xcvr 3	REV 01	740-021310	C10F99169	SFP+-10G-LRM
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LF3	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02597	SFP+-10G-SR

Xcvr 2	REV 01	740-031980	163363A03060	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03057	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LEX	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80FEU	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80FNM	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AJQQQ5G	SFP+-10G-SR
FPC 12	REV 30	750-028467	ZM5111	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ZP6607	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LJA	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MFZ	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NKL	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80KF4	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80FBJ	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MM2	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LJV	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NXV	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80N1H	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NLS	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80FL5	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NL9	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NG2	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80KDU	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80MG1	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80MM0	SFP+-10G-SR
FPC 13	REV 30	750-028467	ABBNO208	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABB11084	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04745	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11F01570	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E04388	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01439	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04739	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11F01869	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01675	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01901	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01346	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11F01288	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01824	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E04312	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02811	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E03847	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01495	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01265	SFP+-10G-USR
FPC 14	REV 23	750-028467	YN2977	MPC 3D 16x 10GE
CPU	REV 10	711-029089	YP1856	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	183363A00875	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	183363A00851	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	183363A00772	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	183363A00882	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	183363A00735	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	183363A00169	SFP+-10G-SR

Xcvr 2	REV 01	740-031980	183363A00726	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	183363A00077	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	183363A00168	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	183363A00676	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	183363A00732	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	183363A00091	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	183363A00725	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	183363A00642	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	183363A00871	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	183363A00853	SFP+-10G-SR
FPC 15	REV 32	750-028467	ABBN6798	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBK6556	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	9ZDZ06A00055	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	183363A00239	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AD0915E003K	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AD0915E003A	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80MRC	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NL5	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NKN	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80N3U	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80N1T	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AJ808DJ	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NG4	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80FND	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80FKQ	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NLT	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NKR	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LKM	SFP+-10G-SR
FPC 16	REV 30	750-028467	ABBN0270	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBJ0966	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NL1	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NXW	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80KD2	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80FMD	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NKQ	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MGH	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80N38	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NL7	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80M5J	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NKD	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80KCY	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LHK	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LEL	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MBE	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NLG	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LFH	SFP+-10G-SR
FPC 17	REV 32	750-028467	ABBN6796	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN7259	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K01856	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K01853	SFP+-10G-SR

Xcvr 2	REV 01	740-031980	B11K01863	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02863	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02668	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02881	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A01671	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02627	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02725	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02692	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02730	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03081	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02736	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02568	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02747	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02579	SFP+-10G-SR
FPC 18	REV 30	750-028467	ABBN0281	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN0526	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01326	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E03973	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E00950	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E00674	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E00775	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E04461	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01074	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E02821	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04501	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E00757	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01623	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E01022	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04359	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02751	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E02736	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E01178	SFP+-10G-USR
FPC 19	REV 32	750-028467	ABBN6813	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN6542	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NA3	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NLF	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80MRH	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80KE4	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	973152A00030	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80L9H	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80ME8	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NLR	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NG1	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MCA	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LFC	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LEM	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80N9X	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80LAC	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LFC2	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80N8T	SFP+-10G-SR

ADC 0	REV 13	750-043596	ABBX5561	Adapter Card
ADC 1	REV 13	750-043596	ABBX5546	Adapter Card
ADC 2	REV 13	750-043596	ABBX5535	Adapter Card
ADC 3	REV 13	750-043596	ABBX5552	Adapter Card
ADC 4	REV 13	750-043596	ABBX5581	Adapter Card
ADC 5	REV 13	750-043596	ABBX5545	Adapter Card
ADC 6	REV 13	750-043596	ABBX5554	Adapter Card
ADC 7	REV 07	750-043596	ABBV7194	Adapter Card
ADC 8	REV 07	750-043596	ABBV7251	Adapter Card
ADC 9	REV 07	750-043596	ABBV7202	Adapter Card
ADC 10	REV 13	750-043596	ABBX5579	Adapter Card
ADC 11	REV 13	750-043596	ABBX5548	Adapter Card
ADC 12	REV 13	750-043596	ABBX5575	Adapter Card
ADC 13	REV 13	750-043596	ABBX5539	Adapter Card
ADC 14	REV 13	750-043596	ABBX5555	Adapter Card
ADC 15	REV 13	750-043596	ABBX5557	Adapter Card
ADC 16	REV 13	750-043596	ABBX5536	Adapter Card
ADC 17	REV 13	750-043596	ABBX5559	Adapter Card
ADC 18	REV 13	750-043596	ABBX5537	Adapter Card
ADC 19	REV 11	750-043596	ABBW5685	Adapter Card
Fan Tray 0	REV 04	760-046960	ACAY0090	172mm FanTray - 6 Fans
Fan Tray 1	REV 04	760-046960	ACAY0088	172mm FanTray - 6 Fans
Fan Tray 2	REV 04	760-046960	ACAY0089	172mm FanTray - 6 Fans
Fan Tray 3	REV 04	760-046960	ACAY0108	172mm FanTray - 6 Fans

show chassis hardware models (MX2020 Router)

```
user@host > show chassis hardware models
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	FRU model number
Midplane	REV 27	750-040240	ABAB9384	750-040240
FPM Board	REV 06	760-040242	ABBT8837	760-040242
PSM 0	REV 01	740-045050	1E02224006G	MX2000-PSM-HC-DC-S-A
PSM 1	REV 01	740-045050	1E022240053	MX2000-PSM-HC-DC-S-A
PSM 2	REV 01	740-045050	1E02224004K	MX2000-PSM-HC-DC-S-A
PSM 3	REV 01	740-045050	1E022240056	MX2000-PSM-HC-DC-S-A
PSM 4	REV 01	740-045050	1E022240054	MX2000-PSM-HC-DC-S-A
PSM 5	REV 01	740-045050	1E02224005H	MX2000-PSM-HC-DC-S-A
PSM 6	REV 01	740-045050	1E02224006S	MX2000-PSM-HC-DC-S-A
PSM 7	REV 01	740-045050	1E02224005M	MX2000-PSM-HC-DC-S-A
PSM 8	REV 01	740-045050	1E022240062	MX2000-PSM-HC-DC-S-A
PSM 9	REV 03	740-045050	1EDB2350095	MX2000-PSM-DC-S-A
PSM 10	REV 03	740-045050	1EDB235009L	MX2000-PSM-DC-S-A
PSM 11	REV 03	740-045050	1EDB2350092	MX2000-PSM-DC-S-A
PSM 12	REV 03	740-045050	1EDB23500AT	MX2000-PSM-DC-S-A
PSM 13	REV 03	740-045050	1EDB2350094	MX2000-PSM-DC-S-A
PSM 15	REV 03	740-045050	1EDB235008X	MX2000-PSM-DC-S-A
PDM 0	REV 01	740-045234	1E012150033	
PDM 1	REV 01	740-045234	1E012150027	
PDM 2	REV 01	740-045234	1E262250072	MX2000-PDM-DC-S-A
Routing Engine 0	REV 02	740-041821	9009094138	RE-S-1800X4-16G-S
Routing Engine 1	REV 02	740-041821	9009089709	RE-S-1800X4-16G-S
CB 0	REV 08	750-040257	CAAB3482	750-040257
CB 1	REV 04	750-040257	ZT2864	750-040257
SFB 0	REV 05	711-044466	ABBT2161	MX2000-SFB-S
SFB 1	REV 05	711-044466	ABBT2159	MX2000-SFB-S
SFB 2	REV 05	711-044466	ABBX3718	MX2000-SFB-S
SFB 4	REV 05	711-044466	ABBT2160	MX2000-SFB-S
SFB 5	REV 05	711-044466	ABBT2145	MX2000-SFB-S
SFB 7	REV 05	711-044466	ABBT2163	MX2000-SFB-S
FPC 0	REV 30	750-028467	ABBN0284	MPC-3D-16XGE-SFPP

FPC 1	REV 30	750-028467	ABBN0308	MPC-3D-16XGE-SFPP
FPC 2	REV 30	750-028467	ABBN0316	MPC-3D-16XGE-SFPP
FPC 3	REV 32	750-028467	ABBN6832	MPC-3D-16XGE-SFPP
FPC 4	REV 32	750-028467	ABBN6811	MPC-3D-16XGE-SFPP
FPC 5	REV 32	750-028467	ABBN6791	MPC-3D-16XGE-SFPP
FPC 6	REV 30	750-028467	ABBM4592	MPC-3D-16XGE-SFPP
FPC 7	REV 32	750-028467	ABBN6810	MPC-3D-16XGE-SFPP
FPC 8	REV 30	750-028467	ABBM4739	MPC-3D-16XGE-SFPP
FPC 9	REV 32	750-028467	ABBN6827	MPC-3D-16XGE-SFPP
FPC 10	REV 30	750-028467	ABBN0302	MPC-3D-16XGE-SFPP
FPC 11	REV 32	750-028467	ABBN6790	MPC-3D-16XGE-SFPP
FPC 12	REV 30	750-028467	ZM5111	MPC-3D-16XGE-SFPP
FPC 13	REV 30	750-028467	ABBN0208	MPC-3D-16XGE-SFPP
FPC 14	REV 23	750-028467	YN2977	MPC-3D-16XGE-SFPP
FPC 15	REV 32	750-028467	ABBN6798	MPC-3D-16XGE-SFPP
FPC 16	REV 30	750-028467	ABBN0270	MPC-3D-16XGE-SFPP
FPC 17	REV 32	750-028467	ABBN6796	MPC-3D-16XGE-SFPP
FPC 18	REV 30	750-028467	ABBN0281	MPC-3D-16XGE-SFPP
FPC 19	REV 32	750-028467	ABBN6813	MPC-3D-16XGE-SFPP
ADC 0	REV 13	750-043596	ABBX5561	PROTO-ASSEMBLY
ADC 1	REV 13	750-043596	ABBX5546	PROTO-ASSEMBLY
ADC 2	REV 13	750-043596	ABBX5535	MX2000-LC-ADAPTER
ADC 3	REV 13	750-043596	ABBX5552	MX2000-LC-ADAPTER
ADC 4	REV 13	750-043596	ABBX5581	MX2000-LC-ADAPTER
ADC 5	REV 13	750-043596	ABBX5545	PROTO-ASSEMBLY
ADC 6	REV 13	750-043596	ABBX5554	PROTO-ASSEMBLY
ADC 7	REV 07	750-043596	ABBV7194	MX2000-LC-ADAPTER
ADC 8	REV 07	750-043596	ABBV7251	MX2000-LC-ADAPTER
ADC 9	REV 07	750-043596	ABBV7202	MX2000-LC-ADAPTER
ADC 10	REV 13	750-043596	ABBX5579	MX2000-LC-ADAPTER
ADC 12	REV 13	750-043596	ABBX5575	MX2000-LC-ADAPTER
ADC 13	REV 13	750-043596	ABBX5539	PROTO-ASSEMBLY
ADC 14	REV 13	750-043596	ABBX5555	PROTO-ASSEMBLY
ADC 15	REV 13	750-043596	ABBX5557	MX2000-LC-ADAPTER
ADC 16	REV 13	750-043596	ABBX5536	PROTO-ASSEMBLY
ADC 17	REV 13	750-043596	ABBX5559	PROTO-ASSEMBLY
ADC 18	REV 13	750-043596	ABBX5537	PROTO-ASSEMBLY
ADC 19	REV 11	750-043596	ABBW5685	PROTO-ASSEMBLY
Fan Tray 0	REV 04	760-046960	ACAY0090	
Fan Tray 1	REV 04	760-046960	ACAY0088	
Fan Tray 2	REV 04	760-046960	ACAY0089	
Fan Tray 3	REV 04	760-046960	ACAY0108	

show chassis hardware clei-models (MX2020 Router)

```
user@ host > show chassis hardware clei-models
```

```
Hardware inventory:
```

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 27	750-040240	PROTOXCLEI	750-040240
FPM Board	REV 06	760-040242	PROTOXCLEI	760-040242
PSM 0	REV 01	740-045050	IPUPAJMKAA	MX2000-PSM-HC-DC-S-A
PSM 1	REV 01	740-045050	IPUPAJMKAA	MX2000-PSM-HC-DC-S-A
PSM 2	REV 01	740-045050	IPUPAJMKAA	MX2000-PSM-HC-DC-S-A
PSM 3	REV 01	740-045050	IPUPAJMKAA	MX2000-PSM-HC-DC-S-A
PSM 4	REV 01	740-045050	IPUPAJMKAA	MX2000-PSM-HC-DC-S-A
PSM 5	REV 01	740-045050	IPUPAJMKAA	MX2000-PSM-HC-DC-S-A
PSM 6	REV 01	740-045050	IPUPAJMKAA	MX2000-PSM-HC-DC-S-A
PSM 7	REV 01	740-045050	IPUPAJMKAA	MX2000-PSM-HC-DC-S-A
PSM 8	REV 01	740-045050	IPUPAJMKAA	MX2000-PSM-HC-DC-S-A
PSM 9	REV 03	740-045050	IPUPAJMKAA	MX2000-PSM-DC-S-A
PSM 10	REV 03	740-045050	IPUPAJMKAA	MX2000-PSM-DC-S-A

PSM 11	REV 03	740-045050	IPUPAJMKAA	MX2000-PSM-DC-S-A
PSM 12	REV 03	740-045050	IPUPAJMKAA	MX2000-PSM-DC-S-A
PSM 13	REV 03	740-045050	IPUPAJMKAA	MX2000-PSM-DC-S-A
PSM 15	REV 03	740-045050	IPUPAJMKAA	MX2000-PSM-DC-S-A
PDM 0	REV 01	740-045234		
PDM 1	REV 01	740-045234		
PDM 2	REV 01	740-045234	IPUPAJSKAA	MX2000-PDM-DC-S-A
Routing Engine 0	REV 02	740-041821		RE-S-1800X4-16G-S
Routing Engine 1	REV 02	740-041821		RE-S-1800X4-16G-S
CB 0	REV 08	750-040257	PROTOXCLEI	750-040257
CB 1	REV 04	750-040257	PROTOXCLEI	750-040257
SFB 0	REV 05	711-044466	IPUCBA6CAA	MX2000-SFB-S
SFB 1	REV 05	711-044466	IPUCBA6CAA	MX2000-SFB-S
SFB 2	REV 05	711-044466	IPUCBA6CAA	MX2000-SFB-S
SFB 4	REV 05	711-044466	IPUCBA6CAA	MX2000-SFB-S
SFB 5	REV 05	711-044466	IPUCBA6CAA	MX2000-SFB-S
SFB 7	REV 05	711-044466	IPUCBA6CAA	MX2000-SFB-S
FPC 0	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 1	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 2	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 3	REV 32	750-028467		MPC-3D-16XGE-SFPP
FPC 4	REV 32	750-028467		MPC-3D-16XGE-SFPP
FPC 5	REV 32	750-028467		MPC-3D-16XGE-SFPP
FPC 6	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 7	REV 32	750-028467		MPC-3D-16XGE-SFPP
FPC 8	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 9	REV 32	750-028467		MPC-3D-16XGE-SFPP
FPC 10	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 11	REV 32	750-028467		MPC-3D-16XGE-SFPP
FPC 12	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 13	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 14	REV 23	750-028467		MPC-3D-16XGE-SFPP
FPC 15	REV 32	750-028467		MPC-3D-16XGE-SFPP
FPC 16	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 17	REV 32	750-028467		MPC-3D-16XGE-SFPP
FPC 18	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 19	REV 32	750-028467		MPC-3D-16XGE-SFPP
ADC 0	REV 13	750-043596	PROTOXCLEI	PROTO-ASSEMBLY
ADC 1	REV 13	750-043596	PROTOXCLEI	PROTO-ASSEMBLY
ADC 2	REV 13	750-043596	IPUCBA8CAA	MX2000-LC-ADAPTER
ADC 3	REV 13	750-043596	IPUCBA8CAA	MX2000-LC-ADAPTER
ADC 4	REV 13	750-043596	IPUCBA8CAA	MX2000-LC-ADAPTER
ADC 5	REV 13	750-043596	PROTOXCLEI	PROTO-ASSEMBLY
ADC 6	REV 13	750-043596	PROTOXCLEI	PROTO-ASSEMBLY
ADC 7	REV 07	750-043596	PROTOXCLEI	MX2000-LC-ADAPTER
ADC 8	REV 07	750-043596	PROTOXCLEI	MX2000-LC-ADAPTER
ADC 9	REV 07	750-043596	PROTOXCLEI	MX2000-LC-ADAPTER
ADC 10	REV 13	750-043596	IPUCBA8CAA	MX2000-LC-ADAPTER
ADC 12	REV 13	750-043596	IPUCBA8CAA	MX2000-LC-ADAPTER
ADC 13	REV 13	750-043596	PROTOXCLEI	PROTO-ASSEMBLY
ADC 14	REV 13	750-043596	PROTOXCLEI	PROTO-ASSEMBLY
ADC 15	REV 13	750-043596	IPUCBA8CAA	MX2000-LC-ADAPTER
ADC 16	REV 13	750-043596	PROTOXCLEI	PROTO-ASSEMBLY
ADC 17	REV 13	750-043596	PROTOXCLEI	PROTO-ASSEMBLY
ADC 18	REV 13	750-043596	PROTOXCLEI	PROTO-ASSEMBLY
ADC 19	REV 11	750-043596	PROTOXCLEI	PROTO-ASSEMBLY
Fan Tray 0	REV 04	760-046960		
Fan Tray 1	REV 04	760-046960		
Fan Tray 2	REV 04	760-046960		
Fan Tray 3	REV 04	760-046960		

show chassis hardware (MX2020 Router with MPC5EQ and MPC6E)

user@host> show chassis hardware

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN120BADBAFJ	MX2020
Midplane	REV 51	750-040240	ABAB9243	Lower Backplane
Midplane 1	REV 04	711-032386	ABAB9399	Upper Backplane
PMP 1	REV 05	711-032428	ACAJ2541	Upper Power Midplane
PMP 0	REV 04	711-032426	ACAJ2194	Lower Power Midplane
FPM Board	REV 13	760-040242	ABCA8835	Front Panel Display
PSM 0	REV 01	740-050037	1EDB32403L5	DC 52V Power Supply
Module				
PSM 1	REV 01	740-050037	1EDB32403L3	DC 52V Power Supply
Module				
PSM 2	REV 01	740-050037	1EDB32403KM	DC 52V Power Supply
Module				
PSM 3	REV 01	740-050037	1EDB3130079	DC 52V Power Supply
Module				
PSM 4	REV 01	740-050037	1EDB3130077	DC 52V Power Supply
Module				
PSM 5	REV 01	740-050037	1EDB3130020	DC 52V Power Supply
Module				
PSM 6	REV 01	740-050037	1EDB313009S	DC 52V Power Supply
Module				
PSM 7	REV 01	740-050037	1EDB313008E	DC 52V Power Supply
Module				
PSM 8	REV 01	740-050037	1EDB3130063	DC 52V Power Supply
Module				
PSM 12	REV 01	740-050037	1EDB3130026	DC 52V Power Supply
Module				
PSM 13	REV 01	740-050037	1EDB3130074	DC 52V Power Supply
Module				
PSM 14	REV 01	740-050037	1EDB313009D	DC 52V Power Supply
Module				
PSM 15	REV 01	740-050037	1EDB3130024	DC 52V Power Supply
Module				
PSM 16	REV 01	740-050037	1EDB3130054	DC 52V Power Supply
Module				
PSM 17	REV 01	740-050037	1EDB3130080	DC 52V Power Supply
Module				
PDM 0	REV 03	740-045234	1EGA3170144	DC Power Dist Module
PDM 1	REV 03	740-045234	1EGA3170158	DC Power Dist Module
PDM 2	REV 03	740-045234	1EGA3170182	DC Power Dist Module
PDM 3	REV 03	740-045234	1EGA3170207	DC Power Dist Module
Routing Engine 0	REV 02	740-041821	9009112112	RE-S-1800x4
Routing Engine 1	REV 02	740-041821	9009112087	RE-S-1800x4
CB 0	REV 23	750-040257	CABA2295	Control Board
CB 1	REV 23	750-040257	CABE8379	Control Board
SPMB 0	REV 02	711-041855	ABCE8851	PMB Board
SPMB 1	REV 02	711-041855	ABCE8839	PMB Board
SFB 0	REV 06	711-044466	ABCD5001	Switch Fabric Board
SFB 1	REV 06	711-044466	ABCD5034	Switch Fabric Board
SFB 2	REV 06	711-044466	ABCH3899	Switch Fabric Board
SFB 3	REV 06	711-044466	ABCD5020	Switch Fabric Board
SFB 4	REV 06	711-044466	ABCD4975	Switch Fabric Board
SFB 5	REV 06	711-044466	ABCH3881	Switch Fabric Board
SFB 6	REV 06	711-044466	ABCD5026	Switch Fabric Board
SFB 7	REV 06	711-044466	ABCD5032	Switch Fabric Board
FPC 0	REV 39	750-045715	CACD1902	MPC5E 3D Q 24XGE+6XLGE
CPU	REV 09	711-045719	CACB1933	RMPC PMB

PIC 0		BUILTIN	BUILTIN	12X10GE SFPP OTN
Xcvr 0	REV 01	740-031980	B11F00361	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	19T511101854	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	19T511100377	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	ANT0878	SFP+-10G-SR
Xcvr 5	REV 01	740-021308	19T511100398	SFP+-10G-SR
Xcvr 6	REV 01	740-021308	AQ4363J	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	19T511101377	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	ANT072M	SFP+-10G-SR
Xcvr 9	REV 01	740-021308	AG90C7N	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	AM30M09	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	B10E01016	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	12X10GE SFPP OTN
Xcvr 0	REV 01	740-031980	B10L04151	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	19T511101379	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQ5036J	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AG90C4M	SFP+-10G-SR
Xcvr 4	REV 01	740-021308	19T511101104	SFP+-10G-SR
Xcvr 5	REV 01	740-021308	AQ502ZM	SFP+-10G-SR
Xcvr 6	REV 01	740-021308	AN10KY2	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	AQ43G41	SFP+-10G-SR
Xcvr 8	REV 01	740-021308	AQ41F04	SFP+-10G-SR
Xcvr 9	REV 01	740-031980	AMS16N3	SFP+-10G-SR
Xcvr 10	REV 01	740-021308	AMH04Y3	SFP+-10G-SR
Xcvr 11	REV 01	740-021308	ANA093E	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	3X40GE QSFPP
PIC 3		BUILTIN	BUILTIN	3X40GE QSFPP
WAN MEZZ	REV 09	750-049136	CABN0410	MPC5E 24XGE OTN Mezz
FPC 1	REV 11	750-045372	CABK8112	MPCE Type 3 3D
CPU	REV 08	711-035209	CABJ6621	HMPC PMB 2G
MIC 0	REV 07	750-033307	CAAZ2897	10X10GE SFPP
PIC 0		BUILTIN	BUILTIN	10X10GE SFPP
Xcvr 0	REV 01	740-021308	AQ501VK	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQ501YC	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQ43HJF	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQ43H8D	SFP+-10G-SR
Xcvr 4	REV 01	740-021308	19T511100370	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	153363A00763	SFP+-10G-SR
Xcvr 6	REV 01	740-021308	APH2LXB	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	AMCOLVV	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	B11F00230	SFP+-10G-SR
MIC 1	REV 14	750-033196	CAAP1390	1X100GE CXP
PIC 2		BUILTIN	BUILTIN	1X100GE CXP
Xcvr 0	REV 01	740-032166	XB11F000M	CFP2-100G-SR10
FPC 2	REV 17	750-037355	CAAS5826	MPC4E 3D 2CGE+8XGE
CPU	REV 08	711-035209	CAAR3986	HMPC PMB 2G
PIC 0		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-021308	T09F43722	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	ALP0KXF	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQ502FG	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQ502T7	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-035329	X12J00571	CFP-100G-SR10
PIC 2		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-031980	AJ71KEH	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11E01355	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11F00249	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	1X100GE CFP
FPC 3	REV 05	750-044444	CAAY9920	MPCE Type 2 3D P
CPU	REV 04	711-038484	CAAW3639	MPCE PMB 2G
MIC 0	REV 28	750-028387	CAAX1083	3D 4x 10GE XFP

PIC 0		BUILTIN	BUILTIN	2x 10GE XFP
Xcvr 0		NON-JNPR	CC07BK05B	XFP-10G-SR
Xcvr 1	REV 01	740-011571	C728XJ00U	XFP-10G-SR
PIC 1		BUILTIN	BUILTIN	2x 10GE XFP
Xcvr 0		NON-JNPR	T12L92339	XFP-10G-SR
QXM 0	REV 06	711-028408	CAAW4915	MPC QXM
QXM 1	REV 06	711-028408	CAAW4894	MPC QXM
FPC 4	REV 18	750-046005	CACH5661	MPC5E 3D Q 2CGE+4XGE
CPU	REV 09	711-045719	CACF2880	RMPC PMB
PIC 0		BUILTIN	BUILTIN	2X10GE SFPP OTN
PIC 1		BUILTIN	BUILTIN	1X100GE CFP2 OTN
Xcvr 0	REV 01	740-046563	XD16FC03Y	CFP2-100G-SR10
PIC 2		BUILTIN	BUILTIN	2X10GE SFPP OTN
PIC 3		BUILTIN	BUILTIN	1X100GE CFP2 OTN
Xcvr 0	REV 01	740-049775	J13K72997	CFP2-100G-LR4-D
FPC 5	REV 35	750-028467	CAAR2623	MPC 3D 16x 10GE
CPU	REV 11	711-029089	CAAR0491	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AQ5027T	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQ502J0	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQ5027S	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQ501Y7	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AQ501YB	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQ503EB	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQ43HJH	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQ43J0Y	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AQ50352	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQ501X6	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQ502NV	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQ502ZJ	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AQ502H4	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQ43HJK	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AJ30CU7	SFP+-10G-SR
FPC 9	REV 30	750-044130	ABCF5773	MPC6E 3D
CPU	REV 09	711-045719	ABCF1270	RMPC PMB
MIC 0	REV 05	750-049457	ABCD7829	2X100GE CFP2 OTN
PIC 0		BUILTIN	BUILTIN	2X100GE CFP2 OTN
Xcvr 0		NON-JNPR	FE13F000K	CFP2-100G-SR10
Xcvr 1	REV 01	740-048813	XD32FE017	CFP2-100G-LR-D
MIC 1	REV 07	750-049457	ABCK2812	2X100GE CFP2 OTN
PIC 1		BUILTIN	BUILTIN	2X100GE CFP2 OTN
Xcvr 0	REV 01	740-048813	XD32FE018	CFP2-100G-SR10
Xcvr 1		NON-JNPR	FE13F000E	CFP2-100G-LR4-D
XLM 0	REV 05.2.00	711-046638	ABCF5915	MPC6E XL
XLM 1	REV 05.2.00	711-046638	ABCF5916	MPC6E XL
FPC 10	REV 36	750-044130	ABCS8602	MPC6E 3D
CPU	REV 09	711-045719	ABCS8779	RMPC PMB
MIC 0	REV 06	750-049979	ABCK2656	24X10GE SFPP OTN
PIC 0		BUILTIN	BUILTIN	24X10GE SFPP OTN
Xcvr 0	REV 01	740-021308	AQ43J08	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQE1Y2E	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQE1UW4	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQE1MQF	SFP+-10G-SR
Xcvr 4	REV 01	740-021308	AQGOMN1	SFP+-10G-SR
Xcvr 5	REV 01	740-021308	AQE1L9M	SFP+-10G-SR
Xcvr 6	REV 01	740-021308	AQGOMPD	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	AQE1Y2B	SFP+-10G-SR
Xcvr 8	REV 01	740-021308	AQGOLT5	SFP+-10G-SR

Xcvr 9	REV 01	740-021308	AQD2ET4	SFP+-10G-SR
Xcvr 10	REV 01	740-021308	AQGOMPC	SFP+-10G-SR
Xcvr 11	REV 01	740-021308	AQGOM63	SFP+-10G-SR
Xcvr 12	REV 01	740-021308	AQGOLT1	SFP+-10G-SR
Xcvr 13	REV 01	740-021308	AQGOM4L	SFP+-10G-SR
Xcvr 14	REV 01	740-021308	AQGOLS7	SFP+-10G-SR
Xcvr 15	REV 01	740-021308	AQE1MQB	SFP+-10G-SR
Xcvr 16	REV 01	740-021308	AQGOLZP	SFP+-10G-SR
Xcvr 17	REV 01	740-021308	AQE1LU9	SFP+-10G-SR
Xcvr 18	REV 01	740-021308	AQGOMRZ	SFP+-10G-SR
Xcvr 19	REV 01	740-021308	AQE1MQ9	SFP+-10G-SR
Xcvr 20	REV 01	740-021308	AQGOLRX	SFP+-10G-SR
Xcvr 21	REV 01	740-021308	AQE1UWD	SFP+-10G-SR
Xcvr 22	REV 01	740-021308	AQGOLT4	SFP+-10G-SR
Xcvr 23	REV 01	740-021308	AQE1MQL	SFP+-10G-SR
MIC 1	REV 12	750-050008	ABCK5372	4X100GE CXP
PIC 1		BUILTIN	BUILTIN	4X100GE CXP
Xcvr 3	REV 01	740-046563	XD16FC02Z	CFP2-100G-SR10
XLM 0	REV 07.2.00	711-046638	ABCK3481	MPC6E XL
XLM 1	REV 07.2.00	711-046638	ABCK4725	MPC6E XL
FPC 17	REV 28	750-044130	ABBZ3873	MPC6E 3D
CPU	REV 08	711-045719	ABBZ3770	RMP C PMB
MIC 0	REV 11	750-046535	ABCC7731	24X10GE SFPP
PIC 0		BUILTIN	BUILTIN	24X10GE SFPP
Xcvr 1	REV 01	740-021308	APK0543	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B10G01119	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQ502SX	SFP+-10G-SR
Xcvr 4	REV 01	740-021308	AQ43H84	SFP+-10G-SR
Xcvr 5	REV 01	740-021308	AQ501TB	SFP+-10G-SR
Xcvr 6	REV 01	740-021308	AQ502JZ	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	AQ502SC	SFP+-10G-SR
Xcvr 8	REV 01	740-021308	AQ502JW	SFP+-10G-SR
Xcvr 9	REV 01	740-021308	AQ502RM	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	AHK013B	SFP+-10G-SR
Xcvr 11	REV 01	740-021308	AQGOMRT	SFP+-10G-SR
Xcvr 13	REV 01	740-031980	AMC0JTC	SFP+-10G-SR
Xcvr 14	REV 01	740-021308	ANAOMQ0	SFP+-10G-SR
Xcvr 15	REV 01	740-021308	AQ502GS	SFP+-10G-SR
Xcvr 16	REV 01	740-021308	AQGOM0J	SFP+-10G-SR
Xcvr 17	REV 01	740-021308	AQGOMUR	SFP+-10G-SR
Xcvr 18	REV 01	740-021308	AQGOMRR	SFP+-10G-SR
Xcvr 19	REV 01	740-021308	AQGOM0F	SFP+-10G-SR
Xcvr 20	REV 01	740-021308	AQ50312	SFP+-10G-SR
Xcvr 21	REV 01	740-021308	AQ5032U	SFP+-10G-SR
Xcvr 22	REV 01	740-021308	APE17B5	SFP+-10G-SR
Xcvr 23	REV 01	740-021309	91D104A00011	SFP+-10G-LR
MIC 1	REV 03	750-050008	ABCC4522	4X100GE CXP
PIC 1		BUILTIN	BUILTIN	4X100GE CXP
Xcvr 0	REV 01	740-046563	XD16FC02U	CFP2-100G-SR10
Xcvr 1	REV 01	740-046563	XC42FC03K	CFP2-100G-SR10
Xcvr 2	REV 01	740-046563	XC42FC01Z	CFP2-100G-SR10
Xcvr 3	REV 01	740-046563	XC42FC02U	CFP2-100G-SR10
XLM 0	REV 04.2.00	711-046638	ABBZ3779	MPC6E XL
XLM 1	REV 04.2.00	711-046638	ABBZ3780	MPC6E XL
FPC 18	REV 39	750-045715	CACD1910	MPC5E 3D Q 24XGE+6XLGE
CPU	REV 09	711-045719	CACD1817	RMP C PMB
PIC 0		BUILTIN	BUILTIN	12X10GE SFPP OTN
PIC 1		BUILTIN	BUILTIN	12X10GE SFPP OTN
PIC 2		BUILTIN	BUILTIN	3X40GE QSFPP
Xcvr 0	REV 01	740-046565	QD130194	QSFP+-40G-SR4
Xcvr 1	REV 01	740-046565	QD130193	QSFP+-40G-SR4

Xcvr 2	REV 01	740-046565	QD130196	QSFP+-40G-SR4
PIC 3		BUILTIN	BUILTIN	3X40GE QSFPP
Xcvr 0	REV 01	740-046565	QD130191	QSFP+-40G-SR4
Xcvr 1	REV 01	740-046565	QD130198	QSFP+-40G-SR4
Xcvr 2	REV 01	740-046565	QD130192	QSFP+-40G-SR4
WAN MEZZ	REV 09	750-049136	CABN0411	MPC5E 24XGE OTN Mezz
FPC 19	REV 39	750-045715	CACD1908	MPC5E 3D Q 24XGE+6XLGE
CPU	REV 09	711-045719	CACD1820	RMPC PMB
PIC 0		BUILTIN	BUILTIN	12X10GE SFPP OTN
Xcvr 0	REV 01	740-021308	AQA0EXJ	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQGOM6D	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQGOLW7	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQA0JKB	SFP+-10G-SR
Xcvr 4	REV 01	740-021308	AQGOMTM	SFP+-10G-SR
Xcvr 5	REV 01	740-021308	AQA07NE	SFP+-10G-SR
Xcvr 6	REV 01	740-021308	AQGOM41	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	AQGOMU7	SFP+-10G-SR
Xcvr 8	REV 01	740-021308	AQGOMUG	SFP+-10G-SR
Xcvr 9	REV 01	740-021308	AQGOMMX	SFP+-10G-SR
Xcvr 10	REV 01	740-021308	AQGOM5K	SFP+-10G-SR
Xcvr 11	REV 01	740-021308	AQGOLVZ	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	12X10GE SFPP OTN
PIC 2		BUILTIN	BUILTIN	3X40GE QSFPP
PIC 3		BUILTIN	BUILTIN	3X40GE QSFPP
Xcvr 0	REV 01	740-046565	QD130242	QSFP+-40G-SR4
Xcvr 1	REV 01	740-046565	QD130245	QSFP+-40G-SR4
Xcvr 2	REV 01	740-046565	QD130613	QSFP+-40G-SR4
WAN MEZZ	REV 09	750-049136	CABN0418	MPC5E 24XGE OTN Mezz
ADC 0	REV 17	750-043596	ABCD5378	Adapter Card
ADC 1	REV 17	750-043596	ABCD5465	Adapter Card
ADC 2	REV 17	750-043596	ABCD5431	Adapter Card
ADC 3	REV 17	750-043596	ABCD5356	Adapter Card
ADC 4	REV 02	750-043596	ZW1545	Adapter Card
ADC 5	REV 17	750-043596	ABCD5517	Adapter Card
ADC 18	REV 17	750-043596	ABCD5535	Adapter Card
ADC 19	REV 01	750-043596	ZV4127	Adapter Card
Fan Tray 0	REV 06	760-046960	ACAY0791	172mm FanTray - 6 Fans
Fan Tray 1	REV 06	760-046960	ACAY0788	172mm FanTray - 6 Fans
Fan Tray 2	REV 06	760-046960	ACAY0755	172mm FanTray - 6 Fans
Fan Tray 3	REV 06	760-046960	ACAY0441	172mm FanTray - 6 Fans

show chassis hardware detail (MX2020 Router with MPC5EQ and MPC6E)

```
user@host>show chassis hardware detail
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN120BADBAFJ	MX2020
Midplane	REV 51	750-040240	ABAB9243	Lower Backplane
Midplane 1	REV 04	711-032386	ABAB9399	Upper Backplane
PMP 1	REV 05	711-032428	ACAJ2541	Upper Power Midplane
PMP 0	REV 04	711-032426	ACAJ2194	Lower Power Midplane
FPM Board	REV 13	760-040242	ABCA8835	Front Panel Display
PSM 0	REV 01	740-050037	1EDB32403L5	DC 52V Power Supply
Module				
PSM 1	REV 01	740-050037	1EDB32403L3	DC 52V Power Supply
Module				
PSM 2	REV 01	740-050037	1EDB32403KM	DC 52V Power Supply
Module				
PSM 3	REV 01	740-050037	1EDB3130079	DC 52V Power Supply
Module				
PSM 4	REV 01	740-050037	1EDB3130077	DC 52V Power Supply

Module				
PSM 5	REV 01	740-050037	1EDB3130020	DC 52V Power Supply
Module				
PSM 6	REV 01	740-050037	1EDB313009S	DC 52V Power Supply
Module				
PSM 7	REV 01	740-050037	1EDB313008E	DC 52V Power Supply
Module				
PSM 8	REV 01	740-050037	1EDB3130063	DC 52V Power Supply
Module				
PSM 12	REV 01	740-050037	1EDB3130026	DC 52V Power Supply
Module				
PSM 13	REV 01	740-050037	1EDB3130074	DC 52V Power Supply
Module				
PSM 14	REV 01	740-050037	1EDB313009D	DC 52V Power Supply
Module				
PSM 15	REV 01	740-050037	1EDB3130024	DC 52V Power Supply
Module				
PSM 16	REV 01	740-050037	1EDB3130054	DC 52V Power Supply
Module				
PSM 17	REV 01	740-050037	1EDB3130080	DC 52V Power Supply
Module				
PDM 0	REV 03	740-045234	1EGA3170144	DC Power Dist Module
PDM 1	REV 03	740-045234	1EGA3170158	DC Power Dist Module
PDM 2	REV 03	740-045234	1EGA3170182	DC Power Dist Module
PDM 3	REV 03	740-045234	1EGA3170207	DC Power Dist Module
Routing Engine 0	REV 02	740-041821	9009112112	RE-S-1800x4
ad0 3998 MB	Virtium - TuffDrive		VCF P1T0200274310822	113 Compact Flash
ad1 30533 MB	UGB94BPH32H0S1-KCI		11000031656	Disk 1
usb0 (addr 1)	EHCI root hub 0		Intel	uhub0
usb0 (addr 2)	product 0x0020 32		vendor 0x8087	uhub1
DIMM 0	SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80			
DIMM 1	SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80			
DIMM 2	SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80			
DIMM 3	SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80			
Routing Engine 1	REV 02	740-041821	9009112087	RE-S-1800x4
ad0 3998 MB	Virtium - TuffDrive		VCF P1T0200274310822	366 Compact Flash
ad1 30533 MB	UGB94BPH32H0S1-KCI		11000039979	Disk 1
CB 0	REV 23	750-040257	CABA2295	Control Board
CB 1	REV 23	750-040257	CABE8379	Control Board
SPMB 0				
SPMB 1				
FPC 0	REV 39	750-045715	CACD1902	MPC5E 3D Q 24XGE+6XLGE
CPU				
FPC 1	REV 11	750-045372	CABK8112	MPCE Type 3 3D
CPU				
FPC 2	REV 17	750-037355	CAAS5826	MPC4E 3D 2CGE+8XGE
CPU				
FPC 3	REV 05	750-044444	CAAY9920	MPCE Type 2 3D P
CPU				
FPC 4	REV 18	750-046005	CACH5661	MPC5E 3D Q 2CGE+4XGE
CPU				
FPC 5	REV 35	750-028467	CAAR2623	MPC 3D 16x 10GE
CPU				
FPC 9	REV 30	750-044130	ABCF5773	MPC6E 3D
CPU				
FPC 10	REV 36	750-044130	ABCS8602	MPC6E 3D
CPU				
FPC 17	REV 28	750-044130	ABBZ3873	MPC6E 3D
CPU				
FPC 18	REV 39	750-045715	CACD1910	MPC5E 3D Q 24XGE+6XLGE
CPU				

FPC 19	REV 39	750-045715	CACD1908	MPC5E 3D Q 24XGE+6XLGE
CPU				
Fan Tray 0	REV 06	760-046960	ACAY0791	172mm FanTray - 6 Fans
Fan Tray 1	REV 06	760-046960	ACAY0788	172mm FanTray - 6 Fans
Fan Tray 2	REV 06	760-046960	ACAY0755	172mm FanTray - 6 Fans
Fan Tray 3	REV 06	760-046960	ACAY0441	172mm FanTray - 6 Fans

show chassis hardware extensive (MX2020 Router with MPC5EQ and MPC6E)

```

user@host> show chassis hardware extensive
Hardware inventory:
Item              Version  Part number  Serial number  Description
Chassis
Jedec Code:       0x7fb0          EEPROM Version: 0x02
S/N:              JN120BADBAFJ
Assembly ID:      0x0557          Assembly Version: 00.00
Date:             00-00-0000      Assembly Flags:  0x00
ID: MX2020
Board Information Record:
Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
Address 0x00: 7f b0 02 ff 05 57 00 00 00 00 00 00 00 00 00 00
Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x20: 4a 4e 31 32 30 42 41 44 42 41 46 4a 00 00 00 00
Address 0x30: 00 00 00 ff 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Midplane          REV 51    750-040240  ABAB9243      Lower Backplane
Jedec Code:       0x7fb0          EEPROM Version: 0x02
P/N:              750-040240      S/N:           ABAB9243
Assembly ID:      0x0b22          Assembly Version: 01.51
Date:             05-30-2013      Assembly Flags: 0x00
Version:          REV 51          CLEI Code:     IPMU710ARA
ID: Lower Backplane
FRU Model Number: CHAS-BP-MX2020-S
Board Information Record:
Address 0x00: ad 01 10 00 4c 96 14 72 30 08 ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 22 01 33 52 45 56 20 35 31 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 34 30 32 34 30 00 00
Address 0x20: 53 2f 4e 20 41 42 41 42 39 32 34 33 00 1e 05 07
Address 0x30: dd ff ff ff ad 01 10 00 4c 96 14 72 30 08 ff ff
Address 0x40: ff ff ff ff 01 49 50 4d 55 37 31 30 41 52 41 43
Address 0x50: 48 41 53 2d 42 50 2d 4d 58 32 30 32 30 2d 53 00
Address 0x60: 00 00 00 00 00 00 41 00 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff d3 ff ff ff ff ff ff ff ff ff ff ff ff
Midplane 1        REV 04    711-032386  ABAB9399      Upper Backplane
Jedec Code:       0x7fb0          EEPROM Version: 0x01
P/N:              711-032386      S/N:           ABAB9399
Assembly ID:      0x0b23          Assembly Version: 01.04
Date:             10-22-2012      Assembly Flags: 0x00
Version:          REV 04
ID: Upper Backplane
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 fe 0b 23 01 04 52 45 56 20 30 34 00 00
Address 0x10: 00 00 00 00 37 31 31 2d 30 33 32 33 38 36 00 00
Address 0x20: 53 2f 4e 20 41 42 41 42 39 33 39 39 00 16 0a 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff

```

```

Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff
Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

PMP 1          REV 05    711-032428    ACAJ2541          Upper Power Midplane
Jedec Code:    0x7fb0          EEPROM Version:    0x01
P/N:           711-032428      S/N:           ACAJ2541
Assembly ID:   0x045c          Assembly Version: 01.05
Date:          04-26-2013      Assembly Flags: 0x00
Version:       REV 05
ID: Upper Power Midplane
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 04 5c 01 05 52 45 56 20 30 35 00 00
Address 0x10: 00 00 00 00 37 31 31 2d 30 33 32 34 32 38 00 00
Address 0x20: 53 2f 4e 20 41 43 41 4a 32 35 34 31 00 1a 04 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff
Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

PMP 0          REV 04    711-032426    ACAJ2194          Lower Power Midplane
Jedec Code:    0x7fb0          EEPROM Version:    0x01
P/N:           711-032426      S/N:           ACAJ2194
Assembly ID:   0x045d          Assembly Version: 01.04
Date:          01-29-2013      Assembly Flags: 0x00
Version:       REV 04
ID: Lower Power Midplane
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 04 5d 01 04 52 45 56 20 30 34 00 00
Address 0x10: 00 00 00 00 37 31 31 2d 30 33 32 34 32 36 00 00
Address 0x20: 53 2f 4e 20 41 43 41 4a 32 31 39 34 00 1d 01 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff
Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

FPM Board      REV 13    760-040242    ABCA8835          Front Panel Display
Jedec Code:    0x7fb0          EEPROM Version:    0x02
P/N:           760-040242      S/N:           ABCA8835
Assembly ID:   0x0b24          Assembly Version: 01.13
Date:          04-13-2013      Assembly Flags: 0x00
Version:       REV 13          CLEI Code:        IPMYAESJRA
ID: Front Panel Display      FRU Model Number: MX2020-CRAFT-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 24 01 0d 52 45 56 20 31 33 00 00
Address 0x10: 00 00 00 00 37 36 30 2d 30 34 30 32 34 32 00 00
Address 0x20: 53 2f 4e 20 41 42 43 41 38 38 33 35 00 0d 04 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 4d 59 41 45 35 4a 52 41 4d
Address 0x50: 58 32 30 32 30 2d 43 52 41 46 54 2d 53 00 00 00
Address 0x60: 00 00 00 00 00 00 41 00 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 95 ff ff ff ff ff ff ff ff ff ff ff ff

PSM 0          REV 01    740-050037    1EDB32403L5      DC 52V Power Supply
Module
Jedec Code:    0x7fb0          EEPROM Version:    0x02

```

```

P/N:          740-050037      S/N:          1EDB32403L5
Assembly ID:  0x0478          Assembly Version: 01.01
Date:         06-21-2013      Assembly Flags: 0x00
Version:      REV 01          CLEI Code:      IPUPAKRKAA
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S
Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
  Address 0x20: 31 45 44 42 33 32 34 30 33 4c 35 00 00 15 06 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
  Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
  Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 1          REV 01  740-050037  1EDB32403L3      DC 52V Power Supply
Module
Jedec Code:    0x7fb0          EEPROM Version: 0x02
P/N:          740-050037      S/N:          1EDB32403L3
Assembly ID:   0x0478          Assembly Version: 01.01
Date:         06-21-2013      Assembly Flags: 0x00
Version:      REV 01          CLEI Code:      IPUPAKRKAA
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S
Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
  Address 0x20: 31 45 44 42 33 32 34 30 33 4c 33 00 00 15 06 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
  Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
  Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 2          REV 01  740-050037  1EDB32403KM      DC 52V Power Supply
Module
Jedec Code:    0x7fb0          EEPROM Version: 0x02
P/N:          740-050037      S/N:          1EDB32403KM
Assembly ID:   0x0478          Assembly Version: 01.01
Date:         06-21-2013      Assembly Flags: 0x00
Version:      REV 01          CLEI Code:      IPUPAKRKAA
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S
Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
  Address 0x20: 31 45 44 42 33 32 34 30 33 4b 4d 00 00 15 06 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
  Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
  Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 3          REV 01  740-050037  1EDB3130079      DC 52V Power Supply
Module
Jedec Code:    0x7fb0          EEPROM Version: 0x02
P/N:          740-050037      S/N:          1EDB3130079
Assembly ID:   0x0478          Assembly Version: 01.01
Date:         05-16-2013      Assembly Flags: 0x00
Version:      REV 01          CLEI Code:      IPUPAKRKAA

```

ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S

Board Information Record:

Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

I2C Hex Data:

Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00

Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00

Address 0x20: 31 45 44 42 33 31 33 30 30 37 39 00 00 10 05 07

Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d

Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00

Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff

Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00

PSM 4 REV 01 740-050037 1EDB3130077 DC 52V Power Supply
Module

Jedec Code: 0x7fb0 EEPROM Version: 0x02

P/N: 740-050037 S/N: 1EDB3130077

Assembly ID: 0x0478 Assembly Version: 01.01

Date: 05-17-2013 Assembly Flags: 0x00

Version: REV 01 CLEI Code: IPUPAKRKAA

ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S

Board Information Record:

Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

I2C Hex Data:

Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00

Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00

Address 0x20: 31 45 44 42 33 31 33 30 30 37 37 00 00 11 05 07

Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d

Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00

Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff

Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00

PSM 5 REV 01 740-050037 1EDB3130020 DC 52V Power Supply
Module

Jedec Code: 0x7fb0 EEPROM Version: 0x02

P/N: 740-050037 S/N: 1EDB3130020

Assembly ID: 0x0478 Assembly Version: 01.01

Date: 05-16-2013 Assembly Flags: 0x00

Version: REV 01 CLEI Code: IPUPAKRKAA

ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S

Board Information Record:

Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

I2C Hex Data:

Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00

Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00

Address 0x20: 31 45 44 42 33 31 33 30 30 32 30 00 00 10 05 07

Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d

Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00

Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff

Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00

PSM 6 REV 01 740-050037 1EDB313009S DC 52V Power Supply
Module

Jedec Code: 0x7fb0 EEPROM Version: 0x02

P/N: 740-050037 S/N: 1EDB313009S

Assembly ID: 0x0478 Assembly Version: 01.01

Date: 05-17-2013 Assembly Flags: 0x00

Version: REV 01 CLEI Code: IPUPAKRKAA

ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S

Board Information Record:

Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

I2C Hex Data:

```

Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
Address 0x20: 31 45 44 42 33 31 33 30 30 39 53 00 00 11 05 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff
Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 7          REV 01   740-050037   1EDB313008E   DC 52V Power Supply
Module
Jedec Code:    0x7fb0          EEPROM Version:  0x02
P/N:           740-050037      S/N:             1EDB313008E
Assembly ID:   0x0478          Assembly Version: 01.01
Date:          05-17-2013      Assembly Flags:   0x00
Version:       REV 01          CLEI Code:        IPUPAKRKAA
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
Address 0x20: 31 45 44 42 33 31 33 30 30 38 45 00 00 11 05 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff
Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 8          REV 01   740-050037   1EDB3130063   DC 52V Power Supply
Module
Jedec Code:    0x7fb0          EEPROM Version:  0x02
P/N:           740-050037      S/N:             1EDB3130063
Assembly ID:   0x0478          Assembly Version: 01.01
Date:          05-17-2013      Assembly Flags:   0x00
Version:       REV 01          CLEI Code:        IPUPAKRKAA
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
Address 0x20: 31 45 44 42 33 31 33 30 30 36 33 00 00 11 05 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff
Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 12         REV 01   740-050037   1EDB3130026   DC 52V Power Supply
Module
Jedec Code:    0x7fb0          EEPROM Version:  0x02
P/N:           740-050037      S/N:             1EDB3130026
Assembly ID:   0x0478          Assembly Version: 01.01
Date:          05-16-2013      Assembly Flags:   0x00
Version:       REV 01          CLEI Code:        IPUPAKRKAA
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
Address 0x20: 31 45 44 42 33 31 33 30 30 32 36 00 00 10 05 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff

```



```

Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 13          REV 01   740-050037   1EDB3130074       DC 52V Power Supply
Module
Jedec Code:    0x7fb0          EEPROM Version:    0x02
P/N:           740-050037      S/N:              1EDB3130074
Assembly ID:   0x0478          Assembly Version:  01.01
Date:          05-17-2013      Assembly Flags:    0x00
Version:       REV 01          CLEI Code:         IPUPAKRKAA
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
Address 0x20: 31 45 44 42 33 31 33 30 30 37 34 00 00 11 05 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 14          REV 01   740-050037   1EDB313009D       DC 52V Power Supply
Module
Jedec Code:    0x7fb0          EEPROM Version:    0x02
P/N:           740-050037      S/N:              1EDB313009D
Assembly ID:   0x0478          Assembly Version:  01.01
Date:          05-17-2013      Assembly Flags:    0x00
Version:       REV 01          CLEI Code:         IPUPAKRKAA
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
Address 0x20: 31 45 44 42 33 31 33 30 30 39 44 00 00 11 05 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 15          REV 01   740-050037   1EDB3130024       DC 52V Power Supply
Module
Jedec Code:    0x7fb0          EEPROM Version:    0x02
P/N:           740-050037      S/N:              1EDB3130024
Assembly ID:   0x0478          Assembly Version:  01.01
Date:          05-16-2013      Assembly Flags:    0x00
Version:       REV 01          CLEI Code:         IPUPAKRKAA
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
...

```

show chassis hardware models (MX2020 Routers with MPC5EQ and MPC6E)

```
user@host> show chassis hardware models
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	FRU model number
------	---------	-------------	---------------	------------------

Midplane	REV 51	750-040240	ABAB9243	CHAS-BP-MX2020-S
FPM Board	REV 13	760-040242	ABCA8835	MX2020-CRAFT-S
PSM 0	REV 01	740-050037	1EDB32403L5	MX2000-PSM-DC-S
PSM 1	REV 01	740-050037	1EDB32403L3	MX2000-PSM-DC-S
PSM 2	REV 01	740-050037	1EDB32403KM	MX2000-PSM-DC-S
PSM 3	REV 01	740-050037	1EDB3130079	MX2000-PSM-DC-S
PSM 4	REV 01	740-050037	1EDB3130077	MX2000-PSM-DC-S
PSM 5	REV 01	740-050037	1EDB3130020	MX2000-PSM-DC-S
PSM 6	REV 01	740-050037	1EDB313009S	MX2000-PSM-DC-S
PSM 7	REV 01	740-050037	1EDB313008E	MX2000-PSM-DC-S
PSM 8	REV 01	740-050037	1EDB3130063	MX2000-PSM-DC-S
PSM 12	REV 01	740-050037	1EDB3130026	MX2000-PSM-DC-S
PSM 13	REV 01	740-050037	1EDB3130074	MX2000-PSM-DC-S
PSM 14	REV 01	740-050037	1EDB313009D	MX2000-PSM-DC-S
PSM 15	REV 01	740-050037	1EDB3130024	MX2000-PSM-DC-S
PSM 16	REV 01	740-050037	1EDB3130054	MX2000-PSM-DC-S
PSM 17	REV 01	740-050037	1EDB3130080	MX2000-PSM-DC-S
PDM 0	REV 03	740-045234	1EGA3170144	MX2000-PDM-DC-S
PDM 1	REV 03	740-045234	1EGA3170158	MX2000-PDM-DC-S
PDM 2	REV 03	740-045234	1EGA3170182	MX2000-PDM-DC-S
PDM 3	REV 03	740-045234	1EGA3170207	MX2000-PDM-DC-S
Routing Engine 0	REV 02	740-041821	9009112112	RE-MX2000-1800X4-S
Routing Engine 1	REV 02	740-041821	9009112087	RE-MX2000-1800X4-S
CB 0	REV 23	750-040257	CABA2295	RE-MX2000-1800X4-S
CB 1	REV 23	750-040257	CABE8379	RE-MX2000-1800X4-S
SFB 0	REV 06	711-044466	ABCD5001	MX2000-SFB-S
SFB 1	REV 06	711-044466	ABCD5034	MX2000-SFB-S
SFB 2	REV 06	711-044466	ABCH3899	MX2000-SFB-S
SFB 3	REV 06	711-044466	ABCD5020	MX2000-SFB-S
SFB 4	REV 06	711-044466	ABCD4975	MX2000-SFB-S
SFB 5	REV 06	711-044466	ABCH3881	MX2000-SFB-S
SFB 6	REV 06	711-044466	ABCD5026	MX2000-SFB-S
SFB 7	REV 06	711-044466	ABCD5032	MX2000-SFB-S
FPC 0	REV 39	750-045715	CACD1902	PROTO-ASSEMBLY
FPC 1	REV 11	750-045372	CABK8112	MX-MPC3E-3D
FPC 2	REV 17	750-037355	CAAS5826	MPC4E-3D-2CGE-8XGE
FPC 3	REV 05	750-044444	CAAY9920	MX-MPC2E-3D-P
FPC 4	REV 18	750-046005	CACH5661	PROTO-ASSEMBLY
FPC 5	REV 35	750-028467	CAAR2623	MPC-3D-16XGE-SFPP
FPC 9	REV 30	750-044130	ABCF5773	PROTO-ASSEMBLY
FPC 10	REV 36	750-044130	ABCS8602	PROTO-ASSEMBLY
FPC 17	REV 28	750-044130	ABBZ3873	PROTO-ASSEMBLY
FPC 18	REV 39	750-045715	CACD1910	PROTO-ASSEMBLY
FPC 19	REV 39	750-045715	CACD1908	PROTO-ASSEMBLY
ADC 0	REV 17	750-043596	ABCD5378	MX2000-LC-ADAPTER
ADC 1	REV 17	750-043596	ABCD5465	MX2000-LC-ADAPTER
ADC 2	REV 17	750-043596	ABCD5431	MX2000-LC-ADAPTER
ADC 3	REV 17	750-043596	ABCD5356	MX2000-LC-ADAPTER
ADC 4	REV 02	750-043596	ZW1545	750-043596
ADC 5	REV 17	750-043596	ABCD5517	MX2000-LC-ADAPTER
ADC 18	REV 17	750-043596	ABCD5535	MX2000-LC-ADAPTER
ADC 19	REV 01	750-043596	ZV4127	750-043596
Fan Tray 0	REV 06	760-046960	ACAY0791	MX2000-FANTRAY-S
Fan Tray 1	REV 06	760-046960	ACAY0788	MX2000-FANTRAY-S
Fan Tray 2	REV 06	760-046960	ACAY0755	MX2000-FANTRAY-S
Fan Tray 3	REV 06	760-046960	ACAY0441	MX2000-FANTRAY-S

show chassis hardware clei-models (MX2020 Router with MPC5EQ and MPC6E)

```
user@host> show chassis hardware clei-models
```

Hardware inventory:

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 51	750-040240	IPMU710ARA	CHAS-BP-MX2020-S
FPM Board	REV 13	760-040242	IPMYAE5JRA	MX2020-CRAFT-S
PSM 0	REV 01	740-050037	IPUPAKRKAA	MX2000-PSM-DC-S
PSM 1	REV 01	740-050037	IPUPAKRKAA	MX2000-PSM-DC-S
PSM 2	REV 01	740-050037	IPUPAKRKAA	MX2000-PSM-DC-S
PSM 3	REV 01	740-050037	IPUPAKRKAA	MX2000-PSM-DC-S
PSM 4	REV 01	740-050037	IPUPAKRKAA	MX2000-PSM-DC-S
PSM 5	REV 01	740-050037	IPUPAKRKAA	MX2000-PSM-DC-S
PSM 6	REV 01	740-050037	IPUPAKRKAA	MX2000-PSM-DC-S
PSM 7	REV 01	740-050037	IPUPAKRKAA	MX2000-PSM-DC-S
PSM 8	REV 01	740-050037	IPUPAKRKAA	MX2000-PSM-DC-S
PSM 12	REV 01	740-050037	IPUPAKRKAA	MX2000-PSM-DC-S
PSM 13	REV 01	740-050037	IPUPAKRKAA	MX2000-PSM-DC-S
PSM 14	REV 01	740-050037	IPUPAKRKAA	MX2000-PSM-DC-S
PSM 15	REV 01	740-050037	IPUPAKRKAA	MX2000-PSM-DC-S
PSM 16	REV 01	740-050037	IPUPAKRKAA	MX2000-PSM-DC-S
PSM 17	REV 01	740-050037	IPUPAKRKAA	MX2000-PSM-DC-S
PDM 0	REV 03	740-045234	IPUPAJSKAA	MX2000-PDM-DC-S
PDM 1	REV 03	740-045234	IPUPAJSKAA	MX2000-PDM-DC-S
PDM 2	REV 03	740-045234	IPUPAJSKAA	MX2000-PDM-DC-S
PDM 3	REV 03	740-045234	IPUPAJSKAA	MX2000-PDM-DC-S
CB 0	REV 23	750-040257	IPUCBA7CTA	RE-MX2000-1800X4-S
CB 1	REV 23	750-040257	IPUCBA7CTA	RE-MX2000-1800X4-S
SFB 0	REV 06	711-044466	IPUCBA6CAA	MX2000-SFB-S
SFB 1	REV 06	711-044466	IPUCBA6CAA	MX2000-SFB-S
SFB 2	REV 06	711-044466	IPUCBA6CAA	MX2000-SFB-S
SFB 3	REV 06	711-044466	IPUCBA6CAA	MX2000-SFB-S
SFB 4	REV 06	711-044466	IPUCBA6CAA	MX2000-SFB-S
SFB 5	REV 06	711-044466	IPUCBA6CAA	MX2000-SFB-S
SFB 6	REV 06	711-044466	IPUCBA6CAA	MX2000-SFB-S
SFB 7	REV 06	711-044466	IPUCBA6CAA	MX2000-SFB-S
FPC 0	REV 39	750-045715	PROTOXCLEI	PROTO-ASSEMBLY
FPC 1	REV 11	750-045372	COUIBBNBAA	MX-MPC3E-3D
FPC 2	REV 17	750-037355	IPU3A4DHAA	MPC4E-3D-2CGE-8XGE
FPC 3	REV 05	750-044444	COUIBBGBAA	MX-MPC2E-3D-P
MIC 0	REV 28	750-028387	COUIA16BAA	MIC-3D-4XGE-XFP
FPC 4	REV 18	750-046005	PROTOXCLEI	PROTO-ASSEMBLY
FPC 5	REV 35	750-028467		MPC-3D-16XGE-SFPP
FPC 9	REV 30	750-044130	PROTOXCLEI	PROTO-ASSEMBLY
MIC 0	REV 05	750-049457	PROTOXCLEI	PROTO-ASSEMBLY
FPC 10	REV 36	750-044130	PROTOXCLEI	PROTO-ASSEMBLY
MIC 0	REV 06	750-049979	PROTOXCLEI	PROTO-ASSEMBLY
MIC 1	REV 12	750-050008	PROTOXCLEI	PROTO-ASSEMBLY
FPC 17	REV 28	750-044130	PROTOXCLEI	PROTO-ASSEMBLY
MIC 1	REV 03	750-050008	PROTOXCLEI	PROTO-ASSEMBLY
FPC 18	REV 39	750-045715	PROTOXCLEI	PROTO-ASSEMBLY
FPC 19	REV 39	750-045715	PROTOXCLEI	PROTO-ASSEMBLY
ADC 0	REV 17	750-043596	IPUCBA8CAA	MX2000-LC-ADAPTER
ADC 1	REV 17	750-043596	IPUCBA8CAA	MX2000-LC-ADAPTER
ADC 2	REV 17	750-043596	IPUCBA8CAA	MX2000-LC-ADAPTER
ADC 3	REV 17	750-043596	IPUCBA8CAA	MX2000-LC-ADAPTER
ADC 4	REV 02	750-043596	PROTOXCLEI	750-043596
ADC 5	REV 17	750-043596	IPUCBA8CAA	MX2000-LC-ADAPTER
ADC 18	REV 17	750-043596	IPUCBA8CAA	MX2000-LC-ADAPTER
ADC 19	REV 01	750-043596	PROTOXCLEI	750-043596
Fan Tray 0	REV 06	760-046960	IPUCBA5CAA	MX2000-FANTRAY-S
Fan Tray 1	REV 06	760-046960	IPUCBA5CAA	MX2000-FANTRAY-S
Fan Tray 2	REV 06	760-046960	IPUCBA5CAA	MX2000-FANTRAY-S

```
Fan Tray 3      REV 06   760-046960   IPUCBA5CAA      MX2000-FANTRAY-S
```

show chassis hardware (MX Series routers with ATM MIC)

```
user@host> show chassis hardware
Hardware inventory:
Item              Version  Part number  Serial number  Description
Chassis           REV 07   760-021404   ABAA5038      MX240
Midplane          REV 03   760-021392   ABBA2758      MX240 Backplane
FPM Board         Rev 01   740-022697   QCS0937C07K   Front Panel Display
AC in
PEM 0             Rev 01   740-022697   QCS0937C04X   PS 1.2-1.7kW; 100-240V
AC in
PEM 1             Rev 01   740-022697   QCS0937C06B   PS 1.2-1.7kW; 100-240V
AC in
PEM 2             Rev 01   740-022697   QCS0937C07U   PS 1.2-1.7kW; 100-240V
AC in
Routing Engine 0  REV 12   740-013063   9009042291    RE-S-2000
Routing Engine 1  REV 12   740-013063   9009042266    RE-S-2000
CB 0              REV 06   710-021523   ABBC1435      MX SCB
CB 1              REV 06   710-021523   ABBC1497      MX SCB
FPC 2             REV 14   750-031088   YH8446        MPC Type 2 3D Q
CPU               REV 06   711-030884   YH9612        MPC PMB 2G
MIC 0
MIC 1             REV 10   750-036132   ZP7062        2x0C12/8x0C3 CC-CE
PIC 2             BUILTIN  BUILTIN       2x0C12/8x0C3 CC-CE

Xcvr 0            NON-JNPR  23393-00492   UNKNOWN
Xcvr 1            NON-JNPR  23393-00500   UNKNOWN
Xcvr 2            NON-JNPR  23393-00912   UNKNOWN
Xcvr 3            REV 01   740-015638   22216-00575   Load SFP
Xcvr 4            REV 01   740-015638   24145-00110   Load SFP
Xcvr 5            REV 01   740-015638   24145-00016   Load SFP
Xcvr 6            REV 01   740-015638   24145-00175   Load SFP
Xcvr 7            NON-JNPR  23393-00627   UNKNOWN
QXM 0             REV 05   711-028408   YF4681        MPC QXM
QXM 1             REV 05   711-028408   YF4817        MPC QXM
Fan Tray 0        REV 01   710-021113   XL3645        MX240 Fan Tray
```

show chassis hardware (MX240, MX480, MX960 routers with Application Services Modular Line Card)

```
user@host> show chassis hardware
Hardware inventory:
Item              Version  Part number  Serial number  Description
Chassis           REV 03   710-013698   ACAA2362      MX960
Midplane          REV 03   710-014974   ZR0639        MX960 Backplane
FPM Board         Rev 03   740-013110   QCS152250SX   Front Panel Display
PDM               Rev 03   740-013683   QCS1512718W   Power Distribution Module
PEM 0             Rev 10   740-013683   QCS1512702Y   DC Power Entry Module
PEM 1             Rev 10   740-013683   QCS1512702Y   DC Power Entry Module
Routing Engine 0  REV 15   740-013063   9012024667    RE-S-2000
Routing Engine 1  REV 15   740-013063   9012024649    RE-S-2000
CB 0              REV 14   750-031391   ZJ7749        Enhanced MX SCB
CB 1              REV 14   750-031391   ZJ7750        Enhanced MX SCB
CB 2              REV 14   750-031391   ZY9233        Enhanced MX SCB
FPC 0             REV 17   750-031089   YR7434        MPC Type 2 3D
CPU
FPC 1             REV 11   750-037207   ZW9727        AS-MCC
CPU               REV 04   711-038173   ZW4817        AS-MCC-PMB
```

MIC 0	REV 01	750-037214	ZH3764	AS-MSC
PIC 0		BUILTIN	BUILTIN	AS-MSC
MIC 1	REV 01	711-028408	JZ9200	AS-MXC
PIC 2		BUILTIN	BUILTIN	AS-MXC
FPC 4	REV 30	750-028467	ABBN0232	MPC 3D 16x 10GE
CPU				
FPC 5	REV 04	750-037207	ZK9074	AS-MCC
CPU				
Fan Tray 0	REV 05	740-014971	VT5683	Fan Tray
Fan Tray 1	REV 05	740-014971	VT5684	Fan Tray

show chassis hardware extensive (MX240, MX480, MX960 Routers with Application Services Modular Line Card)

user@host> show chassis hardware extensive

```

ID: AS-MCC                                FRU Model Number: 750-037207
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 37 01 0b 52 45 56 20 31 31 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 33 37 32 30 37 00 00
Address 0x20: 53 2f 4e 20 5a 57 39 37 32 37 00 00 00 11 02 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 37
Address 0x50: 35 30 2d 30 33 37 32 30 37 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 31 31 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 5e ff ff ff ff ff ff ff ff ff ff ff ff
CPU                                REV 04    711-038173    ZW4817    AS-MCC-PMB
Jedec Code: 0x7fb0                EEPROM Version: 0x02
P/N: 711-038173                  S/N: ZW4817
Assembly ID: 0x0b38              Assembly Version: 01.04
Date: 12-30-2011                 Assembly Flags: 0x00
Version: REV 04
ID: AS-MCC-PMB
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 38 01 04 52 45 56 20 30 34 00 00
Address 0x10: 00 00 00 00 37 31 31 2d 30 33 38 31 37 33 00 00
Address 0x20: 53 2f 4e 20 5a 57 34 38 31 37 00 00 00 1e 0c 07
Address 0x30: db ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 00 50 52 4f 54 4f 58 43 4c 45 49 37
Address 0x50: 31 31 2d 30 33 38 31 37 33 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 30 34 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 60 00 00 00 00 00 00 00 00 00 00 00 00
MIC 0                                REV 01    750-037214    ZH3764    AS-MSC
Jedec Code: 0x7fb0                EEPROM Version: 0x02
P/N: 750-037214                  S/N: ZH3764
Assembly ID: 0x0a44              Assembly Version: 01.01
Date: 07-04-2011                 Assembly Flags: 0x00
Version: REV 01
ID: AS-MSC
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff I2C Hex Data:
Address 0x00: 7f b0 02 ff 0a 44 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 33 37 32 31 34 00 00
Address 0x20: 53 2f 4e 20 5a 48 33 37 36 34 00 00 00 04 07 07
Address 0x30: db ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 00 00 ff ff ff ff ff ff

```

```

Address 0x70: ff ff ff f6 c0 03 e1 bc 00 00 00 00 00 00 00
PIC 0          BUILTIN      BUILTIN      AS-MSC
FPC 4          REV 30      750-028467  ABBN0232  MPC 3D 16x 10GE
Jedec Code:    0x7fb0      EEPROM Version: 0x01

```

show chassis hardware (MX480 Router with MPC4E)

```

user@host> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               JN10FF57BAFB  MX480
Midplane      REV 05   750-047849   Good          MX480 Midplane
FPM Board     REV 02   710-017254   KG2066        Front Panel Display
PEM 0         Rev 03   740-017330   QCS081590BJ   PS 1.2-1.7kW; 100-240V
AC in
PEM 1         Rev 03   740-017330   QCS0815908Z   PS 1.2-1.7kW; 100-240V
AC in
PEM 2         Rev 03   740-029970   QCS1001U001    PS 1.4-2.52kW; 90-264V
AC in
Routing Engine 0 REV 05   740-031116   9009089502     RE-S-1800x4
Routing Engine 1 REV 05   740-031116   9009089624     RE-S-1800x4
CB 0          REV 02   750-031391   YE8506         Enhanced MX SCB
CB 1          REV 14   750-031391   ZK8265         Enhanced MX SCB
FPC 2         REV 05   750-037358   ZT0638         MPC4E 3D 32XGE
CPU           REV 07   711-035209   ZK3187         HMPD PMB 2G
PIC 0         BUILTIN BUILTIN       8X10GE SFPP
PIC 1         BUILTIN BUILTIN       8X10GE SFPP
PIC 2         BUILTIN BUILTIN       8X10GE SFPP
PIC 3         BUILTIN BUILTIN       8X10GE SFPP
FPC 3         REV 06   750-037355   CAAB1144       MPC4E 3D 2CGE+8XGE
CPU           REV 08   711-035209   CAAB1278       HMPD PMB 2G
PIC 0         BUILTIN BUILTIN       4x10GE SFPP
Xcvr 0        REV 01   740-031980   B11E01439      SFP+-10G-SR
Xcvr 1        REV 01   740-031980   B11D05809      SFP+-10G-SR
PIC 1         BUILTIN BUILTIN       1X100GE CFP
Xcvr 0        NON-JNPR D5418          UNKNOWN
PIC 2         BUILTIN BUILTIN       4x10GE SFPP
PIC 3         BUILTIN BUILTIN       1X100GE CFP
Xcvr 0        NON-JNPR X12J00362     CFP-100G-SR10
FPC 4         REV 12.3.10 750-033205   YR9445         MPCE Type 3 3D
CPU
Fan Tray                               Enhanced Left Fan Tray

```

show chassis hardware (MX2020 Router with MPC4E)

```

user@host> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               JN11E188CAFJ  MX2020
Midplane      REV 04   711-032387   ABAC7474       Lower Backplane
Midplane 1    REV 04   711-032386   ABAC7408       Upper Backplane
PMP 1         REV 03   711-032428   ACAJ1137       Upper Power Midplane
PMP 0         REV 03   711-032426   ACAJ1016       Lower Power Midplane
FPM Board     REV 06   760-040242   ABBT8832       Front Panel Display
PSM 3         REV 0C   740-033727   VK00255        DC 52V Power Supply
Module
PSM 4         REV 0C   740-033727   VJ00148        DC 52V Power Supply
Module
PSM 5         REV 0C   740-033727   VK00207        DC 52V Power Supply
Module
PSM 6         REV 0C   740-033727   VK00319        DC 52V Power Supply

```

Module				
PSM 7	REV 0C	740-033727	VK00264	DC 52V Power Supply
Module				
PSM 8	REV 0B	740-033727	VG00025	DC 52V Power Supply
Module				
PSM 13	REV 0C	740-033727	VK00274	DC 52V Power Supply
Module				
PSM 14	REV 0C	740-033727	VJ00167	DC 52V Power Supply
Module				
PSM 15	REV 0C	740-033727	VK00299	DC 52V Power Supply
Module				
PSM 16	REV 0C	740-033727	VK00213	DC 52V Power Supply
Module				
PSM 17	REV 0C	740-033727	VK00253	DC 52V Power Supply
Module				
PDM 0	REV 0B	740-038109	VJ00040	DC Power Dist Module
PDM 2	REV 0B	740-038109	VJ00025	DC Power Dist Module
Routing Engine 0	REV 02	740-041821	9009089735	RE-S-1800x4
Routing Engine 1	REV 02	740-041821	9009089731	RE-S-1800x4
CB 0	REV 04	750-040257	ZT2846	Control Board
CB 1	REV 04	750-040257	ZT2877	Control Board
SPMB 0	REV 01	711-041855	ZS2282	PMB Board
SPMB 1	REV 01	711-041855	ZS2261	PMB Board
SFB 0	REV 07	711-032385	ZZ2582	Switch Fabric Board
SFB 1	REV 04	711-032385	ZV4229	Switch Fabric Board
SFB 2	REV 07	711-032385	CAAB4902	Switch Fabric Board
SFB 3	REV 07	711-032385	CAAB4891	Switch Fabric Board
SFB 4	REV 07	711-032385	CAAB4883	Switch Fabric Board
SFB 5	REV 07	711-032385	CAAB4889	Switch Fabric Board
SFB 6	REV 06	711-032385	ZV1818	Switch Fabric Board
SFB 7	REV 07	711-032385	CAAB4897	Switch Fabric Board
FPC 0	REV 34	750-031090	ZT9799	MPC Type 2 3D EQ
CPU	REV 06	711-030884	ZS1122	MPC PMB 2G
MIC 0	REV 11	750-033535	CAAD7674	MIC-3D-10C192-XFP
PIC 0		BUILTIN	BUILTIN	MIC-3D-10C192-XFP
Xcvr 0	REV 01	740-014279	753019A00404	XFP-OC192-SR
MIC 1	REV 14	750-031967	ZM6103	MIC-3D-80C30C12-40C48
PIC 2		BUILTIN	BUILTIN	MIC-3D-80C30C12-40C48
Xcvr 0	REV 01	740-011615	PEF1AZP	SFP-IR
Xcvr 1	REV 01	740-011615	PEF1AZN	SFP-IR
Xcvr 2	REV 01	740-021308	ANA0N8S	SFP+-10G-SR
QXM 0	REV 06	711-028408	ZT9339	MPC QXM
QXM 1	REV 06	711-028408	ZT9237	MPC QXM
FPC 9	REV 34	750-031090	ZT9770	MPC Type 2 3D EQ
CPU	REV 06	711-030884	ZS1302	MPC PMB 2G
MIC 0	REV 24	750-028387	YJ3950	3D 4x 10GE XFP
PIC 0		BUILTIN	BUILTIN	2x 10GE XFP
Xcvr 0		NON-JNPR	T09M52516	XFP-10G-SR
Xcvr 1		NON-JNPR	CA49BK095	XFP-10G-SR
PIC 1		BUILTIN	BUILTIN	2x 10GE XFP
Xcvr 0	REV 02	740-014289	C834XU01T	XFP-10G-SR
Xcvr 1		NON-JNPR	T09M52515	XFP-10G-SR
MIC 1	REV 11	750-033535	CAAD7681	MIC-3D-10C192-XFP
PIC 2		BUILTIN	BUILTIN	MIC-3D-10C192-XFP
Xcvr 0	REV 01	740-014279	KBQ02BE	XFP-OC192-SR
QXM 0	REV 06	711-028408	ZT9151	MPC QXM
QXM 1	REV 06	711-028408	ZT9116	MPC QXM
FPC 10	REV 27	750-033205	ZL6215	MPCE Type 3 3D
CPU	REV 07	711-035209	ZK9038	HMPC PMB 2G
MIC 0	REV 18	750-028380	YG6885	3D 2x 10GE XFP
PIC 0		BUILTIN	BUILTIN	1x 10GE XFP

Xcvr 0	REV 01	740-014289	C706XU0AG	XFP-10G-SR
PIC 1		BUILTIN	BUILTIN	1x 10GE XFP
Xcvr 0	REV 02	740-014289	T08L84366	XFP-10G-SR
FPC 14	REV 09	750-037355	CAAF1534	MPC4E 3D 2CGE+8XGE
CPU	REV 08	711-035209	CAAB9879	HMPC PMB 2G
PIC 0		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-021308	21T511100436	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AHPOGPM	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	123363A00032	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	19T511100477	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0		NON-JNPR	X12J00260	CFP-100G-SR10
PIC 2		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-021308	21T511104086	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	21T511104627	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	21T511104644	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	1X100GE CFP
FPC 19	REV 32	750-028467	ZR2008	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ZT6933	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	19T511100291	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AMH02VE	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	23T511102128	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AMS15PP	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	123363A00716	SFP+-10G-SR
ADC 0	REV 05	750-043596	CAAC2072	Adapter Card
ADC 9	REV 01	750-043596	ZV4111	Adapter Card
ADC 10	REV 05	750-043596	CAAC2058	Adapter Card
ADC 14	REV 02	750-043596	ZW1561	Adapter Card
ADC 19	REV 01	750-043596	ZV4127	Adapter Card
Fan Tray 0	REV 03	760-046960	ACAY0124	172mm FanTray - 6 Fans
Fan Tray 1	REV 2A	760-046960	ACAY0022	172mm FanTray - 6 Fans
Fan Tray 2	REV 2A	760-046960	ACAY0023	172mm FanTray - 6 Fans
Fan Tray 3	REV 2A	760-046960	ACAY0025	172mm FanTray - 6 Fans

show chassis hardware (MX5, MX10, MX40, MX80, MX240, MX480, and MX960 Routers with Enhanced 20-Port Gigabit Ethernet MIC)

```

user@host> show chassis hardware
Hardware inventory:

```

Item	Version	Part number	Serial number	Description
Chassis			F3434	MX80-P
Midplane	REV 01	711-044315	ZK2681	MX80-P
PEM 0	Rev 04	740-028288	VE05267	AC Power Entry Module
PEM 1	Rev 04	740-028288	VE05270	AC Power Entry Module
Routing Engine		BUILTIN	BUILTIN	Routing Engine
TFEB 0		BUILTIN	BUILTIN	Forwarding Engine
Processor				
QXM 0	REV 05	711-028408	ZK0952	MPC QXM
FPC 0		BUILTIN	BUILTIN	MPC BUILTIN
MIC 0		BUILTIN	BUILTIN	4x 10GE XFP
PIC 0		BUILTIN	BUILTIN	4x 10GE XFP
FPC 1		BUILTIN	BUILTIN	MPC BUILTIN
MIC 0	REV 02	750-049846	CAAV2153	3D 20x 1GE(LAN)-E,SFP
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN) -E SFP
Xcvr 0	REV 01	740-011613	AM0816S9B81	SFP-SX
Xcvr 1	REV 02	740-011613	AM0925SBLK7	SFP-SX
Xcvr 2	REV 01	740-011613	UAQ0005	SFP-SX


```

Xcvr 3      REV 01  740-011613  UAQ000C      SFP-SX
Xcvr 4      REV 01  740-011613  P9F195E      SFP-SX
Xcvr 5      REV 01  740-011613  UAQ0003      SFP-SX
Xcvr 6      REV 01  740-031851  AM1041SU1LD  SFP-SX
Xcvr 8      REV 02  740-013111  B101501      SFP-T
PIC 1       BUILTIN  740-011613  BUILTIN      10x 1GE(LAN) -E SFP
Xcvr 0      REV 01  740-011613  PFM1ML7      SFP-SX
Xcvr 4      REV 01  740-011613  PE729P6      SFP-SX
Xcvr 6      REV 02  740-011613  AM1014SGC84  SFP-SX
Xcvr 9      REV 01  740-011613  AM0812S8UK3  SFP-SX
MIC 1       REV 26  750-028392  ZY0187       3D 20x 1GE(LAN) SFP
PIC 2       BUILTIN  740-011613  BUILTIN      10x 1GE(LAN) SFP
Xcvr 0      REV 01  740-011613  P9F1AN9      SFP-SX
Xcvr 5      REV 02  740-011613  AM1003SFUF4  SFP-SX
Xcvr 9      REV 01  740-031851  AM1041SU1LM  SFP-SX
PIC 3       BUILTIN  740-011613  BUILTIN      10x 1GE(LAN) SFP
Xcvr 4      REV 01  740-011613  PAJ4MYT      SFP-SX
Xcvr 7      +      NON-JNPR     XG32A024     SFP-SX
Xcvr 8      NON-JNPR PFROV6J      SFP-SX
Xcvr 9      REV 01  740-031851  AM1041SU02U  SFP-SX
Fan Tray

```

show chassis hardware models (MX5, MX10, MX40, MX80, MX240, MX480, and MX960 Routers with Enhanced 20-Port Gigabit Ethernet MIC)

```

user@host> show chassis hardware models
Hardware inventory:
Item          Version  Part number  Serial number  FRU model number
PEM 0         Rev 04   740-028288  VE05267       PWR-MX80-AC-S
PEM 1         Rev 04   740-028288  VE05270       PWR-MX80-AC-S
Routing Engine
TFEB 0        BUILTIN  BUILTIN
FPC 0         BUILTIN  BUILTIN
FPC 1         BUILTIN  BUILTIN
MIC 0         REV 02   750-049846  CAAV2153      MIC-3D-20GE-SFP-E
MIC 1         REV 26   750-028392  ZY0187        MIC-3D-20GE-SFP
Fan Tray      FANTRAY-MX80-S

```

show chassis hardware (T320 Router)

```

user@host> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis       19093    T320
Midplane      REV 04   710-004339  BC1436        T320 Backplane
FPM GBUS      REV 03   710-004461  BC1407        T320 FPM Board
FPM Display   REV 04   710-002897  BE0763        FPM Display
CIP           REV 05   710-002895  BB2311        T Series CIP
PEM 0         Rev 01   740-004359  NB12546       Power Entry Module
SCG 0         REV 06   710-004455  AY4522        T320 Sonet
Clock Gen.
Routing Engine 0
CB 0          REV 13   710-002728  BC1577        unknown
Control Board
CB 1          REV 13   710-002728  BC1595        T Series
Control Board
FPC 1         REV 09   710-007531  HS1572        FPC Type 2
CPU           REV 15   710-001726  HR8763        FPC CPU
PIC 0         REV 01   750-010618  CB5579        4x G/E SFP,
1000 BASE
SFP 0         REV 01   740-007326  P5809Z1       SFP-SX

```

```

SFP 1          REV 01  740-007326  P4Q10XU          SFP-SX
SFP 2          NON-JNPR  RA45020031      SFP-SX
SFP 3          NON-JNPR  RA45020032      SFP-SX
PIC 1          REV 01  750-010618  CD9587           4x G/E SFP,
1000 BASE
SFP 0          NON-JNPR  P5A08QZ          SFP-T
SFP 1          REV 01  740-007326  P4Q133K          SFP-SX
SFP 2          REV 01  740-007326  P5809YY          SFP-SX
SFP 3          REV 01  740-007327  4C81704          SFP-LX
MMB 1          REV 03  710-005555  HR9401           MMB-288mbit
PPB 0          REV 04  710-003758  HR2886           PPB Type 2
FPC 2          REV 07  710-005860  HP2392           FPC Type 1
CPU            REV 14  710-001726  HP7797           FPC CPU
PIC 0          REV 02  750-007643  HM0853           1x G/E QPP,
1000 BASE
SFP 0          REV 01  740-007326  P11E9JJ          SFP-SX
MMB 1          REV 02  710-005555  HN2379           MMB-288mbit
PPB 0          REV 04  710-003758  HP8092           PPB Type 2
FPC 3          REV 07  710-005860  HP2393           FPC Type 1
CPU            REV 14  710-001726  HP0968           FPC CPU
PIC 0          REV 01  750-010240  CB5363           1x G/E SFP,
1000 BASE
SFP 0          REV 01  740-007326  P4R0PNH          SFP-SX
PIC 1          REV 03  750-003034  HD2832           4x OC-3 SONET,
SMIR
MMB 1          REV 02  710-005555  HN6307           MMB-288mbit
PPB 0          REV 04  710-003758  HP5051           PPB Type 2
FPC 4          REV 01  710-010845  JD3872           FPC Type 4
CPU            REV 02  710-011481  JB6042           FPC CPU
5              REV 01  710-005802  BC1566           FPC Type 2
CPU            REV 09  710-001726  AY4922           FPC CPU
PIC 0          REV 02  750-008155  BE2114           2x G/E QPP,
1000 BASE
SFP 0          REV 01  740-007326  P4R0PMQ          SFP-SX
SFP 1          REV 01  740-007326  P4R0PN9          SFP-SX
PIC 1          REV 01  750-008155  BE2116           2x G/E QPP,
1000 BASE
SFP 0          REV 01  740-007326  P4R0PNZ          SFP-SX
SFP 1          NON-JNPR  2908             SFP-T
MMB 1          REV 01  710-005555  AZ2246           MMB-288mbit
PPB 0          REV 03  710-003758  AY4839           PPB Type 2
FPC 7          REV 01  710-005803  AZ2123           FPC Type 3
...

```

show chassis hardware (T640 Router)

```
user@host> show chassis hardware
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			19182	T640
Midplane	REV 04	710-002726	AX5608	T640 Backplane
FPM GBUS	REV 02	710-002901	HE3064	T640 FPM Board
FPM Display	REV 02	710-002897	HE7864	FPM Display
CIP	REV 05	710-002895	HA5024	T Series CIP
PEM 0	Rev 02	740-029522	VH26235	AC PEM 10kW US
PEM 1	Rev 02	740-029522	VH26230	AC PEM 10kW US
SCG 0	REV 03	710-003423	HA4508	T640 Sonet Clock Gen.
Routing Engine 0	REV 02	740-005022	210865700483	RE-3.0 (RE-600)
CB 0	REV 01	710-002728	HD3044	T Series Control Board
FPC 2	REV 04	710-001721	HD5572	FPC Type 3
CPU	REV 06	710-001726	HA4712	FPC CPU

PIC 1	REV 03	750-009567	HV2331	1x 10GE(LAN),XENPAK
SFP 0	REV 01	740-009898	USC202R103	XENPAK-SR
PIC 2	REV 03	750-009567	HV2332	1x 10GE(LAN),XENPAK
SFP 0	REV 01	740-011268	USC202R112	XENPAK-ZR
PIC 3	REV 03	750-009567	HX4416	1x 10GE(LAN),XENPAK
SFP 0	REV 01	740-012056	434TC004	XENPAK-CX4
PIC 4	REV 03	750-009567	HX4420	1x 10GE(LAN),XENPAK
SFP 0	REV 01	740-012058	434TC124	XENPAK-LX4
FPC 5	REV 01	710-013553	JE4839	E2-FPC Type 1
CPU	REV 01	710-013569	JW9163	FPC CPU
PIC 0	REV 01	750-009567	HX4419	1x 10GE(LAN),XENPAK
SFP 0	REV 01	740-009898	USC202RT05	XENPAK-LR
PIC 1	REV 03	750-009567	HN7426	1x 10GE(LAN),XENPAK
SFP 0	REV 01	740-009550	03L90051	XENPAK-ER
PIC 2	REV 03	750-009467	HT7423	1x 10GE(LAN),XENPAK
SFP 0		NON-JNPR		UNKNOWN
PIC 3	REV 04	750-005100	AY4850	1x 10GE(LAN),DWDM
FPC 4	REV 01	710-010845	JD3872	FPC Type 4
CPU	REV 02	710-011481	JB6042	FPC CPU
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Fan Tray

show chassis hardware models (T640 Router)

```
user@host> show chassis hardware models
Hardware inventory:
```

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 04	710-002726		CHAS-BP-T640-S
FPM Display	REV 02	710-002897		CRAFT-T640-S
CIP	REV 05	710-002895		CIP-L-T640-S
PEM 0	Rev 01	740-002595		PWR-T-DC-S
SCG 0	REV 04	710-003423		SCG-T-S
SCG 1	REV 04	710-003423		SCG-T-S
Routing Engine 0	REV 01	740-005022		RE-600-2048-S
Routing Engine 1	REV 07	740-005022		RE-600-2048-S
CB 0	REV 06	710-002726		CHAS-BP-T640-S
CB 1	REV 06	710-002728		CB-L-T-S
FPC 5	REV 05	710-007527		T640-FPC2
PIC 0	REV 05	750-002510		PB-2GE-SX
PIC 1	REV 05	750-001901		PB-40C12-SON-SMIR
FPC 6	REV 03	710-001721		T640-FPC3
PIC 1	REV 01	750-009553		PC-40C48-SON-SFP
SIB 4	REV 02	750-005486		SIB-I-T640-S
Fan Tray 0				FANTRAY-T-S
Fan Tray 1				FANTRAY-T-S
Fan Tray 2				FAN-REAR-TX-T640-S

show chassis hardware extensive (T640 Router)

```
user@host> show chassis hardware extensive
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis				T640
Jedec Code:	0x7fb0	EEPROM Version:	0x01	
P/N:	S/N:	
Assembly ID:	0x0507	Assembly Version:	00.00	
Date:	00-00-0000	Assembly Flags:	0x00	
Version:			
ID:	Gibson LCC Chassis			
Board Information Record:				

```

Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
Address 0x00: 7f b0 01 ff 05 07 00 00 00 00 00 00 00 00 00 00
Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x20: ff ff ff ff ff ff ff ff ff ff ff ff ff 00 00 00 00
Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Midplane          REV 04    710-002726    AX5633
Jedec Code:      0x7fb0          EEPROM Version:    0x01
P/N:             710-002726.      S/N:          AX5633.
Assembly ID:     0x0127          Assembly Version: 01.04
Date:           06-27-2001        Assembly Flags:  0x00
Version:         REV 04.....
ID: Gibson Backplane
Board Information Record:
Address 0x00: ad 01 08 00 00 90 69 0e f8 00 ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 01 27 01 04 52 45 56 20 30 34 00 00
Address 0x10: 00 00 00 00 37 31 30 2d 30 30 32 37 32 36 00 00
Address 0x20: 53 2f 4e 20 41 58 35 36 33 33 00 00 00 1b 06 07
Address 0x30: d1 ff ff ff ad 01 08 00 00 90 69 0e f8 00 ff ff
Address 0x40: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
FPM GBUS          REV 02    710-002901    HE3245
...
FPM Display      REV 02    710-002897    HA4873
...
CIP              REV 05    710-002895    HA4729
...
PEM 1            RevX02    740-002595    MD21815          Power Entry Module
...
SCG 0            REV 04    710-003423    HF6023
...
SCG 1            REV 04    710-003423    HF6061
...
Routing Engine 0 REV 01    740-005022    210865700292    RE-3.0
...
CB 0             REV 06    710-002728    HE3614
...
FPC 1            REV 01    710-002385    HE3009          FPC Type 1
...
...             REV 06    710-001726    HC0010

```

show chassis hardware (T4000 Router)

```

user@host> show chassis hardware
Hardware inventory:

```

Item	Version	Part number	Serial number	Description
Chassis			JN1172F25AHA	T4000
Midplane	REV 01	710-027486	RC8355	T-series Backplane
FPM GBUS	REV 13	710-002901	BBAE0927	T640 FPM Board
FPM Display	REV 01	710-021387	EF6764	T1600 FPM Display
CIP	REV 06	710-002895	BBAD9210	T-series CIP
PEM 0	REV 01	740-036442	VA00016	Power Entry Module 6x60
SCG 0	REV 18	710-003423	BBAD7248	T640 Sonet Clock Gen.
SCG 1	REV 18	710-003423	BBAE3874	T640 Sonet Clock Gen.
Routing Engine 0	REV 05	740-026941	P737F-002248	RE-DUO-1800
Routing Engine 1	REV 06	740-026941	P737F-002653	RE-DUO-1800
CB 0	REV 09	710-022597	ED0295	LCC Control Board
CB 1	REV 09	710-022597	EA6050	LCC Control Board
FPC 0	REV 26	750-032819	EK1173	FPC Type 5-3D
CPU	REV 12	711-030686	EJ8584	SNG PMB
PIC 0	REV 07	750-034624	EF6837	12x10GE (LAN/WAN) SFPP

Xcvr 0	REV 01	740-031980	123363A01145	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	123363A01147	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AJJ01P3	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B10M03256	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	AJJ01M2	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	123363A01137	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	AJJ01PN	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	AJJ01NW	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	123363A01139	SFP+-10G-SR
Xcvr 9	REV 01	740-031980	AJJ01KE	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	123363A01336	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	B10M01325	SFP+-10G-SR
PIC 1	REV 07	750-034624	EF6800	12x10GE (LAN/WAN) SFPP
Xcvr 0	REV 01	740-031980	AJJ01SA	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AJJ01QZ	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AJH0217	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AJJ01TE	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	AJJ01KV	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	AJJ01MU	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	AJJ01R0	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	AJJ01TC	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	AJJ0364	SFP+-10G-SR
Xcvr 9	REV 01	740-031980	AJD0GV3	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	B10M03343	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	AJJ01QJ	SFP+-10G-SR
LMB 0	REV 05	711-034381	EJ8490	Type-0 LMB
LMB 1	REV 04	711-035774	EJ8517	Type-1 LMB
LMB 2	REV 05	711-034381	EJ8489	Type-0 LMB
FPC 3	REV 07	750-032819	EG3637	FPC Type 5-3D
CPU	REV 09	711-030686	EG0150	SNG PMB
PIC 0	REV 08	750-035293	EF3657	1x100GE
Xcvr 0	REV 01	740-032210	C22CQNJ	CFP-100G-LR4
PIC 1	REV 10	750-034624	BBAN4098	12x10GE (LAN/WAN) SFPP
Xcvr 0	REV 01	740-031980	B11J04902	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11J04891	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AJJ01MX	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11J04183	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	B11J04894	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	B11J04184	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	B11J04897	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	B11J04899	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	AJJ01TV	SFP+-10G-SR
Xcvr 9	REV 01	740-031980	B11J04057	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	AJJ01M4	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	B11J04905	SFP+-10G-SR
LMB 0	REV 04	711-034381	EG1524	Type-0 LMB
LMB 1	REV 03	711-035774	EG0345	Type-1 LMB
LMB 2	REV 04	711-034381	EG1522	Type-0 LMB
FPC 5	REV 03	710-033871	BBAJ0768	FPC Type 4-ES
CPU	REV 11	710-016744	BBAH9342	ST-PMB2
PIC 0	REV 09	750-029262	EE6789	100GE
PIC 1	REV 03	750-034781	EE6655	100GE CFP
Xcvr 0	REV 01	740-032210	J11A22334	CFP-100G-LR4
BRIDGE 0	REV 03	711-029995	EE6572	100GE Bridge Board
MMB 0	REV 07	710-025563	BBAJ4657	ST-MMB2
MMB 1	REV 07	710-025563	BBAJ3073	ST-MMB2
FPC 6	REV 05	750-010153	EF4936	FPC Type 5-3D
CPU	REV 06	711-030686	EF4189	SNG PMB
PIC 0	REV 10	750-034624	BBAN4109	12x10GE (LAN/WAN) SFPP
Xcvr 0	REV 01	740-031980	B11J04895	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11J04898	SFP+-10G-SR

Xcvr 2	REV 01	740-031980	B11J04021	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11J04903	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	B11J04311	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	B11J04059	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	B11J04016	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	B11J04017	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	B11J04887	SFP+-10G-SR
Xcvr 9	REV 01	740-031980	B11J04297	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	B11J04893	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	B11J04022	SFP+-10G-SR
PIC 1	REV 02	750-034624	EE3711	12x10GE (LAN/WAN) SFPP
Xcvr 0	REV 01	740-031980	AJH033X	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AJJ01N0	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AJJ01SV	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AJJ032L	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	B10M01593	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	AJD0FF1	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	AJJ01NU	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	123363A01305	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	B10M00361	SFP+-10G-SR
Xcvr 9	REV 01	740-031980	AJJ01M7	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	AJJ032X	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	AJJ01PG	SFP+-10G-SR
LMB 0	REV 04	711-034381	EF3838	Type-0 LMB
LMB 1	REV 03	711-035774	EF3821	Type-1 LMB
LMB 2	REV 04	711-034381	EF3834	Type-0 LMB
SPMB 0	REV 05	710-023321	ED1990	LCC Switch CPU
SPMB 1	REV 05	710-023321	EA2768	LCC Switch CPU
SIB 0	REV 02	711-036340	EF8802	SIB-HC-3D
SIB 1	REV 07	711-036340	EG2286	SIB-HC-3D
SIB 2	REV 07	711-036340	EG2252	SIB-HC-3D
SIB 3	REV 02	711-036340	EF1358	SIB-HC-3D
SIB 4	REV 02	711-036340	EF8806	SIB-HC-3D
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
-- Rev 2				
Fan Tray 2				Rear Fan Tray -- Rev 3

show chassis hardware (T4000 Router with 16-GB Line Card Chassis (LCC) Routing Engine)

user@host> show chassis hardware

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN11BDF2CAHA	T1600
Midplane	REV 01	710-027486	ACAJ0774	T640 Backplane
FPM GBUS	REV 13	710-002901	BBAL6812	T640 FPM Board
FPM Display	REV 04	710-021387	BBAP2679	T1600 FPM Display
CIP	REV 06	710-002895	BBAP4758	T-series CIP
PEM 0	Rev 03	740-026384	XF86421	Power Entry Module 3x80
PEM 1	Rev 03	740-026384	XF86429	Power Entry Module 3x80
SCG 0	REV 18	710-003423	BBAP1896	T640 Sonet Clock Gen.
SCG 1	REV 18	710-003423	BBAN8659	T640 Sonet Clock Gen.
Routing Engine 0	REV 01	740-042243	737F-002238	RE-DUO-1800-16G
Routing Engine 1	REV 01	740-042243	737F-002403	RE-DUO-1800-16G
CB 1	REV 11	710-022597	EK4526	LCC Control Board
CB 1	REV 11	710-022597	EK4527	LCC Control Board
FPC 0	REV 05	710-033871	EK5644	FPC Type 4-ES
CPU	REV 11	710-016744	EK3428	ST-PMB2
PIC 0	REV 20	750-017405	EJ3041	4x 10GE (LAN/WAN) XFP
PIC 1	REV 17	750-026962	EH7536	10x10GE(LAN/WAN) SFPP
MMB 0	REV 07	710-025563	EK6039	ST-MMB2

MMB 1	REV 07	710-025563	EK6086	ST-MMB2
FPC 1	REV 05	710-033871	EK6583	FPC Type 4-ES
CPU	REV 11	710-016744	EK3401	ST-PMB2
PIC 0	REV 17	750-026962	EJ8948	10x10GE(LAN/WAN) SFPP
MMB 0	REV 07	710-025563	EK6202	ST-MMB2
MMB 1	REV 07	710-025563	EK6112	ST-MMB2
SPMB 1	REV 05	710-023321	EK4900	LCC Switch CPU
SIB 0	REV 11	710-013074	EK5958	SIB-I8-SF
SIB 1	REV 11	710-013074	EK4606	SIB-I8-SF
SIB 2	REV 11	710-013074	EK5971	SIB-I8-SF
SIB 3	REV 11	710-013074	EK4609	SIB-I8-SF
SIB 4	REV 11	710-013074	EK4602	SIB-I8-SF
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Fan Tray -- Rev 2

show chassis hardware (T4000 Router with LSR FPC)

```
user@host> show chassis hardware
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN1173A24AHA	T4000
FPC 3	REV	750-048373	AN7797	FPC Type 5-LSR
CPU	REV 10	711-030686	AN6649	SNG PMB
PIC 0	REV 07	750-034624	EF6830	12x10GE (LAN/WAN) SFPP

show chassis hardware clei-models (T4000 Router)

```
user@host> show chassis hardware clei-models
Hardware inventory:
```

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 01	710-027486	IPMJ700DRD	CHAS-BP-T1600-S
FPM Display	REV 01	710-021387		CRAFT-T1600-S
CIP	REV 06	710-002895		CIP-L-T640-S
PEM 0	REV 01	740-036442	IPUPAG6KAA	PWR-T-6-60-DC
SCG 0	REV 18	710-003423		SCG-T-S
SCG 1	REV 18	710-003423		SCG-T-S
Routing Engine 0	REV 05	740-026941		RE-DUO-C1800-8G-S
Routing Engine 1	REV 06	740-026941		RE-DUO-C1800-8G-S
CB 0	REV 09	710-022597		CB-LCC-S
CB 1	REV 09	710-022597		CB-LCC-S
FPC 3				
PIC 0	REV 08	750-035293	XXXXXXXXBB	PF-1CGE-CFP
PIC 1	REV 10	750-034624	XXXXXXXXCC	PF-12XGE-SFPP
FPC 5	REV 03	710-033871	IPUCAMBCTD	T1600-FPC4-ES
PIC 1	REV 03	750-034781	IPUIBKLMMA	PD-1CE-CFP-FPC4
FPC 6				
PIC 0	REV 10	750-034624	XXXXXXXXCC	PF-12XGE-SFPP
Fan Tray 0				FANTRAY-T-S
Fan Tray 1				FANTRAY-T4000-S
Fan Tray 2				FANTRAY-TXP-R-S

show chassis hardware detail (T4000 Router)

```
user@host> show chassis hardware detail
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN1172F25AHA	T4000
Midplane	REV 01	710-027486	RC8355	T-series Backplane
FPM GBUS	REV 13	710-002901	BBAE0927	T640 FPM Board
FPM Display	REV 01	710-021387	EF6764	T1600 FPM Display
CIP	REV 06	710-002895	BBAD9210	T-series CIP

PEM 0	REV 01	740-036442	VA00016	Power Entry Module 6x60
SCG 0	REV 18	710-003423	BBAD7248	T640 Sonet Clock Gen.
SCG 1	REV 18	710-003423	BBAE3874	T640 Sonet Clock Gen.
Routing Engine 0	REV 05	740-026941	P737F-002248	RE-DUO-1800
ad0 3823 MB	SMART CF		2009121602A661576157	Compact Flash
ad1 59690 MB	STEC MACH-8 SSD		STM000103FDB	Disk 1
Routing Engine 1	REV 06	740-026941	P737F-002653	RE-DUO-1800
ad0 3823 MB	SMART CF		201011150153F52CF52C	Compact Flash
ad1 62720 MB	SMART Lite SATA Drive		2010110900150A880A88	Disk 1
CB 0	REV 09	710-022597	ED0295	LCC Control Board
CB 1	REV 09	710-022597	EA6050	LCC Control Board
FPC 0	REV 26	750-032819	EK1173	FPC Type 5-3D
CPU	REV 12	711-030686	EJ8584	SNG PMB
PIC 0	REV 07	750-034624	EF6837	12x10GE (LAN/WAN) SFPP
Xcvr 0	REV 01	740-031980	123363A01145	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	123363A01147	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AJJ01P3	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B10M03256	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	AJJ01M2	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	123363A01137	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	AJJ01PN	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	AJJ01NW	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	123363A01139	SFP+-10G-SR
Xcvr 9	REV 01	740-031980	AJJ01KE	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	123363A01336	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	B10M01325	SFP+-10G-SR
PIC 1	REV 07	750-034624	EF6800	12x10GE (LAN/WAN) SFPP
Xcvr 0	REV 01	740-031980	AJJ01SA	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AJJ01QZ	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AJH0217	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AJJ01TE	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	AJJ01KV	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	AJJ01MU	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	AJJ01R0	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	AJJ01TC	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	AJJ0364	SFP+-10G-SR
Xcvr 9	REV 01	740-031980	AJDOG3V3	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	B10M03343	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	AJJ01QJ	SFP+-10G-SR
LMB 0	REV 05	711-034381	EJ8490	Type-0 LMB
LMB 1	REV 04	711-035774	EJ8517	Type-1 LMB
LMB 2	REV 05	711-034381	EJ8489	Type-0 LMB
FPC 3	REV 07	750-032819	EG3637	FPC Type 5-3D
CPU	REV 09	711-030686	EG0150	SNG PMB
PIC 0	REV 08	750-035293	EF3657	1x100GE
Xcvr 0	REV 01	740-032210	C22CQNJ	CFP-100G-LR4
PIC 1	REV 10	750-034624	BBAN4098	12x10GE (LAN/WAN) SFPP
Xcvr 0	REV 01	740-031980	B11J04902	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11J04891	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AJJ01MX	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11J04183	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	B11J04894	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	B11J04184	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	B11J04897	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	B11J04899	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	AJJ01TV	SFP+-10G-SR
Xcvr 9	REV 01	740-031980	B11J04057	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	AJJ01M4	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	B11J04905	SFP+-10G-SR
LMB 0	REV 04	711-034381	EG1524	Type-0 LMB
LMB 1	REV 03	711-035774	EG0345	Type-1 LMB

LMB 2	REV 04	711-034381	EG1522	Type-0 LMB
FPC 5	REV 03	710-033871	BBAJ0768	FPC Type 4-ES
CPU	REV 11	710-016744	BBAH9342	ST-PMB2
PIC 0	REV 09	750-029262	EE6789	100GE
PIC 1	REV 03	750-034781	EE6655	100GE CFP
Xcvr 0	REV 01	740-032210	J11A22334	CFP-100G-LR4
BRIDGE 0	REV 03	711-029995	EE6572	100GE Bridge Board
MMB 0	REV 07	710-025563	BBAJ4657	ST-MMB2
MMB 1	REV 07	710-025563	BBAJ3073	ST-MMB2
FPC 6	REV 05	750-010153	EF4936	FPC Type 5-3D
CPU	REV 06	711-030686	EF4189	SNG PMB
PIC 0	REV 10	750-034624	BBAN4109	12x10GE (LAN/WAN) SFPP
Xcvr 0	REV 01	740-031980	B11J04895	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11J04898	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11J04021	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11J04903	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	B11J04311	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	B11J04059	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	B11J04016	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	B11J04017	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	B11J04887	SFP+-10G-SR
Xcvr 9	REV 01	740-031980	B11J04297	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	B11J04893	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	B11J04022	SFP+-10G-SR
PIC 1	REV 02	750-034624	EE3711	12x10GE (LAN/WAN) SFPP
Xcvr 0	REV 01	740-031980	AJH033X	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AJJ01N0	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AJJ01SV	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AJJ032L	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	B10M01593	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	AJD0FF1	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	AJJ01NU	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	123363A01305	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	B10M00361	SFP+-10G-SR
Xcvr 9	REV 01	740-031980	AJJ01M7	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	AJJ032X	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	AJJ01PG	SFP+-10G-SR
LMB 0	REV 04	711-034381	EF3838	Type-0 LMB
LMB 1	REV 03	711-035774	EF3821	Type-1 LMB
LMB 2	REV 04	711-034381	EF3834	Type-0 LMB
SPMB 0	REV 05	710-023321	ED1990	LCC Switch CPU
SPMB 1	REV 05	710-023321	EA2768	LCC Switch CPU
SIB 0	REV 02	711-036340	EF8802	SIB-HC-3D
SIB 1	REV 07	711-036340	EG2286	SIB-HC-3D
SIB 2	REV 07	711-036340	EG2252	SIB-HC-3D
SIB 3	REV 02	711-036340	EF1358	SIB-HC-3D
SIB 4	REV 02	711-036340	EF8806	SIB-HC-3D
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
-- Rev 2				
Fan Tray 2				Rear Fan Tray -- Rev 3

show chassis hardware models (T4000 Router)

```
user@host> show chassis hardware models
```

Hardware inventory:

Item	Version	Part number	Serial number	FRU model number
Midplane	REV 01	710-027486	RC8355	CHAS-BP-T1600-S
FPM Display	REV 01	710-021387	EF6764	CRAFT-T1600-S
CIP	REV 06	710-002895	BBAD9210	CIP-L-T640-S

PEM 0	REV 01	740-036442	VA00016	PWR-T-6-60-DC
SCG 0	REV 18	710-003423	BBAD7248	SCG-T-S
SCG 1	REV 18	710-003423	BBAE3874	SCG-T-S
Routing Engine 0	REV 05	740-026941	P737F-002248	RE-DUO-C1800-8G-S
Routing Engine 1	REV 06	740-026941	P737F-002653	RE-DUO-C1800-8G-S
CB 0	REV 09	710-022597	ED0295	CB-LCC-S
CB 1	REV 09	710-022597	EA6050	CB-LCC-S
FPC 3				
PIC 0	REV 08	750-035293	EF3657	PF-1CGE-CFP
PIC 1	REV 10	750-034624	BBAN4098	PF-12XGE-SFPP
FPC 5	REV 03	710-033871	BBAJ0768	T1600-FPC4-ES
PIC 1	REV 03	750-034781	EE6655	PD-1CE-CFP-FPC4
FPC 6				
PIC 0	REV 10	750-034624	BBAN4109	PF-12XGE-SFPP
Fan Tray 0				FANTRAY-T-S
Fan Tray 1				FANTRAY-T4000-S
Fan Tray 2				FAN-REAR-TXP-LCC

show chassis hardware lcc (TX Matrix Router)

```
user@host> show chassis hardware lcc 0
lcc0-re0:
```

----- Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			65751	T640
Midplane	REV 03	710-005608	RA1408	T640 Backplane
FPM GBUS	REV 09	710-002901	RA2784	T640 FPM Board
FPM Display	REV 05	710-002897	RA2825	FPM Display
CIP	REV 06	710-002895	HT0684	T Series CIP
PEM 0	Rev 11	740-002595	PM18483	Power Entry Module
PEM 1	Rev 11	740-002595	qb13984	Power Entry Module
SCG 0	REV 11	710-003423	HT0022	T640 Sonet Clock Gen.
Routing Engine 0	REV 13	740-005022	210865700363	RE-3.0 (RE-600)
CB 0	REV 03	710-007655	HW1195	Control Board (CB-T)
CB 1	REV 05	710-007527	HM3245	FPC Type 2
CPU	REV 14	710-001726	HM1084	FPC CPU
PIC 0	REV 02	750-007218	AZ1112	2x OC-12 ATM2 IQ, SMIR
PIC 1	REV 02	750-007745	HG3462	4x OC-3 SONET, SMIR
PIC 2	REV 14	750-001901	BA5390	4x OC-12 SONET, SMIR
PIC 3	REV 09	750-008155	HS3012	2x G/E IQ, 1000 BASE
SFP 0		NON-JNPR	P1186TY	SFP-S
SFP 1	REV 01	740-007326	P11WLTF	SFP-SX
MMB 1	REV 02	710-005555	HL7514	MMB-288mbit
PPB 0	REV 04	710-003758	HM4405	PPB Type 2
PPB 1	REV 04	710-003758	AV1960	PPB Type 2
FPC 2	REV 08	710-010154	HZ3578	E-FPC Type 3
CPU	REV 05	710-010169	HZ3219	FPC CPU-Enhanced
PIC 0	REV 02	750-009567	HX2882	1x 10GE(LAN), XENPAK
SFP 0	REV 01	740-009898	USC202U709	XENPAK-LR
PIC 1	REV 03	750-003336	HJ9954	4x OC-48 SONET, SMSR
PIC 2	REV 01	750-004535	HC0235	1x OC-192 SM SR1
PIC 3	REV 07	750-007141	HX1699	10x 1GE(LAN), 1000 BASE
SFP 0	REV 01	740-007326	2441042	SFP-SX
SFP 1	REV 01	740-007326	2441027	SFP-SX
MMB 0	REV 03	710-010171	HV2365	MMB-5M3-288mbit
MMB 1	REV 03	710-010171	HZ3888	MMB-5M3-288mbit
SPMB 0	REV 09	710-003229	HW5245	T Series Switch CPU
SIB 3	REV 07	710-005781	HR5927	SIB-L8-F16
B Board	REV 06	710-005782	HR5971	SIB-L8-F16 (B)

SIB 4	REV 07	710-005781	HR5903	SIB-L8-F16
B Board	REV 06	710-005782	HZ5275	SIB-L8-F16 (B)

show chassis hardware scc (TX Matrix Router)

```
user@host> show chassis hardware scc
scc-re0:
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis				TX Matrix
Midplane	REV 04	710-004396	RB0014	SCC Midplane
FPM GBUS	REV 04	710-004617	HW9141	SCC FPM Board
FPM Display	REV 04	710-004619	HS5950	SCC FPM
CIP 0	REV 01	710-010218	HV9151	SCC CIP
CIP 1	REV 01	710-010218	HV9152	SCC CIP
PEM 1	Rev 11	740-002595	QB13977	Power Entry Module
Routing Engine 0	REV 05	740-008883	P11123900153	RE-4.0 (RE-1600)
CB 0	REV 01	710-011709	HR5964	Control Board (CB-TX)
SPMB 0	REV 09	710-003229	HW5293	T Series Switch CPU
SIB 3				
SIB 4	REV 01	710-005839	HW1177	SIB-S8-F16
B Board	REV 01	710-005840	HW1202	SIB-S8-F16 (B)

show chassis hardware (T1600 Router)

```
user@host> show chassis hardware
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			B2703	T1600
Midplane	REV 03	710-005608	RC4137	T640 Backplane
FPM GBUS	REV 10	710-002901	DT7062	T640 FPM Board
FPM Display	REV 05	710-002897	DS3067	FPM Display
CIP	REV 06	710-002895	DT3386	T-series CIP
PEM 0	Rev 07	740-017906	UA26344	Power Entry Module 3x80
PEM 1	Rev 18	740-002595	UF38441	Power Entry Module
SCG 0	REV 15	710-003423	DV0941	T640 Sonet Clock Gen.
Routing Engine 0	REV 08	740-014082	9009014502	RE-A-2000
Routing Engine 1	REV 07	740-014082	9009009591	RE-A-2000
CB 0	REV 05	710-007655	JA9360	Control Board (CB-T)
CB 1	REV 03	710-017707	DT3251	Control Board (CB-T)
FPC 0	REV 07	710-013558	DR4253	E2-FPC Type 2
CPU	REV 05	710-013563	DS3902	FPC CPU-Enhanced
PIC 0	REV 01	750-010618	CB5446	4x G/E SFP, 1000 BASE
Xcvr 0	REV 01	740-011613	P9F11CW	SFP-SX
Xcvr 1	REV 01	740-011613	P9F15C2	SFP-SX
Xcvr 2	REV 01	740-011782	PB94K0L	SFP-SX
PIC 1	REV 06	750-001900	HB6399	1x OC-48 SONET, SMSR
PIC 2	REV 14	750-001901	AP1092	4x OC-12 SONET, SMIR
PIC 3	REV 07	750-001900	AR8275	1x OC-48 SONET, SMSR
MMB 1	REV 07	710-010171	DS1524	MMB-5M3-288mbit
FPC 1	REV 06	710-013553	DL9067	E2-FPC Type 1
CPU	REV 04	710-013563	DM1685	FPC CPU-Enhanced
PIC 0	REV 08	750-001072	AB1688	1x G/E, 1000 BASE-SX
PIC 1	REV 10	750-012266	JX5519	4x 1GE(LAN), IQ2
Xcvr 0	REV 01	740-011613	AM0812S8UK6	SFP-SX
Xcvr 2	REV 01	740-011613	AM0812S8UK1	SFP-SX
Xcvr 3	REV 01	740-011782	P8N1YHG	SFP-SX
PIC 2	REV 22	750-005634	DP0083	1x CHOC12 IQ SONET, SMIR
MMB 1	REV 07	710-008923	DN1862	MMB 3M 288-bit

FPC 2	REV 01	710-005548	HJ9899	FPC Type 3
CPU	REV 06	710-001726	HC0586	FPC CPU
PIC 0	REV 16	750-007141	NC9660	10x 1GE(LAN), 1000 BASE
Xcvr 0	REV 01	740-011613	AM0812S8XAR	SFP-SX
Xcvr 1	REV 01	740-011782	P920E7B	SFP-SX
Xcvr 2	REV 01	740-011613	AM0812S8XAU	SFP-SX
Xcvr 4	REV 01	740-011613	AM0812S8XAK	SFP-SX
Xcvr 5	REV 01	740-011613	AM0812S8XAA	SFP-SX
Xcvr 6	REV 01	740-011613	PAJ4NKY	SFP-SX
Xcvr 7	REV 01	740-011613	AM0812S8UJW	SFP-SX
Xcvr 8	REV 01	740-011782	PB81X89	SFP-SX
Xcvr 9	REV 01	740-011613	AM0812S8UJX	SFP-SX
PIC 1	REV 06	750-015217	DK3280	8x 1GE(TYPE3), IQ2
Xcvr 0	REV 01	740-011782	P8P0A3T	SFP-SX
Xcvr 1	REV 01	740-013111	5090002	SFP-T
Xcvr 2	REV 01	740-011613	AM0814S93BQ	SFP-SX
Xcvr 4		NON-JNPR	PDE0FAN	SFP-SX
Xcvr 5	REV 01	740-011782	P8Q20XY	SFP-SX
Xcvr 6	REV 01	740-011613	AM0812S8UJV	SFP-SX
Xcvr 7	REV 01	740-011613	AM0812S8UP7	SFP-SX
PIC 2	REV 05	750-004695	HT4383	1x Tunnel
PIC 3	REV 17	750-009553	RL0204	4x OC-48 SONET
Xcvr 0	REV 01	740-011785	PDS3T23	SFP-SR
Xcvr 1	REV 01	740-011785	P6Q0F3E	SFP-SR
MMB 0	REV 03	710-004047	HD5843	MMB-288mbit
MMB 1	REV 03	710-004047	HE3208	MMB-288mbit
PPB 0	REV 02	710-002845	HA4524	PPB Type 3
PPB 1	REV 02	710-002845	HA4766	PPB Type 3
FPC 3	REV 01	710-010154	HR0863	E-FPC Type 3
CPU	REV 01	710-010169	HN3422	FPC CPU-Enhanced
PIC 0	REV 07	750-012793	WF5096	1x 10GE(LAN/WAN) IQ2
Xcvr 0		NON-JNPR	M64294TP	XFP-10G-LR
PIC 1	REV 25	750-007141	DV2127	10x 1GE(LAN), 1000 BASE
Xcvr 0	REV 01	740-011613	PFA6LTJ	SFP-SX
Xcvr 1	REV 01	740-011782	P9P0XV4	SFP-SX
Xcvr 2	REV 01	740-011782	P9M0TNX	SFP-SX
Xcvr 4	REV 01	740-011782	P9B0TTP	SFP-SX
Xcvr 5		NON-JNPR	PBS4LED	SFP-SX
PIC 2	REV 17	750-009553	RL0212	4x OC-48 SONET
Xcvr 0	REV 01	740-011785	PDS3T8G	SFP-SR
PIC 3	REV 32	750-003700	DL1279	1x OC-192 12xMM VSR
MMB 0	REV 01	710-010171	HR0821	MMB-288mbit
MMB 1	REV 01	710-010171	HR0818	MMB-288mbit
FPC 4	REV 16	710-013037	EB4919	FPC Type 4-ES
CPU	REV 09	710-016744	BBAA4382	ST-PMB2
PIC 0	REV 03	711-029996	EB1569	100GE
PIC 1	REV 05	711-029999	EB9983	100GE CFP
Xcvr 0	REV 0	740-032210	J10G80746	CFP-100G-LR4
BRIDGE 0	REV 02	711-029995	EB2235	100GE Bridge Board
MMB 0	REV 04	710-025563	BBAA7112	ST-MMB2
MMB 1	REV 04	710-025563	BBAA7149	ST-MMB2
FPC 5	REV 02	710-013037	DE3407	FPC Type 4-ES
CPU	REV 04	710-016744	DA2124	ST-PMB2
PIC 0	REV 16	750-012518	DF2554	4x OC-192 SONET XFP
Xcvr 0	REV 01	740-014279	AA0745N1FX8	XFP-OC192-SR
Xcvr 1	REV 01	740-014279	AA0748N1HN5	XFP-OC192-SR
Xcvr 2	REV 01	740-014279	AA0748N1HT6	XFP-OC192-SR
Xcvr 3	REV 01	740-014279	AA0744N1EC9	XFP-OC192-SR
PIC 1	REV 01	750-010850	JA0329	1x OC-768 SONET SR

MMB 0	REV 04	710-016036	DE9577	ST-MMB2
MMB 1	REV 04	710-016036	DK4060	ST-MMB2
FPC 6	REV 14	710-013037	DV1431	FPC Type 4-ES
CPU	REV 09	710-016744	DT9020	ST-PMB2
PIC 0	REV 11	750-017405	DM6261	4x 10GE (LAN/WAN) XFP
Xcvr 0	REV 01	740-014289	C701XU05Q	XFP-10G-SR
Xcvr 1	REV 01	740-014279	AA0748N1HPT	XFP-10G-LR
Xcvr 2	REV 01	740-014289	T08E19189	XFP-10G-SR
Xcvr 3	REV 01	740-014289	C715XU058	XFP-10G-SR
PIC 1	REV 13	750-017405	DP8772	4x 10GE (LAN/WAN) XFP
Xcvr 0	REV 02	740-011571	C850XJ037	XFP-10G-SR
Xcvr 1	REV 02	740-014289	C839XU0L9	XFP-10G-SR
Xcvr 2	REV 02	740-014289	C834XU05A	XFP-10G-SR
Xcvr 3	REV 02	740-014289	C810XU0CE	XFP-10G-SR
MMB 0	REV 01	710-025563	DT8454	ST-MMB2
MMB 1	REV 01	710-025563	DT8366	ST-MMB2
FPC 7	REV 09	710-007529	HZ7624	FPC Type 3
CPU	REV 15	710-001726	HZ1413	FPC CPU
PIC 0	REV 10	750-012793	DM5627	1x 10GE(LAN/WAN) IQ2
Xcvr 0	REV 02	740-011571	C831XJ062	XFP-10G-SR
PIC 1	REV 01	750-015217	JT6762	8x 1GE(TYPE3), IQ2
Xcvr 0	REV 01	740-011782	P8Q25JU	SFP-SX
Xcvr 1	REV 01	740-011782	P9B0U0K	SFP-SX
PIC 2	REV 01	750-015217	JS4268	8x 1GE(TYPE3), IQ2
Xcvr 0	REV 01	740-011613	AM0812S8XBZ	SFP-SX
Xcvr 1	REV 01	740-011613	AM0812S8XAP	SFP-SX
Xcvr 2	REV 01	740-011613	AM0812S8XBY	SFP-SX
Xcvr 3	REV 01	740-011613	AM0812S8XBX	SFP-SX
Xcvr 4	REV 01	740-011613	P9F1652	SFP-SX
Xcvr 5	REV 01	740-011782	P8Q21YC	SFP-SX
Xcvr 6	REV 01	740-011782	P8Q27HQ	SFP-SX
Xcvr 7	REV 01	740-011613	P8E2SSU	SFP-SX
PIC 3	REV 15	750-009450	NB6790	1x OC-192 SM SR2
MMB 0	REV 03	710-005555	HZ3450	MMB-288mbit
MMB 1	REV 03	710-005555	HZ3415	MMB-288mbit
PPB 0	REV 04	710-002845	HP0887	PPB Type 3
PPB 1	REV 04	710-002845	HW5255	PPB Type 3
SPMB 0	REV 10	710-003229	HX3699	T-series Switch CPU
SPMB 1	REV 12	710-003229	DT3091	T-series Switch CPU
SIB 0	REV 07	710-013074	DS4747	SIB-I8-SF
SIB 1	REV 07	710-013074	DS4942	SIB-I8-SF
SIB 2	REV 07	710-013074	DS4965	SIB-I8-SF
SIB 3	REV 07	710-013074	DS4990	SIB-I8-SF
SIB 4	REV 07	710-013074	DS4944	SIB-I8-SF
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Fan Tray -- Rev 2

show chassis hardware (TX Matrix Plus Router)

```
user@host> show chassis hardware
sfc0-re0:
```

```
-----
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN113186EAHB	TXP
Midplane	REV 05	710-022574	TS3822	SFC Midplane
FPM Display	REV 03	710-024027	DW4701	TXP FPM Display
CIP 0	REV 05	710-023792	DW7998	TXP CIP
CIP 1	REV 05	710-023792	DW7999	TXP CIP
PEM 0	Rev 04	740-027463	UM26367	Power Entry Module

PEM 1	Rev 04	740-027463	UM26346	Power Entry Module
Routing Engine 0	REV 06	740-026942	737A-1081	RE-DUO-2600
Routing Engine 1	REV 06	740-026942	737A-1043	RE-DUO-2600
CB 0	REV 05	710-022606	DW4435	SFC Control Board
CB 1	REV 09	710-022606	DW6100	SFC Control Board
SPMB 0		BUILTIN		SFC Switch CPU
SPMB 1		BUILTIN		SFC Switch CPU
SIB F13 0	REV 04	750-024564	DW5764	F13 SIB
B Board	REV 03	710-023431	DW9053	F13 SIB Mezz
SIB F13 3	REV 04	750-024564	DW5785	F13 SIB
B Board	REV 03	710-023431	DW9030	F13 SIB Mezz
SIB F13 6				
SIB F13 8	REV 04	750-024564	DW5752	F13 SIB
B Board	REV 03	710-023431	DW9051	F13 SIB Mezz
SIB F13 11	REV 04	750-024564	DW5782	F13 SIB
B Board	REV 03	710-023431	DW9058	F13 SIB Mezz
SIB F13 12	REV 03	750-024564	DT9466	F13 SIB
B Board	REV 02	710-023431	DT6556	F13 SIB Mezz
SIB F2S 0/0	REV 05	710-022603	DW7898	F2S SIB
B Board	REV 05	710-023787	DW7625	F2S SIB Mezz
SIB F2S 0/2	REV 05	710-022603	DW7811	F2S SIB
B Board	REV 05	710-023787	DW7550	F2S SIB Mezz
SIB F2S 0/4	REV 04	710-022603	DW4873	F2S SIB
B Board	REV 05	710-023787	DW8509	F2S SIB Mezz
SIB F2S 0/6	REV 04	710-022603	DW4867	F2S SIB
B Board	REV 05	710-023787	DW8472	F2S SIB Mezz
SIB F2S 1/0	REV 04	710-022603	DW4871	F2S SIB
B Board	REV 05	710-023787	DW8497	F2S SIB Mezz
SIB F2S 1/2	REV 05	710-022603	DW7868	F2S SIB
B Board	REV 05	710-023787	DW7551	F2S SIB Mezz
SIB F2S 1/4	REV 04	710-022603	DW4854	F2S SIB
B Board	REV 05	710-023787	DW8496	F2S SIB Mezz
SIB F2S 1/6	REV 05	710-022603	DW7889	F2S SIB
B Board	REV 05	710-023787	DW7496	F2S SIB Mezz
SIB F2S 2/0	REV 04	710-022603	DW4852	F2S SIB
B Board	REV 05	710-023787	DW8498	F2S SIB Mezz
SIB F2S 2/2	REV 04	710-022603	DW4845	F2S SIB
B Board	REV 05	710-023787	DW8457	F2S SIB Mezz
SIB F2S 2/4	REV 05	710-022603	DW7802	F2S SIB
B Board	REV 05	710-023787	DW7562	F2S SIB Mezz
SIB F2S 2/6	REV 04	710-022603	DW4822	F2S SIB
B Board	REV 05	710-023787	DW8467	F2S SIB Mezz
SIB F2S 3/0	REV 05	710-022603	DW7815	F2S SIB
B Board	REV 05	710-023787	DW7518	F2S SIB Mezz
SIB F2S 3/2	REV 03	710-022603	DV0068	F2S SIB
B Board	REV 03	710-023787	DT9974	F2S SIB Mezz
SIB F2S 3/4	REV 05	710-022603	DW7874	F2S SIB
B Board	REV 05	710-023787	DW7601	F2S SIB Mezz
SIB F2S 3/6	REV 03	710-022603	DV0033	F2S SIB
B Board	REV 03	710-023787	DT9969	F2S SIB Mezz
SIB F2S 4/0	REV 03	710-022603	DV0043	F2S SIB
B Board	REV 03	710-023787	DT9948	F2S SIB Mezz
SIB F2S 4/2	REV 05	710-022603	DW5446	F2S SIB
B Board	REV 05	710-023787	DW7611	F2S SIB Mezz
SIB F2S 4/4	REV 04	710-022603	DW4826	F2S SIB
B Board	REV 05	710-023787	DW8458	F2S SIB Mezz
SIB F2S 4/6	REV 03	710-022603	DV0026	F2S SIB
B Board	REV 03	710-023787	DT9963	F2S SIB Mezz
Fan Tray 0	REV 02	760-024497	DR8290	Front Fan Tray
Fan Tray 1	REV 02	760-024497	DR8293	Front Fan Tray
Fan Tray 2	REV 05	760-024502	DR8280	Rear Fan Tray

Fan Tray 3				
Fan Tray 4	REV 05	760-024502	DR8276	Rear Fan Tray
Fan Tray 5	REV 02	760-024502	DP5643	Rear Fan Tray

lcc0-re0:

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN11036F8AHA	T1600
Midplane	REV 03	710-017247	RC3799	T-series Backplane
FPM GBUS	REV 10	710-002901	DP7009	T640 FPM Board
FPM Display	REV 01	710-021387	DN7026	T1600 FPM Display
CIP	REV 06	710-002895	DP6024	T-series CIP
PEM 1	Rev 02	740-023211	WA50019	Power Entry Module 4x60A
SCG 0	REV 15	710-003423	DR6757	T640 Sonet Clock Gen.
SCG 1	REV 15	710-003423	DS2225	T640 Sonet Clock Gen.
Routing Engine 0	REV 01	740-026941	737F-1040	RE-DUO-1800
Routing Engine 1	REV 01	740-026941	737F-1016	RE-DUO-1800
CB 0	REV 06	710-022597	DX4011	LCC Control Board
CB 1	REV 06	710-022597	DX4017	LCC Control Board
FPC 1	REV 07	710-013035	DN5847	FPC Type 3-ES
CPU	REV 08	710-016744	DP2570	ST-PMB2
PIC 0	REV 05	750-015217	DB0418	8x 1GE(TYPE3), IQ2
Xcvr 0	REV 01	740-011782	P8Q27ZG	SFP-SX
Xcvr 1		NON-JNPR	PDA1U0D	SFP-SX
Xcvr 2	REV 01	740-011613	P9F1ALW	SFP-SX
Xcvr 3	REV 01	740-011782	PBA403V	SFP-SX
Xcvr 4		NON-JNPR	PDE09DP	SFP-SX
Xcvr 5	REV 01	740-011782	PCH2P4K	SFP-SX
Xcvr 6	REV 01	740-011782	PB94K0F	SFP-SX
Xcvr 7	REV 01	740-011782	PBA2R2A	SFP-SX
PIC 1	REV 03	750-004424	HJ4020	1x 10GE(LAN), DWDM
PIC 2	REV 01	750-003336	HG6073	4x OC-48 SONET, SMSR
MMB 0	REV 04	710-016036	DP3401	ST-MMB2
FPC 3	REV 12	710-013037	DR1169	FPC Type 4-ES
CPU	REV 08	710-016744	DP9429	ST-PMB2
PIC 0	REV 02	750-010850	JA0332	1x OC-768 SONET SR
MMB 0	REV 04	710-016036	DR0628	ST-MMB2
MMB 1	REV 04	710-016036	DR0592	ST-MMB2
FPC 4	REV 05	710-021534	DR7350	FPC Type 1-ES
CPU	REV 08	710-016744	DP8096	ST-PMB2
PIC 0	REV 04	750-014627	DP9171	4x OC-3 1x OC-12 SFP
Xcvr 0	REV 02	740-011615	PDE2RVR	SFP-SR
PIC 1	REV 22	750-005634	DS5815	1x CHOC12 IQ SONET, SMIR
PIC 2	REV 09	750-002911	CF4539	4x F/E, 100 BASE-TX
PIC 3	REV 08	750-021652	DR2827	1x CHOC12 IQE SONET
Xcvr 0		NON-JNPR	8	UNKNOWN
MMB 0	REV 04	710-016036	DR0809	ST-MMB2
FPC 5	REV 07	710-007529	HS5608	FPC Type 3
CPU	REV 15	710-001726	HX4351	FPC CPU
PIC 0	REV 14	750-009567	WJ8961	1x 10GE(LAN), XENPAK
Xcvr 0	REV 01	740-013170	J05K05961	XENPAK-LR
PIC 1	REV 16	750-007141	JJ8146	10x 1GE(LAN), 1000 BASE
Xcvr 1	REV 01	740-011613	P9F117T	SFP-SX
Xcvr 2	REV 01	740-011782	PBA2VCL	SFP-SX
Xcvr 3	REV 01	740-011782	PB83DRB	SFP-SX
Xcvr 4	REV 01	740-011613	AM0812S8UP8	SFP-SX
PIC 2	REV 12	750-009567	WF3566	1x 10GE(LAN), XENPAK
Xcvr 0	REV 02	740-013170	T07C94489	XENPAK-LR

MMB 0	REV 03	710-005555	HZ1907	MMB-288mbit
MMB 1	REV 03	710-005555	HW5283	MMB-288mbit
PPB 0	REV 04	710-002845	HZ7717	PPB Type 3
PPB 1	REV 04	710-002845	HS0110	PPB Type 3
FPC 6	REV 07	710-013035	DP7486	FPC Type 3-ES
CPU	REV 08	710-016744	DP2545	ST-PMB2
PIC 0	REV 09	750-009567	NE6323	1x 10GE(LAN), XENPAK
Xcvr 0	REV 02	740-013170	T09C71959	XENPAK-LR
PIC 1	REV 06	750-015217	DN4775	8x 1GE(TYPE3), IQ2
Xcvr 0	REV 01	740-011782	P7E0T6M	SFP-SX
Xcvr 1	REV 01	740-011613	AM0812S8XAY	SFP-SX
Xcvr 2	REV 01	740-011782	P7E0T6J	SFP-SX
Xcvr 3	REV 01	740-011782	PCH2P7D	SFP-SX
Xcvr 4	REV 01	740-011782	P9B0QYT	SFP-SX
Xcvr 5	REV 01	740-011613	AM0812S8WQJ	SFP-SX
Xcvr 6	REV 02	740-013111	9301220	SFP-T
Xcvr 7	REV 01	740-011782	P9B0TZ5	SFP-SX
PIC 2	REV 06	750-015217	DM6747	8x 1GE(TYPE3), IQ2
Xcvr 0	REV 01	740-011613	PAP0ZB2	SFP-SX
Xcvr 1	REV 01	740-013111	70191002	SFP-T
Xcvr 6	REV 01	740-011782	PBA29H8	SFP-SX
Xcvr 7	REV 01	740-011613	AM0812S8WQG	SFP-SX
MMB 0	REV 04	710-016036	DP3238	ST-MMB2
FPC 7	REV 03	710-021540	DV3154	FPC Type 2-ES
CPU	REV 09	710-016744	DT9053	ST-PMB2
PIC 0	REV 13	750-001901	HB4225	4x OC-12 SONET, SMIR
PIC 1	REV 05	750-001900	AD3644	1x OC-48 SONET, SMSR
PIC 2	REV 10	750-008155	HV0335	2x G/E IQ, 1000 BASE
Xcvr 0	REV 01	740-011782	PCH2UKF	SFP-SX
Xcvr 1	REV 01	740-011782	PCH2V19	SFP-SX
PIC 3	REV 03	750-014638	JS9493	1x OC-48-12-3 SFP
Xcvr 0	REV 01	740-011785	P6Q0ENK	SFP-SR
MMB 0	REV 05	710-016036	DP3323	ST-MMB2
SPMB 0	REV 04	710-023321	DX3004	LCC Switch CPU
SPMB 1	REV 04	710-023321	DX3009	LCC Switch CPU
SIB 0	REV 07	710-022594	DW4195	LCC SIB
B Board	REV 07	710-023185	DW3930	LCC SIB Mezz
SIB 1	REV 07	710-022594	DW4179	LCC SIB
B Board	REV 07	710-023185	DW3919	LCC SIB Mezz
SIB 2				
SIB 3	REV 06	710-022594	DT8251	LCC SIB
B Board	REV 06	710-023185	DT5792	LCC SIB Mezz
SIB 4	REV 08	710-022594	DW8014	LCC SIB
B Board	REV 07	710-023185	DW3917	LCC SIB Mezz
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Fan Tray -- Rev 3

lcc1-re0:

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN1102270AHA	T1600
Midplane	REV 04	710-017247	RC5358	T-series Backplane
FPM GBUS	REV 10	710-002901	DS3443	T640 FPM Board
FPM Display	REV 01	710-021387	DS6411	T1600 FPM Display
CIP	REV 06	710-002895	DS4235	T-series CIP
PEM 0	Rev 02	740-023211	VM82438	Power Entry Module 4x60A
SCG 0	REV 15	710-003423	DS6649	T640 Sonet Clock Gen.
SCG 1	REV 15	710-003423	DR6775	T640 Sonet Clock Gen.
Routing Engine 0	REV 01	740-026941	737F-1083	RE-DU0-1800

Routing Engine 1	REV 01	740-026941	737F-1104	RE-DUO-1800
CB 0	REV 06	710-022597	DW8542	LCC Control Board
CB 1	REV 06	710-022597	DW8530	LCC Control Board
FPC 0	REV 02	710-010845	JE2392	FPC Type 4
CPU	REV 02	710-011481	JF6820	FPC CPU-Enhanced
PIC 0	REV 11	750-017405	DP7259	4x 10GE (LAN/WAN) XFP
Xcvr 0	REV 01	740-014279	AA0741N1C8T	XFP-10G-LR
Xcvr 1	REV 01	740-014279	AA0746N1GAM	XFP-10G-LR
Xcvr 2	REV 01	740-014279	AA0747N1H0B	XFP-10G-LR
Xcvr 3	REV 01	740-014279	AA0748N1HZ5	XFP-10G-LR
MMB 0	REV 03	710-010842	HY7601	ST-MMB
FPC 1	REV 16	710-013037	BBAA7398	FPC Type 4-ES
CPU	REV 09	710-016744	BBAA2329	ST-PMB2
PIC 0	REV 03	711-029996	EB1575	100GE
PIC 1	REV 06	750-034781	EB9980	100GE CFP
MMB 0	REV 04	710-025563	BBAA5325	ST-MMB2
MMB 1	REV 04	710-025563	BBAA5444	ST-MMB2
FPC 2	REV 16	710-013037	BBAA7185	FPC Type 4-ES
CPU	REV 09	710-016744	BBAA3522	ST-PMB2
PIC 0	REV 03	711-029996	EB1557	100GE
PIC 1	REV 05	750-034781	EB4660	100GE CFP
Xcvr 0	REV 0	740-032210	J10F73666	CFP-100G-LR4
BRIDGE 0	REV 02	711-029995	EB2237	100GE Bridge Board
MMB 0	REV 04	710-025563	BBAA5347	ST-MMB2
MMB 1	REV 04	710-025563	BBAA5401	ST-MMB2
FPC 3	REV 10	710-021534	DZ0941	FPC Type 1-ES
CPU	REV 09	710-016744	DY6364	ST-PMB2
PIC 0	REV 13	750-012266	DK9192	4x 1GE(LAN), IQ2
Xcvr 0	REV 01	740-011613	AM0812S8WVD	SFP-SX
Xcvr 1		NON-JNPR	PDD63Q4	SFP-SX
Xcvr 2		NON-JNPR	PDE4G54	SFP-SX
Xcvr 3		NON-JNPR	PD40MAG	SFP-SX
PIC 1	REV 01	750-007641	HJ2003	1x G/E IQ, 1000 BASE
Xcvr 0	REV 01	740-011613	AM0812S8WVG	SFP-SX
PIC 3	REV 17	750-007444	JB6873	1x CHSTM1 IQ SDH, SMIR
MMB 0	REV 04	710-025563	DZ0281	ST-MMB2
FPC 4	REV 06	710-013035	DK0614	FPC Type 3-ES
CPU	REV 07	710-016744	DK1616	ST-PMB2
PIC 0	REV 22	750-007141	DM1870	10x 1GE(LAN), 1000 BASE
Xcvr 0	REV 01	740-011782	PCL3UKW	SFP-SX
Xcvr 1	REV 01	740-011782	P7E0T73	SFP-SX
Xcvr 2	REV 01	740-007326	P4TOWLR	SFP-SX
Xcvr 3	REV 01	740-011782	PAR1LRL	SFP-SX
Xcvr 4	REV 01	740-011782	P9MOU3Z	SFP-SX
Xcvr 5	REV 01	740-011782	P9MOU0C	SFP-SX
Xcvr 6	REV 01	740-011782	P9MOTLG	SFP-SX
Xcvr 7	REV 01	740-011782	P9MOU0F	SFP-SX
Xcvr 8	REV 01	740-011613	PFA6LAP	SFP-SX
Xcvr 9	REV 01	740-011782	PCH2P0U	SFP-SX
PIC 1	REV 16	750-009450	CV2565	1x OC-192 SM SR2
PIC 2	REV 05	750-004424	HH3057	1x 10GE(LAN), 10GBASE-LR
PIC 3	REV 12	750-013423	DP0403	MultiServices 500
MMB 0	REV 04	710-016036	DK1988	ST-MMB2
FPC 5	REV 07	710-013560	DR0004	E2-FPC Type 3
CPU	REV 05	710-013563	DR0089	FPC CPU-Enhanced
PIC 0	REV 11	750-012793	DR6107	1x 10GE(LAN/WAN) IQ2
Xcvr 0	REV 01	740-014289	C743XU074	XFP-10G-SR
PIC 1	REV 01	750-004695	HD5980	1x Tunnel
PIC 2	REV 32	750-003700	DL3770	1x OC-192 12xMM VSR

PIC 3	REV 12	750-009553	WB8901	4x OC-48 SONET
Xcvr 0	REV 01	740-011785	P9D1GTQ	SFP-SR
Xcvr 1	REV 01	740-011785	PDSOMMB	SFP-SR
Xcvr 3	REV 01	740-011785	PDE1KXP	SFP-SR
MMB 0	REV 07	710-010171	DP7374	MMB-5M3-288mbit
MMB 1	REV 07	710-010171	DP7404	MMB-5M3-288mbit
FPC 6	REV 07	710-013035	DM0994	FPC Type 3-ES
CPU	REV 07	710-016744	DM3651	ST-PMB2
PIC 0	REV 07	750-015217	DN4743	8x 1GE(TYPE3), IQ2
Xcvr 3	REV 01	740-011613	AM0812S8XB0	SFP-SX
Xcvr 4	REV 01	740-011782	PB829RB	SFP-SX
Xcvr 5	REV 01	740-011782	P8J1SYX	SFP-SX
PIC 1	REV 03	750-003336	HJ9954	4x OC-48 SONET, SMSR
PIC 3	REV 02	750-012793	JM7665	1x 10GE(LAN/WAN) IQ2
MMB 0	REV 04	710-016036	DN6913	ST-MMB2
FPC 7	REV 08	710-010845	JM3958	FPC Type 4
CPU	REV 04	710-011481	JK3669	FPC CPU-Enhanced
PIC 0	REV 11	750-017405	DP8837	4x 10GE (LAN/WAN) XFP
Xcvr 1	REV 01	740-014279	753019A00277	XFP-10G-LR
Xcvr 2	REV 02	740-011571	C850XJ00P	XFP-10G-SR
Xcvr 3	REV 01	740-014279	AA0813N1RTG	XFP-10G-LR
MMB 0	REV 04	710-010842	JN1971	ST-MMB
SPMB 0	REV 04	710-023321	DW3629	LCC Switch CPU
SPMB 1	REV 04	710-023321	DW3621	LCC Switch CPU
SIB 0	REV 07	710-022594	DW4200	LCC SIB
B Board	REV 07	710-023185	DW3932	LCC SIB Mezz
SIB 1	REV 07	710-022594	DW4193	LCC SIB
B Board	REV 07	710-023185	DW3904	LCC SIB Mezz
SIB 2				
SIB 3	REV 07	710-022594	DW4210	LCC SIB
B Board	REV 06	710-023185	DT5780	LCC SIB Mezz
SIB 4	REV 08	710-022594	DW8019	LCC SIB
B Board	REV 06	710-023185	DT5795	LCC SIB Mezz
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Fan Tray -- Rev 3

show chassis hardware sfc (TX Matrix Plus Router)

```
user@host> show chassis hardware sfc 0
sfc0-re0:
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN112F007AHB	TXP
Midplane	REV 05	710-022574	TS4027	SFC Midplane
FPM Display	REV 03	710-024027	DX0282	TXP FPM Display
CIP 0	REV 04	710-023792	DW4889	TXP CIP
CIP 1	REV 04	710-023792	DW4887	TXP CIP
PEM 0	Rev 07	740-027463	UM26368	Power Entry Module
Routing Engine 0	REV 01	740-026942	737A-1064	SFC RE
Routing Engine 1	REV 01	740-026942	737A-1082	SFC RE
CB 0	REV 09	710-022606	DW6099	SFC Control Board
CB 1	REV 09	710-022606	DW6096	SFC Control Board
SPMB 0		BUILTIN		SFC Switch CPU
SPMB 1		BUILTIN		SFC Switch CPU
SIB F13 0	REV 04	710-022600	DX0841	F13 SIB
B Board	REV 03	710-023431	DX0966	F13 SIB Mezz
SIB F13 1	REV 04	750-024564	DW5776	F13 SIB
B Board	REV 03	710-023431	DW9028	F13 SIB
SIB F13 3	REV 04	750-024564	DW5762	F13 SIB

B Board	REV 03	710-023431	DW9059	F13 SIB
SIB F13 4	REV 04	750-024564	DW5797	F13 SIB
B Board	REV 03	710-023431	DW9041	F13 SIB
SIB F13 6	REV 04	750-024564	DW5770	F13 SIB
B Board	REV 03	710-023431	DW9079	F13 SIB Mezz
SIB F13 7	REV 04	750-024564	DW5758	F13 SIB
B Board	REV 03	710-023431	DW9047	F13 SIB
SIB F13 8	REV 04	750-024564	DW5761	F13 SIB
B Board	REV 03	710-023431	DW9043	F13 SIB Mezz
SIB F13 9	REV 04	750-024564	DW5754	F13 SIB
B Board	REV 03	710-023431	DW9078	F13 SIB Mezz
SIB F13 11	REV 04	710-022600	DX0826	F13 SIB
B Board	REV 03	710-023431	DX0967	F13 SIB Mezz
SIB F13 12	REV 04	750-024564	DW5794	F13 SIB
B Board	REV 03	710-023431	DW9044	F13 SIB Mezz
SIB F2S 0/0	REV 05	710-022603	DW7897	F2S SIB
B Board	REV 05	710-023787	DW7657	NEO PMB
SIB F2S 0/2	REV 05	710-022603	DW7833	F2S SIB
B Board	REV 05	710-023787	DW7526	NEO PMB
SIB F2S 0/4	REV 05	710-022603	DW7875	F2S SIB
B Board	REV 05	710-023787	DW7588	NEO PMB
SIB F2S 0/6	REV 05	710-022603	DW7860	F2S SIB
B Board	REV 05	710-023787	DW7589	NEO PMB
SIB F2S 1/0	REV 04	710-022603	DW4820	F2S SIB
B Board	REV 05	710-023787	DW8510	NEO PMB
SIB F2S 1/2	REV 05	710-022603	DW7849	F2S SIB
B Board	REV 05	710-023787	DW7525	NEO PMB
SIB F2S 1/4	REV 05	710-022603	DW7927	F2S SIB
B Board	REV 05	710-023787	DW7556	F2S SIB Mezz
SIB F2S 1/6	REV 05	710-022603	DW7866	F2S SIB
B Board	REV 05	710-023787	DW7651	NEO PMB
SIB F2S 2/0	REV 05	710-022603	DW7880	F2S SIB
B Board	REV 05	710-023787	DW7523	NEO PMB
SIB F2S 2/2	REV 05	710-022603	DW7895	F2S SIB
B Board	REV 05	710-023787	DW7591	NEO PMB
SIB F2S 2/4	REV 05	710-022603	DW7907	F2S SIB
B Board	REV 05	710-023787	DW7590	NEO PMB
SIB F2S 2/6	REV 05	710-022603	DW7785	F2S SIB
B Board	REV 05	710-023787	DW7524	NEO PMB
SIB F2S 3/0	REV 05	710-022603	DW7782	F2S SIB
B Board	REV 05	710-023787	DW7634	NEO PMB
SIB F2S 3/2	REV 05	710-022603	DW7793	F2S SIB
B Board	REV 05	710-023787	DW7548	NEO PMB
SIB F2S 3/4	REV 05	710-022603	DW7779	F2S SIB
B Board	REV 05	710-023787	DW7587	NEO PMB
SIB F2S 3/6	REV 05	710-022603	DW7930	F2S SIB
B Board	REV 05	710-023787	DW7505	NEO PMB
SIB F2S 4/0	REV 05	710-022603	DW7867	F2S SIB
B Board	REV 05	710-023787	DW7656	NEO PMB
SIB F2S 4/2	REV 05	710-022603	DW7917	F2S SIB
B Board	REV 05	710-023787	DW7640	NEO PMB
SIB F2S 4/4	REV 05	710-022603	DW7929	F2S SIB
B Board	REV 05	710-023787	DW7643	NEO PMB
SIB F2S 4/6	REV 05	710-022603	DW7870	F2S SIB
B Board	REV 05	710-023787	DW7635	NEO PMB
Fan Tray 0	REV 06	760-024497	DV7831	Front Fan Tray
Fan Tray 1	REV 06	760-024497	DV9614	Front Fan Tray
Fan Tray 2	REV 06	760-024502	DV9618	Rear Fan Tray
Fan Tray 3	REV 06	760-024502	DV9616	Rear Fan Tray
Fan Tray 4	REV 06	760-024502	DV7807	Rear Fan Tray
Fan Tray 5	REV 06	760-024502	DV7828	Rear Fan Tray

show chassis hardware extensive (TX Matrix Plus Router)

```

user@host> show chassis hardware extensive
sfc0-re0:
-----
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis
Jedec Code:   0x7fb0          EEPROM Version: 0x02
S/N:          JN112F007AHB
Assembly ID:  0x052c          Assembly Version: 00.00
Date:         00-00-0000      Assembly Flags:  0x00
ID: TXP
Board Information Record:
Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
Address 0x00: 7f b0 02 ff 05 2c 00 00 00 00 00 00 00 00 00 00
Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x20: 4a 4e 31 31 32 46 30 30 37 41 48 42 00 00 00 00
Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Midplane      REV 05      710-022574      TS4027      SFC Midplane
Jedec Code:   0x7fb0          EEPROM Version: 0x01
P/N:          710-022574      S/N:          TS4027
Assembly ID:  0x0962          Assembly Version: 01.05
Date:         03-23-2009      Assembly Flags: 0x00
Version:      REV 05
ID: SFC Midplane
Board Information Record:
Address 0x00: ad 01 ff ff 00 1d b5 14 00 00 ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 09 62 01 05 52 45 56 20 30 35 00 00
Address 0x10: 00 00 00 00 37 31 30 2d 30 32 32 35 37 34 00 00
Address 0x20: 53 2f 4e 20 54 53 34 30 32 37 00 00 00 17 03 07
Address 0x30: d9 ff ff ff ad 01 ff ff 00 1d b5 14 00 00 ff ff
Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
FPM Display   REV 03      710-024027      DX0282      TXP FPM Display
Jedec Code:   0x7fb0          EEPROM Version: 0x01
P/N:          710-024027      S/N:          DX0282
Assembly ID:  0x096c          Assembly Version: 01.03
Date:         02-10-2009      Assembly Flags: 0x00
Version:      REV 03
ID: TXP FPM Display      FRU Model Number: CRAFT-TXP
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 09 6c 01 03 52 45 56 20 30 33 00 00
Address 0x10: 00 00 00 00 37 31 30 2d 30 32 34 30 32 37 00 00
Address 0x20: 53 2f 4e 20 44 58 30 32 38 32 00 00 00 0a 02 07
Address 0x30: d9 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 43
Address 0x50: 52 41 46 54 2d 54 58 50 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 ff ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
CIP 0         REV 04      710-023792      DW4889      TXP CIP

```

```

Jedec Code: 0x7fb0      EEPROM Version: 0x01
P/N: 710-023792        S/N: DW4889
Assembly ID: 0x0969     Assembly Version: 01.04
Date: 01-26-2009       Assembly Flags: 0x00
Version: REV 04
ID: TXP CIP            FRU Model Number: CIP-TXP
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

```

show chassis hardware clei-models (TX Matrix Plus Router)

```

user@host> show chassis hardware clei-models
sfc0-re0:
-----
Hardware inventory:

```

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 05	710-022574		CHAS-BP-TXP-S
FPM Display	REV 03	710-024027		CRAFT-TXP-S
CIP 0	REV 05	710-023792		CIP-TXP-S
CIP 1	REV 05	710-023792		CIP-TXP-S
PEM 0	Rev 04	740-027463	IPUPAFGKTA	PWR-TXP-7-60-DC
PEM 1	Rev 04	740-027463	IPUPAFGKTA	PWR-TXP-7-60-DC
Routing Engine 0	REV 06	740-026942		RE-DUO-C2600-16G-S
Routing Engine 1	REV 06	740-026942		RE-DUO-C2600-16G-S
CB 0	REV 05	710-022606		CB-TXP-S
CB 1	REV 09	710-022606		CB-TXP-S
SIB F13 0	REV 04	750-024564		SIB-TXP-F13
SIB F13 3	REV 04	750-024564		SIB-TXP-F13
SIB F13 8	REV 04	750-024564		SIB-TXP-F13
SIB F13 11	REV 04	750-024564		SIB-TXP-F13
SIB F13 12	REV 03	750-024564		SIB-TXP-F13
SIB F2S 0/0	REV 05	710-022603		SIB-TXP-F2S-S
SIB F2S 0/2	REV 05	710-022603		SIB-TXP-F2S-S
SIB F2S 0/4	REV 04	710-022603		SIB-TXP-F2S-S
SIB F2S 0/6	REV 04	710-022603		SIB-TXP-F2S-S
SIB F2S 1/0	REV 04	710-022603		SIB-TXP-F2S-S
SIB F2S 1/2	REV 05	710-022603		SIB-TXP-F2S-S
SIB F2S 1/4	REV 04	710-022603		SIB-TXP-F2S-S
SIB F2S 1/6	REV 05	710-022603		SIB-TXP-F2S-S
SIB F2S 2/0	REV 04	710-022603		SIB-TXP-F2S-S
SIB F2S 2/2	REV 04	710-022603		SIB-TXP-F2S-S
SIB F2S 2/4	REV 05	710-022603		SIB-TXP-F2S-S
SIB F2S 2/6	REV 04	710-022603		SIB-TXP-F2S-S
SIB F2S 3/0	REV 05	710-022603		SIB-TXP-F2S-S
SIB F2S 3/2	REV 03	710-022603		SIB-TXP-F2S-S
SIB F2S 3/4	REV 05	710-022603		SIB-TXP-F2S-S
SIB F2S 3/6	REV 03	710-022603		SIB-TXP-F2S-S
SIB F2S 4/0	REV 03	710-022603		SIB-TXP-F2S-S
SIB F2S 4/2	REV 05	710-022603		SIB-TXP-F2S-S
SIB F2S 4/4	REV 04	710-022603		SIB-TXP-F2S-S
SIB F2S 4/6	REV 03	710-022603		SIB-TXP-F2S-S
Fan Tray 0	REV 02	760-024497		FANTRAY-TXP-H-S
Fan Tray 1	REV 02	760-024497		FANTRAY-TXP-H-S
Fan Tray 2	REV 05	760-024502		FANTRAY-TXP-V-S
Fan Tray 3				
Fan Tray 4	REV 05	760-024502		FANTRAY-TXP-V-S
Fan Tray 5	REV 02	760-024502		FANTRAY-TXP-V-S

```

lcc0-re0:
-----

```

```

Hardware inventory:

```

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 03	710-017247	IPUPAC8KTA	CHAS-BP-T1600-S
FPM Display	REV 01	710-021387		CRAFT-T1600-S
CIP	REV 06	710-002895		CIP-L-T640-S
PEM 1	Rev 02	740-023211		PWR-T1600-4-60-DC-S
SCG 0	REV 15	710-003423		SCG-T-S
SCG 1	REV 15	710-003423		SCG-T-S
Routing Engine 0	REV 01	740-026941		RE-DUO-C1800-8G-S
Routing Engine 1	REV 01	740-026941		RE-DUO-C1800-8G-S
CB 0	REV 06	710-022597		CB-LCC-S
CB 1	REV 06	710-022597		CB-LCC-S
FPC 1	REV 07	710-013035		T640-FPC3-ES
PIC 0	REV 05	750-015217		PC-8GE-TYPE3-SFP-IQ2
PIC 1	REV 03	750-004424		PC-1XGE-LR
PIC 2	REV 01	750-003336		PC-40C48-SON-SMSR
FPC 3	REV 12	710-013037		T1600-FPC4-ES
PIC 0	REV 02	750-010850		PD-10C768-SON-SR
FPC 4	REV 05	710-021534		T640-FPC1-ES
PIC 0	REV 04	750-014627		PB-40C3-10C12-SON-SFP
PIC 1	REV 22	750-005634		PB-1CHOC12SMIR-QPP
PIC 2	REV 09	750-002911		PB-4FE-TX
PIC 3	REV 08	750-021652		PB-1CHOC12-STM4-IQE-SFP
FPC 5	REV 07	710-007529		T640-FPC3
PIC 0	REV 14	750-009567		PC-1XGE-XENPAK
PIC 1	REV 16	750-007141		PC-10GE-SFP
PIC 2	REV 12	750-009567		PC-1XGE-XENPAK
FPC 6	REV 07	710-013035		T640-FPC3-ES
PIC 0	REV 09	750-009567		PC-1XGE-XENPAK
PIC 1	REV 06	750-015217		PC-8GE-TYPE3-SFP-IQ2
PIC 2	REV 06	750-015217		PC-8GE-TYPE3-SFP-IQ2
FPC 7	REV 03	710-021540		T640-FPC2-ES
PIC 0	REV 13	750-001901		PB-40C12-SON-SMIR
PIC 1	REV 05	750-001900		PB-10C48-SON-SMSR
PIC 2	REV 10	750-008155		PB-2GE-SFP-QPP
PIC 3	REV 03	750-014638		PB-10C48-SON-B-SFP
SIB 0	REV 07	710-022594		SIB-TXP-T1600-S
SIB 1	REV 07	710-022594		SIB-TXP-T1600-S
SIB 3	REV 06	710-022594		SIB-TXP-T1600-S
SIB 4	REV 08	710-022594		SIB-TXP-T1600-S
Fan Tray 0				FANTRAY-T-S
Fan Tray 1				FANTRAY-T-S
Fan Tray 2				FANTRAY-TXP-R-S

lcc1-re0:

Hardware inventory:

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 04	710-017247	IPUPAC8KTA	CHAS-BP-T1600-S
FPM Display	REV 01	710-021387		CRAFT-T1600-S
CIP	REV 06	710-002895		CIP-L-T640-S
PEM 0	Rev 02	740-023211		PWR-T1600-4-60-DC-S
SCG 0	REV 15	710-003423		SCG-T-S
SCG 1	REV 15	710-003423		SCG-T-S
Routing Engine 0	REV 01	740-026941		RE-DUO-C1800-8G-S
Routing Engine 1	REV 01	740-026941		RE-DUO-C1800-8G-S
CB 0	REV 06	710-022597		CB-LCC-S
CB 1	REV 06	710-022597		CB-LCC-S
FPC 0	REV 02	710-010845		T640-FPC4-ES
PIC 0	REV 11	750-017405		PD-4XGE-XFP
FPC 1	REV 16	710-013037		T1600-FPC4-ES
PIC 1	REV 06	750-034781		PD-1CE-CFP

FPC 2	REV 16	710-013037	T1600-FPC4-ES
PIC 1	REV 05	750-034781	PD-1CE-CFP
FPC 3	REV 10	710-021534	T640-FPC1-ES
PIC 0	REV 13	750-012266	PB-4GE-TYPE1-SFP-IQ2
PIC 1	REV 01	750-007641	PE-1GE-SFP-QPP
PIC 3	REV 17	750-007444	PB-1CHSTM1-SMIR-QPP
FPC 4	REV 06	710-013035	T640-FPC3-ES
PIC 0	REV 22	750-007141	PC-10GE-SFP
PIC 1	REV 16	750-009450	PC-10C192-SON-SR2
PIC 2	REV 05	750-004424	PC-1XGE-LR
PIC 3	REV 12	750-013423	PC-MS-500-3
FPC 5	REV 07	710-013560	T640-FPC3-E2
PIC 0	REV 11	750-012793	PC-1XGE-TYPE3-XFP-IQ2
PIC 1	REV 01	750-004695	PC-TUNNEL
PIC 2	REV 32	750-003700	PC-10C192-SON-VSR
PIC 3	REV 12	750-009553	PC-40C48-SON-SFP
FPC 6	REV 07	710-013035	T640-FPC3-ES
PIC 0	REV 07	750-015217	PC-8GE-TYPE3-SFP-IQ2
PIC 1	REV 03	750-003336	PC-40C48-SON-SMSR
PIC 3	REV 02	750-012793	PC-1XGE-TYPE3-XFP-IQ2
FPC 7	REV 08	710-010845	T640-FPC4-ES
PIC 0	REV 11	750-017405	PD-4XGE-XFP
SIB 0	REV 07	710-022594	SIB-TXP-T1600-S
SIB 1	REV 07	710-022594	SIB-TXP-T1600-S
SIB 3	REV 07	710-022594	SIB-TXP-T1600-S
SIB 4	REV 08	710-022594	SIB-TXP-T1600-S
Fan Tray 0			FANTRAY-T-S
Fan Tray 1			FANTRAY-T-S
Fan Tray 2			FANTRAY-TXP-R-S

show chassis hardware detail (TX Matrix Plus Router)

```
user@host> show chassis hardware detail
sfc0-re0:
```

```
-----
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN111B023AHB	TXP
Midplane	REV 01	710-022574	TR7990	SFC Midplane
FPM Display	REV 03	710-024027	DW4699	TXP FPM Display
CIP 0	REV 01	710-023792	DR1437	TXP CIP
CIP 1	REV 02	710-023792	DS4564	TXP CIP
PEM 0	Rev 07	740-027463	UM26360	Power Entry Module
Routing Engine 0	REV 01	740-026942	737A-1024	SFC RE
ad0	3887 MB	SMART CF	200811050193CEB1CEB1	Compact Flash
ad1	30533 MB	SAMSUNG MCBQE32G8MPP-0V	SY814A0762	Disk 1
Routing Engine 1	REV 01	740-026942	737A-1024	SFC RE
ad0	3887 MB	SMART CF	20081105004C19A019A0	Compact Flash
ad1	30533 MB	SAMSUNG MCBQE32G8MPP-0V	SY814A0794	Disk 1
CB 0	REV 03	710-022606	DR7134	SFC Control Board
CB 1	REV 01	710-022606	DP8890	SFC Control Board
SPMB 0		BUILTIN		SFC Switch CPU
SPMB 1		BUILTIN		SFC Switch CPU
SIB F13 0	REV 03	750-024564	DT9478	F13 SIB
B Board	REV 02	710-023431	DT6554	F13 SIB
SIB F13 1	REV 03	750-024564	DT9454	F13 SIB
B Board	REV 02	710-023431	DT6551	F13 SIB
SIB F2S 0/0	REV 02	710-022603	DT2838	F2S SIB
B Board	REV 02	710-023787	DT1725	NEO PMB
SIB F2S 0/2	REV 02	710-022603	DT2824	F2S SIB
B Board	REV 02	710-023787	DT1706	NEO PMB

SIB F2S 0/4	REV 02	710-022603	DT2822	F2S SIB
B Board	REV 02	710-023787	DT1696	NEO PMB
SIB F2S 0/6	REV 02	710-022603	DT2823	F2S SIB
B Board	REV 02	710-023787	DT1717	NEO PMB
SIB F2S 1/0	REV 03	710-022603	DV0059	F2S SIB
B Board	REV 03	710-023787	DT9942	NEO PMB
SIB F2S 1/2	REV 02	710-022603	DT2826	F2S SIB
B Board	REV 02	710-023787	DT1713	NEO PMB
SIB F2S 1/4	REV 03	710-022603	DV0092	F2S SIB
B Board	REV 03	710-023787	DV0000	NEO PMB
SIB F2S 1/6	REV 03	710-022603	DV0079	F2S SIB
B Board	REV 03	710-023787	DT9972	NEO PMB
SIB F2S 2/0	REV 03	710-022603	DV0100	F2S SIB
B Board	REV 03	710-023787	DT9925	NEO PMB
SIB F2S 2/2	REV 03	710-022603	DV0050	F2S SIB
B Board	REV 03	710-023787	DV0005	NEO PMB
SIB F2S 2/4	REV 03	710-022603	DV0097	F2S SIB
B Board	REV 03	710-023787	DT9936	NEO PMB
Fan Tray 0	REV 02	760-024497	DR8286	Front Fan Tray
Fan Tray 1	REV 06	760-024497	DV9624	Front Fan Tray
Fan Tray 2	REV 02	760-024502	DR8259	Rear Fan Tray
Fan Tray 3	REV 02	760-024502	DR8270	Rear Fan Tray
Fan Tray 4	REV 02	760-024502	DR8284	Rear Fan Tray
Fan Tray 5	REV 06	760-024502	DV7813	Rear Fan Tray

1cc0-re0:

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN1101F27AHA	T1600
Midplane	REV 04	710-017247	RC5317	T Series Backplane
FPM GBUS	REV 10	710-002901	DS8197	T640 FPM Board
FPM Display	REV 01	710-021387	DS6433	T1600 FPM Display
CIP	REV 06	710-002895	DS1493	T Series CIP
PEM 0	Rev 08	740-017906	UD26601	Power Entry Module 3x80
SCG 0	REV 15	710-003423	DP5847	T640 Sonet Clock Gen.
SCG 1	REV 15	710-003423	DR0924	T640 Sonet Clock Gen.
Routing Engine 0	REV 01	740-026942	737F-1024	LCC RE
ad0	3887 MB	SMART CF	2008110502B63E513E51	Compact Flash
ad1	30533 MB	SAMSUNG MCBQE32G8MPP-0V	SY814A1208	Disk 1
Routing Engine 1	REV 01	740-026942	737F-1024	LCC RE
ad0	3887 MB	SMART CF	2008110500F9A8A8A8A8	Compact Flash
ad1	30533 MB	SAMSUNG MCBQE32G8MPP-0V	SY814A1076	Disk 1
CB 0	REV 05	710-022597	DV4264	LCC Control Board
CB 1	REV 03	710-022597	DP8558	LCC Control Board
FPC 0	REV 14	710-013037	DS9967	FPC Type 4-ES
CPU	REV 08	710-016744	DS3989	ST-PMB2
PIC 0	REV 12	750-013198	DL7506	1x Tunnel
PIC 1	REV 12	750-013198	DL7505	1x Tunnel
MMB 0	REV 01	710-025563	DS8524	ST-MMB2
MMB 1	REV 01	710-025563	DS8373	ST-MMB2
FPC 1	REV 14	710-013037	DT0027	FPC Type 4-ES
CPU	REV 09	710-016744	DS7684	ST-PMB2
PIC 0	REV 12	750-013198	DL7512	1x Tunnel
PIC 1	REV 12	750-013198	DL7498	1x Tunnel
MMB 0	REV 01	710-025563	DS8494	ST-MMB2
MMB 1	REV 01	710-025563	DS8436	ST-MMB2
SPMB 0	REV 04	710-023321	DV3867	LCC Switch CPU
SPMB 1	REV 02	710-023321	DP0238	LCC Switch CPU
SIB 0	REV 06	710-022594	DT8268	LCC SIB
B Board	REV 06	710-023185	DT5791	LCC SIB Mezz

SIB 1	REV 06	710-022594	DT8261	LCC SIB
B Board	REV 06	710-023185	DT5769	LCC SIB Mezz
SIB 2	REV 04	710-022594	DS2315	LCC SIB
B Board	REV 06	710-023185	DT5788	LCC SIB Mezz
SIB 3	REV 06	710-022594	DT8253	LCC SIB
B Board	REV 06	710-023185	DT5811	LCC SIB Mezz
SIB 4	REV 06	710-022594	DT8248	LCC SIB
B Board	REV 06	710-023185	DT5812	LCC SIB Mezz
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Fan Tray

show chassis hardware models (TX Matrix Plus Router)

```
user@host> show chassis hardware models
sfc0-re0:
```

```
-----
Hardware inventory:
```

Item	Version	Part number	Serial number	FRU model number
FPM Display	REV 03	710-024027	DX0282	CRAFT-TXP
CIP 0	REV 04	710-023792	DW4889	CIP-TXP
CIP 1	REV 04	710-023792	DW4887	CIP-TXP
PEM 0	Rev 07	740-027463	UM26368	yyyyyyyyyyyyyyyyyyyyyyyyyyyyyy
Routing Engine 0	REV 01	740-026942	737A-1064	RE-TXP-SFC-DU0-2600-16G
Routing Engine 1	REV 01	740-026942	737A-1082	RE-TXP-SFC-DU0-2600-16G
CB 0	REV 09	710-022606	DW6099	CB-TXP
CB 1	REV 09	710-022606	DW6096	CB-TXP
SIB F13 1	REV 04	750-024564	DW5776	SIB-TXP-F13
SIB F13 3	REV 04	750-024564	DW5762	SIB-TXP-F13
SIB F13 4	REV 04	750-024564	DW5797	SIB-TXP-F13
SIB F13 6	REV 04	750-024564	DW5770	SIB-TXP-F13
SIB F13 7	REV 04	750-024564	DW5758	SIB-TXP-F13
SIB F13 8	REV 04	750-024564	DW5761	SIB-TXP-F13
SIB F13 9	REV 04	750-024564	DW5754	SIB-TXP-F13
SIB F13 12	REV 04	750-024564	DW5794	SIB-TXP-F13
SIB F2S 0/0	REV 05	710-022603	DW7897	
SIB F2S 0/2	REV 05	710-022603	DW7833	
SIB F2S 0/4	REV 05	710-022603	DW7875	
SIB F2S 0/6	REV 05	710-022603	DW7860	
SIB F2S 1/0	REV 04	710-022603	DW4820	
SIB F2S 1/2	REV 05	710-022603	DW7849	
SIB F2S 1/4	REV 05	710-022603	DW7927	SIB-TXP-F2S
SIB F2S 1/6	REV 05	710-022603	DW7866	
SIB F2S 2/0	REV 05	710-022603	DW7880	
SIB F2S 2/2	REV 05	710-022603	DW7895	
SIB F2S 2/4	REV 05	710-022603	DW7907	
SIB F2S 2/6	REV 05	710-022603	DW7785	
SIB F2S 3/0	REV 05	710-022603	DW7782	
SIB F2S 3/2	REV 05	710-022603	DW7793	
SIB F2S 3/4	REV 05	710-022603	DW7779	
SIB F2S 3/6	REV 05	710-022603	DW7930	
SIB F2S 4/0	REV 05	710-022603	DW7867	
SIB F2S 4/2	REV 05	710-022603	DW7917	
SIB F2S 4/4	REV 05	710-022603	DW7929	
SIB F2S 4/6	REV 05	710-022603	DW7870	
Fan Tray 0	REV 06	760-024497	DV7831	FANTRAY-TXP-F
Fan Tray 1	REV 06	760-024497	DV9614	FANTRAY-TXP-F
Fan Tray 2	REV 06	760-024502	DV9618	FANTRAY-TXP-R
Fan Tray 3	REV 06	760-024502	DV9616	FANTRAY-TXP-R
Fan Tray 4	REV 06	760-024502	DV7807	FANTRAY-TXP-R
Fan Tray 5	REV 06	760-024502	DV7828	FANTRAY-TXP-R

lcc0-re0:

Hardware inventory:

Item	Version	Part number	Serial number	FRU model number
Midplane	REV 03	710-017247	RC3765	CHAS-BP-T1600-S
FPM Display	REV 01	710-021387	DN5441	CRAFT-T1600-S
CIP	REV 06	710-002895	DP6021	CIP-L-T640-S
PEM 0	Rev 07	740-017906	UA26384	PWR-T1600-3-80-DC-S
PEM 1	Rev 07	740-017906	UA26296	PWR-T1600-3-80-DC-S
SCG 0	REV 15	710-003423	DR0875	SCG-T-S
CB 0	REV 06	710-022597	DW8534	CB-LCC
CB 1	REV 06	710-022597	DW8527	CB-LCC
FPC 4	REV 12	710-013037	DJ8717	T1600-FPC4-ES
PIC 0	REV 11	750-017405	DP8795	PD-4XGE-XFP
PIC 1	REV 11	750-017405	DP8794	PD-4XGE-XFP
FPC 6	REV 14	710-013037	DS5335	T1600-FPC4-ES
PIC 0	REV 13	750-017405	DS7634	PD-4XGE-XFP
PIC 1	REV 13	750-017405	DS7637	PD-4XGE-XFP
FPC 7	REV 07	710-013035	DM0990	T1600-FPC3-ES
PIC 0	REV 16	750-007141	JJ8067	PC-10GE-SFP
PIC 1	REV 08	750-015749	WE9598	PC-10C192-SON-XFP
PIC 2	REV 10	750-009450	HX6466	PC-10C192-SON-SR2
SIB 0	REV 08	710-022594	DW8033	SIB-TXP-T1600-S
SIB 1	REV 08	710-022594	DW8044	SIB-TXP-T1600-S
SIB 2	REV 08	710-022594	DW8020	SIB-TXP-T1600-S
SIB 3	REV 08	710-022594	DW8063	SIB-TXP-T1600-S
SIB 4	REV 08	710-022594	DW8064	SIB-TXP-T1600-S
Fan Tray 0				FANTRAY-T-S
Fan Tray 1				FANTRAY-T-S
Fan Tray 2				FANTRAY-TXP-R-S

lcc1-re0:

Hardware inventory:

Item	Version	Part number	Serial number	FRU model number
Midplane	REV 04	710-017247	RC5361	CHAS-BP-T1600-S
FPM Display	REV 01	710-021387	DS6430	CRAFT-T1600-S
CIP	REV 06	710-002895	DS4239	CIP-L-T640-S
PEM 0	Rev 08	740-017906	UD26649	PWR-T1600-3-80-DC-S
SCG 0	REV 15	710-003423	DP5820	SCG-T-S
CB 0	REV 06	710-022597	DW8523	CB-LCC
CB 1	REV 06	710-022597	DW8528	CB-LCC
FPC 4	REV 12	710-013037	DP8509	T1600-FPC4-ES
PIC 0	REV 11	750-017405	DP8808	PD-4XGE-XFP
PIC 1	REV 11	750-017405	DP7263	PD-4XGE-XFP
FPC 6	REV 14	710-013037	DS9961	T1600-FPC4-ES
PIC 0	REV 13	750-017405	DS5532	PD-4XGE-XFP
PIC 1	REV 13	750-017405	DS7639	PD-4XGE-XFP
FPC 7	REV 03	710-013035	DF5564	T1600-FPC3-ES
PIC 0	REV 16	750-007141	JJ8063	PC-10GE-SFP
SIB 0	REV 08	710-022594	DW8035	SIB-TXP-T1600-S
SIB 1	REV 10	710-022594	DX7672	SIB-TXP-T1600-S
SIB 2	REV 08	710-022594	DW8060	SIB-TXP-T1600-S
SIB 3	REV 08	710-022594	DW8072	SIB-TXP-T1600-S
SIB 4	REV 08	710-022594	DW8043	SIB-TXP-T1600-S
Fan Tray 0				FANTRAY-T-S
Fan Tray 1				FANTRAY-T-S
Fan Tray 2				FANTRAY-TXP-R-S

lcc2-re0:

Hardware inventory:

Item	Version	Part number	Serial number	FRU model number
Midplane	REV 03	710-017247	RC3956	CHAS-BP-T1600-S
FPM Display	REV 01	710-021387	DN7030	CRAFT-T1600-S
CIP	REV 06	710-002895	DM3962	CIP-L-T640-S
PEM 0	Rev 08	740-017906	UD26519	PWR-T1600-3-80-DC-S
PEM 1	Rev 07	740-017906	UC26601	PWR-T1600-3-80-DC-S
SCG 0	REV 15	710-003423	DP0277	SCG-T-S
CB 0	REV 06	710-022597	DW8524	CB-LCC
CB 1	REV 06	710-022597	DW8536	CB-LCC
FPC 4	REV 12	710-013037	DR1194	T1600-FPC4-ES
PIC 0	REV 11	750-017405	DP8811	PD-4XGE-XFP
PIC 1	REV 11	750-017405	DP8823	PD-4XGE-XFP
FPC 5	REV 12	710-013037	DR1184	T1600-FPC4-ES
PIC 1	REV 11	750-017405	DP4744	PD-4XGE-XFP
FPC 6	REV 12	710-013037	DN8622	T1600-FPC4-ES
PIC 0	REV 14	750-012518	JY9924	PD-40C192-SON-XFP
PIC 1	REV 11	750-017405	DP8776	PD-4XGE-XFP
FPC 7	REV 04	710-013560	JR3968	T640-FPC3-E2
PIC 0	REV 16	750-007141	NC9330	PC-10GE-SFP
SIB 0	REV 07	710-022594	DW4217	SIB-TXP-T1600-S
SIB 1	REV 07	710-022594	DW4213	SIB-TXP-T1600-S
SIB 2	REV 07	710-022594	DW4189	SIB-TXP-T1600-S
SIB 3	REV 07	710-022594	DW4173	SIB-TXP-T1600-S
SIB 4	REV 07	710-022594	DW4201	SIB-TXP-T1600-S
Fan Tray 0				FANTRAY-T-S
Fan Tray 1				FANTRAY-T-S
Fan Tray 2				FANTRAY-TXP-R-S

lcc3-re0:

Hardware inventory:

Item	Version	Part number	Serial number	FRU model number
Midplane	REV 04	710-017247	RC5319	CHAS-BP-T1600-S
FPM Display	REV 01	710-021387	DS6402	CRAFT-T1600-S
CIP	REV 06	710-002895	DR9973	CIP-L-T640-S
PEM 0	Rev 07	740-017906	UC26496	PWR-T1600-3-80-DC-S
PEM 1	Rev 07	740-017906	UC26599	PWR-T1600-3-80-DC-S
SCG 0	REV 15	710-003423	DP5831	SCG-T-S
CB 0	REV 06	710-022597	DW8533	CB-LCC
CB 1	REV 06	710-022597	DW8538	CB-LCC
FPC 0	REV 14	710-013037	DS5345	T1600-FPC4-ES
PIC 0	REV 13	750-017405	DS7641	PD-4XGE-XFP
PIC 1	REV 13	750-017405	DS5479	PD-4XGE-XFP
FPC 1	REV 14	710-013037	DS7338	T1600-FPC4-ES
PIC 0	REV 13	750-017405	DS7631	PD-4XGE-XFP
PIC 1	REV 13	750-017405	DS7632	PD-4XGE-XFP
FPC 2	REV 14	710-013037	DS9962	T1600-FPC4-ES
PIC 0	REV 13	750-017405	DS7581	PD-4XGE-XFP
PIC 1	REV 13	750-017405	DS7627	PD-4XGE-XFP
FPC 4	REV 10	710-010845	JZ6573	T640-FPC4-ES
PIC 0	REV 14	750-012518	JT5124	PD-40C192-SON-XFP
FPC 5	REV 14	710-013037	DT0016	T1600-FPC4-ES
PIC 0	REV 14	750-012518	JY9918	PD-40C192-SON-XFP
FPC 7	REV 07	710-013035	DM0967	T1600-FPC3-ES
PIC 0	REV 16	750-007141	JJ8059	PC-10GE-SFP
PIC 1	REV 13	750-004695	DM5712	PC-TUNNEL
SIB 0	REV 07	710-022594	DW4174	SIB-TXP-T1600-S
SIB 1	REV 07	710-022594	DW4207	SIB-TXP-T1600-S
SIB 2	REV 06	710-022594	DT8231	SIB-TXP-T1600-S

SIB 3	REV 07	710-022594	DW4175	SIB-TXP-T1600-S
SIB 4	REV 07	710-022594	DW4209	SIB-TXP-T1600-S
Fan Tray 0				FANTRAY-T-S
Fan Tray 1				FANTRAY-T-S
Fan Tray 2				FANTRAY-TXP-R-S

show chassis hardware (TX Matrix Plus Router with 3D SIBs)

```
user@host> show chassis hardware
sfc0-re0:
```

```
-----
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN11CAAA4AHB	TXP
Midplane	REV 05	710-022574	ABAC4696	SFC Midplane
FPM Display	REV 09	710-024027	EH3138	TXP FPM Display
CIP 0	REV 12	710-023792	EF6349	TXP CIP
CIP 1	REV 12	710-023792	EG5294	TXP CIP
PEM 0	Rev 06	740-027463	XH04595	Power Entry Module
PEM 1	Rev 06	740-027463	XH04592	Power Entry Module
Routing Engine 0	REV 07	740-026942	P737A-002541	RE-DUO-2600
Routing Engine 1	REV 07	740-026942	P737A-002602	RE-DUO-2600
CB 0	REV 15	710-022606	EH4376	SFC Control Board
CB 1	REV 15	710-022606	EH4379	SFC Control Board
SPMB 0		BUILTIN		SFC Switch CPU
SPMB 1		BUILTIN		SFC Switch CPU
SIB F13 0	REV 10	750-035002	EM9305	F13 SIB 3D
B Board	REV 06	711-035082	EM9667	F13 SIB 3D Mezz
P Board	REV 05	711-043544	EM9708	F13 SIB 3D Power
Xcvr 0	REV 01	740-047547	XB34FB00S	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB01H	CXP Module
Xcvr 4	REV 01	740-047547	XB34FB02W	CXP Module
Xcvr 6	REV 01	740-047547	XB34FB01T	CXP Module
Xcvr 8	REV 01	740-047547	XB48FB00W	CXP Module
Xcvr 10	REV 01	740-047547	XB34FB01S	CXP Module
Xcvr 12	REV 01	740-047547	XB34FB03H	CXP Module
Xcvr 14	REV 01	740-047547	XB34FB023	CXP Module
SIB F13 3	REV 01	710-035001	EJ2612	F13 SIB 3D
B Board	REV 01	711-035082	EJ3815	F13 SIB 3D Mezz
P Board	REV 01	711-043544	EJ2678	F13 SIB 3D Power
Xcvr 0	REV 01	740-047547	XB48FB04C	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB00Z	CXP Module
Xcvr 4	REV 01	740-047547	XB47FB036	CXP Module
Xcvr 6	REV 01	740-047547	XB47FB029	CXP Module
Xcvr 8	REV 01	740-047547	XB48FB02N	CXP Module
Xcvr 10	REV 01	740-047547	XB42FB0CS	CXP Module
Xcvr 12	REV 01	740-047547	XB47FB01X	CXP Module
Xcvr 14	REV 01	740-047547	XB48FB02F	CXP Module
SIB F13 6	REV 05	750-035002	EK2675	F13 SIB 3D
B Board	REV 03	711-035082	EK2612	F13 SIB 3D Mezz
P Board	REV 04	711-043544	EK1179	F13 SIB 3D Power
Xcvr 0	REV 01	740-047547	XB48FB01T	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB02M	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB031	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB04P	CXP Module
Xcvr 8	REV 01	740-047547	XB48FB02T	CXP Module
Xcvr 10	REV 01	740-047547	XB34FB01V	CXP Module
Xcvr 12	REV 01	740-047547	XB48FB02C	CXP Module
Xcvr 14		NON-JNPR		No Module
SIB F13 12	REV 01	710-035001	EJ2631	F13 SIB 3D
B Board	REV 01	711-035082	EJ3808	F13 SIB 3D Mezz

P Board	REV 01	711-043544	EJ2676	F13 SIB 3D Power
SIB F2S 0/0	REV 01	711-034977	EH9829	F2S SIB 3D
B Board	REV 01	711-034979	EH9927	F2S SIB 3D Mezz
SIB F2S 0/2	REV 01	711-034977	EH9791	F2S SIB 3D
B Board	REV 01	711-034979	EH9852	F2S SIB 3D Mezz
SIB F2S 0/4	REV 01	711-034977	EH9803	F2S SIB 3D
B Board	REV 01	711-034979	EH9915	F2S SIB 3D Mezz
SIB F2S 0/6	REV 01	711-034977	EH9763	F2S SIB 3D
B Board	REV 01	711-034979	EH9880	F2S SIB 3D Mezz
SIB F2S 1/0	REV 01	711-034977	EH9757	F2S SIB 3D
B Board	REV 01	711-034979	EH9889	F2S SIB 3D Mezz
SIB F2S 1/2	REV 01	711-034977	EH9815	F2S SIB 3D
B Board	REV 01	711-034979	EH9890	F2S SIB 3D Mezz
SIB F2S 1/4	REV 08	750-034978	EN1954	F2S SIB 3D
B Board	REV 02	711-034979	EN1436	F2S SIB 3D Mezz
SIB F2S 1/6	REV 01	711-034977	EJ7054	F2S SIB 3D
B Board	REV 01	711-034979	EJ8238	F2S SIB 3D Mezz
SIB F2S 2/0	REV 01	711-034977	EH9830	F2S SIB 3D
B Board	REV 01	711-034979	EH9844	F2S SIB 3D Mezz
SIB F2S 2/2	REV 01	711-034977	EH9818	F2S SIB 3D
B Board	REV 01	711-034979	EH9888	F2S SIB 3D Mezz
SIB F2S 2/4	REV 01	711-034977	EH9795	F2S SIB 3D
B Board	REV 01	711-034979	EH9869	F2S SIB 3D Mezz
SIB F2S 2/6	REV 01	711-034977	EJ7026	F2S SIB 3D
B Board	REV 01	711-034979	EJ8273	F2S SIB 3D Mezz
SIB F2S 3/0	REV 01	711-034977	EH9811	F2S SIB 3D
B Board	REV 01	711-034979	EH9892	F2S SIB 3D Mezz
SIB F2S 3/2	REV 01	711-034977	EH9812	F2S SIB 3D
B Board	REV 01	711-034979	EH9877	F2S SIB 3D Mezz
SIB F2S 3/4	REV 08	750-034978	EN1947	F2S SIB 3D
B Board	REV 02	711-034979	EN1471	F2S SIB 3D Mezz
Fan Tray 0	REV 10	760-024497	EH3313	Front Fan Tray
Fan Tray 1	REV 10	760-024497	EH3290	Front Fan Tray
Fan Tray 2	REV 10	760-024502	EH3292	Rear Fan Tray
Fan Tray 3	REV 10	760-024502	EH3287	Rear Fan Tray
Fan Tray 4	REV 10	760-024502	EH3286	Rear Fan Tray
Fan Tray 5	REV 10	760-024502	EH3285	Rear Fan Tray

lcc0-re0:

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN11B23FEAHA	T1600
Midplane	REV 01	710-027486	RC9787	T-series Backplane
FPM GBUS	REV 13	710-002901	BBAG5132	T640 FPM Board
FPM Display	REV 04	710-021387	BBAL9612	T1600 FPM Display
CIP	REV 06	710-002895	BBAN0605	T-series CIP
PEM 0	REV 05	740-036442	1G022060143	Power Entry Module 6x60
PEM 1	REV 05	740-036442	1G022060011	Power Entry Module 6x60
SCG 0	REV 18	710-003423	BBAL7318	T640 Sonet Clock Gen.
SCG 1	REV 18	710-003423	BBAL7255	T640 Sonet Clock Gen.
Routing Engine 0	REV 07	740-026941	P737F-002933	RE-DUO-1800
Routing Engine 1	REV 06	740-026941	P737F-002749	RE-DUO-1800
CB 0	REV 11	710-022597	EH3611	LCC Control Board
CB 1	REV 11	710-022597	EH4798	LCC Control Board
FPC 5	REV 17	710-013037	BBAC5333	FPC Type 4-ES
CPU	REV 10	710-016744	BBAB7619	ST-PMB2
PIC 0	REV 18	750-017405	BBAE3420	4x 10GE (LAN/WAN) XFP
Xcvr 0	REV 03	740-014289	T10C90659	XFP-10G-SR
MMB 0	REV 05	710-025563	BBAB9538	ST-MMB2
MMB 1	REV 05	710-025563	BBAB9502	ST-MMB2

FPC 7	REV 01	750-045173	BBAV0032	FPC Type 5-3D
CPU				
SPMB 0	REV 05	710-023321	EG9434	LCC Switch CPU
SPMB 1	REV 05	710-023321	EH3878	LCC Switch CPU
SIB 0	REV 01	750-041657	EH7997	LCC SIB 3D
B Board	REV 01	711-042424	EH7674	LCC SIB 3D Mezz
Xcvr 0	REV 01	740-047547	XB48FB014	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB05A	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB052	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB01B	CXP Module
SIB 1	REV 01	750-041657	EH8023	LCC SIB 3D
B Board	REV 01	711-042424	EH7659	LCC SIB 3D Mezz
Xcvr 0	REV 01	740-047547	XB48FB05J	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB01E	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB01J	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB02S	CXP Module
SIB 2	REV 03	750-041657	EJ6554	LCC SIB 3D
B Board	REV 02	711-042424	EJ5756	LCC SIB 3D Mezz
Xcvr 0	REV 01	740-047547	XB34FB01Z	CXP Module
Xcvr 2	REV 01	740-047547	XB34FB013	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB04Z	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB05N	CXP Module
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Fan Tray -- Rev 4

lcc2-re0:

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN11B3975AHA	T1600
Midplane	REV 01	710-027486	RC9826	T-series Backplane
FPM GBUS	REV 13	710-002901	BBAG5124	T640 FPM Board
FPM Display	REV 03	710-021387	BBAJ1112	T1600 FPM Display
CIP	REV 06	710-002895	BBAL3744	T-series CIP
PEM 0	REV 05	740-036442	1G022060081	Power Entry Module 6x60
PEM 1	REV 05	740-036442	1G022060188	Power Entry Module 6x60
SCG 0	REV 18	710-003423	BBAH8775	T640 Sonet Clock Gen.
SCG 1	REV 18	710-003423	BBAL7272	T640 Sonet Clock Gen.
Routing Engine 0	REV 07	740-026941	P737F-002992	RE-DUO-1800
Routing Engine 1	REV 07	740-026941	P737F-002938	RE-DUO-1800
CB 0	REV 11	710-022597	EH4805	LCC Control Board
CB 1	REV 11	710-022597	EH4786	LCC Control Board
FPC 1	REV 01	710-033873	BBAH0320	FPC Type 3-ES
CPU	REV 11	710-016744	BBAF3281	ST-PMB2
MMB 0	REV 06	710-025563	BBAF5061	ST-MMB2
FPC 5	REV 04	710-033871	BBAM5070	FPC Type 4-ES
CPU	REV 11	710-016744	BBAM6653	ST-PMB2
PIC 1	REV 20	750-017405	BBAM1296	4x 10GE (LAN/WAN) XFP
Xcvr 0	REV 03	740-014289	T10B42981	XFP-10G-SR
MMB 0	REV 07	710-025563	BBAN2631	ST-MMB2
MMB 1	REV 07	710-025563	BBAN2538	ST-MMB2
SPMB 0	REV 05	710-023321	EH3903	LCC Switch CPU
SPMB 1	REV 05	710-023321	EH3902	LCC Switch CPU
SIB 0	REV 01	750-041657	EH8019	LCC SIB 3D
B Board	REV 01	711-042424	EH7680	LCC SIB 3D Mezz
Xcvr 0	REV 01	740-047547	XB48FB04F	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB04S	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB04B	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB043	CXP Module
SIB 1	REV 01	750-041657	EH8012	LCC SIB 3D

B Board	REV 01	711-042424	EH7658	LCC SIB 3D Mezz
Xcvr 0	REV 01	740-047547	XB48FB05E	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB01Z	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB018	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB054	CXP Module
SIB 2	REV 01	750-041657	EH7993	LCC SIB 3D
B Board	REV 01	711-042424	EH7678	LCC SIB 3D Mezz
Xcvr 0	REV 01	740-047547	XB48FB05C	CXP Module
Xcvr 2	REV 01	740-047547	XB47FB00N	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB05U	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB05L	CXP Module
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Fan Tray -- Rev 4

show chassis hardware clei-models (TX Matrix Plus Router with 3D SIBs)

```

user@host> show chassis hardware clei-models
sfc0-re0:
-----
Hardware inventory:
Item          Version  Part number  CLEI code  FRU model number
Midplane      REV 05   710-022574
FPM Display   REV 09   710-024027
CIP 0         REV 12   710-023792
CIP 1         REV 12   710-023792
PEM 0         Rev 06   740-027463   IPUPAFGKTA  PWR-TXP-7-60-DC-S
Routing Engine 0 REV 07   740-026942   RE-DUO-C2600-16G-S
Routing Engine 1 REV 07   740-026942   RE-DUO-C2600-16G-S
CB 0          REV 13   710-022606   CB-TXP-S
CB 1          REV 14   710-022606   CB-TXP-S
SIB F13 0     REV 10   750-035002   PROTOXCLEI  SIB-TXP-3D-F13-S
  Xcvr 0       REV 01   740-048813
  Xcvr 1       REV 01   740-048813
  Xcvr 2       REV 01   740-048813
  Xcvr 3       REV 01   740-048813
  Xcvr 4       REV 01   740-048813
  Xcvr 5       REV 01   740-048813
  Xcvr 6       REV 01   740-048813
  Xcvr 7       REV 01   740-048813
  Xcvr 8       REV 01   740-047547   CXP-TXP-3D
  Xcvr 10      REV 01   740-047547   CXP-TXP-3D
  Xcvr 12      REV 01   740-047547   CXP-TXP-3D
  Xcvr 14      REV 01   740-047547   CXP-TXP-3D
SIB F13 1     REV 10   750-035002   PROTOXCLEI  SIB-TXP-3D-F13-S
  Xcvr 0       REV 01   740-047547   CXP-TXP-3D
  Xcvr 1       REV 01   740-047547   CXP-TXP-3D
  Xcvr 2       REV 01   740-047547   CXP-TXP-3D
  Xcvr 3       REV 01   740-047547   CXP-TXP-3D
  Xcvr 4       REV 01   740-047547   CXP-TXP-3D
  Xcvr 5       REV 01   740-047547   CXP-TXP-3D
  Xcvr 6       REV 01   740-047547   CXP-TXP-3D
  Xcvr 7       REV 01   740-047547   CXP-TXP-3D
  Xcvr 8       REV 01   740-047547   CXP-TXP-3D
  Xcvr 10      REV 01   740-047547   CXP-TXP-3D
  Xcvr 12      REV 01   740-047547   CXP-TXP-3D
  Xcvr 14      REV 01   740-047547   CXP-TXP-3D
  Xcvr 0       REV 01   740-048813
  Xcvr 1       REV 01   740-048813
  Xcvr 2       REV 01   740-048813
  Xcvr 3       REV 01   740-048813

```

Xcvr 4	REV 01	740-048813		
Xcvr 5	REV 01	740-048813		
Xcvr 6	REV 01	740-048813		
Xcvr 7	REV 01	740-048813		
Xcvr 8	REV 01	740-048813		
Xcvr 10	REV 01	740-048813		
Xcvr 12	REV 01	740-048813		
Xcvr 14	REV 01	740-048813		
Xcvr 0	REV 01	740-047547		CXP-TXP-3D
Xcvr 1	REV 01	740-047547		CXP-TXP-3D
Xcvr 2	REV 01	740-047547		CXP-TXP-3D
Xcvr 3	REV 01	740-047547		CXP-TXP-3D
Xcvr 4	REV 01	740-047547		CXP-TXP-3D
Xcvr 5	REV 01	740-047547		CXP-TXP-3D
Xcvr 6	REV 01	740-047547		CXP-TXP-3D
Xcvr 7	REV 01	740-047547		CXP-TXP-3D
Xcvr 8	REV 01	740-047547		CXP-TXP-3D
Xcvr 10	REV 01	740-047547		CXP-TXP-3D
Xcvr 12	REV 01	740-047547		CXP-TXP-3D
Xcvr 14	REV 01	740-047547		CXP-TXP-3D
SIB F13 6	REV 16	750-035002	PROTOXCLEI	SIB-TXP-3D-F13
Xcvr 0	REV 01	740-048813		
Xcvr 1	REV 01	740-048813		
Xcvr 2	REV 01	740-048813		
Xcvr 3	REV 01	740-048813		
Xcvr 4	REV 01	740-048813		
Xcvr 5	REV 01	740-048813		
Xcvr 6	REV 01	740-048813		
Xcvr 7	REV 01	740-048813		
Xcvr 8	REV 01	740-047547		CXP-TXP-3D
Xcvr 10	REV 01	740-047547		CXP-TXP-3D
Xcvr 12	REV 01	740-047547		CXP-TXP-3D
Xcvr 14	REV 01	740-047547		CXP-TXP-3D
SIB F13 7	REV 10	750-035002	PROTOXCLEI	SIB-TXP-3D-F13-S
Xcvr 0	REV 01	740-047547		CXP-TXP-3D
Xcvr 1	REV 01	740-047547		CXP-TXP-3D
Xcvr 2	REV 01	740-047547		CXP-TXP-3D
Xcvr 3	REV 01	740-047547		CXP-TXP-3D
Xcvr 4	REV 01	740-047547		CXP-TXP-3D
Xcvr 5	REV 01	740-047547		CXP-TXP-3D
Xcvr 6	REV 01	740-047547		CXP-TXP-3D
Xcvr 7	REV 01	740-047547		CXP-TXP-3D
Xcvr 8	REV 01	740-047547		CXP-TXP-3D
Xcvr 10	REV 01	740-047547		CXP-TXP-3D
Xcvr 12	REV 01	740-047547		CXP-TXP-3D
Xcvr 14	REV 01	740-047547		CXP-TXP-3D
Xcvr 0	REV 01	740-048813		
Xcvr 1	REV 01	740-048813		
Xcvr 2	REV 01	740-048813		
Xcvr 3	REV 01	740-048813		
Xcvr 4	REV 01	740-048813		
Xcvr 5	REV 01	740-047547		CXP-TXP-3D
Xcvr 6	REV 01	740-047547		CXP-TXP-3D
Xcvr 7	REV 01	740-047547		CXP-TXP-3D
Xcvr 8	REV 01	740-047547		CXP-TXP-3D
Xcvr 10	REV 01	740-047547		CXP-TXP-3D
Xcvr 12	REV 01	740-047547		CXP-TXP-3D
Xcvr 14	REV 01	740-047547		CXP-TXP-3D
SIB F13 9	REV 16	750-035002	PROTOXCLEI	SIB-TXP-3D-F13
Xcvr 0	REV 01	740-047547		CXP-TXP-3D
Xcvr 1	REV 01	740-047547		CXP-TXP-3D

Xcvr 2	REV 01	740-047547		CXP-TXP-3D
Xcvr 3	REV 01	740-047547		CXP-TXP-3D
Xcvr 4	REV 01	740-047547		CXP-TXP-3D
Xcvr 5	REV 01	740-047547		CXP-TXP-3D
Xcvr 6	REV 01	740-047547		CXP-TXP-3D
Xcvr 7	REV 01	740-047547		CXP-TXP-3D
Xcvr 8	REV 01	740-047547		CXP-TXP-3D
Xcvr 10	REV 01	740-047547		CXP-TXP-3D
Xcvr 12	REV 01	740-047547		CXP-TXP-3D
Xcvr 14	REV 01	740-047547		CXP-TXP-3D
SIB F13 11	REV 10	750-035002	PROTOXCLEI	750-035002
Xcvr 0	REV 01	740-048813		
Xcvr 1	REV 01	740-048813		
Xcvr 2	REV 01	740-048813		
Xcvr 3	REV 01	740-048813		
Xcvr 4	REV 01	740-048813		
Xcvr 5	REV 01	740-048813		
Xcvr 6	REV 01	740-047547		CXP-TXP-3D
Xcvr 7	REV 01	740-048813		
Xcvr 8	REV 01	740-047547		CXP-TXP-3D
Xcvr 12	REV 01	740-047547		CXP-TXP-3D
Xcvr 14	REV 01	740-047547		CXP-TXP-3D
SIB F13 12	REV 16	750-035002	PROTOXCLEI	SIB-TXP-3D-F13
Xcvr 0	REV 01	740-047547		CXP-TXP-3D
Xcvr 1	REV 01	740-047547		CXP-TXP-3D
Xcvr 2	REV 01	740-047547		CXP-TXP-3D
Xcvr 3	REV 01	740-047547		CXP-TXP-3D
Xcvr 4	REV 01	740-047547		CXP-TXP-3D
Xcvr 5	REV 01	740-047547		CXP-TXP-3D
Xcvr 6	REV 01	740-047547		CXP-TXP-3D
Xcvr 7	REV 01	740-047547		CXP-TXP-3D
Xcvr 8	REV 01	740-047547		CXP-TXP-3D
Xcvr 10	REV 01	740-047547		CXP-TXP-3D
Xcvr 12	REV 01	740-047547		CXP-TXP-3D
Xcvr 14	REV 01	740-047547		CXP-TXP-3D
SIB F2S 0/0	REV 06	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 0/2	REV 07	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 0/4	REV 06	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 0/6	REV 06	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 1/0	REV 07	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 1/2	REV 07	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 1/4	REV 07	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 1/6	REV 08	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 2/0	REV 06	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 2/2	REV 06	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 2/4	REV 07	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 2/6	REV 06	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 3/0	REV 07	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 3/2	REV 06	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 3/4	REV 06	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 3/6	REV 06	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 4/0	REV 07	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 4/2	REV 06	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 4/4	REV 06	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 4/6	REV 06	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
Fan Tray 0	REV 10	760-024497		FANTRAY-TXP-H-S
Fan Tray 1	REV 10	760-024497		FANTRAY-TXP-H-S
Fan Tray 2	REV 10	760-024502		FANTRAY-TXP-V-S
Fan Tray 3	REV 10	760-024502		FANTRAY-TXP-V-S
Fan Tray 4	REV 10	760-024502		FANTRAY-TXP-V-S
Fan Tray 5	REV 10	760-024502		FANTRAY-TXP-V-S

1cc0-re0:

Hardware inventory:

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 01	710-027486	IPMJ700DRD	CHAS-BP-T1600-S
FPM Display	REV 04	710-021387		CRAFT-T1600-S
CIP	REV 06	710-002895		CIP-L-T640-S
PEM 0	REV 05	740-036442	IPUPAG6KAA	PWR-T-6-60-DC-S
PEM 1	REV 05	740-036442	IPUPAG6KAA	PWR-T-6-60-DC-S
SCG 0	REV 18	710-003423		SCG-T-S
SCG 1	REV 18	710-003423		SCG-T-S
Routing Engine 0	REV 10	740-026941		RE-DUO-C1800-8G-S
Routing Engine 1	REV 07	740-026941		RE-DUO-C1800-8G-S
CB 0	REV 11	710-022597		CB-LCC-S
CB 1	REV 11	710-022597		CB-LCC-S
FPC 0	REV 01	750-045173	IP9IAL4DAB	T4000-FPC5-3D
PIC 0	REV 17	750-034624	IP9IAL2DAA	PF-12XGE-SFPP
PIC 1	REV 17	750-034624	IP9IAL2DAA	PF-12XGE-SFPP
FPC 3	REV 01	750-045173	IP9IAL4DAB	T4000-FPC5-3D
PIC 0	REV 13	750-033423	XXXXXXXXDD	PF-12-24XGE-SFPP
FPC 4	REV 02	750-045173	IP9IAL4DAC	T4000-FPC5-3D
PIC 0	REV 17	750-034624	IP9IAL2DAA	PF-12XGE-SFPP
PIC 1	REV 17	750-034624	IP9IAL2DAA	PF-12XGE-SFPP
FPC 5	REV 01	750-045173	IP9IAL4DAB	T4000-FPC5-3D
PIC 0	REV 17	750-034624	IP9IAL2DAA	PF-12XGE-SFPP
PIC 1	REV 17	750-034624	IP9IAL2DAA	PF-12XGE-SFPP
FPC 6	REV 01	750-045173	IP9IAL4DAB	T4000-FPC5-3D
PIC 0	REV 17	750-034624	IP9IAL2DAA	PF-12XGE-SFPP
PIC 1	REV 10	750-035293	IP9IAL3DAA	PF-1CGE-CFP
SIB 0	REV 06	750-041657	PROTOXCLEI	SIB-TXP-3D-LCC
Xcvr 0	REV 01	740-048813		
Xcvr 1	REV 01	740-048813		
Xcvr 2	REV 01	740-048813		
Xcvr 3	REV 01	740-048813		
Xcvr 4	REV 01	740-048813		
Xcvr 5	REV 01	740-048813		
Xcvr 6	REV 01	740-048813		
Xcvr 7	REV 01	740-048813		
SIB 1	REV 06	750-041657	PROTOXCLEI	SIB-TXP-3D-LCC
Xcvr 0	REV 01	740-048813		
Xcvr 1	REV 01	740-048813		
Xcvr 2	REV 01	740-048813		
Xcvr 3	REV 01	740-048813		
Xcvr 4	REV 01	740-048813		
Xcvr 5	REV 01	740-048813		
Xcvr 6	REV 01	740-048813		
Xcvr 7	REV 01	740-048813		
SIB 2	REV 06	750-041657	PROTOXCLEI	SIB-TXP-3D-LCC
Xcvr 0	REV 01	740-048813		
Xcvr 1	REV 01	740-048813		
Xcvr 2	REV 01	740-048813		
Xcvr 3	REV 01	740-048813		
Xcvr 4	REV 01	740-048813		
Xcvr 5	REV 01	740-048813		
Xcvr 6	REV 01	740-048813		
Xcvr 7	REV 01	740-048813		
SIB 3	REV 07	750-041657	PROTOXCLEI	SIB-TXP-3D-LCC
Xcvr 0	REV 01	740-048813		
Xcvr 1	REV 01	740-048813		
Xcvr 2	REV 01	740-048813		

```

Xcvr 3      REV 01  740-048813
Xcvr 4      REV 01  740-048813
Xcvr 5      REV 01  740-048813
Xcvr 6      REV 01  740-048813
Xcvr 7      REV 01  740-048813
SIB 4       REV 06  750-041657  PROTOXCLEI  SIB-TXP-3D-LCC
Xcvr 0      REV 01  740-048813
Xcvr 1      REV 01  740-048813
Xcvr 2      REV 01  740-048813
Xcvr 3      REV 01  740-048813
Xcvr 4      REV 01  740-048813
Xcvr 5      REV 01  740-048813
Xcvr 6      REV 01  740-048813
Xcvr 7      REV 01  740-048813
Fan Tray 0
Fan Tray 1
Fan Tray 2
[Output Truncated]
FANTRAY-T-S
FANTRAY-T-S
FANTRAY-TXP3D-LCC-R-S

```

show chassis hardware detail (TX Matrix Plus Router with 3D SIBs)

```

user@host> show chassis hardware detail
sfc0-re0:
-----
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               JN11CAAA4AHB  TXP
Midplane      REV 05   710-022574  ABAC4696      SFC Midplane
FPM Display   REV 09   710-024027  EH3138      TXP FPM Display
CIP 0         REV 12   710-023792  EF6349      TXP CIP
CIP 1         REV 12   710-023792  EG5294      TXP CIP
PEM 0         Rev 06   740-027463  XH04595      Power Entry Module
PEM 1         Rev 06   740-027463  XH04592      Power Entry Module
Routing Engine 0 REV 07   740-026942  P737A-002541 RE-DUO-2600
  ad0 3823 MB SMART CF 2011030400062C132C13 Compact Flash
  ad1 62720 MB SMART Lite SATA Drive 201105100009A452A452 Disk 1
Routing Engine 1 REV 07   740-026942  P737A-002602 RE-DUO-2600
  ad0 3823 MB SMART CF 20110508085EE471E471 Compact Flash
  ad1 62720 MB SMART Lite SATA Drive 201110210089DF39DF39 Disk 1
CB 0          REV 15   710-022606  EH4376      SFC Control Board
CB 1          REV 15   710-022606  EH4379      SFC Control Board
SPMB 0        BUILTIN
SPMB 1        BUILTIN
SIB F13 0     REV 10   750-035002  EM9305      F13 SIB 3D
  B Board     REV 06   711-035082  EM9667      F13 SIB 3D Mezz
  P Board     REV 05   711-043544  EM9708      F13 SIB 3D Power
Xcvr 0        REV 01   740-047547  XB34FB00S   CXP Module
Xcvr 2        REV 01   740-047547  XB48FB01H   CXP Module
Xcvr 4        REV 01   740-047547  XB34FB02W   CXP Module
Xcvr 6        REV 01   740-047547  XB34FB01T   CXP Module
Xcvr 8        REV 01   740-047547  XB48FB00W   CXP Module
Xcvr 10       REV 01   740-047547  XB34FB01S   CXP Module
Xcvr 12       REV 01   740-047547  XB34FB03H   CXP Module
Xcvr 14       REV 01   740-047547  XB34FB023   CXP Module
SIB F13 3     REV 01   710-035001  EJ2612      F13 SIB 3D
  B Board     REV 01   711-035082  EJ3815      F13 SIB 3D Mezz
  P Board     REV 01   711-043544  EJ2678      F13 SIB 3D Power
Xcvr 0        REV 01   740-047547  XB48FB04C   CXP Module
Xcvr 2        REV 01   740-047547  XB48FB00Z   CXP Module
Xcvr 4        REV 01   740-047547  XB47FB036   CXP Module
Xcvr 6        REV 01   740-047547  XB47FB029   CXP Module

```

Xcvr 8	REV 01	740-047547	XB48FB02N	CXP Module
Xcvr 10	REV 01	740-047547	XB42FB0CS	CXP Module
Xcvr 12	REV 01	740-047547	XB47FB01X	CXP Module
Xcvr 14	REV 01	740-047547	XB48FB02F	CXP Module
SIB F13 6	REV 05	750-035002	EK2675	F13 SIB 3D
B Board	REV 03	711-035082	EK2612	F13 SIB 3D Mezz
P Board	REV 04	711-043544	EK1179	F13 SIB 3D Power
Xcvr 0	REV 01	740-047547	XB48FB01T	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB02M	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB031	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB04P	CXP Module
Xcvr 8	REV 01	740-047547	XB48FB02T	CXP Module
Xcvr 10	REV 01	740-047547	XB34FB01V	CXP Module
Xcvr 12	REV 01	740-047547	XB48FB02C	CXP Module
Xcvr 14		NON-JNPR		No Module
SIB F13 12	REV 01	710-035001	EJ2631	F13 SIB 3D
B Board	REV 01	711-035082	EJ3808	F13 SIB 3D Mezz
P Board	REV 01	711-043544	EJ2676	F13 SIB 3D Power
SIB F2S 0/0	REV 01	711-034977	EH9829	F2S SIB 3D
B Board	REV 01	711-034979	EH9927	F2S SIB 3D Mezz
SIB F2S 0/2	REV 01	711-034977	EH9791	F2S SIB 3D
B Board	REV 01	711-034979	EH9852	F2S SIB 3D Mezz
SIB F2S 0/4	REV 01	711-034977	EH9803	F2S SIB 3D
B Board	REV 01	711-034979	EH9915	F2S SIB 3D Mezz
SIB F2S 0/6	REV 01	711-034977	EH9763	F2S SIB 3D
B Board	REV 01	711-034979	EH9880	F2S SIB 3D Mezz
SIB F2S 1/0	REV 01	711-034977	EH9757	F2S SIB 3D
B Board	REV 01	711-034979	EH9889	F2S SIB 3D Mezz
SIB F2S 1/2	REV 01	711-034977	EH9815	F2S SIB 3D
B Board	REV 01	711-034979	EH9890	F2S SIB 3D Mezz
SIB F2S 1/4	REV 08	750-034978	EN1954	F2S SIB 3D
B Board	REV 02	711-034979	EN1436	F2S SIB 3D Mezz
SIB F2S 1/6	REV 01	711-034977	EJ7054	F2S SIB 3D
B Board	REV 01	711-034979	EJ8238	F2S SIB 3D Mezz
SIB F2S 2/0	REV 01	711-034977	EH9830	F2S SIB 3D
B Board	REV 01	711-034979	EH9844	F2S SIB 3D Mezz
SIB F2S 2/2	REV 01	711-034977	EH9818	F2S SIB 3D
B Board	REV 01	711-034979	EH9888	F2S SIB 3D Mezz
SIB F2S 2/4	REV 01	711-034977	EH9795	F2S SIB 3D
B Board	REV 01	711-034979	EH9869	F2S SIB 3D Mezz
SIB F2S 2/6	REV 01	711-034977	EJ7026	F2S SIB 3D
B Board	REV 01	711-034979	EJ8273	F2S SIB 3D Mezz
SIB F2S 3/0	REV 01	711-034977	EH9811	F2S SIB 3D
B Board	REV 01	711-034979	EH9892	F2S SIB 3D Mezz
SIB F2S 3/2	REV 01	711-034977	EH9812	F2S SIB 3D
B Board	REV 01	711-034979	EH9877	F2S SIB 3D Mezz
SIB F2S 3/4	REV 08	750-034978	EN1947	F2S SIB 3D
B Board	REV 02	711-034979	EN1471	F2S SIB 3D Mezz
Fan Tray 0	REV 10	760-024497	EH3313	Front Fan Tray
Fan Tray 1	REV 10	760-024497	EH3290	Front Fan Tray
Fan Tray 2	REV 10	760-024502	EH3292	Rear Fan Tray
Fan Tray 3	REV 10	760-024502	EH3287	Rear Fan Tray
Fan Tray 4	REV 10	760-024502	EH3286	Rear Fan Tray
Fan Tray 5	REV 10	760-024502	EH3285	Rear Fan Tray

1cc0-re0:

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN11B23FEAHA	T1600
Midplane	REV 01	710-027486	RC9787	T-series Backplane

FPM GBUS	REV 13	710-002901	BBAG5132	T640 FPM Board
FPM Display	REV 04	710-021387	BBAL9612	T1600 FPM Display
CIP	REV 06	710-002895	BBAN0605	T-series CIP
PEM 0	REV 05	740-036442	1G022060143	Power Entry Module 6x60
PEM 1	REV 05	740-036442	1G022060011	Power Entry Module 6x60
SCG 0	REV 18	710-003423	BBAL7318	T640 Sonet Clock Gen.
SCG 1	REV 18	710-003423	BBAL7255	T640 Sonet Clock Gen.
Routing Engine 0	REV 07	740-026941	P737F-002933	RE-DUO-1800
ad0 3823 MB	SMART CF		201103030490604E604E	Compact Flash
ad1 62720 MB	SMART Lite SATA Drive		20110729028B11D411D4	Disk 1
Routing Engine 1	REV 06	740-026941	P737F-002749	RE-DUO-1800
ad0 3823 MB	SMART CF		2011010504EB99649964	Compact Flash
ad1 62720 MB	SMART Lite SATA Drive		201102140058934A934A	Disk 1
CB 0	REV 11	710-022597	EH3611	LCC Control Board
CB 1	REV 11	710-022597	EH4798	LCC Control Board
FPC 5	REV 17	710-013037	BBAC5333	FPC Type 4-ES
CPU	REV 10	710-016744	BBAB7619	ST-PMB2
PIC 0	REV 18	750-017405	BBAE3420	4x 10GE (LAN/WAN) XFP
Xcvr 0	REV 03	740-014289	T10C90659	XFP-10G-SR
MMB 0	REV 05	710-025563	BBAB9538	ST-MMB2
MMB 1	REV 05	710-025563	BBAB9502	ST-MMB2
FPC 7	REV 01	750-045173	BBAV0032	FPC Type 5-3D
CPU				
SPMB 0	REV 05	710-023321	EG9434	LCC Switch CPU
SPMB 1	REV 05	710-023321	EH3878	LCC Switch CPU
SIB 0	REV 01	750-041657	EH7997	LCC SIB 3D
B Board	REV 01	711-042424	EH7674	LCC SIB 3D Mezz
Xcvr 0	REV 01	740-047547	XB48FB014	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB05A	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB052	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB01B	CXP Module
SIB 1	REV 01	750-041657	EH8023	LCC SIB 3D
B Board	REV 01	711-042424	EH7659	LCC SIB 3D Mezz
Xcvr 0	REV 01	740-047547	XB48FB05J	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB01E	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB01J	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB02S	CXP Module
SIB 2	REV 03	750-041657	EJ6554	LCC SIB 3D
B Board	REV 02	711-042424	EJ5756	LCC SIB 3D Mezz
Xcvr 0	REV 01	740-047547	XB34FB01Z	CXP Module
Xcvr 2	REV 01	740-047547	XB34FB013	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB04Z	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB05N	CXP Module
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Fan Tray -- Rev 4

lcc2-re0:

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN11B3975AHA	T1600
Midplane	REV 01	710-027486	RC9826	T-series Backplane
FPM GBUS	REV 13	710-002901	BBAG5124	T640 FPM Board
FPM Display	REV 03	710-021387	BBAJ1112	T1600 FPM Display
CIP	REV 06	710-002895	BBAL3744	T-series CIP
PEM 0	REV 05	740-036442	1G022060081	Power Entry Module 6x60
PEM 1	REV 05	740-036442	1G022060188	Power Entry Module 6x60
SCG 0	REV 18	710-003423	BBAH8775	T640 Sonet Clock Gen.
SCG 1	REV 18	710-003423	BBAL7272	T640 Sonet Clock Gen.
Routing Engine 0	REV 07	740-026941	P737F-002992	RE-DUO-1800

```

ad0      3823 MB SMART CF                201103030356329E329E Compact Flash
ad1      62720 MB SMART Lite SATA Drive 2011051000488D8B8D8B Disk 1
Routing Engine 1 REV 07 740-026941 P737F-002938 RE-DUO-1800
ad0      3823 MB SMART CF                20110304000F02680268 Compact Flash
ad1      62720 MB SMART Lite SATA Drive 201105300A70F325F325 Disk 1
CB 0                      REV 11 710-022597 EH4805 LCC Control Board
CB 1                      REV 11 710-022597 EH4786 LCC Control Board
FPC 1                      REV 01 710-033873 BBAH0320 FPC Type 3-ES
CPU                      REV 11 710-016744 BBAF3281 ST-PMB2
MMB 0                      REV 06 710-025563 BBAF5061 ST-MMB2
FPC 5                      REV 04 710-033871 BBAM5070 FPC Type 4-ES
CPU                      REV 11 710-016744 BBAM6653 ST-PMB2
PIC 1                      REV 20 750-017405 BBAM1296 4x 10GE (LAN/WAN) XFP
Xcvr 0                      REV 03 740-014289 T10B42981 XFP-10G-SR
MMB 0                      REV 07 710-025563 BBAN2631 ST-MMB2
MMB 1                      REV 07 710-025563 BBAN2538 ST-MMB2
SPMB 0                      REV 05 710-023321 EH3903 LCC Switch CPU
SPMB 1                      REV 05 710-023321 EH3902 LCC Switch CPU
SIB 0                      REV 01 750-041657 EH8019 LCC SIB 3D
B Board                      REV 01 711-042424 EH7680 LCC SIB 3D Mezz
Xcvr 0                      REV 01 740-047547 XB48FB04F CXP Module
Xcvr 2                      REV 01 740-047547 XB48FB04S CXP Module
Xcvr 4                      REV 01 740-047547 XB48FB04B CXP Module
Xcvr 6                      REV 01 740-047547 XB48FB043 CXP Module
SIB 1                      REV 01 750-041657 EH8012 LCC SIB 3D
B Board                      REV 01 711-042424 EH7658 LCC SIB 3D Mezz
Xcvr 0                      REV 01 740-047547 XB48FB05E CXP Module
Xcvr 2                      REV 01 740-047547 XB48FB01Z CXP Module
Xcvr 4                      REV 01 740-047547 XB48FB018 CXP Module
Xcvr 6                      REV 01 740-047547 XB48FB054 CXP Module
SIB 2                      REV 01 750-041657 EH7993 LCC SIB 3D
B Board                      REV 01 711-042424 EH7678 LCC SIB 3D Mezz
Xcvr 0                      REV 01 740-047547 XB48FB05C CXP Module
Xcvr 2                      REV 01 740-047547 XB47FB00N CXP Module
Xcvr 4                      REV 01 740-047547 XB48FB05U CXP Module
Xcvr 6                      REV 01 740-047547 XB48FB05L CXP Module
Fan Tray 0                  Front Top Fan Tray
Fan Tray 1                  Front Bottom Fan Tray
Fan Tray 2                  Rear Fan Tray -- Rev 4

```

show chassis hardware lcc (TX Matrix Plus Router with 3D SIBs)

```

user@host> show chassis hardware lcc 0
lcc0-re0:

```

```

-----
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               JN11B23FEAHA  T1600
Midplane      REV 01  710-027486  RC9787         T-series Backplane
FPM GBUS      REV 13  710-002901  BBAG5132       T640 FPM Board
FPM Display   REV 04  710-021387  BBAL9612       T1600 FPM Display
CIP           REV 06  710-002895  BBAN0605       T-series CIP
PEM 0         REV 05  740-036442  1G022060143   Power Entry Module 6x60
PEM 1         REV 05  740-036442  1G022060011   Power Entry Module 6x60
SCG 0         REV 18  710-003423  BBAL7318       T640 Sonet Clock Gen.
SCG 1         REV 18  710-003423  BBAL7255       T640 Sonet Clock Gen.
Routing Engine 0 REV 07  740-026941  P737F-002933  RE-DUO-1800
Routing Engine 1 REV 06  740-026941  P737F-002749  RE-DUO-1800
CB 0          REV 11  710-022597  EH3611         LCC Control Board
CB 1          REV 11  710-022597  EH4798         LCC Control Board
FPC 5         REV 17  710-013037  BBAC5333       FPC Type 4-ES

```

CPU	REV 10	710-016744	BBAB7619	ST-PMB2
PIC 0	REV 18	750-017405	BBAE3420	4x 10GE (LAN/WAN) XFP
Xcvr 0	REV 03	740-014289	T10C90659	XFP-10G-SR
MMB 0	REV 05	710-025563	BBAB9538	ST-MMB2
MMB 1	REV 05	710-025563	BBAB9502	ST-MMB2
FPC 7	REV 01	750-045173	BBAV0032	FPC Type 5-3D
CPU				
SPMB 0	REV 05	710-023321	EG9434	LCC Switch CPU
SPMB 1	REV 05	710-023321	EH3878	LCC Switch CPU
SIB 0	REV 01	750-041657	EH7997	LCC SIB 3D
B Board	REV 01	711-042424	EH7674	LCC SIB 3D Mezz
Xcvr 0	REV 01	740-047547	XB48FB014	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB05A	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB052	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB01B	CXP Module
SIB 1	REV 01	750-041657	EH8023	LCC SIB 3D
B Board	REV 01	711-042424	EH7659	LCC SIB 3D Mezz
Xcvr 0	REV 01	740-047547	XB48FB05J	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB01E	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB01J	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB02S	CXP Module
SIB 2	REV 03	750-041657	EJ6554	LCC SIB 3D
B Board	REV 02	711-042424	EJ5756	LCC SIB 3D Mezz
Xcvr 0	REV 01	740-047547	XB34FB01Z	CXP Module
Xcvr 2	REV 01	740-047547	XB34FB013	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB04Z	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB05N	CXP Module
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Fan Tray -- Rev 4

show chassis hardware sfc (TX Matrix Plus Router with 3D SIBs)

```
user@host> show chassis hardware sfc 0
sfc0-re0:
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN11CAAA4AHB	TXP
Midplane	REV 05	710-022574	ABAC4696	SFC Midplane
FPM Display	REV 09	710-024027	EH3138	TXP FPM Display
CIP 0	REV 12	710-023792	EF6349	TXP CIP
CIP 1	REV 12	710-023792	EG5294	TXP CIP
PEM 0	Rev 06	740-027463	XH04595	Power Entry Module
PEM 1	Rev 06	740-027463	XH04592	Power Entry Module
Routing Engine 0	REV 07	740-026942	P737A-002541	RE-DUO-2600
Routing Engine 1	REV 07	740-026942	P737A-002602	RE-DUO-2600
CB 0	REV 15	710-022606	EH4376	SFC Control Board
CB 1	REV 15	710-022606	EH4379	SFC Control Board
SPMB 0		BUILTIN		SFC Switch CPU
SPMB 1		BUILTIN		SFC Switch CPU
SIB F13 0	REV 10	750-035002	EM9305	F13 SIB 3D
B Board	REV 06	711-035082	EM9667	F13 SIB 3D Mezz
P Board	REV 05	711-043544	EM9708	F13 SIB 3D Power
Xcvr 0	REV 01	740-047547	XB34FB00S	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB01H	CXP Module
Xcvr 4	REV 01	740-047547	XB34FB02W	CXP Module
Xcvr 6	REV 01	740-047547	XB34FB01T	CXP Module
Xcvr 8	REV 01	740-047547	XB48FB00W	CXP Module
Xcvr 10	REV 01	740-047547	XB34FB01S	CXP Module
Xcvr 12	REV 01	740-047547	XB34FB03H	CXP Module

Xcvr 14	REV 01	740-047547	XB34FB023	CXP Module
SIB F13 3	REV 01	710-035001	EJ2612	F13 SIB 3D
B Board	REV 01	711-035082	EJ3815	F13 SIB 3D Mezz
P Board	REV 01	711-043544	EJ2678	F13 SIB 3D Power
Xcvr 0	REV 01	740-047547	XB48FB04C	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB00Z	CXP Module
Xcvr 4	REV 01	740-047547	XB47FB036	CXP Module
Xcvr 6	REV 01	740-047547	XB47FB029	CXP Module
Xcvr 8	REV 01	740-047547	XB48FB02N	CXP Module
Xcvr 10	REV 01	740-047547	XB42FB0CS	CXP Module
Xcvr 12	REV 01	740-047547	XB47FB01X	CXP Module
Xcvr 14	REV 01	740-047547	XB48FB02F	CXP Module
SIB F13 6	REV 05	750-035002	EK2675	F13 SIB 3D
B Board	REV 03	711-035082	EK2612	F13 SIB 3D Mezz
P Board	REV 04	711-043544	EK1179	F13 SIB 3D Power
Xcvr 0	REV 01	740-047547	XB48FB01T	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB02M	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB031	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB04P	CXP Module
Xcvr 8	REV 01	740-047547	XB48FB02T	CXP Module
Xcvr 10	REV 01	740-047547	XB34FB01V	CXP Module
Xcvr 12	REV 01	740-047547	XB48FB02C	CXP Module
Xcvr 14		NON-JNPR		No Module
SIB F13 12	REV 01	710-035001	EJ2631	F13 SIB 3D
B Board	REV 01	711-035082	EJ3808	F13 SIB 3D Mezz
P Board	REV 01	711-043544	EJ2676	F13 SIB 3D Power
SIB F2S 0/0	REV 01	711-034977	EH9829	F2S SIB 3D
B Board	REV 01	711-034979	EH9927	F2S SIB 3D Mezz
SIB F2S 0/2	REV 01	711-034977	EH9791	F2S SIB 3D
B Board	REV 01	711-034979	EH9852	F2S SIB 3D Mezz
SIB F2S 0/4	REV 01	711-034977	EH9803	F2S SIB 3D
B Board	REV 01	711-034979	EH9915	F2S SIB 3D Mezz
SIB F2S 0/6	REV 01	711-034977	EH9763	F2S SIB 3D
B Board	REV 01	711-034979	EH9880	F2S SIB 3D Mezz
SIB F2S 1/0	REV 01	711-034977	EH9757	F2S SIB 3D
B Board	REV 01	711-034979	EH9889	F2S SIB 3D Mezz
SIB F2S 1/2	REV 01	711-034977	EH9815	F2S SIB 3D
B Board	REV 01	711-034979	EH9890	F2S SIB 3D Mezz
SIB F2S 1/4	REV 08	750-034978	EN1954	F2S SIB 3D
B Board	REV 02	711-034979	EN1436	F2S SIB 3D Mezz
SIB F2S 1/6	REV 01	711-034977	EJ7054	F2S SIB 3D
B Board	REV 01	711-034979	EJ8238	F2S SIB 3D Mezz
SIB F2S 2/0	REV 01	711-034977	EH9830	F2S SIB 3D
B Board	REV 01	711-034979	EH9844	F2S SIB 3D Mezz
SIB F2S 2/2	REV 01	711-034977	EH9818	F2S SIB 3D
B Board	REV 01	711-034979	EH9888	F2S SIB 3D Mezz
SIB F2S 2/4	REV 01	711-034977	EH9795	F2S SIB 3D
B Board	REV 01	711-034979	EH9869	F2S SIB 3D Mezz
SIB F2S 2/6	REV 01	711-034977	EJ7026	F2S SIB 3D
B Board	REV 01	711-034979	EJ8273	F2S SIB 3D Mezz
SIB F2S 3/0	REV 01	711-034977	EH9811	F2S SIB 3D
B Board	REV 01	711-034979	EH9892	F2S SIB 3D Mezz
SIB F2S 3/2	REV 01	711-034977	EH9812	F2S SIB 3D
B Board	REV 01	711-034979	EH9877	F2S SIB 3D Mezz
SIB F2S 3/4	REV 08	750-034978	EN1947	F2S SIB 3D
B Board	REV 02	711-034979	EN1471	F2S SIB 3D Mezz
Fan Tray 0	REV 10	760-024497	EH3313	Front Fan Tray
Fan Tray 1	REV 10	760-024497	EH3290	Front Fan Tray
Fan Tray 2	REV 10	760-024502	EH3292	Rear Fan Tray
Fan Tray 3	REV 10	760-024502	EH3287	Rear Fan Tray


```

Fan Tray 4      REV 10  760-024502  EH3286      Rear Fan Tray
Fan Tray 5      REV 10  760-024502  EH3285      Rear Fan Tray

```

show chassis hardware (16-Port 10-Gigabit Ethernet MPC with SFP+ Optics [MX Series Routers])

```
user@host> show chassis hardware
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN112D865AFA	MX960
Midplane	REV 03	710-013698	TS3339	MX960 Backplane
FPM Board	REV 03	710-014974	WW6267	Front Panel Display
PDM	Rev 03	740-013110	QCS12485026	Power Distribution
Module				
PEM 0	Rev 04	740-013682	QCS12434086	PS 1.7kW; 200-240VAC
in				
PEM 1	Rev 04	740-013682	QCS1243408Z	PS 1.7kW; 200-240VAC
in				
PEM 2	Rev 04	740-013682	QCS1243407X	PS 1.7kW; 200-240VAC
in				
Routing Engine 0	REV 07	740-015113	9009009677	RE-S-1300
Routing Engine 1	REV 07	740-015113	9009011510	RE-S-1300
CB 0	REV 03	710-021523	XF0394	MX SCB
CB 1	REV 03	710-021523	XF0550	MX SCB
CB 2	REV 03	710-021523	XD7455	MX SCB
FPC 4	REV 02	750-028467	JR6127	MPC M 16x 10GE
CPU	REV 02	711-029089	JX0129	AS PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Fan Tray 0	REV 05	740-014971	TP9990	Fan Tray
Fan Tray 1	REV 05	740-014971	VS1709	Fan Tray

show chassis hardware (MPC3E [MX Series Routers])

```
user@host> show chassis hardware
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN1101AFEAFB	MX480
Midplane	REV 05	710-017414	TR4444	MX480 Midplane
FPM Board	REV 02	710-017254	KG6056	Front Panel Display
PEM 0	Rev 03	740-017330	QCS082090FC	PS 1.2-1.7kW; 100-240V
PEM 1	Rev 03	740-017330	QCS082090FD	PS 1.2-1.7kW; 100-240V
Routing Engine 0	REV 07	740-013063	9009004124	RE-S-2000
Routing Engine 1	REV 07	740-013063	9009005569	RE-S-2000
CB 0	REV 07	710-021523	XZ3587	MX SCB
CB 1	REV 03	710-021523	KH8306	MX SCB
FPC 1	REV 04.1.07	750-033205	P1240	MPC Type 3
CPU	REV 01	711-035209	YL0504	HMPD PMB 2G
MIC 1	REV 10	750-033199	YX4495	1X100GE CFP
PIC 2		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-032210	C22CQNE	CFP-100G-LR4
FPC 2	REV 26	750-016670	KH0045	DPCE 40x 1GE R EQ
CPU	REV 07	710-013713	KF5448	DPC PMB
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 0	REV 01	740-011613	PF21JHU	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN) EQ

Xcvr 9	REV 01	740-011613	AM0813S8ZL6	SFP-SX
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 0	REV 02	740-011613	PGL2KYF	SFP-SX
Xcvr 2	REV 01	740-011613	AM0806S8N4P	SFP-SX
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 5	REV 01	740-011613	AM0815S967N	SFP-SX
Xcvr 7	REV 01	740-011613	AM0806S8N1X	SFP-SX
Xcvr 8	REV 01	740-011613	AM0815S967J	SFP-SX
Xcvr 9	REV 01	740-011613	AM0815S967M	SFP-SX
FPC 3	REV 12.2.09	750-033205	YR9443	MPC Type 3
CPU	REV 03	711-035209	YL6931	HMPC PMB 2G
MIC 0	REV 05	750-033199	YR3269	1X100GE CFP
PIC 0		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-032210	ULH0KG3	CFP-100G-LR4
MIC 1	REV 02	750-033199	YG3245	1X100GE CFP
PIC 2		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-032210	ULH0KGF	CFP-100G-LR4
FPC 4	REV 12.3.09	750-033205	YR9437	MPC Type 3
CPU	REV 03	711-035209	YT5857	HMPC PMB 2G
MIC 0	REV 05	750-033199	YR3295	1X100GE CFP
PIC 0		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0		NON-JNPR	X12000187	CFP-100G-SR10
MIC 1	REV 10	750-033199	YX4518	1X100GE CFP
PIC 2		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-035329	X12J00008	CFP-100G-SR10
FPC 5	REV 06	750-024884	JW9769	MPC Type 2 3D EQ
CPU	REV 02	711-028401	JR6158	MPC PMB 2G Proto
MIC 0	REV 05	750-028387	JR6197	3D 4x 10GE XFP
PIC 0		BUILTIN	BUILTIN	2x 10GE XFP
Xcvr 0	REV 01	740-014289	T07M71112	XFP-10G-SR
Xcvr 1	REV 02	740-014289	T08L85610	XFP-10G-SR
PIC 1		BUILTIN	BUILTIN	2x 10GE XFP
MIC 1	REV 22	750-028392	YM0053	3D 20x 1GE(LAN) SFP
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 0	REV 01	740-011613	AM0703S005B	SFP-SX
Xcvr 1	REV 01	740-011613	E07L01352	SFP-SX
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 5	REV 01	740-013111	6500217	SFP-T
Xcvr 9	REV 02	740-013111	8499527	SFP-T
Fan Tray				Left Fan Tray

The PIC number for MIC 1 always starts from 2 (even if the first MIC is a 1X100GE CFP or a legacy MIC).

show chassis hardware (QFX3500 Switches)

```
user@switch> show chassis hardware
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis				QFX3500
Routing Engine 0		BUILTIN	BUILTIN	QFX Routing Engine
FPC 0	REV 04	750-044071	BBAR3902	QFX3500-48S4Q-AFI
CPU		BUILTIN	BUILTIN	FPC CPU
PIC 0		BUILTIN	BUILTIN	48x 10G-SFP+
PIC 1		BUILTIN	BUILTIN	15x 10G-SFP+
MGMT BRD	REV 02	750-044063	BBAR0398	QFX3500-MGMT-SFP-AFO
Xcvr 0	REV 01	740-011614	AC0946S0BD1	SFP-LX10
Xcvr 1	REV 02	740-013111	A281922	SFP-T
Power Supply 0	Rev 04	740-032091	UI00677	JPSU-650W-AC-AFI
Power Supply 1	REV 00	740-041741	VJ00162	JPSU-650W-AC-AFO

```

Fan Tray 0                               QFX Fan Tray, Back to
Front Airlfow
Fan Tray 1                               QFX Fan Tray, Back to
Front Airlfow
Fan Tray 2                               QFX Fan Tray, Back to
Front Airlfow

```

show chassis hardware detail (QFX3500 Switches)

```
user@switch> show chassis hardware detail
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN000TEST5	QFX3500
Routing Engine 0		BUILTIN	BUILTIN	QFX Routing Engine
FPC 0	REV 05	750-036931	EE0823	QFX3500-48S4Q-AFI

CPU		BUILTIN	BUILTIN	FPC CPU
PIC 0		BUILTIN	BUILTIN	48x 10G-SFP+
Xcvr 0	REV 01	740-030589	S99E270079	SFP+-10G-LPBK
Xcvr 1	REV 01	740-030589	S9AK450099	SFP+-10G-LPBK
Xcvr 2	REV 01	740-030589	S99E270078	SFP+-10G-LPBK
Xcvr 3	REV 01	740-030589	S9AK450098	SFP+-10G-LPBK
Xcvr 4	REV 01	740-030589	S99E270075	SFP+-10G-LPBK
Xcvr 5	REV 01	740-030589	S9AK450093	SFP+-10G-LPBK
Xcvr 6	REV 01	740-030589	S9AK450097	SFP+-10G-LPBK
Xcvr 7	REV 01	740-030589	S9AK450095	SFP+-10G-LPBK
Xcvr 8	REV 01	740-030589	S99E270072	SFP+-10G-LPBK
Xcvr 9	REV 01	740-030589	S99E270073	SFP+-10G-LPBK
Xcvr 10	REV 01	740-030589	S99E270080	SFP+-10G-LPBK
Xcvr 11	REV 01	740-030589	S9AK450169	SFP+-10G-LPBK
Xcvr 12	REV 01	740-030589	S99E270076	SFP+-10G-LPBK
Xcvr 13	REV 01	740-030589	S9AK450167	SFP+-10G-LPBK
Xcvr 14	REV 01	740-030589	S9AK450170	SFP+-10G-LPBK
Xcvr 15	REV 01	740-030589	S9AK450166	SFP+-10G-LPBK
Xcvr 16	REV 01	740-030589	S9AK450092	SFP+-10G-LPBK
Xcvr 17	REV 01	740-030589	S9AK450163	SFP+-10G-LPBK
Xcvr 18	REV 01	740-030589	S9AK450094	SFP+-10G-LPBK
Xcvr 19	REV 01	740-030589	S9AK450100	SFP+-10G-LPBK
Xcvr 20	REV 01	740-030589	S9AK450168	SFP+-10G-LPBK
Xcvr 21	REV 01	740-030589	S9AK450165	SFP+-10G-LPBK
Xcvr 22	REV 01	740-030589	S9AK450073	SFP+-10G-LPBK
Xcvr 23	REV 01	740-030589	S9AK450164	SFP+-10G-LPBK
Xcvr 24	REV 01	740-030589	S9AK450074	SFP+-10G-LPBK
Xcvr 25	REV 01	740-030589	SA62270195	SFP+-10G-LPBK
Xcvr 26	REV 01	740-030589	S9AK450078	SFP+-10G-LPBK
Xcvr 27	REV 01	740-030589	S9AK450024	SFP+-10G-LPBK
Xcvr 28	REV 01	740-030589	S9AK450027	SFP+-10G-LPBK
Xcvr 29	REV 01	740-030589	S9AK450080	SFP+-10G-LPBK
Xcvr 30	REV 01	740-030589	S9AK450030	SFP+-10G-LPBK
Xcvr 31	REV 01	740-030589	S9AK450025	SFP+-10G-LPBK
Xcvr 32	REV 01	740-030589	S9AK450023	SFP+-10G-LPBK
Xcvr 33	REV 01	740-030589	S9AK450075	SFP+-10G-LPBK
Xcvr 34	REV 01	740-030589	S9AK450161	SFP+-10G-LPBK
Xcvr 35	REV 01	740-030589	S9AK450071	SFP+-10G-LPBK
Xcvr 36	REV 01	740-030589	S9AK450072	SFP+-10G-LPBK
Xcvr 37	REV 01	740-030589	S9AK450022	SFP+-10G-LPBK
Xcvr 38	REV 01	740-030589	S9AK450021	SFP+-10G-LPBK
Xcvr 39	REV 01	740-030589	S9AK450175	SFP+-10G-LPBK
Xcvr 40	REV 01	740-030589	S9AK450162	SFP+-10G-LPBK
Xcvr 41	REV 01	740-030589	S99E270074	SFP+-10G-LPBK

Xcvr 42	REV 01	740-030589	S9AK450174	SFP+-10G-LPBK
Xcvr 43	REV 01	740-030589	S9AK450077	SFP+-10G-LPBK
Xcvr 44	REV 01	740-030589	S9AK450076	SFP+-10G-LPBK
Xcvr 45	REV 01	740-030589	S9AK450026	SFP+-10G-LPBK
Xcvr 46	REV 01	740-030589	S9AK450079	SFP+-10G-LPBK
Xcvr 47	REV 01	740-030589	S9AK450029	SFP+-10G-LPBK
PIC 1		BUILTIN	BUILTIN	15x 10G-SFP+
Xcvr 1	REV 01	740-032986	QA170087	QSFP+-40G-SR4
Xcvr 4	REV 01	740-032986	QA360442	QSFP+-40G-SR4
Xcvr 8	REV 01	740-032986	QA170091	QSFP+-40G-SR4
Xcvr 12	REV 01	740-032986	QA170042	QSFP+-40G-SR4
MGMT BRD	REV 08	750-036946	EE0731	QFX3500-MB
Power Supply 0	Rev 04	740-032091	UI00690	QFX PS 650W AC
Power Supply 1	Rev 04	740-032091	UI00679	QFX PS 650W AC
Fan Tray 0				QFX Fan Tray
Fan Tray 1				QFX Fan Tray

show chassis hardware models (QFX3500 Switches)

```
user@switch> show chassis hardware models
Hardware inventory:
```

Item	Version	Part number	Serial number	FRU model number
Routing Engine 0		BUILTIN	BUILTIN	
FPC 0	REV 02	711-032234	EC4074	
Power Supply 0	PSMI 2C	11-d65800	--	

show chassis hardware clei-models (QFX3500 Switches)

```
user@switch> show chassis hardware clei-models
Hardware inventory:
```

Item	Version	Part number	CLEI code	FRU model number
Routing Engine 0		BUILTIN		
FPC 0	REV 02	711-032234		
Power Supply 0	PSMI 2C	11-d65800		

show chassis hardware clei-models (QFX5100 Switches)

```
user@switch> show chassis hardware clei-models
Hardware inventory:
```

Item	Version	Part number	CLEI code	FRU model number
Routing Engine 0		BUILTIN	CMMNV10BRA	
FPC 0	REV 01	611-053010	CMMNV10BRA	
PIC 0		BUILTIN	CMMNV10BRA	
Power Supply 0	REV 03	740-053352	MUPABHBAA	JPSU-850W-AC-AFO
Power Supply 1	REV 03	740-053352	MUPABHBAA	JPSU-850W-AC-AFO
Fan Tray 0				QFX5100-96S-FANAFO
Fan Tray 1				QFX5100-96S-FANAFO
Fan Tray 2				QFX5100-96S-FANAFO

show chassis hardware interconnect-device (QFabric Systems)

```
user@switch> show chassis hardware interconnect-device interconnect1
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis	REV 07			QFX_olive
Midplane	REV 07	750-021261	BH0208188289	QFX Midplane
CB 0	REV 07	750-021261	BH0208188289	QFXIC08-CB4S

show chassis hardware node-device (QFabric Systems)

```
user@switch> show chassis hardware node-device node1
```

```

Routing Engine 0   BUILTIN   BUILTIN           QFX Routing Engine
node1             REV 05   711-032234   ED3694           QFX3500-48S4Q-AFI

CPU
PIC 0             BUILTIN   BUILTIN           FPC CPU
Xcvr 8           REV 01   740-030658   AD0946A028B     48x 10G-SFP+
SFP+-10G-USR
...

```

show chassis hardware (PTX5000 Packet Transport Router)

```

user@host> show chassis hardware
Hardware inventory:
Item                Version  Part number  Serial number  Description
Chassis              REV 03   711-031896   JN11D1FD7AJA  PTX5000
Midplane             REV 08   760-030647   ABAC5589      Midplane-8S
FPM                  REV 05   740-032019   ZE00006       Front Panel Display
PDU 0                Rev 05   740-032022   ZJ00018       DC Power Dist Unit
  PSM 0              Rev 04   740-032022   ZC00052       DC 12V Power Supply
  PSM 1              Rev 04   740-032022   ZD00051       DC 12V Power Supply
  PSM 2              Rev 05   740-032022   ZJ00060       DC 12V Power Supply
  PSM 3              Rev 04   750-030653   EG3703        Clock Generator
CCG 0                REV 04   750-030653   EG3698        Clock Generator
Routing Engine 0     REV 05   740-026942   P737A-002231  RE-DUO-2600
Routing Engine 1     REV 06   740-026942   P737A-002438  RE-DUO-2600
CB 0                 REV 08   750-030625   EG5519        Control Board
CB 1                 REV 08   750-030625   EG5516        Control Board
FPC 0                REV 18   750-036844   EJ3080        FPC
  CPU                REV 12   711-030686   EJ3260        SNG PMB
FPC 2                REV 13   750-036844   EG5065        FPC
  CPU                REV 09   711-030686   EG4082        SNG PMB
  PIC 0              REV 14   750-031913   EG5127        24x 10GE(LAN) SFP+
    Xcvr 0           REV 01   740-031980   143363A00240  SFP+-10G-SR
    Xcvr 1           REV 01   740-031981   UK90PZ1       SFP+-10G-LR
    Xcvr 2           REV 01   740-031980   AD1141A04XH   SFP+-10G-SR
    Xcvr 3           REV 01   740-031981   UK90Q46       SFP+-10G-LR
    Xcvr 4           REV 01   740-031980   AD1141A04X4   SFP+-10G-SR
    Xcvr 6           REV 01   740-031980   B11H02560     SFP+-10G-SR
    Xcvr 7           REV 01   740-031980   B11C01589     SFP+-10G-SR
    Xcvr 8           REV 01   740-031980   AD1141A04XF   SFP+-10G-SR
    Xcvr 10          REV 01   740-031980   123363A01094  SFP+-10G-SR
    Xcvr 11          REV 01   740-031980   AK80LKF       SFP+-10G-SR
    Xcvr 12          REV 01   740-031980   183363A01528  SFP+-10G-SR
    Xcvr 14          REV 01   740-031980   193363A01079  SFP+-10G-SR
    Xcvr 15          REV 01   740-031980   AK80MC8       SFP+-10G-SR
    Xcvr 16          REV 01   740-031980   AJCOBHC       SFP+-10G-SR
    Xcvr 19          REV 01   740-021309   J08D26856     SFP+-10G-LR
    Xcvr 21          REV 01   740-031980   AK80KCT       SFP+-10G-SR
    Xcvr 22          REV 01   740-031981   UK90PZL       SFP+-10G-LR
    Xcvr 23          REV 01   740-031980   AK80N1V       SFP+-10G-SR
FPC 3                REV 13   750-036844   EG5074        FPC
  CPU                REV 09   711-030686   EG4064        SNG PMB
  PIC 1              REV 10   750-031903   EG0325        SNG Load
FPC 5                REV 06   750-036844   EH3198        FPC
  CPU
  PIC 0              REV 14   750-031913   EG5134        24x 10GE(LAN) SFP+
    Xcvr 0           REV 01   740-031980   AK80LBH       SFP+-10G-SR
    Xcvr 1           REV 01   740-031980   B11B03724     SFP+-10G-SR
    Xcvr 2           REV 01   740-031980   AK80FMH       SFP+-10G-SR
    Xcvr 5           REV 01   740-031980   B11J00818     SFP+-10G-SR

```

Xcvr 6	REV 01	740-031980	193363A00743	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	B11B06125	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	B11H02529	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	AK80LFB	SFP+-10G-SR
Xcvr 12	REV 01	740-031980	193363A01061	SFP+-10G-SR
Xcvr 15	REV 01	740-031980	B11J00687	SFP+-10G-SR
Xcvr 16	REV 01	740-031980	193363A00738	SFP+-10G-SR
Xcvr 18	REV 01	740-031980	AK80MQX	SFP+-10G-SR
Xcvr 19	REV 01	740-021309	J08C17257	SFP+-10G-LR
Xcvr 22	REV 01	740-031980	B11J00730	SFP+-10G-SR
Xcvr 23	REV 01	740-031980	AK80KEE	SFP+-10G-SR
PIC 1	REV 08	750-036710	EG3105	2x 40GE CFP
Xcvr 0	REV 01	740-034554	B260HLT	CFP-40G-LR4
Xcvr 1	REV 01	740-034554	B11C02847	CFP-40G-LR4
FPC 6	REV 18	750-036844	EJ4391	FPC
CPU	REV 12	711-030686	EJ3257	SNG PMB
FPC 7	REV 18	750-036844	EJ4382	FPC
CPU	REV 12	711-030686	EJ3238	SNG PMB
SPMB 0	REV 10	711-030686	EG5418	SNG PMB
SPMB 1	REV 09	711-030686	EG5373	SNG PMB
SIB 0	REV 07	750-030631	EG4858	SIB-I-8S
SIB 1	REV 07	750-030631	EG4872	SIB-I-8S
SIB 2	REV 07	750-030631	EG4866	SIB-I-8S
SIB 3	REV 07	750-030631	EG6011	SIB-I-8S
SIB 4	REV 07	750-030631	EG4907	SIB-I-8S
SIB 5	REV 07	750-030631	EG4879	SIB-I-8S
SIB 6	REV 07	750-030631	EG4864	SIB-I-8S
SIB 7	REV 07	750-030631	EG4899	SIB-I-8S
SIB 8	REV 07	750-030631	EG4880	SIB-I-8S
Fan Tray 0	REV 04	760-032784	EG1496	Vertical Fan Tray
Fan Tray 1	REV 04	760-030642	EG1335	Horizontal Fan Tray
Fan Tray 2	REV 02	760-030642	ED4952	Horizontal Fan Tray

show chassis hardware (PTX5000 Packet Transport Router with AC PSM and PDU)

user@host> show chassis hardware

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN12223A6AJA	PTX5000
Midplane	REV 16	750-035893	ACRA1350	Midplane-8S
FPM	REV 12	760-030647	BBBD5625	Front Panel Display
PDU 0	Rev 01	740-048338	1GB83360005	High Capacity AC WYE PDU
PSM 0	Rev 01	740-048334	1GB43360074	High Capacity AC PSM
PSM 1	Rev 01	740-048334	1GB43360001	High Capacity AC PSM
PSM 2	Rev 01	740-048334	1GB43360104	High Capacity AC PSM
PSM 3	Rev 01	740-048334	1GB43360042	High Capacity AC PSM
PSM 4	Rev 01	740-048334	1GB43360068	High Capacity AC PSM
PSM 5	Rev 01	740-048334	1GB43360080	High Capacity AC PSM
PSM 6	Rev 01	740-048334	1GB43360046	High Capacity AC PSM
PSM 7	Rev 01	740-048334	1GB43360100	High Capacity AC PSM
PDU 1	Rev 01	740-048338	1GB83360006	High Capacity AC WYE PDU
PSM 0	Rev 01	740-048334	1GB43360069	High Capacity AC PSM
PSM 1	Rev 01	740-048334	1GB43360099	High Capacity AC PSM
PSM 2	Rev 01	740-048334	1GB43360050	High Capacity AC PSM
PSM 3	Rev 01	740-048334	1GB43360095	High Capacity AC PSM
PSM 4	Rev 01	740-048334	1GB43360101	High Capacity AC PSM
PSM 5	Rev 01	740-048334	1GB43360075	High Capacity AC PSM
PSM 6	Rev 01	740-048334	1GB43360047	High Capacity AC PSM
PSM 7	Rev 01	740-048334	1GB43360019	High Capacity AC PSM

```

CCG 0          REV 09   750-030653   BBAZ5345      Clock Generator
...

```

show chassis hardware (PTX5000 Packet Transport Router with FPC2-PTX-P1A)

```

user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number  Description
Chassis                               JN1204FC0AJA  PTX5000
Midplane         REV 11   750-035893   ACAB8038      Midplane-8S
FPM              REV 12   760-030647   BBBD5619      Front Panel
Display
PDU 0            Rev 04   740-048336   1GB93470043   High Capacity DC PDU
  PSM 0          Rev 04   740-046988   1GB63500184   High Capacity DC PSM
  PSM 2          Rev 04   740-046988   1GB63500169   High Capacity DC PSM
  PSM 4          Rev 04   740-046988   1GB63500306   High Capacity DC PSM
  PSM 6          Rev 04   740-046988   1GB63500074   High Capacity DC PSM
PDU 1            Rev 04   740-048336   1GB93470045   High Capacity DC PDU
  PSM 1          Rev 04   740-046988   1GB63500193   High Capacity DC PSM
  PSM 3          Rev 04   740-046988   1GB63500143   High Capacity DC PSM
  PSM 5          Rev 04   740-046988   1GB63500146   High Capacity DC PSM
  PSM 7          Rev 04   740-046988   1GB63500192   High Capacity DC PSM
CCG 0            REV 09   750-030653   BBBC1909      Clock Generator
CCG 1            REV 09   750-030653   BBBD2970      Clock Generator
...

```

show chassis hardware clei-models (PTX5000 Packet Transport Router)

```

user@host> show chassis hardware clei-models
Hardware inventory:
Item             Version  Part number  CLEI code      FRU model number
FPM              REV 08   760-030647   PROTOXCLEI     CRAFT-PTX5000-S
PDU 0            Rev 05   740-032019   IPUPAHLKAA     PWR-SAN-PDU-DC
  PSM 0          Rev 05   740-032022   IPUPAHNKAA     PSM-PTX-DC-120-S
  PSM 1          Rev 04   740-032022   032022XXXX     PWR-SAN-12-DC
  PSM 2          Rev 04   740-032022   032022XXXX     PWR-SAN-12-DC
  PSM 3          Rev 05   740-032022   IPUPAHNKAA     PSM-PTX-DC-120-S
CCG 0            REV 04   750-030653   PROTOXCLEI     CCG-PTX-S
CCG 1            REV 04   750-030653   PROTOXCLEI     CCG-PTX-S
Routing Engine 0 REV 05   740-026942                                     RE-DUO-C2600-16G-S
Routing Engine 1 REV 06   740-026942                                     RE-DUO-C2600-16G-S
CB 0             REV 08   750-030625   PROTOXCLEI     CB-PTX-S
CB 1             REV 08   750-030625   PROTOXCLEI     CB-PTX-S
FPC 0            REV 18   750-036844   PROTOXCLEI     FPC-PTX-P1-A
FPC 2            REV 13   750-036844   PROTOXCLEI     FPC-PTX-P1-A
  PIC 0          REV 14   750-031913   PROTOXCLEI     P1-PTX-24-10GE-SFPP
FPC 5            REV 13   750-036844   PROTOXCLEI     FPC-PTX-P1-A
  PIC 0          REV 14   750-031913   PROTOXCLEI     P1-PTX-24-10GE-SFPP
FPC 6            REV 18   750-036844   PROTOXCLEI     FPC-PTX-P1-A
FPC 7            REV 18   750-036844   PROTOXCLEI     FPC-PTX-P1-A
SIB 0            REV 07   750-030631   PROTOXCLEI     SIB-I-PTX5008
SIB 1            REV 07   750-030631   PROTOXCLEI     SIB-I-PTX5008
SIB 2            REV 07   750-030631   PROTOXCLEI     SIB-I-PTX5008
SIB 3            REV 07   750-030631   PROTOXCLEI     SIB-I-PTX5008
SIB 4            REV 07   750-030631   PROTOXCLEI     SIB-I-PTX5008
SIB 5            REV 07   750-030631   PROTOXCLEI     SIB-I-PTX5008
SIB 6            REV 07   750-030631   PROTOXCLEI     SIB-I-PTX5008
SIB 7            REV 07   750-030631   PROTOXCLEI     SIB-I-PTX5008
SIB 8            REV 07   750-030631   PROTOXCLEI     SIB-I-PTX5008
Fan Tray 1       REV 04   760-030642   PROTOXCLEI     FAN-PTX-H-S

```

show chassis hardware clei-models (PTX5000 Packet Transport Router with AC PSM and PDU)

```

user@host> show chassis hardware clei-models
Hardware inventory:
Item          Version  Part number  CLEI code  FRU model number
Midplane      REV 16   750-035893  IPMUN00ARA CHAS-MP-PTX5000-S
FPM           REV 12   760-030647  IPUCA7SCAA CRAFT-PTX5000-S
PDU 0         Rev 01   740-048338  PROTOACPDU PDU2-PTX-AC-W
  PSM 0       Rev 01   740-048334  PROTOACPSM PSM2-PTX-AC
  PSM 1       Rev 01   740-048334  PROTOACPSM PSM2-PTX-AC
  PSM 2       Rev 01   740-048334  PROTOACPSM PSM2-PTX-AC
  PSM 3       Rev 01   740-048334  PROTOACPSM PSM2-PTX-AC
  PSM 4       Rev 01   740-048334  PROTOACPSM PSM2-PTX-AC
  PSM 5       Rev 01   740-048334  PROTOACPSM PSM2-PTX-AC
  PSM 6       Rev 01   740-048334  PROTOACPSM PSM2-PTX-AC
  PSM 7       Rev 01   740-048334  PROTOACPSM PSM2-PTX-AC
PDU 1         Rev 01   740-048338  PROTOACPDU PDU2-PTX-AC-W
  PSM 0       Rev 01   740-048334  PROTOACPSM PSM2-PTX-AC
  PSM 1       Rev 01   740-048334  PROTOACPSM PSM2-PTX-AC
  PSM 2       Rev 01   740-048334  PROTOACPSM PSM2-PTX-AC
  PSM 3       Rev 01   740-048334  PROTOACPSM PSM2-PTX-AC
  PSM 4       Rev 01   740-048334  PROTOACPSM PSM2-PTX-AC
  PSM 5       Rev 01   740-048334  PROTOACPSM PSM2-PTX-AC
  PSM 6       Rev 01   740-048334  PROTOACPSM PSM2-PTX-AC
  PSM 7       Rev 01   740-048334  PROTOACPSM PSM2-PTX-AC
CCG 0         REV 09   750-030653  IPUCA7DCAA CCG-PTX-S
...

```

show chassis hardware clei-models (PTX5000 Packet Transport Router with FPC2-PTX-P1A)

```

user@host> show chassis hardware clei-models
Hardware inventory:
Item          Version  Part number  CLEI code  FRU model number
Midplane      REV 11   750-035893  IPMUN00ARA CHAS-MP-PTX5000-S
FPM           REV 12   760-030647  IPUCA7SCAA CRAFT-PTX5000-S
PDU 0         Rev 04   740-048336  IPUPAL7KAA PDU2-PTX-DC-S
  PSM 0       Rev 04   740-046988  IPUPAL8KAA PSM2-PTX-DC-S
  PSM 2       Rev 04   740-046988  IPUPAL8KAA PSM2-PTX-DC-S
  PSM 4       Rev 04   740-046988  IPUPAL8KAA PSM2-PTX-DC-S
  PSM 6       Rev 04   740-046988  IPUPAL8KAA PSM2-PTX-DC-S
PDU 1         Rev 04   740-048336  IPUPAL7KAA PDU2-PTX-DC-S
  PSM 1       Rev 04   740-046988  IPUPAL8KAA PSM2-PTX-DC-S
  PSM 3       Rev 04   740-046988  IPUPAL8KAA PSM2-PTX-DC-S
  PSM 5       Rev 04   740-046988  IPUPAL8KAA PSM2-PTX-DC-S
  PSM 7       Rev 04   740-046988  IPUPAL8KAA PSM2-PTX-DC-S
CCG 0         REV 09   750-030653  IPUCA7DCAA CCG-PTX-S
CCG 1         REV 09   750-030653  IPUCA7DCAA CCG-PTX-S
...

```

show chassis hardware detail (PTX5000 Packet Transport Router)

```

user@host> show chassis hardware detail
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               JN11D1FD7AJA  PTX5000
Midplane      REV 03   711-031896  ABAC5589      Midplane-8S
FPM           REV 08   760-030647  EG1679        Front Panel Display
PDU 0         Rev 05   740-032019  ZE00006       DC Power Dist Unit
  PSM 0       Rev 05   740-032022  ZJ00018       DC 12V Power Supply
  PSM 1       Rev 04   740-032022  ZC00052       DC 12V Power Supply
  PSM 2       Rev 04   740-032022  ZD00051       DC 12V Power Supply

```


PSM 3	Rev 05	740-032022	ZJ00060	DC 12V Power Supply
CCG 0	REV 04	750-030653	EG3703	Clock Generator
CCG 1	REV 04	750-030653	EG3698	Clock Generator
Routing Engine 0	REV 05	740-026942	P737A-002231	RE-DUO-2600
ad0 3823 MB	SMART CF		201006190039C02DC02D	Compact Flash
ad1 62720 MB	SMART Lite SATA Drive		2011042300CF4C6B4C6B	Disk 1
Routing Engine 1	REV 06	740-026942	P737A-002438	RE-DUO-2600
ad0 3823 MB	SMART CF		20100619053455F055F0	Compact Flash
ad1 62720 MB	SMART Lite SATA Drive		20110423000AE8E7E8E7	Disk 1
CB 0	REV 08	750-030625	EG5519	Control Board
CB 1	REV 08	750-030625	EG5516	Control Board
FPC 0	REV 18	750-036844	EJ3080	FPC
CPU	REV 12	711-030686	EJ3260	SNG PMB
FPC 2	REV 13	750-036844	EG5065	FPC
CPU	REV 09	711-030686	EG4082	SNG PMB
PIC 0	REV 14	750-031913	EG5127	24x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	143363A00240	SFP+-10G-SR
Xcvr 1	REV 01	740-031981	UK90PZ1	SFP+-10G-LR
Xcvr 2	REV 01	740-031980	AD1141A04XH	SFP+-10G-SR
Xcvr 3	REV 01	740-031981	UK90Q46	SFP+-10G-LR
Xcvr 4	REV 01	740-031980	AD1141A04X4	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	B11H02560	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	B11C01589	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	AD1141A04XF	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	123363A01094	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	AK80LKF	SFP+-10G-SR
Xcvr 12	REV 01	740-031980	183363A01528	SFP+-10G-SR
Xcvr 14	REV 01	740-031980	193363A01079	SFP+-10G-SR
Xcvr 15	REV 01	740-031980	AK80MC8	SFP+-10G-SR
Xcvr 16	REV 01	740-031980	AJCOBHC	SFP+-10G-SR
Xcvr 19	REV 01	740-021309	J08D26856	SFP+-10G-LR
Xcvr 21	REV 01	740-031980	AK80KCT	SFP+-10G-SR
Xcvr 22	REV 01	740-031981	UK90PZL	SFP+-10G-LR
Xcvr 23	REV 01	740-031980	AK80N1V	SFP+-10G-SR
FPC 3	REV 13	750-036844	EG5074	FPC
CPU	REV 09	711-030686	EG4064	SNG PMB
PIC 1	REV 10	750-031903	EG0325	SNG Load
FPC 5	REV 06	750-036844	EH3198	FPC
CPU				
PIC 0	REV 14	750-031913	EG5134	24x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LBH	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11B03724	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80FMH	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	B11J00818	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	193363A00743	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	B11B06125	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	B11H02529	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	AK80LFB	SFP+-10G-SR
Xcvr 12	REV 01	740-031980	193363A01061	SFP+-10G-SR
Xcvr 15	REV 01	740-031980	B11J00687	SFP+-10G-SR
Xcvr 16	REV 01	740-031980	193363A00738	SFP+-10G-SR
Xcvr 18	REV 01	740-031980	AK80MQX	SFP+-10G-SR
Xcvr 19	REV 01	740-021309	J08C17257	SFP+-10G-LR
Xcvr 22	REV 01	740-031980	B11J00730	SFP+-10G-SR
Xcvr 23	REV 01	740-031980	AK80KEE	SFP+-10G-SR
PIC 1	REV 08	750-036710	EG3105	2x 40GE CFP
Xcvr 0	REV 01	740-034554	B260HLT	CFP-40G-LR4
Xcvr 1	REV 01	740-034554	B11C02847	CFP-40G-LR4
FPC 6	REV 18	750-036844	EJ4391	FPC
CPU	REV 12	711-030686	EJ3257	SNG PMB
FPC 7	REV 18	750-036844	EJ4382	FPC

CPU	REV 12	711-030686	EJ3238	SNG PMB
SPMB 0	REV 10	711-030686	EG5418	SNG PMB
SPMB 1	REV 09	711-030686	EG5373	SNG PMB
SIB 0	REV 07	750-030631	EG4858	SIB-I-8S
SIB 1	REV 07	750-030631	EG4872	SIB-I-8S
SIB 2	REV 07	750-030631	EG4866	SIB-I-8S
SIB 3	REV 07	750-030631	EG6011	SIB-I-8S
SIB 4	REV 07	750-030631	EG4907	SIB-I-8S
SIB 5	REV 07	750-030631	EG4879	SIB-I-8S
SIB 6	REV 07	750-030631	EG4864	SIB-I-8S
SIB 7	REV 07	750-030631	EG4899	SIB-I-8S
SIB 8	REV 07	750-030631	EG4880	SIB-I-8S
Fan Tray 0	REV 04	760-032784	EG1496	Vertical Fan Tray
Fan Tray 1	REV 04	760-030642	EG1335	Horizontal Fan Tray
Fan Tray 2	REV 02	760-030642	ED4952	Horizontal Fan Tray

show chassis hardware detail (PTX5000 Packet Transport Router with AC PSM and PDU)

```
user@host> show chassis hardware detail
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN12223A6AJA	PTX5000
Midplane	REV 16	750-035893	ACRA1350	Midplane-8S
FPM	REV 12	760-030647	BBBD5625	Front Panel Display
PDU 0	Rev 01	740-048338	1GB83360005	High Capacity AC WYE PDU
PSM 0	Rev 01	740-048334	1GB43360074	High Capacity AC PSM
PSM 1	Rev 01	740-048334	1GB43360001	High Capacity AC PSM
PSM 2	Rev 01	740-048334	1GB43360104	High Capacity AC PSM
PSM 3	Rev 01	740-048334	1GB43360042	High Capacity AC PSM
PSM 4	Rev 01	740-048334	1GB43360068	High Capacity AC PSM
PSM 5	Rev 01	740-048334	1GB43360080	High Capacity AC PSM
PSM 6	Rev 01	740-048334	1GB43360046	High Capacity AC PSM
PSM 7	Rev 01	740-048334	1GB43360100	High Capacity AC PSM
PDU 1	Rev 01	740-048338	1GB83360006	High Capacity AC WYE PDU
PSM 0	Rev 01	740-048334	1GB43360069	High Capacity AC PSM
PSM 1	Rev 01	740-048334	1GB43360099	High Capacity AC PSM
PSM 2	Rev 01	740-048334	1GB43360050	High Capacity AC PSM
PSM 3	Rev 01	740-048334	1GB43360095	High Capacity AC PSM
PSM 4	Rev 01	740-048334	1GB43360101	High Capacity AC PSM
PSM 5	Rev 01	740-048334	1GB43360075	High Capacity AC PSM
PSM 6	Rev 01	740-048334	1GB43360047	High Capacity AC PSM
PSM 7	Rev 01	740-048334	1GB43360019	High Capacity AC PSM
CCG 0	REV 09	750-030653	BBAZ5345	Clock Generator

show chassis hardware detail (PTX5000 Packet Transport Router with FPC2-PTX-P1A)

```
user@host> show chassis hardware detail
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN1204FC0AJA	PTX5000
Midplane	REV 11	750-035893	ACAB8038	Midplane-8S
Display	REV 12	760-030647	BBBD5619	Front Panel
PDU 0	Rev 04	740-048336	1GB93470043	High Capacity DC PDU
PSM 0	Rev 04	740-046988	1GB63500184	High Capacity DC PSM
PSM 2	Rev 04	740-046988	1GB63500169	High Capacity DC PSM
PSM 4	Rev 04	740-046988	1GB63500306	High Capacity DC PSM
PSM 6	Rev 04	740-046988	1GB63500074	High Capacity DC PSM
PDU 1	Rev 04	740-048336	1GB93470045	High Capacity DC PDU

```

PSM 1      Rev 04  740-046988  1GB63500193  High Capacity DC PSM
PSM 3      Rev 04  740-046988  1GB63500143  High Capacity DC PSM
PSM 5      Rev 04  740-046988  1GB63500146  High Capacity DC PSM
PSM 7      Rev 04  740-046988  1GB63500192  High Capacity DC PSM
CCG 0      REV 09  750-030653  BBBC1909     Clock Generator
CCG 1      REV 09  750-030653  BBBD2970     Clock Generator
...

```

show chassis hardware models (PTX5000 Packet Transport Router)

```

user@host> show chassis hardware models
Hardware inventory:
Item                Version  Part number  Serial number  FRU model number
FPM                 REV 08  760-030647  EG1679         CRAFT-PTX5000-S
PDU 0              Rev 05  740-032019  ZE00006        PWR-SAN-PDU-DC
  PSM 0            Rev 05  740-032022  ZJ00018        PSM-PTX-DC-120-S
  PSM 1            Rev 04  740-032022  ZC00052        PWR-SAN-12-DC
  PSM 2            Rev 04  740-032022  ZD00051        PWR-SAN-12-DC
  PSM 3            Rev 05  740-032022  ZJ00060        PSM-PTX-DC-120-S
CCG 0              REV 04  750-030653  EG3703         CCG-PTX-S
CCG 1              REV 04  750-030653  EG3698         CCG-PTX-S
Routing Engine 0   REV 05  740-026942  P737A-002231  RE-DUO-C2600-16G-S
Routing Engine 1   REV 06  740-026942  P737A-002438  RE-DUO-C2600-16G-S
CB 0               REV 08  750-030625  EG5519         CB-PTX-S
CB 1               REV 08  750-030625  EG5516         CB-PTX-S
FPC 0              REV 18  750-036844  EJ3080         FPC-PTX-P1-A
FPC 2              REV 13  750-036844  EG5065         FPC-PTX-P1-A
  PIC 0            REV 14  750-031913  EG5127         P1-PTX-24-10GE-SFPP
FPC 3              REV 13  750-036844  EG5074         FPC-PTX-P1-A
FPC 5
  PIC 0            REV 14  750-031913  EG5134         P1-PTX-24-10GE-SFPP
FPC 6              REV 18  750-036844  EJ4391         FPC-PTX-P1-A
FPC 7              REV 18  750-036844  EJ4382         FPC-PTX-P1-A
SIB 0              REV 07  750-030631  EG4858         SIB-I-PTX5008
SIB 1              REV 07  750-030631  EG4872         SIB-I-PTX5008
SIB 2              REV 07  750-030631  EG4866         SIB-I-PTX5008
SIB 3              REV 07  750-030631  EG6011         SIB-I-PTX5008
SIB 4              REV 07  750-030631  EG4907         SIB-I-PTX5008
SIB 5              REV 07  750-030631  EG4879         SIB-I-PTX5008
SIB 6              REV 07  750-030631  EG4864         SIB-I-PTX5008
SIB 7              REV 07  750-030631  EG4899         SIB-I-PTX5008
SIB 8              REV 07  750-030631  EG4880         SIB-I-PTX5008
Fan Tray 1         REV 04  760-030642  EG1335         FAN-PTX-H-S

```

show chassis hardware models (PTX5000 Packet Transport Router with AC PSM and PDU)

```

user@host> show chassis hardware models
Hardware inventory:
Item                Version  Part number  Serial number  FRU model number
Midplane            REV 16  750-035893  ACRA1350       CHAS-MP-PTX5000-S
FPM                 REV 12  760-030647  BBBD5625       CRAFT-PTX5000-S
PDU 0              Rev 01  740-048338  1GB83360005    PDU2-PTX-AC-W
  PSM 0            Rev 01  740-048334  1GB43360074    PSM2-PTX-AC
  PSM 1            Rev 01  740-048334  1GB43360001    PSM2-PTX-AC
  PSM 2            Rev 01  740-048334  1GB43360104    PSM2-PTX-AC
  PSM 3            Rev 01  740-048334  1GB43360042    PSM2-PTX-AC
  PSM 4            Rev 01  740-048334  1GB43360068    PSM2-PTX-AC
  PSM 5            Rev 01  740-048334  1GB43360080    PSM2-PTX-AC
  PSM 6            Rev 01  740-048334  1GB43360046    PSM2-PTX-AC
  PSM 7            Rev 01  740-048334  1GB43360100    PSM2-PTX-AC
PDU 1              Rev 01  740-048338  1GB83360006    PDU2-PTX-AC-W

```

```

PSM 0      Rev 01  740-048334  1GB43360069  PSM2-PTX-AC
PSM 1      Rev 01  740-048334  1GB43360099  PSM2-PTX-AC
PSM 2      Rev 01  740-048334  1GB43360050  PSM2-PTX-AC
PSM 3      Rev 01  740-048334  1GB43360095  PSM2-PTX-AC
PSM 4      Rev 01  740-048334  1GB43360101  PSM2-PTX-AC
PSM 5      Rev 01  740-048334  1GB43360075  PSM2-PTX-AC
PSM 6      Rev 01  740-048334  1GB43360047  PSM2-PTX-AC
PSM 7      Rev 01  740-048334  1GB43360019  PSM2-PTX-AC
CCG 0      REV 09  750-030653  BBAZ5345     CCG-PTX-S
...

```

show chassis hardware models (PTX5000 Packet Transport Router with FPC2-PTX-P1A)

```

user@host> show chassis hardware models
Hardware inventory:
Item          Version  Part number  Serial number  FRU model number
Midplane      REV 11    750-035893  ACAB8038      CHAS-MP-PTX5000-S
FPM           REV 12    760-030647  BBBD5619      CRAFT-PTX5000-S
PDU 0         Rev 04    740-048336  1GB93470043   PDU2-PTX-DC-S
  PSM 0        Rev 04    740-046988  1GB63500184   PSM2-PTX-DC-S
  PSM 2        Rev 04    740-046988  1GB63500169   PSM2-PTX-DC-S
  PSM 4        Rev 04    740-046988  1GB63500306   PSM2-PTX-DC-S
  PSM 6        Rev 04    740-046988  1GB63500074   PSM2-PTX-DC-S
PDU 1         Rev 04    740-048336  1GB93470045   PDU2-PTX-DC-S
  PSM 1        Rev 04    740-046988  1GB63500193   PSM2-PTX-DC-S
  PSM 3        Rev 04    740-046988  1GB63500143   PSM2-PTX-DC-S
  PSM 5        Rev 04    740-046988  1GB63500146   PSM2-PTX-DC-S
  PSM 7        Rev 04    740-046988  1GB63500192   PSM2-PTX-DC-S
CCG 0         REV 09    750-030653  BBBC1909      CCG-PTX-S
CCG 1         REV 09    750-030653  BBBD2970      CCG-PTX-S
...

```

show chassis hardware extensive (PTX5000 Packet Transport Router)

```

user@host> show chassis hardware extensive
Hardware inventory:
Item          Version  Part number  Serial number  Description
.....
PDU 0         Rev 04    740-032019  UE0003         DC Power Dist Unit
Jedec Code:   0x7fb0          EEPROM Version: 0x02
P/N:          740-032019      S/N:           UE0003
Assembly ID:  0x043d          Assembly Version: 04.00
Date:         11-29-2010      Assembly Flags: 0x00
Version:      Rev 04          CLEI Code:     032022XXXX
ID: DC Power Dist Unit        FRU Model Number: PWR-SAN-PDU-DC
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 3d 04 00 52 65 76 20 30 34 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 33 32 30 31 39 00 00
Address 0x20: 53 2f 4e 20 55 45 30 30 30 33 00 00 00 1d 0b 07
Address 0x30: da ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 30 33 32 30 32 32 58 58 58 58 50
Address 0x50: 57 52 2d 53 41 4e 2d 50 44 55 2d 44 43 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 a3 ff ff ff ff ff ff ff ff ff ff ff ff
PSM 0         Rev 04    740-032022  YG00065        DC 12V Power Supply
Module
Jedec Code:   0x7fb0          EEPROM Version: 0x02
P/N:          740-032022      S/N:           YG00065
Assembly ID:  0x0440          Assembly Version: 04.00

```

```

Date:          07-30-2010      Assembly Flags:  0x00
Version:       Rev 04         CLEI Code:      032022XXXX
ID: DC 12V Power Supply Module FRU Model Number:  PWR-SAN-12-DC
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 40 04 00 52 65 76 20 30 34 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 33 32 30 32 32 00 00
Address 0x20: 53 2f 4e 20 59 47 30 30 30 36 35 00 00 1e 07 07
Address 0x30: da ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 30 33 32 30 32 32 58 58 58 58 50
Address 0x50: 57 52 2d 53 41 4e 2d 31 32 2d 44 43 20 20 20 20
Address 0x60: 20 20 20 20 20 20 01 00 ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff 0c ff ff ff ff ff ff ff ff ff ff ff ff

```

show chassis hardware extensive (PTX1000 Packet Transport Router)

```

user@host> show chassis hardware extensive
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               UNDEFINED      PTX1000
Pseudo CB 0
Routing Engine 0
FPC 0          REV 06    750-053330   ACAM4850       RE-PTX1000
CPU            BUILTIN   BUILTIN      PTX1000-FPC-P2-BUILTIN
PIC 0          BUILTIN   BUILTIN      FPC CPU
                288X10GE/72X40GE/24X100GE

Xcvr 2         REV 01    740-046565   QE240845       QSFP+-40G-SR4
Xcvr 3         REV 01    740-046565   QE240962       QSFP+-40G-SR4
Xcvr 5         REV 01    740-032986   ES400LZ        QSFP+-40G-SR4
Xcvr 12        REV 01    740-054053   QE419452       QSFP+-4X10G-SR
Xcvr 18        REV 01    740-054053   QE419481       QSFP+-4X10G-SR
Xcvr 30        REV 01    740-046565   QE440485       QSFP+-40G-SR4
Xcvr 48        REV 01    740-032986   ES400K3        QSFP+-40G-SR4
Xcvr 68        REV 01    740-046565   QF2805J3       QSFP+-40G-SR4
Mezz           REV 05    711-053333   ACAM4282       Mezzanine Board
Power Supply 2 REV 01    740-054405   1EDN4470131    AC AFO 1600W PSU
Power Supply 3 REV 01    740-054405   1EDN4470112    AC AFO 1600W PSU
Fan Tray 0                                PTX1000 Fan Tray 0, Front
  to Back Airflow - AFO
Fan Tray 1                                PTX1000 Fan Tray 1, Front
  to Back Airflow - AFO
Fan Tray 2                                PTX1000 Fan Tray 2, Front
  to Back Airflow - AFO

```

show chassis hardware (MX Routers with Media Services Blade [MSB])

```

user@switch> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               JN1100FB1AFB   MX480
Midplane      REV 05    710-017414   TR3310         MX480 Midplane
FPM Board     REV 02    710-017254   KG1872         Front Panel Display
PEM 2         Rev 02    740-017343   QCS0812A00N    DC Power Entry Module
PEM 3         Rev 02    740-017343   QCS0812A00U    DC Power Entry Module
Routing Engine 0 REV 07    740-015113   1000740938     RE-S-1300
CB 0          REV 03    710-021523   KF4630         MX SCB
FPC 1         REV 11    750-037207   ZW9726         AS-MCC
CPU           REV 04    711-038173   ZW4819         AS-MCC PMB
MIC 0         REV 06    750-037214   ZW3574         AS-MSC
PIC 0         BUILTIN   BUILTIN      AS-MSC

```

```

MIC 1          REV 00  750-037211          AS-MXC
PIC 2          BUILTIN BUILTIN             AS-MXC

```

show chassis hardware extensive (MX Routers with Media Services Blade [MSB])

```

user@switch> show chassis hardware extensive
FPC 1          REV 11  750-037207  ZW9726          AS-MCC
Jedec Code:    0x7fb0          EEPROM Version:    0x02
P/N:           750-037207      S/N:             ZW9726
Assembly ID:   0x0b37          Assembly Version: 01.11
Date:          02-17-2012      Assembly Flags:   0x00
Version:       REV 11          CLEI Code:        PROTOXCLEI
ID: AS-MCC          FRU Model Number: 750-037207
Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b 37 01 0b 52 45 56 20 31 31 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 33 37 32 30 37 00 00
  Address 0x20: 53 2f 4e 20 5a 57 39 37 32 36 00 00 00 11 02 07
  Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 37
  Address 0x50: 35 30 2d 30 33 37 32 30 37 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 31 31 00 ff ff ff ff ff
  Address 0x70: ff ff ff 5e ff ff ff ff ff ff ff ff ff ff ff
CPU          REV 04  711-038173  ZW4819          AS-MCC-PMB
Jedec Code:    0x7fb0          EEPROM Version:    0x02
P/N:           711-038173      S/N:             ZW4819
Assembly ID:   0x0b38          Assembly Version: 01.04
Date:          12-30-2011      Assembly Flags:   0x00
Version:       REV 04
ID: AS-MCC PMB
Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b 38 01 04 52 45 56 20 30 34 00 00
  Address 0x10: 00 00 00 00 37 31 31 2d 30 33 38 31 37 33 00 00
  Address 0x20: 53 2f 4e 20 5a 57 34 38 31 39 00 00 00 1e 0c 07
  Address 0x30: db ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 00 50 52 4f 54 4f 58 43 4c 45 49 37
  Address 0x50: 31 31 2d 30 33 38 31 37 33 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 30 34 00 ff ff ff ff ff
  Address 0x70: ff ff ff 60 00 00 00 00 00 00 00 00 00 00 00 00
MIC 0          REV 06  750-037214  ZW3574          AS-MSC
Jedec Code:    0x7fb0          EEPROM Version:    0x02
P/N:           750-037214      S/N:             ZW3574
Assembly ID:   0x0a44          Assembly Version: 01.06
Date:          02-19-2012      Assembly Flags:   0x00
Version:       REV 06          CLEI Code:        PROTOXCLEI
ID: AS-MSC          FRU Model Number: 750-037214
Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0a 44 01 06 52 45 56 20 30 36 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 33 37 32 31 34 00 00
  Address 0x20: 53 2f 4e 20 5a 57 33 35 37 34 00 00 00 13 02 07
  Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 37
  Address 0x50: 35 30 2d 30 33 37 32 31 34 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 30 36 00 ff ff ff ff ff
  Address 0x70: ff ff ff 60 c0 03 e5 f4 00 00 00 00 00 00 00 00

```

```

PIC 0          BUILTIN      BUILTIN      AS-MSC
MIC 1          REV 00      750-037211    AS-MXC
Jedec Code:    0x7fb0      EEPROM Version: 0x01
P/N:          750-037211
Assembly ID:   0x0a43      Assembly Version: 01.00
Date:         255-255-65535 Assembly Flags: 0x00
Version:      REV 00
ID: AS-MXC
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 0a 43 01 00 52 45 56 20 30 30 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 33 37 32 31 31 00 00
Address 0x20: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ff ff ff
Address 0x30: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff c0 02 e6 6c 7f b0 02 ff 0a 44 01 06
PIC 2          BUILTIN      BUILTIN      AS-MXC

```

show chassis hardware (QFX3500 Switch running Enhanced Layer 2 Software)

```

user@switch> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis
Pseudo CB 0
Routing Engine 0      BUILTIN      BUILTIN      QFX Routing Engine
FPC 0          REV 16      750-036931    P3566-C      QFX3500-48S4Q
CPU          BUILTIN      BUILTIN      FPC CPU
PIC 0          BUILTIN      BUILTIN      48x 10G-SFP+
  Xcvr 12      REV 01      740-030658    AD1125A0438  SFP+-10G-USR
  Xcvr 13      REV 01      740-030658    AD1125A02GN  SFP+-10G-USR
PIC 1          BUILTIN      BUILTIN      4x 40G-QSFP+
PIC 2
MGMT BRD      REV 10      750-036946    BBAW0328     QFX3500-MGMT-RJ45-AFI
Power Supply 0 Rev 05      740-032091    WA13035      JPSU-650W-AC-AFI
Power Supply 1
Fan Tray 0      QFX3500 Fan Tray, Front
  to Back Airflow
Fan Tray 1      QFX3500 Fan Tray, Front
  to Back Airflow
Fan Tray 2      QFX3500 Fan Tray, Front
  to Back Airflow

```

show chassis hardware (QFX5100 Switch running Enhanced Layer 2 Software)

```

user@switch> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis
Pseudo CB 0
Routing Engine 0      BUILTIN      BUILTIN      QFX Routing Engine
FPC 0          REV 02      650-049942    TB3113280048 QFX5100-24Q-2P
CPU          BUILTIN      BUILTIN      FPC CPU
PIC 0          BUILTIN      BUILTIN      24x 40G-QSFP
  Xcvr 8        REV 01      740-032986    QA470143     QSFP+-40G-SR4
  Xcvr 14       REV 01      740-032986    QB500525     QSFP+-40G-SR4
PIC 1          REV 02      611-049555    RR3113310169 QFX-EM-4Q

```

Xcvr 0	REV 01	740-032986	QC440904	QSFP+-40G-SR4
Xcvr 1	REV 01	740-032986	QB240154	QSFP+-40G-SR4
Xcvr 2	REV 01	740-035085	018110105	QSFP+-40G-LPBK
PIC 2	REV 02	611-049555	RR3113310209	QFX-EM-4Q
Xcvr 0	REV 01	740-032986	QB190270	QSFP+-40G-SR4
Xcvr 1	REV 01	740-035085	018110063	QSFP+-40G-LPBK
Xcvr 2	REV 01	740-032986	QB210034	QSFP+-40G-SR4
Power Supply 0	REV 03	740-041741	1GA23110973	JPSU-650W-AC-AFO
Power Supply 1	REV 03	740-041741	1GA23090878	JPSU-650W-AC-AFO
Fan Tray 0				QFX5100 Fan Tray 0, Front
to Back Airflow - AFO				
Fan Tray 1				QFX5100 Fan Tray 1, Front
to Back Airflow - AFO				
Fan Tray 2				QFX5100 Fan Tray 2, Front
to Back Airflow - AFO				
Fan Tray 3				QFX5100 Fan Tray 3, Front
to Back Airflow - AFO				
Fan Tray 4				QFX5100 Fan Tray 4, Front
to Back Airflow - AFO				

show chassis lcd

List of Syntax	show chassis lcd (EX Series) on page 921 show chassis lcd (QFX Series) on page 921 show chassis lcd (OCX Series) on page 921
show chassis lcd (EX Series)	<pre>show chassis lcd <fpc-slot <i>fpc-slot-number</i>> <menu <(all-members local member <i>member-id</i>)>></pre>
show chassis lcd (QFX Series)	<pre>show chassis lcd <fpc-slot <i>fpc-slot-number</i>> <interconnect-device <i>device-id</i>> <node-device <i>device-id</i>></pre>
show chassis lcd (OCX Series)	<pre>show chassis lcd <fpc-slot <i>fpc-slot-number</i>></pre>
Release Information	<p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>menu option introduced in Junos OS Release 10.2 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 13.1 for QFabric systems.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Display the information that appears on the LCD panel of EX3200, EX3300, EX4200, EX4500, EX6200, and EX8200 switches, XRE200 External Routing Engines, QFX Series standalone switches, OCX Series switches, and Interconnect devices and Node devices within a QFabric system. Display the status of the currently selected port parameter of the Status LED for each network port on the device.</p>
Options	<p>none—Display the information that appears on the LCD panel (for any EX Series member switch in a Virtual Chassis or for XRE200 External Routing Engines, display the information for all Virtual Chassis members). Display the status of the currently selected port parameter of the Status LED for each network port.</p> <p>fpc-slot <<i>fpc-slot-number</i>>—(Optional) Display the information as follows:</p> <ul style="list-style-type: none"> (EX3200, EX3300, EX4200, and EX4500 switches, QFX Series, or OCX Series) Display the information that appears on the LCD panel for either an FPC slot with no <i>fpc-slot-number</i> value specified or for the FPC slot specified by fpc-slot 0. fpc-slot refers to the switch itself and 0 is the only valid value for <i>fpc-slot-number</i>. Output for these options is the same as for the none option. <p>Also display the status of the currently selected port parameter of the Status LED for each network port.</p> <ul style="list-style-type: none"> (EX Series Virtual Chassis member switches or XRE200 External Routing Engines) If no <i>fpc-slot-number</i> value is specified, display the information that appears on the LCD panel for all members of the Virtual Chassis. Output for this option is the same as for the none option. If the <i>fpc-slot-number</i> value is specified (it equals the <i>member-id</i> value), display the information for the specified member.

Also display the status of the currently selected port parameter of the Status LED for each network port.

- (EX6200 or EX8200 switches)—Display the information that appears on the LCD panel for the line card in the line-card slot specified by the *fpc-slot-number* value.

Also display the status of the currently selected port parameter of the Status LED for each network port.

interconnect-device *device-id*—(QFabric systems only) (Optional) Display the front panel contents and LED status of all the ports on the Interconnect device.

menu—(Optional) Display the names of the menus and menu options that are currently enabled on the LCD panel.

menu all-members—(EX Series Virtual Chassis member switches or XRE200 External Routing Engines) (Optional) Display the names of the menus and menu options that are currently enabled on the LCD panel for all Virtual Chassis members.

menu local—(EX Series Virtual Chassis member switches or XRE200 External Routing Engines) (Optional) Display the names of the menus and menu options that are currently enabled on the LCD panel for the Virtual Chassis member from which you issued the command.

menu member *member-id*—(EX Series Virtual Chassis member switches or XRE200 External Routing Engines) (Optional) Display the names of the menus and menu options that are currently enabled on the LCD panel for the specified Virtual Chassis member.

node-device *device-id*—(QFabric systems only) (Optional) Display the front panel contents and LED status of all the ports on the Node device.

Required Privilege Level

view

Related Documentation

- *LCD Panel in EX3200 Switches*
- *LCD Panel in EX4200 Switches*
- *LCD Panel in EX4500 Switches*
- *LCD Panel in an EX8200 Switch*
- *LCD Panel in an XRE200 External Routing Engine*
- *Configuring the LCD Panel on EX Series Switches (CLI Procedure)*
- *set chassis display message*

List of Sample Output

[show chassis lcd \(Two-Member EX4200 Virtual Chassis\) on page 924](#)
[show chassis lcd fpc-slot 1 \(EX4200 Virtual Chassis\) on page 925](#)
[show chassis lcd \(EX8200 Switch\) on page 925](#)
[show chassis lcd fpc-slot 2 \(EX8200 Switch\) on page 927](#)
[show chassis lcd menu \(EX4200 Switch\) on page 927](#)

[show chassis lcd menu \(EX8200 Switch\) on page 928](#)
[show chassis lcd \(QFX3500 Switches\) on page 928](#)
[show chassis lcd \(XRE200 External Routing Engine in EX8200 Virtual Chassis\) on page 928](#)
[show chassis lcd interconnect-device \(QFabric Systems\) on page 931](#)
[show chassis lcd node-device \(QFabric Systems\) on page 933](#)

Output Fields Table 61 lists the output fields for the **show chassis lcd** command. Output fields are listed in the approximate order in which they appear.

Table 61: show chassis lcd Output Fields

Field Name	Field Description
membernumber (XRE200 External Routing Engine)	Member ID of the device whose content is being displayed.
Front panel contents for slot Front panel contents (EX6200, EX8200 switch, XRE200 External Routing Engine, and QFX Series)	<p>FPC slot number of the switch whose content is being displayed. The number is always 0, except for EX4200 switches in a Virtual Chassis, where it is the member ID value.</p> <p>On EX6200 switches, EX8200 switches, and XRE200 External Routing Engines, no slot number is displayed.</p> <p>On XRE200 External Routing Engines, this field appears under the member number field for each member device in the EX8200 Virtual Chassis.</p>
LCD screen	<p>The first line displays the hostname (for Virtual Chassis members, displays the member ID, the current role, and hostname; for EX8200 switches, displays RE and the hostname). The second line displays the currently selected port parameter of the Status LED and the alarms counter. The Status LED port parameters are:</p> <ul style="list-style-type: none"> • ADM—Administrative • SPD—Speed • DPX—Duplex • POE—Power over Ethernet (EX3200 and EX4200 switches only)
LEDs status	Current state of the Alarms, System, and Master LEDs (chassis status LEDs).
Interface	Names of the interfaces on the switch.
LED (ADM/SPD/DPX/POE)	<p>State of the currently selected port parameter of the Status LED for the interface. The Status LED port parameters are:</p> <p>NOTE: The XRE200 External Routing Engine always displays the NA parameter. The QFX Series products do not have any of the port parameters listed below.</p> <ul style="list-style-type: none"> • ADM—Administrative • SPD—Speed • DPX—Duplex • NA—Not applicable. • POE—Power over Ethernet
fpcx	On standalone EX Series and QFX Series switches, always 0 . On EX Series Virtual Chassis member switches, member ID of the Virtual Chassis member whose LCD menu is displayed.

Sample Output

show chassis lcd (Two-Member EX4200 Virtual Chassis)

```

user@switch> show chassis lcd
Front panel contents for slot: 0
-----
LCD screen:
  00:BK switch1
  LED:SPD ALARM 00
LEDs status:
  Alarms LED: Off
  System LED: Green
  Master LED: Off
Interface      LED(ADM/SPD/DPX/POE)
-----
ge-0/0/0       Off
ge-0/0/1       Off
ge-0/0/2       Off
ge-0/0/3       Off
ge-0/0/4       Off
ge-0/0/5       Off
ge-0/0/6       Off
ge-0/0/7       Off
ge-0/0/8       Off
ge-0/0/9       Off
ge-0/0/10      Off
ge-0/0/11      Off
ge-0/0/12      Off
ge-0/0/13      Off
ge-0/0/14      Off
ge-0/0/15      Off
ge-0/0/16      Off
ge-0/0/17      Off
ge-0/0/18      Off
ge-0/0/19      Off
ge-0/0/20      Off
ge-0/0/21      Off
ge-0/0/22      Off
ge-0/0/23      Off
Front panel contents for slot: 1
-----
LCD screen:
  01:RE switch2
  LED:SPD ALARM 01
LEDs status:
  Alarms LED: Yellow
  System LED: Green
  Master LED: Green
Interface      LED(ADM/SPD/DPX/POE)
-----
ge-1/0/0       Off
ge-1/0/1       Off
ge-1/0/2       Off
ge-1/0/3       Off
ge-1/0/4       Off
ge-1/0/5       Off
ge-1/0/6       Off
ge-1/0/7       Off
ge-1/0/8       Off
ge-1/0/9       Off

```

```

ge-1/0/10      Off
ge-1/0/11      Off
ge-1/0/12      Off
ge-1/0/13      Off
ge-1/0/14      Off
ge-1/0/15      Off
ge-1/0/16      Off
ge-1/0/17      Off
ge-1/0/18      Off
ge-1/0/19      Off
ge-1/0/20      Off
ge-1/0/21      Off
ge-1/0/22      Off
ge-1/0/23      Off

```

The output for the **show chassis lcd fpc-slot** command is the same as the output for the **show chassis lcd** command.

show chassis lcd fpc-slot 1 (EX4200 Virtual Chassis)

```

user@switch> show chassis lcd fpc-slot 1
Front panel contents for slot: 1
-----
LCD screen:
  01:RE switch2
  LED:SPD ALARM 01
LEDs status:
  Alarms LED: Yellow
  System LED: Green
  Master LED: Green
Interface      LED (ADM/SPD/DPX/POE)
-----
ge-1/0/0      Off
ge-1/0/1      Off
ge-1/0/2      Off
ge-1/0/3      Off
ge-1/0/4      Off
ge-1/0/5      Off
ge-1/0/6      Off
ge-1/0/7      Off
ge-1/0/8      Off
ge-1/0/9      Off
ge-1/0/10     Off
ge-1/0/11     Off
ge-1/0/12     Off
ge-1/0/13     Off
ge-1/0/14     Off
ge-1/0/15     Off
ge-1/0/16     Off
ge-1/0/17     Off
ge-1/0/18     Off
ge-1/0/19     Off
ge-1/0/20     Off
ge-1/0/21     Off
ge-1/0/22     Off
ge-1/0/23     Off

```

show chassis lcd (EX8200 Switch)

```

user@switch> show chassis lcd

```

Front panel contents:

LCD screen:

RE st-8200-r

LED:ADM ALARM 01

LEDs status:

Alarms LED: Yellow

System LED: Yellow

Master LED: Green

Interface LED(ADM/SPD/DPX)

ge-0/0/0	Off
ge-0/0/1	Off
ge-0/0/2	Off
ge-0/0/3	Off
ge-0/0/4	Off
ge-0/0/5	Off
ge-0/0/6	Off
ge-0/0/7	Off
ge-0/0/8	Off
ge-0/0/9	Off
ge-0/0/10	Off
ge-0/0/11	Off
ge-0/0/12	Off
ge-0/0/13	Off
ge-0/0/14	Off
ge-0/0/15	Off
ge-0/0/16	Off
ge-0/0/17	Off
ge-0/0/18	Off
ge-0/0/19	Off
ge-0/0/20	Off
ge-0/0/21	Off
ge-0/0/22	Off
ge-0/0/23	Off
ge-0/0/24	Off
ge-0/0/25	Off
ge-0/0/26	Off
ge-0/0/27	Off
ge-0/0/28	Off
ge-0/0/29	Off
ge-0/0/30	Off
ge-0/0/31	Off
ge-0/0/32	Off
ge-0/0/33	Off
ge-0/0/34	Off
ge-0/0/35	Off
ge-0/0/36	Off
ge-0/0/37	Off
ge-0/0/38	Off
ge-0/0/39	Off
ge-0/0/40	Off
ge-0/0/41	Off
ge-0/0/42	Off
ge-0/0/43	Off
ge-0/0/44	Off
ge-0/0/45	Off
ge-0/0/46	Off
ge-0/0/47	Off
xe-2/0/0	Off
xe-2/0/1	Off

xe-2/0/2	Off
xe-2/0/3	Off
xe-2/0/4	Off
xe-2/0/5	Off
xe-2/0/6	Off
xe-2/0/7	Off
xe-3/0/0	Off
xe-3/0/1	Off
xe-3/0/2	Off
xe-3/0/3	Off
xe-3/0/4	Off
xe-3/0/5	Off
xe-3/0/6	Off
xe-3/0/7	Off
xe-5/0/0	Off
xe-5/0/1	Off
xe-5/0/2	Off
xe-5/0/3	Off
xe-5/0/4	Off
xe-5/0/5	Off
xe-5/0/6	On
xe-5/0/7	On
xe-7/0/5	Off

show chassis lcd fpc-slot 2 (EX8200 Switch)

```
show chassis lcd fpc-slot 2
```

Interface	LED (ADM/SPD/DPX)
xe-2/0/0	Off
xe-2/0/1	Off
xe-2/0/2	Off
xe-2/0/3	Off
xe-2/0/4	Off
xe-2/0/5	Off
xe-2/0/6	Off
xe-2/0/7	Off

show chassis lcd menu (EX4200 Switch)

```
user@switch> show chassis lcd menu
fpc0:
```

```
-----
status-menu
status-menu vcp-status
status-menu power-status
status-menu environ-menu
status-menu show-version
maintenance-menu
maintenance-menu halt-menu
maintenance-menu system-reboot
maintenance-menu rescue-config
maintenance-menu vc-uplink-config
maintenance-menu factory-default
```

On an EX4200 switch in a Virtual Chassis, the output for the **show chassis lcd menu** **all-members** command is the same as the output for the **show chassis lcd menu** command.

show chassis lcd menu (EX8200 Switch)

```
user@switch> show chassis lcd menu
status-menu
status-menu sf-status1-menu
status-menu sf-status2-menu
status-menu psu-status1-menu
status-menu psu-status2-menu
status-menu environ-menu
status-menu show-version
maintenance-menu
maintenance-menu halt-menu
maintenance-menu system-reboot
maintenance-menu rescue-config
maintenance-menu factory-default
```

show chassis lcd (QFX3500 Switches)

```
user@switch> show chassis lcd
Front panel contents for slot: 0
-----
LCD screen:
00:RE switch
ALARM 01
LEDs status:
Status/Beacon LED: Yellow Blinking
Interface STATUS LED ACTIVITY LED
-----
fte-0/1/0 Off Off
```

show chassis lcd (XRE200 External Routing Engine in EX8200 Virtual Chassis)

```
user@external-routing-engine> show chassis lcd
member0:
-----
Front panel contents:
-----
LCD screen:
  RE ex8200-member0
  LED:ADM ALARM 04
LEDs status:
  Alarms LED: Red
  System LED: Yellow
  Master LED: Green

member1:
-----

member8:
-----
Front panel contents:
-----
LCD screen:
  BACKUP

member9:
-----
Front panel contents:
-----
LCD screen:
  09:RE xre200-member9
```


LED: NA ALARM 01

Interface	LED(ADM/SPD/DPX/POE)
-----------	----------------------

ge-0/0/0	On
ge-0/0/1	On
ge-0/0/2	On
ge-0/0/3	On
ge-0/0/4	Off
ge-0/0/5	Off
ge-0/0/6	Off
ge-0/0/7	Off
ge-0/0/8	Off
ge-0/0/9	Off
ge-0/0/10	On
ge-0/0/11	Off
ge-0/0/12	Off
ge-0/0/13	Off
ge-0/0/14	Off
ge-0/0/15	Off
ge-0/0/16	Off
ge-0/0/17	Off
ge-0/0/18	Off
ge-0/0/19	Off
ge-0/0/20	Off
ge-0/0/21	Off
ge-0/0/22	Off
ge-0/0/23	Off
ge-0/0/24	Off
ge-0/0/25	Off
ge-0/0/26	Off
ge-0/0/27	Off
ge-0/0/28	Off
ge-0/0/29	Off
ge-0/0/30	Off
ge-0/0/31	Off
ge-0/0/32	Off
ge-0/0/33	Off
ge-0/0/34	Off
ge-0/0/35	Off
ge-0/0/36	Off
ge-0/0/37	Off
ge-0/0/38	Off
ge-0/0/39	Off
ge-0/0/40	On
ge-0/0/41	On
ge-0/0/42	On
ge-0/0/43	On
ge-0/0/44	On
ge-0/0/45	On
ge-0/0/46	On
ge-0/0/47	On
ge-16/0/0	On
ge-16/0/1	Off
ge-16/0/2	On
ge-16/0/3	Off
ge-16/0/4	On
ge-16/0/5	Off
ge-16/0/6	On
ge-16/0/7	Off
ge-16/0/8	Off
ge-16/0/9	Off

ge-16/0/10	Off
ge-16/0/11	Off
ge-16/0/12	Off
ge-16/0/13	On
ge-16/0/14	Off
ge-16/0/15	On
ge-16/0/16	Off
ge-16/0/17	On
ge-16/0/18	On
ge-16/0/19	On
ge-16/0/20	On
ge-16/0/21	Off
ge-16/0/22	On
ge-16/0/23	Off
ge-16/0/24	Off
ge-16/0/25	Off
ge-16/0/26	On
ge-16/0/27	Off
ge-16/0/28	Off
ge-16/0/29	Off
ge-16/0/30	On
ge-16/0/31	Off
ge-16/0/32	On
ge-16/0/33	On
ge-16/0/34	On
ge-16/0/35	Off
ge-16/0/36	On
ge-16/0/37	Off
ge-16/0/38	Off
ge-16/0/39	Off
ge-16/0/40	Off
ge-16/0/41	Off
ge-16/0/42	On
ge-16/0/43	Off
ge-16/0/44	Off
ge-16/0/45	Off
ge-16/0/46	Off
ge-16/0/47	Off
xe-19/0/0	Off
xe-19/0/1	On
xe-19/0/2	On
xe-19/0/3	On
xe-19/0/4	On
xe-19/0/5	On
ge-22/0/0	Off
ge-22/0/1	Off
ge-22/0/2	On
ge-22/0/3	Off
ge-22/0/4	On
ge-22/0/5	On
ge-22/0/6	On
ge-22/0/7	On
ge-22/0/8	Off
ge-22/0/9	Off
ge-22/0/10	Off
ge-22/0/11	Off
ge-22/0/12	Off
ge-22/0/13	Off
ge-22/0/14	Off
ge-22/0/15	Off
ge-22/0/16	On

```

ge-22/0/17      Off
ge-22/0/18      On
ge-22/0/19      Off
ge-22/0/20      On
ge-22/0/21      Off
ge-22/0/22      On
ge-22/0/23      Off
ge-22/0/24      On
ge-22/0/25      Off
ge-22/0/26      Off
ge-22/0/27      Off
ge-22/0/28      Off
ge-22/0/29      Off
ge-22/0/30      Off
ge-22/0/31      Off
ge-22/0/32      On
ge-22/0/33      Off
ge-22/0/34      On
ge-22/0/35      Off
ge-22/0/36      Off
ge-22/0/37      Off
ge-22/0/38      Off
ge-22/0/39      Off
ge-22/0/40      Off
ge-22/0/41      Off
ge-22/0/42      Off
ge-22/0/43      Off
ge-22/0/44      Off
ge-22/0/45      Off
ge-22/0/46      Off
ge-22/0/47      Off

```

show chassis lcd interconnect-device (QFabric Systems)

```

show chassis lcd interconnect-device IC-F1012
      Front Panel Module Information
      -----
      LCD screen:
      IC-F1012      3 Alarms active

LEDs status:
  Status LED: Green
  Power LED : Green
  Major Alarm LED: off
  Minor Alarm LED: Yellow
  Fan 0 LED : Green
  Fan 1 LED : Green
  Fan 2 LED : Green
  Fan 3 LED : Green
  Fan 4 LED : Green
  Fan 5 LED : Green
  Fan 6 LED : Green
  Fan 7 LED : Green
  Fan 8 LED : Green
  Fan 9 LED : Green
  PEM 0 LED : Green
  PEM 1 LED : Green
  PEM 2 LED : Green
  PEM 3 LED : off
  PEM 4 LED : off
  PEM 5 LED : off

```

LED info for: CB - 0

LEDs status:

Status LED: Green
Mastership LED: Green

Interface	STATUS LED	LINK/ACTIVITY LED
IC-F1012:pme0 :	Green	N/A
IC-F1012:pme1 :	Green	N/A
IC-F1012:pme2 :	off	N/A
IC-F1012:pme3 :	off	N/A

LED info for: CB - 1

LEDs status:

Status LED: Green
Mastership LED: Amber

Interface	STATUS LED	LINK/ACTIVITY LED
IC-F1012:pme0 :	Green	N/A
IC-F1012:pme1 :	Green	N/A
IC-F1012:pme2 :	off	N/A
IC-F1012:pme3 :	off	N/A

LED info for: FC 0 FPC - 0

LEDs status:

Status LED: Green

Interface	STATUS LED	LINK/ACTIVITY LED
IC-F1012:fte-0/0/0	Green	N/A
IC-F1012:fte-0/0/1	Green	N/A
IC-F1012:fte-0/0/2	Green	N/A
IC-F1012:fte-0/0/3	Green	N/A
IC-F1012:fte-0/0/4	Green	N/A

LED info for: FC 1 FPC - 1

LEDs status:

Status LED: Green

Interface	STATUS LED	LINK/ACTIVITY LED
IC-F1012:fte-1/0/0	Green	N/A
IC-F1012:fte-1/0/1	Green	N/A
IC-F1012:fte-1/0/2	Green	N/A
IC-F1012:fte-1/0/3	Green	N/A
IC-F1012:fte-1/0/4	Green	N/A

LED info for: RC 0 FPC - 8

LEDs status:

Status LED: Green

LED info for: RC 1 FPC - 9

LEDs status:

Status LED: Green

```

LED info for: RC 2 FPC - 10
-----
LEDs status:
  Status LED: Green

LED info for: RC 3 FPC - 11
-----
LEDs status:
  Status LED: Green

LED info for: RC 4 FPC - 12
-----
LEDs status:
  Status LED: Green

LED info for: RC 5 FPC - 13
-----
LEDs status:
  Status LED: Green

LED info for: RC 6 FPC - 14
-----
LEDs status:
  Status LED: Green

LED info for: RC 7 FPC - 15
-----
LEDs status:
  Status LED: Green

```

show chassis lcd node-device (QFabric Systems)

```

show chassis lcd node-device P3774-C
  Front panel contents for: P3774-C
  -----
  LCD screen:
  P3774-C

LEDs status:
  Status/Beacon LED: Yellow Blinking

```

Interface	STATUS LED	LINK/ACTIVITY LED
P3774-C:xe-0/0/6	Green	Green
P3774-C:xe-0/0/7	Green	Green
P3774-C:ge-0/0/10	Green	Green
P3774-C:ge-0/0/11	Green	Green Blinking
P3774-C:ge-0/0/12	Green	Off
P3774-C:ge-0/0/13	Green	Green Blinking
P3774-C:ge-0/0/20	Green	Green
P3774-C:ge-0/0/21	Green	Green
P3774-C:ge-0/0/22	Green	Green Blinking
P3774-C:ge-0/0/23	Green	Off
P3774-C:ge-0/0/30	Green	Green
P3774-C:ge-0/0/31	Green	Green
P3774-C:ge-0/0/32	Green	Green Blinking
P3774-C:ge-0/0/33	Green	Green Blinking
P3774-C:fte-0/1/0	Green	Green
P3774-C:fte-0/1/1	Green	Green Blinking
P3774-C:fte-0/1/2	Green	Green Blinking
P3774-C:fte-0/1/3	Green	Green

show chassis led

List of Syntax	show chassis led (EX Series) on page 935 show chassis led (QFX Series) on page 935 Syntax (OCX Series) on page 935
show chassis led (EX Series)	<pre>show chassis led <fpc-slot <fpc-slot-number>></pre>
show chassis led (QFX Series)	<pre>show chassis led <fpc-slot <fpc-slot-number>> interconnect-device name node-device name</pre>
Syntax (OCX Series)	<pre>show chassis led <fpc-slot <fpc-slot-number>></pre>
Release Information	<p>Command introduced in Junos OS Release 10.1 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Display the status and colors of the chassis LEDs on the front panel of the switch. A major alarm (red) indicates a critical error condition that requires immediate action. A minor alarm (yellow) indicates a noncritical condition that requires monitoring or maintenance. A minor alarm that is left unchecked might cause interruption in service or performance degradation.</p>
Options	<p>none—Display the status of the chassis status LEDs (for EX4200 switches configured as a Virtual Chassis, display the information for all Virtual Chassis members).</p> <p>fpc-slot <fpc-slot-number>—(Optional) (Not on EX2200 switches) Display the information as follows:</p> <ul style="list-style-type: none"> (EX3200, standalone EX4200, standalone QFX3500, EX4500, and OCX Series switches) Display the status of the chassis status LEDs for either an FPC slot with no fpc-slot-number value specified or for the FPC slot specified by fpc-slot 0. fpc-slot refers to the switch itself and 0 is the only valid value for fpc-slot-number. Output for these options is the same as for the none option. (EX4200 switches in a Virtual Chassis with two or more members) If no fpc-slot-number value is specified, display the status of the chassis status LEDs for all members of the Virtual Chassis. Output for this option is the same as for the none option. If the fpc-slot-number value is specified (it equals the member-id value), display the status of the chassis status LEDs for the specified member. (EX8200 switches)—Display the status of the chassis status LEDs for the line card in the line-card slot specified by the fpc-slot-number value. <p>interconnect-device name—</p> <p>— (QFabric systems only) (Optional) Display the status of the chassis and interface status LEDs for the Interconnect device.</p>

node-device name— (QFabric systems only) (Optional) Display the status of the chassis and interface status LEDs for the Node device.

Required Privilege Level view

Related Documentation

- *Chassis Status LEDs in EX2200 Switches*
- *Chassis Status LEDs in EX3200 Switches*
- *Chassis Status LEDs in EX4200 Switches*
- *Chassis Status LEDs in EX4500 Switches*
- *Chassis Status LEDs in an EX8200 Switch*
- *Chassis Status LEDs on a QFX3500 Device*
- *Chassis Status LEDs in the QFX3600 and QFX3600-I Device*
- *Management Port LEDs on a QFX3500 Device*
- *Management Port LEDs in the QFX3600 and QFX3600-I Device*
- *Chassis Status LEDs on a QFX3008-I Interconnect Device*
- *Control Board LEDs on a QFX3008-I Interconnect Device*

List of Sample Output

[show chassis led \(EX2200 Switch\) on page 939](#)
[show chassis led on page 940](#)
[show chassis led fpc-slot 0 on page 941](#)
[show chassis led \(EX Series\) on page 941](#)
[show chassis led node-device \(QFabric System Node Device\) on page 942](#)
[show chassis led interconnect-device \(QFabric System - QFX3600-I Interconnect Device\) on page 942](#)
[show chassis led interconnect-device \(QFabric System - QFX3008-I Interconnect Device\) on page 943](#)

Output Fields [Table 45](#) lists the output fields for the **show chassis led** command. Output fields are listed in the approximate order in which they appear.

Table 62: show chassis led Output Fields

Field Name	Field Description
Front panel contents for slot	FPC slot number of the device whose content is being displayed. The number is always 0, except for EX4200 switches in a Virtual Chassis, where it is the member ID value.
Front panel contents (EX8200 Switches)	
Front Panel Module Information (QFabric system QFX3008-I Interconnect device)	On EX8200 switches, no slot number is displayed.
Front panel contents for (QFabric system Node devices and QFX3600-I Interconnect devices)	On QFabric system Node devices, the name of the Node device whose content is being displayed.

Table 62: show chassis led Output Fields (*continued*)

Field Name	Field Description
Alarms LED	<p>(EX Series switches only) Displays status of the ALM LED:</p> <ul style="list-style-type: none"> • Off—No alarm has been configured. • Green—No alarm has been triggered. • Red—Major alarm. • Yellow—Minor alarm
System LED	<p>(EX Series switches only) Displays status of the SYS LED:</p> <ul style="list-style-type: none"> • Off—Switch is powered off. • Green—Switch is operating normally. • Yellow—Switch is booting.
Master LED:	<p>Displays status of the MST LED (on EX3200, EX4200, and EX8200 switches):</p> <ul style="list-style-type: none"> • Green—On an EX4200 Virtual Chassis switch, indicates the switch is the master in the Virtual Chassis configuration. On other switches, indicates that the Routing Engine is operational. • Off <ul style="list-style-type: none"> • On an EX4200 Virtual Chassis switch, indicates that this switch is not the master in the Virtual Chassis configuration. • On EX3200, standalone EX4200, and EX8200 switches, indicates that the Routing Engine is not operational.
Mode LED:	<p>(EX Series switches only) On an EX2200 switch only, displays the currently selected port parameter of the Status LED:</p> <ul style="list-style-type: none"> • ADM—Administrative • SPD—Speed • DPX—Duplex • POE—Power over Ethernet
Status/Beacon LED	<p>(QFX Series and OCX Series) Displays the system status as indicated by the Status LED on the chassis. For more information, see:</p> <ul style="list-style-type: none"> • <i>Chassis Status LEDs on a QFX3500 Device</i> • <i>Chassis Status LEDs in the QFX3600 and QFX3600-I Device</i>
LINK/SPEED LED	<p>(QFX Series and OCX Series) Displays the link status and speed of a management port. For more information, see:</p> <ul style="list-style-type: none"> • <i>Management Port LEDs on a QFX3500 Device</i> • <i>Management Port LEDs in the QFX3600 and QFX3600-I Device</i>
ACTIVITY LED	<p>(QFX Series and OCX Series) Displays the activity status of a management port. For more information, see:</p> <ul style="list-style-type: none"> • <i>Management Port LEDs on a QFX3500 Device</i> • <i>Management Port LEDs in the QFX3600 and QFX3600-I Device</i>

Table 62: show chassis led Output Fields (*continued*)

Field Name	Field Description
STATUS LED	<p>(QFX Series and OCX Series) Displays the link status of an interface as indicated by the ST LED. For more information, see:</p> <ul style="list-style-type: none"> Control Board LEDs on a QFX3008-I Interconnect Device Access Port and Uplink Port LEDs on a QFX3500 Device Access Port and Uplink Port LEDs on a QFX3600 or QFX3600-I Device
LINK/ACTIVITY LED	<p>(QFX Series and OCX Series) Displays link activity or faults on an interface as indicated by the LA LED. For more information, see:</p> <ul style="list-style-type: none"> Access Port and Uplink Port LEDs on a QFX3500 Device Access Port and Uplink Port LEDs on a QFX3600 or QFX3600-I Device
Status LED	<p>(QFX3008-I Interconnect device only)</p> <ul style="list-style-type: none"> Displays the system status as indicated by the STATUS LED on the front panel of the chassis. For more information, see <i>Chassis Status LEDs on a QFX3008-I Interconnect Device</i>. Displays the status of a Control Board as indicated by the STATUS LED on the Control Board. For more information, see <i>Control Board LEDs on a QFX3008-I Interconnect Device</i>.
Power LED	<p>(QFX3008-I Interconnect device only) Displays the status of system power on the device. For more information, see <i>Chassis Status LEDs on a QFX3008-I Interconnect Device</i>.</p>
Major Alarm LED	<p>(QFX3008-I Interconnect device only) Displays whether a critical error condition that requires immediate action exists on the device. For more information, see <i>Chassis Status LEDs on a QFX3008-I Interconnect Device</i>.</p>
Minor Alarm LED	<p>(QFX3008-I Interconnect device only) Displays whether a noncritical condition that requires monitoring or maintenance exists on the device. For more information, see <i>Chassis Status LEDs on a QFX3008-I Interconnect Device</i>.</p>
Fan 0 LED	<p>(QFX3008-I Interconnect device only) Displays the status of fan trays on the device. For more information, see <i>Chassis Status LEDs on a QFX3008-I Interconnect Device</i>.</p>
Fan 1 LED	
Fan 2 LED	
Fan 3 LED	
Fan 4 LED	
Fan 5 LED	
Fan 6 LED	
Fan 7 LED	
Fan 8 LED	

Table 62: show chassis led Output Fields (*continued*)

Field Name	Field Description
PEM 0 LED	(QFX3008-I Interconnect device only) Displays the status of power supplies on the device. For more information, see <i>Chassis Status LEDs on a QFX3008-I Interconnect Device</i> .
PEM 1 LED	
PEM 2 LED	
PEM 3 LED	
PEM 4 LED	
LED info for	(QFX3008-I Interconnect device only) Displays the LED information for a Control Board.
Mastership LED	(QFX3008-I Interconnect device only) Displays status of the MASTER LED on a Control Board. For more information, see <i>Control Board LEDs on a QFX3008-I Interconnect Device</i> .
Interface	Names of the interfaces on the device.
LED (ADM/SPD/DPX/POE)	<p>(EX Series switches only) State of the currently selected port parameter of the Status LED for the interface. The Status LED port parameters are:</p> <p>NOTE: EX4500 and EX8200 switches do not have the POE port parameter.</p> <ul style="list-style-type: none"> • ADM—Administrative • SPD—Speed • DPX—Duplex • POE—Power over Ethernet

Sample Output

show chassis led (EX2200 Switch)

```

user@switch> show chassis led
Front panel contents for slot: 0
-----
LEDs status:
  Alarms LED: Amber
  System LED: Green
  Mode LED : Duplex
Interface    LED(ADM/SPD/DPX/POE)
-----
ge-0/0/0      Off
ge-0/0/1      Full Duplex
ge-0/0/2      Full Duplex
ge-0/0/3      Off
ge-0/0/4      Off
ge-0/0/5      Full Duplex
ge-0/0/6      Full Duplex
ge-0/0/7      Full Duplex
ge-0/0/8      Full Duplex
ge-0/0/9      Full Duplex
ge-0/0/10     Full Duplex
ge-0/0/11     Full Duplex

```

```

ge-0/0/12      Full Duplex
ge-0/0/13      Full Duplex
ge-0/0/14      Full Duplex
ge-0/0/15      Full Duplex
ge-0/0/16      Full Duplex
ge-0/0/17      Full Duplex
ge-0/0/18      Full Duplex
ge-0/0/19      Full Duplex
ge-0/0/20      Full Duplex
ge-0/0/21      Full Duplex
ge-0/0/22      Off
ge-0/0/23      Off
ge-0/0/24      Full Duplex
ge-0/0/25      Full Duplex
ge-0/0/26      Off
ge-0/0/27      Off
ge-0/0/28      Full Duplex
ge-0/0/29      Full Duplex

```

show chassis led

```
user@switch> show chassis led
```

```
Front panel contents for slot: 0
```

```
-----
LEDs status:
```

```
  Alarms LED: Off
```

```
  System LED: Green
```

```
  Master LED: Green
```

```
Interface      LED (ADM/SPD/DPX/POE)
```

```
-----
ge-0/0/0      Off
ge-0/0/1      Full Duplex
ge-0/0/2      Full Duplex
ge-0/0/3      Off
ge-0/0/4      Off
ge-0/0/5      Full Duplex
ge-0/0/6      Full Duplex
ge-0/0/7      Full Duplex
ge-0/0/8      Full Duplex
ge-0/0/9      Full Duplex
ge-0/0/10     Full Duplex
ge-0/0/11     Full Duplex
ge-0/0/12     Full Duplex
ge-0/0/13     Full Duplex
ge-0/0/14     Full Duplex
ge-0/0/15     Full Duplex
ge-0/0/16     Full Duplex
ge-0/0/17     Full Duplex
ge-0/0/18     Full Duplex
ge-0/0/19     Full Duplex
ge-0/0/20     Full Duplex
ge-0/0/21     Full Duplex
ge-0/0/22     Off
ge-0/0/23     Off
ge-0/0/24     Full Duplex
ge-0/0/25     Full Duplex
ge-0/0/26     Off
ge-0/0/27     Off
ge-0/0/28     Full Duplex
ge-0/0/29     Full Duplex

```

show chassis led fpc-slot 0

```

user@switch> show chassis led fpc-slot 0
Front panel contents for slot: 0
-----
LEDs status:
  Alarms LED: Red
  System LED: Green
  Master LED: Green
Interface      LED(ADM/SPD/DPX/POE)
-----
ge-0/0/0      Off
ge-0/0/1      Off
ge-0/0/2      Off
ge-0/0/3      Off
ge-0/0/4      Off
ge-0/0/5      Off
ge-0/0/6      Off
ge-0/0/7      Off
ge-0/0/8      Off
ge-0/0/9      Off
ge-0/0/10     Off
ge-0/0/11     Off
ge-0/0/12     Off
ge-0/0/13     Off
ge-0/0/14     Off
ge-0/0/15     Off
ge-0/0/16     Off
ge-0/0/17     Off
ge-0/0/18     Off
ge-0/0/19     Off
ge-0/0/20     Off
ge-0/0/21     Off
ge-0/0/22     Off
ge-0/0/23     Off

```

show chassis led (EX Series)

```

user@switch> show chassis led
Front panel contents for slot: 0
-----
LEDs status:
  Alarms LED: Amber
  Status LED: Green
  Mode LED : Duplex
Interface LED(ADM/SPD/DPX/POE)
-----
ge-0/0/0 Off
ge-0/0/1 Full Duplex
ge-0/0/2 Full Duplex
ge-0/0/3 Off
ge-0/0/4 Off
ge-0/0/5 Full Duplex
ge-0/0/6 Full Duplex
ge-0/0/7 Full Duplex
ge-0/0/8 Full Duplex
ge-0/0/9 Full Duplex
ge-0/0/10 Full Duplex
ge-0/0/11 Full Duplex
ge-0/0/12 Full Duplex
ge-0/0/13 Full Duplex

```

```

ge-0/0/14 Full Duplex
ge-0/0/15 Full Duplex
ge-0/0/16 Full Duplex
ge-0/0/17 Full Duplex
ge-0/0/18 Full Duplex
ge-0/0/19 Full Duplex
ge-0/0/20 Full Duplex
ge-0/0/21 Full Duplex
ge-0/0/22 Off
ge-0/0/23 Off
ge-0/0/24 Full Duplex
ge-0/0/25 Full Duplex
ge-0/0/26 Off
ge-0/0/27 Off
ge-0/0/28 Full Duplex
ge-0/0/29 Full Duplex

```

show chassis led node-device (QFabric System Node Device)

```

user@switch> show chassis led node-device node1
Front panel contents for: node1
LEDs status:
  Status/Beacon LED: Yellow Blinking

```

Interface	LINK/SPEED LED	ACTIVITY LED
node1:me5	Green	N/A
node1:me6	Green	N/A

Interface	STATUS LED	LINK/ACTIVITY LED
node1:xe-0/0/8	Green	Green
node1:ge-0/0/10	Green	Green
node1:ge-0/0/12	Green	Green
node1:ge-0/0/24	Green	Green
node1:ge-0/0/25	Green	Green
node1:ge-0/0/26	Green	Green
node1:ge-0/0/27	Green	Green
node1:ge-0/0/28	Green	Green
node1:ge-0/0/29	Green	Green
node1:ge-0/0/30	Green	Green
node1:ge-0/0/31	Green	Green
node1:ge-0/0/32	Green	Green
node1:ge-0/0/33	Green	Green
node1:ge-0/0/34	Green	Green
node1:ge-0/0/35	Green	Green
node1:ge-0/0/36	Green	Green
node1:ge-0/0/37	Green	Green
node1:ge-0/0/38	Green	Green
node1:ge-0/0/39	Green	Green
node1:fte-0/1/0	Green	Green Blinking
node1:fte-0/1/2	Green	Green Blinking

show chassis led interconnect-device (QFabric System - QFX3600-I Interconnect Device)

```

user@switch> show chassis led interconnect-device IC-EG0712
Front panel contents for: FPC 0
-----
LEDs status:
  Status/Beacon LED: Yellow Blinking

```

Interface	LINK/SPEED LED	ACTIVITY LED
IC-EG0712:me5	Green	N/A
IC-EG0712:me6	Green	N/A

Interface	STATUS LED	LINK/ACTIVITY LED
IC-EG0712:fte-0/1/0	Green	Green
IC-EG0712:fte-0/1/1	Green	Green Blinking
IC-EG0712:fte-0/1/2	Green	Green
IC-EG0712:fte-0/1/3	Green	Green Blinking
IC-EG0712:fte-0/1/4	Green	Green
IC-EG0712:fte-0/1/5	Green	Green Blinking
IC-EG0712:fte-0/1/6	Green	Green
IC-EG0712:fte-0/1/7	Green	Green
IC-EG0712:fte-0/1/8	Green	Green Blinking
IC-EG0712:fte-0/1/9	Green	Green Blinking
IC-EG0712:fte-0/1/10	Green	Green Blinking

show chassis led interconnect-device (QFabric System - QFX3008-I Interconnect Device)

```

user@switch> show chassis led interconnect-device IC-EG0712
Front Panel Module Information
-----
LEDs status:
  Status LED: Green
  Power LED : Yellow Blinking
  Major Alarm LED: Red
  Minor Alarm LED: Yellow
  Fan 0 LED : Green
  Fan 1 LED : Green
  Fan 2 LED : Green
  Fan 3 LED : Green
  Fan 4 LED : Green
  Fan 5 LED : Green
  Fan 6 LED : Green
  Fan 7 LED : Green
  Fan 8 LED : Green
  Fan 9 LED : Green
  PEM 0 LED : Green
  PEM 1 LED : Green
  PEM 2 LED : Green
  PEM 3 LED : off
  PEM 4 LED : Yellow Blinking
  PEM 5 LED : off

  LED info for: CB - 0
  -----
LEDs status:
  Status LED: Green
  Mastership LED: Green

Interface          STATUS LED    LINK/ACTIVITY LED
-----
IC-F4899:pme0 :    Green          N/A
IC-F4899:pme1 :    off            N/A
IC-F4899:pme2 :    off            N/A
IC-F4899:pme3 :    off            N/A

  LED info for: CB - 1
  -----

```

LEDs status:

Status LED: Green

Mastership LED: Amber

Interface	STATUS LED	LINK/ACTIVITY LED
IC-F4899:pme0 :	Green	N/A
IC-F4899:pme1 :	off	N/A
IC-F4899:pme2 :	off	N/A
IC-F4899:pme3 :	off	N/A

LED info for: FC 0 FPC - 0

LEDs status:

Status LED: Green

Interface	STATUS LED	LINK/ACTIVITY LED
IC-F4899:fte-0/0/0	Green	N/A
IC-F4899:fte-0/0/1	Green	N/A
IC-F4899:fte-0/0/2	Green	N/A
IC-F4899:fte-0/0/3	Green	N/A
IC-F4899:fte-0/0/4	Green	N/A
IC-F4899:fte-0/0/5	Green	N/A
IC-F4899:fte-0/0/6	Green	N/A
IC-F4899:fte-0/0/7	Green	N/A
IC-F4899:fte-0/0/8	Green	N/A
IC-F4899:fte-0/0/9	Green	N/A
IC-F4899:fte-0/0/10	Green	N/A
IC-F4899:fte-0/0/11	Green	N/A
IC-F4899:fte-0/0/12	Green	N/A
IC-F4899:fte-0/0/13	Green	N/A
IC-F4899:fte-0/0/14	Green	N/A
IC-F4899:fte-0/0/15	Green	N/A

LED info for: FC 1 FPC - 1

LEDs status:

Status LED: Green

Interface	STATUS LED	LINK/ACTIVITY LED
IC-F4899:fte-1/0/0	Green	N/A
IC-F4899:fte-1/0/1	Green	N/A

LED info for: RC 2 FPC - 10

LEDs status:

Status LED: Green

LED info for: RC 3 FPC - 11

LEDs status:

Status LED: Green

show chassis location

List of Syntax	Syntax on page 945 Syntax (TX Matrix Router) on page 945 Syntax (TX Matrix Plus Router) on page 945 Syntax (MX Series Router) on page 945 Syntax (QFX Series) on page 945 Syntax (OCX Series) on page 945
Syntax	show chassis location
Syntax (TX Matrix Router)	show chassis location <fpc interface (by-name <i>name</i> by-slot fpc number lcc number) lcc number scc>
Syntax (TX Matrix Plus Router)	show chassis location <fpc interface (by-name <i>name</i> by-slot fpc number lcc number) lcc number sfc number>
Syntax (MX Series Router)	show chassis location <all-members> <local> <member <i>member-id</i> >
Syntax (QFX Series)	show chassis location <interconnect-device <i>name</i> > <node-device <i>name</i> >
Syntax (OCX Series)	show chassis location
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Display the physical location of the chassis. This command can only be used on the master Routing Engine.
Options	<p>none—Display all information about the physical location of the chassis. On a TX Matrix router, display all information about the physical location of the TX Matrix router and its attached T640 routers. On a TX Matrix Plus router, display all information about the physical location of the TX Matrix Plus router and its attached routers.</p> <p>all-members—(MX Series routers only) (Optional) Display the physical location of the chassis for all the member routers in the Virtual Chassis configuration.</p> <p>fpc—(TX Matrix router and TX Matrix Plus router only) (Optional) Display the physical location of all Flexible PIC Concentrators (FPCs).</p> <p>interconnect-device <i>name</i>—(QFabric systems only) (Optional) Display the physical location of the Interconnect device.</p>

interface by-name *name*—(TX Matrix and TX Matrix Plus routers only) (Optional) Display the physical location of a specified interface name. On a TX Matrix router, this option displays the FPC number and T640 router (line-card chassis) number associated with the specified interface. On a TX Matrix Plus router, this option displays the FPC number and router (line-card chassis) number associated with the specified interface.

interface by-slot *fpc number lcc number*—(TX Matrix and TX Matrix Plus router only) (Optional) On a TX Matrix router, display the global FPC number of an interface by specifying its local FPC number and T640 router (line-card chassis) number. On a TX Matrix Plus router, display the global FPC number of an interface by specifying its local FPC number and router (line-card chassis) number.

- The global FPC number is the FPC slot number when all the FPC slots in the routing matrix are considered: **0** through **31**. On TX Matrix Plus router with 3D SIBs, the value is **0** through **63**. The local FPC number is the FPC slot number on a particular T640 router.
- For **fpc**, replace *number* with a value from **0** through **7**.
- For **lcc**, replace *number* with a value from **0** through **7**.

lcc *number*—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display the physical location of a specified T640 router (line-card chassis) that is connected to a TX Matrix router. On a TX Matrix Plus router, display the physical location of a specified router (line-card chassis) that is connected to a TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(MX Series routers only) (Optional) Display the physical location of the chassis for the local Virtual Chassis member.

member *member-id*—(MX Series routers only) (Optional) Display the physical location of the chassis for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

node-device *name*—(QFabric systems only) (Optional) Display the physical location of the Node device.

scc—(TX Matrix routers only) (Optional) Display the physical location of the TX Matrix router (switch-card chassis).

sfc—(TX Matrix Plus routers only) (Optional) Display the physical location of the TX Matrix Plus router (or switch-fabric chassis).

Required Privilege Level view

Related Documentation • [Displaying Chassis Physical Locations for a Routing Matrix with a TX Matrix Plus Router](#)

List of Sample Output [show chassis location on page 947](#)
[show chassis location fpc \(TX Matrix Router\) on page 948](#)
[show chassis location interface by-slot \(TX Matrix Router\) on page 948](#)
[show chassis location fpc \(TX Matrix Plus Router\) on page 948](#)
[show chassis location interface by-slot \(TX Matrix Plus Router\) on page 948](#)
[show chassis location \(QFX Series and OCX Series\) on page 948](#)
[show chassis location \(QFabric Systems\) on page 948](#)

Output Fields Table 63 lists the output fields for the **show chassis location** command. Output fields are listed in the approximate order in which they appear.

Table 63: show chassis location Output Fields

Field Name	Field Description
country-code	Country code information.
postal-code	Postal code information.
Building	Building information.
Floor	Floor information.
Global FPC	Global FPC number. The FPC slot number, when all FPC slots in the routing matrix are considered. The range of values is 0 through 31. On TX Matrix Plus router with 3D SIBs the value is 0 through 63.
LATA	Local access transport area information.
LCC	Line-card chassis number. On a TX Matrix router, the number of a particular T640 router connected to the TX Matrix router. On a TX Matrix Plus router, the number of a particular router connected to the TX Matrix Plus router.
Local FPC	Local FPC number. On a TX Matrix router, the FPC slot number on a particular T640 router. On a TX Matrix Plus router, the FPC slot number on a particular router.

Sample Output

show chassis location

```
user@host> show chassis location
```

```
country-code: US
postal-code: 94404
Building: Building 2, Floor: 2
```

show chassis location fpc (TX Matrix Router)

```
user@host> show chassis location fpc
Global FPC      LCC      Local FPC
      17         2         1
      21         2         5
```

show chassis location interface by-slot (TX Matrix Router)

```
user@host> show chassis location interface by-slot fpc 1 lcc 1
Global FPC: 9
```

show chassis location fpc (TX Matrix Plus Router)

```
user@host> show chassis location fpc
Global FPC      LCC      Local FPC
      0         0         0
      1         0         1
```

show chassis location interface by-slot (TX Matrix Plus Router)

```
user@host> show chassis location interface by-slot fpc 2 lcc 1
Global FPC: 10
```

show chassis location (QFX Series and OCX Series)

```
user@switch> show chassis location
country-code: US
postal-code: 94404
Building: Building 2, Floor: 2
```

show chassis location (QFabric Systems)

```
user@switch> show chassis location interconnect-device interconnect1
country-code: US
postal-code: 94404
Building: Building 2, Floor: 2
```

show chassis mac-addresses

List of Syntax	Syntax on page 949 Syntax (TX Matrix Router) on page 949 Syntax (TX Matrix Plus Router) on page 949 Syntax (MX Series Router) on page 949 Syntax (MX104, MX2010, and MX2020 3D Universal Edge Routers) on page 949 Syntax (QFX Series) on page 949 Syntax (OCX Series) on page 949 Syntax (ACX Series Universal Access Routers) on page 949
Syntax	show chassis mac-addresses
Syntax (TX Matrix Router)	show chassis mac-addresses <fcc number scc>
Syntax (TX Matrix Plus Router)	show chassis mac-addresses <fcc number sfc number>
Syntax (MX Series Router)	show chassis mac-addresses <all-members> <local> <member member-id>
Syntax (MX104, MX2010, and MX2020 3D Universal Edge Routers)	show chassis mac-addresses
Syntax (QFX Series)	show chassis mac-addresses <interconnect-device name> <node-group name>
Syntax (OCX Series)	show chassis mac-addresses
Syntax (ACX Series Universal Access Routers)	show chassis mac-addresses
Release Information	<p>Command introduced before JUNOS Release 7.4.</p> <p>Command introduced in JUNOS Release 9.0 for EX Series switches.</p> <p>sfc option introduced for the TX Matrix Plus router in JUNOS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for QFX Series.</p> <p>Command introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 13.2 for MX104 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Display the media access control (MAC) addresses for the router, switch chassis, or switch.

Options **none**—(TX Matrix, TX Matrix Plus routers, QFX Series, and OCX Series Switches) Display the MAC addresses for the router chassis or switch. On a TX Matrix router, display MAC addresses on the TX Matrix router and its attached T640 routers. On a TX Matrix Plus router, display MAC addresses on the TX Matrix Plus router and its attached routers.

all-members—(MX Series routers only) (Optional) Display the MAC addresses for all the member routers of the Virtual Chassis configuration.

interconnect-device *name*—(QFabric systems only) (Optional) Display the MAC addresses for the Interconnect device.

lcc *number*—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display MAC addresses for a specified T640 router (line-card chassis) that is connected to the TX Matrix Plus router. On a TX Matrix Plus router, display MAC addresses for a specified router (line-card chassis) that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(MX Series routers only) (Optional) Display the MAC addresses for the local Virtual Chassis member.

member *member-id*—(MX Series routers only) (Optional) Display the MAC addresses for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

node-group *name*—(QFabric systems only) (Optional) Display the MAC addresses for the specified Node group.

scc—(TX Matrix routers only) (Optional) Display MAC addresses for the TX Matrix router (switch-card chassis).

sfc *number*—(TX Matrix Plus routers only) (Optional) Display MAC addresses for the TX Matrix Plus router (or switch-fabric chassis).

Required Privilege Level view

Related Documentation

- *ACX2000 and ACX2100 Routers Hardware and CLI Terminology Mapping*

- List of Sample Output**
- [show chassis mac-addresses on page 951](#)
 - [show chassis mac-addresses \(MX104 Router\) on page 951](#)
 - [show chassis mac-addresses \(MX2010 Router\) on page 951](#)
 - [show chassis mac-addresses \(MX2020 Router\) on page 952](#)
 - [show chassis mac-addresses \(TX Matrix Router\) on page 952](#)
 - [show chassis mac-addresses \(TX Matrix Plus Router\) on page 952](#)
 - [show chassis mac-addresses \(QFX Series and OCX Series \) on page 953](#)
 - [show chassis mac-addresses interconnect-device \(QFabric Systems\) on page 953](#)
 - [show chassis mac-addresses node-group \(QFabric Systems\) on page 953](#)
 - [show chassis mac-addresses \(ACX2000 Universal Access Router\) on page 953](#)

Output Fields Table 64 lists the output fields for the **show chassis mac-addresses** command. Output fields are listed in the approximate order in which they appear.

Table 64: show chassis mac-addresses Output Fields

Field Name	Field Description
MAC address information	
Public base address	Base address of the MAC addresses allocated to this router or switch.
Public count	Number of allocated public addresses.
Private base address	Base address of the private MAC addresses allocated to this router or switch.
Private count	Number of allocated private addresses.

Sample Output

show chassis mac-addresses

```
user@host> show chassis mac-addresses
MAC address information
  Public base address  0:90:69:0:4:0
  Public count         1008
  Private base address 0:90:69:0:7:f0
  Private count        16
```

show chassis mac-addresses (MX104 Router)

```
user@host > show chassis mac-addresses
MAC address information:
  Public base address  b0:a8:6e:a1:e8:58
  Public count         2032
  Private base address b0:a8:6e:a1:f0:48
  Private count        16
```

show chassis mac-addresses (MX2010 Router)

```
user@host> show chassis mac-addresses
MAC address information:
  Public base address  64:87:88:04:50:00
  Public count         1984
```

```
Private base address  64:87:88:04:57:c0
Private count         64
```

show chassis mac-addresses (MX2020 Router)

```
user@host> show chassis mac-addresses
MAC address information:
  Public base address  2c:21:72:70:20:00
  Public count         4032
  Private base address 2c:21:72:70:2f:c0
  Private count        64
```

show chassis mac-addresses (TX Matrix Router)

```
user@host> show chassis mac-addresses
scc-re0:
-----
MAC address information:
  Public base address  00:05:85:9e:cc:00
  Public count         8064
  Private base address 00:05:85:9e:eb:80
  Private count        128
lcc0-re0:
-----
MAC address information:
  Public base address  00:05:85:68:98:00
  Public count         2032
  Private base address 00:05:85:68:9f:f0
  Private count        16
lcc2-re0:
-----
MAC address information:
  Public base address  00:05:85:68:78:00
  Public count         2032
  Private base address 00:05:85:68:7f:f0
  Private count        16
```

show chassis mac-addresses (TX Matrix Plus Router)

```
user@host> show chassis mac-addresses
sfc0-re0:
-----
MAC address information:
  Public base address  00:1d:b5:14:00:00
  Public count         65023
  Private base address 00:1d:b5:14:fd:ff
  Private count        512
lcc0-re0:
-----
MAC address information:
  Public base address  00:1f:12:7a:84:00
  Public count         2032
  Private base address 00:1f:12:7a:8b:f0
  Private count        16
lcc1-re0:
-----
MAC address information:
  Public base address  00:22:83:42:48:00
  Public count         2032
  Private base address 00:22:83:42:4f:f0
```



```

Private count          16

lcc2-re0:
-----
MAC address information:
Public base address    00:1f:12:c3:58:00
Public count           2032
Private base address    00:1f:12:c3:5f:f0
Private count           16

lcc3-re0:
-----
MAC address information:
Public base address    00:21:59:ef:b8:00
Public count           2032
Private base address    00:21:59:ef:bf:f0
Private count           16

```

show chassis mac-addresses (QFX Series and OCX Series)

```

user@switch> show chassis mac-addresses
MAC address information:
Public base address 02:00:08:00:00:00
Public count 512
Private base address 02:00:00:00:00:00
Private count 64

```

show chassis mac-addresses interconnect-device (QFabric Systems)

```

user@switch> show chassis mac-addresses interconnect-device interconnect1
MAC address information:
Public base address    00:1f:12:30:9c:c0
Public count           58
Private base address    00:1f:12:30:9c:fa
Private count           6

```

show chassis mac-addresses node-group (QFabric Systems)

```

user@switch> show chassis mac-addresses node-group NW-NG-0
MAC address information:
-----
RE:
FC MAC base    00:11:00:00:00:00
FC MAC count    2
VLAN MAC       00:11:00:00:00:09
EC6007
Base address    00:00:01:76:00:00
Count           64
EC6008
Base address    00:22:83:22:52:ae
Count           260

```

show chassis mac-addresses (ACX2000 Universal Access Router)

```

user@switch> show chassis mac-addresses
MAC address information:
Public base address    84:18:88:c0:2b:00
Public count           112
Private base address    84:18:88:c0:2b:70
Private count           16

```

show chassis pic

List of Syntax	Syntax on page 954 Syntax (TX Matrix and TX Matrix Plus Routers) on page 954 Syntax (MX Series Routers and EX Series Switches) on page 954 Syntax (MX2010 and MX2010 3D Universal Edge Routers) on page 954 Syntax (PTX Series Packet Transport Router) on page 954 Syntax (QFX Series) on page 954 Syntax (OCX Series) on page 954 Syntax (ACX Series Universal Access Routers) on page 954
Syntax	<code>show chassis pic fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i></code>
Syntax (TX Matrix and TX Matrix Plus Routers)	<code>show chassis pic fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i> <lcc <i>number</i>></code>
Syntax (MX Series Routers and EX Series Switches)	<code>show chassis pic fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i> <all-members></code> <code><local></code> <code><member <i>member-id</i>></code>
Syntax (MX2010 and MX2010 3D Universal Edge Routers)	<code>show chassis pic fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i></code>
Syntax (PTX Series Packet Transport Router)	<code>show chassis pic transport fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i></code>
Syntax (QFX Series)	<code>show chassis pic fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i> <interconnect-device <i>name</i> (fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i>)></code> <code><node-device <i>name</i> pic-slot <i>slot-number</i>></code>
Syntax (OCX Series)	<code>show chassis pic fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i></code>
Syntax (ACX Series Universal Access Routers)	<code>show chassis pic fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for QFX Series. Command introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers. Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers. Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers. Command introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers. Command introduced in Junos OS Release 13.2 for MX104 3D Universal Edge Routers. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Display status information about the PIC installed in the specified Flexible PIC Concentrator (FPC) and PIC slot.

Options **fpc-slot *slot-number***—Display information about the PIC in this particular FPC slot:

- On a TX Matrix router, if you specify the number of the T640 router by using the **lcc *number*** option (the recommended method), replace ***slot-number*** with a value from 0 through 7. Otherwise, replace ***slot-number*** with a value from 0 through 31.

Likewise, on a TX Matrix Plus router, if you specify the number of the T1600 router by using the **lcc *number*** option (the recommended method), replace ***slot-number*** with a value from 0 through 7. Otherwise, replace ***slot-number*** with a value from 0 through 31. For example, the following commands have the same result:

```
user@host> show chassis pic fpc-slot 1 lcc 1 pic-slot 1
user@host> show chassis pic fpc-slot 9 pic-slot 1
```

- M120 routers only—Replace ***slot-number*** with a value from 0 through 5.
- MX80 routers only—Replace ***slot-number*** with a value from 0 through 1.
- MX104 routers only—Replace ***slot-number*** with a value from 0 through 2.
- MX240 routers only—Replace ***slot-number*** with a value from 0 through 2.
- MX480 routers only—Replace ***slot-number*** with a value from 0 through 5.
- MX960 routers only—Replace ***slot-number*** with a value from 0 through 11.
- MX2010 routers only—Replace ***slot-number*** with a value from 0 through 9.
- MX2020 routers only—Replace ***slot-number*** with a value from 0 through 19.
- Other routers—Replace ***slot-number*** with a value from 0 through 7.
- EX Series switches:
 - EX3200 switches and EX4200 standalone switches—Replace ***slot-number*** with 0.
 - EX4200 switches in a Virtual Chassis configuration—Replace ***slot-number*** with a value from 0 through 9 (switch's member ID).
 - EX8208 switches—Replace ***slot-number*** with a value from 0 through 7 (line card).
 - EX8216 switches—Replace ***slot-number*** with a value from 0 through 15 (line card).
- QFX Series:
 - QFX3500, QFX3600, QFX5100, and OCX Series standalone switches—Replace ***slot-number*** with 0. In the command output, FPC refers to a line card. The FPC number equals the slot number for the line card.
 - QFabric systems—Replace ***slot-number*** with any number between 0 and 15. In the command output, FPC refers to a line card. The FPC number equals the slot number for the line card.

all-members—(MX Series routers and EX Series switches only) (Optional) Display PIC information for all member routers in the Virtual Chassis configuration.

interconnect-device *name*—(QFabric systems only) (Optional) Display PIC information for a specified Interconnect device.

lcc *number*—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display PIC information for a specified T640 router (or line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, display PIC information for a specified router (line-card chassis) that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(MX Series routers and EX Series switches only) (Optional) Display PIC information for the local Virtual Chassis member.

member *member-id*—(MX Series routers and EX Series switches only) (Optional) Display PIC information for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

node-device *name*—(QFabric systems only) (Optional) Display PIC information for a specified Node device.

pic-slot *slot-number*—Display information about the PIC in this particular PIC slot. For routers, replace *slot-number* with a value from 0 through 3. For EX3200 and EX4200 switches, replace *slot-number* with 0 for built-in network interfaces and 1 for interfaces on uplink modules. For EX8208 and EX8216 switches, replace *slot-number* with 0. For the QFX3500 standalone switch and the QFabric system, replace *slot-number* with 0 or 1.

transport—Display PIC information for optical transport network.

Required Privilege Level

view

Related Documentation

- [request chassis pic on page 371](#)
- [show chassis hardware on page 743](#)
- *100-Gigabit Ethernet Type 4 PIC with CFP Overview*

List of Sample Output

[show chassis pic fpc-slot pic-slot on page 959](#)
[show chassis pic fpc-slot pic-slot \(PIC Offline\) on page 960](#)

[show chassis pic fpc-slot pic-slot \(FPC Offline\) on page 960](#)
[show chassis pic fpc-slot pic-slot \(FPC Not Present\) on page 960](#)
[show chassis pic fpc-slot pic-slot \(PIC Not Present\) on page 960](#)
[show chassis pic fpc-slot 3 pic-slot 0 \(M120 Router\) on page 960](#)
[show chassis pic fpc-slot pic-slot \(MX104 Router\) on page 960](#)
[show chassis pic fpc-slot pic-slot \(MX960 Router with Bidirectional Optics\) on page 961](#)
[show chassis pic fpc-slot pic-slot \(MX480 Router with 100-Gigabit Ethernet MIC\) on page 961](#)
[show chassis pic fpc-slot pic-slot \(MX240, MX480, MX960 Routers with Application Services Modular Line Card\) on page 961](#)
[show chassis pic fpc-slot pic-slot \(MX960 Router with MPC5EQ\) on page 962](#)
[show chassis pic fpc-slot pic-slot \(MX480 Router with MPC4E\) on page 962](#)
[show chassis pic fpc-slot pic-slot \(MX480 Router with OTN Interface\) on page 962](#)
[show chassis pic fpc-slot pic-slot \(MX2010 Router with OTN Interfaces\) on page 962](#)
[show chassis pic fpc-slot pic-slot \(MX2010 Router\) on page 963](#)
[show chassis pic fpc-slot pic-slot \(MX2020 Router\) on page 963](#)
[show chassis pic fpc-slot pic-slot \(MX2020 Router with MPC5EQ and MPC6E\) on page 963](#)
[show chassis pic fpc-slot pic-slot \(MX2020 Router with MPC6E and OTN MIC\) on page 964](#)
[show chassis pic fpc-slot pic-slot \(MX2020 Router with MPC4E\) on page 964](#)
[show chassis pic fpc-slot pic-slot \(T1600 Router with 100-Gigabit Ethernet PIC\) on page 964](#)
[show chassis pic fpc-slot pic-slot lcc \(TX Matrix Router\) on page 965](#)
[show chassis pic fpc-slot pic-slot lcc \(TX Matrix Plus Router\) on page 965](#)
[show chassis pic fpc-slot pic-slot \(Next-Generation SONET/SDH SFP\) on page 965](#)
[show chassis pic fpc-slot pic-slot \(12-Port T1/E1\) on page 965](#)
[show chassis pic fpc-slot 0 pic-slot 1 \(4x CHOC3 SONET CE SFP\) on page 966](#)
[show chassis pic fpc-slot 0 pic-slot 0 \(SONET/SDH OC3/STM1 \[Multi-Rate\] MIC with SFP\) on page 966](#)
[show chassis pic fpc-slot 3 pic-slot 0 \(8-port Channelized SONET/SDH OC3/STM1 \[Multi-Rate\] MIC with SFP\) on page 966](#)
[show chassis pic fpc-slot 5 pic-slot 0 \(4-Port Channelized SONET/SDH OC3/STM1 \[Multi-Rate\] MIC with SFP\) on page 967](#)
[show chassis pic fpc-slot 1 pic-slot 0 \(1-Port OC192/STM64 MIC with XFP\) on page 967](#)
[show chassis pic fpc-slot 1 pic-slot 2 \(8-Port DS3/E3 MIC\) on page 967](#)
[show chassis pic fpc-slot pic-slot \(OTN\) on page 967](#)
[show chassis pic fpc-slot pic-slot \(QFX3500 Switch\) on page 967](#)
[show chassis pic fpc-slot pic-slot \(QFX5100 Switches and OCX Series\) on page 968](#)
[show chassis pic interconnect-device fpc-slot pic-slot \(QFabric Systems\) on page 968](#)
[show chassis pic node-device fpc-slot pic-slot \(QFabric System\) on page 968](#)
[show chassis pic fpc-slot 0 pic-slot 1 \(ACX2000 Universal Access Router\) on page 969](#)
[show chassis pic FPC-slot 1 PIC-slot 0 \(MX Routers with Media Services Blade \[MSB\]\) on page 969](#)
[show chassis pic FPC slot 1, PIC slot 2 \(MX Routers with Media Services Blade \[MSB\]\) on page 969](#)

Output Fields Table 65 lists the output fields for the **show chassis pic** command. Output fields are listed in the approximate order in which they appear.

Table 65: show chassis pic Output Fields

Field Name	Field Description
Type	<p>PIC type.</p> <p>NOTE: On the 1-port OC192/STM64 MICs with the SDH framing mode, the type is displayed as MIC-3D-1STM64-XFP and with the SONET framing mode, the type is displayed as MIC-3D-1OC192-XFP. By default, the 1-port OC192/STM64 MICs displays the type as MIC-3D-1OC192-XFP.</p>
Account Layer2 Overhead	(MX Series routers) Indicates whether functionality to count the Layer 2 overhead bytes in the interface statistics at the PIC level is enabled or disabled.
ASIC type	Type of ASIC on the PIC.
State	<p>Status of the PIC. State is displayed only when a PIC is in the slot.</p> <ul style="list-style-type: none"> • Online— PIC is online and running. • Offline—PIC is powered down.
PIC version	PIC hardware version.
Uptime	How long the PIC has been online.
Package	(Multiservices PICs only) Services package supported: Layer-2 or Layer-3 .
Port Number	Port number for the PIC.
Cable Type	Type of cable connected to the port: LH , LX , or SX .
PIC Port Information (MX480 Router 100-Gigabit Ethernet CFP)	<p>Port-level information for the PIC.</p> <ul style="list-style-type: none"> • Port—Port number • Cable type—Type of optical transceiver installed. • Fiber type—Type of fiber. SM is single-mode. • Xcvr vendor—Transceiver vendor name. • Xcvr vendor part number—Transceiver vendor part number. • Wavelength—Wavelength of the transmitted signal. Uplinks and downlinks are always 1550 nm. There is a separate fiber for each direction • Xcvr Firmware—Transceiver firmware version.

Table 65: show chassis pic Output Fields (*continued*)

Field Name	Field Description
PIC Port Information (MX960 Router Bidirectional Optics)	<p>Port-level information for the PIC.</p> <ul style="list-style-type: none"> • Port—Port number • Cable type—Type of small form-factor pluggable (SFP) optical transceiver installed. Uplink interfaces display -U. Down link interfaces display -D. • Fiber type—Type of fiber. SM is single-mode. • Xcvr vendor—Transceiver vendor name. • Xcvr vendor part number—Transceiver vendor part number. <ul style="list-style-type: none"> • BX10-10-km bidirectional optics. • BX40-40-km bidirectional optics. • SFP-LX-40-km SFP optics. • Wavelength—Wavelength of the transmitted signal. Uplinks are always 1310 nm. Downlinks are either 1490 nm or 1550 nm.
PIC Port Information (Next-Generation SONET/SDH SFP)	<p>Port-level information for the next-generation SONET/SDH SFP PIC.</p> <ul style="list-style-type: none"> • Port—Port number. • Cable type—Type of small form-factor pluggable (SFP) optical transceiver installed. • Fiber type—Type of fiber: SM (single-mode) or MM (multimode). • Xcvr vendor—Transceiver vendor name. • Xcvr vendor part number—Transceiver vendor part number. • Wavelength—Wavelength of the transmitted signal. Next-generation SONET/SDH SFPs use 1310 nm.
PIC port information (MX104 router)	<p>Port-level information for the PIC.</p> <ul style="list-style-type: none"> • Port—Port number • Cable type—Type of optical transceiver installed. • Fiber type—Type of fiber. SM is single-mode. • Xcvr vendor—Transceiver vendor name. • Xcvr vendor part number—Transceiver vendor part number. • Wavelength—Wavelength of the transmitted signal. • Xcvr Firmware—Firmware version of the transceiver.
Multirate Mode	Rate-selectability status for the MIC: Enabled or Disabled .
Channelization	Indicates whether channelization is enabled or disabled on the DS3/E3 MIC.

Sample Output

show chassis pic fpc-slot pic-slot

```

user@host> show chassis pic fpc-slot 2 pic-slot 0
PIC fpc slot 2 pic slot 0 information:
Type                               10x 1GE(LAN), 1000 BASE

```

```

ASIC type           H chip
State               Online
PIC version         1.1
Uptime              1 day, 50 minutes, 58 seconds
PIC Port Information:
Port      Cable      Xcvr      Xcvr Vendor
Number    Type        Vendor Name Part Number
0         GIGE 1000EX  FINISAR CORP.  FTRJ8519P1BNL-J3
1         GIGE 1000EX  FINISAR CORP.  FTRJ-8519-7D-JUN

```

show chassis pic fpc-slot pic-slot (PIC Offline)

```

user@host> show chassis pic fpc-slot 1 pic-slot 0
PIC fpc slot 1 pic slot 0 information:
State                               Offline

```

show chassis pic fpc-slot pic-slot (FPC Offline)

```

user@host> show chassis pic fpc-slot 1 pic-slot 0
FPC 1 is not online

```

show chassis pic fpc-slot pic-slot (FPC Not Present)

```

user@host> show chassis pic fpc-slot 4 pic-slot 0
FPC slot 4 is empty

```

show chassis pic fpc-slot pic-slot (PIC Not Present)

```

user@host> show chassis pic fpc-slot 5 pic-slot 2
FPC 5, PIC 2 is empty

```

show chassis pic fpc-slot 3 pic-slot 0 (M120 Router)

```

user@host> show chassis pic fpc-slot 3 pic-slot 0
PC slot 3, PIC slot 0 information:
Type                2x G/E IQ, 1000 BASE
ASIC type           IQ GE 2 VLAN-TAG FPGA
State               Online
PIC version         1.16
Uptime              3 hours, 3 minutes

PIC Port Information:
Port      Cable      Xcvr      Xcvr Vendor
Number    Type        Vendor Name Part Number
0         GIGE 1000SX  FINISAR CORP.  FTRJ8519P1BNL-J3
1         GIGE 1000SX  FINISAR CORP.  FTRJ-8519-7D-JUN

```

show chassis pic fpc-slot pic-slot (MX104 Router)

```

user@host> show chassis pic fpc-slot 1 pic-slot 1
FPC slot 1, PIC slot 1 information:
Type                10x 1GE(LAN) -E SFP
State               Online
PIC version         1.1
Uptime              1 hour, 30 minutes, 59 seconds

PIC port information:
Fiber              Xcvr vendor      Wave-      Xcvr
Port Cable type    type Xcvr vendor      part number length
Firmware
3    GIGE 1000T    n/a  Methode Elec.     SP7041-M1-JN  n/a      0.0

```


6	GIGE 1000LX10	SM	FINISAR CORP.	FTLF1318P2BTL-J1	1310 nm	0.0
8	GIGE 1000T	n/a	Methode Elec.	SP7041-M1-JN	n/a	0.0
9	GIGE 1000T	n/a	Methode Elec.	SP7041-M1-JN	n/a	0.0

show chassis pic fpc-slot pic-slot (MX960 Router with Bidirectional Optics)

```

user@host> show chassis pic fpc-slot 4 pic-slot 1
FPC slot 4, PIC slot 1 information:
  Type                10x 1GE(LAN)
  State                Online
  PIC version          0.0
  Uptime               18 days, 5 hours, 41 minutes, 54 seconds

PIC port information:

```

Port	Cable type	Fiber type	Xcvr vendor	Xcvr vendor part number	Wavelength
0	SFP-1000BASE-BX10-D	SM	SumitomoElectric	SBP6H44-J3-BW-49	1490 nm
1	SFP-1000BASE-BX10-D	SM	SumitomoElectric	SBP6H44-J3-BW-49	1490 nm
2	SFP-1000BASE-BX10-D	SM	SumitomoElectric	SBP6H44-J3-BW-49	1490 nm
3	SFP-1000BASE-BX10-D	SM	OCP	TRXBG1LXDBVM2-JW	1490 nm
4	SFP-1000BASE-BX10-D	SM	OCP	TRXBG1LXDBVM2-JW	1490 nm
5	SFP-1000BASE-BX10-U	SM	SumitomoElectric	SBP6H44-J3-BW-31	1310 nm
6	SFP-1000BASE-BX10-U	SM	SumitomoElectric	SBP6H44-J3-BW-31	1310 nm
7	SFP-1000BASE-BX10-U	SM	OCP	TRXBG1LXDBBMH-J1	1310 nm
8	SFP-1000BASE-BX10-U	SM	OCP	TRXBG1LXDBBMH-J1	1310 nm
9	SFP-1000BASE-BX10-U	SM	SumitomoElectric	SBP6H44-J3-BW-31	1310 nm

show chassis pic fpc-slot pic-slot (MX480 Router with 100-Gigabit Ethernet MIC)

```

user@host> show chassis pic fpc-slot 1 pic-slot 2
FPC slot 1, PIC slot 2 information:
  Type                1X100GE CFP
  State                Online
  PIC version          2.10
  Uptime               4 minutes, 48 seconds

PIC port information:

```

Port	Cable type	Fiber type	Xcvr vendor	Xcvr vendor part number	Wavelength
0	100GBASE LR4	SM	FINISAR CORP.	FTLC1181RDNS-J3	1310 nm

```

  Xcvr vendor
  firmware version
  1.8

```

show chassis pic fpc-slot pic-slot (MX240, MX480, MX960 Routers with Application Services Modular Line Card)

```

user@host> show chassis pic fpc-slot 1 pic-slot 2
FPC slot 1, PIC slot 2 information:
  Type                AS-MXC
  State                Online
  PIC version          1.0
  Uptime               11 hours, 18 minutes, 3 seconds

```

show chassis pic fpc-slot pic-slot (MX960 Router with MPC5EQ)

```

user@host> show chassis pic fpc-slot 0 pic-slot 3
FPC slot 0, PIC slot 3 information:
  Type                1X100GE CFP2 OTN
  State               Online
  PIC version         0.0
  Uptime              1 hour, 22 minutes, 42 seconds

PIC port information:

```

		Fiber	Xcvr vendor	Wave-	Xcvr
Port	Cable type	type	Xcvr vendor	part number	length
Firmware					
0	10GBASE LR4	n/a	Oclaro Inc.	TRB5E20FNF-LF150	1309 nm 1.0

show chassis pic fpc-slot pic-slot (MX480 Router with MPC4E)

```

user@host> show chassis pic fpc-slot 3 pic-slot 0
FPC slot 3, PIC slot 0 information:
  Type                4x10GE SFPP
  State               Online
  PIC version         0.0
  Uptime              41 seconds

PIC port information:

```

		Fiber	Xcvr vendor	Wave-	Xcvr
Port	Cable type	type	Xcvr vendor	part number	length
Firmware					
0	10GBASE SR	MM	OPNEXT, INC.	TRS2001EM-0014	850 nm 0.0
1	10GBASE SR	MM	OPNEXT, INC.	TRS2001EM-0014	850 nm 0.0

show chassis pic fpc-slot pic-slot (MX480 Router with OTN Interface)

```

user@host> show chassis pci fpc-slot 4 pic-slot 0
FPC slot 4, PIC slot 0 information:
  Type                12X10GE SFPP OTN
  State               Online
  PIC version         0.0
  Uptime              5 hours, 28 minutes, 23 seconds

PIC port information:

```

		Fiber	Xcvr vendor	Wave-	Xcvr
Port	Cable type	type	Xcvr vendor	part number	length
Firmware					
0	10GBASE SR	MM	FINISAR CORP.	FTLX8571D3BNL-J1	850 nm 0.0
1	10GBASE SR	MM	FINISAR CORP.	FTLX8571D3BCL-J1	850 nm 0.0
2	10GBASE SR	MM	OPNEXT, INC.	TRS2001EM-0014	850 nm 0.0

show chassis pic fpc-slot pic-slot (MX2010 Router with OTN Interfaces)

```

user@host> show chassis pic fpc-slot 9 pic-slot 0

```

```
FPC slot 9, PIC slot 0 information:
Type                2X100GE CFP2 OTN
State                Online
PIC version          1.9
Uptime               3 hours, 56 minutes, 16 seconds
```

```
PIC port information:

Fiber                Xcvr vendor      Wave-    Xcvr
Port Cable type      type  Xcvr vendor      part number  length
Firmware
0    100GBASE LR4-D   SM    FUJITSU          FIM37300/222 1310 nm  1.3
1    100GBASE SR10    MM    AVAGO            AFBR-8420Z    n/a      1.0
```

show chassis pic fpc-slot pic-slot (MX2010 Router)

```
user@host> show chassis pic fpc-slot 9 pic-slot 3
FPC slot 9, PIC slot 3 information:
Type                1X100GE CFP
Account Layer2 Overhead Enabled
State                Online
PIC version          0.0
Uptime               14 hours, 51 seconds
```

show chassis pic fpc-slot pic-slot (MX2020 Router)

```
user@host> show chassis pic fpc-slot 19 pic-slot 3
FPC slot 19, PIC slot 3 information:
Type                4x 10GE(LAN) SFP+
Account Layer2 Overhead Enabled
State                Online
PIC version          0.0
Uptime               1 day, 11 hours, 26 minutes, 36 seconds

PIC port information:

Fiber                Xcvr vendor      Wave-    Xcvr
Port Cable type      type  Xcvr vendor      part number  length
Firmware
0    10GBASE SR        MM    SumitomoElectric SPP5200SR-J6-M 850 nm  0.0
1    10GBASE SR        MM    SumitomoElectric SPP5200SR-J6-M 850 nm  0.0
2    10GBASE SR        MM    SumitomoElectric SPP5200SR-J6-M 850 nm  0.0
3    10GBASE SR        MM    SumitomoElectric SPP5200SR-J6-M 850 nm  0.0
```

show chassis pic fpc-slot pic-slot (MX2020 Router with MPC5EQ and MPC6E)

```
user@host> show chassis pic fpc-slot 18 pic-slot 2
FPC slot 18, PIC slot 2 information:
Type                3X40GE QSFP
State                Online
PIC version          0.0
Uptime               6 minutes, 31 seconds

PIC port information:

Fiber                Xcvr vendor      Wave-    Xcvr
Port Cable type      type  Xcvr vendor      part number  length
```

```

Firmware
 0  40GBASE SR4      MM  AVAGO          AFBR-79E4Z-D-JU2  850 nm  0.0
 1  40GBASE SR4      MM  AVAGO          AFBR-79E4Z-D-JU2  850 nm  0.0
 2  40GBASE SR4      MM  AVAGO          AFBR-79E4Z-D-JU2  850 nm  0.0

```

show chassis pic fpc-slot pic-slot (MX2020 Router with MPC6E and OTN MIC)

```

user@host> show chassis pic fpc-slot 3 pic-slot 0
FPC slot 0, PIC slot 1 information:
  Type                24X10GE SFPP OTN
  State                Online
  PIC version          1.1
  Uptime               1 hour, 33 minutes, 59 seconds

PIC port information:

```

		Fiber		Xcvr vendor	Wave-	Xcvr
Port	Cable type	type	Xcvr vendor	part number	length	
Firmware						
7	10GBASE SR	MM	SumitomoElectric	SPP5200SR-J6-M	850 nm	0.0
9	10GBASE SR	MM	FINISAR CORP.	FTLX8571D3BNL-J1	850 nm	0.0
12	10GBASE LR	SM	FINISAR CORP.	FTLX1472M3BNL-J3	1310 nm	0.0
20	10GBASE ZR	SM	FINISAR CORP.	FTLX1871M3BNL-J3	1550 nm	0.0
21	10GBASE ER	SM	FINISAR CORP.	FTLX1671D3BTL-J4	1550 nm	0.0
22	10GBASE LR	SM	SOURCEPHOTONICS	SPP10SLREDFCJNP	1310 nm	0.0
23	10GBASE LR	SM	FINISAR CORP.	FTLX1471D3BNL-J1	1310 nm	0.0

show chassis pic fpc-slot pic-slot (MX2020 Router with MPC4E)

```

user@host> show chassis pic fpc-slot 14 pic-slot 0
FPC slot 14, PIC slot 1 information:
  Type                1X100GE CFP
  State                Online
  PIC version          0.0
  Uptime               1 day, 2 hours, 19 minutes, 18 seconds

PIC port information:

```

		Fiber		Xcvr vendor	Wave-	Xcvr
Port	Cable type	type	Xcvr vendor	part number	length	
Firmware						
0	100GBASE SR10	MM	Reflex Photonics	CF-X12-C11801-50	860 nm	4.7

show chassis pic fpc-slot pic-slot (T1600 Router with 100-Gigabit Ethernet PIC)

```

user@host> run show chassis pic fpc-slot 3 pic-slot 1
FPC slot 3, PIC slot 1 information:
  Type                100GE SLOT1
  ASIC type            Brooklyn 100GE FPGA
  State                Online
  PIC version          1.3
  Uptime               10 minutes, 44 seconds

```

PIC port information:

Port	Cable type	Fiber type	Xcvr vendor	Xcvr vendor part number	Wavelength
0	100GBASE LR4	SM	Opnext Inc.	TRC5E20ENFSF000F	1310 nm

show chassis pic fpc-slot pic-slot lcc (TX Matrix Router)

```
user@host> show chassis pic fpc-slot 1 pic-slot 1 lcc 0
lcc0-re0:
```

PIC fpc slot 1 pic slot 1 information:

Type	4x OC-3 SONET, SMIR
ASIC type	D chip
State	Online
PIC version	1.2
Uptime	5 days, 2 hours, 12 minutes, 8 seconds

show chassis pic fpc-slot pic-slot lcc (TX Matrix Plus Router)

```
user@host> show chassis pic pic-slot 0 fpc-slot 8
lcc0-re0:
```

FPC slot 8, PIC slot 0 information:

Type	1x 10GE(LAN/WAN)
State	Online
Uptime	2 hours, 46 minutes, 23 seconds

PIC port information:

Port	Cable type	Fiber type	Xcvr vendor	part number	Wavelength
0	10GBASE ZR	SM	Opnext Inc.	TRF7061BN-LF150	1550 nm
0	10GBASE ZR	SM	FINISAR CORP.	FTRX-1811-3-J2	1550 nm

show chassis pic fpc-slot pic-slot (Next-Generation SONET/SDH SFP)

```
user@host> show chassis pic fpc-slot 4 pic-slot 0
```

FPC slot 4, PIC slot 0 information:

Type	4x OC-3 1x OC-12 SFP
ASIC type	D FPGA
State	Online
PIC version	1.3
Uptime	1 day, 50 minutes, 4 seconds

PIC port information:

Port	Cable type	Fiber type	Xcvr vendor	Xcvr vendor part number	Wavelength
0	OC48 short reach	SM	FINISAR CORP.	FTRJ1321P1BTL-J2	1310 nm
1	OC3 short reach	MM	OCF	TRPA03MM3BAS-JE	1310 nm
2	OC3 short reach	MM	OCF	TRXA03MM3BAS-JW	1310 nm
3	OC12 inter reach	SM	FINISAR CORP.	FTLF1322P1BTR	1310 nm

show chassis pic fpc-slot pic-slot (12-Port T1/E1)

```
user@host> show chassis pic fpc-slot 0 pic-slot 3
```

FPC slot 0, PIC slot 3 information:

Type	12x T1/E1 CE
State	Online
PIC version	1.1
CPU load average	1 percent

```

Interrupt load average      0 percent
Total DRAM size            128 MB
Memory buffer utilization   100 percent
Memory heap utilization     4 percent
Uptime                     1 day, 22 hours, 28 minutes, 12 seconds
Internal Clock Synchronization Normal

```

show chassis pic fpc-slot 0 pic-slot 1 (4x CHOC3 SONET CE SFP)

```

user@host> show chassis pic fpc-slot 0 pic-slot 1
FPC slot 0, PIC slot 1 information:
  Type          4x CHOC3 SONET CE SFP
  State          Online
  PIC version    1.3
  CPU load average 1 percent
  Interrupt load average 0 percent
  Total DRAM size 128 MB
  Memory buffer utilization 99 percent
  Memory heap utilization 4 percent
  Uptime         1 day, 22 hours, 55 minutes, 37 seconds
  Internal Clock Synchronization Normal

```

PIC port information:

Port	Cable type	Fiber type	Xcvr vendor	Xcvr vendor part number	Wavelength
0	OC3 short reach	MM	AVAGO	HFBR-57E0P-JU2	n/a
1	OC3 short reach	MM	AVAGO	HFBR-57E0P-JU2	n/a
3	OC3 long reach	SM	OPNEX INC	TRF5456AVLB314	1310 nm

show chassis pic fpc-slot 0 pic-slot 0 (SONET/SDH OC3/STM1 [Multi-Rate] MIC with SFP)

```

user@host> show chassis pic fpc-slot 0 pic-slot 0
FPC slot 0, PIC slot 0 information:
  Type          MIC-3D-80C30C12-40C48
  State          Online
  PIC version    1.8
  Uptime         3 days, 22 hours, 3 minutes, 50 seconds

```

PIC port information:

Port	Cable type	Fiber type	Xcvr vendor	Xcvr vendor part number	Wavelength
1	OC12 inter reach	SM	FINISAR CORP	FTRJ1322P1BTR-J3	1310 nm
7	OC12 inter reach	SM	FINISAR CORP	FTRJ1322P1BTR-J3	1310 nm

Multirate Mode Enabled

show chassis pic fpc-slot 3 pic-slot 0 (8-port Channelized SONET/SDH OC3/STM1 [Multi-Rate] MIC with SFP)

```

user@host> show chassis pic fpc-slot 3 pic-slot 0
FPC slot 3, PIC slot 0 information:
  Type          MIC-3D-8CHOC3-4CHOC12
  State          Online
  PIC version    1.9
  Uptime         1 hour, 21 minutes, 24 seconds

```

PIC port information:

Port	Cable type	Fiber type	Xcvr vendor	Xcvr vendor part number	Wavelength
0	OC12 short reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm
1	OC12 short reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm
2	OC12 inter reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J2	1310 nm
4	OC12 short reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm
5	OC12 short reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm

6	OC12 short reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm
7	OC12 short reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm

show chassis pic fpc-slot 5 pic-slot 0 (4-Port Channelized SONET/SDH OC3/STM1 [Multi-Rate] MIC with SFP)

```
user@host> show chassis pic fpc-slot 5 pic-slot 0
FPC slot 5, PIC slot 0 information:
  Type                MIC-3D-4CHOC3-2CHOC12
  State                Online
  PIC version          1.9
  Uptime               1 hour, 21 minutes

PIC port information:
  Port  Cable type      Fiber type  Xcvr vendor  Xcvr vendor  Wavelength
  1     OC12 inter reach SM         FINISAR CORP. FTRJ1322P1BTR-J3 1310 nm
  2     OC12 inter reach SM         FINISAR CORP. FTRJ1322P1BTR-J3 1310 nm
  3     OC12 short reach SM         FINISAR CORP. FTRJ1322P1BTR-J3 1310 nm
```

show chassis pic fpc-slot 1 pic-slot 0 (1-Port OC192/STM64 MIC with XFP)

```
user@host> show chassis pic fpc-slot 1 pic-slot 0
FPC slot 1, PIC slot 0 information:
  Type                MIC-3D-10C192-XFP
  State                Online
  PIC version          1.2
  Uptime               1 day, 11 hours, 4 minutes, 6 seconds

PIC port information:
  Port  Cable type      Fiber type  Xcvr vendor  Xcvr vendor  Wavelength
  0     OC192 short reach n/a         FINISAR CORP. FTLX1412M3BCL-J3 1310 nm
```

show chassis pic fpc-slot 1 pic-slot 2 (8-Port DS3/E3 MIC)

```
user@host> show chassis pic fpc-slot 1 pic-slot 2
FPC slot 1, PIC slot 2 information:
  Type                MIC-3D-8DS3-E3
  State                Online
  PIC version          1.10
  Uptime               4 days, 1 hour, 29 minutes, 19 seconds
  Channelization Mode Disabled
```

show chassis pic fpc-slot pic-slot (OTN)

```
user@host> show chassis pic fpc-slot 5 pic-slot 0
PIC fpc slot 5 pic slot 0 information:
  Type                1x10GE(LAN),OTN
  ASIC type            H chip
  State                Online
  PIC version          1.0
  Uptime               5 minutes, 50 seconds
```

show chassis pic fpc-slot pic-slot (QFX3500 Switch)

```
user@switch> show chassis pic fpc-slot 0 pic-slot 0
FPC slot 0, PIC slot 0 information:
  Type 48x 10G-SFP+ Builtin
  State Online
  Uptime 3 days, 3 hours, 5 minutes, 20 seconds
```

show chassis pic fpc-slot pic-slot (QFX5100 Switches and OCX Series)

```

user@switch> show chassis pic fpc-slot 0 pic-slot 0
FPC slot 0, PIC slot 0 information:
  Type                Unknown Builtin
  State               Online
  Uptime              1 day, 17 hours, 5 minutes, 9 seconds

```

show chassis pic interconnect-device fpc-slot pic-slot (QFabric Systems)

```

user@switch> show chassis pic interconnect-device interconnect1 fpc-slot 9 pic-slot 0
FPC slot 9, PIC slot 0 information:
  Type                16x 40G-GE Builtin
  State               Online
  Uptime              2 hours, 47 minutes, 40 seconds

```

show chassis pic node-device fpc-slot pic-slot (QFabric System)

```

user@switch> show chassis pic node-device node1 pic-slot 0
FPC slot node1, PIC slot 0 information:
  Type                48x 10G-SFP+ Builtin
  State               Online
  Uptime              2 hours, 52 minutes, 37 seconds

```

PIC port information:

Port	Cable type	Fiber type	Xcvr vendor	Xcvr vendor part number	Wavelength
0	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
1	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
2	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
3	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
4	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
5	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
6	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
7	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
8	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
9	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
10	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
11	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
12	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
13	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
14	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
15	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
16	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
17	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
18	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
19	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
20	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
21	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
22	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
23	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
24	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
25	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
26	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
27	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
28	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
29	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
30	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
31	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
32	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
33	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm

34	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
35	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
36	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
37	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
38	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
39	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
40	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
41	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
42	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
43	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
44	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
45	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
46	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
47	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm

show chassis pic fpc-slot 0 pic-slot 1 (ACX2000 Universal Access Router)

```

user@host> show chassis pic fpc-slot 0 pic-slot 1
FPC slot 0, PIC slot 1 information:
Type                               8x 1GE(LAN) RJ45 Built-in
State                               Online
Uptime                             6 days, 2 hours, 51 minutes, 11 seconds

```

show chassis pic FPC-slot 1 PIC-slot 0 (MX Routers with Media Services Blade [MSB])

```

user@switch> show chassis pic fpc-slot 1 pic-slot 0
FPC slot 1, PIC slot 0 information:
Type                               AS-MSB
State                               Online
PIC version                         1.6
Uptime                             11 hours, 17 minutes, 56 seconds

```

show chassis pic FPC slot 1, PIC slot 2 (MX Routers with Media Services Blade [MSB])

```

user@switch> show chassis pic fpc-slot 1 pic-slot 2
Type                               AS-MXC
State                               Online
PIC version                         1.0
Uptime                             11 hours, 18 minutes, 3 seconds

```

show chassis routing-engine

List of Syntax [Syntax on page 970](#)
 [Syntax \(ACX Series Universal Access Routers\) on page 970](#)
 [Syntax \(EX Series Switches\) on page 970](#)
 [Syntax \(QFX Series\) on page 970](#)
 [Syntax \(MX Series Routers\) on page 970](#)
 [Syntax \(MX2010 3D Universal Edge Routers\) on page 970](#)
 [Syntax \(MX2020 3D Universal Edge Routers\) on page 970](#)
 [Syntax \(MX104 3D Universal Edge Routers\) on page 970](#)
 [Syntax \(PTX Series Packet Transport Routers\) on page 970](#)
 [Syntax \(T Series Routers\) on page 971](#)
 [Syntax \(TX Matrix Routers\) on page 971](#)
 [Syntax \(TX Matrix Plus Routers\) on page 971](#)

Syntax **show chassis routing-engine**
 <bios | slot>

Syntax (ACX Series Universal Access Routers) **show chassis routing-engine**

Syntax (EX Series Switches) **show chassis routing-engine**
 <slot>
 <satellite [slot-id *slot-id* |device-alias *alias-name*>

Syntax (QFX Series) **show chassis routing-engine**
 <interconnect-device *name*>
 <node-device *name*>

Syntax (MX Series Routers) **show chassis routing-engine**
 <all-members>
 <bios | slot>
 <local>
 <member *member-id*>
 <satellite [slot-id *slot-id* |device-alias *alias-name*>

Syntax (MX2010 3D Universal Edge Routers) **show chassis routing-engine**
 <bios | slot>

Syntax (MX2020 3D Universal Edge Routers) **show chassis routing-engine**
 <bios | slot>

Syntax (MX104 3D Universal Edge Routers) **show chassis routing-engine**

Syntax (PTX Series Packet Transport Routers) **show chassis routing-engine**

Syntax (T Series Routers)	show chassis routing-engine <bios <i>slot</i> >
Syntax (TX Matrix Routers)	show chassis routing-engine <bios <i>slot</i> > <lcc <i>number</i> scc>
Syntax (TX Matrix Plus Routers)	show chassis routing-engine <bios <i>slot</i> > <lcc <i>number</i> sfc <i>number</i> >
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>sfc option introduced for the TX Matrix Plus router in Junos OS Release in 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for QFX Series.</p> <p>Command introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 13.2 for MX104 3D Univesral Edge Routers.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> <p>satellite option introduced in Junos OS Release 14.2R3.</p>
Description	Display the status of the Routing Engine.
Options	<p>none—Display information about one or more Routing Engines. On a TX Matrix router, display information about all Routing Engines on the TX Matrix router and its attached T640 routers. On a TX Matrix Plus router, display information about all Routing Engines on the TX Matrix Plus router and its attached routers.</p> <p>all-members—(MX Series routers only) (Optional) Display Routing Engine information for all members of the Virtual Chassis configuration.</p> <p>bios—(Optional) Display the (BIOS) firmware version.</p> <p>interconnect-device <i>number</i>—(QFabric systems only) (Optional) Display Routing Engine information for a specified Interconnect device.</p> <p>lcc <i>number</i>—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display Routing Engine information for a specified T640 router (line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, display Routing Engine information for a specified router (line-card chassis) that is connected to the TX Matrix Plus router.</p> <p>Replace <i>number</i> with the following values depending on the LCC configuration:</p> <ul style="list-style-type: none"> • 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix. • 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.

- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(MX Series routers only) (Optional) Display Routing Engine information for the local Virtual Chassis member.

member *member-id*—(MX Series routers only) (Optional) Display Routing Engine information for the specified member of the Virtual Chassis configuration. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

node-device *number*—(QFabric systems only) (Optional) Display Routing Engine information for a specified Node device.

satellite [*slot-id slot-id* [*device-alias alias-name*]]—(Junos Fusion only) (Optional) Display Routing Engine information for the specified satellite device in a Junos Fusion, or for all satellite devices in the Junos Fusion if no satellite devices are specified.

scc—(TX Matrix routers only) (Optional) Display Routing Engine information for the TX Matrix router (switch-card chassis).

sfc *number*—(TX Matrix Plus routers only) (Optional) Display Routing Engine information for the TX Matrix Plus router (or switch-fabric chassis). Replace *number* with 0.

slot—(Systems with multiple Routing Engines) (Optional) Display information for an individual Routing Engine. Replace *slot* with 0 or 1. For QFX3500 switches, there is only one Routing Engine, so you do not need to specify the slot number.

Required Privilege Level

view

Related Documentation

- [request chassis routing-engine master on page 376](#)
- *Configuring Routing Engine Redundancy*
- *Switching the Global Master and Backup Roles in a Virtual Chassis Configuration*

List of Sample Output

[show chassis routing-engine \(M5 Router\) on page 975](#)
[show chassis routing-engine \(M10 Router\) on page 976](#)
[show chassis routing-engine \(M20 Router\) on page 976](#)
[show chassis routing-engine \(M40 Router\) on page 977](#)
[show chassis routing-engine \(M120 Router\) on page 977](#)
[show chassis routing-engine \(M160 Router\) on page 978](#)
[show chassis routing-engine \(MX104 Router\) on page 979](#)
[show chassis routing-engine \(MX240 Router\) on page 979](#)
[show chassis routing-engine \(MX480 Router\) on page 980](#)
[show chassis routing-engine \(MX960 Router\) on page 981](#)
[show chassis routing-engine \(MX2010 Router\) on page 983](#)
[show chassis routing-engine \(MX2020 Router\) on page 983](#)
[show chassis routing-engine \(T320 Router\) on page 984](#)

[show chassis routing-engine \(T640 Router\) on page 985](#)
[show chassis routing-engine \(T1600 Router\) on page 986](#)
[show chassis routing-engine \(T4000 Router\) on page 986](#)
[show chassis routing-engine \(TX Matrix Router\) on page 987](#)
[show chassis routing-engine lcc \(TX Matrix Router\) on page 988](#)
[show chassis routing-engine bios \(TX Matrix Router\) on page 989](#)
[show chassis routing-engine \(TX Matrix Plus Router\) on page 989](#)
[show chassis routing-engine lcc \(TX Matrix Plus Router\) on page 990](#)
[show chassis routing-engine bios \(TX Matrix Plus Router\) on page 991](#)
[show chassis routing-engine \(QFX Series\) on page 991](#)
[show chassis routing-engine \(OCX Series\) on page 992](#)
[show chassis routing engine interconnect-device \(QFabric Systems\) on page 992](#)
[show chassis routing-engine \(PTX Series Packet Transport Router\) on page 993](#)
[show chassis routing-engine \(EX9200 Switch\) on page 993](#)
[show chassis routing-engine \(ACX2000 Universal Access Router\) on page 994](#)
[show chassis routing-engine \(ACX1000 Universal Access Router\) on page 994](#)

Output Fields Table 66 lists the output fields for the **show chassis routing-engine** command. Output fields are listed in the approximate order in which they appear.

Table 66: show chassis routing-engine Output Fields

Field Name	Field Description
Slot	(Systems with single and multiple Routing Engines) Slot number.
Current state	(Systems with multiple Routing Engines) Current state of the Routing Engine: Master , Backup , or Disabled .
Election priority	(Systems with multiple Routing Engines) Election priority for the Routing Engine: Master or Backup .
Temperature	Temperature of the air flowing past the Routing Engine.
CPU Temperature	Temperature of the CPU.
DRAM	<p>Total DRAM available to the Routing Engine's processor.</p> <p>Starting with Junos OS Release 12.3R1, the DRAM field displays both available memory and installed memory.</p> <p>NOTE: For platforms running Junos OS with upgraded FreeBSD, the way memory utilization is calculated has changed. Starting in Junos OS Release 15.1R1, inactive memory is no longer included in the calculation for memory utilization. For platforms that run Junos OS with upgraded FreeBSD, see "Understanding Junos OS with Upgraded FreeBSD" on page 89.</p>
Memory utilization	Percentage of Routing Engine memory being used.

Table 66: show chassis routing-engine Output Fields (*continued*)

Field Name	Field Description
CPU utilization	Information about the Routing Engine's CPU utilization: <ul style="list-style-type: none"> • User—Percentage of CPU time being used by user processes. • Background—Percentage of CPU time being used by background processes. • Kernel—Percentage of CPU time being used by kernel processes. • Interrupt—Percentage of CPU time being used by interrupts. • Idle—Percentage of CPU time that is idle.
5 sec CPU Utilization NOTE: Supported only on MX240, MX480, MX960, MX2010, and MX2020.	Information about the Routing Engine's CPU utilization in the past 5 seconds: <ul style="list-style-type: none"> • User—Percentage of CPU time being used by user processes. • Background—Percentage of CPU time being used by background processes. • Kernel—Percentage of CPU time being used by kernel processes. • Interrupt—Percentage of CPU time being used by interrupts. • Idle—Percentage of CPU time that is idle.
1 min CPU Utilization NOTE: Supported only on MX240, MX480, MX960, MX2010, and MX2020.	Information about the Routing Engine's CPU utilization in the past 1 minute: <ul style="list-style-type: none"> • User—Percentage of CPU time being used by user processes. • Background—Percentage of CPU time being used by background processes. • Kernel—Percentage of CPU time being used by kernel processes. • Interrupt—Percentage of CPU time being used by interrupts. • Idle—Percentage of CPU time that is idle.
5 min CPU Utilization NOTE: Supported only on MX240, MX480, MX960, MX2010, and MX2020.	Information about the Routing Engine's CPU utilization in the past 5 minutes: <ul style="list-style-type: none"> • User—Percentage of CPU time being used by user processes. • Background—Percentage of CPU time being used by background processes. • Kernel—Percentage of CPU time being used by kernel processes. • Interrupt—Percentage of CPU time being used by interrupts. • Idle—Percentage of CPU time that is idle.
15 min CPU Utilization NOTE: Supported only on MX240, MX480, MX960, MX2010, and MX2020.	Information about the Routing Engine's CPU utilization in the past 15 minutes: <ul style="list-style-type: none"> • User—Percentage of CPU time being used by user processes. • Background—Percentage of CPU time being used by background processes. • Kernel—Percentage of CPU time being used by kernel processes. • Interrupt—Percentage of CPU time being used by interrupts. • Idle—Percentage of CPU time that is idle.
Model	Routing Engine model number.
Serial ID	(Systems with multiple Routing Engines) Identification number of the Routing Engine in this slot.
Start time	Time at which the Routing Engine started running.
Uptime	How long the Routing Engine has been running.

Table 66: show chassis routing-engine Output Fields (*continued*)

Field Name	Field Description
Routing Engine BIOS Version	BIOS version being run by the Routing Engine.
Last reboot reason	Reason for last reboot, including: <ul style="list-style-type: none"> • power cycle/failure—Halt of the Routing Engine using the halt command, powering down using the power button on the chassis or any other method (such as removal of the control board or Routing Engine), and then powering back the Routing Engine. A halt of the operating system also occurs if you enter the request system halt command. You can enter this command to halt the system operations on the chassis or specific Routing Engines. To restart the software, press any key on the keyboard. • watchdog—Reboot due to a hardware watchdog. A watchdog is a hardware monitoring process that examines the health and performance of the router to enable the device to recover from failures. A watchdog checks for problems at certain intervals, and reboots the routing engine if a problem is encountered. • reset-button reset—(Not available on the EX Series switch) Reboot due to pressing of the reset button on the Routing Engine. • power-button hard power off—Reboot due to pressing of the power button on the chassis. A powering down of the software also occurs if you enter the request system power-off command. You can enter this command to power down the chassis or specific Routing Engines; you can then restart the software. • misc hardware reason—Reboot due to miscellaneous hardware reasons. • thermal shutdown—Reboot due to the router or switch reaching a critical temperature at which point it is unsafe to continue operations. • hard disk failure—Reboot due to a hard disk or solid-state drive (SSD) failure. • reset from debugger—Reboot due to reset from the debugger. • chassis control reset—Restart the chassis process that manages PICs, FPCs, and other hardware components. The chassis control module that runs the Routing Engine performs management and monitoring functions, and it provides a single access point for operational and maintenance functions. A reset of the chassis management process occurs when you enter the restart chassis-control command. • bios auto recovery reset—Reboot due to a BIOS auto-recovery reset. • could not be determined—Reboot due to an undetermined reason. • Router rebooted after a normal shutdown—Reboot due to a normal shutdown. This reason is displayed if the Routing Engine is powered down by pushing and holding the online/offline button on the Routing Engine faceplate for 30 seconds, and then powered back. A reboot of the software also occurs if you enter the request system reboot command. You can enter this command to reboot the chassis or specific Routing Engines.
Load averages	Routing Engine load averages for the last 1, 5, and 15 minutes.

Sample Output

show chassis routing-engine (M5 Router)

```

user@host> show chassis routing-engine
Routing Engine status:
  Temperature                25 degrees C / 77 degrees F
  DRAM                       768 MB
  Memory utilization          21 percent
  CPU utilization:

```

```

User                0 percent
Background          0 percent
Kernel              0 percent
Interrupt            0 percent
Idle                100 percent
Model               RE-2.0
Serial ID           31000007349bf701
Start time          2003-12-04 09:42:17 PST
Uptime              26 days, 1 hour, 12 minutes, 27 seconds
Last reboot reason   Router rebooted after a normal shutdown
Load averages:      1 minute   5 minute   15 minute
                    0.00       0.01       0.00

```

show chassis routing-engine (M10 Router)

```

user@host> show chassis routing-engine
Routing Engine status:
  Temperature        25 degrees C / 77 degrees F
  DRAM               768 MB
  Memory utilization  21 percent
  CPU utilization:
    User              0 percent
    Background        0 percent
    Kernel             0 percent
    Interrupt          0 percent
    Idle              100 percent
  Model              RE-2.0
  Serial ID          31000007349bf701
  Start time         2003-12-04 09:42:17 PST
  Uptime             26 days, 1 hour, 12 minutes, 27 seconds
  Last reboot reason  Router rebooted after a normal shutdown
  Load averages:    1 minute   5 minute   15 minute
                    0.00       0.01       0.00

```

show chassis routing-engine (M20 Router)

```

user@host> show chassis routing-engine
Routing Engine status:
  Slot 0:
    Current state      Master
    Election priority  Master (default)
    Temperature        29 degrees C / 84 degrees F
    DRAM               768 MB
    Memory utilization  20 percent
    CPU utilization:
      User              1 percent
      Background        0 percent
      Kernel             2 percent
      Interrupt          0 percent
      Idle              97 percent
    Model              RE-2.0
    Serial ID          58000007348d9a01
    Start time         2003-12-30 07:05:47 PST
    Uptime             3 hours, 41 minutes, 14 seconds
    Last reboot reason  Router rebooted after a normal shutdown
    Load averages:    1 minute   5 minute   15 minute
                      0.00       0.02       0.00
  Routing Engine status:
    Slot 1:
      Current state      Backup
      Election priority  Backup (default)

```



```

Temperature                29 degrees C / 84 degrees F
DRAM                      768 MB
Memory utilization         0 percent
CPU utilization:
  User                    0 percent
  Background              0 percent
  Kernel                  1 percent
  Interrupt               0 percent
  Idle                    99 percent
Model                     RE-2.0
Serial ID                  d800000734745701
Start time                 2003-06-17 16:37:33 PDT
Uptime                     195 days, 18 hours, 47 minutes, 9 seconds
Last reboot reason        Router rebooted after a normal shutdown

```

show chassis routing-engine (M40 Router)

```

user@host> show chassis routing-engine
Routing Engine status:
  Temperature                25 degrees C / 77 degrees F
  DRAM                      768 MB
  Memory utilization         21 percent
  CPU utilization:
    User                    0 percent
    Background              0 percent
    Kernel                  0 percent
    Interrupt               0 percent
    Idle                    100 percent
  Model                     RE-2.0
  Serial ID                  31000007349bf701
  Start time                 2003-12-04 09:42:17 PST
  Uptime                     26 days, 1 hour, 12 minutes, 27 seconds
  Last reboot reason        Router rebooted after a normal shutdown
  Load averages:           1 minute   5 minute   15 minute
                           0.00        0.01        0.00

```

show chassis routing-engine (M120 Router)

```

user@host> show chassis routing-engine
Routing Engine status:
  Slot 0:
    Current state            Master
    Election priority        Master (default)
    Temperature              46 degrees C / 114 degrees F
    CPU temperature          44 degrees C / 111 degrees F
    DRAM                     2048 MB
    Memory utilization        18 percent
    CPU utilization:
      User                    0 percent
      Background              0 percent
      Kernel                  5 percent
      Interrupt               0 percent
      Idle                    95 percent
    Model                    RE-A-1000
    Serial ID                 1000621154
    Start time                2006-10-31 17:10:05 PST
    Uptime                    14 minutes, 31 seconds
    Last reboot reason        Router rebooted after a normal shutdown
    Load averages:           1 minute   5 minute   15 minute
                           0.02        0.07        0.07

Routing Engine status:

```

```

Slot 1:
  Current state          Backup
  Election priority      Backup (default)
  Temperature            45 degrees C / 113 degrees F
  CPU temperature        42 degrees C / 107 degrees F
  DRAM                   2048 MB
  Memory utilization     15 percent
  CPU utilization:
    User                 0 percent
    Background           0 percent
    Kernel               0 percent
    Interrupt            0 percent
    Idle                 100 percent
  Model                  RE-A-1000
  Serial ID              1000621151
  Start time             2006-10-31 17:10:04 PST
  Uptime                 14 minutes, 30 seconds
  Last reboot reason     Router rebooted after a normal shutdown

```

show chassis routing-engine (M160 Router)

```

user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state          Master
  Election priority      Master (default)
  Temperature            43 degrees C / 109 degrees F
  DRAM                   2048 MB
  Memory utilization     11 percent
  CPU utilization:
    User                 1 percent
    Background           0 percent
    Kernel               2 percent
    Interrupt            0 percent
    Idle                 97 percent
  Model                  RE-3.0
  Serial ID              210865700403
  Start time             2003-12-23 12:25:55 PST
  Uptime                 6 days, 22 hours, 33 minutes, 24 seconds
  Last reboot reason     Router rebooted after a normal shutdown
  Load averages:        1 minute   5 minute   15 minute
                        0.24       0.13       0.04

Routing Engine status:
Slot 1:
  Current state          Backup
  Election priority      Backup (default)
  Temperature            40 degrees C / 104 degrees F
  DRAM                   2048 MB
  Memory utilization     9 percent
  CPU utilization:
    User                 0 percent
    Background           0 percent
    Kernel               0 percent
    Interrupt            0 percent
    Idle                 100 percent
  Model                  RE-3.0
  Serial ID              210865700332
  Start time             2003-12-23 12:25:55 PST
  Uptime                 6 days, 22 hours, 33 minutes, 21 seconds
  Last reboot reason     Router rebooted after a normal shutdown

```

show chassis routing-engine (MX104 Router)

```

user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state           Master
  Election priority       Master (default)
  Temperature             32 degrees C / 89 degrees F
  CPU temperature         42 degrees C / 107 degrees F
  DRAM                   3840 MB (3840 MB installed)
  Memory utilization      18 percent
  CPU utilization:
    User                  0 percent
    Background            0 percent
    Kernel                3 percent
    Interrupt             2 percent
    Idle                  94 percent
  Model                   RE-MX-104
  Serial ID               CAAR5925
  Start time              2013-06-05 13:17:08 IST
  Uptime                  1 hour, 15 minutes, 8 seconds
  Last reboot reason      0x200:normal shutdown
  Load averages:         1 minute   5 minute   15 minute
                        0.87       0.90       0.41

Routing Engine status:
Slot 1:
  Current state           Backup
  Election priority       Backup (default)
  Temperature             32 degrees C / 89 degrees F
  CPU temperature         38 degrees C / 100 degrees F
  DRAM                   3840 MB (3840 MB installed)
  Memory utilization      13 percent
  CPU utilization:
    User                  0 percent
    Background            0 percent
    Kernel                1 percent
    Interrupt             2 percent
    Idle                  97 percent
  Model                   RE-MX-104
  Serial ID               CAAM6369
  Start time              2013-06-05 13:07:37 IST
  Uptime                  1 hour, 24 minutes, 34 seconds
  Last reboot reason      0x200:normal shutdown
  Load averages:         1 minute   5 minute   15 minute
                        0.19       0.15       0.06

```

show chassis routing-engine (MX240 Router)

```

user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state           Master
  Election priority       Master (default)
  Temperature             36 degrees C / 96 degrees F
  CPU temperature         35 degrees C / 95 degrees F
  DRAM                   3314 MB (8192 MB installed)
  Memory utilization      37 percent
  5 sec CPU utilization:
    User                  0 percent
    Background            0 percent
    Kernel                1 percent

```

```

        Interrupt          0 percent
        Idle              99 percent
    1 min CPU utilization:
        User              0 percent
        Background        0 percent
        Kernel            1 percent
        Interrupt          0 percent
        Idle              99 percent
    5 min CPU utilization:
        User              0 percent
        Background        0 percent
        Kernel            1 percent
        Interrupt          0 percent
        Idle              99 percent
    15 min CPU utilization:
        User              0 percent
        Background        0 percent
        Kernel            1 percent
        Interrupt          0 percent
        Idle              99 percent
    Model                  RE-S-1800x4
    Serial ID              9009074155
    Start time              2014-10-13 00:35:41 PDT
    Uptime                  98 days, 2 hours, 6 minutes, 35 seconds
    Last reboot reason      Router rebooted after a normal shutdown.
    Load averages:         1 minute   5 minute   15 minute
                           0.12       0.12       0.13

Routing Engine status:
Slot 1:
    Current state          Present

```

show chassis routing-engine (MX480 Router)

```

user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
    Current state          Backup
    Election priority      Master (default)
    Temperature            30 degrees C / 86 degrees F
    CPU temperature        32 degrees C / 89 degrees F
    DRAM                   3314 MB (8192 MB installed)
    Memory utilization      51 percent
    5 sec CPU utilization:
        User              0 percent
        Background        0 percent
        Kernel            0 percent
        Interrupt          0 percent
        Idle              100 percent
    1 min CPU utilization:
        User              0 percent
        Background        0 percent
        Kernel            0 percent
        Interrupt          0 percent
        Idle              0 percent
    5 min CPU utilization:
        User              0 percent
        Background        0 percent
        Kernel            0 percent
        Interrupt          0 percent
        Idle              0 percent
    15 min CPU utilization:

```

```

User                0 percent
Background           0 percent
Kernel              0 percent
Interrupt            0 percent
Idle                0 percent
Model               RE-S-1800x4
Serial ID            9009079817
Start time           2015-01-19 01:45:58 PST
Uptime              7 minutes, 23 seconds
Last reboot reason   Router rebooted after a normal shutdown.
Load averages:       1 minute   5 minute   15 minute
                      0.16      0.16      0.09

Routing Engine status:
Slot 1:
  Current state      Master
  Election priority   Backup (default)
  Temperature        31 degrees C / 87 degrees F
  CPU temperature     32 degrees C / 89 degrees F
  DRAM               8144 MB (8192 MB installed)
  Memory utilization  23 percent
  5 sec CPU utilization:
    User             0 percent
    Background        0 percent
    Kernel            1 percent
    Interrupt         0 percent
    Idle             99 percent
  1 min CPU utilization:
    User             0 percent
    Background        0 percent
    Kernel            1 percent
    Interrupt         0 percent
    Idle             98 percent
  5 min CPU utilization:
    User             0 percent
    Background        0 percent
    Kernel            1 percent
    Interrupt         0 percent
    Idle             98 percent
  15 min CPU utilization:
    User             0 percent
    Background        0 percent
    Kernel            1 percent
    Interrupt         0 percent
    Idle             98 percent
  Model              RE-S-1800x4
  Serial ID           9009079838
  Start time           2015-01-09 10:52:20 PST
  Uptime              9 days, 15 hours, 1 minute, 4 seconds
  Last reboot reason   Router rebooted after a normal shutdown.
  Load averages:       1 minute   5 minute   15 minute
                        0.10      0.16      0.16

```

show chassis routing-engine (MX960 Router)

```

user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state      Master
  Election priority   Master (default)
  Temperature        37 degrees C / 98 degrees F
  CPU temperature     34 degrees C / 93 degrees F

```

```

DRAM                                     3313 MB (16384 MB installed)
Memory utilization                       31 percent
5 sec CPU utilization:
  User                                  0 percent
  Background                           0 percent
  Kernel                               3 percent
  Interrupt                             1 percent
  Idle                                  96 percent
1 min CPU utilization:
  User                                  0 percent
  Background                           0 percent
  Kernel                               4 percent
  Interrupt                             1 percent
  Idle                                  96 percent
5 min CPU utilization:
  User                                  0 percent
  Background                           0 percent
  Kernel                               4 percent
  Interrupt                             1 percent
  Idle                                  95 percent
15 min CPU utilization:
  User                                  0 percent
  Background                           0 percent
  Kernel                               4 percent
  Interrupt                             1 percent
  Idle                                  95 percent
Model                                    RE-S-1800x4
Serial ID                               9013043785
Start time                             2015-01-12 23:37:53 PST
Uptime                                  6 days, 2 hours, 17 minutes, 3 seconds
Last reboot reason                      Router rebooted after a normal shutdown.
Load averages:                         1 minute   5 minute   15 minute
                                         0.00       0.02       0.00

Routing Engine status:
Slot 1:
  Current state                         Backup
  Election priority                     Backup (default)
  Temperature                           37 degrees C / 98 degrees F
  CPU temperature                       34 degrees C / 93 degrees F
  DRAM                                  3313 MB (16384 MB installed)
  Memory utilization                    26 percent
  5 sec CPU utilization:
    User                                0 percent
    Background                          0 percent
    Kernel                              0 percent
    Interrupt                           0 percent
    Idle                                99 percent
  1 min CPU utilization:
    User                                0 percent
    Background                          0 percent
    Kernel                              0 percent
    Interrupt                           0 percent
    Idle                                0 percent
  5 min CPU utilization:
    User                                0 percent
    Background                          0 percent
    Kernel                              0 percent
    Interrupt                           0 percent
    Idle                                0 percent
  15 min CPU utilization:
    User                                0 percent

```

```

Background          0 percent
Kernel              0 percent
Interrupt            0 percent
Idle                 0 percent
Model                RE-S-1800x4
Serial ID            9013037303
Start time           2015-01-12 23:25:29 PST
Uptime               6 days, 2 hours, 29 minutes, 21 seconds
Last reboot reason   Router rebooted after a normal shutdown.
Load averages:       1 minute   5 minute   15 minute
                      0.00       0.00       0.00

```

show chassis routing-engine (MX2010 Router)

```

user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state          Master
  Election priority      Master (default)
  Temperature            41 degrees C / 105 degrees F
  CPU temperature        38 degrees C / 100 degrees F
  DRAM                   3313 MB (16384 MB installed)
  Memory utilization     37 percent
  5 sec CPU utilization:
    User                 0 percent
    Background           0 percent
    Kernel               2 percent
    Interrupt            2 percent
    Idle                 96 percent
  1 min CPU utilization:
    User                 0 percent
    Background           0 percent
    Kernel               2 percent
    Interrupt            2 percent
    Idle                 97 percent
  5 min CPU utilization:
    User                 0 percent
    Background           0 percent
    Kernel               2 percent
    Interrupt            2 percent
    Idle                 97 percent
  15 min CPU utilization:
    User                 0 percent
    Background           0 percent
    Kernel               2 percent
    Interrupt            2 percent
    Idle                 97 percent
  Model                  RE-S-1800x4
  Serial ID              9009146890
  Start time             2015-01-18 21:35:12 PST
  Uptime                 4 hours, 21 minutes, 34 seconds
  Last reboot reason     Router rebooted after a normal shutdown.
  Load averages:        1 minute   5 minute   15 minute
                        0.11       0.14       0.14

```

show chassis routing-engine (MX2020 Router)

```

user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state          Master

```

```

Election priority           Master (default)
Temperature                 2 degrees C / 35 degrees F
CPU temperature             32 degrees C / 89 degrees F
DRAM                       32735 MB (32768 MB installed)
Memory utilization         10 percent
5 sec CPU utilization:
  User                      0 percent
  Background                0 percent
  Kernel                    1 percent
  Interrupt                 1 percent
  Idle                      98 percent
1 min CPU utilization:
  User                      0 percent
  Background                0 percent
  Kernel                    1 percent
  Interrupt                 1 percent
  Idle                      99 percent
5 min CPU utilization:
  User                      0 percent
  Background                0 percent
  Kernel                    1 percent
  Interrupt                 1 percent
  Idle                      99 percent
15 min CPU utilization:
  User                      0 percent
  Background                0 percent
  Kernel                    1 percent
  Interrupt                 1 percent
  Idle                      99 percent
Model                       RE-S-2X00x8
Serial ID                   CADN0309
Start time                  2015-01-08 16:31:15 PST
Uptime                      4 days, 22 hours, 59 minutes, 3 seconds
Last reboot reason          Router rebooted after a normal shutdown.
Load averages:              1 minute   5 minute   15 minute
                             0.39       0.41       0.34

```

show chassis routing-engine (T320 Router)

```

user@host> show chassis routing-engine
Slot 0:
  Current state              Master
  Election priority          Master (default)
  Temperature                 51 degrees C / 123 degrees F
  CPU temperature             55 degrees C / 131 degrees F
  DRAM                       3584 MB
  Memory utilization         11 percent
  CPU utilization:
    User                      0 percent
    Background                0 percent
    Kernel                    2 percent
    Interrupt                 0 percent
    Idle                      97 percent
  Model                       RE-A-2000
  Serial ID                   9009010618
  Start time                  2012-10-10 01:24:05 PDT
  Uptime                      5 days, 10 hours, 49 minutes, 23 seconds
  Last reboot reason          0x1:power cycle/failure
  Load averages:              1 minute   5 minute   15 minute
                             0.00       0.05       0.04

Routing Engine status:

```



```

Slot 1:
  Current state           Backup
  Election priority       Backup (default)
  Temperature             45 degrees C / 113 degrees F
  CPU temperature         48 degrees C / 118 degrees F
  DRAM                   3584 MB
  Memory utilization      9 percent
  CPU utilization:
    User                  0 percent
    Background            0 percent
    Kernel                0 percent
    Interrupt             0 percent
    Idle                  100 percent
  Model                  RE-A-2000
  Serial ID              9009003642
  Start time             2012-10-10 01:24:04 PDT
  Uptime                 5 days, 10 hours, 49 minutes, 28 seconds
  Last reboot reason     0x1:power cycle/failure

```

show chassis routing-engine (T640 Router)

```

user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state           Master
  Election priority       Master (default)
  Temperature             50 degrees C / 122 degrees F
  CPU temperature         58 degrees C / 136 degrees F
  DRAM                   3584 MB
  Memory utilization      14 percent
  CPU utilization:
    User                  1 percent
    Background            0 percent
    Kernel                4 percent
    Interrupt             1 percent
    Idle                  95 percent
  Model                  RE-A-2000
  Serial ID              1000686556
  Start time             2012-10-10 01:24:02 PDT
  Uptime                 5 days, 10 hours, 50 minutes, 27 seconds
  Last reboot reason     0x1:power cycle/failure
  Load averages:        1 minute 5 minute 15 minute
                        1.24      0.33      0.12

Routing Engine status:
Slot 1:
  Current state           Backup
  Election priority       Backup (default)
  Temperature             44 degrees C / 111 degrees F
  CPU temperature         49 degrees C / 120 degrees F
  DRAM                   3584 MB
  Memory utilization      12 percent
  CPU utilization:
    User                  0 percent
    Background            0 percent
    Kernel                0 percent
    Interrupt             1 percent
    Idle                  99 percent
  Model                  RE-A-2000
  Serial ID              1000702739
  Start time             2012-10-10 01:24:02 PDT

```

Uptime	5 days, 10 hours, 50 minutes, 26 seconds
Last reboot reason	0x1:power cycle/failure

show chassis routing-engine (T1600 Router)

```
user@host> show chassis routing-engine
```

```
Routing Engine status:
```

```
Slot 0:
```

Current state	Master
Election priority	Master (default)
Temperature	48 degrees C / 118 degrees F
CPU temperature	58 degrees C / 136 degrees F
DRAM	3584 MB
Memory utilization	13 percent
CPU utilization:	
User	0 percent
Background	0 percent
Kernel	3 percent
Interrupt	1 percent
Idle	96 percent
Model	RE-A-2000
Serial ID	1000704521
Start time	2012-10-10 01:23:41 PDT
Uptime	5 days, 10 hours, 46 minutes, 56 seconds
Last reboot reason	0x1:power cycle/failure
Load averages:	1 minute 5 minute 15 minute
	0.05 0.03 0.01

```
Routing Engine status:
```

```
Slot 1:
```

Current state	Backup
Election priority	Backup (default)
Temperature	44 degrees C / 111 degrees F
CPU temperature	48 degrees C / 118 degrees F
DRAM	3584 MB
Memory utilization	12 percent
CPU utilization:	
User	0 percent
Background	0 percent
Kernel	0 percent
Interrupt	0 percent
Idle	100 percent
Model	RE-A-2000
Serial ID	9009006579
Start time	2012-10-10 01:23:42 PDT
Uptime	5 days, 10 hours, 46 minutes, 54 seconds
Last reboot reason	0x1:power cycle/failure

show chassis routing-engine (T4000 Router)

```
user@host> show chassis routing-engine
```

```
Routing Engine status:
```

```
Slot 0:
```

Current state	Master
Election priority	Master (default)
Temperature	33 degrees C / 91 degrees F
CPU temperature	50 degrees C / 122 degrees F
DRAM	8960 MB
Memory utilization	18 percent
CPU utilization:	
User	0 percent
Background	0 percent

```

Kernel                4 percent
Interrupt             1 percent
Idle                  95 percent
Model                 RE-DUO-1800
Serial ID              P737F-002248
Start time             2012-02-09 22:49:53 PST
Uptime                2 hours, 21 minutes, 35 seconds
Last reboot reason    Router rebooted after a normal shutdown.
Load averages:        1 minute   5 minute   15 minute
                       0.00       0.04       0.00

Routing Engine status:
Slot 1:
  Current state        Backup
  Election priority    Backup (default)
  Temperature          32 degrees C / 89 degrees F
  CPU temperature      46 degrees C / 114 degrees F
  DRAM                 8960 MB
  Memory utilization   24 percent
  CPU utilization:
    User               0 percent
    Background         0 percent
    Kernel              0 percent
    Interrupt          0 percent
    Idle               99 percent
  Model               RE-DUO-1800
  Serial ID           P737F-002653
  Start time          2012-02-08 20:12:51 PST
  Uptime              1 day, 4 hours, 58 minutes, 28 seconds
  Last reboot reason  Router rebooted after a normal shutdown.

```

show chassis routing-engine (TX Matrix Router)

```

user@host> show chassis routing-engine
scc-re0:
-----
Routing Engine status:
Slot 0:
  Current state        Master
  Election priority    Master (default)
  Temperature          34 degrees C / 93 degrees F
  CPU temperature      33 degrees C / 91 degrees F
  DRAM                 2048 MB
  Memory utilization   12 percent
  CPU utilization:
    User               0 percent
    Background         0 percent
    Kernel              2 percent
    Interrupt          0 percent
    Idle               98 percent
  Model               RE-4.0
  Serial ID           P11123900153
  Start time          2004-08-05 18:42:05 PDT
  Uptime              9 days, 22 hours, 49 minutes, 50 seconds
  Last reboot reason  Router rebooted after a normal shutdown
  Load averages:      1 minute   5 minute   15 minute
                       0.00       0.08       0.07

lcc0-re0:
-----
Routing Engine status:
Slot 0:

```

```

Current state           Master
Election priority       Master (default)
Temperature             33 degrees C / 91 degrees F
CPU temperature         30 degrees C / 86 degrees F
DRAM                   2048 MB
Memory utilization      12 percent
CPU utilization:
  User                  0 percent
  Background            0 percent
  Kernel                1 percent
  Interrupt             0 percent
  Idle                  98 percent
Model                  RE-3.0
Serial ID               210865700363
Start time              2004-08-05 18:42:05 PDT
Uptime                  9 days, 22 hours, 48 minutes, 20 seconds
Last reboot reason      Router rebooted after a normal shutdown
Load averages:          1 minute   5 minute  15 minute
                       0.00        0.02    0.00

```

```
lcc2-re0:
```

```
-----
Routing Engine status:
```

```
Slot 0:
```

```

Current state           Master
Election priority       Master (default)
Temperature             34 degrees C / 93 degrees F
CPU temperature         35 degrees C / 95 degrees F
DRAM                   2048 MB
Memory utilization      12 percent
CPU utilization:
  User                  0 percent
  Background            0 percent
  Kernel                2 percent
  Interrupt             0 percent
  Idle                  98 percent
Model                  RE-4.0
Serial ID               P11123900126
Start time              2004-08-05 18:42:05 PDT
Uptime                  9 days, 22 hours, 49 minutes, 4 seconds
Last reboot reason      Router rebooted after a normal shutdown
Load averages:          1 minute   5 minute  15 minute
                       0.01        0.01    0.0

```

show chassis routing-engine lcc (TX Matrix Router)

```
user@host> show chassis routing-engine 0 lcc 0
```

```
lcc0-re0:
```

```
-----
Routing Engine status:
```

```
Slot 0:
```

```

Current state           Master
Election priority       Master (default)
Temperature             33 degrees C / 91 degrees F
CPU temperature         30 degrees C / 86 degrees F
DRAM                   2048 MB
Memory utilization      12 percent
CPU utilization:
  User                  0 percent
  Background            0 percent
  Kernel                1 percent

```

```

        Interrupt          0 percent
        Idle              98 percent
        Model             RE-3.0
        Serial ID         210865700363
        Start time        2004-08-05 18:42:05 PDT
        Uptime            7 days, 22 hours, 49 minutes, 6 seconds
        Last reboot reason Router rebooted after a normal shutdown
        Load averages:    1 minute   5 minute   15 minute
                           0.00       0.00       0.00

```

show chassis routing-engine bios (TX Matrix Router)

```
user@host> show chassis routing-engine bios
```

```
scc-re0:
```

```
-----
Routing Engine BIOS Version: V1.0.0
```

```
lcc0-re0:
```

```
-----
Routing Engine BIOS Version: V1.0.17
```

```
lcc2-re0:
```

```
-----
Routing Engine BIOS Version: V1.0.0
```

show chassis routing-engine (TX Matrix Plus Router)

```
user@host> show chassis routing-engine
```

```
sfc0-re0:
```

```
-----
Routing Engine status:
```

```
Slot 0:
```

```

        Current state      Master
        Election priority   Master (default)
        Temperature         27 degrees C / 80 degrees F
        CPU temperature     42 degrees C / 107 degrees F
        DRAM                3327 MB
        Memory utilization   12 percent
        CPU utilization:
            User            0 percent
            Background      0 percent
            Kernel          2 percent
            Interrupt       0 percent
            Idle            98 percent
        Model              RE-TXP-SFC
        Serial ID          737A-1024
        Start time         2009-05-11 17:39:49 PDT
        Uptime             3 hours, 45 minutes, 25 seconds
        Last reboot reason  Router rebooted after a normal shutdown.
        Load averages:     1 minute   5 minute   15 minute
                           0.00       0.00       0.00

```

```
Routing Engine status:
```

```
Slot 1:
```

```

        Current state      Backup
        Election priority   Backup (default)
        Temperature         29 degrees C / 84 degrees F
        CPU temperature     43 degrees C / 109 degrees F
        DRAM                3327 MB
        Memory utilization   11 percent
        CPU utilization:
            User            0 percent
            Background      0 percent
            Kernel          0 percent

```

```

        Interrupt          0 percent
        Idle              100 percent
        Model             RE-TXP-SFC
        Serial ID         737A-1024
        Start time        2009-05-11 17:08:54 PDT
        Uptime            4 hours, 16 minutes, 52 seconds
        Last reboot reason 0x1:power cycle/failure

lcc0-re0:
-----
Routing Engine status:
Slot 0:
  Current state          Master
  Election priority      Master (default)
  Temperature            30 degrees C / 86 degrees F
  CPU temperature        43 degrees C / 109 degrees F
  DRAM                   3327 MB
  Memory utilization     9 percent
  CPU utilization:
    User                 0 percent
    Background           0 percent
    Kernel               2 percent
    Interrupt            0 percent
    Idle                 98 percent
  Model                  RE-TXP-LCC
  Serial ID              737F-1024
  Start time             2009-05-11 17:40:32 PDT
  Uptime                 3 hours, 44 minutes, 51 seconds
  Last reboot reason     Router rebooted after a normal shutdown.
  Load averages:        1 minute  5 minute 15 minute
                        0.00      0.00   0.00

Routing Engine status:
Slot 1:
  Current state          Backup
  Election priority      Backup (default)
  Temperature            30 degrees C / 86 degrees F
  CPU temperature        43 degrees C / 109 degrees F
  DRAM                   3327 MB
  Memory utilization     9 percent
  CPU utilization:
    User                 0 percent
    Background           0 percent
    Kernel               0 percent
    Interrupt            0 percent
    Idle                 100 percent
  Model                  RE-TXP-LCC
  Serial ID              737F-1024
  Start time             2009-05-06 17:31:32 PDT
  Uptime                 5 days, 3 hours, 54 minutes, 19 seconds
  Last reboot reason     Router rebooted after a normal shutdown.

```

show chassis routing-engine lcc (TX Matrix Plus Router)

```

user@host> show chassis routing-engine 0 lcc 0
lcc0-re0:
-----
Routing Engine status:
Slot 0:
  Current state          Master
  Election priority      Master (default)
  Temperature            30 degrees C / 86 degrees F

```

```

CPU temperature          43 degrees C / 109 degrees F
DRAM                    3327 MB
Memory utilization       9 percent
CPU utilization:
  User                   0 percent
  Background             0 percent
  Kernel                 2 percent
  Interrupt              0 percent
  Idle                   98 percent
Model                   RE-TXP-LCC
Serial ID               737F-1024
Start time              2009-05-11 17:40:32 PDT
Uptime                  3 hours, 45 minutes, 26 seconds
Last reboot reason      Router rebooted after a normal shutdown.
Load averages:          1 minute   5 minute  15 minute
                        0.00       0.00    0.00

Routing Engine status:
Slot 1:
  Current state          Backup
  Election priority      Backup (default)
  Temperature            30 degrees C / 86 degrees F
  CPU temperature        43 degrees C / 109 degrees F
  DRAM                   3327 MB
  Memory utilization     9 percent
  CPU utilization:
    User                 0 percent
    Background           0 percent
    Kernel               0 percent
    Interrupt            0 percent
    Idle                 100 percent
  Model                  RE-TXP-LCC
  Serial ID              737F-1024
  Start time             2009-05-06 17:31:32 PDT
  Uptime                 5 days, 3 hours, 54 minutes, 59 seconds
  Last reboot reason     Router rebooted after a normal shutdown.

```

show chassis routing-engine bios (TX Matrix Plus Router)

```

user@host> show chassis routing-engine bios
sfc0-re0:

```

```

-----
Routing Engine BIOS Version: V0.0.Z

```

```

lcc0-re0:

```

```

-----
Routing Engine BIOS Version: V0.0.N

```

show chassis routing-engine (QFX Series)

```

user@switch> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state Master
  Election priority Master (default)
  DRAM 2820 MB
  Memory utilization 49 percent
  CPU utilization:
    User 1 percent
    Background 0 percent
    Kernel 1 percent
    Interrupt 0 percent

```

```
Idle 97 percent
Model QFX3500-48S4Q
Serial ID S/N ED3709
Uptime 3 days, 4 hours, 29 minutes, 42 seconds
Last reboot reason 0x200:chassis control reset
Load averages: 1 minute 5 minute 15 minute
0.37 0.26 0.19
```

show chassis routing-engine (OCX Series)

```
user@switch> show chassis routing-engine
Routing Engine status:
Slot 0:
Current state Master
Election priority Master (default)
DRAM 2820 MB
Memory utilization 49 percent
CPU utilization:
User 1 percent
Background 0 percent
Kernel 1 percent
Interrupt 0 percent
Idle 97 percent
Model OCX-1100-48SX-AFI
Serial ID S/N ED3709
Uptime 3 days, 4 hours, 29 minutes, 42 seconds
Last reboot reason 0x200:chassis control reset
Load averages: 1 minute 5 minute 15 minute
0.37 0.26 0.19
```

show chassis routing engine interconnect-device (QFabric Systems)

```
user@switch> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state           Master
  Election priority       Master (default)
  Temperature             48 degrees C / 118 degrees F
  DRAM                   3312 MB
  Memory utilization      63 percent
  CPU utilization:
    User                 14 percent
    Background           0 percent
    Kernel               5 percent
    Interrupt            0 percent
    Idle                 81 percent
  Model                  RE-QFXC08-CB4S
  Serial ID              BUILTIN
  Start time             2011-07-06 13:26:15 UTC
  Uptime                 11 hours, 24 minutes, 57 seconds
  Last reboot reason      0x4:reset-button reset
  Load averages:         1 minute   5 minute  15 minute
                        2.62       2.31     2.28

Routing Engine status:
Slot 1:
  Current state           Backup
  Election priority       Backup (default)
  Temperature             39 degrees C / 102 degrees F
  DRAM                   3312 MB
  Memory utilization      59 percent
  CPU utilization:
```



```

User          9 percent
Background    0 percent
Kernel        1 percent
Interrupt     0 percent
Idle          91 percent
Model         RE-QFXC08-CB4S
Serial ID     BUILTIN
Start time    2011-07-06 13:24:58 UTC
Uptime        11 hours, 26 minutes, 18 seconds
Last reboot reason 0x4:reset-button reset

```

show chassis routing-engine (PTX Series Packet Transport Router)

```

user@switch> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state          Master
  Election priority      Master (default)
  Temperature            60 degrees C / 140 degrees F
  CPU temperature        76 degrees C / 168 degrees F
  DRAM                   17152 MB
  Memory utilization      11 percent
  CPU utilization:
    User                  0 percent
    Background            0 percent
    Kernel                 4 percent
    Interrupt              0 percent
    Idle                   95 percent
  Model                  RE-DUO-2600
  Serial ID              P737A-002231
  Start time             2011-12-21 16:54:37 PST
  Uptime                  25 minutes, 44 seconds
  Last reboot reason      Router rebooted after a normal shutdown.
  Load averages:         1 minute   5 minute   15 minute
                        0.01         0.02         0.06

Routing Engine status:
Slot 1:
  Current state          Backup
  Election priority      Backup (default)
  Temperature            50 degrees C / 122 degrees F
  CPU temperature        64 degrees C / 147 degrees F
  DRAM                   17152 MB
  Memory utilization      10 percent
  CPU utilization:
    User                  0 percent
    Background            0 percent
    Kernel                 0 percent
    Interrupt              0 percent
    Idle                   99 percent
  Model                  RE-DUO-2600
  Serial ID              P737A-002438
  Start time             2011-12-21 16:52:26 PST
  Uptime                  27 minutes, 49 seconds
  Last reboot reason      Router rebooted after a normal shutdown.

```

show chassis routing-engine (EX9200 Switch)

```

user@switch> show chassis routing-engine
Routing Engine status:
Slot 0:

```

```

Current state           Master
Election priority       Master (default)
Temperature             35 degrees C / 95 degrees F
CPU temperature         33 degrees C / 91 degrees F
DRAM                   8157 MB
  Installed Memory     8192 MB
  Memory utilization    18 percent
CPU utilization:
  User                 1 percent
  Background           0 percent
  Kernel               4 percent
  Interrupt            1 percent
  Idle                 94 percent
Model                  RE-S-EX9200-1800X4
Serial ID              9009119555
Start time             2014-03-12 14:58:05 UTC
Uptime                 1 hour, 41 minutes, 51 seconds
Last reboot reason     Router rebooted after a normal shutdown.
Load averages:         1 minute  5 minute 15 minute
                       0.02      0.02   0.00

Routing Engine status:
Slot 1:
  Current state        Backup
  Election priority    Backup (default)

[...Output truncated...]

```

show chassis routing-engine (ACX2000 Universal Access Router)

```

user@host> show chassis routing-engine
Routing Engine status:
  Temperature          53 degrees C / 127 degrees F
  DRAM                 1536 MB
  Memory utilization    25 percent
  CPU utilization:
    User               0 percent
    Background         0 percent
    Kernel             0 percent
    Interrupt          1 percent
    Idle               99 percent
  Model               RE-ACX-2000
  Start time          2012-05-09 00:57:07 PDT
  Uptime              5 days, 3 hours, 16 minutes, 15 seconds
  Last reboot reason  Router rebooted after a normal shutdown.
  Load averages:     1 minute  5 minute 15 minute
                     0.00      0.03   0.05

```

show chassis routing-engine (ACX1000 Universal Access Router)

```

user@host> show chassis routing-engine
Routing Engine status:
  Temperature          36 degrees C / 96 degrees F
  DRAM                 768 MB
  Memory utilization    50 percent
  CPU utilization:
    User               3 percent
    Background         0 percent
    Kernel             6 percent
    Interrupt          0 percent
    Idle               91 percent
  Model               RE-ACX-1000

```

Start time	2012-05-10 07:12:23 PDT
Uptime	4 days, 10 hours, 46 minutes, 53 seconds
Last reboot reason	Router rebooted after a normal shutdown.
Load averages:	1 minute 5 minute 15 minute
	0.00 0.00 0.00

show chassis zones

List of Syntax	Syntax on page 996 Syntax (MX Series Routers) on page 996 Syntax (QFX Series) on page 996
Syntax	show chassis zones <detail>
Syntax (MX Series Routers)	show chassis zones <detail> <all-members> <local> <member <i>member-id</i> >
Syntax (QFX Series)	show chassis zones <detail> <interconnect-device <i>name</i> >
Release Information	Command introduced in Junos OS Release 11.3 for the QFX Series. Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers. Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers. all-members , local , and member <i>member-id</i> options introduced in Junos OS Release 15.1 for MX2020 and MX2010 routers.
Description	<p>(QFabric systems only) Display the status of the two cooling system zones on the Interconnect device. Zone 1 consists of eight (0 – 7) front cards, which are cooled by two fan trays. Zone 2 consists of two control boards and eight rear cards, which are cooled by eight (0 – 7) fan trays.</p> <p>(MX2010 and MX2020 routers only) Display the status of the cooling system zones of the chassis. Zone 0 consists of the Control Board, ten (0–9) FPCs, and their respective PICs, Switch Fabric Boards, and Adapter Cards. Zone 1 consists of the Routing Engine, Control Board, and Switch Processor Mezzanine Boards.</p>
Options	<p>all-members—(MX2010 and MX2020 routers only) (Optional) Display the status of the cooling system zones in all members of the Virtual Chassis configuration.</p> <p>detail—(MX2010 and MX2020 routers only) (Optional) Display detailed status of the cooling system zones.</p> <p>detail <i>device-name</i>— (QFabric systems only) (Optional) Display detailed status of the two cooling systems on the Interconnect device.</p> <p>interconnect-device <i>name</i>— (QFabric systems only) (Optional) Display the status of the cooling zones on the Interconnect device.</p> <p>local—(MX2010 and MX2020 routers only) (Optional) Display the status of the cooling system zones in the local member of the Virtual Chassis.</p>

member *member-id*—(MX2010 and MX2020 routers only) (Optional) Display the status of the cooling system zones in the specified member of the Virtual Chassis. Replace *member-id* with the value 0 or 1.

Required Privilege Level view

Related Documentation

- [show chassis fan on page 678](#)
- [show chassis temperature-thresholds](#)

List of Sample Output

- [show chassis zones interconnect-device \(QFabric System\) on page 998](#)
- [show chassis zones \(MX2010 Router\) on page 998](#)
- [show chassis zones detail \(MX2010 Router\) on page 998](#)
- [show chassis zones \(MX2020 Router\) on page 999](#)
- [show chassis zones detail \(MX2020 Router\) on page 999](#)
- [show chassis beacon interconnect-device \(QFabric System\) on page 1001](#)
- [show chassis beacon interconnect-device fpc \(QFabric System\) on page 1001](#)
- [show chassis beacon node-device \(QFabric System\) on page 1001](#)
- [show chassis beacon node-device fpc \(QFabric System\) on page 1001](#)

Output Fields Table 45 lists the output fields for the **show chassis zones** command. Output fields are listed in the approximate order in which they appear.

Table 67: show chassis zones Output Fields

Field Name	Field Description
Slot	FPC slot number of the device whose content is being displayed. On QFX3500 standalone switches, the number is always 0.
Beacon State	Status of the beacon state: <ul style="list-style-type: none"> • Off—The beacon is OFF. • On—The beacon is ON.
show chassis zones command output fields for MX2020 and MX2010 routers:	
Driving FRU	Field replacable unit (FRU).
Temperature	Temperature of the specified FRU in degrees Celsius and degrees Fahrenheit.
Condition	Condition of the specified FRU. Condition can be HIGH TEMP , WARM TEMP , OK , and Offline .
Num Fans Missing	Number of fans or fan trays missing.
Num Fans Failed	Number of fans or fan trays that have failed.
Fan Duty Cycle	Fan duty cycle value.
show chassis zones detail command output fields for MX2020 and MX2010 routers:	

Table 67: show chassis zones Output Fields (*continued*)

Field Name	Field Description
Item	Chassis component: <ul style="list-style-type: none"> Information about the chassis, Routing Engines, Control Boards (CBs), Switch Fabric Boards (SFBs), PICs, Flexible PIC Concentrators (FPCs), and Adapter Cards (ADCs).
Measurement	Fan tray speed utilization in percentage.
Status	Status of the specified item. Status can be OK , Absent , or Offline .

Sample Output

show chassis zones interconnect-device (QFabric System)

```
user@switch> show chassis zones interconnect-device interconnect1
Slot      Beacon State
FPC       0          OFF
```

show chassis zones (MX2010 Router)

```
user@host> show chassis zones
ZONE 0 Status
  Driving FRU          FPC 6
  Temperature          81 degrees C / 177 degrees F
  Condition            HIGH TEMP
  Num Fans Missing     0
  Num Fans Failed      0
  Fan Duty Cycle       30

ZONE 1 Status
  Driving FRU          SFB 0 Exhaust-Zone1
  Temperature          71 degrees C / 159 degrees F
  Condition            WARM TEMP
  Num Fans Missing     0
  Num Fans Failed      0
  Fan Duty Cycle       30
```

show chassis zones detail (MX2010 Router)

```
user@host > show chassis zones
ZONE 0 Status
Item      Status      Measurement
CB 0      WARM TEMP
CB 1      WARM TEMP
FPC 0     HIGH TEMP
FPC 1     HIGH TEMP
FPC 2     WARM TEMP
FPC 3     HIGH TEMP
FPC 4     HIGH TEMP
FPC 5     HIGH TEMP
FPC 6     HIGH TEMP
FPC 7     HIGH TEMP
FPC 8     HIGH TEMP
FPC 9     HIGH TEMP
ADC 0     WARM TEMP
```

ADC 1	WARM TEMP	
ADC 2	WARM TEMP	
ADC 3	WARM TEMP	
ADC 4	WARM TEMP	
ADC 5	WARM TEMP	
ADC 6	WARM TEMP	
ADC 7	WARM TEMP	
ADC 8	WARM TEMP	
ADC 9	WARM TEMP	
SFB 0	WARM TEMP	
SFB 1	WARM TEMP	
SFB 2	WARM TEMP	
SFB 3	Offline	
SFB 4	HIGH TEMP	
SFB 5	WARM TEMP	
SFB 6	HIGH TEMP	
SFB 7	WARM TEMP	
Fan Tray 0	OK	Spinning at 98% fan tray speed
Fan Tray 1	OK	Spinning at 98% fan tray speed
ZONE 1 Status		
Item	Status	Measurement
CB 0	WARM TEMP	
CB 1	WARM TEMP	
Routing Engine 0	OK	
Routing Engine 1	OK	
SFB 0	WARM TEMP	
SFB 1	WARM TEMP	
SFB 2	WARM TEMP	
SFB 3	Offline	
SFB 4	HIGH TEMP	
SFB 5	WARM TEMP	
SFB 6	HIGH TEMP	
SFB 7	WARM TEMP	
SPMB 0	OK	
SPMB 1	OK	
Fan Tray 2	OK	Spinning at 64% fan tray speed
Fan Tray 3	OK	Spinning at 64% fan tray speed

show chassis zones (MX2020 Router)

```

user@host> show chassis zones
ZONE 0 Status
  Driving FRU          FPC 0
  Temperature          31 degrees C / 87 degrees F
  Condition            OK
  Num Fans Missing     0
  Num Fans Failed      0
  Fan Duty Cycle       30

ZONE 1 Status
  Driving FRU          FPC 19
  Temperature          32 degrees C / 89 degrees F
  Condition            OK
  Num Fans Missing     0
  Num Fans Failed      0
  Fan Duty Cycle       30

```

show chassis zones detail (MX2020 Router)

```

user@host> show chassis zones detail

```

ZONE 0 Status

Item	Status	Measurement
CB 0	OK	
CB 1	OK	
FPC 0	OK	
FPC 1	OK	
FPC 2	OK	
FPC 3	OK	
FPC 4	OK	
FPC 5	OK	
FPC 6	OK	
FPC 7	OK	
FPC 8	OK	
FPC 9	OK	
ADC 0	OK	
ADC 1	OK	
ADC 2	OK	
ADC 3	OK	
ADC 4	OK	
ADC 5	OK	
ADC 6	OK	
ADC 7	OK	
ADC 8	OK	
ADC 9	OK	
SFB 0	OK	
SFB 1	OK	
SFB 2	OK	
SFB 3	OK	
SFB 4	OK	
SFB 5	OK	
SFB 6	OK	
SFB 7	OK	
Fan Tray 0	OK	Spinning at 38% fan tray speed
Fan Tray 1	OK	Spinning at 37% fan tray speed

ZONE 1 Status

Item	Status	Measurement
CB 0	OK	
CB 1	OK	
Routing Engine 0	OK	
Routing Engine 1	OK	
FPC 10	OK	
FPC 11	OK	
FPC 12	OK	
FPC 13	OK	
FPC 14	OK	
FPC 15	OK	
FPC 16	OK	
FPC 17	OK	
FPC 18	OK	
FPC 19	OK	
ADC 10	OK	
ADC 11	OK	
ADC 12	OK	
ADC 13	OK	
ADC 14	OK	
ADC 15	OK	
ADC 16	OK	
ADC 17	OK	
ADC 18	OK	
ADC 19	OK	

SFB 0	OK	
SFB 1	OK	
SFB 2	OK	
SFB 3	OK	
SFB 4	OK	
SFB 5	OK	
SFB 6	OK	
SFB 7	OK	
SPMB 0	OK	
SPMB 1	OK	
Fan Tray 2	OK	Spinning at 38% fan tray speed
Fan Tray 3	OK	Spinning at 38% fan tray speed

show chassis beacon interconnect-device (QFabric System)

```
user@switch> show chassis beacon interconnect-device interconnect1
Chassis                OFF
CB 0                   OFF
CB 1                   OFF
FC 0 FPC 0             OFF
FC 1 FPC 1             OFF
RC 0 FPC 8             OFF
RC 1 FPC 9             OFF
```

show chassis beacon interconnect-device fpc (QFabric System)

```
user@switch> show chassis beacon interconnect-device interconnect1 fpc 0
FPC 0                  ON
```

show chassis beacon node-device (QFabric System)

```
user@switch> show chassis beacon node-device node1
node1                  ON
```

show chassis beacon node-device fpc (QFabric System)

```
user@switch> show chassis beacon node-device node1 fpc 0
FPC 0                  ON
```

show host

Syntax	<code>show host <i>hostname</i></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Display Domain Name System (DNS) hostname information.
Options	<i>hostname</i> —Hostname or address.
Additional Information	The <code>show host</code> command displays the raw data received from the DNS server.
Required Privilege Level	view
List of Sample Output	show host on page 1002

Sample Output

show host

```
user@host> show host snark
snark.boojum.net has address 192.168.1.254

user@host> show host 192.168.1.254
Name: snark.boojum.net
Address: 192.168.1.254
Aliases:
```

show log

List of Syntax	Syntax on page 1003 Syntax (QFX Series and OCX Series) on page 1003 Syntax (TX Matrix Router) on page 1003
Syntax	<pre>show log <filename user <username>></pre>
Syntax (QFX Series and OCX Series)	<pre>show log filename <device-type (device-id device-alias)></pre>
Syntax (TX Matrix Router)	<pre>show log <all-lcc lcc number scc> <filename user <username>></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Option <i>device-type (device-id device-alias)</i> is introduced in Junos OS Release 13.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	List log files, display log file contents, or display information about users who have logged in to the router or switch.



NOTE: On MX Series routers, modifying a configuration to replace a service interface with another service interface is treated as a catastrophic event. When you modify a configuration, the entire configuration associated with the service interface—including NAT pools, rules, and service sets—is deleted and then re-created for the newly specified service interface. If there are active sessions associated with the service interface that is being replaced, these sessions are deleted and the NAT pools are then released, which leads to the generation of the NAT_POOL_RELEASE system log messages. However, because NAT pools are already deleted as a result of the catastrophic configuration change and no longer exist, the NAT_POOL_RELEASE system log messages are not generated for the changed configuration.

Options none—List all log files.

<all-lcc | lcc *number* | scc>—(Routing matrix only) (Optional) Display logging information about all T640 routers (or line-card chassis) or a specific T640 router (replace *number* with a value from 0 through 3) connected to a TX Matrix router. Or, display logging information about the TX Matrix router (or switch-card chassis).

device-type—(QFabric system only) (Optional) Display log messages for only one of the following device types:

- **director-device**—Display logs for Director devices.
- **infrastructure-device**—Display logs for the logical components of the QFabric system infrastructure, including the diagnostic Routing Engine, fabric control Routing Engine, fabric manager Routing Engine, and the default network Node group and its backup (NW-NG-0 and NW-NG-0-backup).
- **interconnect-device**—Display logs for Interconnect devices.
- **node-device**—Display logs for Node devices.



NOTE: If you specify the **device-type** optional parameter, you must also specify either the **device-id** or **device-alias** optional parameter.

(device-id | device-alias)—If a device type is specified, display logs for a device of that type. Specify either the device ID or the device alias (if configured).

filename—(Optional) Display the log messages in the specified log file. For the routing matrix, the filename must include the chassis information.

user <username>—(Optional) Display logging information about users who have recently logged in to the router or switch. If you include **username**, display logging information about the specified user.

Required Privilege Level trace

Related Documentation • [syslog \(System\) on page 1704](#)

List of Sample Output [show log on page 1004](#)
[show log filename on page 1005](#)
[show log filename \(QFabric System\) on page 1005](#)
[show log user on page 1006](#)

Sample Output

show log

```
user@host> show log
total 57518
-rw-r--r--  1 root  bin      211663 Oct  1 19:44 dcd
-rw-r--r--  1 root  bin      999947 Oct  1 19:41 dcd.0
-rw-r--r--  1 root  bin      999994 Oct  1 17:48 dcd.1
-rw-r--r--  1 root  bin      238815 Oct  1 19:44 rpd
-rw-r--r--  1 root  bin     1049098 Oct  1 18:00 rpd.0
-rw-r--r--  1 root  bin     1061095 Oct  1 12:13 rpd.1
-rw-r--r--  1 root  bin     1052026 Oct  1 06:08 rpd.2
-rw-r--r--  1 root  bin     1056309 Sep 30 18:21 rpd.3
-rw-r--r--  1 root  bin     1056371 Sep 30 14:36 rpd.4
-rw-r--r--  1 root  bin     1056301 Sep 30 10:50 rpd.5
-rw-r--r--  1 root  bin     1056350 Sep 30 07:04 rpd.6
```

```
-rw-r--r-- 1 root bin 1048876 Sep 30 03:21 rpd.7
-rw-rw-r-- 1 root bin 19656 Oct 1 19:37 wtmp
```

show log filename

```
user@host> show log rpd
Oct 1 18:00:18 trace_on: Tracing to ?/var/log/rpd? started
Oct 1 18:00:18 EVENT <MTU> ds-5/2/0.0 index 24 <Broadcast PointToPoint Multicast
Oct 1 18:00:18
Oct 1 18:00:19 KRT recv len 56 V9 seq 148 op add Type route/if af 2 addr
13.13.13.21 nhop type local nhop 13.13.13.21
Oct 1 18:00:19 KRT recv len 56 V9 seq 149 op add Type route/if af 2 addr
13.13.13.22 nhop type unicast nhop 13.13.13.22
Oct 1 18:00:19 KRT recv len 48 V9 seq 150 op add Type ifaddr index 24 devindex
43
Oct 1 18:00:19 KRT recv len 144 V9 seq 151 op chnge Type ifdev devindex 44
Oct 1 18:00:19 KRT recv len 144 V9 seq 152 op chnge Type ifdev devindex 45
Oct 1 18:00:19 KRT recv len 144 V9 seq 153 op chnge Type ifdev devindex 46
Oct 1 18:00:19 KRT recv len 1272 V9 seq 154 op chnge Type ifdev devindex 47
...
```

show log filename (QFabric System)

```
user@qfabric> show log messages
Mar 28 18:00:06 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:06 ED1486
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 2159)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1486
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 2, jnxFruL3Index 0,
jnxFruName PIC: @ 0/1/*, jnxFruType 11, jnxFruSlot 0, jnxFruOfflineReason 2,
jnxFruLastPowerOff 0, jnxFruLastPowerOn 2191)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 242726)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 2, jnxFruL3Index 0,
jnxFruName PIC: @ 0/1/*, jnxFruType 11, jnxFruSlot 0, jnxFruOfflineReason 2,
jnxFruLastPowerOff 0, jnxFruLastPowerOn 242757)
Mar 28 18:00:16 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:16 ED1486
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:27 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:27 ED1486
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50
_DCF_default__NW-INE-0_RE0_ file: UI_COMMIT: User 'root' requested 'commit'
operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50
_DCF_default__NW-INE-0_RE0_ file: UI_COMMIT: User 'root' requested 'commit'
operation (comment: none)
Mar 28 18:00:55 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:55 ED1492
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:01:10 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:01:10 ED1492
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:02:37 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:02:37 ED1491
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
```

```
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,  
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 33809)
```

show log user

```
user@host> show log user  
darius  mg2546          Thu Oct  1 19:37  still logged in  
darius  mg2529          Thu Oct  1 19:08 - 19:36  (00:28)  
darius  mg2518          Thu Oct  1 18:53 - 18:58  (00:04)  
root    mg1575          Wed Sep 30 18:39 - 18:41  (00:02)  
root    ttyp2    jun.site.per Wed Sep 30 18:39 - 18:41  (00:02)  
alex    ttyp1    192.168.1.2  Wed Sep 30 01:03 - 01:22  (00:19)
```

show subscribers

Syntax show subscribers
 <detail | extensive | terse>
 <aci-interface-set-name *aci-interface-set-name*>
 <address *address*>
 <agent-circuit-identifier *agent-circuit-identifier-substring*>
 <client-type *client-type*>
 <count>
 <id>
 <interface *interface*>
 <logical-system *logical-system*>
 <mac-address *mac-address*>
 <physical-interface *physical-interface-name*>
 <profile-name *profile-name*>
 <routing-instance *routing-instance*>
 <stacked-vlan-id *stacked-vlan-id*>
 <subscriber-state *subscriber-state*>
 <user-name *user-name*>
 <vci *vci-identifier*>
 <vpi *vpi-identifier*>
 <vlan-id *vlan-id*>

Release Information Command introduced in Junos OS Release 9.3.
 Command introduced in Junos OS Release 9.3 for EX Series switches.
client-type, **mac-address**, **subscriber-state**, and **extensive** options introduced in Junos OS Release 10.2.
count option usage with other options introduced in Junos OS Release 10.2.
 Command introduced in Junos OS Release 11.1 for the QFX Series.
 Options **aci-interface-set-name** and **agent-circuit-identifier** introduced in Junos OS Release 12.2.
 The **physical-interface** and **user-name** options introduced in Junos OS Release 12.3.
 Options **vci** and **vpi** introduced in Junos OS Release 12.3R3 and supported in later 12.3Rx releases.
 Options **vci** and **vpi** supported in Junos OS Release 13.2 and later releases. (Not supported in Junos OS Release 13.1.)
 Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
 Command introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.

Description Display information for active subscribers.

Options **detail | extensive | terse**—(Optional) Display the specified level of output.

aci-interface-set-name—(Optional) Display all dynamic subscriber sessions that use the specified agent circuit identifier (ACI) interface set. Use the ACI interface set name generated by the router, such as aci-1003-ge-1/0/0.4001, and not the actual ACI value found in the DHCP or PPPoE control packets.

address—(Optional) Display subscribers whose IP address matches the specified address. You must specify the IPv4 or IPv6 address prefix without a netmask (for example,

192.168.17.1). If you specify the IP address as a prefix with a netmask (for example, 192.168.17.1/32), the router displays a message that the IP address is invalid, and rejects the command.

agent-circuit-identifier-substring—(Optional) Display all dynamic subscriber sessions whose ACI value matches the specified substring.

client-type—(Optional) Display subscribers whose client type matches one of the following client types:

- ***dhcp***—DHCP clients only.
- ***dot1x***—Dot1x clients only.
- ***essm***—ESSM clients only.
- ***fwauth***—FwAuth (authenticated across a firewall) clients only.
- ***l2tp***—L2TP clients only.
- ***mlppp***—MLPPP clients only.
- ***ppp***—PPP clients only.
- ***pppoe***—PPPoE clients only.
- ***static***—Static clients only.
- ***vlan***—VLAN clients only.
- ***vlan-oob***—VLAN out-of-band (ANCP-triggered) clients only.
- ***vpls-pw***—VPLS pseudowire clients only.
- ***xauth***—Xauth clients only.

count—(Optional) Display the count of total subscribers and active subscribers for any specified option. You can use the ***count*** option alone or with the ***address***, ***client-type***, ***interface***, ***logical-system***, ***mac-address***, ***profile-name***, ***routing-instance***, ***stacked-vlan-id***, ***subscriber-state***, or ***vlan-id*** options.

id—(Optional) Display a specific subscriber session whose session id matches the specified subscriber ID. You can display subscriber IDs by using the ***show subscribers extensive*** or the ***show subscribers interface extensive*** commands.

interface—(Optional) Display subscribers whose interface matches the specified interface.

logical-system—(Optional) Display subscribers whose logical system matches the specified logical system.

mac-address—(Optional) Display subscribers whose MAC address matches the specified MAC address.

physical-interface-name—(M120, M320, and MX Series routers only) (Optional) Display subscribers whose physical interface matches the specified physical interface.

profile-name—(Optional) Display subscribers whose dynamic profile matches the specified profile name.

routing-instance—(Optional) Display subscribers whose routing instance matches the specified routing instance.

stacked-vlan-id—(Optional) Display subscribers whose stacked VLAN ID matches the specified stacked VLAN ID.

subscriber-state—(Optional) Display subscribers whose subscriber state matches the specified subscriber state (ACTIVE, CONFIGURED, INIT, TERMINATED, or TERMINATING).

user-name—(M120, M320, and MX Series routers only) (Optional) Display subscribers whose username matches the specified subscriber name.

vci-identifier—(MX Series routers with MPCs and ATM MICs with SFP only) (Optional) Display active ATM subscribers whose ATM virtual circuit identifier (VCI) matches the specified VCI identifier. The range of values is 0 through 255.

vpi-identifier—(MX Series routers with MPCs and ATM MICs with SFP only) (Optional) Display active ATM subscribers whose ATM virtual path identifier (VPI) matches the specified VPI identifier. The range of values is 0 through 65535.

vlan-id—(Optional) Display subscribers whose VLAN ID matches the specified VLAN ID, regardless of whether the subscriber uses a single-tagged or double-tagged VLAN. For subscribers using a double-tagged VLAN, this option displays subscribers where the inner VLAN tag matches the specified VLAN ID. To display only subscribers where the specified value matches only double-tagged VLANs, use the **stacked-vlan-id** option to match the outer VLAN tag.



NOTE: Due to display limitations, logical system and routing instance output values are truncated when necessary.

Required Privilege Level

view

Related Documentation

- [show subscribers summary](#)
- [Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration](#)
- [Verifying and Managing Junos OS Enhanced Subscriber Management](#)

List of Sample Output

[show subscribers \(IPv4\) on page 1014](#)
[show subscribers \(IPv6\) on page 1014](#)
[show subscribers \(IPv4 and IPv6 Dual Stack\) on page 1014](#)
[show subscribers \(LNS on MX Series Routers\) on page 1015](#)
[show subscribers \(L2TP Switched Tunnels\) on page 1015](#)
[show subscribers client-type dhcp detail on page 1015](#)
[show subscribers client-type vlan-oob detail on page 1015](#)
[show subscribers count on page 1016](#)
[show subscribers address detail \(IPv6\) on page 1016](#)

[show subscribers detail \(IPv4\) on page 1016](#)
[show subscribers detail \(IPv6\) on page 1017](#)
[show subscribers detail \(IPv6 Static Demux Interface\) on page 1017](#)
[show subscribers detail \(L2TP LNS Subscribers on MX Series Routers\) on page 1017](#)
[show subscribers detail \(L2TP Switched Tunnels\) on page 1017](#)
[show subscribers detail \(Tunneled Subscriber\) on page 1018](#)
[show subscribers detail \(IPv4 and IPv6 Dual Stack\) on page 1018](#)
[show subscribers detail \(ACI Interface Set Session\) on page 1019](#)
[show subscribers detail \(PPPoE Subscriber Session with ACI Interface Set\) on page 1019](#)
[show subscribers extensive on page 1020](#)
[show subscribers extensive \(RPF Check Fail Filter\) on page 1020](#)
[show subscribers extensive \(L2TP LNS Subscribers on MX Series Routers\) on page 1020](#)
[show subscribers extensive \(IPv4 and IPv6 Dual Stack\) on page 1021](#)
[show subscribers extensive \(ADF Rules \) on page 1022](#)
[show subscribers extensive \(Effective Shaping-Rate\) on page 1022](#)
[show subscribers aci-interface-set-name detail \(Subscriber Sessions Using Specified ACI Interface Set\) on page 1022](#)
[show subscribers agent-circuit-identifier detail \(Subscriber Sessions Using Specified ACI Substring\) on page 1023](#)
[show subscribers interface extensive on page 1023](#)
[show subscribers logical-system terse on page 1024](#)
[show subscribers physical-interface count on page 1024](#)
[show subscribers routing-instance inst1 count on page 1024](#)
[show subscribers stacked-vlan-id detail on page 1024](#)
[show subscribers stacked-vlan-id vlan-id detail \(Combined Output\) on page 1024](#)
[show subscribers stacked-vlan-id vlan-id interface detail \(Combined Output for a Specific Interface\) on page 1025](#)
[show subscribers user-name detail on page 1025](#)
[show subscribers vlan-id on page 1025](#)
[show subscribers vlan-id detail on page 1025](#)
[show subscribers vpi vci extensive \(PPPoE-over-ATM Subscriber Session\) on page 1026](#)
[show subscribers address detail \(Enhanced Subscriber Management\) on page 1026](#)

Output Fields Table 68 lists the output fields for the **show subscribers** command. Output fields are listed in the approximate order in which they appear.

Table 68: show subscribers Output Fields

Field Name	Field Description
Interface	Interface associated with the subscriber. The router or switch displays subscribers whose interface matches or begins with the specified interface. The * character indicates a continuation of addresses for the same session.
IP Address/VLAN ID	Subscriber IP address or VLAN ID associated with the subscriber in the form <i>tpid.vlan-id</i> No IP address or VLAN ID is assigned to an L2TP tunnel-switched session. For these subscriber sessions the value is Tunnel-switched .
User Name	Name of subscriber.

Table 68: show subscribers Output Fields (*continued*)

Field Name	Field Description
LS:RI	Logical system and routing instance associated with the subscriber.
Type	Subscriber client type (DHCP, L2TP, PPP, PPPoE, STATIC-INTERFACE, VLAN).
IP Address	Subscriber IPv4 address.
IP Netmask	Subscriber IP netmask.
Primary DNS Address	IP address of primary DNS server.
Secondary DNS Address	IP address of secondary DNS server.
Primary WINS Address	IP address of primary WINS server.
Secondary WINS Address	IP address of secondary WINS server.
IPv6 Address	Subscriber IPv6 address, or multiple addresses.
IPv6 Prefix	Subscriber IPv6 prefix. If you are using DHCPv6 prefix delegation, this is the delegated prefix.
IPv6 User Prefix	IPv6 prefix obtained through ND/RA.
IPv6 Address Pool	Subscriber IPv6 address pool. The IPv6 address pool is used to allocate IPv6 prefixes to the DHCPv6 clients.
IPv6 Network Prefix Length	Length of the network portion of the IPv6 address.
IPv6 Prefix Length	Length of the subscriber IPv6 prefix.
Logical System	Logical system associated with the subscriber.
Routing Instance	Routing instance associated with the subscriber.
Interface	(Enhanced subscriber management for MX Series routers) Name of the enhanced subscriber management logical interface, in the form demux0.nnnn (for example, demux0.3221225472), to which access-internal and framed subscriber routes are mapped.
Interface Type	Whether the subscriber interface is Static or Dynamic .
Interface Set	Internally generated name of the dynamic ACI interface set used by the subscriber session.
Interface Set Type	Interface type of the ACI interface set: Dynamic . This is the only ACI interface set type currently supported.
Interface Set Session ID	Identifier of the dynamic ACI interface set entry in the session database.

Table 68: show subscribers Output Fields (*continued*)

Field Name	Field Description
Underlying Interface	Name of the underlying interface for the subscriber session.
Dynamic Profile Name	Dynamic profile used for the subscriber.
Dynamic Profile Version	Version number of the dynamic profile used for the subscriber.
MAC Address	MAC address associated with the subscriber.
State	Current state of the subscriber session (Init , Configured , Active , Terminating , Tunneled).
L2TP State	Current state of the L2TP session, Tunneled or Tunnel-switched . When the value is Tunnel-switched , two entries are displayed for the subscriber; the first entry is at the LNS interface on the LTS and the second entry is at the LAC interface on the LTS.
Tunnel switch Profile Name	Name of the L2TP tunnel switch profile that initiates tunnel switching.
Local IP Address	IP address of the local gateway (LAC).
Remote IP Address	IP address of the remote peer (LNS).
VLAN Id	VLAN ID associated with the subscriber in the form <i>tpid.vlan-id</i> .
Stacked VLAN Id	Stacked VLAN ID associated with the subscriber in the form <i>tpid.vlan-id</i> .
RADIUS Accounting ID	RADIUS accounting ID associated with the subscriber.
Agent Circuit ID	<p>For the dhcp client type, option 82 agent circuit ID associated with the subscriber. The ID is displayed as an ASCII string unless the value has nonprintable characters, in which case it is displayed in hexadecimal format.</p> <p>For the vlan-oob client type, the agent circuit ID or access-loop circuit identifier that identifies the subscriber line based on the subscriber-facing DSLAM interface on which the subscriber request originates.</p>
Agent Remote ID	<p>For the dhcp client type, option 82 agent remote ID associated with the subscriber. The ID is displayed as an ASCII string unless the value has nonprintable characters, in which case it is displayed in hexadecimal format.</p> <p>For the vlan-oob client type, the agent remote ID or access-loop remote identifier that identifies the subscriber line based on the NAS-facing DSLAM interface on which the subscriber request originates.</p>
DHCP Relay IP Address	IP address used by the DHCP relay agent.
ATM VPI	(MX Series routers with MPCs and ATM MICs with SFP only) ATM virtual path identifier (VPI) on the subscriber's physical interface.
ATM VCI	(MX Series routers with MPCs and ATM MICs with SFP only) ATM virtual circuit identifier (VCI) for each VPI configured on the subscriber interface.

Table 68: show subscribers Output Fields (*continued*)

Field Name	Field Description
Login Time	Date and time at which the subscriber logged in.
Effective shaping-rate	Actual downstream traffic shaping rate for the subscriber, in kilobits per second.
IPv4 rpf-check Fail Filter Name	Name of the filter applied by the dynamic profile to IPv4 packets that fail the RPF check.
IPv6 rpf-check Fail Filter Name	Name of the filter applied by the dynamic profile to IPv6 packets that fail the RPF check.
DHCP Options	len = number of hex values in the message. The hex values specify the type, length, value (TLV) for DHCP options, as defined in RFC 2132.
Session ID	ID number for a subscriber service session.
Underlying Session ID	For DHCPv6 subscribers on a PPPoE network, displays the session ID of the underlying PPPoE interface.
Service Sessions	Number of service sessions (that is, a service activated using RADIUS CoA) associated with the subscribers.
Service Session Name	Service session profile name.
Session Timeout (seconds)	Number of seconds of access provided to the subscriber before the session is automatically terminated.
Idle Timeout (seconds)	Number of seconds subscriber can be idle before the session is automatically terminated.
IPv6 Delegated Address Pool	Name of the pool used for DHCPv6 prefix delegation.
IPv6 Delegated Network Prefix Length	Length of the prefix configured for the IPv6 delegated address pool.
IPv6 Interface Address	Address assigned by the Framed-Ipv6-Prefix AAA attribute.
IPv6 Framed Interface Id	Interface ID assigned by the Framed-Interface-Id AAA attribute.
ADF IPv4 Input Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv4 input filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
ADF IPv4 Output Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv4 output filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
ADF IPv6 Input Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv6 input filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.

Table 68: show subscribers Output Fields (*continued*)

Field Name	Field Description
ADF IPv6 Output Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv6 output filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
IPv4 Input Filter Name	Name assigned to the IPv4 input filter (client or service session).
IPv4 Output Filter Name	Name assigned to the IPv4 output filter (client or service session).
IPv6 Input Filter Name	Name assigned to the IPv6 input filter (client or service session).
IPv6 Output Filter Name	Name assigned to the IPv6 output filter (client or service session).
IFL Input Filter Name	Name assigned to the logical interface input filter (client or service session).
IFL Output Filter Name	Name assigned to the logical interface output filter (client or service session).

Sample Output

show subscribers (IPv4)

```

user@host> show subscribers
Interface          IP Address/VLAN ID  User Name          LS:RI
ge-1/3/0.1073741824 100                WHOLESALER-CLIENT default:default
demux0.1073741824   100.0.0.10         RETAILER1-CLIENT  test1:retailer1
demux0.1073741825   101.0.0.3          RETAILER1-CLIENT  test1:retailer1
demux0.1073741826   102.0.0.3          RETAILER2-CLIENT  test1:retailer2

```

show subscribers (IPv6)

```

user@host> show subscribers
Interface          IP Address/VLAN ID  User Name          LS:RI
ge-1/0/0.0         2001::c0:0:0:0/74  WHOLESALER-CLIENT default:default
*                  2002::1/128        subscriber-25      default:default

```

show subscribers (IPv4 and IPv6 Dual Stack)

```

user@host> show subscribers
Interface          IP Address/VLAN ID  User Name
LS:RI
demux0.1073741834  0x8100.1002 0x8100.1
default:default
demux0.1073741835  0x8100.1001 0x8100.1
default:default
pp0.1073741836     61.1.1.1        dualstackuser1@ISP1.com
default:ASP-1
*                  2041:1:1::/48
*                  2061:1:1:1::/64
pp0.1073741837     23.1.1.3        dualstackuser2@ISP1.com
default:ASP-1
*                  2001:1:2:5::/64

```

show subscribers (LNS on MX Series Routers)

```

user@host> show subscribers
Interface      IP Address/VLAN ID  User Name      LS:RI
si-4/0/0.1     192.168.4.1         xyz@example.com default:default

```

show subscribers (L2TP Switched Tunnels)

```

user@host> show subscribers
Interface      IP Address/VLAN ID  User Name      LS:RI
si-2/1/0.1073741842 Tunnel-switched    ap@lts.com     default:default

si-2/1/0.1073741843 Tunnel-switched    ap@lts.com     default:default

```

show subscribers client-type dhcp detail

```

user@host> show subscribers client-type dhcp detail
Type: DHCP
IP Address: 100.20.9.7
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: demux0.1073744127
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:10:95:00:00:98
State: Active
Radius Accounting ID: jnpr :2304
Login Time: 2009-08-25 14:43:52 PDT

Type: DHCP
IP Address: 100.20.10.7
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: demux0.1073744383
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:10:94:00:01:f3
State: Active
Radius Accounting ID: jnpr :2560
Login Time: 2009-08-25 14:43:56 PDT

```

show subscribers client-type vlan-oob detail

```

user@host> show subscribers client-type vlan-oob detail
Type: VLAN-OOB
User Name: L2WS.line-aci-1.line-ari-1
Logical System: default
Routing Instance: ISP1
Interface: demux0.1073744127
Interface type: Dynamic
Underlying Interface: ge-1/0/0
Dynamic Profile Name: Prof_L2WS
Dynamic Profile Version: 1
State: Active
Radius Accounting ID: 2304
Session ID: 77
VLAN Id: 126

```

Core-Facing Interface: ge-2/1/1
VLAN Map Id: 6
Inner VLAN Map Id: 2001
Agent Circuit ID: line-aci-1
Agent Remote ID: line-ari-1
Login Time: 2013-10-29 14:43:52 EDT

show subscribers count

```
user@host> show subscribers count
Total Subscribers: 188, Active Subscribers: 188
```

show subscribers address detail (IPv6)

```
user@host> show subscribers address 100.16.12.137 detail
Type: PPPoE
User Name: pppoeTerV6User1Svc
IP Address: 100.16.12.137
IP Netmask: 255.0.0.0
IPv6 User Prefix: 1016:0:0:c88::/64
Logical System: default
Routing Instance: default
Interface: pp0.1073745151
Interface type: Dynamic
Underlying Interface: demux0.8201
Dynamic Profile Name: pppoe-client-profile
MAC Address: 00:0d:02:01:00:01
Session Timeout (seconds): 31622400
Idle Timeout (seconds): 86400
State: Active
Radius Accounting ID: jnpr demux0.8201:6544
Session ID: 6544
Agent Circuit ID: if13720
Agent Remote ID: if13720
Login Time: 2012-05-21 13:37:27 PDT
Service Sessions: 1
```

show subscribers detail (IPv4)

```
user@host> show subscribers detail
Type: DHCP
IP Address: 100.20.9.7
IP Netmask: 255.255.0.0
Primary DNS Address: 192.168.17.1
Secondary DNS Address: 192.168.17.2
Primary WINS Address: 192.168.22.1
Secondary WINS Address: 192.168.22.2
Logical System: default
Routing Instance: default
Interface: demux0.1073744127
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:10:95:00:00:98
State: Active
Radius Accounting ID: jnpr :2304
Idle Timeout (seconds): 600
Login Time: 2009-08-25 14:43:52 PDT
DHCP Options: len 52
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 08 33 04 00 00
00 3c 0c 15 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 36 2f
```



```
33 2d 37 2d 30 37 05 01 06 0f 21 2c
Service Sessions: 2
```

show subscribers detail (IPv6)

```
user@host> show subscribers detail
Type: DHCP
User Name: pd-user1
IPv6 Prefix: 2002:db2:ffff:1::/64
Logical System: default
Routing Instance: default
Interface: ge-3/1/3.2
Interface type: Static
MAC Address: 00:51:ff:ff:00:03
State: Active
Radius Accounting ID: 1
Session ID: 1
Login Time: 2011-08-25 12:12:26 PDT
DHCP Options: len 42
00 08 00 02 00 00 00 01 00 0a 00 03 00 01 00 51 ff ff 00 03
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00
00 00
```

show subscribers detail (IPv6 Static Demux Interface)

```
user@host> show subscribers detail
Type: STATIC-INTERFACE
User Name: demux0.1@jnpr.net
IPv6 Prefix: 1:2:3:4:5:6:7:aa/128
Logical System: default
Routing Instance: default
Interface: demux0.1
Interface type: Static
Dynamic Profile Name: junos-default-profile
State: Active
Radius Accounting ID: 185
Login Time: 2010-05-18 14:33:56 EDT
```

show subscribers detail (L2TP LNS Subscribers on MX Series Routers)

```
user@host> show subscribers detail
Type: L2TP
User Name: user1@jnpr.net
IP Address: 10.1.32.58
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: si-5/2/0.1073749824
Interface type: Dynamic
Dynamic Profile Name: dyn-lns-profile2
Dynamic Profile Version: 1
State: Active
Radius Accounting ID: 8001
Session ID: 8001
Login Time: 2011-04-25 20:27:50 IST
```

show subscribers detail (L2TP Switched Tunnels)

```
user@host> show subscribers detail
Type: L2TP
User Name: ap@example.com
Logical System: default
```

```
Routing Instance: default
Interface: si-2/1/0.1073741842
Interface type: Dynamic
Dynamic Profile Name: dyn-lts-profile
State: Active
L2TP State: Tunnel-switched
Tunnel switch Profile Name: ce-lts-profile
Local IP Address: 10.50.1.1
Remote IP Address: 192.168.20.3
Radius Accounting ID: 21
Session ID: 21
Login Time: 2013-01-18 03:01:11 PST
```

```
Type: L2TP
User Name: ap@example.com
Logical System: default
Routing Instance: default
Interface: si-2/1/0.1073741843
Interface type: Dynamic
Dynamic Profile Name: dyn-lts-profile
State: Active
L2TP State: Tunnel-switched
Tunnel switch Profile Name: ce-lts-profile
Local IP Address: 10.30.1.1
Remote IP Address: 172.20.1.10
Session ID: 22
Login Time: 2013-01-18 03:01:14 PST
```

show subscribers detail (Tunneled Subscriber)

```
user@host> show subscribers detail
Type: PPPoE
User Name: user1@example.com
Logical System: default
Routing Instance: default
Interface: pp0.1
State: Active, Tunneled
Radius Accounting ID: 512
```

show subscribers detail (IPv4 and IPv6 Dual Stack)

```
user@host> show subscribers detail
Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlanProfile
State: Active
Session ID: 1
Stacked VLAN Id: 0x8100.1001
VLAN Id: 0x8100.1
Login Time: 2011-11-30 00:18:04 PST
```

```
Type: PPPoE
User Name: dualstackuser1@ISP1.com
IP Address: 61.1.1.1
IPv6 Prefix: 2041:1:1::/48
IPv6 User Prefix: 2061:1:1:1::/64
Logical System: default
Routing Instance: ASP-1
```

```

Interface: pp0.1073741825
Interface type: Dynamic
Dynamic Profile Name: dualStack-Profile1
MAC Address: 00:00:64:03:01:02
State: Active
Radius Accounting ID: 2
Session ID: 2
Login Time: 2011-11-30 00:18:05 PST

Type: DHCP
IPv6 Prefix: 2041:1:1::/48
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Static
MAC Address: 00:00:64:03:01:02
State: Active
Radius Accounting ID: jnpr :3
Session ID: 3
Underlying Session ID: 2
Login Time: 2011-11-30 00:18:35 PST
DHCP Options: len 42
00 08 00 02 0b b8 00 01 00 0a 00 03 00 01 00 00 64 03 01 02
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00 00
00 00

```

show subscribers detail (ACI Interface Set Session)

```

user@host> show subscribers detail
Type: VLAN
Logical System: default
Routing Instance: default
Interface: ge-1/0/0
Interface Set: aci-1001-ge-1/0/0.2800
Interface Set Session ID: 0
Underlying Interface: ge-1/0/0.2800
Dynamic Profile Name: aci-vlan-set-profile-2
Dynamic Profile Version: 1
State: Active
Session ID: 1
Agent Circuit ID: aci-ppp-dhcp-20
Login Time: 2012-05-26 01:54:08 PDT

```

show subscribers detail (PPPoE Subscriber Session with ACI Interface Set)

```

user@host> show subscribers detail
Type: PPPoE
User Name: ppphint2
IP Address: 10.10.1.5
Logical System: default
Routing Instance: default
Interface: pp0.1073741825
Interface type: Dynamic
Interface Set: aci-1001-demux0.1073741824
Interface Set Type: Dynamic
Interface Set Session ID: 2
Underlying Interface: demux0.1073741824
Dynamic Profile Name: aci-vlan-pppoe-profile
Dynamic Profile Version: 1
MAC Address: 00:00:64:39:01:02

```

```
State: Active
Radius Accounting ID: 3
Session ID: 3
Agent Circuit ID: aci-ppp-dhcp-dvlan-50
Login Time: 2012-03-07 13:46:53 PST
```

show subscribers extensive

```
user@host> show subscribers extensive
Type: DHCP
User Name: pd-user1
IPv6 Prefix: 2002:db2:ffff:1::/64
Logical System: default
Routing Instance: default
Interface: ge-3/1/3.2
Interface type: Static
MAC Address: 00:51:ff:ff:00:03
State: Active
Radius Accounting ID: 1
Session ID: 1
Login Time: 2011-08-25 12:12:26 PDT
DHCP Options: len 42
00 08 00 02 00 00 00 01 00 0a 00 03 00 01 00 51 ff ff 00 03
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00
00 00
IPv6 Address Pool: pd_pool
IPv6 Network Prefix Length: 48
```

show subscribers extensive (RPF Check Fail Filter)

```
user@host> show subscribers extensive
...
Type: VLAN
Logical System: default
Routing Instance: default
Interface: ae0.1073741824
Interface type: Dynamic
Dynamic Profile Name: vlan-prof
State: Active
Session ID: 9
VLAN Id: 100
Login Time: 2011-08-26 08:17:00 PDT
IPv4 rpf-check Fail Filter Name: rpf-allow-dhcp
IPv6 rpf-check Fail Filter Name: rpf-allow-dhcpv6
...
```

show subscribers extensive (L2TP LNS Subscribers on MX Series Routers)

```
user@host> show subscribers extensive
Type: L2TP
User Name: user1@jnpr.net
IP Address: 10.1.32.58
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: si-5/2/0.1073749824
Interface type: Dynamic
Dynamic Profile Name: dyn-lns-profile2
Dynamic Profile Version: 1
State: Active
Radius Accounting ID: 8001
Session ID: 8001
```

```

Login Time: 2011-04-25 20:27:50 IST
IPv4 Input Filter Name: classify-si-5/2/0.1073749824-in
IPv4 Output Filter Name: classify-si-5/2/0.1073749824-out

```

show subscribers extensive (IPv4 and IPv6 Dual Stack)

```

user@host> show subscribers extensive
Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlanProfile
State: Active
Session ID: 1
Stacked VLAN Id: 0x8100.1001
VLAN Id: 0x8100.1
Login Time: 2011-11-30 00:18:04 PST

Type: PPPoE
User Name: dualstackuser1@ISP1.com
IP Address: 61.1.1.1
IPv6 Prefix: 2041:1:1::/48
IPv6 User Prefix: 2061:1:1:1::/64
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Dynamic
Dynamic Profile Name: dualStack-Profile1
MAC Address: 00:00:64:03:01:02
State: Active
Radius Accounting ID: 2
Session ID: 2
Login Time: 2011-11-30 00:18:05 PST
IPv6 Delegated Network Prefix Length: 48
IPv6 Interface Address: 2061:1:1:1::1/64
IPv6 Framed Interface Id: 1:1:2:2
IPv4 Input Filter Name: FILTER-IN-pp0.1073741825-in
IPv4 Output Filter Name: FILTER-OUT-pp0.1073741825-out
IPv6 Input Filter Name: FILTER-IN6-pp0.1073741825-in
IPv6 Output Filter Name: FILTER-OUT6-pp0.1073741825-out

Type: DHCP
IPv6 Prefix: 2041:1:1::/48
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Static
MAC Address: 00:00:64:03:01:02
State: Active
Radius Accounting ID: jnpr :3
Session ID: 3
Underlying Session ID: 2
Login Time: 2011-11-30 00:18:35 PST
DHCP Options: len 42
00 08 00 02 0b b8 00 01 00 0a 00 03 00 01 00 00 64 03 01 02
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00 00
00 00
IPv6 Delegated Network Prefix Length: 48

```

show subscribers extensive (ADF Rules)

```
user@host> show subscribers extensive
...
Service Session ID: 12
Service Session Name: SERVICE-PROFILE
State: Active
Family: inet
  ADF IPv4 Input Filter Name: __junos_adf_12-demux0.3221225474-inet-in
    Rule 0: 010101000b0101020b020200201811
      from {
        source-address 11.1.1.2/32;
        destination-address 11.2.2.0/24;
        protocol 17;
      }
      then {
        accept;
      }
```

show subscribers extensive (Effective Shaping-Rate)

```
user@host> show subscribers extensive
Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.1073741837
Interface type: Dynamic
Interface Set: ifset-1
Underlying Interface: ae1
Dynamic Profile Name: svlan-dhcp-test
State: Active
Session ID: 1
Stacked VLAN Id: 0x8100.201
VLAN Id: 0x8100.201
Login Time: 2011-11-30 00:18:04 PST
Effective shaping-rate: 31000000k
...
```

show subscribers aci-interface-set-name detail (Subscriber Sessions Using Specified ACI Interface Set)

```
user@host> show subscribers aci-interface-set-name aci-1003-ge-1/0/0.4001 detail
Type: VLAN
Logical System: default
Routing Instance: default
Interface: ge-1/0/0.
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-set-profile
Dynamic Profile Version: 1
State: Active
Session ID: 13
Agent Circuit ID: aci-ppp-vlan-10
Login Time: 2012-03-12 10:41:56 PDT

Type: PPPoE
User Name: ppphint2
IP Address: 10.10.1.7
Logical System: default
Routing Instance: default
Interface: pp0.1073741834
Interface type: Dynamic
Interface Set: aci-1003-ge-1/0/0.4001
```

```

Interface Set Type: Dynamic
Interface Set Session ID: 13
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-pppoe-profile
Dynamic Profile Version: 1
MAC Address: 00:00:65:26:01:02
State: Active
Radius Accounting ID: 14
Session ID: 14
Agent Circuit ID: aci-ppp-vlan-10
Login Time: 2012-03-12 10:41:57 PDT

```

show subscribers agent-circuit-identifier detail (Subscriber Sessions Using Specified ACI Substring)

```

user@host> show subscribers agent-circuit-identifier aci-ppp-vlan detail
Type: VLAN
Logical System: default
Routing Instance: default
Interface: ge-1/0/0.
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-set-profile
Dynamic Profile Version: 1
State: Active
Session ID: 13
Agent Circuit ID: aci-ppp-vlan-10
Login Time: 2012-03-12 10:41:56 PDT

Type: PPPoE
User Name: ppphint2
IP Address: 10.10.1.7
Logical System: default
Routing Instance: default
Interface: pp0.1073741834
Interface type: Dynamic
Interface Set: aci-1003-ge-1/0/0.4001
Interface Set Type: Dynamic
Interface Set Session ID: 13
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-pppoe-profile
Dynamic Profile Version: 1
MAC Address: 00:00:65:26:01:02
State: Active
Radius Accounting ID: 14
Session ID: 14
Agent Circuit ID: aci-ppp-vlan-10
Login Time: 2012-03-12 10:41:57 PDT

```

show subscribers interface extensive

```

user@host> show subscribers interface demux0.1073741826 extensive
Type: VLAN
User Name: test1@test.com
Logical System: default
Routing Instance: testnet
Interface: demux0.1073741826
Interface type: Dynamic
Dynamic Profile Name: profile-vdemux-relay-23qos
MAC Address: 00:00:6e:56:01:04
State: Active
Radius Accounting ID: 12
Session ID: 12

```

Stacked VLAN Id: 0x8100.1500
VLAN Id: 0x8100.2902
Login Time: 2011-10-20 16:21:59 EST

Type: DHCP
User Name: test1@test.com
IP Address: 172.16.200.6
IP Netmask: 255.255.255.0
Logical System: default
Routing Instance: testnet
Interface: demux0.1073741826
Interface type: Static
MAC Address: 00:00:6e:56:01:04
State: Active
Radius Accounting ID: 21
Session ID: 21
Login Time: 2011-10-20 16:24:33 EST
Service Sessions: 2

Service Session ID: 25
Service Session Name: SUB-QOS
State: Active

Service Session ID: 26
Service Session Name: service-cb-content
State: Active
IPv4 Input Filter Name: content-cb-in-demux0.1073741826-in
IPv4 Output Filter Name: content-cb-out-demux0.1073741826-out

show subscribers logical-system terse

```
user@host> show subscribers logical-system test1 terse
Interface          IP Address/VLAN ID  User Name          LS:RI
demux0.1073741825  101.0.0.3           RETAILER1-CLIENT  test1:retailer1
demux0.1073741826  102.0.0.3           RETAILER2-CLIENT  test1:retailer2
```

show subscribers physical-interface count

```
user@host> show subscribers physical-interface ge-1/0/0 count
Total subscribers: 3998, Active Subscribers: 3998
```

show subscribers routing-instance inst1 count

```
user@host> show subscribers routing-instance inst1 count
Total Subscribers: 188, Active Subscribers: 183
```

show subscribers stacked-vlan-id detail

```
user@host> show subscribers stacked-vlan-id 101 detail
Type: VLAN
Interface: ge-1/2/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlan-prof
State: Active
Stacked VLAN Id: 0x8100.101
VLAN Id: 0x8100.100
Login Time: 2009-03-27 11:57:19 PDT
```

show subscribers stacked-vlan-id vlan-id detail (Combined Output)

```
user@host> show subscribers stacked-vlan-id 101 vlan-id 100 detail
```



```
Type: VLAN
Interface: ge-1/2/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlan-prof
State: Active
Stacked VLAN Id: 0x8100.101
VLAN Id: 0x8100.100
Login Time: 2009-03-27 11:57:19 PDT
```

show subscribers stacked-vlan-id vlan-id interface detail (Combined Output for a Specific Interface)

```
user@host> show subscribers stacked-vlan-id 101 vlan-id 100 interface ge-1/2/0.* detail
Type: VLAN
Interface: ge-1/2/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlan-prof
State: Active
Stacked VLAN Id: 0x8100.101
VLAN Id: 0x8100.100
Login Time: 2009-03-27 11:57:19 PDT
```

show subscribers user-name detail

```
user@host> show subscribers user-name larry1 detail
Type: DHCP
User Name: larry1
IP Address: 100.0.0.37
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: ge-1/0/0.1
Interface type: Static
Dynamic Profile Name: foo
MAC Address: 00:10:94:00:00:01
State: Active
Radius Accounting ID: 1
Session ID: 1
Login Time: 2011-11-07 08:25:59 PST
DHCP Options: len 52
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 01 33 04 00 00
00 3c 0c 15 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 32 2f
37 2d 30 2d 30 37 05 01 06 0f 21 2c
```

show subscribers vlan-id

```
user@host> show subscribers vlan-id 100
Interface          IP Address          User Name
ge-1/0/0.1073741824
ge-1/2/0.1073741825
```

show subscribers vlan-id detail

```
user@host> show subscribers vlan-id 100 detail
Type: VLAN
Interface: ge-1/0/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: vlan-prof-tpid
State: Active
VLAN Id: 100
Login Time: 2009-03-11 06:48:54 PDT

Type: VLAN
```

```
Interface: ge-1/2/0.1073741825
Interface type: Dynamic
Dynamic Profile Name: vlan-prof-tpid
State: Active
VLAN Id: 100
Login Time: 2009-03-11 06:48:54 PDT
```

show subscribers vpi vci extensive (PPPoE-over-ATM Subscriber Session)

```
user@host> show subscribers vpi 40 vci 50 extensive
Type: PPPoE
User Name: testuser
IP Address: 100.0.0.2
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: pp0.0
Interface type: Static
MAC Address: 00:00:65:23:01:02
State: Active
Radius Accounting ID: 2
Session ID: 2
ATM VPI: 40
ATM VCI: 50
Login Time: 2012-12-03 07:49:26 PST
IP Address Pool: pool_1
IPv6 Framed Interface Id: 200:65ff:fe23:102
```

show subscribers address detail (Enhanced Subscriber Management)

```
user@host> show subscribers address 100.20.0.111 detail
Type: DHCP
User Name: simple_filters_service
IP Address: 5.0.0.2
IP Netmask: 255.0.0.0
Logical System: default
Routing Instance: default
Interface: demux0.3221225482
Interface type: Dynamic
Underlying Interface: demux0.3221225472
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:01:02:03:04:0f
State: Active
Radius Accounting ID: 11
Session ID: 11
PFE Flow ID: 15
Stacked VLAN Id: 210
VLAN Id: 209
Login Time: 2014-03-24 12:53:48 PDT
Service Sessions: 1
DHCP Options: len 3
35 01 01
```

show system alarms

Syntax	show system alarms
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Display active system alarms.
Options	This command has no options.
Additional Information	<p>System alarms are preset. They include a <i>configuration</i> alarm that appears when no rescue configuration alarm is set and a <i>license</i> alarm that appears when a software feature is configured and no valid license is configured for the feature. On EX6200 switches, an alarm can be triggered by an internal link error. For more information about system alarms, see the <i>Junos OS Administration Library for Routing Devices</i>.</p> <p>In Junos OS release 11.1 and later, alarms for fans also show the slot number of the malfunctioning fans in the CLI output.</p> <p>Starting with Junos OS Release 13.2, you can view degraded fabric alarms on a routing matrix based on TX Matrix Plus router with 3D SIBs. The alarm indicates that the source FPC is running with a degraded fabric condition. This alarm is an early warning of a possible fabric black-hole condition. When the degraded fabric alarm is raised on the source FPC, you can take remedial action to avoid a fabric black-hole condition. The degraded fabric alarm is raised on the source FPC if both the following conditions are met:</p> <ul style="list-style-type: none"> • The active Packet Forwarding Engine destinations are reachable on one or no active switching planes. • At least one of the inactive switching planes has a fault that causes the destination Packet Forwarding Engine to become unreachable.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show chassis alarms on page 459
List of Sample Output	<p>show system alarms on page 1028</p> <p>show system alarms (Fan Tray) on page 1028</p> <p>show system alarms (QFX Series and OCX Series) on page 1028</p> <p>show system alarms (EX6200) on page 1028</p> <p>show system alarms (TX Matrix Plus router with 3D SIBs) on page 1028</p>
Output Fields	Table 69 lists the output fields for the show system alarms command. Output fields are listed in the approximate order in which they appear.

Table 69: show system alarms Output Fields

Field Name	Field Description
Alarm time	Date and time the alarm was first recorded.
Class	Severity class for this alarm: Minor or Major .
Description	Information about the alarm.

Sample Output

show system alarms

```

user@host> show system alarms
2 alarms currently active
Alarm time          Class    Description
2005-02-24 17:29:34 UTC  Minor    IPsec VPN tunneling usage requires a
license
2005-02-24 17:29:34 UTC  Minor    Rescue configuration is not sent

```

show system alarms (Fan Tray)

```

user@host> show system alarms
4 alarms currently active
Alarm time          Class    Description
2010-11-11 20:27:38 UTC  Major    Side Fan Tray 7 Failure
2010-11-11 20:27:13 UTC  Minor    Side Fan Tray 7 Overspeed
2010-11-11 20:27:13 UTC  Major    Side Fan Tray 5 Failure
2010-11-11 20:27:13 UTC  Major    Side Fan Tray 0 Failure

```

show system alarms (QFX Series and OCX Series)

```

user@switch> show system alarms
2 alarms currently active
Alarm time Class Description
2005-02-24 17:29:34 UTC Minor Rescue configuration is not sent

```

show system alarms (EX6200)

```

user@switch> show system alarms
2 alarms currently active
Alarm time          Class    Description
2013-04-05 16:51:41 PDT  Major    FPC 8 internal link errors detected
2013-04-04 18:05:35 PDT  Minor    Rescue configuration is not set

```

show system alarms (TX Matrix Plus router with 3D SIBs)

```

user@router> show system alarms

sfc0-re0:
-----
2 alarms currently active
Alarm time          Class    Description
2013-05-08 18:13:58 UTC  Major    LCC 0 Major Errors
2013-05-08 17:48:46 UTC  Major    LCC 7 Major Errors

lcc0-re1:
-----

```

1 alarm currently active

Alarm time	Class	Description
2013-05-08 18:19:24 UTC	Major	FPC 1 degraded fabric condition detected

lcc7-re0:

1 alarm currently active

Alarm time	Class	Description
2013-05-08 18:19:24 UTC	Major	FPC 7 degraded fabric condition detected

show system audit

List of Syntax	Syntax on page 1030 Syntax (EX Series Switch and MX Series Router) on page 1030 Syntax (TX Matrix Router) on page 1030 Syntax (TX Matrix Plus Router) on page 1030 Syntax (QFX Series) on page 1030 Syntax (OCX Series) on page 1030
Syntax	show system audit <root-only>
Syntax (EX Series Switch and MX Series Router)	show system audit <all-members> <local> <member <i>member-id</i> > <root-only>
Syntax (TX Matrix Router)	show system audit <all-lcc lcc <i>number</i> scc> <root-only>
Syntax (TX Matrix Plus Router)	show system audit <all-chassis all-lcc lcc <i>number</i> sfc <i>number</i> > <root-only>
Syntax (QFX Series)	show system audit <infrastructure <i>name</i> interconnect-device <i>name</i> node-group <i>name</i> root-only>
Syntax (OCX Series)	show system audit <root-only>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Display the state and checksum values for file systems.
Options	none —Display the state and checksum values for all file systems. all-chassis —(TX Matrix routers and TX Matrix Plus routers only) (Optional) Display file system MD5 hash and permissions information for all of the chassis. all-lcc —(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display file system MD5 hash and permissions information for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, display file system MD5 hash and permissions information for all T1600 or T4000 routers connected to the TX Matrix Plus router.

all-members—(EX4200 switch, QFX Series, and MX Series routers only) (Optional)

Display file system MD5 hash and permissions information on all members of the Virtual Chassis configuration.

lcc number—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display file system MD5 hash and permissions information for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display file system MD5 hash and permissions information for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

infrastructure name—(QFabric systems only) (Optional) Display file system MD5 hash and permissions information for a fabric control Routing Engine or a fabric control Routing Engine.

interconnect-device name—(QFabric systems only) (Optional) Display file system MD5 hash and permissions information for the Interconnect device.

local—(EX4200 switch, QFX Series, and MX Series routers only) (Optional) Display file system MD5 hash and permissions information on the local Virtual Chassis member.

member member-id—(EX4200 switch, QFX Series, and MX Series routers only) (Optional) Display file system MD5 hash and permissions information on the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

node-group name—(QFabric systems only) (Optional) Display file system MD5 hash and permissions information for the Node group

root-only—(Optional) Check only the root (/) file system. On a QFabric system, you can check the root (/) file system on the infrastructure (fabric manager Routing Engine and fabric control Routing Engine), Interconnect device, or Node group.

scc—(TX Matrix routers only) (Optional) Display file system MD5 hash and permissions information for the TX Matrix router (or switch-card chassis).

sfc number—(TX Matrix Plus routers only) (Optional) Display file system MD5 hash and permissions information for the TX Matrix Plus router (or switch-fabric chassis). Replace *number* with 0.

Additional Information To redirect the output to a file, issue the following command:

```
ssh device-name 'show system audit root-only' > output-file
```

If you save the output of the **show system audit root-only** command to a file, you can compare it to subsequent output from the command to determine whether anything has changed.

By default, when you issue the **show system audit** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

Required Privilege Level admin

List of Sample Output [show system audit root-only on page 1032](#)
[show system audit lcc \(TX Matrix Router\) on page 1033](#)
[show system audit lcc \(TX Matrix Plus Router\) on page 1035](#)
[show system audit root-only \(QFX3500 Switch\) on page 1036](#)

Sample Output

show system audit root-only

```
user@host> show system audit root-only
#          user: root
#          machine: my-host
#          tree: /
date: Fri Feb 11 21:21:46 2000

# .
/set type=file uid=0 gid=0 mode=0755 nlink=1
.          type=dir nlink=23 size=1024 time=950252640.0
.cshrc     uid=3 gid=7 mode=0644 size=177 time=939182975.0 \
           md5digest=f414e06fea6bd646244b98e13d6e6226
.kernel.jkernel.backup \
           mode=0744 size=1934552 time=944688902.0 \
           md5digest=2c343cf0bd9fea8f04f78604feed7aa4
.profile   uid=3 gid=7 mode=0644 nlink=2 size=173 time=939182975.0 \
           md5digest=55a1e3c6c67789c9d3a1cce1ea39f670
COPYRIGHT  uid=3 gid=7 mode=0444 size=3425 time=939182975.0 \
           md5digest=7df8bc77dcee71382ea73eb0ec6a9243
boot.config mode=0644 size=3 time=945902618.0 \
           md5digest=93d722493ed38477338a1405d7dcbb40
boot.help  uid=3 gid=7 mode=0444 size=411 time=939182876.0 \
           md5digest=9b7126385734bcae753f4179ab59d8e5
compat     type=link mode=0777 size=11 time=915149058.0 \
           link=/usr/compat
kernel     mode=0444 size=1947607 time=950230892.0 \
           md5digest=1a2a8aff2fec678a918ba0d6bf063980
kernel.avr uid=1112 size=1947642 time=950252597.0 \
           md5digest=82e1637682d58ec28964dfee7fccb62e
kernel.config \
           mode=0644 size=0 time=915149058.0 \
           md5digest=d41d8cd98f00b204e9800998ecf8427e
```



```
sys          type=link mode=0777 size=11 time=915149029.0 \
link=usr/src/sys
```

show system audit lcc (TX Matrix Router)

```
user@host> show system audit lcc 2
lcc2-re0:
-----
#          user: root
#          machine: rodin-lcc2
#          tree: /
#          date: Mon Sep 13 11:55:33 2004

# .
/set type=file uid=0 gid=0 mode=0555 nlink=1 flags=none
.          type=dir nlink=20 size=512 time=1094982121.0
  COPYRIGHT mode=0644 size=4735 time=986012708.0 \
            md5digest=78396df1404ad742e6eb1be28f0cd63b
  kernel    type=link mode=0700 size=17 time=1090266262.0 \
            link=/packages/jkernel

# ./altconfig
altconfig  type=dir nlink=2 size=512 time=1089801320.0
# ./altconfig
..

# ./altroot
altroot    type=dir nlink=2 size=512 time=1089801320.0
# ./altroot
..

# ./b
b          type=dir mode=0755 nlink=2 size=512 time=1093961429.0
# ./b
..

# ./bin
/set type=file uid=0 gid=0 mode=0700 nlink=1 flags=none
bin        type=dir mode=0755 nlink=2 size=512 time=1089843059.0
  [         type=link size=28 time=1090266270.0 \
            link=/packages/mnt/jbase/bin/test
  cat       type=link size=27 time=1090266270.0 \
            link=/packages/mnt/jbase/bin/cat
  chmod     type=link size=29 time=1090266270.0 \
            link=/packages/mnt/jbase/bin/chmod
  cp        type=link size=26 time=1090266270.0 \
            link=/packages/mnt/jbase/bin/cp
  csh       type=link size=27 time=1090266270.0 \
            link=/packages/mnt/jbase/bin/csh
  date      type=link size=28 time=1090266270.0 \
            link=/packages/mnt/jbase/bin/date
  dd        type=link size=26 time=1090266270.0 \
            link=/packages/mnt/jbase/bin/dd
  df        type=link size=26 time=1090266270.0 \
            link=/packages/mnt/jbase/bin/df
  echo      type=link size=28 time=1090266270.0 \
            link=/packages/mnt/jbase/bin/echo
  ed        type=link size=26 time=1090266270.0 \
            link=/packages/mnt/jbase/bin/ed
```

```

expr      type=link size=28 time=1090266270.0 \
          link=/packages/mnt/jbase/bin/expr
hostname  type=link size=32 time=1090266270.0 \
          link=/packages/mnt/jbase/bin/hostname
kill      type=link size=28 time=1090266270.0 \
          link=/packages/mnt/jbase/bin/kill
ln        type=link size=26 time=1090266270.0 \
          link=/packages/mnt/jbase/bin/ln
ls        type=link size=26 time=1090266270.0 \
          link=/packages/mnt/jbase/bin/ls
mkdir     type=link size=29 time=1090266270.0 \
          link=/packages/mnt/jbase/bin/mkdir
mv        type=link size=26 time=1090266270.0 \
          link=/packages/mnt/jbase/bin/mv
ps        type=link size=26 time=1090266270.0 \
          link=/packages/mnt/jbase/bin/ps
pwd       type=link size=27 time=1090266270.0 \
          link=/packages/mnt/jbase/bin/pwd
rcp       type=link size=27 time=1090266270.0 \
          link=/packages/mnt/jbase/bin/rcp
red       type=link size=26 time=1090266270.0 \
          link=/packages/mnt/jbase/bin/red
rm        type=link size=26 time=1090266270.0 \
          link=/packages/mnt/jbase/bin/rm
rmdir     type=link size=29 time=1090266270.0 \
          link=/packages/mnt/jbase/bin/rmdir
sh        type=link size=26 time=1090266270.0 \
          link=/packages/mnt/jbase/bin/sh
sleep     type=link size=29 time=1090266270.0 \
          link=/packages/mnt/jbase/bin/sleep
stty      type=link size=28 time=1090266270.0 \
          link=/packages/mnt/jbase/bin/stty
sync      type=link size=28 time=1090266270.0 \
          link=/packages/mnt/jbase/bin/sync
tcsh      type=link size=27 time=1090266270.0 \
          link=/packages/mnt/jbase/bin/csh
test      type=link size=28 time=1090266270.0 \
          link=/packages/mnt/jbase/bin/test
# ./bin
..

# ./boot
/set type=file uid=0 gid=0 mode=0444 nlink=1 flags=none
boot      type=dir mode=0555 nlink=3 size=512 time=1095069935.0
  boot0   size=512 time=1094978286.0 \
          md5digest=6f780822dd4ae482a20462b66e542cca
  boot1   mode=0555 size=512 time=1094978294.0 \
          md5digest=8d112b09df342cd0b60fdb9bdcde8e07
  boot2   mode=0555 size=7680 time=1094978294.0 \
          md5digest=28eb58c4068c6b85717e1484f9e028e4
  cdboot  mode=0555 size=165888 time=1094978298.0 \
          md5digest=1474c6b800dfc82ba552d7c36116d07d
  kgzldr.o size=5996 time=1094982121.0 \
          md5digest=c53dc948eb07e2ea4eb0413e4c4634a3
  loader  mode=0555 size=163840 time=1094978298.0 \
          md5digest=82d9dc2d31033476bfb61bb7264c4fed
  loader.4th size=9237 time=986013631.0 \
          md5digest=43144391465ad50267d31e0a320be1de
...
```

show system audit lcc (TX Matrix Plus Router)

```
user@host> show system audit all-chassis
```

```
sfc0-re0:
```

```
-----
#          user: root
#          machine: finalfive
#          tree: /
#          date: Mon May 18 00:13:16 2009

# .
/set type=file uid=0 gid=0 mode=0755 nlink=1 flags=none
.      type=dir nlink=23 size=512 time=1242347096.0
      COPYRIGHT mode=0644 size=6196 time=1168587741.0 \
          md5digest=bbad415e1c29bbdd9b383537100412c
          kernel type=link size=17 time=1242347011.0 link=/packages/jkernel
          staging type=link mode=0777 size=8 time=1242346935.0 link=/var/tmp

# ./snap
.snap type=dir mode=0775 nlink=2 size=512 time=1242346922.0
# ./snap
..

# ./altconfig
altconfig type=dir mode=0500 nlink=2 size=512 time=1242319843.0
# ./altconfig
..

# ./altroot
altroot type=dir mode=0500 nlink=2 size=512 time=1242319843.0
# ./altroot
..

# ./bin
bin type=dir nlink=2 size=512 time=1242346944.0
  \133 type=link size=28 time=1242346942.0 \
      link=/packages/mnt/jbase/bin/test
  cat type=link size=27 time=1242346941.0 \
      link=/packages/mnt/jbase/bin/cat
  chflags type=link size=31 time=1242346941.0 \
      link=/packages/mnt/jbase/bin/chflags
  chmod type=link size=29 time=1242346941.0 \
      link=/packages/mnt/jbase/bin/chmod
  cp type=link size=26 time=1242346941.0 \
      link=/packages/mnt/jbase/bin/cp
  csh type=link size=27 time=1242346941.0 \
      link=/packages/mnt/jbase/bin/csh
  date type=link size=28 time=1242346941.0 \
      link=/packages/mnt/jbase/bin/date
  dd type=link size=26 time=1242346941.0 \
      link=/packages/mnt/jbase/bin/dd
  df type=link size=26 time=1242346941.0 \
      link=/packages/mnt/jbase/bin/df
  echo type=link size=28 time=1242346941.0 \
      link=/packages/mnt/jbase/bin/echo
  ed type=link size=26 time=1242346941.0 \
      link=/packages/mnt/jbase/bin/ed
```

```

expr      type=link size=28 time=1242346941.0 \
          link=/packages/mnt/jbase/bin/expr
hostname  type=link size=32 time=1242346941.0 \
          link=/packages/mnt/jbase/bin/hostname
kill      type=link size=28 time=1242346941.0 \
          link=/packages/mnt/jbase/bin/kill
ln        type=link size=26 time=1242346941.0 \
          link=/packages/mnt/jbase/bin/ln
ls        type=link size=26 time=1242346941.0 \
          link=/packages/mnt/jbase/bin/ls
mkdir     type=link size=29 time=1242346941.0 \
          link=/packages/mnt/jbase/bin/mkdir
mv        type=link size=26 time=1242346941.0 \
          link=/packages/mnt/jbase/bin/mv
pax       type=link size=27 time=1242346944.0 \
          link=/packages/mnt/jbase/bin/pax
ps        type=link size=26 time=1242346941.0 \
          link=/packages/mnt/jbase/bin/ps
pwd       type=link size=27 time=1242346941.0 \
          link=/packages/mnt/jbase/bin/pwd
rcp       type=link size=27 time=1242346942.0 \
          link=/packages/mnt/jbase/bin/rcp
red       type=link size=26 time=1242346941.0 \
          link=/packages/mnt/jbase/bin/red
rm        type=link size=26 time=1242346942.0 \
          link=/packages/mnt/jbase/bin/rm
rmdir     type=link size=29 time=1242346942.0 \
          link=/packages/mnt/jbase/bin/rmdir
sh        type=link size=26 time=1242346942.0 \
          link=/packages/mnt/jbase/bin/sh
sleep     type=link size=29 time=1242346942.0 \
          link=/packages/mnt/jbase/bin/sleep
stty      type=link size=28 time=1242346942.0 \
          link=/packages/mnt/jbase/bin/stty
sync      type=link size=28 time=1242346942.0 \
          link=/packages/mnt/jbase/bin/sync
tcsh      type=link size=27 time=1242346941.0 \
          link=/packages/mnt/jbase/bin/tcsh
test      type=link size=28 time=1242346942.0 \
          link=/packages/mnt/jbase/bin/test
# ./bin
...

```

show system audit root-only (QFX3500 Switch)

```

user@switch> show system audit root-only
#          user: root
#          machine: my-host
#          tree: /
date: Fri Feb 11 21:21:46 2000

# .
/set type=file uid=0 gid=0 mode=0755 nlink=1
.          type=dir nlink=23 size=1024 time=950252640.0
.cshrc     uid=3 gid=7 mode=0644 size=177 time=939182975.0 \
          md5digest=f414e06fea6bd646244b98e13d6e6226
.kernel.jkernel.backup \
          mode=0744 size=1934552 time=944688902.0 \
          md5digest=2c343cf0bd9fea8f04f78604feed7aa4
.profile   uid=3 gid=7 mode=0644 nlink=2 size=173 time=939182975.0 \
          md5digest=55a1e3c6c67789c9d3a1cce1ea39f670

```

```
COPYRIGHT uid=3 gid=7 mode=0444 size=3425 time=939182975.0 \  
md5digest=7df8bc77dcee71382ea73eb0ec6a9243  
boot.config mode=0644 size=3 time=945902618.0 \  
md5digest=93d722493ed38477338a1405d7dcb40  
boot.help uid=3 gid=7 mode=0444 size=411 time=939182876.0 \  
md5digest=9b7126385734bcae753f4179ab59d8e5  
compat type=link mode=0777 size=11 time=915149058.0 \  
link=/usr/compat  
kernel mode=0444 size=1947607 time=950230892.0 \  
md5digest=1a2a8aff2fec678a918ba0d6bf063980  
kernel.avr uid=1112 size=1947642 time=950252597.0 \  
md5digest=82e1637682d58ec28964dfee7fccb62e  
kernel.config \  
mode=0644 size=0 time=915149058.0 \  
md5digest=d41d8cd98f00b204e9800998ecf8427e  
sys type=link mode=0777 size=11 time=915149029.0 \  
link=usr/src/sys
```

show system boot-messages

List of Syntax	Syntax on page 1038 Syntax (EX Series Switches) on page 1038 Syntax (TX Matrix Router) on page 1038 Syntax (TX Matrix Plus Router) on page 1038 Syntax (MX Series Router) on page 1038 Syntax (QFX Series) on page 1038
Syntax	show system boot-messages
Syntax (EX Series Switches)	show system boot-messages <all-members> <local> <member <i>member-id</i> >
Syntax (TX Matrix Router)	show system boot-messages <all-chassis all-lcc lcc <i>number</i> scc>
Syntax (TX Matrix Plus Router)	show system boot-messages <all-chassis all-lcc lcc <i>number</i> sfc <i>number</i> >
Syntax (MX Series Router)	show system boot-messages <all-members> <local> <member <i>member-id</i> >
Syntax (QFX Series)	show system boot-messages infrastructure <i>name</i> interconnect-device <i>name</i> node-group <i>name</i>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Display initial messages generated by the system kernel upon startup. These messages are the contents of <code>/var/run/dmesg.boot</code> .
Options	none —Display all boot time messages. all-chassis —(TX Matrix routers and TX Matrix Plus routers only) (Optional) Display boot time messages for all of the chassis. all-lcc —(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display boot time messages for all T640 routers connected to a TX Matrix router. On a TX Matrix Plus router, display boot time messages for all connected T1600 or T4000 LCCs. all-members —(EX4200 switches and MX Series routers only) (Optional) Display boot time messages on all members of the Virtual Chassis configuration.

infrastructure *name*—(QFabric systems only) (Optional) Display boot time messages on the fabric control Routing Engine or fabric manager Routing engines.

interconnect-device *name*—(QFabric systems only) (Optional) Display boot time messages on the Interconnect device.

lcc *number*—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display boot time messages for a specific T640 router connected to a TX Matrix router. On a TX Matrix Plus router, display boot time messages for a specific router connected to a TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(EX4200 switches and MX Series routers only) (Optional) Display boot time messages on the local Virtual Chassis member.

member *member-id*—(EX4200 switches and MX Series routers only) (Optional) Display boot time messages on the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

node-group *name*—(QFabric systems only) (Optional) Display boot time messages on the Node group.

scc—(TX Matrix routers only) (Optional) Display boot time messages for the TX Matrix router (or switch-card chassis).

sfc *number*—(TX Matrix Plus routers only) (Optional) Display boot time messages for the TX Matrix Plus router. Replace *number* with 0.

Additional Information By default, when you issue the **show system boot-messages** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

Required Privilege Level view

Related Documentation

- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

List of Sample Output

[show system boot-messages \(TX Matrix Router\) on page 1040](#)
[show system boot-messages lcc \(TX Matrix Router\) on page 1041](#)
[show system boot-messages \(TX Matrix Plus Router\) on page 1042](#)
[show system boot-messages \(QFX3500 Switch\) on page 1042](#)

Sample Output**show system boot-messages (TX Matrix Router)**

```
user@host> show system boot-messages
Copyright (c) 1992-1998 FreeBSD Inc.
Copyright (c) 1996-2000 Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1982, 1986, 1989, 1991, 1993
    The Regents of the University of California. All rights reserved.

JUNOS 4.1-20000216-Zf8469 #0: 2000-02-16 12:57:28 UTC
    tlim@single.juniper.net:/p/build/20000216-0905/4.1/release_kernel/sys/compil
e/GENERIC
CPU: Pentium Pro (332.55-MHz 686-class CPU)
    Origin = "GenuineIntel" Id = 0x66a Stepping=10
    Features=0x183f9ff<FPU,VME,DE,PSE,TSC,MSR,PAE,MCE,CX8,SEP,MTRR,PGE,MCA,CMOV,<b
16>,<b17>,MMX,<b24>>
Teknor CPU Card Recognized
real memory = 805306368 (786432K bytes)
avail memory = 786280448 (767852K bytes)
Probing for devices on PCI bus 0:
chip0 <generic PCI bridge (vendor=8086 device=7192 subclass=0)> rev 3 class 6000
0 on pci0:0:0
chip1 <Intel 82371AB PCI-ISA bridge> rev 1 class 60100 on pci0:7:0
chip2 <Intel 82371AB IDE interface> rev 1 class 10180 on pci0:7:1
chip3 <Intel 82371AB USB interface> rev 1 class c0300 int d irq 11 on pci0:7:2
smb0 <Intel 82371AB SMB controller> rev 1 class 68000 on pci0:7:3
pcic0 <TI PCI-1131 PCI-CardBus Bridge> rev 1 class 60700 int a irq 15 on pci0:13
:0
TI1131 PCI Config Reg: [pci only][FUNC0 pci int]
pcic1 <TI PCI-1131 PCI-CardBus Bridge> rev 1 class 60700 int b irq 12 on pci0:13
:1
TI1131 PCI Config Reg: [pci only][FUNC1 pci int]
fxp0 <Intel EtherExpress Pro 10/100B Ethernet> rev 8 class 20000 int a irq 12 on

pci0:16:0
chip4 <generic PCI bridge (vendor=1011 device=0022 subclass=4)> rev 4 class 6040
0 on pci0:17:0
fxp1 <Intel EtherExpress Pro 10/100B Ethernet> rev 8 class 20000 int a irq 10 on

pci0:19:0
Probing for devices on PCI bus 1:
mcs0 <Miscellaneous Control Subsystem> rev 12 class ff0000 int a irq 12 on pci1:
13:0
fxp2 <Intel EtherExpress Pro 10/100B Ethernet> rev 8 class 20000 int a irq 10 on

pci1:14:0
Probing for devices on the ISA bus:
sc0 at 0x60-0x6f irq 1 on motherboard
sc0: EGA color <16 virtual consoles, flags=0x0>
ed0 not found at 0x300
```



```

ed1 not found at 0x280
ed2 not found at 0x340
psm0 not found at 0x60
sio0 at 0x3f8-0x3ff irq 4 flags 0x20010 on isa
sio0: type 16550A, console
sio1 at 0x3e8-0x3ef irq 5 flags 0x20000 on isa
sio1: type 16550A
sio2 at 0x2f8-0x2ff irq 3 flags 0x20000 on isa
sio2: type 16550A
pcic0 at 0x3e0-0x3e1 on isa
PC-Card ctlr(0) TI PCI-1131 [CardBus bridge mode] (5 mem & 2 I/O windows)
pcic0: slot 0 controller I/O address 0x3e0
npx0 flags 0x1 on motherboard
npx0: INT 16 interface
fdc0: direction bit not set
fdc0: cmd 3 failed at out byte 1 of 3
fdc0 not found at 0x3f0
wdc0 at 0x1f0-0x1f7 irq 14 on isa
wdc0: unit 0 (wd0): <SunDisk SQFXB-80>, single-sector-i/o
wd0: 76MB (156672 sectors), 612 cyls, 8 heads, 32 S/T, 512 B/S
wdc0: unit 1 (wd1): <IBM-DCXA-210000>
wd1: 8063MB (16514064 sectors), 16383 cyls, 16 heads, 63 S/T, 512 B/S
wdc1 not found at 0x170
wdc2 not found at 0x180
ep0 not found at 0x300
fxp0: Ethernet address 00:a0:a5:12:05:5a
fxp1: Ethernet address 00:a0:a5:12:05:59
fxp2: Ethernet address 02:00:00:00:00:01
swapon: adding /dev/wd1s1b as swap device
Automatic reboot in progress...
/dev/rwd0s1a: clean, 16599 free (95 frags, 2063 blocks, 0.1% fragmentation)
/dev/rwd0s1e: clean, 9233 free (9 frags, 1153 blocks, 0.1% fragmentation)
/dev/rwd0s1a: clean, 16599 free (95 frags, 2063 blocks, 0.1% fragmentation)
/dev/rwd1s1f: clean, 4301055 free (335 frags, 537590 blocks, 0.0% fragmentation)

```

show system boot-messages lcc (TX Matrix Router)

```

user@host> show system boot-messages lcc 2
lcc2-re0:
-----
Copyright (c) 1996-2001, Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1992-2001 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
    The Regents of the University of California. All rights reserved.
JUNOS 7.0-20040912.0 #0: 2004-09-12 09:16:32 UTC

builder@benten.juniper.net:/build/benten-b/7.0/20040912.0/obj-i386/sys/compile/JUNIPER
Timecounter "i8254" frequency 1193182 Hz
Timecounter "TSC" frequency 601368936 Hz
CPU: Pentium III/Pentium III Xeon/Celeron (601.37-MHz 686-class CPU)
    Origin = "GenuineIntel" Id = 0x68a Stepping = 10

Features=0x387f9ff<FPU,VME,DE,PSE,TSC,MSR,PAE,MCE,CX8,SEP,MTRR,PGE,MCA,CMOV,PAT,PSE36,PN,MMX,FXSR,SSE>
real memory = 2147467264 (2097136K bytes)
sio0: gdb debugging port
avail memory = 2084040704 (2035196K bytes)
Preloaded elf kernel "kernel" at 0xc06d9000.
DEVFS: ready for devices
Pentium Pro MTRR support enabled
md0: Malloc disk

```

```

DRAM Data Integrity Mode: ECC Mode with h/w scrubbing
npx0: <math processor> on motherboard
npx0: INT 16 interface
pcib0: <ServerWorks NB6635 3.0LE host to PCI bridge> on motherboard
pci0: <PCI bus> on pcib0
pcic-pci0: <TI PCI-1410 PCI-CardBus Bridge> irq 15 at device 1.0 on pci0
pcic-pci0: TI12XX PCI Config Reg: [pwr save][pci only]
fxp0: <Intel Embedded 10/100 Ethernet> port 0x1000-0x103f mem
0xfb800000-0xfb81ffff,0xfb820000-0xfb820fff irq 9 at device 3.0 on pci0
fxp1: <Intel Embedded 10/100 Ethernet> port 0x1040-0x107f mem
0xfb840000-0xfb85ffff,0xfb821000-0xfb821fff irq 11 at device 4.0 on pci0
...

```

show system boot-messages (TX Matrix Plus Router)

```

user@host> show system boot-messages
sfc0-re0:
-----
Copyright (c) 1996-2009, Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1992-2006 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
    The Regents of the University of California. All rights reserved.
JUNOS 9.6B3.3 #0: 2009-06-17 19:52:08 UTC

builder@lanath.juniper.net:/volume/build/junos/9.6/release/9.6B3.3/obj-i386/bsd/sys/compile/JUNIPER
MPTable: Timecounter "i8254" frequency 1193182 Hz quality 0 CPU: Intel(R) Xeon(R)
CPU          L5238 @ 2.66GHz (2660.01-MHz 686-class CPU)   Origin =
"GenuineIntel" Id = 0x1067a Stepping = 10   Features=0xbfebfbff
...
lcc1-re0:
-----
Copyright (c) 1996-2009, Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1992-2006 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
    The Regents of the University of California. All rights reserved.
JUNOS 9.6-20090617.0 #0: 2009-06-17 04:15:14 UTC

builder@lanath.juniper.net:/volume/build/junos/9.6/production/20090617.0/obj-i386/bsd/sys/compile/JUNIPER
Timecounter "i8254" frequency 1193182 Hz quality 0
CPU: Intel(R) Xeon(R) CPU          @ 1.86GHz (1862.01-MHz 686-class CPU)

Origin = "GenuineIntel" Id = 0x1067a Stepping = 10
Features=0xbfebfbff
...

```

show system boot-messages (QFX3500 Switch)

```

user@switch> show sytem boot-messages
getmemsize: msgbufp[size=32768] = 0x81d07fe4

System physical memory distribution:
-----
Total physical memory: 4160749568 (3968 MB)
Physical memory used: 3472883712 (3312 MB)
Physical memory allocated to kernel: 2130706432 (2032 MB)
Physical memory allocated to user BTLB: 1342177280 (1280 MB)
-----

Copyright (c) 1996-2010, Juniper Networks, Inc.

```

All rights reserved.

Copyright (c) 1992-2006 The FreeBSD Project.

Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994

The Regents of the University of California. All rights reserved.

JUNOS 11.1I #0: 2010-09-17 19:18:07 UTC

```
ssiano@svl-junos-pool125.juniper.net:/c/ssiano/DEV_QFX_SI_BRANCH/03/20100917.399988/
obj-xlr/bsd/sys/compile/JUNIPER-DCTOR
WARNING: debug.mpsafenet forced to 0 as ipsec requires Giant
JUNOS 11.1I #0: 2010-09-17 19:18:07 UTC
```

```
ssiano@svl-junos-pool125.juniper.net:/c/ssiano/DEV_QFX_SI_BRANCH/03/20100917.399988/
obj-xlr/bsd/sys/compile/JUNIPER-DCTOR
real memory = 3472883712 (3312MB)
avail memory = 1708171264 (1629MB)
cpuid: 0, bt1b_cpumap:0xffffffff8
FreeBSD/SMP: Multiprocessor System Detected: 12 CPUs
ETHERNET SOCKET BRIDGE initialising
Initializing QFX platform properties ..
cpu0 on motherboard
: RMI's XLR CPU Rev. 0.3 with no FPU implemented
  L1 Cache: I size 32kb(32 line), D size 32kb(32 line), eight way.
  L2 Cache: Size 1024kb, eight way
pic_lbus0: <XLR Local Bus>
pic_lbus0: <XLR Local Bus> on motherboard
Enter qfx control ethernet probe addr:0xc5eeec00
gmac4: <XLR GMAC GE Ethernet> on pic_lbus0
me0: Ethernet address 00:1d:b5:f7:68:40
Enter qfx control ethernet probe addr:0xc5eeeb40
gmac5: <XLR GMAC GE Ethernet> on pic_lbus0
me1: Ethernet address 00:1d:b5:f7:68:41
Enter qfx control ethernet probe addr:0xc5eeea80
gmac6: <XLR GMAC GE Ethernet> on pic_lbus0
me1: Ethernet address 00:1d:b5:f7:68:42
sio0 on pic_lbus0
Entering sioattach
sio0: type 16550A, console
xls_setup_intr: skip irq 3, xlr regs are set up somewhere else.
gblmem0 on pic_lbus0
ehci0: <RMI XLS USB 2.0 controller> on pic_lbus0
ehci_bus_attach: allocated resource. tag=1, base=bef24000
xls_ehci_init: endian hardware swapping NOT enabled.
usb0: EHCI version 1.0
usb0 on ehci0
usb0: USB revision 2.0
uhub0: vendor 0x0000 EHCI root hub, class 9/0, rev 2.00/1.00, addr 1
uhub0: 2 ports with 2 removable, self powered
umass0: USB USBFlashDrive, rev 2.00/11.00, addr 2
pcib0: PCIe link 0 up
pcib0: PCIe link 2 up
pcib0: PCIe link 3 up
pcib0: <XLS PCI Host Controller> on pic_lbus0
pci0: <PCI bus> on pcib0
pcib1: <PCI-PCI bridge> at device 0.0 on pci0
pci1: <PCI bus> on pcib1
pci1: <network, ethernet> at device 0.0 (no driver attached)
pcib2: <PCI-PCI bridge> at device 1.0 on pci0
pcib3: <PCI-PCI bridge> at device 2.0 on pci0
pci2: <PCI bus> on pcib3
pci2: <network, ethernet> at device 0.0 (no driver attached)
pcib4: <PCI-PCI bridge> at device 3.0 on pci0
```

```
pci3: <PCI bus> on pcib4
pci3: <network, ethernet> at device 0.0 (no driver attached)
cfi device address space at 0xbc000000
cfi0: <AMD/Fujitsu - 8MB> on pic_lbus0
cfi device address space at 0xbc000000
i2c0: <I2C bus controller> on pic_lbus0
i2c1: <I2C bus controller> on pic_lbus0
qfx_fmn0 on pic_lbus0
pool offset 1503776768
xlr_lbus0: <XLR Local Bus Controller> on motherboard
qfx_bcpld_probe[124]
qfx_bcpld_probe[138]: dev_type=0x0
qfx_bcpld_probe[124]
qfx_bcpld0: QFX BCPLD probe success
qfx_bcpld0qfx_bcpld_attach[174]
qfx_bcpld_attach[207] : bus_space_tag=0x0, bus_space_handle=0xbd900000
qfx_bcpld_probe[124]
qfx_bcpld1: QFX BCPLD probe success
qfx_bcpld1qfx_bcpld_attach[174]
tor_bcpld_slave_attach[1245] : bus_space_tag=0x0, bus_space_handle=0xbda00000
Initializing product: 96 ..
bmeb: bmeb_lib_init done 0xc60a5000, addr 0x809c99a0
bme0:Virtual BME driver initializing
Timecounter "mips" frequency 1200000000 Hz quality 0
Timecounter "xlr_pic_timer" frequency 66666666 Hz quality 1
Timecounters tick every 1.000 msec
Loading the NETPFE fc module
IPsec: Initialized Security Association Processing.
SMP: AP CPU #3 Launched!
SMP: AP CPU #1 Launched!
SMP: AP CPU #2 Launched!
SMP: AP CPU #4 Launched!
SMP: AP CPU #5 Launched!
SMP: AP CPU #7 Launched!
SMP: AP CPU #6 Launched!
SMP: AP CPU #11 Launched!
SMP: AP CPU #10 Launched!
SMP: AP CPU #9 Launched!
SMP: AP CPU #8 Launched!
da0 at umass-sim0 bus 0 target 0 lun 0
da0: <USB USBFlashDrive 1100> Removable Direct Access SCSI-0 device
da0: 40.000MB/s transfers
da0: 3920MB (8028160 512 byte sectors: 255H 63S/T 499C)
Trying to mount root from ufs:/dev/da0s1a
```

show system buffers

List of Syntax	Syntax on page 1045 Syntax (EX Series) on page 1045 Syntax (TX Matrix Router) on page 1045 Syntax (TX Matrix Plus Router) on page 1045 Syntax (MX Series Router) on page 1045 Syntax (QFX Series) on page 1045
Syntax	show system buffers
Syntax (EX Series)	show system buffers <all-members> <local> <member <i>member-id</i> >
Syntax (TX Matrix Router)	show system buffers <all-chassis all-lcc lcc <i>number</i> scc>
Syntax (TX Matrix Plus Router)	show system buffers <all-chassis all-lcc lcc <i>number</i> sfc <i>number</i> >
Syntax (MX Series Router)	show system buffers <all-members> <local> <member <i>member-id</i> >
Syntax (QFX Series)	show system buffers <infrastructure <i>name</i> interconnect-device <i>name</i> node-group <i>name</i> root-only (infrastructure <i>name</i> interconnect-device <i>name</i> node-group <i>name</i>)>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Display information about the buffer pool that the Routing Engine uses for local traffic. Local traffic is the routing and management traffic that is exchanged between the Routing Engine and the Packet Forwarding Engine within the router or switch, as well as the routing and management traffic from IP (that is, from OSPF, BGP, SNMP, ping operations, and so on).
Options	none —Show all buffer statistics. all-lcc —(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, show buffer statistics for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, show buffer statistics for all routers connected to the TX Matrix Plus router.

all-chassis—(TX Matrix routers and TX Matrix Plus routers only) (Optional) Show buffer statistics for all of the chassis.

all-members—(EX4200 switches and MX Series routers only) (Optional) Show buffer statistics for all members of the Virtual Chassis configuration.

infrastructure *name*—(QFabric systems only) (Optional) Show buffer statistics for a fabric control Routing Engine or a fabric control Routing Engine.

interconnect-device *name*—(QFabric systems only) (Optional) Show buffer statistics for the Interconnect device.

lcc *number*—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, show buffer statistics for a specific T640 router (or line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, show buffer statistics for a specific router (line-card chassis) that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(EX4200 switches and MX Series routers only) (Optional) Show buffer statistics for the local Virtual Chassis member.

member *member-id*—(EX4200 switches and MX Series routers only) (Optional) Show buffer statistics for the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

node-group *name*—(QFabric systems only) (Optional) Show buffer statistics for the Node group

sfc—(TX Matrix Plus routers only) (Optional) Show buffer statistics for the TX Matrix Plus router. Replace *number* with 0.

Additional Information By default, when you issue the **show system buffers** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

A special type of memory buffer called a *cluster* is 2 KB in size. For more information, see *The Design and Implementation of the 4.4BSD Operation System* by McKusic, Bostic, Karels, and Quarterman.

Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• Routing Matrix with a TX Matrix Plus Router Solutions Page
List of Sample Output	show system buffers on page 1048 show system buffers scc (TX Matrix Router) on page 1049 show system buffers sfc (TX Matrix Plus Router) on page 1049 show system buffers all-chassis (TX Matrix Plus Router) on page 1049 show system buffers node-group (QFabric System) on page 1050
Output Fields	Table 70 describes the output fields for the show system buffers command. Output fields are listed in the approximate order in which they appear.

Table 70: show system buffers Output Fields

Field Name	Field Description
mbufs in use	Memory buffers (mbufs) are 128-byte buffers that are used for various purposes inside the kernel. Each memory buffer has a type, and the output itemizes the amount allocated for each type. Types with no memory buffers allocated are not displayed.
mbufs allocated to packet headers	Number of memory buffers currently holding packet headers
mbufs allocated to control blocks	Number of memory buffers currently holding the state for sockets.
mbufs allocated to send data	Number of memory buffers currently holding socket send data.
mbufs allocated to pfe refill data	Number of memory buffers currently holding Packet Forwarding Engine refill data.
mbufs allocated to fxp data	Number of memory buffers currently holding fxp data.
mbufs allocated to socket names and addresses	Number of memory buffers currently holding addresses for sockets.
mbuf clusters in use	Allocation statistics for memory buffer clusters.
allocated to network	Total amount of memory in use by the networking and interprocess communication (IPC) code.
requests for memory denied	Number of times a memory allocation request within the IPC and networking code failed.
requests for memory delayed	Number of times a memory allocation request within the IPC and networking code was postponed.
calls to protocol drain routines	Number of times a memory allocation request within the IPC and networking code triggered a memory reclamation attempt.

Sample Output

show system buffers

```

user@host> show system buffers
397/893/1290 mbufs in use (current/cache/total)
395/331/726/30000 mbuf clusters in use (current/cache/total/max)
384/256 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
889K/885K/1774K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/5/1024 sfbufs in use (current/peak/max)

```



```

0 requests for sbufs denied
0 requests for sbufs delayed
0 requests for I/O initiated by sendfile
0 calls to protocol drain routines

```

show system buffers scc (TX Matrix Router)

```

user@host> show system buffers scc
213 mbufs in use:
    11 mbufs allocated to packet headers
    26 mbufs allocated to socket names and addresses
    2 mbufs allocated to socket options
    17 mbufs allocated to socket send data
    2 mbufs allocated to pfe data
    155 mbufs allocated to fxp data (rx)
    511 mbufs allocated to <mbuf type 86>
    256 mbufs allocated to <mbuf type 92>
924/1162 mbuf clusters in use
2788 Kbytes allocated to network (75% in use)
0 requests for memory denied
0 requests for memory delayed
0 calls to protocol drain routines

```

show system buffers sfc (TX Matrix Plus Router)

```

user@host> show system buffers sfc 0

sfc0-re0:
-----
4363/2807/7170 mbufs in use (current/cache/total)
4358/1968/6326/30000 mbuf clusters in use (current/cache/total/max)
256/128 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
9806K/4637K/14444K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/10/1024 sbufs in use (current/peak/max)
0 requests for sbufs denied
0 requests for sbufs delayed
0 requests for I/O initiated by sendfile
0 calls to protocol drain routines

```

show system buffers all-chassis (TX Matrix Plus Router)

```

user@host> show system buffers all-chassis

sfc0-re0:
-----
4363/2807/7170 mbufs in use (current/cache/total)
4358/1968/6326/30000 mbuf clusters in use (current/cache/total/max)
256/128 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
9806K/4637K/14444K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/10/1024 sbufs in use (current/peak/max)
0 requests for sbufs denied
0 requests for sbufs delayed

```

```
0 requests for I/O initiated by sendfile
0 calls to protocol drain routines
```

```
lcc0-re0:
```

```
-----
772/2558/3330 mbufs in use (current/cache/total)
772/598/1370/30000 mbuf clusters in use (current/cache/total/max)
768/512 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
1737K/1835K/3572K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/4/1024 sbufs in use (current/peak/max)
0 requests for sbufs denied
0 requests for sbufs delayed
0 requests for I/O initiated by sendfile
0 calls to protocol drain routines
```

```
lcc1-re0:
```

```
-----
773/2437/3210 mbufs in use (current/cache/total)
773/453/1226/30000 mbuf clusters in use (current/cache/total/max)
768/384 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
1739K/1515K/3254K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/7/1024 sbufs in use (current/peak/max)
0 requests for sbufs denied
0 requests for sbufs delayed
0 requests for I/O initiated by sendfile
0 calls to protocol drain routines
```

```
lcc2-re0:
```

```
-----
816/2514/3330 mbufs in use (current/cache/total)
816/554/1370/30000 mbuf clusters in use (current/cache/total/max)
768/512 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
1836K/1736K/3572K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/4/1024 sbufs in use (current/peak/max)
0 requests for sbufs denied
0 requests for sbufs delayed
0 requests for I/O initiated by sendfile
```

show system buffers node-group (QFabric System)

```
user@switch> show system buffers node-group node1
node-group node1:
```

```
-----
2/2698/2700 mbufs in use (current/cache/total)
2/1520/1522/30000 mbuf clusters in use (current/cache/total/max)
0/1280 mbuf+clusters out of packet secondary zone in use (current/cache)
```

```

0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
4K/3714K/3719K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/6/6656 sfbufs in use (current/peak/max)
0 requests for sfbufs denied
0 requests for sfbufs delayed
0 requests for I/O initiated by sendfile
0 calls to protocol drain routines

```

```
re0:
```

```

-----
516/639/1155 mbufs in use (current/cache/total)
515/147/662/30000 mbuf clusters in use (current/cache/total/max)
512/128 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
1159K/453K/1612K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/4/1024 sfbufs in use (current/peak/max)
0 requests for sfbufs denied
0 requests for sfbufs delayed
0 requests for I/O initiated by sendfile
0 calls to protocol drain routines

```

```
re1:
```

```

-----
519/771/1290 mbufs in use (current/cache/total)
518/176/694/30000 mbuf clusters in use (current/cache/total/max)
512/128 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
1165K/544K/1710K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/4/1024 sfbufs in use (current/peak/max)
0 requests for sfbufs denied
0 requests for sfbufs delayed
0 requests for I/O initiated by sendfile
0 calls to protocol drain routines

```

show system certificate

Syntax	<code>show system certificate</code> <code><certificate-id></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	(Encryption interface on M Series, T Series routers, QFX Series, and OCX Series switches only) Display installed certificates signed by the Juniper Networks certificate authority.
Options	none —Display all installed certificates signed by the Juniper Networks certificate authority. certificate-id —(Optional) Display the details of a particular certificate.
Required Privilege Level	maintenance
List of Sample Output	show system certificate on page 1053 show system certificate (QFX Series) on page 1053
Output Fields	Table 71 lists the output fields for the show system certificate command. Output fields are listed in the approximate order in which they appear.

Table 71: show system certificate Output Fields

Field Name	Field Description
Certificate identifier	Unique identifier associated with a certificate. The certificate identifier is the common name of the subject.
Issuer Subject	Information about the certificate issuer and the distinguished name (DN) of the issuer, respectively: <ul style="list-style-type: none"> • Organization—Name of the owner's organization. • Organizational unit—Name of the owner's department. • Country—Two-character country code in which the owner's system is located. • State—State in the USA in which the owner is using the certificate. • Locality—City in which the owner's system is located. • Common name—Name of the owner of the certificate. • E-mail address—E-mail address of the owner of the certificate.
Validity	When a certificate is valid.
Signature algorithm	Encryption algorithm applied to the installed certificate.
Public key algorithm	Encryption algorithm applied to the public key.

Sample Output

show system certificate

```
user@host> show system certificate
Certificate identifier: Dallas-v3
Issuer:
Organization: Juniper Networks, Organizational unit: Juniper CA,
Country: US, State: CA, Locality: Sunnyvale, Common name: Dallas CA,
E-mail address:ca@juniper.net
Subject:
Organization: Juniper Networks, Organizational unit: Juniper CA,
Country: US, State: CA, Locality: Sunnyvale, Common name: Dallas-v3,
E-mail address:ca@juniper.net
Validity:
Not before: Mar 13 03:23:25 2004 GMT
Not after: Mar 24 03:23:25 2014 GMT
Signature algorithm: sha1WithRSAEncryption
Public key algorithm: dsaEncryption
```

show system certificate (QFX Series)

```
user@host> show system certificate
Certificate identifier: Dallas-v3
Issuer:
Organization: Juniper Networks, Organizational unit: Juniper CA,
Country: US, State: CA, Locality: Sunnyvale, Common name: Dallas CA,
E-mail address:ca@juniper.net
Subject:
Organization: Juniper Networks, Organizational unit: Juniper CA,
Country: US, State: CA, Locality: Sunnyvale, Common name: Dallas-v3,
E-mail address:ca@juniper.net
Validity:
Not before: Mar 13 03:23:25 2004 GMT
Not after: Mar 24 03:23:25 2014 GMT
Signature algorithm: sha1WithRSAEncryption
Public key algorithm: dsaEncryption
```

show system commit


Syntax	show system commit <revision> <server>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Option server introduced in Junos OS Release 12.1 for the PTX Series router.</p> <p>Option revision introduced in Junos OS Release 14.1.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.</p>
Description	Display the system commit history and any pending commit operation.
Options	<p>none—Display the last 50 commit operations listed, most recent to first.</p> <p>revision—(Optional) Display the revision number of the active configuration of the Routing Engine(s).</p> <p>server—(Optional) Display commit server status.</p>
<div>  <p>NOTE: By default, the status of the commit server is “Not running”. The commit server starts running only when a commit job is added to the batch.</p> </div>	
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> clear system commit on page 336 show system commit revision
List of Sample Output	show system commit on page 1056 show system commit (At a Particular Time) on page 1056 show system commit (At the Next Reboot) on page 1056 show system commit (Rollback Pending) on page 1056 show system commit (QFX Series) on page 1056
Output Fields	Table 72 describes the output fields for the show system commit command. Output fields are listed in the approximate order in which they appear.

Table 72: show system commit Output Fields

Field Name	Field Description	Level of Output
<number>	Displays the last 50 commit operations listed, most recent to first. The identifier <number> designates a configuration created for recovery using the request system configuration rescue save command.	none

Table 72: show system commit Output Fields (*continued*)

Field Name	Field Description	Level of Output
<time-stamp>	Date and time of the commit operation.	none
<root>/<username>	User who executed the commit operation.	none
<method>	Method used to execute the commit operation: <ul style="list-style-type: none"> • CLI—CLI interactive user performed the commit operation. • Junos XML protocol—Junos XML protocol client performed the commit operation. • synchronize—The commit synchronize command was performed on the other Routing Engine. • snmp—An SNMP set request caused the commit operation. • button—A button on the router or switch was pressed to commit a rescue configuration for recovery. • autoinstall—A configuration obtained through autoinstallation was committed. • other—When there is no login name associated with the session, the values for user and client default to root and other. For example, during a reboot after package installation, mgd commits the configuration as a system commit, and there is no login associated with the commit. 	none

Sample Output

show system commit

```
user@host> show system commit
0   2003-07-28 19:14:04 PDT by root via other
1   2003-07-25 22:01:36 PDT by regress via cli
2   2003-07-25 22:01:32 PDT by regress via cli
3   2003-07-25 21:30:13 PDT by root via button
4   2003-07-25 13:46:48 PDT by regress via cli
5   2003-07-25 05:33:21 PDT by root via autoinstall
...
rescue 2002-05-10 15:32:03 PDT by root via other
```

show system commit (At a Particular Time)

```
user@host> show system commit
commit requested by root via cli at Tue May  7 15:59:00 2002
```

show system commit (At the Next Reboot)

```
user@host> show system commit
commit requested by root via cli at reboot
```

show system commit (Rollback Pending)

```
user@host> show system commit
0 2005-01-05 15:00:37 PST by root via cli commit confirmed, rollback in 3mins
```

show system commit (QFX Series)

```
user@switch> show system commit
0 2011-11-25 19:17:49 PST by root via cli
```


show system configuration archival

Syntax show system configuration archival

Release Information Introduced in Junos OS Release 7.6.
 Command introduced in Junos OS Release 9.0 for EX Series switches.
 Command introduced in Junos OS Release 11.1 for the QFX Series.
 Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Display directory and number of files queued for archival transfer.



NOTE: The [edit system configuration] hierarchy is not available on QFabric systems.

Options This command has no options.

Required Privilege Level maintenance

List of Sample Output [show system configuration archival on page 1057](#)

Sample Output

show system configuration archival

```
user@host> show system configuration archival
```

```
/var/transfer/config/:
total 8
```

show system configuration rescue

Syntax show system configuration rescue

Release Information Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.0 for EX Series switches.
 Command introduced in Junos OS Release 11.1 for the QFX Series.
 Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.

Description Display a rescue configuration, if one exists.



NOTE: The [edit system configuration] hierarchy is not available on QFabric systems.

Options This command has no options.

Required Privilege Level maintenance

Related Documentation • [show system configuration archival on page 1057](#)

List of Sample Output [show system configuration rescue on page 1058](#)

Sample Output

show system configuration rescue

```
user@switch> show system configuration rescue
version "7.3"; groups {
  global {
    system {
      host-name router1;
      domain-name customer.net;
      domain-search [ customer.net ];
      backup-router 192.168.124.254;
      name-server {
        172.17.28.11;
        172.17.28.101;
        172.17.28.100;
        172.17.28.10;
      }
      login {
        user regress {
          uid 928;
          class ;
          shell csh;
          authentication {
            encrypted-password "$1$kPU..$w.4FGRAGanJ8U4Yq6sbj7."; ##
SECRET-DATA
          }
        }
      }
    }
  }
}
```

```
services {  
    ftp;  
    rlogin;  
    rsh;  
    telnet;  
}  
}  
.....
```

show system connections

List of Syntax [Syntax on page 1060](#)
 [Syntax \(EX Series\) on page 1060](#)
 [Syntax \(TX Matrix Router\) on page 1060](#)
 [Syntax \(TX Matrix Plus Router\) on page 1060](#)
 [Syntax \(MX Series Router\) on page 1060](#)
 [Syntax \(QFX Series\) on page 1060](#)
 [Syntax \(OCX Series\) on page 1060](#)

Syntax `show system connections`
 `<extensive>`
 `<all-chassis | all-lcc | lcc number | scc>`
 `<inet | inet6>`
 `<show-routing-instances>`

Syntax (EX Series) `show system connections`
 `<extensive>`
 `<all-members>`
 `<inet | inet6>`
 `<local>`
 `<member member-id>`
 `<show-routing-instances>`

Syntax (TX Matrix Router) `show system connections`
 `<extensive>`
 `<all-chassis | all-lcc | lcc number | scc>`
 `<inet | inet6>`
 `<show-routing-instances>`

Syntax (TX Matrix Plus Router) `show system connections`
 `<extensive>`
 `<all-chassis | all-lcc | lcc number | sfc number>`
 `<inet | inet6>`
 `<show-routing-instances>`

Syntax (MX Series Router) `show system connections`
 `<extensive>`
 `<all-members>`
 `<inet | inet6>`
 `<local>`
 `<member member-id>`
 `<show-routing-instances>`

Syntax (QFX Series) `show system connections`
 `<extensive>`
 `<inet>`
 `<infrastructure name>`
 `<interconnect-device name>`
 `<node-group name>`
 `<show-routing-instances>`

Syntax (OCX Series) `show system connections`

<extensive>
 <inet>
 <show-routing-instances>

Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Display information about the active IP sockets on the Routing Engine. Use this command to verify which servers are active on a system and what connections are currently in progress.</p>
Options	<p>none—Display information about all active IP sockets on the Routing Engine.</p> <p>extensive—(Optional) Display exhaustive system process information, which, for TCP connections, includes the TCP control block. This option is useful for debugging TCP connections.</p> <p>all-chassis—(TX Matrix routers and TX Matrix Plus routers only) (Optional) Display system connection activity for all the routers in the chassis.</p> <p>all-lcc—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system connection activity for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, display system connection activity for all connected T1600 or T4000 LCCs</p> <p>all-members—(EX4200 switches and MX Series routers only) (Optional) Display system connection activity for all members of the Virtual Chassis configuration.</p> <p>inet inet6—(Optional) Display IPv4 connections or IPv6 connections, respectively.</p> <p>infrastructure <i>name</i>—(QFabric systems only) (Optional) Display system connection activity for the fabric control Routing Engines or fabric manager Routing Engines.</p> <p>interconnect-device <i>name</i>—(QFabric systems only) (Optional) Display system connection activity for the Interconnect device.</p> <p>lcc <i>number</i>—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system connection activity for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display system connection activity for a specific router that is connected to the TX Matrix Plus router. Replace <i>number</i> with the following values depending on the LCC configuration:</p> <ul style="list-style-type: none"> • 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix. • 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.

- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(EX4200 switches and MX Series routers only) (Optional) Display system connection activity for the local Virtual Chassis member.

member *member-id*—(EX4200 switches and MX Series routers only) (Optional) Display system connection activity for the specified member of the Virtual Chassis configuration. For EX4200 switches, replace ***member-id*** with a value from 0 through 9. For an MX Series Virtual Chassis, replace ***member-id*** with a value of 0 or 1.

node-group *name*—(QFabric systems only) (Optional) Display system connection activity for the Node group.

scc—(TX Matrix routers only) (Optional) Display system connection activity for the TX Matrix router (or switch-card chassis).

sfc—(TX Matrix routers only) (Optional) Display system connection activity for the TX Matrix Plus router.

show-routing-instances—(Optional) Display routing instances.

Additional Information By default, when you issue the **show system connections** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

Required Privilege Level view

Related Documentation

- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

List of Sample Output [show system connections on page 1063](#)
[show system connections extensive on page 1064](#)
[show system connections lcc \(TX Matrix Router\) on page 1065](#)
[show system connections show-routing-instances on page 1065](#)
[show system connections \(TX Matrix Plus Router\) on page 1066](#)
[show system connections sfc \(TX Matrix Plus Router\) on page 1069](#)
[show system connections show-routing-instances \(TX Matrix Plus Router\) on page 1071](#)
[show system connections \(QFX3500 Switch\) on page 1076](#)

Output Fields [Table 73](#) describes the output fields for the **show system connections** command. Output fields are listed in the approximate order in which they appear.

Table 73: show system connections Output Fields

Field Name	Field Description
Proto	Protocol of the socket: IP , TCP , or UDP for IPv4 or IPv6.
Recv-Q	Number of input packets received by the protocol and waiting to be processed by the application.
Send-Q	Number of output packets sent by the application and waiting to be processed by the protocol.
Local Address	Local address and port of the socket, separated by a period. An asterisk (*) indicates that the bound address is the wildcard address. Server sockets typically have the wildcard address and a well-known port bound to them.
Foreign Address	Foreign address and port of the socket, separated by a period. An asterisk (*) indicates that the address or port is a wildcard.
Routing Instance	(Displayed only when the show-routing-instance option is used.) Routing instances associated with active IP sockets on the Routing Engine.
(state)	For TCP, the protocol state of the socket.

Sample Output

show system connections

```

user@host> show system connections
Active Internet connections (including servers)
Proto Recv-Q Send-Q Local Address           Foreign Address          (state)
tcp    0      0 192.168.4.16.513        208.197.169.254.894     ESTABLISHED
tcp    0      0 192.168.4.16.513        208.197.169.195.945     ESTABLISHED
tcp    0      0 *.23                    *.*                      LISTEN
tcp    0      0 *.22                    *.*                      LISTEN
tcp    0      0 *.513                   *.*                      LISTEN
tcp00 *.514                *.*                      LISTEN
tcp 0 0*.21                    *.*                      LISTEN
tcp00 *.79                *.*                      LISTEN
tcp 00 *.1023                 *.*                      LISTEN
tcp 00 *.111                  *.*                      LISTEN
udp00192.168.4.16.1634    208.197.169.249.2049
udp00192.168.4.16.1627    208.197.169.254.2049
udp00192.168.4.16.1371    208.197.169.195.2049
udp00*.*                *.*
udp00*.9999              *.*
udp00 *.161              *.*
udp00192.168.4.16.1039    192.168.4.16.1023
udp00192.168.4.16.1038    192.168.4.16.1023
udp 00 192.168.4.16.1037      192.168.4.16.1023
udp00192.168.4.16.1036    192.168.4.16.1023
udp00*.1022              *.*
udp00*.1023              *.*
udp00*.111               *.*
udp00*.*                 *.*
```

show system connections extensive

user@host> show system connections extensive

```

Active Internet connections (including servers)
Proto Recv-Q Send-Q Local Address           Foreign Address
      (state)
tcp4      0      6 192.168.187.15.23
172.27.133.138.3013      ESTABLISHED
    sndsbcc:          6 sndsbmbcnt:          256 sndsbmbmax:      272000
    sndsblowat:       2048 sndsbhiwat:          34000
    rcvsbcc:          0 rcvsbmbcnt:          0 rcvsbmbmax:      533120
    rcvsblowat:       1 rcvsbhiwat:          66640
    proc id:          0 proc name:
    iss: 2566994072    sndup: 2566994491
    snduna: 2566994491 sndnxt: 2566994494    sndwnd:          64094
    sndmax: 2566994494 sndcwnd:          6589 sndssthresh:      2720
    irs: 236981199    rcvup: 236981325
    rcvnxt: 236981327 rcvadv: 237046862    rcvwnd:          66640
    rtt: 140058623    srtt: 15519    rttv:          908
    rxtcur:          1200 rxtshift:          0    rtseq: 2566994491
    rttmin:          1000 mss:          1360
    flags: SACK_PERMIT [0x2000200]
tcp4      0      0 10.255.165.93.179
10.255.165.203.65141      ESTABLISHED
    sndsbcc:          0 sndsbmbcnt:          0 sndsbmbmax:      131072
    sndsblowat:       2048 sndsbhiwat:          16384
    rcvsbcc:          0 rcvsbmbcnt:          0 rcvsbmbmax:      131072
    rcvsblowat:       1 rcvsbhiwat:          16384
    proc id:          0 proc name:
    iss: 2555995917    sndup: 2555995917
    snduna: 2555995917 sndnxt: 2555995917    sndwnd:          16384
    sndmax: 2555995917 sndcwnd:          1000 sndssthresh: 1073725440
    irs: 2123825753    rcvup: 2123860681
    rcvnxt: 2123860681 rcvadv: 2123877065    rcvwnd:          16384
    rtt: 0    srtt: 3309    rttv:          72
    rxtcur:          1200 rxtshift:          0    rtseq: 2555995898
    rttmin:          1000 mss:          500
    flags: REQ_SCALE RCVD_SCALE REQ_TSTMP RCVD_TSTMP SACK_PERMIT [0x3e0]
tcp4      0      0 10.255.165.93.179
10.255.165.203.65141      ESTABLISHED
    sndsbcc:          0 sndsbmbcnt:          0 sndsbmbmax:      131072
    sndsblowat:       2048 sndsbhiwat:          16384
    rcvsbcc:          0 rcvsbmbcnt:          0 rcvsbmbmax:      131072
    rcvsblowat:       1 rcvsbhiwat:          16384
    proc id:          5022 proc name: rpd
    iss: 2123825753    sndup: 2123860662
    snduna: 2123860681 sndnxt: 2123860681    sndwnd:          16384
    sndmax: 2123860681 sndcwnd:          1000 sndssthresh: 1073725440
    irs: 2555995917    rcvup: 2555995917
    rcvnxt: 2555995917 rcvadv: 2556012301    rcvwnd:          16384
    rtt: 0    srtt: 3279    rttv:          22
    rxtcur:          1200 rxtshift:          0    rtseq: 2123860662
    rttmin:          1000 mss:          500
    flags: REQ_SCALE RCVD_SCALE REQ_TSTMP RCVD_TSTMP SACK_PERMIT [0x100003e0]
tcp4      0      0 10.255.165.113.52404
10.255.165.113.52404      ESTABLISHED
    sndsbcc:          0 sndsbmbcnt:          0 sndsbmbmax:      131072
    sndsblowat:       2048 sndsbhiwat:          16384
    rcvsbcc:          0 rcvsbmbcnt:          0 rcvsbmbmax:      131072
    rcvsblowat:       1 rcvsbhiwat:          16384

```



```

proc id:      0  proc name:
  iss: 1109297190  sndup: 1109332099
  snduna: 1109332118  sndnxt: 1109332118  sndwnd: 16384
  sndmax: 1109332118  sndcwnd: 1000  sndssthresh: 1073725440
  irs: 1476831634  rcvup: 1476866449
  rcvnxt: 1476866449  rcvadv: 1476882833  rcvwnd: 16384
  rtt: 0  srtt: 3235  rttv: 18
  rxtcur: 1200  rxtshift: 0  rtseq: 1109332099
  rttmin: 1000  mss: 500
  flags: REQ_SCALE RCVD_SCALE REQ_TSTMP RCVD_TSTMP SACK_PERMIT [0x3e0]

```

show system connections lcc (TX Matrix Router)

```
user@host> show system connections lcc 2
```

```
lcc2-re0:
```

```
-----
Active Internet connections (including servers)
```

Proto	Recv-Q	Send-Q	Local Address	Foreign Address	(state)
tcp4	0	0	192.168.66.131.1342	192.168.66.130.23	ESTABLISHED
tcp4	0	0	192.168.66.131.2059	192.168.66.130.23	ESTABLISHED
tcp4	0	0	192.168.66.131.4571	192.168.66.130.23	ESTABLISHED
tcp4	0	0	192.168.66.131.2496	192.168.66.130.23	ESTABLISHED
tcp4	0	0	*.3221	*.*	LISTEN
tcp4	0	0	*.23	*.*	LISTEN
tcp4	0	0	*.22	*.*	LISTEN
tcp4	0	0	*.514	*.*	LISTEN
tcp4	0	0	*.513	*.*	LISTEN
tcp4	0	0	*.21	*.*	LISTEN
tcp4	0	0	*.79	*.*	LISTEN
tcp4	0	0	*.6234	*.*	LISTEN
udp4	0	0	*.514	*.*	
udp4	0	0	*.6333	*.*	

show system connections show-routing-instances

```
user@host> show system connections show-routing-instances
```

```
Active Internet connections (including servers) (including routing-instances)
Proto Recv-Q Send-Q Local Address Foreign Address Routing Instance
(state)
```

Proto	Recv-Q	Send-Q	Local Address	Foreign Address	Routing Instance
tcp4	0	0	192.168.69.204.23	172.17.28.19.4267	default
			ESTABLISHED		
tcp4	0	0	192.168.69.204.58540	10.209.7.138.23	default
			ESTABLISHED		
tcp4	0	0	192.168.69.204.23	172.17.28.19.1098	default
			ESTABLISHED		
tcp4	0	0	192.168.7.1.57668	192.168.9.1.179	default
			ESTABLISHED		
tcp4	0	0	192.168.7.1.179	192.168.8.1.49209	default
			ESTABLISHED		
tcp4	0	0	128.0.0.1.6234	128.0.3.17.1024	
__juniper_private1__			ESTABLISHED		
tcp4	0	0	128.0.0.4.9000	128.0.0.4.59103	
__juniper_private1__			ESTABLISHED		
tcp4	0	0	128.0.0.4.59103	128.0.0.4.9000	
__juniper_private1__			ESTABLISHED		
tcp4	0	0	*.32012	*.*	
__juniper_private1__			LISTEN		
tcp4	0	0	*.9000	*.*	
__juniper_private1__			LISTEN		
tcp4	0	0	*.33007	*.*	

```

__juniper_private2__ LISTEN
tcp46      0      0 *.179      *.*      default
      LISTEN
tcp4       0      0 *.179      *.*      default
      LISTEN
tcp4       0      0 *.6154     *.*
__juniper_private1__ LISTEN
tcp4       0      0 *.6153     *.*
__juniper_private1__ LISTEN
tcp4       0      0 *.7000     *.*
__juniper_private1__ LISTEN
tcp4       0      0 *.6152     *.*
__juniper_private1__ LISTEN
tcp4       0      0 *.6156     *.*
__juniper_private1__ LISTEN
tcp4       0      0 *.33005    *.*
__juniper_private2__ LISTEN
tcp4       0      0 *.31343    *.*
__juniper_private1__ LISTEN
tcp4       0      0 *.31341    *.*
__juniper_private1__ LISTEN
tcp4       0      0 *.32003    *.*
__juniper_private2__ LISTEN
tcp4       0      0 *.666      *.*
__juniper_private1__ LISTEN
tcp4       0      0 *.38       *.*
__juniper_private1__ LISTEN
tcp4       0      0 *.3221     *.*      default
      LISTEN

```

show system connections (TX Matrix Plus Router)

```

user@host> show system connections
sfc0-re0:

```

```

-----
Active Internet connections (including servers)
Proto Recv-Q Send-Q Local Address           Foreign Address
      (state)
tcp4    0      3 192.168.178.11.23
172.17.28.19.3565        ESTABLISHED
tcp4    0      0 192.168.178.11.23
172.17.28.204.62719      ESTABLISHED
tcp4    0      0 192.168.178.11.23
192.168.69.199.51255    ESTABLISHED
tcp4    0      0 192.168.178.11.23
172.24.26.227.42860     ESTABLISHED
tcp4    0      0 *.6156      *.*
      LISTEN
tcp4    0      0 162.0.0.4.32012
      ESTABLISHED      162.0.0.5.58935
tcp4    0      0 *.32012     *.*
      LISTEN
tcp4    0      0 *.33007     *.*
      LISTEN
tcp4    0      0 *.666       *.*
      LISTEN
tcp4    0      0 162.0.0.4.6161
      ESTABLISHED      162.0.0.5.62026
tcp4    0      0 *.33005     *.*
      LISTEN
tcp4    0      0 162.0.0.4.9000
      162.0.0.4.51611

```

```

                                ESTABLISHED
tcp4      0      0 162.0.0.4.51611              162.0.0.4.9000
                                ESTABLISHED
tcp4      0      0 *.6151                        *.*
                                LISTEN
tcp4      0      0 *.6154                        *.*
                                LISTEN
tcp4      0      0 *.6153                        *.*
                                LISTEN
tcp4      0      0 *.31343                      *.*
                                LISTEN
tcp4      0      0 *.31341                      *.*
                                LISTEN
tcp4      0      0 *.9000                        *.*
                                LISTEN
tcp4      0      0 *.6152                        *.*
                                LISTEN
tcp4      0      0 *.32003                      *.*
                                LISTEN
tcp4      0      0 *.33009                      *.*
                                LISTEN
tcp4      0      0 *.3221                       *.*
                                LISTEN
tcp4      0      0 *.23                         *.*
                                LISTEN
tcp4      0      0 *.22                         *.*
                                LISTEN
tcp4      0      0 *.514                        *.*
                                LISTEN
tcp4      0      0 *.513                        *.*
                                LISTEN
tcp4      0      0 *.21                         *.*
                                LISTEN
tcp4      0      0 *.79                         *.*
                                LISTEN
tcp4      0      0 *.514                        *.*
                                LISTEN
tcp4      0      0 *.513                        *.*
                                LISTEN
tcp4      0      0 *.6234                      *.*
                                LISTEN
udp4      0      0 127.0.0.1.123               *.*
udp4      0      0 10.255.178.11.123           *.*
udp4      0      0 *.123                        *.*
udp46     0      0 *.514                        *.*
udp4      0      0 *.514                        *.*
udp46     0      0 *.62027                      *.*
udp4      0      0 *.59363                      *.*
udp4      0      0 *.31342                      *.*
udp46     0      0 *.161                        *.*
udp4      0      0 *.161                        *.*
udp4      0      0 *.31340                      *.*
udp4      0      0 *.31340                      *.*
udp46     0      0 *.49152                      *.*
udp46     0      0 *.4784                       *.*
udp46     0      0 *.3784                       *.*
udp4      0      0 *.49152                      *.*
udp4      0      0 *.4784                       *.*
udp4      0      0 *.3784                       *.*
udp4      0      0 *.6333                       *.*
ip4       0      0 *.*                          *.*

```

```

ip4          0      0  *.*                               *.*

lcc0-re0:
-----
Active Internet connections (including servers)
Proto Recv-Q Send-Q Local Address           Foreign Address
      (state)
tcp4      0      0 192.168.178.3.23
172.24.26.227.50399      ESTABLISHED
tcp4      0      0 *.*.6234
                                LISTEN
tcp4      0      0 *.*.7000
                                LISTEN
tcp4      0      0 *.*.9000
                                LISTEN
tcp4      0      0 *.*.33009
                                LISTEN
tcp4      0      0 *.*.3221
                                LISTEN
tcp4      0      0 *.*.23
                                LISTEN
tcp4      0      0 *.*.22
                                LISTEN
tcp4      0      0 *.*.514
                                LISTEN
tcp4      0      0 *.*.513
                                LISTEN
tcp4      0      0 *.*.21
                                LISTEN
tcp4      0      0 *.*.79
                                LISTEN
tcp4      0      0 *.*.514
                                LISTEN
tcp4      0      0 *.*.513
                                LISTEN
udp4      0      0 *.*.514
udp4      0      0 *.*.514
udp4      0      0 *.*.59924
udp4      0      0 *.*.59412
udp4      0      0 *.*.161
udp4      0      0 *.*.161
udp4      0      0 *.*.31342
udp4      0      0 *.*.6333

lcc1-re0:
-----
Active Internet connections (including servers)
Proto Recv-Q Send-Q Local Address           Foreign Address
      (state)
tcp4      0      0 *.*.6234
                                LISTEN
tcp4      0      0 *.*.7000
                                LISTEN
tcp4      0      0 *.*.9000
                                LISTEN
tcp4      0      0 *.*.3221
                                LISTEN
tcp4      0      0 *.*.23
                                LISTEN
tcp4      0      0 *.*.22
                                LISTEN

```

```

tcp4      0      0 *.514          LISTEN          *.*
tcp4      0      0 *.513          LISTEN          *.*
tcp4      0      0 *.21           LISTEN          *.*
tcp4      0      0 *.79           LISTEN          *.*
tcp4      0      0 *.514          LISTEN          *.*
tcp4      0      0 *.513          LISTEN          *.*
tcp4      0      0 *.33009        LISTEN          *.*
udp46     0      0 *.514          *.*
udp4      0      0 *.514          *.*
udp46     0      0 *.59924        *.*
udp4      0      0 *.59412        *.*
udp4      0      0 *.31342        *.*
udp46     0      0 *.161          *.*
udp4      0      0 *.161          *.*
udp4      0      0 *.6333         *.*

```

lcc2-re0:

```

-----
Active Internet connections (including servers)
Proto Recv-Q Send-Q Local Address           Foreign Address
      (state)
tcp4      0      0 *.6234          LISTEN          *.*
tcp4      0      0 *.7000          LISTEN          *.*
tcp4      0      0 *.9000          LISTEN          *.*
tcp4      0      0 *.33009         LISTEN          *.*
tcp4      0      0 *.3221          LISTEN          *.*
tcp4      0      0 *.23            LISTEN          *.*
tcp4      0      0 *.22            LISTEN          *.*
tcp4      0      0 *.514           LISTEN          *.*
...

```

show system connections sfc (TX Matrix Plus Router)

```

user@host> show system connections sfc 0
sfc0-re0:

```

```

-----
Active Internet connections (including servers)
Proto Recv-Q Send-Q Local Address           Foreign Address
      (state)
tcp4      0      0 162.0.0.4.514        TIME_WAIT      132.0.0.4.952
tcp4      0      0 162.0.0.4.514        TIME_WAIT      131.0.0.4.694
tcp4      0      0 162.0.0.4.514        TIME_WAIT      130.0.0.4.860
tcp4      0      0 162.0.0.4.514        TIME_WAIT      129.0.0.4.716

```

tcp4	0	0	162.0.0.4.996		132.0.0.4.514
			TIME_WAIT		
tcp4	0	0	162.0.0.4.798		131.0.0.4.514
			TIME_WAIT		
tcp4	0	0	162.0.0.4.995		130.0.0.4.514
			TIME_WAIT		
tcp4	0	0	162.0.0.4.895		129.0.0.4.514
			TIME_WAIT		
tcp4	0	0	192.168.178.11.21		
172.17.28.204.64662				TIME_WAIT	
tcp4	0	0	192.168.178.11.21		
172.17.28.204.51612				TIME_WAIT	
tcp4	0	0	*,6156		*,*
			LISTEN		
tcp4	0	0	*,9000		*,*
			LISTEN		
tcp4	0	0	*,666		*,*
			LISTEN		
tcp4	0	2	192.168.178.11.23		
172.17.28.19.3565				ESTABLISHED	
tcp4	0	0	192.168.178.11.23		
172.17.28.204.62719				ESTABLISHED	
tcp4	0	0	192.168.178.11.23		
192.168.69.199.51255				ESTABLISHED	
tcp4	0	0	192.168.178.11.23		
172.24.26.227.42860				ESTABLISHED	
tcp4	0	0	162.0.0.4.32012		162.0.0.5.58935
			ESTABLISHED		
tcp4	0	0	*,32012		*,*
			LISTEN		
tcp4	0	0	*,33007		*,*
			LISTEN		
tcp4	0	1432	162.0.0.4.6161		162.0.0.5.62026
			ESTABLISHED		
tcp4	0	0	*,33005		*,*
			LISTEN		
tcp4	0	0	162.0.0.4.9000		162.0.0.4.51611
			FIN_WAIT_2		
tcp4	0	0	162.0.0.4.51611		162.0.0.4.9000
			CLOSE_WAIT		
tcp4	0	0	*,6151		*,*
			LISTEN		
tcp4	0	0	*,6154		*,*
			LISTEN		
tcp4	0	0	*,6153		*,*
			LISTEN		
tcp4	0	0	*,31343		*,*
			LISTEN		
tcp4	0	0	*,31341		*,*
			LISTEN		
tcp4	0	0	*,6152		*,*
			LISTEN		
tcp4	0	0	*,32003		*,*
			LISTEN		
tcp4	0	0	*,33009		*,*
			LISTEN		
tcp4	0	0	*,3221		*,*
			LISTEN		
tcp4	0	0	*,23		*,*
			LISTEN		
tcp4	0	0	*,22		*,*

```

tcp4      0      0 *.514      LISTEN      *. *
tcp4      0      0 *.513      LISTEN      *. *
tcp4      0      0 *.21       LISTEN      *. *
tcp4      0      0 *.79       LISTEN      *. *
tcp4      0      0 *.514      LISTEN      *. *
tcp4      0      0 *.513      LISTEN      *. *
tcp4      0      0 *.6234     LISTEN      *. *
udp4      0      0 127.0.0.1.123 LISTEN      *. *
udp4      0      0 10.255.178.11.123 LISTEN      *. *
udp4      0      0 *.123      LISTEN      *. *
udp46     0      0 *.514      LISTEN      *. *
udp4      0      0 *.514      LISTEN      *. *
udp46     0      0 *.50895    LISTEN      *. *
udp4      0      0 *.50794    LISTEN      *. *
udp4      0      0 *.31342    LISTEN      *. *
udp46     0      0 *.161      LISTEN      *. *
udp4      0      0 *.161      LISTEN      *. *
udp4      0      0 *.31340    LISTEN      *. *
udp4      0      0 *.31340    LISTEN      *. *
udp46     0      0 *.49152    LISTEN      *. *
udp46     0      0 *.4784     LISTEN      *. *
udp46     0      0 *.3784     LISTEN      *. *
udp4      0      0 *.49152    LISTEN      *. *
udp4      0      0 *.4784     LISTEN      *. *
udp4      0      0 *.3784     LISTEN      *. *
udp4      0      0 *.6333     LISTEN      *. *
ip4       104    0 *. *      LISTEN      *. *
ip4       0      0 *. *      LISTEN      *. *
ip4       0      0 *. *      LISTEN      *. *

```

show system connections show-routing-instances (TX Matrix Plus Router)

```

user@host> show system connections show-routing-instances
sfc0-re0:
-----
Active Internet connections (including servers) (including routing-instances)
Proto Recv-Q Send-Q Local Address           Foreign Address
      Routing Instance      (state)
tcp4      0      0 *.6156                  __juniper_private1__    LISTEN      *. *
tcp4      0      0 *.9000                  __juniper_private1__    LISTEN      *. *
tcp4      0      0 *.666                   __juniper_private1__    LISTEN      *. *
tcp4      0      2 192.168.178.11.23       default                  ESTABLISHED
172.17.28.19.3565
tcp4      0      0 192.168.178.11.23       default                  ESTABLISHED
172.17.28.204.62719
tcp4      0      0 192.168.178.11.23       default                  ESTABLISHED
192.168.69.199.51255
tcp4      0      0 192.168.178.11.23       default                  ESTABLISHED
172.24.26.227.42860
tcp4      0      0 162.0.0.4.32012         162.0.0.5.58935

```

			__juniper_private1__	ESTABLISHED	
tcp4	0	0 *.32012			*.*
			__juniper_private1__	LISTEN	
tcp4	0	0 *.33007			*.*
			__juniper_private2__	LISTEN	
tcp4	0	0 162.0.0.4.6161			162.0.0.5.62026
			__juniper_private1__	ESTABLISHED	
tcp4	0	0 *.33005			*.*
			__juniper_private2__	LISTEN	
tcp4	0	0 162.0.0.4.9000			162.0.0.4.51611
			__juniper_private1__	FIN_WAIT_2	
tcp4	0	0 162.0.0.4.51611			162.0.0.4.9000
			__juniper_private1__	CLOSE_WAIT	
tcp4	0	0 *.6151			*.*
			__juniper_private1__	LISTEN	
tcp4	0	0 *.6154			*.*
			__juniper_private1__	LISTEN	
tcp4	0	0 *.6153			*.*
			__juniper_private1__	LISTEN	
tcp4	0	0 *.31343			*.*
			__juniper_private1__	LISTEN	
tcp4	0	0 *.31341			*.*
			__juniper_private1__	LISTEN	
tcp4	0	0 *.6152			*.*
			__juniper_private1__	LISTEN	
tcp4	0	0 *.32003			*.*
			__juniper_private2__	LISTEN	
tcp4	0	0 *.33009			*.*
			__juniper_private2__	LISTEN	
tcp4	0	0 *.3221			*.*
			default	LISTEN	
tcp4	0	0 *.23			*.*
			default	LISTEN	
tcp4	0	0 *.22			*.*
			default	LISTEN	
tcp4	0	0 *.514			*.*
			default	LISTEN	
tcp4	0	0 *.513			*.*
			default	LISTEN	
tcp4	0	0 *.21			*.*
			default	LISTEN	
tcp4	0	0 *.79			*.*
			default	LISTEN	
tcp4	0	0 *.514			*.*
			__juniper_private1__	LISTEN	
tcp4	0	0 *.513			*.*
			__juniper_private1__	LISTEN	
tcp4	0	0 *.6234			*.*
			__juniper_private1__	LISTEN	
udp4	0	0 127.0.0.1.123			*.*
			default		
udp4	0	0 10.255.178.11.123			*.*
			default		
udp4	0	0 *.123			*.*
			default		
udp46	0	0 *.514			*.*
			default		
udp4	0	0 *.514			*.*
			default		
udp46	0	0 *.50895			*.*
			default		


```

udp4      0      0 *.50794      *.*
          default
udp4      0      0 *.31342      *.*
          __juniper_private1__
udp46     0      0 *.161        *.*
          default
udp4      0      0 *.161        *.*
          default
udp4      0      0 *.31340      *.*
          __juniper_private2__
udp4      0      0 *.31340      *.*
          __juniper_private1__
udp46     0      0 *.49152      *.*
          default
udp46     0      0 *.4784       *.*
          default
udp46     0      0 *.3784       *.*
          default
udp4      0      0 *.49152      *.*
          default
udp4      0      0 *.4784       *.*
          default
udp4      0      0 *.3784       *.*
          default
udp4      0      0 *.6333       *.*
          __juniper_private1__
ip4       0      0 *.*          *.*
          default
ip4       0      0 *.*          *.*
          default
ip4       0      0 *.*          *.*
          default

```

lcc0-re0:

```

-----
Active Internet connections (including servers) (including routing-instances)
Proto Recv-Q Send-Q Local Address           Foreign Address
                                Routing Instance      (state)
tcp4      0      0 *.7000             *.*
                                __juniper_private1__ LISTEN
tcp4      0      0 192.168.178.3.23   *.*
172.24.26.227.50399 default              ESTABLISHED
tcp4      0      0 *.6234             *.*
                                __juniper_private1__ LISTEN
tcp4      0      0 *.9000             *.*
                                __juniper_private1__ LISTEN
tcp4      0      0 *.33009            *.*
                                __juniper_private2__ LISTEN
tcp4      0      0 *.3221             *.*
                                default              LISTEN
tcp4      0      0 *.23               *.*
                                default              LISTEN
tcp4      0      0 *.22               *.*
                                default              LISTEN
tcp4      0      0 *.514              *.*
                                default              LISTEN
tcp4      0      0 *.513              *.*
                                default              LISTEN
tcp4      0      0 *.21               *.*
                                default              LISTEN
tcp4      0      0 *.79               *.*

```

```

tcp4      0      0 *.514      default    LISTEN     *.*
tcp4      0      0 *.513      __juniper_private1__ LISTEN     *.*
udp46     0      0 *.514      __juniper_private1__ LISTEN     *.*
udp4      0      0 *.514      default    *.*
udp46     0      0 *.59924    default    *.*
udp4      0      0 *.59412    default    *.*
udp46     0      0 *.161      default    *.*
udp4      0      0 *.161      default    *.*
udp4      0      0 *.31342    default    *.*
udp4      0      0 *.6333     __juniper_private1__ *.*
udp4      0      0 *.6333     __juniper_private1__

```

lcc1-re0:

```

-----
Active Internet connections (including servers) (including routing-instances)
Proto Recv-Q Send-Q Local Address           Routing Instance      (state)      Foreign Address
tcp4      0      0 *.7000      default               LISTEN        *.*
tcp4      0      0 *.6234      __juniper_private1__ LISTEN        *.*
tcp4      0      0 *.9000      __juniper_private1__ LISTEN        *.*
tcp4      0      0 *.3221      __juniper_private1__ LISTEN        *.*
tcp4      0      0 *.23        default               LISTEN        *.*
tcp4      0      0 *.22        default               LISTEN        *.*
tcp4      0      0 *.514       default               LISTEN        *.*
tcp4      0      0 *.513       default               LISTEN        *.*
tcp4      0      0 *.21        default               LISTEN        *.*
tcp4      0      0 *.79        default               LISTEN        *.*
tcp4      0      0 *.514       default               LISTEN        *.*
tcp4      0      0 *.513       __juniper_private1__ LISTEN        *.*
tcp4      0      0 *.33009     __juniper_private1__ LISTEN        *.*
tcp4      0      0 *.33009     __juniper_private2__ LISTEN        *.*
udp46     0      0 *.514       default               *.*
udp4      0      0 *.514       default               *.*
udp46     0      0 *.59924    default               *.*
udp4      0      0 *.59412    default               *.*
udp4      0      0 *.59412    default

```

```

udp4      0      0 *.31342      *.*
          __juniper_private1__
udp46     0      0 *.161        *.*
          default
udp4      0      0 *.161        *.*
          default
udp4      0      0 *.6333       *.*
          __juniper_private1__

```

lcc2-re0:

```

-----
Active Internet connections (including servers) (including routing-instances)
Proto Recv-Q Send-Q Local Address           Foreign Address
Routing Instance        (state)
tcp4      0      0 *.7000        *.*
          __juniper_private1__ LISTEN
tcp4      0      0 *.6234        *.*
          __juniper_private1__ LISTEN
tcp4      0      0 *.9000        *.*
          __juniper_private1__ LISTEN
tcp4      0      0 *.33009       *.*
          __juniper_private2__ LISTEN
tcp4      0      0 *.3221        *.*
          default          LISTEN
tcp4      0      0 *.23          *.*
          default          LISTEN
tcp4      0      0 *.22          *.*
          default          LISTEN
tcp4      0      0 *.514         *.*
          default          LISTEN
tcp4      0      0 *.513         *.*
          default          LISTEN
tcp4      0      0 *.21          *.*
          default          LISTEN
tcp4      0      0 *.79          *.*
          default          LISTEN
tcp4      0      0 *.514         *.*
          __juniper_private1__ LISTEN
tcp4      0      0 *.513         *.*
          __juniper_private1__ LISTEN
udp46     0      0 *.514         *.*
          default
udp4      0      0 *.514         *.*
          default
udp4      0      0 *.31342       *.*
          __juniper_private1__
udp46     0      0 *.62103       *.*
          default
udp4      0      0 *.59924       *.*
          default
udp46     0      0 *.161         *.*
          default
udp4      0      0 *.161         *.*
          default
udp4      0      0 *.6333       *.*
          __juniper_private1__

```

lcc3-re0:

```

-----
Active Internet connections (including servers) (including routing-instances)
Proto Recv-Q Send-Q Local Address           Foreign Address

```

				Routing Instance	(state)	
tcp4	0	0	*.7000			*.*
tcp4	0	0	*.6234	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.9000	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.33009	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.3221	__juniper_private2__	LISTEN	*.*
tcp4	0	0	*.23	default	LISTEN	*.*
tcp4	0	0	*.22	default	LISTEN	*.*
tcp4	0	0	*.514	default	LISTEN	*.*
tcp4	0	0	*.513	default	LISTEN	*.*
tcp4	0	0	*.21	default	LISTEN	*.*
tcp4	0	0	*.79	default	LISTEN	*.*
tcp4	0	0	*.514	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.513	__juniper_private1__	LISTEN	*.*
udp46	0	0	*.514			*.*
udp4	0	0	*.514	default		*.*
udp46	0	0	*.62103	default		*.*
udp4	0	0	*.59924	default		*.*
udp4	0	0	*.31342	__juniper_private1__		*.*
udp46	0	0	*.161	default		*.*
udp4	0	0	*.161	default		*.*
udp4	0	0	*.6333	__juniper_private1__		*.*

show system connections (QFX3500 Switch)

```

user@switch> show system connections
Active Internet connections (including servers)
Proto Recv-Q Send-Q Local Address           Foreign Address
          (state)
tcp4      0      0 10.94.204.110.23        172.17.28.19.1308      ESTABLISHED
tcp4      0      0 128.0.0.1.6234          128.0.0.1.65142        ESTABLISHED
tcp4      0      0 128.0.0.1.65142          128.0.0.1.6234        ESTABLISHED
tcp4      0      0 128.0.0.1.33003          128.0.0.1.61441        ESTABLISHED
tcp4      0      0 128.0.0.1.61441          128.0.0.1.33003        ESTABLISHED
tcp46     0      0 *.179                    *.*
```

			LISTEN	
tcp4	0	0 *.179		*.*
			LISTEN	
tcp4	0	0 128.0.0.16.9000		128.0.0.16.50970
			ESTABLISHED	
tcp4	0	0 128.0.0.16.50970		128.0.0.16.9000
			ESTABLISHED	
tcp4	0	0 *.38		*.*
			LISTEN	
tcp4	0	0 *.3491		*.*
			LISTEN	
tcp4	0	0 *.6156		*.*
			LISTEN	
tcp4	0	0 128.0.0.1.33001		128.0.0.1.59437
			ESTABLISHED	
tcp4	0	0 128.0.0.1.59437		128.0.0.1.33001
			ESTABLISHED	
tcp4	0	0 128.0.0.1.33023		128.0.0.1.63605
			ESTABLISHED	
tcp4	0	0 128.0.0.1.63605		128.0.0.1.33023
			ESTABLISHED	
tcp4	0	0 128.0.0.1.33001		128.0.0.1.63830
			ESTABLISHED	
tcp4	0	0 128.0.0.1.63830		128.0.0.1.33001
			ESTABLISHED	
tcp4	0	0 *.667		*.*
			LISTEN	
tcp4	0	0 *.6156		*.*
			LISTEN	
tcp4	0	0 128.0.0.1.7000		128.0.0.1.51580
			ESTABLISHED	
tcp4	0	0 128.0.0.1.51580		128.0.0.1.7000
			ESTABLISHED	
tcp4	0	0 128.0.0.1.6234		128.0.0.1.53646
			ESTABLISHED	
tcp4	0	0 *.33001		*.*
			LISTEN	
tcp4	0	0 *.33003		*.*
			LISTEN	
tcp4	0	0 128.0.0.1.53646		128.0.0.1.6234
			ESTABLISHED	
tcp4	0	0 128.0.0.16.9000		128.0.0.16.63454
			ESTABLISHED	
tcp4	0	0 128.0.0.16.63454		128.0.0.16.9000
			ESTABLISHED	
tcp4	0	0 *.666		*.*
			LISTEN	
tcp4	0	0 *.7000		*.*
			LISTEN	
tcp4	0	0 *.51627		*.*
			LISTEN	
tcp4	0	0 *.3492		*.*
			LISTEN	
tcp4	0	0 *.33023		*.*
			LISTEN	
tcp4	0	0 *.33013		*.*
			LISTEN	
tcp4	0	0 *.7202		*.*
			LISTEN	
tcp4	0	0 *.6151		*.*
			LISTEN	

tcp4	0	0	*.9000		*.*
				LISTEN	
tcp4	0	0	*.6161		*.*
				LISTEN	
tcp4	0	0	*.6011		*.*
				LISTEN	
tcp4	0	0	*.3221		*.*
				LISTEN	
tcp4	0	0	*.23		*.*
				LISTEN	
tcp4	0	0	*.22		*.*
				LISTEN	
tcp4	0	0	*.514		*.*
				LISTEN	
tcp4	0	0	*.513		*.*
				LISTEN	
tcp4	0	0	*.21		*.*
				LISTEN	
tcp4	0	0	*.79		*.*
				LISTEN	
tcp4	0	0	*.514		*.*
				LISTEN	
tcp4	0	0	*.513		*.*
				LISTEN	
tcp4	0	0	*.1127		*.*
				LISTEN	
tcp4	0	0	*.1129		*.*
				LISTEN	
tcp4	0	0	*.1128		*.*
				LISTEN	
tcp4	0	0	*.6234		*.*
				LISTEN	
udp46	0	0	*.514		*.*
udp4	0	0	*.514		*.*
udp4	0	0	128.0.0.1.123		*.*
udp46	0	0	*.53344		*.*
udp4	0	0	*.54261		*.*
udp46	0	0	*.161		*.*
udp4	0	0	*.161		*.*
udp4	0	0	*.31342		*.*
udp4	0	0	*.59137		*.*
udp4	0	0	*.*		*.*
udp46	0	0	*.49152		*.*
udp46	0	0	*.4784		*.*
udp46	0	0	*.3784		*.*
udp4	0	0	*.49152		*.*
udp4	0	0	*.4784		*.*
udp4	0	0	*.3784		*.*
udp4	0	0	10.255.204.110.123		*.*
udp4	0	0	*.123		*.*
udp4	0	0	*.67		*.*
udp4	0	0	*.6333		*.*
udp4	0	0	*.2293		*.*
ip4	0	0	*.*		*.*
ip4	0	0	*.*		*.*
ip4	0	0	*.*		*.*

show system core-dumps

List of Syntax	Syntax on page 1079 Syntax (EX Series Switches) on page 1079 Syntax (TX Matrix Router) on page 1079 Syntax (TX Matrix Plus Router) on page 1079 Syntax (QFX Series and OCX Series) on page 1079
Syntax	<pre>show system core-dumps <brief detail> <core-filename> <core-file-info> <re0> <re1> <routing-engine> <satellite [<i>fpc-slot-id</i> device-alias <i>alias-name</i>]></pre>
Syntax (EX Series Switches)	<pre>show system core-dumps <all-members> <brief detail> <core-filename> <core-file-info> <local> <member <i>member-id</i>></pre>
Syntax (TX Matrix Router)	<pre>show system core-dumps <all-chassis all-lcc lcc <i>number</i> scc> <brief detail> <core-filename> <core-file-info></pre>
Syntax (TX Matrix Plus Router)	<pre>show system core-dumps <all-chassis all-lcc lcc <i>number</i> sfc <i>number</i>> <brief detail> <core-filename> <core-file-info></pre>
Syntax (QFX Series and OCX Series)	<pre>show system core-dumps <brief detail> <component (<i>UUID</i> <i>serial number</i> all)> <core-file-info component (<i>UUID</i> <i>serial number</i>) <i>core-file-name</i>> <display-period (<i>hours</i> <i>minutes</i> <i>seconds</i>)> <display-order> <kernel-crashinfo component (<i>UUID</i> <i>serial number</i>)> <repository (core log)></pre>
Release Information	<p>Command introduced before Junos OS Release 8.5.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>re0, re1, and routing-engine options introduced for dual Routing Engines in Junos OS Release 13.1.</p>

Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
satellite option introduced in Junos OS Release 14.2R3.

Description Show core files on all routers or switches running Junos OS. You can use the **show system core-dumps** command to show a list of system core files created when the router or switch has failed. This command can be useful for diagnostic purposes. Each list item includes the file permissions, number of links, owner, group, size, modification date, and path and filename. If dual Routing Engines are present, you can view core-dump files for either routing engine or both routing engines together. On a QFabric system, you can view core-dump files on individual QFabric system devices as well as on the entire QFabric system.

You can use the option **core-filename** and its options **core-file-info**, **brief**, and **detail** to display more information about the specified core-dump files.

Options **none**—Display a list of all existing core-dump files.



NOTE: If dual Routing Engines are present, then only the core-dump files for the active Routing Engine are listed.

all-chassis—(TX Matrix and TX Matrix Plus routers only) (Optional) On a routing matrix based on a TX Matrix router, display system core files for the TX Matrix router switch-card chassis [SCC] and all the T640 routers [LCCs] connected to the TX Matrix router.

On a routing matrix based on a TX Matrix Plus router, display system core files for the TX Matrix Plus router (switch-fabric chassis [SFC]) and all the T1600 routers [LCCs] connected to the TX Matrix Plus router.

<all-lcc | lcc number>—(TX Matrix and TX Matrix Plus routers only) (Optional) On a routing matrix based on the TX Matrix router, display core dump files for all T640 routers (line-card chassis [LCCs]) or a specific T640 router [LCC] connected to the TX Matrix router.

On a routing matrix based on the TX Matrix Plus router, display logging information for all T1600 routers (line-card chassis [LCCs]) or a specific T1600 router (LCC) connected to the TX Matrix Plus router. When using the **lcc number** option, replace **number** with a value from 0 through 3.



NOTE: The **all-chassis** option displays system core files for the SCC or SFC and the LCCs connected to the SCC or SFC in the routing matrix while the **all-lcc** option only displays system core files for the LCCs in the routing matrix.

all-members—(EX4200 switches) (Optional) Display system core files on all members of the Virtual Chassis configuration.

brief—(Optional) View details of a binary file.

component (*UUID* | *serial number* | *all*)—(QFabric systems only) (Optional) Display a list of core-dump files located on individual QFabric system device or on the entire QFabric system.

core-file-info—(Optional) Display the stack trace of a core file.

core-filename—(Optional) Name of a specific core file to display.

detail—(Optional) View stack trace with details of the binary file.

display-order (*timestamp-sort* | *alphanumeric-sort*)—(QFabric systems only) (Optional) Display list of debug artifacts generated within the specified period—for example, within the last hour, within the last 20 minutes, or within the last 32 seconds—according to their filename.

display-period (*hours* | *minutes* | *seconds*)—(QFabric systems only) (Optional) Display core-dump files generated within the specified period—for example, within the last hour, within the last 20 minutes, or within the last 32 seconds.

kernel-crashinfo component (*UUID* | *serial number*)—(QFabric systems only) (Optional) Display kernel crash information from the EEPROM on a QFabric system device.

local—(EX4200 switches only) (Optional) Display system core files on the local Virtual Chassis member.

member *member-id*—(EX4200 switches only) (Optional) Display system core files on the specified member of the Virtual Chassis configuration. Replace *member-id* with a value from 0 through 9.

re0—(Dual Routing Engines only) Display the core-dump files on re0.

re1—(Dual Routing Engines only) Display the coredump files on re1.

repository (*core* | *log*)—(QFabric systems only) (Optional) Specify either the core or log repository in which to view core-dump files.

routing-engine (*backup* | *both* | *local* | *master* | *other*)—(Dual routing engines only) Display a list of core-dump files for either the backup, local, master, or other routing engine or both routing engines.

satellite [*fpc-slot-id* | *device-alias* *alias-name*]—(Junos Fusion only) (Optional) Display hardware information for the specified satellite device in a Junos Fusion, or for all satellite devices in the Junos Fusion if no satellite devices are specified.

scc—(TX Matrix routers only) (Optional) Display system core files on the TX Matrix router (or switch-card chassis).

sfc—(TX Matrix Plus routers only) (Optional) Display system core files on the TX Matrix Plus router (or switch-fabric chassis).

Required Privilege Level view

List of Sample Output

- [show system core-dumps on page 1084](#)
- [show system core-dumps on page 1084](#)
- [show system core-dumps routing-engine both on page 1084](#)
- [show system core-dumps \(TX Matrix Plus Router\) on page 1084](#)
- [show system core-dumps \(QFX3500 Switch\) on page 1086](#)
- [show system core-dumps \(QFabric Systems\) on page 1086](#)
- [show system core-dumps core-file-info component serial number core-file-name \(QFabric Systems\) on page 1087](#)
- [show system core-dumps component serial number display-order alphanumeric-sort repository core \(QFabric Systems\) on page 1087](#)
- [show system core-dumps display-period \(QFabric Systems\) on page 1088](#)
- [show system core-dumps kernel-crashinfo component serial number \(QFabric Systems\) on page 1090](#)
- [show system core-dumps repository core \(QFabric Systems\) on page 1091](#)
- [show system core-dumps repository log \(QFabric Systems\) on page 1091](#)

Output Fields Table 74 describes the output fields for the **show system core-dumps** command. Output fields are listed in the approximate order in which they appear.

Table 74: show system core-dumps Output Fields

Field Name	Field Description
<i>Permissions</i>	Read/write permissions for the file named.
<i>Links</i>	Number of links to the file.
<i>Owner</i>	Name of the file owner.
<i>Group</i>	Name of the group with file access.
<i>File size</i>	File size in bytes.
<i>Modified</i>	Last file modification date and time.
<i>Path/filename</i>	File path where the file resides and the filename. (MX Series routers only) When you display the core files for an MX Series Virtual Chassis, the show system core-dumps command does not display information about files pertaining to the relayd process.
Repository scope:	Repository where core-dump files and log files are stored. The core-dump files are located in the core repository, and the log files are located in the log repository. The default Repository scope is shared since both the core and log repositories are shared by all of the QFabric system devices.
Repository head:	Path to the top-level repository location.

Table 74: show system core-dumps Output Fields (*continued*)

Field Name	Field Description
Repository name:	Name of the repository: core or log .
List of nodes for core repository:	List of core-dump files associated with a particular QFabric system device located in the core repository.
Node Group	Name of the QFabric system device.
Node Identifier	UUID or serial number of the QFabric system device.
Num	Number of core-dump and log files.
Model	Model number of the QFabric system device.
Usage	Usage of the repository in megabytes.
Total usage of core repository:	Total usage of core-dump files associated with a particular QFabric system device located in the core repository. Usage is specified in megabytes and as a percentage.
Total usage of log repository:	Total usage of log files associated with a particular QFabric system device located in the log repository. Usage is specified in megabytes and as a percentage.
List of nodes for core repository:	List of core-dump files associated with a particular QFabric system device located in the core repository.
List of nodes for log repository:	List of log files associated with a particular QFabric system device located in the log repository.
Filename	Name of the core-dump file.
Date	Last core-dump file modification date and time.
Size	Size of the core-dump file.
Core filename	Filename of the core-dump file.
Process name	Name of the process that is generating a core-dump file or log file.
Release	Junos OS release.
Build server	Junos OS build server.
Build date	Junos OS build date.
Stack trace	Stack trace of the core-dump file.

Sample Output

show system core-dumps

This example shows the command output if core files exist.

```
user@switch> show system core-dumps
-rw----- 1 root wheel 268369920 Jun 18 17:59 /var/crash/vmcore.0
-rw-rw---- 1 root field 3371008 Jun 18 17:53 /var/tmp/rpd.core.0
-rw-r--r-- 1 root wheel 27775914 Jun 18 17:59 /var/crash/kernel.0
```

show system core-dumps

This example shows the command output if core files do not exist.

```
user@host> show system core-dumps
/var/crash/*core*: No such file or directory
/var/tmp/*core*: No such file or directory
/var/crash/kernel.*: No such file or directory
```

show system core-dumps routing-engine both

This example shows the command output if dual Routing Engines are present.

```
user@host> show system core-dumps routing-engine both
re0:
-----
/var/crash/*core*: No such file or directory
/var/tmp/pics/*core*: No such file or directory
/var/crash/kernel.*: No such file or directory

/var/tmp/cores:
total blocks: 496776
-rw-rw---- 1 root field 11910589 Nov 8 13:20 chassisd.core.0.201311081320
...

-rw-rw---- 1 root field 11737227 Oct 28 14:21
rpd.core-tarball.4.tgz.201310281421.3458162
total files: 10

re1:
-----
/var/crash/*core*: No such file or directory
/var/tmp/pics/*core*: No such file or directory
/var/crash/kernel.*: No such file or directory

/var/tmp/cores:
total blocks: 3178420
-rw-rw---- 1 root field 19039721 Nov 8 14:29
chassisd.core.0.201311081429.3485600.gz
-rw-rw---- 1 root field 19039793 Nov 8 14:37
chassisd.core.1.201311081437.3485599.gz
..

-rw-rw---- 1 root field 11710113 Oct 17 15:26
rpd.core-tarball.1.1.tgz.201310171526.3430028
```

show system core-dumps (TX Matrix Plus Router)

```
user@host> show system core-dumps
```

sfc0-re0:

```
-----
/var/crash/kernel.*: No such file or directory
/tftpboot/corefiles/*core*: No such file or directory
```

```
/var/crash/cores:
total 8
```

```
/var/tmp/cores:
total 1627592
-rw-r--r-- 1 root field 535346090 May 15 07:36
rpd.core-tarball.0.090515.0736.tgz
-rw-r--r-- 1 root field 105632057 May 15 07:37
rpd.core-tarball.1.090515.0737.tgz
-rw-r--r-- 1 root field 101981681 May 15 07:38
rpd.core-tarball.2.090515.0738.tgz
-rw-r--r-- 1 root field 85854573 May 15 07:40
rpd.core-tarball.3.090515.0740.tgz
-rw-r--r-- 1 root field 4157845 May 15 08:18
rpd.core-tarball.4.090515.0818.tgz
```

lcc0-re0:

```
-----
/var/crash/kernel.*: No such file or directory
/tftpboot/corefiles/*core*: No such file or directory
```

```
/var/crash/cores:
total 8
```

```
/var/tmp/cores:
total 12
```

lcc1-re0:

```
-----
/var/crash/kernel.*: No such file or directory
/tftpboot/corefiles/*core*: No such file or directory
```

```
/var/crash/cores:
total 8
```

```
/var/tmp/cores:
total 10024
-rw-r--r-- 1 root field 1875794 Apr 22 15:47
chassisd.core-tarball.0.090422.1547.tgz
-rw-r--r-- 1 root field 1894183 Apr 22 19:02
chassisd.core-tarball.0.090422.1902.tgz
-rw-r--r-- 1 root field 1290240 Apr 26 16:01 ksyncd_1558.core.0.090426.1601
```

lcc2-re0:

```
-----
/var/crash/kernel.*: No such file or directory
/tftpboot/corefiles/*core*: No such file or directory
```

```
/var/crash/cores:
total 21124008
-rw-r--r-- 1 root wheel 1022376528 May 2 06:43
core-LCC2-EGFPC7.core.0.090502.0643
-rw-r--r-- 1 root wheel 1022376528 May 2 08:13
core-LCC2-EGFPC7.core.0.090502.0813
-rw-r--r-- 1 root wheel 1022376544 May 5 06:15
core-LCC2-EGFPC7.core.0.090505.0615
```

```

-rw-r--r-- 1 root wheel 1022376544 May 6 10:59
core-LCC2-EGFPC7.core.0.090506.1059
-rw-r--r-- 1 root wheel 1022376528 May 2 06:58
core-LCC2-EGFPC7.core.1.090502.0658
-rw-r--r-- 1 root wheel 754271232 May 5 06:33
core-LCC2-EGFPC7.core.1.090505.0633
-rw-r--r-- 1 root wheel 264897536 May 6 11:12
core-LCC2-EGFPC7.core.1.090506.1112
-rw-r--r-- 1 root wheel 1022376528 May 2 07:22
core-LCC2-EGFPC7.core.2.090502.0722
-rw-r--r-- 1 root wheel 163633152 May 5 06:52
core-LCC2-EGFPC7.core.2.090505.0652
-rw-r--r-- 1 root wheel 171312128 May 6 12:13
core-LCC2-EGFPC7.core.2.090506.1213
-rw-r--r-- 1 root wheel 1022376528 May 2 07:39
core-LCC2-EGFPC7.core.3.090502.0739
-rw-r--r-- 1 root wheel 1022376528 May 2 07:55
core-LCC2-EGFPC7.core.4.090502.0755
-rw-r--r-- 1 root wheel 427277312 May 7 04:47
core-LCC2-STFPC4.core.0.090507.0447
-rw-r--r-- 1 root wheel 419609600 May 7 04:47
core-LCC2-STFPC5.core.0.090507.0447
-rw-r--r-- 1 root wheel 432356352 May 7 04:47
core-LCC2-STFPC6.core.0.090507.0447

/var/tmp/cores:
total 2568
-rw-r--r-- 1 root field 1290240 May 14 14:26 ksyncd_1540.core.0.090514.1426
...
```

show system core-dumps (QFX3500 Switch)

```

user@switch> show system core-dumps
/var/crash/*core*: No such file or directory
-rw-rw---- 1 root field 1545143 Jun 4 2012 /var/tmp/pafxpc.core.0.gz
-rw-rw---- 1 root field 1545146 Jun 4 2012 /var/tmp/pafxpc.core.1.gz
-rw-rw---- 1 root field 1545141 Jun 4 2012 /var/tmp/pafxpc.core.2.gz
-rw-rw---- 1 root field 1545146 Jun 4 2012 /var/tmp/pafxpc.core.3.gz
-rw-rw---- 1 root field 1545142 Jun 5 2012 /var/tmp/pafxpc.core.4.gz
/var/tmp/pics/*core*: No such file or directory
/var/crash/kernel.*: No such file or directory
/tftpboot/corefiles/*core*: No such file or directory
total 5
```

show system core-dumps (QFabric Systems)

```

user@switch> show system core-dumps
Repository scope: shared
Repository head: /pbdata/export
List of nodes for core repository: /pbdata/export/rdumps/
```

Node Group	Node Identifier	Num	Model	Usage
DG-0	BCF7208D-E44F-E011-802F-4171BAAC781D	0	qfx3100	OM
FM-0	73747cd8-0710-11e1-b6a4-00e081c5297e	0	fx-jvre	OM
DRE-0	77116f18-0710-11e1-a2a0-00e081c5297e	0	fx-jvre	OM
NW-NG-0	BBAK0394	0	qfx3500	OM
NW-NG-0	cd78871a-0710-11e1-878e-00e081c5297e	0	fx-jvre	OM
NW-NG-0	d0afda1e-0710-11e1-a1d0-00e081c5297e	0	fx-jvre	OM
FC-0	d31ab7a6-0710-11e1-ad1b-00e081c5297e	0	fx-jvre	OM
FC-1	d4d0f254-0710-11e1-90c3-00e081c5297e	0	fx-jvre	OM
IC-WS001	WS001	0	-	-

```

IC-WS001      WS001/YW3803      0    qfxc08-3008    OM
IC-WS001      WS001/YN5999      0    qfxc08-3008    OM
node-device1  BBAK0372          0    qfx3500        OM
node-device1  EE3093            0    qfx3500        OM

```

Total usage of core repository: 0M of 70000M (0.0%)

List of nodes for log repository: /pbdata/export/rlogs/

Node Group	Node Identifier	Num	Model	Usage
DG-0	BCF7208D-E44F-E011-802F-4171BAAC781D	0	qfx3100	OM
FM-0	73747cd8-0710-11e1-b6a4-00e081c5297e	1	fx-jvre	OM
DRE-0	77116f18-0710-11e1-a2a0-00e081c5297e	1	fx-jvre	OM
NW-NG-0	BBAK0394	1	qfx3500	OM
NW-NG-0	cd78871a-0710-11e1-878e-00e081c5297e	1	fx-jvre	OM
NW-NG-0	d0afdale-0710-11e1-a1d0-00e081c5297e	3	fx-jvre	OM
FC-0	d31ab7a6-0710-11e1-ad1b-00e081c5297e	1	fx-jvre	OM
FC-1	d4d0f254-0710-11e1-90c3-00e081c5297e	1	fx-jvre	OM
IC-WS001	WS001	0	-	-
IC-WS001	WS001/YN5999	1	qfxc08-3008	OM
IC-WS001	WS001/YW3803	1	qfxc08-3008	OM
node-device1	BBAK0372	1	qfx3500	OM
node-device1	EE3093	1	qfx3500	OM

Total usage of log repository: 0M of 70000M (0.0%)

show system core-dumps core-file-info component serial number core-file-name (QFabric Systems)

```

user@switch> show system core-dumps core-file-info component
e8ff4b3e-7d92-11e0-be5d-00e081c1fe0e cosd.core.0.1519.05162011131846.gz
Repository scope: shared
Repository head: /pbstorage
Repository name: core
Core filename: /pbstorage/rdumps/e8ff4b3e-7d92-11e0-be5d-
00e081c1fe0e/5658.cosd.core.0.1519.05162011131846
Process name: cosd
Release: 11.3I0
Build server: /c/ssengupta/dfx_ha_v1/obj-i386-dcp/dcp/usr.sbin/cosd
Build date: 2011-05-14 01:11:44 UTC
Stack trace:
#0 0x8885d183 in select () from /usr/lib/libc.so.6
#0 0x8885d183 in select () from /usr/lib/libc.so.6
#1 0x887d4a45 in pselect () from /usr/lib/libc.so.6
#2 0x88774719 in pselect () from /usr/lib/libthr.so.2
#3 0x885de5db in __evGetNext () from /usr/lib/libisc.so.2
#4 0x885debf0 in __evMainLoop () from /usr/lib/libisc.so.2
#5 0x081125b2 in cosd_loop ()
#6 0x0812e19a in main ()

```

show system core-dumps component serial number display-order alphanumeric-sort repository core (QFabric Systems)

```

user@switch> show system core-dumps component BBAK8891 display-order alphanumeric-sort
repository core
Repository scope: shared
Repository head: /pbdata/export
Repository name: core
List of core dumps for component BBAK8891
Repository location: /pbdata/export/rdumps/BBAK8891

```

Filename	Date	Size
eswd.core.0.1361.11172011214257.gz	Nov 17 21:43:10 2011	4779553
eswd.core.1.80267.11172011214514.gz	Nov 17 21:45:19 2011	3541648

```

eswd.core.2.80682.11172011214535.gz      Nov 17 21:45:43 2011      2156683
vccpd.core.0.1195.11182011151131.gz      Nov 18 15:11:35 2011      375617
Number of core dumps in repository:4

```

show system core-dumps display-period (QFabric Systems)

```

user@switch> show system core-dumps display-period 24h
  show system core-dumps display-period 24h
Repository scope: shared
Repository head: /pbdata/export
List of core dumps at repository: /pbdata/export/rdumps
Delta timespec: Last 24h
Component: BBAK8273
Filename                                     Size                                     Date

vccpd.core.0.1195.11182011151131.gz      Nov 18 15:11:35 2011      375794
Component: cedb7b0e-0025-11e1-9a5f-00e081c52990
Filename                                     Size                                     Date

vccpd.core.0.1461.11182011151131.gz      Nov 18 15:11:31 2011      120951
Component: ee19c4f8-0025-11e1-aef6-00e081c52990
Filename                                     Size                                     Date

vccpd.core.0.1462.11182011151131.gz      Nov 18 15:11:31 2011      109420
Component: BBAK8281
Filename                                     Size                                     Date

vccpd.core.0.1196.11182011151131.gz      Nov 18 15:11:36 2011      375373
Component: BBAK8891
Filename                                     Size                                     Date

vccpd.core.0.1195.11182011151131.gz      Nov 18 15:11:35 2011      375617
Component: BBAK8276
Filename                                     Size                                     Date

vccpd.core.0.1196.11182011151131.gz      Nov 18 15:11:35 2011      375350
Component: BBAK8868
Filename                                     Size                                     Date

vccpd.core.0.1196.11182011151130.gz      Nov 18 15:11:34 2011      376211
Component: BBAK8835
Filename                                     Size                                     Date

vccpd.core.0.1195.11182011151130.gz      Nov 18 15:11:35 2011      375700
Component: BBAK8283
Filename                                     Size                                     Date

vccpd.core.0.1195.11182011151131.gz      Nov 18 15:11:36 2011      368298
Component: YW3781/YW3781
Filename                                     Size                                     Date

vccpd.core.0.1220.11182011151131.gz      Nov 18 15:11:38 2011      380002
Component: 09726be2-0026-11e1-82d9-00e081c52990
Filename                                     Size                                     Date

vccpd.core.0.1461.11182011151130.gz      Nov 18 15:11:31 2011      119965
Component: BBAK8309
Filename                                     Size                                     Date

vccpd.core.0.1196.11182011151131.gz      Nov 18 15:11:36 2011      378930
Component: 303d476a-0026-11e1-abf4-00e081c52990

```


Filename	Size	Date
vccpd.core.0.1460.11182011151131.gz Component: YW3798/YW3798	Nov 18 15:11:31 2011	118385
Filename	Size	Date
vccpd.core.0.1219.11182011151131.gz List of log dumps at repository: /pbdata/export/rlogs Delta timespec: Last 24h Component: BBAK8273	Nov 18 15:11:36 2011	380455
Filename	Size	Date
vccpd.tarball.0.1195.11182011151138.tgz Component: cedb7b0e-0025-11e1-9a5f-00e081c52990	Nov 18 15:11:39 2011	20415
Filename	Size	Date
vccpd.tarball.0.1461.11182011151131.tgz Component: ee19c4f8-0025-11e1-aef6-00e081c52990	Nov 18 15:11:33 2011	19651
Filename	Size	Date
vccpd.tarball.0.1462.11182011151133.tgz Component: BBAK8281	Nov 18 15:11:36 2011	24650
Filename	Size	Date
vccpd.tarball.0.1196.11182011151137.tgz Component: BBAK8891	Nov 18 15:11:41 2011	19445
Filename	Size	Date
vccpd.tarball.0.1195.11182011151138.tgz Component: BBAK8276	Nov 18 15:11:41 2011	21916
Filename	Size	Date
vccpd.tarball.0.1196.11182011151137.tgz Component: BBAK8868	Nov 18 15:11:39 2011	20461
Filename	Size	Date
vccpd.tarball.0.1196.11182011151137.tgz Component: BBAK8835	Nov 18 15:11:41 2011	21924
Filename	Size	Date
vccpd.tarball.0.1195.11182011151137.tgz Component: BBAK8283	Nov 18 15:11:39 2011	19424
Filename	Size	Date
vccpd.tarball.0.1195.11182011151138.tgz Component: YW3781/YW3781	Nov 18 15:11:42 2011	31186
Filename	Size	Date
vccpd.tarball.0.1220.11182011151141.tgz Component: 09726be2-0026-11e1-82d9-00e081c52990	Nov 18 15:11:45 2011	27565
Filename	Size	Date
vccpd.tarball.0.1461.11182011151130.tgz Component: BBAK8309	Nov 18 15:11:34 2011	19613
Filename	Size	Date
vccpd.tarball.0.1196.11182011151138.tgz Component: 303d476a-0026-11e1-abf4-00e081c52990	Nov 18 15:11:46 2011	50362
Filename	Size	Date
vccpd.tarball.0.1460.11182011151133.tgz	Nov 18 15:11:33 2011	19360

Component: YW3798/YW3798		
Filename	Size	Date
vccpd.tarball.0.1219.11182011151140.tgz	Nov 18 15:11:49 2011	24473

show system core-dumps kernel-crashinfo component serial number (QFabric Systems)

```
user@switch> show system core-dumps kernel-crashinfo component A0001/YA0197
Node: A0001/YA0197
```

Information about previous kernel crash:

-- Kernel panic data --

Panic string: kdb_sysctl_panic
 System uptime: 3 day 20 hr 59 min 40 sec Kernel crash time: 2011-11-15 Wed 15:25:17
 Kernel build linkstamp: JUNOS 11.3I #0: 2011-11-10 20:42:27 UTC

-- Stacktrace of panicing context --

Processor 1 (crash monarch):
 savectx+0x0 (c9552800,80214efc,802a7fbc,c88ad05c) ra 801b93a8 sz 0
 kdm_kcore_save_crashinfo+0x254 (c9552800,0,802a7fbc,c88ad05c) ra 801b9f44 sz 784
 kdm_kcore_kern_panic_event_handler+0x4b0 (c9552800,0,802a7fbc,c88ad05c) ra 8022a9b8 sz 88
 panic+0x1d0 (c9552800,0,4,77fed534) ra 802540c0 sz 56
 kdb_sysctl_panic+0x70 (c9552800,0,4,77fed534) ra 80237e58 sz 40 sysctl_root+0x12c (c9552800,0,4,e8bc5cf8) ra 80238e50 sz 48
 userland_sysctl+0x164 (c9552800,0,4,e8bc5cf8) ra 8023956c sz 104
 __sysctl+0xe4 (c9552800,0,4,e8bc5cf8) ra 806d62e8 sz 160
 trap+0xe1c (c9552800,0,4,e8bc5cf8) ra 80896e68 sz 128
 MipsUserGenException+0x1a4 (c9552800,0,4,405cd12c) ra 0 sz 0
 pid 82340, process: sysctl

Processor 0:
 restoreintr+0x14 (1,81bca820,3,0) ra 806cdc3c sz 0
 spinlock_exit+0x30 (1,81bca820,3,0) ra 8025d354 sz 24
 sleepq_release+0x64 (1,81bca820,3,0) ra 8025e670 sz 24
 sleepq_timeout+0x224 (1,81bca820,3,0) ra 80240294 sz 48
 softclock+0x434 (1,81bca820,3,0) ra 802067f8 sz 80
 ithread_loop+0x244 (1,81bca820,3,0) ra 80200e28 sz 64 fork_exit+0xc0 (1,81bca820,3,0) ra 80897c28 sz 48
 MipsNMIException+0x34 (1,81bca820,3,0) ra 0 sz 0
 pid 82340, process: sysctl

Processor 2:
 cpu_idle+0x20 (80960000,51bbc,2031df,81bca1b8) ra 80204948 sz 24 idle_proc+0x130 (80960000,51bbc,2031df,81bca1b8) ra 80200e28 sz 56 fork_exit+0xc0 (80960000,51bbc,2031df,81bca1b8) ra 80897c28 sz 48
 MipsNMIException+0x34 (80960000,51bbc,2031df,81bca1b8) ra 0 sz 0
 pid 82340, process: sysctl

Processor 3:
 cpu_idle+0x20 (80960000,51bbc,2038df,81bca300) ra 80204948 sz 24 idle_proc+0x130 (80960000,51bbc,2038df,81bca300) ra 80200e28 sz 56 fork_exit+0xc0 (80960000,51bbc,2038df,81bca300) ra 80897c28 sz 48
 MipsNMIException+0x34 (80960000,51bbc,2038df,81bca300) ra 0 sz 0
 pid 82340, process: sysctl

Processor 4:
 cpu_idle+0x20 (80960000,51bbc,2037df,81bca448) ra 80204948 sz 24 idle_proc+0x130 (80960000,51bbc,2037df,81bca448) ra 80200e28 sz 56 fork_exit+0xc0

```

(80960000,51bbc,2037df,81bca448) ra 80897c28 sz 48
MipsNMIException+0x34 (80960000,51bbc,2037df,81bca448) ra 0 sz 0
pid 82340, process: sysctl

Processor 5:
restoreintr+0x14 (1,51bbc,203edf,81bca590) ra 806cdc3c sz 0
spinlock_exit+0x30 (1,51bbc,203edf,81bca590) ra 80204a34 sz 24 idle_proc+0x21c
(1,51bbc,203edf,81bca590) ra 80200e28 sz 56 fork_exit+0xc0
(1,51bbc,203edf,81bca590) ra 80897c28 sz 48
MipsNMIException+0x34 (1,51bbc,203edf,81bca590) ra 0 sz 0
pid 82340, process: sysctl

Processor 6:
cpu_idle+0x20 (80960000,51bbc,205cdf,81bca6d8) ra 80204948 sz 24 idle_proc+0x130
(80960000,51bbc,205cdf,81bca6d8) ra 80200e28 sz 56 fork_exit+0xc0
(80960000,51bbc,205cdf,81bca6d8) ra 80897c28 sz 48
MipsNMIException+0x34 (80960000,51bbc,205cdf,81bca6d8) ra 0 sz 0
pid 82340, process: sysctl

Processor 7:
lockmgr+0x5ac (c97e8484,c8dd9800,0,c8dd9800) ra 8c11c81c sz 48
sal_sem_take+0x134 (c97e8484,c8dd9800,0,c8dd9800) ra 8c351108 sz 56
_bcm_esw_linkscan_thread+0x45c (c97e8484,c8dd9800,0,c8dd9800) ra 8c11cdb4 sz 104
sal_thread_start_wrap+0x74 (c97e8484,c8dd9800,0,c8dd9800) ra 80200e28 sz 32
fork_exit+0xc0 (c97e8484,c8dd9800,0,c8dd9800) ra 80897c28 sz 48
MipsNMIException+0x34 (c97e8484,c8dd9800,0,c8dd9800) ra 0 sz 0
pid 82340, process: sysctl
-- End of stacktrace --

```

show system core-dumps repository core (QFabric Systems)

```

user@switch> show system core-dumps repository core
Repository scope: shared
Repository head: /pbdata/export
Repository name: core
List of nodes for core repository: /pbdata/export/rdumps/

```

Node Group	Node Identifier	Num	Model	Usage
DG-0	BCF7208D-E44F-E011-802F-4171BAAC781D	0	qfx3100	0M
FM-0	73747cd8-0710-11e1-b6a4-00e081c5297e	0	fx-jvre	0M
DRE-0	77116f18-0710-11e1-a2a0-00e081c5297e	0	fx-jvre	0M
NW-NG-0	BBAK0394	0	qfx3500	0M
NW-NG-0	cd78871a-0710-11e1-878e-00e081c5297e	0	fx-jvre	0M
NW-NG-0	d0afda1e-0710-11e1-a1d0-00e081c5297e	0	fx-jvre	0M
FC-0	d31ab7a6-0710-11e1-ad1b-00e081c5297e	0	fx-jvre	0M
FC-1	d4d0f254-0710-11e1-90c3-00e081c5297e	0	fx-jvre	0M
IC-WS001	WS001	0	-	-
IC-WS001	WS001/YW3803	0	qfxc08-3008	0M
IC-WS001	WS001/YN5999	0	qfxc08-3008	0M
node-device1	BBAK0372	0	qfx3500	0M
node-device1	EE3093	0	qfx3500	0M

```

Total usage of core repository: 0M of 70000M (0.0%)

```

show system core-dumps repository log (QFabric Systems)

```

user@switch> show system core-dumps repository log
Repository scope: shared
Repository head: /pbdata/export
Repository name: log
List of nodes for log repository: /pbdata/export/rlogs/

```

Node Group	Node Identifier	Num	Model	Usage
------------	-----------------	-----	-------	-------

DG-0	BCF7208D-E44F-E011-802F-4171BAAC781D	0	qfx3100	0M
FM-0	73747cd8-0710-11e1-b6a4-00e081c5297e	1	fx-jvre	0M
DRE-0	77116f18-0710-11e1-a2a0-00e081c5297e	1	fx-jvre	0M
NW-NG-0	BBAK0394	1	qfx3500	0M
NW-NG-0	cd78871a-0710-11e1-878e-00e081c5297e	1	fx-jvre	0M
NW-NG-0	d0afda1e-0710-11e1-a1d0-00e081c5297e	3	fx-jvre	0M
FC-0	d31ab7a6-0710-11e1-ad1b-00e081c5297e	1	fx-jvre	0M
FC-1	d4d0f254-0710-11e1-90c3-00e081c5297e	1	fx-jvre	0M
IC-WS001	WS001	0	-	-
IC-WS001	WS001/YN5999	1	qfxc08-3008	0M
IC-WS001	WS001/YW3803	1	qfxc08-3008	0M
node-device1	BBAK0372	1	qfx3500	0M
node-device1	EE3093	1	qfx3500	0M
Total usage of log repository:0M of 70000M (0.0%)				

show system directory-usage

List of Syntax	Syntax on page 1093 Syntax (EX Series) on page 1093 Syntax (TX Matrix Router) on page 1093 Syntax (TX Matrix Plus Router) on page 1093 Syntax (MX Series Router) on page 1093 Syntax (QFX Series and OCX Series) on page 1093
Syntax	<pre>show system directory-usage <depth <i>number</i>> <path></pre>
Syntax (EX Series)	<pre>show system directory-usage <all-members> <depth <i>number</i>> <local> <member <i>member-id</i>> <path></pre>
Syntax (TX Matrix Router)	<pre>show system directory-usage <all-chassis all-lcc lcc <i>number</i> scc> <depth <i>number</i>> <path></pre>
Syntax (TX Matrix Plus Router)	<pre>show system directory-usage <all-chassis all-lcc lcc <i>number</i> sfc <i>number</i>> <depth <i>number</i>> <path></pre>
Syntax (MX Series Router)	<pre>show system directory-usage <all-members> <depth <i>number</i>> <local> <member <i>member-id</i>> <path></pre>
Syntax (QFX Series and OCX Series)	<pre>show system directory-usage <depth <i>number</i>> <path> <infrastructure <i>name</i>> <interconnect-device <i>name</i>> <node-group <i>name</i>></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Display directory usage information.

Options **none**—Display all directory usage information.

all-chassis—(TX Matrix routers and TX Matrix Plus routers only) (Optional) Display directory usage information about all the T640 routers (in a routing matrix based on a TX Matrix router). Display directory usage information about all the T1600 or T4000 routers (in a routing matrix based on a TX Matrix Plus router) in the chassis.

all-lcc—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display directory information for all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, display directory information for all connected T1600 or T4000 LCCs.

all-members—(EX4200 switches and MX Series routers only) (Optional) Display directory information for all members of the Virtual Chassis configuration.

depth *number*—(Optional) Depth of the directory to traverse. This option is useful when you want to limit the output shown for a large file system.

infrastructure *name*— (QFabric systems only) (Optional) Display directory information for the fabric control Routing Engines and fabric manager Routing Engines.

interconnect-device *name*— (QFabric systems only) (Optional) Display directory information for the Interconnect device.

node-group *name*— (QFabric systems only) (Optional) Display directory information for the Node group.

lcc *number*—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display directory information for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display directory information for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(EX4200 switches and MX Series routers only) (Optional) Display directory information for the local Virtual Chassis member.

member *member-id*—(EX4200 switches and MX Series routers only) (Optional) Display directory information for the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

path—(Optional) Path or root directory to traverse.

scc—(TX Matrix router only) (Optional) Display directory information for the TX Matrix router (or switch-card chassis).

sfc number—(TX Matrix Plus routers only) (Optional) Display directory information for the TX Matrix Plus router. Replace *number* with 0.

- Required Privilege Level
- view
- Related Documentation
- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)
- List of Sample Output
- [show system directory-usage scc \(TX Matrix Router\) on page 1096](#)
[show system directory-usage sfc \(TX Matrix Plus Router\) on page 1096](#)
[show system directory-usage \(QFX Series and OCX Series\) on page 1096](#)
- Output Fields
- Table 75 describes the output fields for the **show system directory-usage** command. Output fields are listed in the approximate order in which they appear.

Table 75: show system directory-usage Output Fields

Field Name	Field Description
<i>bytes</i>	Number of bytes used by files in a directory.
<i>directory-name</i>	Name of the directory.

Sample Output

show system directory-usage scc (TX Matrix Router)

```
user@host> show system directory-usage /var/tmp scc
/var/tmp
1.0K    /var/tmp/vi.recover
2.0K    /var/tmp/instmp.tPMk8u
1.0K    /var/tmp/install
        /var/tmp/instmp.GUMpur
4.8M    /var/tmp/instmp.GUMpur/packages
6.4M    /var/tmp/troy1
297M    /var/tmp/dsw
        /var/tmp/pkg_tmp.2073
83K     /var/tmp/pkg_tmp.2073/bin
        /var/tmp/instmp.oMIDb1
89K     /var/tmp/instmp.oMIDb1/bin
        /var/tmp/instmp.byhMjR
4.6M    /var/tmp/instmp.byhMjR/packages
        /var/tmp/instmp.6fqHf3
1.7M    /var/tmp/instmp.6fqHf3/packages
        /var/tmp/instmp.mljECe
4.6M    /var/tmp/instmp.mljECe/packages
```

show system directory-usage sfc (TX Matrix Plus Router)

```
user@switch> show system directory-usage /var/tmp sfc 0
sfc0-re0:
-----
        /var/tmp
46K     /var/tmp/gres-tp
        /var/tmp/sec-download
2.0K    /var/tmp/sec-download/sub-download
2.0K    /var/tmp/vi.recover
2.0K    /var/tmp/install
795M    /var/tmp/cores
766K    /var/tmp/pr440594
```

show system directory-usage (QFX Series and OCX Series)

```
user@switch> show system directory-usage
/var/tmp
30K     /var/tmp/gres-tp
2.0K    /var/tmp/rtbdb
2.0K    /var/tmp/vi.recover
2.0K    /var/tmp/install
2.0K    /var/tmp/pics
```


show system processes

List of Syntax	Syntax on page 1097 Syntax (EX Series Switches) on page 1097 Syntax (QFX Series Switches) on page 1097 Syntax (MX Series Routers) on page 1097 Syntax (OCX Series) on page 1097 Syntax (TX Matrix Routers) on page 1098 Syntax (TX Matrix Plus Router) on page 1098
Syntax	<pre>show system processes <brief detail extensive summary> <health (pid <i>process-identifer</i> process-name <i>process-name</i>)> <providers> <resource-limits (brief detail) <i>process-name</i>> <wide></pre>
Syntax (EX Series Switches)	<pre>show system processes <all-members> <brief detail extensive summary> <health (pid <i>process-identifer</i> process-name <i>process-name</i>)> <local> <member <i>member-id</i>> <providers> <resource-limits (brief detail) <i>process-name</i>> <wide></pre>
Syntax (QFX Series Switches)	<pre>show system processes <all-members> <brief detail extensive summary> <health (pid <i>process-identifer</i> process-name <i>process-name</i>)> host-processes (brief detail) <local> <member <i>member-id</i>> <providers> <resource-limits (brief detail) <i>process-name</i>> <wide></pre>
Syntax (MX Series Routers)	<pre>show system processes <all-members> <brief detail extensive summary> <health (pid <i>process-identifer</i> process-name <i>process-name</i>)> <local> <member <i>member-id</i>> <providers> <resource-limits (brief detail) <i>process-name</i>> <wide></pre>
Syntax (OCX Series)	<pre>show system processes <brief detail extensive summary > <health (pid <i>process-identifer</i> process-name <i>process-name</i>)> host-processes (brief detail) <providers></pre>

	<code><resource-limits></code> <code><wide></code>
Syntax (TX Matrix Routers)	<code>show system processes</code> <code><brief detail extensive summary></code> <code><all-chassis all-lcc lcc <i>number</i> scc></code> <code><wide></code>
Syntax (TX Matrix Plus Router)	<code>show system processes</code> <code><brief detail extensive summary></code> <code><all-chassis all-lcc lcc <i>number</i> sfc <i>number</i>></code> <code><wide></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Option sfc introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Display information about software processes that are running on the router or switch and that have controlling terminals.
Options	none —Display standard information about system processes. brief detail extensive summary —(Optional) Display the specified level of detail. adaptive-services —(Optional) Display the configuration management process that manages the configuration for stateful firewall, Network Address Translation (NAT), intrusion detection services (IDS), and IP Security (IPsec) services on the Adaptive Services PIC. alarm-control —(Optional) Display the process to configure the system alarm. all-chassis —(TX Matrix routers and TX Matrix Plus routers only) (Optional) Display standard system process information about all the T640 routers (in a routing matrix based on the TX Matrix router) or all the T1600 or T4000 routers (in a routing matrix based on the TX Matrix Plus router) in the chassis. all-lcc —(TX Matrix routers and TX Matrix Plus router only) (Optional) Display standard system process information for all T640 routers (or line-card chassis) connected to the TX Matrix router. Display standard system process information for all connected T1600 or T4000 LCCs. all-members —(EX4200 switches, QFX Series Virtual Chassis, and MX Series routers) (Optional) Display standard system process information for all members of the Virtual Chassis configuration. ancpd-service —Display the Access Node Control Protocol (ANCP) process, which works with a special Internet Group Management Protocol (IGMP) session to collect outgoing interface mapping events in a scalable manner.

application-identification—Display the process that identifies an application using intrusion detection and prevention (IDP) to allow or deny traffic based on applications running on standard or nonstandard ports.

audit-process—(Optional) Display the RADIUS accounting process.

auto-configuration—Display the Interface Auto-Configuration process.

bootp—Display the process that enables a router, switch, or interface to act as a Dynamic Host Configuration Protocol (DHCP) or bootstrap protocol (BOOTP) relay agent. DHCP relaying is disabled.

captive-portal-content-delivery—Display the HTTP redirect service by specifying the location to which a subscriber's initial Web browser session is redirected, enabling initial provisioning and service selection for the subscriber.

ce-l2tp-service—(Optional) (M10, M10i, M7i, and MX Series routers only) Display the Universal Edge Layer 2 Tunneling Protocol (L2TP) process, which establishes L2TP tunnels and Point-to-Point Protocol (PPP) sessions through L2TP tunnels.

cfm—Display Ethernet Operations, Administration, and Maintenance (OAM) connectivity fault management (CFM) process, which can be used to monitor the physical link between two switches.

chassis-control—(Optional) Display the chassis management process.

class-of-service—(Optional) Display the class-of-service (CoS) process, which controls the router's or switch's CoS configuration.

clksyncd-service—Display the external clock synchronization process, which uses synchronous Ethernet (SyncE).

craft-control—Display the process for the I/O of the craft interface.

database-replication—(EX Series switches and MX Series routers only) (Optional) Display the database replication process.

datapath-trace-service—Display the packet path tracing process.

dhcp-service—(EX Series switches and MX Series routers only) (Optional) Display the Dynamic Host Configuration Protocol process, which enables a DHCP server to allocate network IP addresses and deliver configuration settings to client hosts without user intervention.

diameter-service—(Optional) Display the diameter process.

disk-monitoring—(Optional) Display the disk monitoring process, which checks the health of the hard disk drive on the Routing Engine.

dynamic-flow-capture—(Optional) Display the dynamic flow capture (DFC) process, which controls DFC configurations on Monitoring Services III PICs.

- ecc-error-logging**—(Optional) Display the error checking and correction (ECC) process, which logs ECC parity errors in memory on the Routing Engine.
- ethernet-connectivity-fault-management**—Display the process that provides IEEE 802.1ag OAM connectivity fault management (CFM) database information for CFM maintenance association end points (MEPs) in a CFM session.
- ethernet-link-fault-management**—(EX Series switches and MX Series routers only)
(Optional) Display the process that provides the OAM link fault management (LFM) information for Ethernet interfaces.
- event-processing**—(Optional) Display the event process (eventd).
- firewall**—(Optional) Display the firewall management process, which manages the firewall configuration and enables accepting or rejecting packets that are transiting an interface on a router or switch.
- general-authentication-service**—(EX Series switches and MX Series routers only)
(Optional) Display the general authentication process.
- health (pid *process-identifier* | process-name *process-name*)**—(Optional) Display process health information, either by process id (PID) or by process name.
- host-processes**—Display process information of processes running on the host system.
(On OCX Series only) The following options are available:
- **brief | detail**—(Optional) Display the specified level of detail.
- iccp-service**—Display the Inter-Chassis Communication Protocol (ICCP) process.
- idp-policy**—Display the intrusion detection and prevention (IDP) protocol process.
- ilmi**—Display the Integrated Local Management Interface (ILMI) protocol process, which provides bidirectional exchange of management information between two ATM interfaces across a physical connection.
- inet-process**—Display the IP multicast family process.
- init**—Display the process that initializes the USB modem.
- interface-control**—(Optional) Display the interface process, which controls the router's or switch's physical interface devices and logical interfaces.
- kernel-replication**—(Optional) Display the kernel replication process, which replicates the state of the backup Routing Engine when graceful Routing Engine switchover (GRES) is configured.
- l2-learning**—(Optional) Display the Layer 2 address flooding and learning process.
- l2cpd-service**—Display the Layer 2 Control Protocol process, which enables features such as Layer 2 protocol tunneling and nonstop bridging.

lACP—(Optional) Display the Link Aggregation Control Protocol (LACP) process. LACP provides a standardized means for exchanging information between partner systems on a link to allow their link aggregation control instances to reach agreement on the identity of the LAG to which the link belongs, and then to move the link to that LAG, and to enable the transmission and reception processes for the link to function in an orderly manner.

lcc number—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display standard system process information for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display standard system process information for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(EX4200 switches, QFX Series Virtual Chassis, and MX Series routers) (Optional) Display standard system process information for the local Virtual Chassis member.

local-policy-decision-function—Display the process for the Local Policy Decision Function, which regulates collection of statistics related to applications and application groups and tracking of information about dynamic subscribers and static interfaces.

logical-system-mux—Display the logical router multiplexer process (lrmuxd), which manages the multiple instances of the routing protocols process (rpd) on a machine running logical routers.

mac-validation—Display the MAC validation process, which configures MAC address validation for subscriber interfaces created on demux interfaces in dynamic profiles on MX Series routers.

member member-id—(EX4200 switches, QFX Series Virtual Chassis, and MX Series routers) (Optional) Display standard system process information for the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

mib-process—(Optional) Display the MIB II process, which provides the router's MIB II agent.

mobile-ip—(Optional) Display the Mobile IP process, which configures Junos OS Mobile IP features.

mountd-service—(EX Series switches and MX Series routers only) (Optional) Display the service for NFS mounts requests.

mpls-traceroute—(Optional) Display the MPLS Periodic Traceroute process.

mspd—(Optional) Display the Multiservice process.

multicast-snooping—(EX Series switches and MX Series routers only) (Optional) Display the multicast snooping process, which makes Layer 2 devices such as VLAN switches aware of Layer 3 information, such as the media access control (MAC) addresses of members of a multicast group.

named-service—(Optional) Display the DNS Server process, which is used by a router or a switch to resolve hostnames into addresses.

neighbor-liveness—Display the process, which specifies the maximum length of time that the router waits for its neighbor to re-establish an LDP session.

nfsd-service—(Optional) Display the Remote NFS Server process, which provides remote file access for applications that need NFS-based transport.

ntp—Display the Network Time Protocol (NTP) process, which provides the mechanisms to synchronize time and coordinate time distribution in a large, diverse network.

packet-triggered-subscribers—Display the packet-triggered subscribers and policy control (PTSP) process, which allows the application of policies to dynamic subscribers that are controlled by a subscriber termination device.

peer-selection-service—(Optional) Display the Peer Selection Service process.

periodic-packet-services—Display the Periodic packet management process, which is responsible for processing a variety of time-sensitive periodic tasks so that other processes can more optimally direct their resources.

pfe—Display the Packet Forwarding Engine management process.

pgcp-service—(Optional) Display the pgcpd service process running on the Routing Engine.

pgm—Display the Pragmatic General Multicast (PGM) protocol process, which enables a reliable transport layer for multicast applications.

pic-services-logging—(Optional) Display the logging process for some PICs. With this process, also known as fsad (the file system access daemon), PICs send special logging information to the Routing Engine for archiving on the hard disk.

ppp—(Optional) Display the Point-to-Point Protocol (PPP) process, which is the encapsulation protocol process for transporting IP traffic across point-to-point links.

ppp-service—Display the Universal edge PPP process, which is the encapsulation protocol process for transporting IP traffic across universal edge routers.

pppoe—(Optional) Display the Point-to-Point Protocol over Ethernet (PPPoE) process, which combines PPP that typically runs over broadband connections with the Ethernet link-layer protocol that allows users to connect to a network of hosts over a bridge or access concentrator.

process-monitor—Display the process health monitor process (pmond).

providers—(Optional) Display provider processes.

redundancy-interface-process—(Optional) Display the ASP redundancy process.

remote-operations—(Optional) Display the remote operations process, which provides the ping and traceroute MIBs.

resource-cleanup—Display the resource cleanup process.

resource-limits (brief | detail) process-name—(Optional) Display process resource limits.

routing—(Optional) Display the routing protocol process.

sampling—(Optional) Display the sampling process, which performs packet sampling based on particular input interfaces and various fields in the packet header.

sbc-configuration-process—Display the session border controller (SBC) process of the border signaling gateway (BSG).

scc—(TX Matrix routers only) (Optional) Display standard system process information for the TX Matrix router (or switch-card chassis).

sdk-service—Display the SDK Service process, which runs on the Routing Engine and is responsible for communications between the SDK application and Junos OS. Although the SDK Service process is present on the router, it is turned off by default.

secure-neighbor-discovery—(EX Series switches and MX Series routers only) (Optional) Display the secure Neighbor Discovery Protocol (NDP) process, which provides support for protecting NDP messages.

send—(Optional) Display the Secure Neighbor Discovery Protocol (SEND) process, which provides support for protecting Neighbor Discovery Protocol (NDP) messages.

service-deployment—(Optional) Display the service deployment process, which enables Junos OS to work with the Session and Resource Control (SRC) software.

sfc number—(TX Matrix Plus routers only) (Optional) Display system process information for the TX Matrix Plus router. Replace *number* with 0.

snmp—Display the SNMP process, which enables the monitoring of network devices from a central location and provides the router's or switch's SNMP master agent.

sonet-aps—Display the SONET Automatic Protection Switching (APS) process, which monitors any SONET interface that participates in APS.

static-subscribers—(Optional) Display the Static subscribers process, which associates subscribers with statically configured interfaces and provides dynamic service activation and activation for these subscribers.

tunnel-oamd—(Optional) Display the Tunnel OAM process, which enables the Operations, Administration, and Maintenance of Layer 2 tunneled networks. Layer 2 protocol tunneling (L2PT) allows service providers to send Layer 2 protocol data units (PDUs) across the provider's cloud and deliver them to Juniper Networks EX Series Ethernet Switches that are not part of the local broadcast domain.

vrp—(EX Series switches and MX Series routers only) (Optional) Display the Virtual Router Redundancy Protocol (VRRP) process, which enables hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts.

watchdog—Display the watchdog timer process, which enables the watchdog timer when Junos OS encounters a problem.

wide—(Optional) Display process information that might be wider than 80 columns.

Additional Information By default, when you issue the **show system processes** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

Required Privilege Level

view

Related Documentation

- [List of Junos OS Processes](#)
- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

List of Sample Output

[show system processes on page 1107](#)
[show system processes brief on page 1108](#)
[show system processes detail on page 1108](#)
[show system processes extensive on page 1109](#)
[show system processes extensive \(EX9200 Switch\) on page 1109](#)
[show system processes host processes \(OCX1100 Switch\) on page 1110](#)
[show system processes lcc wide \(TX Matrix Routing Matrix\) on page 1110](#)
[show system processes summary on page 1111](#)
[show system processes \(TX Matrix Plus Router\) on page 1111](#)
[show system processes sfc \(TX Matrix Plus Router\) on page 1118](#)
[show system processes lcc wide \(TX Matrix Plus Routing Matrix\) on page 1121](#)
[show system processes \(QFX Series and OCX Series\) on page 1123](#)

Output Fields

Table 76 describes the output fields for the **show system processes** command. Output fields are listed in the approximate order in which they appear.

Table 76: show system processes Output Fields

Field Name	Field Description	Level of Output
last pid	Last process identifier assigned to the process.	brief extensive summary
load averages	Three load averages followed by the current time.	brief extensive summary
processes	Number of existing processes and the number of processes in each state (sleeping , running , starting , zombies , and stopped).	brief extensive summary
Mem	Information about physical and virtual memory allocation.	brief extensive summary
Active	<p>Memory allocated and actively used by the program.</p> <p>When the system is under memory pressure, the pageout process reuses memory from the free, cache, inact and, if necessary, active pages. When the pageout process runs, it scans memory to see which pages are good candidates to be unmapped and freed up. Thus, the distinction between Active and Inact memory is only used by the pageout process to determine which pool of pages to free first at the time of a memory shortage.</p> <p>The pageout process first scans the Inact list, and checks whether the pages on this list have been accessed since the time they have been listed here. The pages that have been accessed are moved from the Inact list to the Active list. On the other hand, pages that have not been accessed become prime candidates to be freed by the pageout process. If the pageout process cannot produce enough free pages from the Inact list, pages from the Active list get freed up.</p> <p>Because the pageout process runs only when the system is under memory pressure, the pages on the Inact list remain untouched – even if they have not been accessed recently – when the amount of Free memory is adequate.</p>	brief extensive summary
Inact	<p>Memory allocated but not recently used or memory freed by the programs. Inactive memory remains mapped in the address space of one or more processes and, therefore, counts toward the RSS value of those processes.</p> <p>Any amount of memory freed by the routing protocol process might still be considered part of the RES value. Generally, the kernel delays the migrating of memory out of the Inact queue into the Cache or Free list unless there is a memory shortage.</p>	brief extensive summary
Wired	Memory that is not eligible to be swapped, usually used for in-kernel memory structures and/or memory physically locked by a process.	brief extensive summary
Cache	Memory that is not associated with any program and does not need to be swapped before being reused.	brief extensive summary
Buf	Size of memory buffer used to hold data recently called from the disk.	brief extensive summary
Free	Memory that is not associated with any programs. Memory freed by a process can become Inactive , Cache , or Free , depending on the method used by the process to free the memory.	brief extensive summary

Table 76: show system processes Output Fields (*continued*)

Field Name	Field Description	Level of Output
Swap	Information about physical and virtual memory allocation. NOTE: Memory can remain swapped out indefinitely if it is not accessed again. Therefore, the show system process extensive command shows that memory is swapped to disk even though there is plenty of free memory, and such a situation is not unusual.	brief extensive summary
PID	Process identifier.	detail extensive summary
TT	Control terminal name.	none detail
STAT	Symbolic process state. The state is given by a sequence of letters. The first letter indicates the run state of the process: <ul style="list-style-type: none"> • D—In disk or other short-term, uninterruptible wait • I—Idle (sleeping longer than about 20 seconds) • R—Runnable • S—Sleeping for less than 20 seconds • T—Stopped • Z—Dead (zombie) • + —The process is in the foreground process group of its control terminal. • <—The process has raised CPU scheduling priority. • >—The process has specified a soft limit on memory requirements and is currently exceeding that limit; such a process is not swapped. • A—The process requested random page replacement. • E—The process is trying to exit. • L—The process has pages locked in core. • N—The process has reduced CPU scheduling priority. • S—The process requested first-in, first-out (FIFO) page replacement. • s—The process is a session leader. • V—The process is temporarily suspended. • W—The process is swapped out. • X—The process is being traced or debugged. 	none detail
UID	User identifier.	detail
USERNAME	Process owner.	extensive summary
PPID	Parent process identifier.	detail
CPU	(D)—Short-term CPU usage. (E and S)—Raw (unweighted) CPU usage. The value of this field is used to sort the processes in the output.	detail extensive summary
RSS	Resident set size.	detail

Table 76: show system processes Output Fields (*continued*)

Field Name	Field Description	Level of Output
WCHAN	Symbolic name of the wait channel.	detail
STARTED	Local time when the process started running.	detail
PRI	Current priority of the process. A lower number indicates a higher priority.	detail extensive summary
NI or NICE	UNIX "niceness" value. A lower number indicates a higher priority.	detail extensive summary
SIZE	Total size of the process (text, data, and stack), in kilobytes.	extensive summary
RES	Current amount of program resident memory, in kilobytes. This is also known as RSS or Resident Set Size. The RES value includes shared library pages used by the process. Any amount of memory freed by the process might still be considered part of the RES value. Generally, the kernel delays the migrating of memory out of the Inact queue into the Cache or Free list unless there is a memory shortage. This can lead to large discrepancies between the values reported by the routing protocol process and the kernel, even after the routing protocol process has freed a large amount of memory.	extensive summary
STATE	Current state of the process (for example, sleep , wait , run , idle , zombie , or stop).	extensive summary
TIME	(S)—Number of system and user CPU seconds that the process has used. (None, D, and E)—Total amount of time that the command has been running.	detail extensive summary
WCPU	Weighted CPU usage.	extensive summary
COMMAND	Command that is currently running. (MX Series routers only) When you display the software processes for an MX Series Virtual Chassis, the show system processes command does not display information about the relayd process.	detail extensive summary
THR	Number of threads in the process	extensive

Sample Output

show system processes

```

user@host> show system processes
PID  TT  STAT  TIME  COMMAND
  0  ??  DLs   0:00.70  (swapper)
  1  ??  Is    0:00.35  /sbin/init --
  2  ??  DL    0:00.00  (pagedaemon)
  3  ??  DL    0:00.00  (vmdaemon)
  4  ??  DL    0:42.37  (update)
  5  ??  DL    0:00.00  (if_jnx)
 80  ??  Ss    0:14.66  syslogd -s

```

```

 96  ??  Is    0:00.01 portmap
128  ??  Is    0:02.70 cron
173  ??  Is    0:02.24 /usr/local/sbin/sshd (sshd1)
189  ??  S      0:03.80 /sbin/watchdog -t180
190  ??  I      0:00.03 /usr/sbin/tftpd -N
191  ??  S      2:24.76 /sbin/iftpd -N
192  ??  S<    0:55.44 /usr/sbin/xntpd -N
195  ??  S      0:53.11 /usr/sbin/snmpd -N
196  ??  S      1:15.73 /usr/sbin/mib2d -N
198  ??  I      0:00.75 /usr/sbin/inetd -N
2677 ??  I      0:00.01 /usr/sbin/mgd -N
2712 ??  Ss     0:00.24 rlogind
2735 ??  R      0:00.00 /bin/ps -ax
1985 p0- S      0:07.41 ./rpd -N
2713 p0  Is    0:00.24 -tcsh (tcsh)
2726 p0  S+    0:00.07 cli

```

show system processes brief

```

user@host> show system processes brief
last pid: 543; load averages: 0.00, 0.00, 0.00 18:29:47
37 processes: 1 running, 36 sleeping

Mem: 25M Active, 3976K Inact, 19M Wired, 8346K Buf, 202M Free
Swap: 528M Total, 64K Used, 528M Free

```

show system processes detail

```

user@host> show system processes detail

```

PID	UID	PPID	CPU	PRI	NI	RSS	WCHAN	STARTED	TT	STAT	TIME	COMMAND
3151	1049	3129	2	28	0	672	-	1:13PM	p0	R+	0:00.00	ps -ax -r
1	0	0	0	10	0	376	wait	1:51PM	??	Is	0:00.29	/sbin/ini
2	0	0	0	-18	0	12	psleep	1:51PM	??	DL	0:00.00	(pagedae
3	0	0	0	28	0	12	psleep	1:51PM	??	DL	0:00.00	(vmdaemon
4	0	0	0	28	0	12	update	1:51PM	??	DL	0:07.15	(update)
5	0	0	0	2	0	12	pfesel	1:51PM	??	IL	0:02.90	(if_pfe)
27	0	1	0	10	0	17936	mfsidl	1:51PM	??	Is	0:00.46	mfs /dev/
81	0	1	0	2	0	496	select	1:52PM	??	Ss	0:31.21	syslogd -
119	1	1	0	2	0	492	select	1:52PM	??	Is	0:00.00	portmap
134	0	1	0	2	0	580	select	1:52PM	??	S	0:02.95	amd -p -a
151	0	1	0	18	0	532	pause	1:52PM	??	Is	0:00.34	cron
183	0	1	0	2	0	420	select	1:52PM	??	Ss	0:00.07	/usr/loca
206	0	1	0	18	0	72	pause	1:52PM	??	S	0:00.51	/sbin/wat
207	0	1	0	2	0	520	select	1:52PM	??	I	0:00.16	/usr/sbin
208	0	1	0	2	0	536	select	1:52PM	??	S	0:08.21	/sbin/dcd
210	0	1	255	2	-12	740	select	1:52PM	??	S<	0:05.83	/usr/sbin
211	0	1	0	2	0	376	select	1:52PM	??	S	0:00.03	/usr/sbin
215	0	1	0	2	0	548	select	1:52PM	??	I	0:00.50	/usr/sbin
219	0	1	0	3	0	540	ttyin	1:52PM	v0	Is+	0:00.02	/usr/libe
220	0	1	0	3	0	540	ttyin	1:52PM	v1	Is+	0:00.01	/usr/libe
221	0	1	0	3	0	540	ttyin	1:52PM	v2	Is+	0:00.01	/usr/libe
222	0	1	0	3	0	540	ttyin	1:52PM	v3	Is+	0:00.01	/usr/libe
735	0	1	0	2	0	468	select	2:47PM	??	S	0:19.14	/usr/sbin
736	0	1	0	2	0	212	select	2:47PM	??	S	0:14.13	/usr/sbin
1380	0	1	0	3	0	888	ttyin	7:32PM	d0	Is+	0:00.46	bash
3019	0	207	0	2	0	636	select	10:49AM	??	Ss	0:02.93	tnp.chass
3122	0	1380	0	2	0	1764	select	12:33PM	d0	S	0:00.77	./rpd -N
3128	0	215	0	2	0	580	select	12:45PM	??	Ss	0:00.12	rlogind
3129	1049	3128	0	18	0	944	pause	12:45PM	p0	Ss	0:00.14	-tcsh (tc
0	0	0	0	-18	0	0	sched	1:51PM	??	DLs	0:00.07	(swapper

show system processes extensive

```
user@host> show system processes extensive
```

```
Mem: 241M Active, 99M Inact, 78M Wired, 325M Cache, 69M Buf, 1251M Free
```

```
Swap: 2048M Total, 2048M Free
```

PID	USERNAME	THR	PRI	NICE	SIZE	RES	STATE	TIME	WCPU	COMMAND
11	root	1	171	52	OK	12K	RUN	807.5H	98.73%	idle
13	root	1	-20	-139	OK	12K	WAIT	36:17	0.00%	swi7: clock sio
1499	root	1	96	0	7212K	3040K	select	34:01	0.00%	license-check
1621	root	1	96	0	20968K	11216K	select	20:25	0.00%	mib2d
1465	root	2	8	-88	115M	11748K	nanslp	14:32	0.00%	chassisd
1478	root	1	96	0	6336K	3816K	select	11:28	0.00%	ppmd
20	root	1	-68	-187	OK	12K	WAIT	10:28	0.00%	irq10: em0 em1+++*
1490	root	1	96	0	11792K	4336K	select	9:44	0.00%	shm-rtssdbd
1618	root	1	96	0	39584K	7464K	select	8:47	0.00%	pfed
1622	root	1	96	0	15268K	10988K	select	6:16	0.00%	snmpd
1466	root	1	96	0	7408K	2896K	select	5:44	0.00%	alarmd
7	root	1	-16	0	OK	12K	client	5:09	0.00%	ifstate notify
1480	root	1	96	0	5388K	2660K	select	4:29	0.00%	ksyncd
12	root	1	-40	-159	OK	12K	WAIT	4:15	0.00%	swi2: netisr 0
1462	root	1	96	0	1836K	1240K	select	3:57	0.00%	bslockd
55	root	1	-16	0	OK	12K	-	3:44	0.00%	schedcpu
1392	root	1	16	0	OK	12K	bcmsem	3:37	0.00%	bcmLINK.0
47	root	1	-16	0	OK	12K	psleep	3:25	0.00%	vmkmemdaemon
36	root	1	20	0	OK	12K	syncer	2:46	0.00%	syncer
1484	root	1	96	0	7484K	3428K	select	2:38	0.00%	clksyncd
1616	root	1	96	0	4848K	2848K	select	2:18	0.00%	irsd
1487	root	1	96	0	32800K	6992K	select	2:10	0.00%	smid
1623	root	1	96	0	34616K	5464K	select	2:01	0.00%	dcd
15	root	1	-16	0	OK	12K	-	1:59	0.00%	yarrow
49	root	1	-16	0	OK	12K	.	1:51	0.00%	ddostasks

show system processes extensive (EX9200 Switch)

```
user@switch> show system processes extensive
```

```
last pid: 3372; load averages: 0.02, 0.02, 0.00 up 0+01:42:22 16:39:57
151 processes: 4 running, 131 sleeping, 1 zombie, 15 waiting
```

```
Mem: 935M Active, 122M Inact, 108M Wired, 838M Cache, 214M Buf, 5872M Free
```

```
Swap: 8192M Total, 8192M Free
```

PID	USERNAME	THR	PRI	NICE	SIZE	RES	STATE	TIME	WCPU	COMMAND
10	root	1	171	52	OK	16K	RUN	96:34	92.19%	idle
3317	root	1	97	0	40412K	30944K	select	0:00	5.13%	mgd
3316	root	1	96	0	26672K	20516K	select	0:00	3.08%	cli
1626	root	2	8	-88	124M	20332K	nanslp	3:19	2.39%	chassisd
260	root	1	-8	0	OK	16K	mdwait	0:16	0.00%	md16
19	root	1	-68	-187	OK	16K	WAIT	0:12	0.00%	irq11: em0 em1 em2*
1642	root	1	96	0	8052K	3936K	RUN	0:10	0.00%	clksyncd
11	root	1	-20	-139	OK	16K	WAIT	0:07	0.00%	swi7: clock sio
154	root	1	-8	0	OK	16K	mdwait	0:06	0.00%	md8
1784	root	1	96	0	98M	33720K	select	0:05	0.00%	authd
1646	root	1	96	0	7776K	2944K	select	0:03	0.00%	license-check
1807	root	1	96	0	41340K	9944K	select	0:02	0.00%	mib2d

```
[...Output truncated...]
```

show system processes host processes (OCX1100 Switch)

```
user@switch> show system processes host processes
fpc0:
```

```
-----
top - 14:14:32 up 2:05, 0 users, load average: 0.11, 0.39, 0.39
Tasks: 101 total, 1 running, 98 sleeping, 0 stopped, 2 zombie
Cpu(s): 3.1%us, 2.2%sy, 0.0%ni, 94.2%id, 0.4%wa, 0.0%hi, 0.0%si, 0.0%st
Mem: 3881300k total, 2667040k used, 1214260k free, 53232k buffers
Swap: 15620k total, 0k used, 15620k free, 808492k cached
```

PID	USER	PR	NI	VIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND
2780	root	20	0	1860m	1.5g	3780	S	14	41.7	20:56.05	kvm
1482	bind	20	0	24676	5912	1944	S	2	0.2	0:00.07	named
4631	root	20	0	648m	94m	13m	S	2	2.5	4:19.59	dcpfe
9230	root	20	0	15208	1092	832	R	2	0.0	0:00.01	top
1	root	20	0	4216	660	576	S	0	0.0	2:09.61	init
2	root	20	0	0	0	0	S	0	0.0	0:00.00	kthreadd
3	root	20	0	0	0	0	S	0	0.0	0:00.21	ksoftirqd/0
4	root	20	0	0	0	0	S	0	0.0	0:00.00	kworker/0:0
5	root	0	-20	0	0	0	S	0	0.0	0:00.00	kworker/0:0H
7	root	RT	0	0	0	0	S	0	0.0	0:00.52	migration/0
8	root	20	0	0	0	0	S	0	0.0	0:04.36	rcu_preempt
9	root	20	0	0	0	0	S	0	0.0	0:00.00	rcu_bh
10	root	20	0	0	0	0	S	0	0.0	0:00.00	rcu_sched
11	root	RT	0	0	0	0	S	0	0.0	0:00.53	migration/1

```
[...Output truncated...]
```

show system processes lcc wide (TX Matrix Routing Matrix)

```
user@host> show system processes lcc 2 wide
lcc2-re0:
```

PID	TT	STAT	TIME	COMMAND
0	??	DLs	0:00.00	(swapper)
1	??	ILs	0:00.10	/sbin/preinit -- (init)
2	??	DL	0:00.00	(pagedaemon)
3	??	DL	0:00.00	(vmddaemon)
4	??	DL	0:00.00	(bufddaemon)
5	??	DL	0:00.04	(syncer)
6	??	DL	0:00.00	(netdaemon)
7	??	IL	0:00.00	(if_pic_listen)
8	??	IL	0:00.00	(scs_housekeeping)
9	??	IL	0:00.00	(if_pfe_listen)
10	??	DL	0:00.00	(vmuncachedaemon)
11	??	SL	0:00.02	(cb_poll)
172	??	ILs	0:00.21	mfs -o noauto /dev/ad1s1b /tmp (newfs)
2909	??	Is	0:00.00	pccardd
2932	??	Ss	0:00.07	syslogd -r -s
3039	??	Is	0:00.00	cron
3217	??	I	0:00.00	/sbin/watchdog -d
3218	??	I	0:00.02	/usr/sbin/tnetd -N
3221	??	S	0:00.11	/usr/sbin/alarmd -N
3222	??	S	0:00.85	/usr/sbin/craftd -N
3223	??	S	0:00.05	/usr/sbin/mgd -N
3224	??	I	0:00.02	/usr/sbin/inetd -N
3225	??	I	0:00.00	/usr/sbin/tnp.sntpd -N
3226	??	I	0:00.01	/usr/sbin/tnp.sntpc -N
3228	??	I	0:00.01	/usr/sbin/smartd -N
3231	??	I	0:00.01	/usr/sbin/eccd -N

```

3425 ?? S      0:00.09 /usr/sbin/dfwd -N
3426 ?? S      0:00.19 /sbin/dcd -N
3427 ?? I      0:00.04 /usr/sbin/pfed -N
3430 ?? S      0:00.10 /usr/sbin/ksyncd -N
3482 ?? S      1:53.63 /usr/sbin/chassisd -N
4285 ?? SL     0:00.01 (peer proxy)
4286 ?? SL     0:00.00 (peer proxy)
4303 ?? Ss     0:00.00 mgd: (mgd) (root) (mgd)
4304 ?? R      0:00.00 /bin/ps -ax -ww
3270 d0 Is+    0:00.00 /usr/libexec/getty std.9600 ttyd0

```

show system processes summary

```

user@host> show system processes summary
last pid: 543; load averages: 0.00, 0.00, 0.00 18:29:47
37 processes: 1 running, 36 sleeping

```

```

Mem: 25M Active, 3976K Inact, 19M Wired, 8346K Buf, 202M Free
Swap: 528M Total, 64K Used, 528M Free

```

PID	USERNAME	PRI	NICE	SIZE	RES	STATE	TIME	WCPU	CPU	COMMAND
527	root	2	0	176K	580K	select	0:00	0.04%	0.04%	rlogind
543	root	30	0	604K	768K	RUN	0:00	0.00%	0.00%	top

show system processes (TX Matrix Plus Router)

```

user@host> show system processes
sfc0-re0:

```

```

-----
PID  TT  STAT      TIME COMMAND
 0  ??  Wls      0:00.00 [swapper]
 1  ??  ILs      0:00.18 /packages/mnt/jbase/sbin/init --
 2  ??  DL       0:00.20 [g_event]
 3  ??  DL       0:00.39 [g_up]
 4  ??  DL       0:00.32 [g_down]
 5  ??  DL       0:00.00 [thread taskq]
 6  ??  DL       0:00.09 [kqueue taskq]
 7  ??  DL       0:00.01 [pagedaemon]
 8  ??  DL       0:00.00 [vmdaemon]
 9  ??  DL       0:06.63 [pagezero]
10  ??  DL       0:00.00 [ktrace]
11  ??  RL      310:52.98 [idle]
12  ??  WL       0:11.03 [swi2: net]
13  ??  WL       0:27.58 [swi7: clock sio]
14  ??  WL       0:00.00 [swi6: vm]
15  ??  DL       0:03.02 [yarrow]
16  ??  WL       0:00.00 [swi9: +]
17  ??  WL       0:00.00 [swi8: +]
18  ??  WL       0:00.00 [swi5: cambio]
19  ??  WL       0:00.00 [swi9: task queue]
20  ??  WL       0:11.41 [irq16: uhci0 uhci*]
21  ??  DL       0:00.00 [usb0]
22  ??  DL       0:00.00 [usbtask]
23  ??  WL       0:39.51 [irq17: uhci1 uhci*]
24  ??  DL       0:00.00 [usb1]
25  ??  WL       0:00.00 [irq18: uhci2 uhci*]
26  ??  DL       0:00.83 [usb2]
27  ??  DL       0:00.00 [usb3]
28  ??  DL       0:00.00 [usb4]
29  ??  DL       0:00.00 [usb5]
30  ??  DL       0:00.73 [usb6]

```

```

31 ?? DL 0:00.00 [usb7]
32 ?? WL 0:00.00 [irq14: ata0]
33 ?? WL 0:00.00 [irq15: ata1]
34 ?? WL 0:00.00 [irq1: atkbd0]
35 ?? WL 0:00.00 [swi0: sio]
36 ?? WL 0:00.00 [irq11: isab0]
37 ?? WL 0:00.00 [swi3: ip6opt ipopt]
38 ?? WL 0:00.00 [swi4: ip6mismatch+]
39 ?? WL 0:00.00 [swi1: ipfwd]
40 ?? DL 0:00.02 [bufdaemon]
41 ?? DL 0:00.02 [vn1ru]
42 ?? DL 0:00.39 [syncer]
43 ?? DL 0:00.05 [softdepflush]
44 ?? DL 0:00.00 [netdaemon]
45 ?? DL 0:00.02 [vmuncachedaemon]
46 ?? DL 0:00.00 [if_pic_listen]
47 ?? DL 0:00.35 [vmkmemdaemon]
48 ?? DL 0:00.00 [cb_poll]
49 ?? DL 0:00.06 [if_pfe_listen]
50 ?? DL 0:00.00 [scs_housekeeping]
51 ?? IL 0:00.00 [kern_dump_proc]
52 ?? IL 0:00.00 [nfsiod 0]
53 ?? IL 0:00.00 [nfsiod 1]
54 ?? IL 0:00.00 [nfsiod 2]
55 ?? IL 0:00.00 [nfsiod 3]
56 ?? DL 0:00.37 [schedcpu]
57 ?? DL 0:00.56 [md0]
79 ?? DL 0:02.58 [md1]
100 ?? DL 0:00.03 [md2]
118 ?? DL 0:00.01 [md3]
139 ?? DL 0:00.95 [md4]
160 ?? DL 0:00.12 [md5]
181 ?? DL 0:00.00 [md6]
217 ?? DL 0:00.02 [md7]
227 ?? DL 0:00.05 [md8]
1341 ?? SL 0:01.34 [bcmTX]
1342 ?? SL 0:01.68 [bcmXGS3AsyncTX]
1343 ?? SL 0:41.40 [bcmLINK.0]
1345 ?? SL 0:33.83 [bcmLINK.1]
1350 ?? Is 0:00.01 /usr/sbin/cron
1502 ?? S 0:00.01 /sbin/watchdog -t-1
1503 ?? S 0:00.86 /usr/libexec/bslockd -mp -N
1504 ?? S 0:00.01 /usr/sbin/tnetd -N
1507 ?? S 0:01.32 /usr/sbin/alarmd -N
1508 ?? S 0:14.54 /usr/sbin/craftd -N
1509 ?? S 0:01.19 /usr/sbin/mgd -N
1512 ?? I 0:00.05 /usr/sbin/inetd -N
1513 ?? S 0:00.10 /usr/sbin/tnp.snmpd -N
1517 ?? S 0:00.11 /usr/sbin/smartd -N
1525 ?? S 0:01.10 /usr/sbin/idpd -N
1526 ?? S 0:01.43 /usr/sbin/license-check -U -M -p 10 -i 10
1527 ?? I 0:00.01 /usr/libexec/getty Pc ttyv0
1616 ?? DL 0:00.30 [peer proxy]
1617 ?? DL 0:00.32 [peer proxy]
1618 ?? DL 0:00.34 [peer proxy]
1619 ?? DL 0:00.30 [peer proxy]
2391 ?? Is 0:00.01 telnetd
7331 ?? Ss 0:00.03 telnetd
9538 ?? DL 0:01.16 [jsr_kkcm]
9613 ?? DL 0:00.18 [peer proxy]
23781 ?? Ss 0:00.01 telnetd

```



```

23926 ?? Ss 0:00.01 mgd: (mgd) (regress)/dev/tty2 (mgd)
36867 ?? S 0:03.14 /usr/sbin/rpd -N
36874 ?? S 0:00.08 /usr/sbin/lmpd
36876 ?? S 0:00.17 /usr/sbin/lacpd -N
36877 ?? S 0:00.15 /usr/sbin/bfdd -N
36878 ?? S 0:05.05 /usr/sbin/ppmd -N
36907 ?? S 0:25.07 /usr/sbin/chassisd -N
37775 ?? S 0:00.01 /usr/sbin/bdbrepd -N
45727 ?? S 0:00.02 /usr/sbin/xntpd -j -N -g (ntpd)
45729 ?? S 0:00.38 /usr/sbin/l2ald -N
45730 ?? S< 0:00.12 /usr/sbin/apspd -N
45731 ?? SN 0:00.10 /usr/sbin/sampled -N
45732 ?? S 0:00.03 /usr/sbin/ilmid -N
45733 ?? S 0:00.09 /usr/sbin/rmopd -N
45734 ?? S 0:00.30 /usr/sbin/cosd
45735 ?? I 0:00.00 /usr/sbin/rtspd -N
45736 ?? S 0:00.06 /usr/sbin/fsad -N
45737 ?? S 0:00.05 /usr/sbin/rdd -N
45738 ?? S 0:00.10 /usr/sbin/pppd -N
45739 ?? S 0:00.05 /usr/sbin/dfcd -N
45740 ?? S 0:00.07 /usr/sbin/lfmd -N
45741 ?? S 0:00.01 /usr/sbin/mpiisoamd -N
45742 ?? I 0:00.01 /usr/sbin/sendd -N
45743 ?? S 0:00.08 /usr/sbin/appidd -N
45744 ?? S 0:00.05 /usr/sbin/mspd -N
45745 ?? S 0:00.25 /usr/sbin/jdiameterd -N
45746 ?? S 0:00.10 /usr/sbin/pfed -N
45747 ?? S 0:00.19 /usr/sbin/lpdfd -N
45748 ?? S 0:00.63 /sbin/dcd -N
45750 ?? S 0:00.45 /usr/sbin/mib2d -N
45751 ?? S 0:00.15 /usr/sbin/dfwd -N
45752 ?? S 0:00.15 /usr/sbin/irsd -N
45764 ?? S 0:20.59 /usr/sbin/snmpd -N
56479 ?? Ss 0:00.00 mgd: (mgd) (root) (mgd)
56480 ?? R 0:00.00 /bin/ps -ax
1142 d0- I 0:00.01 /usr/sbin/usbd -N
1160 d0- S 0:29.17 /usr/sbin/eventd -N -r -s -A
6527 d0 Is+ 0:00.00 /usr/libexec/getty std.9600 ttyd0
2392 p1 Is 0:00.00 login [pam] (login)
2393 p1 I 0:00.00 -csh (csh)
2394 p1 I 0:00.00 su -
2395 p1 I+ 0:00.01 -su (csh)
23782 p2 Is 0:00.00 login [pam] (login)
23881 p2 I 0:00.00 -csh (csh)
23925 p2 S+ 0:00.03 cli
7332 p3 Is 0:00.00 login [pam] (login)
7333 p3 I 0:00.00 -csh (csh)
23780 p3 S+ 0:00.02 telnet aj

```

lcc0-re0:

```

-----
PID TT STAT TIME COMMAND
0 ?? Wls 0:00.00 [swapper]
1 ?? ILs 0:00.16 /packages/mnt/jbase/sbin/init --
2 ?? DL 0:00.01 [g_event]
3 ?? DL 0:00.16 [g_up]
4 ?? DL 0:00.11 [g_down]
5 ?? DL 0:00.00 [thread taskq]
6 ?? DL 0:00.00 [kqueue taskq]
7 ?? DL 0:00.00 [pagedaemon]
8 ?? DL 0:00.00 [vmdaemon]

```

```

 9 ?? DL      0:01.77 [pagezero]
10 ?? DL      0:00.00 [ktrace]
11 ?? RL     17:22.31 [idle]
12 ?? WL      0:00.32 [swi2: net]
13 ?? WL      0:01.21 [swi7: clock sio]
14 ?? WL      0:00.00 [swi6: vm]
15 ?? DL      0:00.10 [yarrow]
16 ?? WL      0:00.00 [swi9: +]
17 ?? WL      0:00.00 [swi8: +]
18 ?? WL      0:00.00 [swi5: cambio]
19 ?? WL      0:00.00 [swi9: task queue]
20 ?? WL      0:02.73 [irq10: bcm0 uhci1*]
21 ?? WL      0:00.02 [irq11: cb0 uhci0+*]
22 ?? DL      0:00.00 [usb0]
23 ?? DL      0:00.00 [usbtask]
24 ?? DL      0:00.00 [usb1]
25 ?? DL      0:00.05 [usb2]
26 ?? DL      0:00.00 [usb3]
27 ?? DL      0:00.00 [usb4]
28 ?? DL      0:00.00 [usb5]
29 ?? DL      0:00.04 [usb6]
30 ?? DL      0:00.00 [usb7]
31 ?? WL      0:00.00 [irq14: ata0]
32 ?? WL      0:00.00 [irq15: ata1]
33 ?? WL      0:00.00 [irq1: atkbd0]
34 ?? WL      0:00.00 [swi0: sio]
35 ?? WL      0:00.00 [swi3: ip6opt ipopt]
36 ?? WL      0:00.00 [swi4: ip6mismatch+]
37 ?? WL      0:00.00 [swi1: ipfwd]
38 ?? DL      0:00.00 [bufdaemon]
39 ?? DL      0:00.00 [vn1ru]
40 ?? DL      0:00.01 [syncer]
41 ?? DL      0:00.00 [softdepflush]
42 ?? DL      0:00.00 [netdaemon]
43 ?? DL      0:00.00 [vmuncachedaemon]
44 ?? DL      0:00.00 [if_pic_listen]
45 ?? DL      0:00.02 [vmkmemdaemon]
46 ?? DL      0:00.01 [cb_poll]
47 ?? DL      0:00.00 [if_pfe_listen]
48 ?? DL      0:00.00 [scs_housekeeping]
49 ?? IL      0:00.00 [kern_dump_proc]
50 ?? IL      0:00.00 [nfsiod 0]
51 ?? IL      0:00.00 [nfsiod 1]
52 ?? IL      0:00.00 [nfsiod 2]
53 ?? IL      0:00.00 [nfsiod 3]
54 ?? DL      0:00.01 [schedcpu]
55 ?? DL      0:00.73 [md0]
77 ?? DL      0:03.54 [md1]
98 ?? DL      0:00.37 [md2]
116 ?? DL     0:00.02 [md3]
137 ?? DL     0:00.56 [md4]
158 ?? DL     0:00.15 [md5]
179 ?? DL     0:00.00 [md6]
215 ?? DL     0:00.03 [md7]
225 ?? DL     0:00.03 [md8]
1078 ?? DL    0:00.00 [jsr_kkcm]
1363 ?? SL    0:00.09 [bcmTX]
1364 ?? SL    0:00.10 [bcmXGS3AsyncTX]
1365 ?? SL    0:03.08 [bcmLINK.0]
1370 ?? Is    0:00.00 /usr/sbin/cron
1522 ?? S     0:00.00 /sbin/watchdog -t-1

```

```

1523 ?? S      0:00.05 /usr/libexec/bslockd -mp -N
1524 ?? I      0:00.01 /usr/sbin/tnetd -N
1526 ?? S      0:04.98 /usr/sbin/chassisd -N
1527 ?? S      0:00.04 /usr/sbin/alarmd -N
1528 ?? I      0:00.40 /usr/sbin/craftd -N
1529 ?? S      0:00.08 /usr/sbin/mgd -N
1532 ?? I      0:00.04 /usr/sbin/inetd -N
1533 ?? I      0:00.00 /usr/sbin/tnp.sntpd -N
1534 ?? I      0:00.00 /usr/sbin/tnp.sntpc -N
1536 ?? S      0:00.01 /usr/sbin/smartd -N
1540 ?? I      0:00.07 /usr/sbin/jcsd -N
1541 ?? S      0:00.11 /usr/sbin/idpd -N
1542 ?? I      0:00.00 /usr/libexec/getty Pc ttyv0
2089 ?? DL     0:00.01 [peer proxy]
2090 ?? DL     0:00.01 [peer proxy]
2091 ?? DL     0:00.01 [peer proxy]
2657 ?? S      0:00.02 /usr/sbin/dfwd -N
2658 ?? S      0:00.02 /sbin/dcd -N
2659 ?? S      0:00.05 /usr/sbin/snmpd -N
2660 ?? S      0:00.01 /usr/sbin/mib2d -N
2661 ?? S      0:00.01 /usr/sbin/pfed -N
2662 ?? S      0:00.01 /usr/sbin/irsd -N
2667 ?? S      0:00.13 /usr/sbin/ksyncd -N
2690 ?? Ss     0:00.00 mgd: (mgd) (root) (mgd)
2691 ?? R      0:00.00 /bin/ps -ax
1164 d0- S     0:00.00 /usr/sbin/usbd -N
1182 d0- S     0:00.34 /usr/sbin/eventd -N -r -s -A
1543 d0 Is+    0:00.00 /usr/libexec/getty std.9600 ttyd0

```

lcc1-re0:

```

-----
PID TT  STAT      TIME COMMAND
  0 ??  Wls      0:00.00 [swapper]
  1 ??  ILs      0:00.17 /packages/mnt/jbase/sbin/init --
  2 ??  DL       0:00.01 [g_event]
  3 ??  DL       0:00.16 [g_up]
  4 ??  DL       0:00.11 [g_down]
  5 ??  DL       0:00.00 [thread taskq]
  6 ??  DL       0:00.00 [kqueue taskq]
  7 ??  DL       0:00.00 [pagedaemon]
  8 ??  DL       0:00.00 [vmdaemon]
  9 ??  DL       0:01.77 [pagezero]
 10 ??  DL       0:00.00 [ktrace]
 11 ??  RL      17:22.83 [idle]
 12 ??  WL       0:00.35 [swi2: net]
 13 ??  WL       0:01.20 [swi7: clock sio]
 14 ??  WL       0:00.00 [swi6: vm]
 15 ??  DL       0:00.10 [yarrow]
 16 ??  WL       0:00.00 [swi9: +]
 17 ??  WL       0:00.00 [swi8: +]
 18 ??  WL       0:00.00 [swi5: cambio]
 19 ??  WL       0:00.00 [swi9: task queue]
 20 ??  WL       0:02.87 [irq10: bcm0 uhci1*]
 21 ??  WL       0:00.02 [irq11: cb0 uhci0+*]
 22 ??  DL       0:00.00 [usb0]
 23 ??  DL       0:00.00 [usbtask]
 24 ??  DL       0:00.00 [usb1]
 25 ??  DL       0:00.05 [usb2]
 26 ??  DL       0:00.00 [usb3]
 27 ??  DL       0:00.00 [usb4]
 28 ??  DL       0:00.00 [usb5]

```

```

29 ?? DL 0:00.04 [usb6]
30 ?? DL 0:00.00 [usb7]
31 ?? WL 0:00.00 [irq14: ata0]
32 ?? WL 0:00.00 [irq15: ata1]
33 ?? WL 0:00.00 [irq1: atkbd0]
34 ?? WL 0:00.00 [swi0: sio]
35 ?? WL 0:00.00 [swi3: ip6opt ipopt]
36 ?? WL 0:00.00 [swi4: ip6mismatch+]
37 ?? WL 0:00.00 [swi1: ipfwd]
38 ?? DL 0:00.00 [bufdaemon]
39 ?? DL 0:00.00 [vn1ru]
40 ?? DL 0:00.01 [syncer]
41 ?? DL 0:00.00 [softdepflush]
42 ?? DL 0:00.00 [netdaemon]
43 ?? DL 0:00.00 [vmuncachedaemon]
44 ?? DL 0:00.00 [if_pic_listen]
45 ?? DL 0:00.02 [vmkmemdaemon]
46 ?? DL 0:00.01 [cb_poll]
47 ?? DL 0:00.00 [if_pfe_listen]
48 ?? DL 0:00.00 [scs_housekeeping]
49 ?? IL 0:00.00 [kern_dump_proc]
50 ?? IL 0:00.00 [nfsiod 0]
51 ?? IL 0:00.00 [nfsiod 1]
52 ?? IL 0:00.00 [nfsiod 2]
53 ?? IL 0:00.00 [nfsiod 3]
54 ?? DL 0:00.02 [schedcpu]
55 ?? DL 0:00.75 [md0]
77 ?? DL 0:03.40 [md1]
98 ?? DL 0:00.37 [md2]
116 ?? DL 0:00.02 [md3]
137 ?? DL 0:00.56 [md4]
158 ?? DL 0:00.15 [md5]
179 ?? DL 0:00.00 [md6]
215 ?? DL 0:00.03 [md7]
225 ?? DL 0:00.03 [md8]
1052 ?? DL 0:00.00 [jsr_kkcm]
1337 ?? SL 0:00.09 [bcmTX]
1338 ?? SL 0:00.10 [bcmXGS3AsyncTX]
1339 ?? SL 0:03.10 [bcmLINK.0]
1344 ?? Is 0:00.00 /usr/sbin/cron
1496 ?? S 0:00.00 /sbin/watchdog -t-1
1497 ?? S 0:00.05 /usr/libexec/bslockd -mp -N
1498 ?? I 0:00.01 /usr/sbin/tnetd -N
1500 ?? S 0:04.97 /usr/sbin/chassisd -N
1501 ?? S 0:00.04 /usr/sbin/alarmd -N
1502 ?? I 0:00.40 /usr/sbin/craftd -N
1503 ?? S 0:00.08 /usr/sbin/mgd -N
1506 ?? I 0:00.04 /usr/sbin/inetd -N
1507 ?? I 0:00.00 /usr/sbin/tnp.sntpd -N
1508 ?? I 0:00.00 /usr/sbin/tnp.sntpc -N
1510 ?? S 0:00.01 /usr/sbin/smartd -N
1514 ?? I 0:00.07 /usr/sbin/jcsd -N
1515 ?? S 0:00.18 /usr/sbin/idpd -N
1516 ?? I 0:00.00 /usr/libexec/getty Pc ttyv0
2068 ?? DL 0:00.01 [peer proxy]
2069 ?? DL 0:00.01 [peer proxy]
2070 ?? DL 0:00.01 [peer proxy]
2666 ?? S 0:00.02 /sbin/dcd -N
2667 ?? S 0:00.01 /usr/sbin/irsd -N
2668 ?? S 0:00.01 /usr/sbin/pfed -N
2669 ?? S 0:00.05 /usr/sbin/snmpd -N

```

```

2670 ?? S      0:00.01 /usr/sbin/mib2d -N
2671 ?? S      0:00.02 /usr/sbin/dfwd -N
2675 ?? S      0:00.13 /usr/sbin/ksyncd -N
2699 ?? Ss     0:00.00 mgd: (mgd) (root) (mgd)
2700 ?? R      0:00.00 /bin/ps -ax
1138 d0- S     0:00.00 /usr/sbin/usbd -N
1156 d0- S     0:00.37 /usr/sbin/eventd -N -r -s -A
1517 d0 Is+    0:00.00 /usr/libexec/getty std.9600 ttyd0

```

```
lcc2-re0:
```

```

-----
PID TT  STAT      TIME COMMAND
  0 ??  Wls      0:00.00 [swapper]
  1 ??  ILs      0:00.18 /packages/mnt/jbase/sbin/init --
  2 ??  DL       0:00.01 [g_event]
  3 ??  DL       0:00.17 [g_up]
  4 ??  DL       0:00.12 [g_down]
  5 ??  DL       0:00.00 [thread taskq]
  6 ??  DL       0:00.00 [kqueue taskq]
  7 ??  DL       0:00.00 [pagedaemon]
  8 ??  DL       0:00.00 [vmdaemon]
  9 ??  DL       0:01.77 [pagezero]
 10 ??  DL       0:00.00 [ktrace]
 11 ??  RL      17:19.13 [idle]
 12 ??  WL       0:00.36 [swi2: net]
 13 ??  WL       0:01.20 [swi7: clock sio]
 14 ??  WL       0:00.00 [swi6: vm]
 15 ??  DL       0:00.13 [yarrow]
 16 ??  WL       0:00.00 [swi9: +]
 17 ??  WL       0:00.00 [swi8: +]
 18 ??  WL       0:00.00 [swi5: cambio]
 19 ??  WL       0:00.00 [swi9: task queue]
 20 ??  WL       0:03.03 [irq10: bcm0 uhci1*]
 21 ??  WL       0:00.02 [irq11: cb0 uhci0+*]
 22 ??  DL       0:00.00 [usb0]
 23 ??  DL       0:00.00 [usbtask]
 24 ??  DL       0:00.00 [usb1]
 25 ??  DL       0:00.05 [usb2]
 26 ??  DL       0:00.00 [usb3]
 27 ??  DL       0:00.00 [usb4]
 28 ??  DL       0:00.00 [usb5]
 29 ??  DL       0:00.04 [usb6]
 30 ??  DL       0:00.00 [usb7]
 31 ??  WL       0:00.00 [irq14: ata0]
 32 ??  WL       0:00.00 [irq15: ata1]
 33 ??  WL       0:00.00 [irq1: atkbd0]
 34 ??  WL       0:00.00 [swi0: sio]
 35 ??  WL       0:00.00 [swi3: ip6opt ipopt]
 36 ??  WL       0:00.00 [swi4: ip6mismatch+]
 37 ??  WL       0:00.00 [swi1: ipfwd]
 38 ??  DL       0:00.00 [bufdaemon]
 39 ??  DL       0:00.00 [vn1ru]
 40 ??  DL       0:00.01 [syncer]
 41 ??  DL       0:00.00 [softdepflush]
 42 ??  DL       0:00.00 [netdaemon]
 43 ??  DL       0:00.00 [vmuncachedaemon]
 44 ??  DL       0:00.00 [if_pic_listen]
 45 ??  DL       0:00.02 [vmkmemdaemon]
 46 ??  DL       0:00.01 [cb_poll]
 47 ??  DL       0:00.00 [if_pfe_listen]
 48 ??  DL       0:00.00 [scs_housekeeping]

```

```

49 ?? IL 0:00.00 [kern_dump_proc]
50 ?? IL 0:00.00 [nfsiod 0]
51 ?? IL 0:00.00 [nfsiod 1]
52 ?? IL 0:00.00 [nfsiod 2]
53 ?? IL 0:00.00 [nfsiod 3]
54 ?? DL 0:00.02 [schedcpu]
55 ?? DL 0:00.75 [md0]
77 ?? DL 0:03.48 [md1]
98 ?? DL 0:00.59 [md2]
116 ?? DL 0:00.02 [md3]
137 ?? DL 0:00.56 [md4]
158 ?? DL 0:00.15 [md5]
179 ?? DL 0:00.00 [md6]
215 ?? DL 0:00.03 [md7]
225 ?? DL 0:00.03 [md8]
1052 ?? DL 0:00.00 [jsr_kkcm]
1337 ?? SL 0:00.09 [bcmTX]
1338 ?? SL 0:00.10 [bcmXGS3AsyncTX]
1339 ?? SL 0:03.22 [bcmLINK.0]
1344 ?? Is 0:00.00 /usr/sbin/cron
1496 ?? S 0:00.00 /sbin/watchdog -t-1
1497 ?? S 0:00.05 /usr/libexec/bslockd -mp -N
1498 ?? S 0:00.01 /usr/sbin/tnetd -N
1500 ?? R 0:05.17 /usr/sbin/chassisd -N
1501 ?? S 0:00.04 /usr/sbin/alarmd -N
1502 ?? I 0:00.39 /usr/sbin/craftd -N
1503 ?? S 0:00.08 /usr/sbin/mgd -N
1506 ?? I 0:00.05 /usr/sbin/inetd -N
1507 ?? I 0:00.00 /usr/sbin/tnp.snptd -N
1508 ?? I 0:00.00 /usr/sbin/tnp.snptc -N
1510 ?? S 0:00.01 /usr/sbin/smartd -N
1514 ?? I 0:00.07 /usr/sbin/jcsd -N
1515 ?? S 0:00.17 /usr/sbin/idpd -N
1516 ?? I 0:00.00 /usr/libexec/getty Pc ttyv0
2591 ?? DL 0:00.01 [peer proxy]
2592 ?? DL 0:00.01 [peer proxy]
2593 ?? DL 0:00.01 [peer proxy]
2597 ?? DL 0:00.00 [peer proxy]
3192 ?? S 0:00.01 /usr/sbin/irsd -N
3193 ?? S 0:00.05 /usr/sbin/snmpd -N
3194 ?? S 0:00.02 /sbin/dcd -N
3195 ?? S 0:00.01 /usr/sbin/pfed -N
3196 ?? S 0:00.01 /usr/sbin/mib2d -N
3197 ?? S 0:00.02 /usr/sbin/dfwd -N
3198 ?? S 0:00.13 /usr/sbin/ksyncd -N
3228 ?? Ss 0:00.00 mgd: (mgd) (root) (mgd)
3229 ?? R 0:00.00 /bin/ps -ax
1138 d0- S 0:00.00 /usr/sbin/usbd -N
1156 d0- S 0:00.42 /usr/sbin/eventd -N -r -s -A
1517 d0 Is+ 0:00.00 /usr/libexec/getty std.9600 ttyd0
...

```

show system processes sfc (TX Matrix Plus Router)

```

user@host> show system processes sfc 0
sfc0-re0:

```

```

-----
PID  TT  STAT    TIME COMMAND
  0  ??  WLS    0:00.00 [swapper]
  1  ??  SLs    0:00.18 /packages/mnt/jbase/sbin/init --
  2  ??  DL     0:00.20 [g_event]

```

```

3 ?? DL 0:00.39 [g_up]
4 ?? DL 0:00.32 [g_down]
5 ?? DL 0:00.00 [thread taskq]
6 ?? DL 0:00.09 [kqueue taskq]
7 ?? DL 0:00.01 [pagedaemon]
8 ?? DL 0:00.00 [vmdaemon]
9 ?? DL 0:06.63 [pagezero]
10 ?? DL 0:00.00 [ktrace]
11 ?? RL 312:09.00 [idle]
12 ?? WL 0:11.07 [swi2: net]
13 ?? WL 0:27.70 [swi7: clock sio]
14 ?? WL 0:00.00 [swi6: vm]
15 ?? DL 0:03.03 [yarrow]
16 ?? WL 0:00.00 [swi9: +]
17 ?? WL 0:00.00 [swi8: +]
18 ?? WL 0:00.00 [swi5: cambio]
19 ?? WL 0:00.00 [swi9: task queue]
20 ?? WL 0:11.46 [irq16: uhci0 uhci*]
21 ?? DL 0:00.00 [usb0]
22 ?? DL 0:00.00 [usbtask]
23 ?? WL 0:39.63 [irq17: uhci1 uhci*]
24 ?? DL 0:00.00 [usb1]
25 ?? WL 0:00.00 [irq18: uhci2 uhci*]
26 ?? DL 0:00.84 [usb2]
27 ?? DL 0:00.00 [usb3]
28 ?? DL 0:00.00 [usb4]
29 ?? DL 0:00.00 [usb5]
30 ?? DL 0:00.73 [usb6]
31 ?? DL 0:00.00 [usb7]
32 ?? WL 0:00.00 [irq14: ata0]
33 ?? WL 0:00.00 [irq15: ata1]
34 ?? WL 0:00.00 [irq1: atkbd0]
35 ?? WL 0:00.00 [swi0: sio]
36 ?? WL 0:00.00 [irq11: isab0]
37 ?? WL 0:00.00 [swi3: ip6opt ipopt]
38 ?? WL 0:00.00 [swi4: ip6mismatch+]
39 ?? WL 0:00.00 [swi1: ipfwd]
40 ?? DL 0:00.02 [bufdaemon]
41 ?? DL 0:00.02 [vn1ru]
42 ?? DL 0:00.39 [syncer]
43 ?? DL 0:00.05 [softdepflush]
44 ?? DL 0:00.00 [netdaemon]
45 ?? DL 0:00.02 [vmuncachedaemon]
46 ?? DL 0:00.00 [if_pic_listen]
47 ?? DL 0:00.35 [vmkmemdaemon]
48 ?? DL 0:00.00 [cb_poll]
49 ?? DL 0:00.06 [if_pfe_listen]
50 ?? DL 0:00.00 [scs_housekeeping]
51 ?? IL 0:00.00 [kern_dump_proc]
52 ?? IL 0:00.00 [nfsiod 0]
53 ?? IL 0:00.00 [nfsiod 1]
54 ?? IL 0:00.00 [nfsiod 2]
55 ?? IL 0:00.00 [nfsiod 3]
56 ?? DL 0:00.37 [schedcpu]
57 ?? DL 0:00.56 [md0]
79 ?? DL 0:02.58 [md1]
100 ?? DL 0:00.03 [md2]
118 ?? DL 0:00.01 [md3]
139 ?? DL 0:00.95 [md4]
160 ?? DL 0:00.12 [md5]
181 ?? DL 0:00.00 [md6]

```

```

217 ?? DL 0:00.02 [md7]
227 ?? DL 0:00.05 [md8]
1341 ?? SL 0:01.35 [bcmTX]
1342 ?? SL 0:01.69 [bcmXGS3AsyncTX]
1343 ?? SL 0:41.57 [bcmLINK.0]
1345 ?? SL 0:33.97 [bcmLINK.1]
1350 ?? Is 0:00.01 /usr/sbin/cron
1502 ?? S 0:00.01 /sbin/watchdog -t-1
1503 ?? S 0:00.86 /usr/libexec/bslockd -mp -N
1504 ?? I 0:00.01 /usr/sbin/tnetd -N
1507 ?? S 0:01.32 /usr/sbin/alarmd -N
1508 ?? S 0:14.54 /usr/sbin/craftd -N
1509 ?? S 0:01.20 /usr/sbin/mgd -N
1512 ?? S 0:00.05 /usr/sbin/inetd -N
1513 ?? S 0:00.10 /usr/sbin/tnp.sntpd -N
1517 ?? S 0:00.11 /usr/sbin/smartd -N
1525 ?? S 0:01.11 /usr/sbin/idpd -N
1526 ?? S 0:01.43 /usr/sbin/license-check -U -M -p 10 -i 10
1527 ?? I 0:00.01 /usr/libexec/getty Pc ttyv0
1616 ?? DL 0:00.30 [peer proxy]
1617 ?? DL 0:00.32 [peer proxy]
1618 ?? DL 0:00.34 [peer proxy]
1619 ?? DL 0:00.30 [peer proxy]
2391 ?? Is 0:00.01 telnetd
7331 ?? Ss 0:00.03 telnetd
9538 ?? DL 0:01.16 [jsr_kkcm]
9613 ?? DL 0:00.18 [peer proxy]
23781 ?? Ss 0:00.01 telnetd
23926 ?? Ss 0:00.03 mgd: (mgd) (regress)/dev/ttyp2 (mgd)
36867 ?? S 0:03.14 /usr/sbin/rpd -N
36874 ?? S 0:00.08 /usr/sbin/lmpd
36876 ?? S 0:00.17 /usr/sbin/lacpd -N
36877 ?? S 0:00.15 /usr/sbin/bfdd -N
36878 ?? S 0:05.05 /usr/sbin/ppmd -N
36907 ?? S 0:26.63 /usr/sbin/chassisd -N
37775 ?? S 0:00.01 /usr/sbin/bdbrepd -N
45727 ?? S 0:00.02 /usr/sbin/xntpd -j -N -g (ntpd)
45729 ?? S 0:00.40 /usr/sbin/l2ald -N
45730 ?? S< 0:00.13 /usr/sbin/apsd -N
45731 ?? SN 0:00.10 /usr/sbin/sampled -N
45732 ?? S 0:00.03 /usr/sbin/ilmid -N
45733 ?? S 0:00.09 /usr/sbin/rmopd -N
45734 ?? S 0:00.31 /usr/sbin/cosd
45735 ?? I 0:00.00 /usr/sbin/rtspd -N
45736 ?? S 0:00.06 /usr/sbin/fsad -N
45737 ?? S 0:00.05 /usr/sbin/rdd -N
45738 ?? S 0:00.10 /usr/sbin/pppd -N
45739 ?? S 0:00.05 /usr/sbin/dfcd -N
45740 ?? S 0:00.08 /usr/sbin/lfmd -N
45741 ?? S 0:00.01 /usr/sbin/mpiisoamd -N
45742 ?? I 0:00.01 /usr/sbin/sendd -N
45743 ?? S 0:00.08 /usr/sbin/appidd -N
45744 ?? S 0:00.05 /usr/sbin/mspd -N
45745 ?? S 0:00.27 /usr/sbin/jdiameterd -N
45746 ?? S 0:00.10 /usr/sbin/pfed -N
45747 ?? S 0:00.19 /usr/sbin/lpdfd -N
45748 ?? S 0:00.64 /sbin/dcd -N
45750 ?? S 0:00.46 /usr/sbin/mib2d -N
45751 ?? S 0:00.16 /usr/sbin/dfwd -N
45752 ?? S 0:00.15 /usr/sbin/irsd -N
45764 ?? S 0:20.60 /usr/sbin/snmpd -N

```



```

56481 ?? Ss 0:00.02 telnetd
56548 ?? Rs 0:00.19 mgd: (mgd) (regress)/dev/tty0 (mgd)
56577 ?? Ss 0:00.00 mgd: (mgd) (root) (mgd)
56578 ?? R 0:00.00 /bin/ps -ax
1142 d0- S 0:00.01 /usr/sbin/usbd -N
1160 d0- S 0:29.71 /usr/sbin/eventd -N -r -s -A
6527 d0 Is+ 0:00.00 /usr/libexec/getty std.9600 ttyd0
56482 p0 Is 0:00.00 login [pam] (login)
56483 p0 S 0:00.01 -csh (csh)
56547 p0 S+ 0:00.02 cli
2392 p1 Is 0:00.00 login [pam] (login)
2393 p1 I 0:00.00 -csh (csh)
2394 p1 I 0:00.00 su -
2395 p1 I+ 0:00.01 -su (csh)
23782 p2 Is 0:00.00 login [pam] (login)
23881 p2 I 0:00.00 -csh (csh)
23925 p2 S+ 0:00.03 cli
7332 p3 Is 0:00.00 login [pam] (login)
7333 p3 I 0:00.00 -csh (csh)
23780 p3 S+ 0:00.02 telnet aj

```

show system processes lcc wide (TX Matrix Plus Routing Matrix)

```

user@host> show system processes lcc 2 wide
lcc2-re0:

```

PID	TT	STAT	TIME	PROVIDER	COMMAND
0	??	Wls	0:00.00	(null)	[swapper]
1	??	ILs	0:00.19		/packages/mnt/jbase/sbin/init --
2	??	DL	0:00.02		[g_event]
3	??	DL	0:00.19		[g_up]
4	??	DL	0:00.13		[g_down]
5	??	DL	0:00.00		[thread taskq]
6	??	DL	0:00.00		[kqueue taskq]
7	??	DL	0:00.00		[pagedaemon]
8	??	DL	0:00.00		[vmdaemon]
9	??	DL	0:01.77		[pagezero]
10	??	DL	0:00.00		[ktrace]
11	??	RL	20:33.81		[idle]
12	??	WL	0:00.38		[swi2: net]
13	??	WL	0:01.43		[swi7: clock sio]
14	??	WL	0:00.00		[swi6: vm]
15	??	DL	0:00.14		[yarrow]
16	??	WL	0:00.00		[swi9: +]
17	??	WL	0:00.00		[swi8: +]
18	??	WL	0:00.00		[swi5: cambio]
19	??	WL	0:00.00		[swi9: task queue]
20	??	WL	0:03.18		[irq10: bcm0 uhci1*]
21	??	WL	0:00.03		[irq11: cb0 uhci0+*]
22	??	DL	0:00.00		[usb0]
23	??	DL	0:00.00		[usbtask]
24	??	DL	0:00.00		[usb1]
25	??	DL	0:00.06		[usb2]
26	??	DL	0:00.00		[usb3]
27	??	DL	0:00.00		[usb4]
28	??	DL	0:00.00		[usb5]
29	??	DL	0:00.05		[usb6]
30	??	DL	0:00.00		[usb7]
31	??	WL	0:00.00		[irq14: ata0]
32	??	WL	0:00.00		[irq15: ata1]
33	??	WL	0:00.00		[irq1: atkbd0]

34	??	WL	0:00.00	[swi0: sio]
35	??	WL	0:00.00	[swi3: ip6opt ipopt]
36	??	WL	0:00.00	[swi4: ip6mismatch+]
37	??	WL	0:00.00	[swi1: ipfwd]
38	??	DL	0:00.00	[bufdaemon]
39	??	DL	0:00.00	[vn1ru]
40	??	DL	0:00.02	[syncer]
41	??	DL	0:00.01	[softdepflush]
42	??	DL	0:00.00	[netdaemon]
43	??	DL	0:00.00	[vmuncachedaemon]
44	??	DL	0:00.00	[if_pic_listen]
45	??	DL	0:00.03	[vmkmemdaemon]
46	??	DL	0:00.01	[cb_poll]
47	??	DL	0:00.00	[if_pfe_listen]
48	??	DL	0:00.00	[scs_housekeeping]
49	??	IL	0:00.00	[kern_dump_proc]
50	??	IL	0:00.00	[nfsiod 0]
51	??	IL	0:00.00	[nfsiod 1]
52	??	IL	0:00.00	[nfsiod 2]
53	??	IL	0:00.00	[nfsiod 3]
54	??	DL	0:00.02	[schedcpu]
55	??	DL	0:00.75	[md0]
77	??	DL	0:03.84	[md1]
98	??	DL	0:00.59	[md2]
116	??	DL	0:00.02	[md3]
137	??	DL	0:00.72	[md4]
158	??	DL	0:00.15	[md5]
179	??	DL	0:00.00	[md6]
215	??	DL	0:00.03	[md7]
225	??	DL	0:00.03	[md8]
1052	??	DL	0:00.00	[jsr_kkcm]
1337	??	SL	0:00.11	[bcmTX]
1338	??	SL	0:00.12	[bcmXGS3AsyncTX]
1339	??	SL	0:03.82	[bcmLINK.0]
1344	??	Is	0:00.00	/usr/sbin/cron
1496	??	I	0:00.00	/sbin/watchdog -t-1
1497	??	S	0:00.06	/usr/libexec/bslockd -mp -N
1498	??	I	0:00.01	/usr/sbin/tnetd -N
1500	??	S	0:09.93	/usr/sbin/chassisd -N
1501	??	S	0:00.05	/usr/sbin/alarmd -N
1502	??	I	0:00.39	/usr/sbin/craftd -N
1503	??	S	0:00.09	/usr/sbin/mgd -N
1506	??	I	0:00.05	/usr/sbin/inetd -N
1507	??	I	0:00.00	/usr/sbin/tnp.sntpd -N
1508	??	I	0:00.00	/usr/sbin/tnp.sntpc -N
1510	??	S	0:00.01	/usr/sbin/smartd -N
1514	??	I	0:00.07	/usr/sbin/jcsd -N
1515	??	S	0:00.17	/usr/sbin/idpd -N
1516	??	I	0:00.00	/usr/libexec/getty Pc ttyv0
2591	??	DL	0:00.01	[peer proxy]
2592	??	DL	0:00.01	[peer proxy]
2593	??	DL	0:00.01	[peer proxy]
2597	??	DL	0:00.01	[peer proxy]
3192	??	S	0:00.02	/usr/sbin/irsd -N
3193	??	S	0:00.05	/usr/sbin/snmpd -N
3194	??	S	0:00.04	/sbin/dcd -N
3195	??	I	0:00.01	/usr/sbin/pfed -N
3196	??	S	0:00.02	/usr/sbin/mib2d -N
3197	??	I	0:00.03	/usr/sbin/dfwd -N
3198	??	S	0:00.15	/usr/sbin/ksyncd -N
3559	??	Ss	0:00.00	mgd: (mgd) (root) (mgd)

```

3560 ?? R      0:00.00      /bin/ps -ax -Jpww
1138 d0- S      0:00.00      /usr/sbin/usbd -N
1156 d0- S      0:00.50      /usr/sbin/eventd -N -r -s -A
1517 d0 Is+    0:00.00      /usr/libexec/getty std.9600 ttyd0

```

show system processes (QFX Series and OCX Series)

```
user@switch> show system processes
```

```

PID  TT  STAT      TIME COMMAND
  0  ??  Wls -2341043:-31.01 [swapper]
  1  ??  SLs  0:01.34 /packages/mnt/jbase/sbin/init --
  2  ??  DL    2:48.31 [g_event]
  3  ??  DL    1:47.44 [g_up]
  4  ??  DL    1:37.82 [g_down]
  5  ??  DL    0:00.00 [kdm_tcp_poller]
  6  ??  DL    0:00.00 [thread taskq]
  7  ??  DL    0:04.86 [kqueue taskq]
  9  ??  DL    0:03.94 [pagedaemon]
 10  ??  DL    0:00.00 [ktrace]
 11  ??  RL    0:00.00 [idle: cpu31]
 12  ??  RL    0:00.00 [idle: cpu30]
 13  ??  RL    0:00.00 [idle: cpu29]
 14  ??  RL    0:00.00 [idle: cpu28]
 15  ??  RL    0:00.00 [idle: cpu27]
 16  ??  RL    0:00.00 [idle: cpu26]
 17  ??  RL    0:00.00 [idle: cpu25]
 18  ??  RL    0:00.00 [idle: cpu24]
 19  ??  RL    0:00.00 [idle: cpu23]
 20  ??  RL    0:00.00 [idle: cpu22]
 21  ??  RL    0:00.00 [idle: cpu21]
 22  ??  RL    0:00.00 [idle: cpu20]
 23  ??  RL    0:00.00 [idle: cpu19]
 24  ??  RL    0:00.00 [idle: cpu18]
 25  ??  RL    0:00.00 [idle: cpu17]
 26  ??  RL    0:00.00 [idle: cpu16]
 27  ??  RL    0:00.00 [idle: cpu15]
 28  ??  RL    0:00.00 [idle: cpu14]
 29  ??  RL    0:00.00 [idle: cpu13]
 30  ??  RL    0:00.00 [idle: cpu12]
 31  ??  RL    0:00.00 [idle: cpu11]
 32  ??  RL    0:00.00 [idle: cpu10]
 33  ??  RL    0:00.00 [idle: cpu9]
 34  ??  RL 18184:07.25 [idle: cpu8]
 35  ??  RL    0:00.00 [idle: cpu7]
 36  ??  RL 17862:11.31 [idle: cpu6]
 37  ??  RL 19343:45.16 [idle: cpu5]
 38  ??  RL 5192:38.30 [idle: cpu4]
 39  ??  RL    0:00.00 [idle: cpu3]
 40  ??  RL 19278:02.24 [idle: cpu2]
 41  ??  RL 19291:00.72 [idle: cpu1]
 42  ??  RL 18910:31.21 [idle: cpu0]
 43  ??  WL   19:03.74 [swi2: net]
 44  ??  WL 261:43.82 [swi7: clock sio]
 45  ??  WL    0:00.00 [swi6: vm]
 46  ??  DL    2:18.57 [yarrow]
 47  ??  WL    0:00.00 [swi9: +]
 48  ??  WL    0:00.00 [swi8: +]
 49  ??  WL    0:12.36 [swi5: cambio]
 50  ??  WL    0:00.00 [swi9: task queue]
 51  ??  WL    0:00.00 [swi0: sio]
 52  ??  WL    0:32.40 [irq39: ehci0]

```

```

53 ?? DL      0:00.21 [usb0]
54 ?? DL      0:00.00 [usbtask]
55 ?? WL      0:00.00 [irq22: xlr_lbus0]
56 ?? WL      0:00.00 [irq38: xlr_lbus0]
57 ?? WL      0:00.00 [swi3: ip6opt ipopt]
58 ?? WL      0:00.00 [swi4: ip6mismatch+]
59 ?? WL      0:00.00 [swi1: ipfwd]
60 ?? DL      0:18.65 [pagezero]
61 ?? DL      0:18.59 [bufdaemon]
62 ?? DL      1:10.44 [vnlr_u_mem]
63 ?? DL      1:51.66 [syncer]
64 ?? DL      0:20.22 [vnlr_u]
65 ?? DL      0:40.48 [softdepflush]
66 ?? DL      0:00.00 [netdaemon]
67 ?? DL      20:47.67 [vmkmemdaemon]
68 ?? DL      0:00.00 [if_pfe_listen]
69 ?? SL      0:02.80 [kdm_checkkcore]
70 ?? SL      0:03.34 [kdm_savekcore]
71 ?? SL      0:04.31 [kdm_livekcore]
72 ?? SL      0:06.14 [kdm_logger]
73 ?? SL      0:04.31 [kdm_kdb]
74 ?? SL      0:00.02 [devrt_kernel_thread]
75 ?? DL      0:21.54 [vmuncachedaemon]
76 ?? DL      0:00.00 [if_pic_listen0]
77 ?? SL      0:00.00 [nfsiod 0]
78 ?? SL      0:00.00 [nfsiod 1]
79 ?? SL      0:00.00 [nfsiod 2]
80 ?? SL      0:00.00 [nfsiod 3]
81 ?? WL      5:59.98 [irq13: +]
82 ?? RL      105:06.81 [pkt_sender: cpu0]
83 ?? DL      0:03.62 [md0]
95 ?? DL      0:37.04 [md1]
115 ?? DL     0:06.01 [md2]
135 ?? DL     0:00.75 [md3]
155 ?? DL     0:21.17 [md4]
175 ?? DL     0:01.90 [md5]
195 ?? DL     0:06.26 [md6]
231 ?? DL     0:00.01 [md7]
755 ?? Ss     0:04.17 /usr/sbin/cron
847 ?? S      0:00.10 /usr/sbin/tinetd -N
849 ?? S      0:06.82 /usr/sbin/mgd -N
850 ?? S      0:00.32 /usr/sbin/inetd -N
852 ?? S      1:05.34 /usr/sbin/dhcpd -N
853 ?? S      0:00.18 /usr/sbin/inetd -p /var/run/inetd_4.pid -N -JU __juni
855 ?? L      1181:02.21 /usr/sbin/dc-pfe -N (pafxpc)
857 ?? S      17:55.86 /usr/sbin/vccpd -N
896 ?? S      93:43.45 /usr/sbin/chassism -N
953 ?? S      0:02.89 /sbin/watchdog -t-1
954 ?? S      3:34.00 /sbin/dcd -N
955 ?? S      10:30.13 /usr/sbin/chassisd -N
956 ?? DL     0:00.21 [peer proxy]
957 ?? S      4:07.43 /usr/sbin/alarmd -N
958 ?? S      0:31.69 /usr/sbin/craftd -N
959 ?? S      0:55.16 /usr/sbin/mib2d -N
960 ?? S      3:40.64 /usr/sbin/rpd -N
961 ?? S      0:00.03 /usr/sbin/tnp.sntpd -N
962 ?? S      0:51.94 /usr/sbin/pfed -N
963 ?? S      0:47.31 /usr/sbin/rmopd -N
964 ?? S      0:33.65 /usr/sbin/cosd
965 ?? S      1:48.41 /usr/sbin/ppmd -N
966 ?? S      0:07.18 /usr/sbin/dfwd -N

```

```

967 ?? S      1:02.56 /usr/sbin/bfdd -N
968 ?? S      0:00.63 /usr/sbin/rdd -N
969 ?? S      0:40.61 /usr/sbin/dfcd -N
971 ?? S      0:07.81 /usr/sbin/bdbrepd -N
972 ?? S      0:00.28 /usr/sbin/sendd -N
973 ?? S      1:37.69 /usr/sbin/xntpd -j -N -g -JU __juniper_private4__ (nt
974 ?? S      5:56.28 /usr/sbin/snmpd -N -JU __juniper_private4__
975 ?? S      16:46.82 /usr/sbin/jdiameterd -N
976 ?? S      2:34.13 /usr/sbin/eswd -N
977 ?? S      1:03.05 /usr/sbin/sflowd -N
978 ?? S      0:22.30 /usr/sbin/fcd -N
979 ?? S      1:07.01 /usr/sbin/vccpdf -N
982 ?? S      0:25.25 /usr/sbin/mcsnoopd -N
983 ?? S      3:45.68 /usr/sbin/rpdf -N
1043 ?? S      0:37.87 /usr/sbin/lacpd -N
1048 ?? DL     0:01.29 [peer proxy]
1111 ?? WL     0:00.00 [swi2: FMNITHRD+]
1112 ?? DL     0:00.03 [peer proxy]
12816 ?? S     15:35.32 /usr/sbin/sfid -N
30893 ?? Ss    0:00.65 sshd: tlewis@tty0 (sshd)
30897 ?? Ss    0:00.15 mgd: (mgd) (tlewis)/dev/tty0 (mgd)
30905 ?? Ss    0:00.64 sshd: tlewis@tty1 (sshd)
30909 ?? Ss    0:00.15 mgd: (mgd) (tlewis)/dev/tty1 (mgd)
30910 ?? Ss    0:01.26 sshd: tcheng@tty2 (sshd)
30914 ?? Ss    0:00.80 mgd: (mgd) (tcheng)/dev/tty2 (mgd)
30937 ?? R      0:00.03 /bin/ps -ax
    661 d0- S    0:21.24 /usr/sbin/eventd -N -r -s -A
    860 d0 Ss+   0:00.07 /usr/libexec/getty std.9600 ttyd0
30896 p0 Ss+   0:00.55 -cli (cli)
30908 p1 Ss+   0:00.50 -cli (cli)
30913 p2 Ss+   0:00.85 -cli (cli)

```

show system reboot

List of Syntax	Syntax on page 1126 Syntax (EX Series Switches) on page 1126 Syntax (TX Matrix Router) on page 1126 Syntax (TX Matrix Plus Router) on page 1126 Syntax (MX Series Router) on page 1126 Syntax (QFX Series and OCX Series) on page 1126
Syntax	show system reboot <both-routing-engines>
Syntax (EX Series Switches)	show system reboot <all-members> <both-routing-engines> <local> <member <i>member-id</i> >
Syntax (TX Matrix Router)	show system reboot <all-chassis all-lcc lcc <i>number</i> scc> <both-routing-engines>
Syntax (TX Matrix Plus Router)	show system reboot <all-chassis all-lcc lcc <i>number</i> sfc <i>number</i> > <both-routing-engines>
Syntax (MX Series Router)	show system reboot <all-members> <both-routing-engines> <local> <member <i>member-id</i> >
Syntax (QFX Series and OCX Series)	show system reboot <both-routing-engines> <infrastructure <i>name</i> > <interconnect-device <i>name</i> > <node-device <i>name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Display pending system reboots or halts.
Options	none —Display pending reboots or halts on the active Routing Engine. all-chassis —(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display halt or reboot request information for all the T640 routers in the chassis that are connected to the TX Matrix router. On a TX Matrix router, display

halt or reboot request information for all the T1600 or T4000 routers in the chassis that are connected to the TX Matrix Plus router.

all-members—(EX4200 switches and MX Series routers only) (Optional) Display halt or reboot request information for all members of the Virtual Chassis configuration.

all-lcc—(TX Matrix routers and TX Matrix Plus router only) (Optional) On a TX Matrix router, display system halt or reboot request information for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, display halt or reboot request information for all connected T1600 or T4000 LCCs.

both-routing-engines—(Systems with multiple Routing Engines) (Optional) Display halt or reboot request information on both Routing Engines.

infrastructure *name*—(QFabric systems only) (Optional) Display reboot request information on the fabric manager Routing Engines and fabric control Routing Engines.

interconnect-device *name*—(QFabric systems only) (Optional) Display reboot request information on the Interconnect device.

lcc *number*—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display halt or reboot request information for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display halt or reboot request information for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(EX4200 switches and MX Series routers only) (Optional) Display halt or reboot request information for the local Virtual Chassis member.

member *member-id*—(EX4200 switches and MX Series routers only) (Optional) Display halt or reboot request information for the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

node-group *name*—(QFabric systems only) (Optional) Display reboot request information on the Node group.

scc—(TX Matrix router only) (Optional) Display halt or reboot request information for the TX Matrix router (or switch-card chassis).

sfc—(TX Matrix Plus router only) (Optional) Display halt or reboot request information for the TX Matrix Plus router.

Additional Information By default, when you issue the **show system reboot** command on a TX Matrix or TX Matrix Plus master Routing Engine, the command is broadcast to all the T640 (in a routing matrix based on the TX Matrix router) or T1600 (in a routing matrix based on the TX Matrix Plus router) master Routing Engines connected to it. Likewise, if you issue the same command on the TX Matrix or TX Matrix Plus backup Routing Engine, the command is broadcast to all the T640 (in a routing matrix based on the TX Matrix router) or T1600 (in a routing matrix based on the TX Matrix Plus router) backup Routing Engines that are connected to it.

Required Privilege Level maintenance

Related Documentation

- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

List of Sample Output [show system reboot on page 1128](#)
[show system reboot all-lcc \(TX Matrix Router\) on page 1128](#)
[show system reboot sfc \(TX Matrix Plus Router\) on page 1128](#)
[show system reboot \(QFX3500 Switch\) on page 1128](#)

Sample Output

show system reboot

```
user@host> show system reboot
reboot requested by root at Wed Feb 10 17:40:46 1999
[process id 17885]
```

show system reboot all-lcc (TX Matrix Router)

```
user@host> show system reboot all-lcc
lcc0-re0:
```

```
-----
No shutdown/reboot scheduled.
```

```
lcc2-re0:
```

```
-----
No shutdown/reboot scheduled.
```

show system reboot sfc (TX Matrix Plus Router)

```
user@host> show system sfc 0
No shutdown/reboot scheduled.
```

show system reboot (QFX3500 Switch)

```
user@switch> show system reboot
No shutdown/reboot scheduled.
```


show system resource-cleanup processes

Syntax	show system resource-cleanup processes <detail> <pid <i>number</i> > <process-name <i>name</i> >
Release Information	Command introduced in Junos OS Release 9.3. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Display the list of processes that have been registered for resource cleanup services.
Options	<p>detail—(Optional) Display the list of processes that have been registered for resource cleanup services, along with the resources that have been requested for cleanup.</p> <p>pid <i>number</i>—(Optional) Display a process that has been registered for resource cleanup services by specifying the Process Identifier number.</p> <p>process-name <i>name</i>—(Optional) Display a process that has been registered for resource cleanup services by name of the process.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> <i>resource-cleanup</i> <i>traceoptions (Resource Cleanup)</i>
List of Sample Output	show system resource-cleanup processes on page 1129 show system resource-cleanup processes detail on page 1130
Output Fields	For a description of the output fields, see Table 77 . Output fields are listed in the approximate order in which they appear.

Table 77: show system resource-cleanup processes Output Fields

Field Name	Field Description
PID	Process ID, a number that identifies a process.
Process name	String that identifies the process.
Resources to clean	Resources that have been registered to be cleaned up.

Sample Output

show system resource-cleanup processes

```
user@host> show system resource-cleanup processes
PID      Process name      Resources to clean
420      jnx-exampld      GENCFG, SYSV shared memory
```

show system resource-cleanup processes detail

```
user@host> show system resource-cleanup processes detail
PID      Process name      Resources to clean
420      jnx-exampld        GENCFG blob major ID 0x8000, minor ID 0x0000
          SYSV shared memory ID 65536, key 1108955839
          SYSV shared memory ID 65537, key 1108955837
```

show system services service-deployment

Syntax	show system services service-deployment
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Display information about a Session and Resource Control (SRC) client.
Options	This command has no options.
Required Privilege Level	system view
List of Sample Output	show system services service-deployment on page 1131
Output Fields	Table 78 lists the output fields for the show system services service-deployment command. Output fields are listed in the approximate order in which they appear.

Table 78: show system services service-deployment Output Fields

Field Name	Field Description
PDT Keepalive settings	Configured PDT keepalive interval, in seconds.
Keepalives sent	Number of keepalives sent.
Notifications sent	Number of notifications sent.
Last update from peer	Time at which the last update from a peer was received.

Sample Output

show system services service-deployment

```
user@host> show system services service-deployment
Connected to 192.4.4.4 port 10288 since 2004-05-03 11:04:34 PDT Keepalive settings:
Interval 15 seconds Keepalives sent: 750 Notifications sent: 0 Last update from
peer: 00:00:06 ago
```

show system software

List of Syntax	Syntax on page 1132 Syntax (EX Series Switches) on page 1132 Syntax (TX Matrix Router) on page 1132 Syntax (TX Matrix Plus Router) on page 1132 Syntax (QFX Series) on page 1132
Syntax	show system software <detail>
Syntax (EX Series Switches)	show system software <all-members> <detail> <local> <member <i>member-id</i> >
Syntax (TX Matrix Router)	show system software <all-chassis all-lcc lcc <i>number</i> scc> <detail>
Syntax (TX Matrix Plus Router)	show system software <all-chassis all-lcc lcc <i>number</i> sfc <i>number</i> > <detail>
Syntax (QFX Series)	show system software <detail> <infrastructure <i>name</i> > <interconnect-device <i>name</i> > <node-group <i>name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Display the Junos OS extensions loaded on your router or switch.
Options	none —Display standard information about all loaded Junos OS extensions. all-chassis —(TX Matrix routers and TX Matrix Plus routers only) (Optional) Display system software information for all the T640 routers (TX Matrix Router) or all the routers (TX Matrix Plus Router) in the chassis. all-lcc —(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system software information for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, display system software information for all connected T1600 or T4000 LCCs. all-members —(EX4200 switches only) (Optional) Display the system software running on all members of the Virtual Chassis configuration.

detail—(Optional) Display detailed information about available Junos OS extensions.

infrastructure *name*—(QFabric systems only) (Optional) Display the system software running on the fabric control Routing Engine and the fabric manager Routing Engine.

interconnect-device *name*—(QFabric systems only) (Optional) Display the system software running on the Interconnect device.

lcc *number*—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system software information for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display system software information for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(EX4200 switches only) (Optional) Display the system software running on the local Virtual Chassis member.

member *member-id*—(EX4200 switches only) (Optional) Display the system software running on the specified member of the Virtual Chassis configuration. Replace *member-id* with a value from 0 through 9.

node-group *name*—(QFabric systems only) (Optional) Display the system software running on the Node group.

scc—(Routing matrix only) (Optional) Display the system software running on a TX Matrix router (or switch-card chassis).

sfc—(TX Matrix Plus routers only) (Optional) Display system software information for the TX Matrix Plus router.

Required Privilege Level maintenance

Related Documentation • [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

List of Sample Output [show system software on page 1134](#)
[show system software \(TX Matrix Plus Router\) on page 1134](#)
[show system software \(QFX Series\) on page 1138](#)

Output Fields When you enter this command, you are provided a list of Junos OS packages installed on the router and their corresponding Junos OS release number.

Sample Output

show system software

```
user@host> show system software
Information for jbase:

Comment:
JUNOS Base OS Software Suite [7.2R1.7]

Information for jcrypto:

Comment:
JUNOS Crypto Software Suite [7.2R1.7]
Information for jdocs:

Comment:
JUNOS Online Documentation [7.2R1.7]

Information for jkernel:

Comment:
JUNOS Kernel Software Suite [7.2R1.7]

Information for jpfe:

Comment:
JUNOS Packet Forwarding Engine Support (M20/M40) [7.2R1.7]

Information for jroute:

Comment:
JUNOS Routing Software Suite [7.2R1.7]

Information for junos:

Comment:
JUNOS Base OS boot [7.2R1.7]
```

show system software (TX Matrix Plus Router)

```
user@host> show system software
sfc0-re0:
-----
Information for jbase:

Comment:
JUNOS Base OS Software Suite [9.6-20090515.0]

Information for jcrypto:
```

Comment:
JUNOS Crypto Software Suite [9.6-20090515.0]

Information for jdocs:

Comment:
JUNOS Online Documentation [9.6-20090515.0]
Information for jkernel:

Comment:
JUNOS Kernel Software Suite [9.6-20090515.0]

Information for jpfe:

Comment:
JUNOS Packet Forwarding Engine Support (T-Series) [9.6-20090515.0]

Information for jpfe-common:

Comment:
JUNOS Packet Forwarding Engine Support (M/T Common) [9.6-20090515.0]

Information for jroute:Comment:
JUNOS Routing Software Suite [9.6-20090515.0]

Information for jservices-aacl:

Comment:
JUNOS Services ACL Container package [9.6-20090515.0]

Information for jservices-appid:

Comment:
JUNOS AppId Services [9.6-20090515.0]

Information for jservices-bgf:

Comment:
JUNOS Border Gateway Function package [9.6-20090515.0]
Information for jservices-idp:

Comment:
JUNOS IDP Services [9.6-20090515.0]

Information for jservices-llpdf:

Comment:

JUNOS Services LL-PDF Container package [9.6-20090515.0]

Information for jservices-sfw:

Comment:

JUNOS Services Stateful Firewall [9.6-20090515.0]

Information for jservices-voice:

Comment:

JUNOS Voice Services Container package [9.6-20090515.0]

Information for junos:

Comment:

JUNOS Base OS boot [9.6-20090515.0]

...

lcc0-re0:

Information for jbase:

Comment:

JUNOS Base OS Software Suite [9.6-20090515.0]

Information for jcrypto:

Comment:

JUNOS Crypto Software Suite [9.6-20090515.0]

Information for jdocs:

Comment:

JUNOS Online Documentation [9.6-20090515.0]

Information for jkernel:

Comment:

JUNOS Kernel Software Suite [9.6-20090515.0]

Information for jpfe:

Comment:

JUNOS Packet Forwarding Engine Support (T-Series) [9.6-20090515.0]

Information for jpfe-common:

Comment:
JUNOS Packet Forwarding Engine Support (M/T Common) [9.6-20090515.0]

Information for jroute:

Comment:
JUNOS Routing Software Suite [9.6-20090515.0]

Information for jservices-aacl:

Comment:
JUNOS Services ACL Container package [9.6-20090515.0]

Information for jservices-appid:

Comment:
JUNOS AppId Services [9.6-20090515.0]

Information for jservices-bgf:

Comment:
JUNOS Border Gateway Function package [9.6-20090515.0]

Information for jservices-idp:

Comment:
JUNOS IDP Services [9.6-20090515.0]

Information for jservices-llpdf:

Comment:
JUNOS Services LL-PDF Container package [9.6-20090515.0]

Information for jservices-sfw:

Comment:
JUNOS Services Stateful Firewall [9.6-20090515.0]

Information for jservices-voice:

Comment:
JUNOS Voice Services Container package [9.6-20090515.0]

Information for junos:

Comment:
JUNOS Base OS boot [9.6-20090515.0]

lcc1-re0:

Information for jbase:

Comment:
JUNOS Base OS Software Suite [9.6-20090515.0]

Information for jcrypto:

Comment:
JUNOS Crypto Software Suite [9.6-20090515.0]
...

show system software (QFX Series)

user@switch> **show system software**
Information for jbase:

Comment:
JUNOS Base OS Software Suite [11.3-20110730.0]

Information for jcrypto:

Comment:
JUNOS Crypto Software Suite [11.3-20110730.0]

Information for jdocs:

Comment:
JUNOS Online Documentation [11.3-20110730.0]

Information for jkernel:

Comment:
JUNOS Kernel Software Suite [11.3-20110730.0]

Information for jpfe:

Comment:
JUNOS Packet Forwarding Engine Support (QFX) [11.3-20110730.0]

Information for jroute:

Comment:

JUNOS Routing Software Suite [11.3-20110730.0]

Information for jswitch:

Comment:

JUNOS Enterprise Software Suite [11.3-20110730.0]

Information for junos:

Comment:

JUNOS Base OS boot [11.3-20110730.0]

Information for jweb:

Comment:

JUNOS Web Management [11.3-20110730.0]

show system statistics

List of Syntax	Syntax on page 1140 Syntax (EX Series Switches) on page 1140 Syntax (TX Matrix Router) on page 1140 Syntax (TX Matrix Plus Router) on page 1140 Syntax (MX Series Router) on page 1140 Syntax (QFX Series) on page 1140
Syntax	show system statistics
Syntax (EX Series Switches)	show system statistics <all-members> <local> <member <i>member-id</i> >
Syntax (TX Matrix Router)	show system statistics <all-chassis all-lcc lcc <i>number</i> scc>
Syntax (TX Matrix Plus Router)	show system statistics <all-chassis all-lcc lcc <i>number</i> sfc <i>number</i> >
Syntax (MX Series Router)	show system statistics <all-members> <local> <member <i>member-id</i> >
Syntax (QFX Series)	show system statistics
Release Information	Command introduced before JUNOS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc option introduced for the TX Matrix Plus router in JUNOS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Display system-wide protocol-related statistics.
Options	none —Display system statistics for all the following protocols: <ul style="list-style-type: none">• arp—Address Resolution Protocol• bridge—IEEE 802.1 Bridging• clns—Connectionless Network Service• esis—End System-to-Intermediate System• ethoamcfm—Ethernet OAM protocol for connectivity fault management• ethoamlfm—Ethernet OAM protocol for link fault management• icmp—Internet Control Message Protocol• icmp6—Internet Control Message Protocol version 6

- **igmp**—Internet Group Management Protocol
- **ip**—Internet Protocol version 4
- **ip6**—Internet Protocol version 6
- **jsr**—Juniper Socket Replication
- **mpls**—Multiprotocol Label Switching
- **rdp**—Reliable Datagram Protocol
- **tcp**—Transmission Control Protocol
- **tnp**—Trivial Network Protocol
- **ttp**—TNP Tunneling Protocol
- **tudp**—Trivial User Datagram Protocol
- **udp**—User Datagram Protocol
- **vpls**—Virtual Private LAN Service

all-chassis—(TX Matrix and TX Matrix Plus routers only) (Optional) Display system statistics for a protocol for all the routers in the chassis.

all-lcc—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system statistics for a protocol for all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for a protocol for all routers (line-card chassis) connected to the TX Matrix Plus router

all-members—(EX4200 switches and MX Series routers only) (Optional) Display system statistics for a protocol for all members of the Virtual Chassis configuration.

lcc number—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system statistics for a protocol for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for a protocol for a specific router that is connected to the TX Matrix Plus router. Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(EX4200 switches and MX Series routers only) (Optional) Display system statistics for a protocol for the local Virtual Chassis member.

member *member-id*—(EX4200 switches and MX Series routers only) (Optional) Display system statistics for a protocol for the specified member of the Virtual Chassis configuration. For EX4200 switches, replace ***member-id*** with a value from 0 through 9. For an MX Series Virtual Chassis, replace ***member-id*** with a value of 0 or 1.

scc—(TX Matrix routers only) (Optional) Display system statistics for a protocol for the TX Matrix router (or switch-card chassis).

sfc *number*—(TX Matrix Plus routers only) (Optional) Display system statistics for a protocol for the TX Matrix Plus router (or switch-fabric chassis). Replace ***number*** with 0.

Additional Information By default, when you issue the **show system statistics** command on a TX Matrix or TX Matrix Plus master Routing Engine, the command is broadcast to all the T640 (in a routing matrix based on the TX Matrix router) or T1600 (in a routing matrix based on the TX Matrix Plus router) master Routing Engines connected to it. Likewise, if you issue the same command on the TX Matrix or TX Matrix Plus backup Routing Engine, the command is broadcast to all the T640 (in a routing matrix based on the TX Matrix router) or T1600 (in a routing matrix based on the TX Matrix Plus router) backup Routing Engines that are connected to it.

Required Privilege Level view

List of Sample Output [show system statistics on page 1142](#)
[show system statistics \(EX Series Switches\) on page 1152](#)
[show system statistics \(TX Matrix Router\) on page 1161](#)
[show system statistics \(QFX Series\) on page 1167](#)

Sample Output

show system statistics

```
user@host> show system statistics
ip:
    3682087 total packets received
    0 bad header checksums
    0 with size smaller than minimum
    0 with data size < data length
    0 with header length < data size
    0 with data length < header length
    0 with incorrect version number
    0 packets destined to dead next hop
    0 fragments received
    0 fragments dropped (dup or out of space)
    0 fragments dropped (queue overflow)
    0 fragments dropped after timeout
    0 fragments dropped due to over limit
    0 packets reassembled ok
    3664774 packets for this host
    17316 packets for unknown/unsupported protocol
    0 packets forwarded
    0 packets not forwardable
    0 redirects sent
    6528 packets sent from this host
```

```

0 packets sent with fabricated ip header
0 output packets dropped due to no bufs
0 output packets discarded due to no route
0 output datagrams fragmented
0 fragments created
0 datagrams that can't be fragmented
0 packets with bad options
1123 packets with options handled without error
0 strict source and record route options
0 loose source and record route options
0 record route options
0 timestamp options
0 timestamp and address options
0 timestamp and prespecified address options
0 option packets dropped due to rate limit
1123 router alert options
0 multicast packets dropped (no iflist)
0 packets dropped (src and int don't match)
icmp:
0 drops due to rate limit
0 calls to icmp_error
0 errors not generated because old message was icmp
Output histogram:
    echo reply: 75
0 messages with bad code fields
0 messages less than the minimum length
0 messages with bad checksum
0 messages with bad source address
0 messages with bad length
0 echo drops with broadcast or multicast destination address
0 timestamp drops with broadcast or multicast destination address
Input histogram:
    echo: 75
    router advertisement: 130
75 message responses generated
tcp:
3844 packets sent
    3618 data packets (1055596 bytes)
    0 data packets (0 bytes) retransmitted
    0 resends initiated by MTU discovery
    205 ack-only packets (148 packets delayed)
    0 URG only packets
    0 window probe packets
    0 window update packets
    1079 control packets
5815 packets received
    3377 acks (for 1055657 bytes)
    24 duplicate acks
    0 acks for unsent data
    2655 packets (15004 bytes) received in-sequence
    1 completely duplicate packet (0 bytes)
    0 old duplicate packets
    0 packets with some dup. data (0 bytes duped)
    0 out-of-order packets (0 bytes)
    0 packets (0 bytes) of data after window
    0 window probes
    7 window update packets
    0 packets received after close
    0 discarded for bad checksums
    0 discarded for bad header offset fields
    0 discarded because packet too short

```

```
1 connection request
32 connection accepts
0 bad connection attempts
0 listen queue overflows
33 connections established (including accepts)
30 connections closed (including 0 drops)
    27 connections updated cached RTT on close
    27 connections updated cached RTT variance on close
    0 connections updated cached ssthresh on close
0 embryonic connections dropped
3374 segments updated rtt (of 3220 attempts)
0 retransmit timeouts
    0 connections dropped by rexmit timeout
0 persist timeouts
    0 connections dropped by persist timeout
344 keepalive timeouts
    0 keepalive probes sent
    0 connections dropped by keepalive
1096 correct ACK header predictions
1314 correct data packet header predictions
32 syncache entries added
    0 retransmitted
    0 dupsyn
    0 dropped
    32 completed
    0 bucket overflow
    0 cache overflow
    0 reset
    0 stale
    0 aborted
    0 badack
    0 unreach
    0 zone failures
0 cookies sent
0 cookies received
0 ACKs sent in response to in-window but not exact RSTs
0 ACKs sent in response to in-window SYNs on established connections
0 rcv packets dropped by TCP due to bad address
0 out-of-sequence segment drops due to insufficient memory
1058 RST packets
0 ICMP packets ignored by TCP
0 send packets dropped by TCP due to auth errors
0 rcv packets dropped by TCP due to auth errors

udp:
3658884 datagrams received
0 with incomplete header
0 with bad data length field
0 with bad checksum
3657342 dropped due to no socket
3657342 broadcast/multicast datagrams dropped due to no socket
0 dropped due to full socket buffers
0 not for hashed pcb
4291311496 delivered
1551 datagrams output

ipsec:
0 inbound packets processed successfully
0 inbound packets violated process security policy
0 inbound packets with no SA available
0 invalid inbound packets
0 inbound packets failed due to insufficient memory
0 inbound packets failed getting SPI
```



```

0 inbound packets failed on AH replay check
0 inbound packets failed on ESP replay check
0 inbound AH packets considered authentic
0 inbound AH packets failed on authentication
0 inbound ESP packets considered authentic
0 inbound ESP packets failed on authentication
0 outbound packets processed successfully
0 outbound packets violated process security policy
0 outbound packets with no SA available
0 invalid outbound packets
0 outbound packets failed due to insufficient memory
0 outbound packets with no route

igmp:
17186 messages received
0 messages received with too few bytes
0 messages received with bad checksum
0 membership queries received
0 membership queries received with invalid field(s)
0 membership reports received
0 membership reports received with invalid field(s)
0 membership reports received for groups to which we belong
0 membership reports sent

arp:
44181302 datagrams received
2 ARP requests received
2028 ARP replies received
3156 resolution requests received
0 unrestricted proxy requests
0 received proxy requests
0 proxy requests not proxied
0 with bogus interface
787 with incorrect length
712 for non-IP protocol
0 with unsupported op code
0 with bad protocol address length
0 with bad hardware address length
0 with multicast source address
7611 with multicast target address
0 with my own hardware address
14241699 for an address not on the interface
0 with a broadcast source address
0 with source address duplicate to mine
29929250 which were not for me
0 packets discarded waiting for resolution
6 packets sent after waiting for resolution
17812 ARP requests sent
2 ARP replies sent
0 requests for memory denied
0 requests dropped on entry
0 requests dropped during retry

ip6:
0 total packets received
0 with size smaller than minimum
0 with data size < data length
0 with bad options
0 with incorrect version number
0 fragments received
0 fragments dropped (dup or out of space)
0 fragments dropped after timeout
0 fragments that exceeded limit
0 packets reassembled ok

```

```
0 packets for this host
0 packets forwarded
0 packets not forwardable
0 redirects sent
0 packets sent from this host
0 packets sent with fabricated ip header
0 output packets dropped due to no bufs, etc.
0 output packets discarded due to no route
0 output datagrams fragmented
0 fragments created
0 datagrams that can't be fragmented
0 packets that violated scope rules
0 multicast packets which we don't join
Mbuf statistics:
0 packets whose headers are not continuous
0 tunneling packets that can't find gif
0 packets discarded due to too many headers
0 failures of source address selection
0 forward cache hit
0 forward cache miss
0 packets destined to dead next hop
0 option packets dropped due to rate limit
0 packets dropped (src and int don't match)
0 packets dropped due to bad protocol

icmp6:
0 calls to icmp_error
0 errors not generated because old message was icmp error or so
0 errors not generated because rate limitation
0 messages with bad code fields
0 messages < minimum length
0 bad checksums
0 messages with bad length
Histogram of error messages to be generated:
    0 no route
    0 administratively prohibited
    0 beyond scope
    0 address unreachable
    0 port unreachable
    0 packet too big
    0 time exceed transit
    0 time exceed reassembly
    0 erroneous header field
    0 unrecognized next header
    0 unrecognized option
    0 redirect
    0 unknown
0 message responses generated
0 messages with too many ND options

ipsec6:
0 inbound packets processed successfully
0 inbound packets violated process security policy
0 inbound packets with no SA available
0 invalid inbound packets
0 inbound packets failed due to insufficient memory
0 inbound packets failed getting SPI
0 inbound packets failed on AH replay check
0 inbound packets failed on ESP replay check
0 inbound AH packets considered authentic
0 inbound AH packets failed on authentication
0 inbound ESP packets considered authentic
0 inbound ESP packets failed on authentication
```

```

0 outbound packets processed successfully
0 outbound packets violated process security policy
0 outbound packets with no SA available
0 invalid outbound packets
0 outbound packets failed due to insufficient memory
0 outbound packets with no route
c1nl:
0 total packets received
0 packets delivered
0 too small
0 bad header length
0 bad checksum
0 bad version
0 unknown or unsupported protocol
0 bogus sdl size
0 no free memory in socket buffer
0 send packets discarded
0 sbappend failure
0 mcopy failure
0 address fields were not reasonable
0 segment information forgotten
0 forwarded packets
0 total packets sent
0 output packets discarded
0 non-forwarded packets
0 packets fragmented
0 fragments sent
0 fragments discarded
0 fragments timed out
0 fragmentation prohibited
0 packets reconstructed
0 packets destined to dead nexthop
0 packets discarded due to no route
0 Error pdu rate drops
0 ER pdu generation failure
esis:
0 total pkts received
0 total packets consumed by protocol
0 pdus received with bad checksum
0 pdus received with bad version number
0 pdus received with bad type field
0 short pdus received
0 bogus sdl size
0 bad header length
0 unknown or unsupported protocol
0 no free memory in socket buffer
0 send packets discarded
0 sbappend failure
0 mcopy failure
0 ISO family not configured
tnp:
146776365 unicast packets received
0 broadcast packets received
0 fragmented packets received
0 hello packets dropped
0 fragments dropped
0 fragment reassembly queue flushes
0 hello packets received
0 control packets received
49681642 rdp packets received
337175 udp packets received

```

```
96757548 tunnel packets received
0 input packets discarded with no protocol
98397591 unicast packets sent
0 broadcast packets sent
0 fragmented packets sent
0 hello packets dropped
0 fragments dropped
0 hello packets sent
0 control packets sent
49681642 rdp packets sent
337175 udp packets sent
48378774 tunnel packets sent
0 packets sent with unknown protocol

rdp:
49681642 input packets
0 discards for bad checksum
0 discards bad sequence number
0 refused connections
2031964 acks received
0 dropped due to full socket buffers
49692 retransmits
49681642 output packets
24815968 acks sent
28 connects
0 closes
22783990 keepalives received
22783990 keepalives sent

tudp:
337175 datagrams received
0 with incomplete header
0 with bad data length field
0 with bad checksum
0 dropped due to no socket
0 broadcast/multicast datagrams dropped due to no socket
0 dropped due to full socket buffers
337175 delivered
337175 datagrams output

ttp:
398749 packets sent
0 packets sent while unconnected
0 packets sent while interface down
0 packets sent couldn't get buffer
0 packets sent couldn't find neighbor
44696687 L2 packets received
0 unknown L3 packets received
3682087 IPv4 L3 packets received
0 MPLS L3 packets received
0 MPLS->IPv4 L3 packets received
0 IPv4->MPLS L3 packets received
0 IPv6 L3 packets received
0 ARP L3 packets received
0 CLNP L3 packets received
0 TNP L3 packets received
0 NULL L3 packets received
0 cyclotron cycle L3 packets received
0 cyclotron send L3 packets received
0 packets received while unconnected
0 packets received from unknown ifl
0 input packets couldn't get buffer
0 input packets with bad type
0 input packets with discard type
```

```

0 Input packets with too many tlvs
0 Input packets with bad tlv header
70633 Input packets with bad tlv type
68877 Input packets dropped based on tlv result
0 input packets for which rt lookup is bypassed

mpls:
0 total mpls packets received
0 packets forwarded
0 packets dropped
0 with header too small
0 after tagging, can't fit link MTU
0 with IPv4 explicit NULL tag
0 with IPv4 explicit NULL cksum errors
0 with router alert tag
0 lsp ping packets (ttl-expired/router alert)
0 with ttl expired
0 with tag encoding error
0 packets discarded, no route

jsr:
  Handle-inf:o
    0 Handles in use
    0 Handles allocated so far
    0 Handles freed so far
    0 Handles in delayed free state
  IHA:
    0 IHA invalid subtype messages
    0 IHA invalid length messages
    0 IHA invalid version messages
    0 IHA too short messages
    0 IHA invalid dst handle messages
    0 IHA invalid src handle messages
    0 IHA unmatched src handle messages
    0 IHA invalid messages for primary
    0 IHA invalid messages for secondary
    0 IHA invalid messages for current state
    0 IHA messages sent for subtype init
    0 IHA messages rcvd for subytpe init
    0 IHA messages sent for subtype init
    0 IHA messages rcvd for subytpe init
    0 IHA messages sent for subtype init
    0 IHA messages rcvd for subytpe init
    0 IHA messages sent for subtype init
    0 IHA messages rcvd for subytpe init
    0 IHA messages sent for subtype init
    0 IHA messages rcvd for subytpe init
    0 IHA message timeouts
    0 IHA socket unreplicate messages
  SDRL:
    0 SDRL socket teardowns
    0 SDRL socket teardown failures
    0 SDRL socket unreplicates
    0 SDRL socket unreplicate failures
    0 SDRL external timeouts
    0 SDRL internal timeouts
    0 SDRL ipc messages sent
    0 SDRL ipc send failures
    0 SDRL ipc messages rcvd
    0 SDRL ipc messages rcvd
    0 SDRL primary replication messages sent
    0 SDRL primary replication message send failures
    0 SDRL primary ack messages received

```

```
0 SDRL primary ack message receive failures
0 SDRL primary sock replication inits
0 SDRL primary sock replication init failures
0 SDRL primary throttle remove messages
0 SDRL primary throttle remove failures
0 SDRL primary init handshake messages
0 SDRL primary init handshake failures
0 SDRL secondary replication messages received
0 SDRL secondary replication message receive failures
0 SDRL secondary replication acks sent
0 SDRL secondary replication ack send failures
0 SDRL secondary sock splits
0 SDRL secondary sock split failures
0 SDRL secondary sock merges
0 SDRL secondary sock merge failures
0 SDRL secondary sockets closed
0 SDRL secondary rcv snoop fd close failures
0 SDRL secondary snd snoop fd close failures
0 SDRL secondary init handshake messages
0 SDRL secondary init handshake failures
PRL:
0 PRL packets enqueued
0 PRL packets failed to enqueue
0 PRL packets dequeued
0 PRL packets failed to dequeue
0 PRL queue entry allocations
0 PRL queue entry frees
0 calls to layer 4 input handlers
0 failed calls to layer 4 input handlers
0 PRL queue drains
0 PRL replication timeouts
0 PRL replication messages sent
0 PRL replication message send failures
0 PRL acknowledgment messages sent
0 PRL acknowledgement message send failures
0 PRL replication messages received
0 PRL replication message receive failures
0 PRL acknowledgement messages received
0 PRL acknowledgement receive failures
0 PRL messages with bad IPC type
0 PRL messages with no handler
2 PRL global state initializations
1 PRL global state cleanups
0 PRL per-socket state creations
0 PRL per-socket state creation failures
0 PRL per-socket state cleanups
0 PRL socket closes
0 PRL socket merges
0 PRL socket unreplicates
0 PRL primary socket replication initializations
0 PRL secondary socket replication initializations
0 PRL primary socket replication activations
0 PRL secondary socket replication activations
0 packets received from peers
0 PRL packets receive operations from peer failed
0 PRL buffer pullup failures
0 new pkts dropped on secondary socket
PSRM:
0 PSRM replication timeouts
0 PSRM replication messages sent
0 PSRM replication message send failures
```

```

0 PSRM acknowledgment messages sent
0 PSRM acknowledgement message send failures
0 PSRM flow control messages sent
0 PSRM flow control message send failures
0 PSRM replication messages received
0 PSRM replication message receive failures
0 PSRM acknowledgment messages received
0 PSRM acknowledgment message receive failures
0 PSRM flow control messages received
0 PSRM flow control message receive failures
0 SRM messages with bad IPC type
0 PSRM messages with no handler
2 PSRM global state initializations
1 PSRM global state cleanups
0 PSRM per-socket state creations
0 PSRM per-socket state creation failures
0 PSRM per-socket state cleanups
0 PSRM socket closes
0 PSRM socket merges
0 PSRM socket unreplicates
0 PSRM primary socket replication initializations
0 psrm-secondary-socket-replication-initializations
0 PSRM primary socket replication activations
0 secondary socket replication activations
0 PSRM tcpcb updates
0 PSRM buffer pullup failures
73 PSRM tcp timestamp msg rcv counters
0 PSRM tcp timestamp msg rcv failures
0 PSRM tcp timestamp msg send counters
0 PSRM tcp timestamp msg send failures

```

TCP:

```

0 TCP out-of-order packets on JSR sockets

```

vpls:

```

0 total packets received
0 with size smaller than minimum
0 with incorrect version number
0 packets for this host
0 packets with no logical interface
0 packets with no family
0 packets with no route table
0 packets with no auxiliary table
0 packets with no corefacing entry
0 packets with no CE-facing entry
0 mac route learning requests
0 mac routes learnt
0 requests to learn an existing route
0 learning requests while learning disabled on interface
0 learning requests over capacity
0 mac routes moved
0 requests to move static route
0 mac route aging requests
0 mac routes aged
0 bogus address in aging requests
0 requests to age static route
0 requests to re-ageout aged route
0 requests involving multiple peer FEs
0 aging acks from PFE
0 aging non-acks from PFE
0 aging requests timed out waiting on FEs
0 aging requests over max-rate
0 errors finding peer FEs

```

show system statistics (EX Series Switches)

```
user@host> show system statistics
```

```
Tcp:
```

```
571779 packets sent
  21517 data packets (1797102 bytes)
  2 data packets retransmitted (20 bytes)
  0 resends initiated by MTU discovery
  3708 ack only packets (531 packets delayed)
  0 URG only packets
  1 window probe packets
  1 window update packets
  1093063 control packets
1132541 packets received
  20961 acks(for 1796102 bytes)
  5861 duplicate acks
  0 acks for unsent data
  19556 packets received in-sequence(232079 bytes)
  3018 completely duplicate packets(0 bytes)
  0 old duplicate packets
  4 packets with some duplicate data(4 bytes duped)
  2 out-of-order packets(2 bytes)
  0 packets of data after window(0 bytes)
  0 window probes
  39 window update packets
  0 packets received after close
  0 discarded for bad checksums
  0 discarded for bad header offset fields
  0 discarded because packet too short
546519 connection requests
78 connection accepts
0 bad connection attempts
0 listen queue overflows
100 connections established (including accepts)
546596 connections closed (including 6 drops)
  47 connections updated cached RTT on close
  47 connections updated cached RTT variance on close
  0 connections updated cached ssthresh on close
546497 embryonic connections dropped
20453 segments updated rtt(of 566914 attempts)
2 retransmit timeouts
  0 connections dropped by retransmit timeout
0 persist timeouts
  0 connections dropped by persist timeout
3028 keepalive timeouts
  3027 keepalive probes sent
  1 connections dropped by keepalive
7515 correct ACK header predictions
12258 correct data packet header predictions
78 syncache entries added
  0 retransmitted
  0 dupsyn
  4 dropped
  78 completed
  0 bucket overflow
  0 cache overflow
  0 reset
  0 stale
  0 aborted
  0 badack
  0 unreachable
```



```

    0 zone failures
    0 cookies sent
    0 cookies received
    1 SACK recovery episodes
    1 segment retransmits in SACK recovery episodes
    1 byte retransmits in SACK recovery episodes
    71 SACK options (SACK blocks) received
    1 SACK options (SACK blocks) sent
    0 SACK scoreboard overflow
    0 ACKs sent in response to in-window but not exact RSTs
    0 ACKs sent in response to in-window SYNs on established connections
    0 rcv packets dropped by TCP due to bad address
    0 out-of-sequence segment drops due to insufficient memory
    546544 RST packets
    0 ICMP packets ignored by TCP
    0 send packets dropped by TCP due to auth errors
    0 rcv packets dropped by TCP due to auth errors
    0 outgoing segments dropped due to policing
udp:
    147 datagrams received
    0 with incomplete header
    0 with bad data length field
    0 with bad checksum
    9 dropped due to no socket
    0 broadcast/multicast datagrams dropped due to no socket
    0 dropped due to full socket buffers
    0 not for hashed pcb
    138 delivered
    0 datagrams output
ip:
    73704 total packets received
    0 bad header checksums
    0 with size smaller than minimum
    0 with data size < data length
    0 with header length < data size
    0 with data length < header length
    0 with incorrect version number
    0 packets destined to dead next hop
    0 fragments received
    0 fragments dropped (dup or out of space)
    0 fragments dropped (queue overflow)
    0 fragments dropped after timeout
    0 fragments dropped due to over limit
    0 packets reassembled ok
    1133057 packets for this host
    0 packets for unknown/unsupported protocol
    40146 packets forwarded
    0 packets not forwardable
    40146 redirects sent
    1121700 packets sent from this host
    0 packets sent with fabricated ip header
    0 output packets dropped due to no bufs
    0 output packets discarded due to no route
    0 output datagrams fragmented
    0 fragments created
    0 datagrams that can't be fragmented
    0 packets with bad options
    0 packets with options handled without error
    0 strict source and record route options
    0 loose source and record route options
    0 record route options

```

```
0 timestamp options
0 timestamp and address options
0 timestamp and prespecified address options
0 option packets dropped due to rate limit
0 router alert options
0 multicast packets dropped (no iflist)
0 packets dropped (src and int don't match)
0 transit re packets dropped on mgmt i/f
0 packets used first nexthop in ecmp unilist
0 incoming ttpoip packets received
0 incoming ttpoip packets dropped
0 outgoing TTPoIP packets sent
0 outgoing TTPoIP packets dropped

icmp:
0 drops due to rate limit
9 calls to icmp_error
0 errors not generated because old message was icmp
Output histogram:
    295 echo reply
    9 destination unreachable
0 messages with bad code fields
0 messages less than the minimum length
0 messages with bad checksum
0 messages with bad source address
0 messages with bad length
0 echo drops with broadcast or multicast destination address
0 timestamp drops with broadcast or multicast destination address
Input histogram:
    295 echo
295 message responses generated

igmp:
0 messages received
0 messages received with too few bytes
0 messages received with bad checksum
0 membership queries received
0 membership queries received with invalid fields
0 membership reports received
0 membership reports received with invalid fields
0 membership reports received for groups to which we belong
0 Membership reports sent

raw_if:
0 RAW packets transmitted
0 PPPOE packets transmitted
0 ISDN packets transmitted
0 DIALER packets transmitted
0 PPP packets transmitted to pppd
0 PPP packets transmitted to jppd
0 IGMPv2 packets transmitted
13 output drops due to tx error
0 MPU packets transmitted
0 PPPOE packets received
0 ISDN packets received
0 DIALER packets received
0 PPP packets received from pppd
0 MPU packets received
0 PPP packets received from jppd
0 IGMPv2 packets received
0 Input drops due to bogus protocol
0 input drops due to no mbufs available
0 input drops due to no space in socket
0 input drops due to no socket
```

```

arp:
    186413 datagrams received
    88 ARP requests received
    88 ARP replies received
    0 resolution request received
    0 unrestricted proxy requests
    0 restricted proxy requests
    0 received proxy requests
    0 proxy requests not proxied
    0 restricted proxy requests not proxied
    0 datagrams with bogus interface
    0 datagrams with incorrect length
    0 datagrams for non-IP protocol
    0 datagrams with unsupported op code
    0 datagrams with bad protocol address length
    0 datagrams with bad hardware address length
    0 datagrams with multicast source address
    0 datagrams with multicast source address
    0 datagrams with my own hardware address
    164 datagrams for an address not on the interface
    0 datagrams with a broadcast source address
    0 datagrams with source address duplicate to mine
    186065 datagrams which were not for me
    0 packets discarded waiting for resolution
    0 packets sent after waiting for resolution
    50 ARP requests sent
    88 ARP replies sent
    0 requests for memory denied
    0 requests dropped on entry
    0 requests dropped during retry
    0 requests dropped due to interface deletion
    0 requests on unnumbered interfaces
    0 new requests on unnumbered interfaces
    0 replies for from unnumbered interfaces
    0 requests on unnumbered interface with non-subnetted donor
    0 replies from unnumbered interface with non-subnetted donor

ip6:
    0 total packets received
    0 packets with size smaller than minimum
    0 packets with data size < data length
    0 packets with bad options
    0 packets with incorrect version number
    0 fragments received
    0 fragments dropped (dup or out of space)
    0 fragments dropped after timeout
    0 fragments that exceeded limit
    0 packets reassembled ok
    0 packets for this host
    0 packets forwarded
    0 packets not forwardable
    0 redirects sent
    0 packets sent from this host
    0 packets sent with fabricated ip header
    0 output packets dropped due to no bufs, etc.
    0 output datagrams fragmented
    0 fragments created
    0 datagrams that can't be fragmented
    0 packets that violated scope rules
    0 multicast packets which we don't join
    0 packets whose headers are not continuous
    0 tunneling packets that can't find gif

```

```
0 packets discarded due to too may headers
0 failures of source address selection
0 forward cache hit
0 forward cache miss
0 Packets destined to dead next hop
0 option packets dropped due to rate limit
0 Packets dropped (src and int don't match)
0 packets dropped due to bad protocol
0 transit re packet(null) dropped on mgmt i/f

icmp6:
0 Calls to icmp_error
0 Errors not generated because old message was icmp error
0 Errors not generated because rate limitation
0 Messages with bad code fields
0 Messages < minimum length
0 Bad checksums
0 Messages with bad length
    0 No route
    0 Administratively prohibited
    0 Beyond scope
    0 Address unreachable
    0 Port unreachable
    0 packet too big
    0 Time exceed transit
    0 Time exceed reassembly
    0 Erroneous header field
    0 Unrecognized next header
    0 Unrecognized option
    0 redirect
    0 Unknown
0 Message responses generated
0 Messages with too many ND options

pfkey:
0 Requests sent from userland
0 Bytes sent from userland
histogram by message type:
    0 reserved
    0 dump
0 Messages with invalid length field
0 Messages with invalid version field
0 Messages with invalid message type field
0 Messages too short
0 Messages with memory allocation failure
0 Messages with duplicate extension
0 Messages with invalid extension type
0 Messages with invalid sa type
0 Messages with invalid address extension
0 Requests sent to userland
0 Bytes sent to userland
histogram by message type:
    0 reserved
    0 dump
0 Messages toward single socket
0 Messages toward all sockets
0 Messages toward registered sockets
0 Messages with memory allocation failure

c1n1:
0 Total packets received
0 Packets delivered
0 Too small packets
0 Packets with bad header length
```

```

0 Packets with bad checksum
0 Bad version packets
0 Unknown or unsupported protocol packets
0 Packets with bogus sdl size
0 No free memory in socket buffer
0 Send packets discarded
0 Sbappend failure
0 Mcopy failure
0 Address fields were not reasonable
0 Segment information forgotten
0 Forwarded packets
0 Total packets sent
0 Output packets discarded
0 Non-forwarded packets
0 Packets fragmented
0 Fragments sent
0 Fragments discarded
0 Fragments timed out
0 Fragmentation prohibited
0 Packets reconstructed
0 Packets destined to dead nexthop
0 Packets discarded due to no route
0 Error pdu rate drops
  0 ER pdu generation failure
esis:
0 Total pkts received
0 Total packets consumed by protocol
0 Pdus received with bad checksum
0 Pdus received with bad version number
0 Pdus received with bad type field
0 Short pdus received
0 Pdus with bogus sdl size
0 Pdus with bad header length
0 Pdus with unknown or unsupported protocol
0 No free memory in socket buffer
0 Send packets discarded
0 Sbappend failure
0 Mcopy failure
0 ISO family not configured
tnp:
0 Unicast packets received
0 Broadcast packets received
0 Fragmented packets received
0 Hello packets dropped
0 Fragments dropped
0 Fragment reassembly queue flushes
0 Packets with tnp src address collision received
0 Hello packets received
0 Control packets received
0 Rdp packets received
0 Udp packets received
0 Tunnel packets received
0 Input packets discarded with no protocol
0 Packets of version unspecified received
0 Packets of version 1 received
0 Packets of version 2 received
0 Packets of version 3 received
0 Unicast packets sent
0 Broadcast packets sent
0 Fragmented packets sent
0 Hello packets dropped

```

```
0 Fragments dropped
0 Hello packets sent
0 Control packets sent
0 Rdp packets sent
0 Udp packets sent
0 Tunnel packets sent
0 Packets sent with unknown protocol
0 Packets of version unspecified sent
0 Packets of version 1 sent
0 Packets of version 2 sent
0 Packets of version 3 sent
rdp:
0 Input packets
0 Packets discarded for bad checksum
0 Packets discarded due to bad sequence number
0 Refused connections
0 Acks received
0 Packets dropped due to full socket buffers
0 Retransmits
0 Output packets
0 Acks sent
0 Connects
0 Closes
0 Keepalives received
0 Keepalives sent
tudp:
67 Datagrams received
0 Datagrams with incomplete header
0 Datagrams with bad data length field
0 Datagrams with bad checksum
0 Datagrams dropped due to no socket
0 Broadcast/multicast datagrams dropped due to no socket
0 Datagrams dropped due to full socket buffers
67 Delivered
68 Datagrams output
ttp:
0 Packets sent
0 Packets sent while unconnected
0 Packets sent while interface down
0 Packets sent couldn't get buffer
0 Packets sent couldn't find neighbor
0 L2 packets received
0 Unknown L3 packets received
0 IPv4 L3 packets received
0 MPLS L3 packets received
0 MPLS->IPv4 L3 packets received
0 IPv4->MPLS L3 packets received
0 IPv6 L3 packets received
0 ARP L3 packets received
0 CLNP L3 packets received
0 TNP L3 packets received
0 NULL L3 packets received
0 Cyclotron cycle L3 packets received
0 Cyclotron send L3 packets received
0 Packets received while unconnected
0 Packets received from unknown ifl
0 Input packets couldn't get buffer
0 Input packets with bad type
0 Input packets with discard type
0 Input packets with too many tlvs
0 Input packets with bad tlv header
```

```

70633 Input packets with bad tlv type
68877 Input packets dropped based on tlv result
0 Input packets for which rt lookup is bypassed

mpls:
0 Total MPLS packets received
0 Packets forwarded
0 Packets dropped
0 Packets with header too small
0 After tagging, packets can't fit link MTU
0 Packets with IPv4 explicit NULL tag
0 Packets with IPv4 explicit NULL cksum errors
0 Packets with router alert tag
0 LSP ping packets (ttl-expired/router alert)
0 Packets with ttl expired
0 Packets with tag encoding error
0 Packets discarded due to no route
0 Packets used first nexthop in ecmp unilist

vpls:
0 Total packets received
0 Packets with size smaller than minimum
0 Packets with incorrect version number
0 Packets for this host
0 Packets with no logical interface
0 Packets with no family
0 Packets with no route table
0 Packets with no auxiliary table
0 Packets with no corefacing entry
0 packets with no CE-facing entry
0 MAC route learning requests
0 MAC routes learnt
0 Requests to learn an existing route
0 Learning requests while learning disabled on interface
0 Learning requests over capacity
0 MAC routes moved
0 Requests to move static route
0 MAC route aging requests
0 MAC routes aged
0 Bogus address in aging requests
0 Requests to age static route
0 Requests to re-ageout aged route
0 Requests involving multiple peer FEs
0 Aging acks from PFE
0 Aging non-acks from PFE
0 Aging requests timed out waiting on FEs
0 Aging requests over max-rate
0 Errors finding peer FEs
0 Unsupported platform
0 Packets dropped due to no l3 route table
0 Packets dropped due to no local ifl
0 Packets punted
0 Packets dropped due to no socket

bridge:
Input:
0 packets received
0 packets forwarded
0 packets failed to forward
0 packets dropped
0 packets with vmember lookup failures
0 packets with vlan lookup failures
0 packets with stp state lookup failures
0 packets dropped due to stp blocked/listening

```

```
0 packets dropped due to stp learning
0 packets with src MAC learning failures
0 packets with input control processing failures
Forward:
0 packets sent successfully
0 packets with send failures
0 packets forwarded to l3 interface
0 packets with l3 send failures
0 packets discarded
0 packets with l2ifl store failures
0 packets with ifl mismatch failures
0 packets with packet duplication failures
0 packets with tag lookup failures
0 packets with no route for DMAC
0 packets with no route table
0 packets with no nexthop
0 packets with dead nexthop
0 packets with eof reached error
Learning:
0 MACs learned
0 packets sent to l3 interface
0 packets with l3 send failures
0 packets hit holdq while learning
0 MAC moves
0 packets discarded
0 packets with no route for SMAC
0 packets with no nexthop
0 packets with dead nexthop
0 packets dropped due to no resolve route
0 packets with l3 ifd lookup failures
0 packets with l3 ifl lookup failures
0 packets with l3 invalid rnh
0 packets with no route for SMAC in clone learning
0 packets with no nexthop in clone learning
0 packets with dead nexthop in clone learning
0 packets dropped due to no resolve nh in clone learning
Output:
0 packets forwarded
0 packets failed to forward
0 packets with vmember lookup failures
0 packets with vlan lookup failures
0 packets with input control processing failures
Send:
0 packets sent successfully
0 packets with send failures
0 packets dropped due to interface down
0 packets with dev output failures
0 blocked ifl discards
0 packets with tag lookup failures
0 packets with stp state lookup failures
0 packets with tag insertion failures
0 packets with tag removal failures
Flood:
0 packets flooded
0 flood failures
IGMP:
0 packets sent successfully
0 packets with send failures
0 packets forwarded
0 packets failed to forward
0 packets with mpull failures
```



```

0 packets with vmember lookup failures
0 packets with vlan lookup failures
0 packets with ifl lookup failures
0 packets with tag lookup failures
Misc:
0 packets with size smaller than minimum
0 packets with double tags
0 packets with no ifl
0 packets with no family
0 packets with no route table

```

show system statistics (TX Matrix Router)

```

user@host> show system statistics
sfc0-re0:

```

```

-----
Tcp:
361694 packets sent
    326507 data packets (103237236 bytes)
    2343 data packets retransmitted (2673324 bytes)
    0 resends initiated by MTU discovery
    33857 ack only packets (31613 packets delayed)
    0 URG only packets
    14 window probe packets
    387 window update packets
    1108 control packets
345879 packets received
    298207 acks(for 103141728 bytes)
    438 duplicate acks
    0 acks for unsent data
    204578 packets received in-sequence(13820995 bytes)
    6 completely duplicate packets(18 bytes)
    0 old duplicate packets
    0 packets with some duplicate data(0 bytes duped)
    0 out-of-order packets(0 bytes)
    0 packets of data after window(0 bytes)
    0 window probes
    899 window update packets
    166 packets received after close
    0 discarded for bad checksums
    0 discarded for bad header offset fields
    0 discarded because packet too short
406 connection requests
233 connection accepts
0 bad connection attempts
0 listen queue overflows
616 connections established (including accepts)
911 connections closed (including 41 drops)
    346 connections updated cached RTT on close
    346 connections updated cached RTT variance on close
    200 connections updated cached ssthresh on close
23 embryonic connections dropped
298155 segments updated rtt(of 287216 attempts)
1163 retransmit timeouts
    27 connections dropped by retransmit timeout
0 persist timeouts
    0 connections dropped by persist timeout
5 keepalive timeouts
    5 keepalive probes sent
    0 connections dropped by keepalive
69922 correct ACK header predictions

```

```
34993 correct data packet header predictions
233 syncache entries added
    0 retransmitted
    0 dupsyn
    0 dropped
    233 completed
    0 bucket overflow
    0 cache overflow
    0 reset
    0 stale
    0 aborted
    0 badack
    0 unreach
    0 zone failures
0 cookies sent
0 cookies received
23 SACK recovery episodes
68 segment retransmits in SACK recovery episodes
71542 byte retransmits in SACK recovery episodes
158 SACK options (SACK blocks) received
0 SACK options (SACK blocks) sent
0 SACK scoreboard overflow
0 ACKs sent in response to in-window but not exact RSTs
0 ACKs sent in response to in-window SYNs on established connections
0 rcv packets dropped by TCP due to bad address
0 out-of-sequence segment drops due to insufficient memory
259 RST packets
0 ICMP packets ignored by TCP
0 send packets dropped by TCP due to auth errors
0 rcv packets dropped by TCP due to auth errors
0 outgoing segments dropped due to policing
```

lcc0-re0:

Tcp:

```
346 packets sent
    222 data packets (22894 bytes)
    0 data packets retransmitted (0 bytes)
    0 resends initiated by MTU discovery
    80 ack only packets (12 packets delayed)
    0 URG only packets
    0 window probe packets
    5 window update packets
    42 control packets
358 packets received
    268 acks(for 22939 bytes)
    9 duplicate acks
    0 acks for unsent data
    203 packets received in-sequence(33820 bytes)
    0 completely duplicate packets(0 bytes)
    0 old duplicate packets
    0 packets with some duplicate data(0 bytes duped)
    0 out-of-order packets(0 bytes)
    0 packets of data after window(0 bytes)
    0 window probes
    6 window update packets
    0 packets received after close
    0 discarded for bad checksums
    0 discarded for bad header offset fields
    0 discarded because packet too short
13 connection requests
```

```

18 connection accepts
0 bad connection attempts
0 listen queue overflows
31 connections established (including accepts)
35 connections closed (including 2 drops)
    3 connections updated cached RTT on close
    3 connections updated cached RTT variance on close
    0 connections updated cached ssthresh on close
0 embryonic connections dropped
268 segments updated rtt(of 247 attempts)
0 retransmit timeouts
    0 connections dropped by retransmit timeout
0 persist timeouts
    0 connections dropped by persist timeout
0 keepalive timeouts
    0 keepalive probes sent
    0 connections dropped by keepalive
0 correct ACK header predictions
42 correct data packet header predictions
18 syncache entries added
    0 retransmitted
    0 dupsyn
    0 dropped
    18 completed
    0 bucket overflow
    0 cache overflow
    0 reset
    0 stale
    0 aborted
    0 badack
    0 unreach
    0 zone failures
0 cookies sent
0 cookies received
0 SACK recovery episodes
0 segment retransmits in SACK recovery episodes
0 byte retransmits in SACK recovery episodes
0 SACK options (SACK blocks) received
0 SACK options (SACK blocks) sent
0 SACK scoreboard overflow
0 ACKs sent in response to in-window but not exact RSTs
0 ACKs sent in response to in-window SYNs on established connections
0 rcv packets dropped by TCP due to bad address
0 out-of-sequence segment drops due to insufficient memory
5 RST packets
0 ICMP packets ignored by TCP
0 send packets dropped by TCP due to auth errors
0 rcv packets dropped by TCP due to auth errors
0 outgoing segments dropped due to policing

```

lcc1-re0:

 Tcp:

```

348 packets sent
    223 data packets (22895 bytes)
    0 data packets retransmitted (0 bytes)
    0 resends initiated by MTU discovery
    81 ack only packets (13 packets delayed)
    0 URG only packets
    0 window probe packets
    5 window update packets

```

```
    42 control packets
360 packets received
    269 acks(for 22940 bytes)
    9 duplicate acks
    0 acks for unsent data
    203 packets received in-sequence(33820 bytes)
    0 completely duplicate packets(0 bytes)
    0 old duplicate packets
    0 packets with some duplicate data(0 bytes duped)
    0 out-of-order packets(0 bytes)
    0 packets of data after window(0 bytes)
    0 window probes
    6 window update packets
    0 packets received after close
    0 discarded for bad checksums
    0 discarded for bad header offset fields
    0 discarded because packet too short
13 connection requests
18 connection accepts
0 bad connection attempts
0 listen queue overflows
31 connections established (including accepts)
36 connections closed (including 2 drops)
    3 connections updated cached RTT on close
    3 connections updated cached RTT variance on close
    0 connections updated cached ssthresh on close
0 embryonic connections dropped
269 segments updated rtt(of 248 attempts)
0 retransmit timeouts
    0 connections dropped by retransmit timeout
0 persist timeouts
    0 connections dropped by persist timeout
0 keepalive timeouts
    0 keepalive probes sent
    0 connections dropped by keepalive
0 correct ACK header predictions
43 correct data packet header predictions
18 syncache entries added
    0 retransmitted
    0 dupsyn
    0 dropped
    18 completed
    0 bucket overflow
    0 cache overflow
    0 reset
    0 stale
    0 aborted
    0 badack
    0 unreach
    0 zone failures
0 cookies sent
0 cookies received
0 SACK recovery episodes
0 segment retransmits in SACK recovery episodes
0 byte retransmits in SACK recovery episodes
0 SACK options (SACK blocks) received
0 SACK options (SACK blocks) sent
0 SACK scoreboard overflow
0 ACKs sent in response to in-window but not exact RSTs
0 ACKs sent in response to in-window SYNs on established connections
0 rcv packets dropped by TCP due to bad address
```

```

0 out-of-sequence segment drops due to insufficient memory
5 RST packets
0 ICMP packets ignored by TCP
0 send packets dropped by TCP due to auth errors
0 rcv packets dropped by TCP due to auth errors
0 outgoing segments dropped due to policing

```

lcc2-re0:

Tcp:

```

405 packets sent
    271 data packets (23926 bytes)
    0 data packets retransmitted (0 bytes)
    0 resends initiated by MTU discovery
    86 ack only packets (13 packets delayed)
    0 URG only packets
    0 window probe packets
    5 window update packets
    46 control packets
418 packets received
    321 acks(for 23975 bytes)
    9 duplicate acks
    0 acks for unsent data
    234 packets received in-sequence(34403 bytes)
    0 completely duplicate packets(0 bytes)
    0 old duplicate packets
    0 packets with some duplicate data(0 bytes duped)
    0 out-of-order packets(0 bytes)
    0 packets of data after window(0 bytes)
    0 window probes
    7 window update packets
    0 packets received after close
    0 discarded for bad checksums
    0 discarded for bad header offset fields
    0 discarded because packet too short
15 connection requests
19 connection accepts
0 bad connection attempts
0 listen queue overflows
34 connections established (including accepts)
39 connections closed (including 2 drops)
    4 connections updated cached RTT on close
    4 connections updated cached RTT variance on close
    0 connections updated cached ssthresh on close
0 embryonic connections dropped
321 segments updated rtt(of 299 attempts)
0 retransmit timeouts
    0 connections dropped by retransmit timeout
0 persist timeouts
    0 connections dropped by persist timeout
0 keepalive timeouts
    0 keepalive probes sent
    0 connections dropped by keepalive
0 correct ACK header predictions
48 correct data packet header predictions
19 syncache entries added
    0 retransmitted
    0 dupsyn
    0 dropped
    19 completed
    0 bucket overflow

```

- 0 cache overflow
- 0 reset
- 0 stale
- 0 aborted
- 0 badack
- 0 unreach
- 0 zone failures
- 0 cookies sent
- 0 cookies received
- 0 SACK recovery episodes
- 0 segment retransmits in SACK recovery episodes
- 0 byte retransmits in SACK recovery episodes
- 0 SACK options (SACK blocks) received
- 0 SACK options (SACK blocks) sent
- 0 SACK scoreboard overflow
- 0 ACKs sent in response to in-window but not exact RSTs
- 0 ACKs sent in response to in-window SYNs on established connections
- 0 rcv packets dropped by TCP due to bad address
- 0 out-of-sequence segment drops due to insufficient memory
- 5 RST packets
- 0 ICMP packets ignored by TCP
- 0 send packets dropped by TCP due to auth errors
- 0 rcv packets dropped by TCP due to auth errors
- 0 outgoing segments dropped due to policing

lcc3-re0:

Tcp:

- 346 packets sent
 - 221 data packets (22895 bytes)
 - 0 data packets retransmitted (0 bytes)
 - 0 resends initiated by MTU discovery
 - 81 ack only packets (13 packets delayed)
 - 0 URG only packets
 - 0 window probe packets
 - 5 window update packets
 - 42 control packets
- 360 packets received
 - 267 acks(for 22940 bytes)
 - 9 duplicate acks
 - 0 acks for unsent data
 - 203 packets received in-sequence(33820 bytes)
 - 0 completely duplicate packets(0 bytes)
 - 0 old duplicate packets
 - 0 packets with some duplicate data(0 bytes duped)
 - 0 out-of-order packets(0 bytes)
 - 0 packets of data after window(0 bytes)
 - 0 window probes
 - 6 window update packets
 - 0 packets received after close
 - 0 discarded for bad checksums
 - 0 discarded for bad header offset fields
 - 0 discarded because packet too short
- 13 connection requests
- 18 connection accepts
- 0 bad connection attempts
- 0 listen queue overflows
- 31 connections established (including accepts)
- 35 connections closed (including 2 drops)
 - 3 connections updated cached RTT on close
 - 3 connections updated cached RTT variance on close

```

    0 connections updated cached ssthresh on close
0 embryonic connections dropped
267 segments updated rtt(of 246 attempts)
0 retransmit timeouts
    0 connections dropped by retransmit timeout
0 persist timeouts
    0 connections dropped by persist timeout
0 keepalive timeouts
    0 keepalive probes sent
    0 connections dropped by keepalive
0 correct ACK header predictions
43 correct data packet header predictions
18 syncache entries added
    0 retransmitted
    0 dupsyn
    0 dropped
    18 completed
    0 bucket overflow
    0 cache overflow
    0 reset
    0 stale
    0 aborted
    0 badack
    0 unreach
    0 zone failures
0 cookies sent
0 cookies received
0 SACK recovery episodes
0 segment retransmits in SACK recovery episodes
0 byte retransmits in SACK recovery episodes
0 SACK options (SACK blocks) received
0 SACK options (SACK blocks) sent
0 SACK scoreboard overflow
0 ACKs sent in response to in-window but not exact RSTs
0 ACKs sent in response to in-window SYNs on established connections
0 rcv packets dropped by TCP due to bad address
0 out-of-sequence segment drops due to insufficient memory
5 RST packets
0 ICMP packets ignored by TCP
0 send packets dropped by TCP due to auth errors
0 rcv packets dropped by TCP due to auth errors
0 outgoing segments dropped due to policing

```

show system statistics (QFX Series)

```

user@switch> show system statistics
Tcp:
571779 packets sent
21517 data packets (1797102 bytes)
2 data packets retransmitted (20 bytes)
0 resends initiated by MTU discovery
3708 ack only packets (531 packets delayed)
0 URG only packets
1 window probe packets
1 window update packets
1093063 control packets
1132541 packets received
20961 acks(for 1796102 bytes)
5861 duplicate acks
0 acks for unsent data
19556 packets received in-sequence(232079 bytes)

```

3018 completely duplicate packets(0 bytes)
0 old duplicate packets
4 packets with some duplicate data(4 bytes duped)
2 out-of-order packets(2 bytes)
0 packets of data after window(0 bytes)
0 window probes
39 window update packets
0 packets received after close
0 discarded for bad checksums
0 discarded for bad header offset fields
0 discarded because packet too short
546519 connection requests
78 connection accepts
0 bad connection attempts
0 listen queue overflows
100 connections established (including accepts)
546596 connections closed (including 6 drops)
47 connections updated cached RTT on close
47 connections updated cached RTT variance on close
0 connections updated cached ssthresh on close
546497 embryonic connections dropped
20453 segments updated rtt(of 566914 attempts)
2 retransmit timeouts
0 connections dropped by retransmit timeout
0 persist timeouts
0 connections dropped by persist timeout
3028 keepalive timeouts
3027 keepalive probes sent
1 connections dropped by keepalive
7515 correct ACK header predictions
12258 correct data packet header predictions
78 syncache entries added
0 retransmitted
0 dupsyn
4 dropped
78 completed
0 bucket overflow
0 cache overflow
0 reset
0 stale
0 aborted
0 badack
0 unreach
0 zone failures
0 cookies sent
0 cookies received
1 SACK recovery episodes
1 segment retransmits in SACK recovery episodes
1 byte retransmits in SACK recovery episodes
71 SACK options (SACK blocks) received
1 SACK options (SACK blocks) sent
0 SACK scoreboard overflow
0 ACKs sent in response to in-window but not exact RSTs
0 ACKs sent in response to in-window SYNs on established connections
0 rcv packets dropped by TCP due to bad address
0 out-of-sequence segment drops due to insufficient memory
546544 RST packets
0 ICMP packets ignored by TCP
0 send packets dropped by TCP due to auth errors
0 rcv packets dropped by TCP due to auth errors
0 outgoing segments dropped due to policing


```

udp:
147 datagrams received
0 with incomplete header
0 with bad data length field
0 with bad checksum
9 dropped due to no socket
0 broadcast/multicast datagrams dropped due to no socket
0 dropped due to full socket buffers
0 not for hashed pcb
138 delivered
0 datagrams output
ip:
73704 total packets received
0 bad header checksums
0 with size smaller than minimum
0 with data size < data length
0 with header length < data size
0 with data length < header length
0 with incorrect version number
0 packets destined to dead next hop
0 fragments received
0 fragments dropped (dup or out of space)
0 fragments dropped (queue overflow)
0 fragments dropped after timeout
0 fragments dropped due to over limit
0 packets reassembled ok
1133057 packets for this host
0 packets for unknown/unsupported protocol
40146 packets forwarded
0 packets not forwardable
40146 redirects sent
1121700 packets sent from this host
0 packets sent with fabricated ip header
0 output packets dropped due to no bufs
0 output packets discarded due to no route
0 output datagrams fragmented
0 fragments created
0 datagrams that can't be fragmented
0 packets with bad options
0 packets with options handled without error
0 strict source and record route options
0 loose source and record route options
0 record route options
0 timestamp options
0 timestamp and address options
0 timestamp and prespecified address options
0 option packets dropped due to rate limit
0 router alert options
0 multicast packets dropped (no iflist)
0 packets dropped (src and int don't match)
0 transit re packets dropped on mgmt i/f
0 packets used first nexthop in ecmp unilist
0 incoming ttpoip packets received
0 incoming ttpoip packets dropped
0 outgoing TTPoIP packets sent
0 outgoing TTPoIP packets dropped
icmp:
0 drops due to rate limit
9 calls to icmp_error
0 errors not generated because old message was icmp
Output histogram:

```

```
295 echo reply
9 destination unreachable
0 messages with bad code fields
0 messages less than the minimum length
0 messages with bad checksum
0 messages with bad source address
0 messages with bad length
0 echo drops with broadcast or multicast destination address
0 timestamp drops with broadcast or multicast destination address
Input histogram:
295 echo
295 message responses generated
igmp:
0 messages received
0 messages received with too few bytes
0 messages received with bad checksum
0 membership queries received
0 membership queries received with invalid fields
0 membership reports received
0 membership reports received with invalid fields
0 membership reports received for groups to which we belong
0 Membership reports sent
raw_if:
0 RAW packets transmitted
0 PPPOE packets transmitted
0 ISDN packets transmitted
0 DIALER packets transmitted
0 PPP packets transmitted to pppd
0 PPP packets transmitted to jppd
0 IGMPv2 packets transmitted
13 output drops due to tx error
0 MPU packets transmitted
0 PPPOE packets received
0 ISDN packets received
0 DIALER packets received
0 PPP packets received from pppd
0 MPU packets received
0 PPP packets received from jppd
0 IGMPv2 packets received
0 Input drops due to bogus protocol
0 input drops due to no mbufs available
0 input drops due to no space in socket
0 input drops due to no socket
arp:
186413 datagrams received
88 ARP requests received
88 ARP replies received
0 resolution request received
0 unrestricted proxy requests
0 restricted proxy requests
0 received proxy requests
0 proxy requests not proxied
0 restricted proxy requests not proxied
0 datagrams with bogus interface
0 datagrams with incorrect length
0 datagrams for non-IP protocol
0 datagrams with unsupported op code
0 datagrams with bad protocol address length
0 datagrams with bad hardware address length
0 datagrams with multicast source address
0 datagrams with multicast source address
```

```

0 datagrams with my own hardware address
164 datagrams for an address not on the interface
0 datagrams with a broadcast source address
0 datagrams with source address duplicate to mine
186065 datagrams which were not for me
0 packets discarded waiting for resolution
0 packets sent after waiting for resolution
50 ARP requests sent
88 ARP replies sent
0 requests for memory denied
0 requests dropped on entry
0 requests dropped during retry
0 requests dropped due to interface deletion
0 requests on unnumbered interfaces
0 new requests on unnumbered interfaces
0 replies for from unnumbered interfaces
0 requests on unnumbered interface with non-subnetted donor
0 replies from unnumbered interface with non-subnetted donor
ip6:
0 total packets received
0 packets with size smaller than minimum
0 packets with data size < data length
0 packets with bad options
0 packets with incorrect version number
0 fragments received
0 fragments dropped (dup or out of space)
0 fragments dropped after timeout
0 fragments that exceeded limit
0 packets reassembled ok
0 packets for this host
0 packets forwarded
0 packets not forwardable
0 redirects sent
0 packets sent from this host
0 packets sent with fabricated ip header
0 output packets dropped due to no bufs, etc.
0 output datagrams fragmented
0 fragments created
0 datagrams that can't be fragmented
0 packets that violated scope rules
0 multicast packets which we don't join
0 packets whose headers are not continuous
0 tunneling packets that can't find gif
0 packets discarded due to too many headers
0 failures of source address selection
0 forward cache hit
0 forward cache miss
0 Packets destined to dead next hop
0 option packets dropped due to rate limit
0 Packets dropped (src and int don't match)
0 packets dropped due to bad protocol
0 transit re packet(null) dropped on mgmt i/f
icmp6:
0 Calls to icmp_error
0 Errors not generated because old message was icmp error
0 Errors not generated because rate limitation
0 Messages with bad code fields
0 Messages < minimum length
0 Bad checksums
0 Messages with bad length
0 No route

```

```
0 Administratively prohibited
0 Beyond scope
0 Address unreachable
0 Port unreachable
0 packet too big
0 Time exceed transit
0 Time exceed reassembly
0 Erroneous header field
0 Unrecognized next header
0 Unrecognized option
0 redirect
0 Unknown
0 Message responses generated
0 Messages with too many ND options
pfkey:
0 Requests sent from userland
0 Bytes sent from userland
histogram by message type:
0 reserved
0 dump
0 Messages with invalid length field
0 Messages with invalid version field
0 Messages with invalid message type field
0 Messages too short
0 Messages with memory allocation failure
0 Messages with duplicate extension
0 Messages with invalid extension type
0 Messages with invalid sa type
0 Messages with invalid address extension
0 Requests sent to userland
0 Bytes sent to userland
histogram by message type:
0 reserved
0 dump
0 Messages toward single socket
0 Messages toward all sockets
0 Messages toward registered sockets
0 Messages with memory allocation failure
c1n1:
0 Total packets received
0 Packets delivered
0 Too small packets
0 Packets with bad header length
0 Packets with bad checksum
0 Bad version packets
0 Unknown or unsupported protocol packets
0 Packets with bogus sdl size
0 No free memory in socket buffer
0 Send packets discarded
0 Sbappend failure
0 Mcopy failure
0 Address fields were not reasonable
0 Segment information forgotten
0 Forwarded packets
0 Total packets sent
0 Output packets discarded
0 Non-forwarded packets
0 Packets fragmented
0 Fragments sent
0 Fragments discarded
0 Fragments timed out
```

```
0 Fragmentation prohibited
0 Packets reconstructed
0 Packets destined to dead nexthop
0 Packets discarded due to no route
0 Error pdu rate drops
0 ER pdu generation failure
esis:
0 Total pkts received
0 Total packets consumed by protocol
0 Pdus received with bad checksum
0 Pdus received with bad version number
0 Pdus received with bad type field
0 Short pdus received
0 Pdus with bogus sdl size
0 Pdus with bad header length
0 Pdus with unknown or unsupported protocol
0 No free memory in socket buffer
0 Send packets discarded
0 Sbappend failure
0 Mcopy failure
0 ISO family not configured
tnp:
0 Unicast packets received
0 Broadcast packets received
0 Fragmented packets received
0 Hello packets dropped
0 Fragments dropped
0 Fragment reassembly queue flushes
0 Packets with tnp src address collision received
0 Hello packets received
0 Control packets received
0 Rdp packets received
0 Udp packets received
0 Tunnel packets received
0 Input packets discarded with no protocol
0 Packets of version unspecified received
0 Packets of version 1 received
0 Packets of version 2 received
0 Packets of version 3 received
0 Unicast packets sent
0 Broadcast packets sent
0 Fragmented packets sent
0 Hello packets dropped
0 Fragments dropped
0 Hello packets sent
0 Control packets sent
0 Rdp packets sent
0 Udp packets sent
0 Tunnel packets sent
0 Packets sent with unknown protocol
0 Packets of version unspecified sent
0 Packets of version 1 sent
0 Packets of version 2 sent
0 Packets of version 3 sent
rdp:
0 Input packets
0 Packets discarded for bad checksum
0 Packets discarded due to bad sequence number
0 Refused connections
0 Acks received
0 Packets dropped due to full socket buffers
```

```
0 Retransmits
0 Output packets
0 Acks sent
0 Connects
0 Closes
0 Keepalives received
0 Keepalives sent
tudp:
67 Datagrams received
0 Datagrams with incomplete header
0 Datagrams with bad data length field
0 Datagrams with bad checksum
0 Datagrams dropped due to no socket
0 Broadcast/multicast datagrams dropped due to no socket
0 Datagrams dropped due to full socket buffers
67 Delivered
68 Datagrams output
ttp:
0 Packets sent
0 Packets sent while unconnected
0 Packets sent while interface down
0 Packets sent couldn't get buffer
0 Packets sent couldn't find neighbor
0 L2 packets received
0 Unknown L3 packets received
0 IPv4 L3 packets received
0 MPLS L3 packets received
0 MPLS->IPv4 L3 packets received
0 IPv4->MPLS L3 packets received
0 IPv6 L3 packets received
0 ARP L3 packets received
0 CLNP L3 packets received
0 TNP L3 packets received
0 NULL L3 packets received
0 Cyclotron cycle L3 packets received
0 Cyclotron send L3 packets received
0 Packets received while unconnected
0 Packets received from unknown ifl
0 Input packets couldn't get buffer
0 Input packets with bad type
0 Input packets with discard type
0 Input packets with too many tlvs
0 Input packets with bad tlv header
70633 Input packets with bad tlv type
68877 Input packets dropped based on tlv result0 Input packets for which rt lookup
  is bypassed
mpls:
0 Total MPLS packets received
0 Packets forwarded
0 Packets dropped
0 Packets with header too small
0 After tagging, packets can't fit link MTU
0 Packets with IPv4 explicit NULL tag
0 Packets with IPv4 explicit NULL cksum errors
0 Packets with router alert tag
0 LSP ping packets (ttl-expired/router alert)
0 Packets with ttl expired
0 Packets with tag encoding error
0 Packets discarded due to no route
0 Packets used first nexthop in ecmp unilist
vpls:
```

```

0 Total packets received
0 Packets with size smaller than minimum
0 Packets with incorrect version number
0 Packets for this host
0 Packets with no logical interface
0 Packets with no family
0 Packets with no route table
582 Copyright © 2010, Juniper Networks, Inc.
0 Packets with no auxiliary table
0 Packets with no corefacing entry
0 packets with no CE-facing entry
0 MAC route learning requests
0 MAC routes learnt
0 Requests to learn an existing route
0 Learning requests while learning disabled on interface
0 Learning requests over capacity
0 MAC routes moved
0 Requests to move static route
0 MAC route aging requests
0 MAC routes aged
0 Bogus address in aging requests
0 Requests to age static route
0 Requests to re-ageout aged route
0 Requests involving multiple peer FEs
0 Aging acks from PFE
0 Aging non-acks from PFE
0 Aging requests timed out waiting on FEs
0 Aging requests over max-rate
0 Errors finding peer FEs
0 Unsupported platform
0 Packets dropped due to no l3 route table
0 Packets dropped due to no local ifl
0 Packets punted
0 Packets dropped due to no socket
bridge:
Input:
0 packets received
0 packets forwarded
0 packets failed to forward
0 packets dropped
0 packets with vmember lookup failures
0 packets with vlan lookup failures
0 packets with stp state lookup failures
0 packets dropped due to stp blocked/listening
0 packets dropped due to stp learning
0 packets with src MAC learning failures
0 packets with input control processing failures
Forward:
0 packets sent successfully
0 packets with send failures
0 packets forwarded to l3 interface
0 packets with l3 send failures
0 packets discarded
0 packets with l2ifl store failures
0 packets with ifl mismatch failures
0 packets with packet duplication failures
0 packets with tag lookup failures
0 packets with no route for DMAC
0 packets with no route table
0 packets with no nexthop
0 packets with dead nexthop

```

```
0 packets with eof reached error
Learning:
0 MACs learned
0 packets sent to l3 interface
0 packets with l3 send failures
0 packets hit holdq while learning
0 MAC moves
0 packets discarded
0 packets with no route for SMAC
0 packets with no nexthop
0 packets with dead nexthop
0 packets dropped due to no resolve route
0 packets with l3 ifd lookup failures
0 packets with l3 ifl lookup failures
0 packets with l3 invalid rnh
0 packets with no route for SMAC in clone learning
0 packets with no nexthop in clone learning
0 packets with dead nexthop in clone learning
0 packets dropped due to no resolve nh in clone learning
Output:
0 packets forwarded
0 packets failed to forward
0 packets with vmember lookup failures
0 packets with vlan lookup failures
0 packets with input control processing failures
Send:
0 packets sent successfully
0 packets with send failures
0 packets dropped due to interface down
0 packets with dev output failures
0 blocked ifl discards
0 packets with tag lookup failures
0 packets with stp state lookup failures
0 packets with tag insertion failures
0 packets with tag removal failures
Flood:
0 packets flooded
0 flood failures
IGMP:
0 packets sent successfully
0 packets with send failures
0 packets forwarded
0 packets failed to forward
0 packets with mpull failures
0 packets with vmember lookup failures
0 packets with vlan lookup failures
0 packets with ifl lookup failures
0 packets with tag lookup failures
Misc:
0 packets with size smaller than minimum
0 packets with double tags
0 packets with no ifl
0 packets with no family
0 packets with no route table
```


show system storage

List of Syntax	Syntax on page 1177 Syntax (EX Series Switches) on page 1177 Syntax (MX Series Router) on page 1177 Syntax (QFX Series) on page 1177 Syntax (SRX Series) on page 1177 Syntax (TX Matrix Router) on page 1177 Syntax (TX Matrix Plus Router and TX Matrix Plus Router with 3D SIBs) on page 1177
Syntax	<pre>show system storage <detail> <invoke-on (all-routing-engines other-routing-engine)></pre>
Syntax (EX Series Switches)	<pre>show system storage <detail> <all-members> <local> <member member-id> <invoke-on (all-routing-engines other-routing-engine)></pre>
Syntax (MX Series Router)	<pre>show system storage <detail> <all-members> <local> <member member-id> <invoke-on (all-routing-engines other-routing-engine)></pre>
Syntax (QFX Series)	<pre>show system storage <detail> <infrastructure name> <interconnect-device name> <node-group name> <invoke-on (all-routing-engines other-routing-engine)></pre>
Syntax (SRX Series)	<pre>show system storage <detail> <partitions> <invoke-on (all-routing-engines other-routing-engine)></pre>
Syntax (TX Matrix Router)	<pre>show system storage <detail> <all-chassis all-lcc lcc number scc> <invoke-on (all-routing-engines other-routing-engine)></pre>
Syntax (TX Matrix Plus Router and TX Matrix Plus Router with 3D SIBs)	<pre>show system storage <detail> <all-chassis all-lcc lcc number sfc number> <invoke-on (all-routing-engines other-routing-engine)></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>

sfc option introduced for the TX Matrix Plus router in JUNOS Release 9.6.
Command introduced in Junos OS Release 11.1 for the QFX Series.
Option **invoke-on (all-routing-engines | other-routing-engine)** introduced in Junos OS Release 14.1
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

- Description** Display statistics about the amount of free disk space in the router's or switch's file systems.
- Options**
- none**—Display standard information about the amount of free disk space in the router's or switch's file systems.
 - detail**—(Optional) Display detailed output.
 - invoke-on all-routing-engines**—(Optional) Display the system storage information on all master and backup Routing Engines on a routing matrix based on the TX Matrix or TX Matrix Plus router or on a router that has dual Routing Engines.
 - invoke-on other-routing-engines**—(Optional) Display the system storage information on the other Routing Engine. For example, if you issue this command on the master Routing Engine on an M320 router, the JUNOS Software displays the system storage information on the backup Routing Engine. On a routing matrix based on the TX Matrix or TX Matrix Plus router, if you issue this command on the TX Matrix or TX Matrix Plus router's master Routing Engine, the JUNOS Software displays all the system storage information on all the backup Routing Engines.
 - all-chassis**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) Display system storage statistics for all the routers in the chassis.
 - all-lcc**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system storage statistics for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, display system storage statistics for all routers connected to the TX Matrix Plus router.
 - all-members**—(EX4200 switches and MX Series routers only) (Optional) Display system storage statistics for all members of the Virtual Chassis configuration.
 - infrastructure *name***—(QFabric systems only) (Optional) Display system storage statistics for the fabric control Routing Engines or fabric manager Routing Engines.
 - interconnect-device *name***—(QFabric systems only) (Optional) Display system storage statistics for the Interconnect device.
 - lcc *number***—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system storage statistics for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display system storage statistics for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(EX4200 switches and MX Series routers only) (Optional) Display system storage statistics for the local Virtual Chassis member.

member *member-id*—(EX4200 switches and MX Series routers only) (Optional) Display system storage statistics for the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

node-group *name*—(QFabric systems only) (Optional) Display system storage statistics for the Node group.

scc—(TX Matrix routers only) (Optional) Display system storage statistics for the TX Matrix router (or switch-card chassis).

sfc *number*—(TX Matrix Plus routers only) (Optional) Display system storage statistics for the TX Matrix Plus router. Replace *number* with 0.

Additional Information By default, when you issue the **show system storage** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

Required Privilege Level view

Related Documentation

- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)
- [show system storage partitions \(View SRX Series\)](#)

List of Sample Output

- [show system storage on page 1180](#)
- [show system storage \(TX Matrix Plus Router\) on page 1180](#)
- [show system storage \(QFX3500 Switch\) on page 1182](#)
- [show system storage invoke-on all-routing-engines on page 1183](#)
- [show system storage invoke-on other-routing-engine on page 1184](#)

Output Fields Table 79 describes the output fields for the **show system storage** command. Output fields are listed in the approximate order in which they appear.

Table 79: show system storage Output Fields

Field Name	Field Description
Filesystem	Name of the filesystem.
Size	Size of the filesystem.
Used	Amount of space used in the filesystem.
Avail	Amount of space available in the filesystem.
Capacity	Percentage of the filesystem space that is being used.
Mounted on	Directory in which the filesystem is mounted.

Sample Output

show system storage

```

user@host> show system storage
Filesystem      Size      Used      Avail  Capacity  Mounted on
/dev/ad0s1a      77M       37M       34M      52%      /
devfs            16K       16K        0B     100%    /dev/
/dev/vn0         12M       12M        0B     100%    /packages/mnt/jbase
/dev/vn1         39M       39M        0B     100%
/packages/mnt/jkernel-7.2R1.7
/dev/vn2         12M       12M        0B     100%
/packages/mnt/jpfe-M40-7.2R1.7
/dev/vn3         2.3M      2.3M        0B     100%
/packages/mnt/jdocs-7.2R1.7
/dev/vn4         14M       14M        0B     100%
/packages/mnt/jroute-7.2R1.7
/dev/vn5         4.5M      4.5M        0B     100%
/packages/mnt/jcrypto-7.2R1.7
mfs:172         1.5G      4.0K       1.3G      0%      /tmp
/dev/ad0s1e      12M       20K        11M      0%      /config
procfs          4.0K      4.0K        0B     100%    /proc
/dev/ad1s1f      9.4G      4.9G       3.7G      57%     /var

```

show system storage (TX Matrix Plus Router)

```

user@host> show system storage
sfc0-re0:
-----
Filesystem      Size      Used      Avail  Capacity  Mounted on
/dev/ad0s1a      3.4G      178M       2.9G      6%      /
devfs            1.0K      1.0K        0B     100%    /dev
devfs            1.0K      1.0K        0B     100%    /dev/
/dev/md0         33M       33M        0B     100%    /packages/mnt/jbase
/dev/md1        216M      216M        0B     100%
/packages/mnt/jkernel-9.6-20090519.0
/dev/md2         66M       66M        0B     100%
/packages/mnt/jpfe-T-9.6-20090519.0

```

/dev/md3	4.1M	4.1M	0B	100%	
/packages/mnt/jdocs-9.6-20090519.0					
/dev/md4	57M	57M	0B	100%	
/packages/mnt/jroute-9.6-20090519.0					
/dev/md5	15M	15M	0B	100%	
/packages/mnt/jcrypto-9.6-20090519.0					
/dev/md6	34M	34M	0B	100%	
/packages/mnt/jpfe-common-9.6-20090519.0					
/dev/md7	2.0G	10.0K	1.8G	0%	/tmp
/dev/md8	2.0G	1.0M	1.8G	0%	/mfs
/dev/ad0s1e	383M	82K	352M	0%	/config
procfs	4.0K	4.0K	0B	100%	/proc
/dev/ad1s1f	52G	7.5G	40G	16%	/var

lcc0-re0:

Filesystem	Size	Used	Avail	Capacity	Mounted on
/dev/ad0s1a	3.4G	178M	2.9G	6%	/
devfs	1.0K	1.0K	0B	100%	/dev
devfs	1.0K	1.0K	0B	100%	/dev/
/dev/md0	33M	33M	0B	100%	/packages/mnt/jbase
/dev/md1	216M	216M	0B	100%	
/packages/mnt/jkernel-9.6-20090519.0					
/dev/md2	66M	66M	0B	100%	
/packages/mnt/jpfe-T-9.6-20090519.0					
/dev/md3	4.1M	4.1M	0B	100%	
/packages/mnt/jdocs-9.6-20090519.0					
/dev/md4	57M	57M	0B	100%	
/packages/mnt/jroute-9.6-20090519.0					
/dev/md5	15M	15M	0B	100%	
/packages/mnt/jcrypto-9.6-20090519.0					
/dev/md6	34M	34M	0B	100%	
/packages/mnt/jpfe-common-9.6-20090519.0					
/dev/md7	2.0G	10.0K	1.8G	0%	/tmp
/dev/md8	2.0G	540K	1.8G	0%	/mfs
/dev/ad0s1e	383M	88K	352M	0%	/config
procfs	4.0K	4.0K	0B	100%	/proc
/dev/ad1s1f	52G	6.3G	41G	13%	/var

lcc1-re0:

Filesystem	Size	Used	Avail	Capacity	Mounted on
/dev/ad0s1a	3.4G	178M	2.9G	6%	/
devfs	1.0K	1.0K	0B	100%	/dev
devfs	1.0K	1.0K	0B	100%	/dev/
/dev/md0	33M	33M	0B	100%	/packages/mnt/jbase
/dev/md1	216M	216M	0B	100%	
/packages/mnt/jkernel-9.6-20090519.0					
/dev/md2	66M	66M	0B	100%	
/packages/mnt/jpfe-T-9.6-20090519.0					
/dev/md3	4.1M	4.1M	0B	100%	
/packages/mnt/jdocs-9.6-20090519.0					
/dev/md4	57M	57M	0B	100%	
/packages/mnt/jroute-9.6-20090519.0					
/dev/md5	15M	15M	0B	100%	
/packages/mnt/jcrypto-9.6-20090519.0					
/dev/md6	34M	34M	0B	100%	
/packages/mnt/jpfe-common-9.6-20090519.0					
/dev/md7	2.0G	10.0K	1.8G	0%	/tmp
/dev/md8	2.0G	540K	1.8G	0%	/mfs
/dev/ad0s1e	383M	88K	352M	0%	/config

```

procfs                4.0K      4.0K      0B      100% /proc
/dev/ad1s1f           23G      13G      7.7G      64% /var

lcc2-re0:
-----
Filesystem            Size      Used      Avail  Capacity  Mounted on
/dev/ad0s1a           3.4G      178M      2.9G        6% /
devfs                 1.0K      1.0K      0B      100% /dev
devfs                 1.0K      1.0K      0B      100% /dev/
/dev/md0               33M       33M       0B      100% /packages/mnt/jbase
/dev/md1              216M      216M       0B      100%
/packages/mnt/jkernel-9.6-20090519.0
/dev/md2               66M       66M       0B      100%
/packages/mnt/jpfe-T-9.6-20090519.0
/dev/md3              4.1M      4.1M       0B      100%
/packages/mnt/jdocs-9.6-20090519.0
/dev/md4               57M       57M       0B      100%
/packages/mnt/jroute-9.6-20090519.0
/dev/md5               15M       15M       0B      100%
/packages/mnt/jcrypto-9.6-20090519.0
/dev/md6               34M       34M       0B      100%
/packages/mnt/jpfe-common-9.6-20090519.0
/dev/md7               2.0G      10.0K      1.8G        0% /tmp
/dev/md8               2.0G      540K      1.8G        0% /mfs
/dev/ad0s1e           383M       64K      352M        0% /config
procfs                4.0K      4.0K      0B      100% /proc
/dev/ad1s1f           23G      3.7G      17G       18% /var

lcc3-re0:
-----
Filesystem            Size      Used      Avail  Capacity  Mounted on
/dev/ad0s1a           3.4G      178M      2.9G        6% /
devfs                 1.0K      1.0K      0B      100% /dev
devfs                 1.0K      1.0K      0B      100% /dev/
/dev/md0               33M       33M       0B      100% /packages/mnt/jbase
/dev/md1              216M      216M       0B      100%
/packages/mnt/jkernel-9.6-20090519.0
/dev/md2               66M       66M       0B      100%
/packages/mnt/jpfe-T-9.6-20090519.0
/dev/md3              4.1M      4.1M       0B      100%
/packages/mnt/jdocs-9.6-20090519.0
/dev/md4               57M       57M       0B      100%
/packages/mnt/jroute-9.6-20090519.0
/dev/md5               15M       15M       0B      100%
/packages/mnt/jcrypto-9.6-20090519.0
/dev/md6               34M       34M       0B      100%
/packages/mnt/jpfe-common-9.6-20090519.0
/dev/md7               2.0G      10.0K      1.8G        0% /tmp
/dev/md8               2.0G      540K      1.8G        0% /mfs
/dev/ad0s1e           383M       34K      352M        0% /config
procfs                4.0K      4.0K      0B      100% /proc
/dev/ad1s1f           23G      18G      3.5G       84% /var

```

show system storage (QFX3500 Switch)

```

user@switch> show system storage
Filesystem            Size      Used      Avail  Capacity  Mounted on
/dev/da0s2a           343M      192M      123M       61% /
devfs                 1.0K      1.0K      0B      100% /dev
/dev/md0              119M      119M       0B      100% /packages/mnt/jbase
/dev/md1              513M      513M       0B      100%

```

```

/packages/mnt/jkernel-qfx-11.1R1.5
/dev/md2          37M          37M          0B          100%
/packages/mnt/jpfe-qfx-e9xxx-11.1R1.5
/dev/md3          6.0M          6.0M          0B          100%
/packages/mnt/jdocs-qfx-11.1R1.5
/dev/md4          216M         216M          0B          100%
/packages/mnt/jroute-qfx-11.1R1.5
/dev/md5          59M          59M          0B          100%
/packages/mnt/jcrypto-qfx-11.1R1.5
/dev/md6          85M          85M          0B          100%
/packages/mnt/jswitch-qfx-11.1R1.5
/dev/md7          63M          8.0K          58M          0% /tmp
/dev/da0s2f       228M          14M         196M          7% /var
/dev/da0s3d       590M          3.0M         540M          1% /var/tmp
/dev/da0s3e       104M          162K          95M          0% /config
procfs            4.0K          4.0K          0B          100% /proc

```

show system storage invoke-on all-routing-engines

```
user@host> show system storage invoke-on all-routing-engines
```

```
re0:
```

```

-----
Filesystem      Size      Used      Avail  Capacity  Mounted on
/dev/ad0s1a     3.3G      440M      2.6G      14%      /
devfs           1.0K      1.0K      0B        100%     /dev
/dev/md0        76M       76M       0B        100%     /packages/mnt/jbase
/dev/md1        40M       40M       0B        100%
/packages/mnt/jkernel64-14.1-20140407.1
/dev/md2        219M      219M      0B        100%
/packages/mnt/jpfe-T-14.1-20140407.1
/dev/md3        5.4M      5.4M      0B        100%
/packages/mnt/jdocs-14.1-20140407.1
/dev/md4        116M      116M      0B        100%
/packages/mnt/jroute-14.1-20140407.1
/dev/md5        44M       44M      0B        100%
/packages/mnt/jcrypto64-14.1-20140407.1
/dev/md6        70M       70M      0B        100%
/packages/mnt/jpfe-common-14.1-20140407.1
/dev/md7        182K      182K      0B        100%
/packages/mnt/jplatform-14.1-20140407.1
/dev/md8        499M      499M      0B        100%
/packages/mnt/jruntime-14.1-20140407.1
/dev/md9        41M       41M      0B        100%
/packages/mnt/jruntime64-14.1-20140407.1
/dev/md10       12M       12M      0B        100%
/packages/mnt/py-base-i386-14.1-20140407.1
/dev/md11       3.2G      8.0K      2.9G       0% /tmp
/dev/md12       3.2G      1.1M      2.9G       0% /mfs
/dev/ad0s1e     376M      220K      346M       0% /config
procfs          4.0K      4.0K      0B        100% /proc
/dev/ad1s1f     50G       43G      3.2G      93% /var

```

```
re1:
```

```

-----
Filesystem      Size      Used      Avail  Capacity  Mounted on
/dev/ad0s1a     3.3G      440M      2.6G      14%      /
devfs           1.0K      1.0K      0B        100%     /dev
/dev/md0        76M       76M       0B        100%     /packages/mnt/jbase
/dev/md1        40M       40M       0B        100%
/packages/mnt/jkernel64-14.1-20140407.1
/dev/md2        219M      219M      0B        100%

```

```

/packages/mnt/jpfe-T-14.1-20140407.1
/dev/md3          5.4M      5.4M      0B      100%
/packages/mnt/jdocs-14.1-20140407.1
/dev/md4          116M     116M      0B      100%
/packages/mnt/jroute-14.1-20140407.1
/dev/md5          44M      44M      0B      100%
/packages/mnt/jcrypto64-14.1-20140407.1
/dev/md6          70M      70M      0B      100%
/packages/mnt/jpfe-common-14.1-20140407.1
/dev/md7          182K     182K      0B      100%
/packages/mnt/jplatform-14.1-20140407.1
/dev/md8          499M     499M      0B      100%
/packages/mnt/jruntime-14.1-20140407.1
/dev/md9          41M      41M      0B      100%
/packages/mnt/jruntime64-14.1-20140407.1
/dev/md10         12M      12M      0B      100%
/packages/mnt/py-base-i386-14.1-20140407.1
/dev/md11         3.2G      8.0K      2.9G      0% /tmp
/dev/md12         3.2G     662K      2.9G      0% /mfs
/dev/ad0s1e       375M     230K     344M      0% /config
procfs           4.0K      4.0K      0B      100% /proc
/dev/ad1s1f       52G      46G      2.2G      95% /var

```

show system storage invoke-on other-routing-engine

```

user@host> show system storage invoke-on other-routing-engine
rel:

```

```

-----
Filesystem      Size      Used      Avail  Capacity  Mounted on
/dev/ad0s1a     3.3G      440M      2.6G      14%      /
devfs           1.0K      1.0K      0B      100%     /dev
/dev/md0        76M      76M      0B      100%     /packages/mnt/jbase
/dev/md1        40M      40M      0B      100%
/packages/mnt/jkernel64-14.1-20140407.1
/dev/md2        219M     219M      0B      100%
/packages/mnt/jpfe-T-14.1-20140407.1
/dev/md3        5.4M      5.4M      0B      100%
/packages/mnt/jdocs-14.1-20140407.1
/dev/md4        116M     116M      0B      100%
/packages/mnt/jroute-14.1-20140407.1
/dev/md5        44M      44M      0B      100%
/packages/mnt/jcrypto64-14.1-20140407.1
/dev/md6        70M      70M      0B      100%
/packages/mnt/jpfe-common-14.1-20140407.1
/dev/md7        182K     182K      0B      100%
/packages/mnt/jplatform-14.1-20140407.1
/dev/md8        499M     499M      0B      100%
/packages/mnt/jruntime-14.1-20140407.1
/dev/md9        41M      41M      0B      100%
/packages/mnt/jruntime64-14.1-20140407.1
/dev/md10       12M      12M      0B      100%
/packages/mnt/py-base-i386-14.1-20140407.1
/dev/md11       3.2G      8.0K      2.9G      0% /tmp
/dev/md12       3.2G     662K      2.9G      0% /mfs
/dev/ad0s1e     375M     230K     344M      0% /config
procfs         4.0K      4.0K      0B      100% /proc
/dev/ad1s1f     52G      46G      2.2G      95% /var

```


show system uptime

List of Syntax	Syntax on page 1185 Syntax (EX Series Switches) on page 1185 Syntax (QFX Series) on page 1185 Syntax (TX Matrix Router) on page 1185 Syntax (TX Matrix Plus Router) on page 1185 Syntax (MX Series Router) on page 1185
Syntax	show system uptime
Syntax (EX Series Switches)	show system uptime <all-members> <local> <member <i>member-id</i> >
Syntax (QFX Series)	show system uptime <director-group <i>name</i> > <infrastructure <i>name</i> > <interconnect-device <i>name</i> > <node-group <i>name</i> >
Syntax (TX Matrix Router)	show system uptime <all-chassis all-lcc lcc <i>number</i> scc>
Syntax (TX Matrix Plus Router)	show system uptime <detail> <all-chassis all-lcc lcc <i>number</i> sfc <i>number</i> >
Syntax (MX Series Router)	show system uptime <all-members> <invoke-on> <local> <member <i>member-id</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc option introduced for the TX Matrix Plus router in JUNOS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Display the current time and information about how long the router or switch, router or switch software, and routing protocols have been running.
Options	none —Show time since the system rebooted and processes started. all-chassis —(TX Matrix routers and TX Matrix Plus routers only) (Optional) Show time since the system rebooted and processes started on all the routers in the chassis. all-lcc —(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, show time since the system rebooted and processes started for all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus

router, show time since the system rebooted and processes started for all connected T1600 or T4000 LCCs.

all-members—(EX4200 switches and MX Series routers only) (Optional) Show time since the system rebooted and processes started on all members of the Virtual Chassis configuration.

director-group *name*—(QFabric systems only) (Optional) Show time since the system rebooted and processes started on the Director group.

infrastructure *name*—(QFabric systems only) (Optional) Show time since the system rebooted and processes started on the fabric control Routing Engine and fabric manager Routing Engine.

interconnect-device *name*—(QFabric systems only) (Optional) Show time since the system rebooted and processes started on the Interconnect device.

invoke-on—(MX Series routers only) (Optional) Display the time since the system rebooted and processes started on the master Routing Engine, backup Routing Engine, or both, on a router with two Routing Engines.

lcc *number*—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, show time since the system rebooted and processes started for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, show time since the system rebooted and processes started for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(EX4200 switches and MX Series routers only) (Optional) Show time since the system rebooted and processes started on the local Virtual Chassis member.

member *member-id*—(EX4200 switches and MX Series routers only) (Optional) Show time since the system rebooted and processes started on the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

node-group *name*—(QFabric systems only) (Optional) Show time since the system rebooted and processes started on the Node group.

scc—(TX Matrix routers only) (Optional) Show time since the system rebooted and processes started for the TX Matrix router (or switch-card chassis).

sfc number—(TX Matrix Plus routers only) (Optional) Show time since the system rebooted and processes started for the TX Matrix Plus router. Replace *number* with 0.

Additional Information By default, when you issue the **show system uptime** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

Required Privilege Level view

Related Documentation

- [10-Gigabit Ethernet LAN/WAN PIC with XFP \(T640 Router\)](#)
- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

List of Sample Output

[show system uptime on page 1188](#)
[show system uptime all-lcc \(TX Matrix Router\) on page 1188](#)
[show system uptime all-lcc \(TX Matrix Plus Router\) on page 1188](#)
[show system uptime \(EX Series\) on page 1189](#)
[show system uptime \(QFX Series\) on page 1189](#)

Output Fields Table 80 describes the output fields for the **show system uptime** command. Output fields are listed in the approximate order in which they appear.

Table 80: show system uptime Output Fields

Field Name	Field Description
Current time	Current system time in UTC.
Time Source	Reference time source that the system is locked to.
System booted	Date and time when the Routing Engine on the router or switch was last booted and how long it has been running.
Protocols started	Date and time when the routing protocols were last started and how long they have been running.
Last configured	Date and time when a configuration was last committed. Also shows the name of the user who issued the last commit command.
time and up	Current time, in the local time zone, and how long the router or switch has been operational.
users	Number of users logged in to the router or switch.
load averages	Load averages for the last 1 minute, 5 minutes, and 15 minutes.

Sample Output

show system uptime

```
user@host> show system uptime
Current time:      1998-10-13 19:45:47 UTC
Time Source:      NTP CLOCK
System booted:     1998-10-12 20:51:41 UTC (22:54:06 ago)
Protocols started: 1998-10-13 19:33:45 UTC (00:12:02 ago)
Last configured:   1998-10-13 19:33:45 UTC (00:12:02 ago) by abc
12:45PM up 22:54, 2 users, load averages: 0.07, 0.02, 0.01
```

show system uptime all-lcc (TX Matrix Router)

```
user@host> show system uptime all-lcc
lcc0-re0:
-----
Current time: 2004-09-13 09:55:35 PDT
Time Source: LOCAL CLOCK
System booted: 2004-09-13 03:13:55 PDT (06:41:40 ago)
Last configured: 2004-09-13 03:17:48 PDT (06:37:47 ago) by root
9:55AM PDT up 6:42, 1 user, load averages: 0.02, 0.03, 0.00
lcc2-re0:
-----
Current time: 2004-09-13 09:55:35 PDT
Time Source: LOCAL CLOCK
System booted: 2004-09-12 03:23:43 PDT (1d 06:31 ago)
Last configured: 2004-09-13 03:05:36 PDT (06:49:59 ago) by root
9:55AM PDT up 1 day, 6:32, 1 user, load averages: 0.02, 0.01, 0.00
```

show system uptime all-lcc (TX Matrix Plus Router)

```
user@host> show system uptime all-lcc
sfc0-re0:
-----
Current time: 2009-05-25 00:24:30 PDT
Time Source: NTP CLOCK
System booted: 2009-05-24 06:39:33 PDT (17:44:57 ago)
Protocols started: 2009-05-24 06:40:30 PDT (17:44:00 ago)
Last configured: 2009-05-24 06:33:27 PDT (17:51:03 ago) by gregdo
12:24AM up 17:45, 2 users, load averages: 0.07, 0.05, 0.01

lcc0-re0:
-----
Current time: 2009-05-25 00:24:30 PDT
Time Source: NTP CLOCK
System booted: 2009-05-24 06:39:46 PDT (17:44:44 ago)
error: the routing subsystem is not running
Last configured: 2009-05-24 06:40:47 PDT (17:43:43 ago) by root
12:24AM up 17:45, 0 users, load averages: 0.00, 0.00, 0.00

lcc1-re0:
-----
Current time: 2009-05-25 00:24:30 PDT
Time Source: NTP CLOCK
System booted: 2009-05-24 06:39:38 PDT (17:44:52 ago)
error: the routing subsystem is not running
Last configured: 2009-05-24 06:40:18 PDT (17:44:12 ago) by root
12:24AM up 17:45, 0 users, load averages: 0.00, 0.00, 0.00

lcc2-re0:
```

```

-----
Current time: 2009-05-25 00:24:30 PDT
Time Source: NTP CLOCK
System booted: 2009-05-24 06:39:48 PDT (17:44:42 ago)
error: the routing subsystem is not running
Last configured: 2009-05-24 06:40:44 PDT (17:43:46 ago) by root
12:24AM up 17:45, 0 users, load averages: 0.00, 0.00, 0.00

```

```
lcc3-re0:
```

```

-----
Current time: 2009-05-25 00:24:30 PDT
Time Source: NTP CLOCK
System booted: 2009-05-24 06:39:44 PDT (17:44:46 ago)
error: the routing subsystem is not running
Last configured: 2009-05-24 06:40:08 PDT (17:44:22 ago) by root
12:24AM up 17:45, 0 users, load averages: 0.00, 0.00, 0.00

```

show system uptime (EX Series)

```

user@switch> show system uptime
Current time: 2014-03-12 16:39:56 UTC
Time Source: NTP CLOCK
System booted: 2014-03-12 14:58:05 UTC (01:41:51 ago)
Protocols started: 2014-03-12 14:59:48 UTC (01:40:08 ago)
Last configured: 2014-03-12 14:58:58 UTC (01:40:58 ago) by root
4:39PM up 1:42, 4 users, load averages: 0.02, 0.02, 0.00

```

show system uptime (QFX Series)

```

user@switch> show system uptime
Current time: 2010-08-27 03:12:30 PDT
Time Source: NTP CLOCK
System booted: 2010-08-13 17:11:54 PDT (1w6d 10:00 ago)
Protocols started: 2010-08-13 17:13:56 PDT (1w6d 09:58 ago)
Last configured: 2010-08-26 05:54:00 PDT (21:18:30 ago) by regress
3:12AM up 13 days, 10:01, 3 users, load averages: 0.00, 0.00, 0.00

```

show system users

List of Syntax	Syntax on page 1190 Syntax (TX Matrix Router) on page 1190 Syntax (TX Matrix Plus Router) on page 1190 Syntax (MX Series Router) on page 1190
Syntax	show system users <no-resolve>
Syntax (TX Matrix Router)	show system users <all-chassis all-lcc lccnumber scc> <no-resolve>
Syntax (TX Matrix Plus Router)	show system users <detail> <all-chassis all-lcc lcc number sfc number> <no-resolve>
Syntax (MX Series Router)	show system users <all-members> <local> <member member-id> <no-resolve>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>sfc option introduced for the TX Matrix Plus router in JUNOS OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	List information about the users who are currently logged in to the router or switch.



NOTE: The **show system users** command lists the information about administrative users that are logged in to a router or switch using the CLI, J-Web, or an SSH client. The output does not list information about web users or automated users that are logged in from a remote client application using Junos XML APIs, such as NETCONF.

- Options**
- none**—List information about the users who are currently logged in to the router or switch.
 - all-chassis**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) Show users currently logged in to all the routers in the chassis.
 - all-lcc**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, show users currently logged in to all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, show users currently logged in to all connected T1600 or T4000 LCCs.

all-members—(MX Series routers only) (Optional) Display users currently logged in to all members of the Virtual Chassis configuration.

lcc *number*—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, show users currently logged in to a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, show users currently logged in to a specific router that is connected to the TX Matrix Plus router. Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(MX Series routers only) (Optional) Display users currently logged in to the local Virtual Chassis member.

member *member-id*—(MX Series routers only) (Optional) Display users currently logged in to the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

no-resolve—(Optional) Do not attempt to resolve IP addresses to hostnames.

scc—(TX Matrix routers only) (Optional) Show users currently logged in to the TX Matrix router (or switch-card chassis).

sfc *number*—(TX Matrix Plus routers only) (Optional) Show users currently logged in to the TX Matrix Plus router. Replace *number* with 0.

Additional Information By default, when you issue the **show system users** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

Required Privilege Level view

Related Documentation • [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

List of Sample Output [show system users on page 1192](#)
[show system users lcc no-resolve \(TX Matrix, TX Matrix Plus Router\) on page 1192](#)
[show system users \(TX Matrix Plus Router\) on page 1192](#)

[show system users \(QFX Series\) on page 1193](#)

[show system users no-resolve \(QFX Series\) on page 1193](#)

Output Fields Table 81 describes the output fields for the **show system users** command. Output fields are listed in the approximate order in which they appear.

Table 81: show system users Output Fields

Field Name	Field Description
<i>time and up</i>	Current time, in the local time zone, and how long the router or switch has been operational.
<i>users</i>	Number of users logged in to the router or switch.
<i>load averages</i>	Load averages for the last 1 minute, 5 minutes, and 15 minutes.
<i>USER</i>	Username.
<i>TTY</i>	Terminal through which the user is logged in.
<i>FROM</i>	System from which the user has logged in. A hyphen indicates that the user is logged in through the console.
<i>LOGIN@</i>	Time when the user logged in.
<i>IDLE</i>	How long the user has been idle.
<i>WHAT</i>	Processes that the user is running.

Sample Output

show system users

```
user@host> show system users
7:30PM up 4 days, 2:26, 2 users, load averages: 0.07, 0.02, 0.01
USER    TTY FROM          LOGIN@  IDLE WHAT
root    d0  -              Fri05PM 4days -csh (csh)
blue    p0  leve15.compan.net 7:30PM  - cli
```

show system users lcc no-resolve (TX Matrix, TX Matrix Plus Router)

```
user@host> show system users lcc 2 no-resolve
```

```
lcc2-re0:
```

```
-----
10:34AM PDT up 1 day, 7:11, 5 users, load averages: 0.03, 0.01, 0.00
USER    TTY FROM          LOGIN@  IDLE WHAT
root    d0  -              3:21AM  7:12 /bin/csh
user1    p0  scc-re0        10:15AM  - telnet hostA
user1    p1  scc-re0        10:16AM  - telnet hostA
user1    p2  scc-re0        10:19AM  - telnet hostA
user1    p3  scc-re0        10:24AM  - telnet hostA
```

show system users (TX Matrix Plus Router)

```
user@host> show system users
```


sfc0-re0:

```

-----
1:41AM up 26 mins, 3 users, load averages: 0.08, 0.04, 0.03
USER   TTY   FROM                               LOGIN@  IDLE WHAT
user2   p0    10.209.208.123                    1:18AM   21 cli
user2   p1    172.17.29.207                    1:37AM    2 cli
user2   p2    172.17.28.19                     1:40AM    - cli

```

lcc0-re0:

```

-----
1:41AM up 26 mins, 0 users, load averages: 0.00, 0.00, 0.03

```

lcc1-re0:

```

-----
1:41AM up 26 mins, 0 users, load averages: 0.00, 0.02, 0.03

```

lcc2-re0:

```

-----
1:41AM up 26 mins, 0 users, load averages: 0.16, 0.06, 0.02

```

lcc3-re0:

```

-----
1:41AM up 26 mins, 0 users, load averages: 0.12, 0.04, 0.04

```

user3@aj> show system users

sfc0-re0:

```

-----
1:42AM up 28 mins, 4 users, load averages: 0.02, 0.03, 0.02
USER   TTY   FROM                               LOGIN@  IDLE WHAT
user3   p0    pssraj-t61.jnpr.net              1:18AM   22 cli
user3   p1    eng-shell14.juniper.net          1:37AM    - cli
user3   p2    bigpink.juniper.net              1:40AM    - cli
user3   p3    sv-cutty-01.englab.juniper.net   1:42AM    - -csh (csh)

```

lcc0-re0:

```

-----
1:42AM up 28 mins, 0 users, load averages: 0.02, 0.01, 0.03

```

lcc1-re0:

```

-----
1:42AM up 28 mins, 0 users, load averages: 0.07, 0.04, 0.03

```

lcc2-re0:

```

-----
1:42AM up 27 mins, 0 users, load averages: 0.07, 0.06, 0.02

```

lcc3-re0:

```

-----
1:42AM up 28 mins, 0 users, load averages: 0.05, 0.04, 0.04

```

show system users (QFX Series)

user@switch> show system users

```

USER   TTY   FROM                               LOGIN@  IDLE WHAT
tlewis  p0    172.22.18.117                    2:54AM   39 -cli (cli)
tlewis  p1    172.22.18.117                    3:01AM    - -cli (cli)
tcheng  p2    172.22.17.197                    3:08AM   11 -cli (cli)

```

show system users no-resolve (QFX Series)

user@switch> show system users no-resolve

USER	TTY	FROM	LOGIN@	IDLE	WHAT
tlewis	p0	172.22.18.117	2:54AM	39	-cli (cli)
tlewis	p1	172.22.18.117	3:01AM	-	-cli (cli)
tcheng	p2	172.22.17.197	3:08AM	11	-cli (cli)

show system virtual-memory

List of Syntax	Syntax on page 1195 Syntax (EX Series) on page 1195 Syntax (TX Matrix Router) on page 1195 Syntax (TX Matrix Plus Router) on page 1195 Syntax (MX Series Router) on page 1195 Syntax (QFX Series) on page 1195
Syntax	show system virtual-memory
Syntax (EX Series)	show system virtual-memory <all-members> <local> <member <i>member-id</i> >
Syntax (TX Matrix Router)	show system virtual-memory <all-chassis all-lcc lcc <i>number</i> scc>
Syntax (TX Matrix Plus Router)	show system virtual-memory <all-chassis all-lcc lcc <i>number</i> sfc <i>number</i> >
Syntax (MX Series Router)	show system virtual-memory <all-members> <local> <member <i>member-id</i> >
Syntax (QFX Series)	show system virtual-memory <infrastructure <i>name</i> > <interconnect-device <i>name</i> > <node-group <i>name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Display the usage of Junos OS kernel memory listed first by size of allocation and then by type of usage. Use the show system virtual-memory command for troubleshooting with Juniper Networks Customer Support.
Options	none —Display kernel dynamic memory usage information. all-chassis —(TX Matrix routers and TX Matrix Plus routers only) (Optional) Display kernel dynamic memory usage information for all chassis. all-lcc —(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display kernel dynamic memory usage information for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, display kernel dynamic memory usage information for all connected T1600 or T4000 LCCs.

all-members—(EX4200 switches and MX Series routers only) (Optional) Display kernel dynamic memory usage information for all members of the Virtual Chassis configuration.

infrastructure *name*—(QFabric systems only) (Optional) Display kernel dynamic memory usage information for the fabric control Routing Engine and fabric manager Routing Engine.

interconnect-device *name*—(QFabric systems only) (Optional) Display kernel dynamic memory usage information for the Interconnect device.

lcc *number*—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display kernel dynamic memory usage information for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display kernel dynamic memory usage information for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(EX4200 switches and MX Series routers only) (Optional) Display kernel dynamic memory usage information for the local Virtual Chassis member.

member *member-id*—(EX4200 switches and MX Series routers only) (Optional) Display kernel dynamic memory usage information for the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

node-group *name*—(QFabric systems only) (Optional) Display kernel dynamic memory usage information for the Node group.

scc—(TX Matrix routers only) (Optional) Display kernel dynamic memory usage information for the TX Matrix router (or switch-card chassis).

sfc *number*—(TX Matrix Plus routers only) (Optional) Display kernel dynamic memory usage information for the TX Matrix Plus router. Replace *number* with 0.

Additional Information By default, when you issue the **show system virtual-memory** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix

or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.



NOTE: The `show system virtual-memory` command with the `| display XML` pipe option now displays XML output for the command in the parent tags: `<vmstat-memstat-malloc>`, `<vmstat-memstat-zone>`, `<vmstat-sumstat>`, `<vmstat-intr>`, and `<vmstat-kernel-state>` with each child element as a separate XML tag. In Junos OS Releases 10.1 and earlier, the `| display XML` option for this command does not have an XML API element and the entire output is displayed in a single `<output>` tag element.

Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• Routing Matrix with a TX Matrix Plus Router Solutions Page
List of Sample Output	show system virtual-memory on page 1199 show system virtual-memory scc (TX Matrix Router) on page 1203 show system virtual-memory sfc (TX Matrix Plus Router) on page 1204 show system virtual-memory display xml on page 1207 show system virtual-memory (QFX Series) on page 1230
Output Fields	Table 82 lists the output fields for the <code>show system virtual-memory</code> command. Output fields are listed in the approximate order in which they appear.

Table 82: show system virtual-memory Output Fields

Field Name	Field Description
Memory statistics by bucket size	
Size	Memory block size (bytes). The kernel memory allocator appropriates blocks of memory whose size is exactly a power of 2.
In Use	Number of memory blocks of this size that are in use (bytes).
Free	Number of memory blocks of this size that are free (bytes).
Requests	Number of memory allocation requests made.
HighWater	Maximum value the free list can have. Once the system starts reclaiming physical memory, it continues until the free list is increased to this value.
Couldfree	Total number of times that the free elements for a bucket size exceed the high-water mark for that bucket size.
Memory usage type by bucket size	
Size	Memory block size (bytes).
Type(s)	Kernel modules that are using these memory blocks. For a definition of each type, refer to a FreeBSD book.
Memory statistics by type	
Type	Kernel module that is using dynamic memory.
InUse	Number of memory blocks used by this type. The number is rounded up.
MemUse	Amount of memory in use, in kilobytes (KB).
HighUse	Maximum memory ever used by this type.
Limit	Maximum memory that can be allocated to this type.
Requests	Total number of dynamic memory allocation requests this type has made.
Type Limit	Number of times requests were blocked for reaching the maximum limit.
Kern Limit	Number of times requests were blocked for the kernel map.
Size(s)	Memory block sizes this type is using.
Memory Totals	
In Use	Total kernel dynamic memory in use (bytes, rounded up).
Free	Total kernel dynamic memory free (bytes, rounded up).

Table 82: show system virtual-memory Output Fields (*continued*)

Field Name	Field Description
Requests	Total number of memory allocation requests.
ITEM	Kernel module that is using memory.
Size	Memory block size (bytes).
Limit	Maximum memory that can be allocated to this type.
Used	Number of memory blocks used by this type. The number is rounded up.
Free	Number of memory blocks available to this type.
Requests	Total number of memory allocation requests this type has made.
interrupt	Timer events and scheduling interruptions.
total	Total number of interruptions for each type.
rate	Interruption rate.
Total	Total for all interruptions.

Sample Output

show system virtual-memory

```

user@host> show system virtual-memory
Memory statistics by bucket size
Size    In Use    Free    Requests  HighWater  Couldfree
16      906      118     154876    1280       0
32      455      313     209956    640        0
64      4412     260     75380     320        20
128     3200     32      19361     160        81
256     1510     10      8844      80         4
512     446      2       5085      40         0
1K      18       2       5901      20         0
2K      1128     2       4445      10        1368
4K      185      1       456       5          0
8K      5        1       2653      5          0
16K     181      0       233       5          0
32K     2        0       1848      5          0
64K     20       0       22        5          0
128K    5        0       5         5          0
256K    2        0       2         5          0
512K    1        0       1         5          0

Memory usage type by bucket size
Size    Type(s)
16      uc_devlist, nexusdev, iftable, temp, devbuf, atexit, COS, BPF,
        DEVFS mount, DEVFS node, vnodes, mount, pcb, soname, proc-args, kld,
        MD disk, rman, ATA generic, bus, sysctl, ippool, pfestat, ifstate,

```

```

pfe_ipc, mkey, rtable, ifmaddr, ipfw, rnode
32 atkbddev, dirrem, mkdir, diradd, freefile, freefrag, indirdep,
bmsafemap, newblk, temp, devbuf, COS, vnodes, cluster_save buffer,
pcb, soname, proc-args, sigio, kld, Gzip trees, taskqueue, SWAP,
eventhandler, bus, sysctl, uidinfo, subproc, pgrp, pfestat, itable32,
ifstate, pfe_ipc, mkey, rtable, ifmaddr, ipfw, rnode, rtnexthop
64 isadev, iftable, MFS node, allocindir, allocdirect, pagedep, temp,
devbuf, lockf, COS, NULLFS hash, DEVFS name, vnodes,
cluster_save buffer, vfscache, pcb, soname, proc-args, file,
AR driver, AD driver, Gzip trees, rman, eventhandler, bus, sysctl,
subproc, pfestat, pic, ifstate, pfe_ipc, mkey, ifaddr, rtable, ipfw
128 ZONE, freeblks, inodedep, temp, devbuf, zombie, COS, DEVFS node,
vnodes, mount, vfscache, pcb, soname, proc-args, ttys, dev_t,
timecounter, kld, Gzip trees, ISOFS node, bus, uidinfo, cred,
session, pic, itable16, ifstate, pfe_ipc, rtable, ifstat, metrics,
rtnexthop, iffamily
256 iflogical, iftable, MFS node, FFS node, newblk, temp, devbuf,
NFS daemon, vnodes, proc-args, kqueue, file desc, Gzip trees, bus,
subproc, itable16, ifstate, pfe_ipc, sysctl, rtnexthop
512 UFS mount, temp, devbuf, mount, BIO buffer, ptys, ttys, AR driver,
Gzip trees, ISOFS mount, msg, ioctlops, ATA generic, bus, proc,
pfestat, lr, ifstate, pfe_ipc, rtable, ipfw, ifstat, rtnexthop
1K iftable, temp, devbuf, NQNFS Lease, kqueue, kld, AD driver,
Gzip trees, sem, MD disk, bus, ifstate, pfe_ipc, ipfw
2K uc_devlist, UFS mount, temp, devbuf, BIO buffer, pcb, AR driver,
Gzip trees, ioctlops, bus, ipfw, ifstat, rcache
4K memdesc, iftable, UFS mount, temp, devbuf, kld, Gzip trees, sem, msg
8K temp, devbuf, syncache, Gzip trees
16K indirdep, temp, devbuf, shm, msg
32K pagedep, kld, Gzip trees
64K VM pgdata, devbuf, MSDOSFS mount
128K UFS ihash, inodedep, NFS hash, kld, ISOFS mount
256K mbuf, vfscache
512K SWAP

```

Memory statistics by type					Type	Kern		
Type	InUse	MemUse	HighUse	Limit	Requests	Limit	Limit	Size(s)
isadev	13	1K	1K127753K	13	0	0	0	64
atkbddev	2	1K	1K127753K	2	0	0	0	32
uc_devlist	24	3K	3K127753K	24	0	0	0	16,2K
nexusdev	3	1K	1K127753K	3	0	0	0	16
memdesc	1	4K	4K127753K	1	0	0	0	4K
mbuf	1	152K	152K127753K	1	0	0	0	256K
iflogical	6	2K	2K127753K	6	0	0	0	256
iftable	17	9K	9K127753K	18	0	0	0	16,64,256,1K,4K
ZONE	15	2K	2K127753K	15	0	0	0	128
VM pgdata	1	64K	64K127753K	1	0	0	0	64K
UFS mount	12	26K	26K127753K	12	0	0	0	512,2K,4K
UFS ihash	1	128K	128K127753K	1	0	0	0	128K
MFS node	6	2K	3K127753K	35	0	0	0	64,256
FFS node	906	227K	227K127753K	1352	0	0	0	256
dirrem	0	0K	4K127753K	500	0	0	0	32
mkdir	0	0K	1K127753K	38	0	0	0	32
diradd	0	0K	6K127753K	521	0	0	0	32
freefile	0	0K	4K127753K	374	0	0	0	32
freeblks	0	0K	8K127753K	219	0	0	0	128
freefrag	0	0K	1K127753K	193	0	0	0	32
allocindir	0	0K	25K127753K	1518	0	0	0	64
indirdep	0	0K	17K127753K	76	0	0	0	32,16K
allocdirect	0	0K	10K127753K	760	0	0	0	64
bmsafemap	0	0K	1K127753K	72	0	0	0	32

newblk	1	1K	1K127753K	2279	0	0	32,256
inodedep	1	128K	175K127753K	2367	0	0	128,128K
pagedep	1	32K	33K127753K	47	0	0	64,32K
temp	1239	92K	96K127753K	8364	0	0	16,32,64K
devbuf	1413	5527K	5527K127753K	1535	0	0	16,32,64,128,256
lockf	38	3K	3K127753K	2906	0	0	64
atexit	1	1K	1K127753K	1	0	0	16
zombie	0	0K	2K127753K	3850	0	0	128
NFS hash	1	128K	128K127753K	1	0	0	128K
NQNFS Lease	1	1K	1K127753K	1	0	0	1K
NFS daemon	1	1K	1K127753K	1	0	0	256
syncache	1	8K	8K127753K	1	0	0	8K
COS	353	44K	44K127753K	353	0	0	16,32,64,128
BPF	189	3K	3K127753K	189	0	0	16
MSDOSFS mount	1	64K	64K127753K	1	0	0	64K
NULLFS hash	1	1K	1K127753K	1	0	0	64
DEVFS mount	2	1K	1K127753K	2	0	0	16
DEVFS name	487	31K	31K127753K	487	0	0	64
DEVFS node	471	58K	58K127753K	479	0	0	16,128
vnodes	28	7K	7K127753K	429	0	0	16,32,64,128,256
mount	15	8K	8K127753K	18	0	0	16,128,512
cluster_save buffer	0	0K	1K127753K	55	0	0	32,64
vfscache	1898	376K	376K127753K	3228	0	0	64,128,256K
BIO buffer	49	98K	398K127753K	495	0	0	512,2K
pcb	159	16K	17K127753K	399	0	0	16,32,64,128,2K
soname	82	10K	10K127753K	42847	0	0	16,32,64,128
proc-args	57	2K	3K127753K	2105	0	0	16,32,64,128,256
ptys	32	16K	16K127753K	32	0	0	512
ttys	254	33K	33K127753K	522	0	0	128,512
kqueue	5	3K	4K127753K	23	0	0	256,1K
sigio	1	1K	1K127753K	27	0	0	32
file	383	24K	24K127753K	16060	0	0	64
file desc	76	19K	20K127753K	3968	0	0	256
shm	1	12K	12K127753K	1	0	0	16K
dev_t	286	36K	36K127753K	286	0	0	128
timecounter	10	2K	2K127753K	10	0	0	128
kld	11	117K	122K127753K	34	0	0	16,32,128,1K,4K
AR driver	1	1K	3K127753K	5	0	0	64,512,2K
AD driver	2	2K	3K127753K	2755	0	0	64,1K
Gzip trees	0	0K	46K127753K	133848	0	0	32,64,128,256
ISOFS node	1136	142K	142K127753K	1189	0	0	128
ISOFS mount	9	132K	132K127753K	10	0	0	512,128K
sem	3	6K	6K127753K	3	0	0	1K,4K
MD disk	2	2K	2K127753K	2	0	0	16,1K
msg	4	25K	25K127753K	4	0	0	512,4K,16K
rman	59	4K	4K127753K	461	0	0	16,64
ioctlops	0	0K	2K127753K	992	0	0	512,2K
taskqueue	2	1K	1K127753K	2	0	0	32
SWAP	2	413K	413K127753K	2	0	0	32,512K
ATA generic	6	3K	3K127753K	6	0	0	16,512
eventhandler	17	1K	1K127753K	17	0	0	32,64
bus	340	30K	31K127753K	794	0	0	16,32,64,128,256
sysctl	0	0K	1K127753K	130262	0	0	16,32,64
uidinfo	4	1K	1K127753K	10	0	0	32,128
cred	22	3K	3K127753K	3450	0	0	128
subproc	156	10K	10K127753K	7882	0	0	32,64,256
proc	2	1K	1K127753K	2	0	0	512
session	12	2K	2K127753K	34	0	0	128
pgrp	16	1K	1K127753K	45	0	0	32
ippool	1	1K	1K127753K	1	0	0	16
pfestat	0	0K	1K127753K	47349	0	0	16,32,64,512

pic	5	1K	1K127753K	5	0	0	64,128
lr	1	1K	1K127753K	1	0	0	512
itable32	110	4K	4K127753K	110	0	0	32
itable16	161	26K	26K127753K	161	0	0	128,256
ifstate	694	159K	160K127753K	1735	0	0	16,32,64,128,1K
pfe_ipc	0	0K	1K127753K	56218	0	0	16,32,64,128,1K
mkey	250	4K	4K127753K	824	0	0	16,32,64
ifaddr	9	1K	1K127753K	9	0	0	64
sysctl	0	0K	1K127753K	30	0	0	256
rtable	49	6K	6K127753K	307	0	0	16,32,64,128,512
ifmaddr	22	1K	1K127753K	22	0	0	16,32
ipfw	23	10K	10K127753K	48	0	0	16,32,64,512,2K
ifstat	698	805K	805K127753K	698	0	0	128,512,2K
rcache	4	8K	8K127753K	4	0	0	2K
rnode	27	1K	1K127753K	285	0	0	16,32
metrics	1	1K	1K127753K	3	0	0	128
rtnexthop	57	9K	9K127753K	312	0	0	32,128,256,512
iffamily	12	2K	2K127753K	12	0	0	128

Memory Totals:	In Use	Free	Requests
	9311K	54K	489068

ITEM	SIZE	LIMIT	USED	FREE	REQUESTS
PIPE:	192,	0,	4,	81,	4422
SWAPMETA:	160,	95814,	0,	0,	0
unpcb:	160,	0,	114,	36,	279
ripcb:	192,	25330,	5,	37,	5
syncache:	128,	15359,	0,	64,	5
tcpcb:	576,	25330,	23,	12,	32
udpcb:	192,	25330,	14,	28,	255
socket:	256,	25330,	246,	26,	819
KNOTE:	96,	0,	27,	57,	71
NFSNODE:	352,	0,	0,	0,	0
NFSMOUNT:	544,	0,	0,	0,	0
VNODE:	224,	0,	2778,	43,	2778
NAMEI:	1024,	0,	0,	8,	40725
VMSPACE:	192,	0,	57,	71,	3906
PROC:	448,	0,	73,	17,	3923
DP fakepg:	64,	0,	0,	0,	0
PV ENTRY:	28,	499566,	44530,	152053,	1525141
MAP ENTRY:	48,	0,	1439,	134,	351075
KMAP ENTRY:	48,	35645,	179,	119,	10904
MAP:	108,	0,	7,	3,	7
VM OBJECT:	92,	0,	2575,	109,	66912

```

792644 cpu context switches
9863474 device interrupts
286510 software interrupts
390851 traps
3596829 system calls
  16 kernel threads created
 3880 fork() calls
   27 vfork() calls
    0 rfork() calls
    0 swap pager pageins
    0 swap pager pages paged in
    0 swap pager pageouts
    0 swap pager pages paged out
  380 vnode pager pageins
  395 vnode pager pages paged in
  122 vnode pager pageouts

```

```

1476 vnode pager pages paged out
    0 page daemon wakeups
    0 pages examined by the page daemon
101 pages reactivated
161722 copy-on-write faults
    0 copy-on-write optimized faults
84623 zero fill pages zeroed
83063 zero fill pages prezeroed
    7 intransit blocking page faults
535606 total VM faults taken
    0 pages affected by kernel thread creation
238254 pages affected by fork()
    2535 pages affected by vfork()
    0 pages affected by rfork()
283379 pages freed
    0 pages freed by daemon
190091 pages freed by exiting processes
17458 pages active
29166 pages inactive
    0 pages in VM cache
10395 pages wired down
134610 pages free
    4096 bytes per page
183419 total name lookups
    cache hits (90% pos + 7% neg) system 0% per-directory
    deletions 0%, falsehits 0%, toolong 0%

```

interrupt	total	rate
ata0 irq14	113338	3
mux irq7	727643	21
fxp1 irq10	1178671	34
sio0 irq4	833	0
clk irq0	3439769	99
rtc irq8	4403221	127
Total	9863475	286

```

Kernel direct memory map:
    4423 pages used
    4057340 pages maximum

```

Note: Kernel direct memory map only displays for 64 bit platform.

show system virtual-memory scc (TX Matrix Router)

```
user@host> show system virtual-memory scc
```

Memory statistics by bucket size

Size	In Use	Free	Requests	HighWater	Couldfree
16	898	126	749493	1280	0
32	2018	1310	980643	640	632
64	3490	13342	935420	320	5365

...

Memory usage type by bucket size

Size	Type(s)
16	uc_devlist, COS, BPF, DEVFS mount, DEVFS node, vnodes, mount, pcb, soname, rman, bus, sysctl, ifstate, pfe_ipc, mkey, socket, rtable, ifmaddr, ipfw, rnode, iftable, temp, devbuf, atexit, proc-args, kld, MD disk
32	atkbddev, Gzip trees, dirrem, mkdir, diradd, freefile, freefrag, indirdep, bmsafemap, newblk, tseg_qent, COS, vnodes,

...

```

Memory statistics by type
      Type InUse MemUse HighUse Limit Requests Limit Limit Size(s)
      isadev 12 1K 1K166400K 12 0 0 64
      atkbdev 2 1K 1K166400K 2 0 0 32
      uc_devlist 24 3K 3K166400K 24 0 0 16,2K
      ....

Memory Totals: In Use Free Requests
                6091K 1554K 2897122

```

show system virtual-memory sfc (TX Matrix Plus Router)

```

user@host> show system virtual-memory sfc 0
sfc0-re0:

```

```

-----
      Type InUse MemUse HighUse Requests Size(s)
CAM dev queue 1 1K - 1 64
  entropy 1024 64K - 1024 64
  linker 487 6272K - 1163 16,32,64,4096,32768,131072
  USB 127 10K - 127 16,32,64,128,256,1024,2048
  lockf 46 3K - 98418 64
  USBdev 10 2K - 34 16,128,2048,16384
ifstateSLLNode 0 0K - 1096 16
  devbuf 21243 15683K - 21810
16,32,64,128,256,512,1024,2048,4096,8192,16384,32768,65536,131072
  temp 1283 151K - 2483472
16,32,64,128,256,512,2048,4096,8192,16384,32768,65536,131072
  ip6ndp 0 0K - 4 64
  in6ifmulti 1 1K - 1 64
  in6grentry 1 1K - 1 64
  iflogical 20 5K - 29 2048
  iffamily 45 6K - 69 32,1024,2048
  rtnexthop 266 46K - 608013 32,256,512,1024,2048,4096
  metrics 31 4K - 54 256
  rnode 212 4K - 607848 16,32
  rcache 4 8K - 4 65536
  iflist 0 0K - 6 16,64
  ifdevice 11 8K - 17 16,32768
  ifstat 424 472K - 427 512,16384,65536
  ipfw 42 23K - 145
16,32,64,128,256,512,1024,16384,32768,65536,131072
  ifmaddr 415 11K - 415 16,32
  rtable 329 28K - 608066 16,32,64,128,1024,16384
  sysctl 0 0K - 887976 16,32,64,4096,16384,32768
  ifaddr 64 5K - 70 32,64,128
  mkey 331 6K - 12528 16,128
  pfe_ipc 0 0K - 7299115
16,32,64,128,256,512,1024,2048,4096,8192,16384,32768,65536,131072
  ifstate 1245054 70088K - 3040437
16,32,64,128,256,512,1024,2048,4096,8192,16384,32768
  idxbucket 1 1K - 1 16
  itable16 5069 1250K - 5103 1024,4096
  itable32 157 10K - 157 64
  itable64 2 1K - 2 128
  lr 1 1K - 4 16384
  pic 37 6K - 37 64,16384
  pfestat 0 0K - 6220 32,64,128,256,131072
  gencfg 1486 424K - 2614 16,32,64,256,512,16384,32768,65536

```

```

        jsr      2      1K      -      22  16
        idl      1      4K      -      165
32,64,128,256,512,1024,2048,8192,16384,32768,65536,131072
        rtmsg    0      0K      -      16  131072
        module  250     16K      -      250  64,128
        mtx_pool 1      8K      -      1   64,128
        DEVFS3   113     13K      -      114  256
        DEVFS1   106     24K      -      106  2048
        pgrp     15      1K      -      8600 64
        session  11      2K      -      2829 512
        proc      2      1K      -      2   16384
        subproc  296     572K     -      24689 2048,131072
        cred      38      5K      -      619244 256
        plimit    18      4K      -      21311 2048
        uidinfo   3      1K      -      10   32,512
        sysctlold 2701     82K     -      2701 16,32,64
        sysctltmp 0      0K      -      15572 16,32,64,1024
        umtx     171     11K      -      171   64
        SWAP      2     277K     -      2     64
        bus      779     125K     -      3072 16,32,64,128,32768
        bus-sc    67      62K     -      1477
16,32,64,512,1024,2048,8192,16384,65536,131072
        devstat   8      17K     -      8   16,131072
        eventhandler 46     2K     -      47   32,128
        kobj      93     186K     -      111  65536
        DEVFS      8      1K     -      9   16,64
        rman     106      7K     -      490 16,32,64
        sbuf       0      0K     -      28234 16,32,32768,131072
...
lcc0-re0:

```

```

-----
      Type InUse MemUse HighUse Requests Size(s)
CAM dev queue    1     1K      -      1     64
      entropy  1024    64K      -     1024     64
      linker   487   6272K     -     1163 16,32,64,4096,32768,131072
      USB     127    10K      -      127 16,32,64,128,256,1024,2048
      lockf    23     2K      -    169585     64
      USBdev   10     2K      -       34 16,128,2048,16384
      devbuf  5128  10760K     -      5310
16,32,64,128,256,512,1024,2048,4096,8192,16384,32768,65536,131072
      temp   1285    151K     -     10770
16,32,64,128,256,512,2048,4096,8192,16384,32768,65536,131072
      ip6ndp    0      0K      -        4     64
      iflogical 20      5K      -       29  2048
      iffamilly 45      6K      -       69 32,1024,2048
      rtnexthop 189     29K     -    1211988 32,256,512,1024,2048,4096
      metrics   11      2K      -       16  256
      rnode    135      3K     -     606391 16,32
      rcache     4      8K      -        4  65536
      iflist     0      0K      -        6  16,64
      ifdevice  11      8K      -       17 16,32768
      ifstat   412    471K     -      415 512,16384,65536
      ipfw      42     23K      -        91
16,32,64,128,256,512,1024,16384,32768,65536,131072
      ifmaddr  415     11K      -       415 16,32
      rtable   225     20K      -     606584 16,32,64,128,1024,16384
      sysctl    0      0K      -    2302479 16,32,64
      ifaddr    53      4K      -        69 32,64,128
      mkey     133      3K      -     8974 16,128
      pfe_ipc    0      0K      -    19035108
16,32,64,128,512,1024,2048,8192,16384,32768,65536,131072

```

```

    ifstate 710270 42176K - 9583703
16,32,64,128,256,512,1024,2048,8192,16384,32768
    idxbucket 1 1K - 1 16
    itable16 5045 1245K - 1825178 1024,4096
    itable32 157 10K - 157 64
    itable64 2 1K - 2 128
        lr 1 1K - 4 16384
        pic 37 6K - 37 64,16384
    pfestat 0 0K - 1682 32,64,128,256,131072
    gencfg 1486 424K - 2812 16,32,64,256,512,16384,32768,65536
    jsr 0 0K - 22 16
    idl 0 0K - 4 32768,131072
    rtsmsg 0 0K - 3 131072
    module 250 16K - 250 64,128
    mtx_pool 1 8K - 1 64,128
    DEVFS3 108 12K - 109 256
    DEVFS1 101 23K - 101 2048
    pgrp 5 1K - 917 64
    session 5 1K - 917 512
    proc 2 1K - 2 16384
    subproc 217 441K - 4867 2048,131072
    cred 21 3K - 48719 256
    plimit 9 2K - 5255 2048
    uidinfo 2 1K - 2 32,512
    sysctluid 2786 85K - 2786 16,32,64
    sysctltmp 0 0K - 1833 16,32,64,1024
    umtx 126 8K - 126 64
    SWAP 2 277K - 2 64
    bus 780 125K - 2734 16,32,64,128,32768
    bus-sc 69 69K - 1194
16,32,64,512,1024,2048,8192,16384,65536,131072
    devstat 8 17K - 8 16,131072
    eventhandler 45 2K - 46 32,128
    kobj 93 186K - 111 65536
    DEVFS 8 1K - 9 16,64
    rman 94 6K - 477 16,32,64
    sbuf 0 0K - 532 16,32,32768,131072
    NULLFS hash 1 1K - 1 64
    taskqueue 5 1K - 5 64
    turnstiles 127 8K - 127 64
    Unitno 6 1K - 44 16,64
    ioctlops 0 0K - 1771718 16,32,64,128,8192,16384,65536,131072

    iov 0 0K - 79425 16,64,128,256,512,1024,2048,131072
    msg 4 25K - 4 32768,131072
    sem 4 7K - 4 16384,32768,131072
    shm 2 13K - 4 32768
    ttys 93 16K - 195 512,32768
    soname 31 3K - 389284 16,32,64,256
    pcb 101 16K - 4374
16,32,64,128,1024,2048,4096,16384,65536
    BIO buffer 40 80K - 750 65536
    vfscache 1 512K - 1 65536
    cluster_save buffer 0 OK - 55 32,64
    VFS hash 1 256K - 1 32,64
    vnodes 1 1K - 1 512
    mount 266 21K - 481 16,32,64,128,256,4096,32768
    vnodemarker 0 0K - 2497 16384
    pfs_nodes 25 3K - 25 128
    pfs_vncache 144 5K - 386 32
    STP 1 1K - 1 64

```

```

        GEOM      173      15K      -      1068
16,32,64,128,256,512,2048,16384,32768,131072
        syncache   1       8K      -       1
16,32,64,128,256,512,2048,16384,32768,131072
        tlv_stat    0       0K      -      223
16,32,64,128,256,512,2048,16384,32768,131072
        NFS daemon  1       8K      -       1
16,32,64,128,256,512,2048,16384,32768,131072
        p1003.1b    1       1K      -       1  16
        MD disk     9      18K      -       9 65536
        ata_generic  2       2K      -      25 16,16384,32768
        ISOFS mount  7       1K      -      13 512
        ISOFS node 1439    135K     -    1453 128
        CAM SIM      1       1K      -       1 64
        CAM XPT      6       1K      -       9 16,64,16384
        CAM periph   1       1K      -       1 128
        ad_driver    2       1K      -       2 256
        pagedep      1      64K      -     105 64
        inodedep     1     256K     -     552 256
        newblk       1       1K      -     327 64,4096
        bmsafemap    0       0K      -      19 64
        allocdirect  0       0K      -     326 128
        freefrag     0       0K      -      31 32
        freeblks     0       0K      -     103 2048
        freefile     0       0K      -     175 32
        diradd       0       0K      -     590 64
        mkdir        0       0K      -     166 32
        dirrem       0       0K      -     382 32
        savedino     0       0K      -     283 512
        UFS mount    15     36K      -      15 2048,65536,131072
        ata_dma      6       1K      -       6 256
        UMAHash      1       4K      -       5 4096,16384,32768,65536,131072
        cdev         26      3K      -      26 256
        file desc    111    25K      -    5199 16,1024,2048,16384
        VM pgdata    2     65K      -       2 64
        sigio        1       1K      -      27 32
        kenv         30      5K      -      33 16,32,64,131072
        atkbddev     2       1K      -       2 32
        kqueue       0       0K      -      88 1024,4096,32768
        proc-args    28      2K      -    3970 32,64,128,256,512,1024
        isadev       23      2K      -      23 64
        zombie       1       1K      -    4651 128
        ithread      92      7K      -      92 16,64,256
        legacydrv     3       1K      -       3 16
        memdesc      1       4K      -       1 131072
        nexusdev     2       1K      -       2 16
        CAM queue    3       1K      -       3 16
        KTRACE       100    10K      -     100 128
        kbdmux       5       9K      -       5 128,2048,65536,131072
ITEM      SIZE      LIMIT      USED      FREE  REQUESTS
UMA Kegs:  136,      0,      71,      1,      71
...
```

show system virtual-memory | display xml

```

user@host> show system virtual-memory | display xml
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/10.2R1/junos">
  <system-virtual-memory-information>
    <vmstat-memstat-malloc>
      <memstat-name>CAM dev queue</memstat-name>
      <inuse>1</inuse>
    </vmstat-memstat-malloc>
  </system-virtual-memory-information>
</rpc-reply>
```

```
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>entropy</memstat-name>
<inuse>1024</inuse>
<memuse>64</memuse>
<high-use>--</high-use>
<memstat-req>1024</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>linker</memstat-name>
<inuse>481</inuse>
<memuse>1871</memuse>
<high-use>--</high-use>
<memstat-req>1145</memstat-req>
<memstat-size>16,32,64,4096,32768,131072</memstat-size>
<memstat-name>lockf</memstat-name>
<inuse>56</inuse>
<memuse>4</memuse>
<high-use>--</high-use>
<memstat-req>5998</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>devbuf</memstat-name>
<inuse>2094</inuse>
<memuse>3877</memuse>
<high-use>--</high-use>
<memstat-req>2099</memstat-req>

<memstat-size>16,32,64,128,512,1024,4096,8192,16384,32768,65536,131072</memstat-size>

<memstat-name>temp</memstat-name>
<inuse>21</inuse>
<memuse>66</memuse>
<high-use>--</high-use>
<memstat-req>3127</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,4096,8192,16384,32768,65536,131072</memstat-size>

<memstat-name>ip6ndp</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>in6ifmulti</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>in6grentry</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>iflogical</memstat-name>
<inuse>13</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>13</memstat-req>
```



```

<memstat-size>64,2048</memstat-size>
<memstat-name>iffamily</memstat-name>
<inuse>28</inuse>
<memuse>4</memuse>
<high-use>--</high-use>
<memstat-req>28</memstat-req>
<memstat-size>32,1024,2048</memstat-size>
<memstat-name>rtnexthop</memstat-name>
<inuse>127</inuse>
<memuse>18</memuse>
<high-use>--</high-use>
<memstat-req>129</memstat-req>
<memstat-size>32,256,512,1024,2048,4096</memstat-size>
<memstat-name>metrics</memstat-name>
<inuse>3</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>5</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>inifmulti</memstat-name>
<inuse>3</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>3</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>ingrentry</memstat-name>
<inuse>6</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>6</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>rnode</memstat-name>
<inuse>68</inuse>
<memuse>2</memuse>
<high-use>--</high-use>
<memstat-req>76</memstat-req>
<memstat-size>16,32</memstat-size>
<memstat-name>rcache</memstat-name>
<inuse>4</inuse>
<memuse>8</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>ifdevice</memstat-name>
<inuse>4</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>ifstat</memstat-name>
<inuse>40</inuse>
<memuse>22</memuse>
<high-use>--</high-use>
<memstat-req>40</memstat-req>
<memstat-size>512,16384,32768</memstat-size>
<memstat-name>ipfw</memstat-name>
<inuse>42</inuse>
<memuse>23</memuse>
<high-use>--</high-use>
<memstat-req>91</memstat-req>

```

```

<memstat-size>16,32,64,128,256,512,1024,16384,32768,65536,131072</memstat-size>
  <memstat-name>ifmaddr</memstat-name>
  <inuse>103</inuse>
  <memuse>3</memuse>
  <high-use>--</high-use>
  <memstat-req>103</memstat-req>
  <memstat-size>16,32</memstat-size>
  <memstat-name>rtable</memstat-name>
  <inuse>129</inuse>
  <memuse>14</memuse>
  <high-use>--</high-use>
  <memstat-req>139</memstat-req>
  <memstat-size>16,32,64,128,1024,16384</memstat-size>
  <memstat-name>sysctl</memstat-name>
  <inuse>0</inuse>
  <memuse>0</memuse>
  <high-use>--</high-use>
  <memstat-req>14847</memstat-req>
  <memstat-size>16,32,64,4096,16384,32768</memstat-size>
  <memstat-name>ifaddr</memstat-name>
  <inuse>29</inuse>
  <memuse>3</memuse>
  <high-use>--</high-use>
  <memstat-req>29</memstat-req>
  <memstat-size>64,128</memstat-size>
  <memstat-name>mkey</memstat-name>
  <inuse>345</inuse>
  <memuse>6</memuse>
  <high-use>--</high-use>
  <memstat-req>2527</memstat-req>
  <memstat-size>16,128</memstat-size>
  <memstat-name>pfe_ipc</memstat-name>
  <inuse>0</inuse>
  <memuse>0</memuse>
  <high-use>--</high-use>
  <memstat-req>1422</memstat-req>

<memstat-size>16,32,64,128,512,1024,2048,8192,16384,32768,65536,131072</memstat-size>
  <memstat-name>ifstate</memstat-name>
  <inuse>594</inuse>
  <memuse>51</memuse>
  <high-use>--</high-use>
  <memstat-req>655</memstat-req>

<memstat-size>16,32,64,128,256,1024,2048,4096,16384,32768</memstat-size>
  <memstat-name>itable16</memstat-name>
  <inuse>276</inuse>
  <memuse>52</memuse>
  <high-use>--</high-use>
  <memstat-req>294</memstat-req>
  <memstat-size>1024,4096</memstat-size>
  <memstat-name>itable32</memstat-name>
  <inuse>160</inuse>
  <memuse>10</memuse>
  <high-use>--</high-use>
  <memstat-req>160</memstat-req>
  <memstat-size>64</memstat-size>
  <memstat-name>itable64</memstat-name>
  <inuse>2</inuse>
  <memuse>1</memuse>

```

```

<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>lr</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>16384</memstat-size>
<memstat-name>pic</memstat-name>
<inuse>5</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>5</memstat-req>
<memstat-size>64,512</memstat-size>
<memstat-name>pfestat</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>162</memstat-req>
<memstat-size>16,32,128,256,16384</memstat-size>
<memstat-name>gencfg</memstat-name>
<inuse>224</inuse>
<memuse>56</memuse>
<high-use>--</high-use>
<memstat-req>540</memstat-req>
<memstat-size>16,32,64,256,512,32768,65536</memstat-size>
<memstat-name>jsr</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>idl</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>13</memstat-req>
<memstat-size>16,32,64,128,256,4096,16384,32768,131072</memstat-size>

<memstat-name>rtsmsg</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>131072</memstat-size>
<memstat-name>module</memstat-name>
<inuse>249</inuse>
<memuse>16</memuse>
<high-use>--</high-use>
<memstat-req>249</memstat-req>
<memstat-size>64,128</memstat-size>
<memstat-name>mtx_pool</memstat-name>
<inuse>1</inuse>
<memuse>8</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64,128</memstat-size>
<memstat-name>DEVFS3</memstat-name>
<inuse>109</inuse>
<memuse>12</memuse>

```

```
<high-use>--</high-use>
<memstat-req>117</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>DEVFS1</memstat-name>
<inuse>102</inuse>
<memuse>23</memuse>
<high-use>--</high-use>
<memstat-req>109</memstat-req>
<memstat-size>2048</memstat-size>
<memstat-name>pgrp</memstat-name>
<inuse>12</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>21</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>session</memstat-name>
<inuse>8</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>15</memstat-req>
<memstat-size>512</memstat-size>
<memstat-name>proc</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>16384</memstat-size>
<memstat-name>subproc</memstat-name>
<inuse>244</inuse>
<memuse>496</memuse>
<high-use>--</high-use>
<memstat-req>1522</memstat-req>
<memstat-size>2048,131072</memstat-size>
<memstat-name>cred</memstat-name>
<inuse>30</inuse>
<memuse>4</memuse>
<high-use>--</high-use>
<memstat-req>11409</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>plimit</memstat-name>
<inuse>17</inuse>
<memuse>4</memuse>
<high-use>--</high-use>
<memstat-req>133</memstat-req>
<memstat-size>2048</memstat-size>
<memstat-name>uidinfo</memstat-name>
<inuse>3</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>6</memstat-req>
<memstat-size>32,512</memstat-size>
<memstat-name>sysctlold</memstat-name>
<inuse>1117</inuse>
<memuse>34</memuse>
<high-use>--</high-use>
<memstat-req>1117</memstat-req>
<memstat-size>16,32,64</memstat-size>
<memstat-name>sysctltmp</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
```

```

<memstat-req>743</memstat-req>
<memstat-size>16,32,64,1024</memstat-size>
<memstat-name>umtx</memstat-name>
<inuse>144</inuse>
<memuse>9</memuse>
<high-use>--</high-use>
<memstat-req>144</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>SWAP</memstat-name>
<inuse>2</inuse>
<memuse>209</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>bus</memstat-name>
<inuse>496</inuse>
<memuse>55</memuse>
<high-use>--</high-use>
<memstat-req>1196</memstat-req>
<memstat-size>16,32,64,128,32768</memstat-size>
<memstat-name>bus-sc</memstat-name>
<inuse>23</inuse>
<memuse>33</memuse>
<high-use>--</high-use>
<memstat-req>335</memstat-req>

<memstat-size>16,32,64,512,1024,2048,8192,16384,65536,131072</memstat-size>
<memstat-name>devstat</memstat-name>
<inuse>10</inuse>
<memuse>21</memuse>
<high-use>--</high-use>
<memstat-req>10</memstat-req>
<memstat-size>16,131072</memstat-size>
<memstat-name>eventhandler</memstat-name>
<inuse>35</inuse>
<memuse>2</memuse>
<high-use>--</high-use>
<memstat-req>36</memstat-req>
<memstat-size>32,128</memstat-size>
<memstat-name>kobj</memstat-name>
<inuse>93</inuse>
<memuse>186</memuse>
<high-use>--</high-use>
<memstat-req>111</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>DEVFS</memstat-name>
<inuse>8</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>9</memstat-req>
<memstat-size>16,64</memstat-size>
<memstat-name>rman</memstat-name>
<inuse>71</inuse>
<memuse>5</memuse>
<high-use>--</high-use>
<memstat-req>433</memstat-req>
<memstat-size>16,32,64</memstat-size>
<memstat-name>sbuf</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>

```

```
<memstat-req>522</memstat-req>
<memstat-size>16,32,32768,131072</memstat-size>
<memstat-name>NULLFS hash</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>taskqueue</memstat-name>
<inuse>5</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>5</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>turnstiles</memstat-name>
<inuse>145</inuse>
<memuse>10</memuse>
<high-use>--</high-use>
<memstat-req>145</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>Unitno</memstat-name>
<inuse>8</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>44</memstat-req>
<memstat-size>16,64</memstat-size>
<memstat-name>ioctlops</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>27622</memstat-req>
<memstat-size>16,64,8192,16384,131072</memstat-size>
<memstat-name>iov</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>18578</memstat-req>
<memstat-size>16,64,128,256,512,1024,2048,131072</memstat-size>
<memstat-name>msg</memstat-name>
<inuse>4</inuse>
<memuse>25</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>32768,131072</memstat-size>
<memstat-name>sem</memstat-name>
<inuse>4</inuse>
<memuse>7</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>16384,32768,131072</memstat-size>
<memstat-name>shm</memstat-name>
<inuse>9</inuse>
<memuse>20</memuse>
<high-use>--</high-use>
<memstat-req>14</memstat-req>
<memstat-size>32768</memstat-size>
<memstat-name>ttys</memstat-name>
<inuse>321</inuse>
<memuse>61</memuse>
<high-use>--</high-use>
<memstat-req>528</memstat-req>
```

```

<memstat-size>512,32768</memstat-size>
<memstat-name>ptys</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>mbuf_tag</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>23383</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>soname</memstat-name>
<inuse>115</inuse>
<memuse>12</memuse>
<high-use>--</high-use>
<memstat-req>24712</memstat-req>
<memstat-size>16,32,64,256</memstat-size>
<memstat-name>pcb</memstat-name>
<inuse>216</inuse>
<memuse>33</memuse>
<high-use>--</high-use>
<memstat-req>484</memstat-req>

<memstat-size>16,32,64,128,1024,2048,4096,16384,32768,65536</memstat-size>
<memstat-name>BIO buffer</memstat-name>
<inuse>43</inuse>
<memuse>86</memuse>
<high-use>--</high-use>
<memstat-req>405</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>vfscache</memstat-name>
<inuse>1</inuse>
<memuse>256</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>cluster_save buffer</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>32,64</memstat-size>
<memstat-name>VFS hash</memstat-name>
<inuse>1</inuse>
<memuse>128</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>32,64</memstat-size>
<memstat-name>vnodes</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>512</memstat-size>
<memstat-name>mount</memstat-name>
<inuse>290</inuse>
<memuse>23</memuse>
<high-use>--</high-use>
<memstat-req>535</memstat-req>

```

```
<memstat-size>16,32,64,128,256,4096,32768</memstat-size>
<memstat-name>vnodemarker</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>498</memstat-req>
<memstat-size>16384</memstat-size>
<memstat-name>pfs_nodes</memstat-name>
<inuse>25</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>25</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>pfs_vncache</memstat-name>
<inuse>27</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>53</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>STP</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>GEOM</memstat-name>
<inuse>146</inuse>
<memuse>11</memuse>
<high-use>--</high-use>
<memstat-req>1042</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,16384,32768,131072</memstat-size>
<memstat-name>syncache</memstat-name>
<inuse>1</inuse>
<memuse>8</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,16384,32768,131072</memstat-size>
<memstat-name>tlv_stat</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>8</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,16384,32768,131072</memstat-size>
<memstat-name>NFS_daemon</memstat-name>
<inuse>1</inuse>
<memuse>8</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,16384,32768,131072</memstat-size>
<memstat-name>p1003.1b</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>MD_disk</memstat-name>
<inuse>10</inuse>
```



```

<memuse>20</memuse>
<high-use>--</high-use>
<memstat-req>10</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>ata_generic</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>6</memstat-req>
<memstat-size>16,16384,32768</memstat-size>
<memstat-name>ISOFs mount</memstat-name>
<inuse>8</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>15</memstat-req>
<memstat-size>512</memstat-size>
<memstat-name>ISOFs node</memstat-name>
<inuse>1440</inuse>
<memuse>135</memuse>
<high-use>--</high-use>
<memstat-req>1457</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>CAM SIM</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>CAM XPT</memstat-name>
<inuse>6</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>9</memstat-req>
<memstat-size>16,64,16384</memstat-size>
<memstat-name>CAM periph</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>ad_driver</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>pagedep</memstat-name>
<inuse>1</inuse>
<memuse>32</memuse>
<high-use>--</high-use>
<memstat-req>106</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>inodedep</memstat-name>
<inuse>1</inuse>
<memuse>128</memuse>
<high-use>--</high-use>
<memstat-req>464</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>newblk</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>

```

```
<high-use>--</high-use>
<memstat-req>336</memstat-req>
<memstat-size>64,4096</memstat-size>
<memstat-name>bmsafemap</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>63</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>allocdirect</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>320</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>indirdep</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>17</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>allocindir</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>15</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>freefrag</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>12</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>freeblks</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>40</memstat-req>
<memstat-size>2048</memstat-size>
<memstat-name>freefile</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>101</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>diradd</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>465</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>mkdir</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>136</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>dirrem</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
```

```

<memstat-req>168</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>newdirblk</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>saveditino</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>157</memstat-req>
<memstat-size>512</memstat-size>
<memstat-name>UFS mount</memstat-name>
<inuse>15</inuse>
<memuse>36</memuse>
<high-use>--</high-use>
<memstat-req>15</memstat-req>
<memstat-size>2048,65536,131072</memstat-size>
<memstat-name>ata_dma</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>UMAHash</memstat-name>
<inuse>1</inuse>
<memuse>2</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>4096,16384,32768,65536</memstat-size>
<memstat-name>cdev</memstat-name>
<inuse>22</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>22</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>file desc</memstat-name>
<inuse>141</inuse>
<memuse>32</memuse>
<high-use>--</high-use>
<memstat-req>1583</memstat-req>
<memstat-size>16,1024,2048,16384</memstat-size>
<memstat-name>VM pgdata</memstat-name>
<inuse>2</inuse>
<memuse>65</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>sigio</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>20</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>kenv</memstat-name>
<inuse>24</inuse>
<memuse>5</memuse>
<high-use>--</high-use>
<memstat-req>27</memstat-req>

```

```
<memstat-size>16,32,64,131072</memstat-size>
<memstat-name>atkbddev</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>kqueue</memstat-name>
<inuse>15</inuse>
<memuse>9</memuse>
<high-use>--</high-use>
<memstat-req>19</memstat-req>
<memstat-size>1024,4096,32768</memstat-size>
<memstat-name>proc-args</memstat-name>
<inuse>57</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>1001</memstat-req>
<memstat-size>16,32,64,128,256,512,1024</memstat-size>
<memstat-name>isadev</memstat-name>
<inuse>21</inuse>
<memuse>2</memuse>
<high-use>--</high-use>
<memstat-req>21</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>zombie</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>1278</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>ithread</memstat-name>
<inuse>69</inuse>
<memuse>5</memuse>
<high-use>--</high-use>
<memstat-req>69</memstat-req>
<memstat-size>16,64,256</memstat-size>
<memstat-name>legacydrv</memstat-name>
<inuse>4</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>memdesc</memstat-name>
<inuse>1</inuse>
<memuse>4</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>131072</memstat-size>
<memstat-name>nexusdev</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>CAM queue</memstat-name>
<inuse>3</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>3</memstat-req>
<memstat-size>16</memstat-size>
```

```

<memstat-name>$PIR</memstat-name>
<inuse>4</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>KTRACE</memstat-name>
<inuse>100</inuse>
<memuse>10</memuse>
<high-use>--</high-use>
<memstat-req>100</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>kbdmux</memstat-name>
<inuse>5</inuse>
<memuse>9</memuse>
<high-use>--</high-use>
<memstat-req>5</memstat-req>
<memstat-size>128,2048,65536,131072</memstat-size>
</vmstat-memstat-malloc>
<vmstat-memstat-zone>
  <zone-name>UMA Kegs:</zone-name>
  <zone-size>136</zone-size>
  <count-limit>0</count-limit>
  <used>71</used>
  <free>1</free>
  <zone-req>71</zone-req>
  <zone-name>UMA Zones:</zone-name>
  <zone-size>120</zone-size>
  <count-limit>0</count-limit>
  <used>71</used>
  <free>19</free>
  <zone-req>71</zone-req>
  <zone-name>UMA Slabs:</zone-name>
  <zone-size>64</zone-size>
  <count-limit>0</count-limit>
  <used>490</used>
  <free>41</free>
  <zone-req>579</zone-req>
  <zone-name>UMA RCntSlabs:</zone-name>
  <zone-size>104</zone-size>
  <count-limit>0</count-limit>
  <used>276</used>
  <free>20</free>
  <zone-req>276</zone-req>
  <zone-name>UMA Hash:</zone-name>
  <zone-size>128</zone-size>
  <count-limit>0</count-limit>
  <used>4</used>
  <free>26</free>
  <zone-req>5</zone-req>
  <zone-name>16 Bucket:</zone-name>
  <zone-size>76</zone-size>
  <count-limit>0</count-limit>
  <used>30</used>
  <free>20</free>
  <zone-req>30</zone-req>
  <zone-name>32 Bucket:</zone-name>
  <zone-size>140</zone-size>
  <count-limit>0</count-limit>
  <used>33</used>
  <free>23</free>

```

```
<zone-req>33</zone-req>
<zone-name>64 Bucket:</zone-name>
<zone-size>268</zone-size>
<count-limit>0</count-limit>
<used>33</used>
<free>9</free>
<zone-req>33</zone-req>
<zone-name>128 Bucket:</zone-name>
<zone-size>524</zone-size>
<count-limit>0</count-limit>
<used>49</used>
<free>0</free>
<zone-req>49</zone-req>
<zone-name>VM OBJECT:</zone-name>
<zone-size>128</zone-size>
<count-limit>0</count-limit>
<used>2111</used>
<free>79</free>
<zone-req>25214</zone-req>
<zone-name>MAP:</zone-name>
<zone-size>160</zone-size>
<count-limit>0</count-limit>
<used>7</used>
<free>41</free>
<zone-req>7</zone-req>
<zone-name>KMAP ENTRY:</zone-name>
<zone-size>68</zone-size>
<count-limit>35336</count-limit>
<used>19</used>
<free>149</free>
<zone-req>2397</zone-req>
<zone-name>MAP ENTRY:</zone-name>
<zone-size>68</zone-size>
<count-limit>0</count-limit>
<used>2031</used>
<free>153</free>
<zone-req>62417</zone-req>
<zone-name>PV ENTRY:</zone-name>
<zone-size>24</zone-size>
<count-limit>509095</count-limit>
<used>57177</used>
<free>6333</free>
<zone-req>1033683</zone-req>
<zone-name>DP fakepg:</zone-name>
<zone-size>72</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>mt_zone:</zone-name>
<zone-size>64</zone-size>
<count-limit>0</count-limit>
<used>238</used>
<free>57</free>
<zone-req>238</zone-req>
<zone-name>16:</zone-name>
<zone-size>16</zone-size>
<count-limit>0</count-limit>
<used>2114</used>
<free>119</free>
<zone-req>80515</zone-req>
```

```
<zone-name>32:</zone-name>
<zone-size>32</zone-size>
<count-limit>0</count-limit>
<used>1335</used>
<free>134</free>
<zone-req>10259</zone-req>
<zone-name>64:</zone-name>
<zone-size>64</zone-size>
<count-limit>0</count-limit>
<used>3529</used>
<free>129</free>
<zone-req>29110</zone-req>
<zone-name>96:</zone-name>
<zone-size>96</zone-size>
<count-limit>0</count-limit>
<used>2062</used>
<free>58</free>
<zone-req>4365</zone-req>
<zone-name>112:</zone-name>
<zone-size>112</zone-size>
<count-limit>0</count-limit>
<used>361</used>
<free>164</free>
<zone-req>24613</zone-req>
<zone-name>128:</zone-name>
<zone-size>128</zone-size>
<count-limit>0</count-limit>
<used>359</used>
<free>61</free>
<zone-req>942</zone-req>
<zone-name>160:</zone-name>
<zone-size>160</zone-size>
<count-limit>0</count-limit>
<used>364</used>
<free>44</free>
<zone-req>577</zone-req>
<zone-name>224:</zone-name>
<zone-size>224</zone-size>
<count-limit>0</count-limit>
<used>422</used>
<free>20</free>
<zone-req>1950</zone-req>
<zone-name>256:</zone-name>
<zone-size>256</zone-size>
<count-limit>0</count-limit>
<used>204</used>
<free>36</free>
<zone-req>1225</zone-req>
<zone-name>288:</zone-name>
<zone-size>288</zone-size>
<count-limit>0</count-limit>
<used>2</used>
<free>24</free>
<zone-req>10</zone-req>
<zone-name>512:</zone-name>
<zone-size>512</zone-size>
<count-limit>0</count-limit>
<used>49</used>
<free>7</free>
<zone-req>911</zone-req>
<zone-name>1024:</zone-name>
```

```
<zone-size>1024</zone-size>
<count-limit>0</count-limit>
<used>213</used>
<free>11</free>
<zone-req>1076</zone-req>
<zone-name>2048:</zone-name>
<zone-size>2048</zone-size>
<count-limit>0</count-limit>
<used>199</used>
<free>113</free>
<zone-req>640</zone-req>
<zone-name>4096:</zone-name>
<zone-size>4096</zone-size>
<count-limit>0</count-limit>
<used>144</used>
<free>7</free>
<zone-req>2249</zone-req>
<zone-name>Files:</zone-name>
<zone-size>72</zone-size>
<count-limit>0</count-limit>
<used>665</used>
<free>77</free>
<zone-req>16457</zone-req>
<zone-name>MAC labels:</zone-name>
<zone-size>20</zone-size>
<count-limit>0</count-limit>
<used>3998</used>
<free>227</free>
<zone-req>21947</zone-req>
<zone-name>PROC:</zone-name>
<zone-size>544</zone-size>
<count-limit>0</count-limit>
<used>116</used>
<free>10</free>
<zone-req>1394</zone-req>
<zone-name>THREAD:</zone-name>
<zone-size>416</zone-size>
<count-limit>0</count-limit>
<used>127</used>
<free>17</free>
<zone-req>131</zone-req>
<zone-name>KSEGRP:</zone-name>
<zone-size>88</zone-size>
<count-limit>0</count-limit>
<used>127</used>
<free>73</free>
<zone-req>131</zone-req>
<zone-name>UPCALL:</zone-name>
<zone-size>44</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>SLEEPQUEUE:</zone-name>
<zone-size>32</zone-size>
<count-limit>0</count-limit>
<used>145</used>
<free>194</free>
<zone-req>145</zone-req>
<zone-name>VMSPACE:</zone-name>
<zone-size>268</zone-size>
```



```
<count-limit>0</count-limit>
<used>57</used>
<free>13</free>
<zone-req>1335</zone-req>
<zone-name>mbuf_packet:</zone-name>
<zone-size>256</zone-size>
<count-limit>180000</count-limit>
<used>256</used>
<free>128</free>
<zone-req>49791</zone-req>
<zone-name>mbuf:</zone-name>
<zone-size>256</zone-size>
<count-limit>180000</count-limit>
<used>50</used>
<free>466</free>
<zone-req>105183</zone-req>
<zone-name>mbuf_cluster:</zone-name>
<zone-size>2048</zone-size>
<count-limit>25190</count-limit>
<used>387</used>
<free>165</free>
<zone-req>5976</zone-req>
<zone-name>mbuf_jumbo_pagesize:</zone-name>
<zone-size>4096</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>mbuf_jumbo_9k:</zone-name>
<zone-size>9216</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>mbuf_jumbo_16k:</zone-name>
<zone-size>16384</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>ACL UMA zone:</zone-name>
<zone-size>388</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>g_bio:</zone-name>
<zone-size>132</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>174</free>
<zone-req>69750</zone-req>
<zone-name>ata_request:</zone-name>
<zone-size>200</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>57</free>
<zone-req>5030</zone-req>
<zone-name>ata_composite:</zone-name>
<zone-size>192</zone-size>
<count-limit>0</count-limit>
```

```
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>GENCFG:</zone-name>
<zone-size>72</zone-size>
<count-limit>1000004</count-limit>
<used>57</used>
<free>102</free>
<zone-req>57</zone-req>
<zone-name>VNODE:</zone-name>
<zone-size>292</zone-size>
<count-limit>0</count-limit>
<used>2718</used>
<free>25</free>
<zone-req>2922</zone-req>
<zone-name>VNODEPOLL:</zone-name>
<zone-size>72</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>S VFS Cache:</zone-name>
<zone-size>68</zone-size>
<count-limit>0</count-limit>
<used>2500</used>
<free>76</free>
<zone-req>3824</zone-req>
<zone-name>L VFS Cache:</zone-name>
<zone-size>291</zone-size>
<count-limit>0</count-limit>
<used>51</used>
<free>14</free>
<zone-req>63</zone-req>
<zone-name>NAMEI:</zone-name>
<zone-size>1024</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>8</free>
<zone-req>53330</zone-req>
<zone-name>NFSMOUNT:</zone-name>
<zone-size>480</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>NFSNODE:</zone-name>
<zone-size>460</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>PIPE:</zone-name>
<zone-size>404</zone-size>
<count-limit>0</count-limit>
<used>27</used>
<free>9</free>
<zone-req>717</zone-req>
<zone-name>KNOTE:</zone-name>
<zone-size>72</zone-size>
<count-limit>0</count-limit>
<used>42</used>
```

```
<free>64</free>
<zone-req>3311</zone-req>
<zone-name>socket:</zone-name>
<zone-size>412</zone-size>
<count-limit>25191</count-limit>
<used>343</used>
<free>8</free>
<zone-req>2524</zone-req>
<zone-name>unpcb:</zone-name>
<zone-size>140</zone-size>
<count-limit>25200</count-limit>
<used>170</used>
<free>26</free>
<zone-req>2157</zone-req>
<zone-name>ipq:</zone-name>
<zone-size>52</zone-size>
<count-limit>216</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>udpcb:</zone-name>
<zone-size>232</zone-size>
<count-limit>25194</count-limit>
<used>19</used>
<free>32</free>
<zone-req>31</zone-req>
<zone-name>inpcb:</zone-name>
<zone-size>232</zone-size>
<count-limit>25194</count-limit>
<used>40</used>
<free>28</free>
<zone-req>105</zone-req>
<zone-name>tcpcb:</zone-name>
<zone-size>520</zone-size>
<count-limit>25193</count-limit>
<used>40</used>
<free>16</free>
<zone-req>105</zone-req>
<zone-name>tcptw:</zone-name>
<zone-size>56</zone-size>
<count-limit>5092</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>syncache:</zone-name>
<zone-size>128</zone-size>
<count-limit>15360</count-limit>
<used>0</used>
<free>60</free>
<zone-req>55</zone-req>
<zone-name>tcpreass:</zone-name>
<zone-size>20</zone-size>
<count-limit>1690</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>sackhole:</zone-name>
<zone-size>20</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
```

```

    <zone-req>0</zone-req>
    <zone-name>ripcb:</zone-name>
    <zone-size>232</zone-size>
    <count-limit>25194</count-limit>
    <used>5</used>
    <free>29</free>
    <zone-req>5</zone-req>
    <zone-name>SWAPMETA:</zone-name>
    <zone-size>276</zone-size>
    <count-limit>94948</count-limit>
    <used>0</used>
    <free>0</free>
    <zone-req>0</zone-req>
    <zone-name>FFS inode:</zone-name>
    <zone-size>132</zone-size>
    <count-limit>0</count-limit>
    <used>1146</used>
    <free>72</free>
    <zone-req>1306</zone-req>
    <zone-name>FFS1 dinode:</zone-name>
    <zone-size>128</zone-size>
    <count-limit>0</count-limit>
    <used>1146</used>
    <free>24</free>
    <zone-req>1306</zone-req>
    <zone-name>FFS2 dinode:</zone-name>
    <zone-size>256</zone-size>
    <count-limit>0</count-limit>
    <used>0</used>
    <free>0</free>
    <zone-req>0</zone-req>
  </vmstat-memstat-zone>
<vmstat-sumstat>
  <cpu-context-switch>934906</cpu-context-switch>
  <dev-intr>1707986</dev-intr>
  <soft-intr>33819</soft-intr>
  <traps>203604</traps>
  <sys-calls>1200636</sys-calls>
  <kernel-thrds>60</kernel-thrds>
  <fork-calls>1313</fork-calls>
  <vfork-calls>21</vfork-calls>
  <rfork-calls>0</rfork-calls>
  <swap-pageins>0</swap-pageins>
  <swap-pagedin>0</swap-pagedin>
  <swap-pageouts>0</swap-pageouts>
  <swap-pagedout>0</swap-pagedout>
  <vnode-pageins>23094</vnode-pageins>
  <vnode-pagedin>23119</vnode-pagedin>
  <vnode-pageouts>226</vnode-pageouts>
  <vnode-pagedout>3143</vnode-pagedout>
  <page-daemon-wakeup>0</page-daemon-wakeup>
  <page-daemon-examined-pages>0</page-daemon-examined-pages>
  <pages-reactivated>8821</pages-reactivated>
  <copy-on-write-faults>48364</copy-on-write-faults>
  <copy-on-write-optimized-faults>31</copy-on-write-optimized-faults>
  <zero-fill-pages-zeroed>74665</zero-fill-pages-zeroed>
  <zero-fill-pages-prezeroed>70061</zero-fill-pages-prezeroed>
  <transit-blocking-page-faults>85</transit-blocking-page-faults>
  <total-vm-faults>191824</total-vm-faults>

<pages-affected-by-kernel-thrd-creat>0</pages-affected-by-kernel-thrd-creat>

```

```

<pages-affected-by-fork>95343</pages-affected-by-fork>
<pages-affected-by-vfork>3526</pages-affected-by-vfork>
<pages-affected-by-rfork>0</pages-affected-by-rfork>
<pages-freed>221502</pages-freed>
<pages-freed-by-daemon>0</pages-freed-by-daemon>
<pages-freed-by-exiting-proc>75630</pages-freed-by-exiting-proc>
<pages-active>45826</pages-active>
<pages-inactive>13227</pages-inactive>
<pages-in-vm-cache>49278</pages-in-vm-cache>
<pages-wired-down>10640</pages-wired-down>
<pages-free>70706</pages-free>
<bytes-per-page>4096</bytes-per-page>
<swap-pages-used>0</swap-pages-used>
<peak-swap-pages-used>0</peak-swap-pages-used>
<total-name-lookups>214496</total-name-lookups>
<positive-cache-hits>92</positive-cache-hits>
<negative-cache-hits>5</negative-cache-hits>
<pass2>0</pass2>
<cache-deletions>0</cache-deletions>
<cache-falsehits>0</cache-falsehits>
<toolong>0</toolong>
</vmstat-sumstat>
<vmstat-intr>
  <intr-name>irq0: clk          </intr-name>
  <intr-cnt>1243455</intr-cnt>
  <intr-rate>999</intr-rate>
  <intr-name>irq4: sio0        </intr-name>
  <intr-cnt>1140</intr-cnt>
  <intr-rate>0</intr-rate>
  <intr-name>irq8: rtc         </intr-name>
  <intr-cnt>159164</intr-cnt>
  <intr-rate>127</intr-rate>
  <intr-name>irq9: cbb1 fxp0   </intr-name>
  <intr-cnt>28490</intr-cnt>
  <intr-rate>22</intr-rate>
  <intr-name>irq10: fxp1       </intr-name>
  <intr-cnt>20593</intr-cnt>
  <intr-rate>16</intr-rate>
  <intr-name>irq14: ata0       </intr-name>
  <intr-cnt>5031</intr-cnt>
  <intr-rate>4</intr-rate>
  <intr-name>Total</intr-name>
  <intr-cnt>1457873</intr-cnt>
  <intr-rate>1171</intr-rate>
</vmstat-intr>
<vm-kernel-state>
  <vm-kmem-map-free>248524800</vm-kmem-map-free>
</vm-kernel-state>
<kernel-direct-mm-size-information>
  <vm-directmm-size-used>4644</vm-directmm-size-used>
  <vm-directmm-size-max>4057334</vm-directmm-size-max>
</kernel-direct-mm-size-information>
</system-virtual-memory-information>
<cli>
  <banner></banner>
</cli>
</rpc-reply>

```

Note: <kernel-direct-mm-size-information> only displays for 64 bit platform.

show system virtual-memory (QFX Series)

```

user@switch> show system virtual-memory | display xml
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/11.1R1/junos">
  <system-virtual-memory-information>
    <vmstat-memstat-malloc>
      <memstat-name>CAM dev queue</memstat-name>
      <inuse>1</inuse>
      <memuse>1</memuse>
      <high-use>-</high-use>
      <memstat-req>1</memstat-req>
      <memstat-size>64</memstat-size>
      <memstat-name>entropy</memstat-name>
      <inuse>1024</inuse>
      <memuse>64</memuse>
      <high-use>-</high-use>
      <memstat-req>1024</memstat-req>
      <memstat-size>64</memstat-size>
      <memstat-name>linker</memstat-name>
      <inuse>481</inuse>
      <memuse>1871</memuse>
      <high-use>-</high-use>
      <memstat-req>1145</memstat-req>
      <memstat-size>16, 32, 64, 4096, 32768, 131072</memstat-size>
      <memstat-name>lockf</memstat-name>
      <inuse>56</inuse>
      <memuse>4</memuse>
      <high-use>-</high-use>
      <memstat-req>5998</memstat-req>
      <memstat-size>64</memstat-size>
      <memstat-name>devbuf</memstat-name>
      <inuse>2094</inuse>
      <memuse>3877</memuse>
      <high-use>-</high-use>
      <memstat-req>2099</memstat-req>

      <memstat-size>16, 32, 64, 128, 512, 1024, 4096, 8192, 16384, 32768, 65536, 131072</memstat-size>

      <memstat-name>temp</memstat-name>
      <inuse>21</inuse>
      <memuse>66</memuse>
      <high-use>-</high-use>
      <memstat-req>3127</memstat-req>

      <memstat-size>16, 32, 64, 128, 256, 512, 2048, 4096, 8192, 16384, 32768, 65536, 131072</memstat-size>

      <memstat-name>ip6ndp</memstat-name>
      <inuse>0</inuse>
      <memuse>0</memuse>
      <high-use>-</high-use>
      <memstat-req>4</memstat-req>
      <memstat-size>64</memstat-size>
      <memstat-name>in6ifmulti</memstat-name>
      <inuse>1</inuse>
      <memuse>1</memuse>
      <high-use>-</high-use>
      <memstat-req>1</memstat-req>
      <memstat-size>64</memstat-size>
      <memstat-name>in6grentry</memstat-name>
      <inuse>1</inuse>
      <memuse>1</memuse>

```

```

<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>iflogical</memstat-name>
<inuse>13</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>13</memstat-req>
<memstat-size>64,2048</memstat-size>
<memstat-name>iffamily</memstat-name>
<inuse>28</inuse>
<memuse>4</memuse>
<high-use>--</high-use>
<memstat-req>28</memstat-req>
<memstat-size>32,1024,2048</memstat-size>
<memstat-name>rtnexthop</memstat-name>
<inuse>127</inuse>
<memuse>18</memuse>
<high-use>--</high-use>
<memstat-req>129</memstat-req>
<memstat-size>32,256,512,1024,2048,4096</memstat-size>
<memstat-name>metrics</memstat-name>
<inuse>3</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>5</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>inifmulti</memstat-name>
<inuse>3</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>3</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>ingrentry</memstat-name>
<inuse>6</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>6</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>rnode</memstat-name>
<inuse>68</inuse>
<memuse>2</memuse>
<high-use>--</high-use>
<memstat-req>76</memstat-req>
<memstat-size>16,32</memstat-size>
<memstat-name>rcache</memstat-name>
<inuse>4</inuse>
<memuse>8</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>ifdevice</memstat-name>
<inuse>4</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>ifstat</memstat-name>
<inuse>40</inuse>
<memuse>22</memuse>
<high-use>--</high-use>

```

```
<memstat-req>40</memstat-req>
<memstat-size>512,16384,32768</memstat-size>
<memstat-name>ipfw</memstat-name>
<inuse>42</inuse>
<memuse>23</memuse>
<high-use>--</high-use>
<memstat-req>91</memstat-req>

<memstat-size>16,32,64,128,256,512,1024,16384,32768,65536,131072</memstat-size>
<memstat-name>ifmaddr</memstat-name>
<inuse>103</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>103</memstat-req>
<memstat-size>16,32</memstat-size>
<memstat-name>rtable</memstat-name>
<inuse>129</inuse>
<memuse>14</memuse>
<high-use>--</high-use>
<memstat-req>139</memstat-req>
<memstat-size>16,32,64,128,1024,16384</memstat-size>
<memstat-name>sysctl</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>14847</memstat-req>
<memstat-size>16,32,64,4096,16384,32768</memstat-size>
<memstat-name>ifaddr</memstat-name>
<inuse>29</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>29</memstat-req>
<memstat-size>64,128</memstat-size>
<memstat-name>mkey</memstat-name>
<inuse>345</inuse>
<memuse>6</memuse>
<high-use>--</high-use>
<memstat-req>2527</memstat-req>
<memstat-size>16,128</memstat-size>
<memstat-name>pfe_ipc</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>1422</memstat-req>

<memstat-size>16,32,64,128,512,1024,2048,8192,16384,32768,65536,131072</memstat-size>

<memstat-name>ifstate</memstat-name>
<inuse>594</inuse>
<memuse>51</memuse>
<high-use>--</high-use>
<memstat-req>655</memstat-req>

<memstat-size>16,32,64,128,256,1024,2048,4096,16384,32768</memstat-size>
<memstat-name>itable16</memstat-name>
<inuse>276</inuse>
<memuse>52</memuse>
<high-use>--</high-use>
<memstat-req>294</memstat-req>
<memstat-size>1024,4096</memstat-size>
<memstat-name>itable32</memstat-name>
```



```

<inuse>160</inuse>
<memuse>10</memuse>
<high-use>--</high-use>
<memstat-req>160</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>itable64</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>lr</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>16384</memstat-size>
<memstat-name>pic</memstat-name>
<inuse>5</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>5</memstat-req>
<memstat-size>64,512</memstat-size>
<memstat-name>pfestat</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>162</memstat-req>
<memstat-size>16,32,128,256,16384</memstat-size>
<memstat-name>gencfg</memstat-name>
<inuse>224</inuse>
<memuse>56</memuse>
<high-use>--</high-use>
<memstat-req>540</memstat-req>
<memstat-size>16,32,64,256,512,32768,65536</memstat-size>
<memstat-name>jsr</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>idl</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>13</memstat-req>
<memstat-size>16,32,64,128,256,4096,16384,32768,131072</memstat-size>

<memstat-name>rtsmsg</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>131072</memstat-size>
<memstat-name>module</memstat-name>
<inuse>249</inuse>
<memuse>16</memuse>
<high-use>--</high-use>
<memstat-req>249</memstat-req>
<memstat-size>64,128</memstat-size>
<memstat-name>mtx_pool</memstat-name>

```

```
<inuse>1</inuse>
<memuse>8</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64,128</memstat-size>
<memstat-name>DEVFS3</memstat-name>
<inuse>109</inuse>
<memuse>12</memuse>
<high-use>--</high-use>
<memstat-req>117</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>DEVFS1</memstat-name>
<inuse>102</inuse>
<memuse>23</memuse>
<high-use>--</high-use>
<memstat-req>109</memstat-req>
<memstat-size>2048</memstat-size>
<memstat-name>pgrp</memstat-name>
<inuse>12</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>21</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>session</memstat-name>
<inuse>8</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>15</memstat-req>
<memstat-size>512</memstat-size>
<memstat-name>proc</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>16384</memstat-size>
<memstat-name>subproc</memstat-name>
<inuse>244</inuse>
<memuse>496</memuse>
<high-use>--</high-use>
<memstat-req>1522</memstat-req>
<memstat-size>2048,131072</memstat-size>
<memstat-name>cred</memstat-name>
<inuse>30</inuse>
<memuse>4</memuse>
<high-use>--</high-use>
<memstat-req>11409</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>plimit</memstat-name>
<inuse>17</inuse>
<memuse>4</memuse>
<high-use>--</high-use>
<memstat-req>133</memstat-req>
<memstat-size>2048</memstat-size>
<memstat-name>uidinfo</memstat-name>
<inuse>3</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>6</memstat-req>
<memstat-size>32,512</memstat-size>
<memstat-name>sysctloid</memstat-name>
<inuse>1117</inuse>
```

```

<memuse>34</memuse>
<high-use>--</high-use>
<memstat-req>1117</memstat-req>
<memstat-size>16,32,64</memstat-size>
<memstat-name>sysctltmp</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>743</memstat-req>
<memstat-size>16,32,64,1024</memstat-size>
<memstat-name>umtx</memstat-name>
<inuse>144</inuse>
<memuse>9</memuse>
<high-use>--</high-use>
<memstat-req>144</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>SWAP</memstat-name>
<inuse>2</inuse>
<memuse>209</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>bus</memstat-name>
<inuse>496</inuse>
<memuse>55</memuse>
<high-use>--</high-use>
<memstat-req>1196</memstat-req>
<memstat-size>16,32,64,128,32768</memstat-size>
<memstat-name>bus-sc</memstat-name>
<inuse>23</inuse>
<memuse>33</memuse>
<high-use>--</high-use>
<memstat-req>335</memstat-req>

<memstat-size>16,32,64,512,1024,2048,8192,16384,65536,131072</memstat-size>
<memstat-name>devstat</memstat-name>
<inuse>10</inuse>
<memuse>21</memuse>
<high-use>--</high-use>
<memstat-req>10</memstat-req>
<memstat-size>16,131072</memstat-size>
<memstat-name>eventhandler</memstat-name>
<inuse>35</inuse>
<memuse>2</memuse>
<high-use>--</high-use>
<memstat-req>36</memstat-req>
<memstat-size>32,128</memstat-size>
<memstat-name>kobj</memstat-name>
<inuse>93</inuse>
<memuse>186</memuse>
<high-use>--</high-use>
<memstat-req>111</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>DEVFS</memstat-name>
<inuse>8</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>9</memstat-req>
<memstat-size>16,64</memstat-size>
<memstat-name>rman</memstat-name>
<inuse>71</inuse>

```

```
<memuse>5</memuse>
<high-use>--</high-use>
<memstat-req>433</memstat-req>
<memstat-size>16,32,64</memstat-size>
<memstat-name>sbuf</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>522</memstat-req>
<memstat-size>16,32,32768,131072</memstat-size>
<memstat-name>NULLFS hash</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>taskqueue</memstat-name>
<inuse>5</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>5</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>turnstiles</memstat-name>
<inuse>145</inuse>
<memuse>10</memuse>
<high-use>--</high-use>
<memstat-req>145</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>Unitno</memstat-name>
<inuse>8</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>44</memstat-req>
<memstat-size>16,64</memstat-size>
<memstat-name>iocltops</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>27622</memstat-req>
<memstat-size>16,64,8192,16384,131072</memstat-size>
<memstat-name>iov</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>18578</memstat-req>
<memstat-size>16,64,128,256,512,1024,2048,131072</memstat-size>
<memstat-name>msg</memstat-name>
<inuse>4</inuse>
<memuse>25</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>32768,131072</memstat-size>
<memstat-name>sem</memstat-name>
<inuse>4</inuse>
<memuse>7</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>16384,32768,131072</memstat-size>
<memstat-name>shm</memstat-name>
<inuse>9</inuse>
<memuse>20</memuse>
```

```

<high-use>--</high-use>
<memstat-req>14</memstat-req>
<memstat-size>32768</memstat-size>
<memstat-name>ttys</memstat-name>
<inuse>321</inuse>
<memuse>61</memuse>
<high-use>--</high-use>
<memstat-req>528</memstat-req>
<memstat-size>512,32768</memstat-size>
<memstat-name>ptys</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>mbuf_tag</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>23383</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>soname</memstat-name>
<inuse>115</inuse>
<memuse>12</memuse>
<high-use>--</high-use>
<memstat-req>24712</memstat-req>
<memstat-size>16,32,64,256</memstat-size>
<memstat-name>pcb</memstat-name>
<inuse>216</inuse>
<memuse>33</memuse>
<high-use>--</high-use>
<memstat-req>484</memstat-req>

<memstat-size>16,32,64,128,1024,2048,4096,16384,32768,65536</memstat-size>
<memstat-name>BIO buffer</memstat-name>
<inuse>43</inuse>
<memuse>86</memuse>
<high-use>--</high-use>
<memstat-req>405</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>vfscache</memstat-name>
<inuse>1</inuse>
<memuse>256</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>cluster_save buffer</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>32,64</memstat-size>
<memstat-name>VFS hash</memstat-name>
<inuse>1</inuse>
<memuse>128</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>32,64</memstat-size>
<memstat-name>vnodes</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>

```

```
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>512</memstat-size>
<memstat-name>mount</memstat-name>
<inuse>290</inuse>
<memuse>23</memuse>
<high-use>--</high-use>
<memstat-req>535</memstat-req>
<memstat-size>16,32,64,128,256,4096,32768</memstat-size>
<memstat-name>vnodemarker</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>498</memstat-req>
<memstat-size>16384</memstat-size>
<memstat-name>pfs_nodes</memstat-name>
<inuse>25</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>25</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>pfs_vncache</memstat-name>
<inuse>27</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>53</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>STP</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>GEOM</memstat-name>
<inuse>146</inuse>
<memuse>11</memuse>
<high-use>--</high-use>
<memstat-req>1042</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,16384,32768,131072</memstat-size>
<memstat-name>syncache</memstat-name>
<inuse>1</inuse>
<memuse>8</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,16384,32768,131072</memstat-size>
<memstat-name>tlv_stat</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>8</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,16384,32768,131072</memstat-size>
<memstat-name>NFS_daemon</memstat-name>
<inuse>1</inuse>
<memuse>8</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,16384,32768,131072</memstat-size>
```

```

<memstat-name>p1003.1b</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>MD disk</memstat-name>
<inuse>10</inuse>
<memuse>20</memuse>
<high-use>--</high-use>
<memstat-req>10</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>ata_generic</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>6</memstat-req>
<memstat-size>16,16384,32768</memstat-size>
<memstat-name>ISofs mount</memstat-name>
<inuse>8</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>15</memstat-req>
<memstat-size>512</memstat-size>
<memstat-name>ISofs node</memstat-name>
<inuse>1440</inuse>
<memuse>135</memuse>
<high-use>--</high-use>
<memstat-req>1457</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>CAM SIM</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>CAM XPT</memstat-name>
<inuse>6</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>9</memstat-req>
<memstat-size>16,64,16384</memstat-size>
<memstat-name>CAM periph</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>ad_driver</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>pagedep</memstat-name>
<inuse>1</inuse>
<memuse>32</memuse>
<high-use>--</high-use>
<memstat-req>106</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>inodedep</memstat-name>

```

```
<inuse>1</inuse>
<memuse>128</memuse>
<high-use>--</high-use>
<memstat-req>464</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>newblk</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>336</memstat-req>
<memstat-size>64,4096</memstat-size>
<memstat-name>bmsafemap</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>63</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>allocdirect</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>320</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>indirdep</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>17</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>allocindir</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>15</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>freefrag</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>12</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>freeblks</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>40</memstat-req>
<memstat-size>2048</memstat-size>
<memstat-name>freefile</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>101</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>diradd</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>465</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>mkdir</memstat-name>
<inuse>0</inuse>
```



```

<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>136</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>dirrem</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>168</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>newdirblk</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>savedino</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>157</memstat-req>
<memstat-size>512</memstat-size>
<memstat-name>UFS mount</memstat-name>
<inuse>15</inuse>
<memuse>36</memuse>
<high-use>--</high-use>
<memstat-req>15</memstat-req>
<memstat-size>2048,65536,131072</memstat-size>
<memstat-name>ata_dma</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>UMAHash</memstat-name>
<inuse>1</inuse>
<memuse>2</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>4096,16384,32768,65536</memstat-size>
<memstat-name>cdev</memstat-name>
<inuse>22</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>22</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>file desc</memstat-name>
<inuse>141</inuse>
<memuse>32</memuse>
<high-use>--</high-use>
<memstat-req>1583</memstat-req>
<memstat-size>16,1024,2048,16384</memstat-size>
<memstat-name>VM pgdata</memstat-name>
<inuse>2</inuse>
<memuse>65</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>sigio</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>

```

```
<high-use>--</high-use>
<memstat-req>20</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>kenv</memstat-name>
<inuse>24</inuse>
<memuse>5</memuse>
<high-use>--</high-use>
<memstat-req>27</memstat-req>
<memstat-size>16,32,64,131072</memstat-size>
<memstat-name>atkbddev</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>kqueue</memstat-name>
<inuse>15</inuse>
<memuse>9</memuse>
<high-use>--</high-use>
<memstat-req>19</memstat-req>
<memstat-size>1024,4096,32768</memstat-size>
<memstat-name>proc-args</memstat-name>
<inuse>57</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>1001</memstat-req>
<memstat-size>16,32,64,128,256,512,1024</memstat-size>
<memstat-name>isadev</memstat-name>
<inuse>21</inuse>
<memuse>2</memuse>
<high-use>--</high-use>
<memstat-req>21</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>zombie</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>1278</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>ithread</memstat-name>
<inuse>69</inuse>
<memuse>5</memuse>
<high-use>--</high-use>
<memstat-req>69</memstat-req>
<memstat-size>16,64,256</memstat-size>
<memstat-name>legacydrv</memstat-name>
<inuse>4</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>memdesc</memstat-name>
<inuse>1</inuse>
<memuse>4</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>131072</memstat-size>
<memstat-name>nexusdev</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
```

```

<memstat-req>2</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>CAM queue</memstat-name>
<inuse>3</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>3</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>$PIR</memstat-name>
<inuse>4</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>KTRACE</memstat-name>
<inuse>100</inuse>
<memuse>10</memuse>
<high-use>--</high-use>
<memstat-req>100</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>kbdmux</memstat-name>
<inuse>5</inuse>
<memuse>9</memuse>
<high-use>--</high-use>
<memstat-req>5</memstat-req>
<memstat-size>128,2048,65536,131072</memstat-size>
</vmstat-memstat-malloc>
<vmstat-memstat-zone>
  <zone-name>UMA Kegs:</zone-name>
  <zone-size>136</zone-size>
  <count-limit>0</count-limit>
  <used>71</used>
  <free>1</free>
  <zone-req>71</zone-req>
  <zone-name>UMA Zones:</zone-name>
  <zone-size>120</zone-size>
  <count-limit>0</count-limit>
  <used>71</used>
  <free>19</free>
  <zone-req>71</zone-req>
  <zone-name>UMA Slabs:</zone-name>
  <zone-size>64</zone-size>
  <count-limit>0</count-limit>
  <used>490</used>
  <free>41</free>
  <zone-req>579</zone-req>
  <zone-name>UMA RCntSlabs:</zone-name>
  <zone-size>104</zone-size>
  <count-limit>0</count-limit>
  <used>276</used>
  <free>20</free>
  <zone-req>276</zone-req>
  <zone-name>UMA Hash:</zone-name>
  <zone-size>128</zone-size>
  <count-limit>0</count-limit>
  <used>4</used>
  <free>26</free>
  <zone-req>5</zone-req>
  <zone-name>16 Bucket:</zone-name>
  <zone-size>76</zone-size>
  <count-limit>0</count-limit>

```

```
<used>30</used>
<free>20</free>
<zone-req>30</zone-req>
<zone-name>32 Bucket:</zone-name>
<zone-size>140</zone-size>
<count-limit>0</count-limit>
<used>33</used>
<free>23</free>
<zone-req>33</zone-req>
<zone-name>64 Bucket:</zone-name>
<zone-size>268</zone-size>
<count-limit>0</count-limit>
<used>33</used>
<free>9</free>
<zone-req>33</zone-req>
<zone-name>128 Bucket:</zone-name>
<zone-size>524</zone-size>
<count-limit>0</count-limit>
<used>49</used>
<free>0</free>
<zone-req>49</zone-req>
<zone-name>VM OBJECT:</zone-name>
<zone-size>128</zone-size>
<count-limit>0</count-limit>
<used>2111</used>
<free>79</free>
<zone-req>25214</zone-req>
<zone-name>MAP:</zone-name>
<zone-size>160</zone-size>
<count-limit>0</count-limit>
<used>7</used>
<free>41</free>
<zone-req>7</zone-req>
<zone-name>KMAP ENTRY:</zone-name>
<zone-size>68</zone-size>
<count-limit>35336</count-limit>
<used>19</used>
<free>149</free>
<zone-req>2397</zone-req>
<zone-name>MAP ENTRY:</zone-name>
<zone-size>68</zone-size>
<count-limit>0</count-limit>
<used>2031</used>
<free>153</free>
<zone-req>62417</zone-req>
<zone-name>PV ENTRY:</zone-name>
<zone-size>24</zone-size>
<count-limit>509095</count-limit>
<used>57177</used>
<free>6333</free>
<zone-req>1033683</zone-req>
<zone-name>DP fakepg:</zone-name>
<zone-size>72</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>mt_zone:</zone-name>
<zone-size>64</zone-size>
<count-limit>0</count-limit>
<used>238</used>
```

```
<free>57</free>
<zone-req>238</zone-req>
<zone-name>16:</zone-name>
<zone-size>16</zone-size>
<count-limit>0</count-limit>
<used>2114</used>
<free>119</free>
<zone-req>80515</zone-req>
<zone-name>32:</zone-name>
<zone-size>32</zone-size>
<count-limit>0</count-limit>
<used>1335</used>
<free>134</free>
<zone-req>10259</zone-req>
<zone-name>64:</zone-name>
<zone-size>64</zone-size>
<count-limit>0</count-limit>
<used>3529</used>
<free>129</free>
<zone-req>29110</zone-req>
<zone-name>96:</zone-name>
<zone-size>96</zone-size>
<count-limit>0</count-limit>
<used>2062</used>
<free>58</free>
<zone-req>4365</zone-req>
<zone-name>112:</zone-name>
<zone-size>112</zone-size>
<count-limit>0</count-limit>
<used>361</used>
<free>164</free>
<zone-req>24613</zone-req>
<zone-name>128:</zone-name>
<zone-size>128</zone-size>
<count-limit>0</count-limit>
<used>359</used>
<free>61</free>
<zone-req>942</zone-req>
<zone-name>160:</zone-name>
<zone-size>160</zone-size>
<count-limit>0</count-limit>
<used>364</used>
<free>44</free>
<zone-req>577</zone-req>
<zone-name>224:</zone-name>
<zone-size>224</zone-size>
<count-limit>0</count-limit>
<used>422</used>
<free>20</free>
<zone-req>1950</zone-req>
<zone-name>256:</zone-name>
<zone-size>256</zone-size>
<count-limit>0</count-limit>
<used>204</used>
<free>36</free>
<zone-req>1225</zone-req>
<zone-name>288:</zone-name>
<zone-size>288</zone-size>
<count-limit>0</count-limit>
<used>2</used>
<free>24</free>
```

```
<zone-req>10</zone-req>
<zone-name>512:</zone-name>
<zone-size>512</zone-size>
<count-limit>0</count-limit>
<used>49</used>
<free>7</free>
<zone-req>911</zone-req>
<zone-name>1024:</zone-name>
<zone-size>1024</zone-size>
<count-limit>0</count-limit>
<used>213</used>
<free>11</free>
<zone-req>1076</zone-req>
<zone-name>2048:</zone-name>
<zone-size>2048</zone-size>
<count-limit>0</count-limit>
<used>199</used>
<free>113</free>
<zone-req>640</zone-req>
<zone-name>4096:</zone-name>
<zone-size>4096</zone-size>
<count-limit>0</count-limit>
<used>144</used>
<free>7</free>
<zone-req>2249</zone-req>
<zone-name>Files:</zone-name>
<zone-size>72</zone-size>
<count-limit>0</count-limit>
<used>665</used>
<free>77</free>
<zone-req>16457</zone-req>
<zone-name>MAC labels:</zone-name>
<zone-size>20</zone-size>
<count-limit>0</count-limit>
<used>3998</used>
<free>227</free>
<zone-req>21947</zone-req>
<zone-name>PROC:</zone-name>
<zone-size>544</zone-size>
<count-limit>0</count-limit>
<used>116</used>
<free>10</free>
<zone-req>1394</zone-req>
<zone-name>THREAD:</zone-name>
<zone-size>416</zone-size>
<count-limit>0</count-limit>
<used>127</used>
<free>17</free>
<zone-req>131</zone-req>
<zone-name>KSEGRP:</zone-name>
<zone-size>88</zone-size>
<count-limit>0</count-limit>
<used>127</used>
<free>73</free>
<zone-req>131</zone-req>
<zone-name>UPCALL:</zone-name>
<zone-size>44</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
```

```

<zone-name>SLEEPQUEUE:</zone-name>
<zone-size>32</zone-size>
<count-limit>0</count-limit>
<used>145</used>
<free>194</free>
<zone-req>145</zone-req>
<zone-name>VMSPACE:</zone-name>
<zone-size>268</zone-size>
<count-limit>0</count-limit>
<used>57</used>
<free>13</free>
<zone-req>1335</zone-req>
<zone-name>mbuf_packet:</zone-name>
<zone-size>256</zone-size>
<count-limit>180000</count-limit>
<used>256</used>
<free>128</free>
<zone-req>49791</zone-req>
<zone-name>mbuf:</zone-name>
<zone-size>256</zone-size>
<count-limit>180000</count-limit>
<used>50</used>
<free>466</free>
<zone-req>105183</zone-req>
<zone-name>mbuf_cluster:</zone-name>
<zone-size>2048</zone-size>
<count-limit>25190</count-limit>
<used>387</used>
<free>165</free>
<zone-req>5976</zone-req>
<zone-name>mbuf_jumbo_pagesize:</zone-name>
<zone-size>4096</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>mbuf_jumbo_9k:</zone-name>
<zone-size>9216</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>mbuf_jumbo_16k:</zone-name>
<zone-size>16384</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>ACL_UMA_zone:</zone-name>
<zone-size>388</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>g_bio:</zone-name>
<zone-size>132</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>174</free>
<zone-req>69750</zone-req>
<zone-name>ata_request:</zone-name>

```

```
<zone-size>200</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>57</free>
<zone-req>5030</zone-req>
<zone-name>ata_composite:</zone-name>
<zone-size>192</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>GENCFG:</zone-name>
<zone-size>72</zone-size>
<count-limit>1000004</count-limit>
<used>57</used>
<free>102</free>
<zone-req>57</zone-req>
<zone-name>VNODE:</zone-name>
<zone-size>292</zone-size>
<count-limit>0</count-limit>
<used>2718</used>
<free>25</free>
<zone-req>2922</zone-req>
<zone-name>VNODEPOLL:</zone-name>
<zone-size>72</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>S VFS Cache:</zone-name>
<zone-size>68</zone-size>
<count-limit>0</count-limit>
<used>2500</used>
<free>76</free>
<zone-req>3824</zone-req>
<zone-name>L VFS Cache:</zone-name>
<zone-size>291</zone-size>
<count-limit>0</count-limit>
<used>51</used>
<free>14</free>
<zone-req>63</zone-req>
<zone-name>NAMEI:</zone-name>
<zone-size>1024</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>8</free>
<zone-req>53330</zone-req>
<zone-name>NFSMOUNT:</zone-name>
<zone-size>480</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>NFSNODE:</zone-name>
<zone-size>460</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>PIPE:</zone-name>
<zone-size>404</zone-size>
```



```

<count-limit>0</count-limit>
<used>27</used>
<free>9</free>
<zone-req>717</zone-req>
<zone-name>KNOTE:</zone-name>
<zone-size>72</zone-size>
<count-limit>0</count-limit>
<used>42</used>
<free>64</free>
<zone-req>3311</zone-req>
<zone-name>socket:</zone-name>
<zone-size>412</zone-size>
<count-limit>25191</count-limit>
<used>343</used>
<free>8</free>
<zone-req>2524</zone-req>
<zone-name>unpcb:</zone-name>
<zone-size>140</zone-size>
<count-limit>25200</count-limit>
<used>170</used>
<free>26</free>
<zone-req>2157</zone-req>
<zone-name>ipq:</zone-name>
<zone-size>52</zone-size>
<count-limit>216</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>udpcb:</zone-name>
<zone-size>232</zone-size>
<count-limit>25194</count-limit>
<used>19</used>
<free>32</free>
<zone-req>31</zone-req>
<zone-name>inpcb:</zone-name>
<zone-size>232</zone-size>
<count-limit>25194</count-limit>
<used>40</used>
<free>28</free>
<zone-req>105</zone-req>
<zone-name>tcpcb:</zone-name>
<zone-size>520</zone-size>
<count-limit>25193</count-limit>
<used>40</used>
<free>16</free>
<zone-req>105</zone-req>
<zone-name>tcptw:</zone-name>
<zone-size>56</zone-size>
<count-limit>5092</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>syncache:</zone-name>
<zone-size>128</zone-size>
<count-limit>15360</count-limit>
<used>0</used>
<free>60</free>
<zone-req>55</zone-req>
<zone-name>tcpreass:</zone-name>
<zone-size>20</zone-size>
<count-limit>1690</count-limit>

```

```

<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>sackhole:</zone-name>
<zone-size>20</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>ripcb:</zone-name>
<zone-size>232</zone-size>
<count-limit>25194</count-limit>
<used>5</used>
<free>29</free>
<zone-req>5</zone-req>
<zone-name>SWAPMETA:</zone-name>
<zone-size>276</zone-size>
<count-limit>94948</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>FFS inode:</zone-name>
<zone-size>132</zone-size>
<count-limit>0</count-limit>
<used>1146</used>
<free>72</free>
<zone-req>1306</zone-req>
<zone-name>FFS1 dinode:</zone-name>
<zone-size>128</zone-size>
<count-limit>0</count-limit>
<used>1146</used>
<free>24</free>
<zone-req>1306</zone-req>
<zone-name>FFS2 dinode:</zone-name>
<zone-size>256</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
</vmstat-memstat-zone>
<vmstat-sumstat>
  <cpu-context-switch>934906</cpu-context-switch>
  <dev-intr>1707986</dev-intr>
  <soft-intr>33819</soft-intr>
  <traps>203604</traps>
  <sys-calls>1200636</sys-calls>
  <kernel-thrds>60</kernel-thrds>
  <fork-calls>1313</fork-calls>
  <vfork-calls>21</vfork-calls>
  <rfork-calls>0</rfork-calls>
  <swap-pageins>0</swap-pageins>
  <swap-pagedin>0</swap-pagedin>
  <swap-pageouts>0</swap-pageouts>
  <swap-pagedout>0</swap-pagedout>
  <vnode-pageins>23094</vnode-pageins>
  <vnode-pagedin>23119</vnode-pagedin>
  <vnode-pageouts>226</vnode-pageouts>
  <vnode-pagedout>3143</vnode-pagedout>
  <page-daemon-wakeup>0</page-daemon-wakeup>
  <page-daemon-examined-pages>0</page-daemon-examined-pages>
  <pages-reactivated>8821</pages-reactivated>

```

```

<copy-on-write-faults>48364</copy-on-write-faults>
<copy-on-write-optimized-faults>31</copy-on-write-optimized-faults>
<zero-fill-pages-zeroed>74665</zero-fill-pages-zeroed>
<zero-fill-pages-prezeroed>70061</zero-fill-pages-prezeroed>
<transit-blocking-page-faults>85</transit-blocking-page-faults>
<total-vm-faults>191824</total-vm-faults>

<pages-affected-by-kernel-thrd-creat>0</pages-affected-by-kernel-thrd-creat>
<pages-affected-by-fork>95343</pages-affected-by-fork>
<pages-affected-by-vfork>3526</pages-affected-by-vfork>
<pages-affected-by-rfork>0</pages-affected-by-rfork>
<pages-freed>221502</pages-freed>
<pages-freed-by-daemon>0</pages-freed-by-daemon>
<pages-freed-by-exiting-proc>75630</pages-freed-by-exiting-proc>
<pages-active>45826</pages-active>
<pages-inactive>13227</pages-inactive>
<pages-in-vm-cache>49278</pages-in-vm-cache>
<pages-wired-down>10640</pages-wired-down>
<pages-free>70706</pages-free>
<bytes-per-page>4096</bytes-per-page>
<swap-pages-used>0</swap-pages-used>
<peak-swap-pages-used>0</peak-swap-pages-used>
<total-name-lookups>214496</total-name-lookups>
<positive-cache-hits>92</positive-cache-hits>
<negative-cache-hits>5</negative-cache-hits>
<pass2>0</pass2>
<cache-deletions>0</cache-deletions>
<cache-falsehits>0</cache-falsehits>
<toolong>0</toolong>
</vmstat-sumstat>
<vmstat-intr>
  <intr-name>irq0: clk          </intr-name>
  <intr-cnt>1243455</intr-cnt>
  <intr-rate>999</intr-rate>
  <intr-name>irq4: sio0        </intr-name>
  <intr-cnt>1140</intr-cnt>
  <intr-rate>0</intr-rate>
  <intr-name>irq8: rtc         </intr-name>
  <intr-cnt>159164</intr-cnt>
  <intr-rate>127</intr-rate>
  <intr-name>irq9: cbb1 fxp0    </intr-name>
  <intr-cnt>28490</intr-cnt>
  <intr-rate>22</intr-rate>
  <intr-name>irq10: fxp1       </intr-name>
  <intr-cnt>20593</intr-cnt>
  <intr-rate>16</intr-rate>
  <intr-name>irq14: ata0       </intr-name>
  <intr-cnt>5031</intr-cnt>
  <intr-rate>4</intr-rate>
  <intr-name>Total</intr-name>
  <intr-cnt>1457873</intr-cnt>
  <intr-rate>1171</intr-rate>
</vmstat-intr>
<vm-kernel-state>
  <vm-kmem-map-free>248524800</vm-kmem-map-free>
</vm-kernel-state>
</system-virtual-memory-information>
<cli>
  <banner></banner>
</cli>
</rpc-reply>

```


show version

List of Syntax	Syntax on page 1253 Syntax (EX Series Switches) on page 1253 Syntax (TX Matrix Router) on page 1253 Syntax (TX Matrix Plus Router) on page 1253 Syntax (MX Series Router) on page 1253 Syntax (QFX Series) on page 1253
Syntax	<pre>show version <brief detail></pre>
Syntax (EX Series Switches)	<pre>show version <all-members> <brief detail> <local> <member member-id></pre>
Syntax (TX Matrix Router)	<pre>show version <brief detail> <all-chassis all-lcc lcc number scc></pre>
Syntax (TX Matrix Plus Router)	<pre>show version <all-chassis all-lcc lcc number sfc number> <brief detail></pre>
Syntax (MX Series Router)	<pre>show version <brief detail> <all-members> <local> <member member-id></pre>
Syntax (QFX Series)	<pre>show version <brief detail> <component component-name all></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Display the hostname and version information about the software running on the router or switch.</p> <p>Beginning in Junos OS Release 13.3, the show version command output includes the Junos field that displays the Junos OS version running on the device. This field provides a consistent means of identifying the Junos OS version, rather than extracting that information from the list of installed sub-packages.</p>

- Options** **none**—Display standard information about the hostname and version of the software running on the router or switch.
- brief | detail**—(Optional) Display the specified level of output.
- all-members**—(EX4200 switches and MX Series routers only) (Optional) Display standard information about the hostname and version of the software running on all members of the Virtual Chassis configuration.
- component all**—(QFabric systems only) (Optional) Display the host name and version information about the software running on all the components on the QFabric system.
- component *component-name***—(QFabric systems only) (Optional) Display the host name and version information about the software running on a specific QFabric system component. Replace *component-name* with the name of the QFabric system component. The *component-name* can be the name of a diagnostics Routing Engine, Director group, fabric control Routing Engine, fabric manager Routing Engine, Interconnect device, or Node group.
- local**—(EX4200 switches and MX Series routers only) (Optional) Display standard information about the hostname and version of the software running on the local Virtual Chassis member.
- member *member-id***—(EX4200 switches and MX Series routers only) (Optional) Display standard information about the hostname and version of the software running on the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.
- scc**—(TX Matrix routers only) (Optional) Display the hostname and version information about the software running on the TX Matrix router (or switch-card chassis).
- lcc *number***—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display the host name and version information about the software running on for a specified T640 router (line-card chassis or LCC) that is connected to the TX Matrix router. On a TX Matrix Plus router, display the host name and version information about the software running for a specified T1600 or T4000 router (LCC) that is connected to the TX Matrix Plus router.
- Replace *number* with the following values depending on the LCC configuration:
- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
 - 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
 - 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
 - 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

sfc *number*—(TX Matrix Plus routers only) (Optional) Display the hostname and version information about the software running on the TX Matrix Plus router (or switch-fabric chassis). Replace *number* with 0.

Additional Information By default, when you issue the **show version** command on a TX Matrix or TX Matrix Plus master Routing Engine, the command is broadcast to all the T640 (in a routing matrix based on a TX Matrix router) or T1600 or T4000 (in a routing matrix based on a TX Matrix Plus router) master Routing Engines connected to it. Likewise, if you issue the same command on the TX Matrix or TX Matrix Plus backup Routing Engine, the command is broadcast to all the T640 (in a routing matrix based on a TX Matrix router) or T1600 or T4000 (in a routing matrix based on a TX Matrix Plus router) backup Routing Engines that are connected to it.

Required Privilege Level view

List of Sample Output [show version \(Devices Running Junos OS Release 13.3 and Later\) on page 1256](#)
[show version on page 1256](#)
[show version \(TX Matrix Plus Router\) on page 1257](#)
[show version \(TX Matrix Plus Router with 3D SIBs\) on page 1259](#)
[show version \(MX Series Router\) on page 1263](#)
[show version \(QFX3500 Switch\) on page 1263](#)
[show version \(QFabric System\) on page 1263](#)
[show version component all \(QFabric System\) on page 1264](#)

Sample Output

show version (Devices Running Junos OS Release 13.3 and Later)

The following output is from the MX240 Router and shows the **Junos** field introduced in Junos OS 13.3. Depending on the platform running Junos OS 13.3, you might see different installed sub-packages, but the **Junos** field is common across all platforms that run Junos OS 13.3 and later.

```
user@host > show version
Hostname: lab
Model: mx240
Junos: 13.3R1.4
JUNOS Base OS boot [13.3R1.4]
JUNOS Base OS Software Suite [13.3R1.4]
JUNOS Kernel Software Suite [13.3R1.4]
JUNOS Crypto Software Suite [13.3R1.4]
JUNOS Packet Forwarding Engine Support (M/T/EX Common) [13.3R1.4]
JUNOS Packet Forwarding Engine Support (MX Common) [13.3R1.4]
JUNOS Online Documentation [13.3R1.4]
JUNOS Services ACL Container package [13.3R1.4]
JUNOS Services Application Level Gateways [13.3R1.4]
JUNOS AppId Services [13.3R1.4]
JUNOS Border Gateway Function package [13.3R1.4]
JUNOS Services Captive Portal and Content Delivery Container package [13.3R1.4]
JUNOS Services HTTP Content Management package [13.3R1.4]
JUNOS IDP Services [13.3R1.4]
JUNOS Services Jflow Container package [13.3R1.4]
JUNOS Services LL-PDF Container package [13.3R1.4]
JUNOS Services MobileNext Software package [13.3R1.4]
JUNOS Services Mobile Subscriber Service Container package [13.3R1.4]
JUNOS Services NAT [13.3R1.4]
JUNOS Services PTSP Container package [13.3R1.4]
JUNOS Services RPM [13.3R1.4]
JUNOS Services Stateful Firewall [13.3R1.4]
JUNOS Voice Services Container package [13.3R1.4]
JUNOS Services Crypto [13.3R1.4]
JUNOS Services SSL [13.3R1.4]
JUNOS Services IPSec [13.3R1.4]
JUNOS platform Software Suite [13.3R1.4]
JUNOS Runtime Software Suite [13.3R1.4]
JUNOS Routing Software Suite [13.3R1.4]
JUNOS py-base-i386 [13.3R1.4]
```

show version

```
user@host> show version
Hostname: router1
Model: m20
JUNOS Base OS boot [7.2-20050312.0]
JUNOS Base OS Software Suite [7.2-20050312.0]
JUNOS Kernel Software Suite [7.2R1.7]
JUNOS Packet Forwarding Engine Support (M20/M40) [7.2R1.7]
JUNOS Routing Software Suite [7.2R1.7]
JUNOS Online Documentation [7.2R1.7]
JUNOS Crypto Software Suite [7.2R1.7]

{master}

user@host> show version psd 1
```



```
psd1-re0:
```

```
-----
Hostname: china
Model: t640
JUNOS Base OS boot [9.1I20080311_1959_builder]
JUNOS Base OS Software Suite [9.1-20080321.0]
JUNOS Kernel Software Suite [9.1-20080321.0]
JUNOS Crypto Software Suite [9.1-20080321.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [9.1-20080321.0]
JUNOS Packet Forwarding Engine Support (T-series) [9.1-20080321.0]
JUNOS Online Documentation [9.1-20080321.0]
JUNOS Routing Software Suite [9.1-20080321.0]
labpkg [7.0]
```

show version (TX Matrix Plus Router)

```
user@host> show version
```

```
sfc0-re0:
```

```
-----
Hostname: host
Model: txp
JUNOS Base OS boot [12.3-20121019.0]
JUNOS Base OS Software Suite [12.3-20121019.0]
JUNOS Kernel Software Suite [12.3-20121019.0]
JUNOS Crypto Software Suite [12.3-20121019.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [12.3-20121019.0]
JUNOS Packet Forwarding Engine Support (T-Series) [12.3-20121019.0]
JUNOS Online Documentation [12.3-20121019.0]
JUNOS Services AACL Container package [12.3-20121019.0]
JUNOS Services Application Level Gateways [12.3-20121019.0]
JUNOS AppId Services [12.3-20121019.0]
JUNOS Border Gateway Function package [12.3-20121019.0]
JUNOS Services Captive Portal and Content Delivery Container package
[12.3-20121019.0]
JUNOS Services HTTP Content Management package [12.3-20121019.0]
JUNOS IDP Services [12.3-20121019.0]
JUNOS Services LL-PDF Container package [12.3-20121019.0]
JUNOS Services NAT [12.3-20121019.0]
JUNOS Services PTSP Container package [12.3-20121019.0]
JUNOS Services RPM [12.3-20121019.0]
JUNOS Services Stateful Firewall [12.3-20121019.0]
JUNOS Voice Services Container package [12.3-20121019.0]
JUNOS Services Example Container package [12.3-20121019.0]
JUNOS Services Crypto [12.3-20121019.0]
JUNOS Services SSL [12.3-20121019.0]
JUNOS Services IPSec [12.3-20121019.0]
JUNOS Runtime Software Suite [12.3-20121019.0]
JUNOS Routing Software Suite [12.3-20121019.0]
```

```
lcc0-re0:
```

```
-----
Hostname: host1
Model: t1600
JUNOS Base OS boot [12.3-20121019.0]
JUNOS Base OS Software Suite [12.3-20121019.0]
JUNOS Kernel Software Suite [12.3-20121019.0]
JUNOS Crypto Software Suite [12.3-20121019.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [12.3-20121019.0]
JUNOS Packet Forwarding Engine Support (T-Series) [12.3-20121019.0]
JUNOS Online Documentation [12.3-20121019.0]
JUNOS Services AACL Container package [12.3-20121019.0]
```

JUNOS Services Application Level Gateways [12.3-20121019.0]
JUNOS AppId Services [12.3-20121019.0]
JUNOS Border Gateway Function package [12.3-20121019.0]
JUNOS Services Captive Portal and Content Delivery Container package [12.3-20121019.0]
JUNOS Services HTTP Content Management package [12.3-20121019.0]
JUNOS IDP Services [12.3-20121019.0]
JUNOS Services LL-PDF Container package [12.3-20121019.0]
JUNOS Services NAT [12.3-20121019.0]
JUNOS Services PTSP Container package [12.3-20121019.0]
JUNOS Services RPM [12.3-20121019.0]
JUNOS Services Stateful Firewall [12.3-20121019.0]
JUNOS Voice Services Container package [12.3-20121019.0]
JUNOS Services Example Container package [12.3-20121019.0]
JUNOS Services Crypto [12.3-20121019.0]
JUNOS Services SSL [12.3-20121019.0]
JUNOS Services IPSec [12.3-20121019.0]
JUNOS Runtime Software Suite [12.3-20121019.0]
JUNOS Routing Software Suite [12.3-20121019.0]

lcc1-re0:

Hostname: host2
Model: t1600
JUNOS Base OS boot [12.3-20121019.0]
JUNOS Base OS Software Suite [12.3-20121019.0]
JUNOS Kernel Software Suite [12.3-20121019.0]
JUNOS Crypto Software Suite [12.3-20121019.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [12.3-20121019.0]
JUNOS Packet Forwarding Engine Support (T-Series) [12.3-20121019.0]
JUNOS Online Documentation [12.3-20121019.0]
JUNOS Services ACL Container package [12.3-20121019.0]
JUNOS Services Application Level Gateways [12.3-20121019.0]
JUNOS AppId Services [12.3-20121019.0]
JUNOS Border Gateway Function package [12.3-20121019.0]
JUNOS Services Captive Portal and Content Delivery Container package [12.3-20121019.0]
JUNOS Services HTTP Content Management package [12.3-20121019.0]
JUNOS IDP Services [12.3-20121019.0]
JUNOS Services LL-PDF Container package [12.3-20121019.0]
JUNOS Services NAT [12.3-20121019.0]
JUNOS Services PTSP Container package [12.3-20121019.0]
JUNOS Services RPM [12.3-20121019.0]
JUNOS Services Stateful Firewall [12.3-20121019.0]
JUNOS Voice Services Container package [12.3-20121019.0]
JUNOS Services Example Container package [12.3-20121019.0]
JUNOS Services Crypto [12.3-20121019.0]
JUNOS Services SSL [12.3-20121019.0]
JUNOS Services IPSec [12.3-20121019.0]
JUNOS Runtime Software Suite [12.3-20121019.0]
JUNOS Routing Software Suite [12.3-20121019.0]

lcc2-re0:

Hostname: host3
Model: t1600
JUNOS Base OS boot [12.3-20121019.0]
JUNOS Base OS Software Suite [12.3-20121019.0]
JUNOS Kernel Software Suite [12.3-20121019.0]
JUNOS Crypto Software Suite [12.3-20121019.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [12.3-20121019.0]

```

JUNOS Packet Forwarding Engine Support (T-Series) [12.3-20121019.0]
JUNOS Online Documentation [12.3-20121019.0]
JUNOS Services AACL Container package [12.3-20121019.0]
JUNOS Services Application Level Gateways [12.3-20121019.0]
JUNOS AppId Services [12.3-20121019.0]
JUNOS Border Gateway Function package [12.3-20121019.0]
JUNOS Services Captive Portal and Content Delivery Container package
[12.3-20121019.0]
JUNOS Services HTTP Content Management package [12.3-20121019.0]
JUNOS IDP Services [12.3-20121019.0]
JUNOS Services LL-PDF Container package [12.3-20121019.0]
JUNOS Services NAT [12.3-20121019.0]
JUNOS Services PTSP Container package [12.3-20121019.0]
JUNOS Services RPM [12.3-20121019.0]
JUNOS Services Stateful Firewall [12.3-20121019.0]
JUNOS Voice Services Container package [12.3-20121019.0]
JUNOS Services Example Container package [12.3-20121019.0]
JUNOS Services Crypto [12.3-20121019.0]
JUNOS Services SSL [12.3-20121019.0]
JUNOS Services IPSec [12.3-20121019.0]
JUNOS Runtime Software Suite [12.3-20121019.0]
JUNOS Routing Software Suite [12.3-20121019.0]

```

```
lcc3-re0:
```

```

-----
Hostname: host4
Model: t1600
JUNOS Base OS boot [12.3-20121019.0]
JUNOS Base OS Software Suite [12.3-20121019.0]
JUNOS Kernel Software Suite [12.3-20121019.0]
JUNOS Crypto Software Suite [12.3-20121019.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [12.3-20121019.0]
JUNOS Packet Forwarding Engine Support (T-Series) [12.3-20121019.0]
JUNOS Online Documentation [12.3-20121019.0]
JUNOS Services AACL Container package [12.3-20121019.0]
JUNOS Services Application Level Gateways [12.3-20121019.0]
JUNOS AppId Services [12.3-20121019.0]
JUNOS Border Gateway Function package [12.3-20121019.0]
JUNOS Services Captive Portal and Content Delivery Container package
[12.3-20121019.0]
JUNOS Services HTTP Content Management package [12.3-20121019.0]
JUNOS IDP Services [12.3-20121019.0]
JUNOS Services LL-PDF Container package [12.3-20121019.0]
JUNOS Services NAT [12.3-20121019.0]
JUNOS Services PTSP Container package [12.3-20121019.0]
JUNOS Services RPM [12.3-20121019.0]
JUNOS Services Stateful Firewall [12.3-20121019.0]
JUNOS Voice Services Container package [12.3-20121019.0]
JUNOS Services Example Container package [12.3-20121019.0]
JUNOS Services Crypto [12.3-20121019.0]
JUNOS Services SSL [12.3-20121019.0]
JUNOS Services IPSec [12.3-20121019.0]
JUNOS Runtime Software Suite [12.3-20121019.0]
JUNOS Routing Software Suite [12.3-20121019.0]

```

show version (TX Matrix Plus Router with 3D SIBs)

```

user@host>show version
sfc0-re0:

```

```

-----
Hostname: sfc0

```

```
Model: txp
JUNOS Base OS boot [13.1-20130306.0]
JUNOS Base OS Software Suite [13.1-20130306.0]
JUNOS Kernel Software Suite [13.1-20130306.0]
JUNOS Crypto Software Suite [13.1-20130306.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [13.1-20130306.0]
JUNOS Packet Forwarding Engine Support (T-Series) [13.1-20130306.0]
JUNOS Online Documentation [13.1-20130306.0]
JUNOS Services ACL Container package [13.1-20130306.0]
JUNOS Services Application Level Gateways [13.1-20130306.0]
JUNOS AppId Services [13.1-20130306.0]
JUNOS Border Gateway Function package [13.1-20130306.0]
JUNOS Services Captive Portal and Content Delivery Container package
[13.1-20130306.0]
JUNOS Services HTTP Content Management package [13.1-20130306.0]
JUNOS IDP Services [13.1-20130306.0]
JUNOS Services Jflow Container package [13.1-20130306.0]
JUNOS Services LL-PDF Container package [13.1-20130306.0]
JUNOS Services MobileNext Software package [13.1-20130306.0]
JUNOS Services Mobile Subscriber Service Container package [13.1-20130306.0]
JUNOS Services NAT [13.1-20130306.0]
JUNOS Services PTSP Container package [13.1-20130306.0]
JUNOS Services RPM [13.1-20130306.0]
JUNOS Services Stateful Firewall [13.1-20130306.0]
JUNOS Voice Services Container package [13.1-20130306.0]
JUNOS Services Example Container package [13.1-20130306.0]
JUNOS Services Crypto [13.1-20130306.0]
JUNOS Services SSL [13.1-20130306.0]
JUNOS Services IPSec [13.1-20130306.0]
JUNOS Runtime Software Suite [13.1-20130306.0]
JUNOS Routing Software Suite [13.1-20130306.0]
```

```
lcc0-re0:
```

```
-----
Hostname: lcc0
Model: t4000
JUNOS Base OS boot [13.1-20130306.0]
JUNOS Base OS Software Suite [13.1-20130306.0]
JUNOS Kernel Software Suite [13.1-20130306.0]
JUNOS Crypto Software Suite [13.1-20130306.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [13.1-20130306.0]
JUNOS Packet Forwarding Engine Support (T-Series) [13.1-20130306.0]
JUNOS Online Documentation [13.1-20130306.0]
JUNOS Services ACL Container package [13.1-20130306.0]
JUNOS Services Application Level Gateways [13.1-20130306.0]
JUNOS AppId Services [13.1-20130306.0]
JUNOS Border Gateway Function package [13.1-20130306.0]
JUNOS Services Captive Portal and Content Delivery Container package
[13.1-20130306.0]
JUNOS Services HTTP Content Management package [13.1-20130306.0]
JUNOS IDP Services [13.1-20130306.0]
JUNOS Services Jflow Container package [13.1-20130306.0]
JUNOS Services LL-PDF Container package [13.1-20130306.0]
JUNOS Services MobileNext Software package [13.1-20130306.0]
JUNOS Services Mobile Subscriber Service Container package [13.1-20130306.0]
JUNOS Services NAT [13.1-20130306.0]
JUNOS Services PTSP Container package [13.1-20130306.0]
JUNOS Services RPM [13.1-20130306.0]
JUNOS Services Stateful Firewall [13.1-20130306.0]
JUNOS Voice Services Container package [13.1-20130306.0]
JUNOS Services Example Container package [13.1-20130306.0]
```

```
JUNOS Services Crypto [13.1-20130306.0]
JUNOS Services SSL [13.1-20130306.0]
JUNOS Services IPSec [13.1-20130306.0]
JUNOS Runtime Software Suite [13.1-20130306.0]
JUNOS Routing Software Suite [13.1-20130306.0]
```

```
lcc2-re0:
```

```
-----
Hostname: lcc2
Model: t4000
JUNOS Base OS boot [13.1-20130306.0]
JUNOS Base OS Software Suite [13.1-20130306.0]
JUNOS Kernel Software Suite [13.1-20130306.0]
JUNOS Crypto Software Suite [13.1-20130306.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [13.1-20130306.0]
JUNOS Packet Forwarding Engine Support (T-Series) [13.1-20130306.0]
JUNOS Online Documentation [13.1-20130306.0]
JUNOS Services AACL Container package [13.1-20130306.0]
JUNOS Services Application Level Gateways [13.1-20130306.0]
JUNOS AppId Services [13.1-20130306.0]
JUNOS Border Gateway Function package [13.1-20130306.0]
JUNOS Services Captive Portal and Content Delivery Container package
[13.1-20130306.0]
JUNOS Services HTTP Content Management package [13.1-20130306.0]
JUNOS IDP Services [13.1-20130306.0]
JUNOS Services Jflow Container package [13.1-20130306.0]
JUNOS Services LL-PDF Container package [13.1-20130306.0]
JUNOS Services MobileNext Software package [13.1-20130306.0]
JUNOS Services Mobile Subscriber Service Container package [13.1-20130306.0]
JUNOS Services NAT [13.1-20130306.0]
JUNOS Services PTSP Container package [13.1-20130306.0]
JUNOS Services RPM [13.1-20130306.0]
JUNOS Services Stateful Firewall [13.1-20130306.0]
JUNOS Voice Services Container package [13.1-20130306.0]
JUNOS Services Example Container package [13.1-20130306.0]
JUNOS Services Crypto [13.1-20130306.0]
JUNOS Services SSL [13.1-20130306.0]
JUNOS Services IPSec [13.1-20130306.0]
JUNOS Runtime Software Suite [13.1-20130306.0]
JUNOS Routing Software Suite [13.1-20130306.0]
```

```
lcc4-re0:
```

```
-----
Hostname: lcc4
Model: t4000
JUNOS Base OS boot [13.1-20130306.0]
JUNOS Base OS Software Suite [13.1-20130306.0]
JUNOS Kernel Software Suite [13.1-20130306.0]
JUNOS Crypto Software Suite [13.1-20130306.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [13.1-20130306.0]
JUNOS Packet Forwarding Engine Support (T-Series) [13.1-20130306.0]
JUNOS Online Documentation [13.1-20130306.0]
JUNOS Services AACL Container package [13.1-20130306.0]
JUNOS Services Application Level Gateways [13.1-20130306.0]
JUNOS AppId Services [13.1-20130306.0]
JUNOS Border Gateway Function package [13.1-20130306.0]
JUNOS Services Captive Portal and Content Delivery Container package
[13.1-20130306.0]
JUNOS Services HTTP Content Management package [13.1-20130306.0]
JUNOS IDP Services [13.1-20130306.0]
```

JUNOS Services Jflow Container package [13.1-20130306.0]
JUNOS Services LL-PDF Container package [13.1-20130306.0]
JUNOS Services MobileNext Software package [13.1-20130306.0]
JUNOS Services Mobile Subscriber Service Container package [13.1-20130306.0]
JUNOS Services NAT [13.1-20130306.0]
JUNOS Services PTSP Container package [13.1-20130306.0]
JUNOS Services RPM [13.1-20130306.0]
JUNOS Services Stateful Firewall [13.1-20130306.0]
JUNOS Voice Services Container package [13.1-20130306.0]
JUNOS Services Example Container package [13.1-20130306.0]
JUNOS Services Crypto [13.1-20130306.0]
JUNOS Services SSL [13.1-20130306.0]
JUNOS Services IPSec [13.1-20130306.0]
JUNOS Runtime Software Suite [13.1-20130306.0]
JUNOS Routing Software Suite [13.1-20130306.0]

lcc6-re0:

Hostname: lcc6
Model: t1600
JUNOS Base OS boot [13.1-20130306.0]
JUNOS Base OS Software Suite [13.1-20130306.0]
JUNOS Kernel Software Suite [13.1-20130306.0]
JUNOS Crypto Software Suite [13.1-20130306.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [13.1-20130306.0]
JUNOS Packet Forwarding Engine Support (T-Series) [13.1-20130306.0]
JUNOS Online Documentation [13.1-20130306.0]
JUNOS Services AACL Container package [13.1-20130306.0]
JUNOS Services Application Level Gateways [13.1-20130306.0]
JUNOS AppId Services [13.1-20130306.0]
JUNOS Border Gateway Function package [13.1-20130306.0]
JUNOS Services Captive Portal and Content Delivery Container package [13.1-20130306.0]
JUNOS Services HTTP Content Management package [13.1-20130306.0]
JUNOS IDP Services [13.1-20130306.0]
JUNOS Services Jflow Container package [13.1-20130306.0]
JUNOS Services LL-PDF Container package [13.1-20130306.0]
JUNOS Services MobileNext Software package [13.1-20130306.0]
JUNOS Services Mobile Subscriber Service Container package [13.1-20130306.0]
JUNOS Services NAT [13.1-20130306.0]
JUNOS Services PTSP Container package [13.1-20130306.0]
JUNOS Services RPM [13.1-20130306.0]
JUNOS Services Stateful Firewall [13.1-20130306.0]
JUNOS Voice Services Container package [13.1-20130306.0]
JUNOS Services Example Container package [13.1-20130306.0]
JUNOS Services Crypto [13.1-20130306.0]
JUNOS Services SSL [13.1-20130306.0]
JUNOS Services IPSec [13.1-20130306.0]
JUNOS Runtime Software Suite [13.1-20130306.0]
JUNOS Routing Software Suite [13.1-20130306.0]

lcc7-re0:

Hostname: lcc7
Model: t1600
JUNOS Base OS boot [13.1-20130306.0]
JUNOS Base OS Software Suite [13.1-20130306.0]
JUNOS Kernel Software Suite [13.1-20130306.0]
JUNOS Crypto Software Suite [13.1-20130306.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [13.1-20130306.0]

```

JUNOS Packet Forwarding Engine Support (T-Series) [13.1-20130306.0]
JUNOS Online Documentation [13.1-20130306.0]
JUNOS Services AACL Container package [13.1-20130306.0]
JUNOS Services Application Level Gateways [13.1-20130306.0]
JUNOS AppId Services [13.1-20130306.0]
JUNOS Border Gateway Function package [13.1-20130306.0]
JUNOS Services Captive Portal and Content Delivery Container package
[13.1-20130306.0]
JUNOS Services HTTP Content Management package [13.1-20130306.0]
JUNOS IDP Services [13.1-20130306.0]
JUNOS Services Jflow Container package [13.1-20130306.0]
JUNOS Services LL-PDF Container package [13.1-20130306.0]
JUNOS Services MobileNext Software package [13.1-20130306.0]
JUNOS Services Mobile Subscriber Service Container package [13.1-20130306.0]
JUNOS Services NAT [13.1-20130306.0]
JUNOS Services PTSP Container package [13.1-20130306.0]
JUNOS Services RPM [13.1-20130306.0]
JUNOS Services Stateful Firewall [13.1-20130306.0]
JUNOS Voice Services Container package [13.1-20130306.0]
JUNOS Services Example Container package [13.1-20130306.0]
JUNOS Services Crypto [13.1-20130306.0]
JUNOS Services SSL [13.1-20130306.0]
JUNOS Services IPSec [13.1-20130306.0]
JUNOS Runtime Software Suite [13.1-20130306.0]
JUNOS Routing Software Suite [13.1-20130306.0]

```

show version (MX Series Router)

```

user@host5> show version
Hostname: host5
Model: mx80
JUNOS Base OS boot [11.3-20110717.0]
JUNOS Base OS Software Suite [11.3-20110717.0]
JUNOS Kernel Software Suite [11.3-20110717.0]
JUNOS Crypto Software Suite [11.3-20110717.0]
JUNOS Packet Forwarding Engine Support (MX80) [11.3-20110717.0]
JUNOS Online Documentation [11.3-20110717.0]
JUNOS Routing Software Suite [11.3-20110717.0]

```

show version (QFX3500 Switch)

```

user@switch> show version
Hostname: switch
Model: qfx_s3500
JUNOS Base OS boot [11.1R1]
JUNOS Base OS Software Suite [11.1R1]
JUNOS Kernel Software Suite [11.1R1]
JUNOS Crypto Software Suite [11.1R1]
JUNOS Online Documentation [11.1R1]
JUNOS Enterprise Software Suite [11.1R1]
JUNOS Packet Forwarding Engine Support (QFX) [11.1R1]
JUNOS Routing Software Suite [11.1R1]

```

show version (QFabric System)

```

user@qfabric> show version
Hostname: qfabric
Model: qfx3000-g
Serial Number: qfsn-0123456789
QFabric System ID: f158527a-f99e-11e0-9fbd-00e081c57cda
JUNOS Base Version [12.2I20111018_0215_dc-builder]

```

show version component all (QFabric System)

```
user@switch> show version component all
dg1:
-
Hostname: qfabric
Model: qfx3100
JUNOS Base Version [11.3R1.6]

dg0:
-
Hostname: qfabric
Model: qfx3100
JUNOS Base Version [11.3R1.6]

NW-NG-0:
-
Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3R1.6]
JUNOS Base OS Software Suite [11.3R1.6]
JUNOS Kernel Software Suite [11.3R1.6]
JUNOS Crypto Software Suite [11.3R1.6]
JUNOS Online Documentation [11.3R1.6]
JUNOS Enterprise Software Suite [11.3R1.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3R1.6]
JUNOS Routing Software Suite [11.3R1.6]

FC-0:
-
Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3R1.6]
JUNOS Base OS Software Suite [11.3R1.6]
JUNOS Kernel Software Suite [11.3R1.6]
JUNOS Crypto Software Suite [11.3R1.6]
JUNOS Online Documentation [11.3R1.6]
JUNOS Enterprise Software Suite [11.3R1.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3R1.6]
JUNOS Routing Software Suite [11.3R1.6]

FC-1:
Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3R1.6]
JUNOS Base OS Software Suite [11.3R1.6]
JUNOS Kernel Software Suite [11.3R1.6]
JUNOS Crypto Software Suite [11.3R1.6]
JUNOS Online Documentation [11.3R1.6]
JUNOS Enterprise Software Suite [11.3R1.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3R1.6]
JUNOS Routing Software Suite [11.3R1.6]

DRE-0:
-
Hostname: dre-0
Model: qfx-jvre
JUNOS Base OS boot [11.3R1.6]
JUNOS Base OS Software Suite [11.3R1.6]
JUNOS Kernel Software Suite [11.3R1.6]
JUNOS Crypto Software Suite [11.3R1.6]
```



```

JUNOS Online Documentation [11.3R1.6]
JUNOS Enterprise Software Suite [11.3R1.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3R1.6]
JUNOS Routing Software Suite [11.3R1.6]

```

```
FM-0:
```

```

-
Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3R1.6]
JUNOS Base OS Software Suite [11.3R1.6]
JUNOS Kernel Software Suite [11.3R1.6]
JUNOS Crypto Software Suite [11.3R1.6]
JUNOS Online Documentation [11.3R1.6]
JUNOS Enterprise Software Suite [11.3R1.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3R1.6]
JUNOS Routing Software Suite [11.3R1.6]

```

```
nodedevice1:
```

```

-
Hostname: qfabric
Model: QFX3500
JUNOS Base OS boot [11.3R1.6]
JUNOS Base OS Software Suite [11.3R1.6]
JUNOS Kernel Software Suite [11.3R1.6]
JUNOS Crypto Software Suite [11.3R1.6]
JUNOS Online Documentation [11.3R1.6]
JUNOS Enterprise Software Suite [11.3R1.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3R1.6]
JUNOS Routing Software Suite [11.3R1.6]

```


```
interconnectdevice1:
```

```

-
Hostname: qfabric
Model: QFX3108
JUNOS Base OS boot [11.3R1.6]
JUNOS Base OS Software Suite [11.3R1.6]
JUNOS Kernel Software Suite [11.3R1.6]
JUNOS Crypto Software Suite [11.3R1.6]
JUNOS Online Documentation [11.3R1.6]
JUNOS Enterprise Software Suite [11.3R1.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3R1.6]
JUNOS Routing Software Suite [11.3R1.6]
warning: from interconnectdevice0: Disconnected

```

start shell

Syntax	<code>start shell (csh sh)</code> <code><user username></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Exit from the CLI environment and create a UNIX-level shell. To return to the CLI, type exit from the shell.
<div> NOTE:<ul style="list-style-type: none">To issue this command, the user must have the required login access privileges configured by including the <code>permissions</code> statement at the <code>[edit system login class class-name]</code> hierarchy level.UNIX wheel group membership or permissions are no longer required to issue this command.</div>	
Options	<code>csh</code> —Create a UNIX C shell. <code>sh</code> —Create a UNIX Bourne shell. <code>user username</code> —(Optional) Start the shell as another user.
Additional Information	When you are in the shell, the shell prompt has the following format: <code>username@hostname%</code> An example of the prompt is: <code>root@host%</code>
Required Privilege Level	shell and maintenance
List of Sample Output	start shell csh on page 1266
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

start shell csh

```
user@host> start shell csh
%
exit
%
```

```
username@hostname% start shell sh
%

exit
user@host>
```

test configuration

Syntax	<code>test configuration <i>filename</i></code> <code>syntax-only</code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. syntax-only option introduced in Junos OS Release 12.1. Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Description	Verify that the syntax of a configuration file is correct. If the configuration contains any syntax or commit check errors, a message is displayed to indicate the line number and column number in which the error was found. This command only accepts text files.
Options	<i>filename</i> —Name of the configuration file. syntax-only —(Optional) Check the syntax of a partial configuration file, without checking for commit errors.
Required Privilege Level	view
List of Sample Output	test configuration on page 1268
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

test configuration

```
user@host> test configuration terminal
[Type ^D to end input]
system {
host-name bluesky;
paris-23;
login;
}
terminal:3:(8) syntax error: paris
[edit system]
  'paris-23;'
    syntax error
terminal:4:(11) statement must contain additional statements: ;
[edit system login]
  'login ;'
    statement must contain additional statements
configuration syntax failed
```

traceroute

List of Syntax [Syntax on page 1269](#)
 [Syntax \(QFX Series and OCX Series\) on page 1269](#)

Syntax `traceroute host`
 `<as-number-lookup>`
 `<bypass-routing>`
 `<clns>`
 `<gateway address>`
 `<inet | inet6>`
 `<interface interface-name>`
 `<logical system logical-system-name>`
 `<monitor host>`
 `<mpls (ldp FEC address | rsvp label-switched-path-name)>`
 `<no-resolve>`
 `<propagate-ttl>`
 `<routing-instance routing-instance-name>`
 `<source source-address>`
 `<tos value>`
 `<ttl value>`
 `<wait seconds>`

Syntax (QFX Series and OCX Series) `traceroute host`
 `<as-number-lookup>`
 `<bypass-routing>`
 `<gateway address>`
 `<inet>`
 `<inet6>`
 `<interface interface-name>`
 `<monitor host>`
 `<no-resolve>`
 `<routing-instance routing-instance-name>`
 `<source source-address>`
 `<tos value>`
 `<ttl value>`
 `<wait seconds>`

Release Information Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.0 for EX Series switches.
 mpls option introduced in Junos OS Release 9.2.
 Command introduced in Junos OS Release 11.1 for the QFX Series.
 propagate-ttl option introduced in Junos OS Release 12.1.
 Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Display the route that packets take to a specified network host. Use **traceroute** as a debugging tool to locate points of failure in a network.

Options **host**—IP address or name of remote host.

as-number-lookup—(Optional) Display the autonomous system (AS) number of each intermediate hop on the path from the host to the destination.

bypass-routing—(Optional) Bypass the normal routing tables and send requests directly to a system on an attached network. If the system is not on a directly attached network, an error is returned. Use this option to display a route to a local system through an interface that has no route through it.

clns—(Optional) Trace the route belonging to the Connectionless Network Service (CLNS).

gateway address—(Optional) Address of a router or switch through which the route transits.

inet | inet6—(Optional) Trace the route belonging to IPv4 or IPv6, respectively.

interface *interface-name*—(Optional) Name of the interface over which to send packets.

logical-system *logical-system-name*—(Optional) Perform this operation on all logical systems or on a particular logical system.

monitor *host*—(Optional) Display real-time monitoring information for the specified host.

mpls (*ldp FEC address* | *rsvp label-switched-path name*)—(Optional) See [traceroute mpls ldp](#) and [traceroute mpls rsvp](#).

no-resolve—(Optional) Do not attempt to determine the hostname that corresponds to the IP address.

propagate-ttl—(Optional) On the PE routing device, use this option to view locally generated Routing Engine transit traffic. This is applicable for MPLS L3VPN traffic only.

Use for troubleshooting, when you want to view hop-by-hop information from the local provider router to the remote provider router, when TTL decrementing is disabled on the core network using the **no-propagate-ttl** configuration statement.



NOTE: Using **propagate-ttl** with **traceroute** on the CE router does not show hop-by-hop information.

routing-instance *routing-instance-name*—(Optional) Name of the routing instance for the traceroute attempt.

source *source-address*—(Optional) Source address of the outgoing traceroute packets.

tos *value*—(Optional) Value to include in the IP type-of-service (ToS) field. The range of values is 0 through 255.

ttl *value*—(Optional) Maximum time-to-live value to include in the traceroute request. The range of values is 0 through 128.

wait *seconds*—(Optional) Maximum time to wait for a response to the traceroute request.

Required Privilege Level network

Related Documentation

- [traceroute monitor on page 1273](#)

List of Sample Output

[traceroute on page 1271](#)
[traceroute as-number-lookup host on page 1271](#)
[traceroute no-resolve on page 1271](#)
[traceroute propagate-ttl on page 1272](#)
[traceroute \(Between CE Routers, Layer 3 VPN\) on page 1272](#)
[traceroute \(Through an MPLS LSP\) on page 1272](#)

Output Fields

Table 83 describes the output fields for the **traceroute** command. Output fields are listed in the approximate order in which they appear.

Table 83: traceroute Output Fields

Field Name	Field Description
traceroute to	IP address of the receiver.
hops max	Maximum number of hops allowed.
byte packets	Size of packets being sent.
<i>number-of-hops</i>	Number of hops from the source to the named router or switch.
<i>router-name</i>	Name of the router or switch for this hop.
<i>address</i>	Address of the router or switch for this hop.
Round trip time	Average round-trip time, in milliseconds (ms).

Sample Output

traceroute

```

user@host> traceroute santacruz
traceroute to green.company.net (10.156.169.254), 30 hops max, 40 byte packets
 1 blue23 (10.168.1.254)  2.370 ms  2.853 ms  0.367 ms
 2 red14 (10.168.255.250) 0.778 ms  2.937 ms  0.446 ms
 3 yellow (10.156.169.254) 7.737 ms  89.905 ms  0.834 ms

```

traceroute as-number-lookup host

```

user@host> traceroute as-number-lookup 10.100.1.1
traceroute to 10.100.1.1 (10.100.1.1), 30 hops max, 40 byte packets
 1 10.39.1.1 (10.39.1.1) 0.779 ms 0.728 ms 0.562 ms
 2 10.39.1.6 (10.39.1.6) [AS 32] 0.657 ms 0.611 ms 0.617 ms
 3 10.100.1.1 (10.100.1.1) [AS 10, 40, 50] 0.880 ms 0.808 ms 0.774 ms

```

traceroute no-resolve

```

user@host> traceroute santacruz no-resolve

```

```
traceroute to green.company.net (10.156.169.254), 30 hops max, 40 byte packets
 1  10.168.1.254  0.458 ms  0.370 ms  0.365 ms
 2  10.168.255.250  0.474 ms  0.450 ms  0.444 ms
 3  10.156.169.254  0.931 ms  0.876 ms  0.862 ms
```

traceroute propagate-ttl

```
user@host> traceroute propagate-ttl 100.200.2.2 routing-instance VPN-A
traceroute to 100.200.2.2 (100.200.2.2) from 1.1.0.2, 30 hops max, 40 byte packets

 1  1.2.0.2 (1.2.0.2)  2.456 ms  1.753 ms  1.672 ms
    MPLS Label=299776 CoS=0 TTL=1 S=0
    MPLS Label=299792 CoS=0 TTL=1 S=1
 2  1.3.0.2 (1.3.0.2)  1.213 ms  1.225 ms  1.166 ms
    MPLS Label=299792 CoS=0 TTL=1 S=1
 3  100.200.2.2 (100.200.2.2)  1.422 ms  1.521 ms  1.443 ms
```

traceroute (Between CE Routers, Layer 3 VPN)

```
user@host> traceroute vpn09
traceroute to vpn09.skybank.net (10.255.14.179), 30 hops max, 40
byte packets
 1  10.39.10.21 (10.39.10.21)  0.598 ms  0.500 ms  0.461 ms
 2  10.39.1.13 (10.39.1.13)  0.796 ms  0.775 ms  0.806 ms
    MPLS Label=100006 CoS=0 TTL=1 S=1
 3  vpn09.skybank.net (10.255.14.179)  0.783 ms  0.716 ms  0.686
```

traceroute (Through an MPLS LSP)

```
user@host> traceroute mpls1
traceroute to 10.168.1.224 (10.168.1.224), 30 hops max, 40 byte packets
 1  mpls1-sr0.company.net (10.168.200.101)  0.555 ms  0.393 ms  0.367 ms
    MPLS Label=1024 CoS=0 TTL=1
 2  mpls5-lo0.company.net (10.168.1.224)  0.420 ms  0.394 ms  0.401 ms
```


traceroute monitor

List of Syntax	Syntax on page 1273 Syntax (QFX Series) on page 1273
Syntax	<pre>traceroute monitor <i>host</i> <count <i>value</i>> <inet inet6> <interval <i>seconds</i>> <no resolve> <size <i>value</i>> <source <i>source-address</i>> <summary></pre>
Syntax (QFX Series)	<pre>traceroute monitor <i>host</i> <count <i>value</i>> <inet> <inet6> <interval <i>seconds</i>> <no resolve> <size <i>value</i>> <source <i>source-address</i>> <summary></pre>
Release Information	<p>Command introduced in Junos OS Release 8.0</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Display live monitoring of each hop in the route that packets take to a specified network host. Use as a debugging tool to locate points of failure in a network.
Options	<p><i>host</i>—IP address or name of remote host.</p> <p><i>count value</i>—Number of ping requests, in packets, to send in summary mode. The default value is 10.</p> <p><i>inet inet6</i>—(Optional) Trace the route belonging to IPv4 or IPv6, respectively.</p> <p><i>interval seconds</i>—(Optional) Number of seconds to wait before sending ping requests. The default value is 1.</p> <p><i>no resolve</i>—(Optional) Do not attempt to display addresses symbolically.</p> <p><i>size value</i>—(Optional) Receive the specified number of bytes for each packet. The range is 0 through 65468 bytes. The default value is 64.</p> <p><i>source source-address</i>—(Optional) Source address of the outgoing ping packets.</p> <p><i>summary</i>—(Optional) Generate and display a summary of live monitoring of each hop on the route that packets take to a specified network host.</p>

Required Privilege Level network

List of Sample Output [traceroute monitor on page 1274](#)

Output Fields Table 84 describes the output fields for the **traceroute monitor** command. Output fields are listed in the approximate order in which they appear.

Table 84: traceroute monitor Output Fields

Field Name	Field Description
Host	Hostname or IP address of the router at each hop.
Loss%	Percent of packet loss. The number of ping responses divided by the number of ping requests, specified as a percentage.
Snt	Number of ping requests sent to the router at this hop.
Last	Most recent round-trip time, in milliseconds, to the router at this hop.
Avg	Average round-trip time, in milliseconds, to the router at this hop.
Best	Shortest round-trip time, in milliseconds, to the router at this hop.
Wrst	Longest round-trip time, in milliseconds, to the router at this hop.
StDev	Standard deviation of round-trip times, in milliseconds, to the router at this hop.

Sample Output

traceroute monitor

```
user@host> traceroute monitor 10.16.0.1
```

	Loss%	Snt	Last	Avg	Best	Wrst	StDev
Host							
1. 10.17.41.254	0.0%	17	0.7	1.0	0.6	5.4	1.2
2. secret.net	0.0%	17	0.6	1.0	0.6	6.6	1.4
3. top-secret.net	0.0%	17	0.6	0.6	0.6	0.6	0.0

CHAPTER 32

Standard Software Installation and Upgrade Operational Commands

- request system software add
- request system software delete
- request system software download
- request system software rollback
- request system software validate
- rollback
- show system rollback

request system software add

List of Syntax [Syntax on page 1276](#)
 [Syntax \(EX Series Switches\) on page 1276](#)
 [Syntax \(TX Matrix Router\) on page 1276](#)
 [Syntax \(TX Matrix Plus Router\) on page 1277](#)
 [Syntax \(MX Series Router\) on page 1277](#)
 [Syntax \(QFX Series\) on page 1277](#)
 [Syntax \(OCX Series\) on page 1277](#)

Syntax request system software add *package-name*
 <best-effort-load>
 <delay-restart>
 <device-alias *alias-name*>
 <force>
 <no-copy>
 <no-validate>
 <re0 | re1>
 <reboot>
 <satellite *slot-id*>
 <set [*package-name package-name*]>
 <unlink>
 <upgrade-group [all | *upgrade-group-name*]>
 <upgrade-with-config>
 <upgrade-with-config-format *format*>
 <satellite *slot-id*>
 <validate>
 <version *version-string*>

Syntax (EX Series Switches) request system software add *package-name*
 <best-effort-load>
 <delay-restart>
 <force>
 <no-copy>
 <no-validate>
 <re0 | re1>
 <reboot>
 <set [*package-name package-name*]>
 <upgrade-with-config>
 <upgrade-with-config-format *format*>
 <validate>

Syntax (TX Matrix Router) request system software add *package-name*
 <best-effort-load>
 <delay-restart>
 <force>
 <lcc *number* | scc>
 <no-copy>
 <no-validate>
 <re0 | re1>
 <reboot>
 <set [*package-name package-name*]>
 <unlink>

	<upgrade-with-config> <upgrade-with-config-format <i>format</i> > <validate>
Syntax (TX Matrix Plus Router)	request system software add <i>package-name</i> <best-effort-load> <delay-restart> <force> <lcc <i>number</i> sfc <i>number</i> > <no-copy> <no-validate> <re0 re1> <reboot> <set [<i>package-name package-name</i>]> <unlink> <upgrade-with-config> <upgrade-with-config-format <i>format</i> > <validate>
Syntax (MX Series Router)	request system software add <i>package-name</i> <best-effort-load> <delay-restart> <device-alias <i>alias-name</i> > <force> <member <i>member-id</i> > <no-copy> <no-validate> <re0 re1> <reboot> <satellite <i>slot-id</i> > <set [<i>package-name package-name</i>]> <upgrade-group [all <i>upgrade-group-name</i>]> <unlink> <upgrade-with-config> <upgrade-with-config-format <i>format</i> > <validate> <version <i>version-string</i> >
Syntax (QFX Series)	request system software add <i>package-name</i> <best-effort-load> <component all> <delay-restart> <force> <force-host> <no-copy> <no-validate> <partition> <reboot> <unlink> <upgrade-with-config> <upgrade-with-config-format <i>format</i> > <validate>
Syntax (OCX Series)	request system software add <i>package-name</i>

<best-effort-load>
<delay-restart>
<force>
<force-host>
<no-copy>
<no-validate>
<reboot>
<unlink>
<upgrade-with-config>
<upgrade-with-config-format *format*>
<validate>

Release Information

Command introduced before Junos OS Release 7.4.

best-effort-load and **unlink** options added in Junos OS Release 7.4.

Command introduced in Junos OS Release 9.0 for EX Series switches.

sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.

Command introduced in Junos OS Release 11.1 for the QFX Series.

set [package-name package-name] option added in Junos OS Release 11.1 for EX Series switches.

set [package-name package-name] option added in Junos OS Release 12.2 for M Series, MX Series, T Series routers, and Branch SRX Series Services Gateways.



NOTE: On EX Series switches, the **set [package-name package-name]** option allows you to install only two software packages on a mixed EX4200 and EX4500 Virtual Chassis, whereas, on M Series, MX Series, T Series routers, and Branch SRX Series Services Gateways, the **set [package-name package-name]** option allows you to install multiple software packages and software add-on packages at the same time.

upgrade-with-config and **upgrade-with-config-format *format*** options added in Junos OS Release 12.3 for M Series routers, MX Series routers, T Series routers, EX Series Ethernet switches, and QFX Series devices.

Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

device-alias, **satellite**, **upgrade-group**, and **version** options introduced in Junos OS Release 14.2R3 for Junos Fusion.

Description

NOTE: We recommend that you always download the software image to **/var/tmp** only. On EX Series and QFX Series switches, you must use the **/var/tmp** directory. Other directories are not supported.

Install a software package or bundle on the router or switch.



WARNING: Any configuration changes performed after inputting the **request system software add** command will be lost when the system reboots with an upgraded version of Junos OS.



NOTE: When graceful Routing Engine switchover (GRES) is enabled on a device, you must perform a unified ISSU operation to update the software running on the device. With GRES enabled, if you attempt to perform a software upgrade by entering the request system software add *package-name* command, an error message is displayed stating that only in-service-software-upgrades are supported when GRES is configured. In such a case, you must either remove the GRES configuration before you attempt the upgrade or perform a unified ISSU.

Options *package-name*—Location from which the software package or bundle is to be installed.
For example:

- */var/tmp/package-name*—For a software package or bundle that is being installed from a local directory on the router or switch.
- *protocol://hostname/pathname/package-name*—For a software package or bundle that is to be downloaded and installed from a remote location. Replace *protocol* with one of the following:
 - **ftp**—File Transfer Protocol.
Use *ftp://hostname/pathname/package-name*. To specify authentication credentials, use *ftp://<username>:<password>@hostname/pathname/package-name*. To have the system prompt you for the password, specify **prompt** in place of the password. If a password is required, and you do not specify the password or **prompt**, an error message is displayed.
 - **http**—Hypertext Transfer Protocol.
Use *http://hostname/pathname/package-name*. To specify authentication credentials, use *http://<username>:<password>@hostname/pathname/package-name*. If a password is required and you omit it, you are prompted for it.
 - **scp**—Secure copy (available only for Canada and U.S. version).
Use *scp://hostname/pathname/package-name*. To specify authentication credentials, use *scp://<username>:<password>@hostname/pathname/package-name*.



NOTE:

- The *pathname* in the protocol is the relative path to the user's home directory on the remote system and not the root directory.
- Do not use the `scp` protocol in the `request system software add` command to download and install a software package or bundle from a remote location. The previous statement does not apply to the QFabric switch. The software upgrade is handled by the MGD process which does not support `scp`.

Use the `file copy` command to copy the software package or bundle from the remote location to the `/var/tmp` directory on the hard disk:

`file copy scp://source/package-name /var/tmp`

Then install the software package or bundle using the `request system software add` command:

`request system software add /var/tmp/package-name`

best-effort-load—(Optional) Activate a partial load and treat parsing errors as warnings instead of errors.

component all—(QFabric systems only) (Optional) Install software package on all of the QFabric components.

delay-restart—(Optional) Install a software package or bundle, but do not restart software processes.

device-alias *alias-name*—(Junos Fusion only) (Optional) Install the satellite software package onto the specified satellite device using the satellite device's alias name.

force—(Optional) Force the addition of the software package or bundle (ignore warnings).

force-host—(Optional) Force the addition of host software package or bundle (ignore warnings) on the QFX5100 device.

lcc *number*—(TX Matrix routers and TX Matrix Plus routers only) (Optional) In a routing matrix based on the TX Matrix router, install a software package or bundle on a T640 router that is connected to the TX Matrix router. In a routing matrix based on the TX Matrix Plus router, install a software package or bundle on a router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.

- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

member *member-id*—(MX Series routers only) (Optional) Install a software package on the specified Virtual Chassis member. Replace *member-id* with a value of 0 or 1.

partition—(QFX3500 switches only) (Optional) Format and repartition the media before installation.

satellite *slot-id*—(Junos Fusion only) (Optional) Install the satellite software package onto the specified satellite device using the satellite devices FPC slot identifier.

scc—(TX Matrix routers only) (Optional) Install a software package or bundle on a Routing Engine on a TX Matrix router (or switch-card chassis).

sfc *number*—(TX Matrix Plus routers only) (Optional) Install a software package or bundle on a Routing Engine on a TX Matrix Plus router. Replace *number* with 0.

no-copy—(Optional) Install a software package or bundle, but do not save copies of the package or bundle files.

no-validate—(Optional) When loading a software package or bundle with a different release, suppress the default behavior of the **validate** option.

re0 | re1—(Optional) On routers or switches that support dual or redundant Routing Engines, load a software package or bundle on the Routing Engine in slot 0 (re0) or the Routing Engine in slot 1 (re1).

reboot—(Optional) After adding the software package or bundle, reboot the system. On a QFabric switch, the software installation is not complete until you reboot the component for which you have installed the software.

set [*package-name package-name*]—(Mixed EX4200 and EX4500 Virtual Chassis only) (Optional) Install two software packages—a package for an EX4200 switch and the same release of the package for an EX4500 switch—to upgrade all member switches in a mixed EX4200 and EX4500 Virtual Chassis.

set [*package-name package-name*]—(M Series, MX Series, T Series routers, and Branch SRX Series Services Gateways only) (Optional) Install multiple software packages and software add-on packages at the same time.

unlink—(Optional) On M Series, T Series, and MX Series routers, use the unlink option to remove the software package from this directory after a successful upgrade is completed.

upgrade-group [all |*upgrade-group-name*]—(Junos Fusion only) (Required to configure a Junos Fusion using autoconversion or manual conversion) Associate a satellite software image with a satellite software upgrade group. The satellite software package is associated with the specified satellite software upgrade group using the

upgrade-group-name, or for all satellite software upgrade groups in a Junos Fusion when the *all* keyword is specified.

A satellite software upgrade group is a group of satellite devices in a Junos Fusion that are designated to upgrade to the same satellite software version using the same satellite software package. See *Understanding Software in a Junos Fusion* and *Managing Satellite Software Upgrade Groups in a Junos Fusion*.

upgrade-with-config—(Optional) Install one or more configuration files.

upgrade-with-config-format *format*—(Optional) Specify the configuration file format, **text** or **xml**. The default format is **text**.



NOTE: The **upgrade-with-config** and **upgrade-with-config-format** options are only available locally on the router or switch. In a routing matrix, the configuration is applied only to the local router and is not propagated to other routers.

The options are validated during the validation process and applied to the router or switch during the upgrade process. If the upgrade process is successful, the options are removed from the configuration. If the upgrade process fails, the configuration file is renamed with the **.failed** suffix.

validate—(Optional) Validate the software package or bundle against the current configuration as a prerequisite to adding the software package or bundle. This is the default behavior when the software package or bundle being added is a different release.



NOTE: The **validate** option only works on systems that do not have graceful-switchover (GRES) enabled. To use the **validate** option on a system with GRES, either disable GRES for the duration of the installation, or install using the command **request system software in-service-upgrade**, which requires nonstop active routing (NSR) to be enabled when using GRES.

version *version-string*—(Junos Fusion only) (Optional) Associate a satellite software package with a satellite software upgrade group by selecting the satellite software package's version. This option can only be used if the specified version of the satellite software has previously been installed on the aggregation device.

Additional Information Before upgrading the software on the router or switch, when you have a known stable system, issue the **request system snapshot** command to back up the software, including the configuration, to the **/altroot** and **/altconfig** file systems. After you have upgraded the software on the router or switch and are satisfied that the new package or bundle is

successfully installed and running, issue the **request system snapshot** command again to back up the new software to the **/altroot** and **/altconfig** file systems.



NOTE: The **request system snapshot** command is currently not supported on the QFabric system. Also, you cannot add or install multiple packages on a QFabric system.

After you run the **request system snapshot** command, you cannot return to the previous version of the software, because the running and backup copies of the software are identical.

If you are upgrading more than one package at the same time, delete the operating system package, **jkernl**, last. Add the operating system package, **jkernl**, first and the routing software package, **jroute**, last. If you are upgrading all packages at once, delete and add them in the following order:

```
user@host> request system software add /var/tmp/jbase
user@host> request system software add /var/tmp/jkernl
user@host> request system software add /var/tmp/jpfe
user@host> request system software add /var/tmp/jdocs
user@host> request system software add /var/tmp/jroute
user@host> request system software add /var/tmp/jcrypto
```

By default, when you issue the **request system software add *package-name*** command on a TX Matrix master Routing Engine, all the T640 master Routing Engines that are connected to it are upgraded to the same version of software. If you issue the same command on the TX Matrix backup Routing Engine, all the T640 backup Routing Engines that are connected to it are upgraded to the same version of software.

Likewise, when you issue the **request system software add *package-name*** command on a TX Matrix Plus master Routing Engine, all the T1600 or T4000 master Routing Engines that are connected to it are upgraded to the same version of software. If you issue the same command on the TX Matrix Plus backup Routing Engine, all the T1600 or T4000 backup Routing Engines that are connected to it are upgraded to the same version of software.

Required Privilege Level maintenance

Related Documentation

- [request system software delete on page 1286](#)
- [request system software rollback on page 1292](#)
- [request system storage cleanup on page 399](#)
- [Upgrading Software on page 189](#)
- [Upgrading Software on a QFabric System](#)
- [Managing Satellite Software Upgrade Groups in a Junos Fusion](#)
- [request system software add \(Maintenance\)](#)

- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

List of Sample Output	request system software add validate on page 1284
	request system software add (Mixed EX4200 and EX4500 Virtual Chassis) on page 1285
	request system software add component all (QFabric Systems) on page 1285
	request system software add upgrade-group (Junos Fusion) on page 1285
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

[request system software add validate](#)

```
user@host> request system software add validate /var/tmp/jinstall-7.2R1.7-domestic-signed.tgz
Checking compatibility with configuration
Initializing...
Using jbase-7.1R2.2
Using /var/tmp/jinstall-7.2R1.7-domestic-signed.tgz
Verified jinstall-7.2R1.7-domestic.tgz signed by PackageProduction_7_2_0
Using /var/validate/tmp/jinstall-signed/jinstall-7.2R1.7-domestic.tgz
Using /var/validate/tmp/jinstall/jbundle-7.2R1.7-domestic.tgz
Checking jbundle requirements on /
Using /var/validate/tmp/jbundle/jbase-7.2R1.7.tgz
Using /var/validate/tmp/jbundle/jkernel-7.2R1.7.tgz
Using /var/validate/tmp/jbundle/jcrypto-7.2R1.7.tgz
Using /var/validate/tmp/jbundle/jpfe-7.2R1.7.tgz
Using /var/validate/tmp/jbundle/jdocs-7.2R1.7.tgz
Using /var/validate/tmp/jbundle/jroute-7.2R1.7.tgz
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
Validating against /config/rescue.conf.gz
mgd: commit complete
Validation succeeded
Installing package '/var/tmp/jinstall-7.2R1.7-domestic-signed.tgz' ...
Verified jinstall-7.2R1.7-domestic.tgz signed by PackageProduction_7_2_0
Adding jinstall...

WARNING: This package will load JUNOS 7.2R1.7 software.
WARNING: It will save JUNOS configuration files, and SSH keys
WARNING: (if configured), but erase all other files and information
WARNING: stored on this machine. It will attempt to preserve dumps
WARNING: and log files, but this can not be guaranteed. This is the
WARNING: pre-installation stage and all the software is loaded when
WARNING: you reboot the system.

Saving the config files ...
Installing the bootstrap installer ...

WARNING: A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING: 'request system reboot' command when software installation is
WARNING: complete. To abort the installation, do not reboot your system,
WARNING: instead use the 'request system software delete jinstall'
WARNING: command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-7.2R1.7-domestic-signed.tgz ...
Saving state for rollback ...
```

Sample Output

request system software add (Mixed EX4200 and EX4500 Virtual Chassis)

```
user@switch> request system software add set
[/var/tmp/jinstall-ex-4200-11.1R1.1-domestic-signed.tgz
/var/tmp/jinstall-ex-4500-11.1R1.1-domestic-signed.tgz]
...
```

request system software add component all (QFabric Systems)

```
user@switch> request system software add /pbdata/packages/jinstall-qfabric-12.2X50-D1.3.rpm
component all
...
```

request system software add upgrade-group (Junos Fusion)

```
user@aggregation-device> request system software add /var/tmp/satellite-1.0R1.1-signed.tgz
upgrade-group group1
```

request system software delete

List of Syntax	Syntax on page 1286 Syntax (TX Matrix Router) on page 1286 Syntax (TX Matrix Plus Router) on page 1286
Syntax	<pre>request system software delete <i>software-package</i> <force> <reboot> <set [<i>package-name package-name</i>]> <upgrade-group [<i>all upgrade-group-name</i>]> <version <i>version-string</i>></pre>
Syntax (TX Matrix Router)	<pre>request system software delete <i>software-package</i> <force> <lcc <i>number</i> scc> <reboot> <set [<i>package-name package-name</i>]></pre>
Syntax (TX Matrix Plus Router)	<pre>request system software delete <i>software-package</i> <force> <lcc <i>number</i> sfc <i>number</i>> <reboot> <set [<i>package-name package-name</i>]></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Option sfc introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Option set [<i>package-name package-name</i>] added in Junos OS Release 12.2 for M Series, MX Series, T Series routers, and Branch SRX Services Gateways.</p> <p>Option reboot introduced in Junos OS Release 12.3.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> <p>Options upgrade-group, and version introduced in Junos OS Release 14.2R3 for Junos Fusion.</p>
Description	Remove a software package or bundle from the router or switch.



CAUTION: Before removing a software package or bundle, make sure that you have already placed the new software package or bundle that you intend to load onto the router or switch.

- Options** *software-package*—Software package or bundle name. You can delete any or all of the following software bundles or packages:
- **jbase**—(Optional) Junos base software suite
 - **crypto**—(Optional, in domestic version only) Junos security software
 - **docs**—(Optional) Junos online documentation file

- **jkernel**—(Optional) Junos kernel software suite
- **jpfe**—(Optional) Junos Packet Forwarding Engine support
- **jroute**—(Optional) Junos routing software suite
- **junos**—(Optional) Junos base software



NOTE: On EX Series switches, some of the package names are different than those listed. To see the list of packages that you can delete on an EX Series switch, enter the command **show system software**.

force—(Optional) Ignore warnings and force removal of the software.

lcc number—(TX Matrix routers and TX Matrix Plus routers only) (Optional) In a routing matrix, delete a software package or bundle on a T640 router indicated by **lcc number** that is connected to the TX Matrix router. In a routing matrix, delete a software package or bundle on a router indicated by **lcc number** that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

re0 | re1—(Optional) On routers or switches that support dual or redundant Routing Engines, delete a software package or bundle on the Routing Engine in slot 0 (re0) or the Routing Engine in slot 1 (re1).

reboot—As of Junos OS 12.3 and greater, automatically reboot upon completing the **request system software delete** command.

scc—(TX Matrix routers only) (Optional) Remove an extension or upgrade package from the TX Matrix router (or switch-card chassis).

set [package-name package-name]—(M Series, MX Series, T Series routers, and Branch SRX Series Services Gateways only) (Optional) Install multiple software packages or software add-on packages at the same time.

sfc number—(TX Matrix Plus routers only) (Optional) Remove an extension or upgrade package from the TX Matrix Plus router. Replace *number* with 0.

upgrade-group [all [*upgrade-group-name*]—(Junos Fusion only) Delete the satellite software image association with the specified satellite software upgrade group.

A satellite software upgrade group is a group of satellite devices in the same Junos Fusion that are designated to upgrade to the same satellite software version using the same satellite software package.

version *version-string*—(Junos Fusion only) (Optional) Delete a satellite software package association with a satellite software upgrade group by selecting the satellite software package's version.

Additional Information Before upgrading the software on the router or switch, when you have a known stable system, issue the **request system snapshot** command to back up the software, including the configuration, to the /altroot and /altconfig file systems (on routers) or the /, /altroot, /config, /var, and /var/tmp file systems (on switches). After you have upgraded the software on the router or switch and are satisfied that the new packages are successfully installed and running, issue the **request system snapshot** command again to back up the new software to the /altroot and /altconfig file systems (on routers) or the /, /altroot, /config, /var, and /var/tmp file systems (on switches). After you run the **request system snapshot** command, you cannot return to the previous version of the software, because the running and backup copies of the software are identical.

Required Privilege Level maintenance

Related Documentation

- [request system software add on page 1276](#)
- [request system software rollback on page 1292](#)
- [request system software validate on page 1297](#)
- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

List of Sample Output [request system software delete jdocs on page 1288](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

[request system software delete jdocs](#)

The following example displays the system software packages before and after the **jdocs** package is deleted through the **request system software delete** command:

```
user@host> show system software
Information for jbase:
```

```
Comment:
JUNOS Base OS Software Suite [7.2R1.7]
```

```
Information for jcrypto:
```

```
Comment:
JUNOS Crypto Software Suite [7.2R1.7]
```


Information for jdocs:

Comment:
JUNOS Online Documentation [7.2R1.7]

Information for jkernel:

Comment:
JUNOS Kernel Software Suite [7.2R1.7]

...

```
user@host> request system software delete jdocs
Removing package 'jdocs' ...
```

```
user@host> show system software
Information for jbase:
```

Comment:
JUNOS Base OS Software Suite [7.2R1.7]

Information for jcrypto:

Comment:
JUNOS Crypto Software Suite [7.2R1.7]

Information for jkernel:

Comment:
JUNOS Kernel Software Suite [7.2R1.7]

...

request system software download

Syntax (QFabric System)	request system software download <i>path package-name</i>
Release Information	Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Download a software package from a location on the Director device, mounted external USB flash drive, remote FTP or SCP location, or other location.
Options	<p>path—Location where the software package is located. For example:</p> <ul style="list-style-type: none">• /pbdata/packages/package-name—For a software package that is being installed from a local directory on the switch.• protocol://hostname/pathname/package-name—For a software package or bundle that is to be downloaded and installed from a remote location. Replace protocol with one of the following:<ul style="list-style-type: none">• ftp—File Transfer Protocol. Use ftp://hostname/pathname/package-name. To specify authentication credentials, use ftp://<username>:<password>@hostname/pathname/package-name. To have the system prompt you for the password, specify prompt in place of the password. If a password is required, and you do not specify the password or prompt, an error message is displayed.• scp—Secure copy (available only for Canada and U.S. version). Use scp://hostname/pathname/package-name. To specify authentication credentials, use scp://<username>:<password>@hostname/pathname/package-name.
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none">• request system software add on page 1276• request system software delete on page 1286• request system software rollback on page 1292• request system storage cleanup on page 399• Upgrading Software on page 189• Upgrading Software on a QFabric System
List of Sample Output	request system software download on page 1291
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system software download

```
user@switch> request system software download
ftp://ftp.install-directory/jinstall-qfabric-11.3X30.6.rpm
% Total      % Received % Xferd  Average Speed   Time    Time     Time  Current
           Dload  Upload   Total     Spent    Left     Speed
100 186M 100 186M    0     0 18.4M      0  0:00:10  0:00:10 --:--:-- 18.6M
```

request system software rollback

List of Syntax	Syntax on page 1292 Syntax (EX Series Switches) on page 1292 Syntax (TX Matrix Router) on page 1292 Syntax (TX Matrix Plus Router) on page 1292 Syntax (MX Series Router) on page 1292
Syntax	request system software rollback
Syntax (EX Series Switches)	request system software rollback <all-members> <local> <member <i>member-id</i> > <reboot>
Syntax (TX Matrix Router)	request system software rollback <lcc <i>number</i> scc> <reboot>
Syntax (TX Matrix Plus Router)	request system software rollback <lcc <i>number</i> sfc <i>number</i> > <reboot>
Syntax (MX Series Router)	request system software rollback <all-members> <device-alias <i>alias-name</i> > <local> <member <i>member-id</i> > <reboot> <satellite <i>slot-id</i> > <upgrade-group [all <i>upgrade-group-name</i>]>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Option sfc introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series. Command behavior changed in Junos OS Release 12.1. Option reboot introduced in Junos OS Release 12.3. Options device-alias , satellite , and upgrade-group introduced in Junos OS Release 14.2R3 for Junos Fusion. Option force deprecated for Junos OS with Upgraded FreeBSD in Junos OS Release 15.1.



NOTE: To determine which platforms run Junos OS with Upgraded FreeBSD, see the table listing the platforms currently running Junos OS with upgraded FreeBSD in “[Understanding Junos OS with Upgraded FreeBSD](#)” on page 89.

Description For all versions of Junos OS up to and including Junos OS 11.4, revert to the software that was loaded at the last successful **request system software add** command.

As of Junos OS 12.1 and greater, revert to the last known good state before the most recent **request system software (add | delete)** command. For example, using **rollback** in Junos OS 12.1 after using **request system software add** restores the system to a known good state prior to using the **add** command. Similarly, using **rollback** in Junos OS 12.1 after using **request system software delete** restores the system to a known good state prior to using the **delete** command.

A software rollback fails if any required package (or a **bundle** package containing the required package) cannot be found in `/var/sw/pkg`.

Additional Information

- On a Junos Fusion, the **request system software rollback** command can be used to roll back the version of satellite software associated with a satellite software upgrade group. Rolling back the version of satellite software associated with a satellite software upgrade group triggers a satellite software upgrade.
- On M Series and T Series routers, if **request system software add <jinstall> reboot** was used for the previous installation, then **request system software rollback** has no effect. In this case, use **jinstall** to reinstall the required package.
- On M Series and T Series routers, if **request system software add <sdk1>** was used for the previous installation, then **request system software rollback** removes the last installed SDK package (**sdk1** in this example).
- On SRX Series devices with dual root systems, when **request system software rollback** is run, the system switches to the alternate root. Each root can have a different version of Junos OS. Roll back takes each root back to the previously installed image.
- On QFX3500 and QFX3600 devices in a mixed Virtual Chassis, when the **request system software rollback** command is issued, the system does not rollback to the image stored in the alternate partition.
- On QFX5100 switches, the **reboot** option has been removed. To reboot the switch after a software rollback, issue the **request system reboot** command as a separate, secondary command.

Options **all-members**—(EX4200 switches and MX Series routers only) (Optional) Attempt to roll back to the previous set of packages on all members of the Virtual Chassis configuration.

device-alias *alias-name*—(Junos Fusion only) (Optional) Rollback the satellite software package onto the specified satellite device using the satellite devices FPC slot identifier.

lcc *number*—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, attempt to roll back to the previous set of packages on a T640 router connected to the TX Matrix router. On a TX Matrix Plus router, attempt to roll back to the previous set of packages on a connected router connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(EX4200 switches and MX Series routers only) (Optional) Attempt to roll back to the previous set of packages on the local Virtual Chassis member.

member *member-id*—(EX4200 switches and MX Series routers only) (Optional) Attempt to roll back to the previous set of packages on the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

none—For all versions of Junos OS up to and including Junos OS 11.4, revert to the set of software as of the last successful **request system software add**. As of Junos OS 12.1 and greater, revert to the last known good state before the most recent **request system software (add | delete)** command.

reboot—As of Junos OS 12.3 and greater, automatically reboot upon completing the **request system software rollback** command.

satellite *slot-id*—(Junos Fusion only) (Optional) Roll back the satellite software package onto the specified satellite device using the satellite devices FPC slot identifier.

scc—(TX Matrix routers only) (Optional) Attempt to roll back to the previous set of packages on the TX Matrix router (or switch-card chassis).

sfc *number*—(TX Matrix Plus routers only) (Optional) Attempt to roll back to the previous set of packages on the TX Matrix Plus router. Replace *number* with 0.

upgrade-group [all | *upgrade-group-name*]—(Junos Fusion only) Roll back the satellite software image associated with the specified satellite software upgrade group, or for all satellite software upgrade groups in the Junos Fusion when **all** is entered.

Required Privilege Level

maintenance

Related Documentation

- [request system software abort](#)
- [request system software add on page 1276](#)
- [request system software delete on page 1286](#)
- [request system software validate on page 1297](#)

- [request system configuration rescue delete on page 382](#)
- [request system configuration rescue save on page 383](#)
- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

List of Sample Output [request system software rollback on page 1296](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system software rollback

```
user@host> request system software rollback
Verified SHA1 checksum of ./jbase-7.2R1.7.tgz
Verified SHA1 checksum of ./jdocs-7.2R1.7.tgz
Verified SHA1 checksum of ./jroute-7.2R1.7.tgz
Installing package './jbase-7.2R1.7.tgz' ...
Available space: 35495 require: 7335
Installing package './jdocs-7.2R1.7.tgz' ...
Available space: 35339 require: 3497
Installing package './jroute-7.2R1.7.tgz' ...
Available space: 35238 require: 6976
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Reloading /config/juniper.conf.gz ...
Activating /config/juniper.conf.gz ...
mgd: commit complete
Restarting mgd ...
Restarting aprobed ...
Restarting apsd ...
Restarting cosd ...
Restarting fsad ...
Restarting fud ...
Restarting gcdrd ...
Restarting ilmid ...
Restarting irsd ...
Restarting l2tpd ...
Restarting mib2d ...
Restarting nasd ...
Restarting pppoed ...
Restarting rdd ...
Restarting rmopd ...
Restarting rtspd ...
Restarting sampled ...
Restarting serviced ...
Restarting snmpd ...
Restarting spd ...
Restarting vrrpd ...

WARNING: cli has been replaced by an updated version:
CLI release 7.2R1.7 built by builder on 2005-04-22 02:03:44 UTC
Restart cli using the new version ? [yes,no] (yes) yes

Restarting cli ...
user@host
```


request system software validate

List of Syntax	Syntax on page 1297 Syntax (TX Matrix Router) on page 1297 Syntax (TX Matrix Plus Router) on page 1297 Syntax (MX Series Router) on page 1297
Syntax	<pre>request system software validate <i>package-name</i> <set [<i>package-name package-name</i>]> <upgrade-with-config> <upgrade-with-config-format <i>format</i>></pre>
Syntax (TX Matrix Router)	<pre>request system software validate <i>package-name</i> <lcc <i>number</i> scc> <set [<i>package-name package-name</i>]> <upgrade-with-config> <upgrade-with-config-format <i>format</i>></pre>
Syntax (TX Matrix Plus Router)	<pre>request system software validate <i>package-name</i> <lcc <i>number</i> sfc <i>number</i>> <set [<i>package-name package-name</i>]> <upgrade-with-config> <upgrade-with-config-format <i>format</i>></pre>
Syntax (MX Series Router)	<pre>request system software validate <i>package-name</i> <member <i>member-id</i>> <set [<i>package-name package-name</i>]> <upgrade-with-config> <upgrade-with-config-format <i>format</i>></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>set [<i>package-name package-name</i>] option added in Junos OS Release 12.2 for M Series, MX Series, T Series routers, and Branch SRX Series Services Gateways.</p> <p>upgrade-with-config and upgrade-with-config-format <i>format</i> options added in Junos OS Release 12.3 for M Series routers, MX Series routers, and T Series routers.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Validate candidate software against the current configuration of the router.
Options	<p>lcc <i>number</i>—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, validate the software bundle or package on a specific T640 router (or line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, validate the software bundle or package for a specific router that is connected to the TX Matrix Plus router.</p>

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

member *member-id*—(MX Series routers only) (Optional) Validate the software bundle or package on the specified member of the Virtual Chassis configuration. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

package-name—Name of the software bundle or package to test.

scc—(TX Matrix routers only) (Optional) Validate the software bundle or package for the TX Matrix router (or switch-card chassis).

set [*package-name package-name*]—(M Series, MX Series, T Series routers, and Branch SRX Series Services Gateways only) (Optional) Install multiple software packages or software add-on packages at the same time.

sfc *number*—(TX Matrix Plus routers only) (Optional) Validate the software bundle or package for the TX Matrix Plus router.

upgrade-with-config—(Optional) Install one or more configuration files.

upgrade-with-config-format *format*—(Optional) Specify the configuration file format, **text** or **xml**. The default format is **text**.



NOTE: The **upgrade-with-config** and **upgrade-with-config-format** options are only available locally on the router or switch. In a routing matrix, the configuration is applied only to the local router and is not propagated to other routers.

The options are validated during the validation process and applied to the router or switch during the upgrade process. If the upgrade process is successful, the options are removed from the configuration. If the upgrade process fails, the configuration file is renamed with the **.failed** suffix.

Additional Information By default, when you issue the **request system software validate** command on a TX Matrix master Routing Engine, all the T640 master Routing Engines that are connected to it are validated. If you issue the same command on the TX Matrix backup Routing Engine, all

the T640 backup Routing Engines that are connected to it are upgraded to the same version of software.

Likewise, if you issue the **request system software validate** command on a TX Matrix Plus master Routing Engine, all the T1600 or T4000 master Routing Engines that are connected to it are validated. If you issue the same command on a TX Matrix Plus backup Routing Engine, all the T1600 or T4000 backup Routing Engines that are connected to it are upgraded to the same version of software.

Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none"> • <i>request system software abort</i> • request system software add on page 1276 • request system software delete on page 1286 • request system software rollback on page 1292 • Routing Matrix with a TX Matrix Plus Router Solutions Page
List of Sample Output	request system software validate (Successful Case) on page 1299 request system software validate (Failure Case) on page 1299
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system software validate (Successful Case)

```

user@host> request system software validate /var/sw/pkg/jbundle-5.3I20020124_0520_sjg.tgz
Checking compatibility with configuration
Initializing...
Using /packages/jbase-5.3I20020122_1901_sjg
Using /var/sw/pkg/jbundle-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jbase-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jkernel-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jcrypto-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jpfe-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jdocs-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jroute-5.3I20020124_0520_sjg.tgz
Validating against /config/juniper.conf.gz
mgd: commit complete

WARNING: cli has been replaced by an updated version:
CLI release 5.3I0 built by sjg on 2002-01-24 05:23:53 UTC
Restart cli using the new version ? [yes,no] (yes)

```

request system software validate (Failure Case)

```

user@host> request system software validate 6.3/
Pushing bundle to lcc0-re0
error: Failed to transfer package to lcc0-re0

user@host> request system software validate test

```


```
Pushing bundle to lcc0-re0
Pushing bundle to lcc2-re0

lcc0-re0:
gzip: stdin: not in gzip format
tar: child returned status 1
ERROR: Not a valid package: /var/tmp/test
```

rollback

Syntax	<code>rollback <number rescue></code>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	<p>Return to a previously committed configuration. The software saves the last 50 committed configurations, including the rollback number, date, time, and name of the user who issued the commit configuration command.</p> <p>The currently operational Junos OS configuration is stored in the file juniper.conf, and the last three committed configurations are stored in the files juniper.conf.1, juniper.conf.2, and juniper.conf.3. These four files are located in the directory /config, which is on the router's flash drive. The remaining 46 previous committed configurations, the files juniper.conf.4 through juniper.conf.49, are stored in the directory /var/db/config, which is on the router's hard disk.</p> <p>During rollback, the configuration you specify is loaded from the associated file. Only objects in the rollback configuration that differ from the previously loaded configuration are marked as changed (equivalent to load update).</p>
Options	<p>none (Optional)—Return to the most recently saved configuration.</p> <p>number—(Optional) Configuration to return to. The range of values is from 0 through 49. The most recently saved configuration is number 0, and the oldest saved configuration is number 49. The default is 0.</p> <p>rescue—(Optional) Return to the rescue configuration.</p>
Required Privilege Level	rollback—To roll back to configurations other than the one most recently committed.
Related Documentation	<ul style="list-style-type: none"> • Returning to a Previously Committed Junos OS Configuration on page 1339 • Creating and Returning to a Rescue Configuration on page 1342

show system rollback

Syntax	<code>show system rollback <i>number</i></code> <code><compare <i>number</i>></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display the contents of a previously committed configuration, or the differences between two previously committed configurations.
<div>  NOTE: The <code>show system rollback</code> command is a purely operational mode command and cannot be issued with <code>run</code> from the configuration mode. </div>	
Options	<p><i>number</i>—Number of a configuration to view. The output displays the configuration. The range of values is 0 through 49.</p> <p><code>compare <i>number</i></code>—(Optional) Number of another previously committed (rollback) configuration to compare to rollback <i>number</i>. The output displays the differences between the two configurations. The range of values is 0 through 49.</p>
Required Privilege Level	view
List of Sample Output	show system rollback compare on page 1302

Sample Output

show system rollback compare

```

user@host> show system rollback 3 compare 1
[edit]
+ interfaces {
+   ge-1/1/1 {
+     unit 0 {
+       family inet {
+         filter {
+           input mf_plp;
+         }
+         address 14.1.1.1/30;
+       }
+     }
+   }
+   ge-1/2/1 {
+     unit 0 {
+       family inet {
+         filter {
+           input mf_plp;
+         }
+         address 13.1.1.1/30;

```

```
+      }  
+    }  
+  }  
+  ge-1/3/0 {  
+    unit 0 {  
+      family inet {  
+        filter {  
+          input mf_plp;  
+        }  
+        address 12.1.1.1/30;  
+      }  
+    }  
+  }  
+}
```


Configuration and File Management Feature Guide for QFX10000 Switches

PART 16

Configuration and File Management Overview

- [Understanding Configuration Files and the configure Command on page 1309](#)
- [Understanding How Configuration Files are Committed and Stored on page 1313](#)

CHAPTER 33

Understanding Configuration Files and the configure Command

- [Configuration File Terms on page 1309](#)
- [Understanding Configuration Files on page 1310](#)
- [Forms of the configure Command on page 1310](#)

Configuration File Terms

Table 85 lists the various configuration file terms and their definitions.

Table 85: Configuration File Terms

Term	Definition
active configuration	Current committed configuration of a switch.
candidate configuration	Working copy of the configuration that allows users to make configurational changes without causing any operational changes until this copy is committed.
configuration group	Group of configuration statements that can be inherited by the rest of the configuration.
commit a configuration	Check configuration for proper syntax, activate and mark as the current configuration file running on the switching platform.
configuration hierarchy	Junos OS configuration consists of a hierarchy of statements. There are two types of statements: container statements, which contain other statements, and leaf statements, which do not contain other statements. All the container and leaf statements together form the configuration hierarchy.
default configuration	Default configuration contains the initial values set for each configuration parameter when a switch is shipped.
rescue configuration	Well-known configuration that recovers a switch from a configuration that denies management access. You set a current committed configuration to be the rescue configuration through the CLI.
roll back a configuration	Return to a previously committed configuration.

- Related Documentation**
- [Loading a Previous Configuration File on page 1330](#)
 - [Reverting to the Rescue Configuration on page 1346](#)
 - [Understanding Configuration Files on page 1310](#)

Understanding Configuration Files

A configuration file stores the complete configuration of a switch. The current configuration of a switch is called the active configuration. You can alter this current configuration and you can also return to a previous configuration or to a rescue configuration.

Juniper Networks Junos OS saves the 50 most recently committed configuration files on a switch so that you can return to a previous configuration. The configuration files are named:

- **juniper.conf.gz**—The current active configuration.
- **juniper.conf.1.gz** to **juniper.conf.49.gz**—Rollback configurations.

To make changes to the configuration file, you have to work in configuration mode in the CLI. When making changes to a configuration file, you are viewing and changing the candidate configuration file. The candidate configuration allows you to make configuration changes without causing operational changes to the active configuration or causing potential damage to your current network operations. Once you commit the changes made to the candidate configuration, the system updates the active configuration.

- Related Documentation**
- [Uploading a Configuration File on page 1332](#)
 - [Loading a Previous Configuration File on page 1330](#)
 - [Reverting to the Rescue Configuration on page 1346](#)
 - [Configuration File Terms on page 1309](#)

Forms of the configure Command

The Junos OS supports three forms of the **configure** command: **configure**, **configure private**, and **configure exclusive**. These forms control how users edit and commit configurations and can be useful when multiple users configure the software. See [Table 86](#).

Table 86: Forms of the configure Command

Command	Edit Access	Commit Access
configure	<ul style="list-style-type: none"> No one can lock the configuration. All users can make configuration changes. <p>When you enter configuration mode, the CLI displays the following information:</p> <ul style="list-style-type: none"> A list of other users editing the configuration. Hierarchy levels the users are viewing or editing. Whether the configuration has been changed, but not committed. When multiple users enter conflicting configurations, the most recent change to be entered takes precedence. 	<ul style="list-style-type: none"> No one can lock the configuration. All users can commit all changes to the configuration. If you and another user make changes and the other user commits changes, your changes are committed as well.
configure exclusive	<ul style="list-style-type: none"> One user locks the configuration and makes changes without interference from other users. Other users can enter and exit configuration mode, but they cannot commit the configuration. If you enter configuration mode while another user has locked the configuration (with the configure exclusive command), the CLI displays the user and the hierarchy level the user is viewing or editing. If you enter configuration mode while another user has locked the configuration, you can forcibly log out that user with the request system logout operational mode command. For details, see the CLI Explorer. 	
configure private	<ul style="list-style-type: none"> Multiple users can edit the configuration at the same time. Each user has a private candidate configuration to edit independently of other users. When multiple users enter conflicting configurations, the first commit operation takes precedence over subsequent commit operations. 	<ul style="list-style-type: none"> When you commit the configuration, the router verifies that the operational (running) configuration has not been modified by another user before accepting your private candidate configuration as the new operational configuration. If the configuration has been modified by another user, you can merge the modifications into your private candidate configuration and attempt to commit again.

**Related
Documentation**

- *Committing a Junos OS Configuration*
- *Using the configure Command*
- *Displaying Users Currently Editing the Junos OS Configuration*
- *Using the configure exclusive Command*
- *Updating the configure private Configuration*
- *Displaying set Commands from the Junos OS Configuration*

Understanding How Configuration Files are Committed and Stored

- [Junos OS Commit Model for Router or Switch Configuration on page 1313](#)
- [Understanding How the Junos OS Configuration Is Stored on page 1314](#)

Junos OS Commit Model for Router or Switch Configuration

The router or switch configuration is saved using a commit model—a candidate configuration is modified as desired and then committed to the system. When a configuration is committed, the router or switch checks the configuration for syntax errors, and if no errors are found, the configuration is saved as **juniper.conf.gz** and activated. The formerly active configuration file is saved as the first rollback configuration file (**juniper.conf.1.gz**), and any other rollback configuration files are incremented by 1. For example, **juniper.conf.1.gz** is incremented to **juniper.conf.2.gz**, making it the second rollback configuration file. The router or switch can have a maximum of 49 rollback configurations (numbered 1 through 49) saved on the system.

On the router or switch, the active configuration file and the first three rollback files (**juniper.conf.gz.1**, **juniper.conf.gz.2**, **juniper.conf.gz.3**) are located in the **/config** directory. If the file **rescue.conf.gz** is saved on the system, this file should also be saved in the **/config** directory. The factory default files are located in the **/etc/config** directory.

There are two mechanisms used to propagate the configurations between Routing Engines within a router or switch:

- Synchronization—Propagates a configuration from one Routing Engine to a second Routing Engine within the same router or switch chassis.



NOTE: The QFX3500 switch has only one Routing Engine.

To synchronize configurations, use the **commit synchronize** CLI command. If one of the Routing Engines is locked, the synchronization fails. If synchronization fails because of a locked configuration file, you can use the **commit synchronize force** command. This command overrides the lock and synchronizes the configuration files.

- Distribution—Propagates a configuration across the routing plane on a multichassis router or switch. Distribution occurs automatically. There is no user command available

to control the distribution process. If a configuration is locked during a distribution of a configuration, the locked configuration does not receive the distributed configuration file, so the synchronization fails. You need to clear the lock before the configuration and resynchronize the routing planes.



NOTE: When you use the `commit synchronize force` CLI command on a multichassis platform, the forced synchronization of the configuration files does not affect the distribution of the configuration file across the routing plane. If a configuration file is locked on a router or switch remote from the router or switch where the command was issued, the synchronization fails on the remote router or switch. You need to clear the lock and reissue the `synchronize` command.

- Related Documentation**
- *Configuring Junos OS for the First Time on a Router or Switch with a Single Routing Engine*
 - [commit on page 329](#)

Understanding How the Junos OS Configuration Is Stored

When you edit a configuration, you work in a copy of the current configuration to create a candidate configuration. The changes you make to the candidate configuration are visible in the CLI immediately, so if multiple users are editing the configuration at the same time, all users can see all changes.

To have a candidate configuration take effect, you *commit* the changes. At this point, the candidate file is checked for proper syntax, activated, and marked as the current, operational software configuration file. If multiple users are editing the configuration, when you commit the candidate configuration, all changes made by all the users take effect.

In addition to saving the current configuration, the CLI saves the current operational version and the previous 49 versions of committed configurations. The most recently committed configuration is version 0, which is the current operational version and the default configuration that the system returns to if you roll back to a previous configuration. The oldest saved configuration is version 49.

By default, Junos OS saves the current configuration and three previous versions of the committed configuration on the CompactFlash card. The currently operational Junos OS configuration is stored in the file `juniper.conf.gz`, and the last three committed configurations are stored in the files `juniper.conf.1.gz`, `juniper.conf.2.gz`, and `conf.3.gz`. These four files are located in the router or switch's CompactFlash card in the directory `/config`.

The remaining 46 previous versions of committed configurations, the files `juniper.conf.4` through `juniper.conf.49`, are stored in the directory `/var/db/config` on the hard disk.

- Related Documentation**
- *Using Junos OS to Specify the Number of Configurations Stored on the CompactFlash Card*

- [Returning to the Most Recently Committed Junos OS Configuration on page 1338](#)
- [Returning to a Previously Committed Junos OS Configuration on page 1339](#)
- [Loading a Configuration from a File on page 1325](#)

PART 17

Managing Configuration Files

- [Using the Default Configuration on page 1319](#)
- [Loading, Uploading, Transferring, and Saving Configuration Files on page 1325](#)
- [Reverting a Configuration or Rolling Back to the Factory Default on page 1337](#)
- [Using Rescue Configurations on page 1345](#)
- [Comparing or Compressing Configuration Files on page 1349](#)
- [Troubleshooting Procedures on page 1353](#)

Using the Default Configuration

- QFX3500 Switch Default Configuration on page 1319

QFX3500 Switch Default Configuration

Each QFX Series product is programmed with a factory default configuration that contains the values set for each configuration parameter when a switch is shipped. The default configuration file sets values for system parameters such as **syslog** and **commit**, configures storm control and Ethernet switching on all interfaces, and enables IGMP snooping, RSTP, and LLDP protocols.

When you commit changes to the configuration, a new configuration file is created, which becomes the active configuration. You can always revert to the factory default configuration if you need to.

The following factory default configuration file is for a QFX3500 switch with 48 ports:



NOTE: In this example, xe-0/0/0 through xe-0/0/47 are the network interface ports.

```
protocols {
  igmp-snooping {
    vlan all;
  }
  rstp;
  lldp {
    interface all;
  }
}
interfaces {
  xe-0/0/0 {
    unit 0 {
      family ethernet-switching;
    }
  }
  xe-0/0/1 {
    unit 0 {
      family ethernet-switching;
    }
  }
}
```

```
}
xe-0/0/2 {
  unit 0 {
    family ethernet-switching;
  }
}
xe-0/0/3 {
  unit 0 {
    family ethernet-switching;
  }
}
xe-0/0/4 {
  unit 0 {
    family ethernet-switching;
  }
}
xe-0/0/5 {
  unit 0 {
    family ethernet-switching;
  }
}
xe-0/0/6 {
  unit 0 {
    family ethernet-switching;
  }
}
xe-0/0/7 {
  unit 0 {
    family ethernet-switching;
  }
}
xe-0/0/8 {
  unit 0 {
    family ethernet-switching;
  }
}
xe-0/0/9 {
  unit 0 {
    family ethernet-switching;
  }
}
xe-0/0/10 {
  unit 0 {
    family ethernet-switching;
  }
}
xe-0/0/11 {
  unit 0 {
    family ethernet-switching;
  }
}
xe-0/0/12 {
  unit 0 {
    family ethernet-switching;
  }
}
```



```
xe-0/0/13 {  
  unit 0 {  
    family ethernet-switching;  
  }  
}  
xe-0/0/14 {  
  unit 0 {  
    family ethernet-switching;  
  }  
}  
xe-0/0/15 {  
  unit 0 {  
    family ethernet-switching;  
  }  
}  
xe-0/0/16 {  
  unit 0 {  
    family ethernet-switching;  
  }  
}  
xe-0/0/17 {  
  unit 0 {  
    family ethernet-switching;  
  }  
}  
xe-0/0/18 {  
  unit 0 {  
    family ethernet-switching;  
  }  
}  
xe-0/0/19 {  
  unit 0 {  
    family ethernet-switching;  
  }  
}  
xe-0/0/20 {  
  unit 0 {  
    family ethernet-switching;  
  }  
}  
xe-0/0/21 {  
  unit 0 {  
    family ethernet-switching;  
  }  
}  
xe-0/0/22 {  
  unit 0 {  
    family ethernet-switching;  
  }  
}  
xe-0/0/23 {  
  unit 0 {  
    family ethernet-switching;  
  }  
}  
xe-0/0/24 {
```

```
    unit 0 {  
        family ethernet-switching;  
    }  
}  
xe-0/0/25 {  
    unit 0 {  
        family ethernet-switching;  
    }  
}  
xe-0/0/26 {  
    unit 0 {  
        family ethernet-switching;  
    }  
}  
xe-0/0/27 {  
    unit 0 {  
        family ethernet-switching;  
    }  
}  
xe-0/0/28 {  
    unit 0 {  
        family ethernet-switching;  
    }  
}  
xe-0/0/29 {  
    unit 0 {  
        family ethernet-switching;  
    }  
}  
xe-0/0/30 {  
    unit 0 {  
        family ethernet-switching;  
    }  
}  
xe-0/0/31 {  
    unit 0 {  
        family ethernet-switching;  
    }  
}  
xe-0/0/32 {  
    unit 0 {  
        family ethernet-switching;  
    }  
}  
xe-0/0/33 {  
    unit 0 {  
        family ethernet-switching;  
    }  
}  
xe-0/0/34 {  
    unit 0 {  
        family ethernet-switching;  
    }  
}  
xe-0/0/35 {  
    unit 0 {
```

```
        family ethernet-switching;
    }
}
xe-0/0/36 {
    unit 0 {
        family ethernet-switching;
    }
}
xe-0/0/37 {
    unit 0 {
        family ethernet-switching;
    }
}
xe-0/0/38 {
    unit 0 {
        family ethernet-switching;
    }
}
xe-0/0/39 {
    unit 0 {
        family ethernet-switching;
    }
}
xe-0/0/40 {
    unit 0 {
        family ethernet-switching;
    }
}
xe-0/0/41 {
    unit 0 {
        family ethernet-switching;
    }
}
xe-0/0/42 {
    unit 0 {
        family ethernet-switching;
    }
}
xe-0/0/43 {
    unit 0 {
        family ethernet-switching;
    }
}
xe-0/0/44 {
    unit 0 {
        family ethernet-switching;
    }
}
xe-0/0/45 {
    unit 0 {
        family ethernet-switching;
    }
}
xe-0/0/46 {
    unit 0 {
        family ethernet-switching;
```

```
    }  
  }  
  xe-0/0/47 {  
    unit 0 {  
      family ethernet-switching;  
    }  
  }  
}  
ethernet-switching-options {  
  storm-control {  
    interface all;  
  }  
}  
system {  
  syslog {  
    archive size 256k;  
    file default-log-messages {  
      structured-data;  
    }  
    user * {  
      any emergency;  
    }  
    file messages {  
      any notice;  
      authorization info;  
    }  
    file interactive-commands {  
      interactive-commands any;  
    }  
  }  
}  
ports {  
  console type vt100;  
}  
compress-configuration-files;  
login {  
  password {  
    minimum-length 6;  
    minimum-changes 1;  
    change-type set transitions;  
    format md5;  
  }  
}  
commit {  
  factory-settings {  
    reset-chassis-lcd-menu;  
  }  
}  
}
```

Related Documentation

- [Reverting to the Default Factory Configuration on page 1344](#)
- [Configuring a QFX3500 Device as a Standalone Switch](#)
- [Understanding Configuration Files on page 1310](#)
- [Interfaces Overview on page 2785](#)

CHAPTER 36

Loading, Uploading, Transferring, and Saving Configuration Files

- Loading a Configuration from a File on page 1325
- Examples: Loading a Configuration from a File on page 1328
- Loading a Previous Configuration File on page 1330
- Saving a Configuration to a File on page 1331
- Uploading a Configuration File on page 1332
- Using Junos OS to Configure a Router or Switch to Transfer Its Configuration to an Archive Site on page 1333

Loading a Configuration from a File

You can create a file, copy the file to the local router, and then load the file into the CLI. After you have loaded the file, you can commit it to activate the configuration on the router, or you can edit the configuration interactively using the CLI and commit it at a later time.

You can also create a configuration while typing at the terminal and then load it. Loading a configuration from the terminal is generally useful when you are cutting existing portions of the configuration and pasting them elsewhere in the configuration.

To load an existing configuration file that is located on the router, use the **load** configuration mode command:

```
[edit]
user@host# load (factory-default | merge | override | patch | replace | set | update)
          filename <relative>
```

For information about specifying the filename, see *Viewing Files and Directories on a Device Running Junos OS*.

To load a configuration from the terminal, use the following version of the **load** configuration mode command. Press Ctrl-d to end the input.

```
[edit]
user@host# load (factory-default | merge | override | patch | replace | set | update)
terminal <relative>
```

To replace an entire configuration, specify the **override** option at any level of the hierarchy. A **load override** operation completely replaces the current candidate configuration with the file you are loading. Thus, if you saved a complete configuration, use this option.

An **override** operation discards the current candidate configuration and loads the configuration in *filename* or the configuration that you type at the terminal. When you use the **override** option and commit the configuration, all system processes reparse the configuration. For an example, see [Figure 7](#).

To replace portions of a configuration, specify the **replace** option. The **load replace** operation looks for **replace:** tags that you added to the loaded file, and replaces the parts of the candidate configuration with whatever is specified after the tag. This is useful when you want more control over exactly what is being changed. For this operation to work, you must include **replace:** tags in the file or configuration you type at the terminal. The software searches for the **replace:** tags, deletes the existing statements of the same name, if any, and replaces them with the incoming configuration. If there is no existing statement of the same name, the **replace** operation adds to the configuration the statements marked with the **replace:** tag. For an example, see [Figure 8](#).

If, in an **override** or **merge** operation, you specify a file or type text that contains **replace:** tags, the **replace:** tags are ignored and the **override** or **merge** operation is performed.

If you are performing a **replace** operation and the file you specify or text you type does not contain any **replace:** tags, the **replace** operation is effectively equivalent to a **merge** operation. This might be useful if you are running automated scripts and cannot know in advance whether the scripts need to perform a **replace** or a **merge** operation. The scripts can use the **replace** operation to cover either case.

The **load merge** operation adds the saved file to the existing candidate configuration. This is useful if you are adding new configuration sections. For example, suppose that you are adding a BGP configuration to the **[edit protocols]** hierarchy level, where there was no BGP configuration before, you can use the **load merge** operation to combine the saved file configuration to the existing candidate configuration. If the existing configuration and the incoming configuration contain conflicting statements, the statements in the incoming configuration override those in the existing configuration.

To replace only the configuration that has changed, specify the **update** option at any level of the hierarchy. The **load update** operation compares the candidate configuration and the file you are loading, and only changes the parts of the candidate configuration that are different from the new configuration. You would use this, for example, if there is an existing BGP configuration and the file you are loading changes it in some way.

To change part of the configuration with a patch file, specify the **patch** option. The **load patch** operation loads a file or terminal input that contains configuration changes. First, on a device that already has the configuration changes, you type the **show | compare** command to output the differences between two configurations. Then you can load the differences on another router. The advantage of the **load patch** command is that it saves you from having to copy snippets from different hierarchy levels into a text file prior to loading them into the target device. This might be a useful time saver if you are configuring several devices with the same options. For example, suppose that you configure a routing

policy on router1 and you want to replicate the policy configuration on router2, router3, and router4. You can use the **load patch** operation.

First, run the **show | compare** command.

```
user@router1# show | compare rollback 3
[edit protocols ospf]
+ export default-static;
- export static-default
[edit policy-options]
+ policy-statement default-static {
+   from protocol static;
+   then accept;
+ }
```

Copy the output of the **show | compare** command to the clipboard, making sure to include the hierarchy levels. On router2, router3, and router4, type **load patch terminal** and paste the output. Press Enter and then press Ctrl-d to end the operation. If the patch input specifies different values for an existing statement, the patch input overrides the existing statement.

To use the **merge**, **replace**, **set**, or **update** option without specifying the full hierarchy level, specify the **relative** option. For example:

```
[edit system]
user@host# show static-host-mapping
bob sysid 987.654.321ab
[edit system]
user@host# load replace terminal relative
[Type ^D at a new line to end input]
replace: static-host-mapping {
  bob sysid 0123.456.789bc;
}
load complete
[edit system]
user@host# show static-host-mapping
bob sysid 0123.456.789bc;
```

To load a configuration that contains the **set** configuration mode command, specify the **set** option. This option executes the configuration instructions line by line as they are stored in a file or from a terminal. The instructions can contain any configuration mode command, such as **set**, **edit**, **exit**, and **top**. For an example, see [Figure 11](#).

To copy a configuration file from another network system to the local router, you can use the SSH and Telnet utilities, as described in the [CLI Explorer](#).



NOTE: If you are using Junos OS in a Common Criteria environment, system log messages are created whenever a secret attribute is changed (for example, password changes or changes to the RADIUS shared secret). These changes are logged during the following configuration load operations:

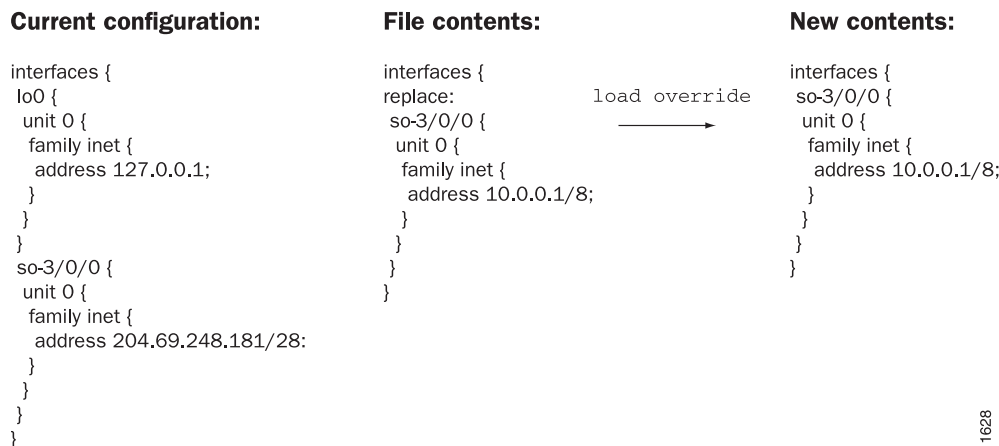
```
load merge
load replace
load override
load update
```

For more information, see the *Secure Configuration Guide for Common Criteria and Junos-FIPS*.

Related Documentation • [Examples: Loading a Configuration from a File on page 1328](#)

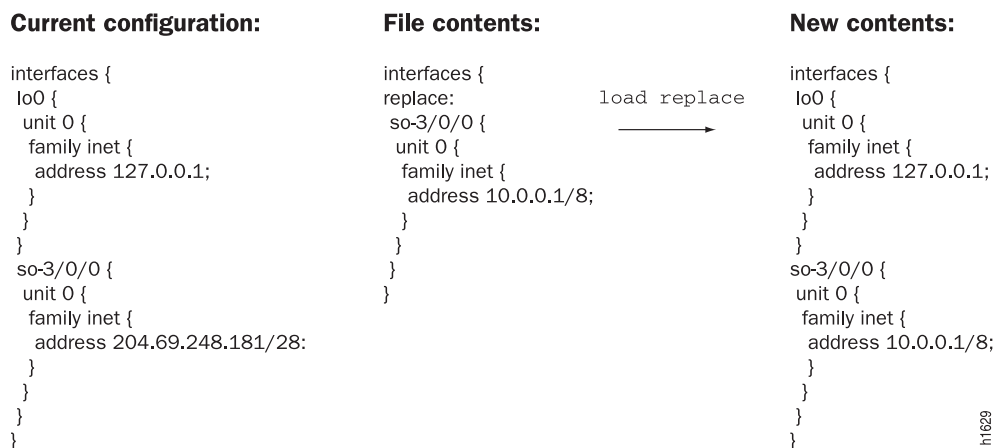
Examples: Loading a Configuration from a File

Figure 7: Overriding the Current Configuration



1628

Figure 8: Using the replace Option



h1629

Figure 9: Using the merge Option




Figure 10: Using a Patch File

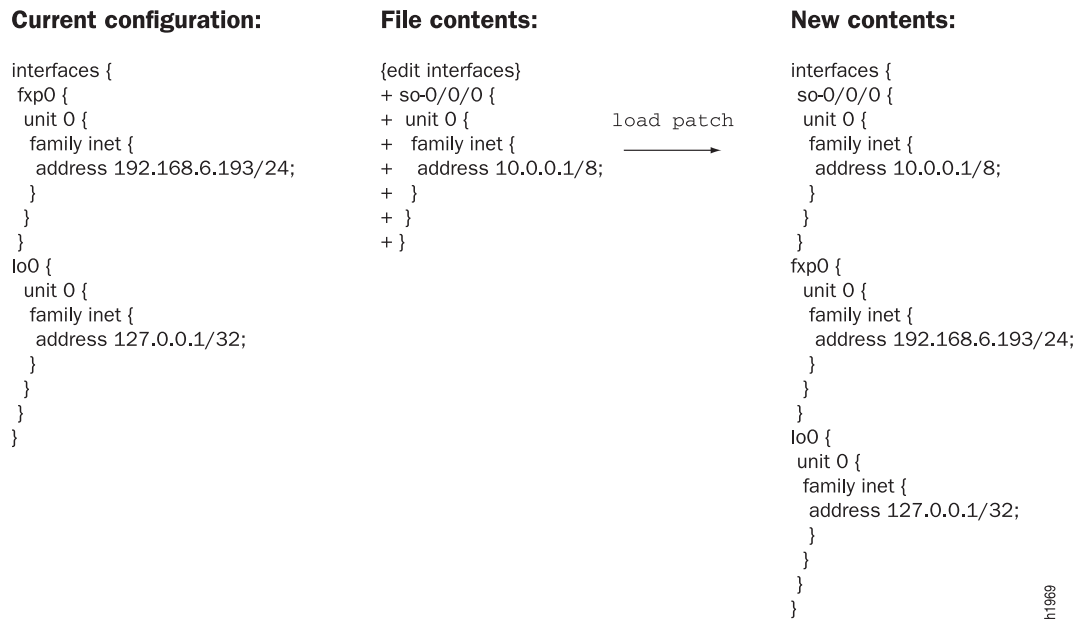


Figure 11: Using the set Option**File contents:**

```
edit access
set profile p1 client cl ike
edit profile p1 client cl ike
set pre-shared-key ascii-text "abcd"
set allowed-proxy-pair local 1.1.1.1 remote 2.2.2.2
exit
deactivate profile p1
top
edit system
set radius-server 1.1.1.1
```

```
load set
```

**New contents:**

```
system {
  radius-server {
    1.1.1.1;
  }
}
access {
  inactive: profile p1 {
    client cl {
      ike {
        allowed-proxy-pair local 1.1.1.1/32 remote 2.2.2.2/32;
        pre-shared-key ascii-text "$9$Ydg4ZDjqf5FVw"; ## SECRET-DATA
      }
    }
  }
}
```

g017215

Related Documentation

- [Loading a Configuration from a File on page 1325](#)

Loading a Previous Configuration File

You can use the **rollback <number>** command to return to a previously committed configuration file. A switch saves the last 50 committed configurations, including the rollback number, date, time, and name of the user who issued the **commit** configuration command.

Syntax

```
rollback <number>
```

Options

- **none**—Return to the most recently saved configuration.
- **number**—Return to the specified configuration.
 - **Range:** 0 through 49. The most recently saved configuration is number 0, and the oldest saved configuration is number 49.
 - **Default:** 0

To return to a configuration prior to the most recently committed one:

1. Specify the rollback number (here, 1 is entered and the configuration returns to the previously committed configuration 0):

```
[edit]
user@switch# rollback 1
load complete
```

2. Activate the configuration you have loaded:

```
[edit]
user@switch# commit
```

Related Documentation • [Configuration File Terms on page 1309](#)

Saving a Configuration to a File

Save Junos OS configuration to a file so that you can edit it with a text editor of your choice. You can save your current configuration to an ASCII file, which saves the configuration in its current form, including any uncommitted changes. If more than one user is modifying the configuration, all changes made by all users are saved.

To save software configuration changes to an ASCII file, use the **save** configuration mode command:

```
[edit]
user@host# save filename
[edit]
user@host#
```

The contents of the current level of the statement hierarchy (and below) are saved, along with the statement hierarchy containing it. This allows a section of the configuration to be saved, while fully specifying the statement hierarchy.

By default, the configuration is saved to a file in your home directory, which is on the flash drive.

When you issue this command from anywhere in the hierarchy (except the top level), a **replace** tag is automatically included at the beginning of the file. You can use the **replace** tag to control how a configuration is loaded from a file.

```
user@host> file show /var/home/user/myconf
replace:
protocols {
  bgp {
    disable;
    group int {
      type internal;
    }
  }
  isis {
    disable;
    interface all {
      level 1 disable;
    }
  }
}
```

```

        interface fxp0.0 {
            disable;
        }
    }
    ospf {
        traffic-engineering;
        reference-bandwidth 4g;
        ...
    }
}

```

Uploading a Configuration File

You can create a configuration file on your local system, copy the file to the switch, and then load the file into the CLI. After you have loaded the configuration file, you can commit it to activate the configuration on the switch. You can also edit the configuration interactively using the CLI and commit it at a later time.

To upload a configuration file from your local system:

1. Create the configuration file using a text editor such as Notepad, making sure that the syntax of the configuration file is correct. For more information about testing the syntax of a configuration file, see the *Junos OS System Basics and Services Command Reference* at <http://www.juniper.net/techpubs/software/junos/index.html>.
2. In the configuration text file, use an option to perform the required action when the file is loaded. [Table 87](#) lists and describes some options for the **load** command.

Table 87: Options for the load Command

Options	Description
merge	Combines the current active configuration and the configuration in the filename you specify or the one that you type at the terminal. A merge operation is useful when you are adding a new section to an existing configuration. If the active configuration and the incoming configuration contain conflicting statements, the statements in the incoming configuration override those in the active configuration.
override	Discards the current candidate configuration and loads the configuration in the filename you specify or the one that you type at the terminal. When you use the override option and commit the configuration, all system processes reparse the configuration. You can use the override option at any level of the hierarchy.
replace	Searches for the replace tags, deletes the existing statements of the same name, if any, and replaces them with the incoming configuration. If there is no existing statement of the same name, the replace operation adds the statements marked with the replace tag to the active configuration. NOTE: For this operation to work, you must include replace tags in the text file or in the configuration you type at the terminal.

3. Press Ctrl+a to select all the text in the configuration file.
4. Press Ctrl+c to copy the contents of the configuration text file to the Clipboard.
5. Log in to the switch using your username and password.

6. Enter configuration mode:
user@switch> **configure**

[edit]
user@switch#
7. Load the configuration file:
[edit]
user@switch# **load merge terminal**
8. At the cursor, paste the contents of the Clipboard using the mouse and the Paste icon:
[edit]
user@switch# **load merge terminal**
[Type ^D at a new line to end input]
>Cursor is here. Paste the contents of the clipboard here<
9. Press Enter.
10. Press Ctrl+d to set the end-of-file marker.

To view results of the configuration steps before committing the configuration, type the **show** command at the user prompt.

To commit these changes to the active configuration, type the **commit** command at the user prompt. You can also edit the configuration interactively using the CLI and commit it at a later time.

Related Documentation

- [Understanding Configuration Files on page 1310](#)

Using Junos OS to Configure a Router or Switch to Transfer Its Configuration to an Archive Site

You can configure a router or switch to transfer its configuration to an archive file periodically. The following tasks describe how to transfer the configuration to an archive site:

1. [Configuring the Transfer of the Currently Active Configuration to an Archive Site on page 1333](#)
2. [Configuring the Periodic Transfer of the Active Configuration to an Archive Site on page 1334](#)
3. [Configuring the Transfer of the Currently Active Configuration When a Configuration Is Committed on page 1334](#)
4. [Configuring Archive Sites for the Transfer of Active Configuration Files on page 1334](#)

Configuring the Transfer of the Currently Active Configuration to an Archive Site

If you want to back up your device's current configuration to an archive site, you can configure the router or switch to transfer its currently active configuration by FTP or secure copy (SCP) periodically or after each commit.

To configure the router or switch to transfer its currently active configuration to an archive site, include statements at the **[edit system archival configuration]** hierarchy level:

```
[edit system archival configuration]
archive-sites {
  ftp://username<:password>@host-address<:port>/url-path;
  scp://username<:password>@host-address<:port>/url-path;
}
transfer-interval interval;
transfer-on-commit;
```



NOTE: When specifying a URL in a Junos OS statement using an IPv6 host address, you must enclose the entire URL in quotation marks ("") and enclose the IPv6 host address in brackets ([]). For example,
"ftp://username<:password>@[ipv6-host-address]<:port>/url-path"

Configuring the Periodic Transfer of the Active Configuration to an Archive Site

To configure the router or switch to periodically transfer its currently active configuration to an archive site, include the **transfer-interval** statement at the **[edit system archival configuration]** hierarchy level:

```
[edit system archival configuration]
transfer-interval interval;
```

The **interval** is a period of time ranging from 15 through 2880 minutes.

Configuring the Transfer of the Currently Active Configuration When a Configuration Is Committed

To configure the router or switch to transfer its currently active configuration to an archive site each time you commit a candidate configuration, include the **transfer-on-commit** statement at the **[edit system archival configuration]** hierarchy level:

```
[edit system archival configuration]
transfer-on-commit;
```



NOTE: When specifying a URL in a Junos OS statement using an IPv6 host address, you must enclose the entire URL in quotation marks ("") and enclose the IPv6 host address in brackets ([]). For example,
"scp://username<:password>@[ipv6-host-address]<:port>/url-path"

Configuring Archive Sites for the Transfer of Active Configuration Files

When you configure the router or switch to transfer its configuration files, you specify an archive site to which the files are transferred. If you specify more than one archive site, the router or switch attempts to transfer files to the first archive site in the list, moving to the next site only if the transfer fails.

When you use the **archive-sites** statement, you can specify a destination as an FTP URL, or SCP-style remote file specification. The URL type **file://** is also supported.

To configure the archive site, include the **archive-sites** statement at the **[edit system archival configuration]** hierarchy level:

```
[edit system archival configuration]
archive-sites {
  ftp://username@host:<port>url-path password password;
  scp://username@host:<port>url-path password password;
  file://<path>/<filename>;
}
```



NOTE: When specifying a URL in a Junos OS statement using an IPv6 host address, you must enclose the entire URL in quotation marks (") and enclose the IPv6 host address in brackets ([]). For example,
 "scp://username<:password>@[ipv6-host-address]<:port>/url-path"

When you specify the archive site, do not add a forward slash (/) to the end of the URL.

The destination filename is saved in the following format, where *n* corresponds to the number of the compressed configuration rollback file that has been archived:

```
<router-name>_YYYYMMDD_HHMMSS_juniper.conf.n.gz
```



NOTE: The time included in the destination filename is always in Coordinated Universal Time (UTC) regardless of whether the time on the router is configured as UTC or the local time zone. The default time zone on the router or switch is UTC.

Reverting a Configuration or Rolling Back to the Factory Default

- [Rolling Back Junos OS Configuration Changes on page 1337](#)
- [Returning to the Most Recently Committed Junos OS Configuration on page 1338](#)
- [Returning to a Previously Committed Junos OS Configuration on page 1339](#)
- [Reverting to the Default Factory Configuration on page 1344](#)

Rolling Back Junos OS Configuration Changes

This topic shows how to use the **rollback** command to return to the most recently committed Junos OS configuration. The **rollback** command is useful if you make configuration changes and then decide not to keep the changes.

The following procedure shows how to configure an SNMP health monitor on a device running Junos OS and then return to the most recently committed configuration that does not include the health monitor. When configured, the SNMP health monitor provides the network management system (NMS) with predefined monitoring for file system usage, CPU usage, and memory usage on the device.

1. Enter configuration mode:

```
user@host> configure
entering configuration mode
[edit]
user@host#
```

2. Show the current configuration (if any) for SNMP:

```
[edit]
user@host# show snmp
```

No **snmp** statements appear because SNMP has not been configured on the device.

3. Configure the health monitor:

```
[edit]
user@host# set snmp health-monitor
```

4. Show the new configuration:

```
[edit]
user@host# show snmp
```

```
health-monitor;
```

The **health-monitor** statement indicates that SNMP health monitoring is configured on the device.

5. Enter the **rollback** configuration mode command to return to the most recently committed configuration:

```
[edit]
user@host# rollback
load complete
```

6. Show the configuration again to make sure your change is no longer present:

```
[edit]
user@host# show snmp
```

No **snmp** configuration statements appear. The health monitor is no longer configured.

7. Enter the **commit** command to activate the configuration to which you rolled back:

```
[edit]
user@host# commit
```

8. Exit configuration mode:

```
[edit]
user@host# exit
Exiting configuration mode
```

You can also use the **rollback** command to return to earlier configurations.

**Related
Documentation**

- [Returning to the Most Recently Committed Junos OS Configuration on page 1338](#)

Returning to the Most Recently Committed Junos OS Configuration

To return to the most recently committed configuration and load it into configuration mode without activating it, use the **rollback** configuration mode command:

```
[edit]
user@host# rollback

load complete
```

To activate the configuration to which you rolled back, use the **commit** command:

```
[edit]
user@host# rollback
load complete
[edit]
user@host# commit
```

**Related
Documentation**

- [Rolling Back Junos OS Configuration Changes on page 1337](#)
- [Returning to a Previously Committed Junos OS Configuration on page 1339](#)
- [Understanding How the Junos OS Configuration Is Stored on page 1314](#)

Returning to a Previously Committed Junos OS Configuration

This topic explains how you can return to a configuration prior to the most recently committed one, and contains the following sections:

- [Returning to a Configuration Prior to the One Most Recently Committed on page 1339](#)
- [Displaying Previous Configurations on page 1339](#)
- [Comparing Configuration Changes with a Prior Version on page 1340](#)
- [Creating and Returning to a Rescue Configuration on page 1342](#)
- [Saving a Configuration to a File on page 1343](#)

Returning to a Configuration Prior to the One Most Recently Committed

To return to a configuration prior to the most recently committed one, include the configuration number, 0 through 49, in the **rollback** command. The most recently saved configuration is number 0 (which is the default configuration to which the system returns), and the oldest saved configuration is number 49.

```
[edit]
user@host# rollback number
load complete
```

Displaying Previous Configurations

To display previous configurations, including the rollback number, date, time, the name of the user who committed changes, and the method of commit, use the **rollback ?** command.

```
[edit]
user@host# rollback ?
Possible completions:
<[Enter]> Execute this command
<number> Numeric argument
0      2005-02-27 12:52:10 PST by abc via cli
1      2005-02-26 14:47:42 PST by def via cli
2      2005-02-14 21:55:45 PST by ghi via cli
3      2005-02-10 16:11:30 PST by jkl via cli
4      2005-02-10 16:02:35 PST by mno via cli
5      2005-03-16 15:10:41 PST by pqr via cli
6      2005-03-16 14:54:21 PST by stu via cli
7      2005-03-16 14:51:38 PST by vwx via cli
8      2005-03-16 14:43:29 PST by yzz via cli
9      2005-03-16 14:15:37 PST by abc via cli
10     2005-03-16 14:13:57 PST by def via cli
11     2005-03-16 12:57:19 PST by root via other
12     2005-03-16 10:45:23 PST by root via other
13     2005-03-16 10:08:13 PST by root via other
14     2005-03-16 01:20:56 PST by root via other
15     2005-03-16 00:40:37 PST by ghi via cli
16     2005-03-16 00:39:29 PST by jkl via cli
17     2005-03-16 00:32:36 PST by mno via cli
18     2005-03-16 00:31:17 PST by pqr via cli
```

```
19      2005-03-15 19:59:00 PST by stu via cli
20      2005-03-15 19:53:39 PST by vwx via cli
21      2005-03-15 18:07:19 PST by yzz via cli
22      2005-03-15 17:59:03 PST by abc via cli
23      2005-03-15 15:05:14 PST by def via cli
24      2005-03-15 15:04:51 PST by ghi via cli
25      2005-03-15 15:03:42 PST by jkl via cli
26      2005-03-15 15:01:52 PST by mno via cli
27      2005-03-15 14:58:34 PST by pqr via cli
28      2005-03-15 13:09:37 PST by root via other
29      2005-03-12 11:01:20 PST by stu via cli
30      2005-03-12 10:57:35 PST by vwx via cli
31      2005-03-11 10:25:07 PST by yzz via cli
32      2005-03-10 23:40:58 PST by abc via cli
33      2005-03-10 23:40:38 PST by def via cli
34      2005-03-10 23:14:27 PST by ghi via cli
35      2005-03-10 23:10:16 PST by jkl via cli
36      2005-03-10 23:01:51 PST by mno via cli
37      2005-03-10 22:49:57 PST by pqr via cli
38      2005-03-10 22:24:07 PST by stu via cli
39      2005-03-10 22:20:14 PST by vwx via cli
40      2005-03-10 22:16:56 PST by yzz via cli
41      2005-03-10 22:16:41 PST by abc via cli
42      2005-03-10 20:44:00 PST by def via cli
43      2005-03-10 20:43:29 PST by ghi via cli
44      2005-03-10 20:39:14 PST by jkl via cli
45      2005-03-10 20:31:30 PST by root via other
46      2005-03-10 18:57:01 PST by mno via cli
47      2005-03-10 18:56:18 PST by pqr via cli
48      2005-03-10 18:47:49 PST by stu via cli
49      2005-03-10 18:47:34 PST by vw via cli
| Pipe through a command
[edit]
```

Comparing Configuration Changes with a Prior Version

In configuration mode only, when you have made changes to the configuration and want to compare the candidate configuration with a prior version, you can use the **compare** command to display the configuration. The **compare** command compares the candidate configuration with either the current committed configuration or a configuration file and displays the differences between the two configurations. To compare configurations, specify the **compare** command after the pipe:

```
[edit]
user@host# show | compare (filename| rollback n)
```

filename is the full path to a configuration file. The file must be in the proper format: a hierarchy of statements.

n is the index into the list of previously committed configurations. The most recently saved configuration is number 0, and the oldest saved configuration is number 49. If you do not specify arguments, the candidate configuration is compared against the active configuration file (**/config/juniper.conf**).

The comparison output uses the following conventions:

- Statements that are only in the candidate configuration are prefixed with a plus sign (+).
- Statements that are only in the comparison file are prefixed with a minus sign (-).
- Statements that are unchanged are prefixed with a single blank space ().

The following example shows various changes, then a comparison of the candidate configuration with the active configuration, showing only the changes made at the **[edit protocols bgp]** hierarchy level:

```
[edit]
user@host# edit protocols bgp
[edit protocols bgp]
user@host# show
group my-group {
    type internal;
    hold-time 60;
    advertise-inactive;
    allow 1.1.1.1/32;
}
group fred {
    type external;
    peer-as 33333;
    allow 2.2.2.2/32;
}
group test-peers {
    type external;
    allow 3.3.3.3/32;
}
[edit protocols bgp]
user@host# set group my-group hold-time 90
[edit protocols bgp]
user@host# delete group my-group advertise-inactive
[edit protocols bgp]
user@host# set group fred advertise-inactive
[edit protocols bgp]
user@host# delete group test-peers
[edit protocols bgp]
user@host# show | compare
[edit protocols bgp group my-group]
-hold-time 60;
+hold-time 90;
-advertise-inactive;
[edit protocols bgp group fred]
+advertise-inactive;
[edit protocols bgp]
-group test-peers {
    -type external;
    -allow 3.3.3.3/32;
}
[edit protocols bgp]
user@host# show
group my-group {
    type internal;
    hold-time 90;
```

```
    allow 1.1.1.1/32;
  }
  group fred {
    type external;
    advertise-inactive;
    peer-as 3333;
    allow 2.2.2.2/32;
  }
```

Creating and Returning to a Rescue Configuration

A rescue configuration allows you to define a known working configuration or a configuration with a known state that you can roll back to at any time. This alleviates the necessity of having to remember the rollback number with the **rollback** command. You use the rescue configuration when you need to roll back to a known configuration or as a last resort if your router or switch configuration and the backup configuration files become damaged beyond repair.

To save the most recently committed configuration as the rescue configuration so that you can return to it at any time, issue the **request system configuration rescue save** command:

```
user@host> request system configuration rescue save
```

To return to the rescue configuration, use the **rollback rescue** configuration mode command:

```
[edit]
user@host# rollback rescue
load complete
```



NOTE: If the rescue configuration does not exist, or if the rescue configuration is not a complete, viable configuration, then the **rollback** command fails, an error message appears, and the current configuration remains active.

To activate the rescue configuration that you have loaded, use the **commit** command:

```
[edit]
user@host# rollback rescue
load complete
[edit]
user@host# commit
```

To delete an existing rescue configuration, issue the **request system configuration rescue delete** command:

```
user@host> request system configuration rescue delete
user@host>
```

For more information about the **request system configuration rescue delete** and **request system configuration rescue save** commands, see the [CLI Explorer](#).

Saving a Configuration to a File

Save Junos OS configuration to a file so that you can edit it with a text editor of your choice. You can save your current configuration to an ASCII file, which saves the configuration in its current form, including any uncommitted changes. If more than one user is modifying the configuration, all changes made by all users are saved.

To save software configuration changes to an ASCII file, use the **save** configuration mode command:

```
[edit]
user@host# save filename
[edit]
user@host#
```

The contents of the current level of the statement hierarchy (and below) are saved, along with the statement hierarchy containing it. This allows a section of the configuration to be saved, while fully specifying the statement hierarchy.

By default, the configuration is saved to a file in your home directory, which is on the flash drive.

When you issue this command from anywhere in the hierarchy (except the top level), a **replace** tag is automatically included at the beginning of the file. You can use the **replace** tag to control how a configuration is loaded from a file.

```
user@host> file show /var/home/user/myconf
replace:
protocols {
  bgp {
    disable;
    group int {
      type internal;
    }
  }
  isis {
    disable;
    interface all {
      level 1 disable;
    }
    interface fxp0.0 {
      disable;
    }
  }
  ospf {
    traffic-engineering;
    reference-bandwidth 4g;
    ...
  }
}
```

Related Documentation

- [Returning to the Most Recently Committed Junos OS Configuration on page 1338](#)
- [Loading a Configuration from a File on page 1325](#)

- *Viewing Files and Directories on a Device Running Junos OS*

Reverting to the Default Factory Configuration

If for any reason the current active configuration fails, you can revert to the default factory configuration. The default factory configuration contains the basic configuration settings. This is the first configuration of the switch, and it is loaded when the switch is first installed and powered on.

The **load factory default** command is a standard Junos OS configuration command. This configuration command replaces the current active configuration with the default factory configuration.

To revert the switch to the rescue configuration:

```
[edit]
user@switch# load factory-default
[edit]
user@switch# delete system commit factory-settings
[edit]
user@switch# commit
```

Related Documentation

- [Understanding Configuration Files on page 1310](#)
- [Loading a Previous Configuration File on page 1330](#)
- [Reverting to the Rescue Configuration on page 1346](#)

Using Rescue Configurations

- [Setting or Deleting the Rescue Configuration on page 1345](#)
- [Creating and Returning to a Rescue Configuration on page 1345](#)
- [Reverting to the Rescue Configuration on page 1346](#)

Setting or Deleting the Rescue Configuration

A *rescue configuration* is a user-defined configuration that restores connectivity to the device. You set a current committed configuration to be the rescue configuration through the CLI. If someone inadvertently commits a configuration that denies management access to a device and the console port is not accessible, you can overwrite the invalid configuration and replace it with the rescue configuration. The rescue configuration is a previously committed, valid configuration. We recommend that the rescue configuration include the IP address (accessible from the network) for the management port.

To set the current active configuration as the rescue configuration:

```
user@switch> request system configuration rescue save
```

To delete an existing rescue configuration:

```
user@switch> request system configuration rescue delete
```

Related Documentation

- [Reverting to the Default Factory Configuration on page 1344](#)
- [Loading a Previous Configuration File on page 1330](#)
- [Configuration File Terms on page 1309](#)
- [CLI Explorer](#)

Creating and Returning to a Rescue Configuration

A rescue configuration allows you to define a known working configuration or a configuration with a known state that you can roll back to at any time. This alleviates the necessity of having to remember the rollback number with the **rollback** command. You use the rescue configuration when you need to roll back to a known configuration or as a last resort if your router or switch configuration and the backup configuration files become damaged beyond repair.

To save the most recently committed configuration as the rescue configuration so that you can return to it at any time, issue the **request system configuration rescue save** command:

```
user@host> request system configuration rescue save
```

To return to the rescue configuration, use the **rollback rescue** configuration mode command:

```
[edit]
user@host# rollback rescue
load complete
```



NOTE: If the rescue configuration does not exist, or if the rescue configuration is not a complete, viable configuration, then the **rollback** command fails, an error message appears, and the current configuration remains active.

To activate the rescue configuration that you have loaded, use the **commit** command:

```
[edit]
user@host# rollback rescue
load complete
[edit]
user@host# commit
```

To delete an existing rescue configuration, issue the **request system configuration rescue delete** command:

```
user@host> request system configuration rescue delete
user@host>
```

For more information about the **request system configuration rescue delete** and **request system configuration rescue save** commands, see the [CLI Explorer](#).

**Related
Documentation**

- [Comparing Configuration Changes with a Prior Version on page 1340](#)
- [Saving a Configuration to a File on page 1331](#)

Reverting to the Rescue Configuration

If someone inadvertently commits a configuration that denies management access to a device and the console port is not accessible, you can overwrite the invalid configuration and replace it with the rescue configuration. The rescue configuration is a previously committed, valid configuration.

To revert the switch to the rescue configuration:

1. Enter the **load override** command.

```
[edit]
user@switch# load override filename
```

2. Commit your changes.

```
[edit]  
user@switch# commit filename
```

Related Documentation

- [Reverting to the Default Factory Configuration on page 1344](#)

Comparing or Compressing Configuration Files

- [Comparing Configuration Changes with a Prior Version on page 1349](#)
- [Compressing the Current Configuration File on page 1351](#)

Comparing Configuration Changes with a Prior Version

In configuration mode only, when you have made changes to the configuration and want to compare the candidate configuration with a prior version, you can use the **compare** command to display the configuration. The **compare** command compares the candidate configuration with either the current committed configuration or a configuration file and displays the differences between the two configurations. To compare configurations, specify the **compare** command after the pipe:

```
[edit]  
user@host# show | compare (filename) rollback n
```

filename is the full path to a configuration file. The file must be in the proper format: a hierarchy of statements.

n is the index into the list of previously committed configurations. The most recently saved configuration is number 0, and the oldest saved configuration is number 49. If you do not specify arguments, the candidate configuration is compared against the active configuration file (**/config/juniper.conf**).

The comparison output uses the following conventions:

- Statements that are only in the candidate configuration are prefixed with a plus sign (+).
- Statements that are only in the comparison file are prefixed with a minus sign (-).
- Statements that are unchanged are prefixed with a single blank space ().

The following example shows various changes, then a comparison of the candidate configuration with the active configuration, showing only the changes made at the **[edit protocols bgp]** hierarchy level:

```
[edit]  
user@host# edit protocols bgp
```

```
[edit protocols bgp]
user@host# show
group my-group {
  type internal;
  hold-time 60;
  advertise-inactive;
  allow 1.1.1.1/32;
}
group fred {
  type external;
  peer-as 33333;
  allow 2.2.2.2/32;
}
group test-peers {
  type external;
  allow 3.3.3.3/32;
}
[edit protocols bgp]
user@host# set group my-group hold-time 90
[edit protocols bgp]
user@host# delete group my-group advertise-inactive
[edit protocols bgp]
user@host# set group fred advertise-inactive
[edit protocols bgp]
user@host# delete group test-peers
[edit protocols bgp]
user@host# show | compare
[edit protocols bgp group my-group]
-hold-time 60;
+hold-time 90;
-advertise-inactive;
[edit protocols bgp group fred]
+advertise-inactive;
[edit protocols bgp]
-group test-peers {
  -type external;
  -allow 3.3.3.3/32;
}
[edit protocols bgp]
user@host# show
group my-group {
  type internal;
  hold-time 90;
  allow 1.1.1.1/32;
}
group fred {
  type external;
  advertise-inactive;
  peer-as 3333;
  allow 2.2.2.2/32;
}
```

Related Documentation

- [Creating and Returning to a Rescue Configuration on page 1342](#)

Compressing the Current Configuration File

By default, the current operational configuration file is compressed and is stored in the file **juniper.conf.gz** in the **/config** file system, along with the last three committed versions of the configuration. If you have large networks, the current configuration file might exceed the available space in the **/config** file system. Compressing the current configuration file enables the file to fit in the file system, typically reducing the size of the file by 90 percent. You might want to compress your current operation configuration files when they reach 3 megabytes (MB) in size.

When you compress the current configuration file, the names of the configuration files change. To determine the size of the files in the **/config** file system, issue the **file list /config detail** command.



NOTE: We recommend that you compress the configuration files (this is the default) to minimize the amount of disk space that they require.

- If you want to compress the current configuration file, include the **compress-configuration-files** statement at the **[edit system]** hierarchy level:

```
[edit system]
compress-configuration-files;
```
- Commit the current configuration file to include the **compression-configuration-files** statement. Commit the configuration again to compress the current configuration file:

```
[edit system]
user@host# set compress-configuration-files
user@host# commit
commit complete
user@host# commit
commit complete
```
- If you do not want to compress the current operational configuration file, include the **no-compress-configuration-files** statement at the **[edit system]** hierarchy level:

```
[edit system]
no-compression-configuration-files;
```
- Commit the current configuration file to include the **no-compress-configuration-files** statement. Commit the configuration again to uncompress the current configuration file:

```
[edit system]
user@host# commit
commit complete
user@host# commit
commit complete
```

Related Documentation

- [Junos OS Commit Model for Router or Switch Configuration on page 1313](#)
- [compress-configuration-files on page 220](#)

Troubleshooting Procedures

- [Loading a Previous Configuration File on page 1353](#)
- [Reverting to the Default Factory Configuration on page 1354](#)
- [Reverting to the Rescue Configuration on page 1354](#)
- [Cleaning Up the System File Storage Space on page 1355](#)

Loading a Previous Configuration File

You can use the **rollback** *<number>* command to return to a previously committed configuration file. A switch saves the last 50 committed configurations, including the rollback number, date, time, and name of the user who issued the **commit** configuration command.

Syntax

rollback *<number>*

Options

- **none**—Return to the most recently saved configuration.
- **number**—Return to the specified configuration.
 - **Range:** 0 through 49. The most recently saved configuration is number 0, and the oldest saved configuration is number 49.
 - **Default:** 0

To return to a configuration prior to the most recently committed one:

1. Specify the rollback number (here, 1 is entered and the configuration returns to the previously committed configuration 0):

```
[edit]
user@switch# rollback 1
load complete
```

2. Activate the configuration you have loaded:

```
[edit]
user@switch# commit
```

- Related Documentation**
- [Configuration File Terms on page 1309](#)

Reverting to the Default Factory Configuration

If for any reason the current active configuration fails, you can revert to the default factory configuration. The default factory configuration contains the basic configuration settings. This is the first configuration of the switch, and it is loaded when the switch is first installed and powered on.

The **load factory default** command is a standard Junos OS configuration command. This configuration command replaces the current active configuration with the default factory configuration.

To revert the switch to the rescue configuration:

```
[edit]
user@switch# load factory-default
[edit]
user@switch# delete system commit factory-settings
[edit]
user@switch# commit
```

- Related Documentation**
- [Understanding Configuration Files on page 1310](#)
 - [Loading a Previous Configuration File on page 1330](#)
 - [Reverting to the Rescue Configuration on page 1346](#)

Reverting to the Rescue Configuration

If someone inadvertently commits a configuration that denies management access to a device and the console port is not accessible, you can overwrite the invalid configuration and replace it with the rescue configuration. The rescue configuration is a previously committed, valid configuration.

To revert the switch to the rescue configuration:

1. Enter the **load override** command.

```
[edit]
user@switch# load override filename
```

2. Commit your changes.

```
[edit]
user@switch# commit filename
```

- Related Documentation**
- [Reverting to the Default Factory Configuration on page 1344](#)

Cleaning Up the System File Storage Space

Problem **Description:** The system file storage space on the switch is full. Rebooting the switch does not solve the problem.

The following error message is displayed during a typical operation on the switch after the file storage space is full.

```
user@switch% cli
user@switch> configure
/var: write failed, filesystem is full
```

Solution Clean up the file storage on the switch by deleting system files.

1. Request to delete system files on the switch.

```
user@switch> request system storage cleanup
```

The list of files to be deleted is displayed.

List of files to delete:

Size	Date	Name
11B	Jul 26 20:55	/var/jail/tmp/alarmd.ts
124B	Aug 4 18:05	/var/log/default-log-messages.0.gz
1301B	Jul 26 20:42	/var/log/install.0.gz
387B	Jun 3 14:37	/var/log/install.1.gz
4920B	Aug 4 18:05	/var/log/messages.0.gz
20.0K	Jul 26 21:00	/var/log/messages.1.gz
16.3K	Jun 25 13:45	/var/log/messages.2.gz
804B	Aug 4 18:05	/var/log/security.0.gz
16.8K	Aug 3 11:15	/var/log/security.1.gz
487B	Aug 4 18:04	/var/log/wtmp.0.gz
855B	Jul 29 22:54	/var/log/wtmp.1.gz
920B	Jun 30 16:32	/var/log/wtmp.2.gz
94B	Jun 3 14:36	/var/log/wtmp.3.gz
353.2K	Jun 3 14:37	/var/sw/pkg/jloader-qfx-11.2I20110303_1117_dc-builder.tgz
124.0K	Jun 3 14:30	/var/tmp/gres-tp/env.dat
0B	Apr 14 16:20	/var/tmp/gres-tp/lock
0B	Apr 14 17:37	/var/tmp/if-rtssdb/env.lock
12.0K	Jul 26 20:55	/var/tmp/if-rtssdb/env.mem
2688.0K	Jul 26 20:55	/var/tmp/if-rtssdb/shm_usr1.mem
132.0K	Jul 26 20:55	/var/tmp/if-rtssdb/shm_usr2.mem
2048.0K	Jul 26 20:55	/var/tmp/if-rtssdb/trace.mem
155B	Jul 26 20:55	/var/tmp/krt_gencfg_filter.txt
0B	Jul 26 20:55	/var/tmp/rtssdb/if-rtssdb
1400.6K	Aug 3 10:13	/var/tmp/sfid.core.0.gz
1398.9K	Aug 3 17:01	/var/tmp/sfid.core.1.gz

Delete these files ? [yes,no] (no)

2. Enter **yes** to delete the files.

3. Reboot the switch.



BEST PRACTICE: We recommend that you regularly request a system file storage cleanup to optimize the performance of the switch.

Related Documentation • [request system storage cleanup on page 399](#)

PART 18

Configuration Statements and Operational Commands

- [Configuration Statements on page 1359](#)
- [Operational Commands on page 1367](#)

CHAPTER 41

Configuration Statements

- [archival on page 1360](#)
- [archive-sites \(Configuration File\) on page 1361](#)
- [configuration on page 1363](#)
- [transfer-interval \(Configuration\) on page 1364](#)
- [transfer-on-commit on page 1365](#)

archival

```
Syntax archival {
    configuration {
        archive-sites {
            file://<path>/<filename>;
            ftp://username@host:<port>url-path password password;
            http://username@host:<port>url-path password password;
            pasvftp://username@host:<port>url-path password password;
            scp://username@host:<port>url-path password password;
        }
        transfer-interval interval;
        transfer-on-commit;
    }
}
```

Hierarchy Level [edit system]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.

Description Configure copying of the currently active configuration to an archive site. An archive site can be a file, or an FTP, HTTP, or SCP location.

Options The remaining statements are explained separately.





NOTE: The [edit system archival] hierarchy is not available on QFabric systems.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- [Using Junos OS to Configure a Router or Switch to Transfer Its Configuration to an Archive Site on page 1333](#)

archive-sites (Configuration File)

Syntax	<pre>archive-sites { file://<path>/<filename>; ftp://username@host:<port>url-path password password; http://username@host:<port>url-path password password; pasvftp://username@host:<port>url-path password password; scp://username@host:<port>url-path password password; }</pre>
Hierarchy Level	[edit system archival configuration]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.</p>
Description	<p>Specify where to transfer the current configuration files. When specifying a URL in a Junos OS statement using an IPv6 host address, you must enclose the entire URL in quotation marks (" ") and enclose the IPv6 host address in brackets ([]). For example, "scp://username<:password>@[ipv6-host-address]<:port>/url-path"</p> <p>If you specify more than one archive site, the router or switch attempts to transfer the configuration files to the first archive site in the list, moving to the next only if the transfer fails.</p> <p>The destination filename is saved in the following format, where <i>n</i> corresponds to the number of the compressed configuration rollback file that has been archived:</p> <p><i>router-name_YYYYMMDD_HHMMSS_juniper.conf.n.gz</i></p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p> NOTE: The time included in the destination filename is always in Coordinated Universal Time (UTC) regardless of whether the time on the router or switch is configured as UTC or the local time zone. The default time zone on the router or switch is UTC.</p> </div> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p> NOTE: The [edit system archival] hierarchy is not available on QFabric systems.</p> </div>
Options	<p>The prefix used in the configuration statement determines the form of transfer:</p> <p>file:// —transfer on a path to a named file</p> <p>ftp:// —transfer using active FTP server</p> <p>http:// —transfer using HTTP server</p>


pasvftp:// —transfer to a device that only accepts passive FTP services

scp:// —transfer to a known host using background SCP file transfers

Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
---------------------------------	---

Related Documentation	<ul style="list-style-type: none">• Configuring Archive Sites for the Transfer of Active Configuration Files on page 1334• Junos OS Commit Model for Router or Switch Configuration on page 1313• configuration on page 1363• transfer-on-commit on page 1365
------------------------------	--

configuration

Syntax	<pre> configuration { transfer-interval interval; transfer-on-commit; archive-sites { file://<path>/<filename>; ftp://username@host:<port>url-path password password; http://username@host:<port>url-path password password; pasvftp://username@host:<port>url-path password password; scp://username@host:<port>url-path password password; } } </pre>
Hierarchy Level	[edit system archival]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.</p>
Description	Configure the router or switch to periodically transfer its currently active configuration (or after each commit).
	<div>  <p>NOTE: The [edit system archival] hierarchy is not available on QFabric systems.</p> </div>
Options	The remaining statements are explained separately.
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Using Junos OS to Configure a Router or Switch to Transfer Its Configuration to an Archive Site on page 1333 • archive on page 1692 • archive-sites on page 1361 • transfer-interval on page 1364 • transfer-on-commit on page 1365

transfer-interval (Configuration)

Syntax	<code>transfer-interval <i>interval</i>;</code>
Hierarchy Level	[edit system archival configuration]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the router or switch to periodically transfer its currently active configuration to an archive site.
Options	interval —Interval at which to transfer the current configuration to an archive site. Range: 15 through 2880 minutes



NOTE: The [edit system archival] hierarchy is not available on QFabric systems.

Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Periodic Transfer of the Active Configuration to an Archive Site on page 1334• archive on page 1692• configuration on page 1363• transfer-on-commit on page 1365

transfer-on-commit

Syntax	transfer-on-commit;
Hierarchy Level	[edit system archival configuration]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Description	Configure the router or switch to transfer its currently active configuration to an archive site each time you commit a candidate configuration.



NOTE: When specifying a URL in a Junos OS statement using an IPv6 host address, you must enclose the entire URL in quotation marks (" ") and enclose the IPv6 host address in brackets ([]). For example, "ftp://username<:password>@[ipv6-host-address]<:port>/url-path" .



NOTE: The [edit system archival] hierarchy is not available on QFabric systems.

Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Transfer of the Currently Active Configuration When a Configuration Is Committed on page 1334 • archive on page 1692 • configuration on page 1363 • transfer-interval on page 1364

CHAPTER 42

Operational Commands

- clear log
- clear system commit
- file archive
- file checksum md5
- file checksum sha1
- file checksum sha-256
- file compare
- file delete
- file list
- file rename
- file show
- request system configuration rescue delete
- request system configuration rescue save
- show system commit
- show system configuration archival
- show system configuration rescue
- show system rollback
- test configuration

clear log

Syntax	<code>clear log <i>filename</i></code> <code><all></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Description	Remove contents of a log file.
Options	<i>filename</i> —Name of the specific log file to delete. all —(Optional) Delete the specified log file and all archived versions of it.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show log on page 1003
List of Sample Output	clear log on page 1368
Output Fields	See file list for an explanation of output fields.

Sample Output

clear log

The following sample commands list log file information, clear the contents of a log file, and then display the updated log file information:

```
user@host> file list lcc0-re0:/var/log/sampled detail
lcc0-re0:
-----
-rw-r-----  1 root  wheel          26450 Jun 23 18:47 /var/log/sampled
total 1

user@host> clear log lcc0-re0:sampled
lcc0-re0:
-----

user@host> file list lcc0-re0:/var/log/sampled detail
lcc0-re0:
-----
-rw-r-----  1 root  wheel           57 Sep 15 03:44 /var/log/sampled
total 1
```


clear system commit

Syntax	clear system commit
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Clear any pending commit operation.
Options	This command has no options.
Required Privilege Level	maintenance (or the actual user who scheduled the commit)
Related Documentation	<ul style="list-style-type: none"> • show system commit on page 1054
List of Sample Output	<p>clear system commit on page 1369</p> <p>clear system commit (None Pending) on page 1369</p> <p>clear system commit (User Does Not Have Required Privilege Level) on page 1369</p>
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear system commit

```
user@host> clear system commit
Pending commit cleared.
```

clear system commit (None Pending)

```
user@host> clear system commit
No commit scheduled.
```

clear system commit (User Does Not Have Required Privilege Level)

```
user@host> clear system commit
error: Permission denied
```

file archive

Syntax	<code>file archive destination <i>destination</i> source <i>source</i> <compress></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Description	Archive, and optionally compress, one or multiple local system files as a single file, locally or at a remote location.
Options	<p>destination <i>destination</i>—Destination of the archived file or files. Specify the destination as a URL or filename. The Junos OS adds one of the following suffixes if the destination filename does not already have it:</p> <ul style="list-style-type: none">• For archived files—The suffix .tar• For archived and compressed files—The suffix .tgz <p>source <i>source</i>—Source of the original file or files. Specify the source as a URL or filename.</p> <p>compress—(Optional) Compress the archived file with the GNU zip (gzip) compression utility. The compressed files have the suffix .tgz.</p>
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none">• Format for Specifying Filenames and URLs in Junos OS CLI Commands on page 40
List of Sample Output	file archive (Multiple Files) on page 1370 file archive (Single File) on page 1370 file archive (with Compression) on page 1371
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

file archive (Multiple Files)

The following sample command archives all message files in the local directory `/var/log/messages` as the single file `messages-archive.tar`.

```
user@host> file archive source /var/log/messages* destination /var/log/messages-archive.tar
/usr/bin/tar: Removing leading / from absolute path names in the archive.
user@host>
```

file archive (Single File)

The following sample command archives one message file in the local directory `/var/log/messages` as the single file `messages-archive.tar`.

```
user@host> file archive source /var/log/messages destination /var/log/messages-archive.tar
/usr/bin/tar: Removing leading / from absolute path names in the archive.
user@host
```

file archive (with Compression)

The following sample command archives and compresses all message files in the local directory **/var/log/messages** as the single file **messages-archive.tgz**.

```
user@host> file archive compress source /var/log/messages* destination
/var/log/messages-archive.tgz
/usr/bin/tar: Removing leading / from absolute path names in the archive.
```

file checksum md5

Syntax	<code>file checksum md5 <pathname> filename</code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Calculate the Message Digest 5 (MD5) checksum of a file.
Options	pathname —(Optional) Path to a filename. filename —Name of a local file for which to calculate the MD5 checksum.
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Checksum Hashes for a Commit Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i>• <i>Configuring Checksum Hashes for an Event Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i>• <i>Configuring Checksum Hashes for an Op Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i>• <i>Executing an Op Script from a Remote Site</i> in the <i>Junos OS Configuration and Operations Automation Guide</i>• file checksum sha-256 on page 346• file checksum sha1 on page 345• <i>op</i>
List of Sample Output	file checksum md5 on page 1372
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

file checksum md5

```
user@host> file checksum md5 jbundle-5.3R2.4-export-signed.tgz
MD5 (jbundle-5.3R2.4-export-signed.tgz) = 2a3b69e43f9bd4893729cc16f505a0f5
```

file checksum sha1

Syntax	<code>file checksum sha1 <pathname> filename</code>
Release Information	<p>Command introduced in Junos OS Release 9.5.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	Calculate the Secure Hash Algorithm (SHA-1) checksum of a file.
Options	<p>pathname—(Optional) Path to a filename.</p> <p>filename—Name of a local file for which to calculate the SHA-1 checksum.</p>
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Checksum Hashes for a Commit Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i> • <i>Configuring Checksum Hashes for an Event Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i> • <i>Configuring Checksum Hashes for an Op Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i> • <i>Executing an Op Script from a Remote Site</i> in the <i>Junos OS Configuration and Operations Automation Guide</i> • file checksum md5 on page 344 • file checksum sha-256 on page 346 • <i>op</i>
List of Sample Output	file checksum sha1 on page 1373
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

file checksum sha1

```
user@host> file checksum sha1 /var/db/scripts/opscript.slax
```

```
SHA1 (/var/db/scripts/commitscript.slax) = ba9e47120c7ce55cff29afd73eacd370e162c676
```

file checksum sha-256

Syntax	<code>file checksum sha-256 <pathname> filename</code>
Release Information	Command introduced in Junos OS Release 9.5. Command introduced in Junos OS Release 9.5 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Calculate the Secure Hash Algorithm 2 family (SHA-256) checksum of a file.
Options	pathname —(Optional) Path to a filename. filename —Name of a local file for which to calculate the SHA-256 checksum.
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Checksum Hashes for a Commit Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i>• <i>Configuring Checksum Hashes for an Event Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i>• <i>Configuring Checksum Hashes for an Op Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i>• <i>Executing an Op Script from a Remote Site</i> in the <i>Junos OS Configuration and Operations Automation Guide</i>• file checksum md5 on page 344• file checksum sha1 on page 345• <i>op</i>
List of Sample Output	file checksum sha-256 on page 1374
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

file checksum sha-256

```
user@host> file checksum sha-256 /var/db/scripts/commitscript.slax

SHA256 (/var/db/scripts/commitscript.slax) =
94c2b061fb55399e15babd2529453815601a602b5c98e5c12ed929c9d343dd71
```

file compare

Syntax	<pre>file compare (files <i>filename filename</i>) <context unified> <ignore-white-space></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.</p>
Description	<p>Compare two local files and describe the differences between them in default, context, or unified output styles:</p> <ul style="list-style-type: none"> • Default—In the first line of output, c means lines were changed between the two files, d means lines were deleted between the two files, and a means lines were added between the two files. The numbers preceding this alphabetical marker represent the first file, and the lines after the alphabetical marker represent the second file. A left angle bracket (<) in front of output lines refers to the first file. A right angle bracket (>) in front of output lines refers to the second file. • Context—The display is divided into two parts. The first part is the first file; the second part is the second file. Output lines preceded by an exclamation point (!) have changed. Additions are marked with a plus sign (+), and deletions are marked with a minus sign (-). • Unified—The display is preceded by the line number from the first and the second file (xx,xxx,x). Before the line number, additions to the file are marked with a plus sign (+), and deletions to the file are marked with a minus sign (-). The body of the output contains the affected lines. Changes are viewed as additions plus deletions.
Options	<p>files <i>filename</i>—Names of two local files to compare.</p> <p>context—(Optional) Display output in context format.</p> <p>ignore-white-space—(Optional) Ignore changes in the amount of white space.</p> <p>unified—(Optional) Display output in unified format.</p>
Required Privilege Level	none
Related Documentation	<ul style="list-style-type: none"> • Format for Specifying Filenames and URLs in Junos OS CLI Commands on page 40 • Viewing Core Files from Junos OS Processes on page 34
List of Sample Output	<p>file compare files on page 1377</p> <p>file compare files context on page 1377</p> <p>file compare files unified on page 1377</p> <p>file compare files unified ignore-white-space on page 1377</p>

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

file compare files

```
user@host> file compare files /tmp/one /tmp/two
100c100
<          full-name "File 1";
---
>          full-name "File 2";
102c102
<          class foo; # 'foo' is not defined
---
>          class super-user;
```

file compare files context

```
user@host> file compare files /tmp/one /tmp/two context
*** /tmp/one   Wed Dec  3 17:12:50 2003
--- /tmp/two   Wed Dec  3 09:13:14 2003
*****
*** 97,104 ****
        }
    }
    user bill {
!         full-name "Bill Smith";
!         class foo; # 'foo' is not defined
        authentication {
            encrypted-password SECRET;
        }
--- 97,105 ----
    }
    user bill {
!         full-name "Bill Smith";
!         uid 1089;
!         class super-user;
        authentication {
            encrypted-password SECRET;
        }
    }
```

file compare files unified

```
user@host> file compare files /tmp/one /tmp/two unified
--- /tmp/one   Wed Dec  3 17:12:50 2003
+++ /tmp/two   Wed Dec  3 09:13:14 2003
@@ -97,8 +97,9 @@
    }
}
user bill {
-     full-name "Bill Smith";
-     class foo; # 'foo' is not defined
+     full-name "Bill Smith";
+     uid 1089;
+     class super-user;
    authentication {
        encrypted-passwordSECRET;
    }
}
```

file compare files unified ignore-white-space

```
user@host> file compare files /tmp/one /tmp/two unified ignore-white-space
```

```
--- /tmp/one    Wed Dec  3 09:13:10 2003
+++ /tmp/two    Wed Dec  3 09:13:14 2003
@@ -99,7 +99,7 @@
     user bill {
         full-name "Bill Smith";
         uid 1089;
-        class foo; # 'foo' is not defined
+        class super-user;
         authentication {
             encrypted-password <SECRET>; # SECRET-DATA
         }
     }
```

file delete

Syntax	<code>file delete <i>filename</i></code> <code><purge></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Description	Delete a file on the local router or switch.
Options	<i>filename</i> —Name of the file to delete. For a routing matrix, include chassis information in the filename if the file to be deleted is not local to the Routing Engine from which the command is issued. <i>purge</i> —(Optional) Overwrite regular files before deleting them.
Required Privilege Level	maintenance
List of Sample Output	file delete on page 1379 file delete (Routing Matrix) on page 1379
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

file delete

```
user@host> file list /var/tmp
dcd.core
rpd.core
snmpd.core

user@host> file delete /var/tmp/snmpd.core
user@host> file list /var/tmp
dcd.core
rpd.core
```

file delete (Routing Matrix)

```
user@host> file list lcc0-re0:/var/tmp
dcd.core
rpd.core
snmpd.core

user@host> file delete lcc0-re0:/var/tmp/snmpd.core
user@host> file list /var/tmp
dcd.core
rpd.core
```

file list

Syntax	<code>file list</code> <code><detail recursive></code> <code><filename></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display a list of files on the local router or switch.
Options	none —Display a list of all files for the current directory. detail recursive —(Optional) Display detailed output or descend recursively through the directory hierarchy, respectively. filename —(Optional) Display a list of files. For a routing matrix, the filename must include the chassis information.
Additional Information	The default directory is the home directory of the user logged in to the router or switch. To view available directories, enter a space and then a backslash (/) after the file list command. To view files within a specific directory, include a backslash followed by the directory and, optionally, subdirectory name after the file list command.
Required Privilege Level	maintenance
List of Sample Output	file list on page 1380 file list (Routing Matrix) on page 1380
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

file list

```
user@host> file list /var/tmp
dcd.core
rpd.core
snmpd.core
```

file list (Routing Matrix)

```
user@host> file list lcc0-re0:var/tmp
lcc0-re0:
-----
/var/tmp/:
.gdbinit
.pccardd
Test/
chassisd*
chassisd.nathan*
check_time*
```

```
cores/  
diagTestPrep*  
diagtest*  
diagtest.regress*  
do_switchovers*  
dump_test*  
err.manoj.log  
esw_clearstats*  
esw_counter*  
esw_debug*  
esw_debug_ge*  
esw_filt_test*  
esw_filter_tnp_addr*  
esw_getstats*  
esw_phy*  
esw_stats*
```

file rename

Syntax	<code>file rename <i>source destination</i></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Description	Rename a file on the local router or switch.
Options	<i>destination</i> —New name for the file. <i>source</i> —Original name of the file. For a routing matrix, the filename must include the chassis information.
Required Privilege Level	maintenance
List of Sample Output	file rename on page 1382 file rename (Routing Matrix) on page 1382
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

file rename

The following example lists the files in `/var/tmp`, renames one of the files, and then displays the list of files again to reveal the newly named file.

```
user@host> file list /var/tmp
dcd.core
rpd.core
snmpd.core

user@host> file rename /var/tmp/dcd.core /var/tmp/dcd.core.990413
user@host> file list /var/tmp
dcd.core.990413
rpd.core
snmpd.core
```

file rename (Routing Matrix)

The following example lists the files in `/var/tmp`, renames one of the files, and then displays the list of files again to reveal the newly named file.

```
user@host> file list lcc0-re1:/var/tmp
lcc0-re1:
-----

/var/tmp:
.pccardd
sartre.conf
snmpd
syslogd.core-tarball.0.tgz
```

```
user@host> file rename lcc0-re0:/var/tmp/snmpd /var/tmp/snmpd.rr
user@host> file list lcc0-re1:/var/tmp
lcc0-re1:
```

```
-----
/var/tmp:
.pccardd
sartre.conf
snmpd.rr
syslogd.core-tarball.0.tgz
```

file show

Syntax	<code>file show <i>filename</i></code> <code><encoding (base64 raw)></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display the contents of a file.
Options	<i>filename</i> —Name of a file. For a routing matrix, the filename must include the chassis information. <code>encoding (base64 raw)</code> —(Optional) Encode file contents with base64 encoding or show raw text.
Required Privilege Level	maintenance
List of Sample Output	file show on page 1384 file show (Routing Matrix) on page 1384
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

file show

```
user@host> file show /var/log/messages
Apr 13 21:00:08 romney /kernel: so-1/1/2: loopback suspected; going to standby.
Apr 13 21:00:40 romney /kernel: so-1/1/2: loopback suspected; going to standby.
Apr 13 21:02:48 romney last message repeated 4 times
Apr 13 21:07:04 romney last message repeated 8 times
Apr 13 21:07:13 romney /kernel: so-1/1/0: Clearing SONET alarm(s) RDI-P
Apr 13 21:07:29 romney /kernel: so-1/1/0: Asserting SONET alarm(s) RDI-P
...
```

file show (Routing Matrix)

```
user@host> file show lcc0-re0:/var/tmp/.gdbinit
lcc0-re0:
-----
#####
# Settings
#####

set print pretty


#####
# Basic stuff
#####

define msgbuf
    printf "%s", msgbufp->msg_ptr
end
```



```
# hex dump of a block of memory
# usage: dump address length
define dump
  p $arg0, $arg1
  set $ch = $arg0
  set $j = 0
  set $n = $arg1
  while ($j < $n)
    #printf "%x %x ",&$ch[$j],$ch[$j]
    printf "%x ",$ch[$j]
    set $j = $j + 1
    if (!($j % 16))
      printf "\n"
    end
  end
end
end
```

request system configuration rescue delete


Syntax	request system configuration rescue delete
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Description	Delete an existing rescue configuration.
	<div> NOTE: The [edit system configuration] hierarchy is not available on QFabric systems.</div>
Options	This command has no options.
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none">• request system configuration rescue save on page 383• request system software rollback on page 1292• show system commit on page 1054
List of Sample Output	request system configuration rescue delete on page 1386
Output Fields	This command produces no output.

Sample Output

request system configuration rescue delete

```
user@host> request system configuration rescue delete
```

request system configuration rescue save

Syntax	request system configuration rescue save
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.</p>
Description	Save the most recently committed configuration as the rescue configuration so that you can return to it at any time by using the rollback command.
<div>  <p>NOTE: The [edit system configuration] hierarchy is not available on QFabric systems.</p> </div>	
Options	This command has no options.
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none"> • request system software delete on page 1286 • request system software rollback on page 1292 • show system commit on page 1054
List of Sample Output	request system configuration rescue save on page 1387
Output Fields	This command produces no output.

Sample Output

request system configuration rescue save

```
user@host> request system configuration rescue save
```

show system commit


Syntax	<pre>show system commit <revision> <server></pre>	
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Option server introduced in Junos OS Release 12.1 for the PTX Series router.</p> <p>Option revision introduced in Junos OS Release 14.1.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.</p>	
Description	Display the system commit history and any pending commit operation.	
Options	<p>none—Display the last 50 commit operations listed, most recent to first.</p> <p>revision—(Optional) Display the revision number of the active configuration of the Routing Engine(s).</p> <p>server—(Optional) Display commit server status.</p>	
	<div>  <p>NOTE: By default, the status of the commit server is “Not running”. The commit server starts running only when a commit job is added to the batch.</p> </div>	
Required Privilege Level	view	
Related Documentation	<ul style="list-style-type: none"> clear system commit on page 336 show system commit revision 	
List of Sample Output	<p>show system commit on page 1390</p> <p>show system commit (At a Particular Time) on page 1390</p> <p>show system commit (At the Next Reboot) on page 1390</p> <p>show system commit (Rollback Pending) on page 1390</p> <p>show system commit (QFX Series) on page 1390</p>	
Output Fields	Table 72 describes the output fields for the show system commit command. Output fields are listed in the approximate order in which they appear.	

Table 88: show system commit Output Fields

Field Name	Field Description	Level of Output
<number>	Displays the last 50 commit operations listed, most recent to first. The identifier <number> designates a configuration created for recovery using the request system configuration rescue save command.	none

Table 88: show system commit Output Fields (*continued*)

Field Name	Field Description	Level of Output
<time-stamp>	Date and time of the commit operation.	none
<root>/<username>	User who executed the commit operation.	none
<method>	Method used to execute the commit operation: <ul style="list-style-type: none"> • CLI—CLI interactive user performed the commit operation. • Junos XML protocol—Junos XML protocol client performed the commit operation. • synchronize—The commit synchronize command was performed on the other Routing Engine. • snmp—An SNMP set request caused the commit operation. • button—A button on the router or switch was pressed to commit a rescue configuration for recovery. • autoinstall—A configuration obtained through autoinstallation was committed. • other—When there is no login name associated with the session, the values for user and client default to root and other. For example, during a reboot after package installation, mgd commits the configuration as a system commit, and there is no login associated with the commit. 	none

Sample Output

show system commit

```
user@host> show system commit
0   2003-07-28 19:14:04 PDT by root via other
1   2003-07-25 22:01:36 PDT by regress via cli
2   2003-07-25 22:01:32 PDT by regress via cli
3   2003-07-25 21:30:13 PDT by root via button
4   2003-07-25 13:46:48 PDT by regress via cli
5   2003-07-25 05:33:21 PDT by root via autoinstall
...
rescue 2002-05-10 15:32:03 PDT by root via other
```

show system commit (At a Particular Time)

```
user@host> show system commit
commit requested by root via cli at Tue May  7 15:59:00 2002
```

show system commit (At the Next Reboot)

```
user@host> show system commit
commit requested by root via cli at reboot
```

show system commit (Rollback Pending)

```
user@host> show system commit
0 2005-01-05 15:00:37 PST by root via cli commit confirmed, rollback in 3mins
```

show system commit (QFX Series)

```
user@switch> show system commit
0 2011-11-25 19:17:49 PST by root via cli
```

show system configuration archival

Syntax show system configuration archival

Release Information Introduced in Junos OS Release 7.6.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Display directory and number of files queued for archival transfer.



NOTE: The [edit system configuration] hierarchy is not available on QFabric systems.

Options This command has no options.

Required Privilege Level maintenance

List of Sample Output [show system configuration archival on page 1391](#)

Sample Output

show system configuration archival

```
user@host> show system configuration archival
```

```
/var/transfer/config/:  
total 8
```

show system configuration rescue

Syntax show system configuration rescue

Release Information Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.0 for EX Series switches.
 Command introduced in Junos OS Release 11.1 for the QFX Series.
 Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.

Description Display a rescue configuration, if one exists.



NOTE: The [edit system configuration] hierarchy is not available on QFabric systems.

Options This command has no options.

Required Privilege Level maintenance

Related Documentation • [show system configuration archival on page 1057](#)

List of Sample Output [show system configuration rescue on page 1392](#)

Sample Output

show system configuration rescue

```
user@switch> show system configuration rescue
version "7.3"; groups {
  global {
    system {
      host-name router1;
      domain-name customer.net;
      domain-search [ customer.net ];
      backup-router 192.168.124.254;
      name-server {
        172.17.28.11;
        172.17.28.101;
        172.17.28.100;
        172.17.28.10;
      }
      login {
        user regress {
          uid 928;
          class ;
          shell csh;
          authentication {
            encrypted-password "$1$kPU..$w.4FGRAGanJ8U4Yq6sbj7."; ##
SECRET-DATA
          }
        }
      }
    }
  }
}
```



```
        services {  
            ftp;  
            rlogin;  
            rsh;  
            telnet;  
        }  
    }  
.....
```

show system rollback

Syntax `show system rollback number`
`<compare number>`

Release Information Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.0 for EX Series switches.
 Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
 Command introduced in Junos OS Release 11.1 for the QFX Series.

Description Display the contents of a previously committed configuration, or the differences between two previously committed configurations.



NOTE: The `show system rollback` command is a purely operational mode command and cannot be issued with `run` from the configuration mode.

Options *number*—Number of a configuration to view. The output displays the configuration. The range of values is 0 through 49.

compare number —(Optional) Number of another previously committed (rollback) configuration to compare to rollback *number*. The output displays the differences between the two configurations. The range of values is 0 through 49.

Required Privilege Level view

List of Sample Output [show system rollback compare on page 1394](#)

Sample Output

show system rollback compare

```
user@host> show system rollback 3 compare 1
[edit]
+ interfaces {
+   ge-1/1/1 {
+     unit 0 {
+       family inet {
+         filter {
+           input mf_plp;
+         }
+         address 14.1.1.1/30;
+       }
+     }
+   }
+   ge-1/2/1 {
+     unit 0 {
+       family inet {
+         filter {
+           input mf_plp;
+         }
+         address 13.1.1.1/30;
```

```
+      }
+    }
+  }
+  ge-1/3/0 {
+    unit 0 {
+      family inet {
+        filter {
+          input mf_plp;
+        }
+        address 12.1.1.1/30;
+      }
+    }
+  }
+}
```

test configuration

Syntax	<code>test configuration <i>filename</i></code> <code>syntax-only</code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. syntax-only option introduced in Junos OS Release 12.1. Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Description	Verify that the syntax of a configuration file is correct. If the configuration contains any syntax or commit check errors, a message is displayed to indicate the line number and column number in which the error was found. This command only accepts text files.
Options	<i>filename</i> —Name of the configuration file. syntax-only —(Optional) Check the syntax of a partial configuration file, without checking for commit errors.
Required Privilege Level	view
List of Sample Output	test configuration on page 1396
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

test configuration

```
user@host> test configuration terminal
[Type ^D to end input]
system {
host-name bluesky;
paris-23;
login;
}
terminal:3:(8) syntax error: paris
[edit system]
    'paris-23;'
    syntax error
terminal:4:(11) statement must contain additional statements: ;
[edit system login]
    'login ;'
    statement must contain additional statements
configuration syntax failed
```

Network Management and Monitoring Feature Guide for QFX10000 Switches

PART 19

Overview

- [Understanding Network Management on page 1401](#)
- [Junos Space on page 1405](#)

Understanding Network Management

- [Understanding Device and Network Management Features on page 1401](#)

Understanding Device and Network Management Features

After you install a QFX Series product, OCX Series device, or EX4600 switch in your network, you need to manage the device. The products support features that you use to manage the device within the network, including the management of configuration, system performance, fault monitoring, and remote access.

[Table 89](#) lists the device and network management features on the QFX Series, OCX Series, and EX4600.

Table 89: Device and Network Management Features on the QFX Series, OCX Series, and EX4600

Feature	Typical Uses	Documentation
AI-Scripts and Advanced Insight Manager (AIM)—Automatically detect and monitor faults on the switch, and depending on the configuration on the AIM application, send notifications of potential problems, and submit problem reports to Juniper Support Systems.	Fault management	Advanced Insight Scripts (AI-Scripts) Release Notes
Alarms and LEDs on the switch—Show status of hardware components and indicate warning or error conditions.	Fault management	Chassis Alarm Messages on a QFX3500 Device
Firewall filters—Control the packets that are sent to and from the network, balance network traffic, and optimize performance.	Performance management	<ul style="list-style-type: none">• Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices• Overview of Firewall Filters on page 5951

Table 89: Device and Network Management Features on the QFX Series, OCX Series, and EX4600 (continued)

Feature	Typical Uses	Documentation
In-band management—Enables connection to the switch using the same interfaces through which customer traffic flows. Communication between the switch and a remote console is typically enabled using SSH and Telnet services. SSH provides secure encrypted communications, whereas Telnet provides unencrypted, and therefore less secure, access to the switch.	Remote access management	<ul style="list-style-type: none"> • Configuring SSH Service for Remote Access to the Router or Switch on page 1412 • Configuring Telnet Service for Remote Access to a Router or Switch
Juniper Networks Junos OS automation scripts—Configuration and operations automation tools provided by Junos OS. These tools include commit scripts, operation scripts, event scripts, and event policies. Commit scripts enforce custom configuration rules, whereas operation scripts, event policies, and event scripts automate network troubleshooting and management.	<ul style="list-style-type: none"> • Configuration management • Performance management • Fault management 	<i>Automation Scripting Feature Guide</i>
Junos OS command-line interface (CLI)—CLI configuration statements that enable you to configure the switch based on your networking requirements, such as security, service, and performance.	<ul style="list-style-type: none"> • Configuration management • Performance management • User access management • Remote access management 	<i>CLI User Guide</i>
Junos Space software—Multipurpose GUI-based network management system that includes a base platform, the Network Application Platform, and other optional applications such as Ethernet Design, Service Now, Service Insight, and Virtual Control. NOTE: Junos Space does not support the OCX Series.	<ul style="list-style-type: none"> • Configuration management • Performance management • Fault management 	<ul style="list-style-type: none"> • Understanding Junos Space Support on page 1405 • Junos Space Network Application Platform User Guide
Junos XML API—XML representation of Junos OS configuration statements and operational mode commands. Junos XML configuration tag elements are the content to which the Junos XML protocol operations apply. Junos XML operational tag elements are equivalent in function to operational mode commands in the CLI, which you can use to retrieve status information for a device. The Junos XML API also includes tag elements that are the counterpart to Junos CLI configuration statements.	<ul style="list-style-type: none"> • Configuration management • Performance management • Fault management 	<ul style="list-style-type: none"> • Junos XML API Configuration Developer Reference • Junos XML API Operational Developer Reference

Table 89: Device and Network Management Features on the QFX Series, OCX Series, and EX4600 (*continued*)

Feature	Typical Uses	Documentation
NETCONF XML management protocol—XML-based management protocol that client applications use to request and change configuration information on routing, switching, and security platforms running Junos OS. The NETCONF XML management protocol defines basic operations that are equivalent to Junos OS CLI configuration mode commands. Client applications use the protocol operations to display, edit, and commit configuration statements (among other operations), just as administrators use CLI configuration mode commands such as show , set , and commit to perform those operations.	<ul style="list-style-type: none"> • Configuration management • Performance management • Fault management 	<i>NETCONF XML Management Protocol Developer Guide</i>
Operational mode commands—May be used to do the following: <ul style="list-style-type: none"> • Monitor switch performance. For example, the show chassis routing-engine command shows the CPU utilization of the Routing Engine. High CPU utilization of the Routing Engine can affect performance of the switch. • View current activity and status of the device or network. For example, you can use the ping command to monitor and diagnose connectivity problems, and the traceroute command to locate points of failure on the network. 	<ul style="list-style-type: none"> • Performance management • Fault management 	CLI Explorer
Out-of-band management—Enables connection to the switch through a management interface. Out-of-band management is supported on two dedicated management Ethernet interfaces as well as on the console and auxiliary ports. The management Ethernet interfaces connect directly to the Routing Engine. No transit traffic is allowed through the interfaces, separating customer and management traffic and ensuring that congestion or failures in the transit network do not affect the management of the switch.	Remote access management	<ul style="list-style-type: none"> • Connecting a QFX3500 Device to a Network for Out-of-Band Management • Connecting a QFX Series Device to a Management Console • Configuring Console and Auxiliary Port Properties on page 15

Table 89: Device and Network Management Features on the QFX Series, OCX Series, and EX4600 (*continued*)

Feature	Typical Uses	Documentation
SNMP Configuration Management MIB—Provides notification for configuration changes in the form of SNMP traps. Each trap contains the time at which the configuration change was committed, the name of the user who made the change, and the method by which the change was made. A history of the last 32 configuration changes is kept in <code>jnxCmChgEventTable</code> .	Configuration management	<i>SNMP MIBs and Traps Reference</i>
SNMP MIBs and traps—Enable the monitoring of network devices from a central location. Use SNMP requests such as get and walk to monitor and view system activity. The QFX3500 switch supports SNMP Version 1 (v1), v2, and v3, and both standard and Juniper Networks enterprise-specific MIBs and traps.	Fault management	<ul style="list-style-type: none"> • <i>SNMP MIBs and Traps Reference</i> • <i>Understanding the Implementation of SNMP</i>
System log messages—Log details of system and user events, including errors. You can specify the severity and type of system log messages you wish to view or save, and configure the output to be sent to local or remote hosts.	<ul style="list-style-type: none"> • Fault management • User access management 	<ul style="list-style-type: none"> • <i>Junos OS System Log Messages Reference</i> • Overview of Junos OS System Log Messages on page 1528 • Overview of Single-Chassis System Logging Configuration on page 1528

Junos Space

- [Understanding Junos Space Support on page 1405](#)

Understanding Junos Space Support

The Juniper Networks Junos Space application, running on a JA1500 appliance or a Junos Space Virtual Appliance, is a comprehensive platform for building and deploying applications for collaboration, productivity, and network infrastructure and operations management. Junos Space provides a runtime environment implemented as a fabric of virtual and physical appliances.

The Junos Space Network Management Platform software comprises various applications for network management and configuration, including:

- Junos Space Administration—Provides management of Junos Space fabric, databases, licenses, applications, authentication servers, tags, permission labels, DMI schemas, and troubleshooting.
- Network Director—Provides unified management of supported Juniper Networks devices in your network. By providing full network life cycle management, Network Director simplifies the discovery, configuration, visualization, monitoring, and administration of large networks.
- Service Automation—Provides an end-to-end solution designed to streamline operations and enable proactive network management for Junos OS devices. The solution consists of Advanced Insight Scripts (AI-Scripts), Junos Space Service Now and Service Insight applications, and Juniper Support Systems (JSS).



NOTE: Do not install Junos Space and AI-Scripts on the control plane network EX4200 switches or EX4200 Virtual Chassis in a QFX3000 QFabric system

Before you can use Junos Space Network Director to manage the QFX Series device, you must ensure that the configuration on the device meets the requirements for all managed devices. For example:

- The device configuration has a static management IP address that is reachable from the Junos Space server.
- There is a user with full administrative privileges for Junos Space administration.
- SNMP is enabled (only if you plan on using SNMP as part of the device discovery).
- In Junos Space, set up a default device management interface (DMI) schema for the QFX Series device.

For more information about Network Director requirements, see the *Network Director Quick Start Guide* at:

http://www.juniper.net/techpubs/en_US/network-director1.5/information-products/pathway-pages/index.html

For more information about Junos Space, go to:

http://www.juniper.net/techpubs/en_US/release-independent/junos-space/index.html

**Related
Documentation**

- [Configuring SNMP on page 1499](#)
- [Configuring SSH Service for Remote Access to the Router or Switch on page 1412](#)

PART 20

Configuration

- [Configuring Network Management on page 1409](#)
- [Configuring Automation on page 1419](#)
- [Configuring sFlow Technology on page 1449](#)
- [Configuring SNMP on page 1459](#)
- [Configuring System Logging on page 1527](#)

Configuring Network Management

- [Understanding Tracing and Logging Operations on page 1409](#)
- [Configuring Console and Auxiliary Port Properties on page 1411](#)
- [Configuring SSH Service for Remote Access to the Router or Switch on page 1412](#)
- [Configuring Telnet Service for Remote Access to a Switch on page 1414](#)
- [Pinging Hosts on page 1414](#)
- [Monitoring Traffic Through the Router or Switch on page 1415](#)

Understanding Tracing and Logging Operations

Tracing and logging operations enable you to track events that occur in the switch—both normal operations and error conditions—and to track the packets that are generated by or passed through the switch. The results of tracing and logging operations are placed in files in the `/var/log` directory on the switch.

The Junos OS supports remote tracing for the following processes:

- **chassisd**—Chassis-control process
- **eventd**—Event-processing process
- **cosd**—Class-of-service process

You configure remote tracing by using the **tracing** statement at the **[edit system]** hierarchy level.



NOTE: The **tracing** statement is not supported on the QFX3000 QFabric system.

If you enabled remote tracing but wish to disable it for specific processes on the switch, use the **no-remote-trace** statement at the **[edit process-name traceoptions]** hierarchy level. This feature does not alter local tracing functionality in any way, and logging files are stored on the switch.

Logging operations use a system logging mechanism similar to the UNIX **syslogd** utility to record systemwide, high-level operations, such as interfaces going up or down and users logging in to or out of the switch. You configure these operations by using the **syslog**

statement at the **[edit system]** hierarchy level and by using the **options** statement at the **[edit ethernet-switching-options]** hierarchy level.

Tracing operations record more detailed information about the operations of the switch, including packet forwarding and routing information. To configure tracing operations, use the **traceoptions** statement.



NOTE: The **traceoptions** statement is not supported on the QFX3000 QFabric system.

You can define tracing operations in different portions of the switch configuration:

- **SNMP agent activity tracing operations**—Define tracing of the activities of SNMP agents on the switch. You configure SNMP agent activity tracing operations at the **[edit snmp]** hierarchy level.
- **Global switching tracing operations**—Define tracing for all switching operations. You configure global switching tracing operations at the **[edit ethernet-switching-options]** hierarchy level of the configuration.
- **Protocol-specific tracing operations**—Define tracing for a specific routing protocol. You configure protocol-specific tracing operations in the **[edit protocols]** hierarchy when configuring the individual routing protocol. Protocol-specific tracing operations override any equivalent operations that you specify in the global **traceoptions** statement. If there are no equivalent operations, they supplement the global tracing options. If you do not specify any protocol-specific tracing, the routing protocol inherits all the global tracing operations.
- **Tracing operations within individual routing protocol entities**—Some protocols allow you to define more granular tracing operations. For example, in Border Gateway Protocol (BGP), you can configure peer-specific tracing operations. These operations override any equivalent BGP-wide operations or, if there are no equivalents, supplement them. If you do not specify any peer-specific tracing operations, the peers inherit, first, all the BGP-wide tracing operations and, second, the global tracing operations.
- **Interface tracing operations**—Define tracing for individual interfaces and for the interface process itself. You define interface tracing operations at the **[edit interfaces]** hierarchy level of the configuration.
- **Remote tracing**—To enable system-wide remote tracing, configure the **destination-override syslog host** statement at the **[edit system tracing]** hierarchy level. This specifies the remote host running the system log process (syslogd), which collects the traces. Traces are written to files on the remote host in accordance with the syslogd configuration in **/etc/syslog.conf**. By default, remote tracing is not configured.

To override the system-wide remote tracing configuration for a particular process, include the **no-remote-trace** statement at the **[edit process-name traceoptions]** hierarchy. When **no-remote-trace** is enabled, the process does local tracing.

To collect traces, use the **local0** facility as the selector in the **/etc/syslog.conf** file on the remote host. To separate traces from various processes into different files, include the process name or trace-file name (if it is specified at the **[edit process-name**

traceoptions file] hierarchy level) in the Program field in the `/etc/syslog.conf` file. If your system log server supports parsing hostname and program name, then you can separate traces from the various processes.



NOTE: During a commit check, warnings about the **traceoptions** configuration (for example, mismatch in trace file sizes or number of trace files) are not displayed on the console. However, these warnings are logged in the system log messages when the new configuration is committed.

Related Documentation

- [Overview of Junos OS System Log Messages on page 1528](#)

Configuring Console and Auxiliary Port Properties

The console port and auxiliary port on a switch provide out-of-band remote access to the switch. You can configure the console and auxiliary ports so that an external data terminal may be connected to the switch. The console port is enabled by default. The console port speed is 9600 baud, except on OCX Series devices, on which it is 115200 baud. The auxiliary port is disabled by default.

By default, terminal connections to the console and auxiliary ports are secure. When you configure the console and auxiliary ports as insecure, root logins are not allowed to establish terminal connections, and superusers and anyone with a user identifier (UID) of 0 are not allowed to establish terminal connections in multiuser mode.

To configure the console and auxiliary port properties on the switch:

1. To specify that the console port session should terminate if the connection to the data carrier is lost:

```
[edit system ports]
user@switch# set console log-out-on-disconnect
```

2. To specify the auxiliary port terminal type:

```
[edit system ports]
user@switch# set auxiliary type (ansi | small-xterm | vt100 | xterm)
```

For example, to specify the auxiliary port terminal type of **xterm** with a display of 80 columns by 65 rows:

```
[edit system ports]
user@switch# set auxiliary type xterm
```

3. To check the configuration:

```
[edit system ports]
user@switch# show
console log-out-on-disconnect;
auxiliary type xterm;
```

- Related Documentation
- [auxiliary on page 218](#)
 - [console \(Physical Port\) on page 221](#)
 - [ports on page 234](#)

Configuring SSH Service for Remote Access to the Router or Switch

To configure the router or switch to accept SSH as an access service, include the **ssh** statement at the **[edit system services]** hierarchy level:

```
[edit system services]
ssh {
  ciphers [ cipher-1 cipher-2 cipher-3 ... ]
  client-alive-count-max number;
  client-alive-interval seconds;
  connection-limit limit;
  hostkey-algorithm <algorithm | no-algorithm>;
  key-exchange algorithm;
  macs algorithm;
  max-sessions-per-connection number;
  no-passwords;
  no-tcp-forwarding;
  protocol-version [v1 v2];
  rate-limit limit;
  root-login <allow | deny | deny-password>;
}
```

By default, the router or switch supports a limited number of simultaneous SSH sessions and connection attempts per minute. Use the following statements to change the defaults:

- **connection-limit *limit***—Maximum number of simultaneous connections per protocol (IPv4 and IPv6). The range is a value from 1 through 250. The default is 75. When you configure a connection limit, the limit is applicable to the number of SSH sessions per protocol (IPv4 and IPv6). For example, a connection limit of 10 allows 10 IPv6 SSH sessions and 10 IPv4 SSH sessions.
- **max-sessions-per-connection *number***—Include this statement to specify the maximum number of SSH sessions allowed per single SSH connection. This allows you to limit the number of cloned sessions tunneled within a single SSH connection. The default value is 10.
- **rate-limit *limit***—Maximum number of connection attempts accepted per minute (a value from 1 through 250). The default is 150. When you configure a rate limit, the limit is applicable to the number of connection attempts per protocol (IPv4 and IPv6). For example, a rate limit of 10 allows 10 IPv6 SSH session connection attempts per minute and 10 IPv4 SSH session connection attempts per minute.

By default, a user can create an SSH tunnel over a CLI session to a router running Junos OS via SSH. This type of tunnel could be used to forward TCP traffic, bypassing any firewall filters or ACLs, allowing access to resources beyond the router. Use the **no-tcp-forwarding** option to prevent a user from creating an SSH tunnel to a router via SSH.

For information about other configuration settings, see the following topics:

- [Configuring the Root Login Through SSH on page 1413](#)
- [Configuring the SSH Protocol Version on page 1413](#)
- [Configuring the Client Alive Mechanism on page 1413](#)

Configuring the Root Login Through SSH

By default, users are allowed to log in to the router or switch as **root** through SSH. To control user access through SSH, include the **root-login** statement at the **[edit system services ssh]** hierarchy level:

```
[edit system services ssh]  
root-login (allow | deny | deny-password);
```

allow—Allows users to log in to the router or switch as root through SSH. The default is **allow**.

deny—Disables users from logging in to the router or switch as root through SSH.

deny-password—Allows users to log in to the router or switch as root through SSH when the authentication method (for example, RSA) does not require a password.

Configuring the SSH Protocol Version

By default, only version 2 of the SSH protocol is enabled. To enable version 1, you must explicitly configure it.

To configure the router or switch to use version 1 and 2 of the SSH protocol, include the **protocol-version** statement and specify **v1** and **v2** at the **[edit system services ssh]** hierarchy level:

```
[edit system services ssh]  
protocol-version [ v1 v2 ];
```

To configure the router or switch to use only version 1 of the SSH protocol, include the **protocol-version** statement and specify **v1** at the **[edit system services ssh]** hierarchy level:

```
[edit system services ssh]  
protocol-version [ v1 ];
```

For J Series Services Routers, the export license software supports SSH version 1 only.

Configuring the Client Alive Mechanism

The client alive mechanism is valuable when the client or server depends on knowing when a connection has become inactive. It differs from the standard keepalive mechanism because the client alive messages are sent through the encrypted channel. The client alive mechanism is not enabled at default. To enable it, configure the **client-alive-count-max** and the **client-alive-interval**. This option applies to SSH protocol version 2 only.

In the following example, unresponsive SSH clients will be disconnected after approximately 100 seconds (20 x 5).

```
[edit system services ssh]
client-alive-count-max 5;
client-alive-interval 20;
```

Configuring Telnet Service for Remote Access to a Switch

Telnet provides unencrypted access to network devices. Configuring Telnet service for a switch enables in-band remote access to the switch.

By default, the switch supports a limited number of simultaneous Telnet sessions and connection attempts per minute. Optionally, you can change the default Telnet settings by configuring the connection limit and rate limit at the **[edit system services telnet]** hierarchy level.

The connection limit is the maximum number of simultaneous connections per protocol (IPv4). The range is from 1 through 250. The default is 75.

The rate limit is the maximum number of connection attempts accepted per minute per protocol. The range is from 1 through 250. The default is 150.

To configure Telnet service:

1. To specify the connection limit:

```
[edit system services]
user@switch# set telnet connection-limit connection-limit
```

2. To specify the rate limit:

```
[edit system services]
user@switch# set telnet rate-limit rate-limit
```

3. Check that the Telnet connection limit and rate limit show the values you specified:

```
[edit system services]
user@switch# show
telnet {
  connection-limit 50;
  rate-limit 100;
}
```

Related Documentation

- *Understanding Telnet on the QFabric System*
- [Example: Limiting the Number of Login Attempts for SSH and Telnet Sessions to Prevent Unauthorized Access on page 1835](#)

Pinging Hosts

Purpose Use the CLI **ping** command to verify that a host can be reached over the network. This command is useful for diagnosing host and network connectivity problems. The switch sends a series of Internet Control Message Protocol (ICMP) echo (ping) requests to a specified host and receives ICMP echo responses.

Action To use the **ping** command to send four requests (ping count) to host3:
ping *host* *count* *number*

Sample Output

```
ping host3 count 4
user@switch> ping host3 count 4
PING host3.site.net (176.26.232.111): 56 data bytes
64 bytes from 176.26.232.111: icmp_seq=0 ttl=122 time=0.661 ms
64 bytes from 176.26.232.111: icmp_seq=1 ttl=122 time=0.619 ms
64 bytes from 176.26.232.111: icmp_seq=2 ttl=122 time=0.621 ms
64 bytes from 176.26.232.111: icmp_seq=3 ttl=122 time=0.634 ms

--- host3.site.net ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.619/0.634/0.661/0.017 ms
```

- Meaning**
- The **ping** results show the following information:
 - Size of the ping response packet (in bytes).
 - IP address of the host from which the response was sent.
 - Sequence number of the ping response packet. You can use this value to match the ping response to the corresponding ping request.
 - Time-to-live (ttl) hop-count value of the ping response packet.
 - Total time between the sending of the ping request packet and the receiving of the ping response packet, in milliseconds. This value is also called round-trip time.
 - Number of ping requests (probes) sent to the host.
 - Number of ping responses received from the host.
 - Packet loss percentage.
 - Round-trip time statistics: minimum, average, maximum, and standard deviation of the round-trip time.

- Related Documentation**
- [Troubleshooting Overview on page 1816](#)
 - [Understanding Troubleshooting Resources on page 1814](#)

Monitoring Traffic Through the Router or Switch

To help with the diagnosis of a problem, display real-time statistics about the traffic passing through physical interfaces on the router or switch.

To display real-time statistics about physical interfaces, perform these tasks:

1. [Displaying Real-Time Statistics About All Interfaces on the Router or Switch on page 1416](#)
2. [Displaying Real-Time Statistics About an Interface on the Router or Switch on page 1416](#)

Displaying Real-Time Statistics About All Interfaces on the Router or Switch

Purpose Display real-time statistics about traffic passing through all interfaces on the router or switch.

Action To display real-time statistics about traffic passing through all interfaces on the router or switch:

```
user@host> monitor interface traffic
```

Sample Output

```
user@host> monitor interface traffic
host name          Seconds: 15          Time: 12:31:09
Interface  Link  Input packets  (pps)  Output packets  (pps)
so-1/0/0    Down      0          (0)      0          (0)
so-1/1/0    Down      0          (0)      0          (0)
so-1/1/1    Down      0          (0)      0          (0)
so-1/1/2    Down      0          (0)      0          (0)
so-1/1/3    Down      0          (0)      0          (0)
t3-1/2/0    Down      0          (0)      0          (0)
t3-1/2/1    Down      0          (0)      0          (0)
t3-1/2/2    Down      0          (0)      0          (0)
t3-1/2/3    Down      0          (0)      0          (0)
so-2/0/0     Up    211035      (1)    36778      (0)
so-2/0/1     Up    192753      (1)    36782      (0)
so-2/0/2     Up    211020      (1)    36779      (0)
so-2/0/3     Up    211029      (1)    36776      (0)
so-2/1/0     Up    189378      (1)    36349      (0)
so-2/1/1    Down      0          (0)    18747      (0)
so-2/1/2    Down      0          (0)    16078      (0)
so-2/1/3     Up      0          (0)    80338      (0)
at-2/3/0     Up      0          (0)      0          (0)
at-2/3/1    Down      0          (0)      0          (0)
Bytes=b, Clear=c, Delta=d, Packets=p, Quit=q or ESC, Rate=r, Up=^U, Down=^D
```

Meaning The sample output displays traffic data for active interfaces and the amount that each field has changed since the command started or since the counters were cleared by using the **C** key. In this example, the **monitor interface** command has been running for 15 seconds since the command was issued or since the counters last returned to zero.

Displaying Real-Time Statistics About an Interface on the Router or Switch

Purpose Display real-time statistics about traffic passing through an interface on the router or switch.

Action To display traffic passing through an interface on the router or switch, use the following Junos OS CLI operational mode command:

```
user@host> monitor interface interface-name
```

Sample Output

```
user@host> monitor interface so-0/0/1
Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'
R1
```



```

Interface: so-0/0/1, Enabled, Link is Up
Encapsulation: PPP, Keepalives, Speed: OC3 Traffic statistics:
  Input bytes:          5856541 (88 bps)
  Output bytes:         6271468 (96 bps)
  Input packets:        157629 (0 pps)
  Output packets:       157024 (0 pps)
Encapsulation statistics:
  Input keepalives:     42353
  Output keepalives:    42320
  LCP state: Opened
Error statistics:
  Input errors:         0
  Input drops:          0
  Input framing errors: 0
  Input runs:           0
  Input giants:         0
  Policed discards:     0
  L3 incompletes:       0
  L2 channel errors:    0
  L2 mismatch timeouts: 0
  Carrier transitions:  1
  Output errors:        0
  Output drops:         0
  Aged packets:         0
Active alarms : None
Active defects: None
SONET error counts/seconds:
  LOS count             1
  LOF count             1
  SEF count             1
  ES-S                  77
  SES-S                 77
SONET statistics:
  BIP-B1                0
  BIP-B2                0
  REI-L                 0
  BIP-B3                0
  REI-P                 0
Received SONET overhead: F1      : 0x00 J0      : 0xZ

```

Meaning The sample output shows the input and output packets for a particular SONET interface (**so-0/0/1**). The information can include common interface failures, such as SONET/SDH and T3 alarms, loopbacks detected, and increases in framing errors. For more information, see *Checklist for Tracking Error Conditions*.

To control the output of the command while it is running, use the keys shown in [Table 90](#).

Table 90: Output Control Keys for the monitor interface Command

Action	Key
Display information about the next interface. The monitor interface command scrolls through the physical or logical interfaces in the same order that they are displayed by the show interfaces terse command.	N
Display information about a different interface. The command prompts you for the name of a specific interface.	I

Table 90: Output Control Keys for the monitor interface Command (*continued*)

Action	Key
Freeze the display, halting the display of updated statistics.	F
Thaw the display, resuming the display of updated statistics.	T
Clear (zero) the current delta counters since monitor interface was started. It does not clear the accumulative counter.	C
Stop the monitor interface command.	Q

See the [CLI Explorer](#) for details on using match conditions with the **monitor traffic** command.

CHAPTER 46

Configuring Automation

- [Overview of Junos Automation Enhancements on page 1419](#)
- [Overview of Python with Junos Automation Enhancements on page 1420](#)
- [Understanding Automation Scripts Support on page 1423](#)
- [How Commit Scripts Work on page 1424](#)
- [Avoiding Potential Conflicts When Using Multiple Commit Scripts on page 1429](#)
- [Overview of Generating Persistent or Transient Configuration Changes on page 1430](#)
- [Required Boilerplate for Commit Scripts on page 1434](#)
- [How Op Scripts Work on page 1436](#)
- [Required Boilerplate for Op Scripts on page 1436](#)
- [Installing Junos OS Software with Junos Automation Enhancements on page 1438](#)
- [Invoking the Python Interpreter on page 1442](#)
- [Controlling the Execution of Commit Scripts on page 1443](#)
- [Displaying Commit Script Output on page 1445](#)

Overview of Junos Automation Enhancements

The Junos Automation Enhancements are designed to support the increasing needs of large data centers for more automation and programmability.

- [Features of the Junos Automation Enhancements on page 1419](#)

Features of the Junos Automation Enhancements

To use the Junos Automation Enhancements, you must install the **jinstall-qfx-5-flex-x.tgz** software bundle. This software bundle is identical to the other software bundle except that Veriexec is disabled, which enables you to run unsigned programs, such as programs that you develop with Python, Chef, and Puppet. The Junos Automation Enhancements include the following features:

- The factory default configuration is a Layer 3 configuration. (The standard default factory configuration on some devices series is Layer 2.)
- Safeguards ensure that you cannot overwrite essential Junos OS files, including system log notifications.

- Zero Touch Provisioning (ZTP) allows you to provision new switches in your network automatically, without manual intervention. See [“Understanding Zero Touch Provisioning” on page 197](#).
- The installation automatically sets up and reserves a 1-gigabit user partition on your system. You can use this partition to store your binaries and additional packages.
- The user partition is not overwritten when you upgrade or downgrade the software to a OS image that does not contain the automation enhancements.



NOTE: If you make changes to the user partition while performing a unified in-service software upgrade (unified ISSU), the changes might be lost.

- The Python interpreter is included by default.
 - You can invoke Python directly from the shell. See [“Invoking the Python Interpreter” on page 1442](#).
 - Starting with Junos OS Release 14.1X53-D10, three Open Source Python modules are pre-installed in the `jinstall-qfx-5-flex-x.tgz` software bundle. See [“Overview of Python with Junos Automation Enhancements” on page 1420](#) for details.
- Chef for Junos OS and Puppet for Junos OS automation tools for provisioning and managing computer networking and storage resources are included.
 - For further information on Chef, see [Chef for Junos Getting Started Guide](#).
 - For further information on Puppet, see [Puppet for Junos OS Documentation](#).



NOTE: For full compatibility, you must use only Chef for Junos OS and Puppet for Junos OS rather than the standard FreeBSD versions of Chef and Puppet software.



CAUTION: Download additional third party packages at your own risk.

Related Documentation

- [Installing Junos OS Software with Junos Automation Enhancements on page 1438](#)
- [Invoking the Python Interpreter on page 1442](#)
- [QFX5100 Switch with Automation Enhancements Frequently Asked Questions](#)

Overview of Python with Junos Automation Enhancements

Python is a programming language that lets you work more quickly and integrate your systems more effectively. The Python interpreter is included within the Junos operating system (Junos OS) `jinstall-qfx-5-flex-x.tgz` software bundle.

Python is also suitable as an extension language for customizable applications.

Starting with Junos OS Release 14.1X53-D10, these Open Source Python modules are pre-installed in the `jinstall-qfx-5-flex-x.tgz` software bundle:

- **ncclient**—Facilitates client scripting and application development through the NETCONF protocol. See <http://ncclient.grnet.gr/0.3.2/> for documentation of some of the external APIs of the ncclient Python module. At the bottom of this list, see examples of usage of some of these APIs with sample scripts.
- **lxml**—Combines the speed and XML feature completeness of the C libraries libxml2 and libxslt with the simplicity of a native Python API. See <http://lxml.de/tutorial.html/> for documentation of some of the external APIs of the lxml Python module.
- **jinja2**—Serves as a fast, secure, designer-friendly templating language. See <http://jinja.pocoo.org/docs/api/> for documentation of some of the external APIs of the jinja2 Python module.

Example usage of some of the APIs of the ncclient Python module follows:

*** Example of "connect" and "command" API:**

```
from ncclient import manager

def connect(host, port, user, password):
    conn = manager.connect(host=host,
                           port=port,
                           username=user,
                           password=password,
                           timeout=10,
                           device_params = {'name': 'junos'},
                           hostkey_verify=False)

    print 'show version'
    print '*' * 30
    result = conn.command('show version', format='text')
    print result.xpath('output')[0].text

if __name__ == '__main__':
    connect('router', '22', 'netconf', 'juniper!')
```

*** Example of "compare_configuration" API:**

```
from ncclient import manager
from ncclient.xml_ import *

import time

def connect(host, port, user, password, source):
    conn = manager.connect(host=host,
                           port=port,
                           username=user,
                           password=password,
                           timeout=10,
                           device_params = {'name': 'junos'},
                           hostkey_verify=False)

    compare_config = conn.compare_configuration(rollback=3)
    print compare_config.tostring

if __name__ == '__main__':
    connect('router', 830, 'netconf', 'juniper!', 'candidate')
```

* Example of "lock", "load_configuration", "validate", "commit", "discard_changes", "unlock" APIs:

```
from ncclient import manager
from ncclient.xml_ import *

import time

def connect(host, port, user, password, source):
    conn = manager.connect(host=host,
                           port=port,
                           username=user,
                           password=password,
                           timeout=10,
                           device_params = {'name': 'junos'},
                           hostkey_verify=False)

    print 'locking configuration'
    lock = conn.lock()

    # build configuration element
    config = new_ele('system')
    sub_ele(config, 'host-name').text = 'foo'
    sub_ele(config, 'domain-name').text = 'bar'

    send_config = conn.load_configuration(config=config)
    print send_config.tostring

    check_config = conn.validate()
    print check_config.tostring

    compare_config = conn.compare_configuration()
    print compare_config.tostring

    print 'commit confirmed 300'
    #commit_config = conn.commit(confirmed=True, timeout='300')
    commit_config = conn.commit()
    print commit_config.tostring

    print 'sleeping for 5 sec...'
    time.sleep(5)

    discard_changes = conn.discard_changes()
    print discard_changes.tostring

    print 'unlocking configuration'
    unlock = conn.unlock()
    print unlock.tostring

if __name__ == '__main__':
    connect('router', 830, 'netconf', 'juniper!', 'candidate')
```



NOTE: For information on using Python, refer to your Python documentation.

**Related
Documentation**

- [Installing Junos OS Software with Junos Automation Enhancements on page 1438](#)
- [Invoking the Python Interpreter on page 1442](#)
- [QFX5100 Switch with Automation Enhancements Frequently Asked Questions](#)

Understanding Automation Scripts Support

This document describes the support for the Junos OS automation scripts on the QFabric system Director devices.

Junos OS automation consists of a suite of tools used to automate operational and configuration tasks on network devices running Junos OS. The automation tools, which leverage the native XML capabilities of the Junos OS, include commit scripts, operation (op) scripts, event policies and event scripts, and macros.



NOTE: Event policies and event scripts are not supported on the QFabric system at this time.

The QFabric system supports Junos OS automation scripts that are written in Stylesheet Language Alternative Syntax (SLAX) version 1.0.

Commit scripts automate the commit process and enforce custom configuration rules. You can use commit scripts to generate specific errors and warnings, and customize configurations and configuration templates. When a candidate configuration is committed, it is inspected by each active commit script. If a configuration violates your custom rules and the scripts generate an error, the commit fails. If the commit is successful, any configuration changes (both transient and permanent) are incorporated into the active configuration before it is passed to the Director software, which distributes the configuration to all applicable QFabric system components, including Node devices and Node servers.

Op scripts automate operational and troubleshooting tasks. Op scripts can be executed manually from the Junos OS CLI or NETCONF XML management protocol, or they can be called from another script.

The QFabric system supports the following automation script features:

- Commit scripts and op scripts are supported.
- Scripts written in SLAX version 1 are supported.
- Scripts are configured and deployed from the Director group. Since there is more than one Director device in a Director group, scripts must be deployed by each Director device or deployed in the shared media space.
- Scripts are stored in the shared media at this location:
`/pbdata/mgd_shared/partition-ip/var/db/scripts`. Under this directory, commit scripts are stored in the **commit** subdirectory, and op scripts are stored in the **op** subdirectory.
- Scripts are not stored in flash memory.

Related Documentation

- [How Commit Scripts Work on page 1424](#)
- [How Op Scripts Work on page 1436](#)
- [Required Boilerplate for Commit Scripts on page 1434](#)

- [Required Boilerplate for Op Scripts on page 1436](#)
- [Controlling the Execution of Commit Scripts on page 1443](#)

How Commit Scripts Work


You enable commit scripts by listing the names of one or more commit script files at the **[edit system scripts commit]** hierarchy level. These scripts contain instructions that enforce custom configuration rules. Commit scripts are invoked during the commit process before the standard Junos OS validity checks are performed.

When you perform a commit operation, Junos OS executes each script in turn, passing the information in the candidate configuration to the scripts. The script inspects the configuration, performs the necessary tests and validations, and generates a set of instructions for performing certain actions. These actions include generating error, warning, and system log messages. If errors are generated, the commit operation fails and the candidate configuration remains unchanged. This is the same behavior that occurs with standard commit errors.

Commit scripts can also generate changes to the system configuration. Because the changes are loaded before the standard validation checks are performed, they are validated for correct syntax, just like statements already present in the configuration before the script is applied. If the syntax is correct, the configuration is activated and becomes the active, operational device configuration.

Figure 12 shows the flow of commit script input and output.

Figure 12: Commit Script Input and Output



Commit scripts cannot make configuration changes to protected statements or within protected hierarchies. If a commit script attempts to modify or delete a protected statement or hierarchy, Junos OS issues a warning that the change cannot be made. Failure to modify a protected configuration element does not halt the commit script or the commit process.

The following sections discuss several important concepts related to the commit script input and output:

- [Commit Script Input on page 1425](#)
- [Commit Script Output on page 1425](#)
- [Commit Scripts and the Junos OS Commit Model on page 1426](#)

Commit Script Input

The input for a commit script is the postinheritance candidate configuration in Junos XML API format. The term *postinheritance* means that all configuration group values have been inherited by their targets in the candidate configuration and the inactive portions of the configuration have been removed. For more information about configuration groups, see the *CLI User Guide*.

When you issue the **commit** command, Junos OS automatically generates the candidate configuration in XML format and reads it into the management (mgd) process, at which time the input is evaluated by any commit scripts.

To display the XML format of the postinheritance configuration, issue the **show | display commit-scripts view** command:

```
[edit]
user@host# show | display commit-scripts view
```

To display all configuration groups data, including script-generated changes to the groups, issue the **show groups | display commit-scripts** command:

```
[edit]
user@host# show groups | display commit-scripts
```

To save the commit script input to a file, add the **save** command to the command line:

```
[edit]
user@host# show | display commit-scripts view | save filename.xml
```

By default, the file is placed in your home directory on the switch, router, or security device.

Commit Script Output

To specify the desired commit script output—including warning, error, and system log messages, persistent changes, and transient changes—the script can contain tags that appear in any order, in any number. The tags for specifying output are as follows:

- **<xnm:warning>**—Generates a warning message
- **<xnm:error>**—Generates an error message.
- **<syslog><message>**—Generates a system log message.
- **<change>**—Generates a persistent change to the configuration.
- **<transient-change>**—Generates a transient change to the configuration.
- **<xsl:call-template name="jcs:emit-change">**
 <xsl:with-param name="content">—Generates a persistent change relative to the current context node as defined by an XPath expression.
- **<xsl:call-template name="jcs:emit-change">**
 <xsl:with-param name="tag" select="'transient-change'"/>
 <xsl:with-param name="content">—Generates a transient change relative to the current context node as defined by an XPath expression.
- **<xsl:call-template name="jcs:emit-change">**

```
<xsl:with-param name="message">
```

```
<xsl:text>
```

—Generates a warning message in conjunction with a configuration change. You can use this set of tags to generate a notification that the configuration has been changed.

Junos OS processes this output and performs the appropriate actions. Errors and warnings are passed back to the Junos OS CLI or to a Junos XML protocol client application. The presence of an error automatically causes the commit operation to fail. Persistent and transient changes are loaded into the appropriate configuration database.

To test the output of error, warning, and system log messages from commit scripts, issue the **commit check | display xml** command:

```
[edit]
user@host# commit check | display xml
```

To display a detailed trace of commit script processing, issue the **commit check | display detail** command:

```
[edit]
user@host# commit check | display detail
```



NOTE: System log messages do not appear in the trace output, so you cannot use the commit check operation to test script-generated system log messages. Furthermore, system log messages are written to the system log during a commit operation, but not during a commit check operation.

**Related
Documentation**

- *Example: Protecting the Junos OS Configuration from Modification or Deletion*
- *jcs:emit-change Template*

Commit Scripts and the Junos OS Commit Model

Junos OS uses a commit model to update the device's configuration. This model allows you to make a series of changes to a candidate configuration without affecting the operation of the device. When the changes are complete, you can commit the configuration. The commit operation saves the candidate configuration changes into the current configuration.


When you commit a set of changes in the candidate configuration, two methods are used to forward these changes to the current configuration:

- Standard commit model—Used when no commit scripts are active on the device.
- Commit script model—Incorporates commit scripts into the commit model.

Standard Commit Model

In the standard commit model, the management (mgd) process validates the candidate configuration based on standard Junos validation rules. If the configuration file is valid, it becomes the current active configuration. [Figure 13](#) and the accompanying discussion explain how the standard commit model works:

Figure 13: Standard Commit Model




In the standard commit model, the software performs the following steps:

1. When the candidate configuration is committed, it is copied to become the checkout configuration.
2. The mgd process validates the checkout configuration.
3. If no error occurs, the checkout configuration is copied as the current active configuration.

Commit Model with Commit Scripts

When commit scripts are added to the standard commit model, the process becomes more complex. The mgd process first passes an XML-formatted checkout configuration to a script driver, which handles the verification of the checkout configuration by the commit scripts. When verification is complete, the script driver returns an XML *action file* to the mgd process. The mgd process follows the instructions in the action file to update the candidate and checkout configurations, issue messages to the CLI, and write information to the system log as required. After processing the action file, the mgd process performs the standard Junos OS validation. Figure 14 and the accompanying discussion explain this process.

Figure 14: Commit Model with Commit Scripts Added



In the commit script model, Junos OS performs the following steps:

1. When the candidate configuration is committed, the mgd process sends the XML-formatted candidate configuration to the script driver.
2. Each enabled commit script is invoked against the candidate configuration, and each script can generate a set of actions for the mgd process to perform. The actions are collected in an XML action file.
3. The mgd process performs the following actions in response to **<error>**, **<warning>**, and **<syslog>** tag elements in the action file:
 - **<error>**—The mgd process halts the commit process (that is, the commit operation fails), returns an error message to the CLI or Junos XML protocol client, and takes no further action.
 - **<warning>**—The mgd process forwards the message to the CLI or the Junos XML protocol client.
 - **<syslog>**—The mgd process forwards the message to the system log process.
4. If the action file includes any **<change>** tag elements, the mgd process loads the requested changes into the candidate configuration.
5. The candidate configuration is copied to become the checkout configuration.
6. If the action file includes any **<transient-change>** tag elements, the mgd process loads the requested changes into the checkout configuration.
7. The mgd process validates the checkout configuration.
8. If there are no validation errors, the checkout configuration is copied to become the current active configuration.



NOTE: Commit scripts cannot make configuration changes to protected statements or within protected hierarchies. If a commit script attempts to modify or delete a protected statement or hierarchy, Junos OS issues a warning that the change cannot be made. Failure to modify a protected configuration element does not halt the commit script or the commit process.

Changes that are made to the candidate configuration during the commit operation are not evaluated by the custom rules during that commit operation. However, persistent changes are maintained in the candidate configuration and are evaluated by the custom rules during subsequent commit operations. For more information about how commit scripts change the candidate configuration, see [“Avoiding Potential Conflicts When Using Multiple Commit Scripts” on page 1429](#).

Transient changes are never evaluated by the custom rules in commit scripts, because they are made to the checkout configuration only after the commit scripts have evaluated the candidate configuration and the candidate is copied to become the checkout configuration. To remove a transient change from the configuration, remove, disable, or deactivate the commit script (as discussed in *Controlling Execution of Commit Scripts During Commit Operations*), or comment out the code that generates the transient change.

For more information about differences between persistent and transient changes, see [“Overview of Generating Persistent or Transient Configuration Changes Using Commit Scripts”](#) on page 1430.


Related Documentation

- [Avoiding Potential Conflicts When Using Multiple Commit Scripts](#) on page 1429

Avoiding Potential Conflicts When Using Multiple Commit Scripts

When you use multiple commit scripts, each script evaluates the original candidate configuration file. Changes made by one script are not evaluated by the other scripts. This means that conflicts between scripts might not be resolved when the scripts are first applied to the configuration. The commit scripts are executed in the order they are listed at the `[edit system scripts commit]` hierarchy level, as illustrated in [Figure 15](#).

Figure 15: Configuration Evaluation by Multiple Commit Scripts



As an example of a conflict between commit scripts, suppose that commit script **A.xsl** is created to ensure that the device uses the domain name server with IP address **192.168.0.255**. Later, the DNS server's address is changed to **192.168.255.255** and a second script, **B.xsl**, is added to check that the device uses the DNS server with that address. However, script **A.xsl** is not removed or disabled.

Because each commit script evaluates the original candidate configuration, the final result of executing both scripts **A.xsl** and **B.xsl** depends on which DNS server address is configured in the original candidate configuration. If the now outdated address of **192.168.0.255** is configured, script **B.xsl** changes it to **192.168.255.255**. However, if the correct address of **192.168.255.255** is configured, script **A.xsl** changes it to the incorrect value **192.168.0.255**.

As another example of a potential conflict between commit scripts, suppose that a commit script protects a hierarchy using the **protect** attribute. If a second commit script attempts to modify or delete the hierarchy or the statements within the hierarchy, Junos OS issues a warning during the commit process and prevents the configuration change.

Exercise care to ensure that you do not introduce conflicts between scripts like those described in the examples. As a method of checking for conflicts with persistent changes, you can issue two separate **commit** commands.

Related Documentation

- [How Commit Scripts Work on page 1424](#)

Overview of Generating Persistent or Transient Configuration Changes

Junos OS commit scripts enforce custom configuration rules. When a candidate configuration includes statements that you have decided must not be included in your configuration, or when the candidate configuration omits statements that you have decided are required, commit scripts can automatically change the configuration and thereby correct the problem.

- [Differences Between Persistent and Transient Changes on page 1430](#)
- [Interaction of Configuration Changes and Configuration Groups on page 1433](#)
- [Tag Elements and Templates for Generating Changes on page 1434](#)

Differences Between Persistent and Transient Changes

Configuration changes made by commit scripts can be *persistent* or *transient*.

A persistent change remains in the candidate configuration and affects routing operations until you explicitly delete it, even if you subsequently remove or disable the commit script that generated the change and reissue the **commit** command. In other words, removing the commit script does not cause a persistent change to be removed from the configuration.

A transient change, in contrast, is made in the *checkout configuration* but not in the candidate configuration. The checkout configuration is the configuration database that is inspected for standard Junos OS syntax just before it is copied to become the active configuration on the device. If you subsequently remove or disable the commit script that made the change and reissue the **commit** command, the change is no longer made to the checkout configuration and so does not affect the active configuration. In other words, removing the commit script effectively removes a transient change from the configuration.

A common use for transient changes is to eliminate the need to repeatedly configure and display well-known policies, thus allowing these policies to be enforced implicitly. For example, if MPLS must be enabled on every interface with an International Organization for Standardization (ISO) protocol enabled, the change can be transient, so that the repetitive or redundant configuration data need not be carried or displayed in the candidate configuration. Furthermore, transient changes allow you to write script instructions that apply the change only if a set of conditions is met.

Persistent and transient changes are loaded into the configuration in the same manner that the **load replace** configuration mode command loads an incoming configuration. When generating a persistent or transient change, adding the **replace="replace"** attribute

to a configuration element produces the same behavior as a **replace** tag in a **load replace** operation.

By default, Junos OS merges the incoming configuration and the candidate configuration. New statements and hierarchies are added, and conflicting statements are overridden. When generating a persistent or transient change, if you add the **replace="replace"** attribute to a configuration element, Junos OS replaces the existing configuration element with the incoming configuration element. If the **replace="replace"** attribute is added to a configuration element, but there is no existing element of the same name in the current configuration, the incoming configuration element is added into the configuration. Elements that do not have the **replace** attribute are merged into the configuration.

Persistent and transient changes are loaded before the standard Junos validation checks are performed. This means any configuration changes introduced by a commit script are validated for correct syntax. If the syntax is correct, the new configuration becomes the active, operational device configuration.

Protected elements in the configuration hierarchy cannot be modified or deleted by either a persistent or a transient change. If a commit script attempts to modify or delete a protected statement or hierarchy, Junos OS issues a warning that the change cannot be made, and proceeds with the commit.

Persistent and transient changes have several important differences, as described in [Table 91](#).

Table 91: Differences Between Persistent and Transient Changes

Persistent Changes	Transient Changes
<p>A persistent change is represented in a commit script by the <change> tag.</p> <p>Another way to represent a persistent change is with the content parameter inside a call to the jcs:emit-change template.</p> <p>The jcs:emit-change template is a helper template contained in the junos.xsl import file.</p>	<p>A transient change is represented in a commit script by the <transient-change> tag.</p> <p>Another way to represent a transient change is to use the content parameter and the tag transient parameter inside a call to the jcs:emit-change template.</p>
<p>You can use persistent changes to perform any Junos XML protocol operation, such as activate, deactivate, delete, insert (reorder), comment (annotate), and replace sections of the configuration.</p>	<p>Like persistent changes, you can use transient changes to perform any Junos XML protocol operation. However, some Junos XML protocol operations do not make sense to use with transient changes, such as generating comments and inactive settings.</p>

Table 91: Differences Between Persistent and Transient Changes (*continued*)

Persistent Changes	Transient Changes
Persistent changes are always loaded during the commit process if no errors are generated by any commit scripts or by the standard Junos OS validity check.	<p>For transient changes to be loaded, you must include the allow-transients statement at the [edit system scripts commit] hierarchy level. If you enable a commit script that generates transient changes and you do not include the allow-transients statement in the configuration, the CLI generates an error message and the commit operation fails.</p> <p>Like persistent changes, transient changes must pass the standard Junos OS validity check.</p> <p>You cannot use a commit script to generate the allow-transients statement at the [edit system scripts commit] hierarchy level. Rather, you must include this statement directly by using the CLI.</p>
<p>Persistent changes work like the load replace configuration mode command, and the change is added to the candidate configuration.</p> <p>When generating a persistent change, if you add the replace="replace" attribute to a configuration element, Junos OS replaces the existing element in the candidate configuration with the incoming configuration element. If there is no existing element of the same name in the candidate configuration, the incoming configuration element is added into the configuration. Elements that do not have the replace attribute are merged into the configuration.</p>	<p>Transient changes work like the load replace configuration mode command, and the change is added to the checkout configuration.</p> <p>When generating a transient change, if you add the replace="replace" attribute to a configuration element, Junos OS replaces the existing element in the checkout configuration with the incoming configuration element. If there is no existing element of the same name in the checkout configuration, the incoming configuration element is added into the configuration. Elements that do not have the replace attribute are merged into the configuration.</p> <p>Transient changes are not copied to the candidate configuration. For this reason, transient changes are not saved in the configuration if the associated commit script is deleted or deactivated.</p>
<p>After a persistent change is committed, the software treats it like a change you make by directly editing and committing the candidate configuration.</p> <p>After the persistent changes are copied to the candidate configuration, they are copied to the checkout configuration. If the changes pass the standard Junos OS validity checks, the changes are propagated to the switch, router, or security device components.</p>	<p>Each time a transient change is committed, the software updates the checkout configuration database. After the transient changes pass the standard Junos OS validity checks, the changes are propagated to the device components.</p>
<p>After committing a script that causes a persistent change to be generated, you can view the persistent change by issuing the show configuration mode command:</p> <pre>user@host# show</pre> <p>This command displays persistent changes only, not transient changes.</p>	<p>After committing a script that causes a transient change to be generated, you can view the transient change by issuing the show display commit-scripts configuration mode command:</p> <pre>user@host# show display commit-scripts</pre> <p>This command displays both persistent and transient changes.</p>

Table 91: Differences Between Persistent and Transient Changes (*continued*)

Persistent Changes	Transient Changes
<p>Persistent changes must conform to your custom configuration design rules as dictated by commit scripts.</p> <p>This does not become apparent until after a second commit operation because persistent changes are not evaluated by commit script rules on the current commit operation. The subsequent commit operation fails if the persistent changes do not conform to the rules imposed by the commit scripts configured during the first commit operation.</p>	<p>Transient changes are never tested by and do not need to conform to your custom rules. This is caused by the order of operations in the Junos OS commit model, which is explained in detail in “Commit Scripts and the Junos OS Commit Model” on page 1426.</p>
<p>A persistent change remains in the configuration even if you delete, disable, or deactivate the commit script instructions that generated the change.</p>	<p>If you delete, disable, or deactivate the commit script instructions that generate a transient change, the change is removed from the configuration after the next commit operation. In short, if the associated instructions or the entire commit script is removed, the transient change is also removed.</p>
<p>As with direct CLI configuration, you can remove a persistent change by rolling back to a previous configuration that did not include the change and issuing the commit command. However, if you do not disable or deactivate the associated commit script, and the problem that originally caused the change to be generated still exists, the change is automatically regenerated when you issue another commit command.</p>	<p>You cannot remove a transient change by rolling back to a previous configuration.</p>
<p>You can alter persistent changes directly by editing the configuration using the CLI.</p>	<p>You cannot directly alter or delete a transient change by using the Junos OS CLI, because the change is not in the candidate configuration.</p> <p>To alter the contents of a transient change, you must alter the statements in the commit script that generates the transient change.</p>

Interaction of Configuration Changes and Configuration Groups

Any configuration change you can make by directly editing the configuration using the Junos OS command-line interface (CLI) can also be generated by a commit script as a persistent or transient change. This includes values specified at a specific hierarchy level or in configuration groups. As with direct CLI configuration, values specified in the *target* override values inherited from a configuration group. The target is the statement to which you apply a configuration group by including the **apply-groups** statement.

If you define persistent or transient changes as belonging to a configuration group, the configuration groups are applied in the order you specify in the **apply-groups** statements, which you can include at any hierarchy level except the top level. You can also disable inheritance of a configuration group by including the **apply-groups-except** statement at any hierarchy level except the top level.



CAUTION: Each commit script inspects the postinheritance view of the configuration. If a candidate configuration contains a configuration group,

be careful when using a commit script to change the related target configuration, because doing so might alter the intended inheritance from the configuration group.

Also be careful when using a commit script to change a configuration group, because the configuration group might be generated by an application that performs a load replace operation on the group during each commit operation.

For more information about configuration groups, see the *CLI User Guide*.

Tag Elements and Templates for Generating Changes

To generate changes, you can use the `jcs:emit-change` template, which implicitly includes `<change>` and `<transient-change>` XML elements; or you can explicitly include `<change>` and `<transient-change>` XML elements. Using the `jcs:emit-change` template allows you to set the hierarchical context of the change once rather than multiple times.

The `<change>` and `<transient-change>` elements are similar to the `<load-configuration>` operation defined by the Junos XML management protocol. The possible contents of the `<change>` and `<transient-change>` elements are the same as the contents of the `<configuration>` tag element used in the Junos XML protocol operation `<load-configuration>`. For complete details about the `<load-configuration>` element, see the *Junos XML Management Protocol Developer Guide*.

Required Boilerplate for Commit Scripts

When you write commit scripts, you use Extensible Stylesheet Language Transformations (XSLT) or Stylesheet Language Alternative Syntax (SLAX) tools provided with Junos OS. These tools include basic boilerplate that you must include in all commit scripts, optional extension functions that accomplish scripting tasks more easily, and named templates that make commit scripts easier to read and write, which you import from a file called `junos.xsl`. For more information about the extension functions and templates, see *Understanding Extension Functions in the jcs and slax Namespaces* and *Named Templates in the jcs Namespace Overview*.

Commit scripts are based on Junos XML and Junos XML protocol tag elements. Like all XML elements, angle brackets enclose the name of a Junos XML or Junos XML protocol tag element in its opening and closing tags. This is an XML convention, and the brackets are a required part of the complete tag element name. They are not to be confused with the angle brackets used in the documentation to indicate optional parts of Junos OS CLI command strings.

You must include either XSLT or SLAX boilerplate as the starting point for all commit scripts that you create. The XSLT boilerplate follows:

XSLT Boilerplate for Commit Scripts

```
1 <?xml version="1.0" standalone="yes"?>
2 <xsl:stylesheet version="1.0"
3   xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
4   xmlns:junos="http://xml.juniper.net/junos/*/junos"
5   xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
```

```

6  xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
7  <xsl:import href="../../import/junos.xsl"/>

8  <xsl:template match="configuration">
    <!-- ... Insert your code here ... -->
9  </xsl:template>
10 </xsl:stylesheet>

```

Line 1 is the Extensible Markup Language (XML) processing instruction (PI). This PI specifies that the code is written in XML using version 1.0. The XML PI, if present, must be the first noncomment token in the script file.

```
1 <?xml version="1.0"?>
```

Lines 2 through 6 set the style sheet element and the associated namespaces. Line 2 sets the style sheet version as 1.0. Lines 3 through 6 list all the namespace mappings commonly used in commit scripts. Not all of these prefixes are used in this example, but it is not an error to list namespace mappings that are not referenced. Listing all namespace mappings prevents errors if the mappings are used in later versions of the script.

```

2 <xsl:stylesheet version="1.0"
3   xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
4   xmlns:junos="http://xml.juniper.net/junos/*/junos"
5   xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
6   xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">

```

Line 7 is an XSLT import statement. It loads the templates and variables from the file referenced as `../import/junos.xsl`, which ships as part of Junos OS. The `junos.xsl` file contains a set of named templates you can call in your scripts. These named templates are discussed in *Named Templates in the jcs Namespace Overview* and *Named Templates in the jcs Namespace Summary*.

```
7 <xsl:import href="../../import/junos.xsl"/>
```

Line 8 defines a template that matches the `<configuration>` element, which is the node selected by the `<xsl:template match="/">` template, contained in the `junos.xsl` import file. The `<xsl:template match="configuration">` element allows you to exclude the `/configuration/` root element from all XML Path Language (XPath) expressions in the script and begin XPath expressions with the top Junos OS hierarchy level. For more information, see *XPath Overview*.

```
8 <xsl:template match="configuration">
```

Add your code between Lines 8 and 9.

Line 9 closes the template.

```
9 </xsl:template>
```

Line 10 closes the style sheet and the commit script.

```
10 </xsl:stylesheet>
```

SLAX Boilerplate for Commit Scripts

The corresponding SLAX boilerplate is as follows:

```

version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";

```

```
import "../import/junos.xml";

match configuration {
  /*
   * Insert your code here
   */
}
```

How Op Scripts Work


Op scripts execute Junos OS operational commands and inspect the resulting output. After inspection, op scripts can automatically correct errors within the device running Junos OS based on this output.

You add op scripts to device operations by listing the filenames of one or more op script files within the **[edit system scripts op]** hierarchy level. These files must be added to the appropriate op script file directory. For more information about op script file directories, see *Storing Scripts in Flash Memory*. Once added to the device, op scripts are invoked from the command line, using the **op filename** command.

You can use op scripts to generate changes to the device configuration by including the **<load-configuration>** tag element. Because the changes are loaded before the standard validation checks are performed, they are validated for correct syntax, just like statements already present in the configuration before the script is applied. If the syntax is correct, the configuration is activated and becomes the active, operational device configuration.

Figure 16 shows a high-level view of the flow of op script input and output.

Figure 16: Op Script Input and Output



Required Boilerplate for Op Scripts

When you write operation (op) scripts, you use Extensible Stylesheet Language Transformations (XSLT) or Stylesheet Language Alternative Syntax (SLAX) tools provided with Junos OS. These tools include basic boilerplate that you must include in all op scripts, optional extension functions that accomplish scripting tasks more easily, and named templates that make scripts easier to read and write, which you import from a file called **junos.xml**. For more information about the extension functions and templates, see *Understanding Extension Functions in the jcs and slax Namespaces* and *Named Templates in the jcs Namespace Overview*.

Op scripts are based on Junos XML and Junos XML protocol tag elements. Like all XML elements, angle brackets enclose the name of a Junos XML or Junos XML protocol tag element in its opening and closing tags. This is an XML convention, and the brackets are

a required part of the complete tag element name. They are not to be confused with the angle brackets used in the documentation to indicate optional parts of Junos OS CLI command strings.

You must include either XSLT or SLAX boilerplate as the starting point for all op scripts that you create. The XSLT boilerplate follows:

XSLT Boilerplate for Op Scripts

```

1  <?xml version="1.0" standalone="yes"?>
2  <xsl:stylesheet version="1.0"
3    xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
4    xmlns:junos="http://xml.juniper.net/junos/*/junos"
5    xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
6    xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
7    <xsl:import href="../../import/junos.xml"/>

8    <xsl:template match="/">
9      <op-script-results>
10       <!-- ... insert your code here ... -->
11     </op-script-results>
12   </xsl:template>
13   <!-- ... insert additional template definitions here ... -->
14 </xsl:stylesheet>

```

Line 1 is the Extensible Markup Language (XML) processing instruction (PI), which marks this file as XML and specifies the version of XML as 1.0. The XML PI, if present, must be the first non-comment token in the script file.

```
1  <?xml version="1.0"?>
```

Line 2 opens the style sheet and specifies the XSLT version as 1.0.

```
2  <xsl:stylesheet version="1.0"
```

Lines 3 through 6 list all the namespace mappings commonly used in operation scripts. Not all of these prefixes are used in this example, but it is not an error to list namespace mappings that are not referenced. Listing all namespace mappings prevents errors if the mappings are used in later versions of the script.

```

3    xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
4    xmlns:junos="http://xml.juniper.net/junos/*/junos"
5    xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
6    xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">

```

Line 7 is an XSLT import statement. It loads the templates and variables from the file referenced as `../import/junos.xml`, which ships as part of Junos OS (in the file `/usr/libdata/cscript/import/junos.xml`). The `junos.xml` file contains a set of named templates you can call in your scripts. These named templates are discussed in *Named Templates in the jcs Namespace Overview* and *Named Templates in the jcs Namespace Summary*.

```
7    <xsl:import href="../../import/junos.xml"/>
```

Line 8 defines a template that matches the `</>` element. The `<xsl:template match="/">` element is the root element and represents the top level of the XML hierarchy. All XML Path Language (XPath) expressions in the script must start at the top level. This allows the script to access all possible Junos XML and Junos XML protocol remote procedure calls (RPCs). For more information, see *XPath Overview*.

```
8    <xsl:template match="/">
```

After the `<xsl:template match="/">` tag element, the `<op-script-results>` and `</op-script-results>` container tags must be the top-level child tags, as shown in Lines 9 and 10.

```
9      <op-script-results>
      <!-- ... insert your code here ... -->
10     </op-script-results>
```

Line 11 closes the template.

```
11    </xsl:template>
```

Between Line 11 and Line 12, you can define additional XSLT templates that are called from within the `<xsl:template match="/">` template.

Line 12 closes the style sheet and the op script.

```
12 </xsl:stylesheet>
```

SLAX Boilerplate for Op Scripts

The corresponding SLAX boilerplate is as follows:

```
version 1.0;

ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xml";

match / {
  <op-script-results> {
    /*
     * Insert your code here
    */
  }
}
```

Installing Junos OS Software with Junos Automation Enhancements

Junos operating system (Junos OS) with Junos Automation Enhancements is a full-featured version of Junos OS with Veriexec disabled, which can only be installed on supported devices.



NOTE: You must install the `jinstall-qfx-5-flex-x.tgz` software bundle in order to use the automation enhancements.

Before you install software, download the Junos OS `jinstall-qfx-5-flex-x.tgz` software bundle. For information on downloading and accessing the files, see [“Upgrading Software” on page 189](#).



BEST PRACTICE: Before you install the software, back up any critical files in `/var/home`. For more information regarding how to back up critical files, contact Customer Support at <http://www.juniper.net/support>.

Install the software:

1. Execute the **request system software add** command with the **validate** option:

- If the installation package resides locally on the switch, execute the **request system software add validate *pathname* source reboot** command, using the following format:

```
user@switch> request system software add validate /var/tmp/jinstall-qfx-5-flex-x.tgz
reboot
```

- If the installation package resides remotely, execute the **request system software add validate *pathname* source reboot** command, using the following format:

```
user@switch> request system software add validate
ftp://ftpserver/directory/jinstall-qfx-5-flex-x.tgz reboot
```

2. After the reboot has finished, verify that the new version of software has been properly installed by executing the **show version** command.

```
user@switch> show version
root@qfx5100-24q-et013> show version
fpc0:
-----
Hostname: qfx5100-24q-et013
Model: qfx5100-24q-2p
JUNOS Base OS Software Suite [13.2X51-D20]
JUNOS Base OS boot [13.2X51-D20]
JUNOS Crypto Software Suite [13.2X51-D20]
JUNOS Online Documentation [13.2X51-D20]
JUNOS Kernel Software Suite [13.2X51-D20]
JUNOS Packet Forwarding Engine Support (qfx-x86-32) [13.2X51-D20]
JUNOS Routing Software Suite [13.2X51-D20]
JUNOS Enterprise Software Suite [13.2X51-D20]
JUNOS py-base-i386 [13.2X51-D20]
Puppet on Junos [2.7.19_1.junos.i386]
Ruby Interpreter [11.10.4_1.junos.i386]
Chef [11.10.4_1.junos.i386]
junos-ez-stdlib [11.10.4_1.junos.i386]
JUNOS Host Software [13.2X51-D20]
JUNOS for Automation Enhancement
```



NOTE: If you are upgrading a device from standard Junos OS to use Junos Automation Enhancements and you are *not* loading the new factory default configuration, you need to use the following procedure.

To upgrade an existing device from standard Junos to use Junos Automation Enhancements:

1. Edit your existing Junos OS configuration to include the following configuration statements:

```
[edit]
user@switch# set system extensions providers juniper license-type juniper
deployment-scope commercial
user@switch# set system extensions providers chef license-type juniper
deployment-scope commercial
```



NOTE: The factory default configuration of the QFX5100 switch `jinstall-qfx-5-flex-x.tgz` software bundle is a Layer 3 configuration, whereas the factory default configuration for QFX5100 switch software bundles is a Layer 2 configuration. Therefore, if you are running the `jinstall-qfx-5-flex-x.tgz` software bundle on a QFX5100 switch and you use the `load factory-default` command, the resulting factory default configuration is set up for Layer 3 interfaces.

This is the factory default configuration for QFX5100 switch `jinstall-qfx-5-flex-x.tgz` software bundle:

```
user@switch> show configuration
```

```
system syslog user * any emergency
system syslog file messages any notice
system syslog file messages authorization info
system syslog file interactive-commands interactive-commands any
system extensions providers juniper license-type juniper deployment-scope
commercial
system extensions providers chef license-type juniper deployment-scope commercial
system commit factory-settings reset-virtual-chassis-configuration
system commit factory-settings reset-chassis-lcd-menu
system processes app-engine-virtual-machine-management-service traceoptions level
notice
system processes app-engine-virtual-machine-management-service traceoptions flag
all
interfaces et-0/0/0 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/0:0 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/0:1 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/0:2 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/0:3 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces et-0/0/1 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/1:0 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/1:1 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/1:2 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/1:3 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces et-0/0/2 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/2:0 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/2:1 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/2:2 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/2:3 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces et-0/0/3 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/3:0 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/3:1 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/3:2 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/3:3 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces et-0/0/4 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/4:0 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/4:1 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/4:2 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/4:3 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces et-0/0/5 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/5:0 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/5:1 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/5:2 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/5:3 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces et-0/0/6 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
```


Copyright © 2016, Juniper Networks, Inc. 1441

```

interfaces xe-0/0/18:1 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/18:2 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/18:3 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces et-0/0/19 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/19:0 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/19:1 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/19:2 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/19:3 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces et-0/0/20 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/20:0 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/20:1 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/20:2 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/20:3 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces et-0/0/21 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/21:0 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/21:1 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/21:2 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/21:3 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces et-0/0/22 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/22:0 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/22:1 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/22:2 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/22:3 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces et-0/0/23 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/23:0 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/23:1 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/23:2 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces xe-0/0/23:3 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces et-0/1/0 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces et-0/1/1 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces et-0/1/2 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces et-0/1/3 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces et-0/2/0 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces et-0/2/1 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces et-0/2/2 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
interfaces et-0/2/3 unit 0 family inet dhcp vendor-id Juniper-qfx5100-24q-2p
forwarding-options storm-control-profiles default all
protocols lldp interface all
protocols lldp-med interface all
protocols igmp-snooping vlan default
vlans default vlan-id 1

```

- Related Documentation**
- [Overview of Junos Automation Enhancements on page 1419](#)
 - [QFX5100 Switch with Automation Enhancements Frequently Asked Questions](#)

Invoking the Python Interpreter

The Python interpreter is available by default with the Junos Automation Enhancements. You can invoke Python by entering the **python** command at the shell script.

To invoke the Python interpreter:

1. Start the shell interface:

```
user@switch> start shell
```
2. Enter the **python** command without any parameters:

```
% python
```



NOTE: The Python interpreter is designated with the prompt `>>>` at the beginning of a line or `...` to indicate the continuation of a line.

Related Documentation

- [Overview of Python with Junos Automation Enhancements on page 1420](#)
- [Overview of Junos Automation Enhancements on page 1419](#)
- [Installing Junos OS Software with Junos Automation Enhancements on page 1438](#)
- [QFX5100 Switch with Automation Enhancements Frequently Asked Questions](#)

Controlling the Execution of Commit Scripts

This document describes the tasks that affect the way commit scripts are executed. In the QFabric system, commit scripts are stored in the in the `/pbdata/mgd_shared/partition-ip/var/db/scripts/commit` directory that is shared among Director devices in a Director group.

To determine which commit scripts are currently enabled on the QFabric system, use the **show** command to display the files included at the **[edit system scripts commit]** hierarchy level. To ensure that the enabled files are on the device, list the contents of the `/pbdata/mgd_shared/partition-ip/var/db/scripts/commit` directory using the **file list** operational mode command.

See the following tasks:

- [Enabling Commit Scripts to Execute on page 1443](#)
- [Removing Commit Scripts from the Configuration on page 1444](#)
- [Deactivating Commit Scripts on page 1445](#)
- [Activating Inactive Commit Scripts on page 1445](#)

Enabling Commit Scripts to Execute

The commit operation requires that all scripts be included in configuration at the **[edit system scripts commit file]** hierarchy level for all QFabric Director devices.

If you need to temporarily remove a script from a commit operation but do not want to remove it from the configuration permanently, you may configure the **optional** statement at the **[edit system scripts commit file filename]** hierarchy level to enable the commit operation to succeed even if a script is missing from the commit script directory.



CAUTION: When you include the **optional** statement at the **[edit system scripts commit file filename]** hierarchy level, no error message is generated during the commit operation if the file does not exist. As a result, you might not be aware that a script has not been executed as expected.

The filename of a commit script written in SLAX must include the **.slax** extension for the script to be executed.

To enable a commit script to execute during a commit operation:

1. Ensure that the commit script is located in the correct directory:
/pbdata/mgd_shared/partition-ip/var/db/scripts/commit directory on the Director device.

2. Configure the commit script.

```
[edit system scripts commit]
user@switch# set file filename <optional>
```

3. Commit the configuration.

```
[edit system scripts commit]
user@switch# top
[edit]
user@switch# commit
```

Removing Commit Scripts from the Configuration

You can prevent commit scripts from executing during a commit operation by removing the scripts from the commit directory in the configuration.



NOTE: You can also deactivate a script using the **deactivate** statement instead of removing it from the configuration. Deactivated scripts may be reactivated later.

To prevent a commit script from executing during a commit operation:

1. Delete the commit script file from the commit directory in the configuration.

```
[edit system scripts commit]
user@switch# delete file filename
```

2. Commit the configuration.

```
[edit system scripts commit]
user@switch# top
[edit]
user@switch# commit
```

3. Remove the commit script from the **/pbdata/mgd_shared/** directory on the Director device.



BEST PRACTICE: Although removing the commit script is not necessary, we recommend deleting unused files from the system.

Deactivating Commit Scripts

Deactivating a commit script results in its being marked as inactive in the configuration. The script is not executed during the commit operation, but you can reactivate the script by using the **activate** statement.

To deactivate the commit script:

1. Deactivate the script.

```
[edit]
user@switch deactivate system scripts commit file filename
```

2. Commit your changes.

```
[edit]
user@switch# commit
```

3. Verify that the commit script is deactivated.

```
[edit]
user@switch# show system scripts commit
inactive: file mycommit.slax
```

Activating Inactive Commit Scripts

Deactivating a commit script results in its being marked as inactive in the configuration and is therefore not executed during the commit operation.

To activate an inactive commit script:

1. Activate the script.

```
[edit]
user@switch# activate system scripts commit file filename
```

2. Commit your changes.

```
[edit]
user@switch# commit
```

Displaying Commit Script Output

Table 92 summarizes the Junos OS command-line interface (CLI) commands you can use to monitor and troubleshoot commit scripts. For more information about the **cscript.log** file, see *Tracing Commit Script Processing*.

Table 92: Commit Script Configuration and Operational Mode Commands

Task	Command
Configuration Mode Commands	
Display errors and warnings generated by commit scripts.	commit or commit check

Table 92: Commit Script Configuration and Operational Mode Commands (*continued*)

Task	Command
Display detailed information.	commit display detail
Display the underlying Extensible Markup Language (XML) data.	commit display xml
Display the postinheritance contents of the configuration database. This view includes transient changes, but does not include changes made in configuration groups.	show display commit-scripts
Display the postinheritance contents of the configuration database. This view excludes transient changes.	show display commit-scripts no-transients
Display the postinheritance configuration in XML format. Viewing the configuration in XML format can be helpful when you are writing XML Path Language (XPath) expressions and configuration element tags.	show display commit-scripts view
Display the postinheritance configuration in XML format, but exclude transient changes.	show display commit-scripts view display commit-scripts no-transients
Display all configuration groups data, including script-generated changes to the groups.	show groups display commit-scripts
Display a particular configuration group, including script-generated changes to the group.	show groups <i>group-name</i> display commit-scripts
Operational Mode Commands	
Display logging data associated with all commit script processing.	show log cscript.log
Display processing for only the most recent commit operation.	show log cscript.log last
Display processing for script errors.	show log cscript.log match error
Display processing for a particular script.	show log cscript.log match <i>filename</i>

Related Documentation

- *Tracing Commit Script Processing*

CHAPTER 47

Configuring sFlow Technology

- [Understanding How to Use sFlow Technology for Network Monitoring on a Switch on page 1449](#)
- [Example: Monitoring Network Traffic Using sFlow Technology on page 1453](#)
- [Configuring sFlow Technology on page 1457](#)

Understanding How to Use sFlow Technology for Network Monitoring on a Switch

The sFlow technology is a monitoring technology for high-speed switched or routed networks. sFlow monitoring technology randomly samples network packets and sends the samples to a monitoring station called a *collector*. You can configure sFlow technology on a Juniper Networks switch to continuously monitor traffic at wire speed on all interfaces simultaneously.

This topic describes:

- [Sampling Mechanism and Architecture of sFlow Technology on Switches on page 1449](#)
- [Adaptive Sampling on page 1451](#)
- [sFlow Agent Address Assignment on page 1452](#)
- [sFlow Limitations on Switches on page 1452](#)

Sampling Mechanism and Architecture of sFlow Technology on Switches

sFlow technology uses the following two sampling mechanisms:

- **Packet-based sampling**—Samples one packet out of a specified number of packets from an interface enabled for sFlow technology. Only the first 128 bytes of each packet are sent to the collector. Data collected include the Ethernet, IP, and TCP headers, along with other application-level headers (if present). Although this type of sampling might not capture infrequent packet flows, the majority of flows are reported over time, allowing the collector to generate a reasonably accurate representation of network activity. To configure packet-based sampling, you must specify a sample rate.
- **Time-based sampling**—Samples interface statistics at a specified interval from an interface enabled for sFlow technology. Statistics such as Ethernet interface errors are captured. To configure time-based sampling, you must specify a polling interval.

The sampling information is used to create a network traffic visibility picture. The Juniper Networks Junos operating system (Junos OS) fully supports the sFlow standard described in RFC 3176, *InMon Corporation's sFlow: A Method for Monitoring Traffic in Switched and Routed Networks* (see <http://faqs.org/rfcs/rfc3176.html>).



NOTE: sFlow technology on the switches samples only raw packet headers. A raw Ethernet packet is the complete Layer 2 network frame.

An sFlow monitoring system consists of an sFlow agent embedded in the switch and a centralized collector. The sFlow agent's two main activities are random sampling and statistics gathering. It combines interface counters and flow samples and sends them across the network to the sFlow collector as UDP datagrams, directing those datagrams to the IP address and UDP destination port of the collector. Each datagram contains the following information:

- The IP address of the sFlow agent
- The number of samples
- The interface through which the packets entered the agent
- The interface through which the packets exited the agent
- The source and destination interface for the packets
- The source and destination VLAN for the packets

EX Series switches, QFX Series switches, and the QFabric systems adopt the distributed sFlow architecture. The sFlow agent has two separate sampling entities that are associated with each Packet Forwarding Engine in case of switches and nodes in case of a QFabric system. These sampling entities are known as subagents. Each subagent has a unique ID that is used by the collector to identify the data source. A subagent has its own independent state and forwards its own sample messages to the sFlow agent. The sFlow agent is responsible for packaging the samples into datagrams and sending them to the sFlow collector. Because sampling is distributed across subagents, the protocol overhead associated with sFlow technology is significantly reduced at the collector.



NOTE: On the QFabric system, an sFlow collector must be reachable through the data network. Because each Node device has all routes stored in the default routing instance, the collector IP address should be included in the default routing instance to ensure the collector's reachability from the Node device.



NOTE: You cannot configure sFlow monitoring on a link aggregation group (LAG), but you can configure it individually on a LAG member interface.

Infrequent sampling flows might not be reported in the sFlow information, but over time the majority of flows are reported. Based on a configured sampling rate N , 1 out of N packets is captured and sent to the collector. This type of sampling does not provide a 100 percent accurate result in the analysis, but it does provide a result with quantifiable accuracy. A user-configured polling interval defines how often the sFlow data for a specific interface are sent to the collector, but an sFlow agent can also schedule polling.



NOTE: We recommend that you configure the same sample rate for all the ports in a line card. If you configure different sample rates, the lowest value is used for all ports on the line card..



NOTE: If the mastership assignment changes in a Virtual Chassis setup, sFlow technology continues to function.

Adaptive Sampling

To ensure sampling accuracy and efficiency, EX Series switches and QFX Series devices use adaptive sFlow sampling. Adaptive sampling monitors the overall incoming traffic rate on the device and provides feedback to the interfaces to dynamically adapt their sampling rate to traffic conditions. The sFlow agent reads the statistics on the interfaces every few seconds (12 seconds for EX Series switches and 5 seconds for QFX Series devices) and identifies five interfaces with the highest number of samples.

On a Flexible PIC Concentrator (FPC), when the CPU processing limit is reached because of sflow sample processing, a binary backoff algorithm is initiated. This reduces the sampling load, arriving through the top five sample-producing interfaces on that FPC by half. The backoff algorithm achieves this by doubling the sampling rate on these five earmarked interfaces. This process is repeated until the CPU-load due to sflow on the given FPC comes down to an acceptable level.

On a QFabric system, sFlow technology monitors the interfaces on each node device as a group, and implements the binary backoff algorithm based on the traffic on that group of interfaces.



NOTE: On the QFX Series standalone switches, if you configure sFlow technology monitoring on multiple interfaces and with a high sampling rate, we recommend that you specify a collector that is on the data network instead of on the management network. Having a high volume of sFlow technology monitoring traffic on the management network might interfere with other management interface traffic.

Using adaptive sampling prevents overloading of the CPU and keeps the device operating at its optimum level even when there is a change in traffic patterns on the interfaces. The reduced sampling rate is used until the device is rebooted or when a new sampling rate is configured.



NOTE: sFlow technology on EX Series switches does not support graceful restart. When a graceful restart occurs, the adaptive sampling rate is set to the user-configured sampling rate.

sFlow Agent Address Assignment

The sFlow collector uses the sFlow agent's IP address to determine the source of the sFlow data. You can configure the IP address of the sFlow agent to ensure that the agent ID of the sFlow agent remains constant. If you do not specify the IP address to be assigned to the agent, an IP address is automatically assigned to the agent based on the following order of priority of interfaces configured on the device:

EX Series Devices	QFX Series Devices
<ol style="list-style-type: none"> 1. Virtual Management Ethernet (VME) interface 2. Management Ethernet interface 	<ol style="list-style-type: none"> 1. Management Ethernet interface me0 IP address 2. Any Layer 3 interface if the me0 IP address is not available

If a particular interface is not configured, the IP address of the next interface in the priority list is used as the IP address for the agent. Once an IP address is assigned to the agent, the agent ID is not modified until the sFlow service is restarted. At least one interface has to be configured for an IP address to be assigned to the agent. When the agent's IP address is assigned automatically, the IP address is dynamic and changes when the switch reboots.

On the QFabric system, the following default values are used if the optional parameters are not configured:

- Agent ID is the management IP address of the default partition.
- Source IP is the management IP address of the default partition.

In addition, the QFabric system subagent ID (which is included in the sFlow datagrams) is the ID of the node group from which the datagram is sent to the collector.

sFlow data can be used to provide network traffic visibility information. You can explicitly configure the source IP address to be assigned to the sFlow datagrams. If you do not explicitly configure the IP address, the IP address of any of the configured Layer 3 network interfaces is used as the source IP address. If a Layer 3 IP address is not configured, then the agent IP address is used as the source IP address.

sFlow Limitations on Switches

On the QFX Series, limitations of sFlow traffic sampling include the following:

- sFlow sampling on ingress interfaces does not capture CPU-bound traffic.
- sFlow sampling on egress interfaces does not support broadcast and multicast packets.
- Egress samples do not contain modifications made to the packet in the egress pipeline.

- If a packet is discarded because of a firewall filter, the reason code for discarding the packet is not sent to the collector.
- The out-priority field for a VLAN is always set to 0 (zero) on ingress and egress samples.
- On QFX5100 standalone switches and the QFX Series Virtual Chassis (including mixed QFX Series Virtual Chassis), egress firewall filters are not applied to sFlow sampling packets. On these platforms, the software architecture is different from that on other QFX Series devices—sFlow packets are sent by the Routing Engine (not the line card on the host) and do not transit the switch. Egress firewall filters affect data packets that are transiting a switch, but do not affect packets sent by the Routing Engine. As a result, sFlow sampling packets are always sent to the sFlow collector.

EX9200 switches support configuration of only one sampling rate (inclusive of ingress and egress rates) on an FPC. To support compatibility with the sflow configuration of other Juniper Networks products, EX9200 switches still accept multiple rate configuration on different interfaces of the same FPC. However, the switch programs the lowest rate as the sampling rate for all the interfaces of that FPC. The sFlow show command (**show sflow interfaces**) displays the configured rate and the actual (effective) rate. However, different rates on different FPCs is still supported on EX9200 switches.

Related Documentation

- [Example: Monitoring Network Traffic Using sFlow Technology on page 1453](#)
- [Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches](#)
- [Configuring sFlow Technology on page 1457](#)
- [Configuring sFlow Technology for Network Monitoring \(CLI Procedure\)](#)
- [Monitoring Interface Status and Traffic](#)

Example: Monitoring Network Traffic Using sFlow Technology

The sFlow technology is a monitoring technology for high-speed switched or routed networks. sFlow monitoring technology collects samples of network packets and sends them in a UDP datagram to a monitoring station called a *collector*. You can configure sFlow technology on a device to monitor traffic continuously at wire speed on all interfaces simultaneously. You must enable sFlow monitoring on each interface individually; you cannot globally enable sFlow monitoring on all interfaces with a single configuration statement. Junos OS fully supports the sFlow technology standard described in RFC 3176, *InMon Corporation's sFlow: A Method for Monitoring Traffic in Switched and Routed Networks*.

This example describes how to configure and use sFlow monitoring on a QFX3500 switch in standalone mode.

- [Requirements on page 1454](#)
- [Overview on page 1454](#)
- [Configuration on page 1455](#)
- [Verification on page 1456](#)

Requirements

This example uses the following hardware and software components:


- Junos OS Release 11.3 or later
- One QFX3500 switch

Overview

An sFlow monitoring system consists of an sFlow agent embedded in the device and a centralized collector on the network. The two main activities of the sFlow agent are random sampling and statistics gathering. The sFlow agent combines interface counters and flow samples and sends them to the IP address and UDP destination port of the sFlow collector in UDP datagrams.

Figure 17 depicts the basic elements of an sFlow system.

Figure 17: sFlow Technology Monitoring System



Configuration

CLI Quick Configuration To quickly configure sFlow technology, copy the following commands and paste them into the terminal window of the switch:

```
[edit protocols sflow]
set collector 10.204.32.46 udp-port 5600
set interfaces xe-0/0/1.0
set polling-interval 20
set sample-rate 1000
```

Step-by-Step Procedure To configure sFlow features using the CLI:

1. Configure the IP address and UDP port of at least one collector:

```
[edit protocols sflow]
user@switch# set collector 10.204.32.46 udp-port 5600
```

The default UDP port assigned is 6343.

2. Enable sFlow technology on a specific interface:

```
[edit protocols sflow]
user@switch# set interfaces xe-0/0/1.0
```



NOTE: You cannot enable sFlow technology on a Layer 3 VLAN-tagged interface.

You cannot enable sFlow technology on a LAG interface (for example, ae0), but you can enable sFlow technology on the member interfaces of the LAG (for example, xe-0/0/1).

3. Specify how often (in seconds) the sFlow agent polls all interfaces at the global level:

```
[edit protocols sflow]
user@switch# set polling-interval 20
```



NOTE: Specify 0 if you do not want to poll the interface.

4. Specify the rate at which packets must be sampled at the global level. The following example sets a sample rate of 1 in 1000 packets:

```
[edit protocols sflow]
user@switch# set sample-rate 1000
```

Results Check the results of the configuration:

```
[edit]
user@switch# show protocols
sflow {
  collector 10.204.32.46 {
    udp-port 5600;
```

```

    }
    interfaces xe-0/0/1.0 {
        polling-interval 20;
        sample-rate 1000;
    }
}

```

Verification

To confirm that the configuration is correct, perform these tasks:

- [Verifying That sFlow Technology Has Been Configured Properly on page 1456](#)
- [Verifying That sFlow Technology Is Enabled on an Interface on page 1456](#)
- [Verifying the sFlow Collector Configuration on page 1457](#)

Verifying That sFlow Technology Has Been Configured Properly

Purpose Verify that sFlow technology has been configured properly.

Action Enter the **show sflow** operational mode command:

```

user@switch> show sflow
sFlow           : Enabled
Sample limit    : 300 packets/second
Polling interval : 20 second
Sample rate     : 1:1000
Agent ID        : 10.1.1.2

```



NOTE: The sample limit cannot be configured and is set to 300 packets per second.

Meaning The output shows that sFlow technology is enabled and specifies the values for the sampling limit, polling interval, and sampling rate.

Verifying That sFlow Technology Is Enabled on an Interface

Purpose Verify that sFlow technology is enabled on interfaces and display the sampling parameters.

Action Enter the **show sflow interface** operational mode command:

```

user@switch> show sflow interface
Interface      Status      Sample    Polling
                rate      interval
xe-0/0/1.0    Enabled      1000       20

```

Meaning The output indicates that sFlow technology is enabled on the **Node1:xe-0/0/1.0** interface on the Node device with a sampling rate of 1000 and a polling interval of 20 seconds.

Verifying the sFlow Collector Configuration

Purpose	Verify the sFlow collector configuration.						
Action	<p>Enter the show sflow collector operational mode command:</p> <pre>user@switch> show sflow collector</pre> <table><tr><th>Collector address</th><th>Udp-port</th><th>No. of samples</th></tr><tr><td>10.204.32.46</td><td>5600</td><td>7516</td></tr></table>	Collector address	Udp-port	No. of samples	10.204.32.46	5600	7516
Collector address	Udp-port	No. of samples					
10.204.32.46	5600	7516					
Meaning	The output displays the IP address of the collector, the UDP port, and the number of samples collected.						
Related Documentation	<ul style="list-style-type: none">• Configuring sFlow Technology on page 1457• <i>Overview of sFlow Technology</i>						

Configuring sFlow Technology

The sFlow technology is a monitoring technology for high-speed switched or routed networks. sFlow monitoring technology collects samples of network packets and sends them in a UDP datagram to a monitoring station called a *collector*. You can configure sFlow technology on a device to monitor traffic continuously at wire speed on all interfaces simultaneously. You must enable sFlow monitoring on each interface individually; you cannot globally enable sFlow monitoring on all interfaces with a single configuration statement. Junos OS fully supports the sFlow technology standard described in RFC 3176, *InMon Corporation's sFlow: A Method for Monitoring Traffic in Switched and Routed Networks*.

On the QFabric system, the sFlow monitoring global configuration that is defined on the Director device is distributed to Node groups that have sFlow sampling configured on the interfaces.

To configure sFlow features using the CLI:

1. Configure the IP address and UDP port of at least one collector:

```
[edit protocols sflow]
user@host# set collector ip-address udp-port port-number
```

The default UDP port assigned is 6343.

2. Enable sFlow technology on a specific interface:

```
[edit protocols sflow]
user@host# set interfaces interface-name
```



NOTE: You cannot enable sFlow technology on a Layer 3 VLAN-tagged interface.

You cannot enable sFlow technology on a LAG interface (for example ae0), but you can enable sFlow technology on the member interfaces of the LAG (for example, xe-0/0/1).

3. Specify how often (in seconds) the sFlow agent polls all interfaces at the global level:

```
[edit protocols sflow]
user@host# set polling-interval seconds
```



NOTE: Specify 0 if you do not want to poll the interface.

4. Specify the rate at which packets are sampled at the global level. For example, configuring a *number* of 1000 sets a sample rate of 1 in 1000 packets.

```
[edit protocols sflow]
user@host# set sample-rate number
```

5. (Optional) You can also configure the polling interval and sample rate at the interface level:

```
[edit protocols sflow]
user@host# set interfaces interface-name polling-interval seconds sample-rate number
```



NOTE: The interface-level configuration overrides the global configuration for the specified interface.

**Related
Documentation**

- [Example: Monitoring Network Traffic Using sFlow Technology on page 1453](#)
- [Overview of sFlow Technology](#)

CHAPTER 48

Configuring SNMP

- [Utility MIB on page 1460](#)
- [SNMPv3 Overview on page 1460](#)
- [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 1461](#)
- [Understanding RMON on page 1463](#)
- [RMON MIB Event, Alarm, Log, and History Control Tables on page 1464](#)
- [Understanding Health Monitoring on page 1467](#)
- [SNMP MIBs Support on page 1468](#)
- [SNMP Traps Support on page 1484](#)
- [MIB Objects for the QFX Series on page 1497](#)
- [Configuring SNMP on page 1499](#)
- [Configuring the SNMP Community String on page 1503](#)
- [Configuring SNMP Trap Groups on page 1504](#)
- [Adding a Group of Clients to an SNMP Community on page 1505](#)
- [Configuring the Interfaces on Which SNMP Requests Can Be Accepted on page 1506](#)
- [Configuring MIB Views on page 1506](#)
- [Configuring RMON Alarms and Events on page 1507](#)
- [Configuring Health Monitoring on page 1509](#)
- [Creating SNMPv3 Users on page 1510](#)
- [Configuring Access Privileges for a Group on page 1511](#)
- [Assigning a Security Name to a Group on page 1513](#)
- [Configuring SNMPv3 Traps on a Device Running Junos OS on page 1514](#)
- [Configuring SNMP Informs on page 1515](#)
- [Monitoring RMON MIB Tables on page 1516](#)
- [Monitoring SNMP on page 1516](#)
- [Tracing SNMP Activity on a Device Running Junos OS on page 1518](#)
- [Using the Enterprise-Specific Utility MIB to Enhance SNMP Coverage on page 1521](#)
- [Example: Configuring SNMP on page 1523](#)

Utility MIB

The Juniper Networks enterprise-specific Utility MIB, whose object ID is `{jnxUtilMibRoot 1}`, defines objects for counters, integers, and strings. The Utility MIB contains one table for each of the following five data types:

- 32-bit counters
- 64-bit counters
- Signed integers
- Unsigned integers
- Octet strings

Each data type has an arbitrary ASCII name, which is defined when the data is populated, and a timestamp that shows the last time when the data instance was modified.

For a downloadable version of this MIB, see http://www.juniper.net/techpubs/en_US/junos15.1x49/topics/reference/mibs/mib-jnx-util.txt.

For information about the enterprise-specific Utility MIB objects, see the following topics:

- *jnxUtilCounter32Table*
- *jnxUtilCounter64Table*
- *jnxUtilIntegerTable*
- *jnxUtilUintTable*
- *jnxUtilStringTable*

Related Documentation

- *Enterprise-Specific SNMP MIBs Supported by Junos OS*
- *Standard SNMP MIBs Supported by Junos OS*
- *Understanding the Implementation of SNMP on the QFabric System*

SNMPv3 Overview

The QFX3500 switch supports SNMP version 3 (SNMPv3). SNMPv3 enhances the functionality of SNMPv1 and SNMPv2c by supporting user authentication and data encryption. SNMPv3 uses the user-based security model (USM) to provide security for SNMP messages, and the view-based access control model (VACM) for user access control.

SNMPv3 features include:

- With USM, the SNMP messages between the SNMP manager and the agent can have the message source authenticated and the data integrity checked. USM reduces messaging delays and message replays by enforcing timeout limits and by checking for duplicate message request IDs.
- VACM complements USM by providing user access control for SNMP queries to the agent. You define access privileges that you wish to extend to a group of one or more users. Access privileges are determined by the security model parameters (**usm**, **v1**, or **v2**) and security level parameters (**authentication**, **privacy**, or **none**). For each security level, you must associate one MIB view for the group. Associating a MIB view with a group grants the read, write, or notify permission to a set of MIB objects for the group.
- You configure security parameters for each user, including the username, authentication type and authentication password, and privacy type and privacy password. The username given to each user is in a format that is dependent on the security model configured for that user.
- To ensure messaging security, another type of username, called the security name, is included in the messaging data that is sent between the local SNMP server and the destination SNMP server. Each user name is mapped to a security name, but the security name is in a format that is independent of the security model.
- Trap entries in SNMPv3 are created by configuring the notify, notify filter, target address, and target parameters. The **notify** statement specifies the type of notification (trap) and contains a single tag that defines a set of target addresses to receive a trap. The notify filter defines access to a collection of trap object identifiers (OIDs). The target address defines the address of an SNMP management application and other attributes used in sending notifications. Target parameters define the message processing and security parameters used in sending notifications to a particular target.

**Related
Documentation**

- [Assigning a Security Name to a Group on page 1513](#)
- [Configuring Access Privileges for a Group on page 1511](#)
- [Configuring SNMP Informs on page 1515](#)
- [Creating SNMPv3 Users on page 1510](#)

Minimum SNMPv3 Configuration on a Device Running Junos OS

To configure the minimum requirements for SNMPv3, include the following statements at the **[edit snmp v3]** and **[edit snmp]** hierarchy levels:



NOTE: You must configure at least one view (notify, read, or write) at the **[edit snmp view-name]** hierarchy level.

```
[edit snmp]
view view-name {
  oid object-identifier (include | exclude);
}
[edit snmp v3]
```

```
notify name {
    tag tag-name;
}
notify-filter profile-name {
    oid object-identifier (include | exclude);
}
snmp-community community-index {
    security-name security-name;
}
target-address target-address-name {
    address address;
    target-parameters target-parameters-name;
}
target-parameters target-parameters-name {
    notify-filter profile-name;
    parameters {
        message-processing-model (v1 | v2c | v3);
        security-level (authentication | none | privacy);
        security-model (usm | v1 | v2c);
        security-name security-name;
    }
}
usm {
    local-engine {
        user username {
        }
    }
}
vacm {
    access {
        group group-name {
            (default-context-prefix | context-prefix context-prefix){
                security-model (any | usm | v1 | v2c) {
                    security-level (authentication | none | privacy) {
                        notify-view view-name;
                        read-view view-name;
                        write-view view-name;
                    }
                }
            }
        }
    }
    security-to-group {
        security-model (usm | v1 | v2c) {
            security-name security-name {
                group group-name;
            }
        }
    }
}
```

- Related Documentation**
- [Creating SNMPv3 Users on page 1510](#)
 - [Configuring MIB Views on page 1506](#)
 - [Defining Access Privileges for an SNMP Group](#)

- [Configuring SNMPv3 Traps on a Device Running Junos OS on page 1514](#)
- [Configuring SNMP Informs on page 1515](#)
- [Complete SNMPv3 Configuration Statements](#)
- [Example: SNMPv3 Configuration](#)

Understanding RMON

- [RMON Overview on page 1463](#)
- [Alarm Thresholds and Events on page 1464](#)

RMON Overview

The Junos OS supports the *Remote Network Monitoring* (RMON) MIB (RFC 2819), which allows a management device to monitor the values of MIB objects, or variables, against configured thresholds. When the value of a variable crosses a threshold, an alarm and its corresponding event are generated. The event can be logged and can generate an SNMP trap.

An operational support system (OSS) or a fault-monitoring system can be used to automatically monitor events that track many different metrics, including performance, availability, faults, and environmental data. For example, an administrator might want to know when the internal temperature of a chassis has risen above a configured threshold, which might indicate that a chassis fan tray is faulty, the chassis air flow is impeded, or the facility cooling system in the vicinity of the chassis is not operating normally.

The RMON MIB also defines tables that store various statistics for Ethernet interfaces, including the **etherStatsTable** and the **etherHistoryTable**. The **etherStatsTable** contains cumulative real-time statistics for Ethernet interfaces, such as the number of unicast, multicast, and broadcast packets received on an interface. The **etherHistoryTable** maintains a historical sample of statistics for Ethernet interfaces. The control of the **etherHistoryTable**, including the interfaces to track and the sampling interval, is defined by the RMON **historyControlTable**.

To enable RMON alarms, you perform the following steps:


1. Configure SNMP, including trap groups. You configure SNMP at the **[edit snmp]** hierarchy level.
2. Configure rising and falling events in the **eventTable**, including the event types and trap groups. You can also configure events using the CLI at the **[edit snmp rmon event]** hierarchy level.
3. Configure alarms in the **alarmTable**, including the variables to monitor, rising and falling thresholds, the sampling types and intervals, and the corresponding events to generate when alarms occur. You can also configure alarms using the CLI at the **[edit snmp rmon alarm]** hierarchy level.

Extensions to the **alarmTable** are defined in the Juniper Networks enterprise-specific MIB **jnxRmon** (mib-jnx-rmon.txt).

Alarm Thresholds and Events

By setting a rising and a falling threshold for a monitored variable, you can be alerted whenever the value of the variable falls outside the allowable operational range (see [Figure 18](#)).

Figure 18: Setting Thresholds



Events are only generated when the alarm threshold is first crossed in any one direction rather than after each sample interval. For example, if a rising threshold alarm, along with its corresponding event, is raised, no more threshold crossing events occur until a corresponding falling alarm occurs. This considerably reduces the quantity of events that are produced by the system, making it easier for operations staff to react when events do occur.

Before you configure remote monitoring, you should identify what variables need to be monitored and their allowable operational range. This requires some period of baselining to determine the allowable operational ranges. An initial baseline period of at least 3 months is not unusual when you first identify the operational ranges and define thresholds, but baseline monitoring should continue over the life span of each monitored variable.

Related Documentation

- [Configuring RMON Alarms and Events on page 1507](#)
- [Juniper Networks Enterprise-Specific MIBs](#)
- [RMON MIB Event, Alarm, Log, and History Control Tables on page 1464](#)

RMON MIB Event, Alarm, Log, and History Control Tables

The Junos OS supports the *Remote Network Monitoring* (RMON) MIB (RFC 2819), which allows a management device to monitor the values of MIB objects, or variables, against configured thresholds. When the value of a variable crosses a threshold, an alarm and

its corresponding event are generated. The event can be logged and can generate an SNMP trap.

Table 93 provides each field in the RMON eventTable, the description of the field, and the corresponding Junos OS statement that you can use to configure the field. The Junos OS statements reside at the `[edit snmp rmon]` hierarchy level.

Table 93: RMON Event Table

Field	Description	Statement [edit snmp rmon]
eventDescription	Text description of this event.	description
eventType	Type of event (for example, log, trap, or log and trap).	type
eventCommunity	Trap group to which to send this event, as defined in the Junos OS configuration. (This is not the same as the SNMP community.)	community
eventOwner	Entity (for example, manager) that created this event.	—
eventStatus	Status of this row (for example, valid, invalid, or createRequest).	—

Table 94 provides each field in the RMON alarmTable, the description of the field, and the corresponding Junos OS statement that you can use to configure the field. The Junos OS statements reside at the `[edit snmp rmon]` hierarchy level.

Table 94: RMON Alarm Table

Field	Description	Statement [edit snmp rmon]
alarmStatus	Status of this row (for example, valid, invalid, or createRequest)	—
alarmInterval	Sampling period (in seconds) of the monitored variable	interval
alarmVariable	Object identifier (OID) and instance of the variable to be monitored	—
alarmValue	Actual value of the sampled variable	—
alarmSampleType	Sample type (absolute or delta changes)	sample-type
alarmStartupAlarm	Initial alarm (rising, falling, or either)	startup-alarm
alarmRisingThreshold	Rising threshold against which to compare the value	rising-threshold
alarmFallingThreshold	Falling threshold against which to compare the value	falling-threshold
alarmRisingEventIndex	Index (row) of the rising event in the event table	rising-event-index
alarmFallingEventIndex	Index (row) of the falling event in the event table	falling-event-index

[Table 95](#) provides each field in the jnxRmon jnxRmonAlarmTable, which is an extension to the RMON alarmTable. You can troubleshoot the RMON agent, rmopd, that runs on a switch by inspecting the contents of the jnxRmonAlarmTable object.

Table 95: jnxRmon Alarm Table

Field	Description
jnxRmonAlarmGetFailCnt	Number of times the internal Get request for the variable failed
jnxRmonAlarmGetFailTime	Value of the sysUpTime object when the last failure occurred
jnxRmonAlarmGetFailReason	Reason why the Get request failed
jnxRmonAlarmGetOkTime	Value of the sysUpTime object when the variable moved out of failure state
jnxRmonAlarmState	Status of this alarm entry

[Table 96](#) provides each field in the RMON historyControlTable, the description of the field, and the corresponding Junos OS statement that you can use to configure the field. The Junos OS statements reside at the **[edit snmp rmon history]** hierarchy level. The historyControlTable controls the RMON etherHistoryTable.

Table 96: RMON History Control Table

Field	Description	Statement [edit snmp rmon history]
historyControlDataSource	Identifies the source of the data for which historical data was collected.	interface
historyControlBucketsRequested	Requested number of discrete time intervals over which data is to be saved.	bucket-size
historyControlBucketsGranted	Number of discrete sampling intervals over which data is to be saved.	—
historyControlInterval	Interval, in seconds, over which the data is sampled for each bucket.	interval
historyControlOwner	Entity that configured this entry.	owner
historyControlStatus	Status of this entry.	—

- Related Documentation**
- [Configuring RMON Alarms and Events on page 1507](#)
 - [Juniper Networks Enterprise-Specific MIBs](#)
 - [Understanding RMON on page 1463](#)

Understanding Health Monitoring

Health monitoring is an SNMP feature that extends the RMON alarm infrastructure to provide monitoring for a predefined set of objects (such as file system usage, CPU usage, and memory usage), and for Junos OS processes.

You enable the health monitor feature using the **health-monitor** statement at the **[edit snmp]** hierarchy level. You can also configure health monitor parameters such as a falling threshold, rising threshold, and interval. If the value of a monitored object exceeds the rising or falling threshold, an alarm is triggered and an event may be logged.

The falling threshold is the lower threshold for the monitored object instance. The rising threshold is the upper threshold for the monitored object instance. Each threshold is expressed as a percentage of the maximum possible value. The interval represents the period of time, in seconds, over which the object instance is sampled and compared with the rising and falling thresholds.

Events are only generated when a threshold is first crossed in any one direction, rather than after each sample interval. For example, if a rising threshold alarm, along with its corresponding event, is raised, no more threshold crossing events occur until a corresponding falling alarm occurs.

System log entries for health monitor events have a corresponding HEALTHMONITOR tag and not a generic SNMPD_RMON_EVENTLOG tag. However, the health monitor sends generic RMON risingThreshold and fallingThreshold traps. You can use the **show snmp health-monitor** operational command to view information about health monitor alarms and logs.

When you configure the health monitor, monitoring information for certain object instances is available, as shown in [Table 97](#).

Table 97: Monitored Object Instances

Object	Description
jnxHrStoragePercentUsed.1	Monitors the /dev/ad0s1a: file system on the switch. This is the root file system mounted on / .
jnxHrStoragePercentUsed.2	Monitors the /dev/ad0s1e: file system on the switch. This is the configuration file system mounted on /config .
jnxOperatingCPU (RE0)	Monitors CPU usage by the Routing Engine (RE0).
jnxOperatingBuffer (RE0)	Monitors the amount of memory available on the Routing Engine (RE0).
sysAppElmtRunCPU	Monitors the CPU usage for each Junos OS process (also called daemon). Multiple instances of the same process are monitored and indexed separately.
sysAppElmtRunMemory	Monitors the memory usage for each Junos OS process. Multiple instances of the same process are monitored and indexed separately.

- Related Documentation**
- [Configuring Health Monitoring on page 1509](#)
 - [falling-threshold \(Health Monitor\) on page 1624](#)
 - [interval \(Health Monitor\) on page 1633](#)
 - [rising-threshold \(Health Monitor\) on page 1653](#)
 - [show snmp health-monitor on page 1749](#)

SNMP MIBs Support

The QFX Series standalone switches, QFX Series Virtual Chassis, and QFabric systems support standard MIBs and Juniper Networks enterprise-specific MIBs.

For more information, see:

- [MIBs Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis on page 1468](#)
- [MIBs Supported on QFabric Systems on page 1478](#)

MIBs Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis

The QFX Series standalone switches and QFX Series Virtual Chassis support both standard MIBs and Juniper Networks enterprise-specific MIBs. For more information, see:

- [Table 98](#) for standard MIBs.
- [Table 99](#) for Juniper Networks enterprise-specific MIBs.

Table 98: Standard MIBs Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis

RFC	Additional Information
IEEE 802.1ab section 12.1, <i>Link Layer Discovery Protocol (LLDP) MIB</i>	<p>Supported tables and objects:</p> <ul style="list-style-type: none"> • lldpRemManAddrOID • lldpLocManAddrOID • lldpReinitDelay • lldpNotificationInterval • lldpStatsRxPortFramesDiscardedTotal • lldpStatsRxPortFramesError • lldpStatsRxPortTLVsDiscardedTotal • lldpStatsRxPortTLVsUnrecognizedTotal • lldpStatsRxPortAgeoutsTotal

Table 98: Standard MIBs Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis (*continued*)

RFC	Additional Information
IEEE 802.3ad, <i>Aggregation of Multiple Link Segments</i>	<p>The following tables and objects are supported:</p> <ul style="list-style-type: none"> dot3adAggPortTable, dot3adAggPortListTable, dot3adAggTable, and dot3adAggPortStatsTable dot3adAggPortDebugTable (only dot3adAggPortDebugRxState, dot3adAggPortDebugMuxState, dot3adAggPortDebugActorSyncTransitionCount, dot3adAggPortDebugPartnerSyncTransitionCount, dot3adAggPortDebugActorChangeCount, and dot3adAggPortDebugPartnerChangeCount) dot3adTablesLastChanged
RFC 1155, <i>Structure and Identification of Management Information for TCP/IP-based Internets</i>	—
RFC 1157, <i>A Simple Network Management Protocol (SNMP)</i>	—
RFC 1212, <i>Concise MIB Definitions</i>	—
RFC 1213, <i>Management Information Base for Network Management of TCP/IP-Based Internets: MIB-II</i>	<p>The following areas are supported:</p> <ul style="list-style-type: none"> MIB II and its SNMP version 2 derivatives, including: <ul style="list-style-type: none"> Statistics counters IP, except for ipRouteTable, which has been replaced by ipCidrRouteTable (RFC 2096, <i>IP Forwarding Table MIB</i>) ipAddrTable SNMP management Interface management SNMPv1 Get, GetNext requests, and SNMPv2 GetBulk request Junos OS-specific secured access list Master configuration keywords Reconfigurations upon SIGHUP
RFC 1215, <i>A Convention for Defining Traps for use with the SNMP</i>	Support is limited to MIB II SNMP version 1 traps and version 2 notifications.
RFC 1286, <i>Definitions of Managed Objects for Bridges</i>	—
RFC 1657, <i>Definitions of Managed Objects for the Fourth Version of the Border Gateway Protocol (BGP-4) using SMIv2</i>	—

Table 98: Standard MIBs Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis (continued)

RFC	Additional Information
RFC 1850, <i>OSPF Version 2 Management Information Base</i>	<p>The following table, objects, and traps are not supported:</p> <ul style="list-style-type: none"> Host Table ospfOriginateNewLsas and ospfRxNewLsas objects ospfOriginateLSA, ospfLsdbOverflow, and ospfLsdbApproachingOverflow traps
RFC 1901, <i>Introduction to Community-based SNMPv2</i>	—
RFC 1905, <i>Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMPv2)</i>	—
RFC 1907, <i>Management Information Base for Version 2 of the Simple Network Management Protocol (SNMPv2)</i>	—
RFC 2011, <i>SNMPv2 Management Information Base for the Internet Protocol Using SMIv2</i>	—
RFC 2012, <i>SNMPv2 Management Information Base for the Transmission Control Protocol Using SMIv2</i>	—
RFC 2013, <i>SNMPv2 Management Information Base for the User Datagram Protocol Using SMIv2</i>	—
RFC 2233, <i>The Interfaces Group MIB Using SMIv2</i>	NOTE: RFC 2233 has been replaced by RFC 2863. However, Junos OS supports both RFC 2233 and RFC 2863.
RFC 2287, <i>Definitions of System-Level Managed Objects for Applications</i>	<p>The following objects are supported:</p> <ul style="list-style-type: none"> sysApplInstallPkgTable sysApplInstallElmtTable sysApplElmtRunTable sysApplMapTable
RFC 2570, <i>Introduction to Version 3 of the Internet-standard Network Management Framework</i>	—
RFC 2571, <i>An Architecture for Describing SNMP Management Frameworks</i> (read-only access)	NOTE: RFC 2571 has been replaced by RFC 3411. However, Junos OS supports both RFC 2571 and RFC 3411.
RFC 2572, <i>Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)</i> (read-only access)	NOTE: RFC 2572 has been replaced by RFC 3412. However, Junos OS supports both RFC 2572 and RFC 3412.

Table 98: Standard MIBs Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis (*continued*)

RFC	Additional Information
RFC 2576, <i>Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework</i>	NOTE: RFC 2576 has been replaced by RFC 3584. However, Junos OS supports both RFC 2576 and RFC 3584.
RFC 2578, <i>Structure of Management Information Version 2 (SMIv2)</i>	—
RFC 2579, <i>Textual Conventions for SMIv2</i>	—
RFC 2580, <i>Conformance Statements for SMIv2</i>	—
RFC 2665, <i>Definitions of Managed Objects for the Ethernet-like Interface Types</i>	—
RFC 2787, <i>Definitions of Managed Objects for the Virtual Router Redundancy Protocol</i>	Support does not include row creation, the Set operation, and the vrrpStatsPacketLengthErrors object.
RFC 2790, <i>Host Resources MIB</i>	Support is limited to the following objects: <ul style="list-style-type: none"> Only hrStorageTable. The file systems <code>/</code>, <code>/config</code>, <code>/var</code>, and <code>/tmp</code> always return the same index number. When SNMP restarts, the index numbers for the remaining file systems might change. Only the objects of the hrSystem and hrSWInstalled groups.
RFC 2819, <i>Remote Network Monitoring Management Information Base</i>	The following objects are supported: <ul style="list-style-type: none"> etherStatsTable (for Ethernet interfaces only), alarmTable, eventTable, and logTable. historyControlTable and etherHistoryTable (except the etherHistoryUtilization object).
RFC 2863, <i>The Interfaces Group MIB</i>	NOTE: RFC 2233 has been replaced by RFC 2863. However, Junos OS supports both RFC 2233 and RFC 2863.
RFC 2932, <i>IPv4 Multicast Routing MIB</i>	—
RFC 2933, <i>Internet Group Management Protocol (IGMP) MIB</i>	—
RFC 2934, <i>Protocol Independent Multicast MIB for IPv4</i>	In Junos OS, RFC 2934 is implemented based on a draft version, <i>pimmib.mib</i> , of the now standard RFC.
RFC 3410, <i>Introduction and Applicability Statements for Internet Standard Management Framework</i>	—

Table 98: Standard MIBs Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis (*continued*)

RFC	Additional Information
RFC 3411, <i>An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks</i>	NOTE: RFC 3411 replaces RFC 2571. However, Junos OS supports both RFC 3411 and RFC 2571.
RFC 3412, <i>Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)</i>	NOTE: RFC 3412 replaces RFC 2572. However, Junos OS supports both RFC 3412 and RFC 2572.
RFC 3413, <i>Simple Network Management Protocol (SNMP) Applications</i>	All MIBs are supported except for the Proxy MIB.
RFC 3414, <i>User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)</i>	—
RFC 3415, <i>View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)</i>	—
RFC 3416, <i>Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)</i>	NOTE: RFC 3416 replaces RFC 1905, which was supported in earlier versions of Junos OS.
RFC 3417, <i>Transport Mappings for the Simple Network Management Protocol (SNMP)</i>	—
RFC 3418, <i>Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)</i>	NOTE: RFC 3418 replaces RFC 1907, which was supported in earlier versions of Junos OS.
RFC 3584, <i>Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework</i>	—
RFC 3826, <i>The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model</i>	—

Table 98: Standard MIBs Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis (*continued*)

RFC	Additional Information
RFC 4188, <i>Definitions of Managed Objects for Bridges</i>	<p>The QFX3500 and QFX3600 switches support 802.1D STP (1998) and the following subtrees and objects only:</p> <ul style="list-style-type: none"> dot1dTp subtree—dot1dTpFdbAddress, dot1dTpFdbPort, and dot1dTpFdbStatus objects from the dot1dTpFdbTable table. dot1dBase subtree—dot1dBasePort and dot1dBasePortIfIndex objects from the dot1dBasePortTable table. <p>NOTE: On QFX3500 and QFX3600 switches, the dot1dTpFdbTable table is populated only with MAC addresses learned on the default VLAN. To see the MAC addresses of all VLANs, specify the dot1qTpFdbTable table (RFC 4363b, <i>Q-Bridge VLAN MIB</i>) when you issue the show snmp mib walk command.</p> <p>Not supported on OCX Series devices.</p>
RFC 4293, <i>Management Information Base for the Internet Protocol (IP)</i>	Supports the ipAddrTable table only.
RFC 4318, <i>Definitions of Managed Objects for Bridges with Rapid Spanning Tree Protocol</i>	<p>Supports 802.1w and 802.1t extensions for RSTP.</p> <p>Not supported on OCX Series devices.</p>
RFC 4363b, <i>Q-Bridge VLAN MIB</i>	<p>NOTE: On QFX3500 and QFX3600 switches, the dot1dTpFdbTable table (RFC 4188, <i>Definitions of Managed Objects for Bridges</i>) is populated only with MAC addresses learned on the default VLAN. To see the MAC addresses of all VLANs, specify the dot1qTpFdbTable table (in this MIB) when you issue the show snmp mib walk command.</p> <p>Not supported on OCX Series devices.</p>
RFC 4444, <i>IS-IS MIB</i>	—
Internet Assigned Numbers Authority, <i>IANAiftype Textual Convention MIB</i> (referenced by RFC 2233)	See http://www.iana.org/assignments/ianaiftype-mib .
Internet draft draft-reeder-snmppv3-usm-3desede-00.txt, <i>Extension to the User-Based Security Model (USM) to Support Triple-DES EDE in 'Outside' CBC Mode</i>	—
Internet draft draft-ietf-idmr-igmp-mib-13.txt, <i>Internet Group Management Protocol (IGMP) MIB</i>	—
ESO Consortium MIB	<p>NOTE: The ESO Consortium MIB has been replaced by RFC 3826. See http://www.snmp.com/eso/.</p>

Table 99: Juniper Networks Enterprise-Specific MIBs Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis

MIB	Description
Alarm MIB (mib-jnx-chassis-alarm)	<p>Provides support for alarms from the switch.</p> <p>For a downloadable version of this MIB, see http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-chassis-alarm.txt.</p> <p>For more information, see <i>Alarm MIB</i>.</p>
Analyzer MIB (mib-jnx-analyzer)	<p>Contains analyzer and remote analyzer data related to port mirroring.</p> <p>For a downloadable version of this MIB, see http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-analyzer.txt.</p> <p>For more information, see <i>Analyzer MIB</i>.</p> <p>Not supported on OCX Series devices.</p>
Chassis MIB (mib-jnx-chassis)	<p>Provides support for environmental monitoring (power supply state, board voltages, fans, temperatures, and airflow) and inventory support for the chassis, Flexible PIC Concentrators (FPCs), and PICs.</p> <p>NOTE: The jnxLEDTable table has been deprecated.</p> <p>For a downloadable version of this MIB, see http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-chassis.txt.</p> <p>For more information, see <i>Chassis MIBs</i>.</p>
Chassis Definitions for Router Model MIB (mib-jnx-chas-defines)	<p>Contains the object identifiers (OIDs) that are used by the Chassis MIB to identify routing and switching platforms and chassis components. The Chassis MIB provides information that changes often, whereas the Chassis Definitions for Router Model MIB provides information that changes less often.</p> <p>For a downloadable version of this MIB, see http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-chas-defines.txt.</p> <p>For more information, see <i>Chassis MIBs</i>.</p>
Class-of-Service MIB (mib-jnx-cos)	<p>Provides support for monitoring interface output queue statistics per interface and per forwarding class.</p> <p>For a downloadable version of this MIB, see http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-cos.txt.</p> <p>For more information, see <i>Class-of-Service MIB</i>.</p>

Table 99: Juniper Networks Enterprise-Specific MIBs Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis (*continued*)

MIB	Description
Configuration Management MIB (mib-jnx-cfgmgmt)	<p>Provides notification for configuration changes and rescue configuration changes in the form of SNMP traps. Each trap contains the time at which the configuration change was committed, the name of the user who made the change, and the method by which the change was made.</p> <p>A history of the last 32 configuration changes is kept in jnxCmChgEventTable.</p> <p>For a downloadable version of this MIB, see http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-cfgmgmt.txt.</p> <p>For more information, see <i>Configuration Management MIB</i>.</p>
Ethernet MAC MIB (mib-jnx-mac)	<p>Monitors media access control (MAC) statistics on Gigabit Ethernet intelligent queuing (IQ) interfaces. It collects MAC statistics; for example, inoctets, inframes, outoctets, and outframes on each source MAC address and virtual LAN (VLAN) ID for each Ethernet port.</p> <p>For a downloadable version of this MIB, see http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-mac.txt.</p> <p>For more information, see <i>Ethernet MAC MIB</i>.</p> <p>Not supported on OCX Series devices.</p>
Event MIB (mib-jnx-event)	<p>Defines a generic trap that can be generated using an operations script or event policy. This MIB provides the ability to specify a system log string and raise a trap if that system log string is found.</p> <p>In Junos OS release 13.2X51-D10 or later, if you configured an event policy to raise a trap when a new SNMP trap target is added, the SNMPD_TRAP_TARGET_ADD_NOTICE trap is generated with information about the new target.</p> <p>For a downloadable version of this MIB, see http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-event.txt.</p> <p>For more information, see <i>Event MIB</i>.</p>
Firewall MIB (mib-jnx-firewall)	<p>Provides support for monitoring firewall filter counters.</p> <p>For a downloadable version of this MIB, see http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-firewall.txt.</p> <p>For more information, see <i>Firewall MIB</i>.</p>

Table 99: Juniper Networks Enterprise-Specific MIBs Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis (*continued*)

MIB	Description
Host Resources MIB (mib-jnx-hostresources)	<p>Extends the hrStorageTable object, providing a measure of the usage of each file system on the switch as a percentage. Previously, the objects in the hrStorageTable measured the usage in allocation units—hrStorageUsed and hrStorageAllocationUnits—only. Using the percentage measurement, you can more easily monitor and apply thresholds on usage.</p> <p>For a downloadable version of this MIB, see http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-hostresources.txt.</p> <p>For more information, see <i>Host Resources MIB</i>.</p>
Interface MIB (Extensions) (mib-jnx-if-extensions)	<p>Extends the standard ifTable (RFC 2863) with additional statistics and Juniper Networks enterprise-specific chassis information in the ifJnxTable and ifChassisTable tables.</p> <p>For a downloadable version of this MIB, see http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-if-extensions.txt.</p> <p>For more information, see <i>Interface MIB</i>.</p>
MPLS MIB (mib-jnx-mpls)	<p>Provides MPLS information and defines MPLS notifications.</p> <p>NOTE: This MIB is not supported on the QFX5100 switch.</p> <p>For a downloadable version of this MIB, see http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-mpls.txt.</p> <p>For more information, see <i>MPLS MIB</i>.</p>
MPLS LDP MIB (mib-jnx-mpls-ldp)	<p>Contains object definitions as described in RFC 3815, <i>Definitions of Managed Objects for the Multiprotocol Label Switching (MPLS), Label Distribution Protocol (LDP)</i>.</p> <p>NOTE: This MIB is not supported on the QFX5100 switch.</p> <p>For a downloadable version of this MIB, see http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-mpls-ldp.txt.</p> <p>For more information, see <i>MPLS LDP MIB</i>.</p>
Ping MIB (mib-jnx-ping)	<p>Extends the standard Ping MIB control table (RFC 2925). Items in this MIB are created when entries are created in pingCtlTable of the Ping MIB. Each item is indexed exactly as it is in the Ping MIB.</p> <p>For a downloadable version of this MIB, see http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-ping.txt.</p> <p>For more information, see <i>PING MIB</i>.</p>

Table 99: Juniper Networks Enterprise-Specific MIBs Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis (*continued*)

MIB	Description
RMON Events and Alarms MIB (mib-jnx-rmon)	<p>Supports Junos OS extensions to the standard Remote Monitoring (RMON) Events and Alarms MIB (RFC 2819). The extension augments the alarmTable object with additional information about each alarm. Two additional traps are also defined to indicate when problems are encountered with an alarm.</p> <p>For a downloadable version of this MIB, see http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-rmon.txt.</p> <p>For more information, see <i>RMON Events and Alarms MIB</i>.</p>
Structure of Management Information MIB (mib-jnx-smi)	<p>Explains how the Juniper Networks enterprise-specific MIBs are structured.</p> <p>For a downloadable version of this MIB, see http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-smi.txt.</p> <p>For more information, see <i>Structure of Management Information MIB</i>.</p>
System Log MIB (mib-jnx-syslog)	<p>Enables notification of an SNMP trap-based application when an important system log message occurs.</p> <p>For a downloadable version of this MIB, see http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-syslog.txt.</p> <p>For more information, see <i>System Log MIB</i>.</p>
Utility MIB (mib-jnx-util)	<p>Provides you with SNMP MIB container objects of the following types: 32-bit counters, 64-bit counters, signed integers, unsigned integers, and octet strings. You can use these objects to store data that can be retrieved using other SNMP operations.</p> <p>For a downloadable version of this MIB, see http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-util.txt.</p> <p>For more information, see “Utility MIB” on page 1460 and “Using the Enterprise-Specific Utility MIB to Enhance SNMP Coverage” on page 1521.</p>
VLAN MIB (mib-jnx-vlan)	<p>Contains information about prestandard IEEE 802.10 VLANs and their association with LAN emulation clients.</p> <p>For a downloadable version of this MIB, see http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-vlan.txt.</p> <p>For more information, see <i>VLAN MIB</i>.</p> <p>Not supported on OCX Series devices.</p>

MIBs Supported on QFabric Systems

The QFabric systems support both standard MIBs and Juniper Networks enterprise-specific MIBs. For more information, see:

- [Table 100](#) for standard MIBs.
- [Table 101](#) for Juniper Networks enterprise-specific MIBs.

Table 100: Standard MIBs Supported on QFabric Systems

RFC	Additional Information
RFC 1155, <i>Structure and Identification of Management Information for TCP/IP-based Internets</i>	—
RFC 1157, <i>A Simple Network Management Protocol (SNMP)</i>	—
RFC 1212, <i>Concise MIB Definitions</i>	—
RFC 1213, <i>Management Information Base for Network Management of TCP/IP-Based Internets: MIB-II</i>	<p>The following areas are supported:</p> <ul style="list-style-type: none"> • MIB II and its SNMP version 2 derivatives, including: <ul style="list-style-type: none"> • Statistics counters • IP, except for ipRouteTable, which has been replaced by ipCidrRouteTable (RFC 2096, <i>IP Forwarding Table MIB</i>) • ipAddrTable • SNMP management • Interface management • SNMPv1 Get, GetNext requests, and version 2 GetBulk request • Junos OS-specific secured access list • Master configuration keywords • Reconfigurations upon SIGHUP
RFC 1215, <i>A Convention for Defining Traps for use with the SNMP</i>	Support is limited to MIB II SNMP version 1 traps and version 2 notifications.
RFC 1286, <i>Definitions of Managed Objects for Bridges</i>	—
RFC 1901, <i>Introduction to Community-based SNMPv2</i>	—
RFC 1905, <i>Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMPv2)</i>	—
RFC 1907, <i>Management Information Base for Version 2 of the Simple Network Management Protocol (SNMPv2)</i>	—

Table 100: Standard MIBs Supported on QFabric Systems (*continued*)

RFC	Additional Information
RFC 2011, <i>SNMPv2 Management Information Base for the Internet Protocol Using SMIv2</i>	NOTE: On the QFabric system, for the SNMP mibwalk request to work, you must configure the IP address of at least one interface besides the management Ethernet interfaces (me0 and me1) in the Director group.
RFC 2012, <i>SNMPv2 Management Information Base for the Transmission Control Protocol Using SMIv2</i>	—
RFC 2013, <i>SNMPv2 Management Information Base for the User Datagram Protocol Using SMIv2</i>	—
RFC 2233, <i>The Interfaces Group MIB Using SMIv2</i>	<p>NOTE: RFC 2233 has been replaced by RFC 2863. However, Junos OS supports both RFC 2233 and RFC 2863.</p> <p>NOTE: The QFabric system supports the following objects only: ifNumber, ifTable, and ifxTable.</p>
RFC 2571, <i>An Architecture for Describing SNMP Management Frameworks</i> (read-only access)	NOTE: RFC 2571 has been replaced by RFC 3411. However, Junos OS supports both RFC 2571 and RFC 3411.
RFC 2572, <i>Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)</i> (read-only access)	NOTE: RFC 2572 has been replaced by RFC 3412. However, Junos OS supports both RFC 2572 and RFC 3412.
RFC 2576, <i>Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework</i>	NOTE: RFC 2576 has been replaced by RFC 3584. However, Junos OS supports both RFC 2576 and RFC 3584.
RFC 2578, <i>Structure of Management Information Version 2 (SMIv2)</i>	—
RFC 2579, <i>Textual Conventions for SMIv2</i>	—
RFC 2580, <i>Conformance Statements for SMIv2</i>	—

Table 100: Standard MIBs Supported on QFabric Systems (*continued*)

RFC	Additional Information
RFC 2665, <i>Definitions of Managed Objects for the Ethernet-like Interface Types</i>	<p>The QFabric system supports the following tables only:</p> <ul style="list-style-type: none"> dot3StatsTable—There is one row with statistics for each Ethernet-like interface in the QFabric system. The dot3StatsIndex is an interface index that is unique across the system. dot3ControlTable—There is one row in this table for each Ethernet-like interface in the QFabric system that implements the MAC control sublayer. OIDs supported are dot3ControlFunctionsSupported and dot3ControlInUnknownOpcode. dot3PauseTable—There is one row in this table for each Ethernet-like interface in the QFabric system that supports the MAC control PAUSE function. OIDs supported are dot3PauseAdminMode, dot3PauseOperMode, dot3InPauseFrames, and dot3OutPauseFrames. <p>NOTE: Scalar variables are not supported on the QFabric system.</p>
RFC 2863, <i>The Interfaces Group MIB</i>	<p>NOTE: RFC 2233 has been replaced by RFC 2863. However, Junos OS supports both RFC 2233 and RFC 2863.</p> <p>NOTE: The QFabric system supports the following objects only: ifNumber, ifTable, and ifxTable.</p>
RFC 2933, <i>Internet Group Management Protocol (IGMP) MIB</i>	—
RFC 3410, <i>Introduction and Applicability Statements for Internet Standard Management Framework</i>	—
RFC 3411, <i>An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks</i>	NOTE: RFC 3411 replaces RFC 2571. However, Junos OS supports both RFC 3411 and RFC 2571.
RFC 3412, <i>Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)</i>	NOTE: RFC 3412 replaces RFC 2572. However, Junos OS supports both RFC 3412 and RFC 2572.
RFC 3416, <i>Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)</i>	NOTE: RFC 3416 replaces RFC 1905, which was supported in earlier versions of Junos OS.
RFC 3417, <i>Transport Mappings for the Simple Network Management Protocol (SNMP)</i>	—
RFC 3418, <i>Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)</i>	NOTE: RFC 3418 replaces RFC 1907, which was supported in earlier versions of Junos OS.
RFC 3584, <i>Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework</i>	—

Table 100: Standard MIBs Supported on QFabric Systems (*continued*)

RFC	Additional Information
RFC 4188, <i>Definitions of Managed Objects for Bridges</i>	<p>The QFabric system support is limited to the following objects:</p> <ul style="list-style-type: none"> Under the dot1dBase OID, the dot1dBasePortTable table supports only the first two columns in the table: dot1dBasePort and dot1dBasePortIfIndex. The system does not implement the optional traps supporting dot1dNotifications (dot1dBridge 0). Under the dot1dStp OID, supports only the dot1dStpPortTable table. Does not support the scalar variables under dot1dStp. The system does not support scalar variables under dot1dTp, but under that, the dot1dTpFdbTable table is supported (dot1dBridge 4). For OIDs with tables support only, scalar values that are returned by the SNMP agent may not be meaningful and are therefore not recommended for use. <p>Not supported on OCX Series devices.</p>
RFC 4293, <i>Management Information Base for the Internet Protocol (IP)</i>	<p>Supports the ipAddrTable table only.</p> <p>On the QFabric system, supported objects in the ipAddrTable table include: ipAdEntAddr, ipAdEntIfIndex, ipAdEntNetMask, ipAdEntBcastAddr, and ipAdEntReasmMaxSize.</p> <p>NOTE: On the QFabric system, for the SNMP mibwalk request to work, you must configure the IP address of at least one interface besides the management Ethernet interfaces (me0 and me1) in the Director group.</p>
RFC 4363b, <i>Q-Bridge VLAN MIB</i>	<p>The QFabric system supports the following tables only:</p> <ul style="list-style-type: none"> dot1qTpFdbTable dot1qVlanStaticTable dot1qPortVlanTable dot1qFdbTable <p>Not supported on OCX Series devices.</p>



NOTE: QFabric-specific MIBs are not supported on OCX Series devices.

Table 101: Juniper Networks Enterprise-Specific MIBs Supported on QFabric Systems

MIB	Description
Analyzer MIB (mib-jnx-analyzer)	<p>Contains analyzer and remote analyzer data related to port mirroring.</p> <p>The QFabric system supports:</p> <ul style="list-style-type: none"> Analyzer table—jnxAnalyzerName, jnxMirroringRatio, jnxLossPriority. Analyzer input table—jnxAnalyzerInputValue, jnxAnalyzerInputOption, jnxAnalyzerInputType. Analyzer output table—jnxAnalyzerOutputValue, jnxAnalyzerOutputType. <p>For a downloadable version of this MIB, see http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-analyzer.txt.</p> <p>For more information, see <i>Analyzer MIB</i>.</p>
Chassis MIB (mib-jnx-chassis)	<p>NOTE: The Chassis MIB has been deprecated for the QFabric system. We recommend that you use the Fabric Chassis MIB (mib-jnx-fabric-chassis) for information about the QFabric system.</p>
Class-of-Service MIB (mib-jnx-cos)	<p>Provides support for monitoring interface output queue statistics per interface and per forwarding class.</p> <p>The QFabric system supports the following tables and objects:</p> <ul style="list-style-type: none"> Jnxcosifstatflagtable—jnxCosIfstatFlags and jnxCosIfIndex. Jnxcosqstattable—jnxCosQstatTxedPkts, jnxCosQstatTxedPktRate, jnxCosQstatTxedBytes, and jnxCosQstatTxedByteRate. Jnxcosfcidtable—jnxCosFcIdToFcName. Jnxcosfctable—jnxCosFcQueueNr. <p>The QFabric system does not support any traps for this MIB.</p> <p>For a downloadable version of this MIB, see http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-cos.txt.</p> <p>For more information, see <i>Class-of-Service MIB</i>.</p>

Table 101: Juniper Networks Enterprise-Specific MIBs Supported on QFabric Systems (*continued*)

MIB	Description
Configuration Management MIB (mib-jnx-cfgmgmt)	<p>Provides notification for configuration changes and rescue configuration changes in the form of SNMP traps. Each trap contains the time at which the configuration change was committed, the name of the user who made the change, and the method by which the change was made.</p> <p>A history of the last 32 configuration changes is kept in jnxCmChgEventTable.</p> <p>NOTE: On the QFabric system, these conditions apply:</p> <ul style="list-style-type: none"> • All scalar variables under the jnxCmCfgChg table are supported. • Supported scalar OIDs are jnxCmCfgChgLatestIndex, jnxCmCfgChgLatestTime, jnxCmCfgChgLatestDate, jnxCmCfgChgLatestSource, jnxCmCfgChgLatestUser, and jnxCmCfgChgMaxEventEntries. • Scalar variables under the jnxCmRescueChg table are not supported. <p>For a downloadable version of this MIB, see http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-cfgmgmt.txt.</p> <p>For more information, see <i>Configuration Management MIB</i>.</p>
Fabric Chassis MIB (mib-jnx-fabric-chassis)	<p>Provides hardware information about the QFabric system and its component devices. This MIB is based on the Juniper Networks enterprise-specific Chassis MIB but adds another level of indexing that provides information for QFabric system component devices.</p> <p>For a downloadable version of this MIB, see http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-fabric-chassis.txt.</p> <p>For more information, see <i>Fabric Chassis MIB</i>.</p>
Interface MIB (Extensions) (mib-jnx-if-extensions)	<p>Extends the standard ifTable (RFC 2863) with additional statistics and Juniper Networks enterprise-specific chassis information in the ifJnxTable and ifChassisTable tables.</p> <p>NOTE: On the QFabric system, scalar variables are not supported.</p> <p>For a downloadable version of this MIB, see http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-if-extensions.txt.</p> <p>For more information, see <i>Interface MIB</i>.</p>
Power Supply Unit MIB (mib-jnx-power-supply-unit)	<p>Provides support for environmental monitoring of the power supply unit for the Interconnect device of the QFabric system.</p> <p>For a downloadable version of this MIB, see http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-power-supply-unit.txt.</p> <p>For more information, see <i>Power Supply Unit MIB</i>.</p> <p>NOTE: On the QFabric system, scalar variables for the jnxPsuObjects 1 object ID in the jnxPsuScalars table are not supported.</p>

Table 101: Juniper Networks Enterprise-Specific MIBs Supported on QFabric Systems (*continued*)

MIB	Description
QFabric MIB (jnx-qf-smi)	<p>Explains how the Juniper Networks enterprise-specific QFabric MIBs are structured. Defines the MIB objects that are reported by the QFabric system and the contents of the traps that can be issued by the QFabric system.</p> <p>For a downloadable version of this MIB, see http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-qf-smi.txt.</p>
Utility MIB (mib-jnx-util)	<p>Provides you with SNMP MIB container objects of the following types: 32-bit counters, 64-bit counters, signed integers, unsigned integers, and octet strings. You can use these objects to store data that can be retrieved using other SNMP operations.</p> <p>For a downloadable version of this MIB, see http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-util.txt.</p> <p>For more information, see “Utility MIB” on page 1460 and “Using the Enterprise-Specific Utility MIB to Enhance SNMP Coverage” on page 1521.</p>

- Related Documentation**
- *SNMP MIBs and Traps Reference*
 - *Understanding the Implementation of SNMP*
 - *Understanding the Implementation of SNMP on the QFabric System*
 - *SNMP Traps Support on page 1484*

SNMP Traps Support

The QFX Series standalone switches, QFX Series Virtual Chassis, and QFabric systems support standard SNMP traps and Juniper Networks enterprise-specific traps.

For more information, see:

- *SNMP Traps Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis on page 1484*
- *SNMP Traps Supported on QFabric Systems on page 1493*

SNMP Traps Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis

QFX Series standalone switches and QFX Series Virtual Chassis support SNMPv1 and v2 traps. For more information, see:

- *SNMPv1 Traps on page 1484*
- *SNMPv2 Traps on page 1489*

SNMPv1 Traps

QFX Series standalone switches and QFX Series Virtual Chassis support both standard SNMPv1 traps and Juniper Networks enterprise-specific SNMPv1 traps. See:

- [Table 102](#) for standard SNMPv1 traps.
- [Table 103](#) for enterprise-specific SNMPv1 traps.

The traps are organized first by trap category and then by trap name. The system logging severity levels are listed for those traps that have them. Traps that do not have corresponding system logging severity levels are marked with an en dash (–).

Table 102: Standard SNMP Version 1 Traps Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis

Defined in	Trap Name	Enterprise ID	Generic Trap Number	Specific Trap Number	System Logging Severity Level	Syslog Tag
Link Notifications						
RFC 1215, <i>Conventions for Defining Traps for Use with the SNMP</i>	linkDown	1.3.6.1.4.1.2636	2	0	Warning	SNMP_TRAP_LINK_DOWN
	linkUp	1.3.6.1.4.1.2636	3	0	Info	SNMP_TRAP_LINK_UP
Remote Operations Notifications						
RFC 2925, <i>Definitions of Managed Objects for Remote Ping, Traceroute, and Lookup Operations</i>	pingProbeFailed	1.3.6.1.2.1.80.0	6	1	Info	SNMP_TRAP_PING_PROBE_FAILED
	pingTestFailed	1.3.6.1.2.1.80.0	6	2	Info	SNMP_TRAP_PING_TEST_FAILED
	pingTestCompleted	1.3.6.1.2.1.80.0	6	3	Info	SNMP_TRAP_PING_TEST_COMPLETED
	traceRoutePathChange	1.3.6.1.2.1.81.0	6	1	Info	SNMP_TRAP_TRACE_ROUTE_PATH_CHANGE
	traceRouteTestFailed	1.3.6.1.2.1.81.0	6	2	Info	SNMP_TRAP_TRACE_ROUTE_TEST_FAILED
	traceRouteTestCompleted	1.3.6.1.2.1.81.0	6	3	Info	SNMP_TRAP_TRACE_ROUTE_TEST_COMPLETED
RMON Alarms						
RFC 2819a, <i>RMON MIB</i>	fallingAlarm	1.3.6.1.2.1.16	6	2	–	–
	risingAlarm	1.3.6.1.2.1.16	6	1	–	–
Routing Notifications						

Table 102: Standard SNMP Version 1 Traps Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis (*continued*)

Defined in	Trap Name	Enterprise ID	Generic Trap Number	Specific Trap Number	System Logging Severity Level	Syslog Tag
<i>BGP 4 MIB</i>	bgpEstablished	1.3.6.1.2.1.15.7	6	1	–	–
	bgpBackwardTransition	1.3.6.1.2.1.15.7	6	2	–	–
<i>OSPF TRAP MIB</i>	ospfVirtIfStateChange	1.3.6.1.2.1.14.16.2	6	1	–	–
	ospfNbrStateChange	1.3.6.1.2.1.14.16.2	6	2	–	–
	ospfVirtNbrStateChange	1.3.6.1.2.1.14.16.2	6	3	–	–
	ospfIfConfigError	1.3.6.1.2.1.14.16.2	6	4	–	–
	ospfVirtIfConfigError	1.3.6.1.2.1.14.16.2	6	5	–	–
	ospfIfAuthFailure	1.3.6.1.2.1.14.16.2	6	6	–	–
	ospfVirtIfAuthFailure	1.3.6.1.2.1.14.16.2	6	7	–	–
	ospfIfRxBadPacket	1.3.6.1.2.1.14.16.2	6	8	–	–
	ospfVirtIfRxBadPacket	1.3.6.1.2.1.14.16.2	6	9	–	–
	ospfTxRetransmit	1.3.6.1.2.1.14.16.2	6	10	–	–
	ospfVirtIfTxRetransmit	1.3.6.1.2.1.14.16.2	6	11	–	–
	ospfMaxAgeLsa	1.3.6.1.2.1.14.16.2	6	13	–	–
	ospfIfStateChange	1.3.6.1.2.1.14.16.2	6	16	–	–
Startup Notifications						
RFC 1215, <i>Conventions for Defining Traps for Use with the SNMP</i>	authenticationFailure	1.3.6.1.4.1.2636	4	0	Notice	SNMPD_TRAP_GEN_FAILURE
	coldStart	1.3.6.1.4.1.2636	0	0	Critical	SNMPD_TRAP_COLD_START
	warmStart	1.3.6.1.4.1.2636	1	0	Error	SNMPD_TRAP_WARM_START
VRRP Notifications						

Table 102: Standard SNMP Version 1 Traps Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis (*continued*)

Defined in	Trap Name	Enterprise ID	Generic Trap Number	Specific Trap Number	System Logging Severity Level	Syslog Tag
RFC 2787, <i>Definitions of Managed Objects for the Virtual Router Redundancy Protocol</i>	vrpTrapNewMaster	1.3.6.1.2.1.68	6	1	Warning	VRRPD_NEW_MASTER_TRAP
	vrpTrapAuthFailure	1.3.6.1.2.1.68	6	2	Warning	VRRPD_AUTH_FAILURE_TRAP

Table 103: Enterprise-Specific SNMPv1 Traps Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis

Defined in	Trap Name	Enterprise ID	Generic Trap Number	Specific Trap Number	System Logging Severity Level	System Log Tag
Chassis Notifications (Alarm Conditions)						
<i>Chassis MIB</i> (jnx-chassis.mib)	jnxPowerSupplyFailure	1.3.6.1.4.1.2636.4.1	6	1	Warning	CHASSISD_SNMP_TRAP
	jnxFanFailure	1.3.6.1.4.1.2636.4.1	6	2	Critical	CHASSISD_SNMP_TRAP
	jnxOverTemperature	1.3.6.1.4.1.2636.4.1	6	3	Alert	CHASSISD_SNMP_TRAP
	jnxFruRemoval	1.3.6.1.4.1.2636.4.1	6	5	Notice	CHASSISD_SNMP_TRAP
	jnxFruInsertion	1.3.6.1.4.1.2636.4.1	6	6	Notice	CHASSISD_SNMP_TRAP
	jnxFruPowerOff	1.3.6.1.4.1.2636.4.1	6	7	Notice	CHASSISD_SNMP_TRAP
	jnxFruPowerOn	1.3.6.1.4.1.2636.4.1	6	8	Notice	CHASSISD_SNMP_TRAP
	jnxFruFailed	1.3.6.1.4.1.2636.4.1	6	9	Warning	CHASSISD_SNMP_TRAP
	jnxFruOffline	1.3.6.1.4.1.2636.4.1	6	10	Notice	CHASSISD_SNMP_TRAP
	jnxFruOnline	1.3.6.1.4.1.2636.4.1	6	11	Notice	CHASSISD_SNMP_TRAP

Table 103: Enterprise-Specific SNMPv1 Traps Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis (*continued*)

Defined in	Trap Name	Enterprise ID	Generic Trap Number	Specific Trap Number	System Logging Severity Level	System Log Tag
	jnxFruCheck	1.3.6.1.4.1.2636.4.1	6	12	Warning	CHASSISD_SNMP_TRAP
	jnxPowerSupplyOk	1.3.6.1.4.1.2636.4.2	6	1	Critical	CHASSISD_SNMP_TRAP
	jnxFanOK	1.3.6.1.4.1.2636.4.2	6	2	Critical	CHASSISD_SNMP_TRAP
	jnxTemperatureOK	1.3.6.1.4.1.2636.4.2	6	3	Alert	CHASSISD_SNMP_TRAP
Configuration Notifications						
<i>Configuration Management MIB</i> (jnx-configmgmt.mib)	jnxCmCfgChange	1.3.6.1.4.1.2636.4.5	6	1	—	—
	jnxCmRescueChange	1.3.6.1.4.1.2636.4.5	6	2	—	—
Remote Operations						
<i>Ping MIB</i> (jnx-ping.mib)	jnxPingRttThresholdExceeded	1.3.6.1.4.1.2636.4.9	6	1	—	—
	jnxPingRttStdDevThreshold Exceeded	1.3.6.1.4.1.2636.4.9	6	2	—	—
	jnxPingRttJitterThreshold Exceeded	1.3.6.1.4.1.2636.4.9	6	3	—	—
	jnxPingEgressThreshold Exceeded	1.3.6.1.4.1.2636.4.9	6	4	—	—
	jnxPingEgressStdDev ThresholdExceeded	1.3.6.1.4.1.2636.4.9	6	5	—	—
	jnxPingEgressJitterThreshold Exceeded	1.3.6.1.4.1.2636.4.9	6	6	—	—
	jnxPingIngressThreshold Exceeded	1.3.6.1.4.1.2636.4.9	6	7	—	—
	jnxPingIngressStddevThreshold Exceeded	1.3.6.1.4.1.2636.4.9	6	8	—	—
	jnxPingIngressJitterThreshold Exceeded	1.3.6.1.4.1.2636.4.9	6	9	—	—

Table 103: Enterprise-Specific SNMPv1 Traps Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis (*continued*)

Defined in	Trap Name	Enterprise ID	Generic Trap Number	Specific Trap Number	System Logging Severity Level	System Log Tag
RMON Alarms						
<i>RMON MIB</i> (jnx-rmon. mib)	jnxRmonAlarmGetFailure	1.3.6.1.4.1.2636.4.3	6	1	–	–
	jnxRmonGetOk	1.3.6.1.4.1.2636.4.3	6	2	–	–

SNMPv2 Traps

- [Table 104](#) lists the standard SNMP traps
- [Table 105](#) lists the Juniper Networks enterprise-specific traps

Table 104: Standard SNMPv2 Traps Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis

Defined in	Trap Name	SNMP Trap OID	System Logging Severity Level	Syslog Tag
Link Notifications				
<i>RFC 2863, The Interfaces Group MIB</i>	linkDown	1.3.6.1.6.3.1.1.5.3	Warning	SNMP_TRAP_LINK_DOWN
	linkUp	1.3.6.1.6.3.1.1.5.4	Info	SNMP_TRAP_LINK_UP
Remote Operations Notifications				
<i>RFC 2925, Definitions of Managed Objects for Remote Ping, Traceroute, and Lookup Operations</i>	pingProbeFailed	1.3.6.1.2.1.80.0.1	Info	SNMP_TRAP_PING_PROBE_FAILED
	pingTestFailed	1.3.6.1.2.1.80.0.2	Info	SNMP_TRAP_PING_TEST_FAILED
	pingTestCompleted	1.3.6.1.2.1.80.0.3	Info	SNMP_TRAP_PING_TEST_COMPLETED
	traceRoutePathChange	1.3.6.1.2.1.81.0.1	Info	SNMP_TRAP_TRACE_ROUTE_PATH_CHANGE
	traceRouteTestFailed	1.3.6.1.2.1.81.0.2	Info	SNMP_TRAP_TRACE_ROUTE_TEST_FAILED
	traceRouteTestCompleted	1.3.6.1.2.1.81.0.3	Info	SNMP_TRAP_TRACE_ROUTE_TEST_COMPLETED

Table 104: Standard SNMPv2 Traps Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis (*continued*)

Defined in	Trap Name	SNMP Trap OID	System Logging Severity Level	Syslog Tag
RMON Alarms				
RFC 2819a, <i>RMON MIB</i>	fallingAlarm	1.3.6.1.2.1.16.0.1	–	–
	risingAlarm	1.3.6.1.2.1.16.0.2	–	–
Routing Notifications				
<i>BGP 4 MIB</i>	bgpEstablished	1.3.6.1.2.1.15.7.1	–	–
	bgpBackwardTransition	1.3.6.1.2.1.15.7.2	–	–
<i>OSPF Trap MIB</i>	ospfVirtIfStateChange	1.3.6.1.2.1.14.16.2.1	–	–
	ospfNbrStateChange	1.3.6.1.2.1.14.16.2.2	–	–
	ospfVirtNbrStateChange	1.3.6.1.2.1.14.16.2.3	–	–
	ospfIfConfigError	1.3.6.1.2.1.14.16.2.4	–	–
	ospfVirtIfConfigError	1.3.6.1.2.1.14.16.2.5	–	–
	ospfIfAuthFailure	1.3.6.1.2.1.14.16.2.6	–	–
	ospfVirtIfAuthFailure	1.3.6.1.2.1.14.16.2.7	–	–
	ospfIfRxBadPacket	1.3.6.1.2.1.14.16.2.8	–	–
	ospfVirtIfRxBadPacket	1.3.6.1.2.1.14.16.2.9	–	–
	ospfTxRetransmit	1.3.6.1.2.1.14.16.2.10	–	–
	ospfVirtIfTxRetransmit	1.3.6.1.2.1.14.16.2.11	–	–
	ospfMaxAgeLsa	1.3.6.1.2.1.14.16.2.13	–	–
	ospfIfStateChange	1.3.6.1.2.1.14.16.2.16	–	–
Startup Notifications				
RFC 1907, <i>Management Information Base for Version 2 of the Simple Network Management Protocol (SNMPv2)</i>	coldStart	1.3.6.1.6.3.1.1.5.1	Critical	SNMPD_TRAP_COLD_START
	warmStart	1.3.6.1.6.3.1.1.5.2	Error	SNMPD_TRAP_WARM_START

Table 104: Standard SNMPv2 Traps Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis (*continued*)

Defined in	Trap Name	SNMP Trap OID	System Logging Severity Level	Syslog Tag
	authenticationFailure	1.3.6.1.6.3.1.1.5.5	Notice	SNMPD_TRAP_GEN_FAILURE
VRRP Notifications				
RFC 2787, <i>Definitions of Managed Objects for the Virtual Router Redundancy Protocol</i>	vrrpTrapNewMaster	1.3.6.1.2.1.68.0.1	Warning	VRRPD_NEWMASTER_TRAP
	vrrpTrapAuthFailure	1.3.6.1.2.1.68.0.2	Warning	VRRPD_AUTH_FAILURE_TRAP

Table 105: Enterprise-Specific SNMPv2 Traps Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis

Source MIB	Trap Name	SNMP Trap OID	System Logging Severity Level	System Log Tag
Chassis (Alarm Conditions) Notifications				
<i>Chassis MIB</i> (mib-jnx-chassis)	jnxPowerSupplyFailure	1.3.6.1.4.1.2636.4.1.1	Alert	CHASSISD_SNMP_TRAP
	jnxFanFailure	1.3.6.1.4.1.2636.4.1.2	Critical	CHASSISD_SNMP_TRAP
	jnxOverTemperature	1.3.6.1.4.1.2636.4.1.3	Critical	CHASSISD_SNMP_TRAP
	jnxFruRemoval	1.3.6.1.4.1.2636.4.1.5	Notice	CHASSISD_SNMP_TRAP
	jnxFruInsertion	1.3.6.1.4.1.2636.4.1.6	Notice	CHASSISD_SNMP_TRAP
	jnxFruPowerOff	1.3.6.1.4.1.2636.4.1.7	Notice	CHASSISD_SNMP_TRAP
	jnxFruPowerOn	1.3.6.1.4.1.2636.4.1.8	Notice	CHASSISD_SNMP_TRAP
	jnxFruFailed	1.3.6.1.4.1.2636.4.1.9	Warning	CHASSISD_SNMP_TRAP
	jnxFruOffline	1.3.6.1.4.1.2636.4.1.10	Notice	CHASSISD_SNMP_TRAP

Table 105: Enterprise-Specific SNMPv2 Traps Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis (*continued*)

Source MIB	Trap Name	SNMP Trap OID	System Logging Severity Level	System Log Tag
	jnxFruOnline	1.3.6.1.4.1.2636.4.1.11	Notice	CHASSISD_SNMP_TRAP
	jnxFruCheck	1.3.6.1.4.1.2636.4.1.12	Notice	CHASSISD_SNMP_TRAP
	jnxPowerSupplyOK	1.3.6.1.4.1.2636.4.2.1	Critical	CHASSISD_SNMP_TRAP
	jnxFanOK	1.3.6.1.4.1.2636.4.2.2	Critical	CHASSISD_SNMP_TRAP
	jnxTemperatureOK	1.3.6.1.4.1.2636.4.2.3	Alert	CHASSISD_SNMP_TRAP
Configuration Notifications				
<i>Configuration Management MIB</i> (mib-jnx-cfgmgmt)	jnxCmCfgChange	1.3.6.1.4.1.2636.4.5.0.1	–	–
	jnxCmRescueChange	1.3.6.1.4.1.2636.4.5.0.2	–	–
Remote Operations Notifications				

Table 105: Enterprise-Specific SNMPv2 Traps Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis (*continued*)

Source MIB	Trap Name	SNMP Trap OID	System Logging Severity Level	System Log Tag
<i>Ping MIB</i> (mib-jnx-ping)	jnxPingRttThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.1	–	–
	jnxPingRttStdDevThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.2	–	–
	jnxPingRttJitterThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.3	–	–
	jnxPingEgressThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.4	–	–
	jnxPingEgressStdDevThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.5	–	–
	jnxPingEgressJitterThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.6	–	–
	jnxPingIngressThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.7	–	–
	jnxPingIngressStddevThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.8	–	–
	jnxPingIngressJitterThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.9	–	–
RMON Alarms				
<i>RMON MIB</i> (mib-jnx-rmon)	jnxRmonAlarmGetFailure	1.3.6.1.4.1.2636.4.3.0.1	–	–
	jnxRmonGetOk	1.3.6.1.4.1.2636.4.3.0.2	–	–

SNMP Traps Supported on QFabric Systems

QFabric systems support standard SNMPv2 traps and Juniper Networks enterprise-specific SNMPv2 traps.



NOTE: QFabric systems do not support SNMPv1 traps.

For more information, see:

- [Table 106](#) for standard SNMPv2 traps

- [Table 107](#) for Juniper Networks enterprise-specific SNMPv2 traps

Table 106: Standard SNMPv2 Traps Supported on QFabric Systems

Defined in	Trap Name	SNMP Trap OID	System Logging Severity Level	Syslog Tag
Link Notifications				
RFC 2863, <i>The Interfaces Group MIB</i>	linkDown	1.3.6.1.6.3.1.1.5.3	Warning	SNMP_TRAP_LINK_DOWN
	linkUp	1.3.6.1.6.3.1.1.5.4	Info	SNMP_TRAP_LINK_UP
Startup Notifications				
RFC 1907, <i>Management Information Base for Version 2 of the Simple Network Management Protocol (SNMPv2)</i>	coldStart	1.3.6.1.6.3.1.1.5.1	Critical	SNMPD_TRAP_COLD_START
	warmStart	1.3.6.1.6.3.1.1.5.2	Error	SNMPD_TRAP_WARM_START
	authenticationFailure	1.3.6.1.6.3.1.1.5.5	Notice	SNMPD_TRAP_GEN_FAILURE

Table 107: Enterprise-Specific SNMPv2 Traps Supported on QFabric Systems

Source MIB	Trap Name	SNMP Trap OID	System Logging Severity Level	System Log Tag
<i>Fabric Chassis MIB</i> (mib-jnx-fabric-chassis)	Fabric Chassis (Alarm Conditions) Notifications			
	jnxFabricPowerSupplyFailure	1.3.6.1.4.1.2636.4.19.1	Warning	–
	jnxFabricFanFailure	1.3.6.1.4.1.2636.4.19.2	Critical	–
	jnxFabricOverTemperature	1.3.6.1.4.1.2636.4.19.3	Alert	–
	jnxFabricRedundancySwitchover	1.3.6.1.4.1.2636.4.19.4	Notice	–
	jnxFabricFruRemoval	1.3.6.1.4.1.2636.4.19.5	Notice	–
	jnxFabricFruInsertion	1.3.6.1.4.1.2636.4.19.6	Notice	–
	jnxFabricFruPowerOff	1.3.6.1.4.1.2636.4.19.7	Notice	–
	jnxFabricFruPowerOn	1.3.6.1.4.1.2636.4.19.8	Notice	–
	jnxFabricFruFailed	1.3.6.1.4.1.2636.4.19.9	Warning	–
	jnxFabricFruOffline	1.3.6.1.4.1.2636.4.19.10	Notice	–
	jnxFabricFruOnline	1.3.6.1.4.1.2636.4.19.11	Notice	–
	jnxFabricFruCheck	1.3.6.1.4.1.2636.4.19.12	Warning	–
	jnxFabricFEBSwitchover	1.3.6.1.4.1.2636.4.19.13	Warning	–
	jnxFabricHardDiskFailed	1.3.6.1.4.1.2636.4.19.14	Warning	–
	jnxFabricHardDiskMissing	1.3.6.1.4.1.2636.4.19.15	Warning	–
	jnxFabricBootFromBackup	1.3.6.1.4.1.2636.4.19.16	Warning	–
	Fabric Chassis (Alarm Cleared Conditions) Notifications			
	jnxFabricPowerSupplyOK	1.3.6.1.4.1.2636.4.20.1	Critical	–
	jnxFabricFanOK	1.3.6.1.4.1.2636.4.20.2	Critical	–
	jnxFabricTemperatureOK	1.3.6.1.4.1.2636.4.20.3	Alert	–
	jnxFabricFruOK	1.3.6.1.4.1.2636.4.20.4	–	–

Table 107: Enterprise-Specific SNMPv2 Traps Supported on QFabric Systems (*continued*)

Source MIB	Trap Name	SNMP Trap OID	System Logging Severity Level	System Log Tag
<i>QFabric MIB</i> (mib-jnx-qf-smi)	QFabric MIB Notifications			
	jnxQFabricDownloadIssued	1.3.6.1.4.1.2636.3.42.1.0.1	–	–
	jnxQFabricDownloadFailed	1.3.6.1.4.1.2636.3.42.1.0.2	–	–
	jnxQFabricDownloadSucceeded	1.3.6.1.4.1.2636.3.42.1.0.3	–	–
	jnxQFabricUpgradeIssued	1.3.6.1.4.1.2636.3.42.1.0.4	–	–
	jnxQFabricUpgradeFailed	1.3.6.1.4.1.2636.3.42.1.0.5	–	–
	jnxQFabricUpgradeSucceeded	1.3.6.1.4.1.2636.3.42.1.0.6	–	–
Configuration Notifications				
<i>Configuration Management MIB</i> (mib-jnx-cfgmgmt)	jnxCmCfgChange	1.3.6.1.4.1.2636.4.5.0.1	–	–
	jnxCmRescueChange	1.3.6.1.4.1.2636.4.5.0.2	–	–
Remote Operations Notifications				
<i>Ping MIB</i> (mib-jnx-ping)	jnxPingRttThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.1	–	–
	jnxPingRttStdDevThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.2	–	–
	jnxPingRttJitterThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.3	–	–
	jnxPingEgressThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.4	–	–
	jnxPingEgressStdDevThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.5	–	–
	jnxPingEgressJitterThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.6	–	–
	jnxPingIngressThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.7	–	–
	jnxPingIngressStddevThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.8	–	–
	jnxPingIngressJitterThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.9	–	–

- Related Documentation**
- [SNMP MIBs and Traps Reference](#)
 - [Understanding the Implementation of SNMP](#)
 - [Understanding the Implementation of SNMP on the QFabric System](#)
 - [SNMP MIBs Support on page 1468](#)

MIB Objects for the QFX Series

This topic lists the Juniper Networks enterprise-specific SNMP Chassis MIB definition objects for the QFX Series:

- [QFX Series Standalone Switches on page 1497](#)
- [QFabric Systems on page 1497](#)
- [QFabric System QFX3100 Director Device on page 1498](#)
- [QFabric System QFX3008-I Interconnect Device on page 1498](#)
- [QFabric System QFX3600-I Interconnect Device on page 1498](#)
- [QFabric System Node Devices on page 1499](#)

QFX Series Standalone Switches

```
jnxProductLineQFXSwitch      OBJECT IDENTIFIER ::= { jnxProductLine      82 }
jnxProductNameQFXSwitch      OBJECT IDENTIFIER ::= { jnxProductName      82 }
jnxProductModelQFXSwitch     OBJECT IDENTIFIER ::= { jnxProductModel      82 }
jnxProductVariationQFXSwitch OBJECT IDENTIFIER ::= { jnxProductVariation 82 }
jnxProductQFX3500s          OBJECT IDENTIFIER ::= { jnxProductVariationQFXSwitch 1 }
jnxProductQFX360016QS       OBJECT IDENTIFIER ::= { jnxProductVariationQFXSwitch 2 }
jnxProductQFX350048T4QS     OBJECT IDENTIFIER ::= { jnxProductVariationQFXSwitch 3 }
jnxProductQFX510024Q        OBJECT IDENTIFIER ::= { jnxProductVariationQFXSwitch 4 }
jnxProductQFX510048S6Q      OBJECT IDENTIFIER ::= { jnxProductVariationQFXSwitch 5 }

jnxChassisQFXSwitch          OBJECT IDENTIFIER ::= { jnxChassis          82 }

jnxSlotQFXSwitch             OBJECT IDENTIFIER ::= { jnxSlot             82 }
jnxQFXSwitchSlotFPC          OBJECT IDENTIFIER ::= { jnxSlotQFXSwitch      1 }
jnxQFXSwitchSlotHM           OBJECT IDENTIFIER ::= { jnxSlotQFXSwitch      2 }
jnxQFXSwitchSlotPower        OBJECT IDENTIFIER ::= { jnxSlotQFXSwitch      3 }
jnxQFXSwitchSlotFan          OBJECT IDENTIFIER ::= { jnxSlotQFXSwitch      4 }
jnxQFXSwitchSlotFPB          OBJECT IDENTIFIER ::= { jnxSlotQFXSwitch      5 }

jnxMediaCardSpaceQFXSwitch   OBJECT IDENTIFIER ::= { jnxMediaCardSpace 82 }
jnxQFXSwitchMediaCardSpacePIC OBJECT IDENTIFIER ::= { jnxMediaCardSpaceQFXSwitch 1 }
```

QFabric Systems

```
jnxProductLineQFX3000        OBJECT IDENTIFIER ::= { jnxProductLine 84 }
jnxProductNameQFX3000         OBJECT IDENTIFIER ::= { jnxProductName 84 }
jnxProductModelQFX3000        OBJECT IDENTIFIER ::= { jnxProductModel 84 }
jnxProductVariationQFX3000    OBJECT IDENTIFIER ::= { jnxProductVariation 84 }
jnxProductQFX3000-G           OBJECT IDENTIFIER ::= { jnxProductVariationQFX3000 1 }
jnxProductQFX3000-M           OBJECT IDENTIFIER ::= { jnxProductVariationQFX3000 2 }
jnxChassisQFX3000             OBJECT IDENTIFIER ::= { jnxChassis      84 }
```

QFabric System QFX3100 Director Device

```
jnxProductLineQFX3100 OBJECT IDENTIFIER ::= { jnxProductLine      100 }
jnxProductNameQFX3100 OBJECT IDENTIFIER ::= { jnxProductName      100 }
jnxProductModelQFX3100 OBJECT IDENTIFIER ::= { jnxProductModel    100 }
jnxProductVariationQFX3100 OBJECT IDENTIFIER ::= { jnxProductVariation 100 }
jnxChassisQFX3100      OBJECT IDENTIFIER ::= { jnxChassis         100 }

jnxSlotQFX3100          OBJECT IDENTIFIER ::= { jnxSlot           100 }
jnxQFX3100SlotCPU       OBJECT IDENTIFIER ::= { jnxSlotQFX3100    1 }
jnxQFX3100SlotMemory    OBJECT IDENTIFIER ::= { jnxSlotQFX3100    2 }
jnxQFX3100SlotPower     OBJECT IDENTIFIER ::= { jnxSlotQFX3100    3 }
jnxQFX3100SlotFan       OBJECT IDENTIFIER ::= { jnxSlotQFX3100    4 }
jnxQFX3100SlotHardDisk  OBJECT IDENTIFIER ::= { jnxSlotQFX3100    5 }
jnxQFX3100SlotNIC       OBJECT IDENTIFIER ::= { jnxSlotQFX3100    6 }
```

QFabric System QFX3008-I Interconnect Device

```
jnxProductLineQFXInterconnect OBJECT IDENTIFIER ::= { jnxProductLine      60 }
jnxProductNameQFXInterconnect OBJECT IDENTIFIER ::= { jnxProductName      60 }
jnxProductModelQFXInterconnect OBJECT IDENTIFIER ::= { jnxProductModel    60 }
jnxProductVariationQFXInterconnect OBJECT IDENTIFIER ::= { jnxProductVariation 60 }
jnxProductQFX3008          OBJECT IDENTIFIER ::= { jnxProductVariationQFXInterconnect 1 }
jnxProductQFXC083008       OBJECT IDENTIFIER ::= { jnxProductVariationQFXInterconnect 2 }
jnxProductQFX3008I         OBJECT IDENTIFIER ::= { jnxProductVariationQFXInterconnect 3 }

jnxChassisQFXInterconnect  OBJECT IDENTIFIER ::= { jnxChassis         60 }

jnxSlotQFXInterconnect     OBJECT IDENTIFIER ::= { jnxSlot           60 }
jnxQFXInterconnectSlotFPC  OBJECT IDENTIFIER ::= { jnxSlotQFXInterconnect  1 }
jnxQFXInterconnectSlotHMM OBJECT IDENTIFIER ::= { jnxSlotQFXInterconnect  2 }
jnxQFXInterconnectSlotPower OBJECT IDENTIFIER ::= { jnxSlotQFXInterconnect  3 }
jnxQFXInterconnectSlotFan  OBJECT IDENTIFIER ::= { jnxSlotQFXInterconnect  4 }
jnxQFXInterconnectSlotCBD  OBJECT IDENTIFIER ::= { jnxSlotQFXInterconnect  5 }
jnxQFXInterconnectSlotFPB  OBJECT IDENTIFIER ::= { jnxSlotQFXInterconnect  6 }

jnxMediaCardSpaceQFXInterconnect OBJECT IDENTIFIER ::= { jnxMediaCardSpace 60 }
jnxQFXInterconnectMediaCardSpacePIC OBJECT IDENTIFIER ::= { jnxMediaCardSpaceQFXInterconnect 1 }

jnxMidplaneQFXInterconnect OBJECT IDENTIFIER ::= { jnxBackplane      60 }
```

QFabric System QFX3600-I Interconnect Device

```
jnxProductLineQFXMInterconnect OBJECT IDENTIFIER ::= { jnxProductLine      91 }
jnxProductNameQFXMInterconnect OBJECT IDENTIFIER ::= { jnxProductName      91 }
jnxProductModelQFXMInterconnect OBJECT IDENTIFIER ::= { jnxProductModel    91 }
jnxProductVariationQFXMInterconnect OBJECT IDENTIFIER ::= { jnxProductVariation 91 }
jnxProductQFX3600I         OBJECT IDENTIFIER ::= { jnxProductVariationQFXMInterconnect 1 }

jnxChassisQFXMInterconnect  OBJECT IDENTIFIER ::= { jnxChassis         91 }

jnxSlotQFXMInterconnect     OBJECT IDENTIFIER ::= { jnxSlot           91 }
jnxQFXMInterconnectSlotFPC  OBJECT IDENTIFIER ::= { jnxSlotQFXMInterconnect  1 }
jnxQFXMInterconnectSlotHMM OBJECT IDENTIFIER ::= { jnxSlotQFXMInterconnect  2 }
jnxQFXMInterconnectSlotPower OBJECT IDENTIFIER ::= { jnxSlotQFXMInterconnect  3 }
jnxQFXMInterconnectSlotFan  OBJECT IDENTIFIER ::= { jnxSlotQFXMInterconnect  4 }
jnxQFXMInterconnectSlotFPB  OBJECT IDENTIFIER ::= { jnxSlotQFXMInterconnect  5 }
```

```
jnxMediaCardSpaceQFXMInterconnect OBJECT IDENTIFIER ::= { jnxMediaCardSpace 91 }
jnxQFXMInterconnectMediaCardSpacePIC OBJECT IDENTIFIER ::= { jnxMediaCardSpaceQFXMInterconnect 1 }
```

QFabric System Node Devices

```
jnxProductLineQFXNode OBJECT IDENTIFIER ::= { jnxProductLine 61 }
jnxProductNameQFXNode OBJECT IDENTIFIER ::= { jnxProductName 61 }
jnxProductModelQFXNode OBJECT IDENTIFIER ::= { jnxProductModel 61 }
jnxProductVariationQFXNode OBJECT IDENTIFIER ::= { jnxProductVariation 61 }
  jnxProductQFX3500 OBJECT IDENTIFIER ::= { jnxProductVariationQFXNode 1 }
  jnxProductQFX360016Q OBJECT IDENTIFIER ::= { jnxProductVariationQFXNode 3 }

jnxChassisQFXNode OBJECT IDENTIFIER ::= { jnxChassis 61 }

jnxSlotQFXNode OBJECT IDENTIFIER ::= { jnxSlot 61 }
  jnxQFXNodeSlotFPC OBJECT IDENTIFIER ::= { jnxSlotQFXNode 1 }
  jnxQFXNodeSlotHMM OBJECT IDENTIFIER ::= { jnxSlotQFXNode 2 }
  jnxQFXNodeSlotPower OBJECT IDENTIFIER ::= { jnxSlotQFXNode 3 }
  jnxQFXNodeSlotFan OBJECT IDENTIFIER ::= { jnxSlotQFXNode 4 }
  jnxQFXNodeSlotFPB OBJECT IDENTIFIER ::= { jnxSlotQFXNode 5 }

jnxMediaCardSpaceQFXNode OBJECT IDENTIFIER ::= { jnxMediaCardSpace 61 }
jnxQFXNodeMediaCardSpacePIC OBJECT IDENTIFIER ::= { jnxMediaCardSpaceQFXNode 1 }
```

- Related Documentation**
- *Understanding the Implementation of SNMP on the QFabric System*
 - *Fabric Chassis MIB*

Configuring SNMP

SNMP is implemented in the Junos OS Software running on the QFX Series and OCX Series products. By default, SNMP is not enabled. To enable SNMP, you must include the SNMP configuration statements at the **[edit]** hierarchy level.

To configure the minimum requirements for SNMP, include the following statements at the **[edit]** hierarchy level of the configuration:

```
[edit]
snmp {
  community public;
}
```

To configure complete SNMP features, include the following statements at the **[edit]** hierarchy level of the configuration:

```
snmp {
  client-list client-list-name {
    ip-addresses;
  }
  community community-name {
    authorization authorization;
    client-list-name client-list-name;
    clients {
      address restrict;
    }
  }
}
```

```
logical-system logical-system-name {  
    routing-instance routing-instance-name {  
        clients {  
            addresses;  
        }  
    }  
}  
routing-instance routing-instance-name {  
    clients {  
        addresses;  
    }  
}  
view view-name;  
}  
contact contact;  
description description;  
filter-duplicates;  
filter-interfaces;  
health-monitor {  
    falling-threshold integer;  
    interval seconds;  
    rising-threshold integer;  
}  
interface [ interface-names ];  
location location;  
name name;  
nonvolatile {  
    commit-delay seconds;  
}  
rmon {  
    alarm index {  
        description description;  
        falling-event-index index;  
        falling-threshold integer;  
        falling-threshold-interval seconds;  
        interval seconds;  
        request-type;  
        rising-event-index index;  
        rising-threshold integer;  
        sample-type (absolute-value | delta-value);  
        startup-alarm (falling-alarm | rising-alarm | rising-or-falling alarm);  
        syslog-subtag syslog-subtag;  
        variable oid-variable;  
    }  
    event index {  
        community community-name;  
        description description;  
        type type;  
    }  
    history history-index {  
        bucket-size number;  
        interface interface-name;  
        interval seconds;  
        owner owner-name;  
    }  
}
```

```

traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable> <match
    regular-expression>;
  flag flag;
}
trap-group group-name {
  categories {
    category;
  }
  destination-port port-number;
  routing-instance routing-instance-name;
  targets {
    address;
  }
  version (all | v1 | v2);
}
trap-options {
  agent-address outgoing-interface;
  source-address address;
}
v3 {
  notify name {
    tag tag-name;
    type trap;
  }
  notify-filter profile-name {
    oid object-identifier (include | exclude);
  }
  snmp-community community-index {
    community-name community-name;
    security-name security-name;
    tag tag-name;
  }
  target-address target-address-name {
    address address;
    address-mask address-mask;
    logical-system logical-system;
    port port-number;
    retry-count number;
    routing-instance routing-instance-name;
    tag-list tag-list;
    target-parameters target-parameters-name;
    timeout seconds;
  }
  target-parameters target-parameters-name {
    notify-filter profile-name;
    parameters {
      message-processing-model (v1 | v2c | V3);
      security-level (authentication | none | privacy);
      security-model (usm | v1 | v2c);
      security-name security-name;
    }
  }
}
usm {
  local-engine {
    user username {

```

```

    authentication-sha {
        authentication-password authentication-password;
    }
    authentication-md5 {
        authentication-password authentication-password;
    }
    authentication-none;
    privacy-aes128 {
        privacy-password privacy-password;
    }
    privacy-des {
        privacy-password privacy-password;
    }
    privacy-3des {
        privacy-password privacy-password;
    }
    privacy-none;
}
}
remote-engine engine-id {
    user username {
        authentication-sha {
            authentication-password authentication-password;
        }
        authentication-md5 {
            authentication-password authentication-password;
        }
        authentication-none;
        privacy-aes128 {
            privacy-password privacy-password;
        }
        privacy-des {
            privacy-password privacy-password;
        }
        privacy-3des {
            privacy-password privacy-password;
        }
        privacy-none {
            privacy-password privacy-password;
        }
    }
}
}
}
vacm {
    access {
        group group-name {
            (default-context-prefix | context-prefix context-prefix) {
                security-model (any | usm | v1 | v2c) {
                    security-level (authentication | none | privacy) {
                        notify-view view-name;
                        read-view view-name;
                        write-view view-name;
                    }
                }
            }
        }
    }
}
}

```

```

    }
    security-to-group {
        security-model (usm | v1 | v2c) {
            security-name security-name {
                group group-name;
            }
        }
    }
}
}
}
view view-name {
    oid object-identifier (include | exclude);
}
}

```

- Related Documentation**
- [Understanding the Implementation of SNMP](#)
 - [snmp on page 1665](#)

Configuring the SNMP Community String

The SNMP community string defines the relationship between an SNMP server system and the client systems. This string acts like a password to control the clients' access to the server. To configure a community string in a Junos OS configuration, include the **community** statement at the **[edit snmp]** hierarchy level:

```

[edit snmp]
community name {
    authorization authorization;
    clients {
        default restrict;
        address restrict;
    }
    view view-name;
}

```

If the community name contains spaces, enclose it in quotation marks (" ").

The default authorization level for a community is **read-only**. To allow **Set** requests within a community, you need to define that community as **authorization read-write**. For **Set** requests, you also need to include the specific MIB objects that are accessible with read-write privileges using the **view** statement. The default view includes all supported MIB objects that are accessible with read-only privileges; no MIB objects are accessible with read-write privileges. For more information about the **view** statement, see [“Configuring MIB Views” on page 1506](#).

The **clients** statement lists the IP addresses of the clients (community members) that are allowed to use this community. If no **clients** statement is present, all clients are allowed. For **address**, you must specify an IPv4 address, not a hostname. Include the **default restrict** option to deny access to all SNMP clients for which access is not explicitly granted. We recommend that you always include the **default restrict** option to limit SNMP client access to the local switch.



NOTE: Community names must be unique within each SNMP system.

Related Documentation

- [Configuring SNMP on page 1499](#)

Configuring SNMP Trap Groups

Before any SNMP traps can be sent, you must configure a trap group, the categories of traps the group can receive, and the targets (systems) that will receive the traps. To create and name an SNMP trap group, include the **trap-group** statement at the **[edit snmp]** hierarchy level:

```
[edit snmp]
trap-group group-name {
  categories {
    category;
  }
  destination-port port-number;
  targets {
    address;
  }
  version (all | v1 | v2);
}
```

The trap group name can be any string and is embedded in the community name field of the trap. To configure your own trap group port, include the **destination-port** statement. The default destination port is port 162.

For each trap group that you define, you must include the **target** statement to define at least one system as the recipient of the SNMP traps in the trap group. Specify the IPv4 address of each recipient and not its hostname.

Specify the types of traps the trap group can receive in the **categories** statement.

A trap group can receive the following categories of traps:

- **authentication**—Authentication failures
- **chassis**—Chassis or environment notifications
- **configuration**—Configuration notifications
- **link**—Link-related notifications such as up-down transitions
- **remote-operations**—Remote operation notifications
- **startup**—System warm and cold starts

The **version** statement allows you to specify the SNMP version of the traps sent to targets of the trap group. If you specify **v1** only, SNMPv1 traps are sent. If you specify **v2** only, SNMPv2 traps are sent. If you specify **all**, both an SNMPv1 and an SNMPv2 trap are sent for every trap condition. For more information about the **version** statement, see [version](#).

Adding a Group of Clients to an SNMP Community

Junos OS enables you to add one or more groups of clients to an SNMP community. You can include the **client-list-name** *name* statement at the **[edit snmp community community-name]** hierarchy level to add all the members of the client list or prefix list to an SNMP community.

To define a list of clients, include the **client-list** statement followed by the IP addresses of the clients at the **[edit snmp]** hierarchy level:

```
[edit snmp]
  client-list client-list-name {
    ip-addresses;
```

You can configure a prefix list at the **[edit policy options]** hierarchy level. Support for prefix lists in the SNMP community configuration enables you to use a single list to configure the SNMP and routing policies. For more information about the **prefix-list** statement, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

To add a client list or prefix list to an SNMP community, include the **client-list-name** statement at the **[edit snmp community community-name]** hierarchy level:

```
[edit snmp community community-name]
  client-list-name client-list-name;
```



NOTE: The client list and prefix list must not have the same name.

The following example shows how to define a client list:

```
[edit]
snmp {
  client-list clentlist1 {
    10.1.1.1/32;
    10.2.2.2/32;
```

The following example shows how to add a client list to an SNMP community:

```
[edit]
snmp {
  community community1 {
    authorization read-only;
    client-list-name clientlist1;
```

The following example shows how to add a prefix list to an SNMP community:

```
[edit]
policy-options {
  prefix-list prefixlist {
```

```
10.3.3.3/32;  
10.5.5.5/32;  
}  
}  
snmp {  
  community community2 {  
    client-list-name prefixlist;  
  }  
}
```

- Related Documentation**
- [client-list on page 1613](#)
 - [client-list-name on page 1614](#)

Configuring the Interfaces on Which SNMP Requests Can Be Accepted

By default, all router or switch interfaces have SNMP access privileges. To limit the access through certain interfaces only, include the **interface** statement at the **[edit snmp]** hierarchy level:

```
[edit snmp]  
interface [ interface-names ];
```

Specify the names of any logical or physical interfaces that should have SNMP access privileges. Any SNMP requests entering the router or switch from interfaces not listed are discarded.

- Related Documentation**
- *Configuring SNMP on a Device Running Junos OS*
 - *Configuration Statements at the [edit snmp] Hierarchy Level*
 - *Example: Configuring Secured Access List Checking*

Configuring MIB Views

SNMPv3 defines the concept of MIB views in RFC 3415, *View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)*. MIB views provide an agent better control over who can access specific branches and objects within its MIB tree. A view consists of a name and a collection of SNMP object identifiers, which are either explicitly included or excluded. Once defined, a view is then assigned to an SNMPv3 group or SNMPv1/v2c community (or multiple communities), automatically masking which parts of the agent's MIB tree members of the group or community can (or cannot) access.

By default, an SNMP community grants read access and denies write access to all supported MIB objects (even communities configured as **authorization read-write**). To restrict or grant read or write access to a set of MIB objects, you must configure a MIB view and associate the view with a community.

To configure MIB views, include the **view** statement at the **[edit snmp]** hierarchy level:

```
[edit snmp]  
view view-name {
```

```
oid object-identifier (include | exclude);
}
```

The **view** statement defines a MIB view and identifies a group of MIB objects. Each MIB object of a view has a common object identifier (OID) prefix. Each object identifier represents a subtree of the MIB object hierarchy. The subtree can be represented either by a sequence of dotted integers (such as **1.3.6.1.2.1.2**) or by its subtree name (such as **interfaces**). A configuration statement uses a view to specify a group of MIB objects on which to define access. You can also use a wildcard character asterisk (*) to include OIDs that match a particular pattern in the SNMP view. To enable a view, you must associate the view with a community.



NOTE: To remove an OID completely, use the **delete view all oid *oid-number*** command but omit the **include** parameter.

To associate MIB views with a community, include the **view** statement at the **[edit snmp community *community-name*]** hierarchy level:

```
[edit snmp community community-name]
view view-name;
```

For more information about the Ping MIB, see RFC 2925 and the *PING MIB* topic.

Related Documentation

- [PING MIB](#)
- [Configuring SNMP on a Device Running Junos OS](#)
- [Configuration Statements at the \[edit snmp\] Hierarchy Level](#)
- [Example: Ping Proxy MIB](#)
- [view \(Configuring a MIB View\)](#)
- [view \(Associating MIB View with a Community\) on page 1690](#)
- [oid](#)

Configuring RMON Alarms and Events

The Junos OS supports the *Remote Network Monitoring* (RMON) MIB (RFC 2819), which allows a management device to monitor the values of MIB objects, or variables, against configured thresholds. When the value of a variable crosses a threshold, an alarm and its corresponding event are generated. The event can be logged and can generate an SNMP trap.

To configure RMON alarms and events using the CLI, perform these tasks:

1. [Configuring SNMP on page 1508](#)
2. [Configuring an Event on page 1508](#)
3. [Configuring an Alarm on page 1509](#)

Configuring SNMP

To configure SNMP:

1. Grant read-only access to all SNMP clients:

```
[edit snmp]
user@switch# set community community-name authorization authorization
```

For example:

```
[edit snmp]
user@switch# set community public authorization read-only
```

2. Grant read-write access to the RMON and jnx-rmon MIBs:

```
[edit snmp]
user@switch# set view view-name oid object-identifier include
user@switch# set view view-name oid object-identifier include
user@switch# set community community-name authorization authorization view view-name
```

For example:

```
[edit snmp]
user@switch# set view rmon-mib-view oid .1.3.6.1.2.1.16 include
user@switch# set view rmon-mib-view oid .1.3.6.1.4.1.2636.13 include
user@switch# set community private authorization read-write view rmon-mib-view
```

OIDs 1.3.6.1.2.1.16 and 1.3.6.1.4.1.2636.13 correspond to the RMON and jnxRmon MIBs.

3. Configure an SNMP trap group:

```
[edit snmp]
user@switch# set trap-group group-name categories category
user@switch# set trap-group group-name targets address
```

For example:

```
[edit snmp]
user@switch# set trap-group rmon-trap-group categories rmon-alarm
user@switch# set trap-group rmon-trap-group targets 192.168.5.5
```

The trap group **rmon-trap-group** is configured to send RMON traps to 192.168.5.5.

Configuring an Event

To configure an event:

1. Configure an event index, community name, and type:

```
[edit snmp rmon]
user@switch# set event index community community-name type type
```

For example:

```
[edit snmp rmon]
user@switch# set event 1 community rmon-trap-group type log-and-trap
```

The event community corresponds to the SNMP trap group and is not the same as an SNMP community. This event generates an SNMP trap and adds an entry to the **logTable** in the RMON MIB.

2. Configure a description for the event:

```
[edit snmp rmon]
user@switch# set event index description description
```

For example:

```
[edit snmp rmon]
user@switch# set event 1 description "rmon event"
```

Configuring an Alarm

To configure an alarm:

1. Configure an alarm index, the variable to monitor, the rising and falling thresholds, and the corresponding rising and falling events:

```
[edit snmp rmon]
user@switch# set alarm index variable oid-variable falling-threshold integer rising-threshold
integer rising-event-index index falling-event-index index
```

For example:

```
[edit snmp rmon]
user@switch# set alarm 5 variable .1.3.6.1.4.1.2636.3.1.13.1.8.9.1.0.0 falling-threshold 75
rising-threshold 90 rising-event-index 1 falling-event-index 1
```

The variable .1.3.6.1.4.1.2636.3.1.13.1.8.9.1.0.0 corresponds to the `jnxRmon` MIB object `jnxOperatingCPU`, which represents the CPU utilization of the Routing Engine. The falling and rising threshold integers are 75 and 90. The rising and falling events both generate the same event (event index 1).

2. Configure the sample interval and type and the alarm type:

```
[edit snmp rmon]
user@switch# set alarm index interval seconds sample-type (absolute-value | delta-value)
startup-alarm (falling-alarm | rising-alarm | rising-or-falling-alarm)
```

For example:

```
[edit snmp rmon]
user@switch# set alarm 5 interval 30 sample-type absolute-value
startup-alarm rising-or-falling-alarm
```

The absolute value of the monitored variable is sampled every 30 seconds. The initial alarm can occur because of rising above the rising threshold or falling below the falling threshold.

Related Documentation

- [Configuring SNMP on page 1499](#)
- [Juniper Networks Enterprise-Specific MIBs](#)
- [Monitoring RMON MIB Tables on page 1516](#)
- [RMON MIB Event, Alarm, Log, and History Control Tables on page 1464](#)
- [Understanding RMON on page 1463](#)

Configuring Health Monitoring

This topic describes how to configure the health monitor feature for QFX Series and OCX Series devices.

The health monitor feature extends the SNMP RMON alarm infrastructure to provide predefined monitoring for a selected set of object instances (such as file system usage,

CPU usage, and memory usage) and dynamic object instances (such as Junos OS processes).

To configure health monitoring:

1. Configure the health monitor:

```
[edit snmp]
user@switch# set health-monitor
```

2. Configure the falling threshold:

```
[edit snmp]
user@switch# set health-monitor falling-threshold percentage
```

For example:

```
user@switch# set health-monitor falling-threshold 85
```

3. Configure the rising threshold:

```
[edit snmp]
user@switch# set health-monitor rising-threshold percentage
```

For example:

```
user@switch# set health-monitor rising-threshold 75
```

4. Configure the interval:

```
[edit snmp]
user@switch# set health-monitor interval seconds
```

For example:

```
user@switch# set health-monitor interval 600
```

Related Documentation

- [Understanding Health Monitoring on page 1467](#)
- [falling-threshold on page 1624](#)
- [interval \(Health Monitor\) on page 1633](#)
- [rising-threshold \(Health Monitor\) on page 1653](#)

Creating SNMPv3 Users

For each SNMPv3 user, you can specify the username, authentication type, authentication password, privacy type, and privacy password. After a user enters a password, a key based on the engine ID and password is generated and is written to the configuration file. After the generation of the key, the password is deleted from this configuration file.



NOTE: You can configure only one encryption type for each SNMPv3 user.

To create users, include the `user` statement at the `[edit snmp v3 usm local-engine]` hierarchy level:

```
[edit snmp v3 usm local-engine]
user username;
```

`username` is the name that identifies the SNMPv3 user.

To configure user authentication and encryption, include the following statements at the `[edit snmp v3 usm local-engine user username]` hierarchy level:

```
[edit snmp v3 usm local-engine user username]
authentication-md5 {
  authentication-password authentication-password;
}
authentication-sha {
  authentication-password authentication-password;
}
authentication-none;
privacy-aes128 {
  privacy-password privacy-password;
}
privacy-des {
  privacy-password privacy-password;
}
privacy-3des {
  privacy-password privacy-password;
}
privacy-none;
```

Related Documentation

- [Complete SNMPv3 Configuration Statements](#)
- [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 1461](#)
- [Example: Creating SNMPv3 Users](#)
- [Example: SNMPv3 Configuration](#)

Configuring Access Privileges for a Group

In SNMPv3, you can configure a group that sets the same access privileges for one or more users. Configuring a group includes defining the security model and security level, and associating one or more MIB view permissions for the group.



NOTE: You must associate at least one MIB view with the group. You can associate multiple MIB views (read, notify, write) to authorize different permissions based on the view. The view name cannot exceed 32 characters.

To configure access privileges for a group:

1. To configure the group:

```
[edit snmp v3 vacm access]
user@switch# edit group group-name
```

2. To configure the context prefix of the SNMP instance for the group:

```
[edit snmp v3 vacm access group group-name]
user@switch# edit (default-context-prefix | context-prefix context-prefix)
```

For example, to configure the default context prefix:

```
[edit snmp v3 vacm access group group-name]
user@switch# edit default-context-prefix
```

3. To configure the security model:

```
[edit snmp v3 vacm access group group-name (default-context-prefix | context-prefix
context-prefix)]
user@switch# edit security-model (any | usm | v1 | v2c)
```

For example, to configure the SNMPv3 user-based security model (USM):

```
[edit snmp v3 vacm access group group-name (default-context-prefix | context-prefix
context-prefix)]
user@switch# edit security-model usm
```

4. To configure the security level:

```
[edit snmp v3 vacm access group group-name (default-context-prefix | context-prefix
context-prefix) security-model (any | usm | v1 | v2c)]
user@switch# edit security-level (authentication | none | privacy)
```

For example, to configure a security level requiring user authentication and encryption:

```
[edit snmp v3 vacm access group group-name (default-context-prefix | context-prefix
context-prefix) security-model (any | usm | v1 | v2c)]
user@switch# edit security-level privacy
```



NOTE: Access privileges are granted to all packets with a security level equal to or greater than that configured. If you are configuring the SNMPv1 or v2c security model, use *none* as your security level. If you are configuring the SNMPv3 security model (USM), use the *authentication*, *none*, or *privacy* security level.

5. (Optional) To associate a read-only MIB view with an SNMP group:

```
[edit snmp v3 vacm access group group-name (default-context-prefix | context-prefix
context-prefix) security-model (any | usm | v1 | v2c) security-level (authentication |
none | privacy)]
user@switch# edit read-view view-name
```

6. (Optional) To associate a MIB view with an SNMP notification permission for an SNMP group:

```
[edit snmp v3 vacm access group group-name (default-context-prefix | context-prefix
context-prefix) security-model (any | usm | v1 | v2c) security-level (authentication |
none | privacy)]
user@switch# edit notify-view view-name
```

7. (Optional) To associate a MIB view with write permission for an SNMP group:


```
[edit snmp v3 vacm access group group-name (default-context-prefix | context-prefix
context-prefix) security-model (any | usm | v1 | v2c) security-level (authentication |
none | privacy)]
user@switch# edit write-view view-name
```

- Related Documentation**
- [SNMPv3 Overview on page 1460](#)
 - [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 1461](#)

Assigning a Security Name to a Group

In SNMPv3, each username is associated with a security name. The security name, together with the SNMP engine ID, is included in SNMP messages to ensure messaging security.

Before you assign a security name to a group, first create the security name. For an SNMPv3 client, the security name is the username configured at the **[edit snmp v3 usm local-engine user *username*]** hierarchy level. For SNMPv1 or v2c clients, the security name is the community string configured at the **[edit snmp v3 snmp-community *community-index*]** hierarchy level.

Assigning a security name to a group includes configuring a security model for the group, assigning the security name to the group, and configuring the group.

To assign an SNMP security name to a group:

1. To configure a security model for the group:

```
[edit snmp v3 vacm security-to-group]
user@switch# edit security-model (usm | v1 | v2c)
```

For example, to configure the SNMPv3 user-based security model (USM):

```
[edit snmp v3 vacm security-to-group]
user@switch# edit security-model usm
```

2. To associate the security name with a group:

```
[edit snmp v3 vacm security-to-group security-model (usm | v1 | v2c)]
user@switch# edit security-name security-name
```

3. To configure a group of SNMPv3 security names with the same security policy:

```
[edit snmp v3 vacm security-to-group security-model (usm | v1 | v2c) security-name
security-name]
user@switch# edit group group-name
```

- Related Documentation**
- [Creating SNMPv3 Users on page 1510](#)
 - [group \(Associating a Security Name\) on page 1627](#)
 - [security-model \(Group\) on page 1660](#)
 - [security-name \(Community String\) on page 1661](#)
 - [security-name \(Security Group\) on page 1662](#)

Configuring SNMPv3 Traps on a Device Running Junos OS

In SNMPv3, you create traps and informs by configuring the **notify**, **target-address**, and **target-parameters** parameters. Traps are unconfirmed notifications, whereas informs are confirmed notifications. This section describes how to configure SNMP traps. For information about configuring SNMP informs, see [“Configuring SNMP Informs” on page 1515](#).

The target address defines a management application’s address and parameters to be used in sending notifications. Target parameters define the message processing and security parameters that are used in sending notifications to a particular management target. SNMPv3 also lets you define SNMPv1 and SNMPv2c traps.



NOTE: When you configure SNMP traps, make sure your configured access privileges allow the traps to be sent. Access privileges are configured at the `[edit snmp v3 vacm access]` and `[edit snmp v3 vacm security-to-group]` hierarchy levels.

To configure SNMP traps, include the following statements at the `[edit snmp v3]` hierarchy level:

```
[edit snmp v3]
  notify name {
    tag tag-name;
    type trap;
  }
  notify-filter name {
    oid object-identifier (include | exclude);
  }
  target-address target-address-name {
    address address;
    address-mask address-mask;
    logical-system (SNMP) logical-system;
    port port-number;
    routing-instance instance;
    tag-list tag-list;
    target-parameters target-parameters-name;
  }
  target-parameters target-parameters-name {
    notify-filter profile-name;
    parameters {
      message-processing-model (v1 | v2c | v3);
      security-level (authentication | none | privacy);
      security-model (usm | v1 | v2c);
      security-name security-name;
    }
  }
}
```

- Related Documentation**
- [Configuring the SNMPv3 Trap Notification](#)
 - [Configuring the Trap Notification Filter](#)
 - [Configuring the Trap Target Address](#)

- *Defining and Configuring the Trap Target Parameters*
- [Configuring SNMP Informs on page 1515](#)
- *Configuring the Remote Engine and Remote User*
- *Configuring the Inform Notification Type and Target Address*

Configuring SNMP Informs


Junos OS supports two types of notifications: traps and informs. With traps, the receiver does not send any acknowledgment when it receives a trap. Therefore, the sender cannot determine if the trap was received. A trap may be lost because a problem occurred during transmission. To increase reliability, an inform is similar to a trap except that the inform is stored and retransmitted at regular intervals until one of these conditions occurs:

- The receiver (target) of the inform returns an acknowledgment to the SNMP agent.
- A specified number of unsuccessful retransmissions have been attempted and the agent discards the inform message.

If the sender never receives a response, the inform can be sent again. Thus, informs are more likely to reach their intended destination than traps are. Informs use the same communications channel as traps (same socket and port) but have different protocol data unit (PDU) types.

Informs are more reliable than traps, but they consume more network, router, and switch resources (see [Figure 19](#)). Unlike a trap, an inform is held in memory until a response is received or the timeout is reached. Also, traps are sent only once, whereas an inform may be retried several times. Use informs when it is important that the SNMP manager receive all notifications. However, if you are more concerned about network traffic, or router and switch memory, use traps.

Figure 19: Inform Request and Response



For information about configuring SNMP traps, see [“Configuring SNMPv3 Traps on a Device Running Junos OS” on page 1514](#).

Related Documentation

- [Configuring SNMPv3 Traps on a Device Running Junos OS on page 1514](#)
- *Configuring the Remote Engine and Remote User*
- *Configuring the Inform Notification Type and Target Address*
- *Complete SNMPv3 Configuration Statements*

- [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 1461](#)

Monitoring RMON MIB Tables

Purpose Monitor remote monitoring (RMON) alarm, event, and log tables.

Action To display the RMON tables:

```
user@switch> show snmp rmon
Alarm
Index  Variable description                               Value State

      5 monitor
      jnxOperatingCPU.9.1.0.0                        5 falling threshold

Event
Index  Type                               Last Event
      1 log and trap                     2010-07-10 11:34:17 PDT
Event Index: 1
  Description: Event 1 triggered by Alarm 5, rising threshold (90) crossed,
(variable: jnxOperatingCPU.9.1.0.0, value: 100)
  Time: 2010-07-10 11:34:07 PDT
  Description: Event 1 triggered by Alarm 5, falling threshold (75) crossed,
(variable: jnxOperatingCPU.9.1.0.0, value: 5)
  Time: 2010-07-10 11:34:17 PDT
```

Meaning The display shows that an alarm has been defined to monitor jnxRmon MIB object jnxOperatingCPU, which represents the CPU utilization of the Routing Engine. The alarm is configured to generate an event that sends an SNMP trap and adds an entry to the logTable in the RMON MIB. The log table shows that two occurrences of the event have been generated—one for rising above a threshold of 90 percent, and one for falling below a threshold of 75 percent.

Related Documentation

- [Configuring RMON Alarms and Events on page 1507](#)
- [show snmp rmon on page 1759](#)
- [show snmp rmon history on page 1763](#)
- [clear snmp statistics on page 1739](#)
- [clear snmp history on page 1738](#)

Monitoring SNMP

There are several commands that you can access in Junos OS operational mode to monitor SNMP information. Some of the commands are:

- **show snmp health-monitor**, which displays the health monitor log and alarm information.
- **show snmp mib**, which displays information from the MIBs, such as device and system information.

- **show snmp statistics**, which displays SNMP statistics such as the number of packets, silent drops, and invalid output values.
- **show snmp rmon**, which displays the RMON alarm, event, history, and log information

The following example provides sample output from the **show snmp health-monitor** command:

```
user@switch> show snmp health-monitor
Alarm
Index  Variable description                               Value State

32768 Health Monitor: root file system utilization
      jnxHrStoragePercentUsed.1                      58 active

32769 Health Monitor: /config file system utilization
      jnxHrStoragePercentUsed.2                      0 active

32770 Health Monitor: RE 0 CPU utilization
      jnxOperatingCPU.9.1.0.0                        0 active

32773 Health Monitor: RE 0 Memory utilization
      jnxOperatingBuffer.9.1.0.0                    35 active

32775 Health Monitor: jkernel daemon CPU utilization
      Init daemon                                   0 active
      Chassis daemon                               50 active
      Firewall daemon                              0 active
      Interface daemon                             5 active
      SNMP daemon                                  11 active
      MIB2 daemon                                  42 active
      ...
```

The following example provides sample output from the **show snmp mib** command:

```
user@switch> show snmp mib walk system

sysDescr.0    = Juniper Networks, Inc. qfx3500s internet router, kernel
JUNOS 11.1-20100926.0 #0: 2010-09-26 06:17:38 UTC builder@abc.juniper.net:
/volume/build/junos/11.1/production/20100926.0/obj-xlr/bsd/sys/compile/JUNIPER-xxxxx

Build date: 2010-09-26 06:00:10 U
sysObjectID.0 = jnxProductQFX3500
sysUpTime.0   = 24444184
sysContact.0  = J Smith
sysName.0     = Lab QFX3500
sysLocation.0 = Lab
sysServices.0 = 4
```

The following example provides sample output from the **show snmp statistics** command:

```
user@switch> show snmp statistics

SNMP statistics:
Input:
  Packets: 0, Bad versions: 0, Bad community names: 0,
  Bad community uses: 0, ASN parse errors: 0,
  Too big: 0, No such names: 0, Bad values: 0,
  Read only: 0, General errors: 0,
```

```
Total request varbinds: 0, Total set varbinds: 0,  
Get requests: 0, Get nexts: 0, Set requests: 0,  
Get responses: 0, Traps: 0,  
Silent drops: 0, Proxy drops: 0, Commit pending drops: 0,  
Throttle drops: 0, Duplicate request drops: 0  
Output:  
Packets: 0, Too bigs: 0, No such names: 0,  
Bad values: 0, General errors: 0,  
Get requests: 0, Get nexts: 0, Set requests: 0,  
Get responses: 0, Traps: 0
```

- Related Documentation
- [health-monitor on page 1629](#)
 - [show snmp mib on page 1756](#)
 - [show snmp statistics on page 1764](#)

Tracing SNMP Activity on a Device Running Junos OS

SNMP tracing operations track activity for SNMP agents and record the information in log files. The logged error descriptions provide detailed information to help you solve problems faster.

By default, Junos OS does not trace any SNMP activity. If you include the **traceoptions** statement at the **[edit snmp]** hierarchy level, the default tracing behavior is:

- Important activities are logged in files located in the **/var/log** directory. Each log is named after the SNMP agent that generates it. Currently, the following log files are created in the **/var/log** directory when the **traceoptions** statement is used:
 - chassisd
 - craftd
 - ilmid
 - mib2d
 - rmopd
 - serviced
 - snmpd
- When a trace file named **filename** reaches its maximum size, it is renamed **filename.0**, then **filename.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. (For more information about how log files are created, see the *Junos OS System Log Messages Reference*.)
- Log files can be accessed only by the user who configured the tracing operation.

You cannot change the directory (**/var/log**) in which trace files are located. However, you can customize the other trace file settings by including the following statements at the **[edit snmp]** hierarchy level:

```
[edit snmp]  
traceoptions {
```

```

file <files number> <match regular-expression> <size size> <world-readable |
no-world-readable>;
flag flag;
memory-trace;
no-remote-trace;
no-default-memory-trace;
}

```

These statements are described in the following sections:

- [Configuring the Number and Size of SNMP Log Files on page 1519](#)
- [Configuring Access to the Log File on page 1519](#)
- [Configuring a Regular Expression for Lines to Be Logged on page 1520](#)
- [Configuring the Trace Operations on page 1520](#)

Configuring the Number and Size of SNMP Log Files

By default, when the trace file reaches 128 kilobytes (KB) in size, it is renamed *filename.0*, then *filename.1*, and so on, until there are three trace files. Then the oldest trace file (*filename.2*) is overwritten.

You can configure the limits on the number and size of trace files by including the following statements at the **[edit snmp traceoptions]** hierarchy level:

```

[edit snmp traceoptions]
file files number size size;

```

For example, set the maximum file size to 2 MB, and the maximum number of files to 20. When the file that receives the output of the tracing operation (*filename*) reaches 2 MB, *filename* is renamed *filename.0*, and a new file called *filename* is created. When the new *filename* reaches 2 MB, *filename.0* is renamed *filename.1* and *filename* is renamed *filename.0*. This process repeats until there are 20 trace files. Then the oldest file (*filename.19*) is overwritten by the newest file (*filename.0*).

The number of files can be from 2 through 1000 files. The file size of each file can be from 10 KB through 1 gigabyte (GB).

Configuring Access to the Log File

By default, log files can be accessed only by the user who configured the tracing operation.

To specify that any user can read all log files, include the **file world-readable** statement at the **[edit snmp traceoptions]** hierarchy level:

```

[edit snmp traceoptions]
file world-readable;

```

To explicitly set the default behavior, include the **file no-world-readable** statement at the **[edit snmp traceoptions]** hierarchy level:

```

[edit snmp traceoptions]
file no-world-readable;

```

Configuring a Regular Expression for Lines to Be Logged

By default, the trace operation output includes all lines relevant to the logged activities.

You can refine the output by including the **match** statement at the **[edit snmp traceoptions file *filename*]** hierarchy level and specifying a regular expression (regex) to be matched:

```
[edit snmp traceoptions]
file filename match regular-expression;
```

Configuring the Trace Operations

By default, only important activities are logged. You can specify which trace operations are to be logged by including the following **flag** statement (with one or more tracing flags) at the **[edit snmp traceoptions]** hierarchy level:

```
[edit snmp traceoptions]
flag {
  all;
  configuration;
  database;
  events;
  general;
  interface-stats;
  nonvolatile-sets;
  pdu;
  policy;
  protocol-timeouts;
  routing-socket;
  server;
  subagent;
  timer;
  varbind-error;
}
```

Table 108 describes the meaning of the SNMP tracing flags.

Table 108: SNMP Tracing Flags

Flag	Description	Default Setting
all	Log all operations.	Off
configuration	Log reading of the configuration at the [edit snmp] hierarchy level.	Off
database	Log events involving storage and retrieval in the events database.	Off
events	Log important events.	Off
general	Log general events.	Off
interface-stats	Log physical and logical interface statistics.	Off

Table 108: SNMP Tracing Flags (*continued*)

Flag	Description	Default Setting
nonvolatile-set	Log nonvolatile SNMP set request handling.	Off
pdu	Log SNMP request and response packets.	Off
policy	Log policy processing.	Off
protocol-timeouts	Log SNMP response timeouts.	Off
routing-socket	Log routing socket calls.	Off
server	Log communication with processes that are generating events.	Off
subagent	Log subagent restarts.	Off
timer	Log internal timer events.	Off
varbind-error	Log variable binding errors.	Off

To display the end of the log for an agent, issue the **show log agentd | last** operational mode command:

```
[edit]
user@host# run show log agentd | last
```

where **agent** is the name of an SNMP agent.

Related Documentation

- [Configuring SNMP on a Device Running Junos OS](#)
- [Configuration Statements at the \[edit snmp\] Hierarchy Level](#)
- [Example: Tracing SNMP Activity](#)
- [Configuring SNMP on page 1499](#)

Using the Enterprise-Specific Utility MIB to Enhance SNMP Coverage

Even though the Junos OS has built-in performance metrics and monitoring options, you might need to have customized performance metrics. To make it easier for you to monitor such customized data through a standard monitoring system, the Junos OS provides you with an enterprise-specific Utility MIB that can store such data and thus extend SNMP support for managing and monitoring the data of your choice.

The enterprise-specific Utility MIB provides you with container objects of the following types: **32-bit counters**, **64-bit counters**, **signed integers**, **unsigned integers**, and **octet strings**. You can use these container MIB objects to store the data that are otherwise not supported for SNMP operations. You can populate data for these objects either by using

CLI commands or with the help of Op scripts and an RPC API that can invoke the CLI commands.

The following CLI commands enable you to set and clear Utility MIB object values:

- **request snmp utility-mib set** instance *name* object-type <counter | counter 64 | integer | string | unsigned integer> object-value *value*
- **request snmp utility-mib clear** instance *name* object-type <counter | counter 64 | integer | string | unsigned integer>

The **instance *name*** option of the **request snmp utility-mib <set | clear>** command specifies the name of the data instance and is the main identifier of the data. The **object-type <counter | counter 64 | integer | string | unsigned integer>** option enables you specify the object type, and the **object-value *value*** option enables you to set the value of the object.

To automate the process of populating Utility MIB data, you can use a combination of an event policy and event script. The following examples show the configuration for an event policy to run **show system buffers** every hour and to store the **show system buffers** data in Utility MIB objects by running an event script (**check-mbufs.slax**).

Event Policy Configuration

To configure an event policy that runs the **show system buffers** command every hour and invokes **check-mbufs.slax** to store the **show system buffers** data into Utility MIB objects, include the following statements at the **[edit]** hierarchy level:

```
event-options {
  generate-event {
    1-HOUR time-interval 3600;
  }
  policy Mbufs {
    events 1-HOUR;
    then {
      event-script check-mbufs.slax; # script stored at /var/db/scripts/event/
    }
  }
  event-script {
    file check-mbufs.slax;
  }
}
```

check-mbufs.slax Script

The following example shows the **check-mbufs.slax** script that is stored under **/var/db/scripts/event/**:

```
----- script START -----
version 1.0;

ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
ns ext = "http://xmlsoft.org/XSLT/namespace";

match / {
  <op-script-results>{
    var $cmd = <command> "show system buffers";
    var $out = jcs:invoke($cmd);
```

```

var $lines = jcs:break_lines($out);
for-each ($lines) {
    if (contains(., "current/peak/max")) {
        var $pattern = "([0-9]+)/([0-9]+)/([0-9]+) mbufs";
        var $split = jcs:regex($pattern, .);
        var $result = $split[2];

        var $rpc = <request-snmp-utility-mib-set> {
            <object-type> "integer";
            <instance> "current-mbufs";
            <object-value> $result;
        }
        var $res = jcs:invoke($rpc);
    }
}
}
----- script END -----

```

You can run the following command to check the data stored in the Utility MIB as a result of the event policy and script shown in the preceding examples:

```

user@host> show snmp mib walk jnxUtilData ascii jnxUtilIntegerValue."current-mbufs"
= 0 jnxUtilIntegerTime."current-mbufs" = 07 da 05 0c 03 14 2c 00 2d 07 00
regress@caramels>

```



NOTE: The `show snmp mib walk` command is not available on the QFabric system, but you can use external SNMP client applications to perform this operation.

Related Documentation

- [Understanding SNMP Implementation in Junos OS](#)
- [Configuring SNMP on Devices Running Junos OS](#)
- [Monitoring SNMP Activity and Tracking Problems That Affect SNMP Performance on a Device Running Junos OS](#)
- [Optimizing the Network Management System Configuration for the Best Results](#)
- [Configuring Options on Managed Devices for Better SNMP Response Time](#)
- [Managing Traps and Informs](#)
- [Understanding the Implementation of SNMP on the QFabric System](#)

Example: Configuring SNMP

By default, SNMP is disabled on devices running Junos OS. This example describes the steps for configuring SNMP on the QFabric system.

- [Requirements on page 1524](#)
- [Overview on page 1524](#)
- [Configuration on page 1524](#)

Requirements

This example uses the following hardware and software components:

- Junos OS Release 12.2
- Network management system (NMS) (running the SNMP manager)
- QFabric system (running the SNMP agent) with multiple Node devices

Overview

Because SNMP is disabled by default on devices running Junos OS, you must enable SNMP on your device by including configuration statements at the **[edit snmp]** hierarchy level. At a minimum, you must configure the **community public** statement. The community defined as public grants read-only access to MIB data to any client.

If no **clients** statement is configured, all clients are allowed. We recommend that you always include the **restrict** option to limit SNMP client access to the switch.

The network topology in this example includes an NMS, a QFabric system with four Node devices, and external SNMP servers that are configured for receiving traps.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set snmp name "snmp qfabric" description "qfabric0 switch"
set snmp location "Lab 4 Row 11" contact "qfabric-admin@qfabric0"
set snmp community public authorization read-only
set snmp client-list list0 192.168.0.0/24
set snmp community public client-list-name list0
set snmp community public clients 192.170.0.0/24 restrict
set snmp trap-group "qf-traps" destination-port 155 targets 192.168.0.100
```

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure SNMP on the QFabric system:



NOTE: If the name, description, location, contact, or community name contains spaces, enclose the text in quotation marks (" ").

1. Configure the SNMP system name:

```
[edit snmp]
user@switch# set name "snmp qfabric"
```

2. Specify a description.

```
[edit snmp]
user@switch# set description "qfabric0 system"
```

This string is placed into the MIB II sysDescription object.

3. Specify the physical location of the QFabric system.

```
[edit snmp]
user@switch# set location "Lab 4 Row 11"
```

This string is placed into the MIB II sysLocation object.

4. Specify an administrative contact for the SNMP system.

```
[edit snmp]
user@switch# set contact "qfabric-admin@qfabric0"
```

This name is placed into the MIB II sysContact object.

5. Specify a unique SNMP community name and the read-only authorization level.



NOTE: The read-write option is not supported on the QFabric system.

```
[edit snmp]
user@switch# set community public authorization read-only
```

6. Create a client list with a set of IP addresses that can use the SNMP community.

```
[edit snmp]
user@switch# set client-list list0 192.168.0.0/24
user@switch# set community public client-list-name list0
```

7. Specify IP addresses of clients that are restricted from using the community.

```
[edit snmp]
user@switch# set community public clients 192.170.0.0/24 restrict
```

8. Configure a trap group, destination port, and a target to receive the SNMP traps in the trap group.

```
[edit snmp]
user@switch# set trap-group "qf-traps" destination-port 155 targets 192.168.0.100
```



NOTE: You do not need to include the destination-port statement if you use the default port 162.

The trap group qf-traps is configured to send traps to 192.168.0.100.

Results From configuration mode, confirm your configuration by entering the **show** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
```

```
user@switch# show
snmp {
  name "snmp qfabric";
  description "qfabric0 system";
  location "Lab 4 Row 11";
  contact "qfabric-admin@qfabric0";
  client-list list0 {
    192.168.0.0/24;
  }
  community public {
    authorization read-only;
    clients {
      197.170.0.0/24 restrict;
    }
  }
  trap-group qf-traps {
    destination-port 155;
    targets {
      192.168.0.100;
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

- Related Documentation**
- *Understanding the Implementation of SNMP on the QFabric System*
 - [snmp on page 1665](#)

Configuring System Logging

- Overview of Junos OS System Log Messages on page 1528
- Overview of Single-Chassis System Logging Configuration on page 1528
- Junos OS Minimum System Logging Configuration on page 1530
- Junos OS System Log Configuration Statements on page 1530
- Adding a Text String to System Log Messages on page 1531
- Directing System Log Messages to a Log File on page 1532
- Directing System Log Messages to a Remote Machine on page 1533
- Directing System Log Messages to a User Terminal on page 1533
- Directing System Log Messages to the Console on page 1534
- Disabling the System Logging of a Facility on page 1534
- Displaying a Log File from a Single-Chassis System on page 1535
- Including Priority Information in System Log Messages on page 1536
- Including the Year or Millisecond in Timestamps on page 1537
- Logging Messages in Structured-Data Format on page 1538
- Interpreting Messages Generated in Structured-Data Format on page 1539
- Interpreting Messages Generated in Standard Format on page 1542
- Specifying Log File Size, Number, and Archiving Properties on page 1543
- Specifying the Facility and Severity of Messages to Include in the Log on page 1544
- Junos OS System Logging Facilities and Message Severity Levels on page 1546
- Default Facilities for System Log Messages Directed to a Remote Destination on page 1548
- Junos OS System Log Alternate Facilities for Remote Logging on page 1548
- Changing the Alternative Facility Name for System Log Messages Directed to a Remote Destination on page 1549
- Using Regular Expressions to Refine the Set of Logged Messages on page 1551
- Managing Host OS System Log and Core Files on page 1553
- Monitoring System Log Messages on page 1555
- Examples: Configuring System Logging on page 1556

- [Examples: Assigning an Alternative Facility on page 1558](#)
- [Example: Configuring System Log Messages on page 1559](#)

Overview of Junos OS System Log Messages

The Junos OS generates system log messages (also called *syslog messages*) to record events that occur on the switch, including the following:

- Routine operations, such as a user login into the configuration database.
- Failure and error conditions, such as failure to access a configuration file.
- Emergency or critical conditions, such as power-down of the switch due to excessive temperature.

Each system log message identifies the Junos OS process that generated the message and briefly describes the operation or error that occurred. For detailed information about specific system log messages, see the *Junos OS System Log Messages Reference*.



NOTE: OCX Series switches comprise both the Junos OS and the host operating system (OS). For information about system logging on the host OS, see [“Managing Host OS System Log and Core Files” on page 1553](#).

Related Documentation

- [Junos OS System Log Configuration Statements on page 1530](#)
- [Junos OS Minimum System Logging Configuration on page 1530](#)

Overview of Single-Chassis System Logging Configuration

The Junos OS system logging utility on the QFX Series is similar to the UNIX **syslogd** utility. This topic describes how to configure system logging for a single-chassis system that runs the Junos OS.

Each system log message belongs to a *facility*, which groups together related messages. Each message is also preassigned a *severity level*, which indicates how seriously the triggering event affects router functions. You always specify the facility and severity of the messages to include in the log. For more information, see [“Specifying the Facility and Severity of Messages to Include in the Log” on page 1544](#).

You direct messages to one or more destinations by including the appropriate statement at the **[edit system syslog]** hierarchy level:

- To a named file in a local file system, by including the **file** statement. See [“Directing System Log Messages to a Log File” on page 1532](#).
- To the terminal session of one or more specific users (or all users) when they are logged in to the switch, by including the **user** statement. See [“Directing System Log Messages to a User Terminal” on page 1533](#).

- To the switch console, by including the **console** statement. See [“Directing System Log Messages to the Console” on page 1534](#).
- To a remote machine that is running the **syslogd** utility, by including the **host** statement. See [“Directing System Log Messages to a Remote Machine” on page 1533](#).

By default, messages are logged in a standard format, which is based on a UNIX system log format; for detailed information about message formatting, see the *Junos OS System Log Messages Reference*. You can alter the content and format of logged messages in the following ways:

- You can log messages to a file in structured-data format instead of the standard Junos OS format. Structured-data format provides more information without adding significant length, and makes it easier for automated applications to extract information from the message. For more information, see [“Logging Messages in Structured-Data Format” on page 1538](#).
- A message’s facility and severity level are together referred to as its *priority*. By default, the standard Junos OS format for messages does not include priority information (structured-data format includes a priority code by default). To include priority information in standard-format messages directed to a file or a remote destination, include the **explicit-priority** statement. For more information, see [“Including Priority Information in System Log Messages” on page 1536](#).
- By default, the standard Junos OS format for messages specifies the month, date, hour, minute, and second when the message was logged. You can modify the timestamp on standard-format system log messages to include the year, the millisecond, or both. (Structured-data format specifies the year and millisecond by default.) For more information, see [“Including the Year or Millisecond in Timestamps” on page 1537](#).
- When directing messages to a remote machine, you can specify the IP address that is reported in messages as their source. You can also configure features that make it easier to separate messages generated by Junos OS or messages generated on particular switches. For more information, see [“Directing System Log Messages to a Remote Machine” on page 1533](#).
- The predefined facilities group together related messages, but you can also use regular expressions to specify more exactly which messages from a facility are logged to a file, a user terminal, or a remote destination. For more information, see [“Using Regular Expressions to Refine the Set of Logged Messages” on page 1551](#).



NOTE: During a commit check, warnings about the **traceoptions** configuration (for example, mismatch in trace file sizes or number of trace files) are not displayed on the console. However, these warnings are logged in the system log messages when the new configuration is committed.

Related Documentation

- [Examples: Configuring System Logging on page 1556](#)
- [Specifying the Facility and Severity of Messages to Include in the Log on page 1544](#)
- [Junos OS System Logging Facilities and Message Severity Levels on page 1546](#)

- [Directing System Log Messages to a Log File on page 1532](#)
- [Directing System Log Messages to a Remote Machine on page 1533](#)
- [Directing System Log Messages to a User Terminal on page 1533](#)
- [Directing System Log Messages to the Console on page 1534](#)

Junos OS Minimum System Logging Configuration

To record or view system log messages, you must include the **syslog** statement at the **[edit system]** hierarchy level. Specify at least one destination for the messages, as described in [Table 109](#). For more information about the configuration statements, see *Single-Chassis System Logging Configuration Overview*.

Table 109: Minimum Configuration Statements for System Logging

Destination	Minimum Configuration Statements
File	<pre>[edit system syslog] file filename { facility severity; }</pre>
Terminal session of one, several, or all users	<pre>[edit system syslog] user (username *) { facility severity; }</pre>
Router or switch console	<pre>[edit system syslog] console { facility severity; }</pre>
Remote machine or the other Routing Engine on the router or switch	<pre>[edit system syslog] host (hostname other-routing-engine) { facility severity; }</pre>

Related Documentation

- [Junos OS System Log Overview](#)

Junos OS System Log Configuration Statements

To configure the switch to log system messages, include the **syslog** statement at the **[edit system]** hierarchy level:

```
[edit system]
syslog {
    archive <files number> <size size> <world-readable | no-world-readable>;
    console {
        facility severity;
    }
    file filename {
```

```

    facility severity;
    archive <archive-sites (ftp-url <password password>)> <files number> <size size>
        <start-time "YYYY-MM-DD.hh:mm"> <transfer-interval minutes> <world-readable |
        no-world-readable>;
    explicit-priority;
    match "regular-expression";
    structured-data {
        brief;
    }
}
host hostname {
    facility severity;
    explicit-priority;
    facility-override facility;
    log-prefix string
    match "regular-expression";
}
source-address source-address;
time-format (year | millisecond | year millisecond);
user (username | *) {
    facility severity;
    match "regular-expression";
}
}

```

Related Documentation • [Overview of Junos OS System Log Messages on page 1528](#)

Adding a Text String to System Log Messages

To add a text string to every system log message directed to a remote machine or to the other Routing Engine, include the **log-prefix** statement at the **[edit system syslog host]** hierarchy level:

```

[edit system syslog host (hostname | other-routing-engine)]
facility severity;
log-prefix string;

```

The string can contain any alphanumeric or special character except the equal sign (=) and the colon (:). It also cannot include the space character; do not enclose the string in quotation marks (" ") in an attempt to include spaces in it.

The Junos OS system logging utility automatically appends a colon and a space to the specified string when the system log messages are written to the log. The string is inserted after the identifier for the Routing Engine that generated the message.

The following example shows how to add the string M120 to all messages to indicate that the router is an M120 router, and direct the messages to the remote machine hardware-logger.mycompany.com:

```

[edit system syslog]
host hardware-logger.mycompany.com {
    any info;
    log-prefix M120;
}

```

When these configuration statements are included on an M120 router called origin1, a message in the system log on hardware-logger.mycompany.com looks like the following:

```
Mar 9 17:33:23 origin1 M120: mgd[477]: UI_CMDLINE_READ_LINE: user 'root', command 'run show version'
```

- Related Documentation**
- [Single-Chassis System Logging Configuration Overview](#)
 - [Specifying Log File Size, Number, and Archiving Properties on page 1543](#)

Directing System Log Messages to a Log File

To direct system log messages to a file in the `/var/log` directory of the local Routing Engine, include the `file` statement at the `[edit system syslog]` hierarchy level:

```
[edit system syslog]
file filename {
  facility severity;
  archive <archive-sites (ftp-url <password password>) > <files number> <size size>
    <start-time "YYYY-MM-DD.hh:mm"> <transfer-interval minutes> <world-readable |
    no-world-readable>;
  explicit-priority;
  match "regular-expression";
  structured-data {
    brief;
  }
}
```

For the list of facilities and severity levels, see [“Specifying the Facility and Severity of Messages to Include in the Log” on page 1544](#).

To prevent log files from growing too large, the Junos OS system logging utility by default writes messages to a sequence of files of a defined size. By including the `archive` statement, you can configure the number of files, their maximum size, and who can read them, either for all log files or for a certain log file. For more information, see [“Specifying Log File Size, Number, and Archiving Properties” on page 1543](#).

For information about the following statements, see the indicated sections:

- **explicit-priority**—See [“Including Priority Information in System Log Messages” on page 1536](#)
- **match**—See [“Using Regular Expressions to Refine the Set of Logged Messages” on page 1551](#)
- **structured-data**—See [Logging Messages in Structured-Data Format](#)

- Related Documentation**
- [Single-Chassis System Logging Configuration Overview](#)
 - [Examples: Configuring System Logging](#)

Directing System Log Messages to a Remote Machine

To direct system log messages to a remote machine, include the **host** statement at the **[edit system syslog]** hierarchy level:

```
[edit system syslog]
host (hostname | other-routing-engine) {
  facility severity;
  explicit-priority;
  facility-override facility;
  log-prefix string;
  match "regular-expression";
}
source-address source-address;
```

To direct system log messages to a remote machine, include the **host hostname** statement to specify the remote machine's IP version 4 (IPv4) address or fully qualified hostname. The remote machine must be running the standard **syslogd** utility. We do not recommend directing messages to another Juniper Networks switch. In each system log message directed to the remote machine, the hostname of the local Routing Engine appears after the timestamp to indicate that it is the source for the message.

For the list of logging facilities and severity levels to configure under the **host** statement, see [“Specifying the Facility and Severity of Messages to Include in the Log” on page 1544](#).

To record facility and severity level information in each message, include the **explicit-priority** statement. For more information, see [“Including Priority Information in System Log Messages” on page 1536](#).

For information about the **match** statement, see [“Using Regular Expressions to Refine the Set of Logged Messages” on page 1551](#).

When directing messages to remote machines, you can include the **source-address** statement to specify the IP address of the switch that is reported in the messages as their source. In each **host** statement, you can also include the **facility-override** statement to assign an alternative facility and the **log-prefix** statement to add a string to each message.

Related Documentation

- [Overview of Single-Chassis System Logging Configuration on page 1528](#)

Directing System Log Messages to a User Terminal

To direct system log messages to the terminal session of one or more specific users (or all users) when they are logged in to the local Routing Engine, include the **user** statement at the **[edit system syslog]** hierarchy level:

```
[edit system syslog]
user (username | *) {
  facility severity;
  match "regular-expression";
}
```

Specify one or more Junos OS usernames, separating multiple values with spaces, or use the asterisk (*) to indicate all users who are logged in to the local Routing Engine.

For the list of logging facilities and severity levels, see [“Specifying the Facility and Severity of Messages to Include in the Log” on page 1544](#). For information about the **match** statement, see [“Using Regular Expressions to Refine the Set of Logged Messages” on page 1551](#).

- Related Documentation**
- *Single-Chassis System Logging Configuration Overview*
 - *Examples: Configuring System Logging*

Directing System Log Messages to the Console

To direct system log messages to the console of the local Routing Engine, include the **console** statement at the **[edit system syslog]** hierarchy level:

```
[edit system syslog]
  console {
    facility severity;
  }
```

For the list of logging facilities and severity levels, see [“Specifying the Facility and Severity of Messages to Include in the Log” on page 1544](#).

- Related Documentation**
- *Single-Chassis System Logging Configuration Overview*
 - *Examples: Configuring System Logging*

Disabling the System Logging of a Facility

To disable the logging of messages that belong to a particular facility, include the **facility none** statement in the configuration. This statement is useful when, for example, you want to log messages that have the same severity level and belong to all but a few facilities. Instead of including a statement for each facility you want to log, you can include the **any severity** statement and then a **facility none** statement for each facility that you do not want to log. For example, the following logs all messages at the **error** level or higher to the console, except for messages from the **daemon** and **kernel** facilities. Messages from those facilities are logged to the file **>/var/log/internals** instead:

```
[edit system syslog]
  console {
    any error;
    daemon none;
    kernel none;
  }
  file internals {
    daemon info;
    kernel info;
  }
```

- Related Documentation**
- *Single-Chassis System Logging Configuration Overview*

Displaying a Log File from a Single-Chassis System

To display a log file stored on a single-chassis system, enter Junos OS CLI operational mode and issue the following commands:

```
user@switch> show log log-filename
user@switch> file show log-file-pathname
```

By default, the commands display the file stored on the local Routing Engine.

The following example shows the output from the **show log messages** command:

```
user@switch1> show log messages
Nov  4 11:30:01 switch1 newsyslog[2283]: logfile turned over due to size>128K
Nov  4 11:30:01 switch1 newsyslog[2283]: logfile turned over due to size>128K
Nov  4 11:30:06 switch1 chassism[952]: CM ENV Monitor: set fan speed is 65 percent
for Fan 1
Nov  4 11:30:06 switch1 chassism[952]: CM ENV Monitor: set fan speed is 65 percent
for Fan 2
Nov  4 11:30:06 switch1 chassism[952]: CM ENV Monitor: set fan speed is 65 percent
for Fan 3
...
Nov  4 11:52:53 switch1 snmpd[944]: SNMPD_HEALTH_MON_INSTANCE: Health Monitor:
jroute daemon memory usage (Management
process): new instance detected (variable: sysApp1ElmtRunMemory.5.6.2293)
Nov  4 11:52:53 switch1 snmpd[944]: SNMPD_HEALTH_MON_INSTANCE: Health Monitor:
jroute daemon memory usage (Command-line
interface): new instance detected (variable: sysApp1ElmtRunMemory.5.8.2292)
...
Nov  4 12:08:30 switch1 rpdf[957]: task_connect: task BGP_100.10.10.1.6+179 addr
10.10.1.6+179: Can't assign requested
address
Nov  4 12:08:30 switch1 rpdf[957]: bgp_connect_start: connect 10.10.1.6 (Internal
AS 100): Can't assign requested address
Nov  4 12:10:24 switch1 mgd[2293]: UI_CMDLINE_READ_LINE: User 'jsmith', command
'exit '
Nov  4 12:10:27 switch1 mgd[2293]: UI_DBASE_LOGOUT_EVENT: User 'jsmith' exiting
configuration mode
Nov  4 12:10:31 switch1 mgd[2293]: UI_CMDLINE_READ_LINE: User 'jsmith', command
'show log messages'
```

The following example shows the output from the **file show** command. The file in the pathname **/var/log/processes** has been previously configured to include messages from the daemon facility.

```
user@switch1> file show /var/log/processes
Feb 22 08:58:24 switch1 snmpd[359]: SNMPD_TRAP_WARM_START: trap_generate_warm:
SNMP trap: warm start
Feb 22 20:35:07 switch1 snmpd[359]: SNMPD_THROTTLE_QUEUE_DRAINED:
trap_throttle_timer_handler: cleared all throttled traps
Feb 23 07:34:56 switch1 snmpd[359]: SNMPD_TRAP_WARM_START: trap_generate_warm:
SNMP trap: warm start
Feb 23 07:38:19 switch1 snmpd[359]: SNMPD_TRAP_COLD_START: trap_generate_cold:
SNMP trap: cold start
...
```

Related Documentation

- [Interpreting Messages Generated in Standard Format on page 1542](#)

- [Interpreting Messages Generated in Structured-Data Format on page 1539](#)

Including Priority Information in System Log Messages

The facility and severity level of a message are together referred to as its *priority*. By default, messages logged in the standard Junos OS format do not include information about priority. To include priority information in standard-format messages directed to a file, include the **explicit-priority** statement at the **[edit system syslog file *filename*]** hierarchy level:

```
[edit system syslog file filename]  
facility severity;  
explicit-priority;
```



NOTE: Messages logged in structured-data format include priority information by default. If you include the **structured-data** statement at the **[edit system syslog file *filename*]** hierarchy level along with the **explicit-priority** statement, the **explicit-priority** statement is ignored and messages are logged in structured-data format.

For information about the **structured-data** statement, see *Logging Messages in Structured-Data Format*. For information about the contents of a structured-data message, see the *Junos OS System Log Reference for Security Devices*.

To include priority information in messages directed to a remote machine or the other Routing Engine, include the **explicit-priority** statement at the **[edit system syslog host (*hostname* | other-routing-engine)]** hierarchy level:

```
[edit system syslog host (hostname | other-routing-engine)]  
facility severity;  
explicit-priority;
```



NOTE: The **other-routing-engine** option does not apply to the QFX Series.

The priority recorded in a message always indicates the original, local facility name. If the **facility-override** statement is included for messages directed to a remote destination, the Junos OS system logging utility still uses the alternative facility name for the messages themselves when directing them to the remote destination. For more information, see [“Changing the Alternative Facility Name for System Log Messages Directed to a Remote Destination” on page 1549](#).

When the **explicit-priority** statement is included, the Junos OS logging utility prepends codes for the facility name and severity level to the message tag name, if the message has one:

```
FACILITY-severity[-TAG]
```


(The tag is a unique identifier assigned to some Junos OS system log messages.)

In the following example, the **CHASSISD_PARSE_COMPLETE** message belongs to the **daemon** facility and is assigned severity **info** (6):

```
Aug 21 12:36:30 router1 chassisd[522]: %DAEMON-6-CHASSISD_PARSE_COMPLETE:
Using new configuration
```

When the **explicit-priority** statement is not included, the priority does not appear in the message:

```
Aug 21 12:36:30 router1 chassisd[522]: CHASSISD_PARSE_COMPLETE: Using new
configuration
```

For more information about message formatting, see the *Junos OS System Log Reference for Security Devices*.

**Related
Documentation**

- *Single-Chassis System Logging Configuration Overview*
- *Examples: Configuring System Logging*

Including the Year or Millisecond in Timestamps

By default, the timestamp recorded in a standard-format system log message specifies the month, date, hour, minute, and second when the message was logged, as in the following example:

```
Aug 21 15:36:30
```

To include the year, the millisecond, or both, in the timestamp, include the **time-format** statement at the **[edit system syslog]** hierarchy level:

```
[edit system syslog]
time-format (year | millisecond | year millisecond);
```

The modified timestamp is used in messages directed to each destination configured by a **file**, **console**, or **user** statement at the **[edit system syslog]** hierarchy level, but not to destinations configured by a **host** statement.

The following example illustrates the format for a timestamp that includes both the millisecond (401) and the year (2010):

```
Aug 21 15:36:30.401 2010
```



NOTE: By default, messages logged in structured-data format include the year and millisecond. If you include the structured-data statement at the [edit system syslog file *filename*] hierarchy level along with the time-format statement, the time-format statement is ignored and messages are logged in structured-data format.

For information about the structured-data statement, see “[Logging Messages in Structured-Data Format](#)” on page 1538. For information about interpreting messages in a structured-data format, see “[Interpreting Messages Generated in Structured-Data Format](#)” on page 1539.

Logging Messages in Structured-Data Format

You can log messages to a file in structured-data format instead of the standard Junos OS format. The structured-data format provides more information without adding significant length, and makes it easier for automated applications to extract information from a message.

The structured-data format complies with Internet draft **draft-ietf-syslog-protocol-21.txt**. The draft establishes a standard message format regardless of the source or transport protocol for logged messages.

To output messages to a file in structured-data format, include the **structured-data** statement at the [edit system syslog file *filename*] hierarchy level:

```
[edit system syslog file filename]  
  facility severity;  
  structured-data {  
    brief;  
  }
```

The optional **brief** statement suppresses the English-language text that appears by default at the end of a message to describe the error or event. For information about the fields in a structured-data-format message, see “[Interpreting Messages Generated in Structured-Data Format](#)” on page 1539.

The structured format is used for all messages logged to the file that are generated by a Junos OS process or software library.



NOTE: If you include either or both of the explicit-priority and time-format statements along with the structured-data statement, they are ignored. These statements apply to the standard Junos OS system log format, not to structured-data format.

Interpreting Messages Generated in Structured-Data Format

By default, Junos OS processes and software libraries write messages to the system log file in structured-data format. For information about the **structured-data** statement, see *Logging Messages in Structured-Data Format*.

Structured-format makes it easier for automated applications to extract information from the message. In particular, the standardized format for reporting the value of variables (elements in the English-language message that vary depending on the circumstances that triggered the message) makes it easy for an application to extract those values.

The structured-data format for a message includes the following fields (which appear here on two lines only for legibility):

```
<priority code>version timestamp hostname process processID TAG [junos@2636.platform
variable-value-pairs] message-text
```

Table 110 describes the fields. If the system logging utility cannot determine the value in a particular field, a hyphen (-) appears instead.

Table 110: Fields in Structured-Data Messages

Field	Description	Examples
<i><priority code></i>	Number that indicates the facility and severity of a message. It is calculated by multiplying the facility number by 8 and then adding the numerical value of the severity. For a mapping of the numerical codes to facility and severity, see “Specifying the Facility and Severity of Messages to Include in the Log” on page 1544 .	<165> for a message from the pfe facility (facility=20) with severity notice (severity=5).
<i>version</i>	Version of the Internet Engineering Task Force (IETF) system logging protocol specification.	1 for the initial version
<i>timestamp</i>	Time when the message was generated, in one of two representations: <ul style="list-style-type: none"> YYYY-MM-DDTHH:MM:SS.MSZ is the year, month, day, hour, minute, second and millisecond in Universal Coordinated Time (UTC) YYYY-MM-DDTHH:MM:SS.MS+/-HH:MM is the year, month, day, hour, minute, second and millisecond in local time; the hour and minute that follows the plus sign (+) or minus sign (-) is the offset of the local time zone from UTC 	2007-02-15T09:17:15.719Z is 9:17 AM UTC on 15 February 2007. 2007-02-15T01:17:15.719-08:00 is the same timestamp expressed as Pacific Standard Time in the United States.
<i>hostname</i>	Name of the host that originally generated the message.	switch1

Table 110: Fields in Structured-Data Messages (*continued*)

Field	Description	Examples
<i>process</i>	Name of the Junos OS process that generated the message.	mgd
<i>processID</i>	UNIX process ID (PID) of the Junos process that generated the message.	3046
<i>TAG</i>	Junos OS system log message tag, which uniquely identifies the message.	UI_DBASE_LOGOUT_EVENT
<i>junos@2636.platform</i>	An identifier for the type of hardware platform that generated the message. The junos@2636 prefix indicates that the platform runs the Junos OS. It is followed by a dot-separated numerical identifier for the platform type.	junos@2636.1.1.1.2.18
<i>variable-value-pairs</i>	A variable-value pair for each element in the <i>message-text</i> string that varies depending on the circumstances that triggered the message. Each pair appears in the format <i>variable</i> = " <i>value</i> ".	username="regress"
<i>message-text</i>	English-language description of the event or error (omitted if the brief statement is included at the [edit system syslog file <i>filename</i> structured-data] hierarchy level).	User 'regress' exiting configuration mode

By default, the structured-data version of a message includes English text at the end, as in the following example (which appears on multiple lines only for legibility):

```
<165>1 2007-02-15T09:17:15.719Z router1 mgd 3046 UI_DBASE_LOGOUT_EVENT
[junos@2636.1.1.1.2.18 username="regress"] User 'regress' exiting configuration mode
```

When the brief statement is included at the [edit system syslog file *filename* structured-data] hierarchy level, the English text is omitted, as in this example:

```
<165>1 2007-02-15T09:17:15.719Z router1 mgd 3046 UI_DBASE_LOGOUT_EVENT
[junos@2636.1.1.1.2.18 username="regress"]
```

Table 111 maps the codes that appear in the *priority-code* field to facility and severity level.



NOTE: Not all of the facilities and severities listed in Table 111 can be included in statements at the [edit system syslog] hierarchy level (some are used by internal processes). For a list of the facilities and severity levels that can be included in the configuration, see “Specifying the Facility and Severity of Messages to Include in the Log” on page 1544.

Table 111: Facility and Severity Codes in the priority-code Field

Facility (number)	Severity emergency	alert	critical	error	warning	notice	info	debug
kernel (0)	1	1	2	3	4	5	6	7
user (1)	8	9	10	11	12	13	14	15
mail (2)	16	17	18	19	20	21	22	23
daemon (3)	24	25	26	27	28	29	30	31
authorization (4)	32	33	34	35	36	37	38	39
syslog (5)	40	41	42	43	44	45	46	47
printer (6)	48	49	50	51	52	53	54	55
news (7)	56	57	58	59	60	61	62	63
uucp (8)	64	65	66	67	68	69	70	71
clock (9)	72	73	74	75	76	77	78	79
authorization-private (10)	80	81	82	83	84	85	86	87
ftp (11)	88	89	90	91	92	93	94	95
ntp (12)	96	97	98	99	100	101	102	103
security (13)	104	105	106	107	108	109	110	111
console (14)	112	113	114	115	116	117	118	119
local0 (16)	128	129	130	131	132	133	134	135
dfc (17)	136	137	138	139	140	141	142	143
local2 (18)	144	145	146	147	148	149	150	151
firewall (19)	152	153	154	155	156	157	158	159
pfe (20)	160	161	162	163	164	165	166	167
conflict-log (21)	168	169	170	171	172	173	174	175
change-log (22)	176	177	178	179	180	181	182	183
interactive-commands (23)	184	185	186	187	188	189	190	191

Interpreting Messages Generated in Standard Format

The syntax of a standard-format message generated by a Junos OS process or subroutine library depends on whether it includes priority information:

- When the **explicit-priority** statement is included at the `[edit system syslog file filename]` or `[edit system syslog host hostname]` hierarchy level, a system log message has the following syntax:

```
timestamp message-source: %facility-severity-TAG: message-text
```

- When directed to the console or to users, or when the **explicit-priority** statement is not included for files or remote hosts, a system log message has the following syntax:

```
timestamp message-source: TAG: message-text
```

Table 112 describes the message fields.

Table 112: Fields in Standard-Format Messages

Field	Description
<i>timestamp</i>	Time at which the message was logged.
<i>message-source</i>	Identifier of the process or component that generated the message and the routing platform on which the message was logged. This field includes two or more subfields: hostname, process and process ID (PID). If the process does not report its PID, the PID is not displayed. The message source subfields are displayed in the following format: <i>hostname process[process-ID]</i>
<i>facility</i>	Code that specifies the facility to which the system log message belongs. For a mapping of codes to facility names, see Table: Facility Codes Reported in Priority Information in “Including Priority Information in System Log Messages” on page 1536.
<i>severity</i>	Numerical code that represents the severity level assigned to the system log message. For a mapping of codes to severity names, see Table: Numerical Codes for Severity Levels Reported in Priority Information in “Including Priority Information in System Log Messages” on page 1536.
<i>TAG</i>	Text string that uniquely identifies the message, in all uppercase letters and using the underscore (_) to separate words. The tag name begins with a prefix that indicates the generating software process or library. The entries in this reference are ordered alphabetically by this prefix. Not all processes on a routing platform use tags, so this field does not always appear.
<i>message-text</i>	Text of the message.

Specifying Log File Size, Number, and Archiving Properties

To prevent log files from growing too large, by default the Junos OS system logging utility writes messages to a sequence of files of a defined size. The files in the sequence are referred to as *archive* files to distinguish them from the *active* file to which messages are currently being written. The default maximum size depends on the platform type:

- 128 kilobytes (KB) for EX Series switches
- 1 megabyte (MB) for M Series, MX Series, and T Series routers
- 10 MB for TX Matrix or TX Matrix Plus routers
- 1 MB for the QFX Series

When an active log file called **logfile** reaches the maximum size, the logging utility closes the file, compresses it, and names the compressed archive file **logfile.0.gz**. The logging utility then opens and writes to a new active file called **logfile**. This process is also known as file rotation. When the new **logfile** reaches the configured maximum size, **logfile.0.gz** is renamed **logfile.1.gz**, and the new **logfile** is closed, compressed, and renamed **logfile.0.gz**. By default, the logging utility creates up to 10 archive files in this manner. When the maximum number of archive files is reached and when the size of the active file reaches the configured maximum size, the contents of the last archived file are overwritten by the current active file. The logging utility by default also limits the users who can read log files to the **root** user and users who have Junos OS **maintenance** permission.

Junos OS provides a configuration statement **log-rotate-frequency** that configures the system log file rotation frequency by configuring the time interval for checking the log file size. The frequency can be set to a value of 1 minute through 59 minutes. The default frequency is 15 minutes.

To configure the log rotation frequency, include the **log-rotate-frequency** statement at the **[edit system syslog]** hierarchy level.

You can include the **archive** statement to change the maximum size of each file, how many archive files are created, and who can read log files.

To configure values that apply to all log files, include the **archive** statement at the **[edit system syslog]** hierarchy level:

```
archive <files number> <size size> <world-readable | no-world-readable>;
```

To configure values that apply to a specific log file, include the **archive** statement at the **[edit system syslog file filename]** hierarchy level:

```
archive <archive-sites (ftp-url <password password>)> <files number> <size size>
  <start-time "YYYY-MM-DD.hh:mm"> <transfer-interval minutes> <world-readable |
  no-world-readable>;
```

archive-sites site-name specifies a list of archive sites that you want to use for storing files. The **site-name** value is any valid FTP URL to a destination. If more than one site name is configured, a list of archive sites for the system log files is created. When a file is archived, the router or switch attempts to transfer the file to the first URL in the list,

moving to the next site only if the transfer does not succeed. The log file is stored at the archive site with the specified log filename. For information about how to specify valid FTP URLs, see [Format for Specifying Filenames and URLs in Junos OS CLI Commands](#).

binary-data Mark file as containing binary data. This allows proper archiving of binary files, such as WTMP files (login records for UNIX based systems). To restore the default setting, include the **no-binary-data** statement.

files number specifies the number of files to create before the oldest file is overwritten. The value can be from 1 through 1000.

size size specifies the maximum size of each file. The value can be from 64 KB (64k) through 1 gigabyte (1g); to represent megabytes, use the letter **m** after the integer. There is no space between the digits and the **k**, **m**, or **g** units letter.

start-time "YYYY-MM-DD.hh:mm" defines the date and time in the local time zone for a one-time transfer of the active log file to the first reachable site in the list of sites specified by the **archive-sites** statement.

transfer-interval interval defines the amount of time the current log file remains open (even if it has not reached the maximum possible size) and receives new statistics before it is closed and transferred to an archive site. This interval value can be from 5 through 2880 minutes.

world-readable enables all users to read log files. To restore the default permissions, include the **no-world-readable** statement.

**Related
Documentation**

- *Single-Chassis System Logging Configuration Overview*
- *Examples: Configuring System Logging*
- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

Specifying the Facility and Severity of Messages to Include in the Log

Each system log message belongs to a facility, which groups together messages that either are generated by the same source (such as a software process) or concern a similar condition or activity (such as authentication attempts). Each message is also preassigned a *severity level*, which indicates how seriously the triggering event affects routing platform functions.

When you configure logging for a facility and destination, you specify a severity level for each facility. Messages from the facility that are rated at that level and higher are logged to the following destination:

```
[edit system syslog]
(console | file filename | host destination | user username) {
  facility severity ;
}
```


For more information about the destinations, see [“Directing System Log Messages to a User Terminal” on page 1533](#), and, [“Directing System Log Messages to the Console” on page 1534](#).

To log messages belonging to more than one facility to a particular destination, specify each facility and associated severity as a separate statement within the set of statements for the destination.

[Table 113](#) lists the Junos OS system logging facilities that you can specify in configuration statements at the `[edit system syslog]` hierarchy level.

Table 113: Junos OS System Logging Facilities

Facility	Type of Event or Error
any	All (messages from all facilities)
authorization	Authentication and authorization attempts
change-log	Changes to the Junos OS configuration
conflict-log	Specified configuration is invalid on the router type
daemon	Actions performed or errors encountered by system processes
dfc	Events related to dynamic flow capture
firewall	Packet filtering actions performed by a firewall filter
ftp	Actions performed or errors encountered by the FTP process
interactive-commands	Commands issued at the Junos OS command-line interface (CLI) prompt or by a client application such as a Junos XML protocol or NETCONF XML client
kernel	Actions performed or errors encountered by the Junos OS kernel
pfe	Actions performed or errors encountered by the Packet Forwarding Engine
user	Actions performed or errors encountered by user-space processes

[Table 114](#) lists the severity levels that you can specify in configuration statements at the `[edit system syslog]` hierarchy level. The levels from **emergency** through **info** are in order from highest severity (greatest effect on functioning) to lowest.

Unlike the other severity levels, the **none** level disables logging of a facility instead of indicating how seriously a triggering event affects routing functions. For more information, see [“Disabling the System Logging of a Facility” on page 1534](#).

Table 114: System Log Message Severity Levels

Severity Level	Description
any	Includes all severity levels
none	Disables logging of the associated facility to a destination
emergency	System panic or other condition that causes the router to stop functioning
alert	Conditions that require immediate correction, such as a corrupted system database
critical	Critical conditions, such as hard errors
error	Error conditions that generally have less serious consequences than errors at the emergency, alert, and critical levels
warning	Conditions that warrant monitoring
notice	Conditions that are not errors but might warrant special handling
info	Events or nonerror conditions of interest

Related Documentation

- [Junos OS System Logging Facilities and Message Severity Levels on page 1546](#)
- *Single-Chassis System Logging Configuration Overview*
- *Examples: Configuring System Logging*

Junos OS System Logging Facilities and Message Severity Levels

Table 113 lists the Junos OS system logging facilities that you can specify in configuration statements at the **[edit system syslog]** hierarchy level.

Table 115: Junos OS System Logging Facilities

Facility	Type of Event or Error
any	All (messages from all facilities)
authorization	Authentication and authorization attempts
change-log	Changes to the Junos OS configuration
conflict-log	Specified configuration is invalid on the router type
daemon	Actions performed or errors encountered by system processes
dfc	Events related to dynamic flow capture

Table 115: Junos OS System Logging Facilities (*continued*)

Facility	Type of Event or Error
firewall	Packet filtering actions performed by a firewall filter
ftp	Actions performed or errors encountered by the FTP process
interactive-commands	Commands issued at the Junos OS command-line interface (CLI) prompt or by a client application such as a Junos XML protocol or NETCONF XML client
kernel	Actions performed or errors encountered by the Junos OS kernel
pfe	Actions performed or errors encountered by the Packet Forwarding Engine
user	Actions performed or errors encountered by user-space processes

Table 114 lists the severity levels that you can specify in configuration statements at the **[edit system syslog]** hierarchy level. The levels from **emergency** through **info** are in order from highest severity (greatest effect on functioning) to lowest.

Unlike the other severity levels, the **none** level disables logging of a facility instead of indicating how seriously a triggering event affects routing functions. For more information, see “Disabling the System Logging of a Facility” on page 1534.

Table 116: System Log Message Severity Levels

Severity Level	Description
any	Includes all severity levels
none	Disables logging of the associated facility to a destination
emergency	System panic or other condition that causes the router to stop functioning
alert	Conditions that require immediate correction, such as a corrupted system database
critical	Critical conditions, such as hard errors
error	Error conditions that generally have less serious consequences than errors at the emergency, alert, and critical levels
warning	Conditions that warrant monitoring
notice	Conditions that are not errors but might warrant special handling
info	Events or nonerror conditions of interest

- Related Documentation**
- *Single-Chassis System Logging Configuration Overview*
 - *Examples: Configuring System Logging*

Default Facilities for System Log Messages Directed to a Remote Destination

Table 117 lists the default alternative facility name next to the Junos OS-specific facility name for which it is used. For facilities that are not listed, the default alternative name is the same as the local facility name.

Table 117: Default Facilities for Messages Directed to a Remote Destination

Junos OS–Specific Local Facility	Default Facility When Directed to Remote Destination
change-log	local6
conflict-log	local5
dfc	local1
firewall	local3
interactive-commands	local7
pfe	local4

- Related Documentation**
- *Single-Chassis System Logging Configuration Overview*

Junos OS System Log Alternate Facilities for Remote Logging

Table 118 lists the facilities that you can specify in the **facility-override** statement.

Table 118: Facilities for the facility-override Statement

Facility	Description
authorization	Authentication and authorization attempts
daemon	Actions performed or errors encountered by system processes
ftp	Actions performed or errors encountered by the FTP process
kernel	Actions performed or errors encountered by the Junos OS kernel
local0	Local facility number 0
local1	Local facility number 1

Table 118: Facilities for the facility-override Statement (*continued*)

Facility	Description
local2	Local facility number 2
local3	Local facility number 3
local4	Local facility number 4
local5	Local facility number 5
local6	Local facility number 6
local7	Local facility number 7
user	Actions performed or errors encountered by user-space processes

We do not recommend including the **facility-override** statement at the **[edit system syslog host other-routing-engine]** hierarchy level. It is not necessary to use alternative facility names when directing messages to the other Routing Engine, because its Junos OS system logging utility can interpret the Junos OS-specific names.

**Related
Documentation**

- *Examples: Assigning an Alternative Facility to System Log Messages Directed to a Remote Destination*
- *Single-Chassis System Logging Configuration Overview*

Changing the Alternative Facility Name for System Log Messages Directed to a Remote Destination

Some facilities assigned to messages logged on the local router or switch have Junos OS-specific names (see “[Junos OS System Logging Facilities and Message Severity Levels](#)” on page 1546). In the recommended configuration, a remote machine designated at the **[edit system syslog host *hostname*]** hierarchy level is not a Juniper Networks router or switch, so its syslogd utility cannot interpret the Junos OS-specific names. To enable the standard syslogd utility to handle messages from these facilities when messages are directed to a remote machine, a standard **localX** facility name is used instead of the Junos OS-specific facility name.

[Table 117](#) lists the default alternative facility name next to the Junos OS-specific facility name it is used for.

The syslogd utility on a remote machine handles all messages that belong to a facility in the same way, regardless of the source of the message (the Juniper Networks router or switch or the remote machine itself). For example, the following statements in the configuration of the router called **local-router** direct messages from the **authorization** facility to the remote machine *monitor.mycompany.com*:

```
[edit system syslog]
```

```
host monitor.mycompany.com {  
    authorization info;  
}
```

The default alternative facility for the local **authorization** facility is also **authorization**. If the syslogd utility on **monitor** is configured to write messages belonging to the **authorization** facility to the file `/var/log/auth-attempts`, then the file contains the messages generated when users log in to **local-router** and the messages generated when users log in to **monitor**. Although the name of the source machine appears in each system log message, the mixing of messages from multiple machines can make it more difficult to analyze the contents of the **auth-attempts** file.

To make it easier to separate the messages from each source, you can assign an alternative facility to all messages generated on **local-router** when they are directed to **monitor**. You can then configure the syslogd utility on **monitor** to write messages with the alternative facility to a different file from messages generated on **monitor** itself.

To change the facility used for all messages directed to a remote machine, include the **facility-override** statement at the `[edit system syslog host hostname]` hierarchy level:

```
[edit system syslog host hostname]  
    facility severity;  
    facility-override facility;
```

In general, it makes sense to specify an alternative facility that is not already in use on the remote machine, such as one of the **localX** facilities. On the remote machine, you must also configure the syslogd utility to handle the messages in the desired manner.

[Table 118](#) lists the facilities that you can specify in the **facility-override** statement.

We do not recommend including the **facility-override** statement at the `[edit system syslog host other-routing-engine]` hierarchy level. It is not necessary to use alternative facility names when directing messages to the other Routing Engine, because its Junos OS system logging utility can interpret the Junos OS-specific names.

The following example shows how to log all messages generated on the local router at the error level or higher to the local0 facility on the remote machine called **monitor.mycompany.com**:

```
[edit system syslog]  
host monitor.mycompany.com {  
    any error;  
    facility-override local0;  
}
```

The following example shows how to configure routers located in California and routers located in New York to send messages to a single remote machine called **central-logger.mycompany.com**. The messages from California are assigned to alternative facility **local0** and the messages from New York are assigned to alternative facility **local2**.

- Configure California routers to aggregate messages in the local0 facility:

```
[edit system syslog]  
host central-logger.mycompany.com {  
    change-log info;
```

```
    facility-override local0;
}
```

- Configure New York routers to aggregate messages in the local2 facility:

```
[edit system syslog]
host central-logger.mycompany.com {
  change-log info;
  facility-override local2;
}
```

On central-logger, you can then configure the system logging utility to write messages from the local0 facility to the file **change-log** and the messages from the local2 facility to the file **new-york-config**.

Related Documentation

- [Table 117](#)
- [Alternate Facilities for System Log Messages Directed to a Remote Destination on page 1548](#)
- *Examples: Assigning an Alternative Facility to System Log Messages Directed to a Remote Destination*

Using Regular Expressions to Refine the Set of Logged Messages

The predefined facilities group together related messages, but you can also use regular expression matching to specify more exactly which messages from a facility are logged to a file, a user terminal, or a remote destination.

To specify the text string that must (or must not) appear in a message for the message to be logged to a destination, include the **match** statement and specify the regular expression which the text string must match:

```
match "regular-expression";
```

You can include this statement at the following hierarchy levels:

- **[edit system syslog file *filename*]** (for a file)
- **[edit system syslog user (*username* | *)]** (for a specific user session or for all user sessions on a terminal)
- **[edit system syslog host (*hostname* | *other-routing-engine*)]** (for a remote destination)

In specifying the regular expression, use the notation defined in POSIX Standard 1003.2 for extended (modern) UNIX regular expressions. Explaining regular expression syntax is beyond the scope of this document, but POSIX standards are available from the Institute of Electrical and Electronics Engineers (IEEE, <http://www.ieee.org>).

[Table 119](#) specifies which character or characters are matched by some of the regular expression operators that you can use in the match statement. In the descriptions, the term *term* refers to either a single alphanumeric character or a set of characters enclosed in square brackets, parentheses, or braces.



NOTE: The match statement is not case-sensitive.

Table 119: Regular Expression Operators for the match Statement

Operator	Matches
. (period)	One instance of any character except the space.
* (asterisk)	Zero or more instances of the immediately preceding term.
+ (plus sign)	One or more instances of the immediately preceding term.
? (question mark)	Zero or one instance of the immediately preceding term.
(pipe)	One of the terms that appears on either side of the pipe operator.
! (exclamation point)	Any string except the one specified by the expression, when the exclamation point appears at the start of the expression. Use of the exclamation point is Junos OS-specific.
^ (caret)	Start of a line, when the caret appears outside square brackets. One instance of any character that does not follow it within square brackets, when the caret is the first character inside square brackets.
\$ (dollar sign)	End of a line.
[] (paired square brackets)	One instance of one of the enclosed alphanumeric characters. To indicate a range of characters, use a hyphen (-) to separate the beginning and ending characters of the range. For example, [a-z0-9] matches any letter or number.
() (paired parentheses)	One instance of the evaluated value of the enclosed term. Parentheses are used to indicate the order of evaluation in the regular expression.

Using Regular Expressions

Filter messages that belong to the **interactive-commands** facility, directing those that include the string **configure** to the terminal of the root user:

```
[edit system syslog]
user root {
  interactive-commands any;
  match ".*configure.*";
}
```

Messages like the following appear on the **root** user's terminal when a user issues a **configure** command to enter configuration mode:

```
timestamp router-name mgd[PID]: UI_CMDLINE_READ_LINE: User 'user', command
'configure private'
```


Filter messages that belong to the **daemon** facility and have a severity of **error** or higher, directing them to the file **/var/log/process-errors**. Omit messages generated by the SNMP process (**snmpd**), instead directing them to the file **/var/log/snmpd-errors**:

```
[edit system syslog]
file process-errors {
  daemon error;
  match "!(.*snmpd.*)";
}
file snmpd-errors {
  daemon error;
  match ".*snmpd.*";
}
```

- Related Documentation**
- *Single-Chassis System Logging Configuration Overview*
 - *Examples: Configuring System Logging*

Managing Host OS System Log and Core Files

On Junos OS switches with a host OS, the Junos OS might generate system log messages (also called *syslog messages*) to record events that occur on the switch, including the following:

- Routine operations, such as a user login into the configuration database.
- Failure and error conditions.
- Emergency or critical conditions, such as power-down of the switch due to excessive temperature.

On OCX Series switches:

- System log messages are logged in the **/var/log/dcpfe.log** file in the host OS in the following scenarios:
 - When the forwarding daemon is initialized.
 - Messages are tagged as emergency (LOG_EMERG). A copy of the message is also sent to the **/var/log** directory on the switch.
- Messages from processes are available on the host system in the **/var/log** directory. System log messages from the host chassis management process are recorded in the **lcmd.log** file in the **/var/log** directory.

On QFX switches with a host OS:

- The Junos OS and host OS record log messages for system and process events, and generate core files upon certain system failures.
- These files are stored in directories such as **/var/log** for log messages, and **/var/tmp** or **/var/crash** for core files, depending on the type of host OS running on the switch.

For diagnostic purposes, you can access these host OS system log and core files from the Junos OS CLI on the switch. You can also clean up directories where the host OS stores temporary log and other files.

This topic includes these sections:

- [Viewing Log Files On the Host OS System on page 1554](#)
- [Copying Log Files From the Host System To the Switch on page 1554](#)
- [Viewing Core Files On the Host OS System on page 1554](#)
- [Copying Core Files From the Host System To the Switch on page 1554](#)
- [Cleaning Up Temporary Files on the Host OS on page 1555](#)

Viewing Log Files On the Host OS System

To view a list of the log files created on the host OS, enter the following command:

```
user@switch> show app-engine logs
```

Copying Log Files From the Host System To the Switch

To copy log files from the host OS to the switch, enter the following command:

```
user@switch> request app-engine file-copy log from-jhost source to-vjunos destination
```

For example, to copy the *lcmd log* file to the switch, enter the following command:

```
user@switch> request app-engine file-copy log from-jhost lcmd.log to-vjunos /var/tmp
```

Viewing Core Files On the Host OS System

To view the list of core files generated and stored on the host OS system, enter the following command:

```
user@switch> show app-engine crash
```

The list might look like this example output:

```
Compute cluster: default-cluster
Compute node: default-node

Crash Info
=====
total 13480
-rw-r--r-- 1 root root 178046 Feb 14 23:08 localhost.lcmd.26653.1455520135.core.tgz

-rw-r--r-- 1 root root 4330343 Feb 15 00:45 localhost.dcpfe.7155.1455525926.core.tgz

-rw-r--r-- 1 root root 4285901 Feb 15 01:49 localhost.dcpfe.25876.1455529782.core.tgz

-rw-r--r-- 1 root root 4288508 Feb 15 02:39 localhost.dcpfe.713.1455532774.core.tgz
-rw-r--r-- 1 root root 264079 Feb 15 17:02 localhost.lcmd.1144.1455584540.core.tgz
```

Copying Core Files From the Host System To the Switch

To copy core files from the host OS to the switch, enter the following command:

```
user@switch> request app-engine file-copy crash from-jhost source to-vjunos
destination-dir-or-file-path
```

When the destination Junos OS path is a directory, the source filename is used by default. To rename the file at the destination, enter the destination argument as a full path including the desired filename.

For example, to copy the *localhost.lcmd.26653.1455520135.core.tgz* core archive file to the switch, enter the following command:

```
user@switch> request app-engine file-copy crash from-jhost
localhost.lcmd.26653.1455520135.core.tgz to-vjunos /var/tmp
```

To see the results on the switch, enter the following command:

```
user@switch> show system core-dumps
re0:
-----
-rw-r--r-- 1 root field 178046 Feb 15 17:15
/var/tmp/localhost.lcmd.26653.1455520135.core.tgz
total files: 1
```

Cleaning Up Temporary Files on the Host OS

To remove temporary files created on the host OS, enter the following command:

```
user@switch> request app-engine cleanup
```

For example, the following sample output on a switch with a Linux host OS shows cleanup of temporary files stored in */var/tmp*:

```
Compute cluster: default-cluster

Compute node: default-node

Cleanup (/var/tmp)
=====
```

Related Documentation • [Overview of Junos OS System Log Messages on page 1528](#)

Monitoring System Log Messages

Purpose Display system log messages about the QFX Series. By looking through a system log file for any entries pertaining to the interface that you are interested in, you can further investigate a problem with an interface on the switch.

Action To view system log messages:

```
user@switch1> show log messages
```

Sample Output

```
Nov  4 11:30:01 switch1 newsyslog[2283]: logfile turned over due to size>128K
Nov  4 11:30:01 switch1 newsyslog[2283]: logfile turned over due to size>128K
Nov  4 11:30:06 switch1 chassism[952]: CM ENV Monitor: set fan speed is 65 percent
```

```
for Fan 1
Nov  4 11:30:06 switch1 chassism[952]: CM ENV Monitor: set fan speed is 65 percent
for Fan 2
Nov  4 11:30:06 switch1 chassism[952]: CM ENV Monitor: set fan speed is 65 percent
for Fan 3
...
Nov  4 11:52:53 switch1 snmpd[944]: SNMPD_HEALTH_MON_INSTANCE: Health Monitor:
jroute daemon
memory usage (Management process): new instance detected (variable:
sysApp1ElmtRunMemory.5.6.2293)
Nov  4 11:52:53 switch1 snmpd[944]: SNMPD_HEALTH_MON_INSTANCE: Health Monitor:
jroute daemon
memory usage (Command-line interface): new instance detected (variable:
sysApp1ElmtRunMemory.5.8.2292)
...
Nov  4 12:10:24 switch1 mgd[2293]: UI_CMDLINE_READ_LINE: User 'jsmith', command
'exit '
Nov  4 12:10:27 switch1 mgd[2293]: UI_DBASE_LOGOUT_EVENT: User 'jsmith' exiting
configuration mode
Nov  4 12:10:31 switch1 mgd[2293]: UI_CMDLINE_READ_LINE: User 'jsmith', command
'show log messages'
```

Meaning The sample output shows the following entries in the **messages** file:

- A new log file was created when the previous file reached the maximum size of 128 kilobytes (KB).
- The fan speed for Fan 1, 2, and 3 is set at 65 percent.
- Health monitoring activity is detected.
- CLI commands were entered by the user jsmith.

**Related
Documentation**

- [Overview of Junos OS System Log Messages on page 1528](#)
- [Understanding the Implementation of System Log Messages on the QFabric System](#)
- [Example: Configuring System Log Messages on page 1559](#)
- [clear log on page 332](#)
- [show log on page 1003](#)
- [syslog on page 1704](#)

Examples: Configuring System Logging

The system log provides an excellent way of tracking all management activity on the switch by recording events such as user authentication, access authorization, and command execution. Logged command executions include commands entered by users at the CLI prompt or by client applications such as the Junos XML protocol or NETCONF XML client. Because system log files contain information about commands executed on the switch and the user who executed the commands, checking system log files for failed authentication events can help identify attempts to hack in to the switch. You can also analyze network activity by correlating executed commands with events and changes that occurred on the network at a particular time.

System log files are stored locally on the switch in the default `/var/log` directory.

The following example shows how to configure system log messages to record all commands entered by users and all authentication or authorization attempts. Logged commands include those entered by users at the CLI prompt and by client applications. Authentication and authorization attempts include events that are saved in the file named `cli-commands` and those that are sent to the terminal of a user who is logged in.

```
[edit system]
syslog {
  file cli-commands {
    interactive-commands info;
    authorization info;
  }
  user * {
    interactive-commands info;
    authorization info;
  }
}
```

The following example shows how to log all alarms state changes to the file `/var/log/alarms`:

```
[edit system]
syslog {
  file alarms {
    kernel warning;
  }
}
```

The following example shows how to configure the handling of messages of various types, as described in the comments. Information is logged to two files, to the terminal of user alex, to a remote machine, and to the console:

```
[edit system]
syslog {
  /* write all security-related messages to file /var/log/security */
  file security {
    authorization info;
    interactive-commands info;
  }
  /* write messages about potential problems to file /var/log/messages: */
  /* messages from "authorization" facility at level "notice"
  and above, */
  /* messages from all other facilities at level "warning" and above */
  file messages {
    authorization notice;
    any warning;
  }
  /* write all messages at level "critical" and above to terminal of user
  "alex" if */
  /* that user is logged in */
  user alex {
    any critical;
  }
}
```

```
/* write all messages from the &ldquo;daemon&rdquo; facility at level &ldquo;info&rdquo;
and above, and */
/* messages from all other facilities at level &ldquo;warning&rdquo; and above, to the
*/
/* machine monitor.mycompany.com */
host monitor.mycompany.com {
    daemon info;
    any warning;
}
/* write all messages at level &ldquo;error&rdquo; and above to the system console */
console {
    any error;
}
}
```

The following example shows how to configure the handling of messages generated when users issue Junos OS CLI commands, by specifying the interactive-commands facility at the info, notice, and warning severity levels:

```
[edit system]
file user-actions {
    interactive-commands info;
}
user philip {
    interactive-commands notice;
}
console {
    interactive-commands warning;
}
}
```

The following list describes the security levels used in the example:

- **info**—Logs a message when users issue any command at the CLI operational or configuration mode prompt. The example writes the messages to the file `/var/log/user-actions`.
- **notice**—Logs a message when users issue the configuration mode command **commit**. The example writes the messages to the terminal of user philip.
- **warning**—Logs a message when users issue a command that restarts a software process. The example writes the messages to the console.

**Related
Documentation**

- [Overview of Single-Chassis System Logging Configuration on page 1528](#)

Examples: Assigning an Alternative Facility

This topic contains examples of configuring system log messages to use an alternative facility for logging.

The following example shows how to log all messages generated on the switch at the **error** level or higher to the **local0** facility on the remote host called **monitor.mycompany.com**:

```
[edit system syslog]
host monitor.mycompany.com {
  any error;
  facility-override local0;
}
```

The following example contains two sets of statements that show how to configure switches located in California and in New York to send messages to a single remote host called **central-logger.mycompany.com**. The messages from California are assigned to alternative facility **local0** and the messages from New York are assigned to alternative facility **local2**.

- The following statements configure the California switch to aggregate messages in the **local0** facility:

```
[edit system syslog]
host central-logger.mycompany.com {
  change-log info;
  facility-override local0;
}
```

- The following statements configure the New York switch to aggregate messages in the **local2** facility:

```
[edit system syslog]
host central-logger.mycompany.com {
  change-log info;
  facility-override local2;
}
```

On the remote host named **central-logger** you can subsequently configure the system logging utility to write messages from the **local0** facility to one file (for example, **california-config**) and the messages from the **local2** facility to another file (for example, **new-york-config**).

Related Documentation

- [Alternate Facilities for System Log Messages Directed to a Remote Destination on page 1548](#)

Example: Configuring System Log Messages

The QFabric system monitors events that occur on its component devices and distributes system log messages about those events to all external system log message servers (hosts) that are configured. Component devices may include Node devices, Interconnect devices, Director devices, and the Virtual Chassis. Messages are stored for viewing only in the QFabric system database. To view the messages, issue the **show log** command.

This example describes how to configure system log messages on the QFabric system.

- [Requirements on page 1560](#)
- [Overview on page 1560](#)
- [Configuration on page 1560](#)

Requirements

This example uses the following hardware and software components:

- Junos OS Release 12.2
- QFabric system
- External servers that can be configured as system log message hosts

Overview

Component devices that generate system log message events may include Node devices, Interconnect devices, Director devices, and the control plane switches. The following configuration example includes these components in the QFabric system:

- Director software running on the Director group
- Control plane switches
- Interconnect device
- Multiple Node devices

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set system syslog host 10.1.1.12 any error
set system syslog file qflogs
set system syslog file qflogs structured-data brief
set system syslog file qflogs archive size 1g
```

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure system messages from the QFabric Director device:

1. Specify a host, any facility, and the **error** severity level.

```
[edit system syslog]
user@switch# set host 10.1.1.12 any error
```



NOTE: You can configure more than one system log message server (host). The QFabric system sends the messages to each server configured.

2. (Optional) Specify a filename to capture log messages.



NOTE: On the QFabric system, a syslog file named `messages` is configured implicitly with facility and severity levels of `any any` and a file size of 100 MBs. Therefore, you cannot specify the filename `messages` in your configuration, and automatic command completion does not work for that filename.

```
[edit system syslog]
user@switch# set file qflogs structured-data brief
user@switch# set file qflogs
```

3. (Optional) Configure the maximum size of your system log message archive file. This example specifies an archive size of 1 GB.

```
[edit system syslog]
user@switch# set file qflogs archive size 1g
```

Results From configuration mode, confirm your configuration by entering the `show system` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@switch# show system
syslog {
  file qflogs {
  }
  host 10.1.1.12 {
    any error;
  }
}
```

If you are done configuring the device, enter `commit` from configuration mode.

- Related Documentation**
- *Understanding the Implementation of System Log Messages on the QFabric System*
 - *syslog (QFabric System)*
 - [show log on page 1003](#)

PART 21

Configuration Statements

- [Network Management on page 1565](#)
- [Automation on page 1573](#)
- [sFlow Technology on page 1591](#)
- [SNMP on page 1601](#)
- [System Logging on page 1691](#)

CHAPTER 50

Network Management

- [connection-limit on page 1566](#)
- [destination-override on page 1567](#)
- [no-remote-trace on page 1567](#)
- [protocol-version on page 1568](#)
- [rate-limit on page 1569](#)
- [ssh on page 1570](#)
- [telnet on page 1571](#)
- [tracing on page 1572](#)

connection-limit

Syntax	<code>connection-limit <i>limit</i>;</code>
Hierarchy Level	<code>[edit system services finger],</code> <code>[edit system services ftp],</code> <code>[edit system services netconf ssh],</code> <code>[edit system services ssh],</code> <code>[edit system services telnet],</code> <code>[edit system services xnm-clear-text],</code> <code>[edit system services xnm-ssl]</code>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	Configure the maximum number of connections sessions for each type of system services (finger, ftp, ssh, telnet, xnm-clear-text, or xnm-ssl) per protocol (either IPv6 or IPv4).
Options	<p>limit—(Optional) Maximum number of established connections per protocol (either IPv6 or IPv4).</p> <p>Range: 1 through 250</p> <p>Default: 75</p>



NOTE: The actual number of maximum connections depends on the availability of system resources, and might be fewer than the configured `connection-limit` value if the system resources are limited.

Required Privilege	system—To view this statement in the configuration.
Level	system-control—To add this statement to the configuration.

Related Documentation	<ul style="list-style-type: none"> • Configuring clear-text or SSL Service for Junos XML Protocol Client Applications • Configuring DTCP-over-SSH Service for the Flow-Tap Application • Configuring Finger Service for Remote Access to the Router • Configuring FTP Service for Remote Access to the Router or Switch • Configuring SSH Service for Remote Access to the Router or Switch on page 1412 • Configuring Telnet Service for Remote Access to a Router or Switch
------------------------------	---

destination-override

Syntax	<code>destination-override { syslog host <i>ip-address</i>; }</code>
Hierarchy Level	[edit system tracing]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Override the system-wide configuration of the switch at the [edit system tracing] hierarchy level. This statement has no effect if system tracing is not configured.
Options	syslog —System process log files to send to the remote tracing host. <ul style="list-style-type: none"> • syslog—System process log files to send to the remote tracing host. • host <i>ip-address</i>—IP address to which to send tracing information.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding Tracing and Logging Operations on page 1409 • tracing on page 1572

no-remote-trace

Syntax	<code>no-remote-trace</code>
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the switch to disable remote tracing after remote tracing has been enabled.
Default	Remote tracing is disabled.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • tracing on page 1572

protocol-version

Syntax	<code>protocol-version <i>version</i>;</code>
Hierarchy Level	[edit system services ssh]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Description	Specify the secure shell (SSH) protocol version.
Default	v2—SSH protocol version 2 is the default, introduced in Junos OS Release 11.4.
Options	<i>version</i> —SSH protocol version: v1, v2, or both.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the SSH Protocol Version on page 1413

rate-limit

Syntax	<code>rate-limit <i>limit</i>;</code>
Hierarchy Level	<code>[edit system services finger],</code> <code>[edit system services ftp],</code> <code>[edit system services netconf ssh],</code> <code>[edit system services ssh],</code> <code>[edit system services telnet],</code> <code>[edit system services tftp-server],</code> <code>[edit system services xnm-clear-text],</code> <code>[edit system services xnm-ssl]</code>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	Configure the maximum number of connections attempts per minute, per protocol (either IPv6 or IPv4) on an access service. For example, a rate limit of 10 allows 10 IPv6 telnet session connection attempts per minute and 10 IPv4 telnet session connection attempts per minute.
Default	150 connections
Options	<p>rate-limit <i>limit</i>—(Optional) Maximum number of connection attempts allowed per minute, per IP protocol (either IPv4 or IPv6).</p> <p>Range: 1 through 250</p> <p>Default: 150</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <i>Configuring clear-text or SSL Service for Junos XML Protocol Client Applications</i>


ssh

Syntax	<pre>ssh { authentication-order [<i>authentication-methods</i>]; ciphers [<i>cipher-1 cipher-2 cipher-3 ...</i>]; client-alive-count-max <i>seconds</i>; client-alive-interval <i>seconds</i>; connection-limit <i>limit</i>; hostkey-algorithm <<i>algorithm</i> <i>no-algorithm</i>>; key-exchange <<i>algorithm</i>>; macs <<i>algorithm</i>>; max-sessions-per-connection <<i>number</i>>; no-passwords; no-public-keys; no-tcp-forwarding; protocol-version [<i>v1 v2</i>]; rate-limit <i>limit</i>; root-login (<i>allow</i> <i>deny</i> <i>deny-password</i>); } tcp-forwarding (JDM)</pre>
Hierarchy Level	[edit system services]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>client-alive-interval and client-alive-max-count statements introduced in Junos OS Release 12.2.</p> <p>no-passwords statement introduced in Junos OS Release 13.3.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.</p> <p>no-public-keys statement introduced in Junos OS release 15.1.</p> <p>tcp-forwarding statement introduced in Junos OS Release 15.1X53-D50 for the NFX250 Network Services Platform.</p>
Description	<p>Allow SSH requests from remote systems to the local router or switch.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring SSH Service for Remote Access to the Router or Switch on page 1412

telnet

Syntax	<pre>telnet { authentication-order [<i>authentication-methods</i>]; connection-limit <i>limit</i>; rate-limit <i>limit</i>; }</pre>
Hierarchy Level	[edit system services]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Provide Telnet connections from remote systems to the local router or switch.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Telnet Service for Remote Access to a Router or Switch</i>

tracing

Syntax	<pre>tracing { destination-override syslog host <i>ip-address</i>; }</pre>
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the switch to enable remote tracing to a specified host IP address.
<hr/>	
<div> NOTE: The tracing statement is not supported on the QFX3000 QFabric system.</div> <hr/>	
The following processes are supported:	
<ul style="list-style-type: none">• chassisd—Chassis-control process• eventd—Event-processing process• cosd—Class-of-service process	
If you enabled remote tracing but wish to disable it for specific processes on the switch, use the no-remote-trace statement at the [edit system <i>process-name</i> traceoptions] hierarchy level.	
Default	Remote tracing is disabled by default.
Options	destination-override syslog host <i>ip-address</i> —Overrides the global configuration for system tracing and has no effect if the tracing statement is not configured.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding Tracing and Logging Operations on page 1409• destination-override on page 1567

CHAPTER 51

Automation

- [allow-transients on page 1574](#)
- [apply-macro on page 1575](#)
- [checksum on page 1576](#)
- [command on page 1577](#)
- [commit on page 1578](#)
- [description on page 1579](#)
- [direct-access on page 1579](#)
- [file \(Commit Scripts\) on page 1580](#)
- [file \(Op Scripts\) on page 1581](#)
- [no-allow-url on page 1582](#)
- [op on page 1583](#)
- [optional on page 1584](#)
- [refresh \(Commit Scripts\) on page 1584](#)
- [refresh \(Op Scripts\) on page 1585](#)
- [refresh-from \(Commit Scripts\) on page 1585](#)
- [refresh-from \(Op Scripts\) on page 1586](#)
- [scripts on page 1587](#)
- [source \(Commit Scripts\) on page 1589](#)
- [source \(Op Scripts\) on page 1589](#)

allow-transients

Syntax	allow-transients;
Hierarchy Level	[edit system scripts commit]
Release Information	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For Junos OS commit scripts, enable transient configuration changes to be committed.
Default	Transient changes are disabled by default. If you do not include the allow-transients statement, and an enabled script generates transient changes, the command-line interface (CLI) generates an error message and the commit operation fails.
Required Privilege Level	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Generating a Persistent or Transient Configuration Change Using Commit Scripts</i>• <i>Creating a Commit Script Macro to Read the Custom Syntax and Generate Related Configuration Statements</i>

apply-macro

Syntax	<code>apply-macro <i>apply-macro-name</i> { <i>parameter-name parameter-value</i>; }</code>
Hierarchy Level	All hierarchy levels
Release Information	Statement introduced in Junos OS Release 7.4.
Description	<p>With commit script macros, use custom syntax in your configuration.</p> <p>Macros work by locating apply-macro statements that you include in the candidate configuration and using the values specified in the apply-macro statement as parameters to a set of instructions (the macro) defined in a commit script. The commit script alters your configuration from one that contains custom syntax into a full configuration containing standard Junos OS statements.</p> <p>In effect, your custom configuration syntax serves a dual purpose. The syntax allows you to simplify your configuration tasks, and it provides data (or <i>hooks</i>) that are used by a commit script macros.</p> <p>You can include the apply-macro statement at any level of the configuration hierarchy. You can include multiple apply-macro statements at each level of the configuration hierarchy; however, each must have a unique name.</p>
Options	<p><i>apply-macro-name</i>—Name of the apply-macro statement.</p> <p><i>parameter-name</i>—One or more parameters. Parameters can be any text you want to include in your configuration.</p> <p><i>parameter-value</i>—A value that corresponds to the parameter name. Parameter values can be any text you want to include in your configuration.</p>
Required Privilege Level	configure—To enter configuration mode; other required privilege levels depend on where the statement is located in the configuration hierarchy.
Related Documentation	<ul style="list-style-type: none"> • <i>Overview of Creating Custom Configuration Syntax with Commit Script Macros</i>

checksum

Syntax	<code>checksum (md5 sha-256 sha1) hash;</code>
Hierarchy Level	[edit event-options event-script file <i>filename</i>], [edit system scripts commit file filename], [edit system scripts op file filename]
Release Information	Statement introduced in Junos OS Release 9.5. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	For Junos OS commit scripts and op scripts, specify the MD5, SHA-1, or SHA-256 checksum hash. When it executes a local event, commit, or op script, Junos OS verifies the authenticity of the script by using the configured checksum hash.
Options	md5 hash —MD5 checksum of this script. sha-256 hash —SHA-256 checksum of this script. sha1 hash —SHA-1 checksum of this script.
Required Privilege Level	maintenance —To view this statement in the configuration. maintenance-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Checksum Hashes for a Commit Script</i>• <i>Configuring Checksum Hashes for an Event Script</i>• <i>Configuring Checksum Hashes for an Op Script</i>• <i>Executing an Op Script from a Remote Site</i>• file checksum md5 on page 344 command in the <i>System Basics and Services Command Reference</i>• file checksum sha-256 on page 346 command in the <i>System Basics and Services Command Reference</i>• file checksum sha1 on page 345 command in the <i>System Basics and Services Command Reference</i>

command

Syntax	<code>command filename-alias;</code>
Hierarchy Level	[edit system scripts op file filename]
Release Information	Statement introduced in Junos OS Release 7.6.
Description	For Junos OS op scripts, configure a filename alias for the script file. This allows you to run the script by referencing either the script filename or the filename alias.
Required Privilege Level	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Enabling an Op Script and Defining a Script Alias</i>

commit

Syntax	<pre>commit { allow-transients; dampen { dampen-options { cpu-factor <i>cpu-factor</i>; line-interval <i>line-interval</i>; time-interval <i>time-interval</i>; } } direct-access; file <i>filename</i> { checksum (md5 sha-256 sha1) <i>hash</i>; optional; refresh; refresh-from <i>url</i>; source <i>url</i>; } max-datasize; refresh; refresh-from <i>url</i>; traceoptions { file <<i>filename</i>> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i>; no-remote-trace; } }</pre>
Hierarchy Level	[edit system scripts]
Release Information	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	For Junos OS commit scripts, configure the commit-time scripting mechanism.
Options	The statements are explained separately.
Required Privilege Level	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"><i>Storing and Enabling Scripts</i>

description

Syntax	<code>description <i>descriptive-text</i>;</code>
Hierarchy Level	[edit system scripts op file filename] [edit system scripts op file filename arguments <i>argument-name</i>]
Release Information	Statement introduced in Junos OS Release 7.6.
Description	For Junos OS op scripts, provide a help-text string that appears in the command-line interface (CLI).
Required Privilege Level	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Help Text for Op Scripts</i> • <i>Declaring Arguments in Op Scripts</i> • file (Op Scripts) on page 1581

direct-access

Syntax	<code>direct-access;</code>
Hierarchy Level	[edit system scripts commit]
Release Information	Statement introduced in Junos OS Release 9.1. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify that commit scripts read input configurations directly from the database when inspecting these scripts for errors.
Required Privilege Level	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Executing Large Commit Scripts</i>

file (Commit Scripts)

Syntax	<pre>file <i>filename</i> { checksum (md5 sha-256 sha1) <i>hash</i>; optional; refresh; refresh-from <i>url</i>; source <i>url</i>; }</pre>
Hierarchy Level	[edit system scripts commit]
Release Information	Statement introduced in Junos OS Release 7.4.
Description	For Junos OS commit scripts, enable a commit script that is located in the <code>/var/db/scripts/commit</code> directory.
Options	<p><i>filename</i>—Name of an Extensible Stylesheet Language Transformations (XSLT) or Stylesheet Language Alternative Syntax (SLAX) file containing a commit script.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>maintenance—To view this statement in the configuration.</p> <p>maintenance-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"><i>Controlling Execution of Commit Scripts During Commit Operations</i>

file (Op Scripts)

Syntax	<pre> file <i>filename</i> { arguments { <i>argument-name</i> { description <i>descriptive-text</i>; } } checksum (md5 sha-256 sha1) <i>hash</i>; command <i>filename-alias</i>; description <i>descriptive-text</i>; refresh; refresh-from <i>url</i>; source <i>url</i>; } </pre>
Hierarchy Level	[edit system scripts op]
Release Information	Statement introduced in Junos OS Release 7.6.
Description	For Junos OS op scripts, enable an op script that is located in the <code>/var/db/scripts/op</code> directory.
Options	<p>filename—The name of an Extensible Stylesheet Language Transformations (XSLT) or Stylesheet Language Alternative Syntax (SLAX) file containing an op script.</p> <p>The statements are explained separately.</p>
Required Privilege Level	<p>maintenance—To view this statement in the configuration.</p> <p>maintenance-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Enabling an Op Script and Defining a Script Alias</i>

no-allow-url

Syntax	no-allow-url;
Hierarchy Level	[edit system scripts op]
Release Information	Statement introduced in Junos OS Release 10.0.
Description	For Junos OS op scripts, prohibit the remote execution of scripts. When you include this configuration statement, the op url operational mode command generates an error and does not permit you to execute the op script from a remote site.
Default	If you do not include the no-allow-url statement, authorized users can issue the op url command to execute op scripts from a remote site.
Required Privilege Level	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• file (Op Scripts) on page 1581• <i>Executing an Op Script from a Remote Site</i>

op

```

Syntax  op {
        dampen {
            dampen-options {
                cpu-factor cpu-factor;
                line-interval line-interval;
                time-interval time-interval;
            }
        }
        file filename {
            allow-commands "regular-expression";
            arguments {
                argument-name {
                    description descriptive-text;
                }
            }
            checksum (md5 | sha-256 | sha1) hash;
            command filename-alias;
            dampen {
                dampen-options {
                    cpu-factor cpu-factor;
                    line-interval line-interval;
                    time-interval time-interval;
                }
            }
            description descriptive-text;
            refresh;
            refresh-from url;
            source url;
        }
        max-datasize;
        no-allow-url
        refresh;
        refresh-from url;
        traceoptions {
            file <filename> <files number> <size size> <world-readable | no-world-readable>;
            flag flag;
            no-remote-trace;
        }
    }

```

Hierarchy Level [edit system [scripts](#)]

Release Information Statement introduced in Junos OS Release 7.6.

Description For Junos OS op scripts, configure an operation scripting mechanism.

Options The statements are explained separately.

Required Privilege Level maintenance—To view this statement in the configuration.
 maintenance-control—To add this statement to the configuration.

Related Documentation

- *Storing and Enabling Scripts*

optional

Syntax	optional;
Hierarchy Level	[edit system scripts commit file <i>filename</i>]
Release Information	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For Junos OS commit scripts, allow a commit operation to succeed even if the script specified in the file statement is missing from the /var/db/scripts/commit directory on the device.
Required Privilege Level	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Controlling Execution of Commit Scripts During Commit Operations</i>

refresh (Commit Scripts)

Syntax	refresh;
Hierarchy Level	[edit system scripts commit], [edit system scripts commit file <i>filename</i>]
Release Information	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For Junos OS commit scripts, overwrite the local copy of all enabled commit scripts or a single enabled script located in the /var/db/scripts/commit directory with the copy located at the source URL, as specified in the source statement at the same hierarchy level.
Required Privilege Level	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Using a Master Source Location for a Script</i>• refresh-from (Commit Scripts) on page 1585• source (Commit Scripts) on page 1589

refresh (Op Scripts)

Syntax	refresh;
Hierarchy Level	[edit system scripts op], [edit system scripts op file filename]
Release Information	Statement introduced in Junos OS Release 7.6.
Description	For Junos OS op scripts, overwrite the local copy of all enabled op scripts or a single enabled script located in the <code>/var/db/scripts/op</code> directory with the copy located at the source URL, specified in the source statement at the same hierarchy level.
Required Privilege Level	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Using a Master Source Location for a Script</i> • refresh-from (Op Scripts) on page 1586 • source (Op Scripts) on page 1589

refresh-from (Commit Scripts)

Syntax	refresh-from <i>url</i> ;
Hierarchy Level	[edit system scripts commit], [edit system scripts commit file filename]
Release Information	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For Junos OS commit scripts, overwrite the local copy of all enabled commit scripts or a single enabled script located in the <code>/var/db/scripts/commit</code> directory with the copy located at a URL other than the URL specified in the source statement.
Options	<i>url</i> —The source specified as a Hypertext Transfer Protocol (HTTP) URL, FTP URL, or secure copy (scp)-style remote file specification.
Required Privilege Level	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Using an Alternate Source Location for a Script</i> • refresh (Commit Scripts) on page 1584 • source (Commit Scripts) on page 1589

refresh-from (Op Scripts)

Syntax	<code>refresh-from url;</code>
Hierarchy Level	[edit system scripts op], [edit system scripts op file filename]
Release Information	Statement introduced in Junos OS Release 7.6.
Description	For Junos OS op scripts, overwrite the local copy of all enabled op scripts or a single enabled script located in the <code>/var/db/scripts/op</code> directory with the copy located at a URL other than the URL specified in the source statement.
Options	url —Source specified as a Hypertext Transfer Protocol (HTTP) URL, FTP URL, or secure copy (scp)-style remote file specification.
Required Privilege Level	maintenance —To view this statement in the configuration. maintenance-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Using an Alternate Source Location for a Script</i>• refresh (Op Scripts) on page 1585• source (Op Scripts) on page 1589

scripts

```
Syntax  scripts {
    commit {
        allow-transients;
        dampen {
            dampen-options {
                cpu-factor cpu-factor;
                line-interval line-interval;
                time-interval time-interval;
            }
        }
        direct-access;
        file filename {
            checksum (md5 | sha-256 | sha1) hash;
            optional;
            refresh;
            refresh-from url;
            source url;
        }
        max-datasize;
        refresh;
        refresh-from url;
        traceoptions {
            file <filename> <files number> <size size> <world-readable | no-world-readable>;
            flag flag;
            no-remote-trace;
        }
    }
    load-scripts-from-flash;
    op {
        dampen {
            dampen-options {
                cpu-factor cpu-factor;
                line-interval line-interval;
                time-interval time-interval;
            }
        }
        file filename {
            allow-commands "regular-expression";
            arguments {
                argument-name {
                    description descriptive-text;
                }
            }
            checksum (md5 | sha-256 | sha1) hash;
            command filename-alias;
            dampen {
                dampen-options {
                    cpu-factor cpu-factor;
                    line-interval line-interval;
                    time-interval time-interval;
                }
            }
        }
    }
}
```

```

    description descriptive-text;
    refresh;
    refresh-from url;
    source url;
  }
  max-datasize;
  no-allow-url
  refresh;
  refresh-from url;
  traceoptions {
    file <filename> <files number> <size size> <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
  }
}
snmp {
  file filename {
    checksum (md5 | sha-256 | sha1) hash;
    oid id {
      priority;
    }
    refresh;
    refresh-from;
    source;
  }
  max-datasize;
  refresh;
  refresh-from;
  traceoptions {
    file <filename> <files number> <size size> <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
  }
}
synchronize;
}

```

Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For Junos OS commit or op scripts, configure scripting mechanisms.
Options	The statements are explained separately.
Required Privilege Level	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> <i>Storing and Enabling Scripts</i>

source (Commit Scripts)

Syntax	<code>source url;</code>
Hierarchy Level	[edit system scripts commit file <i>filename</i>]
Release Information	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For Junos OS commit scripts, specify the location of the source file for an enabled script located in the <code>/var/db/scripts/commit</code> directory. When you include the refresh statement at the same hierarchy level and commit the configuration, the local copy is overwritten by the version stored at the specified URL.
Options	<i>url</i> —The source specified as an HTTP URL, FTP URL, or scp-style remote file specification.
Required Privilege Level	<i>maintenance</i> —To view this statement in the configuration. <i>maintenance-control</i> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Using a Master Source Location for a Script</i> • <i>Overview of Updating Scripts from a Remote Source</i> • refresh (Commit Scripts) on page 1584 • refresh-from (Commit Scripts) on page 1585

source (Op Scripts)

Syntax	<code>source url;</code>
Hierarchy Level	[edit system scripts op file <i>filename</i>]
Release Information	Statement introduced in Junos OS Release 7.6. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For Junos OS op scripts, specify the location of the source file for an enabled script located in the <code>/var/db/scripts/op</code> directory. When you include the refresh statement at the same hierarchy level, the local copy is overwritten by the version stored at the specified URL.
Options	<i>url</i> —Master source file for an op script specified as an HTTP URL, FTP URL, or scp-style remote file specification.
Required Privilege Level	<i>maintenance</i> —To view this statement in the configuration. <i>maintenance-control</i> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Using a Master Source Location for a Script</i> • refresh (Op Scripts) on page 1585 • refresh-from (Op Scripts) on page 1586

CHAPTER 52

sFlow Technology

- [agent-id on page 1591](#)
- [collector \(sFlow Technology\) on page 1592](#)
- [interfaces \(sFlow\) on page 1593](#)
- [polling-interval on page 1594](#)
- [sample-rate on page 1595](#)
- [sflow on page 1596](#)
- [source-ip on page 1597](#)
- [traceoptions \(sFlow Technology\) on page 1598](#)
- [udp-port on page 1599](#)

[agent-id](#)

Syntax	<code>agent-id <i>ip-address</i>;</code>
Hierarchy Level	[edit protocols sflow]
Release Information	Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the IP address of the sFlow agent. If you do not configure the sFlow agent ID, the IP address for the agent is dynamically created using the IP address of an interface configured on the QFX Series device.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring sFlow Technology on page 1457• sflow on page 1596

collector (sFlow Technology)

Syntax	<code>collector <i>ip-address</i> { <i>udp-port</i> <i>port-number</i>; }</code>
Hierarchy Level	[edit protocols sflow]
Release Information	Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Configure a remote collector for sFlow network traffic monitoring. The device sends sFlow UDP datagrams to the configured collector for analysis. You can configure up to four collectors on the device. You specify the IP address for each collector you configure.</p> <p>The remaining statement is explained separately.</p>
Options	<i>ip-address</i> —IP address of the collector.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring sFlow Technology on page 1457• Example: Monitoring Network Traffic Using sFlow Technology on page 1453

interfaces (sFlow)

Syntax	<pre>interfaces <i>interface-name</i> { polling-interval <i>seconds</i>; sample-rate <i>number</i>; }</pre>
Hierarchy Level	[edit protocols sflow]
Release Information	Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Configure sFlow network traffic monitoring on the specified interface on the device. You can configure sFlow parameters (polling interval, sample rate) with different values on different interfaces.</p> <p>The remaining statements are explained separately.</p>
Options	<i>interface-name</i> —Name of the interface on which to configure sFlow parameters.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring sFlow Technology on page 1457• Example: Monitoring Network Traffic Using sFlow Technology on page 1453

polling-interval

Syntax	<code>polling-interval <i>seconds</i>;</code>
Hierarchy Level	[edit protocols sflow], [edit protocols sflow interfaces <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the rate (in seconds) at which successive samples of interface statistics (counters) are taken.
Default	If no polling interval is configured for a particular interface, the device uses the global polling interval configured at the [edit protocols sflow] hierarchy level. If no global interval is configured, the device uses the default polling interval of 20 seconds.
Options	<i>seconds</i> —Number of seconds between successive samples of interface statistics. Specifying a value of 0 (zero) disables the polling. Range: 0 through 3600 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring sFlow Technology on page 1457• Example: Monitoring Network Traffic Using sFlow Technology on page 1453

sample-rate

Syntax	<code>sample-rate <i>number</i>;</code>
Hierarchy Level	[edit protocols sflow], [edit protocols sflow interfaces <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify the denominator (<i>number</i>) of the ratio that is the sample rate in sFlow traffic monitoring. For example, to configure a sample rate of 1 in 1000 packets, you specify a <i>number</i> of 1000.
Default	If no sample rate is configured for a particular interface, the device uses the global sample rate configured at the [edit protocols sflow] hierarchy level. If no global rate is configured, the device uses the default sample rate of 1 in 2000 packets.
Options	<i>number</i> —Denominator of the ratio representing the sample rate (one packet out of <i>number</i>). Range: 1 through 16,777,215
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring sFlow Technology on page 1457 • Example: Monitoring Network Traffic Using sFlow Technology on page 1453

sflow

Syntax	<pre>sflow { agent-id <i>ip-address</i>; collector <i>ip-address</i> { udp-port <i>port-number</i>; } interfaces <i>interface-name</i> { polling-interval <i>number</i>; sample-rate { egress <i>number</i>; ingress <i>number</i>; } } polling-interval <i>number</i>; sample-rate { egress <i>number</i>; ingress <i>number</i>; } source-ip <i>ip-address</i>; traceoptions { file <i>filename</i> <files <i>number</i>> <no-stamp> <replace> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i>; } }</pre>
Hierarchy Level	[edit protocols]
Release Information	Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Configure sFlow technology to monitor traffic continuously on specified interfaces simultaneously. sFlow data can be used to characterize network activity.</p> <p>The remaining statements are explained separately.</p>
Default	The sFlow protocol is disabled by default.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring sFlow Technology on page 1457• Example: Monitoring Network Traffic Using sFlow Technology on page 1453

source-ip

Syntax	<code>source-ip <i>ip-address</i>;</code>
Hierarchy Level	[edit protocols sflow]
Release Information	Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the source IP address to be used for sFlow datagrams. If you do not configure a source IP address, it is dynamically created based on the IP address of an Ethernet interface configured on the QFX Series device.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring sFlow Technology on page 1457• sflow on page 1596

traceoptions (sFlow Technology)

Syntax	<pre>traceoptions { file <i>filename</i> <files <i>number</i>> <no-stamp> <replace> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i>; }</pre>
Hierarchy Level	[edit protocols sflow]
Release Information	Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Define tracing operations for sFlow technology.
Default	The traceoptions feature is disabled.
Options	<p>file <i>filename</i>—Name of the file to receive the tracing operation output. Enclose the name in quotation marks. Output files are located in the <code>/var/log/</code> directory.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0. Incoming trace file data is logged in the now empty trace-file. When trace-file again reaches its maximum size, trace-file.0 is renamed trace-file.1 and trace-file is renamed trace-file.0. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify the maximum number of files, you must also specify the maximum file size using the size option.</p> <p>Range: 2 through 1000 files Default: 1 trace file</p> <p>flag <i>flag</i>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements.</p> <ul style="list-style-type: none">• all—Trace all sFlow monitoring events.• client-server—Trace sFlow monitoring client-server events.• configuration—Trace sFlow monitoring configuration events.• interface—Trace sFlow monitoring interface events.• rtsock—Trace routing socket code events. <p>no-stamp—(Optional) Do not place timestamp information at the beginning of each line in the trace file.</p> <p>no-world-readable—(Optional) Prevent any user from reading the trace file.</p> <p>replace—(Optional) Replace an existing trace file if there is one.</p>

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**. Incoming trace file data is logged in the now empty **trace-file**. When **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option.

Syntax: **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

Range: 10 KB through the maximum file size of 4 GB

Default: 128 KB

world-readable—(Optional) Allow any user to read the trace file.

Required Privilege Level	routing and trace—To view this statement in the configuration.
	routing-control and trace-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Overview of sFlow Technology

udp-port

Syntax	<code>udp-port <i>port-number</i>;</code>
Hierarchy Level	[edit protocols sflow collector]
Release Information	Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the UDP port for a remote collector for sFlow network traffic monitoring. The device sends sFlow UDP datagrams to the collector for analysis.
Default	Port 6343
Options	<i>port-number</i> —UDP port number for this collector.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring sFlow Technology on page 1457 • Example: Monitoring Network Traffic Using sFlow Technology on page 1453

CHAPTER 53

SNMP

- [access \(SNMP\) on page 1604](#)
- [address \(SNMP\) on page 1604](#)
- [address-mask on page 1605](#)
- [agent-address on page 1605](#)
- [alarm \(SNMP RMON\) on page 1606](#)
- [authentication-md5 on page 1607](#)
- [authentication-none on page 1608](#)
- [authentication-password on page 1609](#)
- [authentication-sha on page 1610](#)
- [authorization on page 1611](#)
- [bucket-size on page 1612](#)
- [categories on page 1613](#)
- [client-list on page 1613](#)
- [client-list-name on page 1614](#)
- [clients on page 1614](#)
- [commit-delay on page 1615](#)
- [community \(SNMP\) on page 1616](#)
- [community \(RMON\) on page 1617](#)
- [community-name on page 1618](#)
- [contact on page 1619](#)
- [description on page 1619](#)
- [description \(RMON\) on page 1620](#)
- [destination-port \(SNMP\) on page 1620](#)
- [engine-id on page 1621](#)
- [event on page 1622](#)
- [falling-event-index \(RMON\) on page 1623](#)
- [falling-threshold \(Health Monitor\) on page 1624](#)
- [falling-threshold \(RMON\) on page 1625](#)

- [falling-threshold-interval](#) on page 1626
- [filter-duplicates](#) on page 1626
- [filter-interfaces](#) on page 1627
- [group \(Defining Access Privileges for an SNMPv3 Group\)](#) on page 1627
- [group \(Configuring Group Name\)](#) on page 1628
- [health-monitor](#) on page 1629
- [history](#) on page 1630
- [interface \(SNMP\)](#) on page 1631
- [interface \(RMON\)](#) on page 1632
- [interval \(Health Monitor\)](#) on page 1633
- [interval \(RMON\)](#) on page 1634
- [local-engine](#) on page 1635
- [location](#) on page 1636
- [message-processing-model](#) on page 1636
- [name](#) on page 1637
- [nonvolatile](#) on page 1637
- [notify](#) on page 1638
- [notify-filter \(Applying to the Management Target\)](#) on page 1639
- [notify-filter \(Configuring the Profile Name\)](#) on page 1639
- [notify-view](#) on page 1640
- [oid](#) on page 1640
- [oid \(SNMPv3\)](#) on page 1641
- [owner](#) on page 1642
- [parameters](#) on page 1642
- [port](#) on page 1643
- [privacy-3des](#) on page 1644
- [privacy-aes128](#) on page 1645
- [privacy-des](#) on page 1646
- [privacy-none](#) on page 1646
- [privacy-password](#) on page 1647
- [read-view](#) on page 1648
- [remote-engine](#) on page 1649
- [request-type](#) on page 1650
- [retry-count](#) on page 1651
- [rising-event-index](#) on page 1652
- [rising-threshold \(Health Monitor\)](#) on page 1653
- [rising-threshold \(RMON\)](#) on page 1654

- [rmon](#) on page 1655
- [sample-type](#) on page 1656
- [security-level \(Defining Access Privileges\)](#) on page 1657
- [security-level \(Generating SNMP Notifications\)](#) on page 1658
- [security-model \(Access Privileges\)](#) on page 1659
- [security-model \(Group\)](#) on page 1660
- [security-model \(SNMP Notifications\)](#) on page 1660
- [security-name \(Community String\)](#) on page 1661
- [security-name \(Security Group\)](#) on page 1662
- [security-name \(SNMP Notifications\)](#) on page 1663
- [security-to-group](#) on page 1664
- [snmp](#) on page 1665
- [snmp-community](#) on page 1669
- [source-address \(SNMP\)](#) on page 1669
- [startup-alarm](#) on page 1670
- [syslog-subtag](#) on page 1671
- [tag \(Configuring Notification Targets\)](#) on page 1671
- [tag \(Configuring the SNMP Community\)](#) on page 1672
- [tag-list](#) on page 1672
- [target-address](#) on page 1673
- [target-parameters](#) on page 1674
- [targets](#) on page 1675
- [timeout](#) on page 1675
- [traceoptions \(SNMP\)](#) on page 1676
- [trap-group](#) on page 1678
- [trap-options](#) on page 1679
- [type \(RMON Notification\)](#) on page 1680
- [type](#) on page 1681
- [user](#) on page 1681
- [usm](#) on page 1682
- [v3](#) on page 1684
- [vacm](#) on page 1686
- [variable](#) on page 1687
- [version](#) on page 1688
- [view \(Configuring a MIB View\)](#) on page 1689
- [view \(Associating MIB View with a Community\)](#) on page 1690
- [write-view](#) on page 1690

access (SNMP)

Syntax	<pre>access { group group-name { (default-context-prefix context-prefix context-prefix) { security-model (any usm v1 v2c) { security-level (authentication none privacy) { notify-view view-name; read-view view-name; write-view view-name; } } } } }</pre>
Hierarchy Level	[edit snmp v3 vacm]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Set SNMP access limits. The remaining statements are explained separately.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.

address (SNMP)

Syntax	<pre>address address;</pre>
Hierarchy Level	[edit snmp v3 target-address target-address-name]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify the SNMP target address for receiving traps or informs.
Options	address —IPv4 address of the system to receive traps or informs. You must specify an address, not a hostname.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding the Implementation of SNMP on the QFabric System• Configuring SNMP on page 1499• Example: Configuring SNMP on page 1523

address-mask

Syntax	<code>address-mask <i>address-mask</i>;</code>
Hierarchy Level	<code>[edit snmp v3 target-address <i>target-address-name</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 on the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Verify the source addresses for a group of target addresses.
Options	<i>address-mask</i> combined with the address defines a range of addresses.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the Address Mask</i>

agent-address

Syntax	<code>agent-address outgoing-interface;</code>
Hierarchy Level	<code>[edit snmp trap-options]</code>
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Description	Set the agent address of all SNMPv1 traps generated by this router or switch. Currently, the only option is outgoing-interface , which sets the agent address of each SNMPv1 trap to the address of the outgoing interface of that trap.
Options	outgoing-interface —Value of the agent address of all SNMPv1 traps generated by this router or switch. The outgoing-interface option sets the agent address of each SNMPv1 trap to the address of the outgoing interface of that trap. Default: Disabled (the agent address is not specified in SNMPv1 traps).
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the Agent Address for SNMP Traps</i>

alarm (SNMP RMON)

Syntax alarm *index* {
 description *description*;
 falling-event-index *index*;
 falling-threshold *integer*;
 falling-threshold-interval *seconds*;
 interval *seconds*;
 request-type (get-next-request | get-request | walk-request);
 rising-event-index *index*;
 rising-threshold *integer*;
 sample-type (absolute-value | delta-value);
 startup-alarm (falling-alarm | rising-alarm | rising-or-falling alarm);
 syslog-subtag *syslog-subtag*;
 variable *oid-variable*;
 }

Hierarchy Level [edit snmp rmon]

Release Information Statement introduced before Junos OS Release 7.4.
 Statement introduced in Junos OS Release 9.0 for EX Series switches.
 Statement introduced in Junos OS Release 11.1 for the QFX Series.
 Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure RMON alarm entries.

Options *index*—Identifies this alarm entry as an integer.

 The remaining statements are explained separately.

Required Privilege Level snmp—To view this statement in the configuration.
 snmp-control—To add this statement to the configuration.

Related Documentation

- [Configuring an RMON Alarm Entry and Its Attributes](#)
- [event \(SNMP\)](#)
- [RMON MIB Event, Alarm, Log, and History Control Tables on page 1464](#)
- [Monitoring RMON MIB Tables on page 1516](#)
- [Understanding RMON on page 1463](#)

authentication-md5

Syntax	authentication-md5 { authentication-password authentication-password; }
Hierarchy Level	[edit snmp v3 usm local-engine user <i>username</i>], [edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure MD5 as the authentication type for the SNMPv3 user.



NOTE: You can only configure one authentication type for each SNMPv3 user.

The remaining statement is explained separately.

Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> Configuring MD5 Authentication

authentication-none

Syntax	authentication-none;
Hierarchy Level	[edit snmp v3 usm local-engine user <i>username</i>], [edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure that there should be no authentication for the SNMPv3 user.



NOTE: You can configure only one authentication type for each SNMPv3 user.

Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring No Authentication</i>

authentication-password

Syntax	<code>authentication-password <i>authentication-password</i>;</code>
Hierarchy Level	[edit snmp v3 usm local-engine user <i>username</i> authentication-md5], [edit snmp v3 usm local-engine user <i>username</i> authentication-sha], [edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i> authentication-md5], [edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i> authentication-sha]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the password for user authentication.
Options	<p><i>authentication-password</i>—Password that a user enters. The password is then converted into a key that is used for authentication.</p> <p>SNMPv3 has special requirements when you create plain-text passwords on a router or switch:</p> <ul style="list-style-type: none"> • The password must be at least eight characters long. • The password can include alphabetic, numeric, and special characters, but it cannot include control characters.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring MD5 Authentication</i> • <i>Configuring SHA Authentication</i>

authentication-sha

Syntax	<code>authentication-sha { authentication-password authentication-password; }</code>
Hierarchy Level	[edit snmp v3 usm local-engine user <i>username</i>], [edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the secure hash algorithm (SHA) as the authentication type for the SNMPv3 user.




NOTE: You can configure only one authentication type for each SNMPv3 user.

The remaining statement is explained separately.

Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring SHA Authentication</i>

authorization

Syntax	<code>authorization <i>authorization</i>;</code>
Hierarchy Level	<code>[edit snmp community <i>community-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Description	Set the access authorization for SNMP Get , GetBulk , GetNext , and Set requests.
Options	<p><i>authorization</i>—Access authorization level:</p> <ul style="list-style-type: none"> • read-only—Enable Get, GetNext, and GetBulk requests. • read-write—Enable all requests, including Set requests. You must configure a view to enable Set requests.
	<div>  <p>NOTE: The read-write option is not supported on the QFX3000 QFabric system.</p> </div>
	Default: read-only
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the SNMP Community String on page 1503

bucket-size

Syntax	<code>bucket-size <i>number</i>;</code>
Hierarchy Level	[edit snmp rmon history <i>index</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the sampling of Ethernet statistics for network fault diagnosis, planning, and performance tuning.
Default	50
Options	<i>number</i> —Number of discrete samples of Ethernet statistics requested.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• RMON MIB Event, Alarm, Log, and History Control Tables on page 1464• Configuring RMON Alarms and Events on page 1507• Monitoring RMON MIB Tables on page 1516• Understanding RMON on page 1463• Junos OS Network Management Configuration Guide

categories

Syntax	<pre>categories { category; }</pre>
Hierarchy Level	[edit snmp trap-group <i>group-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Description	Define the types of traps that are sent to the targets of the named trap group.
Default	If you omit the categories statement, all trap types are included in trap notifications.
Options	category —Name of a trap type: authentication , chassis , configuration , link , remote-operations , rmon-alarm , or startup .
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring SNMP Trap Groups on page 1504

client-list

Syntax	<pre>client-list <i>client-list-name</i> { ip-addresses; }</pre>
Hierarchy Level	[edit snmp]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Description	Define a list of SNMP clients.
Options	client-list-name —Name of the client list. ip-addresses —IP addresses of the SNMP clients to be added to the client list,
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Adding a Group of Clients to an SNMP Community on page 1505

client-list-name

Syntax	<code>client-list-name <i>client-list-name</i>;</code>
Hierarchy Level	<code>[edit snmp community <i>community-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Add a client list or prefix list to an SNMP community.
Options	<i>client-list-name</i> —Name of the client list or prefix list.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Adding a Group of Clients to an SNMP Community on page 1505


clients

Syntax	<pre>clients { <i>address</i> <restrict>; }</pre>
Hierarchy Level	<code>[edit snmp community <i>community-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify the IPv4 or IPv6 addresses of the SNMP client hosts that are authorized to use this community.
Default	If you omit the clients statement, all SNMP clients using this community string are authorized to access the switch.
Options	<i>address</i> —Address of an SNMP client that is authorized to access this switch. You must specify an address, not a hostname. To specify more than one client, include multiple <i>address</i> options. <i>restrict</i> —(Optional) Do not allow the specified SNMP client to access the switch.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring SNMP Communities

commit-delay

Syntax	commit-delay <i>seconds</i> ;
Hierarchy Level	[edit snmp nonvolatile]
Release Information	Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the timer for the SNMP Set reply and start of the commit.
Options	seconds —Delay between an affirmative SNMP Set reply and start of the commit operation. Default: 5 seconds
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Commit Delay Timer</i>


community (SNMP)

Syntax	<pre>community <i>community-name</i> { authorization <i>authorization</i>; client-list-name <i>client-list-name</i>; clients { address restrict; } view <i>view-name</i>; }</pre>
Hierarchy Level	[edit snmp]
Release Information	Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Define an SNMP community. An SNMP community authorizes SNMP clients based on the source IP address of incoming SNMP request packets. A community also defines which MIB objects are available and the operations (read-only or read-write) allowed on those objects.
<div> NOTE: The authorization read-write option is not supported on the QFX3000 QFabric system.</div>	
The SNMP client application specifies an SNMP community name in Get , GetBulk , GetNext , and Set SNMP requests.	
Default	If you omit the community statement, all SNMP requests are denied.
Options	<i>community-name</i> —Community string. If the name includes spaces, enclose it in quotation marks (" "). The remaining statements are explained separately.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the SNMP Community String on page 1503

community (RMON)

Syntax	<code>community <i>community-name</i>;</code>
Hierarchy Level	[edit snmp rmon event <i>index</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Configure the SNMP trap group that is used when generating a trap (if the eventType object is configured to send traps). If that trap group has the rmon-alarm trap category configured, a trap is sent to all the targets configured for that trap group. The community string in the trap matches the name of the trap group (and hence, the value of eventCommunity). If nothing is configured, traps are sent to each group that has the rmon-alarm category configured.</p> <p>The event community is not the same as an SNMP community.</p>
Options	<i>community-name</i> —Name of the trap group that is used when generating a trap if the event is configured to send traps.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • RMON MIB Event, Alarm, Log, and History Control Tables on page 1464 • Configuring RMON Alarms and Events on page 1507 • Monitoring RMON MIB Tables on page 1516 • Understanding RMON on page 1463 • Junos OS Network Management Configuration Guide

community-name

Syntax	<code>community-name <i>community-name</i>;</code>
Hierarchy Level	<code>[edit snmp v3 snmp-community <i>community-index</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	The community name defines an SNMP community. The SNMP community authorizes SNMPv1 or SNMPv2 clients. The access privileges associated with the configured security name define which MIB objects are available and the operations (notify, read, or write) allowed on those objects.
Options	<i>community-name</i> —Community string for an SNMPv1 or SNMPv2c community. If unconfigured, it is the same as the community index. If the name includes spaces, enclose it in quotation marks (" ").
<div> NOTE: Community names must be unique. You cannot configure the same community name at the <code>[edit snmp community]</code> and <code>[edit snmp v3 snmp-community <i>community-index</i>]</code> hierarchy levels.</div> <p>The community name at the <code>[edit snmp v3 snmp-community <i>community-index</i>]</code> hierarchy level is encrypted and not displayed in the command-line interface (CLI).</p>	
Required Privilege Level	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the SNMPv3 Community</i>

contact

Syntax	<code>contact <i>contact</i>;</code>
Hierarchy Level	[edit snmp]
Release Information	Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Define the value of the MIB II sysContact object, which is the contact person for the managed system.
Options	contact —Name of the contact person. If the name includes spaces, enclose it in quotation marks (" ").
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> <i>Configuring the System Contact on a Device Running Junos OS</i>

description

Syntax	<code>description <i>description</i>;</code>
Hierarchy Level	[edit snmp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Define the value of the MIB II sysDescription object, which is the description of the system being managed.
Options	description —System description. If the name includes spaces, enclose it in quotation marks (" ").
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> <i>Configuring the System Description on a Device Running Junos OS</i>


description (RMON)

Syntax	<code>description</code> <i>description</i> ;
Hierarchy Level	[edit snmp rmon alarm <i>index</i>], [edit snmp rmon event <i>index</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Text description of alarm or event.
Options	<i>description</i> —Text description of an alarm or event entry. If the description includes spaces, enclose it in quotation marks (" ").
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• RMON MIB Event, Alarm, Log, and History Control Tables on page 1464• Configuring RMON Alarms and Events on page 1507• Monitoring RMON MIB Tables on page 1516• Understanding RMON on page 1463• Junos OS Network Management Configuration Guide

destination-port (SNMP)

Syntax	<code>destination-port</code> <i>port-number</i> ;
Hierarchy Level	[edit snmp trap-group]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Assign a trap port number other than the default.
Default	If you omit this statement, the default port is 162.
Options	<i>port-number</i> —SNMP trap port number.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring SNMP Trap Groups on page 1504

engine-id

Syntax	engine-id { (local <i>engine-id-suffix</i> use-default-ip-address use-mac-address); }
Hierarchy Level	[edit snmp]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Define a unique identifier for an SNMPv3 engine by configuring the suffix of the engine ID. The engine ID is used for identification only and not for addressing. There are two parts of an engine ID: the prefix and the suffix. The prefix is formatted according to the specifications defined in RFC 3411, <i>An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks</i> and cannot be configured. The suffix is configured here.
<div>  <p>NOTE: SNMPv3 authentication and encryption keys are generated based on the associated user passwords and the engine ID. If you configure or change the engine ID, you must commit the user passwords and new engine ID before you configure SNMPv3 users, or the authentication will fail.</p> <p>By default, the engine ID suffix is configured with the MAC address of the management interface (the <i>use-mac-address</i> option) on the QFX Series and OCX Series. You can override this configuration by using the local <i>engine-id-suffix</i> or <i>use-default-ip-address</i> option.</p> </div>	
Default	use-mac-address
Options	<p>local <i>engine-id-suffix</i>—The engine ID suffix is set based on the data entered.</p> <p>use-default-ip-address—The engine ID suffix is generated from the default IP address.</p> <p>use-mac-address—The engine ID suffix is generated from the MAC address of the management interface on the switch.</p>
Required Privilege Level	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • SNMPv3 Overview on page 1460 • Configuring SNMP on page 1499 • Minimum SNMPv3 Configuration on a Device Running Junos OS on page 1461

event

Syntax	<pre>event <i>index</i> { community <i>community-name</i>; description <i>description</i>; type (RMON Notification) <i>type</i>; }</pre>
Hierarchy Level	[edit snmp rmon]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure RMON event entries.
Options	<p><i>index</i>—Identifier for a specific event entry.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• RMON MIB Event, Alarm, Log, and History Control Tables on page 1464• Monitoring RMON MIB Tables on page 1516• Understanding RMON on page 1463• Junos OS Network Management Configuration Guide

falling-event-index (RMON)

Syntax	<code>falling-event-index <i>index</i>;</code>
Hierarchy Level	<code>[edit snmp rmon alarm <i>index</i>]</code>
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Set the index number of the event entry that is used when a falling threshold is crossed. You specify the falling-event index when you configure an SNMP RMON alarm. If this value is zero, no event is triggered.
Options	<i>index</i> —Index of the event entry that is used when a falling threshold is crossed. Range: 0 through 65,535 Default: 0
Required Privilege Level	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• RMON MIB Event, Alarm, Log, and History Control Tables on page 1464• Configuring RMON Alarms and Events on page 1507• Monitoring RMON MIB Tables on page 1516• Understanding RMON on page 1463• Junos OS Network Management Configuration Guide

falling-threshold (Health Monitor)

Syntax	<code>falling-threshold <i>percentage</i>;</code>
Hierarchy Level	[edit snmp health-monitor]
Release Information	Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Set the lower threshold for the monitored object when you configure a health monitor alarm. By setting a rising and a falling threshold for a monitored variable, you can be alerted whenever the value of the variable falls outside the allowable operational range.
Options	<i>percentage</i> —Lower threshold for the alarm entry. Range: 1 through 100 Default: 70 percent of the maximum possible value
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• rising-threshold on page 1653• Configuring Health Monitoring on page 1509

falling-threshold (RMON)

Syntax	<code>falling-threshold <i>integer</i>;</code>
Hierarchy Level	<code>[edit snmp rmon alarm <i>index</i>]</code>
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Set the lower threshold for the sampled variable (monitored object) when you configure an SNMP RMON alarm. By setting a rising and a falling threshold for a variable, you can be alerted whenever the value of the variable falls outside the allowable operational range.
Options	<p><i>integer</i>—Lower threshold for the alarm entry.</p> <p>Range: -2,147,483,648 through 2,147,483,647</p> <p>Default: 20 percent less than the rising-threshold value</p>
Required Privilege Level	<p><code>snmp</code>—To view this statement in the configuration.</p> <p><code>snmp-control</code>—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • RMON MIB Event, Alarm, Log, and History Control Tables on page 1464 • Configuring RMON Alarms and Events on page 1507 • Monitoring RMON MIB Tables on page 1516 • Understanding RMON on page 1463 • Junos OS Network Management Configuration Guide

falling-threshold-interval

Syntax	<code>falling-threshold-interval seconds;</code>
Hierarchy Level	<code>[edit snmp rmon alarm index]</code>
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Set the interval between samples after the rising threshold is exceeded and the value of the sample starts to drop. If the value of the sample drops and exceeds the falling threshold, the regular sampling interval is used.
Options	interval —Time between samples, in seconds. Range: 1 through 2,147,483,647 seconds Default: 60 seconds
Required Privilege Level	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• RMON MIB Event, Alarm, Log, and History Control Tables on page 1464• Configuring RMON Alarms and Events on page 1507• Monitoring RMON MIB Tables on page 1516• Understanding RMON on page 1463• Junos OS Network Management Configuration Guide

filter-duplicates

Syntax	<code>filter-duplicates;</code>
Hierarchy Level	<code>[edit snmp]</code>
Release Information	Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Filter duplicate Get , GetNext , or GetBulk SNMP requests.
Required Privilege Level	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding the Implementation of SNMP on the QFabric System• Example: Configuring SNMP on page 1523

filter-interfaces

Syntax	filter-interfaces { all-internal-interfaces; interfaces <i>interface</i> }
Hierarchy Level	[edit snmp]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Filter out information related to specific interfaces from the output of SNMP Get and GetNext requests performed on interface-related MIBs.
Options	all-internal-interfaces —Filter out information from SNMP Get and GetNext requests for all internal interfaces. interfaces —Filter out information from SNMP Get and GetNext requests for the specified interface.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Filtering Interface Information Out of SNMP Get and GetNext Output</i>

group (Defining Access Privileges for an SNMPv3 Group)

Syntax	group <i>group-name</i> ;
Hierarchy Level	[edit snmp v3 vacm security-to-group security-model (usm v1 v2c) <i>security-name security-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Define access privileges granted to a group.
Options	group-name —Identifies a collection of SNMP security names that belong to the same access policy SNMP.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the Group</i>

group (Configuring Group Name)

Syntax `group group-name {
 (default-context-prefix | context-prefix context-prefix){
 security-model (any | usm | v1 | v2c) {
 security-level (authentication | none | privacy) {
 notify-view view-name;
 read-view view-name;
 write-view view-name;
 }
 }
 }
 }`

Hierarchy Level [edit snmp v3 vacm access]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Assign the security name to a group, and specify the SNMPv3 context applicable to the group. The **default-context-prefix** statement, when included, adds all the contexts configured on the device to the group, whereas the **context-prefix context-prefix** statement enables you to specify a context and to add that particular context to the group.

(Not applicable to the QFX Series and OCX Series.) When the context prefix is specified as default (for example, **context-prefix default**), the context associated with the master routing instance is added to the group. To specify a routing instance that is part of a logical system, specify it as **logical system/routing instance**. For example, to specify routing instance ri1 in logical system ls1, include **context-prefix ls1/ri1**.

The remaining statements under this hierarchy are explained separately.

Options *group-name*—SNMPv3 group name created for the SNMPv3 group.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation • *Configuring the Group*

health-monitor

Syntax	health-monitor { falling-threshold <i>percentage</i> ; interval <i>seconds</i> ; rising-threshold <i>percentage</i> ; }
Hierarchy Level	[edit snmp]
Release Information	Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure health monitoring. The remaining statements are explained separately.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Health Monitoring on page 1509• Understanding Health Monitoring on page 1467

history

Syntax	<pre>history <i>history-index</i> { <i>bucket-size</i> <i>number</i>; <i>interface</i> <i>interface-name</i>; <i>interval</i> <i>seconds</i>; <i>owner</i> <i>owner-name</i>; }</pre>
Hierarchy Level	[edit snmp rmon]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Configure RMON history group entries. This RMON feature can be used with the Simple Network Management Protocol (SNMP) agent on the network to monitor all the traffic flowing among devices on all connected LAN segments. The RMON history feature collects statistics in accordance with user-configurable parameters.</p> <p>The history group controls the periodic statistical sampling of data from various types of networks. This group contains configuration entries that specify an interface, polling period, and other parameters. If you use the history statement, you must also configure the interface <i>interface-name</i> statement.</p>
Options	<p>history-index—Provide a number for this history entry.</p> <p>Range: 1 through 65535</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• RMON MIB Event, Alarm, Log, and History Control Tables on page 1464• Monitoring RMON MIB Tables on page 1516• Understanding RMON on page 1463• Junos OS Network Management Configuration Guide

interface (SNMP)

Syntax	<code>interface [<i>interface-names</i>];</code>
Hierarchy Level	[edit snmp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the interfaces on which SNMP requests can be accepted.
Default	If you omit this statement, SNMP requests entering the router or switch through any interface are accepted.
Options	<i>interface-names</i> —Names of one or more logical interfaces.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Interfaces on Which SNMP Requests Can Be Accepted on page 1506

interface (RMON)

Syntax	<code>interface <i>interface-name</i>;</code>
Hierarchy Level	[edit snmp rmon history <i>history-index</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Specify the interface to be monitored in the specified RMON history entry.</p> <p>Only one interface can be specified for a particular RMON history index. There is a one-to-one relationship between the interface and the history index. The interface must be specified in order for the RMON history to be created.</p>
Options	<i>interface-name</i> —Specify the interface to be monitored within the specified entry of the RMON history of Ethernet statistics.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• RMON MIB Event, Alarm, Log, and History Control Tables on page 1464• Configuring RMON Alarms and Events on page 1507• Monitoring RMON MIB Tables on page 1516• Understanding RMON on page 1463• Junos OS Network Management Configuration Guide

interval (Health Monitor)

Syntax	interval <i>seconds</i> ;
Hierarchy Level	[edit snmp health-monitor]
Release Information	Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the interval between sampling of the object being monitored by the health monitor.
Options	seconds —Time between samples, in seconds. Range: 1 through 2147483647 seconds Default: 300 seconds
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Health Monitoring on page 1509

interval (RMON)

Syntax	<code>interval seconds;</code>
Hierarchy Level	[edit snmp rmon alarm <i>index</i>], [edit snmp rmon history <i>index</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the interval over which data is to be sampled for the specified alarm or interface.
Default	60 sec for alarm sampling. 1800 sec for history sampling.
Options	<i>seconds</i> —Interval at which data is to be sampled for the specified alarm or interface.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• RMON MIB Event, Alarm, Log, and History Control Tables on page 1464• Configuring RMON Alarms and Events on page 1507• Monitoring RMON MIB Tables on page 1516• Understanding RMON on page 1463• Junos OS Network Management Configuration Guide

local-engine

Syntax

```

local-engine {
  user username {
    authentication-md5 {
      authentication-password authentication-password;
    }
    authentication-none;
    authentication-sha {
      authentication-password authentication-password;
    }
    privacy-aes128 {
      privacy-password privacy-password;
    }
    privacy-des {
      privacy-password privacy-password;
    }
    privacy-3des {
      privacy-password privacy-password;
    }
    privacy-none {
      privacy-password privacy-password;
    }
  }
}

```

Hierarchy Level [edit snmp v3 [usm](#)]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure local engine information for the user-based security model (USM).

The remaining statements are explained separately.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- [Creating SNMPv3 Users on page 1510](#)

location

Syntax	<code>location <i>location</i>;</code>
Hierarchy Level	<code>[edit snmp]</code>
Release Information	Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Define the value of the MIB II sysLocation object, which is the physical location of the managed system.
Options	<i>location</i> —Location of the local system. You must enclose the name within quotation marks (" ").
Required Privilege Level	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the System Location for a Device Running Junos OS</i>

message-processing-model

Syntax	<code>message-processing-model (v1 v2c v3);</code>
Hierarchy Level	<code>[edit snmp v3 target-parameters <i>target-parameter-name</i> parameters]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the message processing model to be used when generating SNMP notifications.
Options	<i>v1</i> —SNMPv1 message process model. <i>v2c</i> —SNMPv2c message process model. <i>v3</i> —SNMPv3 message process model.
Required Privilege Level	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Message Processing Model</i>

name

Syntax	<code>name <i>name</i>;</code>
Hierarchy Level	[edit snmp]
Release Information	Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Set the system name from the command-line interface.
Options	<i>name</i> —System name override.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring a Different System Name</i>

nonvolatile

Syntax	<pre>nonvolatile { <i>commit-delay seconds</i>; }</pre>
Hierarchy Level	[edit snmp]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Configure options for SNMP Set requests.</p> <p>The statement is explained separately.</p>
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the Commit Delay Timer</i> • <i>commit-delay</i>

notify

Syntax	<pre>notify <i>name</i> { tag <i>tag-name</i>; type (trap inform); }</pre>
Hierarchy Level	[edit snmp v3]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>type inform option added in Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Select management targets for SNMPv3 notifications as well as the type of notifications. Notifications can be either traps or informs.
Options	<p><i>name</i>—Name assigned to the notification.</p> <p><i>tag-name</i>—Notifications are sent to all targets configured with this tag.</p> <p><i>type</i>—Notification type is trap or inform. Traps are unconfirmed notifications. Informs are confirmed notifications.</p>
Required Privilege Level	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Inform Notification Type and Target Address</i>• <i>Configuring the SNMPv3 Trap Notification</i>

notify-filter (Applying to the Management Target)

Syntax	<code>notify-filter <i>profile-name</i>;</code>
Hierarchy Level	<code>[edit snmp v3 target-parameters <i>target-parameters-name</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify the notify filter to be used by a specific set of target parameters.
Options	<i>profile-name</i> —Name of the notify filter to apply to notifications.
Required Privilege Level	snmp —To view this statement in the configuration. snmp-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Applying the Trap Notification Filter</i>

notify-filter (Configuring the Profile Name)

Syntax	<code>notify-filter <i>profile-name</i> { oid <i>oid</i> (include exclude); }</code>
Hierarchy Level	<code>[edit snmp v3]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify a group of MIB objects for which you define access. The notify filter limits the type of traps or informs sent to the network management system.
Options	<i>profile-name</i> —Name assigned to the notify filter. The remaining statement is explained separately.
Required Privilege Level	snmp —To view this statement in the configuration. snmp-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the Trap Notification Filter</i> • <i>oid (SNMP)</i>

notify-view

Syntax	<code>notify-view view-name;</code>
Hierarchy Level	<code>[edit snmp v3 vacm access group group-name (default-context-prefix context-prefix context-prefix) security-model (any usm v1 v2c) security-level (authentication none privacy)]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Associate the notify view with a community (for SNMPv1 or SNMPv2c clients) or a group name (for SNMPv3 clients).
Options	view-name —Name of the view to which the SNMP user group has access.
Required Privilege Level	snmp —To view this statement in the configuration. snmp-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring MIB Views on page 1506• Configuring the Notify View

oid

Syntax	<code>oid object-identifier (exclude include);</code>
Hierarchy Level	<code>[edit snmp view view-name]</code>
Release Information	Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify an object identifier (OID) used to represent a subtree of MIB objects.
Options	exclude —Exclude the subtree of MIB objects represented by the specified OID. include —Include the subtree of MIB objects represented by the specified OID. object-identifier —OID used to represent a subtree of MIB objects. All MIB objects represented by this statement have the specified OID as a prefix. You can specify the OID using either a sequence of dotted integers or a subtree name.
Required Privilege Level	snmp —To view this statement in the configuration. snmp-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring MIB Views on page 1506

oid (SNMPv3)

Syntax	oid <i>oid</i> (include exclude);
Hierarchy Level	[edit snmp v3 notify-filter <i>profile-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify an object identifier (OID) used to represent a subtree of MIB objects. This OID is a prefix that the represented MIB objects have in common.
Options	<p>exclude—Exclude the subtree of MIB objects represented by the specified OID.</p> <p>include—Include the subtree of MIB objects represented by the specified OID.</p> <p>oid—Object identifier used to represent a subtree of MIB objects. All MIB objects represented by this statement have the specified OID as a prefix. You can specify the OID using either a sequence of dotted integers or a subtree name.</p>
Required Privilege Level	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • SNMPv3 Overview on page 1460 • Minimum SNMPv3 Configuration on a Device Running Junos OS on page 1461 • Configuring SNMP on page 1499 • Configuring the SNMPv3 Trap Notification

owner

Syntax	<code>owner owner-name;</code>
Hierarchy Level	<code>[edit snmp rmon history index]</code>
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify the user or group responsible for this RMON history configuration.
Options	owner-name —User or group responsible for this configuration. Range: 0 through 32 alphanumeric characters
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• RMON MIB Event, Alarm, Log, and History Control Tables on page 1464• Configuring RMON Alarms and Events on page 1507• Monitoring RMON MIB Tables on page 1516• Understanding RMON on page 1463• Junos OS Network Management Configuration Guide

parameters

Syntax	<pre>parameters { message-processing-model (v1 v2c v3); security-level (none authentication privacy); security-model (usm v1 v2c); security-name security-name; }</pre>
Hierarchy Level	<code>[edit snmp v3 target-parameters target-parameters-name]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure a set of target parameters for message processing and security. The remaining statements are explained separately.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Defining and Configuring the Trap Target Parameters

port

Syntax	<code>port <i>port-number</i>;</code>
Hierarchy Level	[edit snmp v3 target-address <i>target-address-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure a UDP port number for an SNMP target.
Default	If you omit this statement, the default port is 162.
Options	<i>port-number</i> —Port number for the SNMP target.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Port</i>

privacy-3des

Syntax	<pre>privacy-3des { privacy-password <i>privacy-password</i>; }</pre>
Hierarchy Level	[edit snmp v3 usm local-engine user <i>username</i>], [edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the triple Data Encryption Standard (3DES) as the privacy type for the SNMPv3 user.
Options	<p>privacy-password <i>privacy-password</i>—Password that a user enters. The password is then converted into a key that is used for encryption.</p> <p>SNMPv3 has special requirements when you create plain-text passwords on a router or switch:</p> <ul style="list-style-type: none">• The password must be at least eight characters long.• The password can include alphabetic, numeric, and special characters, but it cannot include control characters.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the SNMPv3 Encryption Type</i>

privacy-aes128

Syntax	<code>privacy-aes128 { privacy-password <i>privacy-password</i>; }</code>
Hierarchy Level	[edit snmp v3 usm local-engine user <i>username</i>], [edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the Advanced Encryption Standard encryption algorithm (CFB128-AES-128 Privacy Protocol) for the SNMPv3 user.
Options	<p>privacy-password <i>privacy-password</i>—Password that a user enters. The password is then converted into a key that is used for encryption.</p> <p>SNMPv3 has special requirements when you create plain-text passwords on a router or switch:</p> <ul style="list-style-type: none"> • The password must be at least eight characters long. • The password can include alphabetic, numeric, and special characters, but it cannot include control characters.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the SNMPv3 Encryption Type</i>

privacy-des

Syntax	<code>privacy-des { privacy-password <i>privacy-password</i>; }</code>
Hierarchy Level	[edit snmp v3 usm local-engine user <i>username</i>], [edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the Data Encryption Standard (DES) as the privacy type for the SNMPv3 user.
Options	privacy-password <i>privacy-password</i> —Password that a user enters. The password is then converted into a key that is used for encryption. SNMPv3 has special requirements when you create plain-text passwords on a router or switch: <ul style="list-style-type: none">• The password must be at least eight characters long.• The password can include alphabetic, numeric, and special characters, but it cannot include control characters.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the SNMPv3 Encryption Type</i>

privacy-none

Syntax	<code>privacy-none;</code>
Hierarchy Level	[edit snmp v3 usm local-engine user <i>username</i>], [edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure that no encryption be used for the SNMPv3 user.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the SNMPv3 Encryption Type</i>

privacy-password

Syntax	<code>privacy-password <i>privacy-password</i>;</code>
Hierarchy Level	<code>[edit snmp v3 usm local-engine user <i>username</i> privacy-3des],</code> <code>[edit snmp v3 usm local-engine user <i>username</i> privacy-aes128],</code> <code>[edit snmp v3 usm local-engine user <i>username</i> privacy-des],</code> <code>[edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i> privacy-3des],</code> <code>[edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i> privacy-aes128],</code> <code>[edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i> privacy-des]</code>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	Configure a privacy password for the SNMPv3 user.
Options	<p><i>privacy-password</i>—Password that a user enters. The password is then converted into a key that is used for encryption.</p> <p>SNMPv3 has special requirements when you create plain-text passwords on a router or switch:</p> <ul style="list-style-type: none"> • The password must be at least eight characters long. • The password can include alphabetic, numeric, and special characters, but it cannot include control characters.
Required Privilege Level	<p><code>snmp</code>—To view this statement in the configuration.</p> <p><code>snmp-control</code>—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the SNMPv3 Encryption Type</i>

read-view

Syntax	<code>read-view <i>view-name</i>;</code>
Hierarchy Level	[edit snmp v3 vacm access group <i>group-name</i> (default-context-prefix context-prefix <i>context-prefix</i>) security-model (any usm v1 v2c) security-level (authentication none privacy)]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Associate the read-only view with a community (for SNMPv1 or SNMPv2c clients) or a group name (for SNMPv3 clients).
Options	<i>view-name</i> —The name of the view to which the SNMP user group has access.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Read View</i>• Configuring MIB Views on page 1506

remote-engine

Syntax	<pre> remote-engine <i>engine-id</i> { user <i>username</i> { authentication-md5 { authentication-password <i>authentication-password</i>; } authentication-none; authentication-sha { authentication-password <i>authentication-password</i>; } privacy-aes128 { privacy-password <i>privacy-password</i>; } privacy-des { privacy-password <i>privacy-password</i>; } privacy-3des { privacy-password <i>privacy-password</i>; } privacy-none { privacy-password <i>privacy-password</i>; } } } </pre>
Hierarchy Level	[edit snmp v3 usm]
Release Information	<p>Statement introduced in Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Configure the remote engine information for the user-based security model (USM). To send inform messages to an SNMPv3 user on a remote device, you must configure the engine identifier for the SNMP agent on the remote device where the user resides.
Options	<p><i>engine-id</i>—Specify engine identifier in hexadecimal format. Used to compute the security digest for authenticating and encrypting packets sent to a user on the remote host.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <i>Configuring the Remote Engine and Remote User</i>

request-type

Syntax	<code>request-type (get-next-request get-request walk-request);</code>
Hierarchy Level	<code>[edit snmp rmon alarm <i>index</i>]</code>
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Extend monitoring to a specific SNMP object instance (<code>get-request</code>), to all object instances belonging to a MIB branch (<code>walk-request</code>), or to the next object instance after the instance specified in the configuration (<code>get-next-request</code>).
Default	<code>walk-request</code>
Options	<code>get-next-request</code> —Perform an SNMP get next request. <code>get-request</code> —Perform an SNMP get request. <code>walk-request</code> —Perform an SNMP walk request.
Required Privilege Level	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• RMON MIB Event, Alarm, Log, and History Control Tables on page 1464• Configuring RMON Alarms and Events on page 1507• Monitoring RMON MIB Tables on page 1516• Understanding RMON on page 1463• Junos OS Network Management Configuration Guide

retry-count

Syntax	<code>retry-count <i>number</i>;</code>
Hierarchy Level	<code>[edit snmp v3 target-address <i>target-address-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 7.4.
Description	Configure the retry count for SNMP informs.
Options	<i>number</i> —Maximum number of times the inform is transmitted if no acknowledgment is received. If no acknowledgment is received after the inform is transmitted the maximum number of times, the inform message is discarded. Default: 3 times
Required Privilege Level	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring SNMP Informs on page 1515• <i>timeout (SNMP)</i>

rising-event-index

Syntax	<code>rising-event-index <i>index</i>;</code>
Hierarchy Level	[edit snmp rmon alarm index]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Set the index of the event entry that is used when a rising alarm threshold is exceeded. The rising-event index is specified when you configure an SNMP RMON alarm. If this value is zero, no event is triggered.
Options	<i>index</i> —Index of the event entry that is used when a rising threshold is exceeded. Range: 0 through 65,535 Default: 0
Required Privilege Level	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• RMON MIB Event, Alarm, Log, and History Control Tables on page 1464• Configuring RMON Alarms and Events on page 1507• Monitoring RMON MIB Tables on page 1516• Understanding RMON on page 1463• Junos OS Network Management Configuration Guide

rising-threshold (Health Monitor)

Syntax	<code>rising-threshold <i>percentage</i>;</code>
Hierarchy Level	[edit snmp health-monitor]
Release Information	Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Set the upper threshold for the monitored object when you configure a health monitor alarm. By setting a rising and a falling threshold for a monitored object, you can be alerted whenever the value of the variable falls outside the allowable operational range.
Options	<i>percentage</i> —Upper threshold for the alarm entry. Range: 1 through 100 Default: 80 percent of the maximum possible value
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Health Monitoring on page 1509• falling-threshold on page 1624

rising-threshold (RMON)

Syntax	<code>rising-threshold <i>integer</i>;</code>
Hierarchy Level	<code>[edit snmp rmon alarm <i>index</i>]</code>
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Set the upper threshold for the sampled variable (monitored object) when you configure an SNMP RMON alarm. By setting a rising and a falling threshold for a variable, you can be alerted whenever the value of the variable falls outside the allowable operational range.
Options	<i>integer</i> —Upper threshold for the alarm entry. Range: -2,147,483,648 through 2,147,483,647
Required Privilege Level	snmp —To view this statement in the configuration. snmp-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• RMON MIB Event, Alarm, Log, and History Control Tables on page 1464• Configuring RMON Alarms and Events on page 1507• Monitoring RMON MIB Tables on page 1516• Understanding RMON on page 1463• Junos OS Network Management Configuration Guide

rmon

```
Syntax  rmon {
        alarm index {
            description description;
            falling-event-index index;
            falling-threshold integer;
            falling-threshold-interval seconds;
            interval seconds;
            request-type;
            rising-event-index index;
            rising-threshold integer;
            sample-type (absolute-value | delta-value);
            startup-alarm (falling-alarm | rising-alarm | rising-or-falling alarm);
            syslog-subtag syslog-subtag;
            variable oid-variable;
        }
        event index {
            community community-name;
            description description;
            type (RMON Notification) type;
        }
        history history-index {
            bucket-size number;
            interface interface-name;
            interval seconds;
            owner owner-name;
        }
    }
```

Hierarchy Level [edit snmp]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Provide comprehensive network fault diagnosis, planning, and performance tuning information. RMON delivers this information in nine groups of monitoring elements, each providing specific sets of data to meet common network monitoring requirements. Each group is optional, so that vendors do not need to support all the groups within the MIB.

Junos OS supports the RMON statistics, history, alarm, and event groups.

The remaining statements are explained separately.

Default Disabled.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- [RMON MIB Event, Alarm, Log, and History Control Tables on page 1464](#)
- [Monitoring RMON MIB Tables on page 1516](#)

- [Understanding RMON on page 1463](#)
- [Junos OS Network Management Configuration Guide](#)

sample-type

Syntax	sample-type (absolute-value delta-value);
Hierarchy Level	[edit snmp rmon alarm <i>index</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the method of sampling the selected variable (monitored object). When you configure an SNMP RMON alarm, you can specify the sample type.
Options	absolute-value —Actual value of the selected variable is used when comparing against the thresholds. delta-value —Difference between samples of the selected variable is used when comparing against the thresholds.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• RMON MIB Event, Alarm, Log, and History Control Tables on page 1464• Configuring RMON Alarms and Events on page 1507• Monitoring RMON MIB Tables on page 1516• Understanding RMON on page 1463• Junos OS Network Management Configuration Guide

security-level (Defining Access Privileges)

Syntax	<pre>security-level (authentication none privacy) { notify-view view-name; read-view view-name; write-view view-name; }</pre>
Hierarchy Level	[edit snmp v3 vacm access group <i>group-name</i> (default-context-prefix context-prefix <i>context-prefix</i>) security-model (any usm v1 v2c)]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	Define the security level used for access privileges.
Default	none
Options	<p>authentication—Provide authentication but no encryption.</p> <p>none—No authentication and no encryption.</p> <p>privacy—Provide authentication and encryption.</p>
Required Privilege Level	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the Security Level</i>

security-level (Generating SNMP Notifications)

Syntax	security-level (authentication none privacy);
Hierarchy Level	[edit snmp v3 target-parameters <i>target-parameters-name</i> parameters]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the security level to use when generating SNMP notifications.
Default	none
Options	authentication —Provide authentication but no encryption. none —No authentication and no encryption. privacy —Provide authentication and encryption.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Security Level</i>

security-model (Access Privileges)

Syntax	<code>security-model (usm v1 v2c);</code>
Hierarchy Level	<code>[edit snmp v3 vacm access group <i>group-name</i> (default-context-prefix context-prefix <i>context-prefix</i>)]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the security model for an SNMPv3 group. The security model is used to determine access privileges for the group.
Options	<code>usm</code> —SNMPv3 security model. <code>v1</code> —SNMPv1 security model. <code>v2c</code> —SNMPv2c security model.
Required Privilege Level	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Security Model</i>

security-model (Group)

Syntax	<pre>security-model (usm v1 v2c) { security-name security-name { group group-name; } }</pre>
Hierarchy Level	[edit snmp v3 vacm security-to-group]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Define a security model for a group.
Options	usm —SNMPv3 security model. v1 —SNMPv1 security model. v2c —SNMPv2c security model.
Required Privilege Level	snmp —To view this statement in the configuration. snmp-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Security Model</i>

security-model (SNMP Notifications)

Syntax	<pre>security-model (usm v1 v2c);</pre>
Hierarchy Level	[edit snmp v3 target-parameters target-parameters-name parameters]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the security model for an SNMPv3 group. The security model is used for SNMP notifications.
Options	usm —SNMPv3 security model. v1 —SNMPv1 security model. v2c —SNMPv2c security model.
Required Privilege Level	snmp —To view this statement in the configuration. snmp-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Security Model</i>

security-name (Community String)

Syntax	<code>security-name <i>security-name</i>;</code>
Hierarchy Level	<code>[edit snmp v3 <i>snmp-community</i> <i>community-index</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Associate a community string with the security name of a user. The community string, which is used for SNMPv1 and SNMPv2c clients in an SNMPv3 system, is configured at the <code>[edit snmp v3 snmp-community <i>community-index</i>]</code> hierarchy level.
Options	<i>security-name</i> —Name that is used for messaging security and user access control.




NOTE: The security name must match the configured security name at the `[edit snmp v3 target-parameters target-parameters-name parameters]` hierarchy level when you configure traps or informs.

Required Privilege	snmp—To view this statement in the configuration.
Level	snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the Security Names</i>

security-name (Security Group)

Syntax	<code>security-name <i>security-name</i> { group <i>group-name</i>; }</code>
Hierarchy Level	[edit snmp v3 vacm security-to-group security-model (usm v1 v2c)]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Associate a group or a community string with a configured security group.
Options	security-name —Username configured at the [edit snmp v3 usm local-engine user <i>username</i>] hierarchy level. For SNMPv1 and SNMPv2c, the security name is the community string configured at the [edit snmp v3 snmp-community <i>community-index</i>] hierarchy level.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Assigning Security Names to Groups</i>

security-name (SNMP Notifications)

Syntax	<code>security-name <i>security-name</i>;</code>
Hierarchy Level	<code>[edit snmp v3 target-parameters <i>target-parameters-name</i> parameters]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the security name used when generating SNMP notifications.
Options	<i>security-name</i> —If the SNMPv3 USM security model is used, identify the user when generating the SNMP notification. If the v1 or v2c security models are used, identify the SNMP community used when generating the notification.
<div>  <p>NOTE: The access privileges for the group associated with this security name must allow this notification to be sent.</p> <p>If you are using the v1 or v2 security models, the security name at the <code>[edit snmp v3 vacm security-to-group]</code> hierarchy level must match the security name at the <code>[edit snmp v3 snmp-community <i>community-index</i>]</code> hierarchy level.</p> </div>	
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> <i>Configuring the Security Name</i>

security-to-group

Syntax	<pre>security-to-group { security-model (usm v1 v2c) { group group-name; security-name security-name; } }</pre>
Hierarchy Level	[edit snmp v3 vacm]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Configure the group to which a specific SNMPv3 security name belongs. The security name is used for messaging security.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Assigning Security Model and Security Name to a Group</i>

snmp

```

Syntax  snmp {
    client-list client-list-name {
        ip-addresses;
    }
    community community-name {
        authorization authorization;
        client-list-name client-list-name;
        clients {
            address restrict;
        }
        logical-system logical-system-name {
            routing-instance routing-instance-name {
                clients {
                    addresses;
                }
            }
        }
        routing-instance routing-instance-name {
            clients {
                addresses;
            }
        }
        view view-name;
    }
    contact contact;
    description description;
    filter-duplicates;
    filter-interfaces;
    health-monitor {
        falling-threshold integer;
        interval seconds;
        rising-threshold integer;
    }
    interface [ interface-names ];
    location location;
    name name;
    nonvolatile {
        commit-delay seconds;
    }
    rmon {
        alarm index {
            description description;
            falling-event-index index;
            falling-threshold integer;
            falling-threshold-interval seconds;
            interval seconds;
            request-type;
            rising-event-index index;
            rising-threshold integer;
            sample-type (absolute-value | delta-value);
            startup-alarm (falling-alarm | rising-alarm | rising-or-falling alarm);
            syslog-subtag syslog-subtag;
        }
    }
}

```

```
    variable oid-variable;
}
event index {
    community community-name;
    description description;
    type type;
}
history history-index {
    bucket-size number;
    interface interface-name;
    interval seconds;
    owner owner-name;
}
}
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable> <match
        regular-expression>;
    flag flag;
}
trap-group group-name {
    categories {
        category;
    }
    destination-port port-number;
    routing-instance routing-instance-name;
    targets {
        address;
    }
    version (all | v1 | v2);
}
trap-options {
    agent-address outgoing-interface;
    source-address address;
}
}
v3 {
    notify name {
        tag tag-name;
        type trap;
    }
    notify-filter profile-name {
        oid object-identifier (include | exclude);
    }
    snmp-community community-index {
        community-name community-name;
        security-name security-name;
        tag tag-name;
    }
    target-address target-address-name {
        address address;
        address-mask address-mask;
        logical-system logical-system;
        port port-number;
        retry-count number;
        routing-instance routing-instance-name;
        tag-list tag-list;
        target-parameters target-parameters-name;
```

```

    timeout seconds;
}
target-parameters target-parameters-name {
    notify-filter profile-name;
    parameters {
        message-processing-model (v1 | v2c | V3);
        security-level (authentication | none | privacy);
        security-model (usm | v1 | v2c);
        security-name security-name;
    }
}
usm {
    local-engine {
        user username {
            authentication-sha {
                authentication-password authentication-password;
            }
            authentication-md5 {
                authentication-password authentication-password;
            }
            authentication-none;
            privacy-aes128 {
                privacy-password privacy-password;
            }
            privacy-des {
                privacy-password privacy-password;
            }
            privacy-3des {
                privacy-password privacy-password;
            }
            privacy-none;
        }
    }
    remote-engine engine-id {
        user username {
            authentication-sha {
                authentication-password authentication-password;
            }
            authentication-md5 {
                authentication-password authentication-password;
            }
            authentication-none;
            privacy-aes128 {
                privacy-password privacy-password;
            }
            privacy-des {
                privacy-password privacy-password;
            }
            privacy-3des {
                privacy-password privacy-password;
            }
            privacy-none {
                privacy-password privacy-password;
            }
        }
    }
}

```

```

}
vacm {
  access {
    group group-name {
      (default-context-prefix | context-prefix context-prefix) {
        security-model (any | usm | v1 | v2c) {
          security-level (authentication | none | privacy) {
            notify-view view-name;
            read-view view-name;
            write-view view-name;
          }
        }
      }
    }
  }
}
security-to-group {
  security-model (usm | v1 | v2c) {
    security-name security-name {
      group group-name;
    }
  }
}
}
view view-name {
  oid object-identifier (include | exclude);
}
}
}

```

Hierarchy Level [edit]

Release Information Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure SNMP.

The remaining statements are explained separately.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- *Understanding the Implementation of SNMP*
- [Configuring SNMP on page 1499](#)

snmp-community

Syntax	snmp-community <i>community-index</i> { <i>community-name</i> <i>community-name</i> ; <i>security-name</i> <i>security-name</i> ; tag <i>tag-name</i> ; }
Hierarchy Level	[edit snmp v3]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure the SNMP community.
Options	<i>community-index</i> —(Optional) String that identifies an SNMP community. The remaining statements are explained separately.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the SNMPv3 Community</i>

source-address (SNMP)

Syntax	source-address <i>address</i> ;
Hierarchy Level	[edit snmp trap-options]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Set the source address of every SNMP trap packet sent by this switch to a single address regardless of the outgoing interface. If the source address is not specified, the default is to use the address of the outgoing interface as the source address.
Options	<p><i>address</i>—Source address of SNMP traps. You can configure the source address of trap packets two ways: lo0 or a valid IPv4 address configured on one of the interfaces. The value lo0 indicates that the source address of all SNMP trap packets is set to the lowest loopback address configured at interface lo0.</p> <p>Default: Disabled. (The source address is the address of the outgoing interface.)</p>
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the Source Address for SNMP Traps</i>

startup-alarm

Syntax	startup-alarm (falling-alarm rising-alarm rising-or-falling-alarm);
Hierarchy Level	[edit snmp rmon alarm <i>index</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Set an initial alarm that is sent after the configured SNMP RMON alarm becomes active.
Default	rising-or-falling-alarm
Options	<p>falling-alarm—Generated if the first sample after the alarm becomes active is equal to or greater than the falling threshold.</p> <p>rising-alarm—Generated if the first sample after the alarm becomes active is equal to or greater than the rising threshold.</p> <p>rising-or-falling-alarm—Generated if the first sample after the alarm entry becomes active is equal to or greater than either the rising threshold or the falling threshold.</p>
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• RMON MIB Event, Alarm, Log, and History Control Tables on page 1464• Configuring RMON Alarms and Events on page 1507• Monitoring RMON MIB Tables on page 1516• Understanding RMON on page 1463• Junos OS Network Management Configuration Guide

syslog-subtag

Syntax	<code>syslog-subtag <i>syslog-subtag</i>;</code>
Hierarchy Level	[edit snmp rmon alarm <i>index</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Add the syslog-subtag tag to the system log message. The tag should not exceed 80 uppercase characters.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • RMON MIB Event, Alarm, Log, and History Control Tables on page 1464 • Monitoring RMON MIB Tables on page 1516 • Understanding RMON on page 1463 • Junos OS Network Management Configuration Guide

tag (Configuring Notification Targets)

Syntax	<code>tag <i>tag-name</i>;</code>
Hierarchy Level	[edit snmp v3 notify <i>name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure a set of target addresses to receive SNMP traps or informs (for IPv4 packets only).
Options	tag-name —Define the target addresses to which an SNMP notification is sent. Target addresses containing the same tag in their tag list are sent the same notification. The tag-name is not included in the notification.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • SNMPv3 Overview on page 1460 • Minimum SNMPv3 Configuration on a Device Running Junos OS on page 1461 • Configuring SNMP on page 1499 • Configuring the SNMPv3 Trap Notification

tag (Configuring the SNMP Community)

Syntax	<code>tag tag-name;</code>
Hierarchy Level	[edit snmp v3 snmp-community <i>community-index</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure a set of SNMP managers that are authorized to use a community string.
Options	tag-name —Identify the set of addresses for the SNMP managers authorized to use the community string.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• SNMPv3 Overview on page 1460• Minimum SNMPv3 Configuration on a Device Running Junos OS on page 1461• Configuring SNMP on page 1499• Configuring the SNMPv3 Trap Notification

tag-list

Syntax	<code>tag-list tag-list;</code>
Hierarchy Level	[edit snmp v3 target-address <i>target-address-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure an SNMP tag list used to select target addresses.
Options	tag-list —Define sets of target addresses (tags). To specify more than one tag, specify the tag names as a space-separated list enclosed within double quotes.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Trap Target Address

target-address

Syntax	<pre>target-address <i>target-address-name</i> { address <i>address</i>; address-mask <i>address-mask</i>; port <i>port-number</i>; retry-count <i>number</i>; tag-list <i>tag-list</i>; target-parameters <i>target-parameters-name</i>; timeout <i>seconds</i>; }</pre>
Hierarchy Level	[edit snmp v3]
Release Information	<p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Configure the address of an SNMP management application and the parameters to be used in sending notifications.
Options	<p><i>target-address-name</i>—String that identifies the target address.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Understanding the Implementation of SNMP • SNMP MIBs Support on page 1468 • SNMP Traps Support on page 1484 • snmp on page 1665 • Configuring SNMP on page 1499 • Monitoring SNMP on page 1516 • Example: Configuring SNMP on page 1523

target-parameters

Syntax At the `[edit snmp v3]` hierarchy level:

```
target-parameters target-parameters-name {  
  profile-name;  
  parameters {  
    message-processing-model (v1 | v2c | V3);  
    security-level (authentication | none | privacy);  
    security-model (usm | v1 | v2c);  
    security-name security-name;  
  }  
}
```

At the `[edit snmp v3 target-address target-address-name]` hierarchy level:

```
target-parameters target-parameters-name;
```

Hierarchy Level `[edit snmp v3]`
`[edit snmp v3 target-address target-address-name]`

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure the message processing and security parameters for sending notifications to a particular management target. The target parameters are configured at the `[edit snmp v3]` hierarchy level. The remaining statements at this level are explained separately.

Then apply the target parameters configured at the `[edit snmp v3 target-parameters target-parameters-name]` hierarchy level to the target address configuration at the `[edit snmp v3]` hierarchy level.

Required Privilege snmp—To view this statement in the configuration.
Level snmp-control—To add this statement to the configuration.

Related Documentation

- *Defining and Configuring the Trap Target Parameters*
- *Applying Target Parameters*

targets

Syntax	<code>targets { <i>address</i>; }</code>
Hierarchy Level	<code>[edit snmp trap-group <i>group-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure one or more systems to receive SNMP traps.
Options	<i>address</i> —IPv4 or IPv6 address of the system to receive traps. You must specify an address, not a hostname.
Required Privilege Level	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring SNMP Trap Groups on page 1504

timeout

Syntax	<code>timeout <i>seconds</i>;</code>
Hierarchy Level	<code>[edit snmp v3 target-address <i>target-address-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the timeout period (in seconds) for SNMP informs.
Default	15 seconds
Options	<i>seconds</i> —Number of seconds to wait for an inform acknowledgment. If no acknowledgment is received within the timeout period, the inform is retransmitted.
Required Privilege Level	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding the Implementation of SNMP • Configuring SNMP Informs on page 1515 • retry-count (SNMPv3) on page 1651

traceoptions (SNMP)

Syntax	<pre>traceoptions { file <i>filename</i> <files <i>number</i>> <match <i>regular-expression</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i>; no-remote-trace; }</pre>
Hierarchy Level	[edit snmp]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Track the activities of SNMP agents on the switch and record the information in log files.



NOTE: The **traceoptions** statement is not supported on the QFabric system.

The output of the tracing operations is placed into log files in the **/var/log** directory. Each log file is named after the SNMP agent that generates it. The following logs are created in the **/var/log** directory when the **traceoptions** statement is used:

- chassisd
- craftd
- ilmid
- mib2d
- rmopd
- serviced
- snmpd

Options **file *filename***—By default, the name of the log file that records trace output is the name of the process being traced (for example, mib2d or snmpd). Use this option to specify another name.

files *number*—(Optional) Maximum number of trace files per SNMP subagent. When a trace file (for example, snmpd) reaches its maximum size, it is archived by being renamed to snmpd.0. The previous snmpd.1 is renamed to snmpd.2, and so on. The oldest archived file is deleted.

Range: 2 through 1000 files

Default: 10 files

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements:

- **all**—Log all SNMP events.

- **configuration**—Log reading of configuration at the **[edit snmp]** hierarchy level.
- **database**—Log events involving storage and retrieval in the events database.
- **events**—Log important events.
- **general**—Log general events.
- **interface-stats**—Log physical and logical interface statistics.
- **nonvolatile-sets**—Log nonvolatile SNMP set request handling.
- **pdu**—Log SNMP request and response packets.
- **policy**—Log policy processing.
- **protocol-timeouts**—Log SNMP response timeouts.
- **routing-socket**—Log routing socket calls.
- **server**—Log communication with processes that are generating events.
- **subagent**—Log subagent restarts.
- **timer-events**—Log internally generated events.
- **varbind-error**—Log variable binding errors.

match *regular-expression*—(Optional) Refine the output to include lines that contain the regular expression.

size *size*—(Optional) Maximum size, in kilobytes (KB), of each trace file before it is closed and archived.

Range: 10 KB through 1 GB

Default: 1000 KB

world-readable | no-world-readable—(Optional) By default, log files can be accessed only by the user who configures the tracing operation. The **world-readable** option enables any user to read the file. To explicitly set the default behavior, use the **no-world-readable** option.

Required Privilege Level **snmp**—To view this statement in the configuration.
 snmp-control—To add this statement to the configuration.

Related Documentation

- [Understanding Tracing and Logging Operations on page 1409](#)
- [Tracing SNMP Activity on a Device Running Junos OS on page 1518](#)

trap-group

Syntax	<pre>trap-group group-name { categories { category; } destination-port port-number; targets { address; } }</pre>
Hierarchy Level	[edit snmp]
Release Information	Statement introduced in Junos OS Release 11.1 for QFX Series switches. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Create a named group of hosts to receive the specified trap notifications. The name of the trap group is embedded in SNMP trap notification packets as one variable binding (varbind) known as the community name. At least one trap group must be configured for SNMP traps to be sent.
Options	<p>group-name—Name of the trap group. If the name includes spaces, enclose it in quotation marks (" ").</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring SNMP Trap Groups on page 1504

trap-options

Syntax	<pre>trap-options { agent-address outgoing-interface; source-address address; }</pre>
Hierarchy Level	[edit snmp]
Release Information	<p>Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	<p>Using SNMP trap options, you can set the source address of every SNMP trap packet sent by the router or switch to a single address, regardless of the outgoing interface. In addition, you can set the agent address of each SNMPv1 trap. For more information about the contents of SNMPv1 traps, see RFC 1157.</p> <p>The remaining statements are explained separately.</p>
Default	Disabled
Required Privilege Level	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring SNMP Trap Options</i>

type (RMON Notification)

Syntax	<code>type type;</code>
Hierarchy Level	[edit snmp rmon event <i>index</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the type of notification generated when a rising or falling threshold is crossed.
Default	<code>log-and-trap</code>
Options	type —Type of notification. It can be one of the following: <ul style="list-style-type: none">• log—Add an entry to the logTable object.• log-and-trap—Send an SNMP trap and add a log entry.• none—No notifications are sent.• snmptrap—Send an SNMP trap.
Required Privilege Level	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• RMON MIB Event, Alarm, Log, and History Control Tables on page 1464• Configuring RMON Alarms and Events on page 1507• Monitoring RMON MIB Tables on page 1516• Understanding RMON on page 1463• Junos OS Network Management Configuration Guide

type

Syntax	<code>type (inform trap);</code>
Hierarchy Level	<code>[edit snmp v3 notify <i>name</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. inform option added in Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the type of SNMP notification.
Options	inform —Defines the type of notification as an inform. SNMP informs are confirmed notifications. trap —Defines the type of notification as a trap. SNMP traps are unconfirmed notifications.
Required Privilege Level	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring SNMP Informs on page 1515 • Configuring the SNMPv3 Trap Notification

user

Syntax	<code>user <i>username</i>;</code>
Hierarchy Level	<code>[edit snmp v3 usm local-engine],</code> <code>[edit snmp v3 usm remote-engine <i>engine-id</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify a user associated with an SNMPv3 group on a local or remote SNMP engine.
Options	<i>username</i> —SNMPv3 user-based security model (USM) username.
Required Privilege Level	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Creating SNMPv3 Users on page 1510

usm

```

Syntax  usm {
        local-engine {
            user username {
                authentication-md5 {
                    authentication-password authentication-password;
                }
                authentication-none;
                authentication-sha {
                    authentication-password authentication-password;
                }
                privacy-aes128 {
                    privacy-password privacy-password;
                }
                privacy-des {
                    privacy-password privacy-password;
                }
                privacy-3des {
                    privacy-password privacy-password;
                }
                privacy-none {
                    privacy-password privacy-password;
                }
            }
        }
        remote-engine engine-id {
            user username {
                authentication-md5 {
                    authentication-password authentication-password;
                }
                authentication-none;
                authentication-sha {
                    authentication-password authentication-password;
                }
                privacy-aes128 {
                    privacy-password privacy-password;
                }
                privacy-des {
                    privacy-password privacy-password;
                }
                privacy-3des {
                    privacy-password privacy-password;
                }
                privacy-none {
                    privacy-password privacy-password;
                }
            }
        }
    }
}

```

Hierarchy Level [edit snmp v3]

Release Information Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure user-based security model (USM) information.

The remaining statements are explained separately.

Required Privilege snmp—To view this statement in the configuration.
Level snmp-control—To add this statement to the configuration.

Related • [Creating SNMPv3 Users on page 1510](#)
Documentation • *Configuring the Remote Engine and Remote User*

v3

```

Syntax  v3 {
    notify name {
        tag tag-name;
        type trap;
    }
    notify-filter profile-name {
        oid object-identifier (include | exclude);
    }
    snmp-community community-index {
        community-name community-name;
        security-name security-name;
        tag tag-name;
    }
    target-address target-address-name {
        address address;
        address-mask address-mask;
        port port-number;
        retry-count number;
        tag-list tag-list;
        target-parameters target-parameters-name;
        timeout seconds;
    }
    target-parameters target-parameters-name {
        notify-filter profile-name;
        parameters {
            message-processing-model (v1 | v2c | v3);
            security-level (authentication | none | privacy);
            security-model (usm | v1 | v2c);
            security-name security-name;
        }
    }
    usm {
        local-engine {
            user username {
                authentication-md5 {
                    authentication-password authentication-password;
                }
                authentication-sha {
                    authentication-password authentication-password;
                }
                authentication-none;
            }
            privacy-aes128 {
                privacy-password privacy-password;
            }
            privacy-des {
                privacy-password privacy-password;
            }
            privacy-3des {
                privacy-password privacy-password;
            }
            privacy-none;
        }
    }
}

```

```

}
remote-engine engine-id {
  user username {
    authentication-md5 {
      authentication-password authentication-password;
    }
    authentication-sha {
      authentication-password authentication-password;
    }
    authentication-none;
    privacy-aes128 {
      privacy-password privacy-password;
    }
    privacy-des {
      privacy-password privacy-password;
    }
    privacy-3des {
      privacy-password privacy-password;
    }
    privacy-none {
      privacy-password privacy-password;
    }
  }
}
}
vacm {
  access {
    group group-name {
      (default-context-prefix | context-prefix context-prefix) {
        security-model (any | usm | v1 | v2c) {
          security-level (authentication | none | privacy) {
            notify-view view-name;
            read-view view-name;
            write-view view-name;
          }
        }
      }
    }
  }
}
security-to-group {
  security-model (usm | v1 | v2c) {
    security-name security-name {
      group group-name;
    }
  }
}
}
}

```

Hierarchy Level [edit snmp]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description	Configure SNMPv3. The remaining statements are explained separately.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Minimum SNMPv3 Configuration on a Device Running Junos OS on page 1461

vacm

Syntax	<pre>vacm { access { group group-name { (default-context-prefix context-prefix context-prefix){ security-model (any usm v1 v2c) { security-level (authentication none privacy) { notify-view view-name; read-view view-name; write-view view-name; } } } } } security-to-group { security-model (usm v1 v2c); security-name security-name { group group-name; } } }</pre>
Hierarchy Level	[edit snmp v3]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure view-based access control model (VACM) information. The remaining statements are explained separately.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Defining Access Privileges for an SNMP Group</i>


variable

Syntax	<code>variable <i>oid-variable</i>;</code>
Hierarchy Level	[edit snmp rmon alarm <i>index</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Set the object identifier (OID) of the MIB object (also called variable) to be monitored when you configure an SNMP RMON alarm. If the value of the monitored variable exceeds the configured rising threshold or falling threshold, an alarm is triggered and a corresponding event may be generated.
Options	<i>oid-variable</i> —OID of the MIB variable that is being monitored. The OID can be a dotted decimal (for example, 1.3.6.1.2.1.2.1.2.1.10.1) or the name of the MIB object—for example, ifInOctets.1 .
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • RMON MIB Event, Alarm, Log, and History Control Tables on page 1464 • Configuring RMON Alarms and Events on page 1507 • Monitoring RMON MIB Tables on page 1516 • Understanding RMON on page 1463 • Junos OS Network Management Configuration Guide

version

Syntax	version (all v1 v2);
Hierarchy Level	[edit snmp trap-group <i>group-name</i>]
Release Information	Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify the version number of SNMP traps.
Default	all—Send an SNMPv1 and SNMPv2 trap for every trap condition.
Options	all—Send an SNMPv1 and SNMPv2 trap for every trap condition. v1—Send SNMPv1 traps only. v2—Send SNMPv2 traps only.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring SNMP Trap Groups on page 1504

view (Configuring a MIB View)

Syntax	<code>view <i>view-name</i> { oid <i>object-identifier</i> (include exclude); }</code>
Hierarchy Level	[edit snmp]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Define a MIB view. A MIB view identifies a group of MIB objects. Each MIB object in a view has a common OID prefix. Each object identifier represents a subtree of the MIB object hierarchy. The view statement uses a view to specify a group of MIB objects on which to define access. To enable a view, you must associate the view with a community by including the view statement at the [edit snmp community <i>community-name</i>] hierarchy level.
<div>  NOTE: To remove an OID completely, use the <code>delete view all oid oid-number</code> command but omit the <code>include</code> parameter. </div>	
Options	<p><i>view-name</i>—Name of the view.</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring MIB Views on page 1506 • <i>Associating MIB Views with an SNMP User Group</i> • community on page 1616

view (Associating MIB View with a Community)

Syntax	<code>view view-name;</code>
Hierarchy Level	<code>[edit snmp community community-name]</code>
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Associate a view with a community. A view represents a group of MIB objects.
Options	view-name —Name of the view. You must use a view name already configured in the view statement at the <code>[edit snmp]</code> hierarchy level.
Required Privilege Level	snmp —To view this statement in the configuration. snmp-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring SNMP Communities</i>

write-view

Syntax	<code>write-view view-name;</code>
Hierarchy Level	<code>[edit snmp v3 vacm access group group-name (default-context-prefix context-prefix context-prefix) security-model (any usm v1 v2c) security-level (authentication none privacy)]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series switches. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Associate the write view with a community (for SNMPv1 or SNMPv2c clients) or a group name (for SNMPv3 clients).
Options	view-name —Name of the view for which the SNMP user group has write permission.
Required Privilege Level	snmp —To view this statement in the configuration. snmp-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring MIB Views on page 1506• <i>Configuring the Write View</i>

CHAPTER 54

System Logging

- [archive \(All System Log Files\) on page 1692](#)
- [archive \(Individual System Log File\) on page 1694](#)
- [console \(System Logging\) on page 1695](#)
- [explicit-priority on page 1696](#)
- [facility-override on page 1696](#)
- [file \(System Logging\) on page 1697](#)
- [files on page 1698](#)
- [host on page 1699](#)
- [log-prefix \(System\) on page 1701](#)
- [match on page 1701](#)
- [size \(System\) on page 1702](#)
- [structured-data on page 1703](#)
- [syslog \(System\) on page 1704](#)
- [time-format on page 1706](#)
- [user \(System Logging\) on page 1707](#)

archive (All System Log Files)

Syntax	<code>archive <files <i>number</i>> <size <i>size</i>> <start-time <i>time</i>> <transfer-interval <i>interval</i>> <binary-data no-binary-data>; <world-readable no-world-readable> ;</code>
Hierarchy Level	[edit system syslog]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure archiving properties for all system log files.
Options	<p>files <i>number</i>—Maximum number of archived log files to retain. When the Junos OS logging utility has written a defined maximum amount of data to a log file logfile, it closes the file, compresses it, and renames it logfile.0.gz (the amount of data is determined by the size statement at this hierarchy level). The utility then opens and writes to a new file called logfile. When the new file reaches the maximum size, the logfile.0.gz file is renamed to logfile.1.gz, and the new file is closed, compressed, and renamed logfile.0.gz. By default, the logging facility creates up to ten archive files in this manner. Once the maximum number of archive files exists, each time the active log file reaches the maximum size, the contents of the oldest archive file are lost (overwritten by the next oldest file).</p> <p>Range: 1 through 1000</p> <p>Default: 10 files</p> <p>size <i>size</i>—Maximum amount of data that the Junos OS logging utility writes to a log file logfile before archiving it (closing it, compressing it, and changing its name to logfile.0.gz). The utility then opens and writes to a new file called logfile.</p> <p>Syntax: <i>x k</i> to specify the number of kilobytes, <i>x m</i> for the number of megabytes, or <i>x g</i> for the number of gigabytes</p> <p>Range: 64 KB through 1 GB</p> <p>Default:</p> <ul style="list-style-type: none">• 128 KB for EX Series switches• 1 MB for M Series, MX Series, and T Series routers, OCX Series, and the QFX3500 switch• 10 MB for TX Matrix and TX Matrix Plus routers <p>binary-data no-binary-data—Mark file as containing binary data. This allows proper archiving of binary files, such as WTMP files (login records for UNIX based systems)..</p> <p>Default: no-binary-data</p> <p>world-readable no-world-readable—Grant all users permission to read archived log files, or restrict the permission only to the root user and users who have the Junos OS maintenance permission.</p>

Default: no-world-readable

Required Privilege system—To view this statement in the configuration.
Level system-control—To add this statement to the configuration.

Related Documentation • [Specifying Log File Size, Number, and Archiving Properties on page 1543](#)

archive (Individual System Log File)

Syntax	archive <archive-sites (<i>ftp-url</i> <password <i>password</i> >)> <files <i>number</i> > <size <i>size</i> > <start-time "YYYY-MM-DD.hh:mm"> <transfer-interval <i>minutes</i> > <world-readable no-world-readable>;
Hierarchy Level	[edit system syslog file <i>filename</i>]
Release Information	Statement introduced before Junos OS Release 7.4. start-time and transfer-interval statements introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure archiving properties for a specific system log file.
Options	<p>archive-sites <i>site-name</i>—FTP URL representing the destination for the archived log file (for information about how to specify valid FTP URLs, see “Format for Specifying Filenames and URLs in Junos OS CLI Commands” on page 40). If more than one site name is configured, a list of archive sites for the system log files is created. When a file is archived, the router attempts to transfer the file to the first URL in the list, moving to the next site only if the transfer does not succeed. The log file is stored at the archive site with the filename specified at the [edit system syslog] hierarchy level.</p> <p>files <i>number</i>—Maximum number of archived log files to retain. When the Junos OS logging utility has written a defined maximum amount of data to a log file <i>logfile</i>, it closes the file, compresses it, and renames it <i>logfile.0.gz</i> (the amount of data is determined by the size statement at this hierarchy level). The utility then opens and writes to a new file called <i>logfile</i>. When the new file reaches the maximum size, the <i>logfile.0.gz</i> file is renamed to <i>logfile.1.gz</i>, and the new file is closed, compressed, and renamed <i>logfile.0.gz</i>. By default, the logging facility creates up to ten archive files in this manner. Once the maximum number of archive files exists, each time the active log file reaches the maximum size, the contents of the oldest archive file are lost (overwritten by the next oldest file).</p> <p>Range: 1 through 1000</p> <p>Default: 10 files</p> <p>password <i>password</i>—Password for authenticating with the site specified by the archive-sites statement.</p> <p>size <i>size</i>—Maximum amount of data that the Junos OS logging utility writes to a log file <i>logfile</i> before archiving it (closing it, compressing it, and changing its name to <i>logfile.0.gz</i>). The utility then opens and writes to a new file called <i>logfile</i>.</p> <p>Syntax: xk to specify the number of kilobytes, xm for the number of megabytes, or xg for the number of gigabytes</p> <p>Range: 64 KB through 1 GB</p> <p>Default: 128 KB for J Series routers; 1 MB for M Series, MX Series, and T Series routers, and the QFX3500 switch; 10 MB for TX Matrix and TX Matrix Plus routers</p>

start-time "YYYY-MM-DD.hh:mm"—Date and time in the local time zone for a one-time transfer of the active log file to the first reachable site in the list of sites specified by the **archive-sites** statement.

transfer-interval *interval*—Interval at which to transfer the log file to an archive site.

Range: 5 through 2880 minutes

world-readable | no-world-readable—Grant all users permission to read archived log files, or restrict the permission only to the **root** user and users who have the Junos OS **maintenance** permission.

Default: no-world-readable

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- [Specifying Log File Size, Number, and Archiving Properties on page 1543](#)

console (System Logging)

Syntax

```
console {
    facility severity;
}
```

Hierarchy Level [edit system [syslog](#)]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure the logging of system messages to the system console.

Options **facility**—Class of messages to log. To specify multiple classes, include multiple **facility severity** statements. For a list of the facilities, see “[Junos OS System Logging Facilities and Message Severity Levels](#)” on page 1546.

severity—Severity of the messages that belong to the facility specified by the paired **facility** name. Messages with severities of the specified level and higher are logged. For a list of the severities, see “[Junos OS System Logging Facilities and Message Severity Levels](#)” on page 1546.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- [Directing System Log Messages to the Console on page 1534](#)
- [Junos OS System Log Messages Reference](#)

explicit-priority

Syntax	explicit-priority;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> system syslog file <i>filename</i>], [edit logical-systems <i>logical-system-name</i> system syslog host], [edit system syslog file <i>filename</i>], [edit system syslog host]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Record the priority (facility and severity level) in each standard-format system log message directed to a file or remote destination. When the structured-data statement is also included at the [edit system syslog file <i>filename</i>] hierarchy level, this statement is ignored for the file.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Including Priority Information in System Log Messages on page 1536• <i>Junos OS System Log Reference for Security Devices</i>• structured-data on page 1703

facility-override

Syntax	facility-override <i>facility</i> ;
Hierarchy Level	[edit system syslog host]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Substitute an alternate facility for the default facilities used when messages are directed to a remote destination.
Options	<i>facility</i> —Alternate facility to substitute for the default facilities. For a list of the possible facilities, see Table 118 .
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Changing the Alternative Facility Name for System Log Messages Directed to a Remote Destination on page 1549• <i>Junos OS System Log Reference</i>

file (System Logging)

Syntax	<pre>file <i>filename</i> { <i>facility severity</i>; archive { <i>files number</i>; <i>size size</i>; (no-world-readable world-readable); } <i>explicit-priority</i>; <i>match "regular-expression"</i>; <i>structured-data</i> { <i>brief</i>; } }</pre>
Hierarchy Level	[edit system syslog]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	Configure the logging of system messages to a file.
Options	<p><i>facility</i>—Class of messages to log. To specify multiple classes, include multiple <i>facility severity</i> statements. For a list of the facilities, see “Junos OS System Logging Facilities and Message Severity Levels” on page 1546.</p> <p><i>file filename</i>—File in the <code>/var/log</code> directory in which to log messages from the specified facility. To log messages to more than one file, include more than one <i>file</i> statement.</p> <p><i>severity</i>—Severity of the messages that belong to the facility specified by the paired <i>facility</i> name. Messages with severities of the specified level and higher are logged. For a list of the severities, see “Junos OS System Logging Facilities and Message Severity Levels” on page 1546.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Directing System Log Messages to a Log File on page 1532 • Junos OS System Log Reference for Security Devices

files

Syntax	<code>files <i>number</i>;</code>
Hierarchy Level	[edit system syslog archive], [edit system syslog file <i>filename</i> archive]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for EX Series switches.
Description	Configure the maximum number of archived log files to retain. When the Junos OS logging utility has written a defined maximum amount of data to a log file <i>logfile</i> , it closes the file, compresses it, and renames it to <i>logfile.0.gz</i> (for information about the maximum file size, see size). The utility then opens and writes to a new file called <i>logfile</i> . When the new file reaches the maximum size, the <i>logfile.0.gz</i> file is renamed to <i>logfile.1.gz</i> , and the new file is closed, compressed, and renamed <i>logfile.0.gz</i> . By default, the logging facility creates up to ten archive files in this manner. Once the maximum number of archive files exists, each time the active log file reaches the maximum size, the contents of the oldest archive file are lost (overwritten by the next oldest file).
Options	<i>number</i> —Maximum number of archived files. Range: 1 through 1000 Default: 10 files
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Junos OS System Log Reference for Security Devices</i>• size on page 1702

host

Syntax	<pre> host (hostname other-routing-engine) { facility severity; explicit-priority; facility-override facility; log-prefix (System) string; match "regular-expression"; source-address source-address; structured-data { brief; } } </pre>
QFX Series	<pre> host (hostname { facility severity; explicit-priority; facility-override facility; log-prefix (System) string; match "regular-expression"; port; source-address source-address; } </pre>
TX Matrix Router and EX Series Switches	<pre> host (hostname other-routing-engine scc-master) { facility severity; explicit-priority; facility-override facility; log-prefix (System) string; match "regular-expression"; port; source-address source-address; } </pre>
TX Matrix Plus Router	<pre> host (hostname other-routing-engine sfc0-master) { facility severity; explicit-priority; facility-override facility; log-prefix (System) string; match "regular-expression"; port; source-address source-address; } </pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> system syslog], [edit system syslog]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the logging of system messages to a remote destination.

Options *facility*—Class of messages to log. To specify multiple classes, include multiple *facility severity* statements. For a list of the facilities, see [“Junos OS System Logging Facilities and Message Severity Levels” on page 1546](#).

hostname—IPv4 address, IPv6 address, or fully qualified hostname of the remote machine to which to direct messages. To direct messages to multiple remote machines, include a **host** statement for each one.

other-routing-engine—Direct messages to the other Routing Engine on a router or switch with two Routing Engines installed and operational.



NOTE: The **other-routing-engine** option is not applicable to the QFX Series.

port—Port number of the remote syslog server that can be modified.

scc-master—(TX Matrix routers only) On a T640 router that is part of a routing matrix, direct messages to the TX Matrix router.

severity—Severity of the messages that belong to the facility specified by the paired *facility* name. Messages with severities of the specified level and higher are logged. For a list of the severities, see [“Junos OS System Logging Facilities and Message Severity Levels” on page 1546](#).

sfc0-master—(TX Matrix Plus routers only) On a T1600 router that is part of a routing matrix, direct messages to the TX Matrix Plus router.

The remaining statements are explained separately.

Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
---------------------------------	---

Related Documentation	<ul style="list-style-type: none">• <i>Directing System Log Messages to a Remote Machine or the Other Routing Engine</i>• <i>Junos OS System Log Reference</i>
------------------------------	---

log-prefix (System)

Syntax	<code>log-prefix <i>string</i>;</code>
Hierarchy Level	[edit system syslog host]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Include a text string in each message directed to a remote destination.
Options	<i>string</i> —Text string to include in each message.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Adding a Text String to System Log Messages Directed to a Remote Destination on page 1531 • <i>Junos OS System Log Reference for Security Devices</i>


match

Syntax	<code>match "regular-expression";</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> system syslog file <i>filename</i>], [edit logical-systems <i>logical-system-name</i> system syslog user (<i>username</i> *)], [edit system syslog file <i>filename</i>], [edit system syslog host <i>hostname</i> other-routing-engine scc-master)], [edit system syslog user (<i>username</i> *)]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify a text string that must (or must not) appear in a message for the message to be logged to a destination.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Using Regular Expressions to Refine the Set of Logged Messages on page 1551

size (System)

Syntax	<code>size size;</code>
Hierarchy Level	[edit system syslog archive], [edit system syslog file <i>filename</i> archive]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the maximum amount of data that the Junos OS logging utility writes to a log file <i>logfile</i> before archiving it (closing it, compressing it, and changing its name to <i>logfile.0.gz</i>). The utility then opens and writes to a new file called <i>logfile</i> . For information about the number of archive files that the utility creates in this way, see files .
Options	size —Maximum size of each system log file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). Syntax: <i>xk</i> to specify the number of kilobytes, <i>xm</i> for the number of megabytes, or <i>xg</i> for the number of gigabytes Range: 64 KB through 1 GB Default: 1 MB for MX Series routers the QFX Series, and the OCX Series
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Specifying Log File Size, Number, and Archiving Properties on page 1543• Junos OS System Log Messages Reference• files on page 1698

structured-data

Syntax	structured-data { brief; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> system syslog file <i>filename</i>], [edit system syslog file <i>filename</i>]
Release Information	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Write system log messages to the log file in structured-data format, which complies with Internet draft draft-ietf-syslog-protocol-23, <i>The syslog Protocol</i> (http://tools.ietf.org/html/draft-ietf-syslog-protocol-23).
<div>  <p>NOTE: When this statement is included, other statements that specify the format for messages written to the file are ignored (the <code>explicit-priority</code> statement at the [edit system syslog file <i>filename</i>] hierarchy level and the <code>time-format</code> statement at the [edit system syslog] hierarchy level).</p> </div>	
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Logging Messages in Structured-Data Format</i> • <i>Junos OS System Log Reference for Security Devices</i> • explicit-priority on page 1696 • time-format on page 1706


syslog (System)

```
Syntax  syslog {
        allow-duplicates;
        archive {
            (binary-data | no-binary-data);
            files number;
            size maximum-file-size;
            start-time "YYYY-MM-DD.hh:mm";
            transfer-interval minutes;
            (world-readable | no-world-readable);
        }
        console {
            facility severity;
        }
        file filename {
            facility severity;
            explicit-priority;
            match "regular-expression";
            archive {
                (binary-data | no-binary-data);
                files number;
                size maximum-file-size;
                start-time "YYYY-MM-DD.hh:mm";
                transfer-interval minutes;
                (world-readable | no-world-readable);
            }
            structured-data {
                brief;
            }
        }
        host (hostname | other-routing-engine | scc-master) {
            facility severity;
            explicit-priority;
            facility-override facility;
            log-prefix string;
            match "regular-expression";
            source-address source-address;
            structured-data {
                brief;
            }
            port port number;
        }
        log-rotate-frequency frequency;
        server server name;
        source-address source-address;
        time-format (millisecond | year | year millisecond);
        user (username | *) {
            facility severity;
            match "regular-expression";
        }
    }
```

Hierarchy Level [edit system]

Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Configure the types of system log messages to send to files, to a remote destination, to user terminals, or to the system console.</p> <p>The remaining statements are explained separately.</p>
Options	<p>archive—Define parameters for archiving log messages.</p> <p>console—Send log messages of a specified class and severity to the console.</p> <p>file—Send log messages to a named file.</p> <p>host —Remote location to be notified of specific log messages.</p> <p>log-rotate-frequency—Configure the interval for checking logfile size and archiving messages.</p> <p>server—Name of the system log server in the inet.0 routing instance.</p> <p>source-address—Include a specified address as the source address for log messages.</p> <p>time-format—Additional information to include in the system log time stamp.</p> <p>user—Notify a specific user of the log event.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Junos OS System Log Overview</i> • <i>Junos OS System Log Messages Reference</i>

time-format

Syntax	<code>time-format (year millisecond year millisecond);</code>
Hierarchy Level	<code>[edit system syslog]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Include the year, the millisecond, or both, in the timestamp on every standard-format system log message. The additional information is included for messages directed to each destination configured by a file, console, or user statement at the <code>[edit system syslog]</code> hierarchy level. As of Junos OS Release 11.4, the additional time information is also sent to destinations configured by a host statement.</p> <p>By default, the timestamp specifies the month, date, hour, minute, and second when the message was logged—for example, Aug 21 12:36:30. However, the timestamp for traceoption messages is specified in milliseconds by default, and is independent of the <code>[edit system syslog time-format]</code> statement.</p>
	<div> NOTE: When the structured-data statement is included at the <code>[edit system syslog file <i>filename</i>]</code> hierarchy level, this statement is ignored for the file.</div>
Options	<p>millisecond—Include the millisecond in the timestamp.</p> <p>year—Include the year in the timestamp.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Including the Year or Millisecond in Timestamps on page 22• <i>Junos OS System Log Reference for Security Devices</i>• structured-data on page 1703

user (System Logging)

Syntax	<pre> user (username *) { facility severity; match "regular-expression"; } </pre>
Hierarchy Level	[edit system syslog]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	Configure the logging of system messages to user terminals.
Options	<p>* (the asterisk)—Log messages to the terminal sessions of all users who are currently logged in.</p> <p>facility—Class of messages to log. To specify multiple classes, include multiple facility severity statements. For a list of the facilities, see “Junos OS System Logging Facilities and Message Severity Levels” on page 1546.</p> <p>severity—Severity of the messages that belong to the facility specified by the paired facility name. Messages with severities the specified level and higher are logged. For a list of the severities, see “Junos OS System Logging Facilities and Message Severity Levels” on page 1546.</p> <p>username—Junos OS login name of the user whose terminal session is to receive system log messages. To log messages to more than one user's terminal session, include more than one user statement.</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Directing System Log Messages to a User Terminal on page 1533 • Junos OS System Logging Facilities and Message Severity Levels on page 1546 • Junos OS System Log Reference for Security Devices

PART 22

Operational Commands

- [Network Management on page 1711](#)
- [sFlow Technology on page 1729](#)
- [SNMP on page 1737](#)
- [System Logging on page 1775](#)

CHAPTER 55

Network Management

- monitor traffic
- ping

monitor traffic

Syntax `monitor traffic`
 `<brief | detail | extensive>`
 `<absolute-sequence>`
 `<count count>`
 `<interface interface-name>`
 `<layer2-headers>`
 `<matching matching>`
 `<no-domain-names>`
 `<no-promiscuous>`
 `<no-resolve>`
 `<no-timestamp>`
 `<print-ascii>`
 `<print-hex>`
 `<resolve-timeout>`
 `<size size>`

Release Information Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.0 for EX Series switches.
 Command introduced in Junos OS Release 11.1 for the QFX Series.
 Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Display packet headers or packets received and sent from the Routing Engine.



NOTE:

- Using the **monitor-traffic** command can degrade router or switch performance.
 - Delays from DNS resolution can be eliminated by using the **no-resolve** option.
-



NOTE: This command is not supported on the QFabric system.

Options **none**—(Optional) Display packet headers transmitted through **fxp0**. On a TX Matrix Plus router, display packet headers transmitted through **em0**.

brief | detail | extensive—(Optional) Display the specified level of output.

absolute-sequence—(Optional) Display absolute TCP sequence numbers.

count count—(Optional) Specify the number of packet headers to display (0 through 1,000,000). The **monitor traffic** command quits automatically after displaying the number of packets specified.

interface *interface-name*—(Optional) Specify the interface on which the **monitor traffic** command displays packet data. If no interface is specified, the **monitor traffic** command displays packet data arriving on the lowest-numbered interface.

layer2-headers—(Optional) Display the link-level header on each line.

matching *matching*—(Optional) Display packet headers that match a regular expression. Use matching expressions to define the level of detail with which the **monitor traffic** command filters and displays packet data.

no-domain-names—(Optional) Suppress the display of the domain portion of hostnames. With the **no-domain-names** option enabled, the **monitor traffic** command displays only **team** for the hostname **team.company.net**.

no-promiscuous—(Optional) Do not put the interface into promiscuous mode.

no-resolve—(Optional) Suppress reverse lookup of the IP addresses.

no-timestamp—(Optional) Suppress timestamps on displayed packets.

print-ascii—(Optional) Display each packet in ASCII format.

print-hex—(Optional) Display each packet, except the link-level header, in hexadecimal format.

resolve-timeout *timeout*—(Optional) Amount of time the router or switch waits for each reverse lookup before timing out. You can set the timeout for 1 through 4,294,967,295 seconds. The default is 4 seconds. To display each packet, use the **print-ascii**, **print-hex**, or **extensive** option.

size *size*—(Optional) Read but do not display up to the specified number of bytes for each packet. When set to **brief** output, the default packet size is 96 bytes and is adequate for capturing IP, ICMP, UDP, and TCP packet data. When set to **detail** and **extensive** output, the default packet size is 1514. The **monitor traffic** command truncates displayed packets if the matched data exceeds the configured size.

Additional Information In the **monitor traffic** command, you can specify an expression to match by using the **matching** option and including the expression in quotation marks:

```
monitor traffic matching "expression"
```

Replace ***expression*** with one or more of the match conditions listed in [Table 120](#).

Table 120: Match Conditions for the monitor traffic Command

Match Type	Condition	Description
Entity	host [<i>address</i> <i>hostname</i>]	Matches packets that contain the specified address or hostname. The protocol match conditions arp , ip , or rarp , or any of the directional match conditions can be prepended to the host match condition.
	net <i>address</i>	Matches packets with source or destination addresses containing the specified network address.
	net <i>address mask mask</i>	Matches packets containing the specified network address and subnet mask.
	port (<i>port-number</i> <i>port-name</i>)	Matches packets containing the specified source or destination TCP or UDP port number or port name. In place of the numeric port address, you can specify a text synonym, such as bgp (179), dhcp (67), or domain (53) (the port numbers are also listed).
Directional	dst	Matches packets going to the specified destination. This match condition can be prepended to any of the entity type match conditions.
	src	Matches packets from a specified source. This match condition can be prepended to any of the entity type match conditions.
	src and dst	Matches packets that contain the specified source and destination addresses. This match condition can be prepended to any of the entity type match conditions.
	src or dst	Matches packets containing either of the specified addresses. This match condition can be prepended to any of the entity type match conditions.
Packet Length	less <i>value</i>	Matches packets shorter than or equal to the specified value, in bytes.
	greater <i>value</i>	Matches packets longer than or equal to the specified value, in bytes.

Table 120: Match Conditions for the monitor traffic Command (*continued*)

Match Type	Condition	Description
Protocol	amt	Matches all AMT packets. Use the extensive level of output to decode the inner IGMP packets in addition to the AMT outer packet.
	arp	Matches all ARP packets.
	ether	Matches all Ethernet packets.
	ether (broadcast multicast)	Matches broadcast or multicast Ethernet frames. This match condition can be prepended with src and dst .
	ether protocol (address (arp ip rarp))	Matches packets with the specified Ethernet address or Ethernet packets of the specified protocol type. The ether protocol arguments arp , ip , and rarp are also independent match conditions, so they must be preceded by a backslash (\) when used in the ether protocol match condition.
	icmp	Matches all ICMP packets.
	ip	Matches all IP packets.
	ip (broadcast multicast)	Matches broadcast or multicast IP packets.
	ip protocol (address (icmp igmp tcp udp))	Matches packets with the specified address or protocol type. The ip protocol arguments icmp , tcp , and udp are also independent match conditions, so they must be preceded by a backslash (\) when used in the ip protocol match condition.
	isis	Matches all IS-IS routing messages.
	proto ip-protocol-number	Matches packets whose headers contain the specified IP protocol number.
	rarp	Matches all RARP packets.
	tcp	Matches all TCP datagrams.
	udp	Matches all UDP datagrams.

To combine expressions, use the logical operators listed in [Table 121](#).

Table 121: Logical Operators for the monitor traffic Command

Logical Operator (Highest to Lowest Precedence)	Description
!	Logical NOT. If the first condition does not match, the next condition is evaluated.
&&	Logical AND. If the first condition matches, the next condition is evaluated. If the first condition does not match, the next condition is skipped.
	Logical OR. If the first condition matches, the next condition is skipped. If the first condition does not match, the next condition is evaluated.
()	Group operators to override default precedence order. Parentheses are special characters, each of which must be preceded by a backslash (\).

You can use relational operators to compare arithmetic expressions composed of integer constants, binary operators, a length operator, and special packet data accessors. The arithmetic expression matching condition uses the following syntax:

```
monitor traffic matching "ether[0] & 1 != 0"arithmetic_expression relational_operator arithmetic_expression
```

The packet data accessor uses the following syntax:

```
protocol [byte-offset <size>]
```

The optional *size* field represents the number of bytes examined in the packet header. The available values are 1, 2, or 4 bytes. The following sample command captures all multicast traffic:

```
user@host> monitor traffic matching "ether[0] & 1 != 0"
```

To specify match conditions that have a numeric value, use the arithmetic and relational operators listed in [Table 122](#).



NOTE: Because the Packet Forwarding Engine removes Layer 2 header information before sending packets to the Routing Engine:

- The **monitor traffic** command cannot apply match conditions to inbound traffic.
- The **monitor traffic interface** command also cannot apply match conditions for Layer 3 and Layer 4 packet data, resulting in the match pipe option (**| match**) for this command for Layer 3 and Layer 4 packets not working either. Therefore, ensure that you specify match conditions as described in this command summary. For more information about match conditions, see [Table 120](#).
- The 802.1Q VLAN tag information included in the Layer 2 header is removed from all inbound traffic packets. Because the **monitor traffic interface ae[x]** command for aggregated Ethernet interfaces (such as ae0) only shows inbound traffic data, the command does not show VLAN tag information in the output.

Table 122: Arithmetic and Relational Operators for the monitor traffic Command

Arithmetic or Relational Operator	Description
Arithmetic Operator	
+	Addition operator.
-	Subtraction operator.
/	Division operator.
&	Bitwise AND.
*	Bitwise exclusive OR.
	Bitwise inclusive OR.
Relational Operator (Highest to Lowest Precedence)	
<=	If the first expression is less than or equal to the second, the packet matches.
>=	If the first expression is greater than or equal to the second, the packet matches.
<	If the first expression is less than the second, the packet matches.
>	If the first expression is greater than the second, the packet matches.

Table 122: Arithmetic and Relational Operators for the monitor traffic Command (*continued*)

Arithmetic or Relational Operator	Description
=	If the compared expressions are equal, the packet matches.
!=	If the compared expressions are unequal, the packet matches.

Required Privilege Level trace
maintenance

List of Sample Output [monitor traffic count on page 1718](#)
[monitor traffic detail count on page 1718](#)
[monitor traffic extensive \(Absolute Sequence\) on page 1719](#)
[monitor traffic extensive \(Relative Sequence\) on page 1719](#)
[monitor traffic extensive count on page 1719](#)
[monitor traffic interface on page 1719](#)
[monitor traffic matching on page 1720](#)
[monitor traffic \(TX Matrix Plus Router\) on page 1720](#)
[monitor traffic \(QFX3500 Switch\) on page 1721](#)
[monitor traffic matching icmp on page 1721](#)
[monitor traffic matching IP protocol number on page 1722](#)
[monitor traffic matching arp on page 1722](#)
[monitor traffic matching port on page 1723](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

monitor traffic count

```
user@host> monitor traffic count 2
listening on fxp0
04:35:49.814125 In my-server.home.net.1295 > my-server.work.net.telnet: . ack
4122529478 win 16798 (DF)
04:35:49.814185
Out my-server.work.net.telnet > my-server.home.net.1295: P
1:38(37) ack 0 win 17680 (DF) [tos 0x10]
```

monitor traffic detail count

```
user@host> monitor traffic detail count 2
listening on fxp0
04:38:16.265864 In my-server.home.net.1295 > my-server.work.net.telnet: . ack
4122529971 win 17678 (DF) (ttl 121, id 6812)
04:38:16.265926
Out my-server.work.net.telnet.telnet > my-server.home.net.1295: P 1:38(37) ack 0
win 17680 (DF) [tos 0x10] (ttl 6)
```

monitor traffic extensive (Absolute Sequence)

```

user@host> monitor traffic extensive no-domain-names no-resolve no-timestamp count 20
matching "tcp" absolute-sequence
listening on fxp0
In 207.17.136.193.179 > 192.168.4.227.1024: . 4042780859:4042780859(0)
ack 1845421797 win 16384 <nop,nop,timestamp 4935628 965951> [tos 0xc0] (ttl )
In 207.17.136.193.179 > 192.168.4.227.1024: P 4042780859:4042780912(53)
ack 1845421797 win 16384
<nop,nop,timestamp 4935628 965951>:
BGP [|BGP UPDAT)
In 192.168.4.227.1024 > 207.17.136.193.179:
P 1845421797:1845421852(55) ack 4042780912 win 16384 <nop,nop,timestamp 965951
4935628>: BGP [|BGP UPDAT)
...

```

monitor traffic extensive (Relative Sequence)

```

user@host> monitor traffic extensive no-domain-names no-resolve no-timestamp count 20
matching "tcp"
listening on fxp0
In 172.24.248.221.1680 > 192.168.4.210.23: . 396159737:396159737(0)
ack 1664980689 win 17574 (DF) (ttl 121, id 50003)
Out 192.168.4.210.23 > 172.24.248.221.1680: P 1:40(39)
ack 0 win 17680 (DF) [tos 0x10] (ttl 64, id 5394)
In 207.17.136.193.179 > 192.168.4.227.1024: P 4042775817:4042775874(57)
ack 1845416593 win 16384 <nop,nop,timestamp 4935379 965690>: BGP [|BGP UPDAT)
...

```

monitor traffic extensive count

```

user@host> monitor traffic extensive count 5 no-domain-names no-resolve
listening on fxp013:18:17.406933
In 192.168.4.206.2723610880 > 172.17.28.8.2049:
40 null (ttl 64, id 38367)13:18:17.407577
In 172.17.28.8.2049 > 192.168.4.206.2723610880:
reply ok 28 null (ttl 61, id 35495)13:18:17.541140
In 0:e0:1e:42:9c:e0 0:e0:1e:42:9c:e0 9000 60:
0000 0100 0000 0000
0000 0000 0000 0000
0000 0000 0000 0000
0000 0000 0000 0000
0000 0000 0000 0000
0000 0000 000013:18:17.591513
In 172.24.248.156.4139 > 192.168.4.210.23:
3556964918:3556964918(0)
ack 295526518 win 17601 (DF)
(ttl 121, id 14)13:18:17.591568
Out 192.168.4.210.23 >
172.24.248.156.4139: P 1:40(39)
ack 0 win 17680 (DF) [tos 0x10]
(ttl 64, id 52376)

```

monitor traffic interface

```

user@host> monitor traffic interface fxp0
listening on fxp0.0
18:17:28.800650 In server.home.net.723 > host1-0.lab.home.net.log
18:17:28.800733 Out host2-0.lab.home.net.login > server.home.net.7
18:17:28.817813 In host30.lab.home.net.syslog > host40.home0

```

```
18:17:28.817846 In host30.lab.home.net.syslog > host40.home0
...
```

monitor traffic matching

```
user@host> monitor traffic matching "net 192.168.1.0/24"
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Address resolution is ON. Use <no-resolve> to avoid any reverse lookup delay.
Address resolution timeout is 4s.
Listening on fxp0, capture size 96 bytes

Reverse lookup for 192.168.1.255 failed (check DNS reachability).
Other reverse lookup failures will not be reported.
Use no-resolve to avoid reverse lookups on IP addresses.

21:55:54.003511 In IP truncated-ip - 18 bytes missing!
192.168.1.17.netbios-ns > 192.168.1.255.netbios-ns: UDP, length 50
21:55:54.003585 Out IP truncated-ip - 18 bytes missing!
192.168.1.17.netbios-ns > 192.168.1.255.netbios-ns: UDP, length 50
21:55:54.003864 In arp who-has 192.168.1.17 tell 192.168.1.9
...
```

monitor traffic (TX Matrix Plus Router)

```
user@host> monitor traffic
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Address resolution is ON. Use <no-resolve> to avoid any reverse lookup delay.
Address resolution timeout is 4s.
Listening on em0, capture size 96 bytes
04:11:59.862121 Out IP truncated-ip - 25 bytes missing!
summit-em0.englab.juniper.net.syslog > sv-log-01.englab.juniper.net.syslog:
SYSLOG kernel.info, length: 57
04:11:59.862303
Out IP truncated-ip - 25 bytes missing!
summit-em0.englab.juniper.net.syslog >
sv-log-02.englab.juniper.net.syslog: SYSLOG kernel.info, length: 57
04:11:59.923948
In IP aj-em0.englab.juniper.net.65235 >
summit-em0.englab.juniper.net.telnet: .
ack 1087492766 win 33304 <nop,nop,timestamp 42366734 993490>
04:11:59.923983 Out IP truncated-ip - 232 bytes missing!
summit-em0.englab.juniper.net.telnet > aj-em0.englab.juniper.net.65235: P
1:241(240) ack 0 win 33304
<nop,nop,timestamp 993590 42366734>
04:12:00.022900
In IP aj-em0.englab.juniper.net.65235 >
summit-em0.englab.juniper.net.telnet: . ack 241 win 33304 <nop,nop,timestamp
42366834 993590>
04:12:00.141204
In IP truncated-ip - 40 bytes missing!
ipg-lnx-shell1.juniper.net.46182 > summit-em0.englab.juniper.net.telnet: P
2950530356:2950530404(48) ack 485494987 win 63712
<nop,nop,timestamp 1308555294 987086>
04:12:00.141345
Out IP summit-em0.englab.juniper.net.telnet >
ipg-lnx-shell1.juniper.net.46182: P 1:6(5)
ack 48 win 33304
<nop,nop,timestamp 993809 1308555294>
04:12:00.141572
In IP ipg-lnx-shell1.juniper.net.46182 >
summit-em0.englab.juniper.net.telnet: .
```



```

ack 6 win 63712
<nop,nop,timestamp 1308555294 993809>
04:12:00.141597
Out IP summit-em0.englab.juniper.net.telnet >
ipg-lnx-shell11.juniper.net.46182: P 6:10(4) ack 48 win 33304
<nop,nop,timestamp 993810 1308555294>
04:12:00.141821
In IP ipg-lnx-shell11.juniper.net.46182 >
summit-em0.englab.juniper.net.telnet: .
ack 10 win 63712 <nop,nop,timestamp 1308555294 993810>
04:12:00.141837 Out IP truncated-ip - 2 bytes missing!
summit-em0.englab.juniper.net.telnet >
ipg-lnx-shell11.juniper.net.46182: P 10:20(10) ack 48 win 33304
<nop,nop,timestamp 993810 1308555294>
04:12:00.142072
In IP ipg-lnx-shell11.juniper.net.46182 >
summit-em0.englab.juniper.net.telnet: . ack 20 win 63712
<nop,nop,timestamp 1308555294 993810>
04:12:00.142089 Out IP summit-em0.englab.juniper.net.telnet >
ipg-lnx-shell11.juniper.net.46182: P 20:28(8) ack 48 win 33304 <nop,nop,timestamp
 993810 1308555294>
04:12:00.142321
In IP ipg-lnx-shell11.juniper.net.46182 >
summit-em0.englab.juniper.net.telnet: .
ack 28 win 63712 <nop,nop,timestamp 1308555294 993810>
04:12:00.142337
Out IP truncated-ip - 1 bytes missing!
summit-em0.englab.juniper.net.telnet >
ipg-lnx-shell11.juniper.net.46182: P 28:37(9) ack 48 win 33304 <nop,nop,timestamp
993810 1308555294>
...

```

monitor traffic (QFX3500 Switch)

```

user@switch> monitor traffic
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Address resolution is ON. Use <no-resolve> to avoid any reverse lookup delay.
Address resolution timeout is 4s.
Listening on me4, capture size 96 bytes
Reverse lookup for 172.22.16.246 failed (check DNS reachability).
Other reverse lookup failures will not be reported.
Use <no-resolve> to avoid reverse lookups on IP addresses.
16:35:32.240873 Out IP truncated-ip - 112 bytes missing!
labqfx-me0.lab4.juniper.net.ssh >
172.22.16.246.telefinder: P 4200727624:4200727756(132) ack 2889954831 win 65535
16:35:32.240900 Out IP truncated-ip - 176 bytes missing!
labqfx-me0.lab4.juniper.net.ssh >
172.22.16.246.telefinder: P 132:328(196) ack 1 win 65535
...

```

monitor traffic matching icmp

```

user@host> monitor traffic matching "icmp" no-resolve
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Address resolution is OFF.
Listening on me0, capture size 96 bytes

09:23:17.728737 In IP 172.19.10.9 > 10.10.211.93: ICMP echo request, id 1, seq
322, length 40
09:23:17.728780 Out IP 10.10.211.93 > 172.19.10.9: ICMP echo reply, id 1, seq
322, length 40

```

```

09:23:18.735848 In IP 172.19.10.9 > 10.10.211.93: ICMP echo request, id 1, seq
323, length 40
09:23:18.735891 Out IP 10.10.211.93 > 172.19.10.9: ICMP echo reply, id 1, seq
323, length 40
09:23:19.749732 In IP 172.19.10.9 > 10.10.211.93: ICMP echo request, id 1, seq
324, length 40
09:23:19.749775 Out IP 10.10.211.93 > 172.19.10.9: ICMP echo reply, id 1, seq
324, length 40
09:23:20.749747 In IP 172.19.10.9 > 10.10.211.93: ICMP echo request, id 1, seq
325, length 40
09:23:20.749791 Out IP 10.10.211.93 > 172.19.10.9: ICMP echo reply, id 1, seq
325, length 40
...

```

monitor traffic matching IP protocol number

```

user@host> monitor traffic matching "proto 89" no-resolve
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Address resolution is OFF.
Listening on me0, capture size 96 bytes

13:06:14.700311 In IP truncated-ip - 16 bytes missing! 10.94.211.254 > 224.0.0.
5: OSPFv2, Hello, length 56
13:06:16.067010 In IP truncated-ip - 20 bytes missing! 10.94.211.102 > 224.0.0.
5: OSPFv2, Hello, length 60
13:06:16.287566 In IP truncated-ip - 20 bytes missing! 10.94.211.142 > 224.0.0.
5: OSPFv2, Hello, length 60
13:06:20.758500 In IP truncated-ip - 16 bytes missing! 10.200.211.254 > 224.0.0.
5: OSPFv2, Hello, length 56
13:06:24.309882 In IP truncated-ip - 20 bytes missing! 10.94.211.102 > 224.0.0.
5: OSPFv2, Hello, length 60
13:06:24.396699 In IP truncated-ip - 16 bytes missing! 10.94.211.254 > 224.0.0.
5: OSPFv2, Hello, length 56
13:06:25.067386 In IP truncated-ip - 20 bytes missing! 10.94.211.142 > 224.0.0.
5: OSPFv2, Hello, length 60
13:06:29.499988 In IP truncated-ip - 16 bytes missing! 10.200.211.254 > 224.0.0.
5: OSPFv2, Hello, length 56
13:06:32.858753 In IP truncated-ip - 20 bytes missing! 10.94.211.102 > 224.0.0.
5: OSPFv2, Hello, length 60
...

```

monitor traffic matching arp

```

user@host> monitor traffic matching "arp" no-resolve
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Address resolution is OFF.
Listening on me0, capture size 96 bytes

11:57:54.664501 In arp who-has 10.10.213.109 (00:1f:d5:f3:28:30) tell 10.10.213.31
11:57:56.828387 In arp who-has 10.10.213.233 (00:24:9d:06:77:4f) tell 10.10.213.31
11:58:01.735803 In arp who-has 10.10.213.251 (88:e0:f4:1d:41:40) tell 10.10.213.31
11:58:04.663241 In arp who-has 10.10.213.254 tell 10.94.211.170
11:58:28.488191 In arp who-has 10.10.213.149 (00:e0:91:c2:ff:8d) tell 10.10.213.31
11:58:41.858612 In arp who-has 10.10.213.148 tell 10.94.211.254
11:58:42.621533 In arp who-has 10.10.213.254 (5f:5e:ac:79:49:81) tell 10.10.213.31
11:58:44.533391 In arp who-has 10.10.213.186 tell 10.94.211.254
11:58:45.170405 In arp who-has 10.10.213.186 tell 10.94.211.254
11:58:45.770512 In arp who-has 10.10.213.186 tell 10.94.211.254

```

monitor traffic matching port

```

user@host> monitor traffic matching "port 22" no-resolve
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Address resolution is OFF.
Listening on me0, capture size 96 bytes

13:14:19.108089 In IP 172.44.33.22.56714 > 10.19.300.05.22: S
2210742342:2210742342(0) win 65535 <mss 1360,nop,wscale 7,nop,nop,sackOK>
13:14:19.108165 Out IP 10.19.300.05.22 > 172.44.33.22.56714: S 23075150:23075150(0)
ack 2210742343 win 65535 <mss 1460,nop,wscale 1,sackOK,eol>
13:14:19.136883 In IP 172.44.33.22.56714 > 10.19.300.05.22: . ack 1 win 32768
13:14:19.231364 Out IP truncated-ip - 1 bytes missing! 10.19.300.05.22 >
172.29.102.9.56714: P 1:22(21) ack 1 win 33320
13:14:19.260174 In IP truncated-ip - 10 bytes missing! 172.44.33.22.56714 >
10.94.211.93.22: P 1:31(30) ack 22 win 32767
13:14:19.284865 Out IP truncated-ip - 964 bytes missing! 10.19.300.05.22 >
172.29.102.9.56714: P 22:1006(984) ack 31 win 33320
13:14:19.314549 In IP truncated-ip - 652 bytes missing! 172.44.33.22.56714 >
10.94.211.93.22: P 31:703(672) ack 1006 win 32760
13:14:19.414135 Out IP 10.19.300.05.22 > 172.44.33.22.56714: . ack 703 win 33320
13:14:19.443858 In IP 172.44.33.22.56714 > 10.19.300.05.22: P 703:719(16) ack
1006 win 32760
13:14:19.467379 Out IP truncated-ip - 516 bytes missing! 10.19.300.05.22 >
172.29.102.9.56714: P 1006:1542(536) ack 719 win 33320
13:14:19.734097 In IP 172.44.33.22.56714 > 10.19.300.05.22: . ack 1542 win 32768
13:14:19.843574 In IP truncated-ip - 508 bytes missing! 172.44.33.22.56714 >
10.94.211.93.22: P 719:1247(528) ack 1542 win 32768
...

```


ping

List of Syntax [Syntax on page 1724](#)
 [Syntax \(QFX Series\) on page 1724](#)

Syntax `ping host`
 `<bypass-routing>`
 `<count requests>`
 `<detail>`
 `<do-not-fragment>`
 `<inet | inet6>`
 `<interface source-interface>`
 `<interval seconds>`
 `<logical-system logical-system-name>`
 `<loose-source value>`
 `<mac-address mac-address>`
 `<no-resolve>`
 `<pattern string>`
 `<rapid>`
 `<record-route>`
 `<routing-instance routing-instance-name>`
 `<size bytes>`
 `<source source-address>`
 `<strict >`
 `<strict-source value.>`
 `<tos type-of-service>`
 `<ttl value>`
 `<verbose>`
 `<vpls instance-name>`
 `<wait seconds>`

Syntax (QFX Series) `ping host`
 `<bypass-routing>`
 `<count requests>`
 `<detail>`
 `<do-not-fragment>`
 `<inet>`
 `<interface source-interface>`
 `<interval seconds>`
 `<logical-system logical-system-name>`
 `<loose-source value>`
 `<mac-address mac-address>`
 `<no-resolve>`
 `<pattern string>`
 `<rapid>`
 `<record-route>`
 `<routing-instance routing-instance-name>`
 `<size bytes>`
 `<source source-address>`
 `<strict>`
 `< strict-source value>`
 `<tos type-of-service>`
 `<ttl value>`
 `<verbose>`

<wait *seconds*>

Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Check host reachability and network connectivity. The ping command sends Internet Control Message Protocol (ICMP) ECHO_REQUEST messages to elicit ICMP ECHO_RESPONSE messages from the specified host. Press Ctrl+c to interrupt a ping command.</p>
Options	<p>host—IP address or hostname of the remote system to ping.</p> <p>bypass-routing—(Optional) Bypass the normal routing tables and send ping requests directly to a system on an attached network. If the system is not on a directly attached network, an error is returned. Use this option to ping a local system through an interface that has no route through it.</p> <p>count requests—(Optional) Number of ping requests to send. The range of values is 1 through 2,000,000,000. The default value is an unlimited number of requests.</p> <p>detail—(Optional) Include in the output the interface on which the ping reply was received.</p> <p>do-not-fragment—(Optional) Set the do-not-fragment (DF) flag in the IP header of the ping packets. For IPv6 packets, this option disables fragmentation.</p> <div style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;"> <p> NOTE: In Junos OS Release 11.1 and later, when issuing the ping command for an IPv6 route with the do-not-fragment option, the maximum ping packet size is calculated by subtracting 48 bytes (40 bytes for the IPV6 header and 8 bytes for the ICMP header) from the MTU. Therefore, if the ping packet size (including the 48-byte header) is greater than the MTU, the ping operation might fail.</p> </div> <p>inet—(Optional) Ping Packet Forwarding Engine IPv4 routes.</p> <p>inet6—(Optional) Ping Packet Forwarding Engine IPv6 routes.</p> <p>interface source-interface—(Optional) Interface to use to send the ping requests.</p> <p>interval seconds—(Optional) How often to send ping requests. The range of values, in seconds, is 1 through infinity. The default value is 1.</p> <p>logical-system logical-system-name—(Optional) Name of logical system from which to send the ping requests.</p> <p>Alternatively, enter the set cli logical-system logical-system-name command and then run the ping command. To return to the main router or switch, enter the clear cli logical-system command.</p>

loose-source value—(Optional) Intermediate loose source route entry (IPv4). Open a set of values.

mac-address mac-address—(Optional) Ping the physical or hardware address of the remote system you are trying to reach.

no-resolve—(Optional) Do not attempt to determine the hostname that corresponds to the IP address.

pattern string—(Optional) Specify a hexadecimal fill pattern to include in the ping packet.

rapid—(Optional) Send ping requests rapidly. The results are reported in a single message, not in individual messages for each ping request. By default, five ping requests are sent before the results are reported. To change the number of requests, include the **count** option.

record-route—(Optional) Record and report the packet's path (IPv4).

routing-instance routing-instance-name—(Optional) Name of the routing instance for the ping attempt.

size bytes—(Optional) Size of ping request packets. The range of values, in bytes, is 0 through 65,468. The default value is 56, which is effectively 64 bytes because 8 bytes of ICMP header data are added to the packet.

source source-address—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (lo.0).

strict—(Optional) Use the strict source route option (IPv4).

strict-source value—(Optional) Intermediate strict source route entry (IPv4). Open a set of values.

tos type-of-service—(Optional) Set the type-of-service (ToS) field in the IP header of the ping packets. The range of values is 0 through 255.

If the device configuration includes the **dscp-code-point value** statement at the **[edit class-of-service host-outbound-traffic]** hierarchy level, the configured DSCP value overrides the value specified in this command option. In this case, the ToS field of ICMP echo request packets sent on behalf of this command carries the DSCP value specified in the **dscp-code-point** configuration statement instead of the value you specify in this command option.

ttl value—(Optional) Time-to-live (TTL) value to include in the ping request (IPv6). The range of values is 0 through 255.

verbose—(Optional) Display detailed output.

vpls instance-name—(Optional) Ping the instance to which this VPLS belongs.

wait seconds—(Optional) Maximum wait time, in seconds, after the final packet is sent. If this option is not specified, the default delay is 10 seconds. If this option is used without the count option, a default count of 5 packets is used.

Required Privilege Level	network
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Junos OS ICMPv4 Rate Limit for ICMPv4 Routing Engine Messages</i>
List of Sample Output	ping hostname on page 1727 ping hostname rapid on page 1727 ping hostname size count on page 1727
Output Fields	<p>When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. These packets are not counted in the received packets count. They are accounted for separately.</p>

Sample Output

ping hostname

```
user@host> ping skye
PING skye.net (192.168.169.254): 56 data bytes
64 bytes from 192.168.169.254: icmp_seq=0 ttl=253 time=1.028 ms
64 bytes from 192.168.169.254: icmp_seq=1 ttl=253 time=1.053 ms
64 bytes from 192.168.169.254: icmp_seq=2 ttl=253 time=1.025 ms
64 bytes from 192.168.169.254: icmp_seq=3 ttl=253 time=1.098 ms
64 bytes from 192.168.169.254: icmp_seq=4 ttl=253 time=1.032 ms
64 bytes from 192.168.169.254: icmp_seq=5 ttl=253 time=1.044 ms
^C [abort]
```

ping hostname rapid

```
user@host> ping skye rapid
PING skye.net (192.168.169.254): 56 data bytes
!!!!
--- skye.net ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.956/0.974/1.025/0.026 ms
```

ping hostname size count

```
user@host> ping skye size 200 count 5
PING skye.net (192.168.169.254): 200 data bytes
208 bytes from 192.168.169.254: icmp_seq=0 ttl=253 time=1.759 ms
208 bytes from 192.168.169.254: icmp_seq=1 ttl=253 time=2.075 ms
208 bytes from 192.168.169.254: icmp_seq=2 ttl=253 time=1.843 ms
208 bytes from 192.168.169.254: icmp_seq=3 ttl=253 time=1.803 ms
208 bytes from 192.168.169.254: icmp_seq=4 ttl=253 time=17.898 ms

--- skye.net ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max = 1.759/5.075/17.898 ms
```


CHAPTER 56

sFlow Technology

- `clear sflow collector statistics`
- `show sflow`
- `show sflow collector`
- `show sflow interface`

clear sflow collector statistics

Syntax	clear sflow collector statistics
Release Information	Command introduced in Junos OS Release 11.3 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Clear the sample counters for all sFlow collectors.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• Example: Monitoring Network Traffic Using sFlow Technology on page 1453• Configuring sFlow Technology on page 1457• show sflow collector on page 1733
List of Sample Output	clear sflow collector statistics on page 1730

Sample Output

clear sflow collector statistics

The following example shows two output examples for the **show sflow collector** command, one before and one after the **clear sflow collector statistics** command was issued.

```
user@host> show sflow collector
Collector      Udp-port      No. of samples
address
10.1.1.1       6343          3174
10.1.2.1       6343          3562
```

```
user@host> clear sflow collector statistics
```

```
user@host> show sflow collector
Collector      Udp-port      No. of samples
address
10.1.1.1       6343          0
10.1.2.1       6343          0
```

show sflow

Syntax	show sflow <collector> <interface>
Release Information	Command introduced in Junos OS Release 11.3 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Display sFlow configuration information.
Options	<p>none—Display all sFlow configuration information.</p> <p>collector—(Optional) Display a list of configured sFlow collectors and their properties.</p> <p>interface—(Optional) Display the interfaces on which sFlow technology is enabled and the sampling parameters.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show sflow interface on page 1734 • show sflow collector on page 1733 • clear sflow collector statistics on page 1730 • Example: Monitoring Network Traffic Using sFlow Technology on page 1453 • Configuring sFlow Technology on page 1457
List of Sample Output	show sflow on page 1732
Output Fields	Table 123 lists the output fields for the show sflow command. Output fields are listed in the approximate order in which they appear.

Table 123: show sflow Output Fields

Field Name	Field Description	Level of Output
sFlow	Status of the feature: Enabled or Disabled .	All levels
Sample limit	Number of packets sampled per second. This sample limit cannot be configured and is set to 300 packets per second.	All levels
Polling interval	Interval at which the sFlow agent polls the interface.	All levels
Sample rate egress	Rate at which egress packets are sampled.	All levels
Sample rate ingress	Rate at which ingress packets are sampled.	All levels
Agent ID	IP address assigned to the sFlow agent.	All levels

Table 123: show sflow Output Fields (*continued*)

Field Name	Field Description	Level of Output
Source IP address	Source IP address for the sFlow packets.	All levels

Sample Output

show sflow

```
user@host> show sflow
```

```
sFlow           : Enabled
Sample limit    : 300 packets/second
Polling interval : 20 second
Sample rate egress : 1:2048: Disabled
Sample rate ingress : 1:1000: Enabled
Agent ID        : 10.93.54.7
Source IP address : 10.93.54.7
```

show sflow collector

Syntax	show sflow collector
Release Information	Command introduced in Junos OS Release 11.3 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Display a list of configured sFlow collectors and their properties.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear sflow collector statistics on page 1730 • show sflow on page 1731 • show sflow interface on page 1734 • Example: Monitoring Network Traffic Using sFlow Technology on page 1453 • Configuring sFlow Technology on page 1457
List of Sample Output	show sflow collector on page 1733
Output Fields	Table 124 lists the output fields for the show sflow collector command. Output fields are listed in the approximate order in which they appear.

Table 124: show sflow collector Output Fields

Field Name	Field Description	Level of Output
Collector address	IP address of the collector.	All levels
UDP-Port	UDP port number of the collector.	All levels
No. of samples	Number of samples collected.	All levels

Sample Output

show sflow collector

```
user@host> show sflow collector
```

```

Collector      Udp-port      No. of samples
address
10.204.32.46   6343          1000
100.204.32.76 3400          1000
```

show sflow interface

Syntax	show sflow interface
Release Information	Command introduced in Junos OS Release 11.3 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Display the interfaces on which sFlow is enabled and the sampling parameters for the interface.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show sflow on page 1731 • show sflow collector on page 1733 • Example: Monitoring Network Traffic Using sFlow Technology on page 1453 • Configuring sFlow Technology on page 1457
List of Sample Output	show sflow interface (QFX3500 Switch in Standalone Mode) on page 1734 show sflow interface (QFabric System) on page 1735
Output Fields	Table 125 lists the output fields for the show sflow interface command. Output fields are listed in the approximate order in which they appear.

Table 125: show sflow interface Output Fields

Field Name	Field Description	Level of Output
Interface	Interface on which sFlow technology is enabled.	All levels
Status Egress	Indicates whether an egress sample rate is enabled.	All levels
Status Ingress	Indicates whether an ingress sample rate is enabled.	All levels
Sample rate Egress	Rate at which egress packets are sampled.	All levels
Sample rate Ingress	Rate at which ingress packets are sampled.	All levels
Adapted sample rate Egress	Adapted rate at which egress packets are sampled.	All levels
Adapted sample rate Ingress	Adapted rate at which ingress packets are sampled.	All levels
Polling-interval	Interval at which the sFlow agent polls the interface.	All levels

Sample Output

show sflow interface (QFX3500 Switch in Standalone Mode)

```
user@host> show sflow interface
```

Interface	Status	Sample rate		Adapted sample rate			Polling-interval
		Egress	Ingress	Egress	Ingress	Egress	
xe-0/0/0.0	Enabled	Disabled	1000	2048	1000	2048	20
xe-1/0/1.0	Enabled	Disabled	1000	2048	1000	2048	20

Sample Output

show sflow interface (QFabric System)

```

user@host> show sflow interface
Interface  Status      Sample rate    Adapted sample rate  Polling-interval
           Egress Ingress  Egress Ingress  Egress Ingress
node1:xe-0/0/0.0  Enabled Disabled 1000    2048    1000    2048
20
node2:xe-1/0/1.0  Enabled Disabled 1000    2048    1000    2048
20
node4:xe-1/0/0.0  Enabled Disabled 1000    2048    1000    2048
20

```


CHAPTER 57

SNMP

- clear snmp history
- clear snmp statistics
- request snmp spoof-trap
- request snmp utility-mib clear instance
- request snmp utility-mib set instance
- show snmp health-monitor
- show snmp inform-statistics
- show snmp mib
- show snmp rmon
- show snmp rmon history
- show snmp statistics
- show snmp v3

clear snmp history

Syntax	clear snmp history (<i>index</i> all)
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Delete the samples of Ethernet statistics collected for a history group.
Options	all —Clear all the entries in the history index. index —Clear the contents of the specified entry in the history index.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• clear snmp statistics on page 1739

clear snmp statistics

Syntax	clear snmp statistics
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Clear Simple Network Management Protocol (SNMP) statistics.
Options	This command has no options.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • show snmp statistics on page 1764
List of Sample Output	clear snmp statistics on page 1739
Output Fields	See show snmp statistics for an explanation of output fields.

Sample Output

clear snmp statistics

In the following example, SNMP statistics are displayed before and after the **clear snmp statistics** command is issued:

```
user@host> show snmp statistics
SNMP statistics:
  Input:
    Packets: 8, Bad versions: 0, Bad community names: 0,
    Bad community uses: 0, ASN parse errors: 0,
    Too bigs: 0, No such names: 0, Bad values: 0,
    Read onlys: 0, General errors: 0,
    Total request varbinds: 8, Total set varbinds: 0,
    Get requests: 0, Get nexts: 8, Set requests: 0,
    Get responses: 0, Traps: 0,
    Silent drops: 0, Proxy drops 0
  Output:
    Packets: 2298, Too bigs: 0, No such names: 0,
    Bad values: 0, General errors: 0,
    Get requests: 0, Get nexts: 0, Set requests: 0,
    Get responses: 8, Traps: 2290
```

```
user@host> clear snmp statistics
```

```
user@host> show snmp statistics
SNMP statistics:
  Input:
    Packets: 0, Bad versions: 0, Bad community names: 0,
    Bad community uses: 0, ASN parse errors: 0,
    Too bigs: 0, No such names: 0, Bad values: 0,
    Read onlys: 0, General errors: 0,
```

```
Total request varbinds: 0, Total set varbinds: 0,  
Get requests: 0, Get nexts: 0, Set requests: 0,  
Get responses: 0, Traps: 0,  
Silent drops: 0, Proxy drops 0  
Output:  
Packets: 0, Too bigs: 0, No such names: 0,  
Bad values: 0, General errors: 0,  
Get requests: 0, Get nexts: 0, Set requests: 0,  
Get responses: 0, Traps: 0
```

request snmp spoof-trap

Syntax	request snmp spoof-trap <trap> variable-bindings <object> <instance> <value>
Release Information	Command introduced in Junos OS Release 8.2. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Spoof (mimic) the behavior of a Simple Network Management Protocol (SNMP) trap.
Options	<p><trap>—Name of the trap to spoof.</p> <p>variable-bindings <object> <instance> <value>—(Optional) List of variables and values to include in the trap. Each variable binding is specified as an object name, the object instance, and the value (for example, ifIndex[14] = 14). Enclose the list of variable bindings in quotation marks (" ") and use a comma to separate each object name, instance, and value definition (for example, variable-bindings "ifIndex[14] = 14, ifAdminStatus[14] = 1, ifOperStatus[14] = 2"). Objects included in the trap definition that do not have instances and values specified as part of the command are included in the trap and spoofed with automatically generated instances and values.</p> <p><dummy name>—A dummy trap name to display the list of available traps.</p> <p>Question mark (?)—Question mark? to display possible completions.</p>
Required Privilege Level	request
List of Sample Output	request snmp spoof-trap (with Variable Bindings) on page 1741 request snmp spoof-trap (Illegal Trap Name) on page 1741 request snmp spoof-trap (Question Mark ?) on page 1745

Sample Output

request snmp spoof-trap (with Variable Bindings)

```
user@host> request snmp spoof-trap linkUp variable-bindings "ifIndex[14] = 14, ifAdminStatus[14] = 1, ifOperStatus[14] = 2"
Spoof trap request result: trap sent successfully
```

request snmp spoof-trap (Illegal Trap Name)

```
user@host> request snmp spoof-trap xx
Spoof trap request result: trap not found
```

```
Allowed Traps:
ads1AtucInitFailureTrap
ads1AtucPerfESsThreshTrap
ads1AtucPerfLofsThreshTrap
ads1AtucPerfLolsThreshTrap
ads1AtucPerfLossThreshTrap
ads1AtucPerfLprsThreshTrap
ads1AtucRateChangeTrap
```

ads1AturPerfESsThreshTrap
ads1AturPerfLofsThreshTrap
ads1AturPerfLossThreshTrap
ads1AturPerfLprsThreshTrap
ads1AturRateChangeTrap
apsEventChannelMismatch
apsEventFEPLF
apsEventModeMismatch
apsEventPSBF
apsEventSwitchover
authenticationFailure
bfdSessDown
bfdSessUp
bgpBackwardTransition
bgpEstablished
coldStart
d1swTrapCircuitDown
d1swTrapCircuitUp
d1swTrapTConnDown
d1swTrapTConnPartnerReject
d1swTrapTConnProtViolation
d1swTrapTConnUp
dsx1LineStatusChange
dsx3LineStatusChange
entConfigChange
fallingAlarm
frDLCIStatusChange
ggsnTrapChanged
ggsnTrapCleared
ggsnTrapNew
gmp1sTunnelDown
ifMauJabberTrap
ipv6IfStateChange
isisAreaMismatch
isisAttemptToExceedMaxSequence
isisAuthenticationFailure
isisAuthenticationTypeFailure
isisCorruptedLSPDetected
isisDatabaseOverload
isisIDLenMismatch
isisLSPTooLargeToPropagate
isisManualAddressDrops
isisMaxAreaAddressesMismatch
isisOriginatingLSPBufferSizeMismatch
isisOwnLSPPurge
isisProtocolsSupportedMismatch
isisRejectedAdjacency
isisSequenceNumberSkip
isisVersionSkew
jnxAccessAuthServerDisabled
jnxAccessAuthServerEnabled
jnxAccessAuthServiceDown
jnxAccessAuthServiceUp
jnxBfdSessDetectionTimeHigh
jnxBfdSessTxIntervalHigh
jnxBgpM2BackwardTransition
jnxBgpM2Established
jnxCmCfgChange
jnxCmRescueChange
jnxCollFlowOverload
jnxCollFlowOverloadCleared

jnxCollFtpSwitchover
jnxCollMemoryAvailable
jnxCollMemoryUnavailable
jnxCollUnavailableDest
jnxCollUnavailableDestCleared
jnxCollUnsuccessfulTransfer
jnxDfcHardMemThresholdExceeded
jnxDfcHardMemUnderThreshold
jnxDfcHardPpsThresholdExceeded
jnxDfcHardPpsUnderThreshold
jnxDfcSoftMemThresholdExceeded
jnxDfcSoftMemUnderThreshold
jnxDfcSoftPpsThresholdExceeded
jnxDfcSoftPpsUnderThreshold
jnxEventTrap
jnxExampleStartup
jnxFEBSwitchover
jnxFanFailure
jnxFanOK
jnxFruCheck
jnxFruFailed
jnxFruInsertion
jnxFruOK
jnxFruOffline
jnxFruOnline
jnxFruPowerOff
jnxFruPowerOn
jnxFruRemoval
jnxHardDiskFailed
jnxHardDiskMissing
jnxJsAvPatternUpdateTrap
jnxJsChassisClusterSwitchover
jnxJsFwAuthCapacityExceeded
jnxJsFwAuthFailure
jnxJsFwAuthServiceDown
jnxJsFwAuthServiceUp
jnxJsNatAddrPoolThresholdStatus
jnxJsScreenAttack
jnxJsScreenCfgChange
jnxLdpLspDown
jnxLdpLspUp
jnxLdpSesDown
jnxLdpSesUp
jnxMIMstCistPortLoopProtectStateChangeTrap
jnxMIMstCistPortRootProtectStateChangeTrap
jnxMIMstErrTrap
jnxMIMstGenTrap
jnxMIMstInvalidBpduRxdTrap
jnxMIMstMstiPortLoopProtectStateChangeTrap
jnxMIMstMstiPortRootProtectStateChangeTrap
jnxMIMstNewRootTrap
jnxMIMstProtocolMigrationTrap
jnxMIMstRegionConfigChangeTrap
jnxMIMstTopologyChgTrap
jnxMacChangedNotification
jnxMplsLdpInitSesThresholdExceeded
jnxMplsLdpPathVectorLimitMismatch
jnxMplsLdpSessionDown
jnxMplsLdpSessionUp
jnxOspfV3IfConfigError
jnxOspfV3IfRxBadPacket

jnxOspfV3IfStateChange
jnxOspfV3LsdbApproachingOverflow
jnxOspfV3LsdbOverflow
jnxOspfV3NbrRestartHelperStatusChange
jnxOspfV3NbrStateChange
jnxOspfV3NssaTranslatorStatusChange
jnxOspfV3RestartStatusChange
jnxOspfV3VirtIfConfigError
jnxOspfV3VirtIfRxBadPacket
jnxOspfV3VirtIfStateChange
jnxOspfV3VirtNbrRestartHelperStatusChange
jnxOspfV3VirtNbrStateChange
jnxOtnAlarmCleared
jnxOtnAlarmSet
jnxOverTemperature
jnxPmonOverloadCleared
jnxPmonOverloadSet
jnxPingEgressJitterThresholdExceeded
jnxPingEgressStdDevThresholdExceeded
jnxPingEgressThresholdExceeded
jnxPingIngressJitterThresholdExceeded
jnxPingIngressStdDevThresholdExceeded
jnxPingIngressThresholdExceeded
jnxPingRttJitterThresholdExceeded
jnxPingRttStdDevThresholdExceeded
jnxPingRttThresholdExceeded
jnxPortBpduErrorStatusChangeTrap
jnxPortLoopProtectStateChangeTrap
jnxPortRootProtectStateChangeTrap
jnxPowerSupplyFailure
jnxPowerSupplyOK
jnxRedundancySwitchover
jnxRmonAlarmGetFailure
jnxRmonGetOk
jnxSecAccessIfMacLimitExceeded
jnxSecAccessSdsRateLimitCrossed
jnxSonetAlarmCleared
jnxSonetAlarmSet
jnxSpSvcSetCpuExceeded
jnxSpSvcSetCpuOk
jnxSpSvcSetZoneEntered
jnxSpSvcSetZoneExited
jnxStormEventNotification
jnxSyslogTrap
jnxTemperatureOK
jnxVccpPortDown
jnxVccpPortUp
jnxVpnIfDown
jnxVpnIfUp
jnxVpnPwDown
jnxVpnPwUp
jnx12aldGlobalMacLimit
jnx12aldInterfaceMacLimit
jnx12aldRoutingInstMacLimit
linkDown
linkUp
lldpRemTablesChange
mfrMibTrapBundleLinkMismatch
mplsLspChange
mplsLspDown
mplsLspInfoChange


```

mplsLspInfoDown
mplsLspInfoPathDown
mplsLspInfoPathUp
mplsLspInfoUp
mplsLspPathDown
mplsLspPathUp
mplsLspUp
mplsNumVrfRouteMaxThreshExceeded
mplsNumVrfRouteMidThreshExceeded
mplsNumVrfSecIllglLb1ThrshExcd
mplsTunnelDown
mplsTunnelReoptimized
mplsTunnelRerouted
mplsTunnelUp
mplsVrfIfDown
mplsVrfIfUp
mplsXCDown
mplsXCUp
msdpBackwardTransition
msdpEstablished
newRoot
ospfIfAuthFailure
ospfIfConfigError
ospfIfRxBadPacket
ospfIfStateChange
ospfLsdbApproachingOverflow
ospfLsdbOverflow
ospfMaxAgeLsa
ospfNbrStateChange
ospfOriginateLsa
ospfTxRetransmit
ospfVirtIfAuthFailure
ospfVirtIfConfigError
ospfVirtIfRxBadPacket
ospfVirtIfStateChange
ospfVirtIfTxRetransmit
ospfVirtNbrStateChange
pethMainPowerUsageOffNotification
pethMainPowerUsageOnNotification
pethPsePortOnOffNotification
pingProbeFailed
pingTestCompleted
pingTestFailed
ptopoConfigChange
risingAlarm
rpMauJabberTrap
sd1cLSStatusChange
sd1cPortStatusChange
topologyChange
traceRoutePathChange
traceRouteTestCompleted
traceRouteTestFailed
vrrpTrapAuthFailure
vrrpTrapNewMaster
warmStart

```

request snmp spoof-trap (Question Mark ?)

```

user@host> request snmp spoof-trap ?
Possible completions:
<trap>           The name of the trap to spoof

```

```
ads1AtucInitFailureTrap
ads1AtucPerfESsThreshTrap
ads1AtucPerfLofsThreshTrap
ads1AtucPerfLoIsThreshTrap
ads1AtucPerfLossThreshTrap
ads1AtucPerfLprsThreshTrap
ads1AtucRateChangeTrap
ads1AturPerfESsThreshTrap
ads1AturPerfLofsThreshTrap
ads1AturPerfLossThreshTrap
ads1AturPerfLprsThreshTrap
ads1AturRateChangeTrap
apsEventChannelMismatch
apsEventFEPLF
apsEventModeMismatch
apsEventPSBF
apsEventSwitchover
authenticationFailure
bfdSessDown
bfdSessUp
bgpBackwardTransition
bgpEstablished
coldStart
dlswTrapCircuitDown
dlswTrapCircuitUp
---(more 10%)---
```

request snmp utility-mib clear instance

Syntax	request snmp utility-mib clear instance <i>name</i> object-type <i>type</i>
Release Information	Command introduced in Junos OS Release 12.2 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Clear the data stored in the specified container object in the SNMP Utility MIB.
Options	<p><i>name</i>—Name of the SNMP instance that is used to identify the data stored in the container object.</p> <p><i>object-type type</i>—Type of container object in which the data is stored. The following container object types are supported:</p> <ul style="list-style-type: none">• counter—Stores a 32-bit counter value.• counter64—Stores a 64-bit counter value.• integer—Stores a 32-bit signed integer value.• unsigned-integer—Stores a 32-bit unsigned integer value.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• Utility MIB on page 1460• <i>Understanding the Implementation of SNMP on the QFabric System</i>• request snmp utility-mib set instance on page 1748

request snmp utility-mib set instance

Syntax	<code>request snmp utility-mib set instance <i>name</i></code> <code>object-type <i>type</i></code> <code>object-value <i>value</i></code>
Release Information	Command introduced in Junos OS Release 12.2 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Store data in the specified container object in the SNMP Utility MIB. The data may be retrieved by SNMP operations.
Options	<p><i>name</i>—Name of the SNMP instance that is used to identify the data stored in the container object.</p> <p><i>object-type type</i>—Type of container object in which to store data. The following container object types are supported:</p> <ul style="list-style-type: none">• counter—Stores a 32-bit counter value.• counter64—Stores a 64-bit counter value.• integer—Stores a 32-bit signed integer value.• unsigned-integer—Stores a 32-bit unsigned integer value.• string—Stores an octet string value. <p><i>object-value value</i>—Data that is stored in the container object.</p>
Required Privilege Level	request
Related Documentation	<ul style="list-style-type: none">• Utility MIB on page 1460• <i>Understanding the Implementation of SNMP on the QFabric System</i>• request snmp utility-mib clear instance on page 1747

show snmp health-monitor

Syntax	show snmp health-monitor <alarms (brief detail) logs>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Display information about Simple Network Management Protocol (SNMP) health monitor alarms and logs.
Options	<p>none—Display information about all health monitor alarms and logs.</p> <p>alarms (brief detail)—(Optional) Display information about health monitor alarms. Optionally, specify brief or detailed information about the alarms.</p> <p>logs—(Optional) Display information about health monitor logs.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Understanding Health Monitoring on page 1467 • Configuring Health Monitoring on page 1509
List of Sample Output	show snmp health-monitor on page 1751 show snmp health-monitor alarms detail on page 1751
Output Fields	Table 126 describes the output fields for the show snmp health-monitor command. Output fields are listed in the approximate order in which they appear.

Table 126: show snmp health-monitor Output Fields

Field Name	Field Description	Level of Output
Alarm Index	Alarm identifier.	All levels
Variable description	Description of the health monitor object instance being monitored.	All levels
Variable name	Name of the health monitor object instance being monitored.	All levels
Value	Current value of the monitored variable in the most recent sample interval.	All levels

Table 126: show snmp health-monitor Output Fields (*continued*)

Field Name	Field Description	Level of Output
State	<p>State of the alarm or event entry:</p> <ul style="list-style-type: none"> Alarms: <ul style="list-style-type: none"> active—Entry is fully configured and activated. falling threshold crossed—Value of the variable has crossed the lower threshold limit. rising threshold crossed—Value of the variable has crossed the upper threshold limit. under creation—Entry is being configured and is not yet activated. startup—Alarm is waiting for the first sample of the monitored variable. object not available—Monitored variable of that type is not available to the health monitor agent. instance not available—Monitored variable's instance is not available to the health monitor agent. object type invalid—Monitored variable is not a numeric value. object processing errored—An error occurred when the monitored variable was processed. unknown—State is not one of the above. 	All levels
Variable OID	Object ID to which the variable name is resolved. The format is x.x.x.x.	detail
Sample type	Method of sampling the monitored variable and calculating the value to compare against the upper and lower thresholds. It can have the value <i>absolute value</i> or <i>delta value</i> .	detail
Startup alarm	<p>Alarm that might be sent when this entry is first activated, depending on the following criteria:</p> <ul style="list-style-type: none"> Alarm is sent when one of the following situations exists: <ul style="list-style-type: none"> Value of the alarm is above or equal to the rising threshold and the startup type is either rising alarm or rising or falling alarm. <i>falling alarm</i> Value of the alarm is below or equal to the falling threshold and the startup type is either <i>falling alarm</i> or <i>rising or falling alarm</i>. Alarm is <i>not</i> sent when one of the following situations exists: <ul style="list-style-type: none"> Value of the alarm is above or equal to the rising threshold and the startup type is <i>falling alarm</i>. Value of the alarm is below or equal to the falling threshold and the startup type is <i>rising alarm</i>. Value of the alarm is between the thresholds. 	detail
Owner	Name of the entry configured by the user. If the entry was created through the CLI, the owner has monitor prepended to it.	detail
Creator	Mechanism by which the entry was configured (Health Monitor).	detail
Sample interval	Time period between samples (in seconds).	detail
Rising threshold	Upper limit threshold value as a percentage of the maximum possible value.	detail

Table 126: show snmp health-monitor Output Fields (*continued*)

Field Name	Field Description	Level of Output
Falling threshold	Lower limit threshold value as a percentage of the maximum possible value.	detail
Rising event index	Index number of the event triggered when the rising threshold is crossed.	detail
Falling event index	Index number of the event triggered when the falling threshold is crossed. Details include the value of the falling event instance and the state of the falling event instance.	detail

Sample Output

show snmp health-monitor

```
user@switch> show snmp health-monitor
```

```
Alarm
Index  Variable description                                Value State

32768 Health Monitor: root file system utilization
      jnxHrStoragePercentUsed.1                          59 active

32769 Health Monitor: /config file system utilization
      jnxHrStoragePercentUsed.2                          0 active

32770 Health Monitor: RE 0 CPU utilization
      jnxOperatingCPU.9.1.0.0                            9 falling threshold

32772 Health Monitor: RE 0 memory utilization
      jnxOperatingBuffer.9.1.0.0                        23 active

32774 Health Monitor: Max Kernel Memory Used (%)
      jnxBoxKernelMemoryUsedPercent.0                    3 active
Event Index: 32768
Description: Health Monitor: RE 0 CPU utilization crossed falling threshold
70 (value: 5), (variable: jnxOperatingCPU.9.1.0.0)
Time: 2011-01-09 19:18:35 PST
```

show snmp health-monitor alarms detail

```
user@switch> show snmp health-monitor alarms detail
```

```
Alarm Index 32768:
Variable name      jnxHrStoragePercentUsed.1
Variable OID       1.3.6.1.4.1.2636.3.31.1.1.1.1.1
Sample type        absolute value
Startup alarm      rising alarm
Owner              Health Monitor: root file system
                  utilization
Creator            Health Monitor
State              active
Sample interval    300 seconds
Rising threshold   80
```

Falling threshold 70
Rising event index 32768
Falling event index 32768
Instance Value: 59
Instance State: active

Alarm Index 32769:

Variable name jnxHrStoragePercentUsed.2
Variable OID 1.3.6.1.4.1.2636.3.31.1.1.1.2
Sample type absolute value
Startup alarm rising alarm
Owner Health Monitor: /config file system utilization
Creator Health Monitor
State active
Sample interval 300 seconds
Rising threshold 80
Falling threshold 70
Rising event index 32768
Falling event index 32768
Instance Value: 0
Instance State: active

Alarm Index 32770:

Variable name jnxOperatingCPU.9.1.0.0
Variable OID 1.3.6.1.4.1.2636.3.1.13.1.8.9.1.0.0
Sample type absolute value
Startup alarm rising alarm
Owner Health Monitor: RE 0 CPU utilization
Creator Health Monitor
State active
Sample interval 300 seconds
Rising threshold 80
Falling threshold 70
Rising event index 32768
Falling event index 32768
Instance Value: 9
Instance State: falling threshold

Alarm Index 32772:

Variable name jnxOperatingBuffer.9.1.0.0
Variable OID 1.3.6.1.4.1.2636.3.1.13.1.11.9.1.0.0
Sample type absolute value
Startup alarm rising alarm
Owner Health Monitor: RE 0 memory utilization
Creator Health Monitor
State active
Sample interval 300 seconds
Rising threshold 80
Falling threshold 70
Rising event index 32768
Falling event index 32768
Instance Value: 23
Instance State: active

Alarm Index 32774:

Variable name jnxBoxKernelMemoryUsedPercent.0
Variable OID 1.3.6.1.4.1.2636.3.1.16.0
Sample type absolute value

Startup alarm	rising alarm
Owner	Health Monitor: Max Kernel Memory Used (%)
Creator	Health Monitor
State	active
Sample interval	300 seconds
Rising threshold	80
Falling threshold	70
Rising event index	32768
Falling event index	32768
Instance Value: 3	
Instance State: active	

show snmp inform-statistics

Syntax	show snmp inform-statistics
Release Information	<p>Command introduced in Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Display information about Simple Network Management Protocol (SNMP) inform requests.
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show snmp inform-statistics on page 1754
Output Fields	<p>Table 127 describes the output fields for the show snmp inform-statistics command.</p> <p>Output fields are listed in the approximate order in which they appear.</p>

Table 127: show snmp inform-statistics Output Fields

Field Name	Field Description
Target Name	Name of the device configured to receive and respond to SNMP informs.
Address	IP address of the target device.
Sent	Number of informs sent to the target device and acknowledged by the target device.
Pending	Number of informs held in memory pending a response from the target device.
Discarded	Number of informs discarded after the specified number of retransmissions to the target device were attempted.
Timeouts	Number of informs that did not receive an acknowledgement from the target device within the timeout specified.
Probe Failures	Connection failures that occurred (for example, when the target server returned invalid content or you incorrectly configured the target address).

Sample Output

show snmp inform-statistics

```

user@host> show snmp inform-statistics
Inform Request Statistics:
Target Name: TA1_v3_md5_none Address: 172.17.20.184
Sent: 176, Pending: 0
Discarded: 0, Timeouts: 0, Probe Failures: 0

```

Target Name: TA2_v3_sha_none Address: 192.168.110.59
Sent: 0, Pending: 4
Discarded: 84, Timeouts: 0, Probe Failures: 258
Target Name: TA5_v2_none Address: 172.17.20.184
Sent: 0, Pending: 0
Discarded: 2, Timeouts: 10, Probe Failures: 0

show snmp mib

Syntax	<code>show snmp mib (get get-next walk) (ascii decimal) <i>object-id</i></code>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>ascii and decimal options introduced in Junos OS Release 9.6.</p> <p>ascii and decimal options introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Display local Simple Network Management Protocol (SNMP) Management Information Base (MIB) object values.
Options	<p>get—Retrieve and display one or more SNMP object values.</p> <p>get-next—Retrieve and display the next SNMP object values.</p> <p>walk—Retrieve and display the SNMP object values that are associated with the requested object identifier (OID). When you use this option, the Junos OS displays the objects below the subtree that you specify.</p> <p>ascii—Display the SNMP object's string indices as an ASCII-key representation.</p> <p>decimal—Display the SNMP object values in the decimal (default) format. The decimal option is the default option for this command. Therefore, issuing the show snmp mib (get get-next walk) decimal object-id and the show snmp mib (get get-next walk) object-id commands display the same output.</p> <p>object-id—The object can be represented by a sequence of dotted integers (such as 1.3.6.1.2.1.2) or by its subtree name (such as interfaces). When entering multiple objects, enclose the objects in quotation marks.</p>
Required Privilege Level	snmp—To view this statement in the configuration.
List of Sample Output	<p>show snmp mib get on page 1757</p> <p>show snmp mib get (Multiple Objects) on page 1757</p> <p>show snmp mib get (Layer 2 Policer) on page 1757</p> <p>show snmp mib get-next on page 1757</p> <p>show snmp mib get-next (Specify an OID) on page 1757</p> <p>show snmp mib walk on page 1757</p> <p>show snmp mib walk (QFX Series) on page 1757</p> <p>show snmp mib walk decimal on page 1758</p> <p>show snmp mib walk (ASCII) on page 1758</p> <p>show snmp mib walk (Multiple Indices) on page 1758</p> <p>show snmp mib walk decimal (Multiple Indices) on page 1758</p>
Output Fields	Table 128 describes the output fields for the show snmp mib command. Output fields are listed in the approximate order in which they appear.

Table 128: show snmp mib Output Fields

Field Name	Field Description
<i>name</i>	Object name and numeric instance value.
<i>object value</i>	Object value. The Junos OS translates OIDs into the corresponding object names.

Sample Output

show snmp mib get

```
user@host> show snmp mib get sysObjectID.0
sysObjectID.0 = jnxProductNameM20
```

show snmp mib get (Multiple Objects)

```
user@host> show snmp mib get ?sysObjectID.0 sysUpTime.0?
sysObjectID.0 = jnxProductNameM20
sysUpTime.0 = 1640992
```

show snmp mib get (Layer 2 Policer)

```
user@host> show snmp mib get ifInOctets.25970
ifInOctets.25970 = 7545720
```

show snmp mib get-next

```
user@host> show snmp mib get-next jnxMibs
jnxBoxClass.0 = jnxProductLineM20.0
```

show snmp mib get-next (Specify an OID)

```
user@host> show snmp mib get-next 1.3.6.1
sysDescr.0 = Juniper Networks, Inc. m20 internet router, kernel
Junos OS Release: 2004-1 Build date: build date UTC Copyright (c) 1996-2004 Juniper
Networks, Inc.
```

show snmp mib walk

```
user@host> show snmp mib walk system
sysDescr.0 = Juniper Networks, Inc. m20 internet router, kernel
Junos OS Release #0: 2004-1 Build date: build date UTC Copyright (c) 1996-2004
Juniper Networks, Inc.
sysObjectID.0 = jnxProductNameM20
sysUpTime.0 = 1640992
sysContact.0 = Your contact
sysName.0 = my router
sysLocation.0 = building 1
sysServices.0 = 4
```

show snmp mib walk (QFX Series)

```
user@switch> show snmp mib walk system
sysDescr.0 = Juniper Networks, Inc. qfx3500s internet router, kernel JUNOS
11.1-20100926.0 #0: 2010-09-26 06:17:38 UTC Build date: 2010-09-26 06:00:10
sysObjectID.0 = jnxProductQFX3500
sysUpTime.0 = 138980301
sysContact.0 = System Contact
```

```
sysName.0      = LabQFX3500
sysLocation.0  = Lab
sysServices.0  = 4
```

show snmp mib walk decimal

```
user@host show snmp mib walk decimal jnxUtilData
jnxUtilCounter32Value.102.114.101.100 = 100
```

show snmp mib walk (ASCII)

```
show snmp mib walk ascii jnxUtilData
jnxUtilCounter32Value."fred" = 100
```

show snmp mib walk (Multiple Indices)

```
show snmp mib walk ascii jnxFWCounterByteCount
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_BE-fe-1/3/0.0-i".2 = 0
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_CC-fe-1/3/0.0-i".2 = 0
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_RT-fe-1/3/0.0-i".2 = 0
.....
```

show snmp mib walk decimal (Multiple Indices)

```
show snmp mib walk ascii jnxFWCounterByteCount
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_BE-fe-1/3/0.0-i".2 = 0
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_CC-fe-1/3/0.0-i".2 = 0
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_RT-fe-1/3/0.0-i".2 = 0
.....
```

show snmp rmon

Syntax	show snmp rmon <alarms (brief detail)> <events (brief detail)> <logs>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Display information about Simple Network Management Protocol (SNMP) Remote Monitoring (RMON) alarms, events, and logs.
Options	<p>none—Display information about all RMON alarms and events.</p> <p>brief detail—(Optional) Display brief or detailed information about RMON alarms or events.</p> <p>alarms—(Optional) Display information about RMON alarms.</p> <p>events—(Optional) Display information about RMON events.</p> <p>logs—(Optional) Display information about RMON monitoring logs.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • RMON MIB Event, Alarm, Log, and History Control Tables on page 1464 • Monitoring RMON MIB Tables on page 1516 • Configuring RMON Alarms and Events on page 1507 • Understanding RMON on page 1463 • clear snmp statistics on page 1739 • clear snmp history on page 1738 • show snmp rmon history on page 1763
List of Sample Output	show snmp rmon on page 1761 show snmp rmon alarms detail on page 1762 show snmp rmon events detail on page 1762 show snmp rmon logs on page 1762
Output Fields	Table 129 describes the output fields for the show snmp rmon command. Output fields are listed in the approximate order in which they appear.

Table 129: show snmp rmon Output Fields

Field Name	Field Description	Level of Output
Alarm Index	Alarm identifier.	All levels

Table 129: show snmp rmon Output Fields (*continued*)

Field Name	Field Description	Level of Output
State	<p>State of the alarm or event entry:</p> <p>Alarms:</p> <ul style="list-style-type: none"> • active—Entry is fully configured and activated. • falling threshold crossed—Value of the variable has crossed the lower threshold limit. • rising threshold crossed—Value of the variable has crossed the upper threshold limit. • under creation—Entry is being configured and is not yet activated. • startup—Alarm is waiting for the first sample of the monitored variable. • object not available—Monitored variable of that type is not available to the SNMP agent. • instance not available—Monitored variable's instance is not available to the SNMP agent. • object type invalid—Monitored variable is not a numeric value. • object processing errored—An error occurred when the monitored variable was processed. • unknown—State is not one of the above. <p>Events:</p> <ul style="list-style-type: none"> • active—Entry has been fully configured and activated. • under creation—Entry is being configured and is not yet activated. • unknown—State is not one of the above. 	All levels
Variable name	Name of the SNMP object instance being monitored.	All levels
Event Index	Event identifier.	All levels
Type	<p>Type of notification made when an event is triggered. It can be one of the following:</p> <ul style="list-style-type: none"> • log—A system log message is generated and an entry is made to the log table. • snmptrap—An SNMP trap is sent to the configured destination. • log and trap—A system log message is generated, an entry is made to the log table, and an SNMP trap is sent to the configured destination. • none—Neither log nor trap will be sent. 	detail
Last Event	Date and time of the last event. It has the format <i>yyyy-mm-dd hh:mm:ss timezone</i> .	brief
Community	Trap group used for sending the SNMP trap.	detail
Variable OID	Object ID to which the variable name is resolved. The format is x.x.x.x.	detail
Sample type	Method of sampling the monitored variable and calculating the value to compare against the upper and lower thresholds. It can have the value of absolute value or delta value .	detail

Table 129: show snmp rmon Output Fields (*continued*)

Field Name	Field Description	Level of Output
Startup alarm	Alarm that might be sent when this entry is first activated, depending on the following criteria: <ul style="list-style-type: none"> Alarm is sent when one of the following situations exists: <ul style="list-style-type: none"> Value of the alarm is above or equal to the rising threshold and the startup type is either rising alarm or rising or falling alarm. Value of the alarm is below or equal to the falling threshold and the startup type is either falling alarm or rising or falling alarm. Alarm is <i>not</i> sent when one of the following situations exists: <ul style="list-style-type: none"> Value of the alarm is above or equal to the rising threshold and the startup type is falling alarm. Value of the alarm is below or equal to the falling threshold and the startup type is rising alarm. Value of the alarm is between the thresholds. 	detail
Owner	Name of the entry configured by the user. If the entry was created through the CLI, the owner has monitor prepended to it.	detail
Creator	Mechanism by which the entry was configured (CLI or SNMP).	detail
Sample interval	Time period between samples (in seconds).	detail
Rising threshold	Upper limit threshold value configured by the user.	detail
Falling threshold	Lower limit threshold value configured by the user.	detail
Rising event index	Event triggered when the rising threshold is crossed.	detail
Falling event index	Event triggered when the falling threshold is crossed.	detail
Current value	Current value of the monitored variable in the most recent sample interval.	detail

Sample Output

show snmp rmon

```

user@host> show snmp rmon
Alarm
Index  Variable description                               Value State

      5 monitor
      jnxOperatingCPU.9.1.0.0                         5 falling threshold

Event
Index  Type                               Last Event
      1 log and trap                     2009-07-10 11:34:17 PDT
Event Index: 1
      Description: Event 1 triggered by Alarm 5, rising threshold (90) crossed,
      (variable: jnxOperatingCPU.9.1.0.0, value: 100)
      Time: 2009-07-10 11:34:07 PDT

```

Description: Event 1 triggered by Alarm 5, falling threshold (75) crossed,
(variable: jnxOperatingCPU.9.1.0.0, value: 5)
Time: 2009-07-10 11:34:17 PDT

show snmp rmon alarms detail

```
user@host> show snmp rmon alarms detail
Alarm Index 5:
  Variable name           jnxOperatingCPU.9.1.0.0
  Variable OID            1.3.6.1.4.1.2636.3.1.13.1.8.9.1.0.0
  Sample type             absolute value
  Startup alarm           rising or falling alarm
  Owner                   monitor

  Creator                 CLI
  State                   active
  Sample interval         5 seconds
  Rising threshold        90
  Falling threshold       75
  Rising event index      1
  Falling event index     1
  Instance Value: 4
  Instance State: falling threshold
```

show snmp rmon events detail

```
user@host> show snmp rmon events detail
Event Index 1:
  Description             rmon event
  Type                    log and trap
  Community               rmon-trap-group
  Last event              2009-07-10 11:34:17 PDT
  Creator                 CLI
  State                   active
```

show snmp rmon logs

```
user@host> show snmp rmon logs
Event Index: 1
  Description: Event 1 triggered by Alarm 5, rising threshold (90) crossed,
(variable: jnxOperatingCPU.9.1.0.0, value: 100)
  Time: 2009-07-10 11:34:07 PDT
  Description: Event 1 triggered by Alarm 5, falling threshold (75) crossed,
(variable: jnxOperatingCPU.9.1.0.0, value: 5)
  Time: 2009-07-10 11:34:17 PDT
```

show snmp rmon history

Syntax	show snmp rmon history <history-index> sample-index <sample-index>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Display the contents of the RMON history group.
Options	<p>none—Display all the entries in the RMON history group.</p> <p>history-index—(Optional) Display the contents of the specified entry in the RMON history group.</p> <p>sample-index sample-index—(Optional) Display the statistics collected for the specified sample within the specified entry in the RMON history group.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • RMON MIB Event, Alarm, Log, and History Control Tables on page 1464 • Monitoring RMON MIB Tables on page 1516 • Configuring RMON Alarms and Events on page 1507 • Understanding RMON on page 1463 • clear snmp statistics on page 1739 • clear snmp history on page 1738 • show snmp rmon on page 1759

show snmp statistics

Syntax	<code>show snmp statistics</code> <code><subagents></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Option subagents introduced in Junos OS Release 14.2.
Description	Display statistics about Simple Network Management Protocol (SNMP) packets sent and received by the router or switch.
Options	subagents —(Optional) Display the statistics of the protocol data unit (PDU), the number of SNMP requests and responses per subagent, and the SNMP statistics received from each subagent per logical system.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• clear snmp statistics on page 1739
List of Sample Output	show snmp statistics on page 1769 show snmp statistics subagents on page 1769
Output Fields	Table 130 describes the output fields for the show snmp statistics command. Output fields are listed in the approximate order in which they appear.

Table 130: show snmp statistics Output Fields

Field Name	Field Description
Input	<p>Information about received packets:</p> <ul style="list-style-type: none"> • Packets(snmplnPkts)—Total number of messages delivered to the SNMP entity from the transport service. • Bad versions—(snmplnBadVersions) Total number of messages delivered to the SNMP entity that were for an unsupported SNMP version. • Bad community names—(snmplnBadCommunityNames) Total number of messages delivered to the SNMP entity that used an SNMP community name not known to the entity. • Bad community uses—(snmplnBadCommunityUses) Total number of messages delivered to the SNMP entity that represented an SNMP operation that was not allowed by the SNMP community named in the message. • ASN parse errors—(snmplnASNParseErrs) Total number of ASN.1 or BER errors encountered by the SNMP entity when decoding received SNMP messages. • Too bigs—(snmplnTooBigs) Total number of SNMP PDUs delivered to the SNMP entity with an error status field of tooBig. • No such names—(snmplnNoSuchNames) Total number of SNMP PDUs delivered to the SNMP entity with an error status field of noSuchName. • Bad values—(snmplnBadValues) Total number of SNMP PDUs delivered to the SNMP entity with an error status field of badValue. • Read onlys—(snmplnReadOnlys) Total number of valid SNMP PDUs delivered to the SNMP entity with an error status field of readOnly. Only incorrect implementations of SNMP generate this error.

Table 130: show snmp statistics Output Fields (*continued*)

Field Name	Field Description
Input (continued)	<ul style="list-style-type: none"> • General errors—(snmpInGenErrs) Total number of SNMP PDUs delivered to the SNMP entity with an error status field of genErr. • Total requests varbinds—(snmpInTotalReqVars) Total number of MIB objects retrieved successfully by the SNMP entity as a result of receiving valid SNMP GetRequest and GetNext PDUs. • Total set varbinds—(snmpInSetVars) Total number of MIB objects modified successfully by the SNMP entity as a result of receiving valid SNMP SetRequest PDUs. • Get requests—(snmpInGetRequests) Total number of SNMP GetRequest PDUs that have been accepted and processed by the SNMP entity. • Get nexts—(snmpInGetNexts) Total number of SNMP GetNext PDUs that have been accepted and processed by the SNMP entity. • Set requests—(snmpInSetRequests) Total number of SNMP SetRequest PDUs that have been accepted and processed by the SNMP entity. • Get responses—(snmpInGetResponses) Total number of SNMP GetResponse PDUs that have been accepted and processed by the SNMP entity. • Traps—(snmpInTraps) Total number of SNMP traps generated by the SNMP entity. • Silent drops—(snmpSilentDrops) Total number of GetRequest, GetNextRequest, GetBulkRequest, SetRequests, and InformRequest PDUs delivered to the SNMP entity that were silently dropped because the size of a reply containing an alternate response PDU with an empty variable-bindings field was greater than either a local constraint or the maximum message size associated with the originator of the requests. • Proxy drops—(snmpProxyDrops) Total number of GetRequest, GetNextRequest, GetBulkRequest, SetRequests, and InformRequest PDUs delivered to the SNMP entity that were silently dropped because the transmission of the message to a proxy target failed in such a way (other than a timeout) that no response PDU could be returned. • Commit pending drops—Number of SNMP packets for Set requests dropped because of a previous pending SNMP Set request on the committed configuration. • Throttle drops—Number of SNMP packets for any requests dropped reaching the throttle limit.

Table 130: show snmp statistics Output Fields (*continued*)

Field Name	Field Description
V3 Input	<p>Information about SNMP version 3 packets:</p> <ul style="list-style-type: none"> • Unknown security models—(snmpUnknownSecurityModels) Total number of packets received by the SNMP engine that were dropped because they referenced a security model that was not known to or supported by the SNMP engine. • Invalid messages—(snmpInvalidMsgs) Number of packets received by the SNMP engine that were dropped because there were invalid or inconsistent components in the SNMP message. • Unknown pdu handlers—(snmpUnknownPDUHandlers) Number of packets received by the SNMP engine that were dropped because the PDU contained in the packet could not be passed to an application responsible for handling the PDU type. • Unavailable contexts—(snmpUnavailableContexts) Number of requests received for a context that is known to the SNMP engine, but is currently unavailable. • Unknown contexts—(snmpUnknownContexts) Total number of requests received for a context that is unknown to the SNMP engine. • Unsupported security levels—(usmStatsUnsupportedSecLevels) Total number of packets received by the SNMP engine that were dropped because they requested a security level unknown to the SNMP engine (or otherwise unavailable). • Not in time windows—(usmStatsNotInTimeWindows) Total number of packets received by the SNMP engine that were dropped because they appeared outside the authoritative SNMP engine's window. • Unknown user names—(usmStatsUnknownUserNames) Total number of packets received by the SNMP engine that were dropped because they referenced a user that was not known to the SNMP engine. • Unknown engine ids—(usmStatsUnknownEngineIDs) Total number of packets received by the SNMP engine that were dropped because they referenced an SNMP engine ID that was not known to the SNMP engine. • Wrong digests—(usmStatsWrongDigests) Total number of packets received by the SNMP engine that were dropped because they did not contain the expected digest value. • Decryption errors—(usmStatsDecryptionErrors) Total number of packets received by the SNMP engine that were dropped because they could not be decrypted.

Table 130: show snmp statistics Output Fields (*continued*)

Field Name	Field Description
Output	<p>Information about transmitted packets:</p> <ul style="list-style-type: none"> • Packets—(snmpOutPkts) Total number of messages passed from the SNMP entity to the transport service. • Too big—(snmpOutTooBigs) Total number of SNMP PDUs generated by the SNMP entity with an error status field of tooBig. • No such names—(snmpOutNoSuchNames) Total number of SNMP PDUs delivered to the SNMP entity with an error status field of noSuchName. • Bad values—(snmpOutBadValues) Total number of SNMP PDUs generated by the SNMP entity with an error status field of badValue. • General errors—(snmpOutGenErrs) Total number of SNMP PDUs generated by the SNMP entity with an error status field of genErr. • Get requests—(snmpOutGetRequests) Total number of SNMP GetRequest PDUs generated by the SNMP entity. • Get nexts—(snmpOutGetNexts) Total number of SNMP GetNext PDUs generated by the SNMP entity. • Set requests—(snmpOutSetRequests) Total number of SNMP SetRequest PDUs generated by the SNMP entity. • Get responses—(snmpOutGetResponses) Total number of SNMP GetResponse PDUs generated by the SNMP entity. • Traps—(snmpOutTraps) Total number of SNMP traps generated by the SNMP entity.

Table 131 describes the output fields for the **show snmp statistics subagents** command. Output fields are listed in the approximate order in which they appear.

Table 131: show snmp statistics subagents Output Fields

Field Name	Field Description
Subagent	Location of the SNMP subagent.
Request PDUs	Number of PDUs requested by the SNMP manager.
Response PDUs	Number of response PDUs sent by the SNMP subagent.
Request Variables	Number of variable bindings on the PDUs requested by the SNMP manager.
Response Variables	Number of variable bindings on the PDUs sent by the SNMP subagent.
Average Response Time	Average time taken by the SNMP subagent to send statistics response.
Maximum Response Time	Maximum time taken by the SNMP subagent to send the statistics response.

Sample Output

show snmp statistics

```
user@host> show snmp statistics
SNMP statistics:
  Input:
    Packets: 246213, Bad versions: 12, Bad community names: 12,
    Bad community uses: 0, ASN parse errors: 96,
    Too big: 0, No such names: 0, Bad values: 0,
    Read onlys: 0, General errors: 0,
    Total request varbinds: 227084, Total set varbinds: 67,
    Get requests: 44942, Get nexts: 190371, Set requests: 10712,
    Get responses: 0, Traps: 0,
    Silent drops: 0, Proxy drops: 0, Commit pending drops: 0,
    Throttle drops: 0,
  V3 Input:
    Unknown security models: 0, Invalid messages: 0
    Unknown pdu handlers: 0, Unavailable contexts: 0
    Unknown contexts: 0, Unsupported security levels: 1
    Not in time windows: 0, Unknown user names: 0
    Unknown engine ids: 44, Wrong digests: 23, Decryption errors: 0
  Output:
    Packets: 246093, Too big: 0, No such names: 31561,
    Bad values: 0, General errors: 2,
    Get requests: 0, Get nexts: 0, Set requests: 0,
    Get responses: 246025, Traps: 0
```

show snmp statistics subagents

```
user@host> show snmp statistics subagents

Subagent: /var/run/cosd-20
  Request PDUs: 0, Response PDUs: 0,
  Request Variables: 0, Response Variables: 0,
  Average Response Time(ms): 0.00,
  Maximum Response Time(ms): 0.00

Subagent: /var/run/pfed-30
  Request PDUs: 0, Response PDUs: 0,
  Request Variables: 0, Response Variables: 0,
  Average Response Time(ms): 0.00,
  Maximum Response Time(ms): 0.00

Subagent: /var/run/rmopd-15
  Request PDUs: 0, Response PDUs: 0,
  Request Variables: 0, Response Variables: 0,
  Average Response Time(ms): 0.00,
  Maximum Response Time(ms): 0.00

Subagent: /var/run/chassisd-30
  Request PDUs: 33116, Response PDUs: 33116,
  Request Variables: 33116, Response Variables: 33116,
  Average Response Time(ms): 1.83,
  Maximum Response Time(ms): 203.48

Subagent: /var/run/pkid-13
  Request PDUs: 0, Response PDUs: 0,
  Request Variables: 0, Response Variables: 0,
  Average Response Time(ms): 0.00,
  Maximum Response Time(ms): 0.00
```

Subagent: /var/run/apsd-13
Request PDUs: 0, Response PDUs: 0,
Request Variables: 0, Response Variables: 0,
Average Response Time(ms): 0.00,
Maximum Response Time(ms): 0.00

Subagent: /var/run/dfcd-32
Request PDUs: 0, Response PDUs: 0,
Request Variables: 0, Response Variables: 0,
Average Response Time(ms): 0.00,
Maximum Response Time(ms): 0.00

Subagent: /var/run/mib2d-33
Request PDUs: 74211, Response PDUs: 74211,
Request Variables: 74211, Response Variables: 74211,
Average Response Time(ms): 2.30,
Maximum Response Time(ms): 51.04

Subagent: /var/run/license-check-16
Request PDUs: 0, Response PDUs: 0,
Request Variables: 0, Response Variables: 0,
Average Response Time(ms): 0.00,
Maximum Response Time(ms): 0.00

Subagent: /var/run/craftd-14
Request PDUs: 0, Response PDUs: 0,
Request Variables: 0, Response Variables: 0,
Average Response Time(ms): 0.00,
Maximum Response Time(ms): 0.00

Subagent: /var/run/bfdd-19
Request PDUs: 0, Response PDUs: 0,
Request Variables: 0, Response Variables: 0,
Average Response Time(ms): 0.00,
Maximum Response Time(ms): 0.00

Subagent: /var/run/smihelperd-24
Request PDUs: 0, Response PDUs: 0,
Request Variables: 0, Response Variables: 0,
Average Response Time(ms): 0.00,
Maximum Response Time(ms): 0.00

Subagent: /var/run/cfmd-18
Request PDUs: 0, Response PDUs: 0,
Request Variables: 0, Response Variables: 0,
Average Response Time(ms): 0.00,
Maximum Response Time(ms): 0.00

Subagent: /var/run/rpd_snmp
Request PDUs: 0, Response PDUs: 0,
Request Variables: 0, Response Variables: 0,
Average Response Time(ms): 0.00,
Maximum Response Time(ms): 0.00

Subagent: /var/run/l2tpd-18
Request PDUs: 0, Response PDUs: 0,
Request Variables: 0, Response Variables: 0,
Average Response Time(ms): 0.00,
Maximum Response Time(ms): 0.00

show snmp v3

Syntax	<code>show snmp v3</code> <code><access <brief detail> community general groups notify <filter> target <address parameters> users></code>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Display the Simple Network Management Protocol version 3 (SNMPv3) operating configuration.
Options	<p>none—Display all of the SNMPv3 operating configuration.</p> <p>access—(Optional) Display SNMPv3 access information.</p> <p>brief detail—(Optional) Display brief or detailed information about SNMPv3 access information.</p> <p>community—(Optional) Display SNMPv3 community information.</p> <p>general—(Optional) Display SNMPv3 general information.</p> <p>groups—(Optional) Display SNMPv3 security-to-group information.</p> <p>notify <filter>—(Optional) Display SNMPv3 notify information and, optionally, notify filter information.</p> <p>target <address parameters>—(Optional) Display SNMPv3 target information and, optionally, either target address or target parameter information.</p> <p>users—(Optional) Display SNMPv3 user information.</p>
Additional Information	To edit the default display of the show snmp v3 command, specify options in the show statement at the [edit snmp v3] hierarchy level.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• SNMPv3 Overview on page 1460• Minimum SNMPv3 Configuration on a Device Running Junos OS on page 1461• Configuring Access Privileges for a Group on page 1511
List of Sample Output	show snmp v3 on page 1773
Output Fields	Table 132 describes the output fields for the show snmp v3 command. Output fields are listed in the approximate order in which they appear.

Table 132: show snmp v3 Output Fields

Field Name	Field Description
Local engine	<p>Information about the local SNMP engine configuration:</p> <ul style="list-style-type: none"> • Local engine ID—Unique Identifier of the local SNMPv3 engine. • Engine boots—Number of times the local SNMPv3 engine has rebooted or reinitialized since this engine ID was configured. • Engine time—Number of seconds since the local SNMPv3 engine was last rebooted or reinitialized. • Max msg size—Maximum message size the sender can accommodate.
Engine ID (local engine)	<p>Information about the local SNMP engine ID and the associated users:</p> <ul style="list-style-type: none"> • User—SNMPv3 username. • Auth/Priv—Authentication and encryption algorithm that is configured for the user. • Storage—Indicates whether a username is saved to the configuration file (nonvolatile) or not saved (volatile). Applies only to users with active status. • Status—Status of the user as listed in the SNMPv3 user table. Only rows with an active status in the table are used by the SNMPv3 engine.
Engine ID (remote engine)	<p>Information about a remote SNMP engine, associated users, user groups, and user access policies:</p> <ul style="list-style-type: none"> • User—SNMPv3 username. • Auth/Priv—Authentication and encryption algorithm that is configured for the user. • Storage—Indicates whether a username is saved to the configuration file (nonvolatile) or not (volatile). Applies only to users with active status. • Status—Status of a new user that has been activated. Only users with an active status can use SNMPv3. • Group name—Name of a group of users for which the configured access privileges apply. • Security model—Security model (such as usm, v1, v2c, or any) that is configured for the group. The security model is used with the security name to ensure messaging security. • Security name—Security name that is associated with a user, and which is used with the security model to ensure messaging security. • Storage type—Indicates whether a username is saved to the configuration file (nonvolatile) or not saved (volatile). Applies only to users with active status. • Status—Status of a user in a group. Only users with an active status can use SNMPv3.
Access control	<p>Information about access control:</p> <ul style="list-style-type: none"> • Group name—Name of a group of users for which the configured access privileges apply. • Context prefix—SNMPv3 context for which the configured access privileges apply. • Security model/level—Security model and security level combination that is configured for user access privileges. • Read view—Identifies the MIB view used for SNMPv3 read operations. • Write view—Identifies the MIB view used for SNMPv3 write operations. • Notify view—Identifies the MIB view used for outbound SNMP notifications.

Sample Output

show snmp v3

```
user@host> show snmp v3
```

Local engine ID: 80 00 0a 4c e04 31 32 33 34
Engine boots: 38
Engine time: 64583 seconds
Max msg size: 2048 bytes

Engine ID: local

User	Auth/Priv	Storage	Status
user1	md5/des	nonvolatile	active
user2	sha/none	nonvolatile	active
user3	none/none	nonvolatile	active

Engine ID: 81 00 0a 4c 04 64 64 64 64

User	Auth/Priv	Storage	Status
UNEW	md5/none	nonvolatile	active

Group name	Security model	Security name	Storage type	Status
g1	usm	user1	nonvolatile	active
g2	usm	user2	nonvolatile	active
g3	usm	user3	nonvolatile	active

Access control:

Group	Context prefix	Security model/level	Read view	Write view	Notify view
g1		usm/privacy	v1	v1	
g2		usm/authent	v1	v1	
g3		usm/none	v1	v1	

CHAPTER 58

System Logging

- request app-engine file-copy (crash | log) from-jhost to-vjunos
- request app-engine cleanup
- show app-engine crash
- show app-engine logs
- show log

request app-engine file-copy (crash | log) from-jhost to-vjunos

Syntax	<code>request app-engine file-copy (crash log) from-jhost <i>host-os-filename</i> to-vjunos <i>vjunos-filename</i></code>
Release Information	Command introduced in Junos OS Release 13.1X51-D10 for the QFX Series.
Description	<p>In Junos OS environments with a host OS, copies core files (crash option) or system log files (log option) from the host OS to a Junos OS filename.</p> <ul style="list-style-type: none">• crash option—Core files are copied from the directory where the host OS normally stores core files. For example, when the host OS is Linux, the source directory is <code>/var/crash</code>. The show app-engine crash command displays the list of core files available to copy from the host OS source directory.• log option—Log files are copied from the directory where the host OS normally stores system log files. For example, when the host OS is Linux, the source directory is <code>/var/log</code>. The show app-engine logs command displays the list of system log files available to copy from the host OS source directory. <p>Either the crash or log option is required. When the destination Junos OS path for the to-vjunos argument is a directory, the destination filename is the source filename by default. To rename the file at the destination, specify a full path that includes a destination filename in the path with the to-vjunos argument.</p>
Options	<ul style="list-style-type: none">• from-jhost <i>host-os-filename</i>—Source filename to copy from the host OS.• to-vjunos <i>vjunos-filename</i>—Junos OS destination path (with or without a filename) to which the file is copied. If this path is a directory, the source filename is used as the destination filename by default.
Additional Information	This command does not apply to accessing files on guest virtual machines (VMs) on QFX Series devices that support guest VMs.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show app-engine crash on page 1780• show app-engine logs on page 1782• request app-engine cleanup on page 1778
List of Sample Output	request app-engine file-copy crash on page 1777
Output Fields	This command produces no output. To see the results of the copy operation, you can run Junos OS commands for viewing files, such as file list with the destination directory, show system core-dumps , or show log .

Sample Output

request app-engine file-copy crash

```
user@host> request app-engine file-copy crash from-jhost
localhost.dcfpe.6449.1454328456.core.tgz to-junos /var/tmp
user@host> show system core-dumps
re0:
-----
-rw-r--r--  1 user  test    2538949 Feb 2  19:01
/var/tmp/localhost.dcpfe.6449.1454328456.core.tgz
total files: 1
```

request app-engine cleanup

Syntax	request app-engine cleanup <compute-cluster <i>compute-cluster-name</i> > <compute-cluster <i>compute-cluster-name</i> compute-node <i>compute-node-name</i> >
Release Information	Command introduced in Junos OS Release 13.1X51-D10 for the QFX Series.
Description	Cleans up temporary files on the Junos V App Engine or the host OS of QFX Series switches. This command deletes all files from the directories where temporary files are normally stored, such as /var/tmp.
Options	<ul style="list-style-type: none"> compute-cluster <i>compute-cluster-name</i> compute-node <i>compute-node-name</i>—Name of the compute cluster and name of the compute node. On QFX Series switches, the default names of the compute cluster and compute node are default-cluster and default-node, which applies to the host OS.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> show app-engine crash on page 1780 show app-engine logs on page 1782 request app-engine file-copy on page 1776
List of Sample Output	request app-engine cleanup (QFX 5100 Switches) on page 1778
Output Fields	For a description of the output fields, see Table 133 . Output fields are listed in the approximate order in which they appear.

Table 133: request app-engine cleanup Output Fields

Field Name	Field Description
Compute Cluster	Name of compute cluster.
Compute Node	Name of compute node.
Cleanup	Lists the directories that were cleaned up.

Sample Output

request app-engine cleanup (QFX 5100 Switches)

```

user@host> request app-engine cleanup
user@host> request app-engine cleanup
Compute cluster: default-cluster

Compute node: default-node

Cleanup (/var/tmp)

```

=====

show app-engine crash

Syntax	<pre>show app-engine crash <compute-cluster <i>compute-cluster-name</i>> <compute-cluster <i>compute-cluster-name</i> compute-node <i>compute-node-name</i>></pre>
Release Information	Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Description	Displays information regarding crashes on the Junos V App Engine or host OS. The command displays the list of files present in the <code>/var/crash/</code> directory, as well as some core files located in the <code>/var/tmp/</code> directory.
Options	<ul style="list-style-type: none"> compute-cluster <i>compute-cluster-name</i> compute-node <i>compute-node-name</i>—Name of the compute cluster and compute node. On QFX switches with a host OS, the default compute cluster and compute node names are default-cluster and default-node, and the command applies to the host OS by default.
Additional Information	<p>In the operational mode of the CLI when you type <code>?</code> for a name, for example a compute-node name, you would expect to get a list of available compute nodes plus the option to type in a name not listed. This is the auto-complete feature in the CLI. However, in JunosV App Engine, if you specify compute cluster and compute node in the operational command, the auto-complete works only if the compute cluster is put before the compute node.</p> <p>For commands with an optional compute-cluster <i>compute-cluster-name</i> option, if that option is omitted, the command will be executed on all compute nodes of all compute clusters. For commands with an optional compute-node <i>compute-node-name</i> option, if that option is omitted, the command will be executed on all compute nodes of the specified compute cluster.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> request app-engine file-copy on page 1776
List of Sample Output	show app-engine crash (QFX5100 Switches) on page 1781
Output Fields	For a description of the output fields, see Table 134 . Output fields are listed in the approximate order in which they appear.

Table 134: show app-engine crash Output Fields

Field Name	Field Description
Compute Cluster	Name of compute cluster.
Compute Node	Name of compute node.
Crash Info	List of core files.

Sample Output

show app-engine crash (QFX5100 Switches)

```
user@host> show app-engine crash

Compute cluster: default-cluster

Compute node: default-node

Crash Info
=====
total 0
```

show app-engine logs

Syntax	<pre>show app-engine logs <compute-cluster <i>compute-cluster-name</i>> <compute-cluster <i>compute-cluster-name</i> compute-node <i>compute-node-name</i>></pre>
Release Information	Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Description	Displays log files regarding the state of the specified Guest VM or host OS, located in the <code>/var/log/</code> directory.
Options	<ul style="list-style-type: none"> compute-cluster <i>compute-cluster-name</i> compute-node <i>compute-node-name</i>—Name of the compute cluster and name of the compute node. On QFX switches with a host OS, the default compute cluster and compute node names are default-cluster and default-node, and the command applies to the host OS by default.
Additional Information	<p>In the operational mode of the CLI when you type <code>?</code> for a name, for example a compute-node name, you would expect to get a list of available compute nodes plus the option to type in a name not listed. This is the auto-complete feature in the CLI. However, in JunosV App Engine, if you specify compute cluster and compute node in the operational command, the auto-complete works only if the compute cluster is put before the compute node.</p> <p>For commands with an optional compute-cluster <i>compute-cluster-name</i> option, if that option is omitted, the command will be executed on all compute nodes of all compute clusters. For commands with an optional compute-node <i>compute-node-name</i> option, if that option is omitted, the command will be executed on all compute nodes of the specified compute cluster.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> request app-engine file-copy on page 1776
List of Sample Output	show app-engine logs (QFX5100 Switches) on page 1783
Output Fields	For a description of the output fields, see Table 135 . Output fields are listed in the approximate order in which they appear.

Table 135: show app-engine logs Output Fields

Field Name	Field Description
Compute Cluster	Name of compute cluster.
Compute Node	Name of compute node.
Logs Info	List of log files.

Sample Output

show app-engine logs (QFX5100 Switches)

```

user@host> show app-engine logs
Compute cluster: default-cluster

Compute node: default-node

Logs Info
=====
total 1012
drwxr-xr-x 2 root root 4096 Feb 2 13:53 audit
-rw-r--r-- 1 root root 0 Feb 2 13:59 named.log
-rw-r--r-- 1 root root 0 Feb 2 15:30 btmp
drwxr-xr-x 2 root root 4096 Feb 2 17:11 ntpstats
-rw----- 1 root root 0 Feb 2 17:27 tallylog
drwxr-xr-x 2 root root 4096 Feb 7 10:10 stap-server
-rw----- 1 root root 0 Feb 7 10:16 spooler
-rw----- 1 root root 0 Feb 7 10:16 maillog
-rw----- 1 root root 0 Feb 7 10:16 boot.log
drwxr-xr-x 5 root root 4096 Feb 7 10:16 libvirt
drwxr-x-- 2 root root 4096 Feb 7 10:16 watchdog
-rw-r--r-- 1 root root 41693 Feb 10 04:27 dmesg.old
drwxr-xr-x 2 root root 4096 Feb 10 04:30 sa
-rw-r--r-- 1 root root 41693 Feb 10 09:35 dmesg
-rw-rw-r-- 1 root root 1193 Feb 10 09:47 vjunos_install_log
-rw-rw-r-- 1 root utmp 141312 Feb 10 10:08 wtmp
-rw----- 1 root root 25514 Feb 10 10:08 secure
-rw-r--r-- 1 root root 292 Feb 10 10:08 lastlog
-rw----- 1 root root 338593 Feb 10 10:17
cron
-rw-r--r-- 1 root root 383229 Feb 10 10:17 messages

```

show log

List of Syntax	Syntax on page 1784 Syntax (QFX Series and OCX Series) on page 1784 Syntax (TX Matrix Router) on page 1784
Syntax	<code>show log</code> <code><filename user <username>></code>
Syntax (QFX Series and OCX Series)	<code>show log filename</code> <code><device-type (device-id device-alias)></code>
Syntax (TX Matrix Router)	<code>show log</code> <code><all-lcc lcc <i>number</i> scc></code> <code><filename user <username>></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. Option <i>device-type (device-id device-alias)</i> is introduced in Junos OS Release 13.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	List log files, display log file contents, or display information about users who have logged in to the router or switch.



NOTE: On MX Series routers, modifying a configuration to replace a service interface with another service interface is treated as a catastrophic event. When you modify a configuration, the entire configuration associated with the service interface—including NAT pools, rules, and service sets—is deleted and then re-created for the newly specified service interface. If there are active sessions associated with the service interface that is being replaced, these sessions are deleted and the NAT pools are then released, which leads to the generation of the NAT_POOL_RELEASE system log messages. However, because NAT pools are already deleted as a result of the catastrophic configuration change and no longer exist, the NAT_POOL_RELEASE system log messages are not generated for the changed configuration.

Options **none**—List all log files.

<all-lcc | lcc *number* | scc>—(Routing matrix only) (Optional) Display logging information about all T640 routers (or line-card chassis) or a specific T640 router (replace *number* with a value from 0 through 3) connected to a TX Matrix router. Or, display logging information about the TX Matrix router (or switch-card chassis).

device-type—(QFabric system only) (Optional) Display log messages for only one of the following device types:

- **director-device**—Display logs for Director devices.
- **infrastructure-device**—Display logs for the logical components of the QFabric system infrastructure, including the diagnostic Routing Engine, fabric control Routing Engine, fabric manager Routing Engine, and the default network Node group and its backup (NW-NG-0 and NW-NG-0-backup).
- **interconnect-device**—Display logs for Interconnect devices.
- **node-device**—Display logs for Node devices.



NOTE: If you specify the **device-type** optional parameter, you must also specify either the **device-id** or **device-alias** optional parameter.

(device-id | device-alias)—If a device type is specified, display logs for a device of that type. Specify either the device ID or the device alias (if configured).

filename—(Optional) Display the log messages in the specified log file. For the routing matrix, the filename must include the chassis information.

user <username>—(Optional) Display logging information about users who have recently logged in to the router or switch. If you include **username**, display logging information about the specified user.

Required Privilege Level trace

Related Documentation • [syslog \(System\) on page 1704](#)

List of Sample Output [show log on page 1785](#)
[show log filename on page 1786](#)
[show log filename \(QFabric System\) on page 1786](#)
[show log user on page 1787](#)

Sample Output

show log

```
user@host> show log
total 57518
-rw-r--r-- 1 root bin      211663 Oct  1 19:44 dcd
-rw-r--r-- 1 root bin      999947 Oct  1 19:41 dcd.0
-rw-r--r-- 1 root bin      999994 Oct  1 17:48 dcd.1
-rw-r--r-- 1 root bin      238815 Oct  1 19:44 rpd
-rw-r--r-- 1 root bin     1049098 Oct  1 18:00 rpd.0
-rw-r--r-- 1 root bin     1061095 Oct  1 12:13 rpd.1
-rw-r--r-- 1 root bin     1052026 Oct  1 06:08 rpd.2
-rw-r--r-- 1 root bin     1056309 Sep 30 18:21 rpd.3
-rw-r--r-- 1 root bin     1056371 Sep 30 14:36 rpd.4
-rw-r--r-- 1 root bin     1056301 Sep 30 10:50 rpd.5
-rw-r--r-- 1 root bin     1056350 Sep 30 07:04 rpd.6
```

```
-rw-r--r-- 1 root bin      1048876 Sep 30 03:21 rpd.7
-rw-rw-r-- 1 root bin      19656 Oct  1 19:37 wtmp
```

show log filename

```
user@host> show log rpd
Oct  1 18:00:18 trace_on: Tracing to ?/var/log/rpd? started
Oct  1 18:00:18 EVENT <MTU> ds-5/2/0.0 index 24 <Broadcast PointToPoint Multicast
Oct  1 18:00:18
Oct  1 18:00:19 KRT recv len 56 V9 seq 148 op add Type route/if af 2 addr
13.13.13.21 nhop type local nhop 13.13.13.21
Oct  1 18:00:19 KRT recv len 56 V9 seq 149 op add Type route/if af 2 addr
13.13.13.22 nhop type unicast nhop 13.13.13.22
Oct  1 18:00:19 KRT recv len 48 V9 seq 150 op add Type ifaddr index 24 devindex
43
Oct  1 18:00:19 KRT recv len 144 V9 seq 151 op chnge Type ifdev devindex 44
Oct  1 18:00:19 KRT recv len 144 V9 seq 152 op chnge Type ifdev devindex 45
Oct  1 18:00:19 KRT recv len 144 V9 seq 153 op chnge Type ifdev devindex 46
Oct  1 18:00:19 KRT recv len 1272 V9 seq 154 op chnge Type ifdev devindex 47
...
```

show log filename (QFabric System)

```
user@qfabric> show log messages
Mar 28 18:00:06 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:06 ED1486
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 2159)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1486
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 2, jnxFruL3Index 0,
jnxFruName PIC: @ 0/1/*, jnxFruType 11, jnxFruSlot 0, jnxFruOfflineReason 2,
jnxFruLastPowerOff 0, jnxFruLastPowerOn 2191)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 242726)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 2, jnxFruL3Index 0,
jnxFruName PIC: @ 0/1/*, jnxFruType 11, jnxFruSlot 0, jnxFruOfflineReason 2,
jnxFruLastPowerOff 0, jnxFruLastPowerOn 242757)
Mar 28 18:00:16 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:16 ED1486
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:27 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:27 ED1486
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50
_DCF_default__NW-INE-0_RE0_ file: UI_COMMIT: User 'root' requested 'commit'
operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50
_DCF_default__NW-INE-0_RE0_ file: UI_COMMIT: User 'root' requested 'commit'
operation (comment: none)
Mar 28 18:00:55 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:55 ED1492
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:01:10 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:01:10 ED1492
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:02:37 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:02:37 ED1491
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
```

```
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,  
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 33809)
```

show log user

```
user@host> show log user  
darius  mg2546          Thu Oct  1 19:37  still logged in  
darius  mg2529          Thu Oct  1 19:08 - 19:36  (00:28)  
darius  mg2518          Thu Oct  1 18:53 - 18:58  (00:04)  
root    mg1575          Wed Sep 30 18:39 - 18:41  (00:02)  
root    tty2    jun.site.per Wed Sep 30 18:39 - 18:41  (00:02)  
alex    tty1    192.168.1.2  Wed Sep 30 01:03 - 01:22  (00:19)
```


User and Access Management Feature Guide for QFX10000 Switches

PART 23

User and Access Management Overview

- [Understanding the Software on page 1793](#)
- [Understanding Access and Authentication Methods on page 1799](#)

CHAPTER 59

Understanding the Software

- [Understanding Junos OS Infrastructure and Processes on page 1793](#)
- [Understanding LLDP on page 1795](#)
- [Monitoring SNMP on page 1796](#)

Understanding Junos OS Infrastructure and Processes

Junos OS includes processes for Internet Protocol (IP) routing and for managing interfaces, networks, and the switch.

Junos OS runs on the Routing Engine. The Routing Engine kernel coordinates communication among the Junos OS processes and provides a link to the Packet Forwarding Engine.

Using the Junos OS command-line interface (CLI), you configure switching features and set the properties of network interfaces. After activating a software configuration, use either the Junos Space or CLI user interface to monitor, manage operations, and diagnose protocol and network connectivity problems.

- [Routing Engine and Packet Forwarding Engine on page 1793](#)
- [Junos OS Processes on page 1794](#)

Routing Engine and Packet Forwarding Engine

A switch has two primary software processing components:

- **Packet Forwarding Engine**—Processes packets; applies filters, routing policies, and other features; and forwards packets to the next hop along the route to their final destination.
- **Routing Engine**—Provides three main functions:
 - Creates the packet forwarding switch, which provides route lookup, filtering, and switching on incoming data packets, and then directs outbound packets to the appropriate interface for transmission to the network.
 - Maintains the routing tables used by the switch and controls the routing protocols that run on the switch.
 - Provides control and monitoring functions for the switch, including controlling power and monitoring system status.

Junos OS Processes

Junos OS running on the Routing Engine and Packet Forwarding Engine consists of multiple processes that are responsible for individual functions.

The separation of functions provides operational stability, because each process accesses its own protected memory space. In addition, because each process is a separate software package, you can selectively upgrade all or part of the Junos OS for added flexibility.

Table 136 describes the primary Junos OS processes.

Table 136: Junos OS Processes

Process	Name	Description
Chassis process	chassisd	<p>Detects hardware on the system that is used to configure network interfaces.</p> <p>Monitors the physical status of hardware components and field-replaceable units (FRUs), detecting when environment sensors such as temperature sensors are triggered.</p> <p>Relays signals and interrupts—for example, when devices are taken offline, so that the system can close sessions and shut down gracefully.</p>
DNS server process	named-service	Resolves hostnames into addresses.
Dynamic Host Configuration Protocol (DHCP) process	dhcp-service	Enables a DHCP server to allocate network IP addresses and deliver configuration settings to client hosts without user intervention.
Ethernet switching process	eswd	<p>Handles Layer 2 switching functionality such as MAC address learning, Spanning Tree Protocol, and access port security.</p> <p>Manages Ethernet switching interfaces, VLANs, and VLAN interfaces.</p>
Firewall management process	firewall	Manages the firewall configuration and helps accept or reject packets that are transiting an interface on a switch.
Forwarding process	pfem	Defines how routing protocols operate on the partition. The overall performance of the partition is largely determined by the effectiveness of the forwarding process.
Interface process	dcd	Configures and monitors network interfaces by defining physical characteristics such as link encapsulation, hold times, and keepalive timers.
Integrated Local Management Interface (ILMI) process	ilmi	Provides bidirectional exchange of management information between two ATM interfaces across a physical connection.
Link Management Protocol (LMP) process	link-management	Establishes and maintains LMP control channels.

Table 136: Junos OS Processes (*continued*)

Process	Name	Description
Management process	mgd	<p>Provides communication between the other processes and an interface to the configuration database.</p> <p>Populates the configuration database with configuration information and retrieves the information when queried by other processes to ensure that the system operates as configured.</p> <p>Interacts with the other processes when commands are issued through one of the user interfaces on the partition.</p> <p>If a process terminates or fails to start when called, the management process attempts to restart it a limited number of times to prevent thrashing and logs any failure information for further investigation.</p>
Multicast snooping process	multicast-snooping	Makes Layer 2 devices, such as VLAN switches, aware of Layer 3 information, such as the media access control (MAC) addresses of members of a multicast group.
Secure Neighbor Discovery (SEND) protocol process	send	Protects Neighbor Discovery Protocol (NDP) messages.
Simple Network Management Protocol (SNMP) process	snmp	Enables the monitoring of network devices from a central location and provides the switch's SNMP master agent.
Tunnel OAM process	tunnel-oamd	Enables the Operation, Administration, and Maintenance of Layer 2 tunneled networks. Layer 2 protocol tunneling (L2PT) allows service providers to send Layer 2 protocol data units (PDUs) across the provider's cloud and deliver them to Juniper Networks EX Series Ethernet Switches that are not part of the local broadcast domain.
Virtual Router Redundancy Protocol (VRRP) process	vrrp	Enables hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts.

Related Documentation

- *Junos OS Baseline Network Operations Guide*
- *Junos OS Administration Library for Routing Devices*

Understanding LLDP

The device uses Link Layer Discovery Protocol (LLDP) to learn and distribute device information on network links. The information enables the switch to identify a variety of devices quickly. This quick identification results in a LAN that interoperates smoothly and efficiently.

LLDP-capable devices transmit information in type, length, and value (TLV) messages to neighbor devices. Device information can include specifics, such as chassis and port

identification and system name and system capabilities. The TLVs leverage this information from parameters that have already been configured in Junos OS.

The device supports the following basic TLVs:

- **Chassis Identifier**—The MAC address associated with the local system.
- **Port Identifier**—The port identification for the specified port in the local system.
- **Port Description**—The user-configured port description. The port description can be a maximum of 256 characters.
- **System Name**—The user-configured name of the local system. The system name can be a maximum of 256 characters.
- **System Description**—The system description containing information about the software and current image running on the system. This information cannot be configured, but is taken from the software.
- **System Capabilities**—The primary function performed by the system. The capabilities that system supports are defined; for example, bridge or router. This information cannot be configured, but is based on the model of the product.
- **Management Address**—The IP management address of the local system.

The device supports the following 802.3 TLVs:

- **Power via MDI**—A TLV that advertises media dependent interface (MDI) power support, power source equipment (PSE) power pair, and power class information.
- **MAC/PHY Configuration Status**—A TLV that advertises information about the physical interface, such as autonegotiation status and support and MAU type. The information cannot be configured, but is based on the physical interface structure.
- **Link Aggregation**—A TLV that advertises whether the port is aggregated and its aggregated port ID.
- **Maximum Frame Size**—A TLV that advertises the Maximum Transmission Unit (MTU) of the interface sending LLDP frames.
- **Port Vlan**—A TLV that advertises the VLAN name configured on the interface.

**Related
Documentation**

- *Configuring LLDP*

Monitoring SNMP

There are several commands that you can access in Junos OS operational mode to monitor SNMP information. Some of the commands are:

- **show snmp health-monitor**, which displays the health monitor log and alarm information.
- **show snmp mib**, which displays information from the MIBs, such as device and system information.

- **show snmp statistics**, which displays SNMP statistics such as the number of packets, silent drops, and invalid output values.
- **show snmp rmon**, which displays the RMON alarm, event, history, and log information

The following example provides sample output from the **show snmp health-monitor** command:

```
user@switch> show snmp health-monitor
Alarm
Index  Variable description                               Value State

32768 Health Monitor: root file system utilization
      jnxHrStoragePercentUsed.1                      58 active

32769 Health Monitor: /config file system utilization
      jnxHrStoragePercentUsed.2                      0 active

32770 Health Monitor: RE 0 CPU utilization
      jnxOperatingCPU.9.1.0.0                        0 active

32773 Health Monitor: RE 0 Memory utilization
      jnxOperatingBuffer.9.1.0.0                    35 active

32775 Health Monitor: jkernel daemon CPU utilization
      Init daemon                                    0 active
      Chassis daemon                                50 active
      Firewall daemon                               0 active
      Interface daemon                              5 active
      SNMP daemon                                   11 active
      MIB2 daemon                                   42 active
      ...
```

The following example provides sample output from the **show snmp mib** command:

```
user@switch> show snmp mib walk system

sysDescr.0    = Juniper Networks, Inc. qfx3500s internet router, kernel
JUNOS 11.1-20100926.0 #0: 2010-09-26 06:17:38 UTC builder@abc.juniper.net:
/volume/build/junos/11.1/production/20100926.0/obj-xlr/bsd/sys/compile/JUNIPER-xxxxx

Build date: 2010-09-26 06:00:10 U
sysObjectID.0 = jnxProductQFX3500
sysUpTime.0   = 24444184
sysContact.0  = J Smith
sysName.0     = Lab QFX3500
sysLocation.0 = Lab
sysServices.0 = 4
```

The following example provides sample output from the **show snmp statistics** command:

```
user@switch> show snmp statistics

SNMP statistics:
Input:
  Packets: 0, Bad versions: 0, Bad community names: 0,
  Bad community uses: 0, ASN parse errors: 0,
  Too big: 0, No such names: 0, Bad values: 0,
  Read only: 0, General errors: 0,
```

```
Total request varbinds: 0, Total set varbinds: 0,  
Get requests: 0, Get nexts: 0, Set requests: 0,  
Get responses: 0, Traps: 0,  
Silent drops: 0, Proxy drops: 0, Commit pending drops: 0,  
Throttle drops: 0, Duplicate request drops: 0  
Output:  
Packets: 0, Too bigs: 0, No such names: 0,  
Bad values: 0, General errors: 0,  
Get requests: 0, Get nexts: 0, Set requests: 0,  
Get responses: 0, Traps: 0
```

- Related Documentation**
- [health-monitor on page 1629](#)
 - [show snmp mib on page 1756](#)
 - [show snmp statistics on page 1764](#)

CHAPTER 60

Understanding Access and Authentication Methods

- [Understanding Junos OS Access Privilege Levels on page 1799](#)
- [Junos OS User Authentication Methods on page 1804](#)
- [Understanding Login Authentication on page 1804](#)

Understanding Junos OS Access Privilege Levels

Each top-level command-line interface (CLI) command and each configuration statement have an access privilege level associated with them. Users can execute only those commands and configure and view only those statements for which they have access privileges. The access privileges for each login class are defined by one or more *permission flags*.

For each login class, you can explicitly deny or allow the use of operational and configuration mode commands that would otherwise be permitted or not allowed by a privilege level specified in the **permissions** statement.

The following sections provide additional information about permissions:

- [Junos OS Login Class Permission Flags on page 1799](#)
- [Allowing or Denying Individual Commands for Junos OS Login Classes on page 1802](#)

Junos OS Login Class Permission Flags

The **permissions** statement specifies one or more of the permission flags listed in [Table 137](#). Permission flags are not cumulative, so for each class you must list all the permission flags needed, including **view** to display information and **configure** to enter configuration mode. Two forms of permissions control for individual parts of the configuration are:

- "Plain" form—Provides read-only capability for that permission type. An example is **interface**.
- Form that ends in **-control**—Provides read and write capability for that permission type. An example is **interface-control**.

[Table 137](#) lists the Junos OS login class permission flags that you can configure by including the **permissions** statement at the **[edit system login class class-name]** hierarchy level.

Table 137: Login Class Permission Flags

Permission Flag	Description
access	Can view the access configuration in configuration mode and with the show configuration operational mode command.
access-control	Can view and configure access information at the [edit access] hierarchy level.
admin	Can view user account information in configuration mode and with the show configuration operational mode command.
admin-control	Can view user accounts and configure them at the [edit system login] hierarchy level.
all-control	Can access all operational mode commands and configuration mode commands. Can modify configuration in all the configuration hierarchy levels.
clear	Can clear (delete) information learned from the network that is stored in various network databases by using the clear commands.
configure	Can enter configuration mode by using the configure command.
control	Can perform all control-level operations—all operations configured with the -control permission flags.
field	Can view field debug commands. Reserved for debugging support.
firewall	Can view the firewall filter configuration in configuration mode.
firewall-control	Can view and configure firewall filter information at the [edit firewall] hierarchy level.
floppy	Can read from and write to the removable media.
flow-tap	Can view the flow-tap configuration in configuration mode.
flow-tap-control	Can view the flow-tap configuration in configuration mode and can configure flow-tap configuration information at the [edit services flow-tap] hierarchy level.
flow-tap-operation	<p>Can make flow-tap requests to the router or switch. For example, a Dynamic Tasking Control Protocol (DTCP) client must authenticate itself to the Junos OS as an administrative user. That account must have flow-tap-operation permission.</p> <p>NOTE: The flow-tap-operation option is not included in the all-control permissions flag.</p>
idp-profiler-operation	Can view profiler data.

Table 137: Login Class Permission Flags (*continued*)

Permission Flag	Description
interface	Can view the interface configuration in configuration mode and with the show configuration operational mode command.
interface-control	Can view chassis, class of service (CoS), groups, forwarding options, and interfaces configuration information. Can edit configuration at the following hierarchy levels: <ul style="list-style-type: none"> • [edit chassis] • [edit class-of-service] • [edit groups] • [edit forwarding-options] • [edit interfaces]
maintenance	Can perform system maintenance, including starting a local shell on the router or switch and becoming the superuser in the shell by using the su root command, and can halt and reboot the router or switch by using the request system commands.
network	Can access the network by using the ping , ssh , telnet , and traceroute commands.
pgcp-session-mirroring	Can view the pgcp session mirroring configuration.
pgcp-session-mirroring-control	Can modify the pgcp session mirroring configuration.
reset	Can restart software processes by using the restart command and can configure whether software processes are enabled or disabled at the [edit system processes] hierarchy level.
rollback	Can use the rollback command to return to a previously committed configuration other than the most recently committed one.
routing	Can view general routing, routing protocol, and routing policy configuration information in configuration and operational modes.
routing-control	Can view general routing, routing protocol, and routing policy configuration information and can configure general routing at the [edit routing-options] hierarchy level, routing protocols at the [edit protocols] hierarchy level, and routing policy at the [edit policy-options] hierarchy level.
secret	Can view passwords and other authentication keys in the configuration.
secret-control	Can view passwords and other authentication keys in the configuration and can modify them in configuration mode.

Table 137: Login Class Permission Flags (*continued*)

Permission Flag	Description
security	Can view security configuration in configuration mode and with the show configuration operational mode command.
security-control	Can view and configure security information at the [edit security] hierarchy level.
shell	Can start a local shell on the router or switch by using the start shell command.
snmp	Can view Simple Network Management Protocol (SNMP) configuration information in configuration and operational modes.
snmp-control	Can view SNMP configuration information and can modify SNMP configuration at the [edit snmp] hierarchy level.
system	Can view system-level information in configuration and operational modes.
system-control	Can view system-level configuration information and configure it at the [edit system] hierarchy level.
trace	Can view trace file settings and configure trace file properties.
trace-control	Can modify trace file settings and configure trace file properties.
view	Can use various commands to display current system-wide, routing table, and protocol-specific values and statistics. Cannot view the secret configuration.
view-configuration	Can view all of the configuration excluding secrets, system scripts, and event options. NOTE: Only users with the maintenance permission can view commit script, op script, or event script configuration.

Allowing or Denying Individual Commands for Junos OS Login Classes

By default, all top-level CLI commands have associated access privilege levels. Users can execute only those commands and view only those statements for which they have access privileges. For each login class, you can explicitly deny or allow the use of operational and configuration mode commands that would otherwise be permitted or not allowed by a privilege level specified in the **permissions** statement.

Permission flags are used to grant a user access to operational mode commands and configuration hierarchy levels and statements. By specifying a specific permission flag on the user's login class at the **[edit system login class]** hierarchy level, you grant the user access to the corresponding commands and configuration hierarchy levels and statements. To grant access to all commands and configuration statements, use the **all**

permissions flag. For permission flags that grant access to configuration hierarchy levels and statements, the flags grant read-only privilege to that configuration. For example, the **interface** permissions flag grants read-only access to the **[edit interfaces]** hierarchy level. The **-control** form of the flag grants read-write access to that configuration. Using the preceding example, **interface-control** grants read-write access to the **[edit interfaces]** hierarchy level.

- The **all** login class permission bits take precedence over extended regular expressions when a user with **rollback** permission issues the **rollback** command.
- Expressions used to allow and deny commands for users on RADIUS and TACACS+ servers have been simplified. Instead of a single, long expression with multiple commands (**allow-commands=cmd1 cmd2 ... cmdn**), you can specify each command as a separate expression. This new syntax is valid for **allow-configuration-regexps** and **deny-configuration-regexps**, **allow-commands** and **deny-commands**, and all user permission bits.
- Users cannot issue the **load override** command when specifying an extended regular expression. Users can only issue the **merge**, **replace**, and **patch** configuration commands.
- If you allow and deny the same commands, the **allow-commands** permissions take precedence over the permissions specified by the **deny-commands**. For example, if you include **allow-commands "request system software add"** and **deny-commands "request system software add"**, the login class user is allowed to install software using the **request system software add** command.
- Regular expressions for **allow-commands** and **deny-commands** can also include the **commit**, **load**, **rollback**, **save**, **status**, and **update** commands.
- If you specify a regular expression for **allow-commands** and **deny-commands** with two different variants of a command, the longest match is always executed.

For example, if you specify a regular expression for **allow-commands** with the **commit-synchronize** command and a regular expression for **deny-commands** with the **commit** command, users assigned to such a login class would be able to issue the **commit synchronize** command, but not the **commit** command. This is because **commit-synchronize** is the longest match between **commit** and **commit-synchronize** and it is specified for **allow-commands**.

Likewise, if you specify a regular expression for **allow-commands** with the **commit** command and a regular expression for **deny-commands** with the **commit-synchronize** command, users assigned to such a login class would be able to issue the **commit** command, but not the **commit-synchronize** command. This is because **commit-synchronize** is the longest match between **commit** and **commit-synchronize** and it is specified for **deny-commands**.

Related Documentation

- [Configuring Access Privilege Levels on page 1809](#)
- [Access Privilege User Permission Flags Overview](#)

Junos OS User Authentication Methods

The Junos OS supports three methods of user authentication: local password authentication, Remote Authentication Dial-In User Service (RADIUS), and Terminal Access Controller Access Control System Plus (TACACS+).

With local password authentication, you configure a password for each user allowed to log in to the router or switch.

RADIUS and TACACS+ are authentication methods for validating users who attempt to access the router or switch using telnet. They are both distributed client-server systems—the RADIUS and TACACS+ clients run on the router or switch, and the server runs on a remote network system.

You can configure the router or switch to be both a RADIUS and TACACS+ client, and you can also configure authentication passwords in the Junos OS configuration file. You can prioritize the methods to configure the order in which the software tries the different authentication methods when verifying user access.

Related Documentation

- *Configuring RADIUS Server Authentication*
- *Configuring TACACS+ Authentication*
- [Junos OS Authentication Order for RADIUS, TACACS+, and Password Authentication on page 1859](#)

Understanding Login Authentication

You can control access to your network using several different authentication methods—media access control (MAC) RADIUS, for example. Authentication prevents unauthorized devices and users from gaining access to your LAN. For MAC RADIUS authentication, end devices must be authenticated before they receive an IP address from a DHCP server.

You can enable end devices to access the network without authenticating on the RADIUS server by configuring the MAC address of the end device in the static MAC bypass list by configuring the MAC address using the **authentication-whitelist** statement.

You can configure one or more authentication methods on a single interface and thereby enable fallback to the next method if the first or second method is unsuccessful.

On a single interface you can configure one or a combination of several authentication methods.

This topic covers:

- [MAC RADIUS Authentication on page 1805](#)

MAC RADIUS Authentication

You can configure MAC RADIUS authentication on interfaces that are connected to end devices.

The EAP method supported for MAC RADIUS authentication is EAP-MD5.

When you configure the **mac-radius restrict** option, the switch immediately attempts a MAC- RADIUS authentication by sending a request to the RADIUS server for authentication of the MAC address of the end device. If MAC address of the end device is configured for RADIUS authentication, LAN access between the two switches is created.

Related Documentation

- [Configuring RADIUS Authentication \(QFX Series or OCX Series\) on page 1873](#)

PART 24

Configuring Access

- [Configuring and Managing Root Users on page 1809](#)
- [Configuring and Managing User Accounts on page 1821](#)

CHAPTER 61

Configuring and Managing Root Users

- [Configuring Management Access on page 1809](#)
- [Configuring Access Privilege Levels on page 1809](#)
- [Configuring Login Tips on page 1810](#)
- [Recovering the Root Password on page 1810](#)
- [Example: Configuring a Plain-Text Password for Root Logins on page 1812](#)
- [Example: Configuring SSH Authentication for Root Logins on page 1814](#)
- [Understanding Troubleshooting Resources on page 1814](#)
- [Troubleshooting Overview on page 1816](#)
- [Recovering the Root Password on page 1818](#)

Configuring Management Access

To define the management access settings for the routing platform:

1. Next to Allow Telnet Access, select the check box to allow remote Telnet access to the routing platform.
2. Next to Allow SSH Access, selected the check box to allow remote SSH access to the routing platform.
3. Click **Apply** to apply the configuration.

Related Documentation

- [Configuring Junos OS User Accounts on page 125](#)
- [Specifying Access Privileges for Junos OS Operational Mode Commands on page 30](#)
- [Example: Configuring User Permissions with Access Privilege Levels on page 1831](#)

Configuring Access Privilege Levels

Each top-level command-line interface (CLI) command and each configuration statement have an access privilege level associated with it. Users can execute only those commands and configure and view only those statements for which they have access privileges.

To configure access privilege levels, include the **permissions** statement at the **[edit system login class *class-name*]** hierarchy level:

```
[edit system login class class-name]  
permissions [ permissions ];
```

**Related
Documentation**

- [Example: Configuring User Permissions with Access Privilege Levels on page 1831](#)
- [Understanding Junos OS Access Privilege Levels on page 1799](#)
- [Specifying Access Privileges for Junos OS Operational Mode Commands on page 30](#)
- *permissions*

Configuring Login Tips

The Junos OS CLI provides the option of configuring login tips for the user. By default, the **tip** command is not enabled when a user logs in.

- To enable tips, include the **login-tip** statement at the **[edit system login class *class-name*]** hierarchy level:

```
[edit system login class class-name]  
login-tip;
```

Adding this statement enables the **tip** command for the class specified, provided the user logs in using the CLI.

**Related
Documentation**

- [Defining Junos OS Login Classes](#)

Recovering the Root Password

If you forget the root password, you can use the password recovery procedure to reset the root password.



NOTE: The root password cannot be recovered on a QFabric system.



NOTE: You need console access to the switch to recover the root password.

To recover the root password:

1. Power off the switch by switching off the AC power outlet of the device or, if necessary, by pulling the power cords out of the device's power supplies.
2. Turn off the power to the management device, such as a PC or laptop computer, that you want to use to access the CLI.
3. Plug one end of the Ethernet rollover cable supplied with the device into the RJ-45-to-DB-9 serial port adapter supplied with the device.
4. Plug the RJ-45-to-DB-9 serial port adapter into the serial port on the management device.

5. Connect the other end of the Ethernet rollover cable to the console port on the device.
6. Turn on the power to the management device.
7. On the management device, start your asynchronous terminal emulation application (such as Microsoft Windows Hyperterminal) and select the appropriate **COM** port to use (for example, **COM1**).
8. Configure the port settings as follows:

- Bits per second: 9600
- Data bits: 8
- Parity: None
- Stop bits: 1
- Flow control: None

9. Power on the device by (if necessary) plugging the power cords into the device's power supply, or turning on the power to the device by switching on the AC power outlet the device is plugged into.

The terminal emulation screen on your management device displays the device's boot sequence.

10. When the following prompt appears, press the Spacebar to access the device's bootstrap loader command prompt:

```
Hit [Enter] to boot immediately, or space bar for command prompt.  
Booting [kernel] in 9 seconds...
```

11. At the following prompt, enter **boot -s** to start up the system in single-user mode.

```
ok boot -s
```

12. At the following prompt, enter **recovery** to start the root password recovery procedure.

```
Enter full pathname of shell or 'recovery' for root password recovery or RETURN  
for /bin/sh: recovery
```

13. Enter configuration mode in the CLI.

14. Set the root password. For example:

```
user@switch# set system root-authentication plain-text-password
```

15. At the following prompt, enter the new root password. For example:

```
New password: juniper1  
Retype new password:
```

16. At the second prompt, reenter the new root password.

17. After you have finished configuring the password, commit the configuration.

```
root@host# commit  
commit complete
```

18. Exit configuration mode in the CLI.

19. Exit operational mode in the CLI.

20. At the prompt, enter **y** to reboot the device.

```
Reboot the system? [y/n] y
```

Related Documentation

- [Configuring the Root Password](#)

Example: Configuring a Plain-Text Password for Root Logins

This example shows how to configure a plain-text password for the root-level user (whose username is *root*). Configuring a plain-text password is one way to protect access to the root level by unauthorized users. You must prevent unauthorized users from gaining access to superuser commands that can be used to alter your system configuration.

- [Requirements on page 1812](#)
- [Overview on page 1812](#)
- [Configuration on page 1812](#)
- [Verification on page 1813](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.

Make sure that you understand the requirements for a valid plain-text password. For Junos OS, the default requirements for a plain-text password are as follows:

- Must be from 6 up to 128 characters long.
- Can include most character classes (uppercase letters, lowercase letters, numbers, punctuation marks, and other special characters). Control characters are not recommended.
- Must contain at least one change of case or character class.

Overview

Junos OS is preinstalled on the router. When the router is powered on, it is ready to be configured. Initially, you log in as the root-level user with no password. To set the root password, you have several options. This example shows how to enter a plain-text password that Junos OS then encrypts for you.

Configuration

CLI Quick Configuration To quickly configure this example, copy the following command and paste it into the window. When prompted, type the new password, and then when prompted, retype it.

```
set system root-authentication plain-text-password
```

Configuring a Plain-Text Password for User Root

Step-by-Step Procedure To configure a plain-text password for the root-level user:

1. Type the **set** command for the plain-text password and press Enter.

```
[edit]
```

```
user@host# set system root-authentication plain-text-password
```

New password:

2. Type the new password next to the **New password** prompt and press Enter.

New password: *new-password*

Retype new password:

3. Retype the same password next to the **Retype new password** prompt and press Enter.

Results

From configuration mode, confirm your configuration by using the **show** command. It should look something like this:

```
[edit ]
user@host# show system
root-authentication {
  encrypted-password "$1$ASwBkGYd$YUcEwgd0IO4QkRzzlQdmT/"; ## SECRET-DATA
}
```

If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

After you have confirmed that the configuration is correct, enter **commit** from configuration mode.

Verification

Verifying the Configuration of a Plain-Text Password for User Root

Purpose Verify the configuration of a plain-text password for the root-level user.

Action From operational mode, confirm your configuration by entering the **show configuration system** command.

```
user@host> show configuration system
root-authentication {
  encrypted-password "$1$ASwBkGYd$YUcEwgd0IO4QkRzzlQdmT/"; ## SECRET-DATA
}
```

Meaning If you use a clear-text password, Junos OS displays the password as an encrypted string so that users viewing the configuration cannot see the unencrypted password. That is, as you enter the password in plain text, Junos OS encrypts it immediately. You do not have to configure Junos OS to encrypt the password as in some other systems. Plain-text passwords are hidden and marked as ## SECRET-DATA in the configuration.

Related Documentation

- *root-authentication*
- [Special Requirements for Junos OS Plain-Text Passwords on page 1824](#)
- *Configuring Special Requirements for Plain-Text Passwords*
- *Changing the Requirements for Junos OS Plain-Text Passwords*

Example: Configuring SSH Authentication for Root Logins

The following example shows how to configure two public DSA keys for SSH authentication of root logins:

```
[edit system]
root-authentication {
  encrypted-password "$1$1wp5tqMX$uy/u5H7OdXTwfWTmeJWXe/";
  ## SECRET-DATA;
  ssh-dsa "2354 95 9304@boojum.per";
  ssh-dsa "0483 02 8362@ecbatana.per";
}
```

- Related Documentation**
- [Configuring the Root Password on page 145](#)
 - [Special Requirements for Junos OS Plain-Text Passwords on page 1824](#)

Understanding Troubleshooting Resources

This topic describes some of the troubleshooting resources available for the QFX Series or OCX Series. These resources include tools such as the Junos OS CLI, Junos Space applications, and the Advanced Insight Scripts (AI-Scripts).

Table 138 provides a list of some of the troubleshooting resources.

Table 138: Troubleshooting Resources on the QFX and OCX Series

Troubleshooting Resource	Description	Documentation
Chassis alarms	Chassis alarms indicate a failure on the switch or one of its components. A chassis alarm count is displayed on the LCD panel on the front of the switch.	<i>Chassis Alarm Messages on a QFX3500 Device</i>
Chassis Status LEDs and Fan Tray LEDs	A blinking amber Power, Fan, or Fan Tray LED indicates a hardware component error. A blinking amber Status LED indicates a software error.	<i>Chassis Status LEDs on a QFX3500 Device</i>
Interface alarms	A predefined alarm (red or yellow) for an interface type is triggered when an interface of that type goes down.	<i>Interface Alarm Messages</i>
System alarms	A predefined alarm is triggered by a missing rescue configuration or problem with the software license.	<i>Understanding Alarms</i>
System log messages	The system log includes details of system and user events, including errors. Specify the severity and type of system log messages you wish to view or save, and configure the output to be sent to local or remote hosts.	<ul style="list-style-type: none"> • Overview of Single-Chassis System Logging Configuration on page 1528 • Junos OS System Log Configuration Statements on page 1530

Table 138: Troubleshooting Resources on the QFX and OCX Series (*continued*)

Troubleshooting Resource	Description	Documentation
Junos OS operational mode commands	Operational mode commands can be used to monitor switch performance and current activity on the network. For example, use the traceroute monitor command to locate points of failure in a network.	<ul style="list-style-type: none"> • Monitoring System Process Information on page 165 • Monitoring System Properties on page 166 • traceroute monitor
Junos OS automation scripts (event scripts)	Event scripts can be used to automate network troubleshooting and management tasks.	<i>Automation Scripting Feature Guide</i>
Junos OS XML operational tags	XML operational tags are equivalent in function to operational mode commands in the CLI, which you can use to retrieve status information for a device.	<i>Junos XML API Operational Developer Reference</i>
NETCONF XML management protocol	The NETCONF XML management protocol defines basic operations that are equivalent to Junos OS CLI configuration mode commands. Client applications use the protocol operations to display, edit, and commit configuration statements (among other operations), just as administrators use CLI configuration mode commands such as show , set , and commit to perform those operations.	<i>NETCONF XML Management Protocol Developer Guide</i>
SNMP MIBs and traps	MIBs enable the monitoring of network devices from a central location. For example, use the Traceroute MIB to monitor devices remotely.	<ul style="list-style-type: none"> • SNMP MIBs Support on page 1468 • SNMP Traps Support on page 1484 • Using the Traceroute MIB for SNMP Remote Operations
AI-Scripts and Advanced Insight Manager (AIM)	AI-Scripts installed on the switch can automatically detect and monitor faults on the switch, and depending on the configuration on the AIM application, send notifications of potential problems and submit problem reports to Juniper Support Systems.	Advanced Insight Scripts (AI-Scripts) Release Notes
Junos Space Service Now	This application enables you to display and manage information about problem events. When problems are detected on the switch by Advanced Insight Scripts (AI-Scripts) that are installed on the switch, the data is collected and sent to Service Now for your review and action.	<i>Service Automation</i>

Table 138: Troubleshooting Resources on the QFX and OCX Series (*continued*)

Troubleshooting Resource	Description	Documentation
Junos Space Service Insight	This application helps in accelerating operational analysis and managing the exposure to known issues. You can identify devices that are nearing their End Of Life (EOL) and also discover and prevent issues that could occur in your network. The functionality of Service Insight is dependent on the information sent from Service Now.	<i>Service Automation</i>
Juniper Networks Knowledge Base	You can search in this database for Juniper Networks product information, including alerts and troubleshooting tips.	http://kb.juniper.net

Troubleshooting Overview

This topic provides a general guide to troubleshooting some typical problems you may encounter on your QFX Series or OCX Series product.

[Table 139](#) provides a list of problem categories, summary of the symptom or problem, and recommended actions with links to the troubleshooting documentation.

Table 139: Troubleshooting on the QFX Series

Problem Category	Symptom or Problem	Recommended Action
Switch hardware components	LCD panel shows a chassis alarm count.	<i>See Chassis Alarm Messages on a QFX3500 Device.</i>
	Fan tray LED is blinking amber.	<i>See Fan Tray LED on a QFX3500 Device.</i>
	Chassis status LED for the power is blinking amber.	<i>See Chassis Status LEDs on a QFX3500 Device.</i>
	Chassis status LED for the fan (on the management board) is blinking amber.	Replace the management board as soon as possible. <i>See Chassis Status LEDs on a QFX3500 Device.</i>

Table 139: Troubleshooting on the QFX Series (*continued*)

Problem Category	Symptom or Problem	Recommended Action
Port configuration	Cannot configure a port as a Gigabit Ethernet port.	<p>Check whether the port is a valid Gigabit Ethernet port (6 through 41).</p> <p>See <i>QFX3500 Device Overview</i>.</p>
	Cannot configure a port as a Fibre Channel port.	<p>Check whether the port is a valid Fibre Channel port (0 through 5 and 42 through 47).</p> <p>See <i>QFX3500 Device Overview</i>.</p>
	Cannot configure a port as a 10-Gigabit Ethernet port.	<p>If the port is not a 40-Gbps QSFP+ interface, check whether the port is in the range of 0 through 5 or 42 through 47. If one of the ports in that block (0 through 5 or 42 through 47) is configured as a Fibre Channel port, then all ports in that block must also be configured as Fibre Channel ports.</p> <p>If the port is a 40-Gbps QSFP+ interface, make sure the configuration does not exceed the interface limit. Each 40-Gbps QSFP+ interface can be split into four 10-Gigabit Ethernet interfaces, but because port 0 is reserved, so you can only configure an additional fifteen 10-Gigabit Ethernet interfaces.</p> <p>See <i>QFX3500 Device Overview</i>.</p>
	Cannot configure a 40-Gbps QSFP+ interface.	<p>The 40-Gbps QSFP+ interfaces can only be used as 10-Gigabit Ethernet interfaces. Each 40-Gbps QSFP+ interface can be split into four 10-Gigabit Ethernet interfaces using a breakout cable. However, port 0 is reserved, so you can only configure an additional fifteen 10-Gigabit Ethernet interfaces.</p> <p>See <i>QFX3500 Device Overview</i>.</p>
External devices (USB devices)	Upgrading software from a USB device results in an upgrade failure, and the system enters an invalid state.	Unplug the USB device and reboot the switch.
Initial device configuration	Cannot configure management Ethernet ports.	<p>Configure the management ports from the console port. You cannot configure the management ports by directly connecting to them.</p> <p>NOTE: The management ports are on the front panel of the QFX3500 switch. They are labeled C0 and C1 on the front panel. In the CLI they are referred to as me0 and me1.</p> <p>See <i>Configuring a QFX3500 Device as a Standalone Switch</i>.</p>

Table 139: Troubleshooting on the QFX Series (*continued*)

Problem Category	Symptom or Problem	Recommended Action
Software upgrade and configuration	Failed software upgrade.	See “Recovering from a Failed Software Installation” on page 160.
	Active partition becomes inactive after upgrade.	
	Problem with the active configuration file.	See the following topics: <ul style="list-style-type: none"> • Loading a Previous Configuration File on page 1330 • Reverting to the Default Factory Configuration on page 1344 • Reverting to the Rescue Configuration on page 1346 • Performing a Recovery Installation
	Root password is lost or forgotten.	Recover the root password. See “Recovering the Root Password” on page 147.
Network interfaces	An aggregated Ethernet interface is down.	See “Troubleshooting an Aggregated Ethernet Interface” on page 2886.
	Interface on built-in network port is down.	See “Troubleshooting Network Interfaces” on page 2845.
	Interface on port in which SFP or SFP+ transceiver is installed in an SFP+ uplink module is down.	
Ethernet switching	A MAC address entry in the Ethernet switching table is not updated after the device with that MAC address has been moved from one interface to another on the switch.	See “Troubleshooting Ethernet Switching” on page 2134.
Firewall filter	Firewall configuration exceeded available Ternary Content Addressable Memory (TCAM) space.	See “Troubleshooting Firewall Filters” on page 5989.

Recovering the Root Password

If you forget the root password, you can use the password recovery procedure to reset the root password.



NOTE: The root password cannot be recovered on a QFabric system.



NOTE: You need console access to the switch to recover the root password.

To recover the root password:

1. Power off the switch by switching off the AC power outlet of the device or, if necessary, by pulling the power cords out of the device's power supplies.
2. Turn off the power to the management device, such as a PC or laptop computer, that you want to use to access the CLI.
3. Plug one end of the Ethernet rollover cable supplied with the device into the RJ-45-to-DB-9 serial port adapter supplied with the device.
4. Plug the RJ-45-to-DB-9 serial port adapter into the serial port on the management device.
5. Connect the other end of the Ethernet rollover cable to the console port on the device.
6. Turn on the power to the management device.
7. On the management device, start your asynchronous terminal emulation application (such as Microsoft Windows Hyperterminal) and select the appropriate **COM** port to use (for example, **COM1**).
8. Configure the port settings as follows:
 - Bits per second: 9600
 - Data bits: 8
 - Parity: None
 - Stop bits: 1
 - Flow control: None
9. Power on the device by (if necessary) plugging the power cords into the device's power supply, or turning on the power to the device by switching on the AC power outlet the device is plugged into.

The terminal emulation screen on your management device displays the device's boot sequence.

10. When the following prompt appears, press the Spacebar to access the device's bootstrap loader command prompt:

Hit [Enter] to boot immediately, or space bar for command prompt.
Booting [kernel] in 9 seconds...
11. At the following prompt, enter **boot -s** to start up the system in single-user mode.

ok **boot -s**
12. At the following prompt, enter **recovery** to start the root password recovery procedure.

Enter full pathname of shell or 'recovery' for root password recovery or RETURN for /bin/sh: **recovery**
13. Enter configuration mode in the CLI.
14. Set the root password. For example:

user@switch# **set system root-authentication plain-text-password**
15. At the following prompt, enter the new root password. For example:

New password: **juniper1**
Retype new password:

16. At the second prompt, reenter the new root password.
17. After you have finished configuring the password, commit the configuration.

```
root@host# commit
commit complete
```

18. Exit configuration mode in the CLI.
19. Exit operational mode in the CLI.
20. At the prompt, enter **y** to reboot the device.

```
Reboot the system? [y/n] y
```

Related Documentation

- *Configuring the Root Password*

CHAPTER 62

Configuring and Managing User Accounts

- [Junos OS User Accounts Overview on page 1821](#)
- [Junos OS Login Classes Overview on page 1823](#)
- [Special Requirements for Junos OS Plain-Text Passwords on page 1824](#)
- [Regular Expressions for Allowing and Denying Junos OS Configuration Mode Hierarchies on page 1826](#)
- [Regular Expressions for Allowing and Denying Junos OS Operational Mode Commands on page 1827](#)
- [Examples of Defining Access Privileges Using allow/deny-configuration Statements on page 1829](#)
- [Example: Configuring User Accounts on page 1830](#)
- [Example: Configuring User Permissions with Access Privilege Levels on page 1831](#)
- [Example: Configuring User Permissions with Access Privileges for Operational Mode Commands on page 1832](#)
- [Example: Changing the Requirements for Junos OS Plain-Text Passwords on page 1834](#)
- [Example: Limiting the Number of Login Attempts for SSH and Telnet Sessions to Prevent Unauthorized Access on page 1835](#)
- [Understanding Troubleshooting Resources on page 1836](#)
- [Troubleshooting Overview on page 1838](#)
- [Recovering the Root Password on page 1840](#)

Junos OS User Accounts Overview

User accounts provide one way for users to access the switch. (Users can access the switch without accounts if you configured RADIUS or TACACS+ servers, as described in [“Junos OS User Authentication Methods” on page 1804](#).) For each account, you define the login name for the user and, optionally, information that identifies the user. After you have created an account, the software creates a home directory for the user.

For each user account, you can define the following:

- Username—(Optional) Name that identifies the user. It must be unique within the switch. Do not include spaces, colons, or commas in the username. The username can be up to 64 characters long.
- User's full name—(Optional) If the full name contains spaces, enclose it in quotation marks. Do not include colons or commas.
- User identifier (UID)—(Optional) Numeric identifier that is associated with the user account name. The identifier must be in the range from 100 through 64,000 and must be unique within the switch. If you do not assign a UID to a username, the software assigns one when you commit the configuration, preferring the lowest available number.
- You must ensure that the UID is unique. However, it is possible to assign the same UID to different users. If you do this, the CLI displays a warning when you commit the configuration and then assigns the duplicate UID.
- User's access privilege—(Required) One of the login classes you defined in the **class** statement at the **[edit system login]** hierarchy level, or one of the default classes listed in ["Regular Expressions for Allowing and Denying Junos OS Configuration Mode Hierarchies"](#) on page 1826.
- Authentication method or methods and passwords that the user can use to access the switch—(Optional) You can use SSH or a Message Digest 5 (MD5) password, or you can enter a plain-text password that Junos OS encrypts using MD5-style encryption before entering it in the password database. For each method, you can specify the user's password. If you configure the **plain-text-password** option, you are prompted to enter and confirm the password:

```
[edit system login user user-name]  
user@switch# set authentication plain-text-password  
New password: type password here  
Retype new password: retype password here
```

The default requirements for plain-text passwords are:

- The password must be between 6 and 128 characters long
 - You can include most character classes in a password (uppercase letters, lowercase letters, numbers, punctuation marks, and other special characters). Control characters are not recommended.
 - Valid passwords must contain at least one change of case or character class.

For each user account and for root logins, you can configure more than one public RSA or DSA key for user authentication. When a user logs in using a user account or as root, the configured public keys are referenced to determine whether the private key matches any of them.

For SSH authentication, you can also copy the contents of an SSH key file into the configuration.

To load an SSH key file, use the **load-key-file** statement. This statement loads RSA (SSH version 1 and SSH version 2) and DSA (SSH version 2) public keys.

If you load the SSH keys file, the contents of the file are copied into the configuration immediately after you enter the **load-key-file** statement. To view the SSH key entries, use the configuration mode **show** command. For example:

```
[edit system login user boojum]
user@switch# set authentication load-key-file my-host:.ssh/identity.pub
.file.19692 | 0 KB | 0.3 kB/s | ETA: 00:00:00 | 100%
[edit system]
user@switch# show
root-authentication {
  ssh-rsa "1024 35 9727638204084251055468226757249864241630322
207404962528390382038690141584534964170019610608358722961563
475784918273603361276441874265946893207739108344813125957722
625461667999278316123500438660915866283822489746732605661192
181489539813862940327687806538169602027491641637359132693963
44008443 boojum@juniper.net"; # SECRET-DATA
}
```

An account for the user **root** is always present in the configuration. You configure the password for **root** using the **root-authentication** statement, as described in *Configuring the Root Password*.

Junos-FIPS and Common Criteria have special password requirements. FIPS and Common Criteria passwords must be between 10 and 20 characters in length. Passwords must use at least three of the five defined character sets (uppercase letters, lowercase letters, digits, punctuation marks, and other special characters). If Junos-FIPS is installed on the switch, you cannot configure passwords unless they meet this standard.

- Related Documentation**
- [Configuring Junos OS User Accounts on page 125](#)
 - [Junos OS Login Classes Overview on page 1823](#)

Junos OS Login Classes Overview

All users who can log in to the router or switch must be in a login class. With login classes, you define the following:

- Access privileges that users have when they are logged in to the router or switch
- Commands and statements that users can and cannot specify
- How long a login session can be idle before it times out and the user is logged out

You can define any number of login classes and then apply one login class to an individual user account.

The Junos operating system (Junos OS) contains a few predefined login classes, which are listed in [Table 140](#). The predefined login classes cannot be modified.

Table 140: Predefined System Login Classes

Login Class	Permission Flag Set
operator	clear, network, reset, trace, and view

Table 140: Predefined System Login Classes (*continued*)

Login Class	Permission Flag Set
read-only	view
superuser or super-user	all
unauthorized	None

**NOTE:**

- You cannot modify a predefined login class name. If you issue the `set` command on a predefined class name, the Junos OS appends `-local` to the login class name. The following message also appears:

warning: '<class-name>' is a predefined class name; changing to
'<class-name>-local'

- You cannot issue the `rename` or `copy` command on a predefined login class. Doing so results in the following error message:

error: target '<class-name>' is a predefined class

**Related
Documentation**

- [Defining Junos OS Login Classes](#)

Special Requirements for Junos OS Plain-Text Passwords

Junos OS has special requirements when you create plain-text passwords on a router or switch. [Table 141](#) shows the default requirements.

Table 141: Special Requirements for Plain-Text Passwords

Junos OS	Junos-FIPS
The password must be between 6 and 128 characters long.	FIPS passwords must be between 10 and 20 characters long
You can include most character classes in a password (uppercase letters, lowercase letters, numbers, punctuation marks, and other special characters). Control characters are not recommended.	You can include most character classes in a password (uppercase letters, lowercase letters, numbers, punctuation marks, and other special characters). Control characters are not recommended.
Valid passwords must contain at least one change of case or character class.	Passwords must use at least three of the five defined character classes (uppercase letters, lowercase letters, numbers, punctuation marks, and other special characters).

You can change the requirements for plain-text passwords.

Junos OS supports the following five character classes for plain-text passwords:

- Lowercase letters
- Uppercase letters
- Numbers
- Punctuation
- Special characters: ! @ # \$ % ^ & *, + < >



NOTE: "!" and "," are punctuation characters, but are listed under "special characters".

Control characters are not recommended.

You can include the **plain-text-password** statement at the following hierarchy levels:

- [edit system diag-port-authentication]
- [edit system pic-console-authentication]
- [edit system root-authentication]
- [edit system login user *username* authentication]

The **change-type** statement specifies whether the password is checked for the following:

- The total number of character sets used (**character-set**)
- The total number of character set changes (**set-transitions**)

For example, the following password:

MyPassWd@2

has four character sets (uppercase letters, lowercase letters, special characters, and numbers) and seven character set changes (M–y, y–P, P–a, s–W, W–d, d–@, and @–2).

The **change-type** statement is optional. If you omit the **change-type** option, Junos-FIPS plain-text passwords are checked for character sets, and Junos OS plain-text passwords are checked for character set changes.

The **minimum-changes** statement specifies how many character sets or character set changes are required for the password. This statement is optional. If you do not use the **minimum-changes** statement, character sets are not checked for Junos OS. If the **change-type** statement is configured for the **character-set** option, then the **minimum-changes** value must be 5 or less, because Junos OS only supports five character sets.

The **format** statement specifies the hash algorithm (**md5**, **sha1**, **sha256**, **sha512** or **des**) for authenticating plain-text passwords. This statement is optional. For Junos OS, the default format is **md5**. For Junos-FIPS, only **sha1** is supported.



NOTE: Starting with Junos OS Release 13.3, the sha1 does not enable secure, protected specification of passwords and we recommend that you do not use the sha1 algorithm to configure passwords. Instead, you can use the sha256 or sha512 to specify passwords by using the 256-bit and 512-bit cryptographic hash algorithm respectively for a robust and reliable operation.

The **maximum-length** statement specifies the maximum number of characters allowed in a password. This statement is optional. By default, Junos OS passwords have no maximum; however, only the first 128 characters are significant. Junos-FIPS passwords must be 20 characters or less. The range for Junos OS maximum-length passwords is from 20 to 128 characters.

The **minimum-length** statement specifies the minimum number of characters required for a password. This statement is optional. By default, Junos OS passwords must be at least 6 characters long, and Junos-FIPS passwords must be at least 10 characters long. The range is from 6 to 20 characters.

Changes to password requirements do not take effect until the configuration is committed. When requirements change, only newly created, plain-text passwords are checked; existing passwords are not checked against the new requirements.

The default configuration for Junos OS plain-text passwords is:

```
[edit system login]
passwords {
  change-type character-sets;
  format md5;
  minimum-changes 1;
  minimum-length 6;
}
```

The default configuration for Junos-FIPS plain-text passwords is:

```
[edit system login]
passwords {
  change-type set-transitions;
  format sha1;
  maximum-length 20;
  minimum-changes 3;
  minimum-length 10;
}
```

- Related Documentation**
- [Changing the Requirements for Junos OS Plain-Text Passwords](#)
 - [Configuring the Root Password on page 145](#)

Regular Expressions for Allowing and Denying Junos OS Configuration Mode Hierarchies

Use extended regular expressions to specify which configuration mode hierarchies are denied or allowed. You specify these regular expressions in the **allow/deny-configuration-regexps** and **allow/deny-configuration** statements at the **[edit**

system login class] hierarchy level, or by specifying Juniper Networks vendor-specific TACACS+ or RADIUS attributes in your authentication server's configuration. If regular expressions are received during TACACS+ or RADIUS authentication, they merge with any regular expressions configured on the local router or switch.

[Table 142](#) lists common regular expression operators that you can use for allowing or denying configuration mode .

Command regular expressions implement the extended (modern) regular expressions, as defined in POSIX 1003.2.

Table 142: Configuration Mode Hierarchies—Common Regular Expression Operators

Operator	Match
	One of two or more terms separated by the pipe. Each term must be a complete standalone expression enclosed in parentheses (), with no spaces between the pipe and the adjacent parentheses. For example, (show system alarms) (show system software).
^	At the beginning of an expression, used to denote where the command begins, where there might be some ambiguity.
\$	Character at the end of a command. Used to denote a command that must be matched exactly up to that point. For example, allow-commands "show interfaces\$" means that the user can issue the show interfaces command but cannot issue show interfaces detail or show interfaces extensive .
[]	Range of letters or digits. To separate the start and end of a range, use a hyphen (-).
()	A group of commands, indicating a complete, standalone expression to be evaluated the result is then evaluated as part of the overall expression. Parentheses must be used in conjunction with pipe operators as explained .
*	Zero or more terms.
+	One or more terms.
.	Any character except for a space " ".

Related Documentation

- *Specifying Access Privileges for Junos OS Configuration Mode Hierarchies*

Regular Expressions for Allowing and Denying Junos OS Operational Mode Commands

Use extended regular expressions to specify which operational mode commands are denied or allowed. [Table 143](#) lists common regular expression operators that can be used in the operational mode commands. Command regular expressions implement the extended (modern) regular expressions as defined in POSIX 1003.2.

Table 143: Common Regular Expression Operators to Allow or Deny Operational Mode Commands

Operator	Match
	One of two or more terms separated by the pipe () symbol. Each term must be a complete standalone expression enclosed in parentheses (), with no spaces between the pipe and the adjacent parentheses. For example, (show system alarms) (show system software).
^	At the beginning of an expression, used to denote where the command begins, and where there might be some ambiguity.
\$	Character at the end of a command. Used to denote a command that must be matched exactly up to that point. For example, allow-commands "show interfaces\$" means that the user can issue the show interfaces command but cannot issue the show interfaces detail or show interfaces extensive command.
[]	Range of letters or digits. To separate the start and end of a range, use a hyphen (-).
()	A group of commands, indicating a complete, standalone expression to be evaluated; the result is then evaluated as part of the overall expression. Parentheses must always be used in conjunction with pipe operators as explained above.

If a regular expression contains a syntax error, it becomes invalid, and although the user can log in, the permission granted or denied by the regular expression does not take effect. When regular expressions configured on TACACS+ or RADIUS servers merge with regular expressions configured on the router or switch, if the final expression has a syntax error, the overall result is an invalid regular expression. If a regular expression does not contain any operators, all varieties of the command are allowed. For example, if the following statement is included in the configuration, the user can issue the commands **show interfaces detail** and **show interfaces extensive** in addition to showing an individual interface:

```
allow-commands "show interfaces";
```

Related Documentation

- [Specifying Access Privileges for Junos OS Operational Mode Commands on page 30](#)

Examples of Defining Access Privileges Using allow/deny-configuration Statements

You can define access privileges using a combination of the following types of statements:

- permission flags
- **allow-configuration** and **deny-configuration** statements

The permission flags define the larger boundaries of what a person or login class can access and control. The **allow-configuration** and **deny-configuration** statements take precedence over permission flags and give the administrator finer control over exactly what the user has access to.

This topic explains defining access privileges using **allow-configuration** and **deny-configuration** statements by showing a series of examples of login class configuration using these statements. Examples 1 through 3 use both permission flags and **deny-configuration** statements to create login classes that allow users access to all except something. Each **allow-configuration** or **deny-configuration** statement is configured with one or more regular expressions to be allowed or denied.

The following examples show how to configure access privileges for individual configuration mode hierarchy levels.

If the following statement is included in the configuration and the user's login class permission bit is set to **all**, the user cannot configure telnet parameters:

```
[edit system login class class-name]
user@switch# set deny-configuration "system services telnet"
```

If the following statement is included in the configuration and the user's login class permission bit is set to **all**, the user cannot issue login class commands within any login class whose name begins with "m":

```
[edit system login class class-name]
user@switch# set deny-configuration "system login class m.*"
```

Example 3 This next example shows the creation of a login class with the **all** permission bit that prevents the user from editing a configuration or issuing commands (such as **commit**) at the **[edit system login class]** or **[edit system services]** hierarchy levels:

To create a login class that allows the user to configure everything except at the **[edit system login class]** or **[edit system services]** hierarchy levels:

1. Set the user's login class permission bit to **all**.

```
[edit system login]
user@host# set class all-except-login-class-or-system-services permissions all
```

If the following statement is included in the configuration and the user's login class permission bit is set to **all**, the user cannot edit a configuration or issue commands (such as **commit**) at the login class or system services hierarchy levels:

```
[edit system login class class-name]
user@switch# set deny-configuration "(system login class) | (system services)"
```

The following example shows how to configure permissions for individual configuration mode hierarchies:

```
[edit]
system {
  login { # This login class has operator privileges and the additional ability to edit
    # configuration at the system services hierarchy level.
    class only-system-services {
      permissions [ configure ];
      allow-configuration "system services";
    }
    # services commands.
    class all-except-system-services { # This login class has operator privileges but
      # cannot edit any system services configuration.
      permissions [ all ];
      deny-configuration "system services";
    }
  }
}
```

**Related
Documentation**

- *Specifying Access Privileges Using allow/deny-configuration Statements*
- *Specifying Access Privileges for Junos OS Configuration Mode Hierarchies*
- [Regular Expressions for Allowing and Denying Junos OS Configuration Mode Hierarchies on page 1826](#)
- *Defining Junos OS Login Classes*
- [Understanding Junos OS Access Privilege Levels on page 1799](#)

Example: Configuring User Accounts

The following example shows how to create accounts for four router or switch users, and create an account for the template user **remote**. All users use one of the default system login classes. User **alexander** also has two digital signal algorithm (DSA) public keys configured for SSH authentication.

```
[edit]
system {
  login {
    user philip {
      full-name "Philip of Macedonia";
      uid 1001;
      class super-user;
      authentication {
        encrypted-password "$1$poPPeY";
      }
    }
    user alexander {
      full-name "Alexander the Great";
      uid 1002;
      class view;
      authentication {
        encrypted-password "$1$14c5.$sBopasdFFdssdfFFdsdfs0";
      }
    }
  }
}
```

```

        ssh-dsa "8924 37 5678 5678@gaugamela.per";
        ssh-dsa "6273 94 9283@boojum.per";
    }
}
user darius {
    full-name "Darius King of Persia";
    uid 1003;
    class operator;
    authentication {
        ssh-rsa "1024 37 12341234@ecbatana.per";
    }
}
user anonymous {
    class unauthorized;
}
user remote {
    full-name "All remote users";
    uid 9999;
    class read-only;
}
}
}

```

- Related Documentation**
- *Junos OS User Accounts Overview*
 - *Limiting the Number of User Login Attempts for SSH and Telnet Sessions*

Example: Configuring User Permissions with Access Privilege Levels

Create two access privilege classes on the router or switch, one for configuring and viewing user accounts only and the second for configuring and viewing SNMP parameters only:

In this example, you create two custom login classes on the router or switch and assign access privileges to each class through permission flags. The first custom login class is called **user-accounts** and it only includes access privileges for configuring and viewing user accounts. The second custom login class is called **network-mgmt** and only includes access privileges for configuring SNMP parameters.

```

[edit]
system {
  login {
    class user-accounts {
      permissions [ configure admin admin-control ];
    }
    class network-mgmt {
      permissions [ configure snmp snmp-control ];
    }
  }
}

```

1. Create the **user-accounts** custom login class and give it control over user accounts with the **configure admin admin-control** permission flag.

```

[edit system login]
user@router# set class user-accounts permissions configure admin admin-control

```

2. Create the **network-mgmt** custom login class and use the **configure snmp snmp-control** permission flag to assign it SNMP configuration privileges.

```
[edit system login]
user@router# set class network-mgmt permissions configure snmp snmp-control
```

3. Check your configuration by using the **show system login** command.

```
user@router# show system login
class user-accounts {
  permissions [ configure admin admin-control ];
}
class network-mgmt {
  permissions [ configure snmp snmp-control ];
}
```

Related Documentation

- [Configuring Access Privilege Levels on page 1809](#)

Example: Configuring User Permissions with Access Privileges for Operational Mode Commands

Each operational mode command has an access privilege level associated with it. Access privileges control the commands that each custom login class can execute, configure, and view. Custom login classes are groups of users who are assigned with customized levels of access to different commands and statements. This ensures that each group of users can only use commands appropriate to their function, preventing unauthorized users from executing sensitive commands that could potentially cause damage to the network.

In this example, you create three custom login classes on the router or switch and assign access privileges for operational mode commands through the **allow-commands** and **deny-commands** settings. Each custom login class uses the same set of permission flags as the default login class **operator**, but the login class is allowed or denied certain operational mode commands. The first custom login class is called **operator-and-boot** and it has access to the **request system reboot** operational mode command. The second custom login class is called **operator-no-set** and it is denied access to any **set** commands. The third login class is called **operator-and-install-but-no-bgp** and it has access to the **request system software add** and **show route** operational mode commands, but it is denied access to the **show bgp** command.

```
[edit]
system {
  login {
    class operator-and-boot {
      permissions [ clear network reset trace view ];
      allow-commands "request system reboot";
    }
    class operator-no-set {
      permissions [ clear network reset trace view ];
      deny-commands "set";
    }
    class operator-and-install-but-no-bgp {
      permissions [ clear network reset trace view ];
      allow-commands "(request system software add)|(show route$)";
      deny-commands "show bgp";
    }
  }
}
```



```

    }
  }
}

```

1. Create the **operator-and-boot** custom login class, give it **operator** level permission flags, and authorize it to use the **request system reboot** command.

```

[edit system login]
user@router# set class operator-and-boot permissions clear network reset trace view
user@router# set class operator-and-boot allow-commands request system reboot

```

2. Create the **operator-no-set** custom login class, give it **operator** level permission flags, and deny it access to the **set** command.

```

[edit system login]
user@router# set class operator-no-set clear network reset trace view
user@router# set class operator-no-set deny-commands set

```

3. Create the **operator-and-install-but-no-bgp** custom login class, give it **operator** level permission flags, authorize it to use the **request system software add** and **show route** commands, and deny it access to the **show bgp** command.

```

[edit system login]
user@router# set class operator-and-install-but-no-bgp clear network reset trace view
user@router# set class operator-and-install-but-no-bgp request system software add show route
user@router# set class operator-and-install-but-no-bgp show bgp

```

4. Check your configuration by using the **show system login** command.

```

user@router# show system login
class operator-and-boot {
  permissions [ clear network reset trace view ];
  allow-commands "request system reboot";
}
class operator-no-set {
  permissions [ clear network reset trace view ];
  deny-commands "set";
}
class operator-and-install-but-no-bgp {
  permissions [ clear network reset trace view ];
  allow-commands "(request system software add)|(show route$)";
  deny-commands "show bgp";
}

```

Related Documentation

- [Specifying Access Privileges for Junos OS Operational Mode Commands on page 30](#)

Example: Changing the Requirements for Junos OS Plain-Text Passwords

This example shows how to set various maximum and minimum requirements for plain-text passwords to increase password strength.

- [Requirements on page 1834](#)
- [Overview on page 1834](#)
- [Configuration on page 1834](#)

Requirements

This example requires a device running Junos 12.2 or greater. The **minimum-length** and **maximum-length** password requirements statements are available in earlier releases, however, you must have Junos OS Release 12.2 or greater to configure **minimum-lower-cases**, **minimum-numeric**s, **minimum-punctuations**, or **minimum-upper-cases**.

Overview

You can use a variety of requirements to strengthen plain-text passwords for greater security. Junos OS provides a number of possible configurations at the **[edit system login password]** hierarchy level that allow you to require users to create plain-text passwords that conform to a particular set of requirements that may include such things as length, number of changes, type of characters, numbers, or letter case.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set system login password minimum-length 12
set system login password maximum-length 22
set system login password minimum-numeric 1
set system login password minimum-upper-cases 1
set system login password minimum-lower-cases 1
set system login password minimum-punctuations 1
```

Configuring Requirements for Plain-Text Passwords

Step-by-Step Procedure

This example configures password requirements that require the user to create a password that has a minimum length of 12 characters, a maximum length of 22 characters, and that includes at least one lower-case letter, at least one upper-case letter, at least one punctuation character, and at least one numeric character.

1. Navigate to configuration mode in the **[system login password]** hierarchy level.

```
user@host> edit
[edit]
user@host# edit system login password
```

2. Set a minimum length requirement of 12 characters and a maximum length requirement of 22 characters for user passwords.


```
[edit system login password]
user@host# set minimum-length 12
[edit system login password]
user@host# set maximum-length 22
```
3. Require users to set a password that has at least one lower-case letter and at least one upper-case letter.


```
[edit system login password]
user@host# set minimum-lower-cases 1
[edit system login password]
user@host# set minimum-upper-cases 1
```
4. Require users to set a password that has at least one punctuation-class character and at least one number.


```
[edit system login password]
user@host# set minimum-punctuations 1
[edit system login password]
user@host# set minimum-numeric 1
```

Results

From configuration mode, confirm your configuration by entering the show command at the edit system login password hierarchy level. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit system login password]
user@host# show
minimum-length 12;
maximum-length 22;
minimum-numeric 1;
minimum-upper-cases 1;
minimum-lower-cases 1;
```

- Related Documentation**
- [Special Requirements for Junos OS Plain-Text Passwords on page 1824](#)
 - *password (Login)*

Example: Limiting the Number of Login Attempts for SSH and Telnet Sessions to Prevent Unauthorized Access

Limiting the number of SSH and Telnet login attempts per user is one of the most effective methods of stopping brute force attacks from compromising your network security. Brute force attackers execute a large number of login attempts in a short period of time to illegitimately gain access to a private network. By configuring the **retry-options** command, you can create an increasing delay after each failed login attempt, eventually disconnecting any user who passes your set threshold of login attempts.

The following example shows how to limit the user to four attempts when the user enters a password while logging in through SSH or Telnet. Set the **backoff-threshold** to 2, the

back-off-factor to 5 seconds, and the **minimum-time** to 40 seconds. The user experiences a delay of 5 seconds after the second attempt to enter a correct password fails. After each subsequent failed attempt, the delay increases by 5 seconds. After the fourth and final failed attempt to enter a correct password, the user experiences an additional 10-second delay, and the connection closes after a total of 40 seconds.

The additional variables **maximum-time** and **lockout-period** are not set in this example.

```
[edit]
system {
  login {
    retry-options {
      backoff-threshold 2;
      backoff-factor 5;
      minimum-time 40;
      tries-before-disconnect 4;
    }
    password {
    }
  }
}
```



NOTE: This sample only shows the portion of the [edit system login] hierarchy level being modified.

Related Documentation

- *Limiting the Number of User Login Attempts for SSH and Telnet Sessions*
- *login*

Understanding Troubleshooting Resources

This topic describes some of the troubleshooting resources available for the QFX Series or OCX Series. These resources include tools such as the Junos OS CLI, Junos Space applications, and the Advanced Insight Scripts (AI-Scripts).

Table 138 provides a list of some of the troubleshooting resources.

Table 144: Troubleshooting Resources on the QFX and OCX Series

Troubleshooting Resource	Description	Documentation
Chassis alarms	Chassis alarms indicate a failure on the switch or one of its components. A chassis alarm count is displayed on the LCD panel on the front of the switch.	<i>Chassis Alarm Messages on a QFX3500 Device</i>
Chassis Status LEDs and Fan Tray LEDs	A blinking amber Power, Fan, or Fan Tray LED indicates a hardware component error. A blinking amber Status LED indicates a software error.	<i>Chassis Status LEDs on a QFX3500 Device</i>

Table 144: Troubleshooting Resources on the QFX and OCX Series (*continued*)

Troubleshooting Resource	Description	Documentation
Interface alarms	A predefined alarm (red or yellow) for an interface type is triggered when an interface of that type goes down.	<i>Interface Alarm Messages</i>
System alarms	A predefined alarm is triggered by a missing rescue configuration or problem with the software license.	<i>Understanding Alarms</i>
System log messages	The system log includes details of system and user events, including errors. Specify the severity and type of system log messages you wish to view or save, and configure the output to be sent to local or remote hosts.	<ul style="list-style-type: none"> • Overview of Single-Chassis System Logging Configuration on page 1528 • Junos OS System Log Configuration Statements on page 1530
Junos OS operational mode commands	Operational mode commands can be used to monitor switch performance and current activity on the network. For example, use the tracert monitor command to locate points of failure in a network.	<ul style="list-style-type: none"> • Monitoring System Process Information on page 165 • Monitoring System Properties on page 166 • tracert monitor
Junos OS automation scripts (event scripts)	Event scripts can be used to automate network troubleshooting and management tasks.	<i>Automation Scripting Feature Guide</i>
Junos OS XML operational tags	XML operational tags are equivalent in function to operational mode commands in the CLI, which you can use to retrieve status information for a device.	<i>Junos XML API Operational Developer Reference</i>
NETCONF XML management protocol	The NETCONF XML management protocol defines basic operations that are equivalent to Junos OS CLI configuration mode commands. Client applications use the protocol operations to display, edit, and commit configuration statements (among other operations), just as administrators use CLI configuration mode commands such as show , set , and commit to perform those operations.	<i>NETCONF XML Management Protocol Developer Guide</i>
SNMP MIBs and traps	MIBs enable the monitoring of network devices from a central location. For example, use the Traceroute MIB to monitor devices remotely.	<ul style="list-style-type: none"> • SNMP MIBs Support on page 1468 • SNMP Traps Support on page 1484 • Using the Traceroute MIB for SNMP Remote Operations
AI-Scripts and Advanced Insight Manager (AIM)	AI-Scripts installed on the switch can automatically detect and monitor faults on the switch, and depending on the configuration on the AIM application, send notifications of potential problems and submit problem reports to Juniper Support Systems.	Advanced Insight Scripts (AI-Scripts) Release Notes

Table 144: Troubleshooting Resources on the QFX and OCX Series (*continued*)

Troubleshooting Resource	Description	Documentation
Junos Space Service Now	This application enables you to display and manage information about problem events. When problems are detected on the switch by Advanced Insight Scripts (AI-Scripts) that are installed on the switch, the data is collected and sent to Service Now for your review and action.	<i>Service Automation</i>
Junos Space Service Insight	This application helps in accelerating operational analysis and managing the exposure to known issues. You can identify devices that are nearing their End Of Life (EOL) and also discover and prevent issues that could occur in your network. The functionality of Service Insight is dependent on the information sent from Service Now.	<i>Service Automation</i>
Juniper Networks Knowledge Base	You can search in this database for Juniper Networks product information, including alerts and troubleshooting tips.	http://kb.juniper.net

Troubleshooting Overview

This topic provides a general guide to troubleshooting some typical problems you may encounter on your QFX Series or OCX Series product.

[Table 139](#) provides a list of problem categories, summary of the symptom or problem, and recommended actions with links to the troubleshooting documentation.

Table 145: Troubleshooting on the QFX Series

Problem Category	Symptom or Problem	Recommended Action
Switch hardware components	LCD panel shows a chassis alarm count.	<i>See Chassis Alarm Messages on a QFX3500 Device.</i>
	Fan tray LED is blinking amber.	<i>See Fan Tray LED on a QFX3500 Device.</i>
	Chassis status LED for the power is blinking amber.	<i>See Chassis Status LEDs on a QFX3500 Device.</i>
	Chassis status LED for the fan (on the management board) is blinking amber.	Replace the management board as soon as possible. <i>See Chassis Status LEDs on a QFX3500 Device.</i>

Table 145: Troubleshooting on the QFX Series (*continued*)

Problem Category	Symptom or Problem	Recommended Action
Port configuration	Cannot configure a port as a Gigabit Ethernet port.	<p>Check whether the port is a valid Gigabit Ethernet port (6 through 41).</p> <p>See <i>QFX3500 Device Overview</i>.</p>
	Cannot configure a port as a Fibre Channel port.	<p>Check whether the port is a valid Fibre Channel port (0 through 5 and 42 through 47).</p> <p>See <i>QFX3500 Device Overview</i>.</p>
	Cannot configure a port as a 10-Gigabit Ethernet port.	<p>If the port is not a 40-Gbps QSFP+ interface, check whether the port is in the range of 0 through 5 or 42 through 47. If one of the ports in that block (0 through 5 or 42 through 47) is configured as a Fibre Channel port, then all ports in that block must also be configured as Fibre Channel ports.</p> <p>If the port is a 40-Gbps QSFP+ interface, make sure the configuration does not exceed the interface limit. Each 40-Gbps QSFP+ interface can be split into four 10-Gigabit Ethernet interfaces, but because port 0 is reserved, so you can only configure an additional fifteen 10-Gigabit Ethernet interfaces.</p> <p>See <i>QFX3500 Device Overview</i>.</p>
	Cannot configure a 40-Gbps QSFP+ interface.	<p>The 40-Gbps QSFP+ interfaces can only be used as 10-Gigabit Ethernet interfaces. Each 40-Gbps QSFP+ interface can be split into four 10-Gigabit Ethernet interfaces using a breakout cable. However, port 0 is reserved, so you can only configure an additional fifteen 10-Gigabit Ethernet interfaces.</p> <p>See <i>QFX3500 Device Overview</i>.</p>
External devices (USB devices)	Upgrading software from a USB device results in an upgrade failure, and the system enters an invalid state.	Unplug the USB device and reboot the switch.
Initial device configuration	Cannot configure management Ethernet ports.	<p>Configure the management ports from the console port. You cannot configure the management ports by directly connecting to them.</p> <p>NOTE: The management ports are on the front panel of the QFX3500 switch. They are labeled C0 and C1 on the front panel. In the CLI they are referred to as me0 and me1.</p> <p>See <i>Configuring a QFX3500 Device as a Standalone Switch</i>.</p>

Table 145: Troubleshooting on the QFX Series (*continued*)

Problem Category	Symptom or Problem	Recommended Action
Software upgrade and configuration	Failed software upgrade.	See “Recovering from a Failed Software Installation” on page 160 .
	Active partition becomes inactive after upgrade.	
	Problem with the active configuration file.	See the following topics: <ul style="list-style-type: none"> • Loading a Previous Configuration File on page 1330 • Reverting to the Default Factory Configuration on page 1344 • Reverting to the Rescue Configuration on page 1346 • Performing a Recovery Installation
	Root password is lost or forgotten.	Recover the root password. See “Recovering the Root Password” on page 147 .
Network interfaces	An aggregated Ethernet interface is down.	See “Troubleshooting an Aggregated Ethernet Interface” on page 2886 .
	Interface on built-in network port is down.	See “Troubleshooting Network Interfaces” on page 2845 .
	Interface on port in which SFP or SFP+ transceiver is installed in an SFP+ uplink module is down.	
Ethernet switching	A MAC address entry in the Ethernet switching table is not updated after the device with that MAC address has been moved from one interface to another on the switch.	See “Troubleshooting Ethernet Switching” on page 2134 .
Firewall filter	Firewall configuration exceeded available Ternary Content Addressable Memory (TCAM) space.	See “Troubleshooting Firewall Filters” on page 5989 .

Recovering the Root Password

If you forget the root password, you can use the password recovery procedure to reset the root password.



NOTE: The root password cannot be recovered on a QFabric system.



NOTE: You need console access to the switch to recover the root password.

To recover the root password:

1. Power off the switch by switching off the AC power outlet of the device or, if necessary, by pulling the power cords out of the device's power supplies.
2. Turn off the power to the management device, such as a PC or laptop computer, that you want to use to access the CLI.
3. Plug one end of the Ethernet rollover cable supplied with the device into the RJ-45-to-DB-9 serial port adapter supplied with the device.
4. Plug the RJ-45-to-DB-9 serial port adapter into the serial port on the management device.
5. Connect the other end of the Ethernet rollover cable to the console port on the device.
6. Turn on the power to the management device.
7. On the management device, start your asynchronous terminal emulation application (such as Microsoft Windows Hyperterminal) and select the appropriate **COM** port to use (for example, **COM1**).
8. Configure the port settings as follows:
 - Bits per second: 9600
 - Data bits: 8
 - Parity: None
 - Stop bits: 1
 - Flow control: None
9. Power on the device by (if necessary) plugging the power cords into the device's power supply, or turning on the power to the device by switching on the AC power outlet the device is plugged into.

The terminal emulation screen on your management device displays the device's boot sequence.

10. When the following prompt appears, press the Spacebar to access the device's bootstrap loader command prompt:


```
Hit [Enter] to boot immediately, or space bar for command prompt.
Booting [kernel] in 9 seconds...
```
11. At the following prompt, enter **boot -s** to start up the system in single-user mode.


```
ok boot -s
```
12. At the following prompt, enter **recovery** to start the root password recovery procedure.


```
Enter full pathname of shell or 'recovery' for root password recovery or RETURN
for /bin/sh: recovery
```
13. Enter configuration mode in the CLI.
14. Set the root password. For example:


```
user@switch# set system root-authentication plain-text-password
```
15. At the following prompt, enter the new root password. For example:


```
New password: juniper1
Retype new password:
```

16. At the second prompt, reenter the new root password.
17. After you have finished configuring the password, commit the configuration.

```
root@host# commit
commit complete
```

18. Exit configuration mode in the CLI.
19. Exit operational mode in the CLI.
20. At the prompt, enter **y** to reboot the device.

```
Reboot the system? [y/n] y
```

Related Documentation

- *Configuring the Root Password*

PART 25

Configuring Authentication

- [Configuring and Managing Local Password Authentication on page 1845](#)
- [Configuring and Managing TACACS+ Authentication on page 1859](#)
- [Configuring and Managing RADIUS Authentication on page 1869](#)
- [Configuring and Managing RADIUS Accounting on page 1883](#)
- [Configuring and Managing RADIUS Template Accounts on page 1899](#)
- [Configuring and Managing VSAs for RADIUS and TACACS+ on page 1901](#)

Configuring and Managing Local Password Authentication

- [Junos OS User Accounts Overview on page 1845](#)
- [Junos OS User Authentication Methods on page 1847](#)
- [Junos OS Login Classes Overview on page 1848](#)
- [Regular Expressions for Allowing and Denying Junos OS Configuration Mode Hierarchies on page 1848](#)
- [Regular Expressions for Allowing and Denying Junos OS Operational Mode Commands on page 1849](#)
- [Special Requirements for Junos OS Plain-Text Passwords on page 1850](#)
- [Configuring Junos OS User Accounts on page 1853](#)
- [Configuring a Local Administrator Account on page 1853](#)
- [Example: Creating Login Classes with Specific Privileges on page 1854](#)
- [Example: Configuring System Authentication for RADIUS, TACACS+, and Password Authentication on page 1855](#)
- [Example: Changing the Requirements for Junos OS Plain-Text Passwords on page 1857](#)

Junos OS User Accounts Overview

User accounts provide one way for users to access the switch. (Users can access the switch without accounts if you configured RADIUS or TACACS+ servers, as described in [“Junos OS User Authentication Methods” on page 1804](#).) For each account, you define the login name for the user and, optionally, information that identifies the user. After you have created an account, the software creates a home directory for the user.

For each user account, you can define the following:

- **Username**—(Optional) Name that identifies the user. It must be unique within the switch. Do not include spaces, colons, or commas in the username. The username can be up to 64 characters long.
- **User's full name**—(Optional) If the full name contains spaces, enclose it in quotation marks. Do not include colons or commas.

- User identifier (UID)—(Optional) Numeric identifier that is associated with the user account name. The identifier must be in the range from 100 through 64,000 and must be unique within the switch. If you do not assign a UID to a username, the software assigns one when you commit the configuration, preferring the lowest available number.
- You must ensure that the UID is unique. However, it is possible to assign the same UID to different users. If you do this, the CLI displays a warning when you commit the configuration and then assigns the duplicate UID.
- User's access privilege—(Required) One of the login classes you defined in the **class** statement at the **[edit system login]** hierarchy level, or one of the default classes listed in ["Regular Expressions for Allowing and Denying Junos OS Configuration Mode Hierarchies"](#) on page 1826.
- Authentication method or methods and passwords that the user can use to access the switch—(Optional) You can use SSH or a Message Digest 5 (MD5) password, or you can enter a plain-text password that Junos OS encrypts using MD5-style encryption before entering it in the password database. For each method, you can specify the user's password. If you configure the **plain-text-password** option, you are prompted to enter and confirm the password:

```
[edit system login user user-name]
user@switch# set authentication plain-text-password
New password: type password here
Retype new password: retype password here
```

The default requirements for plain-text passwords are:

- The password must be between 6 and 128 characters long
 - You can include most character classes in a password (uppercase letters, lowercase letters, numbers, punctuation marks, and other special characters). Control characters are not recommended.
- Valid passwords must contain at least one change of case or character class.

For each user account and for root logins, you can configure more than one public RSA or DSA key for user authentication. When a user logs in using a user account or as root, the configured public keys are referenced to determine whether the private key matches any of them.

For SSH authentication, you can also copy the contents of an SSH key file into the configuration.

To load an SSH key file, use the **load-key-file** statement. This statement loads RSA (SSH version 1 and SSH version 2) and DSA (SSH version 2) public keys.

If you load the SSH keys file, the contents of the file are copied into the configuration immediately after you enter the **load-key-file** statement. To view the SSH key entries, use the configuration mode **show** command. For example:

```
[edit system login user boojum]
user@switch# set authentication load-key-file my-host:.ssh/identity.pub
.file.19692 | 0 KB | 0.3 kB/s | ETA: 00:00:00 | 100%
```

```
[edit system]
user@switch# show
root-authentication {
  ssh-rsa "1024 35 9727638204084251055468226757249864241630322
207404962528390382038690141584534964170019610608358722961563
475784918273603361276441874265946893207739108344813125957722
625461667999278316123500438660915866283822489746732605661192
181489539813862940327687806538169602027491641637359132693963
44008443 boojum@juniper.net"; # SECRET-DATA
}
```

An account for the user **root** is always present in the configuration. You configure the password for **root** using the **root-authentication** statement, as described in *Configuring the Root Password*.

Junos-FIPS and Common Criteria have special password requirements. FIPS and Common Criteria passwords must be between 10 and 20 characters in length. Passwords must use at least three of the five defined character sets (uppercase letters, lowercase letters, digits, punctuation marks, and other special characters). If Junos-FIPS is installed on the switch, you cannot configure passwords unless they meet this standard.

- Related Documentation**
- [Configuring Junos OS User Accounts on page 125](#)
 - [Junos OS Login Classes Overview on page 1823](#)

Junos OS User Authentication Methods

The Junos OS supports three methods of user authentication: local password authentication, Remote Authentication Dial-In User Service (RADIUS), and Terminal Access Controller Access Control System Plus (TACACS+).

With local password authentication, you configure a password for each user allowed to log in to the router or switch.

RADIUS and TACACS+ are authentication methods for validating users who attempt to access the router or switch using telnet. They are both distributed client-server systems—the RADIUS and TACACS+ clients run on the router or switch, and the server runs on a remote network system.

You can configure the router or switch to be both a RADIUS and TACACS+ client, and you can also configure authentication passwords in the Junos OS configuration file. You can prioritize the methods to configure the order in which the software tries the different authentication methods when verifying user access.

- Related Documentation**
- [Configuring RADIUS Server Authentication](#)
 - [Configuring TACACS+ Authentication](#)
 - [Junos OS Authentication Order for RADIUS, TACACS+, and Password Authentication on page 1859](#)

Junos OS Login Classes Overview

All users who can log in to the router or switch must be in a login class. With login classes, you define the following:

- Access privileges that users have when they are logged in to the router or switch
- Commands and statements that users can and cannot specify
- How long a login session can be idle before it times out and the user is logged out

You can define any number of login classes and then apply one login class to an individual user account.

The Junos operating system (Junos OS) contains a few predefined login classes, which are listed in [Table 140](#). The predefined login classes cannot be modified.

Table 146: Predefined System Login Classes

Login Class	Permission Flag Set
operator	clear, network, reset, trace, and view
read-only	view
superuser or super-user	all
unauthorized	None



NOTE:

- You cannot modify a predefined login class name. If you issue the `set` command on a predefined class name, the Junos OS appends `-local` to the login class name. The following message also appears:

warning: '<class-name>' is a predefined class name; changing to '<class-name>-local'

- You cannot issue the `rename` or `copy` command on a predefined login class. Doing so results in the following error message:

error: target '<class-name>' is a predefined class

Related Documentation

- [Defining Junos OS Login Classes](#)

Regular Expressions for Allowing and Denying Junos OS Configuration Mode Hierarchies

Use extended regular expressions to specify which configuration mode hierarchies are denied or allowed. You specify these regular expressions in the `allow/deny-configuration-regexps` and `allow/deny-configuration` statements at the `[edit`

system login class] hierarchy level, or by specifying Juniper Networks vendor-specific TACACS+ or RADIUS attributes in your authentication server's configuration. If regular expressions are received during TACACS+ or RADIUS authentication, they merge with any regular expressions configured on the local router or switch.

[Table 142](#) lists common regular expression operators that you can use for allowing or denying configuration mode .

Command regular expressions implement the extended (modern) regular expressions, as defined in POSIX 1003.2.

Table 147: Configuration Mode Hierarchies—Common Regular Expression Operators

Operator	Match
	One of two or more terms separated by the pipe. Each term must be a complete standalone expression enclosed in parentheses (), with no spaces between the pipe and the adjacent parentheses. For example, (show system alarms) (show system software).
^	At the beginning of an expression, used to denote where the command begins, where there might be some ambiguity.
\$	Character at the end of a command. Used to denote a command that must be matched exactly up to that point. For example, allow-commands "show interfaces\$" means that the user can issue the show interfaces command but cannot issue show interfaces detail or show interfaces extensive .
[]	Range of letters or digits. To separate the start and end of a range, use a hyphen (-).
()	A group of commands, indicating a complete, standalone expression to be evaluated the result is then evaluated as part of the overall expression. Parentheses must be used in conjunction with pipe operators as explained .
*	Zero or more terms.
+	One or more terms.
.	Any character except for a space " ".

Related Documentation

- *Specifying Access Privileges for Junos OS Configuration Mode Hierarchies*

Regular Expressions for Allowing and Denying Junos OS Operational Mode Commands

Use extended regular expressions to specify which operational mode commands are denied or allowed. [Table 143](#) lists common regular expression operators that can be used in the operational mode commands. Command regular expressions implement the extended (modern) regular expressions as defined in POSIX 1003.2.

Table 148: Common Regular Expression Operators to Allow or Deny Operational Mode Commands

Operator	Match
	One of two or more terms separated by the pipe () symbol. Each term must be a complete standalone expression enclosed in parentheses (), with no spaces between the pipe and the adjacent parentheses. For example, (show system alarms) (show system software).
^	At the beginning of an expression, used to denote where the command begins, and where there might be some ambiguity.
\$	Character at the end of a command. Used to denote a command that must be matched exactly up to that point. For example, allow-commands "show interfaces\$" means that the user can issue the show interfaces command but cannot issue the show interfaces detail or show interfaces extensive command.
[]	Range of letters or digits. To separate the start and end of a range, use a hyphen (-).
()	A group of commands, indicating a complete, standalone expression to be evaluated; the result is then evaluated as part of the overall expression. Parentheses must always be used in conjunction with pipe operators as explained above.

If a regular expression contains a syntax error, it becomes invalid, and although the user can log in, the permission granted or denied by the regular expression does not take effect. When regular expressions configured on TACACS+ or RADIUS servers merge with regular expressions configured on the router or switch, if the final expression has a syntax error, the overall result is an invalid regular expression. If a regular expression does not contain any operators, all varieties of the command are allowed. For example, if the following statement is included in the configuration, the user can issue the commands **show interfaces detail** and **show interfaces extensive** in addition to showing an individual interface:

```
allow-commands "show interfaces";
```

Related Documentation

- [Specifying Access Privileges for Junos OS Operational Mode Commands on page 30](#)

Special Requirements for Junos OS Plain-Text Passwords

Junos OS has special requirements when you create plain-text passwords on a router or switch. [Table 141](#) shows the default requirements.

Table 149: Special Requirements for Plain-Text Passwords

Junos OS	Junos-FIPS
The password must be between 6 and 128 characters long.	FIPS passwords must be between 10 and 20 characters long

Table 149: Special Requirements for Plain-Text Passwords (*continued*)

Junos OS	Junos-FIPS
You can include most character classes in a password (uppercase letters, lowercase letters, numbers, punctuation marks, and other special characters). Control characters are not recommended.	You can include most character classes in a password (uppercase letters, lowercase letters, numbers, punctuation marks, and other special characters). Control characters are not recommended.
Valid passwords must contain at least one change of case or character class.	Passwords must use at least three of the five defined character classes (uppercase letters, lowercase letters, numbers, punctuation marks, and other special characters).

You can change the requirements for plain-text passwords.

Junos OS supports the following five character classes for plain-text passwords:

- Lowercase letters
- Uppercase letters
- Numbers
- Punctuation
- Special characters: ! @ # \$ % ^ & * , + < >



NOTE: "!" and "," are punctuation characters, but are listed under "special characters".

Control characters are not recommended.

You can include the **plain-text-password** statement at the following hierarchy levels:

- [edit system diag-port-authentication]
- [edit system pic-console-authentication]
- [edit system root-authentication]
- [edit system login user *username* authentication]

The **change-type** statement specifies whether the password is checked for the following:

- The total number of character sets used (**character-set**)
- The total number of character set changes (**set-transitions**)

For example, the following password:

MyPassWd@2

has four character sets (uppercase letters, lowercase letters, special characters, and numbers) and seven character set changes (M–y, y–P, P–a, s–W, W–d, d–@, and @–2).

The **change-type** statement is optional. If you omit the **change-type** option, Junos-FIPS plain-text passwords are checked for character sets, and Junos OS plain-text passwords are checked for character set changes.

The **minimum-changes** statement specifies how many character sets or character set changes are required for the password. This statement is optional. If you do not use the **minimum-changes** statement, character sets are not checked for Junos OS. If the **change-type** statement is configured for the **character-set** option, then the **minimum-changes** value must be 5 or less, because Junos OS only supports five character sets.

The **format** statement specifies the hash algorithm (**md5**, **sha1**, **sha256**, **sha512** or **des**) for authenticating plain-text passwords. This statement is optional. For Junos OS, the default format is **md5**. For Junos-FIPS, only **sha1** is supported.



NOTE: Starting with Junos OS Release 13.3, the **sha1** does not enable secure, protected specification of passwords and we recommend that you do not use the **sha1** algorithm to configure passwords. Instead, you can use the **sha256** or **sha512** to specify passwords by using the 256-bit and 512-bit cryptographic hash algorithm respectively for a robust and reliable operation.

The **maximum-length** statement specifies the maximum number of characters allowed in a password. This statement is optional. By default, Junos OS passwords have no maximum; however, only the first 128 characters are significant. Junos-FIPS passwords must be 20 characters or less. The range for Junos OS maximum-length passwords is from 20 to 128 characters.

The **minimum-length** statement specifies the minimum number of characters required for a password. This statement is optional. By default, Junos OS passwords must be at least 6 characters long, and Junos-FIPS passwords must be at least 10 characters long. The range is from 6 to 20 characters.

Changes to password requirements do not take effect until the configuration is committed. When requirements change, only newly created, plain-text passwords are checked; existing passwords are not checked against the new requirements.

The default configuration for Junos OS plain-text passwords is:

```
[edit system login]
passwords {
  change-type character-sets;
  format md5;
  minimum-changes 1;
  minimum-length 6;
}
```

The default configuration for Junos-FIPS plain-text passwords is:

```
[edit system login]
passwords {
  change-type set-transitions;
  format sha1;
```

```
maximum-length 20;  
minimum-changes 3;  
minimum-length 10;  
}
```

- Related Documentation**
- [Changing the Requirements for Junos OS Plain-Text Passwords](#)
 - [Configuring the Root Password on page 145](#)

Configuring Junos OS User Accounts

User accounts provide one way for users to access the router or switch. For each account, you define the login name for the user and, optionally, information that identifies the user. After you have created an account, the software creates a home directory for the user.

To create user accounts, include the **user** statement at the **[edit system login]** hierarchy level:

```
[edit system login]  
user username {  
  class class-name;  
  class {  
    (encrypted-password "password" | plain-text-password);  
    ssh-rsa "public-key";  
    ssh-dsa "public-key";  
  }  
  full-name complete-name;  
  uid uid-value;  
  class class-name;  
}
```

- Related Documentation**
- [Example: Configuring User Accounts on page 1830](#)
 - [Configuring a Local Administrator Account on page 1853](#)
 - [Junos OS User Accounts Overview on page 1821](#)
 - [Limiting the Number of User Login Attempts for SSH and Telnet Sessions](#)

Configuring a Local Administrator Account

The following example shows how to configure a password-protected local administration account called **admin** with superuser privileges. Superuser privileges give a user permission to use any command on the router and are generally reserved for a select few users such as system administrators. It is important to protect the local administrator account with a password to prevent unauthorized users from gaining access to superuser commands that can be used to alter the system configuration. Even users with RADIUS authentication should configure a local password. If RADIUS fails or becomes unreachable, the login process will revert to password authentication on the local administrator account.

```
[edit]
system {
  login {
    user admin {
      uid 1000;
      class superuser;
      authentication {
        encrypted-password "<PASSWORD>"; # SECRET-DATA
      }
    }
  }
}
```

**Related
Documentation**

- [Junos OS Login Classes Overview on page 1823](#)
- *Configuring Junos OS User Accounts by Using a Configuration Group*

Example: Creating Login Classes with Specific Privileges

Login classes are used to assign certain permissions or restrictions to groups of users, ensuring that sensitive commands are only accessible to the appropriate users. By default, Juniper Networks devices have four types of login classes with preset permissions: operator, read-only, superuser or super-user, and unauthorized.

You can create new custom login classes to make different combinations of permissions that are not found in the default login classes. The following example shows how to create three custom login classes, each with specific privileges and timers to disconnect the class members after a period of inactivity. Inactivity timers help protect network security by disconnecting a user from the network if the user is away from his computer for too long, preventing potential security risks created by leaving an unattended account logged in to a switch or router. The permissions and inactivity timers shown here are only examples and should be customized to your organization.

The first class of users is called **observation** and they can only view statistics and configuration. They are not allowed to modify any configuration. The second class of users is called **operation** and they can view and modify the configuration. The third class of users is called **engineering** and they have unlimited access and control. All three login classes use the same inactivity timer of 5 minutes.

```
[edit]
system {
  login {
    class observation {
      idle-timeout 5;
      permissions [ view ];
    }
    class operation {
      idle-timeout 5;
      permissions [ admin clear configure interface interface-control network
        reset routing routing-control snmp snmp-control trace-control
        firewall-control rollback ];
    }
    class engineering {
```

```

        idle-timeout 5;
        permissions all;
    }
}
}

```

**Related
Documentation**

- [Junos OS Login Classes Overview on page 1823](#)
- [Defining Junos OS Login Classes](#)
- [Configuring a Local Administrator Account on page 1853](#)

Example: Configuring System Authentication for RADIUS, TACACS+, and Password Authentication

The following example shows how to configure system authentication for RADIUS, TACACS+, and password authentication.

In this example, only the user Philip and users authenticated by a remote RADIUS server can log in. If a user logs in and is not authenticated by the RADIUS server, the user is denied access to the router or switch. If the RADIUS server is not available, the user is authenticated using the **password** authentication method and allowed access to the router or switch. For more information about the password authentication method, see [“Using Local Password Authentication” on page 1860](#).

When Philip tries to log in to the system, if the RADIUS server authenticates him, he is given access and privileges for the **super-user** class. Local accounts are not configured for other users. When they log in to the system and the RADIUS server authenticates them, they are given access using the same user ID (UID) 9999 and the privileges associated with the **operator** class.

```

[edit]
system {
  authentication-order radius;
  login {
    user philip {
      full-name "Philip";
      uid 1001;
      class super-user;
    }
    user remote {
      full-name "All remote users";
      uid 9999;
      class operator;
    }
  }
}
}

```



NOTE: For authorization purposes, you can use a template account to create a single account that can be shared by a set of users at the same time. For example, when you create a remote template account, a set of remote users can concurrently share a single UID. For more information about template accounts, see [“Overview of Template Accounts for RADIUS and TACACS+ Authentication”](#) on page 1899.

When a user logs in to a device, the user’s login name is used by the RADIUS or TACACS+ server for authentication. If the user is authenticated successfully by the authentication server and the user is not configured at the `[edit system login user]` hierarchy level, the device uses the default remote template user account for the user, provided a remote template account is configured at the `edit system login user remote` hierarchy level. The remote template account serves as a default template user account for all users that are authenticated by the authentication server but not having a locally configured user account on the device. Such users share the same login class and UID.

To configure an alternate template user, specify the **user-name** parameter returned in the RADIUS authentication response packet. Not all RADIUS servers allow you to change this parameter. The following shows a sample Junos OS configuration:

```
[edit]
system {
  authentication-order radius;
  login {
    user philip {
      full-name "Philip";
      uid 1001;
      class super-user;
    }
    user operator {
      full-name "All operators";
      uid 9990;
      class operator;
    }
    user remote {
      full-name "All remote users";
      uid 9999;
      class read-only;
    }
  }
}
```

Assume your RADIUS server is configured with the following information:

- User Philip with password “olympia”
- User Alexander with password “bucephalus” and username “operator”
- User Darius with password “redhead” and username “operator”
- User Roxane with password “athena”

Philip would be given access as a superuser (**super-user**) because he has his own local user account. Alexander and Darius share UID 9990 and have access as operators. Roxane has no template-user override, so she shares access with all the other remote users, getting read-only access.

Related Documentation

- *Configuring the Junos OS Authentication Order for RADIUS, TACACS+, and Local Password Authentication*

Example: Changing the Requirements for Junos OS Plain-Text Passwords

This example shows how to set various maximum and minimum requirements for plain-text passwords to increase password strength.

- [Requirements on page 1857](#)
- [Overview on page 1857](#)
- [Configuration on page 1857](#)

Requirements

This example requires a device running Junos 12.2 or greater. The **minimum-length** and **maximum-length** password requirements statements are available in earlier releases, however, you must have Junos OS Release 12.2 or greater to configure **minimum-lower-cases**, **minimum-numeric**s, **minimum-punctuations**, or **minimum-upper-cases**.

Overview

You can use a variety of requirements to strengthen plain-text passwords for greater security. Junos OS provides a number of possible configurations at the **[edit system login password]** hierarchy level that allow you to require users to create plain-text passwords that conform to a particular set of requirements that may include such things as length, number of changes, type of characters, numbers, or letter case.

Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set system login password minimum-length 12
set system login password maximum-length 22
set system login password minimum-numeric 1
set system login password minimum-upper-cases 1
set system login password minimum-lower-cases 1
set system login password minimum-punctuations 1
```

Configuring Requirements for Plain-Text Passwords

Step-by-Step Procedure This example configures password requirements that require the user to create a password that has a minimum length of 12 characters, a maximum length of 22 characters, and that includes at least one lower-case letter, at least one upper-case letter, at least one punctuation character, and at least one numeric character.

1. Navigate to configuration mode in the [system login password] hierarchy level.

```
user@host> edit
[edit]
user@host# edit system login password
```
2. Set a minimum length requirement of 12 characters and a maximum length requirement of 22 characters for user passwords.

```
[edit system login password]
user@host# set minimum-length 12
[edit system login password]
user@host# set maximum-length 22
```
3. Require users to set a password that has at least one lower-case letter and at least one upper-case letter.

```
[edit system login password]
user@host# set minimum-lower-cases 1
[edit system login password]
user@host# set minimum-upper-cases 1
```
4. Require users to set a password that has at least one punctuation-class character and at least one number.

```
[edit system login password]
user@host# set minimum-punctuations 1
[edit system login password]
user@host# set minimum-numeric 1
```

Results

From configuration mode, confirm your configuration by entering the show command at the edit system login password hierarchy level. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit system login password]
user@host# show
minimum-length 12;
maximum-length 22;
minimum-numeric 1;
minimum-upper-cases 1;
minimum-lower-cases 1;
```

- Related Documentation**
- [Special Requirements for Junos OS Plain-Text Passwords on page 1824](#)
 - *password (Login)*

CHAPTER 64

Configuring and Managing TACACS+ Authentication

- Junos OS Authentication Order for RADIUS, TACACS+, and Password Authentication on page 1859
- Juniper Networks Vendor-Specific TACACS+ Attributes on page 1863
- Example: Configuring System Authentication for RADIUS, TACACS+, and Password Authentication on page 1865

Junos OS Authentication Order for RADIUS, TACACS+, and Password Authentication

Using the **authentication-order** statement, you can prioritize the order in which the Junos OS tries the different authentication methods when verifying user access to a router or switch.

If the **authentication-order** is remote-server then local, Junos OS will retry the local server if the remote-server is unreachable or has timed out. However; if the remote-server rejects the authentication, Junos OS will not retry the authentication.

If none of the configured authentication methods accept the login credentials and if a reject response is received, the login attempt fails. If no response is received from any configured authentication method, the Junos OS consults local password authentication as a last resort.

Using RADIUS or TACACS+ Authentication

You can configure the Junos OS to be both a RADIUS and TACACS+ authentication client.

If an authentication method included in the **[authentication-order]** statement is not available, or if the authentication is available but returns a reject response, the Junos OS tries the next authentication method included in the **authentication-order** statement.

The RADIUS or TACACS+ server authentication might fail because of the following reasons:

- The authentication method is configured, but the corresponding authentication servers are not configured. For instance, the RADIUS and TACACS+ authentication methods are included in the **authentication-order** statement, but the corresponding RADIUS or

TACACS+ servers are not configured at the respective **[edit system radius-server]** and **[edit system tacplus-server]** hierarchy levels.

- The RADIUS or TACACS+ server does not respond within the timeout period configured at the **[edit system radius-server]** or **[edit system tacplus-server]** hierarchy levels.
- The RADIUS or TACACS+ server is not reachable because of a network problem.

The RADIUS or TACACS+ server authentication might return a reject response because of the following reasons:

- The user profiles of users accessing a router or switch might not be configured on the RADIUS or TACACS+ server.
- The user enters incorrect logon credentials.

Using Local Password Authentication

You can explicitly configure the password authentication method or use this method as a fallback mechanism when remote authentication servers fail. The password authentication method consults the local user profiles configured at the **[edit system login]** hierarchy level. Users can log in to a router or switch using their local username and password in the following scenarios:

- The password authentication method (password) is explicitly configured as one of the authentication methods in the **[authentication-order authentication-methods]** statement. In this case, the password authentication method is tried if no previous authentication accepts the logon credentials. This is true whether the previous authentication method fails to respond or returns a reject response because of an incorrect username or password.
- The password authentication method is not explicitly configured as one of the authentication methods in the **authentication-order authentication-methods** statement. In this case, the password authentication method is tried only if all configured authentication methods fail to respond. It is not consulted if any configured authentication method returns a reject response because of an incorrect username or password.

Order of Authentication Attempts

[Table 150](#) describes how the **authentication-order** statement at the **[edit system]** hierarchy level determines the procedure that the Junos OS uses to authenticate users for access to a router or switch.

Table 150: Order of Authentication Attempts

Syntax	Order of Authentication Attempts
authentication-order radius;	<ol style="list-style-type: none"> 1. Try configured RADIUS authentication servers. 2. If RADIUS server is available and authentication is accepted, grant access. 3. If RADIUS server is available but authentication is rejected, deny access. 4. If RADIUS servers are not available, try password authentication. <p>NOTE: If a RADIUS server is available, password authentication is not attempted, because it is not explicitly configured in the authentication order.</p>
authentication-order [radius password];	<ol style="list-style-type: none"> 1. Try configured RADIUS authentication servers. 2. If RADIUS servers fail to respond or return a reject response, try password authentication, because it is explicitly configured in the authentication order.
authentication-order [radius tacplus];	<ol style="list-style-type: none"> 1. Try configured RADIUS authentication servers. 2. If RADIUS server is available and authentication is accepted, grant access. 3. If RADIUS servers fail to respond or return a reject response, try configured TACACS+ servers. 4. If TACACS+ server is available and authentication is accepted, grant access. 5. If TACACS+ server is available but authentication is rejected, deny access. 6. If both RADIUS and TACACS+ servers are not available, try password authentication. <p>NOTE: If either RADIUS or TACACS+ servers are available, password authentication is not attempted, because it is not explicitly configured in the authentication order.</p>
authentication-order [radius tacplus password];	<ol style="list-style-type: none"> 1. Try configured RADIUS authentication servers. 2. If RADIUS server is available and authentication is accepted, grant access. 3. If RADIUS servers fail to respond or return a reject response, try configured TACACS+ servers. 4. If TACACS+ server is available and authentication is accepted, grant access. 5. If TACACS+ servers fail to respond or return a reject response, try password authentication, because it is explicitly configured in the authentication order.

Table 150: Order of Authentication Attempts (*continued*)

Syntax	Order of Authentication Attempts
authentication-order tacplus;	<ol style="list-style-type: none"> 1. Try configured TACACS+ authentication servers. 2. If TACACS+ server is available and authentication is accepted, grant access. 3. If TACACS+ server is available but authentication is rejected, deny access. 4. If TACACS+ servers are not available, try password authentication. <p>NOTE: If a TACACS+ server is available, password authentication is not attempted, because it is not explicitly configured in the authentication order.</p>
authentication-order [tacplus password];	<ol style="list-style-type: none"> 1. Try configured TACACS+ authentication servers. 2. If TACACS+ servers fail to respond or return a reject response, try password authentication, because it is explicitly configured in the authentication order.
authentication-order [tacplus radius];	<ol style="list-style-type: none"> 1. Try configured TACACS+ authentication servers. 2. If TACACS+ server is available and authentication is accepted, grant access. 3. If TACACS+ servers fail to respond or return a reject response, try configured RADIUS servers. 4. If RADIUS server is available and authentication is accepted, grant access. 5. If RADIUS server is available but authentication is rejected, deny access. 6. If both TACACS+ and RADIUS servers are not available, try password authentication. <p>NOTE: If either TACACS+ or RADIUS servers are available, password authentication is not attempted, because it is not explicitly configured in the authentication order.</p>
authentication-order [tacplus radius password];	<ol style="list-style-type: none"> 1. Try configured TACACS+ authentication servers. 2. If TACACS+ server is available and authentication is accepted, grant access. 3. If TACACS+ servers fail to respond or return a reject response, try configured RADIUS servers. 4. If RADIUS server is available and authentication is accepted, grant access. 5. If RADIUS servers fail to respond or return a reject response try password authentication, because it is explicitly configured in the authentication order.

Table 150: Order of Authentication Attempts (*continued*)

Syntax	Order of Authentication Attempts
<code>authentication-order password;</code>	<ol style="list-style-type: none"> 1. Try to authenticate the user, using the password configured at the <code>[edit system login]</code> hierarchy level. 2. If the authentication is accepted, grant access. 3. If the authentication is rejected, deny access.



NOTE: If SSH public keys are configured, SSH user authentication first tries to perform public key authentication before using the authentication methods configured in the `authentication-order` statement. If you want SSH logins to use the authentication methods configured in the `authentication-order` statement without first trying to perform public key authentication, do not configure SSH public keys.

In a routing matrix based on a TX Matrix router, the authentication order must be configured only at the configuration groups `re0` and `re1`. The authentication order must not be configured at the `[edit system]` hierarchy. This is because the authentication order for the routing matrix is controlled on the switch-card chassis (or TX Matrix router) or switch-fabric chassis (for TX Matrix Plus router) only.

In Junos OS Release 10.0 and later, the superuser (belonging to the super-user login class) is also authenticated based on the authentication order that is configured for TACACS+, RADIUS, or password authentication using the `authentication-order` statement. For example, if the only configured authentication order is TACACS+, the superuser can only be authenticated by the TACACS+ server and password authentication cannot be used as an alternative. However, in Junos OS Release 9.6 and earlier, the superuser can use password authentication to login, even if password authentication is not configured explicitly using the `authentication-order` statement.

Related Documentation

- [Overview of Template Accounts for RADIUS and TACACS+ Authentication on page 1899](#)
- [Configuring the Junos OS Authentication Order for RADIUS, TACACS+, and Local Password Authentication](#)
- [Limiting the Number of User Login Attempts for SSH and Telnet Sessions](#)
- [Example: Configuring System Authentication for RADIUS, TACACS+, and Password Authentication on page 1855](#)

Juniper Networks Vendor-Specific TACACS+ Attributes

Junos OS supports the configuration of Juniper Networks TACACS+ vendor-specific attributes (VSAs). These VSAs are encapsulated in a TACACS+ vendor-specific attribute

with the vendor ID set to the Juniper Networks ID number, 2636. [Table 151](#) lists the Juniper Networks VSAs you can configure.

Table 151: Juniper Networks Vendor-Specific TACACS+ Attributes

Name	Description	Length	String
local-user-name	Indicates the name of the user template used by this user when logging in to a device.	≥3	One or more octets containing printable ASCII characters.
allow-commands	Contains an extended regular expression that enables the user to run operational mode commands in addition to those commands authorized by the user's login class permission bits.	≥3	One or more octets containing printable ASCII characters, in the form of an extended regular expression. See Table 143 .
allow-configuration	Contains an extended regular expression that enables the user to run configuration mode commands in addition to those commands authorized by the user's login class permission bits.	≥3	One or more octets containing printable ASCII characters, in the form of an extended regular expression. See “Regular Expressions for Allowing and Denying Junos OS Configuration Mode Hierarchies” on page 1826.
deny-commands	Contains an extended regular expression that denies the user permission to run operational mode commands authorized by the user's login class permission bits.	≥3	One or more octets containing printable ASCII characters, in the form of an extended regular expression. See Table 143 .
deny-configuration	Contains an extended regular expression that denies the user permission to run configuration mode commands authorized by the user's login class permission bits.	≥3	One or more octets containing printable ASCII characters, in the form of an extended regular expression. See Table 142 .
user-permissions	<p>Contains information the server uses to specify user permissions.</p> <p>NOTE: When the user-permissions attribute is configured to grant the Junos OS maintenance or all permissions on an IPv4 or IPv6 TACACS+ server, the UNIX wheel group membership is not automatically added to a user's list of group memberships. Some operations such as running the su root command from a local shell require wheel group membership permissions. However, when a user is configured locally with the permissions maintenance or all, the user is automatically granted membership to the UNIX wheel group. Therefore, we recommend that you create a template user account with the required permissions and associate individual user accounts with the template user account.</p>	≥3	One or more octets containing printable ASCII characters. See Table 137 .

Table 151: Juniper Networks Vendor-Specific TACACS+ Attributes (*continued*)

Name	Description	Length	String
authentication-type	Indicates the authentication method (local database, or TACACS+ server) used to authenticate a user. If the user is authenticated using a local database, the attribute value shows 'local'. If the user is authenticated using TACACS+ server, the attribute value shows 'remote'.	≥5	One or more octets containing printable ASCII characters.
session-port	Indicates the source port number of the established session.	size of integer	Integer

Related Documentation

- *Configuring TACACS+ Authentication*

Example: Configuring System Authentication for RADIUS, TACACS+, and Password Authentication

The following example shows how to configure system authentication for RADIUS, TACACS+, and password authentication.

In this example, only the user Philip and users authenticated by a remote RADIUS server can log in. If a user logs in and is not authenticated by the RADIUS server, the user is denied access to the router or switch. If the RADIUS server is not available, the user is authenticated using the **password** authentication method and allowed access to the router or switch. For more information about the password authentication method, see [“Using Local Password Authentication” on page 1860](#).

When Philip tries to log in to the system, if the RADIUS server authenticates him, he is given access and privileges for the **super-user** class. Local accounts are not configured for other users. When they log in to the system and the RADIUS server authenticates them, they are given access using the same user ID (UID) 9999 and the privileges associated with the **operator** class.

```
[edit]
system {
  authentication-order radius;
  login {
    user philip {
      full-name "Philip";
      uid 1001;
      class super-user;
    }
    user remote {
      full-name "All remote users";
      uid 9999;
      class operator;
    }
  }
}
```



NOTE: For authorization purposes, you can use a template account to create a single account that can be shared by a set of users at the same time. For example, when you create a remote template account, a set of remote users can concurrently share a single UID. For more information about template accounts, see [“Overview of Template Accounts for RADIUS and TACACS+ Authentication”](#) on page 1899.

When a user logs in to a device, the user's login name is used by the RADIUS or TACACS+ server for authentication. If the user is authenticated successfully by the authentication server and the user is not configured at the `[edit system login user]` hierarchy level, the device uses the default remote template user account for the user, provided a remote template account is configured at the `edit system login user remote` hierarchy level. The remote template account serves as a default template user account for all users that are authenticated by the authentication server but not having a locally configured user account on the device. Such users share the same login class and UID.

To configure an alternate template user, specify the **user-name** parameter returned in the RADIUS authentication response packet. Not all RADIUS servers allow you to change this parameter. The following shows a sample Junos OS configuration:

```
[edit]
system {
  authentication-order radius;
  login {
    user philip {
      full-name "Philip";
      uid 1001;
      class super-user;
    }
    user operator {
      full-name "All operators";
      uid 9990;
      class operator;
    }
    user remote {
      full-name "All remote users";
      uid 9999;
      class read-only;
    }
  }
}
```

Assume your RADIUS server is configured with the following information:

- User Philip with password “olympia”
- User Alexander with password “bucephalus” and username “operator”
- User Darius with password “redhead” and username “operator”
- User Roxane with password “athena”

Philip would be given access as a superuser (**super-user**) because he has his own local user account. Alexander and Darius share UID 9990 and have access as operators. Roxane has no template-user override, so she shares access with all the other remote users, getting read-only access.

- Related Documentation**
- *Configuring the Junos OS Authentication Order for RADIUS, TACACS+, and Local Password Authentication*

Configuring and Managing RADIUS Authentication

- Junos OS Authentication Order for RADIUS, TACACS+, and Password Authentication on page 1869
- Configuring RADIUS Authentication (QFX Series or OCX Series) on page 1873
- Using Regular Expressions on a RADIUS or TACACS+ Server to Allow or Deny Access to Commands on page 1876
- Example: Configuring RADIUS Authentication on page 1878
- Example: Configuring RADIUS Template Accounts on page 1879
- Configuring a Local Administrator Account on page 1879
- Example: Configuring System Authentication for RADIUS, TACACS+, and Password Authentication on page 1880

Junos OS Authentication Order for RADIUS, TACACS+, and Password Authentication

Using the **authentication-order** statement, you can prioritize the order in which the Junos OS tries the different authentication methods when verifying user access to a router or switch.

If the **authentication-order** is remote-server then local, Junos OS will retry the local server if the remote-server is unreachable or has timed out. However, if the remote-server rejects the authentication, Junos OS will not retry the authentication.

If none of the configured authentication methods accept the login credentials and if a reject response is received, the login attempt fails. If no response is received from any configured authentication method, the Junos OS consults local password authentication as a last resort.

Using RADIUS or TACACS+ Authentication

You can configure the Junos OS to be both a RADIUS and TACACS+ authentication client.

If an authentication method included in the **[authentication-order]** statement is not available, or if the authentication is available but returns a reject response, the Junos OS tries the next authentication method included in the **authentication-order** statement.

The RADIUS or TACACS+ server authentication might fail because of the following reasons:

- The authentication method is configured, but the corresponding authentication servers are not configured. For instance, the RADIUS and TACACS+ authentication methods are included in the **authentication-order** statement, but the corresponding RADIUS or TACACS+ servers are not configured at the respective **[edit system radius-server]** and **[edit system tacplus-server]** hierarchy levels.
- The RADIUS or TACACS+ server does not respond within the timeout period configured at the **[edit system radius-server]** or **[edit system tacplus-server]** hierarchy levels.
- The RADIUS or TACACS+ server is not reachable because of a network problem.

The RADIUS or TACACS+ server authentication might return a reject response because of the following reasons:

- The user profiles of users accessing a router or switch might not be configured on the RADIUS or TACACS+ server.
- The user enters incorrect logon credentials.

Using Local Password Authentication

You can explicitly configure the password authentication method or use this method as a fallback mechanism when remote authentication servers fail. The password authentication method consults the local user profiles configured at the **[edit system login]** hierarchy level. Users can log in to a router or switch using their local username and password in the following scenarios:

- The password authentication method (password) is explicitly configured as one of the authentication methods in the **[authentication-order authentication-methods]** statement. In this case, the password authentication method is tried if no previous authentication accepts the logon credentials. This is true whether the previous authentication method fails to respond or returns a reject response because of an incorrect username or password.
- The password authentication method is not explicitly configured as one of the authentication methods in the **authentication-order authentication-methods** statement. In this case, the password authentication method is tried only if all configured authentication methods fail to respond. It is not consulted if any configured authentication method returns a reject response because of an incorrect username or password.

Order of Authentication Attempts

[Table 150](#) describes how the **authentication-order** statement at the **[edit system]** hierarchy level determines the procedure that the Junos OS uses to authenticate users for access to a router or switch.

Table 152: Order of Authentication Attempts

Syntax	Order of Authentication Attempts
authentication-order radius;	<ol style="list-style-type: none"> 1. Try configured RADIUS authentication servers. 2. If RADIUS server is available and authentication is accepted, grant access. 3. If RADIUS server is available but authentication is rejected, deny access. 4. If RADIUS servers are not available, try password authentication. <p>NOTE: If a RADIUS server is available, password authentication is not attempted, because it is not explicitly configured in the authentication order.</p>
authentication-order [radius password];	<ol style="list-style-type: none"> 1. Try configured RADIUS authentication servers. 2. If RADIUS servers fail to respond or return a reject response, try password authentication, because it is explicitly configured in the authentication order.
authentication-order [radius tacplus];	<ol style="list-style-type: none"> 1. Try configured RADIUS authentication servers. 2. If RADIUS server is available and authentication is accepted, grant access. 3. If RADIUS servers fail to respond or return a reject response, try configured TACACS+ servers. 4. If TACACS+ server is available and authentication is accepted, grant access. 5. If TACACS+ server is available but authentication is rejected, deny access. 6. If both RADIUS and TACACS+ servers are not available, try password authentication. <p>NOTE: If either RADIUS or TACACS+ servers are available, password authentication is not attempted, because it is not explicitly configured in the authentication order.</p>
authentication-order [radius tacplus password];	<ol style="list-style-type: none"> 1. Try configured RADIUS authentication servers. 2. If RADIUS server is available and authentication is accepted, grant access. 3. If RADIUS servers fail to respond or return a reject response, try configured TACACS+ servers. 4. If TACACS+ server is available and authentication is accepted, grant access. 5. If TACACS+ servers fail to respond or return a reject response, try password authentication, because it is explicitly configured in the authentication order.

Table 152: Order of Authentication Attempts (*continued*)

Syntax	Order of Authentication Attempts
authentication-order tacplus;	<ol style="list-style-type: none"> 1. Try configured TACACS+ authentication servers. 2. If TACACS+ server is available and authentication is accepted, grant access. 3. If TACACS+ server is available but authentication is rejected, deny access. 4. If TACACS+ servers are not available, try password authentication. <p>NOTE: If a TACACS+ server is available, password authentication is not attempted, because it is not explicitly configured in the authentication order.</p>
authentication-order [tacplus password];	<ol style="list-style-type: none"> 1. Try configured TACACS+ authentication servers. 2. If TACACS+ servers fail to respond or return a reject response, try password authentication, because it is explicitly configured in the authentication order.
authentication-order [tacplus radius];	<ol style="list-style-type: none"> 1. Try configured TACACS+ authentication servers. 2. If TACACS+ server is available and authentication is accepted, grant access. 3. If TACACS+ servers fail to respond or return a reject response, try configured RADIUS servers. 4. If RADIUS server is available and authentication is accepted, grant access. 5. If RADIUS server is available but authentication is rejected, deny access. 6. If both TACACS+ and RADIUS servers are not available, try password authentication. <p>NOTE: If either TACACS+ or RADIUS servers are available, password authentication is not attempted, because it is not explicitly configured in the authentication order.</p>
authentication-order [tacplus radius password];	<ol style="list-style-type: none"> 1. Try configured TACACS+ authentication servers. 2. If TACACS+ server is available and authentication is accepted, grant access. 3. If TACACS+ servers fail to respond or return a reject response, try configured RADIUS servers. 4. If RADIUS server is available and authentication is accepted, grant access. 5. If RADIUS servers fail to respond or return a reject response try password authentication, because it is explicitly configured in the authentication order.

Table 152: Order of Authentication Attempts (*continued*)

Syntax	Order of Authentication Attempts
<code>authentication-order password;</code>	<ol style="list-style-type: none"> 1. Try to authenticate the user, using the password configured at the <code>[edit system login]</code> hierarchy level. 2. If the authentication is accepted, grant access. 3. If the authentication is rejected, deny access.



NOTE: If SSH public keys are configured, SSH user authentication first tries to perform public key authentication before using the authentication methods configured in the `authentication-order` statement. If you want SSH logins to use the authentication methods configured in the `authentication-order` statement without first trying to perform public key authentication, do not configure SSH public keys.

In a routing matrix based on a TX Matrix router, the authentication order must be configured only at the configuration groups `re0` and `re1`. The authentication order must not be configured at the `[edit system]` hierarchy. This is because the authentication order for the routing matrix is controlled on the switch-card chassis (or TX Matrix router) or switch-fabric chassis (for TX Matrix Plus router) only.

In Junos OS Release 10.0 and later, the superuser (belonging to the super-user login class) is also authenticated based on the authentication order that is configured for TACACS+, RADIUS, or password authentication using the `authentication-order` statement. For example, if the only configured authentication order is TACACS+, the superuser can only be authenticated by the TACACS+ server and password authentication cannot be used as an alternative. However, in Junos OS Release 9.6 and earlier, the superuser can use password authentication to login, even if password authentication is not configured explicitly using the `authentication-order` statement.

Related Documentation

- [Overview of Template Accounts for RADIUS and TACACS+ Authentication on page 1899](#)
- [Configuring the Junos OS Authentication Order for RADIUS, TACACS+, and Local Password Authentication](#)
- [Limiting the Number of User Login Attempts for SSH and Telnet Sessions](#)
- [Example: Configuring System Authentication for RADIUS, TACACS+, and Password Authentication on page 1855](#)

Configuring RADIUS Authentication (QFX Series or OCX Series)

RADIUS authentication is a method of authenticating users who attempt to access the router or switch. Tasks to configure RADIUS authentication are:



NOTE: The `source-address` statement is not supported at the `[edit system radius-options]` or `[edit system-radius-server name]` hierarchies on the QFabric system.

- [Configuring RADIUS Server Details on page 1874](#)
- [Configuring MS-CHAPv2 for Password-Change Support on page 1875](#)
- [Specifying a Source Address for the Junos OS to Access External RADIUS Servers on page 1876](#)

Configuring RADIUS Server Details

To use RADIUS authentication on the router or switch, configure information about one or more RADIUS servers on the network by including one `radius-server` statement at the `[edit system]` hierarchy level for each RADIUS server:

```
[edit system]
radius-server server-address {
  accounting-port port-number;
  port number;
  retry number;
  secret password;
  source-address source-address;
  timeout seconds;
}
```

`server-address` is the address of the RADIUS server.

You can specify a port on which to contact the RADIUS server. By default, port number **1812** is used (as specified in RFC 2865). You can also specify an accounting port to send accounting packets. The default is **1813** (as specified in RFC 2866).

You must specify a password in the `secret password` statement. If the password contains spaces, enclose it in quotation marks. The secret used by the local router or switch must match that used by the server.

Optionally, you can specify the amount of time that the local router or switch waits to receive a response from a RADIUS server (in the `timeout` statement) and the number of times that the router or switch attempts to contact a RADIUS authentication server (in the `retry` statement). By default, the router or switch waits 3 seconds. You can configure this to be a value from 1 through 90 seconds. By default, the router or switch retries connecting to the server three times. You can configure this to be a value from 1 through 10 times.

You can use the `source-address` statement to specify a logical address for individual or multiple RADIUS servers.

To configure multiple RADIUS servers, include multiple `radius-server` statements.

To configure a set of users that share a single account for authorization purposes, you create a template user. To do this, include the `user` statement at the `[edit system login]`

hierarchy level, as described in [“Overview of Template Accounts for RADIUS and TACACS+ Authentication” on page 1899](#).

You can also configure RADIUS authentication at the **[edit access]** and **[edit access profile]** hierarchy level. Junos OS uses the following search order to determine which set of servers are used for authentication:

1. **[edit access profile *profile-name* radius-server *server-address*]**
2. **[edit access radius-server *server-address*]**
3. **[edit system radius-server *server-address*]**

Configuring MS-CHAPv2 for Password-Change Support

You can configure the Microsoft implementation of the Challenge Handshake Authentication Protocol version 2 (MS-CHAPv2) on the router or switch to support changing of passwords. This feature provides users accessing a router or switch the option of changing the password when the password expires, is reset, or is configured to be changed at the next login.

Before you configure MS-CHAPv2 for password-change support, ensure that you:

- Configure the RADIUS server authentication parameters
- Set the **authentication-order** to use the RADIUS server for the initial password attempt

To configure MS-CHAP-v2, include the following statements at the **[edit system radius-options]** hierarchy level:

```
[edit system radius-options]
password-protocol mschap-v2;
```

The following example shows statements for configuring the MS-CHAPv2 password protocol, password authentication order, and user accounts:

```
[edit]
system {
  authentication-order [ radius password ];
  radius-server {
    192.168.69.149 secret "$9$G-j.5Qz6tpBk.1hrlXxUjiq5Qn/C"; ## SECRET-DATA
  }
  radius-options {
    password-protocol mschap-v2;
  }
  login {
    user bob {
      class operator;
    }
  }
}
```

Specifying a Source Address for the Junos OS to Access External RADIUS Servers

You can specify which source address Junos OS uses when accessing your network to contact an external RADIUS server for authentication. You can also specify which source address Junos OS uses when contacting a RADIUS server for sending accounting information.

To specify a source address for a RADIUS server, include the **source-address** statement at the **[edit system radius-server server-address]** hierarchy level:

```
[edit system radius-server server-address]
source-address source-address;
```

source-address is a valid IP address configured on one of the router or switch interfaces.

Related Documentation

- [Example: Configuring RADIUS Authentication on page 1878](#)
- [Example: Configuring System Authentication for RADIUS, TACACS+, and Password Authentication on page 1855](#)
- [Juniper Networks Vendor-Specific RADIUS Attributes on page 1888](#)
- [Overview of Template Accounts for RADIUS and TACACS+ Authentication on page 1899](#)
- [Example: Configuring RADIUS Template Accounts on page 1879](#)
- [Using Regular Expressions on a RADIUS or TACACS+ Server to Allow or Deny Access to Commands on page 1876](#)
- [Junos OS User Authentication Methods on page 1804](#)

Using Regular Expressions on a RADIUS or TACACS+ Server to Allow or Deny Access to Commands

Use regular expressions to specify which operational or configuration mode commands are allowed or denied when you use a RADIUS or TACACS+ server for user authentication. You can specify the regular expressions using the appropriate Juniper Networks vendor-specific RADIUS or TACACS+ attributes in your authentication server configuration.

You can specify **allow-configuration**, **deny-configuration**, **allow-commands**, or **deny-commands** in a single extended regular expression, enclosing multiple commands in parentheses and separating them using the pipe symbol. For example, you can specify multiple **allow-commands** parameters using: **allow-commands= (cmd1 | cmd2 | cmdn)**. You can specify **user-permissions** as a list of comma-separated values, and not as a regular expression.

On a RADIUS or TACACS+ server, you can also use a simplified version for regular expressions where you specify each individual expression on a separate line. The simplified version is valid for **allow-commands**, **deny-commands**, **allow-configuration**, **deny-configuration**, and **permissions** vendor-specific attributes.

For a RADIUS server, specify the individual regular expressions using the following syntax:

```

Juniper-Allow-Commands+="cmd1"
Juniper-Allow-Commands+="cmd2"
Juniper-Allow-Commands+="cmdn"
Juniper-Deny-Commands+="cmd1"
Juniper-Deny-Commands+="cmd2"
Juniper-Deny-Commands+="cmdn"
Juniper-Allow-Configuration+="regex1"
Juniper-Allow-Configuration+="regex2"
Juniper-Allow-Configuration+="regexn"
Juniper-Deny-Configuration+="regex1"
Juniper-Deny-Configuration+="regex2"
Juniper-Deny-Configuration+="regexn"
Juniper-User-Permissions+="permission-flag1"
Juniper-User-Permissions+="permission-flag2"
Juniper-User-Permissions+="permission-flagn"

```

For TACACS+ server, specify the individual regular expressions using the following syntax:

```

allow-commands1="cmd1"
allow-commands2="cmd2"
allow-commandsn="cmdn"
deny-commands1="cmd1"
deny-commands2="cmd2"
deny-commandsn="cmdn"
allow-configuration1="regex1"
allow-configuration2="regex2"
allow-configurationn="regexn"
deny-configuration1="regex1"
deny-configuration2="regex2"
deny-configurationn="regexn"
user-permissions1="permission-flag1"
user-permissions2="permission-flag2"
user-permissionsn="permission-flagn "

```



NOTE:

- Numeric values 1 to *n* in the syntax (for TACACS+ server) must be unique but need not be sequential. For example, the following syntax is valid:

```

allow-commands1="cmd1"
allow-commands3="cmd3"
allow-commands2="cmd2"
deny-commands3="cmd3"
deny-commands2="cmd2"
deny-commands1="cmd1"

```

- The limit on the number of lines of individual regular expressions is imposed by the TACACS+ or RADIUS server.
- When you issue the `show cli authorization` command, the command output displays the regular expression in a single line, even if you specify each individual expression on a separate line.

For more information about Juniper Networks vendor-specific RADIUS and TACACS+ attributes, see [“Juniper Networks Vendor-Specific RADIUS Attributes” on page 1888](#) and [“Juniper Networks Vendor-Specific TACACS+ Attributes” on page 1863](#).



NOTE: When RADIUS or TACACS+ authentication is configured for a router, regular expressions configured on the RADIUS or TACACS+ server merge with any regular expressions configured on the local router at the [edit system login class] hierarchy level using the allow-commands, deny-commands, allow-configuration, deny-configuration, or permissions statements. If the final expression has a syntax error, the overall result is an invalid regular expression.

**Related
Documentation**

- [Junos OS Authentication Order for RADIUS, TACACS+, and Password Authentication on page 1859](#)

Example: Configuring RADIUS Authentication

The Junos OS supports two protocols for central authentication of users on multiple routers: RADIUS and TACACS+. We recommend RADIUS because it is a multivendor IETF standard, and its features are more widely accepted than those of TACACS+ or other proprietary systems. In addition, we recommend using a one-time-password system for increased security, and all vendors of these systems support RADIUS.

The Junos OS uses one or more template accounts to perform user authentication. You create the template account or accounts, and then configure the user access to use that account. If the RADIUS server is unavailable, the fallback is for the login process to use the local account that set up on the router or switch.

The following example shows how to configure RADIUS authentication:

```
[edit]
system {
  authentication-order [ radius password ];
  root-authentication {
    encrypted-password "$9$aHlj8gqQ1gjyjjhgjgiiii"; # SECRET-DATA
  }
  name-server {
    10.1.1.1;
    10.1.1.2;
  }
}
```

The following example shows how to enable RADIUS authentication and define the shared secret between the client and the server. The secret enables the client and server to determine that they are talking to the trusted peer.

Define a timeout value for each server, so that if there is no response within the specified number of seconds, the router can try either the next server or the next authentication mechanism.

```
[edit]
```

```

system {
  radius-server {
    10.1.2.1 {
      secret "$9$aH1j8gqQ1sdjerrrhser"; # SECRET-DATA
      timeout 5;
    }
    10.1.2.2 {
      secret "$9$aH1j8gqQ1csdoiuardwefoiud"; # SECRET-DATA
      timeout 5;
    }
  }
}

```

Related Documentation

- [Configuring RADIUS Server Authentication](#)

Example: Configuring RADIUS Template Accounts

The following example shows how to configure RADIUS template accounts for different users or groups of users:

```

[edit]
system {
  login {
    user observation {
      uid 1001;
      class observation;
    }
    user operation {
      uid 1002;
      class operation;
    }
    user engineering {
      uid 1003;
      class engineering;
    }
  }
}

```

Related Documentation

- [Overview of Template Accounts for RADIUS and TACACS+ Authentication on page 1899](#)

Configuring a Local Administrator Account

The following example shows how to configure a password-protected local administration account called **admin** with superuser privileges. Superuser privileges give a user permission to use any command on the router and are generally reserved for a select few users such as system administrators. It is important to protect the local administrator account with a password to prevent unauthorized users from gaining access to superuser commands that can be used to alter the system configuration. Even users with RADIUS authentication should configure a local password. If RADIUS fails or becomes unreachable, the login process will revert to password authentication on the local administrator account.

```
[edit]
system {
  login {
    user admin {
      uid 1000;
      class superuser;
      authentication {
        encrypted-password "<PASSWORD>"; # SECRET-DATA
      }
    }
  }
}
```

**Related
Documentation**

- [Junos OS Login Classes Overview on page 1823](#)
- [Configuring Junos OS User Accounts by Using a Configuration Group](#)

Example: Configuring System Authentication for RADIUS, TACACS+, and Password Authentication

The following example shows how to configure system authentication for RADIUS, TACACS+, and password authentication.

In this example, only the user Philip and users authenticated by a remote RADIUS server can log in. If a user logs in and is not authenticated by the RADIUS server, the user is denied access to the router or switch. If the RADIUS server is not available, the user is authenticated using the **password** authentication method and allowed access to the router or switch. For more information about the password authentication method, see [“Using Local Password Authentication” on page 1860](#).

When Philip tries to log in to the system, if the RADIUS server authenticates him, he is given access and privileges for the **super-user** class. Local accounts are not configured for other users. When they log in to the system and the RADIUS server authenticates them, they are given access using the same user ID (UID) 9999 and the privileges associated with the **operator** class.

```
[edit]
system {
  authentication-order radius;
  login {
    user philip {
      full-name "Philip";
      uid 1001;
      class super-user;
    }
    user remote {
      full-name "All remote users";
      uid 9999;
      class operator;
    }
  }
}
```




NOTE: For authorization purposes, you can use a template account to create a single account that can be shared by a set of users at the same time. For example, when you create a remote template account, a set of remote users can concurrently share a single UID. For more information about template accounts, see [“Overview of Template Accounts for RADIUS and TACACS+ Authentication”](#) on page 1899.

When a user logs in to a device, the user’s login name is used by the RADIUS or TACACS+ server for authentication. If the user is authenticated successfully by the authentication server and the user is not configured at the `[edit system login user]` hierarchy level, the device uses the default remote template user account for the user, provided a remote template account is configured at the `edit system login user remote` hierarchy level. The remote template account serves as a default template user account for all users that are authenticated by the authentication server but not having a locally configured user account on the device. Such users share the same login class and UID.

To configure an alternate template user, specify the `user-name` parameter returned in the RADIUS authentication response packet. Not all RADIUS servers allow you to change this parameter. The following shows a sample Junos OS configuration:

```
[edit]
system {
  authentication-order radius;
  login {
    user philip {
      full-name "Philip";
      uid 1001;
      class super-user;
    }
    user operator {
      full-name "All operators";
      uid 9990;
      class operator;
    }
    user remote {
      full-name "All remote users";
      uid 9999;
      class read-only;
    }
  }
}
```

Assume your RADIUS server is configured with the following information:

- User Philip with password “olympia”
- User Alexander with password “bucephalus” and username “operator”
- User Darius with password “redhead” and username “operator”
- User Roxane with password “athena”

Philip would be given access as a superuser (**super-user**) because he has his own local user account. Alexander and Darius share UID 9990 and have access as operators. Roxane has no template-user override, so she shares access with all the other remote users, getting read-only access.

- Related Documentation**
- *Configuring the Junos OS Authentication Order for RADIUS, TACACS+, and Local Password Authentication*

CHAPTER 66

Configuring and Managing RADIUS Accounting

- [Understanding RADIUS Accounting on page 1883](#)
- [Junos OS Authentication Order for RADIUS, TACACS+, and Password Authentication on page 1884](#)
- [Juniper Networks Vendor-Specific RADIUS Attributes on page 1888](#)
- [Configuring RADIUS System Accounting on page 1891](#)
- [Configuring RADIUS Authentication \(QFX Series or OCX Series\) on page 1893](#)
- [Using Regular Expressions on a RADIUS or TACACS+ Server to Allow or Deny Access to Commands on page 1896](#)
- [Example: Configuring RADIUS System Accounting on page 1898](#)

Understanding RADIUS Accounting

Devices support IETF RFC 2866, *RADIUS Accounting*. Configuring RADIUS accounting on the device supports collecting statistical data about users logging in to or out from a LAN and sending the data to a RADIUS accounting server. The statistical data gathered can be used for general network monitoring, analyzing and tracking usage patterns, or billing a user based upon the amount of time or type of services accessed.

To configure RADIUS accounting, specify one or more RADIUS accounting servers to receive the statistical data from the device, and select the type of accounting data to be collected.

The RADIUS accounting server you specify can be the same server used for RADIUS authentication, or it can be a separate RADIUS server. You can specify a list of RADIUS accounting servers. If the primary server (the first one configured) is unavailable, each RADIUS server in the list is tried in the order in which they are configured in the Junos OS.

The RADIUS accounting process between the device and a RADIUS server works like this:

1. A RADIUS accounting server listens for User Datagram Protocol (UDP) packets on a specific port. For example, on FreeRADIUS, the default port is 1813.
2. The device forwards an *accounting-request* packet containing an event record to the accounting server. The event record associated with this supplicant contains an *Acct-Status-Type* attribute whose value indicates the beginning of user service for this supplicant. When the supplicant's session ends, the accounting request contains an *Acct-Status-Type* attribute value indicating the end of user service. The RADIUS accounting server records this as a stop-accounting record containing session information and the length of the session.
3. The RADIUS accounting server logs these events in a file as start-accounting or stop-accounting records. On FreeRADIUS, the filename is the server's address; for example, 122.69.1.250.
4. The accounting server sends an *accounting-response* packet back to the device confirming it has received the accounting request.
5. If the device does not receive a response from the server, it continues to send accounting requests until an accounting response is returned from the accounting server.

The statistics collected through this process can be displayed from the RADIUS server; to see those statistics, the user accesses the log file configured to receive them.

Related Documentation • [Configuring RADIUS System Accounting on page 1891](#)

Junos OS Authentication Order for RADIUS, TACACS+, and Password Authentication

Using the **authentication-order** statement, you can prioritize the order in which the Junos OS tries the different authentication methods when verifying user access to a router or switch.

If the **authentication-order** is remote-server then local, Junos OS will retry the local server if the remote-server is unreachable or has timed out. However, if the remote-server rejects the authentication, Junos OS will not retry the authentication.

If none of the configured authentication methods accept the login credentials and if a reject response is received, the login attempt fails. If no response is received from any configured authentication method, the Junos OS consults local password authentication as a last resort.

Using RADIUS or TACACS+ Authentication

You can configure the Junos OS to be both a RADIUS and TACACS+ authentication client.

If an authentication method included in the **[authentication-order]** statement is not available, or if the authentication is available but returns a reject response, the Junos OS tries the next authentication method included in the **authentication-order** statement.

The RADIUS or TACACS+ server authentication might fail because of the following reasons:

- The authentication method is configured, but the corresponding authentication servers are not configured. For instance, the RADIUS and TACACS+ authentication methods are included in the **authentication-order** statement, but the corresponding RADIUS or TACACS+ servers are not configured at the respective **[edit system radius-server]** and **[edit system tacplus-server]** hierarchy levels.
- The RADIUS or TACACS+ server does not respond within the timeout period configured at the **[edit system radius-server]** or **[edit system tacplus-server]** hierarchy levels.
- The RADIUS or TACACS+ server is not reachable because of a network problem.

The RADIUS or TACACS+ server authentication might return a reject response because of the following reasons:

- The user profiles of users accessing a router or switch might not be configured on the RADIUS or TACACS+ server.
- The user enters incorrect logon credentials.

Using Local Password Authentication

You can explicitly configure the password authentication method or use this method as a fallback mechanism when remote authentication servers fail. The password authentication method consults the local user profiles configured at the **[edit system login]** hierarchy level. Users can log in to a router or switch using their local username and password in the following scenarios:

- The password authentication method (password) is explicitly configured as one of the authentication methods in the **[authentication-order authentication-methods]** statement. In this case, the password authentication method is tried if no previous authentication accepts the logon credentials. This is true whether the previous authentication method fails to respond or returns a reject response because of an incorrect username or password.
- The password authentication method is not explicitly configured as one of the authentication methods in the **authentication-order authentication-methods** statement. In this case, the password authentication method is tried only if all configured authentication methods fail to respond. It is not consulted if any configured authentication method returns a reject response because of an incorrect username or password.

Order of Authentication Attempts

[Table 150](#) describes how the **authentication-order** statement at the **[edit system]** hierarchy level determines the procedure that the Junos OS uses to authenticate users for access to a router or switch.

Table 153: Order of Authentication Attempts

Syntax	Order of Authentication Attempts
authentication-order radius;	<ol style="list-style-type: none"> 1. Try configured RADIUS authentication servers. 2. If RADIUS server is available and authentication is accepted, grant access. 3. If RADIUS server is available but authentication is rejected, deny access. 4. If RADIUS servers are not available, try password authentication. <p>NOTE: If a RADIUS server is available, password authentication is not attempted, because it is not explicitly configured in the authentication order.</p>
authentication-order [radius password];	<ol style="list-style-type: none"> 1. Try configured RADIUS authentication servers. 2. If RADIUS servers fail to respond or return a reject response, try password authentication, because it is explicitly configured in the authentication order.
authentication-order [radius tacplus];	<ol style="list-style-type: none"> 1. Try configured RADIUS authentication servers. 2. If RADIUS server is available and authentication is accepted, grant access. 3. If RADIUS servers fail to respond or return a reject response, try configured TACACS+ servers. 4. If TACACS+ server is available and authentication is accepted, grant access. 5. If TACACS+ server is available but authentication is rejected, deny access. 6. If both RADIUS and TACACS+ servers are not available, try password authentication. <p>NOTE: If either RADIUS or TACACS+ servers are available, password authentication is not attempted, because it is not explicitly configured in the authentication order.</p>
authentication-order [radius tacplus password];	<ol style="list-style-type: none"> 1. Try configured RADIUS authentication servers. 2. If RADIUS server is available and authentication is accepted, grant access. 3. If RADIUS servers fail to respond or return a reject response, try configured TACACS+ servers. 4. If TACACS+ server is available and authentication is accepted, grant access. 5. If TACACS+ servers fail to respond or return a reject response, try password authentication, because it is explicitly configured in the authentication order.

Table 153: Order of Authentication Attempts (*continued*)

Syntax	Order of Authentication Attempts
authentication-order tacplus;	<ol style="list-style-type: none"> 1. Try configured TACACS+ authentication servers. 2. If TACACS+ server is available and authentication is accepted, grant access. 3. If TACACS+ server is available but authentication is rejected, deny access. 4. If TACACS+ servers are not available, try password authentication. <p>NOTE: If a TACACS+ server is available, password authentication is not attempted, because it is not explicitly configured in the authentication order.</p>
authentication-order [tacplus password];	<ol style="list-style-type: none"> 1. Try configured TACACS+ authentication servers. 2. If TACACS+ servers fail to respond or return a reject response, try password authentication, because it is explicitly configured in the authentication order.
authentication-order [tacplus radius];	<ol style="list-style-type: none"> 1. Try configured TACACS+ authentication servers. 2. If TACACS+ server is available and authentication is accepted, grant access. 3. If TACACS+ servers fail to respond or return a reject response, try configured RADIUS servers. 4. If RADIUS server is available and authentication is accepted, grant access. 5. If RADIUS server is available but authentication is rejected, deny access. 6. If both TACACS+ and RADIUS servers are not available, try password authentication. <p>NOTE: If either TACACS+ or RADIUS servers are available, password authentication is not attempted, because it is not explicitly configured in the authentication order.</p>
authentication-order [tacplus radius password];	<ol style="list-style-type: none"> 1. Try configured TACACS+ authentication servers. 2. If TACACS+ server is available and authentication is accepted, grant access. 3. If TACACS+ servers fail to respond or return a reject response, try configured RADIUS servers. 4. If RADIUS server is available and authentication is accepted, grant access. 5. If RADIUS servers fail to respond or return a reject response try password authentication, because it is explicitly configured in the authentication order.

Table 153: Order of Authentication Attempts (*continued*)

Syntax	Order of Authentication Attempts
<code>authentication-order password;</code>	<ol style="list-style-type: none"> 1. Try to authenticate the user, using the password configured at the <code>[edit system login]</code> hierarchy level. 2. If the authentication is accepted, grant access. 3. If the authentication is rejected, deny access.



NOTE: If SSH public keys are configured, SSH user authentication first tries to perform public key authentication before using the authentication methods configured in the `authentication-order` statement. If you want SSH logins to use the authentication methods configured in the `authentication-order` statement without first trying to perform public key authentication, do not configure SSH public keys.

In a routing matrix based on a TX Matrix router, the authentication order must be configured only at the configuration groups `re0` and `re1`. The authentication order must not be configured at the `[edit system]` hierarchy. This is because the authentication order for the routing matrix is controlled on the switch-card chassis (or TX Matrix router) or switch-fabric chassis (for TX Matrix Plus router) only.

In Junos OS Release 10.0 and later, the superuser (belonging to the super-user login class) is also authenticated based on the authentication order that is configured for TACACS+, RADIUS, or password authentication using the `authentication-order` statement. For example, if the only configured authentication order is TACACS+, the superuser can only be authenticated by the TACACS+ server and password authentication cannot be used as an alternative. However, in Junos OS Release 9.6 and earlier, the superuser can use password authentication to login, even if password authentication is not configured explicitly using the `authentication-order` statement.

Related Documentation

- [Overview of Template Accounts for RADIUS and TACACS+ Authentication on page 1899](#)
- [Configuring the Junos OS Authentication Order for RADIUS, TACACS+, and Local Password Authentication](#)
- [Limiting the Number of User Login Attempts for SSH and Telnet Sessions](#)
- [Example: Configuring System Authentication for RADIUS, TACACS+, and Password Authentication on page 1855](#)

Juniper Networks Vendor-Specific RADIUS Attributes

Junos OS supports the configuration of Juniper Networks RADIUS vendor-specific attributes (VSAs). These VSAs are encapsulated in a RADIUS vendor-specific attribute

with the vendor ID set to the Juniper Networks ID number, 2636. [Table 154](#) lists the Juniper Networks VSAs you can configure.

Table 154: Juniper Networks Vendor-Specific RADIUS Attributes

Name	Description	Type	Length	String
Juniper-Local-User-Name	Indicates the name of the user template used by this user when logging in to a device. This attribute is used only in Access-Accept packets.	1	≥3	One or more octets containing printable ASCII characters.
Juniper-Allow-Commands	Contains an extended regular expression that enables the user to run operational mode commands in addition to the commands authorized by the user's login class permission bits. This attribute is used only in Access-Accept packets.	2	≥3	One or more octets containing printable ASCII characters, in the form of an extended regular expression. See "Regular Expressions for Allowing and Denying Junos OS Operational Mode Commands" on page 1827.
Juniper-Deny-Commands	Contains an extended regular expression that denies the user permission to run operation mode commands authorized by the user's login class permission bits. This attribute is used only in Access-Accept packets.	3	≥3	One or more octets containing printable ASCII characters, in the form of an extended regular expression. See "Regular Expressions for Allowing and Denying Junos OS Operational Mode Commands" on page 1827.
Juniper-Allow-Configuration	Contains an extended regular expression that enables the user to run configuration mode commands in addition to the commands authorized by the user's login class permission bits. This attribute is used only in Access-Accept packets.	4	≥3	One or more octets containing printable ASCII characters, in the form of an extended regular expression. See "Regular Expressions for Allowing and Denying Junos OS Configuration Mode Hierarchies" on page 1826.
Juniper-Deny-Configuration	Contains an extended regular expression that denies the user permission to run configuration commands authorized by the user's login class permission bits. This attribute is used only in Access-Accept packets.	5	≥3	One or more octets containing printable ASCII characters, in the form of an extended regular expression. See "Regular Expressions for Allowing and Denying Junos OS Configuration Mode Hierarchies" on page 1826.
Juniper-Interactive-Command	Indicates the interactive command entered by the user. This attribute is used only in Accounting-Request packets.	8	≥3	One or more octets containing printable ASCII characters.

Table 154: Juniper Networks Vendor-Specific RADIUS Attributes (*continued*)

Name	Description	Type	Length	String
Juniper-Configuration-Change	Indicates the interactive command that results in a configuration (database) change. This attribute is used only in Accounting-Request packets.	9	≥3	One or more octets containing printable ASCII characters.
Juniper-User-Permissions	<p>Contains information the server uses to specify user permissions. This attribute is used only in Access-Accept packets.</p> <p>NOTE: When the Juniper-User-Permissions attribute is configured to grant the Junos OS maintenance or all permissions on a RADIUS server, the UNIX wheel group membership is not automatically added to a user's list of group memberships. Some operations such as running the su root command from a local shell require wheel group membership permissions. However, when a user is configured locally with the permissions maintenance or all, the user is automatically granted membership to the UNIX wheel group. Therefore, we recommend that you create a template user account with the required permissions and associate individual user accounts with the template user account.</p>	10	≥3	<p>One or more octets containing printable ASCII characters.</p> <p>The string is a list of permission flags separated by a space. The exact name of each flag must be specified in its entirety. See Table 137.</p>
Juniper-Authentication-Type	Indicates the authentication method (local database, or RADIUS server) used to authenticate a user. If the user is authenticated using a local database, the attribute value shows 'local'. If the user is authenticated using RADIUS server, the attribute value shows 'remote'.	11	≥5	One or more octets containing printable ASCII characters.
Juniper-Session-Port	Indicates the source port number of the established session.	12	size of integer	Integer

For more information about the VSAs, see RFC 2138, *Remote Authentication Dial In User Service (RADIUS)*.

Related Documentation

- [Configuring RADIUS Server Authentication](#)

Configuring RADIUS System Accounting

With RADIUS accounting enabled, Juniper Networks routers or switches, acting as RADIUS clients, can notify the RADIUS server about user activities such as software logins, configuration changes, and interactive commands. The framework for RADIUS accounting is described in RFC 2866.

Tasks for configuring RADIUS system accounting are:

1. [Configuring Auditing of User Events on a RADIUS Server on page 1891](#)
2. [Specifying RADIUS Server Accounting and Auditing Events on page 1891](#)
3. [Configuring RADIUS Server Accounting on page 1892](#)

Configuring Auditing of User Events on a RADIUS Server

To audit user events, include the following statements at the **[edit system accounting]** hierarchy level:

```
[edit system accounting]
destination {
  radius {
    server {
      server-address {
        accounting-port port-number;
        max-outstanding-requests value;
        port port-number;
        retry value;
        secret password;
        source-address address;
        timeout seconds;
      }
    }
  }
}
```

Specifying RADIUS Server Accounting and Auditing Events

To specify the events you want to audit when using a RADIUS server for authentication, include the **events** statement at the **[edit system accounting]** hierarchy level:

```
[edit system accounting]
events [ events ];
```

events is one or more of the following:

- **login**—Audit logins
- **change-log**—Audit configuration changes

- **interactive-commands**—Audit interactive commands (any command-line input)

Configuring RADIUS Server Accounting

To configure RADIUS server accounting, include the **server** statement at the **[edit system accounting destination radius]** hierarchy level:

```
server {  
  server-address {  
    accounting-port port-number;  
    max-outstanding-requests value;  
    port port-number;  
    retry value;  
    secret password;  
    source-address address;  
    timeout seconds;  
  }  
}
```

server-address specifies the address of the RADIUS server. To configure multiple RADIUS servers, include multiple **server** statements.



NOTE: If no RADIUS servers are configured at the **[edit system accounting destination radius]** statement hierarchy level, the Junos OS uses the RADIUS servers configured at the **[edit system radius-server]** hierarchy level.

accounting-port *port-number* specifies the RADIUS server accounting port number.

The default port number is 1813.



NOTE: If you enable RADIUS accounting at the **[edit access profile *profile-name* accounting-order]** hierarchy level, accounting is triggered on the default port of 1813 even if you do not specify a value for the **accounting-port** statement.

You must specify a secret (password) that the local router or switch passes to the RADIUS client by including the **secret** statement. If the password contains spaces, enclose the entire password in quotation marks (" ").

In the **source-address** statement, specify a source address for the RADIUS server. Each RADIUS request sent to a RADIUS server uses the specified source address. The source address is a valid IPv4 address (in case if radius-server address is IPv4) or IPv6 address (in case if radius-server address is IPv6) configured on one of the router or switch interfaces.

Optionally, you can specify the number of times that the router or switch attempts to contact a RADIUS authentication server by including the **retry** statement. By default, the router or switch retries three times. You can configure the router or switch to retry from 1 through 10 times.

Optionally, you can specify the length of time that the local router or switch waits to receive a response from a RADIUS server by including the **timeout** statement. By default, the router or switch waits 3 seconds. You can configure the timeout to be from 1 through 90 seconds.

If you use the **enhanced-accounting** statement at the **[edit system radius-options]** hierarchy level, the RADIUS attributes such as access method, remote port, and access privileges can be audited. You can limit the number of attribute values to be displayed for auditing by using the **enhanced-avs-max <number>** statement at the **[edit system accounting]** hierarchy level.

```
[edit system radius-options]
enhanced-accounting;

[edit system accounting]
enhanced-avs-max <number>;
```

When a Juniper Networks router or switch is configured with RADIUS accounting, it sends **Accounting-Start** and **Accounting-Stop** messages to the RADIUS server. These messages contain information about user activities such as software logins, configuration changes, and interactive commands. This information is typically used for monitoring a network, collecting usage statistics, and ensuring that users are billed properly.

The following example shows three servers (10.5.5.5, 10.6.6.6, and 10.7.7.7) configured for RADIUS accounting:

```
system {
  accounting {
    events [ login change-log interactive-commands ];
    destination {
      radius {
        server {
          10.5.5.5 {
            accounting-port 3333;
            secret $9$dkafeqwrew;
            source-address 10.1.1.1;
            retry 3;
            timeout 3;
          }
          10.6.6.6 secret $9$fe3erqwrez;
          10.7.7.7 secret $9$f34929ftby;
        }
      }
    }
  }
}
```

Configuring RADIUS Authentication (QFX Series or OCX Series)

RADIUS authentication is a method of authenticating users who attempt to access the router or switch. Tasks to configure RADIUS authentication are:



NOTE: The `source-address` statement is not supported at the `[edit system radius-options]` or `[edit system-radius-server name]` hierarchies on the QFabric system.

- [Configuring RADIUS Server Details on page 1894](#)
- [Configuring MS-CHAPv2 for Password-Change Support on page 1895](#)
- [Specifying a Source Address for the Junos OS to Access External RADIUS Servers on page 1896](#)

Configuring RADIUS Server Details

To use RADIUS authentication on the router or switch, configure information about one or more RADIUS servers on the network by including one `radius-server` statement at the `[edit system]` hierarchy level for each RADIUS server:

```
[edit system]
radius-server server-address {
  accounting-port port-number;
  port number;
  retry number;
  secret password;
  source-address source-address;
  timeout seconds;
}
```

`server-address` is the address of the RADIUS server.

You can specify a port on which to contact the RADIUS server. By default, port number **1812** is used (as specified in RFC 2865). You can also specify an accounting port to send accounting packets. The default is **1813** (as specified in RFC 2866).

You must specify a password in the `secret password` statement. If the password contains spaces, enclose it in quotation marks. The secret used by the local router or switch must match that used by the server.

Optionally, you can specify the amount of time that the local router or switch waits to receive a response from a RADIUS server (in the `timeout` statement) and the number of times that the router or switch attempts to contact a RADIUS authentication server (in the `retry` statement). By default, the router or switch waits 3 seconds. You can configure this to be a value from 1 through 90 seconds. By default, the router or switch retries connecting to the server three times. You can configure this to be a value from 1 through 10 times.

You can use the `source-address` statement to specify a logical address for individual or multiple RADIUS servers.

To configure multiple RADIUS servers, include multiple `radius-server` statements.

To configure a set of users that share a single account for authorization purposes, you create a template user. To do this, include the `user` statement at the `[edit system login]`

hierarchy level, as described in [“Overview of Template Accounts for RADIUS and TACACS+ Authentication” on page 1899](#).

You can also configure RADIUS authentication at the **[edit access]** and **[edit access profile]** hierarchy level. Junos OS uses the following search order to determine which set of servers are used for authentication:

1. **[edit access profile *profile-name* radius-server *server-address*]**
2. **[edit access radius-server *server-address*]**
3. **[edit system radius-server *server-address*]**

Configuring MS-CHAPv2 for Password-Change Support

You can configure the Microsoft implementation of the Challenge Handshake Authentication Protocol version 2 (MS-CHAPv2) on the router or switch to support changing of passwords. This feature provides users accessing a router or switch the option of changing the password when the password expires, is reset, or is configured to be changed at the next login.

Before you configure MS-CHAPv2 for password-change support, ensure that you:

- Configure the RADIUS server authentication parameters
- Set the **authentication-order** to use the RADIUS server for the initial password attempt

To configure MS-CHAP-v2, include the following statements at the **[edit system radius-options]** hierarchy level:

```
[edit system radius-options]
password-protocol mschap-v2;
```

The following example shows statements for configuring the MS-CHAPv2 password protocol, password authentication order, and user accounts:

```
[edit]
system {
  authentication-order [ radius password ];
  radius-server {
    192.168.69.149 secret "$9$G-j.5Qz6tpBk.1hrlXxUjiq5Qn/C"; ## SECRET-DATA
  }
  radius-options {
    password-protocol mschap-v2;
  }
  login {
    user bob {
      class operator;
    }
  }
}
```

Specifying a Source Address for the Junos OS to Access External RADIUS Servers

You can specify which source address Junos OS uses when accessing your network to contact an external RADIUS server for authentication. You can also specify which source address Junos OS uses when contacting a RADIUS server for sending accounting information.

To specify a source address for a RADIUS server, include the **source-address** statement at the **[edit system radius-server server-address]** hierarchy level:

```
[edit system radius-server server-address]
source-address source-address;
```

source-address is a valid IP address configured on one of the router or switch interfaces.

Related Documentation

- [Example: Configuring RADIUS Authentication on page 1878](#)
- [Example: Configuring System Authentication for RADIUS, TACACS+, and Password Authentication on page 1855](#)
- [Juniper Networks Vendor-Specific RADIUS Attributes on page 1888](#)
- [Overview of Template Accounts for RADIUS and TACACS+ Authentication on page 1899](#)
- [Example: Configuring RADIUS Template Accounts on page 1879](#)
- [Using Regular Expressions on a RADIUS or TACACS+ Server to Allow or Deny Access to Commands on page 1876](#)
- [Junos OS User Authentication Methods on page 1804](#)

Using Regular Expressions on a RADIUS or TACACS+ Server to Allow or Deny Access to Commands

Use regular expressions to specify which operational or configuration mode commands are allowed or denied when you use a RADIUS or TACACS+ server for user authentication. You can specify the regular expressions using the appropriate Juniper Networks vendor-specific RADIUS or TACACS+ attributes in your authentication server configuration.

You can specify **allow-configuration**, **deny-configuration**, **allow-commands**, or **deny-commands** in a single extended regular expression, enclosing multiple commands in parentheses and separating them using the pipe symbol. For example, you can specify multiple **allow-commands** parameters using: **allow-commands= (cmd1 | cmd2 | cmdn)**. You can specify **user-permissions** as a list of comma-separated values, and not as a regular expression.

On a RADIUS or TACACS+ server, you can also use a simplified version for regular expressions where you specify each individual expression on a separate line. The simplified version is valid for **allow-commands**, **deny-commands**, **allow-configuration**, **deny-configuration**, and **permissions** vendor-specific attributes.

For a RADIUS server, specify the individual regular expressions using the following syntax:


```

Juniper-Allow-Commands+="cmd1"
Juniper-Allow-Commands+="cmd2"
Juniper-Allow-Commands+="cmdn"
Juniper-Deny-Commands+="cmd1"
Juniper-Deny-Commands+="cmd2"
Juniper-Deny-Commands+="cmdn"
Juniper-Allow-Configuration+="regex1"
Juniper-Allow-Configuration+="regex2"
Juniper-Allow-Configuration+="regexn"
Juniper-Deny-Configuration+="regex1"
Juniper-Deny-Configuration+="regex2"
Juniper-Deny-Configuration+="regexn"
Juniper-User-Permissions+="permission-flag1"
Juniper-User-Permissions+="permission-flag2"
Juniper-User-Permissions+="permission-flagn"

```

For TACACS+ server, specify the individual regular expressions using the following syntax:

```

allow-commands1="cmd1"
allow-commands2="cmd2"
allow-commandsn="cmdn"
deny-commands1="cmd1"
deny-commands2="cmd2"
deny-commandsn="cmdn"
allow-configuration1="regex1"
allow-configuration2="regex2"
allow-configurationn="regexn"
deny-configuration1="regex1"
deny-configuration2="regex2"
deny-configurationn="regexn"
user-permissions1="permission-flag1"
user-permissions2="permission-flag2"
user-permissionsn="permission-flagn"

```



NOTE:

- Numeric values 1 to *n* in the syntax (for TACACS+ server) must be unique but need not be sequential. For example, the following syntax is valid:

```

allow-commands1="cmd1"
allow-commands3="cmd3"
allow-commands2="cmd2"
deny-commands3="cmd3"
deny-commands2="cmd2"
deny-commands1="cmd1"

```

- The limit on the number of lines of individual regular expressions is imposed by the TACACS+ or RADIUS server.
- When you issue the `show cli authorization` command, the command output displays the regular expression in a single line, even if you specify each individual expression on a separate line.

For more information about Juniper Networks vendor-specific RADIUS and TACACS+ attributes, see [“Juniper Networks Vendor-Specific RADIUS Attributes” on page 1888](#) and [“Juniper Networks Vendor-Specific TACACS+ Attributes” on page 1863](#).



NOTE: When RADIUS or TACACS+ authentication is configured for a router, regular expressions configured on the RADIUS or TACACS+ server merge with any regular expressions configured on the local router at the [edit system login class] hierarchy level using the allow-commands, deny-commands, allow-configuration, deny-configuration, or permissions statements. If the final expression has a syntax error, the overall result is an invalid regular expression.

**Related
Documentation**

- [Junos OS Authentication Order for RADIUS, TACACS+, and Password Authentication on page 1859](#)

Example: Configuring RADIUS System Accounting

When a Juniper Networks router or switch is configured with RADIUS accounting, it sends **Accounting-Start** and **Accounting-Stop** messages to the RADIUS server. These messages contain information about user activities such as software logins, configuration changes, and interactive commands. This information is typically used for monitoring a network, collecting usage statistics, and ensuring that users are billed properly.

The following example shows three servers (10.5.5.5, 10.6.6.6, and 10.7.7.7) configured for RADIUS accounting:

```
system {
  accounting {
    events [ login change-log interactive-commands ];
    destination {
      radius {
        server {
          10.5.5.5 {
            accounting-port 3333;
            secret $9$dkafeqwrew;
            source-address 10.1.1.1;
            retry 3;
            timeout 3;
          }
          10.6.6.6 secret $9$fe3erqwrez;
          10.7.7.7 secret $9$f34929ftby;
        }
      }
    }
  }
}
```

**Related
Documentation**

- [Configuring RADIUS System Accounting on page 1891](#)

Configuring and Managing RADIUS Template Accounts

- [Overview of Template Accounts for RADIUS and TACACS+ Authentication on page 1899](#)
- [Example: Configuring RADIUS Template Accounts on page 1899](#)

Overview of Template Accounts for RADIUS and TACACS+ Authentication

When you use local password authentication, you must create a local user account for every user who wants to access the system. However, when you are using RADIUS or TACACS+ authentication, you can create single accounts (for authorization purposes) that are shared by a set of users. You create these accounts using the remote and local user template accounts. When a user is using a template account, the command-line interface (CLI) username is the login name; however, the privileges, file ownership, and effective user ID are inherited from the template account.

Related Documentation

- *[Understanding Remote Authentication Servers](#)*
- *[Configuring Remote Template Accounts for User Authentication](#)*
- *[Configuring Local User Template Accounts for User Authentication](#)*

Example: Configuring RADIUS Template Accounts

The following example shows how to configure RADIUS template accounts for different users or groups of users:

```
[edit]
system {
  login {
    user observation {
      uid 1001;
      class observation;
    }
    user operation {
      uid 1002;
      class operation;
    }
    user engineering {
      uid 1003;
```

```
        class engineering;  
    }  
}  
}
```

Related Documentation

- [Overview of Template Accounts for RADIUS and TACACS+ Authentication on page 1899](#)

CHAPTER 68

Configuring and Managing VSAs for RADIUS and TACACS+

- [Understanding Vendor-Specific Attributes \(VSAs\) on page 1901](#)
- [Juniper-Switching-Filter VSA Match Conditions and Actions on page 1902](#)
- [Juniper Networks Vendor-Specific TACACS+ Attributes on page 1904](#)

Understanding Vendor-Specific Attributes (VSAs)

Devices support the configuration of RADIUS server attributes specific to Juniper Networks. These attributes are known as vendor-specific attributes (VSAs) and are described in RFC 2138, *Remote Authentication Dial In User Service* (RADIUS).

Through VSAs, you can configure port-filtering attributes on the RADIUS server. VSAs are cleartext fields sent from the RADIUS server to the device as a result of authentication success or failure. Authentication prevents unauthorized user access by blocking a supplicant at the port until the device is authenticated by the RADIUS server. The VSA attributes are interpreted by the device during authentication, and the device takes appropriate actions. Implementing port-filtering attributes with authentication on the RADIUS server provides a central location for controlling LAN access for supplicants.

These port-filtering attributes specific to Juniper Networks are encapsulated in a RADIUS server VSA with the vendor ID set to the Juniper Networks ID number, 2636.

As well as configuring port-filtering attributes through VSAs, you can apply a port firewall filter that has already been configured on the device directly to the RADIUS server. Like port-filtering attributes, the filter is applied during the authentication process, and its actions are applied at the device port. Adding a port firewall filter to a RADIUS server eliminates the need to add the filter to multiple ports and devices.

Related Documentation

- [Configuring Firewall Filters on page 5978](#)
- [Configuring RADIUS Authentication \(QFX Series or OCX Series\) on page 1873](#)
- [Juniper-Switching-Filter VSA Match Conditions and Actions on page 1902](#)

Juniper-Switching-Filter VSA Match Conditions and Actions

Switching devices support the configuration of RADIUS server attributes specific to Juniper Networks, which are known as vendor-specific attributes (VSAs). The Juniper-Switching-Filter VSA works in conjunction with 802.1X authentication to centrally control access of supplicants to the network. You can use this VSA to configure filters on the RADIUS server, which are sent to the switch and applied to users that have been authenticated using 802.1X authentication.

The Juniper-Switching-Filter VSA can contain one or more filter terms. Filter terms are configured using one or more *match conditions* with a resulting *action*. Match conditions are the criteria that a packet must meet for a configured action to be applied on it. The action is the action that the switch takes if a packet meets the criteria in the match conditions. The action that the switch can take is either accept or deny a packet.

The following guidelines apply when you specify match conditions and actions for VSAs:

- Both **match** and **action** statements are mandatory.
- If no match condition is specified, any packet is considered a match by default.
- If no action is specified, the default action is to deny the packet.
- Any or all options can be included in each **match** and **action** statement.
- The AND operation is performed on fields that are of a different type, which are separated by commas. Fields of the same type cannot be repeated.
- For the **forwarding-class** option to be applied, the forwarding class must be configured on the switch. If the forwarding class is not configured on the switch, this option is ignored.

Table 155 describes the match conditions that you can specify when you configure a VSA attribute as a firewall filter by using the **match** command on the RADIUS server. The string that defines a match condition is called a *match statement*.

Table 155: Match Conditions

Option	Description
destination-mac <i>mac-address</i>	Destination media access control (MAC) address of the packet.
source-vlan <i>source-vlan</i>	Name of the source VLAN.
source-dot1q-tag <i>tag</i>	Tag value in the 802.1Q header, in the range 0 through 4095.
destination-ip <i>ip-address</i>	Address of the final destination node.
ip-protocol <i>protocol-id</i>	IPv4 protocol value. In place of the numeric value, you can specify one of the following text synonyms: ah , egp (8), esp (50), gre (47), icmp (1), igmp (2), ipip (4), ipv6 (41), ospf (89), pim (103), rsvp (46), tcp (6), or udp (17)

Table 155: Match Conditions (*continued*)

Option	Description
source-port <i>port</i>	TCP or User Datagram Protocol (UDP) source port field. Normally, you specify this match statement in conjunction with the ip-protocol match statement to determine which protocol is being used on the port. In place of the numeric field, you can specify one of the text options listed under destination-port .
destination-port <i>port</i>	<p>TCP or UDP destination port field. Normally, you specify this match statement in conjunction with the ip-protocol match statement to determine which protocol is being used on the port. In place of the numeric value, you can specify one of the following text synonyms (the port numbers are also listed):</p> <p>afs (1483), bgp (179), biff (512), bootpc (68), bootps (67), cvspserver (2401), cmd (514), dhcp (67), domain (53), eklogin (2105), ekshell (2106), exec (512), finger (79), ftp (21), ftp-data (20), http (80), https (443), ident (113), imap (143), kerberos-sec (88), klogin (543), kpasswd (761), krb-prop (754), krbupdate (760), kshell (544), ldap (389), login (513), mobileip-agent (434), mobilip-mn (435), msdp (639), netbios-dgm (138), netbios-ns (137), netbios-ssn (139), nfsd (2049), nntp (119), ntalk (518), ntp (123), pop3 (110), pptp (1723), printer (515), radacct (1813), radius (1812), rip (520), rkinit (2108), smtp (25), snmp (161), snmptrap (162), snpp (444), socks (1080), ssh (22), sunrpc (111), syslog (514), telnet (23), tacacs-ds (65), talk (517), tftp (69), timed (525), who (513), xdmcp (177), zephyr-clt (2103), zephyr-hm (2104)</p>

When you define one or more terms that specify the filtering criteria, you also define the action to take if the packet matches all criteria. [Table 156](#) shows the actions that you can specify in a term.

Table 156: Actions for VSAs

Option	Description
(allow deny)	Accept a packet or discard a packet silently without sending an Internet Control Message Protocol (ICMP) message.
forwarding-class <i>class-of-service</i>	<p>(Optional) Classify the packet in one of the following forwarding classes:</p> <ul style="list-style-type: none"> assured-forwarding best-effort expedited-forwarding network-control
loss-priority (low medium high)	(Optional) Set the packet loss priority (PLP) to low , medium , or high . Specify both the forwarding class and the loss priority.

Related Documentation

- [Filtering 802.1X Supplicants by Using RADIUS Server Attributes](#)
- [Understanding Dynamic Filters Based on RADIUS Attributes](#)
- [Understanding Vendor-Specific Attributes \(VSAs\) on page 1901](#)

Juniper Networks Vendor-Specific TACACS+ Attributes

Junos OS supports the configuration of Juniper Networks TACACS+ vendor-specific attributes (VSAs). These VSAs are encapsulated in a TACACS+ vendor-specific attribute with the vendor ID set to the Juniper Networks ID number, 2636. [Table 151](#) lists the Juniper Networks VSAs you can configure.

Table 157: Juniper Networks Vendor-Specific TACACS+ Attributes

Name	Description	Length	String
local-user-name	Indicates the name of the user template used by this user when logging in to a device.	≥3	One or more octets containing printable ASCII characters.
allow-commands	Contains an extended regular expression that enables the user to run operational mode commands in addition to those commands authorized by the user's login class permission bits.	≥3	One or more octets containing printable ASCII characters, in the form of an extended regular expression. See Table 143 .
allow-configuration	Contains an extended regular expression that enables the user to run configuration mode commands in addition to those commands authorized by the user's login class permission bits.	≥3	One or more octets containing printable ASCII characters, in the form of an extended regular expression. See "Regular Expressions for Allowing and Denying Junos OS Configuration Mode Hierarchies" on page 1826.
deny-commands	Contains an extended regular expression that denies the user permission to run operational mode commands authorized by the user's login class permission bits.	≥3	One or more octets containing printable ASCII characters, in the form of an extended regular expression. See Table 143 .
deny-configuration	Contains an extended regular expression that denies the user permission to run configuration mode commands authorized by the user's login class permission bits.	≥3	One or more octets containing printable ASCII characters, in the form of an extended regular expression. See Table 142 .

Table 157: Juniper Networks Vendor-Specific TACACS+ Attributes (*continued*)

Name	Description	Length	String
user-permissions	<p>Contains information the server uses to specify user permissions.</p> <p>NOTE: When the user-permissions attribute is configured to grant the Junos OS maintenance or all permissions on an IPv4 or IPv6 TACACS+ server, the UNIX wheel group membership is not automatically added to a user's list of group memberships. Some operations such as running the su root command from a local shell require wheel group membership permissions. However, when a user is configured locally with the permissions maintenance or all, the user is automatically granted membership to the UNIX wheel group. Therefore, we recommend that you create a template user account with the required permissions and associate individual user accounts with the template user account.</p>	≥3	One or more octets containing printable ASCII characters. See Table 137.
authentication-type	Indicates the authentication method (local database, or TACACS+ server) used to authenticate a user. If the user is authenticated using a local database, the attribute value shows 'local'. If the user is authenticated using TACACS+ server, the attribute value shows 'remote'.	≥5	One or more octets containing printable ASCII characters.
session-port	Indicates the source port number of the established session.	size of integer	Integer

Related Documentation

- *Configuring TACACS+ Authentication*

PART 26

Configuration Statements and Operational Commands

- Configuration Statements on page 1909
- Operational Commands on page 1993

CHAPTER 69

Configuration Statements

- [access](#) on page 1911
- [accounting \(Access Profile\)](#) on page 1912
- [accounting-options](#) on page 1913
- [accounting-server](#) on page 1915
- [accounting-stop-on-access-deny](#) on page 1916
- [accounting-stop-on-failure](#) on page 1917
- [advertisement-interval](#) on page 1918
- [agent-address](#) on page 1919
- [archival](#) on page 1920
- [archive-sites \(Configuration File\)](#) on page 1921
- [authentication-order](#) on page 1922
- [authentication-server](#) on page 1923
- [authorization](#) on page 1924
- [categories](#) on page 1925
- [client-list](#) on page 1925
- [client-list-name](#) on page 1926
- [clients](#) on page 1926
- [commit-delay](#) on page 1927
- [community \(SNMP\)](#) on page 1928
- [configuration](#) on page 1929
- [connection-limit](#) on page 1930
- [contact](#) on page 1931
- [disable \(LLDP\)](#) on page 1931
- [falling-threshold \(Health Monitor\)](#) on page 1932
- [filter-duplicates](#) on page 1932
- [full-name](#) on page 1933
- [health-monitor](#) on page 1933
- [hold-multiplier](#) on page 1934

- [idle-timeout \(Access\) on page 1935](#)
- [interface \(LLDP\) on page 1936](#)
- [interval \(Health Monitor\) on page 1937](#)
- [lldp on page 1938](#)
- [lldp-configuration-notification-interval on page 1940](#)
- [location on page 1940](#)
- [management-address on page 1941](#)
- [name on page 1942](#)
- [nas-ip-address on page 1942](#)
- [nonvolatile on page 1943](#)
- [oid on page 1943](#)
- [order on page 1944](#)
- [port \(RADIUS Server\) on page 1945](#)
- [profile on page 1946](#)
- [protocols on page 1947](#)
- [protocol-version on page 1960](#)
- [ptopo-configuration-maximum-hold-time on page 1961](#)
- [ptopo-configuration-trap-interval on page 1961](#)
- [radius on page 1962](#)
- [radius-options \(edit system\) on page 1963](#)
- [radius-server on page 1964](#)
- [rate-limit on page 1965](#)
- [remote-debug-permission on page 1966](#)
- [retry on page 1967](#)
- [rising-threshold \(Health Monitor\) on page 1968](#)
- [root-login on page 1969](#)
- [services \(Switches\) on page 1970](#)
- [snmp on page 1971](#)
- [ssh on page 1975](#)
- [system on page 1976](#)
- [tacplus-options on page 1982](#)
- [targets on page 1983](#)
- [traceoptions \(LLDP\) on page 1984](#)
- [transfer-interval \(Configuration\) on page 1986](#)
- [transfer-on-commit on page 1987](#)
- [trap-group on page 1988](#)
- [trap-options on page 1989](#)

- [user \(Access\)](#) on page 1990
- [version](#) on page 1991

access

Syntax	<pre> access { address-assignment pool <i>pool-name</i> address-pool <i>pool-name</i> profile <i>profile-name</i> { accounting (Access Profile) { accounting-stop-on-access-deny; accounting-stop-on-failure; (authentication-order (Access Profile) (ldap radius none); order (radius none); } radius { accounting-server [<i>server-addresses</i>]; authentication-server [<i>server-addresses</i>]; } } } </pre>
Hierarchy Level	[edit]
Release Information	<p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.</p>
Description	<p>Configure authentication, authorization, and accounting (AAA) services.</p> <p>The statements are explained separately.</p>
<div>  NOTE: The [edit access] hierarchy is not available on QFabric systems. </div>	
Default	Not enabled
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring 802.1X RADIUS Accounting (CLI Procedure)

accounting (Access Profile)

Syntax	<pre>accounting { accounting-stop-on-access-deny; accounting-stop-on-failure; order (radius none); }</pre>
Hierarchy Level	[edit access profile <i>profile-name</i>]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Description	Configure the authentication order for authentication, authorization, and accounting (AAA) services.
Default	Not enabled
Options	none —Use no authentication for specified subscribers. radius —Use RADIUS authentication for specified subscribers. The remaining statements are explained separately.



NOTE: The [edit access] hierarchy is not available on QFabric systems.

Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch</i>• <i>Configuring 802.1X RADIUS Accounting (CLI Procedure)</i>• <i>Understanding 802.1X and RADIUS Accounting on EX Series Switches</i>• <i>Configuring RADIUS Accounting</i>• Understanding RADIUS Accounting on page 1883

accounting-options

```

Syntax  accounting-options {
    class-usage-profile profile-name {
        destination-classes {
            destination-class-name;
        }
        file filename;
        interval minutes;
        source-classes {
            source-class-name;
        }
    }
    file filename {
        archive-sites {
            site-name;
        }
        files number;
        nonpersistent;
        size bytes;
        start-time time;
        transfer-interval minutes;
    }
    filter-profile profile-name {
        counters {
            counter-name;
        }
        file filename;
        interval minutes;
    }
    interface-profile profile-name {
        fields {
            input-bytes;
            input-errors;
            input-multicast;
            input-packets;
            input-unicast;
            output-bytes;
            output-errors;
            output-multicast;
            output-packets;
            output-unicast;
            rpf-check-bytes;
            rpf-check-packets;
            rpf-check6-bytes;
            rpf-check6-packets;
            unsupported-protocol;
        }
        file filename;
        interval minutes;
    }
    mib-profile profile-name {
        file filename;
        interval minutes;
    }
}

```

```
object-names {
    mib-object-name;
}
operation (get | get-next | walk);
}
policy-decision-statistics-profile profile-name {
    application-aware-access-list-fields {
        address;
        application;
        application-group;
        input-bytes;
        input-interface;
        input-packets;
        mask;
        output-bytes;
        output-packets;
        subscriber-name;
        timestamp;
        vrf-name;
    }
    file filename;
}
routing-engine-profile profile-name {
    fields {
        field-name;
    }
    file filename;
    interval minutes;
}
}
```

Hierarchy Level	[edit]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Description	Configure options for accounting statistics collection.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding RADIUS Accounting on page 1883• Understanding Vendor-Specific Attributes (VSAs) on page 1901• Configuring RADIUS System Accounting on page 1891• Configuring Remote Template Accounts for User Authentication• Configuring Local User Template Accounts for User Authentication

accounting-server

Syntax	<code>accounting-server[server-addresses];</code>
Hierarchy Level	[edit access profile <i>profile-name</i> radius]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Description	Configure the Remote Authentication Dial-In User Service (RADIUS) server for authentication. To configure multiple RADIUS servers, include multiple server addresses. The servers are tried in order and in a round-robin fashion until a valid response is received from one of the servers or until all the configured retry limits are reached.
Default	Not enabled
Options	<i>server-addresses</i> —One or more addresses of RADIUS authentication servers.



NOTE: The [edit access] hierarchy is not available on QFabric systems.

Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> <i>show network-access aaa statistics authentication</i> <i>Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch</i> <i>Understanding 802.1X and RADIUS Accounting on EX Series Switches</i> Understanding RADIUS Accounting on page 1883

accounting-stop-on-access-deny


Syntax	accounting-stop-on-access-deny;
Hierarchy Level	[edit access profile <i>profile-name</i> accounting]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Description	Configure the authentication order for authentication, authorization, and accounting (AAA) services to send an Acct-Stop message if the AAA server denies access to a supplicant.



NOTE: The [edit access] hierarchy is not available on QFabric systems.

Default	Not enabled
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch</i>• <i>Configuring 802.1X RADIUS Accounting (CLI Procedure)</i>• <i>show network-access aaa statistics authentication</i>• <i>Configuring RADIUS Accounting</i>

accounting-stop-on-failure

Syntax	accounting-stop-on-failure;
Hierarchy Level	[edit access profile <i>profile-name</i> accounting]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Description	<p>Configure RADIUS accounting to send an Acct-Stop message when a subscriber session has been successfully authenticated and authorized, but then fails before an Acct-Start message is sent. By default, an Acct-Stop message is sent only if an Acct-Start message has been exchanged with the accounting server.</p> <p>Consider a situation where RADIUS address pools are used to assign IP/IPv6 addresses. After a subscriber session is successfully authenticated, the RADIUS server authorizes the session by assigning an IP address from the RADIUS address pool and conveying that address in the Framed-IP-Address attribute. If a negotiation failure occurs at this point, the session is terminated before activating. The Acct-Start message is never sent because it is initiated by session activation. By default, an Acct-Stop message cannot be sent because the Acct-Start is never sent. However, if the acct-stop-on-failure statement is configured, the negotiation failure causes the Acct-Stop message to be sent, which explicitly notifies the RADIUS server that the session is disconnected and that it can free the allocated IP address back to the pool.</p>
	<p> NOTE: The [edit access] hierarchy is not available on QFabric systems.</p>
Default	Not enabled
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch</i> • <i>Configuring 802.1X RADIUS Accounting (CLI Procedure)</i> • <i>Understanding 802.1X and RADIUS Accounting on EX Series Switches</i> • <i>Configuring RADIUS Accounting</i> • Understanding RADIUS Accounting on page 1883

advertisement-interval

Syntax	<code>advertisement-interval seconds;</code>
Hierarchy Level	[edit protocols lldp], [edit routing-instances <i>routing-instance-name</i> protocols lldp]
Release Information	Statement introduced in Junos OS Release 9.6 for MX Series and T Series routers. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Description	(MX Series and T Series routers only) Configure an interval for LLDP advertisement.
Options	seconds —Interval between LLDP advertisement. Default: 30 Range: 5 through 32768
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring LLDP</i>• show lldp on page 2002• <i>Configuring LLDP (CLI Procedure)</i>• <i>Understanding LLDP and LLDP-MED on EX Series Switches</i>• <i>transmit-delay</i>• Understanding LLDP on page 1795

agent-address

Syntax	agent-address outgoing-interface;
Hierarchy Level	[edit snmp trap-options]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Description	Set the agent address of all SNMPv1 traps generated by this router or switch. Currently, the only option is outgoing-interface , which sets the agent address of each SNMPv1 trap to the address of the outgoing interface of that trap.
Options	outgoing-interface —Value of the agent address of all SNMPv1 traps generated by this router or switch. The outgoing-interface option sets the agent address of each SNMPv1 trap to the address of the outgoing interface of that trap. Default: Disabled (the agent address is not specified in SNMPv1 traps).
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Agent Address for SNMP Traps</i>

archival

```
Syntax archival {
    configuration {
        archive-sites {
            file://<path>/<filename>;
            ftp://username@host:<port>url-path password password;
            http://username@host:<port>url-path password password;
            pasvftp://username@host:<port>url-path password password;
            scp://username@host:<port>url-path password password;
        }
        transfer-interval interval;
        transfer-on-commit;
    }
}
```

Hierarchy Level [edit system]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.

Description Configure copying of the currently active configuration to an archive site. An archive site can be a file, or an FTP, HTTP, or SCP location.

Options The remaining statements are explained separately.





NOTE: The [edit system archival] hierarchy is not available on QFabric systems.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- [Using Junos OS to Configure a Router or Switch to Transfer Its Configuration to an Archive Site on page 1333](#)

archive-sites (Configuration File)

Syntax	<pre>archive-sites { file://<path>/<filename>; ftp://username@host:<port>url-path password password; http://username@host:<port>url-path password password; pasvftp://username@host:<port>url-path password password; scp://username@host:<port>url-path password password; }</pre>
Hierarchy Level	[edit system archival configuration]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.</p>
Description	<p>Specify where to transfer the current configuration files. When specifying a URL in a Junos OS statement using an IPv6 host address, you must enclose the entire URL in quotation marks (" ") and enclose the IPv6 host address in brackets ([]). For example, "scp://username<:password>@[ipv6-host-address]<:port>/url-path"</p> <p>If you specify more than one archive site, the router or switch attempts to transfer the configuration files to the first archive site in the list, moving to the next only if the transfer fails.</p> <p>The destination filename is saved in the following format, where <i>n</i> corresponds to the number of the compressed configuration rollback file that has been archived:</p> <p><i>router-name_YYYYMMDD_HHMMSS_juniper.conf.n.gz</i></p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p> NOTE: The time included in the destination filename is always in Coordinated Universal Time (UTC) regardless of whether the time on the router or switch is configured as UTC or the local time zone. The default time zone on the router or switch is UTC.</p> </div> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p> NOTE: The [edit system archival] hierarchy is not available on QFabric systems.</p> </div>
Options	<p>The prefix used in the configuration statement determines the form of transfer:</p> <p>file:// —transfer on a path to a named file</p> <p>ftp:// —transfer using active FTP server</p> <p>http:// —transfer using HTTP server</p>

pasvftp:// —transfer to a device that only accepts passive FTP services

scp:// —transfer to a known host using background SCP file transfers

Required Privilege Level system—To view this statement in the configuration.
 system-control—To add this statement to the configuration.

Related Documentation

- [Configuring Archive Sites for the Transfer of Active Configuration Files on page 1334](#)
- [Junos OS Commit Model for Router or Switch Configuration on page 1313](#)
- [configuration on page 1363](#)
- [transfer-on-commit on page 1365](#)

authentication-order

Syntax authentication-order [none | password | radius];

Hierarchy Level [edit [access profile](#) *profile-name*],
 [edit [system](#)]

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches.
 Statement introduced in Junos OS Release 11.1 for the QFX Series.
 Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.

Description Configure the order of authentication, authorization, and accounting (AAA) servers to use while sending authentication messages.

Default Not enabled

Options **none**—No authentication for specified subscribers.

password—Password authentication.

radius—RADIUS authentication.




NOTE: The [edit [access](#)] hierarchy is not available on QFabric systems.

Required Privilege Level admin—To view this statement in the configuration.
 admin-control—To add this statement to the configuration.

authentication-server

Syntax	<code>authentication-server [server-addresses];</code>
Hierarchy Level	[edit access profile <i>profile-name</i> radius]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Description	Configure the RADIUS server for authentication. To configure multiple RADIUS servers, include multiple server addresses. The servers are tried in order and in a round-robin fashion until a valid response is received from one of the servers or until all the configured retry limits are reached.
Options	server-addresses —Configure one or more RADIUS server addresses.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch</i>• <i>show network-access aaa statistics authentication</i>

authorization

Syntax	<code>authorization <i>authorization</i>;</code>
Hierarchy Level	[edit snmp community <i>community-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Description	Set the access authorization for SNMP Get , GetBulk , GetNext , and Set requests.
Options	<p><i>authorization</i>—Access authorization level:</p> <ul style="list-style-type: none">• read-only—Enable Get, GetNext, and GetBulk requests.• read-write—Enable all requests, including Set requests. You must configure a view to enable Set requests.
	<div> NOTE: The read-write option is not supported on the QFX3000 QFabric system.</div>
	Default: read-only
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the SNMP Community String on page 1503

categories

Syntax	<code>categories { category; }</code>
Hierarchy Level	<code>[edit snmp trap-group group-name]</code>
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Description	Define the types of traps that are sent to the targets of the named trap group.
Default	If you omit the categories statement, all trap types are included in trap notifications.
Options	category —Name of a trap type: authentication , chassis , configuration , link , remote-operations , rmon-alarm , or startup .
Required Privilege Level	snmp —To view this statement in the configuration. snmp-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring SNMP Trap Groups on page 1504

client-list

Syntax	<code>client-list client-list-name { ip-addresses; }</code>
Hierarchy Level	<code>[edit snmp]</code>
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Description	Define a list of SNMP clients.
Options	client-list-name —Name of the client list. ip-addresses —IP addresses of the SNMP clients to be added to the client list,
Required Privilege Level	snmp —To view this statement in the configuration. snmp-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Adding a Group of Clients to an SNMP Community on page 1505

client-list-name

Syntax	<code>client-list-name <i>client-list-name</i>;</code>
Hierarchy Level	<code>[edit snmp community <i>community-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Add a client list or prefix list to an SNMP community.
Options	<i>client-list-name</i> —Name of the client list or prefix list.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Adding a Group of Clients to an SNMP Community on page 1505


clients

Syntax	<pre>clients { <i>address</i> <restrict>; }</pre>
Hierarchy Level	<code>[edit snmp community <i>community-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify the IPv4 or IPv6 addresses of the SNMP client hosts that are authorized to use this community.
Default	If you omit the clients statement, all SNMP clients using this community string are authorized to access the switch.
Options	<i>address</i> —Address of an SNMP client that is authorized to access this switch. You must specify an address, not a hostname. To specify more than one client, include multiple <i>address</i> options. <i>restrict</i> —(Optional) Do not allow the specified SNMP client to access the switch.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring SNMP Communities


commit-delay

Syntax	<code>commit-delay <i>seconds</i>;</code>
Hierarchy Level	[edit snmp nonvolatile]
Release Information	Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the timer for the SNMP Set reply and start of the commit.
Options	<i>seconds</i> —Delay between an affirmative SNMP Set reply and start of the commit operation. Default: 5 seconds
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Commit Delay Timer</i>

community (SNMP)

Syntax	<pre>community <i>community-name</i> { authorization <i>authorization</i>; client-list-name <i>client-list-name</i>; clients { address restrict; } view <i>view-name</i>; }</pre>
Hierarchy Level	[edit snmp]
Release Information	Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Define an SNMP community. An SNMP community authorizes SNMP clients based on the source IP address of incoming SNMP request packets. A community also defines which MIB objects are available and the operations (read-only or read-write) allowed on those objects.
<div> NOTE: The authorization read-write option is not supported on the QFX3000 QFabric system.</div>	
The SNMP client application specifies an SNMP community name in Get , GetBulk , GetNext , and Set SNMP requests.	
Default	If you omit the community statement, all SNMP requests are denied.
Options	<i>community-name</i> —Community string. If the name includes spaces, enclose it in quotation marks (" "). The remaining statements are explained separately.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the SNMP Community String on page 1503

configuration

Syntax	<pre> configuration { transfer-interval interval; transfer-on-commit; archive-sites { file://<path>/<filename>; ftp://username@host:<port>url-path password password; http://username@host:<port>url-path password password; pasvftp://username@host:<port>url-path password password; scp://username@host:<port>url-path password password; } } </pre>
Hierarchy Level	[edit system archival]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.</p>
Description	Configure the router or switch to periodically transfer its currently active configuration (or after each commit).
<div>  <p>NOTE: The [edit system archival] hierarchy is not available on QFabric systems.</p> </div>	
Options	The remaining statements are explained separately.
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Using Junos OS to Configure a Router or Switch to Transfer Its Configuration to an Archive Site on page 1333 • archive on page 1692 • archive-sites on page 1361 • transfer-interval on page 1364 • transfer-on-commit on page 1365

connection-limit

Syntax	<code>connection-limit <i>limit</i>;</code>
Hierarchy Level	<code>[edit system services finger],</code> <code>[edit system services ftp],</code> <code>[edit system services netconf ssh],</code> <code>[edit system services ssh],</code> <code>[edit system services telnet],</code> <code>[edit system services xnm-clear-text],</code> <code>[edit system services xnm-ssl]</code>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	Configure the maximum number of connections sessions for each type of system services (finger, ftp, ssh, telnet, xnm-clear-text, or xnm-ssl) per protocol (either IPv6 or IPv4).
Options	<p>limit—(Optional) Maximum number of established connections per protocol (either IPv6 or IPv4).</p> <p>Range: 1 through 250</p> <p>Default: 75</p>



NOTE: The actual number of maximum connections depends on the availability of system resources, and might be fewer than the configured `connection-limit` value if the system resources are limited.

Required Privilege	system—To view this statement in the configuration.
Level	system-control—To add this statement to the configuration.

Related Documentation	<ul style="list-style-type: none"> • Configuring clear-text or SSL Service for Junos XML Protocol Client Applications • Configuring DTCP-over-SSH Service for the Flow-Tap Application • Configuring Finger Service for Remote Access to the Router • Configuring FTP Service for Remote Access to the Router or Switch • Configuring SSH Service for Remote Access to the Router or Switch on page 1412 • Configuring Telnet Service for Remote Access to a Router or Switch
------------------------------	---

contact

Syntax	<code>contact <i>contact</i>;</code>
Hierarchy Level	<code>[edit snmp]</code>
Release Information	Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Define the value of the MIB II sysContact object, which is the contact person for the managed system.
Options	contact —Name of the contact person. If the name includes spaces, enclose it in quotation marks (" ").
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the System Contact on a Device Running Junos OS</i>

disable (LLDP)

Syntax	<code>disable;</code>
Hierarchy Level	<code>[edit protocols lldp],</code> <code>[edit protocols interface lldp]</code>
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Disable the LLDP configuration on the switch or on one or more interfaces.
Default	If you do not configure LLDP, it is disabled on the switch and on specific switch interfaces.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • show lldp on page 2002 • <i>Configuring LLDP (CLI Procedure)</i> • <i>Understanding LLDP and LLDP-MED on EX Series Switches</i> • <i>Configuring LLDP</i> • Understanding LLDP on page 1795

falling-threshold (Health Monitor)

Syntax	<code>falling-threshold <i>percentage</i>;</code>
Hierarchy Level	<code>[edit snmp health-monitor]</code>
Release Information	Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Set the lower threshold for the monitored object when you configure a health monitor alarm. By setting a rising and a falling threshold for a monitored variable, you can be alerted whenever the value of the variable falls outside the allowable operational range.
Options	<i>percentage</i> —Lower threshold for the alarm entry. Range: 1 through 100 Default: 70 percent of the maximum possible value
Required Privilege Level	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• rising-threshold on page 1653• Configuring Health Monitoring on page 1509

filter-duplicates

Syntax	<code>filter-duplicates;</code>
Hierarchy Level	<code>[edit snmp]</code>
Release Information	Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Filter duplicate Get , GetNext , or GetBulk SNMP requests.
Required Privilege Level	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding the Implementation of SNMP on the QFabric System• Example: Configuring SNMP on page 1523

full-name

Syntax	<code>full-name <i>complete-name</i>;</code>
Hierarchy Level	[edit system login user]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Description	Configure the complete name of a user.
Options	<i>complete-name</i> —Full name of the user. If the name contains spaces, enclose it in quotation marks.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Junos OS User Accounts by Using a Configuration Group</i> • <i>user</i>

health-monitor

Syntax	<pre>health-monitor { falling-threshold <i>percentage</i>; interval <i>seconds</i>; rising-threshold <i>percentage</i>; }</pre>
Hierarchy Level	[edit snmp]
Release Information	Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure health monitoring. The remaining statements are explained separately.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Health Monitoring on page 1509 • Understanding Health Monitoring on page 1467

hold-multiplier

Syntax	hold-multiplier <i>number</i> ;
Hierarchy Level	[edit protocols lldp]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for QFX Series.
Description	Specify the multiplier used in combination with the advertisement-interval value to determine the length of time LLDP information is held before it is discarded. The default value is 4 (or 120 seconds).
Default	Disabled.
Options	<i>number</i> —A number used as a multiplier. Range: 2 through 10 Default: 4 (or 120 seconds)
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• show lldp on page 2002• <i>Configuring LLDP (CLI Procedure)</i>• <i>Understanding LLDP and LLDP-MED on EX Series Switches</i>• <i>Configuring LLDP</i>• Understanding LLDP on page 1795

idle-timeout (Access)

Syntax	<code>idle-timeout seconds;</code>
Hierarchy Level	<code>[edit access group-profile <i>profile-name</i> ppp],</code> <code>[edit access profile <i>profile-name</i> client <i>client-name</i> ppp]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Description	Configure the idle timeout for a user. The router might consider a PPP session to be idle because of the following reasons: <ul style="list-style-type: none"> • There is no ingress traffic on the PPP session. • There is no egress traffic. • There is neither ingress or egress traffic on the PPP session. • There is no ingress or egress PPP control traffic. This is applicable only if keepalives are enabled.
Options	seconds —Number of seconds a user can remain idle before the session is terminated. Range: 0 through 4,294,967,295 seconds Default: 0




NOTE: The `[edit access]` hierarchy is not available on QFabric systems.

Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
---------------------------------	---

Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the Group Profile for Defining L2TP Attributes</i> • <i>Configuring PPP Properties for a Client-Specific Profile</i> • <i>Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile</i>
------------------------------	--

interface (LLDP)

Syntax	<pre>interface (all <i>interface-name</i>) { disable; power-negotiation { disable; } }</pre>
Hierarchy Level	[edit protocols lldp]
Release Information	<p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	Configure Link Layer Discovery Protocol (LLDP) on all interfaces or on a specific interface.
	<div>  <p>NOTE: On EX4300 switches, LLDP cannot be configured on the me0 or vme interface. Issuing the command <code>set protocols lldp interface me0</code> generates the following error message:</p> <pre>error: name: 'me0': Invalid interface error: statement creation failed: interface</pre> <p>Issuing the command <code>set protocols lldp interface vme</code> generates the following error message:</p> <pre>error: name: 'vme': Invalid interface error: statement creation failed: interface</pre> </div>
Default	None
Options	<p>all—All interfaces on the switch.</p> <p><i>interface-name</i>—Name of a specific interface.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring LLDP (CLI Procedure) • Understanding LLDP and LLDP-MED on EX Series Switches • Configuring LLDP • Understanding LLDP on page 1795

interval (Health Monitor)

Syntax	interval <i>seconds</i> ;
Hierarchy Level	[edit snmp health-monitor]
Release Information	Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the interval between sampling of the object being monitored by the health monitor.
Options	seconds —Time between samples, in seconds. Range: 1 through 2147483647 seconds Default: 300 seconds
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Health Monitoring on page 1509

lldp

```
Syntax  lldp {
    advertisement-interval seconds;
    disable;
    hold-multiplier number;
    interface (all | [interface-name]) {
        disable;
        power-negotiation {
            disable;
        }
    }
    lldp-configuration-notification-interval seconds;
    management-address ip-management-address;
    netbios-snooping;
    no-tagging;
    ptopo-configuration-maximum-hold-time seconds;
    ptopo-configuration-trap-interval seconds;
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>
        <no-stamp> <replace>;
        flag flag <disable>;
    }
    transmit-delay seconds;
}
```

Hierarchy Level [edit protocols]

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Statement introduced in Junos OS Release 11.1 for QFX Series.

Description Configure Link Layer Discovery Protocol (LLDP). The switch uses LLDP to advertise its identity and capabilities on a LAN, as well as to receive information about other network devices. LLDP is defined in the IEEE standard 802.1AB-2005.

The remaining statements are explained separately.



NOTE: The `transmit-delay` and `netbios-snooping` options are not available on QFabric systems.



NOTE: On EX4300 switches, LLDP cannot be configured on the `me0` or `vme` interface. Issuing the command `set protocols lldp interface me0` generates the following error message:

```
error: name: 'me0': Invalid interface
error: statement creation failed: interface
```

Issuing the command `set protocols lldp interface vme` generates the following error message:

```
error: name: 'vme': Invalid interface
error: statement creation failed: interface
```

.....

Default LLDP is enabled.

The following statements have default values:

- **advertisement-interval**—The default value is 30 seconds. The allowable range is from 5 through 32768 seconds.
- **hold-multiplier**—The default values is 4. The allowable range is from 2 through 10.
- **ptopo-configuration-maximum-hold-time**—The default value is 300 seconds. The allowable range is from 1 through 2147483647 seconds.
- **transmit-delay**—The default values is 2 seconds. The allowable range is from 1 through 8192 seconds.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [show lldp on page 2002](#)
- *Configuring LLDP (CLI Procedure)*
- *Configuring LLDP*
- [Understanding LLDP on page 1795](#)
- *Understanding LLDP and LLDP-MED on EX Series Switches*

lldp-configuration-notification-interval

Syntax	lldp-configuration-notification-interval <i>seconds</i> ;
Hierarchy Level	[edit protocols lldp]
Release Information	Statement introduced in Junos OS Release 9.6 for EX Series switches. Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify how often SNMP trap notifications are generated as a result of LLDP database changes. If the interval value is 0, trap notifications of database changes are disabled.
Default	SNMP trap notifications of LLDP database changes are disabled.
Options	seconds —Interval between trap notifications about LLDP database changes. Range: 0 through 3600
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• show lldp on page 2002

location

Syntax	location <i>location</i> ;
Hierarchy Level	[edit snmp]
Release Information	Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Define the value of the MIB II sysLocation object, which is the physical location of the managed system.
Options	location —Location of the local system. You must enclose the name within quotation marks (" ").
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the System Location for a Device Running Junos OS</i>

management-address

Syntax	<code>management-address <i>ip-management-address</i>;</code>
Hierarchy Level	[edit protocols lldp]
Release Information	Statement introduced in Junos OS Release 9.5 for EX Series switches. Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify the management address to be used in LLDP Management Address type, length, and value (TLV) messages. The Management Address TLV typically contains the IPv4 or IPv6 management addresses of the local system. Only out-of-band management addresses must be used for the management-address. Other remote managers can use this address to obtain information related to the local device.
Default	The LLDP Management Address TLV uses the IP address of the switch's management Ethernet interface (me0), or the IP address of the virtual management Ethernet (VME) interface if the switch is a Virtual Chassis member.
Options	<i>ip-management-address</i> —You can specify either an IPv4 or an IPv6 management address for the switch.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • show lldp on page 2002 • <i>Understanding LLDP and LLDP-MED on EX Series Switches</i> • <i>EX Series Switches Interfaces Overview</i> • Understanding LLDP on page 1795

name

Syntax	<code>name <i>name</i>;</code>
Hierarchy Level	[edit snmp]
Release Information	Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Set the system name from the command-line interface.
Options	<i>name</i> —System name override.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring a Different System Name</i>

nas-ip-address

Syntax	<code>nas-ip-address <i>ip-address</i>;</code>
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the NAS-IP address for outgoing RADIUS packets.
Options	<i>ip-address</i> —IP address of the network access server (NAS) that requests user authentication.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring RADIUS Server Authentication</i>• Configuring RADIUS Authentication (QFX Series or OCX Series) on page 1873


nonvolatile

Syntax	nonvolatile { commit-delay <i>seconds</i> ; }
Hierarchy Level	[edit snmp]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure options for SNMP Set requests. The statement is explained separately.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the Commit Delay Timer</i> • <i>commit-delay</i>

oid

Syntax	oid <i>object-identifier</i> (exclude include);
Hierarchy Level	[edit snmp view <i>view-name</i>]
Release Information	Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify an object identifier (OID) used to represent a subtree of MIB objects.
Options	<p>exclude—Exclude the subtree of MIB objects represented by the specified OID.</p> <p>include—Include the subtree of MIB objects represented by the specified OID.</p> <p>object-identifier—OID used to represent a subtree of MIB objects. All MIB objects represented by this statement have the specified OID as a prefix. You can specify the OID using either a sequence of dotted integers or a subtree name.</p>
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring MIB Views on page 1506

order

Syntax	<code>order (radius [<i>accounting-order-data-list</i>]);</code>
Hierarchy Level	[edit access profile <i>profile-name</i> accounting]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the order of authentication, authorization, and accounting (AAA) servers to use while sending accounting messages and updates.
Default	No order specified
Options	radius —RADIUS accounting for specified subscribers. [<i>accounting-order-data-list</i>]— Set of data listing the authentication order to be used, enclosed by brackets. This can be any combination of the authentication methods, up to and including a full list of the entire authentication order.
<hr/> <div> NOTE: The [edit access] hierarchy is not available on QFabric systems.</div> <hr/>	
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch</i>• <i>Configuring 802.1X RADIUS Accounting (CLI Procedure)</i>• <i>Configuring RADIUS Accounting</i>

port (RADIUS Server)

Syntax	<code>port <i>port-number</i>;</code>
Hierarchy Level	[edit system radius-server <i>address</i>], [edit system accounting destination radius server <i>address</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the port number on which to contact the RADIUS server.
Options	<i>number</i> —Port number on which to contact the RADIUS server. Default: 1812 (as specified in RFC 2865)



NOTE: The [edit system accounting] hierarchy is not available on QFabric systems.

Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring RADIUS Server Authentication</i>

profile

Syntax `profile profile-name {
 accounting (Access Profile) {
 accounting-stop-on-access-deny;
 accounting-stop-on-failure;
 order (radius | [accounting-order-data-list]);
 }
 authentication-order (Access Profile) [authentication-method];
 radius {
 accounting-server [server-addresses];
 authentication-server [server-addresses];
 }
 }`

Hierarchy Level [edit access]

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure an access profile. The access profile contains the entire authentication, authorization, and accounting (AAA) configuration that aids in handling AAA requests, including the authentication method and order, AAA server addresses, and AAA accounting.

Default Not enabled

Options *profile-name*—Profile name of up to 32 characters.

The remaining statements are explained separately.



NOTE: The [edit access] hierarchy is not available on QFabric systems.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- *Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch*
- *Configuring 802.1X RADIUS Accounting (CLI Procedure)*
- *Configuring RADIUS Accounting*

protocols

```
Syntax protocols {
    bgp {
        disable;
        accept-remote-nexthop;
        advertise-external <conditional>;
        advertise-inactive;
        (advertise-peer-as | no-advertise-peer-as);
        authentication-algorithm (aes-128-cmac-96 | hmac-sha-1-96 | md5);
        authentication-key key;
        authentication-key-chain key-chain;
        bfd-liveness-detection {
            authentication {
                algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
                    meticulous-keyed-sha-1 | simple-password);
                key-chain key-chain-name;
                loose-check;
            }
            detection-time {
                threshold milliseconds;
            }
            hold-down-interval milliseconds;
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            session-mode (automatic | multihop | single-hop);
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            version (1 | automatic);
        }
        cluster cluster-identifier;
        damping;
        description text-description;
        export [ policy-names ];
        family family-name {
            ... the family subhierarchies appear after the main [edit protocols bgp] hierarchy ...
        }
        graceful-restart {
            disable;
            restart-time seconds;
            stale-routes-time seconds;
        }
        group group-name {
            ... the group subhierarchy appears after the main [edit protocols bgp] hierarchy ...
        }
        hold-time seconds;
        import [ policy-names ];
        include-mp-next-hop;
        keep (all | none);
        local-address address;
```

```
local-as autonomous-system <loops number> < alias> <private>;
local-preference local-preference;
log-updown;
metric-out (metric | igp (delay-med-update | offset) | minimum-igp offset);
mtu-discovery;
multihop {
    no-nexthop-change;
    ttl ttl-value;
}
no-aggregator-id;
no-client-reflect;
out-delay seconds;
outbound-route-filter {
    bgp-orf-cisco-mode;
    prefix-based {
        accept {
            inet;
            inet6;
        }
    }
}
passive;
path-selection {
    always-compare-med;
    as-path-ignore;
    cisco-non-deterministic;
    external-router-id;
    med-plus-igp {
        igp-multiplier number;
        med-multiplier number;
    }
}
peer-as autonomous-system;
preference preference;
remove-private;
tcp-mss segment-size;
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
}
dcbx {
    disable;
    interface (interface-name | all) {
        disable;
        application-map application-map-name;
        applications {
            no-auto-negotiation;
        }
        enhanced-transmission-selection {
            no-auto-negotiation;
            no-recommendation-tlv;
            recommendation-tlv {
                no-auto-negotiation;
            }
        }
    }
}
```

```

    }
    dcbx-version (auto-negotiate | ieee-dcbx | dcbx-version-1.01);
    priority-flow-control {
        no-auto-negotiation;
    }
}
}
iccp {
    authentication-key string;
    local-ip-addr local-ip-addr;
    peer ip-address {
        authentication-key string;
        backup-liveness-detection {
            backup-peer-ip ip-address;
        }
        liveness-detection {
            detection-time {
                threshold milliseconds;
            }
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            version (Liveness Detection) (1 | automatic);
        }
        local-ip-addr ipv4-address;
        session-establishment-hold-time seconds;
    }
    session-establishment-hold-time seconds;
    traceoptions {
        file <filename> <files number> <match regular-expression> <microsecond-stamp>
            <size size> <world-readable | no-world-readable>;
        flag flag;
        no-remote-trace;
    }
}
igmp-snooping {
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable> <match
            regex>;
        flag flag (detail | disable | receive | send);
    }
}
vlan vlan-name {
    disable;
}
interface interface-name {
    group-limit limit;
    multicast-router-interface;
    static {
        group ip-address;
    }
}
}

```

```
    robust-count number;
  }
}
isis {
  disable;
  export [ policy-names ];
  ignore-attached-bit;
  interface interface-name {
    bfd-liveness-detection {
      authentication {
        algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
          meticulous-keyed-sha-1 | simple-password);
        key-chain key-chain-name;
        loose-check;
      }
      detection-time {
        threshold milliseconds;
      }
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
      }
      version (1 | automatic);
    }
  }
  checksum;
  csnp-interval (seconds | disable);
  disable;
  hello-padding (adaptive | loose | strict);
  level (1 | 2) {
    disable;
    hello-authentication-key key;
    hello-authentication-type authentication;
    hello-interval seconds;
    hold-time seconds;
    ipv4-multicast-metric number;
    metric metric;
    passive;
    priority number;
  }
  lsp-interval milliseconds;
  mesh-group (value | blocked);
  no-ipv4-multicast;
  no-unicast-topology;
  passive;
  point-to-point;
}
level (1 | 2) {
  disable;
  authentication-key key;
  authentication-type authentication;
  external-preference preference;
  no-csnp-authentication;
```

```

    no-hello-authentication;
    no-psnp-authentication;
    preference preference;
    prefix-export-limit number;
    wide-metrics-only;
}
loose-authentication-check;
lsp-lifetime seconds;
max-areas number;
no-adjacency-holddown;
no-authentication-check;
no-ipv4-routing;
overload {
    advertise-high-metrics;
    timeout seconds;
}
reference-bandwidth reference-bandwidth;
rib-group {
    inet group-name;
}
topologies {
    ipv4-multicast;
}
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
traffic-engineering {
    disable;
    family inet {
        shortcuts {
            multicast-rpf-routes;
        }
    }
}
}
lldp {
    disable;
    advertisement-interval seconds;
    hold-multiplier number;
    interface (LLDP) (all | interface-name) {
        disable;
    }
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable> <match
            regex>;
        flag flag (detail | disable | receive | send);
    }
}
mstp {
    disable;
    bpdu-timeout-action;
    bridge-priority priority;
    configuration-name name;
    forward-delay seconds;

```

```
hello-time seconds;
interface (all | interface-name) {
    disable;
    bpdu-timeout-action {
        block;
        alarm;
    }
    cost cost;
    edge;
    mode mode;
    no-root-port;
    priority priority;
}
max-age seconds;
max-hops hops;
msti msti-id {
    vlan (vlan-id | vlan-name);
    interface interface-name {
        disable;
        cost cost;
        edge;
        mode mode;
        priority priority;
    }
}
revision-level revision-level;
traceoptions {
    file filename <files number > <size size > <no-stamp | world-readable |
    no-world-readable>;
    flag flag;
}
}
ospf {
    disable;
    area area-id {
        area-range ip-prefix </prefix-length> <exact> <override-metric metric > <restrict>;
        context-identifier identifier
        interface interface-name {
            disable;
            authentication {
                md5 key-id key key-string <start-time YYYY-MM-DD.hh:mm>;
                simple-password key-string;
            }
            bandwidth-based-metrics {
                bandwidth value metric number;
            }
            bfd-liveness-detection {
                authentication {
                    algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
                    meticulous-keyed-sha-1 | simple-password);
                    key-chain key-chain-name;
                    loose-check;
                }
                detection-time {
                    threshold milliseconds;
                }
            }
        }
    }
}
```



```

    full-neighbors-only;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    version (1 | automatic);
}
dead-interval seconds;
dynamic-neighbors;
flood-reduction;
hello-interval seconds;
interface-type (nbma | p2mp | p2p);
metric metric;
neighbor address <eligible>;
no-eligible-backup;
no-interface-state-traps;
no-neighbor-down-notification;
passive {
    traffic-engineering {
        remote-node-id address;
    }
}
poll-interval seconds;
priority number;
retransmit-interval seconds;
secondary;
te-metric metric;
topology (name | default | ipv4-multicast) {
    disable;
    bandwidth-based-metrics {
        bandwidth value;
        metric number;
    }
    metric metric;
}
transit-delay seconds;
}
network-summary-export [ policy-names ];
network-summary-import [ policy-names ];
nssa {
    area-range ip-prefix</prefix-length> <exact> <override-metric metric> <restrict>;
    default-lsa {
        default-metric metric;
        metric-type type;
        type-7;
    }
    (summaries | no-summaries);
}
stub <default-metric metric> <summaries | no-summaries>;
virtual-link neighbor-id router-id transit-area area-id {
    disable;
    authentication {

```

```

        md5 key-id key key-string <start-time YYYY-MM-DD.hh:mm>;
        simple-password key-string;
    }
    dead-interval seconds;
    demand-circuit;
    flood-reduction;
    hello-interval seconds;
    ipsec-sa sa-name;
    no-neighbor-down-notification;
    retransmit-interval seconds;
    topology (name | default | ipv4-multicast) {
        disable;
        metric metric;
    }
    transit-delay seconds;
}
}
database-protection {
    ignore-count number;
    ignore-time seconds;
    maximum-lsa number;
    reset-time seconds;
    warning-only;
    warning-threshold percent;
}
export [ policy-names ];
external-preference preference;
graceful-restart {
    disable;
    helper-disable <both | restart-signaling | standard>;
    no-strict-lsa-checking;
    notify-duration seconds;
    restart-duration seconds;
}
import [ policy-names ];
no-nssa-abr;
no-rfc-1583;
overload <timeout seconds>;
preference preference;
prefix-export-limit number;
reference-bandwidth reference-bandwidth;
rib-group group-name;
topology (default | ipv4-multicast | name) {
    overload;
    prefix-export-limit number;
    topology-id number;
}
}
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
}
traffic-engineering {
    advertise-unnumbered-interfaces;
    credibility-protocol-preference;
    ignore-lsp-metrics;
}

```

```

    multicast-rpf-routes;
    no-topology;
    shortcuts <lsp-metric-into-summary>;
  }
}
pim {
  disable;
  assert-timeout seconds;
  dense-groups {
    addresses;
  }
  dr-election-on-p2p;
  export;
  family (inet | inet6) {
    disable;
  }
  graceful-restart {
    disable;
    restart-duration seconds;
  }
  import [ policy-names ];
  interface interface-name {
    accept-remote-source;
    disable;
    family (inet | inet6) {
      disable;
    }
    hello-interval seconds;
    mode (dense | sparse | sparse-dense);
    neighbor-policy [ policy-names ];
    override-interval milliseconds;
    priority number;
    propagation-delay milliseconds;
    reset-tracking-bit;
    version version;
  }
  join-load-balance;
  join-prune-timeout;
  nonstop-routing;
  override-interval milliseconds;
  propagation-delay milliseconds;
  reset-tracking-bit;
  rib-group group-name;
  rp {
    auto-rp {
      (announce | discovery | mapping);
      (mapping-agent-election | no-mapping-agent-election);
    }
    bootstrap {
      family (inet | inet6) {
        export [ policy-names ];
        import [ policy-names ];
        priority number;
      }
    }
  }
  bootstrap-import [ policy-names ];
}

```

```
bootstrap-export [ policy-names ];
bootstrap-priority number;
dr-register-policy [ policy-names ];
embedded-rp {
    group-ranges {
        destination-ip-prefix</prefix-length>;
    }
    maximum-rps limit;
}
local {
    family (inet | inet6) {
        address address;
        anycast-pim {
            disable;
            rp-set {
                address address <forward-msdp-sa>;
            }
            local-address address;
        }
        group-ranges {
            destination-ip-prefix</prefix-length>;
        }
        hold-time seconds;
        priority number;
    }
}
rp-register-policy [ policy-names ];
spt-threshold {
    infinity [ policy-names ];
}
static {
    address address {
        group-ranges {
            version version;
            destination-ip-prefix</prefix-length>;
        }
    }
}
}
rpf-selection {
    group group-address {
        source source-address {
            next-hop next-hop-address;
        }
        wildcard-source {
            next-hop next-hop-address;
        }
    }
    prefix-list prefix-list-addresses {
        source source-address {
            next-hop next-hop-address;
        }
        wildcard-source {
            next-hop next-hop-address;
        }
    }
}
```

```

traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
tunnel-devices [ mt-fpc/pic/port ];
}
rip {
    authentication-key password;
    authentication-type type;
    (check-zero | no-check-zero);
    group group-name {
        bfd-liveness-detection {
            authentication {
                algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
                    meticulous-keyed-sha-1 | simple-password);
                key-chain key-chain-name;
                loose-check;
            }
            detection-time {
                threshold milliseconds;
            }
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            version (1 | automatic);
        }
    }
    export [ policy-names ];
    import [ policy-names ];
    metric-out metric;
    neighbor neighbor-name {
        any-sender;
        authentication-key password;
        authentication-type type;
        bfd-liveness-detection {
            ... same statements as at the [edit protocols rip group group-name
                bfd-liveness-detection] hierarchy level ...
        }
        (check-zero | no-check-zero);
        import [ policy-names ];
        message-size number;
        metric-in metric;
        receive (both | none | version-1 | version-2);
        route-timeout seconds;
        send (broadcast | multicast | none | version-1);
        update-interval seconds;
    }
    preference preference;
    route-timeout seconds;
    update-interval seconds;
}
holddown seconds;

```

```
import [ policy-names ];
message-size number;
metric-in metric;
receive (both | none | version-1 | version-2);
rib-group group-name;
route-timeout seconds;
send (broadcast | multicast | none | version-1);
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
update-interval seconds;
}
rstp {
    disable;
    bpdu-block-on-edge;
    bridge-priority priority;
    forward-delay seconds;
    hello-time seconds;
    interface (all | interface-name) {
        disable;
        bpdu-timeout-action {
            alarm;
            block;
        }
        cost cost;
        edge;
        mode mode;
        no-root-port;
        priority priority;
    }
    max-age seconds;
}
traceoptions {
    file filename <files number> <size size> <no-stamp> <world-readable |
        no-world-readable>;
    flag flag;
}
}
stp {
    disable;
    bridge-priority priority;
    forward-delay seconds;
    hello-time seconds;
    interface (all | interface-name) {
        disable;
        bpdu-timeout-action {
            alarm;
            block;
        }
        cost cost;
        edge;
        mode mode;
        no-root-port;
        priority priority;
```

```

    }
    max-age seconds;
  }
  traceoptions {
    file filename <files number> <size size> <no-stamp | world-readable |
    no-world-readable>;
    flag flag;
  }
  uplink-failure-detection {
    group group-name {
      link-to-monitor interface-name;
      link-to-disable interface-name;
    }
  }
}
vstp {
  bpdu-block-on-edge;
  disable (Spanning Trees);
  force-version (Spanning Trees) stp;
  vlan (Spanning Trees) vlan-id {
    bridge-priority (Spanning Trees) priority;
    forward-delay (Spanning Trees) seconds;
    hello-time (Spanning Trees) seconds;
    interface (Spanning Trees) (all | interface-name) {
      bpdu-timeout-action (Spanning Trees) {
        block (Spanning Trees);
        log (Spanning Trees);
      }
      cost (Spanning Trees) cost;
      disable (Spanning Trees);
      edge (Spanning Trees);
      mode (Spanning Trees) mode;
      no-root-port (Spanning Trees);
      priority (Spanning Trees) priority;
    }
    max-age (Spanning Trees) seconds;
    traceoptions (Spanning Trees) {
      file filename <files number> <size size> <no-stamp | world-readable |
      no-world-readable>;
      flag flag;
    }
  }
}
}
}

```

Hierarchy Level [edit]

Release Information Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure protocols.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation • [Junos OS Routing Protocols Configuration Guide](#)

protocol-version

Syntax	<code>protocol-version version;</code>
Hierarchy Level	[edit system services ssh]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Description	Specify the secure shell (SSH) protocol version.
Default	v2 —SSH protocol version 2 is the default, introduced in Junos OS Release 11.4.
Options	version —SSH protocol version: v1 , v2 , or both.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	• Configuring the SSH Protocol Version on page 1413


ptopo-configuration-maximum-hold-time

Syntax	<code>ptopo-configuration-maximum-hold-time <i>seconds</i>;</code>
Hierarchy Level	[edit protocols lldp]
Release Information	Statement introduced in Junos OS Release 9.6 for EX Series switches. Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure how long to maintain the physical topology database entries. The physical topology identifies the devices on the network and their physical interconnections.
Options	<i>seconds</i> —Time to maintain physical topology database entries. Default: 300 Range: 1 through 2147483647
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • show lldp on page 2002 • <i>Understanding LLDP and LLDP-MED on EX Series Switches</i> • Understanding LLDP on page 1795


ptopo-configuration-trap-interval

Syntax	<code>ptopo-configuration-trap-interval <i>seconds</i>;</code>
Hierarchy Level	[edit protocols lldp]
Release Information	Statement introduced in Junos OS Release 9.6 for EX Series switches. Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify how often SNMP trap notifications are sent regarding changes in physical topology global statistics.
Default	SNMP trap notifications of changes in physical topology global statistics are disabled.
Options	<i>seconds</i> —Interval between SNMP trap notifications about physical topology global statistics. Range: 0 through 3600
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.


radius

Syntax	<pre>radius { accounting-server [server-addresses]; authentication-server [server-addresses]; }</pre>
Hierarchy Level	[edit access profile <i>profile-name</i>]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Configure the RADIUS servers for authentication and for accounting. To configure multiple RADIUS servers, include multiple radius statements. The servers are tried in order and in a round-robin fashion until a valid response is received from one of the servers or until all the configured retry limits are reached.</p> <p>The statements are explained separately.</p>
<div> NOTE: The [edit access] hierarchy is not available on QFabric systems.</div>	
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch</i>• <i>Configuring 802.1X RADIUS Accounting (CLI Procedure)</i>• <i>Filtering 802.1X Supplicants by Using RADIUS Server Attributes</i>• <i>Configuring RADIUS Accounting</i>

radius-options (edit system)

Syntax	<pre>radius-options { attributes { nas-ip-address <i>ip-address</i>; } enhanced-accounting; password-protocol <i>mschap-v2</i>; }</pre>
Hierarchy Level	[edit system]
Release Information	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>MS-CHAPv2 password protocol configuration option introduced in Junos OS Release 9.2.</p> <p>MS-CHAPv2 password protocol configuration option introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.</p>
	<p> NOTE: The <code>radius-options</code> statement is not available on QFabric systems.</p>
	<p>enhanced-accounting statement introduced in Junos OS Release 14.1.</p>
Description	Configure RADIUS options for the NAS-IP address for outgoing RADIUS packets and password protocol used in RADIUS packets.
Options	<p>enhanced-accounting—View the attribute values of a logged in user.</p> <p>nas-ip-address <i>ip-address</i>—IP address of the network access server (NAS) that requests user authentication.</p> <p>password-protocol <i>mschap-v2</i>—Protocol MS-CHAPv2, used for password authentication and password changing.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring MS-CHAPv2 for Password-Change Support • Configuring RADIUS System Accounting on page 1891 • <i>enhanced-accounting</i>

radius-server

Syntax	<pre>radius-server server-address { accounting-port <i>port-number</i>; port <i>number</i>; retry <i>number</i>; secret <i>password</i>; source-address <i>source-address</i>; timeout <i>seconds</i>; }</pre>
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Configure a RADIUS server for Point-to-Point Protocol (PPP).</p> <p>To configure multiple RADIUS servers, include multiple radius-server statements. The servers are tried in order and in a round-robin fashion until a valid response is received from one of the servers or until all the configured retry limits are reached.</p>
Options	<p>server-address—Address of the RADIUS authentication server.</p> <p>The remaining statements are explained separately.</p>
<div> NOTE: The accounting-port and source-address options are not available on QFabric systems.</div>	
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring RADIUS Authentication (QFX Series or OCX Series) on page 1873• accounting-port• port on page 1945• retry on page 235• secret on page 318• source-address on page 291• timeout on page 321

rate-limit

Syntax	<code>rate-limit <i>limit</i>;</code>
Hierarchy Level	<code>[edit system services finger],</code> <code>[edit system services ftp],</code> <code>[edit system services netconf ssh],</code> <code>[edit system services ssh],</code> <code>[edit system services telnet],</code> <code>[edit system services tftp-server],</code> <code>[edit system services xnm-clear-text],</code> <code>[edit system services xnm-ssl]</code>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	Configure the maximum number of connections attempts per minute, per protocol (either IPv6 or IPv4) on an access service. For example, a rate limit of 10 allows 10 IPv6 telnet session connection attempts per minute and 10 IPv4 telnet session connection attempts per minute.
Default	150 connections
Options	<p>rate-limit <i>limit</i>—(Optional) Maximum number of connection attempts allowed per minute, per IP protocol (either IPv4 or IPv6).</p> <p>Range: 1 through 250</p> <p>Default: 150</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <i>Configuring clear-text or SSL Service for Junos XML Protocol Client Applications</i>

remote-debug-permission

Syntax	remote-debug-permission (qfabric-admin qfabric-operator qfabric-user);
Hierarchy Level	[edit system login user <i>username</i> authentication] [edit system root-authentication]
Release Information	Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	(QFabric systems only) Configure authentication classes that permit or deny user access to individual components of the QFabric system.
Default	qfabric-user
Options	<p>qfabric-admin—Permits a user to log in to individual QFabric system components, view operations, and change component configurations.</p> <p>qfabric-operator—Permits a user to log in to individual QFabric system components and view component operations.</p> <p>qfabric-user—Prevents a user from logging in to individual QFabric system components.</p>
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring QFabric System Login Classes</i>• request component login on page 1996• <i>Understanding QFabric System Login Classes</i>

retry

Syntax	<code>retry <i>number</i>;</code>
Hierarchy Level	[edit system radius server <i>server-address</i>], [edit system accounting destination radius server <i>server-address</i>]
Release Information	Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Number of times the router or switch is allowed to try to contact a RADIUS authentication or accounting server.
Options	<i>number</i> —Number of retries allowed for contacting a RADIUS server. Range: 1 through 10 Default: 3



NOTE: The [edit system accounting] hierarchy is not available on QFabric systems.

Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring RADIUS Authentication (QFX Series or OCX Series) on page 1873 • Configuring RADIUS Accounting • timeout on page 321

rising-threshold (Health Monitor)

Syntax	<code>rising-threshold <i>percentage</i>;</code>
Hierarchy Level	[edit snmp health-monitor]
Release Information	Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Set the upper threshold for the monitored object when you configure a health monitor alarm. By setting a rising and a falling threshold for a monitored object, you can be alerted whenever the value of the variable falls outside the allowable operational range.
Options	<i>percentage</i> —Upper threshold for the alarm entry. Range: 1 through 100 Default: 80 percent of the maximum possible value
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Health Monitoring on page 1509• falling-threshold on page 1624

root-login

Syntax	root-login (allow deny deny-password);
Hierarchy Level	[edit system services ssh]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Control user access through SSH.
Default	Allow user access through SSH.
Options	allow —Allow users to log in to the router or switch as root through SSH. deny —Disable users from logging in to the router or switch as root through SSH. deny-password —Allow users to log in to the router or switch as root through SSH when the authentication method (for example, RSA authentication) does not require a password.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Root Login Through SSH on page 1413

services (Switches)

Syntax

```
services {  
  service-deployment {  
    servers address {  
      port-number port-number;  
    }  
    source-address address;  
  }  
  ssh {  
    connection-limit limit;  
    protocol-version [v1 v2];  
    rate-limit limit;  
    root-login (allow | deny | deny-password);  
  }  
}
```

Hierarchy Level [edit system]

Release Information Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure the switch so that users on remote systems can access the local switch through SSH.

The remaining statements are explained separately.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

snmp

```

Syntax  snmp {
    client-list client-list-name {
        ip-addresses;
    }
    community community-name {
        authorization authorization;
        client-list-name client-list-name;
        clients {
            address restrict;
        }
        logical-system logical-system-name {
            routing-instance routing-instance-name {
                clients {
                    addresses;
                }
            }
        }
        routing-instance routing-instance-name {
            clients {
                addresses;
            }
        }
        view view-name;
    }
    contact contact;
    description description;
    filter-duplicates;
    filter-interfaces;
    health-monitor {
        falling-threshold integer;
        interval seconds;
        rising-threshold integer;
    }
    interface [ interface-names ];
    location location;
    name name;
    nonvolatile {
        commit-delay seconds;
    }
    rmon {
        alarm index {
            description description;
            falling-event-index index;
            falling-threshold integer;
            falling-threshold-interval seconds;
            interval seconds;
            request-type;
            rising-event-index index;
            rising-threshold integer;
            sample-type (absolute-value | delta-value);
            startup-alarm (falling-alarm | rising-alarm | rising-or-falling alarm);
            syslog-subtag syslog-subtag;
        }
    }
}

```

```
    variable oid-variable;
}
event index {
    community community-name;
    description description;
    type type;
}
history history-index {
    bucket-size number;
    interface interface-name;
    interval seconds;
    owner owner-name;
}
}
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable> <match
        regular-expression>;
    flag flag;
}
trap-group group-name {
    categories {
        category;
    }
    destination-port port-number;
    routing-instance routing-instance-name;
    targets {
        address;
    }
    version (all | v1 | v2);
}
trap-options {
    agent-address outgoing-interface;
    source-address address;
}
v3 {
    notify name {
        tag tag-name;
        type trap;
    }
    notify-filter profile-name {
        oid object-identifier (include | exclude);
    }
    snmp-community community-index {
        community-name community-name;
        security-name security-name;
        tag tag-name;
    }
    target-address target-address-name {
        address address;
        address-mask address-mask;
        logical-system logical-system;
        port port-number;
        retry-count number;
        routing-instance routing-instance-name;
        tag-list tag-list;
        target-parameters target-parameters-name;
```

```

    timeout seconds;
  }
  target-parameters target-parameters-name {
    notify-filter profile-name;
    parameters {
      message-processing-model (v1 | v2c | V3);
      security-level (authentication | none | privacy);
      security-model (usm | v1 | v2c);
      security-name security-name;
    }
  }
  usm {
    local-engine {
      user username {
        authentication-sha {
          authentication-password authentication-password;
        }
        authentication-md5 {
          authentication-password authentication-password;
        }
        authentication-none;
        privacy-aes128 {
          privacy-password privacy-password;
        }
        privacy-des {
          privacy-password privacy-password;
        }
        privacy-3des {
          privacy-password privacy-password;
        }
        privacy-none;
      }
    }
    remote-engine engine-id {
      user username {
        authentication-sha {
          authentication-password authentication-password;
        }
        authentication-md5 {
          authentication-password authentication-password;
        }
        authentication-none;
        privacy-aes128 {
          privacy-password privacy-password;
        }
        privacy-des {
          privacy-password privacy-password;
        }
        privacy-3des {
          privacy-password privacy-password;
        }
        privacy-none {
          privacy-password privacy-password;
        }
      }
    }
  }
}

```

```

}
vacm {
  access {
    group group-name {
      (default-context-prefix | context-prefix context-prefix) {
        security-model (any | usm | v1 | v2c) {
          security-level (authentication | none | privacy) {
            notify-view view-name;
            read-view view-name;
            write-view view-name;
          }
        }
      }
    }
  }
}
security-to-group {
  security-model (usm | v1 | v2c) {
    security-name security-name {
      group group-name;
    }
  }
}
}
view view-name {
  oid object-identifier (include | exclude);
}
}
}

```

Hierarchy Level	[edit]
Release Information	Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure SNMP. The remaining statements are explained separately.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Understanding the Implementation of SNMP</i> • Configuring SNMP on page 1499

ssh

Syntax	<pre>ssh { authentication-order [<i>authentication-methods</i>]; ciphers [<i>cipher-1 cipher-2 cipher-3 ...</i>]; client-alive-count-max <i>seconds</i>; client-alive-interval <i>seconds</i>; connection-limit <i>limit</i>; hostkey-algorithm <<i>algorithm</i> no-<i>algorithm</i>>; key-exchange <<i>algorithm</i>>; macs <<i>algorithm</i>>; max-sessions-per-connection <<i>number</i>>; no-passwords; no-public-keys; no-tcp-forwarding; protocol-version [<i>v1 v2</i>]; rate-limit <i>limit</i>; root-login (<i>allow</i> <i>deny</i> <i>deny-password</i>); }</pre> <p>tcp-forwarding (JDM)</p>
Hierarchy Level	[edit system services]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>client-alive-interval and client-alive-max-count statements introduced in Junos OS Release 12.2.</p> <p>no-passwords statement introduced in Junos OS Release 13.3.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.</p> <p>no-public-keys statement introduced in Junos OS release 15.1.</p> <p>tcp-forwarding statement introduced in Junos OS Release 15.1X53-D50 for the NFX250 Network Services Platform.</p>
Description	<p>Allow SSH requests from remote systems to the local router or switch.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Configuring SSH Service for Remote Access to the Router or Switch on page 1412

system

```
Syntax  system {
    accounting {
        events [ login change-log interactive-commands ];
        destination {
            radius {
                server {
                    server-address {
                        accounting-port port-number;
                        retry number;
                        secret password;
                        source-address address;
                        timeout seconds;
                    }
                }
            }
        }
        tacplus {
            server {
                server-address {
                    port port-number;
                    secret password;
                    single-connection;
                    timeout seconds;
                }
            }
        }
    }
    archival {
        configuration {
            archive-sites {
                ftp://<username>:<password>@<host>:<port>/<url-path>;
                ftp://<username>:<password>@<host>:<port>/<url-path>;
            }
            transfer-interval interval;
            transfer-on-commit;
        }
    }
    arp {
        aging-timer minutes;
        interfaces;
    }
    authentication-order [ authentication-methods ];
    (compress-configuration-files | no-compress-configuration-files);
    default-address-selection;
    domain-name domain-name;
    domain-search [ domain-list ];
    host-name hostname;
    internet-options {
        icmpv4-rate-limit bucket-size bucket-size packet-rate packet-rate;
        source-port upper-limit <upper-limit>;
    }
    location {
```



```

altitude feet;
building name;
country-code code;
floor number;
hcoord horizontal-coordinate;
lata service-area;
latitude degrees;
longitude degrees;
npa-nxx number;
postal-code postal-code;
rack number;
vcoord vertical-coordinate;
}
login {
  announcement text;
  class class-name {
    access-end;
    access-start;
    allow-configuration "regular-expression";
    allowed-days "regular-expression";
    deny-commands "regular-expression";
    deny-configuration "regular-expression";
    idle-timeout minutes;
    login-tip;
    permissions [ permissions ];
  }
  message text;
  password {
    change-type (set-transitions | character-set);
    format (md5 | sha1 | des);
    maximum-length length;
    minimum-changes number;
    minimum-length length;
  }
  retry-options {
    backoff-factor seconds;
    backoff-threshold number;
    minimum-time seconds;
    tries-before-disconnect number;
  }
  user username {
    authentication {
      (encrypted-password "password" | plain-text-password);
      load-key-file URL;
      remote-debug-permission (qfabric-admin | qfabric-operator | qfabric-user);
      ssh-rsa "public-key";
      ssh-dsa "public-key";
    }
    uid uid-value;
    class class-name;
    full-name complete-name;
  }
}
name-server {
  address;
}

```

```
no-multicast-echo;
no-redirects;
no-ping-record-route;
no-ping-time-stamp;
ntp {
    authentication-key number type type value password;
    serveraddress <key key-number> <version value> <prefer>;
}
ports {
    auxiliary {
        disable;
        insecure;
        type terminal-type;
    }
    console {
        disable;
        insecure;
        log-out-on-disconnect;
        type terminal-type;
    }
}
radius-server server-address {
    accounting-port port-number;
    port number;
    retry number;
    secret password;
    source-address source-address;
    timeout seconds;
}
radius-options {
    password-protocol mschap-v2;
}
attributes {
    nas-ip-address ip-address;
}
root-authentication {
    (encrypted-password "password" | plain-text-password);
    ssh-rsa "public-key";
    ssh-dsa "public-key";
}
(saved-core-context | no-saved-core-context);
saved-core-files saved-core-files;
services {
    finger {
        connection-limit limit;
        rate-limit limit;
    }
    flow-tap-dtcp {
        ssh {
            connection-limit limit;
            rate-limit limit;
        }
    }
}
ftp {
    connection-limit limit;
    rate-limit limit;
```

```

}
service-deployment {
  servers server-address {
    port port-number;
  }
  source-address source-address;
}
ssh {
  root-login (allow | deny | deny-password);
  protocol-version [v1 v2];
  connection-limit limit;
  rate-limit limit;
}
telnet {
  connection-limit limit;
  rate-limit limit;
}
web-management {
  http {
    interfaces [ interface-names ];
    port port;
  }
  https {
    interfaces [ interface-names ];
    local-certificate name;
    port port;
  }
  session {
    idle-timeout [ minutes ];
    session-limit [ session-limit ];
  }
}
xnm-clear-text {
  connection-limit limit;
  rate-limit limit;
}
xnm-ssl {
  connection-limit limit;
  local-certificate name;
  rate-limit limit;
}
}
static-host-mapping {
  hostname {
    alias [ alias ];
    inet [ address ];
    sysid system-identifier;
  }
}
syslog {
  archive {
    files number;
    size maximum-file-size;
    start-time "YYYY-MM-DD.hh:mm";
    transfer-interval minutes;
    (world-readable | no-world-readable);
  }
}

```

```

}
console {
    facility severity;
}
file filename {
    archive {
        files number;
        size maximum-file-size;
        start-time "YYYY-MM-DD.hh:mm";
        transfer-interval minutes;
        (world-readable | no-world-readable);
    }
    explicit-priority;
    facility severity;
    match "regular-expression";
    structured-data {
        brief;
    }
}
host (hostname | other-routing-engine | scc-master) {
    explicit-priority;
    facility-override facility;
    facility severity;
    log-prefix string;
    match "regular-expression";
}
source-address source-address;
time-format (millisecond | year | year millisecond);
user (username | *) {
    facility severity;
    match "regular-expression";
}
}
tacplus-options {
    service-name service-name;
    (no-cmd-attribute-value | exclude-cmd-attribute);
}
tacplus-server server-address {
    port
    secret password;
    single-connection;
    source-address source-address;
    timeout seconds;
}
time-zone (GMThour-offset | time-zone);
}
tracing {
    destination-override {
        syslog host;
    }
}
use-imported-time-zones;
}

```

Hierarchy Level [edit]

Release Information Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure system management properties.



NOTE: The `radius-server source-address` and `radius-options` statements are not available on the QFabric system.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

tacplus-options

Syntax	<pre>tacplus-options { (exclude-cmd-attribute no-cmd-attribute-value); enhanced-accounting; service-name <i>service-name</i>; timestamp-and-timezone; }</pre>
Hierarchy Level	[edit system]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>no-cmd-attribute-value and exclude-cmd-attribute options introduced in Junos OS Release 9.3.</p> <p>Statement introduced in Junos OS Release 11.1 for QFX Series.</p> <p>timestamp-and-timezone option introduced in Junos OS Release 12.2.</p> <p>enhanced-accounting option introduced in Junos OS Release 14.1.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.</p>
Description	Configure TACACS+ options for authentication and accounting.
Options	<p>enhanced-accounting—View the attribute values of a logged in user.</p> <p>exclude-cmd-attribute—Exclude the cmd attribute value completely from start and stop accounting records to enable logging of accounting records in the correct log file on a TACACS+ server.</p> <p>no-cmd-attribute-value—Set the cmd attribute value to an empty string in the TACACS+ accounting start and stop requests to enable logging of accounting records in the correct log file on a TACACS+ server.</p> <p>service-name <i>service-name</i>—Name of the authentication service used when you configure multiple TACACS+ servers to use the same authentication service.</p> <p>Default: junos-exec</p> <p>timestamp-and-timezone—Include this statement if you want start time, stop time, and timezone attributes included in start/stop accounting records.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Configuring TACACS+ Authentication</i>• <i>Configuring TACACS+ System Accounting</i>• Junos OS Authentication Order for RADIUS, TACACS+, and Password Authentication on page 1859• <i>enhanced-accounting</i>

targets

Syntax	<pre>targets { address; }</pre>
Hierarchy Level	[edit snmp trap-group <i>group-name</i>]
Release Information	Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure one or more systems to receive SNMP traps.
Options	address —IPv4 or IPv6 address of the system to receive traps. You must specify an address, not a hostname.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring SNMP Trap Groups on page 1504

traceoptions (LLDP)

Syntax `traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable> <no-stamp>
 <replace>;
 flag flag <disable>;
 }`

Hierarchy Level [edit protocols [lldp](#)]

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Define tracing operations for the Link Layer Discovery Protocol (LLDP). You can trace messages under LLDP for LLDP and PTOPO MIBs.



NOTE: The traceoptions statement is not supported on the QFX3000 QFabric system.

Default Tracing operations are disabled.

Options **file *filename***—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`.

files *number*—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000

Default: 3 files

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:

- **all**—All tracing operations.
- **configuration**—Trace configuration operations.
- **interface**—Trace interface update events.
- **netbios**—Trace NetBIOS events.
- **packet**—Trace packet events.
- **rtsock**—Trace routing socket operations.
- **snmp**—Trace SNMP configuration operations.

- **vlan**—Trace VLAN update events.

no-stamp—(Optional) Do not timestamp the trace file.

Default: If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

no-world-readable—(Optional) Restrict file access to the user who created the file.

replace—(Optional) Replace an existing trace file if there is one rather than appending output to it.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **files** option.

Syntax: **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

Range: 10 KB through 1 GB

Default: 128 KB

Default: If you do not include this option, tracing output is appended to an existing trace file.

world-readable—(Optional) Enable unrestricted file access.



NOTE: The **traceoptions** statement is not supported on the QFX3000 QFabric system.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring LLDP-MED (CLI Procedure)*
- *Understanding LLDP and LLDP-MED on EX Series Switches*
- *Configuring LLDP*
- [Understanding LLDP on page 1795](#)

transfer-interval (Configuration)

Syntax	<code>transfer-interval <i>interval</i>;</code>
Hierarchy Level	[edit system archival configuration]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the router or switch to periodically transfer its currently active configuration to an archive site.
Options	interval —Interval at which to transfer the current configuration to an archive site. Range: 15 through 2880 minutes



NOTE: The [edit system archival] hierarchy is not available on QFabric systems.

Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Periodic Transfer of the Active Configuration to an Archive Site on page 1334• archive on page 1692• configuration on page 1363• transfer-on-commit on page 1365

transfer-on-commit

Syntax	transfer-on-commit;
Hierarchy Level	[edit system archival configuration]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Description	Configure the router or switch to transfer its currently active configuration to an archive site each time you commit a candidate configuration.



NOTE: When specifying a URL in a Junos OS statement using an IPv6 host address, you must enclose the entire URL in quotation marks (" ") and enclose the IPv6 host address in brackets ([]). For example, "ftp://username<:password>@[ipv6-host-address]<:port>/url-path" .



NOTE: The [edit system archival] hierarchy is not available on QFabric systems.

Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Transfer of the Currently Active Configuration When a Configuration Is Committed on page 1334 • archive on page 1692 • configuration on page 1363 • transfer-interval on page 1364

trap-group

Syntax	<pre>trap-group <i>group-name</i> { categories { <i>category</i>; } destination-port <i>port-number</i>; routing-instance <i>instance</i>; targets { <i>address</i>; } version (all v1 v2); }</pre>
Hierarchy Level	[edit snmp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Description	Create a named group of hosts to receive the specified trap notifications. The name of the trap group is embedded in SNMP trap notification packets as one variable binding (varbind) known as the community name. At least one trap group must be configured for SNMP traps to be sent.
Options	<p><i>group-name</i>—Name of the trap group. If the name includes spaces, enclose it in quotation marks (" ").</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring SNMP Trap Groups</i>

trap-options

Syntax	<pre>trap-options { agent-address outgoing-interface; source-address address; }</pre>
Hierarchy Level	[edit snmp]
Release Information	Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Using SNMP trap options, you can set the source address of every SNMP trap packet sent by the router or switch to a single address, regardless of the outgoing interface. In addition, you can set the agent address of each SNMPv1 trap. For more information about the contents of SNMPv1 traps, see RFC 1157.</p> <p>The remaining statements are explained separately.</p>
Default	Disabled
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring SNMP Trap Options</i>

user (Access)

Syntax `user username {
 authentication {
 (encrypted-password "password" | plain-text-password);
 load-key-file URL;
 remote-debug-permission (qfabric-admin | qfabric-operator | qfabric-user);
 ssh-dsa "public-key" <from hostname>;
 ssh-rsa "public-key" <from hostname>;
 }
 class class-name;
 full-name "complete-name";
 uid uid-value;
 }`

Hierarchy Level [edit system login]

Release Information Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
 Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure access permission for individual users.

Options The remaining statements are explained separately.

Required Privilege Level admin—To view this statement in the configuration.
 admin-control—To add this statement to the configuration.

Related Documentation

- [Configuring Junos OS User Accounts on page 125](#)
- [class on page 275](#)

version

Syntax	version (all v1 v2);
Hierarchy Level	[edit snmp trap-group <i>group-name</i>]
Release Information	Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify the version number of SNMP traps.
Default	all—Send an SNMPv1 and SNMPv2 trap for every trap condition.
Options	all—Send an SNMPv1 and SNMPv2 trap for every trap condition. v1—Send SNMPv1 traps only. v2—Send SNMPv2 traps only.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring SNMP Trap Groups on page 1504

CHAPTER 70

Operational Commands

- `clear lldp neighbors`
- `clear lldp statistics`
- `request component login`
- `show ethernet-switching interfaces`
- `show lldp`
- `show lldp local-information`
- `show lldp neighbors`
- `show lldp statistics`
- `show route instance`
- `show snmp statistics`
- `ssh`

clear lldp neighbors

Syntax	clear lldp neighbors <interface <i>interface</i> >
Release Information	Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Clear the learned remote neighbor information on all or selected interfaces.
Options	none —Clear the remote neighbor information on all interfaces. interface <i>interface</i> —(Optional) Clear the remote neighbor information from the selected interface.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• <i>Configuring LLDP</i>• Understanding LLDP on page 1795
List of Sample Output	clear lldp neighbors on page 1994 clear lldp neighbors interface on page 1994

Sample Output

clear lldp neighbors

```
user@switch> clear lldp neighbors
```

clear lldp neighbors interface

```
user@switch> clear lldp neighbors interface ge-0/1/1.0
```

clear lldp statistics

Syntax	<code>clear lldp statistics</code> <code><interface <i>interface</i>></code>
Release Information	Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Clear LLDP statistics on one or more interfaces.
Options	none —Clears LLDP statistics on all interfaces. interface <i>interface-names</i> —(Optional) Clear LLDP statistics on an interface.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Configuring LLDP • Understanding LLDP on page 1795
List of Sample Output	clear lldp statistics on page 1995 clear lldp statistics interface on page 1995

Sample Output

clear lldp statistics

```
user@switch> clear lldp statistics
```

clear lldp statistics interface

```
user@switch> clear lldp statistics interface ge-0/1/1.0
```

request component login

Syntax	<code>request component login <i>component-name</i></code>
Release Information	Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	(QFabric systems only) Log in to a QFabric system component. To gain access to individual components by way of the request component login command, you must first provide the qfabric-admin or qfabric-operator class privilege to your user (for more information, see: remote-debug-permission).
Options	<i>component-name</i> —Specify the QFabric system component to which you wish to log in.
Required Privilege Level	admin
Related Documentation	<ul style="list-style-type: none">• remote-debug-permission on page 1966
List of Sample Output	request component login (with qfabric-admin Privileges) on page 1996 request component login (with qfabric-operator Privileges) on page 1997 request component login (with qfabric-user Privileges) on page 1997

Sample Output

The three sample output displays show the results of attempts to log in to Node device EE3093. The results differ depending on the privilege level assigned to the user.

request component login (with qfabric-admin Privileges)

```
admin@qfabric> request component login EE3093
Warning: Permanently added 'qfabric-node-ee3093,169.254.128.41' (RSA) to the list
of known hosts.
--- JUNOS 11.3I built 2011-11-04 12:46:16 UTC
{master}
qfabric-admin@node-ee3093> ?
Possible completions:
clear          Clear information in the system
file           Perform file operations
help           Provide help information
load           Load information from file
monitor        Show real-time debugging information
mtrace         Trace multicast path from source to receiver
op             Invoke an operation script
ping           Ping remote target
quit           Exit the management session
request        Make system-level requests
restart        Restart software process
save           Save information to file
set            Set CLI properties, date/time, craft interface message
show           Show system information
ssh            Start secure shell on another host
start          Start shell
telnet         Telnet to another host
test           Perform diagnostic debugging
```

```
tracertoute          Trace route to remote host{master}  
qfabric-admin@node-ee3093>
```

request component login (with qfabric-operator Privileges)

```
operator@qfabric> request component login EE3093  
Warning: Permanently added 'qfabric-node-ee3093,169.254.128.41' (RSA) to the list  
of known hosts.  
--- JUNOS 11.3I built 2011-11-04 12:46:16 UTC  
{master}  
qfabric-operator@node-ee3093> ?  
Possible completions:  
file          Perform file operations  
help          Provide help information  
load          Load information from file  
op            Invoke an operation script  
quit          Exit the management session  
request       Make system-level requests  
save          Save information to file  
set           Set CLI properties, date/time, craft interface message  
show          Show system information  
start         Start shell  
test          Perform diagnostic debugging  
{master}  
qfabric-operator@node-ee3093>
```

request component login (with qfabric-user Privileges)

```
user0@qfabric> request component login EE3093  
error: User user0 does not have sufficient permissions to login to device ee3093
```

show ethernet-switching interfaces

Syntax	show ethernet-switching interfaces <brief detail summary> <interface <i>interface-name</i> >
Release Information	Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display information about switched Ethernet interfaces.
Options	<p>none—(Optional) Display brief information for Ethernet-switching interfaces.</p> <p>brief detail summary—(Optional) Display the specified level of output.</p> <p>interface <i>interface-name</i>—(Optional) Display Ethernet-switching information for a specific interface.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Troubleshooting Ethernet Switching on page 2134Understanding Bridging and VLANs on page 2089 • Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 2096 • Example: Setting Up Bridging with Multiple VLANs • Understanding FCoE on page 6436 • Interfaces Overview on page 2785
List of Sample Output	show ethernet-switching interfaces on page 1999 show ethernet-switching interfaces summary on page 2000 show ethernet-switching interfaces brief on page 2000 show ethernet-switching interfaces detail on page 2000 show ethernet-switching interfaces interface-name on page 2001
Output Fields	Table 158 lists the output fields for the show ethernet-switching interfaces command. Output fields are listed in the approximate order in which they appear.

Table 158: show ethernet-switching interfaces Output Fields

Field Name	Field Description	Level of Output
Interface	Name of a switching interface.	All levels
State	Interface state. Values are up or down .	none, brief , detail , summary
VLAN members	Name of a VLAN.	none, brief , detail , summary

Table 158: show ethernet-switching interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Blocking	Forwarding state of the interface: <ul style="list-style-type: none"> • blocked—Traffic is not being forwarded on the interface. • unblocked—Traffic is forwarded on the interface. • MAC limit exceeded—The interface is temporarily disabled because of a MAC limiting error. The disabled interface is automatically restored to service when the disable timeout expires. • MAC move limit exceeded—The interface is temporarily disabled because of a MAC move limiting error. The disabled interface is automatically restored to service when the disable timeout expires. • Storm control in effect —The interface is temporarily disabled because of a storm control error. The disabled interface is automatically restored to service when the disable timeout expires. • Storm control shutdown in effect —The interface is temporarily disabled because of a storm control shutdown error. The disabled interface is automatically restored to service when the disable timeout expires. 	none, brief , detail , summary
Index	VLAN index internal to Junos OS software.	detail
untagged tagged	Specifies whether the interface forwards IEEE802.1Q-tagged or untagged traffic.	detail

Sample Output

show ethernet-switching interfaces

```
user@switch> show ethernet-switching interfaces
```

Interface	State	VLAN members	Blocking
xe-0/0/0.0	up	T1122	unblocked
xe-0/0/1.0	down	default	– MAC limit exceeded
xe-0/0/2.0	down	default	– MAC move limit exceeded
xe-0/0/3.0	down	default	– Storm control in effect
xe-0/0/4.0	down	default	unblocked
xe-0/0/5.0	down	default	unblocked
xe-0/0/6.0	down	default	unblocked
xe-0/0/7.0	down	default	unblocked
xe-0/0/8.0	down	default	unblocked
xe-0/0/9.0	up	T111	unblocked
xe-0/0/10.0	down	default	unblocked
xe-0/0/11.0	down	default	unblocked
xe-0/0/12.0	down	default	unblocked
xe-0/0/13.0	down	default	unblocked
xe-0/0/14.0	down	default	unblocked
xe-0/0/15.0	down	default	unblocked
xe-0/0/16.0	down	default	unblocked
xe-0/0/17.0	down	default	unblocked
xe-0/0/18.0	down	default	unblocked
xe-0/0/19.0	up	T111	unblocked
xe-0/1/0.0	down	default	unblocked
xe-0/1/1.0	down	default	unblocked
xe-0/1/2.0	down	default	unblocked
xe-0/1/3.0	down	default	unblocked

show ethernet-switching interfaces summary

```
user@switch> show ethernet-switching interfaces summary
xe-0/0/0.0
xe-0/0/1.0
xe-0/0/2.0
xe-0/0/3.0
xe-0/0/8.0
xe-0/0/10.0
xe-0/0/11.0
```

show ethernet-switching interfaces brief

```
user@switch> show ethernet-switching interfaces brief
Interface  State  VLAN members  Blocking
xe-0/0/0.0  down   default        unblocked
xe-0/0/1.0  down   employee-vlan  unblocked
xe-0/0/2.0  down   employee-vlan  unblocked
xe-0/0/3.0  down   employee-vlan  unblocked
xe-0/0/8.0  down   employee-vlan  unblocked
xe-0/0/10.0 down   default        unblocked
xe-0/0/11.0 down   employee-vlan  unblocked
```

show ethernet-switching interfaces detail

```
user@switch> show ethernet-switching interfaces detail
Interface: xe-0/0/0.0 Index: 65
  State: down
  VLANs:
    default                untagged    unblocked

Interface: xe-0/0/1.0 Index: 66
  State: down
  VLANs:
    employee-vlan          untagged    unblocked

Interface: xe-0/0/2.0 Index: 67
  State: down
  VLANs:
    employee-vlan          untagged    unblocked

Interface: xe-0/0/3.0 Index: 68
  State: down
  VLANs:
    employee-vlan          untagged    unblocked

Interface: xe-0/0/8.0 Index: 69
  State: down
  VLANs:
    employee-vlan          untagged    unblocked

Interface: xe-0/0/10.0 Index: 70
  State: down
  VLANs:
    default                untagged    unblocked

Interface: xe-0/0/11.0 Index: 71
  State: down
  VLANs:
    employee-vlan          tagged      unblocked
```


show ethernet-switching interfaces interface-name

```
user@switch> show ethernet-switching interfaces xe-0/0/0.0
  Interface  State   VLAN members  Blocking
xe-0/0/0.0  down    default        unblocked
```

show lldp

Syntax `show lldp`
`<detail>`

Release Information Command introduced in Junos OS Release 9.0 for EX Series switches.
 Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
 Command introduced in Junos OS Release 11.1 for the QFX Series.

Description Display information about Link Layer Discovery Protocol (LLDP) and Link Level Discovery Protocol—Media Endpoint Discovery (LLDP-MED) configuration and capabilities on the switch. LLDP and LLDP-MED are used to learn about and to distribute device information on network links.



NOTE: LLDP-MED is not available on the QFX Series.

Options **none**—Display LLDP information for all interfaces.
detail—(Optional) Display detailed LLDP information for all interfaces.

Required Privilege Level view

Related Documentation

- *Configuring LLDP (CLI Procedure)*
- *Configuring LLDP-MED (CLI Procedure)*
- *Understanding LLDP and LLDP-MED on EX Series Switches*
- *Configuring LLDP*
- [Understanding LLDP on page 1795](#)

List of Sample Output [show lldp \(EX3200 switches\) on page 2005](#)
[show lldp \(EX4300 switches\) on page 2005](#)
[show lldp detail \(EX4300 switches\) on page 2006](#)

Output Fields [Table 159](#) lists the output fields for the **show lldp** command. Output fields are listed in the approximate order in which they appear.

Table 159: show lldp Output Fields

Field Name	Field Description	Level of Output
LLDP	LLDP operating state. The state can be enabled or disabled . NOTE: If a VLAN that has been configured for untagged packets on an interface also has Layer 2 protocol tunneling (L2PT) enabled for LLDP, the LLDP operating state for that interface is displayed as disabled .	All levels

Table 159: show lldp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Advertisement interval	Frequency, in seconds, at which LLDP advertisements are sent. This value is set by the <code>advertisement-interval</code> configuration statement.	All levels
Transmit delay	Seconds of delay before advertisements are sent to neighbors following a change to a TLV (type, length, or value) element in the LLDP protocol or to the state of the local system, such as a change in hostname or management address. You can set this value to reduce the delay in notifying neighbors of a change in the local system. This value is set by the <code>transmit-delay</code> configuration statement.	All levels
Hold timer	On EX4300 switches, the hold timer shows the length of time LLDP information is held before it is discarded. The hold timer value is equal to the advertisement interval multiplied by the hold multiplier. On all other switches, the hold timer shows the value of the hold multiplier. The hold multiplier value is set by the <code>hold-multiplier</code> configuration statement.	All levels
Notification interval	How often LLDP trap notifications are generated as a result of LLDP database changes. If the interval value is 0, LLDP trap notifications on database changes are disabled. This value is set by the <code>lldp-configuration-notification-interval</code> configuration statement.	All levels
Config Trap Interval	How often LLDP trap notifications are generated as a result of changes in topology—for example, when an endpoint connects or disconnects. If the interval value is 0, LLDP trap notifications on topology changes are disabled. This value is set by the <code>ptopo-configuration-trap-interval</code> configuration statement.	All levels
Connection Hold timer	Amount of time the system maintains dynamic topology entries. This value is set by the <code>ptopo-configuration-maximum-hold-time</code> configuration statement.	All levels
LLDP-MED	LLDP-MED operating state. The state can be Enabled or Disabled .	All levels
MED fast start count	Number of advertisements sent from a switch to a device, such as a VoIP telephone, when the device is first detected by the switch. These increased advertisements are temporary. After a device and a switch exchange information and can communicate, advertisements are reduced to one per second. This value is set by using the <code>fast-start</code> configuration statement. NOTE: <code>fast-start</code> is not available on the QFX Series.	All levels
Interface	Name of the interface for which LLDP configuration information is being reported.	All levels
Parent Interface	Name of the aggregated Ethernet interface, if any, to which the interface belongs.	All levels

Table 159: show lldp Output Fields (*continued*)

Field Name	Field Description	Level of Output
LLDP	LLDP operating state. The state can be Enabled or Disabled .	All levels
Power Negotiation	LLDP power negotiation operating state. The state can be Enabled or Disabled .	All levels
Neighbor count	Total number of new LLDP neighbors detected since the last switch reboot.	detail
Interface	Name of the interface that is advertising VLAN information.	All levels
Vlan-id	VLAN tag associated with the interface sending LLDP frames. If the interface is not a member of a VLAN, the VLAN ID is advertised as 0.	detail
Vlan-name	VLAN name associated with the VLAN ID.	detail
LLDP basic TLVs supported	<p>Basic TLVs supported on the switch:</p> <ul style="list-style-type: none"> • Chassis identifier—TLV that advertises the MAC address associated with the local system. • Port identifier—TLV that advertises the port identification for the specified port in the local system. • Port description—Interface name for the port. • System name—TLV that advertises the user-configured name of the local system. • System description—TLV that advertises the system description containing information about the software and current image running on the system. This information is taken from the software and is not configurable. • System capabilities—TLV that advertises the primary functions performed by the system—for example, bridge or router. • Management address—TLV that advertises the IP management address of the local system. 	detail
Supported LLDP 802 TLVs	<p>802.3 TLVs supported on the switch:</p> <ul style="list-style-type: none"> • MAC/PHY configuration status—TLV that advertises information about the physical interface, such as autonegotiation status and support and MAU type. The information is based on the physical interface structure and is not configurable. • Power via MDI—TLV that advertises MDI power support, PSE power pair, and power class information. • Link aggregation—TLV that advertises if the interface is aggregated and its aggregated interface ID. • Maximum frame size—TLV that advertises the maximum transmission unit (MTU) of the interface sending LLDP frames. • Port VLAN tag—TLV that advertises the VLAN tag configured on the interface. • Port VLAN name—TLV that advertises the VLAN name configured on the interface. 	detail

Table 159: show lldp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Supported LLDP MED TLVs	<p>LLDP-MED TLVs supported on the switch:</p> <ul style="list-style-type: none"> • LLDP MED capabilities—TLV that advertises the primary function of the port. The capabilities values range from 0 through 15: <ul style="list-style-type: none"> • 0—Capabilities • 1—Network Policy • 2—Location Identification • 3—Extended Power via MDI-PSE • 4—Inventory • 5–15—Reserved • Network policy—TLV that advertises the port VLAN configuration and associated Layer 2 and Layer 3 attributes. Attributes include the policy identifier, application types—such as voice or streaming video—802.1Q VLAN tagging, and 802.1p priority bits and DiffServ code points. • Endpoint location—TLV that advertises the physical location of the endpoint. • Extended power Via MDI—TLV that advertises the power type, power source, power priority, and power value of the port. It is the responsibility of the PSE device (network connectivity device) to advertise the power priority on a port. 	detail

Sample Output

show lldp (EX3200 switches)

```

user@switch> show lldp
LLDP                               : Enabled
Advertisement interval             : 30 seconds
Transmit delay                     : 2 seconds
Hold timer                        : 4 seconds
Notification interval             : 0 Second(s)
Config Trap Interval              : 0 seconds
Connection Hold timer             : 300 seconds

LLDP MED                           : Disabled
MED fast start count              : 3 Packets

```

Interface	Parent Interface	LLDP	LLDP-MED	Power Negotiation
all	-	Enabled	Enabled	Enabled

show lldp (EX4300 switches)

```

user@switch> show lldp
LLDP                               : Enabled
Advertisement interval             : 30 seconds
Transmit delay                     : 2 seconds
Hold timer                        : 120 seconds
Notification interval             : 0 Second(s)
Config Trap Interval              : 0 seconds
Connection Hold timer             : 300 seconds

LLDP MED                           : Disabled
MED fast start count              : 3 Packets

```

Interface	Parent Interface	LLDP	LLDP-MED	Power Negotiation
all	-	Enabled	Enabled	Enabled

show lldp detail (EX4300 switches)

```

user@switch> show lldp detail
LLDP : Enabled
Advertisement interval : 30 seconds
Transmit delay : 2 seconds
Hold timer : 120 seconds
Notification interval : 0 Second(s)
Config Trap Interval : 0 seconds
Connection Hold timer : 300 seconds

LLDP MED : Disabled
MED fast start count : 3 Packets

```

Interface	Parent Interface	LLDP	LLDP-MED	Power Negotiation
Neighbor count				
all	-	Enabled	Enabled	Enabled
8				

Interface	Parent Interface	Vlan-id	Vlan-name
xe-3/0/0.0	ae31.0	100	v100
xe-3/0/0.0	ae31.0	101	v101
xe-3/0/0.0	ae31.0	4000	v4000
xe-3/0/1.0	ae31.0	100	v100
xe-3/0/1.0	ae31.0	101	v101
xe-3/0/1.0	ae31.0	4000	v4000
xe-3/0/2.0	ae31.0	100	v100
xe-3/0/2.0	ae31.0	101	v101
xe-3/0/2.0	ae31.0	4000	v4000

LLDP basic TLVs supported:

Chassis identifier, Port identifier, Port description, System name, System description, System capabilities, Management address.

Supported LLDP 802 TLVs:

MAC/PHY configuration/status, Power via MDI, Link aggregation, Maximum frame size, Port VLAN tag, Port VLAN name.

Supported LLDP MED TLVs:

LLDP MED capabilities, Network policy, Endpoint location, Extended power Via MDI.

show lldp local-information

Syntax	show lldp local-information
Release Information	Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display the information that the switch provides in Link Layer Discovery Protocol (LLDP) advertisements to its neighbors.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring LLDP (CLI Procedure)</i> • <i>Understanding LLDP and LLDP-MED on EX Series Switches</i> • management-address on page 1941 • <i>Configuring LLDP</i> • Understanding LLDP on page 1795
List of Sample Output	show lldp local-information (EX Series Switch) on page 2008
Output Fields	Table 160 lists the output fields for the show lldp local-information command. Output fields are listed in the approximate order in which they appear.

Table 160: show lldp local-information Output Fields

Field Name	Field Description
LLDP Local Information details	Information about the local system (the switch): <ul style="list-style-type: none"> • Chassis ID—MAC address associated with the switch. • System name—User-configured name of the switch. • System descr—System description containing information about the switch model and the current software image running on the switch. This information is taken from the software and is not configurable.
System Capabilities	Capabilities (such as bridge or router) that are supported or enabled on the system.
Management Information	Details of the management information: Port Name , Port Address (such as 10.204.34.35), Address Type (such as ipv4 or ipv6), Port ID (SNMP interface index), Port ID Subtype , and Port Subtype . <p>The Port Subtype displays:</p> <ul style="list-style-type: none"> • ifindex(2)—IP address of the switch's management Ethernet interface (me0) or virtual management Ethernet (VME) interface address (for a virtual chassis) is used to manage the switch. • unknown(1)—IP management address has been configured with set protocols lldp management-address.

Table 160: show lldp local-information Output Fields (*continued*)

Field Name	Field Description
Interface name	Name of the local interface which is configured for either LLDP or LLDP-MED.
Parent Interface	Name of the aggregated Ethernet interface, if any, to which the local interface belongs.
SNMP Index	SNMP interface index.
Interface description	User-configured port description.
Status	Administrative status of the interface: either up or down .
Tunneling	Status of tunneling on the interface: either enabled or disabled .

Sample Output

show lldp local-information (EX Series Switch)

```
user@switch> show lldp local-information
```

LLDP Local Information details

```
Chassis ID   : 00:1d:b5:aa:b9:f0
System name  : switch
System descr : Juniper Networks, Inc. ex8208 , version 10.4I0 [builder] Build
               date: 2010-11-17 12:38:30 UTC
```

System Capabilities

```
Supported   : Bridge Router
Enabled     : Bridge Router
```

Management Information

```
Port Name    : -
Port Address : 10.93.54.6
Address Type  : IPv4
Port ID       : 34
Port ID Subtype : local(7)
Port Subtype  : ifIndex(2)
```

Interface name	Parent Interface	SNMP Index	Interface description	Status	Tunneling
me0.0	-	34	-	Down	Disabled
xe-3/0/0.0	ae31.0	769	xe-3/0/0.0	Up	Disabled
xe-3/0/1.0	ae31.0	770	xe-3/0/1.0	Up	Disabled
xe-3/0/2.0	ae31.0	771	xe-3/0/2.0	Up	Disabled
xe-3/0/3.0	ae31.0	772	xe-3/0/3.0	Up	Disabled
xe-3/0/4.0	ae31.0	577	xe-3/0/4.0	Up	Disabled
xe-3/0/5.0	ae31.0	578	xe-3/0/5.0	Up	Disabled
xe-3/0/6.0	ae31.0	579	xe-3/0/6.0	Up	Disabled
xe-3/0/7.0	ae31.0	581	xe-3/0/7.0	Up	Disabled

show lldp neighbors

Syntax <show lldp *neighbors*>
<interface *interface-ids*>

Release Information Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Command introduced in Junos OS Release 11.1 for the QFX Series.

Description Display learned information about Link Layer Discovery Protocol (LLDP) on all neighboring interfaces or on selected interfaces.

Options **none**—Display learned LLDP information on all neighboring interfaces and devices.

interface *interface-ids*—(Optional) Display learned LLDP information on the selected interfaces or devices.



NOTE: When a port with DCBX enabled begins to exchange type, length, and value (TLV) entries, optional LLDP TLVs on that port are not advertised to neighbors in order to interoperate with a wider variety of converged network adapters (CNAs). As a result, information for those ports will not be listed in the output for this command.

Required Privilege Level view

Related Documentation

- [Configuring LLDP](#)
- [Understanding LLDP on page 1795](#)

List of Sample Output [show lldp neighbors on page 2012](#)
[show lldp neighbors interface on page 2012](#)

Output Fields [Table 161](#) lists the output fields for the **show lldp neighbors** command. Output fields are listed in the approximate order in which they appear.

Table 161: show lldp neighbors Output Fields

Field Name	Field Description
Local Interface	List of local interfaces for which neighbor information is available.
Parent Interface	List of aggregated Ethernet interfaces, if any, to which the local interfaces belong.
Chassis ID	List of chassis identifiers for neighbors.
Port info	List of port information gathered from neighbors. This could be the port identifier or port description.

Table 161: show lldp neighbors Output Fields (*continued*)

Field Name	Field Description
System name	List of system names gathered from neighbors.
LLDP Neighbor Information	Information about both the local system (the switch) and a neighbor system on the interface (appears when the interface option is used).
Local Information	Information about the local system (appears when the interface option is used).
Index	Local interface index (appears when the interface option is used).
Time to live	Number of seconds for which this information is valid (appears when the interface option is used).
Time mark	Date and timestamp of information (appears when the interface option is used).
Local Interface	Name of the local physical interface (appears when the interface option is used).
Parent Interface	Name of the aggregated Ethernet interface, if any, to which the interface belongs (appears when the interface option is used).
Local Port ID	Local interface SNMP index (appears when the interface option is used).
Ageout Count	Number of times the complete set of information advertised by the neighbor has been deleted from LLDP neighbor information maintained by the local system because the information timeliness interval has expired (appears when the interface option is used).
Neighbor Information	Information about a neighbor system on the interface (appears when the interface option is used).
Chassis type	Type of chassis identifier supplied, such as MAC address (appears when the interface option is used).
Chassis ID	Chassis identifier of the chassis type listed (appears when the interface option is used).
Port type	Type of port identifier supplied, such as locally assigned (appears when the interface option is used).
Port ID	Port identifier of the port type listed (appears when the interface option is used).
Port description	Port description (appears when the interface option is used).
System name	Name supplied by the system on the interface (appears when the interface option is used).

Table 161: show lldp neighbors Output Fields (*continued*)

Field Name	Field Description
System Description	Description supplied by the system on the interface (appears when the interface option is used).
System capabilities	Capabilities (such as Bridge , Router , and Telephone) that are supported or enabled by the system on the interface (appears when the interface option is used).
Management Info	<p>Details of management information: Type (such as ipv4 or ipv6), Address (such as 10.204.34.35), Port ID, Subtype, Interface Subtype, and organization identifier (OID) (appears when the interface option is used).</p> <p>The Interface Subtype displays:</p> <ul style="list-style-type: none"> • ifIndex(2)—IP address of the neighbor's management Ethernet interface (me0) or virtual management Ethernet (VME) interface address (for a virtual chassis) is used to manage the switch. • unknown(1)—Neighbor's IP management address has been configured with set protocols lldp management-address.
Media Info	Additional details about the endpoint device appear when a device that supports LLDP-MED is attached to the interface. The specific details depend upon the capabilities of the device. Details might include Media endpoint class (such as Class 3 for communication devices such as IP phones), MED Hardware revision , MED Firmware revision , MED Software revision , MED Serial number , MED Manufacturer name , or MED Model name .
Organization Info	One or more entries listing remote information by organizationally unique identifier (OUI), Subtype , Index , and Info (appears when the interface option is used).
Age	How long the neighbor has been identified (appears when the interface option is used and NetBIOS snooping is enabled on the switch).
Local Interface	Name of the local physical interface (appears when the interface option is used and NetBIOS snooping is enabled on the switch).
Parent Interface	Name of the aggregated Ethernet interface, if any, to which the interface belongs (appears when the interface option is used and NetBIOS snooping is enabled on the switch).
Chassis ID	Chassis identifier of the chassis type listed (appears when the interface option is used and NetBIOS snooping is enabled on the switch).
Port description	Port description (appears when the interface option is used and NetBIOS snooping is enabled on the switch).
System name	NetBIOS name of the host (appears when the interface option is used and NetBIOS snooping is enabled on the switch).

Sample Output

show lldp neighbors

```
user@switch> show lldp neighbors
```

Local Interface	Parent Interface	Chassis Id	Port info	System Name
xe-3/0/4.0	ae31.0	b0:c6:9a:63:80:40	xe-0/0/0.0	newyork31
xe-3/0/5.0	ae31.0	b0:c6:9a:63:80:40	xe-0/0/1.0	newyork31
xe-3/0/6.0	ae31.0	b0:c6:9a:63:80:40	xe-0/0/2.0	newyork31
xe-3/0/7.0	ae31.0	b0:c6:9a:63:80:40	xe-0/0/3.0	newyork31
xe-3/0/0.0	ae31.0	b0:c6:9a:63:80:40	xe-0/1/0.0	newyork31
xe-3/0/1.0	ae31.0	b0:c6:9a:63:80:40	xe-0/1/1.0	newyork31
xe-3/0/2.0	ae31.0	b0:c6:9a:63:80:40	xe-0/1/2.0	newyork31
xe-3/0/3.0	ae31.0	b0:c6:9a:63:80:40	xe-0/1/3.0	newyork31

show lldp neighbors interface

```
user@switch> show lldp neighbors interface ge-0/0/2
```

LLDP Neighbor Information:

Local Information:

```
Index: 1 Time to live: 240 Time mark: Wed Dec 1 10:23:24 2010 Age: 29 secs
Local Interface   : ge-0/0/2.0
Parent Interface  : -
Local Port ID     : 507
Ageout Count      : 0
```

Neighbour Information:

```
Chassis type      : Mac address
Chassis ID        : 00:1f:12:38:7f:c0
Port type         : Locally assigned
Port ID           : 507
Port description  : ge-0/0/2.0
System name       : bng-148p5-dev
```

```
System Description : Juniper Networks, Inc. ex4200-48p , version 10.4I0 Build
date: 2010-11-30 09:32:17 UTC
```

System capabilities

```
Supported : Bridge Router
Enabled   : Bridge Router
```

Management Info

```
Type           : IPv4
Address        : 10.204.96.235
Port ID        : 34
Subtype        : 1
Interface Subtype : ifIndex(2)
OID            : 1.3.6.1.2.1.31.1.1.1.34
```

```
Media endpoint class: Network Connectivity
```

Organization Info

```
OUI      : 0.12.f
Subtype   : 1
Index     : 1
Info      : 22A8360000
```

Organization Info

```
OUI      : 0.12.f
```

Subtype : 2
Index : 2
Info : 030100

show lldp statistics

Syntax	<code>show lldp statistics</code> <code><interface <i>interface-ids</i>></code>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Description	Display LLDP statistics on all or selected interfaces.
Options	none —Display LLDP statistics on all interfaces and devices. interface <i>interface-ids</i> —(Optional) Display LLDP statistics on the selected devices.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Configuring LLDP • Understanding LLDP on page 1795
List of Sample Output	show lldp statistics on page 2014
Output Fields	Table 162 lists the output fields for the show lldp statistics command. Output fields are listed in the approximate order in which they appear.

Table 162: show lldp statistics Output Fields

Field Name	Field Description	Level of Output
Interface	Name of an interface.	All levels
Received	Total number of LLDP frames received on an interface.	All levels
Unknown-TLVs	Number of unrecognized LLDP TLVs received on an interface.	All levels
With Errors	Number of LLDP frames received that contain errors.	All levels
Discarded TLVs	Number of LLDP TLVs received and then discarded on an interface.	All levels
Transmitted	Total number of LLDP frames transmitted on an interface.	All levels
Untransmitted	Total number of LLDP frames not transmitted on an interface.	All levels

Sample Output

show lldp statistics

```
user@switch> show lldp statistics
```

```
Interface  Received  Unknown TLVs  With Errors  Discarded TLVs  Transmitted
Untransmitted
```

me0.0	0	0	0	0	8003	0
ge-0/0/0.0	8002	0	0	0	8003	0
ge-0/0/1.0	8002	0	0	0	8003	0

show route instance

Syntax	show route instance <brief detail summary> <instance-name> <operational>
Release Information	Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	(QFabric systems only) Display routing instance information.
Options	<p>none—(Same as brief) Display standard information about all routing instances.</p> <p>brief detail summary—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief. (These options are not available with the operational keyword.)</p> <p>instance-name—(Optional) Display information for a specified routing instance.</p> <p>operational—(Optional) Display operational routing instances.</p>
Required Privilege Level	view
List of Sample Output	show route instance on page 2017 show route instance detail on page 2017 show route instance operational on page 2018 show route instance summary on page 2018
Output Fields	Table 163 lists the output fields for the show route instance command. Output fields are listed in the approximate order in which they appear.

Table 163: show route instance Output Fields

Field Name	Field Description	Level of Output
Instance or <i>instance-name</i>	Name of the routing instance.	All levels
Operational Routing Instances	(operational keyword only) Names of all operational routing instances.	—
Type	Type of routing instance: forwarding or virtual-router .	All levels
State	State of the routing instance: active or inactive .	detail
Interfaces	Name of interfaces belonging to this routing instance.	detail
Tables	Tables (and number of routes) associated with this routing instance.	detail
Router ID	Identifier for the router.	detail

Table 163: show route instance Output Fields (*continued*)

Field Name	Field Description	Level of Output
Primary RIB	Primary table for this routing instance.	brief none summary
Active/holddown/hidden	Number of active, hold-down, and hidden routes.	All levels

Sample Output

show route instance

```

user@switch> show route instance
Instance          Type
Primary RIB
master            forwarding
                  inet.0
                  4/0/1

__juniper_private1__ forwarding
                  __juniper_private1__.inet.0
                  1/0/3

__juniper_private2__ forwarding
                  __juniper_private2__.inet.0
                  0/0/1

__juniper_private3__ forwarding
                  __juniper_private3__.inet.0
                  1/0/2

__juniper_private4__ forwarding
                  __juniper_private4__.inet.0
                  4/0/2

__master.anon__   forwarding

r1                virtual-router

r2                virtual-router

```

show route instance detail

```

user@switch> show route instance detail
master:
  Router ID: 3.3.3.7
  Type: forwarding      State: Active
  Tables:
    inet.0              : 5 routes (4 active, 0 holddown, 1 hidden)

__juniper_private1__:
  Router ID: 0.0.0.0
  Type: forwarding      State: Active
  Interfaces:
    lo0.16385
    bme0.0
  Tables:
    __juniper_private1__.inet.0: 6 routes (1 active, 0 holddown, 3 hidden)

__juniper_private2__:
  Router ID: 0.0.0.0
  Type: forwarding      State: Active
  Interfaces:
    lo0.16384

```

```

Tables:
  __juniper_private2__.inet.0: 1 routes (0 active, 0 holddown, 1 hidden)

__juniper_private3__:
Router ID: 0.0.0.0
Type: forwarding      State: Active
Interfaces:
  bme0.1
Tables:
  __juniper_private3__.inet.0: 4 routes (1 active, 0 holddown, 2 hidden)

__juniper_private4__:
Router ID: 0.0.0.0
Type: forwarding      State: Active
Interfaces:
  bme0.2
Tables:
  __juniper_private4__.inet.0: 8 routes (4 active, 0 holddown, 2 hidden)

__master.anon__:
Router ID: 0.0.0.0
Type: forwarding      State: Active

r1:
Router ID: 0.0.0.0
Type: virtual-router  State: Active
Interfaces:
  xe-0/0/0.0

r2:
Router ID: 0.0.0.0
Type: virtual-router  State: Active
Interfaces:
  xe-0/0/3.0

```

show route instance operational

```

user@switch> show route instance operational
Operational Routing Instances:

__juniper_private1__
__juniper_private2__
__juniper_private3__
__juniper_private4__
r1---qfabric
r2---qfabric
master

```

show route instance summary

```

user@switch> show route instance summary
Instance      Type
Primary RIB
master        forwarding
inet.0        4/0/1

__juniper_private1__ forwarding
__juniper_private1__.inet.0 1/0/3

__juniper_private2__ forwarding
__juniper_private2__.inet.0 0/0/1

```

__juniper_private3__ forwarding	
__juniper_private3__.inet.0	1/0/2
__juniper_private4__ forwarding	
__juniper_private4__.inet.0	4/0/2
__master.anon__ forwarding	
r1	virtual-router
r2	virtual-router

show snmp statistics

Syntax	<code>show snmp statistics</code> <code><subagents></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Option subagents introduced in Junos OS Release 14.2.
Description	Display statistics about Simple Network Management Protocol (SNMP) packets sent and received by the router or switch.
Options	subagents —(Optional) Display the statistics of the protocol data unit (PDU), the number of SNMP requests and responses per subagent, and the SNMP statistics received from each subagent per logical system.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• clear snmp statistics on page 1739
List of Sample Output	show snmp statistics on page 2025 show snmp statistics subagents on page 2025
Output Fields	Table 130 describes the output fields for the show snmp statistics command. Output fields are listed in the approximate order in which they appear.

Table 164: show snmp statistics Output Fields

Field Name	Field Description
Input	<p>Information about received packets:</p> <ul style="list-style-type: none"> • Packets(snmplnPkts)—Total number of messages delivered to the SNMP entity from the transport service. • Bad versions—(snmplnBadVersions) Total number of messages delivered to the SNMP entity that were for an unsupported SNMP version. • Bad community names—(snmplnBadCommunityNames) Total number of messages delivered to the SNMP entity that used an SNMP community name not known to the entity. • Bad community uses—(snmplnBadCommunityUses) Total number of messages delivered to the SNMP entity that represented an SNMP operation that was not allowed by the SNMP community named in the message. • ASN parse errors—(snmplnASNParseErrs) Total number of ASN.1 or BER errors encountered by the SNMP entity when decoding received SNMP messages. • Too big—(snmplnTooBigs) Total number of SNMP PDUs delivered to the SNMP entity with an error status field of tooBig. • No such names—(snmplnNoSuchNames) Total number of SNMP PDUs delivered to the SNMP entity with an error status field of noSuchName. • Bad values—(snmplnBadValues) Total number of SNMP PDUs delivered to the SNMP entity with an error status field of badValue. • Read onlys—(snmplnReadOnlys) Total number of valid SNMP PDUs delivered to the SNMP entity with an error status field of readOnly. Only incorrect implementations of SNMP generate this error.

Table 164: show snmp statistics Output Fields (*continued*)

Field Name	Field Description
Input (continued)	<ul style="list-style-type: none"> • General errors—(snmpInGenErrs) Total number of SNMP PDUs delivered to the SNMP entity with an error status field of genErr. • Total requests varbinds—(snmpInTotalReqVars) Total number of MIB objects retrieved successfully by the SNMP entity as a result of receiving valid SNMP GetRequest and GetNext PDUs. • Total set varbinds—(snmpInSetVars) Total number of MIB objects modified successfully by the SNMP entity as a result of receiving valid SNMP SetRequest PDUs. • Get requests—(snmpInGetRequests) Total number of SNMP GetRequest PDUs that have been accepted and processed by the SNMP entity. • Get nexts—(snmpInGetNexts) Total number of SNMP GetNext PDUs that have been accepted and processed by the SNMP entity. • Set requests—(snmpInSetRequests) Total number of SNMP SetRequest PDUs that have been accepted and processed by the SNMP entity. • Get responses—(snmpInGetResponses) Total number of SNMP GetResponse PDUs that have been accepted and processed by the SNMP entity. • Traps—(snmpInTraps) Total number of SNMP traps generated by the SNMP entity. • Silent drops—(snmpSilentDrops) Total number of GetRequest, GetNextRequest, GetBulkRequest, SetRequests, and InformRequest PDUs delivered to the SNMP entity that were silently dropped because the size of a reply containing an alternate response PDU with an empty variable-bindings field was greater than either a local constraint or the maximum message size associated with the originator of the requests. • Proxy drops—(snmpProxyDrops) Total number of GetRequest, GetNextRequest, GetBulkRequest, SetRequests, and InformRequest PDUs delivered to the SNMP entity that were silently dropped because the transmission of the message to a proxy target failed in such a way (other than a timeout) that no response PDU could be returned. • Commit pending drops—Number of SNMP packets for Set requests dropped because of a previous pending SNMP Set request on the committed configuration. • Throttle drops—Number of SNMP packets for any requests dropped reaching the throttle limit.

Table 164: show snmp statistics Output Fields (*continued*)

Field Name	Field Description
V3 Input	<p>Information about SNMP version 3 packets:</p> <ul style="list-style-type: none"> • Unknown security models—(snmpUnknownSecurityModels) Total number of packets received by the SNMP engine that were dropped because they referenced a security model that was not known to or supported by the SNMP engine. • Invalid messages—(snmpInvalidMsgs) Number of packets received by the SNMP engine that were dropped because there were invalid or inconsistent components in the SNMP message. • Unknown pdu handlers—(snmpUnknownPDUHandlers) Number of packets received by the SNMP engine that were dropped because the PDU contained in the packet could not be passed to an application responsible for handling the PDU type. • Unavailable contexts—(snmpUnavailableContexts) Number of requests received for a context that is known to the SNMP engine, but is currently unavailable. • Unknown contexts—(snmpUnknownContexts) Total number of requests received for a context that is unknown to the SNMP engine. • Unsupported security levels—(usmStatsUnsupportedSecLevels) Total number of packets received by the SNMP engine that were dropped because they requested a security level unknown to the SNMP engine (or otherwise unavailable). • Not in time windows—(usmStatsNotInTimeWindows) Total number of packets received by the SNMP engine that were dropped because they appeared outside the authoritative SNMP engine's window. • Unknown user names—(usmStatsUnknownUserNames) Total number of packets received by the SNMP engine that were dropped because they referenced a user that was not known to the SNMP engine. • Unknown engine ids—(usmStatsUnknownEngineIDs) Total number of packets received by the SNMP engine that were dropped because they referenced an SNMP engine ID that was not known to the SNMP engine. • Wrong digests—(usmStatsWrongDigests) Total number of packets received by the SNMP engine that were dropped because they did not contain the expected digest value. • Decryption errors—(usmStatsDecryptionErrors) Total number of packets received by the SNMP engine that were dropped because they could not be decrypted.

Table 164: show snmp statistics Output Fields (*continued*)

Field Name	Field Description
Output	<p>Information about transmitted packets:</p> <ul style="list-style-type: none"> • Packets—(snmpOutPkts) Total number of messages passed from the SNMP entity to the transport service. • Too big—(snmpOutTooBigs) Total number of SNMP PDUs generated by the SNMP entity with an error status field of tooBig. • No such names—(snmpOutNoSuchNames) Total number of SNMP PDUs delivered to the SNMP entity with an error status field of noSuchName. • Bad values—(snmpOutBadValues) Total number of SNMP PDUs generated by the SNMP entity with an error status field of badValue. • General errors—(snmpOutGenErrs) Total number of SNMP PDUs generated by the SNMP entity with an error status field of genErr. • Get requests—(snmpOutGetRequests) Total number of SNMP GetRequest PDUs generated by the SNMP entity. • Get nexts—(snmpOutGetNexts) Total number of SNMP GetNext PDUs generated by the SNMP entity. • Set requests—(snmpOutSetRequests) Total number of SNMP SetRequest PDUs generated by the SNMP entity. • Get responses—(snmpOutGetResponses) Total number of SNMP GetResponse PDUs generated by the SNMP entity. • Traps—(snmpOutTraps) Total number of SNMP traps generated by the SNMP entity.

Table 131 describes the output fields for the **show snmp statistics subagents** command. Output fields are listed in the approximate order in which they appear.

Table 165: show snmp statistics subagents Output Fields

Field Name	Field Description
Subagent	Location of the SNMP subagent.
Request PDUs	Number of PDUs requested by the SNMP manager.
Response PDUs	Number of response PDUs sent by the SNMP subagent.
Request Variables	Number of variable bindings on the PDUs requested by the SNMP manager.
Response Variables	Number of variable bindings on the PDUs sent by the SNMP subagent.
Average Response Time	Average time taken by the SNMP subagent to send statistics response.
Maximum Response Time	Maximum time taken by the SNMP subagent to send the statistics response.

Sample Output

show snmp statistics

```
user@host> show snmp statistics
SNMP statistics:
  Input:
    Packets: 246213, Bad versions: 12, Bad community names: 12,
    Bad community uses: 0, ASN parse errors: 96,
    Too big: 0, No such names: 0, Bad values: 0,
    Read onlys: 0, General errors: 0,
    Total request varbinds: 227084, Total set varbinds: 67,
    Get requests: 44942, Get nexts: 190371, Set requests: 10712,
    Get responses: 0, Traps: 0,
    Silent drops: 0, Proxy drops: 0, Commit pending drops: 0,
    Throttle drops: 0,
  V3 Input:
    Unknown security models: 0, Invalid messages: 0
    Unknown pdu handlers: 0, Unavailable contexts: 0
    Unknown contexts: 0, Unsupported security levels: 1
    Not in time windows: 0, Unknown user names: 0
    Unknown engine ids: 44, Wrong digests: 23, Decryption errors: 0
  Output:
    Packets: 246093, Too big: 0, No such names: 31561,
    Bad values: 0, General errors: 2,
    Get requests: 0, Get nexts: 0, Set requests: 0,
    Get responses: 246025, Traps: 0
```

show snmp statistics subagents

```
user@host> show snmp statistics subagents

Subagent: /var/run/cosd-20
  Request PDUs: 0, Response PDUs: 0,
  Request Variables: 0, Response Variables: 0,
  Average Response Time(ms): 0.00,
  Maximum Response Time(ms): 0.00

Subagent: /var/run/pfed-30
  Request PDUs: 0, Response PDUs: 0,
  Request Variables: 0, Response Variables: 0,
  Average Response Time(ms): 0.00,
  Maximum Response Time(ms): 0.00

Subagent: /var/run/rmopd-15
  Request PDUs: 0, Response PDUs: 0,
  Request Variables: 0, Response Variables: 0,
  Average Response Time(ms): 0.00,
  Maximum Response Time(ms): 0.00

Subagent: /var/run/chassisd-30
  Request PDUs: 33116, Response PDUs: 33116,
  Request Variables: 33116, Response Variables: 33116,
  Average Response Time(ms): 1.83,
  Maximum Response Time(ms): 203.48

Subagent: /var/run/pkid-13
  Request PDUs: 0, Response PDUs: 0,
  Request Variables: 0, Response Variables: 0,
  Average Response Time(ms): 0.00,
  Maximum Response Time(ms): 0.00
```

Subagent: /var/run/apsd-13
Request PDUs: 0, Response PDUs: 0,
Request Variables: 0, Response Variables: 0,
Average Response Time(ms): 0.00,
Maximum Response Time(ms): 0.00

Subagent: /var/run/dfcd-32
Request PDUs: 0, Response PDUs: 0,
Request Variables: 0, Response Variables: 0,
Average Response Time(ms): 0.00,
Maximum Response Time(ms): 0.00

Subagent: /var/run/mib2d-33
Request PDUs: 74211, Response PDUs: 74211,
Request Variables: 74211, Response Variables: 74211,
Average Response Time(ms): 2.30,
Maximum Response Time(ms): 51.04

Subagent: /var/run/license-check-16
Request PDUs: 0, Response PDUs: 0,
Request Variables: 0, Response Variables: 0,
Average Response Time(ms): 0.00,
Maximum Response Time(ms): 0.00

Subagent: /var/run/craftd-14
Request PDUs: 0, Response PDUs: 0,
Request Variables: 0, Response Variables: 0,
Average Response Time(ms): 0.00,
Maximum Response Time(ms): 0.00

Subagent: /var/run/bfdd-19
Request PDUs: 0, Response PDUs: 0,
Request Variables: 0, Response Variables: 0,
Average Response Time(ms): 0.00,
Maximum Response Time(ms): 0.00

Subagent: /var/run/smihelperd-24
Request PDUs: 0, Response PDUs: 0,
Request Variables: 0, Response Variables: 0,
Average Response Time(ms): 0.00,
Maximum Response Time(ms): 0.00

Subagent: /var/run/cfmd-18
Request PDUs: 0, Response PDUs: 0,
Request Variables: 0, Response Variables: 0,
Average Response Time(ms): 0.00,
Maximum Response Time(ms): 0.00

Subagent: /var/run/rpd_snmp
Request PDUs: 0, Response PDUs: 0,
Request Variables: 0, Response Variables: 0,
Average Response Time(ms): 0.00,
Maximum Response Time(ms): 0.00

Subagent: /var/run/l2tpd-18
Request PDUs: 0, Response PDUs: 0,
Request Variables: 0, Response Variables: 0,
Average Response Time(ms): 0.00,
Maximum Response Time(ms): 0.00

ssh

List of Syntax [Syntax on page 2028](#)
 [Syntax \(EX Series Switch and the QFX Series\) on page 2028](#)

Syntax `ssh host`
 `<bypass-routing>`
 `<inet | inet6>`
 `<interface interface-name>`
 `<logical-system logical-system-name>`
 `<routing-instance routing-instance-name>`
 `<source address>`
 `<v1 | v2>`

Syntax (EX Series Switch and the QFX Series) `ssh host`
 `<bypass-routing>`
 `<inet | inet6>`
 `<interface interface-name>`
 `<routing-instance routing-instance-name>`
 `<source address>`
 `<v1 | v2>`

Release Information Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.0 for EX Series switches.
 Command introduced in Junos OS Release 11.1 for the QFX Series.
 Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.

Description Use the SSH program to open a connection between a local router or switch and a remote system and execute commands on the remote system. You can issue the **ssh** command from the Junos OS CLI to log in to a remote system or from a remote system to log in to the local router or switch. When executing this command, you include one or more CLI commands by enclosing them in quotation marks and separating the commands with semicolons:

```
ssh address 'cli-command1 ; cli-command2 '
```

Options **host**—Name or address of the remote system.

bypass-routing—(Optional) Bypass the normal routing tables and send ping requests directly to a system on an attached network. If the system is not on a directly attached network, an error is returned. Use this option to ping a local system through an interface that has no route through it.

inet | inet6—(Optional) Create an IPv4 or IPv6 connection, respectively.

interface interface-name—(Optional) Interface name for the SSH session. (This option does not work when **default-address-selection** is configured at the **[edit system]** hierarchy level, because this configuration uses the loopback interface as the source address for all locally generated IP packets.)

logical-system logical-system-name—(Optional) Name of a particular logical system for the SSH attempt.

routing-instance *routing-instance-name*—(Optional) Name of the routing instance for the SSH attempt.

source address—(Optional) Source address of the SSH connection.

v1 | v2—(Optional) Use SSH version 1 or 2, respectively, when connecting to a remote host.

Additional Information To configure an SSH (version 1) key for your user account, include the **authentication ssh-rsa** statement at the **[edit system login user *user-name*]** hierarchy level. To configure an SSH (version 2) key for your user account, include the **authentication dsa-rsa** statement at the **[edit system login user *user-name*]** hierarchy level.

You can limit the number of times a user can attempt to enter a password while logging in through SSH. To specify the number of times a user can attempt to enter a password to log in through SSH, include the **retry-options** statement at the **[edit system login]** hierarchy level. For details, see the .

Required Privilege Level network

Related Documentation • *Configuring SSH Host Keys for Secure Copying of Data*

List of Sample Output [ssh on page 2029](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

ssh

```
user@switch> ssh cree
Host key not found from the list of known hosts.
Are you sure you want to continue connecting (yes/no)? yes

Host ?cree' added to the list of known hosts.
boojun@cree's password:
Last login: Sun Jun 21 10:43:42 1998 from junos-router
% ...
```


Ethernet Switching Feature Guide for QFX10000 Switches

PART 27

ELS CLI

- [Enhanced Layer 2 Software \(ELS\) CLI on page 2035](#)

Enhanced Layer 2 Software (ELS) CLI

- [Getting Started with Enhanced Layer 2 Software on page 2035](#)

Getting Started with Enhanced Layer 2 Software

- [Understanding Enhanced Layer 2 Software Support on page 2035](#)
- [Using the ELS Translator on page 2036](#)
- [Configuring a VLAN on page 2037](#)
- [Configuring the Native VLAN Identifier on page 2037](#)
- [Configuring Layer 2 Interfaces on page 2038](#)
- [Configuring Layer 3 Interfaces on page 2038](#)
- [Configuring an IRB Interface on page 2039](#)
- [Configuring an Aggregated Ethernet Interface and Configuring LACP on That Interface on page 2039](#)
- [Enhanced Layer 2 CLI Configuration Statement and Command Changes on page 2040](#)

Understanding Enhanced Layer 2 Software Support

Enhanced Layer 2 Software (ELS) is automatically supported if your device is running a Junos OS release that supports it. You do not need to take any action to enable ELS, and you cannot disable ELS.

Table 1 lists the EX Series and QFX Series switches that support ELS and the initial ELS release they support..

Table 166: ELS Support

Device	Initial ELS Release
EX4300 switches	13.2X50-D10
EX4600 switches	13.2X51-D25
EX9200 switches	12.3R2
QFX3500 switches	13.2X50-D15

Table 166: ELS Support (*continued*)

Device	Initial ELS Release
QFX3600 switches	13.2X50-D15
QFX5100 switches	13.2X51-D10
QFX10000 switches	15.1X53-D10

ELS is supported on the EX4300, EX4600, and EX9200 switches for all Junos OS releases, starting with the initial releases shown in [Table 3](#).

ELS support was introduced on QFX3500 and QFX3600 switches in Junos OS Release 13.2X50-D15. ELS is supported only on the software package that supports Virtual Chassis (the **jinstall-qfx-3-*** software package) for QFX3500 and QFX3600 switches.

For QFX5100 switches, ELS support was introduced in Junos OS Release 13.2X51-D10 and is supported on the **jinstall-qfx-5-*** software package.



NOTE: ELS is not supported on software packages that can be installed in a QFabric system.

Using the ELS Translator

ELS Translator is a Web-based tool that converts Junos OS Layer 2 configurations to Enhanced Layer 2 Software (ELS) configurations. This conversion tool supports all Juniper Networks EX Series, MX Series, and QFX Series devices with ELS installed. ELS Translator is hosted on the Juniper Networks Customer Support website for EX Series switches, MX Series routers, and QFX Series switches and is available to registered users, internal users, partners, and premium service contract customers. You need to log in using your Juniper Networks username and password to access the ELS Translator.

[Click here](#) to access the ELS translator.

If you are upgrading from a version of Junos OS that does not support ELS to a version of Junos OS that supports ELS, we recommend that you update your configuration with the ELS Translator using the following procedure:

1. Log in to your device by using the console port.



NOTE: Perform this procedure only from the console port. You will lose connectivity to your device if you perform this procedure from a management port or any other interface.

2. Copy the entire existing configuration to another file. Save the file in a remote location. See [“Saving a Configuration to a File” on page 1331](#).

3. Retain the portion of your existing configuration related to management network connectivity (such as **[edit system]** hierarchy level). Delete all other top-level configuration hierarchy levels (such as the **[edit interfaces]**, **[edit protocols]**, and **[edit vlans]**). Issue the **commit** command to remove the deleted configuration hierarchy levels.
4. Perform the software upgrade. Reboot your device to complete the upgrade. See [“Software Installation Overview” on page 174](#).



NOTE: Ensure that the console port connection is up during the reboot.

5. [Click here](#) to access the ELS Translator in a web browser. Follow the instructions on the page to update your configuration.
6. Return to your console port connection. When the switch has rebooted to complete the software upgrade, copy the configuration from the ELS Translator to your switch. See [“Uploading a Configuration File” on page 1332](#).
7. Commit the new configuration.



NOTE: It is possible that scripts do not translate correctly. Therefore, review translated scripts carefully before loading the converted configuration on your switch or other device.

Configuring a VLAN

You can configure one or more VLANs to perform Layer 2 bridging. The Layer 2 bridging functions include integrated routing and bridging (IRB) for support for Layer 2 bridging and Layer 3 IP routing on the same interface. EX Series and QFX Series switches can function as Layer 2 switches, each with multiple bridging, or broadcast, domains that participate in the same Layer 2 network. You can also configure Layer 3 routing support for a VLAN.

To configure a VLAN:

1. Create the VLAN by setting a unique VLAN name and configuring the VLAN ID:


```
[edit]
user@host# set vlans vlan-name vlan-id vlan-id-number
```
2. Assign at least one interface to the VLAN:


```
[edit]
user@host# set interface interface-name family ethernet-switching vlan members vlan-name
```

Configuring the Native VLAN Identifier

EX Series and QFX Series switches support receiving and forwarding routed or bridged Ethernet frames with 802.1Q VLAN tags. Typically, trunk ports, which connect switches to each other, accept untagged control packets, but do not accept untagged data packets. You can enable a trunk port to accept untagged data packets by configuring a native VLAN ID on the interface on which you want the untagged data packets to be received.

To configure the native VLAN ID:

1. On the interface on which you want untagged data packets to be received, set the interface mode to **trunk**, which specifies that the interface is in multiple VLANs and can multiplex traffic between different VLANs.

```
[edit interfaces]
user@host# set interface-name unit logical-unit-number family ethernet-switching
interface-mode trunk
```

2. Configure the native VLAN ID:

```
[edit interfaces]
user@host# set interface-name native-vlan-id number
```

3. Assign the interface to the native VLAN ID:

```
[edit interfaces]
user@host# set interface-name unit logical-unit-number family ethernet-switching vlan
members native-vlan-id-number
```

Configuring Layer 2 Interfaces

To ensure that your high-traffic network is tuned for optimal performance, explicitly configure some settings on the switch's network interfaces.

To configure a Gigabit Ethernet interface or a 10-Gigabit Ethernet interface as a **trunk** interface:

```
[edit]
user@host# set interfaces interface-name unit logical-unit-number family ethernet-switching
interface-mode trunk
```

To configure a Gigabit Ethernet interface or a 10-Gigabit Ethernet interface as a **access** interface:

```
[edit]
user@host# set interfaces interface-name unit logical-unit-number family ethernet-switching
interface-mode access
```

Configuring Layer 3 Interfaces

To configure a Layer 3 interface, you must assign an IP address to the interface. You assign an address to an interface by specifying the address when you configure the protocol family. For the **inet** or **inet6** family, configure the interface IP address.

You can configure interfaces with a 32-bit IP version 4 (IPv4) address and optionally with a destination prefix, sometimes called a subnet mask. An IPv4 address utilizes a 4-octet dotted decimal address syntax (for example, 192.16.1.1). An IPv4 address with destination prefix utilizes a 4-octet dotted decimal address syntax with a destination prefix appended (for example, 192.16.1.1/30).

To specify an IP4 address for the logical unit:

```
[edit]
user@host# set interfaces interface-name unit logical-unit-number family inet address ip-address
```

You represent IP version 6 (IPv6) addresses in hexadecimal notation by using a colon-separated list of 16-bit values. You assign a 128-bit IPv6 address to an interface.

To specify an IP6 address for the logical unit:

```
[edit]
user@host# set interfaces interface-name unit logical-unit-number family inet6 address ip-address
```

Configuring an IRB Interface

Integrated routing and bridging (IRB) provides support for Layer 2 bridging and Layer 3 IP routing on the same interface. IRB enables you to route packets to another routed interface or to another VLAN that has a Layer 3 protocol configured. IRB interfaces enable the device to recognize packets that are being sent to local addresses so that they are bridged (switched) whenever possible and are routed only when necessary. Whenever packets can be switched instead of routed, several layers of processing are eliminated. An interface named *irb* functions as a logical router on which you can configure a Layer 3 logical interface for VLAN. For redundancy, you can combine an IRB interface with implementations of the Virtual Router Redundancy Protocol (VRRP) in both bridging and virtual private LAN service (VPLS) environments.

To configure an IRB interface:

1. Create a Layer 2 VLAN by assigning it a name and a VLAN ID:

```
[edit]
user@host# set vlans vlan-name vlan-id vlan-id
```

2. Create an IRB logical interface:

```
[edit]
user@host# set interface irb unit logical-unit-number family inet address ip-address
```

3. Associate the IRB interface with the VLAN:

```
[edit]
user@host# set vlans vlan-name l3-interface irb.logical-unit-number
```

Configuring an Aggregated Ethernet Interface and Configuring LACP on That Interface

Use the link aggregation feature to aggregate one or more links to form a virtual link or link aggregation group (LAG). The MAC client can treat this virtual link as if it were a single link to increase bandwidth, provide graceful degradation as failure occurs, and increase availability.

To configure an aggregated Ethernet interface:

1. Specify the number of aggregated Ethernet interfaces to be created:

```
[edit chassis]
user@host# set aggregated-devices ethernet device-count number
```

2. Specify the name of the link aggregation group interface:

```
[edit interfaces]
user@host# set interfaces aex
```

3. Specify the minimum number of links for the aggregated Ethernet interface (*aex*)—that is, the defined bundle—to be labeled *up*:

```
[edit interfaces]
user@host# set aex aggregated-ether-options minimum-links number
```

4. Specify the link speed for the aggregated Ethernet bundle:

```
[edit interfaces]
```

```
user@host# set aex aggregated-ether-options link-speed link-speed
```

5. Specify the members to be included within the aggregated Ethernet bundle:

```
[edit interfaces]
user@host# set interface-name ether-options 802.3ad aex
user@host# set interface-name ether-options 802.3ad aex
```

6. Specify an interface family for the aggregated Ethernet bundle:

```
[edit interfaces]
user@host# set aex unit 0 family inet address ip-address
```

For aggregated Ethernet interfaces on the device, you can configure the Link Aggregation Control Protocol (LACP). LACP bundles several physical interfaces to form one logical interface. You can configure aggregated Ethernet with or without LACP enabled.

When LACP is enabled, the local and remote sides of the aggregated Ethernet links exchange protocol data units (PDUs), containing information about the state of the link. You can configure Ethernet links to actively transmit PDUs, or you can configure the links to passively transmit them, sending out LACP PDUs only when they receive them from another link. One side of the link must be configured as active for the link to be up.

To configure LACP:

1. Enable one side of the aggregated Ethernet link as active:

```
[edit interfaces]
user@host# set aex aggregated-ether-options lacp active
```

2. Specify the interval at which the interfaces send LACP packets:

```
[edit interfaces]
user@host# set aex aggregated-ether-options lacp periodic interval
```

Enhanced Layer 2 CLI Configuration Statement and Command Changes

The enhanced Layer 2 CLI feature is introduced in Junos OS Release 12.3R2. The enhanced Layer 2 CLI feature changes the CLI for some of the Layer 2 features on EX Series switches. This enhanced CLI will be used to configure Layer 2 features on future EX Series hardware platforms, and also to configure Layer 2 features on other Juniper Networks products.

The following sections provide a list of existing commands that were moved to new hierarchy levels or changed on EX Series switches as part of this CLI enhancement effort. These sections are provided as a high-level reference only. For detailed information about these commands, use the links to the configuration statements provided or see the technical documentation.

- [Changes to the ethernet-switching-options Hierarchy Level on page 2041](#)
- [Changes to the Port Mirroring Hierarchy Level on page 2042](#)
- [Changes to the Layer 2 Control Protocol Hierarchy Level on page 2043](#)
- [Changes to the dot1q-tunneling Statement on page 2043](#)
- [Changes to the L2 Learning Protocol on page 2044](#)
- [Changes to Nonstop Bridging on page 2044](#)
- [Changes to Port Security and DHCP Snooping on page 2044](#)
- [Changes to Configuring VLANs on page 2046](#)

- [Changes to Storm Control Profiles on page 2049](#)
- [Changes to the Interfaces Hierarchy on page 2050](#)
- [Changes to IGMP Snooping on page 2051](#)

Changes to the ethernet-switching-options Hierarchy Level

This section outlines the changes to the **ethernet-switching-options** hierarchy level.



NOTE: The **ethernet-switching-options** hierarchy level has been renamed as **switch-options**.

Table 167: Renaming the ethernet-switching-options hierarchy

Original Hierarchy	Changed Hierarchy
<pre>ethernet-switching-options { authentication-whitelist { ... } }</pre>	<pre>switch-options { ... authentication-whitelist { ... } }</pre>
<pre>ethernet-switching-options { interfaces <i>interface-name</i> { no-mac-learning; ... } }</pre>	<pre>switch-options { interfaces <i>interface-name</i> { no-mac-learning; ... } }</pre>
<pre>ethernet-switching-options { unknown-unicast-forwarding { (...) } }</pre>	<pre>switch-options { unknown-unicast-forwarding { (...) } }</pre>
<pre>ethernet-switching-options { voip { interface (all [<i>interface-name</i> access-ports]) { forwarding-class (assured-forwarding best-effort expedited-forwarding network-control); vlan <i>vlan-name</i>; ... } } }</pre>	<pre>switch-options { voip { interface (all [<i>interface-name</i> access-ports]) { forwarding-class (assured-forwarding best-effort expedited-forwarding network-control); vlan <i>vlan-name</i>; ... } } }</pre>

Table 168: RTG Statements

Original Hierarchy	Changed Hierarchy
<pre> ethernet-switching-options { redundant-trunk-group { group name { description; interface interface-name { primary; } } preempt-cutover-timer seconds; ... } } </pre>	<pre> switch-options { redundant-trunk-group { group name { description; interface interface-name { primary; } } preempt-cutover-timer seconds; ... } } </pre>

Table 169: Deleted Statements

Original Hierarchy	Changed Hierarchy
<pre> ethernet-switching-options { mac-notification { notification-interval seconds; ... } } </pre>	The statements have been removed from the switch-options hierarchy.
<pre> ethernet-switching-options { traceoptions { file filename <files number> <no-stamp> <replace> <size size> <world-readable no-world-readable>; flag flag <disable>; ... } } </pre>	The statements have been removed from the switch-options hierarchy.
<pre> ethernet-switching-options { port-error-disable { disable-timeout timeout; ... } } </pre>	<p>NOTE: The port-error-disable statement has been replaced with a new statement.</p> <pre> interfaces interface-name family ethernet-switching { recovery-timeout seconds; } </pre>

Changes to the Port Mirroring Hierarchy Level



NOTE: Statements have moved from the **ethernet-switching-options** hierarchy level to the **forwarding-options** hierarchy level.

Table 170: Port Mirroring hierarchy

Original Hierarchy	Changed Hierarchy
<pre> ethernet-switching-options { analyzer (Port Mirroring) { name { ... } } } </pre>	<pre> forwarding-options { analyzer (Port Mirroring) { name { ... } } } </pre>

Changes to the Layer 2 Control Protocol Hierarchy Level

The Layer 2 control protocol statements have moved from the **ethernet-switching-options** hierarchy to the **protocols** hierarchy.

Table 171: Layer 2 Control Protocol

Original Hierarchy	Changed Hierarchy
<pre> ethernet-switching-options { bpdv-block { ... } } </pre>	<pre> protocols { layer2-control { bpdv-block { ... } } } </pre>

Changes to the dot1q-tunneling Statement

The **dot1q-tunneling** statement has been replaced with a new statement and moved to a different hierarchy level.

Table 172: dot1q-tunneling

Original Hierarchy	Changed Hierarchy
<pre> ethernet-switching-options { dot1q-tunneling { ether-type (0x8100 0x88a8 0x9100); ... } } </pre>	<pre> interfaces interface-name { ether-options { ethernet-switch-profile { tag-protocol-id [tpids]; } } } interfaces interface-name { aggregated-ether-options { ethernet-switch-profile { tag-protocol-id [tpids]; } } } </pre>

Changes to the L2 Learning Protocol

The **mac-table-aging-time** statement has been replaced with a new statement and moved to a different hierarchy level.

Table 173: mac-table-aging-time statement

Original Hierarchy	Changed Hierarchy
<pre> ethernet-switching-options { mac-table-aging-time seconds; ... }</pre>	<pre> protocols { l2-learning { global-mac-table-aging-time seconds; ... } }</pre>

Changes to Nonstop Bridging

The **nonstop-bridging** statement has moved to a different hierarchy level.

Table 174: Nonstop Bridging statement

Original Hierarchy	Changed Hierarchy
<pre> ethernet-switching-options { nonstop-bridging; }</pre>	<pre> protocols { layer2-control { nonstop-bridging { } } }</pre>

Changes to Port Security and DHCP Snooping

Port security and DHCP snooping statements have moved to different hierarchy levels.



NOTE: The statement **examine-dhcp** does not exist in the changed hierarchy. DHCP snooping is now enabled automatically when other DHCP security features are enabled on a VLAN. See *Configuring Port Security Features* for additional information.

Table 175: Port Security statements

Original Hierarchy	Changed Hierarchy
<pre> ethernet-switching-options { secure-access-port { interface (all <i>interface-name</i>) { (dhcp-trusted no-dhcp-trusted); static-ip <i>ip-address</i> { mac <i>mac-address</i>; vlan <i>vlan-name</i>; } } } vlan (all <i>vlan-name</i>) { (arp-inspection no-arp-inspection); dhcp-option82 { disable; circuit-id { prefix <i>hostname</i>; use-interface-description; use-vlan-id; } remote-id { prefix (<i>hostname</i> mac none); use-interface-description; use-string <i>string</i>; } vendor-id [<i>string</i>]; } (examine-dhcp no-examine-dhcp); } (ip-source-guard no-ip-source-guard); } </pre>	<pre> vlans <i>vlan-name</i> forwarding-options{ dhcp-security { arp-inspection; group <i>group-name</i> { interface <i>interface-name</i> { static-ip <i>ip-address</i> { mac <i>mac-address</i>; } } } overrides { no-option82; trusted; } } ip-source-guard; no-dhcp-snooping; option-82 { circuit-id { prefix { host-name; routing-instance-name; } use-interface-description (device logical); use-vlan-id; } remote-id { host-name; use-interface-description (device logical); use-string <i>string</i>; } vendor-id { use-string <i>string</i>; } } } </pre>



NOTE: DHCP snooping statements have moved to a different hierarchy level.

Table 176: DHCP Snooping Statements

Original Hierarchy	Changed Hierarchy
<pre> ethernet-switching-options { secure-access-port { dhcp-snooping-file { location <i>local_pathname</i> <i>remote_URL</i>; timeout <i>seconds</i>; write-interval <i>seconds</i>; } } } </pre>	<pre> system [processes [dhcp-service dhcp-snooping-file <i>local_pathname</i> <i>remote_URL</i>; write-interval <i>interval</i>;]] </pre>

Changes to Configuring VLANs

The statements for configuring VLANs have moved to a different hierarchy level.



NOTE: When configuring xSTP on EX4300 and EX4600 switches, you must add all the interfaces in the applied VLANs in configurations. For MSTP, configure all interfaces in all VLANs at the [edit protocols mstp interface] hierarchy level.

Table 177: VLAN hierarchy

Original Hierarchy	Changed Hierarchy
<pre> ethernet-switching-options { secure-access-port vlan (all <i>vlan-name</i>) { mac-move-limit } </pre>	<pre> vlangs <i>vlan-name</i> switch-options { mac-move-limit } </pre>
<pre> ethernet-switching-options { static { vlan <i>vlan-id</i> { mac <i>mac-address</i> next-hop <i>interface-name</i>; ... } } } </pre>	<p>NOTE: Statement is replaced with a new statement and has moved to a different hierarchy level.</p> <pre> vlangs { <i>vlan-name</i> { switch-options { interface <i>interface-name</i> { static-mac <i>mac-address</i>; ... } } } } </pre>
<pre> vlangs { <i>vlan-name</i> { interface <i>interface-name</i> { egress; ingress; mapping (native (push swap) policy tag (push swap)); pvlan-trunk; ... } } } </pre>	<p>These statements have been removed. You can assign interfaces to a VLAN using the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family ethernet-switching vlan members <i>vlan-name</i>] hierarchy.</p>
<pre> vlangs { <i>vlan-name</i> { isolation-id <i>id-number</i>; ... } } </pre>	<p>Statements have been removed.</p>

Table 177: VLAN hierarchy (*continued*)

Original Hierarchy	Changed Hierarchy
<pre> vlangs { vlan-name { l3-interface vlan.logical-interface-number; ... } } </pre>	<p>NOTE: Syntax is changed.</p> <pre> vlangs { vlan-name { l3-interface irb.logical-interface-number; ... } } </pre>
<pre> vlangs { vlan-name { l3-interface-ingress-counting layer-3-interface-name; ... } } </pre>	Statement is removed. Ingress traffic is automatically tracked.
<pre> vlangs { vlan-name { no-local-switching; ... } } </pre>	Statement is removed.
<pre> vlangs { vlan-name { no-mac-learning; ... } } </pre>	<p>Statement has been moved to different hierarchy.</p> <pre> vlangs { vlan-name { switch-options { no-mac-learning limit ... } } } </pre>
<pre> vlangs { vlan-name { primary-vlan vlan-name; ... } } </pre>	Statement has been removed.
<pre> vlangs { vlan-name { vlan-prune; ... } } </pre>	Statement is removed.
<pre> vlangs { vlan-name { vlan-range vlan-id-low-vlan-id-high; ... } } </pre>	<p>NOTE: Statement has been replaced with a new statement.</p> <pre> vlangs { vlan-name { vlan-id-list [vlan-id-numbers]; ... } } </pre>

Table 177: VLAN hierarchy (*continued*)

Original Hierarchy	Changed Hierarchy
<pre> vlangs { vlan-name { l3-interface vlan.logical-interface-number; ... } } </pre>	<p>NOTE: Syntax is changed.</p> <pre> vlangs { vlan-name { l3-interface irb.logical-interface-number; ... } } </pre>

Table 178: Statements Moved to a Different Hierarchy

Original Hierarchy	Changed Hierarchy
<pre> vlangs { vlan-name { dot1q-tunneling { customer-vlangs (id native range); layer2-protocol-tunneling all protocol-name { drop-threshold number; shutdown-threshold number; ... } } } } </pre>	<pre> interface interface-name { encapsulation extended-vlan-bridge; flexible-vlan-tagging; native-vlan-id number; unit logical-unit-number { input-vlan-map action; output-vlan-map action; vlan-id number; vlan-id-list [vlan-id vlan-id-vlan-id]; } } </pre>
<pre> vlangs { vlan-name { filter { input filter-name output filter-name; ... } } } </pre>	<pre> vlangs { vlan-name { forwarding-options { filter { input filter-name output filter-name; ... } } } } </pre>

Table 178: Statements Moved to a Different Hierarchy (*continued*)

Original Hierarchy	Changed Hierarchy
<pre> vlangs { vlan-name { mac-limit limit action action; ... } } </pre>	<pre> vlangs { vlan-name { switch-options { interface-mac-limit limit { packet-action action; ... } } } } vlangs { vlan-name { switch-options { interface interface-name { interface-mac-limit limit { packet-action action; ... } } } } } </pre>
<pre> vlangs { vlan-name { mac-table-aging-time seconds; ... } } </pre>	<pre> protocols { l2-learning { global-mac-table-aging-time seconds; ... } } </pre>

Changes to Storm Control Profiles

Storm control is configured in two steps. The first step is to create a storm control profile at the **[edit forwarding-options]** hierarchy level, and the second step is to bind the profile to a logical interface at the **[edit interfaces]** hierarchy level. See *Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches* for the changed procedure.

Table 179: Changes to the Storm Control Profile hierarchy level

Original Hierarchy	Changed Hierarchy
<pre> ethernet-switching-options { storm-control { (...) } } </pre>	<pre> forwarding-options { storm-control-profiles profile-name { (...) } } interfaces interface-name unit number family ethernet-switching { storm-control storm-control-profile; } </pre>

Changes to the Interfaces Hierarchy



NOTE: Statements have been moved to a different hierarchy.

Table 180: Changes to the Interfaces hierarchy

Original Hierarchy	Changed Hierarchy
<pre> interfaces <i>interface-name</i> { ether-options { link-mode <i>mode</i>; speed (auto-negotiation <i>speed</i>) } }</pre>	<pre> interfaces <i>interface-name</i> { link-mode <i>mode</i>; speed <i>speed</i>) }</pre>
<pre> interfaces <i>interface-name</i> { unit <i>logical-unit-number</i> { family ethernet-switching { native-vlan-id <i>vlan-id</i> } } }</pre>	<pre> interfaces <i>interface-name</i> { native-vlan-id <i>vlan-id</i> }</pre>
<pre> interfaces <i>interface-name</i> { unit <i>logical-unit-number</i> { family ethernet-switching { port-mode <i>mode</i> } } }</pre>	<p>NOTE: Statement has been replaced with a new statement.</p> <pre> interfaces <i>interface-name</i> { unit <i>logical-unit-number</i> { family ethernet-switching { interface-mode <i>mode</i> } } }</pre>
<pre> interfaces vlan</pre>	<p>NOTE: Statement has been replaced with a new statement.</p> <pre> interfaces irb</pre>

Changes to IGMP Snooping

Table 181: IGMP Snooping hierarchy

Original Hierarchy	Changed Hierarchy
<pre> protocols { igmp-snooping { traceoptions { file <i>filename</i> <files <i>number</i>> <no-stamp> <replace> <size <i>maximum-file-size</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier> <disable>; } vlan (all <i>vlan-identifier</i>) { disable; data-forwarding { receiver { install; source-vlans <i>vlan-name</i>; } source { groups <i>ip-address</i>; } } immediate-leave; interface (all <i>interface-name</i>) { multicast-router-interface; static { group <i>mcast-ip-address</i>; } } proxy { source-address <i>ip-address</i>; } robust-count <i>number</i>; } } } </pre>	<pre> protocols { igmp-snooping { vlan <i>vlan-name</i> { immediate-leave; interface <i>interface-name</i> { group-limit <1..65535> host-only-interface multicast-router-interface; immediate-leave; static { group <i>mcast-ip-address</i> { source <> } } } } l2-querier { source-address <i>ip-address</i>; } proxy { source-address <i>ip-address</i>; } query-interval <i>number</i>; query-last-member-interval <i>number</i>; query-response-interval <i>number</i>; robust-count <i>number</i>; traceoptions { file <i>filename</i> <files <i>number</i>> <no-stamp> <replace> <size <i>maximum-file-size</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier>; } } } </pre>

PART 28

Ethernet Ring Protection

- [Understanding Ethernet Ring Protection on page 2055](#)

Understanding Ethernet Ring Protection

- [Ethernet Ring Protection Switching Overview on page 2055](#)
- [Understanding Ethernet Ring Protection Switching Functionality on page 2056](#)
- [Configuring Ethernet Ring Protection Switching \(CLI Procedure\) on page 2063](#)
- [Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS on page 2066](#)

Ethernet Ring Protection Switching Overview

Ethernet ring protection switching (ERPS) helps achieve high reliability and network stability. Links in the ring will never form loops that fatally affect the network operation and services availability. The basic idea of an Ethernet ring is to use one specific link to protect the whole ring. This special link is called a *ring protection link (RPL)*. If no failure happens in other links of the ring, the RPL blocks the traffic and is not used. The RPL is controlled by a special node called an *RPL owner*. There is only one RPL owner in a ring. The RPL owner is responsible for blocking traffic over the RPL. Under ring failure conditions, the RPL owner is responsible for unblocking traffic over the RPL. A ring failure results in protection switching of the RPL traffic. An automatic protection switching (APS) protocol is used to coordinate the protection actions over the ring. Protection switching blocks traffic on the failed link and unblocks the traffic on the RPL. When the failure clears, revertive protection switching blocks traffic over the RPL and unblocks traffic on the link on which the failure is cleared.

The following standards provide detailed information on Ethernet ring protection switching:

- IEEE 802.1Q - 1998
- IEEE 802.1D - 2004
- IEEE 802.1Q - 2003
- Draft ITU-T Recommendation G.8032/Y.1344, *Ethernet Ring protection switching*
- ITU-T Y.1731, *OAM functions and mechanisms for Ethernet-based networks*

For additional information on configuring Ethernet ring protection switching on EX Series switches, see *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*.

For additional information on configuring Ethernet ring protection switching on MX Series routers, see the *Layer 2 Configuration Guide* for a complete example of Ethernet rings and information about STP loop avoidance and prevention.

**Related
Documentation**

- [Understanding Ethernet Ring Protection Switching Functionality on page 2056](#)
- [Configuring Ethernet Ring Protection Switching](#)
- [Example: Ethernet Ring Protection Switching Configuration on MX Routers](#)
- [Example: Configuring Ethernet Ring Protection Switching on EX Series Switches](#)
- [Ethernet Interfaces Feature Guide for Routing Devices](#)

Understanding Ethernet Ring Protection Switching Functionality

- [Acronyms on page 2056](#)
- [Ring Nodes on page 2057](#)
- [Ring Node States on page 2057](#)
- [Default Logging of Basic State Transitions on EX Series Switches on page 2057](#)
- [Failure Detection on page 2058](#)
- [Logical Ring on page 2058](#)
- [FDB Flush on page 2058](#)
- [Traffic Blocking and Forwarding on page 2058](#)
- [RPL Neighbor Node on page 2059](#)
- [RAPS Message Blocking and Forwarding on page 2059](#)
- [Dedicated Signaling Control Channel on page 2060](#)
- [RAPS Message Termination on page 2060](#)
- [Revertive and Non-revertive Modes on page 2061](#)
- [Multiple Rings on page 2061](#)
- [Node ID on page 2061](#)
- [Ring ID on page 2061](#)
- [Bridge Domains with the Ring Port \(MX Series Routers Only\) on page 2061](#)
- [Wait-to-Block Timer on page 2061](#)
- [Adding and Removing a Node on page 2062](#)

Acronyms

The following acronyms are used in the discussion about Ethernet ring protection switching (ERPS):

- MA—Maintenance association
- MEP—Maintenance association end point
- OAM—Operations, administration, and management (Ethernet ring protection switching uses connectivity fault management daemon)

- FDB—MAC forwarding database
- STP—Spanning Tree Protocol
- RAPS—Ring automatic protection switching
- WTB—Wait to block
- WTR—Wait to restore
- RPL—Ring protection link

Ring Nodes

Multiple nodes are used to form a ring. There are two different node types:

- Normal node—The node has no special role on the ring.
- RPL owner node—The node owns the RPL and blocks or unblocks traffic over the RPL.

Ring Node States

The following are the different states for each node of a specific ring:

- init—Not a participant of a specific ring.
- idle—No failure on the ring; the node is performing normally. For a normal node, traffic is unblocked on both ring ports. For the RPL owner or RPL neighbor, traffic is blocked on the ring port that connects to the RPL and unblocked on the other ring port.
- protection—A failure occurred on the ring. For a normal node, traffic is blocked on the ring port that connects to the failing link and unblocked on working ring ports. For the RPL owner, traffic is unblocked on both ring ports if they connect to non-failure links.
- pending—The node is recovering from failure or its state after a **clear** command is used to remove the previous manual command. When a protection group is configured, the node enters the pending state. When a node is in pending state, the WTR or WTB timer will be running. All nodes are in pending state till WTR or WTB timer expiry.
- force switch—A force switch is issued. When a force switch is issued on a node in the ring all nodes in the ring will move into the force switch state.
- manual switch—A manual switch is issued. When a manual switch is issued on a node in the ring all nodes in the ring will move into the manual switch state.

There can be only one RPL owner for each ring. The user configuration must guarantee this, because the APS protocol cannot check this.

Default Logging of Basic State Transitions on EX Series Switches



NOTE: This section applies only to EX Series switches that run Junos OS that does not support the Enhanced Layer 2 Software (ELS) configuration style.

EX Series switches automatically log basic state transitions for the ERPS protocol. No configuration is required to initiate this logging. Basic state transitions include ERPS

interface transitions from up to down, and down to up; and ERPS state transitions from idle to protection, and protection to idle.

The basic state transitions are logged in a single file named **erp-default**, which resides in the **/var/log** directory of the switch. The maximum size of this file is 15 MB.

Default logging for ERPS can capture initial ERPS interface and state transitions, which can help you troubleshoot issues that occur early in the ERPS protocol startup process. However, if more robust logging is needed, you can enable traceoptions for ERPS by entering the **traceoptions** statement in the **[edit protocols protection-group]** hierarchy.

Be aware that for ERPS, only default logging or traceoptions can be active at a time on the switch. That is, default logging for ERPS is automatically enabled and if you enable traceoptions for ERPS, the switch automatically disables default logging. Conversely, if you disable traceoptions for ERPS, the switch automatically enables default logging.

Failure Detection

Ethernet ring operation depends on quick and accurate failure detection. The failure condition *signal failure (SF)* is supported. For SF detection, an Ethernet continuity check MEP must be configured for each ring link. For fast protection switching, a 10-ms transmission period for this MEP group is supported. OAM monitors the MEP group's MA and reports SF or SF clear events to the Ethernet ring control module. For this MEP group, the action profile must be configured to update the interface device IFF_LINKDOWN flag. OAM updates the IFF_LINKDOWN flag to notify the Ethernet ring control module.

Logical Ring

You can define multiple logical-ring instances on the same physical ring. The logical ring feature currently supports only the physical ring, which means that two adjacent nodes of a ring must be physically connected and the ring must operate on the physical interface, not the VLAN. Multiple ring instances are usually defined with trunk mode ring interfaces.

FDB Flush

When ring protection switching occurs, normally an *FDB flush* is executed. The Ethernet ring control module uses the same mechanism as the STP to trigger the FDB flush. The Ethernet ring control module controls the ring port physical interface's default STP index to execute the FDB flush.

Starting with Junos OS Release 14.2, the FDB flush depends on the RAPS messages received on the both the ports of the ring node.

Traffic Blocking and Forwarding

Ethernet ring control uses the same mechanism as the STP to control forwarding or discarding of user traffic. The Ethernet ring control module sets the ring port physical interface default STP index state to forwarding or discarding in order to control user traffic.

RPL Neighbor Node

Starting with Junos OS Release 14.2, ring protection link neighbor nodes are supported. An RPL neighbor node is adjacent to the RPL and is not the RPL owner. If a node is configured with one interface as the protection-link-end and no protection-link-owner is present in its configuration, the node is an RPL neighbor node.

RAPS Message Blocking and Forwarding

The router or switch treats the ring automatic protection switching (RAPS) message the same as it treats user traffic for forwarding RAPS messages between two ring ports. The ring port physical interface default STP index state also controls forwarding RAPS messages between the two ring ports. Other than forwarding RAPS messages between the two ring ports, as shown in Figure 20, the system also needs to forward the RAPS message between the CPU (Ethernet ring control module) and the ring port. This type of forwarding does not depend on the ring port physical interfaces' STP index state. The RAPS message is always sent by the router or switch through the ring ports, as shown in Figure 21. A RAPS message received from a discarding ring port is sent to the Ethernet ring control module, but is not sent to the other ring port.

Figure 20: Protocol Packets from the Network to the Router





Figure 21: Protocol Packets from the Router or Switch to the Network



Juniper Networks switches and Juniper Networks routers use different methods to achieve these routes.

The switches use forwarding database entries to direct the RAPS messages. The forwarding database entry (keyed by the RAPS multicast address and VLAN) has a composite next hop associated with it—the composite next hop associates the two ring interfaces with the forwarding database entry and uses the split horizon feature to prevent sending the packet out on the interface that it is received on. This is an example of the forwarding database entry relating to the RAPS multicast MAC (a result of the **show ethernet-switching table detail** command):

```

VLAN: v1, Tag: 101, MAC: 01:19:a7:00:00:01, Interface: ERP
Interfaces:      ge-0/0/9.0, ge-0/0/3.0
Type: Static
Action: Mirror
Nexthop index: 1333
  
```

The routers use an implicit filter to achieve ERP routes. Each implicit filter binds to a bridge domain. Therefore, the east ring port control channel and the west ring port control channel of a particular ring instance must be configured to the same bridge domain. For each ring port control channel, a filter term is generated to control RAPS message forwarding. The filter number is the same as the number of bridge domains that contain the ring control channels. If a bridge domain contains control channels from multiple rings, the filter related to this bridge domain will have multiple terms and each term will relate to a control channel. The filter has command parts and control-channel related parts, as follows:

- Common terms:
 - term 1: if [Ethernet type is not OAM Ethernet type (0x8902)]
 { accept packet }
 - term 2: if [source MAC address belongs to this bridge]
 { drop packet, our packet loop through the ring and come back to home }
 - term 3: if [destination is the RAPS PDU multicast address(0x01,0x19,0xa7,0x00,0x00,0x01) AND[ring port STP status is DISCARDING]
 { send to CPU }
- Control channel related terms:
 - if [destination is the RAPS PDU multicast address(0x01,0x19,0xa7,0x00,0x00,0x01) AND[ring port STP status is FORWARDING] AND [Incoming interface IFL equal to control channel IFL]
 { send packet to CPU and send to the other ring port }
 - default term: accept packet.

Dedicated Signaling Control Channel

For each ring port, a dedicated signaling control channel with a dedicated VLAN ID must be configured. In Ethernet ring configuration, only this control logical interface is configured and the underlying physical interface is the physical ring port. Each ring requires that two control physical interfaces be configured. These two logical interfaces must be configured in a bridge domain for routers (or the same VLAN for switches) in order to forward RAPS protocol data units (PDUs) between the two ring control physical interfaces. If the router control channel logical interface is not a trunk port, only control logical interfaces will be configured in ring port configuration. If this router control channel logical interface is a trunk port, in addition to the control channel logical interfaces, a dedicated VLAN ID must be configured for routers. For switches, always specify either a VLAN name or VLAN ID for all links.

RAPS Message Termination

The RAPS message starts from the originating node, travels through the entire ring, and terminates in the originating node unless a failure is present in the ring. The originating node must drop the RAPS message if the source MAC address in the RAPS message belongs to itself. The source MAC address is the node's node ID.

Revertive and Non-revertive Modes

In revertive operation, once the condition causing a switch has cleared, traffic is blocked on the RPL and restored to the working transport entity. In nonrevertive operation, traffic is allowed to use the RPL if it has not failed, even after a switch condition has cleared.

Multiple Rings

The Ethernet ring control module supports multiple rings in each node (two logical interfaces are part of each ring). The ring control module also supports the interconnection of multiple rings. Interconnection of two rings means that two rings might share the same link or share the same node. Ring interconnection is supported only using non-virtual-channel mode. Ring interconnection using virtual channel mode is not supported.

Node ID

For each node in the ring, a unique *node ID* identifies each node. The node ID is the node's MAC address.

For routers only, you can configure this node ID when configuring the ring on the node or automatically select an ID like STP does. In most cases, you will not configure this and the router will select a node ID, like STP does. It should be the manufacturing MAC address. The ring node ID should not be changed, even if you change the manufacturing MAC address. Any MAC address can be used if you make sure each node in the ring has a different node ID. The node ID on switches is selected automatically and is not configurable.

Ring ID

The ring ID is used to determine the value of the last octet of the MAC destination address field of the RAPS protocol data units (PDUs) generated by the ERP control process. The ring ID is also used to discard any RAPS PDU, received by this ERP control process with a non-matching ring ID. Ring ID values 1 through 239 are supported.

Bridge Domains with the Ring Port (MX Series Routers Only)

On the routers, the protection group is seen as an abstract logical port that can be configured to any bridge domain. Therefore, if you configure one ring port or its logical interface in a bridge domain, you must configure the other related ring port or its logical interface to the same bridge domain. The bridge domain that includes the ring port acts as any other bridge domain and supports the IRB Layer 3 interface.

Wait-to-Block Timer

The RPL owner node uses a delay timer before initiating an RPL block in revertive mode of operation or before reverting to IDLE state after clearing manual commands. The Wait-to-Block (WTB) timer is used when clearing **force switch** and **manual switch** commands. As multiple **force switch** commands are allowed to coexist in an Ethernet ring, the WTB timer ensures that clearing of a single **force switch** command does not trigger the re-blocking of the RPL. When clearing a **manual switch** command, the WTB timer prevents the formation of a closed loop due to a possible timing anomaly where

the RPL Owner Node receives an outdated remote **manual switch** request during the recovery process.

When recovering from a **manual switch** command, the delay timer must be long enough to receive any latent remote **force switch**, signal failure, or **manual switch** commands. This delay timer is called the WTB timer and is defined to be 5 seconds longer than the guard timer. This delay timer is activated on the RPL Owner Node. When the WTB timer expires, the RPL Owner Node initiates the reversion process by transmitting an RAPS (NR, RB) message. The WTB timer is deactivated when any higher-priority request preempts it.

Adding and Removing a Node

Starting with Junos OS Release 14.2, you can add or remove a node between two nodes in an Ethernet ring. Nodes are added or removed using the **force switch** command.

Related Documentation

- [Ethernet Ring Protection Switching Overview on page 2055](#)
- *Configuring Ethernet Ring Protection Switching*
- *Example: Ethernet Ring Protection Switching Configuration on MX Routers*
- *Ethernet Interfaces Feature Guide for Routing Devices*
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- [Configuring Ethernet Ring Protection Switching \(CLI Procedure\) on page 2063](#)

Configuring Ethernet Ring Protection Switching (CLI Procedure)

You can configure Ethernet ring protection switching (ERPS) on connected switches to prevent fatal loops from disrupting a network. ERPS is similar to spanning-tree protocols, but ERPS is more efficient than spanning-tree protocols because it is customized for ring topologies. You must configure at least three switches to form a ring. One of the links, called the ring protection link (RPL) end interface, is blocked until another link fails—at this time the RPL link is unblocked, ensuring connectivity.



NOTE: Ethernet OAM connectivity fault management (CFM) can be used with ERPS to detect link faults faster in some cases. See *Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure)*.

The time needed for switchover to the ERPS link is affected by three settings—link failure detection time, the number of nodes in the ring, and the time it takes to unblock the RPL after a failure is detected.



NOTE: Do not configure redundant trunk groups on ERPS interfaces. You can configure VSTP on ERPS interfaces if the VSTP uses a VLAN that is not part of the ERPS control VLAN or data channel VLANs. The total number of ERPS and VSTP or MSTP instances is limited to 253.

Before you begin:

- Configure a VLAN to act as a control channel for ERPS. Two interfaces (east and west) on each switch in the ring must be associated with the control VLAN. See *Configuring VLANs for EX Series Switches (CLI Procedure)*.
- The interfaces on the ERPS control channel are usually (but not required to be) configured as trunk ports. See *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*. Note that if one switch has trunk ports as the ERPS control interfaces, the same must be true of all switches on the ring (the ERPS control interfaces must also be trunk ports).
- Data channels are optional on the ERPS link. If you plan to use them, configure a VLAN for each data channel. If you have multiple ERPS instances, the control VLANs and data channel VLANs must not overlap.

To configure ERPS:



NOTE: You must configure at least three switches, with only one switch designated as the RPL owner node.

1. Spanning tree protocols and ERPS cannot both be configured on the ring ports, so on each ERPS interface, you must disable any configured spanning tree protocols (such

as STP, RSTP, VSTP, or MSTP). Spanning tree protocols are disabled for individual interfaces in two different ways, depending on which Junos OS version and release is running on the switch. RSTP is enabled in the default configuration, so disabling RSTP is shown here.

For switches without Enhanced Layer 2 (ELS) software support, and switches running ELS software with Junos OS release 15.1 or later, use this command to disable RSTP on the individual ERPS interfaces:

```
[edit protocols]
user@switch# set rstp interface interface-name disable
```

For switches running Enhanced Layer 2 (ELS) software with Junos OS releases prior to 15.1, you disable spanning tree protocols on individual interfaces by deleting that configuration item. Use this command to delete the RSTP configuration item on the individual ERPS interfaces:

```
[edit protocols]
user@switch# delete rstp interface interface-name
```

2. Create a node ring on each switch:

```
[edit protocols]
user@switch# set protection-group ethernet-ring ring-name
```

3. Configure a control VLAN for the node ring:

```
[edit protocols protection-group ethernet-ring ring name]
user@switch# set control-vlan vlan-name-or-vlan-id
```

4. Configure the east and west interfaces of the node ring with the control-channel interface.

```
[edit protocols protection-group ethernet-ring ring-name]
user@switch# set east-interface control-channel control-channel-name
user@switch# set west-interface control-channel control-channel-name
```

For switches with ELS support, additionally associate the east and west interfaces with the control VLAN:

```
[edit protocols protection-group ethernet-ring ring-name]
user@switch# set east-interface control-channel vlan vlan-name-or-vlan-id
user@switch# set west-interface control-channel vlan vlan-name-or-vlan-id
```

5. In addition, configure either the east interface or the west interface (but not both) as a link end. For example, configure the east interface:

```
[edit protocols protection-group ethernet-ring ring-name]
user@switch# set east-interface ring-protection-link-end
```

6. Configure only one switch as the RPL owner node:

```
[edit protocols protection-group ethernet-ring ring-name]
user@switch# set ring-protection-link-owner
```

7. The restore interval is the time the RPL owner node waits after the last ring automatic protection switching (RAPS) signal failure (SF) event has been cleared, to see if any further RAPS events occur. During this time interval, the RPL owner continues to process RAPS packets, and the ring remains in protection state with the RPL link unblocked. When this interval expires, if no further RAPS SF events have been reported,

the RPL owner reverts the protection switching, blocks the RPL link, and returns the protection ring to idle state. Optionally, configure a local restore interval for the ERPS ring on each switch:

```
[edit protocols protection-group ethernet-ring ring-name]
user@switch# set restore-interval restore-interval-value
```



NOTE: The restore interval can also be set globally to apply to any ERPS rings configured on the switch. Local per-ring settings take priority over global settings.

8. The guard interval prevents ring nodes from receiving outdated RAPS messages. Optionally, configure the guard interval on each switch:

```
[edit protocols protection-group ethernet-ring ring name]
user@switch# set guard-interval guard-interval-value
```



NOTE: The guard interval can also be set globally to apply to any ERPS rings configured on the switch. Local per-ring settings take priority over global settings.

9. Global restore and guard interval settings are used when no local settings are configured. If these intervals are not configured globally or locally, the default values apply. Optionally configure global interval settings on the switch to apply to all rings that do not have a corresponding interval configured locally for the ring:

- restore interval:

```
[edit protocols protection-group]
user@switch# set restore-interval restore-interval-value
```

- guard interval:

```
[edit protocols protection-group]
user@switch# set guard-interval guard-interval-value
```



NOTE: You can also configure other global settings, such as ERP traceoptions (file, page size, file size, flag name).

10. Optionally, configure VLANs for data channels on the ERPS link:

```
[edit protocols protection-group ethernet-ring ring name]
user@switch# set data-channel vlan-name
```

Related Documentation

- [Example: Configuring Ethernet Ring Protection Switching on EX Series Switches](#)
- [Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS on page 2066](#)
- [Ethernet Ring Protection Switching Overview on page 2055](#)

- [Understanding Ethernet Ring Protection Switching Functionality on page 2056](#)

Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS

You can configure Ethernet ring protection switching (ERPS) on connected EX Series or QFX Series switches to prevent fatal loops from disrupting a network. ERPS is similar to the Spanning Tree Protocol, but ERPS is more efficient because it is customized for ring topologies. You must connect and configure at least three switches to form a ring.

This example shows how to configure Ethernet ring protection switching on four switches with ELS support, connected to one another on a dedicated link in a ring topology. You can include different types of switches in an ERPS ring, including those with and without ELS support. If any of your EX Series switches runs software that does not support ELS, use these configuration directions: *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 41](#).

- [Requirements on page 2066](#)
- [Overview and Topology on page 2067](#)
- [Configuration on page 2068](#)
- [Verification on page 2081](#)

Requirements

This example uses the following hardware and software components:

- Four connected EX Series switches or QFX Series switches that support the Enhanced Layer 2 Software (ELS) to function as nodes in the ring topology. You could use any of these QFX Series switches: QFX5100, QFX5200, and QFX10000. This configuration also applies to EX Series switches that support the Enhanced Layer 2 Software (ELS) configuration style that runs on EX4300 and EX4600 switches.
- Junos OS Release 13.2X50-D10 or later for EX Series switches.
- Junos OS Release 14.1X53-D10 or later for QFX5100 switches. Junos OS Release 15.1X53-D30 or later for QFX5200, and QFX10000 switches.

Before you begin, be sure you have:

- Configured two trunk interfaces on each of the four switches. See [Table 182](#) for a list of the interface names used in this example.
- Configured a VLAN (with name **erp-control-vlan-1** and ID **100**) on all four switches and associated two network interfaces from each of the four switches with the VLAN. See *Configuring VLANs for the QFX Series OR Configuring VLANs for EX Series Switches (CLI Procedure)*. See [Table 182](#) for a list of the interface names used in this example.

- Configured two more VLANs (one with name **erp-data-1** and vlan ID **101** and a second vlan with the name **erp-data-2** and vlan ID **102**) on all four switches and associated both the east and west interfaces on each switch.

Overview and Topology

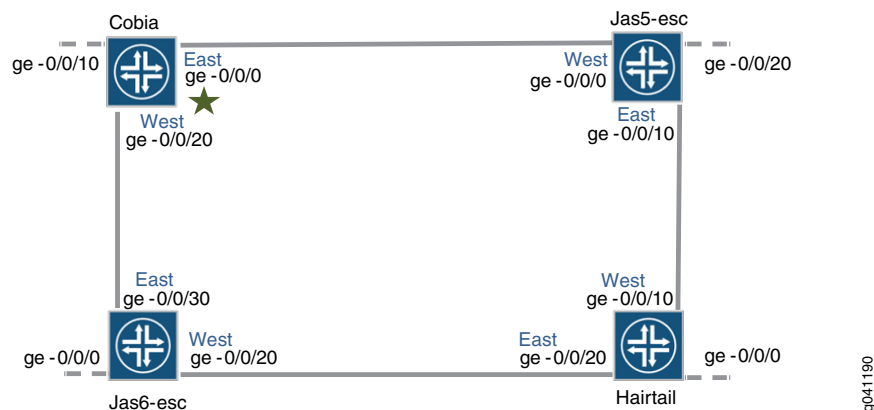
ERPS uses a dedicated physical link, including a control VLAN for trunk ports, between all of the switches to protect the active links. ERPS VLANs are all located on this link and are also blocked by default. When traffic between the switches is flowing with no problems, the active links take care of all traffic. Only if an error occurs on one of the data links would the ERPS control channel take over and start forwarding traffic.



NOTE: Trunk ports on switches use a VLAN to create individual control channels for ERPS. When multiple ERPS instances are configured for a ring, there are multiple sets of ring protection links (RPLs) and RPL owners on the ERPS link, and a different channel is blocked for each instance. Nontrunk ports use the physical link as the control channel and protocol data units (PDUs) are untagged, with no VLAN information in the packet.

This example creates one protection ring (called a node ring) named **erp1** on four switches connected in a ring by trunk ports as shown in [Figure 22](#). Because the links are trunk ports, VLAN 100 is used for **erp1** traffic. The east interface of each switch is connected with the west interface of an adjacent switch. Cobia is the RPL owner, with interface **ge-0/0/0** configured as an RPL end interface. The interface **ge-0/0/0** of **Jas5-esc** is configured as the RPL neighbor interface. In the idle state, the RPL end blocks the control VLAN and data channel VLAN for this particular ERP instance—the blocked port on Cobia is marked with a star in [Figure 22](#).

Figure 22: Ethernet Ring Protection Switching Example



In this example, we configure the four switches with the interfaces indicated in both [Figure 22](#) and [Table 182](#).

Table 182: Components to Configure for This Example

Interfaces	Cobia	Jas5-esc	Jas6-esc	Hairtail
East	ge-0/0/0	ge-0/0/10	ge-0/0/30	ge-0/0/20
West	ge-0/0/20	ge-0/0/0	ge-0/0/20	ge-0/0/10
Third	ge-0/0/10	ge-0/0/20	ge-0/0/0	ge-0/0/0

Configuration

- [Configuring ERPS on Cobia, the RPL Owner Node on page 2068](#)
- [Configuring ERPS on Jas5-esc on page 2071](#)
- [Configuring ERPS on Hairtail on page 2075](#)
- [Configuring ERPS on Jas6-esc on page 2078](#)

Configuring ERPS on Cobia, the RPL Owner Node

CLI Quick Configuration

To quickly configure Cobia, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

ERPS cannot be configured on an interface if any spanning-tree protocol is configured. (RSTP is configured by default.) Therefore, in this example, RSTP is disabled on each ring port before configuring ERPS. Spanning-tree protocols are disabled two different ways, depending on which version of the Junos OS you are running. Therefore, the first two statements in this example vary: Junos release 15.1 or later uses one command to turn off RSTP and Junos releases prior to 15.1 uses another command.

```

Junos OS release 15.1 or later: set protocols rstp interface ge-0/0/0 disable
Junos OS release 15.1 or later: set protocols rstp interface ge-0/0/20 disable
Junos OS release prior to 15.1: delete rstp interface ge-0/0/0
Junos OS release prior to 15.1: delete rstp interface ge-0/0/20
set protocols protection-group ethernet-ring erp1
set protocols protection-group ethernet-ring erp1 ring-protection-link-owner
set protocols protection-group ethernet-ring erp1 data-channel 101
set protocols protection-group ethernet-ring erp1 data-channel 102
set protocols protection-group ethernet-ring erp1 control-vlan 100
set protocols protection-group ethernet-ring erp1 east-interface control-channel
    ge-0/0/0.0
set protocols protection-group ethernet-ring erp1 east-interface ring-protection-link-end
set protocols protection-group ethernet-ring erp1 east-interface control-channel
    ge-0/0/20.0
set protocols protection-group ethernet-ring erp1 west-interface control-channel vlan
    100
set protocols protection-group ethernet-ring erp1 east-interface control-channel vlan
    100

```

**Step-by-Step
Procedure**

To configure ERPS on Cobia:

1. Disable any spanning-tree protocol currently configured on the ERPS interfaces. RSTP, VSTP, and MSTP are all available spanning-tree protocols. RSTP is enabled in the default configuration, so this example shows disabling RSTP. Spanning-tree protocols are disabled two different ways, depending on which version of the Junos OS you are running.

If you are running Junos release 15.1 or later, disable any spanning-tree protocol with these commands. To disable RSTP:

```
[edit protocols]
user@switch# set rstp interface ge-0/0/0 disable
user@switch# set rstp interface ge-0/0/20 disable
```

If you are running a Junos release prior to 15.1, disable any spanning-tree protocol with these commands. To disable RSTP:

```
[edit protocols]
user@switch# delete rstp interface ge-0/0/0
user@switch# delete rstp interface ge-0/0/20
```

2. Create a node ring named erp1:

```
[edit protocols]
user@switch# set protection-group ethernet-ring erp1
```

3. Designate Cobia as the RPL owner node:

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set ring-protection-link-owner
```

4. Configure the VLANs 101 and 102 as data channels:

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set data-channel 101
user@switch# set data-channel 102
```

5. Configure the control vlan 100 for this ERP instance on the trunk interface:

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set control-vlan 100
```

6. Configure the east interface of the node ring erp1 with control channel ge-0/0/0.0 and indicate that this particular ring protection link ends here:

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set east-interface control-channel ge-0/0/0.0
user@switch# set east-interface ring-protection-link-end
```

7. Configure the west interface of the node ring erp1 with control channel ge-0/0/20.0 :

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set west-interface control-channel ge-0/0/20.0
```

8. Every ring instance on a trunk port has one control VLAN in which ERP packets traverse. The control VLAN also controls data VLANs, if any are configured. Assign 100 as the control VLAN on both interfaces:

```
[edit protocols protection-group ethernet-ring erp1]
```

```
user@switch# set west-interface control-channel vlan 100
user@switch# set east-interface control-channel vlan 100
```

Results In configuration mode, check your ERPS configuration by entering the **show configuration protocols** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@switch# show configuration protocols
  protection-group {
    ethernet-ring ERP1 {
      restore-interval 5;
      east-interface {
        control-channel {
          ge-0/0/13.0;
          vlan 4000;
        }
      }
    }
    west-interface {
      control-channel {
        ge-0/0/15.0;
        vlan 4000;
      }
    }
    control-vlan 4000;
    data-channel {
      vlan 2052;
    }
  }
}

user@switch# show configuration protocols | display set
set protocols protection-group ethernet-ring ERP1 restore-interval 5
set protocols protection-group ethernet-ring ERP1 east-interface control-channel
ge-0/0/13.0
set protocols protection-group ethernet-ring ERP1 east-interface control-channel vlan
4000
set protocols protection-group ethernet-ring ERP1 west-interface control-channel
ge-0/0/15.0
set protocols protection-group ethernet-ring ERP1 west-interface control-channel vlan
4000
set protocols protection-group ethernet-ring ERP1 control-vlan 4000
set protocols protection-group ethernet-ring ERP1 data-channel vlan 2052
```

In configuration mode, check your VLAN configuration by entering the **show vlans** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

In configuration mode, check your interface configurations by entering the **show interfaces** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
show vlans
  erp-control-vlan-1 {
    vlan-id 100;
```

```

}
erp-data-1 {
  vlan-id 101;
}
erp-data-2 {
  vlan-id 102;
}
user@switch# show interfaces
ge-0/0/0 {
  unit 0 {
    family ethernet-switching {
      interface-mode trunk;
      vlan {
        members [ 100 101 102 ];
      }
    }
  }
}
ge-0/0/10 {
  unit 0 {
    family ethernet-switching {
      interface-mode trunk;
      vlan {
        members [ 101 102 ];
      }
    }
  }
}
ge-0/0/20 {
  unit 0 {
    family ethernet-switching {
      interface-mode trunk;
      vlan {
        members [ 100 101 102 ];
      }
    }
  }
}

```

If you are finished configuring the device, enter **commit** in configuration mode.

Configuring ERPS on Jas5-esc

CLI Quick Configuration

To quickly configure Jas5-esc, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

ERPS cannot be configured on an interface if any spanning tree protocol is configured. (RSTP is configured by default.) Therefore, in this example, RSTP is disabled on each ring port before configuring ERPS. Spanning tree is disabled two different ways, depending on which version of the Junos OS you are running. Therefore, the first two statements will vary: Junos release 15.1 or later uses one command to turn off RSTP and Junos releases prior to 15.1 uses another command.

Junos OS release 15.1 or later: **set protocols rstp interface ge-0/0/10 disable**

```
Junos OS release 15.1 or later: set protocols rstp interface ge-0/0/0 disable
Junos OS release prior to 15.1: delete rstp interface ge-0/0/10
Junos OS release prior to 15.1: delete rstp interface ge-0/0/0
set protocols protection-group ethernet-ring erp1
set protocols protection-group ethernet-ring erp1 data-channel 101
set protocols protection-group ethernet-ring erp1 data-channel 102
set protocols protection-group ethernet-ring erp1 control-vlan 100
set protocols protection-group ethernet-ring erp1 east-interface control-channel
  ge-0/0/0.0
set protocols protection-group ethernet-ring erp1 west-interface control-channel
  ge-0/0/20.0
set protocols protection-group ethernet-ring erp1 west-interface control-channel vlan
  100
set protocols protection-group ethernet-ring erp1 west-interface control-channel vlan
  100
```

Step-by-Step Procedure

To configure ERPS on Jas5-esc:

1. Disable any spanning-tree protocol currently configured on the ERPS interfaces. RSTP, VSTP, and MSTP are all available spanning-tree protocols. RSTP is enabled in the default configuration, so this example shows disabling RSTP. Spanning-tree protocols are disabled two different ways, depending on which version of the Junos OS you are running.

If you are running Junos release 15.1 or later, disable any spanning-tree protocol with these commands. To disable RSTP:

```
[edit protocols]
user@switch# set rstp interface ge-0/0/10 disable
user@switch# set rstp interface ge-0/0/0 disable
```

If you are running a Junos release prior to 15.1, disable any version of spanning-tree protocol with these commands. To disable RSTP:

```
[edit protocols]
user@switch# delete rstp interface ge-0/0/10
user@switch# delete rstp interface ge-0/0/0
```

2. Create a node ring named erp1:

```
[edit protocols]
user@switch# set protection-group ethernet-ring erp1
```
3. Configure two data channels named erp-data-1 and erp-data-2 to define a set of VLAN IDs that belong to a ring instance.

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set data-channel vlan 101
user@switch# set data-channel vlan 102
```
4. Configure a control VLAN with ID 100 for the node ring erp1:

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set control-vlan 100
```
5. Configure the east interface of the node ring erp1 with the control channel ge-0/0/10.0:

```
[edit protocols protection-group ethernet-ring erp1]
```



```
user@switch# set east-interface control-channel ge-0/0/10.0
```

6. Configure the west interface of the node ring erp1 with the control channel ge-0/0/0.0 vlan 100:

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set west-interface control-channel ge-0/0/0.0
```

7. Every ring instance on a trunk port has one control VLAN in which ERP packets traverse. The control VLAN also controls data VLANs, if any are configured. Assign vlan # 100 as the control VLAN:

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set west-interface control-channel vlan 100
user@switch# set east-interface control-channel vlan 100
```

Results In configuration mode, check your ERPS configuration by entering the **show configuration protocols** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@switch# show configuration protocols
  protection-group {
    ethernet-ring ERP1 {
      restore-interval 5;
      east-interface {
        control-channel {
          ge-0/0/13.0;
          vlan 4000;
        }
      }
      west-interface {
        control-channel {
          ge-0/0/15.0;
          vlan 4000;
        }
      }
      control-vlan 4000;
      data-channel {
        vlan 2052;
      }
    }
  }
user@switch# show configuration protocols | display set
set protocols protection-group ethernet-ring ERP1 restore-interval 5
set protocols protection-group ethernet-ring ERP1 east-interface control-channel
ge-0/0/13.0
set protocols protection-group ethernet-ring ERP1 east-interface control-channel vlan
4000
set protocols protection-group ethernet-ring ERP1 west-interface control-channel
ge-0/0/15.0
set protocols protection-group ethernet-ring ERP1 west-interface control-channel vlan
4000
set protocols protection-group ethernet-ring ERP1 control-vlan 4000
set protocols protection-group ethernet-ring ERP1 data-channel vlan 2052
```

In configuration mode, check your VLAN configuration by entering the **show vlans** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

In configuration mode, check your interface configurations by entering the **show interfaces** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
show vlans
  erp-control-vlan-1 {
    vlan-id 100;
  }
  erp-data-1 {
    vlan-id 101;
  }
  erp-data-2 {
    vlan-id 102;
  }
user@switch# show interfaces
  ge-0/0/0 {
    unit 0 {
      family ethernet-switching {
        interface-mode trunk;
        vlan {
          members [ 100 101 102 ];
        }
      }
    }
  }
  ge-0/0/10 {
    unit 0 {
      family ethernet-switching {
        interface-mode trunk;
        vlan {
          members [ 100 101 102 ];
        }
      }
    }
  }
  ge-0/0/20 {
    unit 0 {
      family ethernet-switching {
        interface-mode trunk;
        vlan {
          members [ 101 102 ];
        }
      }
    }
  }
}
```

If you are finished configuring the device, enter **commit** in configuration mode.

Configuring ERPS on Hairtail

CLI Quick Configuration To quickly configure Hairtail, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

ERPS cannot be configured on an interface if any spanning tree protocol is configured. (RSTP is configured by default.) Therefore, in this example, RSTP is disabled on each ring port before configuring ERPS. Spanning tree is disabled two different ways, depending on which version of the Junos OS you are running. Therefore, the first two statements will vary: Junos release 15.1 or later uses one command to turn off RSTP and Junos releases prior to 15.1 uses another command.

```
Junos OS release 15.1 or later: set protocols rstp interface ge-0/0/10 disable
Junos OS release 15.1 or later: set protocols rstp interface ge-0/0/20 disable
Junos OS release prior to 15.1: delete rstp interface ge-0/0/10
Junos OS release prior to 15.1: delete rstp interface ge-0/0/20
set protocols protection-group ethernet-ring erp1
set protocols protection-group ethernet-ring erp1 data-channel 101
set protocols protection-group ethernet-ring erp1 data-channel 102
set protocols protection-group ethernet-ring erp1 control-vlan 100
Set protocols protection-group ethernet-ring erp1 east-interface control-channel
ge-0/0/0.0
set protocols protection-group ethernet-ring erp1 east-interface control-channel
ge-0/0/20.0
set protocols protection-group ethernet-ring erp1 west-interface control-channel vlan
100
set protocols protection-group ethernet-ring erp1 east-interface control-channel vlan
100
```

Step-by-Step Procedure To configure ERPS on Hairtail:

1. Disable any spanning-tree protocol currently configured on the ERPS interfaces. RSTP, VSTP, and MSTP are all available spanning-tree protocols. RSTP is enabled in the default configuration, so this example shows disabling RSTP. Spanning-tree protocols are disabled two different ways, depending on which version of the Junos OS you are running.

If you are running Junos release 15.1 or later, disable any spanning-tree protocol with these commands. To disable RSTP:

```
[edit protocols]
user@switch# set rstp interface ge-0/0/10 disable
user@switch# set rstp interface ge-0/0/20 disable
```

If you are running a Junos release prior to 15.1, disable any spanning-tree protocol with these commands. To disable RSTP:

```
[edit protocols]
user@switch# delete rstp interface ge-0/0/10
user@switch# delete rstp interface ge-0/0/20
```

2. Create a node ring named erp1:

```
[edit protocols]
```

- ```
user@switch# set protection-group ethernet-ring erp1
```
3. Configure the control vlan 100 for the node ring erp1:  

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set control-vlan 100
```
  4. Configure two data channels numbered 101 and 102 to define a set of VLAN IDs that belong to a ring instance:  

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set data-channel vlan 101
user@switch# set data-channel vlan 102
```
  5. Configure the east interface of the node ring erp1 with the control channel ge-0/0/20.0:  

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set east-interface control-channel ge-0/0/20.0
```
  6. Configure the west interface of the node ring erp1 with the control channel ge-0/0/10.0:  

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set west-interface control-channel ge-0/0/10.0
```
  7. Every ring instance on a trunk port has one control VLAN in which ERP packets traverse. The control VLAN also controls data VLANs, if any are configured. Assign 100 as the control VLAN:  

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set west-interface control-channel vlan 100
user@switch# set east-interface control-channel vlan 100
```

**Results** In configuration mode, check your ERPS configuration by entering the **show configuration protocols** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@switch# show configuration protocols
 protection-group {
 ethernet-ring ERP1 {
 restore-interval 5;
 east-interface {
 control-channel {
 ge-0/0/13.0;
 vlan 4000;
 }
 }
 west-interface {
 control-channel {
 ge-0/0/15.0;
 vlan 4000;
 }
 }
 control-vlan 4000;
 data-channel {
 vlan 2052;
 }
 }
 }
```

```

 }
 }
}
user@switch# show configuration protocols | display set
set protocols protection-group ethernet-ring ERP1 restore-interval 5
set protocols protection-group ethernet-ring ERP1 east-interface control-channel
ge-0/0/13.0
set protocols protection-group ethernet-ring ERP1 east-interface control-channel vlan
4000
set protocols protection-group ethernet-ring ERP1 west-interface control-channel
ge-0/0/15.0
set protocols protection-group ethernet-ring ERP1 west-interface control-channel vlan
4000
set protocols protection-group ethernet-ring ERP1 control-vlan 4000
set protocols protection-group ethernet-ring ERP1 data-channel vlan 2052

```

In configuration mode, check your VLAN configuration by entering the **show vlans** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

In configuration mode, check your interface configurations by entering the **show interfaces** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```

[edit]
show vlans
 erp-control-vlan-1 {
 vlan-id 100;
 }
 erp-data-1 {
 vlan-id 101;
 }
 erp-data-2 {
 vlan-id 102;
 }
user@switch# show interfaces
ge-0/0/0 {
 unit 0 {
 family ethernet-switching {
 interface-mode trunk;
 vlan {
 members [101 102];
 }
 }
 }
}
ge-0/0/10 {
 unit 0 {
 family ethernet-switching {
 interface-mode trunk;
 vlan {
 members [100 101 102];
 }
 }
 }
}
}

```

```
ge-0/0/20 {
 unit 0 {
```

If you are finished configuring the device, enter **commit** in configuration mode.

### Configuring ERPS on Jas6-esc

---

#### CLI Quick Configuration

To quickly configure Jas6-esc, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

ERPS cannot be configured on an interface if any spanning tree protocol is configured. (RSTP is configured by default.) Therefore, in this example, RSTP is disabled on each ring port before configuring ERPS. Spanning tree is disabled two different ways, depending on which version of the Junos OS you are running. Therefore, the first two statements will vary: Junos release 15.1 or later uses one command to turn off RSTP and Junos releases prior to 15.1 uses another command.

```
Junos OS release 15.1 or later: set protocols rstp interface ge-0/0/30 disable
Junos OS release 15.1 or later: set protocols rstp interface ge-0/0/20 disable
Junos OS release prior to 15.1: delete rstp interface ge-0/0/30
Junos OS release prior to 15.1: delete rstp interface ge-0/0/20
set protocols protection-group ethernet-ring erp1
set protocols protection-group ethernet-ring erp1 data-channel 101
set protocols protection-group ethernet-ring erp1 data-channel 102
set protocols protection-group ethernet-ring erp1 control-vlan 100
set protocols protection-group ethernet-ring erp1 east-interface control-channel
 ge-0/0/30.0
set protocols protection-group ethernet-ring erp1 west-interface control-channel
 ge-0/0/20.0
set protocols protection-group ethernet-ring erp1 west-interface control-channel vlan
 100
set protocols protection-group ethernet-ring erp1 east-interface control-channel vlan
 100
```

#### Step-by-Step Procedure

To configure ERPS on Jas6-esc:

1. Disable any spanning-tree protocol currently configured on the ERPS interfaces. RSTP, VSTP, and MSTP are all available spanning-tree protocols. RSTP is enabled in the default configuration, so this example shows disabling RSTP. Spanning-tree protocols are disabled two different ways, depending on which version of the Junos OS you are running.

If you are running Junos release 15.1 or later, disable any spanning-tree protocol with these commands. To disable RSTP:

```
[edit protocols]
user@switch# set rstp interface ge-0/0/30 disable
user@switch# set rstp interface ge-0/0/20 disable
```

If you are running a Junos release prior to 15.1, disable any spanning-tree protocol with these commands. To disable RSTP:

```
[edit protocols]
```

- ```

user@switch# delete rstp interface ge-0/0/30
user@switch# delete rstp interface ge-0/0/20

```
2. Create a node ring named erp1:


```

[edit protocols]
user@switch# set protection-group ethernet-ring erp1

```
 3. Configure the control vlan 100 for the node ring erp1:


```

[edit protocols protection-group ethernet-ring erp1]
user@switch# set control-vlan 100

```
 4. Configure two data channels numbered 101 and 102 to define VLAN IDs that belong to a ring instance.


```

[edit protocols protection-group ethernet-ring erp1]
user@switch# set data-channel 101
user@switch# set data-channel 102

```
 5. Configure the east interface of the node ring erp1 with the control channel ge-0/0/30.0 :


```

[edit protocols protection-group ethernet-ring erp1]
user@switch# set east-interface control-channel ge-0/0/30.0

```
 6. Configure the west interface of the node ring erp1 with the control channel ge-0/0/20.0:


```

[edit protocols protection-group ethernet-ring erp1]
user@switch# set west-interface control-channel ge-0/0/20.0

```
 7. Every ring instance on a trunk port has one control VLAN in which ERP packets traverse. The control VLAN also controls data VLANs, if any are configured. Assign vlan number 100 as the control VLAN:


```

[edit protocols protection-group ethernet-ring erp1]
user@switch# set west-interface control-channel vlan 100
user@switch# set east-interface control-channel vlan 100

```

Results In configuration mode, check your ERPS configuration by entering the **show configuration protocols** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```

[edit]
user@switch# show configuration protocols
}
  protection-group {
    ethernet-ring ERP1 {
      restore-interval 5;
      east-interface {
        control-channel {
          ge-0/0/13.0;
          vlan 4000;
        }
      }
      west-interface {
        control-channel {
          ge-0/0/15.0;

```

```
        vlan 4000;
    }
}
control-vlan 4000;
data-channel {
    vlan 2052;
}
}
}
user@switch# show configuration protocols | display set
set protocols protection-group ethernet-ring ERP1 restore-interval 5
set protocols protection-group ethernet-ring ERP1 east-interface control-channel
ge-0/0/13.0
set protocols protection-group ethernet-ring ERP1 east-interface control-channel vlan
4000
set protocols protection-group ethernet-ring ERP1 west-interface control-channel
ge-0/0/15.0
set protocols protection-group ethernet-ring ERP1 west-interface control-channel vlan
4000
set protocols protection-group ethernet-ring ERP1 control-vlan 4000
set protocols protection-group ethernet-ring ERP1 data-channel vlan 2052
```

In configuration mode, check your VLAN configuration by entering the **show vlans** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

In configuration mode, check your interfaces configuration by entering the **show interfaces** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
show vlans
  erp-control-vlan-1 {
    vlan-id 100;
  }
  erp-data-1 {
    vlan-id 101;
  }
  erp-data-2 {
    vlan-id 102;
  }
}
user@switch# show interfaces
ge-0/0/0 {
  unit 0 {
    family ethernet-switching {
      interface-mode trunk;
      vlan {
        members [ 101 102 ];
      }
    }
  }
}
ge-0/0/20 {
  unit 0 {
    family ethernet-switching {
      interface-mode trunk;

```



```

        vlan {
            members [ 100 101 102 ];
        }
    }
}
ge-0/0/30 {
    unit 0 {
        family ethernet-switching {
            interface-mode trunk;
            vlan {
                members [ 100 101 102 ];
            }
        }
    }
}
}

```

Verification

Verify that ERPS is working correctly.

Verifying That ERPS Is Working Correctly

Purpose Verify that ERPS is working on the four EX switches that function as nodes in the ring topology.

Action Check the state of the ring links in the output of the **show protection-group ethernet-ring interface** command. When the ring is configured but not being used (no error exists on the data links), one ERP interface is forwarding traffic and one is discarding traffic. Discarding blocks the ring.

```
user@switch> show protection-group ethernet-ring interface
```

Ethernet ring port parameters for protection group ERP1

Interface	Control Channel	Forward State	Ring Protection Link End	Signal Failure
ge-0/0/13	ge-0/0/13.0	forwarding	No	Clear
	IFF ready			
ge-0/0/15	ge-0/0/15.0	forwarding	No	Clear
	IFF ready			

To find out what has occurred since the last restart, check the RPS statistics for ring-blocked events. **NR** is a No Request ring block, which means that the switch is not blocking either of the two ERP interfaces. **NR-RB** is a No Request Ring Blocked event, which means that the switch is blocking one of its ERP interfaces and sending a packet out to notify the other switches.

```
user@switch> show protection-group ethernet-ring statistics
```

```

Ethernet Ring statistics for PG erp1
RAPS event sent           : 0
RAPS event received       : 1
Local SF happened:        : 0
Remote SF happened:       : 0
NR event happened:        : 0
NR-RB event happened:     : 1

```

Meaning The **show protection-group ethernet-ring interface** command output from the RPL owner node indicates that one interface is forwarding traffic and one is discarding traffic, meaning that the ERP is ready but not active. If at least one interface in the ring is not forwarding, the ring is blocked and therefore ERP is working.

The **show protection-group ethernet-ring statistics** command output indicates that, since the last reboot, both local and remote signal failures have occurred (**Local SF** and **Remote SF**).

The **NR Event** count is 2, indicating that the NR state was entered into twice. **NR** stands for No Request. This means that the switch either originated NR PDUs or received an NR PDU from another switch and stopped blocking the interface to allow ERP to function.

The three **NR-RB** events indicate that on three occasions, this switch either sent out NR-RB PDUs or received NR-RB PDUs from another switch. This occurs when a network problem is resolved and the switch once again blocks the ERP link at one end.

- Related Documentation**
- [Configuring Ethernet Ring Protection Switching \(CLI Procedure\) on page 2063](#)
 - [Ethernet Ring Protection Switching Overview on page 2055](#)
 - [Understanding Ethernet Ring Protection Switching Functionality on page 2056](#)

PART 29

Bridging and VLANs

- [Using Bridging and VLANs on page 2085](#)

CHAPTER 73

Using Bridging and VLANs

- [Overview of Layer 2 Networking on page 2085](#)
- [Understanding Layer 2 Broadcasting on page 2087](#)
- [Layer 2 Learning and Forwarding for VLANs Overview on page 2088](#)
- [Understanding Bridging and VLANs on page 2089](#)
- [Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 2096](#)
- [Example: Setting Up Bridging with Multiple VLANs on page 2114](#)
- [Configuring VLANs on page 2120](#)
- [Configuring the Native VLAN Identifier \(CLI Procedure\) on page 2121](#)
- [Creating a Series of Tagged VLANs on page 2123](#)
- [Understanding Integrated Routing and Bridging on page 2124](#)
- [Example: Configuring Routing Between VLANs on One Switch on page 2126](#)
- [Configuring IRB Interfaces on page 2131](#)
- [Adding a Static MAC Address Entry to the Ethernet Switching Table \(CLI Procedure\) on page 2133](#)
- [Configuring Static ARP Entries on page 2133](#)
- [Troubleshooting Ethernet Switching on page 2134](#)

Overview of Layer 2 Networking

Layer 2, also known as the Data Link Layer, is the second level in the seven-layer OSI reference model for network protocol design. Layer 2 is equivalent to the link layer (the lowest layer) in the TCP/IP network model. Layer 2 is the network layer used to transfer data between adjacent network nodes in a wide area network or between nodes on the same local area network.

A *frame* is a protocol data unit, the smallest unit of bits on a Layer 2 network. Frames are transmitted to and received from devices on the same local area network (LAN). Unlike bits, frames have a defined structure and can be used for error detection, control plane activities and so forth. Not all frames carry user data. The network uses some frames to control the data link itself..

At Layer 2, *unicast* refers to sending frames from one node to a single other node, whereas *multicast* denotes sending traffic from one node to multiple nodes, and *broadcasting*

refers to the transmission of frames to all nodes in a network. A *broadcast domain* is a logical division of a network in which all nodes of that network can be reached at Layer 2 by a broadcast.

Segments of a LAN can be linked at the frame level using *bridges*. Bridging creates separate broadcast domains on the LAN, creating VLANs, which are independent logical networks that group together related devices into separate network segments. The grouping of devices on a VLAN is independent of where the devices are physically located in the LAN. Without bridging and VLANs, all devices on the Ethernet LAN are in a single broadcast domain, and all the devices detect all the packets on the LAN.

Forwarding is the relaying of packets from one network segment to another by nodes in the network. On a VLAN, a frame whose origin and destination are in the same VLAN are forwarded only within the local VLAN. A network segment is a portion of a computer network wherein every device communicates using the same physical layer.

Layer 2 contains two sublayers:

- Logical link control (LLC) sublayer, which is responsible for managing communications links and handling frame traffic.
- Media access control (MAC) sublayer, which governs protocol access to the physical network medium. By using the MAC addresses that are assigned to all ports on a switch, multiple devices on the same physical link can uniquely identify one another.

The ports, or interfaces, on a switch operate in either access mode, tagged-access, or trunk mode:

- *Access mode* ports connect to a network device such as a desktop computer, an IP telephone, a printer, a file server, or a security camera. The port itself belongs to a single VLAN. The frames transmitted over an access interface are normal Ethernet frames. By default, all ports on a switch are in access mode.
- *Tagged-Access mode* ports connect to a network device such as a desktop computer, an IP telephone, a printer, a file server, or a security camera. The port itself belongs to a single VLAN. The frames transmitted over an access interface are normal Ethernet frames. By default, all ports on a switch are in access mode. Tagged-access mode accommodates cloud computing, specifically scenarios including virtual machines or virtual computers. Because several virtual computers can be included on one physical server, the packets generated by one server can contain an aggregation of VLAN packets from different virtual machines on that server. To accommodate this situation, tagged-access mode reflects packets back to the physical server on the same downstream port when the destination address of the packet was learned on that downstream port. Packets are also reflected back to the physical server on the downstream port when the destination has not yet been learned. Therefore, the third interface mode, tagged access, has some characteristics of access mode and some characteristics of trunk mode:
- *Trunk mode* ports handle traffic for multiple VLANs, multiplexing the traffic for all those VLANs over the same physical connection. Trunk interfaces are generally used to interconnect switches to other devices or switches.

With native VLAN configured, frames that do not carry VLAN tags are sent over the trunk interface. If you have a situation where packets pass from a device to a switch in access mode, and you want to then send those packets from the switch over a trunk port, use native VLAN mode. Configure the single VLAN on the switch's port (which is in access mode) as a native VLAN. The switch's trunk port will then treat those frames differently than the other tagged packets. For example, if a trunk port has three VLANs, 10, 20, and 30, assigned to it with VLAN 10 being the native VLAN, frames on VLAN 10 that leave the trunk port on the other end have no 802.1Q header (tag). There is another native VLAN option. You can have the switch add and remove tags for untagged packets. To do this, you first configure the single VLAN as a native VLAN on a port attached to a device on the edge. Then, assign a VLAN ID tag to the single native VLAN on the port connected to a device. Last, add the VLAN ID to the trunk port. Now, when the switch receives the untagged packet, it adds the ID you specified and sends and receives the tagged packets on the trunk port configured to accept that VLAN.

Including the sublayers, Layer 2 on the QFX Series supports the following functionality:

- Unicast, multicast, and broadcast traffic.
- Bridging.
- VLAN 802.1Q—Also known as *VLAN tagging*, this protocol allows multiple bridged networks to transparently share the same physical network link by adding VLAN tags to an Ethernet frame.
- Extension of Layer 2 VLANs across multiple switches using Spanning Tree Protocol (STP) prevents looping across the network.
- *MAC learning*, including per-VLAN MAC learning and Layer 2 learning suppression—This process obtains the MAC addresses of all the nodes on a network
- Link aggregation—This process groups of Ethernet interfaces at the physical layer to form a single link layer interface, also known as a *link aggregation group (LAG)* or LAG bundle
- Storm control on the physical port for unicast, multicast, and broadcast
- STP support, including 802.1d, RSTP, MSTP, and Root Guard

**Related
Documentation**

- [Understanding Bridging and VLANs on page 2089](#)
- *Understanding Bridging and VLANs*

Understanding Layer 2 Broadcasting

In a Layer 2 network, *broadcasting* refers to sending traffic to all nodes on a network.

Layer 2 broadcast traffic stays within a local area network (LAN) boundary; known as the *broadcast domain*. Layer 2 broadcast traffic is sent to the broadcast domain using a MAC address of FF:FF:FF:FF:FF:FF. Every device in the broadcast domain recognizes this MAC address and passes the broadcast traffic on to other devices in the broadcast

domain, if applicable. Broadcasting can be compared to unicasting (sending traffic to a single node) or multicasting (delivering traffic to a group of nodes simultaneously).

Layer 3 broadcast traffic, however, is sent to all devices in a network using a broadcast network address. For example, if your network address is 192.0.0.0, the broadcast network address is 192.255.255.255. In this case, only devices that belong to the 192.0.0.0 network receive the Layer 3 broadcast traffic. Devices that do not belong to this network drop the traffic.

Broadcasting is used in the following situations:

- Address Resolution Protocol (ARP) uses broadcasting to map MAC addresses to IP addresses. ARP dynamically binds the IP address (the logical address) to the correct MAC address. Before IP unicast packets can be sent, ARP discovers the MAC address used by the Ethernet interface where the IP address is configured.
- Dynamic Host Configuration Protocol (DHCP) uses broadcasting to dynamically assign IP addresses to hosts on a network segment or subnet.
- Routing protocols use broadcasting to advertise routes.

Excessive broadcast traffic can sometimes create a broadcast storm. A broadcast storm occurs when messages are broadcast on a network and each message prompts a receiving node to respond by broadcasting its own messages on the network. This, in turn, prompts further responses that create a snowball effect. The LAN is suddenly flooded with packets, creating unnecessary traffic that leads to poor network performance or even a complete loss of network service.

Related Documentation

- [Overview of Layer 2 Networking on page 2085](#)
- [Understanding Storm Control on page 6045](#)
- [Understanding Bridging and VLANs](#)
- [Understanding Bridging and VLANs on page 2089](#)

Layer 2 Learning and Forwarding for VLANs Overview

When you configure a VLAN, Layer 2 address learning is enabled by default. The VLAN learns unicast media access control (MAC) addresses to avoid flooding the packets to all the ports in the VLAN. Each VLAN creates a source MAC entry in its source and destination MAC tables for each source MAC address learned from packets received on the ports that belong to the VLAN.



NOTE: Traffic is not flooded back onto the interface on which it was received. However, because this “split horizon” occurs at a late stage, the packet statistics displayed by commands such as `show interfaces queue` will include flood traffic.

You can optionally disable MAC learning either for the entire device or for a specific VLAN or logical interface. You can also configure the following Layer 2 learning and forwarding properties:

- Static MAC entries for logical interfaces only
- Limit to the number of MAC addresses learned from a specific logical interface or from all the logical interfaces in a VLAN
- Size of the MAC address table for the VLAN
- MAC accounting for a VLAN

Related Documentation • *Layer 2 Learning and Forwarding Overview*

Understanding Bridging and VLANs

Network switches use Layer 2 bridging protocols to discover the topology of their LAN and to forward traffic toward destinations on the LAN. This topic explains the following concepts regarding bridging and VLANs:

- [History of VLANs on page 2089](#)
- [How Bridging of VLAN Traffic Works on page 2090](#)
- [Packets Are Either Tagged or Untagged on page 2091](#)
- [Switch Interface Modes—Access, Trunk, or Tagged Access on page 2091](#)
- [Additional Advantages of Using VLANs on page 2093](#)
- [Maximum VLANs and VLAN Members Per Switch on page 2094](#)
- [A Default VLAN Is Configured on Most Switches on page 2095](#)
- [Assigning Traffic to VLANs on page 2095](#)
- [Forwarding VLAN Traffic on page 2096](#)
- [VLANs Communicate with Integrated Routing and Bridging Interfaces or Routed VLAN Interfaces on page 2096](#)

History of VLANs

Ethernet LANs were originally designed for small, simple networks that primarily carried text. However, over time, the type of data carried by LANs grew to include voice, graphics, and video. This more complex data, when combined with the ever-increasing speed of transmission, eventually became too much of a load for the original Ethernet LAN design. Multiple packet collisions were significantly slowing down the larger LANs.

The IEEE 802.1D-2004 standard helped evolve Ethernet LANs to cope with the higher data and transmission requirements by defining the concept of *transparent bridging* (generally called simply *bridging*). Bridging divides a single physical LAN (now called a single *broadcast domain*) into two or more virtual LANs, or VLANs. Each VLAN is a collection of some of the LAN nodes grouped together to form individual broadcast domains.

When VLANs are grouped logically by function or organization, a significant percentage of data traffic stays within the VLAN. This relieves the load on the LAN because all traffic no longer has to be forwarded to all nodes on the LAN. A VLAN first transmits packets within the VLAN, thereby reducing the number of packets transmitted on the entire LAN. Because packets whose origin and destination are in the same VLAN are forwarded only within the local VLAN, packets that are not destined for the local VLAN are the only ones forwarded to other broadcast domains. This way, bridging and VLANs limit the amount of traffic flowing across the entire LAN by reducing the possible number of collisions and packet retransmissions within VLANs and on the LAN as a whole.

How Bridging of VLAN Traffic Works

Because the objective of the IEEE 802.1D-2004 standard was to reduce traffic and therefore reduce potential transmission collisions for Ethernet, a system was implemented to reuse information. Instead of having a switch go through a location process every time a frame is sent to a node, the transparent bridging protocol allows a switch to record the location of known nodes. When packets are sent to nodes, those destination node locations are stored in address-lookup tables called *Ethernet switching tables*. Before sending a packet, a switch using bridging first consults the switching tables to see if that node has already been located. If the location of a node is known, the frame is sent directly to that node.

Transparent bridging uses five mechanisms to create and maintain Ethernet switching tables on the switch:

- Learning
- Forwarding
- Flooding
- Filtering
- Aging

The key bridging mechanism used by LANs and VLANs is *learning*. When a switch is first connected to an Ethernet LAN or VLAN, it has no information about other nodes on the network. As packets are sent, the switch learns the embedded MAC addresses of the sending nodes and stores them in the Ethernet switching table, along with two other pieces of information—the interface (or port) on which the traffic was received on the destination node and the time the address was learned.

Learning allows switches to then do *forwarding*. By consulting the Ethernet switching table to see whether the table already contains the frame's destination MAC address, switches save time and resources when forwarding packets to the known MAC addresses. If the Ethernet switching table does not contain an entry for an address, the switch uses flooding to learn that address.

Flooding finds a particular destination MAC address without using the Ethernet switching table. When traffic originates on the switch and the Ethernet switching table does not yet contain the destination MAC address, the switch first floods the traffic to all other interfaces within the VLAN. When the destination node receives the flooded traffic, it

can send an acknowledgment packet back to the switch, allowing it to learn the MAC address of the node and add the address to its Ethernet switching table.

Filtering, the fourth bridging mechanism, is how broadcast traffic is limited to the local VLAN whenever possible. As the number of entries in the Ethernet switching table grows, the switch pieces together an increasingly complete picture of the VLAN and the larger LAN—it learns which nodes are in the local VLAN and which are on other network segments. The switch uses this information to filter traffic. Specifically, for traffic whose source and destination MAC addresses are in the local VLAN, filtering prevents the switch from forwarding this traffic to other network segments.

To keep entries in the Ethernet switching table current, the switch uses a fifth bridging mechanism, *aging*. Aging is the reason that the Ethernet switching table entries include timestamps. Each time the switch detects traffic from a MAC address, it updates the timestamp. A timer on the switch periodically checks the timestamp, and if it is older than a user-configured value, the switch removes the node's MAC address from the Ethernet switching table. This aging process eventually flushes unavailable network nodes out of the Ethernet switching table.

Packets Are Either Tagged or Untagged

When an Ethernet LAN is divided into VLANs, each VLAN is identified by a unique 802.1Q ID. The number of available VLANs and VLAN IDs are listed below:

- On a switch running ELS software, you can configure 4093 VLANs.
- On a switch running non-ELS software, you can configure 4091 VLANs.

Ethernet packets include a tag protocol identifier (TPID) EtherType field, which identifies the protocol being transported. When a device within a VLAN generates a packet, this field includes a value of 0x8100, which indicates that the packet is a VLAN-tagged packet. The packet also has a VLAN ID field that includes the unique 802.1Q ID, which identifies the VLAN to which the packet belongs.

Junos OS switches support the TPID value 0x9100 for Q-in-Q, and switches that run Junos OS that does not support the Enhanced Layer 2 Software (ELS) configuration style also support values of 0x88a8 (Provider Bridging and Shortest Path Bridging).

For a simple network that has only a single VLAN, all packets include a default 802.1Q tag, which is the only VLAN membership that does not mark the packet as tagged. These packets are untagged packets.

Switch Interface Modes—Access, Trunk, or Tagged Access

Ports, or interfaces, on a switch operate in one of three modes:

- Access mode
- Trunk mode
- Tagged-access mode

Access Mode

An interface in access mode connects a switch to a single network device, such as a desktop computer, an IP telephone, a printer, a file server, or a security camera. Access interfaces accept only untagged packets.

By default, when you boot a switch that runs Junos OS that does not support ELS and use the factory default configuration, or when you boot such a switch and do not explicitly configure a port mode, all interfaces on the switch are in access mode and accept only untagged packets from the VLAN named **default**. You can optionally configure another VLAN and use that VLAN instead of **default**.

On a switch that runs Junos OS that supports ELS, the VLAN named **default** is not supported. Therefore, on such switches, you must explicitly configure at least one VLAN, even if your network is simple and you want only one broadcast domain to exist. After you assign an interface to a VLAN, the interface functions in access mode.

For switches that run either type of software, you can also configure a trunk port or interface to accept untagged packets from a user-configured VLAN. For details about this concept (native VLAN), see [“Trunk Mode and Native VLAN” on page 2092](#).

Trunk Mode

Trunk mode interfaces are generally used to connect switches to one another. Traffic sent between switches can then consist of packets from multiple VLANs, with those packets multiplexed so that they can be sent over the same physical connection. Trunk interfaces usually accept only tagged packets and use the VLAN ID tag to determine both the packets' VLAN origin and VLAN destination.

On a switch that runs software that does not support ELS, an untagged packet is not recognized on a trunk port unless you configure additional settings on that port.

On a switch that runs Junos OS that supports ELS, a trunk port recognizes untagged control packets for protocols such as the Link Aggregation Control Protocol (LACP) and the Link Layer Discovery Protocol (LLDP). However, the trunk port does not recognize untagged data packets unless you configure additional settings on that port.

In the rare case where you want untagged packets to be recognized by a trunk port on switches that run either type of software, you must configure the single VLAN on a trunk port as a *native VLAN*. For more information about native VLANs, see [“Trunk Mode and Native VLAN” on page 2092](#).

Trunk Mode and Native VLAN

On a switch that runs Junos OS that does not support ELS, a trunk port does not recognize packets that do not include VLAN tags, which are also known as untagged packets. On a switch that runs Junos OS that supports ELS, a trunk port recognizes untagged control packets, but it does not recognize untagged data packets. With native VLAN configured, untagged packets that a trunk port normally does not recognize are sent over the trunk interface. In a situation where packets pass from a device, such as an IP phone or printer, to a switch in access mode, and you want those packets sent from the switch over a

trunk port, use native VLAN mode. Create a native VLAN by configuring a VLAN ID for it, and specify that the trunk port is a member of the native VLAN.

The switch's trunk port will then treat those packets differently than the other tagged packets. For example, if a trunk port has three VLANs, 10, 20, and 30, assigned to it with VLAN 10 being the native VLAN, packets on VLAN 10 that leave the trunk port on the other end have no 802.1Q header (tag).

There is another native VLAN option for switches that do not support ELS. You can have the switch add and remove tags for untagged packets. To do this, you first configure the single VLAN as a native VLAN on a port attached to a device on the edge. Then, assign a VLAN ID tag to the single native VLAN on the port connected to a device. Last, add the VLAN ID to the trunk port. Now, when the switch receives the untagged packet, it adds the ID you specified and sends and receives the tagged packets on the trunk port configured to accept that VLAN.

Tagged-Access Mode

Only switches that run Junos OS that does not use the ELS configuration style support tagged-access mode. Tagged-access mode accommodates cloud computing, specifically scenarios including virtual machines or virtual computers. Because several virtual computers can be included on one physical server, the packets generated by one server can contain an aggregation of VLAN packets from different virtual machines on that server. To accommodate this situation, tagged-access mode reflects packets back to the physical server on the same downstream port when the destination address of the packet was learned on that downstream port. Packets are also reflected back to the physical server on the downstream port when the destination has not yet been learned. Therefore, the third interface mode, tagged access, has some characteristics of access mode and some characteristics of trunk mode:

- Like access mode, tagged-access mode connects the switch to an access layer device. Unlike access mode, tagged-access mode is capable of accepting VLAN tagged packets.
- Like trunk mode, tagged-access mode accepts VLAN tagged packets from multiple VLANs. Unlike trunk port interfaces, which are connected at the core/distribution layer, tagged-access port interfaces connect devices at the access layer.

Like trunk mode, tagged-access mode also supports native VLAN.



NOTE: Control packets are never reflected back on the downstream port.

Additional Advantages of Using VLANs

In addition to reducing traffic and thereby speeding up the network, VLANs have the following advantages:

- VLANs provide segmentation services traditionally provided by routers in LAN configurations, thereby reducing hardware equipment costs.
- Packets coupled to a VLAN can be reliably identified and sorted into different domains. You can contain broadcasts within parts of the network, thereby freeing up network

resources. For example, when a DHCP server is plugged into a switch and starts broadcasting its presence, you can prevent some hosts from accessing it by using VLANs to split up the network.

- For security issues, VLANs provide granular control of the network because each VLAN is identified by a single IP subnetwork. All packets passing in and out of a VLAN are consistently tagged with the VLAN ID of that VLAN, thereby providing easy identification, because a VLAN ID on a packet cannot be altered. (For a switch that runs Junos OS that does not support ELS, we recommend that you avoid using 1 as a VLAN ID, because that ID is a default value.)
- VLANs react quickly to host relocation—this is also due to the persistent VLAN tag on packets.
- On an Ethernet LAN, all network nodes must be physically connected to the same network. In VLANs, the physical location of nodes is not important—you can group network devices in any way that makes sense for your organization, such as by department or business function, types of network nodes, or physical location.

Maximum VLANs and VLAN Members Per Switch

The number of VLANs supported per switch varies for each switch. Use the configuration-mode command **set vlans *vlan-name* *vlan-id* ?** to determine the maximum number of VLANs allowed on a switch. You cannot exceed this VLAN limit because you have to assign a specific ID number when you create a VLAN—you could overwrite one of the numbers, but you cannot exceed the limit.

You can, however, exceed the recommended VLAN member maximum for a switch.

On a switch that runs Junos OS that does not support the ELS configuration style, the maximum number of VLAN members allowed on the switch is eight times the maximum number of VLANs that the switch supports ($\text{vmember limit} = \text{vlan max} * 8$). If the configuration of the switch exceeds the recommended VLAN member maximum, a warning message appears when you commit the configuration. If you commit the configuration despite the warning, the commit succeeds, but there is a risk of the Ethernet switching process (eswd) failing as a result of memory allocation failure.

On a switch that runs Junos OS that supports ELS, the maximum number of VLAN members allowed on the switch is 24 times the maximum number of VLANs that the switch supports ($\text{vmember limit} = \text{vlan max} * 24$). If the configuration of the switch exceeds the recommended VLAN member maximum, a warning message appears in the system log (syslog).

A QFabric system supports up to 131,008 VLAN members (vmembers) on a single network node group, server node group, or redundant server node group. The number of vmembers is calculated by multiplying the maximum number of VLANs by 32.

For example, to calculate how many interfaces are required to support 4,000 VLANs, divide the maximum number of vmembers (128,000) by the number of configured VLANs (4,000). In this case, 32 interfaces are required.

On network Node groups and server Node groups, you can configure link aggregation groups (LAGs) across multiple interfaces. Each LAG and VLAN combination is considered a vmember.

A Virtual Chassis Fabric supports up to 512,000 vmembers. The number of vmembers is based on the number of VLANs, and the number of interfaces configured in each VLAN.

A Default VLAN Is Configured on Most Switches

Some switches that run Junos OS that do not support the ELS configuration style are preconfigured with a VLAN named **default** that does not tag packets and operates only with untagged packets. On these switches, each interface already belongs to the VLAN named **default** and all traffic uses this VLAN until you configure more VLANs and assign traffic to those VLANs.



NOTE: When a Juniper Networks QFX3500 or QFX3600 switch is interconnected with other switches in a Virtual Chassis configuration, each individual switch that is included as a member of the configuration is identified with a member ID. The member ID functions as an FPC slot number. When you are configuring interfaces for a Virtual Chassis configuration, you specify the appropriate member ID (0 through 9) as the slot element of the interface name. The default factory settings for a Virtual Chassis configuration include FPC 0 as a member of the default VLAN because FPC 0 is configured as part of the ethernet-switching family. In order to include FPC 1 through FPC 9 in the default VLAN, add the ethernet-switching family to the configurations for those interfaces.

Assigning Traffic to VLANs

You can assign traffic on any switch to a particular VLAN by referencing either the interface port of the traffic or the MAC addresses of devices sending traffic.



NOTE: Two logical interfaces that are configured on the same physical interface cannot be mapped to the same VLAN.

Assign VLAN Traffic According to the Interface Port Source

This method is most commonly used to assign traffic to VLANs. In this case, you specify that all traffic received on a particular switch interface is assigned to a specific VLAN. You configure this VLAN assignment when you configure the switch, by using either the VLAN number (called a VLAN ID) or by using the VLAN name, which the switch then translates into a numeric VLAN ID. This method is referred to simply as creating a VLAN because it is the most commonly used method.

Assign VLAN Traffic According to the Source MAC Address

In this case, all traffic received from a specific MAC address is forwarded to a specific egress interface (next hop) on the switch. MAC-based VLANs are either static (named MAC addresses configured one at a time) or dynamic (configured using a RADIUS server).

To configure a static MAC-based VLAN on a switch that supports ELS, see *Adding a Static MAC Address Entry to the Ethernet Switching Table (CLI Procedure)*. To configure a static MAC-based VLAN on a switch that does not support ELS, see *Adding a Static MAC Address Entry to the Ethernet Switching Table (CLI Procedure)*.

Forwarding VLAN Traffic

To pass traffic within a VLAN, the switch uses Layer 2 forwarding protocols, including IEEE 802.1Q spanning-tree protocols.

To pass traffic between two VLANs, the switch uses standard Layer 3 routing protocols, such as static routing, OSPF, and RIP. The same interfaces that support Layer 2 bridging protocols also support Layer 3 routing protocols, providing multilayer switching.

To pass traffic from a single device on an access port to a switch and then pass those packets on a trunk port, use the native mode configuration previously discussed under [“Trunk Mode” on page 2092](#).

VLANs Communicate with Integrated Routing and Bridging Interfaces or Routed VLAN Interfaces

Traditionally, switches sent traffic to hosts that were part of the same broadcast domain (VLAN) but routers were needed to route traffic from one broadcast domain to another. Also, only routers performed other Layer 3 functions such as traffic engineering.

Switches that run Junos OS that supports the ELS configuration style perform inter-VLAN routing functions using an integrated routing and bridging (IRB) interface named `irb`, while switches that run Junos OS that does not support ELS perform these functions using a routed VLAN interface (RVI) named `vlan`. These interfaces detect both MAC addresses and IP addresses and route data to Layer 3 interfaces, thereby frequently eliminating the need to have both a switch and a router.

Related Documentation

- [Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 2096](#)
- [Example: Setting Up Bridging with Multiple VLANs](#)
- [Understanding FCoE on page 6436](#)
- [Interfaces Overview on page 2785](#)

Example: Setting Up Basic Bridging and a VLAN on the QFX Series

The QFX Series products use bridging and virtual LANs (VLANs) to connect network devices—storage devices, file servers, and other LAN components—in a LAN and to segment the LAN into smaller bridging domains.

To segment traffic on a LAN into separate broadcast domains, you create separate virtual LANs (VLANs) on a switch. Each VLAN is a collection of network nodes. When you use VLANs, frames whose origin and destination are in the same VLAN are forwarded only within the local VLAN, and only frames not destined for the local VLAN are forwarded to other broadcast domains. VLANs thus limit the amount of traffic flowing across the entire LAN, reducing the possible number of collisions and packet retransmissions within the LAN.



NOTE: You cannot configure more than one logical interface that belongs to the same physical interface in the same bridge domain.

This example describes how to configure basic bridging and VLANs for the QFX Series:

- [Requirements on page 2097](#)
- [Overview and Topology on page 2097](#)
- [Configuration on page 2098](#)
- [Verification on page 2107](#)

Requirements

This example uses the following software and hardware components:

- Junos OS Release 11.1 or later for the QFX Series
- A configured and provisioned QFX Series product

Overview and Topology

To use a switch to connect network devices on a LAN, you must at a minimum configure bridging and VLANs. By default, bridging is enabled on all switch interfaces, all interfaces are in access mode, and all interfaces belong to a VLAN called **employee-vlan**, which is automatically configured. When you plug in access devices—such as desktop computers, file servers, and printers—they are joined immediately into the **employee-vlan** VLAN, and the LAN is up and running.

The topology used in this example consists of a single QFX3500 switch, with a total of 48 10-Gbps Ethernet ports. (For the purposes of this example, the QSFP+ ports Q0-Q3, which are ports xe-0/1/0 through xe-0/1/15, are excluded.) You use the ports to connect devices that have their own power sources. Table 1 details the topology used in this configuration example.

Table 183: Components of the Basic Bridging Configuration Topology

Property	Settings
Switch hardware	QFX3500 switch, with 48 10-Gbps Ethernet ports
VLAN name	employee-vlan

Table 183: Components of the Basic Bridging Configuration Topology (*continued*)

Property	Settings
VLAN ID	10
Connections to file servers	xe-0/0/17 and xe-0/0/18
Direct connections to desktop PCs and laptops	xe-0/0/0 through xe-0/0/16
Connections to integrated printer/fax/copier machines	xe-0/0/19 through xe-0/0/40
Unused ports	xe-0/0/41 through xe-0/0/47

Configuration

CLI Quick Configuration To quickly configure a VLAN, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans employee-vlan vlan-id 10
set interfaces xe-0/0/0 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/2 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/3 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/4 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/5 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/6 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/7 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/8 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/9 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/10 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/11 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/12 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/13 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/14 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/15 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/16 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/17 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/18 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/19 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/20 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/21 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/22 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/23 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/24 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/25 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/26 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/27 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/28 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/29 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/30 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/31 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/32 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/33 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/34 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/35 unit 0 family ethernet-switching vlan members employee-vlan
```

```

set interfaces xe-0/0/36 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/37 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/38 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/39 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/40 unit 0 family ethernet-switching vlan members employee-vlan

```

Step-by-Step Procedure

To set up basic bridging and a VLAN:

1. Create a VLAN named employee-vlan and specify the VLAN ID of 10 for it:

```

[edit vlans]
user@switch# set employee-vlan vlan-id 10

```

2. Assign interfaces xe-0/0/0 through xe-0/0/40 to the employee-vlan VLAN:

```

[edit interface]
user@switch# set xe-0/0/0 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/1 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/2 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/3 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/4 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/5 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/6 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/7 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/8 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/9 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/10 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/11 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/12 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/13 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/14 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/15 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/16 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/17 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/18 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/19 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/20 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/21 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/22 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/23 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/24 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/25 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/26 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/27 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/28 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/29 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/30 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/31 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/32 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/33 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/34 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/35 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/36 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/37 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/38 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/39 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/40 unit 0 family ethernet-switching vlan members employee-vlan

```

3. Connect the two file servers to ports xe-0/0/17 and xe-0/0/18.

4. Connect the desktop PCs and laptops to ports xe-0/0/0 through xe-0/0/16.
5. Connect the integrated printer/fax/copier machines to ports xe-0/0/19 through xe-0/0/40.

Results Check the results of the configuration:

```
user@switch> show configuration
xe-0/0/0 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members employee-vlan;
      }
    }
  }
}
xe-0/0/1 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members employee-vlan;
      }
    }
  }
}
xe-0/0/2 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members employee-vlan;
      }
    }
  }
}
xe-0/0/3 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members employee-vlan;
      }
    }
  }
}
xe-0/0/4 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members employee-vlan;
      }
    }
  }
}
xe-0/0/5 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members employee-vlan;
      }
    }
  }
}
xe-0/0/6 {
  unit 0 {
    family ethernet-switching {
```

```
        vlan {
            members employee-vlan;
        }
    }
}
xe-0/0/7 {
    unit 0 {
        family ethernet-switching {
            vlan {
                members employee-vlan;
            }
        }
    }
}
xe-0/0/8 {
    unit 0 {
        family ethernet-switching {
            vlan {
                members employee-vlan;
            }
        }
    }
}
xe-0/0/9 {
    unit 0 {
        family ethernet-switching {
            vlan {
                members employee-vlan;
            }
        }
    }
}
xe-0/0/10 {
    unit 0 {
        family ethernet-switching {
            vlan {
                members employee-vlan;
            }
        }
    }
}
xe-0/0/11 {
    unit 0 {
        family ethernet-switching {
            vlan {
                members employee-vlan;
            }
        }
    }
}
xe-0/0/12 {
    unit 0 {
        family ethernet-switching {
            vlan {
                members employee-vlan;
            }
        }
    }
}
```

```
xe-0/0/13 {  
  unit 0 {  
    family ethernet-switching {  
      vlan {  
        members employee-vlan;  
      }  
    }  
  }  
}  
xe-0/0/14 {  
  unit 0 {  
    family ethernet-switching {  
      vlan {  
        members employee-vlan;  
      }  
    }  
  }  
}  
xe-0/0/15 {  
  unit 0 {  
    family ethernet-switching {  
      vlan {  
        members employee-vlan;  
      }  
    }  
  }  
}  
xe-0/0/16 {  
  unit 0 {  
    family ethernet-switching {  
      vlan {  
        members employee-vlan;  
      }  
    }  
  }  
}  
xe-0/0/17 {  
  unit 0 {  
    family ethernet-switching {  
      vlan {  
        members employee-vlan;  
      }  
    }  
  }  
}  
xe-0/0/18 {  
  unit 0 {  
    family ethernet-switching {  
      vlan {  
        members employee-vlan;  
      }  
    }  
  }  
}  
xe-0/0/19 {  
  unit 0 {  
    family ethernet-switching {  
      vlan {  
        members employee-vlan;  
      }  
    }  
  }  
}
```

```
    }  
  }  
  xe-0/0/20 {  
    unit 0 {  
      family ethernet-switching {  
        vlan {  
          members employee-vlan;  
        }  
      }  
    }  
  }  
  xe-0/0/21 {  
    unit 0 {  
      family ethernet-switching {  
        vlan {  
          members employee-vlan;  
        }  
      }  
    }  
  }  
  xe-0/0/22 {  
    unit 0 {  
      family ethernet-switching {  
        vlan {  
          members employee-vlan;  
        }  
      }  
    }  
  }  
  xe-0/0/23 {  
    unit 0 {  
      family ethernet-switching {  
        vlan {  
          members employee-vlan;  
        }  
      }  
    }  
  }  
  xe-0/0/24 {  
    unit 0 {  
      family ethernet-switching {  
        vlan {  
          members employee-vlan;  
        }  
      }  
    }  
  }  
  xe-0/0/25 {  
    unit 0 {  
      family ethernet-switching {  
        vlan {  
          members employee-vlan;  
        }  
      }  
    }  
  }  
  xe-0/0/26 {  
    unit 0 {  
      family ethernet-switching {  
        vlan {  
          members employee-vlan;  
        }  
      }  
    }  
  }
```



```

    }
  }
}
xe-0/0/27 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members employee-vlan;
      }
    }
  }
}
xe-0/0/28 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members employee-vlan;
      }
    }
  }
}
xe-0/0/29 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members employee-vlan;
      }
    }
  }
}
xe-0/0/30 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members employee-vlan;
      }
    }
  }
}
xe-0/0/31 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members employee-vlan;
      }
    }
  }
}
xe-0/0/32 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members employee-vlan;
      }
    }
  }
}
xe-0/0/33 {
  unit 0 {
    family ethernet-switching {

```

```
        vlan {
            members employee-vlan;
        }
    }
}
xe-0/0/34 {
    unit 0 {
        family ethernet-switching {
            vlan {
                members employee-vlan;
            }
        }
    }
}
xe-0/0/35 {
    unit 0 {
        family ethernet-switching {
            vlan {
                members employee-vlan;
            }
        }
    }
}
xe-0/0/36 {
    unit 0 {
        family ethernet-switching {
            vlan {
                members employee-vlan;
            }
        }
    }
}
xe-0/0/37 {
    unit 0 {
        family ethernet-switching {
            vlan {
                members employee-vlan;
            }
        }
    }
}
xe-0/0/38 {
    unit 0 {
        family ethernet-switching {
            vlan {
                members employee-vlan;
            }
        }
    }
}
xe-0/0/39 {
    unit 0 {
        family ethernet-switching {
            vlan {
                members employee-vlan;
            }
        }
    }
}
xe-0/0/40 {
```

```
unit 0 {  
    family ethernet-switching {  
        vlan {  
            members employee-vlan;  
        }  
    }  
}
```

Verification

To verify that switching is operational and that **employee-vlan** has been created, perform these tasks:

- [Verifying That the VLAN Has Been Created on page 2107](#)
- [Verifying That Interfaces Are Associated with the Proper VLANs on page 2108](#)

Verifying That the VLAN Has Been Created

Purpose Verify that the VLAN named **employee-vlan** has been created on the switch.

Action List all VLANs configured on the switch:

```
user@switch> show vlans
Routing instance      VLAN name      Tag      Interfaces
default-switch        employee-vlan   10
                      xe-0/0/0.0
                      xe-0/0/1.0
                      xe-0/0/2.0
                      xe-0/0/3.0
                      xe-0/0/4.0
                      xe-0/0/5.0
                      xe-0/0/6.0
                      xe-0/0/7.0
                      xe-0/0/8.0
                      xe-0/0/9.0
                      xe-0/0/10.0
                      xe-0/0/11.0
                      xe-0/0/12.0
                      xe-0/0/13.0
                      xe-0/0/14.0
                      xe-0/0/15.0
                      xe-0/0/16.0
                      xe-0/0/17.0
                      xe-0/0/18.0
                      xe-0/0/19.0
                      xe-0/0/20.0
                      xe-0/0/21.0
                      xe-0/0/22.0
                      xe-0/0/23.0
                      xe-0/0/24.0
                      xe-0/0/25.0
                      xe-0/0/26.0
                      xe-0/0/27.0
                      xe-0/0/28.0
                      xe-0/0/29.0
                      xe-0/0/30.0
                      xe-0/0/31.0
                      xe-0/0/32.0
                      xe-0/0/33.0
                      xe-0/0/34.0
                      xe-0/0/35.0
                      xe-0/0/36.0
                      xe-0/0/37.0
                      xe-0/0/38.0
                      xe-0/0/39.0
                      xe-0/0/40.0
...

```

Meaning The `show vlans` command lists the VLANs configured on the switch. This output shows that the VLAN `employee-vlan` has been created.

Verifying That Interfaces Are Associated with the Proper VLANs

Purpose Verify that Ethernet switching is enabled on switch interfaces and that all interfaces are included in the VLAN.

Action List all interfaces on which switching is enabled:

```

user@switch> show ethernet-switching interfaces
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state   interface flags
xe-0/0/0.0                65535                    untagged
                        employee-vlan 10
                        65535    Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state   interface flags
xe-0/0/1.0                65535                    untagged
                        employee-vlan 10
                        65535    Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state   interface flags
xe-0/0/2.0                65535                    untagged
                        employee-vlan 10
                        65535    Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state   interface flags
xe-0/0/3.0                65535                    untagged
                        employee-vlan 10
                        65535    Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state   interface flags
xe-0/0/4.0                65535                    untagged
                        employee-vlan 10
                        65535    Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state   interface flags
xe-0/0/5.0                65535                    untagged
                        employee-vlan 10
                        65535    Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state   interface flags
xe-0/0/6.0                65535                    untagged
                        employee-vlan 10
                        65535    Discarding
Routing Instance Name : default-switch

```

```

Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG      MAC      STP      Logical      Tagging
interface    members          limit  state    interface flags
xe-0/0/7.0
      employee-vlan 10
                        65535
                        Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG      MAC      STP      Logical      Tagging
interface    members          limit  state    interface flags
xe-0/0/8.0
      employee-vlan 10
                        65535
                        Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG      MAC      STP      Logical      Tagging
interface    members          limit  state    interface flags
xe-0/0/9.0
      employee-vlan 10
                        65535
                        Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG      MAC      STP      Logical      Tagging
interface    members          limit  state    interface flags
xe-0/0/10.0
      employee-vlan 10
                        65535
                        Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG      MAC      STP      Logical      Tagging
interface    members          limit  state    interface flags
xe-0/0/11.0
      employee-vlan 10
                        65535
                        Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG      MAC      STP      Logical      Tagging
interface    members          limit  state    interface flags
xe-0/0/12.0
      employee-vlan 10
                        65535
                        Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG      MAC      STP      Logical      Tagging
interface    members          limit  state    interface flags
xe-0/0/13.0
      employee-vlan 10
                        65535
                        Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG      MAC      STP      Logical      Tagging
interface    members          limit  state    interface flags
xe-0/0/14.0
      employee-vlan 10
                        65535
                        Discarding

```

```

        employee-vlan 10
            65535    Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG  MAC      STP      Logical      Tagging
interface    members          limit  state      interface flags
xe-0/0/15.0              65535                    untagged
        employee-vlan 10
            65535    Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG  MAC      STP      Logical      Tagging
interface    members          limit  state      interface flags
xe-0/0/16.0              65535                    untagged
        employee-vlan 10
            65535    Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG  MAC      STP      Logical      Tagging
interface    members          limit  state      interface flags
xe-0/0/17.0              65535                    untagged
        employee-vlan 10
            65535    Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG  MAC      STP      Logical      Tagging
interface    members          limit  state      interface flags
xe-0/0/18.0              65535                    untagged
        employee-vlan 10
            65535    Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG  MAC      STP      Logical      Tagging
interface    members          limit  state      interface flags
xe-0/0/19.0              65535                    untagged
        employee-vlan 10
            65535    Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG  MAC      STP      Logical      Tagging
interface    members          limit  state      interface flags
xe-0/0/20.0              65535                    untagged
        employee-vlan 10
            65535    Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG  MAC      STP      Logical      Tagging
interface    members          limit  state      interface flags
xe-0/0/21.0              65535                    untagged
        employee-vlan 10
            65535    Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )

```

```

Logical    Vlan    TAG    MAC    STP    Logical    Tagging
interface  members  state  limit  state  interface  flags
xe-0/0/22.0
    employee-vlan 10
                                65535    Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical    Vlan    TAG    MAC    STP    Logical    Tagging
interface  members  limit  state  interface  flags
xe-0/0/23.0
    employee-vlan 10
                                65535    Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical    Vlan    TAG    MAC    STP    Logical    Tagging
interface  members  limit  state  interface  flags
xe-0/0/24.0
    employee-vlan 10
                                65535    Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical    Vlan    TAG    MAC    STP    Logical    Tagging
interface  members  limit  state  interface  flags
xe-0/0/25.0
    employee-vlan 10
                                65535    Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical    Vlan    TAG    MAC    STP    Logical    Tagging
interface  members  limit  state  interface  flags
xe-0/0/26.0
    employee-vlan 10
                                65535    Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical    Vlan    TAG    MAC    STP    Logical    Tagging
interface  members  limit  state  interface  flags
xe-0/0/27.0
    employee-vlan 10
                                65535    Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical    Vlan    TAG    MAC    STP    Logical    Tagging
interface  members  limit  state  interface  flags
xe-0/0/28.0
    employee-vlan 10
                                65535    Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical    Vlan    TAG    MAC    STP    Logical    Tagging
interface  members  limit  state  interface  flags
xe-0/0/29.0
    employee-vlan 10
                                65535    Discarding

```



```

Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG  MAC      STP      Logical      Tagging
interface    members   limit state   interface flags
xe-0/0/30.0   employee-vlan 10  65535   Discarding
                        65535   Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG  MAC      STP      Logical      Tagging
interface    members   limit state   interface flags
xe-0/0/31.0   employee-vlan 10  65535   Discarding
                        65535   Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG  MAC      STP      Logical      Tagging
interface    members   limit state   interface flags
xe-0/0/32.0   employee-vlan 10  65535   Discarding
                        65535   Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG  MAC      STP      Logical      Tagging
interface    members   limit state   interface flags
xe-0/0/33.0   employee-vlan 10  65535   Discarding
                        65535   Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG  MAC      STP      Logical      Tagging
interface    members   limit state   interface flags
xe-0/0/34.0   employee-vlan 10  65535   Discarding
                        65535   Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG  MAC      STP      Logical      Tagging
interface    members   limit state   interface flags
xe-0/0/35.0   employee-vlan 10  65535   Discarding
                        65535   Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG  MAC      STP      Logical      Tagging
interface    members   limit state   interface flags
xe-0/0/36.0   employee-vlan 10  65535   Discarding
                        65535   Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG  MAC      STP      Logical      Tagging
interface    members   limit state   interface flags

```

```

xe-0/0/37.0          65535          untagged
                    employee-vlan 10
                    65535          Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG  MAC      STP      Logical      Tagging
interface    members   limit state   interface flags
xe-0/0/38.0   65535          untagged
                    employee-vlan 10
                    65535          Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG  MAC      STP      Logical      Tagging
interface    members   limit state   interface flags
xe-0/0/39.0   65535          untagged
                    employee-vlan 10
                    65535          Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG  MAC      STP      Logical      Tagging
interface    members   limit state   interface flags
xe-0/0/40.0   65535          untagged
                    employee-vlan 10
                    65535          Discarding
...

```

Meaning The `show ethernet-switching interfaces` command lists all interfaces on which switching is enabled (in the **Logical interface** column), along with the VLANs that are active on the interfaces (in the **VLAN members** column). The output in this example shows all the connected interfaces, xe-0/0/0 through xe-0/0/40, are all part of VLAN **employee-vlan**. Notice that the interfaces listed are the logical interfaces, not the physical interfaces. For example, the output shows xe-0/0/0.0 instead of xe-0/0/0. This is because Junos OS creates VLANs on logical interfaces, not directly on physical interfaces.

Related Documentation

- [Example: Setting Up Bridging with Multiple VLANs](#)
- [Understanding Bridging and VLANs on page 2089](#)

Example: Setting Up Bridging with Multiple VLANs

The QFX Series products use bridging and virtual LANs (VLANs) to connect network devices in a LAN—storage devices, file servers, and other network components—and to segment the LAN into smaller bridging domains.

To segment traffic on a LAN into separate broadcast domains, you create separate virtual LANs (VLANs) on a switch. Each VLAN is a collection of network nodes. When you use VLANs, frames whose origin and destination are in the same VLAN are forwarded only within the local VLAN, and only frames not destined for the local VLAN are forwarded to other broadcast domains. VLANs thus limit the amount of traffic flowing across the entire LAN, reducing the possible number of collisions and packet retransmissions within the LAN.

This example describes how to configure bridging for the QFX Series and how to create two VLANs to segment the LAN:



NOTE: This task supports the Enhanced Layer 2 Software (ELS) configuration style. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 41](#). If your switch runs software that does not support ELS, see *Example: Setting Up Bridging with Multiple VLANs*.

- [Requirements on page 2115](#)
- [Overview and Topology on page 2115](#)
- [Configuration on page 2116](#)
- [Verification on page 2118](#)

Requirements

This example uses the following hardware and software components:

- A configured and provisioned QFX3500 switch
- Junos OS Release 13.2X50-D15 or later for the QFX Series

Overview and Topology

Switches connect all devices in an office or data center into a single LAN to provide sharing of common resources such as file servers. The default configuration creates a single VLAN, and all traffic on the switch is part of that broadcast domain. Creating separate network segments reduces the span of the broadcast domain and enables you to group related users and network resources without being limited by physical cabling or by the location of a network device in the building or on the LAN.

This example shows a simple configuration to illustrate the basic steps for creating two VLANs on a single switch. One VLAN, called **sales**, is for the sales and marketing group, and a second, called **support**, is for the customer support team. The sales and support groups each have their own dedicated file servers and other resources. For the switch ports to be segmented across the two VLANs, each VLAN must have its own broadcast domain, identified by a unique name and tag (VLAN ID). In addition, each VLAN must be on its own distinct IP subnet.

The topology used in this example consists of a single QFX3500 switch, with a total of 48 10-Gbps Ethernet ports. (For the purposes of this example, the QSFP+ ports Q0-Q3, which are ports xe-0/1/0 through xe-0/1/15, are excluded.)

Table 184: Components of the Multiple VLAN Topology

Property	Settings
Switch hardware	QFX3500 switch configured with 48 10-Gbps Ethernet ports (xe-0/0/0 through xe-0/0/47)

Table 184: Components of the Multiple VLAN Topology (*continued*)

Property	Settings
VLAN names and tag IDs	sales , tag 100 support , tag 200
VLAN subnets	sales : 192.0.2.0/25 (addresses 192.0.2.1 through 192.0.2.126) support : 192.0.2.128/25 (addresses 192.0.2.129 through 192.0.2.254)
Interfaces in VLAN sales	File servers: xe-0/0/20 and xe-0/0/21
Interfaces in VLAN support	File servers: xe-0/0/46 and xe-0/0/47
Unused interfaces	xe-0/0/2 and xe-0/0/25

This configuration example creates two IP subnets, one for the sales VLAN and the second for the support VLAN. The switch bridges traffic within a VLAN. For traffic passing between two VLANs, the switch routes the traffic using a Layer 3 routing interface on which you have configured the address of the IP subnet.

To keep the example simple, the configuration steps show only a few devices in each of the VLANs. Use the same configuration procedure to add more LAN devices.

Configuration

CLI Quick Configuration

To quickly configure Layer 2 switching for the two VLANs (**sales** and **support**) and to quickly configure Layer 3 routing of traffic between the two VLANs, copy the following commands and paste them into the switch terminal window:

```
[edit]
set interfaces xe-0/0/0 unit 0 family ethernet-switching vlan members sales
set interfaces xe-0/0/3 unit 0 family ethernet-switching vlan members sales
set interfaces xe-0/0/22 unit 0 family ethernet-switching vlan members sales
set interfaces xe-0/0/20 unit 0 description "Sales file server port"
set interfaces xe-0/0/20 unit 0 family ethernet-switching vlan members sales
set interfaces xe-0/0/24 unit 0 family ethernet-switching vlan members support
set interfaces xe-0/0/26 unit 0 family ethernet-switching vlan members support
set interfaces xe-0/0/44 unit 0 family ethernet-switching vlan members support
set interfaces xe-0/0/46 unit 0 description "Support file server port"
set interfaces xe-0/0/46 unit 0 family ethernet-switching vlan members support
set interfaces vlan unit 0 family inet address 192.0.2.0/25
set interfaces vlan unit 1 family inet address 192.0.2.128/25
set vlans sales l3-interface irb.0
set vlans sales vlan-id 100
set vlans support vlan-id 200
set vlans support l3-interface irb.1
```

Step-by-Step Procedure Configure the switch interfaces and the VLANs to which they belong. By default, all interfaces are in access mode, so you do not have to configure the port mode.

1. Configure the interface for the file server in the **sales** VLAN:

```
[edit interfaces xe-0/0/20 unit 0]
user@switch# set description "Sales file server port"
user@switch# set family ethernet-switching vlan members sales
```
2. Configure the interface for the file server in the **support** VLAN:

```
[edit interfaces xe-0/0/46 unit 0]
user@switch# set description "Support file server port"
user@switch# set family ethernet-switching vlan members support
```
3. Create the subnet for the **sales** broadcast domain:

```
[edit interfaces]
user@switch# set vlan unit 0 family inet address 192.0.2.1/25
```
4. Create the subnet for the **support** broadcast domain:

```
[edit interfaces]
user@switch# set vlan unit 1 family inet address 192.0.2.129/25
```
5. Configure the VLAN tag IDs for the **sales** and **support** VLANs:

```
[edit vlans]
user@switch# set sales vlan-id 100
user@switch# set support vlan-id 200
```
6. To route traffic between the **sales** and **support** VLANs, define the interfaces that are members of each VLAN and associate a Layer 3 interface:

```
[edit vlans]
user@switch# set sales l3-interface irb.0
user@switch# set support l3-interface irb.1
```

Configuration Results Display the results of the configuration:

```
user@switch> show configuration
interfaces {
  xe-0/0/20 {
    unit 0 {
      description "Sales file server port";
      family ethernet-switching {
        vlan members sales;
      }
    }
  }
  xe-0/0/46 {
    unit 0 {
      description "Support file server port";
      family ethernet-switching {
        vlan members support;
      }
    }
  }
  vlans {
    unit 0 {
      family inet address 192.0.2.1/25;
    }
    unit 1 {
      family inet address 192.0.2.129/25;
    }
  }
}
```

```

    }
  }
}
vllans {
  sales {
    vlan-id 100;
    interface xe-0/0/0.0;
    interface xe-0/0/3.0;
    interface xe-0/0/20.0;
    interface xe-0/0/22.0;
    l3-interface irb0;
  }
  support {
    vlan-id 200;
    interface xe-0/0/24.0;
    interface xe-0/0/26.0;
    interface xe-0/0/44.0;
    interface xe-0/0/46.0;
    l3-interface irb1;
  }
}
}

```



TIP: To quickly configure the sales and support VLAN interfaces, issue the **load merge terminal** command. Then copy the hierarchy and paste it into the switch terminal window.

Verification

Verify that the **sales** and **support** VLANs have been created and are operating properly, perform these tasks:

- [Verifying That the VLANs Have Been Created and Associated with the Correct Interfaces on page 2118](#)
- [Verifying That Traffic Is Being Routed Between the Two VLANs on page 2119](#)
- [Verifying That Traffic Is Being Switched Between the Two VLANs on page 2119](#)

Verifying That the VLANs Have Been Created and Associated with the Correct Interfaces

Purpose Verify that the **sales** and **support** VLANs have been created on the switch and that all connected interfaces on the switch are members of the correct VLAN.

Action To list all VLANs configured on the switch, use the **show vlans** command:

```

user@switch> show vlans
Name      Tag      Interfaces
default
          xe-0/0/1.0, xe-0/0/2.0, xe-0/0/4.0, xe-0/0/5.0,
          xe-0/0/6.0, xe-0/0/7.0, xe-0/0/8.0, xe-0/0/9.0,
          xe-0/0/10.0*, xe-0/0/11.0, xe-0/0/12.0, xe-0/0/13.0*,
          xe-0/0/14.0, xe-0/0/15.0, xe-0/0/16.0, xe-0/0/17.0,

```

```

xe-0/0/18.0, xe-0/0/19.0, xe-0/0/21.0, xe-0/0/23.0*,
xe-0/0/25.0, xe-0/0/27.0, xe-0/0/28.0, xe-0/0/29.0,
xe-0/0/30.0, xe-0/0/31.0, xe-0/0/32.0, xe-0/0/33.0,
xe-0/0/34.0, xe-0/0/35.0, xe-0/0/36.0, xe-0/0/37.0,
xe-0/0/38.0, xe-0/0/39.0, xe-0/0/40.0, xe-0/0/41.0,
xe-0/0/42.0, xe-0/0/43.0, xe-0/0/45.0, xe-0/0/47.0,
xe-0/1/0.0*, xe-0/1/1.0*, xe-0/1/2.0*, xe-0/1/3.0*

sales      100
           xe-0/0/0.0*, xe-0/0/3.0, xe-0/0/20.0, xe-0/0/22.0

support    200
           xe-0/0/0.24, xe-0/0/26.0, xe-0/0/44.0, xe-0/0/46.0*

mgmt
           me0.0*

```

Meaning The **show vlans** command lists all VLANs configured on the switch and which interfaces are members of each VLAN. This command output shows that the **sales** and **support** VLANs have been created. The **sales** VLAN has a tag ID of 100 and is associated with interfaces **xe-0/0/0.0**, **xe-0/0/3.0**, **xe-0/0/20.0**, and **xe-0/0/22.0**. VLAN **support** has a tag ID of 200 and is associated with interfaces **xe-0/0/24.0**, **xe-0/0/26.0**, **xe-0/0/44.0**, and **xe-0/0/46.0**.

Verifying That Traffic Is Being Routed Between the Two VLANs

Purpose Verify routing between the two VLANs.

Action List the Layer 3 routes in the switch Address Resolution Protocol (ARP) table:

```

user@switch> show arp
MAC Address      Address      Name      Flags
00:00:0c:06:2c:0d 192.0.2.3    vlan.0    None
00:13:e2:50:62:e0 192.0.2.11   vlan.1    None

```

Meaning Sending IP packets on a multiaccess network requires mapping from an IP address to a MAC address (the physical or hardware address). The ARP table displays the mapping between the IP address and MAC address for both **vlan.0** (associated with **sales**) and **vlan.1** (associated with **support**). These VLANs can route traffic to each other.

Verifying That Traffic Is Being Switched Between the Two VLANs

Purpose Verify that learned entries are being added to the Ethernet switching table.

Action List the contents of the Ethernet switching table:

```

user@switch> show ethernet-switching table

Ethernet-switching table: 8 entries, 5 learned
VLAN      MAC address      Type      Age      Interfaces
default   *                Flood     -        All-members
default   00:00:05:00:00:01 Learn     -        xe-0/0/10.0
default   00:00:5e:00:01:09 Learn     -        xe-0/0/13.0

```

default	00:19:e2:50:63:e0	Learn	- xe-0/0/23.0
sales	*	Flood	- All-members
sales	00:00:5e:00:07:09	Learn	- xe-0/0/0.0
support	*	Flood	- All-members
support	00:00:5e:00:01:01	Learn	- xe-0/0/46.0

Meaning The output shows that learned entries for the **sales** and **support** VLANs have been added to the Ethernet switching table, and are associated with interfaces **xe-0/0/0.0** and **xe-0/0/46.0**. Even though the VLANs were associated with more than one interface in the configuration, these interfaces are the only ones that are currently operating.

Related Documentation

- [Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 2096](#)
- [Understanding Bridging and VLANs on page 2089](#)

Configuring VLANs

Switches use VLANs to make logical groupings of network nodes with their own broadcast domains. You can use VLANs to limit the traffic flowing across the entire LAN and reduce collisions and packet retransmissions.



NOTE: This task supports the Enhanced Layer 2 Software (ELS) configuration style. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 41](#). If your switch runs software that does not support ELS, see [Configuring VLANs](#).



NOTE: Two logical interfaces that are configured on the same physical interface cannot be mapped to the same VLAN.

For each endpoint on the VLAN, configure the following VLAN parameters on the corresponding interface:

1. Specify the description of the VLAN:

```
[edit interfaces interface-name unit 0]
user@switch# set description vlan-description
```

2. Specify the unique name of the VLAN:



NOTE: Switches that run Junos OS with the ELS configuration style do not support a default VLAN. Therefore, on such switches, you must explicitly configure at least one VLAN, even if your network is simple and you want only one broadcast domain to exist.

```
[edit interfaces interface-name unit 0]
user@switch# set family ethernet-switching vlan members vlan-name
```

3. Create the subnet for the VLAN:

- ```
[edit interfaces]
user@switch# set vlan unit 0 family inet address ip-address
```
4. Configure the VLAN tag ID or VLAN ID list for the VLAN:
- ```
[edit vlans]
user@switch# set vlan-name vlan-id vlan-id-number
or

[edit vlans]
user@switch# set vlan-name vlan-id-list [vlan-ids | vlan-id--vlan-id-]
```
5. Specify a VLAN firewall filter to be applied to incoming or outgoing packets:
- ```
[edit vlans]
user@switch# set vlan-name filter (input | output) filter-name
```

#### Related Documentation

- [Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 2096](#)
- [Configuring IRB Interfaces on page 2131](#)
- [Creating a Series of Tagged VLANs](#)
- [Understanding Bridging and VLANs on page 2089](#)

## Configuring the Native VLAN Identifier (CLI Procedure)



**NOTE:** This task uses Junos OS for EX Series switches and Junos OS for QFX3500 and QFX3600 switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring the Native VLAN Identifier (CLI Procedure)*. For ELS details, see “Getting Started with Enhanced Layer 2 Software” on page 41.

Switches can receive and forward routed or bridged Ethernet frames with 802.1Q VLAN tags. Typically, trunk ports, which connect switches to each other, accept untagged control packets but do not accept untagged data packets. You can enable a trunk port to accept untagged data packets by configuring a native VLAN ID on the interface on which you want the untagged data packets to be received. The logical interface on which untagged packets are to be received must be configured with the same VLAN ID as the native VLAN ID configured on the physical interface.

To configure the native VLAN ID by using the command-line interface (CLI):

1. On the interface on which you want untagged data packets to be received, set the interface mode to **trunk**, which specifies that the interface is in multiple VLANs and can multiplex traffic between different VLANs.:

```
[edit interfaces]
user@switch# set interface-name unit logical-unit-number family
ethernet-switching interface-mode trunk
```

2. Configure the native VLAN ID:

```
[edit interfaces]
user@switch# set interface-name native-vlan-id vlan-id
```

3. Specify that the logical interface that will receive the untagged data packets is a member of the native VLAN:

```
[edit interfaces]
user@switch# set interface-name unit logical-unit-number family
ethernet-switching vlan members vlan-id
```

**Related  
Documentation**

- *Understanding Bridging and VLANs on EX Series Switches*
- *Example: Connecting Access Switches to a Distribution Switch*
- *Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch*
- [Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 2096](#)

---

## Creating a Series of Tagged VLANs

---

When you divide an Ethernet LAN into multiple VLANs, each VLAN is assigned a unique IEEE 802.1Q tag. This tag is associated with each frame in the VLAN, and the network nodes receiving the traffic can use the tag to identify which VLAN a frame is associated with.

Instead of configuring VLANs and 802.1Q tags one at a time for a trunk interface, you can configure a VLAN range to create a series of tagged VLANs.

When an Ethernet LAN is divided into VLANs, each VLAN is identified by a unique 802.1Q tag. The tag is applied to all frames so that the network nodes receiving the frames can detect which VLAN the frames belong to. Trunk ports, which multiplex traffic among a number of VLANs, use the tag to determine the origin of frames and where to forward them.

For example, you could configure the VLAN **employee** and specify a tag range of **10 through 12**. This creates the following VLANs and tags:

- VLAN **employee-10**, tag 10
- VLAN **employee-11**, tag 11
- VLAN **employee-12**, tag 12

Creating tagged VLANs in a series has the following limitations:

- Layer 3 interfaces do not support this feature.
- Because an access interface can only support one VLAN member, access interfaces also do not support this feature.



**NOTE:** This task uses Junos OS for Junos OS for QFX3500 and QFX3600 switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Creating a Series of Tagged VLANs*. For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on page 41.

---

To configure a series of tagged VLANs using the CLI (here, the VLAN is **employee**):

1. Configure the series (here, a VLAN series from 120 through 130):

```
[edit]
user@switch# set vlans employee vlan-id-list [120-130]
```

2. Associate a series of tagged VLANs when you configure an interface in one of two ways:

- Include the name of the series:

```
[edit interfaces]
user@switch# set interfaces xe-0/0/22.0 family ethernet-switching vlanmembers employee
```

- Include the VLAN range:

```
[edit interfaces]
user@switch# set interfaces xe-0/0/22.0 family ethernet-switching vlan members 120-130
```

Associating a series of tagged VLANs to an interface by name or by VLAN range the same result: VLANs **\_\_employee\_120\_\_** through **\_\_employee\_130\_\_** are created.



**NOTE:** When a series of VLANs is created using the `vlan-id-list` command, the VLAN names are preceded and followed by a double underscore.

---

#### Related Documentation

- [Example: Setting Up Bridging with Multiple VLANs on page 2114](#)
- [Understanding Bridging and VLANs on page 2089](#)

---

## Understanding Integrated Routing and Bridging

---

To segment traffic on a LAN into separate broadcast domains, you create separate virtual LANs (VLANs). VLANs limit the amount of traffic flowing across the entire LAN, reducing the possible number of collisions and packet retransmissions within the LAN. For example, you might want to create a VLAN that includes the employees in a department and the resources that they use often, such as printers, servers, and so on.

Of course, you also want to allow these employees to communicate with people and resources in other VLANs. To forward packets between VLANs, you normally you need a router that connects the VLANs. However, you can accomplish this forwarding on a switch without using a router by configuring an integrated routing and bridging (IRB) interface. (These interfaces are also called routed VLAN interfaces, or RVIs). Using this approach reduces complexity and avoids the costs associated with purchasing, installing, managing, powering, and cooling another device.

An IRB is a special type of Layer 3 virtual interface named **vlan**. Like normal Layer 3 interfaces, the **vlan** interface needs a logical unit number with an IP address. In fact, to be useful an IRB needs at least two logical units and two IP addresses—you must create units with addresses in each of the subnets associated with the VLANs between which you want traffic to be routed. That is, if you have two VLANs (for example, VLAN **red** and VLAN **blue**) with corresponding subnets, your IRB must have a logical unit with an address in the subnet for **red** and a logical unit with an address in the subnet for **blue**. The switch

automatically creates direct routes to these subnets and uses these routes to forward traffic between VLANs.



**NOTE:** If you are using a version of Junos OS that supports Enhanced Layer 2 Software (ELS), you can also create a Layer 3 virtual interface named `irb` instead of `vlan`—that is, both statements are supported by ELS

Table 185 shows values you might use when configuring an IRB:

**Table 185: Sample IRB Values**

| Property                      | Settings                                                                                                                        |
|-------------------------------|---------------------------------------------------------------------------------------------------------------------------------|
| VLAN names and tags (IDs)     | blue, ID 100<br>red, ID 200                                                                                                     |
| Subnets associated with VLANs | blue: 192.0.2.0/25 (addresses 192.0.2.1 through 192.0.2.126)<br>red: 192.0.2.128/25 (addresses 192.0.2.129 through 192.0.2.254) |
| IRB name                      | interface <code>irb</code>                                                                                                      |
| IRB units and addresses       | logical unit 100: 192.0.2.1/25<br>logical unit 200: 192.0.2.129/25                                                              |

For the sake of consistency and to avoid confusion, Table 185 shows IRB logical unit numbers that match the IDs of the corresponding VLANs. However, you do not have to assign logical unit numbers that match the VLAN IDs—you can use any values for the units. To bind the logical units of the IRB to the appropriate VLANs, you use the `l3-interface` statement.

Because IRBs operate at Layer 3, you can use Layer 3 services such as firewall filters or CoS rewriting with them.

Table 186 shows the number of IRBs/RVIs that each QFX platform supports.

**Table 186: Number of Supported IRBs/RVIs by Platform**

| Platform  | Number of Supported IRBs/RVIs |
|-----------|-------------------------------|
| QFX3500   | 1200                          |
| QFX3000-G | 1024                          |
| QFX3000-M | 1024                          |

**Related Documentation** • [Example: Configuring Routing Between VLANs on One Switch on page 2126](#)

## Example: Configuring Routing Between VLANs on One Switch

To segment traffic on a LAN into separate broadcast domains, you create separate virtual LANs (VLANs). For example, you might want to create a VLAN that includes the employees in a department and the resources that they use often, such as printers, servers, and so on.

Of course, you also want to allow these employees to communicate with people and resources in other VLANs. To forward packets between VLANs you normally you need a router that connects the VLANs. However, you can accomplish this on a Juniper Networks switch without using a router by configuring an integrated routing and bridging (IRB) interface (also known as a routed VLAN interface—or RVI—in versions of Junos OS that do not support Enhanced Layer 2 Software). Using this approach reduces complexity and avoids the costs associated with purchasing, installing, managing, powering, and cooling another device.

- [Requirements on page 2126](#)
- [Overview and Topology on page 2126](#)
- [Configure Layer 2 switching for two VLANs on page 2127](#)
- [Verification on page 2130](#)

### Requirements


This example uses the following hardware and software components:

- One switch
- Junos OS Release 11.1 or later

### Overview and Topology

This example uses an IRB to route traffic between two VLANs on the same switch. The topology is shown in [Figure 23](#).

**Figure 23: IRB with One Switch**



This example shows a simple configuration to illustrate the basic steps for creating two VLANs on a single switch and configuring an IRB to enable routing between the VLANs. One VLAN, called **blue**, is for the sales and marketing group, and a second, called **red**, is for the customer support team. The sales and support groups each have their own file servers and wireless access points. Each VLAN must have a unique name, tag (VLAN ID), and distinct IP subnet. [Table 187](#) lists the components of the sample topology.

Table 187: Components of the Multiple VLAN Topology

| Property                      | Settings                                                                                                                        |
|-------------------------------|---------------------------------------------------------------------------------------------------------------------------------|
| VLAN names and tag IDs        | blue, ID 100<br>red, ID 200                                                                                                     |
| Subnets associated with VLANs | blue: 192.0.2.0/25 (addresses 192.0.2.1 through 192.0.2.126)<br>red: 192.0.2.128/25 (addresses 192.0.2.129 through 192.0.2.254) |
| Interfaces in VLAN blue       | Sales server port: xe-0/0/4<br>Sales wireless access points: xe-0/0/6                                                           |
| Interfaces in VLAN red        | Support server port: xe-0/0/0<br>Support wireless access points: xe-0/0/2                                                       |
| IRB name                      | interface irb                                                                                                                   |
| IRB units and addresses       | logical unit 100: 192.0.2.1/25<br>logical unit 200: 192.0.2.129/25                                                              |

This configuration example creates two IP subnets, one for the blue VLAN and the second for the red VLAN. The switch bridges traffic within the VLANs. For traffic passing between two VLANs, the switch routes the traffic using an IRB on which you have configured addresses in each IP subnet.

To keep the example simple, the configuration steps show only a few interfaces and VLANs. Use the same configuration procedure to add more interfaces and VLANs. By default, all interfaces are in access mode, so you do not have to configure the port mode.

## Configure Layer 2 switching for two VLANs

**CLI Quick Configuration** To quickly configure Layer 2 switching for the two VLANs (**blue** and **red**) and to quickly configure Layer 3 routing of traffic between the two VLANs, copy the following commands and paste them into the switch terminal window:



**NOTE:** The following example uses a version of Junos OS that supports Enhanced Layer 2 Software (ELS). When you use ELS, you create a Layer 3 virtual interface named **irb**. If you are using a version of Junos OS that does not support ELS, you create a Layer 3 virtual interface named **vlan**.

```
[edit]
set interfaces xe-0/0/4 unit 0 description "Sales server port"
set interfaces xe-0/0/4 unit 0 family ethernet-switching vlan members blue
set interfaces xe-0/0/6 unit 0 description "Sales wireless access point port"
set interfaces xe-0/0/6 unit 0 family ethernet-switching vlan members blue
set interfaces xe-0/0/0 unit 0 description "Support servers"
set interfaces xe-0/0/0 unit 0 family ethernet-switching vlan members red
set interfaces xe-0/0/2 unit 0 description "Support wireless access point port"
set interfaces xe-0/0/2 unit 0 family ethernet-switching vlan members red
set interfaces irb unit 100 family inet address 192.0.2.1/25
```

```
set interfaces irb unit 200 family inet address 192.0.2.129/25
set vlans blue l3-interface irb.100
set vlans blue vlan-id 100
set vlans red vlan-id 200
set vlans red l3-interface irb.200
```

**Step-by-Step  
Procedure**

To configure the switch interfaces and the VLANs to which they belong:

1. Configure the interface for the sales server in the blue VLAN:

```
[edit interfaces xe-0/0/4 unit 0]
user@switch# set description "Sales server port"
user@switch# set family ethernet-switching vlan members blue
```

2. Configure the interface for the wireless access point in the blue VLAN:

```
[edit interfaces xe-0/0/6 unit 0]
user@switch# set description "Sales wireless access point port"
user@switch# set family ethernet-switching vlan members blue
```

3. Configure the interface for the support server in the red VLAN:

```
[edit interfaces xe-0/0/0 unit 0]
user@switch# set description "Support server port"
user@switch# set family ethernet-switching vlan members red
```

4. Configure the interface for the wireless access point in the red VLAN:

```
[edit interfaces xe-0/0/2 unit 0]
user@switch# set description "Support wireless access point port"
user@switch# set family ethernet-switching vlan members red
```

**Step-by-Step  
Procedure**

Now create the VLANs and the IRB. The IRB will have logical units in the broadcast domains of both VLANs.

1. Create the red and blue VLANs by configuring the VLAN IDs for them:

```
[edit vlans]
user@switch# set blue vlan-id 100
user@switch# set red vlan-id 200
```

2. Create the interface named **irb** with a logical unit in the sales broadcast domain (blue VLAN):

```
[edit interfaces]
user@switch# set irb unit 100 family inet address 192.0.2.1/25
```

The unit number is arbitrary and does not have to match the VLAN tag ID. However, configuring the unit number to match the VLAN ID can help avoid confusion.

3. Add a logical unit in the support broadcast domain (red VLAN) to the **irb** interface:

```
[edit interfaces]
user@switch# set irb unit 200 family inet address 192.0.2.129/25
```

4. Complete the IRB configuration by binding the red and blue VLANs (Layer 2) with the appropriate logical units of the **irb** interface (Layer 3):

```
[edit vlans]
user@switch# set blue l3-interface irb.100
user@switch# set red l3-interface irb.200
```

**Configuration Results**

Display the results of the configuration:

```
user@switch> show configuration
interfaces {
```



```

xe-0/0/4 {
 unit 0 {
 description "Sales server port";
 family ethernet-switching {
 vlan members blue;
 }
 }
}
xe-0/0/6 {
 unit 0 {
 description "Sales wireless access point port";
 family ethernet-switching {
 vlan members blue;
 }
 }
}
xe-0/0/0 {
 unit 0 {
 description "Support server port";
 family ethernet-switching {
 vlan members red;
 }
 }
}
xe-0/0/2 {
 unit 0 {
 description "Support wireless access point port";
 family ethernet-switching {
 vlan members red;
 }
 }
}
irb {
 unit 100 {
 family inet address 192.0.2.1/25;
 }
 unit 200 {
 family inet address 192.0.2.129/25;
 }
}
}
vlands {
 blue {
 vlan-id 100;
 interface xe-0/0/4.0;
 interface xe-0/0/6.0;
 l3-interface irb 100;
 }
 red {
 vlan-id 200;
 interface xe-0/0/0.0;
 interface xe-0/0/2.0;
 l3-interface irb 200;
 }
}
}

```



**TIP:** To quickly configure the blue and red VLAN interfaces, issue the **load merge terminal** command, copy the hierarchy, and paste it into the switch terminal window.

## Verification

To verify that the **blue** and **red** VLANs have been created and are operating properly, perform these tasks:

- [Verifying That the VLANs Have Been Created and Associated with the Correct Interfaces on page 2130](#)
- [Verifying That Traffic Can Be Routed Between the Two VLANs on page 2130](#)

### Verifying That the VLANs Have Been Created and Associated with the Correct Interfaces

**Purpose** Verify that the VLANs **blue** and **red** have been created on the switch and that all connected interfaces on the switch are members of the correct VLAN.

**Action** List all VLANs configured on the switch:

```
user@switch> show vlans
Name Tag Interfaces
default 0 xe-0/0/0.0, xe-0/0/2.0, xe-0/0/4.0, xe-0/0/6.0,
blue 100 xe-0/0/4.0, xe-0/0/6.0,
red 200 xe-0/0/0.0, xe-0/0/2.0, *
mgmt * me0.0*
```

**Meaning** The **show vlans** command lists all VLANs configured on the switch and which interfaces are members of each VLAN. This command output shows that the **blue** and **red** VLANs have been created. The **blue** VLAN has a tag ID of 100 and is associated with interfaces **xe-0/0/4.0** and **xe-0/0/6.0**. VLAN **red** has a tag ID of 200 and is associated with interfaces **xe-0/0/0.0** and **xe-0/0/2.0**.

### Verifying That Traffic Can Be Routed Between the Two VLANs

**Purpose** Verify routing between the two VLANs.

**Action** Verify that the IRB logical units are up:

```
user@switch> show interfaces terse
irb.100 up up inet 192.0.2.1/25
irb.200 up up inet 192.0.2.129/25
```



**NOTE:** At least one port (access or trunk) with an appropriate VLAN assigned to it must be up for the irb interface to be up.

Verify that switch has created routes that use the IRB logical units:

```
user@switch> show route
192.0.2.0/25 *[Direct/0] 1d 03:26:45
 > via irb.100
192.0.2.1/32 *[Local/0] 1d 03:26:45
 Local via irb.100
192.0.2.128/25 *[Direct/0] 1d 03:26:45
 > via irb.200
192.0.2.129/32 *[Local/0] 1d 03:26:45
 Local via irb.200
```

List the Layer 3 routes in the switch's Address Resolution Protocol (ARP) table:

```
user@switch> show arp
```

| MAC Address       | Address     | Name    | Flags |
|-------------------|-------------|---------|-------|
| 00:00:0c:06:2c:0d | 192.0.2.7   | irb.100 | None  |
| 00:13:e2:50:62:e0 | 192.0.2.132 | irb.200 | None  |

**Meaning** The output of the **show interfaces** and **show route** commands show that the Layer 3 IRB logical units are working and the switch has used them to create direct routes that it will use to forward traffic between the VLAN subnets. The **show arp** command displays the mappings between the IP addresses and MAC addresses for devices on both **irb.100** (associated with VLAN **blue**) and **irb.200** (associated with VLAN **red**). These two devices can communicate.

**Related Documentation**

- [Understanding Integrated Routing and Bridging on page 2124](#)
- [irb \(Interfaces\) on page 2292](#)
- [l3-interface on page 2295](#)

## Configuring IRB Interfaces

Integrated routing and bridging (IRB) interfaces enable a switch to recognize which packets are being sent to local addresses so that they are bridged whenever possible and are routed only when needed. Whenever packets can be switched instead of routed, several layers of processing are eliminated. Switching also reduces the number of address look-ups.



**NOTE:** In versions of Junos OS that do not support Enhanced Layer 2 Software (ELS), this type of interface is called a routed VLAN interface (RVI).

To configure the routed VLAN interface:

1. Create the VLAN by assigning it a name and a VLAN ID:  

```
[edit]
user@switch# set vlans support vlan-id 111
```
2. Assign an interface to the VLAN by specifying the logical interface (with the **unit** statement) and specifying the VLAN name as the member:

```
[edit]
user@switch# set interfaces ge-0/0/18 unit 0 family ethernet-switching vlan members
support
```

3. Create the subnet for the VLAN's broadcast domain:

```
[edit]
user@switch# set interfaces irb unit 111 family inet address 111.111.111.1/24
```

4. Bind a Layer 3 interface with the VLAN:

```
[edit]
user@switch# set vlans support l3-interface irb.111
```



**NOTE:** If you are using a version of Junos OS that does not support ELS, you create a Layer 3 virtual interface named `vlan`



**NOTE:** Layer 3 interfaces on trunk ports allow the interface to transfer traffic between multiple VLANs. Within a VLAN, traffic is bridged, while across VLANs, traffic is routed.

You can display the configuration settings:

```
user@switch> show interfaces irb terse
Interface Admin Link Proto Local Remote
vlan up up
irb.111 up up inet 111.111.111.1/24
```

```
user@switch> show vlans
Name Tag Interfaces
default
employee-vlan 20 ge-1/0/0.0, ge-1/0/1.0, ge-1/0/2.0
marketing 40 ge-1/0/10.0, ge-1/0/20.0, ge-1/0/30.0
support 111 ge-0/0/18.0
mgmt bme0.32769, bme0.32771*
```

```
user@switch> show ethernet-switching table
Ethernet-switching table: 1 entries, 0 learned
VLAN MAC address Type Age Interfaces
support 00:19:e2:50:95:a0 Static - Router
```

**Related Documentation**

- [Understanding Integrated Routing and Bridging on page 2124](#)

## Adding a Static MAC Address Entry to the Ethernet Switching Table (CLI Procedure)



**NOTE:** This task uses Junos OS for EX Series switches and Junos OS for QFX3500 and QFX3600 switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Adding a Static MAC Address Entry to the Ethernet Switching Table (CLI Procedure)*. For ELS details, see “Getting Started with Enhanced Layer 2 Software” on page 41.

The Ethernet switching table, also known as the forwarding table, specifies the known locations of VLAN nodes and the addresses of devices within those nodes. There are two ways to populate the Ethernet switching table on a switch. The easiest method is to let the switch update the table with MAC addresses.

The second way to populate the Ethernet switching table is to manually insert addresses into the table. You can do this to reduce flooding and speed up the switch’s automatic learning process.

Before configuring a static MAC address, be sure that you have:

- Set up the VLAN. See *Configuring VLANs for EX Series Switches (CLI Procedure)*.

To configure an interface to have a static MAC address:

```
[edit vlans vlan-name switch-options interface interface-name]
user@switch# set static-mac mac-address
```

### Related Documentation

- *Understanding Bridging and VLANs on EX Series Switches*

## Configuring Static ARP Entries

You can create static ARP table entries, which are explicit mappings between IP addresses and MAC addresses.

- To configure a static ARP entry:

```
[edit interfaces interface-name unit logical-unit-number family inet address
address]
user@switch# set arp ip-address (mac | multicast-mac) mac-address
```

The IP address that you specify must be part of the subnet defined in the enclosing **address** statement.

To associate a multicast MAC address with a unicast IP address, use the **multicast-mac** statement.

Specify the MAC address as 6 hexadecimal bytes in one of the following formats: *nnnn.nnnn.nnnn* or *nn:nn:nn:nn:nn:nn*; for example, 0011.2233.4455 or 00:11:22:33:44:55.

- Related Documentation**
- [Understanding Static ARP Entries on page 6041](#)
  - *arp*

## Troubleshooting Ethernet Switching

---

**Problem**    **Description:** Sometimes a MAC address entry in the switch's Ethernet switching table is not updated after the device with that MAC address has been moved from one interface to another on the switch. Typically, the switch does not wait for a MAC address expiration when a MAC move operation occurs. As soon as the switch detects the MAC address on the new interface, it immediately updates the table. Many network devices send a gratuitous ARP packet when switching an IP address from one device to another. The switch updates its ARP cache table after receipt of such gratuitous ARP messages, and then it also updates its Ethernet switching table.

Sometimes silent devices, such as syslog servers or SNMP trap receivers that receive UDP traffic but do not return acknowledgment (ACK) messages to the traffic source, fail to send gratuitous ARP packets when a device moves. If such a move occurs when the system administrator is not available to explicitly clear the affected interfaces by issuing the **clear ethernet-switching table** command, the entry for the moved device in the Ethernet switching table is not updated.

**Solution**    Set up the switch to handle unattended MAC address switchovers.

1. Reduce the system-wide ARP aging timer. (By default, the ARP aging timer is set at 20 minutes. The range of the ARP aging timer is from 1 through 240 minutes.)

```
[edit system arp]
user@switch# set aging-timer 3
```

2. Set the MAC aging timer to the same value as the ARP timer. (By default, the MAC aging timer is set to 300 seconds. The range is 60 to 1,000,000 seconds.)

```
[edit protocols 12-learning]
user@switch# set global-mac-table-aging-time 180
```

The ARP entry and the MAC address entry for the moved device expire within the times specified by the aging timer values. After the entries expire, the switch sends a new ARP message to the IP address of the device. The device responds to the ARP message, thereby refreshing the entries in the switch's ARP cache table and Ethernet switching table.

- Related Documentation**
- *arp*
  - [global-mac-table-aging-time on page 2315](#)

## PART 30

# MAC Addresses

- [Using MAC Addresses on page 2137](#)





## CHAPTER 74

# Using MAC Addresses

- [Introduction to the Media Access Control \(MAC\) Layer 2 Sublayer on page 2137](#)
- [Understanding MAC Learning on page 2138](#)
- [Disabling MAC Learning on page 2138](#)
- [Example: Disabling MAC Learning on page 2139](#)
- [Configuring MAC Notification \(CLI Procedure\) on page 2140](#)
- [Verifying That MAC Notification Is Working Properly on page 2141](#)
- [Configuring MAC Limiting \(CLI Procedure\) on page 2142](#)
- [Configuring MAC Table Aging on page 2144](#)

### Introduction to the Media Access Control (MAC) Layer 2 Sublayer

---

This topic provides an introduction to the MAC sublayer of the data link layer (Layer 2).

In Layer 2 of a network, the Media Access Control (MAC) sublayer provides addressing and channel access control mechanisms that enable several terminals or network nodes to communicate in a network.

The MAC sublayer acts as an interface between the logical link control (LLC) Ethernet sublayer and Layer 1 (the physical layer). The MAC sublayer emulates a full-duplex logical communication channel in a multipoint network. This channel may provide unicast, multicast, or broadcast communication service. The MAC sublayer uses MAC protocols to prevent collisions.

In Layer 2, multiple devices on the same physical link can uniquely identify one another at the data link layer, by using the MAC addresses that are assigned to all ports on a switch. A MAC algorithm accepts as input a secret key and an arbitrary-length message to be authenticated, and outputs a MAC address.

A MAC address is a 12-digit hexadecimal number (48 bits in long). MAC addresses are usually written in one of these formats:

- MM:MM:MM:SS:SS:SS
- MM-MM-MM-SS-SS-SS

The first half of a MAC address contains the ID number of the adapter manufacturer. These IDs are regulated by an Internet standards body. The second half of a MAC address represents the serial number assigned to the adapter by the manufacturer.

Contrast MAC addressing, which works at Layer 2, with IP addressing, which runs at Layer 3 (networking and routing). One way to remember the difference is that the MAC addresses apply to a physical or virtual node, whereas IP addresses apply to the software implementation of that node. MAC addresses are typically fixed on a per-node basis, whereas IP addresses change when the node moves from one part of the network to another.

IP networks maintain a mapping between the IP and MAC addresses of a node using the Address Resolution Protocol (ARP) table. DHCP also typically uses MAC addresses when assigning IP addresses to nodes.

- Related Documentation**
- [Overview of Layer 2 Networking on page 2085](#)
  - [Understanding MAC Learning on page 2138](#)

---

## Understanding MAC Learning

*MAC learning* is the process of obtaining the MAC addresses of all the nodes on a network.

When a node is first connected to an Ethernet LAN or VLAN, it has no information about the other nodes on the network. As data is sent through the network, data packets include a data frame listing their source and destination MAC addresses. The data frame is forwarded to a target port, which is connected to the second device. The MAC address is learned locally at the target port, which facilitates communications for frames that later enter the target port and contain addresses previously learned from a received frame.

By default, MAC learning is enabled on the QFX Series.

- Related Documentation**
- [Introduction to the Media Access Control \(MAC\) Layer 2 Sublayer on page 2137](#)
  - [Overview of Layer 2 Networking on page 2085](#)

---

## Disabling MAC Learning

By default, MAC learning is globally enabled on all node. This topic describes how to disable MAC learning, as well as how to reenable and verify that MAC learning has been enabled or disabled.



**NOTE:** This task supports the Enhanced Layer 2 Software (ELS) configuration style. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 41](#) If your switch runs software that does not support ELS, see *Disabling MAC Learning*.

---

Disabling dynamic MAC learning prevents a node from learning source and destination MAC addresses.

- To disable MAC learning:

```
[edit vlans vlan-name switch-options interface interface-name]
user@switch# set no-mac-learning
```

- To enable MAC learning:

```
[edit vlans vlan-name switch-options interface interface-name]
user@switch# delete no-mac-learning
user@switch# deactivate no-mac-learning
```

- To verify the status of MAC learning, view the Ethernet MAC learning statistics in operational mode.

```
user@switch> show ethernet-switching table
Ethernet-switching table: 2 entries, 1 learned
 VLAN MAC address Type Age Interfaces
 default * Flood - All-members
 default 00:1f:12:39:90:80 Learn 29 xe-0/0/0.0
```

#### Related Documentation

- [Understanding MAC Learning on page 2138](#)
- [Example: Disabling MAC Learning on page 2139](#)
- *no-mac-learning*

## Example: Disabling MAC Learning

By default, MAC learning is enabled on the QFX Series. This topic provides examples for disabling, enabling, and verifying the operation of MAC learning on the QFX Series. These examples require that you be logged in as the root user to the switch on which you wish to modify MAC learning.



**NOTE:** This task uses Junos OS for QFX3500, QFX3600, EX4600, QFX5100, and QFX10002 switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Example: Disabling MAC Learning*. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 41](#).

- To disable MAC learning in a VLAN:

```
[edit]
user@switch# set vlans vlan10 switch-options interface xe-0/0/0.0 no-mac-learning
```

- To reenabling MAC learning:

```
[edit] vlans vlan10 switch-options interface xe-0/0/0.0
user@switch# delete no-mac-learning
```

- To verify the status of MAC learning on the QFX Series:

```
user@switch> show ethernet-switching table
Learning stats: 10 learn msg rcvd, 2 error, 0 forced update
Interface Local pkts Transit pkts Error
xe-0/0/0.0 0 6 1
```

|             |   |   |   |
|-------------|---|---|---|
| xe-0/0/22.0 | 0 | 0 | 0 |
| xe-0/0/1.0  | 0 | 4 | 1 |
| xe-0/0/2.0  | 0 | 0 | 0 |
| xe-0/0/3.0  | 0 | 0 | 0 |
| xe-0/0/4.0  | 0 | 0 | 0 |
| xe-0/0/19.0 | 0 | 0 | 0 |
| xe-0/0/18.0 | 0 | 0 | 0 |
| xe-0/0/9.0  | 0 | 0 | 0 |

- Related Documentation**
- [Understanding MAC Learning on page 2138](#)
  - [Disabling MAC Learning on page 2138](#)

## Configuring MAC Notification (CLI Procedure)



**NOTE:** This task uses the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring MAC Notification (CLI Procedure)* or *Configuring MAC Notification*. For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on page 41.

When a switch learns or unlearns a MAC address, SNMP notifications can be sent to the network management system at regular intervals to record the addition or removal of the MAC address. This process is known as MAC notification.

The MAC notification interval defines how often Simple Network Management Protocol (SNMP) notifications logging the addition or removal of MAC addresses on the switch are sent to the network management system.

MAC notification is disabled by default. When MAC notification is enabled, the default MAC notification interval is 30 seconds.

To enable or disable MAC notification, or to set the MAC notification interval, perform these tasks:

- [Enabling MAC Notification on page 2140](#)
- [Disabling MAC Notification on page 2141](#)
- [Setting the MAC Notification Interval on page 2141](#)

### Enabling MAC Notification

MAC notification is disabled by default. You need to perform this procedure to enable MAC notification.

To enable MAC notification on the switch with the default MAC notification interval of 30 seconds:

```
[edit switch-options]
user@switch# set mac-notification
```

To enable MAC notification on the switch with any other MAC notification interval (here, the MAC notification interval is set to 60 seconds):

```
[edit switch-options]
user@switch# set mac-notification notification-interval 60
```

## Disabling MAC Notification

MAC notification is disabled by default. Perform this procedure only if MAC notification was previously enabled on your switch.

To disable MAC notification on the switch:

```
[edit switch-options]
user@switch# delete mac-notification
```

To disable MAC notification on a specific interface (here, the interface is ge-0/0/3):

```
[edit switch-options]
user@switch# set interface ge-0/0/3 no-mac-notification
```

## Setting the MAC Notification Interval

The default MAC notification interval is 30 seconds. The procedure to change the MAC notification interval to a different interval is identical to the procedure to enable MAC notification on the switch with a nondefault value for the MAC notification interval.

To set the MAC notification interval on the switch (here, the MAC notification interval is set to 5 seconds):

```
[edit switch-options]
user@switch# set mac-notification notification-interval 5
```

**Related Documentation**

- *Verifying That MAC Notification Is Working Properly*

## Verifying That MAC Notification Is Working Properly

**Purpose** Verify that MAC notification is enabled or disabled, and that the MAC notification interval is set to the specified value.

**Action** To verify that MAC notification is enabled or disabled and also to verify the MAC notification interval setting.

```
user@switch> show ethernet-switching mac-notification
Notification Status: Enabled
Notification Interval: 60
Notifications Sent : 0
Notifications Table Maxsize : 256
```

**Meaning** The output in the **Notification Status** field shows that MAC notification is enabled. The output in the **Notification Status** field would display **Disabled** if MAC notification was disabled.

The **Notification Interval** field output shows that the MAC notification interval is set to 60 seconds.

- Related Documentation**
- [Configuring MAC Notification](#)
  - [Configuring MAC Notification \(CLI Procedure\) on page 2140](#)

---

## Configuring MAC Limiting (CLI Procedure)

---



**NOTE:** This task uses Junos OS for EX Series switches and QFX3500 and QFX3600 switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see [Configuring MAC Limiting \(CLI Procedure\)](#). For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 41](#).

This topic describes various ways of configuring a limitation on MAC addresses in packets that are received and forwarded by the switch.



**NOTE:** On a QFX Series Virtual Chassis, if you include the shutdown option at the [edit vlans *vlan-name* switch-options interface *interface-name* interface-mac-limit packet-action] hierarchy level and issue the commit operation, the system generates a commit error. The system does not generate an error if you include the shutdown option at the [edit switch-options interface *interface-name* interface-mac-limit packet-action] hierarchy level.

The different ways of setting a MAC limit are described in the following sections:

- [Limiting the Number of MAC Addresses Learned by an Interface on page 2143](#)
- [Limiting the Number of MAC Addresses Learned by a VLAN on page 2143](#)

## Limiting the Number of MAC Addresses Learned by an Interface

To secure a port, you can set the maximum number of MAC addresses that can be learned by an interface:

- Set the MAC limit on an interface, and specify an action that the switch takes after the specified limit is exceeded:

```
[edit switch-options]
user@switch# set interface interface-name interface-mac-limit limit packet-action
action
```

After you set a new MAC limit for the interface, the system clears existing entries in the MAC address forwarding table associated with the interface.

## Limiting the Number of MAC Addresses Learned by a VLAN

To limit the number of MAC addresses learned by a VLAN, perform both of the following steps:

- Set the maximum number of MAC addresses that can be learned by a VLAN, and specify an action that the switch takes after the specified limit is exceeded:

```
[edit vlans]
user@switch# set vlan-name switch-options mac-table-size limit packet-action
action
```

- Set the maximum number of MAC addresses that can be learned by one or all interfaces in the VLAN, and specify an action that the switch takes after the specified limit is exceeded:

```
[edit vlans]
user@switch# set vlan-name switch-options interface interface-name
interface-mac-limit limit packet-action action
[edit vlans]
user@switch# set vlan-name switch-options interface-mac-limit limit packet-action
action
```



**NOTE:** If you specify a MAC limit and packet action for all interfaces in the VLAN *and* a specific interface in the VLAN, the MAC limit and packet action specified at the specific interface level takes precedence. Also, at the VLAN interface level, only the drop and drop-and-log options are supported.

After you set new MAC limits for a VLAN by using the **mac-table-size** statement or for interfaces associated with a VLAN by using the **interface-mac-limit** statement, the system clears the corresponding existing entries in the MAC address forwarding table.

### Related Documentation

- [Understanding Bridging and VLANs on EX Series Switches](#)
- [Configuring Persistent MAC Learning \(CLI Procedure\) on page 6035](#)

## Configuring MAC Table Aging

---

MAC table aging ensures that a switch tracks only active nodes on the network and that it is able to flush out network nodes that are no longer available.

To manage MAC entries more efficiently, you can configure the maximum time that entries can remain in the MAC address table before being deleted.



**NOTE:** This task uses Junos OS for QFX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 41](#).

You can use the **global-mac-table-aging-time** command to configure how long entries remain in the Ethernet switching table before expiring, as follows:

```
[edit protocols 12-learning]
user@switch# set global-mac-table-aging-time 200
```



**NOTE:** This command applies to all VLANs configured for the switch. You cannot configure separate MAC table aging times for specific VLANs.

### Related Documentation

- [Understanding Bridging and VLANs on page 2089](#)
- [Example: Setting Up Bridging with Multiple VLANs on page 2114](#)



## PART 31

# Spanning Trees

- [Using Spanning Trees on page 2147](#)



## CHAPTER 75

# Using Spanning Trees

- Overview of Spanning-Tree Protocols on page 2148
- Understanding MSTP on page 2149
- Example: Configuring Network Regions for VLANs with MSTP on page 2149
- Configuring MSTP on page 2170
- Understanding RSTP on page 2172
- Example: Configuring Faster Convergence and Improved Network Stability with RSTP on page 2173
- Configuring RSTP (CLI Procedure) on page 2190
- Understanding VSTP on page 2191
- Example: Configuring VSTP on QFX Series Switches and EX4600 Switches on page 2191
- Configuring VSTP (CLI Procedure) on page 2197
- Understanding BPDU Protection for STP, RSTP, and MSTP on page 2199
- Example: Configuring BPDU Protection on Edge Interfaces to Prevent STP Miscalculations on page 2200
- Configuring BPDU Protection on Spanning Tree Interfaces on page 2205
- Understanding Loop Protection for STP, RSTP, VSTP, and MSTP on page 2206
- Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on page 2207
- Understanding Root Protection for STP, RSTP, VSTP, and MSTP on page 2211
- Example: Configuring Root Protection to Enforce Root Bridge Placement in Spanning Trees on page 2212
- Unblocking an Interface That Receives BPDUs in Error (CLI Procedure) on page 2217

## Overview of Spanning-Tree Protocols

---

QFX Series switches provide Layer 2 loop prevention through Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), Multiple Spanning Tree Protocol (MSTP), and VLAN Spanning Tree Protocol (VSTP). The default spanning-tree protocol on the QFX Series is RSTP. RSTP provides faster convergence times than STP. However, some legacy networks require the slower convergence times of basic STP.

The STP support provided for the QFX Series includes:

- IEEE 802.1d
- 802.1w RSTP
- 802.1s MSTP

If your network includes IEEE 802.1D 1998 bridges, you can remove RSTP and explicitly configure STP. When you explicitly configure STP, the QFX Series products use the IEEE 802.1D 2004 specification, force version 0. This configuration runs a version of RSTP that is compatible with the classic, basic STP. If you use virtual LANs (VLANs), you should enable VSTP and use it on your network. See [“Understanding VSTP” on page 2191](#).

You can use the same operational commands (**show spanning-tree bridge** and **show spanning-tree interface**) to check the status of your spanning-tree configuration, regardless of which spanning-tree protocol has been configured.

STP uses bridge protocol data unit (BPDU) packets to exchange information with other switches. BPDUs send hello packets out at regular intervals to exchange information across bridges and detect loops in a network topology. There are two types of BPDUs:

- Configuration BPDUs—These BPDUs contain configuration information about the transmitting switch and its ports, including switch and port MAC addresses, switch priority, port priority, and port cost.
- Topology change notification (TCN) BPDUs—When a bridge needs to signal a topology change, it starts to send TCNs on its root port. The designated bridge receives the TCN, acknowledges it, and generates another one for its own root port. The process continues until the TCN reaches the root bridge.

STP uses the information provided by the BPDUs to elect a root bridge, identify root ports for each switch, identify designated ports for each physical LAN segment, and prune specific redundant links to create a loop-free tree topology. All leaf devices calculate the best path to the root device and place their ports in blocking or forwarding states based on the best path to the root. The resulting tree topology provides a single active Layer 2 data path between any two end stations.

## Understanding Spanning Tree Protocols on a QFabric System

Although there is no need to run STP in a QFabric system, you can connect a QFabric system to another Layer 2 device and use STP. STP traffic can only be processed on network Node groups. Other Node groups, such as redundant server Node groups and server Node groups, discard the STP bridge protocol data units (BPDUs) traffic and disable the

interface automatically. Server Node groups only process host-facing protocols, whereas Network Node groups process all supported protocols.

**Related Documentation**

- [Understanding BPDU Protection for STP, RSTP, and MSTP on page 2199](#)
- [Understanding MSTP on page 2149](#)
- [Understanding RSTP on page 2172](#)
- [Understanding VSTP on page 2191](#)

## Understanding MSTP

Although RSTP provides faster convergence time than STP does, it still does not solve a problem inherent in STP: all VLANs within a LAN must share the same spanning tree. To solve this problem, the QFX Series products use Multiple Spanning Tree Protocol (MSTP) to create a loop-free topology in networks with multiple spanning-tree regions.

An MSTP region allows a group of bridges to be modeled as a single bridge. An MSTP region contains multiple spanning-tree instances (MSTIs). MSTIs provide different paths for different VLANs. This functionality facilitates more efficient load sharing across redundant links.

An MSTP region can support up to 64 MSTIs, and each instance can support from 1 through 4094 VLANs.

**Related Documentation**

- [Overview of Spanning-Tree Protocols on page 2148](#)
- [Understanding RSTP on page 2172](#)
- [Example: Configuring Network Regions for VLANs with MSTP](#)
- [Example: Configuring Network Regions for VLANs with MSTP on page 2149](#)
- [Configuring MSTP on page 2170](#)

## Example: Configuring Network Regions for VLANs with MSTP



**NOTE:** This example uses Junos OS for EX Series and QFX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Example: Configuring Network Regions for VLANs with MSTP on EX Series Switches*. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 41](#).

Multiple Spanning Tree Protocol (MSTP) is used to create a loop-free topology in networks using multiple spanning-tree regions in which each region contains multiple spanning-tree instances (MSTIs). MSTIs provide different paths for different VLANs. This functionality facilitates better load sharing across redundant links.

Up to 64 MSTIs can be created for an EX Series switch, and each MSTI can support up to 4094 VLANs.

This example describes how to configure MSTP on four EX Series switches:

- [Requirements on page 2150](#)
- [Overview and Topology on page 2150](#)
- [Configuring MSTP on Switch 1 on page 2153](#)
- [Configuring MSTP on Switch 2 on page 2156](#)
- [Configuring MSTP on Switch 3 on page 2158](#)
- [Configuring MSTP on Switch 4 on page 2161](#)
- [Verification on page 2163](#)

## Requirements

This example uses the following software and hardware components:

- Junos OS Release 13.2X50-D10 or later for EX Series switches
- Four EX Series switches


Before you configure the switches for MSTP, be sure you have:

- Installed and connected the four switches. See the hardware documentation for your switch.
- Performed the initial software configuration on all switches. See *Connecting and Configuring an EX Series Switch (CLI Procedure)* or *Connecting and Configuring an EX Series Switch (J-Web Procedure)*.

## Overview and Topology

When the number of VLANs grows in a network, MSTP provides an efficient way of creating a loop-free topology by using MSTIs. Each MSTI in the spanning-tree domain maintains its own tree. Each tree can be mapped to different links, utilizing bandwidth that would be unavailable to a single tree. MSTIs reduce the demand on system resources.

Figure 24: Network Topology for MSTP



The interfaces shown in Figure 24 will be configured for MSTP.



**NOTE:** You can configure MSTP only on physical interfaces, not on logical interfaces.

Table 188: Components of the Topology for Configuring MSTP on EX Series Switches

| Property | Settings                                                                                                                                                                                                                                             |
|----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Switch 1 | <p>The following interfaces on Switch 1 are connected in this way:</p> <ul style="list-style-type: none"><li>• ge-0/0/9 is connected to Switch 2</li><li>• ge-0/0/13 is connected to Switch 4</li><li>• ge-0/0/11 is connected to Switch 3</li></ul> |
| Switch 2 | <p>The following interfaces on Switch 2 are connected in this way:</p> <ul style="list-style-type: none"><li>• ge-0/0/14 is connected to Switch 1</li><li>• ge-0/0/18 is connected to Switch 3</li></ul>                                             |

Table 188: Components of the Topology for Configuring MSTP on EX Series Switches (*continued*)

| Property               | Settings                                                                                                                                                                                                                                                                       |
|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Switch 3               | <p>The following interfaces on Switch 3 are connected in this way:</p> <ul style="list-style-type: none"> <li>• <b>ge-0/0/26</b> is connected to Switch 1</li> <li>• <b>ge-0/0/28</b> is connected to Switch 2</li> <li>• <b>ge-0/0/24</b> is connected to Switch 4</li> </ul> |
| Switch 4               | <p>The following interfaces on Switch 4 are connected in this way:</p> <ul style="list-style-type: none"> <li>• <b>ge-0/0/19</b> is connected to Switch 1</li> <li>• <b>ge-0/0/23</b> is connected to Switch 3</li> </ul>                                                      |
| VLAN names and tag IDs | <b>voice-vlan</b> , tag 10<br><b>employee-vlan</b> , tag 20<br><b>guest-vlan</b> , tag 30<br><b>camera-vlan</b> , tag 40                                                                                                                                                       |
| MSTIs                  | 1<br>2                                                                                                                                                                                                                                                                         |
| MSTI region            | <b>region1</b>                                                                                                                                                                                                                                                                 |

The topology in [Figure 24](#) shows a common and internal spanning tree (CIST). The CIST is a single spanning tree connecting all devices in the network. The switch with the lowest bridge priority is elected as the root bridge of the CIST. You can control the election of the root bridge by configuring the bridge priority. Switch 3 is the root bridge of the CIST.

The ports in an MSTP topology have specific roles:

- The *root port* is responsible for forwarding data to the root bridge.
- The *alternate port* is a standby port for the root port. When a root port goes down, the alternate port becomes the active root port.
- The *designated port* forwards data to the downstream network segment or device.
- The *backup port* becomes the active designated port and starts forwarding data when the designated port goes down.

In this example, one MSTP region contains Switch 1, Switch 2, Switch 3, and Switch 4. Within the region, four VLANs are created:

- **voice-vlan** supports voice traffic and has the VLAN tag identifier of **10**.
- **employee-vlan** supports data traffic and has the VLAN tag identifier of **20**.
- **guest-vlan** supports guest VLAN traffic (for supplicants that fail authentication) and has the VLAN tag identifier of **30**.
- **camera-vlan** supports video traffic and has the VLAN tag identifier of **40**.



The VLANs are associated with specific interfaces on each of the four switches. Two MSTIs, 1 and 2, are then associated with the VLAN tag identifiers, and some MSTP parameters, such as cost, are configured on each switch.

## Configuring MSTP on Switch 1

**CLI Quick Configuration** To quickly configure interfaces and MSTP on Switch 1, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans voice-vlan description "Voice VLAN"
set vlans voice-vlan vlan-id 10
set vlans employee-vlan description "Employee VLAN"
set vlans employee-vlan vlan-id 20
set vlans guest-vlan description "Guest VLAN"
set vlans guest-vlan vlan-id 30
set vlans camera-vlan description "Camera VLAN"
set vlans camera-vlan vlan-id 40
set interfaces ge-0/0/13 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/9 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/11 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/13 unit 0 family ethernet-switching interface-mode trunk
set interfaces ge-0/0/9 unit 0 family ethernet-switching interface-mode trunk
set interfaces ge-0/0/11 unit 0 family ethernet-switching interface-mode trunk
set protocols mstp configuration-name region1
set protocols mstp bridge-priority 16k
set protocols mstp interface ge-0/0/13 cost 1000
set protocols mstp interface ge-0/0/13 mode point-to-point
set protocols mstp interface ge-0/0/9 cost 1000
set protocols mstp interface ge-0/0/9 mode point-to-point
set protocols mstp interface ge-0/0/11 cost 1000
set protocols mstp interface ge-0/0/11 mode point-to-point
set protocols mstp msti 1 bridge-priority 16k
set protocols mstp msti 1 vlan [10 20]
set protocols mstp msti 1 interface ge-0/0/11 cost 1000
set protocols mstp msti 2 bridge-priority 8k
set protocols mstp msti 2 vlan [30 40]
```

**Step-by-Step Procedure** To configure interfaces and MSTP on Switch 1:



**NOTE:** Starting with Junos OS Release 15.1 for EX Series and QFX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style, you can configure spanning tree parameters globally on all spanning tree interfaces. See [“Configuring MSTP” on page 2170](#) for additional information.

1. Configure the VLANs **voice-vlan**, **employee-vlan**, **guest-vlan**, and **camera-vlan**:

```
[edit vlans]
user@switch1# set voice-vlan description "Voice VLAN"
user@switch1# set voice-vlan vlan-id 10
user@switch1# set employee-vlan description "Employee VLAN"
user@switch1# set employee-vlan vlan-id 20
user@switch1# set guest-vlan description "Guest VLAN"
user@switch1# set guest-vlan vlan-id 30
user@switch1# set camera-vlan description "Camera VLAN"
user@switch1# set guest-vlan vlan-id 40
```

2. Configure the VLANs on the interfaces, including support for the Ethernet Switching protocol:

```
[edit interfaces]
user@switch1# set ge-0/0/13 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch1# set ge-0/0/9 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch1# set ge-0/0/11 unit 0 family ethernet-switching vlan members [10 20 30 40]
```

3. Configure the port mode for the interfaces:

```
[edit interfaces]
user@switch1# set ge-0/0/13 unit 0 family ethernet-switching interface-mode trunk
user@switch1# set ge-0/0/9 unit 0 family ethernet-switching interface-mode trunk
user@switch1# set ge-0/0/11 unit 0 family ethernet-switching interface-mode trunk
```

4. Configure MSTP on the switch, including the two MSTIs:

```
[edit protocols]
user@switch1# mstp configuration-name region1
user@switch1# mstp bridge-priority 16k
user@switch1# mstp interface ge-0/0/13 cost 1000
user@switch1# mstp interface ge-0/0/13 mode point-to-point
user@switch1# mstp interface ge-0/0/9 cost 1000
user@switch1# mstp interface ge-0/0/9 mode point-to-point
user@switch1# mstp interface ge-0/0/11 cost 1000
user@switch1# mstp interface ge-0/0/11 mode point-to-point
user@switch1# mstp msti 1 bridge-priority 16k
user@switch1# mstp msti 1 vlan [10 20]
user@switch1# mstp msti 1 interface ge-0/0/11 cost 1000
user@switch1# mstp msti 2 bridge-priority 8k
user@switch1# mstp msti 2 vlan [30 40]
```

**Results** Check the results of the configuration:

```
user@switch1> show configuration
interfaces {
 ge-0/0/13 {
 unit 0 {
 family ethernet-switching {
 interface-mode trunk;
 vlan {
 members 10;
 members 20;
 members 30;
 members 40;
 }
 }
 }
 }
 ge-0/0/9 {
 unit 0 {
 family ethernet-switching {
 interface-mode trunk;
 vlan {
 members 10;
 members 20;
 members 30;
 }
 }
 }
 }
}
```

```

 members 40;
 }
}
}
ge-0/0/11 {
 unit 0 {
 family ethernet-switching {
 interface-mode trunk;
 vlan {
 members 10;
 members 20;
 members 30;
 members 40;
 }
 }
 }
}
}
protocols {
 mstp {
 configuration-name region1;
 bridge-priority 16k;
 interface ge-0/0/13 {
 cost 1000;
 mode point-to-point;
 }
 interface ge-0/0/9 {
 cost 1000;
 mode point-to-point;
 }
 interface ge-0/0/11 {
 cost 1000;
 mode point-to-point;
 }
 }
 msti 1 {
 bridge-priority 16k;
 vlan [10 20];
 interface ge-0/0/11 {
 cost 1000;
 }
 }
 msti 2 {
 bridge-priority 8k;
 vlan [30 40];
 }
}
vllans {
 voice-vlan {
 vlan-id 10;
 }
 employee-vlan {
 vlan-id 20;
 }
 guest-vlan {
 vlan-id 30;
 }
}

```

```

 }
 camera-vlan {
 vlan-id 40;
 }
}

```

## Configuring MSTP on Switch 2

**CLI Quick Configuration** To quickly configure interfaces and MSTP on Switch 2, copy the following commands and paste them into the switch terminal window:

```

[edit]
set vlans voice-vlan description "Voice VLAN"
set vlans voice-vlan vlan-id 10
set vlans employee-vlan description "Employee VLAN"
set vlans employee-vlan vlan-id 20
set vlans guest-vlan description "Guest VLAN"
set vlans guest-vlan vlan-id 30
set vlans camera-vlan description "Camera VLAN"
set vlans camera-vlan vlan-id 40
set interfaces ge-0/0/14 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/18 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/14 unit 0 family ethernet-switching interface-mode trunk
set interfaces ge-0/0/18 unit 0 family ethernet-switching interface-mode trunk
set protocols mstp configuration-name region1
set protocols mstp bridge-priority 32k
set protocols mstp interface ge-0/0/14 cost 1000
set protocols mstp interface ge-0/0/14 mode point-to-point
set protocols mstp interface ge-0/0/18 cost 1000
set protocols mstp interface ge-0/0/18 mode point-to-point
set protocols mstp msti 1 bridge-priority 32k
set protocols mstp msti 1 vlan [10 20]
set protocols mstp msti 2 bridge-priority 4k
set protocols mstp msti 2 vlan [30 40]

```

**Step-by-Step Procedure** To configure interfaces and MSTP on Switch 2:

1. Configure the VLANs **voice-vlan**, **employee-vlan**, **guest-vlan**, and **camera-vlan**:
 

```

[edit vlans]
user@switch2# set voice-vlan description "Voice VLAN"
user@switch2# set voice-vlan vlan-id 10
user@switch2# set employee-vlan description "Employee VLAN"
user@switch2# set employee-vlan vlan-id 20
user@switch2# set guest-vlan description "Guest VLAN"
user@switch2# set guest-vlan vlan-id 30
user@switch2# set camera-vlan vlan-description "Camera VLAN"
user@switch2# set guest-vlan vlan-id 40

```
2. Configure the VLANs on the interfaces, including support for the Ethernet Switching protocol:
 

```

[edit interfaces]
user@switch2# set ge-0/0/14 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch2# set ge-0/0/18 unit 0 family ethernet-switching vlan members [10 20 30 40]

```
3. Configure the port mode for the interfaces:
 

```

[edit interfaces]
user@switch2# set ge-0/0/14 unit 0 family ethernet-switching interface-mode trunk

```

- user@switch2# set ge-0/0/18 unit 0 family ethernet-switching interface-mode trunk
4. Configure MSTP on the switch, including the two MSTIs:

```
[edit protocols]
user@switch2# mstp configuration-name region1
user@switch2# mstp bridge-priority 32k
user@switch2# mstp interface ge-0/0/14 cost 1000
user@switch2# mstp interface ge-0/0/14 mode point-to-point
user@switch2# mstp interface ge-0/0/18 cost 1000
user@switch2# mstp interface ge-0/0/18 mode point-to-point
user@switch2# mstp msti 1 bridge-priority 32k
user@switch2# mstp msti 1 vlan [10 20]
user@switch2# mstp msti 2 bridge-priority 4k
user@switch2# mstp msti 2 vlan [30 40]
```

**Results** Check the results of the configuration:

```
user@switch2> show configuration
interfaces {
 ge-0/0/14 {
 unit 0 {
 family ethernet-switching {
 interface-mode trunk;
 vlan {
 members 10;
 members 20;
 members 30;
 members 40;
 }
 }
 }
 }
 ge-0/0/18 {
 unit 0 {
 family ethernet-switching {
 interface-mode trunk;
 vlan {
 members 10;
 members 20;
 members 30;
 members 40;
 }
 }
 }
 }
}
protocols {
 mstp {
 configuration-name region1;
 bridge-priority 32k;
 interface ge-0/0/14 {
 cost 1000;
 mode point-to-point;
 }
 interface ge-0/0/18 {
 cost 1000;
 mode point-to-point;
 }
 }
}
```

```
 }
 msti 1 {
 bridge-priority 32k;
 vlan [10 20];
 }
 msti 2 {
 bridge-priority 4k;
 vlan [30 40];
 }
 }
}
vlangs {
 voice-vlan {
 vlan-id 10;
 }
 employee-vlan {
 vlan-id 20;
 }
 guest-vlan {
 vlan-id 30;
 }
 camera-vlan {
 vlan-id 40;
 }
}
```

### Configuring MSTP on Switch 3

**CLI Quick Configuration** To quickly configure interfaces and MSTP on Switch 3, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans voice-vlan description "Voice VLAN"
set vlans voice-vlan vlan-id 10
set vlans employee-vlan description "Employee VLAN"
set vlans employee-vlan vlan-id 20
set vlans guest-vlan description "Guest VLAN"
set vlans guest-vlan vlan-id 30
set vlans camera-vlan description "Camera VLAN"
set vlans camera-vlan vlan-id 40
set interfaces ge-0/0/26 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/28 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/24 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/26 unit 0 family ethernet-switching interface-mode trunk
set interfaces ge-0/0/28 unit 0 family ethernet-switching interface-mode trunk
set interfaces ge-0/0/24 unit 0 family ethernet-switching interface-mode trunk
set protocols mstp configuration-name region1
set protocols mstp bridge-priority 8k
set protocols mstp interface ge-0/0/26 cost 1000
set protocols mstp interface ge-0/0/26 mode point-to-point
set protocols mstp interface ge-0/0/28 cost 1000
set protocols mstp interface ge-0/0/28 mode point-to-point
set protocols mstp interface ge-0/0/24 cost 1000
set protocols mstp interface ge-0/0/24 mode point-to-point
set protocols mstp msti 1 bridge-priority 4k
set protocols mstp msti 1 vlan [10 20]
set protocols mstp msti 2 bridge-priority 16k
```

```
set protocols mstp msti 2 vlan [30 40]
```

### Step-by-Step Procedure

To configure interfaces and MSTP on Switch 3:

1. Configure the VLANs `voice-vlan`, `employee-vlan`, `guest-vlan`, and `camera-vlan`:

```
[edit vlans]
user@switch3# set voice-vlan description "Voice VLAN"
user@switch3# set voice-vlan vlan-id 10
user@switch3# set employee-vlan description "Employee VLAN"
user@switch3# set employee-vlan vlan-id 20
user@switch3# set guest-vlan description "Guest VLAN"
user@switch3# set guest-vlan vlan-id 30
user@switch3# set camera-vlan description "Camera VLAN"
user@switch3# set guest-vlan vlan-id 40
```

2. Configure the VLANs on the interfaces, including support for the Ethernet Switching protocol:

```
[edit interfaces]
user@switch3# set ge-0/0/26 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch3# set ge-0/0/28 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch3# set ge-0/0/24 unit 0 family ethernet-switching vlan members [10 20 30 40]
```

3. Configure the port mode for the interfaces:

```
[edit interfaces]
user@switch3# set ge-0/0/26 unit 0 family ethernet-switching interface-mode trunk
user@switch3# set ge-0/0/28 unit 0 family ethernet-switching interface-mode trunk
user@switch3# set ge-0/0/24 unit 0 family ethernet-switching interface-mode trunk
```

4. Configure MSTP on the switch, including the two MSTIs:

```
[edit protocols]
user@switch3# mstp configuration-name region1
user@switch3# mstp bridge-priority 8k
user@switch3# mstp interface ge-0/0/26 cost 1000
user@switch3# mstp interface ge-0/0/26 mode point-to-point
user@switch3# mstp interface ge-0/0/28 cost 1000
user@switch3# mstp interface ge-0/0/28 mode point-to-point
user@switch3# mstp interface ge-0/0/24 cost 1000
user@switch3# mstp interface ge-0/0/24 mode point-to-point
user@switch3# mstp msti 1 bridge-priority 4k
user@switch3# mstp msti 1 vlan [10 20]
user@switch3# mstp msti 2 bridge-priority 16k
user@switch3# mstp msti 2 vlan [30 40]
```

**Results** Check the results of the configuration:

```
user@switch3> show configuration
interfaces {
 ge-0/0/26 {
 unit 0 {
 family ethernet-switching {
 interface-mode trunk;
 vlan {
 members 10;
 members 20;
 members 30;
 members 40;
```

```
 }
 }
}
ge-0/0/28 {
 unit 0 {
 family ethernet-switching {
 interface-mode trunk;
 vlan {
 members 10;
 members 20;
 members 30;
 members 40;
 }
 }
 }
}
ge-0/0/24 {
 unit 0 {
 family ethernet-switching {
 interface-mode trunk;
 vlan {
 members 10;
 members 20;
 members 30;
 members 40;
 }
 }
 }
}
}
}
}
protocols {
 mstp {
 configuration-name region1;
 bridge-priority 8k;
 interface ge-0/0/26 {
 cost 1000;
 mode point-to-point;
 }
 interface ge-0/0/28 {
 cost 1000;
 mode point-to-point;
 }
 interface ge-0/0/24 {
 cost 1000;
 mode point-to-point;
 }
 msti 1 {
 bridge-priority 4k;
 vlan [10 20];
 }
 msti 2 {
 bridge-priority 16k;
 vlan [30 40];
 }
 }
}
```



```

 }
 }
 vlans {
 voice-vlan {
 vlan-id 10;
 }
 employee-vlan {
 vlan-id 20;
 }
 guest-vlan {
 vlan-id 30;
 }
 camera-vlan {
 vlan-id 40;
 }
 }
}

```

## Configuring MSTP on Switch 4

**CLI Quick Configuration** To quickly configure interfaces and MSTP on Switch 4, copy the following commands and paste them into the switch terminal window:

```

[edit]
set vlans voice-vlan description "Voice VLAN"
set vlans voice-vlan vlan-id 10
set vlans employee-vlan description "Employee VLAN"
set vlans employee-vlan vlan-id 20
set vlans guest-vlan description "Guest VLAN"
set vlans guest-vlan vlan-id 30
set vlans camera-vlan description "Camera VLAN"
set vlans camera-vlan vlan-id 40
set interfaces ge-0/0/23 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/19 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/23 unit 0 family ethernet-switching interface-mode trunk
set interfaces ge-0/0/19 unit 0 family ethernet-switching interface-mode trunk
set protocols mstp configuration-name region1
set protocols mstp bridge-priority 16k
set protocols mstp interface ge-0/0/23 cost 1000
set protocols mstp interface ge-0/0/23 mode point-to-point
set protocols mstp interface ge-0/0/19 cost 1000
set protocols mstp interface ge-0/0/19 mode point-to-point
set protocols mstp msti 1 bridge-priority 16k
set protocols mstp msti 1 vlan [10 20]
set protocols mstp msti 2 bridge-priority 32k
set protocols mstp msti 2 vlan [30 40]

```

**Step-by-Step Procedure** To configure interfaces and MSTP on Switch 4:

1. Configure the VLANs **voice-vlan**, **employee-vlan**, **guest-vlan**, and **camera-vlan**:

```

[edit vlans]
user@switch4# set voice-vlan description "Voice VLAN"
user@switch4# set voice-vlan vlan-id 10
user@switch4# set employee-vlan description "Employee VLAN"
user@switch4# set employee-vlan vlan-id 20
user@switch4# set guest-vlan description "Guest VLAN"
user@switch4# set guest-vlan vlan-id 30
user@switch4# set camera-vlan description "Camera VLAN"
user@switch4# set camera-vlan vlan-id 40

```

2. Configure the VLANs on the interfaces, including support for the Ethernet Switching protocol:

```
[edit interfaces]
user@switch4# set ge-0/0/23 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch4# set ge-0/0/19 unit 0 family ethernet-switching vlan members [10 20 30 40]
```

3. Configure the port mode for the interfaces:

```
[edit interfaces]
user@switch4# set ge-0/0/23 unit 0 family ethernet-switching interface-mode trunk
user@switch4# set ge-0/0/19 unit 0 family ethernet-switching interface-mode trunk
```

4. Configure MSTP on the switch, including the two MSTIs:

```
[edit protocols]
user@switch4# mstp configuration-name region1
user@switch4# mstp bridge-priority 16k
user@switch4# mstp interface ge-0/0/23 cost 1000
user@switch4# mstp interface ge-0/0/23 mode point-to-point
user@switch4# mstp interface ge-0/0/19 cost 1000
user@switch4# mstp interface ge-0/0/19 mode point-to-point
user@switch4# mstp msti 1 bridge-priority 16k
user@switch4# mstp msti 1 vlan [10 20]
user@switch4# mstp msti 2 bridge-priority 32k
user@switch4# mstp msti 2 vlan [30 40]
```

**Results** Check the results of the configuration:

```
user@switch4> show configuration
interfaces {
 ge-0/0/23 {
 unit 0 {
 family ethernet-switching {
 interface-mode trunk;
 vlan {
 members 10;
 members 20;
 members 30;
 members 40;
 }
 }
 }
 }
 ge-0/0/19 {
 unit 0 {
 family ethernet-switching {
 interface-mode trunk;
 vlan {
 members 10;
 members 20;
 members 30;
 members 40;
 }
 }
 }
 }
}
```

```

protocols {
 mstp {
 configuration-name region1;
 bridge-priority 16k;
 interface ge-0/0/23 {
 cost 1000;
 mode point-to-point;
 }
 interface ge-0/0/19 {
 cost 1000;
 mode point-to-point;
 }
 msti 1 {
 bridge-priority 16k;
 vlan [10 20];
 }
 msti 2 {
 bridge-priority 32k;
 vlan [30 40];
 }
 }
}
vlands {
 voice-vlan {
 vlan-id 10;
 }
 employee-vlan {
 vlan-id 20;
 }
 guest-vlan {
 vlan-id 30;
 }
 camera-vlan {
 vlan-id 40;
 }
}
}

```

## Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verifying MSTP Configuration on Switch 1 on page 2163](#)
- [Verifying MSTP Configuration on Switch 2 on page 2165](#)
- [Verifying MSTP Configuration on Switch 3 on page 2166](#)
- [Verifying MSTP Configuration on Switch 4 on page 2168](#)

### Verifying MSTP Configuration on Switch 1

**Purpose** Verify the MSTP configuration on Switch 1.

**Action** Issue the operational mode commands **show spanning-tree interface** and **show spanning-tree bridge**:

```
user@switch1> show spanning-tree interface
```

## Spanning tree interface parameters for instance 0

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-0/0/9  | 128:490 | 128:490               | 16384.4c9614e9f841      | 1000         | BLK   | DIS  |
| ge-0/0/11 | 128:491 | 128:491               | 16384.4c9614e9f841      | 1000         | BLK   | DIS  |
| ge-0/0/13 | 128:492 | 128:492               | 16384.4c9614e9f841      | 1000         | BLK   | DIS  |

## Spanning tree interface parameters for instance 1

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-0/0/9  | 128:490 | 128:490               | 16385.4c9614e9f841      | 1000         | BLK   | DIS  |
| ge-0/0/11 | 128:491 | 128:491               | 16385.4c9614e9f841      | 1000         | BLK   | DIS  |
| ge-0/0/13 | 128:492 | 128:492               | 16385.4c9614e9f841      | 1000         | BLK   | DIS  |

## Spanning tree interface parameters for instance 2

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-0/0/9  | 128:490 | 128:490               | 8194.4c9614e9f841       | 1000         | BLK   | DIS  |
| ge-0/0/11 | 128:491 | 128:491               | 8194.4c9614e9f841       | 1000         | BLK   | DIS  |
| ge-0/0/13 | 128:492 | 128:492               | 8194.4c9614e9f841       | 1000         | BLK   | DIS  |

## user@switch1&gt; show spanning-tree bridge

## STP bridge parameters

Routing instance name : GLOBAL  
Context ID : 0  
Enabled protocol : MSTP

## STP bridge parameters for CIST

Root ID : 16384.4c:96:14:e9:f8:41  
CIST regional root : 16384.4c:96:14:e9:f8:41  
CIST internal root cost : 0  
Hello time : 2 seconds  
Maximum age : 20 seconds  
Forward delay : 15 seconds  
Number of topology changes : 0  
Local parameters  
Bridge ID : 16384.4c:96:14:e9:f8:41

## STP bridge parameters for MSTI 1

MSTI regional root : 16385.4c:96:14:e9:f8:41  
Hello time : 2 seconds  
Maximum age : 20 seconds  
Forward delay : 15 seconds  
Number of topology changes : 0  
Local parameters  
Bridge ID : 16385.4c:96:14:e9:f8:41

## STP bridge parameters for MSTI 2

MSTI regional root : 8194.4c:96:14:e9:f8:41  
Hello time : 2 seconds  
Maximum age : 20 seconds  
Forward delay : 15 seconds  
Number of topology changes : 0  
Local parameters  
Bridge ID : 8194.4c:96:14:e9:f8:41

**Meaning** The operational mode command **show spanning-tree interface** displays spanning-tree domain information such as the designated port and the port roles.

The operational mode command **show spanning-tree bridge** displays the spanning-tree domain information at either the bridge level or the interface level. If the optional interface name is omitted, all interfaces in the spanning-tree domain are displayed.

### Verifying MSTP Configuration on Switch 2

**Purpose** Verify the MSTP configuration on Switch 2.

**Action** Issue the operational mode commands **show spanning-tree interface** and **show spanning-tree bridge**:

```
user@switch2> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-0/0/14 | 128:513 | 128:513               | 32768.0019e2503d20      | 1000         | FWD   | DESG |
| ge-0/0/18 | 128:519 | 128:515               | 8192.0019e25051e0       | 1000         | FWD   | ROOT |

Spanning tree interface parameters for instance 1

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-0/0/14 | 128:513 | 128:513               | 32769.0019e2503d20      | 1000         | FWD   | DESG |
| ge-0/0/18 | 128:519 | 128:515               | 4097.0019e25051e0       | 1000         | FWD   | ROOT |

Spanning tree interface parameters for instance 2

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-0/0/14 | 128:513 | 128:513               | 4098.0019e2503d20       | 1000         | FWD   | DESG |
| ge-0/0/18 | 128:519 | 128:519               | 4098.0019e2503d20       | 1000         | FWD   | DESG |

```
user@switch2> show spanning-tree bridge
```

STP bridge parameters

```
Context ID : 0
Enabled protocol : MSTP
```

STP bridge parameters for CIST

```
Root ID : 8192.00:19:e2:50:51:e0
Root cost : 1000
Root port : ge-0/0/18
CIST regional root : 8192.00:19:e2:50:51:e0
CIST internal root cost : 1000
Hello time : 2 seconds
Maximum age : 20 seconds
Forward delay : 15 seconds
Hop count : 19
Message age : 0
Number of topology changes : 1
Time since last topology change : 782 seconds
Local parameters
```

```

 Bridge ID : 32768.00:19:e2:50:3d:20
 Extended system ID : 0
 Internal instance ID : 0

STP bridge parameters for MSTI 1
 MSTI regional root : 4096.00:19:e2:50:51:e0
 Root cost : 1000
 Root port : ge-0/0/18
 Hello time : 2 seconds
 Maximum age : 20 seconds
 Forward delay : 15 seconds
 Hop count : 19
 Local parameters
 Bridge ID : 32768.00:19:e2:50:3d:20
 Extended system ID : 0
 Internal instance ID : 1

STP bridge parameters for MSTI 2
 MSTI regional root : 4096.00:19:e2:50:3d:20
 Hello time : 2 seconds
 Maximum age : 20 seconds
 Forward delay : 15 seconds
 Local parameters
 Bridge ID : 4096.00:19:e2:50:3d:20
 Extended system ID : 0
 Internal instance ID : 2

```

**Meaning** The operational mode command **show spanning-tree interface** displays spanning-tree domain information such as the designated port and the port roles. The spanning-tree interface parameters for instance 2 show that both ports are designated ports, which means Switch 2 is the root bridge for this instance.

The operational mode command **show spanning-tree bridge** displays the spanning-tree domain information at either the bridge level or interface level. If the optional interface name is omitted, all interfaces in the spanning-tree domain are displayed.

### Verifying MSTP Configuration on Switch 3

**Purpose** Verify the MSTP configuration on Switch 3.

**Action** Issue the operational mode commands **show spanning-tree interface** and **show spanning-tree bridge**:

```
user@switch3> show spanning-tree interface
```

```
Spanning tree interface parameters for instance 0
```

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-0/0/26 | 128:513 | 128:513               | 8192.0019e25051e0       | 1000         | FWD   | DESG |
| ge-0/0/28 | 128:515 | 128:515               | 8192.0019e25051e0       | 1000         | FWD   | DESG |
| ge-0/0/24 | 128:517 | 128:517               | 8192.0019e25051e0       | 1000         | FWD   | DESG |

```
Spanning tree interface parameters for instance 1
```

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-0/0/26 | 128:513 | 128:513               | 4096.0019e25051e0       | 1000         | FWD   | DESG |

|           |         |         |                   |      |     |      |
|-----------|---------|---------|-------------------|------|-----|------|
| ge-0/0/28 | 128:515 | 128:515 | 4096.0019e25051e0 | 1000 | FWD | DESG |
| ge-0/0/24 | 128:517 | 128:517 | 4096.0019e25051e0 | 1000 | FWD | DESG |

Spanning tree interface parameters for instance 2

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-0/0/26 | 128:513 | 128:531               | 8192.0019e25044e0       | 1000         | BLK   | ALT  |
| ge-0/0/28 | 128:515 | 128:519               | 4096.0019e2503d20       | 1000         | FWD   | ROOT |
| ge-0/0/24 | 128:517 | 128:517               | 16384.0019e25051e0      | 1000         | FWD   | DESG |

user@switch3> **show spanning-tree bridge**

STP bridge parameters

Context ID : 0  
Enabled protocol : MSTP

STP bridge parameters for CIST

Root ID : 8192.00:19:e2:50:51:e0  
CIST regional root : 8192.00:19:e2:50:51:e0  
CIST internal root cost : 0  
Hello time : 2 seconds  
Maximum age : 20 seconds  
Forward delay : 15 seconds  
Number of topology changes : 3  
Time since last topology change : 843 seconds  
Local parameters  
Bridge ID : 8192.00:19:e2:50:51:e0  
Extended system ID : 0  
Internal instance ID : 0

STP bridge parameters for MSTI 1

MSTI regional root : 4096.00:19:e2:50:51:e0  
Hello time : 2 seconds  
Maximum age : 20 seconds  
Forward delay : 15 seconds  
Local parameters  
Bridge ID : 4096.00:19:e2:50:51:e0  
Extended system ID : 0  
Internal instance ID : 1

STP bridge parameters for MSTI 2

MSTI regional root : 4096.00:19:e2:50:3d:20  
Root cost : 1000  
Root port : ge-0/0/28  
Hello time : 2 seconds  
Maximum age : 20 seconds  
Forward delay : 15 seconds  
Hop count : 19  
Local parameters  
Bridge ID : 16384.00:19:e2:50:51:e0  
Extended system ID : 0  
Internal instance ID : 2

**Meaning** The operational mode command **show spanning-tree interface** displays spanning-tree domain information such as the designated port and the port roles. Switch 3 is the root bridge for instance 0, which is the CIST, as well as for instance 1. In both instances, all ports on Switch 3 are designated ports.

The operational mode command **show spanning-tree bridge** displays the spanning-tree domain information at either the bridge level or the interface level. If the optional interface name is omitted, all interfaces in the spanning-tree domain are displayed.

### Verifying MSTP Configuration on Switch 4

**Purpose** Verify the MSTP configuration on Switch 4.

**Action** Issue the operational mode commands **show spanning-tree interface** and **show spanning-tree bridge**:

```
user@switch4> show spanning-tree interface
```

```
Spanning tree interface parameters for instance 0
```

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-0/0/23 | 128:523 | 128:517               | 8192.0019e25051e0       | 1000         | FWD   | ROOT |
| ge-0/0/19 | 128:525 | 128:525               | 16384.0019e25040e0      | 1000         | FWD   | DESG |

```
Spanning tree interface parameters for instance 1
```

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-0/0/23 | 128:523 | 128:517               | 4096.0019e25051e0       | 1000         | FWD   | ROOT |
| ge-0/0/19 | 128:525 | 128:525               | 16384.0019e25040e0      | 1000         | FWD   | DESG |

```
Spanning tree interface parameters for instance 2
```

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-0/0/23 | 128:523 | 128:517               | 16384.0019e25051e0      | 1000         | BLK   | ALT  |
| ge-0/0/19 | 128:525 | 128:527               | 8192.0019e25044e0       | 1000         | FWD   | ROOT |

```
user@switch4> show spanning-tree bridge
```

```
STP bridge parameters
```

```
Context ID : 0
Enabled protocol : MSTP
```

```
STP bridge parameters for CIST
```

```
Root ID : 8192.00:19:e2:50:51:e0
Root cost : 0
Root port : ge-0/0/23
CIST regional root : 8192.00:19:e2:50:51:e0
CIST internal root cost : 1000
Hello time : 2 seconds
Maximum age : 20 seconds
Forward delay : 15 seconds
Hop count : 19
Message age : 0
Number of topology changes : 4
Time since last topology change : 887 seconds
Local parameters
Bridge ID : 16384.00:19:e2:50:40:e0
Extended system ID : 0
Internal instance ID : 0
```



```

STP bridge parameters for MSTI 1
MSTI regional root : 4096.00:19:e2:50:51:e0
Root cost : 1000
Root port : ge-0/0/23
Hello time : 2 seconds
Maximum age : 20 seconds
Forward delay : 15 seconds
Hop count : 19
Local parameters
 Bridge ID : 16384.00:19:e2:50:40:e0
 Extended system ID : 0
 Internal instance ID : 1

STP bridge parameters for MSTI 2
MSTI regional root : 4096.00:19:e2:50:3d:20
Root cost : 2000
Root port : ge-0/0/19
Hello time : 2 seconds
Maximum age : 20 seconds
Forward delay : 15 seconds
Hop count : 18
Local parameters
 Bridge ID : 32768.00:19:e2:50:40:e0
 Extended system ID : 0
 Internal instance ID : 2

```

**Meaning** The operational mode command **show spanning-tree interface** displays spanning-tree domain information such as the designated port and the port roles.

The operational mode command **show spanning-tree bridge** displays the spanning-tree domain information at either the bridge level or the interface level. If the optional interface name is omitted, all interfaces in the spanning-tree domain are displayed.

- Related Documentation**
- [Example: Configuring Faster Convergence and Improved Network Stability with RSTP on page 2173](#)
  - *Understanding MSTP for EX Series and QFX Series Switches*

## Configuring MSTP

You can configure the Multiple Spanning Tree Protocol (MSTP) under **[edit protocols]**.

To configure the Multiple Spanning Tree Protocol:

1. Enable MSTP as the version of spanning-tree protocol to be configured:

```
[edit]
user@switch# edit ... protocols mstp
```

2. Configure the interfaces that participate in the MSTP instance.

- Configure a specific interface:

- a. Enable MSTP on the specified interface:

```
[edit ... protocols mstp]
user@switch# edit interface interface-name
```

- b. Configure the interface priority:

```
[edit ... protocols mstp interface interface-name]
user@switch# set priority interface-priority
```

- c. (Optional) By default, the interface link cost is determined by the link speed. You can configure the interface link cost to control which bridge is the designated bridge and which port is the designated port:

```
[edit ... protocols mstp interface interface-name]
user@switch# set cost interface-link-cost
```

- d. Configure the interface link mode to identify point-to-point links:

```
[edit ... protocols mstp interface interface-name]
user@switch# set mode (p2p | shared)
```

Specify **p2p** if the link is point to point. Specify **shared** if the link is a shared media.

- e. (Optional) Configure the interface as an edge port:

```
[edit ... protocols mstp interface interface-name]
user@switch# set edge
```

Edge ports do not expect to receive bridge protocol data unit (BPDU) packets. If a BPDU packet is received for an edge port, the port becomes a nonedge port.

- f. (Optional) Disable MSTP on a specific interface:

```
[edit ... protocols mstp interface interface-name]
user@switch# set disable
```

- Enable MSTP on all the interfaces:



**NOTE:** You *cannot* disable MSTP on all the interfaces.

- a. Enable MSTP on all interfaces:

```
[edit ... protocols mstp]
user@switch# set interface all
```

You can also enable BPDU root protection for all spanning-tree protocol instances

on the interface. BPDU root protect ensures the port is the spanning-tree designated port. If the port receives superior BPDU packets, root protect moves this port to a root-prevented spanning-tree state. For configuration details, see *Checking the Status of Spanning-Tree Instance Interfaces*.

3. Configure the bridge priority

```
[edit ... protocols mstp]
user@switch# set bridge-priority bridge-priority
```

For more information, see *Bridge Priority for Election of Root Bridge and Designated Bridge*.

4. Configure hello BPDU timers.

a. Configure the maximum expected arrival time of hello BPDUs:

```
[edit ... protocols mstp]
user@switch# set max-age seconds
```

b. Configure the time interval at which the root bridge transmits configuration BPDUs:

```
[edit ... protocols mstp]
user@switch# set hello-time seconds
```

5. (Optional) By default, the bridge port remains in the listening and learning states for 15 seconds before transitioning to the forwarding state. You can specify a delay from 4 through 20 seconds instead:

```
[edit ... protocols mstp]
user@switch# set forward-delay seconds
```

6. Configure MSTP-specific options.

a. Configure the MSTP region configuration name:

```
[edit ... protocols mstp]
user@switch# set configuration-name configuration-name
```

b. Configure the MSTP revision level:

```
[edit ... protocols mstp]
user@switch# set revision-level revision-level
```

c. Configure the maximum number of hops a BPDU can be forwarded in the MSTP region:

```
[edit ... protocols mstp]
user@switch# set max-hops hops
```

**Related  
Documentation**

- *Configuring MST Instances on a Physical Interface*

## Understanding RSTP

---

Juniper Networks QFX Series products use Rapid Spanning Tree Protocol (RSTP) on the network side of the QFX Series to provide quicker convergence time than the base Spanning Tree Protocol (STP) does. RSTP identifies certain links as point to point. When a point-to-point link fails, the alternate link can transition to the forwarding state, which speeds up convergence.

Although STP provides basic loop prevention functionality, it does not provide fast network convergence when there are topology changes. The STP process to determine network state transitions is slower than the RSTP process because it is timer-based. A device must reinitialize every time a topology change occurs. The device must start in the listening state and transition to the learning state and eventually to a forwarding or blocking state. When default values are used for the maximum age (20 seconds) and forward delay (15 seconds), it takes 50 seconds for the device to converge. RSTP converges faster because it uses a handshake mechanism based on point-to-point links instead of the timer-based process used by STP.

For networks with virtual LANs (VLANs), you can use VLAN Spanning Tree Protocol (VSTP), which takes the paths of each VLAN into account when calculating routes. VSTP uses RSTP by default.

An RSTP domain running from the edge outward on a QFX Series product has the following components:

- A *root port*, which is the “best path” to the root device.
- A *designated port*, which indicates that the switch is the designated bridge for the other switch connecting to this port.
- An *alternate port*, which provides an alternate root port.
- A *backup port*, which provides an alternate designated port.

Port assignments change through messages exchanged throughout the domain. An RSTP device generates configuration messages once per hello time interval. If an RSTP device does not receive a configuration message from its neighbor after an interval of three hello times, it determines that the connection with the neighbor is lost. When a *root port* or a *designated port* fails on a device, the device generates a configuration message with the proposal bit set. Once its neighbor device receives this message, it verifies that this configuration message is valid for that port and starts a *synchronizing* operation to ensure that all of its ports are in sync with the new information.

Similar sets of messages propagate through the network, restoring the connectivity very quickly after a topology change (in a well-designed network that uses RSTP, network convergence can take as little as 0.5 seconds). If a device does not receive an agreement to a proposal message it has sent, it returns to the original IEEE 802.D convention.

RSTP was originally defined in the IEEE 802.1w draft specification and later incorporated into the IEEE 802.1D-2004 specification.

VSTP and RSTP can be configured at the same time. If you configure VSTP and RSTP at the same time and the switch has more than 253 VLANs, VSTP is configured only for the first 253 VLANs. For the remaining VLANs, only RSTP is configured. RSTP and VSTP are the only spanning-tree protocols that can be configured at the same time on the QFX Series.

#### Related Documentation

- [Overview of Spanning-Tree Protocols on page 2148](#)
- [Understanding MSTP on page 2149](#)
- [Understanding VSTP on page 2191](#)
- [Example: Configuring Faster Convergence and Improving Network Stability with RSTP](#)
- [Example: Configuring Faster Convergence and Improved Network Stability with RSTP on page 2173](#)
- [Configuring RSTP \(CLI Procedure\) on page 2190](#)

## Example: Configuring Faster Convergence and Improved Network Stability with RSTP



**NOTE:** This example uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches*. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 41](#).

EX Series switches use Rapid Spanning Tree Protocol (RSTP) by default to provide a loop-free topology.

When switches that support redundant Routing Engines use RSTP, it is important to keep RSTP synchronized on both Routing Engines so that no loss of service occurs after a Routing Engine switchover. Nonstop bridging protocol keeps Routing Engines synchronized.

This example describes how to configure RSTP and NSB on four EX Series switches:

- [Requirements on page 2173](#)
- [Overview and Topology on page 2174](#)
- [Configuring RSTP and Nonstop Bridging on Switch 1 on page 2176](#)
- [Configuring RSTP and Nonstop Bridging on Switch 2 on page 2179](#)
- [Configuring RSTP and Nonstop Bridging on Switch 3 on page 2182](#)
- [Configuring RSTP and Nonstop Bridging on Switch 4 on page 2185](#)
- [Verification on page 2188](#)

## Requirements

This example uses the following software and hardware components:

- Junos OS Release 15.1 or later or later for EX Series switches

- Four EX Series switches


Before you configure the switches for RSTP, be sure you have:

- Installed and connected the four switches. See the hardware documentation for your switch.
- Performed the initial software configuration on all switches. See *Connecting and Configuring an EX Series Switch (CLI Procedure)* or *Connecting and Configuring an EX Series Switch (J-Web Procedure)*.

## Overview and Topology

RSTP works by identifying certain links as point to point links and blocking other possible paths. When one of the point-to-point links fails, a designated alternate link transitions to the forwarding state and take over. Configuring nonstop bridging (NSB) on a switch with redundant Routing Engines keeps RSTP synchronized on both Routing Engines. This way, RSTP remains active immediately after a switchover because it is already synchronized to the backup Routing Engine. RSTP does not have to reconverge after a Routing Engine switchover when NSB is enabled because the neighbor devices do not detect an RSTP change on the switch. In this example, four EX Series switches are connected in the topology displayed in [Figure 25](#) to create a loop-free topology with NSB applied to switches with dual Routing Engines.

Figure 25: Network Topology for RSTP



[Table 189](#) shows the components of the topology for this example.



**NOTE:** You can configure RSTP only on physical interfaces, not on logical interfaces.

Table 189: Components of the Topology for Configuring RSTP

| Property               | Settings                                                                                                                                                                                                                                                                              |
|------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Switch 1               | The following interfaces on Switch 1 are connected in this way: <ul style="list-style-type: none"><li>• <code>ge-0/0/9</code> is connected to Switch 2</li><li>• <code>ge-0/0/13</code> is connected to Switch 4</li><li>• <code>ge-0/0/11</code> is connected to Switch 3</li></ul>  |
| Switch 2               | The following interfaces on Switch 2 are connected in this way: <ul style="list-style-type: none"><li>• <code>ge-0/0/14</code> is connected to Switch 1</li><li>• <code>ge-0/0/18</code> is connected to Switch 3</li></ul>                                                           |
| Switch 3               | The following interfaces on Switch 3 are connected in this way: <ul style="list-style-type: none"><li>• <code>ge-0/0/26</code> is connected to Switch 1</li><li>• <code>ge-0/0/28</code> is connected to Switch 2</li><li>• <code>ge-0/0/24</code> is connected to Switch 4</li></ul> |
| Switch 4               | The following interfaces on Switch 4 are connected in this way: <ul style="list-style-type: none"><li>• <code>ge-0/0/19</code> is connected to Switch 1</li><li>• <code>ge-0/0/23</code> is connected to Switch 3</li></ul>                                                           |
| VLAN names and tag IDs | <code>voice-vlan</code> , tag 10<br><code>employee-vlan</code> , tag 20<br><code>guest-vlan</code> , tag 30<br><code>camera-vlan</code> , tag 40                                                                                                                                      |

This configuration example creates a loop-free topology between four EX Series switches using RSTP.

An RSTP topology contains ports that have specific roles:

- The *root port* is responsible for forwarding data to the root bridge.
- The *alternate port* is a standby port for the root port. When a root port goes down, the alternate port becomes the active root port.
- The *designated port* forwards data to the downstream network segment or device.
- The *backup port* is a backup port for the designated port. When a designated port goes down, the backup port becomes the active designated port and starts forwarding data.



**NOTE:** You also can create a loop-free topology between the aggregation layer and the distribution layer using redundant trunk links. For more information about configuring redundant trunk links, see [“Example: Configuring Redundant Trunk Links for Faster Recovery”](#) on page 2903.

## Configuring RSTP and Nonstop Bridging on Switch 1

**CLI Quick Configuration** To quickly configure RSTP and nonstop bridging on Switch 1, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans voice-vlan description "Voice VLAN"
set vlans voice-vlan vlan-id 10
set vlans employee-vlan description "Employee VLAN"
set vlans employee-vlan vlan-id 20
set vlans guest-vlan description "Guest VLAN"
set vlans guest-vlan vlan-id 30
set vlans camera-vlan description "Camera VLAN"
set vlans camera-vlan vlan-id 40
set interfaces ge-0/0/13 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/9 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/11 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/13 unit 0 family ethernet-switching interface-mode trunk
set interfaces ge-0/0/9 unit 0 family ethernet-switching interface-mode trunk
set interfaces ge-0/0/11 unit 0 family ethernet-switching interface-mode trunk
set protocols rstp bridge-priority 16k
set protocols rstp interface all cost 1000
set protocols rstp interface all mode point-to-point
```

If Switch 1 includes dual Routing Engines, configure NSB. To quickly configure nonstop bridging on Switch 1, copy the following commands and paste them into the switch terminal window:

```
set chassis redundancy graceful switchover
set system commit synchronize
set protocols layer2-control nonstop-bridging
```

**Step-by-Step Procedure** To configure RSTP and nonstop bridging on Switch 1:

1. Configure the VLANs **voice-vlan**, **employee-vlan**, **guest-vlan**, and **camera-vlan**:

```
[edit vlans]
user@switch1# set voice-vlan description "Voice VLAN"
user@switch1# set voice-vlan vlan-id 10
user@switch1# set employee-vlan description "Employee VLAN"
user@switch1# set employee-vlan vlan-id 20
user@switch1# set guest-vlan description "Guest VLAN"
user@switch1# set guest-vlan vlan-id 30
user@switch1# set camera-vlan description "Camera VLAN"
user@switch1# set camera-vlan vlan-id 40
```

2. Configure the VLANs on the interfaces, including support for the Ethernet switching protocol:

```
[edit interfaces]
user@switch1# set ge-0/0/13 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch1# set ge-0/0/9 unit 0 family ethernet-switching vlan members [10 20 30 40]
```



```
user@switch1# set ge-0/0/11 unit 0 family ethernet-switching vlan members [10 20 30 40]
```

3. Configure the port mode for the interfaces:

```
[edit interfaces]
user@switch1# set ge-0/0/13 unit 0 family ethernet-switching interface-mode trunk
user@switch1# set ge-0/0/9 unit 0 family ethernet-switching interface-mode trunk
user@switch1# set ge-0/0/11 unit 0 family ethernet-switching interface-mode trunk
```

4. Configure RSTP on the switch:

```
[edit protocols]
user@switch1# rstp bridge-priority 16k
user@switch1# rstp interface all cost 1000
user@switch1# rstp interface all mode point-to-point
```

**Step-by-Step Procedure** If Switch 1 includes dual Routing Engines, configure nonstop bridging. To configure NSB on Switch 1:

1. Enable graceful Routing Engine switchover (GRES):

```
[edit chassis redundancy]
user@switch1# set graceful-switchover
```

2. Configure the switch to always synchronize configuration changes between the Routing Engines:

```
[edit system]
user@switch1# set commit synchronize
```

If you try to commit a configuration in which nonstop bridging is configured but synchronization of configuration changes is not configured, the configuration is not committed.

3. Enable nonstop bridging:

```
[edit protocols layer2-control]
user@switch1# set nonstop-bridging
```



**NOTE:** This process enables NSB for all NSB-supported Layer 2 protocols on the switch, including RSTP.

**Results** Check the results of the configuration:

```
user@switch1> show configuration
interfaces {
 ge-0/0/13 {
 unit 0 {
 family ethernet-switching {
 interface-mode trunk;
 vlan {
 members [10 20 30 40];
 }
 }
 }
 }
 ge-0/0/9 {
 unit 0 {
```

```
 family ethernet-switching {
 interface-mode trunk;
 vlan {
 members [10 20 30 40];
 }
 }
 }
}
ge-0/0/11 {
 unit 0 {
 family ethernet-switching {
 interface-mode trunk;
 vlan {
 members [10 20 30 40];
 }
 }
 }
}
}
protocols {
 layer2-control {
 nonstop-bridging;
 }
 rstp {
 bridge-priority 16k;
 interface ge-0/0/13 {
 cost 1000;
 mode point-to-point;
 }
 interface ge-0/0/9 {
 cost 1000;
 mode point-to-point;
 }
 interface ge-0/0/11 {
 cost 1000;
 mode point-to-point;
 }
 }
}
}
vllans {
 voice-vlan {
 vlan-id 10;
 }
 employee-vlan {
 vlan-id 20;
 }
 guest-vlan {
 vlan-id 30;
 }
 camera-vlan {
 vlan-id 40;
 }
}
}
system {
 commit synchronize;
}
```

```

}
chassis {
 redundancy {
 graceful-switchover;
 }
}

```

## Configuring RSTP and Nonstop Bridging on Switch 2

**CLI Quick Configuration** To quickly configure RSTP and nonstop bridging on Switch 2, copy the following commands and paste them into the switch terminal window:

```

[edit]
set vlans voice-vlan description "Voice VLAN"
set vlans voice-vlan vlan-id 10
set vlans employee-vlan description "Employee VLAN"
set vlans employee-vlan vlan-id 20
set vlans guest-vlan description "Guest VLAN"
set vlans guest-vlan vlan-id 30
set vlans camera-vlan description "Camera VLAN"
set vlans camera-vlan vlan-id 40
set interfaces ge-0/0/14 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/18 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/14 unit 0 family ethernet-switching interface-mode trunk
set interfaces ge-0/0/18 unit 0 family ethernet-switching interface-mode trunk
set protocols rstp bridge-priority 32k
set protocols rstp interface ge-0/0/14 cost 1000
set protocols rstp interface ge-0/0/14 mode point-to-point
set protocols rstp interface ge-0/0/18 cost 1000
set protocols rstp interface ge-0/0/18 mode point-to-point

```



**NOTE:** Starting with Junos OS Release 15.1 for EX Series and QFX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style, you can configure spanning tree parameters globally on all spanning tree interfaces. See [“Configuring RSTP \(CLI Procedure\)” on page 2190](#) for additional information.

If Switch 2 includes dual Routing Engines, configure NSB. To quickly configure nonstop bridging on Switch 2, copy the following commands and paste them into the switch terminal window:

```

set chassis redundancy graceful switchover
set system commit synchronize
set protocols layer2-control nonstop-bridging

```

**Step-by-Step  
Procedure**

To configure RSTP and nonstop bridging on Switch 2:

1. Configure the VLANs **voice-vlan**, **employee-vlan**, **guest-vlan**, and **camera-vlan**:

```
[edit vlans]
user@switch2# set voice-vlan description "Voice VLAN"
user@switch2# set voice-vlan vlan-id 10
user@switch2# set employee-vlan description "Employee VLAN"
user@switch2# set employee-vlan vlan-id 20
user@switch2# set guest-vlan description "Guest VLAN"
user@switch2# set guest-vlan vlan-id 30
user@switch2# set camera-vlan vlan-description "Camera VLAN"
user@switch2# set camera-vlan vlan-id 40
```

2. Configure the VLANs on the interfaces, including support for the Ethernet switching protocol:

```
[edit interfaces]
user@switch2# set ge-0/0/14 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch2# set ge-0/0/18 unit 0 family ethernet-switching vlan members [10 20 30 40]
```

3. Configure the port mode for the interfaces:

```
[edit interfaces]
user@switch2# set ge-0/0/14 unit 0 family ethernet-switching interface-mode trunk
user@switch2# set ge-0/0/18 unit 0 family ethernet-switching interface-mode trunk
```

4. Configure RSTP on the switch:

```
[edit protocols]
user@switch2# rstp bridge-priority 32k
user@switch2# rstp interface ge-0/0/14 cost 1000
user@switch2# rstp interface ge-0/0/14 mode point-to-point
user@switch2# rstp interface ge-0/0/18 cost 1000
user@switch2# rstp interface ge-0/0/18 mode point-to-point
```

**Step-by-Step  
Procedure**

If Switch 2 includes dual Routing Engines, configure nonstop bridging. To configure NSB on Switch 2:

1. Enable graceful Routing Engine switchover (GRES):

```
[edit chassis redundancy]
user@switch2# set graceful-switchover
```

2. Configure the switch to always synchronize configuration changes between the Routing Engines:

```
[edit system]
user@switch2# set commit synchronize
```

If you try to commit a configuration in which nonstop bridging is configured but synchronization of configuration changes is not configured, the configuration is not committed.

3. Enable nonstop bridging:

```
[edit protocols layer2-control]
user@switch2# set nonstop-bridging
```



**NOTE:** This process enables NSB for all NSB-supported Layer 2 protocols on the switch, including RSTP.

**Results** Check the results of the configuration:

```

user@switch2> show configuration
interfaces {
 ge-0/0/14 {
 unit 0 {
 family ethernet-switching {
 interface-mode trunk;
 vlan {
 members [10 20 30 40];
 }
 }
 }
 }
 ge-0/0/18 {
 unit 0 {
 family ethernet-switching {
 interface-mode trunk;
 vlan {
 members [10 20 30 40];
 }
 }
 }
 }
}
protocols {
 layer2-control {
 nonstop-bridging;
 }
 rstp {
 bridge-priority 32k;
 interface ge-0/0/14 {
 cost 1000;
 mode point-to-point;
 }
 interface ge-0/0/18 {
 cost 1000;
 mode point-to-point;
 }
 }
}
vlands {
 voice-vlan {
 vlan-id 10;
 }
 employee-vlan {
 vlan-id 20;
 }
 guest-vlan {
 vlan-id 30;
 }
 camera-vlan {
 vlan-id 40;
 }
}

```

```
system {
 commit synchronize;
}
chassis {
 redundancy {
 graceful-switchover;
 }
}
```

## Configuring RSTP and Nonstop Bridging on Switch 3

**CLI Quick Configuration** To quickly configure RSTP and nonstop bridging on Switch 3, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans voice-vlan description "Voice VLAN"
set vlans voice-vlan vlan-id 10
set vlans employee-vlan description "Employee VLAN"
set vlans employee-vlan vlan-id 20
set vlans guest-vlan description "Guest VLAN"
set vlans guest-vlan vlan-id 30
set vlans camera-vlan description "Camera VLAN"
set vlans camera-vlan vlan-id 40
set interfaces ge-0/0/26 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/28 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/24 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/26 unit 0 family ethernet-switching interface-mode trunk
set interfaces ge-0/0/28 unit 0 family ethernet-switching interface-mode trunk
set interfaces ge-0/0/24 unit 0 family ethernet-switching interface-mode trunk
set protocols rstp bridge-priority 8k
set protocols rstp interface ge-0/0/26 cost 1000
set protocols rstp interface ge-0/0/26 mode point-to-point
set protocols rstp interface ge-0/0/28 cost 1000
set protocols rstp interface ge-0/0/28 mode point-to-point
set protocols rstp interface ge-0/0/24 cost 1000
set protocols rstp interface ge-0/0/24 mode point-to-point
```

If Switch 3 includes dual Routing Engines, configure NSB. To quickly configure nonstop bridging on Switch 3, copy the following commands and paste them into the switch terminal window:

```
set chassis redundancy graceful switchover
set system commit synchronize
set protocols layer2-control nonstop-bridging
```

**Step-by-Step Procedure** To configure RSTP and nonstop bridging on Switch 3:

1. Configure the VLANs **voice-vlan**, **employee-vlan**, **guest-vlan**, and **camera-vlan**:

```
[edit vlans]
user@switch3# set voice-vlan description "Voice VLAN"
user@switch3# set voice-vlan vlan-id 10
user@switch3# set employee-vlan description "Employee VLAN"
user@switch3# set employee-vlan vlan-id 20
user@switch3# set guest-vlan description "Guest VLAN"
user@switch3# set guest-vlan vlan-id 30
user@switch3# set camera-vlan description "Camera VLAN"
user@switch3# set camera-vlan vlan-id 40
```

2. Configure the VLANs on the interfaces, including support for the Ethernet switching protocol:

```
[edit interfaces]
user@switch3# set ge-0/0/26 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch3# set ge-0/0/28 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch3# set ge-0/0/24 unit 0 family ethernet-switching vlan members [10 20 30 40]
```

3. Configure the port mode for the interfaces:

```
[edit interfaces]
user@switch3# set ge-0/0/26 unit 0 family ethernet-switching interface-mode trunk
user@switch3# set ge-0/0/28 unit 0 family ethernet-switching interface-mode trunk
user@switch3# set ge-0/0/24 unit 0 family ethernet-switching interface-mode trunk
```

4. Configure RSTP on the switch:

```
[edit protocols]
user@switch3# rstp bridge-priority 8k
user@switch3# rstp interface ge-0/0/26 cost 1000
user@switch3# rstp interface ge-0/0/26 mode point-to-point
user@switch3# rstp interface ge-0/0/28 cost 1000
user@switch3# rstp interface ge-0/0/28 mode point-to-point
user@switch3# rstp interface ge-0/0/24 cost 1000
user@switch3# rstp interface ge-0/0/24 mode point-to-point
```

#### Step-by-Step Procedure

If Switch 3 includes dual Routing Engines, configure nonstop bridging. To configure NSB on Switch 3:

1. Enable graceful Routing Engine switchover (GRES):

```
[edit chassis redundancy]
user@switch3# set graceful-switchover
```

2. Configure the switch to always synchronize configuration changes between the Routing Engines:

```
[edit system]
user@switch3# set commit synchronize
```

If you try to commit a configuration in which nonstop bridging is configured but synchronization of configuration changes is not configured, the configuration is not committed.

3. Enable nonstop bridging:

```
[edit protocols layer2-control]
user@switch3# set nonstop-bridging
```



**NOTE:** This process enables NSB for all NSB-supported Layer 2 protocols on the switch, including RSTP.

**Results** Check the results of the configuration:

```
user@switch3> show configuration
interfaces {
 ge-0/0/26 {
 unit 0 {
 family ethernet-switching {
```

```
 interface-mode trunk;
 vlan {
 members [10 20 30 40];
 }
 }
}
ge-0/0/28 {
 unit 0 {
 family ethernet-switching {
 interface-mode trunk;
 vlan {
 members [10 20 30 40];
 }
 }
 }
}
ge-0/0/24 {
 unit 0 {
 family ethernet-switching {
 interface-mode trunk;
 vlan {
 members [10 20 30 40];
 }
 }
 }
}
}
}
protocols {
 layer2-control {
 nonstop-bridging;
 }
 rstp {
 bridge-priority 8k;
 interface ge-0/0/26 {
 cost 1000;
 mode point-to-point;
 }
 interface ge-0/0/28 {
 cost 1000;
 mode point-to-point;
 }
 interface ge-0/0/24 {
 cost 1000;
 mode point-to-point;
 }
 bridge-priority 8k;
 }
}
}
vlangs {
 voice-vlan {
 vlan-id 10;
 }
}
```



```

 }
 employee-vlan {
 vlan-id 20;
 }
 guest-vlan {
 vlan-id 30;
 }
 camera-vlan {
 vlan-id 40;
 }
}
system {
 commit synchronize;
}
chassis {
 redundancy {
 graceful-switchover;
 }
}

```

## Configuring RSTP and Nonstop Bridging on Switch 4

**CLI Quick Configuration** To quickly configure RSTP and nonstop bridging on Switch 4, copy the following commands and paste them into the switch terminal window:

```

[edit]
set vlans voice-vlan description "Voice VLAN"
set vlans voice-vlan vlan-id 10
set vlans employee-vlan description "Employee VLAN"
set vlans employee-vlan vlan-id 20
set vlans guest-vlan description "Guest VLAN"
set vlans guest-vlan vlan-id 30
set vlans camera-vlan description "Camera VLAN"
set vlans camera-vlan vlan-id 40
set interfaces ge-0/0/23 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/19 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/23 unit 0 family ethernet-switching interface-mode trunk
set interfaces ge-0/0/19 unit 0 family ethernet-switching interface-mode trunk
set protocols rstp bridge-priority 16k
set protocols rstp interface ge-0/0/23 cost 1000
set protocols rstp interface ge-0/0/23 mode point-to-point
set protocols rstp interface ge-0/0/19 cost 1000
set protocols rstp interface ge-0/0/19 mode point-to-point

```

If Switch 4 includes dual Routing Engines, configure NSB. To quickly configure nonstop bridging on Switch 4, copy the following commands and paste them into the switch terminal window:

```

set chassis redundancy graceful switchover
set system commit synchronize
set protocols layer2-control nonstop-bridging

```

**Step-by-Step  
Procedure**

To configure RSTP and nonstop bridging on Switch 4:

1. Configure the VLANs **voice-vlan**, **employee-vlan**, **guest-vlan**, and **camera-vlan**:

```
[edit vlans]
user@switch4# set voice-vlan description "Voice VLAN"
user@switch4# set voice-vlan vlan-id 10
user@switch4# set employee-vlan description "Employee VLAN"
user@switch4# set employee-vlan vlan-id 20
user@switch4# set guest-vlan description "Guest VLAN"
user@switch4# set guest-vlan vlan-id 30
user@switch4# set camera-vlan description "Camera VLAN"
user@switch4# set camera-vlan vlan-id 40
```

2. Configure the VLANs on the interfaces, including support for the Ethernet switching protocol:

```
[edit interfaces]
user@switch4# set ge-0/0/23 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch4# set ge-0/0/19 unit 0 family ethernet-switching vlan members [10 20 30 40]
```

3. Configure the port mode for the interfaces:

```
[edit interfaces]
user@switch4# set ge-0/0/23 unit 0 family ethernet-switching interface-mode trunk
user@switch4# set ge-0/0/19 unit 0 family ethernet-switching interface-mode trunk
```

4. Configure RSTP on the switch:

```
[edit protocols]
user@switch4# rstp bridge-priority 16k
user@switch4# rstp interface ge-0/0/23 cost 1000
user@switch4# rstp interface ge-0/0/23 mode point-to-point
user@switch4# rstp interface ge-0/0/19 cost 1000
user@switch4# rstp interface ge-0/0/19 mode point-to-point
```

**Step-by-Step  
Procedure**

If Switch 4 includes dual Routing Engines, configure nonstop bridging. To configure NSB on Switch 4:

1. Enable graceful Routing Engine switchover (GRES):

```
[edit chassis redundancy]
user@switch4# set graceful-switchover
```

2. Configure the switch to always synchronize configuration changes between the Routing Engines:

```
[edit system]
user@switch4# set commit synchronize
```

If you try to commit a configuration in which nonstop bridging is configured but synchronization of configuration changes is not configured, the configuration is not committed.

3. Enable nonstop bridging:

```
[edit protocols layer2-control]
user@switch4# set nonstop-bridging
```



**NOTE:** This process enables NSB for all NSB-supported Layer 2 protocols on the switch, including RSTP.

**Results** Check the results of the configuration:

```

user@switch4> show configuration
interfaces {
 ge-0/0/23 {
 unit 0 {
 family ethernet-switching {
 interface-mode trunk;
 vlan {
 members [10 20 30 40];
 }
 }
 }
 }
 ge-0/0/19 {
 unit 0 {
 family ethernet-switching {
 interface-mode trunk;
 vlan {
 members [10 20 30 40];
 }
 }
 }
 }
}
protocols {
 layer2-control {
 nonstop-bridging;
 }
 rstp {
 bridge-priority 16k;
 interface ge-0/0/23 {
 cost 1000;
 mode point-to-point;
 }
 interface ge-0/0/19 {
 cost 1000;
 mode point-to-point;
 }
 }
}
vlands {
 voice-vlan {
 vlan-id 10;
 }
 employee-vlan {
 vlan-id 20;
 }
 guest-vlan {
 vlan-id 30;
 }
 camera-vlan {
 vlan-id 40;
 }
}

```

```
system {
 commit synchronize;
}
chassis {
 redundancy {
 graceful-switchover;
 }
}
```

## Verification

To confirm that the configuration is working properly, perform these tasks on both Routing Engines:

- [Verifying RSTP Configuration on Switch 1 on page 2188](#)
- [Verifying RSTP Configuration on Switch 2 on page 2188](#)
- [Verifying RSTP Configuration on Switch 3 on page 2189](#)
- [Verifying RSTP Configuration on Switch 4 on page 2189](#)

### Verifying RSTP Configuration on Switch 1

**Purpose** Verify the RSTP configuration on Switch 1.

**Action** Use the operational mode command:

```
user@switch1> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-0/0/13 | 128:526 | 128:526               | 16384.0019e25040e0      | 1000         | BLK   | ALT  |
| ge-0/0/9  | 128:522 | 128:522               | 32768.0019e2503d20      | 1000         | BLK   | ALT  |
| ge-0/0/11 | 128:524 | 128:524               | 8192.0019e25051e0       | 1000         | FWD   | ROOT |

**Meaning** Refer to the topology in [Figure 25](#). The operational mode command **show spanning-tree interface** shows that **ge-0/0/13** is in a forwarding state. The other interfaces on Switch 1 are blocking.

### Verifying RSTP Configuration on Switch 2

**Purpose** Use this procedure to verify the RSTP configuration on both Switch 2 Routing Engines.

**Action** Use the operational mode command:

```
user@switch2> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-0/0/14 | 128:527 | 128:527               | 32768.0019e2503d20      | 1000         | FWD   | DESG |
| ge-0/0/18 | 128:529 | 128:529               | 8192.0019e25051e0       | 1000         | FWD   | ROOT |

**Meaning** Refer to the topology in [Figure 25](#). The operational mode command **show spanning-tree interface** shows that **ge-0/0/18** is in a forwarding state and is the root port.

### Verifying RSTP Configuration on Switch 3

**Purpose** Use this procedure to verify the RSTP configuration on both Switch 3 Routing Engines.

**Action** Use the operational mode commands:

```
user@switch3> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-0/0/26 | 128:539 | 128:539               | 8192.0019e25051e0       | 1000         | FWD   | DESG |
| ge-0/0/28 | 128:541 | 128:541               | 8192.0019e25051e0       | 1000         | FWD   | DESG |
| ge-0/0/24 | 128:537 | 128:537               | 8192.0019e25051e0       | 1000         | FWD   | DESG |

**Meaning** Refer to the topology in [Figure 25](#). The operational mode command **show spanning-tree interface** shows that no interface is the root interface.

### Verifying RSTP Configuration on Switch 4

**Purpose** Use this procedure to verify the RSTP configuration on both Switch 4 Routing Engines.

**Action** Use the operational mode commands:

```
user@switch4> show spanning-tree interface
Spanning tree interface parameters for instance 0
```

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-0/0/23 | 128:536 | 128:536               | 8192.0019e25051e0       | 1000         | FWD   | ROOT |
| ge-0/0/19 | 128:532 | 128:532               | 16384.0019e25040e0      | 1000         | FWD   | DESC |

**Meaning** Refer to the topology in [Figure 25](#). The operational mode command **show spanning-tree interface** shows that interface **ge-0/0/23** is the root interface and forwarding.

**Related Documentation**

- 95775 [Configuring RSTP \(CLI Procedure\) on page 2190](#)
- [Understanding RSTP for EX Series and QFX Series Switches](#)

## Configuring RSTP (CLI Procedure)

The default spanning-tree protocol for EX Series switches is Rapid Spanning Tree Protocol (RSTP). RSTP provides faster convergence times than the original Spanning Tree Protocol (STP). Because RSTP is configured by default, you only need to use this procedure if another spanning-tree protocol has been configured. In that case, you can reconfigure RSTP.

To enable RSTP:

1. Disable the other configured spanning-tree protocol (MSTP):

- To disable MSTP:

```
[edit protocols]
user@switch# set mstp disable
```

2. Configure RSTP

- To enable RSTP on a specific interface:

```
[edit protocols]
user@switch# set rstp interface interface-name
```

- To disable RSTP on a specific interface:

```
[edit protocols]
user@switch# set rstp interface interface-name disable
```

- To enable RSTP on a range of interfaces:

```
[edit protocols]
user@switch# set rstp interface interface-range-name
```

- To enable RSTP on all interfaces:

```
[edit protocols]
user@switch# set rstp interface all
```

- Related Documentation**
- [show spanning-tree bridge](#)
  - [show spanning-tree interface](#)
  - [Understanding RSTP for EX Series and QFX Series Switches](#)

## Understanding VSTP

VLAN Spanning Tree Protocol (VSTP) enables Juniper Networks switches to run one or more Spanning Tree Protocol (STP) or Rapid Spanning Tree Protocol (RSTP) instances for each VLAN on which VSTP is enabled. For networks with multiple VLANs, VSTP improves intelligent tree spanning by defining best paths within the VLANs instead of within the entire network.

You can configure VSTP for a maximum of 509 VLANs.

VSTP and RSTP can be configured at the same time. If you configure VSTP and RSTP at the same time and the switch has more than 253 VLANs, VSTP is configured only for the first 253 VLANs. For the remaining VLANs, only RSTP is configured. RSTP and VSTP are the only spanning-tree protocols that can be configured at the same time on a switch.



**NOTE:** We recommend that you enable VSTP on all VLANs that could receive VSTP bridge protocol data units (BPDUs).

- Related Documentation**
- [Example: Configuring VSTP on QFX Series Switches and EX4600 Switches on page 2191](#)
  - [Overview of Spanning-Tree Protocols on page 2148](#)
  - [Understanding RSTP on page 2172](#)
  - [Configuring VSTP \(CLI Procedure\) on page 2197](#)
  - [Configuring VLAN Spanning Tree Protocol](#)

## Example: Configuring VSTP on QFX Series Switches and EX4600 Switches

This example demonstrates configuring VSTP (VLAN Spanning Tree Protocol) on QFX Series switches and EX4600 switches. The default spanning-tree protocol on these switches is Rapid Spanning Tree Protocol (RSTP). VLAN Spanning Tree Protocol (VSTP) is an alternate protocol that allows switches to run one or more Spanning Tree Protocol (STP) or RSTP instances for each VLAN on which VSTP is enabled. For networks with multiple VLANs, VSTP improves network bandwidth utilization by load balancing traffic for each VLAN.

- [Requirements on page 2192](#)
- [Overview on page 2192](#)
- [Configuration on page 2193](#)

## Requirements

This example uses the following hardware and software components:

- QFX5100 switch (Note that any QFX Series switch or the EX4600 switch can be substituted.)
- Junos OS Release 13.2X51-D25 or later for the QFX Series or Junos OS Release 13.2X51-D25 or later for EX Series switches

Before you configure VSTP, be sure you have:

- Configured interfaces ge-0/0/3, ge-0/0/4, ge-0/0/8, ae0. For directions, see [“Configuring Gigabit and 10-Gigabit Ethernet Interfaces” on page 2826](#).
- Configured VLAN 200 with four interfaces. For directions, see [“Configuring VLANs” on page 2120](#).
- Configured VLAN 5 and VLAN 6 with three interfaces each. For directions, see [“Configuring VLANs” on page 2120](#).

## Overview

You can configure VSTP for an interface at the global level (for all configured VLANs) or for a specific VLAN.



NOTE:

- If you configure VSTP on an interface at both the global and the specific VLAN level, the interface configuration that is defined at the specific VLAN level overrides the interface configuration that is defined at the global level.
- If you specify VSTP to be configured on an interface that is not configured to belong to the VLAN (or VLANs), an error message is displayed.

For each VLAN configured with VSTP, a dedicated instance of a spanning tree is created. This approach is useful to optimize network usage in small networks with a limited number of VLANs.

In this example, you configure VSTP at the VLAN level for VLAN 200 and at the global level for VLAN 5 and VLAN 6.



NOTE: You can use Juniper Networks switches with VSTP and Cisco switches with PVST+ and Rapid-PVST+ in the same network. Cisco supports a proprietary Per-VLAN Spanning Tree (PVST) protocol, which maintains a separate spanning tree instance per each VLAN. One Spanning Tree per VLAN allows fine grain load balancing but requires more BPDU CPU processing as the number of VLANs increases. PVST runs on Cisco proprietary ISL trunks which is not supported by Juniper. Juniper switches only inter-operate with PVST+ and Rapid-PVST+.



## Topology

Table 190: Interfaces of the Topology for Configuring VSTP

| Interface | Description             |
|-----------|-------------------------|
| ge-0/0/3  | Promiscuous member port |
| ge-0/0/4  | Promiscuous member port |
| ge-0/0/8  | Promiscuous member port |
| ae0       | Promiscuous member port |

| VLAN ID | Description  |
|---------|--------------|
| 200     | Primary VLAN |
| 5       | Primary VLAN |
| 6       | Primary VLAN |

## Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set protocols vstp vlan 200 interface ge-0/0/3
set protocols vstp vlan 200 interface ge-0/0/4
set protocols vstp vlan 200 interface ge-0/0/8
set protocols vstp vlan 200 interface ae0
set protocols vstp interface ge-0/0/3
set protocols vstp interface ge-0/0/4
set protocols vstp interface ge-0/0/8
set protocols vstp vlan 5
set protocols vstp vlan 6
```

## Configuring

**Step-by-Step Procedure** To configure VSTP:

1. Enable VSTP on VLAN 200 using a single VLAN ID:
 

```
[edit protocols]
user@switch #set vstp vlan 200 interface ge-0/0/3
user@switch #set vstp vlan 200 interface ge-0/0/4
user@switch #set vstp vlan 200 interface ge-0/0/8
user@switch #set vstp vlan 200 interface ae0
```



**NOTE:** When you configure VSTP with the `set protocol vstp vlan all` command, VLAN ID 1 is not set; it is excluded so that the configuration is compatible with Cisco PVST+. If you want VLAN ID 1 to be included in the VSTP configuration on your switch, you must set it separately with the `set protocol vstp vlan 1` command.



**TIP:** You could also enable VSTP on a VLAN using a single VLAN name.

2. Enable VSTP on VLAN 5 and VLAN 6 at the VSTP global level:

```
[edit protocols]
user@switch #set vstp interface ge-0/0/3
user@switch #set vstp interface ge-0/0/4
user@switch #set vstp interface ge-0/0/8
user@switch #set vstp vlan 5
user@switch #set vstp vlan 6
```



**CAUTION:** Ensure that the interface is a member of all VLANs before you add the interface to the VSTP configuration. If the interface is not a member of all VLANs, this VSTP configuration will fail when you try to commit it.

## Verifying VSTP Configuration

**Purpose** View the spanning tree interface and bridge you just created.

**Action** To view a VSTP configuration, use the commands `show spanning-tree bridge` and `show spanning-tree interface`.

## Results

The results of this example are:

```
root@sw-16> show spanning-tree interface
```

Spanning tree interface parameters for VLAN 5

| Interface | Port ID<br>port ID | Designated<br>bridge ID | Designated<br>Cost | Port  | State | Role |
|-----------|--------------------|-------------------------|--------------------|-------|-------|------|
| ge-0/0/3  | 128:490            | 128:490                 | 32773.88e0f31f4040 | 20000 | FWD   | DESG |
| ge-0/0/4  | 128:491            | 128:490                 | 32773.88e0f31f4040 | 20000 | BLK   | BKUP |
| ge-0/0/8  | 128:492            | 128:492                 | 32773.88e0f31f4040 | 20000 | FWD   | DESG |

Spanning tree interface parameters for VLAN 6

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Designated<br>Cost | Port  | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------------|-------|-------|------|
| ge-0/0/3  | 128:490 | 128:490               | 128:490                 | 32774.88e0f31f4040 | 20000 | FWD   | DESG |
| ge-0/0/4  | 128:491 | 128:491               | 128:491                 | 32774.88e0f31f4040 | 20000 | FWD   | DESG |
| ge-0/0/8  | 128:492 | 128:492               | 128:492                 | 32774.88e0f31f4040 | 20000 | FWD   | DESG |

## Spanning tree interface parameters for VLAN 200

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Designated<br>Cost | Port  | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------------|-------|-------|------|
| ge-0/0/3  | 128:490 | 128:490               | 128:490                 | 32968.88e0f31f4040 | 20000 | FWD   | DESG |
| ge-0/0/4  | 128:491 | 128:491               | 128:491                 | 32968.88e0f31f4040 | 20000 | FWD   | DESG |
| ge-0/0/8  | 128:492 | 128:492               | 128:492                 | 32968.88e0f31f4040 | 20000 | FWD   | DESG |
| ae0       | 128:3   | 128:3                 | 128:3                   | 32968.88e0f31f4040 | 20000 | FWD   | DESG |

```
{master:0}
root@sw-16> show spanning-tree bridge
STP bridge parameters
Routing instance name : GLOBAL
Enabled protocol : RSTP
```

## STP bridge parameters for VLAN 5

```
Root ID : 32773.88:e0:f3:1f:40:40
Hello time : 2 seconds
Maximum age : 20 seconds
Forward delay : 15 seconds
Message age : 0
Number of topology changes : 1
Time since last topology change : 52 seconds
Local parameters
Bridge ID : 32773.88:e0:f3:1f:40:40
Extended system ID : 5
```

## STP bridge parameters for VLAN 6

```
Root ID : 32774.88:e0:f3:1f:40:40
Hello time : 2 seconds
Maximum age : 20 seconds
Forward delay : 15 seconds
Message age : 0
Number of topology changes : 0
Local parameters
Bridge ID : 32774.88:e0:f3:1f:40:40
Extended system ID : 6
```

## STP bridge parameters for VLAN 200

```
Root ID : 32968.88:e0:f3:1f:40:40
Hello time : 2 seconds
Maximum age : 20 seconds
Forward delay : 15 seconds
Message age : 0
Number of topology changes : 0
Local parameters
Bridge ID : 32968.88:e0:f3:1f:40:40
Extended system ID : 200
```

```
{master:0}
root@sw-16>
```

**Related  
Documentation**

- *Understanding VSTP for EX Series Switches and QFX Series Switches*

## Configuring VSTP (CLI Procedure)



**NOTE:** This topic applies to Junos OS for EX Series and QFX switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring VSTP (CLI Procedure)*. For ELS details, see “Getting Started with Enhanced Layer 2 Software” on page 41.

The default spanning-tree protocol for EX Series switches is Rapid Spanning Tree Protocol (RSTP). VLAN Spanning Tree Protocol (VSTP) is an alternate protocol that allows EX Series switches to run one or more Spanning Tree Protocol (STP) or RSTP instances for each VLAN on which VSTP is enabled. For networks with multiple VLANs, VSTP improves intelligent tree spanning by defining best paths within the VLANs instead of within the entire network.



**NOTE:** On EX Series (other than EX9200) and QFX switches running Junos OS that supports ELS—VSTP can support up to 510 VLANs. However, on EX9200 switches, VSTP can support only up to 253 VLANs.

You can configure VSTP at the global level:

- For all interfaces on the switch
- For all interfaces within all VLANs
- For all interfaces within a specified VLAN
- For all interfaces within a specified VLAN group

You can configure or disable VSTP for a specific interface:

- For a specific interface on the switch
- For a specific interface within all VLANs
- For a specific interface within a specified VLAN
- For a specific interface within a specified VLAN group



**NOTE:**

- If you configure VSTP on an interface at both the global and the specific VLAN level, the interface configuration that is defined at the specific VLAN level overrides the interface configuration that is defined at the global level.
- If you specify VSTP to be configured on an interface that is not configured to belong to the VLAN (or VLANs), an error message is displayed.

To configure VSTP:

- For all interfaces within any of the following scopes:

- For all interfaces on the switch:

```
[edit protocols vstp]
user@switch# set interface all
```

- For all interfaces within all VLANs:

```
[edit protocols vstp]
user@switch# set vlan all interface all
```

- For all interfaces within a specified VLAN:

```
[edit protocols vstp]
user@switch# set vlan (vlan-id |vlan-range |open-set-of-values) interface all interface
all
```

- For all interfaces within a specified VLAN group:

```
[edit protocols vstp]
user@switch# set vlan-group vlan-group-name vlan (vlan-id |vlan-range |open-set-of-values)
interface all
```

- On a specific interface within any of the following scopes:

- For a specific interface on the switch:

```
[edit protocols vstp]
user@switch# set interface interface-name
```

- For a specific interface within all VLANs:

```
[edit protocols vstp]
user@switch# set vlan all interface interface-name
```



**CAUTION:** Ensure that the interface is a member of all VLANs before you add the interface to the VSTP configuration. If the interface is not a member of all VLANs, this VSTP configuration will fail when you try to commit it.

---

- For a specific interface within a specified VLAN:

```
[edit protocols vstp]
user@switch# set vlan vlan-id-or-vlan-range interface interface-name
```

- For a specific interface within a specified VLAN group:

```
[edit protocols vstp]
user@switch# set vlan-group vlan-group-name vlan (vlan-id |vlan-range |open-set-of-values)
interface interface-name
```

**Related  
Documentation**

- *show spanning-tree bridge*
- *show spanning-tree interface*
- *Understanding VSTP for EX Series Switches and QFX Series Switches*

## Understanding BPDU Protection for STP, RSTP, and MSTP



**NOTE:** Using the original CLI, you can disable BPDU protection on interfaces by issuing the `set ethernet-switching-options bpdu-block interface-name disable` command.

A Juniper Networks device Layer 2 loop prevention through Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), and Multiple Spanning Tree Protocol (MSTP). Bridge protocol data unit (BPDU) protection can help prevent STP misconfigurations that can lead to network outages.

A loop-free network is supported through the exchange of a special type of frame called a BPDU. Receipt of BPDUs on certain interfaces in an STP, RSTP, VSTP, or MSTP topology, however, can lead to network outages. Enable BPDU protection on those interfaces to prevent these outages.

Peer STP applications running on the device interfaces use BPDUs to communicate. Ultimately, the exchange of BPDUs determines which interfaces block traffic and which interfaces become root ports and forward traffic.

However, a user bridge application running on a device connected to the device can also generate BPDUs. If these BPDUs are picked up by STP applications running on the device, they can trigger STP miscalculations, and those miscalculations can lead to network outages.

Enable BPDU protection on device interfaces connected to user devices or on interfaces on which no BPDUs are expected, such as edge ports. If BPDUs are received on a protected interface, the interface is disabled and stops forwarding frames.

Not only can you configure BPDU protection on a device with a spanning tree, but also on a device without a spanning tree. This type of topology typically consists of a non-STP device connected to an STP device through a trunk interface.

To configure BPDU protection on a device with a spanning tree, include the **bpdu-block-on-edge** statement at the `[edit protocols (stp | mstp | rstp)]` hierarchy level. To configure BPDU protection on a device without a spanning tree, include the **bpdu-block** statement at the `[edit ethernet-switching-options interface interface-name]` hierarchy level.

If BPDUs are sent to an interface (indicating that the misconfiguration has been corrected), the interface can be unblocked in one of two ways:

- If the **disable-timeout** statement has been included in the BPDU configuration, the interface automatically returns to service after the timer expires.
- Use the `clear error bpdu interface` operational mode command.

Disabling the BPDU protection configuration does not unblock the interface.

- Related Documentation**
- [Understanding Loop Protection for STP, RSTP, VSTP, and MSTP on page 2206](#)
  - [Understanding Root Protection for STP, RSTP, VSTP, and MSTP on page 2211](#)
  - [Understanding MSTP on page 2149](#)
  - [Understanding RSTP on page 2172](#)
  - [Understanding VSTP on page 2191](#)

---

## Example: Configuring BPDU Protection on Edge Interfaces to Prevent STP Miscalculations

---

EX Series and QFX Series switches provide Layer 2 loop prevention through Rapid Spanning Tree protocol (RSTP) and Multiple Spanning Tree Protocol (MSTP). All spanning-tree protocols use a special type of frame called a bridge protocol data unit (BPDU) to communicate. Other devices—PC bridging applications, for example, also use BPDUs and generate their own BPDUs. These different BPDUs are not compatible. When BPDUs generated by spanning-tree protocols are transmitted to a device that uses another type of BPDU, they can cause problems on the device. Similarly, if switches within a spanning-tree topology receive BPDUs from other devices, network outages can occur because of STP miscalculations.

This example configures BPDU protection on an EX Series switch that uses RSTP. The upstream configuration is done on the edge interfaces, where outside BPDUs are often received from other devices:

- [Requirements on page 2200](#)
- [Overview and Topology on page 2200](#)
- [Configuration on page 2202](#)
- [Verification on page 2202](#)

### Requirements

This example uses the following software and hardware components:

- Two EX Series switches in an RSTP topology
- Junos OS Release 13.2X50-D10 or later or later for EX Series or QFX Series switches

Before you configure the interfaces on Switch 2 for BPDU protection, be sure you have:

- RSTP enabled on the switches.



**NOTE:** By default, RSTP is enabled on all EX Series switches.

---

### Overview and Topology

The switches, being in an RSTP topology, support a loop-free network through the exchange of BPDUs. Receipt of outside BPDUs in an RSTP or MSTP topology, however,



can lead to network outages by triggering an STP misconfiguration. To prevent such outages, enable BPDU protection on spanning tree interfaces that could receive outside BPDUs. If an outside BPDU is received on a BPDU-protected interface, the interface shuts down to prevent the outside BPDU from accessing the spanning tree interface.

Figure 26 shows the topology for this example. In this example, Switch 1 and Switch 2 are configured for RSTP and create a loop-free topology. The interfaces on Switch 2 are edge access ports—edge access ports frequently receive outside BPDUs generated by PC applications.

This example configures interface `ge-0/0/5` and interface `ge-0/0/6` as edge ports on Switch 2, and then configures BPDU protection on those ports. With BPDU protection enabled, these interfaces shut down when they encounter an outside BPDU sent by the PCs connected to Switch 2.

Figure 26: BPDU Protection Topology




Table 191 shows the components that will be configured for BPDU protection.

Table 191: Components of the Topology for Configuring BPDU Protection on EX Series Switches

| Property                      | Settings                                                                                                                                                              |
|-------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Switch 1 (Distribution Layer) | Switch 1 is connected to Switch 2 on a trunk interface.                                                                                                               |
| Switch 2 (Access Layer)       | Switch 2 has these access ports that require BPDU protection: <ul style="list-style-type: none"><li>• <code>ge-0/0/5</code></li><li>• <code>ge-0/0/6</code></li></ul> |

This configuration example uses RSTP topology. You also can configure BPDU protection for MSTP topologies at the `[edit protocols mstp]` hierarchy level.

## Configuration

To configure BPDU protection on two access interfaces:

### CLI Quick Configuration

Quickly configure RSTP on the two Switch 2 interfaces, and then configure BPDU protection on all edge ports on Switch 2 by copying the following commands and pasting them into the switch terminal window:



**NOTE:** This example configures BPDU protection on specific interfaces. However, starting with Junos OS Release 15.1 for EX Series and QFX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style, you can configure BPDU protection globally on all spanning tree interfaces. See [“Configuring BPDU Protection on Spanning Tree Interfaces” on page 2205](#) for additional information.

### Step-by-Step Procedure

[edit]

```
set protocols rstp interface ge-0/0/5 edge
set protocols rstp interface ge-0/0/6 edge
set protocols rstp bpdu-block-on-edge
```

To configure RSTP on the two Switch 2 interfaces, and then configure BPDU protection:

1. Configure RSTP on interface **ge-0/0/5** and interface **ge-0/0/6**, and configure them as edge ports:

```
[edit protocols rstp]
user@switch# set interface ge-0/0/5 edge
user@switch# set interface ge-0/0/6 edge
```

2. Configure BPDU protection on all edge ports on this switch:

```
[edit protocols rstp]
user@switch# set bpdu-block-on-edge
```

**Results** Check the results of the configuration:

```
user@switch> show configuration protocols rstp
interface ge-0/0/5 {
 edge;
}
interface ge-0/0/6 {
 edge;
}
bpdu-block-on-edge;
```

## Verification

To confirm that the configuration is working properly:

- [Displaying the Interface State Before BPDU Protection Is Triggered on page 2203](#)
- [Verifying That BPDU Protection Is Working Correctly on page 2203](#)

### Displaying the Interface State Before BPDU Protection Is Triggered

**Purpose** Before BPDUs can be received from PCs connected to interface **ge-0/0/5** and interface **ge-0/0/6**, confirm the interface state.

**Action** Use the operational mode command:

```
user@switch> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-0/0/0  | 128:513 | 128:513               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| ge-0/0/1  | 128:514 | 128:514               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| ge-0/0/2  | 128:515 | 128:515               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| ge-0/0/3  | 128:516 | 128:516               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| ge-0/0/4  | 128:517 | 128:517               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| ge-0/0/5  | 128:518 | 128:518               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| ge-0/0/6  | 128:519 | 128:519               | 32768.0019e2503f00      | 20000        | FWD   | DESG |

[output truncated]

**Meaning** The output from the operational mode command **show spanning-tree interface** shows that **ge-0/0/5** and interface **ge-0/0/6** are ports in a forwarding state.

### Verifying That BPDU Protection Is Working Correctly

**Purpose** In this example, the PCs connected to Switch 2 start sending BPDUs to interface **ge-0/0/5** and interface **ge-0/0/6**. Verify that BPDU protection is working on the interfaces.

**Action** Use the operational mode command:

```
user@switch> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

| Interface                | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|--------------------------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-0/0/0                 | 128:513 | 128:513               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| ge-0/0/1                 | 128:514 | 128:514               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| ge-0/0/2                 | 128:515 | 128:515               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| ge-0/0/3                 | 128:516 | 128:516               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| ge-0/0/4                 | 128:517 | 128:517               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| ge-0/0/5<br>(Bpdu-Incon) | 128:518 | 128:518               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| ge-0/0/6<br>(Bpdu-Incon) | 128:519 | 128:519               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| ge-0/0/7                 | 128:520 | 128:1                 | 16384.00aabbcc0348      | 20000        | FWD   | ROOT |
| ge-0/0/8                 | 128:521 | 128:521               | 32768.0019e2503f00      | 20000        | FWD   | DESG |

[output truncated]

**Meaning** When BPDUs are sent from the PCs to interface **ge-0/0/5** and interface **ge-0/0/6** on Switch 2, the output from the operational mode command **show spanning-tree interface**

shows that the interfaces have transitioned to a BPDU inconsistent state. The BPDU inconsistent state causes the interfaces to shut down.

Disabling the BPDU protection configuration on an interface does not automatically reenabling the interface. However, if the **disable-timeout (Spanning Trees)** statement has been included in the BPDU configuration, the interface does return to service after the timer expires. Otherwise, you must use the operational mode command **clear error bpdu** to unblock and reenabling the interface.

If the PCs connected to Switch 2 send BPDUs to the interfaces again, BPDU protection is triggered once more and the interfaces transition back to the BPDU inconsistent state, causing them to shut down. In such cases, you need to find and repair the misconfiguration on the PCs that is sending BPDUs to Switch 2.

**Related  
Documentation**

- *Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches*
- *Example: Configuring BPDU Protection on Interfaces to Prevent STP Miscalculations on EX Series Switches*
- *Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on EX Series Switches*
- *Example: Configuring Root Protection to Enforce Root Bridge Placement in Spanning Trees on EX Series Switches*
- *Understanding BPDU Protection for STP, RSTP, and MSTP on EX Series Switches*

## Configuring BPDU Protection on Spanning Tree Interfaces



**NOTE:** This topic applies to Junos OS for EX Series and QFX switches with support for the Enhanced Layer 2 Software (ELS) configuration style. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 41](#).

You can configure BPDU protection to ignore BPDU received on interfaces where none should be expected. If a BPDU is received on a blocked interface, the interface is disabled and stops forwarding frames. By default, all BPDUs are accepted and processed on all interfaces.

To configure BPDU protection for spanning-tree instance interfaces:

- On a specific spanning-tree interface:

1. To enable BPDU protection on a specified spanning-tree interface:

```
[edit protocols layer2-control bpu-block]
user@switch# set interface (aex | (ge-fpc/pic/port | xe-fpc/pic/port)
```

If a BPDU is received on the interface, the system will disable the interface and stop forwarding frames out the interface until the bridging process is restarted.

2. (Optional) Configure the amount of time the system waits before *automatically* unblocking this interface after it has received a BPDU.

```
[edit protocols layer2-control bpu-block interface interface-name]
user@switch# set disable-timeout seconds
```

The range of the *seconds* option value is from 10 through 3600 seconds (one hour). A *seconds* option value of 0 is allowed, but this results in the default behavior (the interface is blocked until the interface is cleared).

- To disable BPDU protection for a specific spanning-tree interface

```
[edit protocols layer2-control bpu-block interface interface-name]
user@switch# set disable-timeout seconds
```

### Related Documentation

- [Understanding BPDU Protection for Spanning-Tree Instance Interfaces](#)
- [BPDU Protection for Individual Spanning-Tree Instance Interfaces](#)
- [clear error bpu interface on page 2418](#)

## Understanding Loop Protection for STP, RSTP, VSTP, and MSTP

---

A Juniper Networks device provides Layer 2 loop prevention through Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), and Multiple Spanning Tree Protocol (MSTP). Loop protection increases the efficiency of STP, RSTP, and MSTP by preventing ports from entering a forwarding state that would cause a loop to open in the network.

A loop-free network in spanning-tree topologies is supported through the exchange of a special type of frame called a bridge protocol data unit (BPDU). Peer STP applications running on the device interfaces use BPDUs to communicate. Ultimately, the exchange of BPDUs determines which interfaces block traffic (preventing loops) and which interfaces become root ports and forward traffic.

However, a blocking interface can mistakenly transition to the forwarding state if the interface stops receiving BPDUs from its designated port on the segment. Such a transition error can occur when there is a hardware error on the device or software configuration error between the device and its neighbor.

When loop protection is enabled, the spanning-tree topology detects root ports and blocked ports and ensures that both keep receiving BPDUs. If a loop-protection-enabled interface stops receiving BPDUs from its designated port, it reacts as it would react to a problem with the physical connection on this interface. It does not transition the interface to a forwarding state, but instead transitions it to a loop-inconsistent state. The interface recovers and it transitions back to the spanning-tree blocking state as soon as it receives a BPDU.

We recommend that you enable loop protection on all device interfaces that have a chance of becoming root or designated ports. Loop protection is most effective when enabled in the entire switched network. When you enable loop protection, you must configure at least one action (**alarm**, **block**, or both).

An interface can be configured for either loop protection or root protection, but not for both.

### Related Documentation

- [Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on page 2207](#)
- [Understanding Root Protection for STP, RSTP, VSTP, and MSTP on page 2211](#)
- [Understanding BPDU Protection for STP, RSTP, and MSTP on page 2199](#)
- [Understanding MSTP on page 2149](#)
- [Understanding RSTP on page 2172](#)
- [Overview of Spanning-Tree Protocols on page 2148](#)
- [Understanding VSTP on page 2191](#)

## Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree

The QFX Series products provide Layer 2 loop prevention through Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), and Multiple Spanning Tree Protocol (MSTP). Loop protection increases the efficiency of STP, RSTP, and MSTP by preventing interfaces from moving into a forwarding state that would create a loop in the network.

This example describes how to configure loop protection for an interface for the QFX Series in an RSTP topology:

- [Requirements on page 2207](#)
- [Overview and Topology on page 2207](#)
- [Configuration on page 2209](#)
- [Verification on page 2209](#)

### Requirements

This example uses the following hardware and software components:

- Junos OS Release 11.1 or later for the QFX Series
- Three switches in an RSTP topology



**NOTE:** By default, RSTP is enabled for the QFX Series.

### Overview and Topology

A loop-free network in spanning-tree topologies is supported through the exchange of a special type of frame called a bridge protocol data unit (BPDU). Peer STP applications running on the switch interfaces use BPDUs to communicate. Ultimately, the exchange of BPDUs determines which interfaces block traffic (preventing loops) and which interfaces become root ports and forward traffic.

A blocking interface can transition to the forwarding state in error if the interface stops receiving BPDUs from its designated port on the segment. Such a transition error can occur when there is a hardware error on the switch or software configuration error between the switch and its neighbor. When this happens, a loop appears in the spanning tree. Loops in a Layer 2 topology cause broadcast, unicast, and multicast frames to continuously circle the looped network. As a switch processes a flood of frames in a looped network, its resources become depleted, and the ultimate result is a network outage.



**NOTE:** An interface can be configured for either loop protection or root protection, but not for both.

Three switches are displayed in Figure 27. In this example, they are configured for RSTP and create a loop-free topology. Interface **xe-0/0/6** is blocking traffic between Switch 3 and Switch 1; thus, traffic is forwarded through interface **xe-0/0/7** on Switch 2. BPDUs are being sent from the root bridge on Switch 1 to both of these interfaces.

This example shows how to configure loop protection on interface **xe-0/0/6** to prevent it from transitioning from a blocking state to a forwarding state and creating a loop in the spanning-tree topology.

Figure 27: Network Topology for Loop Protection




Table 192 shows the components that will be configured for loop protection.

Table 192: Topology for Configuring Loop Protection on the QFX Series

| Components | Settings                                                              |
|------------|-----------------------------------------------------------------------|
| Switch 1   | Switch 1 is the root bridge.                                          |
| Switch 2   | Switch 2 has the root port <b>xe-0/0/7</b> .                          |
| Switch 3   | Switch 3 is connected to Switch 1 through interface <b>xe-0/0/6</b> . |

A spanning-tree topology contains ports that have specific roles:

- The *root port* is responsible for forwarding data to the root bridge.
- The *alternate port* is a standby port for the root port. When a root port goes down, the alternate port becomes the active root port.
- The *designated port* forwards data to the downstream network segment or device.



This configuration example uses an RSTP topology. However, you can also configure loop protection for STP or MSTP topologies at the `[edit protocols (mstp | stp)]` hierarchy level.

## Configuration

|                                |                                                                                                                                                                                               |
|--------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>CLI Quick Configuration</b> | To quickly configure loop protection on interface <b>xe-0/0/6</b> :<br><br>[edit]<br><b>set protocols rstp interface xe-0/0/6 bpdu-timeout-action block</b>                                   |
| <b>Step-by-Step Procedure</b>  | To configure loop protection:<br><br>1. Configure interface <b>xe-0/0/6</b> on Switch 3:<br><br>[edit protocols rstp]<br>user@switch# <b>set interface xe-0/0/6 bpdu-timeout-action block</b> |
| <b>Results</b>                 | Check the results of the configuration:<br><br>user@switch> <b>show configuration protocols rstp</b><br>interface xe-0/0/6.0 {<br>bpdu-timeout-action {<br>block;<br>}<br>}                   |

## Verification

To confirm that the configuration is working properly, perform these tasks:

- [Displaying the Interface State Before Loop Protection Is Triggered on page 2209](#)
- [Verifying That Loop Protection Is Working on an Interface on page 2210](#)

### Displaying the Interface State Before Loop Protection Is Triggered

|                |                                                                                                           |
|----------------|-----------------------------------------------------------------------------------------------------------|
| <b>Purpose</b> | Before loop protection is triggered on interface <b>xe-0/0/6</b> , confirm that the interface is blocked. |
|----------------|-----------------------------------------------------------------------------------------------------------|

**Action** Display the interface state and role before applying root protection:

```
user@switch> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

| Interface  | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|------------|---------|-----------------------|-------------------------|--------------|-------|------|
| xe-0/0/0.0 | 128:513 | 128:513               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| xe-0/0/1.0 | 128:514 | 128:514               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| xe-0/0/2.0 | 128:515 | 128:515               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| xe-0/0/3.0 | 128:516 | 128:516               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| xe-0/0/4.0 | 128:517 | 128:517               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| xe-0/0/5.0 | 128:518 | 128:518               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| xe-0/0/6.0 | 128:519 | 128:2                 | 16384.00aabbcc0348      | 20000        | BLK   | ALT  |

[output truncated]

**Meaning** The output from the operational mode command **show spanning-tree interface** shows that **xe-0/0/6.0** is the alternate port and is blocked.

### Verifying That Loop Protection Is Working on an Interface

**Purpose** Verify that the loop protection configuration on interface **xe-0/0/6**. RSTP has been disabled on interface **xe-0/0/4** on Switch 1. This stops BPDUs from being sent to interface **xe-0/0/6** and triggering loop protection on that interface.

**Action** Display the interface state and role after applying root protection:

```
user@switch> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

| Interface  | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|------------|---------|-----------------------|-------------------------|--------------|-------|------|
| xe-0/0/0.0 | 128:513 | 128:513               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| xe-0/0/1.0 | 128:514 | 128:514               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| xe-0/0/2.0 | 128:515 | 128:515               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| xe-0/0/3.0 | 128:516 | 128:516               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| xe-0/0/4.0 | 128:517 | 128:517               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| xe-0/0/5.0 | 128:518 | 128:518               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| xe-0/0/6.0 | 128:519 | 128:519               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |

(Loop-Incon)

[output truncated]

**Meaning** The operational mode command **show spanning-tree interface** shows that interface **xe-0/0/6.0** has detected that BPDUs are no longer being forwarded to it and has moved into a loop-inconsistent state. The loop-inconsistent state prevents the interface from transitioning to a forwarding state. The interface recovers and transitions back to its original state as soon as it receives BPDUs.

**Related Documentation**

- *Example: Configuring Faster Convergence and Improving Network Stability with RSTP*
- [Example: Configuring Root Protection to Enforce Root Bridge Placement in Spanning Trees on page 2212](#)

- [Example: Configuring BPDU Protection on STP Interfaces to Prevent STP Miscalculations](#)
- [Understanding Loop Protection for STP, RSTP, VSTP, and MSTP on page 2206](#)

## Understanding Root Protection for STP, RSTP, VSTP, and MSTP

A Juniper Networks device provides Layer 2 loop prevention through Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), and Multiple Spanning Tree Protocol (MSTP). A loop-free network is supported through the exchange of a special type of frame called a bridge protocol data unit (BPDU). Peer STP applications running on the device interfaces use BPDUs to communicate. Ultimately, the exchange of BPDUs determines which interfaces block traffic and which interfaces become root ports and forward traffic.

You can also see BPDUs generated when you run a bridge application on a device attached to the device. This can interfere with root port election, which may sometimes lead to the wrong root port being elected through the above process. Root protection allows you to manually enforce the root bridge placement in the network.

Enable root protection on interfaces that should not receive higher-priority BPDUs from the root bridge and should not be elected as the root port. These interfaces become designated ports and are typically located on an administrative boundary. If the bridge receives more STP BPDUs on a port that has root protection enabled, that port transitions to a root-prevented STP state (inconsistency state), and the interface is blocked. This blocking prevents a bridge that should not be the root bridge from being elected the root bridge. After the bridge stops receiving more STP BPDUs on the interface with root protection, the interface returns to a listening state, followed by a learning state, and ultimately back to a forwarding state. Recovery back to the forwarding state is automatic.

When root protection is enabled on an interface, it is enabled for all the STP instances on that interface. The interface is blocked only for instances for which it receives more BPDUs. Otherwise, it participates in the spanning-tree topology.

An interface can be configured for either root protection or loop protection, but not for both.

### Related Documentation

- [Example: Configuring Root Protection to Enforce Root Bridge Placement in Spanning Trees on page 2212](#)
- [Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on page 2207](#)
- [Example: Configuring BPDU Protection on STP Interfaces to Prevent STP Miscalculations](#)
- [Understanding MSTP on page 2149](#)
- [Understanding RSTP on page 2172](#)
- [Overview of Spanning-Tree Protocols on page 2148](#)
- [Understanding VSTP on page 2191](#)

## Example: Configuring Root Protection to Enforce Root Bridge Placement in Spanning Trees

---

QFX Series products provide Layer 2 loop prevention through Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), and Multiple Spanning Tree Protocol (MSTP). Root protection increases the efficiency of STP, RSTP, and MSTP by allowing network administrators to enforce the root bridge placement in the network manually.

This example describes how to configure root protection on an interface for the QFX Series.

- [Requirements on page 2212](#)
- [Overview and Topology on page 2212](#)
- [Configuration on page 2214](#)
- [Verification on page 2215](#)

### Requirements

This example uses the following hardware and software components:

- Junos OS Release 11.1 or later for the QFX Series
- Four switches in an RSTP topology

Before you configure the interface for root protection, be sure you have:

- RSTP operating on the switches.



**NOTE:** By default, RSTP is enabled on the QFX Series.

---

### Overview and Topology

Peer STP applications running on switch interfaces exchange a special type of frame called a bridge protocol data unit (BPDU). Switches communicate interface information using BPDUs to create a loop-free topology that ultimately determines the root bridge and which interfaces block or forward traffic in the spanning tree.

You can also see BPDUs generated when you run a bridge application on a device attached to the switch. This can interfere with root port election, which may sometimes lead to the wrong root port being elected through the above process. Root protection allows you to manually enforce the root bridge placement in the network.

To prevent this from happening, enable root protection on interfaces that should not receive more BPDUs from the root bridge and should not be elected as the root port. These interfaces are typically located on an administrative boundary and are designated ports.

When root protection is enabled on an interface:

- The interface is blocked from becoming the root port.
- Root protection is enabled for all STP instances on that interface.
- The interface is blocked only for instances for which it receives more BPDUs. Otherwise, it participates in the spanning-tree topology.




**NOTE:** An interface can be configured for either root protection or loop protection, but not for both.

Four switches are displayed in Figure 28. In this example, they are configured for RSTP and create a loop-free topology. Interface `xe-0/0/7` on Switch 1 is a designated port on an administrative boundary. It connects to Switch 4. Switch 3 is the root bridge. Interface `xe-0/0/6` on Switch 1 is the root port.

This example shows how to configure root protection on interface `xe-0/0/7` to prevent it from transitioning to become the root port.

Figure 28: Network Topology for Root Protection



g040712

Table 193 shows the components that will be configured for root protection.

**Table 193: Topology for Configuring Root Protection on the QFX Series**

| Component | Settings                                                                                                                                                                                    |
|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Switch 1  | Switch 1 is connected to Switch 4 through interface <b>xe-0/0/7</b> .                                                                                                                       |
| Switch 2  | Switch 2 is connected to Switch 1 and Switch 3. Interface <b>xe-0/0/4</b> is the alternate port in the RSTP topology.                                                                       |
| Switch 3  | Switch 3 is the root bridge and is connected to Switch 1 and Switch 2.                                                                                                                      |
| Switch 4  | Switch 4 is connected to Switch 1. After loop protection is configured on interface <b>xe-0/0/7</b> , Switch 4 sends more BPDUs that trigger loop protection on interface <b>xe-0/0/7</b> . |

A spanning-tree topology contains ports that have specific roles:

- The *root port* is responsible for forwarding data to the root bridge.
- The *alternate port* is a standby port for the root port. When a root port goes down, the alternate port becomes the active root port.
- The *designated port* forwards data to the downstream network segment or device.

This configuration example uses an RSTP topology. However, you can also configure root protection for STP or MSTP topologies at the `[edit protocols (mstp | stp)]` hierarchy level.

## Configuration

**CLI Quick Configuration** To quickly configure root protection on interface **xe-0/0/7**, copy the following command and paste it into the switch terminal window:

```
[edit]
set protocols rstp interface xe-0/0/7 no-root-port
```

**Step-by-Step Procedure** To configure root protection:

1. Configure interface **xe-0/0/7**:
 

```
[edit protocols rstp]
user@switch#
set interface xe-0/0/7 no-root-port
```

**Results** Check the results of the configuration:

```
user@switch> show configuration protocols rstp
interface xe-0/0/7.0 {
 no-root-port;
}
```

## Verification

To confirm that the configuration is working properly, perform these tasks:

- [Displaying the Interface State Before Root Protection Is Triggered on page 2215](#)
- [Verifying That Root Protection Is Working on the Interface on page 2215](#)

### Displaying the Interface State Before Root Protection Is Triggered

**Purpose** Before root protection is triggered on interface **xe-0/0/7**, confirm the interface state.

**Action** Confirm the state of the interfaces before root protection is configured:

```
user@switch> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

| Interface  | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|------------|---------|-----------------------|-------------------------|--------------|-------|------|
| xe-0/0/0.0 | 128:513 | 128:513               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| xe-0/0/1.0 | 128:514 | 128:514               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| xe-0/0/2.0 | 128:515 | 128:515               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| xe-0/0/3.0 | 128:516 | 128:516               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| xe-0/0/4.0 | 128:517 | 128:517               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| xe-0/0/5.0 | 128:518 | 128:2                 | 16384.00aabbcc0348      | 20000        | BLK   | ALT  |
| xe-0/0/6.0 | 128:519 | 128:1                 | 16384.00aabbcc0348      | 20000        | FWD   | ROOT |
| xe-0/0/7.0 | 128:520 | 128:520               | 32768.0019e2503f00      | 20000        | FWD   | DESG |

[output truncated]

**Meaning** The output from the operational mode command **show spanning-tree interface** shows that **xe-0/0/7.0** is a designated port in a forwarding state.

### Verifying That Root Protection Is Working on the Interface

**Purpose** A configuration change takes place on Switch 4. A lower bridge priority on Switch 4 causes it to send more BPDUs to interface **xe-0/0/7**. Receipt of more BPDUs on interface **xe-0/0/7** triggers root protection. Verify that root protection is operating on interface **xe-0/0/7**.

**Action** Verify that root protection has been configured and is operating correctly:

```
user@switch> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

| Interface  | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|------------|---------|-----------------------|-------------------------|--------------|-------|------|
| xe-0/0/0.0 | 128:513 | 128:513               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| xe-0/0/1.0 | 128:514 | 128:514               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| xe-0/0/2.0 | 128:515 | 128:515               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| xe-0/0/3.0 | 128:516 | 128:516               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| xe-0/0/4.0 | 128:517 | 128:517               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| xe-0/0/5.0 | 128:518 | 128:2                 | 16384.00aabbcc0348      | 20000        | BLK   | ALT  |
| xe-0/0/6.0 | 128:519 | 128:1                 | 16384.00aabbcc0348      | 20000        | FWD   | ROOT |
| xe-0/0/7.0 | 128:520 | 128:520               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |

(Root-Incon)  
[output truncated]

**Meaning** The operational mode command [show spanning-tree interface](#) shows that interface **xe-0/0/7.0** has transitioned to a loop inconsistent state. The loop inconsistent state blocks the interface and prevents it from becoming a candidate for the root port. When the root bridge no longer receives more STP BPDUs from the interface, the interface recovers and transitions back to a forwarding state. Recovery is automatic.

- Related Documentation**
- *Example: Configuring Faster Convergence and Improving Network Stability with RSTP*
  - [Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on page 2207](#)
  - *Example: Configuring BPDU Protection on STP Interfaces to Prevent STP Miscalculations*
  - [Understanding Root Protection for STP, RSTP, VSTP, and MSTP on page 2211](#)



## Unblocking an Interface That Receives BPDUs in Error (CLI Procedure)

EX Series and QFX Series switches use bridge protocol data unit (BPDU) protection on interfaces to prevent them from receiving BPDUs that could trigger a spanning-tree misconfiguration. If BPDUs are received on a BPDU-protected interface, the interface either shuts down or transitions to a blocking state and stops forwarding frames. In the latter scenario, after the misconfiguration that triggered the BPDUs being sent to an interface is fixed in the topology, the interface can be unblocked and returned to service.

To unblock an interface and return it to service using the CLI:

- Automatically unblock an interface by configuring a timer that expires:

```
[edit protocols layer 2]
```

```
user@switch# set protocols layer2-control bpdv-block disable-timeout 30
```

All interfaces on the switch will be reenabled (unblocked) after the timer expires.

However, once an interface on the switch receives a new spanning-tree protocol BPDU, the interface returns to the blocked state.

- Manually unblock an interface using the operational mode command:

```
user@switch> clear error bpdv interface ge-0/0/6
```

This command will only reenable an interface but the BPDU configuration for the interface will continue to exist unless you remove the BPDU configuration explicitly.

### Related Documentation

- [Example: Configuring BPDU Protection on Edge Interfaces to Prevent STP Miscalculations on page 2200](#)
- [Example: Configuring BPDU Protection on Interfaces to Prevent STP Miscalculations on EX Series Switches](#)
- [Understanding BPDU Protection for STP, RSTP, and MSTP on EX Series Switches](#)



## PART 32

# Q-in-Q Tunneling

- [Using Q-in-Q Tunneling on page 2221](#)



# Using Q-in-Q Tunneling

- [Understanding Q-in-Q Tunneling on page 2221](#)
- [Configuring Q-in-Q Tunneling on page 2226](#)
- [Configuring All-in-One Bundling on page 2233](#)
- [Configuring Many-to-Many Bundling on page 2235](#)
- [Configuring a Specific Interface Mapping with VLAN ID Translation Option on page 2237](#)

## Understanding Q-in-Q Tunneling

---



**NOTE:** This topic applies to Junos OS switches with support for the Enhanced Layer 2 Software (ELS) configuration style. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 41](#).

Q-in-Q tunneling enables service providers on Ethernet access networks to extend a Layer 2 Ethernet connection between two customer sites. Using Q-in-Q tunneling, providers can also segregate or bundle customer traffic into fewer VLANs or different VLANs by adding another layer of 802.1Q tags. Q-in-Q tunneling is useful when customers have overlapping VLAN IDs because customers' VLAN (C-VLAN) tags are prepended by the service-provider VLAN (S-VLAN) tag, which allows you to preserve each customers' VLAN IDs without conflict. The Juniper Networks Junos operating system (Junos OS) implementation of Q-in-Q tunneling supports the IEEE 802.1ad standard.

This topic describes:

- [How Q-in-Q Tunneling Works on page 2222](#)
- [How VLAN Translation Works on page 2222](#)
- [Sending and Receiving Untagged Packets on page 2222](#)
- [Disabling MAC Address Learning on page 2223](#)
- [Mapping C-VLANs to S-VLANs on page 2223](#)
- [Constraints for Q-in-Q Tunneling and VLAN Translation on page 2224](#)

## How Q-in-Q Tunneling Works

In Q-in-Q tunneling, as a packet travels from a C-VLAN to an S-VLAN, a service-provider-specific 802.1Q tag is added to the packet. This additional tag is used to segregate traffic into S-VLANs. The original customer 802.1Q tag of the packet is retained and is transmitted transparently, passing through the service provider's network. As the packet leaves the S-VLAN in the downstream direction, the additional 802.1Q tag is removed.

When Q-in-Q tunneling is enabled, trunk interfaces are assumed to be part of the service provider or data center network. Access interfaces are assumed to be customer-facing and accept both tagged and untagged frames. This topic refers to trunk interfaces as S-VLAN interfaces. This type of interface is also known as a network-to-network interface (NNI). The topic refers to access interfaces as C-VLAN interfaces. This type of interface is also known as a user-network interface (UNI).



---

### NOTE:

---

An interface can be a member of multiple S-VLANs. You can map one C-VLAN to one S-VLAN (1:1) or many C-VLANs to many S-VLANs (N:N). C-VLAN and S-VLAN tags are unique—for instance, you can have both a C-VLAN tag of 101 and an S-VLAN tag of 101. You can limit the set of accepted customer tags to a range of tags or to discrete values. Class-of-service (CoS) values of C-VLANs are unchanged in the downstream direction. You may copy ingress priority and CoS settings to the S-VLAN.

C-VLAN and S-VLAN interfaces accept priority-tagged packets without any configuration.

## How VLAN Translation Works

VLAN translation replaces an incoming C-VLAN tag with an S-VLAN tag instead of adding an additional tag. The C-VLAN tag is therefore lost, so a single-tagged packet is normally untagged when it leaves the S-VLAN (at the other end of the link). If an incoming packet has had Q-in-Q tunneling applied in advance, VLAN translation replaces the outer tag and the inner tag is retained when the packet leaves the S-VLAN at the other end of the link.

To configure VLAN translation, use the *mapping swap* statement at the **[edit vlans interface]** hierarchy level.



---

**NOTE:** You can configure VLAN translation on access ports only. You cannot configure it on trunk ports, and you cannot configure Q-in-Q tunneling on the same access port.

---

## Sending and Receiving Untagged Packets

To enable an interface to send and receive untagged packets, you must specify a native VLAN for a physical interface. When the interface receives an untagged packet, it adds

the VLAN ID of the native VLAN to the packet and sends the newly tagged packet to the mapped interface.

To specify a native VLAN, use the **native-vlan-id** statement at the **[edit interfaces interface-name]** hierarchy level. The native VLAN ID must match the C-VLAN or S-VLAN ID or be included in the VLAN ID list specified on the logical interface.

For example, on a logical interface for a C-VLAN interface, you might specify a C-VLAN ID list of 100-200. Then, on the C-VLAN physical interface, you could specify a native VLAN ID of 150. This configuration would work because the native VLAN of 150 is included in the C-VLAN ID list of 100-200.

We recommend configuring a native VLAN when using any of the approaches to map C-VLANs to S-VLANs. If you do not configure a native VLAN on an interface, untagged packets received by the interface are discarded. See the Mapping C-VLANs to S-VLANs section in this topic for information about the methods of mapping C-VLANs to S-VLANs.

## Disabling MAC Address Learning

In a Q-in-Q deployment, customer packets interfaces are transported without any changes to source and destination MAC addresses. You can disable MAC address learning at the global, interface, and VLAN levels:

- To disable learning globally, disable MAC address learning for the switch.
- To disable learning for an interface, disable MAC address learning for all VLANs of which the specified interface is a member.
- To disable learning for a VLAN, disable MAC address learning for a specified VLAN.

## Mapping C-VLANs to S-VLANs

There are three ways to map C-VLANs to S-VLANs:

- [All-in-One Bundling on page 2223](#)
- [Many-to-Many Bundling on page 2224](#)
- [Mapping a Specific Interface on page 2224](#)

If you configure multiple mapping methods, the switch gives priority to mapping a specific interface, then to many-to-many bundling, and last to all-in-one bundling. However, for a particular mapping method, setting up overlapping rules for the same C-VLAN is not supported.

### All-in-One Bundling

All-in-one bundling maps all packets from all C-VLAN interfaces to an S-VLAN.

The C-VLAN interface accepts untagged and single-tagged packets. An S-VLAN 802.1Q tag is then added to these packets, and the packets are sent to the S-VLAN interface, which accepts untagged, single-tagged, and double-tagged packets.



**NOTE:** The C-VLAN and S-VLAN interfaces accept untagged packets provided that the `native-vlan-id` statement is configured on these interfaces.

---

### Many-to-Many Bundling

Many-to-many bundling is used to specify which C-VLANs are mapped to which S-VLANs.

Use many-to-many bundling when you want a subset of the C-VLANs on the access switch to be part of multiple S-VLANs. With many-to-many bundling, the C-VLAN interfaces accept untagged and single-tagged packets. An S-VLAN 802.1Q tag is then added to these packets, and the packets are sent to the S-VLAN interfaces, which accept untagged, single-tagged, and double-tagged packets.



**NOTE:** The C-VLAN and S-VLAN interfaces accept untagged packets provided that the `native-vlan-id` statement is configured on these interfaces.

---

### Mapping a Specific Interface

Use specific interface mapping when you want to assign an S-VLAN to a specific C-VLAN on an interface. The configuration applies only to the specific interface, not to all access interfaces.

Specific interface mapping has two suboptions: **push** and **swap**. When traffic that is mapped to a specific interface is pushed, the packet retains its original tag as it moves from the C-VLAN to the S-VLAN and an additional S-VLAN tag is added to the packet. When traffic that is mapped to a specific interface is swapped, the incoming tag is replaced with a new VLAN tag. This is sometimes known as VLAN rewriting or VLAN translation.

Typically, this method is used to keep data from different customers separate or to provide individualized treatment of the packets on a certain interface. You might also use this method to map VLAN traffic from different customers to a single S-VLAN.

When using specific interface mapping, the C-VLAN interfaces accept untagged and single-tagged packets, while the S-VLAN interfaces accept untagged, single-tagged, and double-tagged packets.



**NOTE:** The C-VLAN and S-VLAN interfaces accept untagged packets provided that the `native-vlan-id` statement is configured on these interfaces.

---

## Constraints for Q-in-Q Tunneling and VLAN Translation

Be aware of the following constraints when configuring Q-in-Q tunneling and VLAN translation:



- With releases of Junos OS 13.2X51 previous to 13.2X51-D20, you cannot create a regular VLAN on an interface if you have created an S-VLAN or C-VLAN on that interface for Q-in-Q tunneling. This means that you cannot create an integrated routing and bridging (IRB) interface on that interface because regular VLANs are a required part of IRB configuration. With Junos OS 13.2X51-D25, you can create a regular VLAN on a trunk interface that has an S-VLAN, which means that you can also create an IRB interface on the trunk. In this case, the regular VLAN and S-VLAN on the same trunk interface cannot share the same VLAN ID. Junos OS 13.2X51-D25 does not allow you to create a regular VLAN on an access interface that has a C-VLAN.
- Most access port security features are not supported with Q-in-Q tunneling and VLAN translation.
- Configuring Q-in-Q tunneling and VLAN rewriting/VLAN translation on the same port is not supported.
- You can configure at most one VLAN rewrite/VLAN translation for a given VLAN and interface. For example, you can create no more than one translation for VLAN 100 on interface xe-0/0/0.
- The combined total of VLANs and rules for Q-in-Q tunneling and VLAN translation cannot exceed 6000. For example, you can configure and commit 4000 VLANs and 2000 rules for Q-in-Q tunneling and VLAN translation. However, you cannot configure 4000 VLANs and 2500 rules for Q-in-Q tunneling and VLAN translation. If you try to commit a configuration that exceeds the limit, you see CLI and syslog errors that inform you about the problem.
- MAC addresses are learned from S-VLANs, not C-VLANs.
- Broadcast, unknown unicast, and multicast traffic is forwarded to all members in the S-VLAN.
- The following features are not supported with Q-in-Q tunneling:
  - DHCP relay
  - Fibre Channel over Ethernet
  - IP Source Guard
- The following features are not supported with VLAN rewriting/VLAN translation:
  - Fibre Channel over Ethernet
  - Firewall filter applied to a port or VLAN in the output direction
  - Private VLANs
  - VLAN Spanning Tree Protocol
  - Reflective relay

**Related  
Documentation**

- [Configuring Q-in-Q Tunneling on page 2226](#)

## Configuring Q-in-Q Tunneling

---

Q-in-Q tunneling and VLAN translation allow service providers to create a Layer 2 Ethernet connection between two customer sites. Providers can segregate different customers' VLAN traffic on a link (for example, if the customers use overlapping VLAN IDs) or bundle different customer VLANs into a single service VLAN. Data centers can use Q-in-Q tunneling and VLAN translation to isolate customer traffic within a single site or to enable customer traffic flows between cloud data centers in different geographic locations.

Q-in-Q tunneling adds a service VLAN tag before the customer's 802.1Q VLAN tags. The Juniper Networks Junos operating system implementation of Q-in-Q tunneling supports the IEEE 802.1ad standard.



**NOTE:** This task uses a Junos OS release that supports the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring Q-in-Q Tunneling*.

With releases of Junos OS 13.2X51 previous to 13.2X51-D20, you cannot create a regular VLAN on an interface if you have created an S-VLAN or C-VLAN on that interface for Q-in-Q tunneling. This means that you cannot create an integrated routing and bridging (IRB) interface on that interface because regular VLANs are a required part of IRB configuration. With Junos OS 13.2X51-D25, you can create a regular VLAN on a trunk interface that has an S-VLAN, which means that you can also create an IRB interface on the trunk. In this case, the regular VLAN and S-VLAN on the same trunk interface cannot share the same VLAN ID. Junos OS 13.2X51-D25 does not allow you to create a regular VLAN on an access interface that has a C-VLAN.

Before setting up Q-in-Q tunneling, make sure you have created and configured the necessary customer VLANs on the neighboring switches. See *Configuring VLANs*.

- [Using the Different Mapping Methods on page 2226](#)
- [Configuring All-in-One Bundling on page 2227](#)
- [Configuring Many-to-Many Bundling on page 2228](#)
- [Configuring a Specific Interface Mapping with VLAN ID Translation Option on page 2231](#)

### Using the Different Mapping Methods

Once you have created the required VLANs on the neighboring switches, configure Q-in-Q tunneling using one of the three methods to map customer VLANs (C-VLANs) to service-provider-defined service VLANs (S-VLANs):

- All-in-one bundling maps all packets from all C-VLAN interfaces to an S-VLAN. For information about how to use this method, see [“Configuring All-in-One Bundling” on page 2227](#).
- Use many-to-many bundling when you want a subset of the C-VLANs on the access switch to be part of multiple S-VLANs. For information about how to use this method, see [“Configuring Many-to-Many Bundling” on page 2228](#).

- Use specific interface mapping when you want to assign an S-VLAN to a specific C-VLAN on an interface. For information about how to use this method, see [“Configuring a Specific Interface Mapping with VLAN ID Translation Option”](#) on page 2231.

## Configuring All-in-One Bundling

You can configure Q-in-Q tunneling using the all-in-one bundling method, which forwards all packets that ingress on a C-VLAN interface to an S-VLAN. (Packets are forwarded to the S-VLAN regardless of whether they are tagged or untagged prior to ingress.) Using this approach saves you the effort of specifying a specific mapping for each C-VLAN.

First configure the S-VLAN and its interface:

1. Assign a logical interface (unit) to be a member of the S-VLAN.

```
[edit vlans vlan-name]
user@switch# interface interface-name.unit-number
```



**NOTE:** Do not use logical interface unit 0. You must later bind a VLAN tag ID to the unit you specify in this step, and you cannot bind a VLAN tag ID to unit 0. Also note that you do not create a VLAN ID for the S-VLAN. The ID is created automatically for the appropriate logical interface.

2. Enable the interface to transmit packets with two 802.1Q VLAN tags:

```
[edit interfaces interface-name]
user@switch# flexible-vlan-tagging
```

3. Enable extended VLAN bridge encapsulation on the interface:

```
[edit interfaces interface-name]
user@switch# encapsulation extended-vlan-bridge
```

4. Enable the S-VLAN interface to send and receive untagged packets:

```
[edit interfaces interface-name]
user@switch# native-vlan-id vlan-id
```

5. Bind the logical interface (unit) of the interface that you specified in step 1 to the automatically-created VLAN ID for the S-VLAN:

```
[edit interfaces interface-name unit logical-unit-number]
user@switch# vlan-id number
```

For example, the following configuration makes xe-0/0/0.10 a member of VLAN 10, enables Q-in-Q tunneling on interface xe-0/0/0, enables xe-0/0/0 to accept untagged packets, and binds the VLAN ID of S-VLAN v10 to a logical interface of xe-0/0/0.

```
set vlans v10 interface xe-0/0/0.10
set interfaces xe-0/0/0 flexible-vlan-tagging
set interfaces xe-0/0/0 native-vlan-id 10
set interfaces xe-0/0/0 encapsulation extended-vlan-bridge
set interfaces xe-0/0/0 unit 10 vlan-id 10
```

Now configure all-in-one bundling on a C-VLAN interface:

1. Assign a logical interface (unit) of the C-VLAN interface to be a member of the S-VLAN.

```
[edit vlans vlan-name]
user@switch# interface interface-name.unit-number
```

2. Enable the interface to transmit packets with 802.1Q VLAN tags :

```
[edit interfaces interface-name]
user@switch# flexible-vlan-tagging
```

3. Enable extended VLAN bridge encapsulation on the interface:

```
[edit interfaces interface-name]
user@switch# encapsulation extended-vlan-bridge
```

4. Enable the C-VLAN interface to send and receive untagged packets:

```
[edit interfaces interface-name]
user@switch# native-vlan-id vlan-id
```

5. Configure a logical interface to receive and forward any tagged packet whose VLAN ID tag matches the list of VLAN IDs you specify:

```
[edit interfaces interface-name unit logical-unit-number]
user@switch# vlan-id-list vlan-id-numbers
```



**WARNING:** On some EX and QFX Series switches, you can apply no more than eight VLAN identifier lists to a physical interface.

6. Configure the system to add an S-VLAN tag (outer tag) as packets travel from a C-VLAN interface to the S-VLAN:

```
[edit interfaces interface-name unit logical-unit-number]
user@switch# input-vlan-map push
```

7. Configure the system to remove the S-VLAN tag when packets are forwarded (internally) from the S-VLAN interface to the C-VLAN interface:

```
[edit interfaces interface-name unit logical-unit-number]
user@switch# output-vlan-map pop
```

For example, the following configuration makes xe-0/0/1.10 a member of S-VLAN v10, enables Q-in-Q tunneling, maps packets from C-VLANs 100 through 200 to S-VLAN 10, and enables xe-0/0/1 to accept untagged packets. If a packet originates in C-VLAN 100 and needs to be sent across the S-VLAN, a tag with VLAN ID 10 is added to the packet. When a packet is forwarded (internally) from the S-VLAN interface to interface xe-0/0/1, the tag with VLAN ID 10 is removed.

```
set vlans v10 interface xe-0/0/1.10
set interfaces xe-0/0/1 flexible-vlan-tagging
set interfaces xe-0/0/1 encapsulation extended-vlan-bridge
set interfaces xe-0/0/1 unit 10 vlan-id-list 100-200
set interfaces xe-0/0/1 native-vlan-id 150
set interfaces xe-0/0/1 unit 10 input-vlan-map push
set interfaces xe-0/0/1 unit 10 output-vlan-map pop
```

## Configuring Many-to-Many Bundling

You can configure Q-in-Q tunneling using the many-to-many bundling method, which maps packets from multiple C-VLANs to multiple S-VLANs. This method is convenient

for mapping a range of C-VLANs without having to specify each one individually. (You can also use this method to configure only one C-VLAN to be mapped to an S-VLAN.)

First configure the S-VLANs and assign them to an interface:

1. Assign a logical interface (unit) to be a member of one of the S-VLANs. Do not use logical interface unit 0.

```
[edit vlans vlan-name]
```

```
user@switch# interface interface-name.unit-number
```



**NOTE:** Note that you do not create a VLAN ID for the S-VLAN. The ID is created automatically for the appropriate logical interface.

2. Repeat step 1 for the other S-VLANs.
3. Enable the physical interface to transmit packets with two 802.1Q VLAN tags:

```
[edit interfaces interface-name]
```

```
user@switch# flexible-vlan-tagging
```

4. Enable extended VLAN bridge encapsulation on the interface:

```
[edit interfaces interface-name]
```

```
user@switch# encapsulation extended-vlan-bridge
```

5. Enable the S-VLAN interface to send and receive untagged packets:

```
[edit interfaces interface-name]
```

```
user@switch# native-vlan-id vlan-id
```

6. Bind one of the logical units of the interface to the VLAN ID for one of the S-VLANs.

```
[edit interfaces interface-name unit logical-unit-number]
```

```
user@switch# vlan-id number
```

7. Repeat step 6 to bind the automatically-created VLAN IDs for the other S-VLANs to the other logical units of the interface:

For example, the following configuration creates S-VLANs v10 and v30 and associates them with interface xe-0/0/0. It also enables Q-in-Q tunneling, enables xe-0/0/0 to accept untagged packets, and maps incoming C-VLAN packets to S-VLANs v10 and v30.

```
set vlans v10 interface xe-0/0/0.10
```

```
set vlans v30 interface xe-0/0/0.30
```

```
set interfaces xe-0/0/0 flexible-vlan-tagging
```

```
set interfaces xe-0/0/0 native-vlan-id 10
```

```
set interfaces xe-0/0/0 encapsulation extended-vlan-bridge
```

```
set interfaces xe-0/0/0 unit 10 vlan-id 10
```

```
set interfaces xe-0/0/0 unit 30 vlan-id 30
```

To configure the many-to-many bundling method on a C-VLAN interface, perform the following steps for each customer:

1. Assign a logical interface (unit) of one C-VLAN interface to be a member of one S-VLAN.

```
[edit vlans vlan-name]
```

```
user@switch# interface interface-name.unit-number
```

2. Repeat step 1 to assign another C-VLAN interface (physical interface) to be a member of another S-VLAN.

3. Enable the interface to transmit packets with 802.1Q VLAN tags:

```
[edit interfaces interface-name]
user@switch# flexible-vlan-tagging
```

4. Enable extended VLAN bridge encapsulation on the interface:

```
[edit interfaces interface-name]
user@switch# encapsulation extended-vlan-bridge
```

5. Enable the C-VLAN interface to send and receive untagged packets:

```
[edit interfaces interface-name]
user@switch# native-vlan-id vlan-id
```

6. For each physical interface, configure a logical interface (unit) to receive and forward any tagged packet whose VLAN ID tag matches the list of VLAN IDs you specify:

```
[edit interfaces interface-name unit logical-unit-number]
user@switch# vlan-id-list vlan-id-numbers
```

To configure only one C-VLAN to be mapped to an S-VLAN, specify only one VLAN ID after *vlan-id-list*.



**WARNING:** On some EX and QFX Series switches, you can apply no more than eight VLAN identifier list to a physical interface.

7. For each physical interface, configure the system to add an S-VLAN tag (outer tag) as packets travel from the C-VLAN interface to the S-VLAN:

```
[edit interfaces interface-name unit logical-unit-number]
user@switch# input-vlan-map push
```

8. For each physical interface, configure the system to remove the S-VLAN tag when packets are forwarded from the S-VLAN interface to the C-VLAN interface:

```
[edit interfaces interface-name unit logical-unit-number]
user@switch# output-vlan-map pop
```

For example, the following configuration makes xe-0/0/1.10 a member of S-VLAN v10, enables Q-in-Q tunneling, and maps packets from C-VLANs 10 through 20 to S-VLAN 10. The configuration for customer 2 makes xe-0/0/2.30 a member of S-VLAN v30, enables Q-in-Q tunneling, and maps packets from C-VLANs 30 through 40, 50 through 60, and 70 through 80 to S-VLAN 30. Both interfaces are configured to accept untagged packets.

If a packet originates in C-VLAN 10 and needs to be sent over the S-VLAN, a tag with a VLAN ID 10 is added to the packet. If a packet is forwarded internally from the S-VLAN interface to xe-0/0/1.10, the tag with VLAN ID 10 is removed. The same principles apply to the C-VLANs configured on interface xe-0/0/2.



**NOTE:** Notice that you can use the same tag value for an S-VLAN and C-VLAN. For example, the configuration for customer 1 maps C-VLAN ID 10 to S-VLAN ID 10. C-VLAN and S-VLAN tags use separate name spaces, so this configuration is allowed.

Configuration for customer 1:

```

set vlans v10 interface xe-0/0/1.10
set interfaces xe-0/0/1 flexible-vlan-tagging
set interfaces xe-0/0/1 encapsulation extended-vlan-bridge
set interfaces xe-0/0/1 unit 10 vlan-id-list 10-20
set interfaces xe-0/0/1 native-vlan-id 15
set interfaces xe-0/0/1 unit 10 input-vlan-map push
set interfaces xe-0/0/1 unit 10 output-vlan-map pop

```

Configuration for customer 2:

```

set vlans v30 interface xe-0/0/2.30
set interfaces xe-0/0/2 flexible-vlan-tagging
set interfaces xe-0/0/2 encapsulation extended-vlan-bridge
set interfaces xe-0/0/2 unit 30 vlan-id-list 30-40
set interfaces xe-0/0/2 unit 30 vlan-id-list 50-60
set interfaces xe-0/0/2 unit 30 vlan-id-list 70-80
set interfaces xe-0/0/2 native-vlan-id 75
set interfaces xe-0/0/2 unit 30 input-vlan-map push
set interfaces xe-0/0/2 unit 30 output-vlan-map pop

```

## Configuring a Specific Interface Mapping with VLAN ID Translation Option

You can configure Q-in-Q tunneling by mapping packets from a specified C-VLAN to a specified S-VLAN. In addition, you can configure the system to replace a C-VLAN tag with an S-VLAN tag or replace an S-VLAN tag with a C-VLAN tag (instead of double tagging). This is called VLAN translation or VLAN rewriting. VLAN translation is particularly useful if a service provider's Layer 2 network that connects a customer's sites does not support double tagged packets.

When you use VLAN translation, both ends of the link normally must be able to swap the tags appropriately. That is, both ends of the link must be configured to swap the C-VLAN tag for the S-VLAN tag and swap the S-VLAN tag for the C-VLAN tag so that traffic in both directions is tagged appropriately while in transit and after arrival.

First configure the S-VLAN and its interface:

1. Assign a logical interface to be a member of the S-VLAN. Do not use unit 0.

```
[edit vlans vlan-name]
```

```
user@switch# interface interface-name.unit-number
```



**NOTE:** Note that you do not create a VLAN ID for the S-VLAN. The ID is created automatically for the appropriate logical interface.

2. Enable the interface to transmit packets with 802.1Q VLAN tags:

```
[edit interfaces interface-name]
```

```
user@switch# flexible-vlan-tagging
```

3. Enable the S-VLAN interface to send and receive untagged packets:

```
[edit interfaces interface-name]
```

```
user@switch# native-vlan-id vlan-id
```

4. Enable extended VLAN bridge encapsulation on the interface:

```
[edit interfaces interface-name]
```

```
user@switch# encapsulation extended-vlan-bridge
```

5. Bind the logical interface (unit) of the interface that you specified earlier to the VLAN ID for the S-VLAN:

```
[edit interfaces interface-name unit logical-unit-number]
user@switch# vlan-id number
```

For example, the following configuration creates S-VLAN v200, makes xe-0/0/0.200 a member of that VLAN, enables Q-in-Q tunneling on interface xe-0/0/0, enables xe-0/0/0 to accept untagged packets, and binds a logical interface of xe-0/0/0 to the VLAN ID of VLAN v200.

```
set vlans v200 interface xe-0/0/0.200
set interfaces xe-0/0/0 flexible-vlan-tagging
set interfaces xe-0/0/0 native-vlan-id 10
set interfaces xe-0/0/0 encapsulation extended-vlan-bridge
set interfaces xe-0/0/0 unit 200 vlan-id 200
```

Now configure a specific interface mapping with optional VLAN ID translation on the C-VLAN interface:

1. Assign a logical interface of the C-VLAN interface to be a member of the S-VLAN.

```
[edit vlans vlan-name]
user@switch# interface interface-name.unit-number
```

2. Enable the interface to transmit packets with 802.1Q VLAN tags:

```
[edit interfaces interface-name]
user@switch# flexible-vlan-tagging
```

3. Enable the C-VLAN interface to send and receive untagged packets:

```
[edit interfaces interface-name]
user@switch# native-vlan-id vlan-id
```

4. Enable extended VLAN bridge encapsulation on the interface:

```
[edit interfaces interface-name]
user@switch# encapsulation extended-vlan-bridge
```

5. Configure a logical interface (unit) to receive and forward any tagged packet whose VLAN ID tag matches the VLAN IDs you specify:

```
[edit interfaces interface-name unit logical-unit-number]
user@switch# vlan-id number
```

6. Configure the system to remove the existing C-VLAN tag and replace it with the S-VLAN tag when packets ingress on the C-VLAN interface and are forwarded to the S-VLAN:

```
[edit interfaces interface-name unit logical-unit-number]
user@switch# input-vlan-map swap
```

7. Configure the system to remove the existing S-VLAN tag and replace it with the C-VLAN tag when packets are forwarded from the S-VLAN interface to the C-VLAN interface:

```
[edit interfaces interface-name unit logical-unit-number]
user@switch# output-vlan-map swap
```

8. To configure an S-VLAN and associate it with the appropriate C-VLAN interface:

```
[edit vlans vlan-name]
user@switch# interface interface-name
```

For example, the following configuration on C-VLAN interface xe-0/0/1 enables Q-in-Q tunneling, enables xe-0/0/1 to accept untagged packets, and maps incoming packets



from C-VLAN 150 to logical interface 200, which is a member of S-VLAN 200. Also, when packets egress from C-VLAN interface xe-0/0/1 and travel to the S-VLAN interface, the C-VLAN tag of 150 is removed and replaced with the S-VLAN tag of 200. When packets travel from the S-VLAN interface to the C-VLAN interface, the S-VLAN tag of 200 is removed and replaced with the C-VLAN tag of 150.

```
set vlans v200 interface xe-0/0/1.200
set interfaces xe-0/0/1 flexible-vlan-tagging
set interfaces xe-0/0/1 native-vlan-id 10
set interfaces xe-0/0/1 encapsulation extended-vlan-bridge
set interfaces xe-0/0/1 unit 200 vlan-id 150
set interfaces xe-0/0/1 unit 200 output-vlan-map swap
set interfaces xe-0/0/1 unit 200 input-vlan-map swap
```

**Related Documentation** • [Understanding Q-in-Q Tunneling on page 2221](#)

## Configuring All-in-One Bundling

You can configure Q-in-Q tunneling using the all-in-one bundling method, which forwards all packets that ingress on a C-VLAN interface to an S-VLAN. (Packets are forwarded to the S-VLAN regardless of whether they are tagged or untagged prior to ingress.) Using this approach saves you the effort of specifying a specific mapping for each C-VLAN.

First configure the S-VLAN and its interface:

1. Assign a logical interface (unit) to be a member of the S-VLAN.

```
[edit vlans vlan-name]
user@switch# interface interface-name.unit-number
```



**NOTE:** Do not use logical interface unit 0. You must later bind a VLAN tag ID to the unit you specify in this step, and you cannot bind a VLAN tag ID to unit 0. Also note that you do not create a VLAN ID for the S-VLAN. The ID is created automatically for the appropriate logical interface.

2. Enable the interface to transmit packets with two 802.1Q VLAN tags:

```
[edit interfaces interface-name]
user@switch# flexible-vlan-tagging
```

3. Enable extended VLAN bridge encapsulation on the interface:

```
[edit interfaces interface-name]
user@switch# encapsulation extended-vlan-bridge
```

4. Enable the S-VLAN interface to send and receive untagged packets:

```
[edit interfaces interface-name]
user@switch# native-vlan-id vlan-id
```

5. Bind the logical interface (unit) of the interface that you specified in step 1 to the automatically-created VLAN ID for the S-VLAN:

```
[edit interfaces interface-name unit logical-unit-number]
user@switch# vlan-id number
```

For example, the following configuration makes xe-0/0/0.10 a member of VLAN 10, enables Q-in-Q tunneling on interface xe-0/0/0, enables xe-0/0/0 to accept untagged packets, and binds the VLAN ID of S-VLAN v10 to a logical interface of xe-0/0/0.

```
set vlans v10 interface xe-0/0/0.10
set interfaces xe-0/0/0 flexible-vlan-tagging
set interfaces xe-0/0/0 native-vlan-id 10
set interfaces xe-0/0/0 encapsulation extended-vlan-bridge
set interfaces xe-0/0/0 unit 10 vlan-id 10
```

Now configure all-in-one bundling on a C-VLAN interface:

1. Assign a logical interface (unit) of the C-VLAN interface to be a member of the S-VLAN.

```
[edit vlans vlan-name]
user@switch# interface interface-name.unit-number
```

2. Enable the interface to transmit packets with 802.1Q VLAN tags :

```
[edit interfaces interface-name]
user@switch# flexible-vlan-tagging
```

3. Enable extended VLAN bridge encapsulation on the interface:

```
[edit interfaces interface-name]
user@switch# encapsulation extended-vlan-bridge
```

4. Enable the C-VLAN interface to send and receive untagged packets:

```
[edit interfaces interface-name]
user@switch# native-vlan-id vlan-id
```

5. Configure a logical interface to receive and forward any tagged packet whose VLAN ID tag matches the list of VLAN IDs you specify:

```
[edit interfaces interface-name unit logical-unit-number]
user@switch# vlan-id-list vlan-id-numbers
```



**WARNING:** On some EX and QFX Series switches, you can apply no more than eight VLAN identifier lists to a physical interface.

6. Configure the system to add an S-VLAN tag (outer tag) as packets travel from a C-VLAN interface to the S-VLAN:

```
[edit interfaces interface-name unit logical-unit-number]
user@switch# input-vlan-map push
```

7. Configure the system to remove the S-VLAN tag when packets are forwarded (internally) from the S-VLAN interface to the C-VLAN interface:

```
[edit interfaces interface-name unit logical-unit-number]
user@switch# output-vlan-map pop
```

For example, the following configuration makes xe-0/0/1.10 a member of S-VLAN v10, enables Q-in-Q tunneling, maps packets from C-VLANs 100 through 200 to S-VLAN 10, and enables xe-0/0/1 to accept untagged packets. If a packet originates in C-VLAN 100 and needs to be sent across the S-VLAN, a tag with VLAN ID 10 is added to the packet. When a packet is forwarded (internally) from the S-VLAN interface to interface xe-0/0/1, the tag with VLAN ID 10 is removed.

```
set vlans v10 interface xe-0/0/1.10
set interfaces xe-0/0/1 flexible-vlan-tagging
set interfaces xe-0/0/1 encapsulation extended-vlan-bridge
```

```

set interfaces xe-0/0/1 unit 10 vlan-id-list 100-200
set interfaces xe-0/0/1 native-vlan-id 150
set interfaces xe-0/0/1 unit 10 input-vlan-map push
set interfaces xe-0/0/1 unit 10 output-vlan-map pop

```

#### Related Documentation

- [Understanding Q-in-Q Tunneling on page 2221](#)
- [Configuring Many-to-Many Bundling on page 2228](#)
- [Configuring a Specific Interface Mapping with VLAN ID Translation Option on page 2231](#)

## Configuring Many-to-Many Bundling

You can configure Q-in-Q tunneling using the many-to-many bundling method, which maps packets from multiple C-VLANs to multiple S-VLANs. This method is convenient for mapping a range of C-VLANs without having to specify each one individually. (You can also use this method to configure only one C-VLAN to be mapped to an S-VLAN.)

First configure the S-VLANs and assign them to an interface:

1. Assign a logical interface (unit) to be a member of one of the S-VLANs. Do not use logical interface unit 0.

[edit vlans *vlan-name*]

user@switch# **interface** *interface-name.unit-number*



**NOTE:** Note that you do not create a VLAN ID for the S-VLAN. The ID is created automatically for the appropriate logical interface.

2. Repeat step 1 for the other S-VLANs.
3. Enable the physical interface to transmit packets with two 802.1Q VLAN tags:

[edit interfaces *interface-name*]

user@switch# **flexible-vlan-tagging**

4. Enable extended VLAN bridge encapsulation on the interface:

[edit interfaces *interface-name*]

user@switch# **encapsulation extended-vlan-bridge**

5. Enable the S-VLAN interface to send and receive untagged packets:

[edit interfaces *interface-name*]

user@switch# **native-vlan-id** *vlan-id*

6. Bind one of the logical units of the interface to the VLAN ID for one of the S-VLANs.

[edit interfaces *interface-name* unit *logical-unit-number*]

user@switch# **vlan-id** *number*

7. Repeat step 6 to bind the automatically-created VLAN IDs for the other S-VLANs to the other logical units of the interface:

For example, the following configuration creates S-VLANs v10 and v30 and associates them with interface xe-0/0/0. It also enables Q-in-Q tunneling, enables xe-0/0/0 to accept untagged packets, and maps incoming C-VLAN packets to S-VLANs v10 and v30.

```

set vlans v10 interface xe-0/0/0.10

```

```

set vlans v30 interface xe-0/0/0.30
set interfaces xe-0/0/0 flexible-vlan-tagging
set interfaces xe-0/0/0 native-vlan-id 10
set interfaces xe-0/0/0 encapsulation extended-vlan-bridge
set interfaces xe-0/0/0 unit 10 vlan-id 10
set interfaces xe-0/0/0 unit 30 vlan-id 30

```

To configure the many-to-many bundling method on a C-VLAN interface, perform the following steps for each customer:

1. Assign a logical interface (unit) of one C-VLAN interface to be a member of one S-VLAN.

**[edit vlans *vlan-name*]**

```
user@switch# interface interface-name.unit-number
```

2. Repeat step 1 to assign another C-VLAN interface (physical interface) to be a member of another S-VLAN.

3. Enable the interface to transmit packets with 802.1Q VLAN tags:

**[edit interfaces *interface-name*]**

```
user@switch# flexible-vlan-tagging
```

4. Enable extended VLAN bridge encapsulation on the interface:

**[edit interfaces *interface-name*]**

```
user@switch# encapsulation extended-vlan-bridge
```

5. Enable the C-VLAN interface to send and receive untagged packets:

**[edit interfaces *interface-name*]**

```
user@switch# native-vlan-id vlan-id
```

6. For each physical interface, configure a logical interface (unit) to receive and forward any tagged packet whose VLAN ID tag matches the list of VLAN IDs you specify:

**[edit interfaces *interface-name* unit *logical-unit-number*]**

```
user@switch# vlan-id-list vlan-id-numbers
```

To configure only one C-VLAN to be mapped to an S-VLAN, specify only one VLAN ID after *vlan-id-list*.



**WARNING:** On some EX and QFX Series switches, you can apply no more than eight VLAN identifier list to a physical interface.

7. For each physical interface, configure the system to add an S-VLAN tag (outer tag) as packets travel from the C-VLAN interface to the S-VLAN:

**[edit interfaces *interface-name* unit *logical-unit-number*]**

```
user@switch# input-vlan-map push
```

8. For each physical interface, configure the system to remove the S-VLAN tag when packets are forwarded from the S-VLAN interface to the C-VLAN interface:

**[edit interfaces *interface-name* unit *logical-unit-number*]**

```
user@switch# output-vlan-map pop
```

For example, the following configuration makes xe-0/0/1.10 a member of S-VLAN v10, enables Q-in-Q tunneling, and maps packets from C-VLANs 10 through 20 to S-VLAN 10. The configuration for customer 2 makes xe-0/0/2.30 a member of S-VLAN v30, enables Q-in-Q tunneling, and maps packets from C-VLANs 30 through 40, 50 through

60, and 70 through 80 to S-VLAN 30. Both interfaces are configured to accept untagged packets.

If a packet originates in C-VLAN 10 and needs to be sent over the S-VLAN, a tag with a VLAN ID 10 is added to the packet. If a packet is forwarded internally from the S-VLAN interface to xe-0/0/1.10, the tag with VLAN ID 10 is removed. The same principles apply to the C-VLANs configured on interface xe-0/0/2.



**NOTE:** Notice that you can use the same tag value for an S-VLAN and C-VLAN. For example, the configuration for customer 1 maps C-VLAN ID 10 to S-VLAN ID 10. C-VLAN and S-VLAN tags use separate name spaces, so this configuration is allowed.

Configuration for customer 1:

```
set vlans v10 interface xe-0/0/1.10
set interfaces xe-0/0/1 flexible-vlan-tagging
set interfaces xe-0/0/1 encapsulation extended-vlan-bridge
set interfaces xe-0/0/1 unit 10 vlan-id-list 10-20
set interfaces xe-0/0/1 native-vlan-id 15
set interfaces xe-0/0/1 unit 10 input-vlan-map push
set interfaces xe-0/0/1 unit 10 output-vlan-map pop
```

Configuration for customer 2:

```
set vlans v30 interface xe-0/0/2.30
set interfaces xe-0/0/2 flexible-vlan-tagging
set interfaces xe-0/0/2 encapsulation extended-vlan-bridge
set interfaces xe-0/0/2 unit 30 vlan-id-list 30-40
set interfaces xe-0/0/2 unit 30 vlan-id-list 50-60
set interfaces xe-0/0/2 unit 30 vlan-id-list 70-80
set interfaces xe-0/0/2 native-vlan-id 75
set interfaces xe-0/0/2 unit 30 input-vlan-map push
set interfaces xe-0/0/2 unit 30 output-vlan-map pop
```

#### Related Documentation

- [Understanding Q-in-Q Tunneling on page 2221](#)
- [Configuring All-in-One Bundling on page 2227](#)
- [Configuring a Specific Interface Mapping with VLAN ID Translation Option on page 2231](#)

## Configuring a Specific Interface Mapping with VLAN ID Translation Option

You can configure Q-in-Q tunneling by mapping packets from a specified C-VLAN to a specified S-VLAN. In addition, you can configure the system to replace a C-VLAN tag with an S-VLAN tag or replace an S-VLAN tag with a C-VLAN tag (instead of double tagging). This is called VLAN translation or VLAN rewriting. VLAN translation is particularly useful if a service provider's Layer 2 network that connects a customer's sites does not support double tagged packets.

When you use VLAN translation, both ends of the link normally must be able to swap the tags appropriately. That is, both ends of the link must be configured to swap the

C-VLAN tag for the S-VLAN tag and swap the S-VLAN tag for the C-VLAN tag so that traffic in both directions is tagged appropriately while in transit and after arrival.

First configure the S-VLAN and its interface:

1. Assign a logical interface to be a member of the S-VLAN. Do not use unit 0.

```
[edit vlans vlan-name]
user@switch# interface interface-name.unit-number
```



**NOTE:** Note that you do not create a VLAN ID for the S-VLAN. The ID is created automatically for the appropriate logical interface.

2. Enable the interface to transmit packets with 802.1Q VLAN tags:

```
[edit interfaces interface-name]
user@switch# flexible-vlan-tagging
```

3. Enable the S-VLAN interface to send and receive untagged packets:

```
[edit interfaces interface-name]
user@switch# native-vlan-id vlan-id
```

4. Enable extended VLAN bridge encapsulation on the interface:

```
[edit interfaces interface-name]
user@switch# encapsulation extended-vlan-bridge
```

5. Bind the logical interface (unit) of the interface that you specified earlier to the VLAN ID for the S-VLAN:

```
[edit interfaces interface-name unit logical-unit-number]
user@switch# vlan-id number
```

For example, the following configuration creates S-VLAN v200, makes xe-0/0/0.200 a member of that VLAN, enables Q-in-Q tunneling on interface xe-0/0/0, enables xe-0/0/0 to accept untagged packets, and binds a logical interface of xe-0/0/0 to the VLAN ID of VLAN v200.

```
set vlans v200 interface xe-0/0/0.200
set interfaces xe-0/0/0 flexible-vlan-tagging
set interfaces xe-0/0/0 native-vlan-id 10
set interfaces xe-0/0/0 encapsulation extended-vlan-bridge
set interfaces xe-0/0/0 unit 200 vlan-id 200
```

Now configure a specific interface mapping with optional VLAN ID translation on the C-VLAN interface:

1. Assign a logical interface of the C-VLAN interface to be a member of the S-VLAN.

```
[edit vlans vlan-name]
user@switch# interface interface-name.unit-number
```

2. Enable the interface to transmit packets with 802.1Q VLAN tags:

```
[edit interfaces interface-name]
user@switch# flexible-vlan-tagging
```

3. Enable the C-VLAN interface to send and receive untagged packets:

```
[edit interfaces interface-name]
user@switch# native-vlan-id vlan-id
```

4. Enable extended VLAN bridge encapsulation on the interface:

```
[edit interfaces interface-name]
```

```
user@switch# encapsulation extended-vlan-bridge
```

5. Configure a logical interface (unit) to receive and forward any tagged packet whose VLAN ID tag matches the VLAN IDs you specify:

```
[edit interfaces interface-name unit logical-unit-number]
```

```
user@switch# vlan-id number
```

6. Configure the system to remove the existing C-VLAN tag and replace it with the S-VLAN tag when packets ingress on the C-VLAN interface and are forwarded to the S-VLAN:

```
[edit interfaces interface-name unit logical-unit-number]
```

```
user@switch# input-vlan-map swap
```

7. Configure the system to remove the existing S-VLAN tag and replace it with the C-VLAN tag when packets are forwarded from the S-VLAN interface to the C-VLAN interface:

```
[edit interfaces interface-name unit logical-unit-number]
```

```
user@switch# output-vlan-map swap
```

8. To configure an S-VLAN and associate it with the appropriate C-VLAN interface:

```
[edit vlans vlan-name]
```

```
user@switch# interface interface-name
```

For example, the following configuration on C-VLAN interface xe-0/0/1 enables Q-in-Q tunneling, enables xe-0/0/1 to accept untagged packets, and maps incoming packets from C-VLAN 150 to logical interface 200, which is a member of S-VLAN 200. Also, when packets egress from C-VLAN interface xe-0/0/1 and travel to the S-VLAN interface, the C-VLAN tag of 150 is removed and replaced with the S-VLAN tag of 200. When packets travel from the S-VLAN interface to the C-VLAN interface, the S-VLAN tag of 200 is removed and replaced with the C-VLAN tag of 150.

```
set vlans v200 interface xe-0/0/1.200
```

```
set interfaces xe-0/0/1 flexible-vlan-tagging
```

```
set interfaces xe-0/0/1 native-vlan-id 10
```

```
set interfaces xe-0/0/1 encapsulation extended-vlan-bridge
```

```
set interfaces xe-0/0/1 unit 200 vlan-id 150
```

```
set interfaces xe-0/0/1 unit 200 output-vlan-map swap
```

```
set interfaces xe-0/0/1 unit 200 input-vlan-map swap
```

#### Related Documentation

- [Understanding Q-in-Q Tunneling on page 2221](#)
- [Configuring All-in-One Bundling on page 2227](#)
- [Configuring Many-to-Many Bundling on page 2228](#)





## PART 33

# Proxy ARP

- [Using Proxy ARP on page 2243](#)



## CHAPTER 77

# Using Proxy ARP

- [Understanding Proxy ARP on page 2243](#)
- [Configuring Proxy ARP \(CLI Procedure\) on page 2245](#)
- [Verifying That Proxy ARP Is Working Correctly on page 2245](#)

## Understanding Proxy ARP

---

You can configure proxy Address Resolution Protocol (ARP) to enable the switch to respond to ARP queries for network addresses by offering its own Ethernet media access control (MAC) address. With proxy ARP enabled, the switch captures and routes traffic to the intended destination.

Proxy ARP is useful in situations where hosts are on different physical networks and you do not want to use subnet masking. Because ARP broadcasts are not propagated between hosts on different physical networks, hosts will not receive a response to their ARP request if the destination is on a different subnet. Enabling the switch to act as an ARP proxy allows the hosts to transparently communicate with each other through the switch. Proxy ARP can help hosts on a subnet reach remote subnets without your having to configure routing or a default gateway.

- [What Is ARP? on page 2243](#)
- [Proxy ARP Overview on page 2243](#)
- [Best Practices for Proxy ARP on page 2244](#)

## What Is ARP?

Ethernet LANs use ARP to map Ethernet MAC addresses to IP addresses. Each device maintains a cache containing a mapping of MAC addresses to IP addresses. The switch maintains this mapping in a cache that it consults when forwarding packets to network devices. If the ARP cache does not contain an entry for the destination device, the host (the DHCP client) broadcasts an ARP request for that device's address and stores the response in the cache.

## Proxy ARP Overview

When proxy ARP is enabled, if the switch receives an ARP request for which it has a route to the target (destination) IP address, the switch responds by sending a proxy ARP reply

packet containing its own MAC address. The host that sent the ARP request then sends its packets to the switch, which forwards them to the intended host.



**NOTE:** For security reasons, the source address in an ARP request must be on the same subnet as the interface on which the ARP request is received.

You can configure proxy ARP for each interface. You can also configure proxy ARP for a VLAN by using a routed VLAN interface (RVI).

Two modes of proxy ARP are supported: restricted and unrestricted. Both modes require that the switch have an active route to the destination address of the ARP request.

- **Restricted**—The switch responds to ARP requests in which the physical networks of the source and target are different and does not respond if the source and target IP addresses are on the same subnet. In this mode, hosts on the same subnet communicate without proxy ARP. We recommend that you use this mode on the switch.
- **Unrestricted**—The switch responds to all ARP requests for which it has a route to the destination. This is the default mode (because it is the default mode in Juniper Networks Junos operating system (Junos OS) configurations other than those on the switch). We recommend using restricted mode on the switch.

## Best Practices for Proxy ARP

We recommend these best practices for configuring proxy ARP on the switches:

- Set proxy ARP to restricted mode.
- Use restricted mode when configuring proxy ARP on RVIs.
- If you set proxy ARP to unrestricted, disable gratuitous ARP requests on each interface enabled for proxy ARP.

- Related Documentation**
- *Configuring Proxy ARP*
  - *proxy-arp*

## Configuring Proxy ARP (CLI Procedure)



**NOTE:** This task uses Junos OS for EX Series switches and QFX3500 and QFX3600 switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring Proxy ARP (CLI Procedure)* or *Configuring Proxy ARP*. For ELS details, see “Getting Started with Enhanced Layer 2 Software” on page 41.

You can configure proxy Address Resolution Protocol (ARP) on your switch to enable the switch to respond to ARP queries for network addresses by offering its own media access control (MAC) address. With proxy ARP enabled, the switch captures and routes traffic to the intended destination.

To configure proxy ARP on a single interface:

```
[edit interfaces]
user@switch# set interface-name unit logical-unit-number proxy-arp (restricted |
unrestricted)
```



**BEST PRACTICE:** We recommend that you configure proxy ARP in restricted mode. In restricted mode, the switch does not act as a proxy if the source and target IP addresses are on the same subnet. If you decide to use unrestricted mode, disable gratuitous ARP requests on the interface to avoid a situation wherein the switch's response to a gratuitous ARP request appears to the host to be an indication of an IP conflict.

To configure proxy ARP on an integrated routing and bridging (IRB) interface:

```
[edit interfaces]
user@switch# set irb.logical-unit-number proxy-arp restricted
```

- Related Documentation**
- *Example: Configuring Proxy ARP on an EX Series Switch*
  - [Verifying That Proxy ARP Is Working Correctly on page 2245](#)
  - *Configuring Integrated Routing and Bridging Interfaces (CLI Procedure)*

## Verifying That Proxy ARP Is Working Correctly

**Purpose** Verify that the switch is sending proxy ARP messages.

**Action** List the system statistics for ARP:

```
user@switch> show system statistics arp
arp:
 90060 datagrams received
 34 ARP requests received
 610 ARP replies received
 2 resolution request received
 0 unrestricted proxy requests
 0 restricted proxy requests
```

```
0 received proxy requests
0 unrestricted proxy requests not proxied
0 restricted proxy requests not proxied
0 datagrams with bogus interface
0 datagrams with incorrect length
0 datagrams for non-IP protocol
0 datagrams with unsupported op code
0 datagrams with bad protocol address length
0 datagrams with bad hardware address length
0 datagrams with multicast source address
0 datagrams with multicast target address
0 datagrams with my own hardware address
0 datagrams for an address not on the interface
0 datagrams with a broadcast source address
294 datagrams with source address duplicate to mine
89113 datagrams which were not for me
0 packets discarded waiting for resolution
0 packets sent after waiting for resolution
309 ARP requests sent
35 ARP replies sent
0 requests for memory denied
0 requests dropped on entry
0 requests dropped during retry
0 requests dropped due to interface deletion
0 requests on unnumbered interfaces
0 new requests on unnumbered interfaces
0 replies for from unnumbered interfaces
0 requests on unnumbered interface with non-subnetted donor
0 replies from unnumbered interface with non-subnetted donor
```

**Meaning** The statistics show that two proxy ARP requests were received. The **unrestricted proxy requests not proxied** and **restricted proxy requests not proxied** fields indicate that all the unproxied ARP requests received have been proxied by the switch.

**Related Documentation**

- *Configuring Proxy ARP*
- [Configuring Proxy ARP \(CLI Procedure\) on page 2245](#)

## PART 34

# Reflective Relay

- [Using Reflective Relay on page 2249](#)





## CHAPTER 78

# Using Reflective Relay

- [Understanding Reflective Relay for Use with VEPA Technology on page 2249](#)
- [Configuring Reflective Relay on page 2250](#)
- [Example: Configuring Reflective Relay for Use with VEPA Technology on page 2251](#)

### Understanding Reflective Relay for Use with VEPA Technology

---

Virtual Ethernet Port Aggregator (VEPA) technology aggregates packets generated by virtual machines located on the same server and relays them to a physical switch. The physical switch then provides connectivity between the virtual machines located on the server, so the virtual machines do not communicate with one another. Offloading switching activities from a virtual switch to a physical switch reduces the computing overhead on the virtual servers and takes advantage of the security, filtering, and management features of the physical switch. Reflective relay, also known as “hairpin turn,” enables the physical switch to receive aggregated packets from the virtual machines hosted on the server through the VEPA on the downstream port and send those packets out the same downstream port from which the physical switch received them.

- [VEPA on page 2249](#)
- [Reflective Relay on page 2249](#)

### VEPA

Even though virtual machines are capable of sending packets directly to one another, it is more efficient to pass these aggregated packets from the VEPA to a physical switch. The switch can then send any packets destined for a virtual machine located on the same server to the VEPA.

### Reflective Relay

Reflective relay, also known as a “hairpin turn” or “hairpin mode,” returns aggregated packets to the VEPA by using the same downstream port that initially delivered the aggregated packets from the VEPA to the switch. Reflective relay must be configured on the interface located on the physical switch that receives aggregated packets, such as VEPA packets, because some of these packets might need to be sent back to the server if they are destined for another virtual machine on the same server.

Reflective relay only occurs in two situations:

- When the destination address of the packet was learned on that downstream port
- When the destination has not yet been learned

Reflective relay does not otherwise change the operation of the switch. If the interface to which the virtual machine is connected and the MAC address of the virtual machine packet are not yet included in the Ethernet switching table for the virtual machine's associated VLAN, an entry is added. If the source MAC address of an incoming packet under the respective VLAN is not yet present in the Ethernet switching table, the switch floods the packet on all the other ports that are members of the same VLAN, including the port on which the packet arrived.

**Related  
Documentation**

- [Understanding Bridging and VLANs](#)
- [Understanding Bridging and VLANs on page 2089](#)

---

## Configuring Reflective Relay

Configure reflective relay when a switch port must return packets on a downstream port. For example, configure reflective relay when a switch port receives aggregated virtual machine packets from a technology such as virtual Ethernet port aggregator (VEPA). When these packets are passed through the switch, reflective relay allows the switch to send those packets back on the same interface that was used for delivery.



**NOTE:** This task uses Junos OS for QFX3500 and QFX3600 switches that supports the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring Reflective Relay*.

Before you begin configuring reflective relay, ensure that you have:

- Configured packet aggregation on the server connected to the port. See your server documentation.
- Configured the port for all VLANs that could be included in aggregated packets..

To configure reflective relay:

1. Configure an Ethernet interface with an interface mode of **trunk**:

```
[edit]
user@switch# set interfaces interface-name unit number family family-type interface-mode
trunk
```

For example:

```
[edit]
user@switch# set interfaces xe-0/0/2 unit 0 family ethernet-switching interface-mode trunk
```

2. Configure the interface for reflective relay:

```
[edit]
user@switch# set interfaces interface-name unit number family family-type reflective-relay
```

For example:

```
[edit]
user@switch# set interfaces xe-0/0/2 unit 0 family ethernet-switching reflective-relay
```

3. Configure the interface for the VLANs that exist on the VM server:

```
[edit]
user@switch# set interfaces interface-name unit number family family-type vlan members
vlan-names
```

For example:

```
[edit]
user@switch# set interfaces xe-0/0/2 unit 0 family ethernet-switching vlan members
[VLAN_Purple VLAN_Orange VLAN_Blue]
```

#### Related Documentation

- [Example: Configuring Reflective Relay for Use with VEPA Technology on page 2251](#)
- [Understanding Reflective Relay for Use with VEPA Technology on page 2249](#)

### Example: Configuring Reflective Relay for Use with VEPA Technology

Reflective relay must be configured on a switch that receives virtual machine aggregated packets, such as Virtual Ethernet Port Aggregator (VEPA) packets, because some of these packets might be sent back to the server destined for another virtual machine on the same server. Reflective relay returns those packets to the original device using the same downstream port that delivered the packets to the switch.



**NOTE:** This example uses Junos OS for QFX3500 and QFX3600 switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Example: Configuring Reflective Relay for Use with VEPA Technology*. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 41](#).

This example shows how to configure a switch port interface to return packets sent by VEPA on the downstream interface back to the server using the same downstream interface:

- [Requirements on page 2252](#)
- [Overview and Topology on page 2252](#)
- [Configuration on page 2254](#)
- [Verification on page 2254](#)

## Requirements

This example uses the following hardware and software components:

- One QFX3500 switch
- One server
- Junos OS Release 12.1 or later for the QFX Series


Before you configure reflective relay on a switch port, be sure you have:

- Configured a server with six virtual machines, VM 1 through VM 6.
- Configured the server with three VLANs named VLAN\_Purple, VLAN\_Orange, and VLAN\_Blue and added two virtual machines to each VLAN.
- Configured the same three VLANs named VLAN\_Purple, VLAN\_Orange, and VLAN\_Blue on one interface.
- Installed and configured VEPA to aggregate the virtual machine packets.

## Overview and Topology

In this example, illustrated in [Figure 29](#), a switch is connected to one server that is hosting six virtual machines and is configured with a VEPA for aggregating packets. The server's six virtual machines are VM 1 through VM 6, and each virtual machine belongs to one of the three server VLANs, VLAN\_Purple, VLAN\_Orange, or VLAN\_Blue. Instead of the server directly passing packets between virtual machines, packets from any of the three VLANs that are destined for another one of the three VLANs are aggregated using VEPA technology and passed to the switch for processing. You must configure the switch port to accept these aggregated packets on the downstream interface and to return appropriate packets to the server on the same downstream interface after they are processed. [Figure 29](#) shows the topology for this example.

Figure 29: Reflective Relay Topology



g020996

In this example, you configure the physical Ethernet switch port interface for trunk interface mode and reflective relay. Configuring trunk port mode allows the interface to accept VLAN tagged packets. Configuring reflective relay allows the downstream port to return those packets on the same interface. [Table 194](#) shows the components used in this example.

Table 194: Components of the Topology for Configuring Reflective Relay

| Component        | Description                                                                                                                                             |
|------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
| QFX3500 switch   | Switch that supports reflective relay. For a list of switches that support this feature, see <i>QFX Series Software Features Overview</i> .             |
| <b>xe-0/0/2</b>  | Switch interface to the server.                                                                                                                         |
| Server           | Server with virtual machines and VEPA technology.                                                                                                       |
| Virtual machines | Six virtual machines located on the server: V1, V2, V3, V4, V5, and V6.                                                                                 |
| VLANs            | Three VLANs: VLAN_Purple, VLAN_Orange, and VLAN_Blue. Each VLAN has two virtual machine members.                                                        |
| VEPA             | Virtual Ethernet port aggregator that aggregates virtual machine packets on the server before the resulting single stream is transmitted to the switch. |

## Configuration

To configure reflective relay, perform these tasks:

- [Configuring Reflective Relay on the Port on page 2254](#)

---

### Configuring Reflective Relay on the Port

---

#### CLI Quick Configuration

To quickly configure reflective relay, copy the following commands and paste them into the switch window:

```
[edit]
set interfaces xe-0/0/2 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/2 unit 0 family ethernet-switching reflective-relay
set interfaces xe-0/0/2 unit 0 family ethernet-switching vlan members [VLAN_Blue VLAN_Orange
VLAN_Purple]
```

#### Step-by-Step Procedure

To configure reflective relay:

1. Configure the trunk interface mode on the interface:

```
[edit]
user@switch# set interfaces xe-0/0/2 unit 0 family ethernet-switching interface-mode
trunk
```

2. Configure reflective relay on the interface to allow it to both accept and send packets:

```
[edit]
user@switch# set interfaces xe-0/0/2 unit 0 family ethernet-switching reflective-relay
```

3. Configure the interface for the three VLANs on the server:

```
[edit]
user@switch# set interfaces xe-0/0/2 unit 0 family ethernet-switching vlan members
[VLAN_Purple VLAN_Orange VLAN_Blue]
```

#### Results

Check the results of the configuration:

```
[edit interfaces xe-0/0/2]
user@switch# show
unit 0 {
 family ethernet-switching {
 interface-mode trunk;
 reflective-relay;
 vlan {
 members [VLAN_Purple VLAN_Orange VLAN_Blue];
 }
 }
}
```

## Verification

To confirm that reflective relay is enabled and working correctly, perform these tasks:

- [Verifying That Reflective Relay Is Enabled and Working Correctly on page 2254](#)

---

### Verifying That Reflective Relay Is Enabled and Working Correctly

---

#### Purpose

Verify that reflective relay is enabled and working correctly.

**Action** Use the `show ethernet-switching interfaces detail` command to display the reflective relay status:

```
user@switch> show ethernet-switching interfaces xe-0/0/2 detail
Interface: xe-0/0/2, Index: 66, State: down, Interface mode: Trunk
Reflective Relay Status: Enabled
Ether type for the interface: 0x8100
VLAN membership:
 VLAN_Purple, 802.1Q Tag: 450, tagged, unblocked
 VLAN_Orange, 802.1Q Tag: 460, tagged, unblocked
 VLAN_Blue, 802.1Q Tag: 470, tagged, unblocked
Number of MACs learned on IFL: 0
```

Confirm that reflective relay is working by sending a Layer 2 broadcast message from one virtual machine to another virtual machine located on the same VLAN. Check the switch to verify that the switch sends the packets back on the same interface on which they were received. One way to check this is to set up port mirroring on the switch interface, connect a traffic generator to the mirrored interface, and use the traffic generator to examine packets.

Alternatively, if you do not have a traffic generator available, you can send traffic between two virtual machines with FTP, Telnet, or SSH, while running the `tcpdump` utility on the receiver virtual machine port to capture reflected packets.

**Meaning** The reflective relay status is **Enabled**, meaning that interface `xe-0/0/2` is configured for the trunk interface mode, which accepts VLAN-tagged packets, and for reflective relay, which accepts and returns packets on the same interface.

When the traffic generator shows packets arriving at the switch and returning to the server on the same interface, reflective relay is working.

**Related Documentation**

- [Understanding Reflective Relay for Use with VEPA Technology on page 2249](#)
- [Configuring Port Mirroring](#)
- [interface-mode on page 2290](#)
- [reflective-relay on page 2377](#)





## PART 35

# Configuration Statements and Operational Commands

- [Ethernet Ring Protection Configuration Statements on page 2259](#)
- [VLAN Configuration Statements on page 2271](#)
- [MAC Address Configuration Statements on page 2315](#)
- [STP Configuration Statements on page 2323](#)
- [Q-in-Q Configuration Statements on page 2367](#)
- [Reflective Relay Configuration Statement on page 2377](#)
- [Bridging and VLANs Monitoring Commands on page 2379](#)
- [MAC Address Operational Commands on page 2405](#)
- [Spanning Tree Monitoring Commands on page 2417](#)



## CHAPTER 79

# Ethernet Ring Protection Configuration Statements

- [control-channel on page 2260](#)
- [control-vlan on page 2261](#)
- [data-channel on page 2262](#)
- [east-interface on page 2263](#)
- [ethernet-ring on page 2264](#)
- [guard-interval on page 2265](#)
- [hold-interval \(Protection Group\) on page 2266](#)
- [protection-group on page 2267](#)
- [restore-interval on page 2268](#)
- [ring-protection-link-end on page 2269](#)
- [ring-protection-link-owner on page 2269](#)
- [west-interface on page 2270](#)

## control-channel

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>control-channel <i>channel-name</i> {<br/>    vlan <i>vlan-id</i>;<br/>    interface name <i>interface-name</i><br/>}</code>                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit protocols protection-group <b>ethernet-ring</b> <i>name</i> ( <b>east-interface</b>   <b>west-interface</b> )]                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4.<br>Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Configure the Ethernet RPS control channel logical interface to carry the RAPS PDU. The related physical interface is the physical ring port.                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <b>vlan <i>vlan-id</i></b> —If the control channel logical interface is a trunk port, then a dedicated <b>vlan <i>vlan-id</i></b> defines the dedicated VLAN channel to carry the RAPS traffic. Only configure the <b>vlan-id</b> when the control channel logical interface is the trunk port.<br><br><b>interface name <i>interface-name</i></b> —Interface name of the control channel.                                                                                                  |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Ethernet Ring Protection Switching Overview on page 2055</a></li><li>• <a href="#">Example: Configuring Ethernet Ring Protection Switching on EX Series Switches</a></li><li>• <a href="#">Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS on page 2066</a></li><li>• <a href="#">Configuring Ethernet Ring Protection Switching (CLI Procedure) on page 2063</a></li></ul> |

## control-vlan

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | control-vlan ( <i>vlan-id</i>   <i>vlan-name</i> )                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit protocols protection-group <a href="#">ethernet-ring</a> ]<br>[edit protocols protection-group <a href="#">ethernet-ring</a> name (east-interface  west-interface)]                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4.<br>Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.                                                                                                                                                                                                       |
| <b>Description</b>              | Specify the VLAN that carries the protocol data units (PDUs) between the nodes in the protected Ethernet ring. This is a control VLAN, meaning that it carries data for one instance of an Ethernet ring protection switching (ERPS) in the control channel. Use a control VLAN on trunk port interfaces. One control channel can contain multiple control VLANs.                                            |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Ethernet Ring Protection Switching on EX Series Switches</a></li> <li>• <a href="#">Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS on page 2066</a></li> <li>• <a href="#">Configuring Ethernet Ring Protection Switching (CLI Procedure) on page 2063</a></li> </ul> |

## data-channel

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>data-channel {<br/>    vlan <i>number</i>;<br/>}</code>                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit protocols protection-group <b>ethernet-ring</b> <i>ring-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.2.<br>Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | <p>For Ethernet ring protection, configure a data channel to define a set of VLAN IDs that belong to a ring instance.</p> <p>VLANs specified in the data channel use the same topology used by the ERPS PDU in the control channel. Therefore, if a ring interface is blocked in the control channel, all traffic in the data channel is also blocked on that interface.</p>                                                                                                                  |
| <b>Options</b>                  | <b>vlan <i>number</i></b> —Specify (by VLAN ID) one or more VLANs that belong to a ring instance.                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Ethernet Ring Protection Using Ring Instances for Load Balancing</i></li><li>• <i>Example: Configuring Load Balancing Within Ethernet Ring Protection for MX Series Routers</i></li><li>• <a href="#">Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS on page 2066</a></li><li>• <a href="#">Configuring Ethernet Ring Protection Switching (CLI Procedure) on page 2063</a></li></ul> |

## east-interface

**Syntax**

```
east-interface {
 node-id mac-address;
 control-channel channel-name {
 vlan number;
 interface name interface-name
 }
 interface-none
 ring-protection-link-end;
}
```

**Hierarchy Level** [edit protocols protection-group **ethernet-ring** *ring-name*]

**Release Information** Statement introduced in Junos OS Release 9.4.  
Statement introduced in Junos OS Release 12.1 for EX Series switches.  
Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.

**Description** Define one of the two interface ports for Ethernet ring protection, the other being defined by the **west-interface** statement at the same hierarchy level. The interface must use the control channel's logical interface name. The control channel is a dedicated VLAN channel for the ring port.

EX Series switches do not use the node-id statement--the node ID is automatically configured on the switches using the MAC address.



**NOTE:** Always configure this port first, before configuring the **west-interface** statement.



**NOTE:** The Node ID is not configurable on EX Series switches. The node ID is automatically configured using the MAC address.

The statements are explained separately.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Ethernet Ring Protection Switching Overview on page 2055](#)
- [Ethernet Ring Protection Using Ring Instances for Load Balancing](#)
- [west-interface on page 2270](#)
- [ethernet-ring on page 2264](#)
- [Example: Configuring Ethernet Ring Protection Switching on EX Series Switches](#)

- [Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS on page 2066](#)
- [Configuring Ethernet Ring Protection Switching \(CLI Procedure\) on page 2063](#)

---

## ethernet-ring

---

**Syntax**    ethernet-ring *ring-name* {  
              control-vlan (*vlan-id* | *vlan-name*);  
              data-channel {  
                  vlan *number*  
              }  
              east-interface {  
                  control-channel *channel-name* {  
                      vlan *number*;  
                      interface name *interface-name*  
                  }  
              }  
              guard-interval *number*;  
              node-id *mac-address*;  
              restore-interval *number*;  
              ring-protection-link-owner;  
              west-interface {  
                  control-channel *channel-name* {  
                      vlan *number*;  
                  }  
              }  
          }

**Hierarchy Level**    [edit protocols protection-group]

**Release Information**    Statement introduced in Junos OS Release 9.4.  
                              Statement introduced in Junos OS Release 12.1 for EX Series switches.  
                              Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.

**Description**    For Ethernet PICs on MX Series routers or for EX Series switches, , specify the Ethernet ring in an Ethernet ring protection switching configuration.

**Options**    *ring-name*—Name of the Ethernet protection ring.  
  
              The remaining statements are explained separately.

**Required Privilege Level**    interface—To view this statement in the configuration.  
                                  interface-control—To add this statement to the configuration.

**Related Documentation**

- [Ethernet Ring Protection Switching Overview on page 2055](#)
- [Example: Configuring Ethernet Ring Protection Switching on EX Series Switches](#)
- [Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS on page 2066](#)
- [Configuring Ethernet Ring Protection Switching \(CLI Procedure\) on page 2063](#)



## guard-interval

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>guard-interval <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit protocols protection-group <a href="#">ethernet-ring</a> <i>ring-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4.<br>Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | When a link goes down, the ring protection link (RPL) activates. When the downed link comes back up, the RPL link receives notification, restores the link, and waits for the restore interval before issuing another block on the same link. This configuration is a global configuration and applies to all Ethernet rings if the Ethernet ring does not have a more specific configuration for this value. If no parameter is configured at the protection group level, the global configuration of this parameter uses the default value. |
| <b>Options</b>                  | <i>number</i> —Guard timer interval, in milliseconds.<br><b>Range:</b> 10 through 2000 ms<br><b>Default:</b> 500 ms                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Ethernet Ring Protection Switching Overview on page 2055</a></li> <li>• <a href="#">Example: Configuring Ethernet Ring Protection Switching on EX Series Switches</a></li> <li>• <a href="#">Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS on page 2066</a></li> <li>• <a href="#">Configuring Ethernet Ring Protection Switching (CLI Procedure) on page 2063</a></li> </ul>                                              |

## hold-interval (Protection Group)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | hold-interval <i>number</i> ;                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit protocols protection-group <a href="#">ethernet-ring name</a> ]                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4.                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Specify the hold-off timer interval <i>for all rings</i> in 100 millisecond (ms) increments.                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <p><i>number</i>—Hold-timer interval, in milliseconds.</p> <p><b>Range:</b> 0 through 10,000 ms</p> <p><b>Default:</b> 100 ms</p>                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Ethernet Ring Protection Switching Overview on page 2055</a></li><li>• <a href="#">Example: Configuring Ethernet Ring Protection Switching on EX Series Switches</a></li><li>• <a href="#">Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS on page 2066</a></li></ul> |

## protection-group

```
Syntax protection-group {
 ethernet-ring ring-name {
 control-vlan (vlan-id | vlan-name);
 data-channel {
 vlan number
 }
 east-interface {
 control-channel channel-name {
 vlan number;
 interface name interface-name
 }
 }
 guard-interval number;
 hold-interval number;
 node-id mac-address;
 restore-interval number;
 ring-protection-link-owner RPL owner flag;
 west-interface {
 control-channel channel-name {
 vlan number;
 interface name interface-name
 }
 }
 guard-interval number;
 hold-interval
 node-id mac-address;
 restore-interval number;
 traceoptions {
 file filename <no-stamp> <world-readable | no-world-readable> <replace> <size size>;
 flag flag;
 }
 }
 }
```

**Hierarchy Level** [edit protocols]

**Release Information** Statement introduced in Junos OS Release 12.1 for EX Series switches.

**Description** Configure Ethernet ring protection switching (ERPS).

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Configuring Ethernet Ring Protection Switching on EX Series Switches](#)
- [Configuring Ethernet Ring Protection Switching \(CLI Procedure\) on page 2063](#)
- [Ethernet Ring Protection Switching Overview on page 2055](#)

## restore-interval

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>restore-interval <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit protocols protection-group <b>ethernet-ring</b> <i>ring-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4.<br>Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Configures the number of minutes that the node does not process any Ethernet ring protection (ERP) protocol data units (PDUs).. This configuration is a global configuration and applies to all Ethernet rings if the Ethernet ring does not have a more specific configuration for this value. If no parameter is configured at the protection group level, the global configuration of this parameter uses the default value.                                                             |
| <b>Options</b>                  | <i>number</i> —Specify the restore interval.<br><b>Range:</b> 5 through 12 minutes                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Ethernet Ring Protection Switching Overview on page 2055</a></li><li>• <a href="#">Example: Configuring Ethernet Ring Protection Switching on EX Series Switches</a></li><li>• <a href="#">Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS on page 2066</a></li><li>• <a href="#">Configuring Ethernet Ring Protection Switching (CLI Procedure) on page 2063</a></li></ul> |

## ring-protection-link-end

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | ring-protection-link-end;                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit protocols protection-group <b>ethernet-ring</b> <i>ring-name</i> ( <b>east-interface</b>   <b>west-interface</b> )]                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4.<br>Statement introduced in Junos OS Release 12.1 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Specify that the port is one side of a ring protection link (RPL) by setting the RPL end flag.                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Ethernet Ring Protection Switching Overview on page 2055</a></li> <li>• <a href="#">Example: Configuring Ethernet Ring Protection Switching on EX Series Switches</a></li> <li>• <a href="#">Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS on page 2066</a></li> <li>• <a href="#">Configuring Ethernet Ring Protection Switching (CLI Procedure) on page 2063</a></li> </ul> |

## ring-protection-link-owner

|                                 |                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | ring-protection-link-owner;                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit protocols protection-group <b>ethernet-ring</b> <i>ring-name</i> ]                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4.<br>Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.                                                                                                                                                                                    |
| <b>Description</b>              | Specify the ring protection link (RPL) owner flag in the Ethernet protection ring. Include this statement only once for each ring (only one node can function as the RPL owner).                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Ethernet Ring Protection Switching Overview on page 2055</a></li> <li>• <a href="#">Example: Configuring Ethernet Ring Protection Switching on EX Series Switches</a></li> <li>• <a href="#">Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS on page 2066</a></li> </ul> |

## west-interface

**Syntax**

```
west-interface {
 node-id mac-address;
 control-channel channel-name {
 vlan number;
 interface name interface-name
 }
 interface-name
 ring-protection-link-end;
 virtual-control-channel {
 west-interface name;
 east-interface name;
 }
}
```

**Hierarchy Level** [edit protocols protection-group **ethernet-ring** *ring-name*]

**Release Information** Statement introduced in Junos OS Release 9.5.  
Statement introduced in Junos OS Release 12.1 for EX Series switches.  
Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.

**Description** Define one of the two interface ports for Ethernet ring protection, the other being defined by the **east-interface** statement at the same hierarchy level. The interface must use the control channel's logical interface name. The control channel is a dedicated VLAN channel for the ring port.



**NOTE:** Always configure this port second, after configuring the **east-interface** statement.

The statements are explained separately.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

- Related Documentation**
- [Ethernet Ring Protection Switching Overview on page 2055](#)
  - [Ethernet Ring Protection Using Ring Instances for Load Balancing](#)
  - [east-interface on page 2263](#)
  - [ethernet-ring on page 2264](#)
  - [Example: Configuring Ethernet Ring Protection Switching on EX Series Switches](#)
  - [Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS on page 2066](#)
  - [Configuring Ethernet Ring Protection Switching \(CLI Procedure\) on page 2063](#)

## CHAPTER 80

# VLAN Configuration Statements

- [\[edit vlans\] Configuration Statement Hierarchy on the QFX Series on page 2272](#)
- [description \(VLAN\) on page 2275](#)
- [dhcp-relay on page 2276](#)
- [filter \(VLANs\) on page 2281](#)
- [forwarding-options on page 2282](#)
- [interface \(VLANs\) on page 2287](#)
- [interface-mac-limit on page 2288](#)
- [interface-mode on page 2290](#)
- [irb \(Interfaces\) on page 2292](#)
- [l3-interface \(VLAN\) on page 2295](#)
- [mac \(Static MAC-Based VLANs\) on page 2296](#)
- [members on page 2297](#)
- [native-vlan-id on page 2298](#)
- [packet-action on page 2299](#)
- [port-mode on page 2302](#)
- [service-id on page 2303](#)
- [switch-options on page 2304](#)
- [static \(Static MAC-Based VLANs\) on page 2305](#)
- [static-mac on page 2305](#)
- [vlan-id \(VLANs\) on page 2306](#)
- [vlan-id-list on page 2307](#)
- [vlan-rewrite on page 2308](#)
- [vlan-tagging on page 2309](#)
- [vlan-tags on page 2310](#)
- [vlans on page 2311](#)

## [edit vlans] Configuration Statement Hierarchy on the QFX Series

---

This topic lists supported and unsupported configuration statements in the **[edit vlans]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.

This topic lists:

- [Supported Statements in the \[edit vlans\] Hierarchy Level on page 2272](#)
- [Unsupported Statements in the \[edit vlans\] Hierarchy Level on page 2274](#)

### Supported Statements in the [edit vlans] Hierarchy Level

The following hierarchy shows the **[edit vlans]** configuration statements supported on one or more of the EX Series switches:

```
vlans {
 vlan-name {
 description text-description;
 domain-type bridge;
 forwarding-options {
 dhcp-security {
 arp-inspection;
 group group-name {
 interface interface-name {
 static-ip ip-address {
 mac mac-address;
 }
 }
 }
 overrides {
 no-option82;
 trusted;
 }
 }
 }
 ip-source-guard;
 no-dhcp-snooping;
 option-82 {
 circuit-id {
 prefix {
 host-name;
 logical-system-name;
 routing-instance-name;
 }
 }
 use-interface-description (device | logical);
 use-vlan-id;
 }
 remote-id {
 host-name;
 use-interface-description (device | logical);
 }
 }
}
```



```

 use-string string;
 }
 vendor-id {
 use-string string;
 }
}
filter {
 input filter-name;
 output filter-name;
}
flood {
 input filter-name;
}
}
l3-interface irb.logical-unit-number;
multicast-snooping-options {
 flood-groups [group-names];
 forwarding-cache {
 threshold {
 reuse threshold;
 suppress threshold;
 }
 }
 graceful-restart {
 disable;
 restart-duration duration;
 }
 host-outbound-traffic {
 dot1p bits;
 forwarding-class forwarding-class;
 }
 multichassis-lag-replicate-state;
 nexthop-hold-time time;
 options {
 syslog {
 level level;
 mark interval;
 upto level;
 }
 }
 traceoptions {
 file filename {
 files number;
 no-world-readable;
 size file-size;
 world-readable;
 }
 flag flag {
 disable;
 }
 }
}
switch-options {
 interface interface-name {
 interface-mac-limit limit {

```

```

 packet-action action;
 }
 static-mac mac-address;
}
interface-mac-limit limit {
 packet-action action;
}
mac-move-limit limit {
 packet-action action;
}
mac-table-size limit {
 packet-action drop;
}
no-mac-learning;
}
vlan-id number;
vlan-id-list [vlan-id | vlan-id-vlan-id];
}
}

```

### Unsupported Statements in the [edit vlans] Hierarchy Level

All statements in the **[edit vlans]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

**Table 195: Unsupported [edit vlans] Configuration Statements on EX Series Switches**

| Statement                                                                                         | Hierarchy Level |
|---------------------------------------------------------------------------------------------------|-----------------|
| <b>NOTE:</b> Variables, such as <i>filename</i> , are not shown in the statements or hierarchies. |                 |
| mcae-mac-synchronize                                                                              | [edit vlans]    |
| no-irb-layer-2-copy                                                                               | [edit vlans]    |

**Related Documentation**

- [Understanding Bridging and VLANs on page 2089](#)

## description (VLAN)

---

|                                 |                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>description <i>text-description</i>;</code>                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit vlans <i>vlan-name</i> ]                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                        |
| <b>Description</b>              | Provide a textual description for the VLAN. The text has no effect on the operation of the VLAN or switch.                                                                                                                                                               |
| <b>Options</b>                  | <b><i>text-description</i></b> —Text to describe the interface. It can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. If the text includes spaces, enclose the entire text in quotation marks.                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 2096</a></li><li>• <a href="#">Understanding Bridging and VLANs on page 2089</a></li><li>• <a href="#">show vlans on page 2395</a></li></ul> |

## dhcp-relay

```
Syntax dhcp-relay {
 active-server-group server-group-name;
 authentication {
 password password-string;
 username-include {
 circuit-type;
 delimiter delimiter-character;
 domain-name domain-name-string;
 interface-name;
 logical-system-name;
 mac-address;
 option-60;
 option-82 <circuit-id> <remote-id>;
 routing-instance-name;
 user-prefix user-prefix-string;
 }
 }
 dhcpv6 {
 active-server-group server-group-name;
 authentication {
 password password-string;
 username-include {
 circuit-type;
 client-id;
 delimiter delimiter-character;
 domain-name domain-name-string;
 interface-name;
 logical-system-name;
 relay-agent-interface-id;
 relay-agent-remote-id;
 relay-agent-subscriber-id;
 routing-instance-name;
 user-prefix user-prefix-string;
 }
 }
 dynamic-profile profile-name {
 aggregate-clients (merge | replace);
 use-primary primary-profile-name;
 }
 }
}

group group-name {
 active-server-group server-group-name;
 authentication {
 ...
 }
 dynamic-profile profile-name {
 ...
 }
 interface interface-name {
 exclude;
 liveness-detection {
 failure-action (clear-binding | clear-binding-if-interface-up | log-only);
 }
 }
}
```

```

method {
 bfd {
 version (0 | 1 | automatic);
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 multiplier number;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 detection-time {
 threshold milliseconds;
 }
 session-mode(automatic | multihop | singlehop);
 holddown-interval milliseconds;
 }
}
overrides {
 ...
}
service-profile dynamic-profile-name;
trace;
upto upto-interface-name;
}
service-profile dynamic-profile-name;
}
overrides {
 ...
}
relay-agent-interface-id {
 ...
}
service-profile dynamic-profile-name;
liveness-detection {
 failure-action (clear-binding | clear-binding-if-interface-up | log-only);
 method {
 bfd {
 version (0 | 1 | automatic);
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 multiplier number;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 detection-time {
 threshold milliseconds;
 }
 session-mode(automatic | multihop | singlehop);
 holddown-interval milliseconds;
 }
 }
}
}

```

```
overrides {
 allow-snooped-clients;
 interface-client-limit number;
 no-allow-snooped-clients;
 no-bind-on-request;
 send-release-on-delete;
}
relay-agent-interface-id {
 prefix prefix;
 use-interface-description (logical | device);
}
server-group {
 server-group-name {
 server-ip-address;
 }
}
dynamic-profile profile-name {
 aggregate-clients (merge | replace);
 use-primary primary-profile-name;
}
forward-snooped-clients (all-interfaces | configured-interfaces |
 non-configured-interfaces);
group group-name {
 active-server-group server-group-name;
 authentication {
 ...
 }
}
dynamic-profile profile-name {
 ...
}
interface interface-name {
 exclude;
 liveness-detection {
 failure-action (clear-binding | clear-binding-if-interface-up | log-only);
 method {
 bfd {
 version (0 | 1 | automatic);
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 multiplier number;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 detection-time {
 threshold milliseconds;
 }
 session-mode (automatic | multihop | singlehop);
 holddown-interval milliseconds;
 }
 }
 }
}
overrides {
 ...
}
```

```

 service-profile dynamic-profile-name;
 trace;
 upto upto-interface-name;
 }
 overrides {
 ...
 }
 relay-option-82 {
 ...
 }
 service-profile dynamic-profile-name;
}
liveness-detection {
 failure-action (clear-binding | clear-binding-if-interface-up | log-only);
 method {
 bfd {
 version (0 | 1 | automatic);
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 multiplier number;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 detection-time {
 threshold milliseconds;
 }
 session-mode (automatic | multihop | singlehop);
 holddown-interval milliseconds;
 }
 }
}
overrides {
 allow-snooped-clients;
 always-write-giaddr;
 always-write-option-82;
 client-discover-match <option60-and-option82>;
 disable-relay;
 interface-client-limit number;
 layer2-unicast-replies;
 no-allow-snooped-clients;
 no-bind-on-request;
 proxy-mode;
 replace-ip-source-with;
 send-release-on-delete;
 trust-option-82;
}
relay-option-82 {
 circuit-id {
 prefix prefix;
 use-interface-description (logical | device);
 }
}
server-group {
 server-group-name {

```

```
 server-ip-address;
 }
}
service-profile dynamic-profile-name;
}
```

**Hierarchy Level** [edit forwarding-options],  
[edit vlans forwarding-options]

**Release Information** Statement introduced in Junos OS Release 11.3 for the QFX Series.

**Description** Configure extended Dynamic Host Configuration Protocol (DHCP) relay and DHCPv6 relay options on the switch and enable the switch to function as a DHCP relay agent. A DHCP relay agent forwards DHCP request and reply packets between a DHCP client and a DHCP server.

DHCP relay supports the attachment of dynamic profiles and also interacts with the local AAA Service Framework to use back-end authentication servers, such as RADIUS, to provide subscriber authentication. You can attach dynamic profiles and configure authentication support on a global basis or for a specific group of interfaces.

The extended DHCP and DHCPv6 relay agent options configured with the **dhcp-relay** and **dhcpv6** statements are incompatible with the DHCP/BOOTP relay agent options configured with the **bootp** statement. As a result, the extended DHCP or DHCPv6 relay agent and the DHCP/BOOTP relay agent cannot both be enabled on the router at the same time.

The remaining statements are explained separately.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring DHCP and BOOTP on page 6255](#)



## filter (VLANs)

---

|                                 |                                                                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>filter (input   output) <i>filter-name</i>;</code>                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | <code>[edit vlans <i>vlan-name</i>],</code><br><code>[edit vlans <i>vlan-name</i> forwarding-options]</code>                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.                                                                                                                                                  |
| <b>Description</b>              | Apply a firewall filter to traffic entering or exiting a VLAN.                                                                                                                                                                                                                                     |
| <b>Default</b>                  | All incoming traffic is accepted unmodified to a VLAN, and all outgoing traffic is sent unmodified from a VLAN.                                                                                                                                                                                    |
| <b>Options</b>                  | <p><b><i>filter-name</i></b>—Name of a firewall filter defined at the <code>[edit firewall family <i>family-name</i> filter]</code> hierarchy level.</p> <p><b>input</b>—Apply a firewall filter to VLAN ingress traffic.</p> <p><b>output</b>—Apply a firewall filter to VLAN egress traffic.</p> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Firewall Filters on page 5978</a></li> <li>• <a href="#">Overview of Firewall Filters on page 5951</a></li> </ul>                                                                                                                 |

## forwarding-options

```
Syntax forwarding-options {
 dhcp-relay {
 active-server-group server-group-name;
 authentication {
 password password-string;
 username-include {
 circuit-type;
 delimiter delimiter-character;
 domain-name domain-name-string;
 interface-name;
 logical-system-name;
 mac-address;
 option-60;
 option-82 <circuit-id> <remote-id>;
 routing-instance-name;
 user-prefix user-prefix-string;
 }
 }
 }
 dhcpv6 {
 active-server-group server-group-name;
 authentication {
 password password-string;
 username-include {
 circuit-type;
 client-id;
 delimiter delimiter-character;
 domain-name domain-name-string;
 interface-name;
 logical-system-name;
 relay-agent-interface-id;
 relay-agent-remote-id;
 relay-agent-subscriber-id;
 routing-instance-name;
 user-prefix user-prefix-string;
 }
 }
 }
 dynamic-profile profile-name {
 aggregate-clients (merge | replace);
 use-primary primary-profile-name;
 }
 group group-name {
 active-server-group server-group-name;
 authentication {
 ...
 }
 dynamic-profile profile-name {
 ...
 }
 }
 interface interface-name {
 exclude;
 liveness-detection {
 failure-action (clear-binding | clear-binding-if-interface-up | log-only);
 }
 }
}
```

```

method {
 bfd {
 version (0 | 1 | automatic);
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 multiplier number;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 detection-time {
 threshold milliseconds;
 }
 session-mode(automatic | multihop | singlehop);
 holddown-interval milliseconds;
 }
}
overrides {
 ...
}
service-profile dynamic-profile-name;
trace;
upto upto-interface-name;
}
service-profile dynamic-profile-name;
}
overrides {
 ...
}
relay-agent-interface-id {
 ...
}
service-profile dynamic-profile-name;
liveness-detection {
 failure-action (clear-binding | clear-binding-if-interface-up | log-only);
 method {
 bfd {
 version (0 | 1 | automatic);
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 multiplier number;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 detection-time {
 threshold milliseconds;
 }
 session-mode(automatic | multihop | singlehop);
 holddown-interval milliseconds;
 }
 }
}
}

```

```

overrides {
 allow-snooped-clients;
 interface-client-limit number;
 no-allow-snooped-clients;
 no-bind-on-request;
 send-release-on-delete;
}
relay-agent-interface-id {
 prefix prefix;
 use-interface-description (logical | device);
}
server-group {
 server-group-name {
 server-ip-address;
 }
}
dynamic-profile profile-name {
 aggregate-clients (merge | replace);
 use-primary primary-profile-name;
}
forward-snooped-clients (all-interfaces | configured-interfaces |
 non-configured-interfaces);
group group-name {
 active-server-group server-group-name;
 authentication {
 ...
 }
}
dynamic-profile profile-name {
 ...
}
interface interface-name {
 exclude;
 liveness-detection {
 failure-action (clear-binding | clear-binding-if-interface-up | log-only);
 method {
 bfd {
 version (0 | 1 | automatic);
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 multiplier number;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 detection-time {
 threshold milliseconds;
 }
 }
 session-mode (automatic | multihop | singlehop);
 holddown-interval milliseconds;
 }
 }
}
overrides {
 ...
}

```

```

 service-profile dynamic-profile-name;
 trace;
 upto upto-interface-name;
}
overrides {
 ...
}
relay-option-60 {
 ...
}
relay-option-82 {
 ...
}
service-profile dynamic-profile-name;
}
liveness-detection {
 failure-action (clear-binding | clear-binding-if-interface-up | log-only);
 method {
 bfd {
 version (0 | 1 | automatic);
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 multiplier number;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 detection-time {
 threshold milliseconds;
 }
 session-mode (automatic | multihop | singlehop);
 holddown-interval milliseconds;
 }
 }
}
overrides {
 allow-snooped-clients;
 always-write-giaddr;
 always-write-option-82;
 client-discover-match <option60-and-option82>;
 disable-relay;
 interface-client-limit number;
 layer2-unicast-replies;
 no-allow-snooped-clients;
 no-bind-on-request;
 proxy-mode;
 replace-ip-source-with;
 send-release-on-delete;
 trust-option-82;
}
relay-option-82 {
 circuit-id {
 prefix prefix;
 use-interface-description (logical | device);
 }
}

```

```
 }
 server-group {
 server-group-name {
 server-ip-address;
 }
 }
 service-profile dynamic-profile-name;
}
dhcp-security {
 arp-inspection;
 group group-name {
 interface interface-name {
 static-ip ip-address {
 mac mac-address;
 }
 }
 overrides {
 no-option82;
 trusted;
 untrusted;
 }
 }
}
ip-source-guard;
no-dhcp-snooping;
option-82 {
 circuit-id {
 prefix {
 host-name;
 logical-system-name;
 routing-instance-name;
 }
 use-interface-description (device | logical);
 use-vlan-id;
 }
 remote-id {
 host-name hostname;
 use-interface-description (device | logical);
 use-string string;
 }
 vendor-id {
 use-string string;
 }
}
}
fip-security {
 examine-vn2vf;
 examine-vn2vn {
 beacon-period milliseconds;
 }
 fc-map fc-map-value;
 interface interface-name {
 (fcoe-trusted | no-fcoe-trusted;)
 }
}
}
```

|                                 |                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| <b>Hierarchy Level</b>          | [edit]<br>[edit vlans]                                                                                                      |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 11.3 for QFX Series switches. |
| <b>Description</b>              | Configure traffic forwarding.<br><br>The statements are explained separately.                                               |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.     |

---

## interface (VLANs)

---

|                                 |                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>interface <i>interface-name</i> {<br/>    mapping (native (push   swap)   tag (push   swap));<br/>}</pre>                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit vlans <i>vlan-name</i> ]                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                   |
| <b>Description</b>              | For a specific VLAN, configure an interface.                                                                                                                                                                                        |
| <b>Options</b>                  | <i>interface-name</i> —Name of the interface.<br><br>The remaining statement is explained separately.                                                                                                                               |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 2096</a></li><li>• <i>Configuring VLANs</i></li><li>• <i>Understanding Bridging and VLANs</i></li></ul> |

## interface-mac-limit

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>interface-mac-limit <i>limit</i> {<br/>    <b>packet-action</b> drop;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>     | <pre>[edit bridge-domains <i>bridge-domain-name</i> bridge-options],<br/>[edit bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],<br/>[edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i><br/>    bridge-options],<br/>[edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i><br/>    bridge-options interface <i>interface-name</i>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i><br/>    bridge-domains <i>bridge-domain-name</i> bridge-options],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i><br/>    bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i><br/>    switch-options],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i><br/>    switch-options interface <i>interface-name</i>],<br/>[edit logical-systems <i>logical-system-name</i> switch-options],<br/>[edit logical-systems <i>logical-system-name</i> switch-options interface <i>interface-name</i>],<br/>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i><br/>    bridge-options],<br/>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i><br/>    bridge-options interface <i>interface-name</i>],<br/>[edit routing-instances <i>routing-instance-name</i> switch-options],<br/>[edit routing-instances <i>routing-instance-name</i> switch-options interface <i>interface-name</i>],<br/>[edit switch-options],<br/>[edit switch-options interface <i>interface-name</i>],<br/>[edit switch-options interface <i>interface-name</i>],<br/>[edit vlans <i>vlan-name</i> switch-options],<br/>[edit vlans <i>vlan-name</i> switch-options interface <i>interface-name</i>]</pre> |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 8.4.</p> <p>Support for the <b>switch-options</b> statement added in Junos OS Release 9.2.</p> <p>Support for top-level configuration for the <b>virtual-switch</b> type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p> <p>[edit switch-options], [edit switch-options interface <i>interface-name</i>], [edit vlans <i>vlan-name</i> switch-options], and [edit vlans <i>vlan-name</i> switch-options interface <i>interface-name</i>] hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>         | Configure a limit to the number of MAC addresses that can be learned from a bridge domain, VLAN, virtual switch, or set of bridge domains or VLANs.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |





**NOTE:** For multichassis link aggregation (MC-LAG) peers in active-active mode, configuring the `interface-mac-limit` statement or changing the `interface-mac-limit` configuration when traffic is flowing can cause the MAC entries to be out of synchronization between the two MC-LAG peers, which might result in flooding. To avoid flooding, you must either halt traffic forwarding and then configure the `interface-mac-limit` statement or use the `commit at` configuration statement to commit the changes at the same time in both the peer nodes.

Alternatively, if flooding does occur, you can clear the bridge MAC table on both the routers or switches by using the `clear bridge mac-table` command. Running this command ensures that the MAC entries are re-learned and in synchronization between both the peers.

**Default** For an access port, the default MAC limit is 1024 MAC addresses. For a trunk port, the default MAC limit is 8192 MAC addresses.

**Options** *limit*—Maximum number of MAC addresses learned from an interface.

**Range:** 1 through 524287 MAC addresses per interface

The remaining statement is explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Understanding Layer 2 Learning and Forwarding for Bridge Domains](#)
- [Layer 2 Learning and Forwarding for VLANs Overview on page 2088](#)
- [Understanding Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports](#)
- [Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port](#)

## interface-mode

|                            |                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>interface-mode (access   trunk &lt;inter-switch-link&gt;);</code>                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>     | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family bridge],<br>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family ethernet-switching],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family bridge]                                                      |
| <b>Release Information</b> | Statement introduced in Junos OS Release 9.2.<br>Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for the QFX Series.<br>Statement introduced in Junos OS Release 15.1.<br><b>inter-switch-link</b> option introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers in enhanced LAN mode. |

### Description



**NOTE:** This statement supports the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see [port-mode](#). For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on page 41.

(QFX Series 3500 and 3600 standalone switches)—Determine whether the logical interface accepts or discards packets based on VLAN tags. Specify the **trunk** option to accept packets with a VLAN ID that matches the list of VLAN IDs specified in the **vlan-id** or **vlan-id-list** statement, then forward the packet within the bridge domain or VLAN configured with the matching VLAN ID. Specify the **access** option to accept packets with no VLAN ID, then forward the packet within the bridge domain or VLAN configured with the VLAN ID that matches the VLAN ID specified in the **vlan-id** statement.



**NOTE:** On MX Series routers, if you want IGMP snooping to be functional for a bridge domain, then you should not configure **interface-mode** and **irb** for that bridge. Such a configuration commit succeeds, but IGMP snooping is not functional, and a message informing the same is displayed. For more information, see *Configuring a Trunk Interface on a Bridge Network*.

|                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b> | <p><b>access</b>—Configure a logical interface to accept untagged packets. Specify the VLAN to which this interface belongs using the <b>vlan-id</b> statement.</p> <p><b>trunk</b>—Configure a single logical interface to accept packets tagged with any VLAN ID specified with the <b>vlan-id</b> or <b>vlan-id-list</b> statement.</p> <p><b>trunk inter-switch-link</b>—For a private VLAN, configure the InterSwitch Link protocol (ISL) on a trunk port of the primary VLAN in order to connect the switches composing the PVLAN to each other. You do not need to configure an ISL when a PVLAN is configured</p> |
|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

on a single switch. This configuration specifies whether the particular interface assumes the role of interswitch link for the PVLAN domains of which it is a member. This option is supported only on MX240, MX480, and MX960 routers in enhanced LAN mode.

|                                 |                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Access Mode on a Logical Interface</i></li><li>• <i>Configuring a Logical Interface for Trunk Mode</i></li><li>• <i>Example: Connecting Access Switches to a Distribution Switch</i></li><li>• <i>Tunnel Services Overview</i></li><li>• <i>Configuring Tunnel Interfaces on MX Series Routers</i></li></ul> |

## irb (Interfaces)

---

```
Syntax irb {
 accounting-profile name;
 description text;

 (gratuitous-arp-reply | no-gratuitous-arp-reply);
 hold-time up milliseconds down milliseconds;
 mtu bytes;
 no-gratuitous-arp-request;

 traceoptions {
 flag flag;
 }
 (traps | no-traps);
 unit logical-unit-number {
 accounting-profile name;
 bandwidth rate;
 description text;
 enhanced-convergence;
 disable;
 encapsulation type;
 family inet {
 accounting {
 destination-class-usage;
 source-class-usage {
 input;
 output;
 }
 }
 }
 address ipv4-address {
 arp ip-address (mac | multicast-mac) mac-address <publish>;
 broadcast address;
 preferred;
 primary;
 vrrp-group group-number {
 (accept-data | no-accept-data);
 advertise-interval seconds;
 advertisements-threshold number;
 authentication-key key;
 authentication-type authentication;
 fast-interval milliseconds;
 (preempt | no-preempt) {
 hold-time seconds;
 }
 }
 priority number;
 track {
 interface interface-name {
 bandwidth-threshold bandwidth;
 priority-cost number;
 }
 }
 priority-hold-time seconds;
 route ip-address/mask routing-instance instance-name priority-cost cost;
 }
 }
}
```

```

 virtual-address [addresses];
 vrrp-inherit-from {
 active-group group-number;
 active-interface interface-name;
 }
 }
}
filter {
 input filter-name;
 output filter-name;
}
mtu bytes;
no-neighbor-learn;
no-redirects;
primary;
rpf-check {
 fail-filter filter-name;
 mode {
 loose;
 }
}
targeted-broadcast {
 forward-and-send-to-re;
 forward-only;
}
}
family inet6 {
 accounting {
 destination-class-usage;
 source-class-usage {
 input;
 output;
 }
 }
}
address address {
 eui-64;
 ndp ip-address (mac | multicast-mac) mac-address <publish>;
 preferred;
 primary;
 vrrp-inet6-group group-id {
 accept-data | no-accept-data;
 advertisements-threshold number;
 authentication-key key;
 authentication-type authentication;
 fast-interval milliseconds;
 inet6-advertise-interval milliseconds;
 preempt | no-preempt {
 hold-time seconds;
 }
 priority number;
 track {
 interface interface-name {
 bandwidth-threshold bandwidth priority-cost number;
 priority-cost number;
 }
 priority-hold-time seconds;
 }
 }
}

```

```

 route ip-address/mask routing-instance instance-name priority-cost cost;
 }
 virtual-inet6-address [addresses];
 virtual-link-local-address ipv6-address;
 vrrp-inherit-from {
 active-group group-number;
 active-interface interface-name;
 }
}
}
(dad-disable | no-dad-disable);
filter {
 input filter-name;
 output filter-name;
}
mtu bytes;
nd6-stale-time seconds;
no-neighbor-learn;
no-redirects;
policer {
 input policer-name;
 output policer-name;
}
rpf-check {
 fail-filter filter-name;
 mode {
 loose;
 }
}
}
family iso {
 address interface-address;
 mtu bytes;
}
family mpls {
 filter {
 input filter-name;
 output filter-name;
 }
 mtu bytes;
 policer {
 input policer-name;
 output policer-name;
 }
}
native-inner-vlan-id vlan-id;
proxy-arp (restricted | unrestricted);
(traps | no-traps);
vlan-id-list [vlan-id's];
vlan-id-range [vlan-id-range];
}
}

```

Hierarchy Level [edit interfaces *interface-name*

|                                 |                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.3R2 for EX Series switches.<br><b>irb</b> option introduced in Junos OS Release 13.2 for the QFX Series. |
| <b>Description</b>              | Configure the properties of a specific integrated bridging and routing (IRB) interface.<br><br>The remaining statements are explained separately.    |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                              |

## **l3-interface (VLAN)**

|                            |                                                                                                                                                                                                                                     |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>l3-interface (vlan.<i>logical-interface-number</i>   irb.<i>logical-interface-number</i>);</code>                                                                                                                             |
| <b>Hierarchy Level</b>     | <code>[edit vlans <i>vlan-name</i>]</code>                                                                                                                                                                                          |
| <b>Release Information</b> | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br><b>irb</b> option introduced in Junos OS Release 13.2 for the QFX Series.                                                                                      |
| <b>Description</b>         | Associate a Layer 3 interface with the VLAN. Configure Layer 3 interfaces on trunk ports to allow the interface to transfer traffic between VLANs. Traffic between VLANs must be routed, which requires a common Layer 3 interface. |
| <b>Default</b>             | No Layer 3 (routing) interface is associated with the VLAN.                                                                                                                                                                         |
| <b>Options</b>             | <code>vlan.<i>logical-interface-number</i></code> —Number of the logical interface. Use the <b>unit</b> number that you used when you created the <b>vlan</b> interface with a <b>set interfaces vlan unit</b> statement.           |



**NOTE:** Use this statement with versions of Junos OS that do not support Enhanced Layer 2 Software (ELS).

`irb.logical-interface-number`—Logical interface defined with a **set interfaces irb** statement.



**NOTE:** Use this statement with versions of Junos OS that support Enhanced Layer 2 Software (ELS).

|                                 |                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show ethernet-switching interfaces on page 1998</a></li> <li>• <a href="#">show vlans on page 2395</a></li> </ul> |

## mac (Static MAC-Based VLANs)

---

|                                 |                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>mac <i>mac-address</i> {<br/>    next-hop <i>interface-name</i>;<br/>}</code>                                                        |
| <b>Hierarchy Level</b>          | <code>[edit ethernet-switching-options static vlan <i>vlan-name</i>]</code>                                                                |
| <b>Description</b>              | <p>Specify the MAC address to add to the Ethernet switching table.</p> <p>The remaining statement is explained separately.</p>             |
| <b>Options</b>                  | <i>mac-address</i> —MAC address                                                                                                            |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Adding a Static MAC Address Entry to the Ethernet Switching Table (CLI Procedure)</i></li></ul> |



## members

|                            |                                                                                               |
|----------------------------|-----------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>members [(all   <i>names</i>   <i>vlan-ids</i>)];</code>                                |
| <b>Hierarchy Level</b>     | [edit <a href="#">interfaces</a> <i>interface-name</i> unit 0 family ethernet-switching vlan] |
| <b>Release Information</b> | Statement introduced in Junos OS Release 11.1 for the QFX Series.                             |
| <b>Description</b>         | For trunk interfaces, configure the VLANs for which the interface can carry traffic.          |



**TIP:** To display a list of all configured VLANs on the system, including VLANs that are configured but not committed, type `?` after `vlan` or `vlands` in your configuration mode command line. Note that only one VLAN is displayed for a VLAN range.

**Options** `all`—Specify that this trunk interface be a member of all the VLANs that are configured on this switch. When a new VLAN is configured on the switch, this trunk interface automatically becomes a member of the VLAN.



**NOTE:** Each VLAN that is configured must have a specified VLAN ID when you attempt to commit the configuration; otherwise, the configuration commit fails. Also, `all` cannot be the name of a VLAN on the switch.

*names*—Names of one or more VLANs.

*vlan-ids*—Numeric identifiers of one or more VLANs.

**Required Privilege Level**  
`routing`—To view this statement in the configuration.  
`routing-control`—To add this statement to the configuration.

**Related Documentation**

- [Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 2096](#)
- [Understanding Bridging and VLANs on page 2089](#)
- [show ethernet-switching interfaces on page 1998](#)
- [show vlans on page 2395](#)

## native-vlan-id

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>native-vlan-id <i>vlan-id</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | For platforms without ELS:<br><br>[edit <a href="#">interfaces</a> <i>interface-name</i> unit 0 family ethernet-switching],<br><br>For platforms with ELS:<br><br>[edit <a href="#">interfaces</a> <i>interface-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | <p>Configure the VLAN identifier to associate with untagged packets received on the interface. The logical interface on which untagged packets are received must be configured with the same VLAN ID as the native VLAN ID configured on the physical interface. To configure the logical interface, include the <b>vlan-id</b> statement (matching the <b>native-vlan-id</b> statement on the physical interface) at the [edit <a href="#">interfaces</a> <i>interface-name</i> unit <i>logical-unit-number</i>] hierarchy level.</p> <p>When the <b>native-vlan-id</b> statement is combined with the <a href="#">interface-mode</a> statement, untagged packets are accepted and forwarded within the bridge domain or VLAN that is configured with the matching VLAN ID.</p> <p>When the <b>native-vlan-id</b> statement is combined with the <a href="#">flexible-vlan-tagging</a> statement, untagged packets are accepted on the interfaces that are configured for Q-in-Q tunneling.</p> <p>.</p> |
| <b>Options</b>                  | <p><b>vlan-id</b>—Numeric identifier of the VLAN.</p> <p><b>Range:</b> 1 through 4094</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Junos OS Network Interfaces Configuration Guide</a></li><li>•</li><li>• <a href="#">show ethernet-switching interfaces on page 1998</a></li><li>• <a href="#">show vlans on page 2395</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

## packet-action

**Syntax** `packet-action action;`

**Hierarchy Level** [edit bridge-domains *bridge-domain-name* bridge-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit bridge-domains *bridge-domain-name* bridge-options **interface-mac-limit** *limit*],  
 [edit logical-systems *logical-system-name* bridge-domains *bridge-domain-name* bridge-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit logical-systems *logical-system-name* bridge-domains *bridge-domain-name* bridge-options **interface-mac-limit** *limit*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options **interface-mac-limit** *limit*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* switch-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* switch-options **interface-mac-limit** *limit*],  
 [edit logical-systems *logical-system-name* switch-options **interface-mac-limit** *limit*],  
 [edit protocols l2-learning global-mac-limit *limit*],  
 [edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options **interface-mac-limit** *limit*],  
 [edit routing-instances *routing-instance-name* protocols evpn interface-mac-limit (VPLS)],  
 [edit routing-instances *routing-instance-name* protocols evpn interface *interface-name* interface-mac-limit (VPLS)],  
 [edit routing-instances *routing-instance-name* protocols evpn mac-table-size *limit*],  
 [edit routing-instances *routing-instance-name* switch-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit routing-instances *routing-instance-name* switch-options **interface-mac-limit** *limit*],  
 [edit switch-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit switch-options **interface-mac-limit** *limit*],  
 [edit switch-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit switch-options **interface-mac-limit** *limit*],  
 [edit switch-options **mac-table-size** *limit*],  
 [edit switch-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit vlans *vlan-name* switch-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit vlans *vlan-name* switch-options **interface-mac-limit** *limit*],  
 [edit vlans *vlan-name* switch-options **mac-table-size** *limit*],  
 [edit vlans *vlan-name* switch-options **interface-mac-limit** *limit*],  
 [edit vlans *vlan-name* switch-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit vlans *vlan-name* switch-options **mac-table-size** *limit*]

**Release Information** Statement introduced in Junos OS Release 8.4.  
 Support for the **switch-options** statement added in Junos OS Release 9.2.  
 Support for top-level configuration for the **virtual-switch** type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch.

Support for logical systems added in Junos OS Release 9.6.

[edit switch-options interface *interface-name* interface-mac-limit *limit*], [edit switch-options interface-mac-limit *limit*], [edit switch-options mac-table-size *limit*], [edit vlans *vlan-name* switch-options interface *interface-name* interface-mac-limit *limit*], [edit vlans *vlan-name* switch-options interface-mac-limit *limit*], and [edit vlans *vlan-name* switch-options mac-table-size *limit*] hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.

Support for EVPNs introduced in Junos OS Release 13.2 on MX Series 3D Universal Edge Routers.

Support at the [edit switch-options interface *interface-name* interface-mac-limit *limit*] hierarchy level and hierarchy levels under [edit vlans *vlan-name*] introduced in Junos OS Release 13.2X50-D10 for EX Series switches and Junos OS Release 13.2 for the QFX Series.

**Description** Specify the action taken when packets with new source MAC addresses are received after the MAC address limit is reached. If this statement is not configured, packets with new source MAC addresses are forwarded by default.

**Default**



**NOTE:** On a QFX Series Virtual Chassis, if you include the shutdown option at the [edit vlans *vlan-name* switch-options interface *interface-name* interface-mac-limit packet-action] hierarchy level and issue the commit operation, the system generates a commit error. The system does not generate an error if you include the shutdown option at the [edit switch-options interface *interface-name* interface-mac-limit packet-action] hierarchy level.

Disabled. The default is for packets for new source MAC addresses to be forwarded after the MAC address limit is reached.

**Options**

**drop**—Drop packets with new source MAC addresses, and do not learn the new source MAC addresses.

**drop-and-log**—(EX Series switches and QFX Series only) Drop packets with new source MAC addresses, and generate an alarm, an SNMP trap, or a system log entry.

**log**—(EX Series switches and QFX Series only) Hold packets with new source MAC addresses, and generate an alarm, an SNMP trap, or a system log entry.

**none**—(EX Series switches and QFX Series only) Forward packets with new source MAC addresses, and learn the new source MAC address.

**shutdown**—(EX Series switches and QFX Series only) Disable the specified interface, and generate an alarm, an SNMP trap, or a system log entry.

**Required Privilege Level**

routing—To view this statement in the configuration.


routing-control—To add this statement to the configuration.

**Related  
Documentation**

- *Configuring EVPN Routing Instances*
- *Configuring EVPN Routing Instances on EX9200 Switches*
- [Configuring MAC Limiting \(CLI Procedure\) on page 2142](#)
- [Configuring Persistent MAC Learning \(CLI Procedure\) on page 6035](#)
- *Understanding Layer 2 Learning and Forwarding for Bridge Domains*
- [Layer 2 Learning and Forwarding for VLANs Overview on page 2088](#)
- *Understanding Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports*
- [Layer 2 Learning and Forwarding for VLANs Overview on page 2088](#)
- *Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port*

## port-mode

---

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | port-mode (access   tagged-access   trunk);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Hierarchy Level          | [edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> family ethernet-switching]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Release Information      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Description              | <div> <b>NOTE:</b> This statement does not support the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that supports ELS, see <a href="#">interface-mode</a>. For ELS details, see “<a href="#">Getting Started with Enhanced Layer 2 Software</a>” on page 41.</div> <p>Configure whether an interface on the switch operates in access, tagged access, or trunk mode.</p>                                                                                                                                                                                                                                                         |
| Default                  | All switch interfaces are in access mode.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Options                  | <p><b>access</b>—Have the interface operate in access mode. In this mode, the interface can be in a single VLAN only. Access interfaces typically connect to network devices, such as PCs, printers, IP telephones, and IP cameras.</p> <p><b>tagged-access</b>—Have the interface operate in tagged-access mode. In this mode, the interface can be in multiple VLANs. Tagged access interfaces typically connect to network devices, such as PCs, printers, IP telephones, and IP cameras.</p> <p><b>trunk</b>—Have the interface operate in trunk mode. In this mode, the interface can be in multiple VLANs and can multiplex traffic between different VLANs. Trunk interfaces typically connect to other switches and to routers on the LAN.</p> |
| Required Privilege Level | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Related Documentation    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Reflective Relay for Use with VEPA Technology</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## service-id

---

|                                 |                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>service-id <i>number</i>;</code>                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit switch-options]<br>[edit vlans <i>vlan-name</i> ]                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.3R2 for EX Series switches and MX Series routers.<br>Statement introduced in Junos OS Release 13.2 for the QFX Series. |
| <b>Description</b>              | Specify a service identifier for each multichassis aggregated Ethernet interface that belongs to a link aggregation group (LAG).                                   |
| <b>Options</b>                  | <b>number</b> —A number that identifies a particular service.<br><b>Range:</b> 1 through 65535                                                                     |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system control—To add this statement to the configuration.                                                  |

## switch-options

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>switch-options {<br/>  interface <i>interface-name</i> {<br/>    interface-mac-limit <i>limit</i> {<br/>      packet-action drop;<br/>    }<br/>    no-mac-learning;<br/>    static-mac <i>static-mac-address</i> {<br/>      vlan-id <i>number</i>;<br/>    }<br/>  }<br/>  interface-mac-limit <i>limit</i> {<br/>    packet-action drop;<br/>  }<br/>  mac-statistics;<br/>  mac-table-size <i>limit</i> {<br/>    packet-action drop;<br/>  }<br/>  no-mac-learning;<br/>  service-id <i>number</i>;<br/>  vtep-source-interface<br/>}</pre> |
| <b>Hierarchy Level</b>          | <pre>[edit <i>number</i>],<br/>[edit vlans <i>vlan--name</i>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> vlans<br/>  <i>vlan-name</i>],<br/>[edit routing-instances <i>routing-instance-name</i> vlans <i>vlan-name</i>]</pre>                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches and MX Series routers.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | <p>Configure Layer 2 learning and forwarding properties for a VLAN or a virtual switch.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                        |



## static (Static MAC-Based VLANs)


|                                 |                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>static {   vlan <i>vlan-name</i> {     mac <i>mac-address</i> {       next-hop <i>interface-name</i>;     }   } }</pre>                          |
| <b>Hierarchy Level</b>          | [edit ethernet-switching-options]                                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for EX Series switches.                                                                                 |
| <b>Description</b>              | <p>Specify VLAN and MAC addresses to add to the Ethernet switching table.</p> <p>The remaining statements are explained separately.</p>               |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Adding a Static MAC Address Entry to the Ethernet Switching Table (CLI Procedure)</a></li> </ul> |

## static-mac



|                                 |                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>static-mac <i>mac-address</i>;</code>                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit vlans <i>vlan-name</i> switch-options interface <i>interface-name</i> ]                                                                                      |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 13.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX Series.</p>              |
| <b>Description</b>              | Specify a static MAC address to assign to this interface.                                                                                                          |
| <b>Options</b>                  | <i>mac-address</i> —MAC address                                                                                                                                    |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Adding a Static MAC Address Entry to the Ethernet Switching Table (CLI Procedure)</a> on page 2133</li> </ul> |

## vlan-id (VLANs)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>vlan-id <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | <p>For platforms without ELS:</p> <pre>[edit vlans <i>vlan-name</i> vlan-range]</pre> <p>For platforms without ELS and with ELS:</p> <pre>[edit vlans <i>vlan-name</i>]</pre> <p>For ELS platforms only:</p> <pre>[edit interfaces <i>interface-name</i> unit <i>number</i>]<br/>[edit vlans <i>vlan-name</i> vlan-id-list]</pre>                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Configure an 802.1Q tag to apply to all traffic that originates on the VLAN.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Default</b>                  | <p>On a QFX3500 and QFX3500 switch, if you use the default factory configuration, all traffic originating on the VLAN is untagged and has a VLAN identifier of 1. The number zero is reserved for priority tagging and the number 4093 is also reserved.</p> <p>On a QFX5100 switch, if you use the default factory configuration, all traffic originating on the VLAN is untagged and has a VLAN identifier of 1. The number zero is reserved for priority tagging and the number 4093 is also reserved.</p>                                                                  |
|                                 | <div> <b>NOTE:</b> You can only create up to 4090 VLANs on a QFX5100 switch. If you create more than 4090 VLANs, the interfaces associated with the extra VLANs are not displayed in the <code>show vlans</code> command output. For example, if you create 4094 VLANs, the extra VLANs will not have interfaces associated with the VLANs. The order in which you configure the extra VLANs determines which interfaces are missing from the <code>show vlans</code> command output.</div> |
| <b>Options</b>                  | <p><i>number</i> —VLAN tag identifier.</p> <p><b>Range:</b> 0 through 4093.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Setting Up Bridging with Multiple VLANs</i></li><li>• <i>Understanding Bridging and VLANs</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                    |

## vlan-id-list

|                                                                                                                                                                                                                                                                                                                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                   | <code>vlan-id-list [ <i>vlan-id-numbers</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                          | <p>[edit bridge-domains <i>bridge-domain-name</i>],<br/>         [edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i>],<br/>         [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i>],<br/>         [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i>],<br/>         [edit interfaces <i>interface-name</i> unit 0],<br/>         [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>],<br/>         [edit vlans <i>vlan-name</i>]</p>                                                                                                                                                                                                   |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                      | <p>Statement introduced in Junos OS Release 9.4.<br/>         Support for logical systems added in Junos OS Release 9.6.<br/>         Statement introduced in Junos OS Release 12.3R2 for EX Series switches.<br/>         Statement introduced in Junos OS Release 13.2 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>                                                                                                                                                                                                                                                                                                              | <p>Specify a VLAN identifier list to use for a bridge domain or VLAN in trunk mode.</p> <p>Specify the <b>trunk</b> option in the <b>interface-mode</b> statement to accept packets with a VLAN ID that matches the list of VLAN IDs specified in the <b>vlan-id-list</b> statement to forward the packet within the bridge domain or VLAN configured with the matching VLAN ID. Specify the <b>access</b> option to accept packets with no VLAN ID to forward the packet within the bridge domain or VLAN configured with the VLAN ID that matches the VLAN ID specified in the <b>vlan-id</b> statement.</p> <p>This statement also enables you to bind a logical interface to a list of VLAN IDs, thereby configuring the logical interface to receive and forward a frame with a tag that matches the specified VLAN ID list.</p> |
| <div style="display: flex; align-items: center;">  <div style="margin-left: 10px;"> <p><b>WARNING:</b> On some EX and QFX Series switches, you can apply no more than eight VLAN identifier lists to a physical interface.</p> </div> </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                  | <p><b><i>vlan-id-numbers</i></b>—Valid VLAN identifiers. You can combine individual numbers with range lists by including a hyphen.</p> <p><b>Range:</b> 0 through 4095</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <div style="display: flex; align-items: center;">  <div style="margin-left: 10px;"> <p><b>NOTE:</b> On EX Series switches and the QFX Series, the range is 0 through 4094.</p> </div> </div>                                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                 | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

- Related Documentation**
- *Configuring a Bridge Domain*
  - *Configuring a VLAN*
  - *Configuring VLAN Identifiers for Bridge Domains and VPLS Routing Instances*
  - *Configuring VLAN Identifiers for VLANs and VPLS Routing Instances*

---

## vlan-rewrite

---

|                                 |                                                                                                                                                                                                                                                   |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | vlan-rewrite translate (200 500   201 501)                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> unit <i>number</i> family bridge interface-mode trunk]<br>[edit interfaces <i>interface-name</i> unit <i>number</i> family ethernet-switching interface-mode trunk]                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4.<br>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for the QFX Series.                                                     |
| <b>Description</b>              | Translates an incoming VLAN to a bridge-domain VLAN, corresponding counter translation at egress. Supports translation of VLAN 200 to VLAN 500 and VLAN 201 to VLAN 501. Other valid VLANs pass through without translation.                      |
| <b>Options</b>                  | <b>translate 200 500</b> —Translates incoming packets with VLAN 200 to 500.<br><br><b>translate 201 501</b> —Translates incoming packets with VLAN 201 to 501.<br><br><b>translate 202 502</b> —Translates incoming packets with VLAN 202 to 502. |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Rewriting a VLAN Tag and Adding a New Tag</i></li></ul>                                                                                                                                                |

## vlan-tagging

---

|                                 |                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | vlan-tagging;                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces</a> <i>interface-name</i> ]<br>[edit <a href="#">interfaces</a> <a href="#">interface-range</a> <i>interface-range-name</i> ]            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                     |
| <b>Description</b>              | Enable VLAN tagging. The platform receives and forwards single-tag frames with 802.1Q VLAN tags.                                                                      |
| <b>Default</b>                  | VLAN tagging is disabled by default.                                                                                                                                  |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">vlan-id on page 3036</a></li><li>• <a href="#">Configuring a Layer 3 Logical Interface on page 2860</a></li></ul> |

## vlan-tags

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>vlan-tags outer <i>number</i> inner <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | <code>[edit bridge-domains <i>bridge-domain-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i></code><br><code>bridge-domains <i>bridge-domain-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i>]</code><br><code>[edit vlans <i>vlan-name</i>]</code> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.4.<br>Support for logical systems added in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2X51-D10 for QFX Series switches.                                                                                                                                                                                                                                           |
| <b>Description</b>              | Specify dual VLAN identifier tags for a bridge domain, VLAN, or VPLS routing instance.                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <code>outer <i>number</i></code> —A valid VLAN identifier.<br><br><code>inner <i>number</i></code> —A valid VLAN identifier.                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | <code>routing</code> —To view this statement in the configuration.<br><code>routing-control</code> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring a Bridge Domain</i></li><li>• <i>Configuring a VLAN</i></li><li>• <i>Configuring VLAN Identifiers for Bridge Domains and VPLS Routing Instances</i></li><li>• <i>Configuring VLAN Identifiers for VLANs and VPLS Routing Instances</i></li><li>• <i>Configuring a Layer 2 Virtual Switch .</i></li><li>• <i>Configuring a Layer 2 Virtual Switch</i></li></ul>                                                                                             |

## vlan

```
Syntax vlan {
 vlan-name {
 description text-description;
 domain-type bridge;
 forwarding-options {
 dhcp-security {
 arp-inspection;
 group group-name {
 interface interface-name {
 static-ip ip-address {
 mac mac-address;
 }
 }
 }
 overrides {
 no-option82;
 trusted;
 untrusted;
 }
 }
 }
 ip-source-guard;
 no-dhcp-snooping;
 option-82 {
 circuit-id {
 prefix {
 host-name;
 logical-system-name;
 routing-instance-name;
 }
 use-interface-description (device | logical);
 use-vlan-id;
 }
 remote-id {
 host-name hostname;
 use-interface-description (device | logical);
 use-string string;
 }
 vendor-id {
 use-string string;
 }
 }
 }
 }
 fip-security {
 examine-vn2vf;
 examine-vn2vn {
 beacon-period milliseconds;
 }
 fc-map fc-map-value;
 interface interface-name {
 (fcoe-trusted | no-fcoe-trusted;)
 }
 }
}
```

```
l3-interface irb.logical-unit-number;
multicast-snooping-options {
 flood-groups [group-names];
 forwarding-cache {
 threshold {
 reuse threshold;
 suppress threshold;
 }
 }
 graceful-restart {
 disable;
 restart-duration duration;
 }
 host-outbound-traffic {
 dot1p bits;
 forwarding-class forwarding-class;
 }
 multichassis-lag-replicate-state;
 nexthop-hold-time time;
 options {
 syslog {
 level level;
 mark interval;
 upto level;
 }
 }
 traceoptions {
 file filename {
 files number;
 no-world-readable;
 size file-size;
 world-readable;
 }
 flag flag {
 disable;
 }
 }
}
switch-options {
 interface interface-name {
 interface-mac-limit limit {
 packet-action action;
 }
 static-mac mac-address;
 }
 interface-mac-limit limit {
 packet-action action;
 }
 mac-move-limit limit {
 packet-action action;
 }
 mac-table-size limit {
 packet-action drop;
 }
 no-mac-learning;
}
```



```

 }
 vlan-id number;
 vlan-id-list [vlan-id | vlan-id-vlan-id];
 vlan-tags
 inner value;
 outer value;
 }
 vxlan {
 ingress-node-replication
 ovsdb-managed
 }
 }
}

```

**Hierarchy Level** [edit]

**Release Information** Statement introduced in Junos OS Release 13.2 for the QFX Series.  
Statements for private VLANs and Q-in-Q tunneling introduced in Junos OS Release 12.1 for the QFX Series.

**Description** Configure VLAN properties on the QFX Series.

**Default** If you use the default factory configuration, all switch interfaces become part of the VLAN default.

**Options** *vlan-name*—Name of the VLAN. The name can contain letters, numbers, hyphens (-), and periods (.) and can be up to 255 characters long.

The remaining statements are described separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Understanding Bridging and VLANs on page 2089](#)
- [Configuring VLANs on page 2120](#)



## CHAPTER 81

# MAC Address Configuration Statements

- [global-mac-table-aging-time](#) on page 2315
- [mac-limit](#) on page 2316
- [mac-notification](#) on page 2317
- [mac-statistics](#) on page 2318
- [mac-table-size](#) on page 2320
- [notification-interval](#) on page 2321

### [global-mac-table-aging-time](#)

---

|                                 |                                                                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>global-mac-table-aging-time seconds;</code>                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit protocols l2-learning]                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.2.<br>Support for logical systems added in Junos OS Release 9.6.                                                                                                                                           |
| <b>Description</b>              | Configure the timeout interval for entries in the MAC table.                                                                                                                                                                                          |
| <b>Default</b>                  | 300 seconds                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <b>seconds</b> —Time elapsed before MAC table entries are timed out and entries are deleted from the table.<br><b>Range:</b> For MX Series routers: 10 through 1 million; for EX Series and QFX Series switches: 60 through 1 million                 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the MAC Table Timeout Interval</a></li><li>• <a href="#">Configuring MAC Table Aging (CLI Procedure)</a></li><li>• <a href="#">Configuring MAC Table Aging</a> on page 2144</li></ul> |

## mac-limit

---

|                            |                                                                                  |
|----------------------------|----------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>mac-limit <i>number</i>;</code>                                            |
| <b>Hierarchy Level</b>     | <code>[edit vlans <i>vlan-name</i>]</code>                                       |
| <b>Release Information</b> | Statement introduced in Junos OS Release 11.1 for the QFX Series.                |
| <b>Description</b>         | Configure the number of MAC addresses allowed on a VLAN.                         |
| <b>Default</b>             | MAC limit is disabled.                                                           |
| <b>Options</b>             | <i>number</i> —Maximum number of MAC addresses.<br><b>Range:</b> 1 through 32768 |



**NOTE:** This statement is not supported on QFabric systems.

---

|                                 |                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show vlans on page 2395</a></li><li>• <a href="#">Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 2096</a></li><li>• <a href="#">Configuring MAC Table Aging</a></li><li>• <a href="#">Understanding Bridging and VLANs</a></li></ul> |

## mac-notification

---

|                                 |                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>mac-notification {<br/>  notification-interval <i>seconds</i>;<br/>}</pre>                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit ethernet-switching-options]<br>[edit switch-options]                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Hierarchy level <b>[edit switch-options]</b> added in Junos OS Release 14.1X53-D10 for EX Series and QFX Series.                                                                               |
| <b>Description</b>              | <p>Enable MAC notification for a switch. If you configure this statement without setting a notification interval, MAC notification is enabled with the default MAC notification interval of 30 seconds.</p> <p>The remaining statement is explained separately.</p> |
| <b>Default</b>                  | MAC notification is disabled by default.                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring MAC Notification</i></li><li>• <a href="#">Configuring MAC Notification (CLI Procedure) on page 2140</a></li></ul>                                                                                           |

## mac-statistics

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | mac-statistics;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | <p>[edit bridge-domains <i>bridge-domain-name</i> bridge-options],<br/>[edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> switch-options],<br/>[edit logical-systems <i>logical-system-name</i> switch-options],<br/>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],<br/>[edit routing-instances <i>routing-instance-name</i> switch-options],<br/>[edit routing-instances <i>routing-instance-name</i> protocols evpn],<br/>[edit switch-options],<br/>[edit switch-options],<br/>[edit vlans <i>vlan-name</i> switch-options]</p>      |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.4.</p> <p>Support for the <b>switch-options</b> statement added in Junos OS Release 9.2.</p> <p>Support for top-level configuration for the <b>virtual-switch</b> type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p> <p>[edit switch-options] and [edit vlans <i>vlan-name</i> switch-options] hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Support for EVPNs added in Junos OS Release 13.2 for MX 3D Series routers.</p> <p>[edit switch-options] and [edit vlans <i>vlan-name</i> switch-options] hierarchy levels introduced in Junos OS Release 13.2 for the QFX Series.</p> |
| <b>Description</b>              | (MX Series routers, EX Series switches, and QFX Series only) For bridge domains or VLANs, enable MAC accounting either for a specific bridge domain or VLAN, or for a set of bridge domains or VLANs associated with a Layer 2 trunk port.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Default</b>                  | disabled                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Understanding Layer 2 Learning and Forwarding for Bridge Domains</i></li><li>• <a href="#">Layer 2 Learning and Forwarding for VLANs Overview on page 2088</a></li><li>• <i>Understanding Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports</i></li><li>• <i>Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port</i></li><li>• <i>Configuring EVPN Routing Instances</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                 |

- *Configuring EVPN Routing Instances on EX9200 Switches*

## mac-table-size

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>mac-table-size <i>limit</i> {<br/>    packet-action drop;<br/>}</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>     | <p>[edit bridge-domains <i>bridge-domain-name</i> bridge-options],<br/> [edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],<br/> [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],<br/> [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> switch-options],<br/> [edit logical-systems <i>logical-system-name</i> switch-options],<br/> [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],<br/> [edit routing-instances <i>routing-instance-name</i> switch-options],<br/> [edit switch-options],<br/> [edit vlans <i>vlan-name</i> switch-options]</p>                          |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 8.4.</p> <p>Support for the <b>switch-options</b> statement added in Junos OS Release 9.2.</p> <p>Support for top-level configuration for the <b>virtual-switch</b> type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p> <p><b>[edit switch-options]</b> and <b>[edit vlans <i>vlan-name</i> switch-options]</b> hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Support at the <b>[edit vlans <i>vlan-name</i> switch-options]</b> hierarchy level introduced in Junos OS Release 13.2 for the QFX Series.</p> |
| <b>Description</b>         | Modify the size of the MAC address table for the bridge domain or VLAN, a set of bridge domains or VLANs associated with a trunk port, or a virtual switch. The default is 5120 MAC addresses.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |



**NOTE:** For multichassis link aggregation (MC-LAG) peers in active-active mode, configuring the **mac-table-size** statement or changing the **mac-table-size** configuration when traffic is flowing can cause the MAC entries to be out of synchronization between the two MC-LAG peers, which might result in flooding. To avoid flooding, you must either halt traffic forwarding and then configure the **mac-table-size** statement or use the **commit at** configuration statement to commit the changes at the same time in both the peer nodes.

Alternatively, if flooding does occur, you can clear the bridge MAC table on both the routers by using the **clear bridge mac-table** command. Running this command ensures that the MAC entries are re-learned and in synchronization between both the peers.



**Options** *limit*—Specify the maximum number of addresses in the MAC address table.  
**Range:** 16 through 1,048,575 MAC addresses  
**Default:** 5120 MAC addresses  
 There is no default MAC address limit for the **mac-table-size** statement at the **[edit switch-options]** hierarchy level. The number of MAC addresses that can be learned is only limited by the platform, 65,535 MAC addresses for EX Series switches and 1,048,575 MAC addresses for other devices.

The remaining statement is explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
 routing-control—To add this statement to the configuration.

**Related Documentation**

- *Understanding Layer 2 Learning and Forwarding for Bridge Domains*
- [Layer 2 Learning and Forwarding for VLANs Overview on page 2088](#)
- *Understanding Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports*
- *Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port*

## notification-interval

**Syntax** notification-interval *seconds*;

**Hierarchy Level** [edit ethernet-switching-options mac-notification]  
 [edit switch-options mac-notification]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.  
 Hierarchy level **[edit switch-options]** added in Junos OS Release 14.1X53-D10 for EX Series and QFX Series.

**Description** Configure the MAC notification interval for a switch.

The MAC notification interval is the amount of time the switch waits before sending learned or unlearned MAC address SNMP notifications to the network management server. For instance, if the MAC notification interval is set to 10, all of the MAC address addition and removal SNMP notifications are sent to the network management system every 10 seconds.

**Options** *seconds*—The MAC notification interval, in seconds.  
**Range:** 1 through 60  
**Default:** 30

**Required Privilege Level** routing—To view this statement in the configuration.  
 routing-control—To add this statement to the configuration.

**Related Documentation**

- *Configuring MAC Notification*
- [Configuring MAC Notification \(CLI Procedure\) on page 2140](#)



# STP Configuration Statements

- [bpdu-block on page 2324](#)
- [bpdu-block-on-edge on page 2325](#)
- [bpdu-timeout-action on page 2326](#)
- [bridge-priority \(Spanning Trees\) on page 2327](#)
- [configuration-name on page 2328](#)
- [cost on page 2329](#)
- [disable \(Spanning Trees\) on page 2330](#)
- [bpdu-block on page 2331](#)
- [bpdu-block-on-edge on page 2332](#)
- [bpdu-timeout-action on page 2333](#)
- [bridge-priority \(Spanning Trees\) on page 2334](#)
- [configuration-name on page 2335](#)
- [cost on page 2336](#)
- [disable \(Spanning Trees\) on page 2337](#)
- [disable-timeout \(Spanning Trees\) on page 2338](#)
- [edge on page 2339](#)
- [force-version \(IEEE 802.1D STP\) on page 2340](#)
- [forward-delay on page 2341](#)
- [hello-time on page 2342](#)
- [interface \(BPDU\) on page 2343](#)
- [interface \(Spanning Tree\) on page 2344](#)
- [max-age on page 2345](#)
- [max-hops on page 2346](#)
- [mode on page 2347](#)
- [msti on page 2348](#)
- [mstp on page 2349](#)
- [no-root-port on page 2351](#)
- [priority \(Protocols STP\) on page 2352](#)

- [protocol](#) on page 2353
- [protocols \(STP Type\)](#) on page 2354
- [revision-level](#) on page 2355
- [rstp](#) on page 2356
- [traceoptions \(Spanning Tree\)](#) on page 2357
- [vlan \(MSTP\)](#) on page 2360
- [vlan \(VSTP\)](#) on page 2361
- [vlan \(VSTP\)](#) on page 2362
- [vlan-group](#) on page 2363
- [vstp](#) on page 2364

---

## bpdu-block

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>bpdu-block {<br/>    interface (<i>interface-name</i> disable   all);<br/>    disable-timeout <i>seconds</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit protocols layer2-control ]                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | <p>Enable BPDU blocking on an interface.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Understanding BPDU Protection for Spanning-Tree Instance Interfaces</i></li><li>• <i>BPDU Protection for Individual Spanning-Tree Instance Interfaces</i></li><li>• <i>Configuring BPDU Protection for Spanning-Tree Instance Interfaces</i></li><li>• <i>show spanning-tree bridge</i></li><li>• <i>show spanning-tree interface</i></li><li>• <i>Understanding BPDU Protection for STP, RSTP, and MSTP on EX Series Switches</i></li></ul> |

## bpdu-block-on-edge

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | bpdu-block-on-edge;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols (mstp   rstp   vstp)],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>(mstp   rstp   vstp)],<br>[edit protocols ( mstp   rstp  vstp )],<br>[edit routing-instances <i>routing-instance-name</i> protocols (mstp   rstp   vstp)]                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4.<br>Support for logical systems added in Junos OS Release 9.6.                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Enable BPDU blocking on the edge ports of a virtual switch.                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Understanding BPDU Protection for Spanning-Tree Instance Interfaces</i></li> <li>• <i>BPDU Protection on All Edge Ports of the Bridge</i></li> <li>• <i>Configuring BPDU Protection on All Edge Ports</i></li> <li>• <a href="#">Configuring BPDU Protection on Spanning Tree Interfaces on page 2205</a></li> <li>• <a href="#">rstp on page 2356</a></li> <li>• <a href="#">mstp on page 2349</a></li> <li>• <a href="#">vstp on page 2364</a></li> </ul> |

## bpdu-timeout-action

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | bpdu-timeout-action (log   block);                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols (mstp   rstp   vstp)],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>(mstp   rstp   vstp)],<br>[edit protocols (mstp   rstp   vstp) interface ],<br>[edit routing-instances <i>routing-instance-name</i> protocols (mstp   rstp   vstp)]                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4.<br>Support for logical systems added in Junos OS Release 9.6.                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Provide STP loop protection for a given STP family protocol interface.                                                                                                                                                                                                                                                                                                                                                              |
| <b>Default</b>                  | If the <b>bpdu-timeout-action</b> statement is not configured, an interface that stops receiving BPDUs will transition to the designated port (forwarding) state, creating a potential loop.                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <b>log</b> —The interface logs the fact that it has not received BPDUs during the timeout interval.<br><br><b>block</b> —The interface is blocked and the fact that the interface has not received BPDUs during the timeout interval is logged.                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Understanding Loop Protection for Spanning-Tree Instance Interfaces</i></li><li>• <i>Configuring Loop Protection for a Spanning-Tree Instance Interface</i></li><li>• <i>Example: Enabling Loop Protection for Spanning-Tree Protocols</i></li><li>• <a href="#">rstp on page 2356</a></li><li>• <a href="#">mstp on page 2349</a></li><li>• <a href="#">vstp on page 2364</a></li></ul> |

## bridge-priority (Spanning Trees)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>bridge-priority <i>priority</i>;</code>                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit protocols <b>mstp</b> ],<br>[edit protocols mstp <b>msti</b> <i>msti-id</i> ],<br>[edit protocols <b>rstp</b> ],<br>[edit protocols <b>vstp</b> vlan <i>vlan-id</i> ]                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.                                                                                                                                                                                                                                       |
| <b>Description</b>              | Configure the bridge priority. The bridge priority determines which bridge is elected as the root bridge. If two bridges have the same path cost to the root bridge, the bridge priority determines which bridge becomes the designated bridge for a LAN segment.                                                                                                                                   |
| <b>Default</b>                  | 32,768                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <b>priority</b> —Bridge priority. It can be set only in increments of 4096.<br><b>Range:</b> 0 through 61,440<br><b>Default:</b> 32,768                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>show spanning-tree bridge</i></li> <li>• <i>show spanning-tree interface</i></li> <li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on page 2149</a></li> <li>• <i>Understanding MSTP for EX Series and QFX Series Switches</i></li> <li>• <i>Understanding VSTP for EX Series Switches and QFX Series Switches</i></li> </ul> |

## configuration-name

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>configuration-name <i>configuration-name</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mstp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols mstp],<br>[edit protocols mstp],<br>[edit routing-instances <i>routing-instance-name</i> protocols mstp]                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.4.<br>Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Support for logical systems added in Junos OS Release 9.6.                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Specify the configuration name , which is the MSTP region name carried in the MSTP BPDUs.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Understanding BPDUs Used for Exchanging Information Among Bridges</i></li><li>• <i>Configuring Multiple Spanning Tree Protocol</i></li><li>• <a href="#">Configuring MSTP on page 2170</a></li><li>• <i>show spanning-tree bridge</i></li><li>• <i>show spanning-tree interface</i></li><li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on page 2149</a></li><li>• <a href="#">Example: Configuring Faster Convergence and Improved Network Stability with RSTP on page 2173</a></li><li>• <i>Understanding MSTP for EX Series and QFX Series Switches</i></li></ul> |



## cost

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>cost cost;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols (mstp   rstp   vstp) interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols mstp msti <i>msti-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols vstp <b>vlan</b> <i>vlan-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (mstp   rstp   vstp) interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols mstp msti <i>msti-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vstp <b>vlan</b> <i>vlan-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (mstp   rstp   vstp) interface <i>interface-name</i>],</p> <p>[edit protocols mstp msti <i>msti-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols vstp <b>vlan</b> <i>vlan-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (mstp   rstp   vstp) interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols mstp msti <i>msti-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vstp <b>vlan</b> <i>vlan-id</i> interface <i>interface-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.4.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Configure link cost to control which bridge is the designated bridge and which port is the designated port. By default, the link cost is determined by the link speed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <p><b>cost</b>—(Optional) Link cost associated with the port.</p> <p><b>Range:</b> 1 through 200,000,000</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Spanning-Tree Instance Interface</i></li> <li>• <i>Spanning-Tree Instance Interface Cost</i></li> <li>• <i>show spanning-tree bridge</i></li> <li>• <i>show spanning-tree interface</i></li> <li>• <i>Understanding RSTP for EX Series and QFX Series Switches</i></li> <li>• <i>Understanding MSTP for EX Series and QFX Series Switches</i></li> <li>• <i>Understanding VSTP for EX Series Switches and QFX Series Switches</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

## disable (Spanning Trees)

---

|                            |                                                                                                                                                                                                                                                                                                                                                          |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | disable;                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>     | [edit protocols <b>mstp</b> interface <i>interface-name</i> ,<br>[edit protocols mstp <b>msti</b> <i>msti-id</i> vlan (all <i>vlan-id</i>   <i>vlan-name</i> ) interface <i>interface-name</i> ],<br>[edit protocols <b>rstp</b> interface <i>interface-name</i> ],<br>[edit protocols <b>vstp</b> vlan <i>vlan-id</i> interface <i>interface-name</i> ] |
| <b>Release Information</b> | Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.<br>Statement updated in Junos OS Release 15.1 for EX Series switches.                                                                                                                      |
| <b>Description</b>         | Disable MSTP, RSTP, or VSTP on a specific interface.                                                                                                                                                                                                                                                                                                     |



**NOTE:** You cannot disable spanning tree parameters for all interfaces.

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring RSTP (CLI Procedure) on page 2190</a></li><li>• <a href="#">Configuring MSTP on page 2170</a></li><li>• <a href="#">Configuring MSTP on page 2170</a></li><li>• <a href="#">Configuring VLAN Spanning Tree Protocol</a></li><li>• <a href="#">show spanning-tree bridge</a></li><li>• <a href="#">show spanning-tree interface</a></li></ul> |

## bpdu-block

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>bpdu-block {   interface (<i>interface-name</i> disable   all);   disable-timeout <i>seconds</i>; }</pre>                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit protocols layer2-control ]                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | <p>Enable BPDU blocking on an interface.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Understanding BPDU Protection for Spanning-Tree Instance Interfaces</i></li> <li>• <i>BPDU Protection for Individual Spanning-Tree Instance Interfaces</i></li> <li>• <i>Configuring BPDU Protection for Spanning-Tree Instance Interfaces</i></li> <li>• <i>show spanning-tree bridge</i></li> <li>• <i>show spanning-tree interface</i></li> <li>• <i>Understanding BPDU Protection for STP, RSTP, and MSTP on EX Series Switches</i></li> </ul> |

## bpdu-block-on-edge

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | bpdu-block-on-edge;                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols (mstp   rstp   vstp)],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>(mstp   rstp   vstp)],<br>[edit protocols ( mstp   rstp  vstp )],<br>[edit routing-instances <i>routing-instance-name</i> protocols (mstp   rstp   vstp)]                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4.<br>Support for logical systems added in Junos OS Release 9.6.                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Enable BPDU blocking on the edge ports of a virtual switch.                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Understanding BPDU Protection for Spanning-Tree Instance Interfaces</i></li><li>• <i>BPDU Protection on All Edge Ports of the Bridge</i></li><li>• <i>Configuring BPDU Protection on All Edge Ports</i></li><li>• <a href="#">Configuring BPDU Protection on Spanning Tree Interfaces on page 2205</a></li><li>• <a href="#">rstp on page 2356</a></li><li>• <a href="#">mstp on page 2349</a></li><li>• <a href="#">vstp on page 2364</a></li></ul> |

## bpdu-timeout-action

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | bpdu-timeout-action (log   block);                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols (mstp   rstp   vstp)],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>(mstp   rstp   vstp)],<br>[edit protocols (mstp   rstp   vstp) interface ],<br>[edit routing-instances <i>routing-instance-name</i> protocols (mstp   rstp   vstp)]                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4.<br>Support for logical systems added in Junos OS Release 9.6.                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Provide STP loop protection for a given STP family protocol interface.                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Default</b>                  | If the <b>bpdu-timeout-action</b> statement is not configured, an interface that stops receiving BPDUs will transition to the designated port (forwarding) state, creating a potential loop.                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <b>log</b> —The interface logs the fact that it has not received BPDUs during the timeout interval.<br><br><b>block</b> —The interface is blocked and the fact that the interface has not received BPDUs during the timeout interval is logged.                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Understanding Loop Protection for Spanning-Tree Instance Interfaces</i></li> <li>• <i>Configuring Loop Protection for a Spanning-Tree Instance Interface</i></li> <li>• <i>Example: Enabling Loop Protection for Spanning-Tree Protocols</i></li> <li>• <a href="#">rstp on page 2356</a></li> <li>• <a href="#">mstp on page 2349</a></li> <li>• <a href="#">vstp on page 2364</a></li> </ul> |

## bridge-priority (Spanning Trees)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | bridge-priority <i>priority</i> ;                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit protocols <b>mstp</b> ],<br>[edit protocols mstp <b>msti</b> <i>msti-id</i> ],<br>[edit protocols <b>rstp</b> ],<br>[edit protocols <b>vstp</b> vlan <i>vlan-id</i> ]                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.                                                                                                                                                                                                                                 |
| <b>Description</b>              | Configure the bridge priority. The bridge priority determines which bridge is elected as the root bridge. If two bridges have the same path cost to the root bridge, the bridge priority determines which bridge becomes the designated bridge for a LAN segment.                                                                                                                             |
| <b>Default</b>                  | 32,768                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <b>priority</b> —Bridge priority. It can be set only in increments of 4096.<br><b>Range:</b> 0 through 61,440<br><b>Default:</b> 32,768                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>show spanning-tree bridge</i></li><li>• <i>show spanning-tree interface</i></li><li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on page 2149</a></li><li>• <i>Understanding MSTP for EX Series and QFX Series Switches</i></li><li>• <i>Understanding VSTP for EX Series Switches and QFX Series Switches</i></li></ul> |

## configuration-name

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>configuration-name <i>configuration-name</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mstp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols mstp],<br>[edit protocols mstp],<br>[edit routing-instances <i>routing-instance-name</i> protocols mstp]                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.4.<br>Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Support for logical systems added in Junos OS Release 9.6.                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Specify the configuration name , which is the MSTP region name carried in the MSTP BPDUs.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Understanding BPDUs Used for Exchanging Information Among Bridges</i></li> <li>• <i>Configuring Multiple Spanning Tree Protocol</i></li> <li>• <a href="#">Configuring MSTP on page 2170</a></li> <li>• <i>show spanning-tree bridge</i></li> <li>• <i>show spanning-tree interface</i></li> <li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on page 2149</a></li> <li>• <a href="#">Example: Configuring Faster Convergence and Improved Network Stability with RSTP on page 2173</a></li> <li>• <i>Understanding MSTP for EX Series and QFX Series Switches</i></li> </ul> |

## cost

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>cost cost;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols (mstp   rstp   vstp) interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols mstp msti <i>msti-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols vstp <b>vlan</b> <i>vlan-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (mstp   rstp   vstp) interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols mstp msti <i>msti-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vstp <b>vlan</b> <i>vlan-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (mstp   rstp   vstp) interface <i>interface-name</i>],</p> <p>[edit protocols mstp msti <i>msti-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols vstp <b>vlan</b> <i>vlan-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (mstp   rstp   vstp) interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols mstp msti <i>msti-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vstp <b>vlan</b> <i>vlan-id</i> interface <i>interface-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.4.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Configure link cost to control which bridge is the designated bridge and which port is the designated port. By default, the link cost is determined by the link speed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <p><b>cost</b>—(Optional) Link cost associated with the port.</p> <p><b>Range:</b> 1 through 200,000,000</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Spanning-Tree Instance Interface</i></li> <li>• <i>Spanning-Tree Instance Interface Cost</i></li> <li>• <i>show spanning-tree bridge</i></li> <li>• <i>show spanning-tree interface</i></li> <li>• <i>Understanding RSTP for EX Series and QFX Series Switches</i></li> <li>• <i>Understanding MSTP for EX Series and QFX Series Switches</i></li> <li>• <i>Understanding VSTP for EX Series Switches and QFX Series Switches</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |



## disable (Spanning Trees)

|                            |                                                                                                                                                                                                                                                                                                                                                          |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | disable;                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>     | [edit protocols <b>mstp</b> interface <i>interface-name</i> ,<br>[edit protocols mstp <b>msti</b> <i>msti-id</i> vlan (all <i>vlan-id</i>   <i>vlan-name</i> ) interface <i>interface-name</i> ],<br>[edit protocols <b>rstp</b> interface <i>interface-name</i> ],<br>[edit protocols <b>vstp</b> vlan <i>vlan-id</i> interface <i>interface-name</i> ] |
| <b>Release Information</b> | Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.<br>Statement updated in Junos OS Release 15.1 for EX Series switches.                                                                                                                      |
| <b>Description</b>         | Disable MSTP, RSTP, or VSTP on a specific interface.                                                                                                                                                                                                                                                                                                     |



**NOTE:** You cannot disable spanning tree parameters for all interfaces.

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring RSTP (CLI Procedure) on page 2190</a></li> <li>• <a href="#">Configuring MSTP on page 2170</a></li> <li>• <a href="#">Configuring MSTP on page 2170</a></li> <li>• <a href="#">Configuring VLAN Spanning Tree Protocol</a></li> <li>• <a href="#">show spanning-tree bridge</a></li> <li>• <a href="#">show spanning-tree interface</a></li> </ul> |

## disable-timeout (Spanning Trees)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>disable-timeout <i>seconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit protocols layer2-control <a href="#">bpdu-block</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.1 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | For interfaces configured for BPDU protection, specify the amount of time an interface is disabled by BPDU blocking. If this option is not configured, the interface is not periodically checked and remains disabled.                                                                                                                                                                                                                                                                                                                                      |
| <b>Default</b>                  | The disable timeout is not enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <i>seconds</i> —Amount of time, in seconds, the interface receiving BPDUs protect is disabled. The range is 10 through 3600 seconds.                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>show spanning-tree bridge</i></li><li>• <i>show spanning-tree interface</i></li><li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on page 2149</a></li><li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on page 2149</a></li><li>• <a href="#">Example: Configuring BPDU Protection on Edge Interfaces to Prevent STP Miscalculations on page 2200</a></li><li>• <i>Understanding BPDU Protection for STP, RSTP, and MSTP on EX Series Switches</i></li></ul> |

## edge

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | edge;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols (mstp   rstp   vstp) interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols mstp msti <i>msti-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols vstp <b>vlan</b> <i>vlan-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (mstp   rstp   vstp) interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols mstp msti <i>msti-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vstp <b>vlan</b> <i>vlan-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (mstp   rstp   vstp) interface <i>interface-name</i>],</p> <p>[edit protocols mstp msti <i>msti-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols vstp <b>vlan</b> <i>vlan-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (mstp   rstp   vstp) interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols mstp msti <i>msti-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vstp <b>vlan</b> <i>vlan-id</i> interface <i>interface-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.4.</p> <p>Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Configure interfaces as edge ports. Edge ports do not expect to receive BPDUs. If a BPDU is received, the port becomes a nonedge port.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Spanning-Tree Instance Interface</i></li> <li>• <i>Configuring a Spanning-Tree Instance Interface as an Edge Port for Faster Convergence</i></li> <li>• <i>show spanning-tree bridge</i></li> <li>• <i>show spanning-tree interface</i></li> <li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on page 2149</a></li> <li>• <a href="#">Example: Configuring Faster Convergence and Improved Network Stability with RSTP on page 2173</a></li> <li>• <i>Understanding VSTP for EX Series Switches and QFX Series Switches</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

## force-version (IEEE 802.1D STP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | force-version stp;                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols (rstp   vstp)],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (rstp   vstp)],<br>[edit protocols (rstp   vstp)],<br>[edit routing-instances <i>routing-instance-name</i> protocols (rstp   vstp)]                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.4.<br>Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Support for logical systems added in Junos OS Release 9.6.                                                                                                                                                                                                |
| <b>Description</b>              | Force the spanning-tree version to be the original IEEE 803.1D STP.                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Spanning-Tree Protocols Supported</i></li><li>• <i>RSTP or VSTP Forced to Run as IEEE 802.1D STP</i></li><li>• <i>Reverting to RSTP or VSTP from Forced IEEE 802.1D STP</i></li><li>• <i>Reverting to RSTP or VSTP from Forced IEEE 802.1D STP</i></li><li>• <i>Understanding VSTP for EX Series Switches and QFX Series Switches</i></li></ul> |

## forward-delay

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>forward-delay seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols (mstp   rstp)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols vstp <i>vlan vlan-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (mstp   rstp)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vstp <i>vlan vlan-id</i>],</p> <p>[edit protocols (mstp   rstp)],</p> <p>[edit protocols vstp <i>vlan vlan-id</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (mstp   rstp)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vstp <i>vlan vlan-id</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.4.</p> <p>Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Specify the length of time an STP bridge port remains in the listening and learning states before transitioning to the forwarding state.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <p><b>seconds</b>—(Optional) Number of seconds the bridge port remains in the listening and learning states.</p> <p><b>Range:</b> 4 through 30</p> <p><b>Default:</b> 15 seconds</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Forward Delay Before Ports Transition to Forwarding State</i></li> <li>• <i>show spanning-tree bridge</i></li> <li>• <i>show spanning-tree interface</i></li> <li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on page 2149</a></li> <li>• <a href="#">Example: Configuring Faster Convergence and Improved Network Stability with RSTP on page 2173</a></li> <li>• <i>Understanding MSTP for EX Series and QFX Series Switches</i></li> <li>• <i>Understanding VSTP for EX Series Switches and QFX Series Switches</i></li> </ul>                                                                                                                                   |

## hello-time

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>hello-time seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols (mstp   rstp)],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols vstp <i>vlan</i> <i>vlan-id</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>(mstp   rstp)],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>vstp <i>vlan</i> <i>vlan-id</i>],</code><br><code>[edit protocols (mstp   rstp)],</code><br><code>[edit protocols vstp <i>vlan</i> <i>vlan-id</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols (mstp   rstp)],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols vstp <i>vlan</i> <i>vlan-id</i>]</code> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.4.<br>Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Support for logical systems added in Junos OS Release 9.6.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Specify the number of seconds between transmissions of configuration BPDUs by the root bridge.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <b>seconds</b> —(Optional) Number of seconds between transmissions of configuration BPDUs.<br><b>Range:</b> 1 through 10<br><b>Default:</b> 2 seconds                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Hello Time for Root Bridge to Transmit Hello BPDUs</i></li><li>• <i>show spanning-tree bridge</i></li><li>• <i>show spanning-tree interface</i></li><li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on page 2149</a></li><li>• <a href="#">Example: Configuring Faster Convergence and Improved Network Stability with RSTP on page 2173</a></li><li>• <i>Understanding MSTP for EX Series and QFX Series Switches</i></li><li>• <i>Understanding VSTP for EX Series Switches and QFX Series Switches</i></li></ul>                                                                                                                                                                                                                                                                                   |

## interface (BPDU)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>interface</code> (all   [ <i>interface-name</i> ]) {<br>drop;<br>}                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | <ul style="list-style-type: none"> <li>For platforms with ELS CLI:<br/>    [edit protocols layer2-control]</li> <li>For platforms with Original CLI:<br/>    [edit ethernet-switching-options]</li> </ul>                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Apply BPDU protection to all interfaces or one or more interfaces.                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <p><b>all</b>—All interfaces.</p> <p><b><i>interface-name</i></b>—Name of the interface.</p> <p><b>drop</b>—Drops xSTP BPDUs.</p>                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring Network Regions for VLANs with MSTP</i></li> <li>• <i>Example: Configuring Faster Convergence and Improving Network Stability with RSTP</i></li> <li>• <a href="#">Understanding BPDU Protection for STP, RSTP, and MSTP on page 2199</a></li> <li>• <a href="#">show spanning-tree bridge on page 2420</a></li> <li>• <a href="#">show spanning-tree interface on page 2425</a></li> </ul> |

## interface (Spanning Tree)

|                            |                                                                                                                                                                                                                                                |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>interface (<i>interface-name</i> disable   all){   bpd<del>u</del><del>-</del>timeout-action {     alarm;     block;   }   cost <i>cost</i>;   edge;   mode (p2p   shared);   no-root-port;   priority <i>interface-priority</i>; }</pre> |
| <b>Hierarchy Level</b>     | <pre>[edit protocols (<i>mstp</i>   <i>rstp</i>   <i>vstp</i>)], [edit protocols vstp vlan <i>vlan-id</i>], [edit protocols vstp <i>vlan-group</i> group <i>group-name</i> vlan (<i>vlan-id</i>   <i>vlan-range</i>  open-set-of-values)</pre> |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.</p> <p>Statement updated in Junos OS Release 15.1 for EX Series switches to support configuration of spanning tree parameters globally on all interfaces.</p>  |



**NOTE:** You cannot disable spanning tree parameters globally on all interfaces.

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b>              | Configure the interface to participate in the RSTP, MSTP, or VSTP instance.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <p><i>interface-name</i>—Name of a Gigabit Ethernet or 10-Gigabit Ethernet interface.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring RSTP (CLI Procedure) on page 2190</a></li> <li>• <a href="#">Configuring MSTP on page 2170</a></li> <li>• <a href="#">Configuring MSTP on page 2170</a></li> <li>• <a href="#">Configuring VLAN Spanning Tree Protocol</a></li> <li>• <a href="#">show spanning-tree interface</a></li> <li>• <a href="#">Understanding RSTP for EX Series and QFX Series Switches</a></li> <li>• <a href="#">Understanding MSTP for EX Series and QFX Series Switches</a></li> <li>• <a href="#">Understanding VSTP for EX Series Switches and QFX Series Switches</a></li> </ul> |



## max-age


|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>max-age seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols (mstp   rstp)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols vstp <i>vlan vlan-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (mstp   rstp)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vstp <i>vlan vlan-id</i>],</p> <p>[edit protocols (mstp   rstp)],</p> <p>[edit protocols vstp <i>vlan vlan-id</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (mstp   rstp)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vstp <i>vlan vlan-id</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.4.</p> <p>Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Specify the maximum expected arrival time of hello BPDUs.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <p><b>seconds</b>—(Optional) Number of seconds expected between hello BPDUs.</p> <p><b>Range:</b> 6 through 40</p> <p><b>Default:</b> 20 seconds</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Maximum Age for Awaiting Arrival of Hello BPDUs</i></li> <li>• <i>show spanning-tree bridge</i></li> <li>• <i>show spanning-tree interface</i></li> <li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on page 2149</a></li> <li>• <a href="#">Example: Configuring Faster Convergence and Improved Network Stability with RSTP on page 2173</a></li> <li>• <i>Understanding MSTP for EX Series and QFX Series Switches</i></li> <li>• <i>Understanding VSTP for EX Series Switches and QFX Series Switches</i></li> </ul>                                                                                                                                             |

## max-hops

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>max-hops hops;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mstp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols mstp],<br>[edit protocols mstp],<br>[edit routing-instances <i>routing-instance-name</i> protocols mstp]                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.4.<br>Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Support for logical systems added in Junos OS Release 9.6.                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Configure the maximum number of hops a BPDU can be forwarded in the MSTP region.                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <b>hops</b> —(Optional) Number of hops the BPDU can be forwarded.<br><b>Range:</b> 1 through 255<br><b>Default:</b> 19 hops                                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Multiple Spanning Tree Protocol</a></li><li>• <a href="#">Configuring MSTP on page 2170</a></li><li>• <code>show spanning-tree bridge</code></li><li>• <code>show spanning-tree interface</code></li><li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on page 2149</a></li><li>• <a href="#">Understanding MSTP for EX Series and QFX Series Switches</a></li></ul> |

## mode

|                                                                                                               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                 | <code>mode (p2p   shared);</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>                                                                                        | <p>[edit logical-systems <i>logical-system-name</i> protocols (mstp   rstp   vstp) interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols vstp <b>vlan</b> <i>vlan-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (mstp   rstp   vstp) interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vstp <b>vlan</b> <i>vlan-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (mstp   rstp   vstp) interface <i>interface-name</i>],</p> <p>[edit protocols vstp <b>vlan</b> <i>vlan-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (mstp   rstp   vstp) interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vstp <b>vlan</b> <i>vlan-id</i> interface <i>interface-name</i>]</p> |
| <b>Release Information</b>                                                                                    | <p>Statement introduced in Junos OS Release 8.4.</p> <p>Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>                                                                                            | Configure link mode to identify point-to-point links.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Default</b>                                                                                                | When the link is configured as full-duplex, the default link mode is <b>p2p</b> . When the link is configured half-duplex, the default link mode is <b>shared</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <div>  <b>NOTE:</b> </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                                                                                                | <p><b>p2p</b>—The link is point to point.</p> <p><b>shared</b>—The link is shared media.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b>                                                                               | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>                                                                                  | <ul style="list-style-type: none"> <li>• <i>Spanning-Tree Instance Interface</i></li> <li>• <i>Spanning-Tree Instance Interface Point-to-Point Link Mode</i></li> <li>• <i>show spanning-tree bridge</i></li> <li>• <i>show spanning-tree interface</i></li> <li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on page 2149</a></li> <li>• <a href="#">Example: Configuring Faster Convergence and Improved Network Stability with RSTP on page 2173</a></li> <li>• <i>Understanding VSTP for EX Series Switches and QFX Series Switches</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

## msti

**Syntax** `msti msti-id {  
     bridge-priority priority;  
     vlan (vlan-id | vlan-range|open-set-of-values);  
     interface (interface-name | all) {  
         cost cost;  
         edge;  
         priority interface-priority;  
     }  
 }`

**Hierarchy Level** [edit protocols [mstp](#)],  
 [edit routing-instances *routing-instance-name* protocols mstp]

**Release Information** Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.  
 Statement updated in Junos OS Release 15.1 for EX Series switches to support configuration of spanning tree parameters globally on all interfaces.



**NOTE:** You cannot disable spanning tree parameters globally on all interfaces.

**Description** Configure the multiple spanning-tree instance (MSTI) identifier.

**Options** *msti-id*—MSTI instance identifier.  
**Range:** 1 through 4094

The remaining statements are explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
 routing-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring MST Instances on a Physical Interface](#)
- [Configuring MSTP on page 2170](#)
- [Example: Configuring Network Regions for VLANs with MSTP on page 2149](#)
- [Understanding MSTP for EX Series and QFX Series Switches](#)

## mstp

```
Syntax mstp {
 bpdu-block-on-edge;
 bridge-priority priority;
 configuration-name configuration-name;
 disable;
 forward-delay seconds;
 hello-time seconds;
 interface (interface-name disable | all){
 bpdu-timeout-action {
 alarm;
 block;
 }
 cost cost;
 edge;
 mode (p2p | shared);
 no-root-port;
 priority interface-priority;
 }
 max-age seconds;
 max-hops hops;
 priority-hold-time seconds;
 revision-level revision-level;
 interface (interface-name disable | all){
 bpdu-timeout-action {
 alarm;
 block;
 }
 cost cost;
 disable;
 edge;
 mode (p2p | shared);
 no-root-port;
 priority interface-priority;
 }
 msti msti-id {
 bridge-priority priority;
 interface (interface-name disable | all){
 cost cost;
 edge;
 priority interface-priority;
 }
 vlan (vlan-id | vlan-range | open-set-of-values);
 interface all
 }
 }
 traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
 }
}
```

Hierarchy Level [edit protocols]

**Release Information** Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.  
Statement updated in Junos OS Release 15.1 for EX Series switches to support configuration of spanning tree parameters globally on all interfaces.



**NOTE:** You cannot disable spanning tree parameters globally on all interfaces.

**Description** Configure MSTP parameters.

**Options** The statements are explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring MSTP on page 2170](#)
- [Example: Configuring Network Regions for VLANs with MSTP on page 2149](#)

## no-root-port

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-root-port;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols (mstp   rstp   vstp) interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols vstp <b>vlan</b> <i>vlan-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (mstp   rstp   vstp) interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vstp <b>vlan</b> <i>vlan-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (mstp   rstp   vstp) interface <i>interface-name</i>],</p> <p>[edit protocols vstp <b>vlan</b> <i>vlan-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (mstp   rstp   vstp) interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vstp <b>vlan</b> <i>vlan-id</i> interface <i>interface-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Ensure the port is the spanning-tree designated port. If the port receives superior bridge protocol data unit (BPDU) packets, root protect moves this port to a root-prevented spanning-tree state.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Understanding Root Protection for Spanning-Tree Instance Interfaces in a Layer 2 Switched Network</i></li> <li>• <i>Root Protect for a Spanning-Tree Instance Interface</i></li> <li>• <i>Enabling Root Protection for a Spanning-Tree Instance Interface</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## priority (Protocols STP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>priority interface-priority;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols (mstp   rstp   vstp) interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols mstp msti <i>msti-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols vstp <b>vlan</b> <i>vlan-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (mstp   rstp   vstp) interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols mstp msti <i>msti-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vstp <b>vlan</b> <i>vlan-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (mstp   rstp   vstp) interface <i>interface-name</i>],</p> <p>[edit protocols mstp msti <i>msti-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols vstp <i>vlan-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (mstp   rstp   vstp) interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols mstp msti <i>msti-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vstp <b>vlan</b> <i>vlan-id</i> interface <i>interface-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.4.</p> <p>Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Use the interface priority to control which interface is elected as the root port. The interface priority must be set in increments of 16.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <p><b>priority</b>—(Optional) Interface priority.</p> <p><b>Range:</b> 0 through 240</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Spanning-Tree Instance Interface</i></li> <li>• <i>Configuring a Spanning-Tree Instance Interface as an Edge Port for Faster Convergence</i></li> <li>• <i>Spanning-Tree Instance Interface Priority</i></li> <li>• <i>[edit protocols mstp] Configuration Statement Hierarchy on EX Series Switches</i></li> <li>• <i>[edit protocols rstp] Configuration Statement Hierarchy on EX Series and QFX Series Switches</i></li> <li>• <i>[edit protocols vstp] Configuration Statement Hierarchy on EX Series and QFX Series Switches</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |



## protocol

|                                 |                                                                                                                                                                                                                                                                                                   |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>protocol (cdp   stp   vtp   pvstp);</code>                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit protocols layer2-control mac-rewrite interface <i>interface-name</i> ]                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.1.<br>Statement introduced in Junos OS Release 13.2 for QFX series.<br>Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Support for PVSTP introduced in Junos OS Release 13.3.                                          |
| <b>Description</b>              | Configure the protocol to be tunneled on an interface for Layer 2 protocol tunneling. To tunnel multiple protocols, include multiple <b>protocol</b> statements.                                                                                                                                  |
| <b>Options</b>                  | <p><b>cdp</b>—Tunnel the Cisco discovery protocol.</p> <p><b>stp</b>—Tunnel all versions of the spanning-tree protocol.</p> <p><b>vtp</b>—Tunnel the VLAN trunk protocol.</p> <p><b>pvstp</b>—Tunnel the Per-VLAN Spanning Tree Plus (PVST+) protocol</p>                                         |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Layer 2 Protocol Tunneling Through a Network Overview</i></li> <li>• <i>Layer 2 Protocol Tunneling Configuration Guidelines</i></li> <li>• <i>Layer 2 Protocol to be Tunneled</i></li> <li>• <i>Configuring Layer 2 Protocol Tunneling</i></li> </ul> |

## protocols (STP Type)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>protocols {<br/>  mstp { ... }<br/>  rstp { ... }<br/>  vstp { ... }<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit],<br>[edit logical-systems <i>logical-system-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> ]                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.4.<br>Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Support for logical systems added in Junos OS Release 9.6.                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Configure the Spanning Tree Protocol type as MSTP, RSTP, or VSTP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <b>mstp</b> —Configure the protocol as Multiple Spanning Tree.<br><br><b>rstp</b> —Configure the protocol as Rapid Spanning Tree.<br><br><b>vstp</b> —Configure the protocol as VLAN Spanning Tree.                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring RSTP (CLI Procedure) on page 2190</a></li><li>• <a href="#">Configuring MSTP on page 2170</a></li><li>• <i>Configuring MST Instances on a Physical Interface</i></li><li>• <i>Configuring VLAN Spanning Tree Protocol</i></li><li>• <i>Configuring Rapid Spanning Tree Protocol</i></li><li>• <i>Configuring Multiple Spanning Tree Protocol</i></li><li>• <i>Configuring VLAN Spanning Tree Protocol</i></li><li>• <i>Understanding MSTP for EX Series and QFX Series Switches</i></li></ul> |

## revision-level

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>revision-level <i>revision-level</i>;</code>                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mstp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols mstp],<br>[edit protocols mstp],<br>[edit routing-instances <i>routing-instance-name</i> protocols mstp]                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.4.<br>Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Support for logical systems added in Junos OS Release 9.6.                                                                                                                                                                                   |
| <b>Description</b>              | Set the revision number of the MSTP configuration.                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <i>revision-level</i> —Configure the revision number of the MSTP region configuration.<br><b>Range:</b> 0 through 65,535                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Multiple Spanning Tree Protocol</i></li> <li>• <i>show spanning-tree bridge</i></li> <li>• <i>show spanning-tree interface</i></li> <li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on page 2149</a></li> <li>• <i>Understanding MSTP for EX Series and QFX Series Switches</i></li> </ul> |

## rstp

```
Syntax rstp {
 bpdv-block-on-edge;
 bpdv-destination-mac-address provider-bridge-group;
 bridge-priority priority;
 disable;
 extended-system-id;
 force-version stp;
 forward-delay seconds;
 hello-time seconds;
 interface (interface-name disable | interface-range-name | all){
 bpdv-timeout-action {
 alarm;
 block;
 }
 cost cost;
 edge;
 mode (p2p | shared);
 no-root-port;
 priority interface-priority;
 }
 max-age seconds;
 priority-hold-time seconds;
 traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
 }
}
```

**Hierarchy Level** [edit protocols]

**Release Information** Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.  
Statement updated in Junos OS Release 15.1 for EX Series and QFX Series switches to support configuration of spanning tree parameters globally on all interfaces.



**NOTE:** You cannot disable spanning tree parameters globally on all interfaces.

**Description** Configure RSTP parameters.

**Options** The remaining statements are explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring RSTP \(CLI Procedure\) on page 2190](#)

## traceoptions (Spanning Tree)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;flag-modifier&gt; &lt;disable&gt;; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols (mstp   rstp   vstp)],<br/>         [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br/>         (mstp   rstp   vstp)],<br/>         [edit protocols (mstp   rstp   vstp   vstp vlan <i>vlan-id</i>)],<br/>         [edit routing-instances <i>routing-instance-name</i> protocols (mstp   rstp   vstp)]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 8.4.<br/>         Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br/>         Support for logical systems added in Junos OS Release 9.6.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>         | Set protocol-level tracing options for spanning-tree protocols.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Default</b>             | The default STP protocol-level trace options are inherited from the global <b>traceoptions</b> statement.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>             | <p><b>disable</b>—(Optional) Disable the tracing operation. One use of this option is to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. We recommend that you place STP tracing output in the file <code>/var/log/stp-log</code>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you must also specify a maximum file size with the <b>size</b> option.</p> <p><b>Range:</b> 2 through 1000 files<br/> <b>Default:</b> 1 trace file only</p> <p><b>flag</b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements. The following are the STP-specific tracing options:</p> <ul style="list-style-type: none"> <li><b>all</b>—Trace all operations.</li> <li><b>all-failures</b>—Trace all failure conditions.</li> <li><b>bpdv</b>—Trace BPDU reception and transmission.</li> <li><b>bridge-detection-state-machine</b>—Trace the bridge detection state machine.</li> </ul> |

- **events**—Trace events of the protocol state machine.
- **port-information-state-machine**—Trace the port information state machine.
- **port-migration-state-machine**—Trace the port migration state machine.
- **port-receive-state-machine**—Trace the port receive state machine.
- **port-role-transit-state-machine**—Trace the port role transit state machine.
- **port-role-select-state-machine**—Trace the port role selection state machine.
- **port-state-transit-state-machine**—Trace the port state transit state machine.
- **port-transmit-state-machine**—Trace the port transmit state machine.
- **ppmd**—Trace the state and events for the ppm process.
- **state-machine-variables**—Trace when the state machine variables change.
- **timers**—Trace protocol timers.
- **topology-change-state-machine**—Trace the topology change state machine.

The following are the global tracing options:

- **all**—All tracing operations.
- **config-internal**—Trace configuration internals.
- **general**—Trace general events.
- **normal**—All normal events.

**Default:** If you do not specify this option, only unusual or abnormal operations are traced.

- **parse**—Trace configuration parsing.
- **policy**—Trace policy operations and actions.
- **regex-parse**—Trace regular-expression parsing.
- **route**—Trace routing table changes.
- **state**—Trace state transitions.
- **task**—Trace protocol task processing.
- **timer**—Trace protocol task timer processing.

**no-world-readable**—(Optional) Prevent any user from reading the log file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB) or megabytes (MB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 1 MB

**world-readable**—(Optional) Allow any user to read the log file.

|                                 |                                                             |
|---------------------------------|-------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.        |
|                                 | routing-control—To add this statement to the configuration. |
| <b>Related Documentation</b>    | • <i>Spanning-Tree Protocol Trace Options</i>               |
|                                 | • <i>Tracing Spanning-Tree Operations</i>                   |
|                                 | • <i>Example: Tracing Spanning-Tree Protocol Operations</i> |

## vlan (MSTP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>vlan <i>vlan-id</i>;</code>                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mstp msti <i>msti-id</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols mstp msti <i>msti-id</i> ],<br>[edit protocols mstp msti <i>msti-id</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols mstp msti <i>msti-id</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.4.<br>Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Support for logical systems added in Junos OS Release 9.6.                                                                                                                                                                                 |
| <b>Description</b>              | Configure the VLAN of an MSTI or VSTP instance or configure the VLAN range of an MSTI.                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <b><i>vlan-id</i></b> —The VLAN identifier associated with the MSTI.<br><br><b><i>vlan-id-range</i></b> —Range of VLAN identifiers associated with the MSTI in the form <i>minimum-vlan-id-maximum-vlan-id</i> . VLAN identifier ranges are not supported for VSTP.<br><br><b>Range:</b> 1 through 4096                                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Multiple Spanning Tree Protocol</i></li><li>• <i>Understanding VSTP for EX Series Switches and QFX Series Switches</i></li></ul>                                                                                                                                                                                     |



## vlan (VSTP)

|                                 |                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> vlan <i>vlan-id</i> {     bridge-priority <i>priority</i>;     forward-delay <i>seconds</i>;     hello-time <i>seconds</i>;     max-age <i>seconds</i>;     interface <i>interface-name</i> {         cost <i>cost</i>;         edge;         mode (p2p   shared);         no-root-port;         priority <i>interface-priority</i>;     } } </pre> |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols vstp],<br>[edit protocols vstp]                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0.<br>Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Support for logical systems added in Junos OS Release 9.6.                                                                                                                                                               |
| <b>Description</b>              | Configure VSTP VLAN parameters.                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | The statements are explained separately.                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>Configuring VLAN Spanning Tree Protocol</li> <li>Understanding VSTP for EX Series Switches and QFX Series Switches</li> </ul>                                                                                                                                                                                      |

## vlan (VSTP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>vlan <i>vlan-id</i> {<br/>    bridge-priority <i>priority</i>;<br/>    forward-delay <i>seconds</i>;<br/>    hello-time <i>seconds</i>;<br/>    max-age <i>seconds</i>;<br/>    interface <i>interface-name</i> {<br/>        cost <i>cost</i>;<br/>        edge;<br/>        mode (p2p   shared);<br/>        no-root-port;<br/>        priority <i>interface-priority</i>;<br/>    }<br/>}</pre> |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols vstp],<br>[edit protocols vstp]                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0.<br>Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Support for logical systems added in Junos OS Release 9.6.                                                                                                                                                                                                             |
| <b>Description</b>              | Configure VSTP VLAN parameters.                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | The statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring VSTP (CLI Procedure) on page 2197</a></li><li>• <a href="#">Configuring VLAN Spanning Tree Protocol</a></li><li>• <a href="#">Understanding VSTP for EX Series Switches and QFX Series Switches</a></li></ul>                                                                                                                           |

## vlan-group

**Syntax** `vlan-group group group-name {  
           vlan (vlan-id |vlan-group |all) {  
           }`

**Hierarchy Level**           [edit protocols **vstp**]

**Release Information**   Statement introduced in Junos OS Release 15.1 for EX Series switches.

**Description**           Configure VLAN group for Spanning Tree Protocol (VSTP). VSTP is used to prevent loops in Layer 2 networks on a per-VLAN basis.



**BEST PRACTICE:** Configure RSTP when you configure VSTP. RSTP overhead is minimal and this configuration ensures that a spanning-tree protocol is running on all VLANs on your switch, even when your switch is supporting more than the maximum number of allowed VSTP VLANs.

The remaining statements are explained separately.

**Required Privilege Level**   routing—To view this statement in the configuration.  
                                   routing-control—To add this statement to the configuration.

**Related Documentation**

- [vstp on page 2364](#)
- *show spanning-tree bridge*
- *show spanning-tree interface*
- *Configuring VLAN Spanning Tree Protocol*
- *Understanding VSTP for EX Series Switches and QFX Series Switches*

## vstp

```
Syntax vstp {
 bpd-block-on-edge;
 disable;
 force-version stp;
 interface (interface-name disable | interface-range-name | all){
 bpd-timeout-action {
 alarm;
 block;
 }
 cost cost;
 edge;
 mode (p2p | shared);
 no-root-port;
 priority interface-priority;
 }
 priority-hold-time seconds;
 vlan (vlan-id | all){
 bridge-priority priority;
 forward-delay seconds;
 hello-time seconds;
 max-age seconds;
 interface (interface-name disable | interface-range-name | all){
 bpd-timeout-action {
 alarm;
 block;
 }
 cost cost;
 edge;
 mode (p2p | shared);
 no-root-port;
 priority interface-priority;
 }
 }
 traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
 }
 vlan-group group group-name {
 vlans vlan-name (vlan-id | vlan-range | open-set-of-values) {
 interface all;
 interface interface-name {
 disable;
 }
 }
 }
}
```

Hierarchy Level [edit protocols]

**Release Information** Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches. Statement updated in Junos OS Release 15.1 for EX Series switches to support configuration of spanning tree parameters globally on all interfaces.



NOTE: You cannot disable spanning tree parameters globally on all interfaces.

|                                 |                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------|
| <b>Description</b>              | Configure VSTP parameters.                                                                                          |
| <b>Options</b>                  | The statements are explained separately.                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration. |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring VSTP (CLI Procedure) on page 2197</a></li></ul>     |



## CHAPTER 83

# Q-in-Q Configuration Statements

- [flexible-vlan-tagging on page 2368](#)
- [input-vlan-map on page 2369](#)
- [native-vlan-id on page 2370](#)
- [output-vlan-map \(Gigabit Ethernet IQ and 10-Gigabit Ethernet with SFPP\) on page 2371](#)
- [pop on page 2372](#)
- [push on page 2373](#)
- [swap on page 2374](#)
- [vlan-id-list on page 2375](#)


## flexible-vlan-tagging

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | flexible-vlan-tagging;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit interfaces aex],<br>[edit interfaces ge- <i>fpc/pic/port</i> ],<br>[edit interfaces et- <i>fpc/pic/port</i> ],<br>[edit interfaces ps0],<br>[edit interfaces xe- <i>fpc/pic/port</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.1.<br>Support for aggregated Ethernet added in Junos OS Release 9.0.<br>Statement introduced in Junos OS Release 12.1x48 for PTX Series Packet Transport Routers.<br>Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | <p>Support simultaneous transmission of 802.1Q VLAN single-tag and dual-tag frames on logical interfaces on the same Ethernet port, and on pseudowire logical interfaces.</p> <p>This statement is supported on M Series and T Series routers, for Fast Ethernet and Gigabit Ethernet interfaces only on Gigabit Ethernet IQ2 and IQ2-E, IQ, and IQE PICs, and for aggregated Ethernet interfaces with member links in IQ2, IQ2-E, and IQ PICs or in MX Series DPCs, or on Ethernet interfaces for PTX Series Packet Transport Routers or 100-Gigabit Ethernet Type 5 PIC with CFP. This statement is supported on Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces on EX Series switches.</p> |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Mixed Tagging</i></li><li>• <i>Configuring Flexible VLAN Tagging on PTX Series Packet Transport Routers</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |




## input-vlan-map

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |                                                                                                                                                                                                                                                                                                                                                                              |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <pre>input-vlan-map {   (pop   pop-pop   pop-swap   push   push-push   swap   swap-push   swap-swap);   inner-tag-protocol-id <i>tpid</i>;   inner-vlan-id <i>number</i>;   tag-protocol-id <i>tpid</i>;   vlan-id <i>number</i>; }</pre>                                                                                                                                    |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <pre>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]</pre>                                                                                                                                                                       |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <p>Statement introduced before Junos OS Release 7.4.</p> <p><b>pop-pop</b>, <b>pop-swap</b>, <b>push-push</b>, <b>swap-push</b>, and <b>swap-swap</b> statements introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.</p> |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <p>For Gigabit Ethernet IQ, 10-Gigabit Ethernet SFPP interfaces, 100-Gigabit Ethernet Type 5 PIC with CFP only as well as Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces, define the rewrite profile to be applied to incoming frames on this logical interface.</p> <p>The statements are explained separately.</p>         |
| <div style="display: flex; align-items: flex-start;"> <div style="margin-right: 10px;">  </div> <div> <p><b>NOTE:</b> Connectivity fault management (CFM) sessions for all interfaces in which <b>input-vlan-map</b> is configured are supported only if the interface also has an explicit configuration for <b>output-vlan-map</b> as <b>output-vlan-map pop</b>. See <a href="#">output-vlan-map</a>. This configuration is required for all the interfaces in the topology even when the CFM session is on that interface or on a different interface in the data path of the same topology.</p> </div> </div> |                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <ul style="list-style-type: none"> <li>• <i>Stacking a VLAN Tag</i></li> <li>• <a href="#">output-vlan-map on page 2371</a></li> <li>• <i>Configuring Q-in-Q Tunneling (CLI Procedure)</i></li> </ul>                                                                                                                                                                        |

## native-vlan-id

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>native-vlan-id <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>     | [edit interfaces <i>ge-fpc/pic/port</i> ],<br>[edit interfaces <i>interface-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b> | Statement introduced in Junos OS Release 8.3.<br>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.<br>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>         | <p>Configure mixed tagging support for untagged packets on a port for the following:</p> <ul style="list-style-type: none"><li>• M Series routers with Gigabit Ethernet IQ PICs with SFP and Gigabit Ethernet IQ2 PICs with SFP configured for 802.1Q flexible VLAN tagging</li><li>• MX Series routers with Gigabit Ethernet DPCs and MICs, Tri-Rate Ethernet DPCs and MICs, and 10-Gigabit Ethernet DPCs and MICs and MPCs configured for 802.1Q flexible VLAN tagging</li><li>• T4000 routers with 100-Gigabit Ethernet Type 5 PIC with CFP</li><li>• EX Series switches with Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces</li></ul> <p>When the <b>native-vlan-id</b> statement is included with the <b>flexible-vlan-tagging</b> statement, untagged packets are accepted on the same mixed VLAN-tagged port.</p> <div> <b>NOTE:</b> The logical interface on which untagged packets are received must be configured with the same VLAN ID as the native VLAN ID configured on the physical interface, otherwise the untagged packets are dropped.</div> <p>To configure the logical interface, include the <b>vlan-id</b> statement (matching the <b>native-vlan-id</b> statement on the physical interface) at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] hierarchy level.</p> <p>When the <b>native-vlan-id</b> statement is included with the <b>interface-mode</b> statement, untagged packets are accepted and forwarded within the bridge domain or VLAN that is configured with the matching VLAN ID.</p> |
| <b>Default</b>             | By default, the untagged packets are dropped. That is, if you do not configure the <b>native-vlan-id</b> option, the untagged packets are dropped.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>             | <b><i>number</i></b> —VLAN ID number.<br><b>Range:</b> (ACX Series routers and EX Series switches) 0 through 4094.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- *Configuring Mixed Tagging Support for Untagged Packets*
- *Configuring Access Mode on a Logical Interface*
- [Configuring the Native VLAN Identifier \(CLI Procedure\) on page 2121](#)
- *Understanding Bridging and VLANs on EX Series Switches*
- [flexible-vlan-tagging on page 2368](#)
- *Understanding Q-in-Q Tunneling on EX Series Switches*

## output-vlan-map (Gigabit Ethernet IQ and 10-Gigabit Ethernet with SFPP)

**Syntax** `output-vlan-map {  
    (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);  
    inner-tag-protocol-id tpid;  
    inner-vlan-id number;  
    tag-protocol-id tpid;  
    vlan-id number;  
}`

**Hierarchy Level** [edit interfaces *interface-name* unit *logical-unit-number*],  
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

**Release Information** Statement introduced before Junos OS Release 7.4.  
**pop-pop**, **pop-swap**, **push-push**, **swap-push**, and **swap-swap** statements added in Junos OS Release 8.1.

**Description** For Gigabit Ethernet IQ and 10-Port 10-Gigabit Ethernet SFPP interfaces only, define the rewrite operation to be applied to outgoing frames on this logical interface.

The statements are explained separately.


**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**


- *Stacking and Rewriting Gigabit Ethernet VLAN Tags*
- [input-vlan-map on page 2369](#)

## pop

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | pop;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>input-vlan-map</b> ],<br>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>output-vlan-map</b> ],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>input-vlan-map</b> ],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>output-vlan-map</b> ]                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | <div> <b>NOTE:</b> On EX4300 switches, <b>pop</b> is not supported at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>input-vlan-map</b>] hierarchy level.</div> <p>For Gigabit Ethernet IQ, 10-Gigabit Ethernet IQ2, and IQ2-E interfaces; 10-Gigabit Ethernet LAN/WAN PIC; aggregated Ethernet interfaces using Gigabit Ethernet IQ interfaces; 100-Gigabit Ethernet Type 5 PIC with CFP; and Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces, specify the VLAN rewrite operation to remove a VLAN tag from the top of the VLAN tag stack. The outer VLAN tag of the frame is removed.</p> |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Removing a VLAN Tag</i></li><li>• <i>Configuring Q-in-Q Tunneling (CLI Procedure)</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## push



|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>push;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">input-vlan-map</a>],</code><br><code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">output-vlan-map</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i></code><br><code>  <a href="#">input-vlan-map</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i></code><br><code>  output-vlan-map]</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | <p> <b>NOTE:</b> On EX4300 switches, <code>push</code> is not supported at the <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map]</code> hierarchy level.</p> <p>Specify the VLAN rewrite operation to add a new VLAN tag to the top of the VLAN stack. An outer VLAN tag is pushed in front of the existing VLAN tag.</p> <p>You can use this statement on Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces; 10-Gigabit Ethernet LAN/WAN PIC; aggregated Ethernet interfaces using Gigabit Ethernet IQ interfaces; 100-Gigabit Ethernet Type 5 PIC with CFP; and Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces.</p> <p>If you include the <b>push</b> statement in the configuration, you must also include the <a href="#">pop</a> statement at the <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map]</code> hierarchy level.</p> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Stacking a VLAN Tag</i></li> <li>• <i>Configuring Q-in-Q Tunneling (CLI Procedure)</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## swap

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | swap;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>input-vlan-map</b> ],<br>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>output-vlan-map</b> ],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>input-vlan-map</b> ],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>output-vlan-map</b> ]                                                                                                  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | <p>Specify the VLAN rewrite operation to replace a VLAN tag. The outer VLAN tag of the frame is overwritten with the user-specified VLAN tag information.</p> <p>On MX Series routers, you can enter this statement on Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, aggregated Ethernet using Gigabit Ethernet IQ interfaces, and 100-Gigabit Ethernet Type 5 PIC with CFP. On EX Series switches, you can enter this statement on Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces.</p> |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Rewriting the VLAN Tag on Tagged Frames</i></li><li>• <i>Configuring Q-in-Q Tunneling (CLI Procedure)</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                   |

## vlan-id-list

|                                                                                                                                                                                                                                             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                               | <code>vlan-id-list [ <i>vlan-id-numbers</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                      | <p>[edit bridge-domains <i>bridge-domain-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i>],</p> <p>[edit interfaces <i>interface-name</i> unit 0],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>],</p> <p>[edit vlans <i>vlan-name</i>]</p>                                                                                                                                                                                                                                       |
| <b>Release Information</b>                                                                                                                                                                                                                  | <p>Statement introduced in Junos OS Release 9.4.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>                                                                                                                                                                                                                          | <p>Specify a VLAN identifier list to use for a bridge domain or VLAN in trunk mode.</p> <p>Specify the <b>trunk</b> option in the <b>interface-mode</b> statement to accept packets with a VLAN ID that matches the list of VLAN IDs specified in the <b>vlan-id-list</b> statement to forward the packet within the bridge domain or VLAN configured with the matching VLAN ID. Specify the <b>access</b> option to accept packets with no VLAN ID to forward the packet within the bridge domain or VLAN configured with the VLAN ID that matches the VLAN ID specified in the <b>vlan-id</b> statement.</p> <p>This statement also enables you to bind a logical interface to a list of VLAN IDs, thereby configuring the logical interface to receive and forward a frame with a tag that matches the specified VLAN ID list.</p> |
| <div>  <p><b>WARNING:</b> On some EX and QFX Series switches, you can apply no more than eight VLAN identifier lists to a physical interface.</p> </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                                                                                                                                                                                                                              | <p><b><i>vlan-id-numbers</i></b>—Valid VLAN identifiers. You can combine individual numbers with range lists by including a hyphen.</p> <p><b>Range:</b> 0 through 4095</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <div>  <p><b>NOTE:</b> On EX Series switches and the QFX Series, the range is 0 through 4094.</p> </div>                                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b>                                                                                                                                                                                                             | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

- Related Documentation**
- *Configuring a Bridge Domain*
  - *Configuring a VLAN*
  - *Configuring VLAN Identifiers for Bridge Domains and VPLS Routing Instances*
  - *Configuring VLAN Identifiers for VLANs and VPLS Routing Instances*



# Reflective Relay Configuration Statement

- [reflective-relay on page 2377](#)

## reflective-relay

---

|                                 |                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | reflective-relay;                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> family ethernet-switching]                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D35 for the EX Series.                              |
| <b>Description</b>              | Configure a switch interface to return packets back to a device on the same interface that was used to deliver the packets.                                               |
| <b>Default</b>                  | Switch interfaces are not configured for reflective relay.                                                                                                                |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Reflective Relay for Use with VEPA Technology</i></li><li>• <i>Configuring Reflective Relay</i></li></ul> |




## CHAPTER 85

# Bridging and VLANs Monitoring Commands

- `clear ethernet-switching table`
- `show ethernet-switching interfaces`
- `show ethernet-switching table`
- `show system statistics arp`
- `show vlans`

## clear ethernet-switching table

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | clear ethernet-switching table<br><interface <i>interface-name</i> ><br><mac <i>mac-address</i> ><br><management-vlan><br><persistent-mac < <i>interface</i>   <i>mac-address</i> >><br><vlan <i>vlan-name</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Syntax (QFX Series)</b> | clear ethernet-switching table<br><interface <i>interface-name</i> ><br><mac <i>mac-address</i> ><br><persistent-mac < <i>interface</i>   <i>mac-address</i> >><br><vlan <i>vlan-name</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b> | Command introduced in Junos OS Release 9.3 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>         | <div>  <p><b>NOTE:</b> On a QFabric system, using this command on an FCoE-enabled VLAN when FCoE sessions are active can cause traffic flooding and FCoE traffic drop. The FCoE sessions are not terminated and the traffic reconverges after a short period of time.</p> </div> <p>Clear learned entries, which are media access control (MAC) addresses, in the Ethernet switching table (also called the forwarding database table).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>             | <p><b>none</b>—Clear learned entries in the Ethernet switching table, except for persistent MAC addresses.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Clear all learned MAC addresses for the specified interface from the Ethernet switching table.</p> <p><b>mac <i>mac-address</i></b>—(Optional) Clear the specified learned MAC address from the Ethernet switching table.</p> <p><b>management-vlan</b>—(Optional) Clear all MAC addresses learned for the management VLAN from the Ethernet switching table. Note that you do not specify a VLAN name because only one management VLAN exists.</p> <p><b>persistent-mac &lt;<i>interface</i>   <i>mac-address</i>&gt;</b>—(Optional) Clear all MAC addresses, including persistent MAC addresses. Use the <b>interface</b> option to clear all MAC addresses on an interface, or use the <b>mac-address</b> option to clear all entries for a specific MAC address.</p> <p>Use this command whenever you move a device in your network that has a persistent MAC address on the switch. If you move the device to another port on the switch and do not clear the persistent MAC address from the original port it was learned on, then the new port will not learn the MAC address and the device will not be able to connect. If the original port is down when you move the device, then the new port</p> |

will learn the MAC address and the device can connect—however, unless you cleared the MAC address on the original port, when the port comes back up, the system reinstalls the persistent MAC address in the forwarding table for that port. If this occurs, the address is removed from the new port and the device loses connectivity.

**vlan *vlan-name***—(Optional) Clear all MAC addresses learned for the specified VLAN from the Ethernet switching table.

**Required Privilege Level**

view

**Related Documentation**

- *show ethernet-switching table*
- [show ethernet-switching table on page 2386](#)
- *Verifying That Persistent MAC Learning Is Working Correctly*

**List of Sample Output**    [clear ethernet-switching table on page 2381](#)

**Output Fields**    This command produces no output.

## Sample Output

[clear ethernet-switching table](#)

```
user@switch> clear ethernet-switching table
```

## show ethernet-switching interfaces

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show ethernet-switching interfaces<br><brief   detail   summary><br><interface <i>interface-name</i> >                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Display information about switched Ethernet interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <p><b>none</b>—(Optional) Display brief information for Ethernet-switching interfaces.</p> <p><b>brief   detail   summary</b>—(Optional) Display the specified level of output.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display Ethernet-switching information for a specific interface.</p>                                                                                                                                                                               |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Troubleshooting Ethernet Switching on page 2134</a><a href="#">Understanding Bridging and VLANs on page 2089</a></li> <li>• <a href="#">Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 2096</a></li> <li>• <a href="#">Example: Setting Up Bridging with Multiple VLANs</a></li> <li>• <a href="#">Understanding FCoE on page 6436</a></li> <li>• <a href="#">Interfaces Overview on page 2785</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show ethernet-switching interfaces on page 2383</a><br><a href="#">show ethernet-switching interfaces summary on page 2384</a><br><a href="#">show ethernet-switching interfaces brief on page 2384</a><br><a href="#">show ethernet-switching interfaces detail on page 2384</a><br><a href="#">show ethernet-switching interfaces interface-name on page 2385</a>                                                                                                             |
| <b>Output Fields</b>            | <p><a href="#">Table 158</a> lists the output fields for the <b>show ethernet-switching interfaces</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                 |

**Table 196: show ethernet-switching interfaces Output Fields**

| Field Name          | Field Description                                      | Level of Output                                     |
|---------------------|--------------------------------------------------------|-----------------------------------------------------|
| <b>Interface</b>    | Name of a switching interface.                         | All levels                                          |
| <b>State</b>        | Interface state. Values are <b>up</b> or <b>down</b> . | none, <b>brief</b> , <b>detail</b> , <b>summary</b> |
| <b>VLAN members</b> | Name of a VLAN.                                        | none, <b>brief</b> , <b>detail</b> , <b>summary</b> |

Table 196: show ethernet-switching interfaces Output Fields (*continued*)

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Level of Output                                     |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------|
| <b>Blocking</b>          | Forwarding state of the interface: <ul style="list-style-type: none"> <li>• <b>blocked</b>—Traffic is not being forwarded on the interface.</li> <li>• <b>unblocked</b>—Traffic is forwarded on the interface.</li> <li>• <b>MAC limit exceeded</b>—The interface is temporarily disabled because of a MAC limiting error. The disabled interface is automatically restored to service when the disable timeout expires.</li> <li>• <b>MAC move limit exceeded</b>—The interface is temporarily disabled because of a MAC move limiting error. The disabled interface is automatically restored to service when the disable timeout expires.</li> <li>• <b>Storm control in effect</b> —The interface is temporarily disabled because of a storm control error. The disabled interface is automatically restored to service when the disable timeout expires.</li> <li>• <b>Storm control shutdown in effect</b> —The interface is temporarily disabled because of a storm control shutdown error. The disabled interface is automatically restored to service when the disable timeout expires.</li> </ul> | none, <b>brief</b> , <b>detail</b> , <b>summary</b> |
| <b>Index</b>             | VLAN index internal to Junos OS software.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>detail</b>                                       |
| <b>untagged   tagged</b> | Specifies whether the interface forwards IEEE802.1Q-tagged or untagged traffic.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail</b>                                       |

## Sample Output

### show ethernet-switching interfaces

```
user@switch> show ethernet-switching interfaces
```

```

Interface State VLAN members Blocking
xe-0/0/0.0 up T1122 unblocked
xe-0/0/1.0 down default - MAC limit exceeded
xe-0/0/2.0 down default - MAC move limit exceeded
xe-0/0/3.0 down default - Storm control in effect
xe-0/0/4.0 down default unblocked
xe-0/0/5.0 down default unblocked
xe-0/0/6.0 down default unblocked
xe-0/0/7.0 down default unblocked
xe-0/0/8.0 down default unblocked
xe-0/0/9.0 up T111 unblocked
xe-0/0/10.0 down default unblocked
xe-0/0/11.0 down default unblocked
xe-0/0/12.0 down default unblocked
xe-0/0/13.0 down default unblocked
xe-0/0/14.0 down default unblocked
xe-0/0/15.0 down default unblocked
xe-0/0/16.0 down default unblocked
xe-0/0/17.0 down default unblocked
xe-0/0/18.0 down default unblocked
xe-0/0/19.0 up T111 unblocked
xe-0/1/0.0 down default unblocked
xe-0/1/1.0 down default unblocked
xe-0/1/2.0 down default unblocked
xe-0/1/3.0 down default unblocked

```

### show ethernet-switching interfaces summary

```
user@switch> show ethernet-switching interfaces summary
xe-0/0/0.0
xe-0/0/1.0
xe-0/0/2.0
xe-0/0/3.0
xe-0/0/8.0
xe-0/0/10.0
xe-0/0/11.0
```

### show ethernet-switching interfaces brief

```
user@switch> show ethernet-switching interfaces brief
Interface State VLAN members Blocking
xe-0/0/0.0 down default unblocked
xe-0/0/1.0 down employee-vlan unblocked
xe-0/0/2.0 down employee-vlan unblocked
xe-0/0/3.0 down employee-vlan unblocked
xe-0/0/8.0 down employee-vlan unblocked
xe-0/0/10.0 down default unblocked
xe-0/0/11.0 down employee-vlan unblocked
```

### show ethernet-switching interfaces detail

```
user@switch> show ethernet-switching interfaces detail
Interface: xe-0/0/0.0 Index: 65
State: down
VLANs:
 default untagged unblocked

Interface: xe-0/0/1.0 Index: 66
State: down
VLANs:
 employee-vlan untagged unblocked

Interface: xe-0/0/2.0 Index: 67
State: down
VLANs:
 employee-vlan untagged unblocked

Interface: xe-0/0/3.0 Index: 68
State: down
VLANs:
 employee-vlan untagged unblocked

Interface: xe-0/0/8.0 Index: 69
State: down
VLANs:
 employee-vlan untagged unblocked

Interface: xe-0/0/10.0 Index: 70
State: down
VLANs:
 default untagged unblocked

Interface: xe-0/0/11.0 Index: 71
State: down
VLANs:
 employee-vlan tagged unblocked
```



**show ethernet-switching interfaces interface-name**

```
user@switch> show ethernet-switching interfaces xe-0/0/0.0
Interface State VLAN members Blocking
xe-0/0/0.0 down default unblocked
```

## show ethernet-switching table

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--|
| <b>Syntax</b>                   | <pre>show ethernet-switching table &lt;brief   detail   extensive   summary&gt; &lt;interface <i>interface-name</i>&gt; &lt;management-vlan&gt; &lt;sort-by (<i>name</i>   <i>tag</i>)&gt; &lt;vlan <i>vlan-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                 |  |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Output for private VLANs introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |  |
| <b>Description</b>              | Displays the Ethernet switching table.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |  |
| <b>Options</b>                  | <p><b>none</b>—(Optional) Display brief information about the Ethernet switching table.</p> <p><b>brief   detail   extensive   summary</b>—(Optional) Display the specified level of output.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display the Ethernet switching table for a specific interface.</p> <p><b>management-vlan</b>—(Optional) Display the Ethernet switching table for a management VLAN.</p> <p><b>sort-by (<i>name</i>   <i>tag</i>)</b>—(Optional) Display VLANs in ascending order of VLAN IDs or VLAN names.</p> <p><b>vlan <i>vlan-name</i></b>—(Optional) Display the Ethernet switching table for a specific VLAN.</p> |  |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 2096</a></li> <li>• <a href="#">Example: Setting Up Bridging with Multiple VLANs</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                   |  |
| <b>List of Sample Output</b>    | <p><a href="#">show ethernet-switching table (Enhanced Layer 2 Software) on page 2387</a></p> <p><a href="#">show ethernet-switching table on page 2388</a></p> <p><a href="#">show ethernet-switching table (Private VLANs) on page 2389</a></p> <p><a href="#">show ethernet-switching table brief on page 2389</a></p> <p><a href="#">show ethernet-switching table detail on page 2390</a></p> <p><a href="#">show ethernet-switching table extensive on page 2391</a></p> <p><a href="#">show ethernet-switching table interface on page 2393</a></p>                                                                                                     |  |
| <b>Output Fields</b>            | <p><a href="#">Table 197</a> lists the output fields for the <b>show ethernet-switching table</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |  |

Table 197: show ethernet-switching table Output Fields

| Field Name | Field Description | Level of Output |
|------------|-------------------|-----------------|
| VLAN       | Name of a VLAN.   | All levels      |

Table 197: show ethernet-switching table Output Fields (*continued*)

| Field Name         | Field Description                                                                                                                                                                                                                                                                                             | Level of Output          |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------|
| <b>MAC address</b> | MAC address associated with the VLAN.                                                                                                                                                                                                                                                                         | All levels               |
| <b>Type</b>        | Type of MAC address: <ul style="list-style-type: none"> <li>• <b>static</b>—The MAC address is manually created.</li> <li>• <b>learn</b>—The MAC address is learned dynamically from a packet's source MAC address.</li> <li>• <b>flood</b>—The MAC address is unknown and flooded to all members.</li> </ul> | All levels               |
| <b>Age</b>         | Time remaining before the entry ages out and is removed from the Ethernet switching table.                                                                                                                                                                                                                    | All levels               |
| <b>Interfaces</b>  | Interface associated with learned MAC addresses or with the <b>All-members</b> option (flood entry).                                                                                                                                                                                                          | All levels               |
| <b>Learned</b>     | For learned entries, the time at which the entry was added to the Ethernet switching table.                                                                                                                                                                                                                   | <b>detail, extensive</b> |

## Sample Output

### show ethernet-switching table (Enhanced Layer 2 Software)

```
user@switch> show ethernet-switching table
MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent
static
 SE - statistics enabled, NM - non configured MAC, R - remote PE MAC,
0 - ovsdb MAC)
```

```
Ethernet switching table : 2 entries, 2 learned
```

```
Routing instance : default-switch
```

| Vlan<br>name | MAC<br>address    | MAC<br>flags | Age | Logical<br>interface |
|--------------|-------------------|--------------|-----|----------------------|
| vlan1        | b0:c6:9a:ca:3c:01 | D            | -   | ae1.0                |
| vlan1        | b0:c6:9a:ca:3c:03 | D            | -   | ae1.0                |

```
MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent
static
```

```
 SE - statistics enabled, NM - non configured MAC, R - remote PE MAC,
0 - ovsdb MAC)
```

```
Ethernet switching table : 2 entries, 2 learned
```

```
Routing instance : default-switch
```

| Vlan<br>name | MAC<br>address    | MAC<br>flags | Age | Logical<br>interface |
|--------------|-------------------|--------------|-----|----------------------|
| vlan10       | b0:c6:9a:ca:3c:01 | D            | -   | ae1.0                |
| vlan10       | b0:c6:9a:ca:3c:03 | D            | -   | ae1.0                |

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static

SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)

Ethernet switching table : 2 entries, 2 learned

Routing instance : default-switch

| Vlan name | MAC address       | MAC flags | Age | Logical interface |
|-----------|-------------------|-----------|-----|-------------------|
| vlan2     | b0:c6:9a:ca:3c:01 | D         | -   | ae1.0             |
| vlan2     | b0:c6:9a:ca:3c:03 | D         | -   | ae1.0             |

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static

SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)

Ethernet switching table : 2 entries, 2 learned

Routing instance : default-switch

| Vlan name | MAC address       | MAC flags | Age | Logical interface |
|-----------|-------------------|-----------|-----|-------------------|
| vlan3     | b0:c6:9a:ca:3c:01 | D         | -   | ae1.0             |
| vlan3     | b0:c6:9a:ca:3c:03 | D         | -   | ae1.0             |

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static

SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)

Ethernet switching table : 2 entries, 2 learned

Routing instance : default-switch

| Vlan name | MAC address       | MAC flags | Age | Logical interface |
|-----------|-------------------|-----------|-----|-------------------|
| vlan4     | b0:c6:9a:ca:3c:01 | D         | -   | ae1.0             |
| vlan4     | b0:c6:9a:ca:3c:03 | D         | -   | ae1.0             |

### show ethernet-switching table

user@switch> show ethernet-switching table

Ethernet-switching table: 57 entries, 17 learned

| VLAN  | MAC address       | Type   | Age | Interfaces  |
|-------|-------------------|--------|-----|-------------|
| F2    | *                 | Flood  | -   | All-members |
| F2    | 00:00:05:00:00:03 | Learn  | 0   | xe-0/0/44.0 |
| F2    | 00:19:e2:50:7d:e0 | Static | -   | Router      |
| Linux | *                 | Flood  | -   | All-members |
| Linux | 00:19:e2:50:7d:e0 | Static | -   | Router      |
| Linux | 00:30:48:90:54:89 | Learn  | 0   | xe-0/0/47.0 |
| T1    | *                 | Flood  | -   | All-members |
| T1    | 00:00:05:00:00:01 | Learn  | 0   | xe-0/0/46.0 |
| T1    | 00:00:5e:00:01:00 | Static | -   | Router      |
| T1    | 00:19:e2:50:63:e0 | Learn  | 0   | xe-0/0/46.0 |
| T1    | 00:19:e2:50:7d:e0 | Static | -   | Router      |
| T10   | *                 | Flood  | -   | All-members |

```

T10 00:00:5e:00:01:09 Static - Router
T10 00:19:e2:50:63:e0 Learn 0 xe-0/0/46.0
T10 00:19:e2:50:7d:e0 Static - Router
T111 * Flood - All-members
T111 00:19:e2:50:63:e0 Learn 0 xe-0/0/15.0
T111 00:19:e2:50:7d:e0 Static - Router
T111 00:19:e2:50:ac:00 Learn 0 xe-0/0/15.0
T2 * Flood - All-members
T2 00:00:5e:00:01:01 Static - Router
T2 00:19:e2:50:63:e0 Learn 0 xe-0/0/46.0
T2 00:19:e2:50:7d:e0 Static - Router
T3 * Flood - All-members
T3 00:00:5e:00:01:02 Static - Router
T3 00:19:e2:50:63:e0 Learn 0 xe-0/0/46.0
T3 00:19:e2:50:7d:e0 Static - Router
T4 * Flood - All-members
T4 00:00:5e:00:01:03 Static - Router
T4 00:19:e2:50:63:e0 Learn 0 xe-0/0/46.0
[output truncated]

```

### show ethernet-switching table (Private VLANs)

```

user@switch> show ethernet-switching table
Ethernet-switching table: 10 entries, 3 learned
VLAN MAC address Type Age Interfaces
pvlan * Flood - All-members
pvlan 00:10:94:00:00:02 Replicated - xe-0/0/28.0
pvlan 00:10:94:00:00:35 Replicated - xe-0/0/46.0
pvlan 00:10:94:00:00:46 Replicated - xe-0/0/4.0
c2 * Flood - All-members
c2 00:10:94:00:00:02 Learn 0 xe-0/0/28.0
c1 * Flood - All-members
c1 00:10:94:00:00:46 Learn 0 xe-0/0/4.0
__pvlan_pvlan_xe-0/0/46.0__ * Flood - All-members
__pvlan_pvlan_xe-0/0/46.0__ 00:10:94:00:00:35 Learn 0 xe-0/0/46.0

```

### show ethernet-switching table brief

```

user@switch> show ethernet-switching table brief
Ethernet-switching table: 57 entries, 17 learned
VLAN MAC address Type Age Interfaces
F2 * Flood - All-members
F2 00:00:05:00:00:03 Learn 0 xe-0/0/44.0
F2 00:19:e2:50:7d:e0 Static - Router
Linux * Flood - All-members
Linux 00:19:e2:50:7d:e0 Static - Router
Linux 00:30:48:90:54:89 Learn 0 xe-0/0/47.0
T1 * Flood - All-members
T1 00:00:05:00:00:01 Learn 0 xe-0/0/46.0
T1 00:00:5e:00:01:00 Static - Router
T1 00:19:e2:50:63:e0 Learn 0 xe-0/0/46.0
T1 00:19:e2:50:7d:e0 Static - Router
T10 * Flood - All-members
T10 00:00:5e:00:01:09 Static - Router
T10 00:19:e2:50:63:e0 Learn 0 xe-0/0/46.0
T10 00:19:e2:50:7d:e0 Static - Router
T111 * Flood - All-members
T111 00:19:e2:50:63:e0 Learn 0 xe-0/0/15.0
T111 00:19:e2:50:7d:e0 Static - Router
T111 00:19:e2:50:ac:00 Learn 0 xe-0/0/15.0
T2 * Flood - All-members

```

```

T2 00:00:5e:00:01:01 Static - Router
T2 00:19:e2:50:63:e0 Learn 0 xe-0/0/46.0
T2 00:19:e2:50:7d:e0 Static - Router
T3 * Flood - All-members
T3 00:00:5e:00:01:02 Static - Router
T3 00:19:e2:50:63:e0 Learn 0 xe-0/0/46.0
T3 00:19:e2:50:7d:e0 Static - Router
T4 * Flood - All-members
T4 00:00:5e:00:01:03 Static - Router
T4 00:19:e2:50:63:e0 Learn 0 xe-0/0/46.0
[output truncated]

```

### show ethernet-switching table detail

```

user@switch> show ethernet-switching table detail
Ethernet-switching table: 57 entries, 17 learned
F2, *
 Interface(s): xe-0/0/44.0
 Type: Flood
 Nexthop index: 0

F2, 00:00:05:00:00:03
 Interface(s): xe-0/0/44.0
 Type: Learn, Age: 0, Learned: 2:03:09
 Nexthop index: 0

F2, 00:19:e2:50:7d:e0
 Interface(s): Router
 Type: Static
 Nexthop index: 0

Linux, *
 Interface(s): xe-0/0/47.0
 Type: Flood
 Nexthop index: 0

Linux, 00:19:e2:50:7d:e0
 Interface(s): Router
 Type: Static
 Nexthop index: 0

Linux, 00:30:48:90:54:89
 Interface(s): xe-0/0/47.0
 Type: Learn, Age: 0, Learned: 2:03:08
 Nexthop index: 0

T1, *
 Interface(s): xe-0/0/46.0
 Type: Flood
 Nexthop index: 0

T1, 00:00:05:00:00:01
 Interface(s): xe-0/0/46.0
 Type: Learn, Age: 0, Learned: 2:03:07
 Nexthop index: 0

T1, 00:00:5e:00:01:00
 Interface(s): Router
 Type: Static
 Nexthop index: 0

```

```

T1, 00:19:e2:50:63:e0
 Interface(s): xe-0/0/46.0
 Type: Learn, Age: 0, Learned: 2:03:07
 Nexthop index: 0

T1, 00:19:e2:50:7d:e0
 Interface(s): Router
 Type: Static
 Nexthop index: 0

T10, *
 Interface(s): xe-0/0/46.0
 Type: Flood
 Nexthop index: 0

T10, 00:00:5e:00:01:09
 Interface(s): Router
 Type: Static
 Nexthop index: 0

T10, 00:19:e2:50:63:e0
 Interface(s): xe-0/0/46.0
 Type: Learn, Age: 0, Learned: 2:03:08
 Nexthop index: 0

T10, 00:19:e2:50:7d:e0
 Interface(s): Router
 Type: Static
 Nexthop index: 0

T111, *
 Interface(s): xe-0/0/15.0
 Type: Flood
 Nexthop index: 0
[output truncated]

```

### show ethernet-switching table extensive

```

user@switch> show ethernet-switching table extensive
Ethernet-switching table: 57 entries, 17 learned
F2, *
 Interface(s): xe-0/0/44.0
 Type: Flood
 Nexthop index: 0

F2, 00:00:05:00:00:03
 Interface(s): xe-0/0/44.0
 Type: Learn, Age: 0, Learned: 2:03:09
 Nexthop index: 0

F2, 00:19:e2:50:7d:e0
 Interface(s): Router
 Type: Static
 Nexthop index: 0

Linux, *
 Interface(s): xe-0/0/47.0
 Type: Flood
 Nexthop index: 0

Linux, 00:19:e2:50:7d:e0

```

```
Interface(s): Router
Type: Static
Nexthop index: 0

Linux, 00:30:48:90:54:89
Interface(s): xe-0/0/47.0
Type: Learn, Age: 0, Learned: 2:03:08
Nexthop index: 0

T1, *
Interface(s): xe-0/0/46.0
Type: Flood
Nexthop index: 0

T1, 00:00:05:00:00:01
Interface(s): xe-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:07
Nexthop index: 0

T1, 00:00:5e:00:01:00
Interface(s): Router
Type: Static
Nexthop index: 0

T1, 00:19:e2:50:63:e0
Interface(s): xe-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:07
Nexthop index: 0

T1, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
Nexthop index: 0

T10, *
Interface(s): xe-0/0/46.0
Type: Flood
Nexthop index: 0

T10, 00:00:5e:00:01:09
Interface(s): Router
Type: Static
Nexthop index: 0

T10, 00:19:e2:50:63:e0
Interface(s): xe-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:08
Nexthop index: 0

T10, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
Nexthop index: 0

T111, *
Interface(s): xe-0/0/15.0
Type: Flood
Nexthop index: 0
[output truncated]
```



### show ethernet-switching table interface

```
user@switch> show ethernet-switching table interface xe-0/0/1
Ethernet-switching table: 1 unicast entries
VLAN MAC address Type Age Interfaces
V1 * Flood - All-members
V1 00:00:05:00:00:05 Learn 0 xe-0/0/1.0
```

## show system statistics arp

---

|                                 |                                                                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show system statistics arp                                                                                                                                                                          |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.6 for EX Series switches.                                                                                                                                  |
| <b>Description</b>              | Display system-wide Address Resolution Protocol (ARP) statistics.                                                                                                                                   |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Proxy ARP on an EX Series Switch</i></li><li>• <a href="#">Verifying That Proxy ARP Is Working Correctly on page 2245</a></li></ul> |

## show system statistics arp

```
user@switch> show system statistics arp
arp:
 90060 datagrams received
 34 ARP requests received
 610 ARP replies received
 0 resolution request received
 0 unrestricted proxy requests
 0 restricted proxy requests
 0 received proxy requests
 0 unrestricted proxy requests not proxied
 0 restricted proxy requests not proxied
 0 datagrams with bogus interface
 0 datagrams with incorrect length
 0 datagrams for non-IP protocol
 0 datagrams with unsupported op code
 0 datagrams with bad protocol address length
 0 datagrams with bad hardware address length
 0 datagrams with multicast source address
 0 datagrams with multicast target address
 0 datagrams with my own hardware address
 0 datagrams for an address not on the interface
 0 datagrams with a broadcast source address
 294 datagrams with source address duplicate to mine
 89113 datagrams which were not for me
 0 packets discarded waiting for resolution
 0 packets sent after waiting for resolution
 309 ARP requests sent
 35 ARP replies sent
 0 requests for memory denied
 0 requests dropped on entry
 0 requests dropped during retry
 0 requests dropped due to interface deletion
 0 requests on unnumbered interfaces
 0 new requests on unnumbered interfaces
 0 replies for from unnumbered interfaces
 0 requests on unnumbered interface with non-subnetted donor
 0 replies from unnumbered interface with non-subnetted donor
```

## show vlans

**Syntax** `show vlans`  
`<brief | detail | extensive>`  
`<dot1q-tunneling>`  
`<sort-by (tag | name)>`  
`<vlan-range-name>`

**Release Information** Command introduced in Junos OS Release 11.1 for the QFX Series.  
Option **dot1q-tunneling** added in Junos OS Release 12.1 for the QFX Series.

**Description** Display information about VLANs configured on bridged Ethernet interfaces. For interfaces configured to support a VoIP VLAN and a data VLAN, the **show vlans** command displays both tagged and untagged membership for those VLANs.



**NOTE:** When a series of VLANs is created using the `vlan-range` statement, such VLAN names are preceded and followed by a double underscore. For example, a series of VLANs using the VLAN range 1 through 3 and the base VLAN name `marketing` would be displayed as `__marketing_1__`, `__marketing_2__`, and `__marketing_3__`.



**NOTE:** To display an 802.1X supplicant successfully authenticated in multiple-suppliant mode with dynamic VLAN movement, use the `show vlans vlan-name extensive` operational mode command, where *vlan-name* is the dynamic VLAN.

**Options** **none**—Display information for all VLANs. VLAN information is displayed by VLAN name in ascending order.

**brief | detail | extensive**—(Optional) Display the specified level of output.

**sort-by (tag | name)**—(Optional) Display VLANs in ascending order of VLAN IDs or VLAN names.

**vlan-range-name**—(Optional) Display VLANs in ascending order of VLAN range names.

**Required Privilege Level** `view`

**Related Documentation**

- [Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 2096](#)
- [Example: Setting Up Bridging with Multiple VLANs](#)
- [Understanding Bridging and VLANs](#)
- [show ethernet-switching interfaces on page 1998](#)

**List of Sample Output**

- [show vlans on page 2398](#)
- [show vlans \(Private VLANs\) on page 2398](#)
- [show vlans brief on page 2399](#)
- [show vlans detail on page 2399](#)
- [show vlans extensive \(Port-Based\) on page 2400](#)
- [show vlans \(Q-in-Q Tunneling\) on page 2401](#)
- [show vlans extensive \(Q-in-Q Tunneling\) on page 2401](#)
- [show vlans extensive \(Q-in-Q Tunneling and L2TP\) on page 2401](#)
- [show vlans sort-by tag on page 2401](#)
- [show vlans sort-by name on page 2402](#)
- [show vlans tag on page 2403](#)

**Output Fields** Table 162 lists the output fields for the **show vlans** command. Output fields are listed in the approximate order in which they appear.

**Table 198: show vlans Output Fields**

| Field Name                  | Field Description                                                                                                                                                            | Level of Output          |
|-----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------|
| <b>Name</b>                 | Name of a VLAN.                                                                                                                                                              | none, <b>brief</b>       |
| <b>Tag</b>                  | 802.1Q tag applied to this VLAN. If <b>none</b> is displayed, no tag is applied.                                                                                             | All levels               |
| <b>Interfaces</b>           | Interface associated with learned MAC addresses or <b>All-members</b> option (flood entry). An asterisk (*) beside the interface indicates that the interface is <b>UP</b> . | All levels               |
| <b>Address</b>              | IP address.                                                                                                                                                                  | none, <b>brief</b>       |
| <b>Ports Active /Total</b>  | Number of interfaces associated with a VLAN: <b>Active</b> indicates interfaces that are <b>UP</b> , and <b>Total</b> indicates interfaces that are active and inactive.     | <b>brief</b>             |
| <b>VLAN</b>                 | Name of a VLAN.                                                                                                                                                              | <b>detail, extensive</b> |
| <b>Admin state</b>          | State of the interface. Values are:<br><br><b>enabled</b> —The interface is turned on, and the physical link is operational and can pass packets.                            | <b>detail,extensive</b>  |
| <b>MAC learning Status</b>  | Indicates if MAC learning is disabled.                                                                                                                                       | <b>detail, extensive</b> |
| <b>Description</b>          | Description for the VLAN.                                                                                                                                                    | <b>detail,extensive</b>  |
| <b>Primary IP</b>           | Primary IP address associated with a VLAN.                                                                                                                                   | <b>detail</b>            |
| <b>Number of interfaces</b> | Number of interfaces associated with a VLAN. Both the total number of interfaces and the number of active interfaces associated with a VLAN are displayed.                   | <b>detail, extensive</b> |
| <b>STP</b>                  | Spanning tree associated with a VLAN.                                                                                                                                        | <b>detail,extensive</b>  |
| <b>Tagged interfaces</b>    | Tagged interfaces with which a VLAN is associated.                                                                                                                           | <b>detail,extensive</b>  |

Table 198: show vlans Output Fields (*continued*)

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Level of Output   |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------|
| Untagged interfaces     | Untagged interfaces with which a VLAN is associated.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | detail, extensive |
| Dot1q Tunneling Status  | Indicates if Q-in-Q tunneling is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | extensive         |
| Customer VLAN ranges    | List of customer VLAN (C-VLAN) ranges associated with this service VLAN (S-VLAN).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | extensive         |
| Private VLAN Mode       | The private VLAN mode for this VLAN. Values include <b>Primary</b> , <b>Isolated</b> , and <b>Community</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | extensive         |
| Primary VLAN            | Primary VLAN tag for this secondary VLAN.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | extensive         |
| Internal Index          | VLAN index internal to Junos OS software.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | extensive         |
| Origin                  | Manner in which the VLAN was created: <b>static</b> or <b>learn</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | extensive         |
| Protocol                | Port-based VLAN or MAC-based VLAN. MAC-based protocol is displayed when VLAN assignment is done either statically or dynamically through 802.1X,                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | extensive         |
| IP addresses            | IP address associated with a VLAN.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | extensive         |
| Number of MAC entries   | For MAC-based VLANs created either statically or dynamically, the MAC addresses associated with an interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | extensive         |
| Number of mapping rules | Number of mapping rules for Q-in-Q tunneling ( <b>Push</b> ) and VLAN translation ( <b>Swap</b> ).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |                   |
| Secondary VLANs         | Secondary VLANs associated with a primary VLAN.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | extensive         |
| Isolated VLANs          | Isolated VLANs associated with a primary VLAN.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | extensive         |
| Community VLANs         | Community VLANs associated with a primary VLAN.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | extensive         |
| VLANs summary           | VLAN counts: <ul style="list-style-type: none"> <li>• <b>Total</b>—Total number of VLANs on the switch.</li> <li>• <b>Configured VLANs</b>—Number of VLANs that are based on user-configured settings.</li> <li>• <b>Internal VLANs</b>—Number of VLANs created by the system with no explicit configuration or protocol—for example, the <b>default</b> VLAN and the VLAN created when a trunk interface is not configured with native VLAN membership.</li> <li>• <b>Temporary VLANs</b>—Number of VLANs from the previous configuration that the system retains for a limited time after restart. Temporary VLANs are converted into one of the other types of VLAN, or are removed from the system if the current configuration does not require them.</li> </ul> | All levels        |

Table 198: show vlans Output Fields (*continued*)

| Field Name                          | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Level of Output |
|-------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Dot1q VLANs summary</b>          | 802.1Q VLAN counts: <ul style="list-style-type: none"> <li>• <b>Total</b>—Total number of 802.1Q-tagged and untagged VLANs on the switch.</li> <li>• <b>Tagged VLANs</b>—Number of 802.1Q-tagged VLANs.</li> <li>• <b>Untagged VLANs</b>—Number of untagged 802.1Q VLANs.</li> <li>• <b>Private VLAN</b>—Counts of the following kinds of 802.1Q private VLANs (PVLANS):           <ul style="list-style-type: none"> <li>• <b>Primary VLANs</b>—Number of primary forwarding private VLANs.</li> <li>• <b>Community VLANs</b>—Number of community transporting and forwarding private VLANs.</li> <li>• <b>Isolated VLANs</b>—Number of isolated receiving and forwarding private VLANs.</li> <li>• <b>Inter-switch-isolated VLANs</b>—Number of inter-switch isolated receiving and forwarding private VLANs.</li> </ul> </li> </ul> | All levels      |
| <b>Dot1q Tunneled VLANs summary</b> | Q-in-Q-tunneled VLAN counts: <ul style="list-style-type: none"> <li>• <b>Total</b>—Total number of Q-in-Q-tunneled VLANs on the switch.</li> <li>• <b>Private VLAN</b>—Counts of primary, community, and isolated Q-in-Q-tunneled private VLANs (PVLANS).</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | All levels      |

## Sample Output

### show vlans

```
user@switch> show vlans
```

| Name    | Tag  | Interfaces                                                                                                                                                                                                                                                                                           |
|---------|------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| default | None | xe-0/0/34.0, xe-0/0/33.0, xe-0/0/32.0, xe-0/0/31.0, xe-0/0/30.0, xe-0/0/29.0, xe-0/0/28.0, xe-0/0/27.0, xe-0/0/26.0, xe-0/0/25.0, xe-0/0/19.0, xe-0/0/18.0, xe-0/0/17.0, xe-0/0/16.0, xe-0/0/15.0, xe-0/0/14.0, xe-0/0/13.0, xe-0/0/11.0, xe-0/0/9.0, xe-0/0/8.0, xe-0/0/3.0, xe-0/0/2.0, xe-0/0/1.0 |
| v0001   | 1    | xe-0/0/24.0, xe-0/0/23.0, xe-0/0/22.0, xe-0/0/21.0                                                                                                                                                                                                                                                   |
| v0002   | 2    | None                                                                                                                                                                                                                                                                                                 |
| v0003   | 3    | None                                                                                                                                                                                                                                                                                                 |
| v0004   | 4    | None                                                                                                                                                                                                                                                                                                 |
| v0005   | 5    | None                                                                                                                                                                                                                                                                                                 |

### show vlans (Private VLANs)

```
user@switch> show vlans
```

| Name                        | Tag | Interfaces |
|-----------------------------|-----|------------|
| __pvlan_pvlan_xe-0/0/46.0__ |     |            |

```

c1 xe-0/0/44.0*, xe-0/0/46.0*
c2 xe-0/0/4.0*, xe-0/0/44.0*
default xe-0/0/28.0*, xe-0/0/44.0*
pvlan 500
 None
 xe-0/0/4.0*, xe-0/0/28.0*, xe-0/0/44.0*, xe-0/0/46.0*

```

### show vlans brief

```
user@switch> show vlans brief
```

| Name    | Tag  | Address | Ports<br>Active/Total |
|---------|------|---------|-----------------------|
| default | None |         | 0/23                  |
| v0001   | 1    |         | 0/4                   |
| v0002   | 2    |         | 0/0                   |
| v0003   | 3    |         | 0/0                   |
| v0004   | 4    |         | 0/0                   |
| v0005   | 5    |         | 0/0                   |
| v0006   | 6    |         | 0/0                   |
| v0007   | 7    |         | 0/0                   |
| v0008   | 8    |         | 0/0                   |
| v0009   | 9    |         | 0/0                   |
| v0010   | 10   |         | 0/2                   |
| v0011   | 11   |         | 0/0                   |
| v0012   | 12   |         | 0/0                   |
| v0013   | 13   |         | 0/0                   |
| v0014   | 14   |         | 0/0                   |
| v0015   | 15   |         | 0/0                   |
| v0016   | 16   |         | 0/0                   |

### show vlans detail

```
user@switch> show vlans detail
```

```
VLAN: default, Tag: Untagged, Admin state: Enabled
```

```
Description: None
```

```
Primary IP: None, Number of interfaces: 23 (Active = 0)
```

```
STP: None, RTG: None
```

```
Untagged interfaces: xe-0/0/34.0, xe-0/0/33.0, xe-0/0/32.0, xe-0/0/31.0,
xe-0/0/30.0, xe-0/0/29.0, xe-0/0/28.0, xe-0/0/27.0, xe-0/0/26.0,
xe-0/0/25.0, xe-0/0/19.0, xe-0/0/18.0, xe-0/0/17.0, xe-0/0/16.0,
xe-0/0/15.0, xe-0/0/14.0, xe-0/0/13.0, xe-0/0/11.0, xe-0/0/9.0, xe-0/0/8.0,
xe-0/0/3.0, xe-0/0/2.0, xe-0/0/1.0,
```

```
Tagged interfaces: None
```

```
VLAN: v0001, Tag: 802.1Q Tag 1, Admin state: Enabled
```

```
Description: None
```

```
Primary IP: None, Number of interfaces: 4 (Active = 0)
```

```
Dot1q Tunneling Status: Enabled
```

```
STP: None, RTG: None
```

```
Untagged interfaces: None
```

```
Tagged interfaces: xe-0/0/24.0, xe-0/0/23.0, xe-0/0/22.0, xe-0/0/21.0,
```

```
VLAN: v0002, Tag: 802.1Q Tag 2, Admin state: Enabled
```

```
Description: None
```

```
Primary IP: None, Number of interfaces: 0 (Active = 0)
```

```
STP: None, RTG: None
```

```
Untagged interfaces: None
```

```
Tagged interfaces: None
```

```
VLAN: v0003, Tag: 802.1Q Tag 3, Admin state: Enabled
Description: None
Primary IP: None, Number of interfaces: 0 (Active = 0)
STP: None, RTG: None
Untagged interfaces: None
Tagged interfaces: None

VLAN: vlan4000, 802.1Q Tag: Untagged, Admin State: Enabled
MAC learning Status: Disabled
Number of interfaces: 0 (Active = 0)
```

### show vlans extensive (Port-Based)

```
user@switch> show vlans extensive
VLAN: default, created at Mon Feb 4 12:13:47 2008
Tag: None, Internal index: 0, Admin state: Enabled, Origin: static
Description: None
Customer VLAN ranges:
 1-4100
Protocol: Port based
IP addresses: None
STP: None, RTG: None.
Number of interfaces: Tagged 0 (Active = 0), Untagged 23 (Active = 0)
 xe-0/0/34.0 (untagged, access)
 xe-0/0/33.0 (untagged, access)
 xe-0/0/32.0 (untagged, access)
 xe-0/0/31.0 (untagged, access)
 xe-0/0/30.0 (untagged, access)
 xe-0/0/29.0 (untagged, access)
 xe-0/0/28.0 (untagged, access)
 xe-0/0/27.0 (untagged, access)
 xe-0/0/26.0 (untagged, access)
 xe-0/0/25.0 (untagged, access)
 xe-0/0/19.0 (untagged, access)
 xe-0/0/18.0 (untagged, access)
 xe-0/0/17.0 (untagged, access)
 xe-0/0/16.0 (untagged, access)
 xe-0/0/15.0 (untagged, access)
 xe-0/0/14.0 (untagged, access)
 xe-0/0/13.0 (untagged, access)
 xe-0/0/11.0 (untagged, access)
 xe-0/0/9.0 (untagged, access)
 xe-0/0/8.0 (untagged, access)
 xe-0/0/3.0 (untagged, access)
 xe-0/0/2.0 (untagged, access)
 xe-0/0/1.0 (untagged, access)

Secondary VLANs: Isolated 1, Community 1
Isolated VLANs :
 __pvlan_pvlan_xe-0/0/3.0__
Community VLANs :
 comm1

VLAN: v0001, created at Mon Feb 4 12:13:47 2008
Tag: 1, Internal index: 1, Admin state: Enabled, Origin: static
Description: None
Protocol: Port based, Layer 3 interface: None
IP addresses: None
STP: None, RTG: None.
Number of interfaces: Tagged 4 (Active = 0), Untagged 0 (Active = 0)
```



```

xe-0/0/24.0 (tagged, trunk)
xe-0/0/23.0 (tagged, trunk)
xe-0/0/22.0 (tagged, trunk)
xe-0/0/21.0 (tagged, trunk)

```

```

VLAN: v0002, created at Mon Feb 4 12:13:47 2008
Tag: 2, Internal index: 2, Admin state: Enabled, Origin: static
Description: None
Protocol: Port based, Layer 3 interface: None
IP addresses: None
STP: None, RTG: None.
Number of interfaces: Tagged 0 (Active = 0), Untagged 0 (Active = 0)
None

```

```

VLAN: v0003, created at Mon Feb 4 12:13:47 2008
Tag: 3, Internal index: 3, Admin state: Enabled, Origin: static
Description: None
Protocol: Port based, Layer 3 interface: None
IP addresses: None
STP: None, RTG: None.
Number of interfaces: Tagged 0 (Active = 0), Untagged 0 (Active = 0)
None

```

### show vlans (Q-in-Q Tunneling)

```

user@switch> show vlans dot1q-tunneling
Name Tag Interfaces
sv100 100 xe-0/0/4.0*, xe-0/0/15.0*

```

### show vlans extensive (Q-in-Q Tunneling)

```

user@switch> show vlans sv100 extensive
VLAN: sv100, Created at: Sat Sep 10 12:53:52 2011
802.1Q Tag: 100, Internal index: 2, Admin State: Enabled, Origin: Static
Dot1q Tunneling Status: Enabled
Customer VLAN ranges:
 10-20
 40-50
Protocol: Port Mode
Number of interfaces: Tagged 1 (Active = 1), Untagged 0 (Active = 0)
 ge-0/0/0.0, tagged, trunk

Number of mapping rules:
 Push 1 (Active = 0), Policy 0 (Active = 0), Swap 0 (Active = 0)

 xe-0/0/3.0*, 300, push

```

### show vlans extensive (Q-in-Q Tunneling and L2TP)

```

user@switch> show vlans v1 extensive
VLAN: v1, Created at: Fri Mar 2 05:07:38 2012
802.1Q Tag: 100, Internal index: 4, Admin State: Enabled, Origin: Static
Dot1q Tunneling status: Enabled
Layer2 Protocol Tunneling status: Enabled

```

### show vlans sort-by tag

```

user@switch> show vlans sort-by tag
Name Tag Interfaces
default None
__vlan-x_1__ 1

```

|               |    |      |
|---------------|----|------|
| __vlan-x_2__  | 2  | None |
| __vlan-x_3__  | 3  | None |
| __vlan-x_4__  | 4  | None |
| __vlan-x_5__  | 5  | None |
| __vlan-x_6__  | 6  | None |
| __vlan-x_7__  | 7  | None |
| __vlan-x_8__  | 8  | None |
| __vlan-x_9__  | 9  | None |
| __vlan-x_10__ | 10 | None |
| __vlan-x_11__ | 11 | None |
| __vlan-x_12__ | 12 | None |
| __vlan-x_13__ | 13 | None |
| __vlan-x_14__ | 14 | None |
| __vlan-x_15__ | 15 | None |
| __vlan-x_16__ | 16 | None |
| __vlan-x_17__ | 17 | None |
| __vlan-x_18__ | 18 | None |
| __vlan-x_19__ | 19 | None |
| __vlan-x_20__ | 20 | None |

#### show vlans sort-by name

```
user@switch> show vlans sort-by employee
```

| Name             | Tag | Interfaces   |
|------------------|-----|--------------|
| __employee_120__ | 120 | xe-0/0/22.0* |
| __employee_121__ | 121 | xe-0/0/22.0* |
| __employee_122__ | 122 | xe-0/0/22.0* |
| __employee_123__ | 123 | xe-0/0/22.0* |
| __employee_124__ | 124 | xe-0/0/22.0* |
| __employee_125__ | 125 | xe-0/0/22.0* |
| __employee_126__ | 126 | xe-0/0/22.0* |
| __employee_127__ | 127 |              |

```

__employee_128__ 128 xe-0/0/22.0*
__employee_129__ 129 xe-0/0/22.0*
__employee_130__ 130 xe-0/0/22.0*

```

### show vlans tag

```
user@switch> show vlans employee
```

| Name             | Tag | Interfaces   |
|------------------|-----|--------------|
| __employee_120__ | 120 | xe-0/0/22.0* |
| __employee_121__ | 121 | xe-0/0/22.0* |
| __employee_122__ | 122 | xe-0/0/22.0* |
| __employee_123__ | 123 | xe-0/0/22.0* |
| __employee_124__ | 124 | xe-0/0/22.0* |
| __employee_125__ | 125 | xe-0/0/22.0* |
| __employee_126__ | 126 | xe-0/0/22.0* |
| __employee_127__ | 127 | xe-0/0/22.0* |
| __employee_128__ | 128 | xe-0/0/22.0* |
| __employee_129__ | 129 | xe-0/0/22.0* |
| __employee_130__ | 130 | xe-0/0/22.0* |



## CHAPTER 86

# MAC Address Operational Commands

- `show ethernet-switching mac-learning-log`
- `show ethernet-switching mac-notification`
- `show ethernet-switching statistics aging`
- `show ethernet-switching statistics mac-learning`

## show ethernet-switching mac-learning-log

|                                 |                                                                                                                                                                                           |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show ethernet-switching mac-learning-log                                                                                                                                                  |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                           |
| <b>Description</b>              | Displays the event log of learned MAC addresses.                                                                                                                                          |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show ethernet-switching table on page 2386</a></li> <li>• <a href="#">show ethernet-switching interfaces on page 1998</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show ethernet-switching mac-learning-log on page 2406</a>                                                                                                                     |
| <b>Output Fields</b>            | Table 199 lists the output fields for the <b>show ethernet-switching mac-learning-log</b> command. Output fields are listed in the approximate order in which they appear.                |

Table 199: show ethernet-switching mac-learning-log Output Fields

| Field Name            | Field Description                                                                                                                                                                                                                                                                             |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Date and Time</b>  | Timestamp in UTC when the MAC operation occurred.                                                                                                                                                                                                                                             |
| <b>vlan_name</b>      | VLAN name. A value defined by the user for all user-configured VLANs. The name of the VLAN on which the MAC is learned.                                                                                                                                                                       |
| <b>MAC</b>            | Learned MAC address.                                                                                                                                                                                                                                                                          |
| <b>Event op</b>       | MAC address that are added, learned, deleted, changed or moved from one interface to another interface.                                                                                                                                                                                       |
| <b>Interface Name</b> | The name of the interface on which the MAC address is learned. When a MAC address is moved, there is another field with the name of the interface. The log displays the name of the interface from where the MAC address moved, and the name of the interface to where the MAC address moved. |
| <b>Flags</b>          | Displays the MAC address flags in which the MAC event occurred. This option is for debugging purposes.                                                                                                                                                                                        |

## Sample Output

### show ethernet-switching mac-learning-log

```

user@switch> show ethernet-switching mac-learning-log
Mon Jun 30 13:49:49 2014 vlan_name v11+11 mac 00:10:94:00:00:02 was learned on
ge-1/0/22.0 with flags: 0x2001f << MAC address that as dynamically learned
Mon Jun 30 13:50:29 2014 vlan_name v11+11 mac 00:10:94:00:00:02 was deleted from
ge-1/0/22.0 with flags: 0x1080 << MAC address that was deleted
Mon Jun 30 13:51:28 2014 vlan_name v11+11 mac 00:00:00:01:01:01 was added to
ge-1/0/22.0 with flags: 0x2013f << Static MAC address that was added
Mon Jun 30 13:51:46 2014 vlan_name v11+11 mac 00:00:00:01:01:01 was deleted from
ge-1/0/22.0 with flags: 0x1120 << delete of Static MAC address that was deleted

```

```
Mon Jun 30 13:52:03 2014 vlan_name v11+11 mac 00:10:94:00:00:02 was learned on
ge-1/0/22.0 with flags: 0x2001f << MAC address that was dynamically learned
Mon Jun 30 13:52:11 2014 vlan_name v11+11 mac 00:10:94:00:00:02 was moved from
ge-1/0/22.0 to ge-1/0/21.0 with flags: 0x2101f << MAC address that was moved
Mon Jun 30 13:54:24 2014 vlan_name v11+11 mac 00:10:94:00:00:02 was changed on
ge-1/0/21.0 with flags: 0x2113f << MAC address that changed from a dynamic
address to a static address
```

## show ethernet-switching mac-notification

|                                 |                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show ethernet-switching mac-notification                                                                                                                                                              |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.6 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                 |
| <b>Description</b>              | Display information about MAC notification.                                                                                                                                                           |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Verifying That MAC Notification Is Working Properly</i></li> </ul>                                                                                        |
| <b>List of Sample Output</b>    | <a href="#">show ethernet-switching mac-notification (MAC Notification Enabled) on page 2408</a><br><a href="#">show ethernet-switching mac-notification (MAC Notification Disabled) on page 2408</a> |
| <b>Output Fields</b>            | Table 200 lists the output fields for the <b>show ethernet-switching mac-notification</b> command. Output fields are listed in the order in which they appear.                                        |

Table 200: show ethernet-switching mac-notification Output Fields

| Field Name                         | Field Description                                                                                                                                                                 |
|------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Notification Status</b>         | MAC notification status: <ul style="list-style-type: none"> <li>• <b>Enabled</b>—MAC notification is enabled.</li> <li>• <b>Disabled</b>—MAC notification is disabled.</li> </ul> |
| <b>Notification Interval</b>       | MAC notification interval in seconds.                                                                                                                                             |
| <b>Notifications Sent</b>          | Number of notifications sent to SNMP when MACs are learned or when MACs age out.                                                                                                  |
| <b>Notifications Table Maxsize</b> | Maximum size of the notification table, which is populated when notifications are sent to the SNMP server.                                                                        |

### Sample Output

#### show ethernet-switching mac-notification (MAC Notification Enabled)

```

user@switch> show ethernet-switching mac-notification
Notification Status : Enabled
Notification Interval : 30
Notifications Sent : 0
Notifications Table Maxsize : 256

```

### Sample Output

#### show ethernet-switching mac-notification (MAC Notification Disabled)

```

user@switch> show ethernet-switching mac-notification
Notification Status : Disabled
Notification Interval : 0

```



Notifications Sent : 0  
Notifications Table Maxsize : 256

## show ethernet-switching statistics aging

|                                 |                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show ethernet-switching statistics aging<br><brief   detail>                                                                                                                                         |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                      |
| <b>Description</b>              | Display media access control (MAC) aging statistics.                                                                                                                                                 |
| <b>Options</b>                  | <b>none</b> —(Optional) Display MAC aging statistics.<br><b>brief   detail</b> —(Optional) Display the specified level of output.                                                                    |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show ethernet-switching statistics mac-learning on page 2412</a></li> <li>• <a href="#">Configuring MAC Table Aging on page 2144</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show ethernet-switching statistics aging on page 2411</a>                                                                                                                                |
| <b>Output Fields</b>            | Table 201 lists the output fields for the <b>show ethernet-switching statistics aging</b> command. Output fields are listed in the approximate order in which they appear.                           |

Table 201: show ethernet-switching statistics aging Output Fields

| Field Name                         | Field Description                                                                                                                                                                                                                                                                                                                                                                   | Level of Output |
|------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Total age messages received</b> | Total number of aging messages received from the hardware.                                                                                                                                                                                                                                                                                                                          | All levels      |
| <b>Immediate aging</b>             | Aging message indicating that the entry should be removed immediately.                                                                                                                                                                                                                                                                                                              | All levels      |
| <b>MAC address seen</b>            | Aging message indicating that the MAC address has been detected by hardware and that the aging timer should be stopped.                                                                                                                                                                                                                                                             | All levels      |
| <b>MAC address not seen</b>        | Aging message indicating that the MAC address has not been detected by the hardware and that the aging timer should be started.                                                                                                                                                                                                                                                     | All levels      |
| <b>Error age messages</b>          | The received aging message contains the following errors: <ul style="list-style-type: none"> <li>• <b>Invalid VLAN</b>—The VLAN of the packet does not exist.</li> <li>• <b>No such entry</b>—The MAC address and VLAN pair provided by the aging message does not exist.</li> <li>• <b>Static entry</b>—An unsuccessful attempt was made to age out a static MAC entry.</li> </ul> | All levels      |

## Sample Output

### show ethernet-switching statistics aging

```
user@switch> show ethernet-switching statistics aging
```

```
Total age messages received: 0
```

```
Immediate aging: 0, MAC address seen: 0, MAC address not seen: 0
```

```
Error age messages: 0
```

```
Invalid VLAN: 0, No such entry: 0, Static entry: 0
```

## show ethernet-switching statistics mac-learning

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show ethernet-switching statistics mac-learning</code><br><code>&lt;brief   detail&gt;</code><br><code>&lt;interface <i>interface-name</i>&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.4 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Display media access control (MAC) learning statistics.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <b>none</b> —(Optional) Display MAC learning statistics for all interfaces.<br><br><b>brief   detail</b> —(Optional) Display the specified level of output. The default is <b>brief</b> .<br><br><b>interface <i>interface-name</i></b> —(Optional) Display MAC learning statistics for the specified interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>show ethernet-switching statistics aging</i></li><li>• <i>show ethernet-switching mac-learning-log</i></li><li>• <i>show ethernet-switching table</i></li><li>• <i>show ethernet-switching interfaces</i></li><li>• <i>Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch</i></li><li>• <i>Example: Setting Up Bridging with Multiple VLANs for EX Series Switches</i></li><li>• <a href="#">show ethernet-switching statistics aging on page 2410</a></li><li>• <a href="#">show ethernet-switching mac-learning-log on page 2406</a></li><li>• <a href="#">show ethernet-switching table on page 2386</a></li><li>• <a href="#">show ethernet-switching interfaces on page 1998</a></li><li>• <a href="#">Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 2096</a></li><li>• <a href="#">Example: Setting Up Bridging with Multiple VLANs</a></li></ul> |
| <b>List of Sample Output</b>    | <a href="#">show ethernet-switching statistics mac-learning on page 2413</a><br><a href="#">show ethernet-switching statistics mac-learning detail on page 2414</a><br><a href="#">show ethernet-switching statistics mac-learning interface ge-0/0/28 detail on page 2414</a><br><a href="#">show ethernet-switching statistics mac-learning interface on page 2414</a><br><a href="#">show ethernet-switching statistics mac-learning detail (QFX Series) on page 2414</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Output Fields</b>            | <a href="#">Table 202</a> lists the output fields for the <b>show ethernet-switching statistics mac-learning</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

Table 202: show ethernet-switching statistics mac-learning Output Fields

| Field Name                                   | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Level of Output |
|----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Interface</b>                             | Name of the interface for which statistics are being reported. (Displayed in the output under the heading <b>Interface</b> .)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | All levels      |
| <b>Learning message from local packets</b>   | MAC learning message generated due to packets coming in on the management interface. (Displayed in the output under the heading <b>Local pkts</b> .)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | All levels      |
| <b>Learning message from transit packets</b> | MAC learning message generated due to packets coming in on network interfaces. (Displayed in the output under the heading <b>Transit pkts</b> .)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | All levels      |
| <b>Learning message with error</b>           | <p>MAC learning messages received with errors (Displayed under the heading <b>Error</b>):</p> <ul style="list-style-type: none"> <li>• <b>Invalid VLAN</b>—The VLAN of the packet does not exist.</li> <li>• <b>Invalid MAC</b>—The MAC address is either NULL or a multicast MAC address.</li> <li>• <b>Security violation</b>—The MAC address is not an allowed MAC address.</li> <li>• <b>Interface down</b>—The MAC address is learned on an interface that is down.</li> <li>• <b>Incorrect membership</b>—The MAC address is learned on an interface that is not a member of the VLAN.</li> <li>• <b>Interface limit</b>—The number of MAC addresses learned on the interface has exceeded the limit.</li> <li>• <b>MAC move limit</b>—This MAC address has moved among multiple interfaces too many times in a given interval.</li> <li>• <b>VLAN limit</b>—The number of MAC addresses learned on the VLAN has exceeded the limit.</li> <li>• <b>VLAN membership limit</b>—The number of MAC addresses learned on the interface as a member of the specified VLAN (VLAN membership MAC limit) has exceeded the limit.</li> <li>• <b>Invalid VLAN index</b>—The VLAN of the packet, although configured, does not yet exist in the kernel.</li> <li>• <b>Interface not learning</b>—The MAC address is learned on an interface that does not yet allow learning—for example, the interface is blocked.</li> <li>• <b>No nexthop</b>—The MAC address is learned on an interface that does not have a unicast next hop.</li> <li>• <b>MAC learning disabled</b>—The MAC address is learned on an interface on which MAC learning has been disabled.</li> <li>• <b>Others</b>—The message contains some other error.</li> </ul> | All levels      |

## Sample Output

### show ethernet-switching statistics mac-learning

```
user@switch> show ethernet-switching statistics mac-learning
```

```
Learning stats: 0 learn msg rcvd, 0 error
Interface Local pkts Transit pkts Error
ge-0/0/0.0 0 0 0
ge-0/0/1.0 0 0 0
ge-0/0/2.0 0 0 0
ge-0/0/3.0 0 0 0
```

**show ethernet-switching statistics mac-learning detail**

```
user@switch> show ethernet-switching statistics mac-learning detail
Learning stats: 0 learn msg rcvd, 0 error
```

```
Interface: ge-0/0/0.0
Learning message from local packets: 0
Learning message from transit packets: 1
Learning message with error: 0
 Invalid VLAN: 0 Invalid MAC: 0
 Security violation: 0 Interface down: 0
 Incorrect membership: 0 Interface limit: 0
 MAC move limit: 0 VLAN limit: 0
 Invalid VLAN index: 0 Interface not learning: 0
 No nexthop: 0 MAC learning disabled: 0
 Others: 0
```

```
Interface: ge-0/0/1.0
Learning message from local packets: 0
Learning message from transit packets: 2
Learning message with error: 0
 Invalid VLAN: 0 Invalid MAC: 0
 Security violation: 0 Interface down: 0
 Incorrect membership: 0 Interface limit: 0
 MAC move limit: 0 VLAN limit: 0
 Invalid VLAN index: 0 Interface not learning: 0
 No nexthop: 0 MAC learning disabled: 0
 Others: 0
```

**show ethernet-switching statistics mac-learning interface ge-0/0/28 detail**

```
user@switch> show ethernet-switching statistics mac-learning interface ge-0/0/28 detail
```

```
Interface: ge-0/0/28.0
Learning message from local packets: 0
Learning message from transit packets: 5
Learning message with error: 0
 Invalid VLAN: 0 Invalid MAC: 0
 Security violation: 0 Interface down: 0
 Incorrect membership: 0 Interface limit: 0
 MAC move limit: 0 VLAN limit: 0
 VLAN membership limit: 20
 Invalid VLAN index: 0 Interface not learning: 0
 No nexthop: 0 MAC learning disabled: 0
 Others: 0
```

**show ethernet-switching statistics mac-learning interface**

```
user@switch> show ethernet-switching statistics mac-learning interface ge-0/0/1
```

| Interface  | Local pkts | Transit pkts | Error |
|------------|------------|--------------|-------|
| ge-0/0/1.0 | 0          | 1            | 1     |

**show ethernet-switching statistics mac-learning detail (QFX Series)**

```
user@switch> show ethernet-switching statistics mac-learning detail
Learning stats: 0 learn msg rcvd, 0 error
```

```
Interface: xe-0/0/0.0
Learning message from local packets: 0
Learning message from transit packets: 1
Learning message with error: 0
```

|                       |   |                         |   |
|-----------------------|---|-------------------------|---|
| Invalid VLAN:         | 0 | Invalid MAC:            | 0 |
| Security violation:   | 0 | Interface down:         | 0 |
| Incorrect membership: | 0 | Interface limit:        | 0 |
| MAC move limit:       | 0 | VLAN limit:             | 0 |
| Invalid VLAN index:   | 0 | Interface not learning: | 0 |
| No nexthop:           | 0 | MAC learning disabled:  | 0 |
| Others:               | 0 |                         |   |

Interface: xe-0/0/1.0

Learning message from local packets: 0

Learning message from transit packets: 2

Learning message with error: 0

|                       |   |                         |   |
|-----------------------|---|-------------------------|---|
| Invalid VLAN:         | 0 | Invalid MAC:            | 0 |
| Security violation:   | 0 | Interface down:         | 0 |
| Incorrect membership: | 0 | Interface limit:        | 0 |
| MAC move limit:       | 0 | VLAN limit:             | 0 |
| Invalid VLAN index:   | 0 | Interface not learning: | 0 |
| No nexthop:           | 0 | MAC learning disabled:  | 0 |
| Others:               | 0 |                         |   |





## CHAPTER 87

# Spanning Tree Monitoring Commands

- clear error bpdu interface
- clear spanning-tree statistics
- show spanning-tree bridge
- show spanning-tree interface
- show spanning-tree mstp configuration
- show spanning-tree statistics

## clear error bpdu interface

---


|                                 |                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>clear error bpdu interface (all   <i>interface-name</i>)</code>                                                                                                                                                                              |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.4.<br>Command introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Command supports <b>all</b> option in Junos OS Release 15.1 for EX Series switches.                                   |
| <b>Description</b>              | Clear a bridge protocol data unit (BPDU) error condition caused by the detection of a possible bridging loop from Spanning Tree Protocol (STP) operation.                                                                                          |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring BPDU Protection on Spanning Tree Interfaces on page 2205</a></li><li>• <a href="#">Unblocking an Interface That Receives BPDUs in Error (CLI Procedure) on page 2217</a></li></ul> |
| <b>List of Sample Output</b>    | <a href="#">clear error bpdu interface on page 2418</a>                                                                                                                                                                                            |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                              |

### Sample Output

#### clear error bpdu interface

```
user@switch> clear error bpdu interface ge-1/1/1
```

## clear spanning-tree statistics

|                                                                                                                                                                                     |                                                                                                                                                                                                                                                                                                                         |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                                                                                                                                                               | <a href="#">Syntax on page 2419</a><br><a href="#">Syntax (EX Series Switches and the QFX Series) on page 2419</a>                                                                                                                                                                                                      |
| <b>Syntax</b>                                                                                                                                                                       | clear spanning-tree statistics<br><interface <i>interface-name</i> ><br><logical-system <i>logical-system-name</i> >                                                                                                                                                                                                    |
| <b>Syntax (EX Series Switches and the QFX Series)</b>                                                                                                                               | clear spanning-tree statistics<br><interface <i>interface-name</i> >                                                                                                                                                                                                                                                    |
| <b>Release Information</b>                                                                                                                                                          | Command introduced in Junos OS Release 8.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                    |
| <b>Description</b>                                                                                                                                                                  | Clear Spanning Tree Protocol statistics.                                                                                                                                                                                                                                                                                |
| <b>Options</b>                                                                                                                                                                      | <b>none</b> —Reset STP counters for all interfaces for all routing instances.<br><br><b>interface <i>interface-name</i></b> —(Optional) Clear STP statistics for the specified interface only.<br><br><b>logical-system <i>logical-system-name</i></b> —(Optional) Clear STP statistics on a particular logical system. |
| <div>  <b>NOTE:</b> The <b>logical-system</b> option is not available on QFabric systems. </div> |                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b>                                                                                                                                                     | clear                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>                                                                                                                                                        | <ul style="list-style-type: none"> <li><a href="#">show spanning-tree statistics on page 2433</a></li> </ul>                                                                                                                                                                                                            |
| <b>List of Sample Output</b>                                                                                                                                                        | <a href="#">clear stp statistics on page 2419</a>                                                                                                                                                                                                                                                                       |

### Sample Output

#### clear stp statistics

```
user@host> clear stp statistics
```

## show spanning-tree bridge

**List of Syntax** [Syntax on page 2420](#)  
[Syntax \(QFX Series\) on page 2420](#)

**Syntax** `show spanning-tree bridge`  
`<brief | detail>`  
`<msti msti-id>`  
`<routing-instance routing-instance-name>`  
`<vlan-id vlan-id>`

**Syntax (QFX Series)** `show spanning-tree bridge`  
`<brief | detail>`  
`<msti msti-id>`  
`<vlan-id vlan-id>`

**Release Information** Command introduced in Junos OS Release 8.4.  
 Command introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Display the configured or calculated Spanning Tree Protocol (STP) parameters.

**Options** **none**—(Optional) Display brief STP bridge information for all multiple spanning-tree instances (MSTIs).

**brief | detail**—(Optional) Display the specified level of output.

**msti *msti-id***—(Optional) Display STP bridge information for the specified MSTI.

**routing-instance *routing-instance-name***—(Optional) Display STP bridge information for the specified routing instance.

**vlan-id *vlan-id***—(Optional) Display STP bridge information for the specified VLAN.

**Required Privilege Level** view

**List of Sample Output** [show spanning-tree bridge routing-instance on page 2421](#)  
[show spanning-tree bridge msti on page 2422](#)  
[show spanning-tree bridge vlan-id \(MSTP\) on page 2423](#)  
[show spanning-tree bridge \(RSTP\) on page 2423](#)  
[show spanning-tree bridge vlan-id \(RSTP\) on page 2424](#)

**Output Fields** [Table 203](#) lists the output fields for the **show spanning-tree bridge** command. Output fields are listed in the approximate order in which they appear.

**Table 203: show spanning-tree bridge Output Fields**

| Field Name            | Field Description                                                  |
|-----------------------|--------------------------------------------------------------------|
| Routing instance name | Name of the routing instance under which the bridge is configured. |
| Enabled protocol      | Spanning Tree Protocol type enabled.                               |

Table 203: show spanning-tree bridge Output Fields (*continued*)

| Field Name                             | Field Description                                                                                                                               |
|----------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Root ID</b>                         | Bridge ID of the elected spanning-tree root bridge. The bridge ID consists of a configurable bridge priority and the MAC address of the bridge. |
| <b>Root cost</b>                       | Calculated cost to reach the root bridge from the bridge where the command is entered.                                                          |
| <b>Root port</b>                       | Interface that is the current elected root port for this bridge.                                                                                |
| <b>CIST regional root</b>              | Bridge ID of the elected MSTP regional root bridge.                                                                                             |
| <b>CIST internal root cost</b>         | Calculated cost to reach the regional root bridge from the bridge where the command is entered.                                                 |
| <b>Hello time</b>                      | Configured number of seconds between transmissions of configuration BPDUs.                                                                      |
| <b>Maximum age</b>                     | Configured maximum expected arrival time of hello bridge protocol data units (BPDUs).                                                           |
| <b>Forward delay</b>                   | How long an STP bridge port remains in the listening and learning states before transitioning to the forwarding state.                          |
| <b>Hop count</b>                       | Configured maximum number of hops a BPDU can be forwarded in the MSTP region.                                                                   |
| <b>Message age</b>                     | Number of elapsed seconds since the most recent BPDU was received.                                                                              |
| <b>Number of topology changes</b>      | Total number of STP topology changes detected since the routing device last booted.                                                             |
| <b>Time since last topology change</b> | Number of elapsed seconds since the most recent topology change.                                                                                |
| <b>Bridge ID (Local)</b>               | Locally configured bridge ID. The bridge ID consists of a configurable bridge priority and the MAC address of the bridge.                       |
| <b>Extended system ID</b>              | System identifier.                                                                                                                              |
| <b>MSTI regional root</b>              | Bridge ID of the elected MSTP regional root bridge.                                                                                             |

## Sample Output

### show spanning-tree bridge routing-instance

```

user@host> show spanning-tree bridge routing-instance vs1 detail
STP bridge parameters
Routing instance name : vs1
Enabled protocol : MSTP

```

```
STP bridge parameters for CIST
 Root ID : 32768.00:13:c3:9e:c8:80
 Root cost : 0
 Root port : ge-10/2/0
 CIST regional root : 32768.00:13:c3:9e:c8:80
 CIST internal root cost : 22000
 Hello time : 2 seconds
 Maximum age : 20 seconds
 Forward delay : 15 seconds
 Hop count : 18
 Message age : 0
 Number of topology changes : 1
 Time since last topology change : 1191 seconds
 Local parameters
 Bridge ID : 32768.00:90:69:0b:7f:d1
 Extended system ID : 1

STP bridge parameters for MSTI 1
 MSTI regional root : 32769.00:13:c3:9e:c8:80
 Root cost : 22000
 Root port : ge-10/2/0
 Hello time : 2 seconds
 Maximum age : 20 seconds
 Forward delay : 15 seconds
 Hop count : 18
 Number of topology changes : 1
 Time since last topology change : 1191 seconds
 Local parameters
 Bridge ID : 32769.00:90:69:0b:7f:d1
 Extended system ID : 1

STP bridge parameters for MSTI 2
 MSTI regional root : 32770.00:13:c3:9e:c8:80
 Root cost : 22000
 Root port : ge-10/2/0
 Hello time : 2 seconds
 Maximum age : 20 seconds
 Forward delay : 15 seconds
 Hop count : 18
 Number of topology changes : 1
 Time since last topology change : 1191 seconds
 Local parameters
 Bridge ID : 32770.00:90:69:0b:7f:d1
 Extended system ID : 1
```

### show spanning-tree bridge msti

```
user@host> show spanning-tree bridge msti 1 routing-instance vs1 detail
STP bridge parameters
Routing instance name : vs1
Enabled protocol : MSTP

STP bridge parameters for MSTI 1
 MSTI regional root : 32769.00:13:c3:9e:c8:80
 Root cost : 22000
 Root port : xe-10/2/0
 Hello time : 2 seconds
 Maximum age : 20 seconds
 Forward delay : 15 seconds
 Hop count : 18
```

```

Number of topology changes : 1
Time since last topology change : 1191 seconds
Local parameters
 Bridge ID : 32769.00:90:69:0b:7f:d1
 Extended system ID : 1

```

### show spanning-tree bridge vlan-id (MSTP)

```
user@host> show spanning-tree bridge vlan-id 1101 routing-instance vs1 detail
```

```

STP bridge parameters
Routing instance name : vs1
Enabled protocol : MSTP

STP bridge parameters for CIST
Root ID : 32768.00:13:c3:9e:c8:80
Root cost : 0
Root port : xe-10/2/0
CIST regional root : 32768.00:13:c3:9e:c8:80
CIST internal root cost : 22000
Hello time : 2 seconds
Maximum age : 20 seconds
Forward delay : 15 seconds
Hop count : 18
Message age : 0
Number of topology changes : 0
Local parameters
 Bridge ID : 32768.00:90:69:0b:7f:d1
 Extended system ID : 1
 Hello time : 2 seconds
 Maximum age : 20 seconds
 Forward delay : 15 seconds
 Path cost method : 32 bit
 Maximum hop count : 20

```

### show spanning-tree bridge (RSTP)

```
user@host> show spanning-tree bridge
```

```

STP bridge parameters
Routing instance name : GLOBAL
Enabled protocol : RSTP
Root ID : 28672.00:90:69:0b:3f:d0
Hello time : 2 seconds
Maximum age : 20 seconds
Forward delay : 15 seconds
Message age : 0
Number of topology changes : 58
Time since last topology change : 14127 seconds
Local parameters
 Bridge ID : 28672.00:90:69:0b:3f:d0
 Extended system ID : 0

STP bridge parameters for bridge VLAN 10
Root ID : 28672.00:90:69:0b:3f:d0
Hello time : 2 seconds
Maximum age : 20 seconds
Forward delay : 15 seconds
Message age : 0
Number of topology changes : 58
Time since last topology change : 14127 seconds
Local parameters
 Bridge ID : 28672.00:90:69:0b:3f:d0

```

```
Extended system ID : 0

STP bridge parameters for bridge VLAN 20
Root ID : 28672.00:90:69:0b:3f:d0
Hello time : 2 seconds
Maximum age : 20 seconds
Forward delay : 15 seconds
Message age : 0
Number of topology changes : 58
Time since last topology change : 14127 seconds
Local parameters
 Bridge ID : 28672.00:90:69:0b:3f:d0
 Extended system ID : 0
```

#### show spanning-tree bridge vlan-id (RSTP)

```
user@host> show spanning-tree bridge vlan-id 10
STP bridge parameters
Routing instance name : GLOBAL
Enabled protocol : RSTP

STP bridge parameters for VLAN 10
Root ID : 28672.00:90:69:0b:3f:d0
Hello time : 2 seconds
Maximum age : 20 seconds
Forward delay : 15 seconds
Message age : 0
Number of topology changes : 58
Time since last topology change : 14127 seconds
Local parameters
 Bridge ID : 28672.00:90:69:0b:3f:d0
 Extended system ID : 0
```



## show spanning-tree interface

|                                                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                                 | <a href="#">Syntax on page 2425</a><br><a href="#">Syntax (EX Series Switches and the QFX Series) on page 2425</a>                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Syntax</b>                                         | <pre>show spanning-tree interface &lt;brief   detail&gt; &lt;msti <i>msti-id</i>&gt; &lt;routing-instance <i>routing-instance-name</i>&gt; &lt;vlan-id <i>vlan-id</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                  |
| <b>Syntax (EX Series Switches and the QFX Series)</b> | <pre>show spanning-tree interface &lt;brief   detail&gt; &lt;msti <i>msti-id</i>&gt; &lt;vlan-id <i>vlan-id</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>                            | <p>Command introduced in Junos OS Release 8.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>                                    | Display the configured or calculated interface-level STP parameters.                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                                        | <p><b>none</b>—Display brief STP interface information.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>msti <i>msti-id</i></b>—(Optional) Display STP interface information for the specified MST instance.</p> <p><b>routing-instance <i>routing-instance-name</i></b>—(Optional) Display STP interface information for the specified routing instance.</p> <p><b>vlan-id <i>vlan-id</i></b>—(Optional) Display STP interface information for the specified VLAN.</p> |
| <b>Required Privilege Level</b>                       | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>List of Sample Output</b>                          | <a href="#">show spanning-tree interface on page 2426</a><br><a href="#">show spanning-tree interface (QFX Series) on page 2427</a><br><a href="#">show spanning-tree interface detail on page 2427</a><br><a href="#">show spanning-tree interface msti on page 2429</a><br><a href="#">show spanning-tree interface vlan-id on page 2429</a><br><a href="#">show spanning-tree interface (VSTP) on page 2430</a><br><a href="#">show spanning-tree interface vlan-id (VSTP) on page 2430</a>                   |
| <b>Output Fields</b>                                  | Table 204 lists the output fields for the <b>show spanning-tree interface</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                   |

**Table 204: show spanning-tree Interface Output Fields**

| Field Name            | Field Description                                                             |
|-----------------------|-------------------------------------------------------------------------------|
| <b>Interface name</b> | Interface configured to participate in the STP, RSTP, VSTP, or MSTP instance. |

Table 204: show spanning-tree Interface Output Fields (*continued*)

| Field Name                  | Field Description                                                                                                                                                     |
|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Port ID</b>              | Logical interface identifier configured to participate in the MSTP or VSTP instance.                                                                                  |
| <b>Designated port ID</b>   | Port ID of the designated port for the LAN segment to which this interface is attached.                                                                               |
| <b>Designated bridge ID</b> | Bridge ID of the designated bridge for the LAN segment to which this interface is attached.                                                                           |
| <b>Port Cost</b>            | Configured cost for the interface.                                                                                                                                    |
| <b>Port State</b>           | STP port state: forwarding ( <b>FWD</b> ), blocking ( <b>BLK</b> ), listening, learning, or disabled.                                                                 |
| <b>Port Role</b>            | MSTP, VSTP, or RSTP port role: designated ( <b>DESG</b> ), backup ( <b>BKUP</b> ), alternate ( <b>ALT</b> ), ( <b>ROOT</b> ), or Root Prevented ( <b>Root-Prev</b> ). |
| <b>Link type</b>            | MSTP, VSTP, or RSTP link type. Shared or point-to-point (pt-pt) and edge or nonedge.                                                                                  |
| <b>Alternate</b>            | Identifies the interface as an MSTP, VSTP, or RSTP alternate root port ( <b>Yes</b> ) or nonalternate root port ( <b>No</b> ).                                        |
| <b>Boundary Port</b>        | Identifies the interface as an MSTP regional boundary port ( <b>Yes</b> ) or nonboundary port ( <b>No</b> ).                                                          |

## Sample Output

### show spanning-tree interface

```
user@host> show spanning-tree interface routing-instance vs1 detail
Spanning tree interface parameters for instance 0
```

| Interface | Port ID | Designated port ID | Designated bridge ID | Port Cost | State | Role |
|-----------|---------|--------------------|----------------------|-----------|-------|------|
| ae1       | 128:1   | 128:1              | 32768.0090690b47d1   | 1000      | FWD   | DESG |
| ge-2/1/2  | 128:2   | 128:2              | 32768.0090690b47d1   | 20000     | FWD   | DESG |
| ge-2/1/5  | 128:3   | 128:3              | 32768.0090690b47d1   | 29999     | FWD   | DESG |
| ge-2/2/1  | 128:4   | 128:26             | 32768.0013c39ec880   | 20000     | FWD   | ROOT |
| xe-9/2/0  | 128:5   | 128:5              | 32768.0090690b47d1   | 2000      | FWD   | DESG |
| xe-9/3/0  | 128:6   | 128:6              | 32768.0090690b47d1   | 2000      | FWD   | DESG |

```
Spanning tree interface parameters for instance 1
```

| Interface | Port ID | Designated port ID | Designated bridge ID | Port Cost | State | Role |
|-----------|---------|--------------------|----------------------|-----------|-------|------|
| ae1       | 128:1   | 128:1              | 32769.0090690b47d1   | 1000      | FWD   | DESG |
| ge-2/1/2  | 128:2   | 128:2              | 32769.0090690b47d1   | 20000     | FWD   | DESG |
| ge-2/1/5  | 128:3   | 128:3              | 32769.0090690b47d1   | 29999     | FWD   | DESG |
| ge-2/2/1  | 128:4   | 128:26             | 32769.0013c39ec880   | 20000     | FWD   | ROOT |
| xe-9/2/0  | 128:5   | 128:5              | 32769.0090690b47d1   | 2000      | FWD   | DESG |
| xe-9/3/0  | 128:6   | 128:6              | 32769.0090690b47d1   | 2000      | FWD   | DESG |

Spanning tree interface parameters for instance 2

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------|-------|------|
| ae1       | 128:1   | 128:1                 | 32770.0090690b47d1      | 1000         | FWD   | DESG |
| ge-2/1/2  | 128:2   | 128:2                 | 32770.0090690b47d1      | 20000        | FWD   | DESG |
| ge-2/1/5  | 128:3   | 128:3                 | 32770.0090690b47d1      | 29999        | FWD   | DESG |
| ge-2/2/1  | 128:4   | 128:26                | 32770.0013c39ec880      | 20000        | FWD   | ROOT |
| xe-9/2/0  | 128:5   | 128:5                 | 32770.0090690b47d1      | 2000         | FWD   | DESG |
| xe-9/3/0  | 128:6   | 128:6                 | 32770.0090690b47d1      | 2000         | FWD   | DESG |

### show spanning-tree interface (QFX Series)

```
user@1f0> show spanning-tree interface routing-instance vs1 detail
Spanning tree interface parameters for instance 0
```

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------|-------|------|
| ae1       | 128:1   | 128:1                 | 32768.0090690b47d1      | 1000         | FWD   | DESG |
| ge-2/1/2  | 128:2   | 128:2                 | 32768.0090690b47d1      | 20000        | FWD   | DESG |
| ge-2/1/5  | 128:3   | 128:3                 | 32768.0090690b47d1      | 29999        | FWD   | DESG |
| ge-2/2/1  | 128:4   | 128:26                | 32768.0013c39ec880      | 20000        | FWD   | ROOT |
| xe-9/2/0  | 128:5   | 128:5                 | 32768.0090690b47d1      | 2000         | FWD   | DESG |
| xe-9/3/0  | 128:6   | 128:6                 | 32768.0090690b47d1      | 2000         | FWD   | DESG |

Spanning tree interface parameters for instance 1

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------|-------|------|
| ae1       | 128:1   | 128:1                 | 32769.0090690b47d1      | 1000         | FWD   | DESG |
| ge-2/1/2  | 128:2   | 128:2                 | 32769.0090690b47d1      | 20000        | FWD   | DESG |
| ge-2/1/5  | 128:3   | 128:3                 | 32769.0090690b47d1      | 29999        | FWD   | DESG |
| ge-2/2/1  | 128:4   | 128:26                | 32769.0013c39ec880      | 20000        | FWD   | ROOT |
| xe-9/2/0  | 128:5   | 128:5                 | 32769.0090690b47d1      | 2000         | FWD   | DESG |
| xe-9/3/0  | 128:6   | 128:6                 | 32769.0090690b47d1      | 2000         | FWD   | DESG |

Spanning tree interface parameters for instance 2

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------|-------|------|
| ae1       | 128:1   | 128:1                 | 32770.0090690b47d1      | 1000         | FWD   | DESG |
| ge-2/1/2  | 128:2   | 128:2                 | 32770.0090690b47d1      | 20000        | FWD   | DESG |
| ge-2/1/5  | 128:3   | 128:3                 | 32770.0090690b47d1      | 29999        | FWD   | DESG |
| ge-2/2/1  | 128:4   | 128:26                | 32770.0013c39ec880      | 20000        | FWD   | ROOT |
| xe-9/2/0  | 128:5   | 128:5                 | 32770.0090690b47d1      | 2000         | FWD   | DESG |
| xe-9/3/0  | 128:6   | 128:6                 | 32770.0090690b47d1      | 2000         | FWD   | DESG |

### show spanning-tree interface detail

```
user@host> show spanning-tree interface routing-instance vs1 detail
Spanning tree interface parameters for instance 0
```

```
Interface name : ae1
Port identifier : 128.1
Designated port ID : 128.1
Port cost : 1000
Port state : Forwarding
Designated bridge ID : 32768.00:90:69:0b:47:d1
Port role : Designated
Link type : Pt-Pt/NONEDGE
```

```
Boundary port : No

Interface name : ge-2/1/2
Port identifier : 128.2
Designated port ID : 128.2
Port cost : 20000
Port state : Forwarding
Designated bridge ID : 32768.00:90:69:0b:47:d1
Port role : Designated
Link type : Pt-Pt/NONEDGE
Boundary port : No

Interface name : ge-2/1/5
Port identifier : 128.3
Designated port ID : 128.3
Port cost : 29999
Port state : Forwarding
Designated bridge ID : 32768.00:90:69:0b:47:d1
Port role : Designated
Link type : Pt-Pt/NONEDGE
Boundary port : No

Interface name : ge-2/2/1
Port identifier : 128.4
Designated port ID : 128.26
Port cost : 20000
Port state : Forwarding
Designated bridge ID : 32768.00:13:c3:9e:c8:80
Port role : Root
Link type : Pt-Pt/NONEDGE
Boundary port : No

Interface name : xe-9/2/0
Port identifier : 128.5
Designated port ID : 128.5
Port cost : 2000
Port state : Forwarding
Designated bridge ID : 32768.00:90:69:0b:47:d1
Port role : Designated
Link type : Pt-Pt/NONEDGE
Boundary port : No

Interface name : xe-9/3/0
Port identifier : 128.6
Designated port ID : 128.6
Port cost : 2000
Port state : Forwarding
Designated bridge ID : 32768.00:90:69:0b:47:d1
Port role : Designated
Link type : Pt-Pt/NONEDGE
Boundary port : No
```

#### Spanning tree interface parameters for instance 1

```
Interface name : ae1
Port identifier : 128.1
Designated port ID : 128.1
Port cost : 1000
Port state : Forwarding
Designated bridge ID : 32768.00:90:69:0b:47:d1
```

```

Port role : Designated
Link type : Pt-Pt/NONEDGE
Boundary port : No

Interface name : ge-2/1/2
Port identifier : 128.2
Designated port ID : 128.2
Port cost : 20000
Port state : Forwarding
Designated bridge ID : 32768.00:90:69:0b:47:d1
Port role : Designated
Link type : Pt-Pt/NONEDGE
Boundary port : No

Interface name : ge-2/1/5
Port identifier : 128.3
Designated port ID : 128.3
Port cost : 29999
Port state : Forwarding
Designated bridge ID : 32768.00:90:69:0b:47:d1
Port role : Designated
Link type : Pt-Pt/NONEDGE
Boundary port : No

Interface name : ge-2/2/1
Port identifier : 128.4
Designated port ID : 128.26
Port cost : 20000
Port state : Forwarding
Designated bridge ID : 32768.00:13:c3:9e:c8:80
Port role : Root
Link type : Pt-Pt/NONEDGE
Boundary port : No

...

```

### show spanning-tree interface msti

```

user@host> show spanning-tree interface msti 1 routing-instance vs1 detail
Spanning tree interface parameters for instance 1

```

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------|-------|------|
| xe-7/0/0  | 128:1   | 128:1                 | 32769.0090690b4fd1      | 2000         | FWD   | DESG |
| ge-5/1/0  | 128:2   | 128:2                 | 32769.0090690b4fd1      | 20000        | FWD   | DESG |
| ge-5/1/1  | 128:3   | 128:3                 | 32769.0090690b4fd1      | 20000        | FWD   | DESG |
| ae1       | 128:4   | 128:1                 | 32769.0090690b47d1      | 10000        | BLK   | ALT  |
| ge-5/1/4  | 128:5   | 128:3                 | 32769.0090690b47d1      | 20000        | BLK   | ALT  |
| xe-7/2/0  | 128:6   | 128:6                 | 32769.0090690b47d1      | 2000         | FWD   | ROOT |

### show spanning-tree interface vlan-id

```

user@host> show spanning-tree interface vlan-id 101 routing-instance vs1 detail
Spanning tree interface parameters for instance 0

```

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-11/0/5 | 128:1   | 128:1                 | 32768.0090690b7fd1      | 20000        | FWD   | DESG |
| ge-11/0/6 | 128:2   | 128:1                 | 32768.0090690b7fd1      | 20000        | BLK   | BKUP |
| ge-11/1/0 | 128:3   | 128:2                 | 32768.0090690b4fd1      | 20000        | BLK   | ALT  |
| ge-11/1/1 | 128:4   | 128:3                 | 32768.0090690b4fd1      | 20000        | BLK   | ALT  |

|           |       |       |                    |       |     |      |
|-----------|-------|-------|--------------------|-------|-----|------|
| ge-11/1/4 | 128:5 | 128:1 | 32768.0090690b47d1 | 20000 | BLK | ALT  |
| xe-10/0/0 | 128:6 | 128:5 | 32768.0090690b4fd1 | 2000  | BLK | ALT  |
| xe-10/2/0 | 128:7 | 128:4 | 32768.0090690b47d1 | 2000  | FWD | ROOT |

**show spanning-tree interface (VSTP)**

```
user@host> show spanning-tree interface
```

```
Spanning tree interface parameters for instance 0
```

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Cost  | State | Role |
|-----------|---------|-----------------------|-------------------------|-------|-------|------|
| ge-1/0/1  | 128:1   | 128:1                 | 28672.0090690b3fe0      | 20000 | FWD   | DESG |
| ge-1/0/2  | 128:2   | 128:2                 | 28672.0090690b3fe0      | 20000 | FWD   | DESG |

```
Spanning tree interface parameters for VLAN 10
```

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Cost  | State | Role |
|-----------|---------|-----------------------|-------------------------|-------|-------|------|
| ge-1/0/1  | 128:1   | 128:1                 | 28672.0090690b3fe0      | 20000 | FWD   | DESG |
| ge-1/0/2  | 128:2   | 128:2                 | 28672.0090690b3fe0      | 20000 | FWD   | DESG |

```
Spanning tree interface parameters for VLAN 20
```

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Cost  | State | Role |
|-----------|---------|-----------------------|-------------------------|-------|-------|------|
| ge-1/0/1  | 128:1   | 128:1                 | 28672.0090690b3fe0      | 20000 | FWD   | DESG |
| ge-1/0/2  | 128:2   | 128:2                 | 28672.0090690b3fe0      | 20000 | FWD   | DESG |

**show spanning-tree interface vlan-id (VSTP)**

```
user@host> show spanning-tree interface vlan-id 10
```

```
Spanning tree interface parameters for VLAN 10
```

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Cost  | State | Role |
|-----------|---------|-----------------------|-------------------------|-------|-------|------|
| ge-1/0/1  | 128:1   | 128:1                 | 28672.0090690b3fe0      | 20000 | FWD   | DESG |
| ge-1/0/2  | 128:2   | 128:2                 | 28672.0090690b3fe0      | 20000 | FWD   | DESG |

## show spanning-tree mstp configuration

|                                                     |                                                                                                                                                                                                                                                                                         |
|-----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 2431</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 2431</a>                                                                                                                                                                        |
| <b>Syntax</b>                                       | show spanning-tree mstp configuration<br><brief   detail><br><routing-instance <i>routing-instance-name</i> >                                                                                                                                                                           |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | show spanning-tree mstp configuration<br><brief   detail>                                                                                                                                                                                                                               |
| <b>Release Information</b>                          | Command introduced in Junos OS Release 8.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                    |
| <b>Description</b>                                  | Display the MSTP configuration.                                                                                                                                                                                                                                                         |
| <b>Options</b>                                      | <b>none</b> —Display MSTP configuration information.<br><br><b>brief   detail</b> —(Optional) Display the specified level of output.<br><br><b>routing-instance <i>routing-instance-name</i></b> —(Optional) Display MSTP configuration information for the specified routing instance. |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                    |
| <b>List of Sample Output</b>                        | <a href="#">show spanning-tree mstp configuration detail on page 2432</a><br><a href="#">show spanning-tree mstp configuration detail (QFX Series) on page 2432</a>                                                                                                                     |
| <b>Output Fields</b>                                | Table 205 lists the output fields for the <b>show spanning-tree mstp configuration</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                 |

**Table 205: show spanning-tree mstp configuration Output Fields**

| Field Name                  | Field Description                                                |
|-----------------------------|------------------------------------------------------------------|
| <b>Context id</b>           | Internally generated identifier.                                 |
| <b>Region name</b>          | MSTP region name carried in the MSTP BPDUs.                      |
| <b>Revision</b>             | Revision number of the MSTP configuration.                       |
| <b>Configuration digest</b> | Numerical value derived from the VLAN-to-instance mapping table. |
| <b>MSTI</b>                 | MST instance identifier.                                         |
| <b>Member VLANs</b>         | VLAN identifiers associated with the MSTI.                       |

## Sample Output

### show spanning-tree mstp configuration detail

```
user@host> show spanning-tree mstp configuration routing-instance vs1 detail
MSTP configuration information
Context identifier : 1
Region name : henry
Revision : 3
Configuration digest : 0x6da4b5c4fd587757eef35675365e1

MSTI Member VLANs
 0 0-99,101-199,201-4094
 1 100
 2 200
```

### show spanning-tree mstp configuration detail (QFX Series)

```
user@1f0> show spanning-tree mstp configuration routing-instance vs1 detail
MSTP configuration information
Context identifier : 1
Region name : henry
Revision : 3
Configuration digest : 0x6da4b5c4fd587757eef35675365e1

MSTI Member VLANs
 0 0-99,101-199,201-4094
 1 100
 2 200
```



## show spanning-tree statistics

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                |
|-----------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 2433</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 2433</a>                                                                                                                                                                                                                                                               |
| <b>Syntax</b>                                       | <pre>show spanning-tree statistics &lt;brief   detail&gt; &lt;interface <i>interface-name</i>&gt; &lt;routing-instance <i>routing-instance-name</i>&gt;</pre>                                                                                                                                                                                                                  |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show spanning-tree statistics &lt;brief   detail&gt; &lt;interface <i>interface-name</i>   vlan <i>vlan-id</i>&gt;</pre>                                                                                                                                                                                                                                                  |
| <b>Release Information</b>                          | <p>Command introduced in Junos OS Release 8.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for QFX Series switches.</p>                                                                                                                                                                       |
| <b>Description</b>                                  | Display STP statistics.                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                                      | <p><b>none</b>—Display brief STP statistics.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display STP statistics for the specified interface.</p> <p><b>routing-instance <i>routing-instance-name</i></b>—(Optional) Display STP statistics for the specified routing instance.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                           |
| <b>List of Sample Output</b>                        | <a href="#">show spanning-tree statistics routing-instance on page 2434</a><br><a href="#">show spanning-tree statistics interface routing-instance detail on page 2434</a>                                                                                                                                                                                                    |
| <b>Output Fields</b>                                | Table 206 lists the output fields for the <b>show spanning-tree statistics</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                |

**Table 206: show spanning-tree statistics Output Fields**

| Field Name                      | Field Description                                     |
|---------------------------------|-------------------------------------------------------|
| Message type                    | Type of message being counted.                        |
| BPDUs sent                      | Total number of BPDUs sent.                           |
| BPDUs received                  | Total number of BPDUs received.                       |
| BPDUs sent in last interval     | Number of BPDUs sent within a specified interval.     |
| BPDUs received in last interval | Number of BPDUs received within a specified interval. |

Table 206: show spanning-tree statistics Output Fields (*continued*)

| Field Name                    | Field Description                                              |
|-------------------------------|----------------------------------------------------------------|
| <b>Interface</b>              | Interface for which the statistics are being displayed.        |
| <b>Next BPDU transmission</b> | Number of seconds until the next BPDU is scheduled to be sent. |

## Sample Output

### show spanning-tree statistics routing-instance

```
user@host> show spanning-tree statistics routing-instance vs1 detail
Routing instance level STP statistics
Message type : bpdus
BPDUs sent : 1396
BPDUs received : 1027
BPDUs sent in last interval : 5 (duration: 4 sec)
BPDUs received in last interval: 4 (duration: 4 sec)
```

### show spanning-tree statistics interface routing-instance detail

```
user@host> show spanning-tree statistics interface ge-11/1/4 routing-instance vs1 detail
Interface BPDUs sent BPDUs received Next BPDU
 transmission
ge-11/1/4 7 190 0
```

# OVSDB and VXLAN Feature Guide for QFX10000 Switches (VMware NSX)



## PART 36

# Overview

- [OVSDB and VXLAN Overview on page 2439](#)



# OVSDB and VXLAN Overview

- [Understanding the Junos OS Implementation of OVSDB and VXLAN in a VMware NSX for Multi-Hypervisor Environment on page 2439](#)
- [Understanding VXLANs on page 2442](#)
- [Understanding the OVSDB Protocol Running on Juniper Networks Devices on page 2445](#)
- [OVSDB Support on Juniper Networks Devices on page 2446](#)
- [OVSDB Schema for Physical Devices on page 2447](#)
- [Understanding How Layer 2 BUM and Layer 3 Routed Multicast Traffic Are Handled with OVSDB on page 2450](#)

## Understanding the Junos OS Implementation of OVSDB and VXLAN in a VMware NSX for Multi-Hypervisor Environment

Some Juniper Networks devices support Virtual Extensible LAN (VXLAN) and the Open vSwitch Database (OVSDB) management protocol. (See [“OVSDB Support on Juniper Networks Devices” on page 2446](#).) Support for VXLAN and OVSDB enables the Juniper Networks devices in a physical network to be integrated into a virtual network.

The implementation of VXLAN and OVSDB on Juniper Networks devices is supported in a VMware NSX for Multi-Hypervisor environment for the data center. [Table 207](#) outlines the components that compose this environment and products that are typically deployed for each component.

Table 207: NSX for Multi-Hypervisor Components and Related Products


| Component                       | Products                 |
|---------------------------------|--------------------------|
| Cloud management platform (CMP) | CloudStack               |
|                                 | OpenStack                |
|                                 | Custom CMP               |
| Network virtualization platform | NSX for Multi-Hypervisor |

Table 207: NSX for Multi-Hypervisor Components and Related Products (*continued*)

| Component                                    | Products                                                                                                                  |
|----------------------------------------------|---------------------------------------------------------------------------------------------------------------------------|
| Hypervisor                                   | Kernel-based Virtual Machine (KVM)<br>Red Hat<br>VMware ESXi<br>Xen<br>NOTE: Juniper Networks supports only KVM and ESXi. |
| Virtual switch                               | Open vSwitch (OVS)<br>NSX vSwitch                                                                                         |
| SDN controller                               | NSX for Multi-Hypervisor controller                                                                                       |
| Overlay protocol                             | VXLAN                                                                                                                     |
| Media access control (MAC) learning protocol | OVSDB                                                                                                                     |

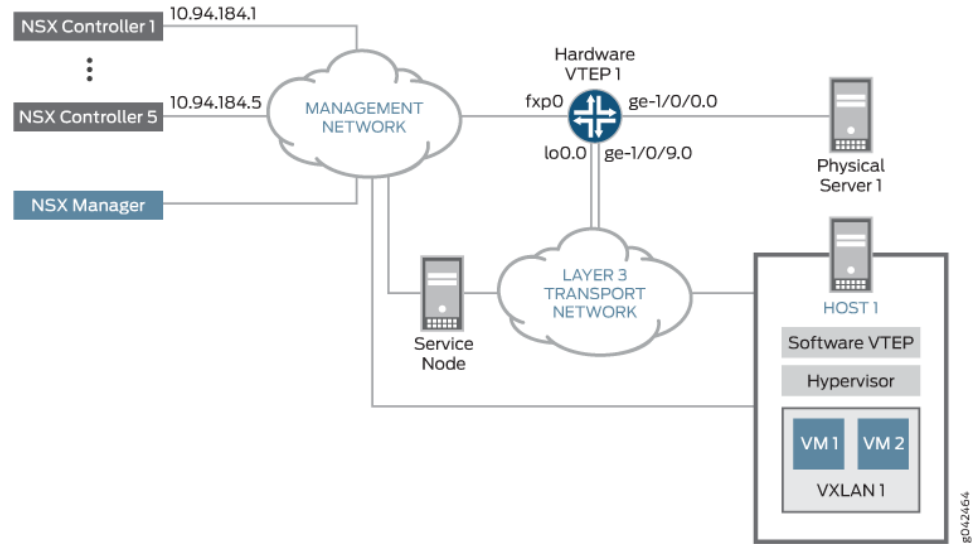
Figure 30 shows a high-level view of the NSX for Multi-Hypervisor platform architecture, while Figure 31 provides a more detailed representation of the components in the virtual and physical networks.

Figure 30: High-Level View of NSX for Multi-Hypervisor Architecture





**Figure 31: Integration of Juniper Networks Device into NSX for Multi-Hypervisor Environment**



In the data center topology shown in [Figure 31](#), the physical and virtual servers need to communicate. To facilitate this communication, a Juniper Networks device that supports VXLAN is strategically deployed so that it serves as a *gateway*, which is also known as a hardware virtual tunnel endpoint (VTEP), at the edge of the physical network. Working in conjunction with the software VTEP, which is deployed at the edge of the virtual network, the hardware VTEP encapsulates packets from resources on Physical Server 1 with a VXLAN header, and after the packets traverse the Layer 3 transport network, the software VTEP removes the VXLAN header from the packets and forwards the packets to the appropriate virtual machines (VMs). In essence, the encapsulation and de-encapsulation of packets by the hardware and software VTEPs enable the components in the physical and virtual networks to coexist without one needing to understand the workings of the other.

The same Juniper Networks device that acts as a hardware VTEP in [Figure 31](#) implements OVSDb, which enables this device to learn the MAC addresses of Physical Server 1 and other physical servers, and publish the addresses in the OVSDb schema, which was defined for physical devices. In the virtual network, one or more NSX controllers collect the MAC addresses of Host 1 and other virtual servers, and publish the addresses in the OVSDb schema. Using the OVSDb schema, components in the physical and virtual networks can exchange MAC addresses, as well as statistical information, enabling the components to learn about and reach each other in their respective networks.

**Related Documentation**

- [Understanding the OVSDb Protocol Running on Juniper Networks Devices on page 2445](#)
- [OVSDb Schema for Physical Devices](#)

## Understanding VXLANs

---

Virtual Extensible LAN protocol (VXLAN) technology allows networks to support more VLANs. According to the IEEE 802.1Q standard, traditional VLAN identifiers are 12 bits long—this naming limits networks to 4094 VLANs. The VXLAN protocol overcomes this limitation by using a longer logical network identifier that allows more VLANs and, therefore, more logical network isolation for large networks such as clouds that typically include many virtual machines.

- [VXLAN Benefits on page 2442](#)
- [How Does VXLAN Work? on page 2443](#)
- [VXLAN Implementation Methods on page 2443](#)
- [Using an MX Series Router or EX9200 Switch as a VTEP on page 2444](#)
- [Manual VXLANs Require PIM on page 2444](#)

### VXLAN Benefits

VXLAN technology allows you to segment your networks (as VLANs do), but it provides benefits that VLANs cannot. Here are the most important benefits of using VXLANs:

- You can theoretically create as many as 16 million VXLANs in an administrative domain (as opposed to 4094 VLANs on a Juniper Networks device).
  - MX Series routers and EX9200 switches support as many as 32,000 VXLANs, 32,000 multicast groups, and 8000 virtual tunnel endpoints (VTEPs). This means that VXLANs based on MX Series routers provide network segmentation at the scale required by cloud builders to support very large numbers of tenants.
  - QFX10000 Series switches support 4000 VXLANs and 2000 VTEPs.
- You can enable migration of virtual machines between servers that exist in separate Layer 2 domains by tunneling the traffic over Layer 3 networks. This functionality allows you to dynamically allocate resources within or between data centers without being constrained by Layer 2 boundaries or being forced to create large or geographically stretched Layer 2 domains.

Using VXLANs to create smaller Layer 2 domains that are connected over a Layer 3 network means that you do not need to use Spanning Tree Protocol (STP) to converge the topology but can use more-robust routing protocols in the Layer 3 network instead. In the absence of STP, none of your links are blocked, which means you can get full value from all the ports that you purchase. Using routing protocols to connect your Layer 2 domains also allows you to load-balance the traffic to ensure that you get the best use of your available bandwidth. Given the amount of east-west traffic that often flows within or between data centers, maximizing your network performance for that traffic is very important.

The video *Why Use an Overlay Network in a Data Center?* presents a brief overview of the advantages of using VXLANs.



Video: [Why Use an Overlay Network in a Data Center?](#)

## How Does VXLAN Work?

VXLAN is often described as an overlay technology because it allows you to stretch Layer 2 connections over an intervening Layer 3 network by encapsulating (tunneling) Ethernet frames in a VXLAN packet that includes IP addresses. Devices that support VXLANs are called virtual tunnel endpoints (VTEPs)—they can be end hosts or network switches or routers. VTEPs encapsulate VXLAN traffic and de-encapsulate that traffic when it leaves the VXLAN tunnel. To encapsulate an Ethernet frame, VTEPs add a number of fields, including the following:


- Outer media access control (MAC) destination address (MAC address of the tunnel endpoint VTEP)
- Outer MAC source address (MAC address of the tunnel source VTEP)
- Outer IP destination address (IP address of the tunnel endpoint VTEP)
- Outer IP source address (IP address of the tunnel source VTEP)
- Outer UDP header
- A VXLAN header that includes a 24-bit field—called the VXLAN network identifier (VNI)—that is used to uniquely identify the VXLAN. The VNI is similar to a VLAN ID, but having 24 bits allows you to create many more VXLANs than VLANs.



**NOTE:** Because VXLAN adds 50 to 54 bytes of additional header information to the original Ethernet frame, you might want to increase the MTU of the underlying network. In this case, configure the MTU of the physical interfaces that participate in the VXLAN network, not the MTU of the logical VTEP source interface, which is ignored.

Figure 32 shows the VXLAN packet format.

**Figure 32: VXLAN Packet Format**



## VXLAN Implementation Methods

Junos OS supports implementing VXLANs in the following environments:

- Without software-defined networking (SDN) controllers. VXLANs implemented in this type of environment are known as manual VXLANs.
- With SDN controllers. In this environment, SDN controllers use the Open vSwitch Database (OVSDB) management protocol to provide a means through which controllers (such as a VMware NSX or Juniper Contrail controller) and Juniper Networks devices that support OVSDB can communicate.



**NOTE:** If you want to implement VXLANs on a QFX10000 Series switch, you must use an SDN controller.

---

## Using an MX Series Router or EX9200 Switch as a VTEP

You can configure an MX Series router or EX9200 switch to act as a VTEP and perform all of the following roles:

- Act as a Layer 2 gateway between virtualized and nonvirtualized networks in the same data center or between data centers. For example, you can use an MX Series router to connect a network that uses VXLANs to one that uses VLANs.
- Act as a Layer 2 gateway between virtualized networks in the same or different data centers and allow virtual machines to move (VMotion) between those networks and data centers.
- Act as a Layer 3 gateway to route traffic between different VXLANs in the same data center.
- Act as a Layer 3 gateway to route traffic between different VXLANs in different data centers over a WAN or the Internet using standard routing protocols or virtual private LAN service (VPLS) tunnels.



**NOTE:** If you want an MX Series router or an EX9200 switch to be a VXLAN Layer 3 gateway, you must configure integrated routing and bridging (IRB) interfaces to connect the VXLANs, just as you do if you want to route traffic between VLANs.

---

## Manual VXLANs Require PIM

In an environment with a controller (such as a VMware NSX or Juniper Contrail controller), you can provision VXLANs on a Juniper Networks device. A controller also provides a control plane that VTEPs use to advertise their reachability and learn about the reachability of other VTEPs. You can also manually create VXLANs on Juniper Networks devices instead of using a controller. If you use this approach, you must also configure Protocol Independent Multicast (PIM) on the VTEPs so that they can create VXLAN tunnels between themselves.

You must also configure each VTEP in a given VXLAN to be a member of the same multicast group. (If possible, you should assign a different multicast group address to each VXLAN, though this is not required. Multiple VXLANs can share the same multicast

group.) The VTEPs can then forward ARP requests they receive from their connected hosts to the multicast group. The other VTEPs in the group de-encapsulate the VXLAN information, and (assuming they are members of the same VXLAN) they forward the ARP request to their connected hosts. When the target host receives the ARP request, it responds with its MAC address, and its VTEP forwards this ARP reply back to the source VTEP. Through this process, the VTEPs learn the IP addresses of the other VTEPs in the VXLAN and the MAC addresses of the hosts connected to the other VTEPs.

The multicast groups and trees are also used to forward broadcast, unknown unicast, and multicast (BUM) traffic between VTEPs. This prevents BUM traffic from being unnecessarily flooded outside the VXLAN.



**NOTE:** Multicast traffic that is forwarded through a VXLAN tunnel is sent only to the remote VTEPs in the VXLAN. That is, the encapsulating VTEP does not copy and send copies of the packets according to the multicast tree—it only forwards the received multicast packets to the remote VTEPs. The remote VTEPs de-encapsulate the encapsulated multicast packets and forward them to the appropriate Layer 2 interfaces. The remote VTEPs also do not copy and send copies of the packets according to the multicast tree.

**Related  
Documentation**

- [OVSDb Support on Juniper Networks Devices on page 2446](#)
- [mtu on page 3024](#)

## Understanding the OVSDb Protocol Running on Juniper Networks Devices

The Juniper Networks Junos OS implementation of the Open vSwitch Database (OVSDb) management protocol provides a means through which Juniper Networks devices that support OVSDb can communicate with software-defined networking (SDN) controllers. Juniper Networks devices exchange control and statistical information with the SDN controllers, thereby enabling virtual machine (VM) traffic from the entities in a virtualized network to be forwarded to entities in a physical network, and vice versa.

The Junos OS implementation of OVSDb includes an OVSDb server and an OVSDb client, both of which run on each Juniper Networks device that supports OVSDb.

The OVSDb server on a Juniper Networks device can communicate with an OVSDb client on an SDN controller. To establish a connection between a Juniper Networks device and an SDN controller, you must specify information about the SDN controller (IP address) and the connection (port over which the connection occurs and the communication protocol to be used) on each Juniper Networks device. After the configuration is successfully committed, the connection is established between the management port of the Juniper Networks device and the SDN controller port that you specify in the Junos OS configuration.

The OVSDb server stores and maintains an OVSDb database schema, which is defined for physical devices. This schema contains control and statistical information provided by the OVSDb client on the Juniper Networks devices and on SDN controllers. This

information is stored in various tables in the schema. The OVSDb client monitors the schema for additions, deletions, and modifications to this information, and the information is used for various purposes, such as learning the media access control (MAC) addresses of virtual hosts and physical servers.

The schema provides a means through which the Juniper Networks devices and the SDN controllers can exchange information. For example, the Juniper Networks devices capture MAC routes to entities in the physical network and push this information to a table in the schema so that SDN controllers with connections to these Juniper Networks devices can access the MAC routes. Conversely, SDN controllers capture MAC routes to entities in the virtualized network and push this information to a table in the schema so that Juniper Networks devices with connections to the SDN controllers can access the MAC routes.

Some of the OVSDb table names include the words *local* or *remote*, for example, *unicast MACs local table* and *unicast MACs remote table*. Information in *local* tables is learned by a Juniper Networks device that functions as a hardware virtual tunnel endpoint (VTEP), while information in *remote* tables is learned from other software or hardware VTEPs.

## OVSDb Support on Juniper Networks Devices

Table 208 lists the Juniper Networks devices that support the Open vSwitch Database (OVSDb) management protocol and the Junos OS releases in which OVSDb is supported.

The OVSDb software is included in the jsdn package. For some Juniper Networks devices, the jsdn package is included in the Junos OS software (jinstall) package. If the jsdn package is not included in the jinstall package for a particular device, a separate jsdn package must be installed on the device in addition to the jinstall package.

For each device and Junos OS release, the table outlines whether or not the jsdn package is included in the jinstall package. If the jsdn package is not included, the table also includes the name of the separate jsdn package.



**NOTE:** The separate jsdn package release number must be the same as the jinstall release number running on the device.

**Table 208: OVSDb Support on Juniper Networks Devices**

| Juniper Networks Device          | Junos OS Release | jsdn Package Included in jinstall Package? | Separate jsdn Package Name             |
|----------------------------------|------------------|--------------------------------------------|----------------------------------------|
| EX9200 Line of Ethernet Switches | 14.2R1 and later | No                                         | jsdn-i386-release<br>jsdn-x86-release* |
| MX80 3D Universal Edge Routers   | 14.1R2 and later | No                                         | jsdn-powerpc-release                   |
| MX104 3D Universal Edge Routers  | 14.2R4 and later | No                                         | jsdn-powerpc-release                   |

Table 208: OVSDB Support on Juniper Networks Devices (*continued*)

| Juniper Networks Device                           | Junos OS Release      | jsdn Package Included in jinstall Package? | Separate jsdn Package Name |
|---------------------------------------------------|-----------------------|--------------------------------------------|----------------------------|
| MX240, MX480, and MX960 3D Universal Edge Routers | 14.1R2 and later      | No                                         | jsdn-i386-release          |
|                                                   |                       |                                            | jsdn-x86-release*          |
| MX2010 and MX2020 3D Universal Edge Routers       | 15.1R2 and later      | No                                         | jsdn-i386-release          |
|                                                   |                       |                                            | jsdn-x86-release           |
| QFX10002 Switches                                 | 15.1X53-D10           | No                                         | jsdn-i386-release          |
|                                                   | 15.1X53-D20 and later | Yes                                        | —                          |
| QFX10008 Switches                                 | 15.1X53-D30 and later | Yes                                        | —                          |

\*This jsdn package is introduced in Junos OS Release 15.1R2.

#### Related Documentation

- *Installing OVSDB on Juniper Networks Devices*

## OVSDB Schema for Physical Devices

An Open vSwitch Database (OVSDB) server runs on a Juniper Networks device that supports the OVSDB management protocol. When this device is connected to one or more SDN controllers, the connections provide a means through which the Juniper Networks device and the SDN controllers can communicate.

Juniper Networks devices that support OVSDB and SDN controllers exchange control and statistical data. This data is stored in the OVSDB database schema defined for physical devices. The schema resides in the OVSDB server. The schema includes several tables. Juniper Networks devices and SDN controllers, both of which have OVSDB clients, can add rows to the tables as well as monitor the tables for the addition, deletion, and modification of rows.

For example, the OVSDB client on a Juniper Networks device and an SDN controller can collect MAC routes learned by entities in the physical or virtualized networks, respectively, and publish the routes to the appropriate table in the schema. By using the MAC routes and other information provided in the table, Juniper Networks devices in the physical network and entities in the virtualized network can determine where to forward virtual machine (VM) traffic.

Some of the OVSDB table names include the words *local* or *remote*—for example, the *unicast MACs local table* and the *unicast MACs remote table*. Information in *local* tables is learned by a Juniper Networks device that functions as a hardware virtual tunnel endpoint (VTEP), whereas information in *remote* tables is learned by other software or hardware VTEPs.

[Table 209](#) describes the tables in the schema, the physical or virtual entity that is the source of the data provided in the table, and the command that you can enter in the CLI of the Juniper Networks device to get similar information.

**Table 209: OVSDb Schema Tables**

| Table Name                       | Description                                                                                                                                                                                                                                                                                                          | Source of Information                                                                             | Command                                              |
|----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------|------------------------------------------------------|
| Global table                     | Includes the top-level configuration for the Juniper Networks device.                                                                                                                                                                                                                                                | Juniper Networks device                                                                           | —                                                    |
| Manager table                    | Includes information about each SDN controller that is connected to the Juniper Networks device.                                                                                                                                                                                                                     | Juniper Networks device                                                                           | <a href="#">show ovssdb controller</a>               |
| Physical switch table            | Includes information about a Juniper Networks device that functions as a hardware VTEP. This table includes information only for the device on which the table resides.                                                                                                                                              | Juniper Networks device                                                                           | —                                                    |
| Physical port table              | Includes information about OVSDb-managed interfaces.                                                                                                                                                                                                                                                                 | Juniper Networks device                                                                           | <a href="#">show ovssdb interface</a>                |
| Logical switch table             | Includes the following information: <ul style="list-style-type: none"> <li>Logical switches, which you configured in a VMware NSX environment, or virtual networks, which you configured in a Contrail environment.</li> <li>The equivalent VXLANs, which were configured on the Juniper Networks device.</li> </ul> | <ul style="list-style-type: none"> <li>SDN controller</li> <li>Juniper Networks device</li> </ul> | <a href="#">show ovssdb logical-switch</a>           |
| Logical binding statistics table | Includes statistics for OVSDb-managed interfaces.                                                                                                                                                                                                                                                                    | Juniper Networks device                                                                           | <a href="#">show ovssdb statistics interface</a>     |
| Physical locator table           | Includes information about Juniper Networks devices configured as hardware VTEPs, software VTEPs, and service nodes in an NSX environment.                                                                                                                                                                           | Juniper Networks device                                                                           | <a href="#">show ovssdb virtual-tunnel-end-point</a> |
| Physical locator set table       | Includes a list of software VTEPs, service nodes, or top-of-rack service nodes (TSNs) for a logical switch.                                                                                                                                                                                                          | Juniper Networks device                                                                           | —                                                    |



Table 209: OVSDb Schema Tables (*continued*)

| Table Name                  | Description                                                                                                                                                                                                                                                                                                                                                                        | Source of Information   | Command                         |
|-----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|---------------------------------|
| Unicast MACs remote table   | Reachability information, including unicast MAC addresses, for entities in the virtualized network.                                                                                                                                                                                                                                                                                | SDN controller          | <a href="#">show ovssdb mac</a> |
| Unicast MACs local table    | Reachability information, including unicast MAC addresses, for entities in the physical network.                                                                                                                                                                                                                                                                                   | Juniper Networks device | <a href="#">show ovssdb mac</a> |
| Multicast MACs remote table | Includes only one row. In this row, the MAC column includes the keyword <b>unknown dst</b> along with a list of software VTEPs, service nodes, or TSNs, which handle multicast traffic.                                                                                                                                                                                            | SDN controller          | <a href="#">show ovssdb mac</a> |
| Multicast MACs local table  | <p><b>NOTE:</b> Only QFX Series switches support this table.</p> <p>Includes one row for each logical switch. In this row, the MAC column includes the keyword <b>unknown dst</b> and a list of hardware VTEPs, which are identified by the IP address assigned to the hardware VTEP loopback interface (lo0). These hardware VTEPs can terminate or originate a VXLAN tunnel.</p> | Juniper Networks device | <a href="#">show ovssdb mac</a> |

**Related  
Documentation**

- [Understanding the OVSDb Protocol Running on Juniper Networks Devices on page 2445](#)
- [Understanding How to Set Up OVSDb Connections Between Juniper Networks Devices and SDN Controllers on page 2458](#)

## Understanding How Layer 2 BUM and Layer 3 Routed Multicast Traffic Are Handled with OVSDB

---

The Juniper Networks Junos OS implementation of the Open vSwitch Database (OVSDB) management protocol provides a means through which software-defined networking (SDN) controllers and Juniper Networks devices that support OVSDB can communicate.

This topic explains how a Juniper Networks device with Virtual Extensible LAN (VXLAN) and OVSDB management protocol capabilities handles the following types of traffic:

- (This scenario applies to all Juniper Networks devices that support VXLAN and OVSDB.) Layer 2 broadcast, unknown unicast, and multicast (BUM) traffic that originates in an OVSDB-managed VXLAN and is forwarded to interfaces within the same VXLAN.



**NOTE:** You must explicitly configure the replication of unknown unicast traffic in a Contrail environment.

- (This scenario applies only to MX Series routers and EX9200 switches that support VXLAN and OVSDB.) Layer 3 multicast traffic that is received by an integrated routing and bridging (IRB) in an OVSDB-managed VXLAN and is forwarded to interfaces in another OVSDB-managed VXLAN.

By default, Layer 2 BUM traffic that originates in an OVSDB-managed VXLAN is handled by one or more software virtual tunnel endpoints (VTEPs), service nodes, or top-of-rack service nodes (TSNs) in the same VXLAN. (This topic refers to the software VTEPs, service nodes, and TSNs collectively as *replicators*.) The table for remote multicast media access control (MAC) addresses in the OVSDB schema for physical devices contains only one entry that has the keyword **unknown-dst** as the MAC string and a list of replicators.

Given the previously described table entry, Layer 2 BUM traffic received on an interface in the OVSDB-managed VXLAN is forwarded to one of the replicators. The replicator to which a BUM packet is forwarded is determined by the Juniper Networks device on which the OVSDB-managed VXLAN is configured. On receiving the BUM packet, the entity replicates the packet and forwards the replicas to all interfaces within the VXLAN.

Instead of using replicators, you can optionally enable ingress node replication to handle Layer 2 BUM traffic on Juniper Networks devices that support OVSDB.



**NOTE:** For VXLAN-OVSDB, ingress node replication is supported on all Juniper Networks devices that support OVSDB except the QFX Series switches.

With ingress node replication enabled, on receiving a Layer 2 BUM packet on an interface in an OVSDB-managed VXLAN, the Juniper Networks device replicates the packet and then forwards the replicas to all software VTEPs included in the unicast MACs remote table in the OVSDB schema. The software VTEPs then forward the replicas to all virtual machines (VMs), except service VMs or nodes, on the same host.



.....

**NOTE:** When Juniper Networks devices replicate Layer 2 BUM packets to a large number of remote software VTEPs, the performance of the Juniper Networks devices can be impacted.

.....

On IRB interfaces that forward Layer 3 multicast traffic from one OVSDB-managed VXLAN to another, ingress node replication is automatically implemented. With ingress node replication, the Juniper Networks device replicates a Layer 3 multicast packet and then the IRB interface forwards the replicas to all hardware and software VTEPs, but not to service nodes, in the other OVSDB-managed VXLAN. For the routing of Layer 3 multicast traffic from one OVSDB-managed VXLAN to another, ingress node replication is the only option and does not need to be configured.



## PART 37

# Configuring OVSDB and VXLAN

- [Configuring OVSDB-Managed VXLANs with an SDN Controller on page 2455](#)
- [Configuring VXLANs Without an SDN Controller on page 2491](#)



## CHAPTER 89

# Configuring OVSDb-Managed VXLANs with an SDN Controller

- [OVSDb and VXLAN Configuration Workflows for VMware NSX Environment on page 2455](#)
- [Understanding How to Set Up OVSDb Connections on a Juniper Networks Device on page 2458](#)
- [Creating and Installing an SSL Key and Certificate on a Juniper Networks Device for Connection with SDN Controllers on page 2459](#)
- [Setting Up OVSDb on Juniper Networks Devices That Support the Dynamic Configuration of VXLANs on page 2461](#)
- [Understanding Dynamically Configured VXLANs in an OVSDb Environment on page 2462](#)
- [VMware NSX Configuration for Juniper Networks Devices Functioning as Virtual Tunnel Endpoints on page 2469](#)
- [Example: Setting Up a VXLAN Layer 2 Gateway and OVSDb Connections in a VMware NSX Environment \(Trunk Interfaces Supporting Untagged Packets\) on page 2472](#)
- [Example: Setting Up a VXLAN Layer 2 Gateway and OVSDb Connections in a VMware NSX Environment \(Trunk Interfaces Supporting Tagged Packets\) on page 2480](#)
- [Verifying That a Logical Switch and Corresponding Junos OS OVSDb-Managed VXLAN Are Working Properly on page 2488](#)

## OVSDb and VXLAN Configuration Workflows for VMware NSX Environment

---

The workflow that you use to configure Open vSwitch Database (OVSDb) and Virtual Extensible LAN (VXLAN) in a VMware NSX environment depends on the Juniper Networks device that you are configuring. This topic provides more information about the following workflows:

- [OVSDb and VXLAN Configuration Workflow for QFX Series Switches on page 2455](#)
- [OVSDb and VXLAN Configuration Workflow for MX Series Routers and EX9200 Switches on page 2457](#)

## OVSDb and VXLAN Configuration Workflow for QFX Series Switches

[Table 210](#) provides a high-level workflow of the tasks that you must perform to configure OVSDb and VXLAN on QFX Series switches. You must perform the tasks in [Table 210](#) for each Juniper Networks switch that you plan to deploy in an OVSDb environment. In

general, the successful completion of a task in this workflow depends on the successful completion of the previous task, so it is important to adhere to the task sequence provided in [Table 210](#).

**Table 210: OVSDB and VXLAN Configuration Workflow for QFX Series Switches**

| Sequence | Task                                                                                                                                                                                                                                                                                                                                                                  | For More Information                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1        | Create and install a Secure Sockets Layer (SSL) key and certificate.                                                                                                                                                                                                                                                                                                  | <a href="#">“Creating and Installing an SSL Key and Certificate on a Juniper Networks Device for Connection with SDN Controllers”</a> on page 2459.                                                                                                                                                                                                                                                                                                                                                          |
| 2        | Enter the <b>set switch-options ovssdb-managed</b> configuration mode command on the Juniper Networks switch.                                                                                                                                                                                                                                                         | —                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 3        | Explicitly configure a connection to at least one VMware NSX controller.                                                                                                                                                                                                                                                                                              | <a href="#">“Setting Up the OVSDB Protocol on Juniper Networks Devices that Support the Dynamic Configuration of VXLANs”</a> on page 2461.                                                                                                                                                                                                                                                                                                                                                                   |
| 4        | Specify that each physical interface associated with a VXLAN is to be managed by OVSDB.                                                                                                                                                                                                                                                                               | <a href="#">“Setting Up the OVSDB Protocol on Juniper Networks Devices that Support the Dynamic Configuration of VXLANs”</a> on page 2461.                                                                                                                                                                                                                                                                                                                                                                   |
| 5        | Configure a logical switch for each OVSDB-managed VXLAN that you plan to implement.                                                                                                                                                                                                                                                                                   | See the VMware documentation that accompanies NSX Manager or the NSX API.                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| 6        | <ul style="list-style-type: none"> <li>For each Juniper Network switch on which OVSDB-managed VXLANs and interfaces are configured, create a gateway.</li> <li>For each OVSDB-managed interface that you configure, create a gateway service.</li> <li>For each logical interface that you plan to implement for a VXLAN, configure a logical switch port.</li> </ul> | <p>For general information about configuring gateways, gateway services, and logical switch ports, see the VMware documentation that accompanies NSX Manager or the NSX API.</p> <p>For key NSX Manager configuration details that help you configure gateways, gateway services, and logical switch ports so they function properly with their physical counterparts, see <a href="#">“VMware NSX Configuration for Juniper Networks Devices Functioning as Virtual Tunnel Endpoints”</a> on page 2469.</p> |
| 7        | <p>Configure the loopback interface (lo0) on the Juniper Networks switch for VXLAN by entering the following configuration mode commands:</p> <ul style="list-style-type: none"> <li><b>set interfaces lo0 unit 0 family inet address <i>ip-address</i> primary</b></li> <li><b>set switch-options vtep-source-Interface lo0.0</b></li> </ul>                         | —                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

After you successfully complete task 6 in [Table 210](#), the Juniper Networks switch dynamically creates a VXLAN for each logical switch that you configured in task 5. The Juniper Networks switch also dynamically creates and associates interfaces with each VXLAN. The dynamically created interface configuration is based on the gateway service and logical switch ports that you configured in task 6. For more information, see [“Understanding Dynamically Configured VXLANs in an OVSDB Environment”](#) on page 2462.



For OVSDb-VXLAN scenarios in which Juniper Networks switches are commonly deployed, see the following topics:

- [Example: Setting Up a VXLAN Layer 2 Gateway and OVSDb Connections in a VMware NSX Environment \(Trunk Interfaces Supporting Untagged Packets\)](#) on page 2472
- [Example: Setting Up a VXLAN Layer 2 Gateway and OVSDb Connections Between Virtual and Physical Entities in a Data Center \(Using Trunk Interfaces\)](#) on page 2480

## OVSDb and VXLAN Configuration Workflow for MX Series Routers and EX9200 Switches

[Table 211](#) provides a high-level workflow of the tasks that you must perform to configure OVSDb and VXLAN on MX Series routers and EX9200 switches. You must perform the tasks in [Table 211](#) for each Juniper Networks device that you plan to deploy in an OVSDb environment. In general, the successful completion of a task in this workflow depends on the successful completion of the previous task, so it is important to adhere to the task sequence provided in [Table 211](#).

**Table 211: OVSDb and VXLAN Configuration Workflow for MX Series Routers and EX9200 Switches**

| Sequence | Task                                                                                                                                                                                                                                                                                                                       | For More Information                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1        | Install the jsdn software package.                                                                                                                                                                                                                                                                                         | <i>Installing OVSDb on Juniper Networks Devices.</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 2        | Create and install an SSL key and certificate.                                                                                                                                                                                                                                                                             | <a href="#">“Creating and Installing an SSL Key and Certificate on a Juniper Networks Device for Connection with SDN Controllers”</a> on page 2459.                                                                                                                                                                                                                                                                                                                                                                |
| 3        | Explicitly configure a connection to at least one NSX controller.                                                                                                                                                                                                                                                          | <i>Setting Up the OVSDb Protocol on Juniper Networks Devices that Support Manual Configuration of VXLANs.</i>                                                                                                                                                                                                                                                                                                                                                                                                      |
| 4        | Specify that each physical interface associated with a VXLAN is to be managed by OVSDb.                                                                                                                                                                                                                                    | <i>Setting Up the OVSDb Protocol on Juniper Networks Devices that Support Manual Configuration of VXLANs.</i>                                                                                                                                                                                                                                                                                                                                                                                                      |
| 5        | Configure a logical switch for each OVSDb-managed VXLAN that you plan to implement.                                                                                                                                                                                                                                        | See the VMware documentation that accompanies NSX Manager or the NSX API.                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 6        | Configure OVSDb-managed VXLANs.                                                                                                                                                                                                                                                                                            | <i>Configuring OVSDb-Managed VXLANs.</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| 7        | <p>For each Juniper Network device on which OVSDb-managed VXLANs and interfaces will be configured, create a gateway.</p> <p>For each OVSDb-managed interface that you configure, create a gateway service.</p> <p>For each logical interface that you plan to implement for a VXLAN, configure a logical switch port.</p> | <p>For general information about configuring gateways, gateway services, and logical switch ports, see the VMware documentation that accompanies NSX Manager or the NSX API.</p> <p>For key NSX Manager configuration details that help you configure gateways, gateway services, and logical switch ports, so that they function properly with their physical counterparts, see <a href="#">“VMware NSX Configuration for Juniper Networks Devices Functioning as Virtual Tunnel Endpoints”</a> on page 2469.</p> |

**Table 211: OVSDB and VXLAN Configuration Workflow for MX Series Routers and EX9200 Switches (continued)**

| Sequence | Task                                                                                                                                                                                                                                                                                                                                              | For More Information |
|----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------|
| 8        | <p>Configure the loopback interface (lo0) on the Juniper Networks device for VXLAN by entering the following configuration mode commands:</p> <ul style="list-style-type: none"> <li>• <b>set interfaces lo0 unit 0 family inet address <i>ip-address</i> primary</b></li> <li>• <b>set switch-options vtep-source-Interface lo0.0</b></li> </ul> | –                    |

For OVSDB-VXLAN scenarios in which these Juniper Networks devices are commonly deployed, see the following topics:

- *Example: Setting Up Inter-VXLAN Unicast Routing and OVSDB Connections in a Data Center*
- *Example: Setting Up Inter-VXLAN Unicast and Multicast Routing and OVSDB Connections in a Data Center*
- *Example: Configuring VXLAN to VPLS Stitching with OVSDB*

**Related Documentation** • [Verifying That a Logical Switch and Corresponding Junos OS OVSDB-Managed VXLAN Are Working Properly on page 2488](#)

## Understanding How to Set Up OVSDB Connections on a Juniper Networks Device

The Juniper Networks Junos OS implementation of the Open vSwitch Database (OVSDB) management protocol provides a means through which Juniper Networks devices that support OVSDB can communicate with software-defined networking (SDN) controllers. A Juniper Networks device exchanges control and statistical data with each SDN controller to which it is connected.

You can connect a Juniper Networks device to more than one SDN controller for redundancy.

In a VMware NSX environment, one cluster of NSX controllers typically includes three or five controllers. To implement the OVSDB management protocol on a Juniper Networks device, you must explicitly configure a connection to one SDN controller, using the Junos OS CLI. If the SDN controller to which you explicitly configure a connection is in a cluster, the controller pushes information about other controllers in the same cluster to the device, and the device establishes connections with the other controllers. However, you can also explicitly configure connections with the other controllers in the cluster, using the Junos OS CLI.

To implement the OVSDB management protocol on a Juniper Networks device in a Contrail environment, you must configure a connection to a Contrail controller, using the Junos OS CLI.

Connections to all SDN controllers are made on the management interface of the Juniper Networks device. To set up a connection between a Juniper Networks device and an SDN controller, you need to configure the following parameters on the Juniper Networks device:

- IP address of the SDN controller.
- The protocol that secures the connection. Secure Sockets Layer (SSL) is the supported protocol.



**NOTE:** The SSL connection requires a private key and certificates, which must be stored in the `/var/db/certs` directory of the Juniper Networks device. See [“Creating and Installing an SSL Key and Certificate on a Juniper Networks Device for Connection with SDN Controllers”](#) on page 2459.

- Number of the port over which the connection is made. The port number of the default port is 6632.

Optionally, you can configure the following connection timers on the Juniper Networks device:

- Inactivity probe duration—The maximum amount of time, in milliseconds, that the connection can be inactive before an inactivity probe is sent. The default value is 0 milliseconds, which means that an inactivity probe is never sent.
- Maximum backoff duration—If an attempt to connect to an SDN controller fails, the maximum amount of time, in milliseconds, before the device can make the next attempt. The default value is 1000 milliseconds.

#### Related Documentation

- [Setting Up the OVSDb Protocol on Juniper Networks Devices that Support Manual Configuration of VXLANs](#)
- [Setting Up the OVSDb Protocol on Juniper Networks Devices that Support the Dynamic Configuration of VXLANs](#) on page 2461

## Creating and Installing an SSL Key and Certificate on a Juniper Networks Device for Connection with SDN Controllers

To secure a connection between a Juniper Networks device that supports the Open vSwitch Database (OVSDb) management protocol and one or more software-defined networking (SDN) controllers, the following Secure Sockets Layer (SSL) files must be present in the `/var/db/certs` directory on the device:

- `vtep-privkey.pem`
- `vtep-cert.pem`
- `ca-cert.pem`

You must create the `vtep-privkey.pem` and `vtep-cert.pem` files for the device and then install the two files in the `/var/db/certs` directory on the device.

Upon initial connection between a Juniper Networks device with OVSDb implemented and an SDN controller, the **ca-cert.pem** file is automatically generated and then installed in the **/var/db/certs** directory on the device.



**NOTE:** The situation at your particular site determines the possible methods that you can use to create the **vtep-privkey.pem** and **vtep-cert.pem** files and install them in the Juniper Networks device. Instead of providing procedures for all possible situations, this topic provides a procedure for one common scenario.

The procedure provided in this topic uses the OpenFlow public key infrastructure (PKI) management utility **ovs-pki** on a Linux computer to initialize a PKI and create the **vtep-privkey.pem** and **vtep-cert.pem** files. (If you have an existing PKI on your Linux computer, you can skip the step to initialize a new one.) By default, the utility initializes the PKI and places these files in the **/usr/local/share/openvswitch/pki** directory of the Linux computer.

To create and install an SSL key and certificate on a Juniper Networks device:

1. Initialize a PKI if one does not already exist on your Linux computer.

```
ovs-pki init
```

2. On the same Linux computer on which the PKI exists, create a new key and certificate for the Juniper Networks device.

```
ovs-pki req+sign vtep
```

3. Copy only the **vtep-privkey.pem** and **vtep-cert.pem** files from the Linux computer to the **/var/db/certs** directory on the Juniper Networks device.

**Related  
Documentation**

- [Understanding How to Set Up OVSDb Connections Between Juniper Networks Devices and SDN Controllers on page 2458](#)
- [OVSDb and VXLAN Configuration Workflows for VMware NSX Environment on page 2455](#)

## Setting Up OVSDB on Juniper Networks Devices That Support the Dynamic Configuration of VXLANs

To implement the Open vSwitch Database (OVSDB) management protocol on a Juniper Networks device, you must configure a connection between the Juniper Networks device and a software-defined networking (SDN) controller using the Junos OS CLI.

All SDN controller connections are made on the management interface of the Juniper Networks device. This connection is secured by using the Secure Sockets Layer (SSL) protocol. The default port number for the connection is 6632.

You must also specify that each physical interface that is connected to a physical server is managed by OVSDB. By performing this configuration, you essentially disable the Juniper Networks device from learning about other Juniper Networks devices that function as hardware virtual tunnel endpoints (VTEPs) and the MAC addresses learned by the hardware VTEPs. Instead, this configuration enables OVSDB to learn about these elements.

Before setting up OVSDB on a Juniper Networks device, you must do the following:

- Create an SSL private key and certificate, if they do not already exist, and install them in the `/var/db/certs` directory of the Juniper Networks device. See [“Creating and Installing an SSL Key and Certificate on a Juniper Networks Device for Connection with SDN Controllers”](#) on page 2459.

To set up OVSDB on a Juniper Networks device:

1. Specify the IP address of the SDN controller.
 

```
[edit protocols ovldb]
user@host# set controller ip-address
```
2. Specify SSL as the protocol that secures the connection between the Juniper Networks device and the SDN controller.
 

```
[edit protocols ovldb]
user@host# set controller ip-address protocol ssl
```
3. Set the number of the port over which the connection to the SDN controller is made.
 

```
[edit protocols ovldb]
user@host# set controller ip-address protocol ssl port number
```
4. (Optional) Specify (in milliseconds) how long the connection can be inactive before an inactivity probe is sent.
 

```
[edit protocols ovldb]
user@host# set controller ip-address inactivity-probe-duration milliseconds
```
5. (Optional) Specify (in milliseconds) how long the device must wait before it can try to connect to the SDN controller again if the previous attempt failed.
 

```
[edit protocols ovldb]
user@host# set controller ip-address maximum-backoff-duration milliseconds
```
6. (Optional) Repeat Steps 1 through 5 to configure a connection to an additional SDN controller in the NSX environment.
7. Specify that each physical interface that is connected to a physical server is managed by OVSDB.

```
[edit protocols ovssdb]
user@host# set interfaces interface-name
```

When specifying the *interface-name*, you do not need to include a logical unit number.

8. Complete the remaining configuration tasks, which are described in [“OVSSDB and VXLAN Configuration Workflows for VMware NSX Environment” on page 2455](#)).

## Understanding Dynamically Configured VXLANs in an OVSSDB Environment

---



**NOTE:** This topic applies only to QFX Series switches, which support the dynamic configuration of Open vSwitch Database (OVSSDB)-managed Virtual Extensible LANs (VXLANs). Although the configuration of OVSSDB-managed VXLANs is automated on these switches, there are tasks that you must perform before and after the dynamic configuration.

On all other Juniper Networks devices that support OVSSDB and VXLAN, you must manually configure OVSSDB-managed VXLANs using the Junos OS CLI. For more information about manually configuring OVSSDB-managed VXLANs, see [Configuring OVSSDB-Managed VXLANs](#).

The Juniper Networks Junos OS implementation of the OVSSDB management protocol provides a means through which Juniper Networks devices that support OVSSDB can communicate with software-defined networking (SDN) controllers. Support for OVSSDB enables the devices in a physical network to be integrated into a virtualized network.

In a Junos OS environment, the concept of an OVSSDB-managed Layer 2 broadcast domain in which data flows are limited to that domain is known as a *VXLAN*. The term used for the same concept in other OVSSDB environments depends on the environment:

- In an NSX environment, the same concept is known as a *logical switch*.
- In a Contrail environment, the same concept is known as a *virtual network*.

Understanding the terminology used in the different environments will help you to better understand the workflow associated with the dynamic configuration of OVSSDB-managed VXLANs, including tasks that you must perform before and after the dynamic configuration.

The following topics describe the dynamic configuration of OVSSDB-managed VXLANs:

- [Performing Tasks Before and After the Dynamic Configuration of OVSSDB-Managed VXLANs on page 2462](#)
- [What the Juniper Networks Switch Actually Creates Dynamically on page 2467](#)

### Performing Tasks Before and After the Dynamic Configuration of OVSSDB-Managed VXLANs

Although the configuration of OVSSDB-managed VXLANs is automated, there are some tasks that you must perform before and after the dynamic configuration. [Table 212](#) includes a sequentially ordered workflow of tasks and events for the dynamic configuration of OVSSDB-managed VXLANs in an NSX environment, while [Table 213](#) includes the

equivalent information for a Contrail environment. Your familiarity with these workflows will ensure that the dynamic configuration of OVSDB-managed VXLANs is properly implemented.

In [Table 212](#), the NSX controller and Juniper Networks switch handle the events described in workflow numbers 4, 6, and 7. You must perform the tasks described in workflow numbers 1, 2, 3, 5, and 8. If you perform a task in a different order than that outlined in [Table 212](#), the dynamic configuration might not work or the dynamically configured OVSDB-managed VXLAN might not become functional.

**Table 212: Workflow of Tasks and Events for the Dynamic Configuration of OVSDB-Managed VXLANs in an NSX Environment**

| Workflow Number | Task or Event                                                                                                                                       | How Task or Event Is Handled                                                                                                                                             | More Information About Task or Event                                                    |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------|
| 1               | Enable the Juniper Networks switch to dynamically configure an OVSDB-managed VXLAN.                                                                 | You must manually enable this capability by entering the <b>set switch-options ovldb-managed</b> configuration mode command on the switch.                               | —                                                                                       |
| 2               | On the Juniper Networks switch, configure each physical interface that is connected to a physical server so that the interface is managed by OVSDB. | For each physical interface, you must manually enter the <b>set protocols ovldb interfaces interface-name</b> configuration mode command.                                | When entering the interface name, you do not need to include a logical unit number.     |
| 3               | For each OVSDB-managed VXLAN that you want to implement, configure a logical switch.                                                                | You must manually configure the logical switch by using NSX Manager or the NSX API. See the documentation that accompanies NSX Manager or the NSX API.                   | A universally unique identifier (UUID) for the logical switch is dynamically generated. |
| 4               | Relevant information about the logical switch is pushed to the Juniper Networks switch.                                                             | The NSX controller pushes relevant information to the logical switch table in the OVSDB schema for physical devices. This schema resides in the Juniper Networks switch. | —                                                                                       |

**Table 212: Workflow of Tasks and Events for the Dynamic Configuration of OVSDB-Managed VXLANs in an NSX Environment (*continued*)**

| Workflow Number | Task or Event                                                                                                                                                                                                                                                                                                                                                                                             | How Task or Event Is Handled                                                                                                                                                                                                                                                              | More Information About Task or Event                                                                                                                                                                                                                           |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 5               | <p>Create the following entities:</p> <ul style="list-style-type: none"> <li>For each Juniper Networks switch that you deploy as a hardware VTEP, you create a gateway.</li> <li>For each OVSDB-managed interface that you configured in workflow number 2, you create a gateway service.</li> <li>For each interface that you plan to implement for a VXLAN, configure a logical switch port.</li> </ul> | You must manually configure these entities by using NSX Manager or the NSX API. See the documentation that accompanies NSX Manager or the NSX API. Also see <a href="#">“VMware NSX Configuration for Juniper Networks Devices Functioning as Virtual Tunnel Endpoints”</a> on page 2469. | –                                                                                                                                                                                                                                                              |
| 6               | Relevant information about the gateway service and logical switch port are pushed to the Juniper Networks switch.                                                                                                                                                                                                                                                                                         | The NSX controller pushes this information to the Juniper Networks switch.                                                                                                                                                                                                                | –                                                                                                                                                                                                                                                              |
| 7               | A corresponding VXLAN is dynamically created. Based on the gateway service and logical switch port configured in NSX Manager or the NSX API, one or more interfaces are also created and associated with the VXLAN.                                                                                                                                                                                       | The Juniper Networks switch dynamically creates the VXLAN and interface configuration.                                                                                                                                                                                                    | For the name of the VXLAN, the Juniper Networks switch uses the UUID of the logical switch.                                                                                                                                                                    |
| 8               | (Recommended) Verify that the logical switch, corresponding VXLAN, and associated interfaces are configured properly and are operational.                                                                                                                                                                                                                                                                 | You can enter the <b>show ovssdb logical-switch</b> operational mode command on the Juniper Networks switch. In the output, check the Flags field for the logical switches that you configured as described in workflow number 3 to ensure that it displays Created by both.              | If the output of the <b>show ovssdb logical-switch</b> operational mode command does not include the Created by both state, see <a href="#">“Troubleshooting a Nonoperational Logical Switch and Corresponding Junos OS OVSDB-Managed VXLAN”</a> on page 2507. |

In [Table 213](#), the Contrail controller and Juniper Networks switch handle the events described in workflow numbers 5, 8, and 9. You must perform all other tasks described in the table. If you perform a task in a different order than that outlined in [Table 213](#), the dynamic configuration might not work or the dynamically configured OVSDB-managed VXLAN might not become functional.





**NOTE:** Although you can perform the Contrail configurations outlined in [Table 213](#) in the Contrail Web user interface or in the Contrail REST API, [Table 213](#) only describes how to perform tasks in the Contrail Web user interface.

**Table 213: Workflow of Tasks and Events for the Dynamic Configuration of OVSDb-Managed VXLANs in a Contrail Environment**

| Workflow Number | Task or Event                                                                                                                                       | How Task or Event Is Handled                                                                                                                                                  | More Information About Task or Event                                                                             |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------|
| 1               | On the Juniper Networks switch, configure a unique hostname for the switch.                                                                         | You must manually enter the <b>set system host-name</b> <i>host-name</i> configuration mode command on the switch.                                                            | If implementing a virtual chassis, be aware that all members of the virtual chassis must have the same hostname. |
| 2               | Enable the Juniper Networks switch to dynamically configure an OVSDb-managed VXLAN.                                                                 | You must manually enable this capability by entering the <b>set switch-options</b> <i>ovsdb-managed</i> configuration mode command on the switch.                             | —                                                                                                                |
| 3               | On the Juniper Networks switch, configure each physical interface that is connected to a physical server so that the interface is managed by OVSDb. | For each physical interface, you must manually enter the <b>set protocols ovsdb interfaces</b> <i>interface-name</i> configuration mode command.                              | When entering the interface name, you do not need to include a logical unit number.                              |
| 4               | For each OVSDb-managed VXLAN that you want to implement, configure a virtual network in the Contrail Web user interface.                            | You must manually configure the virtual network by navigating to Configure > Networking > Networks.<br><br>See <a href="#">Creating a Virtual Network</a> .                   | See <i>Contrail Configuration for Juniper Networks Devices That Function as Hardware VTEPs</i> .                 |
| 5               | Relevant information about the virtual network is pushed to the Juniper Networks switch.                                                            | The Contrail controller pushes relevant information to the logical switch table in the OVSDb schema for physical devices. This schema resides in the Juniper Networks switch. | —                                                                                                                |

Table 213: Workflow of Tasks and Events for the Dynamic Configuration of OVSDB-Managed VXLANs in a Contrail Environment (*continued*)

| Workflow Number | Task or Event                                                                                                                                                                                      | How Task or Event Is Handled                                                                                                                                                                                                                                                                                                       | More Information About Task or Event                                                                                                                                                                                                                           |
|-----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 6               | For each interface that you plan to implement for a VXLAN, configure a logical interface.                                                                                                          | <p>In the Contrail Web user interface, you must manually configure the logical interface by navigating to Configure &gt; Physical Devices &gt; Interfaces.</p> <p>For information about configuring a logical interface, see <a href="#">Using TOR Switches and OVSDB to Extend the Contrail Cluster to Other Instances</a>.</p>   | See <i>Contrail Configuration for Juniper Networks Devices That Function as Hardware VTEPs</i> .                                                                                                                                                               |
| 7               | For each Juniper Networks switch that you deploy as a hardware VTEP, you create a physical router.                                                                                                 | <p>In the Contrail Web user interface, you must manually configure the physical router by navigating to Configure &gt; Physical Devices &gt; Physical Routers.</p> <p>For information about configuring a physical router, see <a href="#">Using TOR Switches and OVSDB to Extend the Contrail Cluster to Other Instances</a>.</p> | See <i>Contrail Configuration for Juniper Networks Devices That Function as Hardware VTEPs</i> .                                                                                                                                                               |
| 8               | Relevant information about the logical interfaces is pushed to the Juniper Networks switch.                                                                                                        | The Contrail controller pushes this information to the Juniper Networks switch.                                                                                                                                                                                                                                                    | —                                                                                                                                                                                                                                                              |
| 9               | A corresponding VXLAN is dynamically created. Based on the logical interface configured in the Contrail Web user interface, one or more interfaces are also created and associated with the VXLAN. | The Juniper Networks switch dynamically creates the VXLAN and interface configurations.                                                                                                                                                                                                                                            | For the name of the VXLAN, the Juniper Networks switch uses the prefix “Contrail-” and the UUID of the virtual network.                                                                                                                                        |
| 10              | (Recommended) Verify that the virtual network, corresponding VXLAN, and interfaces are configured properly and are operational.                                                                    | You can enter the <b>show ovssdb logical-switch</b> operational mode command on the Juniper Networks switch. In the output, check the Flags field for the virtual network that you configured as described in workflow number 4 to ensure that it displays Created by both.                                                        | If the output of the <b>show ovssdb logical-switch</b> operational mode command does not include the Created by both state, see <a href="#">“Troubleshooting a Nonoperational Logical Switch and Corresponding Junos OS OVSDB-Managed VXLAN”</a> on page 2507. |

## What the Juniper Networks Switch Actually Creates Dynamically

When a Juniper Networks switch creates a VXLAN, it sets up a configuration similar to the following sample:

```
set vlans 28805c1d-0122-495d-85df-19abd647d772 vxlan vni 100
```

Note the following meanings for this sample configuration:

- The name of the VXLAN is 28805c1d-0122-495d-85df-19abd647d772. The UUID of the logical switch, which was configured in NSX Manager or in the NSX API, is 28805c1d-0122-495d-85df-19abd647d772. For a VXLAN created in a Contrail environment, the name would be preceded by “Contrail-”.
- For the virtual network identifier (VNI), the Juniper Networks switch uses either the VNI specified in the logical switch configuration (NSX) or the VXLAN identifier specified in the virtual network configuration (Contrail). In this example, VNI 100 is used. If the Juniper Networks switch detects that VNI 100 is a duplicate of a VNI from a VXLAN configured by manually using the **set vlans *vlan-name* vxlan vni (1–16777214)** command in the Junos OS CLI, the switch deletes the manually configured VXLAN. Or, if the Juniper Networks switch detects that VNI 100 is specified in the dynamically configured VXLAN, but for some reason, the VNI is no longer in the equivalent logical switch or virtual network configuration, the Juniper Networks switch deletes VNI 100 from the VXLAN.

If you need to modify or delete an OVSDB-managed VXLAN that was dynamically configured by the Juniper Networks switch, you must modify or delete either the corresponding logical switch configuration (NSX), or the corresponding virtual network configuration (Contrail). After you modify or delete the configuration, the SDN controller pushes the update to the Juniper Networks switch, and the switch modifies or deletes its configuration accordingly.

Depending on either the gateway service and logical switch ports configuration (NSX), or the logical interface configuration (Contrail), the Juniper Networks switch dynamically creates and associates one or more interfaces with the VXLAN. The configuration generated by the switch depends on whether an interface must support untagged or tagged packets. The following sections provide information about the configuration that the switch dynamically generates for each interface:

- [Dynamic Association of a Trunk Interface Supporting Untagged Packets to a Dynamically Created VXLAN on page 2467](#)
- [Dynamic Association of a Trunk Interface Supporting Tagged Packets to a Dynamically Created VXLAN on page 2468](#)

### Dynamic Association of a Trunk Interface Supporting Untagged Packets to a Dynamically Created VXLAN

To determine the type of interface to create and associate with an OVSDB-managed VXLAN, the Juniper Networks switch uses the VLAN ID that you specified when configuring either the logical switch port (NSX), or the logical interface (Contrail). If you specified **0** as the VLAN ID, the switch dynamically configures a trunk interface that can handle

untagged packets. (If you specified a valid VLAN ID other than 0, the switch creates a trunk interface that handles tagged packets.)

After the SDN controller pushes either the NSX or Contrail configurations to the Juniper Networks switch, the switch dynamically creates a configuration similar to the following:

```
set interfaces ge-1/0/0 flexible-vlan-tagging
set interfaces ge-1/0/0 native-vlan-id 4094
set interfaces ge-1/0/0 encapsulation extended-vlan-bridge
set interfaces ge-1/0/0 unit 0 vlan-id 4094
set vlans 28805c1d-0122-495d-85df-19abd647d772 interface ge-1/0/0.0
```

This sample configuration sets up physical interface ge-1/0/0 as a trunk interface. It also configures a native VLAN with an ID of 4094 and specifies that logical interface ge-1/0/0.0 is a member of the native VLAN. As a result, logical interface ge-1/0/0.0 handles incoming untagged packets.



**NOTE:** We reserve VLAN ID 4094 for native VLANs in an OVSDb environment. As a result, when you create either a logical switch port (NSX) or a logical interface (Contrail), if you specify VLAN ID 4094, the Juniper Networks switch does not dynamically configure a corresponding interface. Also, a system log error message is generated.

Instead of dynamically configuring physical interface ge-1/0/0 as an access interface, which typically handles untagged packets, the Juniper Networks switch configures it as a trunk interface. The intent of this configuration is to support the division of physical interface ge-1/0/0 into multiple logical interfaces, some of which are associated with VXLANs that handle untagged packets and some of which are associated with VXLANs that handle tagged packets.

The sample configuration also creates logical interface ge-1/0/0.0 and associates this interface with VXLAN 28805c1d-0122-495d-85df-19abd647d772.

### **Dynamic Association of a Trunk Interface Supporting Tagged Packets to a Dynamically Created VXLAN**

---

In a network that is divided into multiple VXLANs, each VXLAN has a VLAN ID associated with it. Packets associated with a particular VXLAN include the corresponding tag. In this situation, the interface that connects the Juniper Networks switch to a physical server in an OVSDb environment is a trunk interface that handles only tagged packets.

To determine the type of interface to create and associate with an OVSDb-managed VXLAN, the Juniper Networks switch uses the VLAN ID that you specified when configuring either the logical switch port (NSX), or the logical interface (Contrail). If you specified a valid VLAN ID other than 0 in either configuration, the switch creates a trunk interface that can handle tagged packets. (If you specified 0 as the VLAN ID, the switch creates a trunk interface that handles untagged packets.)

After the SDN controller pushes the NSX or Contrail configuration to the Juniper Networks switch, the switch dynamically creates a configuration similar to the following:

```
set interfaces ge-1/0/0 flexible-vlan-tagging
set interfaces ge-1/0/0 encapsulation extended-vlan-bridge
set interfaces ge-1/0/0 unit 10 vlan-id 10
set vlans 28805c1d-0122-495d-85df-19abd647d772 interfaces ge-1/0/0.10
```

The sample configuration sets up physical interface ge-1/0/0 as a trunk interface. It also configures a VLAN with an ID of 10 and specifies that interface ge-1/0/0.10 is a member of the VLAN. With the configuration of VLAN 10, logical interface ge-1/0/0.10 accepts incoming packets with a VLAN tag of 10 and adds a tag of 100 to each packet. Adding a tag of 100 identifies the packets as received by the VXLAN 28805c1d-0122-495d-85df-19abd647d772, which has a VNI of 100. This configuration also associates the trunk interface with VXLAN 28805c1d-0122-495d-85df-19abd647d772.

- Related Documentation**
- [Understanding the OVSDb Protocol Running on Juniper Networks Devices on page 2445](#)
  - [show ovssdb logical-switch on page 2540](#)

---

## VMware NSX Configuration for Juniper Networks Devices Functioning as Virtual Tunnel Endpoints

---

When implementing the Open vSwitch Database (OVSDb) management protocol and Virtual Extensible LANs (VXLANs) on a Juniper Networks device, you must perform the following tasks in VMware NSX Manager or in the NSX API:

- For each Juniper Networks device on which OVSDb-managed VXLANs and physical interfaces are configured, you must create an NSX-equivalent entity, which is known as a *gateway*.
- For each OVSDb-managed physical interface that you configure on a Juniper Networks device, you must configure a gateway service—for example, a VTEP Layer 2 gateway service.
- For each logical interface that you want to implement for a VXLAN, you must configure a logical switch port.

The configurations described in this topic enable connectivity between physical servers in the physical network and virtual machines (VMs) in the virtual network.

This topic provides a high-level summary of the tasks that you must perform to create a gateway, gateway service, and logical switch ports. Although you can create these virtual entities either in NSX Manager or in the NSX API, this topic only describes how to perform the tasks in NSX Manager. Also, this topic does not include a complete procedure for each task. Rather, it includes key NSX Manager configuration details for ensuring the correct configuration of the virtual entities so that they function properly with the physical entities.

For complete information about performing the tasks described in this topic, see the documentation that accompanies NSX Manager.

This topic describes the following tasks:

- [Creating a Gateway on page 2470](#)
- [Creating a Gateway Service on page 2470](#)
- [Creating a Logical Switch Port on page 2471](#)

## Creating a Gateway

In NSX Manager, you must create a gateway for each Juniper Networks device on which OVSDB-managed VXLANs and physical interfaces are configured. [Table 214](#) provides a summary of key configuration fields in NSX Manager and how to configure them when creating a gateway.

**Table 214: Key Configurations to Create a Gateway in NSX Manager**

| NSX Manager Configuration Page or Dialog Box | NSX Manager Configuration Field | How to Configure                                                                       |
|----------------------------------------------|---------------------------------|----------------------------------------------------------------------------------------|
| Type                                         | Transport Node Type             | Select <b>Gateway</b> .                                                                |
| Properties                                   | VTEP Enabled                    | Select <b>VTEP Enabled</b> .                                                           |
| Credential                                   | Type                            | Select <b>Management Address</b> .                                                     |
| Credential                                   | Management Address              | Specify the management IP address of the Juniper Networks device.                      |
| Connections/Create Transport Connector       | Transport Type                  | Select <b>VXLAN</b> .                                                                  |
| Connections/Create Transport Connector       | Transport Zone UUID             | Select the UUID of an existing transport zone, or create a new transport zone.         |
| Connections/Create Transport Connector       | IP Address                      | Specify the IP address of the loopback interface (lo0) of the Juniper Networks device. |

## Creating a Gateway Service

In NSX Manager, you must create a gateway service for each OVSDB-managed physical interface that you configure on a Juniper Networks device. Creating a gateway service essentially does the following for each OVSDB-managed physical interface:

- Specifies a gateway service—for example, a VTEP Layer 2 gateway service.
- Binds the interface to a gateway that you created in [“Creating a Gateway” on page 2470](#).

Before you start this task, you must complete the following configurations:

- A gateway for the Juniper Networks device on which the OVSDB-managed physical interfaces are configured. See [“Creating a Gateway” on page 2470](#).

- The OVSDB-managed physical interfaces on the Juniper Networks device. For information about configuring OVSDB-managed physical interfaces on Juniper Networks devices, see [“Setting Up the OVSDB Protocol on Juniper Networks Devices that Support the Dynamic Configuration of VXLANs” on page 2461](#) or [Setting Up the OVSDB Protocol on Juniper Networks Devices that Support Manual Configuration of VXLANs](#).

[Table 215](#) provides a summary of key configuration fields in NSX Manager and how to configure them when creating a gateway service.

**Table 215: Key Configurations to Create a Gateway Service in NSX Manager**

| NSX Manager Configuration Page or Dialog Box | NSX Manager Configuration Field | How to Configure                                                                      |
|----------------------------------------------|---------------------------------|---------------------------------------------------------------------------------------|
| Type                                         | Gateway Service Type            | Select <b>VTEP L2 Gateway Service</b> .                                               |
| Transport Nodes/Edit Gateway                 | Transport Node                  | Select the gateway that you created for the Juniper Networks device.                  |
| Transport Nodes/Edit Gateway                 | Port ID                         | Select an OVSDB-managed physical interface configured on the Juniper Networks device. |

## Creating a Logical Switch Port

In NSX Manager, you must create a logical switch port for each logical interface that you plan to implement for a VXLAN. Creating the logical switch port essentially does the following for each logical interface:

- Binds the logical switch port to a logical switch that you created in NSX Manager or in the NSX API.
- Binds the logical interface to a gateway service that you configured in [“Creating a Gateway Service” on page 2470](#).

Before you start this task, you must complete the following configurations:

- A logical switch with which this logical port is associated. For information about configuring a logical switch, see the VMware documentation that accompanies NSX Manager or the NSX API.
- A gateway service that specifies the OVSDB-managed physical interface with which the logical interface is associated. See [“Creating a Gateway Service” on page 2470](#).

[Table 216](#) provides a summary of key configuration fields in NSX Manager and how to configure them when creating a logical switch port.

**Table 216: Key Configurations to Create a Logical Switch Port in NSX Manager**

| NSX Manager Configuration Page or Dialog Box | NSX Manager Configuration Field | How to Configure                     |
|----------------------------------------------|---------------------------------|--------------------------------------|
| Logical Switch                               | Logical Switch UUID             | Select the UUID of a logical switch. |

Table 216: Key Configurations to Create a Logical Switch Port in NSX Manager (*continued*)

| NSX Manager Configuration Page or Dialog Box | NSX Manager Configuration Field | How to Configure                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------------------------|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Attachment                                   | Attachment Type                 | Select <b>VTEP L2 Gateway</b> .                                                                                                                                                                                                                                                                                                                                                          |
| Attachment                                   | VTEP L2 Gateway Service UUID    | Select the UUID of a gateway service.                                                                                                                                                                                                                                                                                                                                                    |
| Attachment                                   | VLAN                            | <p>Select <b>0</b> to specify that the port handles untagged packets.</p> <p>Select <b>1</b> through <b>4000</b> to specify that the port handles tagged packets.</p> <p><b>NOTE:</b> VLAN ID 4094 is reserved for a native VLAN in an OVSDb environment. Specifying this VLAN ID results in an error message. Do not specify this VLAN ID or any VLAN ID not in the accepted range.</p> |

**Related Documentation** • [OVSDb and VXLAN Configuration Workflows for VMware NSX Environment on page 2455](#)

### Example: Setting Up a VXLAN Layer 2 Gateway and OVSDb Connections in a VMware NSX Environment (Trunk Interfaces Supporting Untagged Packets)

In a physical network, a Juniper Networks device that supports Virtual Extensible LAN (VXLAN) can function as a hardware virtual tunnel endpoint (VTEP). In this role, the Juniper Networks device encapsulates Layer 2 Ethernet frames received from software applications that run directly on a physical server in VXLAN packets. The VXLAN packets are tunneled over a Layer 3 transport network. Upon receipt of the VXLAN packets, software VTEPs in the virtual network de-encapsulate the packets and forward the packets to virtual machines (VMs).

In this VXLAN environment, you can also include VMware NSX controllers and implement the Open vSwitch Database (OVSDb) management protocol on the Juniper Networks device that functions as a hardware VTEP. The Junos OS implementation of OVSDb provides a means through which VMware NSX controllers and Juniper Networks devices can exchange MAC addresses of entities in the physical and virtual networks. This exchange of MAC addresses enables the Juniper Networks device that functions as a hardware VTEP to forward traffic to software VTEPs in the virtual network and software VTEPs in the virtual network to forward traffic to the Juniper Networks device in the physical network.

This example explains how to configure a QFX Series switch as a hardware VTEP, which serves as a Layer 2 gateway, and set up this device with an OVSDb connection to an NSX controller.

In this example, only one VXLAN is deployed. Given this scenario, the packets exchanged between an application running on a physical server and a VM in the VXLAN are untagged. As a result, the QFX Series switch automatically configures a logical trunk interface for



the connection between the physical server and the switch, as well as a native VLAN. The native VLAN enables the trunk interface to handle the untagged packets.

- [Requirements on page 2473](#)
- [Overview and Topology on page 2473](#)
- [Non-OVSDB and Non-VXLAN Configuration on page 2476](#)
- [OVSDB and VXLAN Configuration on page 2477](#)
- [Verification on page 2478](#)

## Requirements

This example includes the following hardware and software components:

- A physical server on which software applications directly run.
- A QFX10002 switch running Junos OS software 15.1X53-D30 or later.
- On the QFX Series switch, physical interface ge-1/0/0 provides a connection to physical server 1.
- A cluster of five NSX controllers. (In this example, you explicitly configure a connection with one NSX controller.)
- NSX Manager.
- A service node that handles the replication and forwarding of Layer 2 broadcast, unknown unicast, and multicast (BUM) traffic within the VXLAN used in this example.
- A host that includes VMs managed by a hypervisor, which includes a software VTEP.

Before you begin:

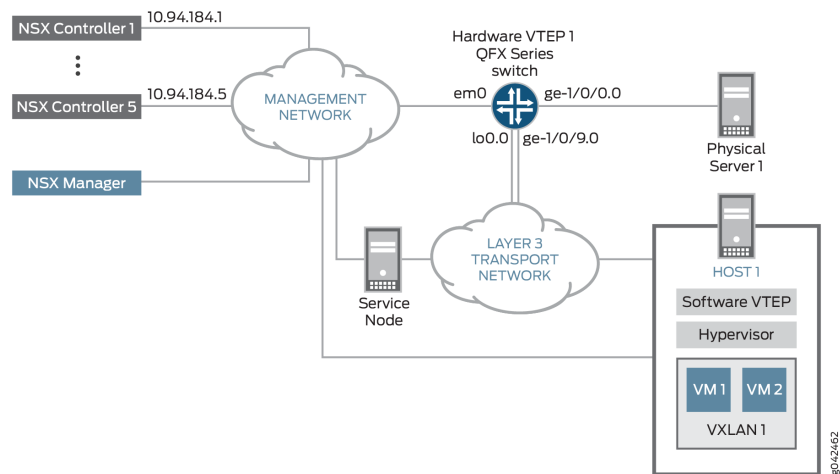
- Create an SSL private key and certificate, and install them in the `/var/db/certs` directory of the QFX Series switch. See [“Creating and Installing an SSL Key and Certificate on a Juniper Networks Device for Connection with SDN Controllers” on page 2459](#).
- Using NSX Manager, specify the IP address of the service node.

For information about using NSX Manager, see the documentation that accompanies these VMware products.

## Overview and Topology

[Figure 33](#) shows a topology in which a software application running directly on physical server 1 in the physical network needs to communicate with virtual machine VM1 in VXLAN 1 and vice versa.

Figure 33: VXLAN-OVSDB Layer 2 Gateway Topology



To establish communication between the software application on physical server 1 and VM 1 in VXLAN 1, a connection with an NSX controller is explicitly configured on the management interface of the QFX Series switch by using the Junos OS CLI.

Also, some entities in the VXLAN-OVSDB topology must be configured in both NSX Manager and on the QFX Series switch. [Table 217](#) provides a summary of the entities that must be configured and where they must be configured.

Table 217: NSX Manager and Junos OS Entities That Must Be Configured

| Entities                                                                      | What Must Be Configured in NSX Manager                                                                                                                                  | What Must Be Configured on a QFX Series Switch                                                                                             |
|-------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------|
| VXLAN 1                                                                       | Logical switch for VXLAN 1                                                                                                                                              | VXLAN 1<br><br><b>NOTE:</b> The QFX Series switch automatically configures this VXLAN.                                                     |
| Physical interface (ge-1/0/0) between physical server 1 and QFX Series switch | A gateway service. For gateway service type, select VTEP L2 Gateway service.                                                                                            | OVSDB management. Specify that interface ge-1/0/0 is managed by OVSDB.                                                                     |
| One logical interface (ge-1/0/0.0) associated with VXLAN 1                    | One logical switch port for VXLAN 1. For this port, specify VLAN number 0.<br><br><b>NOTE:</b> A VLAN number of 0 indicates that the port must handle untagged packets. | One logical interface (ge-1/0/0.0) for VXLAN 1.<br><br><b>NOTE:</b> The QFX Series switch automatically configures this logical interface. |
| QFX Series switch (hardware VTEP 1)                                           | Gateway                                                                                                                                                                 | —                                                                                                                                          |

In NSX Manager, a logical switch for VXLAN 1 is configured. In this configuration, a VXLAN network identifier (VNI) of 100 is specified. Also, the universally unique identifier (UUID) that NSX Manager assigns to the logical switch is 28805c1d-0122-495d-85df-19abd647d772. Based on this configuration, the QFX Series

switch automatically creates the following configuration for a Junos OS-equivalent VXLAN:

```
set vlans 28805c1d-0122-495d-85df-19abd647d772 vxlan vni 100
```

Based on the gateway service and logical switch port configuration (VLAN number 0) in NSX Manager, the QFX Series switch automatically creates the following configuration for a Junos OS-equivalent interface:

```
set interfaces ge-1/0/0 flexible-vlan-tagging
set interfaces ge-1/0/0 native-vlan-id 4094
set interfaces ge-1/0/0 encapsulation extended-vlan-bridge
set interfaces ge-1/0/0 unit 0 vlan-id 4094
set vlans 28805c1d-0122-495d-85df-19abd647d772 interface ge-1/0/0.0
```

This configuration sets physical interface ge-1/0/0 as a trunk interface. It also configures a native VLAN with an ID of 4094. The configuration creates logical interface ge-1/0/0.0 and specifies that it is a member of the native VLAN. As a result, logical interface ge-1/0/0.0 handles incoming untagged packets.

The configuration also associates logical interface ge-1/0/0.0 with VXLAN 28805c1d-0122-495d-85df-19abd647d772.

[Table 218](#) provides a summary of the VXLAN-OVSDB topology components that are configured on the QFX Series switch and the configuration settings for each component.

**Table 218: Components of the Topology for Setting Up a VXLAN Layer 2 Gateway and OVSDB Connections**

| Component                        | Setting                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| NSX controller                   | IP address: 10.94.184.1                                                                                                                                                                                                                                                                                                                                                                                                            |
| OVSDB-managed physical interface | Interface name: ge-1/0/0<br>Native VLAN ID: 4094                                                                                                                                                                                                                                                                                                                                                                                   |
| Logical interface                | <p><b>NOTE:</b> The QFX Series switch automatically creates this logical interface configuration, which is based on the gateway service configuration and logical switch port configuration in NSX Manager. Therefore, no manual configuration is required.</p> <p>Interface name: ge-1/0/0.0</p> <p>Interface type: trunk</p> <p>Member of native VLAN 4094</p> <p>Associated with VXLAN 28805c1d-0122-495d-85df-19abd647d772</p> |

**Table 218: Components of the Topology for Setting Up a VXLAN Layer 2 Gateway and OVSDB Connections** (*continued*)

| Component                                                                     | Setting                                                                                                                                                                                                                                                                                                  |
|-------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| OVSDB-managed VXLAN                                                           | <p><b>NOTE:</b> The QFX Series switch automatically creates this VXLAN configuration, which is based on the logical switch configuration in NSX Manager. Therefore, no manual configuration is required.</p> <p>For VXLAN 1:</p> <p>VXLAN name: 28805c1d-0122-495d-85df-19abd647d772</p> <p>VNI: 100</p> |
| OVSDB tracing operations                                                      | <p>Filename: /var/log/ovsdb</p> <p>File size: 10 MB</p> <p>Flag: All</p>                                                                                                                                                                                                                                 |
| Hardware VTEP source identifier                                               | <p>Source interface: loopback (lo0.0)</p> <p>Source IP address: 10.17.17.17/32</p>                                                                                                                                                                                                                       |
| Handling of Layer 2 BUM traffic in VXLAN 28805c1d-0122-495d-85df-19abd647d772 | <p>Service node</p> <p><b>NOTE:</b> By default, one or more service nodes handle Layer 2 BUM traffic within a VXLAN; therefore, no manual configuration is required.</p>                                                                                                                                 |

## Non-OVSDB and Non-VXLAN Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-1/0/9 unit 0 family inet address 10.40.40.1/24
set routing-options static route 10.19.19.19/32 next-hop 10.40.40.2
set routing-options router-id 10.17.17.17
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface ge-1/0/9.0
```

**Step-by-Step Procedure** To configure the Layer 3 network over which the packets exchanged between the physical server and VMs are tunneled:

1. Configure the Layer 3 interface.
 

```
[edit interfaces]
user@switch# set ge-1/0/9 unit 0 family inet address 10.40.40.1/24
```
2. Set the routing options.
 

```
[edit routing-options]
user@switch# set static route 10.19.19.19/32 next-hop 10.40.40.2
user@switch# set router-id 10.17.17.17
```

3. Configure the routing protocol.

```
[edit protocols]
user@switch# set ospf area 0.0.0.0 interface lo0.0
user@switch# set ospf area 0.0.0.0 interface ge-1/0/9.0
```

## OVSDb and VXLAN Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set switch-options ovbdb-managed
set protocols ovbdb controller 10.94.184.1
set protocols ovbdb interfaces ge-1/0/0
set protocols ovbdb traceoptions file ovbdb
set protocols ovbdb traceoptions file size 10m
set protocols ovbdb traceoptions flag all
set interfaces lo0 unit 0 family inet address 10.17.17.17/32 primary
set interfaces lo0 unit 0 family inet address 10.17.17.17/32 preferred
set switch-options vtep-source-interface lo0.0
```

**Step-by-Step Procedure** To configure the QFX Series switch as a hardware VTEP with an OVSDb connection to an NSX controller:

1. Enable the QFX Series switch to automatically configure OVSDb-managed VXLANs and associated interfaces.

```
[edit switch-options]
user@switch# ovbdb-managed
```

2. Explicitly configure a connection with an NSX controller.

```
[edit protocols]
user@switch# set ovbdb controller 10.94.184.1
```

3. Specify that the interface between hardware VTEP 1 and physical server 1 is managed by OVSDb.

```
[edit protocols]
user@switch# set ovbdb interfaces ge-1/0/0
```

4. Set up OVSDb tracing operations.

```
[edit protocols]
user@switch# set ovbdb traceoptions file ovbdb
user@switch# set ovbdb traceoptions file size 10m
user@switch# set ovbdb traceoptions flag all
```

5. Specify an IP address for the loopback interface. This IP address serves as the source IP address in the outer header of any VXLAN-encapsulated packet.

```
[edit interfaces]
user@switch# set lo0 unit 0 family inet address 10.17.17.17/32 primary
user@switch# set lo0 unit 0 family inet address 10.17.17.17/32 preferred
```

6. Set the loopback interface as the interface that identifies hardware VTEP 1.

[edit switch-options]

user@switch# **set vtep-source-interface lo0.0**

7. In NSX Manager, configure a logical switch for VXLAN 1. See the VMware documentation that accompanies NSX Manager.
8. In NSX Manager, configure a gateway for the QFX Series switch, and configure a gateway service and logical switch port for the logical interface (ge-1/0/0.0). See [“VMware NSX Configuration for Juniper Networks Devices Functioning as Virtual Tunnel Endpoints”](#) on page 2469.

## Verification

Confirm that the configuration is working properly:

- [Verifying the Logical Switch Configuration on page 2478](#)
- [Verifying the MAC Address of VM 1 on page 2478](#)
- [Verifying the NSX Controller Connection on page 2479](#)
- [Verifying the OVSDb-Managed Interface on page 2479](#)

---

### Verifying the Logical Switch Configuration

**Purpose** Verify that the configuration of the logical switch with the UUID of 28805c1d-0122-495d-85df-19abd647d772 is present in the OVSDb schema for physical devices and that the Flags field of the **show ovssdb logical switch** output displays **Created by both**.

**Action** From operational mode, enter the **show ovssdb logical-switch** command.

```
user@switch> show ovssdb logical-switch
Logical switch information:
Logical Switch Name: 28805c1d-0122-495d-85df-19abd647d772
Flags: Created by both
VNI: 100
Num of Remote MAC: 1
Num of Local MAC: 0
```

**Meaning** The output verifies that the configuration for the logical switch is present. The **Created by both** state indicates that the logical switch was configured in NSX Manager, and that the QFX Series switch automatically created the corresponding VXLAN. In this state, the logical switch and the VXLAN are operational.

If the state of the logical switch is something other than **Created by both**, see [“Troubleshooting a Nonoperational Logical Switch and Corresponding Junos OS OVSDb-Managed VXLAN”](#) on page 2507.

---

### Verifying the MAC Address of VM 1

**Purpose** Verify that the MAC address of VM 1 is present in the OVSDb schema.

**Action** From operational mode, enter the **show ovssdb mac remote** command.

```
user@switch> show ovssdb mac remote
Logical Switch Name: 28805c1d-0122-495d-85df-19abd647d772
 Mac IP Encapsulation Vtep
 Address Address Address Address
a8:59:5e:f6:38:90 0.0.0.0 Vxlan over Ipv4 10.17.17.17
```

**Meaning** The output shows that the MAC address for VM 1 is present and is associated with the logical switch with the UUID of 28805c1d-0122-495d-85df-19abd647d772. Given that the MAC address is present, VM 1 is reachable through the QFX Series switch, which functions as a hardware VTEP.

### Verifying the NSX Controller Connection

**Purpose** Verify that the connection with the NSX controller is up.

**Action** From operational mode, enter the **show ovssdb controller** command to verify that the controller connection state is **up**.

```
user@switch> show ovssdb controller
VTEP controller information:
Controller IP address: 10.94.184.1
Controller protocol: ssl
Controller port: 6632
Controller connection: up
Controller seconds-since-connect: 542325
Controller seconds-since-disconnect: 542346
Controller connection status: active
```

**Meaning** The output shows that the connection state of the NSX controller is up, in addition to other information about the controller. The **up** state of the NSX controller indicates that OVSDb is enabled on the QFX Series switch.

### Verifying the OVSDb-Managed Interface

**Purpose** Verify that interface ge-1/0/0.0 is managed by OVSDb.

**Action** From operational mode, enter the **show ovssdb interface** command to verify that interface ge-1/0/0.0 is managed by OVSDb.

```
user@switch> show ovssdb interface
Interface VLAN ID Bridge-domain
ge-1/0/0 0 28805c1d-0122-495d-85df-19abd647d772
```

**Meaning** The output shows that interface ge-1/0/0 is managed by OVSDb. It also indicates that the interface is associated with VXLAN 28805c1d-0122-495d-85df-19abd647d772, which has a VLAN ID of 0.

## Example: Setting Up a VXLAN Layer 2 Gateway and OVSDB Connections in a VMware NSX Environment (Trunk Interfaces Supporting Tagged Packets)

---

In a physical network, a Juniper Networks device that supports Virtual Extensible LAN (VXLAN) can function as a hardware virtual tunnel endpoint (VTEP). In this role, the Juniper Networks device encapsulates Layer 2 Ethernet frames received from software applications that run directly on a physical server in VXLAN packets. The VXLAN packets are tunneled over a Layer 3 transport network. Upon receipt of the VXLAN packets, software VTEPs in the virtual network de-encapsulate the packets and forward the packets to virtual machines (VMs).

In this VXLAN environment, you can also include VMware NSX controllers and implement the Open vSwitch Database (OVSDB) management protocol on the Juniper Networks device that functions as a hardware VTEP. The Junos OS implementation of OVSDB provides a means through which VMware NSX controllers and Juniper Networks devices can exchange MAC addresses of entities in the physical and virtual networks. This exchange of MAC addresses enables the Juniper Networks device that functions as a hardware VTEP to forward traffic to software VTEPs in the virtual network and software VTEPs in the virtual network to forward traffic to the Juniper Networks device in the physical network.

This example explains how to configure a Juniper Networks device that supports VXLAN as a hardware VTEP. (The VTEP serves as a Layer 2 gateway.) This example also explains how to configure this device with an OVSDB connection to an NSX controller.

In this example, an application running directly on a physical server needs to communicate with a VM in a VXLAN, while another application on the physical server needs to communicate with VMs in another VXLAN. Therefore, the packets exchanged between the applications running on the physical server and the respective VMs with which they must communicate are tagged. As a result, a trunk interface is used for the connection between the physical server and the Juniper Networks device.

- [Requirements on page 2480](#)
- [Overview and Topology on page 2481](#)
- [Non-OVSDB and Non-VXLAN Configuration on page 2484](#)
- [OVSDB and VXLAN Configuration on page 2485](#)
- [Verification on page 2486](#)

### Requirements

This example includes the following hardware and software components:

- A physical server on which software applications directly run.
- A Juniper Networks switch that supports VXLAN and OVSDB. This switch can be a QFX10002 switch running Junos OS Release 15.1X53-D10 and later.
- On the Juniper Networks switch, physical interface ge-1/0/0 provides a connection to physical server 1.



- A cluster of five NSX controllers. (In this example, you explicitly configure a connection with one NSX controller.)
- NSX Manager.
- A service node that handles the replication and forwarding of Layer 2 broadcast, unknown unicast, and multicast (BUM) traffic within the VXLANs.
- Two hosts that include VMs. Each host is managed by a hypervisor, and each hypervisor includes a software VTEP.

Before you begin the configuration, you must perform the following tasks:

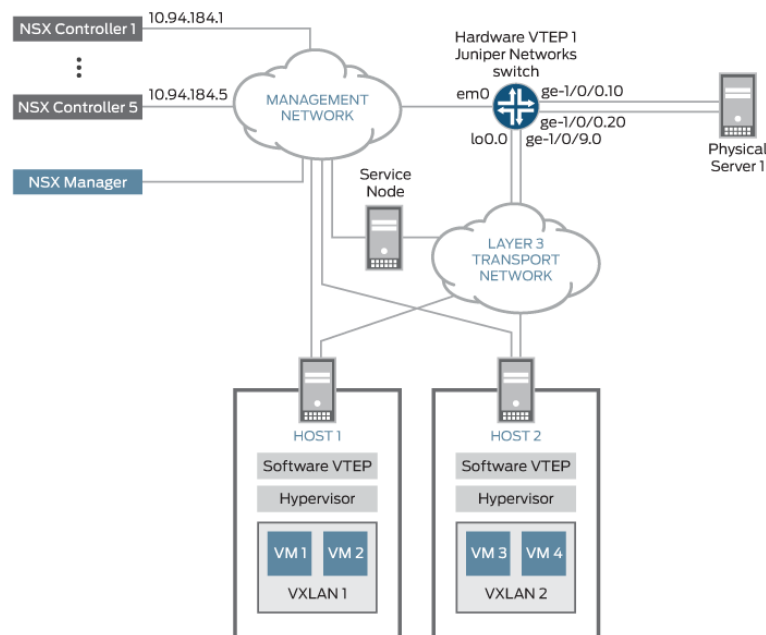
- Create an SSL private key and certificate, and install them in the `/var/db/certs` directory of the Juniper Networks switch. See [“Creating and Installing an SSL Key and Certificate on a Juniper Networks Device for Connection with SDN Controllers”](#) on page 2459.
- Using NSX Manager, specify the IP address of the service node.

For information about using NSX Manager, see the documentation that accompanies NSX Manager.

## Overview and Topology

Figure 33 shows a topology in which a software application running directly on physical server 1 in the physical network needs to communicate with virtual machine VM1 in VXLAN 1 and vice versa, and another software application on physical server 1 needs to communicate with virtual machines VM 3 and VM 4 in VXLAN 2 and vice versa.

Figure 34: VXLAN/OVSDB Layer 2 Gateway Topology



To establish communication between the software applications on physical server 1 and the VMs in VXLANs 1 and 2, some entities in the VXLAN-OVSDB topology must be configured in both NSX Manager and on the Juniper Networks switch. [Table 219](#) provides a summary of the entities that must be configured and where they must be configured.



**NOTE:** The term used for an entity configured in NSX Manager can differ from the term used for essentially the same entity configured on the Junos Network switch. To prevent confusion, [Table 219](#) shows the NSX Manager and Junos OS entities side-by-side.

**Table 219: NSX Manager and Junos OS Entities That Must Be Configured**

| Entities                                                                   | What Must Be Configured In NSX Manager                                                  | What Must Be Configured on Juniper Networks Switch                                        |
|----------------------------------------------------------------------------|-----------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------|
| VXLAN 1                                                                    | Logical switch for VXLAN 1                                                              | VXLAN 1                                                                                   |
| VXLAN 2                                                                    | Logical switch for VXLAN 2                                                              | VXLAN 2                                                                                   |
|                                                                            |                                                                                         | <b>NOTE:</b> The Juniper Networks switch dynamically configures these VXLANs.             |
| Interface (ge-1/0/0) between physical server 1 and Juniper Networks switch | A gateway service. For gateway service type, select VTEP L2 gateway service.            | OVSDb management. Specify that interface ge-1/0/0 is managed by OVSDb.                    |
| One logical interface associated with VXLAN 1                              | One logical switch port for VXLAN 1. For this port, specify VLAN number 10.             | One logical interface (ge-1/0/0.10) for VXLAN 1                                           |
| One logical interface associated with VXLAN 2                              | One logical switch port for VXLAN 2. For this port, specify VLAN number 20.             | One logical interface (ge-1/0/0.20) for VXLAN 1                                           |
|                                                                            | <b>NOTE:</b> A VLAN number from 1 through 4000 indicates that the port is a trunk port. | <b>NOTE:</b> The Juniper Networks switch dynamically configures these logical interfaces. |
| Juniper Networks switch (hardware VTEP 1)                                  | Gateway                                                                                 | —                                                                                         |

Based on the configuration of the entities in NSX Manager as described in [Table 219](#), the Juniper Networks switch dynamically creates VXLANs 1 and 2 and their associated logical interfaces. [Table 220](#) provides the relevant NSX Manager configuration and the resulting VXLANs and associated logical interfaces that the Juniper Networks switch dynamically configures.

**Table 220: NSX Manager Configurations and Dynamic Configurations by Juniper Networks Switch**

| NSX Manager Configuration: Logical Switch and Logical Switch Port                                                                                               | VXLANs and Associated Logical Interfaces Dynamically Configured By Juniper Networks Switch                                                                                                                                                                                                                                                                                                                         |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Logical switch configuration:<br><br>UUID:<br>28805c1d-0122-495d-85df-19abd647d772<br><br>VNI: 100<br><br>Logical switch port configuration:<br><br>VLAN ID: 10 | For VXLAN 1:<br><br><b>set vlans 28805c1d-0122-495d-85df-19abd647d772 vxlan vni 100</b><br><br>For associated logical interface ge-1/0/0.10:<br><br><b>set interfaces ge-1/0/0 flexible-vlan-tagging</b><br><b>set interfaces ge-1/0/0 encapsulation extended-vlan-bridge</b><br><b>set interfaces ge-1/0/0 unit 10 vlan-id 10</b><br><b>set vlans 28805c1d-0122-495d-85df-19abd647d772 interfaces ge-1/0/0.10</b> |
| Logical switch configuration:<br><br>UUID: 9acc24b3-7b0a-4c2e-b572-3370c3e1acff<br><br>VNI: 200<br><br>Logical switch port configuration:<br><br>VLAN ID: 20    | For VXLAN 2:<br><br><b>set vlans 9acc24b3-7b0a-4c2e-b572-3370c3e1acff vxlan vni 200</b><br><br>For associated logical interface ge-1/0/0.20:<br><br><b>set interfaces ge-1/0/0 flexible-vlan-tagging</b><br><b>set interfaces ge-1/0/0 encapsulation extended-vlan-bridge</b><br><b>set interfaces ge-1/0/0 unit 20 vlan-id 20</b><br><b>set vlans 9acc24b3-7b0a-4c2e-b572-3370c3e1acff interfaces ge-1/0/0.20</b> |

For VXLANs 1 and 2, the Juniper Networks switch uses the UUIDs and VNI values that were provided for the corresponding logical switches.

In the logical switch port configurations in NSX Manager, VLAN ID values 10 and 20 and logical switch mappings are specified. As a result, the Juniper Networks switch creates logical interfaces ge-1/0/0.10 and ge-1/0/0.20, respectively. Both of these logical interfaces function as trunk interfaces. The Juniper Networks switch also maps the logical interfaces ge-1/0/0.10 and ge-1/0/0.20 to their respective VXLANs.

Based on the configurations generated by the Juniper Networks switch, the interface ge-1/0/0.10 accepts packets with a VLAN tag of 10 from VXLAN 1, and interface ge-1/0/0.20 accepts packets with a VLAN tag of 20 from VXLAN 2. On receiving packets from VXLAN 1, a VLAN tag of 100 is added to the packets, and a VLAN tag of 200 is added to packets from VXLAN 2. These tags are added to the respective packet streams to map the VLAN ID in a particular VXLAN to the corresponding VNI.

[Table 218](#) provides a summary of the components that are configured on the Juniper Networks switch. Unless noted, all configurations are performed manually in the Junos OS CLI.

**Table 221: Components for Two VXLAN Topologies Configured on a Juniper Networks Switch that Functions as a Hardware VTEP**

| Components     | Settings                |
|----------------|-------------------------|
| NSX controller | IP address: 10.94.184.1 |

**Table 221: Components for Two VXLAN Topologies Configured on a Juniper Networks Switch that Functions as a Hardware VTEP (*continued*)**

| Components                                                                                                                              | Settings                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-----------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| OVSDB-managed interface                                                                                                                 | Interface name: ge-1/0/0                                                                                                                                                                                                                                                                                                                                                                                                           |
| VXLAN 1 and associated logical interface                                                                                                | <p><b>NOTE:</b> The Juniper Networks switch dynamically configures the VXLAN and associated logical interface, which are based on the logical switch and logical switch port configurations in NSX Manager. Therefore, no manual configuration is required.</p> <p>VXLAN name: 28805c1d-0122-495d-85df-19abd647d772</p> <p>VNI: 100</p> <p>Logical interface name: ge-1/0/0.10</p> <p>VLAN ID: 10</p> <p>Interface type: trunk</p> |
| VXLAN 2 and associated logical interface                                                                                                | <p><b>NOTE:</b> The Juniper Networks switch dynamically configures the VXLAN and associated interface, which are based on the logical switch and logical switch port configurations in NSX Manager. Therefore, no manual configuration is required.</p> <p>VXLAN name: VXLAN 9acc24b3-7b0a-4c2e-b572-3370c3e1acff</p> <p>VNI: 200</p> <p>Logical interface name: ge-1/0/0.20</p> <p>VLAN ID: 20</p> <p>Interface type: trunk</p>   |
| OVSDB tracing operations                                                                                                                | <p>Filename: /var/log/ovsdb</p> <p>File size: 10 MB</p> <p>Flag: All</p>                                                                                                                                                                                                                                                                                                                                                           |
| Hardware VTEP source identifier                                                                                                         | <p>Source interface: loopback (lo0.0)</p> <p>Source IP address: 10.17.17.17/32</p>                                                                                                                                                                                                                                                                                                                                                 |
| Handling of Layer 2 BUM traffic within VXLAN 28805c1d-0122-495d-85df-19abd647d772 and within VXLAN 9acc24b3-7b0a-4c2e-b572-3370c3e1acff | <p>Service node</p> <p><b>NOTE:</b> By default, one or more service nodes handle Layer 2 BUM traffic in a VXLAN; therefore, no configuration is required.</p>                                                                                                                                                                                                                                                                      |

## Non-OVSDB and Non-VXLAN Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```

set interfaces ge-1/0/9 unit 0 family inet address 10.40.40.1/24
set routing-options static route 10.19.19.19/32 next-hop 10.40.40.2
set routing-options router-id 10.17.17.17
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface ge-1/0/9.0

```

**Step-by-Step Procedure** To configure the Layer 3 network over which the packets exchanged between physical server 1 and VM1 are tunneled:

1. Configure the Layer 3 interface.

```

[edit interfaces]
user@switch# set ge-1/0/9 unit 0 family inet address 10.40.40.1/24

```

2. Set the routing options.

```

[edit routing-options]
user@switch# set static route 10.19.19.19/32 next-hop 10.40.40.2
user@switch# set router-id 10.17.17.17

```

3. Configure the routing protocol.

```

[edit protocols]
user@switch# set ospf area 0.0.0.0 interface lo0.0
user@switch# set ospf area 0.0.0.0 interface ge-1/0/9.0

```

## OVSDb and VXLAN Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```

set switch-options ovssdb-managed
set protocols ovssdb controller 10.94.184.1
set protocols ovssdb interfaces ge-1/0/0
set protocols ovssdb traceoptions file ovssdb
set protocols ovssdb traceoptions file size 10m
set protocols ovssdb traceoptions flag all
set interfaces lo0 unit 0 family inet address 10.17.17.17/32 primary
set interfaces lo0 unit 0 family inet address 10.17.17.17/32 preferred
set switch-options vtep-source-interface lo0.0

```

**Step-by-Step Procedure** To configure the Juniper Networks switch as hardware VTEP 1 and with an OVSDb connection to an NSX controller:

1. Enable the Juniper Networks switch to dynamically configure OVSDb-managed VXLANs and associated interfaces.

```

[edit switch-options]
user@switch# set ovssdb-managed

```

2. Explicitly configure a connection with an NSX controller.

```

[edit protocols]
user@switch# set ovssdb controller 10.94.184.1

```

3. Specify that interface ge-1/0/0 is managed by OVSDB.

```
[edit protocols]
user@switch# set ovsdb interfaces ge-1/0/0
```

4. Set up OVSDB tracing operations.

```
[edit protocols]
user@switch# set ovsdb traceoptions file ovsdb
user@switch# set ovsdb traceoptions file size 10m
user@switch# set ovsdb traceoptions flag all
```

5. Specify an IP address for the loopback interface. This IP address serves as the source IP address in the outer header of any VXLAN-encapsulated packets.

```
[edit interfaces]
user@switch# set lo0 unit 0 family inet address 10.17.17.17/32 primary
user@switch# set lo0 unit 0 family inet address 10.17.17.17/32 preferred
```

6. Set the loopback interface as the interface that identifies hardware VTEP 1.

```
[edit switch-options]
user@switch# set vtep-source-interface lo0.0
```

7. In NSX Manager, configure a logical switch for VXLAN 1 and a logical switch for VXLAN 2. See the documentation that accompanies NSX Manager.
8. In NSX Manager, configure a gateway for the Juniper Networks switch, a gateway service for OVSDB-managed interface ge-1/0/0, and a logical switch port for logical interface ge-1/0/0.10, which is associated with VXLAN 1, and a logical switch port for logical interface ge-1/0/0.20, which is associated with VXLAN 2.

See “[VMware NSX Configuration for Juniper Networks Devices Functioning as Virtual Tunnel Endpoints](#)” on page 2469.

## Verification

Confirm that the configuration is working properly.

- [Verifying the Logical Switch Configuration on page 2486](#)
- [Verifying the MAC Addresses of VM 1, VM 3, and VM 4 on page 2487](#)
- [Verifying the NSX Controller Connection on page 2487](#)
- [Verifying the OVSDB-Managed Interface on page 2488](#)

---

### Verifying the Logical Switch Configuration

**Purpose** Verify that the configuration of logical switches with the UUIDs of 28805c1d-0122-495d-85df-19abd647d772 and 9acc24b3-7b0a-4c2e-b572-3370c3e1acff are present in the OVSDB schema for physical devices and that the Flags field of the **show ovsdb logical-switch** output is Created by both.

**Action** From operational mode, enter the **show ovssdb logical-switch** command.

```
user@switch> show ovssdb logical-switch
Logical switch information:
Logical Switch Name: 28805c1d-0122-495d-85df-19abd647d772
Flags: Created by both
VNI: 100
Num of Remote MAC: 1
Num of Local MAC: 0
Logical Switch Name: 9acc24b3-7b0a-4c2e-b572-3370c3e1acff
Flags: Created by both
VNI: 200
Num of Remote MAC: 2
Num of Local MAC: 0
```

**Meaning** The output verifies that the configuration for the logical switches is present. The **Created by both** state indicates that the logical switches were configured in NSX Manager, and that the Juniper Networks switch dynamically configured the corresponding VXLANs. In this state, the logical switches and VXLANs are operational.

If the state of the logical switches is something other than **Created by both**, see [“Troubleshooting a Nonoperational Logical Switch and Corresponding Junos OS OVSDB-Managed VXLAN” on page 2507](#).

### Verifying the MAC Addresses of VM 1, VM 3, and VM 4

**Purpose** Verify that the MAC addresses of VM1, VM 3, and VM 4 are present in the OVSDB schema.

**Action** From operational mode, enter the **show ovssdb mac remote** command.

```
user@switch> show ovssdb mac remote
Logical Switch Name: 28805c1d-0122-495d-85df-19abd647d772
 Mac IP Encapsulation Vtep
 Address Address
 a8:59:5e:f6:38:90 0.0.0.0 Vxlan over Ipv4 10.17.17.17
Logical Switch Name: 9acc24b3-7b0a-4c2e-b572-3370c3e1acff
 Mac IP Encapsulation Vtep
 Address Address
 00:23:9c:5e:a7:f0 0.0.0.0 Vxlan over Ipv4 10.17.17.17
 00:23:9c:5e:a7:f0 0.0.0.0 Vxlan over Ipv4 10.17.17.17
```

**Meaning** The output shows that the MAC addresses for VM 1, VM 3, and VM 4 are present and are associated with their respective logical switches. Given that the MAC addresses are present, VM 1, VM 3, and VM 4 are reachable through the Juniper Networks switch, which functions as a hardware VTEP.

### Verifying the NSX Controller Connection

**Purpose** Verify that the connection with the NSX controller is up.

**Action** From operational mode, enter the **show ovssdb controller** command to verify that the controller connection state is **up**.

```
user@switch> show ovssdb controller
VTEP controller information:
Controller IP address: 10.94.184.1
Controller protocol: ssl
Controller port: 6632
Controller connection: up
Controller seconds-since-connect: 542325
Controller seconds-since-disconnect: 542346
Controller connection status: active
```

**Meaning** The output shows that the connection state of the NSX controller is **up**, in addition to other information about the controller. By virtue of this connection being up, OVSSDB is enabled on the Juniper Networks switch.

---

### Verifying the OVSSDB-Managed Interface

---

**Purpose** Verify that interface ge-1/0/0 is managed by OVSSDB.

**Action** From operational mode, enter the **show ovssdb interface** command, and verify that logical interfaces ge-1/0/0.10 and ge-1/0/0.20 are managed by OVSSDB.

```
user@switch> show ovssdb interface
Interface VLAN ID Bridge-domain
ge-1/0/0 10 28805c1d-0122-495d-85df-19abd647d772
ge-1/0/0 20 9acc24b3-7b0a-4c2e-b572-3370c3e1acff
```

**Meaning** The output shows that logical interfaces **ge-1/0/0.10** and **ge-1/0/0.20** are managed by OVSSDB. It also indicates that interface **ge-1/0/0.10** is associated with VXLAN **28805c1d-0122-495d-85df-19abd647d772** and interface **ge-1/0/0.20** is associated with VXLAN **9acc24b3-7b0a-4c2e-b572-3370c3e1acff**.

---

## Verifying That a Logical Switch and Corresponding Junos OS OVSSDB-Managed VXLAN Are Working Properly

---

**Purpose** Verify the following:

- A logical switch, which is configured in an NSX environment, or a virtual network, which is configured in a Contrail environment, is learning MAC addresses in their respective environments.
- The corresponding OVSSDB-managed Virtual Extensible LAN (VXLAN), which is configured on a Juniper Networks device, is learning MAC addresses in the Junos OS environment.
- The logical switch or virtual network and OVSSDB-managed VXLAN are exchanging the MAC addresses learned in their respective environments so that virtual and physical servers can communicate.



**Action** To verify that a logical switch or virtual network and its corresponding OVSDb-managed VXLAN are learning and exchanging MAC addresses in their respective environments, enter the **show ovssdb logical-switch** operational mode command.

```
user@device> show ovssdb logical-switch
Logical switch information:
Logical Switch Name: 28805c1d-0122-495d-85df-19abd647d772
Flags: Created by both
VNI: 100
Num of Remote MAC: 1
Num of Local MAC: 0
```



**NOTE:** In the Open vSwitch Database (OVSDb) schema for physical devices, the logical switch table stores information about the Layer 2 broadcast domain that you configured in a VMware NSX or Contrail environment. In the NSX environment, the Layer 2 broadcast domain is known as a *logical switch*, while in the Contrail environment, the domain is known as a *virtual network*.

In the context of the **show ovssdb logical-switch** command, the term *logical switch* refers to the logical switch or virtual network that was configured in the NSX or Contrail environments, respectively, and the corresponding configuration that was pushed to the OVSDb schema.

**Meaning** The output in the Flags field (**Created by both**) indicates that the logical switch or virtual network and its corresponding OVSDb-managed VXLAN are both properly configured. In this state, the logical switch or virtual network and the VXLAN are learning and exchanging MAC addresses in their respective environments.

If the output in the Flags field displays a state other than **Created by both**, see [“Troubleshooting a Nonoperational Logical Switch and Corresponding Junos OS OVSDb-Managed VXLAN” on page 2507](#).

**Related Documentation**

- [show ovssdb logical-switch on page 2540](#)



# Configuring VXLANs Without an SDN Controller

- [VXLAN Constraints on QFX Series Switches on page 2491](#)
- [Manually Configuring VXLANs on a QFX5100 Switch on page 2492](#)
- [Examples: Manually Configuring VXLANs on QFX Series Switches on page 2494](#)
- [Verifying That a Local VXLAN VTEP Is Configured Correctly on page 2502](#)
- [Verifying MAC Learning from a Remote VTEP on page 2502](#)

## VXLAN Constraints on QFX Series Switches

---

When configuring VXLANs on QFX switches, be aware of the constraints in the following list. In this list, “Layer 3 side” refers to a network-facing interface that performs VXLAN encapsulation and de-encapsulation, and “Layer 2 side” refers to a server-facing interface that is a member of a VLAN that is mapped to a VXLAN.

- You can use VXLANs on a Virtual Chassis or Virtual Chassis Fabric if all of the members are supported QFX switches. You cannot use VXLANs if any of the members is not a supported QFX switch.
- VXLAN configuration is supported only in the default routing instance.
- QFX switches cannot route traffic between different VXLANs.
- A physical interface cannot be a member of a VLAN and a VXLAN. That is, an interface that performs VXLAN encapsulation and de-encapsulation cannot also be a member of a VLAN. For example, if a VLAN that is mapped to a VXLAN is a member of trunk port xe-0/0/0, any other VLAN that is a member of xe-0/0/0 must also be assigned to a VXLAN.
- Multichassis link aggregation groups (MC-LAGS) are not supported with VXLAN.
- IP fragmentation and defragmentation are not supported on the Layer 3 side.
- The following features are not supported on the Layer 2 side:
  - STP (any variant)
  - IGMP snooping
  - Storm control

- Access port security features are not supported with VXLAN. For example, the following features are not supported:
  - DHCP snooping
  - Dynamic ARP inspection
  - MAC limiting and MAC move limiting
- See the following to determine whether ingress node replication is supported:
  - PIM used for control plane: ingress node replication is not supported.
  - Control plane provided by a SDN controller: ingress node replication is not supported.
  - EVPN-VXLAN: ingress node replication is supported.
- PIM-BIDIR and PIM-SSM are not supported with VXLANs.
- Class of service (CoS) features are not supported with VXLANs.
- If you configure a port-mirroring instance to mirror traffic egressing from an interface that performs VXLAN encapsulation, the source and destination MAC addresses of the mirrored packets are invalid. The original VXLAN traffic is not affected.

**Related  
Documentation**

- [Understanding VXLANs on page 2442](#)
- [Examples: Manually Configuring VXLANs on QFX Series Switches on page 2494](#)
- [Manually Configuring VXLANs on a QFX5100 Switch on page 2492](#)

---

## Manually Configuring VXLANs on a QFX5100 Switch

---

Follow these steps to configure a QFX5100 switch to act as a VTEP. (If the switch is acting as a transit Layer 3 switch for downstream VTEPs, you do not need to perform these steps. No special configuration is needed in this case.)

- [Configuring a Source IP Address on page 2493](#)
- [Configuring PIM for VXLANs on page 2493](#)
- [Configuring VXLANs on page 2493](#)

## Configuring a Source IP Address

On a switch that will act as a VTEP, you must configure an IP address that will be used as the source address in the outer IP header of the VXLAN packet. This is the VXLAN tunnel source address.

1. Create a reachable IPv4 address on the loopback interface and configure it to be used as the tunnel source address:

```
[edit]
user@switch# set interfaces lo0.0 unit 0 family inet address ip-address
[edit]
user@switch# set switch-options vtep-interface-source lo0.0
```

## Configuring PIM for VXLANs

If you are not using a controller to create a VXLAN control plane, you must enable PIM on the switch so that the VTEP can use multicast groups to establish reachability with other VTEPs and forward BUM traffic.

1. Enable PIM on the interface that connects to the Layer 3 network. This is the interface the performs the VXLAN encapsulation and de-encapsulation.

```
[edit]
user@switch# set protocols pim interface interface-name
```

2. Configure the address of a PIM rendezvous point:

```
[edit]
user@switch# set protocols pim rp static address ip-address
```

## Configuring VXLANs

You configure VXLANs under the **vlan** stanza (which is why a QFX5100 switch supports 4K VLANs). You must also configure the server-facing interfaces to be VLAN members.

1. Create a VLAN to VXLAN mapping and assign a multicast group address to the VXLAN. All members of a VXLAN must use the same multicast group address:

```
[edit]
user@switch# set vlans name vlan-id ID vxlan vni ID multicast-group multicast-group-address
```

2. (Optional) Configure the switch to retain the original VLAN tag (in the inner Ethernet packet) after VXLAN encapsulation. By default, the original tag is dropped when the packet is encapsulated:

```
[edit]
user@switch# set vlans name vxlan encapsulate-inner-vlan
```

3. (Optional) Configure the switch to de-encapsulate and accept original VLAN tags in VXLAN packets. By default, the original tag is dropped when the packet is encapsulated:

```
[edit]
user@switch# set protocols l2-learning decapsulate-accept-inner-vlan
```

4. Configure server-facing interfaces to support multiple VLANs:

```
[edit]
user@switch# set interfaces interface unit unit family ethernet-switching interface-mode trunk
[edit]
```

```
user@switch# set interfaces interface unit unit family ethernet-switching vlan members all
```

You must create a VLAN to VXLAN mapping for each VLAN that will need Layer 2 connectivity over the Layer 3 network.

- Related Documentation**
- [Understanding VXLANs on page 2442](#)
  - [Examples: Manually Configuring VXLANs on QFX Series Switches on page 2494](#)

---

## Examples: Manually Configuring VXLANs on QFX Series Switches

These examples show how to configure VXLANs on QFX Series Switches for several use cases.

- [Example: Configuring a VXLAN Transit Switch on page 2494](#)
- [Example: Configuring a VXLAN Layer 2 Gateway on page 2496](#)

### Example: Configuring a VXLAN Transit Switch

If a QFX switch acts as a transit switch for downstream devices acting as VTEPs, you do not need to configure any VXLAN information on the QFX switch. You do need to configure PIM on the switch so that it can form the multicast tree required so that the VTEPs can establish reachability with each other.

- [Requirements on page 2494](#)
- [Overview on page 2494](#)
- [Configuring PIM on the Transit Switches on page 2495](#)

---

#### Requirements

This example uses the following hardware and software components:

- Two QFX5100 switches
- Junos OS 14.1X53-D10

---

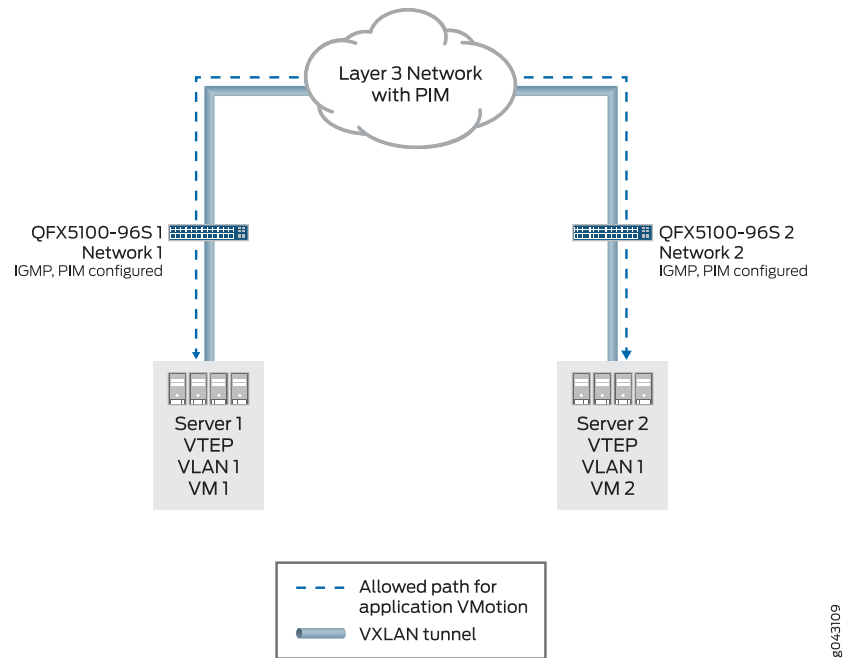
#### Overview

This example shows a simple use case in which QFX switches are connected to downstream servers acting as VTEPs and need to forward VXLAN packets between VM 1 on Server 1 and VM 2 on Server 2. Because this configuration allows Layer 2 connectivity between the VMs through the VXLAN tunnels, applications can VMotion between the VMs.

#### Topology

[Figure 35](#) shows a QFX 5100 switch configured to forward VXLAN packets for downstream VTEPs.

Figure 35: QFX5100 Acting as a VXLAN Transit Switch



### Configuring PIM on the Transit Switches

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set protocols pim interface all
set protocols pim rp static address ip-address
```

#### Step-by-Step Procedure

If you are not using a controller to create a VXLAN control plane, you must enable PIM on each switch so that the VTEP can use multicast groups to advertise its existence and to learn about other VTEPs. (Configuring PIM automatically enables IGMP.) You do not need to perform any VXLAN-specific configuration. Note that you also do not need to configure VLAN 1 or 2 on either switch.

1. Enable PIM:
 

```
[edit]
user@switch# set protocols pim interface all
```
2. Configure the address of a PIM rendezvous point:
 

```
[edit]
user@switch# set protocols pim rp static address ip-address
```

## Example: Configuring a VXLAN Layer 2 Gateway

If a QFX switch is connected to a downstream server that hosts a VM that needs Layer 2 connectivity with another VM that is reachable only through a Layer 3 network, you must configure the switch to act as a VTEP—that is, a Layer 2 gateway for downstream Layer 2 devices. You also need to configure PIM on the switch so that it can form the multicast tree required for reachability with other VTEPs and to allow BUM traffic to be forwarded between the VTEPs.

- [Requirements on page 2496](#)
- [Overview on page 2496](#)
- [Configuring the Switches on page 2497](#)
- [Verification on page 2500](#)

---

### Requirements

This example uses the following hardware and software components:

- Two QFX5100 switches
- Junos OS 14.1X53-D10

---

### Overview

This example shows a use case in which QFX switches are connected to downstream VTEPs and need to allow Layer 2 connectivity between VM 1 on Server 1 and VM 2 on Server 2 so that VMotion can occur between the VMs. The servers in this example can be in the same or different data centers—the only constraint is that there must be Layer 3 connectivity between the QFX switches. This allows your network to be very agile in response to demand for server usage or changes in bandwidth requirements.


Note that because the same VLAN exists in both Layer 2 domains and both switches encapsulate the VLAN traffic into the same VXLAN, you do not need a gateway for the VXLAN traffic in the Layer 3 network. The Layer 3 VXLAN packets are routed normally and no de-encapsulation or re-encapsulation is required..

### Topology

[Figure 36](#) shows a QFX 5100 switch configured to act as a VTEP.



Figure 36: QFX5100 Acting as a VTEP



8043108

### Configuring the Switches

#### CLI Quick Configuration

To quickly configure the QFX5100-96S 1 in this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.

```
set interfaces lo0 unit 0 family inet address 10.1.1.1
set switch-options vtep-source-interface lo0.0
set protocols pim interface lo0.0
set protocols pim interface xe-0/0/0.0
set protocols pim rp static address 10.2.2.2
set vlans VLAN1 vlan-id 100 vxlan vni 100 multicast-group 224.2.2.2
set vlans VLAN1 vxlan encapsulate-inner-vlan
set vlans VLAN1 vxlan unreachable-vtep-aging-timer 600
set protocols l2-learning decapsulate-accept-inner-vlan
set interfaces xe-0/0/0 unit 0 family inet address 10.2.2.100/24
set interfaces xe-0/0/1 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members all
```

The configuration for QFX5100-96S 2 is almost identical. The only changes a few of the addresses:

```
set interfaces lo0 unit 0 family inet address 10.1.1.2
set switch-options vtep-source-interface lo0.0
set protocols pim interface lo0.0
set protocols pim interface xe-0/0/0.0
set protocols pim rp static address 10.2.2.2
set vlans VLAN1 vlan-id 100 vxlan vni 100 multicast-group 224.2.2.2
set vlans VLAN1 vxlan encapsulate-inner-vlan
```

```

set vlans VLAN1 vxlan unreachable-vtep-aging-timer 600
set protocols l2-learning decapsulate-accept-inner-vlan
set interfaces xe-0/0/0 unit 0 family inet address 10.2.2.200/24
set interfaces xe-0/0/1 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members all

```



**NOTE:** You must configure the same multicast group address for VLAN1 on both switches.

### Step-by-Step Procedure

Perform the following procedure on both switches to set up the example configuration. Note that you do not need to configure VLAN 1 or 2 on either switch.

1. Create a reachable IPv4 address on the loopback interface.

```

[edit]
user@switch# set interfaces lo0 unit 0 family inet address 10.1.1.1
For switch QFX5100-96S 2, use address 10.1.1.2.

```

2. Configure the loopback interface—and therefore, its associated address—to be used as the tunnel source address:

```

[edit]
user@switch# set switch-options vtep-source-interface lo0.0

```

3. Enable PIM on the loopback interface:

```

[edit]
user@switch# set protocols pim interface lo0.0

```

4. Enable PIM on the interface that connects to the Layer 3 network:

```

[edit]
user@switch# set protocols pim interface xe-0/0/0.0

```

5. Configure the address of a PIM rendezvous point:

```

[edit]
user@switch# set protocols pim rp static address 10.2.2.2

```

6. Create a VLAN, map it to a VXLAN, and assign a multicast group address to the VXLAN. All members of a VXLAN must use the same multicast group address:

```

[edit]
user@switch# set vlans VLAN1 vlan-id 100 vxlan vni 100 multicast-group 224.2.2.2

```

In this example, the **vlan-id** and **vni** are both set to 100. This is done only for simplicity and clarity. You do not need to set the **vlan-id** and **vni** to the same value.

7. (Optional) Configure the switch to retain the original VLAN tag (in the inner Ethernet packet) after VXLAN encapsulation. By default, the original tag is dropped when the packet is encapsulated:

```

[edit]
user@switch# set vlans VLAN1 vxlan encapsulate-inner-vlan

```

8. (Optional) Configure the system to age out the address for the remote VTEP (the other QFX5100 switch) if all the MAC addresses learned from that VTEP age out. The address for the remote VTEP expires the configured number of seconds after the last learned MAC address expires.

```

[edit]

```

```
user@switch# set vlans VLAN1 vxlan unreachable-vtep-aging-timer 600
```

(Optional) Configure the switch to de-encapsulate and accept original VLAN tags in VXLAN packets. By default, a preserved VLAN tag is dropped when the packet is de-encapsulated:

```
[edit]
```

```
user@switch# set protocols l2-learning decapsulate-accept-inner-vlan
```

9. Configure the interface that connects to the Layer 3 network:

```
[edit]
```

```
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.2.2.100/24
```

For switch QFX5100-96S 2, use address 10.2.2.200.

10. Configure the server-facing interface to support multiple VLANs:

```
[edit]
```

```
user@switch# set interfaces xe-0/0/1 unit 0 family ethernet-switching interface-mode trunk
```

```
[edit]
```

```
user@switch# set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members all
```



**NOTE:** Because this example shows only one VLAN, this step is not required for the example. In a real-world configuration, however, it would be required in order to support multiple VMs connected to multiple VLANs. In this case you would also need to configure additional VLAN to VXLAN mappings.

### Results

From configuration mode, confirm your configuration by entering the following commands on QFX5100-96S 1. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@switch# show switch-options
```

```
 vtep-source-interface lo0.0;
```

```
user@switch# show vlans
```

```
VLAN1 {
 vlan-id 100;
 vxlan {
 vni 100;
 multicast-group 224.2.2.2;
 encapsulate-inner-vlan;
 }
}
```

```
user@switch# show interfaces
```

```
xe-0/0/0 {
 unit 0 {
 family inet {
 address 10.2.2.100/24;
 }
 }
}
```

```

}
xe-0/0/1 {
 unit 0 {
 family ethernet-switching {
 interface-mode trunk;
 vlan {
 members all;
 }
 }
 }
}
lo0 {
 unit 0 {
 family inet {
 address 10.1.1.1/32;
 }
 }
}

```

```
user@switch# show protocols pim
```

```

rp {
 static {
 address 10.2.2.2;
 }
}
interface xe-0/0/0.0

```

## Verification

Confirm that the configuration is working properly.

- [Verifying VXLAN Reachability on page 2500](#)
- [Verifying That the Local VTEP is Configured Correctly on page 2501](#)
- [Verifying MAC Learning from the Remote VTEP on page 2501](#)
- [Monitor the Remote Interface on page 2501](#)

### Verifying VXLAN Reachability

**Purpose** On QFX5100-96S 1, verify that there is connectivity with the remote VTEP (QFX5100-96S 2).

**Action** user@switch> show ethernet-switching vxlan-tunnel-end-point remote

| Logical System Name | Id          | SVTEP-IP | IFL   | L3-Idx |
|---------------------|-------------|----------|-------|--------|
| <default>           | 0           | 10.1.1.2 | lo0.0 | 0      |
| RVTEP-IP            | IFL-Idx     | NH-Id    |       |        |
| 10.1.1.2            | 559         | 1728     |       |        |
| VNID                | MC-Group-IP |          |       |        |
| 100                 | 224.2.2.2   |          |       |        |

**Meaning** The VTEP on QFX5100-96S 2 is reachable because its IP address (the address assigned to the loopback interface) appears in the output. The output also shows that the VXLAN (VNI 100) and corresponding multicast group are configured correctly on the remote VTEP.

***Verifying That the Local VTEP is Configured Correctly***

**Purpose** On QFX5100-96S 1, verify that the tunnel endpoint is correct..

**Action** user@switch> show ethernet-switching vxlan-tunnel-end-point source

|                     |    |               |       |             |
|---------------------|----|---------------|-------|-------------|
| Logical System Name | Id | SVTEP-IP      | IFL   | L3-Idx      |
| <default>           | 0  | 10.1.1.1      | lo0.0 | 0           |
| L2-RTT              |    | Bridge Domain | VNID  | MC-Group-IP |
| default-switch      |    | VLAN1+100     | 100   | 224.2.2.2   |

**Meaning** The VTEP on QFX5100-96S 1 shows the correct tunnel source IP address (assigned to the loopback interface), VLAN, and multicast group for the VXLAN.

***Verifying MAC Learning from the Remote VTEP***

**Purpose** On QFX5100-96S 1, verify that it is learning MAC addresses from the remote VTEP.

**Action** user@switch> show ethernet-switching table

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static

SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Ethernet switching table : 2 entries, 2 learned

Routing instance : default-switch

| Vlan name | MAC address       | MAC flags | Age | Logical interface |
|-----------|-------------------|-----------|-----|-------------------|
| VLAN1     | 00:00:00:ff:ff:ff | D         | -   | vtep.12345        |
| VLAN1     | 00:10:94:00:00:02 | D         | -   | xe-0/0/0.0        |

**Meaning** This shows the MAC addresses learned from the remote VTEP (in addition to those learned on the normal Layer 2 interfaces). It also shows the logical name of the remote VTEP interface (**vtep.12345** in the above output).

***Monitor the Remote Interface***

**Purpose** On QFX5100-96S 1, monitor traffic details for the remote VTEP interface.

**Action** user@switch> show interface vtep.12345 detail

```
M Flags: Up SNMP-Traps Encapsulation: ENET2
 VXLAN Endpoint Type: Remote, VXLAN Endpoint Address: 10.1.1.2, L2 Routing
Instance: default-switch, L3 Routing Instance: default
 Traffic statistics:
 Input bytes : 228851738624
 Output bytes : 0
 Input packets: 714162415
 Output packets: 0
 Local statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
 Transit statistics:
 Input bytes : 228851738624 0 bps
 Output bytes : 0 0 bps
 Input packets: 714162415 0 pps
 Output packets: 0 0 pps
 Protocol eth-switch, MTU: 1600, Generation: 277, Route table: 5
```

**Meaning** This shows traffic details for the remote VTEP interface. To get this information, you must supply the logical name of the remote VTEP interface (vtep.12345 in the above output), which you can learn by using the **show ethernet-switching table** command.

**Related Documentation**

- [Understanding VXLANs on page 2442](#)
- [VXLAN Constraints on QFX Series Switches on page 2491](#)

---

## Verifying That a Local VXLAN VTEP Is Configured Correctly

---

**Purpose** Verify that a local VTEP is correct..

**Action** user@switch> show ethernet-switching vxlan-tunnel-end-point source

| Logical System Name | Id            | SVTEP-IP | IFL   | L3-Idx      |
|---------------------|---------------|----------|-------|-------------|
| <default>           | 0             | 10.1.1.1 | 1o0.0 | 0           |
| L2-RTT              | Bridge Domain |          | VNID  | MC-Group-IP |
| default-switch      | VLAN1+100     |          | 100   | 232.1.1.1   |

**Meaning** The output should show the correct tunnel source IP address (loopback address), VLAN, and multicast group for the VXLAN.

**Related Documentation**

- [Understanding VXLANs on page 2442](#)

---

## Verifying MAC Learning from a Remote VTEP

---

**Purpose** Verify that a local VTEP is learning MAC addresses from a remote VTEP.

**Action** user@switch> show ethernet-switching table

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static

SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Ethernet switching table : 2 entries, 2 learned

Routing instance : default-switch

| Vlan<br>name | MAC<br>address    | MAC<br>flags | Age | Logical<br>interface |
|--------------|-------------------|--------------|-----|----------------------|
| VLAN1        | 00:00:00:ff:ff:ff | D            | -   | vtep.12345           |
| VLAN1        | 00:10:94:00:00:02 | D            | -   | xe-0/0/0.0           |

**Meaning** This shows the MAC addresses learned from the remote VTEP (in addition to those learned on the normal Layer 2 interfaces). It also shows the logical name of the remote VTEP interface (**vtep.12345** in the above output).

**Related Documentation**

- [Understanding VXLANs on page 2442](#)





## PART 38

# Troubleshooting

- [Troubleshooting Tasks on page 2507](#)



# Troubleshooting Tasks

- [Troubleshooting a Nonoperational Logical Switch and Corresponding Junos OS OVSDb-Managed VXLAN on page 2507](#)
- [Verifying VXLAN Reachability on page 2509](#)
- [Monitoring a Remote VTEP Interface on page 2509](#)

## Troubleshooting a Nonoperational Logical Switch and Corresponding Junos OS OVSDb-Managed VXLAN

---

|                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Problem</b>  | <p><b>Description:</b> The Flags field in the <code>show ovssdb logical-switch</code> operational mode command output is one of the following:</p> <ul style="list-style-type: none"><li>• <b>Created by Controller</b></li><li>• <b>Created by L2ALD</b></li><li>• <b>Tunnel key mismatch</b></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Cause</b>    | <ul style="list-style-type: none"><li>• If the Flags field displays <b>Created by Controller</b>, a logical switch is configured in the NSX environment or a virtual network is configured in the Contrail environment. However, an equivalent VXLAN is not configured or is improperly configured on the Juniper Networks device.</li><li>• If the Flags field displays <b>Created by L2ALD</b>, a VXLAN is configured on the Juniper Networks device. However, an equivalent logical switch is not configured in the NSX environment or an equivalent virtual network is not configured in the Contrail environment.</li><li>• If the Flags field displays <b>Tunnel key mismatch</b>, the VXLAN network identifier (VNI) specified in the logical switch configuration or the VXLAN identifier specified in the virtual network configuration do not match the VNI in the equivalent VXLAN configuration.</li></ul> |
| <b>Solution</b> | <p>If the Flags field displays <b>Created by Controller</b>, take the following action:</p> <ul style="list-style-type: none"><li>• On a QFX Series switch, verify that the <code>set switch-options ovssdb-managed</code> configuration command was issued in the Junos OS CLI. Issuing this command and committing the configuration enable the Juniper Networks device to dynamically create OVSDb-managed VXLANs.</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

Another possible cause is that the L2ALD daemon has become nonfunctional. If this is the case, wait for a few seconds, reissue the **show ovsdb logical-switch** operational mode command, and recheck the setting of the Flags field.

Another possible cause is that the Juniper Networks device dynamically configured the VXLAN and its associated logical interface, but there is an error in the configuration of these entities themselves or in an entity that was committed in the same transaction. If there is an issue with one or more of the configurations in a transaction, all configurations in the transaction, even the ones that are correctly configured, remain uncommitted and in a queue until you troubleshoot and resolve the configuration issues. As a result, the Juniper Networks device was unable to commit all configurations in the transaction. For this situation, enter the **show ovsdb commit failures** operational mode command. In the output that is displayed, determine which configurations are erroneous. Issues that can cause commitment errors include but are not limited to the detection of the same VXLAN name or VXLAN network identifier (VNI) in a dynamically configured VXLAN and in a VXLAN that was previously configured using the Junos OS CLI. After resolving the errors, enter the **clear ovsdb commit failures** command to remove the transaction from the queue and then retry committing all configurations in the transaction.

- On all other Juniper Networks devices that support VXLAN and OVSDDB, determine whether a VXLAN equivalent to the logical switch configuration or virtual network configuration exists on the device. If the VXLAN is not configured, configure it using the procedure in *Configuring OVSDDB-Managed VXLANs*. If a VXLAN is configured, check the VXLAN name to make sure that it is the same as the universally unique identifier (UUID) of the logical switch (NSX) or virtual network (Contrail) configuration. Also, check the VNI to make sure that the value is the same as the value in the logical switch (NSX) or virtual network (Contrail) configuration.

If the Flags field displays **Created by L2ALD**, take the following action:

- On a QFX Series switch, two issues exist. First, despite the fact that the Juniper Networks device dynamically creates OVSDDB-managed VXLANs, this VXLAN was configured by using the Junos OS CLI. Second, a corresponding logical switch (NSX) or virtual network (Contrail) was not configured. To resolve both issues, configure a logical switch in the NSX environment or a virtual network in the Contrail environment. After the software-defined networking (SDN) controller pushes relevant logical switch or virtual network information to the Juniper Networks device, the device dynamically creates a corresponding VXLAN and deletes the VXLAN configured using the Junos OS CLI.
- On all other Juniper Networks devices that support VXLAN and OVSDDB, determine whether an equivalent logical switch is configured in the NSX environment or a virtual network is configured in the Contrail environment. If a logical switch or virtual network is not configured, configure one, keeping in mind that a UUID is automatically generated for the logical switch or virtual network and that this UUID must be used as the name of the VXLAN. That is, the VXLAN name must be reconfigured with the logical switch or virtual network UUID.

Another possibility is that the logical switch or virtual network configuration might exist, but the UUID of the entity might not match the VXLAN name. In the NSX or

Contrail environment, check for a logical switch or virtual network, respectively, that has the same configuration as the VXLAN but has a different UUID.

If the Flags field displays **Tunnel key mismatch**, take the following action:

- For a QFX Series switch, check the configuration of the VNI in the NSX environment or the VXLAN identifier in the Contrail environment to see whether it was changed after the Juniper Networks device dynamically created the equivalent VXLAN. If it was changed, update the VNI on the QFX Series switch using the Junos OS CLI.
- On all other Juniper Networks devices that support VXLAN and OVSDb, check the value of the VNI in the NSX environment or the VXLAN identifier in the Contrail environment and the Junos OS CLI. Change the incorrect value.

- Related Documentation**
- [Understanding Dynamically Configured VXLANs in an OVSDb Environment on page 2462](#)
  - [Understanding How to Manually Configure OVSDb-Managed VXLANs](#)
  - [show ovssdb logical-switch on page 2540](#)
  - [show ovssdb commit failures on page 2534](#)
  - [clear ovssdb commit failures on page 2532](#)

## Verifying VXLAN Reachability

**Purpose** On the local VTEP, verify that there is connectivity with the remote VTEP.

**Action** `user@switch> show ethernet-switching vxlan-tunnel-end-point remote`

| Logical System Name | Id          | SVTEP-IP | IFL   | L3-Idx |
|---------------------|-------------|----------|-------|--------|
| <default>           | 0           | 10.1.1.2 | 100.0 | 0      |
| RVTEP-IP            | IFL-Idx     | NH-Id    |       |        |
| 10.1.1.2            | 559         | 1728     |       |        |
| VNID                | MC-Group-IP |          |       |        |
| 100                 | 232.1.1.1   |          |       |        |

**Meaning** The remote VTEP is reachable because its IP address appears in the output. The output also shows that the VXLAN (VNI 100) and corresponding multicast group are configured correctly on the remote VTEP.

- Related Documentation**
- [Understanding VXLANs on page 2442](#)
  - [Manually Configuring VXLANs on a QFX5100 Switch on page 2492](#)
  - [Examples: Manually Configuring VXLANs on QFX Series Switches on page 2494](#)

## Monitoring a Remote VTEP Interface

**Purpose** Monitor traffic details for a remote VTEP interface.

**Action** user@switch> show interface *logical-name* detail

```
M Flags: Up SNMP-Traps Encapsulation: ENET2
 VXLAN Endpoint Type: Remote, VXLAN Endpoint Address: 10.1.1.2, L2 Routing
Instance: default-switch, L3 Routing Instance: default
 Traffic statistics:
 Input bytes : 228851738624
 Output bytes : 0
 Input packets: 714162415
 Output packets: 0
 Local statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
 Transit statistics:
 Input bytes : 228851738624 0 bps
 Output bytes : 0 0 bps
 Input packets: 714162415 0 pps
 Output packets: 0 0 pps
 Protocol eth-switch, MTU: 1600, Generation: 277, Route table: 5
```

**Meaning** This shows traffic details for the remote VTEP interface. To get this information, you must supply the logical name of the remote VTEP interface (vtep.12345 in the above output), which you can learn by using the **show ethernet-switching table** command.

**Related Documentation**

- [Understanding VXLANs on page 2442](#)
- [Manually Configuring VXLANs on a QFX5100 Switch on page 2492](#)
- [Examples: Manually Configuring VXLANs on QFX Series Switches on page 2494](#)

## PART 39

# Configuration Statements and Operational Commands

- [OVSDB Configuration Statements on page 2513](#)
- [VXLAN Configuration Statements on page 2525](#)
- [OVSDB Operational Commands on page 2531](#)
- [VXLAN Operational Commands on page 2557](#)





## CHAPTER 92

# OVSDb Configuration Statements

- [controller \(OVSDb\) on page 2514](#)
- [inactivity-probe-duration on page 2515](#)
- [interfaces \(OVSDb\) on page 2516](#)
- [maximum-backoff-duration on page 2517](#)
- [ovsdb on page 2518](#)
- [ovsdb-managed on page 2519](#)
- [port \(OVSDb\) on page 2520](#)
- [protocol \(OVSDb\) on page 2521](#)
- [traceoptions \(OVSDb\) on page 2522](#)

## controller (OVSDB)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>controller ip-address {<br/>    inactivity-probe-duration milliseconds;<br/>    maximum-backoff-duration milliseconds;<br/>    protocol protocol {<br/>        port number;<br/>    }<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit protocols <b>ovsdb</b> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 14.1R2.</p> <p>Statement introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.</p> <p>Statement introduced in Junos OS Release 14.2 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | <p>Configure a connection between a Juniper Networks device running the Open vSwitch Database (OVSDB) management protocol and a software-defined networking (SDN) controller. You can connect a Juniper Networks device to more than one SDN controller for redundancy.</p> <p>In a VMware NSX environment, one cluster of NSX controllers typically includes three or five controllers. To implement the OVSDB management protocol on a Juniper Networks device, you must explicitly configure a connection to one NSX controller, using the Junos OS CLI. If the NSX controller to which you explicitly configure a connection is in a cluster, the controller pushes information about other controllers in the same cluster to the device, and the device establishes connections with the other controllers. However, you can also explicitly configure connections with the other controllers in the cluster, using the Junos OS CLI.</p> <p>To implement the OVSDB management protocol on a Juniper Networks device in a Contrail environment, you must configure a connection to a Contrail controller, using the Junos OS CLI.</p> <p>Connections to all SDN controllers are made on the management interface of the Juniper Networks device.</p> |
| <b>Options</b>                  | <p><b>ip-address</b>—IPv4 address of the SDN controller.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | <p><b>admin</b>—To view this statement in the configuration.</p> <p><b>admin-control</b>—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Setting Up the OVSDB Protocol on Juniper Networks Devices that Support the Dynamic Configuration of VXLANs on page 2461</a></li><li>• <a href="#">Setting Up the OVSDB Protocol on Juniper Networks Devices that Support Manual Configuration of VXLANs</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

- [Understanding How to Set Up OVSDB Connections Between Juniper Networks Devices and SDN Controllers on page 2458](#)

## inactivity-probe-duration

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>inactivity-probe-duration <i>milliseconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">ovsdb controller</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1R2.<br>Statement introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.<br>Statement introduced in Junos OS Release 14.2 for EX Series switches.                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Configure the maximum amount of time, in milliseconds, that the connection between a Juniper Networks device that supports the Open vSwitch Database (OVSDB) management protocol and a software-defined networking (SDN) controller can be inactive before an inactivity probe is sent.                                                                                                                                                                                 |
| <b>Options</b>                  | <i>milliseconds</i> —Number of milliseconds that the connection can be inactive before an inactivity probe is sent.<br><b>Range:</b> 0 through 4,294,967,295<br><b>Default:</b> 0. This value indicates that an inactivity probe is never sent.                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Setting Up the OVSDB Protocol on Juniper Networks Devices that Support the Dynamic Configuration of VXLANs on page 2461</a></li> <li>• <a href="#">Setting Up the OVSDB Protocol on Juniper Networks Devices that Support Manual Configuration of VXLANs</a></li> <li>• <a href="#">Understanding How to Set Up OVSDB Connections Between Juniper Networks Devices and SDN Controllers on page 2458</a></li> </ul> |

## interfaces (OVSDB)

---

|                                 |                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>interfaces <i>interface-name</i>;</code>                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">ovsdb</a> ]                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1R2.<br>Statement introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.<br>Statement introduced in Junos OS Release 14.2 for EX Series switches.                                                                                                       |
| <b>Description</b>              | Specify the physical interfaces on a Juniper Networks device that you want the Open vSwitch Database (OVSDB) management protocol to manage. Typically, the only interfaces that need to be managed by OVSDB are interfaces that are connected to physical servers.                                               |
| <b>Options</b>                  | <i>interface-name</i> —Name of the interface.                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Setting Up the OVSDB Protocol on Juniper Networks Devices that Support Manual Configuration of VXLANs</i></li><li>• <a href="#">Setting Up the OVSDB Protocol on Juniper Networks Devices that Support the Dynamic Configuration of VXLANs on page 2461</a></li></ul> |

## maximum-backoff-duration

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | maximum-backoff-duration <i>milliseconds</i> ;                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">ovsdb controller</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1R2.<br>Statement introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.<br>Statement introduced in Junos OS Release 14.2 for EX Series switches.                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Specify (in milliseconds) how long a Juniper Networks device that supports the Open vSwitch Database (OVSDb) management protocol waits before it tries again to connect with a software-defined networking (SDN) controller after a previous attempt has failed.                                                                                                                                                                                                        |
| <b>Options</b>                  | <i>milliseconds</i> —Number of milliseconds a Juniper Networks device waits before it tries again to connect with an SDN controller.<br><b>Range:</b> 1000 through 4,294,967,295<br><b>Default:</b> 1000                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Setting Up the OVSDb Protocol on Juniper Networks Devices that Support the Dynamic Configuration of VXLANs on page 2461</a></li> <li>• <a href="#">Setting Up the OVSDb Protocol on Juniper Networks Devices that Support Manual Configuration of VXLANs</a></li> <li>• <a href="#">Understanding How to Set Up OVSDb Connections Between Juniper Networks Devices and SDN Controllers on page 2458</a></li> </ul> |

## ovsdb

---

**Syntax**    ovsdb {  
              controller *ip-address* {  
                  inactivity-probe-duration *milliseconds*;  
                  maximum-backoff-duration *milliseconds*;  
              protocol *protocol* {  
                  port *number*;  
              }  
              }  
              interfaces *interface-name*;  
              traceoptions {  
                  file <*filename*> <files *number*> <match *regular-expression*> <no-world-readable |  
                    world-readable> <size *size*>;  
                  flag *flag*;  
                  no-remote-trace;  
              }  
              }

**Hierarchy Level**    [edit protocols]

**Release Information**    Statement introduced in Junos OS Release 14.1R2.  
                              Statement introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.  
                              Statement introduced in Junos OS Release 14.2 for EX Series switches.

**Description**    Configure support for the Open vSwitch Database (OVSDb) management protocol on a Juniper Networks device.

                      The remaining statements are explained separately.

**Default**    The OVSDb management protocol is disabled on Juniper Networks devices.

**Required Privilege Level**    admin—To view this statement in the configuration.  
                                  admin-control—To add this statement to the configuration.

**Related Documentation**    • [Understanding the OVSDb Protocol Running on Juniper Networks Devices on page 2445](#)

## ovsdb-managed

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | ovsdb-managed;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | <p>[edit bridge-domains <i>bridge-domain-name</i> vxlan],<br/>         [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> vxlan],<br/>         [edit routing-instances <i>routing-instance-name</i> switch-options],<br/>         [edit routing-instances <i>routing-instance-name</i> vlans <i>vlan-name</i> vxlan],<br/>         [edit routing-instances <i>routing-instance-name</i> vxlan],<br/>         [edit switch-options],<br/>         [edit vlans <i>vlan-name</i> vxlan]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 14.1R2.</p> <p>Statement introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.</p> <p>Statement introduced in Junos OS Release 14.2 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | <p>Disable a Juniper Networks device from learning about other Juniper Networks devices that function as hardware virtual tunnel endpoints (VTEPs) in a specified Virtual Extensible LAN (VXLAN) and the MAC addresses learned by the hardware VTEPs. Instead, the Juniper Networks device uses the Open vSwitch Database (OVSDb) management protocol to learn about the hardware VTEPs in the VXLAN and the media access control (MAC) addresses learned by the hardware VTEPs.</p> <p>The specified VXLAN must have a VXLAN network identifier (VNI) configured, using the <b>vni</b> statement in the [edit bridge-domains <i>bridge-domain-name</i> vxlan], [edit routing-instance <i>routing-instance-name</i> vxlan], or [edit vlans <i>vlan-name</i> vxlan] hierarchy.</p> <p>Also, for OVSDb-managed VXLANs, the multicast scheme described in “<a href="#">Understanding How Layer 2 BUM and Layer 3 Routed Multicast Traffic Are Handled with OVSDb</a>” on <a href="#">page 2450</a> is used. Therefore, specifying the <b>multicast-group</b> statement in the [edit bridge-domains <i>bridge-domain-name</i> vxlan], [edit routing-instances <i>routing-instance-name</i> vxlan], or [edit vlans <i>vlan-name</i> vxlan] hierarchy has no effect.</p> |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding Dynamically Configured VXLANs in an OVSDb Environment on page 2462</a></li> <li>• <a href="#">Configuring OVSDb-Managed VXLANs</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |


## port (OVSDB)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>port <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">ovsdb controller protocol</a> ]                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1R2.<br>Statement introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.<br>Statement introduced in Junos OS Release 14.2 for EX Series switches.                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Specify the software-defined networking (SDN) controller port to which a Juniper Networks device that supports the Open vSwitch Database (OVSDB) management protocol connects.                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <i>number</i> —Number of the SDN controller port.<br><b>Range:</b> 1024 through 65,535<br><b>Default:</b> 6632                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Setting Up the OVSDB Protocol on Juniper Networks Devices that Support the Dynamic Configuration of VXLANs on page 2461</a></li><li>• <a href="#">Setting Up the OVSDB Protocol on Juniper Networks Devices that Support Manual Configuration of VXLANs</a></li><li>• <a href="#">Understanding How to Set Up OVSDB Connections Between Juniper Networks Devices and SDN Controllers on page 2458</a></li></ul> |



## protocol (OVSDB)

|                                                                                                                                                                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                        | <code>protocol protocol {<br/>    port number;<br/>}</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>                                                                                                                                               | [edit protocols <a href="#">ovsdb controller</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>                                                                                                                                           | Statement introduced in Junos OS Release 14.1R2.<br>Statement introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.<br>Statement introduced in Junos OS Release 14.2 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>                                                                                                                                                   | <p>Configure the security protocol that protects the connection between a Juniper Networks device that supports the Open vSwitch Database (OVSDB) management protocol and a software-defined networking (SDN) controller.</p> <p>The Secure Sockets Layer (SSL) connection requires a private key and certificates, which must be stored in the <code>/var/db/certs</code> directory of the Juniper Networks device. See <a href="#">“Creating and Installing an SSL Key and Certificate on a Juniper Networks Device for Connection with SDN Controllers”</a> on page 2459.</p> |
| <b>Options</b>                                                                                                                                                       | <i>protocol</i> —Establish a secure connection to the SDN controller, using SSL or TCP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <div>  <b>NOTE:</b> SSL is the only supported connection protocol.         </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Default:</b> <code>ssl</code>                                                                                                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| The remaining statement is explained separately.                                                                                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b>                                                                                                                                      | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>                                                                                                                                         | <ul style="list-style-type: none"> <li>• <a href="#">Setting Up the OVSDB Protocol on Juniper Networks Devices that Support the Dynamic Configuration of VXLANs on page 2461</a></li> <li>• <a href="#">Setting Up the OVSDB Protocol on Juniper Networks Devices that Support Manual Configuration of VXLANs</a></li> <li>• <a href="#">Understanding How to Set Up OVSDB Connections Between Juniper Networks Devices and SDN Controllers on page 2458</a></li> </ul>                                                                                                          |

## traceoptions (OVSDB)

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>traceoptions {<br/>    file &lt;filename&gt; &lt;files number&gt; &lt;match regular-expression&gt; &lt;no-world-readable  <br/>    world-readable&gt; &lt;size size&gt;;<br/>    flag flag;<br/>    no-remote-trace;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>     | [edit protocols <a href="#">ovsdb</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b> | Statement introduced in Junos OS Release 14.1R2.<br>Statement introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.<br>Statement introduced in Junos OS Release 14.2 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>         | Define tracing operations for the Open vSwitch Database (OVSDB) management protocol, which is supported on Juniper Networks devices.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Default</b>             | If you do not include this statement, OVSDB-specific tracing operations are not performed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>             | <p><b>file <i>filename</i></b>—Name of file in which the system places the output of the tracing operations. By default, the system places all files in the <b>/var/log</b> directory.</p> <p><b>Default:</b> <b>/var/log/vgd</b></p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file reaches the size specified by the <b>size</b> option, the filename is appended with 0 and compressed. For example, a trace file named <b>trace-file.gz</b> would be renamed <b>trace-file.0.gz</b>. When <b>trace-file.0.gz</b> reaches the specified size, it is renamed <b>trace-file.1.gz</b> and its contents are compressed to <b>trace-file.0.gz</b>. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the <b>size</b> option and a filename.</p> <p><b>Range:</b> 2 through 1000 files</p> <p><b>Default:</b> 10 files</p> <p><b>flag <i>flag</i></b>—Tracing operation to perform. You can include one or more of the following flags:</p> <ul style="list-style-type: none"><li><b>all</b>—All OVSDB events.</li><li><b>configuration</b>—OVSDB configuration events.</li><li><b>core</b>—OVSDB core events.</li><li><b>function</b>—OVSDB function events.</li><li><b>interface</b>—OVSDB interface events.</li><li><b>l2-client</b>—OVSDB Layer 2 client events.</li></ul> |

**netconf-client**—(QFX Series switches only) Events for the dynamic configuration of Virtual Extensible LANs (VXLANs).

**ovs-client**—OVSDDB client events.

**match *regular-expression***—(Optional) Only log lines that match the regular expression.

**no-remote-trace**—(Optional) Disable tracing and logging operations that track normal operations, error conditions, and packets that are generated by or passed through the Juniper Networks device.

**no-world-readable**—Restrict access to the trace files to the owner.

**Default:** no-world-readable

**size *size***—(Optional) Maximum size of each trace file in bytes, kilobytes (KB), megabytes (MB), or gigabytes (GB). If you do not specify a unit, the default is bytes. If you specify a maximum file size, you also must specify a maximum number of trace files by using the **files** option and a filename by using the **file** option.

**Syntax:** *size* to specify bytes, *sizek* to specify KB, *sizem* to specify MB, or *sizeg* to specify GB.

**Range:** 10,240 through 1,073,741,824 bytes

**Default:** 128 KB

**world-readable**—Enable any user to access the trace files.

|                           |                                                           |
|---------------------------|-----------------------------------------------------------|
| <b>Required Privilege</b> | admin—To view this statement in the configuration.        |
| <b>Level</b>              | admin-control—To add this statement to the configuration. |



## CHAPTER 93

# VXLAN Configuration Statements

- [decapsulate-accept-inner-vlan on page 2525](#)
- [encapsulate-inner-vlan on page 2526](#)
- [multicast-group on page 2526](#)
- [ovsdb-managed on page 2527](#)
- [unreachable-vtep-aging-timer on page 2528](#)
- [vni on page 2528](#)
- [vtep-source-interface on page 2529](#)
- [vxlan on page 2529](#)

### [decapsulate-accept-inner-vlan](#)

---

|                                 |                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>decapsulate-accept-inner-vlan</code>                                                                                                                        |
| <b>Hierarchy Level</b>          | <code>[edit protocols l2-learning]</code>                                                                                                                         |
| <b>Release Information</b>      | Statement modified in Junos OS 14.1X53 for the QFX Series.                                                                                                        |
| <b>Description</b>              | Configure the switch to de-encapsulate and accept original VLAN tags in Virtual Extensible LAN (VXLAN) packets.                                                   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding VXLANs on page 2442</a></li><li>• <a href="#">encapsulate-inner-vlan on page 2526</a></li></ul> |

## encapsulate-inner-vlan

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | encapsulate-inner-vlan                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit <a href="#">vlans</a> <i>VLAN</i> <a href="#">vxlan</a> ]                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1X53-D10.                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Configure the switch to preserve the original VLAN tag (in the inner Ethernet packet) when performing Virtual Extensible LAN (VXLAN) encapsulation.                                                                                                                                                                                                                 |
| <b>Default</b>                  | The original tag is dropped when the packet is encapsulated.                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding VXLANs on page 2442</a></li><li>• <a href="#">Manually Configuring VXLANs on a QFX5100 Switch on page 2492</a></li><li>• <a href="#">Examples: Manually Configuring VXLANs on QFX Series Switches on page 2494</a></li><li>• <a href="#">decapsulate-accept-inner-vlan on page 2525</a></li></ul> |

## multicast-group

---

|                                 |                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | multicast-group                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit <a href="#">vlans</a> <i>VLAN</i> <a href="#">vxlan</a> ]                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1X53-D10.                                                                                                                                                                                                                                          |
| <b>Description</b>              | Assign a multicast group address to a Virtual Extensible LAN (VXLAN). All members of a VXLAN must use the same multicast group address                                                                                                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding VXLANs on page 2442</a></li><li>• <a href="#">Manually Configuring VXLANs on a QFX5100 Switch on page 2492</a></li><li>• <a href="#">Examples: Manually Configuring VXLANs on QFX Series Switches on page 2494</a></li></ul> |

## ovsdb-managed

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | ovsdb-managed;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | <p>[edit bridge-domains <i>bridge-domain-name</i> vxlan],<br/> [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> vxlan],<br/> [edit routing-instances <i>routing-instance-name</i> switch-options],<br/> [edit routing-instances <i>routing-instance-name</i> vlans <i>vlan-name</i> vxlan],<br/> [edit routing-instances <i>routing-instance-name</i> vxlan],<br/> [edit switch-options],<br/> [edit vlans <i>vlan-name</i> vxlan]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 14.1R2.</p> <p>Statement introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.</p> <p>Statement introduced in Junos OS Release 14.2 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | <p>Disable a Juniper Networks device from learning about other Juniper Networks devices that function as hardware virtual tunnel endpoints (VTEPs) in a specified Virtual Extensible LAN (VXLAN) and the MAC addresses learned by the hardware VTEPs. Instead, the Juniper Networks device uses the Open vSwitch Database (OVSDB) management protocol to learn about the hardware VTEPs in the VXLAN and the media access control (MAC) addresses learned by the hardware VTEPs.</p> <p>The specified VXLAN must have a VXLAN network identifier (VNI) configured, using the <b>vni</b> statement in the [edit bridge-domains <i>bridge-domain-name</i> vxlan], [edit routing-instance <i>routing-instance-name</i> vxlan], or [edit vlans <i>vlan-name</i> vxlan] hierarchy.</p> <p>Also, for OVSDB-managed VXLANs, the multicast scheme described in “<a href="#">Understanding How Layer 2 BUM and Layer 3 Routed Multicast Traffic Are Handled with OVSDB</a>” on <a href="#">page 2450</a> is used. Therefore, specifying the <b>multicast-group</b> statement in the [edit bridge-domains <i>bridge-domain-name</i> vxlan], [edit routing-instances <i>routing-instance-name</i> vxlan], or [edit vlans <i>vlan-name</i> vxlan] hierarchy has no effect.</p> |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding Dynamically Configured VXLANs in an OVSDB Environment on page 2462</a></li> <li>• <a href="#">Configuring OVSDB-Managed VXLANs</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

## unreachable-vtep-aging-timer

---

|                                 |                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | unreachable-vtep-aging-timer [300–1800]                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit <a href="#">vlans</a> <i>VLAN</i> <a href="#">vxlan</a> ]                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1X53-D10.                                                                                                                                                                                                                                          |
| <b>Description</b>              | Configure the system to age out the address for the remote VTEP if all the MAC addresses learned from that virtual tunnel endpoint (VTEP) age out. The address for the remote VTEP expires the configured number of seconds after the last learned media access control (MAC) address expires. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding VXLANs on page 2442</a></li><li>• <a href="#">Manually Configuring VXLANs on a QFX5100 Switch on page 2492</a></li><li>• <a href="#">Examples: Manually Configuring VXLANs on QFX Series Switches on page 2494</a></li></ul> |

## vni

---

|                                 |                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | vni [1–16777214]                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit <a href="#">vlans</a> <i>VLAN</i> <a href="#">vxlan</a> ]                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1X53-D10.                                                                                                                                                                                                                                          |
| <b>Description</b>              | Assign a numeric value to identify a Virtual Extensible LAN (VXLAN). All members of a VXLAN must use the same VNI.                                                                                                                                                                             |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding VXLANs on page 2442</a></li><li>• <a href="#">Manually Configuring VXLANs on a QFX5100 Switch on page 2492</a></li><li>• <a href="#">Examples: Manually Configuring VXLANs on QFX Series Switches on page 2494</a></li></ul> |



## vtep-source-interface

|                                 |                                                                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>vtep-source-interface <i>logical-interface</i>;</code>                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit <a href="#">switch-options</a> ,                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1X53-D10.                                                                                                                                                                                                                                              |
| <b>Description</b>              | Configure a source interface for a Virtual Extensible LAN (VXLAN) tunnel. You must provide the name of a logical interface configured on the loopback interface.                                                                                                                                   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding VXLANs on page 2442</a></li> <li>• <a href="#">Manually Configuring VXLANs on a QFX5100 Switch on page 2492</a></li> <li>• <a href="#">Examples: Manually Configuring VXLANs on QFX Series Switches on page 2494</a></li> </ul> |

## vxlan

|                                 |                                                                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>vxlan {   encapsulate-inner-vlan   ingress-node-replication   multicast-group   ovsdb-managed   unreachable-vtep-aging-timer   vni }</pre>                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit <a href="#">vlans</a> ]                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1X53-D10.<br><b>ingress-node-replication</b> option added for QFX Series switches in Junos OS Release 14.1X53-D30.                                                                                                                                     |
| <b>Description</b>              | Configure support for Virtual Extensible LANs (VXLANs) on a Juniper Networks device.                                                                                                                                                                                                               |
| <b>Options</b>                  | The remaining statements are explained separately.                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding VXLANs on page 2442</a></li> <li>• <a href="#">Manually Configuring VXLANs on a QFX5100 Switch on page 2492</a></li> <li>• <a href="#">Examples: Manually Configuring VXLANs on QFX Series Switches on page 2494</a></li> </ul> |



## CHAPTER 94

# OVSDDB Operational Commands

- `clear ovbdb commit failures`
- `show ovbdb commit failures`
- `show ovbdb controller`
- `show ovbdb interface`
- `show ovbdb logical-switch`
- `show ovbdb mac`
- `show ovbdb statistics interface`
- `show ovbdb virtual-tunnel-end-point`
- `show vpls mac-table`

## clear ovssdb commit failures

---

**Syntax**    `clear ovssdb commit failures`  
              `<transaction-id>`

**Release Information**    Command introduced in Junos OS Release 14.1X53-D26 for QFX Series switches.

**Description**    Remove a transaction from a queue maintained by a Juniper Networks switch that supports the Open vSwitch Database (OVSSDB) management protocol and Virtual Extensible LANs (VXLANs). The transaction includes OVSSDB-managed VXLANs and associated logical interfaces that the Juniper Networks switch dynamically configured and tried to commit but was unable to because of an issue with one or more of the configurations. In addition to removing the transaction, entering the **clear ovssdb commit failures** command causes the Juniper Networks switch to automatically retry committing all configurations in the transaction.

If there is an issue with one or more of the configurations in a transaction, this causes all configurations in the transaction, even the ones that are correctly configured, to remain uncommitted and in the queue until you troubleshoot and resolve the configuration issue(s).

You can display an erroneous transaction by entering the **show ovssdb commit failures** command. In the output that appears, you must determine which configuration(s) are erroneous and therefore prevent the Juniper Networks switch from committing the configurations in the transaction.

Issues that can cause commitment errors include but are not limited to the detection of the same VXLAN name or VXLAN network identifier (VNI) in a dynamically configured VXLAN and in a VXLAN that was previously configured using the Junos OS CLI.

To monitor for issues with dynamically configured OVSSDB-managed VXLANs and their associated interfaces, we recommend checking for system log messages and traceoptions files for OVSSDB.

After resolving the error(s), enter the **clear ovssdb commit failures** command to remove the transaction from the queue and retry committing all configurations in the transaction.



**NOTE:** While an erroneous transaction exists in the queue, the Juniper Networks switch cannot commit the dynamic configurations of additional VXLANs and their associated logical interfaces. The commitment of these VXLANs and logical interfaces remain in a pending state until all VXLAN and logical interface configurations in the erroneous transaction are resolved and successfully committed.

---

**Options**    **none**—Remove the transaction that currently appears in the **show ovssdb commit failures** command output, and retry committing all configurations in the transaction.

*transaction-id*—Remove the transaction with the specified numerical ID, and retry committing the configurations in the transaction.

**Required Privilege Level**    clear

**Related Documentation**    • [show ovldb commit failures on page 2534](#)

**List of Sample Output**    [clear ovldb commit failures on page 2533](#)  
                                  [clear ovldb commit failures \(Specific Transaction\) on page 2533](#)

## Sample Output

[clear ovldb commit failures](#)

```
user@host> clear ovldb commit failures
```

[clear ovldb commit failures \(Specific Transaction\)](#)

```
user@host> clear ovldb commit failures 1
```

## show ovssdb commit failures

---

**Syntax**    `show ovssdb commit failures`  
              `<transaction-id>`

**Release Information**    Command introduced in Junos OS Release 14.1X53-D26 for QFX Series switches.

**Description**    Display configurations of Open vSwitch Database (OVSSDB)-managed Virtual Extensible LANs (VXLANS) and associated logical interfaces that the Juniper Networks switch dynamically configured but was unable to commit.

For each OVSSDB-managed VXLAN and associated logical interface that you plan to implement in a Junos OS environment, you must configure equivalent entities in NSX Manager or in the NSX API for an NSX environment, or in the Contrail Web user interface for a Contrail environment. The software-defined networking (SDN) controller pushes these configurations to the connected Juniper Networks switch by way of the OVSSDB schema for physical devices. After the Juniper Networks switch receives these configurations, it dynamically configures a Junos OS-equivalent VXLAN and associated logical interface, and attempts to commit the configurations.

During the commitment of the dynamic configurations, If there is an issue with one or more of the configurations, all configurations in the transaction, even the ones that are correctly configured, remain uncommitted and are saved in a queue. All configurations in the transaction remain uncommitted and in the queue until you troubleshoot and resolve the configuration issues. After you resolve the configuration issues, you must use the [clear ovssdb commit failures](#) command to remove the transaction from the queue and retry committing the configurations.



**NOTE:** While an erroneous transaction exists in the queue, the Juniper Networks switch cannot commit the dynamic configurations of additional VXLANS and their associated logical interfaces. The commitment of these VXLANS and logical interfaces remain in a pending state until all VXLAN and logical interface configurations in the erroneous transaction are resolved and successfully committed.

---

Issues that can cause commitment errors include but are not limited to the detection of the same VXLAN name or VXLAN network identifier (VNI) in a dynamically configured VXLAN and in a VXLAN that was previously configured using the Junos OS CLI.

To monitor for issues with dynamically configured OVSSDB-managed VXLANS and their associated interfaces, we recommend checking for system log messages and traceoptions files for OVSSDB.

**Options**    **none**—Display information about an erroneous transaction.

**transaction-id**—Display information about the transaction with the specified numerical ID.

**Required Privilege Level** admin

**Related Documentation**

- [Understanding Dynamically Configured VXLANs in an OVSDb Environment on page 2462](#)
- [traceoptions \(OVSDb\) on page 2522](#)

**List of Sample Output** [show ovssdb commit failures on page 2535](#)  
[show ovssdb commit failure \(Specific Transaction\) on page 2535](#)

**Output Fields** [Table 222](#) lists the output fields for the **show ovssdb commit failures** command. Output fields are listed in the approximate order in which they appear.

**Table 222: show ovssdb commit failures Output Fields**

| Field Names    | Field Descriptions                                                                                                       |
|----------------|--------------------------------------------------------------------------------------------------------------------------|
| Txn ID         | ID assigned to a transaction by the Juniper Networks switch.                                                             |
| Logical-switch | Name of the VXLAN that the Juniper Networks switch dynamically configured but was unable to commit the configuration of. |
| Port           | Name of an OVSDb-managed physical interface that is associated with the VXLAN.                                           |
| VLAN ID        | ID that is assigned to the VXLAN.                                                                                        |

## Sample Output

### show ovssdb commit failures

```

user@host> show ovssdb commit failures
Txn ID Logical-switch Port VLAN ID
1 28805c1d-0122-495d-85df-19abd647d772 xe-0/0/5:0 1016
1
1 9acc24b3-7b0a-4c2e-b572-3370c3e1acff xe-0/0/5:0 1017
1
...
```

### show ovssdb commit failure (Specific Transaction)

```

user@host> show ovssdb commit failures 1
Txn ID Logical-switch Port VLAN ID
1 28805c1d-0122-495d-85df-19abd647d772 xe-0/0/5:0 1016
1
1 9acc24b3-7b0a-4c2e-b572-3370c3e1acff xe-0/0/5:0 1017
1
...
```

## show ovssdb controller

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show ovssdb controller</code><br><code>&lt;address ip-address&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Command introduced in Junos OS Release 14.1R2.<br>Command introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.<br>Command introduced in Junos OS Release 14.2 for EX Series switches.                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Display information and connection status for software-defined networking (SDN) controllers to which the Juniper Networks device is connected.                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <p><b>none</b>—Display information about all SDN controllers to which the Juniper Networks device is connected.</p> <p><b>address ip-address</b>—Display information about the SDN controller at the specified IP address.</p>                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | admin                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Setting Up the OVSSDB Protocol on Juniper Networks Devices that Support Manual Configuration of VXLANs</a></li> <li>• <a href="#">Setting Up the OVSSDB Protocol on Juniper Networks Devices that Support the Dynamic Configuration of VXLANs on page 2461</a></li> <li>• <a href="#">Understanding How to Set Up OVSSDB Connections Between Juniper Networks Devices and SDN Controllers on page 2458</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show ovssdb controller on page 2537</a><br><a href="#">show ovssdb controller address on page 2537</a>                                                                                                                                                                                                                                                                                                                                                         |
| <b>Output Fields</b>            | Table 223 lists the output fields for the <b>show ovssdb controller</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                   |

**Table 223: show ovssdb controller Output Fields**

| Field Name                       | Field Description                                                                   |
|----------------------------------|-------------------------------------------------------------------------------------|
| Controller IP address            | IP address of the SDN controller to which the Juniper Networks device is connected. |
| Controller protocol              | Protocol used by the Juniper Networks device to initiate the connection.            |
| Controller port                  | Port to which the Juniper Networks device is connected.                             |
| Controller connection            | State of the connection with the SDN controller.                                    |
| Controller seconds-since-connect | Number of seconds since the connection with the SDN controller was established.     |



Table 223: show ovsdb controller Output Fields (*continued*)

| Field Name                          | Field Description                                                           |
|-------------------------------------|-----------------------------------------------------------------------------|
| Controller seconds-since-disconnect | Number of seconds since the connection with the SDN controller was dropped. |
| Controller connection status        | Status of the connection with the SDN controller.                           |

## Sample Output

### show ovsdb controller

```

user@host> show ovsdb controller
VTEP controller information:
Controller IP address: 10.168.66.189
Controller protocol: ssl
Controller port: 6632
Controller connection: up
Controller seconds-since-connect: 56290
Controller seconds-since-disconnect: 0
Controller connection status: active

Controller IP address: 10.168.181.54
Controller protocol: ssl
Controller port: 6632
Controller connection: up
Controller seconds-since-connect: 56292
Controller seconds-since-disconnect: 0
Controller connection status: active

Controller IP address: 10.168.182.45
Controller protocol: ssl
Controller port: 6632
Controller connection: up
Controller seconds-since-connect: 56292
Controller seconds-since-disconnect: 0
Controller connection status: active

```

### show ovsdb controller address

```

user@host> show ovsdb controller address 10.168.182.45
VTEP controller information:
Controller IP address: 192.168.182.45
Controller protocol: ssl
Controller port: 6632
Controller connection: up
Controller seconds-since-connect: 56347
Controller seconds-since-disconnect: 0
Controller connection status: active

```

## show ovssdb interface

|                                 |                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show ovssdb interface</code><br><code>&lt;interface-name&gt;</code>                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Command introduced in Junos OS Release 14.1R2.<br>Command introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.<br>Command introduced in Junos OS Release 14.2 for EX Series switches.                                                                                                                       |
| <b>Description</b>              | Display information about Open vSwitch Database (OVSSDB)-managed interfaces configured by using the <code>interfaces interface-name</code> statement in the <code>[edit protocols ovssdb]</code> hierarchy.                                                                                                                |
| <b>Options</b>                  | <b>none</b> —Display information about all OVSSDB-managed interfaces.<br><br><b>interface-name</b> —Display information about the specified OVSSDB-managed interface.                                                                                                                                                      |
| <b>Required Privilege Level</b> | admin                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">Setting Up the OVSSDB Protocol on Juniper Networks Devices that Support Manual Configuration of VXLANs</a></li> <li><a href="#">Setting Up the OVSSDB Protocol on Juniper Networks Devices that Support the Dynamic Configuration of VXLANs on page 2461</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show ovssdb interface on page 2538</a><br><a href="#">show ovssdb (Specific Interface) on page 2539</a>                                                                                                                                                                                                        |
| <b>Output Fields</b>            | <a href="#">Table 224</a> lists the output fields for the <code>show ovssdb interface</code> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                              |

**Table 224: show ovssdb interface Output Fields**

| Field Name            | Field Description                                                                                                                                                     |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Interface             | Name of interface.                                                                                                                                                    |
| VLAN ID               | ID of Virtual Extensible LAN (VXLAN) with which the interface is associated.<br><br><b>NOTE:</b> This field is not supported by MX Series routers or EX9200 switches. |
| Bridge domain or VLAN | Bridge domain or VLAN under which the VXLAN is created.<br><br><b>NOTE:</b> This field is not supported by MX Series routers or EX9200 switches.                      |

## Sample Output

### show ovssdb interface

```
user@host> show ovssdb interface
```

| Interface   | VLAN ID | Bridge-domain |
|-------------|---------|---------------|
| ge-7/0/9.0  |         |               |
| ge-7/0/9.1  |         |               |
| irb.11      |         |               |
| irb.12      |         |               |
| irb.2       |         |               |
| irb.3       |         |               |
| xe-10/3/0.0 |         |               |
| xe-10/3/0.1 |         |               |

#### show ovbdb (Specific Interface)

```
user@host> show ovbdb interface ge-7/0/9.0
```

| Interface  | VLAN ID | Bridge-domain |
|------------|---------|---------------|
| ge-7/0/9.0 |         |               |

## show ovssdb logical-switch

---

**Syntax**    show ovssdb logical-switch  
              <logical-switch-name>

**Release Information**    Command introduced in Junos OS Release 14.1R2.  
                              Command introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.  
                              Command introduced in Junos OS Release 14.2 for EX Series switches.

**Description**



**NOTE:** In the Open vSwitch Database (OVSSDB) schema for physical devices, the logical switch table stores information about the Layer 2 broadcast domain that you configured in a VMware NSX or Contrail environment. In the NSX environment, the Layer 2 broadcast domain is known as a *logical switch*, while in the Contrail environment, the domain is known as a *virtual network*.

In the context of the **show ovssdb logical-switch** command, the term *logical switch* refers to the logical switch or virtual network that was configured in the NSX or Contrail environments, respectively, and the corresponding configuration that was pushed to the OVSSDB schema.

Display information about logical switches and the corresponding Virtual Extensible LANs (VXLANs), which were configured on the Juniper Networks device.

In the command output, each logical switch is identified by a universally unique identifier (UUID), which in the context of this command, is also known as a logical switch name.

The **show ovssdb logical-switch** command displays the state of the logical switch (**Flags**), which can be one of the following:

**Created by Controller**—A logical switch is configured. However, a corresponding VXLAN is not yet configured. In this state, the logical switch and corresponding VXLAN are not yet operational.

**Created by L2ALD**—A VXLAN is configured. However, a corresponding logical switch is not yet configured. In this state, the logical switch and corresponding VXLAN are not yet operational.

**Created by both**—A logical switch and a corresponding VXLAN are configured. In this state, the logical switch and corresponding VXLAN are operational.

**Tunnel key mismatch**—The VNIs specified in the logical switch and corresponding VXLAN configurations do not match. In this state, the logical switch and corresponding VXLAN are not yet operational.

**Options**    **none**—Display information about all logical switches that are present in the OVSSDB schema for physical devices.

**logical-switch-name**—Display information about the specified logical switch.

**Required Privilege Level** admin

**Related Documentation**

- [OVSDb Schema for Physical Devices on page 2447](#)
- [Troubleshooting a Nonoperational Logical Switch and Corresponding Junos OS OVSDb-Managed VXLAN on page 2507](#)

**List of Sample Output** [show ovssdb logical-switch on page 2541](#)  
[show ovssdb logical-switch \(Specific Logical Switch\) on page 2542](#)

**Output Fields** [Table 225](#) lists the output fields for the **show ovssdb logical-switch** command. Output fields are listed in the approximate order in which they appear.

**Table 225: show ovssdb logical-switch Output Fields**

| Field Name          | Field Description                                                                                                                                                                          |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Logical Switch Name | UUID that is automatically generated and assigned to the logical switch. When you configure the corresponding VXLAN in the Junos OS CLI, you must specify the same UUID as the VXLAN name. |
| Flags               | State of the logical switch. For possible states, see the Description section of this topic.                                                                                               |
| VNI                 | VNI that is configured for the logical switch and corresponding VXLAN.                                                                                                                     |
| Num of Remote MAC   | The total number of remote media access control (MAC) addresses associated with the logical switch. These addresses are learned by software and hardware virtual tunnel endpoints (VTEPs). |
| Num of Local MAC    | The total number of local MAC addresses associated with the logical switch. <i>Local MAC addresses</i> are addresses learned on the local physical ports.                                  |

## Sample Output

### show ovssdb logical-switch

```
user@host> show ovssdb logical-switch
Logical switch information:
Logical Switch Name: 24a76aff-7e61-4520-a78d-3eca26ad7510
Flags: Created by both
VNI: 3
Num of Remote MAC: 13
Num of Local MAC: 12
Logical Switch Name: 9b4f880e-dac8-4612-a832-97ad9dec270f
Flags: Created by Controller
VNI: 50
Num of Remote MAC: 0
Num of Local MAC: 0
Logical Switch Name: bc0da2da-6c16-44bf-b655-442484294ded
Flags: Created by Controller
VNI: 51
Num of Remote MAC: 0
Num of Local MAC: 0
```

### show ovssdb logical-switch (Specific Logical Switch)

```
user@host> show ovssdb logical-switch 24a76aff-7e61-4520-a78d-3eca26ad7510
Logical switch information:
Logical Switch Name: 24a76aff-7e61-4520-a78d-3eca26ad7510
Flags: Created by both
VNI: 3
Num of Remote MAC: 13
Num of Local MAC: 12
```

## show ovssdb mac

**Syntax** show ovssdb mac  
 <address *mac-address*>  
 <local>  
 <logical-switch *logical-switch-uuid*>  
 <multicast>  
 <remote>  
 <unicast>

**Release Information** Command introduced in Junos OS Release 14.1R2.  
 Command introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.  
 Command introduced in Junos OS Release 14.2 for EX Series switches.

**Description** Display media access control (MAC) addresses, as well as information about the media access control (MAC) addresses, learned by a Juniper Networks device that functions as a hardware virtual tunnel endpoint (VTEP). Using the Open vSwitch Database (OVSDb) management protocol, this hardware VTEP can learn about MAC addresses directly or from other software or hardware VTEPs. The MAC addresses learned directly by the hardware VTEP are known as *local addresses*, while the addresses learned from other software or hardware VTEPs are known as *remote addresses*.

**Options** Use one or more of the following options to display a more specific list of MAC addresses and information about the MAC addresses. For example, to display a list of local unicast MAC addresses, you can issue the **show ovssdb mac local unicast** command.

**none**—Display all MAC addresses, which includes all local, remote, unicast, and multicast addresses associated with all logical switches.

**address *mac-address***—Display the specified MAC address.

**count**—(All Juniper Networks devices that support OVSDb except EX9200 switches)  
 Display the number of MAC addresses learned by the Juniper Networks device. Using this option alone, the number includes all local, remote, unicast, and multicast MAC addresses associated with all logical switches in the logical switch table of the OVSDb schema for physical devices. You can use this option with one or more of the other options to display a more specific count of MAC addresses. For example, to display the number of local and remote unicast MAC addresses, you can issue the **show ovssdb mac count local remote unicast** command.

**local**—Display all local MAC addresses.

**logical-switch *logical-switch-uuid***—Display all MAC addresses associated with the specified logical switch in the logical switch table of the OVSDb schema for physical devices.

**multicast**—Display all multicast MAC addresses.

**remote**—Display all remote MAC addresses.

**unicast**—Display all unicast MAC addresses.

**Required Privilege Level** admin

**Related Documentation** • [OVSDb Schema for Physical Devices on page 2447](#)

**List of Sample Output** [show ovssdb mac on page 2544](#)  
[show ovssdb mac address on page 2545](#)  
[show ovssdb mac logical-switch on page 2545](#)  
[show ovssdb mac local unicast on page 2546](#)  
[show ovssdb mac \(Count of All Local, Remote, Unicast, and Multicast MAC Addresses for All Logical Switches\) on page 2546](#)

**Output Fields** Table 226 lists the output fields for the **show ovssdb mac** command. Output fields are listed in the approximate order in which they appear.

**Table 226: show ovssdb mac Output Fields**

| Field Name          | Field Description                                                                                                                                                                                       |
|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Logical Switch Name | Universally unique identifier (UUID) of the logical switch.                                                                                                                                             |
| MAC Address         | MAC addresses of virtual machines (VMs).                                                                                                                                                                |
| IP Address          | IP address of VMs.<br><br><i>NOTE:</i> If the IP addresses of VMs are not published by the SDN controller, this field displays 0.0.0.0.                                                                 |
| Encapsulation       | Encapsulation type.                                                                                                                                                                                     |
| VTEP Address        | IP address of the hardware or software VTEP from which the MAC address was learned. Further, this VTEP can forward VM traffic to the associated host.                                                   |
| MAC Count           | <i>NOTE:</i> This field is supported by all Juniper Networks devices that support OVSDb except EX9200 switches.<br><br>Number of all or specified MAC addresses learned by the Juniper Networks device. |

## Sample Output

### show ovssdb mac

```

user@host> show ovssdb mac
Logical Switch Name: 24a76aff-7e61-4520-a78d-3eca26ad7510
 Mac IP Encapsulation Vtep
 Address Address
02:00:00:00:03:01 0.0.0.0 Vxlan over Ipv4 10.255.18.22
02:00:00:00:03:02 0.0.0.0 Vxlan over Ipv4 10.255.18.22
02:00:00:00:03:03 0.0.0.0 Vxlan over Ipv4 10.255.18.22
02:00:00:00:03:04 0.0.0.0 Vxlan over Ipv4 10.255.18.22
02:00:00:00:03:05 0.0.0.0 Vxlan over Ipv4 10.255.18.22
04:00:00:00:03:05 0.0.0.0 Vxlan over Ipv4 10.255.18.22
06:00:00:00:03:01 0.0.0.0 Vxlan over Ipv4 10.255.18.22
06:00:00:00:03:02 0.0.0.0 Vxlan over Ipv4 10.255.18.22

```



|                   |         |                 |              |
|-------------------|---------|-----------------|--------------|
| 06:00:00:00:03:03 | 0.0.0.0 | Vxlan over Ipv4 | 10.255.18.22 |
| 06:00:00:00:03:04 | 0.0.0.0 | Vxlan over Ipv4 | 10.255.18.22 |
| 06:00:00:00:03:05 | 0.0.0.0 | Vxlan over Ipv4 | 10.255.18.22 |
| 40:b4:f0:06:6f:f0 | 0.0.0.0 | Vxlan over Ipv4 | 10.255.18.22 |
| ff:ff:ff:ff:ff:ff | 0.0.0.0 | Vxlan over Ipv4 | 10.100.100.1 |

Logical Switch Name: bf6d4fd4-f5f6-430c-8c37-4033ef1c55ab

| Mac Address       | IP Address | Encapsulation   | Vtep Address |
|-------------------|------------|-----------------|--------------|
| 02:00:00:00:11:01 | 0.0.0.0    | Vxlan over Ipv4 | 10.1.1.29    |
| 02:00:00:00:11:02 | 0.0.0.0    | Vxlan over Ipv4 | 10.1.1.29    |
| 02:00:00:00:11:03 | 0.0.0.0    | Vxlan over Ipv4 | 10.1.1.29    |
| 02:00:00:00:11:04 | 0.0.0.0    | Vxlan over Ipv4 | 10.1.1.29    |
| 02:00:00:00:11:05 | 0.0.0.0    | Vxlan over Ipv4 | 10.1.1.29    |
| 04:00:00:00:11:05 | 0.0.0.0    | Vxlan over Ipv4 | 10.1.1.29    |
| 06:00:00:00:11:01 | 0.0.0.0    | Vxlan over Ipv4 | 10.1.1.29    |
| 06:00:00:00:11:02 | 0.0.0.0    | Vxlan over Ipv4 | 10.1.1.29    |
| 06:00:00:00:11:03 | 0.0.0.0    | Vxlan over Ipv4 | 10.1.1.29    |
| 06:00:00:00:11:04 | 0.0.0.0    | Vxlan over Ipv4 | 10.1.1.29    |
| 06:00:00:00:11:05 | 0.0.0.0    | Vxlan over Ipv4 | 10.1.1.29    |
| 40:b4:f0:06:6f:f0 | 0.0.0.0    | Vxlan over Ipv4 | 10.1.1.29    |
| 00:23:9c:5e:a7:f0 | 0.0.0.0    | Vxlan over Ipv4 | 10.255.18.22 |
| 08:00:00:00:11:01 | 0.0.0.0    | Vxlan over Ipv4 | 10.255.18.22 |
| 08:00:00:00:11:02 | 0.0.0.0    | Vxlan over Ipv4 | 10.255.18.22 |
| 08:00:00:00:11:03 | 0.0.0.0    | Vxlan over Ipv4 | 10.255.18.22 |
| 08:00:00:00:11:04 | 0.0.0.0    | Vxlan over Ipv4 | 10.255.18.22 |
| 08:00:00:00:11:05 | 0.0.0.0    | Vxlan over Ipv4 | 10.255.18.22 |
| ff:ff:ff:ff:ff:ff | 0.0.0.0    | Vxlan over Ipv4 | 10.110.110.1 |

...

#### show ovssdb mac address

user@host> show ovssdb mac address 02:00:00:00:03:01

| Mac Address       | IP Address | Encapsulation   | Vtep Address |
|-------------------|------------|-----------------|--------------|
| 02:00:00:00:03:01 | 0.0.0.0    | Vxlan over Ipv4 | 10.255.18.22 |

#### show ovssdb mac logical-switch

user@host> show ovssdb mac logical-switch bf6d4fd4-f5f6-430c-8c37-4033ef1c55ab

Logical Switch Name: bf6d4fd4-f5f6-430c-8c37-4033ef1c55ab

| Mac Address       | IP Address | Encapsulation   | Vtep Address |
|-------------------|------------|-----------------|--------------|
| 02:00:00:00:11:01 | 0.0.0.0    | Vxlan over Ipv4 | 10.1.1.29    |
| 02:00:00:00:11:02 | 0.0.0.0    | Vxlan over Ipv4 | 10.1.1.29    |
| 02:00:00:00:11:03 | 0.0.0.0    | Vxlan over Ipv4 | 10.1.1.29    |
| 02:00:00:00:11:04 | 0.0.0.0    | Vxlan over Ipv4 | 10.1.1.29    |
| 02:00:00:00:11:05 | 0.0.0.0    | Vxlan over Ipv4 | 10.1.1.29    |
| 04:00:00:00:11:05 | 0.0.0.0    | Vxlan over Ipv4 | 10.1.1.29    |
| 06:00:00:00:11:01 | 0.0.0.0    | Vxlan over Ipv4 | 10.1.1.29    |
| 06:00:00:00:11:02 | 0.0.0.0    | Vxlan over Ipv4 | 10.1.1.29    |
| 06:00:00:00:11:03 | 0.0.0.0    | Vxlan over Ipv4 | 10.1.1.29    |
| 06:00:00:00:11:04 | 0.0.0.0    | Vxlan over Ipv4 | 10.1.1.29    |
| 06:00:00:00:11:05 | 0.0.0.0    | Vxlan over Ipv4 | 10.1.1.29    |
| 40:b4:f0:06:6f:f0 | 0.0.0.0    | Vxlan over Ipv4 | 10.1.1.29    |
| 00:23:9c:5e:a7:f0 | 0.0.0.0    | Vxlan over Ipv4 | 10.255.18.22 |
| 08:00:00:00:11:01 | 0.0.0.0    | Vxlan over Ipv4 | 10.255.18.22 |
| 08:00:00:00:11:02 | 0.0.0.0    | Vxlan over Ipv4 | 10.255.18.22 |
| 08:00:00:00:11:03 | 0.0.0.0    | Vxlan over Ipv4 | 10.255.18.22 |
| 08:00:00:00:11:04 | 0.0.0.0    | Vxlan over Ipv4 | 10.255.18.22 |

|                   |         |                 |              |
|-------------------|---------|-----------------|--------------|
| 08:00:00:00:11:05 | 0.0.0.0 | Vxlan over Ipv4 | 10.255.18.22 |
| ff:ff:ff:ff:ff:ff | 0.0.0.0 | Vxlan over Ipv4 | 10.110.110.1 |

#### show ovsdb mac local unicast

```
user@host> show ovsdb mac local unicast
Logical Switch Name: 24a76aff-7e61-4520-a78d-3eca26ad7510
Mac IP Encapsulation Vtep
Address Address
02:00:00:00:03:01 0.0.0.0 Vxlan over Ipv4 10.255.181.72
02:00:00:00:03:02 0.0.0.0 Vxlan over Ipv4 10.255.181.72
02:00:00:00:03:03 0.0.0.0 Vxlan over Ipv4 10.255.181.72
02:00:00:00:03:04 0.0.0.0 Vxlan over Ipv4 10.255.181.72
02:00:00:00:03:05 0.0.0.0 Vxlan over Ipv4 10.255.181.72
04:00:00:00:03:05 0.0.0.0 Vxlan over Ipv4 10.255.181.72
06:00:00:00:03:01 0.0.0.0 Vxlan over Ipv4 10.255.181.72
06:00:00:00:03:02 0.0.0.0 Vxlan over Ipv4 10.255.181.72
06:00:00:00:03:03 0.0.0.0 Vxlan over Ipv4 10.255.181.72
06:00:00:00:03:04 0.0.0.0 Vxlan over Ipv4 10.255.181.72
06:00:00:00:03:05 0.0.0.0 Vxlan over Ipv4 10.255.181.72
40:b4:f0:06:6f:f0 0.0.0.0 Vxlan over Ipv4 10.255.181.72
...
```

#### show ovsdb mac (Count of All Local, Remote, Unicast, and Multicast MAC Addresses for All Logical Switches)

```
user@host> show ovsdb mac count
MAC count: 6877
```

## show ovssdb statistics interface

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show ovssdb statistics interface</code><br><code>&lt;interface-name&gt;</code>                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Command introduced in Junos OS Release 14.1R2.<br>Command introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.<br>Command introduced in Junos OS Release 14.2 for EX Series switches.                                                                                                                                                                                                      |
| <b>Description</b>              | Display statistics for Open vSwitch Database (OVSDb)-managed interfaces configured by using the <b>interfaces interface-name</b> statement in the <b>[edit protocols ovssdb]</b> hierarchy.<br><br>When an interface is configured as OVSDb-managed, the collection of statistics for that interface begins, and the statistics displayed at any given time reflects the data collected up to that point. |
| <b>Options</b>                  | <b>none</b> —Display statistics for all configured OVSDb-managed interfaces.<br><br><b>interface-name</b> —Display statistics for the specified interface.                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | admin                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">interfaces on page 2516</a></li> </ul>                                                                                                                                                                                                                                                                                                                 |
| <b>List of Sample Output</b>    | <a href="#">show ovssdb statistics interface on page 2547</a><br><a href="#">show ovssdb statistics interface (Specific Interface) on page 2548</a>                                                                                                                                                                                                                                                       |
| <b>Output Fields</b>            | <a href="#">Table 227</a> lists the output fields for the <b>show ovssdb statistics interface</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                        |

**Table 227: show ovssdb statistics interface Output Fields**

| Field Name      | Field Description                            |
|-----------------|----------------------------------------------|
| Num of rx pkts  | Number of packets received by the interface. |
| Num of tx pkts  | Number of packets sent by the interface.     |
| Num of rx bytes | Number of bytes received by the interface.   |
| Num of tx bytes | Number of bytes sent by the interface.       |

## Sample Output

### show ovssdb statistics interface

```

user@host> show ovssdb statistics interface
Interface Name: ge-7/0/9.0
Num of rx pkts: 945 Num of tx pkts: 113280890
Num of rx bytes: 56700 Num of tx bytes: 57531319540

```

|                             |                              |
|-----------------------------|------------------------------|
| Interface Name: ge-7/0/10.0 |                              |
| Num of rx pkts: 459         | Num of tx pkts: 473840856    |
| Num of rx bytes: 84747      | Num of tx bytes: 45830738532 |
| Interface Name: ge-7/0/11.0 |                              |
| Num of rx pkts: 305         | Num of tx pkts: 367483456    |
| Num of rx bytes: 98974      | Num of tx bytes: 33495468092 |

#### show ovsdb statistics interface (Specific Interface)

```
user@host> show ovsdb statistics interface ge-7/0/9.0
```

|                            |                              |
|----------------------------|------------------------------|
| Interface Name: ge-7/0/9.0 |                              |
| Num of rx pkts: 945        | Num of tx pkts: 113280890    |
| Num of rx bytes: 56700     | Num of tx bytes: 57531319540 |

## show ovssdb virtual-tunnel-end-point

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show ovssdb virtual-tunnel-end-point<br>address <ip-address><br>encapsulation <encapsulation-type>                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 14.1R2.<br>Command introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.<br>Command introduced in Junos OS Release 14.2 for EX Series switches.                                                                                                                                                                                                                 |
| <b>Description</b>              | Display information about the following entities that the Juniper Networks device has learned: <ul style="list-style-type: none"> <li>• Other Juniper Networks devices that function as hardware virtual tunnel endpoints (VTEPs)</li> <li>• Software VTEPs</li> <li>• Service nodes</li> <li>• Top-of-rack service nodes (TSNs)</li> </ul>                                                                          |
| <b>Options</b>                  | <b>none</b> —Display information about all VTEPs, service nodes, and TSNs that the Juniper Networks device has learned.<br><br><b>address ip-address</b> —Display information about the entity with the specified IP address.<br><br><b>encapsulation encapsulation-type</b> —Display information about all entities with the specified encapsulation type.                                                          |
| <b>Required Privilege Level</b> | admin                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>List of Sample Output</b>    | <a href="#">show ovssdb virtual-tunnel-end-point on page 2550</a><br><a href="#">show ovssdb virtual-tunnel-end-point address (Specific Address) on page 2550</a><br><a href="#">show ovssdb virtual-tunnel-end-point encapsulation (Specific Encapsulation) on page 2550</a><br><a href="#">show ovssdb virtual-tunnel-end-point address (Specific Address) encapsulation (Specific Encapsulation) on page 2550</a> |
| <b>Output Fields</b>            | Table 228 lists the output fields for the <b>show ovssdb virtual-tunnel-end-point</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                               |

Table 228: show ovssdb virtual-tunnel-end-point Output Fields

| Field Name    | Field Description                                                     |
|---------------|-----------------------------------------------------------------------|
| Encapsulation | Encapsulation type of entity.                                         |
| IP Address    | IP address of entity.                                                 |
| Num of MACs   | Number of media access control (MAC) addresses learned by the entity. |

## Sample Output

### show ovssdb virtual-tunnel-end-point

```
user@host> show ovssdb virtual-tunnel-end-point
Encapsulation Ip Address Num of MAC's
VXLAN over IPv4 10.255.181.43 24
VXLAN over IPv4 10.255.181.50 12
VXLAN over IPv4 10.255.181.72 24
```

### show ovssdb virtual-tunnel-end-point address (Specific Address)

```
user@host> show ovssdb virtual-tunnel-end-point address 10.255.181.43
Encapsulation Ip Address Num of MAC's
VXLAN over IPv4 10.255.181.43 24
```

### show ovssdb virtual-tunnel-end-point encapsulation (Specific Encapsulation)

```
user@host> show ovssdb virtual-tunnel-end-point encapsulation vxlan-over-ipv4
Encapsulation Ip Address Num of MAC's
VXLAN over IPv4 10.255.181.43 24
VXLAN over IPv4 10.255.181.50 12
VXLAN over IPv4 10.255.181.72 24
```

### show ovssdb virtual-tunnel-end-point address (Specific Address) encapsulation (Specific Encapsulation)

```
user@host> show ovssdb virtual-tunnel-end-point address 10.255.181.43 encapsulation
vxlan-over-ipv4
Encapsulation Ip Address Num of MAC's
VXLAN over IPv4 10.255.181.43 24
```

## show vpls mac-table

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show vpls mac-table &lt;brief   detail   extensive   summary&gt; &lt;bridge-domain <i>bridge-domain-name</i>&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;mac-address&gt; &lt;vlan-id <i>vlan-id-number</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 8.5.</p> <p>Command introduced in Junos OS Release 15.1</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Display learned virtual private LAN service (VPLS) media access control (MAC) address information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <p><b>none</b>—Display all learned VPLS MAC address information.</p> <p><b>brief   detail   extensive   summary</b>—(Optional) Display the specified level of output.</p> <p><b>bridge-domain <i>bridge-domain-name</i></b>—(Optional) Display learned VPLS MAC addresses for the specified bridge domain.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display learned VPLS MAC addresses for the specified instance.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display learned VPLS MAC addresses for the specified instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Display learned VPLS MAC addresses for all logical systems or for the specified logical system.</p> <p><b>mac-address</b>—(Optional) Display the specified learned VPLS MAC address information..</p> <p><b>vlan-id <i>vlan-id-number</i></b>—(Optional) Display learned VPLS MAC addresses for the specified VLAN.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>List of Sample Output</b>    | <p><a href="#">show vpls mac-table on page 2552</a></p> <p><a href="#">show vpls mac-table (with Layer 2 Services over GRE Interfaces) on page 2553</a></p> <p><a href="#">show vpls mac-table (with VXLAN enabled) on page 2553</a></p> <p><a href="#">show vpls mac-table count on page 2553</a></p> <p><a href="#">show vpls mac-table detail on page 2554</a></p> <p><a href="#">show vpls mac-table extensive on page 2554</a></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Output Fields</b>            | Table 229 describes the output fields for the <b>show vpls mac-table</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

Table 229: show vpls mac-table Output fields

| Field Name              | Field Description                                                                                                                                                                                                                                                                                  |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Routing instance        | Name of the routing instance.                                                                                                                                                                                                                                                                      |
| Bridging domain         | Name of the bridging domain.                                                                                                                                                                                                                                                                       |
| MAC address             | MAC address or addresses learned on a logical interface.                                                                                                                                                                                                                                           |
| MAC flags               | Status of MAC address learning properties for each interface: <ul style="list-style-type: none"> <li>• <b>S</b>—Static MAC address configured.</li> <li>• <b>D</b>—Dynamic MAC address learned.</li> <li>• <b>SE</b>—MAC accounting is enabled.</li> <li>• <b>NM</b>—Nonconfigured MAC.</li> </ul> |
| Logical interface       | Name of the logical interface.                                                                                                                                                                                                                                                                     |
| MAC count               | Number of MAC addresses learned on a specific routing instance or interface.                                                                                                                                                                                                                       |
| Learning interface      | Logical interface or logical Label Switched Interface (LSI) the address is learned on.                                                                                                                                                                                                             |
| Base learning interface | Base learning interface of the MAC address. This field is introduced in Junos OS Release 14.2.                                                                                                                                                                                                     |
| Learn VLAN ID/VLAN      | VLAN ID of the routing instance or bridge domain in which the MAC address was learned.                                                                                                                                                                                                             |
| VXLAN ID/VXLAN          | VXLAN Network Identifier (VNI)                                                                                                                                                                                                                                                                     |
| Layer 2 flags           | Debugging flags signifying that the MAC address is present in various lists.                                                                                                                                                                                                                       |
| Epoch                   | Spanning Tree Protocol epoch number identifying when the MAC address was learned. Used for debugging.                                                                                                                                                                                              |
| Sequence number         | Sequence number assigned to this MAC address. Used for debugging.                                                                                                                                                                                                                                  |
| Learning mask           | Mask of Packet Forwarding Engines where this MAC address was learned. Used for debugging.                                                                                                                                                                                                          |
| IPC generation          | Creation time of the logical interface when this MAC address was learned. Used for debugging.                                                                                                                                                                                                      |

## Sample Output

### show vpls mac-table

```

user@host> show vpls mac-table
MAC flags (S -static MAC, D -dynamic MAC,
 SE -Statistics enabled, NM -Non configured MAC)

Routing instance : vpls_ldp1
VLAN : 223
 MAC MAC Logical
 address flags interface
 00:90:69:9c:1c:5d D ge-0/2/5.400

```



MAC flags (S -static MAC, D -dynamic MAC,  
SE -Statistics enabled, NM -Non configured MAC)

Routing instance : vpls\_red  
VLAN : 401

| MAC<br>address    | MAC<br>flags | Logical<br>interface |
|-------------------|--------------|----------------------|
| 00:00:aa:12:12:12 | D            | lsi.1051138          |
| 00:05:85:74:9f:f0 | D            | lsi.1051138          |

#### show vpls mac-table (with Layer 2 Services over GRE Interfaces)

user@host> show vpls mac-table  
MAC flags (S -static MAC, D -dynamic MAC, L -locally learned  
SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

Routing instance : vpls\_4site:1000  
Bridging domain : \_\_vpls\_4site:1000\_\_, MAC

| address           | flags | MAC | Logical<br>interface |
|-------------------|-------|-----|----------------------|
| 00:01:01:00:01:f4 | D,SE  |     | ge-4/2/0.1000        |
| 00:02:01:33:01:f4 | D,SE  |     | lsi.1052004          |
| 00:03:00:32:01:f4 | D,SE  |     | lsi.1048840          |
| 00:04:00:14:01:f4 | D,SE  |     | lsi.1052005          |
| 00:02:01:33:02:f7 | D,SE  |     | gr-1/2/10.10         |

#### show vpls mac-table (with VXLAN enabled)

user@host> show vpls mac-table  
MAC flags (S -static MAC, D -dynamic MAC, L -locally learned  
SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

Routing instance : vpls\_4site:1000  
Bridging domain : \_\_vpls\_4site:1000\_\_, VLAN : 4094,4093  
VXLAN: Id : 300, Multicast group: 226.1.1.3

| MAC<br>address    | MAC<br>flags | Logical<br>interface |
|-------------------|--------------|----------------------|
| 00:01:01:00:01:f4 | D,SE         | ge-4/2/0.1000        |
| 00:02:01:33:01:f4 | D,SE         | lsi.1052004          |
| 00:03:00:32:01:f4 | D,SE         | lsi.1048840          |
| 00:04:00:14:01:f4 | D,SE         | lsi.1052005          |
| 00:02:01:33:02:f7 | D,SE         | vtep.1052010         |
| 00:04:00:14:02:f7 | D,SE         | vtep.1052011         |

#### show vpls mac-table count

user@host> show vpls mac-table count  
0 MAC address learned in routing instance \_\_juniper\_private1\_\_

MAC address count per interface within routing instance:

| Logical interface | MAC count |
|-------------------|-----------|
| lc-0/0/0.32769    | 0         |
| lc-0/1/0.32769    | 0         |
| lc-0/2/0.32769    | 0         |
| lc-2/0/0.32769    | 0         |
| lc-0/3/0.32769    | 0         |
| lc-2/1/0.32769    | 0         |
| lc-9/0/0.32769    | 0         |
| lc-11/0/0.32769   | 0         |
| lc-2/2/0.32769    | 0         |
| lc-9/1/0.32769    | 0         |
| lc-11/1/0.32769   | 0         |

|                 |   |
|-----------------|---|
| 1c-2/3/0.32769  | 0 |
| 1c-9/2/0.32769  | 0 |
| 1c-11/2/0.32769 | 0 |
| 1c-11/3/0.32769 | 0 |
| 1c-9/3/0.32769  | 0 |

MAC address count per learn VLAN within routing instance:

| Learn VLAN ID | MAC count |
|---------------|-----------|
| 0             | 0         |

1 MAC address learned in routing instance vpls\_ldp1

MAC address count per interface within routing instance:

| Logical interface | MAC count |
|-------------------|-----------|
| lsi.1051137       | 0         |
| ge-0/2/5.400      | 1         |

MAC address count per learn VLAN within routing instance:

| Learn VLAN ID | MAC count |
|---------------|-----------|
| 0             | 1         |

1 MAC address learned in routing instance vpls\_red

MAC address count per interface within routing instance:

| Logical interface | MAC count |
|-------------------|-----------|
| ge-0/2/5.300      | 1         |

MAC address count per learn VLAN within routing instance:

| Learn VLAN ID | MAC count |
|---------------|-----------|
| 0             | 1         |

### show vpls mac-table detail

user@host> show vpls mac-table detail

MAC address: 00:90:69:9c:1c:5d

Routing instance: vpls\_ldp1

Learning interface: ge-0/2/5.400

Layer 2 flags: in\_ifd, in\_ifl, in\_vlan, kernel

Epoch: 0

Sequence number: 1

Learning mask: 0x1

IPC generation: 0

MAC address: 00:90:69:9c:1c:5d

Routing instance: vpls\_red

Learning interface: ge-0/2/5.300

Layer 2 flags: in\_ifd, in\_ifl, in\_vlan, kernel

Epoch: 0

Sequence number: 1

Learning mask: 0x1

IPC generation: 0

### show vpls mac-table extensive

user@host> show vpls mac-table extensive

MAC address: 00:10:00:01:00:00

Routing instance: vpls\_1

Bridging domain: \_\_vpls\_1\_\_, VLAN : NA

Learning interface: lsi.1049165

Base learning interface: lsi.1049165

Layer 2 flags: in\_hash, in\_ifd, in\_ifl, in\_vlan, in\_rtt, kernel, in\_ifbd

Epoch: 0

Sequence number: 1

Learning mask: 0x00000001

```
MAC address: 00:10:00:01:00:01
Routing instance: vpls_1
 Bridging domain: __vpls_1__, VLAN : NA
 Learning interface: lsi.1049165
 Base learning interface: lsi.1049165
 Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
 Epoch: 0 Sequence number: 1
 Learning mask: 0x00000001

MAC address: 00:10:00:01:00:02
Routing instance: vpls_1
 Bridging domain: __vpls_1__, VLAN : NA
 Learning interface: lsi.1049165
 Base learning interface: lsi.1049165
 Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
 Epoch: 0 Sequence number: 1
 Learning mask: 0x00000001

MAC address: 00:10:00:01:00:03
Routing instance: vpls_1
 Bridging domain: __vpls_1__, VLAN : NA
 Learning interface: lsi.1049165
 Base learning interface: lsi.1049165
 Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
 Epoch: 0 Sequence number: 1
 Learning mask: 0x00000001
```



## CHAPTER 95

# VXLAN Operational Commands

- `show bridge mac-table`
- `show vpls mac-table`

## show bridge mac-table

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show bridge mac-table</b><br><brief   count   detail   extensive><br><bridge-domain (all   <i>bridge-domain-name</i> )><br><global-count><br><interface <i>interface-name</i> ><br><mac-address><br><vlan-id (all-vlan   <i>vlan-id</i> )>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Command introduced in Junos OS Release 8.4.<br>Command introduced in Junos OS Release 15.1.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | (MX Series routers only) Display Layer 2 media access control (MAC) address information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <b>none</b> —Display all learned Layer 2 MAC address information.<br><br><b>brief   count   detail   extensive</b> —(Optional) Display the specified level of output.<br><br><b>bridge-domain (all   <i>bridge-domain-name</i>)</b> —(Optional) Display learned Layer 2 MAC addresses for all bridging domains or for the specified bridging domain.<br><br><b>global-count</b> —(Optional) Display the total number of learned Layer 2 MAC addresses on the system.<br><br><b>instance <i>instance-name</i></b> —(Optional) Display learned Layer 2 MAC addresses for the specified routing instance.<br><br><b>interface <i>interface-name</i></b> —(Optional) Display learned Layer 2 MAC addresses for the specified interface.<br><br><b>mac-address</b> —(Optional) Display the specified learned Layer 2 MAC address information.<br><br><b>vlan-id (all-vlan   <i>vlan-id</i>)</b> —(Optional) Display learned Layer 2 MAC addresses for all VLANs or for the specified VLAN. |
| <b>Additional Information</b>   | When Layer 2 protocol tunneling is enabled, the tunneling MAC address 01:00:0c:cd:cd:d0 is installed in the MAC table. When the Cisco Discovery Protocol (CDP), Spanning Tree Protocol (STP), or VLAN Trunk Protocol (VTP) is configured for Layer 2 protocol tunneling on an interface, the corresponding protocol MAC address is installed in the MAC table.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>List of Sample Output</b>    | <a href="#">show bridge mac-table on page 2560</a><br><a href="#">show bridge mac-table (with Layer 2 Services over GRE Interfaces) on page 2560</a><br><a href="#">show bridge mac-table (with VXLAN Enabled) on page 2560</a><br><a href="#">show bridge mac-table count on page 2561</a><br><a href="#">show bridge mac-table detail on page 2561</a><br><a href="#">show bridge mac-table instance pbb-evpn on page 2562</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

**Output Fields** Table 230 describes the output fields for the **show bridge mac-table** command. Output fields are listed in the approximate order in which they appear.

**Table 230: show bridge mac-table Output Fields**

| Field Name         | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Routing instance   | Name of the routing instance.                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Bridging domain    | Name of the bridging domain.                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| MAC address        | MAC address or addresses learned on a logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| MAC flags          | Status of MAC address learning properties for each interface: <ul style="list-style-type: none"> <li>• <b>S</b>—Static MAC address is configured.</li> <li>• <b>D</b>—Dynamic MAC address is configured.</li> <li>• <b>L</b>—Locally learned MAC address is configured.</li> <li>• <b>C</b>—Control MAC address is configured.</li> <li>• <b>SE</b>—MAC accounting is enabled.</li> <li>• <b>NM</b>—Non-configured MAC.</li> <li>• <b>R</b>—Remote PE MAC address is configured.</li> </ul> |
| Logical interface  | Name of the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| MAC count          | Number of MAC addresses learned on the specific routing instance or interface.                                                                                                                                                                                                                                                                                                                                                                                                              |
| Learning interface | Name of the logical interface on which the MAC address was learned.                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Learning VLAN      | VLAN ID of the routing instance or bridge domain in which the MAC address was learned.                                                                                                                                                                                                                                                                                                                                                                                                      |
| VXLAN ID/VXLAN     | VXLAN Network Identifier (VNI).                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Layer 2 flags      | Debugging flags signifying that the MAC address is present in various lists.                                                                                                                                                                                                                                                                                                                                                                                                                |
| Epoch              | Spanning Tree Protocol epoch number identifying when the MAC address was learned. Used for debugging.                                                                                                                                                                                                                                                                                                                                                                                       |
| Sequence number    | Sequence number assigned to this MAC address. Used for debugging.                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Learning mask      | Mask of the Packet Forwarding Engines where this MAC address was learned. Used for debugging.                                                                                                                                                                                                                                                                                                                                                                                               |
| IPC generation     | Creation time of the logical interface when this MAC address was learned. Used for debugging.                                                                                                                                                                                                                                                                                                                                                                                               |

## Sample Output

### show bridge mac-table

```
user@host> show bridge mac-table
MAC flags (S -static MAC, D -dynamic MAC, L -locally learned, C -Control MAC
SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

Routing instance : default-switch
Bridging domain : test1, VLAN : 1
 MAC MAC Logical NH RTR
 address flags interface Index ID
01:00:0c:cc:cc:cc S,NM NULL
01:00:0c:cc:cc:cd S,NM NULL
01:00:0c:cd:cd:d0 S,NM NULL
64:87:88:6a:17:d0 D ae0.1
64:87:88:6a:17:f0 D ae0.1
```

### show bridge mac-table (with Layer 2 Services over GRE Interfaces)

```
user@host> show bridge mac-table
MAC flags (S -static MAC, D -dynamic MAC, L -locally learned
SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

Routing instance : default-switch
Bridging domain : vlan-1, VLAN : 1
 MAC MAC Logical
 address flags interface
00:01:01:00:01:f7 D,SE gr-1/2/10.0
00:03:00:32:01:f7 D,SE gr-1/2/10.0
00:00:21:11:11:10 DL ge-1/0/0.0
00:00:21:11:11:11 DL ge-1/1/0.0

Routing instance : default-switch
Bridging domain : vlan-2, VLAN : 2
 MAC MAC Logical
 address flags interface
00:02:01:33:01:f7 D,SE gr-1/2/10.1
00:00:21:11:21:10 DL ge-1/0/0.1
00:00:21:11:21:11 DL ge-1/1/0.1
```

### show bridge mac-table (with VXLAN Enabled)

```
user@host> show bridge mac-table
MAC flags (S -static MAC, D -dynamic MAC, L -locally learned
SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

Routing instance : default-switch
Bridging domain : vlan-1, VLAN : 1
VXLAN: Id : 100, Multicast group: 226.1.1.1
 MAC MAC Logical
 address flags interface
00:01:01:00:01:f7 D,SE vtep.1052010
00:03:00:32:01:f7 D,SE vtep.1052011
00:00:21:11:11:10 DL ge-1/0/0.0
00:00:21:11:11:11 DL ge-1/1/0.0
```



```

Routing instance : default-switch
Bridging domain : vlan-2, VLAN : 2, VXLAN : 200
VXLAN: Id : 200, Multicast group: 226.1.1.2
MAC MAC Logical
address flags interface
00:02:01:33:01:f7 D,SE vtep.1052010
00:04:00:14:01:f7 D,SE vtep.1052011
00:00:21:11:21:10 DL ge-1/0/0.1
00:00:21:11:21:11 DL ge-1/1/0.1

```

### show bridge mac-table count

```

user@host> show bridge mac-table count
2 MAC address learned in routing instance vs1 bridge domain vlan100

MAC address count per interface within routing instance:
Logical interface MAC count
ge-11/0/3.0 1
ge-11/1/4.100 0
ge-11/1/1.100 0
ge-11/1/0.100 0
xe-10/2/0.100 1
xe-10/0/0.100 0

MAC address count per learn VLAN within routing instance:
Learn VLAN ID MAC count
0 2

0 MAC address learned in routing instance vs1 bridge domain vlan200

MAC address count per interface within routing instance:
Logical interface MAC count
ge-11/1/0.200 0
ge-11/1/1.200 0
ge-11/1/4.200 0
xe-10/0/0.200 0
xe-10/2/0.200 0

MAC address count per learn VLAN within routing instance:
Learn VLAN ID MAC count
0 0

```

### show bridge mac-table detail

```

user@host> show bridge mac-table detail
MAC address: 00:00:00:19:1c:db
Routing instance: vs1
Bridging domain: vlan100
Learning interface: ge-11/0/3.0 Learning VLAN: 0
Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
Epoch: 4 Sequence number: 0
Learning mask: 0x800 IPC generation: 0

MAC address: 00:00:00:59:3a:2f
Routing instance: vs1
Bridging domain: vlan100
Learning interface: xe-10/2/0.100 Learning VLAN: 0
Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
Epoch: 7 Sequence number: 0
Learning mask: 0x400 IPC generation: 0

```

**show bridge mac-table instance pbb-evpn**

```
user@host> show bridge mac-table instance pbb-evpn
Routing instance : pbb-evpn
Bridging domain : isid-bd10000, ISID : 10000
MAC MAC Logical NH RTR
address flags interface Index ID
00:19:e2:b0:76:eb D cbp.1000
aa:bb:cc:dd:ee:f2 DC
aa:bb:cc:dd:ee:f3 DC 1048576 1048576
 1048575 1048575
```

## show vpls mac-table

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show vpls mac-table &lt;brief   detail   extensive   summary&gt; &lt;bridge-domain <i>bridge-domain-name</i>&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;mac-address&gt; &lt;vlan-id <i>vlan-id-number</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 8.5.</p> <p>Command introduced in Junos OS Release 15.1</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Display learned virtual private LAN service (VPLS) media access control (MAC) address information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <p><b>none</b>—Display all learned VPLS MAC address information.</p> <p><b>brief   detail   extensive   summary</b>—(Optional) Display the specified level of output.</p> <p><b>bridge-domain <i>bridge-domain-name</i></b>—(Optional) Display learned VPLS MAC addresses for the specified bridge domain.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display learned VPLS MAC addresses for the specified instance.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display learned VPLS MAC addresses for the specified instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Display learned VPLS MAC addresses for all logical systems or for the specified logical system.</p> <p><b>mac-address</b>—(Optional) Display the specified learned VPLS MAC address information..</p> <p><b>vlan-id <i>vlan-id-number</i></b>—(Optional) Display learned VPLS MAC addresses for the specified VLAN.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>List of Sample Output</b>    | <p><a href="#">show vpls mac-table on page 2564</a></p> <p><a href="#">show vpls mac-table (with Layer 2 Services over GRE Interfaces) on page 2565</a></p> <p><a href="#">show vpls mac-table (with VXLAN enabled) on page 2565</a></p> <p><a href="#">show vpls mac-table count on page 2565</a></p> <p><a href="#">show vpls mac-table detail on page 2566</a></p> <p><a href="#">show vpls mac-table extensive on page 2566</a></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Output Fields</b>            | Table 229 describes the output fields for the <b>show vpls mac-table</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

Table 231: show vpls mac-table Output fields

| Field Name              | Field Description                                                                                                                                                                                                                                                                                  |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Routing instance        | Name of the routing instance.                                                                                                                                                                                                                                                                      |
| Bridging domain         | Name of the bridging domain.                                                                                                                                                                                                                                                                       |
| MAC address             | MAC address or addresses learned on a logical interface.                                                                                                                                                                                                                                           |
| MAC flags               | Status of MAC address learning properties for each interface: <ul style="list-style-type: none"> <li>• <b>S</b>—Static MAC address configured.</li> <li>• <b>D</b>—Dynamic MAC address learned.</li> <li>• <b>SE</b>—MAC accounting is enabled.</li> <li>• <b>NM</b>—Nonconfigured MAC.</li> </ul> |
| Logical interface       | Name of the logical interface.                                                                                                                                                                                                                                                                     |
| MAC count               | Number of MAC addresses learned on a specific routing instance or interface.                                                                                                                                                                                                                       |
| Learning interface      | Logical interface or logical Label Switched Interface (LSI) the address is learned on.                                                                                                                                                                                                             |
| Base learning interface | Base learning interface of the MAC address. This field is introduced in Junos OS Release 14.2.                                                                                                                                                                                                     |
| Learn VLAN ID/VLAN      | VLAN ID of the routing instance or bridge domain in which the MAC address was learned.                                                                                                                                                                                                             |
| VXLAN ID/VXLAN          | VXLAN Network Identifier (VNI)                                                                                                                                                                                                                                                                     |
| Layer 2 flags           | Debugging flags signifying that the MAC address is present in various lists.                                                                                                                                                                                                                       |
| Epoch                   | Spanning Tree Protocol epoch number identifying when the MAC address was learned. Used for debugging.                                                                                                                                                                                              |
| Sequence number         | Sequence number assigned to this MAC address. Used for debugging.                                                                                                                                                                                                                                  |
| Learning mask           | Mask of Packet Forwarding Engines where this MAC address was learned. Used for debugging.                                                                                                                                                                                                          |
| IPC generation          | Creation time of the logical interface when this MAC address was learned. Used for debugging.                                                                                                                                                                                                      |

## Sample Output

### show vpls mac-table

```

user@host> show vpls mac-table
MAC flags (S -static MAC, D -dynamic MAC,
 SE -Statistics enabled, NM -Non configured MAC)

Routing instance : vpls_ldp1
VLAN : 223
 MAC MAC Logical
 address flags interface
 00:90:69:9c:1c:5d D ge-0/2/5.400

```

MAC flags (S -static MAC, D -dynamic MAC,  
SE -Statistics enabled, NM -Non configured MAC)

Routing instance : vpls\_red  
VLAN : 401

| MAC<br>address    | MAC<br>flags | Logical<br>interface |
|-------------------|--------------|----------------------|
| 00:00:aa:12:12:12 | D            | lsi.1051138          |
| 00:05:85:74:9f:f0 | D            | lsi.1051138          |

#### show vpls mac-table (with Layer 2 Services over GRE Interfaces)

user@host> show vpls mac-table  
MAC flags (S -static MAC, D -dynamic MAC, L -locally learned  
SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

Routing instance : vpls\_4site:1000  
Bridging domain : \_\_vpls\_4site:1000\_\_, MAC

| address           | flags | MAC | Logical<br>interface |
|-------------------|-------|-----|----------------------|
| 00:01:01:00:01:f4 | D,SE  |     | ge-4/2/0.1000        |
| 00:02:01:33:01:f4 | D,SE  |     | lsi.1052004          |
| 00:03:00:32:01:f4 | D,SE  |     | lsi.1048840          |
| 00:04:00:14:01:f4 | D,SE  |     | lsi.1052005          |
| 00:02:01:33:02:f7 | D,SE  |     | gr-1/2/10.10         |

#### show vpls mac-table (with VXLAN enabled)

user@host> show vpls mac-table  
MAC flags (S -static MAC, D -dynamic MAC, L -locally learned  
SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

Routing instance : vpls\_4site:1000  
Bridging domain : \_\_vpls\_4site:1000\_\_, VLAN : 4094,4093  
VXLAN: Id : 300, Multicast group: 226.1.1.3

| MAC<br>address    | MAC<br>flags | Logical<br>interface |
|-------------------|--------------|----------------------|
| 00:01:01:00:01:f4 | D,SE         | ge-4/2/0.1000        |
| 00:02:01:33:01:f4 | D,SE         | lsi.1052004          |
| 00:03:00:32:01:f4 | D,SE         | lsi.1048840          |
| 00:04:00:14:01:f4 | D,SE         | lsi.1052005          |
| 00:02:01:33:02:f7 | D,SE         | vtep.1052010         |
| 00:04:00:14:02:f7 | D,SE         | vtep.1052011         |

#### show vpls mac-table count

user@host> show vpls mac-table count  
0 MAC address learned in routing instance \_\_juniper\_private1\_\_

MAC address count per interface within routing instance:

| Logical interface | MAC count |
|-------------------|-----------|
| lc-0/0/0.32769    | 0         |
| lc-0/1/0.32769    | 0         |
| lc-0/2/0.32769    | 0         |
| lc-2/0/0.32769    | 0         |
| lc-0/3/0.32769    | 0         |
| lc-2/1/0.32769    | 0         |
| lc-9/0/0.32769    | 0         |
| lc-11/0/0.32769   | 0         |
| lc-2/2/0.32769    | 0         |
| lc-9/1/0.32769    | 0         |
| lc-11/1/0.32769   | 0         |

|                 |   |
|-----------------|---|
| 1c-2/3/0.32769  | 0 |
| 1c-9/2/0.32769  | 0 |
| 1c-11/2/0.32769 | 0 |
| 1c-11/3/0.32769 | 0 |
| 1c-9/3/0.32769  | 0 |

MAC address count per learn VLAN within routing instance:

| Learn VLAN ID | MAC count |
|---------------|-----------|
| 0             | 0         |

1 MAC address learned in routing instance vpls\_ldp1

MAC address count per interface within routing instance:

| Logical interface | MAC count |
|-------------------|-----------|
| lsi.1051137       | 0         |
| ge-0/2/5.400      | 1         |

MAC address count per learn VLAN within routing instance:

| Learn VLAN ID | MAC count |
|---------------|-----------|
| 0             | 1         |

1 MAC address learned in routing instance vpls\_red

MAC address count per interface within routing instance:

| Logical interface | MAC count |
|-------------------|-----------|
| ge-0/2/5.300      | 1         |

MAC address count per learn VLAN within routing instance:

| Learn VLAN ID | MAC count |
|---------------|-----------|
| 0             | 1         |

### show vpls mac-table detail

```

user@host> show vpls mac-table detail
MAC address: 00:90:69:9c:1c:5d
Routing instance: vpls_ldp1
Learning interface: ge-0/2/5.400
Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
Epoch: 0 Sequence number: 1
Learning mask: 0x1 IPC generation: 0

MAC address: 00:90:69:9c:1c:5d
Routing instance: vpls_red
Learning interface: ge-0/2/5.300
Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
Epoch: 0 Sequence number: 1
Learning mask: 0x1 IPC generation: 0

```

### show vpls mac-table extensive

```

user@host> show vpls mac-table extensive

MAC address: 00:10:00:01:00:00
Routing instance: vpls_1
Bridging domain: __vpls_1__, VLAN : NA
Learning interface: lsi.1049165
Base learning interface: lsi.1049165
Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
Epoch: 0 Sequence number: 1
Learning mask: 0x00000001

```

```
MAC address: 00:10:00:01:00:01
Routing instance: vpls_1
 Bridging domain: __vpls_1__, VLAN : NA
 Learning interface: lsi.1049165
 Base learning interface: lsi.1049165
 Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
 Epoch: 0 Sequence number: 1
 Learning mask: 0x00000001

MAC address: 00:10:00:01:00:02
Routing instance: vpls_1
 Bridging domain: __vpls_1__, VLAN : NA
 Learning interface: lsi.1049165
 Base learning interface: lsi.1049165
 Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
 Epoch: 0 Sequence number: 1
 Learning mask: 0x00000001

MAC address: 00:10:00:01:00:03
Routing instance: vpls_1
 Bridging domain: __vpls_1__, VLAN : NA
 Learning interface: lsi.1049165
 Base learning interface: lsi.1049165
 Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
 Epoch: 0 Sequence number: 1
 Learning mask: 0x00000001
```





# High Availability Feature Guide for QFX10000 Switches



## PART 40

# Adaptive Load Balancing

- [Configuring Load Balancing on page 2573](#)



# Configuring Load Balancing

- [Understanding Aggregated Ethernet Load Balancing on page 2573](#)
- [Configuring Adaptive Load Balancing on page 2575](#)

## Understanding Aggregated Ethernet Load Balancing

---

The link aggregation feature is used to bundle several physical aggregated Ethernet interfaces to form one logical interface. One or more links are aggregated to form a virtual link or link aggregation group (LAG). The MAC client treats this virtual link as if it were a single link. Link aggregation increases bandwidth, provides graceful degradation as failure occurs, and increases availability.

In addition to these benefits, an aggregated Ethernet bundle is enhanced to provide load-balancing capabilities that ensure that the link utilization among the member links of the aggregated Ethernet bundle are fully and efficiently utilized.

The load-balancing feature allows a device to divide incoming and outgoing traffic along multiple paths or interfaces in order to reduce congestion in the network. Load balancing improves the utilization of various network paths and provides more effective network bandwidth.

Typically, the applications that use load balancing include:

- Aggregated Interfaces (Layer 2)

Aggregated Interfaces (also called AE for aggregated Ethernet, and AS for aggregated SONET) are a Layer 2 mechanism for load-balancing across multiple interfaces between two devices. Because this is a Layer 2 load-balancing mechanism, all of the individual component links must be between the same two devices on each end. Junos OS supports a non-signaled (static) configuration for Ethernet and SONET, as well as the 802.3ad standardized LACP protocol for negotiation over Ethernet links.

- Equal-Cost Multipath (ECMP) (Layer 3)

By default, when there are multiple equal-cost paths to the same destination for the active route, Junos OS uses a hash algorithm to choose one of the next-hop addresses to install in the forwarding table. Whenever the set of next hops for a destination changes in any way, the next-hop address is rechosen using the hash algorithm. There is also an option that allows multiple next-hop addresses to be installed in the forwarding table, known as per-packet load balancing.

ECMP load balancing can be:

- Across BGP paths (BGP multipath)
- Within a BGP path, across multiple LSPs

In complex Ethernet topologies, traffic imbalances occur due to increased traffic flow, and load balancing becomes challenging for some of the following reasons:

- Incorrect load balancing by aggregate next hops
- Incorrect packet hash computation
- Insufficient variance in the packet flow
- Incorrect pattern selection

As a result of traffic imbalance, the load is not well distributed causing congestion in certain links, whereas some other links are not efficiently utilized.

To overcome these challenges, Junos OS provides the following solutions for resolving the genuine traffic imbalance on aggregated Ethernet bundles (IEEE 802.3ad).

- Adaptive Load Balancing

Adaptive load balancing uses a feedback mechanism to correct a genuine traffic imbalance. To correct the imbalance weights, the bandwidth and packet stream of links are adapted to achieve efficient traffic distribution across the links in an AE bundle.

To configure adaptive load balancing, include the **adaptive** statement at the **[edit interfaces aex aggregated-ether-options load-balance]** hierarchy level.



**NOTE:** Adaptive load balancing is not supported if the VLAN ID is configured on the aggregated Ethernet interface. This limitation affects the PTX Series Packet Transport Routers and QFX10000 switches only.

---

To configure the tolerance value as a percentage, include the **tolerance** optional keyword at the **[edit interfaces aex aggregated-ether-options load-balance adaptive]** hierarchy level.

To configure adaptive load balancing based on packets per second (instead of the default bits per second setting), include the **pps** optional keyword at the **[edit interfaces aex aggregated-ether-options load-balance adaptive]** hierarchy level.

To configure the scan interval for the hash value based on the sample rate for the last two seconds, include the **scan-interval** optional keyword at the **[edit interfaces aex aggregated-ether-options load-balance adaptive]** hierarchy level.



**NOTE:** The **pps** and **scan-interval** optional keywords are supported on PTX Series Packet Transport Routers only.

- Per-Packet Random Spray Load Balancing

When the adaptive load-balancing option fails, per-packet random spray load balancing serves as a last resort. It ensures that the members of an AE bundle are equally loaded without taking bandwidth into consideration. Per packet causes packet reordering and hence is recommended only if the applications absorb reordering. Per-packet random spray eliminates traffic imbalance that occurs as a result of software errors, except for packet hash.

To configure per-packet random spray load balancing, include the **per-packet** statement at the **[edit interfaces aex aggregated-ether-options load-balance]** hierarchy level.



**NOTE:** The Per-Packet option for load balancing is not supported on PTX Series Packet Transport Routers.

The aggregated Ethernet load-balancing solutions are mutually exclusive. When more than one of the load-balancing solutions is configured, the solution that is configured last overrides the previously configured one. You can verify the load-balancing solution being used by issuing the **show interfaces aex aggregated-ether-options load-balance** command.

#### Related Documentation

- *Example: Configuring Aggregated Ethernet Load Balancing*

## Configuring Adaptive Load Balancing

This topic describes how to configure adaptive load balancing. Adaptive load balancing maintains efficient utilization of member link bandwidth for an aggregated Ethernet (AE) bundle. Adaptive load balancing uses a feedback mechanism to correct traffic load imbalance by adjusting the bandwidth and packet streams on links within an AE bundle.

Before you begin:

- Configure a set of interfaces with a protocol family and IP address. These interfaces can make up the membership for the AE bundle.
- Create an AE bundle by configuring a set of router interfaces as aggregated Ethernet and with a specific AE group identifier.

To configure adaptive load balancing for an AE bundles:

1. Enable adaptive load balancing on the AE bundle:

```
[edit interfaces ae-x aggregated-ether-options load-balance]
user@router# set adaptive
```

2. Configure the scan interval value for adaptive load balancing on the AE bundle. The scan interval value determines the length of the traffic scan by multiplying the integer value with a 30-second time period:

```
[edit interfaces ae-x aggregated-ether-options load-balance adaptive]
user@router# set scan-interval multiplier
```

3. Configure the tolerance percentage value. The tolerance value determines the allowed deviation in the traffic rates among the members of the AE bundle before the router triggers an adaptive load balancing update:

```
[edit interfaces ae-x aggregated-ether-options load-balance adaptive]
user@router# set tolerance percentage
```

4. (Optional) Enable packet-per-second-based adaptive load balancing on the AE bundle:

```
[edit interfaces ae-x aggregated-ether-options load-balance adaptive]
user@router# set pps
```

**Related  
Documentation**

- [Understanding Aggregated Ethernet Load Balancing on page 2573](#)
- [Example: Configuring Aggregated Ethernet Load Balancing](#)
- [adaptive on page 2656](#)



## PART 41

# Graceful Restart

- [Configuring Graceful Restart on page 2579](#)



# Configuring Graceful Restart

- [Graceful Restart Concepts on page 2579](#)
- [Configuring Routing Protocols Graceful Restart on page 2580](#)

## Graceful Restart Concepts

---

With routing protocols, any service interruption requires that an affected router recalculate adjacencies with neighboring routers, restore routing table entries, and update other protocol-specific information. An unprotected restart of a router can result in forwarding delays, route flapping, wait times stemming from protocol reconvergence, and even dropped packets. The main benefits of graceful restart are uninterrupted packet forwarding and temporary suppression of all routing protocol updates. Graceful restart enables a router to pass through intermediate convergence states that are hidden from the rest of the network.

Three main types of graceful restart are available on Juniper Networks routing platforms:

- Graceful restart for aggregate and static routes and for routing protocols—Provides protection for aggregate and static routes and for Border Gateway Protocol (BGP), End System-to-Intermediate System (ES-IS), Intermediate System-to-Intermediate System (IS-IS), Open Shortest Path First (OSPF), Routing Information Protocol (RIP), next-generation RIP (RIPng), and Protocol Independent Multicast (PIM) sparse mode routing protocols.
- Graceful restart for MPLS-related protocols—Provides protection for Label Distribution Protocol (LDP), Resource Reservation Protocol (RSVP), circuit cross-connect (CCC), and translational cross-connect (TCC). (Not supported on OCX Series switches.)
- Graceful restart for virtual private networks (VPNs)—Provides protection for Layer 2 and Layer 3 VPNs.

Graceful restart works similarly for routing protocols and MPLS protocols and combines components of these protocol types to enable graceful restart in VPNs. The main benefits of graceful restart are uninterrupted packet forwarding and temporary suppression of all routing protocol updates. Graceful restart thus enables a router to pass through intermediate convergence states that are hidden from the rest of the network.

Most graceful restart implementations define two types of routers—the restarting router and the helper router. The restarting router requires rapid restoration of forwarding state

information so it can resume the forwarding of network traffic. The helper router assists the restarting router in this process. Graceful restart configuration statements typically affect either the restarting router or the helper router.

**Related  
Documentation**

- *Understanding High Availability Features on Juniper Networks Routers*
- *Graceful Restart System Requirements*
- *Graceful Restart for Aggregate and Static Routes*
- *Graceful Restart and Routing Protocols*
- [Graceful Restart and MPLS-Related Protocols on page 4969](#)
- *Graceful Restart and Layer 2 and Layer 3 VPNs*
- *Graceful Restart on Logical Systems*
- *Configuring Graceful Restart*
- *Configuring Graceful Restart for QFabric Systems*

---

## Configuring Routing Protocols Graceful Restart

This topic includes the following sections:

- [Enabling Graceful Restart on page 2580](#)
- [Configuring Graceful Restart Options for BGP on page 2581](#)
- [Configuring Graceful Restart Options for ES-IS on page 2582](#)
- [Configuring Graceful Restart Options for IS-IS on page 2582](#)
- [Configuring Graceful Restart Options for OSPF and OSPFv3 on page 2583](#)
- [Configuring Graceful Restart Options for RIP and RIPng on page 2585](#)
- [Configuring Graceful Restart Options for PIM Sparse Mode on page 2585](#)
- [Tracking Graceful Restart Events on page 2586](#)

### Enabling Graceful Restart

By default, graceful restart is disabled. To enable graceful restart, include the **graceful-restart** statement at the **[edit routing-instance *instance-name* routing-options]** or **[edit routing-options]** hierarchy level.

For example:

```
routing-options {
 graceful-restart;
}
```

To configure the duration of the graceful restart period, include the **restart-duration** at the **[edit routing-options graceful-restart]** hierarchy level.



**NOTE:** Helper mode (the ability to assist a neighboring router attempting a graceful restart) is enabled by default when you start the routing platform, even if graceful restart is not enabled. You can disable helper mode on a per-protocol basis.

```
[edit]
routing-options {
 graceful-restart {
 disable;
 restart-duration seconds;
 }
}
```

To disable graceful restart globally, include the **disable** statement at the **[edit routing-options graceful-restart]** hierarchy level.

When graceful restart is enabled for all routing protocols at the **[edit routing-options graceful-restart]** hierarchy level, you can disable graceful restart on a per-protocol basis.



**NOTE:** If you configure graceful restart after a BGP or LDP session has been established, the BGP or LDP session restarts and the peers negotiate graceful restart capabilities. Also, the BGP peer routing statistics are reset to zero.

## Configuring Graceful Restart Options for BGP

To configure the duration of the BGP graceful restart period, include the **restart-time** statement at the **[edit protocols bgp graceful-restart]** hierarchy level. To set the length of time the router waits to receive messages from restarting neighbors before declaring them down, include the **stale-routes-time** statement at the **[edit protocols bgp graceful-restart]** hierarchy level.

```
[edit]
protocols {
 bgp {
 graceful-restart {
 disable;
 restart-time seconds;
 stale-routes-time seconds;
 }
 }
}
routing-options {
 graceful-restart;
}
```

To disable BGP graceful restart capability for all BGP sessions, include the **disable** statement at the **[edit protocols bgp graceful-restart]** hierarchy level.



**NOTE:** To set BGP graceful restart properties or disable them for a group, include the desired statements at the `[edit protocols bgp group group-name graceful-restart]` hierarchy level.

To set BGP graceful restart properties or disable them for a specific neighbor in a group, include the desired statements at the `[edit protocols bgp group group-name neighbor ip-address graceful-restart]` hierarchy level.



**NOTE:** Configuring graceful restart for BGP resets the BGP peer routing statistics to zero. Also, existing BGP sessions restart, and the peers negotiate graceful restart capabilities.

## Configuring Graceful Restart Options for ES-IS

On J Series Services Routers, to configure the duration of the ES-IS graceful restart period, include the **restart-duration** statement at the `[edit protocols esis graceful-restart]` hierarchy level.

```
[edit]
protocols {
 esis {
 graceful-restart {
 disable;
 restart-duration seconds;
 }
 }
}
routing-options {
 graceful-restart;
}
```

To disable ES-IS graceful restart capability, include the **disable** statement at the `[edit protocols esis graceful-restart]` hierarchy level.

## Configuring Graceful Restart Options for IS-IS

To configure the duration of the IS-IS graceful restart period, include the **restart-duration** statement at the `[edit protocols isis graceful-restart]` hierarchy level.

```
[edit]
protocols {
 isis {
 graceful-restart {
 disable;
 helper-disable;
 restart-duration seconds;
 }
 }
}
routing-options {
 graceful-restart;
}
```

```
}
```

To disable IS-IS graceful restart helper capability, include the **helper-disable** statement at the **[edit protocols isis graceful-restart]** hierarchy level. To disable IS-IS graceful restart capability, include the **disable** statement at the **[edit protocols isis graceful-restart]** hierarchy level.



**NOTE:** If you configure Bidirectional Forwarding Detection (BFD) and graceful restart for IS-IS, graceful restart might not work as expected.



**NOTE:** If adjacencies between the Routing Engine and the neighboring peer 'helper' routers time out, graceful restart protocol extensions are unable to notify the peer 'helper' routers about the impending restart. Graceful restart can then stop and cause interruptions in traffic.

To ensure that these adjacencies are kept, change the hold-time for IS-IS protocols from the default of 27 seconds to a value higher than 40 seconds.



**NOTE:** You can also track graceful restart events with the **traceoptions** statement at the **[edit protocols isis]** hierarchy level. For more information, see [“Tracking Graceful Restart Events” on page 2586](#).

## Configuring Graceful Restart Options for OSPF and OSPFv3

To configure the duration of the OSPF/OSPFv3 graceful restart period, include the **restart-duration** statement at the **[edit protocols (ospf | ospfv3) graceful-restart]** hierarchy level. To specify the length of time for which the router notifies helper routers that it has completed graceful restart, include the **notify-duration** at the **[edit protocols (ospf | ospfv3) graceful-restart]** hierarchy level. Strict OSPF link-state advertisement (LSA) checking results in the termination of graceful restart by a helping router. To disable strict LSA checking, include the **no-strict-lsa-checking** statement at the **[edit protocols (ospf | ospfv3) graceful-restart]** hierarchy level.

```
[edit]
protocols {
 ospf | ospfv3 {
 graceful-restart {
 disable;
 helper-disable
 no-strict-lsa-checking;
 notify-duration seconds;
 restart-duration seconds;
 }
 }
}
routing-options {
 graceful-restart;
}
```

To disable OSPF/OSPFv3 graceful restart, include the **disable** statement at the **[edit protocols (ospf | ospfv3) graceful-restart]** hierarchy level.

Starting with Release 11.3, the Junos OS supports both the standard (based on RFC 3623, *Graceful OSPF Restart*) and the restart signaling-based (as specified in RFC 4811, RFC 4812, and RFC 4813) helper modes for OSPF version 2 graceful restart configurations. Both the standard and restart signaling-based helper modes are enabled by default. To disable the helper mode for OSPF version 2 graceful restart configurations, include the **helper-disable <both | restart-signaling | standard>** statement at the **[edit protocols ospf graceful-restart]** hierarchy level. Note that the last committed statement always takes precedence over the previous one.

```
[edit protocols ospf]
 graceful-restart {
 helper-disable <both | restart-signaling | standard>
 }
```

To reenable the helper mode, delete the **helper-disable** statement from the configuration by using the **delete protocols ospf graceful-restart helper-disable <restart-signaling | standard | both>** command. In this case also, the last executed command takes precedence over the previous ones.



**NOTE:**

Restart signaling-based helper mode is not supported for OSPFv3 configurations. To disable helper mode for OSPFv3 configurations, include the **helper-disable** statement at the **[edit protocols ospfv3 graceful-restart]** hierarchy level.



**TIP:** You can also track graceful restart events with the **traceoptions** statement at the **[edit protocols (ospf | ospfv3)]** hierarchy level. For more information, see [“Tracking Graceful Restart Events” on page 2586](#).



**NOTE:** You cannot enable OSPFv3 graceful restart between a routing platform running Junos OS Release 7.5 and earlier and a routing platform running Junos OS Release 7.6 or later. As a workaround, make sure both routing platforms use the same Junos OS version.



**NOTE:** If you configure BFD and graceful restart for OSPF, graceful restart might not work as expected.

---



## Configuring Graceful Restart Options for RIP and RIPng

To configure the duration of the RIP or RIPng graceful restart period, include the **restart-time** statement at the **[edit protocols (rip | ripng) graceful-restart]** hierarchy level.

```
[edit]
protocols {
 (rip | ripng) {
 graceful-restart {
 disable;
 restart-time seconds;
 }
 }
}
routing-options {
 graceful-restart;
}
```

To disable RIP or RIPng graceful restart capability, include the **disable** statement at the **[edit protocols (rip | ripng) graceful-restart]** hierarchy level.

## Configuring Graceful Restart Options for PIM Sparse Mode

PIM sparse mode continues to forward existing multicast packet streams during a graceful restart, but does not forward new streams until after the restart is complete. After a restart, the routing platform updates the forwarding state with any updates that were received from neighbors and occurred during the restart period. For example, the routing platform relearns the join and prune states of neighbors during the restart, but does not apply the changes to the forwarding table until after the restart.

PIM sparse mode-enabled routing platforms generate a unique 32-bit random number called a generation identifier. Generation identifiers are included by default in PIM hello messages, as specified in the IETF Internet draft *Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)*. When a routing platform receives PIM hellos containing generation identifiers on a point-to-point interface, Junos OS activates an algorithm that optimizes graceful restart.

Before PIM sparse mode graceful restart occurs, each routing platform creates a generation identifier and sends it to its multicast neighbors. If a PIM sparse mode-enabled routing platform restarts, it creates a new generation identifier and sends it to its neighbors. When a neighbor receives the new identifier, it resends multicast updates to the restarting router to allow it to exit graceful restart efficiently. The restart phase completes when either the PIM state becomes stable or when the restart interval timer expires.

If a routing platform does not support generation identifiers or if PIM is enabled on multipoint interfaces, the PIM sparse mode graceful restart algorithm does not activate, and a default restart timer is used as the restart mechanism.

To configure the duration of the PIM graceful restart period, include the **restart-duration** statement at the **[edit protocols pim graceful-restart]** hierarchy level:

```
[edit]
protocols {
```

```
pim {
 graceful-restart {
 disable;
 restart-duration seconds;
 }
}
routing-options {
 graceful-restart;
}
```

To disable PIM sparse mode graceful restart capability, include the **disable** statement at the **[edit protocols pim graceful-restart]** hierarchy level.



**NOTE:** Multicast forwarding can be interrupted in two ways. First, if the underlying routing protocol is unstable, multicast reverse-path-forwarding (RPF) checks can fail and cause an interruption. Second, because the forwarding table is not updated during the graceful restart period, new multicast streams are not forwarded until graceful restart is complete.

---

## Tracking Graceful Restart Events

To track the progress of a graceful restart event, you can configure graceful restart trace options flags for IS-IS and OSPF/OSPFv3. To configure graceful restart trace options, include the **graceful-restart** statement at the **[edit protocols *protocol* traceoptions flag]** hierarchy level:

```
[edit protocols]
isis {
 traceoptions {
 flag graceful-restart;
 }
}
(ospf | ospf3) {
 traceoptions {
 flag graceful-restart;
 }
}
```

### Related Documentation

- [Graceful Restart Concepts on page 2579](#)
- [Graceful Restart System Requirements](#)
- [Graceful Restart and Routing Protocols](#)
- [Verifying Graceful Restart Operation on page 2707](#)
- [Configuring Graceful Restart](#)

## PART 42

# Nonstop Bridging

- [Configuring Nonstop Bridging on page 2589](#)



# Configuring Nonstop Bridging

- [Nonstop Bridging Concepts on page 2589](#)
- [Nonstop Bridging System Requirements on page 2591](#)
- [Configuring Nonstop Bridging on Switches \(CLI Procedure\) on page 2593](#)

## Nonstop Bridging Concepts


---

Nonstop bridging uses the same infrastructure as graceful Routing Engine switchover (GRES) to preserve interface and kernel information. However, nonstop bridging also saves Layer 2 Control Protocol (L2CP) information by running the Layer 2 Control Protocol process (l2cpd) on the backup Routing Engine.



**NOTE:** To use nonstop bridging, you must first enable graceful Routing Engine switchover on your routing (or switching) platform. For more information about graceful Routing Engine switchover, see [“Understanding Graceful Routing Engine Switchover” on page 2617](#).

[Figure 37](#) shows the system architecture of nonstop bridging and the process a routing (or switching) platform follows to prepare for a switchover.


**Figure 37: Nonstop Bridging Switchover Preparation Process**

The switchover preparation process for nonstop bridging follows these steps:

1. The master Routing Engine starts.
2. The routing platform processes on the master Routing Engine (such as the chassis process [chassisd] and the Layer 2 Control Protocol process [l2cpd]) start.
3. The Packet Forwarding Engine starts and connects to the master Routing Engine.
4. All state information is updated in the system.
5. The backup Routing Engine starts, including the chassis process (chassisd) and the Layer 2 Control Protocol process (l2cpd).
6. The system determines whether graceful Routing Engine switchover and nonstop bridging have been enabled.
7. The kernel synchronization process (ksyncd) synchronizes the backup Routing Engine with the master Routing Engine.
8. For supported protocols, state information is updated directly between the l2cpds on the master and backup Routing Engines.

Figure 38 shows the effects of a switchover on the routing platform.

Figure 38: Nonstop Bridging During a Switchover



The switchover process follows these steps:

1. When keepalives from the master Routing Engine are lost, the system switches over gracefully to the backup Routing Engine.
2. The Packet Forwarding Engine connects to the backup Routing Engine, which becomes the new master. Because the Layer 2 Control Protocol process (l2cpd) and chassis process (chassisd) are already running, these processes do not need to restart.
3. State information learned from the point of the switchover is updated in the system. Forwarding and bridging are continued during the switchover, resulting in minimal packet loss.

#### Related Documentation

- [Understanding High Availability Features on Juniper Networks Routers](#)
- [Nonstop Bridging System Requirements on page 2591](#)
- [Configuring Nonstop Bridging](#)
- [Configuring Nonstop Bridging on Switches \(CLI Procedure\) on page 2593](#)

## Nonstop Bridging System Requirements

This topic contains the following sections:

- [Platform Support on page 2591](#)
- [Protocol Support on page 2592](#)

### Platform Support

Nonstop bridging is supported on MX Series 3D Universal Edge Routers. Your system must be running Junos OS Release 8.4 or later.

Nonstop bridging is supported on EX Series switches with redundant Routing Engines in a Virtual Chassis or in a Virtual Chassis Fabric.

Nonstop bridging is supported on QFX Series switches in a Virtual Chassis or in a Virtual Chassis Fabric.

For a list of the EX Series switches and Layer 2 protocols that support nonstop bridging, see *EX Series Switch Software Features Overview*.



**NOTE:** All Routing Engines configured for nonstop bridging must be running the same Junos OS release.

## Protocol Support

Nonstop bridging is supported for the following Layer 2 control protocols:

- Spanning Tree Protocol (STP)
- Rapid Spanning Tree Protocol (RSTP)
- Multiple Spanning Tree Protocol (MSTP)
- VLAN Spanning Tree Protocol (VSTP)

### Related Documentation

- [Nonstop Bridging Concepts on page 2589](#)
- [Configuring Nonstop Bridging](#)
- [Configuring Nonstop Bridging on Switches \(CLI Procedure\) on page 2593](#)



## Configuring Nonstop Bridging on Switches (CLI Procedure)



**NOTE:** This task uses switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring Nonstop Bridging on EX Series Switches (CLI Procedure)*. For ELS details, see “Getting Started with Enhanced Layer 2 Software” on page 41.

You can configure nonstop bridging (NSB) to provide resilience for Layer 2 protocol sessions on a Juniper Networks EX Series switch with multiple Routing Engines or an EX Series or QFX Series switch in a Virtual Chassis or Virtual Chassis Fabric configuration. Limited support for NSB is also provided on QFX5100 and EX4600 standalone switches, but NSB is enabled *only* during an ISSU.

NSB operates by synchronizing all protocol information for NSB-supported Layer 2 protocols between the master and backup Routing Engines. If the switch has a Routing Engine switchover, the NSB-supported Layer 2 protocol sessions remain active because they are already synchronized on the backup Routing Engine. The Routing Engine switchover is transparent to neighbor devices, which do not detect any changes related to the Layer 2 protocol sessions. The neighboring devices and other devices on the network do not, therefore, have to resynchronize their Layer 2 protocol states to respond to the downtime on the switch—a process that adds network overhead and risks disrupting network performance—when a Routing Engine switchover occurs when NSB is enabled.



**NOTE:** If you are using a QFX5100 or EX4600 standalone switch and you want to use ISSU, configure Graceful Routing Engine switchover (GRES), NSB and nonstop active routing (NSR). You must configure NSB, GRES, and NSR in order to run ISSU. However, GRES, NSB and NSR are enabled *only* during the upgrade. During an ISSU, the Junos OS runs in two separate virtual machines (VMs)—one VM is in the master role acting as the master Routing Engine, and the other VM is in the backup role acting as the backup Routing Engine. The Junos OS is upgraded on the backup VM. After a successful software upgrade, the backup VM then becomes the master VM, and the original master VM is no longer needed and is shut down.

To configure NSB:

1. Enable graceful Routing Engine switchover (GRES):
 

```
[edit chassis redundancy]
user@switch# set graceful-switchover
```
2. Enable NSB:
 

```
[edit protocols layer2-control]
user@switch# set nonstop-bridging
```
3. Synchronize configuration changes between the Routing Engines:
 

```
[edit system]
```

user@switch# **set commit** [synchronize](#)

If you try to commit a configuration that includes NSB without including the **commit synchronize** statement, the commit fails.



**NOTE:** There is no requirement to start the two Routing Engines simultaneously. If the backup Routing Engine is not up when you use the **commit synchronize** statement, the candidate configuration is committed in the master Routing Engine. When the backup Routing Engine comes online, its configuration is automatically synchronized with that of the master.



**BEST PRACTICE:** After a graceful Routing Engine switchover, we recommend that you issue the clear interface statistics (*interface-name* | all) command to reset the cumulative values for local statistics on the new master Routing Engine.

**Related  
Documentation**

- *Performing an In-Service Software Upgrade (ISSU)*
- *Understanding Nonstop Bridging on EX Series Switches*
- [Nonstop Bridging Concepts on page 2589](#)
- *Understanding In-Service Software Upgrade (ISSU)*

## PART 43

# Nonstop Active Routing

- [Configuring Nonstop Active Routing on page 2597](#)



# Configuring Nonstop Active Routing

- [Nonstop Active Routing Concepts on page 2597](#)
- [Nonstop Active Routing System Requirements on page 2600](#)
- [Example: Configuring Nonstop Active Routing on Switches on page 2611](#)

## Nonstop Active Routing Concepts

---

Nonstop active routing (NSR) uses the same infrastructure as graceful Routing Engine switchover (GRES) to preserve interface and kernel information. However, NSR also saves routing protocol information by running the routing protocol process (rpd) on the backup Routing Engine. By saving this additional information, NSR is self-contained and does not rely on helper routers (or switches) to assist the routing platform in restoring routing protocol information. NSR is advantageous in networks in which neighbor routers (or switches) do not support graceful restart protocol extensions. As a result of this enhanced functionality, NSR is a natural replacement for graceful restart.

If you have NSR configured, it is never valid to issue the **restart routing** command in any form on the NSR master Routing Engine. Doing so results in a loss of protocol adjacencies and neighbors and a drop in traffic.



.....

**NOTE:** To use NSR, you must first enable GRES on your routing (or switching) platform. For more information about GRES, see [“Understanding Graceful Routing Engine Switchover” on page 2617](#).

.....



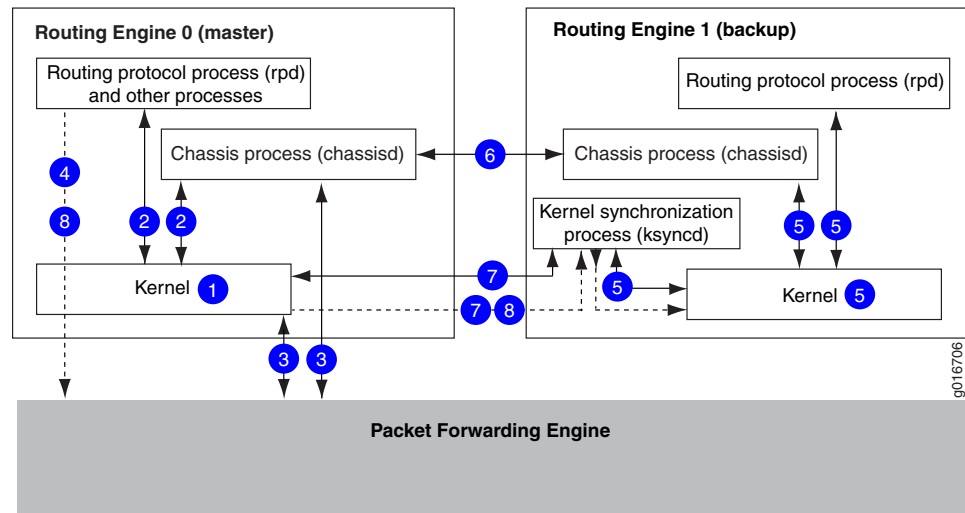
.....

**NOTE:** Due to its synchronization requirements and logic, NSR/GRES performance is limited by the slowest Routing Engine in the system.

.....

Figure 39 shows the system architecture of nonstop active routing and the process a routing (or switching) platform follows to prepare for a switchover.

**Figure 39: Nonstop Active Routing Switchover Preparation Process**




The switchover preparation process for NSR follows these steps:

1. The master Routing Engine starts.
2. The routing (or switching) platform processes on the master Routing Engine (such as the chassis process [chassisd] and the routing protocol process [rpd]) start.
3. The Packet Forwarding Engine starts and connects to the master Routing Engine.
4. All state information is updated in the system.
5. The backup Routing Engine starts, including the chassis process (chassisd) and the routing protocol process (rpd).
6. The system determines whether graceful Routing Engine switchover and nonstop active routing have been enabled.
7. The kernel synchronization process (ksyncd) synchronizes the backup Routing Engine with the master Routing Engine.
8. For supported protocols, state information is updated directly between the routing protocol processes on the master and backup Routing Engines.

Figure 40 shows the effects of a switchover on the routing platform.

Figure 40: Nonstop Active Routing During a Switchover



The switchover process follows these steps:

1. When keepalives from the master Routing Engine are lost, the system switches over gracefully to the backup Routing Engine.
2. The Packet Forwarding Engine connects to the backup Routing Engine, which becomes the new master. Because the routing protocol process (rpd) and chassis process (chassisd) are already running, these processes do not need to restart.
3. State information learned from the point of the switchover is updated in the system. Forwarding and routing are continued during the switchover, resulting in minimal packet loss.
4. Peer routers (or switches) continue to interact with the routing platform as if no change had occurred. Routing adjacencies and session state relying on underlying routing information are preserved and not reset.



**CAUTION:** We recommend that you do not restart the routing protocol process (rpd) on master Routing Engine after enabling NSR, as it disrupts the protocol adjacency/peering sessions, resulting in traffic loss.

#### Related Documentation

- [Understanding High Availability Features on Juniper Networks Routers](#)
- [Nonstop Active Routing System Requirements on page 2600](#)
- [Configuring Nonstop Active Routing](#)

- *Configuring Nonstop Active Routing on Switches*

## Nonstop Active Routing System Requirements

This section contains the following topics:

- [Nonstop Active Routing Platform and Switching Platform Support on page 2600](#)
- [Nonstop Active Routing Protocol and Feature Support on page 2601](#)
- [Nonstop Active Routing BFD Support on page 2604](#)
- [Nonstop Active Routing BGP Support on page 2605](#)
- [Nonstop Active Routing Layer 2 Circuit and VPLS Support on page 2606](#)
- [Nonstop Active Routing PIM Support on page 2606](#)
- [Nonstop Active Routing MSDP Support on page 2609](#)
- [Nonstop Active Routing Support for RSVP-TE LSPs on page 2609](#)

## Nonstop Active Routing Platform and Switching Platform Support

Table 232 lists the platforms that support nonstop active routing (NSR).

**Table 232: Nonstop Active Routing Platform Support**

| Platform                            | Junos OS Release |
|-------------------------------------|------------------|
| M10i router                         | 8.4 or later     |
| M20 router                          | 8.4 or later     |
| M40e router                         | 8.4 or later     |
| M120 router                         | 9.0 or later     |
| M320 router                         | 8.4 or later     |
| MX Series routers                   | 9.0 or later     |
| PTX Series Packet Transport Routers | 12.1R4 or later  |

### NOTE:

Nonstop active routing (NSR) switchover on PTX series is supported only for the following MPLS and VPN protocols and applications using chained composite next hops:

- Labeled BGP
- Layer 2 VPNs excluding Layer 2 interworking (Layer 2 switching)
- Layer 3 VPNs
- LDP
- RSVP



Table 232: Nonstop Active Routing Platform Support (*continued*)

| Platform                                                                                                                                                                                                                                                                                                                                                                                          | Junos OS Release                                               |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------|
| PTX Series Packet Transport Routers                                                                                                                                                                                                                                                                                                                                                               | 12.1R4 or later                                                |
| <p><b>NOTE:</b></p> <p>Nonstop active routing (NSR) switchover on PTX series is supported only for the following MPLS and VPN protocols and applications using chained composite next hops:</p> <ul style="list-style-type: none"> <li>• Labeled BGP</li> <li>• Layer 2 VPNs excluding Layer 2 interworking (Layer 2 stitching)</li> <li>• Layer 3 VPNs</li> <li>• LDP</li> <li>• RSVP</li> </ul> |                                                                |
| PTX Series Packet Transport Routers                                                                                                                                                                                                                                                                                                                                                               | 12.1R4 or later                                                |
| <p><b>NOTE:</b></p> <p>Nonstop active routing (NSR) switchover on PTX series is supported only for the following MPLS and VPN protocols and applications using chained composite next hops:</p> <ul style="list-style-type: none"> <li>• Labeled BGP</li> <li>• Layer 2 VPNs excluding Layer 2 interworking (Layer 2 stitching)</li> <li>• Layer 3 VPNs</li> <li>• LDP</li> <li>• RSVP</li> </ul> |                                                                |
| T320 router, T640 router, and TX Matrix router                                                                                                                                                                                                                                                                                                                                                    | 8.4 or later                                                   |
| Standalone T1600 router                                                                                                                                                                                                                                                                                                                                                                           | 8.5 or later                                                   |
| Standalone T4000 router                                                                                                                                                                                                                                                                                                                                                                           | 12.1R2 or later                                                |
| TX Plus Matrix router                                                                                                                                                                                                                                                                                                                                                                             | 10.0 or later                                                  |
| TX Plus Matrix router with 3D SIBs                                                                                                                                                                                                                                                                                                                                                                | 13.1 or later                                                  |
| EX Series switch with dual Routing Engines or in a Virtual Chassis                                                                                                                                                                                                                                                                                                                                | 10.4 or later for EX Series switches                           |
| EX Series or QFX Series switches in a Virtual Chassis Fabric                                                                                                                                                                                                                                                                                                                                      | 13.2X51-D20 or later for the EX Series and QFX Series switches |



**NOTE:** All Routing Engines configured for nonstop active routing must be running the same Junos OS release.

## Nonstop Active Routing Protocol and Feature Support

Table 233 lists the protocols that are supported by nonstop active routing.

**Table 233: Nonstop Active Routing Protocol and Feature Support**

| Protocol                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Junos OS Release                                                       |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------|
| Aggregated Ethernet interfaces with Link Aggregation Control Protocol (LACP)                                                                                                                                                                                                                                                                                                                                                                                                                                                      | 9.4 or later                                                           |
| Bidirectional Forwarding Detection (BFD)<br><br>For more information, see <a href="#">“Nonstop Active Routing BFD Support” on page 2604.</a>                                                                                                                                                                                                                                                                                                                                                                                      | 8.5 or later                                                           |
| BGP<br><br>For more information, see <a href="#">“Nonstop Active Routing BGP Support” on page 2605.</a>                                                                                                                                                                                                                                                                                                                                                                                                                           | 8.4 or later                                                           |
| Labeled BGP (PTX Series Packet Transport Routers: only)                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | 12.1R4 or later                                                        |
| IS-IS                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 8.4 or later                                                           |
| LDP                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | 8.4 or later                                                           |
| LDP-based virtual private LAN service (VPLS)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | 9.3 or later                                                           |
| LDP OAM (operation, administration, and management) features                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | 9.6 or later                                                           |
| LDP (PTX Series Packet Transport Routers only)<br><br>Nonstop active routing support for LDP includes: <ul style="list-style-type: none"> <li>• LDP unicast transit LSPs</li> <li>• LDP egress LSPs for labeled internal BGP (IBGP) and external BGP (EBGP)</li> <li>• LDP over RSVP transit LSPs</li> <li>• LDP transit LSPs with indexed next hops</li> <li>• LDP transit LSPs with unequal cost load balancing</li> </ul> NOTE: Nonstop active routing is not supported for LDP Point-to-Multipoint LSPs and LDP ingress LSPs. | 12.3R4 or later                                                        |
| Layer 2 circuits                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | (on LDP-based VPLS) 9.2 or later<br><br>(on RSVP-TE LSP) 11.1 or later |
| Layer 2 VPNs                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | 9.1 or later                                                           |
| Layer 2 VPNs (PTX Series Packet Transport Routers only)                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | 12.1R4 or later                                                        |
| NOTE: Nonstop active routing is not supported for Layer 2 interworking (Layer 2 stitching).                                                                                                                                                                                                                                                                                                                                                                                                                                       |                                                                        |

Table 233: Nonstop Active Routing Protocol and Feature Support (*continued*)

| Protocol                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Junos OS Release                                          |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------|
| Layer 3 VPNs (see the first Note after this table for restrictions)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | 9.2 or later                                              |
| Nonstop active routing support for Layer 3 VPNs include:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |                                                           |
| <ul style="list-style-type: none"> <li>• IPv4 labeled-unicast (ingress or egress)</li> <li>• IPv4-vpn unicast (ingress or egress)</li> <li>• IPv6 labeled-unicast (ingress or egress)</li> <li>• IPv6-vpn unicast (ingress or egress)</li> </ul>                                                                                                                                                                                                                                                                                                                                                         |                                                           |
| Layer 3 VPNs (PTX Series Packet Transport Routers only)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | 12.1R4 or later                                           |
| Multicast Source Discovery Protocol (MSDP)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | 12.1 or later                                             |
| For more information, see <a href="#">“Nonstop Active Routing MSDP Support” on page 2609</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |                                                           |
| OSPF/OSPFv3                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | 8.4 or later                                              |
| Protocol Independent Multicast (PIM)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | (for IPv4) 9.3 or later                                   |
| For more information, see <a href="#">“Nonstop Active Routing PIM Support” on page 2606</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | (for IPv6) 10.4 or later                                  |
| RIP and RIP next generation (RIPng)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | 9.0 or later                                              |
| RSVP (PTX Series Packet Transport Routers only)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | 12.1R4 or later                                           |
| Nonstop active routing support for RSVP includes:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |                                                           |
| <ul style="list-style-type: none"> <li>• Point-to-Multipoint LSPs <ul style="list-style-type: none"> <li>• RSVP Point-to-Multipoint ingress, transit, and egress LSPs using existing non-chained next hop.</li> <li>• RSVP Point-to-Multipoint transit LSPs using composite next hops for Point-to-Multipoint label routes.</li> </ul> </li> <li>• Point-to-Point LSPs <ul style="list-style-type: none"> <li>• RSVP Point-to-Point ingress, transit, and egress LSPs using non-chained next hops.</li> <li>• RSVP Point-to-Point transit LSPs using chained composite next hops.</li> </ul> </li> </ul> |                                                           |
| RSVP-TE LSP                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | 9.5 or later                                              |
| For more information, see <a href="#">“Nonstop Active Routing Support for RSVP-TE LSPs” on page 2609</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |                                                           |
| VPLS                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | (LDP-based) 9.1 or later<br>(RSVP-TE-based) 11.2 or later |
| VRRP                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | 13.2 or later                                             |

Table 233: Nonstop Active Routing Protocol and Feature Support (*continued*)

| Protocol | Junos OS Release |
|----------|------------------|
| VRRP     | 13.2 or later    |



**NOTE:** Layer 3 VPN support does not include dynamic GRE tunnels, multicast VPNs, or BGP flow routes.



**NOTE:** If you configure a protocol that is not supported by nonstop active routing, the protocol operates as usual. When a switchover occurs, the state information for the unsupported protocol is not preserved and must be refreshed using the normal recovery mechanisms inherent in the protocol.



**NOTE:** On routers that have logical systems configured on them, only the master logical system supports nonstop active routing.



**NOTE:** On EX9214 switches, the VRRP master state might change during graceful Routing Engine switchover, even when nonstop active routing is enabled.

## Nonstop Active Routing BFD Support

Nonstop active routing supports the Bidirectional Forwarding Detection (BFD) protocol, which uses the topology discovered by routing protocols to monitor neighbors. The BFD protocol is a simple hello mechanism that detects failures in a network. Because BFD is streamlined to be efficient at fast liveness detection, when it is used in conjunction with routing protocols, routing recovery times are improved. With nonstop active routing enabled, BFD session states are not restarted when a Routing Engine switchover occurs.



**NOTE:** BFD session states are saved only for clients using aggregate or static routes or for BGP, IS-IS, OSPF/OSPFv3, or PIM.

When a BFD session is distributed to the Packet Forwarding Engine, BFD packets continue to be sent during a Routing Engine switchover. If nondistributed BFD sessions are to be kept alive during a switchover, you must ensure that the session failure detection time is greater than the Routing Engine switchover time. The following BFD sessions are not distributed to the Packet Forwarding Engine: multihop sessions, tunnel-encapsulated sessions, and sessions over integrated routing and bridging (IRB) interfaces.



**NOTE:** BFD is an intensive protocol that consumes system resources. Specifying a minimum interval for BFD less than 100 ms for Routing Engine-based sessions and 10 ms for distributed BFD sessions can cause undesired BFD flapping. The minimum-interval configuration statement is a BFD liveness detection parameter.

Depending on your network environment, these additional recommendations might apply:

- For large-scale network deployments with a large number of BFD sessions, specify a minimum interval of 300 ms for Routing Engine-based sessions, and 100 ms for distributed BFD sessions.
- For very large-scale network deployments with a large number of BFD sessions, contact Juniper Networks customer support for more information.
- For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing is configured, specify a minimum interval of 10 seconds for Routing Engine-based sessions. For distributed BFD sessions with nonstop active routing configured, the minimum interval recommendations are unchanged and depend only on your network deployment.

## Nonstop Active Routing BGP Support

Nonstop active routing BGP support is subject to the following conditions:

- You must include the **path-selection external-router-ID** statement at the **[edit protocols bgp]** hierarchy level to ensure consistent path selection between the master and backup Routing Engines during and after the nonstop active routing switchover.
- You must include the **advertise-from-main-vpn-tables** statement at the **[edit protocols bgp]** hierarchy level to prevent BGP sessions from going down when route reflector (RR) or autonomous system border router (ASBR) functionality is enabled or disabled on a routing device that has VPN address families configured.
- BGP session uptime and downtime statistics are not synchronized between the primary and backup Routing Engines during Nonstop Active Routing and ISSU. The backup Routing Engine maintains its own session uptime based on the time when the backup first becomes aware of the established sessions. For example, if the backup Routing Engine is rebooted (or if you run **restart routing** on the backup Routing Engine), the backup's uptime is a short duration, because the backup has just learned about the established sessions. If the backup is operating when the BGP sessions first come up on the primary, the uptime on the primary and the uptime on the backup are almost the same duration. After a Routing Engine switchover, the new master continues from the time left on the standby Routing Engine.
- If the BGP peer in the master Routing Engine has negotiated address-family capabilities that are not supported for nonstop active routing, then the corresponding BGP neighbor state on the backup Routing Engine shows as idle. On switchover, the BGP session is reestablished from the new master Routing Engine.

Only the following address families are supported for nonstop active routing.



**NOTE:** Address families are supported only on the main instance of BGP. Only unicast is supported on VRF instances.

- inet unicast
  - inet labeled-unicast
  - inet multicast
  - inet6 labeled-unicast
  - inet6 multicast
  - inet6 unicast
  - route-target
  - l2vpn signaling
  - inet6-vpn unicast
  - inet-vpn unicast
  - inet-mdt
  - iso-vpn
- BGP route dampening does not work on the backup Routing Engine when nonstop active routing is enabled.

## Nonstop Active Routing Layer 2 Circuit and VPLS Support

Nonstop active routing supports Layer 2 circuit and VPLS on both LDP-based and RSVP-TE-based networks. Nonstop active routing support enables the backup Routing Engine to track the label advertised by Layer 2 circuit and VPLS on the primary Routing Engine, and to use the same label after the Routing Engine switchover.

in Junos OS Release 9.6 and later, nonstop active routing support is extended to the Layer 2 circuit and LDP-based VPLS pseudowire redundant configurations.

## Nonstop Active Routing PIM Support

Nonstop active routing supports Protocol Independent Multicast (PIM) with stateful replication on backup Routing Engines. State information replicated on the backup Routing Engine includes information about neighbor relationships, join and prune events, rendezvous point (RP) sets, synchronization between routes and next hops, multicast session states, and the forwarding state between the two Routing Engines.



**NOTE:** Nonstop active routing for PIM is supported for IPv4 on Junos OS Release 9.3 and later, and for IPv6 on Junos OS Release 10.4 and later. Starting with Release 11.1, Junos OS also supports nonstop active routing for PIM on devices that have both IPv4 and IPv6 configured on them.

To configure nonstop active routing for PIM, include the same statements in the configuration as for other protocols: the **nonstop-routing** statement at the **[edit routing-options]** hierarchy level and the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level. To trace PIM nonstop active routing events, include the **flag nsr-synchronization** statement at the **[edit protocols pim traceoptions]** hierarchy level.



**NOTE:** The **clear pim join**, **clear pim register**, and **clear pim statistics** operational mode commands are not supported on the backup Routing Engine when nonstop active routing is enabled.

Nonstop active routing support varies for different PIM features. The features fall into the following three categories: supported features, unsupported features, and incompatible features.

#### Supported features:

- Auto-RP



**NOTE:** Nonstop active routing PIM support on IPv6 does not support auto-RP because IPv6 does not support auto-RP.

- Bootstrap router (BSR)
- Static RPs
- Embedded RP on non-RP IPv6 routers
- Local RP



**NOTE:** RP set information synchronization is supported for local RP and BSR (on IPv4 and IPv6), autoRP (on IPv4), and embedded RP (on IPv6).

- BFD
- Dense mode
- Sparse mode
- Source-specific multicast (SSM)
- Draft Rosen multicast VPNs (MVPNs)

- Anycast RP (anycast RP set information synchronization and anycast RP register state synchronization on IPv4 and IPv6 configurations)
- Flow maps
- Unified ISSU
- Policy features such as neighbor policy, bootstrap router export and import policies, scope policy, flow maps, and reverse path forwarding (RPF) check policies
- Upstream assert synchronization
- PIM join load balancing

Starting with Release 12.2, Junos OS extends the nonstop active routing PIM support to draft Rosen MVPNs. Nonstop active routing PIM support for draft Rosen MVPNs enables nonstop active routing-enabled devices to preserve draft Rosen MPVN-related information—such as default and data multicast distribution tree (MDT) states—across switchovers. In releases earlier than 12.2, nonstop active routing PIM configuration was incompatible with draft Rosen MPVN configuration.

The backup Routing Engine sets up the default MDT based on the configuration and the information it receives from the master Routing Engine, and keeps updating the default MDT state information.

However, for data MDTs, the backup Routing Engine relies on the master Routing Engine to provide updates when data MDTs are created, updated, or deleted. The backup Routing Engine neither monitors data MDT flow rates nor triggers a data MDT switchover based on variations in flow rates. Similarly, the backup Routing Engine does not maintain the data MDT delay timer or timeout timer. It does not send MDT join TLV packets for the data MDTs until it takes over as the master Routing Engine. After the switchover, the new master Routing Engine starts sending MDT join TLV packets for each data MDT, and also resets the data MDT timers. Note that the expiration time for the timers might vary from the original values on the previous master Routing Engine.

Starting with Release 12.3, Junos OS extends the Protocol Independent Multicast (PIM) nonstop active routing support to IGMP-only interfaces.

In Junos OS releases earlier than 12.3, the PIM joins created on IGMP-only interfaces were not replicated on the backup Routing Engine. Thus, the corresponding multicast routes were marked as pruned (meaning discarded) on the backup Routing Engine. Because of this limitation, after a switchover, the new master Routing Engine had to wait for the IGMP module to come up and start receiving reports to create PIM joins and to install multicast routes. This caused traffic loss until the multicast joins and routes were reinstated.

However, in Junos OS Release 12.3 and later, the multicast joins on the IGMP-only interfaces are mapped to PIM states, and these states are replicated on the backup Routing Engine. If the corresponding PIM states are available on the backup, the multicast routes are marked as forwarding on the backup Routing Engine. This enables uninterrupted traffic flow after a switchover. This enhancement covers IGMPv2, IGMPv3, MLDv1, and MLDv2 reports and leaves.



**Unsupported features:** You can configure the following PIM features on a router along with nonstop active routing, but they function as if nonstop active routing is not enabled. In other words, during Routing Engine switchover and other outages, their state information is not preserved, and traffic loss is to be expected.

- Internet Group Management Protocol (IGMP) exclude mode
- IGMP snooping

Nonstop active routing is not supported for next-generation MVPNs with PIM provider tunnels. The commit operation fails if the configuration includes both nonstop active routing and next-generation MVPNs with PIM provider tunnels.

Junos OS provides a configuration statement that disables nonstop active routing for PIM only, so that you can activate incompatible PIM features and continue to use nonstop active routing for the other protocols on the router. Before activating an incompatible PIM feature, include the **nonstop-routing disable** statement at the **[edit protocols pim]** hierarchy level. Note that in this case, nonstop active routing is disabled for all PIM features, not just incompatible features.

## Nonstop Active Routing MSDP Support

Starting with Release 12.1, Junos OS extends nonstop active routing support to the Multicast Source Discovery Protocol (MSDP).

Nonstop active routing support for MSDP preserves the following MSDP-related information across the switchover:

- MSDP configuration and peer information
- MSDP peer socket information
- Source-active and related information

However, note that the following restrictions or limitations apply to nonstop active routing MSDP support:

- Because the backup Routing Engine learns the active source information by processing the source-active messages from the network, synchronizing of source active information between the master and backup Routing Engines might take up to 60 seconds. So, no planned switchover is allowed within 60 seconds of the initial replication of the sockets.
- Similarly, Junos OS does not support two planned switchovers within 240 seconds of each other.

Junos OS enables you to trace MSDP nonstop active routing events by including the **flag nsr-synchronization** statement at the **[edit protocols msdp traceoptions]** hierarchy level.

## Nonstop Active Routing Support for RSVP-TE LSPs

Junos OS extends nonstop active routing support to label-switching routers (LSRs) and Layer 2 Circuits that are part of an RSVP-TE LSP. Nonstop active routing support on LSRs ensures that the master to backup Routing Engine switchover on an LSR remains

transparent to the network neighbors and that the LSP information remains unaltered during and after the switchover.

You can use the **show rsvp version** command to view the nonstop active routing mode and state on an LSR. Similarly, you can use the **show mpls lsp** and **show rsvp session** commands on the standby Routing Engine to view the state recreated on the standby Routing Engine.

The Junos OS nonstop active routing feature is also supported on RSVP point-to-multipoint LSPs. Nonstop active routing support for RSVP point-to-multipoint egress and transit LSPs was added in Junos OS Release 11.4, and for ingress LSPs in Release 12.1. During the switchover, the LSP comes up on the backup Routing Engine that shares and synchronizes the state information with the master Routing Engine before and after the switchover. Nonstop active routing support for point-to-multipoint transit and egress LSPs ensures that the switchover remains transparent to the network neighbors, and preserves the LSP information across the switchover.

However, Junos OS nonstop active routing support for RSVP point-to-multipoint LSPs does not include support for dynamically created point-to-multipoint LSPs, such as VPLS.

Starting with Release 14.1, Junos OS extends nonstop active routing support to the next-generation multicast VPNs.

The **show rsvp session detail** command enables you to check the point-to-multipoint LSP remerge state information (**P2MP LSP re-merge**; possible values are **head**, **member**, and **none**).

However, Junos OS does not support nonstop active routing for the following features:

- Generalized Multiprotocol Label Switching (GMPLS) and LSP hierarchy
- Interdomain or loose-hop expansion LSPs
- BFD liveness detection

Nonstop active routing support for RSVP-TE LSPs is subject to the following limitations and restrictions:

- Detour LSPs are not maintained across a switchover and so, detour LSPs might fail to come back online after the switchover.
- Control plane statistics corresponding to the **show rsvp statistics** and **show rsvp interface detail | extensive** commands are not maintained across Routing Engine switchovers.
- Statistics from the backup Routing Engine are not reported for **show mpls lsp statistics** and **monitor mpls label-switched-path** commands. However, if a switchover occurs, the backup Routing Engine, after taking over as the master, starts reporting statistics. Note that the **clear statistics** command issued on the old master Routing Engine does not have any effect on the new master Routing Engine, which reports statistics, including any uncleared statistics.
- State timeouts might take additional time during nonstop active routing switchover. For example, if a switchover occurs after a neighbor has missed sending two hello

messages to the master, the new master Routing Engine waits for another three hello periods before timing out the neighbor.

- On the RSVP ingress router, if you configure auto-bandwidth functionality, the bandwidth adjustment timers are set in the new master after the switchover. This causes a one-time increase in the length of time required for the bandwidth adjustment after the switchover occurs.
- RSVP ingress LSPs that have BFD liveness detection enabled on them do not come up on the backup Routing Engine during the switchover. Such BFD-enabled LSPs have to be reestablished after the switchover.
- Backup LSPs —LSPs that are established between the point of local repair (PLR) and the merge point after a node or link failure—are not preserved during a Routing Engine switchover.
- When nonstop active routing is enabled, graceful restart is not supported. However, graceful restart helper mode is supported.

#### Related Documentation

- [Nonstop Active Routing Concepts on page 2597](#)
- [Configuring Nonstop Active Routing](#)
- [Configuring Nonstop Active Routing on Switches](#)
- [Example: Configuring Nonstop Active Routing on Switches on page 2611](#)

## Example: Configuring Nonstop Active Routing on Switches

Nonstop active routing (NSR) provides high availability for Routing Engines by enabling transparent switchover of the Routing Engines without necessitating restart of supported routing protocols. Both Routing Engines are fully active in processing protocol sessions, and so each can take over for the other. The switchover is transparent to neighbors.

This example describes how to configure nonstop active routing on switches with multiple Routing Engines or on an EX Series or a QFX series switch in a Virtual Chassis or Virtual Chassis Fabric configuration.

- [Requirements on page 2611](#)
- [Overview and Topology on page 2612](#)
- [Configuration on page 2612](#)
- [Verification on page 2613](#)
- [Troubleshooting on page 2613](#)

### Requirements

This example uses the following hardware and software components:

- An EX Series with multiple Routing Engines or on an EX Series or a QFX series switch in a Virtual Chassis or Virtual Chassis Fabric configuration
- Junos OS Release 10.4 or later for EX Series switches

- Junos OS Release 13.2X51-D20 or later for QFX Series switches

## Overview and Topology

Configure nonstop active routing on any EX Series with multiple Routing Engines or on an EX Series or a QFX series switch in a Virtual Chassis or Virtual Chassis Fabric configuration. Nonstop active routing is advantageous in networks where neighbor routing devices do not support graceful restart protocol extensions.

The topology used in this example consists of an EX8200 switch with redundant Routing Engines connected to neighbor routing devices that are not configured to support graceful restart of protocols.

## Configuration

### CLI Quick Configuration

To quickly configure nonstop active routing, copy the following commands and paste them into the switch terminal window:

```
[edit]
set chassis redundancy graceful-switchover
set routing-options nonstop-routing
set system commit synchronize
```

### Step-by-Step Procedure

To configure nonstop active routing on a switch:

1. Enable graceful Routing Engine switchover (GRES):  

```
[edit chassis redundancy]
user@switch# set graceful-switchover
```
2. Enable nonstop active routing (by default, nonstop active routing is disabled):  

```
[edit routing-options]
user@switch# set nonstop-routing
```
3. Synchronize configuration changes between the Routing Engines:  

```
[edit system]
user@switch# set commit synchronize
```



**NOTE:** If the backup Routing Engine is down when you issue the commit, a warning is displayed and the candidate configuration is committed in the master Routing Engine. When the backup Routing Engine comes up, its configuration is automatically synchronized with that of the master. If you subsequently insert or bring up a backup Routing Engine, it automatically synchronizes its configuration with the master Routing Engine configuration.

---

## Results

Check the results of the configuration:

```
[edit]
user@switch# show
chassis {
 redundancy {
 graceful-switchover;
 }
}
routing-options {
 nonstop-routing;
}
system {
 commit synchronize;
}
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verifying That Nonstop Active Routing Is Working Correctly on the Switch on page 2613](#)

Verifying That Nonstop Active Routing Is Working Correctly on the Switch

**Purpose** Verify that nonstop active routing is enabled.

**Action** Issue the [show task replication](#) command:

```
user@switch# show task replication
Stateful Replication: Enabled
RE mode: Master

Protocol Synchronization Status

OSPF Complete
RIP Complete
PIM Complete
RSVP Complete
```

**Meaning** This output shows that nonstop active routing (Stateful Replication) is enabled on master routing engine. If nonstop routing is not enabled, instead of the output shown above:

- On the backup routing engine the following error message is displayed: “**error: the routing subsystem is not running.**”
- On the master routing engine, the following output is displayed if nonstop routing is not enabled:

```
Stateful Replication: Disabled
RE mode: Master
```

Troubleshooting

To troubleshoot nonstop active routing, perform these tasks:

- [Investigating Problems with Synchronization of Routing Engines When NSR Is Enabled on page 2614](#)

### Investigating Problems with Synchronization of Routing Engines When NSR Is Enabled

---

**Problem** A protocol loses connectivity with neighbors after a graceful Routing Engine switchover (GRES) occurs with nonstop active routing (NSR) enabled.

**Solution** Use trace options to help isolate the problem and gather troubleshooting information. Using the information gathered from trace options, you can confirm or eliminate the synchronization of the Routing Engines as the cause of the loss of connectivity for the protocol. See *Tracing Nonstop Active Routing Synchronization Events*.

**Related Documentation**

- *Configuring Nonstop Active Routing on Switches*
- *Tracing Nonstop Active Routing Synchronization Events*
- *Understanding Nonstop Active Routing on EX Series Switches*
- [Nonstop Active Routing Concepts on page 2597](#)

## PART 44

# Graceful Routing Engine Switchover

- [Configuring Graceful Routing Engine Switchover on page 2617](#)





# Configuring Graceful Routing Engine Switchover

- [Understanding Graceful Routing Engine Switchover on page 2617](#)
- [Graceful Routing Engine Switchover System Requirements on page 2623](#)
- [Configuring Graceful Routing Engine Switchover on page 2627](#)
- [Configuring Graceful Routing Engine Switchover in a Virtual Chassis \(CLI Procedure\) on page 2629](#)
- [Resetting Local Statistics on page 2630](#)

## Understanding Graceful Routing Engine Switchover

---

This topic contains the following sections:

- [Graceful Routing Engine Switchover Concepts on page 2617](#)
- [Effects of a Routing Engine Switchover on page 2622](#)

### Graceful Routing Engine Switchover Concepts

The graceful Routing Engine switchover (GRES) feature in Junos OS enables a routing platform with redundant Routing Engines to continue forwarding packets, even if one Routing Engine fails. GRES preserves interface and kernel information. Traffic is not interrupted. However, GRES does not preserve the control plane.



**NOTE:** On T Series routers, TX Matrix routers, and TX Matrix Plus routers, the control plane is preserved in case of GRES with nonstop active routing (NSR), and nearly 75 percent of line rate worth of traffic per Packet Forwarding Engine remains uninterrupted during GRES.

Neighboring routers detect that the router has experienced a restart and react to the event in a manner prescribed by individual routing protocol specifications.

To preserve routing during a switchover, GRES must be combined with either:

- Graceful restart protocol extensions
- Nonstop active routing

Any updates to the master Routing Engine are replicated to the backup Routing Engine as soon as they occur.



**NOTE:** Due to its synchronization requirements and logic, NSR/GRES performance is limited by the slowest Routing Engine in the system.

Mastership switches to the backup Routing Engine if:

- The master Routing Engine kernel stops operating.
- The master Routing Engine experiences a hardware failure.
- The administrator initiates a manual switchover.



**NOTE:** To quickly restore or to preserve routing protocol state information during a switchover, GRES must be combined with either graceful restart or nonstop active routing, respectively. For more information about graceful restart, see [“Graceful Restart Concepts” on page 2579](#). For more information about nonstop active routing, see [“Nonstop Active Routing Concepts” on page 2597](#).

If the backup Routing Engine does not receive a keepalive from the master Routing Engine after 2 seconds (4 seconds on M20 routers), it determines that the master Routing Engine has failed and: takes mastership.

The Packet Forwarding Engine:

- Seamlessly disconnects from the old master Routing Engine
- Reconnects to the new master Routing Engine
- Does not reboot
- Does not interrupt traffic

The new master Routing Engine and the Packet Forwarding Engine then become synchronized. If the new master Routing Engine detects that the Packet Forwarding Engine state is not up to date, it resends state update messages.



**NOTE:** If adjacencies between the restarting router and the neighboring peer 'helper' routers time out, graceful restart protocol extensions are unable to notify the peer 'helper' routers about the impending restart. Graceful restart can then stop and cause interruptions in traffic.

To ensure that these adjacencies are kept, change the [hold-time](#) for IS-IS protocols from the default of 27 seconds to a value higher than 40 seconds.



**NOTE:** Successive Routing Engine switchover events must be a minimum of 240 seconds (4 minutes) apart after both Routing Engines have come up.

If the router or switch displays a warning message similar to **Standby Routing Engine is not ready for graceful switchover. Packet Forwarding Engines that are not ready for graceful switchover might be reset**. Do not attempt switchover. If you choose to proceed with switchover, only the Packet Forwarding Engines that were not ready for graceful switchover are reset. None of the FPCs should spontaneously restart. We recommend that you wait until the warning no longer appears and then proceed with the switchover.



**NOTE:** Starting from Junos OS Release 14.2, when you perform GRES on MX Series routers, you must execute the **clear synchronous-ethernet wait-to-restore** operational mode command on the new master Routing Engine to clear the wait-to-restore timer on it. This is because the **clear synchronous-ethernet wait-to-restore** operational mode command clears the wait-to-restore timer only on the local Routing Engine.



**NOTE:** In a routing matrix with TX Matrix Plus router with 3D SIBs, for successive Routing Engine switchover, events must be a minimum of 900 seconds (15 minutes) apart after both Routing Engines have come up.

GRES must be performed on one line-card chassis (LCC) (of a TX Matrix router with 3D SIBs) at a time to avoid synchronization issues.




**NOTE:**

- We do *not* recommend performing a commit operation on the backup Routing Engine when GRES is enabled on the router or switch.
- We do *not* recommend enabling GRES on the backup Routing Engine in *any* scenario.

Figure 41 shows the system architecture of graceful Routing Engine switchover and the process a routing platform follows to prepare for a switchover.

Figure 41: Preparing for a Graceful Routing Engine Switchover



**NOTE:** Check GRES readiness by executing both:

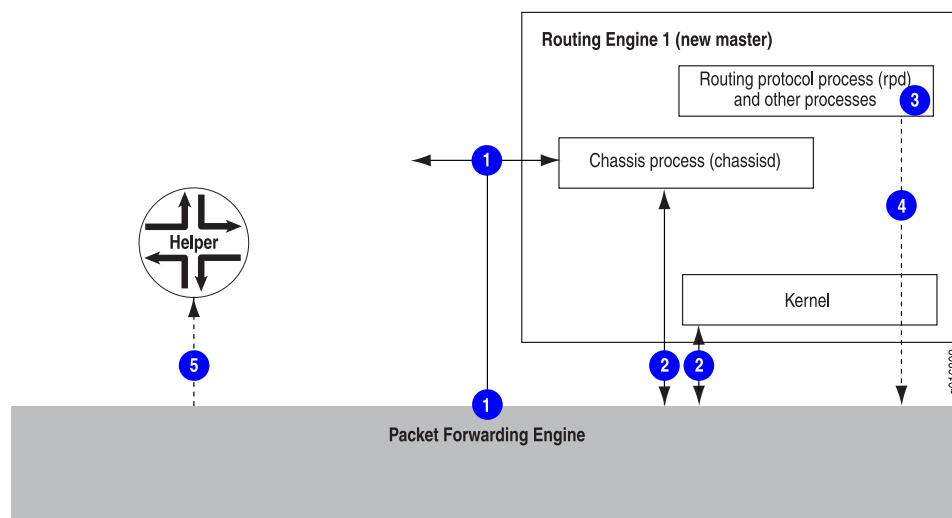
- The `request chassis routing-engine master switch check` command from the master Routing Engine
- The `show system switchover` command from the Backup Routing Engine

The switchover preparation process for GRES is as follows:

1. The master Routing Engine starts.
2. The routing platform processes (such as the chassis process [chassisd]) start.
3. The Packet Forwarding Engine starts and connects to the master Routing Engine.
4. All state information is updated in the system.
5. The backup Routing Engine starts.
6. The system determines whether GRES has been enabled.
7. The kernel synchronization process (ksyncd) synchronizes the backup Routing Engine with the master Routing Engine.
8. After ksyncd completes the synchronization, all state information and the forwarding table are updated.

Figure 42 shows the effects of a switchover on the routing (or switching) platform.

Figure 42: Graceful Routing Engine Switchover Process



When a switchover occurs, the switchover process is as follows:

1. When keepalives from the master Routing Engine are lost, the system switches over gracefully to the backup Routing Engine.
2. The Packet Forwarding Engine connects to the backup Routing Engine, which becomes the new master.
3. Routing platform processes that are not part of GRES (such as the routing protocol process [rpd]) restart.
4. State information learned from the point of the switchover is updated in the system.
5. If configured, graceful restart protocol extensions collect and restore routing information from neighboring peer *helper* routers.



**NOTE:** On T Series and M320 routers during GRES, the Switch Interface Boards (SIBs) are taken offline and restarted one by one. This is done to provide the Switch Processor Mezzanine Board (SPMB) that manages the SIB enough time to populate state information for its associated SIB. However, on a fully populated chassis where all FPCs are sending traffic at full line rate, there might be momentary packet loss during the switchover.



**NOTE:** When GRES is configured and the `restart chassis-control` command is executed on a TX Matrix Plus router with 3D SIBs, we cannot ascertain which Routing Engine becomes the master. This is because the `chassisd` process restarts with the execution of the `restart chassis-control` command. The `chassisd` process is responsible for maintaining and retaining mastership and when it is restarted, the new `chassisd` is processed based on the router or switch load. As a result, any one of the Routing Engines is made the master.

## Effects of a Routing Engine Switchover

Table 234 describes the effects of a Routing Engine switchover when different features are enabled:

- No high availability features
- Graceful Routing Engine switchover
- Graceful restart
- Nonstop active routing

**Table 234: Effects of a Routing Engine Switchover**

| Feature                                         | Benefits                                                                                                                                                                                                               | Considerations                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Dual Routing Engines only (no features enabled) | <ul style="list-style-type: none"> <li>• When the switchover to the new master Routing Engine is complete, routing convergence takes place and traffic is resumed.</li> </ul>                                          | <ul style="list-style-type: none"> <li>• All physical interfaces are taken offline.</li> <li>• Packet Forwarding Engines restart.</li> <li>• The standby Routing Engine restarts the routing protocol process (rpd).</li> <li>• All hardware and interfaces are discovered by the new master Routing Engine.</li> <li>• The switchover takes several minutes.</li> <li>• All of the router's adjacencies are aware of the physical (interface alarms) and routing (topology) changes.</li> </ul> |
| GRES enabled                                    | <ul style="list-style-type: none"> <li>• During the switchover, interface and kernel information is preserved.</li> <li>• The switchover is faster because the Packet Forwarding Engines are not restarted.</li> </ul> | <ul style="list-style-type: none"> <li>• The new master Routing Engine restarts the routing protocol process (rpd).</li> <li>• All hardware and interfaces are acquired by a process that is similar to a warm restart.</li> <li>• All adjacencies are aware of the router's change in state.</li> </ul>                                                                                                                                                                                         |
| GRES <i>and</i> nonstop active routing enabled  | <ul style="list-style-type: none"> <li>• Traffic is not interrupted during the switchover.</li> <li>• Interface and kernel information are preserved.</li> </ul>                                                       | <ul style="list-style-type: none"> <li>• Unsupported protocols must be refreshed using the normal recovery mechanisms inherent in each protocol.</li> </ul>                                                                                                                                                                                                                                                                                                                                      |

Table 234: Effects of a Routing Engine Switchover (*continued*)

| Feature                           | Benefits                                                                                                                                                                                                                                                                                 | Considerations                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| GRES and graceful restart enabled | <ul style="list-style-type: none"> <li>Traffic is not interrupted during the switchover.</li> <li>Interface and kernel information are preserved.</li> <li>Graceful restart protocol extensions quickly collect and restore routing information from the neighboring routers.</li> </ul> | <ul style="list-style-type: none"> <li>Neighbors are required to support graceful restart, and a wait interval is required.</li> <li>The routing protocol process (rpd) restarts.</li> <li>For certain protocols, a significant change in the network can cause graceful restart to stop.</li> <li>If adjacencies between the restarting router and the neighboring peer 'helper' routers time out, graceful restart can stop and cause interruptions in traffic.</li> </ul> |

**Related Documentation**

- [Understanding High Availability Features on Juniper Networks Routers](#)
- [Graceful Routing Engine Switchover System Requirements on page 2623](#)
- [Configuring Graceful Routing Engine Switchover on page 2627](#)
- [Configuring Graceful Routing Engine Switchover in a Virtual Chassis \(CLI Procedure\) on page 2629](#)
- [Requirements for Routers with a Backup Router Configuration](#)
- [Example: Configuring IS-IS for GRES with Graceful Restart](#)
- [hold-time on page 4282](#)

## Graceful Routing Engine Switchover System Requirements

Graceful Routing Engine switchover is supported on all routing (or switching) platforms that contain dual Routing Engines. All Routing Engines configured for graceful Routing Engine switchover must run the same Junos OS release. Hardware and software support for graceful Routing Engine switchover is described in the following sections:

- [Graceful Routing Engine Switchover Platform Support on page 2623](#)
- [Graceful Routing Engine Switchover Feature Support on page 2624](#)
- [Graceful Routing Engine Switchover DPC Support on page 2626](#)
- [Graceful Routing Engine Switchover and Subscriber Access on page 2626](#)
- [Graceful Routing Engine Switchover PIC Support on page 2626](#)

## Graceful Routing Engine Switchover Platform Support

To enable graceful Routing Engine switchover, your system must meet these minimum requirements:

- M20 and M40e routers—Junos OS Release 5.7 or later
- M10i router—Junos OS Release 6.1 or later
- M320 router—Junos OS Release 6.2 or later
- T320 router, T640 router, and TX Matrix router—Junos OS Release 7.0 or later
- M120 router—Junos OS Release 8.2 or later
- MX960 router—Junos OS Release 8.3 or later
- MX480 router—Junos OS Release 8.4 or later (8.4R2 recommended)
- MX240 router—Junos OS Release 9.0 or later
- PTX5000 router—Junos OS Release 12.1X48 or later
- Standalone T1600 router—Junos OS Release 8.5 or later
- Standalone T4000 router—Junos OS Release 12.1R2 or later
- TX Matrix Plus router—Junos OS Release 9.6 or later
- TX Matrix Plus router with 3D SIBs—Junos Release 13.1 or later
- EX Series switches with dual Routing Engines or in a Virtual Chassis — Junos OS Release 9.2 or later for EX Series switches
- QFX Series switches in a Virtual Chassis —Junos OS Release 13.2 or later for the QFX Series
- EX Series or QFX Series switches in a Virtual Chassis Fabric —Junos OS Release 13.2X51-D20 or later for the EX Series and QFX Series switches

For more information about support for graceful Routing Engine switchover, see the sections that follow.

## Graceful Routing Engine Switchover Feature Support

Graceful Routing Engine switchover supports most Junos OS features in Release 5.7 and later. Particular Junos OS features require specific versions of Junos OS. See [Table 235](#).

**Table 235: Graceful Routing Engine Switchover Feature Support**

| Application                                                                                                  | Junos OS Release               |
|--------------------------------------------------------------------------------------------------------------|--------------------------------|
| Aggregated Ethernet interfaces with Link Aggregation Control Protocol (LACP) and aggregated SONET interfaces | 6.2                            |
| Asynchronous Transfer Mode (ATM) virtual circuits (VCs)                                                      | 6.2                            |
| Logical systems                                                                                              | 6.3                            |
| <b>NOTE:</b> In Junos OS Release 9.3 and later, the logical router feature is renamed to logical system.     |                                |
| Multicast                                                                                                    | 6.4 (7.0 for TX Matrix router) |



**Table 235: Graceful Routing Engine Switchover Feature Support (*continued*)**

| Application                                                                                                                                                                                            | Junos OS Release |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| Multilink Point-to-Point Protocol (MLPPP) and Multilink Frame Relay (MLFR)                                                                                                                             | 7.0              |
| Automatic Protection Switching (APS)—The current active interface (either the designated working or the designated protect interface) remains the active interface during a Routing Engine switchover. | 7.4              |
| Point-to-multipoint Multiprotocol Label Switching MPLS LSPs (transit only)                                                                                                                             | 7.4              |
| Compressed Real-Time Transport Protocol (CRTP)                                                                                                                                                         | 7.6              |
| Virtual private LAN service (VPLS)                                                                                                                                                                     | 8.2              |
| Ethernet Operation, Administration, and Management (OAM) as defined by IEEE 802.3ah                                                                                                                    | 8.5              |
| Extended DHCP relay agent                                                                                                                                                                              | 8.5              |
| Ethernet OAM as defined by IEEE 802.1ag                                                                                                                                                                | 9.0              |
| Packet Gateway Control Protocol (PGCP) process (pgcpd) on Multiservices 500 PICs on T640 routers.                                                                                                      | 9.0              |
| Subscriber access                                                                                                                                                                                      | 9.4              |
| Layer 2 Circuit and LDP-based VPLS pseudowire redundant configuration                                                                                                                                  | 9.6              |

The following constraints apply to graceful Routing Engine switchover feature support:

- When graceful Routing Engine switchover and aggregated Ethernet interfaces are configured in the same system, the aggregated Ethernet interfaces must not be configured for fast-polling LACP. When fast polling is configured, the LACP polls time out at the remote end during the Routing Engine mastership switchover. When LACP polling times out, the aggregated link and interface are disabled. The Routing Engine mastership change is fast enough that standard and slow LACP polling do not time out during the procedure. However, note that this restriction does not apply to MX Series Routers that are running Junos OS Release 9.4 or later and have distributed periodic packet management (PPM) enabled—which is the default configuration—on them. In such cases, you can configure graceful Routing Engine switchover and have aggregated Ethernet interfaces configured for fast-polling LACP on the same device.
- When a graceful Routing Engine switchover occurs, the VRRP state does not change. VRRP is supported by graceful Routing Engine switchover only in the case that PPM delegation is enabled (which the default).

## Graceful Routing Engine Switchover DPC Support

Graceful Routing Engine switchover supports all Dense Port Concentrators (DPCs) on the MX Series 3D Universal Edge Routers running the appropriate version of Junos OS as shown in “[Graceful Routing Engine Switchover Platform Support](#)” on page 2623. For more information about DPCs, see the *MX Series DPC Guide*.

## Graceful Routing Engine Switchover and Subscriber Access

Graceful Routing Engine switchover currently supports most of the features directly associated with dynamic DHCP and dynamic PPPoE subscriber access. Graceful Routing Engine switchover also supports the unified in-service software upgrade (ISSU) for the DHCP access model and the PPPoE access model used by subscriber access.

## Graceful Routing Engine Switchover PIC Support

Graceful Routing Engine switchover is supported on most PICs, except for the services PICs listed in this section. The PIC must be on a supported routing platform running the appropriate version of Junos OS. For information about FPC types, FPC/PIC compatibility, and the initial Junos OS Release in which an FPC supported a particular PIC, see the PIC guide for your router platform.

The following constraints apply to graceful Routing Engine switchover support for services PICs:

- You can include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level on a router with Adaptive Services, Multiservices, and Tunnel Services PICs configured on it and successfully commit the configuration. However, all services on these PICs—except the Layer 2 service packages and extension-provider and SDK applications on Multiservices PICs—are reset during a switchover.
- Graceful Routing Engine switchover is not supported on any Monitoring Services PICs or Multilink Services PICs. If you include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level on a router with either of these PIC types configured on it and issue the **commit** command, the commit fails.
- Graceful Routing Engine switchover is not supported on Multiservices 400 PICs configured for monitoring services applications. If you include the **graceful-switchover** statement, the commit fails.



**NOTE:** When an unsupported PIC is online, you cannot enable graceful Routing Engine switchover. If graceful Routing Engine switchover is already enabled, an unsupported PIC cannot come online.

---

### Related Documentation

- [Understanding High Availability Features on Juniper Networks Routers](#)
- [Understanding Graceful Routing Engine Switchover on page 2617](#)
- [Configuring Graceful Routing Engine Switchover on page 2627](#)

- [Configuring Graceful Routing Engine Switchover in a Virtual Chassis \(CLI Procedure\) on page 2629](#)
- *Requirements for Routers with a Backup Router Configuration*

## Configuring Graceful Routing Engine Switchover

---

This section contains the following topics:

- [Enabling Graceful Routing Engine Switchover on page 2627](#)
- [Configuring Graceful Routing Engine Switchover with Graceful Restart on page 2627](#)
- [Synchronizing the Routing Engine Configuration on page 2627](#)
- [Verifying Graceful Routing Engine Switchover Operation on page 2629](#)

### Enabling Graceful Routing Engine Switchover

By default, graceful Routing Engine switchover (GRES) is disabled. To configure GRES, include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level.

```
[edit chassis redundancy]
graceful-switchover;
```

When you enable GRES, the command-line interface (CLI) indicates which Routing Engine you are using. For example:

```
{master} [edit]
user@host#
```

To disable GRES, delete the **graceful-switchover** statement from the **[edit chassis redundancy]** hierarchy level.

### Configuring Graceful Routing Engine Switchover with Graceful Restart

When using GRES with Graceful Restart, if adjacencies between the Routing Engine and the neighboring peer 'helper' routers time out, graceful restart protocol extensions are unable to notify the peer 'helper' routers about the impending restart. Graceful restart can then stop and cause interruptions in traffic.

To ensure that these adjacencies are kept, change the **hold-time** for IS-IS protocols from the default of 27 seconds to a value higher than 40 seconds.

### Synchronizing the Routing Engine Configuration



**NOTE:** A newly inserted backup Routing Engine automatically synchronizes its configuration with the master Routing Engine configuration.

When you configure GRES, you can bring the backup Routing Engine online after the master Routing Engine is already running. There is no requirement to start the two Routing Engines simultaneously.

Only when you enable the graceful Routing Engine switchover, you can copy the running Junos OS version of the master Routing Engine to the backup Routing Engine.



**NOTE:** If the system is in ISSU state, you cannot copy the running Junos OS version of the master Router Engine.

You can enable automatic synchronization of the master Routing Engine configuration with the backup Routing Engine by including the events CHASSISD\_SNMP\_TRAP7 statement at the [edit event-options policy *policy-name*] hierarchy level.

CHASSISD\_SNMP\_TRAP7 is a system event logging message that the chassis process (chassisd) generates a Simple Network Management Protocol (SNMP) trap with the seven indicated argument-value pairs. An example of an event script to trigger automatic synchronization of master to the backup Routing Engine is as follows:

```
[edit event-options]
policy UPGRADE-BACKUPRE {
 events CHASSISD_SNMP_TRAP7;
 attributes-match {
 CHASSISD_SNMP_TRAP7.value5 matches "Routing Engine";
 CHASSISD_SNMP_TRAP7.trap matches "Fru Online";
 CHASSISD_SNMP_TRAP7.argument5 matches jnxFruName;
 }
 then {
 event-script auto-image-upgrade.slax {
 arguments {
 trap "${$.trap}";
 value5 "${$.value5}";
 argument5 "${$.argument5}";
 }
 }
 }
 event-script {
 file auto-image-upgrade.slax;
 }
}
```

After receiving this event, the event policy on the master Router Engine is triggered and the image available in the */var/sw/pkg* path is pushed to the backup Router Engine upgrade. During script execution, the image is copied to the backup Routing Engine's */var/sw/pkg* path.



**NOTE:** If the image is not available in the */var/sw/pkg* path, the script is terminated with an appropriate syslog message.

If the Routing Engine is running at the Junos OS Release 13.2 or later, the Junos automation scripts is synchronized automatically.

After the master Router Engine is rebooted, the event script available at the */usr/libexec/scripts/event/auto-image-upgrade.slax* must be copied to the */var/db/scripts/event* path.

## Verifying Graceful Routing Engine Switchover Operation

To verify whether GRES is enabled on the backup Routing Engine, issue the **show system switchover** command. When the output of the command indicates that the **Graceful switchover** field is set to **On**, GRES is operational. The status of the kernel database and configuration database synchronization between Routing Engines is also provided. For example:

```
Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady state
```



**NOTE:** You must issue the **show system switchover** command on the backup Routing Engine. This command is not supported on the master Routing Engine.

For more information about the **show system switchover** command, see the [CLI Explorer](#).

### Related Documentation

- [Understanding Graceful Routing Engine Switchover on page 2617](#)
- [Graceful Routing Engine Switchover System Requirements on page 2623](#)
- [Requirements for Routers with a Backup Router Configuration](#)
- [Resetting Local Statistics on page 2630](#)
- [graceful-switchover](#)
- [graceful-switchover on page 2673](#)
- [Example: Configuring IS-IS for GRES with Graceful Restart](#)
- [hold-time on page 4282](#)

## Configuring Graceful Routing Engine Switchover in a Virtual Chassis (CLI Procedure)

In a Virtual Chassis, one member switch is assigned the master role and has the master Routing Engine. Another member switch is assigned the backup role and has the backup Routing Engine. Graceful Routing Engine switchover (GRES) enables the master and backup Routing Engines in a Virtual Chassis configuration to switch from the master to backup without interruption to packet forwarding. When you configure graceful Routing Engine switchover, the backup Routing Engine automatically synchronizes with the master Routing Engine to preserve kernel state information and the forwarding state.

To set up the Virtual Chassis configuration to use graceful Routing Engine switchover (GRES):

1. Set up a minimum of two switches in a Virtual Chassis configuration with mastership priority of 255:

```
[edit]
user@switch# set virtual-chassis member 0 mastership-priority 255
[edit]
```

```
user@switch# set virtual-chassis member 1 mastership-priority 255
```

2. Set up graceful Routing Engine switchover:

```
[edit]
```

```
user@switch# set chassis redundancy graceful-switchover
```

Commit the configuration.



**NOTE:** We recommend that you use the `commit synchronize` command to save any configuration changes that you make to a multimember Virtual Chassis.

**Related  
Documentation**

- *Example: Configuring an EX4200 Virtual Chassis with a Master and Backup in a Single Wiring Closet*
- *High Availability Features for EX Series Switches Overview*
- *Understanding EX Series Virtual Chassis Configuration*
- *Understanding QFX Series Virtual Chassis*

---

## Resetting Local Statistics

When you enable graceful Routing Engine switchover, the master Routing Engine configuration is copied and loaded to the backup Routing Engine. User files, accounting information, and trace options information are not replicated to the backup Routing Engine.

When a graceful Routing Engine switchover occurs, local statistics such as process statistics and networking statistics are displayed as a cumulative value from the time the process first came online. Because processes on the master Routing Engine can start at different times from the processes on the backup Routing Engine, the statistics on the two Routing Engines for the same process might differ. After a graceful Routing Engine switchover, we recommend that you issue the **clear interface statistics (*interface-name* | all)** command to reset the cumulative values for local statistics. Forwarding statistics are not affected by graceful Routing Engine switchover.

For information about how to use the **clear** command to clear statistics and protocol database information, see the [CLI Explorer](#).



**NOTE:** The `clear firewall` command cannot be used to clear the Routing Engine filter counters on a backup Routing Engine that is enabled for graceful Routing Engine switchover.

**Related  
Documentation**

- [Understanding Graceful Routing Engine Switchover on page 2617](#)
- [Configuring Graceful Routing Engine Switchover on page 2627](#)

## PART 45

# Virtual Router Redundancy Protocol

- [Configuring Virtual Router Redundancy Protocol on page 2633](#)





# Configuring Virtual Router Redundancy Protocol

- [Understanding VRRP on page 2633](#)
- [Example: Configuring VRRP for Load Sharing on page 2636](#)
- [Configuring Basic VRRP Support for QFX on page 2641](#)
- [Configuring VRRP Authentication \(IPv4 Only\) on page 2642](#)
- [Configuring the Startup Period for VRRP Operations on page 2643](#)
- [Configuring the Advertisement Interval for the VRRP Master on page 2643](#)
- [Configuring VRRP Preemption and Hold Time on page 2644](#)
- [Configuring a Route to Be Tracked on page 2646](#)
- [Configuring a Logical Interface to Be Tracked on page 2646](#)
- [Configuring a Backup to Accept Packets Destined for the Virtual IP Address on page 2648](#)
- [Configuring Passive ARP Learning for VRRP Backups on page 2649](#)
- [Configuring the Silent Period on page 2649](#)
- [Configuring Inheritance for a VRRP Group on page 2650](#)
- [Troubleshooting VRRP on page 2651](#)

## Understanding VRRP

---

Juniper Networks switches support the Virtual Router Redundancy Protocol (VRRP) and VRRPv3 (for IPv6). This topic covers:

- [Overview of VRRP on page 2633](#)
- [Sample VRRP Topology on page 2634](#)

## Overview of VRRP

Configuring end hosts on your network with static default routes minimizes configuration effort and complexity and reduces processing overhead on the end hosts. When hosts are configured with static routes, the failure of the default gateway normally results in a catastrophic event, isolating all hosts that are unable to detect available alternate paths to their gateway. Using Virtual Router Redundancy Protocol (VRRP) enables you to dynamically provide alternative gateways for end hosts if the primary gateway fails.

VRRP (defined in RFC 3768) provides dynamic failover of IP addresses from one router to another in the event of failure. You can implement VRRP to provide a highly available default path to a gateway without needing to configure dynamic routing or router discovery protocols on end hosts.

Switches configured with VRRP share a virtual IP address, which is the address you configure as the default route on the hosts. At any time, one of the switches is the VRRP master, meaning that it owns the virtual IP address and is the active default gateway. The other devices are backups. The switches dynamically assign master and backup roles based on priorities that you configure (**1 through 255**). If the master fails, the backup switch with the highest priority becomes the master within a few seconds. This is done without any interaction with the hosts.

In VRRP operation, the master sends advertisements to the backup switches at regular intervals. The default interval is 1 second. If the backup switches do not receive an advertisement for a set period, the backup with the highest priority takes over as master within a few seconds and begins forwarding packets. This is done without any interaction with the hosts.



**NOTE:** Priority 255 cannot be set for routed VLAN interfaces (RVIs).


---

You can configure two QFabric systems to participate in a VRRP configuration as if they were two standalone switches. One benefit of this configuration is if you use VMware's vMotion, virtual machines can transition between hosts connected to the QFabric systems without updating their default gateway information. For example, a virtual machine running on a host connected to a QFabric system in data center A can transition to a host connected to a QFabric system in data center B without needing to resolve a new gateway IP address and MAC address.

## Sample VRRP Topology

**Figure 43** illustrates a basic VRRP topology. In this example, switches A and B are running VRRP and share the virtual IP address 10.1.1.1. The default gateway for each of the clients is 10.1.1.1.

Figure 43: Basic VRRP Topology



The following illustrates basic VRRP behavior using [Figure 43](#) for reference:

1. When any of the servers wants to send traffic out of the LAN, it sends the traffic to the default gateway address of 10.1.1.1. This is a virtual IP address (VIP) owned by VRRP group 100. Because switch A is the master of the group, the VIP is associated with the “real” address 10.1.1.251 on switch A, and traffic from the servers is actually sent to this address. (Switch A is the master because it has been configured with a higher priority value.)
2. If there is a failure on switch A that prevents it from forwarding traffic to or from the servers—for example, if the interface connected to the LAN fails—switch B becomes the master and assumes ownership of the VIP. The servers continue to send traffic to the VIP, but because the VIP is now associated with the “real” address 10.1.1.252 on switch B (because of change of master), the traffic is sent to switch B instead of switch A.
3. If the problem that caused the failure on switch A is corrected, switch A becomes the master again and reasserts ownership of the VIP. In this case, the servers resume sending traffic to switch A.

Notice that no configuration changes are required on the servers for them to switch between sending traffic to switch A and switch B. When the VIP moves between 10.1.1.251 and 10.1.1.252, the change is detected by normal TCP-IP behavior and no configuration or intervention is required on the servers.

#### Related Documentation

- [Configuring Basic VRRP Support for QFX on page 2641](#)
- [Example: Configuring VRRP for Load Sharing on page 2636](#)

## Example: Configuring VRRP for Load Sharing

---

If you do not want to dedicate a switch to be a VRRP backup (and therefore leave it idle unless the master fails), you can create a load-sharing configuration in which each participating switch simultaneously acts as a master and a backup.

One reason to use a load-sharing (active-active) configuration is that you are more likely to actively monitor and maintain both switches and notice if a problem occurs on either of them. If you use a configuration in which one switch is only a backup (an active-backup configuration), you might be less likely to pay attention to the backup switch while it is idle. In the worst case, this could lead to the backup switch developing an undetected problem and not being able to perform adequately when a failover occurs.

- [Requirements on page 2636](#)
- [Overview and Topology on page 2636](#)
- [Configuring VRRP on Both Switches on page 2637](#)
- [Verification on page 2640](#)

### Requirements


This example uses the following hardware and software components:

- Two switches
- Junos OS Release 11.3 or later
- Static routing or a dynamic routing protocol enabled on both switches.

### Overview and Topology

This example uses two VRRP groups, each of which has its own virtual IP address. Devices on the LAN use one of these virtual IP addresses as their default gateway. If one of the switches fails, the other switch takes over for it. In the topology shown in [Figure 44](#), for example, Switch A is the master for VRRP group 100. If Switch A fails, Switch B takes over and forwards traffic that the end devices send to the default gateway address 10.1.1.1.

Figure 44: VRRP Load-Sharing Configuration



This example shows a simple configuration to illustrate the basic steps for configuring two switches running VRRP to back each other up. [Table 236](#) lists VRRP settings for each switch.

Table 236: Settings for VRRP Load-Sharing Example

| Switch A                                                                                                                                  | Switch B                                                                                                                                  |
|-------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| VRRP Group 100: <ul style="list-style-type: none"><li>Interface address: 10.1.1.251</li><li>VIP: 10.1.1.1</li><li>Priority: 250</li></ul> | VRRP Group 100: <ul style="list-style-type: none"><li>Interface address: 10.1.1.252</li><li>VIP: 10.1.1.1</li><li>Priority: 200</li></ul> |
| VRRP Group 200: <ul style="list-style-type: none"><li>Interface address: 10.1.1.251</li><li>VIP: 10.1.1.2</li><li>Priority: 200</li></ul> | VRRP Group 200: <ul style="list-style-type: none"><li>Interface address: 10.1.1.252</li><li>VIP: 10.1.1.2</li><li>Priority: 250</li></ul> |

In addition to configuring the two switches as shown, you must configure your end devices so that some of them use one of the virtual IP addresses as their default gateway and the remaining end devices use the other virtual IP address as their default gateway.

Note that if a failover occurs, the remaining switch might be unable to handle all of the traffic, depending on the demand.

Configuring VRRP on Both Switches

CLI Quick Configuration

Enter the following on Switch A:  
[edit]  
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group 100 virtual-address 10.1.1.1  
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group 100 priority 250

```
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group 200 virtual-address 10.1.1.2
```

```
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group 200 priority 200
```

Enter the following on Switch B:

```
[edit]
```

```
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group 100 virtual-address 10.1.1.1
```

```
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group 100 priority 200
```

```
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group 200 virtual-address 10.1.1.2
```

```
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group 200 priority 250
```

#### Step-by-Step Procedure

Configure the VRRP groups and priorities on Switch A:

1. Create VRRP group 100 on Switch A and configure the virtual IP address for the group:

```
[edit]
```

```
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group 100 virtual-address 10.1.1.1
```

2. Assign the VRRP priority for this interface in this group:

```
[edit]
```

```
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group 100 priority 250
```

3. Create VRRP group 200 on Switch A and configure the virtual IP address for the group:

```
[edit]
```

```
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group 200 virtual-address 10.1.1.2
```

4. Assign the VRRP priority for this interface in this group:

```
[edit]
```

```
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group 100 priority 200
```

#### Step-by-Step Procedure

Configure the VRRP groups and priorities on Switch B:

1. Create VRRP group 100 on Switch B and configure the virtual IP address for the group:

```
[edit]
```

```
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group 100 virtual-address 10.1.1.1
```

2. Assign the VRRP priority for this interface in this group:

```
[edit]
```

```
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group 100 priority 200
```

Switch A remains the master for group 100 because it has the highest priority for this group.

3. Create VRRP group 200 on Switch A and configure the virtual IP address for the group:

```
[edit]
```

```
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group 200 virtual-address 10.1.1.2
```

4. Assign the VRRP priority for this interface in this group:

```
[edit]
```

```
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group 100 priority 250
```

Switch B becomes the master for group 200 because it has the highest priority for this group.

**Results** Display the results of the configuration on Switch A:

```
user@switch> show configuration
interfaces {
 xe-0/0/0 {
 unit 0 {
 family inet {
 address 10.1.1.251 {
 vrrp-group 100 {
 virtual address 10.1.1.1
 priority 250
 }
 }
 vrrp-group 200 {
 virtual address 10.1.1.2
 priority 200
 }
 }
 }
 }
}
```

Display the results of the configuration on Switch B:

```
user@switch> show configuration
interfaces {
 xe-0/0/0 {
 unit 0 {
 family inet {
 address 10.1.1.252 {
 vrrp-group 100 {
 virtual address 10.1.1.1
 priority 200
 }
 }
 vrrp-group 200 {
 virtual address 10.1.1.2
 priority 250
 }
 }
 }
 }
}
```

## Verification

- [Verifying that VRRP Is Working on Switch A on page 2640](#)
- [Verifying that VRRP Is Working on Switch B on page 2640](#)

### Verifying that VRRP Is Working on Switch A

**Purpose** Verify that VRRP is active on Switch A and that the master and backup roles are correct.

**Action** Use the following command to verify that VRRP is active on Switch A and that the switch is master for group 100 and backup for group 200.

```
user@switch> show vrrp
```

| Interface Address | State | Group | VR state | Timer                                  | Type |
|-------------------|-------|-------|----------|----------------------------------------|------|
| xe-0/0/0.0        | up    | 100   | master   | A .0327 lcl 10.1.1.251<br>vip 10.1.1.1 |      |
| xe-0/0/0.0        | up    | 200   | backup   | A .0327 lcl 10.1.1.251<br>vip 10.1.1.2 |      |

**Meaning** The **show vrrp** command displays fundamental information about the VRRP configuration. This output shows that both VRRP groups are active and that this switch has assumed the correct master and backup roles. The **lcl** address is the physical address of the interface and the **vip** address is the virtual address shared by both switches. The **Timer** value (**A .0327**) indicates the remaining time (in seconds) in which this switch expects to receive a VRRP advertisement from the other switch. If an advertisement for group 200 does not arrive before the timer expires, Switch A asserts itself as the master for this group.

### Verifying that VRRP Is Working on Switch B

**Purpose** Verify that VRRP is active on Switch B and that the master and backup roles are correct.

**Action** Use the following command to verify that VRRP is active on Switch B and that the switch is backup for group 100 and master for group 200.

```
user@switch> show vrrp
```

| Interface Address | State | Group | VR state | Timer                                  | Type |
|-------------------|-------|-------|----------|----------------------------------------|------|
| xe-0/0/0.0        | up    | 100   | backup   | A .0327 lcl 10.1.1.252<br>vip 10.1.1.1 |      |
| xe-0/0/0.0        | up    | 200   | master   | A .0327 lcl 10.1.1.252<br>vip 10.1.1.2 |      |

**Meaning** The **show vrrp** command displays fundamental information about the VRRP configuration. This output shows that both VRRP groups are active and that this switch has assumed the correct master and backup roles. The **lcl** address is the physical address of the interface and the **vip** address is the virtual address shared by both switches. The **Timer** value (**A .0327**) indicates the remaining time (in seconds) in which this switch expects to receive a VRRP advertisement from the other switch. If an advertisement for group 100 does not arrive before the timer expires, Switch B asserts itself as the master for this group.



- Related Documentation**
- [Understanding VRRP on page 2633](#)
  - [Configuring Basic VRRP Support for QFX on page 2641](#)

## Configuring Basic VRRP Support for QFX

To configure basic VRRP support, configure VRRP groups on interfaces by including the **vrrp-group** statement:

```
vrrp-group group-id {
 priority number;
 virtual-address [addresses];
}
```

An interface can be a member of multiple VRRP groups.

You can include this statement at the following hierarchy level:

- **[edit interfaces *interface-name* unit *logical-unit-number* family inet address *address*]**

For each interface, you must configure the following:

- Group identifier—Assign a value from 0 through 255. You must use the same identifier for each switch in the VRRP group.
- Priority—Assign a value from 1 through 255. The switch with the highest priority becomes the VRRP master. Assign different priorities to each switch in the VRRP group. If there are two or more switches with the same priority, the switch with the VRRP interface that has the highest IP address becomes the master.
- Virtual IP address—Normally, you configure only one address per group, but you can configure as many as eight addresses. Do not include a prefix length in a virtual IP address. The following considerations apply to configuring a virtual IP address:
  - You must configure the same address on all the switches in the VRRP group.
  - If you configure a virtual IP address to be the same as a physical interface address, the switch with that interface becomes the master for the group. You must configure the priority to be 255, and you must configure preemption by including the **preempt** statement.
  - If the virtual IP address is not the same as the physical interface address, you must ensure that the address does not appear anywhere else in the switch configuration. For example, verify that you do not use this address for another interface (including an aggregated Ethernet interface) or for a static ARP entry.



**NOTE:** If you enable MAC source address filtering on an interface, you must include the virtual MAC address in the list of source MAC addresses that you specify in the `source-address-filter` statement at the `[edit interfaces interface-name]` hierarchy. MAC addresses ranging from 00:00:5e:00:01:00 through 00:00:5e:00:01:ff are reserved for VRRP, as defined in RFC 3768. The VRRP group number must be the decimal equivalent of the last hexadecimal byte of the virtual MAC address.

---

**Related Documentation**

- [Understanding VRRP on page 2633](#)
- [Configuring the Startup Period for VRRP Operations on page 2643](#)
- [Configuring VRRP Authentication \(IPv4 Only\) on page 2642](#)

---

## Configuring VRRP Authentication (IPv4 Only)

---

VRRP (IPv4 only) protocol exchanges can be authenticated to guarantee that only trusted switches participate in a VRRP group. By default, VRRP authentication is disabled. You can configure one of the following authentication methods for a group, and each switch in the same group must use the same method:

- Simple authentication—Uses a text password included in the transmitted packet. The receiving switch uses an authentication key (password) to verify the packet.
- Message Digest 5 (MD5) algorithm—Adds an authentication header (AH) to the IP packet that encapsulates the VRRP packet. You create an authentication key that is used to create a hash of the packet, and the hash is stored in the AH. A receiving switch recalculates the hash on the incoming packet and compares the hashes. If they are identical, the packet is valid and is accepted. Otherwise the switch drops the incoming packet.

To enable authentication and specify an authentication method, include the **authentication-type** statement.

**authentication-type** *authentication*;

**authentication** can be **simple** or **md5**. The authentication type must be the same for all the switches in the VRRP group.

You can include this statement at the following hierarchy level:

- `[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id]`

If you include the **authentication-type** statement, you can configure a key (password) on each interface by including the **authentication-key** statement:

**authentication-key** *key*;

**key** (the password) is an ASCII string. For simple authentication, it can be from 1 through 8 characters long. For MD5 authentication, it can be from 1 through 16 characters long. If you include spaces, enclose all characters in quotation marks (" ").



**NOTE:** The key must be the same for all switches in the VRRP group.

You can include this statement at the following hierarchy level:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*]

**Related Documentation**

- [Understanding VRRP on page 2633](#)
- [Configuring Basic VRRP Support for QFX on page 2641](#)

## Configuring the Startup Period for VRRP Operations

Configure the startup-silent period interval to avoid alarms caused by the delay or interruption of the incoming VRRP advertisement packets while an interface is coming online. The period starts when the state of a VRRP interface is changed from down to up. During this period, Master Down Events are ignored.

To configure the startup period for VRRP operations, include the **startup-silent-period** statement at the [edit protocols vrrp] hierarchy level:

```
[edit protocols vrrp]
 startup-silent-period seconds;
```

**Related Documentation**

- [Understanding VRRP on page 2633](#)
- [Configuring Basic VRRP Support for QFX on page 2641](#)
- [Example: Configuring VRRP for Load Sharing on page 2636](#)

## Configuring the Advertisement Interval for the VRRP Master

By default, the master switch sends VRRP advertisement packets every second to all members of the VRRP group. These packets indicate that the master switch is still operational. If the master switch fails or becomes unreachable, the backup switch with the highest priority value becomes the new master switch.

You can modify the advertisement interval in seconds or in milliseconds; the interval must be the same for all routing platforms in the VRRP group.

This topic contains the following sections:

- [Modifying the Advertisement Interval in Seconds on page 2644](#)
- [Modifying the Advertisement Interval in Milliseconds on page 2644](#)

## Modifying the Advertisement Interval in Seconds

To modify the time, in seconds, between the sending of VRRP advertisement packets, include the **advertise-interval** statement:

```
advertise-interval seconds;
```

The interval can be from 1 through 255 seconds.

You can include this statement at the following hierarchy level:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*]

## Modifying the Advertisement Interval in Milliseconds

To modify the time, in milliseconds, between the sending of VRRP advertisement packets, include the **fast-interval** statement:

```
fast-interval milliseconds;
```

The interval can be from 100 through 999 milliseconds.

You can include this statement at the following hierarchy level:

- [edit interfaces *interface-name* unit *logical-unit-number* family (inet | inet6) address *address* (vrrp-group | vrrp-inet6-group) *group-id*]



**NOTE:** Junos OS sets the advertisement interval to 0 in VRRP packets. When you configure VRRP with other vendors' equipment, the **fast-interval** statement works correctly only when the other equipment also has the advertisement interval set to 0 in the VRRP packet. Otherwise, Junos OS interprets other routers' settings as advertisement timer errors.

---

### Related Documentation

- [Understanding VRRP on page 2633](#)
- [Configuring Basic VRRP Support for QFX on page 2641](#)
- [Example: Configuring VRRP for Load Sharing on page 2636](#)

---

## Configuring VRRP Preemption and Hold Time

- [Configuring VRRP Preemption on page 2645](#)
- [Configuring the Preemption Hold Time on page 2645](#)
- [Overriding the Hold Time on page 2645](#)

## Configuring VRRP Preemption

By default, a higher-priority VRRP backup switch preempts a lower-priority master switch. To explicitly enable this behavior, include the following statement:

```
preempt;
```

To prohibit a higher-priority VRRP backup switch from preempting a lower-priority master switch, include the following statement on the lower-priority switch:

```
no-preempt;
```

You can include these statements at the following hierarchy level:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*]

## Configuring the Preemption Hold Time

You can also configure a preemption hold time, which is the number of seconds a higher-priority backup router that has just started up waits before preempting the master router. You might want to configure a hold time so that routing protocols or other Junos OS components converge before preemption occurs.

The hold time is applied only on startup. By default, the hold-time value is 0 seconds, meaning that preemption can occur immediately after the backup router starts up.

To modify the preemption hold-time value, configure the following statement:

```
hold-time seconds;
```

The hold time can be from 0 through 3600 seconds.

You can include this statement at the following hierarchy level:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet address vrrp-group *group-id*] preempt

## Overriding the Hold Time

You can use the `asymmetric-hold-time` statement to configure a VRRP master to fail over to the backup immediately—without waiting for the preemption hold time to expire—when a tracked route goes down. Otherwise, the master waits for the hold time to expire before it initiates a failover when a tracked route goes down.

When the tracked route comes up again, the new backup (original master) router waits for the preemption hold time to expire before it reasserts mastership.

You can include this statement at the following hierarchy level:

- [edit protocols vrrp]

Related  
Documentation

- [Understanding VRRP on page 2633](#)

- [Configuring Basic VRRP Support for QFX on page 2641](#)
- [Example: Configuring VRRP for Load Sharing on page 2636](#)

## Configuring a Route to Be Tracked

---

A VRRP master can track a route and dynamically trigger a new master router election if the route becomes unreachable. To enable this behavior, you must configure a cost that will be subtracted from the priority of the master if the tracked route becomes unreachable. The new priority must be less than the priority of one of the backups so that the backup becomes the new master.

To configure a route to be tracked, include the following statements:

```
track {
 priority-hold-time seconds;
 route prefix/prefix-length routing-instance default priority-cost priority;
}
```

You can include these statements at the following hierarchy level:

- `[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id]`

The **prefix** and **prefix-length** values specify the route to be tracked. The **priority-hold-time** statement is the minimum length of time that must elapse between priority changes. If the priority of the master changes because of a tracking event, the priority hold timer begins. If another tracking event or manual configuration change occurs while the timer is running, the new priority update is postponed until the timer expires. You might configure the **priority-hold-time** statement to prevent problems that could occur if there were multiple VRRP transitions in a short period of time.

The **priority-cost** option is the value to be subtracted from the VRRP priority when the tracked route goes down. The value can be 1 through 254. The sum of the costs for all tracked interfaces and routes must be less than or equal to the configured priority (so that subtracting all the costs results in a priority equal to or greater than 0).

### Related Documentation

- [Understanding VRRP on page 2633](#)
- [Configuring Basic VRRP Support for QFX on page 2641](#)
- [Example: Configuring VRRP for Load Sharing on page 2636](#)
- [Configuring a Logical Interface to Be Tracked on page 2646](#)

## Configuring a Logical Interface to Be Tracked

---

VRRP can track whether a logical interface is up, down, or not present, and can change the priority of the switch based on the state of the interface, which might trigger a new master election. VRRP can also track the operational speed of a logical interface and update the priority of the switch when the speed crosses a configured threshold. For each VRRP group, you can track as many as 10 logical interfaces.

When interface tracking is enabled on a switch, you cannot assign the switch a priority of 255 to make it the master for the group.

To configure a logical interface to be tracked, include the following statements:

```
track {
 interface interface-name {
 bandwidth-threshold bits-per-second priority-cost priority;
 priority-cost priority;
 }
 priority-hold-time seconds;
}
```

You can include these statements at the following hierarchy level:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*]

The interface specified is the interface to be tracked for the VRRP group. The **priority-hold-time** statement is the minimum length of time that must elapse between priority changes. If the priority changes because of a tracking event, the priority hold timer begins. If another tracking event or manual configuration change occurs while the timer is running, the new priority update is postponed until the timer expires. You might configure the **priority-hold-time** statement to prevent problems that could occur if there were multiple VRRP transitions in a short period of time.

The **bandwidth-threshold** statement specifies a threshold for the tracked interface. If the bandwidth of the tracked interface drops below the threshold value, the system subtracts the bandwidth threshold **priority-cost** value from the VRRP priority for the switch. You can create as many as five **bandwidth-threshold** statements for each tracked interface.

The interface **priority-cost** statement is the value to be subtracted from the VRRP priority when the tracked route goes down. The value can be 1 through 254. The sum of the costs for all tracked interfaces and routes must be less than or equal to the configured priority (so that subtracting all the costs results in a priority equal to or greater than 0).



**WARNING:** On a QFabric system, do not apply interface tracking to a multichassis link aggregation group (MC-LAG) that includes an interface belonging to a network Node group device and an interface belonging to a server Node group device. If you do apply interface tracking to an MC-LAG configured in this way, a priority update will not occur if the state of the MC-LAG interface changes.

If you configure tracking for more than one interface, Junos OS subtracts the sum of the priority costs for the tracked interfaces from the VRRP priority if all the tracked interfaces fail. However, if you configure the interface **priority-cost** statement and the bandwidth threshold **priority-cost** statement, they are not added together. The switch uses only one priority cost for a tracked interface, as indicated in [Table 237](#):

Table 237: Interface State and Priority Cost Usage

| Tracked Interface State                                      | Priority Cost Usage                                        |
|--------------------------------------------------------------|------------------------------------------------------------|
| Down                                                         | <b>priority-cost</b> <i>priority</i>                       |
| Not down; media speed below one or more bandwidth thresholds | Priority-cost of the lowest applicable bandwidth threshold |

You must configure an interface priority cost only if you do not configure any bandwidth thresholds. If you do not configure an interface **priority-cost** value and the interface fails, Junos OS subtracts the bandwidth threshold **priority-cost** value of the lowest bandwidth threshold from the priority of the switch.

**Related Documentation**

- [Understanding VRRP on page 2633](#)
- [Configuring Basic VRRP Support for QFX on page 2641](#)
- [Example: Configuring VRRP for Load Sharing on page 2636](#)
- [Configuring a Route to Be Tracked on page 2646](#)

## Configuring a Backup to Accept Packets Destined for the Virtual IP Address

By default, a switch configured to be a VRRP backup but acting as the master does not process packets sent to the virtual IP address—that is, packets in which the destination address is the virtual IP address. To configure a backup switch to process packets sent to the virtual IP address while it is acting as the master, include the **accept-data** statement on the backup:

```
accept-data;
```

You can include this statement at the following hierarchy level:

- **[edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group] *group-id***

To explicitly prohibit the backup from accepting packets destined for the virtual IP address while acting as master, include the **no-accept-data** statement:

```
no-accept-data;
```

If you include the **accept-data** statement, configure the connected hosts so that they:

- Process gratuitous ARP requests.
- Do not use packets other than ARP replies to update their ARP cache.

This statement is disabled by default. If you enable it, your configuration does not comply with RFC 3768.

To restrict incoming IP packets to ICMP only, you must configure firewall filters to accept only ICMP packets.



- Related Documentation**
- [Understanding VRRP on page 2633](#)
  - [Configuring Basic VRRP Support for QFX on page 2641](#)
  - [Example: Configuring VRRP for Load Sharing on page 2636](#)

---

## Configuring Passive ARP Learning for VRRP Backups

By default, VRRP backup switches drop ARP requests for the MAC address of the VRRP IP. This means that backups do not learn the ARP mappings (IP address to MAC address mappings) for the hosts sending the requests. If it becomes the master, the configured backup must learn all the entries that were present in the ARP cache of the original master. In environments with many directly attached hosts, the number of ARP entries to learn can be very large. This can cause a significant delay while the backup transitions to the master state, during which traffic transmitted to some of the hosts might be dropped.

Passive ARP learning enables the ARP cache in the backup to learn approximately the same contents as the ARP cache in the master, thus preventing the problem of needing to learn many ARP entries quickly. To enable passive ARP learning, include the **passive-learning** statement at the **[edit system arp]** hierarchy level:

```
[edit system arp]
passive-learning;
```

We recommend setting passive learning on both the backup and master VRRP switches. Doing so prevents the need to manually configure a master that fails and becomes a backup. While a switch operates as the master, the passive learning configuration has no impact. The configuration takes effect only when a switch operates as a backup.

- Related Documentation**
- [Understanding VRRP on page 2633](#)
  - [Configuring Basic VRRP Support for QFX on page 2641](#)
  - [Example: Configuring VRRP for Load Sharing on page 2636](#)

---

## Configuring the Silent Period

When the state of a VRRP interface changes from down to up, a silent period begins. During this period, any master down events are ignored. Configure the silent period interval to avoid problems that can be caused if incoming VRRP advertisement packets are delayed or interrupted while an interface starts up.

To configure the silent period, include the **startup-silent-period** statement at the **[edit protocols vrrp]** hierarchy level:

```
[edit protocols vrrp]
startup-silent-period seconds;
```

- Related Documentation**
- [Understanding VRRP on page 2633](#)
  - [Configuring Basic VRRP Support for QFX on page 2641](#)

- [Example: Configuring VRRP for Load Sharing on page 2636](#)

## Configuring Inheritance for a VRRP Group

---

Junos OS enables you to configure VRRP groups on the various subnets of a VLAN to inherit the state and configuration of one of the groups, which is known as the *active VRRP group*. By configuring inheritance, you can prevent VRRP groups other than the active group from sending out VRRP advertisements. When the **vrrp-inherit-from** configuration statement is included in the configuration, only the active VRRP group from which the other VRRP groups are inheriting the state sends out VRRP advertisements; the groups inheriting the state do not send any VRRP advertisements, because the state is maintained only on the group from which the state is inherited.

If the **vrrp-inherit-from** statement is not configured, each of the VRRP master groups in the various subnets on the VLAN sends out separate VRRP advertisements and adds to the traffic on the VLAN.

To configure inheritance for a VRRP group, include the **vrrp-inherit-from** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* *vrrp-group group-id*]**:

```
[edit interfaces interface-name unit logical-unit-number family inet address address
 vrrp-group group-id]
 vrrp-inherit-from vrrp-group;
```

When you configure a group to inherit a state from another group, note the following conditions:

- Both inheriting groups and active groups must be on the same physical interface and logical system. However, the groups need not necessarily be on the same VLAN or logical interface.
- Both inheriting groups and active groups must be on the same routing instances; however, this limitation does not apply for groups on the integrated routing and bridging (IRB) interfaces.

When you include the **vrrp-inherit-from** statement for a VRRP group, the VRRP group inherits the following parameters from the active group:

- **advertise-interval**
- **authentication-key**
- **authentication-type**
- **fast-interval**
- **preempt | no-preempt**
- **priority**
- **track interfaces**
- **track route**

However, you can configure the **accept-data | no-accept-data** statement for the group to specify whether the interface should accept packets destined for the virtual IP address.

- Related Documentation**
- [Understanding VRRP on page 2633](#)
  - [Configuring Basic VRRP Support for QFX on page 2641](#)
  - [Example: Configuring VRRP for Load Sharing on page 2636](#)

---

## Troubleshooting VRRP

---

**Problem**    **Description:** If you configure multiple VRRP groups on an interface (using multiple VLANs), traffic for some of the groups might be briefly dropped if a failover occurs. This can happen because the new master must send gratuitous ARP replies for each VRRP group to update the ARP tables in the connected devices, and there is a short delay between each gratuitous ARP reply. Traffic sent by devices that have not yet received the gratuitous ARP reply is dropped (until the device receives the reply and learns the MAC address of the new master).

**Solution**    Configure a failover delay so that the new master delays sending gratuitous ARP replies for the period that you set. This allows the new master to send the ARP replies for all of the VRRP groups simultaneously.

- Related Documentation**
- [failover-delay on page 2690](#)



# Configuration Statements and Operational Commands

- [Configuration Statements \(Adaptive Load Balancing\) on page 2655](#)
- [Configuration Statements \(Graceful Restart\) on page 2657](#)
- [Configuration Statements \(Graceful Switchover\) on page 2673](#)
- [Configuration Statements \(Nonstop Bridging and Routing\) on page 2675](#)
- [Configuration Statements \(VRRP\) on page 2683](#)
- [Operational Mode Commands \(Graceful Restart\) on page 2707](#)
- [Operational Mode Commands \(Graceful Switchover\) on page 2737](#)
- [Operational Mode Command \(Nonstop Routing\) on page 2747](#)
- [Operational Mode Commands \(VRRP\) on page 2751](#)



## CHAPTER 102

# Configuration Statements (Adaptive Load Balancing)

- [adaptive on page 2656](#)

## adaptive

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>adaptive {<br/>    pps;<br/>    scan-interval <i>multiple</i>;<br/>    tolerance <i>tolerance-percentage</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit interfaces aex aggregated-ether-options load-balance]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2R3 for MX Series Routers.<br>Statement introduced in Junos OS Release 14.1 for PTX Series Packet Transport Routers.<br>Statement introduced in Junos OS Release 15.1X53-D10 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Correct a genuine traffic imbalance by using a feedback mechanism to distribute the traffic across the links of an aggregated Ethernet bundle.                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <p><b>pps</b>—(PTX Series only) The type of traffic rate among the members of the AE bundle is measured packets per second. The default rate type is bytes per second.</p> <p><b>scan-interval <i>multiple</i></b>—(PTX Series only) Scan interval, as a multiple of a 30-second interval.<br/><b>Range:</b> 1 through 5<br/><b>Default:</b> 1</p> <p><b>tolerance <i>tolerance-percentage</i></b>—(MX Series and PTX Series) Limit to the variance in the packet traffic flow to the aggregated Ethernet links in a percentage.<br/><b>Range:</b> 1 through 100 percent<br/><b>Default:</b> 20 percent</p> |
| <b>Required Privilege Level</b> | interface - To view this statement in the configuration.<br>interface-control - To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Aggregated Ethernet Load Balancing on page 2573</a></li><li>• <a href="#">Example: Configuring Aggregated Ethernet Load Balancing</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                           |



# Configuration Statements (Graceful Restart)

- [disable](#) on page 2658
- [disable \(BGP Graceful Restart\)](#) on page 2659
- [graceful-restart \(Enabling Globally\)](#) on page 2660
- [graceful-restart \(Protocols BGP\)](#) on page 2662
- [graceful-restart](#) on page 2664
- [helper-disable \(OSPF\)](#) on page 2666
- [no-strict-lsa-checking](#) on page 2667
- [notify-duration](#) on page 2668
- [redundancy \(Graceful Switchover\)](#) on page 2669
- [restart-duration](#) on page 2670
- [restart-time \(BGP Graceful Restart\)](#) on page 2671
- [stale-routes-time](#) on page 2672

## disable

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | disable;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols (bgp   isis   ldp   ospf   ospf3   pim   rip   ripng   rsvp) <a href="#">graceful-restart</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (bgp   ldp   ospf   ospf3   pim) <a href="#">graceful-restart</a>],</p> <p>[edit protocols (bgp   esis   isis   ospf   ospf3   ldp   pim   rip   ripng   rsvp) <a href="#">graceful-restart</a>],</p> <p>[edit protocols bgp group <i>group-name</i> <a href="#">graceful-restart</a>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>ip-address</i> <a href="#">graceful-restart</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (bgp   ldp   ospf   ospf3   pim) <a href="#">graceful-restart</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options <a href="#">graceful-restart</a>],</p> <p>[edit routing-options <a href="#">graceful-restart</a>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Disable graceful restart.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Enabling Graceful Restart</a></li> <li>• <a href="#">Configuring Routing Protocols Graceful Restart on page 2580</a></li> <li>• <a href="#">Configuring Graceful Restart for MPLS-Related Protocols</a></li> <li>• <a href="#">Configuring VPN Graceful Restart</a></li> <li>• <a href="#">Configuring Logical System Graceful Restart</a></li> <li>• <a href="#">Graceful Restart Configuration Statements</a></li> <li>• <a href="#">Configuring Graceful Restart for QFabric Systems</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

## disable (BGP Graceful Restart)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | disable;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> protocols bgp graceful-restart],<br>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> graceful-restart],<br>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i> graceful-restart],<br>[edit protocols bgp graceful-restart],<br>[edit protocols bgp <b>group</b> <i>group-name</i> graceful-restart],<br>[edit protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i> graceful-restart] |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>         | Disable graceful restart for BGP. Graceful restart allows a routing device undergoing a restart to inform its adjacent neighbors and peers of its condition.                                                                                                                                                                                                                                                                                                                                                                                                                       |



**NOTE:** When you disable graceful restart at one level in the configuration statement hierarchy, it is also disabled at lower levels in the same hierarchy. For example, if you disable graceful restart at the [edit protocols bgp **group** *group-name*] hierarchy level, it is disabled for all the peers in the group. Therefore, if you want to enable graceful restart for some peers in a group and disable it for others, enable graceful restart at the [edit protocols bgp **group** *group-name*] hierarchy level and disable graceful restart for each peer at the [edit protocols bgp **group** *group-name* **neighbor** *address*] hierarchy level.

|                                 |                                                                                                                                                                                                                                                           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Graceful Restart Options for BGP on page 2581</a></li> <li>• <a href="#">graceful-restart on page 2662</a></li> <li>• <i>restart-time</i></li> <li>• <i>stale-routes-time</i></li> </ul> |

## graceful-restart (Enabling Globally)

---

|                            |                                                                                                                                                                                                                                                                                                                                        |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>graceful-restart {<br/>  disable;<br/>  helper-disable;<br/>  maximum-helper-recovery-time <i>seconds</i>;<br/>  maximum-helper-restart-time <i>seconds</i>;<br/>  notify-duration <i>seconds</i>;<br/>  recovery-time <i>seconds</i>;<br/>  restart-duration <i>seconds</i>;<br/>  stale-routes-time <i>seconds</i>;<br/>}</pre> |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> routing-options],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],<br>[edit routing-options],<br>[edit routing-instances <i>routing-instance-name</i> routing-options]                                            |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                             |
| <b>Description</b>         | Configure graceful restart globally to enable the feature. You cannot enable graceful restart for specific protocols unless graceful restart is also enabled globally. You can, optionally, modify the global settings at the individual protocol level.                                                                               |




### NOTE:

- For VPNs, the `graceful-restart` statement allows a router whose VPN control plane is undergoing a restart to continue to forward traffic while recovering its state from neighboring routers.
  - For BGP, if you configure graceful restart after a BGP session has been established, the BGP session restarts and the peers negotiate graceful restart capabilities.
  - LDP sessions flap when `graceful-restart` configurations change.
- 

|                                 |                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>                  | Graceful restart is disabled by default.                                                                                                                                          |
| <b>Options</b>                  | The remaining statements are explained separately.                                                                                                                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Enabling Graceful Restart</a></li><li>• <a href="#">Configuring Routing Protocols Graceful Restart on page 2580</a></li></ul> |

- *Configuring Graceful Restart for MPLS-Related Protocols*
- *Configuring VPN Graceful Restart*
- *Configuring Logical System Graceful Restart*
- *Graceful Restart Configuration Statements*
- *Configuring Graceful Restart for QFabric Systems*

## graceful-restart (Protocols BGP)

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <pre>graceful-restart {     disable;     restart-time seconds;     stale-routes-time seconds; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Hierarchy Level          | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],<br/>         [edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>],<br/>         [edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],<br/>         [edit protocols bgp],<br/>         [edit protocols bgp <b>group</b> <i>group-name</i>],<br/>         [edit protocols bgp <i>group</i> <i>group-name</i> <b>neighbor</b> <i>address</i>]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Release Information      | <p>Statement introduced before Junos OS Release 7.4.<br/>         Statement introduced in Junos OS Release 9.0 for EX Series switches.<br/>         Statement introduced in Junos OS Release 12.1 for the QFX Series.<br/>         Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Description              | <p>Configure graceful restart for BGP. Graceful restart allows a routing device undergoing a restart to inform its adjacent neighbors and peers of its condition. Graceful restart is disabled by default. However, helper mode, the ability to assist a neighboring router attempting a graceful restart, is enabled by default.</p> <p>To configure the duration of the BGP graceful restart period, include the <b>restart-time</b> statement at the [edit protocols bgp graceful-restart] hierarchy level. To set the length of time the router waits to receive messages from restarting neighbors before declaring them down, include the <b>stale-routes-time</b> statement at the [edit protocols bgp graceful-restart] hierarchy level.</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p> <b>NOTE:</b> If you configure graceful restart after a BGP session has been established, the BGP session restarts and the peers negotiate graceful restart capabilities.</p> </div> <p>Configure graceful restart globally at the [edit routing-options] or [edit routing-instances <i>instance-name</i> routing-options] hierarchy level to enable the feature. You cannot enable graceful restart for specific protocols unless graceful restart is also enabled globally. You can, optionally, modify the global settings at the individual protocol level.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>routing—To view this statement in the configuration.<br/>         routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Related Documentation    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Graceful Restart Options for BGP on page 2581</a></li> <li>• <a href="#">Configuring Graceful Restart for QFabric Systems</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

- *Junos OS High Availability Library for Routing Devices*

## graceful-restart

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>graceful-restart {   disable;   helper-disable (standard   restart-signaling   both);   no-strict-lsa-checking;   notify-duration <i>seconds</i>;   restart-duration <i>seconds</i>; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>     | <pre>[edit logical-systems <i>logical-system-name</i> protocols (<b>ospf</b>   ospf3)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<b>ospf</b>   ospf3)], [edit protocols (<b>ospf</b>   ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols ospf]</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Support for the <b>no-strict-lsa-checking</b> statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the helper mode <b>standard</b>, <b>restart-signaling</b>, and <b>both</b> options introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>         | Configure graceful restart for OSPF.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>             | <p><b>disable</b>—Disable graceful restart for OSPF.</p> <p><b>helper-disable (standard   restart-signaling   both)</b>—Disable helper mode for graceful restart. When helper mode is disabled, a device cannot help a neighboring device that is attempting to restart. Beginning with Junos OS Release 11.4, you can configure restart signaling-based helper mode for OSPFv2 graceful restart configurations. The <b>standard</b>, <b>restart-signaling</b>, and <b>both</b> options are only supported for OSPFv2. Specify <b>standard</b> to disable helper mode for standard graceful restart (based on RFC 3623). Specify <b>restart-signaling</b> to disable helper mode for restart signaling-based graceful restart (based on RFC 4811, RFC 4812, and RFC 4813). Specify <b>both</b> to disable helper mode for both standard and restart signaling-based graceful restart. The last committed statement takes precedence over the previously configured statement.</p> <p><b>Default:</b> Helper mode is enabled by default. For OSPFv2, both standard and restart-signaling based helper modes are enabled by default.</p> <p><b>no-strict-lsa-checking</b>—Disable strict OSPF link-state advertisement (LSA) checking to prevent the termination of graceful restart by a helping router. LSA checking is enabled by default.</p> |



**NOTE:** The **helper-disable** statement and the **no-strict-lsa-checking** statement cannot be configured at the same time. If you attempt to configure both statements at the same time, the routing device displays a warning message when you enter the **show protocols (ospf | ospf3)** command.



**notify-duration *seconds***—Estimated time needed to send out purged grace LSAs over all the interfaces.

**Range:** 1 through 3600 seconds

**Default:** 30 seconds

**restart-duration *seconds***—Estimated time needed to reacquire a full OSPF neighbor from each area.

**Range:** 1 through 3600 seconds


**Default:** 180 seconds

|                                 |                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration. |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------|

|                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Graceful Restart for OSPF on page 4497</a></li><li>• <a href="#">Example: Configuring the Helper Capability Mode for OSPFv2 Graceful Restart on page 4501</a></li><li>• <a href="#">Example: Configuring the Helper Capability Mode for OSPFv3 Graceful Restart on page 4504</a></li><li>• <a href="#">Example: Disabling Strict LSA Checking for OSPF Graceful Restart on page 4507</a></li></ul> |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## helper-disable (OSPF)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | helper-disable < both   restart-signaling   standard >;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ospf <a href="#">graceful-restart</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf <a href="#">graceful-restart</a> ],<br>[edit protocols ospf <a href="#">graceful-restart</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols ospf <a href="#">graceful-restart</a> ]                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Options <b>both</b> , <b>restart-signaling</b> , and <b>standard</b> introduced in Junos OS Release 11.4.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Disable helper mode for graceful restart. When helper mode is disabled, a router cannot help a neighboring router that is attempting to restart. The last committed statement takes precedence over the previously configured statement.                                                                                                                                                                                                                                                                                                                                           |
| <b>Default</b>                  | Helper mode is enabled by default for OSPF.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <b>both</b> —(Optional) Disable helper mode for both standard and restart signaling-based graceful restart.<br><br><b>restart-signaling</b> —(Optional) Disable helper mode for restart signaling-based graceful restart (based on RFC 4811, RFC 4812, and RFC 4813).<br><br><div> <b>NOTE:</b> Restart signaling-based helper mode is not supported for OSPFv3 configurations.</div><br><br><b>standard</b> —(Optional) Disable helper mode for standard graceful restart (based on RFC 3623). |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Routing Protocols Graceful Restart on page 2580</a></li><li>• <a href="#">Configuring Graceful Restart for MPLS-Related Protocols</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                    |

---

## no-strict-lsa-checking

---

|                                 |                                                                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-strict-lsa-checking;                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit protocols (ospf   ospf3) <a href="#">graceful-restart</a> ]                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                 |
| <b>Description</b>              | Disable strict OSPF link-state advertisement (LSA) checking to prevent the termination of graceful restart by a helping router or switch.                                                                                                                                                                      |
| <b>Default</b>                  | By default, LSA checking is enabled.                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Graceful Restart Options for OSPF and OSPFv3 on page 2583</a></li><li>• <i>Configuring Graceful Restart for QFabric Systems</i></li><li>• <a href="#">maximum-neighbor-recovery-time on page 4868</a></li><li>• <i>recovery-time</i></li></ul> |

## notify-duration

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | notify-duration <i>seconds</i> ;                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit protocols (ospf   ospf3) <a href="#">graceful-restart</a> ],<br>[edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) <a href="#">graceful-restart</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols (ospf   ospf3) <a href="#">graceful-restart</a> ],<br>[edit routing-instances <i>instance-name</i> protocols (ospf   ospf3) <a href="#">graceful-restart</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.3.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Specify the length of time the router or switch notifies helper OSPF routers that it has completed graceful restart.                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <b><i>seconds</i></b> —Length of time in the router notifies helper OSPF routers that it has completed graceful restart.<br><b>Range:</b> 1 through 3600<br><b>Default:</b> 30                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Graceful Restart Options for OSPF and OSPFv3 on page 2583</a></li><li>• <a href="#">Configuring Graceful Restart for QFabric Systems</a></li><li>• <a href="#">restart-duration on page 2670</a></li></ul>                                                                                                                                                                              |

## redundancy (Graceful Switchover)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> redundancy {   failover {     on-disk-failure;     on-loss-of-keepalives;   }   graceful-switchover; } </pre>                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit chassis]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | <p>Enable redundant Routing Engines on a Virtual Chassis with two or more member switches or on a Virtual Chassis Fabric, on a standalone EX6200 or EX8200 switch with more than one Routing Engine.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                 |
| <b>Default</b>                  | Redundancy is enabled for the Routing Engines.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">graceful-switchover on page 2673</a></li> <li>• <a href="#">Configuring Graceful Routing Engine Switchover in a Virtual Chassis (CLI Procedure) on page 2629</a></li> <li>• <a href="#">Configuring Graceful Routing Engine Switchover on page 2627</a></li> <li>• <a href="#">Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)</a></li> <li>• <a href="#">High Availability Features for EX Series Switches Overview</a></li> </ul> |

## restart-duration

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>restart-duration seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>     | <code>[edit logical-systems <i>logical-system-name</i> protocols (isis   ospf   ospf3   pim) graceful-restart],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3   pim) graceful-restart],</code><br><code>[edit protocols (esis   isis   ospf   ospf3   pim) graceful-restart],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3   pim) graceful-restart],</code><br><code>[edit routing-options graceful-restart]</code>                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>         | <p>Configure the grace period for graceful restart globally.</p> <p>Additionally, you can individually configure the duration of the graceful restart period for the End System-to-Intermediate System (ES-IS), Intermediate System-to-Intermediate System (IS-IS), Open Shortest Path First (OSPF), and OSPFv3 protocols and for Protocol Independent Multicast (PIM) sparse mode.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>             | <p><b><i>seconds</i></b>—Time for the graceful restart period.</p> <p><b>Range:</b></p> <p>The range of values varies according to whether the graceful restart period is being set globally or for a particular protocol:</p> <ul style="list-style-type: none"><li>• <b>[edit routing-options graceful-restart]</b> (global setting)—120 through 900</li><li>• ES-IS—30 through 300</li><li>• IS-IS—30 through 300</li><li>• OSPF/OSPFv3—1 through 3600</li><li>• PIM—30 through 300</li></ul> <p><b>Default:</b></p> <p>The default value varies according to whether the graceful restart period is being set globally or for a particular protocol:</p> <ul style="list-style-type: none"><li>• <b>[edit routing-options graceful-restart]</b> (global setting)—300</li><li>• ES-IS—180</li><li>• IS-IS—210</li><li>• OSPF/OSPFv3—180</li><li>• PIM—60</li></ul> |

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Enabling Graceful Restart*
- *Configuring Graceful Restart for MPLS-Related Protocols*
- *Configuring VPN Graceful Restart*
- *Configuring Graceful Restart for VPNs*
- *Configuring Logical System Graceful Restart*

## restart-time (BGP Graceful Restart)

**Syntax** restart-time *seconds*;

**Hierarchy Level** [edit protocols (bgp | rip | ripng) [graceful-restart](#)],  
[edit logical-systems *logical-system-name* protocols (bgp | rip | ripng) [graceful-restart](#) ([Enabling Globally](#))],  
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp [graceful-restart](#)],  
[edit routing-instances *routing-instance-name* protocols bgp [graceful-restart](#)]

**Release Information** Statement introduced in Junos OS Release 8.3.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 12.1 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Configure the duration of the BGP, RIP, or next-generation RIP (RIPng) graceful restart period.

**Options** *seconds*—Length of time for the graceful restart period.  
**Range:** 1 through 600 seconds  
**Default:** Varies by protocol:

- BGP—120 seconds
- RIP and RIPng—60 seconds

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring Graceful Restart Options for BGP on page 2581](#)
- [Configuring Graceful Restart Options for RIP and RIPng on page 2585](#)
- *Configuring Graceful Restart for QFabric Systems*
- [stale-routes-time on page 2672](#)

## stale-routes-time

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>stale-routes-time seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-routing-name</i> protocols bgp <a href="#">graceful-restart</a> ],<br>[edit logical-systems <i>logical-routing-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <a href="#">graceful-restart</a> ],<br>[edit protocols bgp <a href="#">graceful-restart</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols bgp <a href="#">graceful-restart</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.3.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.                                                                                                                                                        |
| <b>Description</b>              | Specify the maximum time that stale routes are kept during a restart. The <b>stale-routes-time</b> statement allows you to set the length of time the routing device waits to receive messages from restarting neighbors before declaring them down.                                                                                                                                                                          |
| <b>Options</b>                  | <b>seconds</b> —Time the router device waits to receive messages from restarting neighbors before declaring them down.<br><b>Range:</b> 1 through 600 seconds<br><b>Default:</b> 300 seconds                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Graceful Restart Options for BGP on page 2581</a></li><li>• <a href="#">Configuring Graceful Restart for QFabric Systems</a></li><li>• <a href="#">restart-time (BGP Graceful Restart) on page 2671</a></li></ul>                                                                                                                                             |



# Configuration Statements (Graceful Switchover)

- [graceful-switchover](#) on page 2673
- [redundancy \(Graceful Switchover\)](#) on page 2674

## [graceful-switchover](#)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>graceful-switchover;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit chassis <a href="#">redundancy</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.2 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | For switches with more than one Routing Engine, including those in a Virtual Chassis or a Virtual Chassis Fabric, configure the master Routing Engine to switch over gracefully to a backup Routing Engine without interruption to packet forwarding.                                                                                                                                                                                                                                                                                                      |
| <b>Default</b>                  | Graceful Routing Engine switchover (GRES) is disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Nonstop Active Routing on Switches</a> on page 2611</li> <li>• <a href="#">Configuring Graceful Routing Engine Switchover</a> on page 2627</li> <li>• <a href="#">Configuring Graceful Routing Engine Switchover in a Virtual Chassis (CLI Procedure)</a> on page 2629</li> <li>• <a href="#">Configuring Nonstop Active Routing on Switches</a></li> <li>• <a href="#">Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)</a></li> </ul> |

## redundancy (Graceful Switchover)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>redundancy {<br/>  failover {<br/>    on-disk-failure;<br/>    on-loss-of-keepalives;<br/>  }<br/>  graceful-switchover;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit chassis]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | <p>Enable redundant Routing Engines on a Virtual Chassis with two or more member switches or on a Virtual Chassis Fabric, on a standalone EX6200 or EX8200 switch with more than one Routing Engine.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                           |
| <b>Default</b>                  | Redundancy is enabled for the Routing Engines.                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">graceful-switchover on page 2673</a></li><li>• <a href="#">Configuring Graceful Routing Engine Switchover in a Virtual Chassis (CLI Procedure) on page 2629</a></li><li>• <a href="#">Configuring Graceful Routing Engine Switchover on page 2627</a></li><li>• <a href="#">Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)</a></li><li>• <a href="#">High Availability Features for EX Series Switches Overview</a></li></ul> |

# Configuration Statements (Nonstop Bridging and Routing)

- [nonstop-bridging on page 2675](#)
- [nonstop-routing on page 2676](#)
- [synchronize on page 2677](#)
- [traceoptions on page 2679](#)

## nonstop-bridging

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | nonstop-bridging;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit protocols layer2-control]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.4.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | For platforms with two Routing Engines, configure a master Routing Engine to switch over gracefully to a backup Routing Engine and preserve Layer 2 Control Protocol (L2CP) information.                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Synchronizing the Routing Engine Configuration</i></li> <li>• <i>Configuring Nonstop Bridging</i></li> <li>• For information about configuring NSB on EX Series switches that do not support the Enhanced Layer 2 Software (ELS) CLI style, see <i>Configuring Nonstop Bridging on EX Series Switches (CLI Procedure)</i></li> <li>• For information about configuring NSB on switches that support ELS, see <a href="#">Configuring Nonstop Bridging on Switches (CLI Procedure) on page 2593</a></li> </ul> |

## nonstop-routing

---

**Syntax** nonstop-routing;

**Hierarchy Level** [edit routing-options]



**NOTE:** Although `nonstop-routing` is also a valid keyword at the `logical-systems` hierarchy level, it is not supported.

|                                 |                                                                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.4.<br>Statement introduced in Junos OS Release 10.4 for EX Series switches.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.<br>Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series switches |
| <b>Description</b>              | For routing platforms with two Routing Engines, configure a master Routing Engine to switch over gracefully to a backup Routing Engine and to preserve routing protocol information.                                                                                            |
| <b>Default</b>                  | disabled                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Nonstop Active Routing</i></li></ul>                                                                                                                                                                                     |

## synchronize

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | synchronize;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>     | [edit system commit]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b> | Statement introduced in Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 10.4 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>         | For devices with multiple Routing Engines only. Configure the <b>commit</b> command to automatically perform a <b>commit synchronize</b> action between dual Routing Engines within the same chassis. The Routing Engine on which you execute the <b>commit</b> command (the requesting Routing Engine) copies and loads its candidate configuration to the other (the responding) Routing Engine. Each Routing Engine then performs a syntax check on the candidate configuration file being committed. If no errors are found, the configuration is activated and becomes the current operational configuration on both Routing Engines. |



**NOTE:** If you configure the **commit synchronize** statement at the [edit system] hierarchy level and issue a **commit** in the master Routing Engine, the master configuration is automatically synchronized with the backup. However, if the backup Routing Engine is down when you issue the **commit**, the Junos OS displays a warning and commits the candidate configuration in the master Routing Engine. When the backup Routing Engine comes up, its configuration will automatically be synchronized with the master. A newly inserted backup Routing Engine automatically synchronizes its configuration with the master Routing Engine configuration.



**NOTE:** When you configure nonstop active routing (NSR), you must configure the **commit synchronize** statement. Otherwise, the **commit** operation fails.

On the TX Matrix router, synchronization only occurs between the Routing Engines within the same chassis. When synchronization is complete, the new configuration is then distributed to the Routing Engines on the T640 routers. That is, the master Routing Engine on the TX Matrix router distributes the configuration to the master Routing Engine on each T640 router. Likewise, the backup Routing Engine on the TX Matrix router distributes the configuration to the backup Routing Engine on each T640 router.

On the TX Matrix Plus router, synchronization only occurs between the Routing Engines within the switch-fabric chassis and when synchronization is complete, the new configuration is then distributed to the Routing Engines on the line-card chassis (LCC). That is, the master Routing Engine on the TX Matrix Plus router distributes the configuration to the master Routing Engine on each LCC. Likewise, the backup Routing Engine on the TX Matrix Plus router distributes the configuration to the backup Routing Engine on each LCC.

In EX Series Virtual Chassis configurations:

- On EX4200 switches in Virtual Chassis, synchronization occurs between the switch in the master role and the switch in the backup role.
- On EX8200 switches in a Virtual Chassis, synchronization occurs only between the master and backup XRE200 External Routing Engines.

**Options**    **and-quit**—(Optional) Quit configuration mode if the commit synchronization succeeds.

**at**—(Optional) Time at which to activate configuration changes.

**comment**—(Optional) Write a message to the commit log.

**force**—(Optional) Force a commit synchronization on the other Routing Engine (ignore warnings).

**scripts**—(Optional) Push scripts to the other Routing Engine.

**Required Privilege**    **system**—To view this statement in the configuration.

**Level**    **system-control**—To add this statement to the configuration.

**Related**    • *Synchronizing the Routing Engine Configuration*  
**Documentation**    • *Configuring Multiple Routing Engines to Synchronize Committed Configurations Automatically*

## traceoptions

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre> traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;disable&gt;; } </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit routing-options],</p> <p>[edit routing-options flow],</p> <p>[edit routing-options multicast]</p>                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p><b>nsr-synchronization</b> flag for BGP, IS-IS, LDP, and OSPF added in Junos OS Release 8.4.</p> <p><b>nsr-synchronization</b> and <b>nsr-packet</b> flags for BFD sessions added in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>nsr-synchronization</b> flag for RIP and RIPng added in Junos OS Release 9.0.</p> <p><b>nsr-synchronization</b> flag for Layer 2 VPNs and VPLS added in Junos OS Release 9.1.</p> <p><b>nsr-synchronization</b> flag for PIM added in Junos OS Release 9.3.</p> <p><b>nsr-synchronization</b> flag for MPLS added in Junos OS Release 10.1.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p><b>nsr-synchronization</b> flag for MSDP added in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| <b>Description</b>         | <p>Define tracing operations that track all routing protocol functionality in the routing device.</p> <p>To specify more than one tracing operation, include multiple <b>flag</b> statements.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Default</b>             | If you do not include this statement, no global tracing operations are performed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>             | <p><b>Values:</b></p> <p><b>disable</b>—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>/var/log</b>. We recommend that you place global routing protocol tracing output in the file <b>routing-log</b>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

**files *number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file.0***, then ***trace-file.1***, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

**Range:** 2 through 1000 files

**Default:** 10 files

**flag *flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. These are the global routing protocol tracing options:

- **all**—All tracing operations
- **condition-manager**—Condition-manager events
- **config-internal**—Configuration internals
- **general**—All normal operations and routing table changes (a combination of the **normal** and **route** trace operations)
- **graceful-restart**—Graceful restart operations
- **normal**—All normal operations
- **nsr-packet**—Detailed trace information for BFD nonstop active routing only
- **nsr-synchronization**—Tracing operations for nonstop active routing
- **nsr-synchronization**—Nonstop active routing synchronization
- **parse**—Configuration parsing
- **policy**—Routing policy operations and actions
- **regex-parse**—Regular-expression parsing
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

**no-world-readable**—(Optional) Prevent any user from reading the log file.

**size *size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file.0***. When the ***trace-file*** again reaches its maximum size, ***trace-file.0*** is renamed ***trace-file.1*** and ***trace-file*** is renamed ***trace-file.0***. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

**Syntax:** ***xk*** to specify KB, ***xm*** to specify MB, or ***xg*** to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 128 KB



**world-readable**—(Optional) Allow any user to read the log file.

**Required Privilege** routing and trace—To view this statement in the configuration.  
**Level** routing-control and trace-control—To add this statement to the configuration.

**Related** • *Example: Tracing Global Routing Protocol Operations*  
**Documentation**



## Configuration Statements (VRRP)


- [accept-data on page 2684](#)
- [advertise-interval on page 2685](#)
- [asymmetric-hold-time on page 2686](#)
- [authentication-key on page 2687](#)
- [authentication-type on page 2688](#)
- [bandwidth-threshold on page 2689](#)
- [failover-delay on page 2690](#)
- [fast-interval on page 2691](#)
- [hold-time \(VRRP\) on page 2692](#)
- [interface on page 2693](#)
- [preempt \(VRRP\) on page 2694](#)
- [priority \(Protocols VRRP\) on page 2695](#)
- [priority-cost \(VRRP\) on page 2696](#)
- [priority-hold-time on page 2697](#)
- [route \(Interfaces\) on page 2698](#)
- [startup-silent-period on page 2699](#)
- [traceoptions on page 2700](#)
- [track \(VRRP\) on page 2702](#)
- [virtual-address on page 2703](#)
- [vrrp-group on page 2704](#)

## accept-data

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | (accept-data   no-accept-data);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i> ],<br>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i> ],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i> ],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i> ]                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | <p>In a Virtual Router Redundancy Protocol (VRRP) configuration, determine whether or not an interface accepts packets destined for the virtual IP address:</p> <ul style="list-style-type: none"><li>• <b>accept-data</b>—Enable the interface to accept packets destined for the virtual IP address.</li><li>• <b>no-accept-data</b>—Prevent the interface from accepting packets destined for the virtual IP address.</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Default</b>                  | <p>If the <b>accept-data</b> statement is not configured, the master router responds to ARP requests only.</p> <p>The <b>accept-data</b> statement has the following restrictions and limitations:</p> <ul style="list-style-type: none"><li>• If the master router owns the virtual IP address or if the priority of the master router is set to 255, the <b>accept-data</b> statement becomes inapplicable.</li><li>• If the master router owns the virtual IP address, the master router responds to Internet Control Message Protocol (ICMP) message requests.</li><li>• If you want to restrict the incoming IP packets to ICMP only, you must configure firewall filters to accept only ICMP packets.</li><li>• If you include the <b>accept-data</b> statement, your routing platform configuration does not comply with RFC 3768 (see section 6.4.3 of RFC 3768, <i>Virtual Router Redundancy Protocol (VRRP)</i>).</li></ul> |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring an Interface to Accept All Packets Destined for the Virtual IP Address of a VRRP Group</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## advertise-interval


|                                                                                                                                                                                                                                                                                                                                           |                                                                                                                                                                                                                                                                                                                                       |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                             | <code>advertise-interval seconds;</code>                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                    | <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>]</p> |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.</p>                                                                                                                             |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                        | <p>Configure the interval between Virtual Router Redundancy Protocol (VRRP) IPv4 advertisement packets.</p> <p>All routers in the VRRP group must use the same advertisement interval.</p>                                                                                                                                            |
| <div>  <p><b>NOTE:</b> When VRRPv3 is enabled, the <code>advertise-interval</code> statement cannot be used to configure advertisement intervals. Instead, use the <code>fast-interval</code> statement to configure advertisement intervals.</p> </div> |                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                            | <p><b>seconds</b>—Interval between advertisement packets.</p> <p><b>Range:</b> 1 through 255 seconds</p> <p><b>Default:</b> 1 second</p>                                                                                                                                                                                              |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                           | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                    |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                              | <ul style="list-style-type: none"> <li>• <i>Configuring the Advertisement Interval for the VRRP Master Router</i></li> <li>• <a href="#">fast-interval on page 2691</a></li> <li>• <i>inet6-advertise-interval</i></li> <li>• <i>version-3</i></li> </ul>                                                                             |

## asymmetric-hold-time

---


|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | asymmetric-hold-time;                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit protocols vrrp]                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | <p>Configure a VRRP master to fail over to a backup immediately—without waiting for the preemption hold time to expire—when a tracked route goes down. Otherwise, the master waits for the hold time to expire before it initiates a failover when a tracked route goes down.</p> <p>When the tracked route comes up again, the new backup (original master) router waits for the preemption hold time to expire before it reasserts mastership.</p> |
| <b>Default</b>                  | asymmetric-hold-time is disabled.                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring VRRP Preemption and Hold Time on page 2644</a></li></ul>                                                                                                                                                                                                                                                                                                                             |

## authentication-key

|                                                                                                                                                                                                                                                              |                                                                                                                                                                                                                                                                                                                                       |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                | <code>authentication-key key;</code>                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                       | <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>]</p> |
| <b>Release Information</b>                                                                                                                                                                                                                                   | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS 12.3X48-D10 for the SRX Series.</p>                                                                                                                                     |
| <b>Description</b>                                                                                                                                                                                                                                           | <p>Configure a Virtual Router Redundancy Protocol (VRRP) IPv4 authentication key. You also must specify a VRRP authentication scheme by including the <b>authentication-type</b> statement.</p> <p>All routers in the VRRP group must use the same authentication scheme and password.</p>                                            |
| <div>  <p><b>NOTE:</b> When VRRPv3 is enabled, the <b>authentication-type</b> and <b>authentication-key</b> statements cannot be configured for any VRRP groups.</p> </div> |                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                                                                                                                                                                                                                                               | <p><b>key</b>—Authentication password. For simple authentication, it can be 1 through 8 characters long. For Message Digest 5 (MD5) authentication, it can be 1 through 16 characters long. If you include spaces, enclose all characters in quotation marks (" ").</p>                                                               |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                              | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                    |
| <b>Related Documentation</b>                                                                                                                                                                                                                                 | <ul style="list-style-type: none"> <li>• <a href="#">Configuring VRRP Authentication (IPv4 Only)</a></li> <li>• <a href="#">Configuring VRRP Authentication (IPv4 Only) on page 2642</a></li> <li>• <a href="#">authentication-type on page 2688</a></li> <li>• <a href="#">version-3</a></li> </ul>                                  |

## authentication-type

---

|                                                                                                                                                                                                                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                       | <code>authentication-type authentication;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                              | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i> ],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i> ]                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>                                                                                                                                                                                                                          | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>                                                                                                                                                                                                                                  | <p>Enable Virtual Router Redundancy Protocol (VRRP) IPv4 authentication and specify the authentication scheme for the VRRP group. If you enable authentication, you must specify a password by including the <b>authentication-key</b> statement.</p> <p>All routers in the VRRP group must use the same authentication scheme and password.</p>                                                                                                                                                                                                                                                                                                                |
| <div> <b>NOTE:</b> When VRRPv3 is enabled, the <b>authentication-type</b> and <b>authentication-key</b> statements cannot be configured for any VRRP groups.</div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                                                                                                                                                                                                                                      | <p><b>authentication</b>—Authentication scheme:</p> <ul style="list-style-type: none"><li><b>simple</b>—Use a simple password. The password is included in the transmitted packet, so this method of authentication is relatively insecure.</li><li><b>md5</b>—Use the MD5 algorithm to create an encoded checksum of the packet. The encoded checksum is included in the transmitted packet. The receiving routing platform uses the authentication key to verify the packet, discarding it if the digest does not match. This algorithm provides a more secure authentication scheme.</li></ul> <p><b>Default:</b> none (no authentication is performed).</p> |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                     | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>                                                                                                                                                                                                                        | <ul style="list-style-type: none"><li>• <a href="#">Configuring VRRP Authentication (IPv4 Only)</a></li><li>• <a href="#">Configuring VRRP Authentication (IPv4 Only) on page 2642</a></li><li>• <a href="#">authentication-key on page 2687</a></li><li>• <a href="#">version-3</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                  |



## bandwidth-threshold


|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>bandwidth-threshold <i>bits-per-second</i> priority-cost <i>priority</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <a href="#">vrrp-group <i>group-id</i> track interface <i>interface-name</i></a>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <a href="#">vrrp-inet6-group <i>group-id</i> track interface <i>interface-name</i></a>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <a href="#">vrrp-group <i>group-id</i> track interface <i>interface-name</i></a>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <a href="#">vrrp-inet6-group <i>group-id</i> track interface <i>interface-name</i></a>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Specify the bandwidth threshold for Virtual Router Redundancy Protocol (VRRP) logical interface tracking.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <p><b><i>bits-per-second</i></b>—Bandwidth threshold for the tracked interface. When the bandwidth of the tracked interface drops below the specified value, the VRRP group uses the bandwidth threshold priority cost value. You can include up to five bandwidth threshold statements for each interface you track.</p> <p><b>Range:</b> 1 through 10000000000000 bits per second</p> <p><b><i>priority-cost <i>priority</i></i></b>—The value subtracted from the configured VRRP priority when the tracked interface or route is down, and forces a new master router election. The sum of all the costs for all interfaces or routes that are tracked must be less than or equal to the configured priority of the VRRP group.</p>                                                                                                                                                                            |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring a Logical Interface to Be Tracked for a VRRP Group</a></li> <li>• <a href="#">Configuring a Logical Interface to Be Tracked on page 2646</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## failover-delay

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>failover-delay <i>milliseconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit protocols vrrp]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | <p>If you configure multiple VRRP groups on an interface (using multiple VLANs), traffic for some of the groups might be briefly dropped if a failover occurs. This can happen because the new master must send gratuitous ARP replies for each VRRP group to update the ARP tables in the connected devices, and there is a short delay between each gratuitous ARP reply. Traffic sent by devices that have not yet received the gratuitous ARP reply is dropped (until the device receives the reply and learns the MAC address of the new master).</p> <p>If you configure a failover delay, the new master delays sending gratuitous ARP replies for the period that you set. This allows the new master to send the ARP replies for all of the VRRP groups simultaneously.</p> |
| <b>Options</b>                  | <b><i>milliseconds</i></b> —Specify the failover delay time, in milliseconds.<br><b>Range:</b> 50 through 2000                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Troubleshooting VRRP on page 2651</a></li><li>• <a href="#">show vrrp on page 2752</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## fast-interval

|                                                                                                                                                                                                                                                                                                                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                           | <code>fast-interval milliseconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                  | <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>]</p> |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                              | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS 12.3X48-D10 for the SRX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>                                                                                                                                                                                                                                                                                                      | <p>Configure the interval, in milliseconds, between Virtual Router Redundancy Protocol (VRRP) advertisement packets.</p> <p>All routers in the VRRP group must use the same advertisement interval.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                                                                                                                                                                                                                                                                                                          | <p><i>milliseconds</i>—Interval between advertisement packets.</p> <p><b>Range:</b> 10 through 40,950 milliseconds (range extended from 100–999 to 10–40,950 in Junos OS Release 12.2).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <div>  <p><b>NOTE:</b> When configuring VRRP for IPv4, if you have chosen not to enable VRRPv3, you cannot set a value less than 100 for <i>fast-interval</i>. Commit check fails if a value less than 100 is configured.</p> </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Default:</b> 1 second                                                                                                                                                                                                                                                                                                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                         | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                            | <ul style="list-style-type: none"> <li>• <a href="#">Configuring the Advertisement Interval for the VRRP Master Router</a></li> <li>• <a href="#">Configuring the Advertisement Interval for the VRRP Master on page 2643</a></li> <li>• <a href="#">advertise-interval on page 2685</a></li> <li>• <a href="#">advertise-interval on page 2685</a></li> <li>• <a href="#">inet6-advertise-interval</a></li> <li>• <a href="#">version-3</a></li> </ul>                                                                                                                                                                                                                                    |

## hold-time (VRRP)

---


|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>hold-time seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id preempt</i>],</code><br><code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id preempt</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id preempt</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id preempt</i>]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | In a Virtual Router Redundancy Protocol (VRRP) configuration, set the hold time before a higher-priority backup router preempts the master router.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Default</b>                  | VRRP preemption is not timed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <b><i>seconds</i></b> —Hold-time period.<br><b>Range:</b> 0 through 3600 seconds<br><b>Default:</b> 0 seconds (VRRP preemption is not timed.)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | <b>interface</b> —To view this statement in the configuration.<br><b>interface-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring a Backup Router to Preempt the VRRP Master Router</a></li><li>• <a href="#">Configuring VRRP Preemption and Hold Time on page 2644</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## interface

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>interface <i>interface-name</i> {     <b>bandwidth-threshold</b> <i>bits-per-second</i> <i>priority-cost</i> <i>priority</i>;     <b>priority-cost</b> <i>priority</i>; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <b>vrrp-group</b> <i>group-id</i> <b>track</b>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <b>vrrp-inet6-group</b> <i>group-id</i> <b>track</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <b>vrrp-group</b> <i>group-id</i> <b>track</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <b>vrrp-inet6-group</b> <i>group-id</i> <b>track</b>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p><b>bandwidth-threshold</b> statement added in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Enable logical interface tracking for a Virtual Router Redundancy Protocol (VRRP) group.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <p><b><i>interface-name</i></b>—Interface to be tracked for this VRRP group.</p> <p><b>Range:</b> 1 through 10 interfaces</p> <p>The remaining statements are described separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring a Logical Interface to Be Tracked for a VRRP Group</i></li> <li>• <a href="#">Configuring a Logical Interface to Be Tracked on page 2646</a></li> <li>• <i>Junos OS Services Interfaces Library for Routing Devices</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## preempt (VRRP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | (preempt   no-preempt) {<br>hold-time seconds;<br>}                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i> ],<br>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i> ],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i> ],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i> ]                                                                                                  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | <p>In a Virtual Router Redundancy Protocol (VRRP) configuration, determine whether or not a backup router can preempt a master router:</p> <ul style="list-style-type: none"><li>• <b>preempt</b>—Allow the master router to be preempted.</li></ul> <div> <b>NOTE:</b> By default, a higher-priority backup router can preempt a lower-priority master router.</div> <ul style="list-style-type: none"><li>• <b>no-preempt</b>—Prohibit the preemption of the master router. When <b>no-preempt</b> is configured, the backup router cannot preempt the master router even if the backup router has a higher priority.</li></ul> <p>The remaining statement is explained separately.</p> |
| <b>Default</b>                  | By default the <b>preempt</b> statement is enabled, and a higher-priority backup router preempts a lower-priority master router even if the <b>preempt</b> statement is not explicitly configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring a Backup Router to Preempt the VRRP Master Router</a></li><li>• <a href="#">Configuring VRRP Preemption and Hold Time on page 2644</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

## priority (Protocols VRRP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>priority <i>priority</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Configure a Virtual Router Redundancy Protocol (VRRP) router's priority for becoming the master default router. The router with the highest priority within the group becomes the master.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <p><b><i>priority</i></b>—Router's priority for being elected to be the master router in the VRRP group. A larger value indicates a higher priority for being elected.</p> <p><b>Range:</b> 1 through 255</p> <p><b>Default:</b> 100 (for backup routers)</p>                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Basic VRRP Support</a></li> <li>• <a href="#">Configuring Basic VRRP Support for QFX on page 2641</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |


## priority-cost (VRRP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>priority-cost <i>priority</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i> track interface <i>interface-name</i>],</code><br><code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i> track interface <i>interface-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i> track interface <i>interface-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i> track interface <i>interface-name</i>]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.<br>Statement introduced in Junos OS Release 12.2 for ACX2000 Universal Access Routers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Configure a Virtual Router Redundancy Protocol (VRRP) router's priority cost for becoming the master default router. The router with the highest priority within the group becomes the master.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <b><i>priority</i></b> —The value subtracted from the configured VRRP priority when the tracked interface or route is down to force a new master router election. The sum of all the costs for all interfaces or routes that are tracked must be less than or equal to the configured priority of the VRRP group.<br><b>Range:</b> 1 through 254                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | <code>interface</code> —To view this statement in the configuration.<br><code>interface-control</code> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring a Logical Interface to Be Tracked for a VRRP Group</a></li><li>• <a href="#">Configuring a Logical Interface to Be Tracked on page 2646</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |



## priority-hold-time

|                                                                                                                                                                                                                                                                                                                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                    | <code>priority-hold-time seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                                           | <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id track</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id track</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id track</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id track</i>]</p> |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                       | <p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                               | <p>Configure a Virtual Router Redundancy Protocol (VRRP) router's priority hold time to define the minimum length of time that must elapse between dynamic priority changes. If the dynamic priority changes because of a tracking event, the priority hold timer begins running. If another tracking event or manual configuration change occurs while the timer is running, the new dynamic priority update is postponed until the timer expires.</p>                                                                                                                                                                                                                                                            |
| <div>  <p><b>NOTE:</b> When the track feature is configured, and if VRRP should pre-empt due to the tracking interface or route transition, any configured pre-empt hold time will be ignored. VRRP master will pre-empt according to the configuration of the priority-hold time.</p> </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                                                                   | <p><b>seconds</b>—Minimum length of time that must elapse between dynamic priority changes.</p> <p><b>Range:</b> 0through 3600 seconds</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                                  | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                                                     | <ul style="list-style-type: none"> <li>• <a href="#">Configuring a Logical Interface to Be Tracked for a VRRP Group</a></li> <li>• <a href="#">Configuring a Logical Interface to Be Tracked on page 2646</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## route (Interfaces)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>route <i>prefix</i> routing-instance <i>instance-name</i> priority-cost <i>priority</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id track</i>],</code><br><code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id track</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id track</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id track</i>]</code> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0.<br>Statement introduced in Junos OS 11.3 for QFX Series.<br>Statement introduced in Junos OS 12.1 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Enable route tracking for a Virtual Router Redundancy Protocol (VRRP) group.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <p><b><i>prefix</i></b>—Route to be tracked for this VRRP group.</p> <p><b><i>priority-cost priority</i></b>—The value subtracted from the configured VRRP priority when the tracked interface or route is down, forcing a new master router election. The sum of all the costs for all interfaces or routes that are tracked must be less than or equal to the configured priority of the VRRP group.</p> <p><b><i>routing-instance instance-name</i></b>—Routing instance in which the route is to be tracked. If the route is in the default, or global, routing instance, the value for <b><i>instance-name</i></b> must be <b>default</b>.</p>                                                                                                 |
| <b>Required Privilege Level</b> | <b>interface</b> —To view this statement in the configuration.<br><b>interface-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring a Route to Be Tracked for a VRRP Group</a></li><li>• <a href="#">Configuring a Route to Be Tracked on page 2646</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

---

## startup-silent-period

---

|                                 |                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>startup-silent-period <i>seconds</i>;</code>                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit protocols vrrp]                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.                                                                                                                                        |
| <b>Description</b>              | Instruct the system to ignore the Master Down Event when an interface transitions from the down state to the up state. This statement is used to avoid incorrect error alarms caused by the delay or interruption of incoming Virtual Router Redundancy Protocol (VRRP) advertisement packets during the interface startup phase. |
| <b>Options</b>                  | <b><i>seconds</i></b> —Number of seconds for the startup period.<br><b>Default:</b> 4 seconds<br><b>Range:</b> 1 through 2000 seconds                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the Startup Period for VRRP Operations</i></li><li>• <a href="#">Configuring the Startup Period for VRRP Operations on page 2643</a></li></ul>                                                                                                                             |

## traceoptions

---

**Syntax**    traceoptions {  
              file <filename> <files number> <match *regular-expression*> <microsecond-stamp>  
                  <size size> <world-readable | no-world-readable>;  
              flag *flag*;  
              no-remote-trace;  
              }

**Hierarchy Level**    [edit protocols vrrp]

**Release Information**    Statement introduced in Junos OS Release 11.3 for the QFX Series.  
                              Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.

**Description**    Define tracing operations for the Virtual Router Redundancy Protocol (VRRP) process.

To specify more than one tracing operation, include multiple **flag** statements.

By default, VRRP logs the error, dcd configuration, and routing socket events in a file in the directory **/var/log**.



**NOTE:** The traceoptions statement is not supported on a QFabric system.

---

**Default**    If you do not include this statement, no VRRP-specific tracing operations are performed.

**Options**    **filename filename**—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. By default, VRRP tracing output is placed in the file **vrrpd**.

**files number**—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. When the maximum number is reached, the oldest trace file is overwritten.

**Range:** 0 through 4,294,967,296 files

**Default:** 3 files

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

**flag flag**—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. These are the VRRP-specific tracing options:

- **all**—All VRRP tracing operations
- **database**—Database changes
- **general**—General events

- **interfaces**—Interface changes
- **normal**—Normal events
- **packets**—Packets sent and received
- **state**—State transitions
- **timer**—Timer events

**match *regex***—(Optional) Refine the output to include only those lines that match the given regular expression.

**microsecond-stamp**—(Optional) Provide a timestamp with microsecond granularity.

**no-world-readable**—Restrict users from reading the log file.

**size *size***—(Optional) Maximum size of each trace file, in kilobytes, megabytes, or gigabytes. When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file.0***. When the ***trace-file*** again reaches its maximum size, ***trace-file.0*** is renamed ***trace-file.1*** and ***trace-file*** is renamed ***trace-file.0***. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

**Syntax:** *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

**Range:** 10 KB through the maximum file size supported on your routing platform

**Default:** 1 MB

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

**world-readable**—Allow users to read the log file.

|                                 |                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.                             |
|                                 | interface-control—To add this statement to the configuration.                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Tracing VRRP Operations</i></li> </ul> |

## track (VRRP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>track {<br/>  interface <i>interface-name</i> {<br/>    bandwidth-threshold <i>bits-per-second</i> priority-cost <i>priority</i>;<br/>    priority-cost <i>priority</i>;<br/>  }<br/>  priority-hold-time <i>seconds</i>;<br/>  route <i>prefix/prefix-length</i> routing-instance <i>instance-name</i> priority-cost <i>priority</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <pre>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i><br/>  vrrp-group <i>group-id</i>],<br/>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i><br/>  vrrp-inet6-group <i>group-id</i>],<br/>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i><br/>  family inet address <i>address</i> vrrp-group <i>group-id</i>],<br/>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i><br/>  family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i>]</pre> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p><b>priority-hold-time</b> statement added in Junos OS Release 8.1.</p> <p><b>route</b> statement added in Junos OS Release 9.0.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Enable logical interface tracking, route tracking, or both, for a Virtual Router Redundancy Protocol (VRRP) group.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | The remaining statements are described separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring a Logical Interface to Be Tracked for a VRRP Group</a></li><li>• <a href="#">Configuring a Route to Be Tracked for a VRRP Group</a></li><li>• <a href="#">Configuring a Logical Interface to Be Tracked on page 2646</a></li><li>• <a href="#">Configuring a Route to Be Tracked on page 2646</a></li></ul>                                                                                                                                                                                                                                                                                                                                   |

## virtual-address

---

|                                 |                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>virtual-address [ <i>addresses</i> ];</code>                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i> ],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.                                                                                                                                   |
| <b>Description</b>              | Configure the addresses of the virtual routers in a Virtual Router Redundancy Protocol (VRRP) IPv4 or IPv6 group. You can configure up to eight addresses.                                                                                                                                                                   |
| <b>Options</b>                  | <b><i>addresses</i></b> —Addresses of one or more virtual routers. Do not include a prefix length. If the address is the same as the interface's physical address, the interface becomes the master virtual router for the group.                                                                                            |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Basic VRRP Support</i></li> <li>• <a href="#">Configuring Basic VRRP Support for QFX on page 2641</a></li> </ul>                                                                                                                                                     |

## vrrp-group

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> vrrp-group <i>group-id</i> {   (<b>accept-data</b>   <b>no-accept-data</b>);   <b>advertise-interval</b> <i>seconds</i>;   <b>advertisements-threshold</b> <i>number</i>;   <b>authentication-key</b> <i>key</i>;   <b>authentication-type</b> <i>authentication</i>;   <b>fast-interval</b> <i>milliseconds</i>;   (<b>preempt</b>   <b>no-preempt</b>) {     <b>hold-time</b> <i>seconds</i>;   }   <b>priority</b> <i>number</i>;   <b>track</b> {     <b>interface</b> <i>interface-name</i> {       <b>bandwidth-threshold</b> <i>bits-per-second</i> <b>priority-cost</b> <i>priority</i>;       <b>priority-cost</b> <i>priority</i>;     }     <b>priority-hold-time</b> <i>seconds</i>;     <b>route</b> <i>prefix/prefix-length</i> <b>routing-instance</b> <i>instance-name</i> <b>priority-cost</b> <i>priority</i>;   }   <b>virtual-address</b> [ <i>addresses</i> ];   <b>vrrp-inherit-from</b> <i>vrrp-group</i>; } </pre> |
| <b>Hierarchy Level</b>          | <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | <p>Configure a Virtual Router Redundancy Protocol (VRRP) IPv4 group. As of Junos OS Release 13.2, VRRP nonstop active routing (NSR) is enabled only when you configure the <b>nonstop-routing</b> statement at the [edit routing-options] or [edit logical system <i>logical-system-name</i> routing-options hierarchy level.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <p><b><i>group-id</i></b>—VRRP group identifier. If you enable MAC source address filtering on the interface, you must include the virtual MAC address in the list of source MAC addresses that you specify in the <b>source-address-filter</b> statement. MAC addresses ranging from 00:00:5e:00:01:00 through 00:00:5e:00:01:ff are reserved for VRRP, as defined in RFC 2338. The VRRP group number must be the decimal equivalent of the last hexadecimal byte of the virtual MAC address.</p> <p><b>Range:</b> 0 through 255</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |



- Related Documentation**
- *Configuring Basic VRRP Support*
  - *Configuring VRRP*
  - [Configuring Basic VRRP Support for QFX on page 2641](#)
  - [Example: Configuring VRRP for Load Sharing on page 2636](#)
  - *vrrp-inet6-group*
  - [nonstop-routing on page 2676](#)



# Operational Mode Commands (Graceful Restart)

- [Verifying Graceful Restart Operation on page 2707](#)
- [show bgp neighbor](#)
- [show log](#)
- [show \(ospf | ospf3\) overview](#)

## Verifying Graceful Restart Operation

---

This topic contains the following sections:

- [Graceful Restart Operational Mode Commands on page 2707](#)
- [Verifying BGP Graceful Restart on page 2708](#)
- [Verifying IS-IS and OSPF Graceful Restart on page 2708](#)
- [Verifying CCC and TCC Graceful Restart on page 2709](#)

## Graceful Restart Operational Mode Commands

To verify proper operation of graceful restart, use the following commands:

- **show bgp neighbor** (for BGP graceful restart)
- **show log** (for IS-IS and OSPF/OSPFv3 graceful restart)
- **show (ospf | ospfv3) overview** (for OSPF/OSPFv3 graceful restart)
- **show rsvp neighbor detail** (for RSVP graceful restart—helper router)
- **show rsvp version** (for RSVP graceful restart—restarting router)
- **show ldp session detail** (for LDP graceful restart)
- **show connections** (for CCC and TCC graceful restart)
- **show route instance detail** (for Layer 3 VPN graceful restart and for any protocols using graceful restart in a routing instance)
- **show route protocol l2vpn** (for Layer 2 VPN graceful restart)

For more information about these commands and a description of their output fields, see the [CLI Explorer](#).

## Verifying BGP Graceful Restart

To view graceful restart information for BGP sessions, use the **show bgp neighbor** command:

```
user@PE1> show bgp neighbor 192.255.10.1
Peer: 192.255.10.1+179 AS 64595 Local: 192.255.5.1+1106 AS 64595
 Type: Internal State: Established Flags: <>
 Last State: OpenConfirm Last Event: RecvKeepAlive
 Last Error: None
 Export: [static]
 Options:<Preference LocalAddress HoldTime GracefulRestart Damping PeerAS Refresh>

 Local Address: 192.255.5.1 Holdtime: 90 Preference: 170
 IPsec SA Name: hope
 Number of flaps: 0
 Peer ID: 192.255.10.1 Local ID: 192.255.5.1 Active Holdtime: 90
 Keepalive Interval: 30
 NLRI for restart configured on peer: inet-unicast
 NLRI advertised by peer: inet-unicast
 NLRI for this session: inet-unicast
 Peer supports Refresh capability (2)
 Restart time configured on the peer: 180
 Stale routes from peer are kept for: 180
 Restart time requested by this peer: 300
 NLRI that peer supports restart for: inet-unicast
 NLRI that peer saved forwarding for: inet-unicast
 NLRI that restart is negotiated for: inet-unicast
 NLRI of received end-of-rib markers: inet-unicast
 NLRI of all end-of-rib markers sent: inet-unicast
 Table inet.0 Bit: 10000
 RIB State: restart is complete
 Send state: in sync
 Active prefixes: 0
 Received prefixes: 0
 Suppressed due to damping: 0
 Last traffic (seconds): Received 19 Sent 19 Checked 19
 Input messages: Total 2 Updates 1 Refreshes 0 Octets 42
 Output messages: Total 3 Updates 0 Refreshes 0 Octets 116
 Output Queue[0]: 0
```

## Verifying IS-IS and OSPF Graceful Restart

To view graceful restart information for IS-IS and OSPF, configure traceoptions (see [“Tracking Graceful Restart Events” on page 2586](#)).

Here is the output of a traceoptions log from an OSPF restarting router:

```
Oct 8 05:20:12 Restart mode - sending grace lsas
Oct 8 05:20:12 Restart mode - estimated restart duration timer triggered
Oct 8 05:20:13 Restart mode - Sending more grace lsas
```

Here is the output of a traceoptions log from an OSPF helper router:

```
Oct 8 05:20:14 Helper mode for neighbor 192.255.5.1
Oct 8 05:20:14 Received multiple grace lsa from 192.255.5.1
```

## Verifying CCC and TCC Graceful Restart

To view graceful restart information for CCC and TCC connections, use the **show connections** command. The following example assumes four remote interface CCC connections between CE1 and CE2:

```
user@PE1> show connections
CCC and TCC connections [Link Monitoring On]
Legend for status (St)
UN -- uninitialized
NP -- not present
WE -- wrong encapsulation
DS -- disabled
Dn -- down
-> -- only outbound conn is up
<- -- only inbound conn is up
Up -- operational
RmtDn -- remote CCC down
Restart -- restarting

Legend for connection types
if-sw: interface switching
rmt-if: remote interface switching
lsp-sw: LSP switching

Legend for circuit types
intf -- interface
tlsp -- transmit LSP
rlsp -- receive LSP
```

### CCC Graceful restart : Restarting

| Connection/Circuit | Type   | St      | Time last up | # Up trans |
|--------------------|--------|---------|--------------|------------|
| CE1-CE2-0          | rmt-if | Restart | -----        | 0          |
| fe-1/1/0.0         | intf   | Up      |              |            |
| PE1-PE2-0          | tlsp   | Up      |              |            |
| PE2-PE1-0          | rlsp   | Up      |              |            |
| CE1-CE2-1          | rmt-if | Restart | -----        | 0          |
| fe-1/1/0.1         | intf   | Up      |              |            |
| PE1-PE2-1          | tlsp   | Up      |              |            |
| PE2-PE1-1          | rlsp   | Up      |              |            |
| CE1-CE2-2          | rmt-if | Restart | -----        | 0          |
| fe-1/1/0.2         | intf   | Up      |              |            |
| PE1-PE2-2          | tlsp   | Up      |              |            |
| PE2-PE1-2          | rlsp   | Up      |              |            |
| CE1-CE2-3          | rmt-if | Restart | -----        | 0          |
| fe-1/1/0.3         | intf   | Up      |              |            |
| PE1-PE2-3          | tlsp   | Up      |              |            |
| PE2-PE1-3          | rlsp   | Up      |              |            |

- Related Documentation**
- [Graceful Restart Concepts on page 2579](#)
  - *Configuring Graceful Restart for QFabric Systems*

## show bgp neighbor

---

|                                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                                        | <a href="#">Syntax on page 2710</a><br><a href="#">Syntax (EX Series Switch, QFX Series, and OCX Series) on page 2710</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Syntax</b>                                                | <pre>show bgp neighbor &lt;exact-instance <i>instance-name</i>&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;neighbor-address&gt; &lt;orf (detail   <i>neighbor-address</i>)</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Syntax (EX Series Switch, QFX Series, and OCX Series)</b> | <pre>show bgp neighbor &lt;instance <i>instance-name</i>&gt; &lt;exact-instance <i>instance-name</i>&gt; &lt;neighbor-address&gt; &lt;orf (<i>neighbor-address</i>   detail)</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>                                   | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.<br>Command introduced in Junos OS Release 14.1x53-D20 for the OCX Series.<br><b>orf</b> option introduced in Junos OS Release 9.2.<br><b>exact-instance</b> option introduced in Junos OS Release 11.4.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>                                           | Display information about BGP peers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                                               | <p><b>none</b>—Display information about all BGP peers.</p> <p><b>exact-instance <i>instance-name</i></b>—(Optional) Display information for the specified instance only.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information about BGP peers for all routing instances whose name begins with this string (for example, <b>cust1</b>, <b>cust11</b>, and <b>cust111</b> are all displayed when you run the <b>show bgp neighbor instance cust1</b> command).</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>neighbor-address</b>—(Optional) Display information for only the BGP peer at the specified IP address.</p> <p><b>orf (detail   <i>neighbor-address</i>)</b>—(Optional) Display outbound route-filtering information for all BGP peers or only for the BGP peer at the specified IP address. The default is to display brief output. Use the <b>detail</b> option to display detailed output.</p> |
| <b>Additional Information</b>                                | For information about the <b>local-address</b> , <b>nlri</b> , <b>hold-time</b> , and <b>preference</b> statements, see the <i>Junos OS Routing Protocols Library for Routing Devices</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b>                              | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

- Related Documentation**
- [clear bgp neighbor on page 4093](#)
- List of Sample Output**
- [show bgp neighbor on page 2718](#)
  - [show bgp neighbor \(CLNS\) on page 2719](#)
  - [show bgp neighbor \(Layer 2 VPN\) on page 2720](#)
  - [show bgp neighbor \(Layer 3 VPN\) \(Not supported on the OCX Series.\) on page 2722](#)
  - [show bgp neighbor neighbor-address on page 2723](#)
  - [show bgp neighbor neighbor-address on page 2723](#)
  - [show bgp neighbor neighbor-address \(BGP Graceful Restart Enabled\) on page 2724](#)
  - [show bgp neighbor neighbor-address \(BGP Long-Lived Graceful Restart\) on page 2725](#)
  - [show bgp neighbor orf neighbor-address detail on page 2725](#)
- Output Fields** Table 238 describes the output fields for the **show bgp neighbor** command. Output fields are listed in the approximate order in which they appear.

Table 238: show bgp neighbor Output Fields

| Field Name | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Peer       | Address of the BGP neighbor. The address is followed by the neighbor port number.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| AS         | AS number of the peer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Local      | Address of the local routing device. The address is followed by the peer port number.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Type       | Type of peer: <b>Internal</b> or <b>External</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| State      | <p>Current state of the BGP session:</p> <ul style="list-style-type: none"> <li>• <b>Active</b>—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an Open message.</li> <li>• <b>Connect</b>—BGP is waiting for the transport protocol connection to be completed.</li> <li>• <b>Established</b>—The BGP session has been established, and the peers are exchanging update messages.</li> <li>• <b>Idle</b>—This is the first stage of a connection. BGP is waiting for a Start event.</li> <li>• <b>OpenConfirm</b>—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message.</li> <li>• <b>OpenSent</b>—BGP has sent an open message and is waiting to receive an open message from the peer.</li> <li>• <b>route reflector client</b>—The BGP session is established with a route reflector client.</li> </ul> |

Table 238: show bgp neighbor Output Fields (*continued*)

| Field Name        | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Flags</b>      | <p>Internal BGP flags:</p> <ul style="list-style-type: none"> <li>• <b>Aggregate Label</b>—BGP has aggregated a set of incoming labels (labels received from the peer) into a single forwarding label.</li> <li>• <b>CleanUp</b>—The peer session is being shut down.</li> <li>• <b>Delete</b>—This peer has been deleted.</li> <li>• <b>Idled</b>—This peer has been permanently idled.</li> <li>• <b>ImportEval</b>—At the last commit operation, this peer was identified as needing to reevaluate all received routes.</li> <li>• <b>Initializing</b>—The peer session is initializing.</li> <li>• <b>SendRtn</b>—Messages are being sent to the peer.</li> <li>• <b>Sync</b>—This peer is synchronized with the rest of the peer group.</li> <li>• <b>RSync</b>—This peer in the backup Routing Engine is synchronized with the BGP peer in the master Routing Engine for nonstop active routing.</li> <li>• <b>TryConnect</b>—Another attempt is being made to connect to the peer.</li> <li>• <b>Unconfigured</b>—This peer is not configured.</li> <li>• <b>WriteFailed</b>—An attempt to write to this peer failed.</li> </ul>                                                                                                       |
| <b>Last state</b> | <p>Previous state of the BGP session:</p> <ul style="list-style-type: none"> <li>• <b>Active</b>—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an Open message.</li> <li>• <b>Connect</b>—BGP is waiting for the transport protocol connection to be completed.</li> <li>• <b>Established</b>—The BGP session has been established, and the peers are exchanging update messages.</li> <li>• <b>Idle</b>—This is the first stage of a connection. BGP is waiting for a Start event.</li> <li>• <b>OpenConfirm</b>—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message.</li> <li>• <b>OpenSent</b>—BGP has sent an open message and is waiting to receive an open message from the peer.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Last event</b> | <p>Last activity that occurred in the BGP session:</p> <ul style="list-style-type: none"> <li>• <b>Closed</b>—The BGP session closed.</li> <li>• <b>ConnectRetry</b>—The transport protocol connection failed, and BGP is trying again to connect.</li> <li>• <b>HoldTime</b>—The session ended because the hold timer expired.</li> <li>• <b>KeepAlive</b>—The local routing device sent a BGP keepalive message to the peer.</li> <li>• <b>Open</b>—The local routing device sent a BGP open message to the peer.</li> <li>• <b>OpenFail</b>—The local routing device did not receive an acknowledgment of a BGP open message from the peer.</li> <li>• <b>RecvKeepAlive</b>—The local routing device received a BGP keepalive message from the peer.</li> <li>• <b>RecvNotify</b>—The local routing device received a BGP notification message from the peer.</li> <li>• <b>RecvOpen</b>—The local routing device received a BGP open message from the peer.</li> <li>• <b>RecvUpdate</b>—The local routing device received a BGP update message from the peer.</li> <li>• <b>Start</b>—The peering session started.</li> <li>• <b>Stop</b>—The peering session stopped.</li> <li>• <b>TransportError</b>—A TCP error occurred.</li> </ul> |



Table 238: show bgp neighbor Output Fields (*continued*)

| Field Name                                                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Last error                                                     | <p>Last error that occurred in the BGP session:</p> <ul style="list-style-type: none"> <li>• <b>Cease</b>—An error occurred, such as a version mismatch, that caused the session to close.</li> <li>• <b>Finite State Machine Error</b>—In setting up the session, BGP received a message that it did not understand.</li> <li>• <b>Hold Time Expired</b>—The session's hold time expired.</li> <li>• <b>Message Header Error</b>—The header of a BGP message was malformed.</li> <li>• <b>Open Message Error</b>—A BGP open message contained an error.</li> <li>• <b>None</b>—No errors occurred in the BGP session.</li> <li>• <b>Update Message Error</b>—A BGP update message contained an error.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Export                                                         | Name of the export policy that is configured on the peer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Import                                                         | Name of the import policy that is configured on the peer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Options                                                        | <p>Configured BGP options:</p> <ul style="list-style-type: none"> <li>• <b>AddressFamily</b>—Configured address family: <b>inet</b> or <b>inet-vpn</b>.</li> <li>• <b>AdvertiseBGPStatic</b>—Configured BGP static routes are advertised.</li> <li>• <b>AuthKeyChain</b>—Authentication key change is enabled.</li> <li>• <b>DropPathAttributes</b>—Certain path attributes are configured to be dropped from neighbor updates during inbound processing.</li> <li>• <b>GracefulRestart</b>—Graceful restart is configured.</li> <li>• <b>HoldTime</b>—Hold time configured with the <b>hold-time</b> statement. The hold time is three times the interval at which keepalive messages are sent.</li> <li>• <b>IgnorePathAttributes</b>—Certain path attributes are configured to be ignored in neighbor updates during inbound processing.</li> <li>• <b>Local Address</b>—Address configured with the <b>local-address</b> statement.</li> <li>• <b>LLGR</b>—BGP long-lived graceful restart capability is configured.</li> <li>• <b>LLGRHelperDisabled</b>—BGP long-lived graceful restart is completely disabled for a neighbor.</li> <li>• <b>Multihop</b>—Allow BGP connections to external peers that are not on a directly connected network.</li> <li>• <b>NLRI</b>—Configured MBGP state for the BGP group: <b>multicast</b>, <b>unicast</b>, or both if you have configured <b>nlri any</b>.</li> <li>• <b>Peer AS</b>—Configured peer autonomous system (AS).</li> <li>• <b>Preference</b>—Preference value configured with the <b>preference</b> statement.</li> <li>• <b>Refresh</b>—Configured to refresh automatically when the policy changes.</li> <li>• <b>Rib-group</b>—Configured routing table group.</li> </ul> |
| Path-attributes dropped                                        | Path attribute codes that are dropped from neighbor updates.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Path-attributes ignored                                        | Path attribute codes that are ignored during neighbor updates.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Peer does not support LLGR Restarter or Receiver functionality | BGP neighbor does not support long-lived graceful restart (LLGR) restarter mode completely.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

Table 238: show bgp neighbor Output Fields (*continued*)

| Field Name                                         | Field Description                                                                                                                                                                                                                                                                                                                                                               |
|----------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Peer does not support LLGR Restarter functionality | BGP neighbor does not support long-lived graceful restart (LLGR) restarter mode for any family.                                                                                                                                                                                                                                                                                 |
| Authentication key change                          | (appears only if the <b>authentication-keychain</b> statement has been configured) Name of the authentication keychain enabled.                                                                                                                                                                                                                                                 |
| Authentication algorithm                           | (appears only if the <b>authentication-algorithm</b> statement has been configured) Type of authentication algorithm enabled: <b>hmac</b> or <b>md5</b> .                                                                                                                                                                                                                       |
| Address families configured                        | Names of configured address families for the VPN.                                                                                                                                                                                                                                                                                                                               |
| BGP-Static Advertisement Policy                    | Name of the bgp static policy that is configured on the peer.                                                                                                                                                                                                                                                                                                                   |
| Local Address                                      | Address of the local routing device.                                                                                                                                                                                                                                                                                                                                            |
| Remove-private options                             | Options associated with the <b>remove-private</b> statement.                                                                                                                                                                                                                                                                                                                    |
| Holdtime                                           | Hold time configured with the <b>hold-time</b> statement. The hold time is three times the interval at which keepalive messages are sent.                                                                                                                                                                                                                                       |
| Flags for NLRI inet-label-unicast                  | Flags related to labeled-unicast: <ul style="list-style-type: none"> <li>• <b>TrafficStatistics</b>—Collection of statistics for labeled-unicast traffic is enabled.</li> </ul>                                                                                                                                                                                                 |
| Traffic statistics                                 | Information about labeled-unicast traffic statistics: <ul style="list-style-type: none"> <li>• <b>Options</b>—Options configured for collecting statistics about labeled-unicast traffic.</li> <li>• <b>File</b>—Name and location of statistics log files.</li> <li>• <b>size</b>—Size of all the log files, in bytes.</li> <li>• <b>files</b>—Number of log files.</li> </ul> |
| Traffic Statistics Interval                        | Time between sample periods for labeled-unicast traffic statistics, in seconds.                                                                                                                                                                                                                                                                                                 |
| Preference                                         | Preference value configured with the <b>preference</b> statement.                                                                                                                                                                                                                                                                                                               |
| Outbound Timer                                     | Time for which the route is available in Junos OS routing table before it is exported to BGP. This field is displayed in the output only if the <b>out-delay</b> parameter is configured to a non-zero value.                                                                                                                                                                   |
| Number of flaps                                    | Number of times the BGP session has gone down and then come back up.                                                                                                                                                                                                                                                                                                            |
| Peer ID                                            | Router identifier of the peer.                                                                                                                                                                                                                                                                                                                                                  |
| Group index                                        | Index number for the BGP peer group. The index number differentiates between groups when a single BGP group is split because of different configuration options at the group and peer levels.                                                                                                                                                                                   |

Table 238: show bgp neighbor Output Fields (*continued*)

| Field Name                                                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Peer index                                                     | Index that is unique within the BGP group to which the peer belongs.                                                                                                                                                                                                                                                                                                                                                                                  |
| Local ID                                                       | Router identifier of the local routing device.                                                                                                                                                                                                                                                                                                                                                                                                        |
| Local Interface                                                | Name of the interface on the local routing device.                                                                                                                                                                                                                                                                                                                                                                                                    |
| Active holdtime                                                | Hold time that the local routing device negotiated with the peer.                                                                                                                                                                                                                                                                                                                                                                                     |
| Keepalive Interval                                             | Keepalive interval, in seconds.                                                                                                                                                                                                                                                                                                                                                                                                                       |
| BFD                                                            | Status of BFD failure detection.                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Local Address                                                  | Name of directly connected interface over which direct EBGp peering is established.                                                                                                                                                                                                                                                                                                                                                                   |
| NLRI and times for LLGR configured on peer                     | <p>Names of address families and stale time for BGP long-lived graceful restart configured on the BGP peer or neighbor.</p> <p>Times are displayed using the routing protocol daemon (rpd) %#OT format:</p> <p><b>&lt;weeks&gt;w&lt;days&gt;d &lt;hours&gt;:&lt;minutes&gt;:&lt;seconds&gt;</b></p> <p>Zero leading elements are omitted, for example, a value less than one week do not include the weeks.</p>                                       |
| NLRI and times that peer supports LLGR Restarter for           | <p>Names of address families and stale time that the BGP peer supports for restarter mode for BGP long-lived graceful restart.</p> <p>Times are displayed using the routing protocol daemon (rpd) %#OT format:</p> <p><b>&lt;weeks&gt;w&lt;days&gt;d &lt;hours&gt;:&lt;minutes&gt;:&lt;seconds&gt;</b></p> <p>Zero leading elements are omitted, for example, a value less than one week do not include the weeks.</p>                                |
| NLRI that peer saved LLGR forwarding for                       | Name of the address family for which the BGP peer saved BGP long-lived graceful restart forwarding.                                                                                                                                                                                                                                                                                                                                                   |
| Graceful Restart Details                                       | Amount of time that is remaining until LLGR expires and the time remaining on the GR stale timer, along with RIB details, are displayed while LLGR receiver mode is active (a peer that negotiated LLGR has disconnected and not yet reconnected)                                                                                                                                                                                                     |
| NLRI we are holding stale routes for                           | Names of address families (NLRIs) for which that stale routes are held or preserved when BGP graceful restart receiver mode is active for a neighbor.                                                                                                                                                                                                                                                                                                 |
| Time until end-of-rib is assumed for stale routes              | <p>Amount of time remaining on the stale timer until which end-of-RIB (EoR) markers are assumed when BGP graceful restart receiver mode is active for a neighbor.</p> <p>Time is displayed in Coordinated Universal Time (UTC) format (YYYY-MM-DD-HH:MM:SS). Note that the stale timer display ('Time until end-of-rib is assumed') is also present when a session is active, but the neighbor as not yet sent all of the end-of-rib indications.</p> |
| Time until stale routes are deleted or become long-lived stale | Amount of time up to which stale routes are deleted or become long-lived stale routes when BGP graceful restart receiver mode is active for a neighbor.                                                                                                                                                                                                                                                                                               |

Table 238: show bgp neighbor Output Fields (*continued*)

| Field Name                                    | Field Description                                                                                                                                                                                           |
|-----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| NLRI for restart configured on peer           | Names of address families configured for restart.                                                                                                                                                           |
| NLRI advertised by peer                       | Address families supported by the peer: <b>unicast</b> or <b>multicast</b> .                                                                                                                                |
| NLRI for this session                         | Address families being used for this session.                                                                                                                                                               |
| Peer supports Refresh capability              | Remote peer's ability to send and request full route table readvertisement (route refresh capability). For more information, see RFC 2918, <i>Route Refresh Capability for BGP-4</i> .                      |
| Restart time configured on peer               | Configured time allowed for restart on the neighbor.                                                                                                                                                        |
| Stale routes from peer are kept for           | When graceful restart is negotiated, the maximum time allowed to hold routes from neighbors after the BGP session has gone down.                                                                            |
| Peer does not support Restarter functionality | Graceful restart restarter-mode is disabled on the peer.                                                                                                                                                    |
| Peer does not support Receiver functionality  | Graceful restart helper-mode is disabled on the peer.                                                                                                                                                       |
| Restart time requested by this peer           | Restart time requested by this neighbor during capability negotiation.                                                                                                                                      |
| Restart flag received from the peer           | When this field appears, the BGP speaker has restarted (Restarting), and this peer should not wait for the <b>end-of-rib</b> marker from the speaker before advertising routing information to the speaker. |
| NLRI that peer supports restart for           | Neighbor supports graceful restart for this address family.                                                                                                                                                 |
| NLRI peer can save forwarding state           | Neighbor supporting this address family saves all forwarding states.                                                                                                                                        |
| NLRI that peer saved forwarding for           | Neighbor saves all forwarding states for this address family.                                                                                                                                               |
| NLRI that restart is negotiated for           | Router supports graceful restart for this address family.                                                                                                                                                   |
| NLRI of received end-of-rib markers           | Address families for which end-of-routing-table markers are received from the neighbor.                                                                                                                     |
| NLRI of all end-of-rib markers sent           | Address families for which end-of-routing-table markers are sent to the neighbor.                                                                                                                           |
| Peer supports 4 byte AS extension (peer-as 1) | Peer understands 4-byte AS numbers in BGP messages. The peer is running Junos OS Release 9.1 or later.                                                                                                      |

Table 238: show bgp neighbor Output Fields (*continued*)

| Field Name                                                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| NLRIs for which peer can receive multiple paths            | <p>Appears in the command output of the local router if the downstream peer is configured to receive multiple BGP routes to a single destination, instead of only receiving the active route.</p> <p>Possible value is <b>inet-unicast</b>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| NLRIs for which peer can send multiple paths: inet-unicast | <p>Appears in the command output of the local router if the upstream peer is configured to send multiple BGP routes to a single destination, instead of only sending the active route.</p> <p>Possible value is <b>inet-unicast</b>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Table <i>inet.number</i>                                   | <p>Information about the routing table:</p> <ul style="list-style-type: none"> <li>• <b>RIB State</b>—BGP is in the graceful restart process for this routing table: <b>restart is complete</b> or <b>restart in progress</b>.</li> <li>• <b>Bit</b>—Number that represents the entry in the routing table for this peer.</li> <li>• <b>Send state</b>—State of the BGP group: <b>in sync</b>, <b>not in sync</b>, or <b>not advertising</b>.</li> <li>• <b>Active prefixes</b>—Number of prefixes received from the peer that are active in the routing table.</li> <li>• <b>Received prefixes</b>—Total number of prefixes from the peer, both active and inactive, that are in the routing table.</li> <li>• <b>Accepted prefixes</b>—Total number of prefixes from the peer that have been accepted by a routing policy.</li> <li>• <b>Suppressed due to damping</b>—Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.</li> </ul> |
| Last traffic (seconds)                                     | Last time any traffic was received from the peer or sent to the peer, and the last time the local routing device checked.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Input messages                                             | Messages that BGP has received from the receive socket buffer, showing the total number of messages, number of update messages, number of times a policy is changed and refreshed, and the buffer size in octets. The buffer size is 16 KB.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Output messages                                            | Messages that BGP has written to the transmit socket buffer, showing the total number of messages, number of update messages, number of times a policy is changed and refreshed, and the buffer size in octets. The buffer size is 16 KB.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Input dropped path attributes                              | <p>Information about dropped path attributes:</p> <ul style="list-style-type: none"> <li>• <b>Code</b>—Path attribute code.</li> <li>• <b>Count</b>—Path attribute count.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Input ignored path attributes                              | <p>Information about ignored path attributes:</p> <ul style="list-style-type: none"> <li>• <b>Code</b>—Path attribute code.</li> <li>• <b>Count</b>—Path attribute count.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

Table 238: show bgp neighbor Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Output queue</b>            | <p>Number of BGP packets that are queued to be transmitted to a particular neighbor for a particular routing table. Output queue <b>0</b> is for unicast NLRIs, and queue <b>1</b> is for multicast NLRIs.</p> <p>It also specifies the routing table name and the NLRI that the table was advertised through, in the format <b>(routing table name, NLRI)</b>.</p> <p><b>NOTE:</b> The output queue of routing tables that are not advertised, will only show up at <b>extensive</b> output level.</p> |
| <b>Trace options</b>           | Configured tracing of BGP protocol packets and operations.                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Trace file</b>              | Name of the file to receive the output of the tracing operation.                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Filter Updates rcv</b>      | <p>(<b>orf</b> option only) Number of outbound-route filters received for each configured address family.</p> <p><b>NOTE:</b> The counter is cumulative. For example, the counter is increased after the remote peer either resends or clears the outbound route filtering prefix list.</p>                                                                                                                                                                                                             |
| <b>Immediate</b>               | <p>(<b>orf</b> option only) Number of route updates received with the immediate flag set. The immediate flag indicates that the BGP peer should readvertise the updated routes.</p> <p><b>NOTE:</b> The counter is cumulative. For example, the counter is increased after the remote peer either resends or clears the outbound route filtering prefix list.</p>                                                                                                                                       |
| <b>Filter</b>                  | ( <b>orf</b> option only) Type of prefix filter received: <b>prefix-based</b> or <b>extended-community</b> .                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Received filter entries</b> | ( <b>orf</b> option only) List of received filters displayed.                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>seq</b>                     | ( <b>orf</b> option only) Numerical order assigned to this prefix entry among all the received outbound route filter prefix entries.                                                                                                                                                                                                                                                                                                                                                                    |
| <b>prefix</b>                  | ( <b>orf</b> option only) Address for the prefix entry that matches the filter.                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>minlength</b>               | ( <b>orf</b> option only) Minimum prefix length, in bits, required to match this prefix.                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>maxlength</b>               | ( <b>orf</b> option only) Maximum prefix length, in bits, required to match this prefix.                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>match</b>                   | ( <b>orf</b> option only) For this prefix match, whether to <b>permit</b> or <b>deny</b> route updates.                                                                                                                                                                                                                                                                                                                                                                                                 |

## Sample Output

### show bgp neighbor

```

user@host > show bgp neighbor
Peer: 10.255.7.250+179 AS 10 Local: 10.255.7.248+63740 AS 10
 Type: Internal State: Established Flags: <Sync>
 Last State: OpenConfirm Last Event: RecvKeepAlive
 Last Error: None
 Export: [redist_static]
 Options: <Preference LocalAddress PeerAS Refresh>
 Options: <AdvertiseBGPStatic>

```

```

Local Address: 10.255.7.248 Holdtime: 90 Preference: 170 Outbound Timer: 50
Number of flaps: 0
Peer ID: 10.255.7.250 Local ID: 10.255.7.248 Active Holdtime: 90
Keepalive Interval: 30 Group index: 0 Peer index: 0
BFD: disabled, down
NLRI for restart configured on peer: inet-unicast
NLRI advertised by peer: inet-unicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Stale routes from peer are kept for: 300
Peer does not support Restarter functionality
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Peer supports 4 byte AS extension (peer-as 10)
Peer does not support Addpath
Table inet.0 Bit: 10000
 RIB State: BGP restart is complete
 Send state: in sync
 Active prefixes: 1
 Received prefixes: 1
 Accepted prefixes: 1
 Suppressed due to damping: 0
 Advertised prefixes: 1
Last traffic (seconds): Received 9 Sent 5 Checked 5
Input messages: Total 36 Updates 2 Refreshes 0 Octets 718
Output messages: Total 37 Updates 1 Refreshes 0 Octets 796
Output Queue[0]: 0 (inet.0, inet-unicast)

Peer: 10.255.162.214+52193 AS 100 Local: 10.255.167.205+179 AS 100
 Type: Internal State: Established (route reflector client)Flags: <Sync>
 Last State: OpenConfirm Last Event: RecvKeepAlive
 Last Error: None
 Options: <Preference LocalAddress Cluster AddressFamily Rib-group Refresh>
 Address families configured: inet-unicast inet-vpn-unicast route-target
 Local Address: 10.255.167.205 Holdtime: 90 Preference: 170
 Number of flaps: 0
 Peer ID: 10.255.162.214 Local ID: 10.255.167.205 Active Holdtime: 90
 Keepalive Interval: 30 Group index: 0 Peer index: 1

```

### show bgp neighbor (CLNS)

```

user@host> show bgp neighbor
Peer: 10.245.245.1+179 AS 200 Local: 10.245.245.3+3770 AS 100
 Type: External State: Established Flags: <ImportEval Sync>
 Last State: OpenConfirm Last Event: RecvKeepAlive
 Last Error: None
 Options: <Multihop Preference LocalAddress HoldTime AddressFamily PeerAS
 Rib-group Refresh>
 Address families configured: iso-vpn-unicast
 Local Address: 10.245.245.3 Holdtime: 90 Preference: 170
 Number of flaps: 0
 Peer ID: 10.245.245.1 Local ID: 10.245.245.3 Active Holdtime: 90
 Keepalive Interval: 30 Peer index: 0
 NLRI advertised by peer: iso-vpn-unicast
 NLRI for this session: iso-vpn-unicast
 Peer supports Refresh capability (2)
 Table bgp.isovpn.0 Bit: 10000
 RIB State: BGP restart is complete
 RIB State: VPN restart is complete
 Send state: in sync
 Active prefixes: 3

```

```

Received prefixes: 3
Suppressed due to damping: 0
Advertised prefixes: 3
Table aaaa.iso.0
RIB State: BGP restart is complete
RIB State: VPN restart is complete
Send state: not advertising
Active prefixes: 3
Received prefixes: 3
Suppressed due to damping: 0
Last traffic (seconds): Received 6 Sent 5 Checked 5
Input messages: Total 1736 Updates 4 Refreshes 0 Octets 33385
Output messages: Total 1738 Updates 3 Refreshes 0 Octets 33305
Output Queue[0]: 0 (bgp.isovpn.0, iso-vpn-unicast)
Output Queue[1]: 0 (aaaa.iso.0, iso-vpn-unicast)

```

### show bgp neighbor (Layer 2 VPN)

```

user@host> show bgp neighbor
Peer: 10.69.103.2 AS 65100 Local: 10.69.103.1 AS 65103
Type: External State: Active Flags: <ImportEval>
Last State: Idle Last Event: Start
Last Error: None
Export: [BGP-INET-import]
Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily PeerAS
Refresh>
Address families configured: inet-unicast
Local Address: 10.69.103.1 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer: 10.69.104.2 AS 65100 Local: 10.69.104.1 AS 65104
Type: External State: Active Flags: <ImportEval>
Last State: Idle Last Event: Start
Last Error: None
Export: [BGP-L-import]
Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily PeerAS
Refresh>
Address families configured: inet-labeled-unicast
Local Address: 10.69.104.1 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer: 10.255.14.182+179 AS 69 Local: 10.255.14.176+2131 AS 69
Type: Internal State: Established Flags: <ImportEval>
Last State: OpenConfirm Last Event: RecvKeepAlive
Last Error: None
Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
Rib-group Refresh>
Address families configured: inet-vpn-unicast l2vpn
Local Address: 10.255.14.176 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer ID: 10.255.14.182 Local ID: 10.255.14.176 Active Holdtime: 90
Keepalive Interval: 30
NLRI for restart configured on peer: inet-vpn-unicast l2vpn
NLRI advertised by peer: inet-vpn-unicast l2vpn
NLRI for this session: inet-vpn-unicast l2vpn
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-vpn-unicast l2vpn
NLRI peer can save forwarding state: inet-vpn-unicast l2vpn
NLRI that peer saved forwarding for: inet-vpn-unicast l2vpn
NLRI that restart is negotiated for: inet-vpn-unicast l2vpn

```



```

NLRI of received end-of-rib markers: inet-vpn-unicast 12vpn
Table bgp.13vpn.0 Bit: 10000
 RIB State: BGP restart in progress
 RIB State: VPN restart in progress
 Send state: in sync
 Active prefixes: 10
 Received prefixes: 10
 Suppressed due to damping: 0
Table bgp.12vpn.0 Bit: 20000
 RIB State: BGP restart in progress
 RIB State: VPN restart in progress
 Send state: in sync
 Active prefixes: 1
 Received prefixes: 1
 Suppressed due to damping: 0
Table BGP-INET.inet.0 Bit: 30000
 RIB State: BGP restart in progress
 RIB State: VPN restart in progress
 Send state: in sync
 Active prefixes: 2
 Received prefixes: 2
 Suppressed due to damping: 0
Table BGP-L.inet.0 Bit: 40000
 RIB State: BGP restart in progress
 RIB State: VPN restart in progress
 Send state: in sync
 Active prefixes: 2
 Received prefixes: 2
 Suppressed due to damping: 0
Table LDP.inet.0 Bit: 50000
 RIB State: BGP restart is complete
 RIB State: VPN restart in progress
 Send state: in sync
 Active prefixes: 1
 Received prefixes: 1
 Suppressed due to damping: 0
Table OSPF.inet.0 Bit: 60000
 RIB State: BGP restart is complete
 RIB State: VPN restart in progress
 Send state: in sync
 Active prefixes: 2
 Received prefixes: 2
 Suppressed due to damping: 0
Table RIP.inet.0 Bit: 70000
 RIB State: BGP restart is complete
 RIB State: VPN restart in progress
 Send state: in sync
 Active prefixes: 2
 Received prefixes: 2
 Suppressed due to damping: 0
Table STATIC.inet.0 Bit: 80000
 RIB State: BGP restart is complete
 RIB State: VPN restart in progress
 Send state: in sync
 Active prefixes: 1
 Received prefixes: 1
 Suppressed due to damping: 0
Table L2VPN.12vpn.0 Bit: 90000
 RIB State: BGP restart is complete
 RIB State: VPN restart in progress
 Send state: in sync

```

```

Active prefixes: 1
Received prefixes: 1
Suppressed due to damping: 0
Last traffic (seconds): Received 0 Sent 0 Checked 0
Input messages: Total 14 Updates 13 Refreshes 0 Octets 1053
Output messages: Total 3 Updates 0 Refreshes 0 Octets 105
Output Queue[0]: 0 (bgp.l3vpn.0, inet-vpn-unicast)
Output Queue[1]: 0 (bgp.l2vpn.0, inet-vpn-unicast)
Output Queue[2]: 0 (BGP-INET.inet.0, inet-vpn-unicast)
Output Queue[3]: 0 (BGP-L.inet.0, inet-vpn-unicast)
Output Queue[4]: 0 (LDP.inet.0, inet-vpn-unicast)
Output Queue[5]: 0 (OSPF.inet.0, inet-vpn-unicast)
Output Queue[6]: 0 (RIP.inet.0, inet-vpn-unicast)
Output Queue[7]: 0 (STATIC.inet.0, inet-vpn-unicast)
Output Queue[8]: 0 (L2VPN.l2vpn.0, inet-vpn-unicast)

```

### show bgp neighbor (Layer 3 VPN) (Not supported on the OCX Series.)

```

user@host> show bgp neighbor
Peer: 4.4.4.4+179 AS 10045 Local: 5.5.5.5+1214 AS 10045
Type: Internal State: Established Flags: <ImportEval>
Last State: OpenConfirm Last Event: RecvKeepAlive
Last Error: None
Export: [match-all] Import: [match-all]
Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
Rib-group Refresh>
Address families configured: inet-vpn-unicast
Local Address: 5.5.5.5 Holdtime: 90 Preference: 170
Flags for NLRI inet-labeled-unicast: TrafficStatistics
Traffic Statistics: Options: all File: /var/log/bstat.log
 size 131072 files 10

Traffic Statistics Interval: 60
Number of flaps: 0
Peer ID: 192.168.1.110 Local ID: 192.168.1.111 Active Holdtime: 90
Keepalive Interval: 30
NLRI for restart configured on peer: inet-vpn-unicast
NLRI advertised by peer: inet-vpn-unicast
NLRI for this session: inet-vpn-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-vpn-unicast
NLRI peer can save forwarding state: inet-vpn-unicast
NLRI that peer saved forwarding for: inet-vpn-unicast
NLRI that restart is negotiated for: inet-vpn-unicast
NLRI of received end-of-rib markers: inet-vpn-unicast
NLRI of all end-of-rib markers sent: inet-vpn-unicast
Table bgp.l3vpn.0 Bit: 10000
 RIB State: BGP restart is complete
 RIB State: VPN restart is complete
 Send state: in sync
 Active prefixes: 2
 Received prefixes: 2
 Suppressed due to damping: 0
Table vpn-green.inet.0 Bit: 20001
 RIB State: BGP restart is complete
 RIB State: VPN restart is complete
 Send state: in sync
 Active prefixes: 2
 Received prefixes: 2

```

```

 Suppressed due to damping: 0
 Last traffic (seconds): Received 15 Sent 20 Checked 20
 Input messages: Total 40 Updates 2 Refreshes 0 Octets 856
 Output messages: Total 44 Updates 2 Refreshes 0 Octets 1066
 Output Queue[0]: 0 (bgp.13vpn.0, inet-vpn-unicast)
 Output Queue[1]: 0 (vpn-green.inet.0, inet-vpn-unicast)
 Trace options: detail packets
 Trace file: /var/log/bgpgr.log size 131072 files 10

```

### show bgp neighbor neighbor-address

```

user@host> show bgp neighbor 192.168.1.111
Peer: 10.255.245.12+179 AS 35 Local: 10.255.245.13+2884 AS 35
 Type: Internal State: Established (route reflector client)Flags: <Sync>
 Last State: OpenConfirm Last Event: RecvKeepAlive
 Last Error: None
 Options: <Preference LocalAddress HoldTime Cluster AddressFamily Rib-group
Refresh>
 Address families configured: inet-vpn-unicast inet-labeled-unicast
 Local Address: 10.255.245.13 Holdtime: 90 Preference: 170
 Flags for NLRI inet-vpn-unicast: AggregateLabel
 Flags for NLRI inet-labeled-unicast: AggregateLabel
 Number of flaps: 0
 Peer ID: 10.255.245.12 Local ID: 10.255.245.13 Active Holdtime: 90
 Keepalive Interval: 30
 BFD: disabled
 NLRI advertised by peer: inet-vpn-unicast inet-labeled-unicast
 NLRI for this session: inet-vpn-unicast inet-labeled-unicast
 Peer supports Refresh capability (2)
 Restart time configured on the peer: 300
 Stale routes from peer are kept for: 60
 Restart time requested by this peer: 300
 NLRI that peer supports restart for: inet-unicast inet6-unicast
 NLRI that restart is negotiated for: inet-unicast inet6-unicast
 NLRI of received end-of-rib markers: inet-unicast inet6-unicast
 NLRI of all end-of-rib markers sent: inet-unicast inet6-unicast
 Table inet.0 Bit: 10000
 RIB State: restart is complete
 Send state: in sync
 Active prefixes: 4
 Received prefixes: 6
 Suppressed due to damping: 0
 Table inet6.0 Bit: 20000
 RIB State: restart is complete
 Send state: in sync
 Active prefixes: 0
 Received prefixes: 2
 Suppressed due to damping: 0
 Last traffic (seconds): Received 3 Sent 3 Checked 3
 Input messages: Total 9 Updates 6 Refreshes 0 Octets 403
 Output messages: Total 7 Updates 3 Refreshes 0 Octets 365
 Output Queue[0]: 0 (inet.0, inet-unicast)
 Output Queue[1]: 0 (inet6.0, inet6-unicast)
 Trace options: detail packets
 Trace file: /var/log/bgpgr size 131072 files 10

```

### show bgp neighbor neighbor-address

```

user@host> show bgp neighbor 192.168.4.222
Peer: 192.168.4.222+4902 AS 65501 Local: 192.168.4.221+179 AS 65500
 Type: External State: Established Flags: <Sync>

```

```

Last State: OpenConfirm Last Event: RecvKeepAlive
Last Error: Cease
Export: [export-policy] Import: [import-policy]
Options: <Preference HoldTime AddressFamily PeerAS PrefixLimit Refresh>
Address families configured: inet-unicast inet-multicast
Holdtime: 60000 Preference: 170
Number of flaps: 4
Last flap event: RecvUpdate
Error: 'Cease' Sent: 5 Recv: 0
Peer ID: 10.255.245.6 Local ID: 10.255.245.5 Active Holdtime: 60000
Keepalive Interval: 20000 Peer index: 0
BFD: disabled, down
Local Interface: fxp0.0
NLRI advertised by peer: inet-unicast inet-multicast
NLRI for this session: inet-unicast inet-multicast
Peer supports Refresh capability (2)
Table inet.0 Bit: 10000
 RIB State: BGP restart is complete
 Send state: in sync
 Active prefixes: 8
 Received prefixes: 10
 Accepted prefixes: 10
 Suppressed due to damping: 0
 Advertised prefixes: 3
Table inet.2 Bit: 20000
 RIB State: BGP restart is complete
 Send state: in sync
 Active prefixes: 0
 Received prefixes: 0
 Accepted prefixes: 0
 Suppressed due to damping: 0
 Advertised prefixes: 0
Last traffic (seconds): Received 357 Sent 357 Checked 357
Input messages: Total 4 Updates 2 Refreshes 0 Octets 211
Output messages: Total 4 Updates 1 Refreshes 0 Octets 147
Output Queue[0]: 0 (inet.0, inet-unicast)
Output Queue[1]: 0 (inet.2, inet-multicast)
Trace options: all
Trace file: /var/log/bgp size 10485760 files 10

```

### show bgp neighbor neighbor-address (BGP Graceful Restart Enabled)

```
user@router> show bgp neighbor 10.255.255.16
```

```

Peer: 10.255.255.16 AS 100 Local: 10.255.255.12 AS 100
Type: Internal State: Active Flags: <>
Last State: Idle Last Event: Start
Last Error: None
Options: <Preference LocalAddress AddressFamily Rib-group Refresh>
Options: <LLGR>
Address families configured: 12vpn
Local Address: 10.255.255.12 Holdtime: 90 Preference: 170
NLRI 12vpn:
Number of flaps: 6
Last flap event: Restart
NLRI we are holding stale routes for: inet-vpn-unicast
Time until stale routes are deleted or become long-lived stale: 00:01:57
Time until end-of-rib is assumed for stale routes: 00:04:43
Table bgp.13vpn.0
 RIB State: BGP restart is complete
 RIB State: VPN restart is complete

```

```

Send state: not advertising
Active prefixes: 0
Received prefixes: 7
Accepted prefixes: 7
Suppressed due to damping: 0
Table foo.inet.0 Bit: 30000
RIB State: BGP restart is complete
RIB State: VPN restart is complete
Send state: not in sync
Active prefixes: 0
Received prefixes: 7
Accepted prefixes: 7
Suppressed due to damping: 0

```

### show bgp neighbor neighbor-address (BGP Long-Lived Graceful Restart)

```

user@router> show bgp neighbor 10.4.12.11

Peer: 10.4.12.11 AS 100 Local: 10.6.128.225 AS 100
Type: Internal State: Active Flags: <>
Last State: Idle Last Event: Start
Last Error: None
Export: [foo]
Options: <Preference LocalAddress Refresh GracefulRestart>
Options: <LLGR>
Local Address: 10.6.128.225 Holdtime: 90 Preference: 170
Number of flaps: 3
Last flap event: Restart
Error: 'Cease' Sent: 0 Recv: 1
Time until long-lived stale routes deleted: inet-vpn-unicast 10:00:22
route-target 10:00:22
Table bgp.l3vpn.0
RIB State: BGP restart is complete
RIB State: VPN restart is complete
Send state: not advertising
Active prefixes: 0
Received prefixes: 7
Accepted prefixes: 7
Suppressed due to damping: 0
Table foo.inet.0 Bit: 30000
RIB State: BGP restart is complete
RIB State: VPN restart is complete
Send state: not in sync
Active prefixes: 0
Received prefixes: 7
Accepted prefixes: 7
Suppressed due to damping: 0

```

### show bgp neighbor orf neighbor-address detail

```

user@host > show bgp neighbor orf 192.168.165.56 detail
Peer: 192.168.165.56+179 Type: External
Group: ext1

inet-unicast
Filter updates rcv: 1 Immediate: 1
Filter: prefix-based receive
Received filter entries:
seq 1: prefix 2.2.2.2/32: minlen 32: maxlen 32: match deny:

inet6-unicast

```

```
Filter updates rcv: 0 Immediate: 1
Filter: prefix-based receive
Received filter entries:
:
```

## show log

|                                           |                                                                                                                                                                                                                                                                                                                                                                                                              |
|-------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                     | <a href="#">Syntax on page 2727</a><br><a href="#">Syntax (QFX Series and OCX Series) on page 2727</a><br><a href="#">Syntax (TX Matrix Router) on page 2727</a>                                                                                                                                                                                                                                             |
| <b>Syntax</b>                             | <pre>show log &lt;filename   user &lt;username&gt;&gt;</pre>                                                                                                                                                                                                                                                                                                                                                 |
| <b>Syntax (QFX Series and OCX Series)</b> | <pre>show log filename &lt;device-type (device-id   device-alias)&gt;</pre>                                                                                                                                                                                                                                                                                                                                  |
| <b>Syntax (TX Matrix Router)</b>          | <pre>show log &lt;all-lcc   lcc number   scc&gt; &lt;filename   user &lt;username&gt;&gt;</pre>                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>                | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Option <i>device-type (device-id   device-alias)</i> is introduced in Junos OS Release 13.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| <b>Description</b>                        | List log files, display log file contents, or display information about users who have logged in to the router or switch.                                                                                                                                                                                                                                                                                    |



**NOTE:** On MX Series routers, modifying a configuration to replace a service interface with another service interface is treated as a catastrophic event. When you modify a configuration, the entire configuration associated with the service interface—including NAT pools, rules, and service sets—is deleted and then re-created for the newly specified service interface. If there are active sessions associated with the service interface that is being replaced, these sessions are deleted and the NAT pools are then released, which leads to the generation of the NAT\_POOL\_RELEASE system log messages. However, because NAT pools are already deleted as a result of the catastrophic configuration change and no longer exist, the NAT\_POOL\_RELEASE system log messages are not generated for the changed configuration.

**Options** none—List all log files.

**<all-lcc | lcc *number* | scc>**—(Routing matrix only) (Optional) Display logging information about all T640 routers (or line-card chassis) or a specific T640 router (replace *number* with a value from 0 through 3) connected to a TX Matrix router. Or, display logging information about the TX Matrix router (or switch-card chassis).

**device-type**—(QFabric system only) (Optional) Display log messages for only one of the following device types:

- **director-device**—Display logs for Director devices.
- **infrastructure-device**—Display logs for the logical components of the QFabric system infrastructure, including the diagnostic Routing Engine, fabric control Routing Engine, fabric manager Routing Engine, and the default network Node group and its backup (NW-NG-0 and NW-NG-0-backup).
- **interconnect-device**—Display logs for Interconnect devices.
- **node-device**—Display logs for Node devices.



**NOTE:** If you specify the **device-type** optional parameter, you must also specify either the **device-id** or **device-alias** optional parameter.

**(device-id | device-alias)**—If a device type is specified, display logs for a device of that type. Specify either the device ID or the device alias (if configured).

**filename**—(Optional) Display the log messages in the specified log file. For the routing matrix, the filename must include the chassis information.

**user <username>**—(Optional) Display logging information about users who have recently logged in to the router or switch. If you include **username**, display logging information about the specified user.

**Required Privilege Level** trace

**Related Documentation** • [syslog \(System\) on page 1704](#)

**List of Sample Output** [show log on page 2728](#)  
[show log filename on page 2729](#)  
[show log filename \(QFabric System\) on page 2729](#)  
[show log user on page 2730](#)

## Sample Output

show log

```
user@host> show log
total 57518
-rw-r--r-- 1 root bin 211663 Oct 1 19:44 dcd
-rw-r--r-- 1 root bin 999947 Oct 1 19:41 dcd.0
-rw-r--r-- 1 root bin 999994 Oct 1 17:48 dcd.1
-rw-r--r-- 1 root bin 238815 Oct 1 19:44 rpd
-rw-r--r-- 1 root bin 1049098 Oct 1 18:00 rpd.0
-rw-r--r-- 1 root bin 1061095 Oct 1 12:13 rpd.1
-rw-r--r-- 1 root bin 1052026 Oct 1 06:08 rpd.2
-rw-r--r-- 1 root bin 1056309 Sep 30 18:21 rpd.3
-rw-r--r-- 1 root bin 1056371 Sep 30 14:36 rpd.4
-rw-r--r-- 1 root bin 1056301 Sep 30 10:50 rpd.5
-rw-r--r-- 1 root bin 1056350 Sep 30 07:04 rpd.6
```



```
-rw-r--r-- 1 root bin 1048876 Sep 30 03:21 rpd.7
-rw-rw-r-- 1 root bin 19656 Oct 1 19:37 wtmp
```

### show log filename

```
user@host> show log rpd
Oct 1 18:00:18 trace_on: Tracing to ?/var/log/rpd? started
Oct 1 18:00:18 EVENT <MTU> ds-5/2/0.0 index 24 <Broadcast PointToPoint Multicast
Oct 1 18:00:18
Oct 1 18:00:19 KRT recv len 56 V9 seq 148 op add Type route/if af 2 addr
13.13.13.21 nhop type local nhop 13.13.13.21
Oct 1 18:00:19 KRT recv len 56 V9 seq 149 op add Type route/if af 2 addr
13.13.13.22 nhop type unicast nhop 13.13.13.22
Oct 1 18:00:19 KRT recv len 48 V9 seq 150 op add Type ifaddr index 24 devindex
43
Oct 1 18:00:19 KRT recv len 144 V9 seq 151 op chnge Type ifdev devindex 44
Oct 1 18:00:19 KRT recv len 144 V9 seq 152 op chnge Type ifdev devindex 45
Oct 1 18:00:19 KRT recv len 144 V9 seq 153 op chnge Type ifdev devindex 46
Oct 1 18:00:19 KRT recv len 1272 V9 seq 154 op chnge Type ifdev devindex 47
...
```

### show log filename (QFabric System)

```
user@qfabric> show log messages
Mar 28 18:00:06 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:06 ED1486
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 2159)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1486
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 2, jnxFruL3Index 0,
jnxFruName PIC: @ 0/1/*, jnxFruType 11, jnxFruSlot 0, jnxFruOfflineReason 2,
jnxFruLastPowerOff 0, jnxFruLastPowerOn 2191)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 242726)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 2, jnxFruL3Index 0,
jnxFruName PIC: @ 0/1/*, jnxFruType 11, jnxFruSlot 0, jnxFruOfflineReason 2,
jnxFruLastPowerOff 0, jnxFruLastPowerOn 242757)
Mar 28 18:00:16 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:16 ED1486
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:27 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:27 ED1486
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50
_DCF_default__NW-INE-0_RE0_ file: UI_COMMIT: User 'root' requested 'commit'
operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50
_DCF_default__NW-INE-0_RE0_ file: UI_COMMIT: User 'root' requested 'commit'
operation (comment: none)
Mar 28 18:00:55 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:55 ED1492
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:01:10 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:01:10 ED1492
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:02:37 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:02:37 ED1491
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
```

```
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 33809)
```

### show log user

```
user@host> show log user
darius mg2546 Thu Oct 1 19:37 still logged in
darius mg2529 Thu Oct 1 19:08 - 19:36 (00:28)
darius mg2518 Thu Oct 1 18:53 - 18:58 (00:04)
root mg1575 Wed Sep 30 18:39 - 18:41 (00:02)
root ttyp2 jun.site.per Wed Sep 30 18:39 - 18:41 (00:02)
alex ttyp1 192.168.1.2 Wed Sep 30 01:03 - 01:22 (00:19)
```

## show (ospf | ospf3) overview

|                                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                           | <a href="#">Syntax on page 2731</a><br><a href="#">Syntax (EX Series Switch and QFX Series) on page 2731</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Syntax</b>                                   | <pre>show (ospf   ospf3) overview &lt;brief   extensive&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Syntax (EX Series Switch and QFX Series)</b> | <pre>show (ospf   ospf3) overview &lt;brief   extensive&gt; &lt;instance <i>instance-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>                      | <p>Command introduced in Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>realm</b> option introduced in Junos OS Release 9.2.</p> <p>Database protection introduced in Junos 10.2.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>                              | Display Open Shortest Path First (OSPF) overview information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                                  | <p><b>none</b>—Display standard information about all OSPF neighbors for all routing instances.</p> <p><b>brief   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display all OSPF interfaces under the named routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)</b>—(Optional) (OSPFv3 only) Display information about the specified OSPFv3 realm, or address family. Use the <b>realm</b> option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.</p> |
| <b>Required Privilege Level</b>                 | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>List of Sample Output</b>                    | <a href="#">show ospf overview on page 2733</a><br><a href="#">show ospf overview (With Database Protection) on page 2734</a><br><a href="#">show ospf3 overview (With Database Protection) on page 2734</a><br><a href="#">show ospf overview extensive on page 2734</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Output Fields</b>                            | Table 239 lists the output fields for the <b>show ospf overview</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

Table 239: show ospf overview Output Fields

| Field name                       | Field Description                                                                                                                                                                        | Level of Output |
|----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Instance</b>                  | OSPF routing instance.                                                                                                                                                                   | All levels      |
| <b>Router ID</b>                 | Router ID of the routing device.                                                                                                                                                         | All levels      |
| <b>Route table index</b>         | Route table index.                                                                                                                                                                       | All levels      |
| <b>Configured overload</b>       | Overload capability is enabled. If the overload timer is also configured, display the time that remains before it is set to expire. This field is not displayed after the timer expires. | All levels      |
| <b>Topology</b>                  | Topology identifier.                                                                                                                                                                     | All levels      |
| <b>Prefix export count</b>       | Number of prefixes exported into OSPF.                                                                                                                                                   | All levels      |
| <b>Full SPF runs</b>             | Number of complete Shortest Path First calculations.                                                                                                                                     | All levels      |
| <b>SPF delay</b>                 | Delay before performing consecutive Shortest Path First calculations.                                                                                                                    | All levels      |
| <b>SPF holddown</b>              | Delay before performing additional Shortest Path First (SPF) calculations after the maximum number of consecutive SPF calculations is reached.                                           | All levels      |
| <b>SPF rapid runs</b>            | Maximum number of Shortest Path First calculations that can be performed in succession before the hold-down timer begins.                                                                | All levels      |
| <b>LSA refresh time</b>          | Refresh period for link-state advertisement (in minutes).                                                                                                                                | All levels      |
| <b>Database protection state</b> | Current state of database protection.                                                                                                                                                    | All levels      |
| <b>Warning threshold</b>         | Threshold at which a warning message is logged (percentage of maximum LSA count).                                                                                                        | All levels      |
| <b>Non self-generated LSAs</b>   | Number of LSAs whose router ID is not equal to the local router ID: <b>Current</b> , <b>Warning</b> (threshold), and <b>Allowed</b> .                                                    | All levels      |
| <b>Ignore time</b>               | How long the database has been in the ignore state.                                                                                                                                      | All levels      |
| <b>Reset time</b>                | How long the database must stay out of the ignore or isolated state before it returns to normal operations.                                                                              | All levels      |
| <b>Ignore count</b>              | Number of times the database has been in the ignore state: <b>Current</b> and <b>Allowed</b> .                                                                                           | All levels      |
| <b>Restart</b>                   | Graceful restart capability: <b>enabled</b> or <b>disabled</b> .                                                                                                                         | All levels      |
| <b>Restart duration</b>          | Time period for complete reacquisition of OSPF neighbors.                                                                                                                                | All levels      |
| <b>Restart grace period</b>      | Time period for which the neighbors should consider the restarting routing device as part of the topology.                                                                               | All levels      |

Table 239: show ospf overview Output Fields (*continued*)

| Field name                    | Field Description                                                                                                                                    | Level of Output  |
|-------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| Graceful restart helper mode  | (OSPFv2) Standard graceful restart helper capability (based on RFC 3623): <b>enabled</b> or <b>disabled</b> .                                        | All levels       |
| Restart-signaling helper mode | (OSPFv2) Restart signaling-based graceful restart helper capability (based on RFC 4811, RFC 4812, and RFC 4813): <b>enabled</b> or <b>disabled</b> . | All levels       |
| Helper mode                   | (OSPFv3) Graceful restart helper capability: <b>enabled</b> or <b>disabled</b> .                                                                     | All levels       |
| Trace options                 | OSPF-specific trace options.                                                                                                                         | <b>extensive</b> |
| Trace file                    | Name of the file to receive the output of the tracing operation.                                                                                     | <b>extensive</b> |
| Area                          | Area number. Area 0.0.0.0 is the backbone area.                                                                                                      | All levels       |
| Stub type                     | Stub type of area: <b>Normal Stub</b> , <b>Not Stub</b> , or <b>Not so Stubby Stub</b> .                                                             | All levels       |
| Authentication Type           | Type of authentication: <b>None</b> , <b>Password</b> , or <b>MD5</b> .                                                                              | All levels       |
| Area border routers           | Number of area border routers.                                                                                                                       | All levels       |
| Neighbors                     | Number of autonomous system boundary routers.                                                                                                        | All levels       |

## Sample Output

### show ospf overview

```

user@host> show ospf overview
Instance: master
 Router ID: 10.255.245.6
 Route table index: 0
 Configured overload, expires in 118 seconds
 LSA refresh time: 50 minutes
Restart: Enabled
 Restart duration: 20 sec
 Restart grace period: 40 sec
 Helper mode: enabled
Area: 0.0.0.0
 Stub type: Not Stub
 Authentication Type: None
 Area border routers: 0, AS boundary routers: 0
 Neighbors
 Up (in full state): 0
Topology: default (ID 0)
 Prefix export count: 0
 Full SPF runs: 1
 SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3

```

### show ospf overview (With Database Protection)

```
user@host> show ospf overview
Instance: master
 Router ID: 10.255.112.218
 Route table index: 0
 LSA refresh time: 50 minutes
 Traffic engineering
 Restart: Enabled
 Restart duration: 180 sec
 Restart grace period: 210 sec
 Graceful restart helper mode: Enabled
 Restart-signaling helper mode: Enabled
 Database protection state: Normal
 Warning threshold: 70 percent
 Non self-generated LSAs: Current 582, Warning 700, Allowed 1000
 Ignore time: 30, Reset time: 60
 Ignore count: Current 0, Allowed 1
 Area: 0.0.0.0
 Stub type: Not Stub
 Authentication Type: None
 Area border routers: 0, AS boundary routers: 0
 Neighbors
 Up (in full state): 160
 Topology: default (ID 0)
 Prefix export count: 0
 Full SPF runs: 70
 SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
 Backup SPF: Not Needed
```

### show ospf3 overview (With Database Protection)

```
user@host> show ospf3 overview
Instance: master
 Router ID: 10.255.112.128
 Route table index: 0
 LSA refresh time: 50 minutes
 Database protection state: Normal
 Warning threshold: 80 percent
 Non self-generated LSAs: Current 3, Warning 8, Allowed 10
 Ignore time: 30, Reset time: 60
 Ignore count: Current 0, Allowed 2
 Area: 0.0.0.0
 Stub type: Not Stub
 Area border routers: 0, AS boundary routers: 0
 Neighbors
 Up (in full state): 1
 Topology: default (ID 0)
 Prefix export count: 0
 Full SPF runs: 7
 SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
 Backup SPF: Not Needed
```

### show ospf overview extensive

```
user@host> show ospf overview extensive
Instance: master
 Router ID: 1.1.1.103
 Route table index: 0
 Full SPF runs: 13, SPF delay: 0.200000 sec
 LSA refresh time: 50 minutes
```

```
Restart: Disabled
Trace options: lsa
Trace file: /var/log/ospf size 131072 files 10
Area: 0.0.0.0
 Stub type: Not Stub
 Authentication Type: None
 Area border routers: 0, AS boundary routers: 0
 Neighbors
 Up (in full state): 1
```





## CHAPTER 108

# Operational Mode Commands (Graceful Switchover)

- `show system switchover`
- `show task replication`

## show system switchover

---

|                                       |                                                                                                                                                                                                                                                                                      |
|---------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                 | <a href="#">Syntax on page 2738</a><br><a href="#">Syntax (TX Matrix Router) on page 2738</a><br><a href="#">Syntax (TX Matrix Plus Router) on page 2738</a><br><a href="#">Syntax (MX Series Router) on page 2738</a>                                                               |
| <b>Syntax</b>                         | show system switchover                                                                                                                                                                                                                                                               |
| <b>Syntax (TX Matrix Router)</b>      | show system switchover<br><all-chassis   all-lcc   lcc <i>number</i>   scc>                                                                                                                                                                                                          |
| <b>Syntax (TX Matrix Plus Router)</b> | show system switchover<br><all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i> >                                                                                                                                                                                           |
| <b>Syntax (MX Series Router)</b>      | show system switchover<br><all-members><br><local><br><member <i>member-id</i> >                                                                                                                                                                                                     |
| <b>Release Information</b>            | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.<br>Command introduced in Junos OS Release 13.2X51-D20 for QFX Series switches. |
| <b>Description</b>                    | Display whether graceful Routing Engine switchover is configured, the state of the kernel replication (ready or synchronizing), any replication errors, and whether the primary and standby Routing Engines are using compatible versions of the kernel database.                    |



**NOTE:** Issue the `show system switchover` command *only* on the backup Routing Engine. This command is *not* supported on the master Routing Engine, because the kernel-replication process daemon does not run on the master Routing Engine. This process runs only on the backup Routing Engine.

Beginning Junos OS Release 9.6, the `show system switchover` command has been deprecated on the master Routing Engine on all routers other than a TX Matrix (switch-card chassis) or a TX Matrix Plus (switch-fabric chassis) router.

However, in a routing matrix, if you issue the `show system switchover` command on the master Routing Engine of the TX Matrix router (or switch-card chassis), the CLI displays graceful switchover information for the master Routing Engine of the T640 routers (or line-card chassis) in the routing matrix. Likewise, if you issue the `show system switchover` command on the master Routing Engine of a TX Matrix Plus router (or switch-fabric chassis), the CLI displays output for the master Routing Engine of T1600 or T4000 routers in the routing matrix.

---

**Options** **all-chassis**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display graceful Routing Engine switchover information for all Routing Engines on the TX Matrix router and the T640 routers configured in the routing matrix. On a TX Matrix Plus router, display graceful Routing Engine switchover information for all Routing Engines on the TX Matrix Plus router and the T1600 or T4000 routers configured in the routing matrix.

**all-lcc**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display graceful Routing Engine switchover information for all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, display graceful Routing Engine switchover information for all connected T1600 or T4000 LCCs.

Note that in this instance, packets get dropped. The LCCs perform GRES on their own chassis (GRES cannot be handled by one particular chassis for the entire router) and synchronization is not possible as the LCC plane bringup time varies for each LCC. Therefore, when there is traffic on these planes, there may be a traffic drop.

**all-members**—(MX Series routers only) (Optional) Display graceful Routing Engine switchover information for all Routing Engines on all members of the Virtual Chassis configuration.

**lcc *number***—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display graceful Routing Engine switchover information for a specific T640 router connected to the TX Matrix router. On a TX Matrix Plus router, display graceful Routing Engine switchover information for a specific router connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(MX Series routers only) (Optional) Display graceful Routing Engines switchover information for all Routing Engines on the local Virtual Chassis member.

**member *member-id***—(MX Series routers only) (Optional) Display graceful Routing Engine switchover information for all Routing Engines on the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

**scc**—(TX Matrix router only) (Optional) Display graceful Routing Engine switchover information for the TX Matrix router (or switch-card chassis).

**sfc**—(TX Matrix Plus routers only) (Optional) Display graceful Routing Engine switchover information for the TX Matrix Plus router.

**Additional Information** If you issue the **show system switchover** command on a TX Matrix backup Routing Engine, the command is broadcast to all the T640 backup Routing Engines that are connected to it.

Likewise, if you issue the **show system switchover** command on a TX Matrix Plus backup Routing Engine, the command is broadcast to all the T1600 or T4000 backup Routing Engines that are connected to it.

If you issue the **show system switchover** command on the active Routing Engine in the master router of an MX Series Virtual Chassis, the router displays a message that this command is not applicable on this member of the Virtual Chassis.

**Required Privilege Level** view

**Related Documentation**

- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

**List of Sample Output**

[show system switchover \(Backup Routing Engine - Ready\) on page 2741](#)  
[show system switchover \(Backup Routing Engine - Not Ready\) on page 2741](#)  
[show system switchover \(MX Virtual Chassis\) on page 2741](#)  
[show system switchover \(Routing Matrix and Routing Matrix Plus\) - Master Ready on page 2742](#)  
[show system switchover \(Routing Matrix and Routing Matrix Plus\) - Master Not Ready on page 2742](#)  
[show system switchover \(Routing Matrix and Routing Matrix Plus\) - Backup Ready on page 2742](#)  
[show system switchover \(Routing Matrix and Routing Matrix Plus\) - Backup Not Ready on page 2743](#)  
[show system switchover all-lcc \(Routing Matrix and Routing Matrix Plus\) on page 2743](#)

**Output Fields** [Table 240](#) describes the output fields for the **show system switchover** command. Output fields are listed in the approximate order in which they appear.

**Table 240: show system switchover Output Fields**

| Field Name                    | Field Description                                                                                                                                                                                                                                                                                                                                                           |
|-------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Graceful switchover</b>    | Display graceful Routing Engine switchover status: <ul style="list-style-type: none"> <li>• <b>On</b>—Indicates <b>graceful-switchover</b> is specified for the <b>routing-options</b> configuration command.</li> <li>• <b>Off</b>—Indicates <b>graceful-switchover</b> is not specified for the <b>routing-options</b> configuration command.</li> </ul>                  |
| <b>Configuration database</b> | State of the configuration database: <ul style="list-style-type: none"> <li>• <b>Ready</b>—Configuration database has synchronized.</li> <li>• <b>Synchronizing</b>—Configuration database is synchronizing. Displayed when there are updates within the last 5 seconds.</li> <li>• <b>Synchronize failed</b>—Configuration database synchronize process failed.</li> </ul> |

Table 240: show system switchover Output Fields (*continued*)

| Field Name      | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Kernel database | <p>State of the kernel database:</p> <ul style="list-style-type: none"> <li>• <b>Ready</b>—Kernel database has synchronized. This message implies that the system is ready for GRES.</li> <li>• <b>Synchronizing</b>—Kernel database is synchronizing. Displayed when there are updates within the last 5 seconds.</li> <li>• <b>Version incompatible</b>—The primary and standby Routing Engines are running incompatible kernel database versions.</li> <li>• <b>Replication error</b>—An error occurred when the state was replicated from the primary Routing Engine. Inspect <b>Steady State</b> for possible causes, or notify Juniper Networks customer support.</li> </ul> |
| Peer state      | <p>Routing Engine peer state:</p> <p>This field is displayed only when ksyncd is running in multichassis mode (LCC master).</p> <ul style="list-style-type: none"> <li>• <b>Steady State</b>—Peer completed switchover transition.</li> <li>• <b>Peer Connected</b>—Peer in switchover transition.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                      |
| Switchover      | <p>Switchover status (output of <b>master switch check</b> command):</p> <ul style="list-style-type: none"> <li>• <b>Ready</b>—Message for system being switchover ready.</li> <li>• <b>error: Command aborted. Not ready for mastership switch, try after xxx secs.</b></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                |

## Sample Output

### show system switchover (Backup Routing Engine - Ready)

```
user@host> show system switchover
Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady State
Switchover Ready
```

### show system switchover (Backup Routing Engine - Not Ready)

```
user@host> show system switchover
Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady State
```

```
error: Command aborted. Not ready for mastership switch, try after 174 secs.
```

### show system switchover (MX Virtual Chassis)

```
{master:member1-re1}
user@host> show system switchover
member0:

Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Switchover Ready
```

member1:

-----  
Command is not applicable on this member of the virtual-chassis

#### show system switchover (Routing Matrix and Routing Matrix Plus) - Master Ready

```
user@host> show system switchover
lcc0-re1:
```

-----  
Multichassis replication: On  
Configuration database: Ready  
Kernel database: Ready  
Peer state: Steady State  
Switchover Ready

lcc2-re0:

-----  
Multichassis replication: On  
Configuration database: Ready  
Kernel database: Ready  
Peer state: Steady State  
Switchover Ready

#### show system switchover (Routing Matrix and Routing Matrix Plus) - Master Not Ready

```
user@host> show system switchover
lcc0-re1:
```

-----  
Multichassis replication: On  
Configuration database: Ready  
Kernel database: Ready  
Peer state: Steady State  
Switchover Ready

lcc2-re1:

-----  
Multichassis replication: On  
Configuration database: Ready  
Kernel database: Ready  
Peer state: Steady State  
error: Command aborted. Not ready for mastership switch, try after 228 secs.

#### show system switchover (Routing Matrix and Routing Matrix Plus) - Backup Ready

```
user@host> show system switchover
scc-re0:
```

-----  
Graceful switchover: On  
Configuration database: Ready  
Kernel database: Ready  
Switchover Ready

lcc0-re0:

-----  
Graceful switchover: On  
Configuration database: Ready  
Kernel database: Ready  
Switchover Ready

lcc2-re1:

```
Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Switchover Ready
```

#### show system switchover (Routing Matrix and Routing Matrix Plus) - Backup Not Ready

```
user@host> show system switchover
scc-re0:

Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
error: Command aborted. Not ready for mastership switch, try after 223 secs.

lcc0-re0:

Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Switchover Ready

lcc2-re1:

Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Switchover Ready
```

#### show system switchover all-lcc (Routing Matrix and Routing Matrix Plus)

```
user@host> show system switchover all-lcc

lcc0-re0:

Multichassis replication: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady State
lcc2-re0:

Multichassis replication: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady State
```

## show task replication

|                                 |                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show task replication</b>                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 8.5.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X51-D20 for QFX Series switches.</p> <p>Support for logical systems introduced in Junos OS Release 13.3</p> |
| <b>Description</b>              | Displays nonstop active routing (NSR) status. When you issue this command on the master Routing Engine, the status of nonstop active routing synchronization is also displayed.                                                                                                        |
| <b>Options</b>                  | This command has no options.                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                   |
| <b>List of Sample Output</b>    | <a href="#">show task replication (Issued on the Master Routing Engine) on page 2744</a><br><a href="#">show task replication (Issued on the Backup Routing Engine) on page 2745</a>                                                                                                   |
| <b>Output Fields</b>            | Table 241 lists the output fields for the <b>show task replication</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                |

**Table 241: show task replication Output Fields**

| Field Name                    | Field Description                                                                                                                                                   |
|-------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Stateful replication</b>   | Displays whether or not graceful Routing Engine switchover is configured. The status can be <b>Enabled</b> or <b>Disabled</b> .                                     |
| <b>RE mode</b>                | Displays the Routing Engine on which the command is issued: <b>Master</b> , <b>Backup</b> , or <b>Not applicable</b> (when the router has only one Routing Engine). |
| <b>Protocol</b>               | Protocols that are supported by nonstop active routing.                                                                                                             |
| <b>Synchronization Status</b> | Nonstop active routing synchronization status for the supported protocols. States are <b>NotStarted</b> , <b>InProgress</b> , and <b>Complete</b> .                 |

## Sample Output

### show task replication (Issued on the Master Routing Engine)

```

user@host> show task replication
 Stateful Replication: Enabled
 RE mode: Master

 Protocol Synchronization Status

 OSPF NotStarted
 BGP Complete
 IS-IS NotStarted

```



|     |          |
|-----|----------|
| LDP | Complete |
| PIM | Complete |

#### show task replication (Issued on the Backup Routing Engine)

```
user@host> show task replication
Stateful Replication: Enabled
RE mode: Backup
```



## CHAPTER 109

# Operational Mode Command (Nonstop Routing)

- `show task replication`

## show task replication

|                                 |                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show task replication</b>                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 8.5.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X51-D20 for QFX Series switches.</p> <p>Support for logical systems introduced in Junos OS Release 13.3</p> |
| <b>Description</b>              | Displays nonstop active routing (NSR) status. When you issue this command on the master Routing Engine, the status of nonstop active routing synchronization is also displayed.                                                                                                        |
| <b>Options</b>                  | This command has no options.                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                   |
| <b>List of Sample Output</b>    | <a href="#">show task replication (Issued on the Master Routing Engine) on page 2748</a><br><a href="#">show task replication (Issued on the Backup Routing Engine) on page 2749</a>                                                                                                   |
| <b>Output Fields</b>            | Table 241 lists the output fields for the <b>show task replication</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                |

**Table 242: show task replication Output Fields**

| Field Name                    | Field Description                                                                                                                                                   |
|-------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Stateful replication</b>   | Displays whether or not graceful Routing Engine switchover is configured. The status can be <b>Enabled</b> or <b>Disabled</b> .                                     |
| <b>RE mode</b>                | Displays the Routing Engine on which the command is issued: <b>Master</b> , <b>Backup</b> , or <b>Not applicable</b> (when the router has only one Routing Engine). |
| <b>Protocol</b>               | Protocols that are supported by nonstop active routing.                                                                                                             |
| <b>Synchronization Status</b> | Nonstop active routing synchronization status for the supported protocols. States are <b>NotStarted</b> , <b>InProgress</b> , and <b>Complete</b> .                 |

## Sample Output

### show task replication (Issued on the Master Routing Engine)

```

user@host> show task replication
 Stateful Replication: Enabled
 RE mode: Master

 Protocol Synchronization Status
 OSPF NotStarted
 BGP Complete
 IS-IS NotStarted

```

|     |          |
|-----|----------|
| LDP | Complete |
| PIM | Complete |

#### show task replication (Issued on the Backup Routing Engine)

```
user@host> show task replication
Stateful Replication: Enabled
RE mode: Backup
```



## CHAPTER 110

# Operational Mode Commands (VRRP)

- `show vrrp`

## show vrrp

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show vrrp &lt;brief   detail   extensive   summary&gt; &lt;interface <i>interface-name</i>&gt; &lt;track interfaces&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Display information and status about VRRP groups.                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <p><b>none</b>—(Same as brief) Display brief status information about all VRRP interfaces.</p> <p><b>brief   detail   extensive   summary</b>—(Optional) Display the specified level of output.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display information and status about the specified VRRP interface.</p> <p><b>track interfaces</b>—(Optional) Display information and status about VRRP track interfaces.</p>                                                                                         |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">Configuring VRRP for IPv6 (CLI Procedure)</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>List of Sample Output</b>    | <p><a href="#">show vrrp on page 2757</a></p> <p><a href="#">show vrrp brief on page 2757</a></p> <p><a href="#">show vrrp detail (IPv6) on page 2757</a></p> <p><a href="#">show vrrp detail (Route Track) on page 2758</a></p> <p><a href="#">show vrrp extensive on page 2758</a></p> <p><a href="#">show vrrp interface on page 2759</a></p> <p><a href="#">show vrrp summary on page 2760</a></p> <p><a href="#">show vrrp track detail on page 2760</a></p> <p><a href="#">show vrrp track summary on page 2761</a></p> |
| <b>Output Fields</b>            | Table 243 lists the output fields for the <b>show vrrp</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                   |

**Table 243: show vrrp Output Fields**

| Field Name             | Field Description                                                            | Level of Output                        |
|------------------------|------------------------------------------------------------------------------|----------------------------------------|
| <b>Interface</b>       | Name of the logical interface.                                               | <b>none, brief, extensive, summary</b> |
| <b>Interface index</b> | Physical interface index number, which reflects its initialization sequence. | <b>extensive</b>                       |
| <b>Groups</b>          | Total number of VRRP groups configured on the interface.                     | <b>extensive</b>                       |



Table 243: show vrrp Output Fields (*continued*)

| Field Name                                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Level of Output                       |
|--------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------|
| <b>Active</b>                              | Total number of VRRP groups that are active (that is, whose interface state is either up or down).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>extensive</b>                      |
| <b>Interface VRRP PDU statistics</b>       | Nonerrored statistics for the logical interface: <ul style="list-style-type: none"> <li>• <b>Advertisement sent</b>—Number of VRRP advertisement protocol data units (PDUs) that the interface has transmitted.</li> <li>• <b>Advertisement received</b>—Number of VRRP advertisement PDUs received by the interface.</li> <li>• <b>Packets received</b>—Number of VRRP packets received for VRRP groups on the interface.</li> <li>• <b>No group match received</b>—Number of VRRP packets received for VRRP groups that do not exist on the interface.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>extensive</b>                      |
| <b>Interface VRRP PDU error statistics</b> | Errored statistics for the logical interface: <ul style="list-style-type: none"> <li>• <b>Invalid IPAH next type received</b>—Number of packets received that use the IP Authentication Header protocol (IPAH) and that do not encapsulate VRRP packets.</li> <li>• <b>Invalid VRRP ttl value received</b>—Number of packets received whose IP time-to-live (TTL) value is not 255.</li> <li>• <b>Invalid VRRP version received</b>—Number of packets received whose VRRP version is not 2.</li> <li>• <b>Invalid VRRP pdu type received</b>—Number of packets received whose VRRP PDU type is not 1.</li> <li>• <b>Invalid VRRP authentication type received</b>—Number of packets received whose VRRP authentication is not none, simple, or md5.</li> <li>• <b>Invalid VRRP IP count received</b>—Number of packets received whose VRRP IP count exceeds 8.</li> <li>• <b>Invalid VRRP checksum received</b>—Number of packets received whose VRRP checksum does not match the calculated value.</li> </ul> | <b>extensive</b>                      |
| <b>Physical interface</b>                  | Name of the physical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail, extensive</b>              |
| <b>Unit</b>                                | Logical unit number.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | All levels                            |
| <b>Address</b>                             | Address of the physical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>none, brief, detail, extensive</b> |
| <b>Index</b>                               | Physical interface index number, which reflects its initialization sequence.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>detail, extensive</b>              |
| <b>SNMP ifIndex</b>                        | SNMP index number for the physical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail, extensive</b>              |
| <b>VRRP-Traps</b>                          | Status of VRRP traps: <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail, extensive</b>              |

Table 243: show vrrp Output Fields (*continued*)

| Field Name                          | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Level of Output                        |
|-------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------|
| <b>Type and Address</b>             | Identifier for the address and the address itself: <ul style="list-style-type: none"> <li>• <b>lcl</b>—Configured local interface address.</li> <li>• <b>mas</b>—Address of the master virtual router. This address is displayed only when the local interface is acting as a backup router.</li> <li>• <b>vip</b>—Configured virtual IP addresses.</li> </ul>                                                                                                                                                                                                                                                                                         | <b>none, brief, summary</b>            |
| <b>Interface state or Int state</b> | State of the physical interface: <ul style="list-style-type: none"> <li>• <b>down</b>—The device is present and the link is unavailable.</li> <li>• <b>not present</b>—The interface is configured, but no physical device is present.</li> <li>• <b>unknown</b>—The VRRP process has not had time to query the kernel about the state of the interface.</li> <li>• <b>up</b>—The device is present and the link is established.</li> </ul>                                                                                                                                                                                                            | <b>none, brief, extensive, summary</b> |
| <b>Group</b>                        | VRRP group number.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>none, brief, extensive, summary</b> |
| <b>State</b>                        | VRRP state: <ul style="list-style-type: none"> <li>• <b>backup</b>—The interface is acting as the backup router interface.</li> <li>• <b>bringup</b>—VRRP is just starting, and the physical device is not yet present.</li> <li>• <b>idle</b>—VRRP is configured on the interface and is disabled. This can occur when VRRP is first enabled on an interface whose link is established.</li> <li>• <b>initializing</b>—VRRP is initializing.</li> <li>• <b>master</b>—The interface is acting as the master router interface.</li> <li>• <b>transition</b>—The interface is changing between being the backup and being the master router.</li> </ul> | <b>extensive</b>                       |
| <b>Priority</b>                     | Configured VRRP priority for the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail, extensive</b>               |
| <b>Advertisement interval</b>       | Configured VRRP advertisement interval.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail, extensive</b>               |
| <b>Authentication type</b>          | Configured VRRP authentication type: <b>none</b> , <b>simple</b> , or <b>md5</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>detail, extensive</b>               |
| <b>Preempt</b>                      | Whether preemption is allowed on the interface: <b>yes</b> or <b>no</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail, extensive</b>               |
| <b>Accept-data mode</b>             | Whether the interface is configured to accept packets destined for the virtual IP address: <b>yes</b> or <b>no</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>detail, extensive</b>               |
| <b>VIP count</b>                    | Number of virtual IP addresses that have been configured on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail, extensive</b>               |
| <b>VIP</b>                          | List of virtual IP addresses configured on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail, extensive</b>               |
| <b>Advertisement timer</b>          | Time until the advertisement timer expires.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail, extensive</b>               |

Table 243: show vrrp Output Fields (*continued*)

| Field Name                  | Field Description                                                                                                                                                                                                                                                                                              | Level of Output   |
|-----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------|
| Master router               | IP address of the interface that is acting as the master. If the VRRP interface is down, the output is <b>N/A</b> .                                                                                                                                                                                            | detail, extensive |
| Virtual router uptime       | Time that the virtual router has been up.                                                                                                                                                                                                                                                                      | detail, extensive |
| Master router uptime        | Time that the master router has been up.                                                                                                                                                                                                                                                                       | detail, extensive |
| Virtual MAC                 | MAC address associated with the virtual IP address.                                                                                                                                                                                                                                                            | detail, extensive |
| Tracking                    | Whether tracking is <b>enabled</b> or <b>disabled</b> .                                                                                                                                                                                                                                                        | detail, extensive |
| Current priority            | Current operational priority for being the VRRP master.                                                                                                                                                                                                                                                        | detail, extensive |
| Configured priority         | Configured base priority for being the VRRP master.                                                                                                                                                                                                                                                            | detail, extensive |
| Priority hold-time          | Minimum time interval, in seconds, between successive changes to the current priority. <b>Disabled</b> indicates no minimum interval.                                                                                                                                                                          | detail, extensive |
| Remaining-time              | ( <b>track</b> option only) Displays the time remaining in the priority hold-time interval.                                                                                                                                                                                                                    | detail            |
| Interface tracking          | Whether interface tracking is enabled or disabled. When enabled, the output also displays the number of tracked interfaces.                                                                                                                                                                                    | detail extensive  |
| Interface/Tracked interface | Name of the tracked interface.                                                                                                                                                                                                                                                                                 | detail extensive  |
| Int state/Interface state   | Current operational state of the tracked interface: <b>up</b> or <b>down</b> .                                                                                                                                                                                                                                 | detail, extensive |
| Int speed/Speed             | Current operational speed, in bits per second, of the tracked interface.                                                                                                                                                                                                                                       | detail, extensive |
| Incurred priority cost      | Operational priority cost incurred due to the state and speed of this tracked interface. This cost is applied to the configured priority to obtain the current priority.                                                                                                                                       | detail, extensive |
| Threshold                   | Speed below which the corresponding priority cost is incurred. In other words, when the speed of the interface drops below the threshold speed, the corresponding priority cost is incurred.<br><br>An entry of <b>down</b> means that the corresponding priority cost is incurred when the interface is down. | detail, extensive |
| Route tracking              | Whether route tracking is enabled or disabled. When enabled, the output also displays the number of tracked routes.                                                                                                                                                                                            | detail, extensive |
| Route count                 | The number of routes being tracked.                                                                                                                                                                                                                                                                            | detail, extensive |

Table 243: show vrrp Output Fields (*continued*)

| Field Name                               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Level of Output          |
|------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------|
| <b>Route</b>                             | The IP address of the route being tracked.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>detail, extensive</b> |
| <b>VRF name</b>                          | The VPN routing and forwarding (VRF) routing instance that the tracked route is in.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail, extensive</b> |
| <b>Route state</b>                       | The state of the route being tracked: <b>up</b> , <b>down</b> , or <b>unknown</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail, extensive</b> |
| <b>Priority cost</b>                     | Configured priority cost. This value is incurred when the interface speed drops below the corresponding threshold or when the tracked route goes down.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail, extensive</b> |
| <b>Active</b>                            | Whether the threshold is active (*). If the threshold is active, the corresponding priority cost is incurred.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail, extensive</b> |
| <b>Group VRRP PDU statistics</b>         | Number of VRRP advertisements sent and received by the group.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>extensive</b>         |
| <b>Group VRRP PDU error statistics</b>   | <p>Errored statistics for the VRRP group:</p> <ul style="list-style-type: none"> <li>• <b>Bad authentication type received</b>—Number of VRRP PDUs received with an invalid authentication type. The received authentication can be <b>none</b>, <b>simple</b>, or <b>md5</b> and must be the same for all routers in the VRRP group.</li> <li>• <b>Bad password received</b>—Number of VRRP PDUs received with an invalid key (password). The password for simple authentication must be the same for all routers in the VRRP group</li> <li>• <b>Bad MD5 digest received</b>—Number of VRRP PDUs received for which the MD5 digest computed from the VRRP PDU differs from the digest expected by the VRRP instance configured on the router.</li> <li>• <b>Bad advertisement timer received</b>—Number of VRRP PDUs received with an advertisement time interval that is inconsistent with the one in use among the routers in the VRRP group.</li> <li>• <b>Bad VIP count received</b>—Number of VRRP PDUs whose virtual IP address counts differ from the count that has been configured on the VRRP instance.</li> <li>• <b>Bad VIPADDR received</b>—Number of VRRP PDUs whose virtual IP addresses differ from the list of virtual IP addresses configured on the VRRP instance.</li> </ul> | <b>extensive</b>         |
| <b>Group state transition statistics</b> | <p>State transition statistics for the VRRP group:</p> <ul style="list-style-type: none"> <li>• <b>Idle to master transitions</b>—Number of times that the VRRP instance transitioned from the idle state to the master state.</li> <li>• <b>Idle to backup transitions</b>—Number of times that the VRRP instance transitioned from the idle state to the backup state.</li> <li>• <b>Backup to master transitions</b>—Number of times that the VRRP instance transitioned from the backup state to the master state.</li> <li>• <b>Master to backup transitions</b>—Number of times that the VRRP instance transitioned from the master state to the backup state.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>extensive</b>         |
| <b>Vlan-id</b>                           | ID of Vlan                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>detail</b>            |

Table 243: show vrrp Output Fields (*continued*)

| Field Name | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Level of Output |
|------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| VR state   | VRRP information: <ul style="list-style-type: none"> <li>• <b>backup</b>—The interface is acting as the backup router interface.</li> <li>• <b>bringup</b>—VRRP is just starting, and the physical device is not yet present.</li> <li>• <b>idle</b>—VRRP is configured on the interface and is disabled. This can occur when VRRP is first enabled on an interface whose link is established.</li> <li>• <b>initializing</b>—VRRP is initializing.</li> <li>• <b>master</b>—The interface is acting as the master router interface.</li> <li>• <b>transition</b>—The interface is changing between being the backup and being the master router.</li> </ul> | none, brief     |
| Timer      | VRRP timer information: <ul style="list-style-type: none"> <li>• <b>A</b>—Time, in seconds, until the advertisement timer expires.</li> <li>• <b>D</b>—Time, in seconds, until the Master is Dead timer expires.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                  | none, brief     |

## Sample Output

### show vrrp

```

user@host> show vrrp
Interface State Group VR state Timer Type Address
ge-0/0/0.121 up 1 master A 1.052 1c1 gec0::12:1:1:1
 vip ge80::12:1:1:99
 vip gec0::12:1:1:99
ge-0/0/2.131 up 1 master A 0.364 1c1 gec0::13:1:1:1
 vip ge80::13:1:1:99
 vip gec0::13:1:1:99

```

### show vrrp brief

The output for the **show vrrp brief** command is identical to that for the **show vrrp** command. For sample output, see [show vrrp on page 2757](#).

### show vrrp detail (IPv6)

```

user@host> show vrrp detail
Physical interface: ge-0/0/0, Unit: 121, Vlan-id: 212, Address: gec0::12:1:1:1/120

Index: 67, SNMP ifIndex: 45, VRRP-Traps: enabled
Interface state: up, Group: 1, State: master
Priority: 200, Advertisement interval: 1, Authentication type: none
Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::12:1:1:99,
gec0::12:1:1:99
Advertisement timer: 1.121s, Master router: ge80::12:1:1:1
Virtual router uptime: 00:03:47, Master router uptime: 00:03:41
Virtual MAC: 00:00:5e:00:02:01
Tracking: disabled

```

Physical interface: ge-0/0/2, Unit: 131, Vlan-id: 213, Address: gec0::13:1:1:1/120

Index: 69, SNMP ifIndex: 47, VRRP-Traps: enabled  
 Interface state: up, Group: 1, State: master  
 Priority: 200, Advertisement interval: 1, Authentication type: none  
 Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::13:1:1:99,  
 gec0::13:1:1:99  
 Advertisement timer: 0.327s, Master router: ge80::13:1:1:1  
 Virtual router uptime: 00:03:47, Master router uptime: 00:03:41  
 Virtual MAC: 00:00:5e:00:02:01  
 Tracking: disabled

### show vrrp detail (Route Track)

user@host> show vrrp detail

Physical interface: ge-1/1/0, Unit: 0, Address: 30.30.30.30/24  
 Index: 67, SNMP ifIndex: 379, VRRP-Traps: enabled  
 Interface state: up, Group: 100, State: master  
 Priority: 150, Advertisement interval: 1, Authentication type: none  
 Preempt: yes, Accept-data mode: no, VIP count: 1, VIP: 30.30.30.100  
 Advertisement timer: 1.218s, Master router: 30.30.30.30  
 Virtual router uptime: 00:04:28, Master router uptime: 00:00:13  
 Virtual MAC: 00:00:5e:00:01:64  
 Tracking: enabled  
 Current priority: 150, Configured priority: 150  
 Priority hold-time: disabled  
 Interface tracking: disabled  
 Route tracking: enabled, Route count: 1

| Route           | VRF name | Route state | Priority cost |
|-----------------|----------|-------------|---------------|
| 192.168.40.0/22 | default  | up          | 30            |

### show vrrp extensive

user@host> show vrrp extensive

Interface: ge-0/0/0.121, Interface index: 67, Groups: 1, Active : 1

#### Interface VRRP PDU statistics

|                         |   |     |
|-------------------------|---|-----|
| Advertisement sent      | : | 188 |
| Advertisement received  | : | 0   |
| Packets received        | : | 0   |
| No group match received | : | 0   |

#### Interface VRRP PDU error statistics

|                                           |   |   |
|-------------------------------------------|---|---|
| Invalid IPAH next type received           | : | 0 |
| Invalid VRRP TTL value received           | : | 0 |
| Invalid VRRP version received             | : | 0 |
| Invalid VRRP PDU type received            | : | 0 |
| Invalid VRRP authentication type received | : | 0 |
| Invalid VRRP IP count received            | : | 0 |
| Invalid VRRP checksum received            | : | 0 |

Physical interface: ge-0/0/0, Unit: 121, Vlan-id: 212, Address: gec0::12:1:1:1/120

Index: 67, SNMP ifIndex: 45, VRRP-Traps: enabled  
 Interface state: up, Group: 1, State: master  
 Priority: 200, Advertisement interval: 1, Authentication type: none  
 Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::12:1:1:99,  
 gec0::12:1:1:99  
 Advertisement timer: 1.034s, Master router: ge80::12:1:1:1  
 Virtual router uptime: 00:04:04, Master router uptime: 00:03:58  
 Virtual MAC: 00:00:5e:00:02:01  
 Tracking: disabled  
 Group VRRP PDU statistics

```

 Advertisement sent : 188
 Advertisement received : 0
Group VRRP PDU error statistics
 Bad authentication type received: 0
 Bad password received : 0
 Bad MD5 digest received : 0
 Bad advertisement timer received: 0
 Bad VIP count received : 0
 Bad VIPADDR received : 0
Group state transition statistics
 Idle to master transitions : 0
 Idle to backup transitions : 1
 Backup to master transitions : 1
 Master to backup transitions : 0

Interface: ge-0/0/2.131, Interface index: 69, Groups: 1, Active : 1
Interface VRRP PDU statistics
 Advertisement sent : 186
 Advertisement received : 0
 Packets received : 0
 No group match received : 0
Interface VRRP PDU error statistics
 Invalid IPAH next type received : 0
 Invalid VRRP TTL value received : 0
 Invalid VRRP version received : 0
 Invalid VRRP PDU type received : 0
 Invalid VRRP authentication type received: 0
 Invalid VRRP IP count received : 0
 Invalid VRRP checksum received : 0

Physical interface: ge-0/0/2, Unit: 131, Vlan-id: 213, Address: gec0::13:1:1:1/120

Index: 69, SNMP ifIndex: 47, VRRP-Traps: enabled
Interface state: up, Group: 1, State: master
Priority: 200, Advertisement interval: 1, Authentication type: none
Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::13:1:1:99,
gec0::13:1:1:99
Advertisement timer: 0.396s, Master router: ge80::13:1:1:1
Virtual router uptime: 00:04:04, Master router uptime: 00:03:58
Virtual MAC: 00:00:5e:00:02:01
Tracking: disabled
Group VRRP PDU statistics
 Advertisement sent : 186
 Advertisement received : 0
Group VRRP PDU error statistics
 Bad authentication type received: 0
 Bad password received : 0
 Bad MD5 digest received : 0
 Bad advertisement timer received: 0
 Bad VIP count received : 0
 Bad VIPADDR received : 0
Group state transition statistics
 Idle to master transitions : 0
 Idle to backup transitions : 1
 Backup to master transitions : 1
 Master to backup transitions : 0

```

### show vrrp interface

user@host> show vrrp interface

```

Interface: ge-0/0/0.121, Interface index: 67, Groups: 1, Active : 1
Interface VRRP PDU statistics
 Advertisement sent : 205
 Advertisement received : 0
 Packets received : 0
 No group match received : 0
Interface VRRP PDU error statistics
 Invalid IPAH next type received : 0
 Invalid VRRP TTL value received : 0
 Invalid VRRP version received : 0
 Invalid VRRP PDU type received : 0
 Invalid VRRP authentication type received: 0
 Invalid VRRP IP count received : 0
 Invalid VRRP checksum received : 0

Physical interface: ge-0/0/0, Unit: 121, Vlan-id: 212, Address: gec0::12:1:1:1/120

Index: 67, SNMP ifIndex: 45, VRRP-Traps: enabled
Interface state: up, Group: 1, State: master
Priority: 200, Advertisement interval: 1, Authentication type: none
Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::12:1:1:99,
gec0::12:1:1:99
Advertisement timer: 0.789s, Master router: ge80::12:1:1:1
Virtual router uptime: 00:04:26, Master router uptime: 00:04:20
Virtual MAC: 00:00:5e:00:02:01
Tracking: disabled
Group VRRP PDU statistics
 Advertisement sent : 205
 Advertisement received : 0
Group VRRP PDU error statistics
 Bad authentication type received: 0
 Bad password received : 0
 Bad MD5 digest received : 0
 Bad advertisement timer received: 0
 Bad VIP count received : 0
 Bad VIPADDR received : 0
Group state transition statistics
 Idle to master transitions : 0
 Idle to backup transitions : 1
 Backup to master transitions : 1
 Master to backup transitions : 0

```

### show vrrp summary

```

user@host> show vrrp summary

```

| Interface  | State | Group | VR state | Type | Address     |
|------------|-------|-------|----------|------|-------------|
| ge-4/1/0.0 | up    | 1     | backup   | lcl  | 10.57.0.2   |
|            |       |       |          | vip  | 10.57.0.100 |

### show vrrp track detail

```

user@host> show vrrp track detail
Tracked interface: ae1.211
State: up, Speed: 400m
Incurred priority cost: 0

```

| Threshold | Priority cost | Active |
|-----------|---------------|--------|
| 400m      | 10            |        |
| 300m      | 60            |        |
| 200m      | 110           |        |
| 100m      | 160           |        |
| down      | 190           |        |



```
Tracking VRRP interface: ae0.210, Group: 1
VR State: master
Current priority: 200, Configured priority: 200
Priority hold-time: disabled, Remaining-time: 50.351
```

#### show vrrp track summary

```
user@host> show vrrp track summary
```

| Track if | State | Speed | VRRP if | Group | VR State | Current priority |
|----------|-------|-------|---------|-------|----------|------------------|
| ae1.211  | up    | 400m  | ae0.210 | 1     | master   | 200              |



# Interfaces Feature Guide for QFX10000 Switches



## PART 47

# Ethernet OAM Link Fault Management

- [Understanding Ethernet OAM Link Fault Management on page 2767](#)



# Understanding Ethernet OAM Link Fault Management

- [Understanding Ethernet OAM Link Fault Management for an EX Series Switch on page 2767](#)
- [Configuring Ethernet OAM Link Fault Management \(CLI Procedure\) on page 2768](#)
- [Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 2770](#)

## Understanding Ethernet OAM Link Fault Management for an EX Series Switch

---

Juniper Networks Junos operating system (Junos OS) for Juniper Networks EX Series Ethernet Switches allows the Ethernet interfaces on these switches to support the IEEE 802.3ah standard for the Operation, Administration, and Maintenance (OAM) of Ethernet in access networks. The standard defines OAM link fault management (LFM). You can configure IEEE 802.3ah OAM LFM on point-to-point Ethernet links that are connected either directly or through Ethernet repeaters. The IEEE 802.3ah standard meets the requirement for OAM capabilities even as Ethernet moves from being solely an enterprise technology to a WAN and access technology, and the standard remains backward-compatible with existing Ethernet technology.

Ethernet OAM provides the tools that network management software and network managers can use to determine how a network of Ethernet links is functioning. Ethernet OAM should:

- Rely only on the media access control (MAC) address or virtual LAN identifier for troubleshooting.
- Work independently of the actual Ethernet transport and function over physical Ethernet ports or a virtual service such as pseudowire.
- Isolate faults over a flat (or single operator) network architecture or nested or hierarchical (or multiprovider) networks.

The following OAM LFM features are supported on EX Series switches:

- Discovery and Link Monitoring

The discovery process is triggered automatically when OAM is enabled on the interface. The discovery process permits Ethernet interfaces to discover and monitor the peer

on the link if it also supports the IEEE 802.3ah standard. You can specify the discovery mode used for IEEE 802.3ah OAM support. In active mode, the interface discovers and monitors the peer on the link if the peer also supports IEEE 802.3ah OAM functionality. In passive mode, the peer initiates the discovery process. After the discovery process has been initiated, both sides participate in discovery. The switch performs link monitoring by sending periodic OAM protocol data units (PDUs) to advertise OAM mode, configuration, and capabilities.

You can specify the number of OAM PDUs that an interface can miss before the link between peers is considered down.

- Remote Fault Detection

Remote fault detection uses flags and events. Flags are used to convey the following: Link Fault means a loss of signal, Dying Gasp means an unrecoverable condition such as a power failure, and Critical Event means an unspecified vendor-specific critical event. You can specify the periodic OAM PDU sending interval for fault detection. The EX Series switch uses the Event Notification OAM PDU to notify the remote OAM device when a problem is detected. You can specify the action to be taken by the system when the configured link-fault event occurs.

- Remote Loopback Mode

Remote loopback mode ensures link quality between the switch and a remote peer during installation or troubleshooting. In this mode, when the interface receives a frame that is not an OAM PDU or a pause frame, it sends it back on the same interface on which it was received. The link appears to be in the active state. You can use the returned loopback acknowledgement to test delay, jitter, and throughput.

Junos OS can place a remote DTE into loopback mode (if remote loopback mode is supported by the remote DTE). When you place a remote DTE into loopback mode, the interface receives the remote loopback request and puts the interface into remote loopback mode. When the interface is in remote loopback mode, all frames except OAM PDUs are looped back without any changes made to the frames. OAM PDUs continue to be sent and processed.

**Related  
Documentation**

- [Configuring Ethernet OAM Link Fault Management \(CLI Procedure\) on page 2768](#)
- [Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 2770](#)

---

## Configuring Ethernet OAM Link Fault Management (CLI Procedure)

Ethernet OAM link fault management (LFM) can be used for physical link-level fault detection and management. The IEEE 802.3ah LFM works across point-to-point Ethernet links either directly or through repeaters.

To configure Ethernet OAM LFM using the CLI:



1. Enable IEEE 802.3ah OAM support on an interface:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name
```



**NOTE:** You can configure Ethernet OAM LFM on aggregated interfaces.



**NOTE:** The remaining steps are optional. You can choose which of these features to configure for Ethernet OAM LFM on your switch.

2. Specify whether the interface or the peer initiates the discovery process by configuring the link discovery mode to **active** or **passive** (**active** = interface initiates; **passive** = peer initiates):

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name link-discovery active
```

3. Configure a periodic OAM PDU-sending interval (in milliseconds) for fault detection:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name pdu-interval interval
```

4. Specify the number of OAM PDUs that an interface can miss before the link between peers is considered down:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name pdu-threshold threshold-value
```

5. Configure event threshold values on an interface for the local errors that trigger the sending of link event TLVs:

- Set the threshold value (in seconds) for sending frame-error events or taking the action specified in the action profile:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name event-thresholds frame-error count
```

- Set the threshold value (in seconds) for sending frame-period events or taking the action specified in the action profile:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name event-thresholds frame-period count
```

- Set the threshold value (in seconds) for sending frame-period-summary events or taking the action specified in the action profile:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name event-thresholds frame-period-summary count
```

- Set the threshold value (in seconds) for sending symbol-period events or taking the action specified in the action profile:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name event-thresholds symbol-period count
```



**NOTE:** You can disable the sending of link event TLVs.

To disable the sending of link event TLVs:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name negotiation-options no-allow-link-events
```

6. Create an action profile to define event fault flags and thresholds to be taken when the link fault event occurs. Then apply the action profile to one or more interfaces. (You can also apply multiple action profiles to a single interface.)

- a. Name the action profile:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set action-profile profile-name
```

- b. Specify actions to be taken by the system when the link fault event occurs:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set action-profile profile-name action syslog
user@switch# set action-profile profile-name action link-down
```

- c. Specify events for the action profile:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set action-profile profile-name event link-adjacency-loss
```



**NOTE:** For each action profile, you must specify at least one link event and one action. The actions are taken only when all of the events in the action profile are true. If more than one action is specified, all actions are executed. You can set a low threshold for a specific action such as logging the error and set a high threshold for another action such as system logging.

7. Set a remote interface into loopback mode so that all frames except OAM PDUs are looped back without any changes made to the frames. Set the remote DTE in loopback mode (the remote DTE must support remote-loopback mode) and then enable remote loopback support for the local interface.

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name remote-loopback
user@switch# set interface interface-name negotiation-options allow-remote-loopback
```

#### Related Documentation

- [Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 2770](#)
- [Understanding Ethernet OAM Link Fault Management for an EX Series Switch on page 2767](#)

---

## Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches

---

Junos OS for EX Series switches allows the Ethernet interfaces on these switches to support the IEEE 802.3ah standard for the Operation, Administration, and Maintenance (OAM) of Ethernet in access networks. The standard defines OAM link fault management (LFM). You can configure IEEE 802.3ah OAM LFM on point-to-point Ethernet links that are connected either directly or through Ethernet repeaters.

This example describes how to enable and configure OAM LFM on a Gigabit Ethernet interface:

- [Requirements on page 2771](#)
- [Overview and Topology on page 2771](#)
- [Configuring Ethernet OAM Link Fault Management on Switch 1 on page 2771](#)
- [Configuring Ethernet OAM Link Fault Management on Switch 2 on page 2772](#)
- [Verification on page 2773](#)

## Requirements

This example uses the following hardware and software components:

- Junos OS Release 9.4 or later for EX Series switches
- Two EX3200 or EX4200 switches connected directly

## Overview and Topology

Junos OS for EX Series switches allows the Ethernet interfaces on these switches to support the IEEE 802.3ah standard for the Operation, Administration, and Maintenance (OAM) of Ethernet in access networks. The standard defines OAM link fault management (LFM). You can configure IEEE 802.3ah OAM LFM on point-to-point Ethernet links that are connected either directly or through Ethernet repeaters.

This example uses two EX4200 switches connected directly. Before you begin configuring Ethernet OAM LFM on two switches, connect the two switches directly through a trunk interface.

## Configuring Ethernet OAM Link Fault Management on Switch 1

**CLI Quick Configuration** To quickly configure Ethernet OAM LFM, copy the following commands and paste them into the switch terminal window:

```
[edit protocols oam ethernet link-fault-management]
set interface ge-0/0/0
set interface ge-0/0/0 link-discovery active
set interface ge-0/0/0 pdu-interval 800
set interface ge-0/0/0 remote-loopback
```

**Step-by-Step Procedure** To configure Ethernet OAM LFM on switch 1:

1. Enable IEEE 802.3ah OAM support on an interface:
 

```
[edit protocols oam ethernet link-fault-management]
user@switch1# set interface (OAM LFM) ge-0/0/0
```
2. Specify that the interface initiates the discovery process by configuring the link discovery mode to **active**:
 

```
[edit protocols oam ethernet link-fault-management]
user@switch1# set interface ge-0/0/0 link-discovery active
```
3. Set the periodic OAM PDU-sending interval (in milliseconds) to 800 on switch 1:
 

```
[edit protocols oam ethernet link-fault-management]
user@switch1# set interface pdu-interval 800
```

4. Set a remote interface into loopback mode so that all frames except OAM PDUs are looped back without any changes made to the frames. Ensure that the remote DTE supports remote loopback mode. To set the remote DTE in loopback mode

```
[edit protocols oam ethernet link-fault-management]
user@switch1# set interface ge-0/0/0.0 remote-loopback
```

## Results

---

Check the results of the configuration:

```
[edit]
user@switch1# show

protocols {
 oam {
 ethernet {
 link-fault-management {
 interface ge-0/0/0 {
 pdu-interval 800;
 link-discovery active;
 remote-loopback;
 }
 }
 }
 }
}
```

## Configuring Ethernet OAM Link Fault Management on Switch 2

**CLI Quick Configuration** To quickly configure Ethernet OAM LFM on switch 2, copy the following commands and paste them into the switch terminal window:

```
[edit protocols oam ethernet link-fault-management]
set interface ge-0/0/1
set interface ge-0/0/1 negotiation-options allow-remote-loopback
```

**Step-by-Step Procedure** To configure Ethernet OAM LFM on switch 2:

1. Enable OAM on the peer interface on switch 2:

```
[edit protocols oam ethernet link-fault-management]
user@switch2# set interface ge-0/0/1
```

2. Enable remote loopback support for the local interface:

```
[edit protocols oam ethernet link-fault-management]
user@switch2# set interface ge-0/0/1 negotiation-options allow-remote-loopback
```

**Results** Check the results of the configuration:

```
[edit]
user@switch2# show

protocols {
 oam {
 ethernet {
 link-fault-management {
 interface ge-0/0/1 {
 negotiation-options {
```

```

 allow-remote-loopback;
 }
}
}
}
}

```

## Verification

### Verifying That OAM LFM Has Been Configured Properly

- Purpose** Verify that OAM LFM has been configured properly.
- Action** Use the `show oam ethernet link-fault-management` command:
- ```
user@switch1#show oam ethernet link-fault-management
```

Sample Output

```

Interface: ge-0/0/0.0
Status: Running, Discovery state: Send Any
Peer address: 00:19:e2:50:3b:e1
Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
Remote entity information:
Remote MUX action: forwarding, Remote parser action: forwarding
Discovery mode: active, Unidirectional mode: unsupported
Remote loopback mode: supported, Link events: supported
Variable requests: unsupported

```

- Meaning** When the output displays the MAC address and the discover state is **Send Any**, it means that OAM LFM has been configured properly.

- Related Documentation**
- [Configuring Ethernet OAM Link Fault Management \(CLI Procedure\) on page 2768](#)
 - [Understanding Ethernet OAM Link Fault Management for an EX Series Switch on page 2767](#)

PART 48

Generic Routing Encapsulation (GRE)

- [Understanding GRE on page 2777](#)

CHAPTER 112

Understanding GRE

- [Understanding Generic Routing Encapsulation on page 2777](#)
- [Configuring Generic Routing Encapsulation Tunneling on page 2780](#)
- [Verifying That Generic Routing Encapsulation Tunneling Is Working Correctly on page 2781](#)

Understanding Generic Routing Encapsulation

Generic routing encapsulation (GRE) provides a private, secure path for transporting packets through an otherwise public network by encapsulating (or tunneling) the packets.

This topic describes:

- [Overview of GRE on page 2777](#)
- [GRE Tunneling on page 2778](#)
- [Using a Firewall Filter to De-encapsulate GRE Traffic on a QFX5100 and OCX Series Switches on page 2779](#)
- [Configuration Limitations on page 2780](#)

Overview of GRE

GRE encapsulates data packets and redirects them to a device that de-encapsulates them and routes them to their final destination. This allows the source and destination switches to operate as if they have a virtual point-to-point connection with each other (because the outer header applied by GRE is transparent to the encapsulated payload packet). For example, GRE tunnels allow routing protocols such as RIP and OSPF to forward data packets from one switch to another switch across the Internet. In addition, GRE tunnels can encapsulate multicast data streams for transmission over the Internet.

GRE is described in RFC 2784 (obsoletes earlier RFCs 1701 and 1702). The switches support RFC 2784, but not completely. (For a list of limitations, see [“Configuration Limitations” on page 2780.](#))

As a *tunnel source router*, the switch encapsulates a payload packet for transport through the tunnel to a destination network. The payload packet is first encapsulated in a GRE packet, and then the GRE packet is encapsulated in a delivery protocol. The switch performing the role of a *tunnel remote router* extracts the tunneled packet and forwards the packet to its destination. Note that you can use one firewall term to terminate many GRE tunnels on a QFX5100 switch.

GRE Tunneling

Data is routed by the system to the GRE endpoint over routes established in the route table. (These routes can be statically configured or dynamically learned by routing protocols such as RIP or OSPF.) When a data packet is received by the GRE endpoint, it is de-encapsulated and routed again to its destination address.

GRE tunnels are *stateless*—that is, the endpoint of the tunnel contains no information about the state or availability of the remote tunnel endpoint. Therefore, the switch operating as a tunnel source router cannot change the state of the GRE tunnel interface to down if the remote endpoint is unreachable.

For details about GRE tunneling, see:

- [Encapsulation and De-Encapsulation on the Switch on page 2778](#)
- [Number of Source and Destination Tunnels Allowed on a Switch on page 2778](#)
- [Class of Service on GRE Tunnels on page 2779](#)
- [Applying Firewall Filters to GRE Traffic on page 2779](#)

Encapsulation and De-Encapsulation on the Switch

Encapsulation—A switch operating as a tunnel source router encapsulates and forwards GRE packets as follows:

1. When a switch receives a data packet (payload) to be tunneled, it sends the packet to the tunnel interface.
2. The tunnel interface encapsulates the data in a GRE packet and adds an outer IP header.
3. The IP packet is forwarded on the basis of the destination address in the outer IP header.

De-encapsulation—A switch operating as a tunnel remote router handles GRE packets as follows:

1. When the destination switch receives the IP packet from the tunnel interface, the outer IP header and GRE header are removed.
2. The packet is routed based on the inner IP header.

Number of Source and Destination Tunnels Allowed on a Switch

QFX5100 and OCX Series switches support as many as 512 GRE tunnels, including tunnels created with a firewall filter. That is, you can create a total of 512 GRE tunnels, regardless of which method you use.

EX switches support as many as 500 GRE tunnels between switches transmitting IPv4 or IPv6 payload packets over GRE. If a passenger protocol in addition to IPv4 and IPv6 is used, you can configure up to 333 GRE tunnels between the switches.

An EX switch can have a maximum of 20 tunnel source IP addresses configured, and each tunnel source IP can be configured with up to 20 destination IP addresses on a second switch. As a result, the two connected switches can have a maximum of 400 GRE tunnels. If the first switch is also connected to a third switch, the possible maximum number of tunnels is 500.

Class of Service on GRE Tunnels

When a network experiences congestion and delay, some packets might be dropped. Junos OS class of service (CoS) divides traffic into classes to which you can apply different levels of throughput and packet loss when congestion occurs and thereby set rules for packet loss. For details about CoS, see [Junos OS CoS for EX Series Switches Overview](#).

The following CoS components are available on a switch operating as a GRE tunnel source router or GRE tunnel remote router:

- At the GRE tunnel source—On a switch operating as a tunnel source router, you can apply CoS classifiers on an *ingress port* or on a *GRE port*, with the following results on CoS component support on tunneled packets:
 - Schedulers only—Based on the CoS classification on the ingress port, you can apply CoS schedulers on a GRE port of the switch to define output queues and control the transmission of packets through the tunnel after GRE encapsulation. However, you cannot apply CoS rewrite rules to these packets.
 - Schedulers and rewrite rules—Depending on the CoS classification on the GRE port, you can apply both schedulers and rewrite rules to the encapsulated packets transmitted through the tunnel.
- At the GRE tunnel endpoint—When the switch is a tunnel remote router, you can apply CoS classifiers on the GRE port and schedulers and rewrite rules on the egress port to control the transmission of a de-encapsulated GRE packet out from the egress port.

Applying Firewall Filters to GRE Traffic

Firewall filters provide rules that define whether to permit, deny, or forward packets that are transiting an interface on a switch. (For details, see [Firewall Filters for EX Series Switches Overview](#).) Because of the encapsulation and de-encapsulation performed by GRE, you are constrained as to where you can apply a firewall filter to filter tunneled packets and which header will be affected. [Table 244](#) identifies these constraints.

Table 244: Firewall Filter Application Points for Tunneled Packets

Endpoint Type	Ingress Interface	Egress Interface
Source (encapsulating)	inner header	outer header
Remote (de-encapsulating)	Cannot filter packets on ingress interface	inner header

Using a Firewall Filter to De-encapsulate GRE Traffic on a QFX5100 and OCX Series Switches

You can also use a firewall filter to de-encapsulate GRE traffic on switches. This feature provides significant benefits in terms of scalability, performance, and flexibility because

you don't need to create a tunnel interface to perform the de-encapsulation. For example, you can terminate many tunnels from multiple source IP addresses with one firewall term. See *Configuring a Firewall Filter to De-encapsulate GRE Traffic on a QFX5100 Switch* for information about how to configure a firewall filter for this purpose.

Configuration Limitations

Table 245 lists features that are not supported with GRE.

Table 245: Features Not Supported with GRE

EX Switches	QFX Switches
MPLS over GRE tunnels	MPLS over GRE tunnels
GRE keepalives	GRE keepalives
GRE keys, payload packet fragmentation, and sequence numbers for fragmented packets	GRE keys, payload packet fragmentation, and sequence numbers for fragmented packets
BGP dynamic tunnels	BGP dynamic tunnels
Outer IP address must be IPv4	Outer IP address must be IPv4
Virtual routing instances	
Bidirectional Forwarding Detection (BFD) protocol over GRE distributed mode	
OSPF limitation—Enabling OSPF on a GRE interface creates two equal-cost routes to the destination: one through the Ethernet network or uplink interface and the other through the tunnel interface. If data is routed through the tunnel interface, the tunnel might fail. To keep the interface operational, we recommend that you use a static route, disable OSPF on the tunnel interface, or configure the peer not to advertise the tunnel destination over the tunnel interface.	

- Related Documentation**
- [Configuring Generic Routing Encapsulation Tunneling \(CLI Procedure\)](#)
 - [Configuring Generic Routing Encapsulation Tunneling on page 2780](#)
 - [Configuring a Firewall Filter to De-encapsulate GRE Traffic on a QFX5100 Switch](#)

Configuring Generic Routing Encapsulation Tunneling

Generic routing encapsulation (GRE) provides a private, secure path for transporting packets through an otherwise public network by encapsulating (or tunneling) the packets. GRE tunneling is accomplished through tunnel endpoints that encapsulate or de-encapsulate traffic.

You can also use a firewall filter to de-encapsulate GRE traffic on QFX5100 and OCX Series switches. This feature provides significant benefits in terms of scalability, performance, and flexibility because you don't need to create a tunnel interface to perform the de-encapsulation. For example, you can terminate many tunnels from multiple source

IP addresses with one firewall term. For more information on this feature, see *Configuring a Firewall Filter to De-encapsulate GRE Traffic on a QFX5100 Switch*.

This topic describes:

1. [Configuring a GRE Tunnel on page 2781](#)

Configuring a GRE Tunnel

To configure a GRE tunnel interface:

1. Create a GRE interface with a unit number and address:

```
[edit interfaces]
user@switch# set gr-0/0/0 unit number family inet address
```



NOTE: The base name of the interface must be `gr-0/0/0`.

This is a pseudo interface, and the address you specify can be any IP address. The routing table must specify `gr-0/0/0.x` as the outgoing interface for any packets that will be tunneled.

If you configure a GRE interface on a QFX5100 switch that is a member of a Virtual Chassis and later change the Virtual Chassis member number of the switch, the name of the GRE interface does not change in any way (because it is a pseudo interface). For example, if you change the member number from `0` to `5`, the GRE interface name does *not* change from `gr-0/0/0.x` to `gr-5/0/0.x`.

2. Specify the tunnel source address for the logical interface:

```
[edit interfaces]
user@switch# set gr-0/0/0 unit number tunnel (Legacy Switches) source
source-address
```

3. Specify the destination address:

```
[edit interfaces]
user@switch# set gr-0/0/0 unit number tunnel (Legacy Switches) destination
destination-address
```

The destination address must be reachable through static or dynamic routing. If you use static routing, you must get the destination MAC address (for example, by using `ping`) before user traffic can be forwarded through the tunnel.

Related Documentation

- [Verifying That Generic Routing Encapsulation Tunneling Is Working Correctly on page 2781](#)
- [Understanding Generic Routing Encapsulation on page 2777](#)
- [Configuring a Firewall Filter to De-encapsulate GRE Traffic on a QFX5100 Switch](#)

Verifying That Generic Routing Encapsulation Tunneling Is Working Correctly

Purpose Verify that the generic routing encapsulation (GRE) interface is sending tunneled traffic.

Action Display status information about the specified GRE interface by using the command [show interfaces](#).

```
user@switch> show interfaces gr-0/0/0.0
Physical interface: gr-0/0/0, Enabled, Physical link is Up
Interface index: 132, SNMP ifIndex: 26
  Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

Logical interface gr-0/0/0.0 (Index 68) (SNMP ifIndex 47)
  Flags: Point-To-Point SNMP-Traps 16384
  IP-Header 1.1.1.2:1.1.1.1:47:df:64:0000000000000000 Encapsulation: GRE-NULL
  Input packets : 0
  Output packets: 0
  Protocol inet, MTU: 1476
    Flags: None
    Addresses, Flags: Is-Primary
      Local: 1.10.1.1
```

Meaning The output indicates that the GRE interface gr-0/0/0 is up. The output displays the name of the physical interface and the traffic statistics for this interface---the number of and the rate at which input and output bytes and packets are received and transmitted on the physical interface.

Related Documentation

- *Configuring Generic Routing Encapsulation Tunneling (CLI Procedure)*

PART 49

Interfaces

- [Understanding Interfaces on page 2785](#)

Understanding Interfaces

- [Interfaces Overview on page 2785](#)
- [Understanding Interface Naming Conventions on page 2787](#)
- [Understanding Interface Ranges on page 2795](#)
- [Understanding Management Interfaces on page 2796](#)
- [Understanding Port Ranges and System Modes on page 2798](#)
- [Configuring the Interface Address on page 2824](#)
- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2826](#)
- [Configuring Ethernet Loopback Capability on page 2831](#)
- [Configuring an LPM Table With Junos OS Release 13.2X51-D10 on page 2832](#)
- [Channelizing Interfaces on page 2835](#)
- [Monitoring Interface Status and Traffic on page 2844](#)
- [Troubleshooting Network Interfaces on page 2845](#)

Interfaces Overview

Juniper Networks devices have two types of interfaces: network interfaces and special interfaces. This topic provides brief information about these interfaces. For additional information, see the *Junos OS Network Interfaces Library for Routing Devices*.

- [Network Interfaces on page 2785](#)
- [Special Interfaces on page 2786](#)

Network Interfaces

Network interfaces connect to the network and carry network traffic. [Table 246](#) lists the types of network interfaces supported.

Table 246: Network Interface Types and Purposes

Type	Purpose
Aggregated Ethernet interfaces	You can group Ethernet interfaces at the physical layer to form a single link-layer interface, also known as a <i>link aggregation group (LAG)</i> or <i>bundle</i> . These aggregated Ethernet interfaces help to balance traffic and increase the uplink bandwidth.

Table 246: Network Interface Types and Purposes (*continued*)

Type	Purpose
Channelized Interfaces	<p>Depending on the device and software package, 40-Gbps QSFP+ ports can be configured to operate as the following types of interfaces:</p> <ul style="list-style-type: none"> • 10-Gigabit Ethernet interfaces (<i>xe</i>) • 40-Gigabit Ethernet interfaces (<i>et</i> and <i>xle</i>) • 40-Gigabit data plane uplink interfaces (<i>fte</i>) <p>When an <i>et</i> port is channelized to four <i>xe</i> ports, a colon is used to signify the four separate channels. For example, on a QFX3500 standalone switch with port 2 on PIC 1 configured as four 10-Gigabit Ethernet ports, the interface names are <i>xe-0/1/2:0</i>, <i>xe-0/1/2:1</i>, <i>xe-0/1/2:2</i>, and <i>xe-0/1/2:3</i>.</p> <p>NOTE: You cannot configure channelized interfaces to operate as Virtual Chassis ports.</p>
Ethernet Interfaces	You can configure Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet interfaces to connect to other servers, storage, and switches. You can configure 40-Gigabit data plane uplink ports to connect a Node device to an Interconnect devices as well as for Virtual Chassis ports (VCPs).
Fibre Channel interfaces	You can use Fibre Channel interfaces to connect the switch to a Fibre Channel over Ethernet (FCoE) forwarder or a Fibre Channel switch in a storage area network (SAN). You can configure Fibre Channel interfaces only on ports 0 through 5 and 42 through 47 on QFX3500 devices. Fibre Channel interfaces do not forward Ethernet traffic.
LAN access interfaces	You can use these interfaces to connect to other servers, storage, and switches. When you power on a QFX Series product and use the factory-default configuration, the software automatically configures interfaces in access mode for each of the network ports.
Multichassis aggregated Ethernet (MC-AE) interfaces	You can group a LAG on one standalone switch with a LAG on another standalone switch to create a MC-AE. The MC-AE provides load balancing and redundancy across the two standalone switches.
Tagged-access mode interfaces	You can use tagged-access interfaces to connect a switch to an access layer device. Tagged-access interfaces can accept VLAN-tagged packets from multiple VLANs.
Trunk interfaces	You can use trunk interfaces to connect to other switches or routers. To use a port for this type of connection, you must explicitly configure the port interface for trunk mode. The interfaces from the switches or routers must also be configured for trunk mode. In this mode, the interface can be in multiple VLANs and accept tagged packets from multiple devices. Trunk interfaces typically connect to other switches and to routers on the LAN.
Virtual Chassis ports (VCPs)	You can use Virtual Chassis ports to send and receive Virtual Chassis Control Protocol (VCCP) traffic, and to create, monitor, and maintain the Virtual Chassis. On QFX3500, QFX3600, QFX5100, and EX4600 standalone switches, you can configure 40-Gigabit Ethernet QSFP+ uplink ports (non-channelized) or fixed SFP+ 10-Gigabit Ethernet ports as VCPs by issuing the request virtual-chassis-vc-port-set CLI command.

Special Interfaces

Table 247 lists the types of special interfaces supported.

Table 247: Special Interface Types and Purposes

Type	Purpose
Console port	Each device has a serial console port, labeled CON or CONSOLE , for connecting tty-type terminals to the switch. The console port does not have a physical address or IP address associated with it. However, it is an interface in the sense that it provides access to the switch.
Loopback interface	A software-only virtual interface that is always up. The loopback interface provides a stable and consistent interface and IP address on the switch.
Management interface	<p>The management Ethernet interface provides an out-of-band method for connecting to a standalone switch and QFabric system.</p> <p>NOTE: On OCX Series switches, the em0 management interface always has the status up in show command outputs, even if the physical port is empty. The me0 interface is a virtual interface between Junos and the host operating system, therefore its status is independent from the status of the physical port.</p>
Routed VLAN interfaces (RVI and IRB interfaces)	<p>Layer 3 routed VLAN interfaces (called RVI in the original CLI, and called IRB in Enhanced Layer 2 Software) route traffic from one broadcast domain to another and perform other Layer 3 functions such as traffic engineering. These functions are typically performed by a router interface in a traditional network.</p> <p>The RVI or IRB functions as a logical router, eliminating the need for having both a switch and a router. The RVI or IRB must be configured as part of a broadcast domain or virtual private LAN service (VPLS) routing instance for Layer 3 traffic to be routed out of it.</p>

Related Documentation

- [Understanding Aggregated Ethernet Interfaces and LACP on page 2865](#)
- [Understanding Interface Naming Conventions on page 2787](#)
- [Understanding Layer 3 Logical Interfaces on page 2859](#)
- [Understanding Management Interfaces on page 2796](#)
- [Understanding Integrated Routing and Bridging on page 2124](#)
- [Overview of Fibre Channel](#)

Understanding Interface Naming Conventions

The QFX Series and the EX4600 device uses a naming convention for defining the interfaces that is similar to that of other platforms running under Juniper Networks Junos OS. This topic provides brief information about the naming conventions used for interfaces on the QFX Series and on EX4600 switches.

This topic describes:

- [Physical Part of an Interface Name on page 2788](#)
- [Logical Part of an Interface Name on a Switch Running QFabric Software Package on page 2794](#)
- [Logical Part of a Channelized Interface Name on a Switch Running Enhanced Layer 2 Software on page 2795](#)
- [Wildcard Characters in Interface Names on page 2795](#)

Physical Part of an Interface Name

Interfaces in Junos OS are specified as follows:

device-name:type-fpc/pic/port

The convention is as follows:

- *device-name*—(QFabric systems only) The *device-name* is either the serial number or the alias of the QFabric system component, such as a Node device, Interconnect device, or QFabric infrastructure. The name can contain a maximum of 128 characters and cannot contain any colons.
- *type*—The QFX Series and EX4600 device interfaces use the following media types:
 - **fc**—Fibre Channel interface
 - **ge**—Gigabit Ethernet interface
 - **xe**—10-Gigabit Ethernet interface
 - **xle**—40-Gigabit Ethernet interface (QFX3500, QFX3600, and QFX5100 switches running a QFabric software package)
 - **et**—40-Gigabit Ethernet interface (QFX3500, QFX3600, QFX5100, QFX10000, and EX4600 switches running Enhanced Layer 2 Software)
 - **fte**—40-Gigabit data plane uplink interface (QFX3500, QFX3600, and QFX5100 switches running a QFabric software package)
 - **me**—Management interface
 - **em**—Management interface on QFX5100 and EX4600 switches.
- *fpc*—Flexible PIC Concentrator. QFX Series interfaces use the following convention for the FPC number in interface names:
 - On QFX3500, QFX3600, QFX5100 devices running a QFabric software package, and QFX10002 switches, the FPC number is always **0**.

The FPC number indicates the slot number of the line card that contains the physical interface.

- On QFX3500, QFX3600, QFX5100, EX4600, and QFX10002 switches running Enhanced Layer 2 Software, the member ID of a member in a Virtual Chassis determines the FPC number.



NOTE: Every member in a Virtual Chassis must have a unique member ID, otherwise the Virtual Chassis will not be created.

- On standalone QFX5100, EX4600, and QFX10002 switches, the FPC number is always **0**.
- *pic*—QFX Series and EX4600 device interfaces use the following convention for the PIC (Physical Interface Card) number in interface names:

- On a QFX3500 switch running a QFabric software package, PIC 0 can support 48 ports, PIC 1 can support 16 10-Gigabit Ethernet ports, and PIC 2 can support 4 40-Gigabit Ethernet ports.
- On a QFX3500 switch running Enhanced Layer 2 software, PIC 0 can support 48 ports, and PIC 1 can support 16 10-Gigabit Ethernet ports, and 4 40-Gigabit Ethernet ports.
- On a QFX3500 Node device running a QFabric software package, PIC 0 can support 48 ports and PIC 1 can support four 40-Gigabit data plane uplink ports.
- On a QFX3600 switch running a QFabric software package, PIC 0 can support 64 10-Gigabit Ethernet ports, and PIC 1 can support 16 40-Gigabit Ethernet ports.
- On a QFX3600 switch running Enhanced Layer 2 software, PIC 0 can support 64 10-Gigabit Ethernet ports and can also support 16 40-Gigabit Ethernet ports.
- On a QFX3600 Node device running a QFabric software package, PIC 0 can support 56 10-Gigabit Ethernet ports, and PIC 1 can support 8 40-Gigabit data plane uplink ports, and up to 14 40-Gigabit Ethernet ports.
- On a QFX5100-48S switch running Enhanced Layer 2 software, PIC 0 provides six 40-Gbps QSFP+ ports and 48 10-Gigabit Ethernet interfaces.
- On an EX4600 device running Enhanced Layer 2 software, PIC 0 provides 4 40-Gbps QSFP+ ports and 24 10-Gigabit Ethernet interfaces. There are two expansion bays (PIC 1 and PIC 2), and you can insert QFX-EM-4Q expansion modules and EX4600-EM-8F expansion modules. The QFX-EM-4Q expansion module provide 4 40-Gbps QSFP+ ports. The EX4600-EM-8F expansion module provides 8 40-Gbps QSFP+ ports. You can insert any combination of expansion modules. For example, you can insert two EX4600-EM-8F expansion modules, two QFX-EM-4Q expansion modules, or one of each.
- On a QFX5100-48S switch running a QFabric software package, PIC 1 provides six 40-Gbps QSFP+ ports, and PIC 0 provides 48 10-Gigabit Ethernet interfaces.
- On a QFX5100-24Q switch running Enhanced Layer 2 software, PIC 0 provides 24 40-Gbps QSFP+ ports. PIC 1 and PIC 2 can each contain a QFX-EM-4Q expansion module, and each expansion module provides 4 40-Gbps QSFP+ ports
- On a QFX5100-96S switch running Enhanced Layer 2 software, PIC 0 provides 96 10-Gigabit Ethernet interfaces and 8 40-Gbps QSFP+ ports .
- On a QFX10002-36Q switch running Enhanced Layer 2 software, PIC 0 provides 144 10-Gigabit Ethernet interfaces, and 36 40-Gbps QSFP+ ports, and 12 100-Gigabit Ethernet interfaces.
- On a QFX10002-72Q switch running Enhanced Layer 2 software, PIC 0 provides 288 10-Gigabit Ethernet interfaces, and 72 40-Gbps QSFP+ ports, and 24 100-Gigabit Ethernet interfaces.
- On a QFX10008 switch running Enhanced Layer 2 software, PIC 0 provides one-thousand, one-hundred fifty two 10-Gigabit Ethernet interfaces, two-hundred

eighty-eight 40-Gbps QSFP+ ports, or two-hundred forty 10-Gigabit Ethernet interfaces.

- *port*—Interfaces use the following convention for port numbers:

- On a QFX3500 switch running a QFabric software package, there are 48 network access ports (10-Gigabit Ethernet) labeled 0 through 47 on PIC 0 and, 16 network access ports labeled 0 through 15 on PIC 1, and four 40-Gbps QSFP+ ports labeled Q0 through Q3 on PIC 2. You can use the QSFP+ ports to connect the Node device to Interconnect devices.

By default, the 40-Gbps QSFP+ ports are configured to operate as 10-Gigabit Ethernet ports. You can use QSFP+ to four SFP+ copper breakout cables to connect the 10-Gigabit Ethernet ports to other servers, storage, and switches. Optionally, you can choose to configure the QSFP+ ports as 40-Gigabit Ethernet ports (see *Configuring the QSFP+ Port Type on QFX3500 Standalone Switches*).

- On a QFX3500 switch running Enhanced Layer 2 software, there are 48 network access ports labeled 0 through 47 on PIC 0 and 4 40-Gbps QSFP+ ports labeled Q0 through Q3 on PIC 1. See [“Channelizing Interfaces” on page 2835](#) for information on how to configure and channelize the 40-Gbps QSFP+ ports.
- On a QFX3600 switch running a QFabric software package, there are 64 network access ports (10-Gigabit Ethernet) labeled Q0 through Q15 on PIC 0, and there are 16 network access ports (40-Gigabit Ethernet) labeled Q0 through Q15 on PIC 1.

By default, all the QSFP+ ports are configured to operate as 40-Gigabit Ethernet ports. Optionally, you can choose to configure the QSFP+ ports as 10-Gigabit Ethernet ports (see *Configuring the Port Type on QFX3600 Standalone Switches*) and use QSFP+ to four SFP+ copper breakout cables to connect the 10-Gigabit Ethernet ports to other servers, storage, and switches.

- On a QFX3600 Node device running a QFabric software package, PIC 0 can support up to 56 10-Gigabit Ethernet ports labeled Q2 through Q15, and PIC 1 can support up to 8 40-Gigabit data plane uplink ports labeled Q0 through Q7, and up to 14 40-Gigabit Ethernet ports labeled Q2 through Q15.

On a QFX3600 Node device, by default, four 40-Gbps QSFP+ ports (labeled Q0 through Q3) are configured for uplink connections between your Node device and your Interconnect devices, and twelve 40-Gbps QSFP+ ports (labeled Q4 through Q15) use QSFP+ to four SFP+ copper breakout cables to support up to 48 10-Gigabit Ethernet ports for connections to either endpoint systems (such as servers and storage devices) or external networks. Optionally, you can choose to configure the first eight ports (Q0 through Q7) for uplink connections between your Node device and your Interconnect devices, and ports Q2 through Q15 for 10-Gigabit Ethernet or 40-Gigabit Ethernet connections to either endpoint systems or external networks (see *Configuring the Port Type on QFX3600 Node Devices*).

- On a QFX3600 switch running Enhanced Layer 2 software, PIC 0 can support 64 network access ports (10-Gigabit Ethernet ports) labeled Q0 through Q15 and 16 40-Gigabit Ethernet ports labeled Q0 through Q15. See [“Channelizing Interfaces” on page 2835](#) for information on how to configure and channelize the 40-Gbps QSFP+ ports.

- On a QFX5100-48S switch running Enhanced Layer 2 software, PIC 0 can support 48 network access ports (10-Gigabit Ethernet ports) labeled 0 through 47 and 6 40-Gbps QSFP+ ports labeled 48 through 53. See [“Channelizing Interfaces” on page 2835](#) for information on how to configure and channelize the 40-Gbps QSFP+ ports.
- On an EX4600 switch running Enhanced Layer 2 software, PIC 0 can support 24 network access ports (10-Gigabit Ethernet ports) labeled 0 through 23 and 4 40-Gbps QSFP+ ports labeled 24 through 27. There are two expansion bays (PIC 1 and PIC 2), and you can insert QFX-EM-4Q expansion modules and EX4600-EM-8F expansion modules. The QFX-EM-4Q expansion module provide 4 40-Gbps QSFP+ ports. The EX4600-EM-8F expansion module provides 8 40-Gbps QSFP+ ports. You can insert any combination of expansion modules. For example, you can insert two EX4600-EM-8F expansion modules, two QFX-EM-4Q expansion modules, or one of each. See [“Channelizing Interfaces” on page 2835](#) for information on how to configure and channelize the 40-Gbps QSFP+ ports.
- On a QFX5100-48S switch running a QFabric software package, PIC 0 can support 48 network access ports (10-Gigabit Ethernet ports) labeled 0 through 47, and PIC 1 can support 6 40-Gbps QSFP+ ports labeled 0 through 5. See *Configuring the QSFP+ Port Type on QFX5100 Switches* for information on how to configure the port mode of 40-Gbps QSFP+ ports.
- On a QFX5100-24Q switch running Enhanced Layer 2 software, PIC 0 can support 24 40-Gbps QSFP+ ports labeled 0 through 23. PIC 1 and PIC 2 each support 4 40-Gbps QSFP+ port, for a total of eight 40-Gbps QSFP+ ports. See [“Channelizing Interfaces” on page 2835](#) for information on how to configure and channelize the 40-Gbps QSFP+ ports.



NOTE: You cannot channelize the 40-Gbps QSFP+ ports provided in the two QFX-EM-4Q expansion modules. Also, even though there is a total of 128 physical ports, only 104 logical ports can be channelized.

You can configure different system modes to achieve varying levels of port density on the QFX5100-24Q and QFX5100-96S switches. Depending on the system mode you configure, there are restrictions on which ports you can channelize. If you channelize ports that are restricted, the configuration is ignored. See *Configuring the System Mode* for information on how to configure the system mode.

- On a QFX5100-96S switch running Enhanced Layer 2 software, PIC 0 can support 96 10-Gigabit Ethernet ports labeled 0 through 95, and 8 40-Gbps QSFP+ ports labeled 96 through 103. See [“Channelizing Interfaces” on page 2835](#) for information on how to configure and channelize the 40-Gbps QSFP+ ports.



NOTE: You can only channelize the 40-Gbps QSFP+ ports provided in ports 96 and 100, because only 104 logical ports can be channelized.

You can configure different system modes to achieve varying levels of port density on the QFX5100-24Q and QFX5100-96S switches. Depending on the system mode

you configure, there are restrictions on which ports you can channelize. If you channelize ports that are restricted, the configuration is ignored. See *Configuring the System Mode* for information on how to configure the system mode.

- On a QFX10002-36Q switch running Enhanced Layer 2 software, there are 36 quad small-form factor pluggable plus (QSFP+) ports that support 40-Gigabit Ethernet optical transceivers. Out of these 36 ports, 12 ports are QSFP28 capable, which are dual speed 40- or 100-Gigabit Ethernet optical transceivers.

Each QSFP28 socket can be configured to support:

- 100-Gigabit Ethernet using 28-Gbps QSFP28 optical transceivers. When a QSFP28 transceiver is inserted into the ports marked with a fine black line underneath the socket and the port is configured for 100-Gigabit Ethernet, the two adjacent ports are disabled and the QSFP28 is enabled for 100-Gigabit Ethernet.
- 40-Gigabit Ethernet using QSFP+ optical transceivers.
- 10-Gigabit Ethernet using breakout cables. When configured for channelization, a breakout cable converts the 40-Gigabit Ethernet port into 4 independent 10-Gigabit Ethernet ports.

Any of the 36 ports 0 through 35 can be configured as either uplink or access ports. See [“Channelizing Interfaces” on page 2835](#) for information on how to configure and channelize the 40-Gbps QSFP+ ports.

Each of the 12 QSFP28 ports support:

- 100-Gigabit Ethernet QSFP28 transceivers
- 40-Gigabit Ethernet QSFP+ transceivers

Each of the 36 QSFP+ ports support:

- 40-Gigabit Ethernet QSFP+ transceivers
- Access ports
- On a QFX10002-72Q switch running Enhanced Layer 2 software, there are 72 quad small-form factor pluggable plus (QSFP+) ports that support 40-Gigabit Ethernet optical transceivers. Out of these 72 ports, 24 ports are QSFP28 capable, which are dual speed 40- or 100-Gigabit Ethernet optical transceivers.

Each QSFP28 socket can be configured to support:

- 100-Gigabit Ethernet using 28-Gbps QSFP28 optical transceivers. When a QSFP28 transceiver is inserted into the ports marked with a fine black line underneath the socket and the port is configured for 100-Gigabit Ethernet, the two adjacent ports are disabled and the QSFP28 is enabled for 100-Gigabit Ethernet.
- 40-Gigabit Ethernet using QSFP+ optical transceivers.
- 10-Gigabit Ethernet using breakout cables. When configured for channelization, a breakout cable converts the 40-Gigabit Ethernet port into 4 independent 10-Gigabit Ethernet ports.

Any of the 72 ports 0 through 71 can be configured as either uplink or access ports. See [“Channelizing Interfaces” on page 2835](#) for information on how to configure and channelize the 40-Gbps QSFP+ ports.

Each of the 24 QSFP28 ports support:

- 100-Gigabit Ethernet QSFP28 transceivers

Each of the 72 QSFP+ ports support:

- 40-Gigabit Ethernet QSFP+ transceivers

Each of the 36 QSFP+ ports support:

- 40-Gigabit Ethernet QSFP+ transceivers
- Access ports
- Uplink ports
- On a QFX10008 switch running Enhanced Layer 2 software, there are two line cards available:
 - QFX10000-36Q, a 36-port 40-Gigabit Ethernet quad small form-factor pluggable plus transceiver (QSFP+) or 12-port 100GbE QSFP28 line card
 - QFX10000-30C, a 30-port 100-Gigabit or 40-Gigabit Ethernet QSFP28 line card

The QFX10000-36Q line cards supports

- 36 quad small form-factor pluggable plus (QSFP+) ports that support 40-Gigabit Ethernet optical transceivers. Out of these 36 ports, 12 ports are QSFP28 capable. The QSFP+ ports are dual speed and can support either 40-Gigabit or 100-Gigabit Ethernet optical transceivers. The line card can support 10-Gigabit Ethernet by channelizing the 40-Gigabit ports. Channelization is supported on fiber break-out cable using standard structured cabling techniques.
- Each QSFP28 socket can be configured to support:
 - 100-Gigabit Ethernet using QSFP28 optical transceivers. When a QSFP28 transceiver is inserted into the ports marked with a fine black line underneath the socket and the port is configured for 100-Gigabit Ethernet, the two adjacent ports are disabled and the QSFP28 socket is enabled for 100-Gigabit Ethernet.
 - 40-Gigabit Ethernet using QSFP+ optical transceivers.
 - 10-Gigabit Ethernet using breakout cabling and attached optical transceivers. When configured for channelization, the system converts the 40-Gigabit Ethernet port into 4 independent 10-Gigabit Ethernet ports.

Any of the 36 ports 0 through 35 can be configured as either uplink or access ports. See [“Channelizing Interfaces” on page 2835](#) for information on how to configure and channelize the 40-Gbps QSFP+ ports.

- Each of the 12 QSFP28 ports supports:
 - 100-Gigabit Ethernet QSFP28 transceivers

- 40-Gigabit Ethernet QSFP+ transceivers
- Each of the 12 QSFP28 ports supports:
 - 100-Gigabit Ethernet QSFP28 transceivers
 - 40-Gigabit Ethernet QSFP+ transceivers

Each of the 36 QSFP+ ports support:

- 40-Gigabit Ethernet QSFP+ transceivers
- Access ports
- Uplink ports

The QFX10000-30C line cards supports

- Thirty 28-Gbps QSFP+ Pluggable Solution (QSFP28) cages that support either 40-Gigabit Ethernet or 100-Gigabit Ethernet optical transceivers. The QFX10000-30C ports auto detect the type of transceiver installed and set the configuration to the appropriate speed.
- Each QSFP28 socket can be configured to support:
 - 100-Gigabit Ethernet using QSFP28 optical transceivers. When a QSFP28 transceiver is inserted into the ports marked with a fine black line underneath the socket and the port is configured for 100-Gigabit Ethernet, the two adjacent ports are disabled and the QSFP28 socket is enabled for 100-Gigabit Ethernet.
 - 40-Gigabit Ethernet using QSFP+ optical transceivers.

See [“Channelizing Interfaces” on page 2835](#) for information on how to configure and channelize the 40-Gbps QSFP+ ports.

- Each of the 30 QSFP28 ports supports:
 - 100-Gigabit Ethernet QSFP28 transceivers
 - 40-Gigabit Ethernet QSFP+ transceivers
 - Access ports
 - Uplink ports

Logical Part of an Interface Name on a Switch Running QFabric Software Package

The logical unit part of the interface name corresponds to the logical unit number, which can be a number from 0 through 16384. In the virtual part of the name, a period (.) separates the port and logical unit numbers: *device-name* (QFabric systems only): *type-fpc/pic/port.logical-unit-number*. For example, if you issue the **show ethernet-switching interfaces** command on a system with a default VLAN, the resulting display shows the logical interfaces associated with the VLAN:

Interface	State	VLAN members	Blocking
node-device1:xe-0/0/1.0	down	remote-analyzer	unblocked

```
node-device1:xe-0/0/2.0  down    default  unblocked
node-device1:xe-0/0/3.0  down    default  unblocked
```

When you configure aggregated Ethernet interfaces, you configure a logical interface, which is called a *bundle* or a *LAG*. Each LAG can include up to eight Ethernet interfaces, depending on the switch model.

Logical Part of a Channelized Interface Name on a Switch Running Enhanced Layer 2 Software

Channelizing enables you to configure four 10-Gigabit Ethernet interfaces from a 40-Gigabit Ethernet QSFP+ interface. By default, a 40-Gigabit Ethernet QSFP+ interface is named *et-fpc/pic/port*. The resulting 10-Gigabit Ethernet interfaces appear in the following format: *xe-fpc/pic/port:channel*, where channel can be a value of 0 through 3.

For example, if an *et* interface named **et-0/0/3** is channelized to four 10-Gigabit Ethernet interfaces, the resulting 10-Gigabit Ethernet interface names will be **xe-0/0/3:0**, **xe-0/0/3:1**, **xe-0/0/3:2**, and **xe-0/0/3:3**:

Interface	Admin	Link	Proto	Local	Remote
xe-0/0/3:0	up	down			
xe-0/0/3:1	up	down			
xe-0/0/3:2	up	down			
xe-0/0/3:3	up	down			

Wildcard Characters in Interface Names

In the **show interfaces** and **clear interfaces** commands, you can use wildcard characters in the *interface-name* option to specify groups of interface names without having to type each name individually. You must enclose all wildcard characters except the asterisk (*) in quotation marks (" ").

Related Documentation

- [Interfaces Overview on page 2785](#)
- [Channelizing Interfaces on page 2835](#)
- [Configuring the System Mode](#)
- [Understanding Management Interfaces on page 2796](#)
- [Understanding Port Ranges and System Modes on page 2798](#)
- [Rear Panel of a QFX3500 Device](#)
- [Front Panel of a QFX3600 Device](#)
- [Junos OS Network Interfaces Library for Routing Devices](#)

Understanding Interface Ranges

You can use the interface ranges to group interfaces of the same type that share a common configuration profile. This helps reduce the time and effort in configuring interfaces. The configurations common to all the interfaces can be included in the interface range definition.

The interface range definition contains the name of the interface range defined, the names of the individual member interfaces that do not fall in a series of interfaces, a range of interfaces defined in the member range, and the configuration statements common to all the interfaces. An interface range defined with member ranges and individual members but without any common configurations is also a valid definition.



NOTE: The interface range definition is supported only for Gigabit Ethernet, 10-Gigabit Ethernet, and Fibre Channel interfaces. OCX Series switches do not support Fibre Channel interfaces.

The common configurations defined in the interface range will be overridden by the local configuration.

The defined interface ranges can be used at places where the **interface** statement is used in the following configuration hierarchies:



NOTE: These statements are not supported on OCX Series switches:

- **protocols isis interface**
- **protocols sflow interfaces**

Related Documentation

- [Interfaces Overview on page 2785](#)
- [Interfaces Overview](#)
- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2826](#)
- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces](#)
- [Configuring Link Aggregation on page 2869](#)
- [Configuring a Layer 3 Logical Interface on page 2860](#)
- [Junos OS Network Interfaces Library for Routing Devices](#)
- [interface-range on page 3008](#)

Understanding Management Interfaces

You use management interfaces to access devices remotely. Typically, a management interface is not connected to the in-band network, but is connected to a device in the internal network. Through a management interface, you can access the device over the network using utilities such as **ssh** and **telnet** and configure it from anywhere, regardless of its physical location. As a security feature, users cannot log in as **root** through a management interface. To access the device as **root**, you must use the console port. You can also use **root** to log in using SSH.



NOTE: Before you can use management interfaces, you must configure the logical interfaces with valid IP addresses. Juniper Networks does not support configuring two management interfaces in the same subnet.

Management interface port ranges vary based on device type:

- QFX3500 devices:

The valid port range for a management interface (**me**) on a QFX3500 device is between 0 and 6, with a total of seven available ports. On a QFX3500 standalone switch, however, you can only configure **me0** and **me1** as management interfaces. The management interfaces are labeled **C0** and **C1**, and they correspond to **me0** and **me1**. On a QFX3500 Node device, the RJ-45 management interfaces and SFP management interfaces correspond to **me5** and **me6**.

- QFX3600 devices:

There are two RJ-45 management interfaces (labeled **C0** and **C1**) and two SFP management interfaces (labeled **C0S** and **C1S**). On a QFX3600 standalone switch, the RJ-45 management interfaces and SFP management interfaces correspond to **me0** and **me1**. On a QFX3600 Node device, the RJ-45 management interfaces and SFP management interfaces correspond to **me5** and **me6**. Each pair of management interfaces correspond to one Ethernet interface—for example, both RJ-45 management interfaces (labeled **C0** and **C0s**) can correspond to **me0**, and both SFP management interfaces (labeled **C1** and **C1S**) can correspond to **me1**. By default, both RJ-45 management interfaces are active. If you insert an SFP interface into the SFP management port (**C0S**, for example), the SFP interface would become the active management interface, and the corresponding RJ-45 management interface (**C0**) is disabled.



NOTE: On a QFX3600 device, you can use either the RJ-45 or the SFP management interfaces, but not both at the same time.

- On QFX5100, QFX5200, and EX4600 switches, there is one RJ-45 management interface (labeled **C0** and one SFP management interface (labeled **C1**), and they correspond to **em0** and **em1**. You can use both management interfaces simultaneously.
- On QFX10002 and QFX10008 switches, there is one RJ-45 management interface (labeled **MGMT** and one SFP management interface (labeled **MGMT**), and they correspond to **em0** and **em1**. Although the CLI permits you to configure two management Ethernet interfaces within the same subnet, only one interface is usable and supported.
- On OCX Series switches:

There is one RJ-45 management interface (labeled **MGMT**), which corresponds to **em0**. The **em0** interface always has the status **up** in show command outputs, even if the physical port is empty. The **me0** interface is a virtual interface between Junos and

the host operating system, therefore its status is independent from the status of the physical port.

- QFabric system:

On a QFabric system, there are management interfaces on the Node devices, Interconnect devices, and Director devices. However, you cannot access the management interfaces on the Node devices or Interconnect devices directly. You can only manage and configure these devices using the Director device. You can connect to the management interface over the network using utilities such as SSH.

Related Documentation • [Interfaces Overview on page 2785](#)

Understanding Port Ranges and System Modes

QFX Series devices and EX4600 switches can support different port ranges depending on the device, media type of the interface, the software that is running on the device, and the system mode.

This topic describes:

- [Port Ranges for Different Media Types on page 2798](#)
- [Supported System Modes on page 2822](#)

Port Ranges for Different Media Types

The following media types support the following port ranges:

- On a QFX3500 device:
 - The valid port range for a Fibre Channel (fc) interface is **0** through **5** and **42** through **47** on PIC **0**, with a total of 12 available Fibre Channel ports.



NOTE: Fibre Channel ports are not supported on QFX3500, QFX3600, and QFX5100 switches running Enhanced Layer 2 software.

- The valid port range for a Gigabit Ethernet (ge) interface is **6** through **41** on PIC **0** because the ports between **0** and **5** and **42** and **47** are reserved as Fibre Channel ports. The total number of available Gigabit Ethernet ports is 36, because 12 of the remaining 48 ports are reserved for Fibre Channel and 10-Gigabit Ethernet interfaces. Fibre Channel ports cannot be configured as Gigabit Ethernet ports.
- The valid port range for a 10-Gigabit Ethernet (xe) interface is **0** through **47** on PIC **0**. The valid port range for a 10-Gigabit Ethernet (xe) interface is **0** through **15** on PIC **1**. The total number of available 10-Gigabit Ethernet ports is 64.

- The valid port range for a 40-Gigabit data plane uplink interface is **0** through **3** on PIC 1
- The valid port range for a 40-Gigabit Ethernet interface is **0** through **3** on PIC 2. There are four available ports.
- On a QFX3600 Node device:
 - The valid port range for a 10-Gigabit Ethernet interface is **8** through **63** on PIC 0. There are 56 available ports.
 - The valid port range for a 40-Gigabit Ethernet interface is **2** through **15** on PIC 1. There are 14 available ports.
 - The valid port range for a 40-Gigabit data plane uplink interface is **0** through **7** on PIC 1. There are eight available ports.

See [Table 250](#) for physical port to logical port mappings.

- On a QFX3600 switch running Enhanced Layer 2 Software:
 - The valid port range for a 10-Gigabit Ethernet interface is **0** through **63** on PIC 0. There are 64 available ports.
 - The valid port range for a 40-Gigabit Ethernet interface is **0** through **15** on PIC 0. There are 16 available ports.

See [Table 251](#) for physical port to logical port mappings.

- On QFX5100-48S and QFX5100-48T switches running Enhanced Layer 2 Software:
 - The valid port range for a 10-Gigabit Ethernet interface is **0** through **47** on PIC 0. There are 48 available ports. When you channelize the 6 40-Gbps QSFP+ ports on **0** through **5** on PIC 1, there are 72 available ports.



NOTE: On PIC 1, ports 0 and 1 are reserved for fte ports. You cannot convert these fte ports to xe or xle ports.

- The valid port range for a 40-Gbps QSFP+ port is **0** through **5** on PIC 1. There are six available ports.

See [Table 253](#) for physical port to logical port mappings.

- On EX4600 switches running Enhanced Layer 2 Software:
 - The valid port range for a 10-Gigabit Ethernet interface is **0** through **23** on PIC 0. There are 24 available ports. When you channelize the 4 40-Gbps QSFP+ ports on **24** through **27** on PIC 0. There are 40 available ports.

See [Table 253](#) for physical port to logical port mappings.

- On QFX5100-48S and QFX5100-48T switches running a QFabric software package:
 - The valid port range for a 10-Gigabit Ethernet interface is **0** through **47** on PIC 0. There are 48 available ports.

- The valid port range for a 40-Gbps QSFP+ port is 0 through 5 on PIC 1. There are six available ports.



NOTE: On PIC 1, ports 0 and 1 are reserved for fte ports. You cannot convert these fte ports to xe or xle ports.

See [Table 254](#) for physical port to logical port mappings.

- For QFX5100-24Q and QFX5100-96S switches running Enhanced Layer 2 Software, see [Table 255](#) for physical port to logical port mappings for different system modes.

Table 248: Valid Port Ranges on QFX3500 Switches Running QFabric Software Package

Port Number	Fibre Channel Interfaces (On PIC 0)	Gigabit Ethernet Interfaces (On PIC 0)	10-Gigabit Ethernet Interfaces (On PIC 0 and 1)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)	40-Gigabit Ethernet Interfaces (On PIC 2)
0	fc-0/0/0	Not supported on this port	xe-0/0/0	Not supported on this port	Not supported on this port
1	fc-0/0/1	Not supported on this port	xe-0/0/1	Not supported on this port	Not supported on this port
2	fc-0/0/2	Not supported on this port	xe-0/0/2	Not supported on this port	Not supported on this port
3	fc-0/0/3	Not supported on this port	xe-0/0/3	Not supported on this port	Not supported on this port
4	fc-0/0/4	Not supported on this port	xe-0/0/4	Not supported on this port	Not supported on this port
5	fc-0/0/5	Not supported on this port	xe-0/0/5	Not supported on this port	Not supported on this port
6	Not supported on this port	ge-0/0/6	xe-0/0/6	Not supported on this port	Not supported on this port
7	Not supported on this port	ge-0/0/7	xe-0/0/7	Not supported on this port	Not supported on this port
8	Not supported on this port	ge-0/0/8	xe-0/0/8	Not supported on this port	Not supported on this port
9	Not supported on this port	ge-0/0/9	xe-0/0/9	Not supported on this port	Not supported on this port
10	Not supported on this port	ge-0/0/10	xe-0/0/10	Not supported on this port	Not supported on this port

Table 248: Valid Port Ranges on QFX3500 Switches Running QFabric Software Package (*continued*)

Port Number	Fibre Channel Interfaces (On PIC 0)	Gigabit Ethernet Interfaces (On PIC 0)	10-Gigabit Ethernet Interfaces (On PIC 0 and 1)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)	40-Gigabit Ethernet Interfaces (On PIC 2)
11	Not supported on this port	ge-0/0/11	xe-0/0/11	Not supported on this port	Not supported on this port
12	Not supported on this port	ge-0/0/12	xe-0/0/12	Not supported on this port	Not supported on this port
13	Not supported on this port	ge-0/0/13	xe-0/0/13	Not supported on this port	Not supported on this port
14	Not supported on this port	ge-0/0/14	xe-0/0/14	Not supported on this port	Not supported on this port
15	Not supported on this port	ge-0/0/15	xe-0/0/15	Not supported on this port	Not supported on this port
16	Not supported on this port	ge-0/0/16	xe-0/0/16	Not supported on this port	Not supported on this port
17	Not supported on this port	ge-0/0/17	xe-0/0/17	Not supported on this port	Not supported on this port
18	Not supported on this port	ge-0/0/18	xe-0/0/18	Not supported on this port	Not supported on this port
19	Not supported on this port	ge-0/0/19	xe-0/0/19	Not supported on this port	Not supported on this port
20	Not supported on this port	ge-0/0/20	xe-0/0/20	Not supported on this port	Not supported on this port
21	Not supported on this port	ge-0/0/21	xe-0/0/21	Not supported on this port	Not supported on this port
22	Not supported on this port	ge-0/0/22	xe-0/0/22	Not supported on this port	Not supported on this port
23	Not supported on this port	ge-0/0/23	xe-0/0/23	Not supported on this port	Not supported on this port
24	Not supported on this port	ge-0/0/24	xe-0/0/24	Not supported on this port	Not supported on this port
25	Not supported on this port	ge-0/0/25	xe-0/0/25	Not supported on this port	Not supported on this port

Table 248: Valid Port Ranges on QFX3500 Switches Running QFabric Software Package (*continued*)

Port Number	Fibre Channel Interfaces (On PIC 0)	Gigabit Ethernet Interfaces (On PIC 0)	10-Gigabit Ethernet Interfaces (On PIC 0 and 1)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)	40-Gigabit Ethernet Interfaces (On PIC 2)
26	Not supported on this port	ge-0/0/26	xe-0/0/26	Not supported on this port	Not supported on this port
27	Not supported on this port	ge-0/0/27	xe-0/0/27	Not supported on this port	Not supported on this port
28	Not supported on this port	ge-0/0/28	xe-0/0/28	Not supported on this port	Not supported on this port
29	Not supported on this port	ge-0/0/29	xe-0/0/29	Not supported on this port	Not supported on this port
30	Not supported on this port	ge-0/0/30	xe-0/0/30	Not supported on this port	Not supported on this port
31	Not supported on this port	ge-0/0/31	xe-0/0/31	Not supported on this port	Not supported on this port
32	Not supported on this port	ge-0/0/32	xe-0/0/32	Not supported on this port	Not supported on this port
33	Not supported on this port	ge-0/0/33	xe-0/0/33	Not supported on this port	Not supported on this port
34	Not supported on this port	ge-0/0/34	xe-0/0/34	Not supported on this port	Not supported on this port
35	Not supported on this port	ge-0/0/35	xe-0/0/35	Not supported on this port	Not supported on this port
36	Not supported on this port	ge-0/0/36	xe-0/0/36	Not supported on this port	Not supported on this port
37	Not supported on this port	ge-0/0/37	xe-0/0/37	Not supported on this port	Not supported on this port
38	Not supported on this port	ge-0/0/38	xe-0/0/38	Not supported on this port	Not supported on this port
39	Not supported on this port	ge-0/0/39	xe-0/0/39	Not supported on this port	Not supported on this port
40	Not supported on this port	ge-0/0/40	xe-0/0/40	Not supported on this port	Not supported on this port

Table 248: Valid Port Ranges on QFX3500 Switches Running QFabric Software Package (continued)

Port Number	Fibre Channel Interfaces (On PIC 0)	Gigabit Ethernet Interfaces (On PIC 0)	10-Gigabit Ethernet Interfaces (On PIC 0 and 1)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)	40-Gigabit Ethernet Interfaces (On PIC 2)
41	Not supported on this port	ge-0/0/41	xe-0/0/41	Not supported on this port	Not supported on this port
42	fc-0/0/42	Not supported on this port	xe-0/0/42	Not supported on this port	Not supported on this port
43	fc-0/0/43	Not supported on this port	xe-0/0/43	Not supported on this port	Not supported on this port
44	fc-0/0/44	Not supported on this port	xe-0/0/44	Not supported on this port	Not supported on this port
45	fc-0/0/45	Not supported on this port	xe-0/0/45	Not supported on this port	Not supported on this port
46	fc-0/0/46	Not supported on this port	xe-0/0/46	Not supported on this port	Not supported on this port
47	fc-0/0/47	Not supported on this port	xe-0/0/47	Not supported on this port	Not supported on this port
Q0	Not supported on this port	Not supported on this port	xe-0/1/0 xe-0/1/1 xe-0/1/2 xe-0/1/3 NOTE: Supported on QFX3500 standalone switch only.	fte-0/1/0	xle-0/2/0
Q1	Not supported on this port	Not supported on this port	xe-0/1/4 xe-0/1/5 xe-0/1/6 xe-0/1/7 NOTE: Supported on QFX3500 standalone switch only.	fte-0/1/1	xle-0/2/1

Table 248: Valid Port Ranges on QFX3500 Switches Running QFabric Software Package (*continued*)

Port Number	Fibre Channel Interfaces (On PIC 0)	Gigabit Ethernet Interfaces (On PIC 0)	10-Gigabit Ethernet Interfaces (On PIC 0 and 1)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)	40-Gigabit Ethernet Interfaces (On PIC 2)
Q2	Not supported on this port	Not supported on this port	xe-0/1/8 xe-0/1/9 xe-0/1/10 xe-0/1/11 NOTE: Supported on QFX3500 standalone switch only.	fte-0/1/2	xle-0/2/2
Q3	Not supported on this port	Not supported on this port	xe-0/1/12 xe-0/1/13 xe-0/1/14 xe-0/1/15 NOTE: Supported on QFX3500 standalone switch only.	fte-0/1/3	xle-0/2/3

Table 249: Valid Port Ranges on QFX3500 Switches Running Enhanced Layer 2 Software

Port Number	Gigabit Ethernet Interfaces (On PIC 0)	10-Gigabit Ethernet Interfaces (On PIC 0 and 1)	40-Gigabit Ethernet Interfaces (On PIC 1)
0	Not supported on this port	xe-0/0/0	Not supported on this port
1	Not supported on this port	xe-0/0/1	Not supported on this port
2	Not supported on this port	xe-0/0/2	Not supported on this port
3	Not supported on this port	xe-0/0/3	Not supported on this port
4	Not supported on this port	xe-0/0/4	Not supported on this port
5	Not supported on this port	xe-0/0/5	Not supported on this port
6	ge-0/0/6	xe-0/0/6	Not supported on this port
7	ge-0/0/7	xe-0/0/7	Not supported on this port

Table 249: Valid Port Ranges on QFX3500 Switches Running Enhanced Layer 2 Software (*continued*)

Port Number	Gigabit Ethernet Interfaces (On PIC 0)	10-Gigabit Ethernet Interfaces (On PIC 0 and 1)	40-Gigabit Ethernet Interfaces (On PIC 1)
8	ge-0/0/8	xe-0/0/8	Not supported on this port
9	ge-0/0/9	xe-0/0/9	Not supported on this port
10	ge-0/0/10	xe-0/0/10	Not supported on this port
11	ge-0/0/11	xe-0/0/11	Not supported on this port
12	ge-0/0/12	xe-0/0/12	Not supported on this port
13	ge-0/0/13	xe-0/0/13	Not supported on this port
14	ge-0/0/14	xe-0/0/14	Not supported on this port
15	ge-0/0/15	xe-0/0/15	Not supported on this port
16	ge-0/0/16	xe-0/0/16	Not supported on this port
17	ge-0/0/17	xe-0/0/17	Not supported on this port
18	ge-0/0/18	xe-0/0/18	Not supported on this port
19	ge-0/0/19	xe-0/0/19	Not supported on this port
20	ge-0/0/20	xe-0/0/20	Not supported on this port
21	ge-0/0/21	xe-0/0/21	Not supported on this port
22	ge-0/0/22	xe-0/0/22	Not supported on this port
23	ge-0/0/23	xe-0/0/23	Not supported on this port
24	ge-0/0/24	xe-0/0/24	Not supported on this port
25	ge-0/0/25	xe-0/0/25	Not supported on this port
26	ge-0/0/26	xe-0/0/26	Not supported on this port
27	ge-0/0/27	xe-0/0/27	Not supported on this port
28	ge-0/0/28	xe-0/0/28	Not supported on this port
29	ge-0/0/29	xe-0/0/29	Not supported on this port

Table 249: Valid Port Ranges on QFX3500 Switches Running Enhanced Layer 2 Software (*continued*)

Port Number	Gigabit Ethernet Interfaces (On PIC 0)	10-Gigabit Ethernet Interfaces (On PIC 0 and 1)	40-Gigabit Ethernet Interfaces (On PIC 1)
30	ge-0/0/30	xe-0/0/30	Not supported on this port
31	ge-0/0/31	xe-0/0/31	Not supported on this port
32	ge-0/0/32	xe-0/0/32	Not supported on this port
33	ge-0/0/33	xe-0/0/33	Not supported on this port
34	ge-0/0/34	xe-0/0/34	Not supported on this port
35	ge-0/0/35	xe-0/0/35	Not supported on this port
36	ge-0/0/36	xe-0/0/36	Not supported on this port
37	ge-0/0/37	xe-0/0/37	Not supported on this port
38	ge-0/0/38	xe-0/0/38	Not supported on this port
39	ge-0/0/39	xe-0/0/39	Not supported on this port
40	ge-0/0/40	xe-0/0/40	Not supported on this port
41	ge-0/0/41	xe-0/0/41	Not supported on this port
42	Not supported on this port	xe-0/0/42	Not supported on this port
43	Not supported on this port	xe-0/0/43	Not supported on this port
44	Not supported on this port	xe-0/0/44	Not supported on this port
45	Not supported on this port	xe-0/0/45	Not supported on this port
46	Not supported on this port	xe-0/0/46	Not supported on this port
47	Not supported on this port	xe-0/0/47	Not supported on this port
Q0	Not supported on this port	xe-0/1/0:0 xe-0/1/0:1 xe-0/1/0:2 xe-0/1/0:3	et-0/1/0

Table 249: Valid Port Ranges on QFX3500 Switches Running Enhanced Layer 2 Software (*continued*)

Port Number	Gigabit Ethernet Interfaces (On PIC 0)	10-Gigabit Ethernet Interfaces (On PIC 0 and 1)	40-Gigabit Ethernet Interfaces (On PIC 1)
Q1	Not supported on this port	xe-0/1/1:0 xe-0/1/1:1 xe-0/1/1:2 xe-0/1/1:3	et-0/1/1
Q2	Not supported on this port	xe-0/1/2:0 xe-0/1/2:1 xe-0/1/2:2 xe-0/1/2:3	et-0/1/2
Q3	Not supported on this port	xe-0/1/3:0 xe-0/1/3:1 xe-0/1/3:2 xe-0/1/3:3	et-0/1/3

Table 250: Valid Port Ranges on QFX3600 Switches Running QFabric Software Package

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 1)
Q0	xe-0/0/0 xe-0/0/1 xe-0/0/2 xe-0/0/3	xle-0/1/0
Q1	xe-0/0/4 xe-0/0/5 xe-0/0/6 xe-0/0/7	xle-0/1/1

Table 250: Valid Port Ranges on QFX3600 Switches Running QFabric Software Package (*continued*)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 1)
Q2	xe-0/0/8	xle-0/1/2
	xe-0/0/9	
	xe-0/0/10	
	xe-0/0/11	
Q3	xe-0/0/12	xle-0/1/3
	xe-0/0/13	
	xe-0/0/14	
	xe-0/0/15	
Q4	xe-0/0/16	xle-0/1/4
	xe-0/0/17	
	xe-0/0/18	
	xe-0/0/19	
Q5	xe-0/0/20	xle-0/1/5
	xe-0/0/21	
	xe-0/0/22	
	xe-0/0/23	
Q6	xe-0/0/24	xle-0/1/6
	xe-0/0/25	
	xe-0/0/26	
	xe-0/0/27	
Q7	xe-0/0/28	xle-0/1/7
	xe-0/0/29	
	xe-0/0/30	
	xe-0/0/31	

Table 250: Valid Port Ranges on QFX3600 Switches Running QFabric Software Package (*continued*)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 1)
Q8	xe-0/0/32	xle-0/1/8
	xe-0/0/33	
	xe-0/0/34	
	xe-0/0/35	
Q9	xe-0/0/36	xle-0/1/9
	xe-0/0/37	
	xe-0/0/38	
	xe-0/0/39	
Q10	xe-0/0/40	xle-0/1/10
	xe-0/0/41	
	xe-0/0/42	
	xe-0/0/43	
Q11	xe-0/0/44	xle-0/1/11
	xe-0/0/45	
	xe-0/0/46	
	xe-0/0/47	
Q12	xe-0/0/48	xle-0/1/12
	xe-0/0/49	
	xe-0/0/50	
	xe-0/0/51	
Q13	xe-0/0/52	xle-0/1/13
	xe-0/0/53	
	xe-0/0/54	
	xe-0/0/55	

Table 250: Valid Port Ranges on QFX3600 Switches Running QFabric Software Package (*continued*)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 1)
Q14	xe-0/0/56	xle-0/1/14
	xe-0/0/57	
	xe-0/0/58	
	xe-0/0/59	
Q15	xe-0/0/60	xle-0/1/15
	xe-0/0/61	
	xe-0/0/62	
	xe-0/0/63	

Table 251: Valid Port Ranges on QFX3600 Switches Running Enhanced Layer 2 Software

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 0)
Q0	xe-0/0/0:0	et-0/0/0
	xe-0/0/0:1	
	xe-0/0/0:2	
	xe-0/0/0:3	
Q1	xe-0/0/1:0	et-0/0/1
	xe-0/0/1:1	
	xe-0/0/1:2	
	xe-0/0/1:3	
Q2	xe-0/0/2:0	et-0/0/2
	xe-0/0/2:1	
	xe-0/0/2:2	
	xe-0/0/2:3	

Table 251: Valid Port Ranges on QFX3600 Switches Running Enhanced Layer 2 Software (*continued*)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 0)
Q3	xe-0/0/3:0 xe-0/0/3:1 xe-0/0/3:2 xe-0/0/3:3	et-0/0/3
Q4	xe-0/0/4:0 xe-0/0/4:1 xe-0/0/4:2 xe-0/0/4:3	et-0/0/4
Q5	xe-0/0/5:0 xe-0/0/5:1 xe-0/0/5:2 xe-0/0/5:3	et-0/0/5
Q6	xe-0/0/6:0 xe-0/0/6:1 xe-0/0/6:2 xe-0/0/6:3	et-0/0/6
Q7	xe-0/0/7:0 xe-0/0/7:1 xe-0/0/7:2 xe-0/0/7:3	et-0/0/7
Q8	xe-0/0/8:0 xe-0/0/8:1 xe-0/0/8:2 xe-0/0/8:3	et-0/0/8

Table 251: Valid Port Ranges on QFX3600 Switches Running Enhanced Layer 2 Software (*continued*)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 0)
Q9	xe-0/0/9:0 xe-0/0/9:1 xe-0/0/9:2 xe-0/0/9:3	et-0/0/9
Q10	xe-0/0/10:0 xe-0/0/10:1 xe-0/0/10:2 xe-0/0/10:3	et-0/0/10
Q11	xe-0/0/11:0 xe-0/0/11:1 xe-0/0/11:2 xe-0/0/11:3	et-0/0/11
Q12	xe-0/0/12:0 xe-0/0/12:1 xe-0/0/12:2 xe-0/0/12:3	et-0/0/12
Q13	xe-0/0/13:0 xe-0/0/13:1 xe-0/0/13:2 xe-0/0/13:3	et-0/0/13
Q14	xe-0/0/14:0 xe-0/0/14:1 xe-0/0/14:2 xe-0/0/14:3	et-0/0/14

Table 251: Valid Port Ranges on QFX3600 Switches Running Enhanced Layer 2 Software (*continued*)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 0)
Q15	xe-0/0/15:0 xe-0/0/15:1 xe-0/0/15:2 xe-0/0/15:3	et-0/0/15

Table 252: Valid Port Ranges on QFX3600 Node Devices Running QFabric Software Package

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)	40-Gigabit Ethernet Interfaces (On PIC 1)
Q0	Not supported on this port	fte-0/1/0	xle-0/1/0
Q1	Not supported on this port	fte-0/1/1	xle-0/1/1
Q2	xe-0/0/8 xe-0/0/9 xe-0/0/10 xe-0/0/11	fte-0/1/2	xle-0/1/2
Q3	xe-0/0/12 xe-0/0/13 xe-0/0/14 xe-0/0/15	fte-0/1/3	xle-0/1/3
Q4	xe-0/0/16 xe-0/0/17 xe-0/0/18 xe-0/0/19	fte-0/1/4	xle-0/1/4
Q5	xe-0/0/20 xe-0/0/21 xe-0/0/22 xe-0/0/23	fte-0/1/5	xle-0/1/5

Table 252: Valid Port Ranges on QFX3600 Node Devices Running QFabric Software Package (continued)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)	40-Gigabit Ethernet Interfaces (On PIC 1)
Q6	xe-0/0/24	fte-0/1/6	xle-0/1/6
	xe-0/0/25		
	xe-0/0/26		
	xe-0/0/27		
Q7	xe-0/0/28	fte-0/1/7	xle-0/1/7
	xe-0/0/29		
	xe-0/0/30		
	xe-0/0/31		
Q8	xe-0/0/32	Not supported on this port	xle-0/1/8
	xe-0/0/33		
	xe-0/0/34		
	xe-0/0/35		
Q9	xe-0/0/36	Not supported on this port	xle-0/1/9
	xe-0/0/37		
	xe-0/0/38		
	xe-0/0/39		
Q10	xe-0/0/40	Not supported on this port	xle-0/1/10
	xe-0/0/41		
	xe-0/0/42		
	xe-0/0/43		
Q11	xe-0/0/44	Not supported on this port	xle-0/1/11
	xe-0/0/45		
	xe-0/0/46		
	xe-0/0/47		

Table 252: Valid Port Ranges on QFX3600 Node Devices Running QFabric Software Package (*continued*)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)	40-Gigabit Ethernet Interfaces (On PIC 1)
Q12	xe-0/0/48 xe-0/0/49 xe-0/0/50 xe-0/0/51	Not supported on this port	xle-0/1/12
Q13	xe-0/0/52 xe-0/0/53 xe-0/0/54 xe-0/0/55	Not supported on this port	xle-0/1/13
Q14	xe-0/0/56 xe-0/0/57 xe-0/0/58 xe-0/0/59	Not supported on this port	xle-0/1/14
Q15	xe-0/0/60 xe-0/0/61 xe-0/0/62 xe-0/0/63	Not supported on this port	xle-0/1/15

Table 253: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running Enhanced Layer 2 Software

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 0)
0	xe-0/0/0	Not supported on this port
1	xe-0/0/1	Not supported on this port
2	xe-0/0/2	Not supported on this port
3	xe-0/0/3	Not supported on this port
4	xe-0/0/4	Not supported on this port

Table 253: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running Enhanced Layer 2 Software (*continued*)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 0)
5	xe-0/0/5	Not supported on this port
6	xe-0/0/6	Not supported on this port
7	xe-0/0/7	Not supported on this port
8	xe-0/0/8	Not supported on this port
9	xe-0/0/9	Not supported on this port
10	xe-0/0/10	Not supported on this port
11	xe-0/0/11	Not supported on this port
12	xe-0/0/12	Not supported on this port
13	xe-0/0/13	Not supported on this port
14	xe-0/0/14	Not supported on this port
15	xe-0/0/15	Not supported on this port
16	xe-0/0/16	Not supported on this port
17	xe-0/0/17	Not supported on this port
18	xe-0/0/18	Not supported on this port
19	xe-0/0/19	Not supported on this port
20	xe-0/0/20	Not supported on this port
21	xe-0/0/21	Not supported on this port
22	xe-0/0/22	Not supported on this port
23	xe-0/0/23	Not supported on this port
24	xe-0/0/24	Not supported on this port
25	xe-0/0/25	Not supported on this port
26	xe-0/0/26	Not supported on this port

Table 253: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running Enhanced Layer 2 Software (*continued*)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 0)
27	xe-0/0/27	Not supported on this port
28	xe-0/0/28	Not supported on this port
29	xe-0/0/29	Not supported on this port
30	xe-0/0/30	Not supported on this port
31	xe-0/0/31	Not supported on this port
32	xe-0/0/32	Not supported on this port
33	xe-0/0/33	Not supported on this port
34	xe-0/0/34	Not supported on this port
35	xe-0/0/35	Not supported on this port
36	xe-0/0/36	Not supported on this port
37	xe-0/0/37	Not supported on this port
38	xe-0/0/38	Not supported on this port
39	xe-0/0/39	Not supported on this port
40	xe-0/0/40	Not supported on this port
41	xe-0/0/41	Not supported on this port
42	xe-0/0/42	Not supported on this port
43	xe-0/0/43	Not supported on this port
44	xe-0/0/44	Not supported on this port
45	xe-0/0/45	Not supported on this port
46	xe-0/0/46	Not supported on this port
47	xe-0/0/47	Not supported on this port

Table 253: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running Enhanced Layer 2 Software (*continued*)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 0)
48	xe-0/0/48:0	et-0/0/48
	xe-0/0/48:1	
	xe-0/0/48:2	
	xe-0/0/48:3	
49	xe-0/0/49:0	et-0/0/49
	xe-0/0/49:1	
	xe-0/0/49:2	
	xe-0/0/49:3	
50	xe-0/0/50:0	et-0/0/50
	xe-0/0/50:1	
	xe-0/0/50:2	
	xe-0/0/50:3	
51	xe-0/0/51:0	et-0/0/51
	xe-0/0/51:1	
	xe-0/0/51:2	
	xe-0/0/51:3	
52	xe-0/0/52:0	et-0/0/52
	xe-0/0/52:1	
	xe-0/0/52:2	
	xe-0/0/52:3	
53	xe-0/0/53:0	et-0/0/53
	xe-0/0/53:1	
	xe-0/0/53:2	
	xe-0/0/53:3	

Table 254: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running QFabric Software Package

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 1)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)
0	xe-0/0/0	Not supported on this port	Not supported on this port
1	xe-0/0/1	Not supported on this port	Not supported on this port
2	xe-0/0/2	Not supported on this port	Not supported on this port
3	xe-0/0/3	Not supported on this port	Not supported on this port
4	xe-0/0/4	Not supported on this port	Not supported on this port
5	xe-0/0/5	Not supported on this port	Not supported on this port
6	xe-0/0/6	Not supported on this port	Not supported on this port
7	xe-0/0/7	Not supported on this port	Not supported on this port
8	xe-0/0/8	Not supported on this port	Not supported on this port
9	xe-0/0/9	Not supported on this port	Not supported on this port
10	xe-0/0/10	Not supported on this port	Not supported on this port
11	xe-0/0/11	Not supported on this port	Not supported on this port
12	xe-0/0/12	Not supported on this port	Not supported on this port
13	xe-0/0/13	Not supported on this port	Not supported on this port
14	xe-0/0/14	Not supported on this port	Not supported on this port
15	xe-0/0/15	Not supported on this port	Not supported on this port
16	xe-0/0/16	Not supported on this port	Not supported on this port
17	xe-0/0/17	Not supported on this port	Not supported on this port
18	xe-0/0/18	Not supported on this port	Not supported on this port
19	xe-0/0/19	Not supported on this port	Not supported on this port
20	xe-0/0/20	Not supported on this port	Not supported on this port
21	xe-0/0/21	Not supported on this port	Not supported on this port

Table 254: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running QFabric Software Package (continued)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 1)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)
22	xe-0/0/22	Not supported on this port	Not supported on this port
23	xe-0/0/23	Not supported on this port	Not supported on this port
24	xe-0/0/24	Not supported on this port	Not supported on this port
25	xe-0/0/25	Not supported on this port	Not supported on this port
26	xe-0/0/26	Not supported on this port	Not supported on this port
27	xe-0/0/27	Not supported on this port	Not supported on this port
28	xe-0/0/28	Not supported on this port	Not supported on this port
29	xe-0/0/29	Not supported on this port	Not supported on this port
30	xe-0/0/30	Not supported on this port	Not supported on this port
31	xe-0/0/31	Not supported on this port	Not supported on this port
32	xe-0/0/32	Not supported on this port	Not supported on this port
33	xe-0/0/33	Not supported on this port	Not supported on this port
34	xe-0/0/34	Not supported on this port	Not supported on this port
35	xe-0/0/35	Not supported on this port	Not supported on this port
36	xe-0/0/36	Not supported on this port	Not supported on this port
37	xe-0/0/37	Not supported on this port	Not supported on this port
38	xe-0/0/38	Not supported on this port	Not supported on this port
39	xe-0/0/39	Not supported on this port	Not supported on this port
40	xe-0/0/40	Not supported on this port	Not supported on this port
41	xe-0/0/41	Not supported on this port	Not supported on this port
42	xe-0/0/42	Not supported on this port	Not supported on this port
43	xe-0/0/43	Not supported on this port	Not supported on this port

Table 254: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running QFabric Software Package (*continued*)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 1)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)
44	xe-0/0/44	Not supported on this port	Not supported on this port
45	xe-0/0/45	Not supported on this port	Not supported on this port
46	xe-0/0/46	Not supported on this port	Not supported on this port
47	xe-0/0/47	Not supported on this port	Not supported on this port
48	Not supported on this port	Not supported on this PIC	fte-0/1/0 <i>NOTE:</i> This interface is a fixed fte interface and cannot be changed to xle.
49	Not supported on this port	Not supported on this PIC	fte-0/1/1 <i>NOTE:</i> This interface is a fixed fte interface and cannot be changed to xle.
50	Not supported on this port	xle-0/1/2	fte-0/1/2 <i>NOTE:</i> By default, this interface is an fte interface but can be configured as an xle interface.
51	Not supported on this port	xle-0/1/3	fte-0/1/3 <i>NOTE:</i> By default, this interface is an fte interface but can be configured as an xle interface.
52	Not supported on this port	xle-0/1/4 <i>NOTE:</i> By default, this interface is an xle interface but can be configured as an fte interface.	fte-0/1/4
53	Not supported on this port	xle-0/1/5 <i>NOTE:</i> By default, this interface is an xle interface but can be configured as an fte interface.	fte-0/1/5

Supported System Modes



NOTE: There are restrictions on the ports you can channelize on the QFX5100-24Q and QFX5100-96S switches depending on the system mode you configure. If you try to channelize ports that are restricted, the configuration is ignored.

The following system modes are available on the QFX5100-24Q switch:

- Default mode
- Mode-104-port
- Flexi-PIC mode
- Non-oversubscribed mode

See [Table 255](#) for more information regarding the supported system modes for your switch.

The following system modes are available on the QFX5100-96S switch:

- Default-mode
- Non-oversubscribed mode

See [Table 255](#) for more information regarding the supported system modes for your switch.

Table 255: System Modes Supported on QFX5100 Switches Running Enhanced Layer 2 Software

	Default-mode	Mode-104port	Flexi-pic-mode	Non-oversubscribed-mode
QFX5100-48S and QFX5100-48T	Not supported	Not supported	Not supported	Not supported

Table 255: System Modes Supported on QFX5100 Switches Running Enhanced Layer 2 Software (*continued*)

	Default-mode	Mode-104port	Flexi-pic-mode	Non-oversubscribed-mode
QFX5100-24Q	<p>Supported</p> <p>You do not need to configure the switch to be in this mode. On PIC 0, you can channelize all 24 40-Gbps QSFP+ ports. On PIC 1 and PIC 2, the 40-Gbps QSFP+ ports in the expansion modules are supported but cannot be channelized. In this mode, you can have one of two port combinations: 32 40-Gbps QSFP+ ports, or 96 10-Gigabit Ethernet ports plus 8 40-Gbps QSFP+ ports.</p>	<p>Supported</p> <p>On PIC 0, all 24 40-Gbps QSFP+ ports are channelized by default, which provides 96 10-Gigabit Ethernet ports. 40-Gbps QSFP+ ports contained in an expansion module on PIC 1 are supported. On PIC 1, ports 0 and 2 are channelized by default, and ports 1 and 3 are disabled. If 40-Gbps QSFP+ ports contained in an expansion module are detected on PIC 2, they are ignored.</p>	<p>Supported</p> <p>On PIC 0, the first four ports (ports 0 through 3) cannot be channelized. 40-Gbps QSFP+ ports contained in expansion modules on PIC 1 and PIC 2 are supported but cannot be channelized.</p>	<p>Supported</p> <p>All 24 40-Gbps QSFP+ ports on PIC 0 can be channelized to 96 10-Gigabit Ethernet ports. 40-Gbps QSFP+ ports contained in the expansion modules on PIC 1 and PIC 2 are not supported and cannot be channelized. There is no packet loss for packets of any size in this mode.</p>
QFX5100-96S	<p>Supported</p> <p>You do not need to configure the switch to be in this mode. On PIC 0, all 96 10-Gigabit Ethernet ports are supported. You can only channelize the 40-Gbps QSFP+ interfaces to 10-Gigabit Ethernet interfaces on ports 96 and 100. When you channelize the interfaces on ports 96 and 100, ports 97, 98, 99, 101, 102 and 103 are disabled.</p>	<p>Not supported</p>	<p>Not supported</p>	<p>Supported</p> <p>On PIC 0, all 96 10-Gigabit Ethernet ports are supported. However, the eight 40-Gbps QSFP+ ports are not supported and cannot be channelized. There is no packet loss for packets of any size in this mode.</p>

- Related Documentation**
- [Interfaces Overview on page 2785](#)
 - [Channelizing Interfaces on page 2835](#)
 - [Configuring the System Mode](#)
 - [Understanding Interface Naming Conventions on page 2787](#)
 - [Rear Panel of a QFX3500 Device](#)
 - [Front Panel of a QFX3600 Device](#)

Configuring the Interface Address

You assign an address to an interface by specifying the address when configuring the protocol family. For the **inet** or **inet6** family, configure the interface IP address. For the **iso** family, configure one or more addresses for the loopback interface. For the **ccc**, **ethernet-switching**, **tcc**, **mpls**, **tnp**, and **vppls** families, you never configure an address.



NOTE: The point-to-point (PPP) address is taken from the loopback interface address that has the primary attribute. When the loopback interface is configured as an unnumbered interface, it takes the primary address from the donor interface.

To assign an address to an interface, perform the following steps:

1. Configure the interface address at the **[edit interfaces *interface-name* unit *logical-unit-number* family *family*]** hierarchy level.
 - To configure an IPv4 address on routers and switches running Junos OS, use the **interface *interface-name* unit *number* family inet address *a.b.c.d/nn*** statement at the **[edit interfaces]** hierarchy level.

```
[edit interfaces ]
```

```
user@host# set interface-name unit logical-unit-number family inet address a.b.c.d/nn
```



NOTE:

- Juniper Networks routers and switches support /31 destination prefixes when used in point-to-point Ethernet configurations; however, they are not supported by many other devices, such as hosts, hubs, routers, or switches. You must determine if the peer system also supports /31 destination prefixes before configuration.
- You can configure the same IPv4 address on multiple physical interfaces. When you assign the same IPv4 address to multiple physical interfaces, the operational behavior of those interfaces differs, depending on whether they are implicitly or explicitly point-to-point .
- By default, all interfaces are assumed to be point-to-point (PPP) interfaces. For all interfaces except aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet, you can explicitly configure an interface to be a point-to-point connection.
- If you configure the same IP address on multiple interfaces in the same routing instance, Junos OS uses only the first configuration. The remaining IP address configurations are ignored, leaving some interfaces without an assigned address. Interfaces without an assigned address cannot be used as a donor interface for an unnumbered Ethernet interface.

- To configure an IPv6 address on routers and switches running Junos OS, use the **interface *interface-name* unit *number* family inet6 address *aaaa:bbbb:::zzzz/nn*** statement at the **[edit interfaces]** hierarchy level.

```
[edit interfaces ]
```

```
user@host# set interface-name unit logical-unit-number family inet6 address  
aaaa:bbbb:::zzzz/nn
```



NOTE:

- You represent IP version 6 (IPv6) addresses in hexadecimal notation using a colon-separated list of 16-bit values. The double colon (::) represents all bits set to 0.
- You must manually configure the router or switch advertisement and advertise the default prefix for autoconfiguration to work on a specific interface.

2. [Optional] Set the broadcast address on the network or subnet .

[edit interfaces *interface-name* unit *logical-unit-number* family *family* address *address*],
user@host# **set broadcast address**

.....



NOTE: The broadcast address must have a host portion of either all ones or all zeros. You cannot specify the addresses 0.0.0.0 or 255.255.255.255

.....

3. [Optional] specify the remote address of the connection for the encrypted, PPP-encapsulated, and tunnel interfaces.

[edit logical-systems *logical-system-name* interfaces *interface-name* unit
logical-unit-number family *family* **address** *address*]
user@host# **set destination address**

4. [Optional] For interfaces that carry IP version 6 (IPv6) traffic, configure the host to assign itself a unique 64-Bit IP Version 6 interface identifier (EUI-64).

[edit logical-systems *logical-system-name* interfaces *interface-name* unit
logical-unit-number family *family* **address** *address*]
user@host# **set eui-64**

**Related
Documentation**

- *Configuring Default, Primary, and Preferred Addresses and Interfaces*

Configuring Gigabit and 10-Gigabit Ethernet Interfaces

Devices include a factory default configuration that:

- Enables all 10-Gigabit Ethernet network interfaces on the switch
 - Sets a default port mode (access)
 - Sets default link settings
 - Specifies a logical unit (**unit 0**) and assigns it to **family ethernet-switching**
 - Configures Storm Control on all 10-Gigabit Ethernet network interfaces
 - Provides basic Rapid Spanning Tree Protocol (RSTP) and Link Layer Discovery Protocol (LLDP) configuration
-



NOTE: RSTP and LLDP are not supported on the OCX Series.

.....

The **ether-options** statement enables you to modify the following options:

- **802.3ad**—Specify an aggregated Ethernet bundle for both Gigabit Ethernet and 10-Gigabit Ethernet interfaces.
- **autonegotiation**—Enable or disable autonegotiation of flow control, link mode, and speed for interfaces.
- **link-mode**—Specify **full-duplex**, **half-duplex**, or **automatic** for Gigabit Ethernet interfaces.
- **loopback**—Enable or disable a loopback interface for both Gigabit Ethernet and 10-Gigabit Ethernet interfaces.

To set **ether-options** for both Gigabit Ethernet and 10-Gigabit Ethernet interfaces:

[edit]

```
user@switch# set interfaces interface-name ether-options
```

This topic describes:

- [Configuring Port Mode on QFX5100-48S, QFX5100-48T, QFX5100-24Q, and EX4600 Switches on page 2827](#)
- [Configuring the Link Settings for Gigabit Ethernet Interfaces on QFX5100-48S, QFX5100-24Q Switches, and EX4600 Switches on page 2828](#)
- [Configuring Gigabit Ethernet Interfaces on QFX5100-48T Switches on page 2828](#)
- [Configuring the Link Settings for 10-Gigabit Ethernet Interfaces on QFX5100-48S, QFX5100-24Q Switches, and EX4600 Switches on page 2830](#)
- [Configuring the Link Settings for 10-Gigabit Ethernet Interfaces on QFX5100-48T Switches on page 2830](#)
- [Configuring the IP Options on QFX5100-48S, QFX5100-48T, QFX5100-24Q, and EX4600 Switches on page 2830](#)

Configuring Port Mode on QFX5100-48S, QFX5100-48T, QFX5100-24Q, and EX4600 Switches

If you are connecting a switch to other switches and to routers on the LAN, you need to assign the interface to a logical port and you need to configure the logical port as a trunk port.

To configure a Gigabit Ethernet or 10-Gigabit interface for trunk port mode on the original CLI:

```
[edit]
user@switch# set interfaces interface-name unit logical-unit-number family ethernet-switching
port-mode trunk
```

To configure a Gigabit Ethernet or 10-Gigabit interface for trunk port mode on the Enhanced Layer 2 software (ELS):

```
[edit]
user@switch# set interfaces interface-name unit logical-unit-number family ethernet-switching
interface-mode trunk
```

Configuring the Link Settings for Gigabit Ethernet Interfaces on QFX5100-48S, QFX5100-24Q Switches, and EX4600 Switches

Devices include a factory default configuration that enables Gigabit Ethernet interfaces with applicable link settings.

The following default configurations are available on Gigabit Ethernet interfaces:

- You cannot set the speed on these interfaces.
- Gigabit Ethernet interfaces operate in full-duplex mode.
- Autonegotiation is supported by default.

If for some reason you have disabled autonegotiation, you can enable it by issuing the **set interfaces *name* ether-options auto-negotiate** command.

To disable autonegotiation, issue the **delete interfaces *name* ether-options auto-negotiate** command.



NOTE: Do not use the **set interface *name* ether-options no-auto-negotiate** command to remove the autonegotiation configuration.

Issue the **show chassis interface extensive** command to see if autonegotiation is enabled or disabled and the negotiated speed of the interface.

Configuring Gigabit Ethernet Interfaces on QFX5100-48T Switches

Devices include a factory default configuration that enables Gigabit Ethernet interfaces with applicable link settings.

The following default configurations are available on Gigabit Ethernet interfaces:

- Gigabit Ethernet interfaces operate in full-duplex mode.
- Autonegotiation is enabled by default, and will autonegotiate the speed with the link partner. We recommend that you keep autonegotiation enabled for interfaces operating at 100M, 1G, and 10G.

To disable autonegotiation, issue the **delete interfaces *name* ether-options auto-negotiate** command.



.....

NOTE: Do not use the `set interface name ether-options no-auto-negotiate` command to remove the autonegotiation configuration.

.....

You can reenable autonegotiation it by issuing the `set interfaces name ether-options auto-negotiate` command.

Issue the `show chassis interface extensive` command to see if autonegotiation is enabled or disabled and the negotiated speed of the interface.

Configuring the Link Settings for 10-Gigabit Ethernet Interfaces on QFX5100-48S, QFX5100-24Q Switches, and EX4600 Switches

The following default configurations are available on 10-Gigabit Ethernet interfaces:

- All the 10-Gigabit Ethernet interfaces are set to **auto-negotiation**.
- Flow control for 10-Gigabit Ethernet interfaces is set to **enabled** by default. You can disable flow control by specifying the **no-flow-control** option.
- The speed cannot be configured.
- 10-Gigabit Ethernet interfaces operate in full-duplex mode by default.
- Autonegotiation is supported by default.

If for some reason you have disabled autonegotiation, you can enable it by issuing the **set interfaces *name* ether-options auto-negotiate** command.

To disable autonegotiation, issue the **delete interfaces *name* ether-options auto-negotiate** command.



NOTE: Do not use the **set interface *name* ether-options no-auto-negotiate** command to remove the autonegotiation configuration.

Issue the **show chassis interface extensive** command to see if autonegotiation is enabled or disabled and the negotiated speed of the interface.

Configuring the Link Settings for 10-Gigabit Ethernet Interfaces on QFX5100-48T Switches

The following default configurations are available on 10-Gigabit Ethernet interfaces:

- All the 10-Gigabit Ethernet interfaces are set to **auto-negotiation**.
- Flow control for 10-Gigabit Ethernet interfaces is set to **enabled** by default. You can disable flow control by specifying the **no-flow-control** option.
- The speed cannot be configured.
- 10-Gigabit Ethernet interfaces operate in full-duplex mode by default.
- Autonegotiation is supported by default, and will autonegotiate the speed with the link partner.

If for some reason you have disabled autonegotiation, you can enable it by issuing the **set interfaces *name* ether-options auto-negotiate** command.

Issue the **show chassis interface extensive** command to see if autonegotiation is enabled or disabled and the negotiated speed of the interface.

Configuring the IP Options on QFX5100-48S, QFX5100-48T, QFX5100-24Q, and EX4600 Switches

To specify an IP address for the logical unit:

```
[edit]
user@switch# set interfaces interface-name unit logical-unit-number family inet address ip-address
```

**Related
Documentation**

- [Monitoring Interface Status and Traffic on page 167](#)
- [show interfaces xe on page 3185](#)
- [show interfaces ge-](#)
- [speed on page 3029](#)
- [Understanding Interface Naming Conventions on page 2787](#)

Configuring Ethernet Loopback Capability

To place an interface in loopback mode, include the **loopback** statement:

```
loopback;
```

To return to the default—that is, to disable loopback mode—delete the **loopback** statement from the configuration:

```
[edit]
user@switch# delete interfaces interface-name ether-options loopback
```

To explicitly disable loopback mode, include the **no-loopback** statement:

```
no-loopback;
```

You can include the **loopback** and **no-loopback** statements at the following hierarchy levels:

- **[edit interfaces *interface-name* aggregated-ether-options]**
- **[edit interfaces *interface-name* ether-options]**

**Related
Documentation**

- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2826](#)
- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces](#)

[Configuring an LPM Table With Junos OS Release 13.2X51-D10](#)

In addition to choosing a profile, you can further optimize memory allocation for LPM table entries by configuring how many IPv6 addresses with prefixes in the range /65 through /127 you want to store. If you want to use more than 16 IPv6 addresses with prefixes in this range, you must enter and commit the following statement:

[edit]

```
user@switch# set chassis forwarding-options profile-name num-65-127-prefix value
```

in which *value* can be a value in the range 1 through 128. Each increment adds support for 16 IPv6 addresses with prefixes between /65 and /127, for a maximum of 2048 such addresses (16 x 128 = 2048). The system supports 16 of these addresses by default, so to increase the number of supported addresses, you must enter a value of 2 or greater. For example, if you enter **2**, the system will support 32 IPv6 addresses with prefixes in the range /65 through /127.



NOTE: When you configure the `num-65-127-prefix` value, all the data interfaces on the switch restart. The management interfaces are unaffected.

The LPM table is shared, and each increment that you add for IPv6 addresses with prefixes in the range /65 through /127 reduces the number of forwarding table entries that are available for IPv4 addresses and IPv6 addresses with prefixes less than /65. [Table 256](#) provides examples of valid combinations that the LPM table can store, also using the **l2-profile-one** profile. Once again, each row in the table represents a case in which the table is full and cannot accommodate any more entries.

Table 256: Example LPM Table Combinations Using l2-profile-one With Junos OS 13.2X51-D10

IPv4 entries	IPv6 Entries (prefix <= 64)	IPv6 Entries (prefix >= 65)	num-65-127-prefix
16K	0K	16	1 (default)
0K	8K	16	1 (default)
8K	4K	16	1 (default)
4K	4K	1K	64
2K	5K	1K	64
0K	6K	1K	64
4K	2K	2K	128
2K	3K	2K	128
0K	4K	2K	128

[Table 257](#) provides examples of valid combinations that the LPM table can store when you use the **lpm-profile** profile. As before, each row represents a case in which the table is full and cannot accommodate any more entries.

Table 257: Example LPM Table Combinations Using lpm-profile With Junos OS 13.2X51-D10

IPv4 entries	IPv6 Entries (prefix <= 64)	IPv6 Entries (prefix >= 65)	num-65-127-prefix
128K	0K	16	1 (default)
0K	8K	16	1 (default)
8K	4K	16	1 (default)
4K	4K	1K	64
2K	5K	1K	64
0K	6K	1K	64
4K	2K	2K	128
2K	3K	2K	128
0K	4K	2K	128

Related Documentation

- *Configuring the Unified Forwarding Table*

Channelizing Interfaces



NOTE: On QFX10008 switches, channelization is supported on fiber break-out cables using standard structured cabling techniques. Channelization is not supported on the QFX10000-30C line card.



NOTE: On QFX10002 switches running on Junos OS Release 15.1X53-D10 or Junos OS Release 15.1X53-D15, when you delete and then reapply channelized interfaces, traffic is disrupted and might not be recovered.

The QFX3500, QFX3600, QFX5100, and EX4600, QFX10002, and QFX10008 switches provide 40-Gbps QSFP+ ports that can be channelized. Channelization allows you to configure 40-Gbps QSFP+ ports to operate as four 10-Gigabit Ethernet (xe) interfaces. You can use QSFP+ to four SFP+ breakout cables or QSFP+ transceivers with fiber breakout cables to connect the 10-Gigabit Ethernet ports to other servers, storage, and switches. By default, the four 40-Gbps QSFP+ ports operate as 40-Gigabit Ethernet (et) ports. When an et port is channelized to four xe ports, a colon is used to signify the four separate channels. For example, on a switch with port 2 on PIC 1 configured as four 10-Gigabit Ethernet ports, the interface names are xe-0/1/2:0, xe-0/1/2:1, xe-0/1/2:2, and xe-0/1/2:3.

By default, the 40-Gbps QSFP+ ports on EX4600 and QFX5100 switches are channelized automatically (auto-channelized) if any of the four channels on a 40-Gbps QSFP+ port receive data, unless you have configured channelization either at the chassis level or at the port level. Auto-channelization is not supported on interfaces contained in expansion modules, or on Virtual Chassis ports.

You can disable auto-channelization by including the **disable-auto-speed-detection** statement at the **[edit chassis fpc slot-number pic pic-number (port port-number | port-range port-range-low port-range-high) channel-speed]** hierarchy.

There are restrictions on the ports you can channelize on the QFX5100-24Q and QFX5100-96S switches, depending on the system mode you enable. If you try to channelize ports that are restricted, the configuration is ignored. See *Configuring the System Mode* for more information.

On QFX10002 and QFX10008 switches, there are 100-Gigabit Ethernet ports that work either as 100-Gigabit Ethernet or as 40-Gigabit Ethernet, but are recognized as 40-Gigabit Ethernet by default. You cannot channelize the 100-Gigabit Ethernet ports when they are operating as 100-Gigabit Ethernet interfaces. The 40-Gigabit Ethernet ports can operate independently or be channelized into four 10-Gigabit Ethernet ports as part of a port range. Ports cannot be channelized individually. Only the first and fourth port in each 6XQSFP cage is available to channelize as part of a port range. In a port range, the ports are bundled with the next two consecutive ports. For example, if you want to channelize ports 0 through 2, you would channelize port 0 only. If you try to channelize a port that is not supported, you will receive an error message when you commit the configuration. Auto-channelization is not supported on any ports.

When a 40-Gigabit Ethernet transceiver is inserted into a 100-Gigabit Ethernet port, the port recognizes the 40-Gigabit Ethernet port speed. When a 100-Gigabit Ethernet transceiver is inserted into the port and enabled in the CLI, the port recognizes the 100-Gigabit Ethernet speed and disables two adjacent 40-Gigabit Ethernet ports.

Table 258 provides detailed information on which ports are 100-Gigabit Ethernet, which ports can be channelized, and which ports are disabled when a 100-Gigabit Ethernet is inserted. On the QFX10008 switch with the QFX10000-36Q line card installed, only ports 0 through 35 are available. For more information, see *QFX10002-72Q Port Panel* and *QFX10000-36Q Line Card*.

Table 258: QFX10002-36Q Switch and QFX10000-36Q Line Card Port Mappings

Port Number	4X10 Gigabit Ethernet Port	4X10 Gigabit Channelized Port Group	40-Gigabit Ethernet (Default)	100-Gigabit Ethernet	100-Gigabit Ethernet Disables
0	✓	✓	✓	–	–
1	✓		✓	✓	0, 2
2	✓		✓	–	–
3	✓	✓	✓	–	–
4	✓		✓	–	–
5	✓		✓	✓	3, 4
6	✓	✓	✓	–	–
7	✓		✓	✓	6, 8
8	✓		✓	–	–
9	✓	✓	✓	–	–
10	✓		✓	–	–
11	✓		✓	✓	9, 10
12	✓	✓	✓	–	–
13	✓		✓	✓	12, 14
14	✓		✓	–	–

Table 258: QFX10002-36Q Switch and QFX10000-36Q Line Card Port Mappings (*continued*)

Port Number	4X10 Gigabit Ethernet Port	4X10 Gigabit Channelized Port Group	40-Gigabit Ethernet (Default)	100-Gigabit Ethernet	100-Gigabit Ethernet Disables
15	✓	✓	✓	–	–
16	✓		✓	–	–
17	✓		✓	✓	15, 16
18	✓	✓	✓	–	–
19	✓		✓	✓	18, 20
20	✓		✓	–	–
21	✓	✓	✓	–	–
22	✓		✓	–	–
23	✓		✓	✓	21, 22
24	✓	✓	✓	–	–
25	✓		✓	✓	24, 26
26	✓		✓	–	–
27	✓	✓	✓	–	–
28	✓		✓	–	–
29	✓		✓	✓	27, 28
30	✓	✓	✓	–	–
31	✓		✓	✓	30, 32
32	✓		✓	–	–
33	✓	✓	✓	–	–
34	✓		✓	–	–
35	✓		✓	✓	33, 34

Table 259: QFX10002-72Q Switch Port Mappings

Port Number	4X10 Gigabit Ethernet Port	4X10 Gigabit Channelized Port Group	40-Gigabit Ethernet (Default)	100-Gigabit Ethernet	100-Gigabit Ethernet Disables
0	✓	✓	✓	–	–
1	✓		✓	✓	0, 2
2	✓		✓	–	–
3	✓	✓	✓	–	–
4	✓		✓	–	–
5	✓		✓	✓	3, 4
6	✓	✓	✓	–	–
7	✓		✓	✓	6, 8
8	✓		✓	–	–
9	✓	✓	✓	–	–
10	✓		✓	–	–
11	✓		✓	✓	9, 10
12	✓	✓	✓	–	–
13	✓		✓	✓	12, 14
14	✓		✓	–	–
15	✓	✓	✓	–	–
16	✓		✓	–	–
17	✓		✓	✓	15, 16
18	✓	✓	✓	–	–
19	✓		✓	✓	18, 20
20	✓		✓	–	–

Table 259: QFX10002-72Q Switch Port Mappings (*continued*)

Port Number	4X10 Gigabit Ethernet Port	4X10 Gigabit Channelized Port Group	40-Gigabit Ethernet (Default)	100-Gigabit Ethernet	100-Gigabit Ethernet Disables
21	✓	✓	✓	–	–
22	✓		✓	–	–
23	✓		✓	✓	21, 22
24	✓	✓	✓	–	–
25	✓		✓	✓	24, 26
26	✓		✓	–	–
27	✓	✓	✓	–	–
28	✓		✓	–	–
29	✓		✓	✓	27, 28
30	✓	✓	✓	–	–
31	✓		✓	✓	30, 32
32	✓		✓	–	–
33	✓	✓	✓	–	–
34	✓		✓	–	–
35	✓		✓	✓	33, 34
36	✓	✓	✓	–	–
37	✓		✓	✓	36, 38
38	✓		✓	–	–
39	✓	✓	✓	–	–
40	✓		✓	–	–
41	✓		✓	✓	39, 40

Table 259: QFX10002-72Q Switch Port Mappings (*continued*)

Port Number	4X10 Gigabit Ethernet Port	4X10 Gigabit Channelized Port Group	40-Gigabit Ethernet (Default)	100-Gigabit Ethernet	100-Gigabit Ethernet Disables
42	✓	✓	✓	–	–
43	✓		✓	✓	42, 44
44	✓		✓	–	–
45	✓	✓	✓	–	–
46	✓		✓	–	–
47	✓		✓	✓	45, 46
48	✓	✓	✓	–	–
49	✓		✓	✓	48, 50
50	✓		✓	–	–
51	✓	✓	✓	–	–
52	✓		✓	–	–
53	✓		✓	✓	51, 52
54	✓	✓	✓	–	–
55	✓		✓	✓	54, 56
56	✓		✓	–	–
57	✓	✓	✓	–	–
58	✓		✓	–	–
59	✓		✓	✓	57, 58
60	✓	✓	✓	–	–
61	✓		✓	✓	60, 62
62	✓		✓	–	–

Table 259: QFX10002-72Q Switch Port Mappings (*continued*)

Port Number	4X10 Gigabit Ethernet Port	4X10 Gigabit Channelized Port Group	40-Gigabit Ethernet (Default)	100-Gigabit Ethernet	100-Gigabit Ethernet Disables
63	✓	✓	✓	–	–
64	✓		✓	–	–
65	✓		✓	✓	63, 64
66	✓	✓	✓	–	–
67	✓		✓	✓	66, 68
68	✓		✓	–	–
69	✓	✓	✓	–	–
70	✓		✓	–	–
71	✓		✓	✓	69, 70

The following steps describe how to configure a block of ports or an individual port to operate as 10-Gigabit Ethernet ports.

1. To configure a block of 40-Gigabit Ethernet (*et*) ports on QFX3500, QFX3600, QFX5100, EX4600 switches to operate as 10-Gigabit Ethernet ports, specify a port range and channel speed:

```
[edit chassis fpc fpc-slot pic pic-slot]
user@switch# set port-range port-range-low port-range-high channel-speed speed
```

For example, to configure ports 0 through 3 on PIC 1 to operate as 10-Gigabit Ethernet ports:

```
[edit chassis fpc 0 pic 1]
user@switch# set port-range 0 3 channel-speed 10g
```

2. To configure a range of ports on a QFX10002-72Q or QFX10008 switch to operate as 10-Gigabit Ethernet ports:



NOTE: The `port-range` statement is not available on QFX10002-72Q and QFX10008 switches. Instead, configure the port range using the `port` statement. Starting from port 0, you channelize every third port to channelize a group of three ports. For example, channelize port 0 to channelize ports 0 through 2, port 3 to channelize ports 3 through 5, and so on. See [Table 258](#) for port mapping information.

```
[edit chassis fpc fpc-slot pic pic-slot]
user@switch# set port port-number channel-speed speed
```

For example, to configure ports 0 through 2 on PIC 0 to operate as 10-Gigabit Ethernet ports:



NOTE: When you channelize port 0, ports 1 and 2 are also channelized.

```
[edit chassis fpc 0 pic 1]
user@switch# set port 0 channel-speed 10g
```

3. To configure an individual 40-Gigabit Ethernet (*et*) port on QFX3500, QFX3600, QFX5100, and EX4600 switches to operate as 10-Gigabit Ethernet (*xe*) ports, specify a port number and channel speed:

```
[edit chassis fpc 0 pic 0]
user@switch# set port port-number channel-speed speed
```

For example, to configure port 3 to operate as 10-Gigabit Ethernet ports:

```
[edit chassis fpc 0 pic 0]
user@switch# set port 3 channel-speed 10g
```

4. Review your configuration and issue the **commit** command.

```
[edit]
user@switch# commit
commit complete
```

5. To return a range of ports on QFX3500, QFX3600, QFX5100, and EX4600 switches to the default 40-Gigabit Ethernet configuration, delete the 10g statement:

```
[edit chassis fpc 0 pic 1]
user@switch# delete port-range port-range-low port-range-high channel-speed speed
```

For example, to return ports 0 through 3 to the default 40-Gigabit Ethernet configuration:

```
[edit chassis fpc 0 pic 1]
user@switch# delete port-range 0 3 channel-speed 10g
```

6. To return a range of ports on QFX10002-72Q or QFX10008 switches to the default 40-Gigabit Ethernet configuration, delete the 10g statement:

```
[edit chassis fpc 0 pic 1]
user@switch# delete port port-number channel-speed speed
```

For example, to return ports 0 through 2 to the default 40-Gigabit Ethernet configuration:

```
[edit chassis fpc 0 pic 1]
user@switch# delete port-0 channel-speed 10g
```

7. Review your configuration and issue the **commit** command.

```
[edit]
user@switch# commit
commit complete
```

8. To return a port on QFX3500, QFX3600, QFX5100, and EX4600 switches to the default 40-Gigabit Ethernet configuration, delete the 10g statement:

```
[edit chassis fpc 0 pic 0]
```

```
user@switch# delete port port-number channel-speed speed
```

For example, to return port 2 to the default 40-Gigabit Ethernet configuration:

```
[edit chassis fpc 0 pic 0]  
user@switch# delete port 2 channel-speed 10g
```

9. Review your configuration and issue the **commit** command.

```
[edit]  
user@switch# commit  
commit complete
```

The following steps describe how to disable auto-channelization at the port level on QFX3500, QFX3600, QFX5100, and EX4600 switches.

1. To disable auto-channelization at the port level, include the **disable** statement:

```
[edit]  
user@switch# set chassis fpc slot-number pic pic-number (port port-number |  
    port-range port-range-low port-range-high) channel-speed  
    disable-auto-speed-detection
```

For example, to disable auto-channelization for one port:

```
[edit]  
user@switch# set chassis fpc 0 pic 0 port 2 channel-speed  
    disable-auto-speed-detection
```

For example, to disable auto-channelization for a range of ports:

```
[edit]  
user@switch# set chassis fpc 0 pic 0 port-range 2 4 channel-speed  
    disable-auto-speed-detection
```

2. Review your configuration and issue the **commit** command.

```
[edit]  
user@switch# commit  
commit complete
```

- Related Documentation**
- *Configuring the System Mode*
 - [channel-speed on page 2979](#)
 - [fpc on page 2997](#)
 - [pic on page 3026](#)

Monitoring Interface Status and Traffic

- Purpose** View interface status to monitor interface bandwidth utilization and traffic statistics.
- Action**
- To view interface status for all the interfaces, enter [show interfaces xe](#).
 - To view status and statistics for a specific interface, enter [show interfaces xe interface-name](#).

- To view status and traffic statistics for all interfaces, enter either [show interfaces xe detail](#) or [show interfaces xe extensive](#).

Meaning For details about output from the CLI commands, see [show interfaces xe](#).

Troubleshooting Network Interfaces

The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP or SFP+ module is down

Problem **Description:** The switch has an SFP or SFP+ module installed. The interface on the port in which an SFP or SFP+ transceiver is installed is down.

Symptoms: When you check the status with the CLI command **show interfaces *interface-name***, the disabled port is not listed.

Cause By default, the SFP or SFP+ module operates in the 10-Gigabit Ethernet mode and supports only SFP or SFP+ transceivers. The operating mode for the module is incorrectly set.

Solution Only SFP or SFP+ transceivers can be installed in SFP or SFP+ modules. You must configure the operating mode of the SFP or SFP+ module to match the type of transceiver you want to use. For SFP+ transceivers, configure 10-Gigabit Ethernet operating mode.

PART 50

IP Directed Broadcast

- [Understanding IP Directed Broadcast on page 2849](#)

Understanding IP Directed Broadcast

- [Understanding IP Directed Broadcast on page 2849](#)
- [Configuring IP Directed Broadcast \(CLI Procedure\) on page 2851](#)
- [Example: Configuring IP Directed Broadcast on page 2852](#)

Understanding IP Directed Broadcast

IP directed broadcast helps you implement remote administration tasks such as backups and wake-on-LAN (WOL) application tasks by sending broadcast packets targeted at the hosts in a specified destination subnet. IP directed broadcast packets traverse the network in the same way as unicast IP packets until they reach the destination subnet. When they reach the destination subnet and IP directed broadcast is enabled on the receiving switch, the switch translates (*explodes*) the IP directed broadcast packet into a broadcast that floods the packet on the target subnet. All hosts on the target subnet receive the IP directed broadcast packet.

This topic covers:

- [IP Directed Broadcast Overview on page 2849](#)
- [IP Directed Broadcast Implementation on page 2850](#)
- [When to Enable IP Directed Broadcast on page 2850](#)
- [When Not to Enable IP Directed Broadcast on page 2850](#)

IP Directed Broadcast Overview

IP directed broadcast packets have a destination IP address that is a valid broadcast address for the subnet that is the target of the directed broadcast (the target subnet). The intent of an IP directed broadcast is to flood the target subnet with the broadcast packets without broadcasting to the entire network. IP directed broadcast packets cannot originate from the target subnet.

When you send an IP directed broadcast packet, as it travels to the target subnet, the network forwards it in the same way as it forwards a unicast packet. When the packet reaches a switch that is directly connected to the target subnet, the switch checks to see whether IP directed broadcast is enabled on the interface that is directly connected to the target subnet:

- If IP directed broadcast is enabled on that interface, the switch broadcasts the packet on that subnet by rewriting the destination IP address as the configured broadcast IP address for the subnet. The switch converts the packet to a link-layer broadcast packet that every host on the network processes.
- If IP directed broadcast is disabled on the interface that is directly connected to the target subnet, the switch drops the packet.

IP Directed Broadcast Implementation

You configure IP directed broadcast on a per-subnet basis by enabling IP directed broadcast on the Layer 3 interface of the subnet's VLAN. When the switch that is connected to that subnet receives a packet that has the subnet's broadcast IP address as the destination address, the switch broadcasts the packet to all hosts on the subnet.

By default, IP directed broadcast is disabled.

When to Enable IP Directed Broadcast

IP directed broadcast is disabled by default. Enable IP directed broadcast when you want to perform remote management or administration services such as backups or WOL tasks on hosts in a subnet that does not have a direct connection to the Internet.

Enabling IP directed broadcast on a subnet affects only the hosts within that subnet. Only packets received on the subnet's Layer 3 interface that have the subnet's broadcast IP address as the destination address are flooded on the subnet.

When Not to Enable IP Directed Broadcast

Typically, you do not enable IP directed broadcast on subnets that have direct connections to the Internet. Disabling IP directed broadcast on a subnet's Layer 3 interface affects only that subnet. If you disable IP directed broadcast on a subnet and a packet that has the broadcast IP address of that subnet arrives at the switch, the switch drops the broadcast packet.

If a subnet has a direct connection to the Internet, enabling IP directed broadcast on it increases the network's susceptibility to denial-of-service (DoS) attacks.

For example, a malicious attacker can spoof a source IP address (use a source IP address that is not the actual source of the transmission to deceive a network into identifying the attacker as a legitimate source) and send IP directed broadcasts containing Internet Control Message Protocol (ICMP) echo (ping) packets. When the hosts on the network with IP directed broadcast enabled receive the ICMP echo packets, they all send replies to the victim that has the spoofed source IP address. This creates a flood of ping replies in a DoS attack that can overwhelm the spoofed source address; this is known as a *smurf* attack. Another common DoS attack on exposed networks with IP directed broadcast enabled is a *fraggle* attack, which is similar to a smurf attack except that the malicious packet is a User Datagram Protocol (UDP) echo packet instead of an ICMP echo packet.

Related Documentation

- [Configuring IP Directed Broadcast \(CLI Procedure\) on page 2851](#)
- [Example: Configuring IP Directed Broadcast on page 2852](#)

Configuring IP Directed Broadcast (CLI Procedure)

You can use IP directed broadcast on a switch to facilitate remote network management by sending broadcast packets to hosts on a specified subnet without broadcasting to the entire network. IP directed broadcast packets are broadcast on only the target subnet. The rest of the network treats IP directed broadcast packets as unicast packets and forwards them accordingly.

Before you begin to configure IP directed broadcast:

- Ensure that the subnet on which you want broadcast packets using IP direct broadcast is not directly connected to the Internet.
- Configure an integrated routing and bridging (IRB) interface or routed VLAN interface (RVI) for the subnet that will be enabled for IP direct broadcast. See *Configuring Integrated Routing and Bridging Interfaces (CLI Procedure)*, *Configuring Routed VLAN Interfaces (CLI Procedure)*, or [“Configuring VLANs” on page 2120](#).



NOTE: We recommend that you do not enable IP directed broadcast on subnets that have a direct connection to the Internet because of increased exposure to denial-of-service (DoS) attacks.

To enable IP directed broadcast for a specified subnet:



NOTE: In a mixed Virtual Chassis, when you configure targeted broadcast, you can only configure one interface. Otherwise, targeted broadcast will not work.

1. Add the target subnet's logical interfaces to the VLAN:

```
[edit interfaces]
user@switch# set ge-0/0/0.0 family ethernet-switching vlan members v1
user@switch# set ge-0/0/1.0 family ethernet-switching vlan members v1
```

2. Configure the Layer 3 interface on the VLAN that is the target of the IP directed broadcast packets:

```
[edit interfaces]
user@switch# set irb.1 family inet address 10.1.2.1/24
```

3. Associate a Layer 3 interface with the VLAN:

```
[edit vlans]
user@switch# set v1 l3-interface irb.1
```

4. Enable the Layer 3 interface for the VLAN to receive IP directed broadcasts:

```
[edit interfaces]
user@switch# set irb.1 family inet targeted-broadcast
```

Related Documentation

- [Example: Configuring IP Directed Broadcast on page 2852](#)
- [Understanding IP Directed Broadcast on page 2849](#)

Example: Configuring IP Directed Broadcast

IP directed broadcast provides a method of sending broadcast packets to hosts on a specified subnet without broadcasting those packets to hosts on the entire network.

This example shows how to enable a subnet to receive IP directed broadcast packets so you can perform backups and other network management tasks remotely:

- [Requirements on page 2852](#)
- [Overview and Topology on page 2852](#)
- [Configuration on page 2853](#)

Requirements

This example uses the following software and hardware components:

- Junos OS Release 15.1 or later for QFX Series switches
- One PC
- One QFX Series switch

Before you configure IP directed broadcast for a subnet:

- Ensure that the subnet does not have a direct connection to the Internet.
- Configure routed VLAN interfaces (RVIs) for the ingress and egress VLANs on the switch. See *Configuring Routed VLAN Interfaces (CLI Procedure)* or [“Configuring VLANs” on page 2120](#).

Overview and Topology

You might want to perform remote administration tasks such as backups and wake-on-LAN (WOL) application tasks to manage groups of clients on a subnet. One way to do this is to send IP directed broadcast packets targeted at the hosts in a particular target subnet.

The network forwards IP directed broadcast packets as if they were unicast packets. When the IP directed broadcast packet is received by a VLAN that is enabled for **targeted-broadcast**, the switch broadcasts the packet to all the hosts in its subnet.

In this topology (see [Figure 45](#)), a host is connected to an interface on an EX Series switch to manage the clients in subnet **10.1.2.1/24**. When the switch receives a packet with the broadcast IP address of the target subnet as its destination address, it forwards the packet to the subnet's Layer 3 interface and broadcasts it to all the hosts within the subnet.

Figure 45: Topology for IP Directed Broadcast




Table 260 shows the settings of the components in this example.

Table 260: Components of the IP Directed Broadcast Topology

Property	Settings
Switch hardware	QFX Series switch
Ingress VLAN name	v0
Ingress VLAN IP address	10.1.1.1/24
Egress VLAN name	v1
Egress VLAN IP address	10.1.2.1/24
Interfaces in VLAN v0	ge-0/0/3.0
Interfaces in VLAN v1	ge-0/0/0.0 and ge-0/0/1.0

Configuration

To configure IP directed broadcast on a subnet to enable remote management of its hosts:

CLI Quick Configuration

To quickly configure the switch to accept IP directed broadcasts targeted at subnet 10.1.2.1/24, copy the following commands and paste them into the switch's terminal window:

```
[edit]
set interfaces ge-0/0/0.0 family ethernet-switching vlan members v1
set interfaces ge-0/0/1.0 family ethernet-switching vlan members v1
set interfaces vlan.1 family inet address 10.1.2.1/24
set interfaces ge-0/0/3.0 family ethernet-switching vlan members v0
set interfaces vlan.0 family inet address 10.1.1.1/24
set vlans v1 l3-interface vlan.1
set vlans v0 l3-interface vlan.0
set interfaces vlan.1 family inet targeted-broadcast
```

Step-by-Step Procedure

To configure the switch to accept IP directed broadcasts targeted at subnet 10.1.2.1/24:

1. Add logical interface **ge-0/0/0.0** to VLAN **v1**:

```
[edit interfaces]
user@switch# set ge-0/0/0.0 family ethernet-switching vlan members v1
```
2. Add logical interface **ge-0/0/1.0** to VLAN **v1**:

```
[edit interfaces]
```

- ```
user@switch# set ge-0/0/1.0 family ethernet-switching vlan members v1
```
3. Configure the IP address for the egress VLAN, v1:  

```
[edit interfaces]
user@switch# set vlan.1 family inet address 10.1.2.1/24
```
  4. Add logical interface **ge-0/0/3.0** to VLAN v0:  

```
[edit interfaces]
user@switch# set ge-0/0/3.0 family ethernet-switching vlan members v0
```
  5. Configure the IP address for the ingress VLAN:  

```
[edit interfaces]
user@switch# set vlan.0 family inet address 10.1.1.1/24
```
  6. To route traffic between the ingress and egress VLANs, associate a Layer 3 interface with each VLAN:  

```
[edit vlans]
user@switch# set v1 l3-interfacevlan.1
user@switch# set v0 l3-interface vlan.0
```
  7. Enable the Layer 3 interface for the egress VLAN to receive IP directed broadcasts:  

```
[edit interfaces]
user@switch# set vlan.1 family inet targeted-broadcast
```

**Results** Check the results:

```
user@switch# show
interfaces {
 ge-0/0/0 {
 unit 0 {
 family ethernet-switching {
 vlan {
 members v1;
 }
 }
 }
 }
 ge-0/0/1 {
 unit 0 {
 family ethernet-switching {
 vlan {
 members v1;
 }
 }
 }
 }
 ge-0/0/3 {
 unit 0 {
 family ethernet-switching {
 vlan {
 members v0;
 }
 }
 }
 }
 vlan {
 unit 0 {
 family inet {
```

```
 targeted-broadcast;
 address 10.1.1/24;
 }
}
unit 1 {
 family inet {
 targeted-broadcast;
 address 10.1.2.1/24;
 }
}
}
vllans {
 default;
 v0 {
 l3-interface vllan.0;
 }
 v1 {
 l3-interface vllan.1;
 }
}
```

**Related Documentation**

- [Configuring IP Directed Broadcast \(CLI Procedure\) on page 2851](#)





## PART 51

# Layer 3 Logical Interfaces

- [Understanding Layer 3 Logical Interfaces on page 2859](#)



# Understanding Layer 3 Logical Interfaces

- [Understanding Layer 3 Logical Interfaces on page 2859](#)
- [Configuring a Layer 3 Logical Interface on page 2860](#)
- [Verifying That Layer 3 Logical Interfaces Are Working on page 2860](#)

## Understanding Layer 3 Logical Interfaces

---

A Layer 3 logical interface is a logical division of a physical interface that operates at the network level and therefore can receive and forward 802.1Q VLAN tags. You can use Layer 3 logical interfaces to route traffic among multiple VLANs along a single trunk line that connects a Juniper Networks switch to a Layer 2 switch. Only one physical connection is required between the switches.

To create Layer 3 logical interfaces on a switch, enable VLAN tagging, partition the physical interface into logical partitions, and bind the VLAN ID to the logical interface.

We recommend that you use the VLAN ID as the logical interface number when you configure the logical interface. QFX Series and EX4600 switches support a maximum of 4089 VLANs, which includes the default VLAN. You can, however, assign a VLAN ID in the range of 1 to 4094, but five of these VLAN IDs are reserved for internal use.

VLAN tagging places the VLAN ID in the frame header, allowing each physical interface to handle multiple VLANs. When you configure multiple VLANs on an interface, you must also enable tagging on that interface. Junos OS on switches supports a subset of the 802.1Q standard for receiving and forwarding routed or bridged Ethernet frames with single VLAN tags and running Virtual Router Redundancy Protocol (VRRP) over 802.1Q-tagged interfaces.

### Related Documentation

- [Interfaces Overview on page 2785](#)
- [Configuring a Layer 3 Logical Interface on page 2860](#)
- *Junos OS Network Interfaces Library for Routing Devices*

## Configuring a Layer 3 Logical Interface

Devices use Layer 3 logical interfaces to divide a physical interface into multiple logical interfaces, each corresponding to a VLAN. Layer 3 logical interfaces route traffic between subnets.

To configure Layer 3 logical interfaces, enable VLAN tagging and partition one or more physical ports into multiple logical interfaces, each corresponding to a VLAN ID.

Before you begin, make sure you set up your VLANs. See *Configuring VLANs*.

To configure Layer 3 logical interfaces:

1. Enable VLAN tagging:

```
[edit interfaces interface-name]
user@switch# set vlan-tagging
```

2. Bind each VLAN ID to a logical interface:

```
[edit interfaces interface-name]
user@switch# set unit logical-unit-number vlan-id vlan-id-number
```

### Related Documentation

- [Understanding Layer 3 Logical Interfaces on page 2859](#)
- [Verifying That Layer 3 Logical Interfaces Are Working on page 2860](#)

## Verifying That Layer 3 Logical Interfaces Are Working

**Purpose** After configuring Layer 3 logical interfaces, verify that they are set up properly and transmitting data.

- Action**
1. To determine if you have successfully created the logical interfaces and the links are up:

```
[edit interfaces]
user@switch> show interfaces interface-name terse
```

| Interface      | Admin | Link | Proto | Local      | Remote |
|----------------|-------|------|-------|------------|--------|
| ge-0/0/0       | up    | up   |       |            |        |
| ge-0/0/0.0     | up    | up   | inet  | 1.1.1.1/24 |        |
| ge-0/0/0.1     | up    | up   | inet  | 2.1.1.1/24 |        |
| ge-0/0/0.2     | up    | up   | inet  | 3.1.1.1/24 |        |
| ge-0/0/0.3     | up    | up   | inet  | 4.1.1.1/24 |        |
| ge-0/0/0.4     | up    | up   | inet  | 5.1.1.1/24 |        |
| ge-0/0/0.32767 | up    | up   |       |            |        |

2. Use the **ping** command from a device on one subnet to an address on another subnet to determine if packets were transmitted correctly on the logical interface VLANs:

```
user@switch> ping ip-address
PING 1.1.1.1 (1.1.1.1): 56 data bytes
64 bytes from 1.1.1.1: icmp_seq=0 ttl=64 time=0.157 ms
64 bytes from 1.1.1.1: icmp_seq=1 ttl=64 time=0.238 ms
64 bytes from 1.1.1.1: icmp_seq=2 ttl=64 time=0.255 ms
64 bytes from 1.1.1.1: icmp_seq=3 ttl=64 time=0.128 ms
--- 1.1.1.1 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
```

**Meaning** The output confirms that the logical interfaces have been created and the links are up.

**Related Documentation**

- [Configuring a Layer 3 Logical Interface on page 2860](#)



## PART 52

# Link Aggregation Groups (LAGs) and Link Aggregation Control Protocol (LACP)

- [Understanding LAGs and LACP on page 2865](#)





# Understanding LAGs and LACP

- Understanding Aggregated Ethernet Interfaces and LACP on page 2865
- Configuring Aggregated Ethernet LACP on page 2868
- Configuring Link Aggregation on page 2869
- Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 2872
- Configuring Periodic Rebalancing of Subscribers in an Aggregated Ethernet Interface on page 2875
- Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2876
- Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch on page 2880
- Verifying the Status of a LAG Interface on page 2885
- Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets on page 2885
- Troubleshooting an Aggregated Ethernet Interface on page 2886

## Understanding Aggregated Ethernet Interfaces and LACP

---

IEEE 802.3ad link aggregation enables you to group Ethernet interfaces to form a single, aggregated Ethernet interface, also known as a *link aggregation group (LAG)* or *bundle*.

Link aggregation is used to aggregate Ethernet interfaces between two devices. You can create a LAG between a Juniper Networks device and a router, switch, aggregation switch, server, or other devices. The aggregated Ethernet interfaces that participate in a LAG are called member links. Because a LAG is composed of multiple member links, even if one member link fails, the LAG continues to carry traffic over the remaining links.



**NOTE:** On QFX5100 and EX4600 standalone switches and on a QFX5100 Virtual Chassis and EX4600 Virtual Chassis, you can configure a mixed rate of link speeds for the aggregated Ethernet bundle. Only link speeds of 40G and 10G are supported. Load balancing will not work if you configure link speeds that are not supported.



**NOTE:** The QFX5200 switches do not support mixed rate aggregated Ethernet bundles.

Link Aggregation Control Protocol (LACP) is a subcomponent of the IEEE 802.3ad standard and is used as a discovery protocol.



**NOTE:** To ensure load balancing across the aggregated Ethernet (AE) interfaces on a redundant server Node group, the members of the AE must be equally distributed across the redundant server Node group.



**NOTE:** During a network Node group switchover, traffic might be dropped for a few seconds.

- [Link Aggregation Group on page 2866](#)
- [Link Aggregation Control Protocol \(LACP\) on page 2867](#)

## Link Aggregation Group

To create a LAG:

1. Create a logical aggregated Ethernet interface.
2. Define the parameters associated with the logical aggregated Ethernet interface, such as a logical unit, interface properties, and Link Aggregation Control Protocol (LACP).
3. Define the member links to be contained within the aggregated Ethernet interface—for example, two 10-Gigabit Ethernet interfaces.
4. Configure LACP for link detection.

Keep in mind these hardware and software guidelines:

- Up to 32 Ethernet interfaces can be grouped to form a LAG on a redundant server Node group, a server Node group, and a network Node group on a QFabric system. Up to 48 LAGs are supported on redundant server Node groups and server Node groups on a QFabric system, and up to 128 LAGs are supported on network Node groups on a QFabric system. You can configure LAGs across Node devices in redundant server Node groups, server Node groups, and network Node groups.



**NOTE:** If you try to commit a configuration containing more than 32 Ethernet interfaces in a LAG, you will receive an error message saying that the group limit of 32 has been exceeded, and the configuration checkout has failed.

- Up to 64 Ethernet interfaces can be grouped to form a LAG, and up to 448 LAGs are supported on QFX3500, QFX3600, EX4600, and OCX Series switches, and up to 1,000 LAGs are supported on QFX5100 switches.



**NOTE:** If you try to commit a configuration containing more than 64 Ethernet interfaces in a LAG, you will receive an error message saying that the group limit of 64 has been exceeded, and the configuration checkout has failed.

- Up to 64 Ethernet interfaces can be grouped to form a LAG, and up to 144 LAGs are supported on QFX10002-36Q switches, and up to 288 LAGs are supported on QFX10002-72Q switches.
- The LAG must be configured on both sides of the link.
- The interfaces on either side of the link must be set to the same speed and be in full-duplex mode.



**NOTE:** On a QFX5100, EX4600, QFX10002 standalone switch or QFX5100 Virtual Chassis and EX4600 Virtual Chassis, you can configure mixed rate aggregated Ethernet bundles (LAGs with different link speeds). OCX Series switches do not support LAGs with different speeds.



**NOTE:** Junos OS assigns a unique ID and port priority to each port. The ID and priority are not configurable.

- QFabric systems support a special LAG called an FCoE LAG, which enables you to transport FCoE traffic and regular Ethernet traffic (traffic that is not FCoE traffic) across the same link aggregation bundle. Standard LAGs use a hashing algorithm to determine which physical link in the LAG is used for a transmission, so communication between two devices might use different physical links in the LAG for different transmissions. An FCoE LAG ensures that FCoE traffic uses the same physical link in the LAG for requests and replies in order to preserve the virtual point-to-point link between the FCoE device converged network adapter (CNA) and the FC SAN switch across a QFabric system Node device. An FCoE LAG does not provide load balancing or link redundancy for FCoE traffic. However, regular Ethernet traffic uses the standard hashing algorithm and receives the usual LAG benefits of load balancing and link redundancy in an FCoE LAG. See *Understanding FCoE LAGs* for more information.

## Link Aggregation Control Protocol (LACP)

LACP is one method of bundling several physical interfaces to form one logical aggregated Ethernet interface. The LACP mode can be active or passive. The transmitting link is known as the *actor*, and the receiving link is known as the *partner*. If the actor and partner are both in passive mode, they do not exchange LACP packets, and the aggregated Ethernet links do not come up. If either the actor or partner is active, they do exchange LACP packets. By default, LACP is in passive mode on aggregated Ethernet interfaces. To initiate transmission of LACP packets and response to LACP packets, you must enable LACP active mode. You can configure Ethernet links to actively transmit protocol data units (PDUs), or you can configure the links to passively transmit them, sending out LACP

PDUs only when they receive them from another link. You can configure both VLAN-tagged and untagged aggregated Ethernet interfaces without LACP enabled. LACP is defined in IEEE 802.3ad, *Aggregation of Multiple Link Segments*.

LACP was designed to achieve the following:

- Automatic addition and deletion of individual links to the LAG without user intervention.
- Link monitoring to check whether both ends of the bundle are connected to the correct group.

When a dual-homed server is deployed with a switch, the network interface cards form a LAG with the switch. During a server upgrade, the server may not be able to exchange LACP PDUs. In such a situation you can configure an interface to be in the **up** state even if no PDUs are exchanged. Use the **force-up** statement to configure an interface when the peer has limited LACP capability. The interface selects the associated LAG by default, whether the switch and peer are both in active or passive mode. When there are no received PDUs, the partner is considered to be working in the passive mode. Therefore, LACP PDU transmissions are controlled by the transmitting link.

If the remote end of the LAG link is a security device, LACP might not be supported because security devices require a deterministic configuration. In this case, do not configure LACP. All links in the LAG are permanently operational unless the switch detects a link failure within the Ethernet physical layer or data link layers.

**Related  
Documentation**

- [Configuring Link Aggregation on page 2869](#)
- [Configuring an FCoE LAG](#)
- [Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2876](#)
- [Example: Configuring an FCoE LAG on a Redundant Server Node Group](#)
- [Verifying the Status of a LAG Interface on page 2885](#)
- [Junos OS Network Interfaces Library for Routing Devices](#)

---

## Configuring Aggregated Ethernet LACP

---

For aggregated Ethernet interfaces, you can configure the Link Aggregation Control Protocol (LACP). LACP is one method of bundling several physical interfaces to form one logical interface. You can configure aggregated Ethernet with or without LACP enabled.

Before you configure LACP, be sure you have configured the aggregated Ethernet bundles—also known as link aggregation groups (LAGs).

When LACP is enabled, the local and remote sides of the aggregated Ethernet links exchange protocol data units (PDUs), containing information about the state of the link. You can configure Ethernet links to actively transmit PDUs, or you can configure the links to passively transmit them, sending out LACP PDUs only when they receive them from another link. One side of the link must be configured as **active** for the link to be up.



**NOTE:** Do not add LACP to a LAG if the remote end of the LAG link is a security device, unless the security device supports LACP. Security devices often do not support LACP because they require a deterministic configuration.

To configure LACP:

1. Enable the LACP mode:

```
[edit interfaces]
user@switch# set aex aggregated-ether-options lacp mode
```

For example, to specify the mode as active, execute the following command:

```
[edit interfaces]
user@switch# set aex aggregated-ether-options lacp active
```

2. Specify the interval and speed at which the interfaces send LACP packets:

```
[edit interfaces]
user@switch# set aex aggregated-ether-options lacp periodic interval
```

For example, to specify the interval as fast, execute the following command:

```
[edit interfaces]
user@switch# set aex aggregated-ether-options lacp periodic fast
```

## Configuring Link Aggregation

Use the link aggregation feature to aggregate one or more links to form a virtual link or aggregation group. The MAC client can treat this virtual link as if it were a single link. Link aggregation increases bandwidth, provides graceful degradation as failure occurs, and increases link availability.



**NOTE:** An interface with an already configured IP address cannot form part of the aggregation group.



**NOTE:** On QFX5100, QFX5200, EX4600, QFX10002, and QFX10008 standalone switches and on QFX5100 Virtual Chassis and EX4600 Virtual Chassis, you can configure a mixed rate of link speeds for the aggregated Ethernet bundle. Load balancing will not work if you configure link speeds that are not supported.

1. [Creating an Aggregated Ethernet Interface on page 2870](#)
2. [Configuring the VLAN Name and VLAN ID Number on page 2870](#)
3. [Configuring Aggregated Ethernet LACP on page 2871](#)

## Creating an Aggregated Ethernet Interface

To create an aggregated Ethernet interface:

1. Specify the number of aggregated Ethernet interfaces to be created:

```
[edit chassis]
user@switch# set aggregated-devices interfaces device-count device-count
```

For example, to specify 5:

```
[edit chassis]
user@switch# set aggregated-devices interfaces device-count
```

2. Specify the minimum number of links for the aggregated Ethernet interface (aex), that is, the defined bundle, to be labeled “up”:



**NOTE:** By default only one link must be up for the bundle to be labeled “up”.

```
[edit interfaces]
user@switch# set interface-name aggregated-ether-options minimum-links minimum-links
```

For example, to specify 5:

```
[edit interfaces]
user@switch# set interface-name aggregated-ether-options minimum-links 5
```

3. Specify the link speed for the aggregated Ethernet bundle:

```
[edit interfaces]
user@switch# set interface-name aggregated-ether-options link-speed link-speed
```

For example, to specify 10g:

```
[edit interfaces]
user@switch# set interface-name aggregated-ether-options link-speed 10g
```

4. Specify the members to be included within the aggregated Ethernet bundle:

```
[edit interfaces]
user@switch# set interface-name ether-options 802.3ad aex
user@switch# set interface-name ether-options 802.3ad aex
```

## Configuring the VLAN Name and VLAN ID Number



**NOTE:** VLANs are not supported on OCX Series switches.

```
[edit vlans]
user@switch# set vlan-name vlan-id vlan-id-number
```

For example, 100.

## Configuring Aggregated Ethernet LACP

For aggregated Ethernet interfaces, you can configure the Link Aggregation Control Protocol (LACP). LACP is one method of bundling several physical interfaces to form one logical interface. You can configure aggregated Ethernet with or without LACP enabled.

Before you configure LACP, be sure you have configured the aggregated Ethernet bundles—also known as link aggregation groups (LAGs).

When LACP is enabled, the local and remote sides of the aggregated Ethernet links exchange protocol data units (PDUs), containing information about the state of the link. You can configure Ethernet links to actively transmit PDUs, or you can configure the links to passively transmit them, sending out LACP PDUs only when they receive them from another link. One side of the link must be configured as **active** for the link to be up.



**NOTE:** Do not add LACP to a LAG if the remote end of the LAG link is a security device, unless the security device supports LACP. Security devices often do not support LACP because they require a deterministic configuration.

To configure LACP:

1. Enable the LACP mode:

```
[edit interfaces]
user@switch# set aex aggregated-ether-options lacp mode
```

For example, to specify the mode as active, execute the following command:

```
[edit interfaces]
user@switch# set aex aggregated-ether-options lacp active
```

2. Specify the interval and speed at which the interfaces send LACP packets:

```
[edit interfaces]
user@switch# set aex aggregated-ether-options lacp periodic interval
```

For example, to specify the interval as fast, execute the following command:

```
[edit interfaces]
user@switch# set aex aggregated-ether-options lacp periodic fast
```

### Related Documentation

- [Understanding Interface Naming Conventions on page 2787](#)
- [Configuring an FCoE LAG](#)
- [Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2876](#)
- [Verifying the Status of a LAG Interface on page 2885](#)
- [Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets on page 2885](#)
- [show lacp statistics interfaces \(View\) on page 3209](#)

## Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure)

You can configure LACP link protection and system priority at the global level on QFX10000 switches for a specific aggregated Ethernet interface. When using LACP link protection to protect a single link in the aggregated ethernet bundle, you configure only two member links for an aggregated Ethernet interface: one active and one standby. LACP link protection ensures that only one link—the link with the higher priority—is used for traffic. The other link is forced to stay in a *waiting* state.

When using LACP link protection to protect multiple links in an aggregated ethernet bundle, you configure links into primary and backup subgroups. A link protection subgroup is a collection of ethernet links within the aggregated ethernet bundle. When you use link protection subgroups, you configure a primary subgroup and a backup subgroup. The configuration process includes assigning member links to each subgroup. When the configuration process is complete, the primary subgroup is used to forward traffic until a switchover event, such as a link failure, occurs and causes the backup subgroup to assume control of traffic that was travelling on the links in the primary subgroup within the bundle.

By default LACP link protection reverts to a higher-priority (lower-numbered) link when the higher-priority link becomes operational or when a higher-priority link is added to the aggregated Ethernet bundle. For priority purposes, LACP link protection treats subgroups like links. You can suppress link calculation by adding the **non-revertive** statement to the link protection configuration. In nonrevertive mode, when a link is active in sending and receiving LACP packets, adding a higher-priority link to the bundle does not change the status of the currently active link. It remains active.

If LACP link configuration is specified to be nonrevertive at the global **[edit chassis]** hierarchy level, you can specify the **revertive** statement in the LACP link protection configuration at the aggregated Ethernet interface level to override the nonrevertive setting for the interface. In revertive mode, adding a higher-priority link to the aggregated Ethernet bundle results in LACP recalculating the priority and switching the status from the currently active link to the newly added, higher-priority link.



**NOTE:** When LACP link protection is enabled on both local and remote sides of the link, both sides must use the same mode (either revertive or nonrevertive).

---

Configuring LACP link configuration at the aggregated Ethernet level results in only the configured interfaces using the defined configuration. LACP interface configuration also enables you to override global (chassis) LACP settings.

Before you configure LACP link protection, be sure you have:

- Configured the aggregated Ethernet bundles—also known as link aggregation groups (LAGs). See *Configuring Link Aggregation*.



- Configured LACP for the interface. See [“Configuring Aggregated Ethernet LACP” on page 2868](#).

You can configure LACP link protection for all aggregated Ethernet interfaces on the switch by enabling it at the global level on the switch or configure it for a specific aggregated Ethernet interface by enabling it on that interface.

- [Configuring LACP Link Protection for a Single Link at the Global Level on page 2873](#)
- [Configuring LACP Link Protection for a Single Link at the Aggregated Interface Level on page 2873](#)
- [Configuring Subgroup Bundles to Provide LACP Link Protection to Multiple Links in an Aggregated Ethernet Interface on page 2874](#)

## Configuring LACP Link Protection for a Single Link at the Global Level

To configure LACP link protection for aggregated Ethernet interfaces at the global level:

1. Enable LACP link protection on the switch:  

```
[edit chassis aggregated-devices ethernet lacp]
user@switch# set link-protection
```
2. (Optional) Configure the LACP link protection for the aggregated Ethernet interfaces to be in nonrevertive mode:



**NOTE:** LACP link protection is in revertive mode by default.

- ```
[edit chassis aggregated-devices ethernet lacp link-protection]
user@switch# set non-revertive
```
3. (Optional) To configure LACP system priority for the aggregated Ethernet interfaces:

```
[edit chassis aggregated-devices ethernet lacp]
user@switch# set system-priority
```

Configuring LACP Link Protection for a Single Link at the Aggregated Interface Level

To enable LACP link protection for a specific aggregated Ethernet interface:

1. Enable LACP link protection for the interface:

```
[edit interfaces aeX aggregated-ether-options lacp]
user@switch# set link-protection
```
2. (Optional) Configure the LACP link protection for the aggregated Ethernet interface to be in revertive or nonrevertive mode:
 - To specify revertive mode:

```
[edit interfaces aeX aggregated-ether-options lacp link-protection]
user@switch# set revertive
```
 - To specify nonrevertive mode:

```
[edit interfaces aeX aggregated-ether-options lacp link-protection]
user@switch# set non-revertive
```
3. (Optional) To configure LACP system priority for an aggregated Ethernet interface:

```
[edit interfaces aeX aggregated-ether-options lacp link-protection]
```

```
user@switch# set system-priority
```

4. (Optional) To configure LACP port priority for an aggregated Ethernet interface:

```
[edit interfaces ge-fpc/pic/port ether-options 802.3ad lacp]
user@switch# set port-priority
```

Configuring Subgroup Bundles to Provide LACP Link Protection to Multiple Links in an Aggregated Ethernet Interface

You can configure link protection subgroup bundles to provide link protection for multiple links in an aggregated ethernet bundle.

Link protection subgroups allow you to provide link protection to a collection of Ethernet links within a LAG bundle, instead of providing protection to a single link in the aggregated ethernet bundle only. You can, for instance, configure a primary subgroup with three member links and a backup subgroup with three different member links and use the backup subgroup to provide link protection for the primary subgroup.

To configure link protection using subgroups:

1. Configure the primary link protection subgroup in the aggregated ethernet interface:

```
[edit interfaces aeX aggregated-ether-options]
user@switch# set link-protection-sub-group group-name primary
```

For instance, to create a primary link protection subgroup named **subgroup-primary** for interface **ae0**:

```
[edit interfaces ae0 aggregated-ether-options]
user@switch# set link-protection-sub-group subgroup-primary primary
```

2. Configure the backup link protection subgroup in the aggregated ethernet interface:

```
[edit interfaces aeX aggregated-ether-options]
user@switch# set link-protection-sub-group group-name backup
```

For instance, to create a backup link protection subgroup named **subgroup-backup** for interface **ae0**:

```
[edit interfaces ae0 aggregated-ether-options]
user@switch# set link-protection-sub-group subgroup-backup backup
```



NOTE: You can create one primary and one backup link protection subgroup per aggregated ethernet interface.

3. Attach interfaces to the link protection subgroups:

```
[edit interfaces interface-name ether-options 802.3ad]
user@switch# set link-protection-sub-group group-name
```



NOTE: The primary and backup link protection subgroups must contain the same number of interfaces. For instance, if the primary link protection subgroup contains three interfaces, the backup link protection subgroup must also contain three interfaces.

For instance, to configure interfaces **ge-0/0/0** and **ge-0/0/1** into link protection subgroup **subgroup-primary** and interfaces **ge-0/0/2** and **ge-0/0/3** into link protection subgroup **subgroup-backup**:

```
[edit interfaces ge-0/0/0 ether-options 802.3ad]
user@switch# set link-protection-sub-group subgroup-primary
[edit interfaces ge-0/0/1 ether-options 802.3ad]
user@switch# set link-protection-sub-group subgroup-primary
[edit interfaces ge-0/0/2 ether-options 802.3ad]
user@switch# set link-protection-sub-group subgroup-backup
[edit interfaces ge-0/0/3 ether-options 802.3ad]
user@switch# set link-protection-sub-group subgroup-backup
```

4. (Optional) Configure the port priority for link protection:

```
[edit interfaces interface-name ether-options 802.3ad]
user@switch# set port-priority priority
```

The port priority is used to select the active link.

5. Enable link protection

To enable link protection at the LAG level:

```
[edit interfaces aeX aggregated-ether-options]
user@switch# set link-protection
```

To enable link protection at the LACP level:

```
[edit interfaces aeX aggregated-ether-options lacp]
user@switch# set link-protection
```

For instance, to enable link protection on **ae0** at the LAG level:

```
[edit interfaces ae0 aggregated-ether-options]
user@switch# set link-protection
```

For instance, to enable link protection on **ae0** at the LACP level:

```
[edit interfaces ae0 aggregated-ether-options lacp]
user@switch# set link-protection
```

Related Documentation

- [lACP \(Aggregated Ethernet\) on page 3050](#)

Configuring Periodic Rebalancing of Subscribers in an Aggregated Ethernet Interface

If subscribers are frequently logging in and logging out of your network, you can configure the system to periodically rebalance the links based on a specific time and interval.

To configure periodic rebalancing:

1. Access the aggregated Ethernet interface for which you want to configure periodic rebalancing.

```
edit
user@host# edit interfaces aenumber aggregated-ether-options
```

2. Configure the rebalancing parameters for the interface, including the time and the interval between rebalancing actions.

```
[edit interfaces aenumber aggregated-ether-options]
user@host# rebalance-periodic time hour:minute <interval hours>
```

- Related Documentation**
- *Verifying the Distribution of Demux Subscribers in an Aggregated Ethernet Interface*
 - *Configuring the Distribution Type for Demux Subscribers on Aggregated Ethernet Interfaces*
 - *Distribution of Demux Subscribers in an Aggregated Ethernet Interface*

Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch

A QFX Series product allows you to combine multiple Ethernet links into one logical interface for higher bandwidth and redundancy. The ports that are combined in this manner are referred to as a link aggregation group (LAG) or bundle. The number of Ethernet links you can combine into a LAG depends on your QFX Series product model. You can configure LAGs to connect a QFX Series product to other switches, like aggregation switches, servers, or routers. This example describes how to configure LAGs to connect a QFX3500, QFX3600, EX4600, QFX5100, and QFX10002 switch to an aggregation switch.

- [Requirements on page 2876](#)
- [Overview and Topology on page 2876](#)
- [Configuration on page 2877](#)
- [Verification on page 2879](#)
- [Troubleshooting on page 2880](#)

Requirements

This example uses the following software and hardware components:

- Junos OS Release 11.1 or later for the QFX3500 and QFX3600 switches, Junos OS 13.2 or later for the QFX5100 and EX4600 switch, and Junos OS Release 15.1X53-D10 for QFX10002 switches.
- One QFX3500, QFX3600, EX4600, QFX5100, QFX10002 switch.

Overview and Topology

In this example, the switch has one LAG comprising two 10-Gigabit Ethernet interfaces. This LAG is configured in port mode trunk so that the switch and the VLAN to which it has been assigned can send and receive traffic.

Configuring the Ethernet interfaces as LAGs has the following advantages:

- If one physical port is lost for any reason (a cable is unplugged or a switch port fails), the logical port transparently continues to function over the remaining physical port.
- Link Aggregation Control Protocol (LACP) can optionally be configured for link monitoring and automatic addition and deletion of individual links without user intervention.



NOTE: If the remote end of the LAG link is a security device, LACP might not be supported because security devices require a deterministic configuration. In this case, do not configure LACP. All links in the LAG are permanently operational unless the switch detects a link failure within the Ethernet physical layer or data link layers.

The topology used in this example consists of one switch with a LAG configured between two of its 10-Gigabit Ethernet interfaces. The switch is connected to an aggregation switch.




Table 261 details the topology used in this configuration example.

Table 261: Components of the Topology for Configuring a LAG Between a QFX3500 Switch and Aggregation Switch

Hostname	Base Hardware	Trunk Port
switch	QFX3500, QFX3600, EX4600, QFX5100, or QFX10002 switch	ae0 is configured as a trunk port and combines the following two interfaces: xe-0/0/2 and xe-0/0/3

Configuration

To configure a LAG between two 10-Gigabit Ethernet interfaces:

CLI Quick Configuration

To quickly configure a LAG between two 10-Gigabit Ethernet interfaces on a switch, copy the following commands and paste them into the switch terminal window:



NOTE: If you are configuring a LAG using Enhanced Layer 2 software, use the **interface-mode** statement instead of the **port-mode** statement. For ELS details, see [“Getting Started with Enhanced Layer 2 Software”](#) on page 41.

```

[edit]
set chassis aggregated-devices ethernet device-count 1
set interfaces ae0 aggregated-ether-options minimum-links 1
set interfaces ae0 aggregated-ether-options link-speed 10g
set interfaces ae0 unit 0 family ethernet-switching vlan members green
set interfaces xe-0/0/2 ether-options 802.ad ae0
set interfaces xe-0/0/3 ether-options 802.ad ae0
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp periodic fast
  
```

Step-by-Step Procedure

To configure a LAG between a QFX Series switch and an aggregation switch:

1. Specify the number of LAGs to be created on the switch:

```
[edit chassis]
user@switch# set aggregated-devices ethernet device-count 1
```
2. Specify the number of links that need to be present for the ae0 LAG interface to be up:

```
[edit interfaces]
user@switch# set ae0 aggregated-ether-options minimum-links 1
```
3. Specify the media speed of the ae0 link:

```
[edit interfaces]
user@switch# set ae0 aggregated-ether-options link-speed 10g
```
4. Specify the members to be included within the aggregated Ethernet bundle:

```
[edit interfaces]
user@switch# set interfaces xe-0/0/2 ether-options 802.ad ae0
[edit interfaces]
user@switch# set interfaces xe-0/0/3 ether-options 802.ad ae0
```
5. Assign a port mode of trunk to the ae0 link:



NOTE: If you are configuring a LAG on the QFX5100 switch, use the **interface-mode** statement instead of the **port-mode** statement. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 41](#).

- ```
[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching port-mode trunk
or
[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching interface-mode trunk
```
6. Assign the LAG to a VLAN:  

```
[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching vlan members green vlan-id 200
```
  7. (Optional): Designate one side of the LAG as active for LACP:  

```
[edit interfaces]
user@switch# set ae0 aggregated-ether-options lacp active
```
  8. (Optional): Designate the interval and speed at which the interfaces send LACP packets:  

```
[edit interfaces]
user@switch# set ae0 aggregated-ether-options lacp periodic fast
```

---

**Results**

Display the results of the configuration on a QFX3500 or QFX3600 switch:

```
[edit]
chassis {
 aggregated-devices {
 ethernet {
```

```

 device-count 1;
 }
}
green {
 vlan-id 200;
}
}
interfaces {
 ae0 {
 aggregated-ether-options {
 link-speed 10g;
 minimum-links 1;
 }
 unit 0 {
 family ethernet-switching {
 port-mode trunk;
 vlan {
 members green;
 }
 }
 }
 }
 xe-0/0/2 {
 ether-options {
 802.ad ae0;
 }
 }
 xe-0/0/3 {
 ether-options {
 802.ad ae0;
 }
 }
}
}

```

## Verification

To verify that switching is operational and one LAG has been created, perform these tasks:

- [Verifying That LAG ae0.0 Has Been Created on page 2879](#)
- [Verifying That LAG ae0 Has Been Created on page 2880](#)

### Verifying That LAG ae0.0 Has Been Created

**Purpose** Verify that LAG **ae0.0** has been created on the switch.

**Action** **show interfaces ae0 terse**

| Interface | Admin | Link | Proto      | Local | Remote |
|-----------|-------|------|------------|-------|--------|
| ae0       | up    | up   |            |       |        |
| ae0.0     | up    | up   | eth-switch |       |        |

**Meaning** The output confirms that the **ae0.0** link is up and shows the **family** and IP address assigned to this link.

### Verifying That LAG ae0 Has Been Created

---

**Purpose** Verify that LAG ae0 has been created on the switch

**Action** `show interfaces ae0 terse`

| Interface | Admin | Link | Proto      | Local | Remote |
|-----------|-------|------|------------|-------|--------|
| ae0       | up    | down |            |       |        |
| ae0.0     | up    | down | eth-switch |       |        |

**Meaning** The output shows that the ae0.0 link is down.

## Troubleshooting

### Troubleshooting a LAG That Is Down

---

**Problem** The `show interfaces terse` command shows that the LAG is **down**.

**Solution** Check the following:

- Verify that there is no configuration mismatch.
- Verify that all member ports are up.
- Verify that a LAG is part of family ethernet switching (Layer 2 LAG) or family inet (Layer 3 LAG).
- Verify that the LAG member is connected to the correct LAG at the other end.

- Related Documentation**
- [Configuring Link Aggregation on page 2869](#)
  - [Verifying the Status of a LAG Interface on page 2885](#)
  - [Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets on page 2885](#)
  - [Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch on page 2880](#)
  - [Example: Configuring an FCoE LAG on a Redundant Server Node Group](#)
  - [show lacp statistics interfaces \(View\) on page 3209](#)

## Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch

---

QFX Series products allow you to combine multiple Ethernet links into one logical interface for higher bandwidth and redundancy. The ports that are combined in this manner are referred to as a link aggregation group (LAG) or bundle. The number of Ethernet links you can combine into a LAG depends on your QFX Series product model. On a standalone switch, you can group up to 32 Ethernet interfaces to form a LAG. On a QFabric system, you can group up to 8 Ethernet interfaces to form a LAG. QFX Series products allow you to further enhance these links by configuring Link Aggregation Control Protocol (LACP).



This example describes how to overlay LACP on the LAG configurations that were created in “[Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch](#)” on page 2876:

- [Requirements on page 2881](#)
- [Overview and Topology on page 2881](#)
- [Configuring LACP for the LAG on the QFX Series on page 2881](#)
- [Verification on page 2882](#)
- [Troubleshooting on page 2883](#)

## Requirements

This example uses the following software and hardware components:

- Junos OS Release 11.1 or later for the QFX3500 switch, Junos OS Release 12.1 or later for the QFX3600 switch, Junos OS Release 13.2 or later for the QFX5100 switch, and Junos OS Release 15.1X53-D10 or later for the QFX10002 switch.
- One QFX3500, QFX3600, QFX5100, QFX10002 switch.

Before you configure LACP, be sure you have:

- Configured the ports on the switches as trunk ports.
- Configured the LAG.

## Overview and Topology

The topology in this example is exactly the same as the topology used in the [Configuring a LAG Between a QFX Switch and an Aggregation Switch](#) example. This example shows how to use LACP to enhance the LAG functionality.

LACP exchanges are made between *actors* (the transmitting link) and *partners* (the receiving link). The LACP mode can be either active or passive.



**NOTE:** If the actor and partner are both in passive mode, they do not exchange LACP packets, which results in the aggregated Ethernet links not coming up. By default, LACP is in passive mode. To initiate transmission of LACP packets and responses to LACP packets, you must enable LACP in active mode.

By default, the actor and partner send LACP packets every second. You can configure the interval at which the interfaces send LACP packets by including the **periodic** statement at the **[edit interfaces *interface-name* aggregated-ether-options lacp]** hierarchy level.

The interval can be fast (every second) or slow (every 30 seconds).

## Configuring LACP for the LAG on the QFX Series

To configure LACP for a QFX Series LAG, perform these tasks:

**CLI Quick Configuration** To quickly configure LACP for the access switch LAGs, copy the following commands and paste them into the switch terminal window:

```
[edit]
set interfaces ae0 aggregated-ether-options lacp active periodic fast
```

**Step-by-Step Procedure** To configure LACP for LAG ae0 :

1. Specify the aggregated Ethernet options for the LAG:

```
[edit interfaces]
user@switch# set ae0 aggregated-ether-options lacp active periodic fast
```

**Results** Display the results of the configuration:

```
[edit interfaces]
user@switch# show
ae0 {
 aggregated-ether-options {
 lacp {
 active;
 periodic fast;
 }
 }
}
```

## Verification

To verify that LACP packets are being exchanged, perform the following tasks:

- [Verifying the LACP Settings on page 2882](#)
- [Verifying That the LACP Packets Are Being Exchanged on page 2883](#)

### Verifying the LACP Settings

---

**Purpose** Verify that LACP has been set up correctly.

**Action** Use the **show lacp interfaces *interface-name*** command to check that LACP has been enabled as active on one end.

```
user@switch> show lacp interfaces xe-0/0/2
```

Aggregated interface: ae0

| LACP state:    | Role          | Exp | Def            | Dist | Co1       | Syn | Aggr | Timeout | Activity |
|----------------|---------------|-----|----------------|------|-----------|-----|------|---------|----------|
| xe-0/0/2       | Actor         | No  | Yes            | No   | No        | No  | Yes  | Fast    | Active   |
| xe-0/0/2       | Partner       | No  | Yes            | No   | No        | No  | Yes  | Fast    | Passive  |
| LACP protocol: | Receive State |     | Transmit State |      | Mux State |     |      |         |          |
| xe-0/0/2       | Defaulted     |     | Fast periodic  |      | Detached  |     |      |         |          |

**Meaning** The output indicates that LACP has been set up correctly and is active at one end.

### Verifying That the LACP Packets Are Being Exchanged

**Purpose** Verify that LACP packets are being exchanged.

**Action** Use the `show interfaces aex statistics` command to display LACP information.

```
user@switch> show interfaces ae0 statistics
```

```
Physical interface: ae0, Enabled, Physical link is Down
Interface index: 153, SNMP ifIndex: 30
Link-level type: Ethernet, MTU: 1514, Speed: Unspecified, Loopback: Disabled,
Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1,
Minimum bandwidth needed: 0
Device flags : Present Running
Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
Current address: 02:19:e2:50:45:e0, Hardware address: 02:19:e2:50:45:e0
Last flapped : Never
Statistics last cleared: Never
 Input packets : 0
 Output packets: 0
Input errors: 0, Output errors: 0

Logical interface ae0.0 (Index 71) (SNMP ifIndex 34)
Flags: Hardware-Down Device-Down SNMP-Traps Encapsulation: ENET2
Statistics Packets pps Bytes bps
Bundle:
 Input : 0 0 0 0
 Output: 0 0 0 0
Protocol inet
Flags: None
Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255
```

**Meaning** The output here shows that the link is down and that no PDUs are being exchanged.

## Troubleshooting

To troubleshoot a nonworking LACP link, perform these tasks:

- [Troubleshooting a Nonworking LACP Link on page 2883](#)

### Troubleshooting a Nonworking LACP Link

**Problem** The LACP link is not working.

**Solution** Check the following:

- Remove the LACP configuration and verify whether the static LAG is up.
- Verify that LACP is configured at both ends.
- Verify that LACP is not passive at both ends.
- Verify whether LACP protocol data units (PDUs) are being exchanged by running the `monitor traffic-interface lag-member detail` command.

**Related  
Documentation**

- [Configuring Link Aggregation on page 2869](#)
- [Verifying the Status of a LAG Interface on page 2885](#)
- [Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets on page 2885](#)
- [Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2876](#)
- [Example: Configuring an FCoE LAG on a Redundant Server Node Group](#)
- [show lacp statistics interfaces \(View\) on page 3209](#)

## Verifying the Status of a LAG Interface

**Purpose** Verify that a link aggregation group (LAG) (**ae0**) has been created on the switch.

**Action** To verify that the **ae0** LAG has been created:

```
[edit interfaces]
show interfaces ae0 terse
```

| Interface | Admin | Link | Proto | Local         | Remote |
|-----------|-------|------|-------|---------------|--------|
| ae0       | up    | up   |       |               |        |
| ae0.0     | up    | up   | inet  | 10.10.10.2/24 |        |

**Meaning** The output confirms that the **ae0** link is up and shows the family and IP address assigned to this link.

- Related Documentation**
- [Configuring Link Aggregation on page 2869](#)
  - [Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets on page 2885](#)
  - [Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2876](#)
  - [Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch on page 2880](#)
  - [show lacp statistics interfaces \(View\) on page 3209](#)

## Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets

Verify that LACP has been set up correctly and that the bundle members are transmitting LACP protocol packets.

1. [Verifying the LACP Setup on page 2885](#)
2. [Verifying That LACP Packets Are Being Exchanged on page 2886](#)

### Verifying the LACP Setup

**Purpose** Verify that the LACP has been set up correctly.

**Action** To verify that LACP has been enabled as active on one end:

```
user@switch>show lacp interfaces xe-0/0/0
```

```
Aggregated interface: ae0
```

| LACP state: | Role    | Exp | Def | Dist | Col | Syn | Aggr | Timeout | Activity |
|-------------|---------|-----|-----|------|-----|-----|------|---------|----------|
| xe-0/1/0    | Actor   | No  | Yes | No   | No  | No  | Yes  | Fast    | Active   |
| xe-0/1/0    | Partner | No  | Yes | No   | No  | No  | Yes  | Fast    | Passive  |

|                |               |                |          |       |
|----------------|---------------|----------------|----------|-------|
| LACP protocol: | Receive State | Transmit State | Mux      | State |
| xe-0/1/0       | Defaulted     | Fast periodic  | Detached |       |

**Meaning** This example shows that LACP has been configured with one side as active and the other as passive. When LACP is enabled, one side must be set as active in order for the bundled link to be up.

## Verifying That LACP Packets Are Being Exchanged

**Purpose** Verify that LACP packets are being exchanged between interfaces.

**Action** Use the `show lacp statistics interfaces interface-name` command to display LACP BPDU exchange information.

`show lacp statistics interfaces ae0`

Aggregated interface: ae0

|                  |         |         |            |            |
|------------------|---------|---------|------------|------------|
| LACP Statistics: | LACP Rx | LACP Tx | Unknown Rx | Illegal Rx |
| xe-0/0/2         | 1352    | 2035    | 0          | 0          |
| xe-0/0/3         | 1352    | 2056    | 0          | 0          |

**Meaning** The output here shows that the link is up and that PDUs are being exchanged.

- Related Documentation**
- [Configuring Link Aggregation on page 2869](#)
  - [Verifying the Status of a LAG Interface on page 2885](#)
  - [Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2876](#)
  - [Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch on page 2880](#)
  - [show lacp statistics interfaces \(View\) on page 3209](#)

## Troubleshooting an Aggregated Ethernet Interface

**Problem** **Description:** The `show interfaces terse` command shows that the LAG is down.

**Solution** Check the following:

- Verify that there is no configuration mismatch.
- Verify that all member ports are up.
- Verify that a LAG is part of family ethernet-switching (Layer 2 LAG) or family inet (Layer 3 LAG).



**NOTE:** Layer 2 LAGs are not supported on OCX Series switches.

- Verify that the LAG member is connected to the correct LAG at the other end.
- Verify that the LAG members belong to the same switch.

**Related  
Documentation**

- [Verifying the Status of a LAG Interface on page 2885](#)
- [Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2876](#)





## PART 53

# Load Balancing

- [Understanding Load Balancing on page 2891](#)



# Understanding Load Balancing

- [Configuring Load Balancing Based on MAC Addresses on page 2891](#)

## Configuring Load Balancing Based on MAC Addresses

---

The hash key mechanism for load-balancing uses Layer 2 media access control (MAC) information such as frame source and destination address. To load-balance traffic based on Layer 2 MAC information, include the **multiservice** statement at the **[edit forwarding-options hash-key]** or **[edit chassis fpc slot number pic PIC number hash-key]** hierarchy level:

```
multiservice {
 source-mac;
 destination-mac;
 payload {
 ip {
 layer3-only;
 layer-3 (source-ip-only | destination-ip-only);
 layer-4;
 inner-vlan-id;
 outer-vlan-id;
 }
 }
}
```

To include the destination-address MAC information in the hash key, include the **destination-mac** option. To include the source-address MAC information in the hash key, include the **source-mac** option.



**NOTE:** Any packets that have the same source and destination address will be sent over the same path.



**NOTE:** You can configure per-packet load balancing to optimize EVPN traffic flows across multiple paths.



NOTE: Aggregated Ethernet member links will now use the physical MAC address as the source MAC address in 802.3ah OAM packets.

Related Documentation

- [multiservice on page 3056](#)

## PART 54

# Local Link Bias

- [Understanding Local Link Bias on page 2895](#)



# Understanding Local Link Bias

- [Understanding Local Link Bias on page 2895](#)
- [Configuring Local Link Bias \(CLI Procedure\) on page 2897](#)

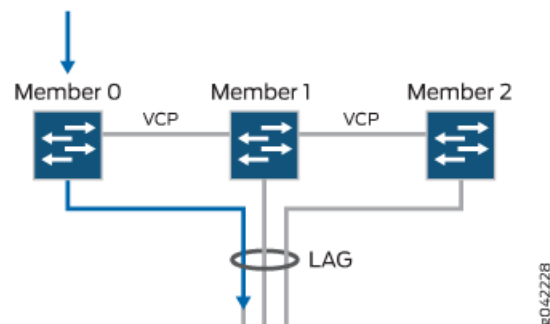
## Understanding Local Link Bias



**NOTE:** The QFX5200 switches do not support Virtual Chassis or Virtual Chassis ports.

Local link bias conserves bandwidth on Virtual Chassis ports (VCPs) by using local links to forward unicast traffic exiting a Virtual Chassis or Virtual Chassis Fabric (VCF) that has a Link Aggregation group (LAG) bundle composed of member links on different member switches in the same Virtual Chassis or VCF. A local link is a member link in the LAG bundle that is on the member switch that received the traffic. Because traffic is received and forwarded on the same member switch when local link bias is enabled, no VCP bandwidth is consumed by traffic traversing the VCPs to exit the Virtual Chassis or VCF using a different member link in the LAG bundle. The traffic flow of traffic exiting a Virtual Chassis or VCF over a LAG bundle when local link bias is enabled is illustrated in [Figure 46](#).


**Figure 46: Egress Traffic Flow with Local Link Bias**



When local link bias is disabled, egress traffic exiting a Virtual Chassis or VCF on a LAG bundle can be forwarded out of any member link in the LAG bundle. Traffic forwarding decisions are made by an internal algorithm that attempts to load-balance traffic between the member links in the bundle. VCP bandwidth is frequently consumed by egress traffic

when local link bias is disabled because the egress traffic traverses the VCPs to reach the destination egress member link in the LAG bundle. The traffic flow of traffic exiting a Virtual Chassis or VCF over a LAG bundle when local link bias is disabled is illustrated in Figure 47.

**Figure 47: Egress Traffic Flow without Local Link Bias**



Local link bias is configured in a LAG bundle. A Virtual Chassis or VCF that has multiple LAG bundles can contain bundles that have and have not enabled local link bias. Local link bias only impacts the forwarding of unicast traffic exiting a Virtual Chassis or VCF; ingress traffic handling is not impacted by the local link bias setting. Egress multicast, unknown unicast, and broadcast traffic exiting a Virtual Chassis or VCF over a LAG bundle is not impacted by the local link bias setting and is always load-balanced among the member links. Local link bias is disabled, by default.

You should enable local link bias if you want to conserve VCP bandwidth by always forwarding egress unicast traffic on a LAG bundle out of a local link. You should not enable local link bias if you want egress traffic load-balanced across the member links in the LAG bundle as it exits the Virtual Chassis or VCF.

#### Related Documentation

- [Configuring Local Link Bias \(CLI Procedure\) on page 2897](#)



---

## Configuring Local Link Bias (CLI Procedure)

---

Local link bias is used to conserve bandwidth on Virtual Chassis ports (VCPs) by using local links to forward unicast traffic exiting a Virtual Chassis or Virtual Chassis Fabric (VCF) that has a Link Aggregation group (LAG) bundle composed of member links on different member switches in the same Virtual Chassis or VCF. A local link is a member link in the LAG bundle that is on the member switch that received the traffic. Because traffic is received and forwarded on the same member switch when local link bias is enabled, no VCP bandwidth is consumed by traffic traversing the VCPs to exit the Virtual Chassis or VCF on a different member link in the LAG bundle.

You should enable local link bias if you want to conserve VCP bandwidth by always forwarding egress unicast traffic on a LAG out of a local link. You should not enable local link bias if you want egress traffic load-balanced as it exits the Virtual Chassis or VCF.

To enable local link bias on a LAG bundle:

```
[edit]
user@switch# set interface aex aggregated-ether-options local-bias
where aex is the name of the aggregated Ethernet link bundle.
```

For instance, to enable local link bias on aggregated Ethernet interface ae0:

```
[edit]
user@switch# set interface ae0 aggregated-ether-options local-bias
```

**Related Documentation**

- [Understanding Local Link Bias on page 2895](#)



## PART 55

# Redundant Trunk Groups

- [Understanding Redundant Trunk Groups on page 2901](#)



# Understanding Redundant Trunk Groups

- [Understanding Redundant Trunk Links on page 2901](#)
- [Example: Configuring Redundant Trunk Links for Faster Recovery on page 2903](#)

## Understanding Redundant Trunk Links

---

In a typical enterprise network composed of distribution and access layers, a redundant trunk link provides a simple solution for network recovery when a trunk port on a switch goes down. In that case, traffic is routed to another trunk port, keeping network convergence time to a minimum.

To configure a redundant trunk link, create a redundant trunk group. The redundant trunk group is configured on the access switch and contains two links: a primary or active link, and a secondary link. If the active link fails, the secondary link automatically starts forwarding data traffic without waiting for normal spanning-tree protocol convergence.

Data traffic is forwarded only on the active link. Data traffic on the secondary link is dropped and shown as dropped packets when you issue the operational mode command **show interfaces *interface-name* extensive**.

While data traffic is blocked on the secondary link, Layer 2 control traffic is still permitted. For example, an LLDP session can be run between two switches on the secondary link.

Rapid Spanning Tree Protocol (RSTP) is enabled by default on the switches to create a loop-free topology, but an interface is not allowed to be in both a redundant trunk group and in a spanning-tree protocol topology at the same time. You must disable RSTP on an interface if a redundant trunk group is configured on that interface. For example, in [Figure 48](#), in addition to disabling RSTP on the Switch 3 interfaces, you must also disable RSTP on the Switch 1 and Switch 2 interfaces connected to Switch 3. Spanning-tree protocols can, however, continue operating on other interfaces on those switches—for example on the link between Switch 1 and Switch 2.

[Figure 48](#) shows three switches in a basic topology for redundant trunk links. Switch 1 and Switch 2 make up the distribution layer, and Switch 3 makes up the access layer. Switch 3 is connected to the distribution layer through trunk ports ge-0/0/9.0 (Link 1) and ge-0/0/10.0 (Link 2). Link 1 and Link 2 are in a redundant trunk group called group1. Link 1 is designated as the primary link. Traffic flows between Switch 3 in the access layer and Switch 1 in the distribution layer through Link 1. While Link 1 is active, Link 2 blocks traffic.

Figure 48: Redundant Trunk Group, Link 1 Active





Figure 49 illustrates how the redundant trunk link topology works when the primary link goes down.

Figure 49: Redundant Trunk Group, Link 2 Active



When Link 1 between Switch 1 and Switch 3 goes down, Link 2 takes over as the active link. Traffic between the access layer and the distribution layer is then automatically switched to Link 2 between Switch 1 and Switch 2.

**Related  
Documentation**

- [Example: Configuring Redundant Trunk Links for Faster Recovery](#)
- [Example: Configuring Redundant Trunk Links for Faster Recovery on page 2903](#)

## Example: Configuring Redundant Trunk Links for Faster Recovery



**NOTE:** This example uses Junos OS for EX Series switches or QFX Series with support for the Enhanced Layer 2 Software (ELS) configuration style. If your EX Series switch runs software that does not support ELS, see *Example: Configuring Redundant Trunk Links for Faster Recovery*. For ELS details, see “Getting Started with Enhanced Layer 2 Software” on page 41.

You can manage network convergence by configuring both a primary link and a secondary link on a switch; this is called a redundant trunk group (RTG). If the primary link in a redundant trunk group fails, it passes its known MAC address locations to the secondary link, which automatically takes over after one minute.

This example describes how to create a redundant trunk group with a primary and a secondary link:

- [Requirements on page 2903](#)
- [Overview and Topology on page 2904](#)
- [Disabling RSTP on Switches 1 and 2 on page 2906](#)
- [Configuring Redundant Trunk Links on Switch 3 on page 2906](#)
- [Verification on page 2907](#)

## Requirements

This example uses the following hardware and software components:

- Two EX Series or QFX Series distribution switches
- One EX Series or QFX Series access switch
- The appropriate software release for your platform:
  - For EX Series switches: Junos OS Release 13.2X50-D10 or later
  - For the QFX Series: Junos OS Release 13.2X50-D15 or later

Before you configure the redundant trunk links network on the access and distribution switches, be sure you have:

- Configured interfaces ge-0/0/9 and ge-0/0/10 on the access switch, Switch 3, as trunk interfaces.

- Configured one trunk interface on each distribution switch, Switch 1 and Switch 2.
- Connected the three switches as shown in the topology for this example (see [Figure 50](#)).

## Overview and Topology

In a typical enterprise network composed of distribution and access layers, a redundant trunk link provides a simple solution for trunk interface network recovery. When a trunk interface fails, data traffic is routed to another trunk interface after one minute, thereby keeping network convergence time to a minimum.

This example shows the configuration of a redundant trunk group that includes one primary link (and its interface) and one unspecified link (and its interface) that serves as the secondary link.

A second type of redundant trunk group, not illustrated in the example, consists of two unspecified links (and their interfaces); in this case, neither of the links is primary. The software selects an active link by comparing the port numbers of the two links and activating the link with the higher port number. For example, if the two link interfaces use interfaces ge-0/1/0 and ge-0/1/1, the software activates ge-0/1/1. (In the interface names, the final number is the port number.)

The two links in a redundant trunk group generally operate the same way, whether they are configured as primary/unspecified or unspecified/unspecified. Data traffic initially passes through the active link but is blocked on the inactive link. While data traffic is blocked on the secondary link, note that Layer 2 control traffic is still permitted if the link is active. For example, an LLDP session can be run between two switches on the secondary link. If the active link either goes down or is disabled administratively, it broadcasts a list of its known MAC addresses for data traffic; the other link immediately picks up and adds the MAC addresses to its address table, becomes active, and begins forwarding traffic.

The one difference in operation between the two types of redundant trunk groups occurs when a primary link is active, goes down, is replaced by the secondary link, and then reactivates. When a primary link is re-enabled while the secondary link is active, the primary link waits 1 second (you can change the time interval by using the preempt cutover timer to accommodate your network) and then takes over as the active link. In other words, the primary link has priority and is always activated if it is available. This differs from the behavior of two unspecified links, both of which act as equals. Because the unspecified links are equal, the active link remains active until it either goes down or is disabled administratively; this is the only time that the other unspecified link learns the MAC addresses and immediately becomes active.

The example given here illustrates a primary/unspecified configuration for a redundant trunk group because that configuration gives you more control and is more commonly used.





**NOTE:** Rapid Spanning Tree Protocol (RSTP) is enabled by default on the switches to create a loop-free topology, but an interface is not allowed to be in both a redundant trunk group and in a spanning-tree protocol topology at the same time. You will need to disable RSTP on the two distribution switches in the example, Switch 1 and Switch 2. Spanning-tree protocols can, however, continue operating in other parts of the network—for example, between the distribution switches and also in links between distribution switches and the enterprise core.

Figure 50 displays an example topology containing three switches. Switch 1 and Switch 2 make up the distribution layer, and Switch 3 makes up the access layer. Switch 3 is connected to the distribution layer through trunk interfaces ge-0/0/9.0 (Link 1) and ge-0/0/10.0 (Link 2).

Table 262 lists the components used in this redundant trunk group.

Because RSTP and RTGs cannot operate simultaneously on a switch, you disable RSTP on Switch 1 and Switch 2 in the first configuration task, and you disable RSTP on Switch 3 in the second task.

The second configuration task creates a redundant trunk group called example 1 on Switch 3. The trunk interfaces ge-0/0/9.0 and ge-0/0/10.0 are the two links configured in the second configuration task. You configure the trunk interface ge-0/0/9.0 as the primary link. You configure the trunk interface ge-0/0/10.0 as an unspecified link, which becomes the secondary link by default.

Figure 50: Topology for Configuring the Redundant Trunk Links




Table 262: Components of the Redundant Trunk Link Topology

| Property              | Settings                                                                                                                                                                                                                           |
|-----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Switch hardware       | <ul style="list-style-type: none"> <li>Switch 1–1 EX Series or QFX Series distribution switch</li> <li>Switch 2–1 EX Series or QFX Series distribution switch</li> <li>Switch 3–1 EX Series or QFX Series access switch</li> </ul> |
| Trunk interfaces      | On Switch 3 (access switch): ge-0/0/9.0 and ge-0/0/10.0                                                                                                                                                                            |
| Redundant trunk group | rtg0                                                                                                                                                                                                                               |

### Disabling RSTP on Switches 1 and 2

To disable RSTP on Switch 1 and Switch 2, perform this task on each switch:

**CLI Quick Configuration** To quickly disable RSTP on Switch 1 and Switch 2, copy the following command and paste it into each switch terminal window:

```
[edit]
set protocols rstp disable
```

**Step-by-Step Procedure** To disable RSTP on Switch 1 and Switch 2:

1. Disable RSTP on Switch 1 and Switch 2:

```
[edit]
user@switch# set protocols rstp disable
```

**Results** Check the results of the configuration:

```
[edit]
user@switch# show
protocols {
 rstp {
 disable;
 }
}
```

### Configuring Redundant Trunk Links on Switch 3

To configure redundant trunk links on Switch 3, perform this task:

**CLI Quick Configuration** To quickly configure the redundant trunk group rtg0 on Switch 3, copy the following commands and paste them into the switch terminal window:

```
[edit]
set protocols rstp disable
set switch-options redundant-trunk-group group rtg0 interface ge-0/0/9.0 primary
set switch-options redundant-trunk-group group rtg0 interface ge-0/0/10.0
set redundant-trunk-group group rtg0 preempt-cutover-timer 60
```

**Step-by-Step Procedure** Configure the redundant trunk group rtg0 on Switch 3.

1. Turn off RSTP:

```
[edit]
user@switch# set protocols rstp disable
```

2. Name the redundant trunk group `rtg0` while configuring trunk interface `ge-0/0/9.0` as the primary link and `ge-0/0/10` as an unspecified link to serve as the secondary link:

```
[edit switch-options]
user@switch# set redundant-trunk-group group rtg0 interface ge-0/0/9.0 primary
user@switch# set redundant-trunk-group group rtg0 interface ge-0/0/10.0
```

3. (Optional) Change the time interval (from the default of 1 second) that a re-enabled primary link waits to take over for an active secondary link:

```
[edit switch-options]
user@switch# set redundant-trunk-group group rtg0 preempt-cutover-timer 60
```

**Results** Check the results of the configuration:

```
[edit]
user@switch# show
switch-options
 redundant-trunk-group {
 group rtg0 {
 preempt-cutover-timer 60;
 interface ge-0/0/9.0 {
 primary;
 }
 interface ge-0/0/10.0;
 }
 }
protocols {
 rstp {
 disable;
 }
}
```

## Verification

To confirm that the configuration is set up correctly, perform this task:

- [Verifying That a Redundant Trunk Group Was Created on page 2907](#)

### Verifying That a Redundant Trunk Group Was Created

**Purpose** Verify that the redundant trunk group `rtg0` has been created on Switch 1 and that trunk interfaces are members of the redundant trunk group.

**Action** List all redundant trunk groups configured on the switch:

```
user@switch> show redundant-trunk-group
```

| Group name | Interface   | State  | Time of last flap | Flap count |
|------------|-------------|--------|-------------------|------------|
| rtg0       | ge-0/0/9.0  | Up/Pri | Never             | 0          |
|            | ge-0/0/10.0 | Up     | Never             | 0          |

**Meaning** The **show redundant-trunk-group** command lists all redundant trunk groups configured on the switch as well as the interface names and their current states (up or down for an unspecified link, and up or down and primary for a primary link). For this configuration example, the output shows that the redundant trunk group **rtg0** is configured on the switch. The **Up** beside the interfaces indicates that both link cables are physically connected. The **Pri** beside trunk interface **ge-0/0/9.0** indicates that it is configured as the primary link.

**Related Documentation**

- [Understanding Redundant Trunk Links on page 2901](#)

## PART 56

# Resilient Hashing

- [Understanding Resilient Hashing on page 2911](#)



# Understanding Resilient Hashing

- [Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 2911](#)
- [Understanding the Use of Resilient Hashing to Minimize Flow Remapping in Trunk/ECMP Groups on page 2917](#)
- [Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic \(QFX 10002 and QFX 10008 Switches\) on page 2919](#)
- [Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic \(CLI Procedure\) on page 2923](#)
- [Configuring Resilient Hashing for Trunk/ECMP Groups on page 2925](#)

## Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic

---

Juniper Networks EX Series and QFX Series use a hashing algorithm to determine how to forward traffic over a link aggregation group (LAG) bundle or to the next-hop device when equal-cost multipath (ECMP) is enabled.

The hashing algorithm makes hashing decisions based on values in various packet fields, as well as on some internal values like source port ID and source device ID. You can configure some of the fields that are used by the hashing algorithm.

This topic contains the following sections:

- [Understanding the Hashing Algorithm on page 2911](#)
- [IP \(IPv4 and IPv6\) on page 2913](#)
- [MPLS on page 2914](#)
- [MAC-in-MAC Packet Hashing on page 2915](#)
- [Layer 2 Header Hashing on page 2916](#)

## Understanding the Hashing Algorithm

The hashing algorithm is used to make traffic-forwarding decisions for traffic entering a LAG bundle or for traffic exiting a switch when ECMP is enabled.

For LAG bundles, the hashing algorithm determines how traffic entering a LAG bundle is placed onto the bundle's member links. The hashing algorithm tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.

For ECMP, the hashing algorithm determines how incoming traffic is forwarded to the next-hop device.

The hashing algorithm makes hashing decisions based on values in various packet fields, as well as on some internal values like source port ID and source device ID. The packet fields used by the hashing algorithm varies by the packet's EtherType and, in some instances, by the configuration on the switch. The hashing algorithm recognizes the following EtherTypes:

- IP (IPv4 and IPv6)
- MPLS
- MAC-in-MAC

Traffic that is not recognized as belonging to any of these EtherTypes is hashed based on the Layer 2 header. IP and MPLS traffic are also hashed based on the Layer 2 header when a user configures the hash mode as Layer 2 header.

You can configure some fields that are used by the hashing algorithm to make traffic forwarding decisions. You cannot, however, configure how certain values within a header are used by the hashing algorithm.

Note the following points regarding the hashing algorithm:

- The fields selected for hashing are based on the packet type only. The fields are not based on any other parameters, including forwarding decision (bridged or routed) or egress LAG bundle configuration (Layer 2 or Layer 3).
- The same fields are used for hashing unicast and multicast packets. Unicast and multicast packets are, however, hashed differently.
- The same fields are used by the hashing algorithm to hash ECMP and LAG traffic, but the hashing algorithm hashes ECMP and LAG traffic differently. LAG traffic uses a trunk hash while ECMP uses ECMP hashing. Both LAG and ECMP use the same RTAG7 seed but use different offsets of that 128B seed to avoid polarization. The initial config of the HASH function to use the trunk and ECMP offset are set at the PFE Init time. The different hashing ensures that traffic is not polarized when a LAG bundle is part of the ECMP next-hop path.
- The same fields are used for hashing regardless of whether the switch is or is not participating in a mixed or non-mixed Virtual Chassis or Virtual Chassis Fabric (VCF).

The fields used for hashing by each EtherType as well as the fields used by the Layer 2 header are discussed in the following sections.



## IP (IPv4 and IPv6)

Payload fields in IPv4 and IPv6 packets are used by the hashing algorithm when IPv4 or IPv6 packets need to be placed onto a member link in a LAG bundle or sent to the next-hop device when ECMP is enabled.

The hash mode is set to Layer 2 payload field, by default. IPv4 and IPv6 payload fields are used for hashing when the hash mode is set to Layer 2 payload.

If the hash mode is configured to Layer 2 header, IPv4, IPv6, and MPLS packets are hashed using the Layer 2 header fields. If you want incoming IPv4, IPv6, and MPLS packets hashed by the source MAC address, destination MAC address, or EtherType fields, you must set the hash mode to Layer 2 header.

Table 263 displays the IPv4 and IPv6 payload fields that are used by the hashing algorithm, by default.

- ✓—Field is used by the hashing algorithm, by default.
- X—Field is not used by the hashing algorithm, by default.
- (configurable)—Field can be configured to be used or not used by the hashing algorithm.

**Table 263: IPv4 and IPv6 Hashing Fields**

| Fields                  | EX4300         |                | QFX5100        |                |
|-------------------------|----------------|----------------|----------------|----------------|
|                         | LAG            | ECMP           | LAG            | ECMP           |
| Source MAC              | X              | X              | X              | X              |
| Destination MAC         | X              | X              | X              | X              |
| EtherType               | X              | X              | X              | X              |
| VLAN ID                 | X              | X              | X              | X              |
|                         | (configurable) | (configurable) | (configurable) | (configurable) |
| Source IP or IPv6       | ✓              | ✓              | ✓              | ✓              |
|                         | (configurable) | (configurable) | (configurable) | (configurable) |
| Destination IP or IPv6  | ✓              | ✓              | ✓              | ✓              |
|                         | (configurable) | (configurable) | (configurable) | (configurable) |
| Protocol (IPv4 only)    | ✓              | ✓              | ✓              | ✓              |
|                         | (configurable) | (configurable) | (configurable) | (configurable) |
| Next header (IPv6 only) | ✓              | ✓              | ✓              | ✓              |
|                         | (configurable) | (configurable) | (configurable) | (configurable) |

**Table 263: IPv4 and IPv6 Hashing Fields (*continued*)**

| Fields                      | EX4300              |                     | QFX5100             |                     |
|-----------------------------|---------------------|---------------------|---------------------|---------------------|
| Layer 4 Source Port         | ✓<br>(configurable) | ✓<br>(configurable) | ✓<br>(configurable) | ✓<br>(configurable) |
| Layer 4 Destination Port    | ✓<br>(configurable) | ✓<br>(configurable) | ✓<br>(configurable) | ✓<br>(configurable) |
| IPv6 Flow label (IPv6 only) | X                   | X                   | X                   | X                   |

## MPLS

The hashing algorithm hashes MPLS packets using the source IP, destination IP, MPLS label 0, MPLS label 1, and MPLS label 2 fields. On the QFX5200 only, LSR routers also support ECMP. ECMP uses these fields for hashing on an LSR router:

- L3VPN: MPLS Labels (top 3 labels), source IP, destination IP, and ingress port ID
- L2Circuit: MPLS Labels (top 3 labels) and ingress port ID

[Table 264](#) displays the MPLS payload fields that are used by the hashing algorithm, by default:

- ✓—Field is used by the hashing algorithm, by default.
- X—Field is not used by the hashing algorithm, by default.

The fields used by the hashing algorithm for MPLS packet hashing are not user-configurable.

The source IP and destination IP fields are not always used for hashing. For non-terminated MPLS packets, the payload is checked if the bottom of stack (BoS) flag is seen in the packet. If the payload is IPv4 or IPv6, then the IP source address and IP destination address fields are used for hashing along with the MPLS labels. If the BoS flag is not seen in the packet, only the MPLS labels are used for hashing.

**Table 264: MPLS Hashing Fields**

| Field           | EX4300 | QFX5100 | QFX5200 |
|-----------------|--------|---------|---------|
| Source MAC      | X      | X       | X       |
| Destination MAC | X      | X       | X       |
| EtherType       | X      | X       | X       |
| VLAN ID         | X      | X       | X       |
| Source IP       | ✓      | ✓       | ✓       |

Table 264: MPLS Hashing Fields (*continued*)

| Field                          | EX4300 | QFX5100 | QFX5200             |
|--------------------------------|--------|---------|---------------------|
| Destination IP                 | ✓      | ✓       | ✓                   |
| Protocol (for IPv4 packets)    | X      | X       | X                   |
| Next header (for IPv6 packets) | X      | X       | X                   |
| Layer 4 Source Port            | X      | X       | X                   |
| Layer 4 Destination Port       | X      | X       | X                   |
| IPv6 Flow lab                  | X      | X       | X                   |
| MPLS label 0                   | ✓      | ✓       | ✓                   |
| MPLS label 1                   | ✓      | ✓       | ✓                   |
| MPLS label 2                   | ✓      | ✓       | ✓                   |
| Ingress Port ID                | X      | X       | ✓                   |
|                                |        |         | (LSR and L2Circuit) |

## MAC-in-MAC Packet Hashing

Packets using the MAC-in-MAC EtherType are hashed by the hashing algorithm using the Layer 2 payload source MAC, Layer 2 payload destination MAC, and Layer 2 payload EtherType fields. See [Table 265](#).

Hashing using the fields in the MAC-in-MAC EtherType packet is first supported on EX4300 switches in Release 13.2X51-D20. Hashing using the fields in the MAC-in-MAC EtherType is not supported on earlier releases.

The fields used by the hashing algorithm for MAC-in-MAC hashing are not user-configurable.

- ✓—Field is used by the hashing algorithm, by default.
- X—Field is not used by the hashing algorithm, by default.

Table 265: MAC-in-MAC Hashing Fields

| Field                           | EX4300 | QFX5100 | QFX5200 |
|---------------------------------|--------|---------|---------|
| Layer 2 Payload Source MAC      | ✓      | ✓       | ✓       |
| Layer 2 Payload Destination MAC | ✓      | ✓       | ✓       |

**Table 265: MAC-in-MAC Hashing Fields (*continued*)**

| Field                      | EX4300 | QFX5100 | QFX5200 |
|----------------------------|--------|---------|---------|
| Layer 2 Payload EtherType  | ✓      | ✓       | ✓       |
| Layer 2 Payload Outer VLAN | X      | X       | X       |

## Layer 2 Header Hashing

Layer 2 header fields are used by the hashing algorithm when a packet's EtherType is not recognized as IP (IPv4 or IPv6), MPLS, or MAC-in-MAC. The Layer 2 header fields are also used for hashing IPv4, IPv6, and MPLS traffic instead of the payload fields when the hash mode is set to Layer 2 header.

- ✓—Field is used by the hashing algorithm, by default.
- X—Field is not used by the hashing algorithm, by default.
- (configurable)—Field can be configured to be used or not used by the hashing algorithm.

**Table 266: Layer 2 Header Hashing Fields**

| Field           | EX4300              | QFX5100             | QFX5200             |
|-----------------|---------------------|---------------------|---------------------|
| Source MAC      | ✓<br>(configurable) | ✓<br>(configurable) | ✓<br>(configurable) |
| Destination MAC | ✓<br>(configurable) | ✓<br>(configurable) | ✓<br>(configurable) |
| EtherType       | ✓<br>(configurable) | ✓<br>(configurable) | ✓<br>(configurable) |
| VLAN ID         | X<br>(configurable) | X<br>(configurable) | ✓<br>(configurable) |

### Related Documentation

- [Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic \(CLI Procedure\) on page 2923](#)
- [Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic \(QFX 10002 and QFX 10008 Switches\) on page 2919](#)

## Understanding the Use of Resilient Hashing to Minimize Flow Remapping in Trunk/ECMP Groups

You use resilient hashing to minimize flow remapping across members of a trunk/ECMP group in a load-balanced system. You can configure resilient hashing in link aggregation groups (LAGs) and in equal cost multipath (ECMP) groups.



**NOTE:** Resilient hashing in LAGs is not supported on QFX10002 and QFX 10008 Switches.

- [Why You Might Want to Use Resilient Hashing and How It Works with Static Hashing on page 2917](#)
- [Limitations and Caveats for Resilient Hashing on page 2918](#)
- [Resilient Hashing on LAGs on page 2918](#)
- [Resilient Hashing on ECMP on page 2919](#)

### Why You Might Want to Use Resilient Hashing and How It Works with Static Hashing

Resilient hashing works in conjunction with the default static hashing algorithm. When members are added to or deleted from a trunk/ECMP group, the static hashing algorithm might remap destination paths. Resilient hashing distributes traffic across all members of a group by tracking the flow's member utilization. When a flow is affected by a member change, the Packet Forwarding Engine rebalances the flow by reprogramming the flow set table.

Resilient hashing thus provides the following benefits:

- Minimizes traffic-distribution imbalances among members of a trunk/ECMP group when members are added to or deleted from the group.
- Minimizes the impact on flows bound to unaffected members when a new member is added or an existing member is deleted from the group.

In normal hash-based load balancing, with the static hashing algorithm used alone, flows are assigned to members through the mathematical mod (%) operation. Any increase or decrease in the number of group members results in a complete remapping of flows to member IDs, as shown in the following example:

- Member ID = Hash (key) mod (number of members in group)
- Example:
  - Hash (key) = 10
  - $10 \bmod 5 = 0$  (member with ID 0 is selected for flow)
  - $10 \bmod 4 = 2$  (member with ID 2 is selected for the same flow when the number of members is decreased by 1)

Resilient hashing minimizes the destination path remapping when a member in the trunk/ECMP group is added or deleted.

When the flow is affected by a member change in the group, resilient hashing rebalances the flow by reprogramming the flow set table.

**Table 267: Destination Path Results for Static Hashing and for Resilient Hashing When Members Are Added to or Deleted from Trunk Groups**

| Trunk Group Size | Normal (Static) Hashing Result                       | Resilient Hashing Result                                                                | Notes                                                                         |
|------------------|------------------------------------------------------|-----------------------------------------------------------------------------------------|-------------------------------------------------------------------------------|
| 4                | Hash(10) % 4 = 2<br>Flow is assigned to member ID 2. | Flow is assigned to one of four group members based on flow set table entries.          | Original trunk/ECMP group size is 4.                                          |
| 3                | Hash(10) % 3 = 1<br>Flow is assigned to member ID 1. | Flow is assigned to same member as in the previous case.                                | Delete one member from original trunk/ECMP group. Trunk/ECMP group size is 3. |
| 5                | Hash(10) % 5 = 0<br>Flow is assigned to member ID 0. | There is minimal redistribution of flows from other members to this newly added member. | Add one member to original trunk group. Trunk/ECMP group size is 5.           |

## Limitations and Caveats for Resilient Hashing

Notice the following limitation and caveats for the resilient hashing feature:

- Resilient hashing applies only to unicast traffic.
- Resilient hashing supports a maximum of 1024 trunk groups, with each group having a maximum of 256 members.
- Resilient hashing does not guarantee that traffic distribution is even across all group members—it depends on the traffic pattern and on the organization of the resilient hashing flow set table in hardware. Resilient hashing *minimizes* remapping of flows to destination links when members are added to or deleted from the group.
- If resilient hashing is enabled on a trunk group or ECMP group and if **set forwarding-options enhanced-hash-key** with one of the options **hash-mode**, **inet**, **inet6**, or **layer2** is used, some flows might change destination links, because the new hash parameters might generate new hash indexes for the flows, and hence the new destination links.
- Resilient hashing is not supported on Virtual Chassis port (VCP) links.

## Resilient Hashing on LAGs

A LAG combines Ethernet interfaces (members) to form a logical point-to-point link that increases bandwidth, provides reliability, and allows load balancing. Resilient hashing minimizes destination remapping behavior when a new member is added or deleted from the LAG.

A resilient hashing configuration on LAGs is per-aggregated-Ethernet-interface–based.

## Resilient Hashing on ECMP

An ECMP group for a route contains multiple next-hop equal cost addresses for the same destination in the routing table. (Routes of equal cost have the same preference and metric values.)

Junos OS uses the static hashing algorithm to choose one of the next-hop addresses in the ECMP group to install in the forwarding table. Resilient hashing enhances ECMPs by minimizing destination remapping behavior when a new member is added or deleted from the ECMP group.

A resilient hashing configuration on ECMP is global—it applies to all ECMP groups.

### Related Documentation

- [Configuring Resilient Hashing for Trunk/ECMP Groups on page 2925](#)

## Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic (QFX 10002 and QFX 10008 Switches)

Juniper Networks QFX 10002 and QFX 10008 switches use a hashing algorithm to determine how to forward traffic over a link aggregation group (LAG) bundle or to the next-hop device when equal-cost multipath (ECMP) is enabled.

The hashing algorithm makes hashing decisions based on values in various packet fields. You can configure some of the fields that are used by the hashing algorithm.

This topic contains the following sections:

- [Understanding the Hashing Algorithm on page 2919](#)
- [IP \(IPv4 and IPv6\) on page 2920](#)
- [MPLS on page 2921](#)
- [Layer 2 Header Hashing on the QFX10002 and QFX 10008 Switches on page 2922](#)

## Understanding the Hashing Algorithm

The hashing algorithm is used to make traffic-forwarding decisions for traffic exiting a switch when ECMP is enabled.

For LAG bundles, the hashing algorithm determines how traffic entering a LAG bundle is placed onto the bundle's member links to manage bandwidth by evenly load-balancing traffic across the outgoing links.

For ECMP, the hashing algorithm determines how incoming traffic is forwarded to the next-hop device.

The hashing algorithm makes hashing decisions based on values in various packet fields, as well as on the hash seed value. The packet fields used by the hashing algorithm varies by the packet's EtherType and, in some instances, by the configuration on the switch.

The hashing algorithm recognizes the following EtherTypes:

- IP (IPv4 and IPv6)
- MPLS
- MAC-in-MAC

You can configure some fields that are used by the hashing algorithm to make traffic forwarding decisions. You cannot, however, configure how certain values within a header are used by the hashing algorithm.

Note the following points regarding the hashing algorithm:

- The fields selected for hashing are based on the packet type only. The fields are not based on any other parameters, including forwarding decision (bridged or routed) or egress LAG bundle configuration (Layer 2 or Layer 3).
- The same fields are used for hashing unicast and multicast packets. Unicast and multicast packets are, however, hashed differently.
- The same fields are used by the hashing algorithm to hash ECMP and LAG traffic. The hashing ensures that traffic is not polarized when a LAG bundle is part of the ECMP next-hop path.

The fields used for hashing by each EtherType as well as the fields used by the Layer 2 header are discussed in the following sections.

## IP (IPv4 and IPv6)

Payload fields in IPv4 and IPv6 packets are used by the hashing algorithm when IPv4 or IPv6 packets need to be placed onto a member link in a LAG bundle or sent to the next-hop device when ECMP is enabled.

The table below displays the IPv4 and IPv6 payload fields that are used by the hashing algorithm, by default. All fields are configurable.

- ✓—Field is used by the hashing algorithm, by default.
- X—Field is not used by the hashing algorithm, by default.
- (configurable)—Field can be configured to be used or not used by the hashing algorithm.

**Table 268: IPv4 and IPv6 Hashing Fields for the QFX10002 Switch and QFX 10008 Switch**

| Fields                 | QFX10002 Switch and QFX 10008 Switch |      |
|------------------------|--------------------------------------|------|
|                        | LAG                                  | ECMP |
| Source MAC (0:47)      | X                                    | X    |
| Destination MAC (0:47) | X                                    | X    |
| EtherType (0:15)       | ✓                                    | ✓    |
| VLAN ID (0:11)         | ✓                                    | ✓    |



**Table 268: IPv4 and IPv6 Hashing Fields for the QFX10002 Switch and QFX 10008 Switch (*continued*)**

| Fields                              | QFX10002 Switch and QFX 10008 Switch |   |
|-------------------------------------|--------------------------------------|---|
| Source IP [1](0:31)                 | ✓                                    | ✓ |
| Destination IP (0:31)               | ✓                                    | ✓ |
| Incoming Port                       | ✓                                    | ✓ |
| Protocol (for IPv4 packets)(0:7)    | X                                    | X |
| Next header (for IPv6 packets)(0:7) | X                                    | X |
| Layer 4 Source Port (0:15)          | ✓                                    | ✓ |
| Layer 4 Destination Port (0:15)     | ✓                                    | ✓ |
| IPv6 Flow label (0:19)              | ✓                                    | ✓ |

## MPLS

The hashing algorithm hashes MPLS packets using the source IP, destination IP, MPLS label 0, MPLS label 1, and MPLS label 2 fields.

The table below displays the MPLS payload fields that are used by the hashing algorithm, by default:

- ✓—Field is used by the hashing algorithm, by default.
- X—Field is not used by the hashing algorithm, by default.

The fields used by the hashing algorithm for MPLS packet hashing are not user-configurable.

The source IP and destination IP fields are not always used for hashing. For non-terminated MPLS packets, the payload is checked if the bottom of stack (BoS) flag is seen in the packet. If the payload is IPv4 or IPv6, then the IP source address and IP destination address fields are used for hashing along with the MPLS labels. If the BoS flag is not seen in the packet, only the MPLS labels are used for hashing.

**Table 269: MPLS Hashing Fields for the QFX10002 Switch and QFX 10008 Switch**

| Fields                 | QFX10002 Switch |      |
|------------------------|-----------------|------|
|                        | LAG             | ECMP |
| Source MAC (0:47)      | X               | X    |
| Destination MAC (0:47) | X               | X    |

**Table 269: MPLS Hashing Fields for the QFX10002 Switch and QFX 10008 Switch (continued)**

| Fields                              | QFX10002 Switch |   |
|-------------------------------------|-----------------|---|
| EtherType (0:15)                    | ✓               | ✓ |
| VLAN ID (0:11)                      | ✓               | ✓ |
| Source IP [1](0:31)                 | ✓               | ✓ |
| Destination IP (0:31)               | ✓               | ✓ |
| Incoming Port                       | ✓               | ✓ |
| Protocol (for IPv4 packets)(0:7)    | X               | X |
| Next header (for IPv6 packets)(0:7) | X               | X |
| Layer 4 Source Port (0:15)          | ✓               | ✓ |
| Layer 4 Destination Port (0:15)     | ✓               | ✓ |
| IPv6 Flow label (0:19)              | ✓               | ✓ |
| MPLS label 0 (0:19)                 | ✓               | ✓ |
| MPLS label 1 (0:19)                 | ✓               | ✓ |
| MPLS label 2 (0:19)                 | ✓               | ✓ |

### Layer 2 Header Hashing on the QFX10002 and QFX 10008 Switches

Layer 2 header fields are used by the hashing algorithm when a packet's EtherType is not recognized as IP (IPv4 or IPv6), MPLS, or MAC-in-MAC. The Layer 2 header fields are also used for hashing IPv4, IPv6, and MPLS traffic instead of the payload fields when the hash mode is set to Layer 2 header.

- ✓—Field is used by the hashing algorithm, by default.
- X—Field is not used by the hashing algorithm, by default.

**Table 270: Layer 2 Header Hashing Fields for the QFX10002 Switch**

| Fields                 | QFX10002 and QFX 10008 Switches |      |
|------------------------|---------------------------------|------|
|                        | LAG                             | ECMP |
| Source MAC (0:47)      | X                               | X    |
| Destination MAC (0:47) | X                               | X    |

**Table 270: Layer 2 Header Hashing Fields for the QFX10002 Switch** (*continued*)

| Fields                | QFX10002 and QFX 10008 Switches |   |
|-----------------------|---------------------------------|---|
| EtherType (0:15)      | ✓                               | ✓ |
| Inner VLAN ID         | X                               | X |
| VLAN ID (0:11)        | ✓                               | ✓ |
| Source IP [1] (0:31)  | X                               | X |
| Destination IP (0:31) | X                               | X |

**Related Documentation**

- [Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic \(CLI Procedure\) on page 2923](#)
- [Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 2911](#)

## Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure)

Juniper Networks EX Series and QFX Series switches use a hashing algorithm to determine how to forward traffic over a Link Aggregation group (LAG) bundle or to the next-hop device when equal-cost multipath (ECMP) is enabled.

The hashing algorithm makes hashing decisions based on values in various packet fields.. You can configure some of the fields that are used by the hashing algorithm.

Configuring the fields used by the hashing algorithm is useful in scenarios where most of the traffic entering the bundle is similar and the traffic needs to be managed in the LAG bundle. For instance, if the only difference in the IP packets for all incoming traffic is the source and destination IP address, you can tune the hashing algorithm to make hashing decisions more efficiently by configuring the algorithm to make hashing decisions using only those fields.



**NOTE:** Configuring the hash mode is not supported on QFX10002 and QFX10008 switches.

- [Configuring the Hashing Algorithm to Use Fields in the Layer 2 Header for Hashing on page 2924](#)
- [Configuring the Hashing Algorithm to Use Fields in the IP Payload for Hashing on page 2924](#)
- [Configuring the Hashing Algorithm to Use Fields in the IPv6 Payload for Hashing on page 2924](#)

## Configuring the Hashing Algorithm to Use Fields in the Layer 2 Header for Hashing

To configure the hashing algorithm to use fields in the Layer 2 header for hashing:

1. Configure the hash mode to Layer 2 header:

```
[edit forwarding-options enhanced-hash-key]
user@switch# set hash-mode layer2-header
```

The default hash mode is Layer 2 payload. Therefore, this step must be performed if you have not previously configured the hash mode.

2. Configure the fields in the Layer 2 header that the hashing algorithm uses for hashing:

```
[edit forwarding-options enhanced-hash-key]
user@switch# set layer2 {no-destination-mac-address | no-ether-type |
no-source-mac-address | vlan-id}
```

By default, the hashing algorithm uses the values in the destination MAC address, Ethertype, and source MAC address fields in the header to hash traffic on the LAG. You can configure the hashing algorithm to not use the values in these fields by configuring **no-destination-mac-address**, **no-ether-type**, or **no-source-mac-address**.

You can also configure the hashing algorithm to include the VLAN ID field in the header by configuring the **vlan-id** option.

If you want the hashing algorithm to not use the Ethertype field for hashing:

```
[edit forwarding-options enhanced-hash-key]
user@switch# set layer2 no-ether-type
```

## Configuring the Hashing Algorithm to Use Fields in the IP Payload for Hashing

To configure the hashing algorithm to use fields in the IP payload for hashing:

1. Configure the hash mode to Layer 2 payload:

```
[edit forwarding-options enhanced-hash-key]
user@switch# set hash-mode layer2-payload
```

The IP payload is not checked by the hashing algorithm unless the hash mode is set to Layer 2 payload. The default hash mode is Layer 2 payload.

2. Configure the fields in the IP payload that the hashing algorithm uses for hashing:

```
[edit forwarding-options enhanced-hash-key]
user@switch# set inet {no-ipv4-destination-address | no-ipv4-source-address |
no-l4-destination-port | no-l4-source-port | no-protocol | vlan-id}
```

For instance, if you want the hashing algorithm to ignore the Layer 4 destination port, Layer 4 source port, and protocol fields and instead hash traffic based only on the IPv4 source and destination addresses:

```
[edit forwarding-options enhanced-hash-key]
user@switch# set inet no-l4-destination-port no-l4-source-port no-protocol
```

## Configuring the Hashing Algorithm to Use Fields in the IPv6 Payload for Hashing

To configure the hashing algorithm to use fields in the IPv6 payload for hashing:

1. Configure the hash mode to Layer 2 payload:

```
[edit forwarding-options enhanced-hash-key]
user@switch# set hash-mode layer2-payload
```

The IPv6 payload is not checked by the hashing algorithm unless the hash mode is set to Layer 2 payload. The default hash mode is Layer 2 payload.

2. Configure the fields in the IPv6 payload that the hashing algorithm uses for hashing:

```
[edit forwarding-options enhanced-hash-key]
user@switch# set inet6 {no-ipv6-destination-address | no-ipv6-source-address |
no-l4-destination-port | no-l4-source-port | no-next-header | vlan-id}
```

For instance, if you want the hashing algorithm to ignore the Layer 4 destination port, Layer 4 source port, and the Next Header fields and instead hash traffic based only on the IPv6 source and IPv6 destination address fields only:

```
[edit forwarding-options enhanced-hash-key]
user@switch# set inet6 no-l4-destination-port no-l4-source-port no-next-header
```

#### Related Documentation

- [Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 2911](#)
- [Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic \(QFX 10002 and QFX 10008 Switches\) on page 2919](#)
- [Understanding Aggregated Ethernet Interfaces and LACP](#)

## Configuring Resilient Hashing for Trunk/ECMP Groups

You use resilient hashing to minimize flow remapping across members of a trunk/ECMP group in a load-balanced system. You can configure resilient hashing in link aggregation groups (LAGs) and in equal cost multipath (ECMP) sets.

This topic includes:

1. [Configuring Resilient Hashing on LAGs on page 2925](#)
2. [Configuring Resilient Hashing on ECMP Groups on page 2926](#)

### Configuring Resilient Hashing on LAGs



**NOTE:** Configuring resilient hashing on LAGs is not supported on QFX10002 and QFX10008 switches.

To enable resilient hashing for a LAG:

- Configure resilient hashing on the aggregated Ethernet interface:

```
[edit interfaces]
user@switch# set aex aggregated-ether-options resilient-hash
```

## Configuring Resilient Hashing on ECMP Groups

To enable resilient hashing for ECMP groups:

- Configure resilient hashing for ECMP:

```
[edit forwarding-options]
user@switch# set enhanced-hash-key ecmp-resilient-hash
```

### Related Documentation

- [Understanding the Use of Resilient Hashing to Minimize Flow Remapping in Trunk/ECMP Groups on page 2917](#)

## PART 57

# Uplink Failure Detection

- [Understanding Uplink Failure Detection on page 2929](#)





# Understanding Uplink Failure Detection

- [Overview of Uplink Failure Detection on page 2929](#)
- [Configuring Interfaces for Uplink Failure Detection on page 2931](#)
- [Example: Configuring Interfaces for Uplink Failure Detection on page 2932](#)

## Overview of Uplink Failure Detection

---

Uplink failure detection allows a switch to detect link failure on uplink interfaces and to propagate this information to the downlink interfaces, so that servers connected to those downlinks can switch over to secondary interfaces.

Uplink failure detection supports network adapter teaming and provides network redundancy. In network adapter teaming, all of the network interface cards (NICs) on a server are configured in a primary or secondary relationship and share the same IP address. When the primary link goes down, the server transparently shifts the connection to the secondary link. With uplink failure detection, the switch monitors uplink interfaces for link failures. When it detects a failure, it disables the downlink interfaces. When the server detects disabled downlink interfaces, it switches over to the secondary link to help ensure that the traffic of the failed link is not dropped.

This topic describes:


- [Uplink Failure Detection Configuration on page 2929](#)
- [Failure Detection Pair on page 2930](#)

## Uplink Failure Detection Configuration

Uplink failure detection allows switches to monitor uplink interfaces to spot link failures. When a switch detects a link failure, it automatically disables the downlink interfaces bound to the uplink interface. A server that is connected to the disabled downlink interface triggers a network adapter failover to a secondary link to avoid any traffic loss.

[Figure 51](#) illustrates a typical setup for uplink failure detection.

Figure 51: Uplink Failure Detection Configuration on Switches



9040577

For uplink failure detection, you specify a group of uplink interfaces to be monitored and downlink interfaces to be brought down when an uplink fails. The downlink interfaces are bound to the uplink interfaces within the group. If all uplink interfaces in a group go down, then the switch brings down all downlink interfaces within that group. If any uplink interface returns to service, then the switch brings all downlink interfaces in that group back to service.

The switch can monitor both physical interface links and logical interface links for uplink failures, but you must put the two types of interfaces into separate groups.



**NOTE:** For logical interfaces, the server must send keepalives between the switch and the server to detect failure of logical links.

## Failure Detection Pair

Uplink failure detection requires that you create pairs of uplink and downlink interfaces in a group. Each pair includes one of each of the following:

- A link-to-monitor interface—The link-to-monitor interfaces specify the uplinks the switch monitors. You can configure a maximum of 48 uplink interfaces as link-to-monitor interfaces for a group.
- A link-to-disable interface—The link-to-disable interfaces specify the downlinks the switch disables when the switch detects an uplink failure. You can configure a maximum of 48 downlinks to disable in the group.

The link-to-disable interfaces are bound to the link-to-monitor interfaces within the group. When a link-to-monitor interface returns to service, the switch automatically enables all link-to-disable interfaces in the group.

- Related Documentation**
- [Configuring Interfaces for Uplink Failure Detection on page 2931](#)
  - [Example: Configuring Interfaces for Uplink Failure Detection on page 2932](#)

## Configuring Interfaces for Uplink Failure Detection

You can configure uplink failure detection to help ensure balanced traffic flow. Using this feature, switches can monitor and detect link failure on uplink interfaces and can propagate the failure information to downlink interfaces, so that servers connected to those downlinks can switch over to secondary interfaces.

Follow these configuration guidelines:

- Configure an interface in only one group.
- Configure a maximum of 48 groups for each switch.
- Configure a maximum of 48 uplinks to monitor and a maximum of 48 downlinks to disable in each group.
- Configure physical links and logical links in separate groups.

To configure uplink failure detection on a switch:

1. Specify a name for an uplink failure detection group:

```
[edit protocols]
user@switch# set uplink-failure-detection group group-name
```

2. Add an uplink interface to the group:

```
[edit protocols]
user@switch# set uplink-failure-detection group group-name link-to-monitor interface-name
```

3. Repeat Step 2 for each uplink interface you add to the group.

4. Add a downlink interface to the group:

```
[edit protocols]
user@switch# set uplink-failure-detection group group-name link-to-disable interface-name
```

5. Repeat Step 4 for each downlink interface you add to the group.



**NOTE:** After you have configured an uplink failure detection group, use the `show uplink-failure-detection group (Uplink Failure Detection) group-name` command to verify that all interfaces in the group are up. If the interfaces are down, uplink failure detection does not work.

- Related Documentation**
- [Overview of Uplink Failure Detection on page 2929](#)
  - [Example: Configuring Interfaces for Uplink Failure Detection on page 2932](#)
  - [Verifying That Uplink Failure Detection Is Working Correctly](#)

## Example: Configuring Interfaces for Uplink Failure Detection

---

Uplink failure detection allows a switch to detect link failure on uplink interfaces and to propagate the failure information to the downlink interfaces. All of the network interface cards (NICs) on a server are configured as being either the primary link or the secondary link and share the same IP address. When the primary link goes down, the server transparently shifts the connection to the secondary link to ensure that the traffic on the failed link is not dropped.

This example describes:

- [Requirements on page 2932](#)
- [Overview and Topology on page 2932](#)
- [Configuring Uplink Failure Detection on Both Switches on page 2933](#)
- [Verification on page 2934](#)

### Requirements

This example uses the following software and hardware components:

- Junos OS Release 12.1 or later for the QFX Series
- Two QFX3500 switches
- Two aggregation switches
- One dual-homed server

### Overview and Topology

The topology in this example illustrates how to configure uplink failure detection on Switch A and Switch B. Switch A and Switch B are both configured with a link-to-monitor interface (the uplink interface to the aggregation switch) and a link-to-disable interface (the downlink interface to the server). For simplicity, only one group of link-to-monitor interfaces and link-to-disable interfaces is configured for each switch. The server is dual-homed to both Switch A and Switch B. In this scenario, if the link-to-monitor interface to Switch A is disabled, the server uses the link-to-monitor interface to Switch B instead.



**NOTE:** This example does not describe how to configure the dual-homed server or the aggregation switches. Please refer to the documentation for each of these devices for more information.

---

Figure 51 illustrates a typical setup for uplink failure detection.

Figure 52: Uplink Failure Detection Configuration on Switches




Table 271 lists uplink failure settings for each QFX3500 switch.

Table 271: Settings for Uplink Failure Protection Example

| Switch A                                                                                                                                                         | Switch B                                                                                                                                                         |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <ul style="list-style-type: none"><li>• Group name: Group1</li><li>• Link-to-monitor interface: xe-0/0/0</li><li>• Link-to-disable interface: xe-0/0/1</li></ul> | <ul style="list-style-type: none"><li>• Group name: Group2</li><li>• Link-to-monitor interface: xe-0/0/0</li><li>• Link-to-disable interface: xe-0/0/1</li></ul> |

Configuring Uplink Failure Detection on Both Switches

To configure uplink failure detection on both switches, perform these tasks:

CLI Quick Configuration

To quickly configure uplink failure protection on Switch A and Switch B, copy the following commands and paste them into the switch terminal window:  
  
[edit protocols]  
set uplink-failure-detection group group1  
set uplink-failure-detection group group2  
set uplink-failure-detection group group1 link-to-monitor xe-0/0/0  
set uplink-failure-detection group group2 link-to-monitor xe-0/0/0  
set uplink-failure-detection group group1 link-to-disable xe-0/0/1  
set uplink-failure-detection group group2 link-to-disable xe-0/0/1

Step-by-Step Procedure

To configure uplink failure protection on both switches:  
  
1. Specify a name for the uplink failure detection group on Switch A:  
  
[edit protocols]  
user@switch# set uplink-failure-detection group group1  
  
2. Add an uplink interface to the group on Switch A:  
  
[edit protocols]  
user@switch# set uplink-failure-detection group group1 link-to-monitor xe-0/0/0

3. Add a downlink interface to the group on Switch A:  

```
[edit protocols]
user@switch# set uplink-failure-detection group group1 link-to-disable xe-0/0/1
```
4. Specify a name for the uplink failure detection group on Switch B:  

```
[edit protocols]
user@switch# set uplink-failure-detection group group2
```
5. Add an uplink interface to the group on Switch B:  

```
[edit protocols]
user@switch# set uplink-failure-detection group group2 link-to-monitor xe-0/0/0
```
6. Add a downlink interface to the group on Switch B:  

```
[edit protocols]
user@switch# set uplink-failure-detection group group2 link-to-disable xe-0/0/1
```

**Results** Display the results of the configuration:

```
uplink-failure-detection {
 group {
 group1 {
 link-to-monitor {
 xe-0/0/0;
 }
 link-to-disable {
 xe-0/0/1;
 }
 }
 group2 {
 link-to-monitor {
 xe-0/0/0;
 }
 link-to-disable {
 xe-0/0/1;
 }
 }
 }
}
```

## Verification

To verify that uplink failure detection is working correctly, perform the following tasks on Switch A and Switch B:

- [Verifying That Uplink Failure Detection is Working Correctly on page 2934](#)

### Verifying That Uplink Failure Detection is Working Correctly

**Purpose** Verify that the switch disables the downlink interface when it detects an uplink failure.

- Action** 1. View the current uplink failure detection status:

```
user@switch> show uplink-failure-detection
Group : group1
Uplink : xe-0/0/0*
Downlink : xe-0/0/1*
Failure Action : Inactive
```



**NOTE:** The asterisk (\*) indicates that the link is up.

2. Disable the uplink interface:

```
[edit]
user@switch# set interface xe-0/0/0 disable
```

3. Save the configuration on the switch.

4. View the current uplink failure detection status:

```
user@switch> show uplink-failure-detection
Group : group1
Uplink : xe-0/0/0
Downlink : xe-0/0/1
Failure Action : Active
```

**Meaning** The output in Step 1 shows that the uplink interface is up, and hence that the downlink interface is also up, and that the status of **Failure Action** is **Inactive**.

The output in Step 4 shows that both the uplink and downlink interfaces are down (there are no asterisks after the interface name) and that the status of **Failure Action** is changed to **Active**. This output shows that uplink failure detection is working.

- Related Documentation**
- [Overview of Uplink Failure Detection on page 2929](#)
  - [Configuring Interfaces for Uplink Failure Detection on page 2931](#)





## PART 58

# Configuration Statements and Operational Commands

- [Ethernet OAM Link Fault Management Configuration Statements on page 2939](#)
- [GRE Configuration Statements on page 2961](#)
- [Interfaces Configuration Statements on page 2971](#)
- [IP Directed Broadcast Configuration Statement on page 3039](#)
- [LAGs and LACP Configuration Statements on page 3041](#)
- [Load Balancing Configuration Statements on page 3053](#)
- [Local Link Bias Configuration Statement on page 3057](#)
- [Redundant Trunk Groups Configuration Statements on page 3059](#)
- [Resilient Hashing Configuration Statements on page 3065](#)
- [Uplink Failure Detection Configuration Statements on page 3079](#)
- [Ethernet OAM Link Fault Management Operational Command on page 3083](#)
- [Interfaces Operational Commands on page 3089](#)
- [LAGs and LACP Operational Commands on page 3203](#)
- [Redundant Trunk Group Operational Command on page 3211](#)
- [Resilient Hashing Operational Command on page 3215](#)
- [Uplink Failure Detection Operational Command on page 3221](#)



# Ethernet OAM Link Fault Management Configuration Statements

- [action \(OAM LFM\) on page 2940](#)
- [action-profile on page 2941](#)
- [allow-remote-loopback on page 2942](#)
- [ethernet \(OAM LFM\) on page 2943](#)
- [event-thresholds on page 2945](#)
- [event \(OAM LFM\) on page 2946](#)
- [frame-error on page 2946](#)
- [frame-period on page 2947](#)
- [frame-period-summary on page 2947](#)
- [oam on page 2948](#)
- [interface \(OAM LFM\) on page 2950](#)
- [link-adjacency-loss on page 2951](#)
- [link-discovery on page 2951](#)
- [link-down on page 2952](#)
- [link-event-rate on page 2952](#)
- [link-fault-management on page 2953](#)
- [negotiation-options on page 2954](#)
- [no-allow-link-events on page 2954](#)
- [pdu-interval on page 2955](#)
- [pdu-threshold on page 2955](#)
- [remote-loopback on page 2956](#)
- [symbol-period on page 2956](#)
- [syslog \(OAM LFM\) on page 2957](#)
- [traceoptions \(OAM LFM\) on page 2958](#)

## action (OAM LFM)

---

|                                 |                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>action {<br/>  syslog;<br/>  link-down;<br/>}</pre>                                                                                                               |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet link-fault-management</a> ]                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.                                                                                                   |
| <b>Description</b>              | <p>Define the action or actions to be taken when the OAM link fault management (LFM) fault event occurs.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure)</a> on page 2768</li></ul>                          |

## action-profile

|                                 |                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> action-profile <i>profile-name</i>;   action {     syslog;     link-down;   }   event {     link-adjacency-loss;     link-event-rate {       frame-error <i>count</i>;       frame-period <i>count</i>;       frame-period-summary <i>count</i>;       symbol-period <i>count</i>;     }   } </pre> |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet link-fault-management</a> ]                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.                                                                                                                                                                                                                                      |
| <b>Description</b>              | <p>Configure an Ethernet OAM link fault management (LFM) action profile by specifying a profile name.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                       |
| <b>Options</b>                  | <i>profile-name</i> —Name of the action profile.                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 2768</a></li> </ul>                                                                                                                                                             |

## allow-remote-loopback

---

|                                 |                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | allow-remote-loopback;                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam</a> <a href="#">ethernet link-fault-management interface</a> <i>interface-name</i> ]                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.                                                                                                                                                                                               |
| <b>Description</b>              | Advertise that the interface is capable of getting into loopback mode. Enable remote loopback in Ethernet OAM link fault management (LFM) on all Ethernet interfaces or the specified interface on the EX Series switch.                                           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 2770</a></li><li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 2768</a></li></ul> |

## ethernet (OAM LFM)

```

Syntax ethernet {
 connectivity-fault-management {
 action-profile profile-name {
 action {
 interface-down;
 }
 default-actions {
 interface-down;
 }
 event {
 adjacency-loss;
 }
 }
 }
 esp-traceoptions {
 file filename <files number> <no-stamp> <replace> <size size> <world-readable |
 no-world-readable>;
 flag (all | error | esp | interface | krt | lib | normal | task | timer);
 }
 linktrace {
 age (30m | 10m | 1m | 30s | 10s);
 path-database-size path-database-size;
 }
 maintenance-domain domain-name {
 level number;
 mip-half-function (none | default | explicit);
 name-format (character-string | none | dns | mac+2oct);
 maintenance-association ma-name {
 continuity-check {
 hold-interval minutes;
 interface-status-tlv;
 interval (10m | 10s | 1m | 1s | 100ms);
 loss-threshold number;
 port-status-tlv;
 }
 mep mep-id {
 auto-discovery;
 direction down;
 interface interface-name;
 priority
 remote-mep mep-id {
 action-profile profile-name;
 sla-iterator-profile profile-name {
 data-tlv-size size;
 iteration-count count-value;
 priority priority-value;
 }
 }
 }
 }
 short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id);
 }
 }
 performance-monitoring {

```

```
sla-iterator-profiles {
 profile-name {
 calculation-weight {
 delay delay-value;
 delay-variation delay-variation-value;
 }
 cycle-time cycle-time-value;
 iteration-period iteration-period-value;
 measurement-type two-way-delay;
 passive;
 }
}
}
traceoptions {
 file filename <files number> <match regex> <size size> <world-readable |
 no-world-readable>;
 flag flag ;
 no-remote-trace;
}
}
link-fault-management {
 action-profile profile-name;
 action {
 syslog;
 link-down;
 }
 event {
 link-adjacency-loss;
 link-event-rate {
 frame-error count;
 frame-period count;
 frame-period-summary count;
 symbol-period count;
 }
 }
}
interface interface-name {
 link-discovery (active | passive);
 pdu-interval interval;
 pdu-threshold threshold-value;
 remote-loopback;
 event-thresholds {
 frame-error count;
 frame-period count;
 frame-period-summary count;
 symbol-period count;
 }
 negotiation-options {
 allow-remote-loopback;
 no-allow-link-events;
 }
}
}
traceoptions {
 file filename <files number> <match regex> <size size> <world-readable |
 no-world-readable>;
 flag flag ;
 no-remote-trace;
```



```

 }
 }
}

```

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.<br><b>connectivity-fault-management</b> introduced in Junos OS Release 10.2 for EX Series switches.                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Provide IEEE 802.3ah Operation, Administration, and Maintenance (OAM) support for Ethernet interfaces on EX Series switches or configure connectivity fault management (CFM) for IEEE 802.1ag Operation, Administration, and Management (OAM) support on the switches.<br><br>The remaining statements are explained separately.                                                                                                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 2770</a></li> <li>• <a href="#">Example: Configuring Ethernet OAM Connectivity Fault Management on EX Series Switches</a></li> <li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 2768</a></li> <li>• <a href="#">Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure)</a></li> </ul> |

## event-thresholds

|                                 |                                                                                                                                                                                                   |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> event-thresholds {   <a href="#">frame-error</a> count;   <a href="#">frame-period</a> count;   <a href="#">frame-period-summary</a> count;   <a href="#">symbol-period</a> count; } </pre> |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet link-fault-management interface</a> <i>interface-name</i> ]                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.                                                                                                                              |
| <b>Description</b>              | Configure threshold limit values for link events in periodic OAM PDUs.<br><br>The remaining statements are explained separately.                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 2768</a></li> </ul>                                                   |

## event (OAM LFM)

---

|                                 |                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>event {<br/>  link-adjacency-loss;<br/>  link-event-rate {<br/>    frame-error <i>count</i>;<br/>    frame-period <i>count</i>;<br/>    frame-period-summary <i>count</i>;<br/>    symbol-period <i>count</i>;<br/>  }<br/>}</pre> |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet link-fault-management action-profile</a> <i>profile-name</i> ]                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.                                                                                                                                                                    |
| <b>Description</b>              | <p>Configure link events in an action profile for Ethernet OAM link fault management (LFM).</p> <p>The remaining statements are explained separately.</p>                                                                               |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 2768</a></li></ul>                                                                                           |

## frame-error

---

|                                 |                                                                                                                                                                                                                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>frame-error <i>count</i>;</pre>                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet link-fault-management event link-event-rate</a> ],<br>[edit protocols <a href="#">oam ethernet link-fault-management interface</a> <i>interface-name</i> <a href="#">event-thresholds</a> ]                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.                                                                                                                                                                                                         |
| <b>Description</b>              | <p>Configure the threshold value for sending frame error events or taking the action specified in the action profile.</p> <p>Frame errors occur on the underlying physical layer. The threshold is reached when the number of frame errors reaches the configured value.</p> |
| <b>Options</b>                  | <p><i>count</i>—Threshold count in seconds for frame error events.</p> <p><b>Range:</b> 1 through 100 seconds</p> <p><b>Default:</b> 1 second</p>                                                                                                                            |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 2768</a></li></ul>                                                                                                                                |

## frame-period

---

|                                 |                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>frame-period count;</code>                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet link-fault-management event link-event-rate</a> ],<br>[edit protocols <a href="#">oam ethernet link-fault-management interface interface-name event-thresholds</a> ]                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.                                                                                                                                                                                   |
| <b>Description</b>              | <p>Configure the number of frame errors within the last N frames that has exceeded a threshold.</p> <p>Frame errors occur on the underlying physical layer. The threshold is reached when the number of frame errors reaches the configured value.</p> |
| <b>Options</b>                  | <p><i>count</i>—Threshold count in seconds for frame error events.</p> <p><b>Range:</b> 1 through 100 seconds</p>                                                                                                                                      |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 2768</a></li> </ul>                                                                                                        |

## frame-period-summary

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>frame-period-summary count;</code>                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet link-fault-management event link-event-rate</a> ],<br>[edit protocols <a href="#">oam ethernet link-fault-management interface interface-name event-thresholds</a> ]                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | <p>Configure the threshold value for sending frame period summary error events or taking the action specified in the action profile.</p> <p>An errored frame second is any 1-second period that has at least one errored frame. This event is generated if the number of errored frame seconds is equal to or greater than the specified threshold for that period.</p> |
| <b>Options</b>                  | <p><i>count</i>—Threshold count in seconds for frame period summary error events.</p> <p><b>Range:</b> 1 through 100 seconds</p>                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 2768</a></li> </ul>                                                                                                                                                                                                                         |

## oam

```

Syntax oam {
 ethernet {
 connectivity-fault-management {
 action-profile profile-name {
 action {
 interface-down;
 }
 default-actions {
 interface-down;
 }
 event {
 adjacency-loss;
 }
 }
 }
 linktrace {
 age (30m | 10m | 1m | 30s | 10s);
 path-database-size path-database-size;
 }
 maintenance-domain domain-name {
 level number;
 mip-half-function (none | default | explicit);
 name-format (character-string | none | dns | mac+2oct);
 maintenance-association ma-name {
 continuity-check {
 hold-interval minutes;
 interface-status-tlv;
 interval (10m | 10s | 1m | 1s | 100ms);
 loss-threshold number;
 port-status-tlv;
 }
 mep mep-id {
 auto-discovery;
 direction down;
 interface interface-name;
 remote-mep mep-id {
 action-profile profile-name;
 }
 }
 }
 }
 }
 performance-monitoring {
 sla-iterator-profiles {
 profile-name {
 calculation-weight {
 delay delay-value;
 delay-variation delay-variation-value;
 }
 cycle-time cycle-time-value;
 iteration-period iteration-period-value;
 measurement-type two-way-delay;
 passive;
 }
 }
 }
}

```

```

 }
 }
}
link-fault-management {
 action-profile profile-name;
 action {
 syslog;
 link-down;
 }
 event {
 link-adjacency-loss;
 link-event-rate {
 frame-error count;
 frame-period count;
 frame-period-summary count;
 symbol-period count;
 }
 }
}
interface interface-name {
 link-discovery (active | passive);
 pdu-interval interval;
 pdu-threshold threshold-value;
 remote-loopback;
 event-thresholds {
 frame-error count;
 frame-period count;
 frame-period-summary count;
 symbol-period count;
 }
 negotiation-options {
 allow-remote-loopback;
 no-allow-link-events;
 }
}
}
}
}

```

**Hierarchy Level** [edit protocols]

**Release Information** Statement introduced in Junos OS Release 9.4 for EX Series switches.  
**connectivity-fault-management** introduced in Junos OS Release 10.2 for EX Series switches.

**Description** Provide IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM) support for Ethernet interfaces on EX Series switches or configure connectivity fault management (CFM) for IEEE 802.lag Operation, Administration, and Management (OAM) support on the switches.

The remaining statements are explained separately.

**Required Privilege Level** interface—To view this statement in the configuration.  
 interface-control—To add this statement to the configuration.

- Related Documentation**
- [Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 2770](#)
  - [Example: Configuring Ethernet OAM Connectivity Fault Management on EX Series Switches](#)
  - [Configuring Ethernet OAM Link Fault Management \(CLI Procedure\) on page 2768](#)
  - [Configuring Ethernet OAM Connectivity Fault Management \(CLI Procedure\)](#)

---

## interface (OAM LFM)

---

**Syntax**

```
interface interface-name {
 link-discovery (active | passive);
 pdu-interval interval;
 pdu-threshold threshold-value;
 remote-loopback;
 event-thresholds {
 frame-error count;
 frame-period count;
 frame-period-summary count;
 symbol-period count;
 }
 negotiation-options {
 allow-remote-loopback;
 no-allow-link-events;
 }
}
```

**Hierarchy Level** [edit protocols [oam ethernet link-fault-management](#)]

**Release Information** Statement introduced in Junos OS Release 9.4 for EX Series switches.

**Description** Configure Ethernet OAM link fault management (LFM) for all interfaces or for specific interfaces.

The remaining statements are explained separately.

**Options** *interface-name*—Name of the interface to be enabled for IEEE 802.3ah OAM link fault management (LFM) support.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

- Related Documentation**
- [Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 2770](#)
  - [Configuring Ethernet OAM Link Fault Management \(CLI Procedure\) on page 2768](#)

## link-adjacency-loss

---

|                                 |                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | link-adjacency-loss;                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet link-fault-management action-profile event</a> ]                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.                                                                                                                                                                                               |
| <b>Description</b>              | Configure <b>loss of adjacency</b> event with the IEEE 802.3ah link fault management (LFM) peer. When included, the loss of adjacency event triggers the action specified under the <a href="#">action</a> statement.                                              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 2770</a></li><li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 2768</a></li></ul> |

## link-discovery

---

|                                 |                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | link-discovery (active   passive);                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet link-fault-management interface interface-name</a> ]                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Specify the discovery mode used for IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM) support. The discovery process is triggered automatically when OAM 802.3ah functionality is enabled on an interface. Link monitoring is done when the interface sends periodic OAM PDUs.         |
| <b>Options</b>                  | <p><i>active</i>—In active mode, the interface discovers and monitors the peer on the link if the peer also supports IEEE 802.3ah OAM functionality.</p> <p><i>passive</i>—In passive mode, the peer initiates the discovery process.</p> <p>Once the discovery process is initiated, both sides participate in discovery.</p> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 2768</a></li></ul>                                                                                                                                                                                  |

## link-down

---

|                                 |                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | link-down;                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet link-fault-management action-profile action</a> ]                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.                                                                          |
| <b>Description</b>              | Mark the interface as down for transit traffic.                                                                                               |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 2768</a></li></ul> |

## link-event-rate

---

|                                 |                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | link-event-rate {<br><a href="#">frame-error</a> <i>count</i> ;<br><a href="#">frame-period</a> <i>count</i> ;<br><a href="#">frame-period-summary</a> <i>count</i> ;<br><a href="#">symbol-period</a> <i>count</i> ;<br>} |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet link-fault-management action-profile event</a> ]                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.                                                                                                                                                       |
| <b>Description</b>              | Configure the number of link fault management (LFM) events per second.<br><br>The remaining statements are explained separately.                                                                                           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 2768</a></li></ul>                                                                              |



## link-fault-management

```
Syntax link-fault-management {
 action-profile profile-name;
 action {
 syslog;
 link-down;
 }
 event {
 link-adjacency-loss;
 link-event-rate {
 frame-error count;
 frame-period count;
 frame-period-summary count;
 symbol-period count;
 }
 }
 interface interface-name {
 link-discovery (active | passive);
 pdu-interval interval;
 pdu-threshold threshold-value;
 remote-loopback;
 event-thresholds {
 frame-error count;
 frame-period count;
 frame-period-summary count;
 symbol-period count;
 }
 negotiation-options {
 allow-remote-loopback;
 no-allow-link-events;
 }
 }
 }
```

**Hierarchy Level** [edit protocols [oam](#) [ethernet](#)]

**Release Information** Statement introduced in Junos OS Release 9.4 for EX Series switches.

**Description** Configure Ethernet OAM link fault management (LFM) for all interfaces or for specific interfaces.

The remaining statements are explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 2770](#)
- [Configuring Ethernet OAM Link Fault Management \(CLI Procedure\) on page 2768](#)

## negotiation-options

---

|                                 |                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>negotiation-options {<br/>    allow-remote-loopback;<br/>    no-allow-link-events;<br/>}</pre>                                                                                                             |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam</a> <a href="#">ethernet</a> <a href="#">link-fault-management</a> <a href="#">interface</a> <i>interface-name</i> ]                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.                                                                                                                                            |
| <b>Description</b>              | <p>Enable and disable IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM) features for Ethernet interfaces.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 2768</a></li></ul>                                                                   |

## no-allow-link-events

---

|                                 |                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>no-allow-link-events;</pre>                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam</a> <a href="#">ethernet</a> <a href="#">link-fault-management</a> <a href="#">interface</a> <i>interface-name</i> <a href="#">negotiation-options</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.                                                                                                                     |
| <b>Description</b>              | Disable the sending of link event TLVs.                                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 2768</a></li></ul>                                            |

## pdu-interval

|                                 |                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>pdu-interval <i>interval</i>;</code>                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit protocols <b>oam</b> <b>ethernet link-fault-management interface</b> <i>interface-name</i> ]                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.                                                                                                                                                                                                  |
| <b>Description</b>              | Specify the periodic OAM PDU sending interval for fault detection. It is used for IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM) support.                                                                                  |
| <b>Options</b>                  | <i>interval</i> —Periodic OAM PDU sending interval.<br><b>Range:</b> 400 through 1000 milliseconds<br><b>Default:</b> 1000 milliseconds                                                                                                                               |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 2770</a></li> <li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 2768</a></li> </ul> |

## pdu-threshold

|                                 |                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>pdu-threshold <i>threshold-value</i>;</code>                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit protocols <b>oam</b> <b>ethernet link-fault-management interface</b> <i>interface-name</i> ]                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.                                                                                                               |
| <b>Description</b>              | Configure how many protocol data units (PDUs) are missed before declaring the peer lost in Ethernet OAM link fault management (LFM) for all interfaces or for specific interfaces. |
| <b>Options</b>                  | <i>threshold-value</i> —Number of PDUs missed before declaring the peer lost.<br><b>Range:</b> 3 through 10 PDUs<br><b>Default:</b> 3 PDUs                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 2768</a></li> </ul>                                    |

## remote-loopback

---

|                                 |                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | remote-loopback;                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam</a> <a href="#">ethernet</a> <a href="#">link-fault-management</a> <a href="#">interface</a> <i>interface-name</i> ]                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.                                                                                                                                                                                               |
| <b>Description</b>              | Set the data terminal equipment (DTE) in loopback mode. Remove the statement from the configuration to take the DTE out of loopback mode. It is used for IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM) support.        |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 2770</a></li><li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 2768</a></li></ul> |

## symbol-period

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | symbol-period <i>count</i> ;                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam</a> <a href="#">ethernet</a> <a href="#">link-fault-management</a> <a href="#">action-profile</a> <i>profile-name</i> ; <a href="#">event</a> <a href="#">link-event-rate</a> ] ,<br>[edit protocols <a href="#">oam</a> <a href="#">ethernet</a> <a href="#">link-fault-management</a> <a href="#">interface</a> <i>interface-name</i> <a href="#">event-thresholds</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Configure the threshold for sending symbol period events or taking the action specified in the action profile.<br><br>Symbol code errors occur on the underlying physical layer. The symbol period threshold is reached when the number of symbol errors reaches the configured value within the period. You cannot configure the default value to a different value.                                      |
| <b>Options</b>                  | <i>count</i> —Threshold count in seconds for symbol period events.<br><b>Range:</b> 1 through 100 seconds                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 2768</a></li></ul>                                                                                                                                                                                                                                                              |

## syslog (OAM LFM)

---

|                                 |                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | syslog;                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet link-fault-management action-profile</a> <i>profile-name</i> ; <a href="#">action</a> ]              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.                                                                          |
| <b>Description</b>              | Generate a system log message for the Ethernet Operation, Administration, and Maintenance (OAM) link fault management (LFM) event.            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure)</a> on page 2768</li></ul> |

## traceoptions (OAM LFM)

---

**Syntax**    traceoptions {  
              file *filename* <files *number*> <match *regex*> <size *size*> <world-readable |  
              no-world-readable>;  
              flag *flag* ;  
              no-remote-trace;  
              }

**Release Information**    Statement introduced in JUNOS Release 10.2 for EX Series switches.

**Description**    Configure tracing options the link fault management.

**Options**    file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.

**files *number***—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

**Range:** 2 through 1000

**Default:** 3 files

**flag *flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:

- **action-profile**—Trace action profile invocation events.
- **all**—Trace all events.
- **configuration**—Trace configuration events.
- **protocol**—Trace protocol processing events.
- **routing socket**—Trace routing socket events.

**match**—(Optional) Refine the output to log only those lines that match the given regular expression.

**no-world-readable**—(Optional) Restrict file access to the user who created the file.

**no-remote-trace**—(Optional) Disable the remote trace.

**size *size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **files** option.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through 1 GB

**Default:** 128 KB

**Default:** If you do not include this option, tracing output is appended to an existing trace file.

**world-readable**—(Optional) Enable unrestricted file access.

|                           |                                                             |
|---------------------------|-------------------------------------------------------------|
| <b>Required Privilege</b> | routing—To view this statement in the configuration.        |
| <b>Level</b>              | routing-control—To add this statement to the configuration. |

|                              |                                                                                                                                                                                                                                                                    |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 2770</a></li><li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 2768</a></li></ul> |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|





# GRE Configuration Statements

- [allow-fragmentation](#) on page 2962
- [copy-tos-to-outer-ip-header](#) on page 2962
- [destination \(Tunnels\)](#) on page 2963
- [do-not-fragment](#) on page 2964
- [family](#) on page 2965
- [routing-instance](#) on page 2966
- [source](#) on page 2966
- [tunnel](#) on page 2967
- [tunnel](#) on page 2968
- [unit \(Interfaces\)](#) on page 2969
- [tunnel-port](#) on page 2970

## allow-fragmentation

---

|                                 |                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | allow-fragmentation;                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit interfaces <i>gr-fpc/pic/port</i> unit <i>logical-unit-number</i> tunnel],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>gr-fpc/pic/port</i> unit <i>logical-unit-number</i> tunnel] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.2.<br>Statement introduced in Junos OS Release 15.1X53-D10 for QFX10000 switches.                                                                                   |
| <b>Description</b>              | Enable fragmentation of generic routing encapsulation (GRE) encapsulated packets regardless of maximum transmission unit (MTU) value.                                                                          |
| <b>Default</b>                  | By default, the GRE-encapsulated packets are dropped if the packet size exceeds the MTU setting of the egress interface.                                                                                       |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>reassemble-packets</i></li><li>• <i>Configuring Unicast Tunnels</i></li><li>• <i>Junos OS Services Interfaces Library for Routing Devices</i></li></ul>             |

## copy-tos-to-outer-ip-header

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | copy-tos-to-outer-ip-header;                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit interfaces <i>gr-fpc/pic/port</i> unit <i>logical-unit-number</i> ],<br>[edit interfaces <i>gre</i> unit <i>logical-unit-number</i> ],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>gr-fpc/pic/port</i> unit <i>logical-unit-number</i> ],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>gre</i> unit <i>logical-unit-number</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.2.<br>Support for GRE interfaces for Generalized MPLS (GMPLS) introduced in Junos OS Release 12.3R7.<br>Statement introduced in Junos OS Release 15.1X53-D10 for QFX10000 switches.                                                                                                                                                    |
| <b>Description</b>              | For GRE tunnel interfaces and GRE interfaces for GMPLS control channels only, enable the inner IP header's ToS bits to be copied to the outer IP packet header.                                                                                                                                                                                                                   |
| <b>Default</b>                  | If you omit this statement, the ToS bits in the outer IP header are set to 0.                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring a GRE Tunnel to Copy ToS Bits to the Outer IP Header</i></li></ul>                                                                                                                                                                                                                                                |

## destination (Tunnels)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>destination address;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet unnumbered-address <i>interface-name</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet unnumbered-address <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | For encrypted, PPP-encapsulated, and tunnel interfaces, specify the remote address of the connection.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <b>address</b> —Address of the remote side of the connection.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring the Interface Address on page 2824</a></li> <li>• <i>Configuring Generic Routing Encapsulation Tunneling (CLI Procedure)</i></li> <li>• <i>Junos OS Services Interfaces Library for Routing Devices</i></li> <li>• <i>point-to-point</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

## do-not-fragment

---

|                                 |                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | do-not-fragment;                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit interfaces <i>gr-fpc/pic/port</i> unit <i>logical-unit-number</i> tunnel],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>gr-fpc/pic/port</i> unit <i>logical-unit-number</i> tunnel] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.2.<br>Statement introduced in Junos OS Release 15.1X53-D10 for QFX10000 switches.                                                                                   |
| <b>Description</b>              | Disable fragmentation of GRE encapsulated packets. Set the do-not-fragment (DF) bit on the packets entering the GRE tunnel so that they do not get fragmented anywhere in the path.                            |
| <b>Default</b>                  | By default, fragmentation is disabled.                                                                                                                                                                         |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>reassemble-packets</i></li><li>• <i>Configuring Unicast Tunnels</i></li><li>• <i>Junos OS Services Interfaces Library for Routing Devices</i></li></ul>             |

## family

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>family <i>family</i> {     address <i>address</i> {         destination <i>address</i>;     } }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 15.1X53-D10 for QFX10000 switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Configure protocol family information for the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <p><b>family</b>—Protocol family:</p> <ul style="list-style-type: none"> <li>• <b>ccc</b>—Circuit cross-connect protocol suite</li> <li>• <b>inet</b>—IP version 4 (IPv4)</li> <li>• <b>inet6</b>—IP version 6 (IPv6)</li> <li>• <b>iso</b>—Open Systems Interconnection (OSI) International Organization for Standardization (ISO) protocol suite</li> <li>• <b>mlfr-end-to-end</b>—Multilink Frame Relay FRF.15</li> <li>• <b>mlfr-uni-nni</b>—Multilink Frame Relay FRF.16</li> <li>• <b>multilink-ppp</b>—Multilink Point-to-Point Protocol</li> <li>• <b>mpls</b>—MPLS</li> <li>• <b>tcc</b>—Translational cross-connect protocol suite</li> <li>• <b>tnp</b>—Trivial Network Protocol</li> <li>• <b>vpls</b>—Virtual private LAN service</li> </ul> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Link and Multilink Services Interfaces Feature Guide for Routing Devices</i></li> <li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

## routing-instance

---

|                                 |                                                                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>routing-instance {<br/>    destination <i>routing-instance-name</i>;<br/>}</code>                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> <b>tunnel</b> ],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> <b>tunnel</b> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 15.1X53-D10 for QFX10000 switches.                                                                                                           |
| <b>Description</b>              | Specify the destination routing instance that points to the routing table containing the tunnel destination address.                                                                                                                       |
| <b>Default</b>                  | The default Internet routing table <b>inet.0</b> .                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Tunnel Interfaces for Routing Table Lookup</i></li></ul>                                                                                                                            |

## source

---

|                                 |                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>source <i>source-address</i>;</code>                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> <b>tunnel</b> ]                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| <b>Description</b>              | Specify the source address of the tunnel.                                                                                                                                                                                                                                   |
| <b>Default</b>                  | If you do not specify a source address, the tunnel uses the unit's primary address as the source address of the tunnel.                                                                                                                                                     |
| <b>Options</b>                  | <b><i>source-address</i></b> —Address of the local side of the tunnel. This is the address that is placed in the outer IP header's source field.                                                                                                                            |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Tunnel Services Overview</i></li></ul>                                                                                                                                                                                           |

## tunnel

---

|                                 |                                                                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>tunnel {   destination destination-address;   source source-address;   ttl number; (not supported on QFX and OCX Series switches) }</pre>                                                                                        |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> ]                                                                                                                                                       |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| <b>Description</b>              | <p>Configure a tunnel. You can use the tunnel for unicast and multicast traffic or just for multicast traffic. You can also use tunnels for encrypted traffic.</p> <p>The remaining statements are explained separately.</p>          |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Generic Routing Encapsulation Tunneling (CLI Procedure)</i></li> </ul>                                                                                                        |

## tunnel

---

|                                 |                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>tunnel {<br/>  allow-fragmentation;<br/>  backup-destination <i>address</i>;<br/>  destination <i>destination-address</i>;<br/>  do-not-fragment;<br/>  key <i>number</i>;<br/>  routing-instance {<br/>    destination <i>routing-instance-name</i>;<br/>  }<br/>  source <i>source-address</i>;<br/>  ttl <i>number</i>;<br/>}</pre> |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> ],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> ]                                                                                                                              |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.1 for EX Series switches.                                                                                                                                                                                                                  |
| <b>Description</b>              | <p>Configure a tunnel. You can use the tunnel for unicast and multicast traffic or just for multicast traffic. You can also use tunnels for encrypted traffic or virtual private networks (VPNs).</p> <p>The remaining statements are explained separately.</p>                                                                             |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Encryption Interfaces</i></li><li>• <i>Junos OS VPNs Library for Routing Devices</i></li></ul>                                                                                                                                                                                       |



## unit (Interfaces)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> unit <i>logical-unit-number</i> {   peer-unit <i>unit-number</i>;   reassemble-packets;   tunnel {     allow-fragmentation;     backup-destination <i>address</i>;     destination <i>destination-address</i>;     do-not-fragment;     key <i>number</i>;     routing-instance {       destination <i>routing-instance-name</i>;     }     source <i>source-address</i>;     ttl <i>number</i>;   } }</pre> |
| <b>Hierarchy Level</b>          | <p>[edit interfaces <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i>]</p>                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 15.1X53-D10 for QFX10000 switches.</p>                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <p><b><i>logical-unit-number</i></b>—Number of the logical unit.</p> <p><b>Range:</b> 0 through 16,384</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Junos OS Network Interfaces Library for Routing Devices</i> for other statements that do not affect services interfaces.</li> </ul>                                                                                                                                                                                                                                    |

## tunnel-port

---

|                                 |                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>tunnel-port <i>port-number</i> tunnel-services;</code>                                                                                                                                                           |
| <b>Hierarchy Level</b>          | <code>[edit chassis fpc slot pic <i>pic-number</i>]</code>                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| <b>Description</b>              | Configure the port number for generic routing encapsulation (GRE) tunneling.                                                                                                                                           |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Generic Routing Encapsulation Tunneling (CLI Procedure)</i></li></ul>                                                                                           |

# Interfaces Configuration Statements

- [address on page 2973](#)
- [alarm on page 2977](#)
- [auto-negotiation on page 2978](#)
- [channel-speed on page 2979](#)
- [configured-flow-control on page 2980](#)
- [craft-lockout on page 2981](#)
- [description on page 2982](#)
- [ethernet \(Alarm\) on page 2983](#)
- [ethernet-switch-profile on page 2984](#)
- [ethernet-switching on page 2986](#)
- [ether-options on page 2987](#)
- [eui-64 on page 2988](#)
- [family on page 2989](#)
- [fibre-channel \(Alarm\) on page 2993](#)
- [filter on page 2994](#)
- [flow-control on page 2996](#)
- [fpc on page 2997](#)
- [gratuitous-arp-reply on page 2998](#)
- [hold-time \(Physical Interface\) on page 2999](#)
- [irb \(Interfaces\) on page 3001](#)
- [inet \(interfaces\) on page 3004](#)
- [inet6 \(interfaces\) on page 3005](#)
- [interface-mode on page 3006](#)
- [interface-range on page 3008](#)
- [interfaces on page 3010](#)
- [link-down on page 3017](#)
- [link-mode on page 3018](#)
- [link-speed on page 3019](#)

- [loopback \(Aggregated Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet\) on page 3020](#)
- [mac on page 3020](#)
- [management-ethernet \(Alarm\) on page 3021](#)
- [member on page 3022](#)
- [member-range on page 3023](#)
- [mtu on page 3024](#)
- [no-gratuitous-arp-request on page 3025](#)
- [pic on page 3026](#)
- [rx-buffers on page 3027](#)
- [source on page 3028](#)
- [speed on page 3029](#)
- [tag-protocol-id \(TPIDs Expected to Be Sent or Received\) on page 3030](#)
- [targeted-broadcast on page 3031](#)
- [traceoptions \(Individual Interfaces\) on page 3032](#)
- [tx-buffers on page 3033](#)
- [unit on page 3035](#)
- [vlan-id on page 3036](#)
- [vlan-tagging on page 3037](#)

## address

```

Syntax address address {
 arp ip-address (mac | multicast-mac) mac-address <publish>;
 broadcast address;
 destination address;
 destination-profile name;
 eui-64;
 master-only;
 multipoint-destination address dlcid dlcid-identifier;
 multipoint-destination address {
 epd-threshold cells;
 inverse-arp;
 oam-liveness {
 up-count cells;
 down-count cells;
 }
 oam-period (disable | seconds);
 shaping {
 (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate burst
 length);
 queue-length number;
 }
 vci vpi-identifier.vci-identifier;
 }
 primary;
 preferred;
 (vrrp-group | vrrp-inet6-group) group-number {
 (accept-data | no-accept-data);
 advertise-interval seconds;
 authentication-type authentication;
 authentication-key key;
 fast-interval milliseconds;
 (preempt | no-preempt) {
 hold-time seconds;
 }
 priority-number number;
 track {
 priority-cost seconds;
 priority-hold-time interface-name {
 interface priority;
 bandwidth-threshold bits-per-second {
 priority;
 }
 }
 }
 route ip-address/mask routing-instance instance-name priority-cost cost;
 }
 virtual-address [addresses];
 }
}

```

**Hierarchy Level** [edit interfaces *interface-name* unit *logical-unit-number* family *family*],  
 [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*  
 family *family*]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 11.1 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Configure the interface address.



.....

**NOTE:** The vrrp High Availability functionality is not available on the QFX Series.

.....

**Options** *address*—Address of the interface.

- In Junos OS Release 13.3 and later, when you configure an IPv6 host address and an IPv6 subnet address on an interface, the commit operation fails.
- In releases earlier than Junos OS Release 13.3, when you use the same configuration on an interface, the commit operation succeeds, but only one of the IPv6 addresses that was entered is assigned to the interface. The other address is not applied.



**NOTE:** If you configure the same address on multiple interfaces in the same routing instance, Junos OS uses only the first configuration. The remaining address configurations are ignored and can leave interfaces without an address. Interfaces that do not have an assigned address cannot be used as a donor interface for an unnumbered Ethernet interface.

For example, in the following configuration the address configuration of interface xe-0/0/1.0 is ignored:

```
interfaces {
 xe-0/0/0 {
 unit 0 {
 family inet {
 address 192.168.1.1/24;
 }
 }
 }
 xe-0/0/1 {
 unit 0 {
 family inet {
 address 192.168.1.1/24;
 }
 }
 }
}
```

For more information on configuring the same address on multiple interfaces, see [“Configuring the Interface Address” on page 2824](#).

The remaining statements are explained separately.



**NOTE:** The `edit logical-systems` hierarchy is not available on QFabric systems.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

- Related Documentation**
- *Configuring the Protocol Family*
  - *Junos OS Administration Library for Routing Devices*
  - *family*
  - *negotiate-address*
  - *unnumbered-address (Ethernet)*
  - *Junos OS Administration Library for Routing Devices*




## alarm

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>alarm {     interface-type {         alarm-name (red   yellow   ignore);     } }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit chassis],<br>[edit chassis interconnect-device <i>name</i> ],<br>[edit chassis node-group <i>name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 12.2 for the ACX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | <p>Configure the chassis alarms and whether they trigger a red or yellow alarm, or whether they are ignored. Red alarm conditions light the <b>RED ALARM</b> LED on either the router's craft interface or the switch's LCD screen and trigger an audible alarm if one is connected to the contact on the craft interface or LCD screen. Yellow alarm conditions light the <b>YELLOW ALARM</b> LED on either the router's craft interface or the switch's LCD screen and trigger an audible alarm if one is connected to the craft interface or LCD screen.</p> <p>To configure more than one alarm, include multiple <i>alarm-name</i> lines.</p> |
| <b>Options</b>                  | <p><i>alarm-name</i>—Alarm condition. For a list of conditions, see <i>System-Wide Alarms and Alarms for Each Interface Type</i>.</p> <p><i>ignore</i>—The specified alarm condition does not set off any alarm.</p> <p><i>interface-type</i>—Type of interface on which you are configuring the alarm: <b>atm</b>, <b>ethernet</b>, <b>sonet</b>, or <b>t3</b>.</p> <p><b>red</b>—The specified alarm condition sets off a red alarm.</p> <p><b>yellow</b>—The specified alarm condition sets off a yellow alarm.</p>                                                                                                                             |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Understanding Alarms</i></li> <li>• <i>Chassis Conditions That Trigger Alarms</i></li> <li>• <i>Chassis Alarm Messages on a QFX3500 Device</i></li> <li>• <i>Interface Alarm Messages</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                   |

## auto-negotiation

---


|                                                                                                                                                                                                                                                                                   |                                                                                                                                                                                                                                                                                                                                     |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                     | (auto-negotiation   no-auto-negotiation);                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                            | [edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">ether-options</a> ]                                                                                                                                                                                                                                              |
| <b>Release Information</b>                                                                                                                                                                                                                                                        | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                       |
| <b>Description</b>                                                                                                                                                                                                                                                                | Explicitly enable or disable autonegotiation. <ul style="list-style-type: none"><li>• <b>auto-negotiation</b>—Enable autonegotiation.</li><li>• <b>no-auto-negotiation</b>—Disable autonegotiation. When autonegotiation is disabled, you must explicitly configure link mode and speed options.</li></ul>                          |
| <b>Default</b>                                                                                                                                                                                                                                                                    | Autonegotiation is automatically enabled for Gigabit Ethernet interfaces. Autonegotiation is not an option for 10-Gigabit Ethernet interfaces. No explicit action is taken after the autonegotiation is complete or if the negotiation fails.                                                                                       |
| <div> <b>NOTE:</b> Autonegotiation is not supported on QFX5100 switches. On QFX5200 switches, autonegotiation is not supported on 40-Gigabit Ethernet and 100-Gigabit Ethernet interfaces</div> |                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                   | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                             |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                      | <ul style="list-style-type: none"><li>• <a href="#">speed on page 3029</a></li><li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2826</a></li><li>• <i>Configuring Gigabit and 10-Gigabit Ethernet Interfaces</i></li><li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li></ul> |

## channel-speed

|                                 |                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>channel-speed (10g   25g   50g;   100g   disable-auto-speed-detection) ;</code>                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | <code>[edit chassis fpc slot-number pic pic-number (port port-number   port-range port-range-low port-range-high)]</code>                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                     |
| <b>Description</b>              | —Enable the specified port on the Physical Interface Card (PIC) to perform in the specified channel speed. Additionally, you can disable auto-speed detection.                                                                                                                                                                                    |
| <b>Default</b>                  | <b>40g</b> (40-Gigabit Ethernet).                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <b>10g</b> —Set the channel speed to 10g (10-Gigabit Ethernet).<br><b>25g</b> —Set the channel speed to 25g (25-Gigabit Ethernet).<br><b>50g</b> —Set the channel speed to 50g (50-Gigabit Ethernet).<br><b>100g</b> —Set the channel speed to 100g (100-Gigabit Ethernet).<br><b>disable-auto-speed-detection</b> —Disable auto-speed detection. |
| <b>Required Privilege Level</b> | <b>interface</b> —To view this statement in the configuration.<br><b>interface-control</b> —To add this statement to the configuration.                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Channelizing Interfaces on page 2835</a></li> <li>• <i>Channelizing Interfaces on QFX5200 Switches</i></li> </ul>                                                                                                                                                                            |

## configured-flow-control

---

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <pre>configured-flow-control {<br/>    rx-buffers (on   off);<br/>    tx-buffers (on   off);<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Hierarchy Level          | [edit <a href="#">interfaces interface-name ether-options</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Release Information      | Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Description              | <p>Configure Ethernet PAUSE asymmetric flow control on an interface. You can set an interface to generate and send PAUSE messages, and you can set an interface to respond to PAUSE messages sent by the connected peer. You must set both the <b>rx-buffers</b> and the <b>tx-buffers</b> values when you configure asymmetric flow control.</p> <p>Use the <b>flow-control</b> and <b>no-flow-control</b> statements to enable and disable symmetric PAUSE on an interface. Symmetric flow control and asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.</p> <div><p><b>NOTE:</b> Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual “lanes” of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC) by applying a congestion notification profile to the interface.</p><p>Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.</p></div> |
| Default                  | Flow control is disabled. You must explicitly configure Ethernet PAUSE flow control on interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Options                  | The statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Required Privilege Level | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Related Documentation    | <ul style="list-style-type: none"><li>• <a href="#">congestion-notification-profile on page 7132</a></li><li>• <a href="#">flow-control on page 2996</a></li><li>• <i>Configuring CoS Asymmetric Ethernet PAUSE Flow Control</i></li><li>• <i>Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control</i></li><li>• <a href="#">Understanding CoS Flow Control (Ethernet PAUSE and PFC) on page 6493</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## craft-lockout

```

Syntax craft-lockout {
 alarm {
 interface-type {
 link-down (red | yellow | ignore);
 }
 }
 container-devices {
 device-count number;
 }
 fpc slot {
 pic pic-number {
 fibre-channel {
 port-range {
 port-range-low port-range-high;
 }
 }
 }
 }
 routing-engine
 on-disk-failure {
 disk-failure-action (halt | reboot);
 }
 }

```

**Hierarchy Level** [edit chassis -interconnect-device]

**Release Information** Statement introduced in Junos Release 11.3 for the QFX Series.

**Description** Disable the physical operation of the craft interface front panel.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- *Configuring the Junos OS to Disable the Physical Operation of the Craft Interface*

## description

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>description text;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> ],<br>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ]                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | <p>Provide a textual description of the interface or the logical unit. Any descriptive text you include is displayed in the output of the <b>show interfaces</b> commands, and is also exposed in the <b>ifAlias</b> Management Information Base (MIB) object. It has no effect on the operation of the interface on the router or switch.</p> <p>The textual description can also be included in the extended DHCP relay option 82 Agent Circuit ID suboption.</p> |
| <b>Options</b>                  | <b>text</b> —Text to describe the interface. If the text includes spaces, enclose the entire text in quotation marks.                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Adding an Interface Description to the Configuration</i></li><li>• <i>Adding a Logical Unit Description to the Configuration</i></li><li>• <i>Configuring Gigabit Ethernet Interfaces (CLI Procedure)</i></li><li>• <i>Using DHCP Relay Agent Option 82 Information</i></li></ul>                                                                                                                                        |

---

## ethernet (Alarm)

---

|                                 |                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | ethernet {<br>link-down (red   yellow   ignore);<br>}                                                                                                 |
| <b>Hierarchy Level</b>          | [edit chassis <b>alarm</b> ],<br>[edit chassis interconnect-device <i>name</i> <b>alarm</b> ],<br>[edit chassis node-group <i>name</i> <b>alarm</b> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.         |
| <b>Description</b>              | Configure alarms for an Ethernet interface.                                                                                                           |
| <b>Options</b>                  | The remaining statement is explained separately.—                                                                                                     |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Understanding Alarms</i></li><li>• <i>Interface Alarm Messages</i></li></ul>                               |

## ethernet-switch-profile

**Syntax**

```
ethernet-switch-profile {
 ethernet-policer-profile {
 input-priority-map {
 ieee802.1p premium [values];
 }
 output-priority-map {
 classifier {
 premium {
 forwarding-class class-name {
 loss-priority (high | low);
 }
 }
 }
 }
 }
 policer cos-policer-name {
 aggregate {
 bandwidth-limit bps;
 burst-size-limit bytes;
 }
 premium {
 bandwidth-limit bps;
 burst-size-limit bytes;
 }
 }
 tag-protocol-id tpid;
}
mac-learn-enable;
```

**Hierarchy Level** [edit interfaces *interface-name* gigether-options],  
[edit interfaces *interface-name* aggregated-ether-options],  
[edit interfaces *interface-name* aggregated-ether-options],  
[edit interfaces *interface-name* ether-options]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.  
Statement introduced in Junos OS Release 13.2 for the QFX Series.  
Statement introduced in Junos OS Release 13.2X50-D15 for the EX Series switches.

**Description**



**NOTE:** On QFX Series standalone switches, the `ethernet-policer-profile` CLI hierarchy and the `mac-learn-enable` statement are supported only on the Enhanced Layer 2 Switching CLI.

For Gigabit Ethernet IQ, 10-Gigabit Ethernet IQ2 and IQ2-E, and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC, aggregated Ethernet with Gigabit Ethernet IQ interfaces, the built-in Gigabit Ethernet port on the M7i router); 100-Gigabit Ethernet Type 5 PIC with CFP; and Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit



Ethernet, and aggregated Ethernet interfaces on EX Series switches, configure VLAN tag and MAC address accounting and filtering properties.

The remaining statements are explained separately.



**NOTE:** When you gather interfaces into a bridge domain, the `no-mac-learn-enable` statement at the [edit interfaces *interface-name* *gather-options* ethernet-switch-profile] hierarchy level is not supported. You must use the `no-mac-learning` statement at the [edit bridge-domains *bridge-domain-name* bridge-options interface *interface-name*] hierarchy level to disable MAC learning on an interface in a bridge domain. For information on disabling MAC learning for a bridge domain, see the *MX Series Layer 2 Configuration Guide*.

|                                 |                                                                                                                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>                  | If the <b>ethernet-switch-profile</b> statement is not configured, Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router) behave like Gigabit Ethernet interfaces.                           |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Gigabit Ethernet Policers</i></li> <li>• <i>Configuring MAC Address Filtering</i></li> <li>• <i>Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview</i></li> <li>• <i>Configuring Q-in-Q Tunneling (CLI Procedure)</i></li> </ul> |

## ethernet-switching

---

**Syntax** ethernet-switching {  
    filter {  
        group *filter-group-number*;  
        input *filter-name*;  
        input-list [ *filter-names* ];  
        output *filter-name*;  
        output-list [ *filter-names* ];  
    }  
    interface-mode (access | trunk);  
    recovery-timeout *seconds*;  
    storm-control *profile-name*;  
    vlan {  
        members (*vlan-name* | [*-vlan-names*] | all);  
    }  
}

**Hierarchy Level** [edit [interfaces](#) ge-chassis/slot/port [unit](#) *logical-unit-number*] family

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure Ethernet switching protocol family information for the logical interface.  
  
The remaining statements are explained separately.

**Default** You must configure a logical interface to be able to use the physical device.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2826](#)
- [JUNOS Software Network Interfaces Configuration Guide](#)

## ether-options

**Syntax** The **auto-negotiation** and **speed** statements are not supported on the OCX Series.

```
ether-options {
 802.3ad aex {
 lacp {
 force-up;
 (primary | backup);
 }
 }
 (auto-negotiation | no-auto-negotiation);
 configured-flow-control {
 rx-buffers (on | off);
 tx-buffers (on | off);
 }
 (flow-control | no-flow-control);
 link-mode mode;
 (loopback | no-loopback);
 speed (auto-negotiation | no-auto-negotiation);
}
```

**Hierarchy Level** [edit **interfaces** *interface-name*]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Configure **ether-options** properties for a Gigabit Ethernet or 10-Gigabit Ethernet interface.



**NOTE:** The **auto-negotiation** and **speed** statements are not supported on the OCX Series.

The statements are explained separately.

**Default** Enabled.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2826](#)
- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces](#)
- [Junos OS Network Interfaces Library for Routing Devices](#)

## eui-64

---

|                                 |                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | eui-64;                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> unit <i>number</i> family inet6 address <i>address</i> ]                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.3 for EX Series switches.<br>Statement introduced in Junos OS Release 12.2 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| <b>Description</b>              | For interfaces that carry IP version 6 (IPv6) traffic, automatically generate the host number portion of interface addresses.                                                                                                                                              |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the Interface Address on page 2824</a></li></ul>                                                                                                                                                           |

## family

**Syntax** The **ethernet-switching** statement and all of its substatements are not supported on OCX Series switches.

```
family {
 ethernet-switching {
 filter {
 group filter-group-number;
 input filter-name;
 input-list [filter-names];
 output filter-name;
 output-list [filter-names];
 }
 interface-mode (access | trunk);
 recovery-timeout seconds;
 storm-control profile-name;
 vlan {
 members (vlan-name [-vlan-names] | all);
 }
 }
 inet {
 accounting {
 destination-class-usage;
 source-class-usage {
 input;
 output;
 }
 }
 }
 address ipv4-address {
 arp ip-address (mac | multicast-mac) mac-address <publish>;
 broadcast address;
 preferred;
 primary;
 vrrp-group group-number {
 (accept-data | no-accept-data);
 advertise-interval seconds;
 advertisements-threshold number;
 authentication-key key;
 authentication-type authentication;
 fast-interval milliseconds;
 (preempt | no-preempt) {
 hold-time seconds;
 }
 priority number;
 track {
 interface interface-name {
 priority-cost number;
 }
 priority-hold-time seconds;
 route ip-address/mask routing-instance instance-name priority-cost cost;
 }
 virtual-address [addresses];
 vrrp-inherit-from {
```

```

 active-group group-number;
 active-interface interface-name;
 }
}
filter {
 group filter-group-number;
 input filter-name;
 input-list [filter-names];
 output filter-name;
 output-list [filter-names];
}
mtu bytes;
no-neighbor-learn;
no-redirects;
primary;
rpf-check {
 fail-filter filter-name;
 mode {
 loose;
 }
}
}
inet6 {
 accounting {
 destination-class-usage;
 source-class-usage {
 input;
 output;
 }
 }
}
address address {
 eui-64;
 ndp ip-address (mac | multicast-mac) mac-address <publish>;
 preferred;
 primary;
 vrrp-inet6-group group-id {
 accept-data | no-accept-data;
 advertisements-threshold number;
 authentication-key key;
 authentication-type authentication;
 fast-interval milliseconds;
 inet6-advertise-interval milliseconds;
 preempt | no-preempt {
 hold-time seconds;
 }
 priority number;
 track {
 interface interface-name {
 priority-cost number;
 }
 priority-hold-time seconds;
 route ip-address/mask routing-instance instance-name priority-cost cost;
 }
 }
 virtual-inet6-address [addresses];
 virtual-link-local-address ipv6-address;
}

```

```

 vrrp-inherit-from {
 active-group group-name;
 active-interface interface-name;
 }
 }
 (dad-disable | no-dad-disable);
 filter {
 group filter-group-number;
 input filter-name;
 input-list [filter-names];
 output filter-name;
 output-list [filter-names];
 }
 mtu bytes;
 nd6-stale-time time;
 no-neighbor-learn;
 no-redirects;
 policer {
 input policer-name;
 output policer-name;
 }
 rpf-check {
 fail-filter filter-name;
 mode {
 loose;
 }
 }
 mpl {
 filter {
 group filter-group-number;
 input filter-name;
 input-list [filter-names];
 output filter-name;
 output-list [filter-names];
 }
 mtu bytes;
 }
}

```

|                            |                                                                                                                                                                                                                                                             |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Hierarchy Level</b>     | [edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> ],<br>[edit <a href="#">interfaces</a> <a href="#">interface-range</a> <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> family] |
| <b>Release Information</b> | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                               |
| <b>Description</b>         | Configure protocol family information for the logical interface on the QFX Series and OCX Series product.                                                                                                                                                   |

**Default**

**NOTE:** The **ethernet-switching** statement and all of its substatements are not supported on OCX Series switches.

Access interfaces on the QFX Series are set to **family ethernet-switching** by default. If you are going to change the family setting for an interface, you might have to delete this default setting or any user-configured family setting first.

You must configure a logical interface to be able to use the physical device.

**Options**

Interface types on the switch are:

- Aggregated Ethernet (**ae**)
- Gigabit Ethernet (**ge**)
- Loopback (**lo0**)
- Management Ethernet (**me0**)
- Routed VLAN interface (RVI) (**vlan**)



**NOTE:** Routed VLAN interfaces, also referred to as integrated routing and bridging (IRB) interfaces, are not supported on OCX Series switches.

- 10-Gigabit Ethernet (**xe**)

Not all interface types support all **family** substatements. Check your switch CLI for supported substatements for a particular protocol family configuration.

**Required Privilege Level**

interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2826](#)
- *Configuring Gigabit and 10-Gigabit Ethernet Interfaces*
- [Configuring Link Aggregation on page 2869](#)
- [Configuring IRB Interfaces on page 2131](#)
- *Junos OS Network Interfaces Library for Routing Devices*




---

## fibre-channel (Alarm)

---

|                                 |                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | fibre-channel {<br>link-down (red   yellow   ignore);<br>}                                                                                            |
| <b>Hierarchy Level</b>          | [edit chassis <b>alarm</b> ],<br>[edit chassis interconnect-device <i>name</i> <b>alarm</b> ],<br>[edit chassis node-group <i>name</i> <b>alarm</b> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                     |
| <b>Description</b>              | Configure alarms for a Fibre Channel interface.                                                                                                       |
| <b>Options</b>                  | The remaining statement is explained separately.—                                                                                                     |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Understanding Alarms</i></li><li>• <i>Interface Alarm Messages</i></li></ul>                               |


## filter

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <pre>filter {   group <i>filter-group-number</i>;   input <i>filter-name</i>;   input-list [ <i>filter-names</i> ];   output <i>filter-name</i>;   output-list [ <i>filter-names</i> ]; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Hierarchy Level          | <pre>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Release Information      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Description              | <p> <b>NOTE:</b> On EX Series switches, the <code>group</code>, <code>input-list</code>, <code>output-filter</code> statements are not supported under the <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]</code>, <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6]</code>, and <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family mpls]</code> hierarchies.</p> <p>Apply a filter to an interface. You can also use filters for encrypted traffic. When you configure filters, you can configure them under the <b>family ethernet-switching</b>, <b>inet</b>, <b>inet6</b>, <b>mpls</b>, or <b>vpls</b> only.</p> |
| Options                  | <p><b>group <i>filter-group-number</i></b>—Define an interface to be part of a filter group. The default filter group number is 0.</p> <p><b>Range:</b> 0 through 255</p> <p><b>input <i>filter-name</i></b>—Name of one filter to evaluate when packets are received on the interface.</p> <p><b>output <i>filter-name</i></b>—Name of one filter to evaluate when packets are transmitted on the interface.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                               |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Related Documentation    | <ul style="list-style-type: none"> <li>• <i>Applying a Filter to an Interface</i></li> <li>• <i>Junos OS Administration Library for Routing Devices</i></li> <li>• <i>Configuring Gigabit Ethernet Interfaces (CLI Procedure)</i></li> <li>• <i>Configuring Firewall Filters (CLI Procedure)</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

- *family*

## flow-control

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | (flow-control   no-flow-control);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">ether-options</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | <p>Explicitly enable or disable symmetric Ethernet PAUSE flow control, which regulates the flow of packets from the switch to the remote side of the connection by pausing all traffic flows on a link during periods of network congestion. Symmetric flow control means that Ethernet PAUSE is enabled in both directions. The interface generates and sends Ethernet PAUSE messages when the receive buffers fill to a certain threshold and the interface responds to PAUSE messages received from the connected peer. By default, flow control is disabled.</p> <p>You can configure asymmetric flow control by including the <b>configured-flow-control</b> statement at the [edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">ether-options</a> hierarchy level. Symmetric flow control and asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.</p> <div><div></div><div><p><b>NOTE:</b> Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual “lanes” of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC).</p><p>Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.</p><p>OCX Series switches do not support PFC.</p></div></div> <ul style="list-style-type: none"><li>• <b>flow-control</b>—Enable flow control; flow control is useful when the remote device is a Gigabit Ethernet switch.</li><li>• <b>no-flow-control</b>—Disable flow control.</li></ul> |
| <b>Default</b>                  | Flow control is disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">configured-flow-control on page 2980</a></li><li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2826</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

- *Configuring Gigabit and 10-Gigabit Ethernet Interfaces*
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 6493](#)
- *Junos OS Network Interfaces Library for Routing Devices*

## fpc

|                                 |                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>fpc slot {   auto-speed-detection disable;   pic <i>pic-number</i> {     <b>tunnel-port</b> <i>port-number</i> tunnel-services;     port <i>port-number</i> {       channel-speed (<i>speed</i> disable-auto-speed-detection) ;     }     port-range <i>port-range-low port-range-high</i> {       channel-speed (<i>speed</i> disable-auto-speed-detection);     }   } }</pre> |
| <b>Hierarchy Level</b>          | [edit chassis]                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                             |
| <b>Description</b>              | <p>Configure the FPC slot number. For QFX3500 switches, the slot is a line card slot.</p> <p>For generic routing encapsulation (GRE) tunneling, use the <b>tunnel-port</b> statement to specify the port that you want to convert to a GRE tunnel port.</p>                                                                                                                          |
| <b>Options</b>                  | <p><b>slot</b>—Number of the FPC slot. For QFX3500, QFX3600, QFX5200, and OCX Series devices, the slot number is always 0.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show chassis fpc on page 703</a></li> <li>• <i>Configuring Generic Routing Encapsulation Tunneling (CLI Procedure)</i></li> </ul>                                                                                                                                                                                               |

## gratuitous-arp-reply

---

|                                 |                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | (gratuitous-arp-reply   no-gratuitous-arp-reply);                                                                       |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> ],<br>[edit interfaces <i>interface-range</i> <i>interface-range-name</i> ]      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                       |
| <b>Description</b>              | Enable processing of ARP updates received via gratuitous ARP reply messages.                                            |
| <b>Default</b>                  | Updating of the ARP cache is disabled on all Ethernet interfaces.                                                       |
| <b>Options</b>                  | <b>gratuitous-arp-reply</b> —Update the ARP cache.<br><b>no-gratuitous-arp-reply</b> —Do not update the ARP cache.      |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration. |

## hold-time (Physical Interface)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>hold-time up <i>milliseconds</i> down <i>milliseconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>     | [edit <b>interfaces</b> <i>interface-name</i> ],<br>[edit <b>interfaces</b> <i>interface-range</i> <i>interface-range-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 10.4R5 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>         | Specify the <b>hold-time</b> value to use to damp shorter interface transitions milliseconds. The hold timer enables interface damping by not advertising interface transitions until the hold timer duration has passed. When a hold-down timer is configured and the interface goes from up to down, the down hold-time timer is triggered. Every interface transition that occurs during the hold-time is ignored. When the timer expires and the interface state is still down, then the router begins to advertise the interface as being down. Similarly, when a hold-up timer is configured and an interface goes from down to up, the up hold-time timer is triggered. Every interface transition that occurs during the hold-time is ignored. When the timer expires and the interface state is still up, then the router begins to advertise the interface as being up. |



### NOTE:

- We recommend that you configure the hold-time value after determining an appropriate value by performing repeated tests in the actual hardware environment. This is because the appropriate value for hold-time depends on the hardware (XFP, SFP, SR, ER, or LR) used in the networking environment.
- The hold-time option is not available for controller interfaces.



**NOTE:** On MX Series routers with MPC3E and MPC4E, we recommend that you do not configure the hold-down timer to be less than 1 second. On MX Series routers with MPC5EQ-100G10G (MPC5EQ) or MPC6E (MX2K-MPC6E) with 100-Gigabit Ethernet MIC with CFP2 OTN interfaces, we recommend that you do not configure the hold-down timer to be less than 3 seconds.

|                |                                                                                                                                                                                            |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b> | Interface transitions are not damped.                                                                                                                                                      |
| <b>Options</b> | <b>down <i>milliseconds</i></b> —Hold time to use when an interface transitions from up to down. Junos OS advertises the transition within 100 milliseconds of the time value you specify. |

**Range:** 0 through 4,294,967,295

**Default:** 0 (interface transitions are not damped)

**up *milliseconds***—Hold time to use when an interface transitions from down to up. Junos OS advertises the transition within 100 milliseconds of the time value you specify.

**Range:** 0 through 4,294,967,295

**Default:** 0 (interface transitions are not damped)

|                           |                                                               |
|---------------------------|---------------------------------------------------------------|
| <b>Required Privilege</b> | interface—To view this statement in the configuration.        |
| <b>Level</b>              | interface-control—To add this statement to the configuration. |

|                              |                                                                                                                                                                                                                                                                                                                      |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <i>advertise-interval</i></li><li>• <i>interfaces (for EX Series switches)</i></li><li>• <i>Physical Interface Damping Overview</i></li><li>• <i>Damping Shorter Physical Interface Transitions</i></li><li>• <i>Damping Longer Physical Interface Transitions</i></li></ul> |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|



## irb (Interfaces)

```

Syntax irb {
 accounting-profile name;
 description text;

 (gratuitous-arp-reply | no-gratuitous-arp-reply);
 hold-time up milliseconds down milliseconds;
 mtu bytes;
 no-gratuitous-arp-request;

 traceoptions {
 flag flag;
 }
 (traps | no-traps);
 unit logical-unit-number {
 accounting-profile name;
 bandwidth rate;
 description text;
 enhanced-convergence;
 disable;
 encapsulation type;
 family inet {
 accounting {
 destination-class-usage;
 source-class-usage {
 input;
 output;
 }
 }
 }
 address ipv4-address {
 arp ip-address (mac | multicast-mac) mac-address <publish>;
 broadcast address;
 preferred;
 primary;
 vrrp-group group-number {
 (accept-data | no-accept-data);
 advertise-interval seconds;
 advertisements-threshold number;
 authentication-key key;
 authentication-type authentication;
 fast-interval milliseconds;
 (preempt | no-preempt) {
 hold-time seconds;
 }
 priority number;
 track {
 interface interface-name {
 bandwidth-threshold bandwidth;
 priority-cost number;
 }
 priority-hold-time seconds;
 route ip-address/mask routing-instance instance-name priority-cost cost;
 }
 }
 }
 }
 }

```

```
 virtual-address [addresses];
 vrrp-inherit-from {
 active-group group-number;
 active-interface interface-name;
 }
}
filter {
 input filter-name;
 output filter-name;
}
mtu bytes;
no-neighbor-learn;
no-redirects;
primary;
rpf-check {
 fail-filter filter-name;
 mode {
 loose;
 }
}
targeted-broadcast {
 forward-and-send-to-re;
 forward-only;
}
}
family inet6 {
 accounting {
 destination-class-usage;
 source-class-usage {
 input;
 output;
 }
 }
}
address address {
 eui-64;
 ndp ip-address (mac | multicast-mac) mac-address <publish>;
 preferred;
 primary;
 vrrp-inet6-group group-id {
 accept-data | no-accept-data;
 advertisements-threshold number;
 authentication-key key;
 authentication-type authentication;
 fast-interval milliseconds;
 inet6-advertise-interval milliseconds;
 preempt | no-preempt {
 hold-time seconds;
 }
 priority number;
 track {
 interface interface-name {
 bandwidth-threshold bandwidth priority-cost number;
 priority-cost number;
 }
 priority-hold-time seconds;
 }
 }
}
```

```

 route ip-address/mask routing-instance instance-name priority-cost cost;
 }
 virtual-inet6-address [addresses];
 virtual-link-local-address ipv6-address;
 vrrp-inherit-from {
 active-group group-number;
 active-interface interface-name;
 }
}
}
(dad-disable | no-dad-disable);
filter {
 input filter-name;
 output filter-name;
}
mtu bytes;
nd6-stale-time seconds;
no-neighbor-learn;
no-redirects;
policer {
 input policer-name;
 output policer-name;
}
rpf-check {
 fail-filter filter-name;
 mode {
 loose;
 }
}
}
family iso {
 address interface-address;
 mtu bytes;
}
family mpls {
 filter {
 input filter-name;
 output filter-name;
 }
 mtu bytes;
 policer {
 input policer-name;
 output policer-name;
 }
}
native-inner-vlan-id vlan-id;
proxy-arp (restricted | unrestricted);
(traps | no-traps);
vlan-id-list [vlan-id's];
vlan-id-range [vlan-id-range];
}
}

```

Hierarchy Level [edit interfaces *interface-name*

|                                 |                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.3R2 for EX Series switches.<br><b>irb</b> option introduced in Junos OS Release 13.2 for the QFX Series. |
| <b>Description</b>              | Configure the properties of a specific integrated bridging and routing (IRB) interface.<br><br>The remaining statements are explained separately.    |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                              |

---

## inet (interfaces)

---

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>inet {<br/>  address <i>address</i> {<br/>    primary;<br/>    filter input <i>filter-name</i>;<br/>    filter output <i>filter-name</i>;<br/>    targeted-broadcast;<br/>  }<br/>}</pre> |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces interface-name unit logical-unit-number</a> family],<br>[edit <a href="#">interfaces interface-range interface-name unit logical-unit-number</a> family]          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                  |
| <b>Description</b>              | Configure the primary IP address for the logical interface.                                                                                                                                    |
| <b>Default</b>                  | You must configure a logical interface to be able to use the physical device.                                                                                                                  |
| <b>Options</b>                  | The remaining statements are explained separately.—                                                                                                                                            |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2826</a></li></ul>                                                          |

## inet6 (interfaces)

|                                 |                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> inet6 {     address address {         eui-64         preferred         primary;         filter input filter-name;         filter output filter-name;     } } </pre>                                                 |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces interface-name unit logical-unit-number</a> family],<br>[edit <a href="#">interfaces interface-range interface-name unit logical-unit-number</a> family]                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                             |
| <b>Description</b>              | Configure the primary IP address for the logical interface.                                                                                                                                                               |
| <b>Default</b>                  | You must configure a logical interface to be able to use the physical device.                                                                                                                                             |
| <b>Options</b>                  | The remaining statements are explained separately.—                                                                                                                                                                       |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2826</a></li> <li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces</a></li> </ul> |

## interface-mode

|                            |                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>interface-mode (access   trunk &lt;inter-switch-link&gt;);</code>                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>     | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family bridge],<br>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family ethernet-switching],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family bridge]                                                      |
| <b>Release Information</b> | Statement introduced in Junos OS Release 9.2.<br>Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for the QFX Series.<br>Statement introduced in Junos OS Release 15.1.<br><b>inter-switch-link</b> option introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers in enhanced LAN mode. |

### Description



**NOTE:** This statement supports the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see [port-mode](#). For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on page 41.

(QFX Series 3500 and 3600 standalone switches)—Determine whether the logical interface accepts or discards packets based on VLAN tags. Specify the **trunk** option to accept packets with a VLAN ID that matches the list of VLAN IDs specified in the **vlan-id** or **vlan-id-list** statement, then forward the packet within the bridge domain or VLAN configured with the matching VLAN ID. Specify the **access** option to accept packets with no VLAN ID, then forward the packet within the bridge domain or VLAN configured with the VLAN ID that matches the VLAN ID specified in the **vlan-id** statement.



**NOTE:** On MX Series routers, if you want IGMP snooping to be functional for a bridge domain, then you should not configure **interface-mode** and **irb** for that bridge. Such a configuration commit succeeds, but IGMP snooping is not functional, and a message informing the same is displayed. For more information, see *Configuring a Trunk Interface on a Bridge Network*.

|                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b> | <p><b>access</b>—Configure a logical interface to accept untagged packets. Specify the VLAN to which this interface belongs using the <b>vlan-id</b> statement.</p> <p><b>trunk</b>—Configure a single logical interface to accept packets tagged with any VLAN ID specified with the <b>vlan-id</b> or <b>vlan-id-list</b> statement.</p> <p><b>trunk inter-switch-link</b>—For a private VLAN, configure the InterSwitch Link protocol (ISL) on a trunk port of the primary VLAN in order to connect the switches composing the PVLAN to each other. You do not need to configure an ISL when a PVLAN is configured</p> |
|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

on a single switch. This configuration specifies whether the particular interface assumes the role of interswitch link for the PVLAN domains of which it is a member. This option is supported only on MX240, MX480, and MX960 routers in enhanced LAN mode.

|                                 |                                                                       |
|---------------------------------|-----------------------------------------------------------------------|
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.                |
|                                 | interface-control—To add this statement to the configuration.         |
| <b>Related Documentation</b>    | • <i>Configuring Access Mode on a Logical Interface</i>               |
|                                 | • <i>Configuring a Logical Interface for Trunk Mode</i>               |
|                                 | • <i>Example: Connecting Access Switches to a Distribution Switch</i> |
|                                 | • <i>Tunnel Services Overview</i>                                     |
|                                 | • <i>Configuring Tunnel Interfaces on MX Series Routers</i>           |

## interface-range

**Syntax** The `vlan-id` statement is not supported on OCX Series switches.

```
interface-range interface-range-name {
 disable;
 description text;
 ether-options {
 802.3ad aex {
 lacp {
 force-up;
 }
 }
 }
 (auto-negotiation| no-auto-negotiation);
 (flow-control | no-flow-control);
 link-mode mode;
 speed (auto-negotiation | speed);
}
hold-time milliseconds down milliseconds;
member interface-name;
member-range starting-interface-name to ending-interface-name;
mtu bytes;
unit logical-unit-number {
 description text;
 disable;
 family family-name {...}
 (traps | no traps);
 vlan-id vlan-id-number;
}
```

**Hierarchy Level** [edit [interfaces](#)]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description**



**NOTE:** The `vlan-id` statement and Fibre Channel interfaces are not supported on OCX Series switches.



**NOTE:** The interface range definition is supported only for Gigabit Ethernet, 10-Gigabit Ethernet, and Fibre Channel interfaces. Interface ranges are not supported on channelized interfaces.

Group interfaces that share a common configuration profile.

**Options** `interface-range-name`—Name of the interface range.





**NOTE:** You can use regular expressions and wildcards to specify the interfaces in the member range configuration. Do not use wildcards for interface types.

The remaining statements are explained separately.

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Interface Ranges on page 2795</a></li><li>• <a href="#">Interfaces Overview on page 2785</a></li><li>• <i>Interfaces Overview</i></li><li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2826</a></li><li>• <i>Configuring Gigabit and 10-Gigabit Ethernet Interfaces</i></li><li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li></ul> |

## interfaces

**Syntax** The following statements and their associated substatements are not supported on OCX Series switches: **auto-negotiation**, **speed**, **ethernet-switching**, **fcoe-lag**, **fibre-channel**, **fibrechannel-options**, **mc-ae**, **vlan**, **vlan-id**, and **vlan-tagging**.

```

interfaces {
 aex {
 disable;
 aggregated-ether-options {
 configured-flow-control {
 rx-buffers (on | off);
 tx-buffers (on | off);
 }
 (fcoe-lag | no-fcoe-lag);
 (flow-control | no-flow-control);
 lacp mode {
 admin-key key;
 force-up;
 periodic interval;
 system-id mac-address;
 }
 link-speed speed;
 local-bias;
 loopback;
 no-loopback;
 minimum-links number;
 }
 mc-ae {
 chassis-id chassis-id;
 mc-ae-id mc-ae-id;
 mode (active-active);
 status-control (active | standby);
 }
 description text;
 gratuitous-arp-reply | no-gratuitous-arp-reply
 hold-time down milliseconds up milliseconds;
 mtu bytes;
 no-gratuitous-arp-request;
 traceoptions;
 (traps | no traps);
 unit logical-unit-number {
 disable;
 description text;
 family {
 ethernet-switching {
 filter input filter-name;
 filter output filter-name;
 native-vlan-id vlan-id;
 port-mode mode;
 reflective-relay;
 vlan {
 members [(all | names | vlan-ids)];
 }
 }
 }
 }
 }
}

```

```

 }
 inet {
 address address {
 primary;
 }
 filter input filter-name;
 filter output filter-name;
 primary;
 targeted-broadcast;
 }
 (traps | no traps);
 vlan-id vlan-id-number;
}
vlan-tagging;
}
interface-range interface-range-name {
 disable;
 description text;
 ether-options {
 802.3ad aex {
 lacp {
 force-up;
 }
 }
 }
 (auto-negotiation | no-auto-negotiation);
 configured-flow-control {
 rx-buffers (on | off);
 tx-buffers (on | off);
 }
 (flow-control | no-flow-control);
 link-mode mode;
 speed (auto-negotiation | speed);
}
hold-time milliseconds down milliseconds;
member interface-name;
member-range starting-interface-name to ending-interface-name;
mtu bytes;
unit logical-unit-number {
 disable;
 description text;
 family family-name {...}
 (traps | no traps);
 vlan-id vlan-id-number;
}
}
}
lo0 {
 disable;
 description text;
 hold-time milliseconds down milliseconds;
 traceoptions;
 (traps | no traps);
 unit logical-unit-number {
 disable;
 description text;
 family {
 inet {

```

```
 address address {
 primary;
 }
 filter input filter-name;
 filter output filter-name;
 primary;
 targeted-broadcast;
 }
 (traps | no traps);
}
}
mex {
 disable;
 description text;
 hold-time milliseconds down milliseconds;
 (gratuitous-arp-reply | no-gratuitous-arp-reply);
 no-gratuitous-arp-request;
 traceoptions;
 traps;
 unit logical-unit-number {
 disable;
 description text;
 family {
 ethernet-switching {
 filter input filter-name;
 filter output filter-name;
 native-vlan-id vlan-id;
 port-mode mode;
 reflective-relay;
 }
 vlan {
 members [(all | names | vlan-ids)];
 }
 }
 inet {
 address address {
 primary;
 filter input filter-name;
 filter output filter-name;
 primary;
 targeted-broadcast;
 }
 }
 }
 traps;
 vlan-id vlan-id-number;
}
vlan-tagging;
vlan {
 disable;
 description text;
 (gratuitous-arp-reply | no-gratuitous-arp-reply);
 hold-time milliseconds down milliseconds;
 mtu bytes;
 no-gratuitous-arp-request;
 traceoptions;
 (traps | no traps);
 unit logical-unit-number {
```

```

description text;
disable;
family {
 inet {
 address address {
 primary;
 }
 filter input filter-name;
 filter output filter-name;
 primary;
 targeted-broadcast;
 }
 (traps | no traps);
}
}
fc-0/0/port {
 fibrechannel-options {
 bb-sc-n;
 (loopback | no-loopback);
 speed (auto-negotiation | 2g | 4g | 8g);
 }
 unit logical-unit-number {
 disable;
 description text;
 family {
 fibre-channel {
 port-mode np-port;
 }
 }
 (traps | no traps);
 }
}
ge-0/0/port {
 disable;
 description text;
 ether-options {
 802.3ad aex {
 lacp {
 force-up;
 primary;
 }
 }
 }
 (auto-negotiation | no-auto-negotiation);
 configured-flow-control {
 rx-buffers (on | off);
 tx-buffers (on | off);
 }
 (flow-control | no-flow-control);
 link-mode mode;
 loopback;
 no-loopback;
 speed (auto-negotiation | speed);
}
gratuitous-arp-reply| no-gratuitous-arp-reply);
hold-time milliseconds down milliseconds;
mac
mtu bytes;
no-gratuitous-arp-request;

```

```

traceoptions;
(traps | no traps);
unit logical-unit-number {
 description text;
 disable;
 family {
 ethernet-switching {
 filter input filter-name;
 filter output filter-name;
 native-vlan-id vlan-id;
 port-mode mode;
 reflective-relay;
 vlan {
 members [(all | names | vlan-ids)];
 }
 }
 inet {
 address address {
 primary;
 }
 filter input filter-name;
 filter output filter-name;
 primary;
 targeted-broadcast;
 }
 }
 (traps | no traps);
 vlan-id vlan-id-number;
}
vlan-tagging;
}
vrrp-group group-id {
 (accept-data | no-accept-data);
 advertise-interval seconds;
 authentication-key key;
 authentication-type authentication;
 fast-interval milliseconds;
 (preempt | no-preempt) {
 hold-time seconds;
 }
}
priority number;
track {
 interface interface-name {
 bandwidth-threshold bits-per-second priority-cost priority;
 priority-cost priority;
 }
 priority-hold-time seconds;
 route prefix/prefix-length routing-instance instance-name priority-cost priority;
}
}
virtual-address [addresses];
}
xe-0/0/port {
 disable;
 description text;
 ether-options {
 802.3ad aex {

```

```

 lacp {
 force-up;
 (primary | backup);
 }
}
configured-flow-control {
 rx-buffers (on | off);
 tx-buffers (on | off);
}
(flow-control | no-flow-control);
loopback;
no-loopback;
}
(gratuitous-arp-reply | no-gratuitous-arp-reply)
hold-time milliseconds down milliseconds;
mac
mtu bytes;
no-gratuitous-arp-request;
traceoptions;
(traps | no traps);
unit logical-unit-number {
 disable;
 description text;
 family {
 ethernet-switching {
 filter input filter-name;
 filter output filter-name;
 native-vlan-id vlan-id;
 port-mode mode;
 reflective-relay;
 vlan {
 members [(all | names | vlan-ids)];
 }
 }
 fibre-channel {
 port-mode (f-port | np-port);
 }
 inet {
 address address {
 primary;
 }
 filter input filter-name;
 filter output filter-name;
 primary;
 targeted-broadcast;
 }
 (traps | no traps);
 vlan-id vlan-id-number;
 }
 vlan-tagging;
}
}

```

Hierarchy Level [\[edit\]](#)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | <p>Configure the interfaces on the QFX Series and OCX Series.</p> <p>The following statements and their associated substatements are not supported on OCX Series switches: <b>auto-negotiation</b>, <b>ethernet-switching</b>, <b>fcoe-lag</b>, <b>fibre-channel</b>, <b>fibrenchannel-options</b>, <b>mc-ae</b>, <b>speed</b>, <b>vlan</b>, <b>vlan-id</b>, and <b>vlan-tagging</b></p> <p>Most standard Junos OS configuration statements are available in the Junos OS for a switch. This topic lists Junos OS statements that you commonly use when configuring a switch as well as statements added to support switches only.</p> |
| <b>Options</b>                  | <p><b>aex</b>—Configure an aggregated Ethernet interface.</p> <p><b>xe-0/0/</b><i>port</i><b>/</b>—Configure a 10-Gigabit Ethernet interface.</p> <p><b>ge-0/0/</b><i>port</i><b>/</b>—Configure a Gigabit Ethernet interface.</p> <p><b>fc-0/0/</b><i>port</i><b>/</b>—Configure a Fibre Channel interface.</p> <p><b>meX</b>/—Configure a management interface.</p> <p><b>mc-ae</b>—Configure a multichassis aggregated Ethernet (MC-AE) interface.</p> <p>The remaining statements are explained separately.</p>                                                                                                                    |
| <b>Required Privilege Level</b> | <p><b>interface</b>—To view this statement in the configuration.</p> <p><b>interface-control</b>—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Interfaces Overview on page 2785</a></li><li>• <a href="#">Understanding Interface Ranges on page 2795</a></li><li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2826</a></li><li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces</a></li><li>• <a href="#">Configuring Link Aggregation on page 2869</a></li><li>• <a href="#">Configuring a Layer 3 Logical Interface on page 2860</a></li></ul>                                                                                                                             |



## link-down

---


|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | link-down (red   yellow   ignore);                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit chassis <b>alarm ethernet</b> ],<br>[edit chassis <b>alarm fibre-channel</b> ],<br>[edit chassis interconnect-device <i>name</i> <b>alarm ethernet</b> ],<br>[edit chassis node-group <i>name</i> <b>alarm fibre-channel</b> ]                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Specify either red, yellow, or ignore to display when the link is down.                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <p><b>red</b>—Indicates that one or more hardware components have failed or exceeded temperature thresholds, or an alarm condition configured on an interface has triggered a critical warning.</p> <p><b>yellow</b>—Indicates a noncritical condition on the device that, if left unchecked, might cause an interruption in service or degradation in performance. A yellow alarm condition requires monitoring or maintenance.</p> <p><b>ignore</b>—Suppresses or ignores the alarm.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                             |

## link-mode

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>link-mode mode;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">ether-options</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Set the device's link-connection characteristic.                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Default</b>                  | The <b>full-duplex</b> mode is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <p><b>mode</b> —Link characteristic:</p> <ul style="list-style-type: none"><li>• <b>full-duplex</b>—Connection is full duplex.</li><li>• <b>half-duplex</b>—Connection is half duplex.</li><li>• <b>automatic</b>—Link mode is negotiated.</li></ul> <p>If <b>no-auto-negotiation</b> is specified in the <b>ether-options</b> option, you can select only <b>full-duplex</b> or <b>half-duplex</b>. If <b>auto-negotiation</b> is specified in the <b>ether-options</b> option, you can select any mode.</p> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2826</a></li><li>• <i>Configuring Gigabit and 10-Gigabit Ethernet Interfaces</i></li><li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li></ul>                                                                                                                                                                                                                        |

## link-speed

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | link-speed <i>speed</i> ;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit interfaces aex <a href="#">aggregated-ether-options</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | For aggregated Ethernet interfaces only, set the required link speed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <p><b>speed</b>—For aggregated Ethernet links, you can specify the speed in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000).</p> <p>On QFX5100 and EX4600 standalone switches and on a QFX5100 Virtual Chassis and EX4600 Virtual Chassis, you can configure a mixed rate of link speeds for the aggregated Ethernet bundle. Only link speeds of 40G and 10G are supported. Load balancing will not work if you configure link speeds that are not supported.</p> <p>Aggregated Ethernet links on the QFX Series can have one of the following speed values:</p> <div style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;"> <p> <b>NOTE:</b> OCX Series switches only support 10g and 40g interfaces. Mixed rate aggregated Ethernet interfaces are not support on the OCX Series.</p> </div> <ul style="list-style-type: none"> <li>• <b>100g</b>—Links are 100 Gbps.</li> <li>• <b>100m</b>—Links are 100 Mbps.</li> <li>• <b>10g</b>—Links are 10 Gbps.</li> <li>• <b>1g</b>—Links are 1 Gbps.</li> <li>• <b>40g</b>—Links are 40 Gbps.</li> <li>• <b>50g</b>—Links are 50 Gbps.</li> <li>• <b>80g</b>—Links are 80 Gbps.</li> <li>• <b>8g</b>—Links are 8 Gbps.</li> <li>• <b>0c192</b>—Links are OC-192.</li> <li>• <b>mixed</b>—Links are 10 Gbps and 40Gbps.</li> </ul> |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Link Aggregation on page 2869</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

## loopback (Aggregated Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet)

---

|                                 |                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | (loopback   no-loopback);                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> aggregated-ether-options],<br>[edit interfaces <i>interface-name</i> ether-options]                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| <b>Description</b>              | For aggregated Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces, enable or disable loopback mode.                               |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Ethernet Loopback Capability on page 2831</a></li></ul>                       |

## mac

---

|                                 |                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | mac <i>mac-address</i> ;                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> ]                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                                                                                                        |
| <b>Description</b>              | Set the MAC address of the interface. You can configure the MAC address on the management Ethernet interface ( <b>fxp0</b> or <b>em0</b> ) only.                                                                         |
| <b>Options</b>                  | <i>mac-address</i> —MAC address. Specify the MAC address as six hexadecimal bytes in one of the following formats: <i>nnnn.nnnn.nnnn</i> or <i>nn:nn:nn:nn:nn:nn</i> . For example, 0011.2233.4455 or 00:11:22:33:44:55. |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the MAC Address on the Management Ethernet Interface</i></li><li>• <i>Configuring a Pseudowire Subscriber Logical Interface Device</i></li></ul>                  |

## management-ethernet (Alarm)

|                            |                                                                                                                                               |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | management-ethernet {<br>link-down (red   yellow   ignore);<br>}                                                                              |
| <b>Hierarchy Level</b>     | [edit chassis alarm],<br>[edit chassis interconnect-device <i>name</i> alarm],<br>[edit chassis node-group <i>name</i> alarm]                 |
| <b>Release Information</b> | Statement introduced in Junos OS Release 12.2 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| <b>Description</b>         | Configure alarms for a management Ethernet interface.                                                                                         |



**NOTE:** If you configure a yellow alarm on the Interconnect device, it will be handled as a red alarm.

|                                 |                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | The remaining statement is explained separately.—                                                                          |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Understanding Alarms</i></li> <li>• <i>Interface Alarm Messages</i></li> </ul> |

## member

---



|                                 |                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>member <i>interface-name</i>;</code>                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces interface-range</a> <i>interface-range-name</i> ]                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                          |
| <b>Description</b>              | Specify the name of the member interface belonging to an interface range on the QFX Series switch.                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <i>interface-name</i> —Name of the interface.                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2826</a></li><li>• <i>Configuring Gigabit and 10-Gigabit Ethernet Interfaces</i></li><li>• <a href="#">Interfaces Overview on page 2785</a></li><li>• <i>Interfaces Overview</i></li><li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li></ul> |

## member-range

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>member-range <i>starting-interface-name ending-interface-name</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces</a> <i>interface-range interface-range-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Specify the names of the first and last members of a sequence of interfaces belonging to an interface range.                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <i>starting interface-name ending interface-name</i> —Name of the first member and the name of the last member in the interface sequence.                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding Interface Ranges on page 2795</a></li> <li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2826</a></li> <li>• <i>Configuring Gigabit and 10-Gigabit Ethernet Interfaces</i></li> <li>• <a href="#">Interfaces Overview on page 2785</a></li> <li>• <i>Interfaces Overview</i></li> <li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li> </ul> |

## mtu

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>mtu bytes;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>     | [edit <a href="#">interfaces interface-name</a> ],<br>[edit <a href="#">interfaces interface-range interface-name</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b> | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>         | <p>Specify the maximum transmission unit (MTU) size for the media. Changing the media MTU size causes an interface to be deleted and added again. On QFX3500, QFX3600, QFX5100, and OCX Series switches, either standalone or as part of the QFabric system, the maximum MTU value on an untagged packet transiting through an ingress Gigabit Ethernet interface must be no more than the currently configured MTU value plus four, whereas the maximum MTU value on a tagged packet transiting through an ingress Gigabit Ethernet interface must be no more than the currently configured MTU value plus eight. The maximum MTU value on an untagged or tagged packet transiting through an ingress 10-Gigabit Ethernet interface must be no more than the currently configured MTU value plus eight.</p> <p>Keep the following points in mind if you are configuring MTU size for jumbo frames on these special types of interfaces:</p> <ul style="list-style-type: none"> <li>• <b>For LAG interfaces</b>—Configuring the jumbo MTU size on a link aggregation group (LAG) interface (<b>aex</b>) automatically configures the jumbo MTU size on the member links.</li> <li>• <b>For RVIs</b>—Jumbo frames of up to 9216 bytes are supported on the routed VLAN interface (RVI), which is named <b>vlan</b>. The RVI functions as a logical router. To route jumbo data packets on the RVI, you must configure the jumbo MTU size on the member physical interfaces of the RVI and not on the RVI itself (the <b>vlan</b> interface). However, for jumbo control packets—for example, to ping the RVI with a packet size of 6000 bytes or more—you must explicitly configure the jumbo MTU size on the interface named <b>vlan</b> (the RVI). On a QFX5100 switch jumbo frames on the RVI are configured on the basis of the interface MTU.</li> </ul> <div style="margin-top: 20px;"> <div style="display: flex; align-items: center;">  <div> <p><b>NOTE:</b> RVIs are not supported on OCX Series switches.</p> </div> </div> <div style="margin-top: 20px;"> <div style="display: flex; align-items: center;">  <div> <p><b>CAUTION:</b> Setting or deleting the jumbo MTU size on the RVI (the <b>vlan</b> interface) while the switch is transmitting packets might result in dropped packets.</p> </div> </div> </div> </div> |
| <b>Options</b>             | <p><b>bytes</b> —MTU size.</p> <p><b>Range:</b> 64 through 9216 bytes</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |



**Default:** 1514 bytes

|                                 |                                                                                                                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2826</a></li> <li>• <i>Configuring Gigabit and 10-Gigabit Ethernet Interfaces</i></li> <li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li> </ul> |

## no-gratuitous-arp-request

---

|                                 |                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-gratuitous-arp-request;                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces interface-name</a> ],<br>[edit <a href="#">interfaces interface-range interface-name</a> ]                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                     |
| <b>Description</b>              | Configure the switch not to respond to gratuitous ARP requests. You can disable responses to gratuitous ARP requests on both Layer 2 Ethernet switching interfaces and routed VLAN interfaces (RVIs). |
| <b>Default</b>                  | Gratuitous ARP responses are enabled on all Ethernet switching interfaces and RVIs.                                                                                                                   |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring IRB Interfaces on page 2131</a></li> </ul>                                                                                           |

## pic

**Syntax** `pic pic-number {  
     tunnel-port port-number tunnel-services;  
     port port-number {  
         channel-speed (speed|disable-auto-speed-detection) ;  
     }  
     port-range port-range-low port-range-high {  
         channel-speed (speed|disable-auto-speed-detection) ;  
     }  
 }`

**Hierarchy Level** [edit chassis fpc *slot*]

**Release Information** Option **channel-speed** introduced in Junos OS Release 13.2 for the QFX Series.



**NOTE:** This statement is not supported on the OCX Series.

**Description** (QFX3500, QFX3600, and QFX5100 standalone switches running Enhanced Layer 2 Software only)—Configure a specific port or a range of ports to operate as 10-Gigabit Ethernet ports or 40-Gigabit Ethernet ports.

**Options** **pic *pic-number***—(QFX3500 standalone switch only) Number of the physical interface card (PIC) on which you want to configure port types. Specify **1** to configure 10-Gigabit Ethernet or 40-Gigabit Ethernet type ports.  
 (QFX3600 standalone switch only) Number of the physical interface card (PIC) on which you want to configure port types. Specify **0** to configure 10-Gigabit Ethernet or 40-Gigabit Ethernet type ports.

**port *physical-port-number***—Port number on which you want to configure the port type.

**port-range-low**—Lowest-numbered port in the range of ports.


**port-range-high**—Highest-numbered port in the range of ports.

**channel-speed (*speed* |disable-auto-speed-detection)** —Configure *10g* for 10-Gigabit Ethernet type ports, and configure *disable-auto-speed-detection* to disable auto-channelization.

**Required Privilege Level** interface—To view this statement in the configuration.  
 interface-control—To add this statement to the configuration.

**Related Documentation** • [Channelizing Interfaces on page 2835](#)

## rx-buffers

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | rx-buffers (on   off);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | [edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">ether-options</a> <a href="#">configured-flow-control</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <p>Enable or disable an interface to generate and send Ethernet PAUSE messages. If you enable the receive buffers to generate and send PAUSE messages, when the receive buffers reach a certain level of fullness, the interface sends a PAUSE message to the connected peer. If the connected peer is properly configured, it stops transmitting frames to the interface on the entire link. When the interface receive buffer empties below a certain threshold, the interface sends a message to the connected peer to resume sending frames.</p> <p>Ethernet PAUSE prevents buffers from overflowing and dropping packets during periods of network congestion. If the other devices in the network are also configured to support PAUSE, PAUSE supports lossless operation. Use the <b>rx-buffers</b> statement with the <b>tx-buffers</b> statement to configure asymmetric Ethernet PAUSE on an interface. (Use the <b>flow-control</b> statement to enable symmetric PAUSE and the <b>no-flow-control</b> statement to disable symmetric PAUSE on an interface. Symmetric flow control and asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.)</p> |
| <div>  <p><b>NOTE:</b> Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual “lanes” of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC).</p> <p>Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.</p> </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Default</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Flow control is disabled. You must explicitly configure Ethernet PAUSE flow control on interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>on   off</b> —Enable or disable an interface to generate and send Ethernet PAUSE messages.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <ul style="list-style-type: none"> <li>• <a href="#">flow-control on page 2996</a></li> <li>• <a href="#">tx-buffers on page 3033</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |


- *Configuring CoS Asymmetric Ethernet PAUSE Flow Control*
- *Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control*
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 6493](#)

---

## source

|                                 |                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>source source-address;</code>                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> <a href="#">tunnel</a> ]                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| <b>Description</b>              | Specify the source address of the tunnel.                                                                                                                                                                                                                                   |
| <b>Default</b>                  | If you do not specify a source address, the tunnel uses the unit's primary address as the source address of the tunnel.                                                                                                                                                     |
| <b>Options</b>                  | <b><i>source-address</i></b> —Address of the local side of the tunnel. This is the address that is placed in the outer IP header's source field.                                                                                                                            |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Tunnel Services Overview</i></li></ul>                                                                                                                                                                                           |

## speed

|                                                                                                                                                                               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                 | speed (10g   1g   100m)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>                                                                                                                                                        | [edit <a href="#">interfaces</a> <i>interface-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>                                                                                                                                                    | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>                                                                                                                                                            | Configure the speed of the interface. On QFX5100 devices using 1-Gigabit Ethernet Copper SFP interfaces, you can configure the speed to be 100 Mbps. To return to the default speed of 1 Gbps on the SFP port, delete the <b>100m</b> statement at the <b>[edit interfaces interface-name speed]</b> CLI hierarchy. Also, on QFX5100-48T devices using 10-Gigabit Ethernet interfaces, you can configure the speed to be 1 Gbps or 100 Mbps. To return to the default speed of 10 Gbps for an interface, simply delete all speed configurations, or set the speed explicitly to 10G at the <b>[edit interfaces interface-name speed]</b> CLI hierarchy. |
| <div>  <b>NOTE:</b> Only 10g and 40g interfaces are supported on OCX Series switches. </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Default</b>                                                                                                                                                                | The speed for 1-Gigabit Ethernet Copper SFP interfaces is set to 1 Gbps by default, but you can configure the speed to be 100 Mbps. The speed for 10-Gigabit Ethernet interfaces is set to 10 Gbps by default and cannot be configured to operate in a different speed.                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                                                                                                                                                                | <ul style="list-style-type: none"> <li>• <b>10g</b>—10 Gbps</li> <li>• <b>1g</b>—1 Gbps</li> <li>• <b>100m</b>—100 Mbps</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b>                                                                                                                                               | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>                                                                                                                                                  | <ul style="list-style-type: none"> <li>• <a href="#">auto-negotiation on page 2978</a></li> <li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2826</a></li> <li>• <i>Configuring Gigabit and 10-Gigabit Ethernet Interfaces</i></li> <li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li> </ul>                                                                                                                                                                                                                                                                                                     |

## tag-protocol-id (TPIDs Expected to Be Sent or Received)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>tag-protocol-id [<i>tpids</i>];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> gigether-options <a href="#">ethernet-switch-profile</a> ],<br>[edit interfaces <i>interface-name</i> aggregated-ether-options <a href="#">ethernet-switch-profile</a> ],<br>[edit interfaces <i>interface-name</i> aggregated-ether-options <a href="#">ethernet-switch-profile</a> ],<br>[edit interfaces <i>interface-name</i> ether-options <a href="#">ethernet-switch-profile</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.<br>Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.<br>Statement introduced in Junos OS Release 14.1X53-D15 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | <p>For Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, aggregated Ethernet with Gigabit Ethernet IQ interfaces, and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC, and the built-in Gigabit Ethernet port on the M7i router), define the TPIDs expected to be sent or received on a particular VLAN. For each Gigabit Ethernet port, you can configure up to eight TPIDs using the <b>tag-protocol-id</b> statement; but only the first four TPIDs are supported on IQ2 and IQ2-E interfaces.</p> <p>For 10-Gigabit Ethernet LAN/WAN PIC interfaces on T Series routers only the default TPID value (<b>0x8100</b>) is supported.</p> <p>For Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces on EX Series switches, define the TPIDs expected to be sent or received on a particular VLAN. The default TPID value is <b>0x8100</b>. Other supported values are <b>0x88a8</b>, <b>0x9100</b>, and <b>0x9200</b>.</p> |
| <b>Options</b>                  | <i>tpids</i> —TPIDs to be accepted on the VLAN. Specify TPIDs in hexadecimal.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li><i>Configuring Frames with Particular TPIDs to Be Processed as Tagged Frames</i></li><li><i>Configuring Q-in-Q Tunneling (CLI Procedure)</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |


## targeted-broadcast

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | targeted-broadcast;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> family inet],<br>[edit <a href="#">interfaces</a> <a href="#">interface-range</a> <i>interface-range-name</i> <a href="#">unit</a> <i>logical-unit-number</i> family inet]                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | Specify whether the IP packets destined for a Layer 3 broadcast need to be forwarded to both an egress interface and the Routing Engine, or to an egress interface only. The packets are broadcast only if the egress interface is a LAN interface.                                                                                                                                                                                                                                                  |
| <b>Default</b>                  | When this statement is not included, broadcast packets are sent to the Routing Engine only.                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring IP Directed Broadcast on a Switch</i></li> <li>• <i>Configuring IP Directed Broadcast (CLI Procedure)</i></li> <li>• <i>Understanding IP Directed Broadcast</i></li> <li>• <a href="#">Understanding IP Directed Broadcast on page 2849</a></li> <li>• <a href="#">Configuring IP Directed Broadcast (CLI Procedure) on page 2851</a></li> <li>• <a href="#">Example: Configuring IP Directed Broadcast on page 2852</a></li> </ul> |


## traceoptions (Individual Interfaces)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>traceoptions {<br/>    flag <i>flag</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | <p>Define tracing operations for individual interfaces.</p> <p>To specify more than one tracing operation, include multiple <b>flag</b> statements.</p> <p>The <b>traceoptions</b> statement for interfaces does not support a trace file. The logging is done by the kernel, so the tracing information is placed in the system <b>syslog</b> file in the directory <b>/var/log</b>.</p> <div> <b>NOTE:</b> The <b>traceoptions</b> statement is not supported on the QFX3000 QFabric system.</div> |
| <b>Default</b>                  | If you do not include this statement, no interface-specific tracing operations are performed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <p><b>flag</b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements. The following are the interface-specific tracing options.</p> <ul style="list-style-type: none"><li>• <b>all</b>—All interface tracing operations</li><li>• <b>event</b>—Interface events</li><li>• <b>ipc</b>—Interface interprocess communication (IPC) messages</li><li>• <b>media</b>—Interface media changes</li><li>• <b>q921</b>—ISDN Q.921 frames</li><li>• <b>q931</b>—ISDN Q.931 frames</li></ul>                                         |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Tracing Operations of an Individual Router or Switch Interface</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                               |



## tx-buffers

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | tx-buffers (on   off);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">ether-options</a> <a href="#">configured-flow-control</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | <p>Enable or disable an interface to respond to received Ethernet PAUSE messages. If you enable the transmit buffers to respond to PAUSE messages, when the interface receives a PAUSE message from the connected peer, the interface stops transmitting frames on the entire link. When the receive buffer on the connected peer empties below a certain threshold, the peer interface sends a message to the paused interface to resume sending frames.</p> <p>Ethernet PAUSE prevents buffers from overflowing and dropping packets during periods of network congestion. If the other devices in the network are also configured to support PAUSE, PAUSE supports lossless operation. Use the <b>tx-buffers</b> statement with the <b>rx-buffers</b> statement to configure asymmetric Ethernet PAUSE on an interface. (Use the <b>flow-control</b> statement to enable symmetric PAUSE and the <b>no-flow-control</b> statement to disable symmetric PAUSE on an interface. Symmetric flow control and asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.)</p> |
|                                 | <div>  <p><b>NOTE:</b> Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual “lanes” of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC).</p> <p>Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.</p> </div>                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Default</b>                  | Flow control is disabled. You must explicitly configure Ethernet PAUSE flow control on interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | on   off—Enable or disable an interface to respond to an Ethernet PAUSE message.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">flow-control on page 2996</a></li> <li>• <a href="#">rx-buffers on page 3027</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

- *Configuring CoS Asymmetric Ethernet PAUSE Flow Control*
- *Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control*
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 6493](#)

## unit

**Syntax** The **ethernet-switching** and **fibre-channel** statements and all of their substatements are not supported on OCX Series switches.

```
unit logical-unit-number {
 family {
 ethernet-switching {
 filter input filter-name;
 filter output filter-name;
 native-vlan-id vlan-id;
 port-mode mode;
 vlan {
 members [(all | names | vlan-ids)];
 }
 }
 fibre-channel {
 port-mode (f-port | np-port);
 }
 inet {
 address address {
 primary;
 }
 filter input filter-name;
 filter output filter-name;
 primary;
 targeted-broadcast;
 }
 }
}
```

**Hierarchy Level** [edit **interfaces** *interface-name*],  
[edit **interfaces** *interface-range* *interface-range-name*]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description**



**NOTE:** The **ethernet-switching** and **fibre-channel** statements and all of their substatements are not supported on OCX Series switches.

Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

**Default** You must configure a logical interface to be able to use the physical device.

**Options** *logical-unit-number*—Number of the logical unit.

**Range:** 0 through 16,384


The remaining statements are explained separately.

|                                 |                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2826</a></li><li>• <i>Configuring Gigabit and 10-Gigabit Ethernet Interfaces</i></li><li>• <a href="#">Configuring Link Aggregation on page 2869</a></li><li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li></ul> |

---

## vlan-id

---

|                                                                                                                                                                                                                           |                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                             | vlan-id <i>vlan-id-number</i> ;                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>                                                                                                                                                                                                    | [edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> ]                                                                                                                                                                                                                                          |
| <b>Release Information</b>                                                                                                                                                                                                | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                 |
| <b>Description</b>                                                                                                                                                                                                        | For 10-Gigabit Ethernet and aggregated Ethernet interfaces only, bind an 802.1Q VLAN tag ID to a logical interface.                                                                                                                                                                                                                               |
| <div> <b>NOTE:</b> The VLAN tag ID cannot be configured on logical interface unit 0. The logical unit number must be 1 or higher.</div> |                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                                                                                                                                                                                                            | <i>vlan-id-number</i> —Valid VLAN identifier.<br><b>Range:</b> 1 through 4094                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b>                                                                                                                                                                                           | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                           |
| <b>Related Documentation</b>                                                                                                                                                                                              | <ul style="list-style-type: none"><li>• <a href="#">vlan-tagging on page 2309</a></li><li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2826</a></li><li>• <a href="#">Configuring a Layer 3 Logical Interface on page 2860</a></li><li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li></ul> |

## vlan-tagging

---

|                                 |                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | vlan-tagging;                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces</a> <i>interface-name</i> ]<br>[edit <a href="#">interfaces</a> <a href="#">interface-range</a> <i>interface-range-name</i> ]            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                     |
| <b>Description</b>              | Enable VLAN tagging. The platform receives and forwards single-tag frames with 802.1Q VLAN tags.                                                                      |
| <b>Default</b>                  | VLAN tagging is disabled by default.                                                                                                                                  |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">vlan-id on page 3036</a></li><li>• <a href="#">Configuring a Layer 3 Logical Interface on page 2860</a></li></ul> |



# IP Directed Broadcast Configuration Statement

- [targeted-broadcast on page 3039](#)

## targeted-broadcast

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | targeted-broadcast;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> family inet],<br>[edit <a href="#">interfaces</a> <a href="#">interface-range</a> <i>interface-range-name</i> <a href="#">unit</a> <i>logical-unit-number</i> family inet]                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | Specify whether the IP packets destined for a Layer 3 broadcast need to be forwarded to both an egress interface and the Routing Engine, or to an egress interface only. The packets are broadcast only if the egress interface is a LAN interface.                                                                                                                                                                                                                                                  |
| <b>Default</b>                  | When this statement is not included, broadcast packets are sent to the Routing Engine only.                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring IP Directed Broadcast on a Switch</i></li> <li>• <i>Configuring IP Directed Broadcast (CLI Procedure)</i></li> <li>• <i>Understanding IP Directed Broadcast</i></li> <li>• <a href="#">Understanding IP Directed Broadcast on page 2849</a></li> <li>• <a href="#">Configuring IP Directed Broadcast (CLI Procedure) on page 2851</a></li> <li>• <a href="#">Example: Configuring IP Directed Broadcast on page 2852</a></li> </ul> |





## CHAPTER 126

# LAGs and LACP Configuration Statements

- [aggregated-devices on page 3042](#)
- [aggregated-ether-options on page 3043](#)
- [chassis on page 3045](#)
- [802.3ad on page 3046](#)
- [device-count on page 3047](#)
- [ethernet on page 3047](#)
- [force-up on page 3048](#)
- [lacp \(802.3ad\) on page 3049](#)
- [lacp \(Aggregated Ethernet\) on page 3050](#)
- [link-protection on page 3051](#)
- [periodic on page 3052](#)

## aggregated-devices

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>aggregated-devices {<br/>    ethernet {<br/>        device-count <i>number</i>;<br/>    }<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit <a href="#">chassis</a> ],                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.<br>Statement introduced in Junos OS Release 14.2R3                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Configure properties for aggregated devices on the switch.<br><br>The remaining statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Default</b>                  | Aggregated devices are disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Understanding Link Aggregation and Link Aggregation Control Protocol in a Junos Fusion</i></li><li>• <a href="#">Understanding Aggregated Ethernet Interfaces and LACP on page 2865</a></li><li>• <a href="#">Configuring Link Aggregation on page 2869</a></li><li>• <i>Configuring Link Aggregation</i></li><li>• <a href="#">Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2876</a></li><li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li></ul> |

## aggregated-ether-options

**Syntax** The **fcoe-lag** and **mc-ae** statements are not supported on OCX Series switches.

```
aggregated-ether-options {
 configured-flow-control {
 rx-buffers (on | off);
 tx-buffers (on | off);
 }
 ethernet-switch-profile {
 tag-protocol-id;
 (fcoe-lag | no-fcoe-lag);
 (flow-control | no-flow-control);
 lacp mode {
 admin-key key;
 periodic interval;
 system-id mac-address;
 force-up;
 }
 }
 (link-protection | no-link-protection);
 link-speed speed;
 local-bias;
 (loopback | no-loopback);
 mc-ae {
 chassis-id chassis-id;
 mc-ae-id mc-ae-id;
 mode (active-active);
 status-control (active | standby);
 }
 minimum-links number;
 rebalance-periodic;
 resilient-hash;
 source-address-filter filter;
 (source-filtering | no-source-filtering);
}
```

**Hierarchy Level** [edit [interfaces](#) aex]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.  
Statements **fcoe-lag** and **no-fcoe-lag** introduced in Junos OS Release 13.2X52-D10 for the QFX Series.  
Statements **force-up**, **lacp**, and **resilient-hash** introduced in Junos OS Release 14.1X53-D10 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Configure properties specific to a specific aggregated Ethernet interface.



**NOTE:** The **fcoe-lag** and **mc-ae** statements are not supported on OCX Series switches.



NOTE: The **force-up** statement is not supported on QFX10002 switches.



NOTE: The **resilient-hash** statement is not supported on QFX10002 switches.

The statements are explained separately.

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>                  | Options are not enabled.                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding Aggregated Ethernet Interfaces and LACP on page 2865</a></li> <li>• <a href="#">Configuring Aggregated Ethernet LACP on page 2868</a></li> <li>• <a href="#">Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch on page 2880</a></li> <li>• <a href="#">Junos OS Network Interfaces Library for Routing Devices</a></li> </ul> |

## chassis

```
Syntax chassis {
 routing-engine
 redundancy {
 failover {
 on-disk-failure {
 disk-failure-action (halt | reboot);
 }
 on-loss-of-keepalives;
 }
 graceful-switchover;
 }
 aggregated-devices {
 ethernet {
 device-count number;
 }
 alarm {
 interface-type {
 alarm-name (red | yellow | ignore);
 }
 }
 }
 forwarding-options profile-name {
 num-65-127-prefix value
 }
 fpc slot {
 auto-speed-detection disable
 pic pic-number{
 port port-number{
 tunnel-port port-number tunnel-services;
 channel-speed speed;
 }
 port-range port-range-low port-range-high {
 channel-speed speed;
 }
 }
 }
 maximum-ecmp next-hops;
 }
```

Hierarchy Level [edit]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.  
Statement introduced in Junos OS Release 14.2R3

**Description** Configure chassis-specific properties for the switch.  
  
The remaining statements are explained separately.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

- Related Documentation**
- [Understanding Link Aggregation and Link Aggregation Control Protocol in a Junos Fusion](#)
  - [Configuring Link Aggregation](#)
  - [Configuring Link Aggregation on page 2869](#)

---

## 802.3ad

**Syntax**

```
802.3ad aex;
 lacp {
 force-up;
 (primary | backup);
 }
 port-priority;
}
```

**Hierarchy Level** [edit [interfaces](#) *interface-name* [ether-options](#)]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Specify the aggregated Ethernet logical interface number.



**NOTE:** The port-priority statement is not supported on QFabric systems.

---



**NOTE:** The force-up statement is not supported on QFX10002 switches.

---

**Options** aex—Aggregated Ethernet logical interface number.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

- Related Documentation**
- [Configuring Link Aggregation on page 2869](#)
  - [Configuring Aggregated Ethernet LACP on page 2868](#)
  - [Understanding Aggregated Ethernet Interfaces and LACP on page 2865](#)
  - [Troubleshooting an Aggregated Ethernet Interface on page 2886](#)
  - [Junos OS Network Interfaces Library for Routing Devices](#)

## device-count

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>device-count <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit <a href="#">chassis aggregated-devices ethernet</a> ],<br>[edit <a href="#">chassis node-group <i>name</i> aggregated-devices ethernet</a> ]                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.<br>Statement introduced in Junos OS Release 14.2R3                                                                                                                                                                                                           |
| <b>Description</b>              | Configure the number of aggregated Ethernet logical devices available to the switch.                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Understanding Link Aggregation and Link Aggregation Control Protocol in a Junos Fusion</i></li> <li>• <i>Configuring Link Aggregation</i></li> <li>• <a href="#">Configuring Link Aggregation on page 2869</a></li> <li>• <a href="#">Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2876</a></li> </ul> |

## ethernet

|                                 |                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>ethernet {<br/>    <a href="#">device-count <i>number</i></a>;<br/>}</code>                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit <a href="#">chassis aggregated-devices</a> ],<br>[edit <a href="#">chassis node-group aggregated-devices</a> ]                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.<br>Statement introduced in Junos OS Release 14.2R3                                                                                                                                                |
| <b>Description</b>              | Configure properties for aggregated Ethernet devices on the switch.<br><br>The remaining statement is explained separately.                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Understanding Link Aggregation and Link Aggregation Control Protocol in a Junos Fusion</i></li> <li>• <i>Configuring Link Aggregation</i></li> <li>• <a href="#">Configuring Link Aggregation on page 2869</a></li> <li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li> </ul> |

## force-up

---

|                            |                                                                                                                                                                                                             |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | force-up;                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>     | [edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">ether-options</a> 802.3ad lacp;<br>[edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">aggregated-ether-options</a> lacp; |
| <b>Release Information</b> | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                               |
| <b>Description</b>         | Configure the state of the interface as up when the peer has limited LACP capability. You can also configure the peer interface (in MC-LAG) to remain up even with limited LACP capability.                 |





**NOTE:** The force-up option is not supported on QFX10002 switches.

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Link Aggregation</a></li><li>• <a href="#">Understanding Aggregated Ethernet Interfaces and LACP on page 2865</a></li><li>• <a href="#">Configuring Aggregated Ethernet LACP on page 2868</a></li><li>• <a href="#">Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch on page 2880</a></li><li>• <a href="#">Junos OS Network Interfaces Library for Routing Devices</a></li><li>• </li></ul> |



## lacp (802.3ad)

|                                                                                                                                                                                     |                                                                                                                                                                                                                                                                                                                                                  |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                       | <pre>lacp {     force-up;     (primary   backup);     port-priority; }</pre>                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>                                                                                                                                                              | [edit <a href="#">interfaces interface-name ether-options 802.3ad</a> ]                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>                                                                                                                                                          | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                    |
| <b>Description</b>                                                                                                                                                                  | Configure the Link Aggregation Control Protocol (LACP) parameters for interfaces. The remaining statement is explained separately.                                                                                                                                                                                                               |
| <div>  <p><b>NOTE:</b> The port-priority statement is not supported on QFabric systems.</p> </div> |                                                                                                                                                                                                                                                                                                                                                  |
| <div>  <p><b>NOTE:</b> The force-up statement is not supported on QFX10002 switches.</p> </div>   |                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b>                                                                                                                                                     | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                          |
| <b>Related Documentation</b>                                                                                                                                                        | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Link Aggregation</a></li> <li>• <a href="#">Configuring Link Aggregation on page 2869</a></li> <li>• <a href="#">Configuring Aggregated Ethernet LACP on page 2868</a></li> <li>• <a href="#">Understanding Aggregated Ethernet Interfaces and LACP on page 2865</a></li> </ul> |

## lacp (Aggregated Ethernet)

---

**Syntax**    lacp (active | passive) {  
              admin-key *key*;  
              fast-failover;  
              link-protection {  
                  disable;  
                  (revertive | non-revertive);  
              }  
              periodic *interval*  
              system-ID *mac-address*;  
              system-priority *priority*;  
              force-up;  
              }

**Hierarchy Level**    [edit [interfaces interface-name](#)[aggregated-ether-options](#)]

**Release Information**    Statement introduced in Junos OS Release 11.1 for the QFX Series.  
                              Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description**    Configure the Link Aggregation Control Protocol (LACP) parameters for interfaces. The remaining statement is explained separately.



**NOTE:** The force-up statement is not supported on QFX10002 switches.

---

**Required Privilege Level**    interface—To view this statement in the configuration.  
                                  interface-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring Link Aggregation](#)
- [Configuring Link Aggregation on page 2869](#)
- [Configuring Aggregated Ethernet LACP on page 2868](#)
- [Configuring LACP Link Protection of Aggregated Ethernet Interfaces \(CLI Procedure\) on page 2872](#)
- [Understanding Aggregated Ethernet Interfaces and LACP on page 2865](#)

## link-protection

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>link-protection {   disable;   (revertive  non-revertive); }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | <p>[edit interfaces aex aggregated-ether-options]<br/> [edit interfaces aex aggregated-ether-options <i>lACP</i>]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.3.<br/> Statement introduced in Junos OS Release 9.0 for EX Series switches.<br/> Support for <b>disable</b>, <b>revertive</b>, and <b>non-revertive</b> statements added in Junos OS Release 9.3.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | <p>On the router, for aggregated Ethernet interfaces only, configure link protection. In addition to enabling link protection, a primary and a secondary (backup) link must be configured to specify what links egress traffic should traverse. To configure primary and secondary links on the router, include the <b>primary</b> and <b>backup</b> statements at the [edit interfaces <i>ge-fpc/pic/port</i> <b>gigether-options 802.3ad aex</b>] hierarchy level or the [edit interfaces <i>fe-fpc/pic/port</i> <b>fastether-options 802.3ad aex</b>] hierarchy level.</p> <p>On the switch, you can configure either Junos OS link protection for aggregated Ethernet interfaces or the LACP standards link protection for aggregated Ethernet interfaces.</p> <p>For Junos OS link protection, specify <b>link-protection</b> at the following hierarchy levels:</p> <ul style="list-style-type: none"> <li>• [edit interfaces <i>ge-fpc/pic/port</i> <b>ether-options 802.3ad aex</b>]</li> <li>• [edit interfaces <i>xe-fpc/pic/port</i> <b>ether-options 802.3ad aex</b>] hierarchy level or at the [edit interfaces <i>xe-fpc/pic/port</i> <b>ether-options 802.3ad aex</b>] hierarchy level.</li> </ul> <p>To disable link protection, use the <b>delete interface ae aggregate-ether-options link-protection</b> statement at the [edit interfaces aex aggregated-ether-options] hierarchy level or the [edit interfaces aex aggregated-ether-options <i>lACP</i>] hierarchy level.</p> |
| <b>Options</b>                  | The statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.<br/> interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Aggregated Ethernet Link Protection</i></li> <li>• <i>Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure)</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## periodic

---

|                                 |                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>periodic (fast   slow);</code>                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces aex aggregated-ether-options lacp</a> ]                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                           |
| <b>Description</b>              | Configure the interval for periodic transmission of LACP packets.                                                                                                                                                                                                                                                                       |
| <b>Default</b>                  | <code>fast</code>                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <b><i>interval</i></b> —Interval at which to periodically transmit LACP packets: <ul style="list-style-type: none"><li>• <b>fast</b>—Receive packets every second. This is the default.</li><li>• <b>slow</b>—Receive packets every 30 seconds.</li></ul>                                                                               |
| <b>Required Privilege Level</b> | <code>interface</code> —To view this statement in the configuration.<br><code>interface-control</code> —To add this statement to the configuration.                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Link Aggregation</i></li><li>• <a href="#">Configuring Aggregated Ethernet LACP on page 2868</a></li><li>• <a href="#">Understanding Aggregated Ethernet Interfaces and LACP on page 2865</a></li><li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li></ul> |

## CHAPTER 127

# Load Balancing Configuration Statements

- [hash-key \(Forwarding Options\) on page 3054](#)
- [multiservice on page 3056](#)

## hash-key (Forwarding Options)

```
Syntax hash-key {
 family {
 inet {
 layer-3;
 layer-4;
 inner-vlan-id;
 outer-vlan-id;
 }
 }
 multiservice {
 source-mac;
 destination-mac;
 payload {
 ip {
 layer3-only;
 layer-3 (source-ip-only | destination-ip-only);
 layer-4;
 inner-vlan-id;
 outer-vlan-id;
 }
 }
 }
 }
```

**Hierarchy Level** [edit forwarding-options]  
[edit chassis fpc slot-number pic pic-number]

**Release Information** Statement introduced in Junos OS Release 15.1X53-D10 for QFX10000 switches.

**Description** (QFX10000 switches only) Select which packet header data to use for per-flow load balancing.



**NOTE:** You can configure either Layer 3 or Layer 4 load balancing, or both at the same time.



**NOTE:** On I chip platforms, an unknown Layer 4 header is excluded from load-balance hashing to avoid undesired packet reordering.

- Options**
- **inet**—IPv4 address family.
  - **layer-3**—Incorporate Layer 3 data into the hash key.
  - **layer-4**—Incorporate Layer 4 data into the hash key.
  - **outer-vlan-id**—Include outer VLAN ID information in the hash key.

- **inner-vlan-id**—Include inner VLAN ID information in the hash key.
- **payload**—Incorporate payload data into the hash key.
- **ip**—Include the IP address of the IPv4 or IPv6 payload into the hash key.
- **layer-3-only**—Include only Layer 3 IP information.

|                           |                                                            |
|---------------------------|------------------------------------------------------------|
| <b>Required Privilege</b> | system—To view this statement in the configuration.        |
| <b>Level</b>              | system-control—To add this statement to the configuration. |

|                      |   |
|----------------------|---|
| <b>Related</b>       | • |
| <b>Documentation</b> |   |

## multiservice

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>multiservice {   source-mac;   destination-mac;   payload {     ip {       layer3-only;       layer-3 (source-ip-only   destination-ip-only);       layer-4;       inner-vlan-id;       outer-vlan-id;     }   } }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit chassis fpc <i>slot-number</i> pic <i>pic-number</i> hash-key family]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 15.1X53-D10 for the QFX10000 switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | (QFX10000 switches only) Configure data used in a hash key for the <b>multiservice</b> protocol family when configuring PIC-level hashing for load balancing on an 802.3ad Link Aggregation Group.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <p><b>destination-mac</b>—Include destination MAC address in the hash key.</p> <p><b>payload</b>—Include payload data in the hash key. This option has the following suboptions:</p> <ul style="list-style-type: none"><li>• <b>ip</b>—Include the IP address of the IPv4 or IPv6 payload into the hash key.</li><li>• <b>layer-3</b>—Include Layer 3 IP information in the hash key.</li><li>• <b>layer-4</b>—Include Layer 4 IP information in the hash key.</li><li>• <b>outer-vlan-id</b>—Include outer VLAN ID information in the hash key.</li><li>• <b>inner-vlan-id</b>—Include inner VLAN ID information in the hash key.</li></ul> <p><b>source-mac</b>—Include source MAC address in the hash key.</p> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring PIC-Level Symmetrical Hashing for Load Balancing on 802.3ad LAGs for MX Series Routers</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |



# Local Link Bias Configuration Statement

- [local-bias on page 3057](#)

## local-bias

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | local-bias;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit interfaces aex aggregated-ether-options]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X51-D20 for EX Series switches and QFX Series devices.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | <p>Enable local link bias for all links in the aggregated Ethernet interface.</p> <p>Local link bias conserves bandwidth on Virtual Chassis ports (VCPs) by using local links to forward unicast traffic exiting a Virtual Chassis or Virtual Chassis Fabric (VCF) that has a Link Aggregation group (LAG) bundle composed of member links on different member switches in the same Virtual Chassis or VCF. A local link is a member link in the LAG bundle that is on the member switch that received the traffic.</p> <p>You should enable local link bias if you want to conserve VCP bandwidth by always forwarding egress unicast traffic on a LAG bundle out of a local link. You should not enable local link bias if you want egress traffic load-balanced as it exits the Virtual Chassis or VCF.</p> |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Local Link Bias (CLI Procedure) on page 2897</a></li> <li>• <a href="#">Understanding Local Link Bias on page 2895</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |



## CHAPTER 129

# Redundant Trunk Groups Configuration Statements

- [group \(Redundant Trunk Groups\) on page 3060](#)
- [interface \(Redundant Trunk Groups\) on page 3061](#)
- [preempt-cutover-timer on page 3062](#)
- [redundant-trunk-group on page 3063](#)

## group (Redundant Trunk Groups)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>group name {<br/>    interface interface-name &lt;primary&gt;;<br/>    interface interface-name;<br/>    preempt-cutover-timer seconds;<br/>}</pre>                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | <ul style="list-style-type: none"><li>• For platforms with ELS:<br/>[edit switch-options <b>redundant-trunk-group</b>]</li><li>• For platforms without ELS:<br/>[edit ethernet-switching-options <b>redundant-trunk-group</b>]</li></ul>                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Hierarchy level [edit switch-options] introduced in Junos OS Release 13.2X50-D10 (ELS). (See <a href="#">“Getting Started with Enhanced Layer 2 Software” on page 41</a> for information about ELS.)</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p>                                                                                                             |
| <b>Description</b>              | Create a redundant trunk group.                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <p><b>name</b>—The name of the redundant trunk group.</p> <ul style="list-style-type: none"><li>• For platforms with ELS:<br/>The group name must be a string “rtg<math>n</math>” where <math>n</math> is a number from 0 through 15, such as “rtg2” or “rtg10”.</li><li>• For platforms without ELS:<br/>The group name must start with a letter and can consist of letters, numbers, dashes, and underscores.</li></ul> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system—control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Redundant Trunk Links for Faster Recovery</i></li><li>• <a href="#">Example: Configuring Redundant Trunk Links for Faster Recovery on page 2903</a></li><li>• <a href="#">Understanding Redundant Trunk Links on page 2901</a></li></ul>                                                                                                                                                                            |

## interface (Redundant Trunk Groups)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>interface <i>interface-name</i> &lt;primary&gt;; interface <i>interface-name</i>;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | <p>For platforms with ELS:</p> <pre>[edit switch-options <b>redundant-trunk-group</b> <i>group name</i>]</pre> <p>For platforms without ELS:</p> <pre>[edit ethernet-switching-options <b>redundant-trunk-group</b> <i>group name</i>]</pre>                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Hierarchy level <b>[edit switch-options]</b> introduced in Junos OS Release 13.2X50-D10 (ELS). (See <a href="#">“Getting Started with Enhanced Layer 2 Software” on page 41</a> for information about ELS.)</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p>                                                                                                                                                                                                                                  |
| <b>Description</b>              | <p>Configure a primary link and secondary link on trunk ports. If the primary link fails, the secondary link automatically takes over as the primary link without waiting for normal STP convergence.</p>                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <p><b>interface</b> <i>interface-name</i>—A logical interface or an aggregated interface containing multiple ports.</p> <p><b>primary</b>—(Optional) Specify one of the interfaces in the redundant group as the primary link. The interface without this option is the secondary link in the redundant group. If a link is not specified as <b>primary</b>, the software compares the two links and selects the link with the highest port number as the active link. For example, if the two interfaces are <b>ge-0/1/0</b> and <b>ge-0/1/1</b>, the software assigns <b>ge-0/1/1</b> as the active link.</p> |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system—control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Redundant Trunk Links for Faster Recovery</a></li> <li>• <a href="#">Example: Configuring Redundant Trunk Links for Faster Recovery on page 2903</a></li> <li>• <a href="#">Understanding Redundant Trunk Links on page 2901</a></li> </ul>                                                                                                                                                                                                                                                                                           |

## preempt-cutover-timer

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>preempt-cutover-timer seconds;</code>                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | <ul style="list-style-type: none"><li>For platforms with ELS:<br/>[edit switch-options <a href="#">redundant-trunk-group group name</a>]</li><li>For platforms without ELS:<br/>[edit ethernet-switching-options <a href="#">redundant-trunk-group group name</a>]</li></ul>                                                                                             |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 11.1 for EX Series switches.</p> <p>Hierarchy level [edit switch-options] introduced in Junos OS Release 13.2X50-D10 (ELS). (See <a href="#">“Getting Started with Enhanced Layer 2 Software” on page 41</a> for information about ELS.)</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> |
| <b>Description</b>              | Change the length of time that a re-enabled primary link waits to take over from an active secondary link in a redundant trunk group.                                                                                                                                                                                                                                    |
| <b>Default</b>                  | If you do not change the time with the <b>preempt-cutover-timer</b> statement, a re-enabled primary link takes over from the active secondary link after 1 second.                                                                                                                                                                                                       |
| <b>Options</b>                  | <p><b>seconds</b>—Number of seconds that the primary link waits to take over from the active secondary link.</p> <p><b>Range:</b> 1 through 600 seconds</p>                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li><i>Example: Configuring Redundant Trunk Links for Faster Recovery</i></li><li><a href="#">Example: Configuring Redundant Trunk Links for Faster Recovery on page 2903</a></li><li><a href="#">Understanding Redundant Trunk Links on page 2901</a></li></ul>                                                                       |

## redundant-trunk-group

|                                 |                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> redundant-trunk-group {   group name {     interface interface-name &lt;primary&gt;;     interface interface-name;     preempt-cutover-timer seconds;   } } </pre>                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <ul style="list-style-type: none"> <li>For platforms with ELS:<br/>[edit switch-options]</li> <li>For platforms without ELS:<br/>[edit ethernet-switching-options]</li> </ul>                                                                                                                                                                                                  |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Hierarchy level <b>[edit switch-options]</b> introduced in Junos OS Release 13.2X50-D10 (ELS). (See <a href="#">“Getting Started with Enhanced Layer 2 Software” on page 41</a> for information about ELS.)</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> |
| <b>Description</b>              | <p>Configure a primary link and secondary link on trunk ports. If the primary link fails, the secondary link automatically takes over without waiting for normal spanning-tree protocol convergence.</p> <p>The remaining statements are explained separately.</p>                                                                                                             |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system—control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Example: Configuring Redundant Trunk Links for Faster Recovery</i></li> <li><a href="#">Example: Configuring Redundant Trunk Links for Faster Recovery on page 2903</a></li> <li><a href="#">Understanding Redundant Trunk Links on page 2901</a></li> </ul>                                                                         |





# Resilient Hashing Configuration Statements

- [ecmp-resilient-hash](#) on page 3065
- [enhanced-hash-key](#) on page 3066
- [hash-mode](#) on page 3069
- [hash-key \(Forwarding Options\)](#) on page 3071
- [hash-seed](#) on page 3073
- [inet \(enhanced-hash-key\)](#) on page 3074
- [inet6 \(enhanced-hash-key\)](#) on page 3076
- [ipv6-flow-label](#) on page 3078
- [resilient-hash](#) on page 3078

## [ecmp-resilient-hash](#)

---

|                                 |                                                                                                                                      |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>ecmp-resilient-hash;</code>                                                                                                    |
| <b>Hierarchy Level</b>          | [edit forwarding-options enhanced-hash-key]                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series.                                                             |
| <b>Description</b>              | Enable resilient hashing for ECMP groups, to minimize remapping of destination paths.                                                |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Resilient Hashing for Trunk/ECMP Groups</a> on page 2925</li> </ul> |

## enhanced-hash-key

---

**List of Syntax**    [Syntax \(EX Series and QFX5100 Switch\) on page 3066](#)  
                          [Syntax \(QFX10002 and QFX10008 Switches\) on page 3066](#)

**Syntax (EX Series and QFX5100 Switch)**    enhanced-hash-key {  
                                  ecmp-resilient-hash;  
                                  fabric-load-balance {  
                                    flowlet {  
                                      inactivity-interval *interval*;  
                                    }  
                                  per-packet;  
                                  }  
                                  hash-mode {  
                                    layer2-header;  
                                    layer2-payload;  
                                  }  
                                  inet {  
                                    no-ipv4-destination-address;  
                                    no-ipv4-source-address;  
                                    no-l4-destination-port;  
                                    no-l4-source-port;  
                                    no-protocol;  
                                    vlan-id;  
                                  }  
                                  inet6 {  
                                    no-ipv6-destination-address;  
                                    no-ipv6-source-address;  
                                    no-l4-destination-port;  
                                    no-l4-source-port;  
                                    no-next-header;  
                                    vlan-id;  
                                  }  
                                  layer2 {  
                                    no-destination-mac-address;  
                                    no-ether-type;  
                                    no-source-mac-address;  
                                    vlan-id;  
                                  }  
                                  }  
                                  }  
                                  }

**Syntax (QFX10002 and QFX10008 Switches)**    enhanced-hash-key {  
                                  hash-seed *seed-value*;  
                                  inet {  
                                    no-ipv4-destination-address;  
                                    no-ipv4-source-address;  
                                    no-l4-destination-port;  
                                    no-l4-source-port;  
                                  }  
                                  inet6 {  
                                    ipv6-flow-label;  
                                    no-ipv6-destination-address;  
                                    no-ipv6-source-address;  
                                    no-l4-destination-port;  
                                  }

```

 no-l4-source-port;
 }
 layer2 {
 destination-mac-address
 inner-vlan-id;
 no-ether-type;
 no-vlan-id;
 source-mac-address;
 }
 no-mpls;
 gre {
 key;
 protocol;
 }
 vxlan-vnid
 }
}

```

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Hierarchy Level</b>          | [edit forwarding-options]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 13.2X51-D15 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.</p> <p>The <b>fabric-load-balance</b> statement introduced in Junos OS Release 14.1X53-D10.</p> <p>The <b>hash-seed</b> statement introduced in Junos OS Release 15.1X53-D30.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | <p>Configure the hashing key used to hash link aggregation group (LAG) and equal-cost multipath (ECMP) traffic, or enable adaptive load balancing (ALB) in a Virtual Chassis Fabric (VCF).</p> <p>The hashing algorithm is used to make traffic-forwarding decisions for traffic entering a LAG bundle or for traffic exiting a switch when ECMP is enabled.</p> <p>For LAG bundles, the hashing algorithm determines how traffic entering a LAG bundle is placed onto the bundle's member links. The hashing algorithm tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.</p> <p>When ECMP is enabled, the hashing algorithm determines how incoming traffic is forwarded to the next-hop device.</p> <p>On QFX10002 and QFX 10008 switches, you can configure the hash seed for load balancing.</p> <p>By default, the QFX10002 and QFX10008 switches use the system MAC address to generate a hash seed value. You can configure the hash seed value using the <b>hash-seed</b> statement at the [edit forwarding-options enhanced-hash-key] hierarchy level. Set a value between 0 and 4294967295. If you do not configure a hash seed value, the system will generate a hash seed value based on the system MAC address.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

**Related  
Documentation**

- [Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic \(CLI Procedure\) on page 2923](#)
- [Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 2911](#)

## hash-mode

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>hash-mode {     layer2-header;     layer2-payload; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit forwarding-options <a href="#">enhanced-hash-key</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 13.2X51-D15 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.</p> <p>Statement is not supported on QFX10002 and QFX 10008 switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | <p>Select the mode for the hashing algorithm.</p> <p>The hashing algorithm is used to make traffic-forwarding decisions for traffic entering a LAG bundle or for traffic exiting a switch when ECMP is enabled.</p> <p>For LAG bundles, the hashing algorithm determines how traffic entering a LAG bundle is placed onto the bundle's member links. The hashing algorithm tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.</p> <p>When ECMP is enabled, the hashing algorithm determines how incoming traffic is forwarded to the next-hop device.</p> <p>The hash mode that is set using this statement determines which fields are inspected by the hashing algorithm. You must set the hash mode to <b>layer2-payload</b> if you want the hashing algorithm to inspect fields in the Layer 2 payload when making hashing decisions. You must set the hash mode to <b>layer2-header</b> if you want the hashing algorithm to inspect fields in the Layer 2 header when making hashing decisions.</p> <p>If the hash mode is set to <b>layer2-payload</b>, you can set the fields used by the hashing algorithm to hash IPv4 traffic using the <b>set forwarding-options enhanced-hash-key inet</b> statement. You can set the fields used by the hashing algorithm to hash IPv6 traffic using the <b>set forwarding-options enhanced-hash-key inet6</b> statement.</p> <p>If the hash mode is set to <b>layer2-header</b>, you can set the fields that the hashing algorithm inspects in the Layer 2 header using the <b>set forwarding-options enhanced-hash-key layer2</b> statement.</p> |
| <b>Default</b>                  | layer2-payload                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <p><b>layer-2-payload</b>—Set the hashing algorithm to use fields in the Layer 2 payload to make hashing decisions.</p> <p><b>layer-2-header</b>—Set the hashing algorithm to use fields in the Layer 2 header to make hashing decisions.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

**Related  
Documentation**

- [Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic \(CLI Procedure\) on page 2923](#)
- [Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 2911](#)
- [enhanced-hash-key on page 3066](#)
- [inet on page 3074](#)
- [inet6 on page 3076](#)
- [layer2](#)

## hash-key (Forwarding Options)

```

Syntax hash-key {
 family inet {
 layer-3;
 layer-4;
 session-id;
 symmetric-hash {
 complement;
 }
 }
 family mpls {
 all-labels;
 bottom-label-1;
 bottom-label-2;
 bottom-label-3;
 label-1;
 label-2;
 label-3;
 no-labels;
 no-label-1-exp;
 payload {
 ether-pseudowire;
 ip {
 disable;
 layer-3-only;
 port-data {
 destination-lsb;
 destination-msb;
 source-lsb;
 source-msb;
 }
 }
 }
 }
 family multiservice {
 destination-mac;
 label-1;
 label-2;
 payload {
 ip {
 layer-3-only;
 layer-3 {
 (source-ip-only | destination-ip-only);
 }
 layer-4;
 }
 }
 source-mac;
 }
 }

```

Hierarchy Level    [edit forwarding-options]

**Release Information** Statement introduced before Junos OS Release 7.4.  
**family multiservice** and **no-label-1-exp** options introduced in Junos OS Release 8.0.  
**label-3** and **no-labels** options introduced in Junos OS Release 8.1.  
**ether-pseudowire** statement introduced in Junos OS Release 9.1 (M320 and T Series routers only); support extended to M120 and MX Series routers in Junos OS Release 9.4.  
**ip**, **label-1**, **label-2**, **layer-3-only**, and **payload** options for the **family multiservice** statement introduced in Junos OS Release 9.4 (M120 and M320 routers only). For MX Series routers, only the **ip** and **payload** statements apply.  
**layer-3**, **source-ip-only**, **destination-ip-only**, and **layer-4** statements introduced for the **family multiservice** statement in Junos OS Release 9.5. (MX Series routers only).  
**all-labels** and **payload ip disable** statements introduced for the **family mpls** statement in Junos OS Release 12.1X48R2. (PTX Series Packet Transport Routers only).  
**bottom-label** statements introduced for the **family mpls** statement in Junos OS Release 14.1 for MX Series routers with DPCs (excluding M7i, M10i, and M120).

**Description** Select which packet header data to use for per-flow load balancing.

The options are explained separately.



**NOTE:** To modify the default hashing mechanism on Modular Port Concentrators (MPCs) and Type 5 FPCs, you need to configure the statements at the [edit forwarding-options enhanced-hash-key] hierarchy level. Statements at the [edit forwarding-options hash-key] hierarchy level do not support MPCs and Type 5 FPCs.



**NOTE:**

The following statements are not supported on T Series routers:

- The **symmetric-hash** and the **session-id** statements at the [edit forwarding-options hash-key family inet] hierarchy level and all statements at the [edit forwarding-options hash-key family multiservice] hierarchy level.
- The **label-1** and **label-2** statements, and the IP configuration at the [edit forwarding-options hash-key family multiservice] hierarchy level.



**NOTE:** The following statements are not supported on Q Series switches:


- The **symmetric-hash** and the **session-id** statements at the [edit forwarding-options hash-key family inet] hierarchy level.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.



- Related Documentation**
- [Configuring Per-Packet Load Balancing](#)
  - [Configuring Load Balancing Based on MPLS Labels on DPC I-Chip-Based Hardware](#)
  - [Configuring Load Balancing Based on MAC Addresses](#)

## hash-seed

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | hash-seed <i>seed-value</i> ;                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit forwarding-options <a href="#">enhanced-hash-key</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 15.1X53-D30 on QFX Series devices.                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | <p>Configure a hash seed for load-balancing functions.</p> <p>By default, the QFX10002 and QFX10008 switches use the system MAC address to generate a hash seed value. You can configure the hash seed value using the <b>hash-seed</b> statement at the [edit forwarding-options <b>enhanced-hash-key</b>] hierarchy. Set a value between 0 and 4294967295. If you do not configure a hash seed value, the system will generate a hash seed value based on the system MAC address.</p> |
|                                 | <p> <b>NOTE:</b> The <b>fabric-load-balance</b> and <b>user-defined-fields</b> statements are not supported at the [edit forwarding-options <b>enhanced-hash-key</b>] hierarchy level.</p>                                                                                                                                                                                                           |
| <b>Options</b>                  | <b>hash-seed</b> <i>seed-value</i> —A hash seed value, in the range from 0 to 4294967295.                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic (QFX 10002 and QFX 10008 Switches) on page 2919</a></li> <li>• <a href="#">Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure) on page 2923</a></li> <li>• <a href="#">show forwarding-options enhanced-hash-key on page 3216</a></li> </ul>                                                     |

## inet (enhanced-hash-key)

---

|                                                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax (EX Series and QFX5100 Switch)</b>   | <pre>inet {<br/>  no-ipv4-destination-address;<br/>  no-ipv4-source-address;<br/>  no-l4-destination-port;<br/>  no-l4-source-port;<br/>  no-protocol;<br/>  vlan-id;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Syntax (QFX10002 and QFX10008 Switches)</b> | <pre>inet {<br/>  no-ipv4-destination-address;<br/>  no-ipv4-source-address;<br/>  no-l4-destination-port;<br/>  no-l4-source-port;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>                         | [edit forwarding-options <a href="#">enhanced-hash-key</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>                     | <p>Statement introduced in Junos OS Release 13.2X51-D15 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.</p> <p>Statement introduced in Junos OS Release 15.1X53-D30 on QFX10002 and QFX10008 Switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>                             | <p>Select the payload fields in IPv4 traffic used by the hashing algorithm to make hashing decisions.</p> <p>When IPv4 traffic enters a LAG and the hash mode is set to Layer 2 payload, the hashing algorithm checks the fields configured using the <b>inet</b> statement and uses the information in the fields to decide how to place traffic onto the LAG bundle's member links or how to forward traffic to the next hop device when ECMP is enabled.</p> <p>The hashing algorithm, when used to hash LAG bundle traffic, always tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.</p> <p>The hashing algorithm only inspects the IPv4 fields in the payload to make hashing decisions when the hash mode is set to <b>layer2-payload</b>. The hash mode is set to Layer 2 payload by default. You can set the hash mode to Layer 2 payload using the <b>set forwarding-options enhanced-hash-key hash-mode layer2-payload</b> statement.</p> |
| <b>Default</b>                                 | <p>The following fields are used by the hashing algorithm to make hashing decisions for IPv4 traffic:</p> <ul style="list-style-type: none"><li>• IP destination address</li><li>• IP source address</li><li>• Layer 4 destination port</li><li>• Layer 4 source port</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

- Protocol

|                |                                                                                                            |
|----------------|------------------------------------------------------------------------------------------------------------|
| <b>Options</b> | <b>no-ipv4-destination-address</b> —Exclude the IPv4 destination address field from the hashing algorithm. |
|                | <b>no-ipv4-source-address</b> —Exclude the IPv4 source address field from the hashing algorithm.           |
|                | <b>no-l4-destination-port</b> —Exclude the Layer 4 destination port field from the hashing algorithm.      |
|                | <b>no-l4-source-port</b> —Exclude the Layer 4 source port field from the hashing algorithm.                |
|                | <b>no-protocol</b> —Exclude the protocol field from the hashing algorithm.                                 |
|                | <b>vlan-id</b> —Include the VLAN ID field in the hashing algorithm.                                        |

|                                 |                                                               |
|---------------------------------|---------------------------------------------------------------|
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.        |
|                                 | interface-control—To add this statement to the configuration. |

|                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure) on page 2923</a></li><li>• <a href="#">Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 2911</a></li><li>• <a href="#">Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic (QFX 10002 and QFX 10008 Switches) on page 2919</a></li><li>• <a href="#">hash-seed on page 3073</a></li><li>• <a href="#">enhanced-hash-key on page 3066</a></li><li>• <a href="#">hash-mode on page 3069</a></li><li>• <a href="#">inet6 on page 3076</a></li></ul> |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## inet6 (enhanced-hash-key)

---

|                                                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                          | <a href="#">Syntax (EX Series and QFX5100 Switch) on page 3076</a><br><a href="#">Syntax (QFX10002 and QFX10008 Switches) on page 3076</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Syntax (EX Series and QFX5100 Switch)</b>   | <pre>inet6 {<br/>  no-ipv6-destination-address;<br/>  no-ipv6-source-address;<br/>  no-l4-destination-port;<br/>  no-l4-source-port;<br/>  no-next-header;<br/>  vlan-id;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Syntax (QFX10002 and QFX10008 Switches)</b> | <pre>inet6 {<br/>  ipv6-flow-label;<br/>  no-ipv6-destination-address;<br/>  no-ipv6-source-address;<br/>  no-l4-destination-port;<br/>  no-l4-source-port;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>                         | [edit forwarding-options <a href="#">enhanced-hash-key</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>                     | Statement introduced in Junos OS Release 13.2X51-D15 on EX Series switches.<br>Statement introduced in Junos OS Release 13.2X51-D20 on QFX Series devices.<br>Statement introduced in Junos OS Release 15.1X53-D30 on QFX10002 and QFX 10008 switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>                             | <p>Select the payload fields in an IPv6 packet used by the hashing algorithm to make hashing decisions.</p> <p>When IPv6 traffic enters a LAG and the hash mode is set to Layer 2 payload, the hashing algorithm checks the fields configured using this statement and uses the information in the fields to decide how to place traffic onto the LAG bundle's member links or to forward traffic to the next hop device when ECMP is enabled.</p> <p>The hashing algorithm, when used to hash LAG traffic, always tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.</p> <p>The hashing algorithm only inspects the IPv6 fields in the payload to make hashing decisions when the hash mode is set to Layer 2 payload. The hash mode is set to Layer 2 payload by default. You can set the hash mode to Layer 2 payload using the <b>set forwarding-options enhanced-hash-key hash-mode layer2-payload</b> statement.</p> |
| <b>Default</b>                                 | <p>The data in the following fields are used by the hashing algorithm to make hashing decisions for IPv6 traffic:</p> <ul style="list-style-type: none"><li>• IP destination address</li><li>• IP source address</li><li>• Layer 4 destination port</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

- Layer 4 source port
- Next header

**Options**    **no-ipv6-destination-address**—Exclude the IPv6 destination address field from the hashing algorithm.

**no-ipv6-source-address**—Exclude the IPv6 source address field from the hashing algorithm.

**no-l4-destination-port**—Exclude the Layer 4 destination port field from the hashing algorithm.

**no-l4-source-port**—Exclude the Layer 4 source port field from the hashing algorithm.

**no-next-header**—Exclude the Next Header field from the hashing algorithm.

**vlan-id**—Include the VLAN ID field in the hashing algorithm.

**Required Privilege Level**    interface—To view this statement in the configuration.  
                                         interface-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic \(CLI Procedure\) on page 2923](#)
- [Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 2911](#)
- [Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic \(QFX 10002 and QFX 10008 Switches\) on page 2919](#)
- [hash-seed on page 3073](#)
- [enhanced-hash-key on page 3066](#)
- [hash-mode on page 3069](#)
- [inet on page 3074](#)


## ipv6-flow-label

---

|                                 |                                                                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | ipv6-flow-label;                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit forwarding-options <a href="#">enhanced-hash-key inet6</a> ]                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 15.1X53-D30 on QFX10002 and 10008 switches..                                                                                                                                                                                                   |
| <b>Description</b>              | Enable IPv6 packet flow labels for hash calculations on QFX10002 and QFX 10008 switches.                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic (QFX 10002 and QFX 10008 Switches) on page 2919</a></li><li>• <a href="#">show forwarding-options enhanced-hash-key on page 3216</a></li></ul> |

## resilient-hash

---

|                                 |                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | resilient-hash;                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit interfaces aex aggregated-ether-options]]                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series.                                                                                                                               |
|                                 | <div> <b>NOTE:</b> Configuring resilient hashing on LAGs is not supported on QFX10002 and QFX 10008 switches.</div> |
| <b>Description</b>              | Enable resilient hashing for a LAG, to minimize remapping of destination paths.                                                                                                                        |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Resilient Hashing for Trunk/ECMP Groups on page 2925</a></li></ul>                                                                     |

# Uplink Failure Detection Configuration Statements

- [group](#) on page 3079
- [link-to-disable](#) on page 3080
- [link-to-monitor](#) on page 3080
- [uplink-failure-detection](#) on page 3081

## group

---

**Syntax** `group group-name {  
     link-to-monitor interface-name;  
     link-to-disable interface-name;  
 }`

**Hierarchy Level** [edit protocols uplink-failure-detection]

**Release Information** Statement introduced in Junos OS Release 12.1 for the QFX Series.

**Description** Configure a group of uplink and downlink interfaces for uplink failure detection.

**Options** *group-name*—Name of the uplink failure detection group.

The remaining statements are explained separately.

**Required Privilege Level** admin—To view this statement in the configuration.  
 admin-control—To add this statement to the configuration.

**Related Documentation**

- [Overview of Uplink Failure Detection](#) on page 2929
- [Configuring Interfaces for Uplink Failure Detection](#) on page 2931
- [Example: Configuring Interfaces for Uplink Failure Detection](#) on page 2932

## link-to-disable

---

|                                 |                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | link-to-disable <i>interface-name</i> ;                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit protocols uplink-failure-detection group <i>group-name</i> ]                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Configure the downlink interfaces to be disabled when the switch detects an uplink failure. The switch can monitor a maximum of eight downlink interfaces in a group.                                                                                                                                              |
| <b>Options</b>                  | <i>interface-name</i> —Name of the downlink interface in an uplink failure detection group. The interface can be a physical interface or a logical interface.                                                                                                                                                      |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Overview of Uplink Failure Detection on page 2929</a></li><li>• <a href="#">Configuring Interfaces for Uplink Failure Detection on page 2931</a></li><li>• <a href="#">Example: Configuring Interfaces for Uplink Failure Detection on page 2932</a></li></ul> |

## link-to-monitor

---

|                                 |                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | link-to-monitor <i>interface-name</i> ;                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit protocols uplink-failure-detection group <i>group-name</i> ]                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Configure the uplink interfaces to be monitored for uplink failure detection. The switch can monitor a maximum of eight uplink interfaces in a group.                                                                                                                                                              |
| <b>Options</b>                  | <i>interface-name</i> —Name of the uplink interface in an uplink failure detection group. The interface can be a physical interface or a logical interface.                                                                                                                                                        |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Overview of Uplink Failure Detection on page 2929</a></li><li>• <a href="#">Configuring Interfaces for Uplink Failure Detection on page 2931</a></li><li>• <a href="#">Example: Configuring Interfaces for Uplink Failure Detection on page 2932</a></li></ul> |



---

## uplink-failure-detection

---

|                                 |                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>uplink-failure-detection {<br/>  group <i>group-name</i> {<br/>    link-to-monitor <i>interface-name</i>;<br/>    link-to-disable <i>interface-name</i>;<br/>  }<br/>}</pre>                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit protocols]                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                  |
| <b>Description</b>              | <p>Configure uplink and downlink interfaces in a group to monitor uplink failures and to propagate uplink failure information to the downlink interfaces.</p> <p>The remaining statements are explained separately.</p>                                                                                            |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Overview of Uplink Failure Detection on page 2929</a></li><li>• <a href="#">Configuring Interfaces for Uplink Failure Detection on page 2931</a></li><li>• <a href="#">Example: Configuring Interfaces for Uplink Failure Detection on page 2932</a></li></ul> |



## CHAPTER 132

# Ethernet OAM Link Fault Management Operational Command

- `show oam ethernet link-fault-management`

## show oam ethernet link-fault-management

|                                 |                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show oam ethernet link-fault-management<br><brief   detail><br><interface-name>                                                                                                                                                                                       |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.4 for EX Series switches.                                                                                                                                                                                                    |
| <b>Description</b>              | Displays Operation, Administration, and Maintenance (OAM) link fault management (LFM) information for Ethernet interfaces.                                                                                                                                            |
| <b>Options</b>                  | <b>brief   detail</b> —(Optional) Display the specified level of output.<br><br><b>interface-name</b> —(Optional) Display link fault management information for the specified Ethernet interface only.                                                                |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 2770</a></li> <li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 2768</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show oam ethernet link-fault-management brief on page 3088</a><br><a href="#">show oam ethernet link-fault-management detail on page 3088</a>                                                                                                             |
| <b>Output Fields</b>            | Table 272 lists the output fields for the <b>show oam ethernet link-fault-management</b> command. Output fields are listed in the approximate order in which they appear.                                                                                             |

Table 272: show oam ethernet link-fault-management Output Fields

| Field Name             | Field Description                                                                                                                                                                                               | Level of Output |
|------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Status</b>          | Indicates the status of the established link. <ul style="list-style-type: none"> <li>• <b>Fail</b>—A link fault condition exists.</li> <li>• <b>Running</b>—A link fault condition does not exist.</li> </ul>   | All levels      |
| <b>Discovery state</b> | State of the discovery mechanism: <ul style="list-style-type: none"> <li>• <b>Passive Wait</b></li> <li>• <b>Send Any</b></li> <li>• <b>Send Local Remote</b></li> <li>• <b>Send Local Remote Ok</b></li> </ul> | All levels      |
| <b>Peer address</b>    | Address of the OAM peer.                                                                                                                                                                                        | All levels      |

Table 272: show oam ethernet link-fault-management Output Fields (*continued*)

| Field Name                       | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Level of Output |
|----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Flags</b>                     | Information about the interface. <ul style="list-style-type: none"> <li><b>Remote-Stable</b>—Indicates remote OAM client acknowledgment of, and satisfaction with local OAM state information. <b>False</b> indicates that remote DTE has either not seen or is unsatisfied with local state information. <b>True</b> indicates that remote DTE has seen and is satisfied with local state information.</li> <li><b>Local-Stable</b>—Indicates local OAM client acknowledgment of, and satisfaction with remote OAM state information. <b>False</b> indicates that local DTE either has not seen or is unsatisfied with remote state information. <b>True</b> indicates that local DTE has seen and is satisfied with remote state information.</li> <li><b>Remote-State-Valid</b>—Indicates the OAM client has received remote state information found within Local Information TLVs of received Information OAM PDUs. <b>False</b> indicates that OAM client has not seen remote state information. <b>True</b> indicates that the OAM client has seen remote state information.</li> </ul>                                                                      | All levels      |
| <b>Remote loopback status</b>    | Indicates the remote loopback status. An OAM entity can put its remote peer into loopback mode using the Loopback control OAM PDU. In loopback mode, every frame received is transmitted back on the same port (except for OAM PDUs, which are needed to maintain the OAM session).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | All levels      |
| <b>Remote entity information</b> | Remote entity information. <ul style="list-style-type: none"> <li><b>Remote MUX action</b>—Indicates the state of the multiplexer functions of the OAM sublayer. Device is forwarding non-OAM PDUs to the lower sublayer or discarding non-OAM PDUs.</li> <li><b>Remote parser action</b>—Indicates the state of the parser function of the OAM sublayer. Device is forwarding non-OAM PDUs to higher sublayer, looping back non-OAM PDUs to the lower sublayer, or discarding non-OAM PDUs.</li> <li><b>Discovery mode</b>—Indicates whether discovery mode is active or inactive.</li> <li><b>Unidirectional mode</b>—Indicates the ability to operate a link in a unidirectional mode for diagnostic purposes.</li> <li><b>Remote loopback mode</b>—Indicates whether remote loopback is supported or not supported.</li> <li><b>Link events</b>—Indicates whether interpreting link events is supported or not supported on the remote peer.</li> <li><b>Variable requests</b>—Indicates whether variable requests are supported or not supported. The Variable Request OAM PDU, is used to request one or more MIB variables from the remote peer.</li> </ul> | All levels      |
| <b>OAM Receive Statistics</b>    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |                 |
| <b>Information</b>               | The number of information PDUs received.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>detail</b>   |
| <b>Event</b>                     | The number of loopback control PDUs received.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail</b>   |
| <b>Variable request</b>          | The number of variable request PDUs received.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail</b>   |
| <b>Variable response</b>         | The number of variable response PDUs received.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>detail</b>   |
| <b>Loopback control</b>          | The number of loopback control PDUs received.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail</b>   |

Table 272: show oam ethernet link-fault-management Output Fields (*continued*)

| Field Name                                         | Field Description                                                                                                                                                                  | Level of Output |
|----------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Organization specific</b>                       | The number of vendor organization specific PDUs received.                                                                                                                          | <b>detail</b>   |
| <b>OAM Transmit Statistics</b>                     |                                                                                                                                                                                    |                 |
| <b>Information</b>                                 | The number of information PDUs transmitted.                                                                                                                                        | <b>detail</b>   |
| <b>Event</b>                                       | The number of event notification PDUs transmitted.                                                                                                                                 | <b>detail</b>   |
| <b>Variable request</b>                            | The number of variable request PDUs transmitted.                                                                                                                                   | <b>detail</b>   |
| <b>Variable response</b>                           | The number of variable response PDUs transmitted.                                                                                                                                  | <b>detail</b>   |
| <b>Loopback control</b>                            | The number of loopback control PDUs transmitted.                                                                                                                                   | <b>detail</b>   |
| <b>Organization specific</b>                       | The number of vendor organization specific PDUs transmitted.                                                                                                                       | <b>detail</b>   |
| <b>OAM Received Symbol Error Event information</b> |                                                                                                                                                                                    |                 |
| <b>Events</b>                                      | The number of symbol error event TLVs that have been received after the OAM sublayer was reset.                                                                                    | <b>detail</b>   |
| <b>Window</b>                                      | The symbol error event window in the received PDU.<br><br>The protocol default value is the number of symbols that can be received in one second on the underlying physical layer. | <b>detail</b>   |
| <b>Threshold</b>                                   | The number of errored symbols in the period required for the event to be generated.                                                                                                | <b>detail</b>   |
| <b>Errors in period</b>                            | The number of symbol errors in the period reported in the received event PDU.                                                                                                      | <b>detail</b>   |
| <b>Total errors</b>                                | The number of errored symbols that have been reported in received event TLVs after the OAM sublayer was reset.<br><br>Symbol errors are coding symbol errors.                      | <b>detail</b>   |
| <b>OAM Received Frame Error Event Information</b>  |                                                                                                                                                                                    |                 |
| <b>Events</b>                                      | The number of errored frame event TLVs that have been received after the OAM sublayer was reset.                                                                                   | <b>detail</b>   |
| <b>Window</b>                                      | The duration of the window in terms of the number of 100 ms period intervals.                                                                                                      | <b>detail</b>   |
| <b>Threshold</b>                                   | The number of detected errored frames required for the event to be generated.                                                                                                      | <b>detail</b>   |
| <b>Errors in period</b>                            | The number of detected errored frames in the period.                                                                                                                               | <b>detail</b>   |

Table 272: show oam ethernet link-fault-management Output Fields (*continued*)

| Field Name                                               | Field Description                                                                                                                                                                       | Level of Output |
|----------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Total errors</b>                                      | The number of errored frames that have been reported in received event TLVs after the OAM sublayer was reset.<br><br>A frame error is any frame error on the underlying physical layer. | <b>detail</b>   |
| <b>OAM Received Frame Period Error Event Information</b> |                                                                                                                                                                                         |                 |
| <b>Events</b>                                            | The number of frame seconds errors event TLVs that have been received after the OAM sublayer was reset.                                                                                 | <b>detail</b>   |
| <b>Window</b>                                            | The duration of the frame seconds window.                                                                                                                                               | <b>detail</b>   |
| <b>Threshold</b>                                         | The number of frame seconds errors in the period.                                                                                                                                       | <b>detail</b>   |
| <b>Errors in period</b>                                  | The number of frame seconds errors in the period.                                                                                                                                       | <b>detail</b>   |
| <b>Total errors</b>                                      | The number of frame seconds errors that have been reported in received event TLVs after the OAM sublayer was reset.                                                                     | <b>detail</b>   |
| <b>OAM Transmitted Symbol Error Event Information</b>    |                                                                                                                                                                                         |                 |
| <b>Events</b>                                            | The number of symbol error event TLVs that have been transmitted after the OAM sublayer was reset.                                                                                      | <b>detail</b>   |
| <b>Window</b>                                            | The symbol error event window in the transmitted PDU.                                                                                                                                   | <b>detail</b>   |
| <b>Threshold</b>                                         | The number of errored symbols in the period required for the event to be generated.                                                                                                     | <b>detail</b>   |
| <b>Errors in period</b>                                  | The number of symbol errors in the period reported in the transmitted event PDU.                                                                                                        | <b>detail</b>   |
| <b>Total errors</b>                                      | The number of errored symbols reported in event TLVs that have been transmitted after the OAM sublayer was reset.                                                                       | <b>detail</b>   |
| <b>OAM Transmitted Frame Error Event Information</b>     |                                                                                                                                                                                         |                 |
| <b>Events</b>                                            | The number of errored frame event TLVs that have been transmitted after the OAM sublayer was reset.                                                                                     | <b>detail</b>   |
| <b>Window</b>                                            | The duration of the window in terms of the number of 100 ms period intervals.                                                                                                           | <b>detail</b>   |
| <b>Threshold</b>                                         | The number of detected errored frames required for the event to be generated.                                                                                                           | <b>detail</b>   |
| <b>Errors in period</b>                                  | The number of detected errored frames in the period.                                                                                                                                    | <b>detail</b>   |
| <b>Total errors</b>                                      | The number of errored frames that have been detected after the OAM sublayer was reset.                                                                                                  | <b>detail</b>   |

## Sample Output

### show oam ethernet link-fault-management brief

```
user@host> show oam ethernet link-fault-management brief
Interface: ge-0/0/1
Status: Running, Discovery state: Send Any
Peer address: 00:90:69:72:2c:83
Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
Remote loopback status: Disabled on local port, Enabled on peer port
Remote entity information:
 Remote MUX action: discarding, Remote parser action: loopback
 Discovery mode: active, Unidirectional mode: unsupported
 Remote loopback mode: supported, Link events: supported
 Variable requests: unsupported
```

### show oam ethernet link-fault-management detail

```
user@host> show oam ethernet link-fault-management detail
Interface: ge-0/0/1
Status: Running, Discovery state: Send Any
Peer address: 00:90:69:0a:07:14
Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
OAM receive statistics:
 Information: 186365, Event: 0, Variable request: 0, Variable response: 0
 Loopback control: 0, Organization specific: 0
OAM transmit statistics:
 Information: 186347, Event: 0, Variable request: 0, Variable response: 0
 Loopback control: 0, Organization specific: 0
OAM received symbol error event information:
 Events: 0, Window: 0, Threshold: 0
 Errors in period: 0, Total errors: 0
OAM received frame error event information:
 Events: 0, Window: 0, Threshold: 0
 Errors in period: 0, Total errors: 0
OAM received frame period error event information:
 Events: 0, Window: 0, Threshold: 0
 Errors in period: 0, Total errors: 0
OAM transmitted symbol error event information:
 Events: 0, Window: 0, Threshold: 1
 Errors in period: 0, Total errors: 0
OAM transmitted frame error event information:
 Events: 0, Window: 0, Threshold: 1
 Errors in period: 0, Total errors: 0
Remote entity information:
 Remote MUX action: forwarding, Remote parser action: forwarding
 Discovery mode: active, Unidirectional mode: unsupported
 Remote loopback mode: supported, Link events: supported
 Variable requests: unsupported
```



## CHAPTER 133

# Interfaces Operational Commands

- `monitor interface`
- `show interfaces diagnostics optics`
- `show interfaces ge`
- `show interfaces (GRE)`
- `show interfaces irb`
- `show interfaces queue`
- `show interfaces xe`

## monitor interface

**Syntax** `monitor interface`  
`<interface-name> | traffic <detail>`

**Release Information** Command introduced before Junos OS Release 7.4.  
 Command introduced in Junos OS Release 9.0 for EX Series switches.  
 Command introduced in Junos OS Release 11.1 for the QFX Series.  
 Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Display real-time statistics about interfaces, updating the statistics every second. Check for and display common interface failures, such as SONET/SDH and T3 alarms, loopbacks detected, and increases in framing errors.



**NOTE:** This command is not supported on the QFX3000 QFabric switch.

**Options** **none**—Display real-time statistics for all interfaces.

**detail**—(Optional) With traffic option only, display detailed output.

**interface-name**—(Optional) Display real-time statistics for the specified interface. In a TX Matrix or TX Matrix Plus router, display real-time statistics for the physical interfaces on the specified line-card chassis (LCC) only.

**traffic**—(Optional) Display traffic data for all active interfaces. In a TX Matrix or TX Matrix Plus router, display real-time statistics for the physical interfaces on the specified LCC only.

**Additional Information** The output of this command shows how much each field has changed since you started the command or since you cleared the counters by pressing the c key. For a description of the statistical information provided in the output of this command, see the **show interfaces extensive** command for a particular interface type in the [CLI Explorer](#). To control the output of the **monitor interface** command while it is running, use the keys listed in [Table 273](#). The keys are not case-sensitive.

**Table 273: Output Control Keys for the monitor interface interface-name Command**

| Key | Action                                                                                                                                                                                                                     |
|-----|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| c   | Clears (returns to zero) the delta counters since <b>monitor interface</b> was started. This does not clear the accumulative counter. To clear the accumulative counter, use the <b>clear interfaces interval</b> command. |
| f   | Freezes the display, halting the display of updated statistics and delta counters.                                                                                                                                         |
| i   | Displays information about a different interface. The command prompts you for the name of a specific interface.                                                                                                            |

**Table 273: Output Control Keys for the monitor interface interface-name Command** (*continued*)

| Key      | Action                                                                                                                                                                                         |
|----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| n        | Displays information about the next interface. The <b>monitor interface</b> command displays the physical or logical interfaces in the same order as the <b>show interfaces terse</b> command. |
| q or Esc | Quits the command and returns to the command prompt.                                                                                                                                           |
| t        | Thaws the display, resuming the update of the statistics and delta counters.                                                                                                                   |

To control the output of the **monitor interface traffic** command while it is running, use the keys listed in [Table 274](#). The keys are not case-sensitive.

**Table 274: Output Control Keys for the monitor interface traffic Command**

| Key      | Action                                                                                                               |
|----------|----------------------------------------------------------------------------------------------------------------------|
| b        | Displays the statistics in units of bytes and bytes per second (bps).                                                |
| c        | Clears (return to 0) the delta counters in the <b>Current Delta</b> column. The statistics counters are not cleared. |
| d        | Displays the <b>Current Delta</b> column (instead of the rate column) in bps or packets per second (pps).            |
| p        | Displays the statistics in units of packets and packets per second (pps).                                            |
| q or Esc | Quits the command and returns to the command prompt.                                                                 |
| r        | Displays the rate column (instead of the <b>Current Delta</b> column) in bps and pps.                                |

**Required Privilege Level** trace

**List of Sample Output**

- [monitor interface \(Physical\) on page 3093](#)
- [monitor interface \(OTN Interface\) on page 3094](#)
- [monitor interface \(MX480 Router with MPC5E and 10-Gigabit Ethernet OTN Interface\) on page 3095](#)
- [monitor interface \(MX480 Router with MPC5E and 100-Gigabit Ethernet Interface\) on page 3096](#)
- [monitor interface \(MX2010 Router with MPC6E and 10-Gigabit Ethernet OTN Interface\) on page 3096](#)
- [monitor interface \(MX2010 Router with MPC6E and 100-Gigabit Ethernet OTN Interface\) on page 3097](#)
- [monitor interface \(MX2020 Router with MPC6E and 10-Gigabit Ethernet OTN Interface\) on page 3098](#)
- [monitor interface \(Logical\) on page 3098](#)
- [monitor interface \(QFX3500 Switch\) on page 3099](#)

[monitor interface traffic on page 3099](#)

[monitor interface traffic \(QFX3500 Switch\) on page 3100](#)

[monitor interface traffic detail \(QFX3500 Switch\) on page 3100](#)

**Output Fields** Table 275 describes the output fields for the **monitor interface** command. Output fields are listed in the approximate order in which they appear.

**Table 275: monitor interface Output Fields**

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Level of Output |
|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>routerl</b>           | Hostname of the router.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | All levels      |
| <b>Seconds</b>           | How long the monitor interface command has been running or how long since you last cleared the counters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | All levels      |
| <b>Time</b>              | Current time (UTC).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | All levels      |
| <b>Delay x/y/z</b>       | Time difference between when the statistics were displayed and the actual clock time. <ul style="list-style-type: none"> <li><b>x</b>—Time taken for the last polling (in milliseconds).</li> <li><b>y</b>—Minimum time taken across all pollings (in milliseconds).</li> <li><b>z</b>—Maximum time taken across all pollings (in milliseconds).</li> </ul>                                                                                                                                                                                                                                                                                                                      | All levels      |
| <b>Interface</b>         | Short description of the interface, including its name, status, and encapsulation.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | All levels      |
| <b>Link</b>              | State of the link: <b>Up</b> , <b>Down</b> , or <b>Test</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | All levels      |
| <b>Current delta</b>     | Cumulative number for the counter in question since the time shown in the Seconds field, which is the time since you started the command or last cleared the counters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | All levels      |
| <b>Local Statistics</b>  | (Logical interfaces only) Number and rate of bytes and packets destined to the router or switch through the specified interface. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It usually takes less than 1 second for this counter to stabilize. <ul style="list-style-type: none"> <li><b>Input bytes</b>—Number of bytes received on the interface.</li> <li><b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li><b>Input packets</b>—Number of packets received on the interface.</li> <li><b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> | All levels      |
| <b>Remote Statistics</b> | (Logical interfaces only) Statistics for traffic transiting the router or switch. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It usually takes less than 1 second for this counter to stabilize. <ul style="list-style-type: none"> <li><b>Input bytes</b>—Number of bytes received on the interface.</li> <li><b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li><b>Input packets</b>—Number of packets received on the interface.</li> <li><b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>                                                | All levels      |

Table 275: monitor interface Output Fields (*continued*)

| Field Name         | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Level of Output |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Traffic statistics | <p>Total number of bytes and packets received and transmitted on the interface. These statistics are the sum of the local and remote statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It usually takes less than 1 second for this counter to stabilize.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> | All levels      |
| Description        | With the <b>traffic</b> option, displays the interface description configured at the <b>[edit interfaces <i>interface-name</i>]</b> hierarchy level.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>detail</b>   |

## Sample Output

### monitor interface (Physical)

```

user@host> monitor interface so-0/0/0
router1 Seconds: 19 Time: 15:46:29

Interface: so-0/0/0, Enabled, Link is Up
Encapsulation: PPP, Keepalives, Speed: 0C48
Traffic statistics:
 Input packets: 6045 (0 pps)
 Input bytes: 6290065 (0 bps)
 Output packets: 10376 (0 pps)
 Output bytes: 10365540 (0 bps)
Encapsulation statistics:
 Input keepalives: 1901
 Output keepalives: 1901
 NCP state: Opened
 LCP state: Opened
Error statistics:
 Input errors: 0
 Input drops: 0
 Input framing errors: 0
 Policed discards: 0
 L3 incompletes: 0
 L2 channel errors: 0
 L2 mismatch timeouts: 0
 Carrier transitions: 1
 Output errors: 0
 Output drops: 0
 Aged packets: 0
Active alarms : None
Active defects: None
SONET error counts/seconds:
 LOS count 1
 LOF count 1
 SEF count 1
 ES-S 0
 SES-S 0
SONET statistics:
 BIP-B1 458871

```

```

BIP-B2 460072 [0]
REI-L 465610 [0]
BIP-B3 458978 [0]
REI-P 458773 [0]

```

## Received SONET overhead:

```

F1 : 0x00 J0 : 0x00 K1 : 0x00
K2 : 0x00 S1 : 0x00 C2 : 0x00
C2(cmp) : 0x00 F2 : 0x00 Z3 : 0x00
Z4 : 0x00 S1(cmp) : 0x00

```

## Transmitted SONET overhead:

```

F1 : 0x00 J0 : 0x01 K1 : 0x00
K2 : 0x00 S1 : 0x00 C2 : 0xcf
F2 : 0x00 Z3 : 0x00 Z4 : 0x00

```

Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'

## monitor interface (OTN Interface)

```
user@host> monitor interface ge-7/0/0
```

```

Interface: ge-7/0/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 10000mbps
Traffic statistics:
 Input bytes: 0 (0 bps)
 Output bytes: 0 (0 bps)
 Input packets: 0 (0 pps)
 Output packets: 0 (0 pps)
Error statistics:
 Input errors: 0
 Input drops: 0
 Input framing errors: 0
 Policed discards: 0
 L3 incompletes: 0
 L2 channel errors: 0
 L2 mismatch timeouts: 0
 Carrier transitions: 5
 Output errors: 0
 Output drops: 0
 Aged packets: 0
Active alarms : None
Active defects: None
Input MAC/Filter statistics:
 Unicast packets 0
 Broadcast packets 0
 Multicast packets 0
 Oversized frames 0
 Packet reject count 0
 DA rejects 0
 SA rejects 0
Output MAC/Filter Statistics:
 Unicast packets 0
 Broadcast packets 0
 Multicast packets 0
 Packet pad count 0
 Packet error count 0
OTN Link 0
 OTN Alarms: OTU_BDI, OTU_TTIM, ODU_BDI
 OTN Defects: OTU_BDI, OTU_TTIM, ODU_BDI, ODU_TTIM
 OTN OC - Seconds
 LOS 2

```

```

LOF 9
OTN OTU - FEC Statistics
 Corr err ratio N/A
 Corr bytes 0
 Uncorr words 0
OTN OTU - Counters
 BIP 0
 BBE 0
 ES 0
 SES 0
 UAS 422
OTN ODU - Counters
 BIP 0
 BBE 0
 ES 0
 SES 0
 UAS 422
OTN ODU - Received Overhead APSPCC 0-3: 0

```

#### monitor interface (MX480 Router with MPC5E and 10-Gigabit Ethernet OTN Interface)

```

user@host> monitor interface xe-0/0/3
Interface: xe-0/0/3, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 10000mbps
Traffic statistics:
 Input bytes: 0 (0 bps)
 Output bytes: 0 (0 bps)
 Input packets: 0 (0 pps)
 Output packets: 0 (0 pps)
Error statistics:
 Input errors: 0
 Input drops: 0
 Input framing errors: 0
 Policed discards: 0
 L3 incompletes: 0
 L2 channel errors: 0
 L2 mismatch timeouts: 0
 Carrier transitions: 5
 Output errors: 0
 Output drops: 0
 Aged packets: 0
Active alarms : None
Active defects: None
PCS statistics:
 Bit Errors 0
 Errored blocks 4
Input MAC/Filter statistics:
 Unicast packets 0
 Broadcast packets 0
 Multicast packets 0
 Oversized frames 0
 Packet reject count 0
 DA rejects 0
 SA rejects 0
Output MAC/Filter Statistics:
 Unicast packets 0
 Broadcast packets 0
 Multicast packets 0
 Packet pad count 0
 Packet error count 0

```

Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'

#### monitor interface (MX480 Router with MPC5E and 100-Gigabit Ethernet Interface)

```

user@host> monitor interface et-2/1/0
Interface: et-2/1/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 100000mbps
Traffic statistics:
Input bytes: 0 (0 bps) [0]
Output bytes: 0 (0 bps) [0]
Input packets: 0 (0 pps) [0]
Output packets: 0 (0 pps) [0]
Error statistics:
Input errors: 0 [0]
Input drops: 0 [0]
Input framing errors: 0 [0]
Policed discards: 0 [0]
L3 incompletes: 0 [0]
L2 channel errors: 0 [0]
L2 mismatch timeouts: 0 [0]
Carrier transitions: 263 [0]
Output errors: 0 [0]
Output drops: 0 [0]
Aged packets: 0 [0]
OTN Link 0
OTN Alarms:
OTN Defects:
OTN OC - Seconds
LOS 129 [0]
LOF 2 [0]
OTN OTU - FEC Statistics
Corr err ratio <8E-5
Corr bytes 169828399453 [0]
Uncorr words 28939961456 [0]
OTN OTU - Counters
BIP 0 [0]
BBE 0 [0]
ES 24 [0]
SES 0 [0]
UAS 1255 [0]
OTN ODU - Counters
BIP 0 [0]
BBE 0 [0]
ES 24 [0]
SES 0 [0]
UAS 1256 [0]
OTN ODU - Received Overhead
APSPCC 0-3: 00 00 00 00 [0]

```

Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'

#### monitor interface (MX2010 Router with MPC6E and 10-Gigabit Ethernet OTN Interface)

```

user@host> monitor interface xe-6/1/0
Interface: xe-6/1/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 10000mbps
Traffic statistics:
Input bytes: 0 (0 bps) [0]

```



```

Output bytes: 0 (0 bps) [0]
Input packets: 0 (0 pps) [0]
Output packets: 0 (0 pps) [0]
Error statistics:
Input errors: 0 [0]
Input drops: 0 [0]
Input framing errors: 0 [0]
Policed discards: 0 [0]
L3 incompletes: 0 [0]
L2 channel errors: 0 [0]
L2 mismatch timeouts: 0 [0]
Carrier transitions: 1 [0]
Output errors: 0 [0]
Output drops: 0 [0]
Aged packets: 0 [0]
Active alarms : None
Active defects: None
PCS statistics:
 Seconds
 Bit Errors 0 [0]
 Errored blocks 1 [0]
Input MAC/Filter statistics:
Unicast packets 0 [0]
Broadcast packets 0 [0]
Multicast packets 0 [0]
Oversized frames 0 [0]
Packet reject count 0 [0]
DA rejects 0 [0]
SA rejects 0 [0]
Output MAC/Filter Statistics:
Unicast packets 0 [0]
Broadcast packets 0 [0]
Multicast packets 0 [0]
Packet pad count 0 [0]
Packet error count 0 [0]

```

Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'

#### monitor interface (MX2010 Router with MPC6E and 100-Gigabit Ethernet OTN Interface)

```

user@host> monitor interface et-9/0/0
Interface: et-9/0/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 100000mbps
Traffic statistics:
Input bytes: 0 (0 bps) [0]
Output bytes: 0 (0 bps) [0]
Input packets: 0 (0 pps) [0]
Output packets: 0 (0 pps) [0]
Error statistics:
Input errors: 0 [0]
Input drops: 0 [0]
Input framing errors: 0 [0]
Policed discards: 0 [0]
L3 incompletes: 0 [0]
L2 channel errors: 0 [0]
L2 mismatch timeouts: 0 [0]
Carrier transitions: 1 [0]
Output errors: 0 [0]
Output drops: 0 [0]
Aged packets: 0 [0]

```

Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'

### monitor interface (MX2020 Router with MPC6E and 10-Gigabit Ethernet OTN Interface)

```

user@host> monitor interface xe-3/0/0
host name Seconds: 67 Time: 23:46:46
 Delay: 0/0/13

Interface: xe-3/0/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 10000mbps
Traffic statistics: Current delta
Input bytes: 0 (0 bps) [0]
Output bytes: 0 (0 bps) [0]
Input packets: 0 (0 pps) [0]
Output packets: 0 (0 pps) [0]
Error statistics:
Input errors: 0 [0]
Input drops: 0 [0]
Input framing errors: 0 [0]
Policed discards: 0 [0]
L3 incompletes: 0 [0]
L2 channel errors: 0 [0]
L2 mismatch timeouts: 0 [0]
Carrier transitions: 3 [0]
Output errors: 0 [0]
Output drops: 0 [0]
Aged packets: 0 [0]
OTN Link 0
OTN Alarms:
OTN Defects:
OTN OC - Seconds
 LOS 0 [0]
 LOF 0 [0]
OTN OTU - FEC Statistics
 Corr err ratio N/A
 Corr bytes 0 [0]
 Uncorr words 0 [0]
OTN OTU - Counters
 BIP 0 [0]
 BBE 0 [0]
 ES 0 [0]
 SES 0 [0]
 UAS 0 [0]
OTN ODU - Counters
 BIP 0 [0]
 BBE 0 [0]
 ES 0 [0]
 SES 0 [0]
 UAS 0 [0]
OTN ODU - Received Overhead
 APSPCC 0-3: 00 00 00 00

```

Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'

### monitor interface (Logical)

```

user@host> monitor interface so-1/0/0.0

```

```

host name Seconds: 16 Time: 15:33:39
 Delay: 0/0/1

Interface: so-1/0/0.0, Enabled, Link is Down
Flags: Hardware-Down Point-To-Point SNMP-Traps
Encapsulation: PPP
Local statistics:
 Input bytes: 0 [0]
 Output bytes: 0 [0]
 Input packets: 0 [0]
 Output packets: 0 [0]
Remote statistics:
 Input bytes: 0 (0 bps) [0]
 Output bytes: 0 (0 bps) [0]
 Input packets: 0 (0 pps) [0]
 Output packets: 0 (0 pps) [0]
Traffic statistics:
 Destination address: 192.168.8.193, Local: 192.168.8.21

Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'

```

### monitor interface (QFX3500 Switch)

```

user@switch> monitor interface ge-0/0/0
Interface: ge-0/0/0, Enabled, Link is Down
Encapsulation: Ethernet, Speed: Unspecified
Traffic statistics:
 Input bytes: 0 (0 bps) [0]
 Output bytes: 0 (0 bps) [0]
 Input packets: 0 (0 pps) [0]
 Output packets: 0 (0 pps) [0]
Error statistics:
 Input errors: 0 [0]
 Input drops: 0 [0]
 Input framing errors: 0 [0]
 Policed discards: 0 [0]
 L3 incompletes: 0 [0]
 L2 channel errors: 0 [0]
 L2 mismatch timeouts: 0 [0]
 Carrier transitions: 0 [0]
 Output errors: 0 [0]
 Output drops: 0 [0]
 Aged packets: 0 [0]
Active alarms : LINK
Active defects: LINK
Input MAC/Filter statistics:
 Unicast packets 0 [0]
 Broadcast packets 0 Multicast packet [0]

Interface warnings:
 o Outstanding LINK alarm

```

### monitor interface traffic

```

user@host> monitor interface traffic
host name Seconds: 15 Time: 12:31:09

Interface Link Input packets (pps) Output packets (pps)
so-1/0/0 Down 0 (0) 0 (0)
so-1/1/0 Down 0 (0) 0 (0)
so-1/1/1 Down 0 (0) 0 (0)
so-1/1/2 Down 0 (0) 0 (0)

```

|          |      |        |     |       |     |
|----------|------|--------|-----|-------|-----|
| so-1/1/3 | Down | 0      | (0) | 0     | (0) |
| t3-1/2/0 | Down | 0      | (0) | 0     | (0) |
| t3-1/2/1 | Down | 0      | (0) | 0     | (0) |
| t3-1/2/2 | Down | 0      | (0) | 0     | (0) |
| t3-1/2/3 | Down | 0      | (0) | 0     | (0) |
| so-2/0/0 | Up   | 211035 | (1) | 36778 | (0) |
| so-2/0/1 | Up   | 192753 | (1) | 36782 | (0) |
| so-2/0/2 | Up   | 211020 | (1) | 36779 | (0) |
| so-2/0/3 | Up   | 211029 | (1) | 36776 | (0) |
| so-2/1/0 | Up   | 189378 | (1) | 36349 | (0) |
| so-2/1/1 | Down | 0      | (0) | 18747 | (0) |
| so-2/1/2 | Down | 0      | (0) | 16078 | (0) |
| so-2/1/3 | Up   | 0      | (0) | 80338 | (0) |
| at-2/3/0 | Up   | 0      | (0) | 0     | (0) |
| at-2/3/1 | Down | 0      | (0) | 0     | (0) |

Bytes=b, Clear=c, Delta=d, Packets=p, Quit=q or ESC, Rate=r, Up=^U, Down=^D

### monitor interface traffic (QFX3500 Switch)

```
user@switch> monitor interface traffic
switch
```

Seconds: 7                      Time: 16:04:37

| Interface | Link | Input packets | (pps) | Output packets | (pps) |
|-----------|------|---------------|-------|----------------|-------|
| ge-0/0/0  | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/1  | Up   | 392187        | (0)   | 392170         | (0)   |
| ge-0/0/2  | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/3  | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/4  | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/5  | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/6  | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/7  | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/8  | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/9  | Up   | 392184        | (0)   | 392171         | (0)   |
| ge-0/0/10 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/11 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/12 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/13 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/14 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/15 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/16 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/17 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/18 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/19 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/20 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/21 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/22 | Up   | 392172        | (0)   | 392187         | (0)   |
| ge-0/0/23 | Up   | 392185        | (0)   | 392173         | (0)   |
| vcp-0     | Down | 0             |       | 0              |       |
| vcp-1     | Down | 0             |       | 0              |       |
| ae0       | Down | 0             | (0)   | 0              | (0)   |
| bme0      | Up   | 0             |       | 1568706        |       |

### monitor interface traffic detail (QFX3500 Switch)

```
user@switch> monitor interface traffic detail
switch
```

Seconds: 74                      Time: 16:03:02

| Interface   | Link | Input packets | (pps) | Output packets | (pps) |
|-------------|------|---------------|-------|----------------|-------|
| Description |      |               |       |                |       |

|           |      |        |     |         |     |
|-----------|------|--------|-----|---------|-----|
| ge-0/0/0  | Down | 0      | (0) | 0       | (0) |
| ge-0/0/1  | Up   | 392183 | (0) | 392166  | (0) |
| ge-0/0/2  | Down | 0      | (0) | 0       | (0) |
| ge-0/0/3  | Down | 0      | (0) | 0       | (0) |
| ge-0/0/4  | Down | 0      | (0) | 0       | (0) |
| ge-0/0/5  | Down | 0      | (0) | 0       | (0) |
| ge-0/0/6  | Down | 0      | (0) | 0       | (0) |
| ge-0/0/7  | Down | 0      | (0) | 0       | (0) |
| ge-0/0/8  | Down | 0      | (0) | 0       | (0) |
| ge-0/0/9  | Up   | 392181 | (0) | 392168  | (0) |
| ge-0/0/10 | Down | 0      | (0) | 0       | (0) |
| ge-0/0/11 | Down | 0      | (0) | 0       | (0) |
| ge-0/0/12 | Down | 0      | (0) | 0       | (0) |
| ge-0/0/13 | Down | 0      | (0) | 0       | (0) |
| ge-0/0/14 | Down | 0      | (0) | 0       | (0) |
| ge-0/0/15 | Down | 0      | (0) | 0       | (0) |
| ge-0/0/16 | Down | 0      | (0) | 0       | (0) |
| ge-0/0/17 | Down | 0      | (0) | 0       | (0) |
| ge-0/0/18 | Down | 0      | (0) | 0       | (0) |
| ge-0/0/19 | Down | 0      | (0) | 0       | (0) |
| ge-0/0/20 | Down | 0      | (0) | 0       | (0) |
| ge-0/0/21 | Down | 0      | (0) | 0       | (0) |
| ge-0/0/22 | Up   | 392169 | (0) | 392184  | (1) |
| ge-0/0/23 | Up   | 392182 | (0) | 392170  | (0) |
| vcp-0     | Down | 0      |     | 0       |     |
| vcp-1     | Down | 0      |     | 0       |     |
| ae0       | Down | 0      | (0) | 0       | (0) |
| bme0      | Up   | 0      |     | 1568693 |     |

## show interfaces diagnostics optics

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show interfaces diagnostics optics <i>interface-name</i></code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | <p>Display diagnostics data and alarms for Gigabit Ethernet optical transceivers (SFP, SFP+, XFP, QSFP+, or CFP) installed in EX Series or QFX Series switches. The information provided by this command is known as digital optical monitoring (DOM) information. For a list of transceivers supported on EX Series switches and their specifications, including DOM support, see <i>Pluggable Transceivers Supported on EX Series Switches</i>.</p> <p>Thresholds that trigger a high alarm, low alarm, high warning, or low warning are set by the transponder vendors. Generally, a high alarm or low alarm indicates that the optics module is not operating properly. This information can be used to diagnose why a transceiver is not working.</p> |
| <b>Options</b>                  | <i>interface-name</i> —Name of the interface associated with the port in which the transceiver is installed: <i>ge-fpc/pic/port</i> , <i>xe-fpc/pic/port</i> , or <i>et-fpc/pic/port</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Monitoring Interface Status and Traffic</i></li> <li>• <a href="#">Monitoring Interface Status and Traffic on page 167</a></li> <li>• <i>Installing a Transceiver in a Switch</i></li> <li>• <i>Installing a Transceiver in a QFX Series Device</i></li> <li>• <i>Removing a Transceiver from a Switch</i></li> <li>• <i>Removing a Transceiver from a QFX Series Device</i></li> <li>• <a href="#">Junos OS Ethernet Interfaces Configuration Guide</a></li> </ul>                                                                                                                                                                                                                                            |
| <b>List of Sample Output</b>    | <p><a href="#">show interfaces diagnostics optics ge-0/1/0 (SFP Transceiver) on page 3109</a></p> <p><a href="#">show interfaces diagnostics optics xe-0/1/0 (SFP+ Transceiver) on page 3110</a></p> <p><a href="#">show interfaces diagnostics optics xe-0/1/0 (XFP Transceiver) on page 3111</a></p> <p><a href="#">show interfaces diagnostics optics et-3/0/0 (QSFP+ Transceiver) on page 3112</a></p> <p><a href="#">show interfaces diagnostics optics et-4/1/0 (CFP Transceiver) on page 3113</a></p>                                                                                                                                                                                                                                               |
| <b>Output Fields</b>            | <p><a href="#">Table 276</a> lists the output fields for the <b>show interfaces diagnostics optics</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

**Table 276: show interfaces diagnostics optics Output Fields**

| Field Name         | Field Description                            |
|--------------------|----------------------------------------------|
| Physical interface | Displays the name of the physical interface. |

Table 276: show interfaces diagnostics optics Output Fields (*continued*)

| Field Name                                                                                           | Field Description                                                                                                                                                  |
|------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Laser bias current</b>                                                                            | Displays the magnitude of the laser bias power setting current, in milliamperes. The laser bias provides direct modulation of laser diodes and modulates currents. |
| <b>Laser output power</b><br>(Not available for QSFP+ transceivers)                                  | Displays the laser output power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).                                                                         |
| <b>Laser temperature</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)               | Displays the laser temperature, in Celsius and Fahrenheit.                                                                                                         |
| <b>Module temperature</b>                                                                            | Displays the temperature, in Celsius and Fahrenheit.                                                                                                               |
| <b>Module voltage</b><br>(Not available for XFP transceivers)                                        | Displays the voltage, in Volts.                                                                                                                                    |
| <b>Laser rx power</b><br>(Not available for SFP, SFP+, QSFP+, and CFP transceivers)                  | Displays the laser received optical power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).                                                               |
| <b>Receiver signal average optical power</b><br>(Not available for XFP, QSFP+, and CFP transceivers) | Displays the receiver signal average optical power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).                                                      |
| <b>Laser bias current high alarm</b>                                                                 | Displays whether the laser bias power setting high alarm is <b>On</b> or <b>Off</b> .                                                                              |
| <b>Laser bias current low alarm</b>                                                                  | Displays whether the laser bias power setting low alarm is <b>On</b> or <b>Off</b> .                                                                               |
| <b>Laser bias current high warning</b>                                                               | Displays whether the laser bias power setting high warning is <b>On</b> or <b>Off</b> .                                                                            |
| <b>Laser bias current low warning</b>                                                                | Displays whether the laser bias power setting low warning is <b>On</b> or <b>Off</b> .                                                                             |
| <b>Laser output power high alarm</b><br>(Not available for QSFP+ transceivers)                       | Displays whether the laser output power high alarm is <b>On</b> or <b>Off</b> .                                                                                    |
| <b>Laser output power low alarm</b><br>(Not available for QSFP+ transceivers)                        | Displays whether the laser output power low alarm is <b>On</b> or <b>Off</b> .                                                                                     |
| <b>Laser output power high warning</b><br>(Not available for QSFP+ transceivers)                     | Displays whether the laser output power high warning is <b>On</b> or <b>Off</b> .                                                                                  |

Table 276: show interfaces diagnostics optics Output Fields (*continued*)

| Field Name                                                                                          | Field Description                                                                 |
|-----------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------|
| <b>Laser output power low warning</b><br>(Not available for QSFP+ transceivers)                     | Displays whether the laser output power low warning is <b>On</b> or <b>Off</b> .  |
| <b>Laser temperature high alarm</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)   | Displays whether the laser temperature high alarm is <b>On</b> or <b>Off</b> .    |
| <b>Laser temperature low alarm</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)    | Displays whether the laser temperature low alarm is <b>On</b> or <b>Off</b> .     |
| <b>Laser temperature high warning</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers) | Displays whether the laser temperature high warning is <b>On</b> or <b>Off</b> .  |
| <b>Laser temperature low warning</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)  | Displays whether the laser temperature low warning is <b>On</b> or <b>Off</b> .   |
| <b>Module temperature high alarm</b><br>(Not available for QSFP+ transceivers)                      | Displays whether the module temperature high alarm is <b>On</b> or <b>Off</b> .   |
| <b>Module temperature low alarm</b><br>(Not available for QSFP+ transceivers)                       | Displays whether the module temperature low alarm is <b>On</b> or <b>Off</b> .    |
| <b>Module temperature high warning</b><br>(Not available for QSFP+ transceivers)                    | Displays whether the module temperature high warning is <b>On</b> or <b>Off</b> . |
| <b>Module temperature low warning</b><br>(Not available for QSFP+ transceivers)                     | Displays whether the module temperature low warning is <b>On</b> or <b>Off</b> .  |
| <b>Module voltage high alarm</b><br>(Not available for XFP and QSFP+ transceivers)                  | Displays whether the module voltage high alarm is <b>On</b> or <b>Off</b> .       |
| <b>Module voltage low alarm</b><br>(Not available for XFP and QSFP+ transceivers)                   | Displays whether the module voltage low alarm is <b>On</b> or <b>Off</b> .        |
| <b>Module voltage high warning</b><br>(Not available for XFP and QSFP+ transceivers)                | Displays whether the module voltage high warning is <b>On</b> or <b>Off</b> .     |
| <b>Module voltage low warning</b><br>(Not available for XFP and QSFP+ transceivers)                 | Displays whether the module voltage low warning is <b>On</b> or <b>Off</b> .      |



Table 276: show interfaces diagnostics optics Output Fields (*continued*)

| Field Name                                                                                                  | Field Description                                                                                                                            |
|-------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Laser rx power high alarm</b><br>(Not available for QSFP+ and CFP transceivers)                          | Displays whether the receive laser power high alarm is <b>On</b> or <b>Off</b> .                                                             |
| <b>Laser rx power low alarm</b><br>(Not available for QSFP+ and CFP transceivers)                           | Displays whether the receive laser power low alarm is <b>On</b> or <b>Off</b> .                                                              |
| <b>Laser rx power high warning</b><br>(Not available for QSFP+ and CFP transceivers)                        | Displays whether the receive laser power high warning is <b>On</b> or <b>Off</b> .                                                           |
| <b>Laser rx power low warning</b><br>(Not available for QSFP+ and CFP transceivers)                         | Displays whether the receive laser power low warning is <b>On</b> or <b>Off</b> .                                                            |
| <b>Laser bias current high alarm threshold</b><br>(Not available for QSFP+ transceivers)                    | Displays the vendor-specified threshold for the laser bias current high alarm.                                                               |
| <b>Module not ready alarm</b><br>(Not available for SFP, SFP+, and QSFP+ transceivers)                      | Displays whether the module not ready alarm is <b>On</b> or <b>Off</b> . When the output is <b>On</b> , the module has an operational fault. |
| <b>Module low power alarm</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)                 | Displays whether the module low power alarm is <b>On</b> or <b>Off</b> .                                                                     |
| <b>Module initialization incomplete alarm</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers) | Displays whether the module initialization incomplete alarm is <b>On</b> or <b>Off</b> .                                                     |
| <b>Module fault alarm</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)                     | Displays whether the module fault alarm is <b>On</b> or <b>Off</b> .                                                                         |
| <b>PLD Flash initialization fault alarm</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)   | Displays whether the PLD Flash initialization fault alarm is <b>On</b> or <b>Off</b> .                                                       |
| <b>Power supply fault alarm</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)               | Displays whether the power supply fault alarm is <b>On</b> or <b>Off</b> .                                                                   |
| <b>Checksum fault alarm</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)                   | Displays whether the checksum fault alarm is <b>On</b> or <b>Off</b> .                                                                       |
| <b>Tx laser disabled alarm</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)                | Displays whether the Tx laser disabled alarm is <b>On</b> or <b>Off</b> .                                                                    |

Table 276: show interfaces diagnostics optics Output Fields (*continued*)

| Field Name                                                                                   | Field Description                                                                                                                                                                                |
|----------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Module power down alarm</b><br>(Not available for SFP, SFP+, QSFP+, and CFP transceivers) | Displays whether the module power down alarm is <b>On</b> or <b>Off</b> . When the output is <b>On</b> , module is in a limited power mode, low for normal operation.                            |
| <b>Tx data not ready alarm</b><br>(Not available for SFP, SFP+, QSFP+, and CFP transceivers) | Any condition leading to invalid data on the transmit path. Displays whether the Tx data not ready alarm is <b>On</b> or <b>Off</b> .                                                            |
| <b>Tx not ready alarm</b><br>(Not available for SFP, SFP+, QSFP+, and CFP transceivers)      | Any condition leading to invalid data on the transmit path. Displays whether the Tx not ready alarm is <b>On</b> or <b>Off</b> .                                                                 |
| <b>Tx laser fault alarm</b><br>(Not available for SFP, SFP+, QSFP+, and CFP transceivers)    | Laser fault condition. Displays whether the Tx laser fault alarm is <b>On</b> or <b>Off</b> .                                                                                                    |
| <b>Tx CDR loss of lock alarm</b><br>(Not available for SFP, SFP+, and QSFP+ transceivers)    | Transmit clock and data recovery (CDR) loss of lock. Loss of lock on the transmit side of the CDR. Displays whether the Tx CDR loss of lock alarm is <b>On</b> or <b>Off</b> .                   |
| <b>Rx not ready alarm</b><br>(Not available for SFP, SFP+, QSFP+, and CFP transceivers)      | Any condition leading to invalid data on the receive path. Displays whether the Rx not ready alarm is <b>On</b> or <b>Off</b> .                                                                  |
| <b>Rx loss of signal alarm</b><br>(Not available for SFP and SFP+ transceivers)              | Receive loss of signal alarm. When the output is <b>On</b> , indicates insufficient optical input power to the module. Displays whether the Rx loss of signal alarm is <b>On</b> or <b>Off</b> . |
| <b>Rx CDR loss of lock alarm</b><br>(Not available for SFP, SFP+, and QSFP+ transceivers)    | Receive CDR loss of lock. Loss of lock on the receive side of the CDR. Displays whether the Rx CDR loss of lock alarm is <b>On</b> or <b>Off</b> .                                               |
| <b>Laser bias current low alarm threshold</b><br>(Not available for QSFP+ transceivers)      | Displays the vendor-specified threshold for the laser bias current low alarm.                                                                                                                    |
| <b>Laser bias current high warning threshold</b><br>(Not available for QSFP+ transceivers)   | Displays the vendor-specified threshold for the laser bias current high warning.                                                                                                                 |
| <b>Laser bias current low warning threshold</b><br>(Not available for QSFP+ transceivers)    | Displays the vendor-specified threshold for the laser bias current low warning.                                                                                                                  |
| <b>Laser output power high alarm threshold</b><br>(Not available for QSFP+ transceivers)     | Displays the vendor-specified threshold for the laser output power high alarm.                                                                                                                   |
| <b>Laser output power low alarm threshold</b><br>(Not available for QSFP+ transceivers)      | Displays the vendor-specified threshold for the laser output power low alarm.                                                                                                                    |

Table 276: show interfaces diagnostics optics Output Fields (*continued*)

| Field Name                                                                                     | Field Description                                                                |
|------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------|
| <b>Laser output power high warning threshold</b><br>(Not available for QSFP+ transceivers)     | Displays the vendor-specified threshold for the laser output power high warning. |
| <b>Laser output power low warning threshold</b><br>(Not available for QSFP+ transceivers)      | Displays the vendor-specified threshold for the laser output power low warning.  |
| <b>Module temperature high alarm threshold</b><br>(Not available for QSFP+ transceivers)       | Displays the vendor-specified threshold for the module temperature high alarm.   |
| <b>Module temperature low alarm threshold</b><br>(Not available for QSFP+ transceivers)        | Displays the vendor-specified threshold for the module temperature low alarm.    |
| <b>Module temperature high warning threshold</b><br>(Not available for QSFP+ transceivers)     | Displays the vendor-specified threshold for the module temperature high warning. |
| <b>Module temperature low warning threshold</b><br>(Not available for QSFP+ transceivers)      | Displays the vendor-specified threshold for the module temperature low warning.  |
| <b>Module voltage high alarm threshold</b><br>(Not available for XFP and QSFP+ transceivers)   | Displays the vendor-specified threshold for the module voltage high alarm.       |
| <b>Module voltage low alarm threshold</b><br>(Not available for XFP and QSFP+ transceivers)    | Displays the vendor-specified threshold for the module voltage low alarm.        |
| <b>Module voltage high warning threshold</b><br>(Not available for XFP and QSFP+ transceivers) | Displays the vendor-specified threshold for the module voltage high warning.     |
| <b>Module voltage low warning threshold</b><br>(Not available for XFP and QSFP+ transceivers)  | Displays the vendor-specified threshold for the module voltage low warning.      |
| <b>Laser rx power high alarm threshold</b><br>(Not available for QSFP+ transceivers)           | Displays the vendor-specified threshold for the laser rx power high alarm.       |
| <b>Laser rx power low alarm threshold</b><br>(Not available for QSFP+ transceivers)            | Displays the vendor-specified threshold for the laser rx power low alarm.        |
| <b>Laser rx power high warning threshold</b><br>(Not available for QSFP+ transceivers)         | Displays the vendor-specified threshold for the laser rx power high warning.     |

Table 276: show interfaces diagnostics optics Output Fields (*continued*)

| Field Name                                                                                                    | Field Description                                                                                          |
|---------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------|
| <b>Laser rx power low warning threshold</b><br>(Not available for QSFP+ transceivers)                         | Displays the vendor-specified threshold for the laser rx power low warning.                                |
| <b>Laser temperature high alarm threshold</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)   | Displays the vendor-specified threshold for the laser temperature high alarm, in Celsius and Fahrenheit.   |
| <b>Laser temperature low alarm threshold</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)    | Displays the vendor-specified threshold for the laser temperature low alarm, in Celsius and Fahrenheit.    |
| <b>Laser temperature high warning threshold</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers) | Displays the vendor-specified threshold for the laser temperature high warning, in Celsius and Fahrenheit. |
| <b>Laser temperature low warning threshold</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)  | Displays the vendor-specified threshold for the laser temperature low warning, in Celsius and Fahrenheit.  |
| <b>SOA bias current high alarm threshold</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)    | Displays the vendor-specified threshold for SOA bias current high alarm.                                   |
| <b>SOA bias current low alarm threshold</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)     | Displays the vendor-specified threshold for SOA bias current low alarm.                                    |
| <b>SOA bias current high warning threshold</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)  | Displays the vendor-specified threshold for SOA bias current high warning.                                 |
| <b>SOA bias current low warning threshold</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)   | Displays the vendor-specified threshold for SOA bias current low warning.                                  |
| <b>Laser receiver power high alarm</b><br>(Not available for SFP, SFP+, and XFP transceivers)                 | Displays whether the laser receiver power high alarm is <b>On</b> or <b>Off</b> .                          |
| <b>Laser receiver power low alarm</b><br>(Not available for SFP, SFP+, and XFP transceivers)                  | Displays whether the laser receiver power low alarm is <b>On</b> or <b>Off</b> .                           |
| <b>Laser receiver power high warning</b><br>(Not available for SFP, SFP+, and XFP transceivers)               | Displays whether the laser receiver power high warning is <b>On</b> or <b>Off</b> .                        |
| <b>Laser receiver power low warning</b><br>(Not available for SFP, SFP+, and XFP transceivers)                | Displays whether the laser receiver power low warning is <b>On</b> or <b>Off</b> .                         |

Table 276: show interfaces diagnostics optics Output Fields (*continued*)

| Field Name                                                                                          | Field Description                                                                            |
|-----------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------|
| <b>Laser receiver power</b><br>(Not available for SFP, SFP+, and XFP transceivers)                  | Displays the laser receiver power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm). |
| <b>Tx loss of signal functionality alarm</b><br>(Not available for SFP, SFP+, and XFP transceivers) | Displays whether the Tx loss of signal functionality alarm is <b>On</b> or <b>Off</b> .      |
| <b>APD supply fault alarm</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)         | Displays whether the APD supply fault alarm is <b>On</b> or <b>Off</b> .                     |
| <b>TEC fault alarm</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)                | Displays whether the TEC fault alarm is <b>On</b> or <b>Off</b> .                            |
| <b>Wavelength unlocked alarm</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)      | Displays whether the Wavelength unlocked alarm is <b>On</b> or <b>Off</b> .                  |

## Sample Output

### show interfaces diagnostics optics ge-0/1/0 (SFP Transceiver)

```

user@switch> show interfaces diagnostics optics ge-0/1/0
Physical interface: ge-0/1/0
 Laser bias current : 5.444 mA
 Laser output power : 0.3130 mW / -5.04 dBm
 Module temperature : 36 degrees C / 97 degrees F
 Module voltage : 3.2120 V
 Receiver signal average optical power : 0.3840 mW / -4.16 dBm
 Laser bias current high alarm : Off
 Laser bias current low alarm : Off
 Laser bias current high warning : Off
 Laser bias current low warning : Off
 Laser output power high alarm : Off
 Laser output power low alarm : Off
 Laser output power high warning : Off
 Laser output power low warning : Off
 Module temperature high alarm : Off
 Module temperature low alarm : Off
 Module temperature high warning : Off
 Module temperature low warning : Off
 Module voltage high alarm : Off
 Module voltage low alarm : Off
 Module voltage high warning : Off
 Module voltage low warning : Off
 Laser rx power high alarm : Off
 Laser rx power low alarm : Off
 Laser rx power high warning : Off
 Laser rx power low warning : Off
 Laser bias current high alarm threshold : 15.000 mA
 Laser bias current low alarm threshold : 1.000 mA
 Laser bias current high warning threshold : 12.000 mA

```

```

Laser bias current low warning threshold : 2.000 mA
Laser output power high alarm threshold : 0.6300 mW / -2.01 dBm
Laser output power low alarm threshold : 0.0660 mW / -11.80 dBm
Laser output power high warning threshold : 0.6300 mW / -2.01 dBm
Laser output power low warning threshold : 0.0780 mW / -11.08 dBm
Module temperature high alarm threshold : 109 degrees C / 228 degrees F
Module temperature low alarm threshold : -29 degrees C / -20 degrees F
Module temperature high warning threshold : 103 degrees C / 217 degrees F
Module temperature low warning threshold : -13 degrees C / 9 degrees F
Module voltage high alarm threshold : 3.900 V
Module voltage low alarm threshold : 2.700 V
Module voltage high warning threshold : 3.700 V
Module voltage low warning threshold : 2.900 V
Laser rx power high alarm threshold : 1.2589 mW / 1.00 dBm
Laser rx power low alarm threshold : 0.0100 mW / -20.00 dBm
Laser rx power high warning threshold : 0.7939 mW / -1.00 dBm
Laser rx power low warning threshold : 0.0157 mW / -18.04 dBm

```

## Sample Output

### show interfaces diagnostics optics xe-0/1/0 (SFP+ Transceiver)

```

user@switch> show interfaces diagnostics optics xe-0/1/0
Physical interface: xe-0/1/0
 Laser bias current : 4.968 mA
 Laser output power : 0.4940 mW / -3.06 dBm
 Module temperature : 27 degrees C / 81 degrees F
 Module voltage : 3.2310 V
 Receiver signal average optical power : 0.0000
 Laser bias current high alarm : Off
 Laser bias current low alarm : Off
 Laser bias current high warning : Off
 Laser bias current low warning : Off
 Laser output power high alarm : Off
 Laser output power low alarm : Off
 Laser output power high warning : Off
 Laser output power low warning : Off
 Module temperature high alarm : Off
 Module temperature low alarm : Off
 Module temperature high warning : Off
 Module temperature low warning : Off
 Module voltage high alarm : Off
 Module voltage low alarm : Off
 Module voltage high warning : Off
 Module voltage low warning : Off
 Laser rx power high alarm : Off
 Laser rx power low alarm : On
 Laser rx power high warning : Off
 Laser rx power low warning : On
 Laser bias current high alarm threshold : 10.500 mA
 Laser bias current low alarm threshold : 2.000 mA
 Laser bias current high warning threshold : 9.000 mA
 Laser bias current low warning threshold : 2.500 mA
 Laser output power high alarm threshold : 1.4120 mW / 1.50 dBm
 Laser output power low alarm threshold : 0.0740 mW / -11.31 dBm
 Laser output power high warning threshold : 0.7070 mW / -1.51 dBm
 Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
 Module temperature high alarm threshold : 75 degrees C / 167 degrees F
 Module temperature low alarm threshold : -5 degrees C / 23 degrees F
 Module temperature high warning threshold : 70 degrees C / 158 degrees F
 Module temperature low warning threshold : 0 degrees C / 32 degrees F

```

```

Module voltage high alarm threshold : 3.630 V
Module voltage low alarm threshold : 2.970 V
Module voltage high warning threshold : 3.465 V
Module voltage low warning threshold : 3.135 V
Laser rx power high alarm threshold : 1.5849 mW / 2.00 dBm
Laser rx power low alarm threshold : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold : 0.1023 mW / -9.90 dBm

```

## Sample Output

### show interfaces diagnostics optics xe-0/1/0 (XFP Transceiver)

```
user@switch> show interfaces diagnostics optics xe-0/1/0
```

```
Physical interface: xe-0/1/0
```

```

Laser bias current : 8.029 mA
Laser output power : 0.6430 mW / -1.92 dBm
Module temperature : 4 degrees C / 39 degrees F
Laser rx power : 0.0012 mW / -29.21 dBm
Laser bias current high alarm : Off
Laser bias current low alarm : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser output power high alarm : Off
Laser output power low alarm : Off
Laser output power high warning : Off
Laser output power low warning : Off
Module temperature high alarm : Off
Module temperature low alarm : Off
Module temperature high warning : Off
Module temperature low warning : Off
Laser rx power high alarm : Off
Laser rx power low alarm : On
Laser rx power high warning : Off
Laser rx power low warning : On
Module not ready alarm : On
Module power down alarm : Off
Tx data not ready alarm : Off
Tx not ready alarm : Off
Tx laser fault alarm : Off
Tx CDR loss of lock alarm : Off
Rx not ready alarm : On
Rx loss of signal alarm : On
Rx CDR loss of lock alarm : On
Laser bias current high alarm threshold : 13.000 mA
Laser bias current low alarm threshold : 2.000 mA
Laser bias current high warning threshold : 12.000 mA
Laser bias current low warning threshold : 3.000 mA
Laser output power high alarm threshold : 0.8310 mW / -0.80 dBm
Laser output power low alarm threshold : 0.1650 mW / -7.83 dBm
Laser output power high warning threshold : 0.7410 mW / -1.30 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold : 90 degrees C / 194 degrees F
Module temperature low alarm threshold : 0 degrees C / 32 degrees F
Module temperature high warning threshold : 85 degrees C / 185 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Laser rx power high alarm threshold : 0.8912 mW / -0.50 dBm
Laser rx power low alarm threshold : 0.0912 mW / -10.40 dBm
Laser rx power high warning threshold : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold : 0.1023 mW / -9.90 dBm

```

## Sample Output

### show interfaces diagnostics optics et-3/0/0 (QSFP+ Transceiver)

```

user@switch> show interfaces diagnostics optics et-3/0/0
Physical interface: et-3/0/0
 Module temperature : 33 degrees C / 92 degrees F
 Module voltage : 3.3060 V
Lane 0
 Laser bias current : 7.182 mA
 Laser receiver power : 0.743 mW / -1.29 dBm
 Laser bias current high alarm : Off
 Laser bias current low alarm : Off
 Laser bias current high warning : Off
 Laser bias current low warning : Off
 Laser receiver power high alarm : Off
 Laser receiver power low alarm : Off
 Laser receiver power high warning : Off
 Laser receiver power low warning : Off
 Tx loss of signal functionality alarm : Off
 Rx loss of signal alarm : Off
Lane 1
 Laser bias current : 7.326 mA
 Laser receiver power : 0.752 mW / -1.24 dBm
 Laser bias current high alarm : Off
 Laser bias current low alarm : Off
 Laser bias current high warning : Off
 Laser bias current low warning : Off
 Laser receiver power high alarm : Off
 Laser receiver power low alarm : Off
 Laser receiver power high warning : Off
 Laser receiver power low warning : Off
 Tx loss of signal functionality alarm : Off
 Rx loss of signal alarm : Off
Lane 2
 Laser bias current : 7.447 mA
 Laser receiver power : 0.790 mW / -1.03 dBm
 Laser bias current high alarm : Off
 Laser bias current low alarm : Off
 Laser bias current high warning : Off
 Laser bias current low warning : Off
 Laser receiver power high alarm : Off
 Laser receiver power low alarm : Off
 Laser receiver power high warning : Off
 Laser receiver power low warning : Off
 Tx loss of signal functionality alarm : Off
 Rx loss of signal alarm : Off
Lane 3
 Laser bias current : 7.734 mA
 Laser receiver power : 0.768 mW / -1.15 dBm
 Laser bias current high alarm : Off
 Laser bias current low alarm : Off
 Laser bias current high warning : Off
 Laser bias current low warning : Off
 Laser receiver power high alarm : Off
 Laser receiver power low alarm : Off
 Laser receiver power high warning : Off
 Laser receiver power low warning : Off
 Tx loss of signal functionality alarm : Off
 Rx loss of signal alarm : Off

```



## Sample Output

### show interfaces diagnostics optics et-4/1/0 (CFP Transceiver)

```

user@switch> show interfaces diagnostics optics et-4/1/0
Physical interface: et-4/1/0
 Module temperature : 38 degrees C / 101 degrees F
 Module voltage : 3.2500 V
 Module temperature high alarm : Off
 Module temperature low alarm : Off
 Module temperature high warning : Off
 Module temperature low warning : Off
 Module voltage high alarm : Off
 Module voltage low alarm : Off
 Module voltage high warning : Off
 Module voltage low warning : Off
 Module not ready alarm : Off
 Module low power alarm : Off
 Module initialization incomplete alarm : Off
 Module fault alarm : Off
 PLD Flash initialization fault alarm : Off
 Power supply fault alarm : Off
 Checksum fault alarm : Off
 Tx laser disabled alarm : Off
 Tx loss of signal functionality alarm : Off
 Tx CDR loss of lock alarm : Off
 Rx loss of signal alarm : Off
 Rx CDR loss of lock alarm : Off
 Module temperature high alarm threshold : 75 degrees C / 167 degrees F
 Module temperature low alarm threshold : -5 degrees C / 23 degrees F
 Module temperature high warning threshold : 70 degrees C / 158 degrees F
 Module temperature low warning threshold : 0 degrees C / 32 degrees F
 Module voltage high alarm threshold : 3.5000 V
 Module voltage low alarm threshold : 3.0990 V
 Module voltage high warning threshold : 3.4000 V
 Module voltage low warning threshold : 3.2000 V
 Laser bias current high alarm threshold : 250.000 mA
 Laser bias current low alarm threshold : 37.500 mA
 Laser bias current high warning threshold : 225.000 mA
 Laser bias current low warning threshold : 50.000 mA
 Laser output power high alarm threshold : 3.9800 mW / 6.00 dBm
 Laser output power low alarm threshold : 0.4670 mW / -3.31 dBm
 Laser output power high warning threshold : 3.5480 mW / 5.50 dBm
 Laser output power low warning threshold : 0.5240 mW / -2.81 dBm
 Laser rx power high alarm threshold : 3.5481 mW / 5.50 dBm
 Laser rx power low alarm threshold : 0.0616 mW / -12.10 dBm
 Laser rx power high warning threshold : 3.1622 mW / 5.00 dBm
 Laser rx power low warning threshold : 0.0691 mW / -11.61 dBm
 Laser temperature high alarm threshold : 67 degrees C / 153 degrees F
 Laser temperature low alarm threshold : 35 degrees C / 95 degrees F
 Laser temperature high warning threshold : 62 degrees C / 144 degrees F
 Laser temperature low warning threshold : 40 degrees C / 104 degrees F
 SOA bias current high alarm threshold : 0.000 mA
 SOA bias current low alarm threshold : 0.000 mA
 SOA bias current high warning threshold : 0.000 mA
 SOA bias current low warning threshold : 0.000 mA
Lane 0
 Laser bias current : 131.684 mA
 Laser output power : 1.002 mW / 0.01 dBm
 Laser temperature : 54 degrees C / 128 degrees F
 Laser receiver power : 0.497 mW / -3.03 dBm

```

```

Laser bias current high alarm : Off
Laser bias current low alarm : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser output power high alarm : Off
Laser output power low alarm : Off
Laser output power high warning : Off
Laser output power low warning : Off
Laser temperature high alarm : Off
Laser temperature low alarm : Off
Laser temperature high warning : Off
Laser temperature low warning : Off
Laser receiver power high alarm : Off
Laser receiver power low alarm : Off
Laser receiver power high warning : Off
Laser receiver power low warning : Off
Tx loss of signal functionality alarm : Off
Rx CDR loss of lock alarm : Off
Rx loss of signal alarm : Off
Rx CDR loss of lock alarm : Off
APD supply fault alarm : Off
TEC fault alarm : Off
Wavelength unlocked alarm : Off

Lane 1
Laser bias current : 122.345 mA
Laser output power : 1.002 mW / 0.01 dBm
Laser temperature : 51 degrees C / 124 degrees F
Laser receiver power : 0.611 mW / -2.14 dBm
Laser bias current high alarm : Off
Laser bias current low alarm : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser output power high alarm : Off
Laser output power low alarm : Off
Laser output power high warning : Off
Laser output power low warning : Off
Laser temperature high alarm : Off
Laser temperature low alarm : Off
Laser temperature high warning : Off
Laser temperature low warning : Off
Laser receiver power high alarm : Off
Laser receiver power low alarm : Off
Laser receiver power high warning : Off
Laser receiver power low warning : Off
Tx loss of signal functionality alarm : Off
Tx CDR loss of lock alarm : Off
Rx loss of signal alarm : Off
Rx CDR loss of lock alarm : Off
APD supply fault alarm : Off
TEC fault alarm : Off
Wavelength unlocked alarm : Off

Lane 2
Laser bias current : 112.819 mA
Laser output power : 1.000 mW / 0.00 dBm
Laser temperature : 50 degrees C / 122 degrees F
Laser receiver power : 0.540 mW / -2.67 dBm
Laser bias current high alarm : Off
Laser bias current low alarm : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser output power high alarm : Off

```

```

Laser output power low alarm : Off
Laser output power high warning : Off
Laser output power low warning : Off
Laser temperature high alarm : Off
Laser temperature low alarm : Off
Laser temperature high warning : Off
Laser temperature low warning : Off
Laser receiver power high alarm : Off
Laser receiver power low alarm : Off
Laser receiver power high warning : Off
Laser receiver power low warning : Off
Tx loss of signal functionality alarm : Off
Tx CDR loss of lock alarm : Off
Rx loss of signal alarm : Off
Rx CDR loss of lock alarm : Off
APD supply fault alarm : Off
TEC fault alarm : Off
Wavelength unlocked alarm : Off

Lane 3
Laser bias current : 100.735 mA
Laser output power : 1.002 mW / 0.01 dBm
Laser temperature : 50 degrees C / 122 degrees F
Laser receiver power : 0.637 mW / -1.96 dBm
Laser bias current high alarm : Off
Laser bias current low alarm : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser output power high alarm : Off
Laser output power low alarm : Off
Laser output power high warning : Off
Laser output power low warning : Off
Laser temperature high alarm : Off
Laser temperature low alarm : Off
Laser temperature high warning : Off
Laser temperature low warning : Off
Laser receiver power high alarm : Off
Laser receiver power low alarm : Off
Laser receiver power high warning : Off
Laser receiver power low warning : Off
Tx loss of signal functionality alarm : Off
Tx CDR loss of lock alarm : Off
Rx loss of signal alarm : Off
Rx CDR loss of lock alarm : Off
APD supply fault alarm : Off
TEC fault alarm : Off
Wavelength unlocked alarm : Off

```

## show interfaces ge

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show interfaces <i>device-name:type-fpc/pic/port</i></code><br><code>&lt;brief   detail   extensive   terse&gt;</code><br><code>&lt;descriptions&gt;</code><br><code>&lt;media&gt;</code><br><code>&lt;routing-instance (all   <i>instance-name</i>)&gt;</code><br><code>&lt;snmp-index <i>snmp-index</i>&gt;</code><br><code>&lt;statistics&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.<br>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Display status information about the specified Gigabit Ethernet interface. This command does not display statistics for routed VLAN interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output.</p> <p><b><i>device-name:type-fpc/pic/port</i></b>—The device name is either the serial number or the alias of the QFabric system component, such as a Node device, Interconnect device, or QFabric infrastructure. The name can contain a maximum of 128 characters and cannot contain any colons.</p> <p><b>descriptions</b>—(Optional) Display interface description strings.</p> <p><b>media</b>—(Optional) Display media-specific information about network interfaces.</p> <p><b>routing instance (all   <i>instance-name</i>)</b>—(Optional) Display the name of an individual routing-instance or display all routing-instances.</p> <p><b>snmp-index <i>snmp-index</i></b>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><b>statistics</b>—(Optional) Display static interface statistics.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Monitoring Interface Status and Traffic on page 167</a></li><li>• <a href="#">Troubleshooting Network Interfaces on page 2845</a></li><li>• <a href="#">Troubleshooting an Aggregated Ethernet Interface on page 2886</a></li><li>• <a href="#">Junos OS Network Interfaces Library for Routing Devices</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>List of Sample Output</b>    | <a href="#">show interfaces on page 3124</a><br><a href="#">show interfaces brief on page 3124</a><br><a href="#">show interfaces detail (Symmetric Flow Control and Autonegotiation Enabled) on page 3124</a><br><a href="#">show interfaces detail (Asymmetric Flow Control and Autonegotiation Enabled) on page 3125</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

[show interfaces extensive \(Symmetric Flow Control and Autonegotiation Enabled\) on page 3126](#)

[show interfaces extensive \(Asymmetric Flow Control and Autonegotiation Enabled\) on page 3128](#)

[show interfaces terse on page 3130](#)

[show interfaces terse \(QFabric Systems\) on page 3130](#)

**Output Fields** Table 277 lists the output fields for the **show interfaces ge** command. Output fields are listed in the approximate order in which they appear.

**Table 277: show interfaces ge Output Fields**

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Level of Output               |
|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------|
| <b>Physical Interface</b>      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |                               |
| <b>Physical interface</b>      | Name of the physical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | All levels                    |
| <b>Enabled</b>                 | State of the interface: <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | All levels                    |
| <b>Interface index</b>         | Index number of the physical interface, which reflects its initialization sequence.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail extensive none</b>  |
| <b>SNMP ifIndex</b>            | SNMP index number for the physical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive none</b>  |
| <b>Generation</b>              | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail extensive</b>       |
| <b>Description</b>             | Optional user-specified description.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>brief detail extensive</b> |
| <b>Link-level type</b>         | Encapsulation being used on the physical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | All levels                    |
| <b>MTU</b>                     | Maximum transmission unit size on the physical interface. The default is 1514.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | All levels                    |
| <b>Speed</b>                   | Speed at which the interface is running.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | All levels                    |
| <b>Loopback</b>                | Loopback status: <b>Enabled</b> or <b>Disabled</b> . If loopback is enabled, type of loopback: <b>Local</b> or <b>Remote</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | All levels                    |
| <b>Source filtering</b>        | Source filtering status: <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | All levels                    |
| <b>Flow control</b>            | Flow control status: <b>Enabled</b> or <b>Disabled</b> .<br><br><i>NOTE:</i> This field is only displayed if asymmetric flow control is not configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail extensive</b>       |
| <b>Configured-flow-control</b> | Configured flow control for the interface transmit buffers ( <b>tx-buffers</b> ) and receive buffers ( <b>rx-buffers</b> ):<br><br><ul style="list-style-type: none"> <li><b>tx-buffers</b>—<b>On</b> if the interface is configured to respond to Ethernet PAUSE messages received from the connected peer.<br/> <b>Off</b> if the interface is not configured to respond to received PAUSE messages.</li> <li><b>rx-buffers</b>—<b>On</b> if the interface is configured to generate and send Ethernet PAUSE messages to the connected peer.<br/> <b>Off</b> if the interface is not configured to generate and send PAUSE messages.</li> </ul><br><i>NOTE:</i> This field is only displayed if asymmetric flow control is configured. | <b>detail extensive</b>       |

Table 277: show interfaces ge Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Level of Output              |
|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Auto-negotiation</b>        | Autonegotiation status: <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | All levels                   |
| <b>Remote-fault</b>            | Remote fault status: <ul style="list-style-type: none"> <li>• <b>Online</b>—Autonegotiation is manually configured as online.</li> <li>• <b>Offline</b>—Autonegotiation is manually configured as offline.</li> </ul>                                                                                                                                                                                                                                                                                                            | All levels                   |
| <b>Device flags</b>            | Information about the physical device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | All levels                   |
| <b>Interface flags</b>         | Information about the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | All levels                   |
| <b>Link flags</b>              | Information about the link.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | All levels                   |
| <b>CoS queues</b>              | Number of CoS queues configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>detail extensive none</b> |
| <b>Hold-times</b>              | Current interface hold-time up and hold-time down, in milliseconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail extensive</b>      |
| <b>Current address</b>         | Configured MAC address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>detail extensive none</b> |
| <b>Hardware address</b>        | MAC address of the hardware.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>detail extensive none</b> |
| <b>Last flapped</b>            | Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: year-month-day hour:minute:second timezone (hour:minute:second ago)</b> . For example, <b>Last flapped: 2008-01-16 10:52:40 UTC (3d 22:58 ago)</b> .                                                                                                                                                                                                                                                                             | <b>detail extensive none</b> |
| <b>Statistics last cleared</b> | Time when the statistics for the interface were last set to zero.                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive</b>      |
| <b>Traffic statistics</b>      | Number and rate of bytes and packets received and transmitted on the physical interface. <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> <p><b>NOTE:</b> The bandwidth bps counter is not enabled on the switch.</p> | <b>detail extensive</b>      |

Table 277: show interfaces ge Output Fields (*continued*)

| Field Name          | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Level of Output  |
|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>Input errors</b> | <p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Errors</b>—Sum of the incoming frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li>• <b>Runts</b>—Number of frames received that are smaller than the runt threshold.</li> <li>• <b>Policed discards</b>—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle.</li> <li>• <b>L3 incompletes</b>—Number of incoming packets discarded because they failed Layer 3 sanity checks of the headers. For example, a frame with less than 20 bytes of available IP header is discarded.</li> <li>• <b>L2 channel errors</b>—Number of times the software did not find a valid logical interface for an incoming frame.</li> <li>• <b>L2 mismatch timeouts</b>—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable.</li> <li>• <b>FIFO errors</b>—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul> | <b>extensive</b> |

Table 277: show interfaces ge Output Fields (*continued*)

| Field Name                              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Level of Output              |
|-----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Output errors</b>                    | <p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Carrier transitions</b>—Number of times the interface has gone from <b>down</b> to <b>up</b>. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning.</li> <li>• <b>Errors</b>—Sum of the outgoing frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Collisions</b>—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug.</li> <li>• <b>Aged packets</b>—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware.</li> <li>• <b>FIFO errors</b>—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>HS link CRC errors</b>—Number of errors on the high-speed links between the ASICs responsible for handling the switch interfaces.</li> <li>• <b>MTU errors</b>—Number of packets whose size exceeded the MTU of the interface.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul> | <b>extensive</b>             |
| <b>Egress queues</b>                    | Total number of egress queues supported on the specified interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b>      |
| <b>Queue counters (Egress )</b>         | <p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b>      |
| <b>Queue Number</b>                     | The CoS queue number and the forwarding classes mapped to the queue number. The <b>Mapped forwarding class</b> column lists the forwarding classes mapped to each CoS queue.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>detail extensive</b>      |
| <b>Active alarms and Active defects</b> | <p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the switch configuration, an alarm can ring the red or yellow alarm bell on the switch or turn on the red or yellow alarm LED on the front of the switch. These fields can contain the value <b>None</b> or <b>Link</b>.</p> <ul style="list-style-type: none"> <li>• <b>None</b>—There are no active defects or alarms.</li> <li>• <b>Link</b>—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive none</b> |



Table 277: show interfaces ge Output Fields (*continued*)

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Level of Output  |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>MAC statistics</b>    | <p>Receive and Transmit statistics reported by the PIC's MAC subsystem.</p> <ul style="list-style-type: none"> <li>• <b>Total octets</b> and <b>total packets</b>—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type.</li> <li>• <b>Unicast packets, Broadcast packets, and Multicast packets</b>—Number of unicast, broadcast, and multicast packets.</li> <li>• <b>CRC/Align errors</b>—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).</li> <li>• <b>FIFO error</b>—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>MAC control frames</b>—Number of MAC control frames.</li> <li>• <b>MAC pause frames</b>—Number of MAC control frames with <b>pause</b> operational code.</li> <li>• <b>Oversized frames</b>—Number of packets that exceeds the configured MTU.</li> <li>• <b>Jabber frames</b>—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms.</li> <li>• <b>Fragment frames</b>—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted.</li> <li>• <b>VLAN tagged frames</b>—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. This counter is not supported on EX Series switches and is always displayed as 0.</li> <li>• <b>Code violations</b>—Number of times an event caused the PHY to indicate “Data reception error” or “invalid data symbol error.”</li> </ul> | <b>extensive</b> |
| <b>Filter Statistics</b> | Receive and Transmit statistics reported by the PIC's MAC address filter subsystem.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>extensive</b> |

Table 277: show interfaces ge Output Fields (*continued*)

| Field Name                                    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Level of Output              |
|-----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Autonegotiation information</b>            | <p>Information about link autonegotiation:</p> <ul style="list-style-type: none"> <li>• <b>Negotiation status:</b> <ul style="list-style-type: none"> <li>• <b>Incomplete</b>—Ethernet interface has the speed or link mode configured.</li> <li>• <b>No autonegotiation</b>—Remote Ethernet interface has the speed or link mode configured or does not perform autonegotiation.</li> <li>• <b>Complete</b>—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.</li> </ul> </li> <li>• <b>Link partner status</b>—OK when the Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.</li> <li>• <b>Link partner:</b> <ul style="list-style-type: none"> <li>• <b>Link mode</b>—Depending on the capability of the attached Ethernet device, either <b>Full-duplex</b> or <b>Half-duplex</b>.</li> <li>• <b>Flow control</b>—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are <b>Symmetric</b> (link partner supports PAUSE on receive and transmit), <b>Asymmetric</b> (link partner supports PAUSE on transmit), and <b>Symmetric/Asymmetric</b> (link partner supports PAUSE on both receive and transmit or PAUSE only on receive).</li> <li>• <b>Remote fault</b>—Remote fault information from the link partner—<b>Failure</b> indicates a receive link error. <b>OK</b> indicates that the link partner is receiving. <b>Negotiation error</b> indicates a negotiation error. <b>Offline</b> indicates that the link partner is going offline.</li> <li>• <b>Link partner speed</b>—Speed of the link partner.</li> </ul> </li> <li>• <b>Local resolution:</b> <ul style="list-style-type: none"> <li>• <b>Flow control</b>—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are <b>Symmetric</b> (link partner supports PAUSE on receive and transmit), <b>Asymmetric</b> (link partner supports PAUSE on transmit), and <b>Symmetric/Asymmetric</b> (link partner supports PAUSE on both receive and transmit or PAUSE only on receive). For asymmetric PAUSE, shows if the PAUSE transmit and PAUSE receive states on the interface are <b>enable</b> or <b>disable</b>.</li> <li>• <b>Remote fault</b>—Remote fault information. <b>Link OK</b> (no error detected on receive), <b>Offline</b> (local interface is offline), and <b>Link Failure</b> (link error detected on receive).</li> </ul> </li> </ul> | <b>extensive</b>             |
| <b>Packet Forwarding Engine configuration</b> | <p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> <li>• <b>Destination slot</b>—FPC slot number.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>extensive</b>             |
| <b>Logical Interface</b>                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |                              |
| <b>Logical interface</b>                      | Name of the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | All levels                   |
| <b>Index</b>                                  | Index number of the logical interface, which reflects its initialization sequence.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b> none |
| <b>SNMP ifIndex</b>                           | SNMP interface index number for the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail extensive</b> none |
| <b>Generation</b>                             | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail extensive</b>      |
| <b>Flags</b>                                  | Information about the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | All levels                   |

Table 277: show interfaces ge Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                     | Level of Output              |
|--------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Encapsulation</b>           | Encapsulation on the logical interface.                                                                                                                                                                                                                                                                                                                                                                               | All levels                   |
| <b>Protocol</b>                | Protocol family.                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail extensive none</b> |
| <b>Traffic statistics</b>      | Number and rate of bytes and packets received (input) and transmitted (output) on the specified interface.                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b>      |
| <b>IPv6 transit statistics</b> | If IPv6 statistics tracking is enabled, number of IPv6 bytes and packets received and transmitted on the logical interface.                                                                                                                                                                                                                                                                                           | <b>extensive</b>             |
| <b>Local statistics</b>        | Number and rate of bytes and packets destined to and from the switch.                                                                                                                                                                                                                                                                                                                                                 | <b>extensive</b>             |
| <b>Transit statistics</b>      | Number and rate of bytes and packets transiting the switch.                                                                                                                                                                                                                                                                                                                                                           | <b>extensive</b>             |
| <b>Generation</b>              | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                     | <b>detail extensive</b>      |
| <b>Route Table</b>             | Route table in which the logical interface address is located. For example, <b>0</b> refers to the routing table <b>inet.0</b> .                                                                                                                                                                                                                                                                                      | <b>detail extensive none</b> |
| <b>Input Filters</b>           | Names of any input filters applied to this interface.                                                                                                                                                                                                                                                                                                                                                                 | <b>detail extensive</b>      |
| <b>Output Filters</b>          | Names of any output filters applied to this interface.                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive</b>      |
| <b>Flags</b>                   | Information about protocol family flags.<br><br>If unicast reverse-path forwarding (RPF) is explicitly configured on the specified interface, the uRPF flag is displayed. If unicast RPF was configured on a different interface (and therefore is enabled on all switch interfaces) but was not explicitly configured on the specified interface, the uRPF flag is not displayed even though unicast RPF is enabled. | <b>detail extensive</b>      |
| <b><i>protocol-family</i></b>  | Protocol family configured on the logical interface. If the protocol is <b>inet</b> , the IP address of the interface is also displayed.                                                                                                                                                                                                                                                                              | <b>brief</b>                 |
| <b>Flags</b>                   | Information about the address flags.                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail extensive none</b> |
| <b>Destination</b>             | IP address of the remote side of the connection.                                                                                                                                                                                                                                                                                                                                                                      | <b>detail extensive none</b> |
| <b>Local</b>                   | IP address of the logical interface.                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail extensive none</b> |
| <b>Broadcast</b>               | Broadcast address of the logical interlace.                                                                                                                                                                                                                                                                                                                                                                           | <b>detail extensive none</b> |
| <b>Generation</b>              | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                     | <b>detail extensive</b>      |

## Sample Output

### show interfaces

```
user@switch> show interfaces ge-0/0/9
Physical interface: ge-0/0/9, Enabled, Physical link is Down
 Interface index: 129, SNMP ifIndex: 21
 Link-level type: Ethernet, MTU: 1514, Speed: Unspecified, Loopback: Disabled,
 Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled
 Remote fault: Online
 Device flags : Present Running Down
 Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
 CoS queues : 8 supported, 8 maximum usable queues
 Hold-times : Up 0 ms, Down 0 ms
 Current address: 00:19:e2:50:3f:41, Hardware address: 00:19:e2:50:3f:41
 Last flapped : 2008-01-16 11:40:53 UTC (4d 02:30 ago)
 Input rate : 0 bps (0 pps)
 Output rate : 0 bps (0 pps)
 Ingress rate at Packet Forwarding Engine : 0 bps (0 pps)
 Ingress drop rate at Packet Forwarding Engine : 0 bps (0 pps)
 Active alarms : None
 Active defects : None

Logical interface ge-0/0/9.0 (Index 65) (SNMP ifIndex 22)
 Flags: SNMP-Traps
 Encapsulation: ENET2
 Input packets : 0
 Output packets: 0
 Protocol eth-switch
 Flags: None
```

### show interfaces brief

```
user@switch> show interfaces ge-0/0/9 brief
Physical interface: ge-0/0/9, Enabled, Physical link is Down
 Description: voice priority and tcp and icmp traffic rate-limiting filter at i
 ngress port
 Link-level type: Ethernet, MTU: 1514, Speed: Unspecified, Loopback: Disabled,
 Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
 Remote fault: Online
 Device flags : Present Running Down
 Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
 Link flags : None

Logical interface ge-0/0/9.0
 Flags: Device-Down SNMP-Traps Encapsulation: ENET2
 eth-switch
```

### show interfaces detail (Symmetric Flow Control and Autonegotiation Enabled)

```
user@switch> show interfaces ge-0/0/9 detail
Physical interface: ge-0/0/9, Enabled, Physical link is Up
 Interface index: 193, SNMP ifIndex: 206, Generation: 196
 Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: Auto,
 BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
 Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
 Remote fault: Online
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0
 Link flags : None
 CoS queues : 8 supported, 8 maximum usable queues
```

```

Hold-times : Up 0 ms, Down 0 ms
Current address: 00:1f:12:30:ff:40, Hardware address: 00:1f:12:30:ff:40
Last flapped : 2009-05-05 06:03:05 UTC (00:22:13 ago)
Statistics last cleared: Never
Traffic statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets : 0 0 pps
 Output packets: 0 0 pps
IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets : 0
 Output packets: 0
Egress queues: 8 supported, 4 in use
Queue counters:
 Queued packets Transmitted packets Dropped packets

 0 best-effort 0 0 0
 1 assured-forw 0 0 0
 5 expedited-fo 0 0 0
 7 network-cont 0 0 0

Active alarms : None
Active defects : None

Logical interface ge-0/0/9.0 (Index 65) (SNMP ifIndex 235) (Generation 130)
Flags: SNMP-Traps Encapsulation: ENET2
Bandwidth: 0
Traffic statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets : 0
 Output packets: 0
Local statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets : 0
 Output packets: 0
Transit statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets : 0 0 pps
 Output packets: 0 0 pps
Protocol eth-switch, Generation: 146, Route table: 0
Flags: Is-Primary
Input Filters: f1,
Output Filters: f2,,,

```

#### show interfaces detail (Asymmetric Flow Control and Autonegotiation Enabled)

```

user@switch> show interfaces ge-0/0/9 detail
Physical interface: ge-0/0/9, Enabled, Physical link is Up
Interface index: 193, SNMP ifIndex: 206, Generation: 196
Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: Auto,
BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
Source filtering: Disabled, Configured-flow-control tx-buffers: off
rx-buffers: on ,
Auto-negotiation: Enabled,

```

```

Remote fault: Online
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x0
Link flags : None
CoS queues : 8 supported, 8 maximum usable queues
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:1f:12:30:ff:40, Hardware address: 00:1f:12:30:ff:40
Last flapped : 2009-05-05 06:03:05 UTC (00:22:13 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets : 0 0 pps
Output packets: 0 0 pps
IPv6 transit statistics:
Input bytes : 0
Output bytes : 0
Input packets : 0
Output packets: 0
Egress queues: 8 supported, 4 in use
Queue counters:

```

|                | Queued packets | Transmitted packets | Dropped packets |
|----------------|----------------|---------------------|-----------------|
| 0 best-effort  | 0              | 0                   | 0               |
| 1 assured-forw | 0              | 0                   | 0               |
| 5 expedited-fo | 0              | 0                   | 0               |
| 7 network-cont | 0              | 0                   | 0               |

```

Active alarms : None
Active defects : None

Logical interface ge-0/0/9.0 (Index 65) (SNMP ifIndex 235) (Generation 130)
Flags: SNMP-Traps Encapsulation: ENET2
Bandwidth: 0
Traffic statistics:
Input bytes : 0
Output bytes : 0
Input packets : 0
Output packets: 0
Local statistics:
Input bytes : 0
Output bytes : 0
Input packets : 0
Output packets: 0
Transit statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets : 0 0 pps
Output packets: 0 0 pps
Protocol eth-switch, Generation: 146, Route table: 0
Flags: Is-Primary
Input Filters: f1,
Output Filters: f2,,,,

```

#### show interfaces extensive (Symmetric Flow Control and Autonegotiation Enabled)

```

user@switch> show interfaces ge-0/0/12 extensive
interface: ge-0/0/12, Enabled, Physical link is Down
Interface index: 49164, SNMP ifIndex: 574, Generation: 142

```

```

Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Duplex: Full-Duplex,
BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
Device flags : Present Running Down
Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
Link flags : None
CoS queues : 8 supported, 8 maximum usable queues
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:22:83:2a:d8:dc, Hardware address: 00:22:83:2a:d8:dc
Last flapped : 2011-02-25 00:45:03 UTC (22:42:48 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets : 0 0 pps
Output packets: 0 0 pps
IPv6 transit statistics:
Input bytes : 0
Output bytes : 0
Input packets : 0
Output packets: 0
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
FIFO errors: 0, Resource errors: 0
Output errors:
Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 8 in use
Queue counters:

```

|                | Queued packets | Transmitted packets | Dropped packets |
|----------------|----------------|---------------------|-----------------|
| 0 best-effort  | 0              | 0                   | 0               |
| 2 no-loss      | 0              | 0                   | 0               |
| 3 fcoe         | 0              | 0                   | 0               |
| 7 network-cont | 0              | 0                   | 0               |

```

Queue number: Mapped forwarding classes
0 best-effort
2 no-loss
3 fcoe
7 network-control
Active alarms : LINK
Active defects : LINK
MAC statistics:

```

|                    | Receive | Transmit |
|--------------------|---------|----------|
| Total octets       | 0       | 0        |
| Total packets      | 0       | 0        |
| Unicast packets    | 0       | 0        |
| Broadcast packets  | 0       | 0        |
| Multicast packets  | 0       | 0        |
| CRC/Align errors   | 0       | 0        |
| FIFO errors        | 0       | 0        |
| MAC control frames | 0       | 0        |
| MAC pause frames   | 0       | 0        |
| Oversized frames   | 0       |          |
| Jabber frames      | 0       |          |
| Fragment frames    | 0       |          |

```

VLAN tagged frames 0
Code violations 0
MAC Priority Flow Control Statistics:
 Priority : 0 0 0
 Priority : 1 0 0
 Priority : 2 0 0
 Priority : 3 0 0
 Priority : 4 0 0
 Priority : 5 0 0
 Priority : 6 0 0
 Priority : 7 0 0
Filter statistics:
 Input packet count 0
 Input packet rejects 0
 Input DA rejects 0
 Input SA rejects 0
 Output packet count 0
 Output packet pad count 0
 Output packet error count 0
 CAM destination filters: 1, CAM source filters: 0
Autonegotiation information:
 Negotiation status: Incomplete
Packet Forwarding Engine configuration:
 Destination slot: 0
CoS information:
 Direction : Output
 CoS transmit queue Bandwidth Buffer Priority
Limit
 0 best-effort 75 750000000 75 0 low
none
 7 network-control 5 500000000 5 0 low
none
 8 mcast-be 15 1500000000 15 0 low
none
 11 mcast-nc 5 500000000 5 0 low
none

```

#### show interfaces extensive (Asymmetric Flow Control and Autonegotiation Enabled)

```

user@switch> show interfaces ge-0/0/12 extensive
interface: ge-0/0/12, Enabled, Physical link is Down
 Interface index: 49164, SNMP ifIndex: 574, Generation: 142
 Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Duplex: Full-Duplex,
 BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
 Source filtering: Disabled, Configured-flow-control tx-buffers: off
rx-buffers: on
 Auto-negotiation: Enabled,
 Remote fault: Online
 Device flags : Present Running Down
 Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
 Link flags : None
 CoS queues : 8 supported, 8 maximum usable queues
 Hold-times : Up 0 ms, Down 0 ms
 Current address: 00:22:83:2a:d8:dc, Hardware address: 00:22:83:2a:d8:dc
 Last flapped : 2011-02-25 00:45:03 UTC (22:42:48 ago)
 Statistics last cleared: Never
Traffic statistics:
 Input bytes : 0 0 bps
 Output bytes: 0 0 bps
 Input packets: 0 0 pps

```



```

Output packets: 0 0 pps
IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Input errors:
 Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
 L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
 FIFO errors: 0, Resource errors: 0
Output errors:
 Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

 FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 8 in use
Queue counters: Queued packets Transmitted packets Dropped packets

 0 best-effort 0 0 0
 2 no-loss 0 0 0
 3 fcoe 0 0 0
 7 network-cont 0 0 0

Queue number: Mapped forwarding classes
 0 best-effort
 2 no-loss
 3 fcoe
 7 network-control
Active alarms : LINK
Active defects : LINK
MAC statistics:
 Total octets Receive Transmit
 Total packets 0 0
 Unicast packets 0 0
 Broadcast packets 0 0
 Multicast packets 0 0
 CRC/Align errors 0 0
 FIFO errors 0 0
 MAC control frames 0 0
 MAC pause frames 0 0
 Oversized frames 0
 Jabber frames 0
 Fragment frames 0
 VLAN tagged frames 0
 Code violations 0
MAC Priority Flow Control Statistics:
 Priority : 0 0 0
 Priority : 1 0 0
 Priority : 2 0 0
 Priority : 3 0 0
 Priority : 4 0 0
 Priority : 5 0 0
 Priority : 6 0 0
 Priority : 7 0 0
Filter statistics:
 Input packet count 0
 Input packet rejects 0
 Input DA rejects 0
 Input SA rejects 0

```

```

Output packet count 0
Output packet pad count 0
Output packet error count 0
CAM destination filters: 1, CAM source filters: 0
Autonegotiation information:
Negotiation status: Complete
Link Partner:
 Link mode: Full-duplex, Flow control: None, Remote fault: OK,
 Link partner Speed: 1000 Mbps
Local resolution:
 Flow control: enable PAUSE transmit and Disable PAUSE receive, Remote
 fault: Link OK
Packet Forwarding Engine configuration:
 Destination slot: 0
CoS information:
 Direction : Output
 CoS transmit queue Bandwidth Buffer Priority
Limit
 % bps % usec
0 best-effort 75 750000000 75 0 low
none
7 network-control 5 50000000 5 0 low
none
8 mcast-be 15 150000000 15 0 low
none
11 mcast-nc 5 50000000 5 0 low
none

```

#### show interfaces terse

```

user@switch> show interfaces ge-0/0/12 terse
Interface Admin Link Proto Local Remote
ge-0/0/12 up up

```


#### show interfaces terse (QFabric Systems)

```

user@switch> show interfaces node1:ge-0/0/0 terse
Physical interface: node1:ge-0/0/0, Enabled, Physical link is Down
 Interface index: 129, SNMP ifIndex: 2884086
 Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
 Error: None, MAC-REWRITE Error: None,
 Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled
 Interface flags: Internal: 0x4000
 CoS queues : 8 supported, 8 maximum usable queues
 Current address: 02:00:09:03:00:00, Hardware address: 02:00:09:03:00:00
 Last flapped : Never
 Input rate : 0 bps (0 pps)
 Output rate : 0 bps (0 pps)

```

## show interfaces (GRE)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show interfaces <i>interface-type</i> &lt;brief   detail   extensive   terse&gt; &lt;descriptions&gt; &lt;media&gt; &lt;snmp-index <i>snmp-index</i>&gt; &lt;statistics&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 12.1 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Display status information about the specified generic routing encapsulation (GRE) interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <p><b><i>interface-type</i></b>—On M Series and T Series routers and EX Series switches, the interface type is <b><i>gr-fpc/pic/port</i></b>.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified output level of interface information.</p> <p><b>descriptions</b>—(Optional) Display interface description strings.</p> <p><b>media</b>—(Optional) Display media-specific information about network interfaces.</p> <p><b>snmp-index <i>snmp-index</i></b>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><b>statistics</b>—(Optional) Display static interface statistics.</p> |
|                                 | <div>  <p><b>NOTE:</b> You can configure generic routing encapsulation (GRE) interfaces (<i>gre-x/y/z</i>) only for GMPLS control channels. GRE interfaces are not supported or configurable for other applications. For more information about GMPLS, see the <i>MPLS Applications Feature Guide for Routing Devices</i> and the <i>Junos OS, Release 15.1</i>.</p> </div>                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>List of Sample Output</b>    | <p><a href="#">show interfaces (GRE) on page 3136</a></p> <p><a href="#">show interfaces brief (GRE) on page 3136</a></p> <p><a href="#">show interfaces detail (GRE) on page 3136</a></p> <p><a href="#">show interfaces (Layer 2 Services Over GRE Interfaces) on page 3137</a></p> <p><a href="#">show interfaces extensive (Layer 2 Services Over GRE Interfaces) on page 3137</a></p> <p><a href="#">show interfaces detail (GRE) on an EX4200 Virtual Chassis Member Switch on page 3138</a></p> <p><a href="#">show interfaces extensive (GRE) on page 3139</a></p>                                                                            |

**Output Fields** Table 278 lists the output fields for the **show interfaces** (GRE) command. Output fields are listed in the approximate order in which they appear.

**Table 278: GRE show interfaces Output Fields**

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Level of Output              |
|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Physical Interface</b>      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |                              |
| <b>Physical interface</b>      | Name of the physical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | All levels                   |
| <b>Enabled</b>                 | State of the interface. Possible values are described in the "Enabled Field" section under <i>Common Output Fields Description</i> .                                                                                                                                                                                                                                                                                                                                                               | All levels                   |
| <b>Interface index</b>         | Physical interface's index number, which reflects its initialization sequence.                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>detail extensive none</b> |
| <b>SNMP ifIndex</b>            | SNMP index number for the physical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail extensive none</b> |
| <b>Generation</b>              | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail extensive</b>      |
| <b>Type</b>                    | Type of interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | All levels                   |
| <b>Link-level type</b>         | Encapsulation used on the physical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                      | All levels                   |
| <b>MTU</b>                     | MTU size on the physical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                | All levels                   |
| <b>Speed</b>                   | Speed at which the interface is running.                                                                                                                                                                                                                                                                                                                                                                                                                                                           | All levels                   |
| <b>Hold-times</b>              | Current interface hold-time up and hold-time down, in milliseconds.                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive</b>      |
| <b>Device Flags</b>            | Information about the physical device. Possible values are described in the "Device Flags" section under <i>Common Output Fields Description</i> .                                                                                                                                                                                                                                                                                                                                                 | All levels                   |
| <b>Interface Flags</b>         | Information about the interface. Possible values are described in the "Interface Flags" section under <i>Common Output Fields Description</i> .                                                                                                                                                                                                                                                                                                                                                    | All levels                   |
| <b>Input rate</b>              | Input rate in bits per second (bps) and packets per second (pps).                                                                                                                                                                                                                                                                                                                                                                                                                                  | None specified               |
| <b>Output rate</b>             | Output rate in bps and pps.                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | None specified               |
| <b>Statistics last cleared</b> | Time when the statistics for the interface were last set to zero.                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail extensive</b>      |
| <b>Traffic statistics</b>      | <p>The number of and the rate at which input and output bytes and packets are received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> | <b>detail extensive</b>      |
| <b>Logical Interface</b>       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |                              |

Table 278: GRE show interfaces Output Fields (*continued*)

| Field Name                         | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Level of Output              |
|------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Logical interface</b>           | Name of the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | All levels                   |
| <b>Index</b>                       | Logical interface index number, which reflects its initialization sequence.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>detail extensive none</b> |
| <b>SNMP ifIndex</b>                | Logical interface SNMP interface index number.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail extensive none</b> |
| <b>Generation</b>                  | Unique number for use by Juniper Networks technical support.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>detail extensive</b>      |
| <b>Flags</b>                       | <p>Information about the logical interface. Possible values listed in the “Logical Interface Flags” section under <i>Common Output Fields Description</i>. describe general information about the logical interface.</p> <p>GRE-specific information about the logical interface is indicated by the presence or absence of the following value in this field:</p> <ul style="list-style-type: none"> <li>• <b>Reassemble-Pkts</b>—If the <b>Flags</b> field includes this string, the GRE tunnel is configured to reassemble tunnel packets that were fragmented after tunnel encapsulation.</li> </ul>                              | All levels                   |
| <b>IP-Header</b>                   | <p>IP header of the logical interface. If the <b>tunnel key</b> statement is configured, this information is included in the <b>IP Header</b> entry.</p> <p>GRE-specific information about the logical interface is indicated by the presence or absence of the following value in this field:</p> <ul style="list-style-type: none"> <li>• <b>df</b>—If the <b>IP-Header</b> field includes this string immediately following the 16 bits of identification information (that is, if <b>:df</b> displays after the twelfth byte), the GRE tunnel is configured to allow fragmentation of GRE packets after encapsulation.</li> </ul> | All levels                   |
| <b>Encapsulation</b>               | Encapsulation on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | All levels                   |
| <b>L2 Routing Instance</b>         | Name of the Layer 2 routing instance associated with the GRE interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | All levels                   |
| <b>L3 Routing Instance</b>         | Name of the Layer 3 routing instance associated with the GRE interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | All levels                   |
| <b>Copy-tos-to-outer-ip-header</b> | <p>Status of type of service (ToS) bits in the GRE packet header:</p> <ul style="list-style-type: none"> <li>• <b>On</b>—ToS bits were copied from the payload packet header into the header of the IP packet sent through the GRE tunnel.</li> <li>• <b>Off</b>—ToS bits were not copied from the payload packet header and are set to 0 in the GRE packet header.</li> </ul> <p><b>NOTE:</b> EX Series switches do not support copying ToS bits to the encapsulated packet, so the value of this field is always <b>Off</b> in switch output.</p>                                                                                   | <b>detail extensive</b>      |
| <b>Gre keepalives configured</b>   | <p>Indicates whether a GRE keepalive time and hold time are configured for the GRE tunnel.</p> <p><b>NOTE:</b> EX Series switches do not support configuration of GRE tunnel keepalive times and hold times, so the value of this field is always <b>Off</b> in switch output.</p>                                                                                                                                                                                                                                                                                                                                                    | <b>detail extensive</b>      |

Table 278: GRE show interfaces Output Fields (*continued*)

| Field Name                            | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Level of Output              |
|---------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Gre keepalives adjacency state</b> | Status of the other end of the GRE tunnel: <b>Up</b> or <b>Down</b> . If keepalive messages are not received by either end of the GRE tunnel within the hold-time period, the GRE keepalive adjacency state is down even when the GRE tunnel is up.                                                                                                                                                                                                                                                                                                                          | <b>detail extensive</b>      |
| <b>Input packets</b>                  | Number of packets received on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | None specified               |
| <b>Output packets</b>                 | Number of packets transmitted on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | None specified               |
| <b>Traffic statistics</b>             | Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize. <ul style="list-style-type: none"> <li>• <b>Input rate</b>—Rate of bits and packets received on the interface.</li> <li>• <b>Output rate</b>—Rate of bits and packets transmitted on the interface.</li> </ul> | <b>detail extensive</b>      |
| <b>Local statistics</b>               | Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.                                                                                                                                                                                                                                                                                           | <b>detail extensive</b>      |
| <b>Transit statistics</b>             | Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.                                                                                                                                                                                                                                                                                                                         | <b>detail extensive none</b> |
| <b>Protocol</b>                       | Protocol family configured on the logical interface, such as <b>iso</b> , <b>inet6</b> , or <b>mpls</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail extensive none</b> |
| <b><i>protocol-family</i></b>         | Protocol family configured on the logical interface. If the protocol is <b>inet</b> , the IP address of the interface is also displayed.                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>brief</b>                 |
| <b>MTU</b>                            | MTU size on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>detail extensive none</b> |
| <b>Generation</b>                     | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b>      |
| <b>Route table</b>                    | Routing table in which the logical interface address is located. For example, <b>0</b> refers to the routing table <b>inet.0</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>detail extensive</b>      |
| <b>Flags</b>                          | Information about the protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> .                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>detail extensive none</b> |
| <b>Addresses, Flags</b>               | Information about the address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>detail extensive none</b> |
| <b>Destination</b>                    | IP address of the remote side of the connection.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail extensive none</b> |
| <b>Local</b>                          | IP address of the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>detail extensive none</b> |
| <b>Broadcast</b>                      | Broadcast address of the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail extensive none</b> |

Table 278: GRE show interfaces Output Fields (*continued*)

| Field Name        | Field Description                                                 | Level of Output         |
|-------------------|-------------------------------------------------------------------|-------------------------|
| <b>Generation</b> | Unique number for use by Juniper Networks technical support only. | <b>detail extensive</b> |

## Sample Output

### show interfaces (GRE)

```
user@host> show interfaces gr-1/2/0
Physical interface: gr-0/0/0, Enabled, Physical link is Up
 Interface index: 132, SNMP ifIndex: 26
 Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
 Device flags : Present Running
 Interface flags: Point-To-Point SNMP-Traps
 Input rate : 0 bps (0 pps)
 Output rate : 0 bps (0 pps)

Logical interface gr-0/0/0.0 (Index 68) (SNMP ifIndex 47)
 Flags: Point-To-Point SNMP-Traps 16384
 IP-Header 1.1.1.2:1.1.1.1:47:df:64:0000000000000000 Encapsulation: GRE-NULL
 Input packets : 0
 Output packets: 0
 Protocol inet, MTU: 1476
 Flags: None
 Addresses, Flags: Is-Primary
 Local: 1.10.1.1
```

### show interfaces brief (GRE)

```
user@host> show interfaces gr-1/2/0 brief
Physical interface: gr-1/2/0, Enabled, Physical link is Up
 Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
 Device flags : Present Running
 Interface flags: Point-To-Point SNMP-Traps

Logical interface gr-1/2/0.0
 Flags: Hardware-Down Point-To-Point SNMP-Traps 0x4000
 IP-Header 10.10.0.2:10.10.0.1:47:df:64:0000000000000000
 Encapsulation: GRE-NULL
 inet 10.100.0.1/30
 mpls
```

### show interfaces detail (GRE)

```
user@host> show interfaces gr-1/2/0 detail
Physical interface: gr-0/0/0, Enabled, Physical link is Up
 Interface index: 132, SNMP ifIndex: 26, Generation: 13
 Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
 Hold-times : Up 0 ms, Down 0 ms
 Device flags : Present Running
 Interface flags: Point-To-Point SNMP-Traps
 Statistics last cleared: Never
 Traffic statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps

Logical interface gr-0/0/0.0 (Index 68) (SNMP ifIndex 47) (Generation 8)
 Flags: Point-To-Point SNMP-Traps 16384
 IP-Header 1.1.1.2:1.1.1.1:47:df:64:0000000000000000 Encapsulation: GRE-NULL
 Traffic statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
```



```

Output packets: 0
Local statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Transit statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
Protocol inet, MTU: 1476, Generation: 12, Route table: 0
Flags: None
Addresses, Flags: Is-Primary
Destination: Unspecified, Local: 1.10.1.1, Broadcast: Unspecified,
Generation: 15

```

### show interfaces (Layer 2 Services Over GRE Interfaces)

```

user@host> show interfaces gr-2/2/10
show interfaces gr-2/2/10
Physical interface: gr-2/2/10, Enabled, Physical link is Up
Interface index: 214, SNMP ifIndex: 690
Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 1000mbps
Device flags : Present Running
Interface flags: Point-To-Point SNMP-Traps
Input rate : 0 bps (0 pps)
Output rate : 0 bps (0 pps)

Logical interface gr-2/2/10.0 (Index 342) (SNMP ifIndex 10834)
Flags: Up Point-To-Point SNMP-Traps 0x4000 IP-Header
3.0.0.1:3.0.0.254:47:df:64:0000000000000000 Encapsulation: GRE-NULL
L2 Routing Instance: vs1, L3 Routing Instance: default
Copy-tos-to-outer-ip-header: Off
Gre keepalives configured: Off, Gre keepalives adjacency state: down
Input packets : 2
Output packets: 0
Protocol bridge, MTU: 1476
Flags: Sendbroadcast-pkt-to-re
Addresses, Flags: Is-Preferred Is-Primary
Destination: 6/8, Local: 6.0.0.1, Broadcast: 6.255.255.255

```

### show interfaces extensive (Layer 2 Services Over GRE Interfaces)

```

user@host> show interfaces gr-2/2/10.0 extensive

Flags: SNMP-Traps Encapsulation: ENET2
L2 Routing Instance: vs1, L3 Routing Instance: default
Traffic statistics:
Input bytes : 58851250
Output bytes : 0
Input packets: 1279375
Output packets: 0
Local statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Transit statistics:
Input bytes : 58851250 75136 bps

```

```

Output bytes : 0 0 bps
Input packets: 1279375 204 pps
Output packets: 0 0 pps
Protocol bridge, MTU: 1476, Generation: 175, Route table: 7
Flags: Access-Mode

```

### show interfaces detail (GRE) on an EX4200 Virtual Chassis Member Switch

```

user@switch> show interfaces gr-2/0/15 detail
Physical interface: gr-2/0/15, Enabled, Physical link is Up
Interface index: 195, SNMP ifIndex: 846, Generation: 198
Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 1000mbps
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:1f:12:38:0f:d2, Hardware address: 00:1f:12:38:0f:d2
Device flags : Present Running
Interface flags: Point-To-Point SNMP-Traps
Statistics last cleared: 2011-09-14 17:43:15 UTC (00:00:18 ago)
Traffic statistics:
Input bytes : 5600636 0 bps
Output bytes : 5600636 0 bps
Input packets: 20007 0 pps
Output packets: 20007 0 pps
IPv6 transit statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

Logical interface gr-2/0/15.0 (Index 75) (SNMP ifIndex 847) (HW Token 4093)
(Generation 140)
Flags: Point-To-Point SNMP-Traps 0x0
IP-Header 180.20.30.2:180.20.20.3:47:df:64:0000000000000000
Encapsulation: GRE-NULL
Copy-tos-to-outer-ip-header: Off
Gre keepalives configured: Off, Gre keepalives adjacency state: down
Traffic statistics:
Input bytes : 5600886
Output bytes : 2881784
Input packets: 20010
Output packets: 10018
Local statistics:
Input bytes : 398
Output bytes : 264
Input packets: 5
Output packets: 3
Transit statistics:
Input bytes : 5600488 0 bps
Output bytes : 2881520 0 bps
Input packets: 20005 0 pps
Output packets: 10015 0 pps
Protocol inet, Generation: 159, Route table: 0
Flags: None
Addresses, Flags: Is-Preferred Is-Primary
Destination: 90.90.90/24, Local: 90.90.90.10, Broadcast: 90.90.90.255,
Generation: 144

Logical interface gr-2/0/15.1 (Index 80) (SNMP ifIndex 848) (HW Token 4088)
(Generation 150)
Flags: Point-To-Point SNMP-Traps 0x0
IP-Header 160.20.40.2:160.20.30.1:47:df:64:0000000000000000
Encapsulation: GRE-NULL

```

```

Copy-tos-to-outer-ip-header: Off
Gre keepalives configured: Off, Gre keepalives adjacency state: down
Traffic statistics:
 Input bytes : 260
 Output bytes : 2880148
 Input packets: 4
 Output packets: 10002
Local statistics:
 Input bytes : 112
 Output bytes : 0
 Input packets: 2
 Output packets: 0
Transit statistics:
 Input bytes : 148 0 bps
 Output bytes : 2880148 0 bps
 Input packets: 2 0 pps
 Output packets: 10002 0 pps
Protocol inet, Generation: 171, Route table: 0
Flags: None
Addresses, Flags: Is-Preferred Is-Primary
 Destination: 70.70.70/24, Local: 70.70.70.10, Broadcast: 70.70.70.255,
 Generation: 160

```

#### show interfaces extensive (GRE)

The output for the **show interfaces extensive** command is identical to that for the **show interfaces detail** command. For sample output, see [show interfaces detail \(GRE\) on page 3136](#) and [show interfaces detail \(GRE\) on an EX4200 Virtual Chassis Member Switch on page 3138](#).

## show interfaces irb

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show interfaces irb &lt;brief   detail   extensive   terse&gt; &lt;descriptions&gt; &lt;media&gt; &lt;routing-instance <i>instance-name</i>&gt; &lt;snmp-index <i>snmp-index</i>&gt; &lt;statistics&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 12.3R2.</p> <p>Command introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2 for the QFX Series</p>                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Display integrated routing and bridging interfaces information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output.</p> <p><b>descriptions</b>—(Optional) Display interface description strings.</p> <p><b>media</b>—(Optional) Display media-specific information about network interfaces.</p> <p><b>routing-instance <i>instance-name</i></b>—(Optional) Display information for the interface with the specified SNMP index.</p> <p><b>snmp-index <i>snmp-index</i></b>—(Optional) Display information for the interface with the specified SNMP index.</p> <p><b>statistics</b>—(Optional) Display static interface statistics.</p> |
| <b>Additional Information</b>   | Integrated routing and bridging (IRB) provides simultaneous support for Layer 2 bridging and Layer 3 IP routing on the same interface. IRB enables you to route local packets to another routed interface or to another VLAN that has a Layer 3 protocol configured.                                                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>List of Sample Output</b>    | <p><a href="#">show interfaces irb extensive on page 3144</a></p> <p><a href="#">show interfaces irb snmp-index on page 3145</a></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Output Fields</b>            | <a href="#">Table 279</a> lists the output fields for the <b>show interfaces irb</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                               |

**Table 279: show interfaces irb Output Fields**

| Field Name                | Field Description                                                                                                                             | Level of Output |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Physical Interface</b> |                                                                                                                                               |                 |
| <b>Physical interface</b> | Name of the physical interface.                                                                                                               | All levels      |
| <b>Enabled</b>            | State of the physical interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> . | All levels      |

Table 279: show interfaces irb Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                          | Level of Output                    |
|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------|
| <b>Proto</b>                   | Protocol configured on the interface.                                                                                                                                                                                                                      | <b>terse</b>                       |
| <b>Interface index</b>         | Physical interface index number, which reflects its initialization sequence.                                                                                                                                                                               | <b>detail extensive none</b>       |
| <b>SNMP ifIndex</b>            | SNMP index number for the physical interface.                                                                                                                                                                                                              | <b>detail extensive none</b>       |
| <b>Type</b>                    | Physical interface type.                                                                                                                                                                                                                                   | <b>detail extensive none</b>       |
| <b>Link-level type</b>         | Encapsulation being used on the physical interface.                                                                                                                                                                                                        | <b>detail extensive brief none</b> |
| <b>MTU</b>                     | MTU size on the physical interface.                                                                                                                                                                                                                        | <b>detail extensive brief none</b> |
| <b>Clocking</b>                | Reference clock source: <b>Internal</b> or <b>External</b> . Always unspecified on IRB interfaces.                                                                                                                                                         | <b>detail extensive brief</b>      |
| <b>Speed</b>                   | Speed at which the interface is running. Always unspecified on IRB interfaces.                                                                                                                                                                             | <b>detail extensive brief</b>      |
| <b>Device flags</b>            | Information about the physical device. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> .                                                                                                         | <b>detail extensive brief none</b> |
| <b>Interface flags</b>         | Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> .                                                                                                            | <b>detail extensive brief none</b> |
| <b>Link type</b>               | Physical interface link type: <b>full duplex</b> or <b>half duplex</b> .                                                                                                                                                                                   | <b>detail extensive none</b>       |
| <b>Link flags</b>              | Information about the link. Possible values are described in the “Links Flags” section under <i>Common Output Fields Description</i> .                                                                                                                     | <b>detail extensive none</b>       |
| <b>Physical Info</b>           | Physical interface information.                                                                                                                                                                                                                            | All levels                         |
| <b>Hold-times</b>              | Current interface hold-time up and hold-time down, in milliseconds.                                                                                                                                                                                        | <b>detail extensive</b>            |
| <b>Current address</b>         | Configured MAC address.                                                                                                                                                                                                                                    | <b>detail extensive none</b>       |
| <b>Hardware address</b>        | MAC address of the hardware.                                                                                                                                                                                                                               | <b>detail extensive none</b>       |
| <b>Alternate link address</b>  | Backup address of the link.                                                                                                                                                                                                                                | <b>detail extensive</b>            |
| <b>Last flapped</b>            | Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: year-month-day hours:minutes:seconds timezone (hours:minutes:seconds ago)</b> . For example, <b>Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)</b> . | <b>detail extensive none</b>       |
| <b>Statistics last cleared</b> | Time when the statistics for the interface were last set to zero.                                                                                                                                                                                          | <b>detail extensive</b>            |

Table 279: show interfaces irb Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Level of Output         |
|--------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Traffic statistics</b>      | <p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>detail extensive</b> |
| <b>IPv6 transit statistics</b> | <p>Number of IPv6 transit bytes and packets received and transmitted on the physical interface if IPv6 statistics tracking is enabled.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>detail extensive</b> |
| <b>Input errors</b>            | <p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Errors</b>—Sum of the incoming frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li>• <b>Runts</b>—Number of frames received that are smaller than the runt threshold.</li> <li>• <b>Giants</b>—Number of frames received that are larger than the giant threshold.</li> <li>• <b>Policed discards</b>—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul>           | <b>detail extensive</b> |
| <b>Output errors</b>           | <p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Carrier transitions</b>—Number of times the interface has gone from <b>down</b> to <b>up</b>. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the DPC is malfunctioning.</li> <li>• <b>Errors</b>—Sum of the outgoing frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>MTU errors</b>—Number of packets whose size exceeded the MTU of the interface.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul> | <b>detail extensive</b> |

---

#### Logical Interface

---

Table 279: show interfaces irb Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Level of Output                 |
|--------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Logical interface</b>       | Name of the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | All levels                      |
| <b>Index</b>                   | Index number of the logical interface (which reflects its initialization sequence).                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail extensive</b><br>none |
| <b>SNMP ifIndex</b>            | SNMP interface index number of the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>detail extensive</b><br>none |
| <b>Generation</b>              | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive</b>         |
| <b>Flags</b>                   | Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under <i>Common Output Fields Description</i> .                                                                                                                                                                                                                                                                                                                                 | <b>detail extensive</b>         |
| <b>Encapsulation</b>           | Encapsulation on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>detail extensive</b>         |
| <b>Bandwidth</b>               | Speed at which the interface is running.                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail extensive</b>         |
| <b>Routing Instance</b>        | Routing instance IRB is configured under.                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail extensive</b>         |
| <b>Bridging Domain</b>         | Bridging domain IRB is participating in.                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail extensive</b>         |
| <b>Traffic statistics</b>      | Number and rate of bytes and packets received and transmitted on the logical interface. <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>                                             | <b>detail extensive</b>         |
| <b>IPv6 transit statistics</b> | Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled. <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> | <b>detail extensive</b>         |
| <b>Local statistics</b>        | Statistics for traffic received from and transmitted to the Routing Engine.                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>detail extensive</b>         |
| <b>Transit statistics</b>      | Statistics for traffic transiting the router.                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>detail extensive</b>         |
| <b>Protocol</b>                | Protocol family configured on the local interface. Possible values are described in the "Protocol Field" section under <i>Common Output Fields Description</i> .                                                                                                                                                                                                                                                                                                                                | <b>detail extensive</b>         |
| <b>MTU</b>                     | Maximum transmission unit size on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail extensive</b>         |
| <b>Maximum labels</b>          | Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                 | <b>detail extensive</b><br>none |

Table 279: show interfaces irb Output Fields (*continued*)

| Field Name              | Field Description                                                                                                                                               | Level of Output         |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Generation</b>       | Unique number for use by Juniper Networks technical support only.                                                                                               | <b>detail extensive</b> |
| <b>Route table</b>      | Routing table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.                                             | <b>detail extensive</b> |
| <b>Addresses, Flags</b> | Information about address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .                 | <b>detail extensive</b> |
| <b>Policer</b>          | The policer that is to be evaluated when packets are received or transmitted on the interface.                                                                  | <b>detail extensive</b> |
| <b>Flags</b>            | Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> . | <b>detail extensive</b> |

## Sample Output

### show interfaces irb extensive

```

user@host> show interfaces irb extensive
Physical interface: irb, Enabled, Physical link is Up
 Interface index: 129, SNMP ifIndex: 23, Generation: 130
 Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
 Speed: Unspecified
 Device flags : Present Running
 Interface flags: SNMP-Traps
 Link type : Full-Duplex
 Link flags : None
 Physical info : Unspecified
 Hold-times : Up 0 ms, Down 0 ms
 Current address: 02:00:00:00:00:30, Hardware address: 02:00:00:00:00:30
 Alternate link address: Unspecified
 Last flapped : Never
 Statistics last cleared: Never
 Traffic statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets : 0
 Output packets: 0
 IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets : 0
 Output packets: 0
 Input errors:
 Errors: 0, Drops: 0, Framing errors: 0, Runt: 0, Giants: 0, Policed discards:
0, Resource errors: 0
 Output errors:
 Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
0

Logical interface irb.0 (Index 68) (SNMP ifIndex 70) (Generation 143)
 Flags: Hardware-Down SNMP-Traps 0x4000 Encapsulation: ENET2
 Bandwidth: 1000mbps
 Routing Instance: customer_0 Bridging Domain: bd0

```



```

Traffic statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Local statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Transit statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps
IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Protocol inet, MTU: 1500, Generation: 154, Route table: 0
 Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
 Destination: 10.51.1/24, Local: 10.51.1.2, Broadcast: 10.51.1.255,
 Generation: 155
Protocol multiservice, MTU: 1500, Generation: 155, Route table: 0
 Flags: Is-Primary
 Policer: Input: __default_arp_policer

```

#### show interfaces irb snmp-index

```

user@host> show interfaces irb snmp-index 25
Physical interface: irb, Enabled, Physical link is Up
 Interface index: 128, SNMP ifIndex: 25
 Type: Ethernet, Link-level type: Ethernet, MTU: 1514
 Device flags : Present Running
 Interface flags: SNMP-Traps
 Link type : Full-Duplex
 Link flags : None
 Current address: 02:00:00:00:00:30, Hardware address: 02:00:00:00:00:30
 Last flapped : Never
 Input packets : 0
 Output packets: 0

Logical interface irb.0 (Index 68) (SNMP ifIndex 70)
 Flags: Hardware-Down SNMP-Traps 0x4000 Encapsulation: ENET2
 Bandwidth: 1000mbps
 Routing Instance: customer_0 Bridging Domain: bd0
 Input packets : 0
 Output packets: 0
 Protocol inet, MTU: 1500
 Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
 Destination: 10.51.1/24, Local: 10.51.1.2, Broadcast: 10.51.1.255
 Protocol multiservice, MTU: 1500
 Flags: Is-Primary

```

## show interfaces queue

---

**Syntax**    show interfaces queue  
              <aggregate | remaining-traffic>  
              <both-ingress-egress>  
              <egress>  
              <forwarding-class *forwarding-class*>  
              <ingress>  
              <interface-name *interface-name*>  
              <l2-statistics>

**Release Information**    Command introduced before Junos OS Release 7.4.  
                              **both-ingress-egress**, **egress**, and **ingress** options introduced in Junos OS Release 7.6.  
                              Command introduced in Junos OS Release 11.1 for the QFX Series.  
                              **l2-statistics** option introduced in Junos OS Release 12.1.  
                              Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description**    Display class-of-service (CoS) queue information for physical interfaces.

**Options**    **none**—Show detailed CoS queue statistics for all physical interfaces.

**aggregate**—(Optional) Display the aggregated queuing statistics of all logical interfaces that have traffic-control profiles configured. (Not on the QFX Series.)

**both-ingress-egress**—(Optional) On Gigabit Ethernet Intelligent Queuing 2 (IQ2) PICs, display both ingress and egress queue statistics. (Not on the QFX Series.)

**egress**—(Optional) Display egress queue statistics.

**forwarding-class *forwarding-class***—(Optional) Forwarding class name for this queue. Shows detailed CoS statistics for the queue associated with the specified forwarding class.

**ingress**—(Optional) On Gigabit Ethernet IQ2 PICs, display ingress queue statistics. (Not on the QFX Series.)

**interface-name *interface-name***—(Optional) Show detailed CoS queue statistics for the specified interface.

**l2-statistics**—(Optional) Display Layer 2 statistics for MLPPP, FRF.15, and FRF.16 bundles

**remaining-traffic**—(Optional) Display the remaining-traffic queue statistics of all logical interfaces that have traffic-control profiles configured.

### Overhead for Layer 2 Statistics

Transmitted packets and transmitted byte counts are displayed for the Layer 2 level with the addition of encapsulation overheads applied for fragmentation, as shown in [Table 280](#). Others counters, such as packets and bytes queued (input) and drop counters, are displayed at the Layer 3 level. In the case of link fragmentation and interleaving (LFI) for which fragmentation is not applied, corresponding Layer 2 overheads are added, as shown in [Table 280](#).

Table 280: Layer 2 Overhead and Transmitted Packets or Byte Counts

| Protocol       | Fragmentation       |                                   | LFI |
|----------------|---------------------|-----------------------------------|-----|
|                | First fragmentation | Second to <i>n</i> fragmentations |     |
|                | Bytes               | Bytes                             |     |
| MLPPP (Long)   | 13                  | 12                                | 8   |
| MLPPP (short)  | 11                  | 10                                | 8   |
| MLFR (FRF15)   | 12                  | 10                                | 8   |
| MFR (FRF16)    | 10                  | 8                                 | -   |
| MCMLPPP(Long)  | 13                  | 12                                | -   |
| MCMLPPP(Short) | 11                  | 10                                | -   |

## Layer 2 Statistics—Fragmentation Overhead Calculation

## MLPPP/MC-MLPPP Overhead details:

=====

## Fragment 1:

```

Outer PPP header : 4 bytes
Long or short sequence MLPPP header : 4 bytes or 2 bytes
Inner PPP header : 1 byte
HDLC flag and FCS bytes : 4 bytes

```

## Fragments 2 .. n :

```

Outer PPP header : 4 bytes
Long or short sequence MLPPP header : 4 bytes or 2 bytes
HDLC flag and FCS bytes : 4 bytes

```

## MLFR (FRF15) Overhead details:

=====

## Fragment 1:

```

Framereley header : 2 bytes
Control,NLPID : 2 bytes
Fragmentaion header : 2 bytes
Inner proto : 2 bytes
HDLC flag and FCS : 4 bytes

```

## Fragments 2 ...n :

```

Framereley header : 2 bytes
Control,NLPID : 2 bytes
Fragmentaion header : 2 bytes
HDLC flag and FCS : 4 bytes

```

## MFR (FRF16) Overhead details:

=====

```
Fragment 1:
 Fragmentation header : 2 bytes
 Framereelay header : 2 bytes
 Inner proto : 2 bytes
 HDLC flag and FCS : 4 bytes

Fragments 2 ...n :
 Fragmentation header : 2 bytes
 Framereelay header : 2 bytes
 HDLC flag and FCS : 4 bytes
```

## Overhead with LFI

```
MLPPP(Long & short sequence):
=====
 Outer PPP header : 4 bytes
 HDLC flag and FCS : 4 bytes

MLFR (FRF15):
=====
 Framereelay header : 2 bytes
 Control,NLPID : 2 bytes
 HDLC flag and FCS : 4 bytes
```

The following examples show overhead for different cases:

- A 1000-byte packet is sent to a mlppp bundle without any fragmentation. At the Layer 2 level, bytes transmitted is 1013 in 1 packet. This overhead is for MLPPP long sequence encap.
- A 1000-byte packet is sent to a mlppp bundle with a fragment threshold of 250byte. At the Layer 2 level, bytes transmitted is 1061 bytes in 5 packets.
- A 1000-byte LFI packet is sent to an mlppp bundle. At the Layer 2 level, bytes transmitted is 1008 in 1 packet.

**remaining-traffic**—(Optional) Display the queuing statistics of all logical interfaces that do not have traffic-control profiles configured. (Not on the QFX Series.)

## Additional Information

For rate-limited interfaces hosted on Modular Interface Cards (MICs), Modular Port Concentrators (MPCs), or Enhanced Queuing DPCs, rate-limit packet-drop operations occur *before* packets are queued for transmission scheduling. For such interfaces, the statistics for queued traffic do not include the packets that have already been dropped due to rate limiting, and consequently the displayed statistics for queued traffic are the same as the displayed statistics for transmitted traffic.



**NOTE:** For rate-limited interfaces hosted on other types of hardware, rate-limit packet-drop operations occur *after* packets are queued for transmission scheduling. For these other interface types, the statistics for queued traffic include the packets that are later dropped due to rate limiting, and consequently the displayed statistics for queued traffic equals the sum of the statistics for transmitted and rate-limited traffic.

---

On M Series routers (except for the M320 and M120 routers), this command is valid only for a PIC installed on an enhanced Flexible PIC Concentrator (FPC).

Queue statistics for aggregated interfaces are supported on the M Series and T Series routers only. Statistics for an aggregated interface are the summation of the queue statistics of the child links of that aggregated interface. You can view the statistics for a child interface by using the **show interfaces statistics** command for that child interface.

When you configure tricolor marking on a 10-port 1-Gigabit Ethernet PIC, for queues 6 and 7 only, the output does not display the number of queued bytes and packets, or the number of bytes and packets dropped because of RED. If you do not configure tricolor marking on the interface, these statistics are available for all queues.

For the 4-port Channelized OC12 IQE PIC and 1-port Channelized OC48 IQE PIC, the **Packet Forwarding Engine Chassis Queues** field represents traffic bound for a particular physical interface on the PIC. For all other PICs, the **Packet Forwarding Engine Chassis Queues** field represents the total traffic bound for the PIC.

For Gigabit Ethernet IQ2 PICs, the **show interfaces queue** command output does not display the number of tail-dropped packets. This limitation does not apply to Packet Forwarding Engine chassis queues.

When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (under the **Packet Forwarding Engine Chassis Queues** field) shows the prefragmentation values.

The behavior of the **egress** queues for the **Routing Engine-Generated Traffic** is not same as the configured queue for MLPPP and MFR configurations.

For information about how to configure CoS, see the *Junos OS Network Interfaces Library for Routing Devices*. For related CoS operational mode commands, see the [CLI Explorer](#).

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>List of Sample Output</b>    | <a href="#">show interfaces queue (Rate-Limited Interface on a Gigabit Ethernet MIC in an MPC) on page 3154</a><br><a href="#">show interfaces queue (Aggregated Ethernet on a T320 Router) on page 3155</a><br><a href="#">show interfaces queue (Gigabit Ethernet on a T640 Router) on page 3157</a><br><a href="#">show interfaces queue aggregate (Gigabit Ethernet Enhanced DPC) on page 3157</a><br><a href="#">show interfaces queue (Gigabit Ethernet IQ2 PIC) on page 3161</a><br><a href="#">show interfaces queue both-ingress-egress (Gigabit Ethernet IQ2 PIC) on page 3164</a><br><a href="#">show interfaces queue ingress (Gigabit Ethernet IQ2 PIC) on page 3166</a><br><a href="#">show interfaces queue egress (Gigabit Ethernet IQ2 PIC) on page 3167</a><br><a href="#">show interfaces queue remaining-traffic (Gigabit Ethernet Enhanced DPC) on page 3169</a><br><a href="#">show interfaces queue (Channelized OC12 IQE Type 3 PIC in SONET Mode) on page 3171</a><br><a href="#">show interfaces queue (QFX Series) on page 3181</a><br><a href="#">show interfaces queue l2-statistics (lsq interface) on page 3182</a><br><a href="#">show interfaces queue lsq (lsq-ifd) on page 3183</a> |
| <b>Output Fields</b>            | Table 281 lists the output fields for the <b>show interfaces queue</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

Table 281: show interfaces queue Output Fields

| Field Name                                                                                                                                          | Field Description                                                                                                                                                                                                                                                                                                                                                         |
|-----------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Physical interface</b>                                                                                                                           | Name of the physical interface.                                                                                                                                                                                                                                                                                                                                           |
| <b>Enabled</b>                                                                                                                                      | State of the interface. Possible values are described in the "Enabled Field" section under <i>Common Output Fields Description</i> .                                                                                                                                                                                                                                      |
| <b>Interface index</b>                                                                                                                              | Physical interface's index number, which reflects its initialization sequence.                                                                                                                                                                                                                                                                                            |
| <b>SNMP ifIndex</b>                                                                                                                                 | SNMP index number for the interface.                                                                                                                                                                                                                                                                                                                                      |
| <b>Forwarding classes supported</b>                                                                                                                 | Total number of forwarding classes supported on the specified interface.                                                                                                                                                                                                                                                                                                  |
| <b>Forwarding classes in use</b>                                                                                                                    | Total number of forwarding classes in use on the specified interface.                                                                                                                                                                                                                                                                                                     |
| <b>Ingress queues supported</b>                                                                                                                     | On Gigabit Ethernet IQ2 PICs only, total number of ingress queues supported on the specified interface.                                                                                                                                                                                                                                                                   |
| <b>Ingress queues in use</b>                                                                                                                        | On Gigabit Ethernet IQ2 PICs only, total number of ingress queues in use on the specified interface.                                                                                                                                                                                                                                                                      |
| <b>Output queues supported</b>                                                                                                                      | Total number of output queues supported on the specified interface.                                                                                                                                                                                                                                                                                                       |
| <b>Output queues in use</b>                                                                                                                         | Total number of output queues in use on the specified interface.                                                                                                                                                                                                                                                                                                          |
| <b>Egress queues supported</b>                                                                                                                      | Total number of egress queues supported on the specified interface.                                                                                                                                                                                                                                                                                                       |
| <b>Egress queues in use</b>                                                                                                                         | Total number of egress queues in use on the specified interface.                                                                                                                                                                                                                                                                                                          |
| <b>Queue counters (Ingress)</b>                                                                                                                     | CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces. <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul> |
| <b>Burst size</b>                                                                                                                                   | (Logical interfaces on IQ PICs only) Maximum number of bytes up to which the logical interface can burst. The burst size is based on the shaping rate applied to the interface.                                                                                                                                                                                           |
| The following output fields are applicable to both interface component and Packet Forwarding component in the <b>show interfaces queue</b> command: |                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Queue</b>                                                                                                                                        | Queue number.                                                                                                                                                                                                                                                                                                                                                             |
| <b>Forwarding classes</b>                                                                                                                           | Forwarding class name.                                                                                                                                                                                                                                                                                                                                                    |

Table 281: show interfaces queue Output Fields (*continued*)

| Field Name                  | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Queued Packets</b>       | <p>Number of packets queued to this queue.</p> <p><b>NOTE:</b> For Gigabit Ethernet IQ2 interfaces, the Queued Packets count is calculated by the Junos OS interpreting one frame buffer as one packet. If the queued packets are very large or very small, the calculation might not be completely accurate for transit traffic. The count is completely accurate for traffic terminated on the router.</p> <p>For rate-limited interfaces hosted on MICs or MPCs only, this statistic does not include traffic dropped due to rate limiting. For more information, see <a href="#">“Additional Information” on page 3148</a>.</p> |
| <b>Queued Bytes</b>         | <p>Number of bytes queued to this queue. The byte counts vary by interface hardware. For more information, see <a href="#">Table 282</a>.</p> <p>For rate-limited interfaces hosted on MICs or MPCs only, this statistic does not include traffic dropped due to rate limiting. For more information, see <a href="#">“Additional Information” on page 3148</a>.</p>                                                                                                                                                                                                                                                                |
| <b>Transmitted Packets</b>  | <p>Number of packets transmitted by this queue. When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (displayed under the <b>Packet Forwarding Engine Chassis Queues</b> field) shows the prefragmentation values.</p> <p><b>NOTE:</b> For Layer 2 statistics, see <a href="#">“Overhead for Layer 2 Statistics” on page 3146</a></p>                                                                                                                                                                                          |
| <b>Transmitted Bytes</b>    | <p>Number of bytes transmitted by this queue. The byte counts vary by interface hardware. For more information, see <a href="#">Table 282</a>.</p> <p><b>NOTE:</b> On MX Series routers, this number can be inaccurate when you issue the command for a physical interface repeatedly and in quick succession, because the statistics for the child nodes are collected infrequently. Wait ten seconds between successive iterations to avoid this situation.</p> <p><b>NOTE:</b> For Layer 2 statistics, see <a href="#">“Overhead for Layer 2 Statistics” on page 3146</a></p>                                                    |
| <b>Tail-dropped packets</b> | <p>Number of packets dropped because of tail drop.</p> <p><b>NOTE:</b> The <b>Tail-dropped packets</b> counter is not supported on the PTX Series Packet Transport Routers.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>RL-dropped packets</b>   | <p>Number of packets dropped due to rate limiting.</p> <p>For rate-limited interfaces hosted on MICs, MPCs, and Enhanced Queuing DPCs only, this statistic is not included in the queued traffic statistics. For more information, see <a href="#">“Additional Information” on page 3148</a>.</p> <p><b>NOTE:</b> The <b>RL-dropped packets</b> counter is not supported on the PTX Series Packet Transport Routers, and is omitted from the output.</p>                                                                                                                                                                            |
| <b>RL-dropped bytes</b>     | <p>Number of bytes dropped due to rate limiting.</p> <p>For rate-limited interfaces hosted on MICs, MPCs, and Enhanced Queuing DPCs only, this statistic is not included in the queued traffic statistics. For more information, see <a href="#">“Additional Information” on page 3148</a>.</p>                                                                                                                                                                                                                                                                                                                                     |

Table 281: show interfaces queue Output Fields (*continued*)

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>RED-dropped packets</b> | <p>Number of packets dropped because of random early detection (RED).</p> <ul style="list-style-type: none"> <li>(M Series and T Series routers only) On M320 and M120 routers and the T Series routers, the total number of dropped packets is displayed. On all other M Series routers, the output classifies dropped packets into the following categories: <ul style="list-style-type: none"> <li><b>Low, non-TCP</b>—Number of low-loss priority non-TCP packets dropped because of RED.</li> <li><b>Low, TCP</b>—Number of low-loss priority TCP packets dropped because of RED.</li> <li><b>High, non-TCP</b>—Number of high-loss priority non-TCP packets dropped because of RED.</li> <li><b>High, TCP</b>—Number of high-loss priority TCP packets dropped because of RED.</li> </ul> </li> <li>(MX Series routers with enhanced DPCs, and T Series routers with enhanced FPCs only) The output classifies dropped packets into the following categories: <ul style="list-style-type: none"> <li><b>Low</b>—Number of low-loss priority packets dropped because of RED.</li> <li><b>Medium-low</b>—Number of medium-low loss priority packets dropped because of RED.</li> <li><b>Medium-high</b>—Number of medium-high loss priority packets dropped because of RED.</li> <li><b>High</b>—Number of high-loss priority packets dropped because of RED.</li> </ul> </li> </ul> <p><b>NOTE:</b> Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), this field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p> |
| <b>RED-dropped bytes</b>   | <p>Number of bytes dropped because of RED. The byte counts vary by interface hardware. For more information, see <a href="#">Table 282</a>.</p> <ul style="list-style-type: none"> <li>(M Series and T Series routers only) On M320 and M120 routers and the T Series routers, only the total number of dropped bytes is displayed. On all other M Series routers, the output classifies dropped bytes into the following categories: <ul style="list-style-type: none"> <li><b>Low, non-TCP</b>—Number of low-loss priority non-TCP bytes dropped because of RED.</li> <li><b>Low, TCP</b>—Number of low-loss priority TCP bytes dropped because of RED.</li> <li><b>High, non-TCP</b>—Number of high-loss priority non-TCP bytes dropped because of RED.</li> <li><b>High, TCP</b>—Number of high-loss priority TCP bytes dropped because of RED.</li> </ul> </li> </ul> <p><b>NOTE:</b> Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), this field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

Byte counts vary by interface hardware. [Table 282](#) shows how the byte counts on the outbound interfaces vary depending on the interface hardware. [Table 282](#) is based on the assumption that outbound interfaces are sending IP traffic with 478 bytes per packet.



Table 282: Byte Count by Interface Hardware

| Interface Hardware               | Output Level                | Byte Count Includes                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Comments                                                                                                                                                                                                     |
|----------------------------------|-----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Gigabit Ethernet IQ and IQE PICs | Interface                   | <p>Queued: 490 bytes per packet, representing 478 bytes of Layer 3 packet + 12 bytes</p> <p>Transmitted: 490 bytes per packet, representing 478 bytes of Layer 3 packet + 12 bytes</p> <p>RED dropped: 496 bytes per packet representing 478 bytes of Layer 3 packet + 18 bytes</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <p>The 12 additional bytes include 6 bytes for the destination MAC address + 4 bytes for the VLAN + 2 bytes for the Ethernet type.</p> <p>For RED dropped, 6 bytes are added for the source MAC address.</p> |
|                                  | Packet forwarding component | <p>Queued: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p> <p>Transmitted: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | —                                                                                                                                                                                                            |
| Non-IQ PIC                       | Interface                   | <p>T Series, TX Series, T1600, and MX Series routers:</p> <ul style="list-style-type: none"> <li>• Queued: 478 bytes of Layer 3 packet.</li> <li>• Transmitted: 478 bytes of Layer 3 packet.</li> </ul> <p>T4000 routers with Type 5 FPCs :</p> <ul style="list-style-type: none"> <li>• Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead 8 bytes preamble + 12 bytes Inter frame Gap.</li> <li>• Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead 8 bytes preamble + 12 bytes Interframe Gap.</li> </ul> <p>M Series routers:</p> <ul style="list-style-type: none"> <li>• Queued: 478 bytes of Layer 3 packet.</li> <li>• Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead.</li> </ul> <p>PTX Series Packet Transport Routers:</p> <ul style="list-style-type: none"> <li>• Queued: The sum of the transmitted bytes and the RED dropped bytes.</li> <li>• Transmitted: Full Layer 2 overhead (including all L2 encapsulation and CRC) + 12 inter-packet gap + 8 for the preamble.</li> <li>• RED dropped: Full Layer 2 overhead (including all L2 encapsulation and CRC) + 12 inter-packet gap + 8 for the preamble (does not include the VLAN header or MPLS pushed bytes).</li> </ul> | <p>The Layer 2 overhead is 14 bytes for non-VLAN traffic and 18 bytes for VLAN traffic.</p>                                                                                                                  |

Table 282: Byte Count by Interface Hardware (*continued*)

| Interface Hardware                                   | Output Level                | Byte Count Includes                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Comments                                                                                                                           |
|------------------------------------------------------|-----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------|
| IQ and IQE PICs with a SONET/SDH interface           | Interface                   | <p>Queued: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes</p> <p>Transmitted: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes</p> <p>RED dropped: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes</p>                                                                                                                                                                                                                       | The additional 4 bytes are for the Layer 2 Point-to-Point Protocol (PPP) header.                                                   |
|                                                      | Packet forwarding component | <p>Queued: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p> <p>Transmitted: 486 bytes per packet, representing 478 bytes of Layer 3 packet + 8 bytes</p>                                                                                                                                                                                                                                                                                                                              | For transmitted packets, the additional 8 bytes includes 4 bytes for the PPP header and 4 bytes for a cookie.                      |
| Non-IQ PIC with a SONET/SDH interface                | Interface                   | <p>T Series, TX Series, T1600, and MX Series routers:</p> <ul style="list-style-type: none"> <li>Queued: 478 bytes of Layer 3 packet.</li> <li>Transmitted: 478 bytes of Layer 3 packet.</li> </ul> <p>M Series routers:</p> <ul style="list-style-type: none"> <li>Queued: 478 bytes of Layer 3 packet.</li> <li>Transmitted: 483 bytes per packet, representing 478 bytes of Layer 3 packet + 5 bytes</li> <li>RED dropped: 478 bytes per packet, representing 478 bytes of Layer 3 packet</li> </ul> | For transmitted packets, the additional 5 bytes includes 4 bytes for the PPP header and 1 byte for the packet loss priority (PLP). |
| Interfaces configured with Frame Relay Encapsulation | Interface                   | The default Frame Relay overhead is 7 bytes. If you configure the Frame Check Sequence (FCS) to 4 bytes, then the overhead increases to 10 bytes.                                                                                                                                                                                                                                                                                                                                                       |                                                                                                                                    |
| 1-port 10-Gigabit Ethernet IQ2 and IQ2-E PICs        | Interface                   | <p>Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including CRC.</p> <p>Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including CRC.</p>                                                                                                                                                                                                                                                                                                                        | The Layer 2 overhead is 18 bytes for non-VLAN traffic and 22 bytes for VLAN traffic.                                               |
| 4-port 1G IQ2 and IQ2-E PICs                         | Packet forwarding component | Queued: 478 bytes of Layer 3 packet.                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | —                                                                                                                                  |
| 8-port 1G IQ2 and IQ2-E PICs                         |                             | Transmitted: 478 bytes of Layer 3 packet.                                                                                                                                                                                                                                                                                                                                                                                                                                                               |                                                                                                                                    |

## Sample Output

### show interfaces queue (Rate-Limited Interface on a Gigabit Ethernet MIC in an MPC)

The following example shows queue information for the rate-limited interface ge-4/2/0 on a Gigabit Ethernet MIC in an MPC. For rate-limited queues for interfaces hosted on MICs or MPCs, rate-limit packet drops occur prior to packet output queuing. In the

command output, the nonzero statistics displayed in the **RL-dropped packets** and **RL-dropped bytes** fields quantify the traffic dropped to rate-limit queue 0 output to 10 percent of 1 gigabyte (100 megabits) per second. Because the RL-dropped traffic is not included in the **Queued** statistics, the statistics displayed for queued traffic are the same as the statistics for transmitted traffic.

```
user@host> show interfaces queue ge-4/2/0
Physical interface: ge-4/2/0, Enabled, Physical link is Up
 Interface index: 203, SNMP ifIndex: 1054
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
 Queued:
 Packets : 131300649 141751 pps
 Bytes : 11287964840 99793248 bps
 Transmitted:
 Packets : 131300649 141751 pps
 Bytes : 11287964840 99793248 bps
 Tail-dropped packets : 0 0 pps
 RL-dropped packets : 205050862 602295 pps
 RL-dropped bytes : 13595326612 327648832 bps
 RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
```

### show interfaces queue (Aggregated Ethernet on a T320 Router)

The following example shows that the aggregated Ethernet interface, **ae1**, has traffic on queues **af1** and **af12**:

```
user@host> show interfaces queue ae1
Physical interface: ae1, Enabled, Physical link is Up
 Interface index: 158, SNMP ifIndex: 33 Forwarding classes: 8 supported, 8 in use
Output queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: be
 Queued:
 Packets : 5 0 pps
 Bytes : 242 0 bps
 Transmitted:
 Packets : 5 0 pps
 Bytes : 242 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: af1
 Queued:
 Packets : 42603765 595484 pps
```

```

 Bytes : 5453281920 609776496 bps
 Transmitted:
 Packets : 42603765 595484 pps
 Bytes : 5453281920 609776496 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: ef1
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: nc
 Queued:
 Packets : 45 0 pps
 Bytes : 3930 0 bps
 Transmitted:
 Packets : 45 0 pps
 Bytes : 3930 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 4, Forwarding classes: af11
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 5, Forwarding classes: ef11
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 6, Forwarding classes: af12
 Queued:
 Packets : 31296413 437436 pps
 Bytes : 4005940864 447935200 bps
 Transmitted:
 Packets : 31296413 437436 pps
 Bytes : 4005940864 447935200 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 7, Forwarding classes: nc2
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps

```

```

Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

#### show interfaces queue (Gigabit Ethernet on a T640 Router)

```

user@host> show interfaces queue
Physical interface: ge-7/0/1, Enabled, Physical link is Up
 Interface index: 150, SNMP ifIndex: 42
 Forwarding classes: 8 supported, 8 in use
 Output queues: 8 supported, 8 in use
 Queue: 0, Forwarding classes: be
 Queued:
 Packets : 13 0 pps
 Bytes : 622 0 bps
 Transmitted:
 Packets : 13 0 pps
 Bytes : 622 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
 Queue: 1, Forwarding classes: af1
 Queued:
 Packets : 1725947945 372178 pps
 Bytes : 220921336960 381110432 bps
 Transmitted:
 Packets : 1725947945 372178 pps
 Bytes : 220921336960 381110432 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
 Queue: 2, Forwarding classes: ef1
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
 Queue: 3, Forwarding classes: nc
 Queued:
 Packets : 571 0 pps
 Bytes : 49318 336 bps
 Transmitted:
 Packets : 571 0 pps
 Bytes : 49318 336 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

```

#### show interfaces queue aggregate (Gigabit Ethernet Enhanced DPC)

```

user@host> show interfaces queue ge-2/2/9 aggregate

```

```

Physical interface: ge-2/2/9, Enabled, Physical link is Up
 Interface index: 238, SNMP ifIndex: 71
 Forwarding classes: 16 supported, 4 in use
 Ingress queues: 4 supported, 4 in use
 Queue: 0, Forwarding classes: best-effort
 Queued:
 Packets : 148450735 947295 pps
 Bytes : 8016344944 409228848 bps
 Transmitted:
 Packets : 76397439 487512 pps
 Bytes : 4125461868 210602376 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 72053285 459783 pps
 Low : 72053285 459783 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes : 3890877444 198626472 bps
 Low : 3890877444 198626472 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
 Queue: 1, Forwarding classes: expedited-forwarding
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
 Queue: 2, Forwarding classes: assured-forwarding
 Queued:
 Packets : 410278257 473940 pps
 Bytes : 22156199518 204742296 bps
 Transmitted:
 Packets : 4850003 4033 pps
 Bytes : 261900162 1742256 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 405425693 469907 pps
 Low : 405425693 469907 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes : 21892988124 203000040 bps
 Low : 21892988124 203000040 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
 Queue: 3, Forwarding classes: network-control
 Queued:
 Packets : 0 0 pps

```

```

Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
Low : 0 0 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 0 0 pps
RED-dropped bytes : 0 0 bps
Low : 0 0 bps
Medium-low : 0 0 bps
Medium-high : 0 0 bps
High : 0 0 bps
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets : 76605230 485376 pps
Bytes : 5209211400 264044560 bps
Transmitted:
Packets : 76444631 484336 pps
Bytes : 5198235612 263478800 bps
Tail-dropped packets : Not Available
RED-dropped packets : 160475 1040 pps
Low : 160475 1040 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 0 0 pps
RED-dropped bytes : 10912300 565760 bps
Low : 10912300 565760 bps
Medium-low : 0 0 bps
Medium-high : 0 0 bps
High : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
Low : 0 0 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 0 0 pps
RED-dropped bytes : 0 0 bps
Low : 0 0 bps
Medium-low : 0 0 bps
Medium-high : 0 0 bps
High : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets : 4836136 3912 pps
Bytes : 333402032 2139056 bps
Transmitted:
Packets : 3600866 1459 pps
Bytes : 244858888 793696 bps
Tail-dropped packets : Not Available

```

```

RED-dropped packets : 1225034 2450 pps
 Low : 1225034 2450 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 83302312 1333072 bps
 Low : 83302312 1333072 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps

```

#### Packet Forwarding Engine Chassis Queues:

Queues: 4 supported, 4 in use

Queue: 0, Forwarding classes: best-effort

```

Queued:
 Packets : 77059796 486384 pps
 Bytes : 3544750624 178989576 bps
Transmitted:
 Packets : 77059797 486381 pps
 Bytes : 3544750670 178988248 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps

```

Queue: 1, Forwarding classes: expedited-forwarding

```

Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps

```



```

 High : 0 0 pps
 RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
 Queued:
 Packets : 4846580 3934 pps
 Bytes : 222942680 1447768 bps
 Transmitted:
 Packets : 4846580 3934 pps
 Bytes : 222942680 1447768 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 3, Forwarding classes: network-control
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps

```

#### show interfaces queue (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-7/1/3
Physical interface: ge-7/1/3, Enabled, Physical link is Up
 Interface index: 170, SNMP ifIndex: 70 Forwarding classes: 16 supported, 4 in
 use Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
 Queued:
 Packets : 418390039 10 pps
 Bytes : 38910269752 7440 bps
 Transmitted:
 Packets : 418390039 10 pps
 Bytes : 38910269752 7440 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding

```

```

Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
 Packets : 7055 1 pps
 Bytes : 451552 512 bps
Transmitted:
 Packets : 7055 1 pps
 Bytes : 451552 512 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Forwarding classes: 16 supported, 4 in use Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
 Packets : 1031 0 pps
 Bytes : 143292 0 bps
Transmitted:
 Packets : 1031 0 pps
 Bytes : 143292 0 bps
 Tail-dropped packets : Not Available
 RL-dropped packets : 0 0 pps
 RL-dropped bytes : 0 0 bps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RL-dropped packets : 0 0 pps
 RL-dropped bytes : 0 0 bps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps

```

```

Tail-dropped packets : Not Available
RL-dropped packets : 0 0 pps
RL-dropped bytes : 0 0 bps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets : 77009 11 pps
Bytes : 6894286 7888 bps
Transmitted:
Packets : 77009 11 pps
Bytes : 6894286 7888 bps
Tail-dropped packets : Not Available
RL-dropped packets : 0 0 pps
RL-dropped bytes : 0 0 bps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

#### Packet Forwarding Engine Chassis Queues:

Queues: 4 supported, 4 in use

Queue: 0, Forwarding classes: best-effort

```

Queued:
Packets : 1031 0 pps
Bytes : 147328 0 bps
Transmitted:
Packets : 1031 0 pps
Bytes : 147328 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
Low, non-TCP : 0 0 pps
Low, TCP : 0 0 pps
High, non-TCP : 0 0 pps
High, TCP : 0 0 pps
RED-dropped bytes : 0 0 bps
Low, non-TCP : 0 0 bps
Low, TCP : 0 0 bps
High, non-TCP : 0 0 bps
High, TCP : 0 0 bps

```

Queue: 1, Forwarding classes: expedited-forwarding

```

Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
Low, non-TCP : 0 0 pps
Low, TCP : 0 0 pps
High, non-TCP : 0 0 pps
High, TCP : 0 0 pps
RED-dropped bytes : 0 0 bps
Low, non-TCP : 0 0 bps
Low, TCP : 0 0 bps
High, non-TCP : 0 0 bps
High, TCP : 0 0 bps

```

Queue: 2, Forwarding classes: assured-forwarding

```

Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:

```

```

Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
 Low, non-TCP : 0 0 pps
 Low, TCP : 0 0 pps
 High, non-TCP : 0 0 pps
 High, TCP : 0 0 pps
RED-dropped bytes : 0 0 bps
 Low, non-TCP : 0 0 bps
 Low, TCP : 0 0 bps
 High, non-TCP : 0 0 bps
 High, TCP : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
 Packets : 94386 12 pps
 Bytes : 13756799 9568 bps
Transmitted:
 Packets : 94386 12 pps
 Bytes : 13756799 9568 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 Low, non-TCP : 0 0 pps
 Low, TCP : 0 0 pps
 High, non-TCP : 0 0 pps
 High, TCP : 0 0 pps
 RED-dropped bytes : 0 0 bps
 Low, non-TCP : 0 0 bps
 Low, TCP : 0 0 bps
 High, non-TCP : 0 0 bps
 High, TCP : 0 0 bps

```

#### show interfaces queue both-ingress-egress (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-6/2/0 both-ingress-egress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
 Packets : Not Available
 Bytes : 0 0 bps
Transmitted:
 Packets : 254 0 pps
 Bytes : 16274 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
 Packets : Not Available
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding

```

```

Queued:
 Packets : Not Available
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
 Packets : Not Available
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
 Packets : Not Available
 Bytes : 0 0 bps
Transmitted:
 Packets : 3 0 pps
 Bytes : 126 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
 Packets : Not Available
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
 Packets : Not Available
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
 Packets : Not Available
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

```

```

Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
 Queued:
 Packets : 80564692 0 pps
 Bytes : 3383717100 0 bps
 Transmitted:
 Packets : 80564692 0 pps
 Bytes : 3383717100 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
 Queued:
 Packets : 80564685 0 pps
 Bytes : 3383716770 0 bps
 Transmitted:
 Packets : 80564685 0 pps
 Bytes : 3383716770 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
 Queued:
 Packets : 9397 0 pps
 Bytes : 3809052 232 bps
 Transmitted:
 Packets : 9397 0 pps
 Bytes : 3809052 232 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

```

#### show interfaces queue ingress (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-6/2/0 ingress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
 Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
 Queued:
 Packets : Not Available
 Bytes : 0 0 bps
 Transmitted:
 Packets : 288 0 pps
 Bytes : 18450 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

```

```

Queue: 1, Forwarding classes: expedited-forwarding
 Queued:
 Packets : Not Available
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
 Queued:
 Packets : Not Available
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
 Queued:
 Packets : Not Available
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

```

#### show interfaces queue egress (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-6/2/0 egress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
 Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
 Queued:
 Packets : Not Available
 Bytes : 0 0 bps
 Transmitted:
 Packets : 3 0 pps
 Bytes : 126 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
 Queued:
 Packets : Not Available
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
 Queued:

```

```

Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets : 80564692 0 pps
Bytes : 3383717100 0 bps
Transmitted:
Packets : 80564692 0 pps
Bytes : 3383717100 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
Packets : 80564685 0 pps
Bytes : 3383716770 0 bps
Transmitted:
Packets : 80564685 0 pps
Bytes : 3383716770 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets : 9538 0 pps
Bytes : 3819840 0 bps
Transmitted:
Packets : 9538 0 pps
Bytes : 3819840 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```



## show interfaces queue remaining-traffic (Gigabit Ethernet Enhanced DPC)

```

user@host> show interfaces queue ge-2/2/9 remaining-traffic
Physical interface: ge-2/2/9, Enabled, Physical link is Up
 Interface index: 238, SNMP ifIndex: 71
Forwarding classes: 16 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
 Queued:
 Packets : 110208969 472875 pps
 Bytes : 5951284434 204282000 bps
 Transmitted:
 Packets : 110208969 472875 pps
 Bytes : 5951284434 204282000 bps
 Tail-dropped packets : Not Available
 RED-dropped packets :
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes :
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets :
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes :
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets :
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes :
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps

```

```

 High : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
 Packets : 109355853 471736 pps
 Bytes : 7436199152 256627968 bps
Transmitted:
 Packets : 109355852 471736 pps
 Bytes : 7436198640 256627968 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps

```

```

Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps

```

#### show interfaces queue (Channelized OC12 IQE Type 3 PIC in SONET Mode)

```

user@host> show interfaces queue t3-1/1/0:7
Physical interface: t3-1/1/0:7, Enabled, Physical link is Up

 Interface index: 192, SNMP ifIndex: 1948

 Description: full T3 interface connect to 6ce13 t3-3/1/0:7 for FR testing -
 Lam

 Forwarding classes: 16 supported, 9 in use

 Egress queues: 8 supported, 8 in use

 Queue: 0, Forwarding classes: DEFAULT

 Queued:

 Packets : 214886 13449 pps

 Bytes : 9884756 5164536 bps

 Transmitted:

 Packets : 214886 13449 pps

 Bytes : 9884756 5164536 bps

```

|                        |   |       |
|------------------------|---|-------|
| Tail-dropped packets : | 0 | 0 pps |
| RED-dropped packets :  | 0 | 0 pps |
| Low :                  | 0 | 0 pps |
| Medium-low :           | 0 | 0 pps |
| Medium-high :          | 0 | 0 pps |
| High :                 | 0 | 0 pps |
| RED-dropped bytes :    | 0 | 0 bps |
| Low :                  | 0 | 0 bps |
| Medium-low :           | 0 | 0 bps |
| Medium-high :          | 0 | 0 bps |
| High :                 | 0 | 0 bps |

Queue: 1, Forwarding classes: REALTIME

Queued:

|           |   |       |
|-----------|---|-------|
| Packets : | 0 | 0 pps |
| Bytes :   | 0 | 0 bps |

Transmitted:

|                        |   |       |
|------------------------|---|-------|
| Packets :              | 0 | 0 pps |
| Bytes :                | 0 | 0 bps |
| Tail-dropped packets : | 0 | 0 pps |
| RED-dropped packets :  | 0 | 0 pps |
| Low :                  | 0 | 0 pps |
| Medium-low :           | 0 | 0 pps |
| Medium-high :          | 0 | 0 pps |
| High :                 | 0 | 0 pps |
| RED-dropped bytes :    | 0 | 0 bps |
| Low :                  | 0 | 0 bps |
| Medium-low :           | 0 | 0 bps |
| Medium-high :          | 0 | 0 bps |
| High :                 | 0 | 0 bps |

Queue: 2, Forwarding classes: PRIVATE

## Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

## Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |
| Low                  | : | 0 | 0 bps |
| Medium-low           | : | 0 | 0 bps |
| Medium-high          | : | 0 | 0 bps |
| High                 | : | 0 | 0 bps |

## Queue: 3, Forwarding classes: CONTROL

## Queued:

|         |   |      |       |
|---------|---|------|-------|
| Packets | : | 60   | 0 pps |
| Bytes   | : | 4560 | 0 bps |

## Transmitted:

|                      |   |      |       |
|----------------------|---|------|-------|
| Packets              | : | 60   | 0 pps |
| Bytes                | : | 4560 | 0 bps |
| Tail-dropped packets | : | 0    | 0 pps |
| RED-dropped packets  | : | 0    | 0 pps |
| Low                  | : | 0    | 0 pps |
| Medium-low           | : | 0    | 0 pps |
| Medium-high          | : | 0    | 0 pps |
| High                 | : | 0    | 0 pps |
| RED-dropped bytes    | : | 0    | 0 bps |

|             |   |   |       |
|-------------|---|---|-------|
| Low         | : | 0 | 0 bps |
| Medium-low  | : | 0 | 0 bps |
| Medium-high | : | 0 | 0 bps |
| High        | : | 0 | 0 bps |

Queue: 4, Forwarding classes: CLASS\_B\_OUTPUT

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |
| Low                  | : | 0 | 0 bps |
| Medium-low           | : | 0 | 0 bps |
| Medium-high          | : | 0 | 0 bps |
| High                 | : | 0 | 0 bps |

Queue: 5, Forwarding classes: CLASS\_C\_OUTPUT

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |

|                     |   |   |       |
|---------------------|---|---|-------|
| RED-dropped packets | : | 0 | 0 pps |
| Low                 | : | 0 | 0 pps |
| Medium-low          | : | 0 | 0 pps |
| Medium-high         | : | 0 | 0 pps |
| High                | : | 0 | 0 pps |
| RED-dropped bytes   | : | 0 | 0 bps |
| Low                 | : | 0 | 0 bps |
| Medium-low          | : | 0 | 0 bps |
| Medium-high         | : | 0 | 0 bps |
| High                | : | 0 | 0 bps |

Queue: 6, Forwarding classes: CLASS\_V\_OUTPUT

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |
| Low                  | : | 0 | 0 bps |
| Medium-low           | : | 0 | 0 bps |
| Medium-high          | : | 0 | 0 bps |
| High                 | : | 0 | 0 bps |

Queue: 7, Forwarding classes: CLASS\_S\_OUTPUT, GETS

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
|---------|---|---|-------|

|                      |   |   |       |
|----------------------|---|---|-------|
| Bytes                | : | 0 | 0 bps |
| Transmitted:         |   |   |       |
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |
| Low                  | : | 0 | 0 bps |
| Medium-low           | : | 0 | 0 bps |
| Medium-high          | : | 0 | 0 bps |
| High                 | : | 0 | 0 bps |

Packet Forwarding Engine Chassis Queues:

Queues: 8 supported, 8 in use

Queue: 0, Forwarding classes: DEFAULT

|                      |   |          |             |
|----------------------|---|----------|-------------|
| Queued:              |   |          |             |
| Packets              | : | 371365   | 23620 pps   |
| Bytes                | : | 15597330 | 7936368 bps |
| Transmitted:         |   |          |             |
| Packets              | : | 371365   | 23620 pps   |
| Bytes                | : | 15597330 | 7936368 bps |
| Tail-dropped packets | : | 0        | 0 pps       |
| RED-dropped packets  | : | 0        | 0 pps       |
| Low                  | : | 0        | 0 pps       |
| Medium-low           | : | 0        | 0 pps       |
| Medium-high          | : | 0        | 0 pps       |



|                   |   |   |       |
|-------------------|---|---|-------|
| High              | : | 0 | 0 pps |
| RED-dropped bytes | : | 0 | 0 bps |
| Low               | : | 0 | 0 bps |
| Medium-low        | : | 0 | 0 bps |
| Medium-high       | : | 0 | 0 bps |
| High              | : | 0 | 0 bps |

Queue: 1, Forwarding classes: REALTIME

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |
| Low                  | : | 0 | 0 bps |
| Medium-low           | : | 0 | 0 bps |
| Medium-high          | : | 0 | 0 bps |
| High                 | : | 0 | 0 bps |

Queue: 2, Forwarding classes: PRIVATE

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Transmitted:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

|                        |   |       |
|------------------------|---|-------|
| Tail-dropped packets : | 0 | 0 pps |
| RED-dropped packets :  | 0 | 0 pps |
| Low :                  | 0 | 0 pps |
| Medium-low :           | 0 | 0 pps |
| Medium-high :          | 0 | 0 pps |
| High :                 | 0 | 0 pps |
| RED-dropped bytes :    | 0 | 0 bps |
| Low :                  | 0 | 0 bps |
| Medium-low :           | 0 | 0 bps |
| Medium-high :          | 0 | 0 bps |
| High :                 | 0 | 0 bps |

Queue: 3, Forwarding classes: CONTROL

Queued:

|           |         |        |
|-----------|---------|--------|
| Packets : | 32843   | 0 pps  |
| Bytes :   | 2641754 | 56 bps |

Transmitted:

|                        |         |        |
|------------------------|---------|--------|
| Packets :              | 32843   | 0 pps  |
| Bytes :                | 2641754 | 56 bps |
| Tail-dropped packets : | 0       | 0 pps  |
| RED-dropped packets :  | 0       | 0 pps  |
| Low :                  | 0       | 0 pps  |
| Medium-low :           | 0       | 0 pps  |
| Medium-high :          | 0       | 0 pps  |
| High :                 | 0       | 0 pps  |
| RED-dropped bytes :    | 0       | 0 bps  |
| Low :                  | 0       | 0 bps  |
| Medium-low :           | 0       | 0 bps  |
| Medium-high :          | 0       | 0 bps  |
| High :                 | 0       | 0 bps  |

Queue: 4, Forwarding classes: CLASS\_B\_OUTPUT

## Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

## Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |
| Low                  | : | 0 | 0 bps |
| Medium-low           | : | 0 | 0 bps |
| Medium-high          | : | 0 | 0 bps |
| High                 | : | 0 | 0 bps |

## Queue: 5, Forwarding classes: CLASS\_C\_OUTPUT

## Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

## Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |

|             |   |   |       |
|-------------|---|---|-------|
| Low         | : | 0 | 0 bps |
| Medium-low  | : | 0 | 0 bps |
| Medium-high | : | 0 | 0 bps |
| High        | : | 0 | 0 bps |

Queue: 6, Forwarding classes: CLASS\_V\_OUTPUT

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |
| Low                  | : | 0 | 0 bps |
| Medium-low           | : | 0 | 0 bps |
| Medium-high          | : | 0 | 0 bps |
| High                 | : | 0 | 0 bps |

Queue: 7, Forwarding classes: CLASS\_S\_OUTPUT, GETS

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |

|                     |   |   |       |
|---------------------|---|---|-------|
| RED-dropped packets | : | 0 | 0 pps |
| Low                 | : | 0 | 0 pps |
| Medium-low          | : | 0 | 0 pps |
| Medium-high         | : | 0 | 0 pps |
| High                | : | 0 | 0 pps |
| RED-dropped bytes   | : | 0 | 0 bps |
| Low                 | : | 0 | 0 bps |
| Medium-low          | : | 0 | 0 bps |
| Medium-high         | : | 0 | 0 bps |
| High                | : | 0 | 0 bps |

#### show interfaces queue (QFX Series)

```

user@switch> show interfaces queue xe-0/0/15
Physical interface: xe-0/0/15, Enabled, Physical link is Up
Interface index: 49165, SNMP ifIndex: 539
Forwarding classes: 12 supported, 8 in use
Egress queues: 12 supported, 8 in use
Queue: 0, Forwarding classes: best-effort
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped packets: 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: fcoe
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped packets: 0 0 pps
 Total-dropped bytes : 0 0 bps
0 bps
Queue: 4, Forwarding classes: no-loss
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped packets: 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 7, Forwarding classes: network-control

```

```

Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped packets: 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 8, Forwarding classes: mcast
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped packets: 0 0 pps
 Total-dropped bytes : 0 0 bps

```

#### show interfaces queue l2-statistics (lsq interface)

```

user@switch> show interfaces queue lsq-2/2/0.2 l2-statistics
Logical interface lsq-2/2/0.2 (Index 69) (SNMP ifIndex 1598)
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Burst size: 0
Queue: 0, Forwarding classes: be
Queued:
 Packets : 1 0 pps
 Bytes : 1001 0 bps
Transmitted:
 Packets : 5 0 pps
 Bytes : 1062 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: ef
Queued:
 Packets : 1 0 pps
 Bytes : 1500 0 bps
Transmitted:
 Packets : 6 0 pps
 Bytes : 1573 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: af
Queued:
 Packets : 1 0 pps
 Bytes : 512 0 bps
Transmitted:
 Packets : 3 0 pps
 Bytes : 549 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: nc
Queued:
 Packets : 0 0 pps

```

```

Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
=====

```

### show interfaces queue lsq (lsq-ifd)

```

user@switch> show interfaces queue lsq-1/0/0
Logical interface lsq-1/0/0 (Index 348) (SNMP ifIndex 660)
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Burst size: 0
Queue: 0, Forwarding classes: be
 Queued:
 Packets : 55576 1206 pps
 Bytes : 29622008 5145472 bps
 Transmitted:
 Packets : 55576 1206 pps
 Bytes : 29622008 5145472 bps
 Tail-dropped packets : 0 0 pps
 RL-dropped packets : 0 0 pps
 RL-dropped bytes : 0 0 bps
 RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 1, Forwarding classes: ef
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : 0 0 pps
 RL-dropped packets : 0 0 pps
 RL-dropped bytes : 0 0 bps
 RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 2, Forwarding classes: af
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps

```

```
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : 0 0 pps
RL-dropped packets : 0 0 pps
RL-dropped bytes : 0 0 bps
RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 3, Forwarding classes: nc
Queued:
Packets : 22231 482 pps
Bytes : 11849123 2057600 bps
Transmitted:
Packets : 22231 482 pps
Bytes : 11849123 2057600 bps
Tail-dropped packets : 0 0 pps
RL-dropped packets : 0 0 pps
RL-dropped bytes : 0 0 bps
RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
```



## show interfaces xe

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show interfaces <i>device-name:type-fpc/pic/port</i> &lt;brief   detail   extensive   terse&gt; &lt;descriptions&gt; &lt;media&gt; &lt;routing-instance (all   <i>instance-name</i>)&gt; &lt;snmp-index <i>snmp-index</i>&gt; &lt;statistics&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Display status information about the specified 10-Gigabit Ethernet interface. This command does not display statistics for routed VLAN interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><b><i>device-name:type-fpc/pic/port</i></b>—(QFabric systems only) The device name is either the serial number or the alias of the QFabric system component, such as a Node device, Interconnect device, or QFabric infrastructure. The name must contain a maximum of 128 characters and not contain any colons.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output.</p> <p><b>descriptions</b>—(Optional) Display interface description strings.</p> <p><b>media</b>—(Optional) Display media-specific information about network interfaces.</p> <p><b>routing-instance (all   <i>instance-name</i>)</b>—(Optional) Display the name of an individual routing instance or display all routing instances.</p> <p><b>snmp-index <i>snmp-index</i></b>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><b>statistics</b>—(Optional) Display static interface statistics.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Monitoring Interface Status and Traffic on page 167</a></li> <li>• <a href="#">Troubleshooting Network Interfaces on page 2845</a></li> <li>• <a href="#">Troubleshooting an Aggregated Ethernet Interface on page 2886</a></li> <li>• <a href="#">Junos OS Network Interfaces Library for Routing Devices</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>List of Sample Output</b>    | <p><a href="#">show interfaces on page 3193</a></p> <p><a href="#">show interfaces (Asymmetric Flow Control) on page 3194</a></p> <p><a href="#">show interfaces brief on page 3194</a></p> <p><a href="#">show interfaces detail on page 3194</a></p> <p><a href="#">show interfaces detail (Asymmetric Flow Control) on page 3196</a></p> <p><a href="#">show interfaces extensive on page 3197</a></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

[show interfaces extensive \(Asymmetric Flow Control\) on page 3199](#)

[show interfaces terse on page 3201](#)

[show interfaces \(QFabric System\) on page 3201](#)

**Output Fields** Table 283 lists the output fields for the **show interfaces xe** command. Output fields are listed in the approximate order in which they appear.

**Table 283: show interfaces xe Output Fields**

| Field Name                                                                              | Field Description                                                                                                                                                                             | Level of Output              |
|-----------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Physical Interface</b>                                                               |                                                                                                                                                                                               |                              |
| <b>Physical interface</b>                                                               | Name of the physical interface.                                                                                                                                                               | All levels                   |
| <b>Enabled</b>                                                                          | State of the interface.                                                                                                                                                                       | All levels                   |
| <b>Interface index</b>                                                                  | Index number of the physical interface, which reflects its initialization sequence.                                                                                                           | <b>detail extensive none</b> |
| <b>SNMP ifIndex</b>                                                                     | SNMP index number for the physical interface.                                                                                                                                                 | <b>detail extensive none</b> |
| <b>Generation</b>                                                                       | Unique number for use by Juniper Networks technical support only.                                                                                                                             | <b>detail extensive</b>      |
| <b>Link-level type</b>                                                                  | Encapsulation being used on the physical interface.                                                                                                                                           | All levels                   |
| <b>MTU</b>                                                                              | Maximum transmission unit size on the physical interface.                                                                                                                                     | All levels                   |
| <b>Speed</b>                                                                            | Speed at which the interface is running.                                                                                                                                                      | All levels                   |
| <b>Duplex</b>                                                                           | Duplex mode of the interface, either <b>Full-Duplex</b> or <b>Half-Duplex</b> .                                                                                                               | All levels                   |
| <b>Loopback</b>                                                                         | Loopback status: <b>Enabled</b> or <b>Disabled</b> . If loopback is enabled, type of loopback: <b>Local</b> or <b>Remote</b> .                                                                | All levels                   |
| <b>Source filtering</b>                                                                 | Source filtering status: <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                  | All levels                   |
| <b>LAN-PHY mode</b>                                                                     | 10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications. | All levels                   |
| <b>Unidirectional</b>                                                                   | Unidirectional link mode status for 10-Gigabit Ethernet interface: <b>Enabled</b> or <b>Disabled</b> for parent interface; <b>Rx-only</b> or <b>Tx-only</b> for child interfaces.             | All levels                   |
| <b>Flow control</b>                                                                     | Flow control status: <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                      | All levels                   |
| <b>NOTE:</b> This field is only displayed if asymmetric flow control is not configured. |                                                                                                                                                                                               |                              |

Table 283: show interfaces xe Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Level of Output              |
|--------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Configured-flow-control</b> | Configured flow control for the interface transmit buffers ( <b>tx-buffers</b> ) and receive buffers ( <b>rx-buffers</b> ): <ul style="list-style-type: none"> <li><b>tx-buffers</b>—<b>On</b> if the interface is configured to respond to Ethernet PAUSE messages received from the connected peer.<br/><b>Off</b> if the interface is not configured to respond to received PAUSE messages.</li> <li><b>rx-buffers</b>—<b>On</b> if the interface is configured to generate and send Ethernet PAUSE messages to the connected peer.<br/><b>Off</b> if the interface is not configured to generate and send PAUSE messages.</li> </ul> <p><b>NOTE:</b> This field is only displayed if asymmetric flow control is configured.</p> | All levels                   |
| <b>Auto-negotiation</b>        | Autonegotiation status: <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | All levels                   |
| <b>Remote-fault</b>            | Remote fault status: <ul style="list-style-type: none"> <li><b>Online</b>—Autonegotiation is manually configured as online.</li> <li><b>Offline</b>—Autonegotiation is manually configured as offline.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | All levels                   |
| <b>Device flags</b>            | Information about the physical device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | All levels                   |
| <b>Interface flags</b>         | Information about the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | All levels                   |
| <b>Link flags</b>              | Information about the link.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | All levels                   |
| <b>Wavelength</b>              | Configured wavelength, in nanometers (nm).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | All levels                   |
| <b>Frequency</b>               | Frequency associated with the configured wavelength, in terahertz (THz).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | All levels                   |
| <b>CoS queues</b>              | Number of CoS queues configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail extensive none</b> |
| <b>Schedulers</b>              | Number of CoS schedulers configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>extensive</b>             |
| <b>Hold-times</b>              | Current interface hold-time up and hold-time down, in milliseconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>detail extensive</b>      |
| <b>Current address</b>         | Configured MAC address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail extensive none</b> |
| <b>Hardware address</b>        | Hardware MAC address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive none</b> |
| <b>Last flapped</b>            | Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago)</b> . For example, <b>Last flapped: 2008-01-16 10:52:40 UTC (3d 22:58 ago)</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive none</b> |
| <b>Input Rate</b>              | Input rate in bits per second (bps) and packets per second (pps).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | None specified               |
| <b>Output Rate</b>             | Output rate in bps and pps.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | None specified               |
| <b>Statistics last cleared</b> | Time when the statistics for the interface were last set to zero.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>detail extensive</b>      |

Table 283: show interfaces xe Output Fields (*continued*)

| Field Name                | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Level of Output         |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Traffic statistics</b> | <p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> <p><b>NOTE:</b> The bandwidth bps counter is not enabled.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail extensive</b> |
| <b>Input errors</b>       | <p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Errors</b>—Sum of the incoming frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li>• <b>Runts</b>—Number of frames received that are smaller than the runt threshold.</li> <li>• <b>Policed discards</b>—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle.</li> <li>• <b>L3 incompletes</b>—Number of incoming packets discarded because they failed Layer 3 sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored if you configure the <b>ignore-l3-incompletes</b> statement.</li> <li>• <b>L2 channel errors</b>—Number of times the software did not find a valid logical interface for an incoming frame.</li> <li>• <b>L2 mismatch timeouts</b>—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable.</li> <li>• <b>FIFO errors</b>—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul> | <b>extensive</b>        |

Table 283: show interfaces xe Output Fields (*continued*)

| Field Name                      | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Level of Output         |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Output errors</b>            | <p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Carrier transitions</b>—Number of times the interface has gone from <b>down</b> to <b>up</b>. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning.</li> <li>• <b>Errors</b>—Sum of the outgoing frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Collisions</b>—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug.</li> <li>• <b>Aged packets</b>—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware.</li> <li>• <b>FIFO errors</b>—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>HS link CRC errors</b>—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces.</li> <li>• <b>MTU errors</b>—Number of packets whose size exceeded the MTU of the interface.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul> | <b>extensive</b>        |
| <b>Egress queues</b>            | Total number of egress queues supported on the specified interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b> |
| <b>Queue counters (Egress)</b>  | <p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b> |
| <b>Queue Number</b>             | The CoS queue number and the forwarding classes mapped to the queue number. The <b>Mapped forwarding class</b> column lists the forwarding classes mapped to each CoS queue.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>detail extensive</b> |
| <b>Ingress queues</b>           | Total number of ingress queues supported on the specified interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>extensive</b>        |
| <b>Queue counters (Ingress)</b> | <p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>extensive</b>        |

Table 283: show interfaces xe Output Fields (*continued*)

| Field Name                              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Level of Output              |
|-----------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Active alarms and Active defects</b> | <p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the switch configuration, an alarm can ring the red or yellow alarm bell on the switch, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value <b>None</b> or <b>Link</b>.</p> <ul style="list-style-type: none"> <li>• <b>None</b>—There are no active defects or alarms.</li> <li>• <b>Link</b>—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail extensive none</b> |
| <b>PCS statistics</b>                   | Physical Coding Sublayer (PCS) fault conditions from the LAN PHY device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>detail extensive</b>      |
| <b>MAC statistics</b>                   | <p>Receive and Transmit statistics reported by the PIC's MAC subsystem.</p> <ul style="list-style-type: none"> <li>• <b>Total octets and total packets</b>—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type.</li> <li>• <b>Unicast packets, Broadcast packets, and Multicast packets</b>—Number of unicast, broadcast, and multicast packets.</li> <li>• <b>CRC/Align errors</b>—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).</li> <li>• <b>FIFO error</b>—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>MAC control frames</b>—Number of MAC control frames.</li> <li>• <b>MAC pause frames</b>—Number of MAC control frames with <b>pause</b> operational code.</li> <li>• <b>Oversized frames</b>—Number of packets that exceeds the configured MTU.</li> <li>• <b>Jabber frames</b>—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms.</li> <li>• <b>Fragment frames</b>—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runs (which are normal occurrences caused by collisions) and noise hits are counted.</li> <li>• <b>VLAN tagged frames</b>—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. This counter is not supported on EX Series switches and is always displayed as 0.</li> <li>• <b>Code violations</b>—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error."</li> </ul> | <b>extensive</b>             |
| <b>Filter statistics</b>                | Receive and Transmit statistics reported by the PIC's MAC address filter subsystem.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>extensive</b>             |

Table 283: show interfaces xe Output Fields (*continued*)

| Field Name                  | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Level of Output |
|-----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Autonegotiation information | <p>Information about link autonegotiation.</p> <ul style="list-style-type: none"> <li>• <b>Negotiation status:</b> <ul style="list-style-type: none"> <li>• <b>Incomplete</b>—Ethernet interface has the speed or link mode configured.</li> <li>• <b>No autonegotiation</b>—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.</li> <li>• <b>Complete</b>—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.</li> </ul> </li> <li>• <b>Link partner status</b>—OK when the Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.</li> <li>• <b>Link partner:</b> <ul style="list-style-type: none"> <li>• <b>Link mode</b>—Depending on the capability of the attached Ethernet device, either <b>Full-duplex</b> or <b>Half-duplex</b>.</li> <li>• <b>Flow control</b>—Types of flow control supported by the remote Ethernet device. For Fast Ethernet interfaces, the type is <b>None</b>. For Gigabit Ethernet interfaces, types are <b>Symmetric</b> (link partner supports <b>PAUSE</b> on receive and transmit), <b>Asymmetric</b> (link partner supports <b>PAUSE</b> on transmit), and <b>Symmetric/Asymmetric</b> (link partner supports both <b>PAUSE</b> on receive and transmit or only <b>PAUSE</b> receive).</li> <li>• <b>Remote fault</b>—Remote fault information from the link partner—<b>Failure</b> indicates a receive link error. <b>OK</b> indicates that the link partner is receiving. <b>Negotiation error</b> indicates a negotiation error. <b>Offline</b> indicates that the link partner is going offline.</li> </ul> </li> <li>• <b>Local resolution:</b> <ul style="list-style-type: none"> <li>• <b>Flow control</b>—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are <b>Symmetric</b> (link partner supports <b>PAUSE</b> on receive and transmit), <b>Asymmetric</b> (link partner supports <b>PAUSE</b> on transmit), and <b>Symmetric/Asymmetric</b> (link partner supports both <b>PAUSE</b> on receive and transmit or only <b>PAUSE</b> receive). For asymmetric <b>PAUSE</b>, shows if the <b>PAUSE</b> transmit and <b>PAUSE</b> receive states on the interface are <b>enable</b> or <b>disable</b>.</li> <li>• <b>Remote fault</b>—Remote fault information. <b>Link OK</b> (no error detected on receive), <b>Offline</b> (local interface is offline), and <b>Link Failure</b> (link error detected on receive).</li> </ul> </li> </ul> | extensive       |

Table 283: show interfaces xe Output Fields (*continued*)

| Field Name                                    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Level of Output              |
|-----------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Packet Forwarding Engine configuration</b> | Information about the configuration of the Packet Forwarding Engine: <ul style="list-style-type: none"> <li><b>Destination slot</b>—FPC slot number.</li> <li><b>CoS transmit queue</b>—Queue number and its associated user-configured forwarding class name.</li> <li><b>Bandwidth %</b>—Percentage of bandwidth allocated to the queue.</li> <li><b>Bandwidth bps</b>—Bandwidth allocated to the queue (in bps).</li> <li><b>Buffer %</b>—Percentage of buffer space allocated to the queue.</li> <li><b>Buffer usec</b>—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time.</li> <li><b>Priority</b>—Queue priority: <b>low</b> or <b>high</b>.</li> <li><b>Limit</b>—Displayed if rate limiting is configured for the queue. Possible values are <b>none</b> and <b>exact</b>. If <b>exact</b> is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If <b>none</b> is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.</li> </ul> | <b>extensive</b>             |
| <b>Logical Interface</b>                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |                              |
| <b>Logical interface</b>                      | Name of the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | All levels                   |
| <b>Index</b>                                  | Index number of the logical interface, which reflects its initialization sequence.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>detail extensive none</b> |
| <b>SNMP ifIndex</b>                           | SNMP interface index number for the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive none</b> |
| <b>Generation</b>                             | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail extensive</b>      |
| <b>Flags</b>                                  | Information about the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | All levels                   |
| <b>Encapsulation</b>                          | Encapsulation on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | All levels                   |
| <b>Protocol</b>                               | Protocol family.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>detail extensive none</b> |
| <b>Traffic statistics</b>                     | Number and rate of bytes and packets received (input) and transmitted (output) on the specified interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>detail extensive</b>      |
| <b>IPv6 transit statistics</b>                | If IPv6 statics tracking is enabled, number of IPv6 bytes and packets received and transmitted on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>extensive</b>             |
| <b>Local statistics</b>                       | Number and rate of bytes and packets destined to and from the switch.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>extensive</b>             |
| <b>Transit statistics</b>                     | Number and rate of bytes and packets transiting the switch.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>extensive</b>             |
| <b>Generation</b>                             | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail extensive</b>      |
| <b>Route Table</b>                            | Route table in which the logical interface address is located. For example, <b>0</b> refers to the routing table inet.0.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail extensive none</b> |



Table 283: show interfaces xe Output Fields (*continued*)

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                                                                                                  | Level of Output              |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Input Filters</b>    | Names of any input filters applied to this interface.                                                                                                                                                                                                                                                                                                                                              | <b>detail extensive</b>      |
| <b>Output Filters</b>   | Names of any output filters applied to this interface.                                                                                                                                                                                                                                                                                                                                             | <b>detail extensive</b>      |
| <b>Flags</b>            | Information about protocol family flags.<br><br>If unicast Reverse Path Forwarding (uRPF) is explicitly configured on the specified interface, the uRPF flag appears. If uRPF was configured on a different interface (and therefore is enabled on all switch interfaces) but was not explicitly configured on the specified interface, the uRPF flag does not appear even though uRPF is enabled. | <b>detail extensive</b>      |
| <b>Addresses, Flags</b> | Information about the address flags.                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive none</b> |
| <i>protocol-family</i>  | Protocol family configured on the logical interface. If the protocol is <b>inet</b> , the IP address of the interface is also displayed.                                                                                                                                                                                                                                                           | <b>brief</b>                 |
| <b>Flags</b>            | Information about the address flag.                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive none</b> |
| <b>Destination</b>      | IP address of the remote side of the connection.                                                                                                                                                                                                                                                                                                                                                   | <b>detail extensive none</b> |
| <b>Local</b>            | IP address of the logical interface.                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive none</b> |
| <b>Broadcast</b>        | Broadcast address of the logical interface.                                                                                                                                                                                                                                                                                                                                                        | <b>detail extensive none</b> |
| <b>Generation</b>       | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                  | <b>detail extensive</b>      |

## Sample Output

### show interfaces

```

user@switch> show interfaces xe-0/0/1
Physical interface: xe-0/0/1, Enabled, Physical link is Up
 Interface index: 49195, SNMP ifIndex: 591
 Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
 Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
 Disabled,
 Flow control: Disabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0
 Link flags : None
 CoS queues : 12 supported, 12 maximum usable queues
 Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
 Last flapped : 2011-06-01 00:42:03 PDT (00:02:42 ago)
 Input rate : 0 bps (0 pps)
 Output rate : 0 bps (0 pps)
 Active alarms : None
 Active defects : None

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523)
 Flags: SNMP-Traps 0x0 Encapsulation: ENET2
 Input packets : 0

```

```
Output packets: 0
Protocol eth-switch, MTU: 0
Flags: Trunk-Mode
```

### show interfaces (Asymmetric Flow Control)

```
user@switch> show interfaces xe-0/0/1
Physical interface: xe-0/0/1, Enabled, Physical link is Up
 Interface index: 49195, SNMP ifIndex: 591
 Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
 Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled,
 Configured-flow-control tx-buffers: off rx-buffers: on
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0
 Link flags : None
 CoS queues : 12 supported, 12 maximum usable queues
 Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
 Last flapped : 2011-06-01 00:42:03 PDT (00:02:42 ago)
 Input rate : 0 bps (0 pps)
 Output rate : 0 bps (0 pps)
 Active alarms : None
 Active defects : None

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523)
 Flags: SNMP-Traps 0x0 Encapsulation: ENET2
 Input packets : 0
 Output packets: 0
 Protocol eth-switch, MTU: 0
 Flags: Trunk-Mode
```

### show interfaces brief

```
user@switch> show interfaces xe-0/0/1 brief
Physical interface: xe-0/0/1, Enabled, Physical link is Up
 Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
 Source filtering: Disabled, Flow control: Enabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0
 Link flags : None

Logical interface xe-0/0/1.0
 Flags: SNMP-Traps Encapsulation: ENET2
 eth-switch
```

### show interfaces detail

```
user@switch> show interfaces xe-0/0/1 detail
Physical interface: xe-0/0/1, Enabled, Physical link is Up
 Interface index: 49195, SNMP ifIndex: 591, Generation: 169
 Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
 Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled,
 Flow control: Disabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0
 Link flags : None
 CoS queues : 12 supported, 12 maximum usable queues
 Hold-times : Up 0 ms, Down 0 ms
 Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
```

Last flapped : 2011-06-01 00:42:03 PDT (00:02:50 ago)  
 Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:14 ago)

Traffic statistics:

|                 |   |       |
|-----------------|---|-------|
| Input bytes :   | 0 | 0 bps |
| Output bytes :  | 0 | 0 bps |
| Input packets:  | 0 | 0 pps |
| Output packets: | 0 | 0 pps |

IPv6 transit statistics:

|                 |   |
|-----------------|---|
| Input bytes :   | 0 |
| Output bytes :  | 0 |
| Input packets:  | 0 |
| Output packets: | 0 |

Egress queues: 12 supported, 9 in use

| Queue counters: | Queued packets | Transmitted packets | Dropped packets |
|-----------------|----------------|---------------------|-----------------|
| 0 best-effort   | 0              | 0                   | 0               |
| 1 fc7           | 0              | 0                   | 0               |
| 2 no-loss       | 0              | 0                   | 0               |
| 3 fcoe          | 0              | 0                   | 0               |
| 4 fc4           | 0              | 0                   | 0               |
| 5 fc5           | 0              | 0                   | 0               |
| 6 fc6           | 0              | 0                   | 0               |
| 7 network-cont  | 0              | 0                   | 0               |
| 8 mcast         | 0              | 0                   | 0               |

| Queue number: | Mapped forwarding classes |
|---------------|---------------------------|
| 0             | best-effort               |
| 1             | fc7                       |
| 2             | no-loss                   |
| 3             | fcoe                      |
| 4             | fc4                       |
| 5             | fc5                       |
| 6             | fc6                       |
| 7             | network-control           |
| 8             | mcast                     |

Active alarms : None

Active defects : None

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)

Flags: SNMP-Traps 0x0 Encapsulation: ENET2

Traffic statistics:

|                 |   |
|-----------------|---|
| Input bytes :   | 0 |
| Output bytes :  | 0 |
| Input packets:  | 0 |
| Output packets: | 0 |

Local statistics:

|                 |   |
|-----------------|---|
| Input bytes :   | 0 |
| Output bytes :  | 0 |
| Input packets:  | 0 |
| Output packets: | 0 |

Transit statistics:

|                |   |       |
|----------------|---|-------|
| Input bytes :  | 0 | 0 bps |
| Output bytes : | 0 | 0 bps |

```

Input packets: 0 0 pps
Output packets: 0 0 pps
Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
Flags: Trunk-Mode

```

### show interfaces detail (Asymmetric Flow Control)

```

user@switch> show interfaces xe-0/0/1 detail
Physical interface: xe-0/0/1, Enabled, Physical link is Up
 Interface index: 49195, SNMP ifIndex: 591, Generation: 169
 Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
 Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
 Disabled,
 Configured-flow-control tx-buffers: off rx-buffers: on
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0
 Link flags : None
 CoS queues : 12 supported, 12 maximum usable queues
 Hold-times : Up 0 ms, Down 0 ms
 Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
 Last flapped : 2011-06-01 00:42:03 PDT (00:02:50 ago)
 Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:14 ago)
 Traffic statistics:
 Input bytes : 0 0 bps
 Output bytes: 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps
 IPv6 transit statistics:
 Input bytes : 0
 Output bytes: 0
 Input packets: 0
 Output packets: 0
 Egress queues: 12 supported, 9 in use
 Queue counters:
 Queued packets Transmitted packets Dropped packets

 0 best-effort 0 0 0
 1 fc7 0 0 0
 2 no-loss 0 0 0
 3 fcoe 0 0 0
 4 fc4 0 0 0
 5 fc5 0 0 0
 6 fc6 0 0 0
 7 network-cont 0 0 0
 8 mcast 0 0 0

 Queue number: Mapped forwarding classes
 0 best-effort
 1 fc7
 2 no-loss
 3 fcoe
 4 fc4
 5 fc5
 6 fc6

```

```

7 network-control
8 mcast
Active alarms : None
Active defects : None

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)
Flags: SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Local statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Transit statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps
Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
Flags: Trunk-Mode

```

#### show interfaces extensive

```

user@switch> show interfaces xe-0/0/1 extensive
Physical interface: xe-0/0/1, Enabled, Physical link is Up
Interface index: 49195, SNMP ifIndex: 591, Generation: 169
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled,
Flow control: Disabled
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x0
Link flags : None
CoS queues : 12 supported, 12 maximum usable queues
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
Last flapped : 2011-06-01 00:42:03 PDT (00:03:08 ago)
Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:32 ago)
Traffic statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps
IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Input errors:
 Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
Resource errors: 0
Output errors:
 Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 12 supported, 9 in use
Queue counters: Queued packets Transmitted packets Dropped packets

```

|                |   |   |   |
|----------------|---|---|---|
| 0 best-effort  | 0 | 0 | 0 |
| 1 fc7          | 0 | 0 | 0 |
| 2 no-loss      | 0 | 0 | 0 |
| 3 fcoe         | 0 | 0 | 0 |
| 4 fc4          | 0 | 0 | 0 |
| 5 fc5          | 0 | 0 | 0 |
| 6 fc6          | 0 | 0 | 0 |
| 7 network-cont | 0 | 0 | 0 |
| 8 mcast        | 0 | 0 | 0 |

Queue number:            Mapped forwarding classes

|   |                 |
|---|-----------------|
| 0 | best-effort     |
| 1 | fc7             |
| 2 | no-loss         |
| 3 | fcoe            |
| 4 | fc4             |
| 5 | fc5             |
| 6 | fc6             |
| 7 | network-control |
| 8 | mcast           |

Active alarms : None

Active defects : None

MAC statistics:

|                    | Receive | Transmit |
|--------------------|---------|----------|
| Total octets       | 0       | 0        |
| Total packets      | 0       | 0        |
| Unicast packets    | 0       | 0        |
| Broadcast packets  | 0       | 0        |
| Multicast packets  | 0       | 0        |
| CRC/Align errors   | 0       | 0        |
| FIFO errors        | 0       | 0        |
| MAC control frames | 0       | 0        |
| MAC pause frames   | 0       | 0        |
| Oversized frames   | 0       |          |
| Jabber frames      | 0       |          |
| Fragment frames    | 0       |          |
| VLAN tagged frames | 0       |          |
| Code violations    | 0       |          |

MAC Priority Flow Control Statistics:

|              |   |   |
|--------------|---|---|
| Priority : 0 | 0 | 0 |
| Priority : 1 | 0 | 0 |
| Priority : 2 | 0 | 0 |
| Priority : 3 | 0 | 0 |
| Priority : 4 | 0 | 0 |
| Priority : 5 | 0 | 0 |
| Priority : 6 | 0 | 0 |
| Priority : 7 | 0 | 0 |

Filter statistics:

|                      |   |   |
|----------------------|---|---|
| Input packet count   | 0 |   |
| Input packet rejects | 0 |   |
| Input DA rejects     | 0 |   |
| Input SA rejects     | 0 |   |
| Output packet count  |   | 0 |

```

Output packet pad count 0
Output packet error count 0
CAM destination filters: 1, CAM source filters: 0
Packet Forwarding Engine configuration:
 Destination slot: 0
CoS information:
 Direction : Output
 CoS transmit queue Bandwidth Buffer Priority
Limit
 % bps % usec
0 best-effort 75 7500000000 75 0 low
none
7 network-control 5 500000000 5 0 low
none
8 mcast 20 2000000000 20 0 low
none

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)
Flags: SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Local statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Transit statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps
Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
Flags: Trunk-Mode

```

#### show interfaces extensive (Asymmetric Flow Control)

```

user@switch> show interfaces xe-0/0/1 extensive
Physical interface: xe-0/0/1, Enabled, Physical link is Up
Interface index: 49195, SNMP ifIndex: 591, Generation: 169
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled,
Configured-flow-control tx-buffers: off rx-buffers: on
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x0
Link flags : None
CoS queues : 12 supported, 12 maximum usable queues
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
Last flapped : 2011-06-01 00:42:03 PDT (00:03:08 ago)
Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:32 ago)
Traffic statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps
IPv6 transit statistics:
 Input bytes : 0

```

```

Output bytes : 0
Input packets: 0
Output packets: 0
Input errors:
 Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
Resource errors: 0
Output errors:
 Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 12 supported, 9 in use
Queue counters: Queued packets Transmitted packets Dropped packets

 0 best-effort 0 0 0
 1 fc7 0 0 0
 2 no-loss 0 0 0
 3 fcoe 0 0 0
 4 fc4 0 0 0
 5 fc5 0 0 0
 6 fc6 0 0 0
 7 network-cont 0 0 0
 8 mcast 0 0 0

Queue number: Mapped forwarding classes
 0 best-effort
 1 fc7
 2 no-loss
 3 fcoe
 4 fc4
 5 fc5
 6 fc6
 7 network-control
 8 mcast

Active alarms : None
Active defects : None
MAC statistics:
 Total octets Receive Transmit
 Total packets 0 0
 Unicast packets 0 0
 Broadcast packets 0 0
 Multicast packets 0 0
 CRC/Align errors 0 0
 FIFO errors 0 0
 MAC control frames 0 0
 MAC pause frames 0 0
 Oversized frames 0
 Jabber frames 0
 Fragment frames 0
 VLAN tagged frames 0
 Code violations 0
MAC Priority Flow Control Statistics:
 Priority : 0 0 0
 Priority : 1 0 0

```



```

Priority : 2 0 0
Priority : 3 0 0
Priority : 4 0 0
Priority : 5 0 0
Priority : 6 0 0
Priority : 7 0 0
Filter statistics:
Input packet count 0
Input packet rejects 0
Input DA rejects 0
Input SA rejects 0
Output packet count 0
Output packet pad count 0
Output packet error count 0
CAM destination filters: 1, CAM source filters: 0
Packet Forwarding Engine configuration:
Destination slot: 0
CoS information:
Direction : Output
CoS transmit queue Bandwidth Buffer Priority Limit
 % bps % usec
0 best-effort 75 7500000000 75 0 low none
7 network-control 5 500000000 5 0 low none
8 mcast 20 2000000000 20 0 low none

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)
Flags: SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Local statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Transit statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
Flags: Trunk-Mode

```

### show interfaces terse

```

user@switch> show interfaces xe-0/0/1 terse
Interface Admin Link Proto Local Remote

xe-0/0/1 up up
xe-0/0/1.0 up up eth-switch

```

### show interfaces (QFabric System)

```

user@switch> show interfaces node1:xe-0/0/0
Physical interface: node1:xe-0/0/0, Enabled, Physical link is Down
Interface index: 129, SNMP ifIndex: 2884086
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None,
Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled

```

Interface flags: Internal: 0x4000  
CoS queues : 8 supported, 8 maximum usable queues  
Current address: 02:00:09:03:00:00, Hardware address: 02:00:09:03:00:00  
Last flapped : Never  
Input rate : 0 bps (0 pps)  
Output rate : 0 bps (0 pps)

## CHAPTER 134

# LAGs and LACP Operational Commands

- [show lacp interfaces](#)
- [show lacp statistics interfaces \(View\)](#)

## show lacp interfaces

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show lacp interfaces</code><br><code>&lt;interface-name&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.<br>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.<br>Command introduced in Junos OS Release 14.2R3                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Display Link Aggregation Control Protocol (LACP) information about the specified aggregated Ethernet or Gigabit Ethernet interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <code>none</code> —Display LACP information for all interfaces.<br><br><code>interface-name</code> —(Optional) Display LACP information for the specified interface: <ul style="list-style-type: none"><li>• Aggregated Ethernet—<code>aex</code></li><li>• Gigabit Ethernet—<code>ge-fpc/pic/port</code></li><li>• 10-Gigabit Ethernet—<code>xe-fpc/pic/port</code></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch</i></li><li>• <i>Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch</i></li><li>• <a href="#">Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2876</a></li><li>• <a href="#">Configuring Aggregated Ethernet Links (CLI Procedure)</a></li><li>• <a href="#">Configuring Link Aggregation on page 2869</a></li><li>• <a href="#">Configuring Aggregated Ethernet LACP (CLI Procedure)</a></li><li>• <a href="#">Configuring Aggregated Ethernet LACP on page 2868</a></li><li>• <a href="#">Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure)</a></li><li>• <a href="#">Understanding Aggregated Ethernet Interfaces and LACP</a></li><li>• <a href="#">Understanding Aggregated Ethernet Interfaces and LACP on page 2865</a></li><li>• <a href="#">Junos OS Interfaces Fundamentals Configuration Guide</a></li></ul> |
| <b>List of Sample Output</b>    | <a href="#">show lacp interfaces (EX Series Switches) on page 3206</a><br><a href="#">show lacp interfaces (QFX Series) on page 3207</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

**Output Fields** Table 284 lists the output fields for the **show lacp interfaces** command. Output fields are listed in the approximate order in which they appear.

**Table 284: show lacp interfaces Output Fields**

| Field Name           | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Aggregated interface | Aggregated Ethernet interface name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| LACP State           | <p>LACP state information for each aggregated Ethernet interface:</p> <ul style="list-style-type: none"> <li>For a child interface configured with the <b>force-up</b> statement, LACP state displays <b>FUP</b> along with the interface name.</li> <li><b>Role</b>—Role played by the interface. It can be one of the following: <ul style="list-style-type: none"> <li><b>Actor</b>—Local device participating in the LACP negotiation.</li> <li><b>Partner</b>—Remote device participating in the LACP negotiation.</li> </ul> </li> <li><b>Exp</b>—Expired state. <b>Yes</b> indicates that the actor or partner is in an expired state. <b>No</b> indicates that the actor or partner is not in an expired state.</li> <li><b>Def</b>—Default. <b>Yes</b> indicates that the actor's receive machine is using the default operational partner information, which is administratively configured for the partner. <b>No</b> indicates that the operational partner information in use has been received in an LACP PDU.</li> <li><b>Dist</b>—Distribution of outgoing frames. <b>No</b> indicates that the distribution of outgoing frames on the link is currently disabled and is not expected to be enabled. Otherwise, the value is <b>Yes</b>.</li> <li><b>Col</b>—Collection of incoming frames. <b>Yes</b> indicates that the collection of incoming frames on the link is currently enabled and is not expected to be disabled. Otherwise, the value is <b>No</b>.</li> <li><b>Syn</b>—Synchronization. If the value is <b>Yes</b>, the link is considered to be synchronized. The link has been allocated to the correct link aggregation group, the group has been associated with a compatible aggregator, and the identity of the link aggregation group is consistent with the system ID and operational key information transmitted. If the value is <b>No</b>, the link is not synchronized. The link is currently not in the right aggregation.</li> <li><b>Aggr</b>—Ability of the aggregation port to aggregate (<b>Yes</b>) or to operate only as an individual link (<b>No</b>).</li> <li><b>Timeout</b>—LACP timeout preference. Periodic transmissions of LACP PDUs occur at either a slow or a fast transmission rate, depending upon the expressed LACP timeout preference (<b>Long Timeout</b> or <b>Short Timeout</b>).</li> <li><b>Activity</b>—Actor's or partner's port activity. <b>Passive</b> indicates the port's preference for not transmitting LAC PDUs unless its partner's control value is <b>Active</b>. <b>Active</b> indicates the port's preference to participate in the protocol regardless of the partner's control value.</li> </ul> |

Table 284: show lacp interfaces Output Fields (*continued*)

| Field Name    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| LACP Protocol | <p>LACP protocol information for each aggregated interface:</p> <ul style="list-style-type: none"> <li>Link state (active or standby) indicated in parentheses next to the interface when link protection is configured.</li> <li><b>Receive State</b>—One of the following values: <ul style="list-style-type: none"> <li><b>Current</b>—The state machine receives an LACP PDU and enters the <b>Current</b> state.</li> <li><b>Defaulted</b>—If no LACP PDU is received before the timer for the <b>Current</b> state expires a second time, the state machine enters the <b>Defaulted</b> state.</li> <li><b>Expired</b>—If no LACP PDU is received before the timer for the <b>Current</b> state expires once, the state machine enters the <b>Expired</b> state.</li> <li><b>Initialize</b>—When the physical connectivity of a link changes or a Begin event occurs, the state machine enters the <b>Initialize</b> state.</li> <li><b>LACP Disabled</b>—If the port is operating in half duplex, the operation of LACP is disabled on the port, forcing the state to <b>LACP Disabled</b>. This state is similar to the <b>Defaulted</b> state, except that the port is forced to operate as an individual port.</li> <li><b>Port Disabled</b>—If the port becomes inoperable and a Begin event has not occurred, the state machine enters the <b>Port Disabled</b> state.</li> </ul> </li> <li><b>Transmit State</b>—Transmit state of the state machine. The transmit state is one of the following values: <ul style="list-style-type: none"> <li><b>Fast periodic</b>—Periodic transmissions are enabled at a fast transmission rate.</li> <li><b>No periodic</b>—Periodic transmissions are disabled.</li> <li><b>Periodic timer</b>—Transitory state entered when the periodic timer expires.</li> <li><b>Slow periodic</b>—Periodic transmissions are enabled at a slow transmission rate.</li> </ul> </li> <li><b>Mux State</b>—State of the multiplexer state machine for the aggregation port. The state is one of the following values: <ul style="list-style-type: none"> <li><b>Attached</b>—The multiplexer state machine initiates the process of attaching the port to the selected aggregator.</li> <li><b>Collecting—Yes</b> indicates that the receive function of this link is enabled with respect to its participation in an aggregation. Received frames are passed to the aggregator for collection. <b>No</b> indicates the receive function of this link is not enabled.</li> <li><b>Collecting distributing</b>—Collecting and distributing states are merged together to form a combined state (coupled control). Because independent control is not possible, the coupled control state machine does not wait for the partner to signal that collection has started before enabling both collection and distribution.</li> <li><b>Detached</b>—Process of detaching the port from the aggregator is in progress.</li> <li><b>Distributing—Yes</b> indicates that the transmit function of this link is enabled with respect to its participation in an aggregation. Frames can be passed down from the aggregator's distribution function for transmission. <b>No</b> indicates the transmit function of this link is not enabled.</li> <li><b>Waiting</b>—The multiplexer state machine is in a holding process, awaiting an outcome.</li> </ul> </li> </ul> |

## Sample Output

### show lacp interfaces (EX Series Switches)

```

user@switch> show lacp interfaces ae5
Aggregated interface: ae5
 LACP state: Role Exp Def Dist Co1 Syn Aggr Timeout Activity
 xe-2/0/7 Actor No No Yes Yes Yes Yes Fast Active
 xe-2/0/7 Partner No No Yes Yes Yes Yes Fast Passive

```

|          |         |    |    |    |     |     |     |      |         |
|----------|---------|----|----|----|-----|-----|-----|------|---------|
| xe-4/0/7 | Actor   | No | No | No | No  | No  | Yes | Fast | Active  |
| xe-4/0/7 | Partner | No | No | No | Yes | Yes | Yes | Fast | Passive |

| LACP protocol:     | Receive State | Transmit State | Mux State               |
|--------------------|---------------|----------------|-------------------------|
| xe-2/0/7(Active)   | Current       | Fast periodic  | Collecting distributing |
| xe-34/0/7(Standby) | Current       | Fast periodic  | Waiting                 |

### show lacp interfaces (QFX Series)

```

user@switch> show lacp interfaces nodegroup1:ae0 extensive
Aggregated interface: nodegroup1:ae0
LACP state: Role Exp Def Dist Col Syn Aggr Timeout Activity

node1:xe-0/0/1FUP Actor No Yes No No No No Yes Fast
Active
node1xe-0/0/1FUP Partner No Yes No No No No Yes Fast
Passive
node2:xe-0/0/2 Actor No Yes No No No No Yes Fast
Active
node2:xe-0/0/2 Partner No Yes No No No No Yes Fast
Passive

```

|              | LACP protocol:           | Receive State | Transmit State | Mux State  |
|--------------|--------------------------|---------------|----------------|------------|
|              | node1:xe-0/0/1FUP        | Current       | Fast periodic  | Collecting |
| distributing |                          |               |                |            |
|              | node2:xe-0/0/2           | Current       | Fast periodic  | Collecting |
| distributing |                          |               |                |            |
|              | node1:xe-0/0/1 (active)  | Current       | Fast periodic  | Collecting |
| distributing |                          |               |                |            |
|              | node2:xe-0/0/2 (standby) | Current       | Fast periodic  | WAITING    |



## show lacp statistics interfaces (View)

|                                 |                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show lacp statistics interfaces <i>interface-name</i></code>                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Command modified in Junos OS Release 10.2.                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Display Link Aggregation Control Protocol (LACP) statistics about the specified aggregated Ethernet interface or redundant Ethernet interface. If you do not specify an interface name, LACP statistics for all interfaces are displayed.                                                                                                       |
| <b>Options</b>                  | <i>interface-name</i> —(Optional) Name of an interface.                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Verifying LACP on Redundant Ethernet Interfaces</a></li> <li>• <a href="#">Verifying the Status of a LAG Interface on page 2885</a></li> <li>• <a href="#">Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets on page 2885</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show lacp statistics interfaces on page 3209</a>                                                                                                                                                                                                                                                                                    |
| <b>Output Fields</b>            | Table 285 lists the output fields for the <code>show lacp statistics interfaces</code> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                         |

Table 285: show lacp statistics interfaces Output Fields

| Field Name           | Field Description                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Aggregated interface | Aggregated interface value.                                                                                                                                                                                                                                                                                                                                                                                  |
| LACP Statistics      | <p>LACP statistics provide the following information:</p> <ul style="list-style-type: none"> <li>• <b>LACP Rx</b>—LACP received counter that increments for each normal hello.</li> <li>• <b>LACP Tx</b>—Number of LACP transmit packet errors logged.</li> <li>• <b>Unknown Rx</b>—Number of unrecognized packet errors logged.</li> <li>• <b>Illegal Rx</b>—Number of invalid packets received.</li> </ul> |

## Sample Output

### show lacp statistics interfaces

```

user@host> show lacp statistics interfaces ae0
Aggregated interface: ae0
 LACP Statistics: LACP Rx LACP Tx Unknown Rx Illegal Rx
 ge-2/0/0 1352 2035 0 0
 ge-2/0/1 1352 2056 0 0
 ge-2/2/0 1352 2045 0 0
 ge-2/2/1 1352 2043 0 0

```



## CHAPTER 135

# Redundant Trunk Group Operational Command

- `show redundant-trunk-group`

## show redundant-trunk-group

|                                 |                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show redundant-trunk-group &lt;group-name group-name&gt;</b>                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.                                                                                                                                                                          |
| <b>Description</b>              | Display information about redundant trunk groups.                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <b>group-name group-name</b> —Display information about the specified redundant trunk group.                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Redundant Trunk Links for Faster Recovery</a></li> <li>• <a href="#">Example: Configuring Redundant Trunk Links for Faster Recovery on page 2903</a></li> <li>• <a href="#">Understanding Redundant Trunk Links on page 2901</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show redundant-trunk-group group-name Group1 on page 3212</a>                                                                                                                                                                                                                                             |
| <b>Output Fields</b>            | <a href="#">Table 286</a> lists the output fields for the <b>show redundant-trunk-group</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                          |

**Table 286: show redundant-trunk-group Output Fields**

| Field Name               | Field Description                                                                                                                                                                                                                                                                        |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Group name</b>        | Name of the redundant trunk port group.                                                                                                                                                                                                                                                  |
| <b>Interface</b>         | Name of an interface belonging to the trunk port group.                                                                                                                                                                                                                                  |
| <b>State</b>             | Operating state of the interface. <ul style="list-style-type: none"> <li>• <b>Up</b> denotes the interface is up.</li> <li>• <b>Down</b> denotes the interface is down.</li> <li>• <b>Pri</b> denotes a primary interface.</li> <li>• <b>Act</b> denotes an active interface.</li> </ul> |
| <b>Time of last flap</b> | Date and time at which the advertised link became unavailable, and then, available again.                                                                                                                                                                                                |
| <b>Flap count</b>        | Total number of flaps since the last switch reboot.                                                                                                                                                                                                                                      |

## Sample Output

### show redundant-trunk-group group-name Group1

```
user@switch> show redundant-trunk-group group-name Group1
```

| Group name | Interface | State | Time of last flap | Flap Count |
|------------|-----------|-------|-------------------|------------|
|------------|-----------|-------|-------------------|------------|

|        |             |            |       |   |
|--------|-------------|------------|-------|---|
| Group1 | ge-0/0/45.0 | UP/Pri/Act | Never | 0 |
|        | ge-0/0/47.0 | UP         | Never | 0 |



## CHAPTER 136

# Resilient Hashing Operational Command

- `show forwarding-options enhanced-hash-key`

## show forwarding-options enhanced-hash-key

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show forwarding-options enhanced-hash-key</b>                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 13.2X51-D15 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.</p> <p><b>Fabric Load Balancing Options</b> output fields introduced in Junos OS Release 14.1X53-D10.</p>                                                                                                                                                                        |
| <b>Description</b>              | <p>Display information about which packet fields are used by the hashing algorithm to make hashing decisions.</p> <p>You can configure the fields that are inspected by the hashing algorithm to make hashing decisions for traffic entering a LAG bundle using the <b>forwarding-options enhanced-hash-key</b> statement.</p>                                                                                                                   |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure) on page 2923</a></li> <li>• <a href="#">Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 2911</a></li> <li>• <a href="#">enhanced-hash-key on page 3066</a></li> </ul>                                                                         |
| <b>List of Sample Output</b>    | <p><a href="#">show forwarding-options enhanced-hash-key (Layer 2 Payload Hash Mode) on page 3218</a></p> <p><a href="#">show forwarding-options enhanced-hash-key (Layer 2 Header Hash Mode) on page 3218</a></p> <p><a href="#">show forwarding-options enhanced-hash-key (Fabric Load Balancing Options) on page 3219</a></p> <p><a href="#">show forwarding-options enhanced-hash-key (QFX10002 and QFX 10008 Switches) on page 3219</a></p> |
| <b>Output Fields</b>            | <p>Table 287 lists the output fields for the <b>show forwarding-options enhanced-hash-key</b> command. Output fields are listed in the approximate order in which they first appear. Output fields vary by platform.</p>                                                                                                                                                                                                                         |

**Table 287: show forwarding-options enhanced-hash-key Output Fields**

| Field Name                 | Field Description                                                                                      |
|----------------------------|--------------------------------------------------------------------------------------------------------|
| <b>Hash-Mode</b>           | Current hash mode: Layer 2 header or Layer 2 payload.                                                  |
| <b>Protocol</b>            | Indicates whether the Protocol field is or is not used by the hashing algorithm: Yes or No.            |
| <b>Destination L4 Port</b> | Indicates whether the Destination L4 Port field is or is not used by the hashing algorithm: Yes or No. |
| <b>Source L4 Port</b>      | Indicates whether the Source L4 Port field is or is not used by the hashing algorithm: Yes or No.      |



Table 287: show forwarding-options enhanced-hash-key Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                |
|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Destination IPv4 Addr</b>   | Indicates whether the Destination IPv4 Addr field is or is not used by the hashing algorithm: Yes or No.                                                                                                                                                                         |
| <b>Source IPv4 Addr</b>        | Indicates whether the Source IPv4 Addr field is or is not used by the hashing algorithm: Yes or No.                                                                                                                                                                              |
| <b>Vlan id</b>                 | Indicates whether the Vlan ID field is or is not used by the hashing algorithm: Yes or No.                                                                                                                                                                                       |
| <b>Inner-Vlan ID</b>           | Indicates whether the inner Vlan field is or is not used by the hashing algorithm: Yes or No.                                                                                                                                                                                    |
| <b>Next Hdr</b>                | Indicates whether the Next Hdr field is or is not used by the hashing algorithm: Yes or No.                                                                                                                                                                                      |
| <b>Destination IPv6 Addr</b>   | Indicates whether the Destination IPv6 Addr field is or is not used by the hashing algorithm: Yes or No.                                                                                                                                                                         |
| <b>Source IPv6 Addr</b>        | Indicates whether the Source IPv6 Addr field is or is not used by the hashing algorithm: Yes or No.                                                                                                                                                                              |
| <b>Ether Type</b>              | Indicates whether the Ether Type field is or is not used by the hashing algorithm: Yes or No.                                                                                                                                                                                    |
| <b>Destination MAC Address</b> | Indicates whether the Destination MAC Address field is or is not used by the hashing algorithm: Yes or No.                                                                                                                                                                       |
| <b>Source MAC Address</b>      | Indicates whether the Source MAC Address field is or is not used by the hashing algorithm: Yes or No.                                                                                                                                                                            |
| <b>Load Balancing Method</b>   | Indicates the load balancing method for adaptive load balancing (ALB): flowlet or per-packet.<br><br>The load balancing method is flowlet by default, and can be configured using the <b>fabric-load-balance</b> statement.                                                      |
| <b>Fabric Link Scale</b>       | Indicates the fabric link scale, in mbps.                                                                                                                                                                                                                                        |
| <b>Inactivity Interval</b>     | Indicates the fabric load balance inactivity interval, in microseconds (us).<br><br>The inactivity interval is 16 microseconds by default, and can be configured using the <b>inactivity-interval</b> statement.                                                                 |
| <b>Hash Region Size/Trunk</b>  | Indicates the hash region size, in buckets per fabric trunk.                                                                                                                                                                                                                     |
| <b>Seed</b>                    | A hash seed value, between 0 and 4294967295. If a hash-seed value is not configured it is automatically assigned on the QFX10002 and QFX10008 switches. A hash-seed prevents traffic polarization to same links on the next hop QFX switch when two are connected with LAG/ECMP. |

**Table 287: show forwarding-options enhanced-hash-key Output Fields (*continued*)**

| Field Name          | Field Description                                                                                                                         |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Key</b>          | Indicates whether the GRE key field is or is not used by the hashing algorithm: Yes or No.                                                |
| <b>Protocol</b>     | Indicates if a Generic Router Encapsulation (GRE) endpoint over routes was dynamically learned by a routing protocol such as RIP or OSPF. |
| <b>MPLS Enabled</b> | Indicates if MPLS is enabled under L2 switching.                                                                                          |
| <b>VXLAN VNID</b>   | A 24-bit virtual network identifier (VNID) that uniquely identifies the Virtual Extensible Local Area Networks (VXLAN) segment.           |

## Sample Output

### show forwarding-options enhanced-hash-key (Layer 2 Payload Hash Mode)

```

user@switch> show forwarding-options enhanced-hash-key
Slot 0

Current Hash Settings

Hash-Mode :layer2-payload

inet Hash settings-

inet packet fields
 Protocol : Yes
 Destination L4 Port : Yes
 Source L4 Port : Yes
 Destination IPv4 Addr : Yes
 Source IPv4 Addr : Yes
 Vlan id : No

inet6 Hash settings-

inet6 packet fields
 Next Hdr : Yes
 Destination L4 Port : Yes
 Source L4 Port : Yes
 Destination IPv6 Addr : Yes
 Source IPv6 Addr : Yes
 Vlan id : No

```

### show forwarding-options enhanced-hash-key (Layer 2 Header Hash Mode)

```

user@switch> show forwarding-options enhanced-hash-key
Slot 0

```

```

Current Hash Settings

```

```
Hash-Mode : layer2-header
```

```
layer2 Hash settings-
```

```

```

```
layer2 packet fields
```

```

Ether Type : Yes
Destination MAC Address : Yes
Source MAC Address : Yes
VLAN ID : No
```

### show forwarding-options enhanced-hash-key (Fabric Load Balancing Options)

```
user@switch> show forwarding-options enhanced-hash-key
<some output removed for brevity>
```

```
Fabric Load Balancing Options
```

```

```

```

Load Balancing Method : Flowlet
Fabric Link Scale : 40960 (mbps)
Inactivity Interval : 16 (us)
Hash Region Size/Trunk : 1024 (buckets)
```

### show forwarding-options enhanced-hash-key (QFX10002 and QFX 10008 Switches)

```
user@switch> show forwarding-options enhanced-hash-key
Slot 0
```

```

Seed value for Hash function 0: 3626023417
Seed value for Hash function 1: 3626023417
Seed value for Hash function 2: 3626023417
Seed value for Hash function 3: 3626023417
```

```
Inet settings:
```

```

```

```

IPV4 dest address: Yes
IPV4 source address: Yes
L4 Dest Port: Yes
L4 Source Port: Yes
```

```
Inet6 settings:
```

```

```

```

IPV6 dest address: Yes
IPV6 source address: Yes
L4 Dest Port: Yes
L4 Source Port: Yes
```

```
L2 settings:
```

```

```

```

Dest Mac address: No
Source Mac address: No
Vlan Id: Yes
Inner-vlan Id: No
Incoming port: Yes
```

```
GRE settings:
```

```

```

```

Key: No
Protocol: No
```

```
MPLS settings:
```

```

```

```
MPLS Enabled: Yes
```

```
VXLAN settings:
```

-----  
VXLAN VNID: No

## CHAPTER 137

# Uplink Failure Detection Operational Command

- `show uplink-failure-detection`

## show uplink-failure-detection

|                                 |                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show uplink-failure-detection</code><br><code>&lt;group group-name&gt;</code>                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for EX Series switches.                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Display information about the uplink-failure-detection group, the member interfaces, and their status.                                                                                                                                                                                                                 |
| <b>Options</b>                  | <b>none</b> —Display information about all groups configured for uplink failure detection.<br><b>group group-name</b> —(Optional) Display information about the specified group only.                                                                                                                                  |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Overview of Uplink Failure Detection on page 2929</a></li> <li>• <a href="#">Configuring Interfaces for Uplink Failure Detection on page 2931</a></li> <li>• <a href="#">Example: Configuring Interfaces for Uplink Failure Detection on page 2932</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show uplink-failure-detection on page 3222</a>                                                                                                                                                                                                                                                             |
| <b>Output Fields</b>            | Table 288 lists the output fields for the <b>show uplink-failure-detection</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                        |

**Table 288: show uplink-failure-detection Output Fields**

| Field Name     | Field Description                                                                                                                                                                                                             |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Group          | Name of the group.                                                                                                                                                                                                            |
| Uplink         | The uplink interface or interfaces configured as link-to-monitor.<br><b>NOTE:</b> The asterisk (*) indicates that the link is up.                                                                                             |
| Downlink       | The downlink interface or interfaces configured as link-to-disable.<br><b>NOTE:</b> The asterisk (*) indicates that the link is up.                                                                                           |
| Failure Action | Status of uplink failure detection: <ul style="list-style-type: none"> <li>• Active—The switch has detected an uplink failure and has brought the downlink down.</li> <li>• Inactive—The uplink or uplinks are up.</li> </ul> |

## Sample Output

### show uplink-failure-detection

```
user@switch> show uplink-failure-detection
```

Group : group1  
Uplink : ge-0/0/0\*  
Downlink : ge-0/0/1\*  
Failure Action : Inactive

Group : group2  
Uplink : ge-0/0/3.0  
Downlink : ge-0/0/4.0  
Failure Action : Active





# Routing Options Feature Guide for QFX10000 Switches



# Configuring Routing Options

- [Overview of Routing Options on page 3227](#)
- [Configuring Static Routing on page 3228](#)
- [Configuring Per-Packet Load Balancing on page 3228](#)
- [Examples: Configuring Per-Packet Load Balancing on page 3230](#)
- [Using a QFX10000 Switch with 6PE on page 3231](#)
- [Understanding Virtual Router Routing Instances on page 3232](#)
- [Configuring Virtual Router Routing Instances on page 3232](#)
- [Understanding Distributed Periodic Packet Management on page 3233](#)
- [Configuring Distributed Periodic Packet Management on page 3234](#)
- [Understanding Bidirectional Forwarding Detection \(BFD\) on page 3236](#)
- [Examples: Configuring BFD for Static Routes on page 3237](#)
- [Example: Configuring BFD Authentication for Static Routes on page 3253](#)
- [Monitoring Routing Information on page ?](#)
- [Verifying That Virtual Router Routing Instances Are Working on page ?](#)
- [Troubleshooting Virtual Routing Instances on page 3263](#)

## Overview of Routing Options

---

In addition to dynamic routing protocols, you can configure static routing on QFX Series switches. You can also configure a variety of protocol-independent routing properties, such as

- Per-packet load balancing (equal cost multipath routing)
- Autonomous system numbers
- Autonomous system confederation members
- Router identifiers
- Routing table groups
- Multicast scoping

### Related Documentation

- [Understanding Distributed Periodic Packet Management on page 3233](#)

- [Understanding Virtual Router Routing Instances on page 3232](#)

## Configuring Static Routing

---

Static routes are routes that are manually configured and entered into the routing table.

The switch uses static routes:

- When the switch does not have a route to a destination that has a better (lower) *preference* value. The preference is an arbitrary value in the range from 0 through 255 that the software uses to rank routes received from different protocols, interfaces, or remote systems. The routing protocol process generally determines the active route by selecting the route with the lowest preference value. In the given range, **0** is the lowest and **255** is the highest.
- When the switch cannot determine the route to a destination.
- When the switch is forwarding unroutable packets.

To configure basic static route options using the CLI:

- To configure the switch's default gateway:

```
[edit]
user@switch# set routing-options static route 0.0.0.0/0 next-hop 10.0.1.1
```

- To configure a static route and specify the next address to be used when routing traffic to the static route:

```
[edit]
user@switch# set routing-options static route 20.0.0.0/24 next-hop 10.0.0.2.1
```

- To always keep the static route in the forwarding table:

```
[edit]
user@switch# set routing-options static route 20.0.0.0/24 retain
```

- To prevent the static route from being readvertised:

```
[edit]
user@switch# set routing-options static route 20.0.0.0/24 no-readvertise
```

- To remove inactive routes from the forwarding table:

```
[edit]
user@switch# set routing-options static route 20.0.0.0/24 active
```

Related  
Documentation

- [Monitoring Routing Information on page ?](#)

## Configuring Per-Packet Load Balancing

---

By default, when there are multiple equal-cost paths to the same destination for the active route, Junos OS chooses one of the next-hop addresses to install into the forwarding table in a random fashion. Whenever the set of next hops for a destination changes in any way, the next-hop address is chosen again, also in a random fashion.

You can configure Junos OS so that, for the active route, all next-hop addresses for a destination are installed in the forwarding table. This is called per-packet load balancing. You can use this feature to spread traffic across multiple paths.

On a QFX3500 standalone switch, with static routing configured, whenever a route pointing to an ECMP next-hop changes to a new ECMP next-hop with a different member list but contains the exact member count as before, the location of the retained members in the new member list is the same as in the old member list.

For example, if you have the following configuration on the switch:

```
set routing-options static route 0.0.0.0/0 next-hop 11.8.12.2
set routing-options static route 0.0.0.0/0 next-hop 11.9.12.2
set routing-options static route 0.0.0.0/0 next-hop 11.10.12.2
set routing-options static route 0.0.0.0/0 next-hop 11.11.12.2
set routing-options static route 0.0.0.0/0 next-hop 11.12.12.2
set routing-options static route 0.0.0.0/0 next-hop 11.13.12.2
set routing-options static route 0.0.0.0/0 next-hop 11.14.12.2
set routing-options static route 0.0.0.0/0 next-hop 11.15.12.2
```

and want to change the first and eighth route to point to another location, you can issue the following commands:

```
delete routing-options static route 0.0.0.0/0 next-hop 11.8.12.2
delete routing-options static route 0.0.0.0/0 next-hop 11.15.12.2
set routing-options static route 0.0.0.0/0 next-hop 11.16.12.2
set routing-options static route 0.0.0.0/0 next-hop 11.17.12.2
```

This configuration does not affect the second next-hop through the seventh next-hop.

When per-packet load balancing is configured, traffic is divided into individual flows (up to a maximum of 16). Packets for an individual flow are sent out a single interface. To determine flows, the switch examines each of the following packet fields:

- Source IP address
- Destination IP address
- Protocol
- Source port number
- Destination port number
- Source interface index
- Type of service (ToS)

The switch recognizes packets in which all of these parameters are identical and ensures that these packets are sent out through the same interface. This prevents problems that might otherwise occur with packets arriving at their destination out of their original sequence.



**NOTE:** Load balancing is not supported on management interfaces.

The following steps show how to configure per-packet load balancing:

1. Define a load-balancing routing policy by including one or more **policy-statement** statements at the **[edit policy-options]** hierarchy level, defining an action of **load-balance per-packet**:

```
policy-statement policy-name {
 from {
 match-conditions;
 route-filter destination-prefix match-type <actions>;
 prefix-list name;
 }
 then {
 load-balance per-packet;
 }
}
```

2. Apply the policy to routes exported from the routing table to the forwarding table. To do this, include the **forwarding-table** and **export** statements:

```
forwarding-table {
 export policy-name;
}
```

When you enable per-packet load balancing on a QFabric system, packets might be switched across the fabric even though there is a local port with a same-cost route to the destination. For example, if per-packet load balancing is enabled and a packet arrives at network Node device A, it might be switched to network Node device B and forwarded from there even if there is a same-cost route through a port on Node device A to the destination. In this case, traffic transits the fabric needlessly. You can configure a QFabric system to choose a locally switched route if one is available. To enable this feature, include the **ecmp-do-local-lookup** statement at the **[edit forwarding-options]** hierarchy level.

#### Related Documentation

- [Examples: Configuring Per-Packet Load Balancing on page 3230](#)
- [Understanding the Use of Resilient Hashing to Minimize Flow Remapping in Trunk/ECMP Groups on page 2917](#)

---

## Examples: Configuring Per-Packet Load Balancing

Perform per-packet load balancing for all routes:

```
[edit]
policy-options {
 policy-statement load-balancing-policy {
 then {
 load-balance per-packet;
 }
 }
}
routing-options {
 forwarding-table {
 export load-balancing-policy;
 }
}
```

Perform per-packet load balancing for a limited set of routes:

```
[edit]
policy-options {
 policy-statement load-balancing-policy {
 from {
 route-filter 192.168.10/24 orlonger;
 route-filter 9.114/16 orlonger;
 }
 then {
 load-balance per-packet;
 }
 }
}
routing-options {
 forwarding-table {
 export load-balancing-policy;
 }
}
```

**Related  
Documentation**

- [Configuring Per-Packet Load Balancing on page 3228](#)
- [Understanding the Use of Resilient Hashing to Minimize Flow Remapping in Trunk/ECMP Groups on page 2917](#)

## Using a QFX10000 Switch with 6PE

Using 6PE allows service providers to provide new revenue-generating services to customers by offering native IPv6 connectivity over an MPLS core network without having to configure or upgrade the core routers or change the existing IPv4 and MPLS services. To use 6PE you configure dual-stack (IPv4 and IPv6) provider edge (PE) devices to tunnel IPv6 traffic over MPLS.

When you use a QFX10000 switch as a PE device in a 6PE topology, you must enable **family inet6** on the core-facing interfaces that will be used for the IPv6 traffic. If you do not do so, these interfaces drop the IPv6 traffic.

**Related  
Documentation**

- [MPLS Feature Guide for QFX10000 Switches](#)

## Understanding Virtual Router Routing Instances

---

Virtual router routing instances allow administrators to divide a QFX Series switch into multiple independent virtual routers, each with its own routing table. Virtual router routing instances enable you to isolate traffic without using multiple devices to segment your network. You can create routing instances for unicast routing protocols and PIM sparse mode.

Each virtual router routing instance consists of sets of the following:

- Routing tables
- Interfaces that belong to these routing tables
- Routing protocol configurations
- Routing option configurations

You can use virtual router routing instances to isolate customer traffic on a network and to bind customer-specific routing instances to customer-owned interfaces. Each interface can belong to only one routing instance. QFX 3500 and QFX3600 switches and QFabric systems support as many as 256 virtual router routing instances. QFX 5100 switches support as many as 512 virtual router routing instances.

### Related Documentation

- [Configuring Virtual Router Routing Instances on page 3232](#)

## Configuring Virtual Router Routing Instances

---

Use virtual router routing instances to divide a QFX Series switch into multiple independent virtual routers, each with its own routing table. Virtual router routing instances enable you to isolate traffic without using multiple devices to segment your network. You can create routing instances for unicast routing protocols and PIM sparse mode.

To configure virtual router routing instances:

1. Create a routing instance:

```
[edit routing-instances]
user@switch# set routing-instance-name instance-type virtual-router
```



**NOTE:** The default routing instance, *master*, refers to the main *inet.0* routing table. The *master* routing instance is reserved and cannot be specified as a routing instance.

2. Bind each routing instance to the corresponding interfaces:

```
[edit routing-instances]
user@switch# set routing-instance-name interface
device-name:type-fpc/pic/port.logical-unit-number
```



**NOTE:**

- You must bind routing instances to interfaces from the Node devices assigned to the network Node group only. If you try to bind routing instances to interfaces from the Node devices assigned to server Node groups, the configuration does not commit.
- You can bind an interface to one routing instance only.

3. Create each of the logical interfaces bound to each routing instance:

```
[edit interfaces]
user@switch# set device-name: type-fpc/pic/port unit logical-unit-number family inet address
ip-address
```



**NOTE:** Do not create a logical interface using the family ethernet-switching option in this step. Binding an interface using the family ethernet-switching option to a routing instance can cause the interface to shut down.

4. (Optional) Configure routing protocols for the routing instance at the **[edit routing-instances routing-instance-name protocols]** hierarchy level.
5. (Optional) Configure routing options for the routing instance at the **[edit routing-instances routing-instance-name routing-options]** hierarchy level.

**Related  
Documentation**

- [Understanding Virtual Router Routing Instances on page 3232](#)
- [Understanding Node Groups](#)
- [Verifying That Virtual Router Routing Instances Are Working on page ?](#)

## Understanding Distributed Periodic Packet Management

Periodic packet management (PPM) is responsible for processing a variety of time-sensitive periodic tasks for particular processes so that other processes can more optimally direct their resources. PPM is responsible for the periodic transmission of packets on behalf of its various client processes, which include the processes that control the Link Aggregation Control Protocol (LACP) and Bidirectional Forwarding Detection (BFD) protocol, and also for receiving packets on behalf of these client processes. PPM also gathers some statistics and sends process-specific packets. PPM cannot be disabled and is always running on any operational switch.

The responsibility for PPM processing on the switch is distributed between the Routing Engine and the access interfaces for all protocols that use PPM by default. This distributed model provides a faster response time for protocols that use PPM than the response time provided by the nondistributed model.

If distributed PPM is disabled, the PPM process runs on the Routing Engine only.

You can disable distributed PPM for all protocols that use PPM. You can also disable distributed PPM for LACP packets only.



**BEST PRACTICE:** We generally recommend that you disable distributed PPM only if Juniper Networks Customer Service advises you to do so. You should disable distributed PPM only if you have a compelling reason to disable it.

**Related  
Documentation**

- [Configuring Distributed Periodic Packet Management on page 3234](#)

---

## Configuring Distributed Periodic Packet Management

Periodic packet management (PPM) is responsible for processing a variety of time-sensitive periodic tasks so that other processes can more optimally direct their resources.

The responsibility for PPM processing on the switch is distributed between the Routing Engine and the access interfaces for all protocols that use PPM by default. This distributed model provides a faster response time for protocols that use PPM than the response time provided by the nondistributed model.

If distributed PPM is disabled, the PPM process runs on the Routing Engine only.

You can disable distributed PPM for all protocols that use PPM. You can also disable distributed PPM for Link Aggregation Control Protocol (LACP) packets only.



**BEST PRACTICE:** We generally recommend that you disable distributed PPM only if Juniper Networks Customer Service advises you to do so. You should disable distributed PPM only if you have a compelling reason to disable it.

This topic describes:

- [Disabling or Enabling Distributed Periodic Packet Management Globally on page 3234](#)
- [Disabling or Enabling Distributed Periodic Packet Management for LACP Packets on page 3235](#)

### Disabling or Enabling Distributed Periodic Packet Management Globally

Distributed PPM is enabled by default. Disable distributed PPM if you need to move all PPM processing to the Routing Engine. Enable distributed PPM if it was previously disabled and you need to run distributed PPM.

To disable distributed PPM:

```
[edit routing-options]
user@switch# set ppm no-delegate-processing
```

To enable distributed PPM if it was previously disabled:

```
[edit routing-options]
```

```
user@switch# delete ppm no-delegate-processing
```

## Disabling or Enabling Distributed Periodic Packet Management for LACP Packets

Distributed PPM is enabled by default. Disable distributed PPM for only LACP packets if you need to move all PPM processing for LACP packets to the Routing Engine.

To disable distributed PPM for LACP packets:

```
[edit protocols]
user@switch# set lacp ppm centralized
```

To enable distributed PPM for LACP packets if it was previously disabled:

```
[edit protocols]
user@switch# delete lacp ppm centralized
```

### Related Documentation

- [Understanding Distributed Periodic Packet Management on page 3233](#)
- [Understanding Aggregated Ethernet Interfaces and LACP on page 2865](#)

## Understanding Bidirectional Forwarding Detection (BFD)

---

The Bidirectional Forwarding Detection (BFD) protocol is a simple mechanism that detects failures in a network and works in a wide variety of network environments and topologies. In BFD operation, switches exchange BFD hello packets at a specified interval and detect a neighbor failure if they do not receive a reply after a specified interval. The BFD failure detection timers support shorter time limits than the static route failure detection mechanisms, so they can provide faster detection of failures.

To configure faster failure detection, use lower BFD timer values. The timers can automatically adapt to a higher value if an adjacency fails, and they also adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. In this case, a back-off algorithm increases the receive interval by two if the local BFD instance is the reason for the session flap and increases the transmission interval by two if the remote BFD instance is the reason for the session flap.

You can use the **clear bfd adaptation** command to return BFD interval timers to their configured values. This command is hitless, meaning that it does not affect traffic flow.



NOTE: QFX10000 series switches support BFD timer values of less than 1 second.



NOTE: BFD is not supported with ISIS for IPV6 on QFX10000 series switches.



NOTE: EX3300 supports BFD over static routes only.

### Related Documentation

- [Examples: Configuring BFD for Static Routes on page 3237](#)
- [Example: Configuring BFD Authentication for Static Routes on page 3253](#)

# Examples: Configuring BFD for Static Routes

- [Understanding BFD for Static Routes for Faster Network Failure Detection on page 3237](#)
- [Example: Configuring BFD for Static Routes for Faster Network Failure Detection on page 3241](#)
- [Example: Enabling BFD on Qualified Next Hops in Static Routes for Route Selection on page 3246](#)

## Understanding BFD for Static Routes for Faster Network Failure Detection

---

The Bidirectional Forwarding Detection (BFD) protocol is a simple hello mechanism that detects failures in a network. BFD works with a wide variety of network environments and topologies. A pair of routing devices exchanges BFD packets. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the routing device stops receiving a reply after a specified interval. The BFD failure detection timers have shorter time limits than the static route failure detection mechanisms, so they provide faster detection.



**NOTE:** EX3300 supports BFD over static routes only.

The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the **clear bfd adaptation** command to return BFD interval timers to their configured values. The **clear bfd adaptation** command is hitless, meaning that the command does not affect traffic flow on the routing device.

By default, BFD is supported on single-hop static routes.

To enable failure detection, include the **bfd-liveness-detection** statement in the static route configuration.

In Junos OS Release 9.1 and later, the BFD protocol is supported for IPv6 static routes. Global unicast and link-local IPv6 addresses are supported for static routes. The BFD protocol is not supported on multicast or anycast IPv6 addresses. For IPv6, the BFD protocol supports only static routes and only in Junos OS Release 9.3 and later. IPv6 for BFD is also supported for the eBGP protocol.



**NOTE:**

Inline BFD is supported on PTX routers with Third generation FPCs starting in Junos OS Release 15.1F3.

There are three types of BFD sessions based on the source from which BFD packets are sent to the neighbors. Different types of BFD sessions and their descriptions are given in the table below:

| Type of BFD session | Description                                            |
|---------------------|--------------------------------------------------------|
| Non-distributed BFD | BFD sessions running completely on the Routing Engine. |
| Distributed BFD     | BFD sessions running on the Packet Forwarding Engine.  |
| Inline BFD          | BFD sessions running on the FPC hardware.              |

**NOTE:** Supported only on static routers starting from Junos OS Release 13.3 and supported on PTX Series routers starting from Junos OS Release 15.1F3.

To configure the BFD protocol for IPv6 static routes, include the **bfd-liveness-detection** statement at the **[edit routing-options rib inet6.0 static route destination-prefix]** hierarchy level.

In Junos OS Release 8.5 and later, you can configure a hold-down interval to specify how long the BFD session must remain up before a state change notification is sent.

To specify the hold-down interval, include the **holddown-interval** statement in the BFD configuration.

You can configure a number in the range from 0 through 255,000 milliseconds. The default is 0. If the BFD session goes down and then comes back up during the hold-down interval, the timer is restarted.



**NOTE:** If a single BFD session includes multiple static routes, the hold-down interval with the highest value is used.

To specify the minimum transmit and receive intervals for failure detection, include the **minimum-interval** statement in the BFD configuration.

This value represents both the minimum interval after which the local routing device transmits hello packets and the minimum interval after which the routing device expects to receive a reply from the neighbor with which it has established a BFD session. You can configure a number in the range from 1 through 255,000 milliseconds. Optionally, instead of using this statement, you can configure the minimum transmit and receive intervals separately using the **transmit-interval**, **minimum-interval**, and **minimum-receive-interval** statements.



**NOTE:** BFD is an intensive protocol that consumes system resources. Specifying a minimum interval for BFD of less than 100 ms for Routing Engine-based sessions and 10 ms for distributed BFD sessions can cause undesired BFD flapping.

Depending on your network environment, these additional recommendations might apply:

- For large-scale network deployments with a large number of BFD sessions, specify a minimum interval of 300 ms for Routing Engine-based sessions and 100 ms for distributed BFD sessions.
- For very large-scale network deployments with a large number of BFD sessions, contact Juniper Networks customer support for more information.
- For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing (NSR) is configured, specify a minimum interval of 2500 ms for Routing Engine-based sessions. For distributed BFD sessions with NSR configured, the minimum interval recommendations are unchanged and depend only on your network deployment.



**NOTE:** SRX Series devices do not support distributed BFD.

To specify the minimum receive interval for failure detection, include the **minimum-receive-interval** statement in the BFD configuration. This value represents the minimum interval after which the routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a number in the range from 1 through 255,000 milliseconds. Optionally, instead of using this statement, you can configure the minimum receive interval using the **minimum-interval** statement at the **[edit routing-options static route destination-prefix bfd-liveness-detection]** hierarchy level.

To specify the number of hello packets not received by the neighbor that causes the originating interface to be declared down, include the **multiplier** statement in the BFD configuration.

The default value is 3. You can configure a number in the range from 1 through 255.

To specify a threshold for detecting the adaptation of the detection time, include the **threshold** statement in the BFD configuration.

When the BFD session detection time adapts to a value equal to or higher than the threshold, a single trap and a system log message are sent. The detection time is based on the multiplier of the **minimum-interval** or the **minimum-receive-interval** value. The threshold must be a higher value than the multiplier for either of these configured values. For example if the **minimum-receive-interval** is 300 ms and the **multiplier** is 3, the total detection time is 900 ms. Therefore, the detection time threshold must have a value higher than 900.

To specify the minimum transmit interval for failure detection, include the **transmit-interval minimum-interval** statement in the BFD configuration.

This value represents the minimum interval after which the local routing device transmits hello packets to the neighbor with which it has established a BFD session. You can configure a value in the range from 1 through 255,000 milliseconds. Optionally, instead of using this statement, you can configure the minimum transmit interval using the **minimum-interval** statement at the **[edit routing-options static route destination-prefix bfd-liveness-detection]** hierarchy level.

To specify the threshold for the adaptation of the transmit interval, include the **transmit-interval threshold** statement in the BFD configuration.

The threshold value must be greater than the transmit interval. When the BFD session transmit time adapts to a value greater than the threshold, a single trap and a system log message are sent. The detection time is based on the multiplier of the value for the **minimum-interval** or the **minimum-receive-interval** statement at the **[edit routing-options static route destination-prefix bfd-liveness-detection]** hierarchy level. The threshold must be a higher value than the multiplier for either of these configured values.

To specify the BFD version, include the **version** statement in the BFD configuration. The default is to have the version detected automatically.

To include an IP address for the next hop of the BFD session, include the **neighbor** statement in the BFD configuration.



**NOTE:** You must configure the **neighbor** statement if the next hop specified is an interface name. If you specify an IP address as the next hop, that address is used as the neighbor address for the BFD session.

---

In Junos OS Release 9.0 and later, you can configure BFD sessions not to adapt to changing network conditions.

To disable BFD adaptation, include the **no-adaptation** statement in the BFD configuration.



**NOTE:** We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation in your network.

---





**NOTE:** If BFD is configured only on one end of a static route, the route is removed from the routing table. BFD establishes a session when BFD is configured on both ends of the static route.

BFD is not supported on ISO address families in static routes. BFD does support IS-IS.

If you configure graceful Routing Engine switchover (GRES) at the same time as BFD, GRES does not preserve the BFD state information during a failover.

#### Related Documentation

- [Example: Configuring BFD for Static Routes for Faster Network Failure Detection on page 3241](#)
- [Example: Enabling BFD on Qualified Next Hops in Static Routes for Route Selection on page 3246](#)

## Example: Configuring BFD for Static Routes for Faster Network Failure Detection

This example shows how to configure Bidirectional Forwarding Detection (BFD) for static routes.

- [Requirements on page 3241](#)
- [Overview on page 3241](#)
- [Configuration on page 3242](#)
- [Verification on page 3245](#)

### Requirements

In this example, no special configuration beyond device initialization is required.

### Overview


There are many practical applications for static routes. Static routing is often used at the network edge to support attachment to stub networks, which, given their single point of entry and egress, are well suited to the simplicity of a static route. In Junos OS, static routes have a global preference of 5. Static routes are activated if the specified next hop is reachable.

In this example, you configure the static route 192.168.47.0/24 from the provider network to the customer network, using the next-hop address of 172.16.1.2. You also configure a static default route of 0.0.0.0/0 from the customer network to the provider network, using a next-hop address of 172.16.1.1.

For demonstration purposes, some loopback interfaces are configured on Device B and Device D. These loopback interfaces provide addresses to ping and thus verify that the static routes are working.

[Figure 53](#) shows the sample network.

Figure 53: Customer Routes Connected to a Service Provider



## Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Device B**

```

set interfaces ge-1/2/0 unit 0 description B->D
set interfaces ge-1/2/0 unit 0 family inet address 172.16.1.1/24
set interfaces lo0 unit 57 family inet address 10.0.0.1/32
set interfaces lo0 unit 57 family inet address 10.0.0.2/32
set routing-options static route 192.168.47.0/24 next-hop 172.16.1.2
set routing-options static route 192.168.47.0/24 bfd-liveness-detection minimum-interval 1000
set protocols bfd traceoptions file bfd-trace
set protocols bfd traceoptions flag all

```

**Device D**

```

set interfaces ge-1/2/0 unit 1 description D->B
set interfaces ge-1/2/0 unit 1 family inet address 172.16.1.2/24
set interfaces lo0 unit 2 family inet address 192.168.47.5/32
set interfaces lo0 unit 2 family inet address 192.168.47.6/32
set routing-options static route 0.0.0.0/0 next-hop 172.16.1.1
set routing-options static route 0.0.0.0/0 bfd-liveness-detection minimum-interval 1000
set protocols bfd traceoptions file bfd-trace
set protocols bfd traceoptions flag all

```

**Step-by-Step Procedure** The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure BFD for static routes:

1. On Device B, configure the interfaces.

- ```
[edit interfaces]
user@B# set ge-1/2/0 unit 0 description B->D
user@B# set ge-1/2/0 unit 0 family inet address 172.16.1.1/24
user@B# set lo0 unit 57 family inet address 10.0.0.1/32
user@B# set lo0 unit 57 family inet address 10.0.0.2/32
```
2. On Device B, create a static route and set the next-hop address.


```
[edit routing-options]
user@B# set static route 192.168.47.0/24 next-hop 172.16.1.2
```
 3. On Device B, configure BFD for the static route.


```
[edit routing-options]
user@B# set static route 192.168.47.0/24 bfd-liveness-detection minimum-interval 1000
```
 4. On Device B, configure tracing operations for BFD.


```
[edit protocols]
user@B# set bfd traceoptions file bfd-trace
user@B# set bfd traceoptions flag all
```
 5. If you are done configuring Device B, commit the configuration.


```
[edit]
user@B# commit
```
 6. On Device D, configure the interfaces.


```
[edit interfaces]
user@D# set ge-1/2/0 unit 1 description D->B
user@D# set ge-1/2/0 unit 1 family inet address 172.16.1.2/24
user@D# set lo0 unit 2 family inet address 192.168.47.5/32
user@D# set lo0 unit 2 family inet address 192.168.47.6/32
```
 7. On Device D, create a static route and set the next-hop address.


```
[edit routing-options]
user@D# set static route 0.0.0.0/0 next-hop 172.16.1.1
```
 8. On Device D, configure BFD for the static route.


```
[edit routing-options]
user@D# set static route 0.0.0.0/0 bfd-liveness-detection minimum-interval 1000
```
 9. On Device D, configure tracing operations for BFD.


```
[edit protocols]
user@D# set bfd traceoptions file bfd-trace
user@D# set bfd traceoptions flag all
```
 10. If you are done configuring Device D, commit the configuration.


```
[edit]
user@D# commit
```

Results

Confirm your configuration by issuing the **show interfaces**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
Device B    user@B# show interfaces
            ge-1/2/0 {
              unit 0 {
                description B->D;
                family inet {
                  address 172.16.1.1/24;
                }
              }
            }
            lo0 {
              unit 57 {
                family inet {
                  address 10.0.0.1/32;
                  address 10.0.0.2/32;
                }
              }
            }

            user@D# show protocols
            bfd {
              traceoptions {
                file bfd-trace;
                flag all;
              }
            }

            user@B# show routing-options
            static {
              route 192.168.47.0/24 {
                next-hop 172.16.1.2;
                bfd-liveness-detection {
                  minimum-interval 1000;
                }
              }
            }

Device D    user@D# show interfaces
            ge-1/2/0 {
              unit 1 {
                description D->B;
                family inet {
                  address 172.16.1.2/24;
                }
              }
            }
            lo0 {
              unit 2 {
                family inet {
                  address 192.168.47.5/32;
                  address 192.168.47.6/32;
                }
              }
            }

            user@D# show routing-options
            static {
```

```

route 0.0.0.0/0 {
  next-hop 172.16.1.1;
  bfd-liveness-detection {
    minimum-interval 1000;
  }
}

```

Verification

Confirm that the configuration is working properly.

- [Verifying That BFD Sessions Are Up on page 3245](#)
- [Viewing Detailed BFD Events on page 3246](#)

Verifying That BFD Sessions Are Up

Purpose Verify that the BFD sessions are up, and view details about the BFD sessions.

Action From operational mode, enter the **show bfd session extensive** command.

```
user@B> show bfd session extensive
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
172.16.1.2	Up	lt-1/2/0.0	3.000	1.000	3

Client Static, TX interval 1.000, RX interval 1.000
 Session up time 00:14:30
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 Replicated, routing table index 172
 Min async interval 1.000, min slow interval 1.000
 Adaptive async TX interval 1.000, RX interval 1.000
 Local min TX interval 1.000, minimum RX interval 1.000, multiplier 3
 Remote min TX interval 1.000, min RX interval 1.000, multiplier 3
 Local discriminator 2, remote discriminator 1
 Echo mode disabled/inactive

1 sessions, 1 clients

Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps

```
user@D> show bfd session extensive
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
172.16.1.1	Up	lt-1/2/0.1	3.000	1.000	3

Client Static, TX interval 1.000, RX interval 1.000
 Session up time 00:14:35
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 Replicated, routing table index 170
 Min async interval 1.000, min slow interval 1.000
 Adaptive async TX interval 1.000, RX interval 1.000
 Local min TX interval 1.000, minimum RX interval 1.000, multiplier 3
 Remote min TX interval 1.000, min RX interval 1.000, multiplier 3
 Local discriminator 1, remote discriminator 2
 Echo mode disabled/inactive

1 sessions, 1 clients

Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps

Meaning The TX interval 1.000, RX interval 1.000 output represents the setting configured with the **minimum-interval** statement. All of the other output represents the default settings for BFD. To modify the default settings, include the optional statements under the **bfd-liveness-detection** statement.

Viewing Detailed BFD Events

Purpose View the contents of the BFD trace file to assist in troubleshooting, if needed.

Action From operational mode, enter the **file show /var/log/bfd-trace** command.

```
user@B> file show /var/log/bfd-trace
Nov 23 14:26:55      Data (9) len 35: (hex) 42 46 44 20 70 65 72 69 6f 64 69 63 20
78 6d 69 74 20 72
Nov 23 14:26:55 PPM Trace: BFD periodic xmit rt tbl index 172
Nov 23 14:26:55 Received Downstream TraceMsg (22) len 108:
Nov 23 14:26:55      IfIndex (3) len 4: 0
Nov 23 14:26:55      Protocol (1) len 1: BFD
Nov 23 14:26:55      Data (9) len 83: (hex) 70 70 6d 64 5f 62 66 64 5f 73 65 6e 64
6d 73 67 20 3a 20
Nov 23 14:26:55 PPM Trace: pcmd_bfd_sendmsg : socket 12 len 24, ifl 78 src
172.16.1.1 dst 172.16.1.2 errno 65
Nov 23 14:26:55 Received Downstream TraceMsg (22) len 93:
Nov 23 14:26:55      IfIndex (3) len 4: 0
Nov 23 14:26:55      Protocol (1) len 1: BFD
Nov 23 14:26:55      Data (9) len 68: (hex) 42 46 44 20 70 65 72 69 6f 64 69 63 20
78 6d 69 74 20 74
```

Meaning BFD messages are being written to the trace file.

Related Documentation

- [Understanding BFD for Static Routes for Faster Network Failure Detection on page 3237](#)

Example: Enabling BFD on Qualified Next Hops in Static Routes for Route Selection

This example shows how to configure a static route with multiple possible next hops. Each next hop has Bidirectional Forwarding Detection (BFD) enabled.

- [Requirements on page 3246](#)
- [Overview on page 3246](#)
- [Configuration on page 3247](#)
- [Verification on page 3250](#)

Requirements

In this example, no special configuration beyond device initialization is required.

Overview


In this example, Device B has the static route **192.168.47.0/24** with two possible next hops. The two next hops are defined using two **qualified-next-hop** statements. Each next hop has BFD enabled.

BFD is also enabled on Device D because BFD must be enabled on both ends of the connection.

A next hop is included in the routing table if the BFD session is up. The next hop is removed from the routing table if the BFD session is down.

See [Figure 54](#).

Figure 54: BFD Enabled on Qualified Next Hops



Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device B

```
set interfaces fe-0/1/0 unit 2 description secondary-B->D
set interfaces fe-0/1/0 unit 2 family inet address 192.168.2.1/24
set interfaces ge-1/2/0 unit 0 description B->D
set interfaces ge-1/2/0 unit 0 family inet address 172.16.1.1/24
set routing-options static route 192.168.47.0/24 qualified-next-hop 192.168.2.2
  bfd-liveness-detection minimum-interval 60
set routing-options static route 192.168.47.0/24 qualified-next-hop 172.16.1.2
  bfd-liveness-detection minimum-interval 60
```

Device D

```
set interfaces fe-0/1/0 unit 3 description secondary-D->B
set interfaces fe-0/1/0 unit 3 family inet address 192.168.2.2/24
set interfaces ge-1/2/0 unit 1 description D->B
set interfaces ge-1/2/0 unit 1 family inet address 172.16.1.2/24
set routing-options static route 0.0.0.0/0 qualified-next-hop 192.168.2.1
set routing-options static route 0.0.0.0/0 qualified-next-hop 172.16.1.1
set routing-options static route 0.0.0.0/0 bfd-liveness-detection minimum-interval 60
```

Step-by-Step Procedure The following example requires that you navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a static route with two possible next hops, both with BFD enabled:

1. On Device B, configure the interfaces.

```
[edit interfaces fe-0/1/0]
user@B# set unit 2 description secondary-B->D
user@B# set unit 2 family inet address 192.168.2.1/24
```

```
[edit interfaces ge-1/2/0]
user@B# set unit 0 description B->D
user@B# set unit 0 family inet address 172.16.1.1/24
```

2. On Device B, configure the static route with two next hops, both with BFD enabled.

```
[edit routing-options static route 192.168.47.0/24]
user@B# set qualified-next-hop 192.168.2.2 bfd-liveness-detection minimum-interval
60
user@B# set qualified-next-hop 172.16.1.2 bfd-liveness-detection minimum-interval
60
```

3. On Device D, configure the interfaces.

```
[edit interfaces fe-0/1/0]
user@D# set unit 3 description secondary-D->B
user@D# set unit 3 family inet address 192.168.2.2/24
```

```
[edit interfaces ge-1/2/0]
user@D# set unit 1 description D->B
user@D# set unit 1 family inet address 172.16.1.2/24
```

4. On Device D, configure a BFD-enabled default static route with two next hops to the provider network.

In this case, BFD is enabled on the route, not on the next hops.

```
[edit routing-options static route 0.0.0.0/0]
user@D# set qualified-next-hop 192.168.2.1
user@D# set qualified-next-hop 172.16.1.1
user@D# set bfd-liveness-detection minimum-interval 60
```

Results Confirm your configuration by issuing the **show interfaces** and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@B# show interfaces
fe-0/1/0 {
  unit 2 {
    description secondary-B->D;
    family inet {
      address 192.168.2.1/24;
    }
  }
}
```



```

ge-1/2/0 {
  unit 0 {
    description B->D;
    family inet {
      address 172.16.1.1/24;
    }
  }
}

user@B# show routing-options
static {
  route 192.168.47.0/24 {
    qualified-next-hop 192.168.2.2 {
      bfd-liveness-detection {
        minimum-interval 60;
      }
    }
    qualified-next-hop 172.16.1.2 {
      bfd-liveness-detection {
        minimum-interval 60;
      }
    }
  }
}

user@D# show interfaces
fe-0/1/0 {
  unit 3 {
    description secondary-D->B;
    family inet {
      address 192.168.2.2/24;
    }
  }
}
ge-1/2/0 {
  unit 1 {
    description D->B;
    family inet {
      address 172.16.1.2/24;
    }
  }
}

user@D# show routing-options
static {
  route 0.0.0.0/0 {
    qualified-next-hop 192.168.2.1;
    qualified-next-hop 172.16.1.1;
    bfd-liveness-detection {
      minimum-interval 60;
    }
  }
}

```

If you are done configuring the devices, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Checking the Routing Tables on page 3250](#)
- [Verifying the BFD Sessions on page 3250](#)
- [Removing BFD from Device D on page 3250](#)
- [Removing BFD from One Next Hop on page 3251](#)

Checking the Routing Tables

Purpose Make sure that the static route appears in the routing table on Device B with two possible next hops.

Action `user@B> show route 192.168.47.0 extensive`
inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
192.168.47.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 192.168.47.0/24 -> {192.168.2.2}
 *Static Preference: 5
 Next hop type: Router
 Address: 0x9334010
 Next-hop reference count: 1
 Next hop: 172.16.1.2 via ge-1/2/0.0
 Next hop: 192.168.2.2 via fe-0/1/0.2, selected
 State: <Active Int Ext>
 Age: 9
 Task: RT
 Announcement bits (1): 3-KRT
 AS path: I

Meaning Both next hops are listed. The next hop 192.168.2.2 is the selected route.

Verifying the BFD Sessions

Purpose Make sure that the BFD sessions are up.

Action `user@B> show bfd session`

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
172.16.1.2	Up	ge-1/2/0.0	0.720	0.240	3
192.168.2.2	Up	fe-0/1/0.2	0.720	0.240	3

2 sessions, 2 clients
Cumulative transmit rate 8.3 pps, cumulative receive rate 8.3 pps

Meaning The output shows that the BFD sessions are up.

Removing BFD from Device D

Purpose Demonstrate what happens when the BFD session is down for both next hops.

- Action** 1. Deactivate BFD on Device D.

```
[edit routing-options static route 0.0.0.0/0]
user@D# deactivate bfd-liveness-detection
user@D# commit
```

2. Rerun the **show bfd session** command on Device B.

```
user@B> show bfd session
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
172.16.1.2	Down	ge-1/2/0.0	3.000	1.000	3
192.168.2.2	Down	fe-0/1/0.2	3.000	1.000	3

```
2 sessions, 2 clients
```

```
Cumulative transmit rate 2.0 pps, cumulative receive rate 2.0 pps
```

3. Rerun the **show route 192.168.47.0** command on Device B.

```
user@B> show route 192.168.47.0
```

Meaning As expected, when the BFD sessions are down, the static route is removed from the routing table.

Removing BFD from One Next Hop

Purpose Demonstrate what happens when only one next hop has BFD enabled.

- Action** 1. If it is not already deactivated, deactivate BFD on Device D.

```
[edit routing-options static route 0.0.0.0/0]
user@D# deactivate bfd-liveness-detection
user@D# commit
```

2. Deactivate BFD on one of the next hops on Device B.

```
[edit routing-options static route 192.168.47.0/24 qualified-next-hop 172.16.1.2]
user@B# deactivate bfd-liveness-detection
user@B# commit
```

3. Rerun the **show bfd session** command on Device B.

```
user@B> show bfd session
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
192.168.2.2	Down	fe-0/1/0.2	3.000	1.000	3

4. Rerun the **show route 192.168.47.0 extensive** command on Device B.

```
user@B> show route 192.168.47.0 extensive
```

```
inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
```

```
192.168.47.0/24 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 192.168.47.0/24 -> {172.16.1.2}
```

```
*Static Preference: 5
```

```
Next hop type: Router, Next hop index: 624
```

```
Address: 0x92f0178
```

```
Next-hop reference count: 3
```

Next hop: 172.16.1.2 via ge-1/2/0.0, selected
State: <Active Int Ext>
Age: 2:36
Task: RT
Announcement bits (1): 3-KRT
AS path: I

Meaning As expected, the BFD session is down for the 192.168.2.2 next hop. The 172.16.1.2 next hop remains in the routing table, and the route remains active, because BFD is not a condition for this next hop to remain valid.

Related Documentation

- *Example: Configuring Static Route Preferences and Qualified Next Hops to Control Static Route Selection*
- *Understanding Static Route Preferences and Qualified Next Hops*
- [Understanding BFD for Static Routes for Faster Network Failure Detection on page 3237](#)
- *Verifying the Static Route Configuration*

Related Documentation

- [Example: Configuring BFD Authentication for Static Routes on page 3253](#)
- [Example: Configuring BFD for OSPF on page 4482](#)
- [Example: Configuring BFD for BGP on page 3811](#)
- *Example: Configuring BFD for IS-IS*
- *Configuring PIM and the Bidirectional Forwarding Detection (BFD) Protocol*

Example: Configuring BFD Authentication for Static Routes

- [Understanding BFD Authentication for Static Route Security on page 3253](#)
- [Example: Configuring BFD Authentication for Securing Static Routes on page 3255](#)

Understanding BFD Authentication for Static Route Security

Bidirectional Forwarding Detection (BFD) enables rapid detection of communication failures between adjacent systems. By default, authentication for BFD sessions is disabled. However, when you run BFD over Network Layer protocols, the risk of service attacks can be significant.



NOTE: We strongly recommend using authentication if you are running BFD over multiple hops or through insecure tunnels.

Beginning with Junos OS Release 9.6, Junos OS supports authentication for BFD sessions running over IPv4 and IPv6 static routes. BFD authentication is not supported on MPLS OAM sessions. BFD authentication is only supported in the Canada and United States version of the Junos OS image and is not available in the export version.



NOTE: EX3300 supports BFD over static routes only.

You authenticate BFD sessions by specifying an authentication algorithm and keychain, and then associating that configuration information with a security authentication keychain using the keychain name.

The following sections describe the supported authentication algorithms, security keychains, and level of authentication that can be configured:

- [BFD Authentication Algorithms on page 3254](#)
- [Security Authentication Keychains on page 3254](#)
- [Strict Versus Loose Authentication on page 3255](#)

BFD Authentication Algorithms

Junos OS supports the following algorithms for BFD authentication:

- **simple-password**—Plain-text password. One to 16 bytes of plain text are used to authenticate the BFD session. One or more passwords can be configured. This method is the least secure and should be used only when BFD sessions are not subject to packet interception.
- **keyed-md5**—Keyed Message Digest 5 hash algorithm for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed MD5 uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than or equal to the last sequence number received. Although more secure than a simple password, this method is vulnerable to replay attacks. Increasing the rate at which the sequence number is updated can reduce this risk.
- **meticulous-keyed-md5**—Meticulous keyed Message Digest 5 hash algorithm. This method works in the same manner as keyed MD5, but the sequence number is updated with every packet. Although more secure than keyed MD5 and simple passwords, this method might take additional time to authenticate the session.
- **keyed-sha-1**—Keyed Secure Hash Algorithm I for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed SHA uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. The key is not carried within the packets. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than the last sequence number received.
- **meticulous-keyed-sha-1**—Meticulous keyed Secure Hash Algorithm I. This method works in the same manner as keyed SHA, but the sequence number is updated with every packet. Although more secure than keyed SHA and simple passwords, this method might take additional time to authenticate the session.



NOTE: Nonstop active routing (NSR) is not supported with meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

Security Authentication Keychains

The security authentication keychain defines the authentication attributes used for authentication key updates. When the security authentication keychain is configured and associated with a protocol through the keychain name, authentication key updates can occur without interrupting routing and signaling protocols.

The authentication keychain contains one or more keychains. Each keychain contains one or more keys. Each key holds the secret data and the time at which the key becomes valid. The algorithm and keychain must be configured on both ends of the BFD session,

and they must match. Any mismatch in configuration prevents the BFD session from being created.

BFD allows multiple clients per session, and each client can have its own keychain and algorithm defined. To avoid confusion, we recommend specifying only one security authentication keychain.

Strict Versus Loose Authentication

By default, strict authentication is enabled, and authentication is checked at both ends of each BFD session. Optionally, to smooth migration from nonauthenticated sessions to authenticated sessions, you can configure *loose checking*. When loose checking is configured, packets are accepted without authentication being checked at each end of the session. This feature is intended for transitional periods only.

Related Documentation

- [Example: Configuring BFD Authentication for Securing Static Routes on page 3255](#)

Example: Configuring BFD Authentication for Securing Static Routes

This example shows how to configure Bidirectional Forwarding Detection (BFD) authentication for static routes.

- [Requirements on page 3255](#)
- [Overview on page 3255](#)
- [Configuration on page 3256](#)
- [Verification on page 3259](#)

Requirements

Junos OS Release 9.6 or later (Canda and United States version).

BFD authentication is only supported in the Canada and United States version of the Junos OS image and is not available in the export version.

Overview

You can configure authentication for BFD sessions running over IPv4 and IPv6 static routes. Routing instances and logical systems are also supported.

The following steps are needed to configure authentication on a BFD session:

1. Specify the BFD authentication algorithm for the static route.
2. Associate the authentication keychain with the static route.
3. Configure the related security authentication keychain. This must be configured on the main router.




TIP: We recommend that you specify loose authentication checking if you are transitioning from nonauthenticated sessions to authenticated sessions.

```
[edit]
user@host> set routing-options static route ipv4 bfd-liveness-detection
authentication loose-check
```

Figure 55 shows the sample network.

Figure 55: Customer Routes Connected to a Service Provider



Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device B

```
set interfaces ge-1/2/0 unit 0 description B->D
set interfaces ge-1/2/0 unit 0 family inet address 172.16.1.1/24
set interfaces lo0 unit 57 family inet address 10.0.0.1/32
set interfaces lo0 unit 57 family inet address 10.0.0.2/32
set routing-options static route 192.168.47.0/24 next-hop 172.16.1.2
set routing-options static route 192.168.47.0/24 bfd-liveness-detection minimum-interval 1000
set routing-options static route 192.168.47.0/24 bfd-liveness-detection authentication key-chain bfd-kc4
set routing-options static route 192.168.47.0/24 bfd-liveness-detection authentication algorithm keyed-sha-1
set security authentication-key-chains key-chain bfd-kc4 key 5 secret "$9$JhZHmn6Ap0In/9ApOcSs24oaZikPfT3wY24ZG.mz36AtOIEyMWxSrlKvM-dbs2aDkP5FtOIQFcleV7N"
set security authentication-key-chains key-chain bfd-kc4 key 5 start-time "2011-1-1.12:00:00 -0800"
```

Device D

```
set interfaces ge-1/2/0 unit 1 description D->B
set interfaces ge-1/2/0 unit 1 family inet address 172.16.1.2/24
```



```

set interfaces lo0 unit 2 family inet address 192.168.47.5/32
set interfaces lo0 unit 2 family inet address 192.168.47.6/32
set routing-options static route 0.0.0.0/0 next-hop 172.16.1.1
set routing-options static route 0.0.0.0/0 bfd-liveness-detection minimum-interval 1000
set routing-options static route 0.0.0.0/0 bfd-liveness-detection authentication key-chain
  bfd-kc4
set routing-options static route 0.0.0.0/0 bfd-liveness-detection authentication algorithm
  keyed-sha-1
set security authentication-key-chains key-chain bfd-kc4 key 5 secret
  "$9$JhZHmn6Ap0In/9ApOcSs24oaZikPfT3wY24ZG.mz36AtOIEyMWxSrlKvM-dbs2a
  DkP5FtOIQFclv7N"
set security authentication-key-chains key-chain bfd-kc4 key 5 start-time
  "2011-1-1.12:00:00 -0800"

```

Step-by-Step Procedure The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure BFD for static routes:

1. On Device B, configure the interfaces.

```

[edit interfaces]
user@B# set ge-1/2/0 unit 0 description B->D
user@B# set ge-1/2/0 unit 0 family inet address 172.16.1.1/24

```

```

user@B# set lo0 unit 57 family inet address 10.0.0.1/32
user@B# set lo0 unit 57 family inet address 10.0.0.2/32

```

2. On Device B, create a static route and set the next-hop address.

```

[edit routing-options]
user@B# set static route 192.168.47.0/24 next-hop 172.16.1.2

```

3. On Device B, configure BFD for the static route.

```

[edit routing-options]
user@B# set static route 192.168.47.0/24 bfd-liveness-detection minimum-interval
  1000

```

4. On Device B, specify the algorithm (**keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**, or **simple-password**) to use for BFD authentication on the static route.

```

[edit routing-options]
user@B# set static route 192.168.47.0/24 bfd-liveness-detection authentication
  algorithm keyed-sha-1

```



NOTE: Nonstop active routing (NSR) is not supported with the **meticulous-keyed-md5** and **meticulous-keyed-sha-1** authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

5. On Device B, specify the keychain to be used to associate BFD sessions on the specified route with the unique security authentication keychain attributes.

This should match the keychain name configured at the **[edit security authentication key-chains]** hierarchy level.

```
[edit routing-options]
user@B# set static route 192.168.47.0/24 bfd-liveness-detection authentication
key-chain bfd-kc4
```

6. On Device B, specify the unique security authentication information for BFD sessions:

- The matching keychain name as specified in Step 5.
- At least one key, a unique integer between 0 and 63. Creating multiple keys allows multiple clients to use the BFD session.
- The secret data used to allow access to the session.
- The time at which the authentication key becomes active, in the format *yyyy-mm-dd.hh:mm:ss*.

```
[edit security authentication-key-chains key-chain bfd-kc4]
user@B# set key 5 secret
"$9$JhZHmn6Ap0In/9ApOcSs24oaZikPfT3wY24ZG.mz36AtOIeYMWxSrlKvM-dbs2a
DkP5Ft0IQFclev7N"
user@B# set key 5 start-time "2011-1-1.12:00:00 -0800"
```

7. If you are done configuring Device B, commit the configuration.

```
[edit]
user@B# commit
```

8. Repeat the configuration on Device D.

The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

Results

Confirm your configuration by issuing the **show interfaces**, **show routing-options**, and **show security** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
Device B user@B# show interfaces
ge-1/2/0 {
  unit 0 {
    description B->D;
    family inet {
      address 172.16.1.1/24;
    }
  }
}
lo0 {
  unit 57 {
    family inet {
      address 10.0.0.1/32;
    }
  }
}
```

```

        address 10.0.0.2/32;
    }
}

user@B# show routing-options
static {
  route 192.168.47.0/24 {
    next-hop 172.16.1.2;
    bfd-liveness-detection {
      minimum-interval 1000;
      authentication {
        key-chain bfd-kc4;
        algorithm keyed-sha-1;
      }
    }
  }
}

user@B# show security
authentication-key-chains {
  key-chain bfd-kc4 {
    key 5 {
      secret
        "$9$JhZHmn6Ap0ln/9ApOcSs24oaZikPft3wY24ZG.mz36AtOIEyMWxSrlKvM-dbs2a
        DkP5FtOIQFclev7N"; ## SECRET-DATA
      start-time "2011-1-1.12:00:00 -0800";
    }
  }
}

```

Verification

Confirm that the configuration is working properly.

- [Verifying That BFD Sessions Are Up on page 3259](#)
- [Viewing Details About the BFD Session on page 3260](#)
- [Viewing Extensive BFD Session Information on page 3260](#)

Verifying That BFD Sessions Are Up

Purpose Verify that the BFD sessions are up.

Action From operational mode, enter the **show bfd session** command.

```

user@B> show bfd session

```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
172.16.1.2	Up	ge-1/2/0.0	3.000	1.000	3

```

1 sessions, 1 clients
Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps

```

Meaning The command output shows that the BFD session is up.

Viewing Details About the BFD Session

Purpose View details about the BFD sessions and make sure that authentication is configured.

Action From operational mode, enter the **show bfd session detail** command.

```
user@B> show bfd session detail
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
172.16.1.2	Up	ge-1/2/0.0	3.000	1.000	3

Client Static, TX interval 1.000, RX interval 1.000, **Authenticate**
 Session up time 00:53:58
 Local diagnostic NbrSignal, remote diagnostic None
 Remote state Up, version 1
 Logical system 9, routing table index 22

1 sessions, 1 clients
 Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps

Meaning In the command output, **Authenticate** is displayed to indicate that BFD authentication is configured.

Viewing Extensive BFD Session Information

Purpose View more detailed information about the BFD sessions.

Action From operational mode, enter the **show bfd session extensive** command.

```
user@B> show bfd session extensive
```

Address	State	Interface	Time	Interval	Multiplier
172.16.1.2	Up	ge-1/2/0.0	3.000	1.000	3

Client Static, TX interval 1.000, RX interval 1.000, **Authenticate**
 keychain bfd-kc4, algo keyed-sha-1, mode strict
 Session up time 01:39:45
 Local diagnostic NbrSignal, remote diagnostic None
 Remote state Up, version 1
 Logical system 9, routing table index 22
 Min async interval 1.000, min slow interval 1.000
 Adaptive async TX interval 1.000, RX interval 1.000
 Local min TX interval 1.000, minimum RX interval 1.000, multiplier 3
 Remote min TX interval 1.000, min RX interval 1.000, multiplier 3
 Local discriminator 3, remote discriminator 4
 Echo mode disabled/inactive
 Authentication enabled/active, keychain bfd-kc4, algo keyed-sha-1, mode strict

1 sessions, 1 clients
 Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps

Meaning In the command output, **Authenticate** is displayed to indicate that BFD authentication is configured. The output for the **extensive** command provides the keychain name, the authentication algorithm, and the mode for each client in the session.

Related Documentation

- [Understanding BFD Authentication for Static Route Security on page 3253](#)

- Related Documentation**
- [Examples: Configuring BFD for Static Routes on page 3237](#)

Troubleshooting Virtual Routing Instances

- [Direct Routes Not Leaked Between Routing Instances on page 3263](#)

Direct Routes Not Leaked Between Routing Instances

Problem **Description:** Direct routes are not exported (leaked) between virtual routing instances. For example, consider the following scenario:

- Switch with two virtual routing instances:
 - Routing instance 1 connects to downstream device through interface xe-0/0/1.
 - Routing instance 2 connects to upstream device through interface xe-0/0/2.

If you enable route leaking between the routing instances (by using the **rib-group** statement, for example), the downstream device cannot connect to the upstream device because the switch connects to the upstream device over a direct route and these routes are not leaked between instances.



NOTE: You can see a route to the upstream device in the routing table of the downstream device, but this route is not functional.

Indirect routes *are* leaked between routing instances, so the downstream device can connect to any upstream devices that are connected to the switch over indirect routes.

Solution This is expected behavior.

Related Documentation

- [Understanding Virtual Router Routing Instances on page 3232](#)
- [Configuring Virtual Router Routing Instances on page 3232](#)
- [rib-group on page 3364](#)

Related Documentation

- [Understanding Virtual Router Routing Instances on page 3232](#)
- [Configuring Virtual Router Routing Instances on page 3232](#)
- [rib-group on page 3364](#)

PART 60

Configuration Statements and Operational Commands

- [Configuration Statements \(Routing Options\) on page 3267](#)
- [Operational Commands \(Routing Options\) on page 3389](#)

Configuration Statements (Routing Options)

- [active](#) on page 3270
- [aggregate \(Routing\)](#) on page 3272
- [as-path \(Routing Options\)](#) on page 3274
- [autonomous-system](#) on page 3276
- [backup-pe-group](#) on page 3278
- [backups](#) on page 3279
- [bandwidth \(Multicast Flow Map\)](#) on page 3280
- [bfd-liveness-detection \(Routing Options Static Route\)](#) on page 3281
- [bgp-orf-cisco-mode](#) on page 3285
- [bmp](#) on page 3287
- [brief](#) on page 3289
- [centralized](#) on page 3290
- [community \(Routing Options\)](#) on page 3291
- [confederation](#) on page 3293
- [disable \(Routing Options\)](#) on page 3294
- [description \(Routing Instances\)](#) on page 3294
- [discard](#) on page 3295
- [export](#) on page 3296
- [export-rib](#) on page 3297
- [fate-sharing](#) on page 3300
- [flow](#) on page 3301
- [flow-map](#) on page 3302
- [forwarding-cache \(Flow Maps\)](#) on page 3303
- [forwarding-cache \(Multicast\)](#) on page 3304
- [forwarding-options \(chassis\)](#) on page 3306
- [forwarding-table](#) on page 3308

- [generate](#) on page 3309
- [graceful-restart \(Enabling Globally\)](#) on page 3311
- [import](#) on page 3312
- [import-policy](#) on page 3313
- [import-rib](#) on page 3314
- [indirect-next-hop](#) on page 3315
- [install \(Routing Options\)](#) on page 3316
- [instance-export](#) on page 3317
- [instance-import](#) on page 3317
- [instance-type](#) on page 3318
- [interface \(Multicast Static Routes\)](#) on page 3319
- [interface \(Routing Instances\)](#) on page 3320
- [interface \(Routing Options\)](#) on page 3321
- [interface-routes](#) on page 3322
- [local-address \(Routing Options\)](#) on page 3323
- [martians](#) on page 3324
- [maximum-bandwidth \(Routing Options\)](#) on page 3325
- [maximum-paths](#) on page 3326
- [maximum-prefixes](#) on page 3328
- [med-igp-update-interval](#) on page 3329
- [metric \(Aggregate, Generated, or Static Route\)](#) on page 3330
- [multicast \(Routing Options\)](#) on page 3331
- [no-qos-adjust](#) on page 3332
- [options \(Routing Options\)](#) on page 3333
- [pim-to-igmp-proxy](#) on page 3334
- [pim-to-mld-proxy](#) on page 3335
- [policy \(Aggregate and Generated Routes\)](#) on page 3336
- [policy \(Flow Maps\)](#) on page 3337
- [policy-options](#) on page 3338
- [policy-statement](#) on page 3339
- [ppm](#) on page 3343
- [ppm \(Ethernet Switching\)](#) on page 3344
- [preference \(Routing Options\)](#) on page 3345
- [prefix](#) on page 3346
- [protocols](#) on page 3347
- [qualified-next-hop \(Static Routes\)](#) on page 3350
- [readvertise](#) on page 3352

- [redundant-sources](#) on page 3353
- [resolution](#) on page 3354
- [resolution-ribs](#) on page 3355
- [resolve](#) on page 3356
- [restart-duration \(Routing Options\)](#) on page 3357
- [retain](#) on page 3358
- [reverse-oif-mapping](#) on page 3359
- [rpf-check-policy \(Routing Options RPF\)](#) on page 3360
- [rib \(General\)](#) on page 3361
- [rib \(Route Resolution\)](#) on page 3363
- [rib-group \(Routing Options\)](#) on page 3364
- [rib-groups](#) on page 3365
- [route-distinguisher-id](#) on page 3367
- [route-record](#) on page 3368
- [router-id](#) on page 3369
- [routing-instances](#) on page 3370
- [routing-options](#) on page 3371
- [scope](#) on page 3371
- [scope-policy](#) on page 3372
- [source \(Source-Specific Multicast\)](#) on page 3373
- [source-routing](#) on page 3374
- [ssm-groups](#) on page 3375
- [ssm-map \(Routing Options Multicast\)](#) on page 3376
- [static \(Routes\)](#) on page 3377
- [subscriber-leave-timer](#) on page 3379
- [tag \(Routing Options\)](#) on page 3380
- [threshold \(Multicast Forwarding Cache\)](#) on page 3381
- [timeout \(Flow Maps\)](#) on page 3382
- [timeout \(Multicast\)](#) on page 3383
- [traceoptions](#) on page 3384
- [upstream-interface](#) on page 3387

active

Syntax	(active passive);
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (aggregate generate static) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options (aggregate generate static) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options (aggregate generate static) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</p> <p>[edit routing-options (aggregate generate static) (defaults route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Determine whether static, aggregate, or generated routes are removed from the routing and forwarding tables when they become inactive. Static routes are only removed from the routing table if the next hop becomes unreachable. This can occur if the local or neighbor interface goes down. Routes that have been configured to remain continually installed in the routing and forwarding tables are marked with reject next hops when they are inactive.</p> <ul style="list-style-type: none"> • active—Remove a route from the routing and forwarding tables when it becomes inactive. • passive—Have a route remain continually installed in the routing and forwarding tables even when it becomes inactive. <p>Include the active statement when configuring an individual route in the route portion of the static statement to override a passive option specified in the defaults portion of the statement.</p>
Default	active
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Examples: Configuring Static Routes</i> • <i>Example: Summarizing Static Routes Through Route Aggregation</i> • <i>Example: Configuring a Conditional Default Route Policy</i>

- *Understanding Conditionally Generated Routes*

aggregate (Routing)

Syntax	<pre> aggregate { defaults { ... aggregate-options ... } route destination-prefix { policy policy-name; ... aggregate-options ... } } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i>],</p> <p>[edit routing-options],</p> <p>[edit routing-options rib <i>routing-table-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Configure aggregate routes.
Options	<p>aggregate-options—Additional information about aggregate routes that is included with the route when it is installed in the routing table. Specify zero or more of the following options in aggregate-options. Each option is explained separately.</p> <ul style="list-style-type: none"> • (active passive); • as-path <as-path> <origin (egp igp incomplete)> <atomic-aggregate> <aggregator as-number ip-address>; • (brief full); • community [<i>community-ids</i>]; • discard; • (metric metric2 metric3 metric4) <i>value</i> <type type>; • (preference preference2 color color2) <i>preference</i> <type type>; • tag <i>metric type number</i>; <p>defaults—Specify global aggregate route options. These options only set default attributes inherited by all newly created aggregate routes. These are treated as global defaults</p>

and apply to all the aggregate routes you configure in the **aggregate** statement. This part of the **aggregate** statement is optional.

route *destination-prefix*—Configure a nondefault aggregate route:

- **default**—For the default route to the destination. This is equivalent to specifying an IP address of **0.0.0.0/0**.
- ***destination-prefix/prefix-length***—***destination-prefix*** is the network portion of the IP address, and ***prefix-length*** is the destination prefix length.

The **policy** statement is explained separately.

Required Privilege Level	routing—To view this statement in the configuration.
	routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Summarizing Static Routes Through Route Aggregation</i>

as-path (Routing Options)

Syntax	<code>as-path <as-path> <aggregator as-number ip-address> <atomic-aggregate> <origin (egp igp incomplete)>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (aggregate generate static) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options (aggregate generate static) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options (aggregate generate static) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</p> <p>[edit routing-options (aggregate generate static) (defaults route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Associate BGP autonomous system (AS) path information with a static, aggregate, or generated route.</p> <p>In Junos OS Release 9.1 and later, the numeric range for the AS number is extended to provide BGP support for 4-byte AS numbers as defined in RFC 4893, <i>BGP Support for Four-octet AS Number Space</i>. RFC 4893 introduces two new optional transitive BGP attributes, AS4_PATH and AS4_AGGREGATOR. These new attributes are used to propagate 4-byte AS path information across BGP speakers that do not support 4-byte AS numbers. RFC 4893 also introduces a reserved, well-known, 2-byte AS number, AS 23456. This reserved AS number is called AS_TRANS in RFC 4893. All releases of Junos OS support 2-byte AS numbers.</p> <p>In Junos OS Release 9.2 and later, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: <i><16-bit high-order value in decimal>.<16-bit low-order value in decimal></i>. For example, the 4-byte AS number of 65,546 in plain-number format is represented as 1.10 in the AS-dot notation format. You can specify a value in the range from 0.0 through 65535.65535 in AS-dot notation format.</p>
Default	No AS path information is associated with static routes.
Options	<p>aggregator—(Optional) Attach the BGP aggregator path attribute to the aggregate route. You must specify the last AS number that formed the aggregate route (encoded as two octets) for as-number, followed by the IP address of the BGP system that formed the aggregate route for ip-address.</p>

as-path—(Optional) AS path to include with the route. It can include a combination of individual AS path numbers and AS sets. Enclose sets in brackets ([]). The first AS number in the path represents the AS immediately adjacent to the local AS. Each subsequent number represents an AS that is progressively farther from the local AS, heading toward the origin of the path. You cannot specify a regular expression for **as-path**. You must use a complete, valid AS path.

atomic-aggregate—(Optional) Attach the BGP **atomic-aggregate** path attribute to the aggregate route. This path attribute indicates that the local system selected a less specific route instead of a more specific route.

origin egp—(Optional) BGP origin attribute that indicates that the path information originated in another AS.

origin igp—(Optional) BGP origin attribute that indicates that the path information originated within the local AS.

origin incomplete—(Optional) BGP origin attribute that indicates that the path information was learned by some other means.

Required Privilege	routing—To view this statement in the configuration.
Level	routing-control—To add this statement to the configuration.

Related Documentation	<ul style="list-style-type: none"> • <i>Examples: Configuring Static Routes</i> • <i>Example: Summarizing Static Routes Through Route Aggregation</i> • <i>Example: Conditionally Generating Static Routes</i> • <i>Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview</i>
------------------------------	--

autonomous-system

Syntax	<pre>autonomous-system <i>autonomous-system</i> <asdot-notation> <loops <i>number</i>> { independent-domain <no-attrset>; }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options], [edit routing-options]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>asdot-notation option introduced in Junos OS Release 9.3.</p> <p>asdot-notation option introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>no-attrset option introduced in Junos OS Release 10.4.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Specify the routing device's AS number.</p> <p>An autonomous system (AS) is a set of routing devices that are under a single technical administration and that generally use a single interior gateway protocol (IGP) and metrics to propagate routing information within the set of routing devices. An AS appears to other ASs to have a single, coherent interior routing plan and presents a consistent picture of what destinations are reachable through it. ASs are identified by a number that is assigned by the Network Information Center (NIC) in the United States (http://www.isi.edu).</p> <p>If you are using BGP on the routing device, you must configure an AS number.</p> <p>The AS path attribute is modified when a route is advertised to an EBGP peer. Each time a route is advertised to an EBGP peer, the local routing device prepends its AS number to the existing path attribute, and a value of 1 is added to the AS number.</p> <p>In Junos OS Release 9.1 and later, the numeric range is extended to provide BGP support for 4-byte AS numbers as defined in RFC 4893, <i>BGP Support for Four-octet AS Number Space</i>. RFC 4893 introduces two new optional transitive BGP attributes, AS4_PATH and AS4_AGGREGATOR. These new attributes are used to propagate 4-byte AS path information across BGP speakers that do not support 4-byte AS numbers. RFC 4893 also introduces a reserved, well-known, 2-byte AS number, AS 23456. This reserved AS number is called AS_TRANS in RFC 4893. All releases of Junos OS support 2-byte AS numbers.</p> <p>In Junos OS Release 9.3 and later, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: <i><16-bit high-order value in decimal>.<16-bit low-order value in decimal></i>. For example, the 4-byte AS number of 65,546 in plain-number format is represented as 1.10 in the AS-dot notation format.</p>
Options	<p><i>autonomous-system</i>—AS number. Use a number assigned to you by the NIC.</p>

Range: 1 through 4,294,967,295 ($2^{32} - 1$) in plain-number format for 4-byte AS numbers

In this example, the 4-byte AS number 65,546 is represented in plain-number format:

```
[edit]
routing-options {
  autonomous-system 65546;
}
```

Range: 0.0 through 65535.65535 in AS-dot notation format for 4-byte numbers

In this example, 1.10 is the AS-dot notation format for 65,546:

```
[edit]
routing-options {
  autonomous-system 1.10;
}
```

Range: 1 through 65,535 in plain-number format for 2-byte AS numbers (this is a subset of the 4-byte range)

In this example, the 2-byte AS number 60,000 is represented in plain-number format:

```
[edit]
routing-options {
  autonomous-system 60000;
}
```

asdot-notation—(Optional) Display the configured 4-byte autonomous system number in the AS-dot notation format.

Default: Even if a 4-byte AS number is configured in the AS-dot notation format, the default is to display the AS number in the plain-number format.

loops number—(Optional) Specify the number of times detection of the AS number in the AS_PATH attribute causes the route to be discarded or hidden. For example, if you configure **loops 1**, the route is hidden if the AS number is detected in the path one or more times. This is the default behavior. If you configure **loops 2**, the route is hidden if the AS number is detected in the path two or more times.

Range: 1 through 10

Default: 1



NOTE: When you specify the same AS number in more than one routing instance on the local routing device, you must configure the same number of loops for the AS number in each instance. For example, if you configure a value of 3 for the loops statement in a VRF routing instance that uses the same AS number as that of the master instance, you must also configure a value of 3 loops for the AS number in the master instance.

Use the **independent-domain** option if the loops statement must be enabled only on a subset of routing instances.

The remaining statement is explained separately.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Examples: Configuring External BGP Peering on page 3601• Examples: Configuring Internal BGP Peering on page 3624

backup-pe-group

Syntax	<pre>backup-pe-group <i>group-name</i> { backups [<i>addresses</i>]; local-address <i>address</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast], [edit logical-systems <i>logical-system-name</i> routing-options multicast], [edit routing-instances <i>routing-instance-name</i> routing-options multicast], [edit routing-options multicast]
Release Information	Statement introduced in Junos OS Release 9.0. Statement introduced in Junos OS Release 9.5 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure a backup provider edge (PE) group for ingress PE redundancy when point-to-multipoint label-switched paths (LSPs) are used for multicast distribution.
Options	<p>backups <i>addresses</i>—Specify the address of backup PE routers for ingress PE redundancy when point-to-multipoint LSPs are used for multicast distribution.</p> <p>local-address <i>address</i>—Specify the address of the local PE router for ingress PE redundancy when point-to-multipoint LSPs are used for multicast distribution.</p> <p>pe-group-name—Specify the name for the group of PE routers that provide ingress PE router redundancy for point-to-multipoint LSPs.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Ingress PE Redundancy• Configuring Ingress PE Router Redundancy for Point-to-Multipoint LSPs

backups

Syntax	<code>backups [<i>addresses</i>];</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast backup-pe-group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast backup-pe-group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast backup-pe-group <i>group-name</i>],</p> <p>[edit routing-options multicast backup-pe-group <i>group-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.0.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Configure the address of backup PEs for ingress PE redundancy when point-to-multipoint label-switched paths (LSPs) are used for multicast distribution.
Options	<i>addresses</i> —Addresses of other PEs in the backup group.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Ingress PE Redundancy</i>

bandwidth (Multicast Flow Map)

Syntax	<code>bandwidth (<i>bps</i> <i>adaptive</i>);</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast flow-map],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-options multicast flow-map],</code> <code>[edit routing-instances <i>routing-instance-name</i> routing-options multicast flow-map],</code> <code>[edit routing-options multicast flow-map]</code>
Release Information	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the bandwidth property for multicast flow maps.
Options	<i>adaptive</i> —Specify that the bandwidth is measured for the flows that are matched by the flow map. <i>bps</i> —Bandwidth, in bits per second, for the flow map. Range: 0 through any amount of bandwidth Default: 2 Mbps
Required Privilege Level	routing —To view this statement in the configuration. routing-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring a Multicast Flow Map</i>

bfd-liveness-detection (Routing Options Static Route)

Syntax `bfd-liveness-detection {`
 `authentication {`
 `algorithm` *algorithm-name*;
 `key-chain` *key-chain-name*;
 `loose-check`;
 `}`
 `detection-time {`
 `threshold` *milliseconds*;
 `}`
 `holddown-interval` *milliseconds*;
 `local-address` *ip-address*;
 `minimum-interval` *milliseconds*;
 `minimum-receive-interval` *milliseconds*;
 `minimum-receive-ttl` *number*;
 `multiplier` *number*;
 `neighbor` *address*;
 `no-adaptation`;
 `transmit-interval {`
 `minimum-interval` *milliseconds*;
 `threshold` *milliseconds*;
 `}`
 `version` (1 | automatic);
`}`

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* static route *destination-prefix*],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options static route *destination-prefix*],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)],
 [edit logical-systems *logical-system-name* routing-options rib *routing-table-name* static route *destination-prefix*],
 [edit logical-systems *logical-system-name* routing-options rib *routing-table-name* static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)],
 [edit logical-systems *logical-system-name* routing-options static route *destination-prefix*],
 [edit logical-systems *logical-system-name* routing-options static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)],
 [edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* static route *destination-prefix*],
 [edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)],
 [edit routing-instances *routing-instance-name* routing-options static route *destination-prefix*],
 [edit routing-instances *routing-instance-name* routing-options static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)],
 [edit routing-options rib *routing-table-name* static route *destination-prefix*],
 [edit routing-options rib *routing-table-name* static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)],
 [edit routing-options static route *destination-prefix*],

[edit routing-options static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)]

Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>detection-time threshold and transmit-interval threshold options introduced in Junos OS Release 8.2.</p> <p>local-address statement introduced in Junos OS Release 8.2.</p> <p>minimum-receive-ttl statement introduced in Junos OS Release 8.2.</p> <p>Support for logical routers introduced in Junos OS Release 8.3.</p> <p>holddown-interval statement introduced in Junos OS Release 8.5.</p> <p>no-adaptation statement introduced in Junos OS Release 9.0.</p> <p>Support for IPv6 static routes introduced in Junos OS Release 9.1.</p> <p>authentication algorithm, authentication key-chain, and authentication loose-check statements introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Configure bidirectional failure detection timers and authentication criteria for static routes.</p>

Options **authentication algorithm** *algorithm-name* —Configure the algorithm used to authenticate the specified BFD session: **simple-password**, **keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, or **meticulous-keyed-sha-1**.

authentication key-chain *key-chain-name* —Associate a security key with the specified BFD session using the name of the security keychain. The name you specify must match one of the keychains configured in the **authentication-key-chains key-chain** statement at the **[edit security]** hierarchy level.

authentication loose-check —(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication may not be configured at both ends of the BFD session.

detection-time threshold *milliseconds* —Configure a threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

holddown-interval *milliseconds* —Configure an interval specifying how long a BFD session must remain up before a state change notification is sent. If the BFD session goes down and then comes back up during the hold-down interval, the timer is restarted.

Range: 0 through 255,000

Default: 0

local-address *ip-address* —Enable a multihop BFD session and configure the source address for the BFD session.

minimum-interval *milliseconds* —Configure the minimum interval after which the local routing device transmits a hello packet and then expects to receive a reply from the neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit and receive intervals separately using the **transmit-interval**, **minimum-interval**, and **minimum-receive-interval** statements.

Range: 1 through 255,000

minimum-receive-interval *milliseconds* —Configure the minimum interval after which the routing device expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the **minimum-interval** statement at the **[edit routing-options static route destination-prefix bfd-liveness-detection]** hierarchy level.

Range: 1 through 255,000

minimum-receive-ttl *number* —Configure the time to live (TTL) for the multihop BFD session.

Range: 1 through 255

Default: 255

multiplier *number* —Configure number of hello packets not received by the neighbor that causes the originating interface to be declared down.

Range: 1 through 255

Default: 3

neighbor *address*—Configure a next-hop address for the BFD session for a next hop specified as an interface name.

no-adaptation—Specify for BFD sessions not to adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

transmit-interval threshold *milliseconds*—Configure the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

Range: 0 through 4,294,967,295

transmit-interval minimum-interval *milliseconds*—Configure the minimum interval at which the routing device transmits hello packets to a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit interval using the **minimum-interval** statement at the **[edit routing-options static route *destination-prefix* bfd-liveness-detection]** hierarchy level.

Range: 1 through 255,000

version—Configure the BFD version to detect: **1** (BFD version 1) or **automatic** (autodetect the BFD version).

Default: automatic

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
---------------------------------	---

Related Documentation	<ul style="list-style-type: none">• Example: Configuring BFD for Static Routes for Faster Network Failure Detection on page 3241• Example: Configuring BFD Authentication for Securing Static Routes on page 3255
------------------------------	--

bgp-orf-cisco-mode

Syntax	<code>bgp-orf-cisco-mode;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options outbound-route-filter],</p> <p>[edit protocols bgp outbound-route-filter],</p> <p>[edit protocols bgp group <i>group-name</i> outbound-route-filter],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> outbound-route-filter],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp outbound-route-filter],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> outbound-route-filter],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> outbound-route-filter],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options outbound-route-filter],</p> <p>[edit routing-options outbound-route-filter]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.2.</p> <p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Support for the BGP group and neighbor hierarchy levels introduced in Junos OS Release 9.2.</p> <p>Support for the BGP group and neighbor hierarchy levels introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Enable interoperability with routing devices that use the vendor-specific outbound route filter compatibility code of 130 and code type of 128.



NOTE: To enable interoperability for all BGP peers configured on the routing device, include the statement at the [edit routing-options outbound-route-filter] hierarchy level.

Default Disabled

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring BGP Prefix-Based Outbound Route Filtering on page 3773

bmp

```
Syntax  bmp {
    authentication-algorithm (aes-128-cmac-96 | hmac-sha-1-96 | md5);
    authentication-key key;
    authentication-key-chain authentication-key-chain;
    connection-mode (active | passive);
    hold-down {
        seconds;
        flaps flaps;
        period seconds;
    }
    initiation-message text;
    local-address address;
    local-port port;
    monitor (disable | enable);
    priority (high | low | medium);
    route-monitoring {
        none;
        post-policy {
            exclude-non-eligible;
        }
        pre-policy {
            exclude-non-feasible;
        }
    }
}
station station-name {
    authentication-algorithm (aes-128-cmac-96 | hmac-sha-1-96 | md5);
    authentication-key key;
    authentication-key-chain authentication-key-chain;
    connection-mode (active | passive);
    hold-down {
        seconds;
        flaps flaps;
        period seconds;
    }
    initiation-message text;
    local-address address;
    local-port port;
    monitor (disable | enable);
    priority (high | low | medium);
    route-monitoring {
        none;
        post-policy {
            exclude-non-eligible;
        }
        pre-policy {
            exclude-non-feasible;
        }
    }
}
station-address (ip-address | name);
station-port port-number;
statistics-timeout seconds;
traceoptions {
```

```

        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier>;
    }
}
station-address (ip-address | name);
station-port port-number;
statistics-timeout seconds;
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier>;
}
}

```

Hierarchy Level [edit routing-options]
 [edit logical-systems *logical-system-name* routing-options],
 [edit logical-systems *logical-system-name* protocols *bgp*],
 [edit logical-systems *logical-system-name* protocols *bgp* *group* *group-name*],
 [edit logical-systems *logical-system-name* protocols *bgp* *group* *group-name* *neighbor* *address*],
 [edit protocols *bgp*],
 [edit protocols *bgp* *group* *group-name*],
 [edit protocols *bgp* *group* *group-name* *neighbor* *address*],



NOTE: 1. Complete BMP configuration, as mentioned in the syntax, can be done under the first two hierarchy levels only

2. Under other hierarchy levels, only the following configurations are supported:


- Either we can inherit or not inherit the configuration data
- Enable/disable monitoring
- Control route monitoring settings

Release Information	Statement introduced in Junos OS Release 9.5. Statement introduced in Junos OS Release 9.5 for EX Series switches. Statement introduced in Junos OS Release 12.3 for ACX Series routers. Support for BMP version 3 introduced in Junos OS Release 13.3.
Description	Configure the BGP Monitoring Protocol (BMP), which enables the routing device to collect data from the BGP Adjacency-RIB-In routing tables and periodically send that data to a monitoring station.
Options	The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring the BGP Monitoring Protocol</i>

brief

Syntax	(brief full);
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (aggregate generate) (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate) (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-options (aggregate generate) (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate) (defaults route)], [edit routing-instances <i>routing-instance-name</i> routing-options (aggregate generate) (defaults route)], [edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate) (defaults route)], [edit routing-options (aggregate generate) (defaults route)], [edit routing-options rib <i>routing-table-name</i> (aggregate generate) (defaults route)]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Configure all AS numbers from all contributing paths to be included in the aggregate or generated route's path.</p> <ul style="list-style-type: none"> • brief—Include only the longest common leading sequences from the contributing AS paths. If this results in AS numbers being omitted from the aggregate route, the BGP ATOMIC_ATTRIBUTE path attribute is included with the aggregate route. • full—Include all AS numbers from all contributing paths in the aggregate or generated route's path. Include this option when configuring an individual route in the route portion of the generate statement to override a retain option specified in the defaults portion of the statement.
Default	full
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Summarizing Static Routes Through Route Aggregation</i> • <i>Example: Configuring a Conditional Default Route Policy</i> • <i>Understanding Conditionally Generated Routes</i> • aggregate on page 3272 • generate on page 3309

centralized

Syntax	centralized;
Hierarchy Level	[edit protocols lacp ppm]
Release Information	Statement introduced in Junos OS Release 9.4 for MX Series routers. Statement introduced in Junos OS Release 10.2 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Disable distributed periodic packet management (PPM) processing for Link Aggregation Control Protocol (LACP) packets and run all PPM processing for LACP packets on the Routing Engine.</p> <p>This statement disables distributed PPM processing for only LACP packets. You can disable distributed PPM processing for all packets that use PPM and run all PPM processing on the Routing Engine by configuring the no-delegate-processing statement in the [edit routing-options ppm] hierarchy.</p>
<div>  <p>BEST PRACTICE: We generally recommend that you disable distributed PPM only if Juniper Networks Customer Service advises you to do so. You should disable distributed PPM only if you have a compelling reason to disable it.</p> </div>	
Default	Distributed PPM processing is enabled for all packets that use PPM.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Distributed Periodic Packet Management on an EX Series Switch (CLI Procedure)</i> • <i>Configuring Aggregated Ethernet LACP (CLI Procedure)</i> • Configuring Distributed Periodic Packet Management on page 3234 • Configuring Link Aggregation on page 2869

community (Routing Options)

Syntax	<code>community ([<i>community-ids</i>] no-advertise no-export no-export-subconfed none llgr-stale no-llgr);</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (aggregate generate static) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options (aggregate generate static) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options (aggregate generate static) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</p> <p>[edit routing-options (aggregate generate static) (defaults route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> <p>llgr-stale and no-llgr options added in Junos OS Release 15.1.</p>
Description	Associate BGP community information with a static, aggregate, or generated route.
Default	No BGP community information is associated with static routes.
Options	<p><i>community-ids</i>—One or more community identifiers. The <i>community-ids</i> format varies according to the type of attribute that you use.</p> <p>The BGP community attribute format is <i>as-number:community-value</i>:</p> <ul style="list-style-type: none"> • <i>as-number</i>—AS number of the community member. It can be a value from 1 through 65,535. The AS number can be a decimal or hexadecimal value. • <i>community-value</i>—Identifier of the community member. It can be a number from 0 through 65,535. <p>For more information about BGP community attributes, see the “Configuring the Extended Communities Attribute” section in the <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i>.</p> <p>For specifying the BGP community attribute only, you also can specify <i>community-ids</i> as one of the following well-known community names defined in RFC 1997:</p> <ul style="list-style-type: none"> • no-advertise—Routes containing this community name are not advertised to other BGP peers.

- **no-export**—Routes containing this community name are not advertised outside a BGP confederation boundary.
- **no-export-subconfed**—Routes containing this community are advertised to IBGP peers with the same AS number, but not to members of other confederations.
- **llgr-stale**—Adds a community to a long-lived stale route when it is readvertised.
- **no-llgr**—Marks routes which a BGP speaker does not want to be retained by LLGR. The Notification message feature does not have any associated configuration parameters.



NOTE: Extended community attributes are not supported at the [edit routing-options] hierarchy level. You must configure extended communities at the [edit policy-options] hierarchy level. For information about configuring extended communities, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Examples: Configuring Static Routes*
- *Example: Summarizing Static Routes Through Route Aggregation*
- *Example: Conditionally Generating Static Routes*
- [aggregate on page 3272](#)
- [generate on page 3309](#)
- *static*

confederation

Syntax	<code>confederation <i>confederation-autonomous-system</i> members [<i>autonomous-systems</i>];</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Specify the routing device's confederation AS number.</p> <p>If you administer multiple ASs that contain a very large number of BGP systems, you can group them into one or more <i>confederations</i>. Each confederation is identified by its own AS number, which is called a <i>confederation AS number</i>. To external ASs, a confederation appears to be a single AS. Thus, the internal topology of the ASs making up the confederation is hidden.</p> <p>The BGP path attributes NEXT_HOP, LOCAL_PREF, and MULTI_EXIT_DISC, which normally are restricted to a single AS, are allowed to be propagated throughout the ASs that are members of the same confederation.</p> <p>Because each confederation is treated as if it were a single AS, you can apply the same routing policy to all the ASs that make up the confederation.</p> <p>Grouping ASs into confederations reduces the number of BGP connections required to interconnect ASs.</p> <p>If you are using BGP, you can enable the local routing device to participate as a member of an AS confederation. To do this, include the confederation statement.</p> <p>Specify the AS confederation identifier, along with the peer AS numbers that are members of the confederation.</p> <p>Note that peer adjacencies do not form if two BGP neighbors disagree about whether an adjacency falls within a particular confederation.</p>
Options	<p><i>autonomous-systems</i>—AS numbers of the confederation members. Range: 1 through 65,535</p> <p><i>confederation-autonomous-system</i>—Confederation AS number. Use one of the numbers assigned to you by the NIC. Range: 1 through 65,535</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

- Related Documentation**
- [Example: Configuring BGP Confederations on page 3891](#)
 - [Understanding BGP Confederations on page 3890](#)


disable (Routing Options)

Syntax	disable;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options graceful-restart], [edit logical-systems <i>logical-system-name</i> routing-options graceful-restart], [edit routing-instances <i>routing-instance-name</i> routing-options graceful-restart], [edit routing-options graceful-restart]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Disable graceful restart.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	• Junos OS High Availability Library for Routing Devices

description (Routing Instances)

Syntax	description text;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>], [edit routing-instances <i>routing-instance-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Provide a text description for the routing instance. If the text includes one or more spaces, enclose it in quotation marks (" "). Any descriptive text you include is displayed in the output of the show route instance detail command and has no effect on the operation of the routing instance.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	• Configuring Routing Instances on PE Routers in VPNs • show route instance on page 3496

discard

Syntax	<code>discard;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (aggregate generate) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options (aggregate generate) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options (aggregate generate) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate) (defaults route)],</p> <p>[edit routing-options (aggregate generate) (defaults route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> (aggregate generate) (defaults route)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Do not forward packets addressed to this destination. Instead, drop the packets, do not send ICMP unreachable messages to the packets' originators, and install a reject route for this destination into the routing table.</p> <p>To propagate static routes into the routing protocols, include the discard statement when you define the route, along with a routing policy.</p>
	<p> NOTE: In other vendors' software, a common way to propagate static routes into routing protocols is to configure the routes so that the next-hop routing device is the loopback address (commonly, 127.0.0.1). However, configuring static routes in this way (by including a statement such as route address/mask-length next-hop 127.0.0.1) does not propagate the static routes, because the forwarding table ignores static routes whose next-hop routing device is the loopback address.</p>
Default	When an aggregate route becomes active, it is installed in the routing table with a reject next hop, which means that ICMP unreachable messages are sent.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <i>Example: Summarizing Static Routes Through Route Aggregation</i>

- [Example: Conditionally Generating Static Routes](#)
- [aggregate on page 3272](#)
- [generate on page 3309](#)

export

Syntax	<code>export [<i>policy-name</i>];</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options forwarding-table], [edit logical-systems <i>logical-system-name</i> routing-options forwarding-table], [edit routing-instances <i>routing-instance-name</i> routing-options forwarding-table], [edit routing-options forwarding-table]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Apply one or more policies to routes being exported from the routing table into the forwarding table.</p> <p>In the export statement, list the name of the routing policy to be evaluated when routes are being exported from the routing table into the forwarding table. Only active routes are exported from the routing table.</p> <p>You can reference the same routing policy one or more times in the same or a different export statement.</p> <p>You can apply export policies to routes being exported from the routing table into the forwarding table for the following features:</p> <ul style="list-style-type: none">• Per-packet load balancing• Class of service (CoS)
Options	<i>policy-name</i> —Name of one or more policies.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Load Balancing BGP Traffic on page 3828

export-rib

Syntax	<code>export-rib routing-table-name;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib-groups <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-options rib-groups <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> routing-options rib-groups <i>group-name</i>], [edit routing-options rib-groups <i>group-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify the name of the routing table from which Junos OS should export routing information. For any individual RIB group, only one table can be specified in the export-rib statement.

The **export-rib** statement specifies the source table from which routing information is advertised.

One common use of the **export-rib** statement is interdomain routing. The export RIB is the table used when BGP extracts routes to advertise to peers. In multicast interdomain routing, for example, the export RIB is likely to be inet.2.

Another use of **export-rib** is dynamic route leaking between the global routing table (inet.0) and a VRF routing table (*instance.inet.0*). For example, you can use a RIB group to copy routes learned in the VRF into the global routing table, inet.0, or copy routes learned in inet.0 into a VRF. You define the use of this RIB group in the VRF's BGP configuration. In a routing policy you can do dynamic filtering of routes. For instance, you can use an import policy to only copy routes with certain communities into the global routing table.

For example:

```
rib-groups {
  rib-interface-routes-v4 {
    import-rib [ inet.0 VRF.inet.0 ];
  }
  rib-import-VRF-routes-to-inet0-v4 {
    export-rib VRF.inet.0;
    import-rib [ VRF.inet.0 inet.0 ];
    import-policy rib-import-VRF-routes-to-inet0-v4;
  }
  rib-import-inet0-routes-to-VRF-v4 {
    export-rib inet.0;
    import-rib [ inet.0 VRF.inet.0 ];
    import-policy rib-import-inet0-routes-to-VRF-v4;
  }
}
routing-options {
  interface-routes {
```

```
        rib-group {
            inet rib-interface-routes-v4;
        }
    }
}
protocols {
    bgp {
        group iBGP-peers {
            type internal;
            family inet {
                unicast {
                    rib-group rib-import-inet0-routes-to-VRF-v4;
                }
            }
        }
    }
}
routing-instances {
    VRF {
        routing-options {
            interface-routes {
                rib-group {
                    inet rib-interface-routes-v4;
                }
            }
        }
        protocols {
            bgp {
                group peersin-VRF {
                    family inet {
                        unicast {
                            rib-group rib-import-VRF-routes-to-inet0-v4;
                        }
                    }
                }
            }
        }
    }
}
```

Options *routing-table-name*—Routing table group name.

Required Privilege routing—To view this statement in the configuration.
Level routing-control—To add this statement to the configuration.

**Related
Documentation**

- *Example: Exporting Specific Routes from One Routing Table Into Another Routing Table*
- *Example: Configuring a PIM RPF Routing Table*
- *Example: Configuring DVMRP to Announce Unicast Routes*
- *Example: Configuring a Dedicated PIM RPF Routing Table*
- *Example: Configuring Any-Source Multicast for Draft-Rosen VPNs*
- [import-rib on page 3314](#)
- *passive*

fate-sharing

Syntax	<pre>fate-sharing { group <i>group-name</i> { cost <i>value</i>; from <i>address</i> <to <i>address</i>>; } }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options], [edit routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Specify a backup path in case the primary path becomes unusable.</p> <p>You specify one or more objects with common characteristics within a group. All objects are treated as /32 host addresses. The objects can be a LAN interface, a router ID, or a point-to-point link. Sequence is insignificant.</p> <p>Changing the fate-sharing database does not affect existing established LSPs until the next CSPF reoptimization. The fate-sharing database does affect fast-reroute detour path computations.</p>
Options	<p>cost <i>value</i>—Cost assigned to the group. Range: 1 through 65,535 Default: 1</p> <p>from <i>address</i>—Address of the router or address of the LAN/NBMA interface. For example, an Ethernet network with four hosts in the same fate-sharing group would require you to list all four of the separate from addresses in the group.</p> <p>group <i>group-name</i>—Each fate-sharing group must have a name, which can have a maximum of 32 characters, including letters, numbers, periods (.), and hyphens (-). You can define up to 512 groups.</p> <p>to <i>address</i>—(Optional) Address of egress router. For point-to-point link objects, you must specify both a from and a to address.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Ingress Router for MPLS-Signaled LSPs</i>• <i>MPLS Applications Feature Guide for Routing Devices</i>

flow

Syntax	<pre> flow { route <i>name</i> { match { <i>match-conditions</i>; } term-order (legacy standard); then { <i>actions</i>; } } firewall-install-disable; term-order (legacy standard); validation { traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier> <disable>; } } } </pre>
Hierarchy Level	<p>[edit routing-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>term-order statement introduced in Junos OS Release 10.0</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>firewall-install-disable statement introduced in Junos OS Releases 12.1X48 and 12.3 for PTX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Configure a flow route.
Default	legacy
Options	<p><i>actions</i>—An action to take if conditions match.</p> <p>firewall-install-disable—(PTX Series routers only) Disable installing flow-specification firewall filters in the firewall process (dfwd).</p> <p>Default: If you omit the firewall-install-disable statement, the default behavior is firewall-install-disable mode.</p> <p><i>match-conditions</i>—Match packets to these conditions.</p> <p>route <i>name</i>—Name of the flow route.</p> <p>standard—Specify to use version 7 or later of the flow-specification algorithm.</p> <p>term-order (legacy standard)—Specify the version of the flow-specification algorithm.</p>

- **legacy**—Use version 6 of the flow-specification algorithm.
- **standard**—Use version 7 of the flow-specification algorithm.

then—Actions to take on matching packets.

The remaining statements are explained separately.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Enabling BGP to Carry Flow-Specification Routes</i>• <i>Understanding BGP Flow Routes for Traffic Filtering</i>

flow-map

Syntax	<pre>flow-map <i>flow-map-name</i> { bandwidth (<i>bps</i> adaptive); forwarding-cache { timeout (never non-discard-entry-only <i>minutes</i>); } policy [<i>policy-names</i>]; redundant-sources [<i>addresses</i>]; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast], [edit logical-systems <i>logical-system-name</i> routing-options multicast], [edit routing-instances <i>routing-instance-name</i> routing-options multicast], [edit routing-options multicast]
Release Information	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure multicast flow maps.
Options	<i>flow-map-name</i> —Name of the flow-map. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring a Multicast Flow Map</i>

forwarding-cache (Flow Maps)

Syntax	forwarding-cache { timeout (minutes never non-discard-entry-only); }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast flow-map <i>flow-map-name</i>], [edit logical-systems <i>logical-system-name</i> routing-options multicast flow-map <i>flow-map-name</i>], [edit routing-instances <i>routing-instance-name</i> routing-options multicast flow-map <i>flow-map-name</i>], [edit routing-options multicast flow-map <i>flow-map-name</i>]
Release Information	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure multicast forwarding cache properties for the flow map.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

forwarding-cache (Multicast)

Syntax	<pre>forwarding-cache { allow-maximum; family (inet inet6) { threshold { log-warning value; suppress value <reuse value>; } } threshold { log-warning value; suppress value <reuse value>; } timeout minutes; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast], [edit logical-systems <i>logical-system-name</i> routing-options multicast], [edit routing-instances <i>routing-instance-name</i> routing-options multicast], [edit routing-options multicast]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Configure multicast forwarding cache properties. These properties include threshold suppression and reuse limits, the threshold at which a warning message is logged, and timeout values.</p> <p>Specify a value for the threshold at which to suppress new multicast forwarding cache entries and an optional reuse value for the threshold at which the router begins to create new multicast forwarding cache entries. The range for both is from 1 through 200,000. If configured, the reuse value should be less than the suppression threshold value. The suppression value is mandatory. If you do not specify the optional reuse value, then the number of multicast forwarding cache entries is limited to the suppression value. A new entry is created as soon as the number of multicast forwarding cache entries falls below the suppression value.</p> <p>You can configure the thresholds globally for the multicast forwarding cache or individually for the IPv4 and IPv6 multicast forwarding caches. Configuring the threshold statement globally for the multicast forwarding cache or including the family statement to configure the thresholds for the IPv4 and IPv6 multicast forwarding caches are mutually exclusive.</p>
Default	By default, there are no limits on the number of multicast forwarding cache entries.
Options	family (inet inet6) —(Optional) Apply the configured thresholds to either IPv4 or IPv6 multicast forwarding cache entries.

Default: By default, the configured thresholds are applied to both IPv4 and IPv6 multicast forwarding cache entries.

The remaining statements are explained separately.

Required Privilege	routing—To view this statement in the configuration.
Level	routing-control—To add this statement to the configuration.

Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring the Multicast Forwarding Cache</i>
------------------------------	--

forwarding-options (chassis)

Syntax forwarding options *profile-name* {
 num-65-127-prefix *number*;
}

forwarding-options lpm-profile {
 prefix-65-127-disable;
 unicast-in-lpm
}

Syntax forwarding-options custom-profile {
 l2-entries | l3-entries | lpm-entries {
 num-banks *number*;
 }
}

Hierarchy Level [edit [chassis](#)]

Release Information Statement introduced in Junos OS Release 13.2.
custom-profile option introduced in Junos OS Release 15.1x53-D30 for QFX5200 Series switches only.

Description Configure a unified forwarding table profile to allocate the amount a memory available for the following:

- MAC addresses.
- Layer 3 host entries.
- Longest prefix match table entries.

This feature enables you to select a profile that optimizes the amount of memory available for various types of forwarding-table entries based on the needs of your network. For example, for a switch that handles a great deal of Layer 2 traffic, such as a virtualized network with many servers and virtualized machines, you would choose the **l2-profile-one**, which allocates the highest amount of memory to MAC addresses.

You configure the memory allocation for LPM table entries differently, depending on whether you using Junos OS Release 13.2X51-D10 or Junos OS Release 13.2X51-D15 and later. For more information about configuring memory allocation for LPM table entries see *Configuring the Unified Forwarding Table*.

The **num-65-127-prefix *number*** statement is not supported on the **custom-profile** and the **lpm-profile**. The **prefix-65-127-disable** and **unicast-in-lpm** statements are supported only on the **lpm-profile**.

Options **profile-name**—name of the profile to use for memory allocation in the unified forwarding table. [Table 291](#) lists the profiles you can choose that have set values and the associated values for each type of entry.

On QFX5200 Series switches only, you can also select **custom-profile**. This profile enables you to allocate from one to four banks of shared hash memory to a specific

type of forwarding-table entry. Each shared hash memory bank can store a maximum of the equivalent of 32,000 IPv4 unicast addresses.

Table 291: Unified Forwarding Table Profiles

Profile Name	MAC Table	Host Table (unicast and multicast addresses)					
	MAC Addresses	IPv4 unicast	IPv6 unicast	IPv4 (*, G)	IPv4 (S, G)	IPv6 (*, G)	IPv6 (S, G)
l2-profile-one	288K	16K	8K	8K	8K	4K	4K
l2-profile-two	224K	80K	40K	40K	40K	20K	20K
l2-profile-three (default)	160K	144K	72K	72K	72K	36K	36K
l3-profile	96K	208K	104K	104K	104K	52K	52K
lpm-profile*	32K	16K	8K	8K	8K	4K	4K

* This profile supports only IPv4 in Junos OS Release 13.2X51-D10. Starting in Junos OS Release 13.2X51-D15, the **lpm-profile** supports IPv4 and IPv6 entries.



NOTE: If the host stores the maximum number of entries for any given type, the entire table is full and is unable to accommodate *any* entries of any other type. For information about valid combinations of table entries see *Understanding the Unified Forwarding Table*.

l2-entries | l3-entries | lpm-entries—(custom-profile only) Select a type of forwarding-table entry—Layer 2, Layer 3, or LPM—to allocate a specific number of shared memory banks. You configure the amount of memory to allocate for each type of entry separately.

num-banks *number*—(custom-profile only) Specify the number of shared memory banks to allocate for a specific type of forwarding-table entry. Each shared memory bank stores the equivalent of 32,000 IPv4 unicast addresses.

Range: 0 through 4.



NOTE: There are four shared memory banks, which can be allocated flexibly among the three types of forwarding-table entries. To allocate no shared memory for a particular entry type, specify the number 0. When you commit the configuration, the system issues a commit check to ensure that you have not configured more than four memory banks. You do not have to configure all four shared memory banks. By default, each entry type is allocated the equivalent of 32,000 IPv4 unicast addresses in shared memory.

Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Understanding the Unified Forwarding Table</i>• <i>Example: Configuring a Unified Forwarding Table Custom Profile on QFX Series Switches</i>

forwarding-table

Syntax	<pre>forwarding-table { export [policy--names]; (indirect-next-hop no-indirect-next-hop); }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-options]
Release Information	Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure information about the routing device's forwarding table. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Per-Packet Load Balancing on page 3228

generate

Syntax	<pre>generate { defaults { generate-options; } route destination-prefix { policy policy-name; generate-options; } }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i>],</p> <p>[edit routing-options],</p> <p>[edit routing-options rib <i>routing-table-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Configure generated routes, which are used as routes of last resort.</p>
Options	<p>defaults—(Optional) Specify global generated route options. These options only set default attributes inherited by all newly created generated routes. These are treated as global defaults and apply to all the generated routes you configure in the generate statement.</p> <p>generate-options—Additional information about generated routes, which is included with the route when it is installed in the routing table. Specify zero or more of the following options in generate-options. Each option is explained separately.</p> <ul style="list-style-type: none"> • (active passive); • as-path <<i>as-path</i>> <origin (egp igp incomplete)> <atomic-aggregate> <aggregator <i>as-number in-address</i>>; • (brief full); • community [<i>community-ids</i>]; • discard; • (metric metric2 metric3 metric4) <i>value</i> <type <i>type</i>>; • (preference preference2 color color2) <i>preference</i> <type <i>type</i>>; • tag <i>metric type number</i>; <p>route destination-prefix—Configure a non-default generated route:</p>


- **default**—For the default route to the destination. This is equivalent to specifying an IP address of **0.0.0.0/0**.
- **destination-prefix/prefix-length**—/**destination-prefix** is the network portion of the IP address, and **prefix-length** is the destination prefix length.

The [policy](#) statement is explained separately.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
---------------------------------	---

Related Documentation	<ul style="list-style-type: none">• <i>Example: Conditionally Generating Static Routes</i>
------------------------------	--

graceful-restart (Enabling Globally)


Syntax	<pre> graceful-restart { disable; helper-disable; maximum-helper-recovery-time <i>seconds</i>; maximum-helper-restart-time <i>seconds</i>; notify-duration <i>seconds</i>; recovery-time <i>seconds</i>; restart-duration <i>seconds</i>; stale-routes-time <i>seconds</i>; } </pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options], [edit routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure graceful restart globally to enable the feature. You cannot enable graceful restart for specific protocols unless graceful restart is also enabled globally. You can, optionally, modify the global settings at the individual protocol level.
<div>  NOTE: <ul style="list-style-type: none"> For VPNs, the <code>graceful-restart</code> statement allows a router whose VPN control plane is undergoing a restart to continue to forward traffic while recovering its state from neighboring routers. For BGP, if you configure graceful restart after a BGP session has been established, the BGP session restarts and the peers negotiate graceful restart capabilities. LDP sessions flap when <code>graceful-restart</code> configurations change. </div>	
Default	Graceful restart is disabled by default.
Options	The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> Enabling Graceful Restart Configuring Routing Protocols Graceful Restart on page 2580

- *Configuring Graceful Restart for MPLS-Related Protocols*
- *Configuring VPN Graceful Restart*
- *Configuring Logical System Graceful Restart*
- *Graceful Restart Configuration Statements*
- *Configuring Graceful Restart for QFabric Systems*


import

Syntax	<code>import [<i>policy-names</i>];</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options resolution rib], [edit logical-systems <i>logical-system-name</i> routing-options resolution rib], [edit routing-instances <i>routing-instance-name</i> routing-options resolution rib], [edit routing-options resolution rib]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify one or more import policies to use for route resolution.
Options	<i>policy-names</i> —Name of one or more import policies.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring Route Resolution on PE Routers</i>


import-policy

Syntax	<code>import-policy [<i>policy-names</i>];</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib-groups <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib-groups <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib-groups <i>group-name</i>],</p> <p>[edit routing-options rib-groups <i>group-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Apply one or more policies to routes imported into the routing table group. The import-policy statement complements the import-rib statement and cannot be used unless you first specify the routing tables to which routes are being imported.</p>
<div>  <p>NOTE: On EX Series switches, only dynamically learned routes can be imported from one routing table group to another.</p> </div>	
Options	<i>policy-names</i> —Name of one or more policies.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Exporting Specific Routes from One Routing Table Into Another Routing Table</i> • export-rib on page 3297 • <i>passive</i>

import-rib

Syntax	<code>import-rib [<i>routing-table-names</i>];</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib-groups <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib-groups <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib-groups <i>group-name</i>],</p> <p>[edit routing-options rib-groups <i>group-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Specify the name of the routing table into which Junos OS should import routing information. The first routing table name you enter is the primary routing table. Any additional names you enter identify secondary routing tables. When a protocol imports routes, it imports them into the primary and any secondary routing tables. If the primary route is deleted, the secondary route also is deleted. For IPv4 import routing tables, the primary routing table must be inet.0 or routing-instance-name.inet.0. For IPv6 import routing tables, the primary routing table must be inet6.0.</p> <p>In Junos OS Release 9.5 and later, you can configure an IPv4 import routing table that includes both IPv4 and IPv6 routing tables. Including both types of routing tables permits you, for example, to populate an IPv6 routing table with IPv6 addresses that are compatible with IPv4. In releases prior to Junos OS Release 9.5, you could configure an import routing table with only either IPv4 or IPv6 routing tables.</p>
	<p> NOTE: On EX Series switches, only dynamically learned routes can be imported from one routing table group to another.</p>
Options	<i>routing-table-names</i> —Name of one or more routing tables.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Exporting Specific Routes from One Routing Table Into Another Routing Table</i> • export-rib on page 3297 • <i>passive</i>

indirect-next-hop

Syntax	(indirect-next-hop no-indirect-next-hop);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options forwarding-table], [edit routing-options forwarding-table]
Release Information	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Enable indirectly connected next hops for route convergence. This statement is implemented on the Packet Forward Engine to speed up forwarding information base (FIB) updates. Configuring this statement significantly speeds convergence times. The only downside of configuring this statement is that some additional FIB memory overhead is required. Unless routes have an extremely high number of next hops, this increased memory usage should not be noticeable.
<div>  <div> <p>NOTE:</p> <ul style="list-style-type: none"> When virtual private LAN service (VPLS) is configured on the routing device, the indirect-next-hop statement is configurable at the [edit routing-options forwarding-table] hierarchy level. However, this configuration is not applicable to indirect nexthops specific to VPLS routing instances. By default, the Junos Trio Modular Port Concentrator (MPC) chipset on MX Series routers is enabled with indirectly connected next hops, and this cannot be disabled using the no-indirect-next-hop statement. By default, indirectly connected next hops are enabled on PTX Series routers. </div> </div>	
Default	Disabled.
Options	indirect-next-hop —Enable indirectly connected next hops. no-indirect-next-hop —Explicitly disable indirect next hops.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> <i>Example: Optimizing Route Reconvergence by Enabling Indirect Next Hops on the Packet Forwarding Engine</i>

install (Routing Options)

Syntax	(install no-install);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options static (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> static (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> static (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-options static (defaults route)], [edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> static (defaults route)], [edit routing-instances <i>routing-instance-name</i> routing-options static (defaults route)], [edit routing-options rib <i>routing-table-name</i> static (defaults route)] [edit routing-options static (defaults route)]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure whether Junos OS installs all static routes into the forwarding table. Even if you configure a route so it is not installed in the forwarding table, the route is still eligible to be exported from the routing table to other protocols.
Options	install —Explicitly install all static routes into the forwarding table. Include this statement when configuring an individual route in the route portion of the static statement to override a no-install option specified in the defaults portion of the statement. no-install —Do not install the route into the forwarding table, even if it is the route with the lowest preference. Default: install
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Examples: Configuring Static Routes</i>• <i>static</i>

instance-export

Syntax	<code>instance-export [<i>policy-names</i>];</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options], [edit routing-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Apply one or more policies to routes being exported from a routing instance.
Options	<i>policy-names</i> —Name of one or more export policies.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i>

instance-import

Syntax	<code>instance-import [<i>policy-names</i>];</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options], [edit routing-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Apply one or more policies to routes being imported into a routing instance.
Options	<i>policy-names</i> —Name of one or more import policies.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i>

instance-type

Syntax	instance-type virtual-router
Hierarchy Level	[edit routing-instances]
Release Information	Statement introduced in Junos OS Release 9.2 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify the type of routing instance.
Options	virtual-router —Virtual router routing instance.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Using Virtual Routing Instances to Route Among VLANs on EX Series Switches</i>• <i>Configuring Virtual Routing Instances (CLI Procedure)</i>• Configuring Virtual Router Routing Instances on page 3232

interface (Multicast Static Routes)

Syntax	<pre> interface <i>interface-names</i> { disable; maximum-bandwidth <i>bps</i>; no-qos-adjust; reverse-oif-mapping { no-qos-adjust; } subscriber-leave-timer <i>seconds</i>; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit routing-options multicast]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Enable multicast traffic on an interface.</p> <p>By default, multicast packets are forwarded by enabling Protocol Independent Multicast (PIM) on an interface. PIM adds multicast routes into the routing table.</p> <p>You can also configure multicast packets to be forwarded over a static route, such as a static route associated with an LSP next hop. Multicast packets are accepted on an interface and forwarded over a static route in the forwarding table. This is useful when you want to enable multicast traffic on a specific interface without configuring PIM on the interface.</p> <p>You cannot enable multicast traffic on an interface and configure PIM on the same interface simultaneously.</p> <p>Static routes must be configured before you can enable multicast on an interface. Configuring the interface statement alone does not install any routes into the routing table. This feature relies on the static route configuration.</p>
Options	<p><i>interface-names</i>—Name of one or more interfaces on which to enable multicast traffic.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Defining Interface Bandwidth Maximums</i> • <i>Example: Configuring Multicast with Subscriber VLANs</i>

interface (Routing Instances)

Syntax	<code>interface <i>interface-name</i>;</code>
Hierarchy Level	[edit routing-instances]
Release Information	Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	For virtual router routing instances, configure an interface.

**NOTE:**

- You must configure only interfaces from the Node devices assigned to the network Node group. If you try to configure interfaces from the Node devices assigned to server Node groups, the configuration does not commit.
 - You can configure an interface for one routing instance only.
-

Options	<i>interface-name</i> —Name of an interface.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Virtual Router Routing Instances on page 3232

interface (Routing Options)

Syntax	<pre> interface <i>interface-names</i> { maximum-bandwidth <i>bps</i>; no-qos-adjust; reverse-oif-mapping { no-qos-adjust; } subscriber-leave-timer <i>seconds</i>; } </pre>
Hierarchy Level	<pre> [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast], [edit logical-systems <i>logical-system-name</i> routing-options multicast], [edit routing-instances <i>routing-instance-name</i> routing-options multicast], [edit routing-options multicast] </pre>
Release Information	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Enable multicast traffic on an interface.



TIP: You cannot enable multicast traffic on an interface by using the **routing-options multicast interface** statement and configure PIM on the interface.

Options	<p><i>interface-name</i>—Names of the physical or logical interface.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Defining Interface Bandwidth Maximums</i> • <i>Example: Configuring Multicast with Subscriber VLANs</i>

interface-routes

Syntax

```
interface-routes {  
    family (inet | inet6) {  
        export {  
            lan;  
            point-to-point;  
        }  
    }  
    rib-group group-name;  
}
```

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.



NOTE: On EX Series switches, only dynamically learned routes can be imported from one routing table group to another.

Description Associate a routing table group with the routing device's interfaces, and specify routing table groups into which interface routes are imported.

By default, IPv4 interface routes (also called direct routes) are imported into routing table **inet.0**, and IPv6 interface routes are imported into routing table **inet6.0**. If you are configuring alternate routing tables for use by some routing protocols, it might be necessary to import the interface routes into the alternate routing tables. To define the routing tables into which interface routes are imported, you create a routing table group and associate it with the routing device's interfaces.

To create the routing table groups, include the **passive** statement at the **[edit routing-options]** hierarchy level.

If you have configured a routing table, configure the OSPF primary instance at the **[edit protocols ospf]** hierarchy level with the statements needed for your network so that routes are installed in **inet.0** and in the forwarding table. Make sure to include the routing table group.

To export local routes, include the **export** statement.

To export LAN routes, include the **lan** option. To export point-to-point routes, include the **point-to-point** option.

Only local routes on point-to-point interfaces configured with a destination address are exportable.

- Options**
- inet**—Specify the IPv4 address family.
 - inet6**—Specify the IPv6 address family.
 - lan**—Export LAN routes.
 - point-to-point**—Export point-to-point routes.

The remaining statements are explained separately. See [CLI Explorer](#).

- Required Privilege Level**
- routing—To view this statement in the configuration.
 - routing-control—To add this statement to the configuration.

- Related Documentation**
- *Example: Importing Direct and Static Routes Into a Routing Instance*
 - *Example: Configuring Multiple Routing Instances of OSPF*
 - *passive*

local-address (Routing Options)

Syntax `local-address address;`

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name
  routing-options multicast backup-pe-group group-name],
[edit logical-systems logical-system-name routing-options multicast backup-pe-group
  group-name],
[edit routing-instances routing-instance-name routing-options multicast backup-pe-group
  group-name],
[edit routing-options multicast backup-pe-group group-name]
```

Release Information

Statement introduced in Junos OS Release 9.0.
 Statement introduced in Junos OS Release 9.0 for EX Series switches.
 Statement introduced in Junos OS Release 11.3 for the QFX Series.
 Statement added to the multicast hierarchy in Junos OS Release 13.2.

Description Configure the address of the local PE for ingress PE redundancy when point-to-multipoint LSPs are used for multicast distribution.

Options **address**—Address of local PEs in the backup group.

Required Privilege Level

- routing—To view this statement in the configuration.
- routing-control—To add this statement to the configuration.

- Related Documentation**
- *Example: Configuring Ingress PE Redundancy*

martians

Syntax	<pre>martians { destination-prefix match-type <allow>; }</pre>
Hierarchy Level	<pre>[edit logical-systems logical-system-name routing-instances routing-instance-name routing-options], [edit logical-systems logical-system-name routing-instances routing-instance-name routing-options rib routing-table-name], [edit logical-systems logical-system-name routing-options], [edit logical-systems logical-system-name routing-options rib routing-table-name], [edit routing-instances routing-instance-name routing-options], [edit routing-instances routing-instance-name routing-options rib routing-table-name], [edit routing-options], [edit routing-options rib routing-table-name]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Configure martian addresses.
Options	<p>allow—(Optional) Explicitly allow a subset of a range of addresses that has been disallowed. The allow option is the only supported action.</p> <p>destination-prefix—Destination route you are configuring:</p> <ul style="list-style-type: none">• destination-prefix/prefix-length—destination-prefix is the network portion of the IP address, and prefix-length is the destination prefix length.• default—Default route to use when routing packets do not match a network or host in the routing table. This is equivalent to specifying the IP address 0.0.0.0/0. <p>match-type—Criteria that the destination must match:</p> <ul style="list-style-type: none">• exact—Exactly match the route's mask length.• longer—The route's mask length is greater than the specified mask length.• orlonger—The route's mask length is equal to or greater than the specified mask length.• through destination-prefix—The route matches the first prefix, the route matches the second prefix for the number of bits in the route, and the number of bits in the route is less than or equal to the number of bits in the second prefix.• upto prefix-length—The route's mask length falls between the two destination prefix lengths, inclusive.

Required Privilege routing—To view this statement in the configuration.
Level routing-control—To add this statement to the configuration.

Related Documentation • *Example: Configuring Class E Martian Addresses for Routing*

maximum-bandwidth (Routing Options)

Syntax	maximum-bandwidth <i>bps</i> ;
Hierarchy Level	<p>[edit dynamic-profiles <i>profile-name</i> routing-instances <i>instance-name</i> routing-options multicast interface <i>interface-name</i>],</p> <p>[edit dynamic-profiles <i>profile-name</i> routing-options multicast interface <i>interface-name</i>]</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast interface <i>interface-name</i>],</p> <p>[edit routing-options multicast interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>dynamic-profiles hierarchy level added in Junos OS Release 11.2.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Configure the multicast bandwidth for the interface.
Options	<p>bps—Bandwidth rate, in bits per second, for the multicast interface.</p> <p>Range: 0 through any amount of bandwidth</p>
Required Privilege	routing—To view this statement in the configuration.
Level	routing-control—To add this statement to the configuration.
Related Documentation	• <i>Example: Defining Interface Bandwidth Maximums</i>

maximum-paths

Syntax	<code>maximum-paths <i>path-limit</i> <log-interval <i>seconds</i>> <log-only threshold <i>value</i>>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options], [edit routing-options]
Release Information	Statement introduced in Junos OS Release 8.0. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure a limit for the number of routes installed in a routing table based upon the route path.



NOTE: The `maximum-paths` statement is similar to the `maximum-prefixes` statement. The `maximum-prefixes` statement limits the number of unique destinations in a routing instance. For example, suppose a routing instance has the following routes:

```
OSPF 10.10.10.0/24
ISIS 10.10.10.0/24
```

These are two routes, but only one destination (prefix). The `maximum-paths` limit applies the total number of routes (two). The `maximum-prefixes` limit applies to the total number of unique prefixes (one).

Options	<p><code>log-interval <i>seconds</i></code>—(Optional) Minimum time interval (in seconds) between log messages.</p> <p>Range: 5 through 86,400</p> <p><code>log-only</code>—(Optional) Sets the route limit as an advisory limit. An advisory limit triggers only a warning, and additional routes are not rejected.</p> <p><code><i>path-limit</i></code>—Maximum number of routes. If this limit is reached, a warning is triggered and additional routes are rejected.</p> <p>Range: 1 through 4,294,967,295 ($2^{32} - 1$)</p> <p>Default: No default</p> <p><code>threshold <i>value</i></code>—(Optional) Percentage of the maximum number of routes that starts triggering a warning. You can configure a percentage of the <code><i>path-limit</i></code> value that starts triggering the warnings.</p> <p>Range: 1 through 100</p>
----------------	---




NOTE: When the number of routes reaches the **threshold** value, routes are still installed into the routing table while warning messages are sent. When the number of routes reaches the *path-limit* value, then additional routes are rejected.

Required Privilege	routing—To view this statement in the configuration.
Level	routing-control—To add this statement to the configuration.

Related Documentation	<ul style="list-style-type: none">• <i>Limiting the Number of Paths and Prefixes Accepted from CE Routers in Layer 3 VPNs</i>
------------------------------	---

maximum-prefixes

Syntax	<code>maximum-prefixes <i>prefix-limit</i> <log-interval <i>seconds</i>> <log-only threshold <i>percentage</i>>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options], [edit routing-options]
Release Information	Statement introduced in Junos OS Release 8.0. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Configure a limit for the number of routes installed in a routing table based upon the route prefix.</p> <p>Using a prefix limit, you can curtail the number of prefixes received from a CE router in a VPN. Prefix limits apply only to dynamic routing protocols and are not applicable to static or interface routes.</p>
<div>  <p>NOTE: The <code>maximum-prefixes</code> statement is similar to the <code>maximum-paths</code> statement. The <code>maximum-prefixes</code> statement limits the number of unique destinations in a routing instance. For example, suppose a routing instance has the following routes:</p> <pre> OSPF 10.10.10.0/24 ISIS 10.10.10.0/24 </pre> <p>These are two routes, but only one destination (prefix). The <code>maximum-paths</code> limit applies the total number of routes (two). The <code>maximum-prefixes</code> limit applies to the total number of unique prefixes (one).</p> </div>	
Options	<p>log-interval <i>seconds</i>—(Optional) Minimum time interval (in seconds) between log messages.</p> <p>Range: 5 through 86,400</p> <p>log-only—(Optional) Sets the prefix limit as an advisory limit. An advisory limit triggers only a warning, and additional routes are not rejected.</p> <p><i>prefix-limit</i>—Maximum number of route prefixes. If this limit is reached, a warning is triggered and any additional routes are rejected.</p> <p>Range: 1 through 4,294,967,295</p> <p>Default: No default</p>

threshold value—(Optional) Percentage of the maximum number of prefixes that starts triggering a warning. You can configure a percentage of the **prefix-limit** value that starts triggering the warnings.

Range: 1 through 100



NOTE: When the number of routes reaches the threshold value, routes are still installed into the routing table while warning messages are sent. When the number of routes reaches the **prefix-limit** value, then additional routes are rejected.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [Limiting the Number of Paths and Prefixes Accepted from CE Routers in Layer 3 VPNs](#)

med-igp-update-interval

Syntax med-igp-update-interval *minutes*;

Hierarchy Level [edit routing-options]

Release Information Statement introduced in Junos OS Release 9.0
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure a timer for how long to delay updates for the multiple exit discriminator (MED) path attribute for BGP groups and peers configured with the **metric-out igp offset delay-med-update** statement. The timer delays MED updates for the interval configured unless the MED is lower than the previously advertised attribute or another attribute associated with the route has changed or if the BGP peer is responding to a refresh route request.

Options **minutes**—Interval to delay MED updates.
Range: 10 through 600
Default: 10 minutes

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.


Related Documentation

- [Example: Associating the MED Path Attribute with the IGP Metric and Delaying MED Updates on page 3694](#)
- [metric-out on page 4032](#)

metric (Aggregate, Generated, or Static Route)

Syntax	(metric metric2 metric3 metric4) <i>metric</i> <type <i>type</i> >;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options (aggregate generate static) (defaults route)], [edit routing-options (aggregate generate static) (defaults route)]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify the metric value for an aggregate, generated, or static route. You can specify up to four metric values, starting with metric (for the first metric value) and continuing with metric2 , metric3 , and metric4 .
Options	metric —Metric value. Range: 0 through 4,294,967,295 ($2^{32} - 1$) type type —(Optional) Type of route. When routes are exported to OSPF, type 1 routes are advertised in type 1 externals, and routes of any other type are advertised in type 2 externals. Note that if a qualified-next-hop metric value is configured, this value overrides the route metric. Range: 1 through 16
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• aggregate on page 3272• generate on page 3309• <i>static</i>• <i>Example: Conditionally Generating Static Routes</i>• <i>Example: Summarizing Static Routes Through Route Aggregation</i>• <i>Understanding Route Aggregation</i>

multicast (Routing Options)

Syntax	<pre> multicast { forwarding-cache { threshold suppress value <reuse value>; } interface interface-name { enable; } local-address address scope scope-name { interface [interface-names]; prefix destination-prefix; } ssm-groups { address; } } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit routing-options]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	<p>Configure generic multicast properties.</p>
<div>  <p>NOTE: You cannot apply a scoping policy to a specific routing instance. All scoping policies are applied to all routing instances. However, you can apply the <code>scope</code> statement to a specific routing instance.</p> </div>	
<p>The remaining statements are explained separately.</p>	
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Examples: Configuring Administrative Scoping</i> • Example: Configuring Source-Specific Multicast Groups with Any-Source Override on page 5537 • <i>Examples: Configuring the Multicast Forwarding Cache</i> • <i>Multicast Protocols Feature Guide for Routing Devices</i> • (indirect-next-hop on page 3315 no-indirect-next-hop)

no-qos-adjust

Syntax	no-qos-adjust;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast interface <i>interface-name</i> reverse-oif-mapping],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast interface <i>interface-name</i> reverse-oif-mapping],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast interface <i>interface-name</i> reverse-oif-mapping],</p> <p>[edit routing-options multicast interface <i>interface-name</i>],</p> <p>[edit routing-options multicast interface <i>interface-name</i> reverse-oif-mapping]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.5.</p> <p>Statement introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Statement added to [edit routing-instances <i>routing-instance-name</i> routing-options multicast interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast interface <i>interface-name</i>], and [edit routing-options multicast interface <i>interface-name</i>] hierarchy levels in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Disable hierarchical bandwidth adjustment for all subscriber interfaces that are identified by their MLD or IGMP request from a specific multicast interface.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

options (Routing Options)

Syntax	<pre>options { syslog (level <i>level</i> upto level <i>level</i>); }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit routing-options]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Configure the types of system logging messages sent about the routing protocols process to the system message logging file. These messages are also displayed on the system console. You can log messages at a particular level, or up to and including a particular level.</p>
Options	<p>level <i>level</i>—Severity of the message. It can be one or more of the following levels, in order of decreasing urgency:</p> <ul style="list-style-type: none"> • alert—Conditions that should be corrected immediately, such as a corrupted system database. • critical—Critical conditions, such as hard drive errors. • debug—Software debugging messages. • emergency—Panic or other conditions that cause the system to become unusable. • error—Standard error conditions. • info—Informational messages. • notice—Conditions that are not error conditions, but might warrant special handling. • warning—System warning messages. <p>upto level <i>level</i>—Log all messages up to a particular level.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • syslog in the <i>Junos OS Administration Library for Routing Devices</i>

pim-to-igmp-proxy

Syntax	<code>pim-to-igmp-proxy { upstream-interface [interface-names]; }</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast], [edit logical-systems <i>logical-system-name</i> routing-options multicast], [edit routing-instances <i>routing-instance-name</i> routing-options multicast], [edit routing-options multicast]
Release Information	Statement introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 9.6 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Configure the rendezvous point (RP) routing device that resides between a customer edge-facing Protocol Independent Multicast (PIM) domain and a core-facing PIM domain to translate PIM join or prune messages into corresponding Internet Group Management Protocol (IGMP) report or leave messages. The routing device then transmits the report or leave messages by proxying them to one or two upstream interfaces that you configure on the RP routing device. Including the pim-to-igmp-proxy statement enables you to use IGMP to forward IPv4 multicast traffic across the PIM sparse mode domains.</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Configuring PIM-to-IGMP Message Translation

pim-to-mld-proxy

Syntax	<code>pim-to-mld-proxy { upstream-interface [interface-names]; }</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast], [edit logical-systems <i>logical-system-name</i> routing-options multicast], [edit routing-instances <i>routing-instance-name</i> routing-options multicast], [edit routing-options multicast]
Release Information	Statement introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 9.6 for EX Series switches. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Configure the rendezvous point (RP) routing device that resides between a customer edge-facing Protocol Independent Multicast (PIM) domain and a core-facing PIM domain to translate PIM join or prune messages into corresponding Multicast Listener Discovery (MLD) report or leave messages. The routing device then transmits the report or leave messages by proxying them to one or two upstream interfaces that you configure on the RP routing device. Including the pim-to-mld-proxy statement enables you to use MLD to forward IPv6 multicast traffic across the PIM sparse mode domains. The remaining statement is explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> Configuring PIM-to-MLD Message Translation

policy (Aggregate and Generated Routes)

Syntax	<code>policy <i>policy-name</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (aggregate generate) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options (aggregate generate) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options (aggregate generate) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate) (defaults route)],</p> <p>[edit routing-options (aggregate generate) (defaults route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> (aggregate generate) (defaults route)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Associate a routing policy when configuring an aggregate or generated route's destination prefix in the routes part of the aggregate or generate statement. This provides the equivalent of an import routing policy filter for the destination prefix. That is, each potential contributor to an aggregate route, along with any aggregate options, is passed through the policy filter. The policy then can accept or reject the route as a contributor to the aggregate route.</p> <p>If the contributor is accepted, the policy can modify the default preferences. The contributor with the numerically smallest prefix becomes the most preferred, or <i>primary</i>, contributor. A rejected contributor still can contribute to a less specific aggregate route. If you do not specify a policy filter, all candidate routes contribute to an aggregate route.</p> <p>The following algorithm is used to compare two generated contributing routes in order to determine which one is the primary or preferred contributor:</p> <ol style="list-style-type: none"> 1. Compare the protocol's preference of the contributing routes. The lower the preference, the better the route. This is similar to the comparison that is done while determining the best route for the routing table. 2. Compare the protocol's preference2 of the contributing routes. The lower preference2 value is better. If only one route has preference2, then this route is preferred. 3. The preference values are the same. Proceed with a numerical comparison of the prefixes' values. <ol style="list-style-type: none"> a. The primary contributor is the numerically smallest prefix value.

- b. If the two prefixes are numerically equal, the primary contributor is the route that has the smallest prefix length value.

At this point, the two routes are the same. The primary contributor does not change. An additional next hop is available for the existing primary contributor.

A rejected contributor still can contribute to less specific generated route. If you do not specify a policy filter, all candidate routes contribute to a generated route.

Options	<i>policy-name</i> —Name of a routing policy.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Summarizing Static Routes Through Route Aggregation</i> • <i>Example: Conditionally Generating Static Routes</i> • aggregate on page 3272 • generate on page 3309

policy (Flow Maps)

Syntax	<code>policy [<i>policy-names</i>];</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast flow-map <i>flow-map-name</i>], [edit logical-systems <i>logical-system-name</i> routing-options multicast flow-map <i>flow-map-name</i>], [edit routing-instances <i>routing-instance-name</i> routing-options multicast flow-map <i>flow-map-name</i>], [edit routing-options multicast flow-map <i>flow-map-name</i>]
Release Information	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure a flow map policy.
Options	<i>policy-names</i> —Name of one or more policies for flow mapping.
Required Privilege Level	routing—To view this statement in the configuration.

policy-options

```
Syntax  policy-options
        application-maps application-map-name {
            application application-name {
                code-points [ aliases ] [ bit-patterns ];
            }
        }
        policy-statement policy-name {
            term term-name {
                from {
                    family family-name;
                    match-conditions;
                    policy subroutine-policy-name;
                    prefix-list prefix-list-name;
                    prefix-list-filter prefix-list-name match-type <actions>;
                    route-filter destination-prefix match-type <actions>;
                    source-address-filter source-prefix match-type <actions>;
                }
                to {
                    match-conditions;
                    policy subroutine-policy-name;
                }
                then actions;
            }
        }
```

Hierarchy Level [edit]

Release Information Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 12.1 for the EX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure options such as application maps for DCBX application protocol exchange and policy statements.

Required Privilege Level storage—To view this statement in the configuration.
storage-control—To add this statement to the configuration.

Related Documentation

- [Defining an Application for DCBX Application Protocol TLV Exchange on page 6532](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 6535](#)
- [Example: Configuring DCBX to Support an iSCSI Application](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 6528](#)
- [Understanding DCBX Application Protocol TLV Exchange on EX Series Switches](#)

policy-statement

Syntax	<pre> policy-statement <i>policy-name</i> { term <i>term-name</i> { from { family <i>family-name</i>; match-conditions; policy <i>subroutine-policy-name</i>; prefix-list <i>prefix-list-name</i>; prefix-list-filter <i>prefix-list-name</i> <i>match-type</i> <<i>actions</i>>; protocol <i>protocol-name</i>; route-filter <i>destination-prefix</i> <i>match-type</i> <<i>actions</i>>; source-address-filter <i>source-prefix</i> <i>match-type</i> <<i>actions</i>>; } to { match-conditions; policy <i>subroutine-policy-name</i>; } then <i>actions</i>; } then { no-entropy-label-capability; } } </pre>
Hierarchy Level	[edit dynamic policy-options], [edit logical-systems <i>logical-system-name</i> policy-options], [edit policy-options]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for configuration in the dynamic database introduced in Junos OS Release 9.5.</p> <p>Support for configuration in the dynamic database introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>inet-mdt option introduced in Junos OS Release 10.0R2.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>route-target option introduced in Junos OS Release 12.2.</p> <p>Statement introduced in Junos OS 14.1X53-D20 for the OCX Series.</p> <p>protocol and traffic-engineering options introduced in Junos OS Release 14.2.</p> <p>no-entropy-label-capability option introduced in Junos OS Release 15.1.</p>
Description	<p>Define a routing policy, including subroutine policies.</p> <p>A <i>term</i> is a named structure in which match conditions and actions are defined. Routing policies are made up of one or more terms. Each routing policy term is identified by a term name. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose the entire name in double quotation marks.</p> <p>Each term contains a set of match conditions and a set of actions:</p>

- Match conditions are criteria that a route must match before the actions can be applied. If a route matches all criteria, one or more actions are applied to the route.
- Actions specify whether to accept or reject the route, control how a series of policies are evaluated, and manipulate the characteristics associated with a route.

Generally, a router compares a route against the match conditions of each term in a routing policy, starting with the first and moving through the terms in the order in which they are defined, until a match is made and an explicitly configured or default action of **accept** or **reject** is taken. If none of the terms in the policy match the route, the router compares the route against the next policy, and so on, until either an action is taken or the default policy is evaluated.

If none of the match conditions of each term evaluates to true, the final action is executed. The final action is defined in an unnamed term. Additionally, you can define a default action (either **accept** or **reject**) that overrides any action intrinsic to the protocol.

The order of match conditions in a term is not relevant, because a route must match all match conditions in a term for an action to be taken.

To list the routing policies under the **[edit policy-options]** hierarchy level by **policy-statement *policy-name*** in alphabetical order, enter the **show policy-options** configuration command.

Options **actions**—(Optional) One or more actions to take if the conditions match. The actions are described in *Configuring Flow Control Actions*.

family **family-name**—(Optional) Specify an address family protocol. Specify **inet** for IPv4. Specify **inet6** for 128-bit IPv6, and to enable interpretation of IPv6 router filter addresses. For IS-IS traffic, specify **iso**. For IPv4 multicast VPN traffic, specify **inet-mvpn**. For IPv6 multicast VPN traffic, specify **inet6-mvpn**. For multicast-distribution-tree (MDT) IPv4 traffic, specify **inet-mdt**. For BGP route target VPN traffic, specify **route-target**. For traffic engineering, specify **traffic-engineering**.



NOTE: When **family** is not specified, the routing device or routing instance uses the address family or families carried by BGP. If multiprotocol BGP (MP-BGP) is enabled, the policy defaults to the protocol family or families carried in the network layer reachability information (NLRI) as configured in the **family** statement for BGP. If MP-BGP is not enabled, the policy uses the default BGP address family unicast IPv4.

from—(Optional) Match a route based on its source address.

match-conditions—(Optional in **from** statement; required in **to** statement) One or more conditions to use to make a match. The qualifiers are described in *Routing Policy Match Conditions*.

policy **subroutine-policy-name**—Use another policy as a match condition within this policy. The name identifying the subroutine policy can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" "). Policy names cannot take the form **__.*-internal__**, as this form is reserved. For information about how to configure subroutines, see *Understanding Policy Subroutines in Routing Policy Match Conditions*.

no-entropy-label-capability—(Optional) Disable the entropy label capability advertisement at egress or transit routes specified in the policy.

policy **subroutine-policy-name**—Use another policy as a match condition within this policy. The name identifying the subroutine policy can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" "). Policy names cannot take the form **__.*-internal__**, as this form is reserved. For information about how to configure subroutines, see *Understanding Policy Subroutines in Routing Policy Match Conditions*.

policy-name—Name that identifies the policy. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" ").

prefix-list **prefix-list-name**—Name of a list of IPv4 or IPv6 prefixes.

prefix-list-filter **prefix-list-name**—Name of a prefix list to evaluate using qualifiers; **match-type** is the type of match (see *Configuring Prefix List Filters*), and **actions** is the action to take if the prefixes match.

protocol *protocol-name*—Name of the protocol used to control traffic engineering database import at the originating point.

route-filter *destination-prefix match-type <actions>*—(Optional) List of routes on which to perform an immediate match; *destination-prefix* is the IPv4 or IPv6 route prefix to match, *match-type* is the type of match (see *Configuring Route Lists*), and *actions* is the action to take if the *destination-prefix* matches.

source-address-filter *source-prefix match-type <actions>*—(Optional) Unicast source addresses in multiprotocol BGP (MBGP) and Multicast Source Discovery Protocol (MSDP) environments on which to perform an immediate match. *source-prefix* is the IPv4 or IPv6 route prefix to match, *match-type* is the type of match (see *Configuring Route Lists*), and *actions* is the action to take if the *source-prefix* matches.

term *term-name*—Name that identifies the term. The term name must be unique in the policy. It can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. To include spaces in the name, enclose the entire name in quotation marks (" "). A policy statement can include multiple terms. We recommend that you name all terms. However, you do have the option to include an unnamed term which must be the final term in the policy. To configure an unnamed term, omit the **term** statement when defining match conditions and actions.

to—(Optional) Match a route based on its destination address or the protocols into which the route is being advertised.

then—(Optional) Actions to take on matching routes. The actions are described in *Configuring Flow Control Actions* and *Configuring Actions That Manipulate Route Characteristics*.

Required Privilege	routing—To view this statement in the configuration.
Level	routing-control—To add this statement to the configuration.

Related Documentation	<ul style="list-style-type: none">• <i>dynamic-db</i>
------------------------------	---

ppm

Syntax	<pre>ppm { no-delegate-processing; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-options]
Release Information	<p>Statement introduced in Junos OS Release 9.4.</p> <p>Statement introduced in Junos OS Release 10.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>(M120, M320, MX Series, T Series, TX Matrix routers, M7i and M10i routers with Enhanced CFEB [CFEB-E], EX Series switches, and QFX Series only) Disable distributed periodic packet management (PPM) to the Packet Forwarding Engine (on routers), to access ports (on EX3200 and EX4200 switches, and QFX Series), or to line cards (on EX6200 and EX8200 switches).</p> <p>After you disable PPM, PPM processing continues to run on the Routing Engine.</p> <p>In Junos OS Release 8.2, PPM was moved from the Routing Engine to the Packet Forwarding Engine, access ports, or line cards. The no-delegate-processing statement disables the default behavior and restores the legacy behavior.</p>
Default	Distributed PPM processing is enabled for all protocols that use PPM.
Options	no-delegate-processing —Disable PPM to the Packet Forwarding Engine, access ports, or line cards. Distributed PPM is enabled by default.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <i>Ensuring That Distributed ppm Is Not Disabled</i>

ppm (Ethernet Switching)

Syntax	<pre>ppm { centralized; }</pre>
Hierarchy Level	[edit protocols lacp]
Release Information	<p>Statement introduced in Junos OS Release 9.4 for MX Series routers.</p> <p>Statement introduced in Junos OS Release 10.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.1 for T Series devices.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Configure PPM processing options for Link Aggregation Control Protocol (LACP) packets.</p> <p>This command configures the PPM processing options for LACP packets only. You can disable distributed PPM processing for all packets that use PPM and run all PPM processing on the Routing Engine by configuring the no-delegate-processing configuration statement in the [edit routing-options ppm] statement hierarchy.</p>
Default	Distributed PPM processing is enabled for all packets that use PPM.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Distributed Periodic Packet Management on an EX Series Switch (CLI Procedure)</i>• Configuring Distributed Periodic Packet Management on page 3234

preference (Routing Options)

Syntax	<pre> preference { metric-value; <type metric_type> } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (aggregate generate static) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options (aggregate generate static) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options (aggregate generate static) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</p> <p>[edit routing-options (aggregate generate static) (defaults route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Preference value for a static, aggregate, or generated route. You also can specify a secondary preference value, as well as color values, which are even finer-grained preference values.</p> <p>You can specify a primary route preference (by including the preference statement in the configuration), and a secondary preference that is used as a tiebreaker (by including the preference2 statement). You can also mark route preferences with additional route tiebreaker information by specifying a color and a tiebreaker color (by including the color and color2 statements in the configuration).</p> <p>If the Junos OS routing table contains a dynamic route to a destination that has a better (lower) preference value than the static, aggregate, or generated route, the dynamic route is chosen as the active route and is installed in the forwarding table.</p>
Options	<p>metric_value—The metric value for an aggregate, a generated, or a static route to determine the best route among multiple routes to a destination</p> <p>Range: 0 through 4,294,967,295 ($2^{32} - 1$)</p> <p>Default: 5 (for static routes), 130 (for aggregate and generated routes)</p> <p>type metric_type—(Optional) External metric type for routes exported by OSPF. When routes are exported to OSPF, type 1 routes are advertised in type 1 externals, and routes of any other type are advertised in type 2 externals. Note that if a qualified-next-hop metric value is configured, this value overrides the route metric.</p>

Range: 1 through 16

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Examples: Configuring Static Routes*
- *Example: Summarizing Static Routes Through Route Aggregation*
- *Example: Conditionally Generating Static Routes*
- [aggregate on page 3272](#)
- [generate on page 3309](#)
- *static*
- *color*

prefix

Syntax `prefix destination-prefix;`

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast [scope scope-name](#)],
[edit logical-systems *logical-system-name* routing-options multicast [scope scope-name](#)],
[edit routing-instances *routing-instance-name* routing-options multicast [scope scope-name](#)],
[edit routing-options multicast [scope scope-name](#)]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure the prefix for multicast scopes.

Options *destination-prefix*—Address range for the multicast scope.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Examples: Configuring Administrative Scoping*
- *Example: Creating a Named Scope for Multicast Scoping*
- *multicast*

protocols

```

Syntax protocols {
    bgp {
        ... bgp-configuration ...
    }
    isis {
        ... isis-configuration ...
    }
    ldp {
        ... ldp-configuration ...
    }
    mpls {
        ... mpls-configuration ...
    }
    msdp {
        ... msdp-configuration ...
    }
    mstp {
        ... mstp-configuration ...
    }
    ospf {
        domain-id domain-id;
        domain-vpn-tag number;
        route-type-community (iana | vendor);
        traffic-engineering {
            <advertise-unnumbered-interfaces>;
            <credibility-protocol-preference>;
            ignore-lsp-metrics;
            multicast-rpf-routes;
            no-topology;
            shortcuts {
                lsp-metric-into-summary;
            }
        }
        ... ospf-configuration ...
    }
    ospf3 {
        domain-id domain-id;
        domain-vpn-tag number;
        route-type-community (iana | vendor);
        traffic-engineering {
            <advertise-unnumbered-interfaces>;
            <credibility-protocol-preference>;
            ignore-lsp-metrics;
            multicast-rpf-routes;
            no-topology;
            shortcuts {
                lsp-metric-into-summary;
            }
        }
        ... ospf3-configuration ...
    }
    pim {

```

```
    ...pim-configuration ...
  }
  rip {
    ...rip-configuration ...
  }
  ripng {
    ...ripng-configuration ...
  }
  rstp {
    rstp-configuration;
  }
  rsvp{
    ...rsvp-configuration ...
  }
  vstp {
    vstp configuration;
  }
  vpls {
    vpls configuration;
  }
}
```

Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>], [edit routing-instances <i>routing-instance-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Support for RIPng introduced in Junos OS Release 9.0. Statement introduced in Junos OS Release 11.1 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. mpls and rsvp options added in Junos OS Release 15.1. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify the protocol for a routing instance. You can configure multiple instances of many protocol types. Not all protocols are supported on the switches. See the switch CLI.

- Options**
- bgp**—Specify BGP as the protocol for a routing instance.
 - isis**—Specify IS-IS as the protocol for a routing instance.
 - ldp**—Specify LDP as the protocol for a routing instance or for a virtual router instance.
 - l2vpn**—Specify Layer 2 VPN as the protocol for a routing instance.
 - mpls**—Specify MPLS as the protocol for a routing instance.
 - msdp**—Specify the Multicast Source Discovery Protocol (MSDP) for a routing instance.
 - mstp**—Specify the Multiple Spanning Tree Protocol (MSTP) for a virtual switch routing instance.
 - ospf**—Specify OSPF as the protocol for a routing instance.
 - ospf3**—Specify OSPF version 3 (OSPFv3) as the protocol for a routing instance.



NOTE: OSPFv3 supports the **no-forwarding**, **virtual-router**, and **vrf** routing instance types only.

- pim**—Specify the Protocol Independent Multicast (PIM) protocol for a routing instance.
- rip**—Specify RIP as the protocol for a routing instance.
- ripng**—Specify RIP next generation (RIPng) as the protocol for a routing instance.
- rstp**—Specify the Rapid Spanning Tree Protocol (RSTP) for a virtual switch routing instance.
- rsvp**—Specify the RSVP for a routing instance.
- vstp**—Specify the VLAN Spanning Tree Protocol (VSTP) for a virtual switch routing instance.
- vpls**—Specify VPLS as the protocol for a routing instance.

- | | |
|---------------------------------|---|
| Required Privilege Level | routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • <i>Example: Configuring Multiple Routing Instances of OSPF</i> |

qualified-next-hop (Static Routes)

Syntax `qualified-next-hop (address | interface-name) {
 bfd-liveness-detection {
 authentication {
 algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 | meticulous-keyed-sha-1 |
 simple-password);
 key-chain key-chain-name;
 loose-check;
 }
 detection-time {
 threshold milliseconds;
 }
 holddown-interval milliseconds;
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 multiplier number;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 version (1 | automatic);
 }
 interface interface-name;
 metric metric;
 preference preference;
}`

Hierarchy Level `[edit logical-systems logical-system-name routing-instances routing-instance-name
 routing-options static route destination-prefix],
 [edit logical-systems logical-system-name routing-options rib inet6.0 static route
 destination-prefix],
 [edit logical-systems logical-system-name routing-options static route destination-prefix],
 [edit routing-instances routing-instance-name routing-options static route destination-prefix],
 [edit routing-options rib inet6.0 static route destination-prefix],
 [edit routing-options static route destination-prefix]`

Release Information Statement introduced before Junos OS Release 7.4.
 Statement introduced in Junos OS Release 9.0 for EX Series switches.
 Statement introduced in Junos OS Release 11.3 for the QFX Series.
 Statement introduced in Junos OS Release 12.3 for ACX Series routers.
 Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure a static route with multiple possible next hops, each of which can have its own preference value, IGP metric that is used when the route is exported into an IGP, and Bidirectional Forwarding Detection (BFD) settings. If multiple links are operational, the one with the most preferred next hop is used. The most preferred next hop is the one with the lowest preference value.

Options *address*—IPv4, IPv6, or ISO network address of the next hop.

interface-name—Name of the interface on which to configure an independent metric or preference for a static route. To configure an unnumbered interface as the next-hop interface for a static route, specify **qualified-next-hop interface-name**, where **interface-name** is the name of the IPv4 or IPv6 unnumbered interface.



NOTE: For an Ethernet interface to be configured as the qualified next hop for a static route, it must be an unnumbered interface.

To configure an Ethernet interface as an unnumbered interface, configure the **unnumbered-address <interface-name>** statement at the [edit interfaces <interface-name> unit <logical-unit-number> family <family-name>] hierarchy level as described in *Configuring an Unnumbered Interface*.

The remaining statements are explained separately.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
---------------------------------	---

Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Static Route Preferences and Qualified Next Hops • Example: Enabling BFD on Qualified Next Hops in Static Routes for Route Selection on page 3246
------------------------------	---


readvertise

Syntax	(readvertise no-readvertise);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> static (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options static (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> static (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-options static (defaults route)], [edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> static (defaults route)], [edit routing-instances <i>routing-instance-name</i> routing-options static (defaults route)], [edit routing-options rib <i>routing-table-name</i> static (defaults route)], [edit routing-options static (defaults route)]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure whether static routes are eligible to be readvertised by routing protocols:
Default	Static routes are eligible to be readvertised (that is, exported from the routing table into dynamic routing protocols) if a policy to do so is configured. To mark an IPv4 static route as being ineligible for readvertisement, include the no-readvertise statement.
Options	readvertise —Readvertise static routes. Include the readvertise statement when configuring an individual route in the route portion of the static statement to override a no-readvertise option specified in the defaults portion of the statement. no-readvertise —Mark a static route as being ineligible for readvertisement. Include the no-readvertise option when configuring the route.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Preventing a Static Route from Being Readvertised</i>• <i>Understanding Static Route Importation in Routing and Forwarding Tables</i>• <i>static</i>

redundant-sources

Syntax	<code>redundant-sources [<i>addresses</i>];</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast flow-map <i>flow-map-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast flow-map <i>flow-map-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast flow-map <i>flow-map-name</i>],</p> <p>[edit routing-options multicast flow-map <i>flow-map-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Configure a list of redundant sources for multicast flows defined by a flow map.
Options	<i>addresses</i> —List of IPv4 or IPv6 addresses for use as redundant (backup) sources for multicast flows defined by a flow map.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring a Multicast Flow Map</i>

resolution

Syntax	<pre> resolution { rib routing-table-name { import [policy-names]; resolution-ribs [routing-table-names]; } } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit routing-options]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Configure the router to perform custom route resolution on protocol next hops of routes in a certain routing table. The protocol next hop is used to determine the forwarding next-hop.</p> <p>For example, you might want to direct inet.2 route resolution to use topology routing tables :red.inet.0 and :blue.inet.0 for protocol next-hop IP address lookups. Or you might want to direct bgp.l3vpn.0 to use the information in inet.0 to resolve routes, thus overriding the default behavior, which is to use inet.3.</p> <p>You can specify up to two routing tables in the resolution-ribs statement. The route resolution scheme first checks the first-listed routing table for the protocol next-hop address. If the address is found, it uses this entry. If it is not found, the resolution scheme checks second-listed routing table. Hence, only one routing table is used for each protocol nexthop address. For example, if you configure resolution rib bgp.l3vpn.0 resolution-ribs [inet.0 inet.3], inet.0 is checked first and then inet.3 is checked.</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 20px;"> <p> NOTE: Customizing route resolution might cause the routing protocol process (rpd) to consume more memory resources than it ordinarily would. When you customize route resolution, we recommend that you check the memory resources by running the show system processes and the show task memory commands. For more information, see <i>Routing Protocol Process Overview</i>.</p> </div> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

- Related Documentation**
- *Example: Configuring Route Resolution on PE Routers*
 - *Example: Configuring Route Resolution on Route Reflectors*
 - *Understanding Multitopology Routing in Conjunction with PIM*
 - *Example: Configuring Multitopology Routing to Provide Redundancy for Multicast Traffic over Separate Network Paths*

resolution-ribs

Syntax	<code>resolution-ribs [<i>routing-table-names</i>];</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options resolution <i>rib</i>], [edit logical-systems <i>logical-system-name</i> routing-options resolution <i>rib</i>], [edit routing-instances <i>routing-instance-name</i> routing-options resolution <i>rib</i>], [edit routing-options resolution <i>rib</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify one or more routing tables to use for route resolution. This statement enables you to override the default routing tables that Junos OS uses for route resolution. For example, suppose that the resolution routing table is inet.3 , but you want to allow fallback resolution through inet.0 . One example use case is overriding the bgp.rtarget.0 (family route-target) routing table resolution from using only inet.3 to using both inet.3 and inet.0 .
Options	<i>routing-table-names</i> —Name of one or more routing tables.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Route Resolution on PE Routers</i> • <i>Example: Configuring Multitopology Routing to Provide Redundancy for Multicast Traffic over Separate Network Paths</i> • <i>Understanding Multitopology Routing in Conjunction with PIM</i>

resolve

Syntax	resolve;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> static (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options static (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> static (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-options static (defaults route)], [edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> static (defaults route)], [edit routing-instances <i>routing-instance-name</i> routing-options static (defaults route)], [edit routing-options rib <i>routing-table-name</i> static (defaults route)], [edit routing-options static (defaults route)]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Statically configure routes to be resolved to a next hop that is not directly connected. The route is resolved through the inet.0 and inet.3 routing tables.
Default	Static routes can point only to a directly connected next hop.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>static</i>

restart-duration (Routing Options)

Syntax	<code>restart-duration <i>seconds</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options graceful-restart], [edit logical-systems <i>logical-system-name</i> routing-options graceful-restart], [edit routing-instances <i>routing-instance-name</i> routing-options graceful-restart], [edit routing-options graceful-restart]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure the restart timer for graceful restart.
Options	<i>seconds</i> —Configure the time period for the restart to last. Range: 120 through 900 seconds Default: 300 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Junos OS High Availability Library for Routing Devices</i>

retain

Syntax	(no-retain retain);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> static (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options static (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> static (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-options static (defaults route)], [edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> static (defaults route)], [edit routing-instances <i>routing-instance-name</i> routing-options static (defaults route)], [edit routing-options rib <i>routing-table-name</i> static (defaults route)], [edit routing-options static (defaults route)]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure statically configured routes to be deleted from or retained in the forwarding table when the routing protocol process shuts down normally:
Default	Statically configured routes are deleted from the forwarding table when the routing protocol process shuts down normally.
Options	no-retain —Delete statically configured routes from the forwarding table when the routing protocol process shuts down normally. To explicitly specify that routes be deleted from the forwarding table, include the no-retain statement. Include this statement when configuring an individual route in the route portion of the static statement to override a retain option specified in the defaults portion of the statement. retain —Have a static route remain in the forwarding table when the routing protocol process shuts down normally. Doing this greatly reduces the time required to restart a system that has a large number of routes in its routing table.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Examples: Configuring Static Routes</i>• <i>static</i>

reverse-oif-mapping

Syntax	reverse-oif-mapping { no-qos-adjust; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast <i>interface interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-options multicast <i>interface interface-name</i>], [edit routing-instances <i>routing-instance-name</i> routing-options multicast <i>interface interface-name</i>], [edit routing-options multicast <i>interface interface-name</i>]
Release Information	Statement introduced in Junos OS Release 9.2. Statement introduced in Junos OS Release 9.2 for EX Series switches. The no-qos-adjust statement added in Junos OS Release 9.5. The no-qos-adjust statement introduced in Junos OS Release 9.5 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Enable the routing device to identify a subscriber VLAN or interface based on an IGMP or MLD request it receives over the multicast VLAN. The remaining statement is explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

rpf-check-policy (Routing Options RPF)

Syntax	<code>rpf-check-policy [<i>policy-names</i>];</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast], [edit logical-systems <i>logical-system-name</i> routing-options multicast], [edit routing-instances <i>routing-instance-name</i> routing-options multicast], [edit routing-options multicast]
Release Information	Statement introduced in Junos OS Release 8.1. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Apply policies for disabling RPF checks on arriving multicast packets. The policies must be correctly configured.
Options	<i>policy-names</i> —Name of one or more multicast RPF check policies.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring RPF Policies</i>

rib (General)

```
Syntax  rib routing-table-name {
        aggregate {
            defaults {
                ... aggregate-options ...
            }
            route destination-prefix {
                policy policy-name;
                ... aggregate-options ...
            }
        }
        generate {
            defaults {
                generate-options;
            }
            route destination-prefix {
                policy policy-name;
                generate-options;
            }
        }
        martians {
            destination-prefix match-type <allow>;
        }
    }
    static {
        defaults {
            static-options;
        }
        rib-group group-name;
        route destination-prefix {
            next-hop;
            static-options;
        }
    }
}
```

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Create a routing table.

Explicitly creating a routing table with ***routing-table-name*** is optional if you are not adding any static, martian, aggregate, or generated routes to the routing table and if you also are creating a routing table group.



NOTE: The IPv4 multicast routing table (`inet.1`) and the IPv6 multicast routing table (`inet6.1`) are not supported for this statement.

Default If you do not specify a routing table name with the ***routing-table-name*** option, the software uses the default routing tables, which are `inet.0` for unicast routes and `inet.1` for the multicast cache.

Options ***routing-table-name***—Name of the routing table, in the following format:
protocol [.identifier].

In a routing instance, the routing table name must include the routing instance name.

For example, if the routing instance name is `link0`, the routing table name might be `link0.inet6.0`.

- ***protocol*** is the protocol family. It can be `inet6` for the IPv6 family, `inet` for the IPv4 family, `iso` for the ISO protocol family, or ***instance-name.iso.0*** for an ISO routing instance.
- ***identifier*** is a positive integer that specifies the instance of the routing table.

Default: `inet.0`

The remaining statements are explained separately.

Required Privilege Level `routing`—To view this statement in the configuration.
`routing-control`—To add this statement to the configuration.

Related Documentation

- *Example: Creating Routing Tables*
- *passive*

rib (Route Resolution)

Syntax	<pre> rib <i>routing-table-name</i> { import [<i>policy-names</i>]; resolution-ribs [<i>routing-table-names</i>]; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options resolution],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options resolution],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options resolution],</p> <p>[edit routing-options resolution]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Specify a routing table name for route resolution.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Route Resolution on PE Routers</i>

rib-group (Routing Options)

Syntax	<code>rib-group group-name;</code>
Hierarchy Level	<code>[edit logical-systems logical-system-name routing-instances routing-instance-name routing-options interface-routes],</code> <code>[edit logical-systems logical-system-name routing-options interface-routes],</code> <code>[edit logical-systems logical-system-name routing-options rib routing-table-name static],</code> <code>[edit logical-systems logical-system-name routing-options static],</code> <code>[edit routing-instances routing-instance-name routing-options interface-routes],</code> <code>[edit routing-options interface-routes],</code> <code>[edit routing-options rib routing-table-name static],</code> <code>[edit routing-options static]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure which routing table groups interface routes are imported into.
Options	group-name —Name of the routing table group. The name must start with a letter and can include letters, numbers, and hyphens. It generally does not make sense to specify more than a single routing table group.
Required Privilege Level	<code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Importing Direct and Static Routes Into a Routing Instance</i>• <i>Example: Exporting Specific Routes from One Routing Table Into Another Routing Table</i>• interface-routes on page 3322• rib-groups on page 3365

rib-groups

Syntax	<pre> rib-groups { group-name { export-rib group-name; import-policy [policy-names]; import-rib [group-names]; } } </pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-options]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Group one or more routing tables to form a routing table group. A routing protocol can import routes into all the routing tables in the group and can export routes from a single routing table.</p> <p>Each routing table group must contain one or more routing tables that Junos OS uses when importing routes (specified in the import-rib statement) and optionally can contain one routing table group that Junos OS uses when exporting routes to the routing protocols (specified in the export-rib statement).</p> <p>The first routing table you specify is the <i>primary routing table</i>, and any additional routing tables are the <i>secondary routing tables</i>.</p> <p>The primary routing table determines the address family of the routing table group. To configure an IP version 4 (IPv4) routing table group, specify inet.0 as the primary routing table. To configure an IP version 6 (IPv6) routing table group, specify inet6.0 as the primary routing table. If you configure an IPv6 routing table group, the primary and all secondary routing tables must be IPv6 routing tables (inet6.x).</p> <p>In Junos OS Release 9.5 and later, you can include both IPv4 and IPv6 routing tables in an IPv4 import routing table group using the import-rib statement. In releases prior to Junos OS Release 9.5, you can only include either IPv4 or IPv6 routing tables in the same import-rib statement. The ability to configure an import routing table group with both IPv4 and IPv6 routing tables enables you, for example, to populate the inet6.3 routing table with IPv6 addresses that are compatible with IPv4. Specify inet.0 as the primary routing table, and specify inet6.3 as a secondary routing table.</p>



NOTE: On EX Series switches, only dynamically learned routes can be imported from one routing table group to another.



NOTE: If you configure an import routing table group that includes both IPv4 and IPv6 routing tables, any corresponding export routing table group must include only IPv4 routing tables.

If you have configured a routing table, configure the OSPF primary instance at the **[edit protocols ospf]** hierarchy level with the statements needed for your network so that routes are installed in **inet.0** and in the forwarding table. Make sure to include the routing table group. For more information, see *Example: Configuring Multiple Routing Instances of OSPF*.

After specifying the routing table from which to import routes, you can apply one or more policies to control which routes are installed in the routing table group. To apply a policy to routes being imported into the routing table group, include the **import-policy** statement.

Options *group-name*—Name of the routing table group. The name must start with a letter and can include letters, numbers, and hyphens.


The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.


Related Documentation

- *Example: Exporting Specific Routes from One Routing Table Into Another Routing Table*
- [rib-group on page 3364](#)

route-distinguisher-id

Syntax	<code>route-distinguisher-id <i>ip-address</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	<p>Automatically assign a route distinguisher to the routing instance.</p> <p>If you configure the route-distinguisher statement in addition to the route-distinguisher-id statement, the value configured for route-distinguisher supersedes the value generated from route-distinguisher-id.</p>
<div style="display: flex; align-items: flex-start;"> <div style="flex: 1; text-align: center; margin-right: 10px;">  </div> <div> <p>NOTE: To avoid a conflict in the two route distinguisher values, it is recommended to ensure that the first half of the route distinguisher obtained by configuring the route-distinguisher statement is different from the first half of the route distinguisher obtained by configuring the route-distinguisher-id statement.</p> </div> </div>	
Options	<i>ip-address</i> —Address for routing instance.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring BGP Route Target Filtering for VPNs</i> • <i>Configuring Routing Instances on PE Routers in VPNs</i>

route-record

Syntax	route-record;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Export the AS path and routing information to the traffic sampling process.</p> <p>Before you can perform flow aggregation, the routing protocol process must export the AS path and routing information to the sampling process.</p> <div><div></div><div><p>NOTE: Starting with Junos OS Release 15.1, when you commit a minor configuration change, the routing protocol process sends only AS paths that are active routes to the FPCs. Not all known AS paths are sent to the FPC, thereby considerably reducing the memory and CPU usage, resulting in a faster route record database update.</p></div></div>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Enabling Flow Aggregation</i>• <i>Junos OS Services Interfaces Library for Routing Devices</i>

router-id

Syntax	<code>router-id address;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit routing-options]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Specify the routing device's IP address.</p> <p>The router identifier is used by BGP and OSPF to identify the routing device from which a packet originated. The router identifier usually is the IP address of the local routing device. If you do not configure a router identifier, the IP address of the first interface to come online is used. This is usually the loopback interface. Otherwise, the first hardware interface with an IP address is used.</p>



NOTE: We strongly recommend that you configure the router identifier under the [edit routing-options] hierarchy level to avoid unpredictable behavior if the interface address on a loopback interface changes.

You must configure a router-id in order for BGP and OSPF to function in a routing instance. Use the **show route instance detail** command to display the router-id value for a routing instance. If the router-id is 0.0.0.0, then the routing instance has no router-id.

For more information about the router identifier in OSPF, see [“Example: Configuring an OSPF Router Identifier” on page 4400](#).



NOTE: If you run OSPF for IPv6 or BGP for IPv6 in a routing instance, you must configure an IPv4 router identifier (**router-id**) in the routing instance itself. In other words, the IPv4 router-id in the main routing instance is not inherited by other routing instances. Even if you run *only* IPv6 OSPF or BGP in a routing instance, the IPv4 router-id must be configured because OSPF and BGP, even when used exclusively with IPv6, use the IPv4 router-id for handshaking. If you do not configure the IPv4 router-id in the IPv6 OSPF or BGP routing instance, then the IPv6 protocols will use invalid IPv4 address 0.0.0.0 and the adjacencies and connections will fail.

Options	address —IP address of the routing device. Default: Address of the first interface encountered by Junos OS
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Examples: Configuring External BGP Peering on page 3601• Examples: Configuring Internal BGP Peering on page 3624

routing-instances

Syntax	<pre>routing-instances <i>routing-instance-name</i> { <i>description</i>; instance-type virtual-router; <i>interface</i> <i>interface-name</i>; <i>protocols</i>; <i>routing-options</i> }</pre>
Hierarchy Level	[edit]
Release Information	Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	(QFabric switches only) Configure a virtual router routing instance.
Options	<i>routing-instance-name</i> —Name of this routing instance. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Virtual Router Routing Instances on page 3232

routing-options

Syntax	routing-options { ... }
Hierarchy Level	[edit], [edit routing-instances <i>routing-instance-name</i>]
Release Information	Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure protocol-independent routing properties.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Overview of Routing Options on page 3227 • Understanding Distributed Periodic Packet Management on page 3233 • Example: Configuring SSM Maps for Different Groups to Different Sources on page 5541

scope

Syntax	scope <i>scope-name</i> { interface [<i>interface-names</i>]; prefix <i>destination-prefix</i> ; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast], [edit logical-systems <i>logical-system-name</i> routing-options multicast], [edit routing-instances <i>routing-instance-name</i> routing-options multicast], [edit routing-options multicast]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure multicast scoping.
Options	<i>scope-name</i> —Name of the multicast scope. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Multicast Snooping

scope-policy

Syntax `scope-policy [policy-names];`

Hierarchy Level `[edit logical-systems logical-system-name routing-options multicast],`
`[edit routing-options multicast]`




NOTE: You can configure a scope policy at these two hierarchy levels only. You cannot apply a scope policy to a specific routing instance, because all scoping policies are applied to all routing instances. However, you can apply the `scope` statement to a specific routing instance at the `[edit routing-instances routing-instance-name routing-options multicast]` or `[edit logical-systems logical-system-name routing-instances routing-instance-name routing-options multicast]` hierarchy level.

Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Apply policies for scoping. The policy must be correctly configured at the edit policy-options policy-statement hierarchy level.
Options	<i>policy-names</i> —Name of one or more multicast scope policies.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• scope on page 3371• <i>Example: Using a Scope Policy for Multicast Scoping</i>

source (Source-Specific Multicast)

Syntax	<code>source [<i>addresses</i>];</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast ssm-map <i>ssm-map-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast ssm-map <i>ssm-map-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast ssm-map <i>ssm-map-name</i>],</p> <p>[edit routing-options multicast ssm-map <i>ssm-map-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	Specify IPv4 or IPv6 source addresses for an SSM map.
Options	<i>addresses</i> —IPv4 or IPv6 source addresses.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To view this statement in the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring SSM Mapping on page 5535

source-routing

Syntax	source-routing { (ip ipv6) }
Hierarchy Level	[edit routing-options]
Release Information	Statement for IPv6 introduced in Junos OS Release 8.2. Statement for IPv4 introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	<p>Enable source routing.</p> <p>Source routing allows a sender of a packet to partially or completely specify the route the packet takes through the network. In contrast, in non-source routing protocols, routers in the network determine the path based on the packet's destination.</p> <div> NOTE: We recommend that you not use source routing. Instead, we recommend that you use policy-based routing or filter-based forwarding to route packets based on source addresses.</div>
Default	Disabled
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"><i>Example: Configuring Filter-Based Forwarding on the Source Address</i>

ssm-groups

Syntax	<code>ssm-groups [<i>ip-addresses</i>];</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit routing-options multicast]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Configure source-specific multicast (SSM) groups.</p> <p>By default, the SSM group multicast address is limited to the IP address range from 232.0.0.0 through 232.255.255.255. However, you can extend SSM operations into another Class D range by including the ssm-groups statement in the configuration. The default SSM address range from 232.0.0.0 through 232.255.255.255 cannot be used in the ssm-groups statement. This statement is for adding other multicast addresses to the default SSM group addresses. This statement does not override the default SSM group address range.</p> <p>IGMPv3 supports SSM groups. By utilizing inclusion lists, only sources that are specified send to the SSM group.</p>
Options	<i>ip-addresses</i> —List of one or more additional SSM group addresses separated by a space.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Source-Specific Multicast Groups with Any-Source Override on page 5537

ssm-map (Routing Options Multicast)

Syntax	<code>ssm-map <i>ssm-map-name</i> { <i>policy</i> [<i>policy-names</i>]; <i>source</i> [<i>addresses</i>]; }</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast], [edit logical-systems <i>logical-system-name</i> routing-options multicast], [edit routing-instances <i>routing-instance-name</i> routing-options multicast], [edit routing-options multicast]
Release Information	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure SSM mapping.
Options	<i>ssm-map-name</i> —Name of the SSM map. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring SSM Mapping on page 5535

static (Routes)

Syntax	<pre> static { defaults { static-options; } rib-group <i>group-name</i>; route <i>destination-prefix</i> { next-hop <i>address</i>; next-hop <i>options</i>; qualified-next-hop <i>address</i> { metric <i>metric</i>; preference <i>preference</i>; } static-options; } } </pre>
Hierarchy Level	[edit routing-options], [edit routing-options rib <i>routing-table-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure static routes to be installed in the routing table. You can specify any number of routes within a single static statement, and you can specify any number of static options in the configuration.
Options	<p>defaults—Specify global static route options. These options only set default attributes inherited by all newly created static routes. These are treated as global defaults and apply to all the static routes you configure in the static statement. This part of the static statement is optional.</p> <p>route <i>destination-prefix</i>—Destination of the static route.</p> <ul style="list-style-type: none"> defaults—For the default route to the destination. This is equivalent to specifying an IP address of 0.0.0.0/0. <i>destination-prefix/prefix-length—destination-prefix</i> is the network portion of the IP address, and <i>prefix-length</i> is the destination prefix length. next-hop <i>address</i>—Reach the next-hop routing device by specifying an IP address, an interface name, or an ISO network entity title (NET). <i>nsap-prefix—nsap-prefix</i> is the network service access point (NSAP) address for ISO. <p>next-hop <i>options</i>—Additional information for how to manage forwarding of packets to the next hop.</p> <ul style="list-style-type: none"> discard—Do not forward packets addressed to this destination. Instead, drop the packets, do not send ICMP unreachable messages to the packets' originators, and install a reject route for this destination into the routing table.

- **iso-net**—Reach the next-hop routing device by specifying an ISO NSAP.
- **next-table *routing-table-name***—Name of the next routing table to the destination.
- **receive**—Install a receive route for this destination into the routing table.
- **reject**—Do not forward packets addressed to this destination. Instead, drop the packets, send ICMP unreachable messages to the packets' originators, and install a reject route for this destination into the routing table.

static-options—(Optional under **route**) Additional information about static routes, which is included with the route when it is installed in the routing table.

You can specify one or more of the following in **static-options**. Each of the options is explained separately.

- **(active | passive);**
- **(install | no-install);**
- **(metric | metric2 | metric3 | metric4) *value* <type type>;**
- **(preference | preference2 | color | color2) *preference* <type type>;**
- **(resolve | no-resolve);**
- **(no-retain | retain);**

The remaining statements are explained separately.

Required Privilege Level	routing—To view this statement in the configuration.
	routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Static Routing on page 3228

subscriber-leave-timer

Syntax	<code>subscriber-leave-timer <i>seconds</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast interface <i>interface-name</i>],</p> <p>[edit routing-options multicast interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.2.</p> <p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Length of time before the multicast VLAN updates QoS data (for example, available bandwidth) for subscriber interfaces after it receives an IGMP leave message.
Options	<p>seconds—Length of time before the multicast VLAN updates QoS data (for example, available bandwidth) for subscriber interfaces after it receives an IGMP leave message. Specifying a value of 0 results in an immediate update. This is the same as if the statement were not configured.</p> <p>Range: 0 through 30</p> <p>Default: 0 seconds</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

tag (Routing Options)

Syntax	<code>tag metric type number;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (aggregate generate static) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options (aggregate generate static) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options aggregate generate static) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</p> <p>[edit routing-options (aggregate generate static) (defaults route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Associate a tag with a static, aggregate, or generated route.
Default	No tag strings are associated with routes.
Options	<p><i>metric</i>—Tag metric.</p> <p>Range: 0 through 4,294,967,295</p> <p><i>type number</i>—Tag type.</p> <p>Range: 1 through 16</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Examples: Configuring Static Routes</i> • <i>Example: Summarizing Static Routes Through Route Aggregation</i> • <i>Example: Conditionally Generating Static Routes</i> • aggregate on page 3272 • generate on page 3309 • <i>static</i>

threshold (Multicast Forwarding Cache)

Syntax	<pre>threshold { log-warning <i>value</i>; suppress <i>value</i> <reuse <i>value</i>>; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast forwarding-cache],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast forwarding-cache family (inet inet6)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast forwarding-cache],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast forwarding-cache family (inet inet6)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast forwarding-cache],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast forwarding-cache (inet inet6)],</p> <p>[edit routing-options multicast forwarding-cache],</p> <p>[edit routing-options multicast forwarding-cache family (inet inet6)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Configure the global suppression, reuse, and warning log message thresholds for multicast forwarding cache limits. You can configure the thresholds globally for the multicast forwarding cache or individually for the IPv4 and IPv6 multicast forwarding caches. Configuring the threshold statement globally for the multicast forwarding cache or including the family statement to configure the thresholds for the IPv4 and IPv6 multicast forwarding caches are mutually exclusive.</p> <p>To confirm the configured threshold values, use the show multicast forwarding-cache statistics command.</p>
Options	<p>reuse <i>value</i>—(Optional) Value at which to begin creating new multicast forwarding cache entries. If configured, this number should be less than the suppress value.</p> <p>Range: 1 through 200,000</p> <p>suppress <i>value</i>—Value at which to begin suppressing new multicast forwarding cache entries. This value is mandatory. This number should be greater than the reuse value.</p> <p>Range: 1 through 200,000</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Examples: Configuring the Multicast Forwarding Cache</i>

timeout (Flow Maps)

Syntax	timeout (never non-discard-entry-only <i>minutes</i>);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast flow-map <i>flow-map-name</i>], [edit logical-systems <i>logical-system-name</i> routing-options multicast flow-map <i>flow-map-name</i>], [edit routing-instances <i>routing-instance-name</i> routing-options multicast flow-map <i>flow-map-name</i>], [edit routing-options multicast flow-map <i>flow-map-name</i>]
Release Information	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the timeout value for multicast forwarding cache entries associated with the flow map.
Options	minutes —Length of time that the forwarding cache entry remains active. Range: 1 through 720 never non-discard-entry-only —Specify that the forwarding cache entry always remain active. If you omit the non-discard-entry-only option, all multicast forwarding entries, including those in forwarding and pruned states, are kept forever. If you include the non-discard-entry-only option, entries with forwarding states are kept forever, and entries with pruned states time out.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

timeout (Multicast)

Syntax	<code>timeout <i>minutes</i> <family (inet inet6)>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast forwarding-cache],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast forwarding-cache],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast forwarding-cache],</p> <p>[edit routing-options multicast forwarding-cache]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.2.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Configure the timeout value for multicast forwarding cache entries.
Options	<p><i>minutes</i>—Length of time that the forwarding cache limit remains active.</p> <p>Range: 1 through 720</p> <p><i>family (inet inet6)</i>—(Optional) Apply the configured timeout to either IPv4 or IPv6 multicast forwarding cache entries. Configuring the timeout statement globally for the multicast forwarding cache or including the family statement to configure the timeout value for the IPv4 and IPv6 multicast forwarding caches are mutually exclusive.</p> <p>Default: By default, the configured timeout applies to both IPv4 and IPv6 multicast forwarding cache entries.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring the Multicast Forwarding Cache</i>

traceoptions

Syntax	<pre>traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <disable>; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit routing-options],</p> <p>[edit routing-options flow],</p> <p>[edit routing-options multicast]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>nsr-synchronization flag for BGP, IS-IS, LDP, and OSPF added in Junos OS Release 8.4.</p> <p>nsr-synchronization and nsr-packet flags for BFD sessions added in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>nsr-synchronization flag for RIP and RIPng added in Junos OS Release 9.0.</p> <p>nsr-synchronization flag for Layer 2 VPNs and VPLS added in Junos OS Release 9.1.</p> <p>nsr-synchronization flag for PIM added in Junos OS Release 9.3.</p> <p>nsr-synchronization flag for MPLS added in Junos OS Release 10.1.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>nsr-synchronization flag for MSDP added in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Define tracing operations that track all routing protocol functionality in the routing device.</p> <p>To specify more than one tracing operation, include multiple flag statements.</p>
Default	If you do not include this statement, no global tracing operations are performed.
Options	<p>Values:</p> <p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. We recommend that you place global routing protocol tracing output in the file routing-log.</p>

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file.0***, then ***trace-file.1***, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 10 files

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. These are the global routing protocol tracing options:

- **all**—All tracing operations
- **condition-manager**—Condition-manager events
- **config-internal**—Configuration internals
- **general**—All normal operations and routing table changes (a combination of the **normal** and **route** trace operations)
- **graceful-restart**—Graceful restart operations
- **normal**—All normal operations
- **nsr-packet**—Detailed trace information for BFD nonstop active routing only
- **nsr-synchronization**—Tracing operations for nonstop active routing
- **nsr-synchronization**—Nonstop active routing synchronization
- **parse**—Configuration parsing
- **policy**—Routing policy operations and actions
- **regex-parse**—Regular-expression parsing
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

no-world-readable—(Optional) Prevent any user from reading the log file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file.0***. When the ***trace-file*** again reaches its maximum size, ***trace-file.0*** is renamed ***trace-file.1*** and ***trace-file*** is renamed ***trace-file.0***. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 128 KB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege	routing and trace—To view this statement in the configuration.
Level	routing-control and trace-control—To add this statement to the configuration.

Related Documentation	<ul style="list-style-type: none">• <i>Example: Tracing Global Routing Protocol Operations</i>
------------------------------	--

upstream-interface

Syntax	<code>upstream-interface [<i>interface-names</i>];</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast pim-to-igmp-proxy],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast pim-to-mld-proxy],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast pim-to-igmp-proxy],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast pim-to-mld-proxy],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast pim-to-igmp-proxy],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast pim-to-mld-proxy],</p> <p>[edit routing-options multicast pim-to-igmp-proxy],</p> <p>[edit routing-options multicast pim-to-mld-proxy]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Configure at least one, but not more than two, upstream interfaces on the rendezvous point (RP) routing device that resides between a customer edge-facing Protocol Independent Multicast (PIM) domain and a core-facing PIM domain. The RP routing device translates PIM join or prune messages into corresponding IGMP report or leave messages (if you include the pim-to-igmp-proxy statement), or into corresponding MLD report or leave messages (if you include the pim-to-mld-proxy statement). The routing device then proxies the IGMP or MLD report or leave messages to one or both upstream interfaces to forward IPv4 multicast traffic (for IGMP) or IPv6 multicast traffic (for MLD) across the PIM domains.</p>
Options	<p><i>interface-names</i>—Names of one or two upstream interfaces to which the RP routing device proxies IGMP or MLD report or leave messages for transmission of multicast traffic across PIM domains. You can specify a maximum of two upstream interfaces on the RP routing device. To configure a set of two upstream interfaces, specify the full interface names, including all physical and logical address components, within square brackets ([]).</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring PIM-to-IGMP Message Translation</i> • <i>Configuring PIM-to-MLD Message Translation</i>

Operational Commands (Routing Options)

- `clear ipv6 neighbors`
- `show as-path`
- `show as-path domain`
- `show as-path summary`
- `show ipv6 neighbors`
- `show ipv6 router-advertisement`
- `show route`
- `show route active-path`
- `show route all`
- `show route aspath-regex`
- `show route best`
- `show route brief`
- `show route community`
- `show route community-name`
- `show route damping`
- `show route detail`
- `show route exact`
- `show route export`
- `show route extensive`
- `show route flow validation`
- `show route forwarding-table`
- `show route inactive-path`
- `show route inactive-prefix`
- `show route instance`
- `show route label`
- `show route label-switched-path`

- `show route martians`
- `show route next-hop`
- `show route no-community`
- `show route protocol`
- `show route range`
- `show route receive-protocol`
- `show route resolution`
- `show route snooping`
- `show route source-gateway`
- `show route summary`
- `show route table`
- `show route terse`

clear ipv6 neighbors

Syntax	clear ipv6 neighbors <all host <i>hostname</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.3 for EX Series switches. Command introduced in Junos OS Release 12.2 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Clear IPv6 neighbor cache information.
Options	none —Clear all IPv6 neighbor cache information. all —(Optional) Clear all IPv6 neighbor cache information. host <i>hostname</i> —(Optional) Clear the information for the specified IPv6 neighbors.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show ipv6 neighbors on page 3401
List of Sample Output	clear ipv6 neighbors on page 3391
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ipv6 neighbors

```
user@host> clear ipv6 neighbors
```

show as-path

List of Syntax	Syntax on page 3392 Syntax (EX Series Switches) on page 3392
Syntax	<code>show as-path</code> <code><brief detail></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switches)	<code>show as-path</code> <code><brief detail></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Display the distribution of autonomous system (AS) paths that the local routing device is using (usually through the routing table). Use this command to debug problems for AS paths and to understand how AS paths have been manipulated through a policy (through the as-path-prepend action) or through aggregation.</p> <p>AS paths are stored in a hash table. A hash table is one method for fast lookup. Each entry in the table is called a bucket. Junos OS computes a hash value that indicates in which bucket the AS path is stored. The AS paths are dispersed among the hash buckets so that a manageable number of AS paths is stored in each bucket. Only unique AS paths are stored. Duplicate AS paths increase a reference count, but do not increase the number of AS paths stored in the hash table.</p>
Options	<p>none—Display basic information about AS paths that the local routing device is using (same as brief).</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show as-path summary on page 3399
List of Sample Output	show as-path on page 3393 show as-path detail on page 3394
Output Fields	Table 292 lists the output fields for the show as-path command. Output fields are listed in the approximate order in which they appear.

Table 292: show as-path Output Fields

Field Name	Field Description	Level of Output
Total AS paths	Total number of AS paths.	brief none
Bucket	Bucket number.	All levels
Count	Number of AS path entries in this bucket.	All levels
AS path	<p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> • I—IGP. • E—EGP. • ?—Incomplete; typically, the AS path was aggregated. • Atomic—Route is an aggregate of several route prefixes. • Aggregat—Routing device has summarized a range of prefixes. 	All levels
domain	Number of independent AS domains. The AS paths of an independent AS domain are not shared with the AS paths and AS path attributes of other domains, including the master routing instance domain.	detail
neighbor as	AS peer address.	detail
length	Length of the AS path.	detail
segments	Length of the AS segment descriptor.	detail
references	Path reference count.	detail

Sample Output

show as-path

```

user@host> show as-path
Total AS paths: 30382
  Bucket 0      Count: 36
    I
    14203 2914 174 31752 I
    14203 2914 701 21512 I
    14203 2914 1239 26632 I
    14203 2914 1239 29704 I
    14203 2914 4323 10248 I
    14203 2914 4766 23560 I
    14203 2914 6395 32776 I
    14203 2914 7911 11272 I
    14203 2914 12180 18440 I
    14203 2914 17408 17416 I
    14203 2914 701 702 24586 I
    14203 2914 1239 4657 9226 I
    14203 2914 1239 7132 16394 I
    14203 2914 1299 8308 34826 I
    14203 2914 3320 5603 28682 I

```

```

14203 2914 3491 1680 33802 I
14203 2914 3549 7908 27658 I
14203 2914 3549 20804 30730 I
14203 2914 7018 2687 9226 I
14203 2914 174 9318 9318 23564 I
14203 2914 701 3786 3786 23564 I
14203 2914 701 4761 4795 9228 I
14203 2914 1239 7132 5673 18444 I
14203 2914 3491 20485 24588 24588 I
14203 2914 5511 2200 1945 2060 I
14203 2914 7911 14325 14325 14348 I
14203 2914 701 4637 9230 9230 9230 I
14203 2914 6395 14 14 14 14 I
14203 2914 9299 6163 6163 6163 9232 I
14203 2914 3356 3356 3356 3356 11955 21522 I
14203 2914 9837 9837 9219 I Aggregator: 9219 202.27.91.253
14203 2914 174 30209 30222 30222 30222 ?
14203 2914 1299 5377 I (Atomic) Aggregator: 5377 193.219.192.22
14203 2914 4323 36097 I (Atomic) Aggregator: 36097 216.69.252.254
14203 2914 209 2516 17676 23813 I (Atomic) Aggregator: 23813 219.127.233.66
Bucket 1    Count: 28
14203 2914 35847 I
14203 2914 174 19465 I
14203 2914 174 35849 I
14203 2914 2828 32777 I
14203 2914 4323 14345 I
14203 2914 4323 29705 I
14203 2914 6395 32777 I

```

...

show as-path detail

```

user@host> show as-path detail
Total AS paths: 30410
Bucket 0    Count: 36
AS path: I
  domain 0, length 0, segments 0, references 54
AS path: 14203 2914 174 31752 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 701 21512 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 1239 26632 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 1239 29704 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 4323 10248 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 4766 23560 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 6395 32776 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 3
AS path: 14203 2914 7911 11272 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 12180 18440 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 3
AS path: 14203 2914 17408 17416 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 3
AS path: 14203 2914 701 702 24586 I
  domain 1, neighbor as: 14203, length 5, segments 1, references 3
AS path: 14203 2914 1239 4657 9226 I

```

```

    domain 1, neighbor as: 14203, length 5, segments 1, references 7
AS path: 14203 2914 1239 7132 16394 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 1299 8308 34826 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 3320 5603 28682 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 3491 1680 33802 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 3549 7908 27658 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 3549 20804 30730 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 7018 2687 9226 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 3
AS path: 14203 2914 174 9318 9318 23564 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 701 3786 3786 23564 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 701 4761 4795 9228 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 14
AS path: 14203 2914 1239 7132 5673 18444 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 3491 20485 24588 24588 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 4
AS path: 14203 2914 5511 2200 1945 2060 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 7911 14325 14325 14348 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 701 4637 9230 9230 9230 I
    domain 1, neighbor as: 14203, length 7, segments 1, references 3
AS path: 14203 2914 6395 14 14 14 14 I
    domain 1, neighbor as: 14203, length 7, segments 1, references 10
...

```

show as-path domain

List of Syntax	Syntax on page 3396 Syntax (EX Series Switches) on page 3396
Syntax	show as-path domain <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	show as-path domain
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display autonomous system (AS) path domain information.
Options	none —(Optional) Display AS path domain information for all routing instances. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show as-path domain on page 3398
Output Fields	Table 293 lists the output fields for the show as-path domain command. Output fields are listed in the approximate order in which they appear

Table 293: show as-path domain Output Fields

Field Name	Field Description
Domain	Number of independent AS domains. The AS paths of an independent AS domain are not shared with the AS paths and AS path attributes of other domains, including the master routing instance domain.
Primary	Primary AS number.
References	Path reference count.
Number Paths	Number of known AS paths.
Flags	Information about the AS path: <ul style="list-style-type: none"> • ASLoop—Path contains an AS loop. • Atomic—Path includes the ATOMIC_AGGREGATE path attribute. • Local—Path was created by local aggregation. • Master—Path was created by the master routing instance.
Local AS	AS number of the local routing device.

Table 293: show as-path domain Output Fields (*continued*)

Field Name	Field Description
Loops	How many times this AS number can appear in an AS path.

Sample Output

show as-path domain

```
user@host> show as-path domain
Domain: 1          Primary: 10458
References:        3 Paths:      30383
Flags: Master
Local AS: 10458   Loops: 1
```

show as-path summary

List of Syntax	Syntax on page 3399 Syntax (EX Series Switches) on page 3399
Syntax	<pre>show as-path summary <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switches)	show as-path summary
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	<p>Display autonomous system (AS) path summary information.</p> <p>AS paths are stored in a hash table. A hash table is one method for fast lookup. Each entry in the table is called a bucket. Junos OS computes a hash value that indicates in which bucket the AS path is stored. The AS paths are dispersed among the hash buckets so that a manageable number of AS paths is stored in each bucket. Only unique AS paths are stored. Duplicate AS paths increase a reference count, but do not increase the number of AS paths stored in the hash table.</p>
Options	<p>none—(Optional) Display AS path summary information for all routing instances.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show as-path on page 3392
List of Sample Output	show as-path summary on page 3400
Output Fields	<p>Table 294 lists the output fields for the show as-path summary command. Output fields are listed in the approximate order in which they appear.</p>

Table 294: show as-path summary Output Fields

Field Name	Field Description
AS Paths	Number of AS paths.
Buckets	Number of hash buckets in use.
Max	Maximum number of AS path entries per bucket.
Min	Minimum number of AS path entries per bucket.
Avg	Average number of AS path entries per bucket.

Table 294: show as-path summary Output Fields (*continued*)

Field Name	Field Description
Std deviation	Standard deviation of AS path entries per bucket.

Sample Output

show as-path summary

```
user@host> show as-path summary
AS Paths Buckets Max Min Avg Std deviation
30425    1024    95  12  29    6.481419
```


show ipv6 neighbors

Syntax	show ipv6 neighbors
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.2 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Display information about the IPv6 neighbor cache.
Options	This command has no options.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> clear ipv6 neighbors on page 3391
List of Sample Output	show ipv6 neighbors on page 3401
Output Fields	Table 295 describes the output fields for the show ipv6 neighbors command. Output fields are listed in the approximate order in which they appear.

Table 295: show ipv6 neighbors Output Fields

Field Name	Field Description
IPv6 Address	Name of the IPv6 interface.
Linklayer Address	Link-layer address.
State	State of the link: up , down , incomplete , reachable , stale , or unreachable .
Exp	Number of seconds until the entry expires.
Rtr	Whether the neighbor is a routing device: yes or no .
Secure	Whether this entry was created using the Secure Neighbor Discovery (SEND) protocol: yes or no .
Interface	Name of the interface.

Sample Output

show ipv6 neighbors

```

user@host> show ipv6 neighbors
IPv6 Address          Linklayer Address  State      Exp Rtr Secure
Interface
2001:db8:0:1:2a0:a514:0:24c  00:05:85:8f:c8:bd  stale      546 yes no

```

fe-1/2/0.1				
fe80::2a0:a514:0:24c	00:05:85:8f:c8:bd	stale	258	yes no
fe-1/2/0.1				
fe80::2a0:a514:0:64c	00:05:85:8f:c8:bd	stale	111	yes no
fe-1/2/1.5				
fe80::2a0:a514:0:a4c	00:05:85:8f:c8:bd	stale	327	yes no
fe-1/2/2.9				

show ipv6 router-advertisement

Syntax	<pre>show ipv6 router-advertisement <conflicts> <interface <i>interface</i>> <logical-system (all <i>logical-system-name</i>)> <prefix <i>prefix/prefix length</i>></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 12.2 for the QFX Series.</p>
Description	Display information about IPv6 router advertisements, including statistics about messages sent and received on interfaces, and information received from advertisements from other routers.
Options	<p>none—Display all IPv6 router advertisement information for all interfaces.</p> <p>conflicts—(Optional) Display only the IPv6 router advertisement information that is conflicting.</p> <p>interface <i>interface</i>—(Optional) Display IPv6 router advertisement information for the specified interface.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>prefix <i>prefix/prefix length</i>—(Optional) Display IPv6 router advertisement information for the specified prefix.</p>
Additional Information	The display identifies conflicting information by enclosing the value the router is advertising in brackets.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear ipv6 router-advertisement
List of Sample Output	<p>show ipv6 router-advertisement on page 3404</p> <p>show ipv6 router-advertisement conflicts on page 3405</p> <p>show ipv6 router-advertisement prefix on page 3405</p>
Output Fields	<p>Table 296 describes the output fields for the show ipv6 router-advertisement command. Output fields are listed in the approximate order in which they appear.</p>

Table 296: show ipv6 router-advertisement Output Fields

Field Name	Field Description
Interface	Name of the interface.
Advertisements sent	Number of router advertisements sent and elapsed time since they were sent.

Table 296: show ipv6 router-advertisement Output Fields (*continued*)

Field Name	Field Description
Solicits received	Number of solicitation messages received.
Advertisements received	Number of router advertisements received.
Advertisements from	Names of interfaces from which router advertisements have been received and elapsed time since the last one was received.
Managed	Managed address configuration flag: 0 (stateless) or 1 (stateful).
Other configuration	Other stateful configuration flag: 0 (stateless) or 1 (stateful).
Reachable time	Time that a node identifies a neighbor as reachable after receiving a reachability confirmation, in milliseconds.
Default lifetime	Default lifetime, in seconds: from 0 seconds to 18.2 hours. A setting of 0 indicates that the router is not a default router.
Retransmit timer	Time between retransmitted Neighbor Solicitation messages, in milliseconds.
Current hop limit	Configured current hop limit.
Prefix	Name and length of the prefix.
Valid lifetime	How long the prefix remains valid for onlink determination.
Preferred lifetime	How long the prefix generated by stateless autoconfiguration remains preferred.
On link	Onlink flag: 0 (not onlink) or 1 (onlink).
Autonomous	Autonomous address configuration flag: 0 (not autonomous) or 1 (autonomous).

Sample Output

show ipv6 router-advertisement

```

user@host> show ipv6 router-advertisement
Interface: fe-0/1/1.0
  Advertisements sent: 0
  Solicits received: 0
  Advertisements received: 0
Interface: fxp0.0
  Advertisements sent: 0
  Solicits received: 0
  Advertisements received: 1
  Advertisement from fe80::2d0:b7ff:fe1e:7b0e, heard 00:00:13 ago
  Managed: 0
  Other configuration: 0 [1]
  Reachable time: 0 ms
  Default lifetime: 1800 sec

```

```
Retransmit timer: 0 ms  
Current hop limit: 64
```

show ipv6 router-advertisement conflicts

```
user@host> show ipv6 router-advertisement conflicts  
Interface: fxp0.0  
Advertisement from fe80::2d0:b7ff:fe1e:7b0e, heard 00:01:08 ago  
Other configuration: 0 [1]
```

show ipv6 router-advertisement prefix

```
user@host> show ipv6 router-advertisement prefix 8040::/16  
Interface: fe-0/1/3.0  
Advertisements sent: 3, last sent 00:04:11 ago  
Solicits received: 0  
Advertisements received: 3  
Advertisement from fe80::290:69ff:fe9a:5403, heard 00:00:05 ago  
Managed: 0  
Other configuration: 0  
Reachable time: 0 ms  
Default lifetime: 180 sec [1800 sec]  
Retransmit timer: 0 ms  
Current hop limit: 64  
Prefix: 8040:1::/64  
Valid lifetime: 2592000 sec  
Preferred lifetime: 604800 sec  
On link: 1  
Autonomous: 1
```

show route

List of Syntax [Syntax on page 3406](#)
 [Syntax \(EX Series Switches\) on page 3406](#)

Syntax show route
 <all>
 <*destination-prefix*>
 <logical-system (all | *logical-system-name*)>
 <private>

Syntax (EX Series Switches) show route
 <all>
 <*destination-prefix*>
 <private>

Release Information Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.0 for EX Series switches.
 Option **private** introduced in Junos OS Release 9.5.
 Option **private** introduced in Junos OS Release 9.5 for EX Series switches.
 Command introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.

Description Display the active entries in the routing tables.

Options **none**—Display brief information about all active entries in the routing tables.

all—(Optional) Display information about all routing tables, including private, or internal, routing tables.

destination-prefix—(Optional) Display active entries for the specified address or range of addresses.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

private—(Optional) Display information only about all private, or internal, routing tables.

Required Privilege Level view

Related Documentation

- [Example: Configuring RIP on page 4685](#)
- [Example: Configuring RIPng](#)
- [Example: Configuring IS-IS](#)
- [Examples: Configuring Internal BGP Peering on page 3624](#)
- [Examples: Configuring External BGP Peering on page 3601](#)
- [Examples: Configuring OSPF Routing Policy on page 4527](#)
- [Verifying and Managing Junos OS Enhanced Subscriber Management](#)

List of Sample Output [show route on page 3409](#)
[show route on page 3410](#)
[show route \(with Destination Prefix\) on page 3410](#)
[show route destination-prefix detail on page 3410](#)
[show route extensive on page 3411](#)
[show route \(Enhanced Subscriber Management\) on page 3411](#)

Output Fields Table 297 describes the output fields for the **show route** command. Output fields are listed in the approximate order in which they appear.

Table 297: show route Output Fields

Field Name	Field Description
<i>routing-table-name</i>	Name of the routing table (for example, inet.0).
<i>number destinations</i>	Number of destinations for which there are routes in the routing table.
<i>number routes</i>	<p>Number of routes in the routing table and total number of routes in the following states:</p> <ul style="list-style-type: none"> • active (routes that are active). • holddown (routes that are in the pending state before being declared inactive). A holddown route was once the active route and is no longer the active route. The route is in the holddown state because a protocol still has interest in the route, meaning that the interest bit is set. A protocol might have its interest bit set on the previously active route because the protocol is still advertising the route. The route will be deleted after all protocols withdraw their advertisement of the route and remove their interest bit. A persistent holddown state often means that the interested protocol is not releasing its interest bit properly. <p>However, if you have configured advertisement of multiple routes (with the add-path or advertise-inactive statement), the holddown bit is most likely set because BGP is advertising the route as an active route. In this case, you can ignore the holddown state because nothing is wrong.</p> <ul style="list-style-type: none"> • hidden (routes that are not used because of a routing policy).
<i>destination-prefix</i>	<p>Route destination (for example:10.0.0.1/24). Sometimes the route information is presented in another format, such as:</p> <ul style="list-style-type: none"> • MPLS-label (for example, 80001). • interface-name (for example, ge-1/0/2). • neighbor-address:control-word-status:encapsulation type:vc-id:source (Layer 2 circuit only. For example, 10.1.1.195:NoCtrlWord:1:1:Local/96): <ul style="list-style-type: none"> • neighbor-address—Address of the neighbor. • control-word-status—Whether the use of the control word has been negotiated for this virtual circuit: NoCtrlWord or CtrlWord. • encapsulation type—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport. • vc-id—Virtual circuit identifier. • source—Source of the advertisement: Local or Remote.

Table 297: show route Output Fields (*continued*)

Field Name	Field Description
[<i>protocol, preference</i>]	<p>Protocol from which the route was learned and the preference value for the route.</p> <ul style="list-style-type: none"> • +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • - —A hyphen indicates the last active route. • *—An asterisk indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route. <p>In every routing metric except for the BGP LocalPref attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the LocalPref value in the Preference2 field. For example, if the LocalPref value for Route 1 is 100, the Preference2 value is -101. If the LocalPref value for Route 2 is 155, the Preference2 value is -156. Route 2 is preferred because it has a higher LocalPref value and a lower Preference2 value.</p>
<i>weeks:days</i> <i>hours:minutes:seconds</i>	How long the route been known (for example, 2w4d 13:11:14 , or 2 weeks, 4 days, 13 hours, 11 minutes, and 14 seconds).
metric	Cost value of the indicated route. For routes within an AS, the cost is determined by the IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
localpref	Local preference value included in the route.
from	Interface from which the route was received.
AS path	<p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> • I—IGP. • E—EGP. • ?—Incomplete; typically, the AS path was aggregated. <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> • []—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device, or if AS path prepending is configured. • { }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • ()—Parentheses enclose a confederation. • ([])—Parentheses and brackets enclose a confederation set. <p>NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p>

Table 297: show route Output Fields (*continued*)

Field Name	Field Description
validation-state	<p>(BGP-learned routes) Validation status of the route:</p> <ul style="list-style-type: none"> • Invalid—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • Unknown—Indicates that the prefix is not among the prefixes or prefix ranges in the database. • Unverified—Indicates that the origin of the prefix is not verified against the database. This is because the database got populated and the validation is not called for in the BGP import policy, although origin validation is enabled, or the origin validation is not enabled for the BGP peers. • Valid—Indicates that the prefix and autonomous system pair are found in the database.
to	<p>Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.</p> <p>If the destination is Discard, traffic is dropped.</p>
via	<p>Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word Selected. This field can also contain the following information:</p> <ul style="list-style-type: none"> • Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing. • lsp-path-name—Name of the LSP used to reach the next hop. • label-action—MPLS label and operation occurring at the next hop. The operation can be pop (where a label is removed from the top of the stack), push (where another label is added to the label stack), or swap (where a label is replaced by another label). For VPNs, expect to see multiple push operations, corresponding to the inner and outer labels required for VPN routes (in the case of a direct PE-to-PE connection, the VPN route would have the inner label push only).
Private unicast	<p>(Enhanced subscriber management for MX Series routers) Indicates that an access-internal route is managed by enhanced subscriber management. By contrast, access-internal routes <i>not</i> managed by enhanced subscriber management are displayed with associated next-hop and media access control (MAC) address information.</p>

Sample Output

show route

```

user@host> show route
inet.0: 11 destinations, 12 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:65500:1:10.0.0.20/240
    * [MVPN/70] 19:53:41, metric2 1
    Indirect
1:65500:1:10.0.0.40/240
    * [BGP/170] 19:53:29, localpref 100, from 10.0.0.30
    AS path: I
    > to 10.0.24.4 via 1t-0/3/0.24, label-switched-path toD

```

```

[BGP/170] 19:53:26, localpref 100, from 10.0.0.33
AS path: I
> to 10.0.24.4 via lt-0/3/0.24, label-switched-path toD
1:65500:1:10.0.0.60/240
*[BGP/170] 19:53:29, localpref 100, from 10.0.0.30
AS path: I
> to 10.0.28.8 via lt-0/3/0.28, label-switched-path toF
[BGP/170] 19:53:25, localpref 100, from 10.0.0.33
AS path: I
> to 10.0.28.8 via lt-0/3/0.28, label-switched-path toF

```

show route

The following sample output shows a VPN route with composite next hops enabled. The first **Push** operation corresponds to the outer label. The second **Push** operation corresponds to the inner label.

```

user@host> show route 70.0.0.0

13979:665001.inet.0: 871 destinations, 3556 routes (871 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both

70.0.0.0/24      @[BGP/170] 00:28:32, localpref 100, from 10.9.9.160
                  AS path: 13980 ?, validation-state: unverified
                  > to 10.100.0.42 via ae2.0, Push 16, Push 300368(top)
                  [BGP/170] 00:28:28, localpref 100, from 10.9.9.169
                  AS path: 13980 ?, validation-state: unverified
                  > to 10.100.0.42 via ae2.0, Push 126016, Push 300368(top)
                  #[Multipath/255] 00:28:28, metric2 102
                  > to 10.100.0.42 via ae2.0, Push 16, Push 300368(top)
                  to 10.100.0.42 via ae2.0, Push 16, Push 300368(top)

```

show route (with Destination Prefix)

```

user@host> show route 172.16.0.0/12

inet.0: 10 destinations, 10 routes (9 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.0.0/12    *[Static/5] 2w4d 12:54:27
                  > to 192.168.167.254 via fxp0.0

```

show route destination-prefix detail

```

user@host> show route 5.5.5.0 detail

inet.0: 15 destinations, 20 routes (15 active, 0 holddown, 0 hidden)
5.5.5.0/24 (2 entries, 2 announced)
  *BGP      Preference: 170/-101
  ...
  BGP-Static Preference: 4294967292
    Next hop type: Discard
    Address: 0x9041ae4
    Next-hop reference count: 2
    State: <NoReadvrt Int Ext AlwaysFlash>
  Inactive reason: Route Preference
  Local AS: 200
  Age: 4d 1:40:40
  Validation State: unverified
  Task: RT

```

```
Announcement bits (1): 2-BGP_RT_Background
AS path: 4 5 6 I
```

show route extensive

```
user@host> show route extensive
v1.mvpn.0: 5 destinations, 8 routes (5 active, 1 holddown, 0 hidden)
1:65500:1:10.0.0.40/240 (1 entry, 1 announced)
  *BGP Preference: 170/-101
    PMSI: Flags 0x0: Label[0:0:0]: PIM-SM: Sender 10.0.0.40 Group 225.1.1.1

    Next hop type: Indirect
    Address: 0x92455b8
    Next-hop reference count: 2
    Source: 10.0.0.30
    Protocol next hop: 10.0.0.40
    Indirect next hop: 2 no-forward
    State: <Active Int Ext>
      Local AS: 65500 Peer AS: 65500
    Age: 3 Metric2: 1
    Validation State: unverified
    Task: BGP_65500.10.0.0.30+179
    Announcement bits (2): 0-PIM.v1 1-mvpn global task
    AS path: I (Originator) Cluster list: 10.0.0.30
    AS path: Originator ID: 10.0.0.40
    Communities: target:65520:100
    Import Accepted
    Localpref: 100
    Router ID: 10.0.0.30
    Primary Routing Table bgp.mvpn.0
    Indirect next hops: 1
      Protocol next hop: 10.0.0.40 Metric: 1
      Indirect next hop: 2 no-forward
      Indirect path forwarding next hops: 1
        Next hop type: Router
        Next hop: 10.0.24.4 via lt-0/3/0.24 weight 0x1
      10.0.0.40/32 Originating RIB: inet.3
        Metric: 1 Node path count: 1
        Forwarding nexthops: 1
        Nexthop: 10.0.24.4 via lt-0/3/0.24
```

show route (Enhanced Subscriber Management)

```
user@host> show route
inet.0: 41 destinations, 41 routes (40 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

100.20.0.111/32    *[Access-internal/12] 00:00:08
                  > to #0 10.0.0.1.93.65 via demux0.1073741824
100.20.0.112/32    *[Access-internal/12] 00:00:08
                  Private unicast
.
.
.
```

show route active-path

List of Syntax	Syntax on page 3412 Syntax (EX Series Switches) on page 3412
Syntax	<code>show route active-path</code> <code><brief detail extensive terse></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switches)	<code>show route active-path</code> <code><brief detail extensive terse></code>
Release Information	Command introduced in Junos OS Release 8.0. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display all active routes for destinations. An active route is a route that is selected as the best path. Inactive routes are not displayed.
Options	none —Display all active routes. brief detail extensive terse —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show route active-path on page 3412 show route active-path brief on page 3413 show route active-path detail on page 3413 show route active-path extensive on page 3414 show route active-path terse on page 3416
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route active-path

```
user@host> show route active-path

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.70.19/32    *[Direct/0] 21:33:52
                  > via lo0.0
10.255.71.50/32   *[IS-IS/15] 00:18:13, metric 10
                  > to 100.1.2.1 via so-2/1/3.0
100.1.2.0/24      *[Direct/0] 00:18:36
                  > via so-2/1/3.0
```

```

100.1.2.2/32      *[Local/0] 00:18:41
                  Local via so-2/1/3.0
192.168.64.0/21  *[Direct/0] 21:33:52
                  > via fxp0.0
192.168.70.19/32 *[Local/0] 21:33:52
                  Local via fxp0.0

```

show route active-path brief

The output for the **show route active-path brief** command is identical to that for the **show route active-path** command. For sample output, see [show route active-path on page 3412](#).

show route active-path detail

```

user@host> show route active-path detail

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)

10.255.70.19/32 (1 entry, 1 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 3
    Next hop: via lo0.0, selected
    State: <Active Int>
    Local AS: 200
    Age: 21:37:10
    Task: IF
    Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3
    AS path: I

10.255.71.50/32 (1 entry, 1 announced)
  *IS-IS Preference: 15
    Level: 1
    Next hop type: Router, Next hop index: 397
    Next-hop reference count: 4
    Next hop: 100.1.2.1 via so-2/1/3.0, selected
    State: <Active Int>
    Local AS: 200
    Age: 21:31 Metric: 10
    Task: IS-IS
    Announcement bits (4): 0-KRT 2-IS-IS 5-Resolve tree 2 6-Resolve
tree 3
    AS path: I

100.1.2.0/24 (1 entry, 1 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 3
    Next hop: via so-2/1/3.0, selected
    State: <Active Int>
    Local AS: 200
    Age: 21:54
    Task: IF
    Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3
    AS path: I

100.1.2.2/32 (1 entry, 1 announced)
  *Local Preference: 0
    Next hop type: Local

```

```

Next-hop reference count: 11
Interface: so-2/1/3.0
State: <Active NoReadvrt Int>
Local AS: 200
Age: 21:59
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

192.168.64.0/21 (1 entry, 1 announced)
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 3
Next hop: via fxp0.0, selected
State: <Active Int>
Local AS: 200
Age: 21:37:10
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

192.168.70.19/32 (1 entry, 1 announced)
*Local Preference: 0
Next hop type: Local
Next-hop reference count: 11
Interface: fxp0.0
State: <Active NoReadvrt Int>
Local AS: 200
Age: 21:37:10
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

```

show route active-path extensive

```

user@host> show route active-path extensive

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
10.255.70.19/32 (1 entry, 1 announced)
TSI:
IS-IS level 1, LSP fragment 0
IS-IS level 2, LSP fragment 0
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 3
Next hop: via lo0.0, selected
State: <Active Int>
Local AS: 200
Age: 21:39:47
Task: IF
Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3

AS path: I

10.255.71.50/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.255.71.50/32 -> {100.1.2.1}
IS-IS level 2, LSP fragment 0
*IS-IS Preference: 15
Level: 1
Next hop type: Router, Next hop index: 397

```

```

Next-hop reference count: 4
Next hop: 100.1.2.1 via so-2/1/3.0, selected
State: <Active Int>
Local AS: 200
Age: 24:08 Metric: 10
Task: IS-IS
Announcement bits (4): 0-KRT 2-IS-IS 5-Resolve tree 2 6-Resolve
tree 3
AS path: I

100.1.2.0/24 (1 entry, 1 announced)
TSI:
IS-IS level 1, LSP fragment 0
IS-IS level 2, LSP fragment 0
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 3
Next hop: via so-2/1/3.0, selected
State: <Active Int>
Local AS: 200
Age: 24:31
Task: IF
Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3
AS path: I

100.1.2.2/32 (1 entry, 1 announced)
*Local Preference: 0
Next hop type: Local
Next-hop reference count: 11
Interface: so-2/1/3.0
State: <Active NoReadvrt Int>
Local AS: 200
Age: 24:36
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

192.168.64.0/21 (1 entry, 1 announced)
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 3
Next hop: via fxp0.0, selected
State: <Active Int>
Local AS: 200
Age: 21:39:47
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

192.168.70.19/32 (1 entry, 1 announced)
*Local Preference: 0
Next hop type: Local
Next-hop reference count: 11
Interface: fxp0.0
State: <Active NoReadvrt Int>
Local AS: 200
Age: 21:39:47
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3

```

AS path: I

show route active-path terse

```
user@host> show route active-path terse
```

```
inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
```

```
+ = Active Route, - = Last Active, * = Both
```

A	Destination	P	Prf	Metric 1	Metric 2	Next hop	AS path
*	10.255.70.19/32	D	0			>1o0.0	
*	10.255.71.50/32	I	15	10		>100.1.2.1	
*	100.1.2.0/24	D	0			>so-2/1/3.0	
*	100.1.2.2/32	L	0			Local	
*	192.168.64.0/21	D	0			>fxp0.0	
*	192.168.70.19/32	L	0			Local	

show route all

List of Syntax	Syntax on page 3417 Syntax (EX Series Switches) on page 3417
Syntax	show route all <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	show route all
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display information about all routes in all routing tables, including private, or internal, tables.
Options	<p>none—Display information about all routes in all routing tables, including private, or internal, tables.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show route all on page 3417
Output Fields	In Junos OS Release 9.5 and later, only the output fields for the show route all command display all routing tables, including private, or hidden, routing tables. The output field table of the show route command does not display entries for private, or hidden, routing tables in Junos OS Release 9.5 and later.

Sample Output

show route all

The following example displays a snippet of output from the **show route** command and then displays the same snippet of output from the **show route all** command:

```

user@host> show route
mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
0          *[MPLS/0] 2d 02:24:39, metric 1
            Receive
1          *[MPLS/0] 2d 02:24:39, metric 1
            Receive
2          *[MPLS/0] 2d 02:24:39, metric 1
            Receive
800017     *[VPLS/7] 1d 14:00:16
            > via vt-3/2/0.32769, Pop
800018     *[VPLS/7] 1d 14:00:26
            > via vt-3/2/0.32772, Pop

```

```
user@host> show route all
mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
0          *[MPLS/0] 2d 02:19:12, metric 1
            Receive
1          *[MPLS/0] 2d 02:19:12, metric 1
            Receive
2          *[MPLS/0] 2d 02:19:12, metric 1
            Receive
800017     *[VPLS/7] 1d 13:54:49
            > via vt-3/2/0.32769, Pop
800018     *[VPLS/7] 1d 13:54:59
            > via vt-3/2/0.32772, Pop
vt-3/2/0.32769 [VPLS/7] 1d 13:54:49
              Unusable
vt-3/2/0.32772 [VPLS/7] 1d 13:54:59
              Unusable
```

show route aspath-regex

List of Syntax	Syntax on page 3419 Syntax (EX Series Switches) on page 3419
Syntax	<pre>show route aspath-regex <i>regular-expression</i> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switches)	<pre>show route aspath-regex <i>regular-expression</i></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	<p>Display the entries in the routing table that match the specified autonomous system (AS) path regular expression.</p>
Options	<p><i>regular-expression</i>—Regular expression that matches an entire AS path.</p> <p><i>logical-system (all <i>logical-system-name</i>)</i>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Additional Information	<p>You can specify a regular expression as:</p> <ul style="list-style-type: none"> • An individual AS number • A period wildcard used in place of an AS number • An AS path regular expression that is enclosed in parentheses <p>You also can include the operators described in the table of AS path regular expression operators in the <i>Junos Policy Framework Configuration Guide</i>. The following list summarizes these operators:</p> <ul style="list-style-type: none"> • <i>{m,n}</i>—At least <i>m</i> and at most <i>n</i> repetitions of the AS path term. • <i>{m}</i>—Exactly <i>m</i> repetitions of the AS path term. • <i>{m,}</i>—<i>m</i> or more repetitions of the AS path term. • <i>*</i>—Zero or more repetitions of an AS path term. • <i>+</i>—One or more repetitions of an AS path term. • <i>?</i>—Zero or one repetition of an AS path term. • <i>aspath_term aspath_term</i>—Match one of the two AS path terms. <p>When you specify more than one AS number or path term, or when you include an operator in the regular expression, enclose the entire regular expression in quotation marks. For example, to match any path that contains AS number 234, specify the following command:</p> <pre>show route aspath-regex ". * 234 . *"</pre>

Required Privilege Level	view
List of Sample Output	show route aspath-regex (Matching a Specific AS Number) on page 3420 show route aspath-regex (Matching Any Path with Two AS Numbers) on page 3420
Output Fields	For information about output fields, see the output field table for the show route command.

Sample Output

[show route aspath-regex \(Matching a Specific AS Number\)](#)

```
user@host> show route aspath-regex 65477
inet.0: 46411 destinations, 46411 routes (46409 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

111.222.1.0/25      *[BGP/170] 00:08:48, localpref 100, from 111.222.2.24
                   AS Path: [65477] ({65488 65535}) IGP
                   to 111.222.18.225 via fpa0.0(111.222.18.233)
111.222.1.128/25   *[IS-IS/15] 09:15:37, metric 37, tag 1
                   to 111.222.18.225 via fpa0.0(111.222.18.233)
                   [BGP/170] 00:08:48, localpref 100, from 111.222.2.24
                   AS Path: [65477] ({65488 65535}) IGP
                   to 111.222.18.225 via fpa0.0(111.222.18.233)
...
```

[show route aspath-regex \(Matching Any Path with Two AS Numbers\)](#)

```
user@host> show route aspath-regex ?.* 234 3561.*?

inet.0: 46351 destinations, 46351 routes (46349 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

9.20.0.0/17        *[BGP/170] 01:35:00, localpref 100, from 131.103.20.49
                   AS Path: [666] 234 3561 2685 2686 Incomplete
                   to 192.156.169.1 via 192.156.169.14(so-0/0/0)
12.10.231.0/24     *[BGP/170] 01:35:00, localpref 100, from 131.103.20.49
                   AS Path: [666] 234 3561 5696 7369 IGP
                   to 192.156.169.1 via 192.156.169.14(so-0/0/0)
24.64.32.0/19      *[BGP/170] 01:34:59, localpref 100, from 131.103.20.49
                   AS Path: [666] 234 3561 6327 IGP
                   to 192.156.169.1 via 192.156.169.14(so-0/0/0)
...
```

show route best

List of Syntax	Syntax on page 3421 Syntax (EX Series Switches) on page 3421
Syntax	show route best <i>destination-prefix</i> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	show route best <i>destination-prefix</i> <brief detail extensive terse>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the route in the routing table that is the best route to the specified address or range of addresses. The best route is the longest matching route.
Options	brief detail extensive terse —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief. <i>destination-prefix</i> —Address or range of addresses. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show route best on page 3421 show route best detail on page 3422 show route best extensive on page 3423 show route best terse on page 3423
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route best

```

user@host> show route best 10.255.70.103
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.255.70.103/32    *[OSPF/10] 1d 13:19:20, metric 2
                  > to 10.31.1.6 via ge-3/1/0.0
                  via so-0/3/0.0

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.255.70.103/32    *[RSVP/7] 1d 13:20:13, metric 2

```

```

> via so-0/3/0.0, label-switched-path green-r1-r3

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.0.0.0/8          *[Direct/0] 2d 01:43:34
                    > via fxp2.0
                    [Direct/0] 2d 01:43:34
                    > via fxp1.0

```

show route best detail

```

user@host> show route best 10.255.70.103 detail
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
10.255.70.103/32 (1 entry, 1 announced)
  *OSPF   Preference: 10
          Next-hop reference count: 9
          Next hop: 10.31.1.6 via ge-3/1/0.0, selected
          Next hop: via so-0/3/0.0
          State: <Active Int>
          Local AS: 69
          Age: 1d 13:20:06      Metric: 2
          Area: 0.0.0.0
          Task: OSPF
          Announcement bits (2): 0-KRT 3-Resolve tree 2
          AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
10.255.70.103/32 (1 entry, 1 announced)
  State: <FlashAll>
  *RSVP   Preference: 7
          Next-hop reference count: 5
          Next hop: via so-0/3/0.0 weight 0x1, selected
          Label-switched-path green-r1-r3
          Label operation: Push 100016
          State: <Active Int>
          Local AS: 69
          Age: 1d 13:20:59      Metric: 2
          Task: RSVP
          Announcement bits (1): 1-Resolve tree 2
          AS path: I

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
10.0.0.0/8 (2 entries, 0 announced)
  *Direct Preference: 0
          Next hop type: Interface
          Next-hop reference count: 1
          Next hop: via fxp2.0, selected
          State: <Active Int>
          Age: 2d 1:44:20
          Task: IF
          AS path: I
  Direct Preference: 0
          Next hop type: Interface
          Next-hop reference count: 1
          Next hop: via fxp1.0, selected
          State: <NotBest Int>
          Inactive reason: No difference
          Age: 2d 1:44:20

```

Task: IF
AS path: I

show route best extensive

The output for the **show route best extensive** command is identical to that for the **show route best detail** command. For sample output, see [show route best detail on page 3422](#).

show route best terse

```
user@host> show route best 10.255.70.103 terse
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1   Metric 2   Next hop      AS path
* 10.255.70.103/32  0 10           2           >10.31.1.6
                                   so-0/3/0.0

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1   Metric 2   Next hop      AS path
* 10.255.70.103/32  R   7           2           >so-0/3/0.0

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1   Metric 2   Next hop      AS path
* 10.0.0.0/8        D   0           0           >fxp2.0
                    D   0           0           >fxp1.0
```

show route brief

List of Syntax	Syntax on page 3424 Syntax (EX Series Switches) on page 3424
Syntax	show route brief <destination-prefix> <logical-system (all logical-system-name)>
Syntax (EX Series Switches)	show route brief <destination-prefix>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display brief information about the active entries in the routing tables.
Options	none —Display all active entries in the routing table. destination-prefix —(Optional) Display active entries for the specified address or range of addresses. logical-system (all logical-system-name) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show route brief on page 3424
Output Fields	For information about output fields, see the Output Field table of the show route command.

Sample Output

show route brief

```

user@host> show route brief
inet.0: 10 destinations, 10 routes (9 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 1w5d 20:30:29
                   Discard
10.255.245.51/32   *[Direct/0] 2w4d 13:11:14
                   > via lo0.0
172.16.0.0/12      *[Static/5] 2w4d 13:11:14
                   > to 192.168.167.254 via fxp0.0
192.168.0.0/18     *[Static/5] 1w5d 20:30:29
                   > to 192.168.167.254 via fxp0.0
192.168.40.0/22    *[Static/5] 2w4d 13:11:14
                   > to 192.168.167.254 via fxp0.0
192.168.64.0/18    *[Static/5] 2w4d 13:11:14
                   > to 192.168.167.254 via fxp0.0
192.168.164.0/22   *[Direct/0] 2w4d 13:11:14
                   > via fxp0.0

```



```
192.168.164.51/32 *[Local/0] 2w4d 13:11:14
                  Local via fxp0.0
207.17.136.192/32 *[Static/5] 2w4d 13:11:14
                  > to 192.168.167.254 via fxp0.0
green.inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
100.101.0.0/16    *[Direct/0] 1w5d 20:30:28
                  > via fe-0/0/3.0
100.101.2.3/32   *[Local/0] 1w5d 20:30:28
                  Local via fe-0/0/3.0
224.0.0.5/32     *[OSPF/10] 1w5d 20:30:29, metric 1
                  MultiRecv
```

show route community

List of Syntax	Syntax on page 3426 Syntax (EX Series Switches) on page 3426
Syntax	<code>show route community <i>as-number:community-value</i></code> <code><brief detail extensive terse></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switches)	<code>show route community <i>as-number:community-value</i></code> <code><brief detail extensive terse></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the route entries in each routing table that are members of a Border Gateway Protocol (BGP) community.
Options	<p><i>as-number:community-value</i>—One or more community identifiers. <i>as-number</i> is the AS number, and <i>community-value</i> is the community identifier. When you specify more than one community identifier, enclose the identifiers in double quotation marks. Community identifiers can include wildcards.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Additional Information	Specifying the community option displays all routes matching the community found within the routing table. The community option does not limit the output to only the routes being advertised to the neighbor after any egress routing policy.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show route detail on page 3435
List of Sample Output	show route community on page 3426
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route community

```
user@host> show route community 234:80
inet.0: 46511 destinations, 46511 routes (46509 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
4.0.0.0/8      *[BGP/170] 03:33:07, localpref 100, from 131.103.20.49
                AS Path: {666} 234 2548 1 IGP
                to 192.156.169.1 via 192.156.169.14(so-0/0/0)
6.0.0.0/8      *[BGP/170] 03:33:07, localpref 100, from 131.103.20.49
                AS Path: {666} 234 2548 568 721 Incomplete
                to 192.156.169.1 via 192.156.169.14(so-0/0/0)
9.2.0.0/16     *[BGP/170] 03:33:06, localpref 100, from 131.103.20.49
                AS Path: {666} 234 2548 1673 1675 1747 IGP
                to 192.156.169.1 via 192.156.169.14(so-0/0/0)
```

show route community-name

List of Syntax	Syntax on page 3428 Syntax (EX Series Switches) on page 3428
Syntax	show route community-name <i>community-name</i> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	show route community-name <i>community-name</i> <brief detail extensive terse>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the route entries in each routing table that are members of a Border Gateway Protocol (BGP) community, specified by a community name.
Options	<i>community-name</i> —Name of the community. brief detail extensive terse —(Optional) Display the specified level of output. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show route community-name on page 3428
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route community-name

```

user@host> show route community-name red-com
inet.0: 17 destinations, 17 routes (16 active, 0 holddown, 1 hidden)

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

instance1.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.245.212/32  *[BGP/170] 00:04:40, localpref 100, from 10.255.245.204
                   AS path: 300 I
                   > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix
20.20.20.20/32    *[BGP/170] 00:04:40, localpref 100, from 10.255.245.204
                   AS path: I
                   > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix
100.1.4.0/24     *[BGP/170] 00:04:40, localpref 100, from 10.255.245.204

```

```

AS path: I
> to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

bgp.l3vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.245.204:10:10.255.245.212/32
    *[BGP/170] 00:06:40, localpref 100, from 10.255.245.204
    AS path: 300 I
    > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix
10.255.245.204:10:20.20.20.20/32
    *[BGP/170] 00:36:02, localpref 100, from 10.255.245.204
    AS path: I
    > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix
10.255.245.204:10:100.1.4.0/24
    *[BGP/170] 00:36:02, localpref 100, from 10.255.245.204
    AS path: I
    > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

instance1.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

show route damping

List of Syntax	Syntax on page 3430 Syntax (EX Series Switch and QFX Series) on page 3430
Syntax	<code>show route damping (decayed history suppressed)</code> <code><brief detail extensive terse></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switch and QFX Series)	<code>show route damping (decayed history suppressed)</code> <code><brief detail extensive terse></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display the BGP routes for which updates might have been reduced because of route flap damping.
Options	brief detail extensive terse —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief. decayed —Display route damping entries that might no longer be valid, but are not suppressed. history —Display entries that have already been withdrawn, but have been logged. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. suppressed —Display entries that have been suppressed and are no longer being installed into the forwarding table or exported by routing protocols.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• clear bgp damping on page 4092• show policy damping on page 4130
List of Sample Output	show route damping decayed detail on page 3433 show route damping history on page 3434 show route damping history detail on page 3434
Output Fields	Table 298 lists the output fields for the show route damping command. Output fields are listed in the approximate order in which they appear.

Table 298: show route damping Output Fields

Field Name	Field Description	Level of Output
<i>routing-table-name</i>	Name of the routing table—for example, inet.0 .	All levels
destinations	Number of destinations for which there are routes in the routing table.	All levels
number routes	Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> • active • holdddown (routes that are in a pending state before being declared inactive) • hidden (the routes are not used because of a routing policy) 	All levels
destination-prefix (entry, announced)	Destination prefix. The entry value is the number of routes for this destination, and the announced value is the number of routes being announced for this destination.	detail extensive
[protocol, preference]	Protocol from which the route was learned and the preference value for the route. <ul style="list-style-type: none"> • +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • -—A hyphen indicates the last active route. • *—An asterisk indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route. <p>In every routing metric except for the BGP LocalPref attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the LocalPref value in the Preference2 field. For example, if the LocalPref value for Route 1 is 100, the Preference2 value is -101. If the LocalPref value for Route 2 is 155, the Preference2 value is -156. Route 2 is preferred because it has a higher LocalPref value and a lower Preference2 value.</p>	All levels
Next-hop reference count	Number of references made to the next hop.	detail extensive
Source	IP address of the route source.	detail extensive
Next hop	Network layer address of the directly reachable neighboring system.	detail extensive
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word Selected .	detail extensive
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to derive a forwarding next hop.	detail extensive
Indirect next hop	Index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.	detail extensive
State	Flags for this route. For a description of possible values for this field, see the output field table for the show route detail command.	detail extensive

Table 298: show route damping Output Fields (*continued*)

Field Name	Field Description	Level of Output
Local AS	AS number of the local routing device.	detail extensive
Peer AS	AS number of the peer routing device.	detail extensive
Age	How long the route has been known.	detail extensive
Metric	Metric for the route.	detail extensive
Task	Name of the protocol that has added the route.	detail extensive
Announcement bits	List of protocols that announce this route. <i>n-Resolve inet</i> indicates that the route is used for route resolution for next hops found in the routing table. <i>n</i> is an index used by Juniper Networks customer support only.	detail extensive
AS path	<p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> I—IGP. E—EGP. ?—Incomplete; typically, the AS path was aggregated. <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> []—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device or if AS path prepending is configured. { }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. ()—Parentheses enclose a confederation. ([])—Parentheses and brackets enclose a confederation set. <p>NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p>	All levels
to	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.	brief none
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word Selected .	brief none
Communities	Community path attribute for the route. See the output field table for the show route detail command.	detail extensive
Localpref	Local preference value included in the route.	All levels
Router ID	BGP router ID as advertised by the neighbor in the open message.	detail extensive

Table 298: show route damping Output Fields (*continued*)

Field Name	Field Description	Level of Output
Merit (last update/now)	Last updated and current figure-of-merit value.	detail extensive
damping-parameters	Name that identifies the damping parameters used, which is defined in the damping statement at the [edit policy-options] hierarchy level.	detail extensive
Last update	Time of most recent change in path attributes.	detail extensive
First update	Time of first change in path attributes, which started the route damping process.	detail extensive
Flaps	Number of times the route has gone up or down or its path attributes have changed.	detail extensive
Suppressed	(suppressed keyword only) This route is currently suppressed. A suppressed route does not appear in the forwarding table and routing protocols do not export it.	All levels
Reusable in	(suppressed keyword only) Time when a suppressed route will again be available.	All levels
Preference will be	(suppressed keyword only) Preference value that will be applied to the route when it is again active.	All levels

Sample Output

show route damping decayed detail

```

user@host> show route damping decayed detail
inet.0: 173319 destinations, 1533668 routes (172625 active, 4 holddown, 108083
hidden)
10.0.111.0/24 (7 entries, 1 announced)
  *BGP      Preference: 170/-101
            Next-hop reference count: 151973
            Source: 172.23.2.129
            Next hop: via so-1/2/0.0
            Next hop: via so-5/1/0.0, selected
            Next hop: via so-6/0/0.0
            Protocol next hop: 172.23.2.129
            Indirect next hop: 89a1a00 264185
            State: <Active Ext>
            Local AS: 65000 Peer AS: 65490
            Age: 3:28      Metric2: 0
            Task: BGP_65490.172.23.2.129+179
            Announcement bits (6): 0-KRT 1-RT 4-KRT 5-BGP.0.0.0.0+179

  6-Resolve tree 2 7-Resolve tree 3
    AS path: 65490 65520 65525 65525 65525 65525 I ()
    Communities: 65501:390 65501:2000 65501:3000 65504:701
    Localpref: 100
    Router ID: 172.23.2.129
    Merit (last update/now): 1934/1790
    damping-parameters: damping-high

```

```
      Last update:      00:03:28 First update:      00:06:40
      Flaps: 2
```

show route damping history

```
user@host> show route damping history
inet.0: 173320 destinations, 1533529 routes (172624 active, 6 holddown, 108122
hidden)
+ = Active Route, - = Last Active, * = Both

10.108.0.0/15      [BGP ] 2d 22:47:58, localpref 100
                  AS path: 65220 65501 65502 I
                  > to 192.168.60.85 via so-3/1/0.0
```

show route damping history detail

```
user@host> show route damping history detail
inet.0: 173319 destinations, 1533435 routes (172627 active, 2 holddown, 108105
hidden)
10.108.0.0/15 (3 entries, 1 announced)
    BGP                /-101
        Next-hop reference count: 69058
        Source: 192.168.60.85
        Next hop: 192.168.60.85 via so-3/1/0.0, selected
        State: <Hidden Ext>
        Inactive reason: Unusable path
        Local AS: 65000 Peer AS: 65220
        Age: 2d 22:48:10
        Task: BGP_65220.192.168.60.85+179
        AS path: 65220 65501 65502 I ()
        Communities: 65501:390 65501:2000 65501:3000 65504:3561
        Localpref: 100
        Router ID: 192.168.80.25
        Merit (last update/now): 1000/932
        damping-parameters: set-normal
        Last update:      00:01:05 First update:      00:01:05
        Flaps: 1
```

show route detail

List of Syntax	Syntax on page 3435 Syntax (EX Series Switches) on page 3435
Syntax	show route detail <destination-prefix> <logical-system (all logical-system-name)>
Syntax (EX Series Switches)	show route detail <destination-prefix>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display detailed information about the active entries in the routing tables.
Options	<p>none—Display all active entries in the routing table on all systems.</p> <p>destination-prefix—(Optional) Display active entries for the specified address or range of addresses.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show route detail on page 3444 show route detail (with BGP Multipath) on page 3450 show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 3451 show route label detail (Multipoint LDP with Multicast-Only Fast Reroute) on page 3451
Output Fields	Table 299 describes the output fields for the show route detail command. Output fields are listed in the approximate order in which they appear.

Table 299: show route detail Output Fields

Field Name	Field Description
<i>routing-table-name</i>	Name of the routing table (for example, inet.0).
<i>number destinations</i>	Number of destinations for which there are routes in the routing table.
<i>number routes</i>	Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> active (routes that are active) holddown (routes that are in the pending state before being declared inactive) hidden (routes that are not used because of a routing policy)

Table 299: show route detail Output Fields (*continued*)

Field Name	Field Description
<i>route-destination</i> (entry, announced)	<p>Route destination (for example:10.0.0.1/24). The entry value is the number of routes for this destination, and the announced value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:</p> <ul style="list-style-type: none"> • MPLS-label (for example, 80001). • interface-name (for example, ge-1/0/2). • neighbor-address:control-word-status:encapsulation type:vc-id:source (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96). <ul style="list-style-type: none"> • neighbor-address—Address of the neighbor. • control-word-status—Whether the use of the control word has been negotiated for this virtual circuit: NoCtrlWord or CtrlWord. • encapsulation type—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport. • vc-id—Virtual circuit identifier. • source—Source of the advertisement: Local or Remote. • source—Source of the advertisement: Local or Remote.
label stacking	<p>(Next-to-the-last-hop routing device for MPLS only) Depth of the MPLS label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels.</p> <ul style="list-style-type: none"> • S=0 route indicates that a packet with an incoming label stack depth of 2 or more exits this routing device with one fewer label (the label-popping operation is performed). • If there is no S= information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).
[<i>protocol, preference</i>]	<p>Protocol from which the route was learned and the preference value for the route.</p> <ul style="list-style-type: none"> • +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • - —A hyphen indicates the last active route. • *—An asterisk indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route. <p>In every routing metric except for the BGP LocalPref attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the LocalPref value in the Preference2 field. For example, if the LocalPref value for Route 1 is 100, the Preference2 value is -101. If the LocalPref value for Route 2 is 155, the Preference2 value is -156. Route 2 is preferred because it has a higher LocalPref value and a lower Preference2 value.</p>
Level	<p>(IS-IS only). In IS-IS, a single AS can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.</p>
Route Distinguisher	IP subnet augmented with a 64-bit prefix.
PMSI	Provider multicast service interface (MVPN routing table).

Table 299: show route detail Output Fields (*continued*)

Field Name	Field Description
Next-hop type	Type of next hop. For a description of possible values for this field, see Table 300 .
Next-hop reference count	Number of references made to the next hop.
Flood nexthop branches exceed maximum message	Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.
Source	IP address of the route source.
Next hop	Network layer address of the directly reachable neighboring system.
via	<p>Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word Selected. This field can also contain the following information:</p> <ul style="list-style-type: none"> • Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.
Label-switched-path lsp-path-name	Name of the LSP used to reach the next hop.
Label operation	MPLS label and operation occurring at this routing device. The operation can be pop (where a label is removed from the top of the stack), push (where another label is added to the label stack), or swap (where a label is replaced by another label).
Interface	(Local only) Local interface name.
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to derive a forwarding next hop.
Indirect next hop	Index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.
State	State of the route (a route can be in more than one state). See Table 301 .
Local AS	AS number of the local routing device.
Age	How long the route has been known.
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.

Table 299: show route detail Output Fields (*continued*)

Field Name	Field Description
Metric	Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
MED-plus-IGP	Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.
TTL-Action	For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signaled and LDP-signaled LSPs or for specific VRF routing instances. For sample output, see show route table .
Task	Name of the protocol that has added the route.
Announcement bits	The number of BGP peers or protocols to which Junos OS has announced this route, followed by the list of the recipients of the announcement. Junos OS can also announce the route to the KRT for installing the route into the Packet Forwarding Engine, to a resolve tree, a L2 VC, or even a VPN. For example, <i>n-Resolve inet</i> indicates that the specified route is used for route resolution for next hops found in the routing table. <ul style="list-style-type: none"> <i>n</i>—An index used by Juniper Networks customer support only.
AS path	<p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> I—IGP. E—EGP. Recorded—The AS path is recorded by the sample process (sampled). ?—Incomplete; typically, the AS path was aggregated. <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> []—Brackets enclose the number that precedes the AS path. This number represents the number of ASs present in the AS path, when calculated as defined in RFC 4271. This value is used in the AS-path merge process, as defined in RFC 4893. []—If more than one AS number is configured on the routing device, or if AS path prepending is configured, brackets enclose the local AS number associated with the AS path. { }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. ()—Parentheses enclose a confederation. ([])—Parentheses and brackets enclose a confederation set. <p>NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p>

Table 299: show route detail Output Fields (*continued*)

Field Name	Field Description
validation-state	<p>(BGP-learned routes) Validation status of the route:</p> <ul style="list-style-type: none"> • Invalid—Indicates that the prefix is found, but either the corresponding AS received from the EBGp peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • Unknown—Indicates that the prefix is not among the prefixes or prefix ranges in the database. • Unverified—Indicates that the origin of the prefix is not verified against the database. This is because the database got populated and the validation is not called for in the BGP import policy, although origin validation is enabled, or the origin validation is not enabled for the BGP peers. • Valid—Indicates that the prefix and autonomous system pair are found in the database.
FECs bound to route	Point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.
Primary Upstream	When multipoint LDP with multicast-only fast reroute (MoFRR) is configured, the primary upstream path. MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path.
RPF Nexthops	When multipoint LDP with MoFRR is configured, the reverse-path forwarding (RPF) next-hop information. Data packets are received from both the primary path and the secondary paths. The redundant packets are discarded at topology merge points due to the RPF checks.
Label	Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.
weight	Value used to distinguish MoFRR primary and backup routes. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.
VC Label	MPLS label assigned to the Layer 2 circuit virtual connection.
MTU	Maximum transmission unit (MTU) of the Layer 2 circuit.
VLAN ID	VLAN identifier of the Layer 2 circuit.
Prefixes bound to route	Forwarding equivalent class (FEC) bound to this route. Applicable only to routes installed by LDP.
Communities	Community path attribute for the route. See Table 302 for all possible values for this field.
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).
control flags	Control flags: none or Site Down .
mtu	Maximum transmission unit (MTU) information.
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.
status vector	Layer 2 VPN and VPLS network layer reachability information (NLRI).

Table 299: show route detail Output Fields (*continued*)

Field Name	Field Description
Accepted Multipath	Current active path when BGP multipath is configured.
Accepted LongLivedStale	The LongLivedStale flag indicates that the route was marked LLGR-stale by this router, as part of the operation of LLGR receiver mode. Either this flag or the LongLivedStaleImport flag may be displayed for a route. Neither of these flags are displayed at the same time as the Stale (ordinary GR stale) flag.
Accepted LongLivedStaleImport	<p>The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy. Either this flag or the LongLivedStale flag may be displayed for a route. Neither of these flags are displayed at the same time as the Stale (ordinary GR stale) flag.</p> <p>Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and import into the inet.0 routing table</p>
ImportAccepted LongLivedStaleImport	<p>Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and imported into the inet.0 routing table</p> <p>The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy.</p>
Accepted MultipathContrib	Path currently contributing to BGP multipath.
Localpref	Local preference value included in the route.
Router ID	BGP router ID as advertised by the neighbor in the open message.
Primary Routing Table	In a routing table group, the name of the primary routing table in which the route resides.
Secondary Tables	In a routing table group, the name of one or more secondary tables in which the route resides.

Table 300 describes all possible values for the Next-hop Types output field.

Table 300: Next-Hop Types Output Field Values

Next-Hop Type	Description
Broadcast (bcast)	Broadcast next hop.
Deny	Deny next hop.
Discard	Discard next hop.
Flood	Flood next hop. Consists of components called branches, up to a maximum of 32 branches. Each flood next-hop branch sends a copy of the traffic to the forwarding interface. Used by P2MP RSVP, P2MP LDP, P2MP CCC, and multicast.
Hold	Next hop is waiting to be resolved into a unicast or multicast type.

Table 300: Next-Hop Types Output Field Values (*continued*)

Next-Hop Type	Description
Indexed (idxd)	Indexed next hop.
Indirect (indr)	Used with applications that have a protocol next hop address that is remote. You are likely to see this next-hop type for internal BGP (IBGP) routes when the BGP next hop is a BGP neighbor that is not directly connected.
Interface	Used for a network address assigned to an interface. Unlike the router next hop, the interface next hop does not reference any specific node on the network.
Local (locl)	Local address on an interface. This next-hop type causes packets with this destination address to be received locally.
Multicast (mcst)	Wire multicast next hop (limited to the LAN).
Multicast discard (mdsc)	Multicast discard.
Multicast group (mgrp)	Multicast group member.
Receive (recv)	Receive.
Reject (rjct)	Discard. An ICMP unreachable message was sent.
Resolve (rslv)	Resolving next hop.
Routed multicast (mcrtr)	Regular multicast next hop.
Router	<p>A specific node or set of nodes to which the routing device forwards packets that match the route prefix.</p> <p>To qualify as next-hop type router, the route must meet the following criteria:</p> <ul style="list-style-type: none"> • Must not be a direct or local subnet for the routing device. • Must have a next hop that is directly connected to the routing device.
Table	Routing table next hop.
Unicast (ucst)	Unicast.
Unilist (ulst)	List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.

Table 301 describes all possible values for the State output field. A route can be in more than one state (for example, <Active NoReadvrt Int Ext>).

Table 301: State Output Field Values

Value	Description
Accounting	Route needs accounting.
Active	Route is active.
Always Compare MED	Path with a lower multiple exit discriminator (MED) is available.
AS path	Shorter AS path is available.
Cisco Non-deterministic MED selection	Route is a clone.
Clone	Cisco nondeterministic MED is enabled and a path with a lower MED is available.
Cluster list length	Length of cluster list sent by the route reflector.
Delete	Route has been deleted.
Ex	Exterior route.
Ext	BGP route received from an external BGP neighbor.
FlashAll	Forces all protocols to be notified of a change to any route, active or inactive, for a prefix. When not set, protocols are informed of a prefix only when the active route changes.
Hidden	Route not used because of routing policy.
IfCheck	Route needs forwarding RPF check.
IGP metric	Path through next hop with lower IGP metric is available.
Inactive reason	Flags for this route, which was not selected as best for a particular destination.
Initial	Route being added.
Int	Interior route.
Int Ext	BGP route received from an internal BGP peer or a BGP confederation peer.
Interior > Exterior > Exterior via Interior	Direct, static, IGP, or EBGp path is available.
Local Preference	Path with a higher local preference value is available.

Table 301: State Output Field Values (*continued*)

Value	Description
Martian	Route is a martian (ignored because it is obviously invalid).
MartianOK	Route exempt from martian filtering.
Next hop address	Path with lower metric next hop is available.
No difference	Path from neighbor with lower IP address is available.
NoReadvrt	Route not to be advertised.
NotBest	Route not chosen because it does not have the lowest MED.
Not Best in its group	Incoming BGP AS is not the best of a group (only one AS can be the best).
NotInstall	Route not to be installed in the forwarding table.
Number of gateways	Path with a greater number of next hops is available.
Origin	Path with a lower origin code is available.
Pending	Route pending because of a hold-down configured on another route.
Release	Route scheduled for release.
RIB preference	Route from a higher-numbered routing table is available.
Route Distinguisher	64-bit prefix added to IP subnets to make them unique.
Route Metric or MED comparison	Route with a lower metric or MED is available.
Route Preference	Route with lower preference value is available
Router ID	Path through a neighbor with lower ID is available.
Secondary	Route not a primary route.
Unusable path	Path is not usable because of one of the following conditions: <ul style="list-style-type: none"> • The route is damped. • The route is rejected by an import policy. • The route is unresolved.
Update source	Last tiebreaker is the lowest IP address value.

Table 302 describes the possible values for the Communities output field.

Table 302: Communities Output Field Values

Value	Description
<i>area-number</i>	4 bytes, encoding a 32-bit area number. For AS-external routes, the value is 0 . A nonzero value identifies the route as internal to the OSPF domain, and as within the identified area. Area numbers are relative to a particular OSPF domain.
bandwidth: local AS number:link-bandwidth-number	Link-bandwidth community value used for unequal-cost load balancing. When BGP has several candidate paths available for multipath purposes, it does not perform unequal-cost load balancing according to the link-bandwidth community unless all candidate paths have this attribute.
domain-id	Unique configurable number that identifies the OSPF domain.
domain-id-vendor	Unique configurable number that further identifies the OSPF domain.
<i>link-bandwidth-number</i>	Link-bandwidth number: from 0 through 4,294,967,295 (bytes per second).
<i>local AS number</i>	Local AS number: from 1 through 65,535 .
<i>options</i>	1 byte. Currently this is only used if the route type is 5 or 7 . Setting the least significant bit in the field indicates that the route carries a type 2 metric.
origin	(Used with VPNs) Identifies where the route came from.
<i>ospf-route-type</i>	1 byte, encoded as 1 or 2 for intra-area routes (depending on whether the route came from a type 1 or a type 2 LSA); 3 for summary routes; 5 for external routes (area number must be 0); 7 for NSSA routes; or 129 for sham link endpoint addresses.
route-type-vendor	Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute 0x8000 . The format is area-number:ospf-route-type:options .
rte-type	Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute 0x0306 . The format is area-number:ospf-route-type:options .
target	Defines which VPN the route participates in; target has the format 32-bit IP address:16-bit number . For example, 10.19.0.0:100.
unknown IANA	Incoming IANA codes with a value between 0x1 and 0x7fff . This code of the BGP extended community attribute is accepted, but it is not recognized.
unknown OSPF vendor community	Incoming IANA codes with a value above 0x8000 . This code of the BGP extended community attribute is accepted, but it is not recognized.

Sample Output

show route detail

```
user@host> show route detail
```

```
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
```

```

10.10.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 29
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 1:31:43
    Task: RT
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I

10.31.1.0/30 (2 entries, 1 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 2
    Next hop: via so-0/3/0.0, selected
    State: <Active Int>
    Local AS: 69
    Age: 1:30:17
    Task: IF
    Announcement bits (1): 3-Resolve tree 2
    AS path: I
  OSPF Preference: 10
    Next-hop reference count: 1
    Next hop: via so-0/3/0.0, selected
    State: <Int>
    Inactive reason: Route Preference
    Local AS: 69
    Age: 1:30:17 Metric: 1
    Area: 0.0.0.0
    Task: OSPF
    AS path: I

10.31.1.1/32 (1 entry, 1 announced)
  *Local Preference: 0
    Next hop type: Local
    Next-hop reference count: 7
    Interface: so-0/3/0.0
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:30:20
    Task: IF
    Announcement bits (1): 3-Resolve tree 2
    AS path: I

...

10.31.2.0/30 (1 entry, 1 announced)
  *OSPF Preference: 10
    Next-hop reference count: 9
    Next hop: via so-0/3/0.0
    Next hop: 10.31.1.6 via ge-3/1/0.0, selected
    State: <Active Int>
    Local AS: 69
    Age: 1:29:56 Metric: 2
    Area: 0.0.0.0
    Task: OSPF
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I

...

```

```
224.0.0.2/32 (1 entry, 1 announced)
  *PIM    Preference: 0
          Next-hop reference count: 18
          State: <Active NoReadvrt Int>
          Local AS:    69
          Age: 1:31:45
          Task: PIM Recv
          Announcement bits (2): 0-KRT 3-Resolve tree 2
          AS path: I

...

224.0.0.22/32 (1 entry, 1 announced)
  *IGMP   Preference: 0
          Next-hop reference count: 18
          State: <Active NoReadvrt Int>
          Local AS:    69
          Age: 1:31:43
          Task: IGMP
          Announcement bits (2): 0-KRT 3-Resolve tree 2
          AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.255.70.103/32 (1 entry, 1 announced)
  State: <FlashAll>
  *RSVP   Preference: 7
          Next-hop reference count: 6
          Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
          Label-switched-path green-r1-r3
          Label operation: Push 100096
          State: <Active Int>
          Local AS:    69
          Age: 1:25:49   Metric: 2
          Task: RSVP
          Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
          AS path: I

10.255.71.238/32 (1 entry, 1 announced)
  State: <FlashAll>
  *RSVP   Preference: 7
          Next-hop reference count: 6
          Next hop: via so-0/3/0.0 weight 0x1, selected
          Label-switched-path green-r1-r2
          State: <Active Int>
          Local AS:    69
          Age: 1:25:49   Metric: 1
          Task: RSVP
          Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
          AS path: I

private__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

47.0005.80ff.f800.0000.0108.0001.0102.5507.1052/152 (1 entry, 0 announced)
  *Direct Preference: 0
          Next hop type: Interface
          Next-hop reference count: 1
          Next hop: via lo0.0, selected
```

```

        State: <Active Int>
        Local AS: 69
        Age: 1:31:44
        Task: IF
        AS path: I

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
0 (1 entry, 1 announced)
    *MPLS Preference: 0
        Next hop type: Receive
        Next-hop reference count: 6
        State: <Active Int>
        Local AS: 69
        Age: 1:31:45 Metric: 1
        Task: MPLS
        Announcement bits (1): 0-KRT
        AS path: I

...

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
299776 (1 entry, 1 announced)
TSI:
KRT in-kerne1 299776 /52 -> {Flood}
    *RSVP Preference: 7
        Next hop type: Flood
        Next-hop reference count: 130
        Flood nexthop branches exceed maximum
        Address: 0x8ea65d0

...

299840 (1 entry, 1 announced)
TSI:
KRT in-kerne1 299840 /52 -> {indirect(1048575)}
    *RSVP Preference: 7/2
        Next hop type: Flood
        Address: 0x9174a30
        Next-hop reference count: 4
        Next hop type: Router, Next hop index: 798
        Address: 0x9174c28
        Next-hop reference count: 2
        Next hop: 8.0.0.2 via lt-1/2/0.9 weight 0x1
        Label-switched-path R2-to-R4-2p2mp
        Label operation: Pop
        Next hop type: Router, Next hop index: 1048574
        Address: 0x92544f0
        Next-hop reference count: 2
        Next hop: 7.0.0.2 via lt-1/2/0.7 weight 0x1
        Label-switched-path R2-to-R200-p2mp
        Label operation: Pop
        Next hop: 6.0.0.2 via lt-1/2/0.5 weight 0x8001
        Label operation: Pop
        State: <Active Int>
        Age: 1:29 Metric: 1
        Task: RSVP
        Announcement bits (1): 0-KRT
        AS path: I...

800010 (1 entry, 1 announced)
    *VPLS Preference: 7
        Next-hop reference count: 2

```

```
Next hop: via vt-3/2/0.32769, selected
Label operation: Pop
State: <Active Int>
Age: 1:29:30
Task: Common L2 VC
Announcement bits (1): 0-KRT
AS path: I

vt-3/2/0.32769 (1 entry, 1 announced)
  *VPLS Preference: 7
    Next-hop reference count: 2
    Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
    Label-switched-path green-r1-r3
    Label operation: Push 800012, Push 100096(top)
    Protocol next hop: 10.255.70.103
    Push 800012
    Indirect next hop: 87272e4 1048574
    State: <Active Int>
    Age: 1:29:30 Metric2: 2
    Task: Common L2 VC
    Announcement bits (2): 0-KRT 1-Common L2 VC
    AS path: I
    Communities: target:11111:1 Layer2-info: encaps:VPLS,
    control flags:, mtu: 0

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

abcd::10:255:71:52/128 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
    State: <Active Int>
    Local AS: 69
    Age: 1:31:44
    Task: IF
    AS path: I

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:31:44
    Task: IF
    AS path: I

ff02::2/128 (1 entry, 1 announced)
  *PIM Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:31:45
    Task: PIM Recv6
    Announcement bits (1): 0-KRT
    AS path: I

ff02::d/128 (1 entry, 1 announced)
  *PIM Preference: 0
```



```

Next-hop reference count: 18
State: <Active NoReadvrt Int>
Local AS: 69
Age: 1:31:45
Task: PIM Recv6
Announcement bits (1): 0-KRT
AS path: I

ff02::16/128 (1 entry, 1 announced)
*MLD Preference: 0
Next-hop reference count: 18
State: <Active NoReadvrt Int>
Local AS: 69
Age: 1:31:43
Task: MLD
Announcement bits (1): 0-KRT
AS path: I

private.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 1
Next hop: via lo0.16385, selected
State: <Active NoReadvrt Int>
Age: 1:31:44
Task: IF
AS path: I

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

10.255.70.103:1:3:1/96 (1 entry, 1 announced)
*BGP Preference: 170/-101
Route Distinguisher: 10.255.70.103:1
Next-hop reference count: 7
Source: 10.255.70.103
Protocol next hop: 10.255.70.103
Indirect next hop: 2 no-forward
State: <Secondary Active Int Ext>
Local AS: 69 Peer AS: 69
Age: 1:25:49 Metric2: 1
AIGP 210
Task: BGP_69.10.255.70.103+179
Announcement bits (1): 0-green-l2vpn
AS path: I
Communities: target:11111:1 Layer2-info: encaps:VPLS,
control flags:, mtu: 0
Label-base: 800008, range: 8
Localpref: 100
Router ID: 10.255.70.103
Primary Routing Table bgp.l2vpn.0

10.255.71.52:1:1:1/96 (1 entry, 1 announced)
*L2VPN Preference: 170/-1
Next-hop reference count: 5
Protocol next hop: 10.255.71.52
Indirect next hop: 0 -
State: <Active Int Ext>
Age: 1:31:40 Metric2: 1
Task: green-l2vpn

```

```

Announcement bits (1): 1-BGP.0.0.0+179
AS path: I
Communities: Layer2-info: encaps:VPLS, control flags:Site-Down,
mtu: 0
Label-base: 800016, range: 8, status-vector: 0x9F

10.255.71.52:1:5:1/96 (1 entry, 1 announced)
  *L2VPN Preference: 170/-101
    Next-hop reference count: 5
    Protocol next hop: 10.255.71.52
    Indirect next hop: 0 -
    State: <Active Int Ext>
    Age: 1:31:40 Metric2: 1
    Task: green-l2vpn
    Announcement bits (1): 1-BGP.0.0.0+179
    AS path: I
    Communities: Layer2-info: encaps:VPLS, control flags:, mtu: 0
    Label-base: 800008, range: 8, status-vector: 0x9F

...

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
  *L2CKT Preference: 7
    Next hop: via so-1/1/2.0 weight 1, selected
    Label-switched-path my-lsp
    Label operation: Push 100000[0]
    Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
    State: <Active Int>
    Local AS: 99
    Age: 10:21
    Task: l2 circuit
    Announcement bits (1): 0-LDP
    AS path: I
    VC Label 100000, MTU 1500, VLAN ID 512

```

show route detail (with BGP Multipath)

```

user@host> show route detail

10.1.1.8/30 (2 entries, 1 announced)
  *BGP Preference: 170/-101
    Next hop type: Router, Next hop index: 262142
    Address: 0x901a010
    Next-hop reference count: 2
    Source: 10.1.1.2
    Next hop: 10.1.1.2 via lt-0/3/0.1, selected
    Next hop: 10.1.1.6 via lt-0/3/0.5
    State: <Active Ext>
    Local AS: 1 Peer AS: 2
    Age: 5:04:43
    Task: BGP_2.10.1.1.2+59955
    Announcement bits (1): 0-KRT
    AS path: 2 I
    Accepted Multipath
    Localpref: 100
    Router ID: 1.1.1.2
  BGP Preference: 170/-101
    Next hop type: Router, Next hop index: 678
    Address: 0x8f97520
    Next-hop reference count: 9

```

```

Source: 10.1.1.6
Next hop: 10.1.1.6 via lt-0/3/0.5, selected
State: <NotBest Ext>
Inactive reason: Not Best in its group - Active preferred
Local AS: 1 Peer AS: 2
Age: 5:04:43
Task: BGP_2.10.1.1.6+58198
AS path: 2 I
Accepted MultipathContrib
Localpref: 100
Router ID: 1.1.1.3

```

show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```

user@host> show route label 299872 detail
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
299872 (1 entry, 1 announced)
    *LDP      Preference: 9
              Next hop type: Flood
              Next-hop reference count: 3
              Address: 0x9097d90
              Next hop: via vt-0/1/0.1
              Next-hop index: 661
              Label operation: Pop
              Address: 0x9172130
              Next hop: via so-0/0/3.0
              Next-hop index: 654
              Label operation: Swap 299872
              State: **Active Int>
              Local AS: 1001
              Age: 8:20      Metric: 1
              Task: LDP
              Announcement bits (1): 0-KRT
              AS path: I
              FECs bound to route: P2MP root-addr 10.255.72.166, grp 232.1.1.1,
src 192.168.142.2

```

show route label detail (Multipoint LDP with Multicast-Only Fast Reroute)

```

user@host> show route label 301568 detail
mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
301568 (1 entry, 1 announced)
    *LDP      Preference: 9
              Next hop type: Flood
              Address: 0x2735208
              Next-hop reference count: 3
              Next hop type: Router, Next hop index: 1397
              Address: 0x2735d2c
              Next-hop reference count: 3
              Next hop: 1.3.8.2 via ge-1/2/22.0
              Label operation: Pop
              Load balance label: None;
              Next hop type: Router, Next hop index: 1395
              Address: 0x2736290
              Next-hop reference count: 3
              Next hop: 1.3.4.2 via ge-1/2/18.0
              Label operation: Pop
              Load balance label: None;
              State: <Active Int AckRequest MulticastRPF>
              Local AS: 10

```

```
Age: 54:05      Metric: 1
Validation State: unverified
Task: LDP
Announcement bits (1): 0-KRT
AS path: I
FECs bound to route: P2MP root-addr 1.1.1.1, grp: 232.1.1.1, src:
192.168.219.11
Primary Upstream : 1.1.1.3:0--1.1.1.2:0
  RPF Nexthops :
    ge-1/2/15.0, 1.2.94.1, Label: 301568, weight: 0x1
    ge-1/2/14.0, 1.2.3.1, Label: 301568, weight: 0x1
Backup Upstream : 1.1.1.3:0--1.1.1.6:0
  RPF Nexthops :
    ge-1/2/20.0, 1.2.96.1, Label: 301584, weight: 0xffffe
    ge-1/2/19.0, 1.3.6.1, Label: 301584, weight: 0xffffe
```

show route exact

List of Syntax	Syntax on page 3453 Syntax (EX Series Switches) on page 3453
Syntax	show route exact <i>destination-prefix</i> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	show route exact <i>destination-prefix</i> <brief detail extensive terse>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display only the routes that exactly match the specified address or range of addresses.
Options	brief detail extensive terse —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief. <i>destination-prefix</i> —Address or range of addresses. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show route exact on page 3453 show route exact detail on page 3453 show route exact extensive on page 3454 show route exact terse on page 3454
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route exact

```

user@host> show route exact 207.17.136.0/24

inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
207.17.136.0/24    *[Static/5] 2d 03:30:22
                  > to 192.168.71.254 via fxp0.0

```

show route exact detail

```

user@host> show route exact 207.17.136.0/24 detail

inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)

```

```
Restart Complete
207.17.136.0/24 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 29
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 2d 3:30:26
    Task: RT
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I
```

show route exact extensive

```
user@host> show route exact 207.17.136.0/24 extensive
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
207.17.136.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.0/24 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 29
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 1:25:18
    Task: RT
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I
```

show route exact terse

```
user@host> show route exact 207.17.136.0/24 terse

inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
* 207.17.136.0/24  S  5                >192.168.71.254
```

show route export

List of Syntax	Syntax on page 3455 Syntax (EX Series Switches) on page 3455	
Syntax	<pre>show route export <brief detail> <instance <instance-name> routing-table-name> <logical-system (all logical-system-name)></pre>	
Syntax (EX Series Switches)	<pre>show route export <brief detail> <instance <instance-name> routing-table-name></pre>	
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>	
Description	<p>Display policy-based route export information. Policy-based export simplifies the process of exchanging route information between routing instances.</p>	
Options	<p>none—(Same as brief.) Display standard information about policy-based export for all instances and routing tables on all systems.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>instance <instance-name>—(Optional) Display a particular routing instance for which policy-based export is currently enabled.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>routing-table-name—(Optional) Display information about policy-based export for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the show route export inet command).</p>	
Required Privilege Level	view	
List of Sample Output	show route export on page 3456 show route export detail on page 3456 show route export instance detail on page 3456	
Output Fields	<p>Table 303 lists the output fields for the show route export command. Output fields are listed in the approximate order in which they appear.</p>	

Table 303: show route export Output Fields

Field Name	Field Description	Level of Output
Table or table-name	Name of the routing tables that either import or export routes.	All levels
Routes	Number of routes exported from this table into other tables. If a particular route is exported to different tables, the counter will only increment by one.	brief none

Table 303: show route export Output Fields (*continued*)

Field Name	Field Description	Level of Output
Export	Whether the table is currently exporting routes to other tables: Y or N (Yes or No).	brief none
Import	Tables currently importing routes from the originator table. (Not displayed for tables that are not exporting any routes.)	detail
Flags	(instance keyword only) Flags for this feature on this instance: <ul style="list-style-type: none"> config auto-policy—The policy was deduced from the configured IGP export policies. cleanup—Configuration information for this instance is no longer valid. config—The instance was explicitly configured. 	detail
Options	(instance keyword only) Configured option displays the type of routing tables the feature handles: <ul style="list-style-type: none"> unicast—Indicates <i>instance.inet.0</i>. multicast—Indicates <i>instance.inet.2</i>. unicast multicast—Indicates <i>instance.inet.0</i> and <i>instance.inet.2</i>. 	detail
Import policy	(instance keyword only) Policy that route export uses to construct the import-export matrix. Not displayed if the instance type is vrf .	detail
Instance	(instance keyword only) Name of the routing instance.	detail
Type	(instance keyword only) Type of routing instance: forwarding , non-forwarding , or vrf .	detail

Sample Output

show route export

```

user@host> show route export
Table      Export      Routes
inet.0     N            0
black.inet.0 Y           3
red.inet.0 Y            4

```

show route export detail

```

user@host> show route export detail
inet.0                                Routes:      0
black.inet.0                          Routes:      3
  Import: [ inet.0 ]
red.inet.0                            Routes:      4
  Import: [ inet.0 ]

```

show route export instance detail

```

user@host> show route export instance detail
Instance: master                      Type: forwarding
Flags: <config auto-policy> Options: <unicast multicast>
Import policy: [ (ospf-master-from-red || isis-master-from-black) ]

```


Instance: black
Instance: red

Type: non-forwarding
Type: non-forwarding

show route extensive

List of Syntax	Syntax on page 3458 Syntax (EX Series Switches) on page 3458
Syntax	show route extensive <destination-prefix> <logical-system (all logical-system-name)>
Syntax (EX Series Switches)	show route extensive <destination-prefix>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display extensive information about the active entries in the routing tables.
Options	none —Display all active entries in the routing table. destination-prefix —(Optional) Display active entries for the specified address or range of addresses. logical-system (all logical-system-name) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show route extensive on page 3465 show route extensive (Access Route) on page 3471 show route extensive (BGP PIC Edge) on page 3472 show route extensive (FRR and LFA) on page 3472 show route extensive (Route Reflector) on page 3473 show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 3473 show route label detail (Multipoint LDP with Multicast-Only Fast Reroute) on page 3474
Output Fields	Table 304 describes the output fields for the show route extensive command. Output fields are listed in the approximate order in which they appear.

Table 304: show route extensive Output Fields

Field Name	Field Description
<i>routing-table-name</i>	Name of the routing table (for example, inet.0).
<i>number destinations</i>	Number of destinations for which there are routes in the routing table.

Table 304: show route extensive Output Fields (*continued*)

Field Name	Field Description
<i>number routes</i>	<p>Number of routes in the routing table and total number of routes in the following states:</p> <ul style="list-style-type: none"> • active (routes that are active). • holddown (routes that are in the pending state before being declared inactive). • hidden (routes that are not used because of a routing policy).
<i>route-destination</i> (entry, announced)	<p>Route destination (for example: 10.0.0.1/24). The entry value is the number of route for this destination, and the announced value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:</p> <ul style="list-style-type: none"> • MPLS-label (for example, 80001). • interface-name (for example, ge-1/0/2). • neighbor-address:control-word-status:encapsulation type:vc-id:source (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96). <ul style="list-style-type: none"> • neighbor-address—Address of the neighbor. • control-word-status—Whether the use of the control word has been negotiated for this virtual circuit: NoCtrlWord or CtrlWord. • encapsulation type—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport. • vc-id—Virtual circuit identifier. • source—Source of the advertisement: Local or Remote.
TSI	Protocol header information.
label stacking	<p>(Next-to-the-last-hop routing device for MPLS only) Depth of the Multiprotocol Label Switching (MPLS) label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels.</p> <ul style="list-style-type: none"> • S=0 route indicates that a packet with an incoming label stack depth of two or more exits this router with one fewer label (the label-popping operation is performed). • If there is no S= information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).
[protocol, preference]	<p>Protocol from which the route was learned and the preference value for the route.</p> <ul style="list-style-type: none"> • +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • —A hyphen indicates the last active route. • *—An asterisk indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route. <p>In every routing metric except for the BGP LocalPref attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the LocalPref value in the Preference2 field. For example, if the LocalPref value for Route 1 is 100, the Preference2 value is -101. If the LocalPref value for Route 2 is 155, the Preference2 value is -156. Route 2 is preferred because it has a higher LocalPref value and a lower Preference2 value.</p>

Table 304: show route extensive Output Fields (*continued*)

Field Name	Field Description
Level	(IS-IS only). In IS-IS, a single autonomous system (AS) can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.
Route Distinguisher	IP subnet augmented with a 64-bit prefix.
PMSI	Provider multicast service interface (MVPN routing table).
Next-hop type	Type of next hop. For a description of possible values for this field, see the Output Field table in the show route detail command.
Next-hop reference count	Number of references made to the next hop.
Flood nexthop branches exceed maximum message	Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.
Source	IP address of the route source.
Next hop	Network layer address of the directly reachable neighboring system.
via	<p>Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word Selected. This field can also contain the following information:</p> <ul style="list-style-type: none"> • Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when Multiprotocol Label Switching (MPLS) label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable Border Gateway Protocol (BGP) multipath load balancing.
Label-switched-path lsp-path-name	Name of the label-switched path (LSP) used to reach the next hop.
Label operation	MPLS label and operation occurring at this routing device. The operation can be pop (where a label is removed from the top of the stack), push (where another label is added to the label stack), or swap (where a label is replaced by another label).
Offset	Whether the metric has been increased or decreased by an offset value.
Interface	(Local only) Local interface name.
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to recursively derive a forwarding next hop.

Table 304: show route extensive Output Fields (*continued*)

Field Name	Field Description
<i>label-operation</i>	MPLS label and operation occurring at this routing device. The operation can be pop (where a label is removed from the top of the stack), push (where another label is added to the label stack), or swap (where a label is replaced by another label).
Indirect next hops	<p>When present, a list of nodes that are used to resolve the path to the next-hop destination, in the order that they are resolved.</p> <p>When BGP PIC Edge is enabled, the output lines that contain Indirect next hop: weight follow next hops that the software can use to repair paths where a link failure occurs. The next-hop weight has one of the following values:</p> <ul style="list-style-type: none"> • 0x1 indicates active next hops. • 0x4000 indicates passive next hops.
State	State of the route (a route can be in more than one state). See the Output Field table in the show route detail command.
Session ID	The BFD session ID number that represents the protection using MPLS fast reroute (FRR) and loop-free alternate (LFA).
Weight	<p>Weight for the backup path. If the weight of an indirect next hop is larger than zero, the weight value is shown.</p> <p>For sample output, see show route table.</p>

Table 304: show route extensive Output Fields (*continued*)

Field Name	Field Description
Inactive reason	<p>If the route is inactive, the reason for its current state is indicated. Typical reasons include:</p> <ul style="list-style-type: none"> • Active preferred—Currently active route was selected over this route. • Always compare MED—Path with a lower multiple exit discriminator (MED) is available. • AS path—Shorter AS path is available. • Cisco Non-deterministic MED selection—Cisco nondeterministic MED is enabled and a path with a lower MED is available. • Cluster list length—Path with a shorter cluster list length is available. • Forwarding use only—Path is only available for forwarding purposes. • IGP metric—Path through the next hop with a lower IGP metric is available. • IGP metric type—Path with a lower OSPF link-state advertisement type is available. • Interior > Exterior > Exterior via Interior—Direct, static, IGP, or EBGP path is available. • Local preference—Path with a higher local preference value is available. • Next hop address—Path with a lower metric next hop is available. • No difference—Path from a neighbor with a lower IP address is available. • Not Best in its group—Occurs when multiple peers of the same external AS advertise the same prefix and are grouped together in the selection process. When this reason is displayed, an additional reason is provided (typically one of the other reasons listed). • Number of gateways—Path with a higher number of next hops is available. • Origin—Path with a lower origin code is available. • OSPF version—Path does not support the indicated OSPF version. • RIB preference—Route from a higher-numbered routing table is available. • Route distinguisher—64-bit prefix added to IP subnets to make them unique. • Route metric or MED comparison—Route with a lower metric or MED is available. • Route preference—Route with a lower preference value is available. • Router ID—Path through a neighbor with a lower ID is available. • Unusable path—Path is not usable because of one of the following conditions: the route is damped, the route is rejected by an import policy, or the route is unresolved. • Update source—Last tiebreaker is the lowest IP address value.
Local AS	Autonomous system (AS) number of the local routing device.
Age	How long the route has been known.
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.
Metric	Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
MED-plus-IGP	Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.
TTL-Action	<p>For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signalled and LDP-signalled LSPs or for specific VRF routing instances.</p> <p>For sample output, see show route table.</p>

Table 304: show route extensive Output Fields (*continued*)

Field Name	Field Description
Task	Name of the protocol that has added the route.
Announcement bits	List of protocols that announce this route. n-Resolve inet indicates that the route is used for route resolution for next hops found in the routing table. n is an index used by Juniper Networks customer support only.
AS path	<p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> • I—IGP. • E—EGP. • Recorded—The AS path is recorded by the sample process (sampled). • ?—Incomplete; typically, the AS path was aggregated. <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> • []—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device, or if AS path prepending is configured. • { }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • ()—Parentheses enclose a confederation. • ([])—Parentheses and brackets enclose a confederation set. <p>NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p>
validation-state	<p>(BGP-learned routes) Validation status of the route:</p> <ul style="list-style-type: none"> • Invalid—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • Unknown—Indicates that the prefix is not among the prefixes or prefix ranges in the database. • Unverified—Indicates that origin validation is not enabled for the BGP peers. • Valid—Indicates that the prefix and autonomous system pair are found in the database.
FECs bound to route	Point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.
AS path: I <Originator>	(For route reflected output only) Originator ID attribute set by the route reflector.

Table 304: show route extensive Output Fields (*continued*)

Field Name	Field Description
route status	<p>Indicates the status of a BGP route:</p> <ul style="list-style-type: none"> • Accepted—The specified BGP route is imported by the default BGP policy. • Import—The route is imported into a Layer 3 VPN routing instance. • Import-Protect—A remote instance egress that is protected. • Multipath—A BGP multipath active route. • MultipathContrib—The route is not active but contributes to the BGP multipath. • Protect—An egress route that is protected. • Stale—A route that is marked stale due to graceful restart.
Primary Upstream	When multipoint LDP with multicast-only fast reroute (MoFRR) is configured, the primary upstream path. MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path.
RPF Nexthops	When multipoint LDP with MoFRR is configured, the reverse-path forwarding (RPF) next-hop information. Data packets are received from both the primary path and the secondary paths. The redundant packets are discarded at topology merge points due to the RPF checks.
Label	Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.
weight	Value used to distinguish MoFRR primary and backup routes. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.
VC Label	MPLS label assigned to the Layer 2 circuit virtual connection.
MTU	Maximum transmission unit (MTU) of the Layer 2 circuit.
VLAN ID	VLAN identifier of the Layer 2 circuit.
Cluster list	(For route reflected output only) Cluster ID sent by the route reflector.
Originator ID	(For route reflected output only) Address of router that originally sent the route to the route reflector.
Prefixes bound to route	Forwarding Equivalent Class (FEC) bound to this route. Applicable only to routes installed by LDP.
Communities	Community path attribute for the route. See the Output Field table in the show route detail command for all possible values for this field.
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).
control flags	Control flags: none or Site Down.
mtu	Maximum transmission unit (MTU) information.
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.

Table 304: show route extensive Output Fields (*continued*)

Field Name	Field Description
status vector	Layer 2 VPN and VPLS network layer reachability information (NLRI).
Localpref	Local preference value included in the route.
Router ID	BGP router ID as advertised by the neighbor in the open message.
Primary Routing Table	In a routing table group, the name of the primary routing table in which the route resides.
Secondary Tables	In a routing table group, the name of one or more secondary tables in which the route resides.
Originating RIB	Name of the routing table whose active route was used to determine the forwarding next-hop entry in the resolution database. For example, in the case of inet.0 resolving through inet.0 and inet.3, this field indicates which routing table, inet.0 or inet.3, provided the best path for a particular prefix.
Node path count	Number of nodes in the path.
Forwarding nexthops	Number of forwarding next hops. The forwarding next hop is the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.

Sample Output

show route extensive

```

user@host> show route extensive
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 29
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 1:34:06
    Task: RT
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I

10.31.1.0/30 (2 entries, 1 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 2
    Next hop: via so-0/3/0.0, selected
    State: <Active Int>
    Local AS: 69
    Age: 1:32:40
    Task: IF
    Announcement bits (1): 3-Resolve tree 2
    AS path: I
  OSPF Preference: 10
    Next-hop reference count: 1
    Next hop: via so-0/3/0.0, selected

```

```
State: <Int>
Inactive reason: Route Preference
Local AS: 69
Age: 1:32:40 Metric: 1
Area: 0.0.0.0
Task: OSPF
AS path: I

10.31.1.1/32 (1 entry, 1 announced)
*Local Preference: 0
Next hop type: Local
Next-hop reference count: 7
Interface: so-0/3/0.0
State: <Active NoReadvrt Int>
Local AS: 69
Age: 1:32:43
Task: IF
Announcement bits (1): 3-Resolve tree 2
AS path: I

...

10.31.2.0/30 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.31.2.0/30 -> {10.31.1.6}
*OSPF Preference: 10
Next-hop reference count: 9
Next hop: via so-0/3/0.0
Next hop: 10.31.1.6 via ge-3/1/0.0, selected
State: <Active Int>
Local AS: 69
Age: 1:32:19 Metric: 2
Area: 0.0.0.0
Task: OSPF
Announcement bits (2): 0-KRT 3-Resolve tree 2
AS path: I

...

224.0.0.2/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 224.0.0.2/32 -> {}
*PIM Preference: 0
Next-hop reference count: 18
State: <Active NoReadvrt Int>
Local AS: 69
Age: 1:34:08
Task: PIM Recv
Announcement bits (2): 0-KRT 3-Resolve tree 2
AS path: I

...

224.0.0.22/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 224.0.0.22/32 -> {}
*IGMP Preference: 0
Next-hop reference count: 18
State: <Active NoReadvrt Int>
Local AS: 69
Age: 1:34:06
```

```

Task: IGMP
Announcement bits (2): 0-KRT 3-Resolve tree 2
AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.255.70.103/32 (1 entry, 1 announced)
State: <FlashAll>
*RSVP Preference: 7
Next-hop reference count: 6
Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
Label-switched-path green-r1-r3
Label operation: Push 100096
State: <Active Int>
Local AS: 69
Age: 1:28:12 Metric: 2
Task: RSVP
Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
AS path: I

10.255.71.238/32 (1 entry, 1 announced)
State: <FlashAll>
*RSVP Preference: 7
Next-hop reference count: 6
Next hop: via so-0/3/0.0 weight 0x1, selected
Label-switched-path green-r1-r2
State: <Active Int>
Local AS: 69
Age: 1:28:12 Metric: 1
Task: RSVP
Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
AS path: I

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

...

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

47.0005.80ff.f800.0000.0108.0001.0102.5507.1052/152 (1 entry, 0 announced)
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 1
Next hop: via lo0.0, selected
State: <Active Int>
Local AS: 69
Age: 1:34:07
Task: IF
AS path: I

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

0 (1 entry, 1 announced)
TSI:
KRT in-kernel 0 /36 -> {}
*MPLS Preference: 0
Next hop type: Receive
Next-hop reference count: 6
State: <Active Int>
Local AS: 69
Age: 1:34:08 Metric: 1

```

```

Task: MPLS
Announcement bits (1): 0-KRT
AS path: I

...

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
299776 (1 entry, 1 announced)
TSI:
KRT in-kernel 299776 /52 -> {Flood}
    *RSVP    Preference: 7
             Next hop type: Flood
             Next-hop reference count: 130
             Flood nexthop branches exceed maximum
             Address: 0x8ea65d0

...

800010 (1 entry, 1 announced)
TSI:
KRT in-kernel 800010 /36 -> {vt-3/2/0.32769}
    *VPLS    Preference: 7
             Next-hop reference count: 2
             Next hop: via vt-3/2/0.32769, selected
             Label operation: Pop
             State: <Active Int>
             Age: 1:31:53
             Task: Common L2 VC
             Announcement bits (1): 0-KRT
             AS path: I

vt-3/2/0.32769 (1 entry, 1 announced)
TSI:
KRT in-kernel vt-3/2/0.32769.0      /16 -> {indirect(1048574)}
    *VPLS    Preference: 7
             Next-hop reference count: 2
             Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
             Label-switched-path green-r1-r3
             Label operation: Push 800012, Push 100096(top)
             Protocol next hop: 10.255.70.103
             Push 800012
             Indirect next hop: 87272e4 1048574
             State: <Active Int>
             Age: 1:31:53    Metric2: 2
             Task: Common L2 VC
             Announcement bits (2): 0-KRT 1-Common L2 VC
             AS path: I
             Communities: target:11111:1 Layer2-info: encaps:VPLS,
             control flags:, mtu: 0
             Indirect next hops: 1
                 Protocol next hop: 10.255.70.103 Metric: 2
                 Push 800012
                 Indirect next hop: 87272e4 1048574
                 Indirect path forwarding next hops: 1
                     Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1
                     10.255.70.103/32 Originating RIB: inet.3
                     Metric: 2                                Node path count: 1
                     Forwarding nexthops: 1
                         Nexthop: 10.31.1.6 via ge-3/1/0.0

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

```

```

abcd::10:255:71:52/128 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
    State: <Active Int>
    Local AS: 69
    Age: 1:34:07
    Task: IF
    AS path: I

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:34:07
    Task: IF
    AS path: I

ff02::2/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::2/128 -> {}
  *PIM Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:34:08
    Task: PIM Recv6
    Announcement bits (1): 0-KRT
    AS path: I

ff02::d/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::d/128 -> {}
  *PIM Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:34:08
    Task: PIM Recv6
    Announcement bits (1): 0-KRT
    AS path: I

ff02::16/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::16/128 -> {}
  *MLD Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:34:06
    Task: MLD
    Announcement bits (1): 0-KRT
    AS path: I

private.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

```
fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.16385, selected
    State: <Active NoReadvrt Int>
    Age: 1:34:07
    Task: IF
    AS path: I

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

10.255.70.103:1:3:1/96 (1 entry, 1 announced)
  *BGP Preference: 170/-101
    Route Distinguisher: 10.255.70.103:1
    Next-hop reference count: 7
    Source: 10.255.70.103
    Protocol next hop: 10.255.70.103
    Indirect next hop: 2 no-forward
    State: <Secondary Active Int Ext>
    Local AS: 69 Peer AS: 69
    Age: 1:28:12 Metric2: 1
    Task: BGP_69.10.255.70.103+179
    Announcement bits (1): 0-green-l2vpn
    AS path: I
    Communities: target:11111:1 Layer2-info: encaps:VPLS,
    control flags:, mtu: 0
    Label-base: 800008, range: 8
    Localpref: 100
    Router ID: 10.255.70.103
    Primary Routing Table bgp.l2vpn.0

10.255.71.52:1:1:1/96 (1 entry, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 8699540
  *L2VPN Preference: 170/-1
    Next-hop reference count: 5
    Protocol next hop: 10.255.71.52
    Indirect next hop: 0 -
    State: <Active Int Ext>
    Age: 1:34:03 Metric2: 1
    Task: green-l2vpn
    Announcement bits (1): 1-BGP.0.0.0.0+179
    AS path: I
    Communities: Layer2-info: encaps:VPLS, control flags:Site-Down,
    mtu: 0
    Label-base: 800016, range: 8, status-vector: 0x9F

10.255.71.52:1:5:1/96 (1 entry, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 8699528
  *L2VPN Preference: 170/-101
    Next-hop reference count: 5
    Protocol next hop: 10.255.71.52
    Indirect next hop: 0 -
    State: <Active Int Ext>
    Age: 1:34:03 Metric2: 1
    Task: green-l2vpn
    Announcement bits (1): 1-BGP.0.0.0.0+179
    AS path: I
    Communities: Layer2-info: encaps:VPLS, control flags:, mtu: 0
```

```

Label-base: 800008, range: 8, status-vector: 0x9F

...

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
TSI:
10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
  *L2CKT Preference: 7
    Next hop: via so-1/1/2.0 weight 1, selected
    Label-switched-path my-lsp
    Label operation: Push 100000[0]
    Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
    State: <Active Int>
    Local AS: 99
    Age: 10:21
    Task: l2 circuit
    Announcement bits (1): 0-LDP
    AS path: I
    VC Label 100000, MTU 1500, VLAN ID 512

55.0.0.0/24 (1 entry, 1 announced)
TSI:
KRT queued (pending) add
  55.0.0.0/24 -> {Push 300112}
    *BGP Preference: 170/-101
      Next hop type: Router
      Address: 0x925c208
      Next-hop reference count: 2
      Source: 10.0.0.9
      Next hop: 10.0.0.9 via lt-1/2/0.15, selected
      Label operation: Push 300112
      Label TTL action: prop-ttl
      State: <Active Ext>
      Local AS: 7019 Peer AS: 13979
      Age: 1w0d 23:06:56
      AIGP: 25
      Task: BGP_13979.10.0.0.9+56732
      Announcement bits (1): 0-KRT
      AS path: 13979 7018 I
      Accepted
      Route Label: 300112
      Localpref: 100
      Router ID: 10.9.9.1

```

show route extensive (Access Route)

```

user@host> show route 13.160.0.102 extensive
inet.0: 39256 destinations, 39258 routes (39255 active, 0 holddown, 1 hidden)
13.160.0.102/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 13.160.0.102/32 -> {13.160.0.2}
OSPF area : 0.0.0.0, LSA ID : 13.160.0.102, LSA type : Extern
  *Access Preference: 13
    Next-hop reference count: 78472
    Next hop: 13.160.0.2 via fe-0/0/0.0, selected
    State: <Active Int>
  Age: 12
    Task: RPD Unix Domain Server./var/run/rpd_serv.local

```

```
Announcement bits (2): 0-KRT 1-OSPFv2
AS path: I
```

show route extensive (BGP PIC Edge)

```
user@host> show route 1.1.1.6 extensive
ed.inet.0: 6 destinations, 9 routes (6 active, 0 holddown, 0 hidden)
  1.1.1.6/32 (3 entries, 2 announced)
    State: <CalcForwarding>
    TSI:
    KRT in-kernel 1.1.1.6/32 -> {indirect(1048574), indirect(1048577)}
    Page 0 idx 0 Type 1 val 9219e30
      Nexthop: Self
      AS path: [2] 3 I
      Communities: target:2:1
    Path 1.1.1.6 from 1.1.1.4 Vector len 4. Val: 0
  ..
    #Multipath Preference: 255
      Next hop type: Indirect
      Address: 0x93f4010
      Next-hop reference count: 2
  ..
    Protocol next hop: 1.1.1.4
    Push 299824
    Indirect next hop: 944c000 1048574 INH Session ID: 0x3
    Indirect next hop: weight 0x1
    Protocol next hop: 1.1.1.5
    Push 299824
    Indirect next hop: 944c1d8 1048577 INH Session ID: 0x4
    Indirect next hop: weight 0x4000
    State: <ForwardingOnly Int Ext>
    Inactive reason: Forwarding use only
    Age: 25      Metric2: 15
    Validation State: unverified
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: 3 I
    Communities: target:2:1
```

show route extensive (FRR and LFA)

```
user@host> show route 20.31.2.0 extensive
inet.0: 46 destinations, 49 routes (45 active, 0 holddown, 1 hidden)
  20.31.2.0/24 (2 entries, 1 announced)
    State: FlashAll
    TSI:
    KRT in-kernel 20.31.2.0/24 -> {Push 299776, Push 299792}
      *RSVP Preference: 7/1
        Next hop type: Router, Next hop index: 1048574
        Address: 0xbbbc010
        Next-hop reference count: 5
        Next hop: 10.31.1.2 via ge-2/1/8.0 weight 0x1, selected
        Label-switched-path europa-d-to-europa-e
        Label operation: Push 299776
        Label TTL action: prop-ttl
        Session Id: 0x201
        Next hop: 10.31.2.2 via ge-2/1/4.0 weight 0x4001
        Label-switched-path europa-d-to-europa-e
        Label operation: Push 299792
        Label TTL action: prop-ttl
        Session Id: 0x202
```



```

State: Active Int
Local AS: 100
Age: 5:31 Metric: 2
Task: RSVP
Announcement bits (1): 0-KRT
AS path: I
OSPF Preference: 10
Next hop type: Router, Next hop index: 615
Address: 0xb9d78c4
Next-hop reference count: 7
Next hop: 10.31.1.2 via ge-2/1/8.0, selected
Session Id: 0x201
State: Int
Inactive reason: Route Preference
Local AS: 100
Age: 5:35 Metric: 3
Area: 0.0.0.0
Task: OSPF
AS path: I

```

show route extensive (Route Reflector)

```

user@host> show route extensive
1.0.0.0/8 (1 entry, 1 announced)

TSI:
KRT in-kernel 1.0.0.0/8 -> {indirect(40)}
*BGP Preference: 170/-101
Source: 192.168.4.214
Protocol next hop: 207.17.136.192 Indirect next hop: 84ac908 40
State: <Active Int Ext>
Local AS: 10458 Peer AS: 10458
Age: 3:09 Metric: 0 Metric2: 0
Task: BGP_10458.192.168.4.214+1033
Announcement bits (2): 0-KRT 4-Resolve inet.0
AS path: 3944 7777 I <Originator>
Cluster list: 1.1.1.1
Originator ID: 10.255.245.88
Communities: 7777:7777
Localpref: 100
Router ID: 4.4.4.4
Indirect next hops: 1
    Protocol next hop: 207.17.136.192 Metric: 0
    Indirect next hop: 84ac908 40
    Indirect path forwarding next hops: 0
    Next hop type: Discard

```

show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```

user@host> show route label 299872 detail
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
299872 (1 entry, 1 announced)
*LDP Preference: 9
Next hop type: Flood
Next-hop reference count: 3
Address: 0x9097d90
Next hop: via vt-0/1/0.1
Next-hop index: 661
Label operation: Pop
Address: 0x9172130
Next hop: via so-0/0/3.0

```

```

Next-hop index: 654
Label operation: Swap 299872
State: **Active Int>
Local AS: 1001
Age: 8:20      Metric: 1
Task: LDP
Announcement bits (1): 0-KRT
AS path: I
FECs bound to route: P2MP root-addr 10.255.72.166, grp 232.1.1.1,
src 192.168.142.2

```

show route label detail (Multipoint LDP with Multicast-Only Fast Reroute)

```
user@host> show route label 301568 detail
```

```

mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
301568 (1 entry, 1 announced)
  *LDP      Preference: 9
    Next hop type: Flood
    Address: 0x2735208
    Next-hop reference count: 3
    Next hop type: Router, Next hop index: 1397
    Address: 0x2735d2c
    Next-hop reference count: 3
    Next hop: 1.3.8.2 via ge-1/2/22.0
    Label operation: Pop
    Load balance label: None;
    Next hop type: Router, Next hop index: 1395
    Address: 0x2736290
    Next-hop reference count: 3
    Next hop: 1.3.4.2 via ge-1/2/18.0
    Label operation: Pop
    Load balance label: None;
    State: <Active Int AckRequest MulticastRPF>
    Local AS: 10
    Age: 54:05      Metric: 1
    Validation State: unverified
    Task: LDP
    Announcement bits (1): 0-KRT
    AS path: I
    FECs bound to route: P2MP root-addr 1.1.1.1, grp: 232.1.1.1, src:
192.168.219.11
      Primary Upstream : 1.1.1.3:0--1.1.1.2:0
        RPF Nexthops :
          ge-1/2/15.0, 1.2.94.1, Label: 301568, weight: 0x1
          ge-1/2/14.0, 1.2.3.1, Label: 301568, weight: 0x1
      Backup Upstream : 1.1.1.3:0--1.1.1.6:0
        RPF Nexthops :
          ge-1/2/20.0, 1.2.96.1, Label: 301584, weight: 0xffffe
          ge-1/2/19.0, 1.3.6.1, Label: 301584, weight: 0xffffe

```

show route flow validation

List of Syntax [Syntax on page 3475](#)
[Syntax \(EX Series Switches\) on page 3475](#)

Syntax show route flow validation
 <brief | detail>
 <ip-prefix>
 <table *table-name*>
 <logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches) show route flow validation
 <brief | detail>
 <ip-prefix>
 <table *table-name*>

Release Information Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.0 for EX Series switches.

Description Display flow route information.

Options **none**—Display flow route information.

brief | detail—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.

ip-prefix—(Optional) IP address for the flow route.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

table *table-name*—(Optional) Display flow route information for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the **show route flow validation inet** command).

Required Privilege Level view

List of Sample Output [show route flow validation on page 3476](#)

Output Fields [Table 305](#) lists the output fields for the **show route flow validation** command. Output fields are listed in the approximate order in which they appear.

Table 305: show route flow validation Output Fields

Field Name	Field Description	Level of Output
<i>routing-table-name</i>	Name of the routing table (for example, inet.0).	All levels
<i>prefix</i>	Route address.	All levels
Active unicast route	Active route in the routing table.	All levels

Table 305: show route flow validation Output Fields (*continued*)

Field Name	Field Description	Level of Output
Dependent flow destinations	Number of flows for which there are routes in the routing table.	All levels
Origin	Source of the route flow.	All levels
Neighbor AS	Autonomous system identifier of the neighbor.	All levels
Flow destination	Number of entries and number of destinations that match the route flow.	All levels
Unicast best match	Destination that is the best match for the route flow.	All levels
Flags	Information about the route flow.	All levels

Sample Output

show route flow validation

```
user@host> show route flow validation
inet.0:
10.0.5.0/24Active unicast route
Dependent flow destinations: 1
Origin: 192.168.224.218, Neighbor AS: 65001
Flow destination (3 entries, 1 match origin)
Unicast best match: 10.0.5.0/24
Flags: SubtreeApex Consistent
```

show route forwarding-table

List of Syntax	Syntax on page 3477 Syntax (MX Series Routers) on page 3477 Syntax (TX Matrix and TX Matrix Plus Routers) on page 3477
Syntax	<pre>show route forwarding-table <detail extensive summary> <all> <ccc interface-name> <destination destination-prefix> <family family matching matching> <interface-name interface-name> <label name> <matching matching> <multicast> <table (default logical-system-name/routing-instance-name routing-instance-name)> <vlan (all vlan-name)> <vpn vpn></pre>
Syntax (MX Series Routers)	<pre>show route forwarding-table <detail extensive summary> <all> <bridge-domain (all domain-name)> <ccc interface-name> <destination destination-prefix> <family family matching matching> <interface-name interface-name> <label name> <learning-vlan-id learning-vlan-id> <matching matching> <multicast> <table (default logical-system-name/routing-instance-name routing-instance-name)> <vlan (all vlan-name)> <vpn vpn></pre>
Syntax (TX Matrix and TX Matrix Plus Routers)	<pre>show route forwarding-table <detail extensive summary> <all> <ccc interface-name> <destination destination-prefix> <family family matching matching> <interface-name interface-name> <matching matching> <label name> <lcc number> <multicast> <table routing-instance-name> <vpn vpn></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Option bridge-domain introduced in Junos OS Release 7.5</p> <p>Option learning-vlan-id introduced in Junos OS Release 8.4</p>

Options **all** and **vlan** introduced in Junos OS Release 9.6.

Command introduced in Junos OS Release 11.3 for the QFX Series.

Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Display the Routing Engine's forwarding table, including the network-layer prefixes and their next hops. This command is used to help verify that the routing protocol process has relayed the correction information to the forwarding table. The Routing Engine constructs and maintains one or more routing tables. From the routing tables, the Routing Engine derives a table of active routes, called the forwarding table.



NOTE: The Routing Engine copies the forwarding table to the Packet Forwarding Engine, the part of the router that is responsible for forwarding packets. To display the entries in the Packet Forwarding Engine's forwarding table, use the **show pfe route** command.

Options **none**—Display the routes in the forwarding tables. By default, the **show route forwarding-table** command does not display information about private, or internal, forwarding tables.

detail | extensive | summary—(Optional) Display the specified level of output.

all—(Optional) Display routing table entries for all forwarding tables, including private, or internal, tables.

bridge-domain (all | bridge-domain-name)—(MX Series routers only) (Optional) Display route entries for all bridge domains or the specified bridge domain.

ccc interface-name—(Optional) Display route entries for the specified circuit cross-connect interface.

destination destination-prefix—(Optional) Destination prefix.

family family—(Optional) Display routing table entries for the specified family: **fibre-channel**, **fmembers**, **inet**, **inet6**, **iso**, **mpls**, **tnp**, **unix**, **vpls**, or **vlan-classification**.

interface-name interface-name—(Optional) Display routing table entries for the specified interface.

label name—(Optional) Display route entries for the specified label.

lcc number—(TX Matrix and TX matrix Plus routers only) (Optional) On a routing matrix composed of a TX Matrix router and T640 routers, display information for the specified T640 router (or line-card chassis) connected to the TX Matrix router. On a routing matrix composed of the TX Matrix Plus router and T1600 or T4000 routers, display information for the specified router (line-card chassis) connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

learning-vlan-id *learning-vlan-id*—(MX Series routers only) (Optional) Display learned information for all VLANs or for the specified VLAN.

matching *matching*—(Optional) Display routing table entries matching the specified prefix or prefix length.

multicast—(Optional) Display routing table entries for multicast routes.

table (*default* | *logical-system-name/routing-instance-name* | *routing-instance-name*)—(Optional) Display route entries for all the routing tables in the main routing instance or for the specified routing instance. If your device supports logical systems, you can also display route entries for the specified logical system and routing instance. To view the routing instances on your device, use the [show route instance](#) command.

vlan (*all* | *vlan-name*)—(Optional) Display information for all VLANs or for the specified VLAN.

vpn *vpn*—(Optional) Display routing table entries for a specified VPN.

Required Privilege Level

view

List of Sample Output

[show route forwarding-table on page 3482](#)
[show route forwarding-table detail on page 3483](#)
[show route forwarding-table destination extensive \(Weights and Balances\) on page 3483](#)
[show route forwarding-table extensive on page 3484](#)
[show route forwarding-table extensive \(RPF\) on page 3485](#)
[show route forwarding-table family mpls on page 3486](#)
[show route forwarding-table family vpls on page 3486](#)
[show route forwarding-table vpls \(Broadcast, unknown unicast, and multicast \(BUM\) hashing is enabled\) on page 3486](#)
[show route forwarding-table vpls \(Broadcast, unknown unicast, and multicast \(BUM\) hashing is enabled with MAC Statistics\) on page 3487](#)
[show route forwarding-table family vpls extensive on page 3487](#)
[show route forwarding-table table default on page 3488](#)
[show route forwarding-table table logical-system-name/routing-instance-name on page 3489](#)

[show route forwarding-table vpn on page 3490](#)

Output Fields Table 306 lists the output fields for the **show route forwarding-table** command. Output fields are listed in the approximate order in which they appear. Field names might be abbreviated (as shown in parentheses) when no level of output is specified, or when the **detail** keyword is used instead of the **extensive** keyword.

Table 306: show route forwarding-table Output Fields

Field Name	Field Description	Level of Output
Logical system	Name of the logical system. This field is displayed if you specify the table logical-system-name/routing-instance-name option on a device that is configured for and supports logical systems.	All levels
Routing table	Name of the routing table (for example, inet, inet6, mpls).	All levels
Address family	Address family (for example, IP, IPv6, ISO, MPLS, and VPLS).	All levels
Destination	Destination of the route.	detail extensive
Route Type (Type)	How the route was placed into the forwarding table. When the detail keyword is used, the route type might be abbreviated (as shown in parentheses): <ul style="list-style-type: none"> cloned (clon)—(TCP or multicast only) Cloned route. destination (dest)—Remote addresses directly reachable through an interface. destination down (iddn)—Destination route for which the interface is unreachable. interface cloned (ifcl)—Cloned route for which the interface is unreachable. route down (ifdn)—Interface route for which the interface is unreachable. ignore (ignr)—Ignore this route. interface (intf)—Installed as a result of configuring an interface. permanent (perm)—Routes installed by the kernel when the routing table is initialized. user—Routes installed by the routing protocol process or as a result of the configuration. 	All levels
Route Reference (RtRef)	Number of routes to reference.	detail extensive
Flags	Route type flags: <ul style="list-style-type: none"> none—No flags are enabled. accounting—Route has accounting enabled. cached—Cache route. incoming-iface interface-number—Check against incoming interface. prefix load balance—Load balancing is enabled for this prefix. rt nh decoupled—Route has been decoupled from the next hop to the destination. sent to PFE—Route has been sent to the Packet Forwarding Engine. static—Static route. 	extensive
Next hop	IP address of the next hop to the destination.	detail extensive

Table 306: show route forwarding-table Output Fields (*continued*)

Field Name	Field Description	Level of Output
Next hop Type (Type)	<p>Next-hop type. When the detail keyword is used, the next-hop type might be abbreviated (as indicated in parentheses):</p> <ul style="list-style-type: none"> • broadcast (bcst)—Broadcast. • deny—Deny. • discard (dscd)—Discard. • hold—Next hop is waiting to be resolved into a unicast or multicast type. • indexed (idxd)—Indexed next hop. • indirect (indr)—Indirect next hop. • local (locl)—Local address on an interface. • routed multicast (mcrst)—Regular multicast next hop. • multicast (mcst)—Wire multicast next hop (limited to the LAN). • multicast discard (mdsc)—Multicast discard. • multicast group (mgrp)—Multicast group member. • receive (rcv)—Receive. • reject (rjct)—Discard. An ICMP unreachable message was sent. • resolve (rslv)—Resolving the next hop. • unicast (ucst)—Unicast. • unilist (ulst)—List of unicast next hops. A packet sent to this next hop goes to any next hop in the list. 	detail extensive
Index	Software index of the next hop that is used to route the traffic for a given prefix.	detail extensive none
Route interface-index	Logical interface index from which the route is learned. For example, for interface routes, this is the logical interface index of the route itself. For static routes, this field is zero. For routes learned through routing protocols, this is the logical interface index from which the route is learned.	extensive
Reference (NhRef)	Number of routes that refer to this next hop.	detail extensive none
Next-hop interface (Netif)	Interface used to reach the next hop.	detail extensive none
Weight	Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible (see the Balance field description).	extensive
Balance	Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a router is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.	extensive
RPF interface	List of interfaces from which the prefix can be accepted. Reverse path forwarding (RPF) information is displayed only when rpf-check is configured on the interface.	extensive

Sample Output

show route forwarding-table

```

user@host> show route forwarding-table
Routing table: default.inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                               rjct  46   4
0.0.0.0/32       perm  0                               dscd  44   1
1.1.1.0/24       ifdn  0                               rslv  608  1 ge-2/0/1.0
1.1.1.0/32       iddn  0 1.1.1.0             recv  606  1 ge-2/0/1.0
1.1.1.1/32       user  0                               rjct  46   4
1.1.1.1/32       intf  0 1.1.1.1             locl  607  2
1.1.1.1/32       iddn  0 1.1.1.1             locl  607  2
1.1.1.255/32     iddn  0 ff:ff:ff:ff:ff:ff   bcst  605  1 ge-2/0/1.0
10.0.0.0/24       intf  0                               rslv  616  1 ge-2/0/0.0
10.0.0.0/32       dest  0 10.0.0.0            recv  614  1 ge-2/0/0.0
10.0.0.1/32       intf  0 10.0.0.1            locl  615  2
10.0.0.1/32       dest  0 10.0.0.1            locl  615  2
10.0.0.255/32     dest  0 10.0.0.255          bcst  613  1 ge-2/0/0.0
10.1.1.0/24       ifdn  0                               rslv  612  1 ge-2/0/1.0
10.1.1.0/32       iddn  0 10.1.1.0            recv  610  1 ge-2/0/1.0
10.1.1.1/32       user  0                               rjct  46   4
10.1.1.1/32       intf  0 10.1.1.1            locl  611  2
10.1.1.1/32       iddn  0 10.1.1.1            locl  611  2
10.1.1.255/32     iddn  0 ff:ff:ff:ff:ff:ff   bcst  609  1 ge-2/0/1.0
10.209.0.0/16     user  0 10.209.63.254        ucst  419  20 fxp0.0
10.209.0.0/16     user  1 0:12:1e:ca:98:0      ucst  419  20 fxp0.0
10.209.0.0/18     intf  0                               rslv  418  1 fxp0.0
10.209.0.0/32     dest  0 10.209.0.0          recv  416  1 fxp0.0
10.209.2.131/32   intf  0 10.209.2.131         locl  417  2
10.209.2.131/32   dest  0 10.209.2.131         locl  417  2
10.209.17.55/32   dest  0 0:30:48:5b:78:d2     ucst  435  1 fxp0.0
10.209.63.42/32   dest  0 0:23:7d:58:92:ca     ucst  434  1 fxp0.0
10.209.63.254/32   dest  0 0:12:1e:ca:98:0      ucst  419  20 fxp0.0
10.209.63.255/32   dest  0 10.209.63.255        bcst  415  1 fxp0.0
10.227.0.0/16     user  0 10.209.63.254        ucst  419  20 fxp0.0

...

Routing table: iso
ISO:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                               rjct  27   1
47.0005.80ff.f800.0000.0108.0003.0102.5524.5220.00
intf  0                               locl  28   1

Routing table: inet6
Internet6:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                               rjct  6    1
ff00::/8         perm  0                               mdsc  4    1
ff02::1/128       perm  0 ff02::1             mcst  3    1

Routing table: ccc
MPLS:
Interface.Label  Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                               rjct  16   1
100004(top)fe-0/0/1.0

```

show route forwarding-table detail

```

user@host> show route forwarding-table detail
Routing table: inet
Internet:
Destination      Type RtRef Next hop      Type Index NhRef Netif
default          user  2 0:90:69:8e:b1:1b ucst  132  4 fxp0.0
default          perm  0                               rjct  14  1
10.1.1.0/24      intf  0 ff.3.0.21      ucst  322  1 so-5/3/0.0
10.1.1.0/32      dest  0 10.1.1.0       recv  324  1 so-5/3/0.0
10.1.1.1/32      intf  0 10.1.1.1       locl  321  1
10.1.1.255/32    dest  0 10.1.1.255     bcst  323  1 so-5/3/0.0
10.21.21.0/24    intf  0 ff.3.0.21      ucst  326  1 so-5/3/0.0
10.21.21.0/32    dest  0 10.21.21.0     recv  328  1 so-5/3/0.0
10.21.21.1/32    intf  0 10.21.21.1     locl  325  1
10.21.21.255/32  dest  0 10.21.21.255   bcst  327  1 so-5/3/0.0
127.0.0.1/32     intf  0 127.0.0.1      locl  320  1
172.17.28.19/32  clon  1 192.168.4.254   ucst  132  4 fxp0.0
172.17.28.44/32  clon  1 192.168.4.254   ucst  132  4 fxp0.0

...

Routing table: private1__inet
Internet:
Destination      Type RtRef Next hop      Type Index NhRef Netif
default          perm  0                               rjct  46  1
10.0.0.0/8       intf  0                               rslv  136  1 fxp1.0
10.0.0.0/32      dest  0 10.0.0.0       recv  134  1 fxp1.0
10.0.0.4/32      intf  0 10.0.0.4       locl  135  2
10.0.0.4/32      dest  0 10.0.0.4       locl  135  2

...

Routing table: iso
ISO:
Destination      Type RtRef Next hop      Type Index NhRef Netif
default          perm  0                               rjct  38  1

Routing table: inet6
Internet6:
Destination      Type RtRef Next hop      Type Index NhRef Netif
default          perm  0                               rjct  22  1
ff00::/8         perm  0                               mdsc  21  1
ff02::1/128      perm  0 ff02::1       mcst  17  1

...

Routing table: mpls
MPLS:
Destination      Type RtRef Next hop      Type Index NhRef Netif
default          perm  0                               rjct  28  1

```

show route forwarding-table destination extensive (Weights and Balances)

```

user@host> show route forwarding-table destination 3.4.2.1 extensive
Routing table: inet [Index 0]
Internet:

Destination: 3.4.2.1/32
Route type: user
Route reference: 0                               Route interface-index: 0

```

Flags: sent to PFE		
Next-hop type: unicast	Index: 262143	Reference: 1
Nexthop: 4.4.4.4		
Next-hop type: unicast	Index: 335	Reference: 2
Next-hop interface: so-1/1/0.0	Weight: 22	Balance: 3
Nexthop: 145.12.1.2		
Next-hop type: unicast	Index: 337	Reference: 2
Next-hop interface: so-0/1/2.0	Weight: 33	Balance: 33

show route forwarding-table extensive

user@host> show route forwarding-table extensive

Routing table: inet [Index 0]

Internet:

Destination: default

Route type: user

Route reference: 2

Route interface-index: 0

Flags: sent to PFE

Nexthop: 0:90:69:8e:b1:1b

Next-hop type: unicast

Index: 132 Reference: 4

Next-hop interface: fxp0.0

Destination: default

Route type: permanent

Route reference: 0

Route interface-index: 0

Flags: none

Next-hop type: reject

Index: 14 Reference: 1

Destination: 127.0.0.1/32

Route type: interface

Route reference: 0

Route interface-index: 0

Flags: sent to PFE

Nexthop: 127.0.0.1

Next-hop type: local

Index: 320 Reference: 1

...

Routing table: private1__inet [Index 1]

Internet:

Destination: default

Route type: permanent

Route reference: 0

Route interface-index: 0

Flags: sent to PFE

Next-hop type: reject

Index: 46 Reference: 1

Destination: 10.0.0.0/8

Route type: interface

Route reference: 0

Route interface-index: 3

Flags: sent to PFE

Next-hop type: resolve

Next-hop interface: fxp1.0

Index: 136 Reference: 1

...

Routing table: iso [Index 0]

ISO:

Destination: default

Route type: permanent

```

Route reference: 0
Flags: sent to PFE
Next-hop type: reject
Route interface-index: 0
Index: 38      Reference: 1

Routing table: inet6 [Index 0]
Internet6:

Destination: default
Route type: permanent
Route reference: 0
Flags: sent to PFE
Next-hop type: reject
Route interface-index: 0
Index: 22      Reference: 1

Destination: ff00::/8
Route type: permanent
Route reference: 0
Flags: sent to PFE
Next-hop type: multicast discard
Route interface-index: 0
Index: 21      Reference: 1

...

Routing table: private1__inet6 [Index 1]
Internet6:

Destination: default
Route type: permanent
Route reference: 0
Flags: sent to PFE
Next-hop type: reject
Route interface-index: 0
Index: 54      Reference: 1

Destination: fe80::2a0:a5ff:fe3d:375/128
Route type: interface
Route reference: 0
Flags: sent to PFE
Next-hop: fe80::2a0:a5ff:fe3d:375
Next-hop type: local
Route interface-index: 0
Index: 75      Reference: 1

...

```

show route forwarding-table extensive (RPF)

The next example is based on the following configuration, which enables an RPF check on all routes that are learned from this interface, including the interface route:

```

so-1/1/0 {
  unit 0 {
    family inet {
      rpf-check;
      address 15.95.1.2/30;
    }
  }
}

```

```

user@host> show route forwarding-table extensive
Routing table: inet [Index 0]
Internet:
...
...
Destination: 15.95.1.3/32
Route type: destination
Route reference: 0
Route interface-index: 67

```

```

Flags: sent to PFE
Nexthop: 15.95.1.3
Next-hop type: broadcast          Index: 328      Reference: 1
Next-hop interface: so-1/1/0.0
RPF interface: so-1/1/0.0

```

show route forwarding-table family mpls

```

user@host> show route forwarding-table family mpls
Routing table: mpls
MPLS:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0
0                user  0
1                user  0
2                user  0
100000           user  0 10.31.1.6          swap 100001      fe-1/1/0.0
800002           user  0                  Pop                                vt-0/3/0.32770

vt-0/3/0.32770 (VPLS)
                    user  0                  indr  351      4
                    Push 800000, Push 100002(top)

so-0/0/0.0

```

show route forwarding-table family vpls

```

user@host> show route forwarding-table family vpls
Routing table: green.vpls
VPLS:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          dymn  0
default          perm  0
fe-0/1/0.0       dymn  0
00:90:69:0c:20:1f/48      <<<<<Remote CE

                    dymn  0                  indr  351      4
                    Push 800000, Push 100002(top)

so-0/0/0.0
00:90:69:85:b0:1f/48      <<<<<Local CE

                    dymn  0                  ucst  354      2 fe-0/1/0.0

```

show route forwarding-table vpls (Broadcast, unknown unicast, and multicast (BUM) hashing is enabled)

```

user@host> show route forwarding-table vpls
Routing table: green.vpls
VPLS:
Enabled protocols: BUM hashing
Destination      Type RtRef Next hop          Type Index      NhRef Netif
default          perm  0
lsi.1048832      intf  0
                    4.4.3.2          indr 1048574    4      2
                    Push 262145
ge-3/0/0.0
00:19:e2:25:d0:01/48 user  0                  ucst  590      5 ge-2/3/9.0
0x30003/51       user  0                  comp  627      2
ge-2/3/9.0       intf  0                  ucst  590      5 ge-2/3/9.0
ge-3/1/3.0       intf  0                  ucst  619      4 ge-3/1/3.0
0x30002/51       user  0                  comp  600      2
0x30001/51       user  0                  comp  597      2

```

show route forwarding-table vpls (Broadcast, unknown unicast, and multicast (BUM) hashing is enabled with MAC Statistics)

```

user@host> show route forwarding-table vpls
Routing table: green.vpls
VPLS:
Enabled protocols: BUM hashing, MAC Stats
Destination      Type RtRef Next hop      Type Index  NhRef Netif
default          perm  0         4.4.3.2      dscd   519      1
1si.1048834      intf  0         4.4.3.2      indr  1048574  4
Push 262145      592      2
ge-3/0/0.0
00:19:e2:25:d0:01/48 user  0         ucst   590      5 ge-2/3/9.0
0x30003/51      user  0         comp   630      2
ge-2/3/9.0      intf  0         ucst   590      5 ge-2/3/9.0
ge-3/1/3.0      intf  0         ucst   591      4 ge-3/1/3.0
0x30002/51      user  0         comp   627      2
0x30001/51      user  0         comp   624      2

```

show route forwarding-table family vpls extensive

```

user@host> show route forwarding-table family vpls extensive
Routing table: green.vpls [Index 2]
VPLS:

Destination: default
Route type: dynamic
Route reference: 0
Flags: sent to PFE
Next-hop type: flood
Next-hop type: unicast
Next-hop interface: fe-0/1/3.0
Next-hop type: unicast
Next-hop interface: fe-0/1/2.0
Route interface-index: 72
Index: 289      Reference: 1
Index: 291      Reference: 3
Index: 290      Reference: 3

Destination: default
Route type: permanent
Route reference: 0
Flags: none
Next-hop type: discard
Route interface-index: 0
Index: 341      Reference: 1

Destination: fe-0/1/2.0
Route type: dynamic
Route reference: 0
Flags: sent to PFE
Next-hop type: flood
Next-hop type: indirect
Next-hop type: Push 800016
Next-hop interface: at-1/0/1.0
Next-hop type: indirect
Next hop: 10.31.3.2
Next-hop type: Push 800000
Next-hop interface: fe-0/1/1.0
Next-hop type: unicast
Next-hop interface: fe-0/1/3.0
Route interface-index: 69
Index: 293      Reference: 1
Index: 363      Reference: 4
Index: 301      Reference: 5
Index: 291      Reference: 3

Destination: fe-0/1/3.0
Route type: dynamic
Route reference: 0
Flags: sent to PFE
Next-hop type: flood
Route interface-index: 70
Index: 292      Reference: 1

```

```

Next-hop type: indirect          Index: 363      Reference: 4
Next-hop type: Push 800016
Next-hop interface: at-1/0/1.0
Next-hop type: indirect          Index: 301      Reference: 5
Next hop: 10.31.3.2
Next-hop type: Push 800000
Next-hop interface: fe-0/1/1.0
Next-hop type: unicast           Index: 290      Reference: 3
Next-hop interface: fe-0/1/2.0

Destination: 10:00:00:01:01:01/48
Route type: dynamic
Route reference: 0                Route interface-index: 70
Flags: sent to PFE, prefix load balance
Next-hop type: unicast           Index: 291      Reference: 3
Next-hop interface: fe-0/1/3.0
Route used as destination:
  Packet count:      6640    Byte count:      675786
Route used as source
  Packet count:      6894    Byte count:      696424

Destination: 10:00:00:01:01:04/48
Route type: dynamic
Route reference: 0                Route interface-index: 69
Flags: sent to PFE, prefix load balance
Next-hop type: unicast           Index: 290      Reference: 3
Next-hop interface: fe-0/1/2.0
Route used as destination:
  Packet count:      96      Byte count:      8079
Route used as source:
  Packet count:      296      Byte count:      24955

Destination: 10:00:00:01:03:05/48
Route type: dynamic
Route reference: 0                Route interface-index: 74
Flags: sent to PFE, prefix load balance
Next-hop type: indirect          Index: 301      Reference: 5
Next hop: 10.31.3.2
Next-hop type: Push 800000
Next-hop interface: fe-0/1/1.0

```

show route forwarding-table table default

```

user@host> show route forwarding-table table default
Routing table: default.inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0
0.0.0.0/32       perm  0
10.0.60.0/30     user  0 10.0.60.13             ucst  713  5 fe-0/1/3.0
10.0.60.12/30    intf  0                       rslv  688  1 fe-0/1/3.0
10.0.60.12/32    dest  0 10.0.60.12             recv  686  1 fe-0/1/3.0
10.0.60.13/32    dest  0 0:5:85:8b:bc:22        ucst  713  5 fe-0/1/3.0
10.0.60.14/32    intf  0 10.0.60.14             locl  687  2
10.0.60.14/32    dest  0 10.0.60.14             locl  687  2
10.0.60.15/32    dest  0 10.0.60.15             bcst  685  1 fe-0/1/3.0
10.0.67.12/30    user  0 10.0.60.13             ucst  713  5 fe-0/1/3.0
10.0.80.0/30     ifdn  0 ff.3.0.21             ucst  676  1 so-0/0/1.0
10.0.80.0/32     dest  0 10.0.80.0             recv  678  1 so-0/0/1.0
10.0.80.2/32     user  0                       rjct  36   2
10.0.80.2/32     intf  0 10.0.80.2             locl  675  1

```



```

10.0.80.3/32      dest    0 10.0.80.3      bcst   677    1 so-0/0/1.0
10.0.90.12/30     intf    0                rslv   684    1 fe-0/1/0.0
10.0.90.12/32     dest    0 10.0.90.12    recv   682    1 fe-0/1/0.0
10.0.90.14/32     intf    0 10.0.90.14     locl   683    2
10.0.90.14/32     dest    0 10.0.90.14     locl   683    2
10.0.90.15/32     dest    0 10.0.90.15     bcst   681    1 fe-0/1/0.0
10.5.0.0/16       user    0 192.168.187.126 ucst   324    15 fxp0.0
10.10.0.0/16      user    0 192.168.187.126 ucst   324    15 fxp0.0
10.13.10.0/23     user    0 192.168.187.126 ucst   324    15 fxp0.0
10.84.0.0/16      user    0 192.168.187.126 ucst   324    15 fxp0.0
10.150.0.0/16     user    0 192.168.187.126 ucst   324    15 fxp0.0
10.157.64.0/19    user    0 192.168.187.126 ucst   324    15 fxp0.0
10.209.0.0/16     user    0 192.168.187.126 ucst   324    15 fxp0.0

```

...

Routing table: default.iso

ISO:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		rjct	60	1	

Routing table: default.inet6

Internet6:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		rjct	44	1	
::/128	perm	0		dscd	42	1	
ff00::/8	perm	0		mdsc	43	1	
ff02::1/128	perm	0	ff02::1	mcst	39	1	

Routing table: default.mpls

MPLS:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		dscd	50	1	

show route forwarding-table table logical-system-name/routing-instance-name

```
user@host> show route forwarding-table table R4/vpn-red
```

Logical system: R4

Routing table: vpn-red.inet

Internet:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		rjct	563	1	
0.0.0.0/32	perm	0		dscd	561	2	
1.0.0.1/32	user	0		dscd	561	2	
2.0.2.0/24	intf	0		rslv	771	1	ge-1/2/0.3
2.0.2.0/32	dest	0	2.0.2.0	recv	769	1	ge-1/2/0.3
2.0.2.1/32	intf	0	2.0.2.1	locl	770	2	
2.0.2.1/32	dest	0	2.0.2.1	locl	770	2	
2.0.2.2/32	dest	0	0.4.80.3.0.1b.c0.d5.e4.bd.0.1b.c0.d5.e4.bc.8.0	ucst	789	1	ge-1/2/0.3
2.0.2.255/32	dest	0	2.0.2.255	bcst	768	1	ge-1/2/0.3
224.0.0.0/4	perm	1		mdsc	562	1	
224.0.0.1/32	perm	0	224.0.0.1	mcst	558	1	
255.255.255.255/32	perm	0		bcst	559	1	

Logical system: R4

Routing table: vpn-red.iso

ISO:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		rjct	608	1	

```

Logical system: R4
Routing table: vpn-red.inet6
Internet6:
Destination      Type RtRef Next hop      Type Index NhRef Netif
default          perm  0              rjct  708   1
::/128           perm  0              dscd  706   1
ff00::/8         perm  0              mdsc  707   1
ff02::1/128     perm  0 ff02::1      mcst  704   1

```

```

Logical system: R4
Routing table: vpn-red.mpls
MPLS:
Destination      Type RtRef Next hop      Type Index NhRef Netif
default          perm  0              dscd  638

```

show route forwarding-table vpn

```

user@host> show route forwarding-table vpn VPN-A
Routing table:: VPN-A.inet
Internet:
Destination      Type RtRef Nexthop      Type Index NhRef Netif
default          perm  0              rjct   4    4
10.39.10.20/30   intf  0 ff.3.0.21      ucst   40    1
so-0/0/0.0
10.39.10.21/32   intf  0 10.39.10.21     locl   36    1
10.255.14.172/32 user   0              ucst   69    2
so-0/0/0.0
10.255.14.175/32 user   0              indr   81    3
Push 100004, Push
100004(top) so-1/0/0.0
224.0.0.0/4      perm  2              mdsc   5    3
224.0.0.1/32     perm  0 224.0.0.1      mcst   1    8
224.0.0.5/32     user   1 224.0.0.5      mcst   1    8
255.255.255.255/32 perm  0              bcst   2    3

```

show route inactive-path

List of Syntax	Syntax on page 3491 Syntax (EX Series Switches) on page 3491
Syntax	<pre>show route inactive-path <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switches)	<pre>show route inactive-path <brief detail extensive terse></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	Display routes for destinations that have no active route. An inactive route is a route that was not selected as the best path.
Options	<p>none—Display all inactive routes.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show route inactive-path on page 3491 show route inactive-path detail on page 3492 show route inactive-path extensive on page 3493 show route inactive-path terse on page 3493
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route inactive-path

```
user@host> show route inactive-path

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.12.100.12/30      [OSPF/10] 03:57:28, metric 1
> via so-0/3/0.0

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```

10.0.0.0/8          [Direct/0] 04:39:56
                    > via fxp1.0

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.12.80.0/30       [BGP/170] 04:38:17, localpref 100
                    AS path: 100 I
                    > to 10.12.80.1 via ge-6/3/2.0

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

bgp.l3vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

show route inactive-path detail

```

user@host> show route inactive-path detail

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete

10.12.100.12/30 (2 entries, 1 announced)
  OSPF   Preference: 10
         Next-hop reference count: 1
         Next hop: via so-0/3/0.0, selected
         State: <Int>
         Inactive reason: Route Preference
         Local AS:      1
         Age: 3:58:24   Metric: 1
         Area: 0.0.0.0
         Task: OSPF
         AS path: I

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

10.0.0.0/8 (2 entries, 0 announced)
  Direct Preference: 0
         Next hop type: Interface
         Next-hop reference count: 1
         Next hop: via fxp1.0, selected
         State: <NotBest Int>
         Inactive reason: No difference
         Age: 4:40:52
         Task: IF
         AS path: I

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete

10.12.80.0/30 (2 entries, 1 announced)

```

```

BGP      Preference: 170/-101
        Next-hop reference count: 6
        Source: 10.12.80.1
        Next hop: 10.12.80.1 via ge-6/3/2.0, selected
        State: <Ext>
        Inactive reason: Route Preference
        Peer AS: 100
        Age: 4:39:13
        Task: BGP_100.10.12.80.1+179
        AS path: 100 I
        Localpref: 100
        Router ID: 10.0.0.0

```

show route inactive-path extensive

The output for the **show route inactive-path extensive** command is identical to that of the **show route inactive-path detail** command. For sample output, see [show route inactive-path detail on page 3492](#).

show route inactive-path terse

```

user@host> show route inactive-path terse

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
  10.12.100.12/30   0 10      1          >so-0/3/0.0

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
  10.0.0.0/8        D  0          >fxp1.0

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
  10.12.80.0/30     B 170      100      >10.12.80.1    100 I

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

bgp.l3vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

show route inactive-prefix

List of Syntax	Syntax on page 3494 Syntax (EX Series Switches) on page 3494
Syntax	<code>show route inactive-prefix</code> <code><brief detail extensive terse></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switches)	<code>show route inactive-prefix</code> <code><brief detail extensive terse></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display inactive route destinations in each routing table.
Options	none —Display all inactive route destination. brief detail extensive terse —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show route inactive-prefix on page 3494 show route inactive-prefix detail on page 3494 show route inactive-prefix extensive on page 3495 show route inactive-prefix terse on page 3495
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route inactive-prefix

```
user@host> show route inactive-prefix

inet.0: 14 destinations, 14 routes (13 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

127.0.0.1/32          [Direct/0] 00:04:54
> via lo0.0
```

show route inactive-prefix detail

```
user@host> show route inactive-prefix detail

inet.0: 14 destinations, 14 routes (13 active, 0 holddown, 1 hidden)
```

```

127.0.0.1/32 (1 entry, 0 announced)
  Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
    State: <Hidden Martian Int>
    Age: 4:51
    Task: IF
    AS path: I00:04:54
      > via lo0.0

```

show route inactive-prefix extensive

The output for the **show route inactive-prefix extensive** command is identical to that of the **show route inactive-path detail** command. For sample output, see [show route inactive-prefix detail on page 3494](#).

show route inactive-prefix terse

```
user@host> show route inactive-prefix terse
```

```
inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
```

A Destination	P Prf	Metric 1	Metric 2	Next hop	AS path
127.0.0.1/32	D 0			>lo0.0	

show route instance

List of Syntax	Syntax on page 3496 Syntax (EX Series Switch and QFX Series) on page 3496
Syntax	<pre>show route instance <brief detail summary> <instance-name> <logical-system (all logical-system-name)> <operational></pre>
Syntax (EX Series Switch and QFX Series)	<pre>show route instance <brief detail summary> <instance-name> <operational></pre>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display routing instance information.
Options	<p>none—(Same as brief) Display standard information about all routing instances.</p> <p>brief detail summary—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief. (These options are not available with the operational keyword.)</p> <p>instance-name—(Optional) Display information for all routing instances whose name begins with this string (for example, cust1, cust11, and cust111 are all displayed when you run the show route instance cust1 command).</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>operational—(Optional) Display operational routing instances.</p>
Required Privilege Level	view
List of Sample Output	show route instance on page 3497 show route instance detail (Graceful Restart Complete) on page 3498 show route instance detail (Graceful Restart Incomplete) on page 3499 show route instance detail (VPLS Routing Instance) on page 3501 show route instance operational on page 3502 show route instance summary on page 3502
Output Fields	Table 163 lists the output fields for the show route instance command. Output fields are listed in the approximate order in which they appear.

Table 307: show route instance Output Fields

Field Name	Field Description	Level of Output
Instance or <i>instance-name</i>	Name of the routing instance.	All levels
Operational Routing Instances	(operational keyword only) Names of all operational routing instances.	—
Type	Type of routing instance: forwarding , l2vpn , no-forwarding , vpls , virtual-router , or vrf .	All levels
State	State of the routing instance: active or inactive .	brief detail none
Interfaces	Name of interfaces belonging to this routing instance.	brief detail none
Restart State	Status of graceful restart for this instance: Pending or Complete .	detail
Path selection timeout	Maximum amount of time, in seconds, remaining until graceful restart is declared complete. The default is 300 .	detail
Tables	Tables (and number of routes) associated with this routing instance.	brief detail none
Route-distinguisher	Unique route distinguisher associated with this routing instance.	detail
Vrf-import	VPN routing and forwarding instance import policy name.	detail
Vrf-export	VPN routing and forwarding instance export policy name.	detail
Vrf-import-target	VPN routing and forwarding instance import target community name.	detail
Vrf-export-target	VPN routing and forwarding instance export target community name.	detail
Fast-reroute-priority	Fast reroute priority setting for a VPLS routing instance: high , medium , or low . The default is low .	detail
Restart State	Restart state: <ul style="list-style-type: none"> Pending:protocol-name—List of protocols that have not yet completed graceful restart for this routing table. Complete—All protocols have restarted for this routing table. 	detail
Primary rib	Primary table for this routing instance.	brief none summary
Active/holddown/hidden	Number of active, hold-down, and hidden routes.	All levels

Sample Output

show route instance

```

user@host> show route instance
Instance          Type
      Primary RIB
master            forwarding
Active/holddown/hidden

```

```

inet.0                                16/0/1
iso.0                                 1/0/0
mpls.0                               0/0/0
inet6.0                              2/0/0
l2circuit.0                          0/0/0
__juniper_private1__ forwarding
__juniper_private1__.inet.0          12/0/0
__juniper_private1__.inet6.0         1/0/0

```

show route instance detail (Graceful Restart Complete)

```

user@host> show route instance detail
master:
  Router ID: 10.255.14.176
  Type: forwarding      State: Active
  Restart State: Complete Path selection timeout: 300
  Tables:
    inet.0                : 17 routes (15 active, 0 holddown, 1 hidden)
    Restart Complete
    inet.3                : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    iso.0                 : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
    mpls.0                : 19 routes (19 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l3vpn.0           : 10 routes (10 active, 0 holddown, 0 hidden)
    Restart Complete
    inet6.0               : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l2vpn.0           : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
  BGP-INET:
    Router ID: 10.69.103.1
    Type: vrf            State: Active
    Restart State: Complete Path selection timeout: 300
    Interfaces:
      t3-0/0/0.103
    Route-distinguisher: 10.255.14.176:103
    Vrf-import: [ BGP-INET-import ]
    Vrf-export: [ BGP-INET-export ]
    Tables:
      BGP-INET.inet.0      : 4 routes (4 active, 0 holddown, 0 hidden)
      Restart Complete
  BGP-L:
    Router ID: 10.69.104.1
    Type: vrf            State: Active
    Restart State: Complete Path selection timeout: 300
    Interfaces:
      t3-0/0/0.104
    Route-distinguisher: 10.255.14.176:104
    Vrf-import: [ BGP-L-import ]
    Vrf-export: [ BGP-L-export ]
    Tables:
      BGP-L.inet.0         : 4 routes (4 active, 0 holddown, 0 hidden)
      Restart Complete
      BGP-L.mpls.0         : 3 routes (3 active, 0 holddown, 0 hidden)
      Restart Complete
  L2VPN:
    Router ID: 0.0.0.0
    Type: l2vpn          State: Active
    Restart State: Complete Path selection timeout: 300

```

```

Interfaces:
  t3-0/0/0.512
Route-distinguisher: 10.255.14.176:512
Vrf-import: [ L2VPN-import ]
Vrf-export: [ L2VPN-export ]
Tables:
  L2VPN.l2vpn.0          : 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
LDP:
Router ID: 10.69.105.1
Type: vrf                State: Active
Restart State: Complete Path selection timeout: 300
Interfaces:
  t3-0/0/0.105
Route-distinguisher: 10.255.14.176:105
Vrf-import: [ LDP-import ]
Vrf-export: [ LDP-export ]
Tables:
  LDP.inet.0             : 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
OSPF:
Router ID: 10.69.101.1
Type: vrf                State: Active
Restart State: Complete Path selection timeout: 300
Interfaces:
  t3-0/0/0.101
Route-distinguisher: 10.255.14.176:101
Vrf-import: [ OSPF-import ]
Vrf-export: [ OSPF-export ]
Vrf-import-target: [ target:11111
Tables:
  OSPF.inet.0            : 8 routes (7 active, 0 holddown, 0 hidden)
Restart Complete
RIP:
Router ID: 10.69.102.1
Type: vrf                State: Active
Restart State: Complete Path selection timeout: 300
Interfaces:
  t3-0/0/0.102
Route-distinguisher: 10.255.14.176:102
Vrf-import: [ RIP-import ]
Vrf-export: [ RIP-export ]
Tables:
  RIP.inet.0              : 6 routes (6 active, 0 holddown, 0 hidden)
Restart Complete
STATIC:
Router ID: 10.69.100.1
Type: vrf                State: Active
Restart State: Complete Path selection timeout: 300
Interfaces:
  t3-0/0/0.100
Route-distinguisher: 10.255.14.176:100
Vrf-import: [ STATIC-import ]
Vrf-export: [ STATIC-export ]
Tables:
  STATIC.inet.0           : 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

```

show route instance detail (Graceful Restart Incomplete)

```
user@host> show route instance detail
```

```

master:
  Router ID: 10.255.14.176
  Type: forwarding          State: Active
  Restart State: Pending    Path selection timeout: 300
  Tables:
    inet.0                  : 17 routes (15 active, 1 holddown, 1 hidden)
    Restart Pending: OSPF LDP
    inet.3                  : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Pending: OSPF LDP
    iso.0                   : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
    mpls.0                  : 23 routes (23 active, 0 holddown, 0 hidden)
    Restart Pending: LDP VPN
    bgp.l3vpn.0             : 10 routes (10 active, 0 holddown, 0 hidden)
    Restart Pending: BGP VPN
    inet6.0                 : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l2vpn.0             : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Pending: BGP VPN
BGP-INET:
  Router ID: 10.69.103.1
  Type: vrf                  State: Active
  Restart State: Pending    Path selection timeout: 300
  Interfaces:
    t3-0/0/0.103
  Route-distinguisher: 10.255.14.176:103
  Vrf-import: [ BGP-INET-import ]
  Vrf-export: [ BGP-INET-export ]
  Tables:
    BGP-INET.inet.0        : 6 routes (5 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
BGP-L:
  Router ID: 10.69.104.1
  Type: vrf                  State: Active
  Restart State: Pending    Path selection timeout: 300
  Interfaces:
    t3-0/0/0.104
  Route-distinguisher: 10.255.14.176:104
  Vrf-import: [ BGP-L-import ]
  Vrf-export: [ BGP-L-export ]
  Tables:
    BGP-L.inet.0           : 6 routes (5 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
    BGP-L.mpls.0           : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
L2VPN:
  Router ID: 0.0.0.0
  Type: l2vpn                State: Active
  Restart State: Pending    Path selection timeout: 300
  Interfaces:
    t3-0/0/0.512
  Route-distinguisher: 10.255.14.176:512
  Vrf-import: [ L2VPN-import ]
  Vrf-export: [ L2VPN-export ]
  Tables:
    L2VPN.l2vpn.0          : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Pending: VPN L2VPN
LDP:
  Router ID: 10.69.105.1
  Type: vrf                  State: Active
  Restart State: Pending    Path selection timeout: 300

```

```

Interfaces:
  t3-0/0/0.105
Route-distinguisher: 10.255.14.176:105
Vrf-import: [ LDP-import ]
Vrf-export: [ LDP-export ]
Tables:
  LDP.inet.0          : 5 routes (4 active, 1 holddown, 0 hidden)
Restart Pending: OSPF LDP VPN
OSPF:
  Router ID: 10.69.101.1
  Type: vrf           State: Active
  Restart State: Pending Path selection timeout: 300
  Interfaces:
    t3-0/0/0.101
  Route-distinguisher: 10.255.14.176:101
  Vrf-import: [ OSPF-import ]
  Vrf-export: [ OSPF-export ]
  Tables:
    OSPF.inet.0       : 8 routes (7 active, 1 holddown, 0 hidden)
Restart Pending: OSPF VPN
RIP:
  Router ID: 10.69.102.1
  Type: vrf           State: Active
  Restart State: Pending Path selection timeout: 300
  Interfaces:
    t3-0/0/0.102
  Route-distinguisher: 10.255.14.176:102
  Vrf-import: [ RIP-import ]
  Vrf-export: [ RIP-export ]
  Tables:
    RIP.inet.0        : 8 routes (6 active, 2 holddown, 0 hidden)
Restart Pending: RIP VPN
STATIC:
  Router ID: 10.69.100.1
  Type: vrf           State: Active
  Restart State: Pending Path selection timeout: 300
  Interfaces:
    t3-0/0/0.100
  Route-distinguisher: 10.255.14.176:100
  Vrf-import: [ STATIC-import ]
  Vrf-export: [ STATIC-export ]
  Tables:
    STATIC.inet.0     : 4 routes (4 active, 0 holddown, 0 hidden)
Restart Pending: VPN

```

show route instance detail (VPLS Routing Instance)

```

user@host> show route instance detail test-vpls
test-vpls:
  Router ID: 0.0.0.0
  Type: vpls           State: Active
  Interfaces:
    lsi.1048833
    lsi.1048832
    fe-0/1/0.513
  Route-distinguisher: 10.255.37.65:1
  Vrf-import: [ __vrf-import-test-vpls-internal__ ]
  Vrf-export: [ __vrf-export-test-vpls-internal__ ]
  Vrf-import-target: [ target:300:1 ]
  Vrf-export-target: [ target:300:1 ]
  Fast-reroute-priority: high

```

Tables:
 test-vpls.12vpn.0 : 3 routes (3 active, 0 holddown, 0 hidden)

show route instance operational

```
user@host> show route instance operational
Operational Routing Instances:
```

```
master
default
```

show route instance summary

```
user@host> show route instance summary
```

Instance	Type	Primary rib	Active/holddown/hidden
master	forwarding		
		inet.0	15/0/1
		iso.0	1/0/0
		mpls.0	35/0/0
		13vpn.0	0/0/0
		inet6.0	2/0/0
		12vpn.0	0/0/0
		12circuit.0	0/0/0
BGP-INET	vrf		
		BGP-INET.inet.0	5/0/0
		BGP-INET.iso.0	0/0/0
		BGP-INET.inet6.0	0/0/0
BGP-L	vrf		
		BGP-L.inet.0	5/0/0
		BGP-L.iso.0	0/0/0
		BGP-L.mpls.0	4/0/0
		BGP-L.inet6.0	0/0/0
L2VPN	12vpn		
		L2VPN.inet.0	0/0/0
		L2VPN.iso.0	0/0/0
		L2VPN.inet6.0	0/0/0
		L2VPN.12vpn.0	2/0/0
LDP	vrf		
		LDP.inet.0	4/0/0
		LDP.iso.0	0/0/0
		LDP.mpls.0	0/0/0
		LDP.inet6.0	0/0/0
		LDP.12circuit.0	0/0/0
OSPF	vrf		
		OSPF.inet.0	7/0/0
		OSPF.iso.0	0/0/0
		OSPF.inet6.0	0/0/0
RIP	vrf		
		RIP.inet.0	6/0/0
		RIP.iso.0	0/0/0
		RIP.inet6.0	0/0/0
STATIC	vrf		
		STATIC.inet.0	4/0/0
		STATIC.iso.0	0/0/0
		STATIC.inet6.0	0/0/0

show route label

List of Syntax	Syntax on page 3503 Syntax (EX Series Switches) on page 3503
Syntax	show route label <i>label</i> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	show route label <i>label</i> <brief detail extensive terse>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
Description	Display the routes based on a specified Multiprotocol Label Switching (MPLS) label value.
Options	<p><i>label</i>—Value of the MPLS label.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> Example: Configuring Multipoint LDP In-Band Signaling for Point-to-Multipoint LSPs
List of Sample Output	show route label terse on page 3503 show route label on page 3504 show route label detail on page 3504 show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 3504 show route label detail (Multipoint LDP with Multicast-Only Fast Reroute) on page 3505 show route label extensive on page 3505
Output Fields	For information about output fields, see the output field table for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route label terse

```

user@host> show route label 100016 terse

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

```

A Destination	P Prf	Metric 1	Metric 2	Next hop	AS path
* 100016	V 170			>10.12.80.1	

show route label

```
user@host> show route label 100016
```

```
mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
100016          *[VPN/170] 03:25:41
                > to 10.12.80.1 via ge-6/3/2.0, Pop
```

show route label detail

```
user@host> show route label 100016 detail
```

```
mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
100016 (1 entry, 1 announced)
    *VPN      Preference: 170
              Next-hop reference count: 2
              Source: 10.12.80.1
              Next hop: 10.12.80.1 via ge-6/3/2.0, selected
              Label operation: Pop
              State: <Active Int Ext>
              Local AS: 1
              Age: 3:23:31
              Task: BGP.0.0.0.0+179
              Announcement bits (1): 0-KRT
              AS path: 100 I
              Ref Cnt: 2
```

show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show route label 299872 detail
```

```
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
299872 (1 entry, 1 announced)
    *LDP      Preference: 9
              Next hop type: Flood
              Next-hop reference count: 3
              Address: 0x9097d90
              Next hop: via vt-0/1/0.1
              Next-hop index: 661
              Label operation: Pop
              Address: 0x9172130
              Next hop: via so-0/0/3.0
              Next-hop index: 654
              Label operation: Swap 299872
              State: **Active Int>
              Local AS: 1001
              Age: 8:20      Metric: 1
              Task: LDP
              Announcement bits (1): 0-KRT
              AS path: I
              FECs bound to route: P2MP root-addr 10.255.72.166, grp 232.1.1.1,
src 192.168.142.2
```


show route label detail (Multipoint LDP with Multicast-Only Fast Reroute)

```

user@host> show route label 301568 detail

mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
301568 (1 entry, 1 announced)
  *LDP    Preference: 9
          Next hop type: Flood
          Address: 0x2735208
          Next-hop reference count: 3
          Next hop type: Router, Next hop index: 1397
          Address: 0x2735d2c
          Next-hop reference count: 3
          Next hop: 1.3.8.2 via ge-1/2/22.0
          Label operation: Pop
          Load balance label: None;
          Next hop type: Router, Next hop index: 1395
          Address: 0x2736290
          Next-hop reference count: 3
          Next hop: 1.3.4.2 via ge-1/2/18.0
          Label operation: Pop
          Load balance label: None;
          State: <Active Int AckRequest MulticastRPF>
          Local AS: 10
          Age: 54:05      Metric: 1
          Validation State: unverified
          Task: LDP
          Announcement bits (1): 0-KRT
          AS path: I
          FECs bound to route: P2MP root-addr 1.1.1.1, grp: 232.1.1.1, src:
192.168.219.11
          Primary Upstream : 1.1.1.3:0--1.1.1.2:0
          RPF Nexthops :
              ge-1/2/15.0, 1.2.94.1, Label: 301568, weight: 0x1
              ge-1/2/14.0, 1.2.3.1, Label: 301568, weight: 0x1
          Backup Upstream : 1.1.1.3:0--1.1.1.6:0
          RPF Nexthops :
              ge-1/2/20.0, 1.2.96.1, Label: 301584, weight: 0xffffe
              ge-1/2/19.0, 1.3.6.1, Label: 301584, weight: 0xffffe

```

show route label extensive

The output for the **show route label extensive** command is identical to that of the **show route label detail** command. For sample output, see [show route label detail on page 3504](#).

show route label-switched-path

List of Syntax	Syntax on page 3506 Syntax (EX Series Switches) on page 3506
Syntax	<code>show route label-switched-path <i>path-name</i></code> <code><brief detail extensive terse></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switches)	<code>show route label-switched-path <i>path-name</i></code> <code><brief detail extensive terse></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
Description	Display the routes used in an MPLS label-switched path (LSP).
Options	<code>brief detail extensive terse</code> —(Optional) Display the specified level of output. <code><i>path-name</i></code> —LSP tunnel name. <code>logical-system (all <i>logical-system-name</i>)</code> —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show route label-switched-path on page 3506
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route label-switched-path

```
user@host> show route label-switched-path sf-to-ny
inet.0: 29 destinations, 29 routes (29 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.1/32          [MPLS/7] 00:00:06, metric 0
> to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny
3.3.3.3/32          *[MPLS/7] 00:00:06, metric 0
> to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2.2.2.2/32          *[MPLS/7] 00:00:06, metric 0
> to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny
4.4.4.4/32          *[MPLS/7] 00:00:06, metric 0
> to 111.222.1.9 via s0-0/0/0, label-switched-path abc
> to 111.222.1.9 via s0-0/0/0, label-switched-path xyz
> to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny
```

```
111.222.1.9/32      [MPLS/7] 00:00:06, metric 0  
                  > to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny
```

```
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)  
+ = Active Route, - = Last Active, * = Both
```

```
mpls.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)  
+ = Active Route, - = Last Active, * = Both
```

show route martians

List of Syntax [Syntax on page 3508](#)
[Syntax \(EX Series Switches\) on page 3508](#)

Syntax show route martians
 <logical-system (all | *logical-system-name*)>
 <table *routing-table-name*>

Syntax (EX Series Switches) show route martians
 <table *routing-table-name*>

Release Information Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.0 for EX Series switches.

Description Display the martian (invalid and ignored) entries associated with each routing table.

Options **none**—Display standard information about route martians for all routing tables.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

table *routing-table-name*—(Optional) Display information about route martians for all routing tables whose name begins with this string (for example, **inet.0** and **inet6.0** are both displayed when you run the **show route martians table inet** command).

Required Privilege Level view

Related Documentation

- *Example: Configuring Class E Martian Addresses for Routing*
- *Understanding Martian Addresses*

List of Sample Output [show route martians on page 3509](#)

Output Fields [Table 308](#) lists the output fields for the **show route martians** command. Output fields are listed in the approximate order in which they appear

Table 308: show route martians Output Fields

Field Name	Field Description
<i>table-name</i>	Name of the route table in which the route martians reside.
<i>destination-prefix</i>	Route destination.
<i>match value</i>	Route match parameter.
<i>status</i>	Status of the route: allowed or disallowed .

Sample Output

show route martians

```

user@host> show route martians

inet.0:
    0.0.0.0/0 exact -- allowed
    0.0.0.0/8 orlonger -- disallowed
    127.0.0.0/8 orlonger -- disallowed
    192.0.0.0/24 orlonger -- disallowed
    240.0.0.0/4 orlonger -- disallowed
    224.0.0.0/4 exact -- disallowed
    224.0.0.0/24 exact -- disallowed

inet.1:
    0.0.0.0/0 exact -- allowed
    0.0.0.0/8 orlonger -- disallowed
    127.0.0.0/8 orlonger -- disallowed
    192.0.0.0/24 orlonger -- disallowed
    240.0.0.0/4 orlonger -- disallowed

inet.2:
    0.0.0.0/0 exact -- allowed
    0.0.0.0/8 orlonger -- disallowed
    127.0.0.0/8 orlonger -- disallowed
    192.0.0.0/24 orlonger -- disallowed
    240.0.0.0/4 orlonger -- disallowed
    224.0.0.0/4 exact -- disallowed
    224.0.0.0/24 exact -- disallowed

inet.3:
    0.0.0.0/0 exact -- allowed
    0.0.0.0/8 orlonger -- disallowed
    127.0.0.0/8 orlonger -- disallowed
    192.0.0.0/24 orlonger -- disallowed
    240.0.0.0/4 orlonger -- disallowed
    224.0.0.0/4 exact -- disallowed
    224.0.0.0/24 exact -- disallowed

...

inet6.0:
    ::1/128 exact -- disallowed
    ff00::/8 exact -- disallowed
    ff02::/16 exact -- disallowed

inet6.1:
    ::1/128 exact -- disallowed

inet6.2:
    ::1/128 exact -- disallowed
    ff00::/8 exact -- disallowed
    ff02::/16 exact -- disallowed

inet6.3:
    ::1/128 exact -- disallowed
    ff00::/8 exact -- disallowed
    ff02::/16 exact -- disallowed

...

```

show route next-hop

List of Syntax	Syntax on page 3510 Syntax (EX Series Switches) on page 3510
Syntax	<code>show route next-hop <i>next-hop</i></code> <code><brief detail extensive terse></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switches)	<code>show route next-hop <i>next-hop</i></code> <code><brief detail extensive terse></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the entries in the routing table that are being sent to the specified next-hop address.
Options	<code>brief detail extensive terse</code> —(Optional) Display the specified level of output. <code>logical-system (all <i>logical-system-name</i>)</code> —(Optional) Perform this operation on all logical systems or on a particular logical system. <code>next-hop</code> —Next-hop address.
Required Privilege Level	view
List of Sample Output	show route next-hop on page 3510 show route next-hop detail on page 3511 show route next-hop extensive on page 3513 show route next-hop terse on page 3514
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route next-hop

```
user@host> show route next-hop 192.168.71.254

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.10.0.0/16      *[Static/5] 06:26:25
                  > to 192.168.71.254 via fxp0.0
10.209.0.0/16    *[Static/5] 06:26:25
                  > to 192.168.71.254 via fxp0.0
172.16.0.0/12    *[Static/5] 06:26:25
                  > to 192.168.71.254 via fxp0.0
192.168.0.0/16   *[Static/5] 06:26:25
```

```

> to 192.168.71.254 via fxp0.0
192.168.102.0/23  *[Static/5] 06:26:25
> to 192.168.71.254 via fxp0.0
207.17.136.0/24  *[Static/5] 06:26:25
> to 192.168.71.254 via fxp0.0
207.17.136.192/32 *[Static/5] 06:26:25
> to 192.168.71.254 via fxp0.0

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

show route next-hop detail

```

user@host> show route next-hop 192.168.71.254 detail

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
Restart Complete
10.10.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 36
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 1
    Age: 6:27:41
    Task: RT
    Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
    AS path: I

10.209.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 36
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 1
    Age: 6:27:41
    Task: RT
    Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
    AS path: I

172.16.0.0/12 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 36
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 1
    Age: 6:27:41
    Task: RT
    Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2

```

```
AS path: I

192.168.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 36
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 1
    Age: 6:27:41
    Task: RT
    Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
    AS path: I

192.168.102.0/23 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 36
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 1
    Age: 6:27:41
    Task: RT
    Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
    AS path: I

207.17.136.0/24 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 36
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 1
    Age: 6:27:41
    Task: RT
    Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
    AS path: I

207.17.136.192/32 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 36
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 1
    Age: 6:27:41
    Task: RT
    Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
    AS path: I

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```


show route next-hop extensive

```
user@host> show route next-hop 192.168.71.254 extensive
```

```
inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
```

```
10.10.0.0/16 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
```

```
*Static Preference: 5
```

```
Next-hop reference count: 22
```

```
Next hop: 192.168.71.254 via fxp0.0, selected
```

```
State: <Active NoReadvrt Int Ext>
```

```
Local AS: 69
```

```
Age: 2:02:28
```

```
Task: RT
```

```
Announcement bits (1): 0-KRT
```

```
AS path: I
```

```
10.209.0.0/16 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 10.209.0.0/16 -> {192.168.71.254}
```

```
*Static Preference: 5
```

```
Next-hop reference count: 22
```

```
Next hop: 192.168.71.254 via fxp0.0, selected
```

```
State: <Active NoReadvrt Int Ext>
```

```
Local AS: 69
```

```
Age: 2:02:28
```

```
Task: RT
```

```
Announcement bits (1): 0-KRT
```

```
AS path: I
```

```
172.16.0.0/12 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 172.16.0.0/12 -> {192.168.71.254}
```

```
*Static Preference: 5
```

```
Next-hop reference count: 22
```

```
Next hop: 192.168.71.254 via fxp0.0, selected
```

```
State: <Active NoReadvrt Int Ext>
```

```
Local AS: 69
```

```
Age: 2:02:28
```

```
Task: RT
```

```
Announcement bits (1): 0-KRT
```

```
AS path: I
```

```
192.168.0.0/16 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 192.168.0.0/16 -> {192.168.71.254}
```

```
*Static Preference: 5
```

```
Next-hop reference count: 22
```

```
Next hop: 192.168.71.254 via fxp0.0, selected
```

```
State: <Active NoReadvrt Int Ext>
```

```
Local AS: 69
```

```
Age: 2:02:28
```

```
Task: RT
```

```
Announcement bits (1): 0-KRT
```

```
AS path: I
```

```
192.168.102.0/23 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 192.168.102.0/23 -> {192.168.71.254}
```

```
*Static Preference: 5
```

```

Next-hop reference count: 22
Next hop: 192.168.71.254 via fxp0.0, selected
State: <Active NoReadvrt Int Ext>
Local AS: 69
Age: 2:02:28
Task: RT
Announcement bits (1): 0-KRT
AS path: I

207.17.136.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.0/24 -> {192.168.71.254}
*Static Preference: 5
Next-hop reference count: 22
Next hop: 192.168.71.254 via fxp0.0, selected
State: <Active NoReadvrt Int Ext>
Local AS: 69
Age: 2:02:28
Task: RT
Announcement bits (1): 0-KRT
AS path: I

207.17.136.192/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.192/32 -> {192.168.71.254}
*Static Preference: 5
Next-hop reference count: 22
Next hop: 192.168.71.254 via fxp0.0, selected
State: <Active NoReadvrt Int Ext>
Local AS: 69
Age: 2:02:28
Task: RT
Announcement bits (1): 0-KRT
AS path: I

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

green.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

red.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

show route next-hop terse

```

user@host> show route next-hop 192.168.71.254 terse

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
* 10.10.0.0/16     S  5         0         0  >192.168.71.254
* 10.209.0.0/16    S  5         0         0  >192.168.71.254
* 172.16.0.0/12    S  5         0         0  >192.168.71.254

```

```
* 192.168.0.0/16      S   5                >192.168.71.254
* 192.168.102.0/23   S   5                >192.168.71.254
* 207.17.136.0/24    S   5                >192.168.71.254
* 207.17.136.192/32 S   5                >192.168.71.254

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route no-community

List of Syntax	Syntax on page 3516 Syntax (EX Series Switches) on page 3516
Syntax	show route no-community <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	show route no-community <brief detail extensive terse>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the route entries in each routing table that are not associated with any community.
Options	<p>none—(Same as brief) Display the route entries in each routing table that are not associated with any community.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show route no-community on page 3516 show route no-community detail on page 3517 show route no-community extensive on page 3517 show route no-community terse on page 3518
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route no-community

```

user@host> show route no-community
inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.0.0/16      *[Static/5] 00:36:27
> to 192.168.71.254 via fxp0.0
10.209.0.0/16    *[Static/5] 00:36:27
> to 192.168.71.254 via fxp0.0
10.255.71.52/32  *[Direct/0] 00:36:27
> via lo0.0
10.255.71.63/32  *[OSPF/10] 00:04:39, metric 1
> to 35.1.1.2 via ge-3/1/0.0
10.255.71.64/32  *[OSPF/10] 00:00:08, metric 2

```

```

> to 35.1.1.2 via ge-3/1/0.0
10.255.71.240/32 * [OSPF/10] 00:05:04, metric 2
                  via so-0/1/2.0
> via so-0/3/2.0
10.255.71.241/32 * [OSPF/10] 00:05:14, metric 1
> via so-0/1/2.0
10.255.71.242/32 * [OSPF/10] 00:05:19, metric 1
> via so-0/3/2.0
12.1.1.0/24      * [OSPF/10] 00:05:14, metric 2
> via so-0/3/2.0
14.1.1.0/24      * [OSPF/10] 00:00:08, metric 3
> to 35.1.1.2 via ge-3/1/0.0
                  via so-0/1/2.0
                  via so-0/3/2.0
16.1.1.0/24      * [OSPF/10] 00:05:14, metric 2
> via so-0/1/2.0
.....

```

show route no-community detail

```

user@host> show route no-community detail

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 38:08
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

10.209.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 38:08
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

....

```

show route no-community extensive

```

user@host> show route no-community extensive

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 2:03:33
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

```

```

10.209.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.209.0.0/16 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 2:03:33
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

```

show route no-community terse

```
user@host> show route no-community terse
```

```

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

```

A	Destination	P	Prf	Metric 1	Metric 2	Next hop	AS path
*	10.10.0.0/16	S	5			>192.168.71.254	
*	10.209.0.0/16	S	5			>192.168.71.254	
*	10.255.71.52/32	D	0			>100.0	
*	10.255.71.63/32	0	10	1		>35.1.1.2	
*	10.255.71.64/32	0	10	2		>35.1.1.2	
*	10.255.71.240/32	0	10	2		so-0/1/2.0	
						>so-0/3/2.0	
*	10.255.71.241/32	0	10	1		>so-0/1/2.0	
*	10.255.71.242/32	0	10	1		>so-0/3/2.0	
*	12.1.1.0/24	0	10	2		>so-0/3/2.0	
*	14.1.1.0/24	0	10	3		>35.1.1.2	
						so-0/1/2.0	
						so-0/3/2.0	
*	16.1.1.0/24	0	10	2		>so-0/1/2.0	
...							

show route protocol

List of Syntax	Syntax on page 3519 Syntax (EX Series Switches) on page 3519
Syntax	<pre>show route protocol <i>protocol</i> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switches)	<pre>show route protocol <i>protocol</i> <brief detail extensive terse></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Options ospf2 and ospf3 introduced in Junos OS Release 9.2.</p> <p>Options ospf2 and ospf3 introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Option flow introduced in Junos OS Release 10.0.</p> <p>Option flow introduced in Junos OS Release 10.0 for EX Series switches.</p>
Description	Display the route entries in the routing table that were learned from a particular protocol.
Options	<p>brief detail extensive terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><i>protocol</i>—Protocol from which the route was learned:</p> <ul style="list-style-type: none"> • access—Access route for use by DHCP application • access-internal—Access-internal route for use by DHCP application • aggregate—Locally generated aggregate route • atmvpn—Asynchronous Transfer Mode virtual private network • bgp—Border Gateway Protocol • ccc—Circuit cross-connect • direct—Directly connected route • dvmrp—Distance Vector Multicast Routing Protocol • esis—End System-to-Intermediate System • flow—Locally defined flow-specification route. • isis—Intermediate System-to-Intermediate System • ldp—Label Distribution Protocol • l2circuit—Layer 2 circuit • l2vpn—Layer 2 virtual private network • local—Local address

- **mpls**—Multiprotocol Label Switching
- **msdp**—Multicast Source Discovery Protocol
- **ospf**—Open Shortest Path First versions 2 and 3
- **ospf2**—Open Shortest Path First version 2 only
- **ospf3**—Open Shortest Path First version 3 only
- **pim**—Protocol Independent Multicast
- **rip**—Routing Information Protocol
- **ripng**—Routing Information Protocol next generation
- **rsvp**—Resource Reservation Protocol
- **rtarget**—Local route target virtual private network
- **static**—Statically defined route
- **tunnel**—Dynamic tunnel
- **vpn**—Virtual private network



NOTE: EX Series switches run a subset of these protocols. See the switch CLI for details.

Required Privilege Level	view
List of Sample Output	show route protocol access on page 3521 show route protocol access-internal extensive on page 3521 show route protocol bgp on page 3521 show route protocol bgp detail on page 3521 show route protocol bgp extensive on page 3522 show route protocol bgp terse on page 3522 show route protocol direct on page 3522 show route protocol l2circuit detail on page 3523 show route protocol l2vpn extensive on page 3524 show route protocol ldp on page 3524 show route protocol ldp extensive on page 3525 show route protocol ospf (Layer 3 VPN) on page 3526 show route protocol ospf detail on page 3527 show route protocol rip on page 3527 show route protocol rip detail on page 3527 show route protocol ripng table inet6 on page 3527
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route protocol access

```
user@host> show route protocol access
inet.0: 30380 destinations, 30382 routes (30379 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

13.160.0.3/32      *[Access/13] 00:00:09
                  > to 13.160.0.2 via fe-0/0/0.0
13.160.0.4/32      *[Access/13] 00:00:09
                  > to 13.160.0.2 via fe-0/0/0.0
13.160.0.5/32      *[Access/13] 00:00:09
                  > to 13.160.0.2 via fe-0/0/0.0
```

show route protocol access-internal extensive

```
user@host> show route protocol access-internal 13.160.0.19 extensive
inet.0: 100020 destinations, 100022 routes (100019 active, 0 holddown, 1 hidden)
13.160.0.19/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 13.160.0.19/32 -> {13.160.0.2}
    *Access-internal Preference: 12
      Next-hop reference count: 200000
      Next hop: 13.160.0.2 via fe-0/0/0.0, selected
      State: <Active Int>
    Age: 36
      Task: RPD Unix Domain Server./var/run/rpd_serv.local
      Announcement bits (1): 0-KRT
      AS path: I
```

show route protocol bgp

```
user@host> show route protocol bgp 192.168.64.0/21
inet.0: 335832 destinations, 335833 routes (335383 active, 0 holddown, 450 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.64.0/21    *[BGP/170] 6d 10:41:16, localpref 100, from 192.168.69.71
                  AS path: 10458 14203 2914 4788 4788 I
                  > to 192.168.167.254 via fxp0.0
```

show route protocol bgp detail

```
show route protocol bgp 66.117.63.0/24 exact detail
inet.0: 335805 destinations, 335806 routes (335356 active, 0 holddown, 450 hidden)
66.117.63.0/24 (1 entry, 1 announced)
    *BGP Preference: 170/-101
      Next hop type: Indirect
      Next-hop reference count: 1006436
      Source: 192.168.69.71
      Next hop type: Router, Next hop index: 324
      Next hop: 192.168.167.254 via fxp0.0, selected
      Protocol next hop: 192.168.69.71
      Indirect next hop: 8e166c0 342
      State: <Active Ext>
      Local AS: 69 Peer AS: 10458
      Age: 6d 10:42:42 Metric2: 0
      Task: BGP_10458.192.168.69.71+179
      Announcement bits (3): 0-KRT 2-BGP RT Background 3-Resolve tree
1
      AS path: 10458 14203 2914 4788 4788 I
```

```

Communities: 2914:410 2914:2403 2914:3400
Accepted
Localpref: 100
Router ID: 207.17.136.192

```

show route protocol bgp extensive

```
user@host> show route protocol bgp 192.168.64.0/21 extensive
```

```

inet.0: 335827 destinations, 335828 routes (335378 active, 0 holddown, 450 hidden)
192.168.64.0/21 (1 entry, 1 announced)
TSI:
KRT in-kernel 1.9.0.0/16 -> {indirect(342)}
Page 0 idx 1 Type 1 val db31a80
  Nexthop: Self
  AS path: [69] 10458 14203 2914 4788 4788 I
  Communities: 2914:410 2914:2403 2914:3400
Path 1.9.0.0 from 192.168.69.71 Vector len 4. Val: 1
  *BGP Preference: 170/-101
    Next hop type: Indirect
    Next-hop reference count: 1006502
    Source: 192.168.69.71
    Next hop type: Router, Next hop index: 324
    Next hop: 192.168.167.254 via fxp0.0, selected
    Protocol next hop: 192.168.69.71
    Indirect next hop: 8e166c0 342
    State: <Active Ext>
    Local AS: 69 Peer AS: 10458
    Age: 6d 10:44:45 Metric2: 0
    Task: BGP_10458.192.168.69.71+179
    Announcement bits (3): 0-KRT 2-BGP RT Background 3-Resolve tree

```

```
1
```

```

AS path: 10458 14203 2914 4788 4788 I
Communities: 2914:410 2914:2403 2914:3400
Accepted
Localpref: 100
Router ID: 207.17.136.192
Indirect next hops: 1
  Protocol next hop: 192.168.69.71
  Indirect next hop: 8e166c0 342
  Indirect path forwarding next hops: 1
    Next hop type: Router
    Next hop: 192.168.167.254 via fxp0.0
192.168.0.0/16 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
    Nexthop: 192.168.167.254 via fxp0.0

```

show route protocol bgp terse

```
user@host> show route protocol bgp 192.168.64.0/21 terse
```

```

inet.0: 24 destinations, 32 routes (23 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

```

A Destination	P Prf	Metric 1	Metric 2	Next hop	AS path
192.168.64.0/21	B 170	100		>100.1.3.2	10023 21 I

show route protocol direct

```
user@host> show route protocol direct
```

```
inet.0: 335843 destinations, 335844 routes (335394 active, 0 holddown, 450 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
8.8.8.0/24      *[Direct/0] 17w0d 10:31:49
                 > via fe-1/3/1.0
10.255.165.1/32 *[Direct/0] 25w4d 04:13:18
                 > via lo0.0
30.30.30.0/24   *[Direct/0] 17w0d 23:06:26
                 > via fe-1/3/2.0
192.168.164.0/22 *[Direct/0] 25w4d 04:13:20
                 > via fxp0.0
```

```
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
47.0005.80ff.f800.0000.0108.0001.0102.5516.5001/152
                 *[Direct/0] 25w4d 04:13:21
                 > via lo0.0
```

```
inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
abcd::10:255:165:1/128
                 *[Direct/0] 25w4d 04:13:21
                 > via lo0.0
fe80::2a0:a5ff:fe12:ad7/128
                 *[Direct/0] 25w4d 04:13:21
                 > via lo0.0
```

show route protocol l2circuit detail

```
user@host> show route protocol l2circuit detail
```

```
mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
100000 (1 entry, 1 announced)
```

```
*L2CKT Preference: 7
Next hop: via ge-2/0/0.0, selected
Label operation: Pop      Offset: 4
State: <Active Int>
Local AS: 99
Age: 9:52
Task: Common L2 VC
Announcement bits (1): 0-KRT
AS path: I
```

```
ge-2/0/0.0 (1 entry, 1 announced)
```

```
*L2CKT Preference: 7
Next hop: via so-1/1/2.0 weight 1, selected
Label-switched-path my-lsp
Label operation: Push 100000, Push 100000(top)[0] Offset: -4
Protocol next hop: 10.245.255.63
Push 100000 Offset: -4
Indirect next hop: 86af0c0 298
State: <Active Int>
Local AS: 99
Age: 9:52
Task: Common L2 VC
Announcement bits (2): 0-KRT 1-Common L2 VC
AS path: I
```

```

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
  *L2CKT Preference: 7
    Next hop: via so-1/1/2.0 weight 1, selected
    Label-switched-path my-lsp
    Label operation: Push 100000[0]
    Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
    State: <Active Int>
    Local AS: 99
    Age: 10:21
    Task: l2 circuit
    Announcement bits (1): 0-LDP
    AS path: I
    VC Label 100000, MTU 1500, VLAN ID 512

```

show route protocol l2vpn extensive

```

user@host> show route protocol l2vpn extensive

inet.0: 14 destinations, 15 routes (13 active, 0 holddown, 1 hidden)

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
800001 (1 entry, 1 announced)
TSI:
KRT in-kernel 800001 /36 -> {so-0/0/0.0}
  *L2VPN Preference: 7
    Next hop: via so-0/0/0.0 weight 49087 balance 97%, selected
    Label operation: Pop Offset: 4
    State: <Active Int>
    Local AS: 69
    Age: 7:48
    Task: Common L2 VC
    Announcement bits (1): 0-KRT
    AS path: I

so-0/0/0.0 (1 entry, 1 announced)
TSI:
KRT in-kernel so-0/0/0.0 /16 -> {indirect(288)}
  *L2VPN Preference: 7
    Next hop: via so-0/0/1.0, selected
    Label operation: Push 800000 Offset: -4
    Protocol next hop: 10.255.14.220
    Push 800000 Offset: -4
    Indirect next hop: 85142a0 288
    State: <Active Int>
    Local AS: 69
    Age: 7:48
    Task: Common L2 VC
    Announcement bits (2): 0-KRT 1-Common L2 VC
    AS path: I
    Communities: target:69:1 Layer2-info: encaps:PPP,
    control flags:2, mtu: 0

```

show route protocol ldp

```

user@host> show route protocol ldp

```

```

inet.0: 12 destinations, 13 routes (12 active, 0 holddown, 0 hidden)

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.16.1/32    *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Push 100000
192.168.17.1/32    *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0

private1___.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100064            *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Pop
100064(S=0)        *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Pop
100080            *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Swap 100000

```

show route protocol ldp extensive

```

user@host> show route protocol ldp extensive
192.168.16.1/32 (1 entry, 1 announced)
  State: <FlashAll>
  *LDP    Preference: 9
          Next-hop reference count: 3
          Next hop: via t1-4/0/0.0, selected
          Label operation: Push 100000
          State: <Active Int>
          Local AS: 65500
          Age: 1d 23:03:58      Metric: 1
          Task: LDP
          Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 2
          AS path: I

192.168.17.1/32 (1 entry, 1 announced)
  State: <FlashAll>
  *LDP    Preference: 9
          Next-hop reference count: 3
          Next hop: via t1-4/0/0.0, selected
          State: <Active Int>
          Local AS: 65500
          Age: 1d 23:03:58      Metric: 1
          Task: LDP
          Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 2
          AS path: I

private1___.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)

100064 (1 entry, 1 announced)
TSI:
KRT in-kernel 100064 /36 -> {t1-4/0/0.0}
  *LDP    Preference: 9
          Next-hop reference count: 2
          Next hop: via t1-4/0/0.0, selected
          State: <Active Int>

```

```

Local AS: 65500
Age: 1d 23:03:58      Metric: 1
Task: LDP
Announcement bits (1): 0-KRT
AS path: I
Prefixes bound to route: 192.168.17.1/32

100064(S=0) (1 entry, 1 announced)
TSI:
KRT in-kernel 100064 /40 -> {t1-4/0/0.0}
    *LDP      Preference: 9
              Next-hop reference count: 2
              Next hop: via t1-4/0/0.0, selected
              Label operation: Pop
              State: <Active Int>
              Local AS: 65500
              Age: 1d 23:03:58      Metric: 1
              Task: LDP
              Announcement bits (1): 0-KRT
              AS path: I

100080 (1 entry, 1 announced)
TSI:
KRT in-kernel 100080 /36 -> {t1-4/0/0.0}
    *LDP      Preference: 9
              Next-hop reference count: 2
              Next hop: via t1-4/0/0.0, selected
              Label operation: Swap 100000
              State: <Active Int>
              Local AS: 65500
              Age: 1d 23:03:58      Metric: 1
              Task: LDP
              Announcement bits (1): 0-KRT
              AS path: I
              Prefixes bound to route: 192.168.16.1/32

```

show route protocol ospf (Layer 3 VPN)

```

user@host> show route protocol ospf
inet.0: 40 destinations, 40 routes (39 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.4/30      *[OSPF/10] 00:05:18, metric 4
                  > via t3-3/2/0.0
10.39.1.8/30      [OSPF/10] 00:05:18, metric 2
                  > via t3-3/2/0.0
10.255.14.171/32 *[OSPF/10] 00:05:18, metric 4
                  > via t3-3/2/0.0
10.255.14.179/32 *[OSPF/10] 00:05:18, metric 2
                  > via t3-3/2/0.0
224.0.0.5/32     *[OSPF/10] 20:25:55, metric 1

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.16/30     [OSPF/10] 00:05:43, metric 1
                  > via so-0/2/2.0
10.255.14.173/32 *[OSPF/10] 00:05:43, metric 1
                  > via so-0/2/2.0
224.0.0.5/32     *[OSPF/10] 20:26:20, metric 1

```

show route protocol ospf detail

```

user@host> show route protocol ospf detail
VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.16/30 (2 entries, 0 announced)
    OSPF   Preference: 10
           Nexthop: via so-0/2/2.0, selected
           State: <Int>
           Inactive reason: Route Preference
           Age: 6:25      Metric: 1
           Area: 0.0.0.0
           Task: VPN-AB-OSPF
           AS path: I
           Communities: Route-Type:0.0.0.0:1:0

...

```

show route protocol rip

```

user@host> show route protocol rip
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32  * [RIP/100] 20:24:34, metric 2
                  > to 10.39.1.22 via t3-0/2/2.0
224.0.0.9/32      * [RIP/100] 00:03:59, metric 1

```

show route protocol rip detail

```

user@host> show route protocol rip detail
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32 (1 entry, 1 announced)
    *RIP   Preference: 100
           Nexthop: 10.39.1.22 via t3-0/2/2.0, selected
           State: <Active Int>
           Age: 20:25:02  Metric: 2
           Task: VPN-AB-RIPv2
           Announcement bits (2): 0-KRT 2-BGP.0.0.0.0+179
           AS path: I
           Route learned from 10.39.1.22 expires in 96 seconds

```

show route protocol ripng table inet6

```

user@host> show route protocol ripng table inet6
inet6.0: 4215 destinations, 4215 routes (4214 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

1111::1/128      * [RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::2/128      * [RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::3/128      * [RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0

```

```
1111::4/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::5/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::6/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
```


show route range

List of Syntax	Syntax on page 3529 Syntax (EX Series Switches) on page 3529
Syntax	<pre>show route range <brief detail extensive terse> <destination-prefix> <logical-system (all logical-system-name)></pre>
Syntax (EX Series Switches)	<pre>show route range <brief detail extensive terse> <destination-prefix></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	Display routing table entries using a prefix range.
Options	<p>none—Display standard information about all routing table entries using a prefix range.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>destination-prefix—Destination and prefix mask for the range.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show route range on page 3529 show route range destination-prefix on page 3530 show route range detail on page 3530 show route range extensive on page 3531 show route range terse on page 3532
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route range

```
user@host> show route range
```

```
inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
10.10.0.0/16      *[Static/5] 00:30:01
                  > to 192.168.71.254 via fxp0.0
10.209.0.0/16    *[Static/5] 00:30:01
```

```
10.255.71.14/32      > to 192.168.71.254 via fxp0.0
                    *[Direct/0] 00:30:01
                    > via lo0.0
172.16.0.0/12       *[Static/5] 00:30:01
                    > to 192.168.71.254 via fxp0.0
192.168.0.0/16      *[Static/5] 00:30:01
                    > to 192.168.71.254 via fxp0.0
192.168.64.0/21     *[Direct/0] 00:30:01
                    > via fxp0.0
192.168.71.14/32    *[Local/0] 00:30:01
                    Local via fxp0.0
192.168.102.0/23    *[Static/5] 00:30:01
                    > to 192.168.71.254 via fxp0.0
...
```

show route range destination-prefix

```
user@host> show route range 192.168.0.0/16

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.0.0/16      *[Static/5] 00:31:14
                    > to 192.168.71.254 via fxp0.0
192.168.64.0/21     *[Direct/0] 00:31:14
                    > via fxp0.0
192.168.71.14/32    *[Local/0] 00:31:14
                    Local via fxp0.0
192.168.102.0/23    *[Static/5] 00:31:14
                    > to 192.168.71.254 via fxp0.0
```

show route range detail

```
user@host> show route range detail

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 30:05
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

10.209.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 30:05
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

10.255.71.14/32 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
```

```

State: <Active Int>
Age: 30:05
Task: IF
AS path: I

172.16.0.0/12 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 30:05
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

...

```

show route range extensive

```

user@host> show route range extensive

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 30:17
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

10.209.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.209.0.0/16 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 30:17
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

10.255.71.14/32 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
    State: <Active Int>
    Age: 30:17
    Task: IF
    AS path: I

172.16.0.0/12 (1 entry, 1 announced)
TSI:
KRT in-kernel 172.16.0.0/12 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22

```

```

Next hop: 192.168.71.254 via fxp0.0, selected
State: <Active NoReadvrt Int Ext>
Age: 30:17
Task: RT
Announcement bits (1): 0-KRT
AS path: I

```

```
...
```

show route range terse

```
user@host> show route range terse
```

```
inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
```

A	Destination	P	Prf	Metric 1	Metric 2	Next hop	AS path
*	10.10.0.0/16	S	5			>192.168.71.254	
*	10.209.0.0/16	S	5			>192.168.71.254	
*	10.255.71.14/32	D	0			>lo0.0	
*	172.16.0.0/12	S	5			>192.168.71.254	
*	192.168.0.0/16	S	5			>192.168.71.254	
*	192.168.64.0/21	D	0			>fxp0.0	
*	192.168.71.14/32	L	0			Local	
*	192.168.102.0/23	S	5			>192.168.71.254	
*	207.17.136.0/24	S	5			>192.168.71.254	
*	207.17.136.192/32	S	5			>192.168.71.254	

```
__juniper_private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

A	Destination	P	Prf	Metric 1	Metric 2	Next hop	AS path
*	10.0.0.0/8	D	0			>fxp2.0	
		D	0			>fxp1.0	
*	10.0.0.4/32	L	0			Local	

```
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

A	Destination	P	Prf	Metric 1	Metric 2	Next hop	AS path
	47.0005.80ff.f800.0000.0108.0001.0102.5507.1014/152						
*		D	0			>lo0.0	

```
inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

A	Destination	P	Prf	Metric 1	Metric 2	Next hop	AS path
	abcd::10:255:71:14/128						
*		D	0			>lo0.0	
	fe80::280:42ff:fe11:226f/128						
*		D	0			>lo0.0	

```
__juniper_private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

A	Destination	P	Prf	Metric 1	Metric 2	Next hop	AS path
	fe80::280:42ff:fe11:226f/128						
*		D	0			>lo0.16385	

show route receive-protocol

List of Syntax	Syntax on page 3533 Syntax (EX Series Switches) on page 3533
Syntax	show route receive-protocol <i>protocol neighbor-address</i> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	show route receive-protocol <i>protocol neighbor-address</i> <brief detail extensive terse>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the routing information as it was received through a particular neighbor using a particular dynamic routing protocol.
Options	<p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><i>protocol neighbor-address</i>—Protocol transmitting the route (bgp, dvmrp, msdp, pim, rip, or ripng) and address of the neighboring router from which the route entry was received.</p>
Additional Information	The output displays the selected routes and the attributes with which they were received, but does not show the effects of import policy on the routing attributes.
Required Privilege Level	view
List of Sample Output	show route receive-protocol bgp on page 3536 show route receive-protocol bgp extensive on page 3536 show route receive-protocol bgp table extensive on page 3536 show route receive-protocol bgp logical-system extensive on page 3537 show route receive-protocol bgp detail (Layer 2 VPN) on page 3538 show route receive-protocol bgp extensive (Layer 2 VPN) on page 3538 show route receive-protocol bgp (Layer 3 VPN) on page 3539 show route receive-protocol bgp detail (Layer 3 VPN) on page 3539 show route receive-protocol bgp detail (Long-Lived Graceful Restart) on page 3540 show route receive-protocol bgp extensive (Layer 3 VPN) on page 3540
Output Fields	Table 309 describes the output fields for the show route receive-protocol command. Output fields are listed in the approximate order in which they appear.

Table 309: show route receive-protocol Output Fields

Field Name	Field Description	Level of Output
<i>routing-table-name</i>	Name of the routing table—for example, inet.0.	All levels
<i>number destinations</i>	Number of destinations for which there are routes in the routing table.	All levels
<i>number routes</i>	Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> • active • holdddown (routes that are in pending state before being declared inactive) • hidden (routes that are not used because of a routing policy) 	All levels
Prefix	Destination prefix.	none brief
MED	Multiple exit discriminator value included in the route.	none brief
<i>destination-prefix</i> (entry, announced)	Destination prefix. The entry value is the number of routes for this destination, and the announced value is the number of routes being announced for this destination.	detail extensive
Accepted LongLivedStale	The LongLivedStale flag indicates that the route was marked LLGR-stale by this router, as part of the operation of LLGR receiver mode. Either this flag or the LongLivedStaleImport flag may be displayed for a route. Neither of these flags are displayed at the same time as the Stale (ordinary GR stale) flag.	detail extensive
Accepted LongLivedStaleImport	The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy. Either this flag or the LongLivedStale flag may be displayed for a route. Neither of these flags are displayed at the same time as the Stale (ordinary GR stale) flag. Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and import into the inet.0 routing table	detail extensive
ImportAccepted LongLivedStaleImport	Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and imported into the inet.0 routing table The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy.	detail extensive
Route Distinguisher	64-bit prefix added to IP subnets to make them unique.	detail extensive
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.	detail extensive
VPN Label	Virtual private network (VPN) label. Packets are sent between CE and PE routing devices by advertising VPN labels. VPN labels transit over either an RSVP or an LDP label-switched path (LSP) tunnel.	detail extensive

Table 309: show route receive-protocol Output Fields (*continued*)

Field Name	Field Description	Level of Output
Next hop	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.	All levels
Localpref or Lclpref	Local preference value included in the route.	All levels
AS path	<p>Autonomous system (AS) path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> • I—IGP. • E—EGP. • ?—Incomplete; typically, the AS path was aggregated. <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> • []—Brackets enclose the number that precedes the AS path. This number represents the number of ASs present in the AS path, when calculated as defined in RFC 4271. This value is used the AS-path merge process, as defined in RFC 4893. • []—If more than one AS number is configured on the router, or if AS path prepending is configured, brackets enclose the local AS number associated with the AS path. • { }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • ()—Parentheses enclose a confederation. • ([])—Parentheses and brackets enclose a confederation set. <p>NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p>	All levels
Cluster list	(For route reflected output only) Cluster ID sent by the route reflector.	detail extensive
Originator ID	(For route reflected output only) Address of routing device that originally sent the route to the route reflector.	detail extensive
Communities	Community path attribute for the route. See the Output Field table in the show route detail command for all possible values for this field.	detail extensive
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.	detail extensive
Attrset AS	Number, local preference, and path of the AS that originated the route. These values are stored in the Attrset attribute at the originating routing device.	detail extensive
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).	detail extensive
control flags	Control flags: none or Site Down .	detail extensive

Table 309: show route receive-protocol Output Fields (*continued*)

Field Name	Field Description	Level of Output
mtu	Maximum transmission unit (MTU) of the Layer 2 circuit.	detail extensive

Sample Output

show route receive-protocol bgp

```

user@host> show route receive-protocol bgp 10.255.245.215

inet.0: 28 destinations, 33 routes (27 active, 0 holddown, 1 hidden)
Prefix          Next hop          MED      Lclpref  AS path
10.22.1.0/24     10.255.245.215    0        100      I
10.22.2.0/24     10.255.245.215    0        100      I

```

show route receive-protocol bgp extensive

```

user@host> show route receive-protocol bgp 10.255.245.63 extensive
inet.0: 244 destinations, 244 routes (243 active, 0 holddown, 1 hidden)
Prefix          Next hop          MED      Lclpref  AS path
1.1.1.0/24 (1 entry, 1 announced)
  Next hop: 10.0.50.3
  Localpref: 100
  AS path: I <Originator>
  Cluster list: 10.2.3.1
  Originator ID: 10.255.245.45
165.3.0.0/16 (1 entry, 1 announced)
  Next hop: 111.222.5.254
  Localpref: 100
  AS path: I <Originator>
  Cluster list: 10.2.3.1
  Originator ID: 10.255.245.68
165.4.0.0/16 (1 entry, 1 announced)
  Next hop: 111.222.5.254
  Localpref: 100
  AS path: I <Originator>
  Cluster list: 10.2.3.1
  Originator ID: 10.255.245.45
195.1.2.0/24 (1 entry, 1 announced)
  Next hop: 111.222.5.254
  Localpref: 100
  AS path: I <Originator>
  Cluster list: 10.2.3.1
  Originator ID: 10.255.245.68
inet.2: 63 destinations, 63 routes (63 active, 0 holddown, 0 hidden)
Prefix          Next hop          MED      Lclpref  AS path
inet.3: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix          Next hop          MED      Lclpref  AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Next hop          MED      Lclpref  AS path
mpls.0: 48 destinations, 48 routes (48 active, 0 holddown, 0 hidden)

```

show route receive-protocol bgp table extensive

```

user@host> show route receive-protocol bgp 207.17.136.192 table inet.0 66.117.68.0/24 extensive
inet.0: 227315 destinations, 227316 routes (227302 active, 0 holddown, 13 hidden)
* 66.117.63.0/24 (1 entry, 1 announced)
  Nexthop: 207.17.136.29

```



```

Localpref: 100
AS path: AS2 PA[6]: 14203 2914 3356 29748 33437 AS_TRANS
AS path: AS4 PA[2]: 33437 393219
AS path: Merged[6]: 14203 2914 3356 29748 33437 393219 I
Communities: 2914:420

```

show route receive-protocol bgp logical-system extensive

```

user@host> show route receive-protocol bgp 10.0.0.9 logical-system PE4 extensive
inet.0: 12 destinations, 13 routes (12 active, 0 holddown, 0 hidden)
* 10.0.0.0/30 (1 entry, 1 announced)
    Accepted
    Route Label: 3
    Nexthop: 10.0.0.9
    AS path: 13979 I

* 10.0.0.4/30 (1 entry, 1 announced)
    Accepted
    Route Label: 3
    Nexthop: 10.0.0.9
    AS path: 13979 I

10.0.0.8/30 (2 entries, 1 announced)
    Accepted
    Route Label: 3
    Nexthop: 10.0.0.9
    AS path: 13979 I

* 10.9.9.1/32 (1 entry, 1 announced)
    Accepted
    Route Label: 3
    Nexthop: 10.0.0.9
    AS path: 13979 I

* 10.100.1.1/32 (1 entry, 1 announced)
    Accepted
    Route Label: 3
    Nexthop: 10.0.0.9
    AS path: 13979 I

* 44.0.0.0/24 (1 entry, 1 announced)
    Accepted
    Route Label: 300096
    Nexthop: 10.0.0.9
    AS path: 13979 I
    AIGP: 203

* 55.0.0.0/24 (1 entry, 1 announced)
    Accepted
    Route Label: 300112
    Nexthop: 10.0.0.9
    AS path: 13979 7018 I
    AIGP: 25

* 66.0.0.0/24 (1 entry, 1 announced)
    Accepted
    Route Label: 300144
    Nexthop: 10.0.0.9
    AS path: 13979 7018 I

* 99.0.0.0/24 (1 entry, 1 announced)

```

```

Accepted
Route Label: 300160
Nexthop: 10.0.0.9
AS path: 13979 7018 I

```

show route receive-protocol bgp detail (Layer 2 VPN)

```

user@host> show route receive-protocol bgp 10.255.14.171 detail
inet.0: 68 destinations, 68 routes (67 active, 0 holddown, 1 hidden)
Prefix          Nexthop          MED    Lclpref AS path
inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
frame-vpn.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0
hidden)
Prefix          Nexthop          MED    Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 1 announced)
  Route Distinguisher: 10.255.245.35:1
  Label-base : 800000, range : 4, status-vector : 0x0
  Nexthop: 10.255.245.35
  Localpref: 100
  AS path: I
  Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
  control flags: 0, mtu: 0
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 0 announced)
  Route Distinguisher: 10.255.245.35:1
  Label-base : 800000, range : 4, status-vector : 0x0
  Nexthop: 10.255.245.35
  Localpref: 100
  AS path: I
  Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
  control flags: 0, mtu: 0

```

show route receive-protocol bgp extensive (Layer 2 VPN)

```

user@host> show route receive-protocol bgp 10.255.14.171 extensive
inet.0: 68 destinations, 68 routes (67 active, 0 holddown, 1 hidden)
Prefix          Nexthop          MED    Lclpref AS path
inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
frame-vpn.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 1 announced)
  Route Distinguisher: 10.255.245.35:1
  Label-base : 800000, range : 4, status-vector : 0x0
  Nexthop: 10.255.245.35
  Localpref: 100
  AS path: I
  Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
  control flags: 0, mtu: 0
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path

```

```

10.255.245.35:1:5:1/96 (1 entry, 0 announced)
  Route Distinguisher: 10.255.245.35:1
  Label-base : 800000, range : 4, status-vector : 0x0
  Nexthop: 10.255.245.35
  Localpref: 100
  AS path: I
  Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
  control flags:0, mtu: 0

```

show route receive-protocol bgp (Layer 3 VPN)

```

user@host> show route receive-protocol bgp 10.255.14.171
inet.0: 33 destinations, 33 routes (32 active, 0 holddown, 1 hidden)
Prefix          Nexthop          MED      Lclpref AS path
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
VPN-A.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
10.255.14.175/32 10.255.14.171          100 2 I
10.255.14.179/32 10.255.14.171          2    100 I
VPN-B.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
10.255.14.175/32 10.255.14.171          100 2 I
10.255.14.177/32 10.255.14.171          100 I
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
mpls.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
bgp.l3vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
10.255.14.171:300:10.255.14.177/32
                  10.255.14.171          100 I
10.255.14.171:100:10.255.14.179/32
                  10.255.14.171          2    100 I
10.255.14.171:200:10.255.14.175/32
                  10.255.14.171          100 2 I

```

show route receive-protocol bgp detail (Layer 3 VPN)

```

user@host> show route receive-protocol bgp 10.255.14.174 detail
inet.0: 16 destinations, 17 routes (15 active, 0 holddown, 1 hidden)
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
vpna.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
* 10.49.0.0/30 (1 entry, 1 announced)
  Route Distinguisher: 10.255.14.176:2
  VPN Label: 101264
  Nexthop: 10.255.14.174
  Localpref: 100
  AS path: I
  Communities: target:200:100
  AttrSet AS: 100
    Localpref: 100
    AS path: I
* 10.255.14.172/32 (1 entry, 1 announced)
  Route Distinguisher: 10.255.14.176:2
  VPN Label: 101280
  Nexthop: 10.255.14.174
  Localpref: 100
  AS path: I
  Communities: target:200:100
  AttrSet AS: 100

```

```

        Localpref: 100
        AS path: I
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
* 10.255.14.174:2:10.49.0.0/30 (1 entry, 0 announced)
    Route Distinguisher: 10.255.14.174:2
    VPN Label: 101264
    Nexthop: 10.255.14.174
    Localpref: 100
    AS path: I
    Communities: target:200:100
    AttrSet AS: 100
        Localpref: 100
        AS path: I
* 10.255.14.174:2:10.255.14.172/32 (1 entry, 0 announced)
    Route Distinguisher: 10.255.14.174:2
    VPN Label: 101280
    Nexthop: 10.255.14.174
    Localpref: 100
    AS path: I
    Communities: target:200:100
    AttrSet AS: 100
        Localpref: 100
        AS path: I
inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

```

show route receive-protocol bgp detail (Long-Lived Graceful Restart)

```

user@host> show route receive-protocol bgp 10.4.12.11 detail

bgp.l2vpn.0: 38 destinations, 39 routes (37 active, 0 holddown, 1 hidden)
* 1.1.1.4:100:1.1.1.4/96 AD (1 entry, 1 announced)
    Accepted LongLivedStale LongLivedStaleImport
    Nexthop: 10.4.12.11
    Localpref: 100
    AS path: I

```

show route receive-protocol bgp extensive (Layer 3 VPN)

```

user@host> show route receive-protocol bgp 10.255.245.63 extensive
inet.0: 244 destinations, 244 routes (243 active, 0 holddown, 1 hidden)
  Prefix                Nexthop                MED    Lclpref AS path
  1.1.1.0/24 (1 entry, 1 announced)
    Nexthop: 10.0.50.3
    Localpref: 100
    AS path: I <Originator>
    Cluster list: 10.2.3.1
    Originator ID: 10.255.245.45
  165.3.0.0/16 (1 entry, 1 announced)
    Nexthop: 111.222.5.254
    Localpref: 100
    AS path: I <Originator>
    Cluster list: 10.2.3.1
    Originator ID: 10.255.245.68
  165.4.0.0/16 (1 entry, 1 announced)
    Nexthop: 111.222.5.254
    Localpref: 100
    AS path: I <Originator>
    Cluster list: 10.2.3.1
    Originator ID: 10.255.245.45

```

```
195.1.2.0/24 (1 entry, 1 announced)
  Nexthop: 111.222.5.254
  Localpref: 100
  AS path: I <Originator>
  Cluster list: 10.2.3.1
  Originator ID: 10.255.245.68
inet.2: 63 destinations, 63 routes (63 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
inet.3: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
mpls.0: 48 destinations, 48 routes (48 active, 0 holddown, 0 hidden)
```

show route resolution

List of Syntax	Syntax on page 3542 Syntax (EX Series Switches) on page 3542
Syntax	<pre>show route resolution <brief detail extensive summary> <index <i>index</i>> <logical-system (all <i>logical-system-name</i>)> <prefix> <table <i>routing-table-name</i>> <unresolved></pre>
Syntax (EX Series Switches)	<pre>show route resolution <brief detail extensive summary> <index <i>index</i>> <prefix> <table <i>routing-table-name</i>> <unresolved></pre>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the entries in the next-hop resolution database. This database provides for recursive resolution of next hops through other prefixes in the routing table.
Options	<p>none—Display standard information about all entries in the next-hop resolution database.</p> <p>brief detail extensive summary—(Optional) Display the specified level of output.</p> <p>index <i>index</i>—(Optional) Show the index of the resolution tree.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>prefix <i>network/destination-prefix</i>—(Optional) Display database entries for the specified address.</p> <p>table <i>routing-table-name</i>—(Optional) Display information about a particular routing table (for example, inet.0) where policy-based export is currently enabled.</p> <p>unresolved—(Optional) Display routes that could not be resolved.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Route Resolution on PE Routers
List of Sample Output	show route resolution detail on page 3543 show route resolution summary on page 3544 show route resolution unresolved on page 3544

Output Fields Table 310 describes the output fields for the **show route resolution** command. Output fields are listed in the approximate order in which they appear.

Table 310: show route resolution Output Fields

Field Name	Field Description
<i>routing-table-name</i>	Name of the routing table whose prefixes are resolved using the entries in the route resolution database. For routing table groups, this is the name of the primary routing table whose prefixes are resolved using the entries in the route resolution database.
Tree index	Tree index identifier.
Nodes	Number of nodes in the tree.
Reference count	Number of references made to the next hop.
Contributing routing tables	Routing tables used for next-hop resolution.
Originating RIB	Name of the routing table whose active route was used to determine the forwarding next-hop entry in the resolution database. For example, in the case of inet.0 resolving through inet.0 and inet.3 , this field indicates which routing table, inet.0 or inet.3 , provided the best path for a particular prefix.
Metric	Metric associated with the forwarding next hop.
Node path count	Number of nodes in the path.
Forwarding next hops	Number of forwarding next hops. The forwarding next hop is the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.

Sample Output

show route resolution detail

```

user@host> show route resolution detail
Tree Index: 1, Nodes 0, Reference Count 1
Contributing routing tables: inet.3
Tree Index: 2, Nodes 23, Reference Count 1
Contributing routing tables: inet.0 inet.3
10.10.0.0/16 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
10.31.1.0/30 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
10.31.1.1/32 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 0
10.31.1.4/30 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
10.31.1.5/32 Originating RIB: inet.0

```

```
Node path count: 1
Forwarding nexthops: 0
10.31.2.0/30 Originating RIB: inet.0
Metric: 2 Node path count: 1
Forwarding nexthops: 2
10.31.11.0/24 Originating RIB: inet.0
Node path count: 1
Forwarding nexthops: 1
```

show route resolution summary

```
user@host> show route resolution summary
Tree Index: 1, Nodes 24, Reference Count 1
Contributing routing tables: :voice.inet.0 :voice.inet.3
Tree Index: 2, Nodes 2, Reference Count 1
Contributing routing tables: inet.3
Tree Index: 3, Nodes 43, Reference Count 1
Contributing routing tables: inet.0 inet.3
```

show route resolution unresolved

```
user@host> show route resolution unresolved
Tree Index 1
vt-3/2/0.32769.0      /16
  Protocol Nexthop: 10.255.71.238 Push 800000
  Indirect nexthop: 0 -
vt-3/2/0.32772.0      /16
  Protocol Nexthop: 10.255.70.103 Push 800008
  Indirect nexthop: 0 -
Tree Index 2
```


show route snooping

Syntax	<pre>show route snooping <brief detail extensive terse> <all> <best address/prefix> <exact address> <logical-system logical-system-name> <range prefix-range> <summary> <table table-name></pre>
Release Information	<p>Command introduced in Junos OS Release 8.5.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	Display the entries in the routing table that were learned from snooping.
Options	<p>none—Display the entries in the routing table that were learned from snooping.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>all—(Optional) Display all entries, including hidden entries.</p> <p>best address/prefix—(Optional) Display the longest match for the provided address and optional prefix.</p> <p>exact address/prefix—(Optional) Display exact matches for the provided address and optional prefix.</p> <p>logical-system logical-system-name—(Optional) Display information about a particular logical system, or type 'all'.</p> <p>range prefix-range—(Optional) Display information for the provided address range.</p> <p>summary—(Optional) Display route snooping summary statistics.</p> <p>table table-name—(Optional) Display information for the named table.</p>
Required Privilege Level	view
List of Sample Output	<p>show route snooping detail on page 3545</p> <p>show route snooping logical-system all on page 3546</p>
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route snooping detail

```
user@host> show route snooping detail
```

```

__+domainAll__.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

224.0.0.2/32 (1 entry, 1 announced)
  *IGMP   Preference: 0
          Next hop type: MultiRecv
          Next-hop reference count: 4
          State: <Active NoReadvrt Int>
          Age: 2:24
          Task: IGMP
          Announcement bits (1): 0-KRT
          AS path: I

224.0.0.22/32 (1 entry, 1 announced)
  *IGMP   Preference: 0
          Next hop type: MultiRecv
          Next-hop reference count: 4
          State: <Active NoReadvrt Int>
          Age: 2:24
          Task: IGMP
          Announcement bits (1): 0-KRT
          AS path: I

__+domainAll__.inet.1: 36 destinations, 36 routes (36 active, 0 holddown, 0 hidden)

224.0.0.0.0.0.0.0/24 (1 entry, 1 announced)
  *Multicast Preference: 180
          Next hop type: Multicast (IPv4), Next hop index: 1048584
          Next-hop reference count: 4
          State: <Active Int>
          Age: 2:24
          Task: MC
          Announcement bits (1): 0-KRT
          AS path: I

<snip>

```

show route snooping logical-system all

```

user@host> show route snooping logical-system all

logical-system: default

inet.1: 20 destinations, 20 routes (20 active, 0 holddown, 0 hidden)
Restart Unsupported
+ = Active Route, - = Last Active, * = Both

0.0,0.1,0.0,232.1.1.65,100.1.1.2/112*[Multicast/180] 00:07:36
      Multicast (IPv4) Composite
0.0,0.1,0.0,232.1.1.66,100.1.1.2/112*[Multicast/180] 00:07:36
      Multicast (IPv4) Composite
0.0,0.1,0.0,232.1.1.67,100.1.1.2/112*[Multicast/180] 00:07:36

<snip>

default-switch.inet.1: 237 dest, 237 rts (237 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

0.15,0.1,0.0,0.0.0.0,0.0.0.0,2/120*[Multicast/180] 00:08:21
      Multicast (IPv4) Composite
0.15,0.1,0.0,0.0.0.0,0.0.0.0,2,17/128*[Multicast/180] 00:08:21

```

Multicast (IPv4) Composite

<snip>

show route source-gateway

List of Syntax	Syntax on page 3548 Syntax (EX Series Switches) on page 3548
Syntax	<code>show route source-gateway address</code> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	<code>show route source-gateway address</code> <brief detail extensive terse>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the entries in the routing table that were learned from a particular address. The Source field in the <code>show route detail</code> command output lists the source for each route, if known.
Options	brief detail extensive terse —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief . address —IP address of the system. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show route source-gateway on page 3548 show route source-gateway detail on page 3549 show route source-gateway extensive on page 3551
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route source-gateway

```
user@host> show route source-gateway 10.255.70.103
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete
```

```

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.255.70.103:1:3:1/96
    *[BGP/170] 12:12:24, localpref 100, from 10.255.70.103
    AS path: I
    > via so-0/3/0.0, label-switched-path green-r1-r3

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.255.70.103:2:3:1/96
    *[BGP/170] 12:12:24, localpref 0, from 10.255.70.103
    AS path: I
    > via so-0/3/0.0, label-switched-path green-r1-r3

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.255.70.103:1:3:1/96
    *[BGP/170] 12:12:24, localpref 100, from 10.255.70.103
    AS path: I
    > via so-0/3/0.0, label-switched-path green-r1-r3

10.255.70.103:2:3:1/96
    *[BGP/170] 12:12:24, localpref 0, from 10.255.70.103
    AS path: I
    > via so-0/3/0.0, label-switched-path green-r1-r3

```

show route source-gateway detail

```

user@host> show route source-gateway 10.255.70.103 detail
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete
green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

```

Restart Complete

10.255.70.103:1:3:1/96 (1 entry, 1 announced)

```
*BGP    Preference: 170/-101
        Route Distinguisher: 10.255.70.103:1
        Next-hop reference count: 7
        Source: 10.255.70.103
        Protocol next hop: 10.255.70.103
        Indirect next hop: 2 no-forward
        State: <Secondary Active Int Ext>
        Local AS:    69 Peer AS:    69
        Age: 12:14:00  Metric2: 1
        Task: BGP_69.10.255.70.103+179
        Announcement bits (1): 0-green-12vpn
        AS path: I
        Communities: target:11111:1 Layer2-info: encaps:VPLS,
        control flags:, mtu: 0
        Label-base: 800008, range: 8
        Localpref: 100
        Router ID: 10.255.70.103
        Primary Routing Table bgp.12vpn.0
```

red.12vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

Restart Complete

10.255.70.103:2:3:1/96 (1 entry, 1 announced)

```
*BGP    Preference: 170/-1
        Route Distinguisher: 10.255.70.103:2
        Next-hop reference count: 7
        Source: 10.255.70.103
        Protocol next hop: 10.255.70.103
        Indirect next hop: 2 no-forward
        State: <Secondary Active Int Ext>
        Local AS:    69 Peer AS:    69
        Age: 12:14:00  Metric2: 1
        Task: BGP_69.10.255.70.103+179
        Announcement bits (1): 0-red-12vpn
        AS path: I
        Communities: target:11111:2 Layer2-info: encaps:VPLS,
        control flags:Site-Down, mtu: 0
        Label-base: 800016, range: 8
        Localpref: 0
        Router ID: 10.255.70.103
        Primary Routing Table bgp.12vpn.0
```

bgp.12vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

Restart Complete

10.255.70.103:1:3:1/96 (1 entry, 0 announced)

```
*BGP    Preference: 170/-101
        Route Distinguisher: 10.255.70.103:1
        Next-hop reference count: 7
        Source: 10.255.70.103
        Protocol next hop: 10.255.70.103
        Indirect next hop: 2 no-forward
        State: <Active Int Ext>
        Local AS:    69 Peer AS:    69
        Age: 12:14:00  Metric2: 1
        Task: BGP_69.10.255.70.103+179
        AS path: I
        Communities: target:11111:1 Layer2-info: encaps:VPLS, control
flags:, mtu: 0
```

```

Label-base: 800008, range: 8
Localpref: 100
Router ID: 10.255.70.103
Secondary Tables: green.l2vpn.0
10.255.70.103:2:3:1/96 (1 entry, 0 announced)
  *BGP Preference: 170/-1
    Route Distinguisher: 10.255.70.103:2
    Next-hop reference count: 7
    Source: 10.255.70.103
    Protocol next hop: 10.255.70.103
    Indirect next hop: 2 no-forward
    State: <Active Int Ext>
    Local AS: 69 Peer AS: 69
    Age: 12:14:00 Metric2: 1
    Task: BGP_69.10.255.70.103+179
    AS path: I
    Communities: target:11111:2 Layer2-info: encaps:VPLS,
    control flags:Site-Down,
    mtu: 0
    Label-base: 800016, range: 8
    Localpref: 0
    Router ID: 10.255.70.103
    Secondary Tables: red.l2vpn.0

```

show route source-gateway extensive

```

user@host> show route source-gateway 10.255.70.103 extensive
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
10.255.70.103:1:3:1/96 (1 entry, 1 announced)
  *BGP Preference: 170/-101
    Route Distinguisher: 10.255.70.103:1
    Next-hop reference count: 7
    Source: 10.255.70.103
    Protocol next hop: 10.255.70.103
    Indirect next hop: 2 no-forward
    State: <Secondary Active Int Ext>
    Local AS: 69 Peer AS: 69
    Age: 12:15:24 Metric2: 1
    Task: BGP_69.10.255.70.103+179
    Announcement bits (1): 0-green-l2vpn
    AS path: I
    Communities: target:11111:1 Layer2-info: encaps:VPLS,
    control flags:, mtu: 0

```

```
Label-base: 800008, range: 8
Localpref: 100
Router ID: 10.255.70.103
Primary Routing Table bgp.l2vpn.0

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:2:3:1/96 (1 entry, 1 announced)
*BGP Preference: 170/-1
Route Distinguisher: 10.255.70.103:2
Next-hop reference count: 7
Source: 10.255.70.103
Protocol next hop: 10.255.70.103
Indirect next hop: 2 no-forward
State: <Secondary Active Int Ext>
Local AS: 69 Peer AS: 69
Age: 12:15:24 Metric2: 1
Task: BGP_69.10.255.70.103+179
Announcement bits (1): 0-red-l2vpn
AS path: I
Communities: target:11111:2 Layer2-info: encaps:VPLS,
control flags:Site-Down, mtu: 0
Label-base: 800016, range: 8
Localpref: 0
Router ID: 10.255.70.103
Primary Routing Table bgp.l2vpn.0

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:1:3:1/96 (1 entry, 0 announced)
*BGP Preference: 170/-101
Route Distinguisher: 10.255.70.103:1
Next-hop reference count: 7
Source: 10.255.70.103
Protocol next hop: 10.255.70.103
Indirect next hop: 2 no-forward
State: <Active Int Ext>
Local AS: 69 Peer AS: 69
Age: 12:15:24 Metric2: 1
Task: BGP_69.10.255.70.103+179
AS path: I
Communities: target:11111:1 Layer2-info: encaps:VPLS,
control flags:, mtu: 0
Label-base: 800008, range: 8
Localpref: 100
Router ID: 10.255.70.103
Secondary Tables: green.l2vpn.0
Indirect next hops: 1
    Protocol next hop: 10.255.70.103 Metric: 2
    Indirect next hop: 2 no-forward
    Indirect path forwarding next hops: 1
Next hop: via so-0/3/0.0 weight 0x1
    10.255.70.103/32 Originating RIB: inet.3
    Metric: 2 Node path count: 1
    Forwarding nexthops: 1
    Nexthop: via so-0/3/0.0

10.255.70.103:2:3:1/96 (1 entry, 0 announced)
*BGP Preference: 170/-1
```



```
Route Distinguisher: 10.255.70.103:2
Next-hop reference count: 7
Source: 10.255.70.103
Protocol next hop: 10.255.70.103
Indirect next hop: 2 no-forward
State: <Active Int Ext>
Local AS: 69 Peer AS: 69
Age: 12:15:24 Metric2: 1
Task: BGP_69.10.255.70.103+179
AS path: I
Communities: target:11111:2 Layer2-info: encaps:VPLS,
control flags:Site-Down,
mtu: 0
Label-base: 800016, range: 8
Localpref: 0
Router ID: 10.255.70.103
Secondary Tables: red.12vpn.0
Indirect next hops: 1
    Protocol next hop: 10.255.70.103 Metric: 2
    Indirect next hop: 2 no-forward
    Indirect path forwarding next hops: 1
Next hop: via so-0/3/0.0 weight 0x1
    10.255.70.103/32 Originating RIB: inet.3
    Metric: 2 Node path count: 1
    Forwarding nexthops: 1
    Nexthop: via so-0/3/0.0
```

show route summary

List of Syntax	Syntax on page 3554 Syntax (EX Series Switches) on page 3554
Syntax	<pre>show route summary <logical-system (all <i>logical-system-name</i>)> <table <i>routing-table-name</i>></pre>
Syntax (EX Series Switches)	show route summary
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	<p>Display summary statistics about the entries in the routing table.</p> <p>CPU utilization might increase while the device learns routes. We recommend that you use the show route summary command after the device learns and enters the routes into the routing table. Depending on the size of your network, this might take several minutes. If you receive a “timeout communicating with routing daemon” error when using the show route summary command, wait several minutes before attempting to use the command again. This is not a critical system error, but you might experience a delay in using the command-line interface (CLI).</p>
Options	<p>none—Display summary statistics about the entries in the routing table.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>table <i>routing-table-name</i>—(Optional) Display summary statistics for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the show route summary table inet command). If you only want to display statistics for a specific routing table, make sure to enter the exact name of that routing table.</p>
Required Privilege Level	view
List of Sample Output	show route summary on page 3556 show route summary table on page 3556 show route summary table (with Route Limits Configured for the Routing Table) on page 3557
Output Fields	<p>Table 311 lists the output fields for the show route summary command. Output fields are listed in the approximate order in which they appear.</p>

Table 311: show route summary Output Fields

Field Name	Field Description
Router ID	Address of the local routing device.

Table 311: show route summary Output Fields (*continued*)

Field Name	Field Description
<i>routing-table-name</i>	Name of the routing table (for example, inet.0).
destinations	Number of destinations for which there are routes in the routing table.
routes	<p>Number of routes in the routing table:</p> <ul style="list-style-type: none"> • active—Number of routes that are active. • holddown—Number of routes that are in the hold-down state before being declared inactive. • hidden—Number of routes that are not used because of routing policy.
Restart complete	<p>All protocols have restarted for this routing table.</p> <p>Restart state:</p> <ul style="list-style-type: none"> • Pending:protocol-name—List of protocols that have not yet completed graceful restart for this routing table. • Complete—All protocols have restarted for this routing table. <p>For example, if the output shows-</p> <ul style="list-style-type: none"> • LDP.inet.0: 5 routes (4 active, 1 holddown, 0 hidden) Restart Pending: OSPF LDP VPN <p>This indicates that OSPF, LDP, and VPN protocols did not restart for LDP.inet.0 routing table.</p> <ul style="list-style-type: none"> • vpls_1.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden) Restart Complete <p>This indicates that all protocols have restarted for vpls_1.l2vpn.0 routing table.</p>
Limit/Threshold	<p>Displays the configured route limits for the routing table set with the maximum-prefixes and the maximum-paths statements. If you do not configure route limits for the routing table, the show output does not display this information.</p> <ul style="list-style-type: none"> • destinations—The first number represents the maximum number of route prefixes installed in the routing table. The second number represents the number of route prefixes that trigger a warning message. • routes—The first number represents the maximum number of routes. The second number represents the number of routes that trigger a warning message.
Direct	Routes on the directly connected network.
Local	Local routes.
<i>protocol-name</i>	Name of the protocol from which the route was learned. For example, OSPF , RSVP , and Static .

Sample Output

show route summary

```

user@host> show route summary
Autonomous system number: 69
Router ID: 10.255.71.52
Maximum-ECMP: 32
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
    Direct:    6 routes,    5 active
    Local:    4 routes,    4 active
    OSPF:     5 routes,    4 active
    Static:   7 routes,    7 active
    IGMP:     1 routes,    1 active
    PIM:      2 routes,    2 active

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
    RSVP:      2 routes,    2 active

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete
    Direct:      1 routes,    1 active

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
    MPLS:       3 routes,    3 active
    VPLS:       4 routes,    2 active

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete
    Direct:      2 routes,    2 active
    PIM:         2 routes,    2 active
    MLD:         1 routes,    1 active

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
    BGP:        2 routes,    2 active
    L2VPN:      2 routes,    2 active

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete
    BGP:        2 routes,    2 active
    L2VPN:      1 routes,    1 active

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
    BGP:        4 routes,    4 active

```

show route summary table

```

user@host> show route summary table inet
Router ID: 192.168.0.1

inet.0: 32 destinations, 34 routes (31 active, 0 holddown, 1 hidden)
    Direct:    6 routes,    5 active
    Local:     9 routes,    9 active
    OSPF:      3 routes,    1 active
    Static:   13 routes,   13 active
    IGMP:      1 routes,    1 active

```

```

PIM:      2 routes,      2 active

inet.1: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Multicast: 1 routes,      1 active

inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Local:    1 routes,      1 active
PIM:      2 routes,      2 active

inet6.1: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Multicast: 1 routes,      1 active
```

show route summary table (with Route Limits Configured for the Routing Table)

```

user@host> show route summary table VPN-A.inet.0
Autonomous system number: 100
Router ID: 10.255.182.142

VPN-A.inet.0: 13 destinations, 14 routes (13 active, 0 holddown, 0 hidden)
Limit/Threshold: 2000/200 destinations 20/12 routes
Direct:      2 routes,      2 active
Local:       1 routes,      1 active
OSPF:        4 routes,      3 active
BGP:         4 routes,      4 active
IGMP:        1 routes,      1 active
PIM:         2 routes,      2 active
```

show route table

List of Syntax	Syntax on page 3558 Syntax (EX Series Switches and QFX Series Switches) on page 3558
Syntax	<code>show route table <i>routing-table-name</i></code> <code><brief detail extensive terse></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switches and QFX Series Switches)	<code>show route table <i>routing-table-name</i></code> <code><brief detail extensive terse></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 14.1X53-D15 for QFX Series switches. Show route table evpn statement introduced in Junos OS Release 15.1X53-D30 for QFX Series switches.
Description	Display the route entries in a particular routing table.
Options	brief detail extensive terse —(Optional) Display the specified level of output. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. <i>routing-table-name</i> —Display route entries for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the show route table inet command).
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show route summary on page 3554
List of Sample Output	show route table bgp.l2.vpn on page 3569 show route table bgp.l3vpn.0 on page 3569 show route table bgp.l3vpn.0 detail on page 3569 show route table bgp.rtarget.0 (When Proxy BGP Route Target Filtering Is Configured) on page 3570 show route table bgp.evpn.0 on page 3571 show route table evpna.evpn.0 on page 3571 show route table inet.0 on page 3571 show route table inet.3 on page 3572 show route table inet6.0 on page 3572 show route table inet6.3 on page 3572 show route table inetflow detail on page 3573 show route table l2circuit.0 on page 3573 show route table mpls on page 3573 show route table mpls extensive on page 3574

[show route table mpls.0 on page 3574](#)
[show route table mpls.0 detail \(PTX Series\) on page 3574](#)
[show route table mpls.0 extensive \(PTX Series\) on page 3575](#)
[show route table mpls.0 \(RSVP Route—Transit LSP\) on page 3576](#)
[show route table vpls_1 detail on page 3576](#)
[show route table vpn-a on page 3576](#)
[show route table vpn-a.mdt.0 on page 3577](#)
[show route table VPN-A detail on page 3577](#)
[show route table VPN-AB.inet.0 on page 3578](#)
[show route table VPN_blue.mvpn-inet6.0 on page 3578](#)
[show route table vrf1.mvpn.0 extensive on page 3578](#)
[show route table MVPN.mvpn.0 on page 3579](#)
[show route table inetflow detail on page 3579](#)
[show route table bgp.evpn.0 extensive |no-more \(EVPN\) on page 3582](#)

Output Fields Table 297 describes the output fields for the **show route table** command. Output fields are listed in the approximate order in which they appear.

Table 312: show route table Output Fields

Field Name	Field Description
<i>routing-table-name</i>	Name of the routing table (for example, inet.0).
Restart complete	<p>All protocols have restarted for this routing table.</p> <p>Restart state:</p> <ul style="list-style-type: none"> • Pending:<i>protocol-name</i>—List of protocols that have not yet completed graceful restart for this routing table. • Complete—All protocols have restarted for this routing table. <p>For example, if the output shows-</p> <ul style="list-style-type: none"> • LDP.inet.0 : 5 routes (4 active, 1 holddown, 0 hidden) Restart Pending: OSPF LDP VPN <p>This indicates that OSPF, LDP, and VPN protocols did not restart for LDP.inet.0 routing table.</p> <ul style="list-style-type: none"> • vpls_1.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden) Restart Complete <p>This indicates that all protocols have restarted for vpls_1.l2vpn.0 routing table.</p>
<i>number destinations</i>	Number of destinations for which there are routes in the routing table.
<i>number routes</i>	<p>Number of routes in the routing table and total number of routes in the following states:</p> <ul style="list-style-type: none"> • active (routes that are active) • holddown (routes that are in the pending state before being declared inactive) • hidden (routes that are not used because of a routing policy)

Table 312: show route table Output Fields (*continued*)

Field Name	Field Description
<i>route-destination</i> (entry, announced)	<p>Route destination (for example:10.0.0.1/24). The entry value is the number of routes for this destination, and the announced value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:</p> <ul style="list-style-type: none"> • MPLS-label (for example, 80001). • interface-name (for example, ge-1/0/2). • neighbor-address:control-word-status:encapsulation type:vc-id:source (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96). <ul style="list-style-type: none"> • neighbor-address—Address of the neighbor. • control-word-status—Whether the use of the control word has been negotiated for this virtual circuit: NoCtrlWord or CtrlWord. • encapsulation type—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport. • vc-id—Virtual circuit identifier. • source—Source of the advertisement: Local or Remote. • inclusive multicast Ethernet tag route—Type of route destination represented by (for example, 3:100.100.100.10:100::0::10::100.100.100.10/384): <ul style="list-style-type: none"> • route distinguisher—(8 octets) Route distinguisher (RD) must be the RD of the EVPN instance (EVI) that is advertising the NLRI. • Ethernet tag ID—(4 octets) Identifier of the Ethernet tag. Can set to 0 or to a valid Ethernet tag value. • IP address length—(1 octet) Length of IP address in bits. • originating router's IP address—(4 or 16 octets) Must set to the provider edge (PE) device's IP address. This address should be common for all EVIs on the PE device, and may be the PE device's loopback address.
label stacking	<p>(Next-to-the-last-hop routing device for MPLS only) Depth of the MPLS label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels.</p> <ul style="list-style-type: none"> • S=0 route indicates that a packet with an incoming label stack depth of 2 or more exits this routing device with one fewer label (the label-popping operation is performed). • If there is no S= information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).
[<i>protocol, preference</i>]	<p>Protocol from which the route was learned and the preference value for the route.</p> <ul style="list-style-type: none"> • +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • -—A hyphen indicates the last active route. • *—An asterisk indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route. <p>In every routing metric except for the BGP LocalPref attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the LocalPref value in the Preference2 field. For example, if the LocalPref value for Route 1 is 100, the Preference2 value is -101. If the LocalPref value for Route 2 is 155, the Preference2 value is -156. Route 2 is preferred because it has a higher LocalPref value and a lower Preference2 value.</p>

Table 312: show route table Output Fields (*continued*)

Field Name	Field Description
Level	(IS-IS only). In IS-IS, a single AS can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.
Route Distinguisher	IP subnet augmented with a 64-bit prefix.
PMSI	Provider multicast service interface (MVPN routing table).
Next-hop type	Type of next hop. For a description of possible values for this field, see Table 300 .
Next-hop reference count	Number of references made to the next hop.
Flood nexthop branches exceed maximum message	Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.
Source	IP address of the route source.
Next hop	Network layer address of the directly reachable neighboring system.
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word Selected . This field can also contain the following information: <ul style="list-style-type: none"> • Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.
Label-switched-path <i>lsp-path-name</i>	Name of the LSP used to reach the next hop.
Label operation	MPLS label and operation occurring at this routing device. The operation can be pop (where a label is removed from the top of the stack), push (where another label is added to the label stack), or swap (where a label is replaced by another label).
Interface	(Local only) Local interface name.
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to derive a forwarding next hop.
Indirect next hop	Index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.
State	State of the route (a route can be in more than one state). See Table 301 .

Table 312: show route table Output Fields (*continued*)

Field Name	Field Description
Local AS	AS number of the local routing device.
Age	How long the route has been known.
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.
Metric	Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
MED-plus-IGP	Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.
TTL-Action	For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signaled and LDP-signaled LSPs or for specific VRF routing instances.
Task	Name of the protocol that has added the route.
Announcement bits	<p>The number of BGP peers or protocols to which Junos OS has announced this route, followed by the list of the recipients of the announcement. Junos OS can also announce the route to the KRT for installing the route into the Packet Forwarding Engine, to a resolve tree, a L2 VC, or even a VPN. For example, <i>n-Resolve inet</i> indicates that the specified route is used for route resolution for next hops found in the routing table.</p> <ul style="list-style-type: none"> <i>n</i>—An index used by Juniper Networks customer support only.
AS path	<p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> I—IGP. E—EGP. Recorded—The AS path is recorded by the sample process (sampled). ?—Incomplete; typically, the AS path was aggregated. <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> []—Brackets enclose the number that precedes the AS path. This number represents the number of ASs present in the AS path, when calculated as defined in RFC 4271. This value is used in the AS-path merge process, as defined in RFC 4893. []—If more than one AS number is configured on the routing device, or if AS path prepending is configured, brackets enclose the local AS number associated with the AS path. { }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. ()—Parentheses enclose a confederation. ([])—Parentheses and brackets enclose a confederation set. <p>NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p>

Table 312: show route table Output Fields (*continued*)

Field Name	Field Description
validation-state	<p>(BGP-learned routes) Validation status of the route:</p> <ul style="list-style-type: none"> • Invalid—Indicates that the prefix is found, but either the corresponding AS received from the EBGp peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • Unknown—Indicates that the prefix is not among the prefixes or prefix ranges in the database. • Unverified—Indicates that the origin of the prefix is not verified against the database. This is because the database got populated and the validation is not called for in the BGP import policy, although origin validation is enabled, or the origin validation is not enabled for the BGP peers. • Valid—Indicates that the prefix and autonomous system pair are found in the database.
FECs bound to route	Point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.
Primary Upstream	When multipoint LDP with multicast-only fast reroute (MoFRR) is configured, the primary upstream path. MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path.
RPF Nexthops	When multipoint LDP with MoFRR is configured, the reverse-path forwarding (RPF) next-hop information. Data packets are received from both the primary path and the secondary paths. The redundant packets are discarded at topology merge points due to the RPF checks.
Label	Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.
weight	Value used to distinguish MoFRR primary and backup routes. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.
VC Label	MPLS label assigned to the Layer 2 circuit virtual connection.
MTU	Maximum transmission unit (MTU) of the Layer 2 circuit.
VLAN ID	VLAN identifier of the Layer 2 circuit.
Prefixes bound to route	Forwarding equivalent class (FEC) bound to this route. Applicable only to routes installed by LDP.
Communities	Community path attribute for the route. See Table 302 for all possible values for this field.
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).
control flags	Control flags: none or Site Down .
mtu	Maximum transmission unit (MTU) information.
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.
status vector	Layer 2 VPN and VPLS network layer reachability information (NLRI).

Table 312: show route table Output Fields (*continued*)

Field Name	Field Description
Accepted Multipath	Current active path when BGP multipath is configured.
Accepted LongLivedStale	The LongLivedStale flag indicates that the route was marked LLGR-stale by this router, as part of the operation of LLGR receiver mode. Either this flag or the LongLivedStaleImport flag might be displayed for a route. Neither of these flags is displayed at the same time as the Stale (ordinary GR stale) flag.
Accepted LongLivedStaleImport	<p>The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy. Either this flag or the LongLivedStale flag might be displayed for a route. Neither of these flags is displayed at the same time as the Stale (ordinary GR stale) flag.</p> <p>Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and import into the inet.0 routing table</p>
ImportAccepted LongLivedStaleImport	<p>Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and imported into the inet.0 routing table</p> <p>The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy.</p>
Accepted MultipathContrib	Path currently contributing to BGP multipath.
Localpref	Local preference value included in the route.
Router ID	BGP router ID as advertised by the neighbor in the open message.
Primary Routing Table	In a routing table group, the name of the primary routing table in which the route resides.
Secondary Tables	In a routing table group, the name of one or more secondary tables in which the route resides.

Table 300 describes all possible values for the Next-hop Types output field.

Table 313: Next-hop Types Output Field Values

Next-Hop Type	Description
Broadcast (bcast)	Broadcast next hop.
Deny	Deny next hop.
Discard	Discard next hop.
Flood	Flood next hop. Consists of components called branches, up to a maximum of 32 branches. Each flood next-hop branch sends a copy of the traffic to the forwarding interface. Used by point-to-multipoint RSVP, point-to-multipoint LDP, point-to-multipoint CCC, and multicast.

Table 313: Next-hop Types Output Field Values (*continued*)

Next-Hop Type	Description
Hold	Next hop is waiting to be resolved into a unicast or multicast type.
Indexed (idxd)	Indexed next hop.
Indirect (indr)	Used with applications that have a protocol next hop address that is remote. You are likely to see this next-hop type for internal BGP (IBGP) routes when the BGP next hop is a BGP neighbor that is not directly connected.
Interface	Used for a network address assigned to an interface. Unlike the router next hop, the interface next hop does not reference any specific node on the network.
Local (locl)	Local address on an interface. This next-hop type causes packets with this destination address to be received locally.
Multicast (mcst)	Wire multicast next hop (limited to the LAN).
Multicast discard (mdsc)	Multicast discard.
Multicast group (mgrp)	Multicast group member.
Receive (recv)	Receive.
Reject (rjct)	Discard. An ICMP unreachable message was sent.
Resolve (rslv)	Resolving next hop.
Routed multicast (mcrst)	Regular multicast next hop.
Router	<p>A specific node or set of nodes to which the routing device forwards packets that match the route prefix.</p> <p>To qualify as next-hop type router, the route must meet the following criteria:</p> <ul style="list-style-type: none"> • Must not be a direct or local subnet for the routing device. • Must have a next hop that is directly connected to the routing device.
Table	Routing table next hop.
Unicast (ucst)	Unicast.
Unilist (ulst)	List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.

Table 301 describes all possible values for the State output field. A route can be in more than one state (for example, <Active NoReadvrt Int Ext>).

Table 314: State Output Field Values

Value	Description
Accounting	Route needs accounting.
Active	Route is active.
Always Compare MED	Path with a lower multiple exit discriminator (MED) is available.
AS path	Shorter AS path is available.
Cisco Non-deterministic MED selection	Cisco nondeterministic MED is enabled, and a path with a lower MED is available.
Clone	Route is a clone.
Cluster list length	Length of cluster list sent by the route reflector.
Delete	Route has been deleted.
Ex	Exterior route.
Ext	BGP route received from an external BGP neighbor.
FlashAll	Forces all protocols to be notified of a change to any route, active or inactive, for a prefix. When not set, protocols are informed of a prefix only when the active route changes.
Hidden	Route not used because of routing policy.
IfCheck	Route needs forwarding RPF check.
IGP metric	Path through next hop with lower IGP metric is available.
Inactive reason	Flags for this route, which was not selected as best for a particular destination.
Initial	Route being added.
Int	Interior route.
Int Ext	BGP route received from an internal BGP peer or a BGP confederation peer.
Interior > Exterior > Exterior via Interior	Direct, static, IGP, or EBGp path is available.

Table 314: State Output Field Values (*continued*)

Value	Description
Local Preference	Path with a higher local preference value is available.
Martian	Route is a martian (ignored because it is obviously invalid).
MartianOK	Route exempt from martian filtering.
Next hop address	Path with lower metric next hop is available.
No difference	Path from neighbor with lower IP address is available.
NoReadvrt	Route not to be advertised.
NotBest	Route not chosen because it does not have the lowest MED.
Not Best in its group	Incoming BGP AS is not the best of a group (only one AS can be the best).
NotInstall	Route not to be installed in the forwarding table.
Number of gateways	Path with a greater number of next hops is available.
Origin	Path with a lower origin code is available.
Pending	Route pending because of a hold-down configured on another route.
Release	Route scheduled for release.
RIB preference	Route from a higher-numbered routing table is available.
Route Distinguisher	64-bit prefix added to IP subnets to make them unique.
Route Metric or MED comparison	Route with a lower metric or MED is available.
Route Preference	Route with lower preference value is available.
Router ID	Path through a neighbor with lower ID is available.
Secondary	Route not a primary route.
Unusable path	Path is not usable because of one of the following conditions: <ul style="list-style-type: none"> • The route is damped. • The route is rejected by an import policy. • The route is unresolved.
Update source	Last tiebreaker is the lowest IP address value.

Table 302 describes the possible values for the Communities output field.

Table 315: Communities Output Field Values

Value	Description
<i>area-number</i>	4 bytes, encoding a 32-bit area number. For AS-external routes, the value is 0. A nonzero value identifies the route as internal to the OSPF domain, and as within the identified area. Area numbers are relative to a particular OSPF domain.
bandwidth: local AS number:link-bandwidth-number	Link-bandwidth community value used for unequal-cost load balancing. When BGP has several candidate paths available for multipath purposes, it does not perform unequal-cost load balancing according to the link-bandwidth community unless all candidate paths have this attribute.
domain-id	Unique configurable number that identifies the OSPF domain.
domain-id-vendor	Unique configurable number that further identifies the OSPF domain.
<i>link-bandwidth-number</i>	Link-bandwidth number: from 0 through 4,294,967,295 (bytes per second).
<i>local AS number</i>	Local AS number: from 1 through 65,535.
<i>options</i>	1 byte. Currently this is only used if the route type is 5 or 7. Setting the least significant bit in the field indicates that the route carries a type 2 metric.
origin	(Used with VPNs) Identifies where the route came from.
<i>ospf-route-type</i>	1 byte, encoded as 1 or 2 for intra-area routes (depending on whether the route came from a type 1 or a type 2 LSA); 3 for summary routes; 5 for external routes (area number must be 0); 7 for NSSA routes; or 129 for sham link endpoint addresses.
route-type-vendor	Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute 0x8000. The format is area-number:ospf-route-type:options .
rte-type	Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute 0x0306. The format is area-number:ospf-route-type:options .
target	Defines which VPN the route participates in; target has the format 32-bit IP address:16-bit number . For example, 10.19.0.0:100.
unknown IANA	Incoming IANA codes with a value between 0x1 and 0x7fff. This code of the BGP extended community attribute is accepted, but it is not recognized.
unknown OSPF vendor community	Incoming IANA codes with a value above 0x8000. This code of the BGP extended community attribute is accepted, but it is not recognized.

Sample Output

show route table bgp.l2vpn

```
user@host> show route table bgp.l2vpn
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.24.1:1:4:1/96
    *[BGP/170] 01:08:58, localpref 100, from 192.168.24.1
    AS path: I
    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
```

show route table bgp.l3vpn.0

```
user@host> show route table bgp.l3vpn.0
bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.71.15:100:10.255.71.17/32
    *[BGP/170] 00:03:59, MED 1, localpref 100, from
10.255.71.15
    AS path: I
    > via so-2/1/0.0, Push 100020, Push 100011(top)
10.255.71.15:200:10.255.71.18/32
    *[BGP/170] 00:03:59, MED 1, localpref 100, from
10.255.71.15
    AS path: I
    > via so-2/1/0.0, Push 100021, Push 100011(top)
```

show route table bgp.l3vpn.0 detail

```
user@host> show route table bgp.l3vpn.0 detail
bgp.l3vpn.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)

10.255.245.12:1:4.0.0.0/8 (1 entry, 1 announced)
    *BGP Preference: 170/-101
    Route Distinguisher: 10.255.245.12:1
    Source: 10.255.245.12
    Next hop: 192.168.208.66 via fe-0/0/0.0, selected
    Label operation: Push 182449
    Protocol next hop: 10.255.245.12
    Push 182449
    Indirect next hop: 863a630 297
    State: <Active Int Ext>
    Local AS: 35 Peer AS: 35
    Age: 12:19 Metric2: 1
    Task: BGP_35.10.255.245.12+179
    Announcement bits (1): 0-BGP.0.0.0.0+179
    AS path: 30 10458 14203 2914 3356 I (Atomic) Aggregator: 3356 4.68.0.11

    Communities: 2914:420 target:11111:1 origin:56:78
    VPN Label: 182449
    Localpref: 100
    Router ID: 10.255.245.12

10.255.245.12:1:4.17.225.0/24 (1 entry, 1 announced)
    *BGP Preference: 170/-101
    Route Distinguisher: 10.255.245.12:1
    Source: 10.255.245.12
    Next hop: 192.168.208.66 via fe-0/0/0.0, selected
```

```
Label operation: Push 182465
Protocol next hop: 10.255.245.12
Push 182465
Indirect next hop: 863a8f0 305
State: <Active Int Ext>
Local AS: 35 Peer AS: 35
Age: 12:19 Metric2: 1
Task: BGP_35.10.255.245.12+179
Announcement bits (1): 0-BGP.0.0.0.0+179
AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496 6496 I
Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
VPN Label: 182465
Localpref: 100
Router ID: 10.255.245.12

10.255.245.12:1:4.17.226.0/23 (1 entry, 1 announced)
*BGP Preference: 170/-101
Route Distinguisher: 10.255.245.12:1
Source: 10.255.245.12
Next hop: 192.168.208.66 via fe-0/0/0.0, selected
Label operation: Push 182465
Protocol next hop: 10.255.245.12
Push 182465
Indirect next hop: 86bd210 330
State: <Active Int Ext>
Local AS: 35 Peer AS: 35
Age: 12:19 Metric2: 1
Task: BGP_35.10.255.245.12+179
Announcement bits (1): 0-BGP.0.0.0.0+179
AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496
6496 I
Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
VPN Label: 182465
Localpref: 100
Router ID: 10.255.245.12

10.255.245.12:1:4.17.251.0/24 (1 entry, 1 announced)
*BGP Preference: 170/-101
Route Distinguisher: 10.255.245.12:1
Source: 10.255.245.12
Next hop: 192.168.208.66 via fe-0/0/0.0, selected
Label operation: Push 182465
Protocol next hop: 10.255.245.12
Push 182465
Indirect next hop: 86bd210 330
State: <Active Int Ext>
Local AS: 35 Peer AS: 35
Age: 12:19 Metric2: 1
Task: BGP_35.10.255.245.12+179
Announcement bits (1): 0-BGP.0.0.0.0+179
AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496
6496 I
Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
VPN Label: 182465
Localpref: 100
```

show route table bgp.rtarget.0 (When Proxy BGP Route Target Filtering Is Configured)

```
user@host> show route table bgp.rtarget.0
```

```

bgp.rtarget.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

```

```

100:100:100/96
    * [RTarget/5] 00:03:14
      Type Proxy
      for 10.255.165.103
      for 10.255.166.124
      Local

```

show route table bgp.evpn.0

```

user@host> show route table bgp.evpn.0
bgp.evpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2:100.100.100.2:100::0::00:26:88:5f:67:b0/304
    * [BGP/170] 11:00:05, localpref 100, from 100.100.100.2
      AS path: I, validation-state: unverified
      > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
2:100.100.100.2:100::0::00:51:51:51:51:51/304
    * [BGP/170] 11:00:05, localpref 100, from 100.100.100.2
      AS path: I, validation-state: unverified
      > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
2:100.100.100.3:100::0::00:52:52:52:52:52/304
    * [BGP/170] 10:59:58, localpref 100, from 100.100.100.3
      AS path: I, validation-state: unverified
      > to 100.1.13.3 via ge-2/0/8.0, label-switched-path R0toR2
2:100.100.100.3:100::0::a8:d0:e5:5b:01:c8/304
    * [BGP/170] 10:59:58, localpref 100, from 100.100.100.3
      AS path: I, validation-state: unverified
      > to 100.1.13.3 via ge-2/0/8.0, label-switched-path R0toR2
3:100.100.100.2:100::1000::100.100.100.2/304
    * [BGP/170] 11:00:16, localpref 100, from 100.100.100.2
      AS path: I, validation-state: unverified
      > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
3:100.100.100.2:100::2000::100.100.100.2/304
    * [BGP/170] 11:00:16, localpref 100, from 100.100.100.2
      AS path: I, validation-state: unverified
      > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1

```

show route table evpna.evpn.0

```

user@host> show route table evpna.evpn.0
evpna.evpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

3:100.100.100.10:100::0::10::100.100.100.10/384
    * [EVPN/170] 01:37:09
      Indirect
3:100.100.100.2:100::2000::100.100.100.2/304
    * [EVPN/170] 01:37:12
      Indirect

```

show route table inet.0

```

user@host> show route table inet.0
inet.0: 12 destinations, 12 routes (11 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0
    * [Static/5] 00:51:57
      > to 111.222.5.254 via fxp0.0

```

```

1.0.0.1/32      *[Direct/0] 00:51:58
                 > via at-5/3/0.0
1.0.0.2/32      *[Local/0] 00:51:58
                 Local
12.12.12.21/32  *[Local/0] 00:51:57
                 Reject
13.13.13.13/32  *[Direct/0] 00:51:58
                 > via t3-5/2/1.0
13.13.13.14/32  *[Local/0] 00:51:58
                 Local
13.13.13.21/32  *[Local/0] 00:51:58
                 Local
13.13.13.22/32  *[Direct/0] 00:33:59
                 > via t3-5/2/0.0
127.0.0.1/32    [Direct/0] 00:51:58
                 > via lo0.0
111.222.5.0/24  *[Direct/0] 00:51:58
                 > via fxp0.0
111.222.5.81/32 *[Local/0] 00:51:58
                 Local

```

show route table inet.3

```

user@host> show route table inet.3
inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

22.0.0.5/32      *[LDP/9] 00:25:43, metric 10, tag 200
                  to 1.2.94.2 via lt-1/2/0.49
                  > to 1.2.3.2 via lt-1/2/0.23

```

show route table inet6.0

```

user@host> show route table inet6.0
inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Route, * = Both

fec0:0:0:3::/64 *[Direct/0] 00:01:34
>via fe-0/1/0.0

fec0:0:0:3::/128 *[Local/0] 00:01:34
>Local

fec0:0:0:4::/64 *[Static/5] 00:01:34
>to fec0:0:0:3::ffff via fe-0/1/0.0

```

show route table inet6.3

```

user@router> show route table inet6.3
inet6.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

::10.255.245.195/128
                  *[LDP/9] 00:00:22, metric 1
                  > via so-1/0/0.0
::10.255.245.196/128
                  *[LDP/9] 00:00:08, metric 1
                  > via so-1/0/0.0, Push 100008

```

show route table inetflow detail

```

user@host> show route table inetflow detail
inetflow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.12.44.1,*/48 (1 entry, 1 announced)
    *BGP    Preference: 170/-101
            Next-hop reference count: 2
            State: <Active Ext>
            Local AS: 65002 Peer AS: 65000
            Age: 4
            Task: BGP_65000.10.12.99.5+3792
            Announcement bits (1): 0-Flow
            AS path: 65000 I
            Communities: traffic-rate:0:0
            Validation state: Accept, Originator: 10.12.99.5
            Via: 10.12.44.0/24, Active
            Localpref: 100
            Router ID: 10.255.71.161

10.12.56.1,*/48 (1 entry, 1 announced)
    *Flow    Preference: 5
            Next-hop reference count: 2
            State: <Active>
            Local AS: 65002
            Age: 6:30
            Task: RT Flow
            Announcement bits (2): 0-Flow 1-BGP.0.0.0.0+179
            AS path: I
            Communities: 1:1

```

show route table l2circuit.0

```

user@host> show route table l2circuit.0
l2circuit.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.1.1.195:NoCtrlWord:1:1:Local/96
    *[L2CKT/7] 00:50:47
    > via so-0/1/2.0, Push 100049
    via so-0/1/3.0, Push 100049
10.1.1.195:NoCtrlWord:1:1:Remote/96
    *[LDP/9] 00:50:14
    Discard
10.1.1.195:CtrlWord:1:2:Local/96
    *[L2CKT/7] 00:50:47
    > via so-0/1/2.0, Push 100049
    via so-0/1/3.0, Push 100049
10.1.1.195:CtrlWord:1:2:Remote/96
    *[LDP/9] 00:50:14
    Discard

```

show route table mpls

```

user@host> show route table mpls
mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0          *[MPLS/0] 00:13:55, metric 1
           Receive
1          *[MPLS/0] 00:13:55, metric 1
           Receive

```

```

2          *[MPLS/0] 00:13:55, metric 1
           Receive
1024       *[VPN/0] 00:04:18
           to table red.inet.0, Pop

```

show route table mpls extensive

```

user@host> show route table mpls extensive
100000 (1 entry, 1 announced)
TSI:
KRT in-kernel 100000 /36 -> {so-1/0/0.0}
    *LDP    Preference: 9
           Next hop: via so-1/0/0.0, selected
           Pop
           State: <Active Int>
           Age: 29:50      Metric: 1
           Task: LDP
           Announcement bits (1): 0-KRT
           AS path: I
           Prefixes bound to route: 10.0.0.194/32

```

show route table mpls.0

```

user@host> show route table mpls.0
mpls.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0          *[MPLS/0] 00:45:09, metric 1
           Receive
1          *[MPLS/0] 00:45:09, metric 1
           Receive
2          *[MPLS/0] 00:45:09, metric 1
           Receive
100000     *[L2VPN/7] 00:43:04
           > via so-0/1/0.1, Pop
100001     *[L2VPN/7] 00:43:03
           > via so-0/1/0.2, Pop      Offset: 4
100002     *[LDP/9] 00:43:22, metric 1
           via so-0/1/2.0, Pop
           > via so-0/1/3.0, Pop
100002(S=0) *[LDP/9] 00:43:22, metric 1
           via so-0/1/2.0, Pop
           > via so-0/1/3.0, Pop
100003     *[LDP/9] 00:43:22, metric 1
           > via so-0/1/2.0, Swap 100002
           via so-0/1/3.0, Swap 100002
100004     *[LDP/9] 00:43:16, metric 1
           via so-0/1/2.0, Swap 100049
           > via so-0/1/3.0, Swap 100049
so-0/1/0.1 *[L2VPN/7] 00:43:04
           > via so-0/1/2.0, Push 100001, Push 100049(top)
           via so-0/1/3.0, Push 100001, Push 100049(top)
so-0/1/0.2 *[L2VPN/7] 00:43:03
           via so-0/1/2.0, Push 100000, Push 100049(top) Offset: -4
           > via so-0/1/3.0, Push 100000, Push 100049(top) Offset: -4

```

show route table mpls.0 detail (PTX Series)

```

user@host> show route table mpls.0 detail
ge-0/0/2.600 (1 entry, 1 announced)
    *L2VPN Preference: 7
           Next hop type: Indirect

```

```

Address: 0x9438f34
Next-hop reference count: 2
Next hop type: Router, Next hop index: 567
Next hop: 3.0.0.1 via ge-0/0/1.0, selected
Label operation: Push 299808
Label TTL action: prop-ttl
Load balance label: Label 299808:None;
Session Id: 0x1
Protocol next hop: 10.255.255.1
Label operation: Push 299872 Offset: 252
Label TTL action: no-prop-ttl
Load balance label: Label 299872:Flow label PUSH;
Composite next hop: 0x9438ed8 570 INH Session ID: 0x2
Indirect next hop: 0x9448208 262142 INH Session ID: 0x2
State: <Active Int>
Age: 21          Metric2: 1
Validation State: unverified
Task: Common L2 VC
Announcement bits (2): 0-KRT 2-Common L2 VC
AS path: I

```

show route table mpls.0 extensive (PTX Series)

```

user@host> show route table mpls.0 extensive
ge-0/0/2.600 (1 entry, 1 announced)
TSI:
KRT in-kernel ge-0/0/2.600.0      /32 -> {composite(570)}
    *L2VPN Preference: 7
      Next hop type: Indirect
      Address: 0x9438f34
      Next-hop reference count: 2
      Next hop type: Router, Next hop index: 567
      Next hop: 3.0.0.1 via ge-0/0/1.0, selected
      Label operation: Push 299808
      Label TTL action: prop-ttl
      Load balance label: Label 299808:None;
      Session Id: 0x1
      Protocol next hop: 10.255.255.1
      Label operation: Push 299872 Offset: 252
      Label TTL action: no-prop-ttl
      Load balance label: Label 299872:Flow label PUSH;
      Composite next hop: 0x9438ed8 570 INH Session ID: 0x2
      Indirect next hop: 0x9448208 262142 INH Session ID: 0x2
      State: <Active Int>
      Age: 47          Metric2: 1
      Validation State: unverified
      Task: Common L2 VC
      Announcement bits (2): 0-KRT 2-Common L2 VC
      AS path: I
      Composite next hops: 1
        Protocol next hop: 10.255.255.1 Metric: 1
        Label operation: Push 299872 Offset: 252
        Label TTL action: no-prop-ttl
        Load balance label: Label 299872:Flow label PUSH;
        Composite next hop: 0x9438ed8 570 INH Session ID: 0x2
        Indirect next hop: 0x9448208 262142 INH Session ID: 0x2
        Indirect path forwarding next hops: 1
          Next hop type: Router
          Next hop: 3.0.0.1 via ge-0/0/1.0
          Session Id: 0x1
          10.255.255.1/32 Originating RIB: inet.3

```

```

Metric: 1
Forwarding nexthops: 1
Node path count: 1
Nexthop: 3.0.0.1 via ge-0/0/1.0

```

show route table mpls.0 (RSVP Route—Transit LSP)

In the sample output, the 1 in [RSVP/7/1] indicates the secondary preference value. The secondary preference value becomes significant when multiple RSVP LSPs of different types are signaled to the destination. The possible values of RSVP secondary preferences are:

1—Normal Point-to-Point RSVP-TE LSP

2—Point-to-Multipoint (P2MP) RSVP-TE LSP

3—Dynamic RSVP-TE LSP

```
user@host> show route table mpls.0
```

```

mpls.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

```

```

0          *[MPLS/0] 00:37:31, metric 1
            Receive
1          *[MPLS/0] 00:37:31, metric 1
            Receive
2          *[MPLS/0] 00:37:31, metric 1
            Receive
13         *[MPLS/0] 00:37:31, metric 1
            Receive
300352     *[RSVP/7/1] 00:08:00, metric 1
            > to 8.64.0.106 via ge-1/0/1.0, label-switched-path lsp1_p2p
300352(S=0) *[RSVP/7/1] 00:08:00, metric 1
            > to 8.64.0.106 via ge-1/0/1.0, label-switched-path lsp1_p2p
300384     *[RSVP/7/2] 00:05:20, metric 1
            > to 8.64.1.106 via ge-1/0/0.0, Pop
300384(S=0) *[RSVP/7/2] 00:05:20, metric 1
            > to 8.64.1.106 via ge-1/0/0.0, Pop

```

show route table vpls_1 detail

```
user@host> show route table vpls_1 detail
```

```

vpls_1.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

```

```

1.1.1.11:1000:1:1/96 (1 entry, 1 announced)
*L2VPN Preference: 170/-1
Receive table: vpls_1.l2vpn.0
Next-hop reference count: 2
State: <Active Int Ext>
Age: 4:29:47 Metric2: 1
Task: vpls_1-l2vpn
Announcement bits (1): 1-BGP.0.0.0.0+179
AS path: I
Communities: Layer2-info: encaps:VPLS, control flags:Site-Down
Label-base: 800000, range: 8, status-vector: 0xFF

```

show route table vpn-a

```
user@host> show route table vpn-a
```



```

vpn-a.12vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
192.168.16.1:1:1:1/96
    *[VPN/7] 05:48:27
    Discard
192.168.24.1:1:2:1/96
    *[BGP/170] 00:02:53, localpref 100, from 192.168.24.1
    AS path: I
    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
192.168.24.1:1:3:1/96
    *[BGP/170] 00:02:53, localpref 100, from 192.168.24.1
    AS path: I
    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am

```

show route table vpn-a.mdt.0

```

user@host> show route table vpn-a.mdt.0
vpn-a.mdt.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:1:0:10.255.14.216:232.1.1.1/144
    *[MVPN/70] 01:23:05, metric2 1
    Indirect
1:1:1:10.255.14.218:232.1.1.1/144
    *[BGP/170] 00:57:49, localpref 100, from 10.255.14.218
    AS path: I
    > via so-0/0/0.0, label-switched-path r0e-to-r1
1:1:2:10.255.14.217:232.1.1.1/144
    *[BGP/170] 00:57:49, localpref 100, from 10.255.14.217
    AS path: I
    > via so-0/0/1.0, label-switched-path r0-to-r2

```

show route table VPN-A detail

```

user@host> show route table VPN-A detail
VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
10.255.179.9/32 (1 entry, 1 announced)
    *BGP Preference: 170/-101
    Route Distinguisher: 10.255.179.13:200
    Next hop type: Indirect
    Next-hop reference count: 5
    Source: 10.255.179.13
    Next hop type: Router, Next hop index: 732
    Next hop: 10.39.1.14 via fe-0/3/0.0, selected
    Label operation: Push 299824, Push 299824(top)
    Protocol next hop: 10.255.179.13
    Push 299824
    Indirect next hop: 8f275a0 1048574
    State: (Secondary Active Int Ext)
    Local AS: 1 Peer AS: 1
    Age: 3:41:06 Metric: 1 Metric2: 1
    Task: BGP_1.10.255.179.13+64309
    Announcement bits (2): 0-KRT 1-BGP RT Background
    AS path: I
    Communities: target:1:200 rte-type:0.0.0.0:1:0
    Import Accepted
    VPN Label: 299824 TTL Action: vrf-ttl-propagate
    Localpref: 100
    Router ID: 10.255.179.13
    Primary Routing Table bgp.13vpn.0

```

show route table VPN-AB.inet.0

```

user@host> show route table VPN-AB.inet.0
VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.0/30      *[OSPF/10] 00:07:24, metric 1
                  > via so-7/3/1.0
10.39.1.4/30      *[Direct/0] 00:08:42
                  > via so-5/1/0.0
10.39.1.6/32      *[Local/0] 00:08:46
                  Local
10.255.71.16/32   *[Static/5] 00:07:24
                  > via so-2/0/0.0
10.255.71.17/32   *[BGP/170] 00:07:24, MED 1, localpref 100, from
10.255.71.15
                  AS path: I
                  > via so-2/1/0.0, Push 100020, Push 100011(top)
10.255.71.18/32   *[BGP/170] 00:07:24, MED 1, localpref 100, from
10.255.71.15
                  AS path: I
                  > via so-2/1/0.0, Push 100021, Push 100011(top)
10.255.245.245/32 *[BGP/170] 00:08:35, localpref 100
                  AS path: 2 I
                  > to 10.39.1.5 via so-5/1/0.0
10.255.245.246/32 *[OSPF/10] 00:07:24, metric 1
                  > via so-7/3/1.0

```

show route table VPN_blue.mvpn-inet6.0

```

user@host> show route table VPN_blue.mvpn-inet6.0
vpn_blue.mvpn-inet6.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:10.255.2.202:65535:10.255.2.202/432
                  *[BGP/170] 00:02:37, localpref 100, from 10.255.2.202
                  AS path: I
                  > via so-0/1/3.0
1:10.255.2.203:65535:10.255.2.203/432
                  *[BGP/170] 00:02:37, localpref 100, from 10.255.2.203
                  AS path: I
                  > via so-0/1/0.0
1:10.255.2.204:65535:10.255.2.204/432
                  *[MVPN/70] 00:57:23, metric2 1
                  Indirect
5:10.255.2.202:65535:128::192.168.90.2:128:ffff::1/432
                  *[BGP/170] 00:02:37, localpref 100, from 10.255.2.202
                  AS path: I
                  > via so-0/1/3.0
6:10.255.2.203:65535:65000:128::10.12.53.12:128:ffff::1/432
                  *[PIM/105] 00:02:37
                  Multicast (IPv6)
7:10.255.2.202:65535:65000:128::192.168.90.2:128:ffff::1/432
                  *[MVPN/70] 00:02:37, metric2 1
                  Indirect

```

show route table vrf1.mvpn.0 extensive

```

user@host> show route table vrf1.mvpn.0 extensive
1:10.255.50.77:1:10.255.50.77/240 (1 entry, 1 announced)
    *MVPN    Preference: 70

```

```

PMSI: Flags 0x0: Label 0: RSVP-TE:
Session_13[10.255.50.77:0:25624:10.255.50.77]
  Next hop type: Indirect
  Address: 0xbb2c944
  Next-hop reference count: 360
  Protocol next hop: 10.255.50.77
  Indirect next hop: 0x0 - INH Session ID: 0x0
  State: <Active Int Ext>
  Age: 53:03      Metric2: 1
  Validation State: unverified
  Task: mvpn global task
  Announcement bits (3): 0-PIM.vrf1 1-mvpn global task 2-rt-export

AS path: I

```

show route table MVPN.mvpn.0

Starting in Junos OS Release 15.1, multicast routes on the locally originated type 7 customer multicast routes are added exclusively by PIM. The functionality of the BGP-MVPN service (which, internally, depends on contributions of state from both the MVPN and PIM protocol components of Junos OS) remains unchanged. MVPN, however, no longer appears as the originator of the locally advertised route. Routes advertised by remote PEs are, as usual, always learned locally from their respective [BGP/...] protocol.

```

user@host> show route table MVPN.mvpn.0
MVPN.mvpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

7:10.255.2.202:65535:65000:128:::192.168.90.2:128:ffff::1/432
    *[PIM/70] 00:02:37, metric2 1
    Indirect
5:100:32:192.168.1.9:32:239.1.1.1/240
    *[PIM/105] 01:51:21
    Multicast (IPv4)
7:100:1:100.32.192.168.5:32:237.1.1.1/240
    *[PIM/105] 01:51:21
    Multicast (IPv4)

```

show route table inetflow detail

```

user@host> show route table inetflow detail
inetflow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.12.44.1,*/48 (1 entry, 1 announced)
    *BGP      Preference: 170/-101
    Next-hop reference count: 2
    State: <Active Ext>
    Local AS: 65002 Peer AS: 65000
    Age: 4
    Task: BGP_65000.10.12.99.5+3792
    Announcement bits (1): 0-Flow
    AS path: 65000 I
    Communities: traffic-rate:0:0
    Validation state: Accept, Originator: 10.12.99.5
    Via: 10.12.44.0/24, Active
    Localpref: 100
    Router ID: 10.255.71.161

10.12.56.1,*/48 (1 entry, 1 announced)
    *Flow      Preference: 5
    Next-hop reference count: 2

```

```

State: <Active>
Local AS: 65002
Age: 6:30
Task: RT Flow
Announcement bits (2): 0-Flow 1-BGP.0.0.0.0+179
AS path: I
Communities: 1:1

user@PE1> show route table green.l2vpn.0 (VPLS Multihoming with FEC 129)
green.l2vpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.2:100:1.1.1.2/96 AD
    *[VPLS/170] 1d 03:11:03, metric2 1
    Indirect
1.1.1.4:100:1.1.1.4/96 AD
    *[BGP/170] 1d 03:11:02, localpref 100, from 1.1.1.4
    AS path: I, validation-state: unverified
    > via ge-1/2/1.5
1.1.1.2:100:1:0/96 MH
    *[VPLS/170] 1d 03:11:03, metric2 1
    Indirect
1.1.1.4:100:1:0/96 MH
    *[BGP/170] 1d 03:11:02, localpref 100, from 1.1.1.4
    AS path: I, validation-state: unverified
    > via ge-1/2/1.5
1.1.1.4:NoCtrlWord:5:100:100:1.1.1.2:1.1.1.4/176
    *[VPLS/7] 1d 03:11:02, metric2 1
    > via ge-1/2/1.5
1.1.1.4:NoCtrlWord:5:100:100:1.1.1.4:1.1.1.2/176
    *[LDP/9] 1d 03:11:02
    Discard

user@host> show route table red extensive
red.inet.0: 364481 destinations, 714087 routes (364480 active, 48448 holddown, 1
hidden)
22.0.0.0/32 (3 entries, 1 announced)
    State: <OnList CalcForwarding>
TSI:
KRT in-kernel 22.0.0.0/32 -> {composite(1048575)} Page 0 idx 1 Type 1 val 0x934342c

    Nexthop: Self
    AS path: [2] I
    Communities: target:2:1
Path 22.0.0.0 from 2.3.0.0 Vector len 4. Val: 1
    @BGP Preference: 170/-1
    Route Distinguisher: 2:1
    Next hop type: Indirect
    Address: 0x258059e4
    Next-hop reference count: 2
    Source: 2.2.0.0
    Next hop type: Router
    Next hop: 10.1.1.1 via ge-1/1/9.0, selected
    Label operation: Push 707633
    Label TTL action: prop-ttl
    Session Id: 0x17d8
    Protocol next hop: 2.2.0.0
    Push 16
    Composite next hop: 0x25805988 - INH Session ID: 0x193c
    Indirect next hop: 0x23eea900 - INH Session ID: 0x193c
    State: <Secondary Active Int Ext ProtectionPath ProtectionCand>

```

```

Local AS:      2 Peer AS:      2
Age: 23        Metric2: 35
Validation State: unverified
Task: BGP_2.2.2.0.0+34549
AS path: I
Communities: target:2:1
Import Accepted
VPN Label: 16
Localpref: 0
Router ID: 2.2.0.0
Primary Routing Table bgp.13vpn.0
Composite next hops: 1
    Protocol next hop: 2.2.0.0 Metric: 35
    Push 16
    Composite next hop: 0x25805988 - INH Session ID: 0x193c
    Indirect next hop: 0x23eea900 - INH Session ID: 0x193c
    Indirect path forwarding next hops: 1
        Next hop type: Router
        Next hop: 10.1.1.1 via ge-1/1/9.0
        Session Id: 0x17d8
    2.2.0.0/32 Originating RIB: inet.3
    Metric: 35                               Node path count: 1
    Forwarding nexthops: 1
        Nexthop: 10.1.1.1 via ge-1/1/9.0
BGP Preference: 170/-1
Route Distinguisher: 2:1
Next hop type: Indirect
Address: 0x9347028
Next-hop reference count: 3
Source: 2.3.0.0
Next hop type: Router, Next hop index: 702
Next hop: 10.1.4.2 via ge-1/0/0.0, selected
Label operation: Push 634278
Label TTL action: prop-ttl
Session Id: 0x17d9
Protocol next hop: 2.3.0.0
Push 16
Composite next hop: 0x93463a0 1048575 INH Session ID: 0x17da
Indirect next hop: 0x91e8800 1048574 INH Session ID: 0x17da
State: <Secondary NotBest Int Ext ProtectionPath ProtectionCand>

Inactive reason: Not Best in its group - IGP metric
Local AS:      2 Peer AS:      2
Age: 3:34      Metric2: 70
Validation State: unverified
Task: BGP_2.2.3.0.0+32805
Announcement bits (2): 0-KRT 1-BGP_RT_Background
AS path: I
Communities: target:2:1
Import Accepted
VPN Label: 16
Localpref: 0
Router ID: 2.3.0.0
Primary Routing Table bgp.13vpn.0
Composite next hops: 1
    Protocol next hop: 2.3.0.0 Metric: 70
    Push 16
    Composite next hop: 0x93463a0 1048575 INH Session ID:
0x17da
    Indirect next hop: 0x91e8800 1048574 INH Session ID:
0x17da

```

```

        Indirect path forwarding next hops: 1
            Next hop type: Router
            Next hop: 10.1.4.2 via ge-1/0/0.0
            Session Id: 0x17d9
        2.3.0.0/32 Originating RIB: inet.3
            Metric: 70
            Node path count: 1
            Forwarding nexthops: 1
            Nexthop: 10.1.4.2 via ge-1/0/0.0
#Multipath Preference: 255
    Next hop type: Indirect
    Address: 0x24afca30
    Next-hop reference count: 1
    Next hop type: Router
    Next hop: 10.1.1.1 via ge-1/1/9.0, selected
    Label operation: Push 707633
    Label TTL action: prop-ttl
    Session Id: 0x17d8
    Next hop type: Router, Next hop index: 702
    Next hop: 10.1.4.2 via ge-1/0/0.0
    Label operation: Push 634278
    Label TTL action: prop-ttl
    Session Id: 0x17d9
    Protocol next hop: 2.2.0.0
    Push 16
    Composite next hop: 0x25805988 - INH Session ID: 0x193c
    Indirect next hop: 0x23eea900 - INH Session ID: 0x193c Weight 0x1

    Protocol next hop: 2.3.0.0
    Push 16
    Composite next hop: 0x93463a0 1048575 INH Session ID: 0x17da
    Indirect next hop: 0x91e8800 1048574 INH Session ID: 0x17da Weight

0x4000
    State: <ForwardingOnly Int Ext>
    Inactive reason: Forwarding use only
    Age: 23
    Metric2: 35
    Validation State: unverified
    Task: RT
    AS path: I
    Communities: target:2:1

```

show route table bgp.evpn.0 extensive |no-more (EVPN)

```

show route table bgp.evpn.0 extensive | no-more
bgp.evpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
2:1000:10::100::00:aa:aa:aa:aa:aa/304 (1 entry, 0 announced)
    *BGP
        Preference: 170/-101
        Route Distinguisher: 1000:10
        Next hop type: Indirect
        Address: 0x9420fd0
        Next-hop reference count: 12
        Source: 1.2.3.4
        Protocol next hop: 1.2.3.4
        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
    State: Local AS: 17 Peer AS:17 Age:21:12 Metric2:1 Validation State:
unverified
        Task: BGP_17.1.2.3.4+50756
        AS path: I
        Communities: target:1111:8388708 encapsulation0:0:0:0:3
        Import Accepted
        Route Label: 100
        ESI: 00:00:00:00:00:00:00:00:00:00

```

```

Localpref: 100
Router ID: 1.2.3.4
Secondary Tables: default-switch.evpn.0
Indirect next hops: 1
    Protocol next hop: 1.2.3.4 Metric: 1
    Indirect next hop: 0x2 no-forward INH Session ID: 0x0
    Indirect path forwarding next hops: 1
        Next hop type: Router
        Next hop: 10.10.10.1 via xe-0/0/1.0
        Session Id: 0x2
    1.2.3.4/32 Originating RIB: inet.0
        Metric: 1                      Node path count: 1
        Forwarding nexthops: 2
        Nexthop: 10.92.78.102 via em0.0

2:1000:10::200::00:bb:bb:bb:bb:bb/304 (1 entry, 0 announced)
  *BGP   Preference: 170/-101
        Route Distinguisher: 1000:10
        Next hop type: Indirect
        Address: 0x9420fd0
        Next-hop reference count: 12
        Source: 1.2.3.4
        Protocol next hop: 1.2.3.4
        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
        State: Local AS:17 Peer AS:17 Age:19:43 Metric2:1 Validation
State:unverified
        Task: BGP_17.1.2.3.4+50756
        AS path: I
        Communities: target:2222:22 encapsulation0:0:0:0:3
        Import Accepted
        Route Label: 200
        ESI: 00:00:00:00:00:00:00:00:00:00
        Localpref: 100
        Router ID: 1.2.3.4
        Secondary Tables: default-switch.evpn.0
        Indirect next hops: 1
            Protocol next hop: 1.2.3.4 Metric: 1
            Indirect next hop: 0x2 no-forward INH Session ID: 0x0
            Indirect path forwarding next hops: 1
                Next hop type: Router
                Next hop: 10.10.10.1 via xe-0/0/1.0
                Session Id: 0x2
            1.2.3.4/32 Originating RIB: inet.0
                Metric: 1                      Node path count: 1
                Forwarding nexthops: 2
                Nexthop: 10.92.78.102 via em0.0

2:1000:10::300::00:cc:cc:cc:cc:cc/304 (1 entry, 0 announced)
  *BGP   Preference: 170/-101
        Route Distinguisher: 1000:10
        Next hop type: Indirect
        Address: 0x9420fd0
        Next-hop reference count: 12
        Source: 1.2.3.4
        Protocol next hop: 1.2.3.4
        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
        State: Local AS:17 Peer AS:17 Age:17:21 Metric2:1 Validation State:
unverified Task: BGP 17,1,2,3,4+50756
        AS path: I
        Communities: target:3333:33 encapsulation0:0:0:0:3
        Import Accepted

```

```

Route Label: 300
ESI: 00:00:00:00:00:00:00:00:00
Localpref: 100
Router ID: 1.2.3.4
Secondary Tables: default-switch.evpn.0
Indirect next hops: 1
    Protocol next hop: 1.2.3.4 Metric: 1
    Indirect next hop: 0x2 no-forward INH Session ID: 0x0
    Indirect path forwarding next hops: 1
        Next hop type: Router
        Next hop: 10.10.10.1 via xe-0/0/1.0
        Session Id: 0x2
    1.2.3.4/32 Originating RIB: inet.0
        Metric: 1                      Node path count: 1
        Forwarding nexthops: 2
        Nexthop: 10.92.78.102 via em0.0

3:1000:10::100::1.2.3.4/304 (1 entry, 0 announced)
*BGP   Preference: 170/-101
Route Distinguisher: 1000:10
PMSI: Flags 0x0: Label 100: Type INGRESS-REPLICATION 1.2.3.4
Next hop type: Indirect
Address: 0x9420fd0
Next-hop reference count: 12
Source: 1.2.3.4
Protocol next hop: 1.2.3.4
Indirect next hop: 0x2 no-forward INH Session ID: 0x0
State: Local AS:17 Peer AS:17 Age:37:01 Metric2:1 Validation State:
unverified Task: BGP 17.1.2.3.4+50756
AS path: I
Communities: target:1111:8388708 encapsulation0:0:0:0:3
Import Accepted
Localpref: 100
Router ID: 1.2.3.4
Secondary Tables: default-switch.evpn.0
Indirect next hops: 1
    Protocol next hop: 1.2.3.4 Metric: 1
    Indirect next hop: 0x2 no-forward INH Session ID: 0x0
    Indirect path forwarding next hops: 1
        Next hop type: Router
        Next hop: 10.10.10.1 via xe-0/0/1.0
        Session Id: 0x2
    1.2.3.4/32 Originating RIB: inet.0
        Metric: 1                      Node path count: 1
        Forwarding nexthops: 2
        Nexthop: 10.92.78.102 via em0.0

3:1000:10::200::1.2.3.4/304 (1 entry, 0 announced)
*BGP   Preference: 170/-101
Route Distinguisher: 1000:10
PMSI: Flags 0x0: Label 200: Type INGRESS-REPLICATION 1.2.3.4
Next hop type: Indirect
Address: 0x9420fd0
Next-hop reference count: 12
Source: 1.2.3.4
Protocol next hop: 1.2.3.4
Indirect next hop: 0x2 no-forward INH Session ID: 0x0
State: Local AS: 17 Peer AS: 17 Age:35:22 Metric2:1 Validation
State:unverified Task: BGP 17.1.2.3.4+50756
AS path:I Communities: target:2222:22 encapsulation):0:0:0:0:3

```



```

Import Accepted
  Localpref: 100
  Router ID: 1.2.3.4
  Secondary Tables: default-switch.evpn.0
  Indirect next hops: 1
    Protocol next hop: 1.2.3.4 Metric: 1
    Indirect next hop: 0x2 no-forward INH Session ID: 0x0
    Indirect path forwarding next hops: 1
      Next hop type: Router
      Next hop: 10.10.10.1 via xe-0/0/1.0
      Session Id: 0x2
    1.2.3.4/32 Originating RIB: inet.0
      Metric: 1
      Forwarding nexthops: 2
      Nexthop: 10.92.78.102 via em0.0
      Node path count: 1

3:1000:10::300::1.2.3.4/304 (1 entry, 0 announced)
  *BGP Preference: 170/-101
    Route Distinguisher: 1000:10
    PMSI: Flags 0x0: Label 300: Type INGRESS-REPLICATION 1.2.3.4
    Next hop type: Indirect
    Address: 0x9420fd0
    Next-hop reference count: 12
    Source: 1.2.3.4
    Protocol next hop: 1.2.3.4
    Indirect next hop: 0x2 no-forward INH Session ID: 0x0
    State: Local AS: 17 Peer AS: 17 Age 35:22 Metric2:1 Validation State:
unverified Task: BGP 17.1.2.3.4+5075
  6 AS path: I Communities: target:3333:33 encapsulation0:0:0:0:3
Import Accepted Localpref:100
  Router ID: 1.2.3.4
  Secondary Tables: default-switch.evpn.0
  Indirect next hops: 1
    Protocol next hop: 1.2.3.4 Metric: 1
    Indirect next hop: 0x2 no-forward INH Session ID: 0x0
    Indirect path forwarding next hops: 1
      Next hop type: Router
      Next hop: 10.10.10.1 via xe-0/0/1.0
      Session Id: 0x2
    1.2.3.4/32 Originating RIB: inet.0
      Metric: 1
      Forwarding nexthops: 2
      Nexthop: 10.92.78.102 via em0.0
      Node path count: 1

```

show route terse

List of Syntax [Syntax on page 3586](#)
[Syntax \(EX Series Switches\) on page 3586](#)

Syntax show route terse
 <logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches) show route terse

Release Information Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.0 for EX Series switches.

Description Display a high-level summary of the routes in the routing table.



NOTE: For BGP routes, the **show route terse** command displays the local preference attribute and MED instead of metric1 and metric2 values. This is mostly due to historical reasons.

To display the metric1 and metric2 value of a BGP route, use the [show route extensive](#) command.

Options none—Display a high-level summary of the routes in the routing table.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level view

List of Sample Output [show route terse on page 3588](#)

Output Fields [Table 316](#) describes the output fields for the **show route terse** command. Output fields are listed in the approximate order in which they appear.

Table 316: show route terse Output Fields

Field Name	Field Description
<i>routing-table-name</i>	Name of the routing table (for example, <i>inet.0</i>).
<i>number destinations</i>	Number of destinations for which there are routes in the routing table.
<i>number routes</i>	Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> active (routes that are active) holddown (routes that are in the pending state before being declared inactive) hidden (routes that are not used because of a routing policy)

Table 316: show route terse Output Fields (*continued*)

Field Name	Field Description
route key	Key for the state of the route: <ul style="list-style-type: none"> • +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • -—A hyphen indicates the last active route. • *—An asterisk indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route.
A	Active route. An asterisk (*) indicates this is the active route.
Destination	Destination of the route.
P	Protocol through which the route was learned: <ul style="list-style-type: none"> • A—Aggregate • B—BGP • C—CCC • D—Direct • G—GMPLS • I—IS-IS • L—L2CKT, L2VPN, LDP, Local • K—Kernel • M—MPLS, MSDP • O—OSPF • P—PIM • R—RIP, RIPng • S—Static • T—Tunnel
Prf	Preference value of the route. In every routing metric except for the BGP LocalPref attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the LocalPref value in the Preference2 field. For example, if the LocalPref value for Route 1 is 100, the Preference2 value is -101. If the LocalPref value for Route 2 is 155, the Preference2 value is -156. Route 2 is preferred because it has a higher LocalPref value and a lower Preference2 value.
Metric 1	First metric value in the route. For routes learned from BGP, this is the MED metric.
Metric 2	Second metric value in the route. For routes learned from BGP, this is the IGP metric.
Next hop	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: <ul style="list-style-type: none"> • I—IGP. • E—EGP. • ?—Incomplete; typically, the AS path was aggregated.

Sample Output

show route terse

```
user@host> show route terse
inet.0: 12 destinations, 12 routes (11 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf Metric 1    Metric 2    Next hop      AS path
* 0.0.0.0/0        S   5                      >111.222.5.254
* 1.0.0.1/32       D   0                      >at-5/3/0.0
* 1.0.0.2/32       L   0                      Local
* 12.12.12.21/32   L   0                      Reject
* 13.13.13.13/32   D   0                      >t3-5/2/1.0
* 13.13.13.14/32   L   0                      Local
* 13.13.13.21/32   L   0                      Local
* 13.13.13.22/32   D   0                      >t3-5/2/0.0
 127.0.0.1/32      D   0                      >lo0.0
* 111.222.5.0/24   D   0                      >fxp0.0
* 111.222.5.81/32  L   0                      Local
* 224.0.0.5/32     O  10                      1          MultiRecv
```

BGP Feature Guide for QFX10000 Switches

PART 61

Overview

- [BGP Overview on page 3593](#)

BGP Overview

- [Understanding BGP on page 3593](#)
- [BGP Routes Overview on page 3596](#)
- [BGP Messages Overview on page 3597](#)

Understanding BGP

BGP is an exterior gateway protocol (EGP) that is used to exchange routing information among routers in different autonomous systems (ASs). BGP routing information includes the complete route to each destination. BGP uses the routing information to maintain a database of network reachability information, which it exchanges with other BGP systems. BGP uses the network reachability information to construct a graph of AS connectivity, which enables BGP to remove routing loops and enforce policy decisions at the AS level.

Multiprotocol BGP (MBGP) extensions enable BGP to support IP version 6 (IPv6). MBGP defines the attributes MP_REACH_NLRI and MP_UNREACH_NLRI, which are used to carry IPv6 reachability information. Network layer reachability information (NLRI) update messages carry IPv6 address prefixes of feasible routes.

BGP allows for policy-based routing. You can use routing policies to choose among multiple paths to a destination and to control the redistribution of routing information.

BGP uses TCP as its transport protocol, using port 179 for establishing connections. Running over a reliable transport protocol eliminates the need for BGP to implement update fragmentation, retransmission, acknowledgment, and sequencing.

The Junos OS routing protocol software supports BGP version 4. This version of BGP adds support for Classless Interdomain Routing (CIDR), which eliminates the concept of network classes. Instead of assuming which bits of an address represent the network by looking at the first octet, CIDR allows you to explicitly specify the number of bits in the network address, thus providing a means to decrease the size of the routing tables. BGP version 4 also supports aggregation of routes, including the aggregation of AS paths.

This section discusses the following topics:

- [Autonomous Systems on page 3594](#)
- [AS Paths and Attributes on page 3594](#)

- [External and Internal BGP on page 3594](#)
- [Multiple Instances of BGP on page 3595](#)

Autonomous Systems

An *autonomous system* (AS) is a set of routers that are under a single technical administration and normally use a single interior gateway protocol and a common set of metrics to propagate routing information within the set of routers. To other ASs, an AS appears to have a single, coherent interior routing plan and presents a consistent picture of what destinations are reachable through it.


AS Paths and Attributes

The routing information that BGP systems exchange includes the complete route to each destination, as well as additional information about the route. The route to each destination is called the *AS path*, and the additional route information is included in *path attributes*. BGP uses the AS path and the path attributes to completely determine the network topology. Once BGP understands the topology, it can detect and eliminate routing loops and select among groups of routes to enforce administrative preferences and routing policy decisions.

External and Internal BGP

BGP supports two types of exchanges of routing information: exchanges among different ASs and exchanges within a single AS. When used among ASs, BGP is called *external BGP* (EBGP) and BGP sessions perform *inter-AS routing*. When used within an AS, BGP is called *internal BGP* (IBGP) and BGP sessions perform *intra-AS routing*. [Figure 56](#) illustrates ASs, IBGP, and EBGP.

Figure 56: ASs, EBGP, and IBGP



A BGP system shares network reachability information with adjacent BGP systems, which are referred to as *neighbors* or *peers*.

BGP systems are arranged into *groups*. In an IBGP group, all peers in the group—called *internal peers*—are in the same AS. Internal peers can be anywhere in the local AS and do not have to be directly connected to one another. Internal groups use routes from an

IGP to resolve forwarding addresses. They also propagate external routes among all other internal routers running IBGP, computing the next hop by taking the BGP next hop received with the route and resolving it using information from one of the interior gateway protocols.

In an EBGP group, the peers in the group—called *external peers*—are in different ASs and normally share a subnet. In an external group, the next hop is computed with respect to the interface that is shared between the external peer and the local router.

Multiple Instances of BGP

You can configure multiple instances of BGP at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols]

Multiple instances of BGP are primarily used for Layer 3 VPN support.

IGP peers and external BGP (EBGP) peers (both nonmultihop and multihop) are all supported for routing instances. BGP peering is established over one of the interfaces configured under the **routing-instances** hierarchy.



NOTE: When a BGP neighbor sends BGP messages to the local routing device, the incoming interface on which these messages are received must be configured in the same routing instance that the BGP neighbor configuration exists in. This is true for neighbors that are a single hop away or multiple hops away.

Routes learned from the BGP peer are added to the **instance-name.inet.0** table by default. You can configure import and export policies to control the flow of information into and out of the instance routing table.

For Layer 3 VPN support, configure BGP on the provider edge (PE) router to receive routes from the customer edge (CE) router and to send the instances' routes to the CE router if necessary. You can use multiple instances of BGP to maintain separate per-site forwarding tables for keeping VPN traffic separate on the PE router.

You can configure import and export policies that allow the service provider to control and rate-limit traffic to and from the customer.

You can configure an EBGP multihop session for a VRF routing instance. Also, you can set up the EBGP peer between the PE and CE routers by using the loopback address of the CE router instead of the interface addresses.

Related Documentation

- [BGP Routes Overview on page 3596](#)
- [BGP Messages Overview on page 3597](#)

BGP Routes Overview

A BGP route is a destination, described as an IP address prefix, and information that describes the path to the destination.

The following information describes the path:

- AS path, which is a list of numbers of the ASs that a route passes through to reach the local router. The first number in the path is that of the last AS in the path—the AS closest to the local router. The last number in the path is the AS farthest from the local router, which is generally the origin of the path.
- Path attributes, which contain additional information about the AS path that is used in routing policy.

BGP peers advertise routes to each other in update messages.

BGP stores its routes in the Junos OS routing table (**inet.0**). The routing table stores the following information about BGP routes:

- Routing information learned from update messages received from peers
- Local routing information that BGP applies to routes because of local policies
- Information that BGP advertises to BGP peers in update messages

For each prefix in the routing table, the routing protocol process selects a single best path, called the active path. Unless you configure BGP to advertise multiple paths to the same destination, BGP advertises only the active path.

The BGP router that first advertises a route assigns it one of the following values to identify its origin. During route selection, the lowest origin value is preferred.

- **0**—The router originally learned the route through an IGP (OSPF, IS-IS, or a static route).
- **1**—The router originally learned the route through an EGP (most likely BGP).
- **2**—The route's origin is unknown.

Related Documentation

- [Understanding BGP Path Selection on page 3793](#)
- [Example: Advertising Multiple Paths in BGP on page 3844](#)

BGP Messages Overview

All BGP messages have the same fixed-size header, which contains a marker field that is used for both synchronization and authentication, a length field that indicates the length of the packet, and a type field that indicates the message type (for example, open, update, notification, keepalive, and so on).

This section discusses the following topics:

- [Open Messages on page 3597](#)
- [Update Messages on page 3597](#)
- [Keepalive Messages on page 3598](#)
- [Notification Messages on page 3598](#)
- [Route-Refresh Messages on page 3598](#)

Open Messages

After a TCP connection is established between two BGP systems, they exchange BGP open messages to create a BGP connection between them. Once the connection is established, the two systems can exchange BGP messages and data traffic.

Open messages consist of the BGP header plus the following fields:

- Version—The current BGP version number is 4.
- Local AS number—You configure this by including the **autonomous-system** statement at the **[edit routing-options]** or **[edit logical-systems *logical-system-name* routing-options]** hierarchy level.
- Hold time—Proposed hold-time value. You configure the local hold time with the BGP **hold-time** statement.
- BGP identifier—IP address of the BGP system. This address is determined when the system starts and is the same for every local interface and every BGP peer. You can configure the BGP identifier by including the **router-id** statement at the **[edit routing-options]** or **[edit logical-systems *logical-system-name* routing-options]** hierarchy level. By default, BGP uses the IP address of the first interface it finds in the router.
- Parameter field length and the parameter itself—These are optional fields.

Update Messages

BGP systems send update messages to exchange network reachability information. BGP systems use this information to construct a graph that describes the relationships among all known ASs.

Update messages consist of the BGP header plus the following optional fields:

- Unfeasible routes length—Length of the withdrawn routes field
- Withdrawn routes—IP address prefixes for the routes being withdrawn from service because they are no longer deemed reachable

- Total path attribute length—Length of the path attributes field; it lists the path attributes for a feasible route to a destination
- Path attributes—Properties of the routes, including the path origin, the multiple exit discriminator (MED), the originating system's preference for the route, and information about aggregation, communities, confederations, and route reflection
- Network layer reachability information (NLRI)—IP address prefixes of feasible routes being advertised in the update message

Keepalive Messages

BGP systems exchange keepalive messages to determine whether a link or host has failed or is no longer available. Keepalive messages are exchanged often enough so that the hold timer does not expire. These messages consist only of the BGP header.

Notification Messages

BGP systems send notification messages when an error condition is detected. After the message is sent, the BGP session and the TCP connection between the BGP systems are closed. Notification messages consist of the BGP header plus the error code and subcode, and data that describes the error.

Route-Refresh Messages

BGP systems send route-refresh messages to a peer only if they have received the route refresh capability advertisement from the peer. A BGP system must advertise the route refresh capability to its peers using BGP capabilities advertisement if it wants to receive route-refresh messages. This optional message is sent to request dynamic, inbound, BGP route updates from BGP peers or to send outbound route updates to a BGP peer.

Route-refresh messages consist of the following fields:

- AFI—Address Family Identifier (16-bit).
- Res—Reserved (8-bit) field, which must be set to 0 by the sender and ignored by the receiver.
- SAFI—Subsequent Address Family Identifier (8-bit).

If a peer without the route-refresh capability receives a route-refresh request message from a remote peer, the receiver ignores the message.

Related Documentation

- [Understanding BGP on page 3593](#)
- [BGP Routes Overview on page 3596](#)

PART 62

Configuring BGP

- [Basic BGP Configuration on page 3601](#)
- [BGP Path Attribute Configuration on page 3651](#)
- [BGP Policy Configuration on page 3765](#)
- [BGP BFD Configuration on page 3811](#)
- [BGP Load Balancing Configuration on page 3827](#)
- [IBGP Scaling Configuration on page 3873](#)
- [BGP Security Configuration on page 3897](#)
- [BGP Flap Configuration on page 3919](#)

Basic BGP Configuration

- [Examples: Configuring External BGP Peering on page 3601](#)
- [Examples: Configuring Internal BGP Peering on page 3624](#)
- [Configuring BGP Monitoring Protocol Version 3 on page 3647](#)


Examples: Configuring External BGP Peering

- [Understanding External BGP Peering Sessions on page 3601](#)
- [Example: Configuring External BGP Point-to-Point Peer Sessions on page 3602](#)
- [Example: Configuring External BGP on Logical Systems with IPv6 Interfaces on page 3609](#)

Understanding External BGP Peering Sessions

To establish point-to-point connections between peer autonomous systems (ASs), you configure a BGP session on each interface of a point-to-point link. Generally, such sessions are made at network exit points with neighboring hosts outside the AS. [Figure 57](#) shows an example of a BGP peering session.

Figure 57: BGP Peering Session



In [Figure 57](#), Router A is a gateway router for AS 3, and Router B is a gateway router for AS 10. For traffic internal to either AS, an interior gateway protocol (IGP) is used (OSPF, for instance). To route traffic between peer ASs, a BGP session is used.

You arrange BGP routing devices into groups of peers. Different peer groups must have different group types, AS numbers, or route reflector cluster identifiers.

To define a BGP group that recognizes only the specified BGP systems as peers, statically configure all the system's peers by including one or more **neighbor** statements. The peer neighbor's address can be either an IPv6 or IPv4 address.



NOTE: On SRX Series devices, the default mode for processing traffic is flow mode. To configure an SRX Series device as a border router, you must change the mode from flow-based processing to packet-based processing. Use the **set security forwarding-options family mpls mode packet-based** statement to configure the SRX device to packet mode. You must reboot the device for the configuration to take effect.

As the number of external BGP (EBGP) groups increases, the ability to support a large number of BGP sessions might become a scaling issue. The preferred way to configure a large number of BGP neighbors is to configure a few groups consisting of multiple neighbors per group. Supporting fewer EBGP groups generally scales better than supporting a large number of EBGP groups. This becomes more evident in the case of hundreds of EBGP groups when compared with a few EBGP groups with multiple peers in each group.

After the BGP peers are established, BGP routes are not automatically advertised by the BGP peers. At each BGP-enabled device, policy configuration is required to export the local, static, or IGP-learned routes into the BGP RIB and then advertise them as BGP routes to the other peers. BGP's advertisement policy, by default, does not advertise any non-BGP routes (such as local routes) to peers.

Example: Configuring External BGP Point-to-Point Peer Sessions

This example shows how to configure BGP point-to-point peer sessions.

- [Requirements on page 3602](#)
- [Overview on page 3602](#)
- [Configuration on page 3603](#)
- [Verification on page 3605](#)


Requirements

Before you begin, if the default BGP policy is not adequate for your network, configure routing policies to filter incoming BGP routes and to advertise BGP routes.

Overview

Figure 58 shows a network with BGP peer sessions. In the sample network, Device E in AS 17 has BGP peer sessions to a group of peers called **external-peers**. Peers A, B, and C reside in AS 22 and have IP addresses 10.10.10.2, 10.10.10.6, and 10.10.10.10. Peer D resides in AS 79, at IP address 10.21.7.2. This example shows the configuration on Device E.

Figure 58: Typical Network with BGP Peer Sessions



Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set interfaces ge-1/2/0 unit 0 description to-A
set interfaces ge-1/2/0 unit 0 family inet address 10.10.10.1/30
set interfaces ge-0/0/1 unit 5 description to-B
set interfaces ge-0/0/1 unit 5 family inet address 10.10.10.5/30
set interfaces ge-0/1/0 unit 9 description to-C
set interfaces ge-0/1/0 unit 9 family inet address 10.10.10.9/30
set interfaces ge-1/2/1 unit 21 description to-D
set interfaces ge-1/2/1 unit 21 family inet address 10.21.7.1/30
set protocols bgp group external-peers type external
set protocols bgp group external-peers peer-as 22
set protocols bgp group external-peers neighbor 10.10.10.2
set protocols bgp group external-peers neighbor 10.10.10.6
set protocols bgp group external-peers neighbor 10.10.10.10
set protocols bgp group external-peers neighbor 10.21.7.2 peer-as 79
set routing-options autonomous-system 17
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the BGP peer sessions:

1. Configure the interfaces to Peers A, B, C, and D.

[edit interfaces]

```
user@E# set ge-1/2/0 unit 0 description to-A
user@E# set ge-1/2/0 unit 0 family inet address 10.10.10.1/30
user@E# set ge-0/0/1 unit 5 description to-B
user@E# set ge-0/0/1 unit 5 family inet address 10.10.10.5/30
user@E# set ge-0/1/0 unit 9 description to-C
user@E# set ge-0/1/0 unit 9 family inet address 10.10.10.9/30
user@E# set ge-1/2/1 unit 21 description to-D
user@E# set ge-1/2/1 unit 21 family inet address 10.21.7.1/30
```

2. Set the autonomous system (AS) number.

```
[edit routing-options]
user@E# set autonomous-system 17
```

3. Create the BGP group, and add the external neighbor addresses.

```
[edit protocols bgp group external-peers]
user@E# set neighbor 10.10.10.2
user@E# set neighbor 10.10.10.6
user@E# set neighbor 10.10.10.10
```

4. Specify the autonomous system (AS) number of the external AS.

```
[edit protocols bgp group external-peers]
user@E# set peer-as 22
```

5. Add Peer D, and set the AS number at the individual neighbor level.

The neighbor configuration overrides the group configuration. So, while **peer-as 22** is set for all the other neighbors in the group, **peer-as 79** is set for neighbor 10.21.7.2.

```
[edit protocols bgp group external-peers]
user@E# set neighbor 10.21.7.2 peer-as 79
```

6. Set the peer type to external BGP (EBGP).

```
[edit protocols bgp group external-peers]
user@E# set type external
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@E# show interfaces
ge-1/2/0 {
  unit 0 {
    description to-A;
    family inet {
      address 10.10.10.1/30;
    }
  }
}
ge-0/0/1 {
  unit 5 {
    description to-B;
    family inet {
      address 10.10.10.5/30;
    }
  }
}
```

```

    }
  }
  ge-0/1/0 {
    unit 9 {
      description to-C;
      family inet {
        address 10.10.10.9/30;
      }
    }
  }
  ge-1/2/1 {
    unit 21 {
      description to-D;
      family inet {
        address 10.21.7.1/30;
      }
    }
  }
}

[edit]
user@E# show protocols
bgp {
  group external-peers {
    type external;
    peer-as 22;
    neighbor 10.10.10.2;
    neighbor 10.10.10.6;
    neighbor 10.10.10.10;
    neighbor 10.21.7.2 {
      peer-as 79;
    }
  }
}

[edit]
user@E# show routing-options
autonomous-system 17;

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying BGP Neighbors on page 3605](#)
- [Verifying BGP Groups on page 3608](#)
- [Verifying BGP Summary Information on page 3608](#)

Verifying BGP Neighbors

Purpose Verify that BGP is running on configured interfaces and that the BGP session is active for each neighbor address.

Action From operational mode, run the **show bgp neighbor** command.

```
user@E> show bgp neighbor
```

```

Peer: 10.10.10.2+179 AS 22    Local: 10.10.10.1+65406 AS 17
Type: External    State: Established    Flags: <Sync>
Last State: OpenConfirm    Last Event: RecvKeepAlive
Last Error: None
Options: <Preference PeerAS Refresh>
Holdtime: 90 Preference: 170
Number of flaps: 0
Peer ID: 10.10.10.2    Local ID: 10.10.10.1    Active Holdtime: 90
Keepalive Interval: 30    Peer index: 0
BFD: disabled, down
Local Interface: ge-1/2/0.0
NLRI for restart configured on peer: inet-unicast
NLRI advertised by peer: inet-unicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-unicast
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Peer supports 4 byte AS extension (peer-as 22)
Peer does not support Addpath
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:      0
  Received prefixes:    0
  Accepted prefixes:    0
  Suppressed due to damping: 0
  Advertised prefixes:  0
Last traffic (seconds): Received 10    Sent 6    Checked 1
Input messages: Total 8522    Updates 1    Refreshes 0    Octets 161922
Output messages: Total 8433    Updates 0    Refreshes 0    Octets 160290
Output Queue[0]: 0

Peer: 10.10.10.6+54781 AS 22    Local: 10.10.10.5+179 AS 17
Type: External    State: Established    Flags: <Sync>
Last State: OpenConfirm    Last Event: RecvKeepAlive
Last Error: None
Options: <Preference PeerAS Refresh>
Holdtime: 90 Preference: 170
Number of flaps: 0
Peer ID: 10.10.10.6    Local ID: 10.10.10.1    Active Holdtime: 90
Keepalive Interval: 30    Peer index: 1
BFD: disabled, down
Local Interface: ge-0/0/1.5
NLRI for restart configured on peer: inet-unicast
NLRI advertised by peer: inet-unicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-unicast
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Peer supports 4 byte AS extension (peer-as 22)
Peer does not support Addpath

```

```

Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        0
  Accepted prefixes:        0
  Suppressed due to damping: 0
  Advertised prefixes:      0
Last traffic (seconds): Received 12   Sent 6   Checked 33
Input messages:  Total 8527   Updates 1   Refreshes 0   Octets 162057
Output messages: Total 8430   Updates 0   Refreshes 0   Octets 160233
Output Queue[0]: 0

Peer: 10.10.10.10+55012 AS 22  Local: 10.10.10.9+179 AS 17
  Type: External  State: Established  Flags: <Sync>
  Last State: OpenConfirm  Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference PeerAS Refresh>
  Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.10.10.10      Local ID: 10.10.10.1      Active Holdtime: 90
  Keepalive Interval: 30    Peer index: 2
  BFD: disabled, down
  Local Interface: fe-0/1/0.9
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 120
  Stale routes from peer are kept for: 300
  Restart time requested by this peer: 120
  NLRI that peer supports restart for: inet-unicast
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Peer supports 4 byte AS extension (peer-as 22)
  Peer does not support Addpath
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        0
  Accepted prefixes:        0
  Suppressed due to damping: 0
  Advertised prefixes:      0
Last traffic (seconds): Received 15   Sent 6   Checked 37
Input messages:  Total 8527   Updates 1   Refreshes 0   Octets 162057
Output messages: Total 8429   Updates 0   Refreshes 0   Octets 160214
Output Queue[0]: 0

Peer: 10.21.7.2+61867 AS 79  Local: 10.21.7.1+179 AS 17
  Type: External  State: Established  Flags: <ImportEval Sync>
  Last State: OpenConfirm  Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference PeerAS Refresh>
  Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.21.7.2      Local ID: 10.10.10.1      Active Holdtime: 90
  Keepalive Interval: 30    Peer index: 3
  BFD: disabled, down
  Local Interface: ge-1/2/1.21

```

```

NLRI for restart configured on peer: inet-unicast
NLRI advertised by peer: inet-unicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-unicast
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Peer supports 4 byte AS extension (peer-as 79)
Peer does not support Addpath
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        0
  Accepted prefixes:        0
  Suppressed due to damping: 0
  Advertised prefixes:      0
Last traffic (seconds): Received 28   Sent 24   Checked 47
Input messages: Total 8521   Updates 1   Refreshes 0   Octets 161943
Output messages: Total 8427   Updates 0   Refreshes 0   Octets 160176
Output Queue[0]: 0

```

Verifying BGP Groups

Purpose Verify that the BGP groups are configured correctly.

Action From operational mode, run the **show bgp group** command.

```

user@E> show bgp group
Group Type: External                               Local AS: 17
Name: external-peers  Index: 0                     Flags: <>
Holdtime: 0
Total peers: 4      Established: 4
10.10.10.2+179
10.10.10.6+54781
10.10.10.10+55012
10.21.7.2+61867
inet.0: 0/0/0/0

Groups: 1  Peers: 4   External: 4   Internal: 0   Down peers: 0   Flaps: 0
Table      Tot Paths  Act Paths Suppressed  History Damp State   Pending
inet.0      0          0          0          0         0      0       0

```

Verifying BGP Summary Information

Purpose Verify that the BGP configuration is correct.

Action From operational mode, run the **show bgp summary** command.

```

user@E> show bgp summary
Groups: 1 Peers: 4 Down peers: 0
Table      Tot Paths  Act Paths Suppressed  History Damp State   Pending
inet.0      0          0          0          0         0      0       0
Peer      AS      InPkt    OutPkt    OutQ    Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.10.10.2      22      8559      8470      0        0 2d 16:12:56

```


0/0/0/0	0/0/0/0				
10.10.10.6	22	8566	8468	0	0 2d 16:12:12
0/0/0/0	0/0/0/0				
10.10.10.10	22	8565	8466	0	0 2d 16:11:31
0/0/0/0	0/0/0/0				
10.21.7.2	79	8560	8465	0	0 2d 16:10:58
0/0/0/0	0/0/0/0				

Example: Configuring External BGP on Logical Systems with IPv6 Interfaces

This example shows how to configure external BGP (EBGP) point-to-point peer sessions on logical systems with IPv6 interfaces.

- [Requirements on page 3609](#)
- [Overview on page 3609](#)
- [Configuration on page 3610](#)
- [Verification on page 3619](#)

Requirements

In this example, no special configuration beyond device initialization is required.

Overview

Junos OS supports EBGP peer sessions by means of IPv6 addresses. An IPv6 peer session can be configured when an IPv6 address is specified in the **neighbor** statement. This example uses EUI-64 to generate IPv6 addresses that are automatically applied to the interfaces. An EUI-64 address is an IPv6 address that uses the IEEE EUI-64 format for the interface identifier portion of the address (the last 64 bits).



NOTE: Alternatively, you can configure EBGP sessions using manually assigned 128-bit IPv6 addresses.

If you use 128-bit link-local addresses for the interfaces, you must include the **local-interface** statement. This statement is valid only for 128-bit IPv6 link-local addresses and is mandatory for configuring an IPv6 EBGP link-local peer session.


Configuring EBGP peering using link-local addresses is only applicable for directly connected interfaces. There is no support for multihop peering.

After your interfaces are up, you can use the **show interfaces terse** command to view the EUI-64-generated IPv6 addresses on the interfaces. You must use these generated addresses in the BGP **neighbor** statements. This example demonstrates the full end-to-end procedure.

In this example, Frame Relay interface encapsulation is applied to the logical tunnel (**lt**) interfaces. This is a requirement because only Frame Relay encapsulation is supported when IPv6 addresses are configured on the **lt** interfaces.

Figure 59 shows a network with BGP peer sessions. In the sample network, Router R1 has five logical systems configured. Device E in autonomous system (AS) 17 has BGP peer sessions to a group of peers called **external-peers**. Peers A, B, and C reside in AS 22. This example shows the step-by-step configuration on Logical System A and Logical System E.

Figure 59: Typical Network with BGP Peer Sessions



Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

Device A

```
set logical-systems A interfaces lt-0/1/0 unit 1 description to-E
set logical-systems A interfaces lt-0/1/0 unit 1 encapsulation frame-relay
set logical-systems A interfaces lt-0/1/0 unit 1 dlci 1
set logical-systems A interfaces lt-0/1/0 unit 1 peer-unit 25
set logical-systems A interfaces lt-0/1/0 unit 1 family inet6 address 2001:db8:0:1::/64
  eui-64
set logical-systems A interfaces lo0 unit 1 family inet6 address 2001:db8::1/128
set logical-systems A protocols bgp group external-peers type external
set logical-systems A protocols bgp group external-peers peer-as 17
```

```

set logical-systems A protocols bgp group external-peers neighbor
  2001:db8:0:1:2a0:a502:0:19da
set logical-systems A routing-options router-id 1.1.1.1
set logical-systems A routing-options autonomous-system 22

Device B
set logical-systems B interfaces lt-0/1/0 unit 6 description to-E
set logical-systems B interfaces lt-0/1/0 unit 6 encapsulation frame-relay
set logical-systems B interfaces lt-0/1/0 unit 6 dlci 6
set logical-systems B interfaces lt-0/1/0 unit 6 peer-unit 5
set logical-systems B interfaces lt-0/1/0 unit 6 family inet6 address 2001:db8:0:2::/64
  eui-64
set logical-systems B interfaces lo0 unit 2 family inet6 address 2001:db8::2/128
set logical-systems B protocols bgp group external-peers type external
set logical-systems B protocols bgp group external-peers peer-as 17
set logical-systems B protocols bgp group external-peers neighbor
  2001:db8:0:2:2a0:a502:0:5da
set logical-systems B routing-options router-id 2.2.2.2
set logical-systems B routing-options autonomous-system 22

Device C
set logical-systems C interfaces lt-0/1/0 unit 10 description to-E
set logical-systems C interfaces lt-0/1/0 unit 10 encapsulation frame-relay
set logical-systems C interfaces lt-0/1/0 unit 10 dlci 10
set logical-systems C interfaces lt-0/1/0 unit 10 peer-unit 9
set logical-systems C interfaces lt-0/1/0 unit 10 family inet6 address 2001:db8:0:3::/64
  eui-64
set logical-systems C interfaces lo0 unit 3 family inet6 address 2001:db8::3/128
set logical-systems C protocols bgp group external-peers type external
set logical-systems C protocols bgp group external-peers peer-as 17
set logical-systems C protocols bgp group external-peers neighbor
  2001:db8:0:3:2a0:a502:0:9da
set logical-systems C routing-options router-id 3.3.3.3
set logical-systems C routing-options autonomous-system 22

Device D
set logical-systems D interfaces lt-0/1/0 unit 7 description to-E
set logical-systems D interfaces lt-0/1/0 unit 7 encapsulation frame-relay
set logical-systems D interfaces lt-0/1/0 unit 7 dlci 7
set logical-systems D interfaces lt-0/1/0 unit 7 peer-unit 21
set logical-systems D interfaces lt-0/1/0 unit 7 family inet6 address 2001:db8:0:4::/64
  eui-64
set logical-systems D interfaces lo0 unit 4 family inet6 address 2001:db8::4/128
set logical-systems D protocols bgp group external-peers type external
set logical-systems D protocols bgp group external-peers peer-as 17
set logical-systems D protocols bgp group external-peers neighbor
  2001:db8:0:4:2a0:a502:0:15da
set logical-systems D routing-options router-id 4.4.4.4
set logical-systems D routing-options autonomous-system 79

Device E
set logical-systems E interfaces lt-0/1/0 unit 5 description to-B
set logical-systems E interfaces lt-0/1/0 unit 5 encapsulation frame-relay
set logical-systems E interfaces lt-0/1/0 unit 5 dlci 6
set logical-systems E interfaces lt-0/1/0 unit 5 peer-unit 6
set logical-systems E interfaces lt-0/1/0 unit 5 family inet6 address 2001:db8:0:2::/64
  eui-64
set logical-systems E interfaces lt-0/1/0 unit 9 description to-C
set logical-systems E interfaces lt-0/1/0 unit 9 encapsulation frame-relay

```

```

set logical-systems E interfaces lt-0/1/0 unit 9 dlci 10
set logical-systems E interfaces lt-0/1/0 unit 9 peer-unit 10
set logical-systems E interfaces lt-0/1/0 unit 9 family inet6 address 2001:db8:0:3::/64
  eui-64
set logical-systems E interfaces lt-0/1/0 unit 21 description to-D
set logical-systems E interfaces lt-0/1/0 unit 21 encapsulation frame-relay
set logical-systems E interfaces lt-0/1/0 unit 21 dlci 7
set logical-systems E interfaces lt-0/1/0 unit 21 peer-unit 7
set logical-systems E interfaces lt-0/1/0 unit 21 family inet6 address 2001:db8:0:4::/64
  eui-64
set logical-systems E interfaces lt-0/1/0 unit 25 description to-A
set logical-systems E interfaces lt-0/1/0 unit 25 encapsulation frame-relay
set logical-systems E interfaces lt-0/1/0 unit 25 dlci 1
set logical-systems E interfaces lt-0/1/0 unit 25 peer-unit 1
set logical-systems E interfaces lt-0/1/0 unit 25 family inet6 address 2001:db8:0:1::/64
  eui-64
set logical-systems E interfaces lo0 unit 5 family inet6 address 2001:db8::5/128
set logical-systems E protocols bgp group external-peers type external
set logical-systems E protocols bgp group external-peers peer-as 22
set logical-systems E protocols bgp group external-peers neighbor
  2001:db8:0:1:2a0:a502:0:1da
set logical-systems E protocols bgp group external-peers neighbor
  2001:db8:0:2:2a0:a502:0:6da
set logical-systems E protocols bgp group external-peers neighbor
  2001:db8:0:3:2a0:a502:0:ada
set logical-systems E protocols bgp group external-peers neighbor
  2001:db8:0:4:2a0:a502:0:7da peer-as 79
set logical-systems E routing-options router-id 5.5.5.5
set logical-systems E routing-options autonomous-system 17

```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the BGP peer sessions:

1. Run the **show interfaces terse** command to verify that the physical router has a logical tunnel (lt) interface.
2. On Logical System A, configure the interface encapsulation, peer-unit number, and DLCI to reach Logical System E.

```

user@R1> show interfaces terse
Interface                Admin Link Proto  Local                Remote
...
lt-0/1/0                  up    up
...

```

```

user@R1> set cli logical-system A
Logical system: A
[edit]
user@R1:A> edit
Entering configuration mode
[edit]
user@R1:A# edit interfaces
[edit interfaces]
user@R1:A# set lt-0/1/0 unit 1 encapsulation frame-relay

```

```
user@R1:A# set lt-0/1/0 unit 1 dlci 1
user@R1:A# set lt-0/1/0 unit 1 peer-unit 25
```

3. On Logical System A, configure the network address for the link to Peer E, and configure a loopback interface.

```
[edit interfaces]
user@R1:A# set lt-0/1/0 unit 1 description to-E
user@R1:A# set lt-0/1/0 unit 1 family inet6 address 2001:db8:0:1::/64 eui-64
user@R1:A# set lo0 unit 1 family inet6 address 2001:db8::1/128
```

4. On Logical System E, configure the interface encapsulation, peer-unit number, and DLCI to reach Logical System A.

```
user@R1> set cli logical-system E
Logical system: E
[edit]
user@R1:E> edit
Entering configuration mode
[edit]
user@R1:E# edit interfaces
[edit interfaces]
user@R1:E# set lt-0/1/0 unit 25 encapsulation frame-relay
user@R1:E# set lt-0/1/0 unit 25 dlci 1
user@R1:E# set lt-0/1/0 unit 25 peer-unit 1
```

5. On Logical System E, configure the network address for the link to Peer A, and configure a loopback interface.

```
[edit interfaces]
user@R1:E# set lt-0/1/0 unit 25 description to-A
user@R1:E# set lt-0/1/0 unit 25 family inet6 address 2001:db8:0:1::/64 eui-64
user@R1:E# set lo0 unit 5 family inet6 address 2001:db8::5/128
```

6. Run the **show interfaces terse** command to see the IPv6 addresses that are generated by EUI-64.

The 2001 addresses are used in this example in the BGP **neighbor** statements.



NOTE: The fe80 addresses are link-local addresses and are not used in this example.

```
user@R1:A> show interfaces terse
Interface          Admin Link Proto  Local                               Remote
Logical system: A

betsy@tp8:A> show interfaces terse
Interface          Admin Link Proto  Local                               Remote
lt-0/1/0
lt-0/1/0.1          up    up    inet6  2001:db8:0:1:2a0:a502:0:1da/64
                                     fe80::2a0:a502:0:1da/64
lo0
lo0.1               up    up    inet6  2001:db8::1
                                     fe80::2a0:a50f:fc56:1da

user@R1:E> show interfaces terse
```

Interface	Admin	Link	Proto	Local	Remote
lt-0/1/0					
lt-0/1/0.25	up	up	inet6	2001:db8:0:1:2a0:a502:0:19da/64	
				fe80::2a0:a502:0:19da/64	
lo0					
lo0.5	up	up	inet6	2001:db8::5	
				fe80::2a0:a50f:fc56:1da	

- Repeat the interface configuration on the other logical systems.

Configuring the External BGP Sessions

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the BGP peer sessions:

- On Logical System A, create the BGP group, and add the external neighbor address.

```
[edit protocols bgp group external-peers]
user@R1:A# set neighbor 2001:db8:0:1:2a0:a502:0:19da
```
- On Logical System E, create the BGP group, and add the external neighbor address.

```
[edit protocols bgp group external-peers]
user@R1:E# set neighbor 2001:db8:0:1:2a0:a502:0:1da
```
- On Logical System A, specify the autonomous system (AS) number of the external AS.

```
[edit protocols bgp group external-peers]
user@R1:A# set peer-as 17
```
- On Logical System E, specify the autonomous system (AS) number of the external AS.

```
[edit protocols bgp group external-peers]
user@R1:E# set peer-as 22
```
- On Logical System A, set the peer type to EBGp.

```
[edit protocols bgp group external-peers]
user@R1:A# set type external
```
- On Logical System E, set the peer type to EBGp.

```
[edit protocols bgp group external-peers]
user@R1:E# set type external
```
- On Logical System A, set the autonomous system (AS) number and router ID.

```
[edit routing-options]
user@R1:A# set router-id 1.1.1.1
user@R1:A# set autonomous-system 22
```
- On Logical System E, set the AS number and router ID.

```
[edit routing-options]
user@R1:E# set router-id 5.5.5.5
```

```
user@R1:E# set autonomous-system 17
```

9. Repeat these steps for Peers A, B, C, and D.

Results From configuration mode, confirm your configuration by entering the **show logical-systems** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@R1# show logical-systems
A {
  interfaces {
    lt-0/1/0 {
      unit 1 {
        description to-E;
        encapsulation frame-relay;
        dlci 1;
        peer-unit 25;
        family inet6 {
          address 2001:db8:0:1::/64 {
            eui-64;
          }
        }
      }
    }
  }
  lo0 {
    unit 1 {
      family inet6 {
        address 2001:db8::1/128;
      }
    }
  }
  protocols {
    bgp {
      group external-peers {
        type external;
        peer-as 17;
        neighbor 2001:db8:0:1:2a0:a502:0:19da;
      }
    }
    routing-options {
      router-id 1.1.1.1;
      autonomous-system 22;
    }
  }
}
B {
  interfaces {
    lt-0/1/0 {
      unit 6 {
        description to-E;
        encapsulation frame-relay;
        dlci 6;
        peer-unit 5;
        family inet6 {
```

```
        address 2001:db8:0:2::/64 {
            eui-64;
        }
    }
}
lo0 {
    unit 2 {
        family inet6 {
            address 2001:db8::2/128;
        }
    }
}
protocols {
    bgp {
        group external-peers {
            type external;
            peer-as 17;
            neighbor 2001:db8:0:2:2a0:a502:0:5da;
        }
    }
    routing-options {
        router-id 2.2.2.2;
        autonomous-system 22;
    }
}
C {
    interfaces {
        lt-0/1/0 {
            unit 10 {
                description to-E;
                encapsulation frame-relay;
                dlci 10;
                peer-unit 9;
                family inet6 {
                    address 2001:db8:0:3::/64 {
                        eui-64;
                    }
                }
            }
        }
    }
    lo0 {
        unit 3 {
            family inet6 {
                address 2001:db8::3/128;
            }
        }
    }
}
protocols {
    bgp {
        group external-peers {
            type external;
            peer-as 17;
            neighbor 2001:db8:0:3:2a0:a502:0:9da;
```



```

    }
  }
}
routing-options {
  router-id 3.3.3.3;
  autonomous-system 22;
}
}
D {
  interfaces {
    lt-0/1/0 {
      unit 7 {
        description to-E;
        encapsulation frame-relay;
        dlci 7;
        peer-unit 21;
        family inet6 {
          address 2001:db8:0:4::/64 {
            eui-64;
          }
        }
      }
    }
  }
  lo0 {
    unit 4 {
      family inet6 {
        address 2001:db8::4/128;
      }
    }
  }
}
protocols {
  bgp {
    group external-peers {
      type external;
      peer-as 17;
      neighbor 2001:db8:0:4:2a0:a502:0:15da;
    }
  }
  routing-options {
    router-id 4.4.4.4;
    autonomous-system 79;
  }
}
E {
  interfaces {
    lt-0/1/0 {
      unit 5 {
        description to-B;
        encapsulation frame-relay;
        dlci 6;
        peer-unit 6;
        family inet6 {
          address 2001:db8:0:2::/64 {
            eui-64;
          }
        }
      }
    }
  }
}

```

```
    }
  }
  unit 9 {
    description to-C;
    encapsulation frame-relay;
    dlci 10;
    peer-unit 10;
    family inet6 {
      address 2001:db8:0:3::/64 {
        eui-64;
      }
    }
  }
  unit 21 {
    description to-D;
    encapsulation frame-relay;
    dlci 7;
    peer-unit 7;
    family inet6 {
      address 2001:db8:0:4::/64 {
        eui-64;
      }
    }
  }
  unit 25 {
    description to-A;
    encapsulation frame-relay;
    dlci 1;
    peer-unit 1;
    family inet6 {
      address 2001:db8:0:1::/64 {
        eui-64;
      }
    }
  }
}
lo0 {
  unit 5 {
    family inet6 {
      address 2001:db8::5/128;
    }
  }
}
}
protocols {
  bgp {
    group external-peers {
      type external;
      peer-as 22;
      neighbor 2001:db8:0:1:2a0:a502:0:1da;
      neighbor 2001:db8:0:2:2a0:a502:0:6da;
      neighbor 2001:db8:0:3:2a0:a502:0:ada;
      neighbor 2001:db8:0:4:2a0:a502:0:7da {
        peer-as 79;
      }
    }
  }
}
```

```

    }
  }
  routing-options {
    router-id 5.5.5.5;
    autonomous-system 17;
  }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying BGP Neighbors on page 3619](#)
- [Verifying BGP Groups on page 3622](#)
- [Verifying BGP Summary Information on page 3622](#)
- [Checking the Routing Table on page 3622](#)

Verifying BGP Neighbors

Purpose Verify that BGP is running on configured interfaces and that the BGP session is active for each neighbor address.

Action From operational mode, run the **show bgp neighbor** command.

```

user@R1:E> show bgp neighbor
Peer: 2001:db8:0:1:2a0:a502:0:1da+54987 AS 22 Local:
2001:db8:0:1:2a0:a502:0:19da+179 AS 17
  Type: External   State: Established   Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: Open Message Error
  Options: <Preference PeerAS Refresh>
  Holdtime: 90 Preference: 170
  Number of flaps: 0
  Error: 'Open Message Error' Sent: 20 Recv: 0
  Peer ID: 1.1.1.1      Local ID: 5.5.5.5      Active Holdtime: 90
  Keepalive Interval: 30      Peer index: 0
  BFD: disabled, down
  Local Interface: lt-0/1/0.25
  NLRI for restart configured on peer: inet6-unicast
  NLRI advertised by peer: inet6-unicast
  NLRI for this session: inet6-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  NLRI that restart is negotiated for: inet6-unicast
  NLRI of received end-of-rib markers: inet6-unicast
  NLRI of all end-of-rib markers sent: inet6-unicast
  Peer supports 4 byte AS extension (peer-as 22)
  Peer does not support Addpath
  Table inet6.0 Bit: 10000
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:          0
    Received prefixes:        0
    Accepted prefixes:        0

```

```

    Suppressed due to damping:    0
    Advertised prefixes:          0
    Last traffic (seconds): Received 7    Sent 18    Checked 81
    Input messages:  Total 1611    Updates 1      Refreshes 0      Octets 30660
    Output messages: Total 1594    Updates 0      Refreshes 0      Octets 30356
    Output Queue[0]: 0

Peer: 2001:db8:0:2:2a0:a502:0:6da+179 AS 22 Local:
2001:db8:0:2:2a0:a502:0:5da+55502 AS 17
  Type: External    State: Established    Flags: <Sync>
  Last State: OpenConfirm    Last Event: RecvKeepAlive
  Last Error: Open Message Error
  Options: <Preference PeerAS Refresh>
  Holdtime: 90 Preference: 170
  Number of flaps: 0
  Error: 'Open Message Error' Sent: 26 Recv: 0
  Peer ID: 2.2.2.2      Local ID: 5.5.5.5      Active Holdtime: 90
  Keepalive Interval: 30      Peer index: 2
  BFD: disabled, down
  Local Interface: lt-0/1/0.5
  NLRI for restart configured on peer: inet6-unicast
  NLRI advertised by peer: inet6-unicast
  NLRI for this session: inet6-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  NLRI that restart is negotiated for: inet6-unicast
  NLRI of received end-of-rib markers: inet6-unicast
  NLRI of all end-of-rib markers sent: inet6-unicast
  Peer supports 4 byte AS extension (peer-as 22)
  Peer does not support Addpath
  Table inet6.0 Bit: 10000
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:          0
    Received prefixes:        0
    Accepted prefixes:        0
    Suppressed due to damping: 0
    Advertised prefixes:      0
    Last traffic (seconds): Received 15    Sent 8      Checked 8
    Input messages:  Total 1610    Updates 1      Refreshes 0      Octets 30601
    Output messages: Total 1645    Updates 0      Refreshes 0      Octets 32417
    Output Queue[0]: 0

Peer: 2001:db8:0:3:2a0:a502:0:ada+55983 AS 22 Local:
2001:db8:0:3:2a0:a502:0:9da+179 AS 17
  Type: External    State: Established    Flags: <Sync>
  Last State: OpenConfirm    Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference PeerAS Refresh>
  Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 3.3.3.3      Local ID: 5.5.5.5      Active Holdtime: 90
  Keepalive Interval: 30      Peer index: 3
  BFD: disabled, down
  Local Interface: lt-0/1/0.9
  NLRI for restart configured on peer: inet6-unicast
  NLRI advertised by peer: inet6-unicast
  NLRI for this session: inet6-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300

```

```

Peer does not support Restarter functionality
NLRI that restart is negotiated for: inet6-unicast
NLRI of received end-of-rib markers: inet6-unicast
NLRI of all end-of-rib markers sent: inet6-unicast
Peer supports 4 byte AS extension (peer-as 22)
Peer does not support Addpath
Table inet6.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        0
  Accepted prefixes:        0
  Suppressed due to damping: 0
  Advertised prefixes:      0
Last traffic (seconds): Received 21   Sent 21   Checked 67
Input messages:  Total 1610   Updates 1     Refreshes 0     Octets 30641
Output messages: Total 1587   Updates 0     Refreshes 0     Octets 30223
Output Queue[0]: 0

Peer: 2001:db8:0:4:2a0:a502:0:7da+49255 AS 79 Local:
2001:db8:0:4:2a0:a502:0:15da+179 AS 17
  Type: External   State: Established   Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference PeerAS Refresh>
  Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 4.4.4.4      Local ID: 5.5.5.5      Active Holdtime: 90
  Keepalive Interval: 30      Peer index: 1
  BFD: disabled, down
  Local Interface: lt-0/1/0.21
  NLRI for restart configured on peer: inet6-unicast
  NLRI advertised by peer: inet6-unicast
  NLRI for this session: inet6-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  NLRI that restart is negotiated for: inet6-unicast
  NLRI of received end-of-rib markers: inet6-unicast
  NLRI of all end-of-rib markers sent: inet6-unicast
  Peer supports 4 byte AS extension (peer-as 79)
  Peer does not support Addpath
  Table inet6.0 Bit: 10000
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:          0
    Received prefixes:        0
    Accepted prefixes:        0
    Suppressed due to damping: 0
    Advertised prefixes:      0
  Last traffic (seconds): Received 6    Sent 17    Checked 25
  Input messages:  Total 1615   Updates 1     Refreshes 0     Octets 30736
  Output messages: Total 1593   Updates 0     Refreshes 0     Octets 30337
  Output Queue[0]: 0

```

Meaning IPv6 unicast network layer reachability information (NLRI) is being exchanged between the neighbors.

Verifying BGP Groups

Purpose Verify that the BGP groups are configured correctly.

Action From operational mode, run the **show bgp group** command.

```
user@R1:~> show bgp group
Group Type: External                               Local AS: 17
  Name: external-peers  Index: 0                   Flags: <>
  Holdtime: 0
  Total peers: 4      Established: 4
  2001:db8:0:1:2a0:a502:0:1da+54987
  2001:db8:0:2:2a0:a502:0:6da+179
  2001:db8:0:3:2a0:a502:0:ada+55983
  2001:db8:0:4:2a0:a502:0:7da+49255
  inet6.0: 0/0/0/0

Groups: 1 Peers: 4 External: 4 Internal: 0 Down peers: 0 Flaps: 0
Table      Tot Paths  Act Paths Suppressed  History Damp State  Pending
inet6.0           0          0          0          0          0          0
inet6.2           0          0          0          0          0          0
```

Meaning The group type is external, and the group has four peers.

Verifying BGP Summary Information

Purpose Verify that the BGP that the peer relationships are established.

Action From operational mode, run the **show bgp summary** command.

```
user@R1:~> show bgp summary
Groups: 1 Peers: 4 Down peers: 0
Table      Tot Paths  Act Paths Suppressed  History Damp State  Pending
inet6.0           0          0          0          0          0          0
inet6.2           0          0          0          0          0          0
Peer      AS      InPkt    OutPkt    OutQ    Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
2001:db8:0:1:2a0:a502:0:1da      22    1617    1600      0      0
  12:07:00 Establ
    inet6.0: 0/0/0/0
2001:db8:0:2:2a0:a502:0:6da      22    1616    1651      0      0
  12:06:56 Establ
    inet6.0: 0/0/0/0
2001:db8:0:3:2a0:a502:0:ada      22    1617    1594      0      0
  12:04:32 Establ
    inet6.0: 0/0/0/0
2001:db8:0:4:2a0:a502:0:7da      79    1621    1599      0      0
  12:07:00 Establ
    inet6.0: 0/0/0/0
```

Meaning The Down peers: 0 output shows that the BGP peers are in the established state.

Checking the Routing Table

Purpose Verify that the inet6.0 routing table is populated with local and direct routes.

Action From operational mode, run the **show route** command.

```

user@R1:E> show route
inet6.0: 15 destinations, 18 routes (15 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8::5/128    *[Direct/0] 12:41:18
                  > via lo0.5
2001:db8:0:1::/64  *[Direct/0] 14:40:01
                  > via lt-0/1/0.25
2001:db8:0:1:2a0:a502:0:19da/128
                  *[Local/0] 14:40:01
                  Local via lt-0/1/0.25
2001:db8:0:2::/64  *[Direct/0] 14:40:02
                  > via lt-0/1/0.5
2001:db8:0:2:2a0:a502:0:5da/128
                  *[Local/0] 14:40:02
                  Local via lt-0/1/0.5
2001:db8:0:3::/64  *[Direct/0] 14:40:02
                  > via lt-0/1/0.9
2001:db8:0:3:2a0:a502:0:9da/128
                  *[Local/0] 14:40:02
                  Local via lt-0/1/0.9
2001:db8:0:4::/64  *[Direct/0] 14:40:01
                  > via lt-0/1/0.21
2001:db8:0:4:2a0:a502:0:15da/128
                  *[Local/0] 14:40:01
                  Local via lt-0/1/0.21
fe80::/64         *[Direct/0] 14:40:02
                  > via lt-0/1/0.5
                  [Direct/0] 14:40:02
                  > via lt-0/1/0.9
                  [Direct/0] 14:40:01
                  > via lt-0/1/0.21
                  [Direct/0] 14:40:01
                  > via lt-0/1/0.25
fe80::2a0:a502:0:5da/128
                  *[Local/0] 14:40:02
                  Local via lt-0/1/0.5
fe80::2a0:a502:0:9da/128
                  *[Local/0] 14:40:02
                  Local via lt-0/1/0.9
fe80::2a0:a502:0:15da/128
                  *[Local/0] 14:40:01
                  Local via lt-0/1/0.21
fe80::2a0:a502:0:19da/128
                  *[Local/0] 14:40:01
                  Local via lt-0/1/0.25
fe80::2a0:a50f:fc56:1da/128
                  *[Direct/0] 12:41:18
                  > via lo0.5

```

Meaning The inet6.0 routing table contains local and direct routes. To populate the routing table with other types of routes, you must configure routing policies.

Related Documentation

- [Examples: Configuring Internal BGP Peering on page 3624](#)
- [BGP Configuration Overview](#)


Examples: Configuring Internal BGP Peering

- [Understanding Internal BGP Peering Sessions on page 3624](#)
- [Example: Configuring Internal BGP Peer Sessions on page 3625](#)
- [Example: Configuring Internal BGP Peering Sessions on Logical Systems on page 3636](#)

Understanding Internal BGP Peering Sessions

When two BGP-enabled devices are in the same autonomous system (AS), the BGP session is called an *internal* BGP session, or IBGP session. BGP uses the same message types on IBGP and external BGP (EBGP) sessions, but the rules for when to send each message and how to interpret each message differ slightly. For this reason, some people refer to IBGP and EBGP as two separate protocols.

Figure 60: Internal and External BGP



In [Figure 60](#), Device Jackson, Device Memphis, and Device Biloxi have IBGP peer sessions with each other. Likewise, Device Miami and Device Atlanta have IBGP peer sessions between each other.

The purpose of IBGP is to provide a means by which EBGP route advertisements can be forwarded throughout the network. In theory, to accomplish this task you could redistribute all of your EBGP routes into an interior gateway protocol (IGP), such as OSPF or IS-IS. This, however, is not recommended in a production environment because of the large number of EBGP routes in the Internet and because of the way that IGPs operate. In short, with that many routes the IGP churns or crashes.

Generally, the loopback interface (lo0) is used to establish connections between IBGP peers. The loopback interface is always up as long as the device is operating. If there is a route to the loopback address, the IBGP peering session stays up. If a physical interface address is used instead and that interface goes up and down, the IBGP peering session also goes up and down. Thus the loopback interface provides fault tolerance in case the physical interface or the link goes down, if the device has link redundancy.

While IBGP neighbors do not need to be directly connected, they do need to be fully meshed. In this case, fully meshed means that each device is logically connected to every

other device through neighbor peer relationships. The **neighbor** statement creates the mesh. Because of the full mesh requirement of IBGP, you must configure individual peering sessions between all IBGP devices in the AS. The full mesh need not be physical links. Rather, the configuration on each routing device must create a full mesh of peer sessions (using multiple **neighbor** statements).



NOTE: The requirement for a full mesh is waived if you configure a confederation or route reflection.

To understand the full-mesh requirement, consider that an IBGP-learned route cannot be readvertised to another IBGP peer. The reason for preventing the readvertisement of IBGP routes and requiring the full mesh is to avoid routing loops within an AS. The AS path attribute is the means by which BGP routing devices avoid loops. The path information is examined for the local AS number only when the route is received from an EBGP peer. Because the attribute is only modified across AS boundaries, this system works well. However, the fact that the attribute is only modified across AS boundaries presents an issue inside the AS. For example, suppose that routing devices A, B, and C are all in the same AS. Device A receives a route from an EBGP peer and sends the route to Device B, which installs it as the active route. The route is then sent to Device C, which installs it locally and sends it back to Device A. If Device A installs the route, a loop is formed within the AS. The routing devices are not able to detect the loop because the AS path attribute is not modified during these advertisements. Therefore, the BGP protocol designers decided that the only assurance of never forming a routing loop was to prevent an IBGP peer from advertising an IBGP-learned route within the AS. For route reachability, the IBGP peers are fully meshed.

IBGP supports multihop connections, so IBGP neighbors can be located anywhere within the AS and often do not share a link. A recursive route lookup resolves the loopback peering address to an IP forwarding next hop. The lookup service is provided by static routes or an IGP such as OSPF, or BGP routes.

Example: Configuring Internal BGP Peer Sessions

This example shows how to configure internal BGP peer sessions.

- [Requirements on page 3625](#)
- [Overview on page 3625](#)
- [Configuration on page 3627](#)
- [Verification on page 3634](#)

Requirements

No special configuration beyond device initialization is required before you configure this example.

Overview

In this example, you configure internal BGP (IBGP) peer sessions. The loopback interface (lo0) is used to establish connections between IBGP peers. The loopback interface is

always up as long as the device is operating. If there is a route to the loopback address, the IBGP peer session stays up. If a physical interface address is used instead and that interface goes up and down, the IBGP peer session also goes up and down. Thus, if the device has link redundancy, the loopback interface provides fault tolerance in case the physical interface or one of the links goes down.

When a device peers with a remote device's loopback interface address, the local device expects BGP update messages to come from (be sourced by) the remote device's loopback interface address. The **local-address** statement enables you to specify the source information in BGP update messages. If you omit the **local-address** statement, the expected source of BGP update messages is based on the device's source address selection rules, which normally results in the egress interface address being the expected source of update messages. When this happens, the peer session is not established because a mismatch exists between the expected source address (the egress interface of the peer) and the actual source (the loopback interface of the peer). To make sure that the expected source address matches the actual source address, specify the loopback interface address in the **local-address** statement.

Because IBGP supports multihop connections, IBGP neighbors can be located anywhere within the autonomous system (AS) and often do not share a link. A recursive route lookup resolves the loopback peer address to an IP forwarding next hop. In this example, this service is provided by OSPF. Although interior gateway protocol (IGP) neighbors do not need to be directly connected, they do need to be fully meshed. In this case, fully meshed means that each device is logically connected to every other device through neighbor peer relationships. The **neighbor** statement creates the mesh.




NOTE: The requirement for a full mesh is waived if you configure a confederation or route reflection.

After the BGP peers are established, local routes are not automatically advertised by the BGP peers. At each BGP-enabled device, policy configuration is required to export the local, static, or IGP-learned routes into the BGP routing information base (RIB) and then advertise them as BGP routes to the other peers. BGP's advertisement policy, by default, does not advertise any non-BGP routes (such as local routes) to peers.

In the sample network, the devices in AS 17 are fully meshed in the group **internal-peers**. The devices have loopback addresses 192.168.6.5, 192.163.6.4, and 192.168.40.4.

Figure 61 shows a typical network with internal peer sessions.

Figure 61: Typical Network with IBGP Sessions



Configuration

- [Configuring Device A on page 3628](#)
- [Configuring Device B on page 3630](#)
- [Configuring Device C on page 3632](#)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device A

```
set interfaces ge-0/1/0 unit 1 description to-B
set interfaces ge-0/1/0 unit 1 family inet address 10.10.10.1/30
set interfaces lo0 unit 1 family inet address 192.168.6.5/32
set protocols bgp group internal-peers type internal
set protocols bgp group internal-peers description "connections to B and C"
set protocols bgp group internal-peers local-address 192.168.6.5
set protocols bgp group internal-peers export send-direct
set protocols bgp group internal-peers neighbor 192.163.6.4
set protocols bgp group internal-peers neighbor 192.168.40.4
set protocols ospf area 0.0.0.0 interface lo0.1 passive
set protocols ospf area 0.0.0.0 interface ge-0/1/0.1
set policy-options policy-statement send-direct term 2 from protocol direct
set policy-options policy-statement send-direct term 2 then accept
set routing-options router-id 192.168.6.5
set routing-options autonomous-system 17
```

Device B

```
set interfaces ge-0/1/0 unit 2 description to-A
set interfaces ge-0/1/0 unit 2 family inet address 10.10.10.2/30
set interfaces ge-0/1/1 unit 5 description to-C
set interfaces ge-0/1/1 unit 5 family inet address 10.10.10.5/30
set interfaces lo0 unit 2 family inet address 192.163.6.4/32
set protocols bgp group internal-peers type internal
set protocols bgp group internal-peers description "connections to A and C"
set protocols bgp group internal-peers local-address 192.163.6.4
set protocols bgp group internal-peers export send-direct
set protocols bgp group internal-peers neighbor 192.168.40.4
set protocols bgp group internal-peers neighbor 192.168.6.5
```

```
set protocols ospf area 0.0.0.0 interface lo0.2 passive
set protocols ospf area 0.0.0.0 interface ge-0/1/0.2
set protocols ospf area 0.0.0.0 interface ge-0/1/1.5
set policy-options policy-statement send-direct term 2 from protocol direct
set policy-options policy-statement send-direct term 2 then accept
set routing-options router-id 192.163.6.4
set routing-options autonomous-system 17
```

Device C

```
set interfaces ge-0/1/0 unit 6 description to-B
set interfaces ge-0/1/0 unit 6 family inet address 10.10.10.6/30
set interfaces lo0 unit 3 family inet address 192.168.40.4/32
set protocols bgp group internal-peers type internal
set protocols bgp group internal-peers description "connections to A and B"
set protocols bgp group internal-peers local-address 192.168.40.4
set protocols bgp group internal-peers export send-direct
set protocols bgp group internal-peers neighbor 192.163.6.4
set protocols bgp group internal-peers neighbor 192.168.6.5
set protocols ospf area 0.0.0.0 interface lo0.3 passive
set protocols ospf area 0.0.0.0 interface ge-0/1/0.6
set policy-options policy-statement send-direct term 2 from protocol direct
set policy-options policy-statement send-direct term 2 then accept
set routing-options router-id 192.168.40.4
set routing-options autonomous-system 17
```

Configuring Device A

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure internal BGP peer sessions on Device A:

1. Configure the interfaces.

```
[edit interfaces ge-0/1/0 unit 1]
user@A# set description to-B
user@A# set family inet address 10.10.10.1/30
```

```
[edit interfaces]
user@A# set lo0 unit 1 family inet address 192.168.6.5/32
```

2. Configure BGP.

The **neighbor** statements are included for both Device B and Device C, even though Device A is not directly connected to Device C.

```
[edit protocols bgp group internal-peers]
user@A# set type internal
user@A# set description "connections to B and C"
user@A# set local-address 192.168.6.5
user@A# set export send-direct
user@A# set neighbor 192.163.6.4
user@A# set neighbor 192.168.40.4
```

3. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
```

```

user@A# set interface lo0.1 passive
user@A# set interface ge-0/1/0.1

```

4. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```

[edit policy-options policy-statement send-direct term 2]
user@A# set from protocol direct
user@A# set then accept

```

5. Configure the router ID and the AS number.

```

[edit routing-options]
user@A# set router-id 192.168.6.5
user@A# set autonomous-system 17

```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@A# show interfaces
ge-0/1/0 {
  unit 1 {
    description to-B;
    family inet {
      address 10.10.10.1/30;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 192.168.6.5/32;
    }
  }
}

user@A# show policy-options
policy-statement send-direct {
  term 2 {
    from protocol direct;
    then accept;
  }
}

user@A# show protocols
bgp {
  group internal-peers {
    type internal;
    description "connections to B and C";
    local-address 192.168.6.5;
    export send-direct;
    neighbor 192.163.6.4;
    neighbor 192.168.40.4;
  }
}

```

```
}  
}  
ospf {  
  area 0.0.0.0 {  
    interface lo0.1 {  
      passive;  
    }  
    interface ge-0/1/0.1;  
  }  
}  
  
user@A# show routing-options  
router-id 192.168.6.5;  
autonomous-system 17;
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Device B

Step-by-Step Procedure The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.

To configure internal BGP peer sessions on Device B:

1. Configure the interfaces.

```
[edit interfaces ge-0/1/0 unit 2]  
user@B# set description to-A  
user@B# set family inet address 10.10.10.2/30
```

```
[edit interfaces ge-0/1/1]  
user@B# set unit 5 description to-C  
user@B# set unit 5 family inet address 10.10.10.5/30
```

```
[edit interfaces]  
user@B# set lo0 unit 2 family inet address 192.163.6.4/32
```

2. Configure BGP.

The **neighbor** statements are included for both Device B and Device C, even though Device A is not directly connected to Device C.

```
[edit protocols bgp group internal-peers]  
user@B# set type internal  
user@B# set description "connections to A and C"  
user@B# set local-address 192.163.6.4  
user@B# set export send-direct  
user@B# set neighbor 192.168.40.4  
user@B# set neighbor 192.168.6.5
```

3. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]  
user@B# set interface lo0.2 passive  
user@B# set interface ge-0/1/0.2  
user@B# set interface ge-0/1/1.5
```

4. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 2]
user@B# set from protocol direct
user@B# set then accept
```

5. Configure the router ID and the AS number.

```
[edit routing-options]
user@B# set router-id 192.163.6.4
user@B# set autonomous-system 17
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@B# show interfaces
ge-0/1/0 {
  unit 2 {
    description to-A;
    family inet {
      address 10.10.10.2/30;
    }
  }
}
ge-0/1/1 {
  unit 5 {
    description to-C;
    family inet {
      address 10.10.10.5/30;
    }
  }
}
lo0 {
  unit 2 {
    family inet {
      address 192.163.6.4/32;
    }
  }
}

user@B# show policy-options
policy-statement send-direct {
  term 2 {
    from protocol direct;
    then accept;
  }
}

user@B# show protocols
bgp {
  group internal-peers {
    type internal;
```

```
        description "connections to A and C";
        local-address 192.163.6.4;
        export send-direct;
        neighbor 192.168.40.4;
        neighbor 192.168.6.5;
    }
}
ospf {
    area 0.0.0.0 {
        interface lo0.2 {
            passive;
        }
        interface ge-0/1/0.2;
        interface ge-0/1/1.5;
    }
}
```

```
user@B# show routing-options
router-id 192.163.6.4;
autonomous-system 17;
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Device C

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure internal BGP peer sessions on Device C:

1. Configure the interfaces.

```
[edit interfaces ge-0/1/0 unit 6]
user@C# set description to-B
user@C# set family inet address 10.10.10.6/30

[edit interfaces]
user@C# set lo0 unit 3 family inet address 192.168.40.4/32
```

2. Configure BGP.

The **neighbor** statements are included for both Device B and Device C, even though Device A is not directly connected to Device C.

```
[edit protocols bgp group internal-peers]
user@C# set type internal
user@C# set description "connections to A and B"
user@C# set local-address 192.168.40.4
user@C# set export send-direct
user@C# set neighbor 192.163.6.4
user@C# set neighbor 192.168.6.5
```

3. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
user@C# set interface lo0.3 passive
user@C# set interface ge-0/1/0.6
```


4. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 2]
user@C# set from protocol direct
user@C# set then accept
```

5. Configure the router ID and the AS number.

```
[edit routing-options]
user@C# set router-id 192.168.40.4
user@C# set autonomous-system 17
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@C# show interfaces
ge-0/1/0 {
  unit 6 {
    description to-B;
    family inet {
      address 10.10.10.6/30;
    }
  }
}
lo0 {
  unit 3 {
    family inet {
      address 192.168.40.4/32;
    }
  }
}

user@C# show policy-options
policy-statement send-direct {
  term 2 {
    from protocol direct;
    then accept;
  }
}

user@C# show protocols
bgp {
  group internal-peers {
    type internal;
    description "connections to A and B";
    local-address 192.168.40.4;
    export send-direct;
    neighbor 192.163.6.4;
    neighbor 192.168.6.5;
  }
}
ospf {
```

```
area 0.0.0.0 {  
  interface lo0.3 {  
    passive;  
  }  
  interface ge-0/1/0.6;  
}  
}  
  
user@C# show routing-options  
router-id 192.168.40.4;  
autonomous-system 17;
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying BGP Neighbors on page 3634](#)
- [Verifying BGP Groups on page 3635](#)
- [Verifying BGP Summary Information on page 3636](#)
- [Verifying That BGP Routes Are Installed in the Routing Table on page 3636](#)

Verifying BGP Neighbors

Purpose Verify that BGP is running on configured interfaces and that the BGP session is active for each neighbor address.

Action From operational mode, enter the **show bgp neighbor** command.

```
user@A> show bgp neighbor  
Peer: 192.163.6.4+179 AS 17    Local: 192.168.6.5+58852 AS 17  
  Type: Internal    State: Established    Flags: Sync  
  Last State: OpenConfirm    Last Event: RecvKeepAlive  
  Last Error: None  
  Export: [ send-direct ]  
  Options: Preference LocalAddress Refresh  
  Local Address: 192.168.6.5 Holdtime: 90 Preference: 170  
  Number of flaps: 0  
  Peer ID: 192.163.6.4    Local ID: 192.168.6.5    Active Holdtime: 90  
  Keepalive Interval: 30    Peer index: 0  
  BFD: disabled, down  
  NLRI for restart configured on peer: inet-unicast  
  NLRI advertised by peer: inet-unicast  
  NLRI for this session: inet-unicast  
  Peer supports Refresh capability (2)  
  Restart time configured on the peer: 120  
  Stale routes from peer are kept for: 300  
  Restart time requested by this peer: 120  
  NLRI that peer supports restart for: inet-unicast  
  NLRI that restart is negotiated for: inet-unicast  
  NLRI of received end-of-rib markers: inet-unicast  
  NLRI of all end-of-rib markers sent: inet-unicast  
  Peer supports 4 byte AS extension (peer-as 17)  
  Peer does not support Addpath  
  Table inet.0 Bit: 10000  
    RIB State: BGP restart is complete
```

```

Send state: in sync
Active prefixes:          0
Received prefixes:        3
Accepted prefixes:        3
Suppressed due to damping: 0
Advertised prefixes:      2
Last traffic (seconds): Received 25   Sent 19   Checked 67
Input messages:  Total 2420   Updates 4     Refreshes 0     Octets 46055
Output messages: Total 2411   Updates 2     Refreshes 0     Octets 45921
Output Queue[0]: 0

Peer: 192.168.40.4+179 AS 17   Local: 192.168.6.5+56466 AS 17
Type: Internal   State: Established   Flags: Sync
Last State: OpenConfirm   Last Event: RecvKeepAlive
Last Error: None
Export: [ send-direct ]
Options: Preference LocalAddress Refresh
Local Address: 192.168.6.5 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer ID: 192.168.40.4   Local ID: 192.168.6.5   Active Holdtime: 90
Keepalive Interval: 30   Peer index: 1
BFD: disabled, down
NLRI for restart configured on peer: inet-unicast
NLRI advertised by peer: inet-unicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-unicast
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Peer supports 4 byte AS extension (peer-as 17)
Peer does not support Addpath
Table inet.0 Bit: 10000
RIB State: BGP restart is complete
Send state: in sync
Active prefixes:          0
Received prefixes:        2
Accepted prefixes:        2
Suppressed due to damping: 0
Advertised prefixes:      2
Last traffic (seconds): Received 7     Sent 21   Checked 24
Input messages:  Total 2412   Updates 2     Refreshes 0     Octets 45867
Output messages: Total 2409   Updates 2     Refreshes 0     Octets 45883
Output Queue[0]: 0

```

Verifying BGP Groups

Purpose Verify that the BGP groups are configured correctly.

Action From operational mode, enter the **show bgp group** command.

```

user@A> show bgp group
Group Type: Internal   AS: 17                               Local AS: 17
Name: internal-peers   Index: 0                               Flags: <Export Eval>
Export: [ send-direct ]
Holdtime: 0
Total peers: 2         Established: 2
192.163.6.4+179

```

```
192.168.40.4+179
inet.0: 0/5/5/0
```

Groups:	1	Peers:	2	External:	0	Internal:	2	Down peers:	0	Flaps:	0
Table		Tot	Paths	Act	Paths	Suppressed		History	Damp	State	Pending
inet.0		5		0		0		0		0	0

Verifying BGP Summary Information

Purpose Verify that the BGP configuration is correct.

Action From operational mode, enter the **show bgp summary** command.

```
user@A> show bgp summary
Groups: 1 Peers: 2 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet.0 5 0 0 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
192.163.6.4 17 2441 2432 0 0 18:18:52
0/3/3/0 0/0/0/0
192.168.40.4 17 2432 2430 0 0 18:18:48
0/2/2/0 0/0/0/0
```

Verifying That BGP Routes Are Installed in the Routing Table

Purpose Verify that the export policy configuration is causing the BGP routes to be installed in the routing tables of the peers.

Action From operational mode, enter the **show route protocol bgp** command.

```
user@A> show route protocol bgp
inet.0: 7 destinations, 12 routes (7 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.10.0/30 [BGP/170] 07:09:57, localpref 100, from 192.163.6.4
AS path: I
> to 10.10.10.2 via ge-0/1/0.1
10.10.10.4/30 [BGP/170] 07:09:57, localpref 100, from 192.163.6.4
AS path: I
> to 10.10.10.2 via ge-0/1/0.1
[BGP/170] 07:07:12, localpref 100, from 192.168.40.4
AS path: I
> to 10.10.10.2 via ge-0/1/0.1
192.163.6.4/32 [BGP/170] 07:09:57, localpref 100, from 192.163.6.4
AS path: I
> to 10.10.10.2 via ge-0/1/0.1
192.168.40.4/32 [BGP/170] 07:07:12, localpref 100, from 192.168.40.4
AS path: I
> to 10.10.10.2 via ge-0/1/0.1
```

Example: Configuring Internal BGP Peering Sessions on Logical Systems

This example shows how to configure internal BGP peer sessions on logical systems.

- [Requirements on page 3637](#)
- [Overview on page 3637](#)

- [Configuration on page 3637](#)
- [Verification on page 3644](#)

Requirements

In this example, no special configuration beyond device initialization is required.


Overview

In this example, you configure internal BGP (IBGP) peering sessions.

In the sample network, the devices in AS 17 are fully meshed in the group **internal-peers**. The devices have loopback addresses 192.168.6.5, 192.163.6.4, and 192.168.40.4.

[Figure 62](#) shows a typical network with internal peer sessions.

Figure 62: Typical Network with IBGP Sessions



Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set logical-systems A interfaces lt-0/1/0 unit 1 description to-B
set logical-systems A interfaces lt-0/1/0 unit 1 encapsulation ethernet
set logical-systems A interfaces lt-0/1/0 unit 1 peer-unit 2
set logical-systems A interfaces lt-0/1/0 unit 1 family inet address 10.10.10.1/30
set logical-systems A interfaces lo0 unit 1 family inet address 192.168.6.5/32
set logical-systems A protocols bgp group internal-peers type internal
set logical-systems A protocols bgp group internal-peers local-address 192.168.6.5
set logical-systems A protocols bgp group internal-peers export send-direct
set logical-systems A protocols bgp group internal-peers neighbor 192.163.6.4
set logical-systems A protocols bgp group internal-peers neighbor 192.168.40.4
set logical-systems A protocols ospf area 0.0.0.0 interface lo0.1 passive
set logical-systems A protocols ospf area 0.0.0.0 interface lt-0/1/0.1
set logical-systems A policy-options policy-statement send-direct term 2 from protocol direct
set logical-systems A policy-options policy-statement send-direct term 2 then accept
```

```
set logical-systems A routing-options router-id 192.168.6.5
set logical-systems A routing-options autonomous-system 17
set logical-systems B interfaces lt-0/1/0 unit 2 description to-A
set logical-systems B interfaces lt-0/1/0 unit 2 encapsulation ethernet
set logical-systems B interfaces lt-0/1/0 unit 2 peer-unit 1
set logical-systems B interfaces lt-0/1/0 unit 2 family inet address 10.10.10.2/30
set logical-systems B interfaces lt-0/1/0 unit 5 description to-C
set logical-systems B interfaces lt-0/1/0 unit 5 encapsulation ethernet
set logical-systems B interfaces lt-0/1/0 unit 5 peer-unit 6
set logical-systems B interfaces lt-0/1/0 unit 5 family inet address 10.10.10.5/30
set logical-systems B interfaces lo0 unit 2 family inet address 192.163.6.4/32
set logical-systems B protocols bgp group internal-peers type internal
set logical-systems B protocols bgp group internal-peers local-address 192.163.6.4
set logical-systems B protocols bgp group internal-peers export send-direct
set logical-systems B protocols bgp group internal-peers neighbor 192.168.40.4
set logical-systems B protocols bgp group internal-peers neighbor 192.168.6.5
set logical-systems B protocols ospf area 0.0.0.0 interface lo0.2 passive
set logical-systems B protocols ospf area 0.0.0.0 interface lt-0/1/0.2
set logical-systems B protocols ospf area 0.0.0.0 interface lt-0/1/0.5
set logical-systems B policy-options policy-statement send-direct term 2 from protocol
    direct
set logical-systems B policy-options policy-statement send-direct term 2 then accept
set logical-systems B routing-options router-id 192.163.6.4
set logical-systems B routing-options autonomous-system 17
set logical-systems C interfaces lt-0/1/0 unit 6 description to-B
set logical-systems C interfaces lt-0/1/0 unit 6 encapsulation ethernet
set logical-systems C interfaces lt-0/1/0 unit 6 peer-unit 5
set logical-systems C interfaces lt-0/1/0 unit 6 family inet address 10.10.10.6/30
set logical-systems C interfaces lo0 unit 3 family inet address 192.168.40.4/32
set logical-systems C protocols bgp group internal-peers type internal
set logical-systems C protocols bgp group internal-peers local-address 192.168.40.4
set logical-systems C protocols bgp group internal-peers export send-direct
set logical-systems C protocols bgp group internal-peers neighbor 192.163.6.4
set logical-systems C protocols bgp group internal-peers neighbor 192.168.6.5
set logical-systems C protocols ospf area 0.0.0.0 interface lo0.3 passive
set logical-systems C protocols ospf area 0.0.0.0 interface lt-0/1/0.6
set logical-systems C policy-options policy-statement send-direct term 2 from protocol
    direct
set logical-systems C policy-options policy-statement send-direct term 2 then accept
set logical-systems C routing-options router-id 192.168.40.4
set logical-systems C routing-options autonomous-system 17
```

Device A

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure internal BGP peer sessions on Device A:

1. Configure the interfaces.

```
[edit logical-systems A interfaces lt-0/1/0 unit 1]
user@R1# set description to-B
user@R1# set encapsulation ethernet
user@R1# set peer-unit 2
```

```

user@R1# set family inet address 10.10.10.1/30
user@R1# set family inet address 192.168.6.5/32
user@R1# up
user@R1# up
[edit logical-systems A interfaces]
user@R1# set lo0 unit 1 family inet address 192.168.6.5/32
user@R1# exit
[edit]
user@R1# edit logical-systems B interfaces lt-0/1/0
[edit logical-systems B interfaces lt-0/1/0]
user@R1# set unit 2 description to-A
user@R1# set unit 2 encapsulation ethernet
user@R1# set unit 2 peer-unit 1
user@R1# set unit 2 family inet address 10.10.10.2/30
user@R1# set unit 5 description to-C
user@R1# set unit 5 encapsulation ethernet
user@R1# set unit 5 peer-unit 6
user@R1# set family inet address 10.10.10.5/30
user@R1# up
[edit logical-systems B interfaces]
user@R1# set lo0 unit 2 family inet address 192.163.6.4/32
user@R1# exit
[edit]
user@R1# edit logical-systems C interfaces lt-0/1/0 unit 6
[edit logical-systems C interfaces lt-0/1/0 unit 6]
set description to-B
set encapsulation ethernet
set peer-unit 5
set family inet address 10.10.10.6/30
user@R1# up
user@R1# up
[edit logical-systems C interfaces]
set lo0 unit 3 family inet address 192.168.40.4/32

```

2. Configure BGP.

On Logical System A, the **neighbor** statements are included for both Device B and Device C, even though Logical System A is not directly connected to Device C.

```

[edit logical-systems A protocols bgp group internal-peers]
user@R1# set type internal
user@R1# set local-address 192.168.6.5
user@R1# set export send-direct
user@R1# set neighbor 192.163.6.4
user@R1# set neighbor 192.168.40.4

```

```

[edit logical-systems B protocols bgp group internal-peers]
user@R1# set type internal
user@R1# set local-address 192.163.6.4
user@R1# set export send-direct
user@R1# set neighbor 192.168.40.4
user@R1# set neighbor 192.168.6.5

```

```

[edit logical-systems C protocols bgp group internal-peers]
user@R1# set type internal
user@R1# set local-address 192.168.40.4

```

```
user@R1# set export send-direct
user@R1# set neighbor 192.163.6.4
user@R1# set neighbor 192.168.6.5
```

3. Configure OSPF.

```
[edit logical-systems A protocols ospf area 0.0.0.0]
user@R1# set interface lo0.1 passive
user@R1# set interface lt-0/1/0.1
```

```
[edit logical-systems A protocols ospf area 0.0.0.0]
user@R1# set interface lo0.2 passive
user@R1# set interface lt-0/1/0.2
user@R1# set interface lt-0/1/0.5
```

```
[edit logical-systems A protocols ospf area 0.0.0.0]
user@R1# set interface lo0.3 passive
user@R1# set interface lt-0/1/0.6
```

4. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit logical-systems A policy-options policy-statement send-direct term 2]
user@R1# set from protocol direct
user@R1# set then accept
```

```
[edit logical-systems B policy-options policy-statement send-direct term 2]
user@R1# set from protocol direct
user@R1# set then accept
```

```
[edit logical-systems C policy-options policy-statement send-direct term 2]
user@R1# set from protocol direct
user@R1# set then accept
```

5. Configure the router ID and the autonomous system (AS) number.

```
[edit logical-systems A routing-options]
user@R1# set router-id 192.168.6.5
user@R1# set autonomous-system 17
```

```
[edit logical-systems B routing-options]
user@R1# set router-id 192.163.6.4
user@R1# set autonomous-system 17
```

```
[edit logical-systems C routing-options]
user@R1# set router-id 192.168.40.4
user@R1# set autonomous-system 17
```

Results From configuration mode, confirm your configuration by entering the **show logical-systems** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show logical-systems
```



```

A {
  interfaces {
    lt-0/1/0 {
      unit 1 {
        description to-B;
        encapsulation ethernet;
        peer-unit 2;
        family inet {
          address 10.10.10.1/30;
        }
      }
    }
  }
  lo0 {
    unit 1 {
      family inet {
        address 192.168.6.5/32;
      }
    }
  }
}
protocols {
  bgp {
    group internal-peers {
      type internal;
      local-address 192.168.6.5;
      export send-direct;
      neighbor 192.163.6.4;
      neighbor 192.168.40.4;
    }
  }
  ospf {
    area 0.0.0.0 {
      interface lo0.1 {
        passive;
      }
      interface lt-0/1/0.1;
    }
  }
}
policy-options {
  policy-statement send-direct {
    term 2 {
      from protocol direct;
      then accept;
    }
  }
}
routing-options {
  router-id 192.168.6.5;
  autonomous-system 17;
}
}
B {
  interfaces {
    lt-0/1/0 {
      unit 2 {

```

```
        description to-A;
        encapsulation ethernet;
        peer-unit 1;
        family inet {
            address 10.10.10.2/30;
        }
    }
    unit 5 {
        description to-C;
        encapsulation ethernet;
        peer-unit 6;
        family inet {
            address 10.10.10.5/30;
        }
    }
}
lo0 {
    unit 2 {
        family inet {
            address 192.163.6.4/32;
        }
    }
}
}
protocols {
    bgp {
        group internal-peers {
            type internal;
            local-address 192.163.6.4;
            export send-direct;
            neighbor 192.168.40.4;
            neighbor 192.168.6.5;
        }
    }
    ospf {
        area 0.0.0.0 {
            interface lo0.2 {
                passive;
            }
            interface lt-0/1/0.2;
            interface lt-0/1/0.5;
        }
    }
}
policy-options {
    policy-statement send-direct {
        term 2 {
            from protocol direct;
            then accept;
        }
    }
}
routing-options {
    router-id 192.163.6.4;
    autonomous-system 17;
}
```

```

}
C {
  interfaces {
    lt-0/1/0 {
      unit 6 {
        description to-B;
        encapsulation ethernet;
        peer-unit 5;
        family inet {
          address 10.10.10.6/30;
        }
      }
    }
    lo0 {
      unit 3 {
        family inet {
          address 192.168.40.4/32;
        }
      }
    }
  }
  protocols {
    bgp {
      group internal-peers {
        type internal;
        local-address 192.168.40.4;
        export send-direct;
        neighbor 192.163.6.4;
        neighbor 192.168.6.5;
      }
    }
    ospf {
      area 0.0.0.0 {
        interface lo0.3 {
          passive;
        }
        interface lt-0/1/0.6;
      }
    }
  }
  policy-options {
    policy-statement send-direct {
      term 2 {
        from protocol direct;
        then accept;
      }
    }
  }
  routing-options {
    router-id 192.168.40.4;
    autonomous-system 17;
  }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying BGP Neighbors on page 3644](#)
- [Verifying BGP Groups on page 3645](#)
- [Verifying BGP Summary Information on page 3645](#)
- [Verifying That BGP Routes Are Installed in the Routing Table on page 3646](#)

Verifying BGP Neighbors

Purpose Verify that BGP is running on configured interfaces and that the BGP session is active for each neighbor address.

Action From the operational mode, enter the **show bgp neighbor** command.

```
user@R1> show bgp neighbor logical-system A
Peer: 192.163.6.4+179 AS 17      Local: 192.168.6.5+58852 AS 17
  Type: Internal    State: Established    Flags: <Sync>
  Last State: OpenConfirm    Last Event: RecvKeepAlive
  Last Error: None
  Export: [ send-direct ]
  Options: <Preference LocalAddress Refresh>
  Local Address: 192.168.6.5 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 192.163.6.4      Local ID: 192.168.6.5      Active Holdtime: 90
  Keepalive Interval: 30      Peer index: 0
  BFD: disabled, down
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 120
  Stale routes from peer are kept for: 300
  Restart time requested by this peer: 120
  NLRI that peer supports restart for: inet-unicast
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Peer supports 4 byte AS extension (peer-as 17)
  Peer does not support Addpath
  Table inet.0 Bit: 10000
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:          0
    Received prefixes:        3
    Accepted prefixes:        3
    Suppressed due to damping: 0
    Advertised prefixes:      2
  Last traffic (seconds): Received 16    Sent 1    Checked 63
  Input messages:  Total 15713  Updates 4    Refreshes 0    Octets 298622
  Output messages: Total 15690  Updates 2    Refreshes 0    Octets 298222
  Output Queue[0]: 0

Peer: 192.168.40.4+179 AS 17      Local: 192.168.6.5+56466 AS 17
  Type: Internal    State: Established    Flags: <Sync>
  Last State: OpenConfirm    Last Event: RecvKeepAlive
  Last Error: None
```

```

Export: [ send-direct ]
Options: <Preference LocalAddress Refresh>
Local Address: 192.168.6.5 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer ID: 192.168.40.4    Local ID: 192.168.6.5    Active Holdtime: 90
Keepalive Interval: 30    Peer index: 1
BFD: disabled, down
NLRI for restart configured on peer: inet-unicast
NLRI advertised by peer: inet-unicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-unicast
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Peer supports 4 byte AS extension (peer-as 17)
Peer does not support Addpath
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:      0
  Received prefixes:    2
  Accepted prefixes:    2
  Suppressed due to damping: 0
  Advertised prefixes:  2
Last traffic (seconds): Received 15    Sent 22    Checked 68
Input messages: Total 15688 Updates 2    Refreshes 0    Octets 298111
Output messages: Total 15688 Updates 2    Refreshes 0    Octets 298184
Output Queue[0]: 0

```

Verifying BGP Groups

Purpose Verify that the BGP groups are configured correctly.

Action From the operational mode, enter the **show bgp group** command.

```

user@A> show bgp group logical-system A
Group Type: Internal    AS: 17                      Local AS: 17
Name: internal-peers   Index: 0                    Flags: <Export Eval>
Export: [ send-direct ]
Holdtime: 0
Total peers: 2          Established: 2
192.163.6.4+179
192.168.40.4+179
inet.0: 0/5/5/0

Groups: 1  Peers: 2  External: 0  Internal: 2  Down peers: 0  Flaps: 0
Table      Tot Paths  Act Paths  Suppressed  History  Damp State  Pending
inet.0      5          0          0          0        0        0        0

```

Verifying BGP Summary Information

Purpose Verify that the BGP configuration is correct.

Action From the operational mode, enter the **show bgp summary** command.

```

user@A> show bgp summary logical-system A

```

```

Groups: 1 Peers: 2 Down peers: 0
Table      Tot Paths  Act Paths Suppressed  History  Damp State  Pending
inet.0      5          0          0          0        0      0      0
Peer        AS      InPkt   OutPkt   OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
192.163.6.4  17     15723    15700     0      0 4d 22:13:15
0/3/3/0      0/0/0/0
192.168.40.4 17     15698    15699     0      0 4d 22:13:11
0/2/2/0      0/0/0/0

```

Verifying That BGP Routes Are Installed in the Routing Table

Purpose Verify that the export policy configuration is working.

Action From the operational mode, enter the **show route protocol bgp** command.

```

user@A> show route protocol bgp logical-system A
inet.0: 7 destinations, 12 routes (7 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.10.0/30      [BGP/170] 4d 11:05:55, localpref 100, from 192.163.6.4
                  AS path: I
                  > to 10.10.10.2 via lt-0/1/0.1
10.10.10.4/30      [BGP/170] 4d 11:05:55, localpref 100, from 192.163.6.4
                  AS path: I
                  > to 10.10.10.2 via lt-0/1/0.1
                  [BGP/170] 4d 11:03:10, localpref 100, from 192.168.40.4
                  AS path: I
                  > to 10.10.10.2 via lt-0/1/0.1
192.163.6.4/32     [BGP/170] 4d 11:05:55, localpref 100, from 192.163.6.4
                  AS path: I
                  > to 10.10.10.2 via lt-0/1/0.1
192.168.40.4/32    [BGP/170] 4d 11:03:10, localpref 100, from 192.168.40.4
                  AS path: I
                  > to 10.10.10.2 via lt-0/1/0.1

```

Related Documentation

- [Examples: Configuring External BGP Peering on page 3601](#)

Configuring BGP Monitoring Protocol Version 3

BGP Monitoring Protocol (BMP) allows the Junos OS to send the BGP route information from the router to a monitoring application on a separate device. The monitoring application is called the BMP monitoring station or BMP station. To deploy BMP in your network, you need to configure BMP on each router and you also need to configure at least one BMP station. This procedure describes how to configure BMP on a router.

You can specify these settings for all BMP stations by configuring the statements described here at the **[edit routing-options bmp]** hierarchy level. You can also configure settings for specific BMP stations by configuring these statements at the **[edit routing-options bmp station *station-name*]** hierarchy level.

The following procedure describes how to configure BMP version 3 on the router:

1. Specify the name or address for the BMP monitoring station by configuring the **station-address** statement. You can specify one or the other but not both. The address must be a valid IPv4 or IPv6 address.

```
station-address (station-address | station-name);
```

2. Specify the authentication algorithm used to encrypt authentication between the BMP-enabled router and the BMP station using the **authentication-algorithm** statement.

```
authentication-algorithm algorithm;
```

You can specify one of the following types of authentication algorithms:

- **aes-128-cmac-96**—Cipher-based message authentication code (AES128, 96 bits).
- **hmac-sha-1-96**—Hash-based message authentication code (SHA1, 96 bits).
- **md5**—Message digest 5.

3. Specify an MD5 authentication key (password) using the **authentication-key** statement.

```
authentication-key key;
```

4. Specify the authentication key chain using the **authentication-key-chain** statement.

```
authentication-key-chain key-chain;
```

The authentication key chain itself needs to be configured at the **[edit security authentication-key-chains *key-chain*]** hierarchy level. For a detailed example, see [“Example: Configuring Route Authentication for BGP” on page 3898](#).

5. Specify how to handle a BMP station flap by configuring the **hold-down** statement. A flap is when the TCP session unexpectedly switches from established to non-established. The BMP station can be prevented from attempting to reconnect to the device for a specified period of time.

```
hold-down {
  seconds;
  flaps number;
  period seconds;
}
```

You can specify the following options for the **hold-down** statement:

- **seconds**—Specify the time in seconds to wait before allowing the BMP station to reconnect to the device.
 - **flaps number**—Specify the number of BMP station flaps allowed before terminating the connection to the BMP station and triggering the hold down timer.
 - **period seconds**—Specify the time in seconds between BMP station flaps before terminating the connection to the BMP station and triggering the hold down timer.
6. (Optional) Specify an initiation message to be sent to the BMP station using the **initiation-message** statement. This statement allows you to provide some information to the BMP station system administrator (for example, a contact phone number).

initiation-message *text*;

7. Specify the connection mode for the connection between the BMP-enabled router and the BMP station using the **connection-mode** statement. The connection mode can be **active** or **passive**:

- **active**—BMP initiates the connection to the BMP station. If you configure active mode, you must also configure a station port using the **station-port** statement. However, you must not configure a local port (active mode).
- **passive**—BMP does not initiate a connection the BMP station. However, it does listen for a connection request from active BMP stations and will connect if a station is available. If you configure passive mode, you must not configure a station port. However, you must configure a local port using the **local-port** statement (passive mode).

connection-mode (active | passive);

8. Specify the port number for the BMP monitoring station by configuring the **station-port** statement. See also **connection-mode**.

station-port *port*;

9. Specify the listening port for the BMP station connection using the **local-port** statement. See also **connection-mode**.

If you change the local port, the BMP station connection flaps when you commit the configuration.

local-port

10. (Optional) Specify the IPv4 or IPv6 address for the BMP connection on the device using the **local-address** statement. For both active and passive connections, configure a loopback local address. This provides a consistent local endpoint, is useful for debugging, and assures greater reliability for the BMP connection since it is not tied to a single router interface.

For passive mode, specifying a local address is required. It also provides some security against a malicious BMP connection. For active mode, we also recommend configuring a local address to help ensure reliability.

If you change the local address, the BMP station connection flaps when you commit the configuration.

local-address *address*;

11. BMP monitoring is enabled by default. You can explicitly enable BMP monitoring or disable it. You can also selectively enable or disable BMP monitoring at various hierarchy levels (for example, `[edit protocols bgp group group-name]` or `[edit protocols bgp group group-name neighbor address]`). If you disable BMP monitoring, withdrawal messages are sent for any previously advertised routes. These are followed by a down message. If you enable BMP monitoring, an up message is sent first and then the route advertisements follow.

`monitor` (enable | disable);

12. Specify the dispatch priority for BMP by configuring the `priority` statement. The dispatch priority controls the frequency with which the device is able to forward BMP messages to BMP stations. You can configure the dispatch priority as either **high**, **medium**, or **low**.

`priority` (high | medium | low);

13. Specify whether BMP should send pre-policy route monitoring messages, post-policy route monitoring messages, both types of messages, or none at all. The pre-policy can be configured to exclude routes that are non-feasible for the decision process (for example, a route loop) by including the **non-feasible** option for the **pre-policy** statement. This represents the view of the BGP routes before running the import policy.

The post-policy can be configured to exclude routes that are not eligible for the decision process (for example, protocol nexthop not resolved) by including the **exclude-non-eligible** option for the **post-policy** statement. This represents the view of the BGP routes after running the import policy. If the import policy has rejected the BGP route, the route does not exist in the post policy view.

You can explicitly disable route monitoring by specifying the **none** option for the **route-monitoring** statement. This is the default behavior.

```
route-monitoring {
  none;
  post-policy {
    exclude-non-eligible;
  }
  pre-policy {
    exclude-non-feasible;
  }
}
```

14. Configure how often statistics messages are sent to the BMP monitoring station by specifying the number of seconds between message transmissions using **statistics-timeout** statement. If you configure a value of 0, no statistics messages are sent.

`statistics-timeout` *seconds*;

Related Documentation

- [Example: Configuring Route Authentication for BGP on page 3898](#)

BGP Path Attribute Configuration

- [Understanding the Advertisement of Multiple Paths to a Single Destination in BGP on page 3651](#)
- [Example: Configuring BGP Local Preference on page 3652](#)
- [Examples: Configuring BGP MED on page 3665](#)
- [Examples: Configuring BGP Local AS on page 3704](#)
- [Example: Configuring the Accumulated IGP Attribute for BGP on page 3724](#)

Understanding the Advertisement of Multiple Paths to a Single Destination in BGP

BGP peers advertise routes to each other in update messages. BGP stores its routes in the Junos OS routing table (**inet.0**). For each prefix in the routing table, the routing protocol process selects a single best path, called the active path. Unless you configure BGP to advertise multiple paths to the same destination, BGP advertises only the active path.

Instead of advertising only the active path to a destination, you can configure BGP to advertise multiple paths to the destination. Within an autonomous system (AS), the availability of multiple exit points to reach a destination provides the following benefits:

- **Fault tolerance**—Path diversity leads to reduction in restoration time after failure. For instance, a border after receiving multiple paths to the same destination can precompute a backup path and have it ready so that when the primary path becomes invalid, the border routing device can use the backup to quickly restore connectivity. Without a backup path, the restoration time depends on BGP reconvergence, which includes withdraw and advertisement messages in the network before a new best path can be learned.
- **Load balancing**—The availability of multiple paths to reach the same destination enables load balancing of traffic, if the routing within the AS meets certain constraints.
- **Maintenance**—The availability of alternate exit points allows for graceful maintenance operation of routers.

The following limitations apply to advertising multiple routes in BGP:

- Address families supported:
 - IPv4 unicast (**family inet unicast**)
 - IPv6 unicast (**family inet6 unicast**)

- IPv4 labeled unicast (**family inet labeled-unicast**)
- IPv6 labeled unicast (**family inet6 labeled-unicast**)
- Internal BGP (IBGP) peers only. No support on external BGP (EBGP) peers.
- Master instance only. No support for routing instances.
- Graceful restart and nonstop active routing (NSR) are supported.
- No BGP Monitoring Protocol (BMP) support.
- No support for EBGP sessions between confederations.
- Prefix policies enable you to filter routes on a router that is configured to advertise multiple paths to a destination. Prefix policies can only match prefixes. They cannot match route attributes, and they cannot change the attributes of routes.

**Related
Documentation**

- [Understanding BGP Path Selection on page 3793](#)
- [Example: Advertising Multiple Paths in BGP on page 3844](#)

Example: Configuring BGP Local Preference

- [Understanding the Local Preference Metric for Internal BGP Routes on page 3652](#)
- [Example: Configuring the Local Preference Value for BGP Routes on page 3653](#)

Understanding the Local Preference Metric for Internal BGP Routes

Internal BGP (IBGP) sessions use a metric called the *local preference*, which is carried in IBGP update packets in the path attribute LOCAL_PREF. When an autonomous system (AS) has multiple routes to another AS, the local preference indicates the degree of preference for one route over the other routes. The route with the highest local preference value is preferred.

The LOCAL_PREF path attribute is always advertised to IBGP peers and to neighboring confederations. It is never advertised to external BGP (EBGP) peers. The default behavior is to not modify the LOCAL_PREF path attribute if it is present.

The LOCAL_PREF path attribute applies at export time only, when the routes are exported from the routing table into BGP.

If a BGP route is received without a LOCAL_PREF attribute, the route is stored in the routing table and advertised by BGP as if it were received with a LOCAL_PREF value of 100. A non-BGP route that is advertised by BGP is advertised with a LOCAL_PREF value of 100 by default.

Example: Configuring the Local Preference Value for BGP Routes

This example shows how to configure local preference in internal BGP (IBGP) peer sessions.

- [Requirements on page 3653](#)
- [Overview on page 3653](#)
- [Configuration on page 3654](#)
- [Verification on page 3663](#)

Requirements

No special configuration beyond device initialization is required before you configure this example.


Overview

To change the local preference metric advertised in the path attribute, you must include the `local-preference` statement, specifying a value from 0 through 4,294,967,295 ($2^{32} - 1$).

There are several reasons you might want to prefer one path over another. For example, compared to other paths, one path might be less expensive to use, might have higher bandwidth, or might be more stable.

Figure 63 shows a typical network with internal peer sessions and multiple exit points to a neighboring AS.

Figure 63: Typical Network with IBGP Sessions and Multiple Exit Points



To reach Device R4, Device R1 can take a path through either Device R2 or Device R3. By default, the local preference is 100 for either route. When the local preferences are equal, Junos OS has rules for breaking the tie and choosing a path. (See [“Understanding BGP](#)

[Path Selection](#)" on page 3793.) In this example, the active route is through Device R2 because the router ID of Device R2 is lower than the router ID of Device R3. The following example shows how to override the default behavior with an explicit setting for the local preference. The example configures a local preference of 300 on Device R3, thereby making Device R3 the preferred path to reach Device R4.

Configuration

- [Configuring Device R1 on page 3655](#)
- [Configuring Device R2 on page 3657](#)
- [Configuring Device R3 on page 3659](#)
- [Configuring Device R4 on page 3662](#)

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device R1

```
set interfaces fe-1/2/0 unit 1 family inet address 12.12.12.1/24
set interfaces fe-1/2/1 unit 2 family inet address 13.13.13.1/24
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.1.1
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.2.1
set protocols bgp group internal neighbor 192.168.3.1
set protocols ospf area 0.0.0.0 interface lo0.1 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/0.1
set protocols ospf area 0.0.0.0 interface fe-1/2/1.2
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 123
set routing-options router-id 192.168.1.1
```

Device R2

```
set interfaces fe-1/2/0 unit 3 family inet address 12.12.12.2/24
set interfaces fe-1/2/1 unit 4 family inet address 24.24.24.2/24
set interfaces lo0 unit 2 family inet address 192.168.2.1/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.2.1
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.1.1
set protocols bgp group internal neighbor 192.168.3.1
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 4
set protocols bgp group external neighbor 24.24.24.4
set protocols ospf area 0.0.0.0 interface lo0.2 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/0.3
set protocols ospf area 0.0.0.0 interface fe-1/2/1.4
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 123
set routing-options router-id 192.168.2.1
```

Device R3

```

set interfaces fe-1/2/0 unit 5 family inet address 13.13.13.3/24
set interfaces fe-1/2/1 unit 6 family inet address 34.34.34.3/24
set interfaces lo0 unit 3 family inet address 192.168.3.1/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.3.1
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.1.1
set protocols bgp group internal neighbor 192.168.2.1
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 4
set protocols bgp group external neighbor 34.34.34.4
set protocols ospf area 0.0.0.0 interface lo0.3 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/0.5
set protocols ospf area 0.0.0.0 interface fe-1/2/1.6
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 123
set routing-options router-id 192.168.3.1

```

Device R4

```

set interfaces fe-1/2/0 unit 7 family inet address 24.24.24.4/24
set interfaces fe-1/2/1 unit 8 family inet address 34.34.34.4/24
set interfaces lo0 unit 4 family inet address 192.168.4.1/32
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 123
set protocols bgp group external neighbor 34.34.34.3
set protocols bgp group external neighbor 24.24.24.2
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 4
set routing-options router-id 192.168.4.1

```

Configuring Device R1

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R1:

1. Configure the interfaces.


```

[edit interfaces fe-1/2/0 unit 1]
user@R1# set family inet address 12.12.12.1/24

[edit interfaces fe-1/2/1 unit 2]
user@R1# set family inet address 13.13.13.1/24

[edit interfaces lo0 unit 1]
user@R1# set family inet address 192.168.1.1/32

```
2. Configure BGP.


```

[edit protocols bgp group internal]
user@R1# set type internal

```

```
user@R1# set local-address 192.168.1.1
user@R1# set export send-direct
user@R1# set neighbor 192.168.2.1
user@R1# set neighbor 192.168.3.1
```

3. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
user@R1# set interface lo0.1 passive
user@R1# set interface fe-1/2/0.1
user@R1# set interface fe-1/2/1.2
```

4. Configure a policy that accepts direct routes.



NOTE: Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 1]
user@R1# set from protocol direct
user@R1# set then accept
```

5. Configure the router ID and autonomous system (AS) number.

```
[edit routing-options]
user@R1# set autonomous-system 123
user@R1# set router-id 192.168.1.1
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show interfaces
fe-1/2/0 {
  unit 1 {
    family inet {
      address 12.12.12.1/24;
    }
  }
}
fe-1/2/1 {
  unit 2 {
    family inet {
      address 13.13.13.1/24;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 192.168.1.1/32;
    }
  }
}
```



```

user@R1# show policy-options
policy-statement send-direct {
  term 1 {
    from protocol direct;
    then accept;
  }
}

user@R1# show protocols
bgp {
  group internal {
    type internal;
    local-address 192.168.1.1;
    export send-direct;
    neighbor 192.168.2.1;
    neighbor 192.168.3.1;
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.1 {
      passive;
    }
    interface fe-1/2/0.1;
    interface fe-1/2/1.2;
  }
}

user@R1# show routing-options
autonomous-system 123;
router-id 192.168.1.1;

```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Device R2

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R2:

1. Configure the interfaces.


```

[edit interfaces fe-1/2/0 unit 3]
user@R2# set family inet address 12.12.12.21/24

[edit interfaces fe-1/2/1 unit 4]
user@R2# set family inet address 24.24.24.2/24

[edit interfaces lo0 unit 2]
user@R2# set family inet address 192.168.2.1/32

```
2. Configure BGP.


```

[edit protocols bgp group internal]
user@R2# set type internal

```

```
user@R2# set local-address 192.168.2.1
user@R2# set export send-direct
user@R2# set neighbor 192.168.1.1
user@R2# set neighbor 192.168.3.1
```

```
[edit protocols bgp group external]
user@R2# set type external
user@R2# set export send-direct
user@R2# set peer-as 4
user@R2# set neighbor 24.24.24.4
```

3. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
user@R2# set interface lo0.2 passive
user@R2# set interface fe-1/2/0.3
user@R2# set interface fe-1/2/1.4
```

4. Configure a policy that accepts direct routes.



NOTE: Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 1]
user@R2# set from protocol direct
user@R2# set then accept
```

5. Configure the router ID and autonomous system (AS) number.

```
[edit routing-options]
user@R2# set autonomous-system 123
user@R2# set router-id 192.168.2.1
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R2# show interfaces
fe-1/2/0 {
  unit 3 {
    family inet {
      address 12.12.12.2/24;
    }
  }
}
fe-1/2/1 {
  unit 4 {
    family inet {
      address 24.24.24.2/24;
    }
  }
}
lo0 {
```

```

    unit 2 {
        family inet {
            address 192.168.2.1/32;
        }
    }
}

user@R2# show policy-options
policy-statement send-direct {
    term 1 {
        from protocol direct;
        then accept;
    }
}

user@R2# show protocols
bgp {
    group internal {
        type internal;
        local-address 192.168.2.1;
        export send-direct;
        neighbor 192.168.1.1;
        neighbor 192.168.3.1;
    }
    group external {
        type external;
        export send-direct;
        peer-as 4;
        neighbor 24.24.24.4;
    }
}
ospf {
    area 0.0.0.0 {
        interface lo0.2 {
            passive;
        }
        interface fe-1/2/0.3;
        interface fe-1/2/1.4;
    }
}

user@R2# show routing-options
autonomous-system 123;
router-id 192.168.2.1;

```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Device R3

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R3:

1. Configure the interfaces.

```
[edit interfaces fe-1/2/0 unit 5]
user@R3# set family inet address 13.13.13.3/24
```

```
[edit interfaces fe-1/2/1 unit 6]
user@R3# set family inet address 34.34.34.3/24
```

```
[edit interfaces lo0 unit 3]
user@R3# set family inet address 192.168.3.1/32
```

2. Configure BGP.

```
[edit protocols bgp group internal]
user@R3# set type internal
user@R3# set local-address 192.168.3.1
user@R3# set export send-direct
user@R3# set neighbor 192.168.1.1
user@R3# set neighbor 192.168.2.1
```

```
[edit protocols bgp group external]
user@R3# set type external
user@R3# set export send-direct
user@R3# set peer-as 4
user@R3# set neighbor 34.34.34.4
```

3. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
user@R3# set interface lo0.3 passive
user@R3# set interface fe-1/2/0.5
user@R3# set interface fe-1/2/1.6
```

4. Configure a policy that accepts direct routes.



NOTE: Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 1]
user@R3# set from protocol direct
user@R3# set then accept
```

5. Configure the router ID and autonomous system (AS) number.

```
[edit routing-options]
user@R3# set autonomous-system 123
user@R3# set router-id 192.168.3.1
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R3# show interfaces
fe-1/2/0 {
  unit 5 {
```

```

        family inet {
            address 13.13.13.3/24;
        }
    }
}
fe-1/2/1 {
    unit 6 {
        family inet {
            address 34.34.34.3/24;
        }
    }
}
lo0 {
    unit 3 {
        family inet {
            address 192.168.3.1/32;
        }
    }
}

user@R3# show policy-options
policy-statement send-direct {
    term 1 {
        from protocol direct;
        then accept;
    }
}

user@R3# show protocols
bgp {
    group internal {
        type internal;
        local-address 192.168.3.1;
        export send-direct;
        neighbor 192.168.1.1;
        neighbor 192.168.2.1;
    }
    group external {
        type external;
        export send-direct;
        peer-as 4;
        neighbor 34.34.34.4;
    }
}
ospf {
    area 0.0.0.0 {
        interface lo0.3 {
            passive;
        }
        interface fe-1/2/0.5;
        interface fe-1/2/1.6;
    }
}

user@R3# show routing-options
autonomous-system 123;
router-id 192.168.3.1;

```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Device R4

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R4:

1. Configure the interfaces.

[edit interfaces fe-1/2/0 unit 7]
user@R4# set family inet address 24.24.24.4/24

[edit interfaces fe-1/2/1 unit 8]
user@R4# set family inet address 34.34.34.4/24

[edit interfaces lo0 unit 4]
user@R4# set family inet address 192.168.4.1/32
2. Configure BGP.

[edit protocols bgp group external]
user@R4# set type external
user@R4# set export send-direct
user@R4# set peer-as 123
user@R4# set neighbor 34.34.34.3
user@R4# set neighbor 24.24.24.2
3. Configure a policy that accepts direct routes.



NOTE: Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

- ```
[edit policy-options policy-statement send-direct term 1]
user@R4# set from protocol direct
user@R4# set then accept
```
4. Configure the router ID and autonomous system (AS) number.  
  
[edit routing-options]  
user@R4# set autonomous-system 4  
user@R4# set router-id 192.168.4.1

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R4# show interfaces
fe-1/2/0 {
 unit 7 {
```

```

 family inet {
 address 24.24.24.4/24;
 }
 }
}
fe-1/2/1 {
 unit 8 {
 family inet {
 address 34.34.34.4/24;
 }
 }
}
lo0 {
 unit 4 {
 family inet {
 address 192.168.4.1/32;
 }
 }
}

user@R4# show policy-options
policy-statement send-direct {
 term 1 {
 from protocol direct;
 then accept;
 }
}

user@R4# show protocols
bgp {
 group external {
 type external;
 export send-direct;
 peer-as 123;
 neighbor 34.34.34.3;
 neighbor 24.24.24.2;
 }
}

user@R4# show routing-options
autonomous-system 4;
router-id 192.168.4.1;

```

If you are done configuring the device, enter **commit** from configuration mode.

### Verification

Confirm that the configuration is working properly.

- [Checking the Active Path From Device R1 to Device R4 on page 3663](#)
- [Altering the Local Preference to Change the Path Selection on page 3664](#)
- [Rechecking the Active Path From Device R1 to Device R4 on page 3664](#)

#### *Checking the Active Path From Device R1 to Device R4*

**Purpose** Verify that the active path from Device R1 to Device R4 goes through Device R2.

**Action** From operational mode, enter the **show route protocol bgp** command.

```
user@R1> show route protocol bgp
inet.0: 11 destinations, 18 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

12.12.12.0/24 [BGP/170] 00:11:48, localpref 100, from 192.168.2.1
 AS path: I
 > to 12.12.12.2 via fe-1/2/0.1
13.13.13.0/24 [BGP/170] 00:11:48, localpref 100, from 192.168.3.1
 AS path: I
 > to 13.13.13.3 via fe-1/2/1.2
24.24.24.0/24 [BGP/170] 00:11:48, localpref 100, from 192.168.2.1
 AS path: I
 > to 12.12.12.2 via fe-1/2/0.1
34.34.34.0/24 [BGP/170] 00:11:48, localpref 100, from 192.168.3.1
 AS path: I
 > to 13.13.13.3 via fe-1/2/1.2
192.168.2.1/32 [BGP/170] 00:11:48, localpref 100, from 192.168.2.1
 AS path: I
 > to 12.12.12.2 via fe-1/2/0.1
192.168.3.1/32 [BGP/170] 00:11:48, localpref 100, from 192.168.3.1
 AS path: I
 > to 13.13.13.3 via fe-1/2/1.2
192.168.4.1/32 *[BGP/170] 00:05:14, localpref 100, from 192.168.2.1
 AS path: 4 I
 > to 12.12.12.2 via fe-1/2/0.1
 [BGP/170] 00:05:14, localpref 100, from 192.168.3.1
 AS path: 4 I
 > to 13.13.13.3 via fe-1/2/1.2
```

**Meaning** The asterisk (\*) shows that the preferred path is through Device R2. In the default configuration, Device R2 has a lower router ID than Device R3. The router ID is controlling the path selection.

#### *Altering the Local Preference to Change the Path Selection*

**Purpose** Change the path so that it goes through Device R3.

**Action** From configuration mode, enter the **set local-preference 300** command.

```
[edit protocols bgp group internal]
user@R3# set local-preference 300
user@R3# commit
```

#### *Rechecking the Active Path From Device R1 to Device R4*

**Purpose** Verify that the active path from Device R1 to Device R4 goes through Device R3.

**Action** From operational mode, enter the **show route protocol bgp** command.

```
user@R1> show route protocol bgp
inet.0: 11 destinations, 17 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

12.12.12.0/24 [BGP/170] 00:16:48, localpref 100, from 192.168.2.1
 AS path: I
 > to 12.12.12.2 via fe-1/2/0.1
13.13.13.0/24 [BGP/170] 00:00:22, localpref 300, from 192.168.3.1
```



```

 AS path: I
 > to 13.13.13.3 via fe-1/2/1.2
24.24.24.0/24 [BGP/170] 00:16:48, localpref 100, from 192.168.2.1
 AS path: I
 > to 12.12.12.2 via fe-1/2/0.1
34.34.34.0/24 [BGP/170] 00:00:22, localpref 300, from 192.168.3.1
 AS path: I
 > to 13.13.13.3 via fe-1/2/1.2
192.168.2.1/32 [BGP/170] 00:16:48, localpref 100, from 192.168.2.1
 AS path: I
 > to 12.12.12.2 via fe-1/2/0.1
192.168.3.1/32 [BGP/170] 00:00:22, localpref 300, from 192.168.3.1
 AS path: I
 > to 13.13.13.3 via fe-1/2/1.2
192.168.4.1/32 * [BGP/170] 00:00:21, localpref 300, from 192.168.3.1
 AS path: 4 I
 > to 13.13.13.3 via fe-1/2/1.2

```

**Meaning** The asterisk (\*) shows that the preferred path is through Device R3. In the altered configuration, Device R3 has a higher local preference than Device R2. The local preference is controlling the path selection.

**Related Documentation**

- [Examples: Configuring Internal BGP Peering on page 3624](#)
- [BGP Configuration Overview](#)

## Examples: Configuring BGP MED

- [Understanding the MED Attribute That Determines the Exit Point in an AS on page 3665](#)
- [Example: Configuring the MED Attribute That Determines the Exit Point in an AS on page 3668](#)
- [Example: Configuring the MED Using Route Filters on page 3680](#)
- [Example: Configuring the MED Using Communities on page 3693](#)
- [Example: Associating the MED Path Attribute with the IGP Metric and Delaying MED Updates on page 3694](#)

## Understanding the MED Attribute That Determines the Exit Point in an AS

The BGP multiple exit discriminator (MED, or MULTI\_EXIT\_DISC) is a non-transitive attribute, meaning that it is not propagated throughout the Internet, but only to adjacent autonomous systems (ASs). The MED attribute is optional, meaning that it is not always sent with the BGP updates. The purpose of MED is to influence how other ASs enter your AS to reach a certain prefix.

The MED attribute has a value that is referred to as a *metric*. If all other factors in determining an exit point are equal, the exit point with the lowest metric is preferred.

If a MED is received over an external BGP link, it is propagated over internal links to other BGP-enabled devices within the AS.

BGP update messages include a MED metric if the route was learned from BGP and already had a MED metric associated with it, or if you configure the MED metric in the configuration file.

A MED metric is advertised with a route according to the following general rules:

- A more specific metric overrides a less specific metric. That is, a group-specific metric overrides a global BGP metric, and a peer-specific metric overrides a global BGP or group-specific metric.
- A metric defined with a routing policy overrides a metric defined with the **metric-out** statement.
- If any metric is defined, it overrides a metric received in a route.
- If the received route does not have an associated MED metric, and if you do not explicitly configure a metric value, no metric is advertised. When you do not explicitly configure a metric value, the MED value is equivalent to zero (0) when advertising an active route.

Because the AS path rather than the number of hops between hosts is the primary criterion for BGP route selection, an AS with multiple connections to a peer AS can have multiple equivalent AS paths. When the routing table contains two routes to the same host in a neighboring AS, a MED metric assigned to each route can determine which to include in the forwarding table. The MED metric you assign can force traffic through a particular exit point in an AS.

Figure 64 illustrates how MED metrics are used to determine route selection.

**Figure 64: Default MED Example**




Figure 64 shows AS 1 and AS 2 connected by two separate BGP links to Routers C and D. Host E in AS 1 is located nearer to Router C. Host F, also in AS 1, is located nearer to Router D. Because the AS paths are equivalent, two routes exist for each host, one through Router C and one through Router D. To force all traffic destined for Host E through Router C, the network administrator for AS 1 assigns a MED metric for each router to Host E at its exit point. A MED metric of 10 is assigned to the route to Host E through

Router C, and a MED metric of 20 is assigned to the route to Host E through Router D. BGP routers in AS 2 select the route with the lower MED metric for the forwarding table.

By default, only the MEDs of routes that have the same peer ASs are compared. However, you can configure the routing table path selection options listed in [Table 317](#) to compare MEDs in different ways. The MED options are not mutually exclusive and can be configured in combination or independently. For the MED options to take effect, you must configure them uniformly all through your network. The MED option or options you configure determine the route selected. Thus we recommend that you carefully evaluate your network for preferred routes before configuring the MED options.

**Table 317: MED Options for Routing Table Path Selection**

| Option (Name)                                                                   | Function                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Use                                                                                                                                                                                                                                            |
|---------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Always comparing MEDs ( <b>always-compare-med</b> )                             | Ensures that the MEDs for paths from peers in different ASs are always compared in the route selection process.                                                                                                                                                                                                                                                                                                                                                                                                                      | Useful when all enterprises participating in a network agree on a uniform policy for setting MEDs. For example, in a network shared by two ISPs, both must agree that a certain path is the better path to configure the MED values correctly. |
| Adding IGP cost to MED ( <b>med-plus-igp</b> )                                  | <p>Before comparing MED values for path selection, adds to the MED the cost of the IGP route to the BGP next-hop destination.</p> <p>This option replaces the MED value for the router, but does not affect the IGP metric comparison. As a result, when multiple routes have the same value after the MED-plus-IGP comparison, and route selection continues, the IGP route metric is also compared, even though it was added to the MED value and compared earlier in the selection process.</p>                                   | Useful when the downstream AS requires the complete cost of a certain route that is received across multiple ASs.                                                                                                                              |
| Applying Cisco IOS nondeterministic behavior ( <b>cisco-non-deterministic</b> ) | <p>Specifies the nondeterministic behavior of the Cisco IOS software:</p> <ul style="list-style-type: none"> <li>The active path is always first. All nonactive but eligible paths follow the active path and are maintained in the order in which they were received. Ineligible paths remain at the end of the list.</li> <li>When a new path is added to the routing table, path comparisons are made among all routes, including those paths that must never be selected because they lose the MED tie-breaking rule.</li> </ul> | We recommend that you do not configure this option, because the nondeterministic behavior sometimes prevents the system from properly comparing the MEDs between paths.                                                                        |

## Example: Configuring the MED Attribute That Determines the Exit Point in an AS

This example shows how to configure a multiple exit discriminator (MED) metric to advertise in BGP update messages.

- [Requirements on page 3668](#)
- [Overview on page 3668](#)
- [Configuration on page 3669](#)
- [Verification on page 3679](#)

---

### Requirements

No special configuration beyond device initialization is required before you configure this example.

---

### Overview

To directly configure a MED metric to advertise in BGP update messages, include the **metric-out** statement:

**metric-out** (*metric* | *minimum-igp offset* | *igp delay-med-update* | *offset*);

**metric** is the primary metric on all routes sent to peers. It can be a value in the range from 0 through 4,294,967,295 ( $2^{32} - 1$ ).


The following optional settings are also supported:

- **minimum-igp**—Sets the metric to the minimum metric value calculated in the interior gateway protocol (IGP) to get to the BGP next hop. If a newly calculated metric is greater than the minimum metric value, the metric value remains unchanged. If a newly calculated metric is lower, the metric value is lowered to that value.
- **igp**—Sets the metric to the most recent metric value calculated in the IGP to get to the BGP next hop.
- **delay-med-update**—Delays sending MED updates when the MED value increases. Include the **delay-med-update** statement when you configure the **igp** statement. The default interval to delay sending updates, unless the MED is lower or another attribute associated with the route has changed is 10 minutes. Include the **med-igp-update-interval minutes** statement at the **[edit routing-options]** hierarchy level to modify the default interval.
- **offset**—Specifies a value for **offset** to increase or decrease the metric that is used from the metric value calculated in the IGP. The metric value is offset by the value specified. The metric calculated in the IGP (by specifying either **igp** or **igp-minimum**) is increased if the **offset** value is positive. The metric calculated in the IGP (by specifying either **igp** or **igp-minimum**) is decreased if the **offset** value is negative.

**offset** can be a value in the range from  $-2^{31}$  through  $2^{31} - 1$ . Note that the adjusted metric can never go below 0 or above  $2^{32} - 1$ .

[Figure 65](#) shows a typical network with internal peer sessions and multiple exit points to a neighboring autonomous system (AS).

Figure 65: Typical Network with IBGP Sessions and Multiple Exit Points



Device R4 has multiple loopback interfaces configured to simulate advertised prefixes. The extra loopback interface addresses are 44.44.44.44/32 and 144.144.144.144/32. This example shows how to configure Device R4 to advertise a MED value of 30 to Device R3 and a MED value of 20 to Device R2. This causes all of the devices in AS 123 to prefer the path through Device R2 to reach AS 4.

### Configuration

- [Configuring Device R1 on page 3671](#)
- [Configuring Device R2 on page 3673](#)
- [Configuring Device R3 on page 3675](#)
- [Configuring Device R4 on page 3677](#)

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

#### Device R1

```
set interfaces fe-1/2/0 unit 1 family inet address 12.12.12.1/24
set interfaces fe-1/2/1 unit 2 family inet address 13.13.13.1/24
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.1.1
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.2.1
set protocols bgp group internal neighbor 192.168.3.1
set protocols ospf area 0.0.0.0 interface lo0.1 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/0.1
set protocols ospf area 0.0.0.0 interface fe-1/2/1.2
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
```

```
set routing-options autonomous-system 123
set routing-options router-id 192.168.1.1
```

**Device R2**

```
set interfaces fe-1/2/0 unit 3 family inet address 12.12.12.2/24
set interfaces fe-1/2/1 unit 4 family inet address 24.24.24.2/24
set interfaces lo0 unit 2 family inet address 192.168.2.1/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.2.1
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.1.1
set protocols bgp group internal neighbor 192.168.3.1
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 4
set protocols bgp group external neighbor 24.24.24.4
set protocols ospf area 0.0.0.0 interface lo0.2 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/0.3
set protocols ospf area 0.0.0.0 interface fe-1/2/1.4
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 123
set routing-options router-id 192.168.2.1
```

**Device R3**

```
set interfaces fe-1/2/0 unit 5 family inet address 13.13.13.3/24
set interfaces fe-1/2/1 unit 6 family inet address 34.34.34.3/24
set interfaces lo0 unit 3 family inet address 192.168.3.1/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.3.1
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.1.1
set protocols bgp group internal neighbor 192.168.2.1
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 4
set protocols bgp group external neighbor 34.34.34.4
set protocols ospf area 0.0.0.0 interface lo0.3 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/0.5
set protocols ospf area 0.0.0.0 interface fe-1/2/1.6
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 123
set routing-options router-id 192.168.3.1
```

**Device R4**

```
set interfaces fe-1/2/0 unit 7 family inet address 24.24.24.4/24
set interfaces fe-1/2/1 unit 8 family inet address 34.34.34.4/24
set interfaces lo0 unit 4 family inet address 192.168.4.1/32
set interfaces lo0 unit 4 family inet address 44.44.44.44/32
set interfaces lo0 unit 4 family inet address 144.144.144.144/32
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 123
set protocols bgp group external neighbor 34.34.34.3 metric-out 30
set protocols bgp group external neighbor 24.24.24.2 metric-out 20
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
```

```
set routing-options autonomous-system 4
set routing-options router-id 192.168.4.1
```

### Configuring Device R1

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R1:

1. Configure the interfaces.

```
[edit interfaces fe-1/2/0 unit 1]
user@R1# set family inet address 12.12.12.1/24
```

```
[edit interfaces fe-1/2/1 unit 2]
user@R1# set family inet address 13.13.13.1/24
```

```
[edit interfaces lo0 unit 1]
user@R1# set family inet address 192.168.1.1/32
```

2. Configure BGP.

```
[edit protocols bgp group internal]
user@R1# set type internal
user@R1# set local-address 192.168.1.1
user@R1# set export send-direct
user@R1# set neighbor 192.168.2.1
user@R1# set neighbor 192.168.3.1
```

3. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
user@R1# set interface lo0.1 passive
user@R1# set interface fe-1/2/0.1
user@R1# set interface fe-1/2/1.2
```

4. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 1]
user@R1# set from protocol direct
user@R1# set then accept
```

5. Configure the router ID and autonomous system (AS) number.

```
[edit routing-options]
user@R1# set autonomous-system 123
user@R1# set router-id 192.168.1.1
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show interfaces
fe-1/2/0 {
 unit 1 {
 family inet {
 address 12.12.12.1/24;
 }
 }
}
fe-1/2/1 {
 unit 2 {
 family inet {
 address 13.13.13.1/24;
 }
 }
}
lo0 {
 unit 1 {
 family inet {
 address 192.168.1.1/32;
 }
 }
}

user@R1# show policy-options
policy-statement send-direct {
 term 1 {
 from protocol direct;
 then accept;
 }
}

user@R1# show protocols
bgp {
 group internal {
 type internal;
 local-address 192.168.1.1;
 export send-direct;
 neighbor 192.168.2.1;
 neighbor 192.168.3.1;
 }
}
ospf {
 area 0.0.0.0 {
 interface lo0.1 {
 passive;
 }
 interface fe-1/2/0.1;
 interface fe-1/2/1.2;
 }
}

user@R1# show routing-options
autonomous-system 123;
router-id 192.168.1.1;
```

If you are done configuring the device, enter **commit** from configuration mode.



**Configuring Device R2**

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R2:

1. Configure the interfaces.

```
[edit interfaces fe-1/2/0 unit 3]
user@R2# set family inet address 12.12.12.21/24
```

```
[edit interfaces fe-1/2/1 unit 4]
user@R2# set family inet address 24.24.24.2/24
```

```
[edit interfaces lo0 unit 2]
user@R2# set family inet address 192.168.2.1/32
```

2. Configure BGP.

```
[edit protocols bgp group internal]
user@R2# set type internal
user@R2# set local-address 192.168.2.1
user@R2# set export send-direct
user@R2# set neighbor 192.168.1.1
user@R2# set neighbor 192.168.3.1
```

```
[edit protocols bgp group external]
user@R2# set type external
user@R2# set export send-direct
user@R2# set peer-as 4
user@R2# set neighbor 24.24.24.4
```

3. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
user@R2# set interface lo0.2 passive
user@R2# set interface fe-1/2/0.3
user@R2# set interface fe-1/2/1.4
```

4. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 1]
user@R2# set from protocol direct
user@R2# set then accept
```

5. Configure the router ID and autonomous system (AS) number.

```
[edit routing-options]
user@R2# set autonomous-system 123
user@R2# set router-id 192.168.2.1
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R2# show interfaces
fe-1/2/0 {
 unit 3 {
 family inet {
 address 12.12.12.2/24;
 }
 }
}
fe-1/2/1 {
 unit 4 {
 family inet {
 address 24.24.24.2/24;
 }
 }
}
lo0 {
 unit 2 {
 family inet {
 address 192.168.2.1/32;
 }
 }
}

user@R2# show policy-options
policy-statement send-direct {
 term 1 {
 from protocol direct;
 then accept;
 }
}

user@R2# show protocols
bgp {
 group internal {
 type internal;
 local-address 192.168.2.1;
 export send-direct;
 neighbor 192.168.1.1;
 neighbor 192.168.3.1;
 }
 group external {
 type external;
 export send-direct;
 peer-as 4;
 neighbor 24.24.24.4;
 }
}
ospf {
 area 0.0.0.0 {
 interface lo0.2 {
 passive;
 }
 }
}
```

```

 }
 interface fe-1/2/0.3;
 interface fe-1/2/1.4;
 }
}

user@R2# show routing-options
autonomous-system 123;
router-id 192.168.2.1;

```

If you are done configuring the device, enter **commit** from configuration mode.

### Configuring Device R3

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R3:

1. Configure the interfaces.

```

[edit interfaces fe-1/2/0 unit 5]
user@R3# set family inet address 13.13.13.3/24

```

```

[edit interfaces fe-1/2/1 unit 6]
user@R3# set family inet address 34.34.34.3/24

```

```

[edit interfaces lo0 unit 3]
user@R3# set family inet address 192.168.3.1/32

```

2. Configure BGP.

```

[edit protocols bgp group internal]
user@R3# set type internal
user@R3# set local-address 192.168.3.1
user@R3# set export send-direct
user@R3# set neighbor 192.168.1.1
user@R3# set neighbor 192.168.2.1

```

```

[edit protocols bgp group external]
user@R3# set type external
user@R3# set export send-direct
user@R3# set peer-as 4
user@R3# set neighbor 34.34.34.4

```

3. Configure OSPF.

```

[edit protocols ospf area 0.0.0.0]
user@R3# set interface lo0.3 passive
user@R3# set interface fe-1/2/0.5
user@R3# set interface fe-1/2/1.6

```

4. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 1]
user@R3# set from protocol direct
user@R3# set then accept
```

5. Configure the router ID and autonomous system (AS) number.

```
[edit routing-options]
user@R3# set autonomous-system 123
user@R3# set router-id 192.168.3.1
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R3# show interfaces
fe-1/2/0 {
 unit 5 {
 family inet {
 address 13.13.13.3/24;
 }
 }
}
fe-1/2/1 {
 unit 6 {
 family inet {
 address 34.34.34.3/24;
 }
 }
}
lo0 {
 unit 3 {
 family inet {
 address 192.168.3.1/32;
 }
 }
}

user@R3# show policy-options
policy-statement send-direct {
 term 1 {
 from protocol direct;
 then accept;
 }
}

user@R3# show protocols
bgp {
 group internal {
 type internal;
 local-address 192.168.3.1;
 export send-direct;
 neighbor 192.168.1.1;
 neighbor 192.168.2.1;
 }
 group external {
```

```

 type external;
 export send-direct;
 peer-as 4;
 neighbor 34.34.34.4;
 }
}
ospf {
 area 0.0.0.0 {
 interface lo0.3 {
 passive;
 }
 interface fe-1/2/0.5;
 interface fe-1/2/1.6;
 }
}

user@R3# show routing-options
autonomous-system 123;
router-id 192.168.3.1;

```

If you are done configuring the device, enter **commit** from configuration mode.

### Configuring Device R4

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R4:

1. Configure the interfaces.

```

[edit interfaces fe-1/2/0 unit 7]
user@R4# set family inet address 24.24.24.4/24

```

```

[edit interfaces fe-1/2/1 unit 8]
user@R4# set family inet address 34.34.34.4/24

```

```

[edit interfaces lo0 unit 4]
user@R4# set family inet address 192.168.4.1/32
user@R4# set family inet address 44.44.44.44/32
user@R4# set family inet address 144.144.144.144/32

```

Device R4 has multiple loopback interface addresses to simulate advertised prefixes.

2. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```

[edit policy-options policy-statement send-direct term 1]
user@R4# set from protocol direct
user@R4# set then accept

```

3. Configure BGP.

```

[edit protocols bgp group external]
user@R4# set type external

```

```
user@R4# set export send-direct
user@R4# set peer-as 123
```

4. Configure a MED value of 30 for neighbor Device R3, and a MED value of 20 for neighbor Device R2.

```
[edit protocols bgp group external]
user@R4# set neighbor 34.34.34.3 metric-out 30
user@R4# set neighbor 24.24.24.2 metric-out 20
```

This configuration causes autonomous system (AS) 123 (of which Device R1, Device R2, and Device R3 are members) to prefer the path through Device R2 to reach AS 4.

5. Configure the router ID and AS number.

```
[edit routing-options]
user@R4# set autonomous-system 4
user@R4# set router-id 192.168.4.1
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R4# show interfaces
fe-1/2/0 {
 unit 7 {
 family inet {
 address 24.24.24.4/24;
 }
 }
}
fe-1/2/1 {
 unit 8 {
 family inet {
 address 34.34.34.4/24;
 }
 }
}
lo0 {
 unit 4 {
 family inet {
 address 192.168.4.1/32;
 address 44.44.44.44/32;
 address 144.144.144.144/32;
 }
 }
}

user@R4# show policy-options
policy-statement send-direct {
 term 1 {
 from protocol direct;
 then accept;
 }
}
```

```

user@R4# show protocols
bgp {
 group external {
 type external;
 export send-direct;
 peer-as 123;
 neighbor 34.34.34.3 {
 metric-out 30;
 }
 neighbor 24.24.24.2 {
 metric-out 20;
 }
 }
}

user@R4# show routing-options
autonomous-system 4;
router-id 192.168.4.1;

```

If you are done configuring the device, enter **commit** from configuration mode.

### Verification

Confirm that the configuration is working properly.

- [Checking the Active Path From Device R1 to Device R4 on page 3679](#)
- [Verifying That Device R4 Is Sending Its Routes Correctly on page 3680](#)

#### *Checking the Active Path From Device R1 to Device R4*

**Purpose** Verify that the active path goes through Device R2.

**Action** From operational mode, enter the **show route protocol bgp** command.

```

user@R1> show route protocol bgp
inet.0: 13 destinations, 19 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

12.12.12.0/24 [BGP/170] 3d 22:52:38, localpref 100, from 192.168.2.1
 AS path: I
 > to 12.12.12.2 via fe-1/2/0.1
13.13.13.0/24 [BGP/170] 3d 03:15:16, localpref 100, from 192.168.3.1
 AS path: I
 > to 13.13.13.3 via fe-1/2/1.2
24.24.24.0/24 [BGP/170] 3d 22:52:38, localpref 100, from 192.168.2.1
 AS path: I
 > to 12.12.12.2 via fe-1/2/0.1
34.34.34.0/24 [BGP/170] 3d 03:15:16, localpref 100, from 192.168.3.1
 AS path: I
 > to 13.13.13.3 via fe-1/2/1.2
44.44.44.44/32 *[BGP/170] 01:41:11, MED 20, localpref 100, from 192.168.2.1
 AS path: 4 I
 > to 12.12.12.2 via fe-1/2/0.1
144.144.144.144/32 *[BGP/170] 00:08:13, MED 20, localpref 100, from 192.168.2.1
 AS path: 4 I
 > to 12.12.12.2 via fe-1/2/0.1
192.168.2.1/32 [BGP/170] 3d 22:52:38, localpref 100, from 192.168.2.1
 AS path: I

```

```

> to 12.12.12.2 via fe-1/2/0.1
192.168.3.1/32 [BGP/170] 3d 03:15:16, localpref 100, from 192.168.3.1
AS path: I
> to 13.13.13.3 via fe-1/2/1.2
192.168.4.1/32 * [BGP/170] 01:41:11, MED 20, localpref 100, from 192.168.2.1
AS path: 4 I
> to 12.12.12.2 via fe-1/2/0.1

```

**Meaning** The asterisk (\*) shows that the preferred path is through Device R2. The reason for the path selection is listed as MED 20.

#### *Verifying That Device R4 Is Sending Its Routes Correctly*

**Purpose** Make sure that Device R4 is sending update messages with a value of 20 to Device R2 and a value of 30 to Device R3.

**Action** From operational mode, enter the **show route advertising-protocol bgp 24.24.24.2** command.

```

user@R4> show route advertising-protocol bgp 24.24.24.2
inet.0: 11 destinations, 13 routes (11 active, 0 holddown, 0 hidden)
 Prefix Nexthop MED Lclpref AS path
* 24.24.24.0/24 Self 20 I
* 34.34.34.0/24 Self 20 I
* 44.44.44.44/32 Self 20 I
* 144.144.144.144/32 Self 20 I
* 192.168.4.1/32 Self 20 I

user@R4> show route advertising-protocol bgp 34.34.34.3
inet.0: 11 destinations, 13 routes (11 active, 0 holddown, 0 hidden)
 Prefix Nexthop MED Lclpref AS path
* 24.24.24.0/24 Self 30 I
* 34.34.34.0/24 Self 30 I
* 44.44.44.44/32 Self 30 I
* 144.144.144.144/32 Self 30 I
* 192.168.4.1/32 Self 30 I

```

**Meaning** The MED column shows that Device R4 is sending the correct MED values to its two external BGP (EBGP) neighbors.

### Example: Configuring the MED Using Route Filters

This example shows how to configure a policy that uses route filters to modify the multiple exit discriminator (MED) metric to advertise in BGP update messages.

- [Requirements on page 3680](#)
- [Overview on page 3681](#)
- [Configuration on page 3681](#)
- [Verification on page 3692](#)

#### Requirements

No special configuration beyond device initialization is required before you configure this example.




## Overview

To configure a route-filter policy that modifies the advertised MED metric in BGP update messages, include the **metric** statement in the policy action.

Figure 66 shows a typical network with internal peer sessions and multiple exit points to a neighboring autonomous system (AS).

Figure 66: Typical Network with IBGP Sessions and Multiple Exit Points



Device R4 has multiple loopback interfaces configured to simulate advertised prefixes. The extra loopback interface addresses are 44.44.44.44/32 and 144.144.144.144/32. This example shows how to configure Device R4 to advertise a MED value of 30 to Device R3 for all routes except 144.144.144.144. For 144.144.144.144, a MED value of 10 is advertised to Device 3. A MED value of 20 is advertised to Device R2, regardless of the route prefix.

## Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Device R1**

```

set interfaces fe-1/2/0 unit 1 family inet address 12.12.12.1/24
set interfaces fe-1/2/1 unit 2 family inet address 13.13.13.1/24
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.1.1
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.2.1
set protocols bgp group internal neighbor 192.168.3.1
set protocols ospf area 0.0.0.0 interface lo0.1 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/0.1

```

```
set protocols ospf area 0.0.0.0 interface fe-1/2/1.2
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 123
set routing-options router-id 192.168.1.1
```

Device R2

```
set interfaces fe-1/2/0 unit 3 family inet address 12.12.12.2/24
set interfaces fe-1/2/1 unit 4 family inet address 24.24.24.2/24
set interfaces lo0 unit 2 family inet address 192.168.2.1/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.2.1
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.1.1
set protocols bgp group internal neighbor 192.168.3.1
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 4
set protocols bgp group external neighbor 24.24.24.4
set protocols ospf area 0.0.0.0 interface lo0.2 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/0.3
set protocols ospf area 0.0.0.0 interface fe-1/2/1.4
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 123
set routing-options router-id 192.168.2.1
```

Device R3

```
set interfaces fe-1/2/0 unit 5 family inet address 13.13.13.3/24
set interfaces fe-1/2/1 unit 6 family inet address 34.34.34.3/24
set interfaces lo0 unit 3 family inet address 192.168.3.1/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.3.1
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.1.1
set protocols bgp group internal neighbor 192.168.2.1
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 4
set protocols bgp group external neighbor 34.34.34.4
set protocols ospf area 0.0.0.0 interface lo0.3 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/0.5
set protocols ospf area 0.0.0.0 interface fe-1/2/1.6
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 123
set routing-options router-id 192.168.3.1
```

Device R4

```
set interfaces fe-1/2/0 unit 7 family inet address 24.24.24.4/24
set interfaces fe-1/2/1 unit 8 family inet address 34.34.34.4/24
set interfaces lo0 unit 4 family inet address 192.168.4.1/32
set interfaces lo0 unit 4 family inet address 44.44.44.44/32
set interfaces lo0 unit 4 family inet address 144.144.144.144/32
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 123
set protocols bgp group external neighbor 34.34.34.3 export med-10
```

```

set protocols bgp group external neighbor 34.34.34.3 export med-30
set protocols bgp group external neighbor 24.24.24.2 metric-out 20
set policy-options policy-statement med-10 from route-filter 144.144.144.144/32 exact
set policy-options policy-statement med-10 then metric 10
set policy-options policy-statement med-10 then accept
set policy-options policy-statement med-30 from route-filter 0.0.0.0/0 longer
set policy-options policy-statement med-30 then metric 30
set policy-options policy-statement med-30 then accept
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 4
set routing-options router-id 192.168.4.1

```

### Configuring Device R1

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R1:

1. Configure the device interfaces.

```

[edit interfaces fe-1/2/0 unit 1]
user@R1# set family inet address 12.12.12.1/24

```

```

[edit interfaces fe-1/2/1 unit 2]
user@R1# set family inet address 13.13.13.1/24

```

```

[edit interfaces lo0 unit 1]
user@R1# set family inet address 192.168.1.1/32

```

2. Configure BGP.

```

[edit protocols bgp group internal]
user@R1# set type internal
user@R1# set local-address 192.168.1.1
user@R1# set export send-direct
user@R1# set neighbor 192.168.2.1
user@R1# set neighbor 192.168.3.1

```

3. Configure OSPF.

```

[edit protocols ospf area 0.0.0.0]
user@R1# set interface lo0.1 passive
user@R1# set interface fe-1/2/0.1
user@R1# set interface fe-1/2/1.2

```

4. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```

[edit policy-options policy-statement send-direct term 1]
user@R1# set from protocol direct
user@R1# set then accept

```

5. Configure the router ID and autonomous system (AS) number.

```
[edit routing-options]
user@R1# set autonomous-system 123
user@R1# set router-id 192.168.1.1
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show interfaces
fe-1/2/0 {
 unit 1 {
 family inet {
 address 12.12.12.1/24;
 }
 }
}
fe-1/2/1 {
 unit 2 {
 family inet {
 address 13.13.13.1/24;
 }
 }
}
lo0 {
 unit 1 {
 family inet {
 address 192.168.1.1/32;
 }
 }
}

user@R1# show protocols
bgp {
 group internal {
 type internal;
 local-address 192.168.1.1;
 export send-direct;
 neighbor 192.168.2.1;
 neighbor 192.168.3.1;
 }
}
ospf {
 area 0.0.0.0 {
 interface lo0.1 {
 passive;
 }
 interface fe-1/2/0.1;
 interface fe-1/2/1.2;
 }
}

user@R1# show policy-options
policy-statement send-direct {
 term 1 {
 from protocol direct;
```

```

 then accept;
 }
}

user@R1# show routing-options
autonomous-system 123;
router-id 192.168.1.1;

```

If you are done configuring the device, enter **commit** from configuration mode.

### Configuring Device R2

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R2:

1. Configure the device interfaces.

```

[edit interfaces fe-1/2/0 unit 3]
user@R2# set family inet address 12.12.12.21/24

```

```

[edit interfaces fe-1/2/1 unit 4]
user@R2# set family inet address 24.24.24.2/24

```

```

[edit interfaces lo0 unit 2]
user@R2# set family inet address 192.168.2.1/32

```

2. Configure BGP.

```

[edit protocols bgp group internal]
user@R2# set type internal
user@R2# set local-address 192.168.2.1
user@R2# set export send-direct
user@R2# set neighbor 192.168.1.1
user@R2# set neighbor 192.168.3.1

```

```

[edit protocols bgp group external]
user@R2# set type external
user@R2# set export send-direct
user@R2# set peer-as 4
user@R2# set neighbor 24.24.24.4

```

3. Configure OSPF.

```

[edit protocols ospf area 0.0.0.0]
user@R2# set interface lo0.2 passive
user@R2# set interface fe-1/2/0.3
user@R2# set interface fe-1/2/1.4

```

4. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```

[edit policy-options policy-statement send-direct term 1]
user@R2# set from protocol direct

```

```
user@R2# set then accept
```

5. Configure the router ID and autonomous system (AS) number.

```
[edit routing-options]
user@R2# set autonomous-system 123
user@R2# set router-id 192.168.2.1
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R2# show interfaces
fe-1/2/0 {
 unit 3 {
 family inet {
 address 12.12.12.2/24;
 }
 }
}
fe-1/2/1 {
 unit 4 {
 family inet {
 address 24.24.24.2/24;
 }
 }
}
lo0 {
 unit 2 {
 family inet {
 address 192.168.2.1/32;
 }
 }
}
```

```
user@R2# show protocols
bgp {
 group internal {
 type internal;
 local-address 192.168.2.1;
 export send-direct;
 neighbor 192.168.1.1;
 neighbor 192.168.3.1;
 }
 group external {
 type external;
 export send-direct;
 peer-as 4;
 neighbor 24.24.24.4;
 }
}
ospf {
 area 0.0.0.0 {
 interface lo0.2 {
 passive;
 }
 }
}
```

```

 }
 interface fe-1/2/0.3;
 interface fe-1/2/1.4;
 }
}

user@R2# show policy-options
policy-statement send-direct {
 term 1 {
 from protocol direct;
 then accept;
 }
}

user@R2# show routing-options
autonomous-system 123;
router-id 192.168.2.1;

```

If you are done configuring the device, enter **commit** from configuration mode.

### Configuring Device R3

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R3:

1. Configure the device interfaces.
 

```

[edit interfaces fe-1/2/0 unit 5]
user@R3# set family inet address 13.13.13.3/24

[edit interfaces fe-1/2/1 unit 6]
user@R3# set family inet address 34.34.34.3/24

[edit interfaces lo0 unit 3]
user@R3# set family inet address 192.168.3.1/32

```
2. Configure BGP.
 

```

[edit protocols bgp group internal]
user@R3# set type internal
user@R3# set local-address 192.168.3.1
user@R3# set export send-direct
user@R3# set neighbor 192.168.1.1
user@R3# set neighbor 192.168.2.1

[edit protocols bgp group external]
user@R3# set type external
user@R3# set export send-direct
user@R3# set peer-as 4
user@R3# set neighbor 34.34.34.4

```
3. Configure OSPF.
 

```

[edit protocols ospf area 0.0.0.0]

```

```
user@R3# set interface lo0.3 passive
user@R3# set interface fe-1/2/0.5
user@R3# set interface fe-1/2/1.6
```

4. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 1]
user@R3# set from protocol direct
user@R3# set then accept
```

5. Configure the router ID and autonomous system (AS) number.

```
[edit routing-options]
user@R3# set autonomous-system 123
user@R3# set router-id 192.168.3.1
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R3# show interfaces
fe-1/2/0 {
 unit 5 {
 family inet {
 address 13.13.13.3/24;
 }
 }
}
fe-1/2/1 {
 unit 6 {
 family inet {
 address 34.34.34.3/24;
 }
 }
}
lo0 {
 unit 3 {
 family inet {
 address 192.168.3.1/32;
 }
 }
}
```

```
user@R3# show protocols
bgp {
 group internal {
 type internal;
 local-address 192.168.3.1;
 export send-direct;
 neighbor 192.168.1.1;
 neighbor 192.168.2.1;
 }
 group external {
```



```

 type external;
 export send-direct;
 peer-as 4;
 neighbor 34.34.34.4;
 }
}
ospf {
 area 0.0.0.0 {
 interface lo0.3 {
 passive;
 }
 interface fe-1/2/0.5;
 interface fe-1/2/1.6;
 }
}

user@R3# show policy-options
policy-statement send-direct {
 term 1 {
 from protocol direct;
 then accept;
 }
}

user@R3# show routing-options
autonomous-system 123;
router-id 192.168.3.1;

```

If you are done configuring the device, enter **commit** from configuration mode.

### Configuring Device R4

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R4:

1. Configure the device interfaces.

```

[edit interfaces fe-1/2/0 unit 7]
user@R4# set family inet address 24.24.24.4/24

```

```

[edit interfaces fe-1/2/1 unit 8]
user@R4# set family inet address 34.34.34.4/24

```

```

[edit interfaces lo0 unit 4]
user@R4# set family inet address 192.168.4.1/32
user@R4# set family inet address 44.44.44.44/32
user@R4# set family inet address 144.144.144.144/32

```

Device R4 has multiple loopback interface addresses to simulate advertised prefixes.

2. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 1]
user@R4# set from protocol direct
user@R4# set then accept
```

3. Configure BGP.

```
[edit protocols bgp group external]
user@R4# set type external
user@R4# set export send-direct
user@R4# set peer-as 123
```

4. Configure the two MED policies.

```
[edit policy-options]
set policy-statement med-10 from route-filter 144.144.144.144/32 exact
set policy-statement med-10 then metric 10
set policy-statement med-10 then accept
```

```
set policy-statement med-30 from route-filter 0.0.0.0/0 longer
set policy-statement med-30 then metric 30
set policy-statement med-30 then accept
```

5. Configure the two EBGP neighbors, applying the two MED policies to Device R3, and a MED value of 20 to Device R2.

```
[edit protocols bgp group external]
user@R4# set neighbor 34.34.34.3 export med-10
user@R4# set neighbor 34.34.34.3 export med-30
user@R4# set neighbor 24.24.24.2 metric-out 20
```

6. Configure the router ID and autonomous system (AS) number.

```
[edit routing-options]
user@R4# set autonomous-system 4
user@R4# set router-id 192.168.4.1
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R4# show interfaces
fe-1/2/0 {
 unit 7 {
 family inet {
 address 24.24.24.4/24;
 }
 }
}
fe-1/2/1 {
 unit 8 {
 family inet {
 address 34.34.34.4/24;
 }
 }
}
lo0 {
```

```

unit 4 {
 family inet {
 address 192.168.4.1/32;
 address 44.44.44.44/32;
 address 144.144.144.144/32;
 }
}

user@R4# show protocols
bgp {
 group external {
 type external;
 export send-direct;
 peer-as 123;
 neighbor 24.24.24.2 {
 metric-out 20;
 }
 neighbor 34.34.34.3 {
 export [med-10 med-30];
 }
 }
}

user@R4# show policy-options
policy-statement med-10 {
 from {
 route-filter 144.144.144.144/32 exact;
 }
 then {
 metric 10;
 accept;
 }
}
policy-statement med-30 {
 from {
 route-filter 0.0.0.0/0 longer;
 }
 then {
 metric 30;
 accept;
 }
}
policy-statement send-direct {
 term 1 {
 from protocol direct;
 then accept;
 }
}

user@R4# show routing-options
autonomous-system 4;
router-id 192.168.4.1;

```

If you are done configuring the device, enter **commit** from configuration mode.

## Verification

Confirm that the configuration is working properly.

- [Checking the Active Path from Device R1 to Device R4 on page 3692](#)
- [Verifying That Device R4 Is Sending Its Routes Correctly on page 3692](#)

### *Checking the Active Path from Device R1 to Device R4*

**Purpose** Verify that the active path goes through Device R2.

**Action** From operational mode, enter the **show route protocol bgp** command.

```
user@R1> show route protocol bgp
inet.0: 13 destinations, 19 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

12.12.12.0/24 [BGP/170] 4d 01:13:32, localpref 100, from 192.168.2.1
 AS path: I
 > to 12.12.12.2 via fe-1/2/0.1
13.13.13.0/24 [BGP/170] 3d 05:36:10, localpref 100, from 192.168.3.1
 AS path: I
 > to 13.13.13.3 via fe-1/2/1.2
24.24.24.0/24 [BGP/170] 4d 01:13:32, localpref 100, from 192.168.2.1
 AS path: I
 > to 12.12.12.2 via fe-1/2/0.1
34.34.34.0/24 [BGP/170] 3d 05:36:10, localpref 100, from 192.168.3.1
 AS path: I
 > to 13.13.13.3 via fe-1/2/1.2
44.44.44.44/32 *[BGP/170] 00:06:03, MED 20, localpref 100, from 192.168.2.1
 AS path: 4 I
 > to 12.12.12.2 via fe-1/2/0.1
144.144.144.144/32 *[BGP/170] 00:06:03, MED 10, localpref 100, from 192.168.3.1
 AS path: 4 I
 > to 13.13.13.3 via fe-1/2/1.2
192.168.2.1/32 [BGP/170] 4d 01:13:32, localpref 100, from 192.168.2.1
 AS path: I
 > to 12.12.12.2 via fe-1/2/0.1
192.168.3.1/32 [BGP/170] 3d 05:36:10, localpref 100, from 192.168.3.1
 AS path: I
 > to 13.13.13.3 via fe-1/2/1.2
192.168.4.1/32 *[BGP/170] 00:06:03, MED 20, localpref 100, from 192.168.2.1
 AS path: 4 I
 > to 12.12.12.2 via fe-1/2/0.1
```

**Meaning** The output shows that the preferred path to the routes advertised by Device R4 is through Device R2 for all routes except 144.144.144.144/32. For 144.144.144.144/32, the preferred path is through Device R3.

### *Verifying That Device R4 Is Sending Its Routes Correctly*

**Purpose** Make sure that Device R4 is sending update messages with a value of 20 to Device R2 and a value of 30 to Device R3.

**Action** From operational mode, enter the **show route advertising-protocol bgp** command.

```
user@R4> show route advertising-protocol bgp 24.24.24.2
```

```
inet.0: 11 destinations, 13 routes (11 active, 0 holddown, 0 hidden)
 Prefix Nexthop MED Lclpref AS path
* 24.24.24.0/24 Self 20 I
* 34.34.34.0/24 Self 20 I
* 44.44.44.44/32 Self 20 I
* 144.144.144.144/32 Self 20 I
* 192.168.4.1/32 Self 20 I
```

```
user@R4> show route advertising-protocol bgp 34.34.34.3
inet.0: 11 destinations, 13 routes (11 active, 0 holddown, 0 hidden)
 Prefix Nexthop MED Lclpref AS path
* 24.24.24.0/24 Self 30 I
* 34.34.34.0/24 Self 30 I
* 44.44.44.44/32 Self 30 I
* 144.144.144.144/32 Self 10 I
* 192.168.4.1/32 Self 30 I
```

**Meaning** The MED column shows that Device R4 is sending the correct MED values to its two EBGp neighbors.

### Example: Configuring the MED Using Communities

Set the multiple exit discriminator (MED) metric to 20 for all routes from a particular community.

```
[edit]
routing-options {
 router-id 10.0.0.1;
 autonomous-system 23;
}
policy-options {
 policy-statement from-otago {
 from community otago;
 then metric 20;
 }
 community otago members [56:2379 23:46944];
}
protocols {
 bgp {
 import from-otago;
 group 23 {
 type external;
 peer-as 56;
 neighbor 192.168.0.1 {
 traceoptions {
 file bgp-log-peer;
 flag packets;
 }
 log-updown;
 }
 }
 }
}
```

## Example: Associating the MED Path Attribute with the IGP Metric and Delaying MED Updates

This example shows how to associate the multiple exit discriminator (MED) path attribute with the interior gateway protocol (IGP) metric, and configure a timer to delay update of the MED attribute.

- [Requirements on page 3694](#)
- [Overview on page 3694](#)
- [Configuration on page 3696](#)
- [Verification on page 3702](#)

---

### Requirements

No special configuration beyond device initialization is required before you configure this example.

---

### Overview

BGP can be configured to advertise the MED attribute for a route based on the IGP distance of its internal BGP (IBGP) route next-hop. The IGP metric enables internal routing to follow the shortest path according to the administrative setup. In some deployments, it might be ideal to communicate IGP shortest-path knowledge to external BGP (EBGP) peers in a neighboring autonomous system (AS). This allows those EBGP peers to forward traffic into your AS using the shortest paths possible.

Routes learned from an EBGP peer usually have a next hop on a directly connected interface, and thus the IGP value is equal to zero. Zero is the value advertised. The IGP metric is a nonzero value when a BGP peer sends third-party next hops that require the local system to perform next-hop resolution—IBGP configurations, configurations within confederation peers, or EBGP configurations that include the **multihop** command. In these scenarios, it might make sense to associate the MED value with the IGP metric by including the **metric-out minimum-igp** or **metric-out igp** option.

The drawback of associating the MED with the IGP metric is the risk of excessive route advertisements when there are IGP instabilities in the network. Configuring a delay for the MED update provides a mechanism to reduce route advertisements in such scenarios. The delay works by slowing down MED updates when the IGP metric for the next hop changes. The approach uses a timer to periodically advertise MED updates. When the timer expires, the MED attribute for routes with **metric-out igp delay-updates** configured is updated to the current IGP metric of the next hop. The BGP-enabled device sends out advertisements for routes for which the MED attribute has changed.

The **delay-updates** option identifies the BGP groups (or peers) for which the MED updates must be suppressed. The time for advertising MED updates is set to 10 minutes by default. You can increase the interval up to 600 minutes by including the **med-igp-update-interval** statement in the **routing-options** configuration.



**NOTE:** If you have nonstop active routing (NSR) enabled and a switchover occurs, the delayed MED updates might be advertised as soon as the switchover occurs.


When you configure the **metric-out igp** option, the IGP metric directly tracks the IGP cost to the IBGP peer. When the IGP cost goes down, so does the advertised MED value. Conversely, when the IGP cost goes up, the MED value goes up as well.

When you configure the **metric-out minimum-igp** option, the advertised MED value changes only when the IGP cost to the IBGP peer goes down. An increase in the IGP cost does not affect the MED value. The router monitors and remembers the lowest IGP cost until the routing process (rpd) is restarted. The BGP peer sends an update only if the MED is lower than the previously advertised value or another attribute associated with the route has changed, or if the BGP peer is responding to a refresh route request.

This example uses the **metric** statement in the OSPF configuration to demonstrate that when the IGP metric changes, the MED also changes after the configured delay interval. The OSPF metric can range from 1 through 65,535.

Figure 67 shows the sample topology.

Figure 67: Topology for Delaying the MED Update



In this example, the MED value advertised by Device R1 is associated with the IGP running in AS 1. The MED value advertised by Device R1 impacts the decisions of the neighboring AS (AS 2) when AS 2 is forwarding traffic into AS 1.

### Configuration

- [Configuring Device R1 on page 3700](#)

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Device R1**

```

set interfaces fe-1/2/0 unit 2 description R1->R2
set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.1/30
set interfaces fe-1/2/1 unit 7 description R1->R4
set interfaces fe-1/2/1 unit 7 family inet address 172.16.0.1/30
set interfaces lo0 unit 1 family inet address 192.168.0.1/32

```



```

set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.0.1
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.0.2
set protocols bgp group internal neighbor 192.168.0.3
set protocols bgp group external type external
set protocols bgp group external metric-out igp delay-med-update
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 2
set protocols bgp group external neighbor 172.16.0.2
set protocols ospf area 0.0.0.0 interface fe-1/2/0.2 metric 600
set protocols ospf area 0.0.0.0 interface lo0.1 passive
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options med-igp-update-interval 12
set routing-options router-id 192.168.0.1
set routing-options autonomous-system 1

```

Device R2

```

set interfaces fe-1/2/0 unit 1 description R2->R1
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 4 description R2->R3
set interfaces fe-1/2/1 unit 4 family inet address 10.0.2.2/30
set interfaces lo0 unit 2 family inet address 192.168.0.2/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.0.2
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.0.1
set protocols bgp group internal neighbor 192.168.0.3
set protocols ospf area 0.0.0.0 interface fe-1/2/0.1
set protocols ospf area 0.0.0.0 interface fe-1/2/1.4
set protocols ospf area 0.0.0.0 interface lo0.2 passive
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options router-id 192.168.0.2
set routing-options autonomous-system 1

```

Device R3

```

set interfaces fe-1/2/0 unit 3 description R3->R2
set interfaces fe-1/2/0 unit 3 family inet address 10.0.2.1/30
set interfaces fe-1/2/1 unit 5 description R3->R5
set interfaces fe-1/2/1 unit 5 family inet address 172.16.0.5/30
set interfaces lo0 unit 3 family inet address 192.168.0.3/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.0.3
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.0.1
set protocols bgp group internal neighbor 192.168.0.2
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 2
set protocols bgp group external neighbor 172.16.0.6
set protocols ospf area 0.0.0.0 interface fe-1/2/0.3
set protocols ospf area 0.0.0.0 interface lo0.3 passive
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options router-id 192.168.0.3

```

```
set routing-options autonomous-system 1
```

```
Device R4 set interfaces fe-1/2/0 unit 8 description R4->R1
 set interfaces fe-1/2/0 unit 8 family inet address 172.16.0.2/30
 set interfaces fe-1/2/1 unit 9 description R4->R5
 set interfaces fe-1/2/1 unit 9 family inet address 10.0.4.1/30
 set interfaces fe-1/2/2 unit 13 description R4->R6
 set interfaces fe-1/2/2 unit 13 family inet address 172.16.0.9/30
 set interfaces lo0 unit 4 family inet address 192.168.0.4/32
 set protocols bgp group internal type internal
 set protocols bgp group internal local-address 192.168.0.4
 set protocols bgp group internal export send-direct
 set protocols bgp group internal neighbor 192.168.0.5
 set protocols bgp group external type external
 set protocols bgp group external export send-direct
 set protocols bgp group external neighbor 172.16.0.10 peer-as 3
 set protocols bgp group external neighbor 172.16.0.1 peer-as 1
 set protocols ospf area 0.0.0.0 interface fe-1/2/1.9
 set protocols ospf area 0.0.0.0 interface lo0.4 passive
 set policy-options policy-statement send-direct term 1 from protocol direct
 set policy-options policy-statement send-direct term 1 then accept
 set routing-options router-id 192.168.0.4
 set routing-options autonomous-system 2
```

```
Device R5 set interfaces fe-1/2/0 unit 6 description R5->R3
 set interfaces fe-1/2/0 unit 6 family inet address 172.16.0.6/30
 set interfaces fe-1/2/1 unit 10 description R5->R4
 set interfaces fe-1/2/1 unit 10 family inet address 10.0.4.2/30
 set interfaces fe-1/2/2 unit 11 description R5->R8
 set interfaces fe-1/2/2 unit 11 family inet address 172.16.0.13/30
 set interfaces lo0 unit 5 family inet address 192.168.0.5/32
 set protocols bgp group internal type internal
 set protocols bgp group internal local-address 192.168.0.5
 set protocols bgp group internal export send-direct
 set protocols bgp group internal neighbor 192.168.0.4
 set protocols bgp group external type external
 set protocols bgp group external export send-direct
 set protocols bgp group external neighbor 172.16.0.5 peer-as 1
 set protocols bgp group external neighbor 172.16.0.14 peer-as 3
 set protocols ospf area 0.0.0.0 interface fe-1/2/1.10
 set protocols ospf area 0.0.0.0 interface lo0.5 passive
 set policy-options policy-statement send-direct term 1 from protocol direct
 set policy-options policy-statement send-direct term 1 then accept
 set routing-options router-id 192.168.0.5
 set routing-options autonomous-system 2
```

```
Device R6 set interfaces fe-1/2/0 unit 14 description R6->R4
 set interfaces fe-1/2/0 unit 14 family inet address 172.16.0.10/30
 set interfaces fe-1/2/1 unit 15 description R6->R7
 set interfaces fe-1/2/1 unit 15 family inet address 10.0.6.1/30
 set interfaces lo0 unit 6 family inet address 192.168.0.6/32
 set protocols bgp group internal type internal
 set protocols bgp group internal local-address 192.168.0.6
 set protocols bgp group internal export send-direct
 set protocols bgp group internal neighbor 192.168.0.7
```

```

set protocols bgp group internal neighbor 192.168.0.8
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 2
set protocols bgp group external neighbor 172.16.0.9 peer-as 2
set protocols ospf area 0.0.0.0 interface fe-1/2/1.15
set protocols ospf area 0.0.0.0 interface lo0.6 passive
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options router-id 192.168.0.6
set routing-options autonomous-system 3

```

**Device R7**

```

set interfaces fe-1/2/0 unit 16 description R7->R6
set interfaces fe-1/2/0 unit 16 family inet address 10.0.6.2/30
set interfaces fe-1/2/1 unit 17 description R7->R8
set interfaces fe-1/2/1 unit 17 family inet address 10.0.7.2/30
set interfaces lo0 unit 7 family inet address 192.168.0.7/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.0.7
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.0.6
set protocols bgp group internal neighbor 192.168.0.8
set protocols ospf area 0.0.0.0 interface fe-1/2/0.16
set protocols ospf area 0.0.0.0 interface fe-1/2/1.17
set protocols ospf area 0.0.0.0 interface lo0.7 passive
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options router-id 192.168.0.7
set routing-options autonomous-system 3

```

**Device R8**

```

set interfaces fe-1/2/0 unit 12 description R8->R5
set interfaces fe-1/2/0 unit 12 family inet address 172.16.0.14/30
set interfaces fe-1/2/1 unit 18 description R8->R7
set interfaces fe-1/2/1 unit 18 family inet address 10.0.7.1/30
set interfaces lo0 unit 8 family inet address 192.168.0.8/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.0.8
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.0.6
set protocols bgp group internal neighbor 192.168.0.7
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 2
set protocols bgp group external neighbor 172.16.0.13 peer-as 2
set protocols ospf area 0.0.0.0 interface fe-1/2/1.18
set protocols ospf area 0.0.0.0 interface lo0.8 passive
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options router-id 192.168.0.8
set routing-options autonomous-system 3

```

### Configuring Device R1

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R1:

1. Configure the interfaces.

```
[edit interfaces fe-1/2/0 unit 2]
user@R1# set description R1->R2
user@R1# set family inet address 10.0.0.1/30
```

```
[edit interfaces fe-1/2/1 unit 7]
user@R1# set description R1->R4
user@R1# set family inet address 172.16.0.1/30
```

```
[edit interfaces lo0 unit 1]
user@R1# set family inet address 192.168.0.1/32
```

2. Configure IBGP.

```
[edit protocols bgp group internal]
user@R1# set type internal
user@R1# set local-address 192.168.0.1
user@R1# set export send-direct
user@R1# set neighbor 192.168.0.2
user@R1# set neighbor 192.168.0.3
```

3. Configure EBGp.

```
[edit protocols bgp group external]
user@R1# set type external
user@R1# set export send-direct
user@R1# set peer-as 2
user@R1# set neighbor 172.16.0.2
```

4. Associate the MED value with the IGP metric.

```
[edit protocols bgp group external]
user@R1# set metric-out igp delay-med-update
```

The default for the MED update is 10 minutes when you include the **delay-med-update** option. When you exclude the **delay-med-update** option, the MED update occurs immediately after the IGP metric changes.

5. (Optional) Configure the update interval for the MED update.

```
[edit routing-options]
user@R1# set med-igp-update-interval 12
```

You can configure the interval from 10 minutes through 600 minutes.

6. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
user@R1# set interface fe-1/2/0.2 metric 600
user@R1# set interface lo0.1 passive
```

The **metric** statement is used here to demonstrate what happens when the IGP metric changes.

7. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 1]
user@R1# set from protocol direct
user@R1# set then accept
```

8. Configure the router ID and autonomous system (AS) number.

```
[edit routing-options]
user@R1# set router-id 192.168.0.1
user@R1# set autonomous-system 1
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show interfaces
fe-1/2/0 {
 unit 2 {
 description R1->R2;
 family inet {
 address 10.0.0.1/30;
 }
 }
}
fe-1/2/1 {
 unit 7 {
 description R1->R4;
 family inet {
 address 172.16.0.1/30;
 }
 }
}
lo0 {
 unit 1 {
 family inet {
 address 192.168.0.1/32;
 }
 }
}

user@R1# show policy-options
policy-statement send-direct {
 term 1 {
 from protocol direct;
 then accept;
 }
}

user@R1# show protocols
```

```
bgp {
 group internal {
 type internal;
 local-address 192.168.0.1;
 export send-direct;
 neighbor 192.168.0.2;
 neighbor 192.168.0.3;
 }
 group external {
 type external;
 metric-out igp delay-med-update;
 export send-direct;
 peer-as 2;
 neighbor 172.16.0.2;
 }
}
ospf {
 area 0.0.0.0 {
 interface fe-1/2/0.2 {
 metric 600;
 }
 interface lo0.1 {
 passive;
 }
 }
}
```

```
user@R1# show routing-options
med-igp-update-interval 12;
router-id 192.168.0.1;
autonomous-system 1;
```

If you are done configuring the device, enter **commit** from configuration mode. Repeat the configuration steps on the other devices in the topology, as needed for your network.

---

## Verification

Confirm that the configuration is working properly.

- [Checking the BGP Advertisements on page 3702](#)
- [Verifying That the MED Value Changes When the OSPF Metric Changes on page 3703](#)
- [Testing the minimum-igp Setting on page 3703](#)

### *Checking the BGP Advertisements*

**Purpose** Verify that Device R1 is advertising to Device R4 a BGP MED value that reflects the IGP metric.

**Action** From operational mode, enter the **show route advertising-protocol bgp 172.16.0.2** command.

```
user@R1> show route advertising-protocol bgp 172.16.0.2
inet.0: 19 destinations, 33 routes (19 active, 0 holddown, 0 hidden)
 Prefix Nexthop MED Lc1pref AS path
* 10.0.0.0/30 Self 0 I I
* 172.16.0.0/30 Self 0 I I
```

|                  |      |     |   |
|------------------|------|-----|---|
| * 172.16.0.4/30  | Self | 601 | I |
| * 192.168.0.1/32 | Self | 0   | I |

**Meaning** The 601 value in the MED column shows that the MED value has been updated to reflect the configured OSPF metric.

#### *Verifying That the MED Value Changes When the OSPF Metric Changes*

**Purpose** Make sure that when you raise the OSPF metric to 700, the MED value is updated to reflect this change.

**Action** From configuration mode, enter the **set protocols ospf area 0 interface fe-1/2/0.2 metric 700** command.

```
user@R1# set protocols ospf area 0 interface fe-1/2/0.2 metric 700
user@R1# commit
```

After waiting 12 minutes (the configured delay period), enter the **show route advertising-protocol bgp** command from operational mode.

```
user@R1> show route advertising-protocol bgp 172.16.0.2
inet.0: 19 destinations, 33 routes (19 active, 0 holddown, 0 hidden)
 Prefix Nexthop MED Lclpref AS path
* 10.0.0.0/30 Self 0 I
* 172.16.0.0/30 Self 0 I
* 172.16.0.4/30 Self 701 I
* 192.168.0.1/32 Self 0 I
```

**Meaning** The 701 value in the MED column shows that the MED value has been updated to reflect the configured OSPF metric.

#### *Testing the minimum-igp Setting*

**Purpose** Change the configuration to use the **minimum-igp** statement instead of the **igp** statement. When you increase the OSPF metric, the MED value remains unchanged, but when you decrease the OSPF metric, the MED value reflects the new OSPF metric.

**Action** From configuration mode, delete the **igp** statement, add the **minimum-igp** statement, and increase the OSPF metric.

```
user@R1# delete protocols bgp group external metric-out igp
user@R1# set protocols bgp group external metric-out minimum-igp
user@R1# set protocols ospf area 0 interface fe-1/2/0.2 metric 800
user@R1# commit
```

From operational mode, enter the **show route advertising-protocol bgp** command to make sure that the MED value does not change.

```
user@R1> show route advertising-protocol bgp 172.16.0.2
inet.0: 19 destinations, 33 routes (19 active, 0 holddown, 0 hidden)
 Prefix Nexthop MED Lclpref AS path
* 10.0.0.0/30 Self 0 I
* 172.16.0.0/30 Self 0 I
* 172.16.0.4/30 Self 701 I
* 192.168.0.1/32 Self 0 I
```

From configuration mode, decrease the OSPF metric.

```
user@R1# set protocols ospf area 0 interface fe-1/2/0.2 metric 20
user@R1# commit
```

From operational mode, enter the **show route advertising-protocol bgp** command to make sure that the MED value does change.

```
user@R1> show route advertising-protocol bgp 172.16.0.2
inet.0: 19 destinations, 33 routes (19 active, 0 holddown, 0 hidden)
 Prefix Nexthop MED Lc1pref AS path
* 10.0.0.0/30 Self 0 I
* 172.16.0.0/30 Self 0 I
* 172.16.0.4/30 Self 21 I
* 192.168.0.1/32 Self 0 I
```

**Meaning** When the **minimum-igp** statement is configured, the MED value changes only when a shorter path is available.

**Related Documentation**

- [Examples: Configuring External BGP Peering on page 3601](#)
- [BGP Configuration Overview](#)

## Examples: Configuring BGP Local AS

- [Understanding the BGP Local AS Attribute on page 3704](#)
- [Example: Configuring a Local AS for EBGp Sessions on page 3709](#)
- [Example: Configuring a Private Local AS for EBGp Sessions on page 3719](#)

### Understanding the BGP Local AS Attribute

When an Internet service provider (ISP) acquires a network that belongs to a different autonomous system (AS), there is no seamless method for moving the BGP peers of the acquired network to the AS of the acquiring ISP. The process of configuring the BGP peers with the new AS number can be time-consuming and cumbersome. Sometimes customers do not want to or are not immediately able to modify their peer arrangements or configuration. During this kind of transition period, it can be useful to configure BGP-enabled devices in the new AS to use the former AS number in BGP updates. This former AS number is called a *local AS*.

Using a local AS number permits the routing devices in an acquired network to appear to belong to the former AS.

For example, ISP A, with an AS of 200, acquires ISP B, with an AS of 250. ISP B has a customer, ISP C, that does not want to change its configuration. After ISP B becomes part of ISP A, a local AS number of 250 is configured for use in EBGp peer sessions with ISP C. Consequently, the local AS number of 250 is either prepended before or used instead of the global AS number of 200 in the AS path used to export routes to direct external peers in ISP C.

If the route is received from an internal BGP (IBGP) peer, the AS path includes the local AS number prepended before the global AS number.



The local AS number is used instead of the global AS number if the route is an external route, such as a static route or an interior gateway protocol (IGP) route that is imported into BGP. If the route is external and you want the global AS number to be included in the AS path, you can apply a routing policy that uses **as-path-expand** or **as-path-prepend**. Use the **as-path-expand** policy action to place the global AS number behind the local AS number. Use the **as-path-prepend** policy action to place the global AS number in front of the local AS number.

For example:

```

user@R2# show policy-options
policy-statement prepend-global {
 term 1 {
 from protocol static;
 then {
 as-path-prepend 200; # or use as-path-expand
 accept;
 }
 }
}

user@R2# show protocols bgp
group ext {
 export prepend-global;
 type external;
 local-as 250;
 neighbor 10.0.0.1 {
 peer-as 100;
 }
 neighbor 10.1.0.2 {
 peer-as 300;
 }
}

user@R2# show routing-options
static {
 route 1.1.1.1/32 next-hop 10.0.0.1;
}
autonomous-system 200;

user@R3# run show route 1.1.1.1 protocol bgp
inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.1/32 *[BGP/170] 00:05:11, localpref 100
 AS path: 200 250 I, validation-state: unverified
 > to 10.1.0.1 via 1t-1/2/0.4

```

In a Layer 3 VPN scenario, in which a provider edge (PE) device uses external BGP (EBGP) to peer with a customer edge (CE) device, the **local-as** statement behaves differently than in the non-VPN scenario. In the VPN scenario, the global AS number defined in the master instance is prepended to the AS path by default. To override this behavior, you can configure the **no-prepend-global-as** in the routing-instance BGP configuration on the PE device, as shown here:

```

user@R2# show routing-instances

```

```
red {
 instance-type vrf;
 interface fe-1/2/0.2;
 route-distinguisher 2:1;
 vrf-target target:2:1;
 protocols {
 bgp {
 group toR1 {
 type external;
 peer-as 1;
 local-as 200 no-prepend-global-as;
 neighbor 10.1.1.1;
 }
 }
 }
}
```

The Junos operating system (Junos OS) implementation of the local AS attribute supports the following options:


- **Local AS with private option**—When you use the **private** option, the local AS is used during the establishment of the BGP session with an EBGP neighbor but is hidden in the AS path sent to other EBGP peers. Only the global AS is included in the AS path sent to external peers.

The **private** option is useful for establishing local peering with routing devices that remain configured with their former AS or with a specific customer that has not yet modified its peer arrangements. The local AS is used to establish the BGP session with the EBGP neighbor but is hidden in the AS path sent to external peers in another AS.

Include the **private** option so that the local AS is not prepended before the global AS in the AS path sent to external peers. When you specify the **private** option, the local AS is prepended only in the AS path sent to the EBGP neighbor.

For example, in [Figure 68](#), Router 1 and Router 2 are in AS 64496, Router 4 is in AS 64511, and Router 3 is in AS 64510. Router 2 formerly belonged to AS 64497, which has merged with another network and now belongs to AS 64496. Because Router 3 still peers with Router 2 using its former AS (64497), Router 2 needs to be configured with a local AS of 64497 in order to maintain peering with Router 3. Configuring a local AS of 64497 permits Router 2 to add AS 64497 when advertising routes to Router 3. Router 3 sees an AS path of 64497 64496 for the prefix 10/8.

**Figure 68: Local AS Configuration**



To prevent Router 2 from adding the local AS number in its announcements to other peers, use the **local-as 64497 private** statement. This statement configures Router 2 to not include local AS 64497 when announcing routes to Router 1 and to Router 4. In this case, Router 4 sees an AS path of 64496 64510 for the prefix 10.222/16.

- **Local AS with alias option**—In Junos OS Release 9.5 and later, you can configure a local AS as an alias. During the establishment of the BGP open session, the AS used in the open message alternates between the local AS and the global AS. If the local AS is used to connect with the EBGP neighbor, then only the local AS is prepended to the AS path when the BGP peer session is established. If the global AS is used to connect with the EBGP neighbor, then only the global AS is prepended to the AS path when the BGP peer session is established. The use of the **alias** option also means that

the local AS is not prepended to the AS path for any routes learned from that EBGp neighbor. Therefore, the local AS remains hidden from other external peers.

Configuring a local AS with the **alias** option is especially useful when you are migrating the routing devices in an acquired network to the new AS. During the migration process, some routing devices might be configured with the new AS while others remain configured with the former AS. For example, it is good practice to start by first migrating to the new AS any routing devices that function as route reflectors. However, as you migrate the route reflector clients incrementally, each route reflector has to peer with routing devices configured with the former AS, as well as peer with routing devices configured with the new AS. To establish local peer sessions, it can be useful for the BGP peers in the network to use both the local AS and the global AS. At the same time, you want to hide this local AS from external peers and use only the global AS in the AS path when exporting routes to another AS. In this kind of situation, configure the **alias** option.

Include the **alias** option to configure the local AS as an alias to the global AS configured at the **[edit routing-options]** hierarchy level. When you configure a local AS as an alias, during the establishment of the BGP open session, the AS used in the open message alternates between the local AS and the global AS. The local AS is prepended to the AS path only when the peer session with an EBGp neighbor is established using that local AS. The local AS is hidden in the AS path sent to any other external peers. Only the global AS is prepended to the AS path when the BGP session is established using the global AS.



**NOTE:** The **private** and **alias** options are mutually exclusive. You cannot configure both options with the same **local-as** statement.

---

- **Local AS with option not to prepend the global AS**—In Junos OS Release 9.6 and later, you can configure a local AS with the option not to prepend the global AS. Only the local AS is included in the AS path sent to external peers.

Use the **no-prepend-global-as** option when you want to strip the global AS number from outbound BGP updates in a virtual private network (VPN) scenario. This option is useful in a VPN scenario in which you want to hide the global AS from the VPN.

Include the **no-prepend-global-as** option to have the global AS configured at the **[edit routing-options]** hierarchy level removed from the AS path sent to external peers. When you use this option, only the local AS is included in the AS path for the routes sent to a customer edge (CE) device.

- **Number of loops option**—The local AS feature also supports specifying the number of times that detection of the AS number in the AS\_PATH attribute causes the route to be discarded or hidden. For example, if you configure **loops 1**, the route is hidden if the AS number is detected in the path one or more times. This is the default behavior. If you configure **loops 2**, the route is hidden if the AS number is detected in the path two or more times.

For the **loops number** statement, you can configure 1 through 10.



**NOTE:** If you configure the local AS values for any BGP group, the detection of routing loops is performed using both the AS and the local AS values for all BGP groups.

If the local AS for the EBGP or IBGP peer is the same as the current AS, do not use the `local-as` statement to specify the local AS number.

When you configure the local AS within a VRF, this impacts the AS path loop-detection mechanism. All of the `local-as` statements configured on the device are part of a single AS domain. The AS path loop-detection mechanism is based on looking for a matching AS present in the domain.

## Example: Configuring a Local AS for EBGP Sessions

This example shows how to configure a local autonomous system (AS) for a BGP peer so that both the global AS and the local AS are used in BGP inbound and outbound updates.

- [Requirements on page 3709](#)
- [Overview on page 3709](#)
- [Configuration on page 3710](#)
- [Verification on page 3716](#)

### Requirements

No special configuration beyond device initialization is required before you configure this example.

### Overview

Use the **local-as** statement when ISPs merge and want to preserve a customer's configuration, particularly the AS with which the customer is configured to establish a peer relationship. The **local-as** statement simulates the AS number already in place in customer routers, even if the ISP's router has moved to a different AS.


This example shows how to use the **local-as** statement to configure a local AS. The **local-as** statement is supported for BGP at the global, group, and neighbor hierarchy levels.

When you configure the **local-as** statement, you must specify an AS number. You can specify a number from 1 through 4,294,967,295 in plain-number format. In Junos OS Release 9.1 and later, the range for AS numbers is extended to provide BGP support for 4-byte AS numbers as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*. In Junos OS Release 9.3 and later, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: *<16-bit high-order value in decimal>.<16-bit low-order value in decimal>*. For example, the 4-byte AS number of 65,546 in plain-number format is represented as 1.10 in the AS-dot notation format. You can specify a value from 0.0 through 65535.65535 in AS-dot notation format. Junos

OS continues to support 2-byte AS numbers. The 2-byte AS number range is 1 through 65,535 (this is a subset of the 4-byte range).

Figure 69 shows the sample topology.

**Figure 69: Topology for Configuring the Local AS**



In this example, Device R2 formerly belonged to AS 250 and now is in AS 200. Device R1 and Device R3 are configured to peer with AS 250 instead of with the new AS number (AS 200). Device R2 has the new AS number configured with the **autonomous-system 200** statement. To enable the peering sessions to work, the **local-as 250** statement is added in the BGP configuration. Because **local-as 250** is configured, Device R2 includes both the global AS (200) and the local AS (250) in its BGP inbound and outbound updates.

### Configuration

- [Configuring Device R1 on page 3711](#)
- [Configuring Device R2 on page 3713](#)
- [Configuring Device R3 on page 3715](#)

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
Device R1 set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
 set interfaces lo0 unit 1 family inet address 192.168.0.1/32
 set protocols bgp group ext type external
 set protocols bgp group ext export send-direct
 set protocols bgp group ext export send-static
 set protocols bgp group ext peer-as 250
 set protocols bgp group ext neighbor 10.0.0.2
 set policy-options policy-statement send-direct term 1 from protocol direct
 set policy-options policy-statement send-direct term 1 then accept
 set policy-options policy-statement send-static term 1 from protocol static
 set policy-options policy-statement send-static term 1 then accept
 set routing-options static route 10.1.0.0/30 next-hop 10.0.0.2
 set routing-options autonomous-system 100

Device R2 set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
 set interfaces fe-1/2/1 unit 3 family inet address 10.1.0.1/30
 set interfaces lo0 unit 2 family inet address 192.168.0.2/32
```

```

set protocols bgp group ext type external
set protocols bgp group ext export send-direct
set protocols bgp group ext export send-static
set protocols bgp group ext local-as 250
set protocols bgp group ext neighbor 10.0.0.1 peer-as 100
set protocols bgp group ext neighbor 10.1.0.2 peer-as 300
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set policy-options policy-statement send-static term 1 from protocol static
set policy-options policy-statement send-static term 1 then accept
set routing-options autonomous-system 200

```

**Device R3**

```

set interfaces fe-1/2/0 unit 4 family inet address 10.1.0.2/30
set interfaces lo0 unit 3 family inet address 192.168.0.3/32
set protocols bgp group ext type external
set protocols bgp group ext export send-direct
set protocols bgp group ext export send-static
set protocols bgp group ext peer-as 250
set protocols bgp group ext neighbor 10.1.0.1
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set policy-options policy-statement send-static term 1 from protocol static
set policy-options policy-statement send-static term 1 then accept
set routing-options static route 10.0.0.0/30 next-hop 10.1.0.1
set routing-options autonomous-system 300

```

### Configuring Device R1

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R1:

1. Configure the interfaces.
 

```

[edit interfaces]
user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30

user@R1# set lo0 unit 1 family inet address 192.168.0.1/32

```
2. Configure external BGP (EBGP).
 

```

[edit protocols bgp group ext]
user@R1# set type external
user@R1# set export send-direct
user@R1# set export send-static
user@R1# set peer-as 250
user@R1# set neighbor 10.0.0.2

```
3. Configure the routing policy.
 

```

[edit policy-options]
user@R1# set policy-statement send-direct term 1 from protocol direct
user@R1# set policy-statement send-direct term 1 then accept
user@R1# set policy-statement send-static term 1 from protocol static
user@R1# set policy-statement send-static term 1 then accept

```

4. Configure a static route to the remote network between Device R2 and Device R3.

```
[edit routing-options]
user@R1# set static route 10.1.0.0/30 next-hop 10.0.0.2
```

5. Configure the global AS number.

```
[edit routing-options]
user@R1# set autonomous-system 100
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show interfaces
fe-1/2/0 {
 unit 1 {
 family inet {
 address 10.0.0.1/30;
 }
 }
}
lo0 {
 unit 1 {
 family inet {
 address 192.168.0.1/32;
 }
 }
}

user@R1# show policy-options
policy-statement send-direct {
 term 1 {
 from protocol direct;
 then accept;
 }
}
policy-statement send-static {
 term 1 {
 from protocol static;
 then accept;
 }
}

user@R1# show protocols
bgp {
 group ext {
 type external;
 export [send-direct send-static];
 peer-as 250;
 neighbor 10.0.0.2;
 }
}

user@R1# show routing-options
static {
```



```

 route 10.1.0.0/30 next-hop 10.0.0.2;
 }
 autonomous-system 100;

```

When you are done configuring the device, enter **commit** from configuration mode.

### Configuring Device R2

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R2:

1. Configure the interfaces.
 

```

[edit interfaces]
user@R2# set fe-1/2/0 unit 2 family inet address 10.0.0.2/30

user@R2# set fe-1/2/1 unit 3 family inet address 10.1.0.1/30

user@R2# set lo0 unit 2 family inet address 192.168.0.2/32

```
2. Configure EBGP.
 

```

[edit protocols bgp group ext]
user@R2# set type external
user@R2# set export send-direct
user@R2# set export send-static
user@R2# set neighbor 10.0.0.1 peer-as 100
user@R2# set neighbor 10.1.0.2 peer-as 300

```
3. Configure the local autonomous system (AS) number.
 

```

[edit protocols bgp group ext]
user@R2# set local-as 250

```
4. Configure the global AS number.
 

```

[edit routing-options]
user@R2# set autonomous-system 200

```
5. Configure the routing policy.
 

```

[edit policy-options]
user@R2# set policy-statement send-direct term 1 from protocol direct
user@R2# set policy-statement send-direct term 1 then accept
user@R2# set policy-statement send-static term 1 from protocol static
user@R2# set policy-statement send-static term 1 then accept

```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@R2# show interfaces
fe-1/2/0 {
 unit 2 {

```

```
 family inet {
 address 10.0.0.2/30;
 }
 }
}
fe-1/2/1 {
 unit 3 {
 family inet {
 address 10.1.0.1/30;
 }
 }
}
lo0 {
 unit 2 {
 family inet {
 address 192.168.0.2/32;
 }
 }
}

user@R2# show policy-options
policy-statement send-direct {
 term 1 {
 from protocol direct;
 then accept;
 }
}
policy-statement send-static {
 term 1 {
 from protocol static;
 then accept;
 }
}

user@R2# show protocols
bgp {
 group ext {
 type external;
 export [send-direct send-static];
 local-as 250;
 neighbor 10.0.0.1 {
 peer-as 100;
 }
 neighbor 10.1.0.2 {
 peer-as 300;
 }
 }
}

user@R2# show routing-options
autonomous-system 200;
```

When you are done configuring the device, enter **commit** from configuration mode.

**Configuring Device R3**

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R3:

1. Configure the interfaces.  

```
[edit interfaces]
user@R3# set fe-1/2/0 unit 4 family inet address 10.1.0.2/30

user@R3# set lo0 unit 3 family inet address 192.168.0.3/32
```
2. Configure EBGP.  

```
[edit protocols bgp group ext]
user@R3# set type external
user@R3# set export send-direct
user@R3# set export send-static
user@R3# set peer-as 250
user@R3# set neighbor 10.1.0.1
```
3. Configure the global autonomous system (AS) number.  

```
[edit routing-options]
user@R3# set autonomous-system 300
```
4. Configure a static route to the remote network between Device R1 and Device R2.  

```
[edit routing-options]
user@R3# set static route 10.0.0.0/30 next-hop 10.1.0.1
```
5. Configure the routing policy.  

```
[edit policy-options]
user@R3# set policy-statement send-direct term 1 from protocol direct
user@R3# set policy-statement send-direct term 1 then accept
user@R3# set policy-statement send-static term 1 from protocol static
user@R3# set policy-statement send-static term 1 then accept
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R3# show interfaces
fe-1/2/0 {
 unit 4 {
 family inet {
 address 10.1.0.2/30;
 }
 }
}
lo0 {
 unit 3 {
```

```
family inet {
 address 192.168.0.3/32;
}
}

user@R3# show policy-options
policy-statement send-direct {
 term 1 {
 from protocol direct;
 then accept;
 }
}
policy-statement send-static {
 term 1 {
 from protocol static;
 then accept;
 }
}

user@R3# show protocols
bgp {
 group ext {
 type external;
 export [send-direct send-static];
 peer-as 250;
 neighbor 10.1.0.1;
 }
}

user@R3# show routing-options
static {
 route 10.0.0.0/30 next-hop 10.1.0.1;
}
autonomous-system 300;
```

When you are done configuring the device, enter **commit** from configuration mode.

---

### Verification

Confirm that the configuration is working properly.

- [Checking the Local and Global AS Settings on page 3716](#)
- [Checking the BGP Peering Sessions on page 3718](#)
- [Verifying the BGP AS Paths on page 3718](#)

#### ***Checking the Local and Global AS Settings***

**Purpose** Make sure that Device R2 has the local and global AS settings configured.

**Action** From operational mode, enter the **show bgp neighbors** command.

```
user@R2> show bgp neighbors
Peer: 10.0.0.1+179 AS 100 Local: 10.0.0.2+61036 AS 250
Type: External State: Established Flags: <Sync>
Last State: OpenConfirm Last Event: RecvKeepAlive
Last Error: None
```

```

Export: [send-direct send-static]
Options: <Preference PeerAS LocalAS Refresh>
Holdtime: 90 Preference: 170 Local AS: 250 Local System AS: 200
Number of flaps: 0
Peer ID: 192.168.0.1 Local ID: 192.168.0.2 Active Holdtime: 90
Keepalive Interval: 30 Peer index: 0
BFD: disabled, down
Local Interface: fe-1/2/0.2
NLRI for restart configured on peer: inet-unicast
NLRI advertised by peer: inet-unicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Stale routes from peer are kept for: 300
Peer does not support Restarter functionality
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Peer supports 4 byte AS extension (peer-as 100)
Peer does not support Addpath
Table inet.0 Bit: 10000
 RIB State: BGP restart is complete
 Send state: in sync
 Active prefixes: 1
 Received prefixes: 3
 Accepted prefixes: 2
 Suppressed due to damping: 0
 Advertised prefixes: 4
Last traffic (seconds): Received 6 Sent 14 Checked 47
Input messages: Total 258 Updates 3 Refreshes 0 Octets 4969
Output messages: Total 258 Updates 2 Refreshes 0 Octets 5037
Output Queue[0]: 0

Peer: 10.1.0.2+179 AS 300 Local: 10.1.0.1+52296 AS 250
Type: External State: Established Flags: <Sync>
Last State: OpenConfirm Last Event: RecvKeepAlive
Last Error: None
Export: [send-direct send-static]
Options: <Preference PeerAS LocalAS Refresh>
Holdtime: 90 Preference: 170 Local AS: 250 Local System AS: 200
Number of flaps: 0
Peer ID: 192.168.0.3 Local ID: 192.168.0.2 Active Holdtime: 90
Keepalive Interval: 30 Peer index: 1
BFD: disabled, down
Local Interface: fe-1/2/1.3
NLRI for restart configured on peer: inet-unicast
NLRI advertised by peer: inet-unicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Stale routes from peer are kept for: 300
Peer does not support Restarter functionality
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Peer supports 4 byte AS extension (peer-as 300)
Peer does not support Addpath
Table inet.0 Bit: 10000
 RIB State: BGP restart is complete
 Send state: in sync
 Active prefixes: 1
 Received prefixes: 3
 Accepted prefixes: 2

```

```

Suppressed due to damping: 0
Advertised prefixes: 4
Last traffic (seconds): Received 19 Sent 26 Checked 9
Input messages: Total 256 Updates 3 Refreshes 0 Octets 4931
Output messages: Total 256 Updates 2 Refreshes 0 Octets 4999
Output Queue[0]: 0

```

**Meaning** The Local AS: 250 and Local System AS: 200 output shows that Device R2 has the expected settings. Additionally, the output shows that the options list includes LocalAS.

### *Checking the BGP Peering Sessions*

**Purpose** Ensure that the sessions are established and that the local AS number 250 is displayed.

**Action** From operational mode, enter the **show bgp summary** command.

```

user@R1> show bgp summary
Groups: 1 Peers: 1 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet.0 4 2 0 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.0.0.2 250 232 233 0 4 1:42:37
2/4/4/0 0/0/0/0

```

```

user@R3> show bgp summary
Groups: 1 Peers: 1 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet.0 4 2 0 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.1.0.1 250 235 236 0 4 1:44:25
2/4/4/0 0/0/0/0

```

**Meaning** Device R1 and Device R3 appear to be peering with a device in AS 250, even though Device R2 is actually in AS 200.

### *Verifying the BGP AS Paths*

**Purpose** Make sure that the routes are in the routing tables and that the AS paths show the local AS number 250.

**Action** From configuration mode, enter the **set route protocol bgp** command.

```

user@R1> show route protocol bgp
inet.0: 6 destinations, 8 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.0/30 [BGP/170] 01:46:44, localpref 100
AS path: 250 I
> to 10.0.0.2 via fe-1/2/0.1
10.1.0.0/30 [BGP/170] 01:46:44, localpref 100
AS path: 250 I
> to 10.0.0.2 via fe-1/2/0.1
192.168.0.2/32 *[BGP/170] 01:46:44, localpref 100
AS path: 250 I
> to 10.0.0.2 via fe-1/2/0.1

```

```

192.168.0.3/32 *[BGP/170] 01:46:40, localpref 100
 AS path: 250 300 I
 > to 10.0.0.2 via fe-1/2/0.1

user@R3> show route protocol bgp

inet.0: 6 destinations, 8 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.0/30 [BGP/170] 01:47:10, localpref 100
 AS path: 250 I
 > to 10.1.0.1 via fe-1/2/0.4
10.1.0.0/30 [BGP/170] 01:47:10, localpref 100
 AS path: 250 I
 > to 10.1.0.1 via fe-1/2/0.4
192.168.0.1/32 *[BGP/170] 01:47:10, localpref 100
 AS path: 250 100 I
 > to 10.1.0.1 via fe-1/2/0.4
192.168.0.2/32 *[BGP/170] 01:47:10, localpref 100
 AS path: 250 I
 > to 10.1.0.1 via fe-1/2/0.4

```

**Meaning** The output shows that Device R1 and Device R3 appear to have routes with AS paths that include AS 250, even though Device R2 is actually in AS 200.

## Example: Configuring a Private Local AS for EBGp Sessions

This example shows how to configure a private local autonomous system (AS) number. The local AS is considered to be private because it is advertised to peers that use the local AS number for peering, but is hidden in the announcements to peers that can use the global AS number for peering.

- [Requirements on page 3719](#)
- [Overview on page 3719](#)
- [Configuration on page 3721](#)
- [Verification on page 3723](#)

### Requirements

No special configuration beyond device initialization is required before you configure this example.

### Overview

Use the **local-as** statement when ISPs merge and want to preserve a customer's configuration, particularly the AS with which the customer is configured to establish a peer relationship. The **local-as** statement simulates the AS number already in place in customer routers, even if the ISP's router has moved to a different AS.


When you use the **private** option, the local AS is used during the establishment of the BGP session with an external BGP (EBGP) neighbor, but is hidden in the AS path sent to other EBGp peers. Only the global AS is included in the AS path sent to external peers.

The **private** option is useful for establishing local peering with routing devices that remain configured with their former AS or with a specific customer that has not yet modified its peer arrangements. The local AS is used to establish the BGP session with the EBGP neighbor, but is hidden in the AS path sent to external peers in another AS.

Include the **private** option so that the local AS is not prepended before the global AS in the AS path sent to external peers. When you specify the **private** option, the local AS is prepended only in the AS path sent to the EBGP neighbor.

Figure 70 shows the sample topology.

Figure 70: Topology for Configuring a Private Local AS



Device R1 is in AS 64496. Device R2 is in AS 64510. Device R3 is in AS 64511. Device R4 is in AS 64512. Device R1 formerly belonged to AS 64497, which has merged with another network and now belongs to AS 64496. Because Device R3 still peers with Device R1, using its former AS, 64497, Device R1 needs to be configured with a local AS of 64497 in order to maintain peering with Device R3. Configuring a local AS of 64497 permits Device R1 to add AS 64497 when advertising routes to Device R3. Device R3 sees an AS path of 64497 64496 for the prefix 10.1.1.2/32, which is Device R2's loopback interface. Device R4, which is behind Device R3, sees an AS path of 64511 64497 64496 64510 to Device R2's loopback interface. To prevent Device R1 from adding the local AS number in its announcements to other peers, this example includes the **local-as 64497 private** statement. The **private** option configures Device R1 to not include the local AS 64497 when announcing routes to Device R2. Device R2 sees an AS path of 64496 64511 to Device R3 and an AS path of 64496 64511 64512 to Device R4. The **private** option in Device R1's configuration causes the AS number 64497 to be missing from the AS paths that Device R1 readvertises to Device R2.

Device R2 is hiding the private local AS from all the routers, except Device R3. The **private** option applies to the routes that Device R1 receives (learns) from Device R3 and that Device R1, in turn, readvertises to other routers. When these routes learned from Device R3 are readvertised by Device R1 to Device R2, the private local AS is missing from the AS path advertised to Device R2.



## Configuration

|                                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>CLI Quick Configuration</b> | To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the <b>[edit]</b> hierarchy level.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Device R1</b>               | <pre> set interfaces fe-1/2/0 unit 3 family inet address 192.168.1.1/24 set interfaces fe-1/2/1 unit 5 family inet address 192.168.10.1/24 set interfaces lo0 unit 2 family inet address 10.1.1.1/32 set protocols bgp group external-AS64511 type external set protocols bgp group external-AS64511 peer-as 64511 set protocols bgp group external-AS64511 local-as 64497 set protocols bgp group external-AS64511 local-as private set protocols bgp group external-AS64511 neighbor 192.168.1.2 set protocols bgp group external-AS64510 type external set protocols bgp group external-AS64510 peer-as 64510 set protocols bgp group external-AS64510 neighbor 192.168.10.2 set policy-options policy-statement send-direct term 1 from protocol direct set policy-options policy-statement send-direct term 1 then accept set routing-options autonomous-system 64496 </pre> |
| <b>Device R2</b>               | <pre> set interfaces fe-1/2/0 unit 6 family inet address 192.168.10.2/24 set interfaces lo0 unit 3 family inet address 10.1.1.2/32 set protocols bgp group external type external set protocols bgp group external export send-direct set protocols bgp group external peer-as 64496 set protocols bgp group external neighbor 192.168.10.1 set policy-options policy-statement send-direct term 1 from protocol direct set policy-options policy-statement send-direct term 1 then accept set routing-options autonomous-system 64510 </pre>                                                                                                                                                                                                                                                                                                                                     |
| <b>Device R3</b>               | <pre> set interfaces fe-1/2/0 unit 4 family inet address 192.168.1.2/24 set interfaces fe-1/2/1 unit 7 family inet address 192.168.5.1/24 set interfaces lo0 unit 4 family inet address 10.1.1.3/32 set protocols bgp group external type external set protocols bgp group external export send-direct set protocols bgp group external neighbor 192.168.1.1 peer-as 64497 set protocols bgp group external neighbor 192.168.5.2 peer-as 64512 set policy-options policy-statement send-direct term 1 from protocol direct set policy-options policy-statement send-direct term 1 then accept set routing-options autonomous-system 64511 </pre>                                                                                                                                                                                                                                  |
| <b>Device R4</b>               | <pre> set interfaces fe-1/2/0 unit 8 family inet address 192.168.5.2/24 set interfaces lo0 unit 5 family inet address 10.1.1.4/32 set protocols bgp group external type external set protocols bgp group external export send-direct set protocols bgp group external peer-as 64511 set protocols bgp group external neighbor 192.168.5.1 set policy-options policy-statement send-direct term 1 from protocol direct set policy-options policy-statement send-direct term 1 then accept set routing-options autonomous-system 64512 </pre>                                                                                                                                                                                                                                                                                                                                       |

### Configuring Device R1

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R1:

1. Configure the interfaces.  

```
[edit interfaces fe-1/2/0 unit 3]
user@R1# set family inet address 192.168.1.1/24

[edit interfaces fe-1/2/1 unit 5]
user@R1# set family inet address 192.168.10.1/24

[edit interfaces lo0 unit 2]
user@R1# set family inet address 10.1.1.1/32
```
2. Configure the EBGP peering session with Device R2.  

```
[edit protocols bgp group external-AS64510]
user@R1# set type external
user@R1# set peer-as 64510
user@R1# set neighbor 192.168.10.2
```
3. Configure the EBGP peering session with Device R3.  

```
[edit protocols bgp group external-AS64511]
user@R1# set type external
user@R1# set peer-as 64511
user@R1# set local-as 64497
user@R1# set local-as private
user@R1# set neighbor 192.168.1.2
```
4. Configure the routing policy.  

```
[edit policy-options policy-statement send-direct term 1]
user@R1# set from protocol direct
user@R1# set then accept
```
5. Configure the global autonomous system (AS) number.  

```
[edit routing-options]
user@R1# set autonomous-system 64496
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show interfaces
fe-1/2/0 {
 unit 3 {
 family inet {
 address 192.168.1.1/24;
 }
 }
}
```

```

 }
 }
 fe-1/2/1 {
 unit 5 {
 family inet {
 address 192.168.10.1/24;
 }
 }
 }
 lo0 {
 unit 2 {
 family inet {
 address 10.1.1.1/32;
 }
 }
 }
}

user@R1# show policy-options
policy-statement send-direct {
 term 1 {
 from protocol direct;
 then accept;
 }
}

user@R1# show protocols
bgp {
 group external-AS64511 {
 type external;
 peer-as 64511;
 local-as 64497 private;
 neighbor 192.168.1.2;
 }
 group external-AS64510 {
 type external;
 peer-as 64510;
 neighbor 192.168.10.2;
 }
}

user@R1# show routing-options
autonomous-system 64496;

```

If you are done configuring the device, enter **commit** from configuration mode.

Repeat the configuration as needed for the other devices in the topology.

### Verification

Confirm that the configuration is working properly.

- [Checking Device R2's AS Paths on page 3724](#)
- [Checking Device R3's AS Paths on page 3724](#)

### ***Checking Device R2's AS Paths***

**Purpose** Make sure that Device R2 does not have AS 64497 in its AS paths to Device R3 and Device R4.

**Action** From operational mode, enter the **show route protocol bgp** command.

```
user@R2> show route protocol bgp
inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.1.1.3/32 *[BGP/170] 01:33:11, localpref 100
 AS path: 64496 64511 I
 > to 192.168.10.1 via fe-1/2/0.6
10.1.1.4/32 *[BGP/170] 01:33:11, localpref 100
 AS path: 64496 64511 64512 I
 > to 192.168.10.1 via fe-1/2/0.6
192.168.5.0/24 *[BGP/170] 01:49:15, localpref 100
 AS path: 64496 64511 I
 > to 192.168.10.1 via fe-1/2/0.6
```

**Meaning** Device R2's AS paths do not include AS 64497.

### ***Checking Device R3's AS Paths***

**Purpose** Make sure that Device R3 does not have AS 64497 in its AS path to Device R4.

**Action** From operational mode, enter the **show route protocol bgp** command.

```
user@R3> show route protocol bgp
inet.0: 7 destinations, 8 routes (7 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.1.1.2/32 *[BGP/170] 01:35:11, localpref 100
 AS path: 64497 64496 64510 I
 > to 192.168.1.1 via fe-1/2/0.4
10.1.1.4/32 *[BGP/170] 01:35:11, localpref 100
 AS path: 64512 I
 > to 192.168.5.2 via fe-1/2/1.7
192.168.5.0/24 [BGP/170] 01:51:15, localpref 100
 AS path: 64512 I
 > to 192.168.5.2 via fe-1/2/1.7
```

**Meaning** Device R3's route to Device R2 (prefix 10.1.1.2) includes both the local and the global AS configured on Device R1 (64497 and 64496, respectively).

**Related Documentation**

- [Examples: Configuring External BGP Peering on page 3601](#)
- [BGP Configuration Overview](#)

---

## **Example: Configuring the Accumulated IGP Attribute for BGP**

- [Understanding the Accumulated IGP Attribute for BGP on page 3725](#)
- [Example: Configuring the Accumulated IGP Attribute for BGP on page 3725](#)

## Understanding the Accumulated IGP Attribute for BGP

The interior gateway protocols (IGPs) are designed to handle routing within a single domain or an autonomous system (AS). Each link is assigned a particular value called a metric. The distance between the two nodes is calculated as a sum of all the metric values of links along the path. The IGP selects the shortest path between two nodes based on distance.

BGP is designed to provide routing over a large number of independent ASs with limited or no coordination among respective administrations. BGP does not use metrics in the path selection decisions.

The accumulated IGP (AIGP) metric attribute for BGP enables deployment in which a single administration can run several contiguous BGP ASs. Such deployments allow BGP to make routing decisions based on the IGP metric. In such networks, it is possible for BGP to select paths based on metrics as is done by IGPs. In this case, BGP chooses the shortest path between two nodes, even though the nodes might be in two different ASs.

The AIGP attribute is particularly useful in networks that use tunneling to deliver a packet to its BGP next hop. The Juniper Networks® Junos® operating system (Junos OS) currently supports the AIGP attribute for two BGP address families, **family inet labeled-unicast** and **family inet6 labeled-unicast**.

AIGP impacts the BGP best-route decision process. The AIGP attribute preference rule is applied after the local-preference rule. The AIGP distance is compared to break a tie. The BGP best-route decision process also impacts the way the interior cost rule is applied if the resolving next hop has an AIGP attribute. Without AIGP enabled, the interior cost of a route is based on the calculation of the metric to the next hop for the route. With AIGP enabled, the resolving AIGP distance is added to the interior cost.

The AIGP attribute is an optional non-transitive BGP path attribute and is specified in Internet draft draft-ietf-idr-aigp-06, *The Accumulated IGP Metric Attribute for BGP*.

## Example: Configuring the Accumulated IGP Attribute for BGP

This example shows how to configure the accumulated IGP (AIGP) metric attribute for BGP.

- [Requirements on page 3725](#)
- [Overview on page 3726](#)
- [Configuration on page 3727](#)
- [Verification on page 3757](#)

### Requirements

This example uses the following hardware and software components:

- Seven BGP-speaking devices.
- Junos OS Release 12.1 or later.

## Overview

The AIGP attribute enables deployments in which a single administration can run several contiguous BGP autonomous systems (ASs). Such deployments allow BGP to make routing decisions based on the IGP metric. With AIGP enabled, BGP can select paths based on IGP metrics. This enables BGP to choose the shortest path between two nodes, even though the nodes might be in different ASs. The AIGP attribute is particularly useful in networks that use tunneling to deliver a packet to its BGP next hop. This example shows AIGP configured with MPLS label-switched paths.

To enable AIGP, you include the **aigp** statement in the BGP configuration on a protocol family basis. Configuring AIGP on a particular family enables sending and receiving of the AIGP attribute on that family. By default, AIGP is disabled. An AIGP-disabled neighbor does not send an AIGP attribute and silently discards a received AIGP attribute.

Junos OS supports AIGP for **family inet labeled-unicast** and **family inet6 labeled-unicast**. The **aigp** statement can be configured for a given family at the global BGP, group, or neighbor level.

By default, the value of the AIGP attribute for a local prefix is zero. An AIGP-enabled neighbor can originate an AIGP attribute for a given prefix by export policy, using the **aigp-originate** policy action. The value of the AIGP attribute reflects the IGP distance to the prefix. Alternatively, you can specify a value, by using the **aigp-originate distance distance** policy action. The configurable range is 0 through 4,294,967,295. Only one node needs to originate an AIGP attribute. The AIGP attribute is retained and readvertised if the neighbors are AIGP enabled with the **aigp** statement in the BGP configuration.

The policy action to originate the AIGP attribute has the following requirements:

- Neighbor must be AIGP enabled.
- Policy must be applied as an export policy.
- Prefix must have no current AIGP attribute.
- Prefix must export with next-hop self.
- Prefix must reside within the AIGP domain. Typically, a loopback IP address is the prefix to originate.


The policy is ignored if these requirements are not met.

### Topology Diagram

Figure 71 shows the topology used in this example. OSPF is used as the interior gateway protocol (IGP). Internal BGP (IBGP) is configured between Device PE1 and Device PE4. External BGP (EBGP) is configured between Device PE7 and Device PE1, between Device PE4 and Device PE3, and between Device PE4 and Device PE2. Devices PE4, PE2, and PE3 are configured for multihop. Device PE4 selects a path based on the AIGP value and then readvertises the AIGP value based on the AIGP and policy configuration. Device PE1 readvertises the AIGP value to Device PE7, which is in another administrative domain. Every device has two loopback interface addresses: 10.9.9.x is used for BGP peering and the router ID, and 10.100.1.x is used for the BGP next hop.

The network between Device PE1 and PE3 has IBGP peering and multiple OSPF areas. The external link to Device PE7 is configured to show that the AIGP attribute is readadvertised to a neighbor outside of the administrative domain, if that neighbor is AIGP enabled.

**Figure 71: Advertisement of Multiple Paths in BGP**



For origination of an AIGP attribute, the BGP next hop is required to be itself. If the BGP next hop remains unchanged, the received AIGP attribute is readadvertised, as is, to another AIGP neighbor. If the next hop changes, the received AIGP attribute is readadvertised with an increased value to another AIGP neighbor. The increase in value reflects the IGP distance to the previous BGP next hop. To demonstrate, this example uses loopback interface addresses for Device PE4's EBGP peering sessions with Device PE2 and Device PE3. Multihop is enabled on these sessions so that a recursive lookup is performed to determine the point-to-point interface. Because the next hop changes, the IGP distance is added to the AIGP distance.

### Configuration

- [Configuring Device P1 on page 3733](#)
- [Configuring Device P2 on page 3736](#)
- [Configuring Device PE4 on page 3739](#)
- [Configuring Device PE1 on page 3744](#)
- [Configuring Device PE2 on page 3748](#)
- [Configuring Device PE3 on page 3752](#)
- [Configuring Device PE7 on page 3755](#)

### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
Device P1 set interfaces fe-1/2/0 unit 1 description P1-to-PE1
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.2/30
set interfaces fe-1/2/0 unit 1 family mpls
```

```
set interfaces fe-1/2/1 unit 4 description P1-to-P2
set interfaces fe-1/2/1 unit 4 family inet address 10.0.0.29/30
set interfaces fe-1/2/1 unit 4 family mpls
set interfaces fe-1/2/2 unit 8 description P1-to-PE4
set interfaces fe-1/2/2 unit 8 family inet address 10.0.0.17/30
set interfaces fe-1/2/2 unit 8 family mpls
set interfaces lo0 unit 3 family inet address 10.9.9.2/32
set interfaces lo0 unit 3 family inet address 10.100.1.2/32
set protocols rsvp interface fe-1/2/0.1
set protocols rsvp interface fe-1/2/2.8
set protocols rsvp interface fe-1/2/1.4
set protocols mpls label-switched-path P1-to-P2 to 10.9.9.3
set protocols mpls label-switched-path P1-to-PE1 to 10.9.9.1
set protocols mpls label-switched-path P1-to-PE4 to 10.9.9.4
set protocols mpls interface fe-1/2/0.1
set protocols mpls interface fe-1/2/2.8
set protocols mpls interface fe-1/2/1.4
set protocols bgp group internal type internal
set protocols bgp group internal local-address 10.9.9.2
set protocols bgp group internal family inet labeled-unicast aigp
set protocols bgp group internal neighbor 10.9.9.1
set protocols bgp group internal neighbor 10.9.9.3
set protocols bgp group internal neighbor 10.9.9.4
set protocols ospf area 0.0.0.1 interface fe-1/2/0.1 metric 1
set protocols ospf area 0.0.0.1 interface fe-1/2/1.4 metric 1
set protocols ospf area 0.0.0.0 interface fe-1/2/2.8 metric 1
set protocols ospf area 0.0.0.0 interface 10.9.9.2 passive
set protocols ospf area 0.0.0.0 interface 10.9.9.2 metric 1
set protocols ospf area 0.0.0.0 interface 10.100.1.2 passive
set protocols ospf area 0.0.0.0 interface 10.100.1.2 metric 1
set routing-options router-id 10.9.9.2
set routing-options autonomous-system 13979
```

**Device P2**

```
set interfaces fe-1/2/0 unit 3 description P2-to-PE1
set interfaces fe-1/2/0 unit 3 family inet address 10.0.0.6/30
set interfaces fe-1/2/0 unit 3 family mpls
set interfaces fe-1/2/1 unit 5 description P2-to-P1
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.30/30
set interfaces fe-1/2/1 unit 5 family mpls
set interfaces fe-1/2/2 unit 6 description P2-to-PE4
set interfaces fe-1/2/2 unit 6 family inet address 10.0.0.13/30
set interfaces fe-1/2/2 unit 6 family mpls
set interfaces lo0 unit 5 family inet address 10.9.9.3/32
set interfaces lo0 unit 5 family inet address 10.100.1.3/32
set protocols rsvp interface fe-1/2/1.5
set protocols rsvp interface fe-1/2/2.6
set protocols rsvp interface fe-1/2/0.3
set protocols mpls label-switched-path P2-to-PE1 to 10.9.9.1
set protocols mpls label-switched-path P2-to-P1 to 10.9.9.2
set protocols mpls label-switched-path P2-to-PE4 to 10.9.9.4
set protocols mpls interface fe-1/2/1.5
set protocols mpls interface fe-1/2/2.6
set protocols mpls interface fe-1/2/0.3
set protocols bgp group internal type internal
set protocols bgp group internal local-address 10.9.9.3
set protocols bgp group internal family inet labeled-unicast aigp
```



```

set protocols bgp group internal neighbor 10.9.9.1
set protocols bgp group internal neighbor 10.9.9.2
set protocols bgp group internal neighbor 10.9.9.4
set protocols ospf area 0.0.0.0 interface fe-1/2/2.6 metric 1
set protocols ospf area 0.0.0.0 interface 10.9.9.3 passive
set protocols ospf area 0.0.0.0 interface 10.9.9.3 metric 1
set protocols ospf area 0.0.0.0 interface 10.100.1.3 passive
set protocols ospf area 0.0.0.0 interface 10.100.1.3 metric 1
set routing-options router-id 10.9.9.3
set routing-options autonomous-system 13979

```

**Device PE4**

```

set interfaces fe-1/2/0 unit 7 description PE4-to-P2
set interfaces fe-1/2/0 unit 7 family inet address 10.0.0.14/30
set interfaces fe-1/2/0 unit 7 family mpls
set interfaces fe-1/2/1 unit 9 description PE4-to-P1
set interfaces fe-1/2/1 unit 9 family inet address 10.0.0.18/30
set interfaces fe-1/2/1 unit 9 family mpls
set interfaces fe-1/2/2 unit 10 description PE4-to-PE2
set interfaces fe-1/2/2 unit 10 family inet address 10.0.0.21/30
set interfaces fe-1/2/2 unit 10 family mpls
set interfaces fe-1/0/2 unit 12 description PE4-to-PE3
set interfaces fe-1/0/2 unit 12 family inet address 10.0.0.25/30
set interfaces fe-1/0/2 unit 12 family mpls
set interfaces lo0 unit 7 family inet address 10.9.9.4/32
set interfaces lo0 unit 7 family inet address 10.100.1.4/32
set protocols rsvp interface fe-1/2/0.7
set protocols rsvp interface fe-1/2/1.9
set protocols rsvp interface fe-1/2/2.10
set protocols rsvp interface fe-1/0/2.12
set protocols mpls label-switched-path PE4-to-PE2 to 10.9.9.5
set protocols mpls label-switched-path PE4-to-PE3 to 10.9.9.6
set protocols mpls label-switched-path PE4-to-P1 to 10.9.9.2
set protocols mpls label-switched-path PE4-to-P2 to 10.9.9.3
set protocols mpls interface fe-1/2/0.7
set protocols mpls interface fe-1/2/1.9
set protocols mpls interface fe-1/2/2.10
set protocols mpls interface fe-1/0/2.12
set protocols bgp export next-hop
set protocols bgp export aigp
set protocols bgp group internal type internal
set protocols bgp group internal local-address 10.9.9.4
set protocols bgp group internal family inet labeled-unicast aigp
set protocols bgp group internal neighbor 10.9.9.1
set protocols bgp group internal neighbor 10.9.9.3
set protocols bgp group internal neighbor 10.9.9.2
set protocols bgp group external type external
set protocols bgp group external multihop ttl 2
set protocols bgp group external local-address 10.9.9.4
set protocols bgp group external family inet labeled-unicast aigp
set protocols bgp group external peer-as 7018
set protocols bgp group external neighbor 10.9.9.5
set protocols bgp group external neighbor 10.9.9.6
set protocols ospf area 0.0.0.0 interface fe-1/2/1.9 metric 1
set protocols ospf area 0.0.0.0 interface fe-1/2/0.7 metric 1
set protocols ospf area 0.0.0.0 interface 10.9.9.4 passive
set protocols ospf area 0.0.0.0 interface 10.9.9.4 metric 1

```

```
set protocols ospf area 0.0.0.0 interface 10.100.1.4 passive
set protocols ospf area 0.0.0.0 interface 10.100.1.4 metric 1
set protocols ospf area 0.0.0.2 interface fe-1/2/2.10 metric 1
set protocols ospf area 0.0.0.3 interface fe-1/0/2.12 metric 1
set policy-options policy-statement aigp term 10 from protocol static
set policy-options policy-statement aigp term 10 from route-filter 44.0.0.0/24 exact
set policy-options policy-statement aigp term 10 then aigp-originate distance 200
set policy-options policy-statement aigp term 10 then next-hop 10.100.1.4
set policy-options policy-statement aigp term 10 then accept
set policy-options policy-statement next-hop term 10 from protocol bgp
set policy-options policy-statement next-hop term 10 then next-hop 10.100.1.4
set policy-options policy-statement next-hop term 10 then accept
set policy-options policy-statement next-hop term 20 from protocol direct
set policy-options policy-statement next-hop term 20 from route-filter 10.9.9.4/32 exact
set policy-options policy-statement next-hop term 20 from route-filter 10.100.1.4/32
 exact
set policy-options policy-statement next-hop term 20 then next-hop 10.100.1.4
set policy-options policy-statement next-hop term 20 then accept
set routing-options static route 44.0.0.0/24 discard
set routing-options router-id 10.9.9.4
set routing-options autonomous-system 13979
```

**Device PE1**

```
set interfaces fe-1/2/0 unit 0 description PE1-to-P1
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.1/30
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/1 unit 2 description PE1-to-P2
set interfaces fe-1/2/1 unit 2 family inet address 10.0.0.5/30
set interfaces fe-1/2/1 unit 2 family mpls
set interfaces fe-1/2/2 unit 14 description PE1-to-PE7
set interfaces fe-1/2/2 unit 14 family inet address 10.0.0.9/30
set interfaces lo0 unit 1 family inet address 10.9.9.1/32
set interfaces lo0 unit 1 family inet address 10.100.1.1/32
set protocols rsvp interface fe-1/2/0.0
set protocols rsvp interface fe-1/2/1.2
set protocols rsvp interface fe-1/2/2.14
set protocols mpls label-switched-path PE1-to-P1 to 10.9.9.2
set protocols mpls label-switched-path PE1-to-P2 to 10.9.9.3
set protocols mpls interface fe-1/2/0.0
set protocols mpls interface fe-1/2/1.2
set protocols mpls interface fe-1/2/2.14
set protocols bgp group internal type internal
set protocols bgp group internal local-address 10.9.9.1
set protocols bgp group internal family inet labeled-unicast aigp
set protocols bgp group internal export SET_EXPORT_ROUTES
set protocols bgp group internal vpn-apply-export
set protocols bgp group internal neighbor 10.9.9.4
set protocols bgp group internal neighbor 10.9.9.2
set protocols bgp group internal neighbor 10.9.9.3
set protocols bgp group external type external
set protocols bgp group external family inet labeled-unicast aigp
set protocols bgp group external export SET_EXPORT_ROUTES
set protocols bgp group external peer-as 7019
set protocols bgp group external neighbor 10.0.0.10
set protocols ospf area 0.0.0.1 interface fe-1/2/0.0 metric 1
set protocols ospf area 0.0.0.1 interface fe-1/2/1.2 metric 1
set protocols ospf area 0.0.0.1 interface 10.9.9.1 passive
```

```

set protocols ospf area 0.0.0.1 interface 10.9.9.1 metric 1
set protocols ospf area 0.0.0.1 interface 10.100.1.1 passive
set protocols ospf area 0.0.0.1 interface 10.100.1.1 metric 1
set policy-options policy-statement SET_EXPORT_ROUTES term 10 from protocol direct
set policy-options policy-statement SET_EXPORT_ROUTES term 10 from protocol bgp
set policy-options policy-statement SET_EXPORT_ROUTES term 10 then next-hop
 10.100.1.1
set policy-options policy-statement SET_EXPORT_ROUTES term 10 then accept
set routing-options router-id 10.9.9.1
set routing-options autonomous-system 13979

```

#### Device PE2

```

set interfaces fe-1/2/0 unit 11 description PE2-to-PE4
set interfaces fe-1/2/0 unit 11 family inet address 10.0.0.22/30
set interfaces fe-1/2/0 unit 11 family mpls
set interfaces lo0 unit 9 family inet address 10.9.9.5/32 primary
set interfaces lo0 unit 9 family inet address 10.100.1.5/32
set protocols rsvp interface fe-1/2/0.11
set protocols mpls label-switched-path PE2-to-PE4 to 10.9.9.4
set protocols mpls interface fe-1/2/0.11
set protocols bgp group external type external
set protocols bgp group external multihop ttl 2
set protocols bgp group external local-address 10.9.9.5
set protocols bgp group external family inet labeled-unicast aigp
set protocols bgp group external export next-hop
set protocols bgp group external export aigp
set protocols bgp group external export SET_EXPORT_ROUTES
set protocols bgp group external vpn-apply-export
set protocols bgp group external peer-as 13979
set protocols bgp group external neighbor 10.9.9.4
set protocols ospf area 0.0.0.2 interface 10.9.9.5 passive
set protocols ospf area 0.0.0.2 interface 10.9.9.5 metric 1
set protocols ospf area 0.0.0.2 interface 10.100.1.5 passive
set protocols ospf area 0.0.0.2 interface 10.100.1.5 metric 1
set protocols ospf area 0.0.0.2 interface fe-1/2/0.11 metric 1
set policy-options policy-statement SET_EXPORT_ROUTES term 10 from protocol direct
set policy-options policy-statement SET_EXPORT_ROUTES term 10 from protocol static
set policy-options policy-statement SET_EXPORT_ROUTES term 10 from protocol bgp
set policy-options policy-statement SET_EXPORT_ROUTES term 10 then next-hop
 10.100.1.5
set policy-options policy-statement SET_EXPORT_ROUTES term 10 then accept
set policy-options policy-statement aigp term 10 from route-filter 55.0.0.0/24 exact
set policy-options policy-statement aigp term 10 then aigp-originate distance 20
set policy-options policy-statement aigp term 10 then next-hop 10.100.1.5
set policy-options policy-statement aigp term 10 then accept
set policy-options policy-statement aigp term 20 from route-filter 99.0.0.0/24 exact
set policy-options policy-statement aigp term 20 then aigp-originate distance 30
set policy-options policy-statement aigp term 20 then next-hop 10.100.1.5
set policy-options policy-statement aigp term 20 then accept
set policy-options policy-statement next-hop term 10 from protocol bgp
set policy-options policy-statement next-hop term 10 then next-hop 10.100.1.5
set policy-options policy-statement next-hop term 10 then accept
set policy-options policy-statement next-hop term 20 from protocol direct
set policy-options policy-statement next-hop term 20 from route-filter 10.9.9.5/32 exact
set policy-options policy-statement next-hop term 20 from route-filter 10.100.1.5/32
 exact
set policy-options policy-statement next-hop term 20 then next-hop 10.100.1.5

```

```
set policy-options policy-statement next-hop term 20 then accept
set routing-options static route 99.0.0.0/24 discard
set routing-options static route 55.0.0.0/24 discard
set routing-options router-id 10.9.9.5
set routing-options autonomous-system 7018
```

**Device PE3**

```
set interfaces fe-1/2/0 unit 13 description PE3-to-PE4
set interfaces fe-1/2/0 unit 13 family inet address 10.0.0.26/30
set interfaces fe-1/2/0 unit 13 family mpls
set interfaces lo0 unit 11 family inet address 10.9.9.6/32
set interfaces lo0 unit 11 family inet address 10.100.1.6/32
set protocols rsvp interface fe-1/2/0.13
set protocols mpls label-switched-path PE3-to-PE4 to 10.9.9.4
set protocols mpls interface fe-1/2/0.13
set protocols bgp group external type external
set protocols bgp group external multihop ttl 2
set protocols bgp group external local-address 10.9.9.6
set protocols bgp group external family inet labeled-unicast aigp
set protocols bgp group external export next-hop
set protocols bgp group external export SET_EXPORT_ROUTES
set protocols bgp group external vpn-apply-export
set protocols bgp group external peer-as 13979
set protocols bgp group external neighbor 10.9.9.4
set protocols ospf area 0.0.0.3 interface 10.9.9.6 passive
set protocols ospf area 0.0.0.3 interface 10.9.9.6 metric 1
set protocols ospf area 0.0.0.3 interface 10.100.1.6 passive
set protocols ospf area 0.0.0.3 interface 10.100.1.6 metric 1
set protocols ospf area 0.0.0.3 interface fe-1/2/0.13 metric 1
set policy-options policy-statement SET_EXPORT_ROUTES term 10 from protocol direct
set policy-options policy-statement SET_EXPORT_ROUTES term 10 from protocol static
set policy-options policy-statement SET_EXPORT_ROUTES term 10 from protocol bgp
set policy-options policy-statement SET_EXPORT_ROUTES term 10 then next-hop
 10.100.1.6
set policy-options policy-statement SET_EXPORT_ROUTES term 10 then accept
set policy-options policy-statement next-hop term 10 from protocol bgp
set policy-options policy-statement next-hop term 10 then next-hop 10.100.1.6
set policy-options policy-statement next-hop term 10 then accept
set policy-options policy-statement next-hop term 20 from protocol direct
set policy-options policy-statement next-hop term 20 from route-filter 10.9.9.6/32 exact
set policy-options policy-statement next-hop term 20 from route-filter 10.100.1.6/32
 exact
set policy-options policy-statement next-hop term 20 then next-hop 10.100.1.6
set policy-options policy-statement next-hop term 20 then accept
set routing-options router-id 10.9.9.6
set routing-options autonomous-system 7018
```

**Device PE7**

```
set interfaces fe-1/2/0 unit 15 description PE7-to-PE1
set interfaces fe-1/2/0 unit 15 family inet address 10.0.0.10/30
set interfaces lo0 unit 13 family inet address 10.9.9.7/32
set interfaces lo0 unit 13 family inet address 10.100.1.7/32
set protocols bgp group external type external
set protocols bgp group external family inet labeled-unicast aigp
set protocols bgp group external export SET_EXPORT_ROUTES
set protocols bgp group external peer-as 13979
set protocols bgp group external neighbor 10.0.0.9
```

```

set policy-options policy-statement SET_EXPORT_ROUTES term 10 from protocol direct
set policy-options policy-statement SET_EXPORT_ROUTES term 10 from protocol bgp
set policy-options policy-statement SET_EXPORT_ROUTES term 10 then next-hop
 10.100.1.7
set policy-options policy-statement SET_EXPORT_ROUTES term 10 then accept
set routing-options router-id 10.9.9.7
set routing-options autonomous-system 7019

```

### Configuring Device P1

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device P1:

1. Configure the interfaces.

```

[edit interfaces]
user@P1# set fe-1/2/0 unit 1 description P1-to-PE1
user@P1# set fe-1/2/0 unit 1 family inet address 10.0.0.2/30
user@P1# set fe-1/2/0 unit 1 family mpls
user@P1# set fe-1/2/1 unit 4 description P1-to-P2
user@P1# set fe-1/2/1 unit 4 family inet address 10.0.0.29/30
user@P1# set fe-1/2/1 unit 4 family mpls
user@P1# set fe-1/2/2 unit 8 description P1-to-PE4
user@P1# set fe-1/2/2 unit 8 family inet address 10.0.0.17/30
user@P1# set fe-1/2/2 unit 8 family mpls
user@P1# set lo0 unit 3 family inet address 10.9.9.2/32
user@P1# set lo0 unit 3 family inet address 10.100.1.2/32

```

2. Configure MPLS and a signaling protocol, such as RSVP or LDP.

```

[edit protocols]
user@P1# set rsvp interface fe-1/2/0.1
user@P1# set rsvp interface fe-1/2/2.8
user@P1# set rsvp interface fe-1/2/1.4
user@P1# set mpls label-switched-path P1-to-P2 to 10.9.9.3
user@P1# set mpls label-switched-path P1-to-PE1 to 10.9.9.1
user@P1# set mpls label-switched-path P1-to-PE4 to 10.9.9.4
user@P1# set mpls interface fe-1/2/0.1
user@P1# set mpls interface fe-1/2/2.8
user@P1# set mpls interface fe-1/2/1.4

```

3. Configure BGP.

```

[edit protocols bgp group internal]
user@P1# set type internal
user@P1# set local-address 10.9.9.2
user@P1# set neighbor 10.9.9.1
user@P1# set neighbor 10.9.9.3
user@P1# set neighbor 10.9.9.4

```

4. Enable AIGP.

```

[edit protocols bgp group internal]
user@P1# set family inet labeled-unicast aigp

```

5. Configure an IGP, such as OSPF, RIP, or IS-IS.

```
[edit protocols ospf]
user@P1# set area 0.0.0.1 interface fe-1/2/0.1 metric 1
user@P1# set area 0.0.0.1 interface fe-1/2/1.4 metric 1
user@P1# set area 0.0.0.0 interface fe-1/2/2.8 metric 1
user@P1# set area 0.0.0.0 interface 10.9.9.2 passive
user@P1# set area 0.0.0.0 interface 10.9.9.2 metric 1
user@P1# set area 0.0.0.0 interface 10.100.1.2 passive
user@P1# set area 0.0.0.0 interface 10.100.1.2 metric 1
```

6. Configure the router ID and the autonomous system number.

```
[edit routing-options]
user@P1# set router-id 10.9.9.2
user@P1# set autonomous-system 13979
```

7. If you are done configuring the device, commit the configuration.

```
user@P1# commit
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@P1# show interfaces
fe-1/2/0 {
 unit 1 {
 description P1-to-PE1;
 family inet {
 address 10.0.0.2/30;
 }
 family mpls;
 }
}
fe-1/2/1 {
 unit 4 {
 description P1-to-P2;
 family inet {
 address 10.0.0.29/30;
 }
 family mpls;
 }
}
fe-1/2/2 {
 unit 8 {
 description P1-to-PE4;
 family inet {
 address 10.0.0.17/30;
 }
 family mpls;
 }
}
lo0 {
 unit 3 {
 family inet {
```

```

 address 10.9.9.2/32;
 address 10.100.1.2/32;
 }
}
}

user@P1# show protocols
rsvp {
 interface fe-1/2/0.1;
 interface fe-1/2/2.8;
 interface fe-1/2/1.4;
}
mpls {
 label-switched-path P1-to-P2 {
 to 10.9.9.3;
 }
 label-switched-path P1-to-PE1 {
 to 10.9.9.1;
 }
 label-switched-path P1-to-PE4 {
 to 10.9.9.4;
 }
 interface fe-1/2/0.1;
 interface fe-1/2/2.8;
 interface fe-1/2/1.4;
}
bgp {
 group internal {
 type internal;
 local-address 10.9.9.2;
 family inet {
 labeled-unicast {
 aigp;
 }
 }
 neighbor 10.9.9.1;
 neighbor 10.9.9.3;
 neighbor 10.9.9.4;
 }
}
ospf {
 area 0.0.0.1 {
 interface fe-1/2/0.1 {
 metric 1;
 }
 interface fe-1/2/1.4 {
 metric 1;
 }
 }
 area 0.0.0.0 {
 interface fe-1/2/2.8 {
 metric 1;
 }
 interface 10.9.9.2 {
 passive;
 metric 1;
 }
 }
}

```

```
}
interface 10.100.1.2 {
 passive;
 metric 1;
}
}
}

user@P1# show routing-options
router-id 10.9.9.2;
autonomous-system 13979;
```

### Configuring Device P2

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device P2:

1. Configure the interfaces.

```
[edit interfaces]
user@P2# set fe-1/2/0 unit 3 description P2-to-PE1
user@P2# set fe-1/2/0 unit 3 family inet address 10.0.0.6/30
user@P2# set fe-1/2/0 unit 3 family mpls
user@P2# set fe-1/2/1 unit 5 description P2-to-P1
user@P2# set fe-1/2/1 unit 5 family inet address 10.0.0.30/30
user@P2# set fe-1/2/1 unit 5 family mpls
user@P2# set fe-1/2/2 unit 6 description P2-to-PE4
user@P2# set fe-1/2/2 unit 6 family inet address 10.0.0.13/30
user@P2# set fe-1/2/2 unit 6 family mpls
user@P2# set lo0 unit 5 family inet address 10.9.9.3/32
user@P2# set lo0 unit 5 family inet address 10.100.1.3/32
```

2. Configure MPLS and a signaling protocol, such as RSVP or LDP.

```
[edit protocols]
user@P2# set rsvp interface fe-1/2/1.5
user@P2# set rsvp interface fe-1/2/2.6
user@P2# set rsvp interface fe-1/2/0.3
user@P2# set mpls label-switched-path P2-to-PE1 to 10.9.9.1
user@P2# set mpls label-switched-path P2-to-P1 to 10.9.9.2
user@P2# set mpls label-switched-path P2-to-PE4 to 10.9.9.4
user@P2# set mpls interface fe-1/2/1.5
user@P2# set mpls interface fe-1/2/2.6
user@P2# set mpls interface fe-1/2/0.3
```

3. Configure BGP.

```
[edit protocols bgp group internal]
user@P2# set type internal
user@P2# set local-address 10.9.9.3
user@P2# set neighbor 10.9.9.1
user@P2# set neighbor 10.9.9.2
user@P2# set neighbor 10.9.9.4
```



4. Enable AIGP.

```
[edit protocols bgp group internal]
user@P2# set family inet labeled-unicast aigp
```

5. Configure an IGP, such as OSPF, RIP, or IS-IS.

```
[edit protocols ospf]
user@P2# set area 0.0.0.0 interface fe-1/2/2.6 metric 1
user@P2# set area 0.0.0.0 interface 10.9.9.3 passive
user@P2# set area 0.0.0.0 interface 10.9.9.3 metric 1
user@P2# set area 0.0.0.0 interface 10.100.1.3 passive
user@P2# set area 0.0.0.0 interface 10.100.1.3 metric 1
```

6. Configure the router ID and the autonomous system number.

```
[edit routing-options]
user@P2# set router-id 10.9.9.3
user@P2# set autonomous-system 13979
```

7. If you are done configuring the device, commit the configuration.

```
user@P2# commit
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@P2# show interfaces
fe-1/2/0 {
 unit 3 {
 description P2-to-PE1;
 family inet {
 address 10.0.0.6/30;
 }
 family mpls;
 }
}
fe-1/2/1 {
 unit 5 {
 description P2-to-P1;
 family inet {
 address 10.0.0.30/30;
 }
 family mpls;
 }
}
fe-1/2/2 {
 unit 6 {
 description P2-to-PE4;
 family inet {
 address 10.0.0.13/30;
 }
 family mpls;
 }
}
lo0 {
```

```
unit 5 {
 family inet {
 address 10.9.9.3/32;
 address 10.100.1.3/32;
 }
}

user@P2# show protocols
rsvp {
 interface fe-1/2/1.5;
 interface fe-1/2/2.6;
 interface fe-1/2/0.3;
}
mpls {
 label-switched-path P2-to-PE1 {
 to 10.9.9.1;
 }
 label-switched-path P2-to-P1 {
 to 10.9.9.2;
 }
 label-switched-path P2-to-PE4 {
 to 10.9.9.4;
 }
 interface fe-1/2/1.5;
 interface fe-1/2/2.6;
 interface fe-1/2/0.3;
}
bgp {
 group internal {
 type internal;
 local-address 10.9.9.3;
 family inet {
 labeled-unicast {
 aigp;
 }
 }
 neighbor 10.9.9.1;
 neighbor 10.9.9.2;
 neighbor 10.9.9.4;
 }
}
ospf {
 area 0.0.0.0 {
 interface fe-1/2/2.6 {
 metric 1;
 }
 interface 10.9.9.3 {
 passive;
 metric 1;
 }
 interface 10.100.1.3 {
 passive;
 metric 1;
 }
 }
}
```

```

}
user@P2# show routing-options
router-id 10.9.9.3;
autonomous-system 13979;

```

### Configuring Device PE4

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device PE4:

1. Configure the interfaces.

```

[edit interfaces]
user@PE4# set fe-1/2/0 unit 7 description PE4-to-P2
user@PE4# set fe-1/2/0 unit 7 family inet address 10.0.0.14/30
user@PE4# set fe-1/2/0 unit 7 family mpls
user@PE4# set fe-1/2/1 unit 9 description PE4-to-P1
user@PE4# set fe-1/2/1 unit 9 family inet address 10.0.0.18/30
user@PE4# set fe-1/2/1 unit 9 family mpls
user@PE4# set fe-1/2/2 unit 10 description PE4-to-PE2
user@PE4# set fe-1/2/2 unit 10 family inet address 10.0.0.21/30
user@PE4# set fe-1/2/2 unit 10 family mpls
user@PE4# set fe-1/0/2 unit 12 description PE4-to-PE3
user@PE4# set fe-1/0/2 unit 12 family inet address 10.0.0.25/30
user@PE4# set fe-1/0/2 unit 12 family mpls
user@PE4# set lo0 unit 7 family inet address 10.9.9.4/32
user@PE4# set lo0 unit 7 family inet address 10.100.1.4/32

```

2. Configure MPLS and a signaling protocol, such as RSVP or LDP.

```

[edit protocols]
user@PE4# set rsvp interface fe-1/2/0.7
user@PE4# set rsvp interface fe-1/2/1.9
user@PE4# set rsvp interface fe-1/2/2.10
user@PE4# set rsvp interface fe-1/0/2.12
user@PE4# set mpls label-switched-path PE4-to-PE2 to 10.9.9.5
user@PE4# set mpls label-switched-path PE4-to-PE3 to 10.9.9.6
user@PE4# set mpls label-switched-path PE4-to-P1 to 10.9.9.2
user@PE4# set mpls label-switched-path PE4-to-P2 to 10.9.9.3
user@PE4# set mpls interface fe-1/2/0.7
user@PE4# set mpls interface fe-1/2/1.9
user@PE4# set mpls interface fe-1/2/2.10
user@PE4# set mpls interface fe-1/0/2.12

```

3. Configure BGP.

```

[edit protocols bgp]
user@PE4# set export next-hop
user@PE4# set export aigp
user@PE4# set group internal type internal
user@PE4# set group internal local-address 10.9.9.4
user@PE4# set group internal neighbor 10.9.9.1
user@PE4# set group internal neighbor 10.9.9.3

```

```
user@PE4# set group internal neighbor 10.9.9.2
user@PE4# set group external type external
user@PE4# set group external multihop ttl 2
user@PE4# set group external local-address 10.9.9.4
user@PE4# set group external peer-as 7018
user@PE4# set group external neighbor 10.9.9.5
user@PE4# set group external neighbor 10.9.9.6
```

4. Enable AIGP.

```
[edit protocols bgp]
user@PE4# set group external family inet labeled-unicast aigp
user@PE4# set group internal family inet labeled-unicast aigp
```

5. Originate a prefix, and configure an AIGP distance.

By default, a prefix is originated using the current IGP distance. Optionally, you can configure a distance for the AIGP attribute, using the **distance** option, as shown here.

```
[edit policy-options policy-statement aigp term 10]
user@PE4# set from protocol static
user@PE4# set from route-filter 44.0.0.0/24 exact
user@PE4# set then aigp-originate distance 200
user@PE4# set then next-hop 10.100.1.4
user@PE4# set then accept
```

6. Enable the policies.

```
[edit policy-options policy-statement next-hop]
user@PE4# set term 10 from protocol bgp
user@PE4# set term 10 then next-hop 10.100.1.4
user@PE4# set term 10 then accept
user@PE4# set term 20 from protocol direct
user@PE4# set term 20 from route-filter 10.9.9.4/32 exact
user@PE4# set term 20 from route-filter 10.100.1.4/32 exact
user@PE4# set term 20 then next-hop 10.100.1.4
user@PE4# set term 20 then accept
```

7. Configure a static route.

```
[edit routing-options]
user@PE4# set static route 44.0.0.0/24 discard
```

8. Configure an IGP, such as OSPF, RIP, or IS-IS.

```
[edit protocols ospf]
user@PE4# set area 0.0.0.0 interface fe-1/2/1.9 metric 1
user@PE4# set area 0.0.0.0 interface fe-1/2/0.7 metric 1
user@PE4# set area 0.0.0.0 interface 10.9.9.4 passive
user@PE4# set area 0.0.0.0 interface 10.9.9.4 metric 1
user@PE4# set area 0.0.0.0 interface 10.100.1.4 passive
user@PE4# set area 0.0.0.0 interface 10.100.1.4 metric 1
user@PE4# set area 0.0.0.2 interface fe-1/2/2.10 metric 1
user@PE4# set area 0.0.0.3 interface fe-1/0/2.12 metric 1
```

9. Configure the router ID and the autonomous system number.

```
[edit routing-options]
```

```

user@PE4# set router-id 10.9.9.4
user@PE4# set autonomous-system 13979

```

10. If you are done configuring the device, commit the configuration.

```

user@PE4# commit

```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@PE4# show interfaces
fe-1/0/2 {
 unit 12 {
 description PE4-to-PE3;
 family inet {
 address 10.0.0.25/30;
 }
 family mpls;
 }
}
fe-1/2/0 {
 unit 7 {
 description PE4-to-P2;
 family inet {
 address 10.0.0.14/30;
 }
 family mpls;
 }
}
fe-1/2/1 {
 unit 9 {
 description PE4-to-P1;
 family inet {
 address 10.0.0.18/30;
 }
 family mpls;
 }
}
fe-1/2/2 {
 unit 10 {
 description PE4-to-PE2;
 family inet {
 address 10.0.0.21/30;
 }
 family mpls;
 }
}
lo0 {
 unit 7 {
 family inet {
 address 10.9.9.4/32;
 address 10.100.1.4/32;
 }
 }
}

```

```
 }
 }
user@PE4# show policy-options
policy-statement aigp {
 term 10 {
 from {
 protocol static;
 route-filter 44.0.0.0/24 exact;
 }
 then {
 aigp-originate distance 200;
 next-hop 10.100.1.4;
 accept;
 }
 }
}
policy-statement next-hop {
 term 10 {
 from protocol bgp;
 then {
 next-hop 10.100.1.4;
 accept;
 }
 }
 term 20 {
 from {
 protocol direct;
 route-filter 10.9.9.4/32 exact;
 route-filter 10.100.1.4/32 exact;
 }
 then {
 next-hop 10.100.1.4;
 accept;
 }
 }
}
user@PE4# show protocols
rsvp {
 interface fe-1/2/0.7;
 interface fe-1/2/1.9;
 interface fe-1/2/2.10;
 interface fe-1/0/2.12;
}
mpls {
 label-switched-path PE4-to-PE2 {
 to 10.9.9.5;
 }
 label-switched-path PE4-to-PE3 {
 to 10.9.9.6;
 }
 label-switched-path PE4-to-P1 {
 to 10.9.9.2;
 }
 label-switched-path PE4-to-P2 {
 to 10.9.9.3;
 }
}
```

```
}
interface fe-1/2/0.7;
interface fe-1/2/1.9;
interface fe-1/2/2.10;
interface fe-1/0/2.12;
}
bgp {
 export [next-hop aigp];
 group internal {
 type internal;
 local-address 10.9.9.4;
 family inet {
 labeled-unicast {
 aigp;
 }
 }
 neighbor 10.9.9.1;
 neighbor 10.9.9.3;
 neighbor 10.9.9.2;
 }
 group external {
 type external;
 multihop {
 ttl 2;
 }
 local-address 10.9.9.4;
 family inet {
 labeled-unicast {
 aigp;
 }
 }
 peer-as 7018;
 neighbor 10.9.9.5;
 neighbor 10.9.9.6;
 }
}
ospf {
 area 0.0.0.0 {
 interface fe-1/2/1.9 {
 metric 1;
 }
 interface fe-1/2/0.7 {
 metric 1;
 }
 interface 10.9.9.4 {
 passive;
 metric 1;
 }
 interface 10.100.1.4 {
 passive;
 metric 1;
 }
 }
 area 0.0.0.2 {
 interface fe-1/2/2.10 {
 metric 1;
 }
 }
}
```

```
 }
 }
 area 0.0.0.3 {
 interface fe-1/0/2.12 {
 metric 1;
 }
 }
}

user@PE4# show routing-options
static {
 route 44.0.0.0/24 discard;
}
router-id 10.9.9.4;
autonomous-system 13979;
```

### *Configuring Device PE1*

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device PE1:

1. Configure the interfaces.

```
[edit interfaces]
user@PE1# set fe-1/2/0 unit 0 description PE1-to-P1
user@PE1# set fe-1/2/0 unit 0 family inet address 10.0.0.1/30
user@PE1# set fe-1/2/0 unit 0 family mpls
user@PE1# set fe-1/2/1 unit 2 description PE1-to-P2
user@PE1# set fe-1/2/1 unit 2 family inet address 10.0.0.5/30
user@PE1# set fe-1/2/1 unit 2 family mpls
user@PE1# set fe-1/2/2 unit 14 description PE1-to-PE7
user@PE1# set fe-1/2/2 unit 14 family inet address 10.0.0.9/30
user@PE1# set lo0 unit 1 family inet address 10.9.9.1/32
user@PE1# set lo0 unit 1 family inet address 10.100.1.1/32
```

2. Configure MPLS and a signaling protocol, such as RSVP or LDP.

```
[edit protocols]
user@PE1# set rsvp interface fe-1/2/0.0
user@PE1# set rsvp interface fe-1/2/1.2
user@PE1# set rsvp interface fe-1/2/2.14
user@PE1# set mpls label-switched-path PE1-to-P1 to 10.9.9.2
user@PE1# set mpls label-switched-path PE1-to-P2 to 10.9.9.3
user@PE1# set mpls interface fe-1/2/0.0
user@PE1# set mpls interface fe-1/2/1.2
user@PE1# set mpls interface fe-1/2/2.14
```

3. Configure BGP.

```
[edit protocols bgp]
user@PE1# set group internal type internal
user@PE1# set group internal local-address 10.9.9.1
user@PE1# set group internal export SET_EXPORT_ROUTES
user@PE1# set group internal vpn-apply-export
```



```

user@PE1# set group internal neighbor 10.9.9.4
user@PE1# set group internal neighbor 10.9.9.2
user@PE1# set group internal neighbor 10.9.9.3
user@PE1# set group external type external
user@PE1# set group external export SET_EXPORT_ROUTES
user@PE1# set group external peer-as 7019
user@PE1# set group external neighbor 10.0.0.10

```

4. Enable AIGP.

```

[edit protocols bgp]
user@PE1# set group internal family inet labeled-unicast aigp
user@PE1# set group external family inet labeled-unicast aigp

```

5. Enable the policies.

```

[edit policy-options policy-statement SET_EXPORT_ROUTES term 10]
user@PE1# set from protocol direct
user@PE1# set from protocol bgp
user@PE1# set then next-hop 10.100.1.1
user@PE1# set then accept

```

6. Configure an IGP, such as OSPF, RIP, or IS-IS.

```

[edit protocols ospf area 0.0.0.1]
user@PE1# set interface fe-1/2/0.0 metric 1
user@PE1# set interface fe-1/2/1.2 metric 1
user@PE1# set interface 10.9.9.1 passive
user@PE1# set interface 10.9.9.1 metric 1
user@PE1# set interface 10.100.1.1 passive
user@PE1# set interface 10.100.1.1 metric 1

```

7. Configure the router ID and the autonomous system number.

```

[edit routing-options]
user@PE1# set router-id 10.9.9.1
user@PE1# set autonomous-system 13979

```

8. If you are done configuring the device, commit the configuration.

```

user@PE1# commit

```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@PE1# show interfaces
fe-1/2/0 {
 unit 0 {
 description PE1-to-P1;
 family inet {
 address 10.0.0.1/30;
 }
 family mpls;
 }
}
fe-1/2/1 {

```

```
 unit 2 {
 description PE1-to-P2;
 family inet {
 address 10.0.0.5/30;
 }
 family mpls;
 }
 }
 fe-1/2/2 {
 unit 14 {
 description PE1-to-PE7;
 family inet {
 address 10.0.0.9/30;
 }
 }
 }
}
lo0 {
 unit 1 {
 family inet {
 address 10.9.9.1/32;
 address 10.100.1.1/32;
 }
 }
}

user@PE1# show policy-options
policy-statement SET_EXPORT_ROUTES {
 term 10 {
 from protocol [direct bgp];
 then {
 next-hop 10.100.1.1;
 accept;
 }
 }
}

user@PE1# show protocols
rsvp {
 interface fe-1/2/0.0;
 interface fe-1/2/1.2;
 interface fe-1/2/2.14;
}
mpls {
 label-switched-path PE1-to-P1 {
 to 10.9.9.2;
 }
 label-switched-path PE1-to-P2 {
 to 10.9.9.3;
 }
 interface fe-1/2/0.0;
 interface fe-1/2/1.2;
 interface fe-1/2/2.14;
}
bgp {
 group internal {
 type internal;
 local-address 10.9.9.1;
```

```
family inet {
 labeled-unicast {
 aigp;
 }
}
export SET_EXPORT_ROUTES;
vpn-apply-export;
neighbor 10.9.9.4;
neighbor 10.9.9.2;
neighbor 10.9.9.3;
}
group external {
 type external;
 family inet {
 labeled-unicast {
 aigp;
 }
 }
 export SET_EXPORT_ROUTES;
 peer-as 7019;
 neighbor 10.0.0.10;
}
}
ospf {
 area 0.0.0.1 {
 interface fe-1/2/0.0 {
 metric 1;
 }
 interface fe-1/2/1.2 {
 metric 1;
 }
 interface 10.9.9.1 {
 passive;
 metric 1;
 }
 interface 10.100.1.1 {
 passive;
 metric 1;
 }
 }
}
}
```

```
user@PE1# show routing-options
router-id 10.9.9.1;
autonomous-system 13979;
```

### Configuring Device PE2

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device PE2:

1. Configure the interfaces.

```
[edit interfaces]
user@PE2# set fe-1/2/0 unit 11 description PE2-to-PE4
user@PE2# set fe-1/2/0 unit 11 family inet address 10.0.0.22/30
user@PE2# set fe-1/2/0 unit 11 family mpls
user@PE2# set lo0 unit 9 family inet address 10.9.9.5/32 primary
user@PE2# set lo0 unit 9 family inet address 10.100.1.5/32
```

2. Configure MPLS and a signaling protocol, such as RSVP or LDP.

```
[edit protocols]
user@PE2# set rsvp interface fe-1/2/0.11
user@PE2# set mpls label-switched-path PE2-to-PE4 to 10.9.9.4
user@PE2# set mpls interface fe-1/2/0.11
```

3. Configure BGP.

```
[edit protocols bgp]
user@PE2# set group external type external
user@PE2# set group external multihop ttl 2
user@PE2# set group external local-address 10.9.9.5
user@PE2# set group external export next-hop
user@PE2# set group external export aigp
user@PE2# set group external export SET_EXPORT_ROUTES
user@PE2# set group external vpn-apply-export
user@PE2# set group external peer-as 13979
user@PE2# set group external neighbor 10.9.9.4
```

4. Enable AIGP.

```
[edit protocols bgp]
user@PE2# set group external family inet labeled-unicast aigp
```

5. Originate a prefix, and configure an AIGP distance.

By default, a prefix is originated using the current IGP distance. Optionally, you can configure a distance for the AIGP attribute, using the **distance** option, as shown here.

```
[edit policy-options policy-statement aigp]
user@PE2# set term 10 from route-filter 55.0.0.0/24 exact
user@PE2# set term 10 then aigp-originate distance 20
user@PE2# set term 10 then next-hop 10.100.1.5
user@PE2# set term 10 then accept
user@PE2# set term 20 from route-filter 99.0.0.0/24 exact
user@PE2# set term 20 then aigp-originate distance 30
user@PE2# set term 20 then next-hop 10.100.1.5
user@PE2# set term 20 then accept
```

6. Enable the policies.

```
[edit policy-options]
user@PE2# set policy-statement SET_EXPORT_ROUTES term 10 from protocol
direct
user@PE2# set policy-statement SET_EXPORT_ROUTES term 10 from protocol
static
user@PE2# set policy-statement SET_EXPORT_ROUTES term 10 from protocol
bgp
user@PE2# set policy-statement SET_EXPORT_ROUTES term 10 then next-hop
10.100.1.5
user@PE2# set policy-statement SET_EXPORT_ROUTES term 10 then accept
user@PE2# set policy-statement next-hop term 10 from protocol bgp
user@PE2# set policy-statement next-hop term 10 then next-hop 10.100.1.5
user@PE2# set policy-statement next-hop term 10 then accept
user@PE2# set policy-statement next-hop term 20 from protocol direct
user@PE2# set policy-statement next-hop term 20 from route-filter 10.9.9.5/32
exact
user@PE2# set policy-statement next-hop term 20 from route-filter 10.100.1.5/32
exact
user@PE2# set policy-statement next-hop term 20 then next-hop 10.100.1.5
user@PE2# set policy-statement next-hop term 20 then accept
```

7. Enable some static routes.

```
[edit routing-options]
user@PE2# set static route 99.0.0.0/24 discard
user@PE2# set static route 55.0.0.0/24 discard
```

8. Configure an IGP, such as OSPF, RIP, or IS-IS.

```
[edit protocols ospf area 0.0.0.2]
user@PE2# set interface 10.9.9.5 passive
user@PE2# set interface 10.9.9.5 metric 1
user@PE2# set interface 10.100.1.5 passive
user@PE2# set interface 10.100.1.5 metric 1
user@PE2# set interface fe-1/2/0.11 metric 1
```

9. Configure the router ID and the autonomous system number.

```
[edit routing-options]
user@PE2# set router-id 10.9.9.5
user@PE2# set autonomous-system 7018
```

10. If you are done configuring the device, commit the configuration.

```
user@PE2# commit
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PE2# show interfaces
fe-1/2/0 {
 unit 11 {
 description PE2-to-PE4;
```

```
 family inet {
 address 10.0.0.22/30;
 }
 family mpls;
 }
}
lo0 {
 unit 9 {
 family inet {
 address 10.9.9.5/32 {
 primary;
 }
 address 10.100.1.5/32;
 }
 }
}

user@PE2# show policy-options
policy-statement SET_EXPORT_ROUTES {
 term 10 {
 from protocol [direct static bgp];
 then {
 next-hop 10.100.1.5;
 accept;
 }
 }
}

policy-statement aigp {
 term 10 {
 from {
 route-filter 55.0.0.0/24 exact;
 }
 then {
 aigp-originate distance 20;
 next-hop 10.100.1.5;
 accept;
 }
 }
 term 20 {
 from {
 route-filter 99.0.0.0/24 exact;
 }
 then {
 aigp-originate distance 30;
 next-hop 10.100.1.5;
 accept;
 }
 }
}

policy-statement next-hop {
 term 10 {
 from protocol bgp;
 then {
 next-hop 10.100.1.5;
 accept;
 }
 }
}
```

```

 }
 term 20 {
 from {
 protocol direct;
 route-filter 10.9.9.5/32 exact;
 route-filter 10.100.1.5/32 exact;
 }
 then {
 next-hop 10.100.1.5;
 accept;
 }
 }
}

user@PE2# show protocols
rsvp {
 interface fe-1/2/0.11;
}
mpls {
 label-switched-path PE2-to-PE4 {
 to 10.9.9.4;
 }
 interface fe-1/2/0.11;
}
bgp {
 group external {
 type external;
 multihop {
 ttl 2;
 }
 local-address 10.9.9.5;
 family inet {
 labeled-unicast {
 aigp;
 }
 }
 export [next-hop aigp SET_EXPORT_ROUTES];
 vpn-apply-export;
 peer-as 13979;
 neighbor 10.9.9.4;
 }
}
ospf {
 area 0.0.0.2 {
 interface 10.9.9.5 {
 passive;
 metric 1;
 }
 interface 10.100.1.5 {
 passive;
 metric 1;
 }
 interface fe-1/2/0.11 {
 metric 1;
 }
 }
}

```

```
}
user@PE2# show routing-options
static {
 route 99.0.0.0/24 discard;
 route 55.0.0.0/24 discard;
}
router-id 10.9.9.5;
autonomous-system 7018;
```

### ***Configuring Device PE3***

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device PE3:

1. Configure the interfaces.

```
[edit interfaces]
user@PE3# set fe-1/2/0 unit 13 description PE3-to-PE4
user@PE3# set fe-1/2/0 unit 13 family inet address 10.0.0.26/30
user@PE3# set fe-1/2/0 unit 13 family mpls
user@PE3# set lo0 unit 11 family inet address 10.9.9.6/32
user@PE3# set lo0 unit 11 family inet address 10.100.1.6/32
```

2. Configure MPLS and a signaling protocol, such as RSVP or LDP.

```
[edit protocols]
user@PE3# set rsvp interface fe-1/2/0.13
user@PE3# set mpls label-switched-path PE3-to-PE4 to 10.9.9.4
user@PE3# set mpls interface fe-1/2/0.13
```

3. Configure BGP.

```
[edit protocols bgp group external]
user@PE3# set type external
user@PE3# set multihop ttl 2
user@PE3# set local-address 10.9.9.6
user@PE3# set export next-hop
user@PE3# set export SET_EXPORT_ROUTES
user@PE3# set vpn-apply-export
user@PE3# set peer-as 13979
user@PE3# set neighbor 10.9.9.4
```

4. Enable AIGP.

```
[edit protocols bgp group external]
user@PE3# set family inet labeled-unicast aigp
```

5. Enable the policies.

```
[edit policy-options]
user@PE3# set policy-statement SET_EXPORT_ROUTES term 10 from protocol
 direct
user@PE3# set policy-statement SET_EXPORT_ROUTES term 10 from protocol
 static
```



```

user@PE3# set policy-statement SET_EXPORT_ROUTES term 10 from protocol
bgp
user@PE3# set policy-statement SET_EXPORT_ROUTES term 10 then next-hop
10.100.1.6
user@PE3# set policy-statement SET_EXPORT_ROUTES term 10 then accept
user@PE3# set policy-statement next-hop term 10 from protocol bgp
user@PE3# set policy-statement next-hop term 10 then next-hop 10.100.1.6
user@PE3# set policy-statement next-hop term 10 then accept
user@PE3# set policy-statement next-hop term 20 from protocol direct
user@PE3# set policy-statement next-hop term 20 from route-filter 10.9.9.6/32
exact
user@PE3# set policy-statement next-hop term 20 from route-filter 10.100.1.6/32
exact
user@PE3# set policy-statement next-hop term 20 then next-hop 10.100.1.6
user@PE3# set policy-statement next-hop term 20 then accept

```

6. Configure an IGP, such as OSPF, RIP, or IS-IS.

```

[edit protocols ospf area 0.0.0.3]
user@PE3# set interface 10.9.9.6 passive
user@PE3# set interface 10.9.9.6 metric 1
user@PE3# set interface 10.100.1.6 passive
user@PE3# set interface 10.100.1.6 metric 1
user@PE3# set interface fe-1/2/0.13 metric 1

```

7. Configure the router ID and the autonomous system number.

```

[edit routing-options]
user@PE3# set router-id 10.9.9.6
user@PE3# set autonomous-system 7018

```

8. If you are done configuring the device, commit the configuration.

```

user@PE3# commit

```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@PE3# show interfaces
fe-1/2/0 {
 unit 13 {
 description PE3-to-PE4;
 family inet {
 address 10.0.0.26/30;
 }
 family mpls;
 }
}
lo0 {
 unit 11 {
 family inet {
 address 10.9.9.6/32;
 address 10.100.1.6/32;
 }
 }
}

```

```
 }
 }
user@PE3# show policy-options
policy-statement SET_EXPORT_ROUTES {
 term 10 {
 from protocol [direct static bgp];
 then {
 next-hop 10.100.1.6;
 accept;
 }
 }
}
policy-statement next-hop {
 term 10 {
 from protocol bgp;
 then {
 next-hop 10.100.1.6;
 accept;
 }
 }
 term 20 {
 from {
 protocol direct;
 route-filter 10.9.9.6/32 exact;
 route-filter 10.100.1.6/32 exact;
 }
 then {
 next-hop 10.100.1.6;
 accept;
 }
 }
}
user@PE3# show protocols
rsvp {
 interface fe-1/2/0.13;
}
mpls {
 label-switched-path PE3-to-PE4 {
 to 10.9.9.4;
 }
 interface fe-1/2/0.13;
}
bgp {
 group external {
 type external;
 multihop {
 ttl 2;
 }
 local-address 10.9.9.6;
 family inet {
 labeled-unicast {
 aigp;
 }
 }
 }
 export [next-hop SET_EXPORT_ROUTES];
```

```

 vpn-apply-export;
 peer-as 13979;
 neighbor 10.9.9.4;
 }
}
ospf {
 area 0.0.0.3 {
 interface 10.9.9.6 {
 passive;
 metric 1;
 }
 interface 10.100.1.6 {
 passive;
 metric 1;
 }
 interface fe-1/2/0.13 {
 metric 1;
 }
 }
}

user@PE3# show routing-options
router-id 10.9.9.6;
autonomous-system 7018;

```

### Configuring Device PE7

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device PE7:

1. Configure the interfaces.

```

[edit interfaces]
user@PE7# set fe-1/2/0 unit 15 description PE7-to-PE1
user@PE7# set fe-1/2/0 unit 15 family inet address 10.0.0.10/30
user@PE7# set lo0 unit 13 family inet address 10.9.9.7/32
user@PE7# set lo0 unit 13 family inet address 10.100.1.7/32

```

2. Configure BGP.

```

[edit protocols bgp group external]
user@PE7# set type external
user@PE7# set export SET_EXPORT_ROUTES
user@PE7# set peer-as 13979
user@PE7# set neighbor 10.0.0.9

```

3. Enable AIGP.

```

[edit protocols bgp group external]
user@PE7# set family inet labeled-unicast aigp

```

4. Configure the routing policy.

```

[edit policy-options policy-statement SET_EXPORT_ROUTES term 10]
user@PE7# set from protocol direct

```

```
user@PE7# set from protocol bgp
user@PE7# set then next-hop 10.100.1.7
user@PE7# set then accept
```

5. Configure the router ID and the autonomous system number.

```
[edit routing-options]
user@PE7# set router-id 10.9.9.7
user@PE7# set autonomous-system 7019
```

6. If you are done configuring the device, commit the configuration.

```
user@PE7# commit
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PE7# show interfaces
interfaces {
 fe-1/2/0 {
 unit 15 {
 description PE7-to-PE1;
 family inet {
 address 10.0.0.10/30;
 }
 }
 }
 lo0 {
 unit 13 {
 family inet {
 address 10.9.9.7/32;
 address 10.100.1.7/32;
 }
 }
 }
}

user@PE7# show policy-options
policy-statement SET_EXPORT_ROUTES {
 term 10 {
 from protocol [direct bgp];
 then {
 next-hop 10.100.1.7;
 accept;
 }
 }
}

user@PE7# show protocols
bgp {
 group external {
 type external;
 family inet {
 labeled-unicast {
 aigp;
 }
 }
 }
}
```

```

 }
 }
 export SET_EXPORT_ROUTES;
 peer-as 13979;
 neighbor 10.0.0.9;
}
}

user@PE7# show routing-options
router-id 10.9.9.7;
autonomous-system 7019;

```

### Verification

Confirm that the configuration is working properly.

- [Verifying That Device PE4 Is Receiving the AIGP Attribute from Its EBGp Neighbor PE2 on page 3757](#)
- [Checking the IGP Metric on page 3757](#)
- [Verifying That Device PE4 Adds the IGP Metric to the AIGP Attribute on page 3758](#)
- [Verifying That Device PE7 Is Receiving the AIGP Attribute from Its EBGp Neighbor PE1 on page 3758](#)
- [Verifying the Resolving AIGP Metric on page 3759](#)
- [Verifying the Presence of AIGP Attributes in BGP Updates on page 3762](#)

#### *Verifying That Device PE4 Is Receiving the AIGP Attribute from Its EBGp Neighbor PE2*

**Purpose** Make sure that the AIGP policy on Device PE2 is working.

**Action** user@PE4> show route receive-protocol bgp 10.9.9.5 extensive

```

* 55.0.0.0/24 (1 entry, 1 announced)
 Accepted
 Route Label: 299888
 Nexthop: 10.100.1.5
 AS path: 7018 I
 AIGP: 20

* 99.0.0.0/24 (1 entry, 1 announced)
 Accepted
 Route Label: 299888
 Nexthop: 10.100.1.5
 AS path: 7018 I
 AIGP: 30

```

**Meaning** On Device PE2, the **aigp-originate** statement is configured with a distance of 20 (**aigp-originate distance 20**). This statement is applied to route 55.0.0.0/24. Likewise, the **aigp-originate distance 30** statement is applied to route 99.0.0.0/24. Thus, when Device PE4 receives these routes, the AIGP attribute is attached with the configured metrics.

#### *Checking the IGP Metric*

**Purpose** From Device PE4, check the IGP metric to the BGP next hop 10.100.1.5.

**Action** user@PE4> show route 10.100.1.5  
inet.0: 30 destinations, 40 routes (30 active, 0 holddown, 0 hidden)  
+ = Active Route, - = Last Active, \* = Both

10.100.1.5/32      \*[OSPF/10] 05:35:50, metric 2  
                         > to 10.0.0.22 via fe-1/2/2.10  
                         [BGP/170] 03:45:07, localpref 100, from 10.9.9.5  
                         AS path: 7018 I  
                         > to 10.0.0.22 via fe-1/2/2.10

**Meaning** The IGP metric for this route is 2.

***Verifying That Device PE4 Adds the IGP Metric to the AIGP Attribute***

**Purpose** Make sure that Device PE4 adds the IGP metric to the AIGP attribute when it readvertises routes to its IBGP neighbor, Device PE1.

**Action** user@PE4> show route advertising-protocol bgp 10.9.9.1 extensive

\* 55.0.0.0/24 (1 entry, 1 announced)  
BGP group internal type Internal  
Route Label: 300544  
Nexthop: 10.100.1.4  
Flags: Nexthop Change  
Localpref: 100  
AS path: [13979] 7018 I  
AIGP: 22

\* 99.0.0.0/24 (1 entry, 1 announced)  
BGP group internal type Internal  
Route Label: 300544  
Nexthop: 10.100.1.4  
Flags: Nexthop Change  
Localpref: 100  
AS path: [13979] 7018 I  
AIGP: 32

**Meaning** The IGP metric is added to the AIGP metric ( $20 + 2 = 22$  and  $30 + 2 = 32$ ), because the next hop is changed for these routes.

***Verifying That Device PE7 Is Receiving the AIGP Attribute from Its EBGP Neighbor PE1***

**Purpose** Make sure that the AIGP policy on Device PE1 is working.

**Action** user@PE7> show route receive-protocol bgp 10.0.0.9 extensive

\* 44.0.0.0/24 (1 entry, 1 announced)

Accepted  
Route Label: 300096  
Nexthop: 10.0.0.9  
AS path: 13979 I  
AIGP: 203

\* 55.0.0.0/24 (1 entry, 1 announced)

Accepted  
Route Label: 300112  
Nexthop: 10.0.0.9  
AS path: 13979 7018 I  
AIGP: 25

\* 99.0.0.0/24 (1 entry, 1 announced)

Accepted  
Route Label: 300112  
Nexthop: 10.0.0.9  
AS path: 13979 7018 I  
AIGP: 35

**Meaning** The 44.0.0.0/24 route is originated at Device PE4. The 55.0.0.0/24 and 99.0.0.0/24 routes are originated at Device PE2. The IGP distances are added to the configured AIGP distances.

#### *Verifying the Resolving AIGP Metric*

**Purpose** Confirm that if the prefix is resolved through recursion and the recursive next hops have AIGP metrics, the prefix has the sum of the AIGP values that are on the recursive BGP next hops.

**Action** 1. Add a static route to 66.0.0.0/24.

```
[edit routing-options]
user@PE2# set static route 66.0.0.0/24 discard
```

2. Delete the existing terms in the **aigp** policy statement on Device PE2.

```
[edit policy-options policy-statement aigp]
user@PE2# delete term 10
user@PE2# delete term 20
```

3. Configure a recursive route lookup for the route to 66.0.0.0.

The policy shows the AIGP metric for prefix 66.0.0.0/24 (none) and its recursive next hop. Prefix 66.0.0.0/24 is resolved by 55.0.0.1. Prefix 66.0.0.0/24 does not have its own AIGP metric being originated, but its recursive next hop, 55.0.0.1, has an AIGP value.

```
[edit policy-options policy-statement aigp]
user@PE2# set term 10 from route-filter 55.0.0.1/24 exact
user@PE2# set term 10 then aigp-originate distance 20
user@PE2# set term 10 then next-hop 10.100.1.5
user@PE2# set term 10 then accept
user@PE2# set term 20 from route-filter 66.0.0.0/24 exact
user@PE2# set term 20 then next-hop 55.0.0.1
```

**user@PE2# set term 20 then accept**

4. On Device PE4, run the **show route 55.0.0.0 extensive** command.

The value of Metric2 is the IGP metric to the BGP next hop. When Device PE4 readvertises these routes to its IBGP peer, Device PE1, the AIGP metric is the sum of AIGP + its Resolving AIGP metric + Metric2.

Prefix 55.0.0.0 shows its own IGP metric 20, as defined and advertised by Device PE2. It does not show a resolving AIGP value because it does not have a recursive BGP next hop. The value of Metric2 is 2.

```
user@PE4> show route 55.0.0.0 extensive
inet.0: 31 destinations, 41 routes (31 active, 0 holddown, 0 hidden)
55.0.0.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 55.0.0.0/24 -> {indirect(262151)}
Page 0 idx 0 Type 1 val 928d1b8
 Flags: Nexthop Change
 Nexthop: 10.100.1.4
 Localpref: 100
 AS path: [13979] 7018 I
 Communities:
 AIGP: 22
Path 55.0.0.0 from 10.9.9.5 Vector len 4. Val: 0
 *BGP Preference: 170/-101
 Next hop type: Indirect
 Address: 0x925da38
 Next-hop reference count: 4
 Source: 10.9.9.5
 Next hop type: Router, Next hop index: 1004
 Next hop: 10.0.0.22 via fe-1/2/2.10, selected
 Label operation: Push 299888
 Label TTL action: prop-ttl
 Protocol next hop: 10.100.1.5
 Push 299888
 Indirect next hop: 93514d8 262151
 State: <Active Ext>
 Local AS: 13979 Peer AS: 7018
 Age: 22:03:26 Metric2: 2
 AIGP: 20
 Task: BGP_7018.10.9.9.5+58560
 Announcement bits (3): 3-KRT 4-BGP_RT_Background 5-Resolve tree 1
 AS path: 7018 I
 Accepted
 Route Label: 299888
 Localpref: 100
 Router ID: 10.9.9.5
 Indirect next hops: 1
 Protocol next hop: 10.100.1.5 Metric: 2
 Push 299888
 Indirect next hop: 93514d8 262151
 Indirect path forwarding next hops: 1
 Next hop type: Router
 Next hop: 10.0.0.22 via fe-1/2/2.10
 10.100.1.5/32 Originating RIB: inet.0
 Metric: 2 Node path count: 1
 Forwarding nexthops: 1
 Nexthop: 10.0.0.22 via fe-1/2/2.10
```

5. On Device PE4, run the **show route 66.0.0.0 extensive** command.



Prefix 66.0.0.0/24 shows the Resolving AIGP, which is the sum of its own AIGP metric and its recursive BGP next hop:

66.0.0.1 = 0, 55.0.0.1 = 20, 0+20 = 20

```

user@PE4> show route 66.0.0.0 extensive
inet.0: 31 destinations, 41 routes (31 active, 0 holddown, 0 hidden)
66.0.0.0/24 (1 entry, 1 announced)
TSI:
KRT in-kerne1 66.0.0.0/24 -> {indirect(262162)}
Page 0 idx 0 Type 1 val 928cefc
 Flags: Nexthop Change
 Nexthop: 10.100.1.4
 Localpref: 100
 AS path: [13979] 7018 I
 Communities:
Path 66.0.0.0 from 10.9.9.5 Vector len 4. Val: 0
 *BGP Preference: 170/-101
 Next hop type: Indirect
 Address: 0x925d4e0
 Next-hop reference count: 4
 Source: 10.9.9.5
 Next hop type: Router, Next hop index: 1006
 Next hop: 10.0.0.22 via fe-1/2/2.10, selected
 Label operation: Push 299888, Push 299888(top)
 Label TTL action: prop-ttl, prop-ttl(top)
 Protocol next hop: 55.0.0.1
 Push 299888
 Indirect next hop: 9353e88 262162
 State: <Active Ext>
 Local AS: 13979 Peer AS: 7018
 Age: 31:42 Metric2: 2
 Resolving-AIGP: 20
 Task: BGP_7018.10.9.9.5+58560
 Announcement bits (3): 3-KRT 4-BGP_RT_Background 5-Resolve tree 1
 AS path: 7018 I
 Accepted
 Route Label: 299888
 Localpref: 100
 Router ID: 10.9.9.5
 Indirect next hops: 1
 Protocol next hop: 55.0.0.1 Metric: 2 AIGP: 20
 Push 299888
 Indirect next hop: 9353e88 262162
 Indirect path forwarding next hops: 1
 Next hop type: Router
 Next hop: 10.0.0.22 via fe-1/2/2.10
 55.0.0.0/24 Originating RIB: inet.0
 Metric: 2 Node path count: 1
 Indirect nexthops: 1
 Protocol Nexthop: 10.100.1.5 Metric: 2 Push 299888
 Indirect nexthop: 93514d8 262151
 Indirect path forwarding nexthops: 1
 Nexthop: 10.0.0.22 via fe-1/2/2.10
 10.100.1.5/32 Originating RIB: inet.0
 Metric: 2 Node path count: 1
 Forwarding nexthops: 1
 Nexthop: 10.0.0.22 via fe-1/2/2.10

```

*Verifying the Presence of AIGP Attributes in BGP Updates*

**Purpose** If the AIGP attribute is not enabled under BGP (or the **group** or **neighbor** hierarchies), the AIGP attribute is silently discarded. Enable **traceoptions** and include the **packets** flag in the **detail** option in the configuration to confirm the presence of the AIGP attribute in transmitted or received BGP updates. This is useful when debugging AIGP issues.

**Action** 1. Configure Device PE2 and Device PE4 for **traceoptions**.

```
user@host> show protocols bgp
 traceoptions {
 file bgp size 1m files 5;
 flag packets detail;
 }
```

2. Check the **traceoptions** file on Device PE2.

The following sample shows Device PE2 advertising prefix 99.0.0.0/24 to Device PE4 (10.9.9.4) with an AIGP metric of 20:

```
user@PE2> show log bgp
Mar 22 09:27:18.982150 BGP SEND 10.9.9.5+49652 -> 10.9.9.4+179
Mar 22 09:27:18.982178 BGP SEND message type 2 (Update) length 70
Mar 22 09:27:18.982198 BGP SEND Update PDU length 70
Mar 22 09:27:18.982248 BGP SEND flags 0x40 code Origin(1): IGP
Mar 22 09:27:18.982273 BGP SEND flags 0x40 code ASPath(2) length 6: 7018
Mar 22 09:27:18.982295 BGP SEND flags 0x80 code AIGP(26): AIGP: 20
Mar 22 09:27:18.982316 BGP SEND flags 0x90 code MP_reach(14): AFI/SAFI 1/4
Mar 22 09:27:18.982341 BGP SEND nhop 10.100.1.5 len 4
Mar 22 09:27:18.982372 BGP SEND 99.0.0.0/24 (label 301664)
Mar 22 09:27:33.665412 bgp_send: sending 19 bytes to abcd::10:255:170:84
(External AS 13979)
```

3. Verify that the route was received on Device PE4 using the **show route receive-protocol** command.

AIGP is not enabled on Device PE4, so the AIGP attribute is silently discarded for prefix 99.0.0.0/24 and does not appear in the following output:

```
user@PE4> show route receive-protocol bgp 10.9.9.5 extensive | find 55.0.0.0
* 99.0.0.0/24 (2 entries, 1 announced)
 Accepted
 Route Label: 301728
 Nexthop: 10.100.1.5
 AS path: 7018 I
```

4. Check the **traceoptions** file on Device PE4.

The following output from the **traceoptions** log shows that the 99.0.0.0/24 prefix was received with the AIGP attribute attached:

```
user@PE4> show log bgp
Mar 22 09:41:39.650295 BGP RECV 10.9.9.5+64690 -> 10.9.9.4+179
Mar 22 09:41:39.650331 BGP RECV message type 2 (Update) length 70
Mar 22 09:41:39.650350 BGP RECV Update PDU length 70
Mar 22 09:41:39.650370 BGP RECV flags 0x40 code Origin(1): IGP
Mar 22 09:41:39.650394 BGP RECV flags 0x40 code ASPath(2) length 6: 7018
Mar 22 09:41:39.650415 BGP RECV flags 0x80 code AIGP(26): AIGP: 20
Mar 22 09:41:39.650436 BGP RECV flags 0x90 code MP_reach(14): AFI/SAFI 1/4
Mar 22 09:41:39.650459 BGP RECV nhop 10.100.1.5 len 4
```

```
Mar 22 09:41:39.650495 BGP RECV 99.0.0.0/24 (label 301728)
Mar 22 09:41:39.650574 bgp_rcv_nlri: 99.0.0.0/24
Mar 22 09:41:39.650607 bgp_rcv_nlri: 99.0.0.0/24 belongs to meshgroup
Mar 22 09:41:39.650629 bgp_rcv_nlri: 99.0.0.0/24 qualified bnp->ribact 0x0
12afcb 0x0
```

**Meaning** Performing this verification helps with AIGP troubleshooting and debugging issues. It enables you to verify which devices in your network send and receive AIGP attributes.

- Related Documentation**
- [Understanding BGP Path Selection on page 3793](#)
  - [Examples: Configuring Internal BGP Peering on page 3624](#)



# BGP Policy Configuration

- [Example: Configuring BGP Interactions with IGPs on page 3765](#)
- [Example: Configuring BGP Route Advertisement on page 3769](#)
- [Example: Configuring EBGP Multihop on page 3777](#)
- [Example: Configuring BGP Route Preference \(Administrative Distance\) on page 3786](#)
- [Example: Configuring BGP Path Selection on page 3793](#)
- [Example: Removing Private AS Numbers on page 3803](#)

## Example: Configuring BGP Interactions with IGPs

---

- [Understanding Routing Policies on page 3765](#)
- [Example: Injecting OSPF Routes into the BGP Routing Table on page 3766](#)

### Understanding Routing Policies

Each routing policy is identified by a policy name. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose the entire name in double quotation marks. Each routing policy name must be unique within a configuration.

Once a policy is created and named, it must be applied before it is active. You apply routing policies using the **import** and **export** statements at the **protocols>protocol-name** level in the configuration hierarchy.

In the **import** statement, you list the name of the routing policy to be evaluated when routes are imported into the routing table from the routing protocol.

In the **export** statement, you list the name of the routing policy to be evaluated when routes are being exported from the routing table into a dynamic routing protocol. Only active routes are exported from the routing table.

To specify more than one policy and create a policy chain, you list the policies using a space as a separator. If multiple policies are specified, the policies are evaluated in the order in which they are specified. As soon as an accept or reject action is executed, the policy chain evaluation ends.

## Example: Injecting OSPF Routes into the BGP Routing Table

This example shows how to create a policy that injects OSPF routes into the BGP routing table.

- [Requirements on page 3766](#)
- [Overview on page 3766](#)
- [Configuration on page 3766](#)
- [Verification on page 3768](#)
- [Troubleshooting on page 3768](#)

---

### Requirements

Before you begin:

- Configure network interfaces.
- Configure external peer sessions. See [“Example: Configuring External BGP Point-to-Point Peer Sessions” on page 3602](#).
- Configure interior gateway protocol (IGP) sessions between peers.

---

### Overview

In this example, you create a routing policy called **injectpolicy1** and a routing term called **injectterm1**. The policy injects OSPF routes into the BGP routing table.

---

### Configuration

- [Configuring the Routing Policy on page 3766](#)
- [Configuring Tracing for the Routing Policy on page 3767](#)

#### *Configuring the Routing Policy*

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

```
set policy-options policy-statement injectpolicy1 term injectterm1 from protocol ospf
set policy-options policy-statement injectpolicy1 term injectterm1 from area 0.0.0.1
set policy-options policy-statement injectpolicy1 term injectterm1 then accept
set protocols bgp export injectpolicy1
```

#### Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To inject OSPF routes into a BGP routing table:

1. Create the policy term.

```
[edit policy-options policy-statement injectpolicy1]
```

```
user@host# set term injectterm1
```

2. Specify OSPF as a match condition.

```
[edit policy-options policy-statement injectpolicy1 term injectterm1]
user@host# set from protocol ospf
```

3. Specify the routes from an OSPF area as a match condition.

```
[edit policy-options policy-statement injectpolicy1 term injectterm1]
user@host# set from area 0.0.0.1
```

4. Specify that the route is to be accepted if the previous conditions are matched.

```
[edit policy-options policy-statement injectpolicy1 term injectterm1]
user@host# set then accept
```

5. Apply the routing policy to BGP.

```
[edit]
user@host# set protocols bgp export injectpolicy1
```

**Results** Confirm your configuration by entering the **show policy-options** and **show protocols bgp** commands from configuration mode. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show policy-options
policy-statement injectpolicy1 {
 term injectterm1 {
 from {
 protocol ospf;
 area 0.0.0.1;
 }
 then accept;
 }
}
```

```
user@host# show protocols bgp
export injectpolicy1;
```

If you are done configuring the device, enter **commit** from configuration mode.

### *Configuring Tracing for the Routing Policy*

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

```
set policy-options policy-statement injectpolicy1 term injectterm1 then trace
set routing-options traceoptions file ospf-bgp-policy-log
set routing-options traceoptions file size 5m
set routing-options traceoptions file files 5
set routing-options traceoptions flag policy
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Include a trace action in the policy.

```
[edit policy-options policy-statement injectpolicy] term injectterm1]
user@host# then trace
```

2. Configure the tracing file for the output.

```
[edit routing-options traceoptions]
user@host# set file ospf-bgp-policy-log
user@host# set file size 5m
user@host# set file files 5
user@host# set flag policy
```

**Results** Confirm your configuration by entering the **show policy-options** and **show routing-options** commands from configuration mode. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show policy-options
policy-statement injectpolicy1 {
 term injectterm1 {
 then {
 trace;
 }
 }
}

user@host# show routing-options
traceoptions {
 file ospf-bgp-policy-log size 5m files 5;
 flag policy;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

---

### Verification

Confirm that the configuration is working properly.

#### *Verifying That the Expected BGP Routes Are Present*

**Purpose** Verify the effect of the export policy.

**Action** From operational mode, enter the **show route** command.

---

### Troubleshooting

- [Using the show log Command to Examine the Actions of the Routing Policy on page 3768](#)

#### *Using the show log Command to Examine the Actions of the Routing Policy*

**Problem** The routing table contains unexpected routes, or routes are missing from the routing table.



**Solution** If you configure policy tracing as shown in this example, you can run the **show log ospf-bgp-policy-log** command to diagnose problems with the routing policy. The **show log ospf-bgp-policy-log** command displays information about the routes that the **injectpolicy1** policy term analyzes and acts upon.

**Related Documentation**

- [Understanding External BGP Peering Sessions on page 3601](#)
- [BGP Configuration Overview](#)

## Example: Configuring BGP Route Advertisement

- [Configuring Routing Policies to Control BGP Route Advertisements on page 3769](#)
- [Example: Configuring BGP Prefix-Based Outbound Route Filtering on page 3773](#)

### Configuring Routing Policies to Control BGP Route Advertisements

All routing protocols use the Junos OS routing table to store the routes that they learn and to determine which routes they should advertise in their protocol packets. Routing policy allows you to control which routes the routing protocols store in and retrieve from the routing table. For information about routing policy, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

When configuring BGP routing policy, you can perform the following tasks:

- [Applying Routing Policy on page 3769](#)
- [Setting BGP to Advertise Inactive Routes on page 3770](#)
- [Configuring BGP to Advertise the Best External Route to Internal Peers on page 3770](#)
- [Configuring How Often BGP Exchanges Routes with the Routing Table on page 3772](#)
- [Disabling Suppression of Route Advertisements on page 3773](#)

#### Applying Routing Policy

You define routing policy at the **[edit policy-options]** hierarchy level. To apply policies you have defined for BGP, include the **import** and **export** statements within the BGP configuration.

You can apply policies as follows:

- BGP global **import** and **export** statements—Include these statements at the **[edit protocols bgp]** hierarchy level (for routing instances, include these statements at the **[edit routing-instances routing-instance-name protocols bgp]** hierarchy level).
- Group **import** and **export** statements—Include these statements at the **[edit protocols bgp group group-name]** hierarchy level (for routing instances, include these statements at the **[edit routing-instances routing-instance-name protocols bgp group group-name]** hierarchy level).
- Peer **import** and **export** statements—Include these statements at the **[edit protocols bgp group group-name neighbor address]** hierarchy level (for routing instances, include

these statements at the `[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]` hierarchy level).

A peer-level **import** or **export** statement overrides a group **import** or **export** statement. A group-level **import** or **export** statement overrides a global BGP **import** or **export** statement.

To apply policies, see the following sections:

- [Applying Policies to Routes Being Imported into the Routing Table from BGP on page 3770](#)
- [Applying Policies to Routes Being Exported from the Routing Table into BGP on page 3770](#)

#### ***Applying Policies to Routes Being Imported into the Routing Table from BGP***

To apply policy to routes being imported into the routing table from BGP, include the **import** statement, listing the names of one or more policies to be evaluated:

```
import [policy-names];
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

If you specify more than one policy, they are evaluated in the order specified, from first to last, and the first matching filter is applied to the route. If no match is found, BGP places into the routing table only those routes that were learned from BGP routing devices.

#### ***Applying Policies to Routes Being Exported from the Routing Table into BGP***

To apply policy to routes being exported from the routing table into BGP, include the **export** statement, listing the names of one or more policies to be evaluated:

```
export [policy-names];
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

If you specify more than one policy, they are evaluated in the order specified, from first to last, and the first matching filter is applied to the route. If no routes match the filters, the routing table exports into BGP only the routes that it learned from BGP.

---

#### **Setting BGP to Advertise Inactive Routes**

By default, BGP stores the route information it receives from update messages in the Junos OS routing table, and the routing table exports only active routes into BGP, which BGP then advertises to its peers. To have the routing table export to BGP the best route learned by BGP even if Junos OS did not select it to be an active route, include the **advertise-inactive** statement:

```
advertise-inactive;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

---

#### **Configuring BGP to Advertise the Best External Route to Internal Peers**

In general, deployed BGP implementations do not advertise the external route with the highest local preference value to internal peers unless it is the best route. Although this

behavior was required by an earlier version of the BGP version 4 specification, RFC 1771, it was typically not followed in order to minimize the amount of advertised information and to prevent routing loops. However, there are scenarios in which advertising the best external route is beneficial, in particular, situations that can result in IBGP route oscillation.

In Junos OS Release 9.3 and later, you can configure BGP to advertise the best external route into an internal BGP (IBGP) mesh group, a route reflector cluster, or an autonomous system (AS) confederation, even when the best route is an internal route.



**NOTE:** In order to configure the `advertise-external` statement on a route reflector, you must disable intracluster reflection with the `no-client-reflect` statement.

When a routing device is configured as a route reflector for a cluster, a route advertised by the route reflector is considered internal if it is received from an internal peer with the same cluster identifier or if both peers have no cluster identifier configured. A route received from an internal peer that belongs to another cluster, that is, with a different cluster identifier, is considered external.

In a confederation, when advertising a route to a confederation border router, any route from a different confederation sub-AS is considered external.

You can also configure BGP to advertise the external route only if the route selection process reaches the point where the multiple exit discriminator (MED) metric is evaluated. As a result, an external route with an AS path worse (that is, longer) than that of the active path is not advertised.

Junos OS also provides support for configuring a BGP export policy that matches on the state of an advertised route. You can match on either active or inactive routes. For more information, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

To configure BGP to advertise the best external path to internal peers, include the `advertise-external` statement:

```
advertise-external;
```



**NOTE:** The `advertise-external` statement is supported at both the group and neighbor level. If you configure the statement at the neighbor level, you must configure it for all neighbors in a group. Otherwise, the group is automatically split into different groups.

For a complete list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

To configure BGP to advertise the best external path only if the route selection process reaches the point where the MED value is evaluated, include the `conditional` statement:

```
advertise-external {
```

```
conditional;
}
```

### Configuring How Often BGP Exchanges Routes with the Routing Table

---

BGP stores the route information it receives from update messages in the routing table, and the routing table exports active routes from the routing table into BGP. BGP then advertises the exported routes to its peers. By default, the exchange of route information between BGP and the routing table occurs immediately after the routes are received. This immediate exchange of route information might cause instabilities in the network reachability information. To guard against this, you can delay the time between when BGP and the routing table exchange route information.

To configure how often BGP and the routing table exchange route information, include the **out-delay** statement:

```
out-delay seconds;
```

By default, the routing table retains some of the route information learned from BGP. To have the routing table retain all or none of this information, include the **keep** statement:

```
keep (all | none);
```

For a list of hierarchy levels at which you can include these statements, see the statement summary sections for these statements.

The routing table can retain the route information learned from BGP in one of the following ways:

- Default (omit the **keep** statement)—Keep all route information that was learned from BGP, except for routes whose AS path is looped and whose loop includes the local AS.
- **keep all**—Keep all route information that was learned from BGP.
- **keep none**—Discard routes that were received from a peer and that were rejected by import policy or other sanity checking, such as AS path or next hop. When you configure **keep none** for the BGP session and the inbound policy changes, Junos OS forces readvertisement of the full set of routes advertised by the peer.

In an AS path healing situation, routes with looped paths theoretically could become usable during a soft reconfiguration when the AS path loop limit is changed. However, there is a significant memory usage difference between the default and **keep all**.

Consider the following scenarios:

- A peer readvertises routes back to the peer from which it learned them.

This can happen in the following cases:

- Another vendor's routing device advertises the routes back to the sending peer.
- The Junos OS peer's default behavior of not readvertising routes back to the sending peer is overridden by configuring **advertise-peer-as**.
- A provider edge (PE) routing device discards any VPN route that does not have any of the expected route targets.

When **keep all** is configured, the behavior of discarding routes received in the above scenarios is overridden.

### Disabling Suppression of Route Advertisements

Junos OS does not advertise the routes learned from one EBGP peer back to the same external BGP (EBGP) peer. In addition, the software does not advertise those routes back to any EBGP peers that are in the same AS as the originating peer, regardless of the routing instance. You can modify this behavior by including the **advertise-peer-as** statement in the configuration. To disable the default advertisement suppression, include the **advertise-peer-as** statement:

```
advertise-peer-as;
```



**NOTE:** The route suppression default behavior is disabled if the **as-override** statement is included in the configuration.

If you include the **advertise-peer-as** statement in the configuration, BGP advertises the route regardless of this check.

To restore the default behavior, include the **no-advertise-peer-as** statement in the configuration:

```
no-advertise-peer-as;
```

If you include both the **as-override** and **no-advertise-peer-as** statements in the configuration, the **no-advertise-peer-as** statement is ignored. You can include these statements at multiple hierarchy levels.

For a list of hierarchy levels at which you can include these statements, see the statement summary section for these statements.

### Example: Configuring BGP Prefix-Based Outbound Route Filtering

This example shows how to configure a Juniper Networks router to accept route filters from remote peers and perform outbound route filtering using the received filters.

- [Requirements on page 3773](#)
- [Overview on page 3774](#)
- [Configuration on page 3774](#)
- [Verification on page 3776](#)

#### Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol (IGP).

## Overview

You can configure a BGP peer to accept route filters from remote peers and perform outbound route filtering using the received filters. By filtering out unwanted updates, the sending peer saves resources needed to generate and transmit updates, and the receiving peer saves resources needed to process updates. This feature can be useful, for example, in a virtual private network (VPN) in which subsets of customer edge (CE) devices are not capable of processing all the routes in the VPN. The CE devices can use prefix-based outbound route filtering to communicate to the provider edge (PE) routing device to transmit only a subset of routes, such as routes to the main data centers only.

The maximum number of prefix-based outbound route filters that a BGP peer can accept is 5000. If a remote peer sends more than 5000 outbound route filters to a peer address, the additional filters are discarded, and a system log message is generated.


You can configure interoperability for the routing device as a whole or for specific BGP groups or peers only.

## Topology

In the sample network, Device CE1 is a router from another vendor. The configuration shown in this example is on Juniper Networks Router PE1.

Figure 72 shows the sample network.

Figure 72: BGP Prefix-Based Outbound Route Filtering



## Configuration

### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
PE1 set protocols bgp group cisco-peers type external
 set protocols bgp group cisco-peers description "to CE1"
 set protocols bgp group cisco-peers local-address 192.168.165.58
 set protocols bgp group cisco-peers peer-as 35
 set protocols bgp group cisco-peers outbound-route-filter bgp-orf-cisco-mode
 set protocols bgp group cisco-peers outbound-route-filter prefix-based accept inet
 set protocols bgp group cisco-peers neighbor 192.168.165.56
 set routing-options autonomous-system 65500
```

**Step-by-Step Procedure** The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router PE1 to accept route filters from Device CE1 and perform outbound route filtering using the received filters:

1. Configure the local autonomous system.

```
[edit routing-options]
user@PE1# set autonomous-system 65500
```

2. Configure external peering with Device CE1.

```
[edit protocols bgp group cisco-peers]
user@PE1# set type external
user@PE1# set description "to CE1"
user@PE1# set local-address 192.168.165.58
user@PE1# set peer-as 35
user@PE1# set neighbor 192.168.165.56
```

3. Configure Router PE1 to accept IPv4 route filters from Device CE1 and perform outbound route filtering using the received filters.

```
[edit protocols bgp group cisco-peers]
user@PE1# set outbound-route-filter prefix-based accept inet
```

4. (Optional) Enable interoperability with routing devices that use the vendor-specific compatibility code of 130 for outbound route filters and the code type of 128.

The IANA standard code is 3, and the standard code type is 64.

```
[edit protocols bgp group cisco-peers]
user@PE1# set outbound-route-filter bgp-orf-cisco-mode
```

**Results** From configuration mode, confirm your configuration by entering the **show protocols** and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PE1# show protocols
group cisco-peers {
 type external;
 description "to CE1";
 local-address 192.168.165.58;
 peer-as 35;
 outbound-route-filter {
 bgp-orf-cisco-mode;
 prefix-based {
 accept {
 inet;
 }
 }
 }
 neighbor 192.168.165.56;
}

user@PE1# show routing-options
autonomous-system 65500;
```

If you are done configuring the device, enter **commit** from configuration mode.

### Verification

---

Confirm that the configuration is working properly.

- [Verifying the Outbound Route Filter on page 3776](#)
- [Verifying the BGP Neighbor Mode on page 3776](#)

#### *Verifying the Outbound Route Filter*

**Purpose** Display information about the prefix-based outbound route filter received from Device CE1.

**Action** From operational mode, enter the **show bgp neighbor orf detail** command.

```
user@PE1> show bgp neighbor orf 192.168.165.56 detail
Peer: 192.168.165.56 Type: External
Group: cisco-peers

inet-unicast
Filter updates rcv: 4 Immediate: 0
Filter: prefix-based receive
Updates rcv: 4
Received filter entries:
 seq 10 2.2.0.0/16 deny minlen 0 maxlen 0
 seq 20 3.3.0.0/16 deny minlen 24 maxlen 0
 seq 30 4.4.0.0/16 deny minlen 0 maxlen 28
 seq 40 5.5.0.0/16 deny minlen 24 maxlen 28
```

#### *Verifying the BGP Neighbor Mode*

**Purpose** Verify that the **bgp-orf-cisco-mode** setting is enabled for the peer by making sure that the **ORFCiscoMode** option is displayed in the **show bgp neighbor** command output.

**Action** From operational mode, enter the **show bgp neighbor** command.

```
user@PE1> show bgp neighbor
Peer: 192.168.165.56 AS 35 Local: 192.168.165.58 AS 65500
Type: External State: Active Flags: <>
Last State: Idle Last Event: Start
Last Error: None
Export: [adv_stat]
Options: <Preference LocalAddress AddressFamily PeerAS Refresh>
Options: <ORF ORFCiscoMode>
Address families configured: inet-unicast
Local Address: 192.168.165.58 Holdtime: 90 Preference: 170
Number of flaps: 0
Trace options: detail open detail refresh
Trace file: /var/log/orf size 5242880 files 20
```

**Related Documentation**

- [Understanding External BGP Peering Sessions on page 3601](#)
- [BGP Configuration Overview](#)
- [Example: Configuring a Routing Policy to Advertise the Best External Route to Internal Peers](#)



- [Example: Configuring BGP to Advertise Inactive Routes](#)

---

## Example: Configuring EBGp Multihop

---

- [Understanding BGP Multihop for Peering on page 3777](#)
- [Example: Configuring EBGp Multihop Sessions on page 3777](#)

### Understanding BGP Multihop for Peering

When external BGP (EBGP) peers are not directly connected to each other, they must cross one or more non-BGP routers to reach each other. Configuring multihop EBGp enables the peers to pass through the other routers to form peer relationships and exchange update messages. This type of configuration is typically used when a Juniper Networks routing device needs to run EBGp with a third-party router that does not allow direct connection of the two EBGp peers. EBGp multihop enables a neighbor connection between two EBGp peers that do not have a direct connection.

### Example: Configuring EBGp Multihop Sessions

This example shows how to configure an external BGP (EBGP) peer that is more than one hop away from the local router. This type of session is called a *multihop* BGP session.

- [Requirements on page 3777](#)
- [Overview on page 3777](#)
- [Configuration on page 3778](#)
- [Verification on page 3784](#)

---

#### Requirements

---

No special configuration beyond device initialization is required before you configure this example.

---

#### Overview

---

The configuration to enable multihop EBGp sessions requires connectivity between the two EBGp peers. This example uses static routes to provide connectivity between the devices.


Unlike directly connected EBGp sessions in which physical address are typically used in the **neighbor** statements, you must use loopback interface addresses for multihop EBGp by specifying the loopback interface address of the indirectly connected peer. In this way, EBGp multihop is similar to internal BGP (IBGP).

Finally, you must add the **multihop** statement. Optionally, you can set a maximum time-to-live (TTL) value with the **ttl** statement. The TTL is carried in the IP header of BGP packets. If you do not specify a TTL value, the system's default maximum TTL value is used. The default TTL value is 64 for multihop EBGp sessions. Another option is to retain the BGP next-hop value for route advertisements by including the **no-next-hop-change** statement.

Figure 73 shows a typical EBGp multihop network.

Device C and Device E have an established EBGp session. Device D is not a BGP-enabled device. All of the devices have connectivity via static routes.

Figure 73: Typical Network with EBGp Multihop Sessions



### Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Device C**

```
set interfaces fe-1/2/0 unit 9 description to-D
set interfaces fe-1/2/0 unit 9 family inet address 10.10.10.9/30
set interfaces lo0 unit 3 family inet address 192.168.40.4/32
set protocols bgp group external-peers type external
set protocols bgp group external-peers multihop ttl 2
set protocols bgp group external-peers local-address 192.168.40.4
set protocols bgp group external-peers export send-static
set protocols bgp group external-peers peer-as 18
set protocols bgp group external-peers neighbor 192.168.6.7
set policy-options policy-statement send-static term 1 from protocol static
set policy-options policy-statement send-static term 1 then accept
set routing-options static route 10.10.10.14/32 next-hop 10.10.10.10
set routing-options static route 192.168.6.7/32 next-hop 10.10.10.10
set routing-options router-id 192.168.40.4
set routing-options autonomous-system 17
```

**Device D**

```
set interfaces fe-1/2/0 unit 10 description to-C
set interfaces fe-1/2/0 unit 10 family inet address 10.10.10.10/30
set interfaces fe-1/2/1 unit 13 description to-E
set interfaces fe-1/2/1 unit 13 family inet address 10.10.10.13/30
set interfaces lo0 unit 4 family inet address 192.168.6.6/32
set routing-options static route 192.168.40.4/32 next-hop 10.10.10.9
set routing-options static route 192.168.6.7/32 next-hop 10.10.10.14
set routing-options router-id 192.168.6.6
```

**Device E**

```
set interfaces fe-1/2/0 unit 14 description to-D
set interfaces fe-1/2/0 unit 14 family inet address 10.10.10.14/30
set interfaces lo0 unit 5 family inet address 192.168.6.7/32
set protocols bgp group external-peers multihop ttl 2
set protocols bgp group external-peers local-address 192.168.6.7
```

```

set protocols bgp group external-peers export send-static
set protocols bgp group external-peers peer-as 17
set protocols bgp group external-peers neighbor 192.168.40.4
set policy-options policy-statement send-static term 1 from protocol static
set policy-options policy-statement send-static term 1 then accept
set routing-options static route 10.10.10.8/30 next-hop 10.10.10.13
set routing-options static route 192.168.40.4/32 next-hop 10.10.10.13
set routing-options router-id 192.168.6.7
set routing-options autonomous-system 18

```

### Device C

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device C:

1. Configure the interface to the directly connected device (to-D), and configure the loopback interface.

```

[edit interfaces fe-1/2/0 unit 9]
user@C# set description to-D
user@C# set family inet address 10.10.10.9/30

```

```

[edit interfaces lo0 unit 3]
user@C# set family inet address 192.168.40.4/32

```

2. Configure an EBGp session with Device E.

The **neighbor** statement points to the loopback interface on Device E.

```

[edit protocols bgp group external-peers]
user@C# set type external
user@C# set local-address 192.168.40.4
user@C# set export send-static
user@C# set peer-as 18
user@C# set neighbor 192.168.6.7

```

3. Configure the multihop statement to enable Device C and Device E to become EBGp peers.

Because the peers are two hops away from each other, the example uses the **ttl 2** statement.

```

[edit protocols bgp group external-peers]
user@C# set multihop ttl 2

```

4. Configure connectivity to Device E, using static routes.

You must configure a route to both the loopback interface address and to the address on the physical interface.

```

[edit routing-options]
user@C# set static route 10.10.10.14/32 next-hop 10.10.10.10
user@C# set static route 192.168.6.7/32 next-hop 10.10.10.10

```

5. Configure the local router ID and the autonomous system (AS) number.

```
[edit routing-options]
user@C# set router-id 192.168.40.4
user@C# set autonomous-system 17
```

6. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-static term 1]
user@C# set from protocol static
user@C# set then accept
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@C# show interfaces
fe-1/2/0 {
 unit 9 {
 description to-D;
 family inet {
 address 10.10.10.9/30;
 }
 }
}
lo0 {
 unit 3 {
 family inet {
 address 192.168.40.4/32;
 }
 }
}

user@C# show protocols
bgp {
 group external-peers {
 type external;
 multihop {
 ttl 2;
 }
 local-address 192.168.40.4;
 export send-static;
 peer-as 18;
 neighbor 192.168.6.7;
 }
}

user@C# show policy-options
policy-statement send-static {
 term 1 {
 from protocol static;
 then accept;
 }
}
```

```

user@C# show routing-options
static {
 route 10.10.10.14/32 next-hop 10.10.10.10;
 route 192.168.6.7/32 next-hop 10.10.10.10;
}
router-id 192.168.40.4;
autonomous-system 17;

```

If you are done configuring the device, enter **commit** from configuration mode. Repeat these steps for all BFD sessions in the topology.

### Configuring Device D

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device D:

1. Set the CLI to Device D.

```
user@host> set cli logical-system D
```

2. Configure the interfaces to the directly connected devices, and configure a loopback interface.

```

[edit interfaces fe-1/2/0 unit 10]
user@D# set description to-C
user@D# set family inet address 10.10.10.10/30

```

```

[edit interfaces fe-1/2/1 unit 13]
user@D# set description to-E
user@D# set family inet address 10.10.10.13/30

```

```

[edit interfaces lo0 unit 4]
user@D# set family inet address 192.168.6.6/32

```

3. Configure connectivity to the other devices using static routes to the loopback interface addresses.

On Device D, you do not need static routes to the physical addresses because Device D is directly connected to Device C and Device E.

```

[edit routing-options]
user@D# set static route 192.168.40.4/32 next-hop 10.10.10.9
user@D# set static route 192.168.6.7/32 next-hop 10.10.10.14

```

4. Configure the local router ID.

```

[edit routing-options]
user@D# set router-id 192.168.6.6

```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces** and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@D# show interfaces
```

```
fe-1/2/0 {
 unit 10 {
 description to-C;
 family inet {
 address 10.10.10.10/30;
 }
 }
}
fe-1/2/1 {
 unit 13 {
 description to-E;
 family inet {
 address 10.10.10.13/30;
 }
 }
}
lo0 {
 unit 4 {
 family inet {
 address 192.168.6.6/32;
 }
 }
}

user@D# show protocols

user@D# show routing-options
static {
 route 192.168.40.4/32 next-hop 10.10.10.9;
 route 192.168.6.7/32 next-hop 10.10.10.14;
}
router-id 192.168.6.6;
```

If you are done configuring the device, enter **commit** from configuration mode.  
Repeat these steps for all BFD sessions in the topology.

### ***Configuring Device E***

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device E:

1. Set the CLI to Device E.  

```
user@host> set cli logical-system E
```
2. Configure the interface to the directly connected device (to-D), and configure the loopback interface.  

```
[edit interfaces fe-1/2/0 unit 14]
user@E# set description to-D
user@E# set family inet address 10.10.10.14/30

[edit interfaces lo0 unit 5]
user@E# set family inet address 192.168.6.7/32
```

3. Configure an EBGP session with Device E.

The **neighbor** statement points to the loopback interface on Device C.

```
[edit protocols bgp group external-peers]
user@E# set local-address 192.168.6.7
user@E# set export send-static
user@E# set peer-as 17
user@E# set neighbor 192.168.40.4
```

4. Configure the **multihop** statement to enable Device C and Device E to become EBGP peers.

Because the peers are two hops away from each other, the example uses the **ttl 2** statement.

```
[edit protocols bgp group external-peers]
user@E# set multihop ttl 2
```

5. Configure connectivity to Device E, using static routes.

You must configure a route to both the loopback interface address and to the address on the physical interface.

```
[edit routing-options]
user@E# set static route 10.10.10.8/30 next-hop 10.10.10.13
user@E# set static route 192.168.40.4/32 next-hop 10.10.10.13
```

6. Configure the local router ID and the autonomous system (AS) number.

```
[edit routing-options]
user@E# set router-id 192.168.6.7
user@E# set autonomous-system 18
```

7. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-static term 1]
user@E# set from protocol static
user@E# set then accept
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@E# show interfaces
fe-1/2/0 {
 unit 14 {
 description to-D;
 family inet {
 address 10.10.10.14/30;
 }
 }
}
lo0 {
 unit 5 {
```

```
 family inet {
 address 192.168.6.7/32;
 }
 }
}

user@E# show protocols
bgp {
 group external-peers {
 multihop {
 ttl 2;
 }
 local-address 192.168.6.7;
 export send-static;
 peer-as 17;
 neighbor 192.168.40.4;
 }
}

user@E# show policy-options
policy-statement send-static {
 term 1 {
 from protocol static;
 then accept;
 }
}

user@E# show routing-options
static {
 route 10.10.10.8/30 next-hop 10.10.10.13;
 route 192.168.40.4/32 next-hop 10.10.10.13;
}
router-id 192.168.6.7;
autonomous-system 18;
```

If you are done configuring the device, enter **commit** from configuration mode.

---

## Verification

Confirm that the configuration is working properly.

- [Verifying Connectivity on page 3784](#)
- [Verifying That BGP Sessions Are Established on page 3785](#)
- [Viewing Advertised Routes on page 3785](#)

### *Verifying Connectivity*

**Purpose** Make sure that Device C can ping Device E, specifying the loopback interface address as the source of the ping request.

The loopback interface address is the source address that BGP will use.

**Action** From operational mode, enter the **ping 10.10.10.14 source 192.168.40.4** command from Device C, and enter the **ping 10.10.10.9 source 192.168.6.7** command from Device E.

```
user@C> ping 10.10.10.14 source 192.168.40.4
```



```

PING 10.10.10.14 (10.10.10.14): 56 data bytes
64 bytes from 10.10.10.14: icmp_seq=0 ttl=63 time=1.262 ms
64 bytes from 10.10.10.14: icmp_seq=1 ttl=63 time=1.202 ms
^C
--- 10.10.10.14 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.202/1.232/1.262/0.030 ms

user@E> ping 10.10.10.9 source 192.168.6.7

PING 10.10.10.9 (10.10.10.9): 56 data bytes
64 bytes from 10.10.10.9: icmp_seq=0 ttl=63 time=1.255 ms
64 bytes from 10.10.10.9: icmp_seq=1 ttl=63 time=1.158 ms
^C
--- 10.10.10.9 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.158/1.206/1.255/0.049 ms

```

**Meaning** The static routes are working if the pings work.

### *Verifying That BGP Sessions Are Established*

**Purpose** Verify that the BGP sessions are up.

**Action** From operational mode, enter the `show bgp summary` command.

```
user@C> show bgp summary
```

```

Groups: 1 Peers: 1 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet.0 2 0 0 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
192.168.6.7 18 147 147 0 1 1:04:27
0/2/2/0 0/0/0/0

```

```
user@E> show bgp summary
```

```

Groups: 1 Peers: 1 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet.0 2 0 0 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
192.168.40.4 17 202 202 0 1 1:02:18
0/2/2/0 0/0/0/0

```

**Meaning** The output shows that both devices have one peer each. No peers are down.

### *Viewing Advertised Routes*

**Purpose** Check to make sure that routes are being advertised by BGP.

**Action** From operational mode, enter the `show route advertising-protocol bgp neighbor` command.

```
user@C> show route advertising-protocol bgp 192.168.6.7
```

```
inet.0: 5 destinations, 7 routes (5 active, 0 holddown, 0 hidden)
```

```

 Prefix Nexthop MED Lclpref AS path
* 10.10.10.14/32 Self
* 192.168.6.7/32 Self

```

```
user@E> show route advertising-protocol bgp 192.168.40.4
```

```

inet.0: 5 destinations, 7 routes (5 active, 0 holddown, 0 hidden)
 Prefix Nexthop MED Lclpref AS path
* 10.10.10.8/30 Self
* 192.168.40.4/32 Self

```

**Meaning** The **send-static** routing policy is exporting the static routes from the routing table into BGP. BGP is advertising these routes between the peers because the BGP peer session is established.

**Related Documentation**

- [Examples: Configuring External BGP Peering on page 3601](#)
- [BGP Configuration Overview](#)

## Example: Configuring BGP Route Preference (Administrative Distance)

- [Understanding Route Preference Values \(Administrative Distance\) on page 3786](#)
- [Example: Configuring the Preference Value for BGP Routes on page 3787](#)

## Understanding Route Preference Values (Administrative Distance)

The Junos OS routing protocol process assigns a default preference value (also known as an *administrative distance*) to each route that the routing table receives. The default value depends on the source of the route. The preference value is a value from 0 through 4,294,967,295 ( $2^{32} - 1$ ), with a lower value indicating a more preferred route. [Table 318](#) lists the default preference values.

**Table 318: Default Route Preference Values**

| How Route Is Learned         | Default Preference | Statement to Modify Default Preference                                                                    |
|------------------------------|--------------------|-----------------------------------------------------------------------------------------------------------|
| Directly connected network   | 0                  | —                                                                                                         |
| System routes                | 4                  | —                                                                                                         |
| Static and Static LSPs       | 5                  | <i>static</i>                                                                                             |
| RSVP-signaled LSPs           | 7                  | RSVP <b>preference</b> as described in the <i>Junos OS MPLS Applications Library for Routing Devices</i>  |
| LDP-signaled LSPs            | 9                  | LDP <b>preference</b> , as described in the <i>Junos OS MPLS Applications Library for Routing Devices</i> |
| OSPF internal route          | 10                 | OSPF <b>preference</b>                                                                                    |
| IS-IS Level 1 internal route | 15                 | IS-IS <b>preference</b>                                                                                   |

Table 318: Default Route Preference Values (*continued*)

| How Route Is Learned         | Default Preference | Statement to Modify Default Preference                                           |
|------------------------------|--------------------|----------------------------------------------------------------------------------|
| IS-IS Level 2 internal route | 18                 | IS-IS <a href="#">preference</a>                                                 |
| Redirects                    | 30                 | –                                                                                |
| Kernel                       | 40                 | –                                                                                |
| SNMP                         | 50                 | –                                                                                |
| Router discovery             | 55                 | –                                                                                |
| RIP                          | 100                | RIP <a href="#">preference</a>                                                   |
| RIPng                        | 100                | RIPng <i>preference</i>                                                          |
| PIM                          | 105                | <i>Multicast Protocols Feature Guide for Routing Devices</i>                     |
| DVMRP                        | 110                | <i>Multicast Protocols Feature Guide for Routing Devices</i>                     |
| Aggregate                    | 130                | <a href="#">aggregate</a>                                                        |
| OSPF AS external routes      | 150                | OSPF <a href="#">external-preference</a>                                         |
| IS-IS Level 1 external route | 160                | IS-IS <a href="#">external-preference</a>                                        |
| IS-IS Level 2 external route | 165                | IS-IS <a href="#">external-preference</a>                                        |
| BGP                          | 170                | BGP <a href="#">preference</a> , <a href="#">export</a> , <a href="#">import</a> |
| MSDP                         | 175                | <i>Multicast Protocols Feature Guide for Routing Devices</i>                     |

In general, the narrower the scope of the statement, the higher precedence its preference value is given, but the smaller the set of routes it affects. To modify the default preference value for routes learned by routing protocols, you generally apply routing policy when configuring the individual routing protocols. You also can modify some preferences with other configuration statements, which are indicated in the table.

### Example: Configuring the Preference Value for BGP Routes

This example shows how to specify the preference for routes learned from BGP. Routing information can be learned from multiple sources. To break ties among equally specific routes learned from multiple sources, each source has a preference value. Routes that are learned through explicit administrative action, such as static routes, are preferred

over routes learned from a routing protocol, such as BGP or OSPF. This concept is called *administrative distance* by some vendors.

- [Requirements on page 3788](#)
- [Overview on page 3788](#)
- [Configuration on page 3789](#)
- [Verification on page 3792](#)

---

## Requirements

No special configuration beyond device initialization is required before you configure this example.

---

## Overview

Routing information can be learned from multiple sources, such as through static configuration, BGP, or an interior gateway protocol (IGP). When Junos OS determines a route's preference to become the active route, it selects the route with the lowest preference as the active route and installs this route into the forwarding table. By default, the routing software assigns a preference of 170 to routes that originated from BGP. Of all the routing protocols, BGP has the highest default preference value, which means that routes learned by BGP are the least likely to become the active route.

Some vendors have a preference (distance) of 20 for external BGP (EBGP) and a distance of 200 for internal BGP (IBGP). Junos OS uses the same value (170) for both EBGP and IBGP. However, this difference between vendors has no operational impact because Junos OS always prefers EBGP routes over IBGP routes.

Another area in which vendors differ is in regard to IGP distance compared to BGP distance. For example, some vendors assign a distance of 110 to OSPF routes. This is higher than the EBGP distance of 20, and results in the selection of an EBGP route over an equivalent OSPF route. In the same scenario, Junos OS chooses the OSPF route, because of the default preference 10 for an internal OSPF route and 150 for an external OSPF route, which are both lower than the 170 preference assigned to all BGP routes.

In a multivendor environment, you might want to change the preference value for BGP routes so that Junos OS chooses an EBGP route instead of an OSPF route. To accomplish this goal, one option is to include the **preference** statement in the EBGP configuration. To modify the default BGP preference value, include the **preference** statement, specifying a value from 0 through 4,294,967,295 ( $2^{32} - 1$ ).



**TIP:** Another way to achieve multivendor compatibility is to include the **advertise-inactive** statement in the EBGP configuration. This causes the routing table to export to BGP the best route learned by BGP even if Junos OS did not select it to be an active route. By default, BGP stores the route information it receives from update messages in the Junos OS routing table, and the routing table exports only active routes into BGP, which BGP then advertises to its peers. The **advertise-inactive** statement causes Junos OS to advertise the best BGP route that is inactive because of IGP preference. When

you use the `advertise-inactive` statement, the Junos OS device uses the OSPF route for forwarding, and the other vendor's device uses the EBGP route for forwarding. However, from the perspective of an EBGP peer in a neighboring AS, both vendors' devices appear to behave the same way.

### Topology


In the sample network, Device R1 and Device R2 have EBGP routes to each other and also OSPF routes to each other.

This example shows the routing tables in the following cases:

- Accept the default preference values of 170 for BGP and 10 for OSPF.
- Change the BGP preference to 8.

Figure 74 shows the sample network.

**Figure 74: BGP Preference Value Topology**



### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```

Device R1 set interfaces fe-1/2/0 unit 4 family inet address 1.12.0.1/30
 set interfaces lo0 unit 2 family inet address 10.255.71.24/32
 set protocols bgp export send-direct

```

```
set protocols bgp group ext type external
set protocols bgp group ext preference 8
set protocols bgp group ext peer-as 65000
set protocols bgp group ext neighbor 1.12.0.2
set protocols ospf area 0.0.0.0 interface fe-1/2/0.4
set protocols ospf area 0.0.0.0 interface 10.255.71.24
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 65500
```

**Device R2**

```
set interfaces fe-1/2/0 unit 6 family inet address 1.12.0.2/30
set interfaces lo0 unit 3 family inet address 10.255.14.177/32
set protocols bgp export send-direct
set protocols bgp group ext type external
set protocols bgp group ext peer-as 65500
set protocols bgp group ext neighbor 1.12.0.1
set protocols ospf area 0.0.0.0 interface fe-1/2/0.6
set protocols ospf area 0.0.0.0 interface 10.255.14.177
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 65000
```

**Step-by-Step Procedure** The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R1:

1. Configure the interfaces.  

```
[edit interfaces]
user@R1# set fe-1/2/0 unit 4 family inet address 1.12.0.1/30
user@R1# set lo0 unit 2 family inet address 10.255.71.24/32
```
2. Configure the local autonomous system.  

```
[edit routing-options]
user@R1# set autonomous-system 65500
```
3. Configure the external peering with Device R2.  

```
[edit protocols bgp]
user@R1# set export send-direct
user@R1# set group ext type external
user@R1# set group ext preference 8
user@R1# set group ext peer-as 65000
user@R1# set group ext neighbor 1.12.0.2
```
4. Configure OSPF.  

```
[edit protocols ospf area 0.0.0.0]
user@R1# set interface fe-1/2/0.4
user@R1# set interface 10.255.71.24
```
5. Configure the routing policy.  

```
[edit policy-options policy-statement send-direct term 1]
user@R1# set from protocol direct
user@R1# set then accept
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@R1# show interfaces
fe-1/2/0 {
 unit 4 {
 family inet {
 address 1.12.0.1/30;
 }
 }
}
lo0 {
 unit 2 {
 family inet {
 address 10.255.71.24/32;
 }
 }
}

user@R1# show policy-options
policy-statement send-direct {
 term 1 {
 from protocol direct;
 then accept;
 }
}

user@R1# show protocols
protocols {
 bgp {
 export send-direct;
 group ext {
 type external;
 preference 8;
 peer-as 65000;
 neighbor 1.12.0.2;
 }
 }
 ospf {
 area 0.0.0.0 {
 interface fe-1/2/0.4;
 interface 10.255.71.24;
 }
 }
}

user@R1# show routing-options
autonomous-system 65500;

```

If you are done configuring the device, enter **commit** from configuration mode. Repeat these steps on Device R2.

## Verification

Confirm that the configuration is working properly.

### Verifying the Preference

**Purpose** Make sure that the routing tables on Device R1 and Device R2 reflect the fact that Device R1 is using the configured EBGP preference of 8, and Device R2 is using the default EBGP preference of 170.

**Action** From operational mode, enter the **show route** command.

```
user@R1> show route
inet.0: 5 destinations, 7 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
1.12.0.0/30 *[Direct/0] 3d 07:03:01
 > via fe-1/2/0.4
 [BGP/8] 01:04:49, localpref 100
 AS path: 65000 I
 > to 1.12.0.2 via fe-1/2/0.4
1.12.0.1/32 *[Local/0] 3d 07:03:01
 Local via fe-1/2/0.4
10.255.14.177/32 *[BGP/8] 01:04:49, localpref 100
 AS path: 65000 I
 > to 1.12.0.2 via fe-1/2/0.4
 [OSPF/10] 3d 07:02:16, metric 1
 > to 1.12.0.2 via fe-1/2/0.4
10.255.71.24/32 *[Direct/0] 3d 07:03:01
 > via lo0.2
224.0.0.5/32 *[OSPF/10] 5d 03:42:16, metric 1
 MultiRecv
```

```
user@R2> show route
inet.0: 5 destinations, 7 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
1.12.0.0/30 *[Direct/0] 3d 07:03:30
 > via fe-1/2/0.6
 [BGP/170] 00:45:36, localpref 100
 AS path: 65500 I
 > to 1.12.0.1 via fe-1/2/0.6
1.12.0.2/32 *[Local/0] 3d 07:03:30
 Local via fe-1/2/0.6
10.255.14.177/32 *[Direct/0] 3d 07:03:30
 > via lo0.3
10.255.71.24/32 *[OSPF/10] 3d 07:02:45, metric 1
 > to 1.12.0.1 via fe-1/2/0.6
 [BGP/170] 00:45:36, localpref 100
 AS path: 65500 I
 > to 1.12.0.1 via fe-1/2/0.6
224.0.0.5/32 *[OSPF/10] 5d 03:42:45, metric 1
 MultiRecv
```

**Meaning** The output shows that on Device R1, the active path to Device R2's loopback interface (10.255.14.177/32) is a BGP route. The output also shows that on Device R2, the active path to Device R1's loopback interface (10.255.71.24/32) is an OSPF route.



- Related Documentation**
- [Route Preferences Overview](#)
  - [Understanding External BGP Peering Sessions on page 3601](#)
  - [BGP Configuration Overview](#)

## Example: Configuring BGP Path Selection

- [Understanding BGP Path Selection on page 3793](#)
- [Example: Ignoring the AS Path Attribute When Selecting the Best Path on page 3796](#)

### Understanding BGP Path Selection

For each prefix in the routing table, the routing protocol process selects a single best path. After the best path is selected, the route is installed in the routing table. The best path becomes the active route if the same prefix is not learned by a protocol with a lower (more preferred) global preference value, also known as the administrative distance. The algorithm for determining the active route is as follows:

1. Verify that the next hop can be resolved.
2. Choose the path with the lowest preference value (routing protocol process preference).

Routes that are not eligible to be used for forwarding (for example, because they were rejected by routing policy or because a next hop is inaccessible) have a preference of  $-1$  and are never chosen.

3. Prefer the path with higher local preference.  
For non-BGP paths, choose the path with the lowest **preference2** value.
4. If the accumulated interior gateway protocol (AIGP) attribute is enabled, prefer the path with the lower AIGP attribute.
5. Prefer the path with the shortest autonomous system (AS) path value (skipped if the **as-path-ignore** statement is configured).

A confederation segment (sequence or set) has a path length of 0. An AS set has a path length of 1.

6. Prefer the route with the lower origin code.

Routes learned from an IGP have a lower origin code than those learned from an exterior gateway protocol (EGP), and both have lower origin codes than incomplete routes (routes whose origin is unknown).

7. Prefer the path with the lowest multiple exit discriminator (MED) metric.

Depending on whether nondeterministic routing table path selection behavior is configured, there are two possible cases:

- If nondeterministic routing table path selection behavior is not configured (that is, if the **path-selection cisco-nondeterministic** statement is not included in the BGP configuration), for paths with the same neighboring AS numbers at the front of the

AS path, prefer the path with the lowest MED metric. To always compare MEDs whether or not the peer ASs of the compared routes are the same, include the **path-selection always-compare-med** statement.

- If nondeterministic routing table path selection behavior is configured (that is, the **path-selection cisco-nondeterministic** statement is included in the BGP configuration), prefer the path with the lowest MED metric.

Confederations are not considered when determining neighboring ASs. A missing MED metric is treated as if a MED were present but zero.



**NOTE:** MED comparison works for single path selection within an AS (when the route does not include an AS path), though this usage is uncommon.

By default, only the MEDs of routes that have the same peer autonomous systems (ASs) are compared. You can configure routing table path selection options to obtain different behaviors.

8. Prefer strictly internal paths, which include IGP routes and locally generated routes (static, direct, local, and so forth).
9. Prefer strictly external BGP (EBGP) paths over external paths learned through internal BGP (IBGP) sessions.
10. Prefer the path whose next hop is resolved through the IGP route with the lowest metric.



**NOTE:** A path is considered a BGP equal-cost path (and will be used for forwarding) if a tie-break is performed after the previous step. All paths with the same neighboring AS, learned by a multipath-enabled BGP neighbor, are considered.

BGP multipath does not apply to paths that share the same MED-plus-IGP cost yet differ in IGP cost. Multipath path selection is based on the IGP cost metric, even if two paths have the same MED-plus-IGP cost.

11. If both paths are external, prefer the currently active path to minimize route-flapping. This rule is not used if any one of the following conditions is true:
  - **path-selection external-router-id** is configured.
  - Both peers have the same router ID.
  - Either peer is a confederation peer.
  - Neither path is the current active path.
12. Prefer a primary route over a secondary route. A primary route is one that belongs to the routing table. A secondary route is one that is added to the routing table through an export policy.

13. Prefer the path from the peer with the lowest router ID. For any path with an originator ID attribute, substitute the originator ID for the router ID during router ID comparison.
14. Prefer the path with the shortest cluster list length. The length is 0 for no list.
15. Prefer the path from the peer with the lowest peer IP address.

### Routing Table Path Selection

The shortest AS path step of the algorithm, by default, evaluates the length of the AS path and determines the active path. You can configure an option that enables Junos OS to skip this step of the algorithm by including the **as-path-ignore** option.



**NOTE:** The **as-path-ignore** option is not supported for routing instances.

To configure routing table path selection behavior, include the **path-selection** statement:

```
path-selection {
 (always-compare-med | cisco-non-deterministic | external-router-id);
 as-path-ignore;
 med-plus-igp {
 igp-multiplier number;
 med-multiplier number;
 }
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Routing table path selection can be configured in one of the following ways:

- Emulate the Cisco IOS default behavior (**cisco-non-deterministic**). This mode evaluates routes in the order that they are received and does not group them according to their neighboring AS. With **cisco-non-deterministic** mode, the active path is always first. All inactive, but eligible, paths follow the active path and are maintained in the order in which they were received, with the most recent path first. Ineligible paths remain at the end of the list.

As an example, suppose you have three path advertisements for the 192.168.1.0 /24 route:

- Path 1—learned through EBGp; AS Path of 65010; MED of 200
- Path 2—learned through IBGP; AS Path of 65020; MED of 150; IGP cost of 5
- Path 3—learned through IBGP; AS Path of 65010; MED of 100; IGP cost of 10

These advertisements are received in quick succession, within a second, in the order listed. Path 3 is received most recently, so the routing device compares it against path 2, the next most recent advertisement. The cost to the IBGP peer is better for path 2, so the routing device eliminates path 3 from contention. When comparing paths 1 and 2, the routing device prefers path 1 because it is received from an EBGp peer. This allows the routing device to install path 1 as the active path for the route.



**NOTE:** We do not recommend using this configuration option in your network. It is provided solely for interoperability to allow all routing devices in the network to make consistent route selections.

- Always comparing MEDs whether or not the peer ASs of the compared routes are the same (**always-compare-med**).
- Override the rule that If both paths are external, the currently active path is preferred (**external-router-id**). Continue with the next step (Step 12) in the path-selection process.
- Adding the IGP cost to the next-hop destination to the MED value before comparing MED values for path selection (**med-plus-igp**).

BGP multipath does not apply to paths that share the same MED-plus-IGP cost, yet differ in IGP cost. Multipath path selection is based on the IGP cost metric, even if two paths have the same MED-plus-IGP cost.

---

### Effects of Advertising Multiple Paths to a Destination

BGP advertises only the active path, unless you configure BGP to advertise multiple paths to a destination.

Suppose a routing device has in its routing table four paths to a destination and is configured to advertise up to three paths (**add-path send path-count 3**). The three paths are chosen based on path selection criteria. That is, the three best paths are chosen in path-selection order. The best path is the active path. This path is removed from consideration and a new best path is chosen. This process is repeated until the specified number of paths is reached.

### Example: Ignoring the AS Path Attribute When Selecting the Best Path

If multiple BGP routes to the same destination exist, BGP selects the best path based on the route attributes of the paths. One of the route attributes that affects the best-path decision is the length of the AS paths of each route. Routes with shorter AS paths are preferred over those with longer AS paths. Although not typically practical, some scenarios might require that the AS path length be ignored in the route selection process. This example shows how to configure a routing device to ignore the AS path attribute.

- [Requirements on page 3796](#)
- [Overview on page 3797](#)
- [Configuration on page 3798](#)
- [Verification on page 3803](#)

---

### Requirements

No special configuration beyond device initialization is required before you configure this example.

## Overview

On externally connected routing devices, the purpose of skipping the AS path comparison might be to force an external BGP (EBGP) versus internal BGP (IBGP) decision to remove traffic from your network as soon as possible. On internally connected routing devices, you might want your IBGP-only routers to default to the local externally connected gateway. The local IBGP-only (internal) routers skip the AS path comparison and move down the decision tree to use the closest interior gateway protocol (IGP) gateway (lowest IGP metric). Doing this might be an effective way to force these routers to use a LAN connection instead of their WAN connection.



**CAUTION:** When you include the `as-path-ignore` statement on a routing device in your network, you might need to include it on all other BGP-enabled devices in your network to prevent routing loops and convergence issues. This is especially true for IBGP path comparisons.


In this example, Device R2 is learning about the loopback interface address on Device R4 (4.4.4.4/32) from Device R1 and Device R3. Device R1 is advertising 4.4.4.4/32 with an AS-path of 1 5 4, and Device R3 is advertising 4.4.4.4/32 with an AS-path of 3 4. Device R2 selects the path for 4.4.4.4/32 from Device R3 as the best path because the AS path is shorter than the AS path from Device R1.

This example modifies the BGP configuration on Device R2 so that the AS-path length is not used in the best-path selection.

Device R1 has a lower router ID (1.1.1.1) than Device R3 (1.1.1.1). If all other path selection criteria are equal (or, as in this case, ignored), the route learned from Device R1 is used. Because the AS-path attribute is being ignored, the best path is toward Device R1 because of its lower router ID value.

Figure 75 shows the sample topology.

Figure 75: Topology for Ignoring the AS-Path Length



g041166

### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

#### Device R1

```
set interfaces fe-1/2/0 unit 1 family inet address 192.168.10.1/24
set interfaces fe-1/2/1 unit 10 family inet address 192.168.50.2/24
set interfaces lo0 unit 1 family inet address 1.1.1.1/32
set protocols bgp group ext type external
set protocols bgp group ext export send-direct
set protocols bgp group ext export send-static
set protocols bgp group ext export send-local
set protocols bgp group ext neighbor 192.168.10.2 peer-as 2
set protocols bgp group ext neighbor 192.168.50.1 peer-as 5
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set policy-options policy-statement send-local term 1 from protocol local
```

```

set policy-options policy-statement send-local term 1 then accept
set policy-options policy-statement send-static term 1 from protocol static
set policy-options policy-statement send-static term 1 then accept
set routing-options static route 192.168.20.0/24 next-hop 192.168.10.2
set routing-options static route 192.168.30.0/24 next-hop 192.168.10.2
set routing-options static route 192.168.40.0/24 next-hop 192.168.50.1
set routing-options router-id 1.1.1.1
set routing-options autonomous-system 1

```

**Device R2**

```

set interfaces fe-1/2/0 unit 2 family inet address 192.168.10.2/24
set interfaces fe-1/2/1 unit 3 family inet address 192.168.20.2/24
set interfaces lo0 unit 2 family inet address 2.2.2.2/32
set protocols bgp path-selection as-path-ignore
set protocols bgp group ext type external
set protocols bgp group ext export send-direct
set protocols bgp group ext export send-static
set protocols bgp group ext export send-local
set protocols bgp group ext neighbor 192.168.10.1 peer-as 1
set protocols bgp group ext neighbor 192.168.20.1 peer-as 3
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set policy-options policy-statement send-local term 1 from protocol local
set policy-options policy-statement send-local term 1 then accept
set policy-options policy-statement send-static term 1 from protocol static
set policy-options policy-statement send-static term 1 then accept
set routing-options static route 192.168.50.0/24 next-hop 192.168.10.1
set routing-options static route 192.168.40.0/24 next-hop 192.168.10.1
set routing-options static route 192.168.30.0/24 next-hop 192.168.20.1
set routing-options router-id 2.2.2.2
set routing-options autonomous-system 2

```

**Device R3**

```

set interfaces fe-1/2/0 unit 4 family inet address 192.168.20.1/24
set interfaces fe-1/2/1 unit 5 family inet address 192.168.30.1/24
set interfaces lo0 unit 3 family inet address 1.1.1.1/32
set protocols bgp group ext type external
set protocols bgp group ext export send-direct
set protocols bgp group ext export send-static
set protocols bgp group ext export send-local
set protocols bgp group ext neighbor 192.168.20.2 peer-as 2
set protocols bgp group ext neighbor 192.168.30.2 peer-as 4
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set policy-options policy-statement send-local term 1 from protocol local
set policy-options policy-statement send-local term 1 then accept
set policy-options policy-statement send-static term 1 from protocol static
set policy-options policy-statement send-static term 1 then accept
set routing-options static route 192.168.10.0/24 next-hop 192.168.20.2
set routing-options static route 192.168.50.0/24 next-hop 192.168.20.2
set routing-options static route 192.168.40.0/24 next-hop 192.168.30.2
set routing-options router-id 3.3.3.3
set routing-options autonomous-system 3

```

**Device R4**

```

set interfaces fe-1/2/0 unit 6 family inet address 192.168.30.2/24
set interfaces fe-1/2/1 unit 7 family inet address 192.168.40.1/24
set interfaces lo0 unit 4 family inet address 4.4.4.4/32

```

```
set protocols bgp group ext type external
set protocols bgp group ext export send-direct
set protocols bgp group ext export send-static
set protocols bgp group ext export send-local
set protocols bgp group ext neighbor 192.168.30.1 peer-as 3
set protocols bgp group ext neighbor 192.168.40.2 peer-as 5
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set policy-options policy-statement send-local term 1 from protocol local
set policy-options policy-statement send-local term 1 then accept
set policy-options policy-statement send-static term 1 from protocol static
set policy-options policy-statement send-static term 1 then accept
set routing-options static route 192.168.10.0/24 next-hop 192.168.40.2
set routing-options static route 192.168.50.0/24 next-hop 192.168.40.2
set routing-options static route 192.168.40.0/24 next-hop 192.168.30.1
set routing-options router-id 4.4.4.4
set routing-options autonomous-system 4
```

**Device R5**

```
set interfaces fe-1/2/0 unit 8 family inet address 192.168.40.2/24
set interfaces fe-1/2/1 unit 9 family inet address 192.168.50.1/24
set interfaces lo0 unit 5 family inet address 5.5.5.5/32
set protocols bgp group ext type external
set protocols bgp group ext export send-direct
set protocols bgp group ext export send-static
set protocols bgp group ext export send-local
set protocols bgp group ext neighbor 192.168.40.1 peer-as 4
set protocols bgp group ext neighbor 192.168.50.2 peer-as 1
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set policy-options policy-statement send-local term 1 from protocol local
set policy-options policy-statement send-local term 1 then accept
set policy-options policy-statement send-static term 1 from protocol static
set policy-options policy-statement send-static term 1 then accept
set routing-options static route 192.168.10.0/24 next-hop 192.168.50.2
set routing-options static route 192.168.20.0/24 next-hop 192.168.50.2
set routing-options static route 192.168.30.0/24 next-hop 192.168.40.1
set routing-options router-id 5.5.5.5
set routing-options autonomous-system 5
```

### *Configuring Device R2*

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R2:

1. Configure the interfaces.  
  
[edit interfaces]  
user@R2# set fe-1/2/0 unit 2 family inet address 192.168.10.2/24  
user@R2# set fe-1/2/1 unit 3 family inet address 192.168.20.2/24  
user@R2# set lo0 unit 2 family inet address 2.2.2.2/32
2. Configure EBGp.  
  
[edit protocols bgp group ext]



```

user@R2# set type external
user@R2# set export send-direct
user@R2# set export send-static
user@R2# set export send-local
user@R2# set neighbor 192.168.10.1 peer-as 1
user@R2# set neighbor 192.168.20.1 peer-as 3

```

3. Configure the autonomous system (AS) path attribute to be ignored in the Junos OS path selection algorithm.

```

[edit protocols bgp]
user@R2# set path-selection as-path-ignore

```

4. Configure the routing policy.

```

[edit policy-options]
user@R2# set policy-statement send-direct term 1 from protocol direct
user@R2# set policy-statement send-direct term 1 then accept
user@R2# set policy-statement send-local term 1 from protocol local
user@R2# set policy-statement send-local term 1 then accept
user@R2# set policy-statement send-static term 1 from protocol static
user@R2# set policy-statement send-static term 1 then accept

```

5. Configure some static routes.

```

[edit routing-options static]
user@R2# set route 192.168.50.0/24 next-hop 192.168.10.1
user@R2# set route 192.168.40.0/24 next-hop 192.168.10.1
user@R2# set route 192.168.30.0/24 next-hop 192.168.20.1

```

6. Configure the autonomous system (AS) number and the router ID.

```

[edit routing-options]
user@R2# set router-id 2.2.2.2
user@R2# set autonomous-system 2

```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@R2# show interfaces
fe-1/2/0 {
 unit 2 {
 family inet {
 address 192.168.10.2/24;
 }
 }
}
fe-1/2/1 {
 unit 3 {
 family inet {
 address 192.168.20.2/24;
 }
 }
}
lo0 {
 unit 2 {

```

```
 family inet {
 address 2.2.2.2/32;
 }
 }
}

user@R2# show policy-options
policy-statement send-direct {
 term 1 {
 from protocol direct;
 then accept;
 }
}
policy-statement send-local {
 term 1 {
 from protocol local;
 then accept;
 }
}
policy-statement send-static {
 term 1 {
 from protocol static;
 then accept;
 }
}

user@R2# show protocols
bgp {
 path-selection as-path-ignore;
 group ext {
 type external;
 export [send-direct send-static send-local];
 neighbor 192.168.10.1 {
 peer-as 1;
 }
 neighbor 192.168.20.1 {
 peer-as 3;
 }
 }
}

user@R2# show routing-options
static {
 route 192.168.50.0/24 next-hop 192.168.10.1;
 route 192.168.40.0/24 next-hop 192.168.10.1;
 route 192.168.30.0/24 next-hop 192.168.20.1;
}
router-id 2.2.2.2;
autonomous-system 2;
```

If you are done configuring the device, enter **commit** from configuration mode. Repeat the configuration on the other devices in the network, changing the interface names and IP addresses, as needed.

## Verification

Confirm that the configuration is working properly.

- [Checking the Neighbor Status on page 3803](#)

### Checking the Neighbor Status

**Purpose** Make sure that from Device R2, the active path to get to AS 4 is through AS 1 and AS 5, not through AS 3.



**NOTE:** To verify the functionality of the `as-path-ignore` statement, you might need to run the `restart routing` command to force reevaluation of the active path. This is because for BGP, if both paths are external, the Junos OS behavior is to prefer the currently active path. This behavior helps to minimize route-flapping. Use caution when restarting the routing protocol process in a production network.

**Action** From operational mode, enter the `restart routing` command.

```
user@R2> restart routing
Routing protocols process started, pid 49396
```

From operational mode, enter the `show route 4.4.4.4 protocol bgp` command.

```
user@R2> show route 4.4.4.4 protocol bgp
inet.0: 12 destinations, 25 routes (12 active, 0 holddown, 4 hidden)
+ = Active Route, - = Last Active, * = Both

4.4.4.4/32 *[BGP/170] 00:00:12, localpref 100
 AS path: 1 5 4 I
 > to 192.168.10.1 via fe-1/2/0.2
 [BGP/170] 00:00:08, localpref 100
 AS path: 3 4 I
 > to 192.168.20.1 via fe-1/2/1.3
```

**Meaning** The asterisk (\*) is next to the path learned from R1, meaning that this is the active path. The AS path for the active path is 1 5 4, which is longer than the AS path (3 4) for the nonactive path learned from Router R3.

**Related Documentation**

- [Understanding External BGP Peering Sessions on page 3601](#)
- [BGP Configuration Overview](#)

## Example: Removing Private AS Numbers

- [Understanding Private AS Number Removal from AS Paths on page 3804](#)
- [Example: Removing Private AS Numbers from AS Paths on page 3805](#)

## Understanding Private AS Number Removal from AS Paths

By default, when BGP advertises AS paths to remote systems, it includes all AS numbers, including private AS numbers. You can configure the software so that it removes private AS numbers from AS paths. Doing this is useful when any of the following circumstances are true:

- A remote AS for which you provide connectivity is multihomed, but only to the local AS.
- The remote AS does not have an officially allocated AS number.
- It is not appropriate to make the remote AS a confederation member AS of the local AS.

Most companies acquire their own AS number. Some companies also use private AS numbers to connect to their public AS network. These companies might use a different private AS number for each region in which their company does business. In any implementation, announcing a private AS number to the Internet must be avoided. Service providers can use the **remove-private** statement to prevent advertising private AS numbers to the Internet.

In an enterprise scenario, suppose that you have multiple AS numbers in your company, some of which are private AS numbers, and one with a public AS number. The one with a public AS number has a direct connection to the service provider. In the AS that connects directly to the service provider, you can use the **remove-private** statement to filter out any private AS numbers in the advertisements that are sent to the service provider.



**CAUTION:** Changing configuration statements that affect BGP peers, such as enabling or disabling **remove-private** or renaming a BGP group, resets the BGP sessions. Changes that affect BGP peers should only be made when resetting a BGP session is acceptable.

The AS numbers are stripped from the AS path starting at the left end of the AS path (the end where AS paths have been most recently added). The routing device stops searching for private ASs when it finds the first nonprivate AS or a peer's private AS. If the AS path contains the AS number of the external BGP (EBGP) neighbor, BGP does not remove the private AS number.



**NOTE:** As of Junos OS 10.0R2 and later, if there is a need to send prefixes to an EBGP peer that has an AS number that matches an AS number in the AS path, consider using the **as-override** statement instead of the **remove-private** statement.

The operation takes place after any confederation member ASs have already been removed from the AS path, if applicable.

The software is preconfigured with knowledge of the set of AS numbers that is considered private, a range that is defined in the Internet Assigned Numbers Authority (IANA) assigned numbers document. The set of AS numbers reserved as private are in the range from 64,512 through 65,534, inclusive.

### Example: Removing Private AS Numbers from AS Paths

This example demonstrates the removal of a private AS number from the advertised AS path to avoid announcing the private AS number to the Internet.

- [Requirements on page 3805](#)
- [Overview on page 3805](#)
- [Configuration on page 3806](#)
- [Verification on page 3808](#)

#### Requirements


No special configuration beyond device initialization is required before you configure this example.

#### Overview

Service providers and enterprise networks use the **remove-private** statement to prevent advertising private AS numbers to the Internet. The **remove-private** statement works in the outbound direction. You configure the **remove-private** statement on a device that has a public AS number and that is connected to one or more devices that have private AS numbers. Generally, you would not configure this statement on a device that has a private AS number.

Figure 76 shows the sample topology.

**Figure 76: Topology for Removing a Private AS from the Advertised AS Path**



In this example, Device R1 is connected to its service provider using private AS number 65530. The example shows the **remove-private** statement configured on Device ISP to prevent Device R1's private AS number from being announced to Device R2. Device R2 sees only the AS number of the service provider.

## Configuration

---

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Device R1**

```
set interfaces fe-1/2/0 unit 1 family inet address 192.168.10.1/24
set interfaces lo0 unit 1 family inet address 10.10.10.1/32
set protocols bgp group ext type external
set protocols bgp group ext export send-direct
set protocols bgp group ext export send-static
set protocols bgp group ext peer-as 100
set protocols bgp group ext neighbor 192.168.10.10
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set policy-options policy-statement send-static term 1 from protocol static
set policy-options policy-statement send-static term 1 then accept
set routing-options static route 192.168.20.0/24 next-hop 192.168.10.10
set routing-options autonomous-system 65530
```

**Device ISP**

```
set interfaces fe-1/2/0 unit 2 family inet address 192.168.10.10/24
set interfaces fe-1/2/1 unit 3 family inet address 192.168.20.20/24
set interfaces lo0 unit 2 family inet address 10.10.0.1/32
set protocols bgp group ext type external
set protocols bgp group ext neighbor 192.168.10.1 peer-as 65530
set protocols bgp group ext neighbor 192.168.20.1 remove-private
set protocols bgp group ext neighbor 192.168.20.1 peer-as 200
set routing-options autonomous-system 100
```

**Device R2**

```
set interfaces fe-1/2/0 unit 4 family inet address 192.168.20.1/24
set interfaces lo0 unit 3 family inet address 10.10.20.1/32
set protocols bgp group ext type external
set protocols bgp group ext export send-direct
set protocols bgp group ext export send-static
set protocols bgp group ext peer-as 100
set protocols bgp group ext neighbor 192.168.20.20
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set policy-options policy-statement send-static term 1 from protocol static
set policy-options policy-statement send-static term 1 then accept
set routing-options static route 192.168.10.0/24 next-hop 192.168.20.20
set routing-options autonomous-system 200
```

### Device ISP

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device ISP:

1. Configure the interfaces.

[edit interfaces]

```

user@ISP# set fe-1/2/0 unit 2 family inet address 192.168.10.10/24
user@ISP# set fe-1/2/1 unit 3 family inet address 192.168.20.20/24
user@ISP# set lo0 unit 2 family inet address 10.10.0.1/32

```

2. Configure EBGP.

```

[edit protocols bgp group ext]
user@ISP# set type external
user@ISP# set neighbor 192.168.10.1 peer-as 65530
user@ISP# set neighbor 192.168.20.1 peer-as 200

```

3. For the neighbor in autonomous system (AS) 200 (Device R2), remove private AS numbers from the advertised AS paths.

```

[edit protocols bgp group ext]
user@ISP# set neighbor 192.168.20.1 remove-private

```

4. Configure the AS number.

```

[edit routing-options]
user@ISP# set autonomous-system 100

```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@ISP# show interfaces
fe-1/2/0 {
 unit 2 {
 family inet {
 address 192.168.10.10/24;
 }
 }
}
fe-1/2/1 {
 unit 3 {
 family inet {
 address 192.168.20.20/24;
 }
 }
}
lo0 {
 unit 2 {
 family inet {
 address 10.10.0.1/32;
 }
 }
}
}

user@ISP# show protocols
bgp {
 group ext {
 type external;
 neighbor 192.168.10.1 {
 peer-as 65530;
 }
 neighbor 192.168.20.1 {
 remove-private;
 }
 }
}

```

```
 peer-as 200;
 }
}

user@ISP# show routing-options
autonomous-system 100;
```

If you are done configuring the device, enter **commit** from configuration mode. Repeat the configuration on Device R1 and Device R2, changing the interface names and IP address, as needed, and adding the routing policy configuration.

---

### Verification

Confirm that the configuration is working properly.

- [Checking the Neighbor Status on page 3808](#)
- [Checking the Routing Tables on page 3809](#)
- [Checking the AS Path When the remove-private Statement Is Deactivated on page 3809](#)

#### *Checking the Neighbor Status*

**Purpose** Make sure that Device ISP has the **remove-private** setting enabled in its neighbor session with Device R2.

**Action** From operational mode, enter the **show bgp neighbor 192.168.20.1** command.

```
user@ISP> show bgp neighbor 192.168.20.1
Peer: 192.168.20.1+179 AS 200 Local: 192.168.20.20+60216 AS 100
 Type: External State: Established Flags: <ImportEval Sync>
 Last State: OpenConfirm Last Event: RecvKeepAlive
 Last Error: None
 Options: <Preference RemovePrivateAS PeerAS Refresh>
 Holdtime: 90 Preference: 170
 Number of flaps: 0
 Peer ID: 10.10.20.1 Local ID: 10.10.0.1 Active Holdtime: 90
 Keepalive Interval: 30 Peer index: 0
 BFD: disabled, down
 Local Interface: fe-1/2/1.3
 NLRI for restart configured on peer: inet-unicast
 NLRI advertised by peer: inet-unicast
 NLRI for this session: inet-unicast
 Peer supports Refresh capability (2)
 Stale routes from peer are kept for: 300
 Peer does not support Restarter functionality
 NLRI that restart is negotiated for: inet-unicast
 NLRI of received end-of-rib markers: inet-unicast
 NLRI of all end-of-rib markers sent: inet-unicast
 Peer supports 4 byte AS extension (peer-as 200)
 Peer does not support Addpath
 Table inet.0 Bit: 10001
 RIB State: BGP restart is complete
 Send state: in sync
 Active prefixes: 1
 Received prefixes: 3
 Accepted prefixes: 2
 Suppressed due to damping: 0
 Advertised prefixes: 1
```



```

Last traffic (seconds): Received 10 Sent 16 Checked 55
Input messages: Total 54 Updates 3 Refreshes 0 Octets 1091
Output messages: Total 54 Updates 1 Refreshes 0 Octets 1118
Output Queue[0]: 0

```

**Meaning** The `RemovePrivateAS` option shows that Device ISP has the expected setting.

### *Checking the Routing Tables*

**Purpose** Make sure that the devices have the expected routes and AS paths.

**Action** From operational mode, enter the `show route protocol bgp` command.

```

user@R1> show route protocol bgp
inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.20.1/32 *[BGP/170] 00:28:57, localpref 100
 AS path: 100 200 I
 > to 192.168.10.10 via fe-1/2/0.1

user@ISP> show route protocol bgp

inet.0: 7 destinations, 11 routes (7 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.10.1/32 *[BGP/170] 00:29:40, localpref 100
 AS path: 65530 I
 > to 192.168.10.1 via fe-1/2/0.2
10.10.20.1/32 *[BGP/170] 00:29:36, localpref 100
 AS path: 200 I
 > to 192.168.20.1 via fe-1/2/1.3
192.168.10.0/24 [BGP/170] 00:29:40, localpref 100
 AS path: 65530 I
 > to 192.168.10.1 via fe-1/2/0.2
192.168.20.0/24 [BGP/170] 00:29:36, localpref 100
 AS path: 200 I
 > to 192.168.20.1 via fe-1/2/1.3

user@R2> show route protocol bgp
inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.10.1/32 *[BGP/170] 00:29:53, localpref 100
 AS path: 100 I
 > to 192.168.20.20 via fe-1/2/0.4

```

**Meaning** Device ISP has the private AS number 65530 in its AS path to Device R1. However, Device ISP does not advertise this private AS number to Device R2. This is shown in the routing table of Device R2. Device R2's path to Device R1 contains only the AS number for Device ISP.

### *Checking the AS Path When the remove-private Statement Is Deactivated*

**Purpose** Verify that without the `remove-private` statement, the private AS number appears in Device R2's routing table.

**Action** From configuration mode on Device ISP, enter the **deactivate remove-private** command and then recheck the routing table on Device R2.

```
[protocols bgp group ext neighbor 192.168.20.1]
user@ISP# deactivate remove-private
user@ISP# commit

user@R2> show route protocol bgp
inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.10.1/32 *[BGP/170] 00:00:54, localpref 100
 AS path: 100 65530 I
 > to 192.168.20.20 via fe-1/2/0.4
```

**Meaning** Private AS number 65530 appears in Device R2's AS path to Device R1.

**Related Documentation**

- [Understanding External BGP Peering Sessions on page 3601](#)
- *BGP Configuration Overview*

# BGP BFD Configuration

- [Example: Configuring BFD for BGP on page 3811](#)
- [Example: Configuring BFD Authentication for BGP on page 3820](#)

## Example: Configuring BFD for BGP

---

- [Understanding BFD for BGP on page 3811](#)
- [Example: Configuring BFD on Internal BGP Peer Sessions on page 3812](#)

## Understanding BFD for BGP

The Bidirectional Forwarding Detection (BFD) protocol is a simple hello mechanism that detects failures in a network. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the routing device stops receiving a reply after a specified interval. BFD works with a wide variety of network environments and topologies. The failure detection timers for BFD have shorter time limits than default failure detection mechanisms for BGP, so they provide faster detection.

The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the **clear bfd adaptation** command to return BFD interval timers to their configured values. The **clear bfd adaptation** command is hitless, meaning that the command does not affect traffic flow on the routing device.

In Junos OS Release 8.3 and later, BFD is supported on internal BGP (IBGP) and multihop external BGP (EBGP) sessions as well as on single-hop EBGP sessions. In Junos OS Release 9.1 through Junos OS Release 11.1, BFD supports IPv6 interfaces in static routes only. In Junos OS Release 11.2 and later, BFD supports IPv6 interfaces with BGP.

## Example: Configuring BFD on Internal BGP Peer Sessions

This example shows how to configure internal BGP (IBGP) peer sessions with the Bidirectional Forwarding Detection (BFD) protocol to detect failures in a network.

- [Requirements on page 3812](#)
- [Overview on page 3812](#)
- [Configuration on page 3813](#)
- [Verification on page 3817](#)

---

### Requirements

No special configuration beyond device initialization is required before you configure this example.

---

### Overview

The minimum configuration to enable BFD on IBGP sessions is to include the **bfd-liveness-detection minimum-interval** statement in the BGP configuration of all neighbors participating in the BFD session. The **minimum-interval** statement specifies the minimum transmit and receive intervals for failure detection. Specifically, this value represents the minimum interval after which the local routing device transmits hello packets as well as the minimum interval that the routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a value from 1 through 255,000 milliseconds.

Optionally, you can specify the minimum transmit and receive intervals separately using the **transmit-interval minimum-interval** and **minimum-receive-interval** statements. For information about these and other optional BFD configuration statements, see [bfd-liveness-detection](#).



**NOTE:** BFD is an intensive protocol that consumes system resources. Specifying a minimum interval for BFD less than 100 ms for Routing Engine-based sessions and less than 10 ms for distributed BFD sessions can cause undesired BFD flapping.


Depending on your network environment, these additional recommendations might apply:

- For large-scale network deployments with a large number of BFD sessions, specify a minimum interval of 300 ms for Routing Engine-based sessions and 100 ms for distributed BFD sessions.
- For very large-scale network deployments with a large number of BFD sessions, contact Juniper Networks customer support for more information.
- For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing (NSR) is configured, specify a minimum interval of 2500 ms for Routing Engine-based sessions. For distributed BFD sessions with NSR configured, the minimum interval recommendations are unchanged and depend only on your network deployment.

BFD is supported on the default routing instance (the main router), routing instances, and logical systems. This example shows BFD on logical systems.

Figure 77 shows a typical network with internal peer sessions.

**Figure 77: Typical Network with IBGP Sessions**



### Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Device A**     **set logical-systems A interfaces lt-1/2/0 unit 1 description to-B**  
                  **set logical-systems A interfaces lt-1/2/0 unit 1 encapsulation ethernet**  
                  **set logical-systems A interfaces lt-1/2/0 unit 1 peer-unit 2**

```
set logical-systems A interfaces lt-1/2/0 unit 1 family inet address 10.10.10.1/30
set logical-systems A interfaces lo0 unit 1 family inet address 192.168.6.5/32
set logical-systems A protocols bgp group internal-peers type internal
set logical-systems A protocols bgp group internal-peers traceoptions file bgp-bfd
set logical-systems A protocols bgp group internal-peers traceoptions flag bfd detail
set logical-systems A protocols bgp group internal-peers local-address 192.168.6.5
set logical-systems A protocols bgp group internal-peers export send-direct
set logical-systems A protocols bgp group internal-peers bfd-liveness-detection
 minimum-interval 1000
set logical-systems A protocols bgp group internal-peers neighbor 192.163.6.4
set logical-systems A protocols bgp group internal-peers neighbor 192.168.40.4
set logical-systems A protocols ospf area 0.0.0.0 interface lo0.1 passive
set logical-systems A protocols ospf area 0.0.0.0 interface lt-1/2/0.1
set logical-systems A policy-options policy-statement send-direct term 2 from protocol
 direct
set logical-systems A policy-options policy-statement send-direct term 2 then accept
set logical-systems A routing-options router-id 192.168.6.5
set logical-systems A routing-options autonomous-system 17
```

**Device B**

```
set logical-systems B interfaces lt-1/2/0 unit 2 description to-A
set logical-systems B interfaces lt-1/2/0 unit 2 encapsulation ethernet
set logical-systems B interfaces lt-1/2/0 unit 2 peer-unit 1
set logical-systems B interfaces lt-1/2/0 unit 2 family inet address 10.10.10.2/30
set logical-systems B interfaces lt-1/2/0 unit 5 description to-C
set logical-systems B interfaces lt-1/2/0 unit 5 encapsulation ethernet
set logical-systems B interfaces lt-1/2/0 unit 5 peer-unit 6
set logical-systems B interfaces lt-1/2/0 unit 5 family inet address 10.10.10.5/30
set logical-systems B interfaces lo0 unit 2 family inet address 192.163.6.4/32
set logical-systems B protocols bgp group internal-peers type internal
set logical-systems B protocols bgp group internal-peers local-address 192.163.6.4
set logical-systems B protocols bgp group internal-peers export send-direct
set logical-systems B protocols bgp group internal-peers bfd-liveness-detection
 minimum-interval 1000
set logical-systems B protocols bgp group internal-peers neighbor 192.168.40.4
set logical-systems B protocols bgp group internal-peers neighbor 192.168.6.5
set logical-systems B protocols ospf area 0.0.0.0 interface lo0.2 passive
set logical-systems B protocols ospf area 0.0.0.0 interface lt-1/2/0.2
set logical-systems B protocols ospf area 0.0.0.0 interface lt-1/2/0.5
set logical-systems B policy-options policy-statement send-direct term 2 from protocol
 direct
set logical-systems B policy-options policy-statement send-direct term 2 then accept
set logical-systems B routing-options router-id 192.163.6.4
set logical-systems B routing-options autonomous-system 17
```

**Device C**

```
set logical-systems C interfaces lt-1/2/0 unit 6 description to-B
set logical-systems C interfaces lt-1/2/0 unit 6 encapsulation ethernet
set logical-systems C interfaces lt-1/2/0 unit 6 peer-unit 5
set logical-systems C interfaces lt-1/2/0 unit 6 family inet address 10.10.10.6/30
set logical-systems C interfaces lo0 unit 3 family inet address 192.168.40.4/32
set logical-systems C protocols bgp group internal-peers type internal
set logical-systems C protocols bgp group internal-peers local-address 192.168.40.4
set logical-systems C protocols bgp group internal-peers export send-direct
set logical-systems C protocols bgp group internal-peers bfd-liveness-detection
 minimum-interval 1000
set logical-systems C protocols bgp group internal-peers neighbor 192.163.6.4
```

```

set logical-systems C protocols bgp group internal-peers neighbor 192.168.6.5
set logical-systems C protocols ospf area 0.0.0.0 interface lo0.3 passive
set logical-systems C protocols ospf area 0.0.0.0 interface lt-1/2/0.6
set logical-systems C policy-options policy-statement send-direct term 2 from protocol
 direct
set logical-systems C policy-options policy-statement send-direct term 2 then accept
set logical-systems C routing-options router-id 192.168.40.4
set logical-systems C routing-options autonomous-system 17

```

### Configuring Device A

**Step-by-Step Procedure** The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device A:

1. Set the CLI to Logical System A.

```
user@host> set cli logical-system A
```

2. Configure the interfaces.

```

[edit interfaces lt-1/2/0 unit 1]
user@host:A# set description to-B
user@host:A# set encapsulation ethernet
user@host:A# set peer-unit 2
user@host:A# set family inet address 10.10.10.1/30

```

```

[edit interfaces lo0 unit 1]
user@host:A# set family inet address 192.168.6.5/32

```

3. Configure BGP.

The **neighbor** statements are included for both Device B and Device C, even though Device A is not directly connected to Device C.

```

[edit protocols bgp group internal-peers]
user@host:A# set type internal
user@host:A# set local-address 192.168.6.5
user@host:A# set export send-direct
user@host:A# set neighbor 192.163.6.4
user@host:A# set neighbor 192.168.40.4

```

4. Configure BFD.

```

[edit protocols bgp group internal-peers]
user@host:A# set bfd-liveness-detection minimum-interval 1000

```

You must configure the same minimum interval on the connecting peer.

5. (Optional) Configure BFD tracing.

```

[edit protocols bgp group internal-peers]
user@host:A# set traceoptions file bgp-bfd
user@host:A# set traceoptions flag bfd detail

```

6. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
```

```
user@host:A# set interface lo0.1 passive
user@host:A# set interface lt-1/2/0.1
```

7. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 2]
user@host:A# set from protocol direct
user@host:A# set then accept
```

8. Configure the router ID and the autonomous system (AS) number.

```
[edit routing-options]
user@host:A# set router-id 192.168.6.5
user@host:A# set autonomous-system 17
```

9. If you are done configuring the device, enter **commit** from configuration mode. Repeat these steps to configure Device B and Device C.

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host:A# show interfaces
lt-1/2/0 {
 unit 1 {
 description to-B;
 encapsulation ethernet;
 peer-unit 2;
 family inet {
 address 10.10.10.1/30;
 }
 }
}
lo0 {
 unit 1 {
 family inet {
 address 192.168.6.5/32;
 }
 }
}

user@host:A# show policy-options
policy-statement send-direct {
 term 2 {
 from protocol direct;
 then accept;
 }
}

user@host:A# show protocols
bgp {
 group internal-peers {
 type internal;
 traceoptions {
```



```

 file bgp-bfd;
 flag bfd detail;
 }
 local-address 192.168.6.5;
 export send-direct;
 bfd-liveness-detection {
 minimum-interval 1000;
 }
 neighbor 192.163.6.4;
 neighbor 192.168.40.4;
}
}
ospf {
 area 0.0.0.0 {
 interface lo0.1 {
 passive;
 }
 interface lt-1/2/0.1;
 }
}

user@host:A# show routing-options
router-id 192.168.6.5;
autonomous-system 17;

```

### Verification

Confirm that the configuration is working properly.

- [Verifying That BFD Is Enabled on page 3817](#)
- [Verifying That BFD Sessions Are Up on page 3818](#)
- [Viewing Detailed BFD Events on page 3818](#)
- [Viewing Detailed BFD Events After Deactivating and Reactivating a Loopback Interface on page 3819](#)

#### *Verifying That BFD Is Enabled*

**Purpose** Verify that BFD is enabled between the IBGP peers.

**Action** From operational mode, enter the **show bgp neighbor** command. You can use the **| match bfd** filter to narrow the output.

```

user@host:A> show bgp neighbor | match bfd
Options: <BfdEnabled>
BFD: enabled, up
Trace file: /var/log/A/bgp-bfd size 131072 files 10
Options: <BfdEnabled>
BFD: enabled, up
Trace file: /var/log/A/bgp-bfd size 131072 files 10

```

**Meaning** The output shows that Logical System A has two neighbors with BFD enabled. When BFD is not enabled, the output displays **BFD: disabled, down**, and the **<BfdEnabled>** option is absent. If BFD is enabled and the session is down, the output displays **BFD: enabled**,

**down.** The output also shows that BFD-related events are being written to a log file because trace operations are configured.

### *Verifying That BFD Sessions Are Up*

**Purpose** Verify that the BFD sessions are up, and view details about the BFD sessions.

**Action** From operational mode, enter the **show bfd session extensive** command.

```
user@host: A> show bfd session extensive
```

| Address                                                              | State | Interface | Detect Time | Transmit Interval | Multiplier |
|----------------------------------------------------------------------|-------|-----------|-------------|-------------------|------------|
| 192.163.6.4                                                          | Up    |           | 3.000       | 1.000             | 3          |
| Client BGP, TX interval 1.000, RX interval 1.000                     |       |           |             |                   |            |
| Session up time 00:54:40                                             |       |           |             |                   |            |
| Local diagnostic None, remote diagnostic None                        |       |           |             |                   |            |
| Remote state Up, version 1                                           |       |           |             |                   |            |
| Logical system 12, routing table index 25                            |       |           |             |                   |            |
| Min async interval 1.000, min slow interval 1.000                    |       |           |             |                   |            |
| Adaptive async TX interval 1.000, RX interval 1.000                  |       |           |             |                   |            |
| Local min TX interval 1.000, minimum RX interval 1.000, multiplier 3 |       |           |             |                   |            |
| Remote min TX interval 1.000, min RX interval 1.000, multiplier 3    |       |           |             |                   |            |
| Local discriminator 10, remote discriminator 9                       |       |           |             |                   |            |
| Echo mode disabled/inactive                                          |       |           |             |                   |            |
| Multi-hop route table 25, local-address 192.168.6.5                  |       |           |             |                   |            |

| Address                                                              | State | Interface | Detect Time | Transmit Interval | Multiplier |
|----------------------------------------------------------------------|-------|-----------|-------------|-------------------|------------|
| 192.168.40.4                                                         | Up    |           | 3.000       | 1.000             | 3          |
| Client BGP, TX interval 1.000, RX interval 1.000                     |       |           |             |                   |            |
| Session up time 00:48:03                                             |       |           |             |                   |            |
| Local diagnostic None, remote diagnostic None                        |       |           |             |                   |            |
| Remote state Up, version 1                                           |       |           |             |                   |            |
| Logical system 12, routing table index 25                            |       |           |             |                   |            |
| Min async interval 1.000, min slow interval 1.000                    |       |           |             |                   |            |
| Adaptive async TX interval 1.000, RX interval 1.000                  |       |           |             |                   |            |
| Local min TX interval 1.000, minimum RX interval 1.000, multiplier 3 |       |           |             |                   |            |
| Remote min TX interval 1.000, min RX interval 1.000, multiplier 3    |       |           |             |                   |            |
| Local discriminator 14, remote discriminator 13                      |       |           |             |                   |            |
| Echo mode disabled/inactive                                          |       |           |             |                   |            |
| Multi-hop route table 25, local-address 192.168.6.5                  |       |           |             |                   |            |

2 sessions, 2 clients

Cumulative transmit rate 2.0 pps, cumulative receive rate 2.0 pps

**Meaning** The TX interval 1.000, RX interval 1.000 output represents the setting configured with the **minimum-interval** statement. All of the other output represents the default settings for BFD. To modify the default settings, include the optional statements under the [bfd-liveness-detection](#) statement.

### *Viewing Detailed BFD Events*

**Purpose** View the contents of the BFD trace file to assist in troubleshooting, if needed.

**Action** From operational mode, enter the **file show /var/log/A/bgp-bfd** command.

```
user@host: A> file show /var/log/A/bgp-bfd
```

```

Aug 15 17:07:25 trace_on: Tracing to "/var/log/A/bgp-bfd" started
Aug 15 17:07:26.492190 bgp_peer_init: BGP peer 192.163.6.4 (Internal AS 17) local
address 192.168.6.5 not found. Leaving peer idled
Aug 15 17:07:26.493176 bgp_peer_init: BGP peer 192.168.40.4 (Internal AS 17) local
address 192.168.6.5 not found. Leaving peer idled
Aug 15 17:07:32.597979 task_connect: task BGP_17.192.163.6.4+179 addr
192.163.6.4+179: No route to host
Aug 15 17:07:32.599623 bgp_connect_start: connect 192.163.6.4 (Internal AS 17):
No route to host
Aug 15 17:07:36.869394 task_connect: task BGP_17.192.168.40.4+179 addr
192.168.40.4+179: No route to host
Aug 15 17:07:36.870624 bgp_connect_start: connect 192.168.40.4 (Internal AS 17):
No route to host
Aug 15 17:08:04.599220 task_connect: task BGP_17.192.163.6.4+179 addr
192.163.6.4+179: No route to host
Aug 15 17:08:04.601135 bgp_connect_start: connect 192.163.6.4 (Internal AS 17):
No route to host
Aug 15 17:08:08.869717 task_connect: task BGP_17.192.168.40.4+179 addr
192.168.40.4+179: No route to host
Aug 15 17:08:08.869934 bgp_connect_start: connect 192.168.40.4 (Internal AS 17):
No route to host
Aug 15 17:08:36.603544 advertising receiving-speaker only capability to neighbor
192.163.6.4 (Internal AS 17)
Aug 15 17:08:36.606726 bgp_read_message: 192.163.6.4 (Internal AS 17): 0 bytes
buffered
Aug 15 17:08:36.609119 Initiated BFD session to peer 192.163.6.4 (Internal AS
17): address=192.163.6.4 ifindex=0 ifname=(none) txivl=1000 rxivl=1000 mult=3
ver=255
Aug 15 17:08:36.734033 advertising receiving-speaker only capability to neighbor
192.168.40.4 (Internal AS 17)
Aug 15 17:08:36.738436 Initiated BFD session to peer 192.168.40.4 (Internal AS
17): address=192.168.40.4 ifindex=0 ifname=(none) txivl=1000 rxivl=1000 mult=3
ver=255
Aug 15 17:08:40.537552 BFD session to peer 192.163.6.4 (Internal AS 17) up
Aug 15 17:08:40.694410 BFD session to peer 192.168.40.4 (Internal AS 17) up

```

**Meaning** Before the routes are established, the **No route to host** message appears in the output. After the routes are established, the last two lines show that both BFD sessions come up.

#### *Viewing Detailed BFD Events After Deactivating and Reactivating a Loopback Interface*

**Purpose** Check to see what happens after bringing down a router or switch and then bringing it back up. To simulate bringing down a router or switch, deactivate the loopback interface on Logical System B.

**Action** 1. From configuration mode, enter the **deactivate logical-systems B interfaces lo0 unit 2 family inet** command.

```

user@host:A# deactivate logical-systems B interfaces lo0 unit 2 family inet
user@host:A# commit

```

2. From operational mode, enter the **file show /var/log/A/bgp-bfd** command.

```

user@host:A> file show /var/log/A/bgp-bfd
...
Aug 15 17:20:55.995648 bgp_read_v4_message:9747: NOTIFICATION received from
192.163.6.4 (Internal AS 17): code 6 (Cease) subcode 6 (Other Configuration
Change)

```

```
Aug 15 17:20:56.004508 Terminated BFD session to peer 192.163.6.4 (Internal AS 17)
Aug 15 17:21:28.007755 task_connect: task BGP_17.192.163.6.4+179 addr 192.163.6.4+179: No route to host
Aug 15 17:21:28.008597 bgp_connect_start: connect 192.163.6.4 (Internal AS 17): No route to host
```

3. From configuration mode, enter the **activate logical-systems B interfaces lo0 unit 2 family inet** command.

```
user@host:A# activate logical-systems B interfaces lo0 unit 2 family inet
user@host:A# commit
```

4. From operational mode, enter the **file show /var/log/A/bgp-bfd** command.

```
user@host:A> file show /var/log/A/bgp-bfd
...
Aug 15 17:25:53.623743 advertising receiving-speaker only capability to neighbor 192.163.6.4 (Internal AS 17)
Aug 15 17:25:53.631314 Initiated BFD session to peer 192.163.6.4 (Internal AS 17): address=192.163.6.4 ifindex=0 ifname=(none) txivl=1000 rxivl=1000 mult=3 ver=255
Aug 15 17:25:57.570932 BFD session to peer 192.163.6.4 (Internal AS 17) up
```

- Related Documentation**
- [Understanding External BGP Peering Sessions on page 3601](#)
  - [BGP Configuration Overview](#)

---

## Example: Configuring BFD Authentication for BGP

---

- [Understanding BFD Authentication for BGP on page 3820](#)
- [Example: Configuring BFD Authentication for BGP on page 3822](#)

### Understanding BFD Authentication for BGP

Bidirectional Forwarding Detection protocol (BFD) enables rapid detection of communication failures between adjacent systems. By default, authentication for BFD sessions is disabled. However, when you run BFD over Network Layer protocols, the risk of service attacks can be significant. We strongly recommend using authentication if you are running BFD over multiple hops or through insecure tunnels. Beginning with Junos OS Release 9.6, Junos OS supports authentication for BFD sessions running over BGP. BFD authentication is not supported on MPLS OAM sessions. BFD authentication is only supported in the Canada and United States version of the Junos OS image and is not available in the export version.

You authenticate BFD sessions by specifying an authentication algorithm and keychain, and then associating that configuration information with a security authentication keychain using the keychain name.

The following sections describe the supported authentication algorithms, security keychains, and level of authentication that can be configured:

- [BFD Authentication Algorithms on page 3821](#)
- [Security Authentication Keychains on page 3822](#)
- [Strict Versus Loose Authentication on page 3822](#)

### BFD Authentication Algorithms

---

Junos OS supports the following algorithms for BFD authentication:

- **simple-password**—Plain-text password. One to 16 bytes of plain text are used to authenticate the BFD session. One or more passwords can be configured. This method is the least secure and should be used only when BFD sessions are not subject to packet interception.
- **keyed-md5**—Keyed Message Digest 5 hash algorithm for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed MD5 uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than or equal to the last sequence number received. Although more secure than a simple password, this method is vulnerable to replay attacks. Increasing the rate at which the sequence number is updated can reduce this risk.
- **meticulous-keyed-md5**—Meticulous keyed Message Digest 5 hash algorithm. This method works in the same manner as keyed MD5, but the sequence number is updated with every packet. Although more secure than keyed MD5 and simple passwords, this method might take additional time to authenticate the session.
- **keyed-sha-1**—Keyed Secure Hash Algorithm I for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed SHA uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. The key is not carried within the packets. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than the last sequence number received.
- **meticulous-keyed-sha-1**—Meticulous keyed Secure Hash Algorithm I. This method works in the same manner as keyed SHA, but the sequence number is updated with every packet. Although more secure than keyed SHA and simple passwords, this method might take additional time to authenticate the session.



**NOTE:** Nonstop active routing (NSR) is not supported with meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

---

## Security Authentication Keychains

---

The security authentication keychain defines the authentication attributes used for authentication key updates. When the security authentication keychain is configured and associated with a protocol through the keychain name, authentication key updates can occur without interrupting routing and signaling protocols.

The authentication keychain contains one or more keychains. Each keychain contains one or more keys. Each key holds the secret data and the time at which the key becomes valid. The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

BFD allows multiple clients per session, and each client can have its own keychain and algorithm defined. To avoid confusion, we recommend specifying only one security authentication keychain.

## Strict Versus Loose Authentication

---

By default, strict authentication is enabled and authentication is checked at both ends of each BFD session. Optionally, to smooth migration from nonauthenticated sessions to authenticated sessions, you can configure *loose checking*. When loose checking is configured, packets are accepted without authentication being checked at each end of the session. This feature is intended for transitional periods only.

## Example: Configuring BFD Authentication for BGP

Beginning with Junos OS Release 9.6, you can configure authentication for BFD sessions running over BGP. Only three steps are needed to configure authentication on a BFD session:

1. Specify the BFD authentication algorithm for the BGP protocol.
2. Associate the authentication keychain with the BGP protocol.
3. Configure the related security authentication keychain.

The following sections provide instructions for configuring and viewing BFD authentication on BGP:

- [Configuring BFD Authentication Parameters on page 3822](#)
- [Viewing Authentication Information for BFD Sessions on page 3824](#)

## Configuring BFD Authentication Parameters

---

BFD authentication can be configured for the entire BGP protocol, or a specific BGP group, neighbor, or routing instance.

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure BFD authentication:

1. Specify the algorithm (**keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**, or **simple-password**) to use.

```
[edit]
user@host# set protocols bgp bfd-liveness-detection authentication algorithm
keyed-sha-1
user@host# set protocols bgp group bgp-gr1 bfd-liveness-detection authentication
algorithm keyed-sha-1
user@host# set protocols bgp group bgp-gr1 neighbor 10.10.10.7 bfd-liveness-detection
authentication algorithm keyed-sha-1
```



**NOTE:** Nonstop active routing is not supported with meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

2. Specify the keychain to be used to associate BFD sessions on BGP with the unique security authentication keychain attributes.

The keychain name you specify must match a keychain name configured at the **[edit security authentication key-chains]** hierarchy level.

```
[edit]
user@host# set protocols bgp bfd-liveness-detection authentication keychain bfd-bgp
user@host# set protocols bgp group bgp-gr1 bfd-liveness-detection authentication
keychain bfd-bgp
user@host# set protocols bgp group bgp-gr1 neighbor 10.10.10.7 bfd-liveness-detection
authentication keychain bfd-bgp
```



**NOTE:** The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

3. Specify the unique security authentication information for BFD sessions:
  - The matching keychain name as specified in Step 2.
  - At least one key, a unique integer between **0** and **63**. Creating multiple keys allows multiple clients to use the BFD session.
  - The secret data used to allow access to the session.
  - The time at which the authentication key becomes active, in the format *yyyy-mm-dd.hh:mm:ss*.

```
[edit security]
```

```
user@host# set authentication-key-chains key-chain bfd-bgp key 53 secret
9ggaJDmPQ6/tJgF/AtREvsyPsnCtUHm start-time 2009-06-14.10:00:00
```

- (Optional) Specify loose authentication checking if you are transitioning from nonauthenticated sessions to authenticated sessions.

```
[edit]
user@host# set protocols bgp bfd-liveness-detection authentication loose-check
user@host# set protocols bgp group bgp-gr1 bfd-liveness-detection authentication
loose-check
user@host# set protocols bgp group bgp-gr1 neighbor 10.10.10.7 bfd-liveness-detection
authentication loose-check
```

- (Optional) View your configuration using the **show bfd session detail** or **show bfd session extensive** command.
- Repeat these steps to configure the other end of the BFD session.



**NOTE:** BFD authentication is only supported in the Canada and United States version of the Junos OS image and is not available in the export version.

### Viewing Authentication Information for BFD Sessions

You can view the existing BFD authentication configuration using the **show bfd session detail** and **show bfd session extensive** commands.

The following example shows BFD authentication configured for the **bgp-gr1** BGP group. It specifies the keyed SHA-1 authentication algorithm and a keychain name of **bfd-bgp**. The authentication keychain is configured with two keys. Key 1 contains the secret data “\$9\$ggaJDmPQ6/tJgF/AtREvsyPsnCtUHm” and a start time of June 1, 2009, at 9:46:02 AM PST. Key 2 contains the secret data “\$9\$a5jiKW9l.reP38ny.TszF2/9” and a start time of June 1, 2009, at 3:29:20 PM PST.

```
[edit protocols bgp]
group bgp-gr1 {
 bfd-liveness-detection {
 authentication {
 algorithm keyed-sha-1;
 key-chain bfd-bgp;
 }
 }
}
[edit security]
authentication key-chains {
 key-chain bfd-bgp {
 key 1 {
 secret "9ggaJDmPQ6/tJgF/AtREvsyPsnCtUHm";
 start-time "2009-6-1.09:46:02 -0700";
 }
 key 2 {
 secret "9a5jiKW9l.reP38ny.TszF2/9";
 start-time "2009-6-1.15:29:20 -0700";
 }
 }
}
```



```
}

```

If you commit these updates to your configuration, you see output similar to the following. In the output for the **show bfd session detail** command, **Authenticate** is displayed to indicate that BFD authentication is configured. For more information about the configuration, use the **show bfd session extensive** command. The output for this command provides the keychain name, the authentication algorithm and mode for each client in the session, and the overall BFD authentication configuration status, keychain name, and authentication algorithm and mode.

#### show bfd session detail

```
user@host# show bfd session detail
```

| Address  | State | Interface  | Detect Time | Transmit Interval | Multiplier |
|----------|-------|------------|-------------|-------------------|------------|
| 50.0.0.2 | Up    | ge-0/1/5.0 | 0.900       | 0.300             | 3          |

Client BGP, TX interval 0.300, RX interval 0.300, **Authenticate**  
 Session up time 3d 00:34  
 Local diagnostic None, remote diagnostic NbrSignal  
 Remote state Up, version 1  
 Replicated

#### show bfd session extensive

```
user@host# show bfd session extensive
```

| Address  | State | Interface  | Detect Time | Transmit Interval | Multiplier |
|----------|-------|------------|-------------|-------------------|------------|
| 50.0.0.2 | Up    | ge-0/1/5.0 | 0.900       | 0.300             | 3          |

Client BGP, TX interval 0.300, RX interval 0.300, **Authenticate**  
**keychain bfd-bgp, algo keyed-sha-1, mode strict**  
 Session up time 00:04:42  
 Local diagnostic None, remote diagnostic NbrSignal  
 Remote state Up, version 1  
 Replicated  
 Min async interval 0.300, min slow interval 1.000  
 Adaptive async TX interval 0.300, RX interval 0.300  
 Local min TX interval 0.300, minimum RX interval 0.300, multiplier 3  
 Remote min TX interval 0.300, min RX interval 0.300, multiplier 3  
 Local discriminator 2, remote discriminator 2  
 Echo mode disabled/inactive  
**Authentication enabled/active, keychain bfd-bgp, algo keyed-sha-1, mode strict**

- Related Documentation**
- [Understanding External BGP Peering Sessions on page 3601](#)
  - [BGP Configuration Overview](#)



# BGP Load Balancing Configuration

- [Examples: Configuring BGP Multipath on page 3827](#)
- [Example: Advertising Multiple Paths in BGP on page 3844](#)
- [Configuring ECMP Next Hops for RSVP and LDP LSPs for Load Balancing on page 3870](#)

## Examples: Configuring BGP Multipath

---

- [Understanding BGP Multipath on page 3827](#)
- [Example: Load Balancing BGP Traffic on page 3828](#)
- [Example: Configuring Single-Hop EBGP Peers to Accept Remote Next Hops on page 3833](#)

## Understanding BGP Multipath

BGP multipath allows you to install multiple internal BGP paths and multiple external BGP paths to the forwarding table. Selecting multiple paths enables BGP to load-balance traffic across multiple links.

A path is considered a BGP equal-cost path (and is used for forwarding) if the BGP path selection process performs a tie-break after comparing the IGP cost to the next-hop. By default, all paths with the same neighboring AS, learned by a multipath-enabled BGP neighbor are considered in the multipath selection process.

BGP, typically selects only one best path for each prefix and installs that route in the forwarding table. When BGP multipath is enabled, the device selects multiple equal-cost BGP paths to reach a given destination, and all these paths are installed in the forwarding table. BGP advertises only the active path to its neighbors, unless add-path is in use.

The Junos OS BGP multipath feature supports the following applications:

- Load balancing across multiple links between two routing devices belonging to different autonomous systems (ASs)
- Load balancing across a common subnet or multiple subnets to different routing devices belonging to the same peer AS
- Load balancing across multiple links between two routing devices belonging to different external confederation peers
- Load balancing across a common subnet or multiple subnets to different routing devices belonging to external confederation peers

In a common scenario for load balancing, a customer is multihomed to multiple routers in a point of presence (POP). The default behavior is to send all traffic across only one of the available links. Load balancing causes traffic to use two or more of the links.

BGP multipath does not apply to paths that share the same MED-plus-IGP cost, yet differ in IGP cost. Multipath path selection is based on the IGP cost metric, even if two paths have the same MED-plus-IGP cost.

## Example: Load Balancing BGP Traffic

This example shows how to configure BGP to select multiple equal-cost external BGP (EBGP) or internal BGP (IBGP) paths as active paths.

- [Requirements on page 3828](#)
- [Overview on page 3828](#)
- [Configuration on page 3829](#)
- [Verification on page 3831](#)

---

### Requirements

Before you begin:

- Configure the device interfaces.
- Configure an interior gateway protocol (IGP).
- Configure BGP.
- Configure a routing policy that exports routes (such as direct routes or IGP routes) from the routing table into BGP.

---

### Overview

The following steps show how to configure per-packet load balancing:

1. Define a load-balancing routing policy by including one or more **policy-statement** statements at the **[edit policy-options]** hierarchy level, defining an action of **load-balance per-packet**:

```
policy-statement policy-name {
 from {
 match-conditions;
 route-filter destination-prefix match-type <actions>;
 prefix-list name;
 }
 then {
 load-balance per-packet;
 }
}
```

2. Apply the policy to routes exported from the routing table to the forwarding table. To do this, include the **forwarding-table** and **export** statements:

```
forwarding-table {
 export policy-name;
```

```
}
```

You cannot apply the export policy to VRF routing instances.

- Specify all next hops of that route, if more than one exists, when allocating a label corresponding to a route that is being advertised.
- Configure the forwarding-options hash key for MPLS to include the IP payload.




**NOTE:** On some platforms, you can increase the number of paths that are load balanced by using the `chassis maximum-ecmp` statement. With this statement, you can change the maximum number of equal-cost load-balanced paths to 32 or 64.

In this example, Device R1 is in AS 65000 and is connected to both Device R2 and Device R3, which are in AS 65001. This example shows the configuration on Device R1.

### Topology

Figure 78 shows the topology used in this example.

Figure 78: BGP Load Balancing



### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```

set protocols bgp group external type external
set protocols bgp group external peer-as 65001
set protocols bgp group external multipath
set protocols bgp group external neighbor 10.0.1.1
set protocols bgp group external neighbor 10.0.0.2
set policy-options policy-statement loadbal from route-filter 10.0.0.0/16 orlonger

```

```
set policy-options policy-statement loadbal then load-balance per-packet
set routing-options forwarding-table export loadbal
set routing-options autonomous-system 65000
```

**Step-by-Step Procedure** The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the BGP peer sessions:

1. Configure the BGP group.  

```
[edit protocols bgp group external]
user@R1# set type external
user@R1# set peer-as 65001
user@R1# set neighbor 10.0.1.1
user@R1# set neighbor 10.0.0.2
```
2. Enable the BGP group to use multiple paths.



**NOTE:** To disable the default check requiring that paths accepted by BGP multipath must have the same neighboring autonomous system (AS), include the `multiple-as` option.

- ```
[edit protocols bgp group external]
user@R1# set multipath
```
3. Configure the load-balancing policy.

```
[edit policy-options policy-statement loadbal]
user@R1# set from route-filter 10.0.0.0/16 orlonger
user@R1# set then load-balance per-packet
```
 4. Apply the load-balancing policy.

```
[edit routing-options]
user@R1# set forwarding-table export loadbal
```
 5. Configure the local autonomous system (AS) number.

```
[edit routing-options]
user@R1# set autonomous-system 65000
```

Results From configuration mode, confirm your configuration by entering the **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@R1# show protocols
bgp {
  group external {
    type external;
    peer-as 65001;
    multipath;
```

```

        neighbor 10.0.1.1;
        neighbor 10.0.0.2;
    }
}

[edit]
user@R1# show policy-options
policy-statement loadbal {
    from {
        route-filter 10.0.0.0/16 orlonger;
    }
    then {
        load-balance per-packet;
    }
}

[edit]
user@R1# show routing-options
autonomous-system 65000;
forwarding-table {
    export loadbal;
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly:

- [Verifying Routes on page 3831](#)
- [Verifying Forwarding on page 3832](#)

Verifying Routes

Purpose Verify that routes are learned from both routers in the neighboring AS.

Action From operational mode, run the **show route** command.

```

user@R1> show route 10.0.2.0
inet.0: 12 destinations, 15 routes (12 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

```

```

10.0.2.0/30          *[BGP/170] 03:12:32, localpref 100
                    AS path: 65001 I
                    to 10.0.1.1 via ge-1/2/0.0
                    > to 10.0.0.2 via ge-1/2/1.0
                    [BGP/170] 03:12:32, localpref 100
                    AS path: 65001 I
                    > to 10.0.1.1 via ge-1/2/0.0

```

```

user@R1> show route 10.0.2.0 detail
inet.0: 12 destinations, 15 routes (12 active, 0 holddown, 0 hidden)
10.0.2.0/30 (2 entries, 1 announced)
    *BGP      Preference: 170/-101
              Next hop type: Router, Next hop index: 262142
              Next-hop reference count: 3
              Source: 10.0.0.2
              Next hop: 10.0.1.1 via ge-1/2/0.0

```

```
Next hop: 10.0.0.2 via ge-1/2/1.0, selected
State: <Active Ext>
Local AS: 65000 Peer AS: 65001
Age: 3:18:30
Task: BGP_65001.10.0.0.2+55402
Announcement bits (1): 2-KRT
AS path: 65001 I
Accepted Multipath
Localpref: 100
Router ID: 192.168.2.1
BGP Preference: 170/-101
Next hop type: Router, Next hop index: 602
Next-hop reference count: 5
Source: 10.0.1.1
Next hop: 10.0.1.1 via ge-1/2/0.0, selected
State: <NotBest Ext>
Inactive reason: Not Best in its group - Active preferred
Local AS: 65000 Peer AS: 65001
Age: 3:18:30
Task: BGP_65001.10.0.1.1+53135
AS path: 65001 I
Accepted
Localpref: 100
Router ID: 192.168.3.1
```

Meaning The active path, denoted with an asterisk (*), has two next hops: 10.0.1.1 and 10.0.0.2 to the 10.0.2.0 destination. The 10.0.1.1 next hop is copied from the inactive path to the active path.



NOTE: The `show route detail` command output designates one gateway as selected. This output is potentially confusing in the context of load balancing. The selected gateway is used for many purposes in addition to deciding which gateway to install into the kernel when Junos OS is not performing per-packet load-balancing. For instance, the `ping mpls` command uses the selected gateway when sending packets. Multicast protocols use the selected gateway in some cases to determine the upstream interface. Therefore, even when Junos OS is performing per-packet load-balancing by way of a forwarding-table policy, the selected gateway information is still required for other purposes. It is useful to display the selected gateway for troubleshooting purposes. Additionally, it is possible to use forwarding-table policy to override what is installed into the kernel (for example, by using the `install-nexthop` action). In this case, the next-hop gateway installed in the forwarding table might be a subset of the total gateways displayed in the `show route` command.

Verifying Forwarding

Purpose Verify that both next hops are installed in the forwarding table.

Action From operational mode, run the `show route forwarding-table` command.

```
user@R1> show route forwarding-table destination 10.0.2.0
```



```

Routing table: default.inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
10.0.2.0/30      user  0          10.0.1.1          ucst  602    5 ge-1/2/0.0
                  10.0.0.2          ucst  522    6 ge-1/2/1.0

```

Example: Configuring Single-Hop EBGPeers to Accept Remote Next Hops

This example shows how to configure a single-hop external BGP (EBGP) peer to accept a remote next hop with which it does not share a common subnet.

- [Requirements on page 3833](#)
- [Overview on page 3833](#)
- [Configuration on page 3834](#)
- [Verification on page 3842](#)

Requirements

No special configuration beyond device initialization is required before you configure this example.

Overview

In some situations, it is necessary to configure a single-hop EBGPeer to accept a remote next hop with which it does not share a common subnet. The default behavior is for any next-hop address received from a single-hop EBGPeer that is not recognized as sharing a common subnet to be discarded. The ability to have a single-hop EBGPeer accept a remote next hop to which it is not directly connected also prevents you from having to configure the single-hop EBGPeer neighbor as a multihop session. When you configure a multihop session in this situation, all next-hop routes learned through this EBGPeer are labeled indirect even when they do share a common subnet. This situation breaks multipath functionality for routes that are recursively resolved over routes that include these next-hop addresses. Configuring the `accept-remote-nexthop` statement allows a single-hop EBGPeer to accept a remote next hop, which restores multipath functionality for routes that are resolved over these next-hop addresses. You can configure this statement at the global, group, and neighbor hierarchy levels for BGP. The statement is also supported on logical systems and the VPN routing and forwarding (VRF) routing instance type. Both the remote next-hop and the EBGPeer must support BGP route refresh as defined in RFC 2918, *Route Refresh Capability in BGP-4*. If the remote peer does not support BGP route refresh, the session is reset.



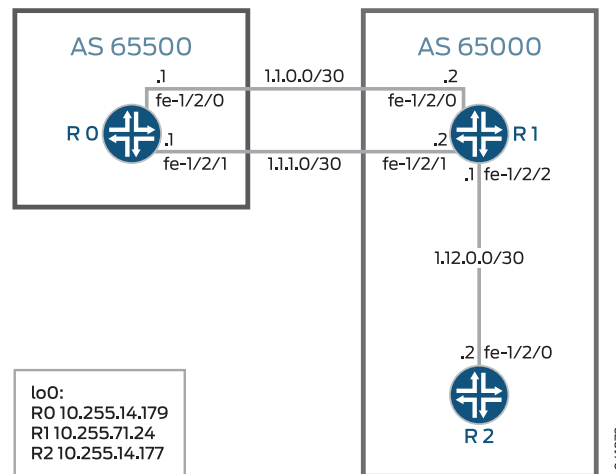
NOTE: You cannot configure both the `multihop` and `accept-remote-nexthop` statements for the same EBGPeer.

When you enable a single-hop EBGPeer to accept a remote next hop, you must also configure an import routing policy on the EBGPeer that specifies the remote next-hop address.

This example includes an import routing policy, **agg_route**, that enables a single-hop external BGP peer (Device R1) to accept the remote next-hop 1.1.10.10 for the route to the 1.1.230.0/23 network. At the **[edit protocols bgp]** hierarchy level, the example includes the **import agg_route** statement to apply the policy to the external BGP peer and includes the **accept-remote-nexthop** statement to enable the single-hop EBGP peer to accept the remote next hop.

Figure 79 shows the sample topology.

Figure 79: Topology for Accepting a Remote Next Hop



Configuration

- [Device R0 on page 3835](#)
- [Configuring Device R1 on page 3838](#)
- [Configuring Device R2 on page 3840](#)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
Device R0
set interfaces fe-1/2/0 unit 1 family inet address 1.1.0.1/30
set interfaces fe-1/2/1 unit 2 family inet address 1.1.1.1/30
set interfaces lo0 unit 1 family inet address 10.255.14.179/32
set protocols bgp group ext type external
set protocols bgp group ext export test_route
set protocols bgp group ext export agg_route
set protocols bgp group ext peer-as 65000
set protocols bgp group ext multipath
set protocols bgp group ext neighbor 1.1.0.2
set protocols bgp group ext neighbor 1.1.1.2
set policy-options policy-statement agg_route term 1 from protocol static
set policy-options policy-statement agg_route term 1 from route-filter 1.1.230.0/23 exact
set policy-options policy-statement agg_route term 1 then accept
set policy-options policy-statement test_route term 1 from protocol static
set policy-options policy-statement test_route term 1 from route-filter 1.1.10.10/32 exact
set policy-options policy-statement test_route term 1 then accept
```

```

set routing-options static route 1.1.10.10/32 reject
set routing-options static route 1.1.230.0/23 reject
set routing-options autonomous-system 65500

```

Device R1

```

set interfaces fe-1/2/0 unit 3 family inet address 1.1.0.2/30
set interfaces fe-1/2/1 unit 4 family inet address 1.12.0.1/30
set interfaces fe-1/2/2 unit 5 family inet address 1.1.1.2/30
set interfaces lo0 unit 2 family inet address 10.255.71.24/32
set protocols bgp accept-remote-nexthop
set protocols bgp group ext type external
set protocols bgp group ext import agg_route
set protocols bgp group ext peer-as 65500
set protocols bgp group ext multipath
set protocols bgp group ext neighbor 1.1.0.1
set protocols bgp group ext neighbor 1.1.1.1
set protocols bgp group int type internal
set protocols bgp group int local-address 10.255.71.24
set protocols bgp group int neighbor 10.255.14.177
set protocols ospf area 0.0.0.0 interface fe-1/2/1.4
set protocols ospf area 0.0.0.0 interface 10.255.71.24
set policy-options policy-statement agg_route term 1 from protocol bgp
set policy-options policy-statement agg_route term 1 from route-filter 1.1.230.0/23 exact
set policy-options policy-statement agg_route term 1 then next-hop 1.1.10.10
set policy-options policy-statement agg_route term 1 then accept
set routing-options autonomous-system 65000

```

Device R2

```

set interfaces fe-1/2/0 unit 6 family inet address 1.12.0.2/30
set interfaces lo0 unit 3 family inet address 10.255.14.177/32
set protocols bgp group int type internal
set protocols bgp group int local-address 10.255.14.177
set protocols bgp group int neighbor 10.255.71.24
set protocols ospf area 0.0.0.0 interface fe-1/2/0.6
set protocols ospf area 0.0.0.0 interface 10.255.14.177
set routing-options autonomous-system 65000

```

Device R0

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R0:

1. Configure the interfaces.


```

[edit interfaces fe-1/2/0 unit 1]
user@R0# set family inet address 1.1.0.1/30

[edit interfaces fe-1/2/1 unit 2]
user@R0# set family inet address 1.1.1.1/30

[edit interfaces lo0 unit 1]
user@R0# set family inet address 10.255.14.179/32

```
2. Configure EBGP.

```
[edit protocols bgp group ext]
user@R0# set type external
user@R0# set peer-as 65000
user@R0# set neighbor 1.1.0.2
user@R0# set neighbor 1.1.1.2
```

3. Enable multipath BGP between Device R0 and Device R1.

```
[edit protocols bgp group ext]
user@R0# set multipath
```

4. Configure static routes to remote networks.
These routes are not part of the topology. The purpose of these routes is to demonstrate the functionality in this example.

```
[edit routing-options]
user@R0# set static route 1.1.10.10/32 reject
user@R0# set static route 1.1.230.0/23 reject
```

5. Configure routing policies that accept the static routes.

```
[edit policy-options policy-statement agg_route term 1]
user@R0# set from protocol static
user@R0# set from route-filter 1.1.230.0/23 exact
user@R0# set then accept
```

```
[edit policy-options policy-statement test_route term 1]
user@R0# set from protocol static
user@R0# set from route-filter 1.1.10.10/32 exact
user@R0# set then accept
```

6. Export the **agg_route** and **test_route** policies from the routing table into BGP.

```
[edit protocols bgp group ext]
user@R0# set export test_route
user@R0# set export agg_route
```

7. Configure the autonomous system (AS) number.

```
[edit routing-options]
user@R0# set autonomous-system 65500
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R0# show interfaces
fe-1/2/0 {
  unit 1 {
    family inet {
      address 1.1.0.1/30;
    }
  }
}
fe-1/2/1 {
  unit 2 {
    family inet {
```

```

        address 1.1.1.1/30;
    }
}
lo0 {
    unit 1 {
        family inet {
            address 10.255.14.179/32;
        }
    }
}

user@R0# show policy-options
policy-statement agg_route {
    term 1 {
        from {
            protocol static;
            route-filter 1.1.230.0/23 exact;
        }
        then accept;
    }
}
policy-statement test_route {
    term 1 {
        from {
            protocol static;
            route-filter 1.1.10.10/32 exact;
        }
        then accept;
    }
}

user@R0# show protocols
bgp {
    group ext {
        type external;
        export [ test_route agg_route ];
        peer-as 65000;
        multipath;
        neighbor 1.1.0.2;
        neighbor 1.1.1.2;
    }
}

user@R0# show routing-options
static {
    route 1.1.10.10/32 reject;
    route 1.1.230.0/23 reject;
}
autonomous-system 65500;

```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Device R1

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R1:

1. Configure the interfaces.

```
[edit interfaces fe-1/2/0 unit 3]
user@R1# set family inet address 1.1.0.2/30

[edit interfaces fe-1/2/1 unit 4]
user@R1# set family inet address 1.12.0.1/30

[edit interfaces fe-1/2/2 unit 5]
user@R1# set family inet address 1.1.1.2/30

[edit interfaces lo0 unit 2]
user@R1# set family inet address 10.255.71.24/32
```
2. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
user@R1# set interface fe-1/2/1.4
user@R1# set interface 10.255.71.24
```
3. Enable Device R1 to accept the remote next hop.

```
[edit protocols bgp]
user@R1# set accept-remote-nexthop
```
4. Configure IBGP.

```
[edit protocols bgp group int]
user@R1# set type internal
user@R1# set local-address 10.255.71.24
user@R1# set neighbor 10.255.14.177
```
5. Configure EBGP.

```
[edit protocols bgp group ext]
user@R1# set type external
user@R1# set peer-as 65500
user@R1# set neighbor 1.1.0.1
user@R1# set neighbor 1.1.1.1
```
6. Enable multipath BGP between Device R0 and Device R1.

```
[edit protocols bgp group ext]
user@R1# set multipath
```
7. Configure a routing policy that enables a single-hop external BGP peer (Device R1) to accept the remote next-hop 1.1.10.10 for the route to the 1.1.230.0/23 network.

```
[edit policy-options policy-statement agg_route term 1]
user@R1# set from protocol bgp
user@R1# set from route-filter 1.1.230.0/23 exact
```

```

user@R1# set then next-hop 1.1.10.10
user@R1# set then accept

```

8. Import the **agg_route** policy into the routing table on Device R1.

```

[edit protocols bgp group ext]
user@R1# set import agg_route

```

9. Configure the autonomous system (AS) number.

```

[edit routing-options]
user@R1# set autonomous-system 65000

```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@R1# show interfaces
fe-1/2/0 {
  unit 3 {
    family inet {
      address 1.1.0.2/30;
    }
  }
}
fe-1/2/1 {
  unit 4 {
    family inet {
      address 1.12.0.1/30;
    }
  }
}
fe-1/2/2 {
  unit 5 {
    family inet {
      address 1.1.1.2/30;
    }
  }
}
lo0 {
  unit 2 {
    family inet {
      address 10.255.71.24/32;
    }
  }
}
}

user@R1# show policy-options
policy-statement agg_route {
  term 1 {
    from {
      protocol bgp;
      route-filter 1.1.230.0/23 exact;
    }
    then {
      next-hop 1.1.10.10;
    }
  }
}

```

```
        accept;
      }
    }
  }

user@R1# show protocols
bgp {
  accept-remote-nexthop;
  group ext {
    type external;
    import agg_route;
    peer-as 65500;
    multipath;
    neighbor 1.1.0.1;
    neighbor 1.1.1.1;
  }
  group int {
    type internal;
    local-address 10.255.71.24;
    neighbor 10.255.14.177;
  }
}
ospf {
  area 0.0.0.0 {
    interface fe-1/2/1.4;
    interface 10.255.71.24;
  }
}

user@R1# show routing-options
autonomous-system 65000;
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Device R2

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R2:

1. Configure the interfaces.

[edit interfaces fe-1/2/0 unit 6]
user@R2# set family inet address 1.12.0.2/30

[edit interfaces lo0 unit 3]
user@R2# set family inet address 10.255.14.177/32
2. Configure OSPF.

[edit protocols ospf area 0.0.0.0]
user@R2# set interface fe-1/2/0.6
user@R2# set interface 10.255.14.177
3. Configure IBGP.


```
[edit protocols bgp group int]
user@R2# set type internal
user@R2# set local-address 10.255.14.177
user@R2# set neighbor 10.255.71.24
```

4. Configure the autonomous system (AS) number.

```
[edit routing-options]
user@R1# set autonomous-system 65000
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R2# show interfaces
fe-1/2/0 {
  unit 6 {
    family inet {
      address 1.12.0.2/30;
    }
  }
}
lo0 {
  unit 3 {
    family inet {
      address 10.255.14.177/32;
    }
  }
}

user@R2# show protocols
bgp {
  group int {
    type internal;
    local-address 10.255.14.177;
    neighbor 10.255.71.24;
  }
}
ospf {
  area 0.0.0.0 {
    interface fe-1/2/0.6;
    interface 10.255.14.177;
  }
}

user@R2# show routing-options
autonomous-system 65000;
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying That the Multipath Route with the Indirect Next Hop Is in the Routing Table on page 3842](#)
- [Deactivating and Reactivating the accept-remote-nexthop Statement on page 3843](#)

Verifying That the Multipath Route with the Indirect Next Hop Is in the Routing Table

Purpose Verify that Device R1 has a route to the 1.1.230.0/23 network.

Action From operational mode, enter the **show route 1.1.230.0 extensive** command.

```

user@R1> show route 1.1.230.0 extensive
inet.0: 11 destinations, 13 routes (11 active, 0 holddown, 0 hidden)
Restart Complete
1.1.230.0/23 (2 entries, 1 announced)
TSI:
KRT in-kernel 1.1.230.0/23 -> {indirect(262142)}
Page 0 idx 1 Type 1 val 9168f6c
  Nexthop: 1.1.10.10
  Localpref: 100
  AS path: [65000] 65500 I
  Communities:
Path 1.1.230.0 from 1.1.0.1 Vector len 4. Val: 1
  *BGP Preference: 170/-101
    Next hop type: Indirect
    Address: 0x90c44d8
    Next-hop reference count: 4
    Source: 1.1.0.1
    Next hop type: Router, Next hop index: 262143
    Next hop: 1.1.0.1 via fe-1/2/0.3, selected
    Next hop: 1.1.1.1 via fe-1/2/2.5
    Protocol next hop: 1.1.10.10
    Indirect next hop: 91c0000 262142
    State: <Active Ext>
    Local AS: 65000 Peer AS: 65500
    Age: 2:55:31 Metric2: 0
    Task: BGP_65500.1.1.0.1+64631
    Announcement bits (3): 2-KRT 3-BGP_RT_Background 4-Resolve tree
1
    AS path: 65500 I
    Accepted Multipath
    Localpref: 100
    Router ID: 10.255.14.179
    Indirect next hops: 1
      Protocol next hop: 1.1.10.10
      Indirect next hop: 91c0000 262142
      Indirect path forwarding next hops: 2
        Next hop type: Router
        Next hop: 1.1.0.1 via fe-1/2/0.3
        Next hop: 1.1.1.1 via fe-1/2/2.5
      1.1.10.10/32 Originating RIB: inet.0
        Node path count: 1
        Forwarding nexthops: 2
          Nexthop: 1.1.0.1 via fe-1/2/0.3
          Nexthop: 1.1.1.1 via fe-1/2/2.5
  BGP Preference: 170/-101

```

```

Next hop type: Indirect
Address: 0x90c44d8
Next-hop reference count: 4
Source: 1.1.1.1
Next hop type: Router, Next hop index: 262143
Next hop: 1.1.0.1 via fe-1/2/0.3, selected
Next hop: 1.1.1.1 via fe-1/2/2.5
Protocol next hop: 1.1.10.10
Indirect next hop: 91c0000 262142
State: <NotBest Ext>
Inactive reason: Not Best in its group - Update source
Local AS: 65500 Peer AS: 65500
Age: 2:55:27 Metric2: 0
Task: BGP_65500.1.1.1.1+53260
AS path: 65500 I
Accepted
Localpref: 100
Router ID: 10.255.14.179
Indirect next hops: 1
  Protocol next hop: 1.1.10.10
  Indirect next hop: 91c0000 262142
  Indirect path forwarding next hops: 2
    Next hop type: Router
    Next hop: 1.1.0.1 via fe-1/2/0.3
    Next hop: 1.1.1.1 via fe-1/2/2.5
  1.1.10.10/32 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 2
    Nexthop: 1.1.0.1 via fe-1/2/0.3
    Nexthop: 1.1.1.1 via fe-1/2/2.5

```

Meaning The output shows that Device R1 has a route to the 1.1.230.0 network with the multipath feature enabled (**Accepted Multipath**). The output also shows that the route has an indirect next hop of 1.1.10.10.

Deactivating and Reactivating the accept-remote-nexthop Statement

Purpose Make sure that the multipath route with the indirect next hop is removed from the routing table when you deactivate the **accept-remote-nexthop** statement.

Action 1. From configuration mode, enter the **deactivate protocols bgp accept-remote-nexthop** command.

```

user@R1# deactivate protocols bgp accept-remote-nexthop
user@R1# commit

```

2. From operational mode, enter the **show route 1.1.230.0** command.

```

user@R1> show route 1.1.230.0

```

3. From configuration mode, reactivate the statement by entering the **activate protocols bgp accept-remote-nexthop** command.

```

user@R1# activate protocols bgp accept-remote-nexthop
user@R1# commit

```

4. From operational mode, reenter the **show route 1.1.230.0** command.

```

user@R1> show route 1.1.230.0

```

```

inet.0: 11 destinations, 13 routes (11 active, 0 holddown, 0 hidden)

```

Restart Complete

+ = Active Route, - = Last Active, * = Both

```
1.1.230.0/23      *[BGP/170] 03:13:19, localpref 100
                  AS path: 65500 I
                  > to 1.1.0.1 via fe-1/2/0.3
                  to 1.1.1.1 via fe-1/2/2.5
                  [BGP/170] 03:13:15, localpref 100, from 1.1.1.1
                  AS path: 65500 I
                  > to 1.1.0.1 via fe-1/2/0.3
                  to 1.1.1.1 via fe-1/2/2.5
```

Meaning When the **accept-remote-nexthop** statement is deactivated, the multipath route to the 1.1.230.0 network is removed from the routing table .

- Related Documentation**
- *Example: Overriding the Default BGP Routing Policy on PTX Series Packet Transport Routers*
 - *Example: Load Balancing BGP Traffic with Unequal Bandwidth Allocated to the Paths*

Example: Advertising Multiple Paths in BGP

In this example, BGP routers are configured to advertise multiple paths instead of advertising only the active path. Advertising multiple paths in BGP is specified in Internet draft draft-ietf-idr-add-paths-04, *Advertisement of Multiple Paths in BGP*.

- [Requirements on page 3844](#)
- [Overview on page 3844](#)
- [Configuration on page 3846](#)
- [Verification on page 3864](#)

Requirements

This example uses the following hardware and software components:

- Eight BGP-enabled devices.
- Five of the BGP-enabled devices do not necessarily need to be routers. For example, they can be EX Series Ethernet Switches.
- Three of the BGP-enabled devices are configured to send multiple paths or receive multiple paths (or both send and receive multiple paths). These three BGP-enabled devices must be M Series Multiservice Edge Routers, MX Series 3D Universal Edge Routers, or T Series Core Routers.
- The three routers must be running Junos OS Release 11.4 or later.

Overview

The following statements are used for configuring multiple paths to a destination:

```
[edit protocols bgp group group-name family family]
add-path {
```

```

receive;
send {
    path-count number;
    prefix-policy [ policy-names ];
}
}

```

In this example, Router R5, Router R6, and Router R7 redistribute static routes into BGP. Router R1 and Router R4 are route reflectors. Router R2 and Router R3 are clients to Route Reflector R1. Router R8 is a client to Route Reflector R4.

Route reflection is optional when multiple-path advertisement is enabled in BGP.

With the **add-path send path-count 6** configuration, Router R1 is configured to send up to six paths (per destination) to Router R4.

With the **add-path receive** configuration, Router R4 is configured to receive multiple paths from Router R1.

With the **add-path send path-count 6** configuration, Router R4 is configured to send up to six paths to Router R8.

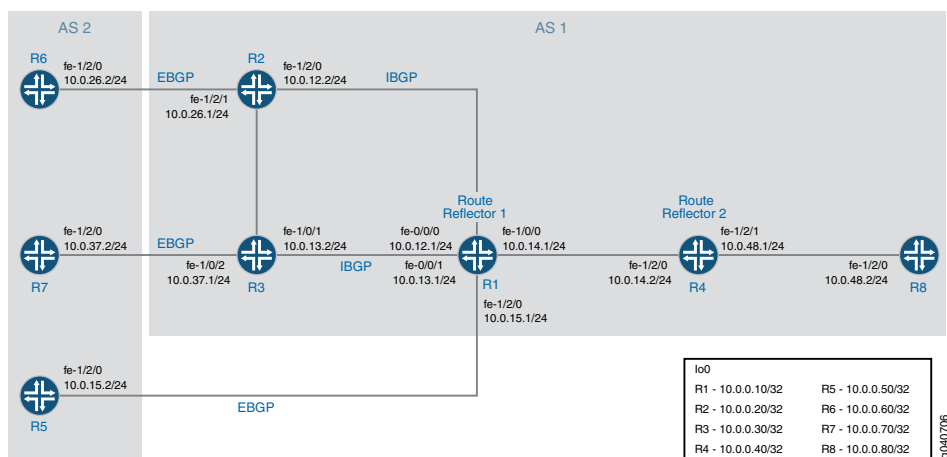
With the **add-path receive** configuration, Router R8 is configured to receive multiple paths from Router R4.

The **add-path send prefix-policy allow_199** policy configuration (along with the corresponding route filter) limits Router R4 to sending multiple paths for only the 199.1.1.1/32 route.

Topology Diagram

Figure 80 shows the topology used in this example.

Figure 80: Advertisement of Multiple Paths in BGP



Configuration

- [Configuring Router R1 on page 3848](#)
- [Configuring Router R2 on page 3851](#)
- [Configuring Router R3 on page 3853](#)
- [Configuring Router R4 on page 3855](#)
- [Configuring Router R5 on page 3858](#)
- [Configuring Router R6 on page 3859](#)
- [Configuring Router R7 on page 3861](#)
- [Configuring Router R8 on page 3862](#)
- [Results on page 3863](#)

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Router R1

```
set interfaces fe-0/0/0 unit 12 family inet address 10.0.12.1/24
set interfaces fe-0/0/1 unit 13 family inet address 10.0.13.1/24
set interfaces fe-1/0/0 unit 14 family inet address 10.0.14.1/24
set interfaces fe-1/2/0 unit 15 family inet address 10.0.15.1/24
set interfaces lo0 unit 10 family inet address 10.0.0.10/32
set protocols bgp group rr type internal
set protocols bgp group rr local-address 10.0.0.10
set protocols bgp group rr cluster 10.0.0.10
set protocols bgp group rr neighbor 10.0.0.20
set protocols bgp group rr neighbor 10.0.0.30
set protocols bgp group e1 type external
set protocols bgp group e1 neighbor 10.0.15.2 local-address 10.0.15.1
set protocols bgp group e1 neighbor 10.0.15.2 peer-as 2
set protocols bgp group rr_rr type internal
set protocols bgp group rr_rr local-address 10.0.0.10
set protocols bgp group rr_rr neighbor 10.0.0.40 family inet unicast add-path send
  path-count 6
set protocols ospf area 0.0.0.0 interface lo0.10 passive
set protocols ospf area 0.0.0.0 interface fe-0/0/0.12
set protocols ospf area 0.0.0.0 interface fe-0/0/1.13
set protocols ospf area 0.0.0.0 interface fe-1/0/0.14
set protocols ospf area 0.0.0.0 interface fe-1/2/0.15
set routing-options router-id 10.0.0.10
set routing-options autonomous-system 1
```

Router R2

```
set interfaces fe-1/2/0 unit 21 family inet address 10.0.12.2/24
set interfaces fe-1/2/1 unit 26 family inet address 10.0.26.1/24
set interfaces lo0 unit 20 family inet address 10.0.0.20/32
set protocols bgp group rr type internal
set protocols bgp group rr local-address 10.0.0.20
set protocols bgp group rr neighbor 10.0.0.10 export set_nh_self
set protocols bgp group e1 type external
set protocols bgp group e1 neighbor 10.0.26.2 peer-as 2
set protocols ospf area 0.0.0.0 interface lo0.20 passive
```

```

set protocols ospf area 0.0.0.0 interface fe-1/2/0.21
set protocols ospf area 0.0.0.0 interface fe-1/2/1.28
set policy-options policy-statement set_nh_self then next-hop self
set routing-options autonomous-system 1

```

Router R3

```

set interfaces fe-1/0/1 unit 31 family inet address 10.0.13.2/24
set interfaces fe-1/0/2 unit 37 family inet address 10.0.37.1/24
set interfaces lo0 unit 30 family inet address 10.0.0.30/32
set protocols bgp group rr type internal
set protocols bgp group rr local-address 10.0.0.30
set protocols bgp group rr neighbor 10.0.0.10 export set_nh_self
set protocols bgp group e1 type external
set protocols bgp group e1 neighbor 10.0.37.2 peer-as 2
set protocols ospf area 0.0.0.0 interface lo0.30 passive
set protocols ospf area 0.0.0.0 interface fe-1/0/1.31
set protocols ospf area 0.0.0.0 interface fe-1/0/2.37
set policy-options policy-statement set_nh_self then next-hop self
set routing-options autonomous-system 1

```

Router R4

```

set interfaces fe-1/2/0 unit 41 family inet address 10.0.14.2/24
set interfaces fe-1/2/1 unit 48 family inet address 10.0.48.1/24
set interfaces lo0 unit 40 family inet address 10.0.0.40/32
set protocols bgp group rr type internal
set protocols bgp group rr local-address 10.0.0.40
set protocols bgp group rr family inet unicast add-path receive
set protocols bgp group rr neighbor 10.0.0.10
set protocols bgp group rr_client type internal
set protocols bgp group rr_client local-address 10.0.0.40
set protocols bgp group rr_client cluster 10.0.0.40
set protocols bgp group rr_client neighbor 10.0.0.80 family inet unicast add-path send
    path-count 6
set protocols bgp group rr_client neighbor 10.0.0.80 family inet unicast add-path send
    prefix-policy allow_199
set protocols ospf area 0.0.0.0 interface fe-1/2/0.41
set protocols ospf area 0.0.0.0 interface lo0.40 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/1.48
set policy-options policy-statement allow_199 from route-filter 199.1.1.1/32 exact
set policy-options policy-statement allow_199 term match_199 from prefix-list match_199
set policy-options policy-statement allow_199 then add-path send-count 20
set policy-options policy-statement allow_199 then accept
set routing-options autonomous-system 1

```

Router R5

```

set interfaces fe-1/2/0 unit 51 family inet address 10.0.15.2/24
set interfaces lo0 unit 50 family inet address 10.0.0.50/32
set protocols bgp group e1 type external
set protocols bgp group e1 neighbor 10.0.15.1 export s2b
set protocols bgp group e1 neighbor 10.0.15.1 peer-as 1
set policy-options policy-statement s2b from protocol static
set policy-options policy-statement s2b from protocol direct
set policy-options policy-statement s2b then as-path-expand 2
set policy-options policy-statement s2b then accept
set routing-options autonomous-system 2
set routing-options static route 199.1.1.1/32 reject
set routing-options static route 198.1.1.1/32 reject

```

Router R6 `set interfaces fe-1/2/0 unit 62 family inet address 10.0.26.2/24`
`set interfaces lo0 unit 60 family inet address 10.0.0.60/32`
`set protocols bgp group e1 type external`
`set protocols bgp group e1 neighbor 10.0.26.1 export s2b`
`set protocols bgp group e1 neighbor 10.0.26.1 peer-as 1`
`set policy-options policy-statement s2b from protocol static`
`set policy-options policy-statement s2b from protocol direct`
`set policy-options policy-statement s2b then accept`
`set routing-options autonomous-system 2`
`set routing-options static route 199.1.1.1/32 reject`
`set routing-options static route 198.1.1.1/32 reject`

Router R7 `set interfaces fe-1/2/0 unit 73 family inet address 10.0.37.2/24`
`set interfaces lo0 unit 70 family inet address 10.0.0.70/32`
`set protocols bgp group e1 type external`
`set protocols bgp group e1 neighbor 10.0.37.1 export s2b`
`set protocols bgp group e1 neighbor 10.0.37.1 peer-as 1`
`set policy-options policy-statement s2b from protocol static`
`set policy-options policy-statement s2b from protocol direct`
`set policy-options policy-statement s2b then accept`
`set routing-options autonomous-system 2`
`set routing-options static route 199.1.1.1/32 reject`

Router R8 `set interfaces fe-1/2/0 unit 84 family inet address 10.0.48.2/24`
`set interfaces lo0 unit 80 family inet address 10.0.0.80/32`
`set protocols bgp group rr type internal`
`set protocols bgp group rr local-address 10.0.0.80`
`set protocols bgp group rr neighbor 10.0.0.40 family inet unicast add-path receive`
`set protocols ospf area 0.0.0.0 interface lo0.80 passive`
`set protocols ospf area 0.0.0.0 interface fe-1/2/0.84`
`set routing-options autonomous-system 1`

Configuring Router R1

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router R1:

1. Configure the interfaces to Router R2, Router R3, Router R4, and Router R5, and configure the loopback (lo0) interface.

[edit interfaces]

user@R1# `set fe-0/0/0 unit 12 family inet address 10.0.12.1/24`

user@R1# `set fe-0/0/1 unit 13 family inet address 10.0.13.1/24`

user@R1# `set fe-1/0/0 unit 14 family inet address 10.0.14.1/24`

user@R1# `set fe-1/2/0 unit 15 family inet address 10.0.15.1/24`

user@R1# `set lo0 unit 10 family inet address 10.0.0.10/32`

2. Configure BGP on the interfaces, and configure IBGP route reflection.

```
[edit protocols bgp]
user@R1# set group rr type internal
user@R1# set group rr local-address 10.0.0.10
user@R1# set group rr cluster 10.0.0.10
user@R1# set group rr neighbor 10.0.0.20
user@R1# set group rr neighbor 10.0.0.30

user@R1# set group rr_rr type internal
user@R1# set group rr_rr local-address 10.0.0.10

user@R1# set group e1 type external
user@R1# set group e1 neighbor 10.0.15.2 local-address 10.0.15.1
user@R1# set group e1 neighbor 10.0.15.2 peer-as 2
```

3. Configure Router R1 to send up to six paths to its neighbor, Router R4.

The destination of the paths can be any destination that Router R1 can reach through multiple paths.

```
[edit protocols bgp]
user@R1# set group rr_rr neighbor 10.0.0.40 family inet unicast add-path send
path-count 6
```

4. Configure OSPF on the interfaces.

```
[edit protocols ospf]
user@R1# set area 0.0.0.0 interface lo0.10 passive
user@R1# set area 0.0.0.0 interface fe-0/0/0.12
user@R1# set area 0.0.0.0 interface fe-0/0/1.13
user@R1# set area 0.0.0.0 interface fe-1/0/0.14
user@R1# set area 0.0.0.0 interface fe-1/2/0.15
```

5. Configure the router ID and the autonomous system number.

```
[edit routing-options]
user@R1# set router-id 10.0.0.10
user@R1# set autonomous-system 1
```

6. If you are done configuring the device, commit the configuration.

```
user@R1# commit
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show interfaces
fe-0/0/0 {
  unit 12 {
    family inet {
      address 10.0.12.1/24;
    }
  }
}
```

```
fe-0/0/1 {
  unit 13 {
    family inet {
      address 10.0.13.1/24;
    }
  }
}
fe-1/0/0 {
  unit 14 {
    family inet {
      address 10.0.14.1/24;
    }
  }
}
fe-1/2/0 {
  unit 15 {
    family inet {
      address 10.0.15.1/24;
    }
  }
}
lo0 {
  unit 10 {
    family inet {
      address 10.0.0.10/32;
    }
  }
}

user@R1# show protocols
bgp {
  group rr {
    type internal;
    local-address 10.0.0.10;
    cluster 10.0.0.10;
    neighbor 10.0.0.20;
    neighbor 10.0.0.30;
  }
  group e1 {
    type external;
    neighbor 10.0.15.2 {
      local-address 10.0.15.1;
      peer-as 2;
    }
  }
  group rr_rr {
    type internal;
    local-address 10.0.0.10;
    neighbor 10.0.0.40 {
      family inet {
        unicast {
          add-path {
            send {
              path-count 6;
            }
          }
        }
      }
    }
  }
}
```

```

    }
  }
}
}
ospf {
  area 0.0.0.0 {
    interface lo0.10 {
      passive;
    }
    interface fe-0/0/0.12;
    interface fe-0/0/1.13;
    interface fe-1/0/0.14;
    interface fe-1/2/0.15;
  }
}

user@R1# show routing-options
router-id 10.0.0.10;
autonomous-system 1;

```

Configuring Router R2

Step-by-Step Procedure

To configure Router R2:

1. Configure the loopback (lo0) interface and the interfaces to Router R6 and Router R1.


```

[edit interfaces]
user@R2# set fe-1/2/0 unit 21 family inet address 10.0.12.2/24

user@R2# set fe-1/2/1 unit 26 family inet address 10.0.26.1/24

user@R2# set lo0 unit 20 family inet address 10.0.0.20/32

```
2. Configure BGP and OSPF on Router R2's interfaces.


```

[edit protocols]
user@R2# set bgp group rr type internal
user@R2# set bgp group rr local-address 10.0.0.20

user@R2# set bgp group e1 type external
user@R2# set bgp group e1 neighbor 10.0.26.2 peer-as 2

user@R2# set ospf area 0.0.0.0 interface lo0.20 passive
user@R2# set ospf area 0.0.0.0 interface fe-1/2/0.21
user@R2# set ospf area 0.0.0.0 interface fe-1/2/1.28

```
3. For routes sent from Router R2 to Router R1, advertise Router R2 as the next hop, because Router R1 does not have a route to Router R6's address on the 10.0.26.0/24 network.


```

[edit]
user@R2# set policy-options policy-statement set_nh_self then next-hop self

```

```
user@R2# set protocols bgp group rr neighbor 10.0.0.10 export set_nh_self
```

4. Configure the autonomous system number.

```
[edit]  
user@R2# set routing-options autonomous-system 1
```

5. If you are done configuring the device, commit the configuration.

```
user@R2# commit
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R2# show interfaces  
fe-1/2/0 {  
  unit 21 {  
    family inet {  
      address 10.0.12.2/24;  
    }  
  }  
}  
fe-1/2/1 {  
  unit 26 {  
    family inet {  
      address 10.0.26.1/24;  
    }  
  }  
}  
lo0 {  
  unit 20 {  
    family inet {  
      address 10.0.0.20/32;  
    }  
  }  
}  
  
user@R2# show protocols  
bgp {  
  group rr {  
    type internal;  
    local-address 10.0.0.20;  
    neighbor 10.0.0.10 {  
      export set_nh_self;  
    }  
  }  
  group e1 {  
    type external;  
    neighbor 10.0.26.2 {  
      peer-as 2;  
    }  
  }  
}  
ospf {
```

```

area 0.0.0.0 {
  interface lo0.20 {
    passive;
  }
  interface fe-1/2/0.21;
  interface fe-1/2/1.28;
}
}

user@R2# show policy-options
policy-statement set_nh_self {
  then {
    next-hop self;
  }
}

user@R2# show routing-options
autonomous-system 1;

```

Configuring Router R3

Step-by-Step Procedure

To configure Router R3:

1. Configure the loopback (lo0) interface and the interfaces to Router R7 and Router R1.

```

[edit interfaces]
user@R3# set fe-1/0/1 unit 31 family inet address 10.0.13.2/24

user@R3# set fe-1/0/2 unit 37 family inet address 10.0.37.1/24

user@R3# set lo0 unit 30 family inet address 10.0.0.30/32

```

2. Configure BGP and OSPF on Router R3's interfaces.

```

[edit protocols]
user@R3# set bgp group rr type internal
user@R3# set bgp group rr local-address 10.0.0.30

user@R3# set bgp group e1 type external
user@R3# set bgp group e1 neighbor 10.0.37.2 peer-as 2

user@R3# set ospf area 0.0.0.0 interface lo0.30 passive
user@R3# set ospf area 0.0.0.0 interface fe-1/0/1.31
user@R3# set ospf area 0.0.0.0 interface fe-1/0/2.37

```

3. For routes sent from Router R3 to Router R1, advertise Router R3 as the next hop, because Router R1 does not have a route to Router R7's address on the 10.0.37.0/24 network.

```

[edit]
user@R3# set policy-options policy-statement set_nh_self then next-hop self

user@R3# set protocols bgp group rr neighbor 10.0.0.10 export set_nh_self

```

4. Configure the autonomous system number.

```
[edit]
user@R3# set routing-options autonomous-system 1
```

5. If you are done configuring the device, commit the configuration.

```
user@R3# commit
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R3# show interfaces
fe-1/0/1 {
  unit 31 {
    family inet {
      address 10.0.13.2/24;
    }
  }
}
fe-1/0/2 {
  unit 37 {
    family inet {
      address 10.0.37.1/24;
    }
  }
}
lo0 {
  unit 30 {
    family inet {
      address 10.0.0.30/32;
    }
  }
}

user@R3# show protocols
bgp {
  group rr {
    type internal;
    local-address 10.0.0.30;
    neighbor 10.0.0.10 {
      export set_nh_self;
    }
  }
  group e1 {
    type external;
    neighbor 10.0.37.2 {
      peer-as 2;
    }
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.30 {
```

```

        passive;
    }
    interface fe-1/0/1.31;
    interface fe-1/0/2.37;
}
}
user@R3# show policy-options
policy-statement set_nh_self {
    then {
        next-hop self;
    }
}

user@R3# show routing-options
autonomous-system 1;

```

Configuring Router R4

Step-by-Step Procedure

To configure Router R4:

1. Configure the interfaces to Router R1 and Router R8, and configure the loopback (lo0) interface.

```
[edit interfaces]
```

```
user@R4# set fe-1/2/0 unit 41 family inet address 10.0.14.2/24
```

```
user@R4# set fe-1/2/1 unit 48 family inet address 10.0.48.1/24
```

```
user@R4# set lo0 unit 40 family inet address 10.0.0.40/32
```

2. Configure BGP on the interfaces, and configure IBGP route reflection.

```
[edit protocols bgp]
```

```
user@R4# set group rr type internal
```

```
user@R4# set group rr local-address 10.0.0.40
```

```
user@R4# set group rr neighbor 10.0.0.10
```

```
user@R4# set group rr_client type internal
```

```
user@R4# set group rr_client local-address 10.0.0.40
```

```
user@R4# set group rr_client cluster 10.0.0.40
```

3. Configure Router R4 to send up to six paths to its neighbor, Router R8.

The destination of the paths can be any destination that Router R4 can reach through multiple paths.

```
[edit protocols bgp]
```

```
user@R4# set group rr_client neighbor 10.0.0.80 family inet unicast add-path send
path-count 6
```

4. Configure Router R4 to receive multiple paths from its neighbor, Router R1.

The destination of the paths can be any destination that Router R1 can reach through multiple paths.

```
[edit protocols bgp group rr family inet unicast]
```

```
user@R4# set add-path receive
```

5. Configure OSPF on the interfaces.

```
[edit protocols ospf area 0.0.0.0]
user@R4# set interface fe-1/2/0.41
user@R4# set interface lo0.40 passive
user@R4# set interface fe-1/2/1.48
```

6. Configure a policy that allows Router R4 to send Router R8 multiple paths to the 199.1.1.1/32 route.

- Router R4 receives multiple paths for the 198.1.1.1/32 route and the 199.1.1.1/32 route. However, because of this policy, Router R4 only sends multiple paths for the 199.1.1.1/32 route.

```
[edit protocols bgp group rr_client neighbor 10.0.0.80 family inet unicast]
user@R4# set add-path send prefix-policy allow_199
[edit policy-options policy-statement allow_199]
user@R4# set from route-filter 199.1.1.1/32 exact
user@R4# set then accept
```

- Router R4 can also be configured to send up-to 20 BGP **add-path** routes for a subset of *add-path advertised prefixes*.

```
[edit policy-options policy-statement allow_199]
user@R4# set term match_199 from prefix-list match_199
user@R4# set then add-path send-count 20
```

7. Configure the autonomous system number.

```
[edit routing-options]
user@R4# set autonomous-system 1
```

8. If you are done configuring the device, commit the configuration.

```
user@R4# commit
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R4# show interfaces
fe-1/2/0 {
  unit 41 {
    family inet {
      address 10.0.14.2/24;
    }
  }
}
fe-1/2/1 {
  unit 48 {
    family inet {
      address 10.0.48.1/24;
    }
  }
}
lo0 {
```



```
unit 40 {
  family inet {
    address 10.0.0.40/32;
  }
}

user@R4# show protocols
bgp {
  group rr {
    type internal;
    local-address 10.0.0.40;
    family inet {
      unicast {
        add-path {
          receive;
        }
      }
    }
    neighbor 10.0.0.10;
  }
  group rr_client {
    type internal;
    local-address 10.0.0.40;
    cluster 10.0.0.40;
    neighbor 10.0.0.80 {
      family inet {
        unicast {
          add-path {
            send {
              path-count 6;
              prefix-policy allow_199;
            }
          }
        }
      }
    }
  }
}

ospf {
  area 0.0.0.0 {
    interface lo0.40 {
      passive;
    }
    interface fe-1/2/0.41;
    interface fe-1/2/1.48;
  }
}

user@R4# show policy-options
policy-statement allow_199 {
  from {
    route-filter 199.1.1.1/32 exact;
  }
  from term match_199 {
    prefix-list match_199;
  }
}
```

```
    then add-path send-count 20;
  then accept;
}
```

```
user@R4# show routing-options
autonomous-system 1;
```

Configuring Router R5

Step-by-Step Procedure

To configure Router R5:

1. Configure the loopback (lo0) interface and the interface to Router R1.

```
[edit interfaces]
user@R5# set fe-1/2/0 unit 51 family inet address 10.0.15.2/24
```

```
user@R5# set lo0 unit 50 family inet address 10.0.0.50/32
```

2. Configure BGP on Router R5's interface.

```
[edit protocols bgp group e1]
user@R5# set type external
user@R5# set neighbor 10.0.15.1 peer-as 1
```

3. Create static routes for redistribution into BGP.

```
[edit routing-options]
user@R5# set static route 199.1.1.1/32 reject
user@R5# set static route 198.1.1.1/32 reject
```

4. Redistribute static and direct routes into BGP.

```
[edit protocols bgp group e1 neighbor 10.0.15.1]
user@R5# set export s2b
```

```
[edit policy-options policy-statement s2b]
user@R5# set from protocol static
user@R5# set from protocol direct
user@R5# set then as-path-expand 2
user@R5# set then accept
```

5. Configure the autonomous system number.

```
[edit routing-options]
user@R5# set autonomous-system 2
```

6. If you are done configuring the device, commit the configuration.

```
user@R5# commit
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R5# show interfaces
fe-1/2/0 {
```

```

    unit 51 {
      family inet {
        address 10.0.15.2/24;
      }
    }
  }
  lo0 {
    unit 50 {
      family inet {
        address 10.0.0.50/32;
      }
    }
  }
}

user@R5# show protocols
bgp {
  group e1 {
    type external;
    neighbor 10.0.15.1 {
      export s2b;
      peer-as 1;
    }
  }
}

user@R5# show policy-options
policy-statement s2b {
  from protocol [ static direct ];
  then {
    as-path-expand 2;
    accept;
  }
}

user@R5# show routing-options
static {
  route 198.1.1.1/32 reject;
  route 199.1.1.1/32 reject;
}
autonomous-system 2;

```

Configuring Router R6

Step-by-Step Procedure

To configure Router R6:

1. Configure the loopback (lo0) interface and the interface to Router R2.

[edit interfaces]

```
user@R6# set fe-1/2/0 unit 62 family inet address 10.0.26.2/24
```

```
user@R6# set lo0 unit 60 family inet address 10.0.0.60/32
```

2. Configure BGP on Router R6's interface.

[edit protocols]

```
user@R6# set bgp group e1 type external
```

```
user@R6# set bgp group e1 neighbor 10.0.26.1 peer-as 1
```

3. Create static routes for redistribution into BGP.

```
[edit]
user@R6# set routing-options static route 199.1.1.1/32 reject
user@R6# set routing-options static route 198.1.1.1/32 reject
```

4. Redistribute static and direct routes from Router R6's routing table into BGP.

```
[edit protocols bgp group e1 neighbor 10.0.26.1]
user@R6# set export s2b
```

```
[edit policy-options policy-statement s2b]
user@R6# set from protocol static
user@R6# set from protocol direct
user@R6# set then accept
```

5. Configure the autonomous system number.

```
[edit routing-options]
user@R6# set autonomous-system 2
```

6. If you are done configuring the device, commit the configuration.

```
user@R6# commit
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R6# show interfaces
fe-1/2/0 {
  unit 62 {
    family inet {
      address 10.0.26.2/24;
    }
  }
}
lo0 {
  unit 60 {
    family inet {
      address 10.0.0.60/32;
    }
  }
}

user@R6# show protocols
bgp {
  group e1 {
    type external;
    neighbor 10.0.26.1 {
      export s2b;
      peer-as 1;
    }
  }
}
```

```

user@R6# show policy-options
policy-statement s2b {
  from protocol [ static direct ];
  then accept;
}

user@R6# show routing-options
static {
  route 198.1.1.1/32 reject;
  route 199.1.1.1/32 reject;
}
autonomous-system 2;

```

Configuring Router R7

Step-by-Step Procedure

To configure Router R7:

1. Configure the loopback (lo0) interface and the interface to Router R3.


```

[edit interfaces]
user@R7# set fe-1/2/0 unit 73 family inet address 10.0.37.2/24

user@R7# set lo0 unit 70 family inet address 10.0.0.70/32

```
2. Configure BGP on Router R7's interface.


```

[edit protocols bgp group e1]
user@R7# set type external
user@R7# set neighbor 10.0.37.1 peer-as 1

```
3. Create a static route for redistribution into BGP.


```

[edit]
user@R7# set routing-options static route 199.1.1.1/32 reject

```
4. Redistribute static and direct routes from Router R7's routing table into BGP.


```

[edit protocols bgp group e1 neighbor 10.0.37.1]
user@R7# set export s2b

[edit policy-options policy-statement s2b]
user@R7# set from protocol static
user@R7# set from protocol direct
user@R7# set then accept

```
5. Configure the autonomous system number.


```

[edit routing-options]
user@R7# set autonomous-system 2

```
6. If you are done configuring the device, commit the configuration.


```

user@R7# commit

```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output

does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R7# show interfaces
fe-1/2/0 {
  unit 73 {
    family inet {
      address 10.0.37.2/24;
    }
  }
}
lo0 {
  unit 70 {
    family inet {
      address 10.0.0.70/32;
    }
  }
}

user@R7# show protocols
bgp {
  group e1 {
    type external;
    neighbor 10.0.37.1 {
      export s2b;
      peer-as 1;
    }
  }
}

user@R7# show policy-options
policy-statement s2b {
  from protocol [ static direct ];
  then accept;
}

user@R7# show routing-options
static {
  route 199.1.1.1/32 reject;
}
autonomous-system 2;
```

Configuring Router R8

Step-by-Step Procedure

To configure Router R8:

1. Configure the loopback (lo0) interface and the interface to Router R4.

[edit interfaces]

```
user@R8# set fe-1/2/0 unit 84 family inet address 10.0.48.2/24
```

```
user@R8# set lo0 unit 80 family inet address 10.0.0.80/32
```

2. Configure BGP and OSPF on Router R8's interface.

[edit protocols]

```
user@R8# set bgp group rr type internal
```

```
user@R8# set bgp group rr local-address 10.0.0.80
```

```
user@R8# set ospf area 0.0.0.0 interface lo0.80 passive
```

```
user@R8# set ospf area 0.0.0.0 interface fe-1/2/0.84
```

3. Configure Router R8 to receive multiple paths from its neighbor, Router R4.

The destination of the paths can be any destination that Router R4 can reach through multiple paths.

```
[edit protocols]
```

```
user@R8# set bgp group rr neighbor 10.0.0.40 family inet unicast add-path receive
```

4. Configure the autonomous system number.

```
[edit]
```

```
user@R8# set routing-options autonomous-system 1
```

5. If you are done configuring the device, commit the configuration.

```
user@R8# commit
```

Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R8# show interfaces
```

```
fe-1/2/0 {
  unit 84 {
    family inet {
      address 10.0.48.2/24;
    }
  }
}
lo0 {
  unit 80 {
    family inet {
      address 10.0.0.80/32;
    }
  }
}
```

```
user@R8# show protocols
```

```
bgp {
  group rr {
    type internal;
    local-address 10.0.0.80;
    neighbor 10.0.0.40 {
      family inet {
        unicast {
          add-path {
            receive;
          }
        }
      }
    }
  }
}
```

```
    }  
  }  
}  
ospf {  
  area 0.0.0.0 {  
    interface lo0.80 {  
      passive;  
    }  
    interface fe-1/2/0.84;  
  }  
}
```

user@R8# show routing-options
autonomous-system 1;

Verification

Confirm that the configuration is working properly.

- [Verifying That the BGP Peers Have the Ability to Send and Receive Multiple Paths on page 3864](#)
- [Verifying That Router R1 Is Advertising Multiple Paths on page 3865](#)
- [Verifying That Router R4 Is Receiving and Advertising Multiple Paths on page 3865](#)
- [Verifying That Router R8 Is Receiving Multiple Paths on page 3866](#)
- [Checking the Path ID on page 3867](#)

Verifying That the BGP Peers Have the Ability to Send and Receive Multiple Paths

Purpose Make sure that one or both of the following strings appear in the output of the **show bgp neighbor** command:

- NLRI's for which peer can receive multiple paths: inet-unicast
- NLRI's for which peer can send multiple paths: inet-unicast

Action user@R1> show bgp neighbor 10.0.0.40
Peer: 10.0.0.40+179 AS 1 Local: 10.0.0.10+65237 AS 1
Type: Internal State: Established Flags: <Sync>
... NLRI's for which peer can receive multiple paths: inet-unicast
...
user@R4> show bgp neighbor 10.0.0.10
Peer: 10.0.0.10+65237 AS 1 Local: 10.0.0.40+179 AS 1
Type: Internal State: Established Flags: <Sync>
...
NLRI's for which peer can send multiple paths: inet-unicast
...
user@R4> show bgp neighbor 10.0.0.80
Peer: 10.0.0.80+55416 AS 1 Local: 10.0.0.40+179 AS 1
Type: Internal State: Established (route reflector client)Flags: <Sync>
...
NLRI's for which peer can receive multiple paths: inet-unicast
...
user@R8> show bgp neighbor 10.0.0.40
Peer: 10.0.0.40+179 AS 1 Local: 10.0.0.80+55416 AS 1
Type: Internal State: Established Flags: <Sync>
...
NLRI's for which peer can send multiple paths: inet-unicast
...

Verifying That Router R1 Is Advertising Multiple Paths

Purpose Make sure that multiple paths to the 198.1.1.1/32 destination and multiple paths to the 199.1.1.1/32 destination are advertised to Router R4.

Action user@R1> show route advertising-protocol bgp 10.0.0.40
inet.0: 21 destinations, 25 routes (21 active, 0 holddown, 0 hidden)

Prefix	Nexthop	MED	Lc1pref	AS path
* 10.0.0.50/32	10.0.15.2		100	2 2 I
* 10.0.0.60/32	10.0.0.20		100	2 I
* 10.0.0.70/32	10.0.0.30		100	2 I
* 198.1.1.1/32	10.0.0.20		100	2 I
	10.0.15.2		100	2 2 I
* 199.1.1.1/32	10.0.0.20		100	2 I
	10.0.0.30		100	2 I
	10.0.15.2		100	2 2 I
* 200.1.1.0/30	10.0.0.20		100	2 I

Meaning When you see one prefix and more than one next hop, it means that multiple paths are advertised to Router R4.

Verifying That Router R4 Is Receiving and Advertising Multiple Paths

Purpose Make sure that multiple paths to the 199.1.1.1/32 destination are received from Router R1 and advertised to Router R8. Make sure that multiple paths to the 198.1.1.1/32 destination are received from Router R1, but only one path to this destination is advertised to Router R8.

Action user@R4> **show route receive-protocol bgp 10.0.0.10**
inet.0: 19 destinations, 22 routes (19 active, 0 holddown, 0 hidden)
Prefix Nexthop MED Lclpref AS path
* 10.0.0.50/32 10.0.15.2 100 2 2 I
* 10.0.0.60/32 10.0.0.20 100 2 I
* 10.0.0.70/32 10.0.0.30 100 2 I
* 198.1.1.1/32 10.0.0.20 100 2 I
10.0.15.2 100 2 2 I
* 199.1.1.1/32 10.0.0.20 100 2 I
10.0.0.30 100 2 I
10.0.15.2 100 2 2 I
* 200.1.1.0/30 10.0.0.20 100 2 I

user@R4> **show route advertising-protocol bgp 10.0.0.80**
inet.0: 19 destinations, 22 routes (19 active, 0 holddown, 0 hidden)
Prefix Nexthop MED Lclpref AS path
* 10.0.0.50/32 10.0.15.2 100 2 2 I
* 10.0.0.60/32 10.0.0.20 100 2 I
* 10.0.0.70/32 10.0.0.30 100 2 I
* 198.1.1.1/32 10.0.0.20 100 2 I
* 199.1.1.1/32 10.0.0.20 100 2 I
10.0.0.30 100 2 I
10.0.15.2 100 2 2 I
* 200.1.1.0/30 10.0.0.20 100 2 I

Meaning The **show route receive-protocol** command shows that Router R4 receives two paths to the 198.1.1.1/32 destination and three paths to the 199.1.1.1/32 destination. The **show route advertising-protocol** command shows that Router R4 advertises only one path to the 198.1.1.1/32 destination and advertises all three paths to the 199.1.1.1/32 destination.

Because of the prefix policy that is applied to Router R4, Router R4 does not advertise multiple paths to the 198.1.1.1/32 destination. Router R4 advertises only one path to the 198.1.1.1/32 destination even though it receives multiple paths to this destination.

Verifying That Router R8 Is Receiving Multiple Paths

Purpose Make sure that Router R8 receives multiple paths to the 199.1.1.1/32 destination through Router R4. Make sure that Router R8 receives only one path to the 198.1.1.1/32 destination through Router R4.

Action user@R8> show route receive-protocol bgp 10.0.0.40

inet.0: 18 destinations, 20 routes (18 active, 0 holddown, 0 hidden)

Prefix	Nexthop	MED	Lc1pref	AS path
* 10.0.0.50/32	10.0.15.2		100	2 2 I
* 10.0.0.60/32	10.0.0.20		100	2 I
* 10.0.0.70/32	10.0.0.30		100	2 I
* 198.1.1.1/32	10.0.0.20		100	2 I
* 199.1.1.1/32	10.0.0.20		100	2 I
	10.0.0.30		100	2 I
	10.0.15.2		100	2 2 I
* 200.1.1.0/30	10.0.0.20		100	2 I

Checking the Path ID

Purpose On the downstream devices, Router R4 and Router R8, verify that a path ID uniquely identifies the path. Look for the **Addpath Path ID:** string.

Action user@R4> show route 199.1.1.1/32 detail

```
inet.0: 18 destinations, 20 routes (18 active, 0 holddown, 0 hidden)
199.1.1.1/32 (3 entries, 3 announced)
  *BGP    Preference: 170/-101
          Next hop type: Indirect
          Next-hop reference count: 9
          Source: 10.0.0.10
          Next hop type: Router, Next hop index: 676
          Next hop: 10.0.14.1 via lt-1/2/0.41, selected
          Protocol next hop: 10.0.0.20
          Indirect next hop: 92041c8 262146
          State: <Active Int Ext>
          Local AS:      1 Peer AS:      1
          Age: 1:44:37    Metric2: 2
          Task: BGP_1.10.0.0.10+65237
          Announcement bits (3): 2-KRT 3-BGP RT Background 4-Resolve tree

1
  AS path: 2 I (Originator) Cluster list: 10.0.0.10
  AS path: Originator ID: 10.0.0.20
  Accepted
  Localpref: 100
  Router ID: 10.0.0.10
  Addpath Path ID: 1
  BGP    Preference: 170/-101
          Next hop type: Indirect
          Next-hop reference count: 4
          Source: 10.0.0.10
          Next hop type: Router, Next hop index: 676
          Next hop: 10.0.14.1 via lt-1/2/0.41, selected
          Protocol next hop: 10.0.0.30
          Indirect next hop: 92042ac 262151
          State: <NotBest Int Ext>
          Inactive reason: Not Best in its group - Router ID
          Local AS:      1 Peer AS:      1
          Age: 1:44:37    Metric2: 2
          Task: BGP_1.10.0.0.10+65237
          Announcement bits (1): 3-BGP RT Background
          AS path: 2 I (Originator) Cluster list: 10.0.0.10
          AS path: Originator ID: 10.0.0.30
          Accepted
          Localpref: 100
          Router ID: 10.0.0.10
          Addpath Path ID: 2
  BGP    Preference: 170/-101
          Next hop type: Indirect
          Next-hop reference count: 4
          Source: 10.0.0.10
          Next hop type: Router, Next hop index: 676
          Next hop: 10.0.14.1 via lt-1/2/0.41, selected
          Protocol next hop: 10.0.15.2
          Indirect next hop: 92040e4 262150
          State: <Int Ext>
          Inactive reason: AS path
          Local AS:      1 Peer AS:      1
          Age: 1:44:37    Metric2: 2
          Task: BGP_1.10.0.0.10+65237
          Announcement bits (1): 3-BGP RT Background
          AS path: 2 2 I
          Accepted
```

```

Localpref: 100
Router ID: 10.0.0.10
Addpath Path ID: 3

```

```
user@R8> show route 199.1.1.1/32 detail
```

```

inet.0: 17 destinations, 19 routes (17 active, 0 holddown, 0 hidden)
199.1.1.1/32 (3 entries, 1 announced)
*BGP   Preference: 170/-101
      Next hop type: Indirect
      Next-hop reference count: 9
      Source: 10.0.0.40
      Next hop type: Router, Next hop index: 1045
      Next hop: 10.0.48.1 via lt-1/2/0.84, selected
      Protocol next hop: 10.0.0.20
      Indirect next hop: 91fc0e4 262148
      State: <Active Int Ext>
      Local AS:      1 Peer AS:      1
      Age: 1:56:51   Metric2: 3
      Task: BGP_1.10.0.0.40+179
      Announcement bits (2): 2-KRT 4-Resolve tree 1
      AS path: 2 I (Originator) Cluster list: 10.0.0.40 10.0.0.10
      AS path: Originator ID: 10.0.0.20
      Accepted
      Localpref: 100
      Router ID: 10.0.0.40
      Addpath Path ID: 1
BGP   Preference: 170/-101
      Next hop type: Indirect
      Next-hop reference count: 4
      Source: 10.0.0.40
      Next hop type: Router, Next hop index: 1045
      Next hop: 10.0.48.1 via lt-1/2/0.84, selected
      Protocol next hop: 10.0.0.30
      Indirect next hop: 91fc1c8 262152
      State: <NotBest Int Ext>
      Inactive reason: Not Best in its group - Router ID
      Local AS:      1 Peer AS:      1
      Age: 1:56:51   Metric2: 3
      Task: BGP_1.10.0.0.40+179
      AS path: 2 I (Originator) Cluster list: 10.0.0.40 10.0.0.10
      AS path: Originator ID: 10.0.0.30
      Accepted
      Localpref: 100
      Router ID: 10.0.0.40
      Addpath Path ID: 2
BGP   Preference: 170/-101
      Next hop type: Indirect
      Next-hop reference count: 4
      Source: 10.0.0.40
      Next hop type: Router, Next hop index: 1045
      Next hop: 10.0.48.1 via lt-1/2/0.84, selected
      Protocol next hop: 10.0.15.2
      Indirect next hop: 91fc2ac 262153
      State: <Int Ext>
      Inactive reason: AS path
      Local AS:      1 Peer AS:      1
      Age: 1:56:51   Metric2: 3
      Task: BGP_1.10.0.0.40+179
      AS path: 2 2 I (Originator) Cluster list: 10.0.0.40
      AS path: Originator ID: 10.0.0.10

```

Accepted
Localpref: 100
Router ID: 10.0.0.40
Addpath Path ID: 3

**Related
Documentation**

- [Understanding the Advertisement of Multiple Paths to a Single Destination in BGP on page 3651](#)
- [Understanding Adding AS Numbers to BGP AS Paths](#)

Configuring ECMP Next Hops for RSVP and LDP LSPs for Load Balancing

The Junos OS supports configurations of 16, 32, or 64 equal-cost multipath (ECMP) next hops for RSVP and LDP LSPs on M10i routers with an Enhanced CFEB, M320, M120, MX Series, and T Series routers, and routing devices. For networks with high-volume traffic, this provides more flexibility to load-balance the traffic over as many as 64 LSPs.

To configure the maximum limit for ECMP next hops, include the **maximum-ecmp next-hops** statement at the **[edit chassis]** hierarchy level:

```
[edit chassis]
maximum-ecmp next-hops;
```

You can configure a maximum ECMP next-hop limit of 16, 32, or 64 using this statement. The default limit is 16.



NOTE: MX Series routers with one or more Modular Port Concentrator (MPC) cards and with Junos OS 11.4 or earlier installed, support the configuration of the **maximum-ecmp** statement with only 16 next hops. You should *not* configure the **maximum-ecmp** statement with 32 or 64 next hops. When you commit the configuration with 32 or 64 next hops, the following warning message appears:

Error: Number of members in Unilist NH exceeds the maximum supported 16 on Trio.

The following types of routes support the ECMP maximum next-hop configuration for as many as 64 ECMP gateways:

- Static IPv4 and IPv6 routes with direct and indirect next-hop ECMPs
- LDP ingress and transit routes learned through associated IGP routes
- RSVP ECMP next hops created for LSPs
- OSPF IPv4 and IPv6 route ECMPs
- ISIS IPv4 and IPv6 route ECMPs
- EBGp IPv4 and IPv6 route ECMPs
- IBGP (resolving over IGP routes) IPv4 and IPv6 route ECMPs

The enhanced ECMP limit of up to 64 ECMP next hops is also applicable for Layer 3 VPNs, Layer 2 VPNs, Layer 2 circuits, and VPLS services that resolve over an MPLS route, because the available ECMP paths in the MPLS route can also be used by such traffic.



NOTE:

The following FPCs on M320, T640, and T1600 routers only support 16 ECMP next hops:

- (M320, T640, and T1600 routers only) Enhanced II FPC1
- (M320, T640, and T1600 routers only) Enhanced II FPC2
- (M320 and T640 routers only) Enhanced II FPC3
- (T640 and T1600 routers only) FPC2
- (T640 and T1600 routers only) FPC3

If a maximum ECMP next-hop limit of 32 or 64 is configured on an M320, T640, or T1600 router with any of these FPCs installed, the Packet Forwarding Engines on these FPCs use only the first 16 ECMP next hops. For Packet Forwarding Engines on FPCs that support only 16 ECMP next hops, the Junos OS generates a system log message if a maximum ECMP next-hop limit of 32 or 64 is configured. However, for Packet Forwarding Engines on other FPCs installed on the router, a maximum configured ECMP limit of 32 or 64 ECMP next hops is applicable.



NOTE: If RSVP LSPs are configured with bandwidth allocation, for ECMP next hops with more than 16 LSPs, traffic is not distributed optimally based on bandwidths configured. Some LSPs with smaller allocated bandwidths receive more traffic than the ones configured with higher bandwidths. Traffic distribution does not strictly comply with the configured bandwidth allocation. This caveat is applicable to the following routers:

- T1600 and T640 routers with Enhanced Scaling FPC1, Enhanced Scaling FPC2, Enhanced Scaling FPC3, Enhanced Scaling FPC 4, and all Type 4 FPCs
- M320 routers with Enhanced III FPC1, Enhanced III FPC2, and Enhanced III FPC3
- MX Series routers with all types of FPCs and DPCs, excluding MPCs. This caveat is not applicable to MX Series routers with line cards based on the Junos Trio chipset.
- M120 routers with Type 1, Type 2, and Type 3 FPCs
- M10i routers with Enhanced CFEB

Next-hop cloning and permutations are disabled on T Series routers with Enhanced Scaling FPCs (Enhanced Scaling FPC1, Enhanced Scaling FPC2, Enhanced Scaling FPC3,

and Enhanced Scaling FPC 4) that support enhanced load-balancing capability. As a result, memory utilization is reduced for a highly scaled system with a high number of next hops on ECMP or aggregated interfaces. Next-hop cloning and permutations are also disabled on T Series routers with Type-4 FPCs.

To view the details of the ECMP next hops, issue the **show route** command. The **show route summary command** also shows the current configuration for the maximum ECMP limit. To view details of the ECMP LDP paths, issue the **traceroute mpls ldp** command.

Related Documentation

- *maximum-ecmp*

IBGP Scaling Configuration

- [Example: Configuring BGP Route Reflectors on page 3873](#)
- [Example: Configuring BGP Confederations on page 3890](#)

Example: Configuring BGP Route Reflectors

- [Understanding BGP Route Reflectors on page 3873](#)
- [Example: Configuring a Route Reflector on page 3876](#)

Understanding BGP Route Reflectors

This topic discusses using route reflectors to simplify configuration and aid in scaling. A further way to reduce the workload on a route reflector that is not in the traffic-forwarding path is to use the **no-install** statement at the **[edit protocols bgp family *family-name*]** hierarchy level. The **no-install** statement, introduced in Junos OS Release 15.1, eliminates interaction between the routing protocols daemon (rpd) and other components in the Junos system such as the kernel or the distributed firewall daemon (dfwd). This interaction is eliminated by prohibiting any routes in the associated rpd routing information bases (RIBs), also known as routing tables, from being published to those components.



NOTE: In releases previous to Junos OS 15.1, you can reduce the workload on a route reflector that is not in the traffic-forwarding path by using a forwarding-table export policy that rejects routes learned from BGP.

Because of the internal BGP (IBGP) full-mesh requirement, most networks use route reflectors to simplify configuration. The formula to compute the number of sessions required for a full mesh is $v * (v - 1) / 2$, where v is the number of BGP-enabled devices. The full-mesh model does not scale well. Using a route reflector, you group routers into clusters, which are identified by numeric identifiers unique to the autonomous system (AS). Within the cluster, you must configure a BGP session from a single router (the route reflector) to each internal peer. With this configuration, the IBGP full-mesh requirement is met.


To use route reflection in an AS, you designate one or more routers as a route reflector—typically, one per point of presence (POP). Route reflectors have the special BGP ability to readvertise routes learned from an internal peer to other internal peers. So rather than requiring all internal peers to be fully meshed with each other, route

reflection requires only that the route reflector be fully meshed with all internal peers. The route reflector and all of its internal peers form a cluster, as shown in [Figure 81](#).



NOTE: For some Juniper Networks devices, you must have an Advanced BGP Feature license installed on each device that uses a route reflector. For license details, see the *Installation and Upgrade Guide*.

Figure 81: Simple Route Reflector Topology (One Cluster)



[Figure 81](#) shows Router RR configured as the route reflector for Cluster 127. The other routers are designated internal peers within the cluster. BGP routes are advertised to Router RR by any of the internal peers. RR then readvertises those routes to all other peers within the cluster.

You can configure multiple clusters and link them by configuring a full mesh of route reflectors (see [Figure 82](#)).

Figure 82: Basic Route Reflection (Multiple Clusters)




Figure 82 shows Route Reflectors RR 1, RR 2, RR 3, and RR 4 as fully meshed internal peers. When a router advertises a route to RR 1, RR 1 readvertises the route to the other route reflectors, which, in turn, readvertise the route to the remaining routers within the AS. Route reflection allows the route to be propagated throughout the AS without the scaling problems created by the full mesh requirement.

However, as clusters become large, a full mesh with a route reflector becomes difficult to scale, as does a full mesh between route reflectors. To help offset this problem, you can group clusters of routers together into clusters of clusters for hierarchical route reflection (see Figure 83).

Figure 83: Hierarchical Route Reflection (Clusters of Clusters)

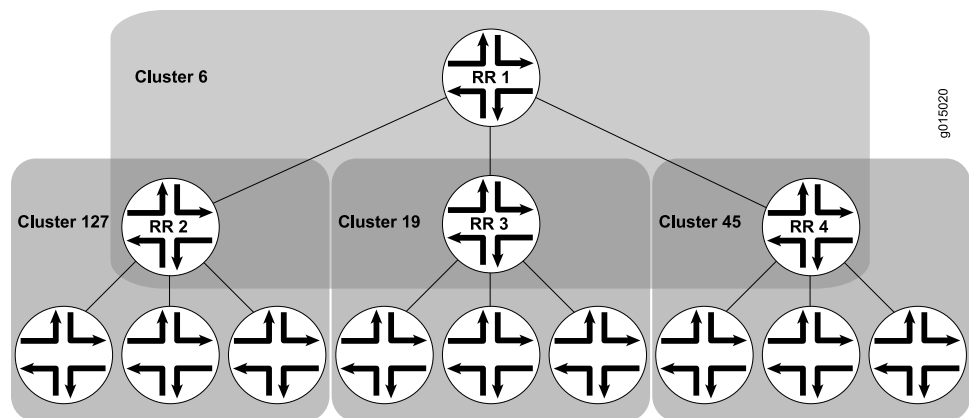


Figure 83 shows RR 2, RR 3, and RR 4 as the route reflectors for Clusters 127, 19, and 45, respectively. Rather than fully mesh those route reflectors, the network administrator has configured them as part of another cluster (Cluster 6) for which RR 1 is the route

reflector. When a router advertises a route to RR 2, RR 2 readvertises the route to all the routers within its own cluster, and then readvertises the route to RR 1. RR 1 readvertises the route to the routers in its cluster, and those routers propagate the route down through their clusters.

Example: Configuring a Route Reflector

This example shows how to configure a route reflector.

- [Requirements on page 3876](#)
- [Overview on page 3876](#)
- [Configuration on page 3877](#)
- [Verification on page 3885](#)

Requirements

No special configuration beyond device initialization is required before you configure this example.

Overview

Generally, internal BGP (IBGP)-enabled devices need to be fully meshed, because IBGP does not readvertise updates to other IBGP-enabled devices. The full mesh is a logical mesh achieved through configuration of multiple **neighbor** statements on each IBGP-enabled device. The full mesh is not necessarily a physical full mesh. Maintaining a full mesh (logical or physical) does not scale well in large deployments.

[Figure 84](#) shows an IBGP network with Device A acting as a route reflector. Device B and Device C are clients of the route reflector. Device D and Device E are outside the cluster, so they are nonclients of the route reflector.

On Device A (the route reflector), you must form peer relationships with all of the IBGP-enabled devices by including the **neighbor** statement for the clients (Device B and Device C) and the nonclients (Device D and Device E). You must also include the **cluster** statement and a cluster identifier. The cluster identifier can be any 32-bit value. This example uses the loopback interface IP address of the route reflector.


On Device B and Device C, the route reflector clients, you only need one **neighbor** statement that forms a peer relationship with the route reflector, Device A.

On Device D and Device E, the nonclients, you need a **neighbor** statement for each nonclient device (D-to-E and E-to-D). You also need a **neighbor** statement for the route reflector (D-to-A and E-to-A). Device D and Device E do not need **neighbor** statements for the client devices (Device B and Device C).



TIP: Device D and Device E are considered to be nonclients because they have explicitly configured peer relationships with each other. To make them RR route reflector clients, remove the **neighbor 192.168.5.5** statement from the configuration on Device D, and remove the **neighbor 192.168.0.1** statement from the configuration on Device E.

Figure 84: IBGP Network Using a Route Reflector



Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device A

```
set interfaces fe-0/0/0 unit 1 description to-B
set interfaces fe-0/0/0 unit 1 family inet address 10.10.10.1/30
set interfaces fe-0/0/1 unit 3 description to-D
set interfaces fe-0/0/1 unit 3 family inet address 10.10.10.9/30
set interfaces lo0 unit 1 family inet address 192.168.6.5/32
set protocols bgp group internal-peers type internal
set protocols bgp group internal-peers local-address 192.168.6.5
set protocols bgp group internal-peers export send-ospf
set protocols bgp group internal-peers cluster 192.168.6.5
set protocols bgp group internal-peers neighbor 192.163.6.4
set protocols bgp group internal-peers neighbor 192.168.40.4
set protocols bgp group internal-peers neighbor 192.168.0.1
set protocols bgp group internal-peers neighbor 192.168.5.5
set protocols ospf area 0.0.0.0 interface lo0.1 passive
set protocols ospf area 0.0.0.0 interface fe-0/0/0.1
set protocols ospf area 0.0.0.0 interface fe-0/0/1.3
set policy-options policy-statement send-ospf term 2 from protocol ospf
```

```
set policy-options policy-statement send-ospf term 2 then accept
set routing-options router-id 192.168.6.5
set routing-options autonomous-system 17
```

Device B

```
set interfaces fe-0/0/0 unit 2 description to-A
set interfaces fe-0/0/0 unit 2 family inet address 10.10.10.2/30
set interfaces fe-0/0/1 unit 5 description to-C
set interfaces fe-0/0/1 unit 5 family inet address 10.10.10.5/30
set interfaces lo0 unit 2 family inet address 192.163.6.4/32
set protocols bgp group internal-peers type internal
set protocols bgp group internal-peers local-address 192.163.6.4
set protocols bgp group internal-peers export send-ospf
set protocols bgp group internal-peers neighbor 192.168.6.5
set protocols ospf area 0.0.0.0 interface lo0.2 passive
set protocols ospf area 0.0.0.0 interface fe-0/0/0.2
set protocols ospf area 0.0.0.0 interface fe-0/0/1.5
set policy-options policy-statement send-ospf term 2 from protocol ospf
set policy-options policy-statement send-ospf term 2 then accept
set routing-options router-id 192.163.6.4
set routing-options autonomous-system 17
```

Device C

```
set interfaces fe-0/0/0 unit 6 description to-B
set interfaces fe-0/0/0 unit 6 family inet address 10.10.10.6/30
set interfaces lo0 unit 3 family inet address 192.168.40.4/32
set protocols bgp group internal-peers type internal
set protocols bgp group internal-peers local-address 192.168.40.4
set protocols bgp group internal-peers export send-ospf
set protocols bgp group internal-peers neighbor 192.168.6.5
set protocols ospf area 0.0.0.0 interface lo0.3 passive
set protocols ospf area 0.0.0.0 interface fe-0/0/0.6
set policy-options policy-statement send-ospf term 2 from protocol ospf
set policy-options policy-statement send-ospf term 2 then accept
set routing-options router-id 192.168.40.4
set routing-options autonomous-system 17
```

Device D

```
set interfaces fe-0/0/0 unit 4 description to-A
set interfaces fe-0/0/0 unit 4 family inet address 10.10.10.10/30
set interfaces fe-0/0/1 unit 7 description to-E
set interfaces fe-0/0/1 unit 7 family inet address 10.10.10.13/30
set interfaces lo0 unit 4 family inet address 192.168.0.1/32
set protocols bgp group internal-peers type internal
set protocols bgp group internal-peers local-address 192.168.0.1
set protocols bgp group internal-peers export send-ospf
set protocols bgp group internal-peers neighbor 192.168.6.5
set protocols bgp group internal-peers neighbor 192.168.5.5
set protocols ospf area 0.0.0.0 interface lo0.4 passive
set protocols ospf area 0.0.0.0 interface fe-0/0/0.4
set protocols ospf area 0.0.0.0 interface fe-0/0/1.7
set policy-options policy-statement send-ospf term 2 from protocol ospf
set policy-options policy-statement send-ospf term 2 then accept
set routing-options router-id 192.168.0.1
set routing-options autonomous-system 17
```

Device E

```
set interfaces fe-0/0/0 unit 8 description to-D
set interfaces fe-0/0/0 unit 8 family inet address 10.10.10.14/30
```

```

set interfaces lo0 unit 5 family inet address 192.168.5.5/32
set protocols bgp group internal-peers type internal
set protocols bgp group internal-peers local-address 192.168.5.5
set protocols bgp group internal-peers export send-ospf
set protocols bgp group internal-peers neighbor 192.168.0.1
set protocols bgp group internal-peers neighbor 192.168.6.5
set protocols ospf area 0.0.0.0 interface lo0.5 passive
set protocols ospf area 0.0.0.0 interface fe-0/0/0.8
set policy-options policy-statement send-ospf term 2 from protocol ospf
set policy-options policy-statement send-ospf term 2 then accept
set routing-options router-id 192.168.5.5
set routing-options autonomous-system 17

```

Configuring the Route Reflector

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IBGP in the network using Juniper Networks Device A as a route reflector:

1. Configure the interfaces.

```

[edit interfaces]
user@A# set fe-0/0/0 unit 1 description to-B
user@A# set fe-0/0/0 unit 1 family inet address 10.10.10.1/30
user@A# set fe-0/0/1 unit 3 description to-D
user@A# set fe-0/0/1 unit 3 family inet address 10.10.10.9/30
user@A# set lo0 unit 1 family inet address 192.168.6.5/32

```

2. Configure BGP, including the cluster identifier and neighbor relationships with all IBGP-enabled devices in the autonomous system (AS).

Also apply the policy that redistributes OSPF routes into BGP.

```

[edit protocols bgp group internal-peers]
user@A# set type internal
user@A# set local-address 192.168.6.5
user@A# set export send-ospf
user@A# set cluster 192.168.6.5
user@A# set neighbor 192.163.6.4
user@A# set neighbor 192.168.40.4
user@A# set neighbor 192.168.0.1
user@A# set neighbor 192.168.5.5

```

3. Configure static routing or an interior gateway protocol (IGP).

This example uses OSPF.

```

[edit protocols ospf area 0.0.0.0]
user@A# set interface lo0.1 passive
user@A# set interface fe-0/0/0.1
user@A# set interface fe-0/0/1.3

```

4. Configure the policy that redistributes OSPF routes into BGP.

```

[edit policy-options policy-statement send-ospf term 2]
user@A# set from protocol ospf
user@A# set then accept

```

5. Configure the router ID and the autonomous system (AS) number.

```
[edit routing-options]
user@A# set router-id 192.168.6.5
user@A# set autonomous-system 17
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@A# show interfaces
fe-0/0/0 {
  unit 1 {
    description to-B;
    family inet {
      address 10.10.10.1/30;
    }
  }
}
fe-0/0/1 {
  unit 3 {
    description to-D;
    family inet {
      address 10.10.10.9/30;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 192.168.6.5/32;
    }
  }
}

user@A# show protocols
bgp {
  group internal-peers {
    type internal;
    local-address 192.168.6.5;
    export send-ospf;
    cluster 192.168.6.5;
    neighbor 192.163.6.4;
    neighbor 192.168.40.4;
    neighbor 192.168.0.1;
    neighbor 192.168.5.5;
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.1 {
      passive;
    }
    interface fe-0/0/0.1;
    interface fe-0/0/1.3;
```



```

    }
  }

user@A# show policy-options
policy-statement send-ospf {
  term 2 {
    from protocol ospf;
    then accept;
  }
}

user@A# show routing-options
router-id 192.168.6.5;
autonomous-system 17;

```

If you are done configuring the device, enter **commit** from configuration mode.



NOTE: Repeat these steps for each nonclient BGP peer within the cluster that you are configuring, if the other nonclient devices are from Juniper Networks. Otherwise, consult the device's documentation for instructions.

Configuring Client Peers

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure client peers:

1. Configure the interfaces.


```

[edit interfaces]
user@B# set fe-0/0/0 unit 2 description to-A
user@B# set fe-0/0/0 unit 2 family inet address 10.10.10.2/30
user@B# set fe-0/0/1 unit 5 description to-C
user@B# set fe-0/0/1 unit 5 family inet address 10.10.10.5/30
user@B# set lo0 unit 2 family inet address 192.163.6.4/32

```

2. Configure the BGP neighbor relationship with the route reflector.

Also apply the policy that redistributes OSPF routes into BGP.

```

[edit protocols bgp group internal-peers]
user@B# set type internal
user@B# set local-address 192.163.6.4
user@B# set export send-ospf
user@B# set neighbor 192.168.6.5

```

3. Configure OSPF.


```

[edit protocols ospf area 0.0.0.0]
user@B# set interface lo0.2 passive
user@B# set interface fe-0/0/0.2
user@B# set interface fe-0/0/1.5

```
4. Configure the policy that redistributes OSPF routes into BGP.

```
[edit policy-options policy-statement send-ospf term 2]
user@B# set from protocol ospf
user@B# set then accept
```

5. Configure the router ID and the AS number.

```
[edit routing-options]
user@B# set router-id 192.163.6.4
user@B# set autonomous-system 17
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@B# show interfaces
fe-0/0/0 {
  unit 2 {
    description to-A;
    family inet {
      address 10.10.10.2/30;
    }
  }
}
fe-0/0/1 {
  unit 5 {
    description to-C;
    family inet {
      address 10.10.10.5/30;
    }
  }
}
lo0 {
  unit 2 {
    family inet {
      address 192.163.6.4/32;
    }
  }
}
```

```
user@B# show protocols
bgp {
  group internal-peers {
    type internal;
    local-address 192.163.6.4;
    export send-ospf;
    neighbor 192.168.6.5;
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.2 {
      passive;
    }
    interface fe-0/0/0.2;
    interface fe-0/0/1.5;
```

```

    }
  }

user@B# show policy-options
policy-statement send-ospf {
  term 2 {
    from protocol ospf;
    then accept;
  }
}

user@B# show routing-options
router-id 192.163.6.4;
autonomous-system 17;

```

If you are done configuring the device, enter **commit** from configuration mode.



NOTE: Repeat these steps for each client BGP peer within the cluster that you are configuring if the other client devices are from Juniper Networks. Otherwise, consult the device's documentation for instructions.

Configuring Nonclient Peers

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure nonclient peers:

1. Configure the interfaces.


```

[edit interfaces]
user@D# set fe-0/0/0 unit 4 description to-A
user@D# set fe-0/0/0 unit 4 family inet address 10.10.10.10/30
user@D# set fe-0/0/1 unit 7 description to-E
user@D# set fe-0/0/1 unit 7 family inet address 10.10.10.13/30
user@D# set lo0 unit 4 family inet address 192.168.0.1/32
      
```
2. Configure the BGP neighbor relationships with the RRroute reflector and with the other nonclient peers.

Also apply the policy that redistributes OSPF routes into BGP.

```

[edit protocols bgp group internal-peers]
user@D# set type internal
user@D# set local-address 192.168.0.1
user@D# set export send-ospf
user@D# set neighbor 192.168.6.5
user@D# set neighbor 192.168.5.5

```

3. Configure OSPF.


```

[edit protocols ospf area 0.0.0.0]
user@D# set interface lo0.4 passive
user@D# set interface fe-0/0/0.4
      
```

```
user@D# set interface fe-0/0/1.7
```

4. Configure the policy that redistributes OSPF routes into BGP.

```
[edit policy-options policy-statement send-ospf term 2]
user@D# set from protocol ospf
user@D# set then accept
```

5. Configure the router ID and the AS number.

```
[edit routing-options]
user@D# set router-id 192.168.0.1
user@D# set autonomous-system 17
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@D# show interfaces
fe-0/0/0 {
  unit 4 {
    description to-A;
    family inet {
      address 10.10.10.10/30;
    }
  }
}
fe-0/0/1 {
  unit 7 {
    description to-E;
    family inet {
      address 10.10.10.13/30;
    }
  }
}
lo0 {
  unit 4 {
    family inet {
      address 192.168.0.1/32;
    }
  }
}
```

```
user@D# show protocols
bgp {
  group internal-peers {
    type internal;
    local-address 192.168.0.1;
    export send-ospf;
    neighbor 192.168.6.5;
    neighbor 192.168.5.5;
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.4 {
```

```

        passive;
    }
    interface fe-0/0/0.4;
    interface fe-0/0/1.7;
}
}

user@D# show policy-options
policy-statement send-ospf {
    term 2 {
        from protocol ospf;
        then accept;
    }
}

user@D# show routing-options
router-id 192.168.0.1;
autonomous-system 17;

```

If you are done configuring the device, enter **commit** from configuration mode.



NOTE: Repeat these steps for each nonclient BGP peer within the cluster that you are configuring if the other nonclient devices are from Juniper Networks. Otherwise, consult the device's documentation for instructions.

Verification

Confirm that the configuration is working properly.

- [Verifying BGP Neighbors on page 3885](#)
- [Verifying BGP Groups on page 3888](#)
- [Verifying BGP Summary Information on page 3888](#)
- [Verifying Routing Table Information on page 3888](#)

Verifying BGP Neighbors

Purpose Verify that BGP is running on configured interfaces and that the BGP session is established for each neighbor address.

Action From operational mode, enter the **show bgp neighbor** command.

```

user@A> show bgp neighbor
Peer: 192.163.6.4+179 AS 17    Local: 192.168.6.5+62857 AS 17
  Type: Internal    State: Established (route reflector client)Flags: <Sync>
  Last State: OpenConfirm    Last Event: RecvKeepAlive
  Last Error: None
  Export: [ send-ospf ]
  Options: <Preference LocalAddress Cluster Refresh>
  Local Address: 192.168.6.5 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 192.163.6.4    Local ID: 192.168.6.5    Active Holdtime: 90
  Keepalive Interval: 30    Peer index: 0
  BFD: disabled, down

```

```

NLRI for restart configured on peer: inet-unicast
NLRI advertised by peer: inet-unicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-unicast
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Peer supports 4 byte AS extension (peer-as 17)
Peer does not support Addpath
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        6
  Accepted prefixes:        1
  Suppressed due to damping: 0
  Advertised prefixes:      6
Last traffic (seconds): Received 5    Sent 3    Checked 19
Input messages: Total 2961    Updates 7    Refreshes 0    Octets 56480
Output messages: Total 2945    Updates 6    Refreshes 0    Octets 56235
Output Queue[0]: 0

Peer: 192.168.0.1+179 AS 17    Local: 192.168.6.5+60068 AS 17
Type: Internal    State: Established (route reflector client)Flags: <Sync>
Last State: OpenConfirm    Last Event: RecvKeepAlive
Last Error: None
Export: [ send-ospf ]
Options: <Preference LocalAddress Cluster Refresh>
Local Address: 192.168.6.5 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer ID: 192.168.0.1    Local ID: 192.168.6.5    Active Holdtime: 90
Keepalive Interval: 30    Peer index: 3
BFD: disabled, down
NLRI for restart configured on peer: inet-unicast
NLRI advertised by peer: inet-unicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-unicast
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Peer supports 4 byte AS extension (peer-as 17)
Peer does not support Addpath
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        6
  Accepted prefixes:        1
  Suppressed due to damping: 0
  Advertised prefixes:      6
Last traffic (seconds): Received 18    Sent 20    Checked 12
Input messages: Total 15    Updates 5    Refreshes 0    Octets 447
Output messages: Total 554    Updates 4    Refreshes 0    Octets 32307

```

Output Queue[0]: 0

```

Peer: 192.168.5.5+57458 AS 17 Local: 192.168.6.5+179 AS 17
Type: Internal State: Established (route reflector client)Flags: <Sync>
Last State: OpenConfirm Last Event: RecvKeepAlive
Last Error: None
Export: [ send-ospf ]
Options: <Preference LocalAddress Cluster Refresh>
Local Address: 192.168.6.5 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer ID: 192.168.5.5 Local ID: 192.168.6.5 Active Holdtime: 90
Keepalive Interval: 30 Peer index: 2
BFD: disabled, down
NLRI for restart configured on peer: inet-unicast
NLRI advertised by peer: inet-unicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-unicast
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Peer supports 4 byte AS extension (peer-as 17)
Peer does not support Addpath
Table inet.0 Bit: 10000
RIB State: BGP restart is complete
Send state: in sync
Active prefixes: 0
Received prefixes: 7
Accepted prefixes: 7
Suppressed due to damping: 0
Advertised prefixes: 6
Last traffic (seconds): Received 17 Sent 3 Checked 9
Input messages: Total 2967 Updates 7 Refreshes 0 Octets 56629
Output messages: Total 2943 Updates 6 Refreshes 0 Octets 56197
Output Queue[0]: 0

```

```

Peer: 192.168.40.4+53990 AS 17 Local: 192.168.6.5+179 AS 17
Type: Internal State: Established (route reflector client)Flags: <Sync>
Last State: OpenConfirm Last Event: RecvKeepAlive
Last Error: None
Export: [ send-ospf ]
Options: <Preference LocalAddress Cluster Refresh>
Local Address: 192.168.6.5 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer ID: 192.168.40.4 Local ID: 192.168.6.5 Active Holdtime: 90
Keepalive Interval: 30 Peer index: 1
BFD: disabled, down
NLRI for restart configured on peer: inet-unicast
NLRI advertised by peer: inet-unicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-unicast
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast

```

```

Peer supports 4 byte AS extension (peer-as 17)
Peer does not support Addpath
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        7
  Accepted prefixes:        7
  Suppressed due to damping: 0
  Advertised prefixes:      6
Last traffic (seconds): Received 5   Sent 23   Checked 52
Input messages: Total 2960   Updates 7   Refreshes 0   Octets 56496
Output messages: Total 2943   Updates 6   Refreshes 0   Octets 56197
Output Queue[0]: 0

```

Verifying BGP Groups

Purpose Verify that the BGP groups are configured correctly.

Action From operational mode, enter the **show bgp group** command.

```

user@A> show bgp group
Group Type: Internal      AS: 17                      Local AS: 17
Name: internal-peers     Index: 0                    Flags: <>
Export: [ send-ospf ]
Options: <Cluster>
Holdtime: 0
Total peers: 4           Established: 4
192.163.6.4+179
192.168.40.4+53990
192.168.0.1+179
192.168.5.5+57458
inet.0: 0/26/16/0

Groups: 1  Peers: 4   External: 0   Internal: 4   Down peers: 0   Flaps: 0
Table      Tot Paths  Act Paths Suppressed  History Damp State   Pending
inet.0          26         0         0         0         0         0         0

```

Verifying BGP Summary Information

Purpose Verify that the BGP configuration is correct.

Action From operational mode, enter the **show bgp summary** command.

```

user@A> show bgp summary

Groups: 1 Peers: 4 Down peers: 0
Table      Tot Paths  Act Paths Suppressed  History Damp State   Pending
inet.0          26         0         0         0         0         0         0
Peer      AS      InPkt    OutPkt    OutQ   Flaps  Last Up/Dwn
State|#Active/Received/Accepted/Damped...
192.163.6.4      17      2981      2965        0        0   22:19:15 0/6/1/0      0/0/0/0
192.168.0.1      17        36        575        0        0    13:43 0/6/1/0      0/0/0/0
192.168.5.5      17      2988      2964        0        0   22:19:10 0/7/7/0      0/0/0/0
192.168.40.4     17      2980      2964        0        0   22:19:14 0/7/7/0      0/0/0/0

```

Verifying Routing Table Information

Purpose Verify that the routing table contains the IBGP routes.

Action From operational mode, enter the **show route** command.

```

user@A> show route
inet.0: 12 destinations, 38 routes (12 active, 0 holddown, 10 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.10.0/30      * [Direct/0] 22:22:03
                  > via fe-0/0/0.1
                  [BGP/170] 22:20:55, MED 2, localpref 100, from 192.168.40.4
                  AS path: I
                  > to 10.10.10.2 via fe-0/0/0.1
                  [BGP/170] 22:20:51, MED 3, localpref 100, from 192.168.5.5
                  AS path: I
                  > to 10.10.10.10 via fe-0/0/1.3
10.10.10.1/32     * [Local/0] 22:22:03
                  Local via fe-0/0/0.1
10.10.10.4/30     * [OSPF/10] 22:21:13, metric 2
                  > to 10.10.10.2 via fe-0/0/0.1
                  [BGP/170] 22:20:51, MED 4, localpref 100, from 192.168.5.5
                  AS path: I
                  > to 10.10.10.10 via fe-0/0/1.3
10.10.10.8/30     * [Direct/0] 22:22:03
                  > via fe-0/0/1.3
                  [BGP/170] 22:20:51, MED 2, localpref 100, from 192.168.5.5
                  AS path: I
                  > to 10.10.10.10 via fe-0/0/1.3
                  [BGP/170] 22:20:55, MED 3, localpref 100, from 192.168.40.4
                  AS path: I
                  > to 10.10.10.2 via fe-0/0/0.1
10.10.10.9/32     * [Local/0] 22:22:03
                  Local via fe-0/0/1.3
10.10.10.12/30    * [OSPF/10] 22:21:08, metric 2
                  > to 10.10.10.10 via fe-0/0/1.3
                  [BGP/170] 22:20:55, MED 4, localpref 100, from 192.168.40.4
                  AS path: I
                  > to 10.10.10.2 via fe-0/0/0.1
192.163.6.4/32    * [OSPF/10] 22:21:13, metric 1
                  > to 10.10.10.2 via fe-0/0/0.1
                  [BGP/170] 22:20:55, MED 1, localpref 100, from 192.168.40.4
                  AS path: I
                  > to 10.10.10.2 via fe-0/0/0.1
                  [BGP/170] 22:20:51, MED 3, localpref 100, from 192.168.5.5
                  AS path: I
                  > to 10.10.10.10 via fe-0/0/1.3
192.168.0.1/32    * [OSPF/10] 22:21:08, metric 1
                  > to 10.10.10.10 via fe-0/0/1.3
                  [BGP/170] 22:20:51, MED 1, localpref 100, from 192.168.5.5
                  AS path: I
                  > to 10.10.10.10 via fe-0/0/1.3
                  [BGP/170] 22:20:55, MED 3, localpref 100, from 192.168.40.4
                  AS path: I
                  > to 10.10.10.2 via fe-0/0/0.1
192.168.5.5/32    * [OSPF/10] 22:21:08, metric 2
                  > to 10.10.10.10 via fe-0/0/1.3
                  [BGP/170] 00:15:24, MED 1, localpref 100, from 192.168.0.1
                  AS path: I
                  > to 10.10.10.10 via fe-0/0/1.3
                  [BGP/170] 22:20:55, MED 4, localpref 100, from 192.168.40.4
                  AS path: I
                  > to 10.10.10.2 via fe-0/0/0.1
192.168.6.5/32    * [Direct/0] 22:22:04

```

```
> via lo0.1
[BGP/170] 22:20:51, MED 2, localpref 100, from 192.168.5.5
AS path: I
> to 10.10.10.10 via fe-0/0/1.3
[BGP/170] 22:20:55, MED 2, localpref 100, from 192.168.40.4
AS path: I
192.168.40.4/32 > to 10.10.10.2 via fe-0/0/0.1
*[OSPF/10] 22:21:13, metric 2
> to 10.10.10.2 via fe-0/0/0.1
[BGP/170] 22:20:55, MED 1, localpref 100, from 192.163.6.4
AS path: I
> to 10.10.10.2 via fe-0/0/0.1
[BGP/170] 22:20:51, MED 4, localpref 100, from 192.168.5.5
AS path: I
224.0.0.5/32 > to 10.10.10.10 via fe-0/0/1.3
*[OSPF/10] 22:22:07, metric 1
MultiRecv
```

- Related Documentation**
- [Understanding External BGP Peering Sessions on page 3601](#)
 - [BGP Configuration Overview](#)

Example: Configuring BGP Confederations

- [Understanding BGP Confederations on page 3890](#)
- [Example: Configuring BGP Confederations on page 3891](#)

Understanding BGP Confederations

BGP confederations are another way to solve the scaling problems created by the BGP full mesh requirement. BGP confederations effectively break up a large autonomous system (AS) into subautonomous systems (sub-ASs). Each sub-AS must be uniquely identified within the confederation AS by a sub-AS number. Typically, sub-AS numbers are taken from the private AS numbers between 64,512 and 65,535.

Within a sub-AS, the same internal BGP (IBGP) full mesh requirement exists. Connections to other confederations are made with standard external BGP (EBGP), and peers outside the sub-AS are treated as external. To avoid routing loops, a sub-AS uses a confederation sequence, which operates like an AS path but uses only the privately assigned sub-AS numbers.

The confederation AS appears whole to other confederation ASs. The AS path received by other ASs shows only the globally assigned AS number. It does not include the confederation sequence or the privately assigned sub-AS numbers. The sub-AS numbers are removed when the route is advertised out of the confederation AS. [Figure 85](#) shows an AS divided into four confederations.

Figure 85: BGP Confederations




Figure 85 shows AS 3 divided into four sub-ASs, 64517, 64550, 65300, and 65410, which are linked through EBGP sessions. Because the confederations are connected by EBGP, they do not need to be fully meshed. EBGP routes are readvertised to other sub-ASs.

Example: Configuring BGP Confederations

This example shows how to configure BGP confederations.

- [Requirements on page 3891](#)
- [Overview on page 3891](#)
- [Configuration on page 3892](#)
- [Verification on page 3894](#)

Requirements

- Configure network interfaces.
- Configure external peer sessions. See [“Example: Configuring External BGP Point-to-Point Peer Sessions” on page 3602](#).
- Configure interior gateway protocol (IGP) sessions between peers.
- Configure a routing policy to advertise the BGP routes.

Overview


Within a BGP confederation, the links between the confederation member autonomous systems (ASs) must be external BGP (EBGP) links, not internal BGP (IBGP) links.

Similar to route reflectors, BGP confederations reduce the number of peer sessions and TCP sessions to maintain connections between IBGP routing devices. BGP confederation is one method used to solve the scaling problems created by the IBGP full mesh requirement. BGP confederations effectively break up a large AS into subautonomous

systems. Each sub-AS must be uniquely identified within the confederation AS by a sub-AS number. Typically, sub-AS numbers are taken from the private AS numbers between 64512 and 65535. Within a sub-AS, the same IBGP full mesh requirement exists. Connections to other confederations are made with standard EBGP, and peers outside the sub-AS are treated as external. To avoid routing loops, a sub-AS uses a confederation sequence, which operates like an AS path but uses only the privately assigned sub-AS numbers.

Figure 86 shows a sample network in which AS 17 has two separate confederations: sub-AS 64512 and sub-AS 64513, each of which has multiple routers. Within a sub-AS, an IGP is used to establish network connectivity with internal peers. Between sub-ASs, an EBGP peer session is established.

Figure 86: Typical Network Using BGP Confederations



Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

All Devices in Sub-AS 64512	<pre> set routing-options autonomous-system 64512 set routing-options confederation 17 members 64512 set routing-options confederation 17 members 64513 set protocols bgp group sub-AS-64512 type internal set protocols bgp group sub-AS-64512 local-address 192.168.5.1 set protocols bgp group sub-AS-64512 neighbor 192.168.8.1 set protocols bgp group sub-AS-64512 neighbor 192.168.15.1 </pre>
Border Device in Sub-AS 64512	<pre> set protocols bgp group to-sub-AS-64513 type external set protocols bgp group to-sub-AS-64513 peer-as 64513 set protocols bgp group to-sub-AS-64513 neighbor 192.168.5.2 </pre>
All Devices in Sub-AS 64513	<pre> set routing-options autonomous-system 64513 set routing-options confederation 17 members 64512 set routing-options confederation 17 members 64513 set protocols bgp group sub-AS-64513 type internal set protocols bgp group sub-AS-64513 local-address 192.168.5.2 </pre>

	<pre> set protocols bgp group sub-AS-64513 neighbor 192.168.9.1 set protocols bgp group sub-AS-64513 neighbor 192.168.16.1 </pre>
Border Device in Sub-AS 64513	<pre> set protocols bgp group to-sub-AS-64512 type external set protocols bgp group to-sub-AS-64512 peer-as 64512 set protocols bgp group to-sub-AS-64512 neighbor 192.168.5.1 </pre>
Step-by-Step Procedure	<p>This procedure shows the steps for the devices that are in sub-AS 64512.</p> <p>The autonomous-system statement sets the sub-AS number of the device.</p> <p>The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see <i>Using the CLI Editor in Configuration Mode</i> in the <i>CLI User Guide</i>.</p> <p>To configure BGP confederations:</p> <ol style="list-style-type: none"> Set the sub-AS number for the device. <pre> [edit routing-options] user@host# set autonomous-system 64512 </pre> In the confederation, include all sub-ASs in the main AS. <p>The number 17 represents the main AS. The members statement lists all the sub-ASs in the main AS.</p> <pre> [edit routing-options confederation] user@host# set 17 members 64512 user@host# set 17 members 64513 </pre> On the border device in sub-AS 64512, configure an EBGP connection to the border device in AS 64513. <pre> [edit protocols bgp group to-sub-AS-64513] user@host# set type external user@host# set neighbor 192.168.5.2 user@host# set peer-as 64513 </pre> Configure an IBGP group for peering with the devices within sub-AS 64512. <pre> [edit protocols bgp group sub-AS-64512] user@host# set type internal user@host# set local-address 192.168.5.1 user@host# neighbor 192.168.8.1 user@host# neighbor 192.168.15.1 </pre>
Results	<p>From configuration mode, confirm your configuration by entering the show routing-options and show protocols commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.</p> <pre> user@host# show routing-options autonomous-system 64512; confederation 17 members [64512 64513]; user@host# show protocols bgp { group to-sub-AS-64513 { # On the border devices only </pre>

```
    type external;
    peer-as 64513;
    neighbor 192.168.5.2;
  }
  group sub-AS-64512 {
    type internal;
    local-address 192.168.5.1;
    neighbor 192.168.8.1;
    neighbor 192.168.15.1;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.
Repeat these steps for sSub-AS 64513.

Verification

Confirm that the configuration is working properly.

- [Verifying BGP Neighbors on page 3894](#)
- [Verifying BGP Groups on page 3895](#)
- [Verifying BGP Summary Information on page 3896](#)

Verifying BGP Neighbors

Purpose Verify that BGP is running on configured interfaces and that the BGP session is active for each neighbor address.

Action From the CLI, enter the **show bgp neighbor** command.

Sample Output

```
user@host> show bgp neighbor
Peer: 10.255.245.12+179 AS 35 Local: 10.255.245.13+2884 AS 35
  Type: Internal    State: Established (route reflector client)Flags: Sync
  Last State: OpenConfirm  Last Event: RecvKeepAlive
  Last Error: None
  Options: Preference LocalAddress HoldTime Cluster AddressFamily Rib-group Refresh

  Address families configured: inet-vpn-unicast inet-labeled-unicast
  Local Address: 10.255.245.13 Holdtime: 90 Preference: 170
  Flags for NLRI inet-vpn-unicast: AggregateLabel
  Flags for NLRI inet-labeled-unicast: AggregateLabel
  Number of flaps: 0
  Peer ID: 10.255.245.12 Local ID: 10.255.245.13 Active Holdtime: 90
  Keepalive Interval: 30
  NLRI advertised by peer: inet-vpn-unicast inet-labeled-unicast
  NLRI for this session: inet-vpn-unicast inet-labeled-unicast
  Peer supports Refresh capability (2)
Restart time configured on the peer: 300
  Stale routes from peer are kept for: 60
  Restart time requested by this peer: 300
  NLRI that peer supports restart for: inet-unicast inet6-unicast
  NLRI that restart is negotiated for: inet-unicast inet6-unicast
  NLRI of received end-of-rib markers: inet-unicast inet6-unicast
  NLRI of all end-of-rib markers sent: inet-unicast inet6-unicast
  Table inet.0 Bit: 10000
```

```

RIB State: restart is complete
Send state: in sync
Active prefixes: 4
Received prefixes: 6
Suppressed due to damping: 0
Table inet6.0 Bit: 20000
RIB State: restart is complete
Send state: in sync
Active prefixes: 0
Received prefixes: 2
Suppressed due to damping: 0
Last traffic (seconds): Received 3    Sent 3    Checked 3
Input messages: Total 9    Updates 6    Refreshes 0    Octets 403
Output messages: Total 7    Updates 3    Refreshes 0    Octets 365
Output Queue[0]: 0
Output Queue[1]: 0
Trace options: detail packets
Trace file: /var/log/bgpr size 131072 files 10

```

Meaning The output shows a list of the BGP neighbors with detailed session information. Verify the following information:

- Each configured peering neighbor is listed.
- For **State**, each BGP session is **Established**.
- For **Type**, each peer is configured as the correct type (either internal or external).
- For **AS**, the AS number of the BGP neighbor is correct.

Verifying BGP Groups

Purpose Verify that the BGP groups are configured correctly.

Action From the CLI, enter the **show bgp group** command.

Sample Output

```

user@host> show bgp group
Group Type: Internal  AS: 10045      Local AS: 10045
Name: pe-to-asbr2                      Flags: Export Eval
Export: [ match-all ]
Total peers: 1      Established: 1
10.0.0.4+179
bgp.13vpn.0: 1/1/0
vpn-green.inet.0: 1/1/0

Groups: 1  Peers: 1  External: 0  Internal: 1  Down peers: 0  Flaps: 0
Table      Tot Paths  Act Paths  Suppressed  History Damp State  Pending
bgp.13vpn.0      1          1          0          0          0          0

```

Meaning The output shows a list of the BGP groups with detailed group information. Verify the following information:

- Each configured group is listed.
- For **AS**, each group's remote AS is configured correctly.
- For **Local AS**, each group's local AS is configured correctly.

- For **Group Type**, each group has the correct type (either internal or external).
- For **Total peers**, the expected number of peers within the group is shown.
- For **Established**, the expected number of peers within the group have BGP sessions in the **Established** state.
- The IP addresses of all the peers within the group are present.

Verifying BGP Summary Information

Purpose Verify that the BGP configuration is correct.

Action From the CLI, enter the **show bgp summary** command.

Sample Output

```
user@host> show bgp summary
Groups: 1 Peers: 3 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet.0 6 4 0 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Damped...
10.0.0.2 65002 88675 88652 0 2 42:38 2/4/0
0/0/0
10.0.0.3 65002 54528 54532 0 1 2w4d22h 0/0/0
0/0/0
10.0.0.4 65002 51597 51584 0 0 2w3d22h 2/2/0
0/0/0
```

Meaning The output shows a summary of BGP session information. Verify the following information:

- For **Groups**, the total number of configured groups is shown.
- For **Peers**, the total number of BGP peers is shown.
- For **Down Peers**, the total number of unestablished peers is 0. If this value is not zero, one or more peering sessions are not yet established.
- Under **Peer**, the IP address for each configured peer is shown.
- Under **AS**, the peer AS for each configured peer is correct.
- Under **Up/Dwn State**, the BGP state reflects the number of paths received from the neighbor, the number of these paths that have been accepted, and the number of routes being dampened (such as 0/0/0). If the field is **Active**, it indicates a problem in the establishment of the BGP session.

Related Documentation

- [Understanding External BGP Peering Sessions on page 3601](#)
- [BGP Configuration Overview](#)

BGP Security Configuration

- [Example: Configuring BGP Route Authentication on page 3897](#)
- [Examples: Configuring TCP and BGP Security on page 3904](#)

Example: Configuring BGP Route Authentication

- [Understanding Route Authentication on page 3897](#)
- [Example: Configuring Route Authentication for BGP on page 3898](#)

Understanding Route Authentication

The use of router and route authentication and route integrity greatly mitigates the risk of being attacked by a machine or router that has been configured to share incorrect routing information with another router. In this kind of attack, the attacked router can be tricked into creating a routing loop, or the attacked router's routing table can be greatly increased thus impacting performance, or routing information can be redirected to a place in the network for the attacker to analyze it. Bogus route advertisements can be sent out on a segment. These updates can be accepted into the routing tables of neighbor routers unless an authentication mechanism is in place to verify the source of the routes.

Router and route authentication enables routers to share information only if they can verify that they are talking to a trusted source, based on a password (key). In this method, a hashed key is sent along with the route being sent to another router. The receiving router compares the sent key to its own configured key. If they are the same, it accepts the route. By using a hashing algorithm, the key is not sent over the wire in plain text. Instead, a hash is calculated using the configured key. The routing update is used as the input text, along with the key, into the hashing function. This hash is sent along with the route update to the receiving router. The receiving router compares the received hash with a hash it generates on the route update using the preshared key configured on it. If the two hashes are the same, the route is assumed to be from a trusted source. The key is known only to the sending and receiving routers.

To further strengthen security, you can configure a series of authentication keys (a *keychain*). Each key has a unique start time within the keychain. Keychain authentication allows you to change the password information periodically without bringing down peering sessions. This keychain authentication method is referred to as *hitless* because the keys roll over from one to the next without resetting any peering sessions or interrupting the routing protocol.

The sending peer uses the following rules to identify the active authentication key:

- The start time is less than or equal to the current time (in other words, not in the future).
- The start time is greater than that of all other keys in the chain whose start time is less than the current time (in other words, closest to the current time).

The receiving peer determines the key with which it authenticates based on the incoming key identifier.

The sending peer identifies the current authentication key based on a configured start time and then generates a hash value using the current key. The sending peer then inserts a TCP-enhanced authentication option object into the BGP update message. The object contains an object ID (assigned by IANA), the object length, the current key, and a hash value.

The receiving peer examines the incoming TCP-enhanced authentication option, looks up the received authentication key, and determines whether the key is acceptable based on the start time, the system time, and the tolerance parameter. If the key is accepted, the receiving peer calculates a hash and authenticates the update message.

Initial application of a keychain to a TCP session causes the session to reset. However, once the keychain is applied, the addition or removal of a password from the keychain does not cause the TCP session to reset. Also, the TCP session does not reset when the keychain changes from one authentication algorithm to another.

Example: Configuring Route Authentication for BGP

All BGP protocol exchanges can be authenticated to guarantee that only trusted routing devices participate in autonomous system (AS) routing updates. By default, authentication is disabled.

- [Requirements on page 3898](#)
- [Overview on page 3898](#)
- [Configuration on page 3900](#)
- [Verification on page 3902](#)

Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol (IGP).

Overview

When you configure authentication, the algorithm creates an encoded checksum that is included in the transmitted packet. The receiving routing device uses an authentication key (password) to verify the packet's checksum.

This example includes the following statements for configuring and applying the keychain:

- **key**—A keychain can have multiple keys. Each key within a keychain must be identified by a unique integer value. The range of valid identifier values is from 0 through 63.
The key can be up to 126 characters long. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").
- **tolerance**—(Optional) For each keychain, you can configure a clock-skew tolerance value in seconds. The clock-skew tolerance is applicable to the receiver accepting keys for BGP updates. The configurable range is 0 through 999,999,999 seconds. During the tolerance period, either the current or previous password is acceptable.
- **key-chain**—For each keychain, you must specify a name. This example defines one keychain: **bgp-auth**. You can have multiple keychains on a routing device. For example, you can have a keychain for BGP, a keychain for OSPF, and a keychain for LDP.
- **secret**—For each key in the keychain, you must set a secret password. This password can be entered in either encrypted or plain text format in the **secret** statement. It is always displayed in encrypted format.
- **start-time**—Each key must specify a start time in UTC format. Control gets passed from one key to the next. When a configured start time arrives (based on the routing device's clock), the key with that start time becomes active. Start times are specified in the local time zone for a routing device and must be unique within the keychain.
- **authentication-key-chain**—Enables you to apply a keychain at the global BGP level for all peers, for a group, or for a neighbor. This example applies the keychain to the peers defined in the external BGP (EBGP) group called **ext**.
- **authentication-algorithm**—For each keychain, you can specify a hashing algorithm. The algorithm can be AES-128, MD5, or SHA-1.


You associate a keychain and an authentication algorithm with a BGP neighboring session.

This example configures a keychain named **bgp-auth**. Key 0 will be sent and accepted starting at 2011-6-23.20:19:33 -0700, and will stop being sent and accepted when the next key in the keychain (key 1) becomes active. Key 1 becomes active one year later at 2012-6-23.20:19:33 -0700, and will not stop being sent and accepted unless another key is configured with a start time that is later than the start time of key 1. A clock-skew tolerance of 30 seconds applies to the receiver accepting the keys. During the tolerance period, either the current or previous key is acceptable. The keys are shared-secret passwords. This means that the neighbors receiving the authenticated routing updates must have the same authentication keychain configuration, including the same keys (passwords). So Router R0 and Router R1 must have the same authentication-key-chain configuration if they are configured as peers. This example shows the configuration on only one of the routing devices.

Topology Diagram

Figure 87 shows the topology used in this example.

Figure 87: Authentication for BGP



Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```

set protocols bgp group ext type external
set protocols bgp group ext peer-as 65530
set protocols bgp group ext neighbor 172.16.2.1
set routing-options autonomous-system 65533
set protocols bgp group ext authentication-key-chain bgp-auth
set protocols bgp group ext authentication-algorithm md5
set security authentication-key-chains key-chain bgp-auth tolerance 30
set security authentication-key-chains key-chain bgp-auth key 0 secret
  this-is-the-secret-password
set security authentication-key-chains key-chain bgp-auth key 0 start-time
  2011-6-23.20:19:33-0700
set security authentication-key-chains key-chain bgp-auth key 1 secret
  this-is-another-secret-password
set security authentication-key-chains key-chain bgp-auth key 1 start-time
  2012-6-23.20:19:33-0700
  
```

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router R1 to accept route filters from Device CE1 and perform outbound route filtering using the received filters:

1. Configure the local autonomous system.


```

[edit routing-options]
user@R1# set autonomous-system 65533
      
```
2. Configure one or more BGP groups.


```

[edit protocols bgp group ext]
user@R1# set type external
user@R1# set peer-as 65530
user@R1# set neighbor 172.16.2.1
      
```
3. Configure authentication with multiple keys.


```

[edit security authentication-key-chains key-chain bgp-auth]
user@R1# set key 0 secret this-is-the-secret-password
user@R1# set key 0 start-time 2011-6-23.20:19:33-0700
user@R1# set key 1 secret this-is-another-secret-password
user@R1# set key 1 start-time 2012-6-23.20:19:33-0700
      
```

The start time of each key must be unique within the keychain.

4. Apply the authentication keychain to BGP, and set the hashing algorithm.

```
[edit protocols bgp group ext]
user@R1# set authentication-key-chain bgp-auth
user@R1# set authentication-algorithm md5
```

5. (Optional) Apply a clock-skew tolerance value in seconds.

```
[edit security authentication-key-chains key-chain bgp-auth]
user@R1# set tolerance 30
```

Results From configuration mode, confirm your configuration by entering the **show protocols**, **show routing-options**, and **show security** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show protocols
bgp {
  group ext {
    type external;
    peer-as 65530;
    neighbor 172.16.2.1;
    authentication-key-chain bgp-auth;
    authentication-algorithm md5;
  }
}

user@R1# show routing-options
autonomous-system 65533;

user@R1# show security
authentication-key-chains {
  key-chain bgp-auth {
    tolerance 30;
    key 0 {
      secret
        "$9$5T6AREYk8RhXNdwaJn/CtO1cykWx9AyIMWdVgoJDjqP5FCA0z3IEhcmWLxNbgJDiF6A";
      ## SECRET-DATA
      start-time "2011-6-23.20:19:33 -0700";
    }
    key 1 {
      secret "$9$UyD.59CuO1h9AylKW-dqmfT369CuRhSP5hrvMN-JGDiqfu0lleWpuh.";
      ## SECRET-DATA
      start-time "2012-6-23.20:19:33 -0700";
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Repeat the procedure for every BGP-enabled device in the network, using the appropriate interface names and addresses for each BGP-enabled device.

Verification

Confirm that the configuration is working properly.

- [Verifying Authentication for the Neighbor on page 3902](#)
- [Verifying That Authorization Messages Are Sent on page 3902](#)
- [Checking Authentication Errors on page 3903](#)
- [Verifying the Operation of the Keychain on page 3903](#)

Verifying Authentication for the Neighbor

Purpose Make sure that the **AuthKeyChain** option appears in the output of the **show bgp neighbor** command.

Action From operational mode, enter the **show bgp neighbor** command.

```
user@R1> show bgp neighbor
Peer: 172.16.2.1+179 AS 65530 Local: 172.16.2.2+1222 AS 65533
  Type: External State: Established Flags: <Sync>
  Last State: OpenConfirm Last Event: RecvKeepAlive
  Last Error: None
  Export: [ direct-lo0 ]
  Options: <Preference PeerAS Refresh>
  Options: <AuthKeyChain>
  Authentication key is configured
  Authentication key chain: jni
  Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 172.16.2.1 Local ID: 10.255.124.35 Active Holdtime: 90
  Keepalive Interval: 30 Peer index: 0
  Local Interface: fe-0/0/1.0
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Table inet.0 Bit: 10000
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes: 2
    Received prefixes: 2
    Suppressed due to damping: 0
    Advertised prefixes: 1
  Last traffic (seconds): Received 2 Sent 2 Checked 2
  Input messages: Total 21 Updates 2 Refreshes 0 Octets 477
  Output messages: Total 22 Updates 1 Refreshes 0 Octets 471
  Output Queue[0]: 0
```

Verifying That Authorization Messages Are Sent

Purpose Confirm that BGP has the enhanced authorization option.

Action From operational mode, enter the **monitor traffic interface fe-0/0/1** command.

```
user@R1> monitor traffic interface fe-0/0/1
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Listening on fe-0/0/1, capture size 96 bytes
```

```
13:08:00.618402 In arp who-has 172.16.2.66 tell 172.16.2.69
```

```

13:08:02.408249 Out IP 172.16.2.2.1122 > 172.16.2.1.646: P
1889289217:1889289235(18) ack 2215740969 win 58486 <nop,nop,timestamp 167557
1465469,nop,Enhanced Auth keyid 0 diglen 12 digest: fe3366001f45767165f17037>:
13:08:02.418396 In IP 172.16.2.1.646 > 172.16.2.2.1122: P 1:19(18) ack 18 win
57100 <nop,nop,timestamp 1466460 167557,nop,Enhanced Auth keyid 0 diglen 12
digest: a18c31eda1b14b2900921675>:
13:08:02.518146 Out IP 172.16.2.2.1122 > 172.16.2.1.646: . ack 19 win 58468
<nop,nop,timestamp 167568 1466460,nop,Enhanced Auth keyid 0 diglen 12 digest:
c3b6422eb6bd3fd9cf79742b>
13:08:28.199557 Out IP 172.16.2.2.nerv > 172.16.2.1.bgp: P
286842489:286842508(19) ack 931203976 win 57200 <nop,Enhanced Auth keyid 0
diglen 12 digest: fc0e42900a73736bcc07c1a4>: BGP, length: 19
13:08:28.209661 In IP 172.16.2.1.bgp > 172.16.2.2.nerv: P 1:20(19) ack 19 win
56835 <nop,Enhanced Auth keyid 0 diglen 12 digest: 0fc8578c489fabce63aeb2c3>:
BGP, length: 19
13:08:28.309525 Out IP 172.16.2.2.nerv > 172.16.2.1.bgp: . ack 20 win 57181
<nop,Enhanced Auth keyid 0 diglen 12 digest: ef03f282fb2ece0039491df8>
13:08:32.439708 Out IP 172.16.2.2.1122 > 172.16.2.1.646: P 54:72(18) ack 55 win
58432 <nop,nop,timestamp 170560 1468472,nop,Enhanced Auth keyid 0 diglen 12
digest: 76e0cf926f348b726c631944>:
13:08:32.449795 In IP 172.16.2.1.646 > 172.16.2.2.1122: P 55:73(18) ack 72 win
57046 <nop,nop,timestamp 1469463 170560,nop,Enhanced Auth keyid 0 diglen 12
digest: dae3eec390d18a114431f4d8>:
13:08:32.549726 Out IP 172.16.2.2.1122 > 172.16.2.1.646: . ack 73 win 58414
<nop,nop,timestamp 170571 1469463,nop,Enhanced Auth keyid 0 diglen 12 digest:
851df771aee2ea7a43a0c46c>
13:08:33.719880 In arp who-has 172.16.2.66 tell 172.16.2.69
^C
35 packets received by filter
0 packets dropped by kernel

```

Checking Authentication Errors

Purpose Check the number of packets dropped by TCP because of authentication errors.

Action From operational mode, enter the **show system statistics tcp | match auth** command.

```

user@R1> show system statistics tcp | match auth
      0 send packets dropped by TCP due to auth errors
      58 rcv packets dropped by TCP due to auth errors

```

Verifying the Operation of the Keychain

Purpose Check the number of packets dropped by TCP because of authentication errors.

Action From operational mode, enter the **show security keychain detail** command.

```

user@R1> show security keychain detail
keychain          Active-ID      Next-ID      Transition  Tolerance
                  Send Receive  Send Receive
bgp-auth          3      3      1      1      1d 23:58      30
Id 3, Algorithm hmac-md5, State send-receive, Option basic
Start-time Wed Aug 11 16:28:00 2010, Mode send-receive
Id 1, Algorithm hmac-md5, State inactive, Option basic
Start-time Fri Aug 20 11:30:57 2010, Mode send-receive

```

Related Documentation

- [Understanding External BGP Peering Sessions on page 3601](#)
- [BGP Configuration Overview](#)

Examples: Configuring TCP and BGP Security

- [Understanding Security Options for BGP with TCP on page 3904](#)
- [Example: Configuring a Filter to Block TCP Access to a Port Except from Specified BGP Peers on page 3904](#)
- [Example: Configuring a Filter to Limit TCP Access to a Port Based On a Prefix List on page 3909](#)
- [Example: Limiting TCP Segment Size for BGP on page 3912](#)

Understanding Security Options for BGP with TCP

Among routing protocols, BGP is unique in using TCP as its transport protocol. BGP peers are established by manual configuration between routing devices to create a TCP session on port 179. A BGP-enabled device periodically sends keepalive messages to maintain the connection.

Over time, BGP has become the dominant interdomain routing protocol on the Internet. However, it has limited guarantees of stability and security. Configuring security options for BGP must balance suitable security measures with acceptable costs. No one method has emerged as superior to other methods. Each network administrator must configure security measures that meet the needs of the network being used.

For detailed information about the security issues associated with BGP's use of TCP as a transport protocol, see RFC 4272, *BGP Security Vulnerabilities Analysis*.

Example: Configuring a Filter to Block TCP Access to a Port Except from Specified BGP Peers

This example shows how to configure a standard stateless firewall filter that blocks all TCP connection attempts to port 179 from all requesters except from specified BGP peers.

- [Requirements on page 3904](#)
- [Overview on page 3904](#)
- [Configuration on page 3905](#)
- [Verification on page 3908](#)

Requirements

No special configuration beyond device initialization is required before you configure this example.


Overview

In this example, you create a stateless firewall filter that blocks all TCP connection attempts to port 179 from all requesters except the specified BGP peers.

The stateless firewall filter **filter_bgp179** matches all packets from the directly connected interfaces on Device A and Device B to the destination port number 179.

Figure 88 shows the topology used in this example. Device C attempts to make a TCP connection to Device E. Device E blocks the connection attempt. This example shows the configuration on Device E.

Figure 88: Typical Network with BGP Peer Sessions



Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device C

```
set interfaces ge-1/2/0 unit 10 description to-E
set interfaces ge-1/2/0 unit 10 family inet address 10.10.10.10/30
set protocols bgp group external-peers type external
set protocols bgp group external-peers peer-as 17
set protocols bgp group external-peers neighbor 10.10.10.9
set routing-options autonomous-system 22
```

Device E

```
set interfaces ge-1/2/0 unit 0 description to-A
set interfaces ge-1/2/0 unit 0 family inet address 10.10.10.1/30
set interfaces ge-1/2/1 unit 5 description to-B
set interfaces ge-1/2/1 unit 5 family inet address 10.10.10.5/30
set interfaces ge-1/0/0 unit 9 description to-C
set interfaces ge-1/0/0 unit 9 family inet address 10.10.10.9/30
set interfaces lo0 unit 2 family inet filter input filter_bgp179
set interfaces lo0 unit 2 family inet address 192.168.0.1/32
set protocols bgp group external-peers type external
set protocols bgp group external-peers peer-as 22
set protocols bgp group external-peers neighbor 10.10.10.2
set protocols bgp group external-peers neighbor 10.10.10.6
set protocols bgp group external-peers neighbor 10.10.10.10
set routing-options autonomous-system 17
set firewall family inet filter filter_bgp179 term 1 from source-address 10.10.10.2/32
set firewall family inet filter filter_bgp179 term 1 from source-address 10.10.10.6/32
set firewall family inet filter filter_bgp179 term 1 from destination-port bgp
set firewall family inet filter filter_bgp179 term 1 then accept
set firewall family inet filter filter_bgp179 term 2 then reject
```

Configuring Device E

Step-by-Step Procedure The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device E with a stateless firewall filter that blocks all TCP connection attempts to port 179 from all requestors except specified BGP peers:

1. Configure the interfaces.

```
user@E# set interfaces ge-1/2/0 unit 0 description to-A
user@E# set interfaces ge-1/2/0 unit 0 family inet address 10.10.10.1/30
```

```
user@E# set interfaces ge-1/2/1 unit 5 description to-B
user@E# set interfaces ge-1/2/1 unit 5 family inet address 10.10.10.5/30
```

```
user@E# set interfaces ge-1/0/0 unit 9 description to-C
user@E# set interfaces ge-1/0/0 unit 9 family inet address 10.10.10.9/30
```

2. Configure BGP.

```
[edit protocols bgp group external-peers]
user@E# set type external
user@E# set peer-as 22
user@E# set neighbor 10.10.10.2
user@E# set neighbor 10.10.10.6
user@E# set neighbor 10.10.10.10
```

3. Configure the autonomous system number.

```
[edit routing-options]
user@E# set autonomous-system 17
```

4. Define the filter term that accepts TCP connection attempts to port 179 from the specified BGP peers.

```
[edit firewall family inet filter filter_bgp179]
user@E# set term 1 from source-address 10.10.10.2/32
user@E# set term 1 from source-address 10.10.10.6/32
user@E# set term 1 from destination-port bgp
user@E# set term 1 then accept
```

5. Define the other filter term to reject packets from other sources.

```
[edit firewall family inet filter filter_bgp179]
user@E# set term 2 then reject
```

6. Apply the firewall filter to the loopback interface.

```
[edit interfaces lo0 unit 2 family inet]
user@E# set filter input filter_bgp179
user@E# set address 192.168.0.1/32
```

Results From configuration mode, confirm your configuration by entering the **show firewall**, **show interfaces**, **show protocols**, and **show routing-options** commands. If the output does not

display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@E# show firewall
family inet {
  filter filter_bgp179 {
    term 1 {
      from {
        source-address {
          10.10.10.2/32;
          10.10.10.6/32;
        }
        destination-port bgp;
      }
      then accept;
    }
    term 2 {
      then {
        reject;
      }
    }
  }
}

user@E# show interfaces
lo0 {
  unit 2 {
    family inet {
      filter {
        input filter_bgp179;
      }
      address 192.168.0.1/32;
    }
  }
}
ge-1/2/0 {
  unit 0 {
    description to-A;
    family inet {
      address 10.10.10.1/30;
    }
  }
}
ge-1/2/1 {
  unit 5 {
    description to-B;
    family inet {
      address 10.10.10.5/30;
    }
  }
}
ge-1/0/0 {
  unit 9 {
    description to-C;
    family inet {
      address 10.10.10.9/30;
    }
  }
}

```

```

    }
  }
}

user@E# show protocols
bgp {
  group external-peers {
    type external;
    peer-as 22;
    neighbor 10.10.10.2;
    neighbor 10.10.10.6;
    neighbor 10.10.10.10;
  }
}

user@E# show routing-options
autonomous-system 17;

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying That the Filter Is Configured on page 3908](#)
- [Verifying the TCP Connections on page 3908](#)
- [Monitoring Traffic on the Interfaces on page 3909](#)

Verifying That the Filter Is Configured

Purpose Make sure that the filter is listed in output of the **show firewall filter** command.

Action user@E> show firewall filter filter_bgp179
Filter: filter_bgp179

Verifying the TCP Connections

Purpose Verify the TCP connections.

Action From operational mode, run the **show system connections extensive** command on Device C and Device E.

The output on Device C shows the attempt to establish a TCP connection. The output on Device E shows that connections are established with Device A and Device B only.

user@C> show system connections extensive | match 10.10.10

```
tcp4      0      0  10.10.10.9.51872      10.10.10.10.179      SYN_SENT
```

user@E> show system connections extensive | match 10.10.10

```
tcp4      0      0  10.10.10.5.179        10.10.10.6.62096     ESTABLISHED
tcp4      0      0  10.10.10.6.62096      10.10.10.5.179       ESTABLISHED
tcp4      0      0  10.10.10.1.179        10.10.10.2.61506     ESTABLISHED
tcp4      0      0  10.10.10.2.61506      10.10.10.1.179       ESTABLISHED
```

Monitoring Traffic on the Interfaces

Purpose Use the **monitor traffic** command to compare the traffic on an interface that establishes a TCP connection with the traffic on an interface that does not establish a TCP connection.

Action From operational mode, run the **monitor traffic** command on the Device E interface to Device B and on the Device E interface to Device C. The following sample output verifies that in the first example, acknowledgment (**ack**) messages are received. In the second example, **ack** messages are not received.

```
user@E> monitor traffic size 1500 interface ge-1/2/1.5
19:02:49.700912 Out IP 10.10.10.5.bgp > 10.10.10.6.62096: P
3330573561:3330573580(19) ack 915601686 win 16384 <nop,nop,timestamp 1869518816
1869504850>: BGP, length: 19
19:02:49.801244 In IP 10.10.10.6.62096 > 10.10.10.5.bgp: . ack 19 win 16384
<nop,nop,timestamp 1869518916 1869518816>
19:03:03.323018 In IP 10.10.10.6.62096 > 10.10.10.5.bgp: P 1:20(19) ack 19 win
16384 <nop,nop,timestamp 1869532439 1869518816>: BGP, length: 19
19:03:03.422418 Out IP 10.10.10.5.bgp > 10.10.10.6.62096: . ack 20 win 16384
<nop,nop,timestamp 1869532539 1869532439>
19:03:17.220162 Out IP 10.10.10.5.bgp > 10.10.10.6.62096: P 19:38(19) ack 20 win
16384 <nop,nop,timestamp 1869546338 1869532439>: BGP, length: 19
19:03:17.320501 In IP 10.10.10.6.62096 > 10.10.10.5.bgp: . ack 38 win 16384
<nop,nop,timestamp 1869546438 1869546338>
```

```
user@E> monitor traffic size 1500 interface ge-1/0/0.9
```

```
18:54:20.175471 Out IP 10.10.10.9.61335 > 10.10.10.10.bgp: S 573929123:573929123(0)
win 16384 <mss 1460,nop,wscale 0,nop,nop,timestamp 1869009240 0,sackOK,eol>
18:54:23.174422 Out IP 10.10.10.9.61335 > 10.10.10.10.bgp: S 573929123:573929123(0)
win 16384 <mss 1460,nop,wscale 0,nop,nop,timestamp 1869012240 0,sackOK,eol>
18:54:26.374118 Out IP 10.10.10.9.61335 > 10.10.10.10.bgp: S 573929123:573929123(0)
win 16384 <mss 1460,nop,wscale 0,nop,nop,timestamp 1869015440 0,sackOK,eol>
18:54:29.573799 Out IP 10.10.10.9.61335 > 10.10.10.10.bgp: S 573929123:573929123(0)
win 16384 <mss 1460,sackOK,eol>
18:54:32.773493 Out IP 10.10.10.9.61335 > 10.10.10.10.bgp: S 573929123:573929123(0)
win 16384 <mss 1460,sackOK,eol>
18:54:35.973185 Out IP 10.10.10.9.61335 > 10.10.10.10.bgp: S 573929123:573929123(0)
win 16384 <mss 1460,sackOK,eol>
```

Example: Configuring a Filter to Limit TCP Access to a Port Based On a Prefix List

This example shows how to configure a standard stateless firewall filter that limits certain TCP and Internet Control Message Protocol (ICMP) traffic destined for the Routing Engine by specifying a list of prefix sources that contain allowed BGP peers.

- [Requirements on page 3909](#)
- [Overview on page 3910](#)
- [Configuration on page 3910](#)
- [Verification on page 3912](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

In this example, you create a stateless firewall filter that blocks all TCP connection attempts to port 179 from all requesters except BGP peers that have a specified prefix.

A source prefix list, **plist_bgp179**, is created that specifies the list of source prefixes that contain allowed BGP peers.

The stateless firewall filter **filter_bgp179** matches all packets from the source prefix list **plist_bgp179** to the destination port number 179.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set policy-options prefix-list plist_bgp179 apply-path "protocols bgp group <*> neighbor <*>"
set firewall family inet filter filter_bgp179 term 1 from source-address 0.0.0.0/0
set firewall family inet filter filter_bgp179 term 1 from source-prefix-list plist_bgp179 except
set firewall family inet filter filter_bgp179 term 1 from destination-port bgp
set firewall family inet filter filter_bgp179 term 1 then reject
set firewall family inet filter filter_bgp179 term 2 then accept
set interfaces lo0 unit 0 family inet filter input filter_bgp179
set interfaces lo0 unit 0 family inet address 127.0.0.1/32
```

Configure the Filter

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the filter:

1. Expand the prefix list **bgp179** to include all prefixes pointed to by the BGP peer group defined by **protocols bgp group <*> neighbor <*>**.

```
[edit policy-options prefix-list plist_bgp179]
user@host# set apply-path " protocolsbgp group <*> neighbor <*>"
```

2. Define the filter term that rejects TCP connection attempts to port 179 from all requesters except the specified BGP peers.

```
[edit firewall family inet filter filter_bgp179]
user@host# set term term1 from source-address 0.0.0.0/0
user@host# set term term1 from source-prefix-list bgp179 except
user@host# set term term1 from destination-port bgp
user@host# set term term1 then reject
```

3. Define the other filter term to accept all packets.

```
[edit firewall family inet filter filter_bgp179]
user@host# set term term2 then accept
```

4. Apply the firewall filter to the loopback interface.

```
[edit interfaces lo0 unit 0 family inet]
user@host# set filter input filter_bgp179
user@host# set address 127.0.0.1/32
```

Results

From configuration mode, confirm your configuration by entering the **show firewall**, **show interfaces**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show firewall
family inet {
  filter filter_bgp179 {
    term 1 {
      from {
        source-address {
          0.0.0.0/0;
        }
        source-prefix-list {
          plist_bgp179 except;
        }
        destination-port bgp;
      }
      then {
        reject;
      }
    }
    term 2 {
      then {
        accept;
      }
    }
  }
}

user@host# show interfaces
lo0 {
  unit 0 {
    family inet {
      filter {
        input filter_bgp179;
      }
      address 127.0.0.1/32;
    }
  }
}

user@host# show policy-options
prefix-list plist_bgp179 {
  apply-path "protocols bgp group <*> neighbor <*>";
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

Displaying the Firewall Filter Applied to the Loopback Interface

Purpose Verify that the firewall filter **filter_bgp179** is applied to the IPv4 input traffic at logical interface **lo0.0**.

Action Use the **show interfaces statistics operational mode** command for logical interface **lo0.0**, and include the **detail** option. Under the **Protocol inet** section of the command output section, the **Input Filters** field displays the name of the stateless firewall filter applied to the logical interface in the input direction.

[edit]

```
user@host> show interfaces statistics lo0.0 detail
Logical interface lo0.0 (Index 321) (SNMP ifIndex 16) (Generation 130)
Flags: SNMP-Traps Encapsulation: Unspecified
Traffic statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Local statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Transit statistics:
  Input bytes : 0 0 bps
  Output bytes : 0 0 bps
  Input packets: 0 0 pps
  Output packets: 0 0 pps
Protocol inet, MTU: Unlimited, Generation: 145, Route table: 0
Flags: Sendbroadcast-pkt-to-re
Input Filters: filter_bgp179
Addresses, Flags: Primary
  Destination: Unspecified, Local: 127.0.0.1, Broadcast: Unspecified,
Generation: 138
```

Example: Limiting TCP Segment Size for BGP

This example shows how to avoid Internet Control Message Protocol (ICMP) vulnerability issues by limiting TCP segment size when you are using maximum transmission unit (MTU) discovery. Using MTU discovery on TCP paths is one method of avoiding BGP packet fragmentation.

- [Requirements on page 3913](#)
- [Overview on page 3913](#)
- [Configuration on page 3913](#)
- [Verification on page 3915](#)
- [Troubleshooting on page 3916](#)

Requirements

No special configuration beyond device initialization is required before you configure this example.

Overview

TCP negotiates a maximum segment size (MSS) value during session connection establishment between two peers. The MSS value negotiated is primarily based on the maximum transmission unit (MTU) of the interfaces to which the communicating peers are directly connected. However, due to variations in link MTU on the path taken by the TCP packets, some packets in the network that are well within the MSS value might be fragmented when the packet size exceeds the link's MTU.

To configure the TCP MSS value, include the `tcp-mss` statement with a segment size from 1 through 4096.

If the router receives a TCP packet with the SYN bit and the MSS option set, and the MSS option specified in the packet is larger than the MSS value specified by the `tcp-mss` statement, the router replaces the MSS value in the packet with the lower value specified by the `tcp-mss` statement.


The configured MSS value is used as the maximum segment size for the sender. The assumption is that the TCP MSS value used by the sender to communicate with the BGP neighbor is the same as the TCP MSS value that the sender can accept from the BGP neighbor. If the MSS value from the BGP neighbor is less than the MSS value configured, the MSS value from the BGP neighbor is used as the maximum segment size for the sender.

This feature is supported with TCP over IPv4 and TCP over IPv6.

Topology Diagram

Figure 89 shows the topology used in this example.

Figure 89: TCP Maximum Segment Size for BGP



Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the `[edit]` hierarchy level.

```

R0    set interfaces fe-1/2/0 unit 1 family inet address 1.1.0.1/30
      set interfaces lo0 unit 1 family inet address 10.255.14.179/32
      set protocols bgp group-int tcp-mss 2020
      set protocols bgp group int type internal
  
```

```
set protocols bgp group int local-address 10.255.14.179
set protocols bgp group int mtu-discovery
set protocols bgp group int neighbor 10.255.71.24 tcp-mss 2000
set protocols bgp group int neighbor 10.255.14.177
set protocols bgp group int neighbor 10.0.14.4 tcp-mss 4000
set protocols ospf area 0.0.0.0 interface fe-1/2/0.1
set protocols ospf area 0.0.0.0 interface 10.255.14.179
set routing-options autonomous-system 65000
```

Step-by-Step Procedure The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router R0:

1. Configure the interfaces.

```
[edit interfaces]
user@R0# set fe-1/2/0 unit 1 family inet address 1.1.0.1/30
user@R0# set lo0 unit 1 family inet address 10.255.14.179/32
```

2. Configure an interior gateway protocol (IGP), OSPF in this example.

```
[edit protocols ospf area 0.0.0.0]
user@R0# set interface fe-1/2/0.1
user@R0# set interface 10.255.14.179
```

3. Configure one or more BGP groups.

```
[edit protocols bgp group int]
user@R0# set type internal
user@R0# set local-address 10.255.14.179
```

4. Configure MTU discovery to prevent packet fragmentation.

```
[edit protocols bgp group int]
user@R0# set mtu-discovery
```

5. Configure the BGP neighbors, with the TCP MSS set globally for the group or specifically for the various neighbors.

```
[edit protocols bgo group int]
user@R0# set tcp-mss 2020
user@R0# set neighbor 10.255.14.177
user@R0# set neighbor 10.255.71.24 tcp-mss 2000
user@R0# set neighbor 10.0.14.4 tcp-mss 4000
```



NOTE: The TCP MSS neighbor setting overrides the group setting.

6. Configure the local autonomous system.

```
[edit routing-options]
user@R0# set autonomous-system 65000
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@R0# show interfaces
fe-1/2/0 {
  unit 1 {
    family inet {
      address 1.1.0.1/30;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 10.255.14.179/32;
    }
  }
}

user@R0# show protocols
bgp {
  group int {
    type internal;
    local-address 10.255.14.179;
    mtu-discovery;
    tcp-mss 2020;
    neighbor 10.255.71.24 {
      tcp-mss 2000;
    }
    neighbor 10.255.14.177;
    neighbor 10.0.14.4 {
      tcp-mss 4000;
    }
  }
}
ospf {
  area 0.0.0.0 {
    interface fe-1/2/0.1;
    interface 10.255.14.179;
  }
}

user@R0# show routing-options
autonomous-system 65000;

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, run the following commands:

- **show system connections extensive | find <neighbor-address>**, to check the negotiated TCP MSS value.

- **monitor traffic interface**, to monitor BGP traffic and to make sure that the configured TCP MSS value is used as the MSS option in the TCP SYN packet.

Troubleshooting

- [MSS Calculation with MTU Discovery on page 3916](#)

MSS Calculation with MTU Discovery

Problem Consider an example in which two routing devices (R1 and R2) have an internal BGP (IBGP) connection. On both of the routers, the connected interfaces have 4034 as the IPv4 MTU.

```
user@R1# show protocols bgp | display set
[edit]
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 45.45.45.2
set protocols bgp group ibgp mtu-discovery
set protocols bgp group ibgp neighbor 45.45.45.1
```

```
user@R1# run show interfaces xe-0/0/3 extensive | match mtu
```

```
Link-level type: Ethernet, MTU: 4048, LAN-PHY mode, Speed: 10Gbps,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Protocol inet, MTU: 4034, Generation: 180, Route table: 0
Protocol multiservice, MTU: Unlimited, Generation: 181, Route table: 0
```

In the following packet capture on Device R1, the negotiated MSS is 3994. In the **show system connections extensive** information for MSS, it is set to 2048.

```
05:50:01.575218 Out
  Juniper PCAP Flags [Ext], PCAP Extension(s) total length 16
    Device Media Type Extension TLV #3, length 1, value: Ethernet (1)
    Logical Interface Encapsulation Extension TLV #6, length 1, value:
Ethernet (14)
    Device Interface Index Extension TLV #1, length 2, value: 137
    Logical Interface Index Extension TLV #4, length 4, value: 69
    -----original packet-----
    00:21:59:e1:e8:03 > 00:19:e2:20:79:01, ethertype IPv4 (0x0800), length
78: (tos 0xc0, ttl 64, id 53193, offset 0, flags [DF], proto: TCP (6), length:
64) 45.45.45.2.62840 > 45.45.45.1.bgp: S 2939345813:2939345813(0) win 16384 **mss
3994,nop,wscale 0,nop,nop,timestamp 70559970 0,sackOK,eol>
05:50:01.575875 In
  Juniper PCAP Flags [Ext, no-L2, In], PCAP Extension(s) total length 16
    Device Media Type Extension TLV #3, length 1, value: Ethernet (1)
    Logical Interface Encapsulation Extension TLV #6, length 1, value:
Ethernet (14)
    Device Interface Index Extension TLV #1, length 2, value: 137
    Logical Interface Index Extension TLV #4, length 4, value: 69
    -----original packet-----
    PFE proto 2 (ipv4): (tos 0xc0, ttl 255, id 37709, offset 0, flags [DF], proto:
TCP (6), length: 64) 45.45.45.1.bgp > 45.45.45.2.62840: S 2634967984:2634967984(0)
ack 2939345814 win 16384 **mss 3994,nop,wscale 0,nop,nop,timestamp 174167273
70559970,sackOK,eol>
```

```
user@R1# run show system connections extensive | find 45.45
```

```
tcp4      0      0  45.45.45.2.62840          45.45.45.1.179
                        ESTABLISHED
```

```

sndsbcc:          0 sndsbmbcnt:          0 sndsbmbmax:      131072
sndsblowat:       2048 sndsbhiwat:       16384
rcvsbcc:          0 rcvsbmbcnt:          0 rcvsbmbmax:      131072
rcvsblowat:       1 rcvsbhiwat:       16384
proc id:         19725 proc name:         rpd
  iss: 2939345813   sndup: 2939345972
snduna: 2939345991   sndnxt: 2939345991   sndwnd:      16384
sndmax: 2939345991   sndcwnd:      10240 sndssthresh: 1073725440
  irs: 2634967984   rcvup: 2634968162
rcvnxt: 2634968162   rcvadv: 2634984546   rcvwnd:      16384
  rtt:      0       srtt:      1538       rttv:      1040
rxtcur:      1200   rxtshift:      0       rtseq: 2939345972
rttmin:      1000   mss:          2048

```

Solution This is expected behavior with Junos OS. The MSS value is equal to the MTU value minus the IP or IPv6 and TCP headers. This means that the MSS value is generally 40 bytes less than the MTU (for IPv4) and 60 bytes less than the MTU (for IPv6). This value is negotiated between the peers. In this example, it is $4034 - 40 = 3994$. Junos OS then rounds this value to a multiple of 2 KB. The value is $3994 / 2048 * 2048 = 2048$. So it is not necessary to see same MSS value with in the **show system connections** output.

$3994 / 2048 = 1.95$

1.95 is rounded to 1.

$1 * 2048 = 2048$

- Related Documentation**
- [Understanding External BGP Peering Sessions on page 3601](#)
 - [BGP Configuration Overview](#)

BGP Flap Configuration

- [Example: Preventing BGP Session Resets on page 3919](#)
- [Examples: Configuring BGP Flap Damping on page 3926](#)

Example: Preventing BGP Session Resets

- [Understanding BGP Session Resets on page 3919](#)
- [Example: Preventing BGP Session Flaps When VPN Families Are Configured on page 3919](#)

Understanding BGP Session Resets

Certain configuration actions and events cause BGP sessions to be reset (dropped and then reestablished).

If you configure both route reflection and VPNs on the same routing device, the following modifications to the route reflection configuration cause current BGP sessions to be reset:

- Adding a cluster ID—If a BGP session shares the same autonomous system (AS) number with the group where you add the cluster ID, all BGP sessions are reset regardless of whether the BGP sessions are contained in the same group.
- Creating a new route reflector—If you have an internal BGP (IBGP) group with an AS number and create a new route reflector group with the same AS number, all BGP sessions in the IBGP group and the new route reflector group are reset.
- Changing configuration statements that affect BGP peers, such as renaming a BGP group, resets the BGP sessions.
- If you change the address family specified in the **[edit protocols bgp family]** hierarchy level, all current BGP sessions on the routing device are dropped and then reestablished.

Example: Preventing BGP Session Flaps When VPN Families Are Configured

This example shows a workaround for a known issue in which BGP sessions sometimes go down and then come back up (in other words, flap) when virtual private network (VPN) families are configured. If any VPN family (for example, **inet-vpn**, **inet6-vpn**, **inet-mpvn**, **inet-mdt**, **inet6-mpvn**, **l2vpn**, **iso-vpn**, and so on) is configured on a BGP master instance, a flap of either a route reflector (RR) internal BGP (IBGP) session or an external

BGP (EBGP) session causes flaps of other BGP sessions configured with the same VPN family.

- [Requirements on page 3920](#)
- [Overview on page 3921](#)
- [Configuration on page 3922](#)
- [Verification on page 3925](#)

Requirements

Before you begin:

- Configure router interfaces.
- Configure an interior gateway protocol (IGP).
- Configure BGP.
- Configure VPNs.

Overview

When a router or switch is configured as either a route reflector (RR) or an AS boundary router (an external BGP peer) and a VPN family (for example, the **family inet-vpn unicast** statement) is configured, a flap of either the RR IBGP session or the EBGP session causes flaps of all other BGP sessions that are configured with the **family inet-vpn unicast** statement. This example shows how to prevent these unnecessary session flaps.

The reason for the flapping behavior is related to BGP operation in Junos OS when originating VPN routes.

BGP has the following two modes of operation with respect to originating VPN routes:

- If BGP does not need to propagate VPN routes because the session has no EBGP peer and no RR clients, BGP exports VPN routes directly from the *instance.inet.0* routing table to other PE routers. This behavior is efficient in that it avoids the creation of two copies of many routes (one in the *instance.inet.0* table and one in the *bgp.l3vpn.0* table).
- If BGP does need to propagate VPN routes because the session has an EBGP peer or RR clients, BGP first exports the VPN routes from the *instance.inet.0* table to the *bgp.l3vpn.0* table. Then BGP exports the routes to other PE routers. In this scenario, two copies of the route are needed to enable best-route selection. A PE router might receive the same VPN route from a CE device and also from an RR client or EBGP peer.



NOTE: The route export is not performed if the route in *instance.inet.0* is a secondary route. In Junos OS, a route is only exported one time from one routing table as a primary route to another routing table as a secondary route. Because the route in *instance.inet.0* is already a secondary route, it is not allowed to be moved again to the *bgp.l3vpn.0* table, as needed to be advertised. The route does not reach the *bgp.l3vpn.0* table and thus is not advertised. One workaround is to send the routes that should be advertised to *inet.0* so that they are advertised.

When, because of a configuration change, BGP transitions from needing two copies of a route to not needing two copies of a route (or the reverse), all sessions over which VPN routes are exchanged go down and then come back up. Although this example focuses on the **family inet-vpn unicast** statement, the concept applies to all VPN network layer reachability information (NLRI) families. This issue impacts logical systems as well. All BGP sessions in the master instance related to the VPN NLRI family are brought down to implement the table advertisement change for the VPN NLRI family. Changing an RR to a non-RR or the reverse (by adding or removing the **cluster** statement) causes the table advertisement change. Also, configuring the first EBGP session or removing the EBGP session from the configuration in the master instance for a VPN NLRI family causes the table advertisement change.

The way to prevent these unnecessary session flaps is to configure an extra RR client or EBGP session as a passive session with a neighbor address that does not exist. This example focuses on the EBGP case, but the same workaround works for the RR case.

When a session is passive, the routing device does not send Open requests to a peer. Once you configure the routing device to be passive, the routing device does not originate the TCP connection. However, when the routing device receives a connection from the peer and an Open message, it replies with another BGP Open message. Each routing device declares its own capabilities.

Figure 90 shows the topology for the EBGp case. Router R1 has an IBGP session with Routers R2 and R3 and an EBGp session with Router R4. All sessions have the **family inet-vpn unicast** statement configured. If the R1-R4 EBGp session flaps, the R1-R2 and R1-R3 BGP sessions flap also.

Figure 90: Topology for the EBGp Case





Figure 91 shows the topology for the RR case. Router R1 is the RR, and Router R3 is the client. Router R1 has IBGP sessions with Routers R2 and R3. All sessions have the **family inet-vpn unicast** statement configured. If the R1-R3 session flaps, the R1-R2 and R1-R4 sessions flap also.

Figure 91: Topology for the RR Case



Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network

configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set protocols bgp family inet-vpn unicast
set protocols bgp family l2vpn signaling
set protocols bgp group R1-R4 type external
set protocols bgp group R1-R4 local-address 4.4.4.2
set protocols bgp group R1-R4 neighbor 4.4.4.1 peer-as 200
set protocols bgp group R1-R2-R3 type internal
set protocols bgp group R1-R2-R3 log-updown
set protocols bgp group R1-R2-R3 local-address 15.15.15.15
set protocols bgp group R1-R2-R3 neighbor 12.12.12.12
set protocols bgp group R1-R2-R3 neighbor 13.13.13.13
set protocols bgp group Fake type external
set protocols bgp group Fake passive
set protocols bgp group Fake neighbor 100.100.100.100 peer-as 500
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the EBGp scenario:

1. Configure one or more VPN families.

```
[edit protocols bgp]
user@R1# set family inet-vpn unicast
user@R1# set family l2vpn signaling
```

2. Configure the EBGp session.

```
[edit protocols bgp]
user@R1# set group R1-R4 type external
user@R1# set group R1-R4 local-address 4.4.4.2
user@R1# set group R1-R4 neighbor 4.4.4.1 peer-as 200
```

3. Configure the IBGP sessions.

```
[edit protocols bgp]
user@R1# set group R1-R2-R3 type internal
user@R1# set group R1-R2-R3 local-address 15.15.15.15
user@R1# set group R1-R2-R3 neighbor 12.12.12.12
user@R1# set group R1-R2-R3 neighbor 13.13.13.13
```

4. (Optional) Configure BGP so that it generates a **syslog** message whenever a BGP peer makes a state transition.

```
[edit protocols bgp]
user@R1# set group R1-R2-R3 log-updown
```

Enabling the **log-updown** statement causes BGP state transitions to be logged at **warning** level.

Step-by-Step Procedure To verify that unnecessary session flaps are occurring:

1. Run the **show bgp summary** command to verify that the sessions have been established.

```
user@R1> show bgp summary
```

```
Groups: 2 Peers: 3 Down peers: 0
```

Table	Tot	Paths	Act	Paths	Suppressed	History	Damp	State	Pending
bgp.13vpn.0	0		0		0	0	0		0
bgp.12vpn.0	0		0		0	0	0		0
inet.0	0		0		0	0	0		0

Peer	AS	InPkt	OutPkt	OutQ	Flaps	Last	Up/Dwn
4.4.4.1	200	6	5	0	0	1:08	Establ
bgp.13vpn.0: 0/0/0/0							
bgp.12vpn.0: 0/0/0/0							
12.12.12.12	100	3	7	0	0	1:18	Establ
bgp.13vpn.0: 0/0/0/0							
bgp.12vpn.0: 0/0/0/0							
13.13.13.13	100	3	6	0	0	1:14	Establ
bgp.13vpn.0: 0/0/0/0							
bgp.12vpn.0: 0/0/0/0							

- Deactivate the EBGp session.

```
user@R1# deactivate group R1-R4
```

```
user@R1# commit
```

```
Mar 10 18:27:40 R1: rpd[1464]: bgp_peer_delete:6589: NOTIFICATION sent to 4.4.4.1 (External AS 200): code 6 (Cease) subcode 3 (Peer Unconfigured), Reason: Peer Deletion
```

```
Mar 10 18:27:40 R1: rpd[1464]: bgp_adv_main_update:7253: NOTIFICATION sent to 12.12.12.12 (Internal AS 100): code 6 (Cease) subcode 6 (Other Configuration Change), Reason: Configuration change - VPN table advertise
```

```
Mar 10 18:27:40 R1: rpd[1464]: bgp_adv_main_update:7253: NOTIFICATION sent to 13.13.13.13 (Internal AS 100): code 6 (Cease) subcode 6 (Other Configuration Change), Reason: Configuration change - VPN table advertise
```

- Run the **show bgp summary** command to view the session flaps.

```
user@R1> show bgp summary
```

```
Groups: 1 Peers: 2 Down peers: 2
```

Table	Tot	Paths	Act	Paths	Suppressed	History	Damp	State	Pending
bgp.13vpn.0	0		0		0	0	0		0
bgp.12vpn.0	0		0		0	0	0		0
inet.0	0		0		0	0	0		0

Peer	AS	InPkt	OutPkt	OutQ	Flaps	Last	Up/Dwn
12.12.12.12	100	4	9	0	1	19	Active
13.13.13.13	100	4	8	0	1	19	Active

```
user@R1> show bgp summary
```

```
Groups: 1 Peers: 2 Down peers: 0
```

Table	Tot	Paths	Act	Paths	Suppressed	History	Damp	State	Pending
bgp.13vpn.0	0		0		0	0	0		0
bgp.12vpn.0	0		0		0	0	0		0
inet.0	0		0		0	0	0		0

Peer	AS	InPkt	OutPkt	OutQ	Flaps	Last	Up/Dwn
12.12.12.12	100	2	3	0	1	0	Establ
bgp.13vpn.0: 0/0/0/0							
bgp.12vpn.0: 0/0/0/0							
13.13.13.13	100	2	3	0	1	0	Establ
bgp.13vpn.0: 0/0/0/0							
bgp.12vpn.0: 0/0/0/0							

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To prevent unnecessary BGP session flaps:

1. Add a passive EBGP session with a neighbor address that does not exist in the peer autonomous system (AS).

```
[edit protocols bgp]
user@R1# set group Fake type external
user@R1# set group Fake passive
user@R1# set neighbor 100.100.100.100 peer-as 500
```

2. Run the **show bgp summary** command to verify that the real sessions have been established and the passive session is idle.

```
user@R1> show bgp summary
Groups: 3 Peers: 4 Down peers: 1
Table Tot Paths Act Paths Suppressed History Damp State Pending
bgp.13vpn.0 0 0 0 0 0 0
bgp.12vpn.0 0 0 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
4.4.4.1 200 9500 9439 0 0 2d 23:14:23 Estab1
bgp.13vpn.0: 0/0/0/0
bgp.12vpn.0: 0/0/0/0
12.12.12.12 100 10309 10239 0 0 3d 5:17:49 Estab1
bgp.13vpn.0: 0/0/0/0
13.13.13.13 100 10306 10241 0 0 3d 5:18:25 Estab1
bgp.13vpn.0: 0/0/0/0
100.100.100.100 500 0 0 0 0 2d 23:38:52 Idle
```

Verification

Confirm that the configuration is working properly.

- [Bringing Down the EBGp Session on page 3925](#)
- [Verifying That the IBGP Sessions Remain Up on page 3925](#)

Bringing Down the EBGp Session

Purpose Try to cause the flap issue after the workaround is configured.

Action user@R1# deactivate group R1-R4
user@R1# commit

Verifying That the IBGP Sessions Remain Up

Purpose Make sure that the IBGP sessions do not flap after the EBGp session is deactivated.

Action user@R1> show bgp summary

```

Groups: 2 Peers: 3 Down peers: 1
Table      Tot Paths Act Paths Suppressed History Damp State Pending
bgp.13vpn.0 0      0      0      0      0      0      0
bgp.12vpn.0 0      0      0      0      0      0      0
Peer      AS  InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
12.12.12.12 100 10312 10242 0 0 3d 5:19:01 Establ
bgp.13vpn.0: 0/0/0/0
13.13.13.13 100 10309 10244 0 0 3d 5:19:37 Establ
bgp.13vpn.0: 0/0/0/0
100.100.100.100 500 0 0 0 2d 23:40:04 Idle

```

user@R1> show bgp summary

```

Groups: 3 Peers: 4 Down peers: 1
Table      Tot Paths Act Paths Suppressed History Damp State Pending
bgp.13vpn.0 0      0      0      0      0      0      0
bgp.12vpn.0 0      0      0      0      0      0      0
Peer      AS  InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
4.4.4.1    200 5 4 0 0 28      Establ
bgp.13vpn.0: 0/0/0/0
bgp.12vpn.0: 0/0/0/0
12.12.12.12 100 10314 10244 0 0 3d 5:19:55 Establ
bgp.13vpn.0: 0/0/0/0
13.13.13.13 100 10311 10246 0 0 3d 5:20:31 Establ
bgp.13vpn.0: 0/0/0/0
100.100.100.100 500 0 0 0 2d 23:40:58 Idle

```

- Related Documentation**
- [Understanding External BGP Peering Sessions on page 3601](#)
 - [BGP Configuration Overview](#)

Examples: Configuring BGP Flap Damping

- [Understanding Damping Parameters on page 3926](#)
- [Example: Configuring BGP Route Flap Damping Parameters on page 3927](#)
- [Example: Configuring BGP Route Flap Damping Based on the MBGP MVPN Address Family on page 3936](#)

Understanding Damping Parameters

BGP *route flapping* describes the situation in which BGP systems send an excessive number of update messages to advertise network reachability information. BGP *flap damping* is a method of reducing the number of update messages sent between BGP peers, thereby reducing the load on these peers, without adversely affecting the route convergence time for stable routes.

Flap damping reduces the number of update messages by marking routes as ineligible for selection as the active or preferable route. Marking routes in this way leads to some delay, or *suppression*, in the propagation of route information, but the result is increased network stability. You typically apply flap damping to external BGP (EBGP) routes (routes in different ASs). You can also apply flap damping within a confederation, between

confederation member ASs. Because routing consistency within an AS is important, do not apply flap damping to internal BGP (IBGP) routes. (If you do, it is ignored.) The exception to this rule is when flap damping is applied at the address family level, which is supported in Junos OS Release 12.2 and later. When you apply flap damping at the address family level, it works for both IBGP and EBGP.

By default, route flap damping is not enabled. Damping is applied to external peers and to peers at confederation boundaries.

When you enable damping, default parameters are applied, as summarized in [Table 319](#).

Table 319: Damping Parameters

Damping Parameter	Description	Default Value	Possible Values
half-life <i>minutes</i>	Decay half-life—Number of minutes after which an arbitrary value is halved if a route stays stable.	15 (minutes)	1 through 45
max-suppress <i>minutes</i>	Maximum hold-down time for a route, in minutes.	60 (minutes)	1 through 720
reuse	Reuse threshold—Arbitrary value below which a suppressed route can be used again.	750	1 through 20,000
suppress	Cutoff (suppression) threshold—Arbitrary value above which a route can no longer be used or included in advertisements.	3000	1 through 20,000

To change the default BGP flap damping values, you define actions by creating a named set of damping parameters and including it in a routing policy with the damping action. For the damping routing policy to work, you also must enable BGP route flap damping.

Example: Configuring BGP Route Flap Damping Parameters

This example shows how to configure damping parameters.

- [Requirements on page 3927](#)
- [Overview on page 3927](#)
- [Configuration on page 3928](#)
- [Verification on page 3932](#)

Requirements

Before you begin, configure router interfaces and configure routing protocols.

Overview

This example has three routing devices. Device R2 has external BGP (EBGP) connections with Device R1 and Device R3.

Device R1 and Device R3 have some static routes configured for testing purposes, and these static routes are advertised through BGP to Device R2.


Device R2 damps routes received from Device R1 and Device R3 according to these criteria:

- Damp all prefixes with a mask length equal to or greater than 17 more aggressively than routes with a mask length between 9 and 16.
- Damp routes with a mask length between 0 and 8, inclusive, less than routes with a mask length greater than 8.
- Do not damp the 10.128.0.0/9 prefix at all.

The routing policy is evaluated when routes are being exported from the routing table into the forwarding table. Only the active routes are exported from the routing table.

Figure 92 shows the sample network.

Figure 92: BGP Flap Damping Topology



“CLI Quick Configuration” on page 3928 shows the configuration for all of the devices in Figure 92.

The section “Step-by-Step Procedure” on page 3929 describes the steps on Device R2.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device R1

```

set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.1/30
set interfaces lo0 unit 0 family inet address 192.168.0.1/32
set protocols bgp group ext type external
set protocols bgp group ext export send-direct-and-static
set protocols bgp group ext peer-as 200
set protocols bgp group ext neighbor 10.0.0.2
set policy-options policy-statement send-direct-and-static term 1 from protocol direct
set policy-options policy-statement send-direct-and-static term 1 from protocol static
set policy-options policy-statement send-direct-and-static term 1 then accept
set routing-options static route 172.16.0.0/16 reject
set routing-options static route 172.16.128.0/17 reject
set routing-options static route 172.16.192.0/20 reject
set routing-options static route 10.0.0.0/9 reject
set routing-options static route 224.0.0.0/7 reject
set routing-options static route 10.224.0.0/11 reject
set routing-options static route 0.0.0.0/0 reject
set routing-options autonomous-system 100
  
```

Device R2

```

set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 0 family inet address 10.1.0.1/30
set interfaces lo0 unit 0 family inet address 192.168.0.2/32
  
```



```

set protocols bgp damping
set protocols bgp group ext type external
set protocols bgp group ext import damp
set protocols bgp group ext export send-direct
set protocols bgp group ext neighbor 10.0.0.1 peer-as 100
set protocols bgp group ext neighbor 10.1.0.2 peer-as 300
set policy-options policy-statement damp term 1 from route-filter 10.128.0.0/9 exact
  damping dry
set policy-options policy-statement damp term 1 from route-filter 0.0.0.0/0
  prefix-length-range /0-/8 damping timid
set policy-options policy-statement damp term 1 from route-filter 0.0.0.0/0
  prefix-length-range /17-/32 damping aggressive
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set policy-options damping aggressive half-life 30
set policy-options damping aggressive suppress 2500
set policy-options damping timid half-life 5
set policy-options damping dry disable
set routing-options autonomous-system 200

```

Device R3

```

set interfaces fe-1/2/1 unit 0 family inet address 10.1.0.2/30
set interfaces lo0 unit 0 family inet address 192.168.0.3/32
set protocols bgp group ext type external
set protocols bgp group ext export send-direct-and-static
set protocols bgp group ext peer-as 200
set protocols bgp group ext neighbor 10.1.0.1
set policy-options policy-statement send-direct-and-static term 1 from protocol direct
set policy-options policy-statement send-direct-and-static term 1 from protocol static
set policy-options policy-statement send-direct-and-static term 1 then accept
set routing-options static route 10.128.0.0/9 reject
set routing-options autonomous-system 300

```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure damping parameters:

1. Configure the interfaces.

```

[edit interfaces]
user@R2# set fe-1/2/0 unit 0 family inet address 10.0.0.2/30

user@R2# set fe-1/2/1 unit 0 family inet address 10.1.0.1/30

user@R2# set lo0 unit 0 family inet address 192.168.0.2/32

```

2. Configure the BGP neighbors.

```

[edit protocols bgp group ext]
user@R2# set type external
user@R2# set neighbor 10.0.0.1 peer-as 100
user@R2# set neighbor 10.1.0.2 peer-as 300

```

3. Create and configure the damping parameter groups.

```
[edit policy-options]
user@R2# set damping aggressive half-life 30
user@R2# set damping aggressive suppress 2500
user@R2# set damping timid half-life 5
user@R2# set damping dry disable
```

4. Configure the damping policy.

```
[edit policy-options policy-statement damp term 1]
user@R2# set from route-filter 10.128.0.0/9 exact damping dry
user@R2# set from route-filter 0.0.0.0/0 prefix-length-range /0-/8 damping timid
user@R2# set from route-filter 0.0.0.0/0 prefix-length-range /17-/32 damping
    aggressive
```

5. Enable damping for BGP.

```
[edit protocols bgp]
user@R2# set damping
```

6. Apply the policy as an import policy for the BGP neighbor.

```
[edit protocols bgp group ext]
user@R2# set import damp
```



NOTE: You can refer to the same routing policy one or more times in the same or different import statements.

7. Configure an export policy.

```
[edit policy-options policy-statement send-direct term 1]
user@R2# set from protocol direct
user@R2# set then accept
```

8. Apply the export policy.

```
[edit protocols bgp group ext]
user@R2# set export send-direct
```

9. Configure the autonomous system (AS) number.

```
[edit routing-options]
user@R2# set autonomous-system 200
```

Results From configuration mode, confirm your configuration by issuing the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R2# show interfaces
fe-1/2/0 {
  unit 0 {
    family inet {
      address 10.0.0.2/30;
    }
  }
}
```

```

    }
  }
  fe-1/2/1 {
    unit 0 {
      family inet {
        address 10.1.0.1/30;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 192.168.0.2/32;
      }
    }
  }
}

user@R2# show protocols
bgp {
  damping;
  group ext {
    type external;
    import damp;
    export send-direct;
    neighbor 10.0.0.1 {
      peer-as 100;
    }
    neighbor 10.1.0.2 {
      peer-as 300;
    }
  }
}

user@R2# show policy-options
policy-statement damp {
  term 1 {
    from {
      route-filter 10.128.0.0/9 exact damping dry;
      route-filter 0.0.0.0/0 prefix-length-range /0-/8 damping timid;
      route-filter 0.0.0.0/0 prefix-length-range /17-/32 damping aggressive;
    }
  }
}

policy-statement send-direct {
  term 1 {
    from protocol direct;
    then accept;
  }
}

damping aggressive {
  half-life 30;
  suppress 2500;
}

damping timid {
  half-life 5;
}

damping dry {

```

```
disable;
}

user@R2# show routing-options
autonomous-system 200;
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Causing Some Routes to Flap on page 3932](#)
- [Checking the Route Flaps on page 3932](#)
- [Verifying Route Flap Damping on page 3933](#)
- [Displaying the Details of a Damped Route on page 3934](#)
- [Verifying That Default Damping Parameters Are in Effect on page 3934](#)
- [Filtering the Damping Information on page 3935](#)

Causing Some Routes to Flap

Purpose To verify your route flap damping policy, some routes must flap. Having a live Internet feed almost guarantees that a certain number of route flaps will be present. If you have control over a remote system that is advertising the routes, you can modify the advertising router's policy to effect the advertisement and withdrawal of all routes or of a given prefix. In a test environment, you can cause routes to flap by clearing the BGP neighbors or by restarting the routing process on the BGP neighbors, as shown here.

Action From operational mode on Device R1 and Device R3, enter the **restart routing** command.



CAUTION: Use this command cautiously in a production network.

```
user@R1> restart routing
```

```
R1 started, pid 10474
```

```
user@R3> restart routing
```

```
R3 started, pid 10478
```

Meaning On Device R2, all of the routes from the neighbors are withdrawn and re-advertised.

Checking the Route Flaps

Purpose View the number of neighbor flaps.

Action From operational mode, enter the **show bgp summary** command.

```
user@R2> show bgp summary
```

```

Groups: 1 Peers: 2 Down peers: 0
Table      Tot Paths  Act Paths Suppressed  History Damp State  Pending
inet.0
          12         1         11         0         11         0
Peer      AS      InPkt    OutPkt    OutQ    Flaps  Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.0.0.1   100         10         10         0         4        2:50
0/9/0/9    0/0/0/0
10.1.0.2   300         10         10         0         4        2:53
1/3/1/2    0/0/0/0

```

Meaning This output was captured after the routing process was restarted on Device R2's neighbors four times.

Verifying Route Flap Damping

Purpose Verify that routes are being hidden due to damping.

Action From operational mode, enter the **show route damping suppressed** command.

```
user@R2> show route damping suppressed
```

```
inet.0: 15 destinations, 17 routes (6 active, 0 holddown, 11 hidden)
+ = Active Route, - = Last Active, * = Both
```

```

0.0.0.0/0      [BGP ] 00:00:12, localpref 100
                AS path: 100 I, validation-state: unverified
                > to 10.0.0.1 via fe-1/2/0.0
10.0.0.0/9     [BGP ] 00:00:12, localpref 100
                AS path: 100 I, validation-state: unverified
                > to 10.0.0.1 via fe-1/2/0.0
10.0.0.0/30    [BGP ] 00:00:12, localpref 100
                AS path: 100 I, validation-state: unverified
                > to 10.0.0.1 via fe-1/2/0.0
10.1.0.0/30    [BGP ] 00:00:15, localpref 100
                AS path: 300 I, validation-state: unverified
                > to 10.1.0.2 via fe-1/2/1.0
10.224.0.0/11  [BGP ] 00:00:12, localpref 100
                AS path: 100 I, validation-state: unverified
                > to 10.0.0.1 via fe-1/2/0.0
172.16.0.0/16  [BGP ] 00:00:12, localpref 100
                AS path: 100 I, validation-state: unverified
                > to 10.0.0.1 via fe-1/2/0.0
172.16.128.0/17 [BGP ] 00:00:12, localpref 100
                AS path: 100 I, validation-state: unverified
                > to 10.0.0.1 via fe-1/2/0.0
172.16.192.0/20 [BGP ] 00:00:12, localpref 100
                AS path: 100 I, validation-state: unverified
                > to 10.0.0.1 via fe-1/2/0.0
192.168.0.1/32 [BGP ] 00:00:12, localpref 100
                AS path: 100 I, validation-state: unverified
                > to 10.0.0.1 via fe-1/2/0.0
192.168.0.3/32 [BGP ] 00:00:15, localpref 100
                AS path: 300 I, validation-state: unverified
                > to 10.1.0.2 via fe-1/2/1.0
224.0.0.0/7    [BGP ] 00:00:12, localpref 100
                AS path: 100 I, validation-state: unverified
                > to 10.0.0.1 via fe-1/2/0.0

```

Meaning The output shows some routing instability. Eleven routes are hidden due to damping.

Displaying the Details of a Damped Route

Purpose Display the details of damped routes.

Action From operational mode, enter the **show route damping suppressed 172.16.192.0/20 detail** command.

```
user@R2> show route damping suppressed 172.16.192.0/20 detail

inet.0: 15 destinations, 17 routes (6 active, 0 holddown, 11 hidden)
172.16.192.0/20 (1 entry, 0 announced)
    BGP                /-101
        Next hop type: Router, Next hop index: 758
        Address: 0x9414484
        Next-hop reference count: 9
        Source: 10.0.0.1
        Next hop: 10.0.0.1 via fe-1/2/0.0, selected
        Session Id: 0x100201
        State: <Hidden Ext>
        Local AS: 200 Peer AS: 100
        Age: 52
        Validation State: unverified
        Task: BGP_100.10.0.0.1+55922
        AS path: 100 I
        Localpref: 100
        Router ID: 192.168.0.1
        Merit (last update/now): 4278/4196
        damping-parameters: aggressive
        Last update: 00:00:52 First update: 01:01:55
        Flaps: 8
        Suppressed. Reusable in: 01:14:40
        Preference will be: 170
```

Meaning This output indicates that the displayed route has a mask length that is equal to or greater than /17, and confirms that it has been correctly mapped to the aggressive damping profile. You can also see the route's current (and last) figure of merit value, and when the route is expected to become active if it remains stable.

Verifying That Default Damping Parameters Are in Effect

Purpose Locating a damped route with a /16 mask confirms that the default parameters are in effect.

Action From operational mode, enter the **show route damping suppressed detail | match 0/16** command.

```
user@R2> show route damping suppressed detail | match 0/16

172.16.0.0/16 (1 entry, 0 announced)

user@R2> show route damping suppressed 172.16.0.0/16 detail

inet.0: 15 destinations, 17 routes (6 active, 0 holddown, 11 hidden)
172.16.0.0/16 (1 entry, 0 announced)
```

```

BGP                               /-101
Next hop type: Router, Next hop index: 758
Address: 0x9414484
Next-hop reference count: 9
Source: 10.0.0.1
Next hop: 10.0.0.1 via fe-1/2/0.0, selected
Session Id: 0x100201
State: <Hidden Ext>
Local AS: 200 Peer AS: 100
Age: 1:58
Validation State: unverified
Task: BGP_100.10.0.0.1+55922
AS path: 100 I
Localpref: 100
Router ID: 192.168.0.1
Merit (last update/now): 3486/3202
Default damping parameters used
Last update: 00:01:58 First update: 01:03:01
Flaps: 8
Suppressed. Reusable in: 00:31:40
Preference will be: 170

```

Meaning Routes with a /16 mask are not impacted by the custom damping rules. Therefore, the default damping rules are in effect.

To repeat, the custom rules are as follows:

- Damp all prefixes with a mask length equal to or greater than 17 more aggressively than routes with a mask length between 9 and 16.
- Damp routes with a mask length between 0 and 8, inclusive, less than routes with a mask length greater than 8.
- Do not damp the 10.128.0.0/9 prefix at all.

Filtering the Damping Information

Purpose Use OR groupings or cascaded piping to simplify the determination of what damping profile is being used for routes with a given mask length.

Action From operational mode, enter the **show route damping suppressed** command.

```
user@R2> show route damping suppressed detail | match "0 announced | damp"
```

```

0.0.0.0/0 (1 entry, 0 announced)
    damping-parameters: timid
10.0.0.0/9 (1 entry, 0 announced)
    Default damping parameters used
    damping-parameters: aggressive
    damping-parameters: aggressive
10.224.0.0/11 (1 entry, 0 announced)
    Default damping parameters used
172.16.0.0/16 (1 entry, 0 announced)
    Default damping parameters used
172.16.128.0/17 (1 entry, 0 announced)
    damping-parameters: aggressive
172.16.192.0/20 (1 entry, 0 announced)
    damping-parameters: aggressive
192.168.0.1/32 (1 entry, 0 announced)

```

```

damping-parameters: aggressive
192.168.0.3/32 (1 entry, 0 announced)
damping-parameters: aggressive
224.0.0.0/7 (1 entry, 0 announced)
damping-parameters: timid

```

Meaning When you are satisfied that your EBGp routes are correctly associated with a damping profile, you can issue the **clear bgp damping** operational mode command to restore an active status to your damped routes, which will return your connectivity to normal operation.

Example: Configuring BGP Route Flap Damping Based on the MBGP MVPN Address Family

This example shows how to configure an multiprotocol BGP multicast VPN (also called Next-Generation MVPN) with BGP route flap damping.

- [Requirements on page 3936](#)
- [Overview on page 3936](#)
- [Configuration on page 3937](#)
- [Verification on page 3944](#)

Requirements

This example uses Junos OS Release 12.2. BGP route flap damping support for MBGP MVPN, specifically, and on an address family basis, in general, is introduced in Junos OS Release 12.2.


Overview

BGP route flap damping helps to diminish route instability caused by routes being repeatedly withdrawn and readvertised when a link is intermittently failing.

This example uses the default damping parameters and demonstrates an MBGP MVPN scenario with three provider edge (PE) routing devices, three customer edge (CE) routing devices, and one provider (P) routing device.

Figure 93 shows the topology used in this example.

Figure 93: MBGP MVPN with BGP Route Flap Damping



On PE Device R4, BGP route flap damping is configured for address family **inet-mvpn**. A routing policy called **dampPolicy** uses the **nlri-route-type** match condition to damp only MVPN route types 3, 4, and 5. All other MVPN route types are not damped.

This example shows the full configuration on all devices in the “[CLI Quick Configuration](#)” on page 3937 section. The “[Configuring Device R4](#)” on page 3940 section shows the step-by-step configuration for PE Device R4.

Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```

Device R1    set interfaces ge-1/2/0 unit 1 family inet address 10.1.1.1/30
              set interfaces ge-1/2/0 unit 1 family mpls
              set interfaces lo0 unit 1 family inet address 1.1.1.1/32
              set protocols ospf area 0.0.0.0 interface lo0.1 passive
              set protocols ospf area 0.0.0.0 interface ge-1/2/0.1
              set protocols pim rp static address 100.1.1.2
              set protocols pim interface all
              set routing-options router-id 1.1.1.1

Device R2    set interfaces ge-1/2/0 unit 2 family inet address 10.1.1.2/30
              set interfaces ge-1/2/0 unit 2 family mpls
              set interfaces ge-1/2/1 unit 5 family inet address 10.1.1.5/30
              set interfaces ge-1/2/1 unit 5 family mpls
              set interfaces vt-1/2/0 unit 2 family inet
              set interfaces lo0 unit 2 family inet address 1.1.1.2/32
              set interfaces lo0 unit 102 family inet address 100.1.1.2/32
              set protocols mpls interface ge-1/2/1.5
              set protocols bgp group ibgp type internal
              set protocols bgp group ibgp local-address 1.1.1.2
              set protocols bgp group ibgp family inet-vpn any
              set protocols bgp group ibgp family inet-mvpn signaling
              set protocols bgp group ibgp neighbor 1.1.1.4
              set protocols bgp group ibgp neighbor 1.1.1.5
              set protocols ospf area 0.0.0.0 interface lo0.2 passive
              set protocols ospf area 0.0.0.0 interface ge-1/2/1.5
              set protocols ldp interface ge-1/2/1.5
              set protocols ldp p2mp
              set policy-options policy-statement parent_vpn_routes from protocol bgp
              set policy-options policy-statement parent_vpn_routes then accept
              set routing-instances vpn-1 instance-type vrf
              set routing-instances vpn-1 interface ge-1/2/0.2
              set routing-instances vpn-1 interface vt-1/2/0.2
              set routing-instances vpn-1 interface lo0.102
              set routing-instances vpn-1 route-distinguisher 100:100
              set routing-instances vpn-1 provider-tunnel ldp-p2mp
              set routing-instances vpn-1 vrf-target target:1:1
              set routing-instances vpn-1 protocols ospf export parent_vpn_routes
              set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface lo0.102 passive
              set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface ge-1/2/0.2
              set routing-instances vpn-1 protocols pim rp static address 100.1.1.2

```

```
set routing-instances vpn-1 protocols pim interface ge-1/2/0.2 mode sparse
set routing-instances vpn-1 protocols mvpn
set routing-options router-id 1.1.1.2
set routing-options autonomous-system 1001
```

Device R3

```
set interfaces ge-1/2/0 unit 6 family inet address 10.1.1.6/30
set interfaces ge-1/2/0 unit 6 family mpls
set interfaces ge-1/2/1 unit 9 family inet address 10.1.1.9/30
set interfaces ge-1/2/1 unit 9 family mpls
set interfaces ge-1/2/2 unit 13 family inet address 10.1.1.13/30
set interfaces ge-1/2/2 unit 13 family mpls
set interfaces lo0 unit 3 family inet address 1.1.1.3/32
set protocols mpls interface ge-1/2/0.6
set protocols mpls interface ge-1/2/1.9
set protocols mpls interface ge-1/2/2.13
set protocols ospf area 0.0.0.0 interface lo0.3 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.6
set protocols ospf area 0.0.0.0 interface ge-1/2/1.9
set protocols ospf area 0.0.0.0 interface ge-1/2/2.13
set protocols ldp interface ge-1/2/0.6
set protocols ldp interface ge-1/2/1.9
set protocols ldp interface ge-1/2/2.13
set protocols ldp p2mp
set routing-options router-id 1.1.1.3
```

Device R4

```
set interfaces ge-1/2/0 unit 10 family inet address 10.1.1.10/30
set interfaces ge-1/2/0 unit 10 family mpls
set interfaces ge-1/2/1 unit 17 family inet address 10.1.1.17/30
set interfaces ge-1/2/1 unit 17 family mpls
set interfaces vt-1/2/0 unit 4 family inet
set interfaces lo0 unit 4 family inet address 1.1.1.4/32
set interfaces lo0 unit 104 family inet address 100.1.1.4/32
set protocols rsvp interface all aggregate
set protocols mpls interface all
set protocols mpls interface ge-1/2/0.10
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 1.1.1.4
set protocols bgp group ibgp family inet-vpn unicast
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp family inet-mvpn signaling damping
set protocols bgp group ibgp neighbor 1.1.1.2 import dampPolicy
set protocols bgp group ibgp neighbor 1.1.1.5
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface lo0.4 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.10
set protocols ldp interface ge-1/2/0.10
set protocols ldp p2mp
set policy-options policy-statement dampPolicy term term1 from family inet-mvpn
set policy-options policy-statement dampPolicy term term1 from nlri-route-type 3
set policy-options policy-statement dampPolicy term term1 from nlri-route-type 4
set policy-options policy-statement dampPolicy term term1 from nlri-route-type 5
set policy-options policy-statement dampPolicy term term1 then accept
set policy-options policy-statement dampPolicy then damping no-damp
set policy-options policy-statement dampPolicy then accept
```

```

set policy-options policy-statement parent_vpn_routes from protocol bgp
set policy-options policy-statement parent_vpn_routes then accept
set policy-options damping no-damp disable
set routing-instances vpn-1 instance-type vrf
set routing-instances vpn-1 interface vt-1/2/0.4
set routing-instances vpn-1 interface ge-1/2/1.17
set routing-instances vpn-1 interface lo0.104
set routing-instances vpn-1 route-distinguisher 100:100
set routing-instances vpn-1 vrf-target target:1:1
set routing-instances vpn-1 protocols ospf export parent_vpn_routes
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface lo0.104 passive
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface ge-1/2/1.17
set routing-instances vpn-1 protocols pim rp static address 100.1.1.2
set routing-instances vpn-1 protocols pim interface ge-1/2/1.17 mode sparse
set routing-instances vpn-1 protocols mvpn
set routing-options router-id 1.1.1.4
set routing-options autonomous-system 1001

```

Device R5

```

set interfaces ge-1/2/0 unit 14 family inet address 10.1.1.14/30
set interfaces ge-1/2/0 unit 14 family mpls
set interfaces ge-1/2/1 unit 21 family inet address 10.1.1.21/30
set interfaces ge-1/2/1 unit 21 family mpls
set interfaces vt-1/2/0 unit 5 family inet
set interfaces lo0 unit 5 family inet address 1.1.1.5/32
set interfaces lo0 unit 105 family inet address 100.1.1.5/32
set protocols mpls interface ge-1/2/0.14
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 1.1.1.5
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp family inet-mvpn signaling
set protocols bgp group ibgp neighbor 1.1.1.2
set protocols bgp group ibgp neighbor 1.1.1.4
set protocols ospf area 0.0.0.0 interface lo0.5 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.14
set protocols ldp interface ge-1/2/0.14
set protocols ldp p2mp
set policy-options policy-statement parent_vpn_routes from protocol bgp
set policy-options policy-statement parent_vpn_routes then accept
set routing-instances vpn-1 instance-type vrf
set routing-instances vpn-1 interface vt-1/2/0.5
set routing-instances vpn-1 interface ge-1/2/1.21
set routing-instances vpn-1 interface lo0.105
set routing-instances vpn-1 route-distinguisher 100:100
set routing-instances vpn-1 vrf-target target:1:1
set routing-instances vpn-1 protocols ospf export parent_vpn_routes
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface lo0.105 passive
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface ge-1/2/1.21
set routing-instances vpn-1 protocols pim rp static address 100.1.1.2
set routing-instances vpn-1 protocols pim interface ge-1/2/1.21 mode sparse
set routing-instances vpn-1 protocols mvpn
set routing-options router-id 1.1.1.5
set routing-options autonomous-system 1001

```

Device R6

```

set interfaces ge-1/2/0 unit 18 family inet address 10.1.1.18/30
set interfaces ge-1/2/0 unit 18 family mpls

```

```
set interfaces lo0 unit 6 family inet address 1.1.1.6/32
set protocols sap listen 224.1.1.1
set protocols ospf area 0.0.0.0 interface lo0.6 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.18
set protocols pim rp static address 100.1.1.2
set protocols pim interface all
set routing-options router-id 1.1.1.6
```

Device R7

```
set interfaces ge-1/2/0 unit 22 family inet address 10.1.1.22/30
set interfaces ge-1/2/0 unit 22 family mpls
set interfaces lo0 unit 7 family inet address 1.1.1.7/32
set protocols ospf area 0.0.0.0 interface lo0.7 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.22
set protocols pim rp static address 100.1.1.2
set protocols pim interface all
set routing-options router-id 1.1.1.7
```

Configuring Device R4

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R4:

1. Configure the interfaces.

```
[edit interfaces]
user@R4# set ge-1/2/0 unit 10 family inet address 10.1.1.10/30
user@R4# set ge-1/2/0 unit 10 family mpls

user@R4# set ge-1/2/1 unit 17 family inet address 10.1.1.17/30
user@R4# set ge-1/2/1 unit 17 family mpls

user@R4# set vt-1/2/0 unit 4 family inet

user@R4# set lo0 unit 4 family inet address 1.1.1.4/32
user@R4# set lo0 unit 104 family inet address 100.1.1.4/32
```

2. Configure MPLS and the signaling protocols on the interfaces.

```
[edit protocols]
user@R4# set mpls interface all
user@R4# set mpls interface ge-1/2/0.10
user@R4# set rsvp interface all aggregate
user@R4# set ldp interface ge-1/2/0.10
user@R4# set ldp p2mp
```

3. Configure BGP.

The BGP configuration enables BGP route flap damping for the **inet-mvpn** address family. The BGP configuration also imports into the routing table the routing policy called **dampPolicy**. This policy is applied to neighbor PE Device R2.

```
[edit protocols bgp group ibgp]
user@R4# set type internal
```

```

user@R4# set local-address 1.1.1.4
user@R4# set family inet-vpn unicast
user@R4# set family inet-vpn any
user@R4# set family inet-mvpn signaling damping
user@R4# set neighbor 1.1.1.2 import dampPolicy
user@R4# set neighbor 1.1.1.5

```

4. Configure an interior gateway protocol.

```

[edit protocols ospf]
user@R4# set traffic-engineering

[edit protocols ospf area 0.0.0.0]
user@R4# set interface all
user@R4# set interface lo0.4 passive
user@R4# set interface ge-1/2/0.10

```

5. Configure a damping policy that uses the **nlri-route-type** match condition to damp only MVPN route types 3, 4, and 5.

```

[edit policy-options policy-statement dampPolicy term term1]
user@R4# set from family inet-mvpn
user@R4# set from nlri-route-type 3
user@R4# set from nlri-route-type 4
user@R4# set from nlri-route-type 5
user@R4# set then accept

```

6. Configure the **damping** policy to disable BGP route flap damping.

The **no-damp** policy (**damping no-damp disable**) causes any damping state that is present in the routing table to be deleted. The **then damping no-damp** statement applies the **no-damp** policy as an action and has no **from** match conditions. Therefore, all routes that are not matched by **term1** are matched by this term, with the result that all other MVPN route types are not damped.

```

[edit policy-options policy-statement dampPolicy]
user@R4# set then damping no-damp
user@R4# set then accept

```

```

[edit policy-options]
user@R4# set damping no-damp disable

```

7. Configure the **parent_vpn_routes** to accept all other BGP routes that are not from the **inet-mvpn** address family.

This policy is applied as an OSPF export policy in the routing instance.

```

[edit policy-options policy-statement parent_vpn_routes]
user@R4# set from protocol bgp
user@R4# set then accept

```

8. Configure the VPN routing and forwarding (VRF) instance.

```

[edit routing-instances vpn-1]
user@R4# set instance-type vrf
user@R4# set interface vt-1/2/0.4
user@R4# set interface ge-1/2/1.17
user@R4# set interface lo0.104

```

```
user@R4# set route-distinguisher 100:100
user@R4# set vrf-target target:1:1
user@R4# set protocols ospf export parent_vpn_routes
user@R4# set protocols ospf area 0.0.0.0 interface lo0.104 passive
user@R4# set protocols ospf area 0.0.0.0 interface ge-1/2/1.17
user@R4# set protocols pim rp static address 100.1.1.2
user@R4# set protocols pim interface ge-1/2/1.17 mode sparse
user@R4# set protocols mvpn
```

9. Configure the router ID and the autonomous system (AS) number.

```
[edit routing-options]
user@R4# set router-id 1.1.1.4
user@R4# set autonomous-system 1001
```

10. If you are done configuring the device, commit the configuration.

```
user@R4# commit
```

Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, **show routing-instances**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R4# show interfaces
ge-1/2/0 {
  unit 10 {
    family inet {
      address 10.1.1.10/30;
    }
    family mpls;
  }
}
ge-1/2/1 {
  unit 17 {
    family inet {
      address 10.1.1.17/30;
    }
    family mpls;
  }
}
vt-1/2/0 {
  unit 4 {
    family inet;
  }
}
lo0 {
  unit 4 {
    family inet {
      address 1.1.1.4/32;
    }
  }
}
unit 104 {
  family inet {
```

```

        address 100.1.1.4/32;
    }
}
}
user@R4# show protocols
rsvp {
    interface all {
        aggregate;
    }
}
mpls {
    interface all;
    interface ge-1/2/0.10;
}
bgp {
    group ibgp {
        type internal;
        local-address 1.1.1.4;
        family inet-vpn {
            unicast;
            any;
        }
        family inet-mvpn {
            signaling {
                damping;
            }
        }
        neighbor 1.1.1.2 {
            import dampPolicy;
        }
        neighbor 1.1.1.5;
    }
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface all;
        interface lo0.4 {
            passive;
        }
        interface ge-1/2/0.10;
    }
}
ldp {
    interface ge-1/2/0.10;
    p2mp;
}
user@R4# show policy-options
policy-statement dampPolicy {
    term term1 {
        from {
            family inet-mvpn;
            nlri-route-type [ 3 4 5 ];
        }
        then accept;
    }
}

```

```
    }
    then {
        damping no-damp;
        accept;
    }
}
policy-statement parent_vpn_routes {
    from protocol bgp;
    then accept;
}
damping no-damp {
    disable;
}
```

user@R4# show routing-instances

```
vpn-1 {
    instance-type vrf;
    interface vt-1/2/0.4;
    interface ge-1/2/1.17;
    interface lo0.104;
    route-distinguisher 100:100;
    vrf-target target:1:1;
    protocols {
        ospf {
            export parent_vpn_routes;
            area 0.0.0.0 {
                interface lo0.104 {
                    passive;
                }
                interface ge-1/2/1.17;
            }
        }
        pim {
            rp {
                static {
                    address 100.1.1.2;
                }
            }
            interface ge-1/2/1.17 {
                mode sparse;
            }
        }
        mvpn;
    }
}
```

user@R4# show routing-options

```
router-id 1.1.1.4;
autonomous-system 1001;
```

Verification

Confirm that the configuration is working properly.

- [Verifying That Route Flap Damping Is Disabled on page 3945](#)
- [Verifying Route Flap Damping on page 3945](#)

Verifying That Route Flap Damping Is Disabled

Purpose Verify the presence of the **no-damp** policy, which disables damping for MVPN route types other than 3, 4, and 5.

Action From operational mode, enter the **show policy damping** command.

```
user@R4> show policy damping
Default damping information:
  Halflife: 15 minutes
  Reuse merit: 750 Suppress/cutoff merit: 3000
  Maximum suppress time: 60 minutes
  Computed values:
    Merit ceiling: 12110
    Maximum decay: 6193
Damping information for "no-damp":
  Damping disabled
```

Meaning The output shows that the default damping parameters are in effect and that the **no-damp** policy is also in effect for the specified route types.

Verifying Route Flap Damping

Purpose Check whether BGP routes have been damped.

Action From operational mode, enter the **show bgp summary** command.

```
user@R4> show bgp summary
Groups: 1 Peers: 2 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
bgp.l3vpn.0
      6      6      0      0      0      0
bgp.l3vpn.2
      0      0      0      0      0      0
bgp.mvpn.0
      2      2      0      0      0      0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
1.1.1.2 1001 3159 3155 0 0 23:43:47
Establ
  bgp.l3vpn.0: 3/3/3/0
  bgp.l3vpn.2: 0/0/0/0
  bgp.mvpn.0: 1/1/1/0
  vpn-1.inet.0: 3/3/3/0
  vpn-1.mvpn.0: 1/1/1/0
1.1.1.5 1001 3157 3154 0 0 23:43:40
Establ
  bgp.l3vpn.0: 3/3/3/0
  bgp.l3vpn.2: 0/0/0/0
  bgp.mvpn.0: 1/1/1/0
  vpn-1.inet.0: 3/3/3/0
  vpn-1.mvpn.0: 1/1/1/0
```

Meaning The Damp State field shows that zero routes in the bgp.mvpn.0 routing table have been damped. Further down, the last number in the State field shows that zero routes have been damped for BGP peer 1.1.1.2.

- Related Documentation**
- [Understanding External BGP Peering Sessions on page 3601](#)
 - *BGP Configuration Overview*

PART 63

Monitoring and Troubleshooting

- [BGP Monitoring Configuration on page 3949](#)
- [Routine Monitoring on page 3959](#)

BGP Monitoring Configuration

- [Example: Configuring BGP Trace Operations on page 3949](#)
- [Tracing BMP Operations on page 3956](#)

Example: Configuring BGP Trace Operations

- [Understanding Trace Operations for BGP Protocol Traffic on page 3949](#)
- [Example: Viewing BGP Trace Files on Logical Systems on page 3951](#)

Understanding Trace Operations for BGP Protocol Traffic

You can trace various BGP protocol traffic to help you debug BGP protocol issues. To trace BGP protocol traffic, include the **traceoptions** statement at the **[edit protocols bgp]** hierarchy level. For routing instances, include the **traceoptions** statement at the **[edit routing-instances routing-instance-name protocols bgp]** hierarchy level.

```
traceoptions {  
    file filename <files number> <size size> <world-readable | no-world-readable>;  
    flag flag <flag-modifier> <disable>;  
}
```

You can specify the following BGP protocol-specific trace options using the **flag** statement:

- **4byte-as**—4-byte AS events.
- **bfd**—BFD protocol events.
- **damping**—Damping operations.
- **graceful-restart**—Graceful restart events.
- **keepalive**—BGP keepalive messages.
- **nsr-synchronization**—Nonstop active routing synchronization events.
- **open**—BGP open packets. These packets are sent between peers when they are establishing a connection.
- **packets**—All BGP protocol packets.
- **refresh**—BGP refresh packets.
- **update**—BGP update packets. These packets provide routing updates to BGP systems.

Global tracing options are inherited from the configuration set by the **traceoptions** statement at the **[edit routing-options]** hierarchy level. You can override the following global trace options for the BGP protocol using the **traceoptions flag** statement included at the **[edit protocols bgp]** hierarchy level:

- **all**—All tracing operations
- **general**—All normal operations and routing table changes (a combination of the normal and route trace operations)
- **normal**—Normal events
- **policy**—Policy processing
- **route**—Routing information
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing

You can optionally specify one or more of the following flag modifiers:

- **detail**—Detailed trace information.
- **filter**—Filter trace information. Applies only to **route** and **damping** tracing flags.
- **receive**—Packets being received.
- **send**—Packets being transmitted.



NOTE: Use the **all** trace flag and the **detail** flag modifier with caution because these might cause the CPU to become very busy.



NOTE: If you only enable the **update** flag, received keepalive messages do not generate a trace message.

You can filter trace statements and display only the statement information that passes through the filter by specifying the **filter** flag modifier. The **filter** modifier is only supported for the **route** and **damping** tracing flags.

The **match-on** statement specifies filter matches based on prefixes. It is used to match on route filters.



NOTE: Per-neighbor trace filtering is not supported on a BGP per-neighbor level for **route** and **damping** flags. Trace option filtering support is on a peer group level.

Example: Viewing BGP Trace Files on Logical Systems

This example shows how to list and view files that are stored on a logical system.

- [Requirements on page 3951](#)
- [Overview on page 3951](#)
- [Configuration on page 3952](#)
- [Verification on page 3955](#)

Requirements

- You must have the **view** privilege for the logical system.
- Configure a network, such as the BGP network shown in “[Example: Configuring Internal BGP Peering Sessions on Logical Systems](#)” on page 3636.

Overview

Logical systems have their individual directory structure created in the `/var/logical-systems/logical-system-name` directory. It contains the following subdirectories:

- `/config`—Contains the active configuration specific to the logical system.
- `/log`—Contains system log and tracing files specific to the logical system.

To maintain backward compatibility for the log files with previous versions of Junos OS, a symbolic link (symlink) from the `/var/logs/logical-system-name` directory to the `/var/logical-systems/logical-system-name` directory is created when a logical system is configured.

- `/tmp`—Contains temporary files specific to the logical system.

The file system for each logical system enables logical system users to view trace logs and modify logical system files. Logical system administrators have full access to view and modify all files specific to the logical system.

Logical system users and administrators can save and load configuration files at the logical-system level using the **save** and **load** configuration mode commands. In addition, they can also issue the **show log**, **monitor**, and **file** operational mode commands at the logical-system level.

This example shows how to configure and view a BGP trace file on a logical system. The steps can be adapted to apply to trace operations for any Junos OS hierarchy level that supports trace operations.



TIP: To view a list of hierarchy levels that support tracing operations, enter the `help apropos traceoptions` command in configuration mode.

Configuration

- [Configuring Trace Operations on page 3952](#)
- [Viewing the Trace File on page 3952](#)
- [Deactivating and Reactivating Trace Logging on page 3954](#)
- [Results on page 3955](#)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set logical-systems A protocols bgp group internal-peers traceoptions file bgp-log
set logical-systems A protocols bgp group internal-peers traceoptions file size 10k
set logical-systems A protocols bgp group internal-peers traceoptions file files 2
set logical-systems A protocols bgp group internal-peers traceoptions flag update detail
```

Configuring Trace Operations

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the trace operations:

1. Configure trace operations on the logical system.

```
[edit logical-systems A protocols bgp group internal-peers]
user@host# set traceoptions file bgp-log
user@host# set traceoptions file size 10k
user@host# set traceoptions file files 2
user@host# set traceoptions flag update detail
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Viewing the Trace File

Step-by-Step Procedure

To view the trace file:

1. In operational mode on the main router, list the directories on the logical system.

```
user@host> file list /var/logical-systems/A
/var/logical-systems/A:
config/
log/
tmp/
```

2. In operational mode on the main router, list the log files on the logical system.

```
user@host> file list /var/logical-systems/A/log/
/var/logical-systems/A/log:
bgp-log
```


3. View the contents of the **bgp-log** file.

```
user@host> file show /var/logical-systems/A/log/bgp-log
Aug 10 17:12:01 trace_on: Tracing to "/var/log/A/bgp-log" started
Aug 10 17:14:22.826182 bgp_peer_mgmt_clear:5829: NOTIFICATION sent to
192.163.6.4 (Internal AS 17): code 6 (Cease) subcode 4 (Administratively
Reset), Reason: Management session cleared BGP neighbor
Aug 10 17:14:22.826445 bgp_send: sending 21 bytes to 192.163.6.4 (Internal
AS 17)
Aug 10 17:14:22.826499
Aug 10 17:14:22.826499 BGP SEND 192.168.6.5+64965 -> 192.163.6.4+179
Aug 10 17:14:22.826559 BGP SEND message type 3 (Notification) length 21
Aug 10 17:14:22.826598 BGP SEND Notification code 6 (Cease) subcode 4
(Administratively Reset)
Aug 10 17:14:22.831756 bgp_peer_mgmt_clear:5829: NOTIFICATION sent to
192.168.40.4 (Internal AS 17): code 6 (Cease) subcode 4 (Administratively
Reset), Reason: Management session cleared BGP neighbor
Aug 10 17:14:22.831851 bgp_send: sending 21 bytes to 192.168.40.4 (Internal
AS 17)
Aug 10 17:14:22.831901
Aug 10 17:14:22.831901 BGP SEND 192.168.6.5+53889 -> 192.168.40.4+179
Aug 10 17:14:22.831959 BGP SEND message type 3 (Notification) length 21
Aug 10 17:14:22.831999 BGP SEND Notification code 6 (Cease) subcode 4
(Administratively Reset)
...
```

4. Filter the output of the log file.

```
user@host> file show /var/logical-systems/A/log/bgp-log | match "flags 0x40"
Aug 10 17:14:54.867460 BGP SEND flags 0x40 code Origin(1): IGP
Aug 10 17:14:54.867595 BGP SEND flags 0x40 code ASPath(2) length 0: <null>
Aug 10 17:14:54.867650 BGP SEND flags 0x40 code NextHop(3): 192.168.6.5
Aug 10 17:14:54.867692 BGP SEND flags 0x40 code LocalPref(5): 100
Aug 10 17:14:54.884529 BGP RECV flags 0x40 code Origin(1): IGP
Aug 10 17:14:54.884581 BGP RECV flags 0x40 code ASPath(2) length 0: <null>
Aug 10 17:14:54.884628 BGP RECV flags 0x40 code NextHop(3): 192.163.6.4
Aug 10 17:14:54.884667 BGP RECV flags 0x40 code LocalPref(5): 100
Aug 10 17:14:54.911377 BGP RECV flags 0x40 code Origin(1): IGP
Aug 10 17:14:54.911422 BGP RECV flags 0x40 code ASPath(2) length 0: <null>
Aug 10 17:14:54.911466 BGP RECV flags 0x40 code NextHop(3): 192.168.40.4
Aug 10 17:14:54.911507 BGP RECV flags 0x40 code LocalPref(5): 100
Aug 10 17:14:54.916008 BGP SEND flags 0x40 code Origin(1): IGP
Aug 10 17:14:54.916054 BGP SEND flags 0x40 code ASPath(2) length 0: <null>
Aug 10 17:14:54.916100 BGP SEND flags 0x40 code NextHop(3): 192.168.6.5
Aug 10 17:14:54.916143 BGP SEND flags 0x40 code LocalPref(5): 100
Aug 10 17:14:54.920304 BGP RECV flags 0x40 code Origin(1): IGP
Aug 10 17:14:54.920348 BGP RECV flags 0x40 code ASPath(2) length 0: <null>
Aug 10 17:14:54.920393 BGP RECV flags 0x40 code NextHop(3): 10.0.0.10
Aug 10 17:14:54.920434 BGP RECV flags 0x40 code LocalPref(5): 100
```

5. View the tracing operations in real time.

```
user@host> clear bgp neighbor logical-system A
Cleared 2 connections
```



CAUTION: Clearing the BGP neighbor table is disruptive in a production environment.

6. Run the **monitor start** command with an optional **match** condition.

```
user@host> monitor start A/bgp-log | match 0.0.0.0/0
Aug 10 19:21:40.773467 BGP RECV          0.0.0.0/0
Aug 10 19:21:40.773685 bgp_rcv_nlri: 0.0.0.0/0
Aug 10 19:21:40.773778 bgp_rcv_nlri: 0.0.0.0/0 belongs to meshgroup
Aug 10 19:21:40.773832 bgp_rcv_nlri: 0.0.0.0/0 qualified bnp->ribact 0x0
12afcb 0x0
```

7. Pause the **monitor** command by pressing Esc-Q.
To unpause the output, press Esc-Q again.
8. Halt the **monitor** command by pressing Enter and typing **monitor stop**.
[Enter]
user@host> **monitor stop**
9. When you are finished troubleshooting, consider deactivating trace logging to avoid any unnecessary impact to system resources.

```
[edit protocols bgp group internal-peers]
user@host:A# deactivate traceoptions
user@host:A# commit
```

When configuration is deactivated, it appears in the configuration with the **inactive** tag. To reactivate trace operations, use the **activate** configuration-mode statement.

```
[edit protocols bgp group internal-peers]
user@host:A# show
```

```
type internal;
inactive: traceoptions {
    file bgp-log size 10k files 2;
    flag update detail;
    flag all;
}
local-address 192.168.6.5;
export send-direct;
neighbor 192.163.6.4;
neighbor 192.168.40.4;
```

10. To reactivate trace operations, use the **activate** configuration-mode statement.

```
[edit protocols bgp group internal-peers]
user@host:A# activate traceoptions
user@host:A# commit
```

Deactivating and Reactivating Trace Logging

Step-by-Step Procedure

To deactivate and reactivate the trace file:

1. When you are finished troubleshooting, consider deactivating trace logging to avoid an unnecessary impact to system resources.

```
[edit protocols bgp group internal-peers]
user@host:A# deactivate traceoptions
user@host:A# commit
```

When configuration is deactivated, the statement appears in the configuration with the **inactive** tag.

```
[edit protocols bgp group internal-peers]
user@host:A# show

type internal;
inactive: traceoptions {
    file bgp-log size 10k files 2;
    flag update detail;
    flag all;
}
local-address 192.168.6.5;
export send-direct;
neighbor 192.163.6.4;
neighbor 192.168.40.4;
```

2. To reactivate logging, use the **activate** configuration-mode statement.

```
[edit protocols bgp group internal-peers]
user@host:A# activate traceoptions
user@host:A# commit
```

Results

From configuration mode, confirm your configuration by entering the **show logical-systems A protocols bgp group internal-peers** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host:# show logical-systems A protocols bgp group internal-peers
traceoptions {
    file bgp-log size 10k files 2;
    flag update detail;
}
```

Verification

Confirm that the configuration is working properly.

Verifying That the Trace Log File Is Operating

Purpose Make sure that events are being written to the log file.

Action user@host:A> **show log bgp-log**
 Aug 12 11:20:57 trace_on: Tracing to "/var/log/A/bgp-log" started

Related Documentation

- [Understanding External BGP Peering Sessions on page 3601](#)
- [BGP Configuration Overview](#)

Tracing BMP Operations

You can trace BMP operations for all BMP stations by configuring the **traceoptions** statement at the **[edit routing-options bmp]** hierarchy level or for specific BMP stations at the **[edit routing-options bmp station station-name]** hierarchy level.

To trace BMP operations, complete the following steps:

1. Configure the **traceoptions** statement:

```
traceoptions {  
    file filename <files number> <size size> <world-readable | no-world-readable>;  
    flag flag <flag-modifier> <disable>;  
}
```

2. Specify the name of the file to receive the output of the tracing operation using the **file** option. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. We recommend that you place BMP tracing output in the file **bmp-log**.
3. (Optional) Specify the maximum number of trace files using the **files** option. When a trace file named **trace-file.0** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum number of files, you must also specify a maximum file size with the **size** option.
4. (Optional) Specify the maximum size of each trace file using the **size** option in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.
5. (Optional) You can specify that the log files are either **world-readable** (accessible to all users on the device) or **no-world-readable** (not accessible to all users on the device).
6. You can specify the following BMP-specific trace options using the **flag** statement:
 - **all**—Trace all BMP monitoring operations.
 - **down**—Down messages.
 - **error**—Error conditions.
 - **event**—Major events, station establishment, errors, and events.
 - **general**—General events.
 - **normal**—Normal events.
 - **packets**—All messages.
 - **policy**—Policy processing.
 - **route**—Routing information.
 - **route-monitoring**—Route monitoring messages.

- **state**—State transitions.
- **statistics**—Statistics messages.
- **task**—Routing protocol task processing.
- **timer**—Routing protocol timer processing.
- **up**—Up messages.
- **write**—Writing of messages.

You can optionally specify one or more of the following flag modifiers:

- **detail**—Provide detailed trace information.
- **disable**—Disable the tracing flag.
- **receive**—Trace the packets being received.
- **send**—Trace the packets being transmitted.



NOTE: Use the all trace flag and the detail flag modifier with caution due to the increased computer processing power required.

**Related
Documentation**

- [Configuring BGP Monitoring Protocol Version 3 on page 3647](#)

Routine Monitoring

- [Monitoring BGP Routing Information on page 3959](#)

Monitoring BGP Routing Information

Purpose	Use the monitoring functionality to monitor BGP routing information on the routing device.
Action	To view BGP routing information in the CLI, enter the following commands: <ul style="list-style-type: none">• <code>show bgp summary</code>• <code>show bgp neighbor</code>
Related Documentation	<ul style="list-style-type: none">• show bgp neighbor on page 2710• show bgp summary on page 4124

PART 64

Configuration Statement and Operational Commands

- [Configuration Statements on page 3963](#)
- [Operational Commands on page 4091](#)

Configuration Statements

- [accept-remote-nexthop](#) on page 3966
- [advertise-external](#) on page 3967
- [advertise-inactive](#) on page 3969
- [advertise-peer-as](#) on page 3970
- [algorithm \(BGP BFD Authentication\)](#) on page 3971
- [apply-groups](#) on page 3973
- [apply-groups-except](#) on page 3973
- [authentication \(BGP BFD Liveness Detection\)](#) on page 3974
- [authentication-algorithm](#) on page 3976
- [authentication-key \(Protocols BGP and BMP\)](#) on page 3979
- [authentication-key-chain \(Protocols BGP and BMP\)](#) on page 3980
- [bfd-liveness-detection \(Protocols BGP\)](#) on page 3981
- [bgp](#) on page 3985
- [bgp-orf-cisco-mode](#) on page 3986
- [cluster](#) on page 3988
- [connection-mode](#) on page 3989
- [damping \(Protocols BGP\)](#) on page 3990
- [description \(Protocols BGP\)](#) on page 3992
- [detection-time \(BFD Liveness Detection\)](#) on page 3993
- [disable \(Protocols BGP\)](#) on page 3994
- [disable \(BGP Graceful Restart\)](#) on page 3995
- [export \(Protocols BGP\)](#) on page 3996
- [family](#) on page 3997
- [graceful-restart \(Protocols BGP\)](#) on page 4001
- [group \(Protocols BGP\)](#) on page 4003
- [hold-down](#) on page 4006
- [hold-down-interval \(BGP BFD Liveness Detection\)](#) on page 4008
- [hold-time \(Protocols BGP\)](#) on page 4010

- [import](#) on page 4012
- [include-mp-next-hop](#) on page 4014
- [initiation-message](#) on page 4015
- [keep](#) on page 4016
- [key-chain](#) (BGP BFD Authentication) on page 4018
- [local-address](#) (Protocols BGP) on page 4020
- [local-address](#) (Protocols BMP) on page 4022
- [local-as](#) on page 4023
- [local-port](#) on page 4025
- [local-preference](#) on page 4026
- [log-updown](#) on page 4027
- [loops](#) (BGP Address Family) on page 4028
- [loose-check](#) (BGP BFD Authentication) on page 4030
- [maximum-ecmp](#) on page 4031
- [metric-out](#) on page 4032
- [minimum-interval](#) (BFD Liveness Detection) on page 4034
- [minimum-interval](#) (transmit-interval) on page 4036
- [minimum-receive-interval](#) (BFD Liveness Detection) on page 4038
- [monitor](#) (Protocols BMP) on page 4039
- [mtu-discovery](#) on page 4040
- [multihop](#) on page 4042
- [multiplier](#) (BFD Liveness Detection) on page 4044
- [neighbor](#) (Protocols BGP) on page 4046
- [no-adaptation](#) (BFD Liveness Detection) on page 4050
- [no-advertise-peer-as](#) on page 4051
- [no-aggregator-id](#) on page 4052
- [no-client-reflect](#) on page 4053
- [out-delay](#) on page 4054
- [outbound-route-filter](#) on page 4056
- [passive](#) (Protocols BGP) on page 4057
- [path-selection](#) on page 4058
- [peer-as](#) (Protocols BGP) on page 4060
- [post-policy](#) on page 4061
- [pre-policy](#) on page 4062
- [preference](#) (Protocols BGP) on page 4063
- [priority](#) (Protocols BMP) on page 4064
- [remove-private](#) on page 4065

- [restart-time \(BGP Graceful Restart\) on page 4067](#)
- [route-monitoring on page 4068](#)
- [session-mode on page 4070](#)
- [stale-routes-time on page 4071](#)
- [station on page 4072](#)
- [station-address on page 4073](#)
- [station-port on page 4074](#)
- [statistics-timeout on page 4075](#)
- [tcp-mss \(Protocols BGP\) on page 4076](#)
- [threshold \(detection-time\) on page 4077](#)
- [threshold \(transmit-interval\) on page 4079](#)
- [traceoptions \(Protocols BGP\) on page 4081](#)
- [traceoptions \(Protocols BMP\) on page 4084](#)
- [transmit-interval \(BFD Liveness Detection\) on page 4086](#)
- [version \(BFD Liveness Detection\) on page 4088](#)

accept-remote-nexthop

Syntax	accept-remote-nexthop;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Specify that a single-hop EBGp peer accepts a remote next hop with which it does not share a common subnet. Configure a separate import policy on the EBGp peer to specify the remote next hop.</p> <p>For Junos OS Release 13.3 and later releases, specify that a multihop EBGp peer accepts a remote next hop with which it does not share a common subnet. This allows working around current resolver limitations to realize multipath forwarding in recursive next-hop resolution scenarios.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Single-Hop EBGp Peers to Accept Remote Next Hops on page 3833• Configuring Routing Policies to Control BGP Route Advertisements on page 3769• <i>multipath</i>

advertise-external

Syntax	<code>advertise-external {conditional};</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>neighbor-address</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.3.</p> <p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Specify BGP to advertise the best external route into an IBGP mesh group, a route reflector cluster, or an AS confederation even if the best route is an internal route.</p> <p>In general, deployed BGP implementations do not advertise the external route with the highest local preference value to internal peers unless it is the best route. Although this behavior was required by an earlier version of the BGP version 4 specification, RFC 1771, it was typically not followed in order to minimize the amount of advertised information and to prevent routing loops. However, there are scenarios in which advertising the best external route is beneficial, in particular, situations that can result in IBGP route oscillation.</p> <p>The advertise-external statement is supported at both the group and neighbor level. If you configure the statement at the neighbor level, you must configure it for all neighbors in a group. Otherwise, the group is automatically split into different groups.</p> <p>In a confederation, when advertising a route to a confederation border router, any route from a different confederation sub-AS is considered external. When configuring the advertise-external statement for an AS confederation, it is recommended that EBGP peers belonging to different autonomous systems are configured in a separate EBGP peer group. This ensures consistency while BGP sends the best external route to peers in the configured peer group.</p> <p>To configure the advertise-external statement on a route reflector, you must disable intracluster reflection with the no-client-reflect statement.</p> <p>When a routing device is configured as a route reflector for a cluster, a route advertised by the route reflector is considered internal if it is received from an internal peer with the same cluster identifier or if both peers have no cluster identifier configured. A route received from an internal peer that belongs to another cluster, that is, with a different cluster identifier, is considered external.</p>

The **conditional** option causes BGP to advertise the external route only if the route selection process reaches the point where the multiple exit discriminator (MED) metric is evaluated. As a result, an external route with an AS path longer than that of the active path is not advertised.

Junos OS also provides support for configuring a BGP export policy that matches on the state of an advertised route. You can match on either active or inactive routes.

Default BGP does not advertise the external route with the highest local preference value to internal peers unless it is the best route.

Options **conditional**—(Optional) Advertise the best external path only if the route selection process reaches the point at which the multiple exit discriminator (MED) metric is evaluated. The **conditional** option restricts advertisement to when the best external path and the active path are equal until the MED step of the route selection process. This implies that external routes with a longer AS path length than the active path, for instance, are not advertised. The criteria used for selecting the best external path is the same whether or not the **conditional** option is configured.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring a Routing Policy to Advertise the Best External Route to Internal Peers*
- [advertise-inactive on page 3969](#)

advertise-inactive

Syntax	advertise-inactive;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Configure the routing table to export to BGP the best route learned by BGP even if Junos OS did not select this route to be an active route.</p> <p>One way to achieve multivendor compatibility is to include the advertise-inactive statement in the external BGP (EBGP) configuration. By default, BGP stores the route information it receives from update messages in the Junos OS routing table, and the routing table exports only active routes into BGP, which BGP then advertises to its peers. The advertise-inactive statement causes Junos OS to advertise the best BGP route that is inactive because of IGP preference. When you use the advertise-inactive statement, the Junos OS device uses, for example, the OSPF route for forwarding, and the other vendor's device uses the EBGP route for forwarding. However, from the perspective of an EBGP peer in a neighboring AS, both vendors' devices appear to behave the same way.</p>
Default	By default, BGP stores the route information it receives from update messages in the Junos OS routing table, and the routing table exports only active routes into BGP, which BGP then advertises to its peers.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring the Preference Value for BGP Routes on page 3787 • Configuring Routing Policies to Control BGP Route Advertisements on page 3769 • advertise-external on page 3967

advertise-peer-as

Syntax	advertise-peer-as;
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group</i> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp <i>group</i> <i>group-name</i>], [edit protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Disable the default behavior of suppressing AS routes.</p> <p>If you include the advertise-peer-as statement in the configuration, BGP advertises routes learned from one external BGP (EBGP) peer back to another EBGP peer in the same autonomous system (AS) but not back to the originating peer.</p> <p>Another way to disable the route suppression default behavior is with the as-override statement. If you include both the as-override and no-advertise-peer-as statements in the configuration, the no-advertise-peer-as statement is ignored.</p>
Default	By default, Junos OS does not advertise the routes learned from one EBGP peer back to the same external BGP (EBGP) peer. In addition, the software does not advertise those routes back to any EBGP peers that are in the same AS as the originating peer, regardless of the routing instance.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Enabling BGP Route Advertisements</i> • <i>Example: Configuring a Layer 3 VPN with Route Reflection and AS Override</i> • no-advertise-peer-as on page 4051

algorithm (BGP BFD Authentication)

Syntax	<code>algorithm <i>algorithm-name</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection authentication],</p> <p>[edit protocols bgp bgp bfd-liveness-detection authentication],</p> <p>[edit protocols bgp group <i>group-name</i> bfd-liveness-detection authentication],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection authentication],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp bgp bfd-liveness-detection authentication],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection authentication],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection authentication]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for BFD authentication introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Configure the algorithm used to authenticate the specified BFD session.
Options	<p><i>algorithm-name</i>—Authentication algorithm name: simple-password, keyed-md5, keyed-sha-1, meticulous-keyed-md5, meticulous-keyed-sha-1.</p> <p>simple-password—Plain-text password. One to 16 bytes of plain text are used to authenticate the BFD session. One or more passwords can be configured. This method is the least secure and should be used only when BFD sessions are not subject to packet interception.</p> <p>keyed-md5—Keyed Message Digest 5 hash algorithm for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed MD5 uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than or equal to the last sequence number received. Although more secure than a simple password, this method is vulnerable to replay attacks. Increasing the rate at which the sequence number is updated can reduce this risk.</p>

meticulous-keyed-md5—Meticulous keyed Message Digest 5 hash algorithm. This method works in the same manner as keyed MD5, but the sequence number is updated with every packet. Although more secure than keyed MD5 and simple passwords, this method can take additional time to authenticate the session.

keyed-sha-1—Keyed Secure Hash Algorithm I for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed SHA uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. The key is not carried within the packets. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than the last sequence number received.

meticulous-keyed-sha-1—Meticulous keyed Secure Hash Algorithm I. This method works in the same manner as keyed SHA, but the sequence number is updated with every packet. Although more secure than keyed SHA and simple passwords, this method can take additional time to authenticate the session.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
---------------------------------	---

Related Documentation	<ul style="list-style-type: none">• Example: Configuring BFD Authentication for Securing Static Routes on page 3255• Example: Configuring BGP Route Authentication on page 3897• Example: Configuring EBGp Multihop Sessions on page 3777• authentication on page 3974• bfd-liveness-detection on page 3981• key-chain on page 4018• loose-check on page 4030
------------------------------	---

apply-groups

Syntax	<code>apply-groups [<i>group-names</i>];</code>
Hierarchy Level	All hierarchy levels
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>Apply a configuration group to a specific hierarchy level in a configuration, to have a configuration inherit the statements in the configuration group.</p> <p>You can specify more than one group name. You must list them in order of inheritance priority. The configuration data in the first group takes priority over the data in subsequent groups.</p>
Options	<i>group-names</i> —One or more names specified in the groups statement.
Required Privilege Level	configure—To enter configuration mode, but other required privilege levels depend on where the statement is located in the configuration hierarchy.
Related Documentation	<ul style="list-style-type: none">• <i>Applying a Junos OS Configuration Group</i>• <i>groups</i>

apply-groups-except

Syntax	<code>apply-groups-except [<i>group-names</i>];</code>
Hierarchy Level	All hierarchy levels except the top level
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Disable inheritance of a configuration group.
Options	<i>group-names</i> —One or more names specified in the groups statement.
Required Privilege Level	configure—To enter configuration mode, but other required privilege levels depend on where the statement is located in the configuration hierarchy.
Related Documentation	<ul style="list-style-type: none">• <i>groups</i>• <i>Disabling Inheritance of a Junos OS Configuration Group</i>

authentication (BGP BFD Liveness Detection)

Syntax	<pre>authentication { algorithm <i>algorithm-name</i>; key-chain <i>key-chain-name</i>; loose-check ; }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection], [edit protocols bgp bgp bfd-liveness-detection], [edit protocols bgp group <i>group-name</i> bfd-liveness-detection], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection]</pre>
Release Information	<p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for BFD authentication introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Specify the router and route authentication to mitigate the risk of being attacked by a machine or router that has been configured to share incorrect routing information with another router. Router and route authentication enables routers to share information only if they can verify that they are talking to a trusted source, based on a password (key). In this method, a hashed key is sent along with the route being sent to another router. The receiving router compares the sent key to its own configured key. If they are the same, the receiving router accepts the route.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Example: Configuring BFD for Static Routes for Faster Network Failure Detection on page 3241• Example: Configuring BFD Authentication for Securing Static Routes on page 3255• Example: Configuring BGP Route Authentication on page 3897

- [algorithm on page 3971](#)
- [bfd-liveness-detection on page 3981](#)
- [key-chain on page 4018](#)
- [loose-check on page 4030](#)

authentication-algorithm

Syntax authentication-algorithm *algorithm*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols bgp],
 [edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
 [edit logical-systems *logical-system-name* protocols bgp **group** *group-name* neighbor *address*],
 [edit logical-systems *logical-system-name* protocols ldp session *session-address*],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp session *session-address*],
 [edit logical-systems *logical-system-name* routing-options **bmp**],
 [edit logical-systems *logical-system-name* routing-options bmp **station** *station-name*],
 [edit protocols bgp],
 [edit protocols bgp **group** *group-name*],
 [edit protocols bgp group *group-name* **neighbor** *address*],
 [edit protocols ldp session *session-address*],
 [edit routing-instances *routing-instance-name* protocols bgp],
 [edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
 [edit routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*],
 [edit routing-instances *routing-instance-name* protocols ldp session *session-address*],
 [edit routing-options **bmp**],
 [edit routing-options bmp **station** *station-name*]

Release Information Statement introduced in Junos OS Release 7.6.
 Statement introduced for BGP in Junos OS Release 8.0.
 Statement introduced in Junos OS Release 9.0 for EX Series switches.
 Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
 Statement introduced for BMP in Junos OS Release 13.2X51-D15 for the QFX Series.
 Statement introduced for BMP in Junos OS Release 13.3.
 Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure an authentication algorithm type.



NOTE: Keep the following points in mind when you configure the authentication algorithm in an IPsec proposal:

- When both ends of an IPsec VPN tunnel contain the same IKE proposal but different IPsec proposals, an error occurs and the tunnel is not established in this scenario. For example, if one end of the tunnel contains router 1 configured with the authentication algorithm as hmac-sha-256-128 and the other end of the tunnel contains router 2 configured with the authentication algorithm as hmac-md5-96, the VPN tunnel is not established.

- When both ends of an IPsec VPN tunnel contain the same IKE proposal but different IPsec proposals, and when one end of the tunnel contains two IPsec proposals to check whether a less secure algorithm is selected or not, an error occurs and the tunnel is not established. For example, if you configure two authentication algorithms for an IPsec proposal as `hmac-sha-256-128` and `hmac-md5-96` on one end of the tunnel, router 1, and if you configure the algorithm for an IPsec proposal as `hmac-md5-96` on the other end of the tunnel, router 2, the tunnel is not established and the number of proposals mismatch.
 - When you configure two IPsec proposals at both ends of a tunnel, such as the `authentication-algorithm hmac-sha-256-128` and `authentication-algorithm hmac-md5-96` statements at the `[edit services ipsec-vpn ipsec proposal proposal-name]` hierarchy level on one of the tunnel, router 1 (with the algorithms in two successive statements to specify the order), and the `authentication-algorithm hmac-md5-96` and `authentication-algorithm hmac-sha-256-128` statements at the `[edit services ipsec-vpn ipsec proposal proposal-name]` hierarchy level on one of the tunnel, router 2 (with the algorithms in two successive statements to specify the order, which is the reverse order of router 1), the tunnel is established in this combination as expected because the number of proposals is the same on both ends and they contain the same set of algorithms. However, the authentication algorithm selected is `hmac-md5-96` and not the stronger algorithm of `hmac-sha-256-128`. This method of selection of the algorithm occurs because the first matching proposal is selected. Also, for a default proposal, regardless of whether the router supports the Advanced Encryption Standard (AES) encryption algorithm, the `3des-cbc` algorithm is chosen and not the `aes-cfb` algorithm, which is because of the first algorithm in the default proposal being selected. In the sample scenario described here, on router 2, if you reverse the order of the algorithm configuration in the proposal so that it is the same order as the one specified on router 1, `hmac-sha-256-128` is selected as the authentication method.
 - You must be aware of the order of proposals in an IPsec policy at the time of configuration if you want the matching of proposals to happen in a certain order of preference, such as the strongest algorithm to be considered first when a match is made when both policies from the two peers have a proposal.
-

Options *algorithm*—Specify one of the following types of authentication algorithms:

- **aes-128-cmac-96**—Cipher-based message authentication code (AES128, 96 bits).
- **hmac-sha-1-96**—Hash-based message authentication code (SHA1, 96 bits).
- **md5**—Message digest 5.

Default: hmac-sha-1-96



NOTE: The default is not displayed in the output of the `show bgp bmp` command unless a key or key-chain is also configured.

Required Privilege Level routing—To view this statement in the configuration.
 routing-control—To add this statement to the configuration.

Related Documentation • [Example: Configuring Route Authentication for BGP on page 3898](#)
 • [Configuring BGP Monitoring Protocol Version 3 on page 3647](#)

authentication-key (Protocols BGP and BMP)

Syntax	authentication-key <i>key</i> ;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options bmp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options bmp station <i>station-name</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-options bmp],</p> <p>[edit routing-options bmp station <i>station-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced for BMP in Junos OS Release 13.2X51-D15 for the QFX Series.</p> <p>Statement introduced for BMP version 3 in Junos OS Release 13.3.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Configure an MD5 authentication key (password). Neighboring routing devices use the same password to verify the authenticity of BGP packets sent from this system.
Options	key —Authentication password. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Route Authentication for BGP on page 3898 • Configuring BGP Monitoring Protocol Version 3 on page 3647

authentication-key-chain (Protocols BGP and BMP)

Syntax	<code>authentication-key-chain <i>key-chain</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options bmp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options bmp station <i>station-name</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-options bmp],</p> <p>[edit routing-options bmp station <i>station-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.0.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced for BMP in Junos OS Release 13.2X51-D15 for the QFX Series.</p> <p>Statement introduced for BMP in Junos OS Release 13.3.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Apply and enable an authentication keychain to the routing device. Note that the referenced key chain must be defined. When configuring the authentication key update feature for BGP, you cannot commit the 0.0.0.0/allow statement with authentication keys or key chains. The CLI issues a warning and fails to commit the configuration.
Options	key-chain —Authentication keychain name. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Route Authentication for BGP on page 3898 • Example: Configuring BFD Authentication for Securing Static Routes on page 3255 • Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols • Configuring BGP Monitoring Protocol Version 3 on page 3647

bfd-liveness-detection (Protocols BGP)

Syntax

```

bfd-liveness-detection {
    authentication {
        algorithm algorithm-name;
        key-chain key-chain-name;
        loose-check;
    }
    detection-time {
        threshold milliseconds;
    }
    hold-down-interval milliseconds;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    session-mode (automatic | multihop | single-hop);
    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    version (1 | automatic);
}

```

Hierarchy Level

```

[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
bgp group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
bgp group group-name neighbor address],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor
address]

```

Release Information

Statement introduced in Junos OS Release 8.1.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

detection-time threshold and **transmit-interval threshold** options introduced in Junos OS Release 8.2

Support for logical routers introduced in Junos OS Release 8.3.

Support for IBGP and multihop EBGP sessions introduced in Junos OS Release 8.3.

holddown-interval statement introduced in Junos OS Release 8.5. You can configure this statement only for EBGP peers at the **[edit protocols bgp group *group-name* neighbor *address*]** hierarchy level.

no-adaptation statement introduced in Junos OS Release 9.0.

Support for BFD authentication introduced in Junos OS Release 9.6.

Support for BFD on IPv6 interfaces with BGP introduced in Junos OS Release 11.2.
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure bidirectional failure detection (BFD) timers and authentication for BGP.

For IBGP and multihop EBGP support, configure the **bfd-liveness-detection** statement at the global **[edit bgp protocols]** hierarchy level. You can also configure IBGP and multihop support for a routing instance or a logical system.

Options **authentication algorithm** *algorithm-name* (Optional)—Configure the algorithm used to authenticate the specified BFD session: **simple-password**, **keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**.

authentication key-chain *key-chain-name* (Optional)—Associate a security key with the specified BFD session using the name of the security keychain. The keychain name must match one of the keychains configured in the **authentication-key-chains key-chain** statement at the **[edit security]** hierarchy level.

authentication loose-check—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication may not be configured at both ends of the BFD session.

detection-time threshold *milliseconds* (Optional)—Configure a threshold. When the BFD session detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

holddown-interval *milliseconds* (Optional)—Configure an interval specifying how long a BFD session must remain up before a state change notification is sent. When you configure the hold-down interval for the BFD protocol for EBGp, the BFD session is unaware of the BGP session during this time. In this case, if the BGP session goes down during the configured hold-down interval, BFD already assumes it is down and does not send a state change notification. The **holddown-interval** statement is supported only for EBGp peers at the **[edit protocols bgp group group-name neighbor address]** hierarchy level. If the BFD session goes down and then comes back up during the configured hold-down interval, the timer is restarted. You must configure the hold-down interval on both EBGp peers. If you configure the hold-down interval for a multihop EBGp session, you must also configure a local IP address by including the **local-address** statement at the **[edit protocols bgp group group-name]** hierarchy level.

Range: 0 through 255,000

Default: 0

minimum-interval *milliseconds* (Required)—Configure the minimum intervals at which the local routing device transmits hello packets and then expects to receive a reply from a neighbor with which it has established a BFD session. This value represents the minimum interval at which the local routing device transmits hello packets as well as the minimum interval that the routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a value in the range from 1 through 255,000 milliseconds. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately (using the **minimum-receive-interval** and **transmit-interval** statements).

Range: 1 through 255,000

minimum-receive-interval *milliseconds* (Optional)—Configure only the minimum interval at which the local routing device expects to receive a reply from a neighbor with which it has established a BFD session.

Range: 1 through 255,000

multiplier *number* (Optional)—Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

Range: 1 through 255

Default: 3

no-adaptation (Optional)—Configure BFD sessions not to adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable to not to have BFD adaptation enabled in your network.

transmit-interval threshold *milliseconds* (Optional)—Configure a threshold. When the BFD session transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

transmit-interval minimum-interval *milliseconds* (Optional)—Configure only the minimum interval at which the local routing device transmits hello packets to a neighbor with which it has established a BFD session.

Range: 1 through 255,000

version (Optional)—Configure the BFD version to detect.

Range: 1 or **automatic** (autodetect the BFD version)

Default: **automatic**

The remaining statements are explained separately.


Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
---------------------------------	---

Related Documentation	<ul style="list-style-type: none">• Example: Configuring BFD for Static Routes for Faster Network Failure Detection on page 3241• Example: Configuring BFD Authentication for Securing Static Routes on page 3255• Example: Configuring BFD on Internal BGP Peer Sessions on page 3812• Example: Configuring BFD Authentication for BGP on page 3822• Understanding BFD for BGP on page 3811
------------------------------	--

bgp

Syntax	<code>bgp { ... }</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit protocols], [edit routing-instances <i>routing-instance-name</i> protocols]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Enable BGP on the routing device or for a routing instance.
Default	BGP is disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	

bgp-orf-cisco-mode

Syntax	<code>bgp-orf-cisco-mode;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options outbound-route-filter],</p> <p>[edit protocols bgp outbound-route-filter],</p> <p>[edit protocols bgp group <i>group-name</i> outbound-route-filter],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> outbound-route-filter],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp outbound-route-filter],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> outbound-route-filter],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> outbound-route-filter],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options outbound-route-filter],</p> <p>[edit routing-options outbound-route-filter]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.2.</p> <p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Support for the BGP group and neighbor hierarchy levels introduced in Junos OS Release 9.2.</p> <p>Support for the BGP group and neighbor hierarchy levels introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Enable interoperability with routing devices that use the vendor-specific outbound route filter compatibility code of 130 and code type of 128.
	<p> NOTE: To enable interoperability for all BGP peers configured on the routing device, include the statement at the [edit routing-options outbound-route-filter] hierarchy level.</p>
Default	Disabled

Required Privilege routing—To view this statement in the configuration.
Level routing-control—To add this statement to the configuration.

Related Documentation • [Example: Configuring BGP Prefix-Based Outbound Route Filtering on page 3773](#)

cluster

Syntax	<code>cluster <i>cluster-identifier</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Specify the cluster identifier to be used by the route reflector cluster in an internal BGP group.



CAUTION:

If you configure both route reflection and VPNs on the same routing device, the following modifications to the route reflection configuration cause current BGP sessions to be reset:

- Adding a cluster ID—If a BGP session shares the same AS number with the group where you add the cluster ID, all BGP sessions are reset regardless of whether the BGP sessions are contained in the same group.
- Creating a new route reflector—If you have an IBGP group with an AS number and create a new route reflector group with the same AS number, all BGP sessions in the IBGP group and the new route reflector group are reset.



NOTE: If you change the address family specified in the [edit protocols bgp family] hierarchy level, all current BGP sessions on the routing device are dropped and then reestablished.

Options	<i>cluster-identifier</i> —4-byte number (such as an IPv4 address).
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring BGP Route Reflectors on page 3873 • Understanding External BGP Peering Sessions on page 3601 • no-client-reflect on page 4053

connection-mode

Syntax	connection-mode (active passive);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options bmp], [edit logical-systems <i>logical-system-name</i> routing-options bmp station station-name], [edit routing-options bmp], [edit routing-options bmp station station-name]
Release Information	Statement introduced for BMP in Junos OS Release 13.2X51-D15 for the QFX Series. Statement introduced for BMP in Junos OS Release 13.3.
Description	Specifies whether the BMP station connection is active or passive .
Options	<p>active—BMP initiates the connection to the BMP station.</p> <p>passive—BMP does not initiate a connection the BMP station. However, it does listen for a connection request from active BMP stations and will connect if a station is available.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring BGP Monitoring Protocol Version 3 on page 3647

damping (Protocols BGP)

Syntax	damping;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp family <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp family <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> family <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp family <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp family <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family <i>family</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> family <i>family</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> family <i>family</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp family <i>family</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family <i>family</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family <i>family</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for flap damping at the address family level introduced in Junos OS Release 12.2.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Enable route flap damping. BGP route flapping describes the situation in which BGP systems send an excessive number of update messages to advertise network reachability

information. Flap damping reduces the number of update messages sent between BGP peers, thereby reducing the load on these peers, without adversely affecting the route convergence time for stable routes.

You typically apply flap damping to external BGP (EBGP) routes (that is, to routes in different ASs). You can also apply it within a confederation, between confederation member ASs. Because routing consistency within an AS is important, do not apply flap damping to internal BGP (IBGP) routes. (If you do, it is ignored.) The exception to this rule is when flap damping is applied at the address family level. When you apply flap damping at the address family level, it works for both IBGP and EBGP.

Default Flap damping is disabled on the routing device.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [Examples: Configuring BGP Flap Damping on page 3926](#)
- [Example: Configuring BGP Route Flap Damping Based on the MBGP MVPN Address Family on page 3936](#)

description (Protocols BGP)

Syntax	<code>description text-description;</code>
Hierarchy Level	<code>[edit logical-systems logical-system-name protocols bgp],</code> <code>[edit logical-systems logical-system-name protocols bgp group group-name],</code> <code>[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],</code> <code>[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],</code> <code>[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp group group-name],</code> <code>[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp group group-name neighbor address],</code> <code>[edit protocols bgp],</code> <code>[edit protocols bgp group group-name],</code> <code>[edit protocols bgp group group-name neighbor address],</code> <code>[edit routing-instances routing-instance-name protocols bgp],</code> <code>[edit routing-instances routing-instance-name protocols bgp group group-name],</code> <code>[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Provide a description of the global, group, or neighbor configuration. If the text includes one or more spaces, enclose it in quotation marks (" "). The text is displayed in the output of the show command and has no effect on the configuration.
Options	<i>text-description</i> —Text description of the configuration. It is limited to 255 characters.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>BGP Feature Guide for Routing Devices</i>

detection-time (BFD Liveness Detection)

Syntax	<pre> detection-time { threshold milliseconds; } </pre>
Hierarchy Level	<pre> [edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection], [edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols l2vpn oam bfd-liveness-detection], [edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam bfd-liveness-detection], [edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam bfd-liveness-detection], [edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls oam bfd-liveness-detection], [edit protocols bgp bfd-liveness-detection], [edit protocols bgp group <i>group-name</i> bfd-liveness-detection], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bgp bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols l2vpn oam bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols vpls oam bfd-liveness-detection] </pre>
Release Information	<p>Statement introduced in Junos OS Release 8.2.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for BFD authentication introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2 for Layer 2 VPNs and VPLS.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Enable BFD failure detection. The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds.</p>

A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the **clear bfd adaptation** command to return BFD interval timers to their configured values. The **clear bfd adaptation** command is hitless, meaning that the command does not affect traffic flow on the routing device.

The remaining statement is explained separately.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring BFD for Layer 2 VPN and VPLS• Example: Configuring BFD for BGP on page 3811• bfd-liveness-detection on page 3981• threshold on page 4077

disable (Protocols BGP)

Syntax	disable;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Disable BGP on the system.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

disable (BGP Graceful Restart)

Syntax	disable;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols bgp graceful-restart], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> graceful-restart], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> graceful-restart], [edit protocols bgp graceful-restart], [edit protocols bgp group <i>group-name</i> graceful-restart], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> graceful-restart]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Disable graceful restart for BGP. Graceful restart allows a routing device undergoing a restart to inform its adjacent neighbors and peers of its condition.



NOTE: When you disable graceful restart at one level in the configuration statement hierarchy, it is also disabled at lower levels in the same hierarchy. For example, if you disable graceful restart at the [edit protocols bgp group *group-name*] hierarchy level, it is disabled for all the peers in the group. Therefore, if you want to enable graceful restart for some peers in a group and disable it for others, enable graceful restart at the [edit protocols bgp group *group-name*] hierarchy level and disable graceful restart for each peer at the [edit protocols bgp group *group-name* neighbor *address*] hierarchy level.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Graceful Restart Options for BGP on page 2581 • graceful-restart on page 2662 • <i>restart-time</i> • <i>stale-routes-time</i>

export (Protocols BGP)

Syntax	<code>export [<i>policy-names</i>];</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols bgp],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i></code> <code> <i>neighbor address</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp <i>group group-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp <i>group group-name</i> <i>neighbor address</i>],</code> <code>[edit protocols bgp],</code> <code>[edit protocols bgp <i>group group-name</i>],</code> <code>[edit protocols bgp <i>group group-name</i> <i>neighbor address</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i></code> <code> <i>neighbor address</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Apply one or more policies to routes being exported from the routing table into BGP.</p> <p>If you specify more than one policy, they are evaluated in the order specified, from left to right, and the first matching filter is applied to the route. If no routes match the filters, the routing table exports into BGP only the routes that it learned from BGP. If an action specified in one of the policies manipulates a route characteristic, the policy framework software carries the new route characteristic forward during the evaluation of the remaining policies. For example, if the action specified in the first policy of a chain sets a route's metric to 500, this route matches the criterion of metric 500 defined in the next policy.</p>
Options	<i>policy-names</i> —Name of one or more policies.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Routing Policies to Control BGP Route Advertisements on page 3769• <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i>• import on page 4012

family

```

Syntax  family {
    (inet | inet6 | inet-vpn | inet6-vpn | iso-vpn) {
        (any | flow | labeled-unicast | multicast | unicast) {
            accepted-prefix-limit {
                maximum number;
                teardown <percentage-threshold> idle-timeout (forever | minutes);
            }
            add-path {
                send {
                    path-count number;
                    prefix-policy [ policy-names ];
                }
                receive;
            }
            algp [disable];
            loops number;
            prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            protection;
            rib-group group-name;
            topology name {
                community {
                    target identifier;
                }
            }
            flow {
                no-install;
                no-validate policy-name;
            }
            labeled-unicast {
                accepted-prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                }
                aggregate-label {
                    community community-name;
                }
                explicit-null {
                    connected-only;
                }
                prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                }
                resolve-vpn;
                rib inet.3;
                rib-group group-name;
                traffic-statistics {
                    file filename <world-readable | no-world-readable>;
                    interval seconds;
                }
            }
        }
    }
}

```

```
    }
  }
}
route-target {
  accepted-prefix-limit {
    maximum number;
    proxy-generate <route-target-policy route-target-policy-name>;
    teardown <percentage> <idle-timeout (forever | minutes)>;
  }
  advertise-default;
  external-paths number;
  prefix-limit {
    maximum number;
    teardown <percentage> <idle-timeout (forever | minutes)>;
  }
}
(inet-mdt | inet-mvpn | inet6-mvpn | l2vpn) {
  signaling {
    accepted-prefix-limit {
      maximum number;
      teardown <percentage-threshold> idle-timeout (forever | minutes);
    }
  }
  add-path {
    send {
      path-count number;
      prefix-policy [ policy-names ];
    }
    receive;
  }
  aigp [disable];
  damping;
  loops number;
  prefix-limit {
    maximum number;
    teardown <percentage> <idle-timeout (forever | minutes)>;
  }
  rib-group group-name;
}
}
traffic-engineering;
}
```

Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> <p>inet-mvpn and inet6-mvpn statements introduced in Junos OS Release 8.4.</p> <p>inet-mdt statement introduced in Junos OS Release 9.4.</p> <p>Support for the loops statement introduced in Junos OS Release 9.6.</p> <p>evpn statement introduced in Junos OS Release 13.2.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> <p>traffic-engineering statement introduced in Junos OS Release 14.2.</p>
Description	<p>Enable multiprotocol BGP (MP-BGP) by configuring BGP to carry network layer reachability information (NLRI) for address families other than unicast IPv4, to specify MP-BGP to carry NLRI for the IPv6 address family, or to carry NLRI for VPNs.</p>

- Options**
- any**—Configure the family type to be both unicast and multicast.
 - evpn**—Configure NLRI parameters for Ethernet VPNs (EVPNs).
 - inet**—Configure NLRI parameters for IPv4.
 - inet6**—Configure NLRI parameters for IPv6.
 - inet-mdt**—Configure NLRI parameters for the multicast distribution tree (MDT) subaddress family identifier (SAFI) for IPv4 traffic in Layer 3 VPNs.
 - inet-mvpn**—Configure NLRI parameters for IPv4 for multicast VPNs.
 - inet6-mvpn**—Configure NLRI parameters for IPv6 for multicast VPNs.
 - inet-vpn**—Configure NLRI parameters for IPv4 for Layer 3 VPNs.
 - inet6-vpn**—Configure NLRI parameters for IPv6 for Layer 3 VPNs.
 - inet6-vpn**—Configure NLRI parameters for IPv6 for Layer 3 VPNs.
 - iso-vpn**—Configure NLRI parameters for IS-IS for Layer 3 VPNs.
 - l2vpn**—Configure NLRI parameters for IPv4 for MPLS-based Layer 2 VPNs and VPLS.
 - labeled-unicast**—Configure the family type to be labeled-unicast. This means that the BGP peers are being used only to carry the unicast routes that are being used by labeled-unicast for resolving the labeled-unicast routes. This statement is supported only with **inet** and **inet6**.
 - multicast**—Configure the family type to be multicast. This means that the BGP peers are being used only to carry the unicast routes that are being used by multicast for resolving the multicast routes.
 - unicast**—Configure the family type to be unicast. This means that the BGP peers only carry the unicast routes that are being used for unicast forwarding purposes. The default family type is **unicast**.


The remaining statements are explained separately.

Required Privilege Level

- routing—To view this statement in the configuration.
- routing-control—To add this statement to the configuration.

- Related Documentation**
- *Configuring IBGP Sessions Between PE Routers in VPNs*
 - *Understanding Multiprotocol BGP*
 - [autonomous-system on page 3276](#)
 - [local-as on page 4023](#)
 - *Understanding Multiprotocol BGP*

graceful-restart (Protocols BGP)

Syntax	<pre> graceful-restart { disable; restart-time seconds; stale-routes-time seconds; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Configure graceful restart for BGP. Graceful restart allows a routing device undergoing a restart to inform its adjacent neighbors and peers of its condition. Graceful restart is disabled by default. However, helper mode, the ability to assist a neighboring router attempting a graceful restart, is enabled by default.</p> <p>To configure the duration of the BGP graceful restart period, include the restart-time statement at the [edit protocols bgp graceful-restart] hierarchy level. To set the length of time the router waits to receive messages from restarting neighbors before declaring them down, include the stale-routes-time statement at the [edit protocols bgp graceful-restart] hierarchy level.</p> <hr/> <div>  <p>NOTE: If you configure graceful restart after a BGP session has been established, the BGP session restarts and the peers negotiate graceful restart capabilities.</p> </div> <hr/> <p>Configure graceful restart globally at the [edit routing-options] or [edit routing-instances <i>instance-name</i> routing-options] hierarchy level to enable the feature. You cannot enable graceful restart for specific protocols unless graceful restart is also enabled globally. You can, optionally, modify the global settings at the individual protocol level.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Graceful Restart Options for BGP on page 2581 • Configuring Graceful Restart for QFabric Systems

- *Junos OS High Availability Library for Routing Devices*

group (Protocols BGP)

```
Syntax  group group-name {
    advertise-inactive;
    allow [ network/mask-length ];
    authentication-key key;
    cluster cluster-identifier;
    damping;
    description text-description;
    enforce-first-as;
    export [ policy-names ];
    family {
        (inet | inet6 | inet-vpn | inet6-vpn | l2-vpn) {
            (any | multicast | unicast | signaling) {
                accepted-prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                }
            }
            add-path {
                send {
                    path-count number;
                    prefix-policy [ policy-names ];
                }
                receive;
            }
            aigp [disable];
            damping;
            prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            rib-group group-name;
            topology name {
                community {
                    target identifier;
                }
            }
        }
    }
    flow {
        no-validate policy-name;
    }
    labeled-unicast {
        accepted-prefix-limit {
            maximum number;
            teardown <percentage> <idle-timeout (forever | minutes)>;
        }
        explicit-null {
            connected-only;
        }
        prefix-limit {
            maximum number;
            teardown <percentage> <idle-timeout (forever | minutes)>;
        }
    }
    resolve-vpn;
}
```

```

        rib inet.3;
        rib-group group-name;
    }
}
route-target {
    accepted-prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    advertise-default;
    external-paths number;
    prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
    }
}
}
graceful-restart {
    long-lived {
        receiver {
            enable;
            disable;
        }
        advertise-to-non-llgr-neighbor {
            omit-no-export;
        }
    }
}
graceful-restart {
    long-lived {
        disable-notification-flag;
        disable-notification-extensions {
            omit-no-export;
        }
        forwarding-state-bit (from-fib | set); /* Configurable to be common for all address
        families */
        forwarding-state-bit (as-rr-client | from-fib); /* Configurable for each address family
        */
        restarter {
            disable;
            stale-time interval;
        }
    }
}
hold-time seconds;
import [ policy-names ];
ipsec-sa ipsec-sa;
keep (all | none);
local-address address;
local-as autonomous-system <private>;
local-preference local-preference;
log-updown;
metric-out metric;
multihop <ttl-value>;
multipath {
    multiple-as;

```

```

}
mvpn-iana-rt-import;
no-aggregator-id;
no-client-reflect;
out-delay seconds;
passive;
peer-as autonomous-system;
preference preference;
remove-private;
tcp-aggressive-transmission;
tcp-mss segment-size;
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
type type;
neighbor address {
    ... peer-specific-options ...
}
}

```

Hierarchy Level [edit logical-systems *logical-system-name* protocols bgp],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols
 bgp],
 [edit protocols bgp],
 [edit routing-instances *routing-instance-name* protocols bgp]

Release Information Statement introduced before Junos OS Release 7.4.
 Statement introduced in Junos OS Release 9.0 for EX Series switches.
 Statement introduced in Junos OS Release 11.3 for the QFX Series.
 Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Define a BGP peer group. BGP peer groups share a common type, peer autonomous system (AS) number, and cluster ID, if present. To configure multiple BGP groups, include multiple **group** statements.

By default, the group's options are identical to the global BGP options. To override the global options, include group-specific options within the **group** statement.

The **group** statement is one of the statements you must include in the configuration to run BGP on the routing device.

Each group must contain at least one peer.

Options *group-name*—Name of the BGP group.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
 routing-control—To add this statement to the configuration.

Related Documentation • [Understanding BGP on page 3593](#)

hold-down

Syntax	<pre>hold-down { seconds; flaps number; period seconds; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options bmp], [edit logical-systems <i>logical-system-name</i> routing-options bmp station station-name], [edit routing-options bmp], [edit routing-options bmp station station-name]
Release Information	Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series. Statement introduced in Junos OS Release 13.3. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>If the connection to a BMP station flaps and the hold-down statement is configured, the station is prevented from reconnecting to the device for the specified period of time. A flap is when the TCP session unexpectedly switches from established to non-established. If you alter the configuration of the hold-down statement, the hold down timer and flap counter are reset.</p> <p>You can effectively disable the hold-down statement by setting the flaps option to 10 and the period option to 30 seconds.</p>
Options	<p>seconds—Specify the time in seconds to wait before allowing the BMP station to reconnect to the device.</p> <p>Default: 600 seconds</p> <p>Range: 30 through 65,535 seconds</p> <p>flaps number—Specify the number of BMP station flaps allowed before terminating the connection to the BMP station and triggering the hold down timer.</p> <p>Default: 3 flaps</p> <p>Range: 2 to 10 flaps</p> <p>period seconds—Specify the time in seconds for the BGP station flaps (specified using the flaps option) to occur before triggering the hold down timer. Every time a flap occurs, the number of flaps in the last time period is checked to see if the criteria is met.</p> <p>For example, if you defined the period as 60 seconds and the flaps as 4 and the BGP station flaps just 2 times in a 60 second period, the hold down timer would not be triggered. However, if the BGP station flaps 4 times in a 60 second period, the hold down timer would be triggered.</p> <p>Default: 300 seconds</p> <p>Range: 30 through 65,535 seconds</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

- Related Documentation**
- [Configuring BGP Monitoring Protocol Version 3 on page 3647](#)

hold-down-interval (BGP BFD Liveness Detection)

Syntax	<code>holddown-interval milliseconds;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i></code> <code> bfd-liveness-detection],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i></code> <code> bfd-liveness-detection],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp bfd-liveness-detection],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp group <i>group-name</i> bfd-liveness-detection],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</code> <code>[edit protocols bgp bfd-liveness-detection],</code> <code>[edit protocols bgp group <i>group-name</i> bfd-liveness-detection],</code> <code>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bgp bfd-liveness-detection],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i></code> <code> bfd-liveness-detection],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor</code> <code> <i>address</i> bfd-liveness-detection]</code>
Release Information	Statement introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for BFD authentication introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Configure an interval specifying how long a BFD session must remain up before a state change notification is sent.</p> <p>When you configure the hold-down interval for the BFD protocol for EBGp, the BFD session is unaware of the BGP session during this time. In this case, if the BGP session goes down during the configured hold-down interval, BFD already assumes the BGP session is down and does not send a state change notification. The holddown-interval statement is supported only for EBGp peers at the [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>] hierarchy level. If the BFD session goes down and then comes back up during the configured hold-down interval, the timer is restarted. You must configure the hold-down interval on both EBGp peers. If you configure the hold-down interval for a multihop EBGp session, you must also configure a local IP address by including the local-address statement at the [edit protocols bgp group <i>group-name</i>] hierarchy level.</p>
Options	milliseconds —Specify the hold-down interval value. Range: 0 through 255,000 Default: 0
Required Privilege Level	routing —To view this statement in the configuration. routing-control —To add this statement to the configuration.

**Related
Documentation**

- [Example: Configuring BFD for Static Routes for Faster Network Failure Detection on page 3241](#)
- [bfd-liveness-detection on page 3981](#)

hold-time (Protocols BGP)

Syntax	<code>hold-time seconds;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols bgp],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i></code> <code> <i>neighbor address</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp <i>group group-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp <i>group group-name neighbor address</i>],</code> <code>[edit protocols bgp],</code> <code>[edit protocols bgp <i>group group-name</i>],</code> <code>[edit protocols bgp <i>group group-name neighbor address</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i></code> <code> <i>neighbor address</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Specify the hold-time value to use when negotiating a connection with the peer. The hold-time value is advertised in open packets and indicates to the peer the length of time that it should consider the sender valid. If the peer does not receive a keepalive, update, or notification message within the specified hold time, the BGP connection to the peer is closed and routing devices through that peer become unavailable.</p> <p>The hold time is three times the interval at which keepalive messages are sent.</p> <p>BGP on the local routing device uses the smaller of either the local hold-time value or the peer's hold-time value received in the open message as the hold time for the BGP connection between the two peers.</p> <p>Starting in Junos OS Release 12.3, the BGP hold-time value can be zero (0). This implies that the speaker does not expect keepalive messages from its peer to maintain the BGP session. When negotiating between two peers, if one side requests a nonzero hold time and the other requests a zero hold time, the negotiation settles on the nonzero value and keepalive intervals are determined accordingly. Both sides must be set to zero for keepalive messages to stop being sent.</p>
Options	<i>seconds</i> —Hold time. Range: 10 through 65,535 seconds (or 0 for infinite hold time) Default: 90 seconds



TIP: When you set a hold-time value of 1 through 19 seconds, we recommend that you also configure the BGP `precision-timers` statement. The `precision-timers` statement ensures that if scheduler slip messages occur, the routing device continues to send keepalive messages. When the `precision-timers` statement is included, keepalive message generation is performed in a dedicated kernel thread, which helps to prevent BGP session flaps.

Required Privilege Level	routing—To view this statement in the configuration.
	routing-control—To add this statement to the configuration.
Related Documentation	• BGP Messages Overview on page 3597
	• <i>precision-timers</i>

import

Syntax	<code>import [<i>policy-names</i>];</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Apply one or more routing policies to routes being imported into the Junos OS routing table from BGP.</p> <p>If you specify more than one policy, they are evaluated in the order specified, from left to right, and the first matching filter is applied to the route. If no match is found, BGP places into the routing table only those routes that were learned from BGP routing devices. The policy framework software evaluates the routing policies in a chain sequentially. If an action specified in one of the policies manipulates a route characteristic, the policy framework software carries the new route characteristic forward during the evaluation of the remaining policies. For example, if the action specified in the first policy of a chain sets a route's metric to 500, this route matches the criterion of metric 500 defined in the next policy.</p> <p>It is also important to understand that in Junos OS, although an import policy (inbound route filter) might reject a route, not use it for traffic forwarding, and not include it in an advertisement to other peers, the router retains these routes as hidden routes. These hidden routes are not available for policy or routing purposes. However, they do occupy memory space on the router. A service provider filtering routes to control the amount of information being kept in memory and processed by a router might want the router to entirely drop the routes being rejected by the import policy.</p> <p>Hidden routes can be viewed by using the show route receive-protocol bgp neighbor-address hidden command. The hidden routes can then be retained or dropped from the routing</p>

table by configuring the **keep all | none** statement at the **[edit protocols bgp]** or **[edit protocols bgp group *group-name*]** hierarchy level.

The rules of BGP route retention are as follows:

- By default, all routes learned from BGP are retained, except those where the AS path is looped. (The AS path includes the local AS.)
- By configuring the **keep all** statement, all routes learned from BGP are retained, even those with the local AS in the AS path.
- By configuring the **keep none** statement, all routes received are discarded. When this statement is configured and the inbound policy changes, Junos OS re-advertises all the routes advertised by the peer.

Options *policy-names*—Name of one or more policies.

Required Privilege Level routing—To view this statement in the configuration.
 routing-control—To add this statement to the configuration.

Related Documentation

- [Example: Configuring BGP Interactions with IGP's on page 3765](#)
- [Configuring Routing Policies to Control BGP Route Advertisements on page 3769](#)
- [Understanding Routing Policies on page 3765](#)
- [export on page 3996](#)

include-mp-next-hop

Syntax	include-mp-next-hop;
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i> <i>neighbor address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name neighbor address</i>], [edit protocols bgp], [edit protocols bgp <i>group group-name</i>], [edit protocols bgp <i>group group-name neighbor address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i> <i>neighbor address</i>]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Enable multiprotocol updates to contain next-hop reachability information.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring IPv6 BGP Routes over IPv4 Transport</i>• <i>Enabling Layer 2 VPN and VPLS Signaling</i>• <i>Understanding Multiprotocol BGP</i>

initiation-message

Syntax	<code>initiation-message text;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-options bmp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options bmp station station-name],</p> <p>[edit routing-options bmp],</p> <p>[edit routing-options bmp station station-name]</p>
Release Information	<p>Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.3.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>(Optional) Allows you to specify an initiation message for a type 0 TLV to be sent to the BMP station. The message is transmitted when a BMP station establishes a connection to the device. You can provide some information to the BMP station system administrator (for example, a contact phone number). The initiation message includes a type 1 TLV containing the SNMP sysDescr value specified in RFC 1213 <i>Management Information Base for Network Management of TCP/IP-based internets: MIB-II</i> and a type 2 TLV containing the SNMP sysName value also from RFC 1213. The string in the initiation-message message is UTF-8.</p> <p>The normal time for sending an initiation message is when the BMP session is first established. However, an initiation message change also triggers the transmission of an initiation message to current BMP sessions.</p> <p>Another event that triggers the transmission of an initiation message is when you change in the sysName or sysDescr values in the SNMP configuration. The initiation message is sent to current BMP sessions.</p>
Options	<p>text—Specify a character string for a type 0 TLV to send with the initiation message.</p> <p>Range: 1 through 255 characters</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring BGP Monitoring Protocol Version 3 on page 3647

keep

Syntax	keep (all none);
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i> <i>neighbor address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <i>neighbor address</i>], [edit protocols bgp], [edit protocols bgp <i>group group-name</i>], [edit protocols bgp group <i>group-name</i> <i>neighbor address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <i>neighbor address</i>]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Control whether or not Junos OS keeps in memory and hides certain routes.</p> <p>If the keep none statement is used, Junos OS does not retain in memory and hide routes that are rejected because of a BGP import policy. Nor does BGP keep in memory and hide routes that are declared unfeasible due to BGP sanity checks. The keep none statement causes Junos OS to discard from memory the routes that are rejected due to BGP-specific logic or BGP evaluation. When a route is rejected because of some non-BGP-specific reason, the keep none statement has no effect on this route. This rejected route is retained in memory and hidden even though keep none is configured. An example of this type of hidden route is a route for which the protocol nexthop is unresolved.</p> <p>The routing table can retain the route information learned from BGP in one of the following ways:</p> <ul style="list-style-type: none">• Default (omit the keep statement)—Keep all route information that was learned from BGP, except for routes whose AS path is looped and whose loop includes the local AS.• keep all—Keep all route information that was learned from BGP.• keep none—Discard routes that were received from a peer and that were rejected by import policy or other sanity checking, such as AS path or next hop. When you configure keep none for the BGP session and the inbound policy changes, Junos OS forces readvertisement of the full set of routes advertised by the peer.

In an AS path healing situation, routes with looped paths theoretically could become usable during a soft reconfiguration when the AS path loop limit is changed. However, there is a significant memory usage difference between the default and **keep all**.

Consider the following scenarios:

- A peer readadvertises routes back to the peer from which it learned them.

This can happen in the following cases:

- Another vendor's routing device advertises the routes back to the sending peer.
- The Junos OS peer's default behavior of not readvertising routes back to the sending peer is overridden by configuring **advertise-peer-as**.
- A provider edge (PE) routing device discards any VPN route that does not have any of the expected route targets.

When **keep all** is configured, the behavior of discarding routes received in the above scenarios is overridden.



CAUTION: If you add or remove **keep all** or **keep none** and the peer does not support session restart, the associated BGP sessions are restarted (flapped). To determine if a peer supports refresh, check for **Peer supports Refresh capability** in the output of the **show bgp neighbor** command.

Default	By default, BGP retains incoming rejected routes in memory and hides them. If you do not include the keep statement, most routes are retained in the routing table. BGP keeps all route information that was learned from BGP, except for routes whose AS path is looped and whose loop includes the local AS.
Options	<p>all—Retain all routes.</p> <p>none—Discard routes that were received from a peer and that were rejected by import policy or other sanity checking. When keep none is configured for the BGP session and the inbound policy changes, Junos OS forces readvertisement of the full set of routes advertised by the peer.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • out-delay on page 4054 • <i>Interprovider VPN Example—MP-EBGP Between ISP Peer Routers</i> • <i>Example: Configuring a Routing Policy for Conditional Advertisement of Prefixes in a Routing Table</i>

key-chain (BGP BFD Authentication)

Syntax	<code>key-chain key-chain-name;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection authentication],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection authentication],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection authentication],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection authentication],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection authentication],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection authentication],</code> <code>[edit protocols bgp bgp bfd-liveness-detection authentication],</code> <code>[edit protocols bgp group <i>group-name</i> bgp bfd-liveness-detection authentication],</code> <code>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection authentication],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection authentication],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection authentication],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection authentication]</code>
Release Information	Statement introduced in Junos OS Release 8.1. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for BFD authentication introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Associate a security key with the specified BFD session using the name of the security keychain. Each key has a unique start time within the keychain. Keychain authentication allows you to change the password information periodically without bringing down peering sessions. This keychain authentication method is referred to as <i>hitless</i> because the keys roll over from one to the next without resetting any peering sessions or interrupting the routing protocol.
Options	key-chain-name —Name of the authentication keychain. The keychain name must match one of the keychains configured with the key-chain key-chain-name statement at the [edit security authentication-key-chain] hierarchy level.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring BFD for Static Routes for Faster Network Failure Detection on page 3241• Example: Configuring BFD Authentication for Securing Static Routes on page 3255

- [Example: Configuring BFD on Internal BGP Peer Sessions on page 3812](#)
- [Example: Configuring BGP Route Authentication on page 3897](#)
- [Example: Configuring EBGP Multihop Sessions on page 3777](#)

local-address (Protocols BGP)

Syntax	<code>local-address address;</code>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i> <i>neighbor address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name neighbor address</i>], [edit protocols bgp], [edit protocols bgp <i>group group-name</i>], [edit protocols bgp <i>group group-name neighbor address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i> <i>neighbor address</i>]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Specify the address of the local end of a BGP session. This address is used to accept incoming connections to the peer and to establish connections to the remote peer. When none of the operational interfaces are configured with the specified local address, a session with a BGP peer is placed in the idle state.</p> <p>You generally configure a local address to explicitly configure the system's IP address from BGP's point of view. This IP address can be either an IPv6 or IPv4 address. Typically, an IP address is assigned to a loopback interface, and that IP address is configured here.</p> <p>For internal BGP (IBGP) peering sessions, generally the loopback interface (lo0) is used to establish connections between the IBGP peers. The loopback interface is always up as long as the device is operating. If there is a route to the loopback address, the IBGP peering session stays up. If a physical interface address is used instead and that interface goes up and down, the IBGP peering session also goes up and down. Thus, the loopback interface provides fault tolerance in case the physical interface or the link goes down, if the device has link redundancy.</p> <p>When a device peers with a remote device's loopback interface address, the local device expects BGP update messages to come from (be sourced by) the remote device's loopback interface address. The local-address statement enables you to specify the source information in BGP update messages. If you omit the local-address statement, the expected source of BGP update messages is based on the device's source address selection rules, which normally result in the egress interface address being the expected source of update messages. When this happens, the peering session is not established</p>

because a mismatch exists between the expected source address (the egress interface of the peer) and the actual source (the loopback interface of the peer). To ensure that the expected source address matches the actual source address, specify the loopback interface address in the **local-address** statement.



NOTE: Although a BGP session can be established when only one of the paired routing devices has **local-address** configured, we strongly recommend that you configure **local-address** on both paired routing devices for IBGP and multihop EBGP sessions. The **local-address** statement ensures that deterministic fixed addresses are used for the BGP session end-points.

If you include the **default-address-selection** statement in the configuration, the software chooses the system default address as the source for most locally generated IP packets. For protocols in which the local address is unconstrained by the protocol specification, for example IBGP and multihop EBGP, if you do not configure a specific local address when configuring the protocol, the local address is chosen using the same methods as other locally generated IP packets.

Default If you do not configure a local address, BGP uses the routing device's source address selection rules to set the local address.

Options **address**—IPv6 or IPv4 address of the local end of the connection.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [Example: Configuring Internal BGP Peering Sessions on Logical Systems on page 3636](#)
- [Example: Configuring Internal BGP Peer Sessions on page 3625](#)
- [Understanding Internal BGP Peering Sessions on page 3624](#)
- [router-id on page 3369](#)

local-address (Protocols BMP)

Syntax	<code>local-address address;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options bmp], [edit logical-systems <i>logical-system-name</i> routing-options bmp station station-name], [edit routing-options bmp], [edit routing-options bmp station station-name]
Release Information	Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series. Statement introduced in Junos OS Release 13.3. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>(Optional) Specifies the IPv4 or IPv6 address for the BMP connection on the device. We recommend that you configure a local address. For both active and passive modes, configure a loopback local address. This provides a consistent local endpoint, is useful for debugging, and assures greater reliability for the BMP connection since it is not tied to a single router interface.</p> <p>For passive mode, specifying a local address is required. It also provides some security against a malicious BMP connection. For active mode, we also recommend configuring a local address to help ensure reliability.</p> <p>If you change the local address, the BMP station connection flaps when you commit the configuration.</p>
Options	address —Specify the IPv4 or IPv6 address for the BMP connection on the local device.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring BGP Monitoring Protocol Version 3 on page 3647

local-as

Syntax	<code>local-as <i>autonomous-system</i> <loops number> <private alias> <no-prepend-global-as>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <i>group group-name</i>],</p> <p>[edit protocols bgp <i>group group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>alias option introduced in Junos OS Release 9.5.</p> <p>no-prepend-global-as option introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Specify the local autonomous system (AS) number. An AS is a set of routing devices that are under a single technical administration and generally use a single interior gateway protocol (IGP) and metrics to propagate routing information within the set of routing devices.</p> <p>Internet service providers (ISPs) sometimes acquire networks that belong to a different AS. When this occur, there is no seamless method for moving the BGP peers of the acquired network to the AS of the acquiring ISP. The process of configuring the BGP peers with the new AS number can be time-consuming and cumbersome. In this case, it might not be desirable to modify peer arrangements or configuration. During this kind of transition period, it can be useful to configure BGP-enabled devices in the new AS to use the former AS number in BGP updates. This former AS number is called a <i>local</i> AS.</p>



NOTE: If you are using BGP on the routing device, you must configure an AS number before you specify the local as number.

In Junos OS Release 9.1 and later, the AS numeric range in plain-number format is extended to provide BGP support for 4-byte AS numbers, as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*.

In Junos OS Release 9.3 and later, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period:

<16-bit high-order value in decimal>.<16-bit low-order value in decimal>. For example, the 4-byte AS number of 65546 in plain-number format is represented as 1.10 in the AS-dot notation format.

Options **alias**—(Optional) Configure the local AS as an alias of the global AS number configured for the router at the **[edit routing-options]** hierarchy level. As a result, a BGP peer considers any local AS to which it is assigned as equivalent to the primary AS number configured for the routing device. When you use the **alias** option, only the AS (global or local) used to establish the BGP session is prepended in the AS path sent to the BGP neighbor.

autonomous-system—AS number.

Range: 1 through 4,294,967,295 ($2^{32} - 1$) in plain-number format

Range: 0.0 through 65535.65535 in AS-dot notation format

loops number—(Optional) Specify the number of times detection of the AS number in the AS_PATH attribute causes the route to be discarded or hidden. For example, if you configure **loops 1**, the route is hidden if the AS number is detected in the path one or more times. This is the default behavior. If you configure **loops 2**, the route is hidden if the AS number is detected in the path two or more times.



NOTE: If you configure the local AS values for any BGP group, the detection of routing loops is performed using both the AS and the local AS values for all BGP groups.

If the local AS for the EBGp or IBGP peer is the same as the current AS, do not use the **local-as** statement to specify the local AS number.

When you configure the local AS within a VRF, this impacts the AS path loop-detection mechanism. All of the **local-as** statements configured on the device are part of a single AS domain. The AS path loop-detection mechanism is based on looking for a matching AS present in the domain.

Range: 1 through 10

Default: 1

no-prepend-global-as—(Optional) Specify to strip the global AS and to prepend only the local AS in AS paths sent to external peers.

private—(Optional) Configure to use the local AS only during the establishment of the BGP session with a BGP neighbor but to hide it in the AS path sent to external BGP peers. Only the global AS is included in the AS path sent to external peers.



NOTE: The **private** and **alias** options are mutually exclusive. You cannot configure both options with the same **local-as** statement.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Examples: Configuring BGP Local AS on page 3704 • Example: Configuring a Local AS for EBGp Sessions on page 3709 • autonomous-system on page 3276 • family on page 3997

local-port

Syntax	<code>local-port port;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options bmp], [edit logical-systems <i>logical-system-name</i> routing-options bmp station station-name], [edit routing-options bmp], [edit routing-options bmp station station-name]
Release Information	Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series. Statement introduced in Junos OS Release 13.3. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Specifies the listening port for the BMP station connection.</p> <p>If you configure the connection-mode statement as active, do not configure the local-port statement. If you configure the connection-mode statement as passive, you must configure local-port statement.</p> <p>If you change the local port, the BMP station connection flaps when you commit the configuration.</p>
Options	<p>port—Specify the local port for the BMP station connection.</p> <p>Range: 1 through 65,535</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring BGP Monitoring Protocol Version 3 on page 3647

local-preference

Syntax	<code>local-preference local-preference;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Modify the value of the LOCAL_PREF path attribute, which is a metric used by IBGP sessions to indicate the degree of preference for an external route. The route with the highest local preference value is preferred.</p> <p>The LOCAL_PREF path attribute always is advertised to internal BGP peers and to neighboring confederations. It is never advertised to external BGP peers.</p>
Default	If you omit this statement, the LOCAL_PREF path attribute, if present, is not modified.
Options	<p>local-preference—Preference to assign to routes learned from BGP or from the group or peer.</p> <p>Range: 0 through 4,294,967,295 ($2^{32} - 1$)</p> <p>Default: If the LOCAL_PREF path attribute is present, do not modify its value. If a BGP route is received without a LOCAL_PREF attribute, the route is handled locally (it is stored in the routing table and advertised by BGP) as if it were received with a LOCAL_PREF value of 100. By default, non-BGP routes that are advertised by BGP are advertised with a LOCAL_PREF value of 100.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring the Local Preference Value for BGP Routes on page 3653 • Understanding Internal BGP Peering Sessions on page 3624

- [preference on page 4063](#)

log-updown

Syntax	log-updown;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Specify to generate a log a message whenever a BGP peer makes a state transition. Messages are logged using the system logging mechanism located at the [edit system syslog] hierarchy level.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Preventing BGP Session Resets on page 3919 • Configuring System Logging of BGP Peer State Transitions • traceoptions on page 4081

loops (BGP Address Family)

Syntax	<code>loops <i>number</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols bgp family <i>address-family</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> family <i>address-family</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family <i>address-family</i>],</code> <code>[edit protocols bgp family <i>address-family</i>],</code> <code>[edit protocols bgp group <i>group-name</i> family <i>address-family</i>],</code> <code>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> family <i>address-family</i>]</code>
Release Information	Statement introduced in Junos OS Release 9.6.
Description	<p>For the specified BGP address family, allow the local device's AS number in the received AS paths and specify the number of times the detection of the local device's AS in the AS_PATH attribute is allowed. If the count exceeds the specified loop count, the system discards this route. For example, if you configure loops 1, the route is discarded if the neighbor's local AS is detected in the path more than once. This prevents routing loops and is the default behavior. If you configure loops 2, the route is discarded if the neighbor's local AS is detected more than 2 times.</p>

For debugging, you can configure the **keep all** option. If you want to hide this route.

Some examples of BGP address families are as follows:

- **inet unicast**
- **inet-vpn multicast**
- **inet6 any**
- **l2vpn auto-discovery-only**
- ...

This list is truncated for brevity. For a complete list of protocol families for which you can specify the **loops** statement, enter the **help apropos loops** configuration command at the **[edit protocols bgp]** hierarchy level on your device.

```
[edit protocols bgp]
user@host# help apropos loops
set family inet unicast loops
    Allow local AS in received AS paths
set family inet unicast loops <loops>
    AS-Path loop count
set family inet multicast loops
    Allow local AS in received AS paths
set family inet multicast loops <loops>
    AS-Path loop count
set family inet flow loops
    Allow local AS in received AS paths
set family inet flow loops <loops>
    AS-Path loop count
set family inet any loops
```

```

    Allow local AS in received AS paths
set family inet any loops <loops>
    AS-Path loop count
set family inet labeled-unicast loops
    Allow local AS in received AS paths
...

```



NOTE: The behavior of this statement is slightly different from the *loops (Autonomous System)* statement.

Options *number*—Maximum number of times that the local device's AS number is allowed in the AS_PATH attribute to accept the route.

Range: 1 through 10

Default: None. The system does not take any action unless the *loops (BGP Address Family)* statement is configured.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Enabling BGP Route Advertisements*
- [autonomous-system on page 3276](#)
- [family on page 3997](#)
- [local-as on page 4023](#)
- *loops (Autonomous System)*

loose-check (BGP BFD Authentication)

Syntax	loose-check ;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection authentication],</p> <p>[edit protocols bgp bgp bfd-liveness-detection authentication],</p> <p>[edit protocols bgp group <i>group-name</i> bfd-liveness-detection authentication],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection authentication],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection authentication],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection authentication],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection authentication]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for BFD authentication introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Specify loose authentication checking on the BFD session. Use loose authentication for transitional periods only when authentication might not be configured at both ends of the BFD session.</p> <p>By default, strict authentication is enabled and authentication is checked at both ends of each BFD session. Optionally, to smooth migration from nonauthenticated sessions to authenticated sessions, you can configure <i>loose checking</i>. When loose checking is configured, packets are accepted without authentication being checked at each end of the session.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Example: Configuring BFD for Static Routes for Faster Network Failure Detection on page 3241• Example: Configuring BFD Authentication for Securing Static Routes on page 3255• Example: Configuring BFD on Internal BGP Peer Sessions on page 3812

- [Example: Configuring BGP Route Authentication on page 3897](#)
- [Example: Configuring EBGp Multihop Sessions on page 3777](#)

maximum-ecmp

Syntax	<code>maximum-ecmp <i>next-hops</i>;</code>
Hierarchy Level	[edit chassis]
Release Information	Statement introduced in Junos OS Release 13.2 for QFX switches.
Description	Configure 16, 32, or 64 ECMP next hops for RSVP or LDP LSPs, or MPLS static LSPs that are configured using <code>set protocols mpls static-label-switched-path</code> .
Default	16
Options	<code>next-hops</code> —Specify the number of next hops (16, 32, or 64) for RSVP or LDP LSPs, or MPLS static LSPs
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

metric-out

Syntax	<code>metric-out (<i>metric</i> <i>minimum-igp offset</i> <i>igp</i> (<i>delay-med-update</i> <i>offset</i>);</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name neighbor address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name neighbor address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <i>group group-name</i>],</p> <p>[edit protocols bgp <i>group group-name neighbor address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name neighbor address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Option delay-med-update introduced in Junos OS Release 9.0.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Specify the metric for all routes sent using the multiple exit discriminator (MED, or MULTI_EXIT_DISC) path attribute in update messages. This path attribute is used to discriminate among multiple exit points to a neighboring AS. If all other factors are equal, the exit point with the lowest metric is preferred.</p> <p>You can specify a constant metric value by including the metric option. For configurations in which a BGP peer sends third-party next hops that require the local system to perform next-hop resolution—IBGP configurations, configurations within confederation peers, or EBGP configurations that include the multihop command—you can specify a variable metric by including the minimum-igp or igp option.</p> <p>You can increase or decrease the variable metric calculated from the IGP metric (either from the igp or minimum-igp statement) by specifying a value for offset. The metric is increased by specifying a positive value for offset, and decreased by specifying a negative value for offset.</p> <p>In Junos OS Release 9.0 and later, you can specify that a BGP group or peer not advertise updates for the MED path attributes used to calculate IGP costs for BGP next hops unless the MED is lower. You can also configure an interval to delay when MED updates are sent by including the med-igp-update-interval minutes statement at the [edit routing-options] hierarchy level.</p>

Options **delay-med-update**—Specify that a BGP group or peer configured with the **metric-out igp** statement not advertise MED updates unless the current MED value is lower than the previously advertised MED value, or another attribute associated with the route has changed, or the BGP peer is responding to a refresh route request.



NOTE: You cannot configure the **delay-med-update** statement at the global BGP level.

igp—Set the metric to the most recent metric value calculated in the IGP to get to the BGP next hop. Routes learned from an EBGP peer usually have a next hop on a directly connected interface and thus the IGP value is equal to zero. This is the value advertised.

metric—Primary metric on all routes sent to peers.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

Default: No metric is sent.

minimum-igp—Set the metric to the minimum metric value calculated in the IGP to get to the BGP next hop. If a newly calculated metric is greater than the minimum metric value, the metric value remains unchanged. If a newly calculated metric is lower, the metric value is lowered to that value. When you change a neighbor's export policy from any configuration to a configuration that sets the minimum IGP offset on an exported route, the advertised MED is not updated if the value would increase as a result, even if the previous configuration does not use a minimum IGP-based MED value. This behavior helps to prevent unnecessary route flapping when an IGP cost changes, by not forcing a route update if the metric value increases past the previous lowest known value.

offset—Increases or decreases the metric by this value.

Range: -2^{31} through $2^{31} - 1$

Default: None

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [Example: Associating the MED Path Attribute with the IGP Metric and Delaying MED Updates on page 3694](#)
- [Examples: Configuring BGP MED on page 3665](#)
- [Example: Configuring the MED Attribute Directly on page 3668](#)
- [Understanding the MED Attribute That Determines the Exit Point in an AS on page 3665](#)
- [med-igp-update-interval on page 3329](#)

minimum-interval (BFD Liveness Detection)

Syntax	<code>minimum-interval <i>milliseconds</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols l2vpn oam bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls oam bfd-liveness-detection],</p> <p>[edit protocols bgp bfd-liveness-detection],</p> <p>[edit protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bgp bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols l2vpn oam bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls oam bfd-liveness-detection]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2 for Layer 2 VPN and VPLS.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Configure the minimum interval after which the local routing device transmits hello packets and then expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately using the minimum-interval (specified under the transmit-interval statement) and minimum-receive-interval statements.</p>
Options	<p><i>milliseconds</i>—Specify the minimum interval value for BFD liveliness detection.</p> <p>Range: 1 through 255,000</p>

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring BFD for Layer 2 VPN and VPLS• Example: Configuring BFD for Static Routes for Faster Network Failure Detection on page 3241• bfd-liveness-detection on page 3981• minimum-receive-interval on page 4038• transmit-interval on page 4086

minimum-interval (transmit-interval)

Syntax	<code>minimum-interval <i>milliseconds</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection transmit-interval],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection transmit-interval],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection transmit-interval],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection transmit-interval],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection transmit-interval],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection transmit-interval],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols l2vpn oam bfd-liveness-detection transmit-interval],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam bfd-liveness-detection transmit-interval],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam bfd-liveness-detection transmit-interval],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls oam bfd-liveness-detection transmit-interval],</p> <p>[edit protocols bgp bfd-liveness-detection transmit-interval],</p> <p>[edit protocols bgp group <i>group-name</i> bfd-liveness-detection transmit-interval],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bgp bfd-liveness-detection transmit-interval],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection transmit-interval],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection transmit-interval],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection transmit-interval],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols l2vpn oam bfd-liveness-detection transmit-interval],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam bfd-liveness-detection transmit-interval],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam bfd-liveness-detection transmit-interval],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls oam bfd-liveness-detection transmit-interval]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.2.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for BFD authentication introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2 for Layer 2 VPN and VPLS.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Configure the minimum interval at which the local routing device transmits hello packets to a neighbor with which it has established a BFD session. Optionally, instead of using

this statement at this hierarchy level, you can configure the minimum transmit interval using the [minimum-interval](#) statement at the **bfd-liveness-detection** hierarchy level.

Options *milliseconds*—Minimum transmit interval value.

Range: 1 through 255,000



NOTE: The threshold value specified in the **threshold** statement must be greater than the value specified in the **minimum-interval** statement for the **transmit-interval** statement.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [Configuring BFD for Layer 2 VPN and VPLS](#)
- [Example: Configuring BFD for Static Routes for Faster Network Failure Detection on page 3241](#)
- [bfd-liveness-detection on page 3981](#)
- [minimum-interval on page 4034](#)
- [threshold on page 4079](#)

minimum-receive-interval (BFD Liveness Detection)

Syntax	<code>minimum-receive-interval <i>milliseconds</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols l2vpn oam bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls oam bfd-liveness-detection],</p> <p>[edit protocols bgp bfd-liveness-detection],</p> <p>[edit protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bgp bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols l2vpn oam bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls oam bfd-liveness-detection]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for BFD authentication introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2 for Layer 2 VPN and VPLS.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Configure the minimum interval after which the local routing device must receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the minimum-interval statement.
Options	<p><i>milliseconds</i>—Specify the minimum receive interval value.</p> <p>Range: 1 through 255,000</p>

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [Configuring BFD for Layer 2 VPN and VPLS](#)
- [Example: Configuring BFD for Static Routes for Faster Network Failure Detection on page 3241](#)
- [bfd-liveness-detection on page 3981](#)
- [minimum-interval on page 4034](#)
- [transmit-interval on page 4086](#)

monitor (Protocols BMP)

Syntax monitor (enable | disable);

Hierarchy Level [edit logical-systems *logical-system-name* protocols bgp **bmp**],
[edit logical-systems *logical-system-name* protocols bgp group *group-name* bmp],
[edit logical-systems *logical-system-name* protocols bgp group *group-name* neighbor *address* bmp],
[edit logical-systems *logical-system-name* routing-options **bmp**],
[edit logical-systems *logical-system-name* routing-options bmp **station** *station-name*],
[edit protocols bgp **bmp**],
[edit protocols bgp group *group-name* bmp],
[edit protocols bgp group *group-name* neighbor *address* bmp],
[edit routing-options **bmp**],
[edit routing-options bmp **station** *station-name*]

Release Information Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 13.3.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description BMP monitoring is enabled by default. You can explicitly enable BMP monitoring or disable it. You can also selectively enable or disable BMP monitoring at various hierarchy levels (for example, [edit protocols bgp group *group-name*] or [edit protocols bgp group *group-name* neighbor *address*]). If you disable BMP monitoring, withdrawal messages are sent for any previously advertised routes. These are followed by a down message. If you enable BMP monitoring, an up message is sent first and then the route advertisements follow.

Options **enable**—Enable BMP monitoring.
Default: BMP monitoring is enabled by default.
disable—Disable BMP monitoring.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

mtu-discovery

Syntax	mtu-discovery;
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i> <i>neighbor address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name neighbor address</i>], [edit protocols bgp], [edit protocols bgp <i>group group-name</i>], [edit protocols bgp <i>group group-name neighbor address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i> <i>neighbor address</i>]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Configure TCP path maximum transmission unit (MTU) discovery.</p> <p>TCP path MTU discovery enables BGP to automatically discover the best TCP path MTU for each BGP session. In Junos OS, TCP path MTU discovery is disabled by default for all BGP neighbor sessions.</p> <p>When MTU discovery is disabled, TCP sessions that are not directly connected transmit packets of 512-byte maximum segment size (MSS). These small packets minimize the chances of packet fragmentation at a device along the path to the destination. However, because most links use an MTU of at least 1500 bytes, 512-byte packets do not result in the most efficient use of link bandwidth. For directly connected EBGP sessions, MTU mismatches prevent the BGP session from being established. As a workaround, enable path MTU discovery within the EBGP group.</p> <p>Path MTU discovery dynamically determines the MTU size on the network path between the source and the destination, with the goal of avoiding IP fragmentation. Path MTU discovery works by setting the Don't Fragment (DF) bit in the IP headers of outgoing packets. When a device along the path has an MTU that is smaller than the packet, the device drops the packet. The device also sends back an ICMP Fragmentation Needed (Type 3, Code 4) message that contains the device's MTU, thus allowing the source to reduce its path MTU appropriately. The process repeats until the MTU is small enough to traverse the entire path without fragmentation.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

- Related Documentation**
- [Example: Limiting TCP Segment Size for BGP on page 3912](#)
 - *Configuring Junos OS for IPv6 Path MTU Discovery*
 - *Configuring the Junos OS for Path MTU Discovery on Outgoing GRE Tunnel Connections*

multihop

Syntax multihop {
 no-nexthop-change;
 ttl *ttl-value*;
 }

Hierarchy Level [edit logical-systems *logical-system-name* protocols bgp],
 [edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
 [edit logical-systems *logical-system-name* protocols bgp **group** *group-name*
 neighbor *address*],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols
 bgp],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols
 bgp **group** *group-name*],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols
 bgp group *group-name* **neighbor** *address*],
 [edit protocols bgp],
 [edit protocols bgp **group** *group-name*],
 [edit protocols bgp group *group-name* **neighbor** *address*],
 [edit routing-instances *routing-instance-name* protocols bgp],
 [edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
 [edit routing-instances *routing-instance-name* protocols bgp group *group-name*
 neighbor *address*]

Release Information Statement introduced before Junos OS Release 7.4.
 Statement introduced in Junos OS Release 9.0 for EX Series switches.
 Statement introduced in Junos OS Release 11.3 for the QFX Series.
 Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure an EBGp multihop session.

For Layer 3 VPNs, you configure the EBGp multihop session between the PE and CE routing devices. This allows you to configure one or more routing devices between the PE and CE routing devices.

An external confederation peer is a special case that allows unconnected third-party next hops. You do not need to configure multihop sessions explicitly in this particular case because multihop behavior is implied.

If you have external BGP confederation peer-to-loopback addresses, you still need the multihop configuration.



NOTE: You cannot configure the `accept-remote-nexthop` statement at the same time.

Default If you omit this statement, all EBGp peers are assumed to be directly connected (that is, you are establishing a nonmultihop, or “regular,” BGP session), and the default time-to-live (TTL) value is 1.

The remaining statements are explained separately.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring EBGp Multihop Sessions on page 3777• <i>Configuring EBGp Multihop Sessions Between PE and CE Routers in Layer 3 VPNs</i>• accept-remote-nextthop on page 3966• <i>no-nextthop-change</i>• <i>tth</i>

multiplier (BFD Liveness Detection)

Syntax	<code>multiplier <i>number</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols l2vpn oam bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls oam bfd-liveness-detection],</p> <p>[edit protocols bgp bfd-liveness-detection],</p> <p>[edit protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bgp bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols l2vpn oam bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls oam bfd-liveness-detection]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for BFD authentication introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2 for Layer 2 VPN and VPLS.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.
Options	<p><i>number</i>—Number of hello packets.</p> <p>Range: 1 through 255</p> <p>Default: 3</p>

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring BFD for Layer 2 VPN and VPLS</i>• Example: Configuring BFD for Static Routes for Faster Network Failure Detection on page 3241• bfd-liveness-detection on page 3981

neighbor (Protocols BGP)

```
Syntax  neighbor address {
    accept-remote-nexthop;
    advertise-bgp-static
    advertise-external <conditional>;
    advertise-inactive;
    (advertise-peer-as | no-advertise-peer-as);
    as-override;
    authentication-algorithm algorithm;
    authentication-key key;
    authentication-key-chain key-chain;
    cluster cluster-identifier;
    damping;
    description text-description;
    export [ policy-names ];
    enforce-first-as;
    family {
        (inet | inet6 | inet-mvpn | inet6-mpvn | inet-vpn | inet6-vpn | iso-vpn | l2-vpn) {
            (any | flow | multicast | unicast | signaling) {
                accepted-prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                }
                damping;
                prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                }
                rib-group group-name;
                topology name {
                    community {
                        target identifier;
                    }
                }
            }
        }
        flow {
            no-validate policy-name;
        }
        labeled-unicast {
            accepted-prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            aggregate-label {
                community community-name;
            }
            explicit-null {
                connected-only;
            }
            prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
        }
    }
}
```

```

        resolve-vpn;
        rib inet.3;
        rib-group group-name;
        topology name {
            community {
                target identifier;
            }
        }
    }
}
route-target {
    advertise-default;
    external-paths number;
    accepted-prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
    }
}
signaling {
    prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
    }
}
}
graceful-restart {
    disable;
    restart-time seconds;
    stale-routes-time seconds;
}
graceful-restart {
    long-lived {
        receiver {
            enable;
            disable;
        }
        advertise-to-non-llgr-neighbor {
            omit-no-export;
        }
    }
}
graceful-restart {
    disable-notification-flag;
    disable-notification-extensions {
        omit-no-export;
    }
}
forwarding-state-bit (from-fib | set); /* Configurable to be common for all address
    families */
forwarding-state-bit (as-rr-client | from-fib); /* Configurable for each address family */
long-lived {
    restarter {
        disable;
    }
}

```

```

        stale-time interval;
    }
}
}
hold-time seconds;
import [ policy-names ];
ipsec-sa ipsec-sa;
keep (all | none);
local-address address;
local-as autonomous-system <private>;
local-interface interface-name;
local-preference preference;
log-updown;
metric-out (metric | minimum-igp <offset> | igp <offset>);
mtu-discovery;
multihop <ttl-value>;
multipath {
    multiple-as;
}
no-aggregator-id;
no-client-reflect;
out-delay seconds;
passive;
peer-as autonomous-system;
preference preference;
tcp-aggressive-transmission;
tcp-mss segment-size;
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
vpn-apply-export;
}

```

Hierarchy Level [edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols
bgp **group** *group-name*],
[edit protocols bgp **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Explicitly configure a neighbor (peer). To configure multiple BGP peers, include multiple **neighbor** statements.

By default, the peer's options are identical to those of the group. You can override these options by including peer-specific option statements within the **neighbor** statement.

The **neighbor** statement is one of the statements you can include in the configuration to define a minimal BGP configuration on the routing device. (You can include an **allow all** statement in place of a **neighbor** statement.)

Options *address*—IPv6 or IPv4 address of a single peer.

The remaining statements are explained separately.

Required Privilege routing—To view this statement in the configuration.

Level routing-control—To add this statement to the configuration.

Related Documentation • *BGP Feature Guide for Routing Devices*

no-adaptation (BFD Liveness Detection)

Syntax	no-adaptation;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols l2vpn oam bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls oam bfd-liveness-detection],</p> <p>[edit protocols bgp bfd-liveness-detection],</p> <p>[edit protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bgp bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols l2vpn oam bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls oam bfd-liveness-detection]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.0</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for BFD authentication introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2 for Layer 2 VPN and VPLS.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Configure BFD sessions not to adapt to changing network conditions. We recommend that you <i>do not</i> disable BFD adaptation unless it is preferable to have BFD adaptation disabled in your network.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

- Related Documentation**
- [Configuring BFD for Layer 2 VPN and VPLS](#)
 - [Example: Configuring BFD for Static Routes for Faster Network Failure Detection on page 3241](#)
 - [bfd-liveness-detection on page 3981](#)

no advertise-peer-as

Syntax	no-advertise-peer-as;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Enable the default behavior of suppressing AS routes.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring BGP Route Advertisement on page 3769 • Configuring Routing Policies to Control BGP Route Advertisements on page 3769 • advertise-peer-as on page 3970

no-aggregator-id

Syntax	no-aggregator-id;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Prevent different routing devices within an AS from creating aggregate routes that contain different AS paths.</p> <p>Junos OS performs route aggregation, which is the process of combining the characteristics of different routes so that only a single route is advertised. Aggregation reduces the amount of information that BGP must store and exchange with other BGP systems. When aggregation occurs, the local routing device adds the local AS number and the router ID to the aggregator path attribute. The no-aggregator-id statement causes Junos OS to place a 0 in the router ID field and thus eliminate the possibility of having multiple aggregate advertisements in the network, each with different path information.</p>
Default	If you omit this statement, the router ID is included in the BGP aggregator path attribute.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Update Messages on page 3597

no-client-reflect

Syntax	no-client-reflect;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Disable intracluster route redistribution by the system acting as the route reflector. Include this statement when the client cluster is fully meshed to prevent the sending of redundant route advertisements. Route reflection provides a way to decrease BGP control traffic and minimizing the number of update messages sent within the AS.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring BGP Route Reflectors on page 3873 • cluster on page 3988

out-delay

Syntax	<code>out-delay seconds;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols bgp],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i></code> <code> <i>neighbor address</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp <i>group group-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp <i>group group-name</i> <i>neighbor address</i>],</code> <code>[edit protocols bgp],</code> <code>[edit protocols bgp <i>group group-name</i>],</code> <code>[edit protocols bgp <i>group group-name</i> <i>neighbor address</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i></code> <code> <i>neighbor address</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Control how often BGP and the routing table exchange route information by specifying how long a route must be present in the Junos OS routing table before it is exported to BGP. Use this time delay to help bundle routing updates and to avoid sending updates too often.</p> <p>Alternatively or in addition, external BGP (EBGP) sessions can also use the route-flap damping mechanism upon the reception of BGP messages coming from an external neighbor.</p> <p>BGP stores the route information it receives from update messages in the routing table, and the routing table exports active routes from the routing table into BGP. BGP then advertises the exported routes to its peers. The out-delay statement enables a form of rate limiting. The delay is added to each update for each prefix individually. When a routing device changes its best path to a destination prefix, the device does not inform its peer about the change unless the route has been present in its routing table for the specified out-delay. If you use out-delay to perform rate-limiting, you can expect a less bursty pattern of updates. You will see a pattern in which updates arrive in a steady flow, and two updates for the same prefix are always spaced by at least the out-delay timer value (for example, 30 seconds). Thus, the out-delay setting is useful for limiting oscillation (sometimes called <i>churn</i>) in a network. Keep in mind that, regardless of the out-delay setting, BGP peers exchange routes immediately after neighbor establishment. The out-delay setting is only designed to delay the exchange of routes between BGP and the local routing table.</p>

Caution is warranted because an **out-delay** can delay convergence. If your network is configured in a way that avoids oscillation, setting an **out-delay** is not necessary.

When configured, the **out-delay** value displays as **Outbound Timer** when using **show bgp group** or **show bgp group neighbor** commands.


Default By default, the exchange of route information between BGP and the routing table occurs immediately after the routes are received. This immediate exchange of route information might cause instabilities in the network reachability information. If you omit this statement, routes are exported to BGP immediately after they have been added to the routing table.

Options *seconds*—Output delay time.
Range: 0 through 65,535 seconds
Default: 0 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation • [keep on page 4016](#)

outbound-route-filter

Syntax	<pre> outbound-route-filter { bgp-orf-cisco-mode; prefix-based { accept { (inet inet6); } } } </pre>
Hierarchy Level	<pre> [edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>] </pre>
Release Information	<p>Statement introduced in Junos OS Release 9.2.</p> <p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Configure a BGP peer to accept outbound route filters from a remote peer.
Options	<p>accept—Specify that outbound route filters from a BGP peer be accepted.</p> <p>inet—Specify that IPv4 prefix-based outbound route filters be accepted.</p> <p>inet6—Specify that IPv6 prefix-based outbound route filters be accepted.</p>
	<p> NOTE: You can specify that both IPv4 and IPv6 outbound route filters be accepted.</p>
	<p>prefix-based—Specify that prefix-based filters be accepted.</p> <p>The bgp-orf-cisco-mode statement is explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

- Related Documentation**
- [Example: Configuring BGP Prefix-Based Outbound Route Filtering on page 3773](#)

passive (Protocols BGP)

Syntax	passive;
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i> <i>neighbor address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name neighbor address</i>], [edit protocols bgp], [edit protocols bgp <i>group group-name</i>], [edit protocols bgp <i>group group-name neighbor address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i> <i>neighbor address</i>]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Configure the routing device so that active open messages are not sent to the peer. Once you configure the routing device to be passive, the routing device will wait for the peer to issue an open request before a message is sent.
Default	If you omit this statement, all explicitly configured peers are active, and each peer periodically sends open requests until its peer responds.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Preventing BGP Session Flaps When VPN Families Are Configured on page 3919

path-selection

Syntax	<pre> path-selection { (always-compare-med cisco-non-deterministic external-router-id); as-path-ignore; l2vpn-use-bgp-rules; med-plus-igp { igp-multiplier <i>number</i>; med-multiplier <i>number</i>; } } </pre>
Hierarchy Level	<pre> [edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp] </pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>med-plus-igp option introduced in Junos OS Release 8.1.</p> <p>as-path-ignore and l2vpn-use-bgp-rules options introduced in Junos OS Release 10.2.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Configure BGP path selection.
Default	If the path-selection statement is not included in the configuration, only the multiple exit discriminators (MEDs) of routes that have the same peer ASs are compared.
Options	always-compare-med —Always compare MEDs whether or not the peer ASs of the compared routes are the same.



NOTE: We recommend that you configure the **always-compare-med** option.

as-path-ignore—In the best-path algorithm, skip the step that compares the autonomous system (AS) path lengths. By default, the best-path algorithm evaluates the length of the AS paths and prefers the route with the shortest AS path length.



NOTE: The **as-path-ignore** statement is not supported with routing instances.

cisco-non-deterministic—Emulate the Cisco IOS default behavior. This mode evaluates routes in the order that they are received and does not group them according to their neighboring AS. With **cisco-non-deterministic** mode, the active path is always first. All inactive, but eligible, paths follow the active path and are maintained in the order

in which they were received, with the most recent path first. Ineligible paths remain at the end of the list.

As an example, suppose you have three path advertisements for the 192.168.1.0 /24 route:

- Path 1—learned through EBGP; AS Path of 65010; MED of 200
- Path 2—learned through IBGP; AS Path of 65020; MED of 150; IGP cost of 5
- Path 3—learned through IBGP; AS Path of 65010; MED of 100; IGP cost of 10

These advertisements are received in quick succession, within a second, in the order listed. Path 3 is received most recently, so the routing device compares it against path 2, the next most recent advertisement. The cost to the IBGP peer is better for path 2, so the routing device eliminates path 3 from contention. When comparing paths 1 and 2, the routing device prefers path 1 because it is received from an EBGP peer. This allows the routing device to install path 1 as the active path for the route.



NOTE: We do not recommend using this configuration option in your network. It is provided solely for interoperability to allow all routing devices in the network to make consistent route selections.

external-router-id—Compare the router ID between external BGP paths to determine the active path.

igp-multiplier *number*—The multiplier value for the IGP cost to a next-hop address. This option is useful for making the MED and IGP cost comparable.

Range: 1 through 1000

Default: 1

med-multiplier *number*—The multiplier value for the MED calculation. This option is useful for making the MED and IGP cost comparable.

Range: 1 through 1000

Default: 1

med-plus-igp—Add the IGP cost to the indirect next-hop destination to the MED before comparing MED values for path selection. This statement only affects best-path selection. It does not affect the advertised MED.

The other option is explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [Understanding BGP Path Selection on page 3793](#)
- [Example: Ignoring the AS Path Attribute When Selecting the Best Path on page 3796](#)

peer-as (Protocols BGP)

Syntax	<code>peer-as <i>autonomous-system</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols bgp],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i> neighbor <i>address</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code>bgp],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code>bgp <i>group group-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code>bgp <i>group group-name</i> neighbor <i>address</i>],</code> <code>[edit protocols bgp],</code> <code>[edit protocols bgp <i>group group-name</i>],</code> <code>[edit protocols bgp <i>group group-name</i> neighbor <i>address</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i></code> <code>neighbor <i>address</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Specify the neighbor (peer) autonomous system (AS) number.</p> <p>For EBGp, the peer is in another AS, so the AS number you specify in the peer-as statement must be different from the local router's AS number, which you specify in the autonomous-system statement. For IBGP, the peer is in the same AS, so the two AS numbers that you specify in the autonomous-system and peer-as statements must be the same.</p> <p>The AS numeric range in plain-number format has been extended in Junos OS Release 9.1 to provide BGP support for 4-byte AS numbers, as defined in RFC 4893, <i>BGP Support for Four-octet AS Number Space</i>. RFC 4893 introduces two new optional transitive BGP attributes, AS4_PATH and AS4_AGGREGATOR. These new attributes are used to propagate 4-byte AS path information across BGP speakers that do not support 4-byte AS numbers. RFC 4893 also introduces a reserved, well-known, 2-byte AS number, AS 23456. This reserved AS number is called AS_TRANS in RFC 4893. All releases of the Junos OS support 2-byte AS numbers.</p> <p>In Junos OS Release 9.2 and later, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: <i><16-bit high-order value in decimal>.<16-bit low-order value in decimal></i>. For example, the 4-byte AS number of 65,546 in plain-number format is represented as 1.10 in the AS-dot notation format.</p> <p>With the introduction of 4-byte AS numbers, you might have a combination of routers that support 4-byte AS numbers and 2-byte AS numbers. For more information about what happens when establishing BGP peer relationships between 4-byte and 2-byte capable routers, see the following topics:</p>

- *Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview.*

Options *autonomous-system*—AS number.
Range: 1 through 4,294,967,295 ($2^{32} - 1$) in plain-number format for 4-byte AS numbers
Range: 1 through 65,535 in plain-number format for 2-byte AS numbers (this is a subset of the 4-byte range)
Range: 0.0 through 65535.65535 in AS-dot notation format for 4-byte AS numbers

Required Privilege routing—To view this statement in the configuration.
Level routing-control—To add this statement to the configuration.

post-policy

Syntax `post-policy {
 exclude-non-eligible;
}`

Hierarchy Level [edit protocols bgp bmp [route-monitoring](#)],
 [edit protocols bgp group *group-name* bmp route-monitoring],
 [edit protocols bgp group neighbor *group-name* neighbor *address* bmp route-monitoring],
 [edit routing-options bmp route-monitoring],
 [edit routing-options bmp station *station-name* route-monitoring]

Release Information Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
 Statement introduced in Junos OS Release 13.3.
 Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description For BMP route monitoring, allows you to excludes routes that are non-eligible for the decision process (for example, protocol nexthop not resolved). This represents the view of the BGP routes after running the import policy. If the import policy has rejected the BGP route, the route does not exist in the post policy view.

Options *exclude-non-eligible*—Exclude routes that are non-eligible for the decision process.


Required Privilege routing—To view this statement in the configuration.
Level routing-control—To add this statement to the configuration.

Related Documentation • [Configuring BGP Monitoring Protocol Version 3 on page 3647](#)


pre-policy

Syntax	<pre>pre-policy { exclude-non-feasible; }</pre>
Hierarchy Level	[edit protocols bgp bmp route-monitoring], [edit protocols bgp group <i>group-name</i> bmp route-monitoring], [edit protocols bgp group neighbor <i>group-name</i> neighbor <i>address</i> bmp route-monitoring], [edit routing-options bmp route-monitoring], [edit routing-options bmp station <i>station-name</i> route-monitoring]
Release Information	Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series. Statement introduced in Junos OS Release 13.3. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Excludes routes that are non-feasible from the BMP route monitoring decision process (for example, a route loop). This represents the view of the BGP routes before running the import policy.
Options	exclude-non-feasible —Exclude routes that are non-feasible for the decision process.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring BGP Monitoring Protocol Version 3 on page 3647

preference (Protocols BGP)

Syntax	<code>preference <i>preference</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Specify the preference for routes learned from BGP.</p> <p>At the BGP global level, the preference statement sets the preference for routes learned from BGP. You can override this preference in a BGP group or peer preference statement.</p> <p>At the group or peer level, the preference statement sets the preference for routes learned from the group or peer. Use this statement to override the preference set in the BGP global preference statement when you want to favor routes from one group or peer over those of another.</p>
<div>  NOTE: Do not set preference2 for BGP route-policy. </div>	
Options	<p>preference—Preference to assign to routes learned from BGP or from the group or peer.</p> <p>Range: 0 through 4,294,967,295 ($2^{32} - 1$)</p> <p>Default: 170 for the primary preference</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • local-preference on page 4026 • Example: Configuring the Preference Value for BGP Routes on page 3787

priority (Protocols BMP)

Syntax	priority (high medium low);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options bmp], [edit logical-systems <i>logical-system-name</i> routing-options bmp station station-name], [edit routing-options bmp], [edit routing-options bmp station station-name]
Release Information	Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series. Statement introduced in Junos OS Release 13.3. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specifies the dispatch priority for BMP. The dispatch priority controls the frequency with which the device is able to forward BMP messages to BMP stations.
Options	high —Specifies that the routing protocol process handle BMP requests with high urgency. medium —Specifies that the routing protocol process handle BMP requests with medium urgency. low —Specifies that the routing protocol process handle BMP requests with low urgency. Default: The default dispatch priority is low to minimize interference with other routing protocol process priorities and to match the behavior of previous versions of BMP.
<div> NOTE: Setting high or medium priority may reduce the performance of the routing protocol process in its handling route convergence or other work.</div>	
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring BGP Monitoring Protocol Version 3 on page 3647

remove-private

Syntax `remove-private;`

Hierarchy Level [edit logical-systems *logical-system-name* protocols bgp],
 [edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
 [edit logical-systems *logical-system-name* protocols bgp **group** *group-name* neighbor *address*],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name* neighbor *address*],
 [edit protocols bgp],
 [edit protocols bgp **group** *group-name*],
 [edit protocols bgp **group** *group-name* neighbor *address*],
 [edit routing-instances *routing-instance-name* protocols bgp],
 [edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
 [edit routing-instances *routing-instance-name* protocols bgp **group** *group-name* neighbor *address*]

Release Information Statement introduced before Junos OS Release 7.4.
 Statement introduced in Junos OS Release 9.0 for EX Series switches.
 Statement introduced in Junos OS Release 11.3 for the QFX Series.
 Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
no-peer-loop-check option added in Junos OS Release 15.1.

Description When advertising AS paths to remote systems, have the local system strip private AS numbers from the AS path. The numbers are stripped from the AS path starting at the left end of the AS path (the end where AS paths have been most recently added). The routing device stops searching for private ASs when it finds the first nonprivate AS or a peer's private AS. If the AS path contains the AS number of the external BGP (EBGP) neighbor, BGP does not remove the private AS number.



NOTE: As of Junos OS 10.0R2 and higher, if there is a need to send prefixes to an EBGP peer that has an AS number that matches an AS number in the AS path, consider using the **as-override** statement instead of the **remove-private** statement.

The operation takes place after any confederation member ASs have already been removed from the AS path, if applicable.

Junos OS recognizes the set of AS numbers that is considered private, a range that is defined in the Internet Assigned Numbers Authority (IANA) assigned numbers document.

The set of reserved AS numbers is in the range from 64,512 through 65,535.

Options **all**—Remove all private AS numbers from the original path. Do not stop the process of removing private AS numbers, even if a public AS number is encountered.

nearest—When you use the **all** and **replace** options, choose the last (right-most) public AS number encountered in the original AS path for the replacement value, as the AS path is processed from left to right. If no public AS number is encountered, the default replacement value is used. (See the **replace** option for information about the default replacement value.)

replace—When you use the **all** option, instead of removing private AS numbers, perform a replace operation. The default replacement value for the private AS number is the local AS number at the BGP group level for the BGP peer. If you are unsure about the replacement value, check the local AS value displayed in the output of the **show bgp group group-name** command.

no-peer-loop-check—Peer loop check is removed. By default, the **remove-private** statement has a peer loop check restriction. If a private AS in the AS path has the same value as the configured **peer-as** for the neighbor, **remove-private** does not remove or replace this private AS number. This restriction provides peer-as loop protection. However, you can remove this restriction using the **no-peer-loop-check** option.

Required Privilege	routing—To view this statement in the configuration.
Level	routing-control—To add this statement to the configuration.

Related Documentation	<ul style="list-style-type: none">• Example: Removing Private AS Numbers from AS Paths on page 3805
------------------------------	---

restart-time (BGP Graceful Restart)

Syntax	<code>restart-time seconds;</code>
Hierarchy Level	<p>[edit protocols (bgp rip ripng) graceful-restart],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (bgp rip ripng) graceful-restart (Enabling Globally)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp graceful-restart],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp graceful-restart]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Configure the duration of the BGP, RIP, or next-generation RIP (RIPng) graceful restart period.
Options	<p>seconds—Length of time for the graceful restart period.</p> <p>Range: 1 through 600 seconds</p> <p>Default: Varies by protocol:</p> <ul style="list-style-type: none"> • BGP—120 seconds • RIP and RIPng—60 seconds
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Graceful Restart Options for BGP on page 2581 • Configuring Graceful Restart Options for RIP and RIPng on page 2585 • Configuring Graceful Restart for QFabric Systems • stale-routes-time on page 2672

route-monitoring

Syntax	<pre>route-monitoring { none; post-policy { exclude-non-eligible; } pre-policy { exclude-non-feasible; } }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp bmp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bmp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bmp], [edit logical-systems <i>logical-system-name</i> routing-options bmp], [edit logical-systems <i>logical-system-name</i> routing-options bmp station <i>station-name</i>], [edit protocols bgp bmp], [edit protocols bgp group <i>group-name</i> bmp], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bmp], [edit routing-options bmp], [edit routing-options bmp station <i>station-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.3.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Specify whether BMP should send pre-policy route monitoring messages, post-policy route monitoring messages, both types of messages, or none at all. The pre-policy can be configured to exclude routes that are non-feasible for the decision process (for example, a route loop). The post-policy can be configured to exclude routes that are not eligible for the decision process (for example, protocol nexthop not resolved).</p> <p>You can also selectively enable or disable BMP route monitoring at various hierarchy levels (for example, [edit protocols bgp group <i>group-name</i>] or [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>]).</p>
Options	<p>none—Explicitly disables BMP route monitoring.</p> <p>Default: If you configure the route-monitoring statement at the [edit routing-options bmp] hierarchy level, the default option is pre-policy. If you configure the route-monitoring statement at any of the [edit protocols bgp] hierarchy levels, the default option is to inherit the configuration from the route-monitoring statement configured at the [edit routing-options bmp] hierarchy level.</p> <p>The other statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

- Related Documentation**
- [Configuring BGP Monitoring Protocol Version 3 on page 3647](#)

session-mode

Syntax	<code>session-mode (automatic multihop single-hop);</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection]</p>
Release Information	<p>Statement introduced in Junos OS Release 11.1.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Configure BFD session mode to be single-hop or multihop. By default, BGP uses single-hop BFD sessions if the peer is directly connected to the router's interface. BGP uses multihop BFD sessions if the peer is not directly connected to the router's interface. If the peer session's local-address option is configured, the directly connected check is based partly on the source address that would be used for BGP and BFD.</p> <p>For backward compatibility, you can override the default behavior by configuring the single-hop or multihop option. Before Junos OS Release 11.1, the behavior was to assume that IBGP peer sessions were multihop.</p>
Options	<p>automatic—Configure BGP to use single-hop BFD sessions if the peer is directly connected to the router's interface, and multihop BFD sessions if the peer is not directly connected to the router's interface</p> <p>multihop—Configure BGP to use multihop BFD sessions.</p> <p>single-hop—Configure BGP to use single-hop BFD sessions.</p> <p>Default: automatic</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring BFD Authentication for BGP on page 3822 • Example: Configuring BFD on Internal BGP Peer Sessions on page 3812

- [Example: Configuring BFD Authentication for BGP on page 3822](#)
- [Understanding BFD Authentication for BGP on page 3820](#)

stale-routes-time

Syntax	<code>stale-routes-time <i>seconds</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-routing-name</i> protocols bgp graceful-restart], [edit logical-systems <i>logical-routing-name</i> routing-instances <i>routing-instance-name</i> protocols bgp graceful-restart], [edit protocols bgp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols bgp graceful-restart]
Release Information	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.
Description	Specify the maximum time that stale routes are kept during a restart. The stale-routes-time statement allows you to set the length of time the routing device waits to receive messages from restarting neighbors before declaring them down.
Options	seconds —Time the router device waits to receive messages from restarting neighbors before declaring them down. Range: 1 through 600 seconds Default: 300 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Graceful Restart Options for BGP on page 2581 • Configuring Graceful Restart for QFabric Systems • restart-time (BGP Graceful Restart) on page 2671

station

Syntax	<pre> station <i>station-name</i> { authentication-algorithm (aes-128-cmac-96 hmac-sha-1-96 md5); authentication-key <i>key</i>; authentication-key-chain <i>authentication-key-chain</i>; connection-mode (active passive); hold-down { seconds; flaps <i>flaps</i>; period <i>seconds</i>; } initiation-message <i>text</i>; local-address <i>address</i>; local-port <i>port</i>; monitor (disable enable); priority (high low medium); route-monitoring { none; post-policy { exclude-non-eligible; } pre-policy { exclude-non-feasible; } } station-address (<i>ip-address</i> <i>name</i>); station-port <i>port-number</i>; statistics-timeout <i>seconds</i>; traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier>; } } </pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options bmp], [edit routing-options bmp]
Release Information	Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series. Statement introduced in Junos OS Release 13.3. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify and configure a BMP monitoring station. Be aware that each BMP monitoring station can use a significant amount of a device's resources. You can configure up to 3 BMP monitoring stations.
Options	<p><i>station-name</i>—Specify a name for the BMP station.</p> <p>The other statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

station-address

Syntax	<code>station-address (address station-name);</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options bmp], [edit logical-systems <i>logical-system-name</i> routing-options bmp station <i>station-name</i>], [edit routing-options bmp], [edit routing-options bmp station <i>station-name</i>]
Release Information	Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series. Statement introduced in Junos OS Release 13.3. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify the name or address for the BMP monitoring station. You can specify one or the other but not both.
Options	<i>station-address</i> —Specify the address for the BMP station. The address should be a valid IPv4 or IPv6 address. <i>station-name</i> —Specify the name for the BMP station.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring BGP Monitoring Protocol Version 3 on page 3647

station-port

Syntax	<code>station-port port;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options bmp], [edit logical-systems <i>logical-system-name</i> routing-options bmp station station-name], [edit routing-options bmp], [edit routing-options bmp station station-name]
Release Information	Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series. Statement introduced in Junos OS Release 13.3. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify the port number for the BMP monitoring station.
Options	port —Specify the port number for the BMP monitoring station. If the connection-mode statement is configured as active a station port number is required. If the connection-mode statement is configured as passive , you must not configure a station port number. Range: 1 though 65535
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring BGP Monitoring Protocol Version 3 on page 3647• connection-mode on page 3989

statistics-timeout

Syntax	<code>statistics-timeout <i>seconds</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options bmp], [edit logical-systems <i>logical-system-name</i> routing-options bmp station <i>station-name</i>], [edit routing-options bmp], [edit routing-options bmp station <i>station-name</i>]
Release Information	Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series. Statement introduced in Junos OS Release 13.3. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify how often statistics messages are sent to the BMP monitoring station. If you configure a value of 0, no statistics messages are sent.
Options	<i>seconds</i> —Specify the number for the BMP monitoring station. Default: 3600 seconds Range: 15 though 65535 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring BGP Monitoring Protocol Version 3 on page 3647

tcp-mss (Protocols BGP)

Syntax	<code>tcp-mss <i>segment-size</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols bgp],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor</code> <code> <i>neighbor-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp group <i>group-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp group <i>group-name</i> neighbor <i>neighbor-name</i>],</code> <code>[edit protocols bgp],</code> <code>[edit protocol bgp group <i>group-name</i>],</code> <code>[edit protocols bgp group <i>group-name</i> neighbor <i>neighbor-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor</code> <code> <i>neighbor-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 8.1. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the maximum segment size (MSS) for the TCP connection for BGP neighbors. The MSS is only valid in increments of 2 KB. The value used is based on the value set, but is rounded down to the nearest multiple of 2048.
Options	<i>segment-size</i> —MSS for the TCP connection. Range: 1 through 4096
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Limiting TCP Segment Size for BGP on page 3912

threshold (detection-time)

Syntax	<code>threshold <i>milliseconds</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection detection-time],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection detection-time],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection detection-time],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection detection-time],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection detection-time],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection detection-time],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols l2vpn oam bfd-liveness-detection detection-time],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam bfd-liveness-detection detection-time],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam bfd-liveness-detection detection-time],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls oam bfd-liveness-detection detection-time],</p> <p>[edit protocols bgp bfd-liveness-detection detection-time],</p> <p>[edit protocols bgp group <i>group-name</i> bfd-liveness-detection detection-time],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bgp bfd-liveness-detection detection-time],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection detection-time],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection detection-time],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection detection-time]</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols l2vpn oam bfd-liveness-detection detection-time],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam bfd-liveness-detection detection-time],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam bfd-liveness-detection detection-time],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls oam bfd-liveness-detection detection-time]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.2.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for BFD authentication introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2 for Layer 2 VPNs and VPLS.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>

Description Specify the threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.



NOTE: The threshold value must be equal to or greater than the transmit interval.

The threshold time must be equal to or greater than the value specified in the `minimum-interval` or the `minimum-receive-interval` statement.

Options *milliseconds*—Value for the detection time adaptation threshold.
Range: 1 through 255,000

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [Configuring BFD for Layer 2 VPN and VPLS](#)
- [Example: Configuring BFD for Static Routes for Faster Network Failure Detection on page 3241](#)

threshold (transmit-interval)

Syntax	<code>threshold <i>milliseconds</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection transmit-interval],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection transmit-interval],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection transmit-interval],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection transmit-interval],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection transmit-interval],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection transmit-interval],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols l2vpn oam bfd-liveness-detection transmit-interval],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam bfd-liveness-detection transmit-interval],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam bfd-liveness-detection transmit-interval],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls oam bfd-liveness-detection transmit-interval],</p> <p>[edit protocols bgp bfd-liveness-detection transmit-interval],</p> <p>[edit protocols bgp group <i>group-name</i> bfd-liveness-detection transmit-interval],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bgp bfd-liveness-detection transmit-interval],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection transmit-interval],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection transmit-interval],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection transmit-interval],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols l2vpn oam bfd-liveness-detection transmit-interval],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam bfd-liveness-detection transmit-interval],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam bfd-liveness-detection transmit-interval],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls oam bfd-liveness-detection transmit-interval]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.2.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2 for Layer 2 VPN and VPLS.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Specify the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent.

Options *milliseconds*—Value for the transmit interval adaptation threshold.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)




NOTE: The threshold value specified in the `threshold` statement must be greater than the value specified in the `minimum-interval` statement for the `transmit-interval` statement.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [Configuring BFD for Layer 2 VPN and VPLS](#)
- [Example: Configuring BFD for Static Routes for Faster Network Failure Detection on page 3241](#)
- [bfd-liveness-detection on page 3981](#)

traceoptions (Protocols BGP)

Syntax	<pre> traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier> <disable>; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>4byte-as statement introduced in Junos OS Release 9.2.</p> <p>4byte-as statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Configure BGP protocol-level tracing options. To specify more than one tracing operation, include multiple flag statements.
<div>  NOTE: The traceoptions statement is not supported on QFabric systems. </div>	
Default	<p>The default BGP protocol-level tracing options are inherited from the routing protocols traceoptions statement included at the [edit routing-options] hierarchy level. The default group-level trace options are inherited from the BGP protocol-level traceoptions statement. The default peer-level trace options are inherited from the group-level traceoptions statement.</p>
Options	<p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p>

file name—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`. We recommend that you place BGP tracing output in the file `bgp-log`.

files number—(Optional) Maximum number of trace files. When a trace file named `trace-file.0` reaches its maximum size, it is renamed `trace-file.0`, then `trace-file.1`, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum number of files, you must also specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 10 files

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

BGP Tracing Flags

- **4byte-as**—4-byte AS events.
- **bfd**—BFD protocol events.
- **damping**—Damping operations.
- **graceful-restart**—Graceful restart events.
- **keepalive**—BGP keepalive messages. If you enable the the BGP **update** flag only, received keepalive messages do not generate a trace message.
- **nsr-synchronization**—Nonstop routing synchronization events.
- **open**—Open packets. These packets are sent between peers when they are establishing a connection.
- **packets**—All BGP protocol packets.
- **refresh**—BGP refresh packets.
- **update**—Update packets. These packets provide routing updates to BGP systems. If you enable only this flag, received keepalive messages do not generate a trace message. Use the **keepalive** flag to generate a trace message for keepalive messages.

Global Tracing Flags

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

Default: If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information.
- **filter**—Provide filter trace information. Applies only to **route**, **damping**, and **update** tracing flags.
- **receive**—Trace the packets being received.
- **send**—Trace the packets being transmitted.

no-world-readable—(Optional) Prevent any user from reading the log file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 128 KB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level	<p>routing and trace—To view this statement in the configuration.</p> <p>routing-control and trace-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • log-updown on page 4027 statement • <i>Tracing Nonstop Active Routing Synchronization Events</i> • Understanding Trace Operations for BGP Protocol Traffic on page 3949 • Configuring OSPF Refresh and Flooding Reduction in Stable Topologies on page 4459

traceoptions (Protocols BMP)

Syntax	<pre>traceoptions { file <i>file-name</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <<i>flag-modifier</i>> <disable>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options bmp], [edit logical-systems <i>logical-system-name</i> routing-options bmp station <i>station-name</i>], [edit routing-options bmp], [edit routing-options bmp station <i>station-name</i>]
Release Information	Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series. Statement introduced in Junos OS Release 13.3. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure tracing options for BMP monitoring. To specify more than one tracing operation, include multiple flag statements.
Options	<p>file <i>file-name</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. We recommend that you place BMP tracing output in the file bmp-log.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named trace-file.0 reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum number of files, you must also specify a maximum file size with the size option.</p> <p>Range: 2 through 1000 files</p> <p>Default: 10 files</p> <p>flag—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements.</p> <ul style="list-style-type: none">• all—Trace all BMP monitoring operations.• down—Down messages.• error—Error conditions.• event—Major events, station establishment, errors, and events.• general—General events.• normal—Normal events.• packets—All messages.• policy—Policy processing.• route—Routing information.• route-monitoring—Route monitoring messages.• state—State transitions.

- **statistics**—Statistics messages.
- **task**—Routing protocol task processing.
- **timer**—Routing protocol timer processing.
- **up**—Up messages.
- **write**—Writing of messages.

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information.
- **disable**—Disable the tracing flag.
- **receive**—Trace the packets being received.
- **send**—Trace the packets being transmitted.

no-world-readable—(Optional) Prevent any user from reading the log file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 128 KB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level	routing and trace—To view this statement in the configuration. routing-control and trace-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Tracing BMP Operations on page 3956 • Understanding Trace Operations for BGP Protocol Traffic on page 3949 • Configuring OSPF Refresh and Flooding Reduction in Stable Topologies on page 4459

transmit-interval (BFD Liveness Detection)

Syntax	<pre>transmit-interval { minimum-interval milliseconds; threshold milliseconds; }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection], [edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols l2vpn oam bfd-liveness-detection], [edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam bfd-liveness-detection], [edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam bfd-liveness-detection], [edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls oam bfd-liveness-detection], [edit protocols bgp bfd-liveness-detection], [edit protocols bgp group <i>group-name</i> bfd-liveness-detection], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bgp bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols l2vpn oam bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols vpls oam bfd-liveness-detection]</pre>
Release Information	<p>Statement introduced in Junos OS Release 8.2.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for BFD authentication introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2 for Layer 2 VPN and VPLS.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Specify the transmit interval for the bfd-liveness-detection statement. The negotiated transmit interval for a peer is the interval between the sending of BFD packets to peers. The receive interval for a peer is the minimum time that it requires between packets sent from its peer; the receive interval is not negotiated between peers. To determine the transmit interval, each peer compares its configured minimum transmit interval with its</p>

peer's minimum receive interval. The larger of the two numbers is accepted as the transmit interval for that peer.

The remaining statements are explained separately.

Required Privilege	routing—To view this statement in the configuration.
Level	routing-control—To add this statement to the configuration.

Related Documentation	<ul style="list-style-type: none">• <i>Configuring BFD for Layer 2 VPN and VPLS</i>• Example: Configuring BFD for Static Routes for Faster Network Failure Detection on page 3241• bfd-liveness-detection on page 3981• threshold on page 4079• minimum-interval on page 4036• minimum-receive-interval on page 4038
------------------------------	---

version (BFD Liveness Detection)

Syntax	version (0 1 automatic);
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols l2vpn oam bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls oam bfd-liveness-detection],</p> <p>[edit protocols bgp bfd-liveness-detection],</p> <p>[edit protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bgp bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols l2vpn oam bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls oam bfd-liveness-detection]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.1</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2 for Layer 2 VPN and VPLS.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Specify the BFD version for detection. You can explicitly configure BFD version 0, version 1, or the routing device can automatically detect the BFD version. By default, the routing device automatically detects the BFD version, which is either 0 or 1.
Options	<p>Configure the BFD version to detect: 0 (BFD version 0), 1 (BFD version 1), or automatic (autodetect the BFD version)</p> <p>Default: automatic</p>

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring BFD for Layer 2 VPN and VPLS• Example: Configuring BFD Authentication for BGP on page 3822• Example: Configuring BFD on Internal BGP Peer Sessions on page 3812• Example: Configuring BFD Authentication for BGP on page 3822• Understanding BFD Authentication for BGP on page 3820

CHAPTER 155

Operational Commands

- clear bgp damping
- clear bgp neighbor
- clear bgp table
- show bgp bmp
- show bgp group
- show bgp neighbor
- show bgp summary
- show policy damping
- show route damping
- show route detail

clear bgp damping

List of Syntax	Syntax on page 4092 Syntax (EX Series Switch and QFX Series) on page 4092
Syntax	<code>clear bgp damping</code> <code><logical-system (all <i>logical-system-name</i>)></code> <code><prefix></code>
Syntax (EX Series Switch and QFX Series)	<code>clear bgp damping</code> <code><prefix></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Clear BGP route flap damping information.
Options	none —Clear all BGP route flap damping information. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. prefix —(Optional) Clear route flap damping information for only the specified destination prefix.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show policy damping on page 4130• show route damping on page 3430
List of Sample Output	clear bgp damping on page 4092
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear bgp damping

```
user@host> clear bgp damping
```

clear bgp neighbor

List of Syntax [Syntax on page 4093](#)
 [Syntax \(EX Series Switch and QFX Series\) on page 4093](#)

Syntax clear bgp neighbor
 <all>
 <as *as-number*>
 <gracefully>
 <instance *instance-name*>
 <logical-system (all | *logical-system-name*)>
 <malformed-route>
 <neighbor>
 <soft | soft-inbound>
 <soft-minimum-igp>
 <stale-routes>

Syntax (EX Series Switch and QFX Series) clear bgp neighbor
 <all>
 <as *as-number*>
 <instance *instance-name*>
 <malformed-route>
 <neighbor>
 <soft | soft-inbound>
 <soft-minimum-igp>

Release Information Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.0 for EX Series switches.
 Command introduced in Junos OS Release 11.3 for the QFX Series.
 Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
 malformed-route option introduced in Junos OS Release 13.2.
 all option introduced in Junos OS Release 14.2.
 gracefully and **stale-routes** options introduced in Junos OS Release 15.1.

Description Perform one of the following tasks:

- Change the state of one or more BGP neighbors to **IDLE**. For neighbors in the **ESTABLISHED** state, this command drops the TCP connection to the neighbors and then reestablishes the connection.
- (**soft** keyword only) Reapply export policies and send refresh updates to one or more BGP neighbors without changing their state.
- (**soft-inbound** keyword only) Send a route-refresh message to one or more BGP neighbors without changing their state, and reapply import policies on the received updates.

Options **none | all**—(Optional) Change the state of all BGP neighbors to **IDLE**.
 Both **clear bgp neighbor** and **clear bgp neighbor all** function identically.

as *as-number*—(Optional) Apply this command only to neighbors in the specified autonomous system (AS).

gracefully—(Optional) Enable the BGP peer to start graceful-restart receiving-speaker mode. The receiving speaker also sends its own routes to the restarted speaker, and sends an End-of-RIB marker when it completes the update. The **clear bgp neighbor neighbor-address gracefully** command is the same as **clear bgp neighbor hard** (the default for **clear bgp neighbor**), but it does not use the new Hard Reset subcode on the Notify and Cease messages that are sent. This allows the neighbor to enter GR or LLGR helper mode, if negotiated. The session is still cleared on this router, and this router does not enter GR or LLGR helper mode.

instance instance-name—(Optional) Apply this command only to neighbors for the specified routing instance.

logical-system (all | logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.

malformed-route—(Optional) Remove malformed routes. If a specific neighbor is provided, Junos OS removes malformed routes for that particular neighbor. Otherwise, Junos OS removes malformed routes for all BGP neighbors. To find routes that have malformed attributes, run the **show route hidden** command, and look for routes marked with **MalformedAttr** in the AS path field.

neighbor—(Optional) IP address of a BGP peer. Apply this command only to the specified neighbor.

soft—(Optional) Reapply any export policies and send refresh updates to neighbors without clearing the state.

soft-inbound—(Optional) Send a route-refresh message to BGP neighbors and reapply import policies on the route updates received from the BGP neighbors without clearing the BGP state.

soft-minimum-igp—(Optional) Provide soft refresh of the outbound state when the interior gateway protocol (IGP) metric is reset.

stale-routes—(Optional) Any stale route currently being held for the specified neighbor because of BGP graceful restart (GR) or long-lived graceful restart (LLGR) receiver mode operations.

Required Privilege Level

clear

Related Documentation

- [show bgp neighbor on page 2710](#)

List of Sample Output

[clear bgp neighbor on page 4094](#)

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear bgp neighbor

```
user@host> clear bgp neighbor
```


clear bgp table

Syntax	<code>clear bgp table <i>table-name</i></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switch and QFX Series)	<code>clear bgp table <i>table-name</i></code>
Release Information	Command introduced in Junos OS Release 9.0. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Request that BGP refresh routes in a specified routing table.
Options	logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. <i>table-name</i> —Request that BGP refresh routes in the specified table.
Additional Information	In some cases, a prefix limit is associated with a routing table for a VPN instance. When this limit is exceeded (for example, because of a network misconfiguration), some routes might not be inserted in the table. Such routes need to be added to the table after the network issue is resolved. Use the clear bgp table command to request that BGP refresh routes in a VPN instance table.
Required Privilege Level	clear
List of Sample Output	clear bgp table private.inet.0 on page 4096 clear bgp table inet.6 logical-system all on page 4096 clear bgp table private.inet.6 logical-system ls1 on page 4096 clear bgp table logical-system all inet.0 on page 4096 clear bgp table logical-system ls2 private.inet.0 on page 4097
Output Fields	This command produces no output.

Sample Output

[clear bgp table private.inet.0](#)

```
user@host> clear bgp table private.inet.0
```

[clear bgp table inet.6 logical-system all](#)

```
user@host> clear bgp table inet.6 logical-system all
```

[clear bgp table private.inet.6 logical-system ls1](#)

```
user@host> clear bgp table private.inet.6 logical-system ls1
```

[clear bgp table logical-system all inet.0](#)

```
user@host> clear bgp table logical-system all inet.0
```


`clear bgp table logical-system ls2 private.inet.0`

`user@host> clear bgp table logical-system ls2 private.inet.0`

show bgp bmp

Syntax	show bgp bmp
Release Information	<p>Command introduced in Junos OS Release 9.5.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Display information about the BGP Monitoring Protocol (BMP).
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show bgp bmp on page 4098
Output Fields	Table 320 lists the output fields for the show bgp bmp command. Output fields are listed in the approximate order in which they appear.

Table 320: show bgp bmp Output Fields

Field Name	Field Description
BMP station address/port	IP address and port number of the monitoring station to which BGP Monitoring Protocol (BMP) statistics are sent.
BMP session state	Status of the BMP session: UP or DOWN .
Memory consumed by BMP	Memory used by the active BMP session.
Statistics timeout	Amount of time, in seconds, between transmissions of BMP data to the monitoring station.
Memory limit	Threshold, in bytes, at which the routing device stops collecting BMP data.
Memory-connect retry timeout	Amount of time, in seconds, after which the routing device attempts to resume a BMP session that was ended after the configured memory threshold was exceeded.

Sample Output

show bgp bmp

```

user@host> show bgp bmp
  BMP station address/port: 172.24.24.157+5454
  BMP session state: DOWN
  Memory consumed by BMP: 0
  Statistics timeout: 15
  Memory limit: 10485760
  Memory connect retry timeout: 600

```


show bgp group

List of Syntax	Syntax on page 4100 Syntax (EX Series Switch and QFX Series) on page 4100
Syntax	<pre>show bgp group <brief detail summary> <group-name> <exact-instance instance-name> <instance instance-name> <logical-system (all logical-system-name)> <rtf></pre>
Syntax (EX Series Switch and QFX Series)	<pre>show bgp group <brief detail summary> <group-name> <exact-instance instance-name> <instance instance-name></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> <p>exact-instance option introduced in Junos OS Release 11.4.</p>
Description	Display information about the configured BGP groups.
Options	<p>none—Display group information about all BGP groups.</p> <p>brief detail summary—(Optional) Display the specified level of output.</p> <p>group-name—(Optional) Display group information for the specified group.</p> <p>exact-instance instance-name—(Optional) Display information for the specified instance only.</p> <p>instance instance-name—(Optional) Display information about BGP groups for all routing instances whose name begins with this string (for example, cust1, cust11, and cust111 are all displayed when you run the show bgp group instance cust1 command). The instance name can be master for the main instance, or any valid configured instance name or its prefix.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>rtf—(Optional) Display BGP group route targeting information.</p>
Required Privilege Level	view
List of Sample Output	show bgp group on page 4104 show bgp group brief on page 4104

[show bgp group detail on page 4105](#)
[show bgp group rtf detail on page 4106](#)
[show bgp group summary on page 4106](#)

Output Fields Table 321 describes the output fields for the **show bgp group** command. Output fields are listed in the approximate order in which they appear.

Table 321: show bgp group Output Fields

Field Name	Field Description	Level of Output
Group Type or Group	Type of BGP group: Internal or External .	All levels
group-index	Index number for the BGP peer group. The index number differentiates between groups when a single BGP group is split because of different configuration options at the group and peer levels.	rtf detail
AS	AS number of the peer. For internal BGP (IBGP), this number is the same as Local AS .	brief detail none
Local AS	AS number of the local routing device.	brief detail none
Name	Name of a specific BGP group.	brief detail none
Index	Unique index number of a BGP group.	brief detail none
Flags	Flags associated with the BGP group. This field is used by Juniper Networks customer support.	brief detail none
BGP-Static Advertisement Policy	Policies configured for the BGP group with the advertise-bgp-static policy statement.	brief none
Remove-private options	Options associated with the remove-private statement.	brief detail none
Holdtime	Maximum number of seconds allowed to elapse between successive keepalive or update messages that BGP receives from a peer in the BGP group, after which the connection to the peer is closed and routing devices through that peer become unavailable.	brief detail none
Export	Export policies configured for the BGP group with the export statement.	brief detail none
MED tracks IGP metric update delay	Time, in seconds, that updates to multiple exit discriminator (MED) are delayed. Also displays the time remaining before the interval is set to expire	All levels
Traffic Statistics Interval	Time between sample periods for labeled-unicast traffic statistics, in seconds.	brief detail none

Table 321: show bgp group Output Fields (*continued*)

Field Name	Field Description	Level of Output
Total peers	Total number of peers in the group.	brief detail none
Established	Number of peers in the group that are in the established state.	All levels
Active/Received/Accepted/Damped	<p>Multipurpose field that displays information about BGP peer sessions. The field's contents depend upon whether a session is established and whether it was established in the main routing device or in a routing instance.</p> <ul style="list-style-type: none"> If a peer is not established, the field shows the state of the peer session: Active, Connect, or Idle. If a BGP session is established in the main routing device, the field shows the number of active, received, accepted, and damped routes that are received from a neighbor and appear in the inet.0 (main) and inet.2 (multicast) routing tables. For example, 8/10/10/2 and 2/4/4/0 indicate the following: <ul style="list-style-type: none"> 8 active routes, 10 received routes, 10 accepted routes, and 2 damped routes from a BGP peer appear in the inet.0 routing table. 2 active routes, 4 received routes, 4 accepted routes, and no damped routes from a BGP peer appear in the inet.2 routing table. 	summary
ip-addresses	List of peers who are members of the group. The address is followed by the peer's port number.	All levels
Route Queue Timer	Number of seconds until queued routes are sent. If this time has already elapsed, this field displays the number of seconds by which the updates are delayed.	detail
Route Queue	Number of prefixes that are queued up for sending to the peers in the group.	detail
inet.number	<p>Number of active, received, accepted, and damped routes in the routing table. For example, inet.0: 7/10/9/0 indicates the following:</p> <ul style="list-style-type: none"> 7 active routes, 10 received routes, 9 accepted routes, and no damped routes from a BGP peer appear in the inet.0 routing table. 	none

Table 321: show bgp group Output Fields (*continued*)

Field Name	Field Description	Level of Output
Table <i>inet.number</i>	Information about the routing table. <ul style="list-style-type: none"> • Received prefixes—Total number of prefixes from the peer, both active and inactive, that are in the routing table. • Active prefixes—Number of prefixes received from the peer that are active in the routing table. • Suppressed due to damping—Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols. • Advertised prefixes—Number of prefixes advertised to a peer. • Received external prefixes—Total number of prefixes from the external BGP (EBGP) peers, both active and inactive, that are in the routing table. • Active external prefixes—Number of prefixes received from the EBGP peers that are active in the routing table. • Externals suppressed—Number of routes received from EBGP peers currently inactive because of damping or other reasons. • Received internal prefixes—Total number of prefixes from the IBGP peers, both active and inactive, that are in the routing table. • Active internal prefixes—Number of prefixes received from the IBGP peers that are active in the routing table. • Internals suppressed—Number of routes received from IBGP peers currently inactive because of damping or other reasons. • RIB State—Status of the graceful restart process for this routing table: BGP restart is complete, BGP restart in progress, VPN restart in progress, or VPN restart is complete. 	detail
Groups	Total number of groups.	All levels
Peers	Total number of peers.	All levels
External	Total number of external peers.	All levels
Internal	Total number of internal peers.	All levels
Down peers	Total number of unavailable peers.	All levels
Flaps	Total number of flaps that occurred.	All levels
Table	Name of a routing table.	brief , none
Tot Paths	Total number of routes.	brief , none
Act Paths	Number of active routes.	brief , none
Suppressed	Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.	brief , none

Table 321: show bgp group Output Fields (*continued*)

Field Name	Field Description	Level of Output
History	Number of withdrawn routes stored locally to keep track of damping history.	brief, none
Damp State	Number of active routes with a figure of merit greater than zero, but lower than the threshold at which suppression occurs.	brief, none
Pending	Routes being processed by the BGP import policy.	brief, none
Group	Group the peer belongs to in the BGP configuration.	detail
Receive mask	Mask of the received target included in the advertised route.	detail
Entries	Number of route entries received.	detail
Target	Route target that is to be passed by route-target filtering. If a route advertised from the provider edge (PE) routing device matches an entry in the route-target filter, the route is passed to the peer.	detail
Mask	Mask which specifies that the peer receive routes with the given route target.	detail

Sample Output

show bgp group

```
user@host> show bgp group
```

```

Groups: 2  Peers: 2   External: 1   Internal: 2   Down peers: 1   Flaps: 0
Table      Tot Paths  Act Paths  Suppressed    History Damp State   Pending

inet.0
          0         0         0         0         0         0
bgp.13vpn.0
          0         0         0         0         0         0
bgp.rtarget.0
          2         0         0         0         0         0

```

show bgp group brief

```
user@host> show bgp group brief
```

```

Groups: 2  Peers: 2   External: 0   Internal: 2   Down peers: 1   Flaps: 0
Table      Tot Paths  Act Paths  Suppressed    History Damp State   Pending

inet.0
          0         0         0         0         0         0
bgp.13vpn.0
          0         0         0         0         0         0

```



```

bgp.rtarget.0
                2          0          0          0          0          0

```

show bgp group detail

```

user@host> show bgp group detail
Group Type: Internal  AS: 1                      Local AS: 1
Name: ibgp           Index: 0                  Flags: <Export Eval>
Holdtime: 0
Total peers: 3       Established: 0
22.0.0.2
22.0.0.8
22.0.0.5

Groups: 1  Peers: 3  External: 0  Internal: 3  Down peers: 3  Flaps: 3
Table bgp.l3vpn.0
  Received prefixes:      0
  Accepted prefixes:      0
  Active prefixes:        0
  Suppressed due to damping: 0
  Received external prefixes: 0
  Active external prefixes: 0
  Externals suppressed:   0
  Received internal prefixes: 0
  Active internal prefixes: 0
  Internals suppressed:   0
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
Table bgp.mdt.0
  Received prefixes:      0
  Accepted prefixes:      0
  Active prefixes:        0
  Suppressed due to damping: 0
  Received external prefixes: 0
  Active external prefixes: 0
  Externals suppressed:   0
  Received internal prefixes: 0
  Active internal prefixes: 0
  Internals suppressed:   0
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
Table VPN-A.inet.0
  Received prefixes:      0
  Accepted prefixes:      0
  Active prefixes:        0
  Suppressed due to damping: 0
  Received external prefixes: 0
  Active external prefixes: 0
  Externals suppressed:   0
  Received internal prefixes: 0
  Active internal prefixes: 0
  Internals suppressed:   0
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
Table VPN-A.mdt.0
  Received prefixes:      0
  Accepted prefixes:      0
  Active prefixes:        0
  Suppressed due to damping: 0
  Received external prefixes: 0
  Active external prefixes: 0

```

```
Externals suppressed:      0
Received internal prefixes: 0
Active internal prefixes:  0
Internals suppressed:      0
RIB State: BGP restart is complete
RIB State: VPN restart is complete
```

show bgp group rtf detail

```
user@host> show bgp group rtf detail
Group: internal (group-index: 0)
  Receive mask: 00000002
  Table: bgp.rtarget.0
    Target      Mask      Entries: 2
    100:100/64  00000002
    200:201/64  (Group)
Group: internal (group-index: 1)
  Table: bgp.rtarget.0
    Target      Mask      Entries: 1
    200:201/64  (Group)
```

show bgp group summary

```
user@host> show bgp group summary
Group      Type      Peers  Established  Active/Received/Accepted/Damped
ibgp       Internal  3      0
Groups: 1  Peers: 3  External: 0  Internal: 3  Down peers: 3  Flaps: 3
bgp.l3vpn.0 : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0
bgp.mdt.0   : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0
VPN-A.inet.0 : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0
VPN-A.mdt.0 : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0
```

show bgp neighbor

List of Syntax	Syntax on page 4107 Syntax (EX Series Switch, QFX Series, and OCX Series) on page 4107
Syntax	<pre>show bgp neighbor <exact-instance <i>instance-name</i>> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <neighbor-address> <orf (detail <i>neighbor-address</i>)</pre>
Syntax (EX Series Switch, QFX Series, and OCX Series)	<pre>show bgp neighbor <instance <i>instance-name</i>> <exact-instance <i>instance-name</i>> <neighbor-address> <orf (<i>neighbor-address</i> detail)</pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1x53-D20 for the OCX Series.</p> <p>orf option introduced in Junos OS Release 9.2.</p> <p>exact-instance option introduced in Junos OS Release 11.4.</p>
Description	Display information about BGP peers.
Options	<p>none—Display information about all BGP peers.</p> <p>exact-instance <i>instance-name</i>—(Optional) Display information for the specified instance only.</p> <p>instance <i>instance-name</i>—(Optional) Display information about BGP peers for all routing instances whose name begins with this string (for example, cust1, cust11, and cust111 are all displayed when you run the show bgp neighbor instance cust1 command).</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>neighbor-address—(Optional) Display information for only the BGP peer at the specified IP address.</p> <p>orf (detail <i>neighbor-address</i>)—(Optional) Display outbound route-filtering information for all BGP peers or only for the BGP peer at the specified IP address. The default is to display brief output. Use the detail option to display detailed output.</p>
Additional Information	For information about the local-address , nlri , hold-time , and preference statements, see the <i>Junos OS Routing Protocols Library for Routing Devices</i> .
Required Privilege Level	view

Related Documentation • [clear bgp neighbor on page 4093](#)

List of Sample Output [show bgp neighbor on page 4115](#)
[show bgp neighbor \(CLNS\) on page 4116](#)
[show bgp neighbor \(Layer 2 VPN\) on page 4117](#)
[show bgp neighbor \(Layer 3 VPN\) \(Not supported on the OCX Series.\) on page 4119](#)
[show bgp neighbor neighbor-address on page 4120](#)
[show bgp neighbor neighbor-address on page 4120](#)
[show bgp neighbor neighbor-address \(BGP Graceful Restart Enabled\) on page 4121](#)
[show bgp neighbor neighbor-address \(BGP Long-Lived Graceful Restart\) on page 4122](#)
[show bgp neighbor orf neighbor-address detail on page 4122](#)

Output Fields [Table 238](#) describes the output fields for the **show bgp neighbor** command. Output fields are listed in the approximate order in which they appear.

Table 322: show bgp neighbor Output Fields

Field Name	Field Description
Peer	Address of the BGP neighbor. The address is followed by the neighbor port number.
AS	AS number of the peer.
Local	Address of the local routing device. The address is followed by the peer port number.
Type	Type of peer: Internal or External .
State	Current state of the BGP session: <ul style="list-style-type: none"> • Active—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an Open message. • Connect—BGP is waiting for the transport protocol connection to be completed. • Established—The BGP session has been established, and the peers are exchanging update messages. • Idle—This is the first stage of a connection. BGP is waiting for a Start event. • OpenConfirm—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message. • OpenSent—BGP has sent an open message and is waiting to receive an open message from the peer. • route reflector client—The BGP session is established with a route reflector client.

Table 322: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Flags	<p>Internal BGP flags:</p> <ul style="list-style-type: none"> • Aggregate Label—BGP has aggregated a set of incoming labels (labels received from the peer) into a single forwarding label. • CleanUp—The peer session is being shut down. • Delete—This peer has been deleted. • Idled—This peer has been permanently idled. • ImportEval—At the last commit operation, this peer was identified as needing to reevaluate all received routes. • Initializing—The peer session is initializing. • SendRtn—Messages are being sent to the peer. • Sync—This peer is synchronized with the rest of the peer group. • RSync—This peer in the backup Routing Engine is synchronized with the BGP peer in the master Routing Engine for nonstop active routing. • TryConnect—Another attempt is being made to connect to the peer. • Unconfigured—This peer is not configured. • WriteFailed—An attempt to write to this peer failed.
Last state	<p>Previous state of the BGP session:</p> <ul style="list-style-type: none"> • Active—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an Open message. • Connect—BGP is waiting for the transport protocol connection to be completed. • Established—The BGP session has been established, and the peers are exchanging update messages. • Idle—This is the first stage of a connection. BGP is waiting for a Start event. • OpenConfirm—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message. • OpenSent—BGP has sent an open message and is waiting to receive an open message from the peer.
Last event	<p>Last activity that occurred in the BGP session:</p> <ul style="list-style-type: none"> • Closed—The BGP session closed. • ConnectRetry—The transport protocol connection failed, and BGP is trying again to connect. • HoldTime—The session ended because the hold timer expired. • KeepAlive—The local routing device sent a BGP keepalive message to the peer. • Open—The local routing device sent a BGP open message to the peer. • OpenFail—The local routing device did not receive an acknowledgment of a BGP open message from the peer. • RecvKeepAlive—The local routing device received a BGP keepalive message from the peer. • RecvNotify—The local routing device received a BGP notification message from the peer. • RecvOpen—The local routing device received a BGP open message from the peer. • RecvUpdate—The local routing device received a BGP update message from the peer. • Start—The peering session started. • Stop—The peering session stopped. • TransportError—A TCP error occurred.

Table 322: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Last error	<p>Last error that occurred in the BGP session:</p> <ul style="list-style-type: none"> • Cease—An error occurred, such as a version mismatch, that caused the session to close. • Finite State Machine Error—In setting up the session, BGP received a message that it did not understand. • Hold Time Expired—The session's hold time expired. • Message Header Error—The header of a BGP message was malformed. • Open Message Error—A BGP open message contained an error. • None—No errors occurred in the BGP session. • Update Message Error—A BGP update message contained an error.
Export	Name of the export policy that is configured on the peer.
Import	Name of the import policy that is configured on the peer.
Options	<p>Configured BGP options:</p> <ul style="list-style-type: none"> • AddressFamily—Configured address family: inet or inet-vpn. • AdvertiseBGPStatic—Configured BGP static routes are advertised. • AuthKeyChain—Authentication key change is enabled. • DropPathAttributes—Certain path attributes are configured to be dropped from neighbor updates during inbound processing. • GracefulRestart—Graceful restart is configured. • HoldTime—Hold time configured with the hold-time statement. The hold time is three times the interval at which keepalive messages are sent. • IgnorePathAttributes—Certain path attributes are configured to be ignored in neighbor updates during inbound processing. • Local Address—Address configured with the local-address statement. • LLGR—BGP long-lived graceful restart capability is configured. • LLGRHelperDisabled—BGP long-lived graceful restart is completely disabled for a neighbor. • Multihop—Allow BGP connections to external peers that are not on a directly connected network. • NLRI—Configured MBGP state for the BGP group: multicast, unicast, or both if you have configured nlri any. • Peer AS—Configured peer autonomous system (AS). • Preference—Preference value configured with the preference statement. • Refresh—Configured to refresh automatically when the policy changes. • Rib-group—Configured routing table group.
Path-attributes dropped	Path attribute codes that are dropped from neighbor updates.
Path-attributes ignored	Path attribute codes that are ignored during neighbor updates.
Peer does not support LLGR Restarter or Receiver functionality	BGP neighbor does not support long-lived graceful restart (LLGR) restarter mode completely.

Table 322: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Peer does not support LLGR Restarter functionality	BGP neighbor does not support long-lived graceful restart (LLGR) restarter mode for any family.
Authentication key change	(appears only if the authentication-keychain statement has been configured) Name of the authentication keychain enabled.
Authentication algorithm	(appears only if the authentication-algorithm statement has been configured) Type of authentication algorithm enabled: hmac or md5 .
Address families configured	Names of configured address families for the VPN.
BGP-Static Advertisement Policy	Name of the bgp static policy that is configured on the peer.
Local Address	Address of the local routing device.
Remove-private options	Options associated with the remove-private statement.
Holdtime	Hold time configured with the hold-time statement. The hold time is three times the interval at which keepalive messages are sent.
Flags for NLRI inet-label-unicast	Flags related to labeled-unicast: <ul style="list-style-type: none"> • TrafficStatistics—Collection of statistics for labeled-unicast traffic is enabled.
Traffic statistics	Information about labeled-unicast traffic statistics: <ul style="list-style-type: none"> • Options—Options configured for collecting statistics about labeled-unicast traffic. • File—Name and location of statistics log files. • size—Size of all the log files, in bytes. • files—Number of log files.
Traffic Statistics Interval	Time between sample periods for labeled-unicast traffic statistics, in seconds.
Preference	Preference value configured with the preference statement.
Outbound Timer	Time for which the route is available in Junos OS routing table before it is exported to BGP. This field is displayed in the output only if the out-delay parameter is configured to a non-zero value.
Number of flaps	Number of times the BGP session has gone down and then come back up.
Peer ID	Router identifier of the peer.
Group index	Index number for the BGP peer group. The index number differentiates between groups when a single BGP group is split because of different configuration options at the group and peer levels.

Table 322: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Peer index	Index that is unique within the BGP group to which the peer belongs.
Local ID	Router identifier of the local routing device.
Local Interface	Name of the interface on the local routing device.
Active holdtime	Hold time that the local routing device negotiated with the peer.
Keepalive Interval	Keepalive interval, in seconds.
BFD	Status of BFD failure detection.
Local Address	Name of directly connected interface over which direct EBGP peering is established.
NLRI and times for LLGR configured on peer	<p>Names of address families and stale time for BGP long-lived graceful restart configured on the BGP peer or neighbor.</p> <p>Times are displayed using the routing protocol daemon (rpd) %#OT format:</p> <p><weeks>w<days>d <hours>:<minutes>:<seconds></p> <p>Zero leading elements are omitted, for example, a value less than one week do not include the weeks.</p>
NLRI and times that peer supports LLGR Restarter for	<p>Names of address families and stale time that the BGP peer supports for restarter mode for BGP long-lived graceful restart.</p> <p>Times are displayed using the routing protocol daemon (rpd) %#OT format:</p> <p><weeks>w<days>d <hours>:<minutes>:<seconds></p> <p>Zero leading elements are omitted, for example, a value less than one week do not include the weeks.</p>
NLRI that peer saved LLGR forwarding for	Name of the address family for which the BGP peer saved BGP long-lived graceful restart forwarding.
Graceful Restart Details	Amount of time that is remaining until LLGR expires and the time remaining on the GR stale timer, along with RIB details, are displayed while LLGR receiver mode is active (a peer that negotiated LLGR has disconnected and not yet reconnected)
NLRI we are holding stale routes for	Names of address families (NLRIs) for which that stale routes are held or preserved when BGP graceful restart receiver mode is active for a neighbor.
Time until end-of-rib is assumed for stale routes	<p>Amount of time remaining on the stale timer until which end-of-RIB (EoR) markers are assumed when BGP graceful restart receiver mode is active for a neighbor.</p> <p>Time is displayed in Coordinated Universal Time (UTC) format (YYYY-MM-DD-HH:MM:SS). Note that the stale timer display ('Time until end-of-rib is assumed') is also present when a session is active, but the neighbor as not yet sent all of the end-of-rib indications.</p>
Time until stale routes are deleted or become long-lived stale	Amount of time up to which stale routes are deleted or become long-lived stale routes when BGP graceful restart receiver mode is active for a neighbor.

Table 322: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
NLRI for restart configured on peer	Names of address families configured for restart.
NLRI advertised by peer	Address families supported by the peer: unicast or multicast .
NLRI for this session	Address families being used for this session.
Peer supports Refresh capability	Remote peer's ability to send and request full route table readvertisement (route refresh capability). For more information, see RFC 2918, <i>Route Refresh Capability for BGP-4</i> .
Restart time configured on peer	Configured time allowed for restart on the neighbor.
Stale routes from peer are kept for	When graceful restart is negotiated, the maximum time allowed to hold routes from neighbors after the BGP session has gone down.
Peer does not support Restarter functionality	Graceful restart restarter-mode is disabled on the peer.
Peer does not support Receiver functionality	Graceful restart helper-mode is disabled on the peer.
Restart time requested by this peer	Restart time requested by this neighbor during capability negotiation.
Restart flag received from the peer	When this field appears, the BGP speaker has restarted (Restarting), and this peer should not wait for the end-of-rib marker from the speaker before advertising routing information to the speaker.
NLRI that peer supports restart for	Neighbor supports graceful restart for this address family.
NLRI peer can save forwarding state	Neighbor supporting this address family saves all forwarding states.
NLRI that peer saved forwarding for	Neighbor saves all forwarding states for this address family.
NLRI that restart is negotiated for	Router supports graceful restart for this address family.
NLRI of received end-of-rib markers	Address families for which end-of-routing-table markers are received from the neighbor.
NLRI of all end-of-rib markers sent	Address families for which end-of-routing-table markers are sent to the neighbor.
Peer supports 4 byte AS extension (peer-as 1)	Peer understands 4-byte AS numbers in BGP messages. The peer is running Junos OS Release 9.1 or later.

Table 322: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
NLRIs for which peer can receive multiple paths	<p>Appears in the command output of the local router if the downstream peer is configured to receive multiple BGP routes to a single destination, instead of only receiving the active route.</p> <p>Possible value is inet-unicast.</p>
NLRIs for which peer can send multiple paths: inet-unicast	<p>Appears in the command output of the local router if the upstream peer is configured to send multiple BGP routes to a single destination, instead of only sending the active route.</p> <p>Possible value is inet-unicast.</p>
Table inet.number	<p>Information about the routing table:</p> <ul style="list-style-type: none"> • RIB State—BGP is in the graceful restart process for this routing table: restart is complete or restart in progress. • Bit—Number that represents the entry in the routing table for this peer. • Send state—State of the BGP group: in sync, not in sync, or not advertising. • Active prefixes—Number of prefixes received from the peer that are active in the routing table. • Received prefixes—Total number of prefixes from the peer, both active and inactive, that are in the routing table. • Accepted prefixes—Total number of prefixes from the peer that have been accepted by a routing policy. • Suppressed due to damping—Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.
Last traffic (seconds)	Last time any traffic was received from the peer or sent to the peer, and the last time the local routing device checked.
Input messages	Messages that BGP has received from the receive socket buffer, showing the total number of messages, number of update messages, number of times a policy is changed and refreshed, and the buffer size in octets. The buffer size is 16 KB.
Output messages	Messages that BGP has written to the transmit socket buffer, showing the total number of messages, number of update messages, number of times a policy is changed and refreshed, and the buffer size in octets. The buffer size is 16 KB.
Input dropped path attributes	<p>Information about dropped path attributes:</p> <ul style="list-style-type: none"> • Code—Path attribute code. • Count—Path attribute count.
Input ignored path attributes	<p>Information about ignored path attributes:</p> <ul style="list-style-type: none"> • Code—Path attribute code. • Count—Path attribute count.

Table 322: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Output queue	<p>Number of BGP packets that are queued to be transmitted to a particular neighbor for a particular routing table. Output queue 0 is for unicast NLRI, and queue 1 is for multicast NLRI.</p> <p>It also specifies the routing table name and the NLRI that the table was advertised through, in the format (routing table name, NLRI).</p> <p>NOTE: The output queue of routing tables that are not advertised, will only show up at extensive output level.</p>
Trace options	Configured tracing of BGP protocol packets and operations.
Trace file	Name of the file to receive the output of the tracing operation.
Filter Updates rcv	<p>(orf option only) Number of outbound-route filters received for each configured address family.</p> <p>NOTE: The counter is cumulative. For example, the counter is increased after the remote peer either resends or clears the outbound route filtering prefix list.</p>
Immediate	<p>(orf option only) Number of route updates received with the immediate flag set. The immediate flag indicates that the BGP peer should readvertise the updated routes.</p> <p>NOTE: The counter is cumulative. For example, the counter is increased after the remote peer either resends or clears the outbound route filtering prefix list.</p>
Filter	(orf option only) Type of prefix filter received: prefix-based or extended-community .
Received filter entries	(orf option only) List of received filters displayed.
seq	(orf option only) Numerical order assigned to this prefix entry among all the received outbound route filter prefix entries.
prefix	(orf option only) Address for the prefix entry that matches the filter.
minlength	(orf option only) Minimum prefix length, in bits, required to match this prefix.
maxlength	(orf option only) Maximum prefix length, in bits, required to match this prefix.
match	(orf option only) For this prefix match, whether to permit or deny route updates.

Sample Output

show bgp neighbor

```

user@host > show bgp neighbor
Peer: 10.255.7.250+179 AS 10   Local: 10.255.7.248+63740 AS 10
  Type: Internal   State: Established   Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Export: [ redist_static ]
  Options: <Preference LocalAddress PeerAS Refresh>
  Options: <AdvertiseBGPStatic>

```

```

Local Address: 10.255.7.248 Holdtime: 90 Preference: 170 Outbound Timer: 50
Number of flaps: 0
Peer ID: 10.255.7.250    Local ID: 10.255.7.248    Active Holdtime: 90
Keepalive Interval: 30    Group index: 0    Peer index: 0
BFD: disabled, down
NLRI for restart configured on peer: inet-unicast
NLRI advertised by peer: inet-unicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Stale routes from peer are kept for: 300
Peer does not support Restarter functionality
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Peer supports 4 byte AS extension (peer-as 10)
Peer does not support Addpath
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        1
  Accepted prefixes:        1
  Suppressed due to damping: 0
  Advertised prefixes:      1
Last traffic (seconds): Received 9    Sent 5    Checked 5
Input messages: Total 36    Updates 2    Refreshes 0    Octets 718
Output messages: Total 37    Updates 1    Refreshes 0    Octets 796
Output Queue[0]: 0 (inet.0, inet-unicast)

Peer: 10.255.162.214+52193 AS 100 Local: 10.255.167.205+179 AS 100
  Type: Internal    State: Established (route reflector client)Flags: <Sync>
  Last State: OpenConfirm    Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress Cluster AddressFamily Rib-group Refresh>
  Address families configured: inet-unicast inet-vpn-unicast route-target
  Local Address: 10.255.167.205 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.255.162.214    Local ID: 10.255.167.205    Active Holdtime: 90
  Keepalive Interval: 30    Group index: 0    Peer index: 1

```

show bgp neighbor (CLNS)

```

user@host> show bgp neighbor
Peer: 10.245.245.1+179 AS 200 Local: 10.245.245.3+3770 AS 100
  Type: External    State: Established    Flags: <ImportEval Sync>
  Last State: OpenConfirm    Last Event: RecvKeepAlive
  Last Error: None
  Options: <Multihop Preference LocalAddress HoldTime AddressFamily PeerAS
  Rib-group Refresh>
  Address families configured: iso-vpn-unicast
  Local Address: 10.245.245.3 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.245.245.1    Local ID: 10.245.245.3    Active Holdtime: 90
  Keepalive Interval: 30    Peer index: 0
  NLRI advertised by peer: iso-vpn-unicast
  NLRI for this session: iso-vpn-unicast
  Peer supports Refresh capability (2)
Table bgp.isovpn.0 Bit: 10000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          3

```

```

Received prefixes:          3
Suppressed due to damping:  0
Advertised prefixes:       3
Table aaaa.iso.0
RIB State: BGP restart is complete
RIB State: VPN restart is complete
Send state: not advertising
Active prefixes:           3
Received prefixes:         3
Suppressed due to damping:  0
Last traffic (seconds): Received 6    Sent 5    Checked 5
Input messages: Total 1736    Updates 4    Refreshes 0    Octets 33385
Output messages: Total 1738    Updates 3    Refreshes 0    Octets 33305
Output Queue[0]: 0 (bgp.isovpn.0, iso-vpn-unicast)
Output Queue[1]: 0 (aaaa.iso.0, iso-vpn-unicast)

```

show bgp neighbor (Layer 2 VPN)

```

user@host> show bgp neighbor
Peer: 10.69.103.2      AS 65100 Local: 10.69.103.1      AS 65103
Type: External      State: Active      Flags: <ImportEval>
Last State: Idle      Last Event: Start
Last Error: None
Export: [ BGP-INET-import ]
Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily PeerAS
Refresh>
Address families configured: inet-unicast
Local Address: 10.69.103.1 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer: 10.69.104.2      AS 65100 Local: 10.69.104.1      AS 65104
Type: External      State: Active      Flags: <ImportEval>
Last State: Idle      Last Event: Start
Last Error: None
Export: [ BGP-L-import ]
Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily PeerAS
Refresh>
Address families configured: inet-labeled-unicast
Local Address: 10.69.104.1 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer: 10.255.14.182+179 AS 69      Local: 10.255.14.176+2131 AS 69
Type: Internal      State: Established  Flags: <ImportEval>
Last State: OpenConfirm Last Event: RecvKeepAlive
Last Error: None
Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
Rib-group Refresh>
Address families configured: inet-vpn-unicast l2vpn
Local Address: 10.255.14.176 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer ID: 10.255.14.182      Local ID: 10.255.14.176      Active Holdtime: 90
Keepalive Interval: 30
NLRI for restart configured on peer: inet-vpn-unicast l2vpn
NLRI advertised by peer: inet-vpn-unicast l2vpn
NLRI for this session: inet-vpn-unicast l2vpn
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-vpn-unicast l2vpn
NLRI peer can save forwarding state: inet-vpn-unicast l2vpn
NLRI that peer saved forwarding for: inet-vpn-unicast l2vpn
NLRI that restart is negotiated for: inet-vpn-unicast l2vpn

```

NLRI of received end-of-rib markers: inet-vpn-unicast l2vpn

Table bgp.l3vpn.0 Bit: 10000

RIB State: BGP restart in progress

RIB State: VPN restart in progress

Send state: in sync

Active prefixes: 10

Received prefixes: 10

Suppressed due to damping: 0

Table bgp.l2vpn.0 Bit: 20000

RIB State: BGP restart in progress

RIB State: VPN restart in progress

Send state: in sync

Active prefixes: 1

Received prefixes: 1

Suppressed due to damping: 0

Table BGP-INET.inet.0 Bit: 30000

RIB State: BGP restart in progress

RIB State: VPN restart in progress

Send state: in sync

Active prefixes: 2

Received prefixes: 2

Suppressed due to damping: 0

Table BGP-L.inet.0 Bit: 40000

RIB State: BGP restart in progress

RIB State: VPN restart in progress

Send state: in sync

Active prefixes: 2

Received prefixes: 2

Suppressed due to damping: 0

Table LDP.inet.0 Bit: 50000

RIB State: BGP restart is complete

RIB State: VPN restart in progress

Send state: in sync

Active prefixes: 1

Received prefixes: 1

Suppressed due to damping: 0

Table OSPF.inet.0 Bit: 60000

RIB State: BGP restart is complete

RIB State: VPN restart in progress

Send state: in sync

Active prefixes: 2

Received prefixes: 2

Suppressed due to damping: 0

Table RIP.inet.0 Bit: 70000

RIB State: BGP restart is complete

RIB State: VPN restart in progress

Send state: in sync

Active prefixes: 2

Received prefixes: 2

Suppressed due to damping: 0

Table STATIC.inet.0 Bit: 80000

RIB State: BGP restart is complete

RIB State: VPN restart in progress

Send state: in sync

Active prefixes: 1

Received prefixes: 1

Suppressed due to damping: 0

Table L2VPN.l2vpn.0 Bit: 90000

RIB State: BGP restart is complete

RIB State: VPN restart in progress

Send state: in sync

```

Active prefixes:          1
Received prefixes:        1
Suppressed due to damping: 0
Last traffic (seconds): Received 0    Sent 0    Checked 0
Input messages:  Total 14    Updates 13    Refreshes 0    Octets 1053
Output messages: Total 3    Updates 0    Refreshes 0    Octets 105
Output Queue[0]: 0 (bgp.l3vpn.0, inet-vpn-unicast)
Output Queue[1]: 0 (bgp.l2vpn.0, inet-vpn-unicast)
Output Queue[2]: 0 (BGP-INET.inet.0, inet-vpn-unicast)
Output Queue[3]: 0 (BGP-L.inet.0, inet-vpn-unicast)
Output Queue[4]: 0 (LDP.inet.0, inet-vpn-unicast)
Output Queue[5]: 0 (OSPF.inet.0, inet-vpn-unicast)
Output Queue[6]: 0 (RIP.inet.0, inet-vpn-unicast)
Output Queue[7]: 0 (STATIC.inet.0, inet-vpn-unicast)
Output Queue[8]: 0 (L2VPN.l2vpn.0, inet-vpn-unicast)

```

show bgp neighbor (Layer 3 VPN) (Not supported on the OCX Series.)

```

user@host> show bgp neighbor
Peer: 4.4.4.4+179      AS 10045 Local: 5.5.5.5+1214      AS 10045
Type: Internal      State: Established      Flags: <ImportEval>
Last State: OpenConfirm      Last Event: RecvKeepAlive
Last Error: None
Export: [ match-all ] Import: [ match-all ]
Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
Rib-group Refresh>
Address families configured: inet-vpn-unicast
Local Address: 5.5.5.5 Holdtime: 90 Preference: 170
Flags for NLRI inet-labeled-unicast: TrafficStatistics
Traffic Statistics: Options: all File: /var/log/bstat.log
                                size 131072 files 10

Traffic Statistics Interval: 60
Number of flaps: 0
Peer ID: 192.168.1.110      Local ID: 192.168.1.111      Active Holdtime: 90
Keepalive Interval: 30
NLRI for restart configured on peer: inet-vpn-unicast
NLRI advertised by peer: inet-vpn-unicast
NLRI for this session: inet-vpn-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-vpn-unicast
NLRI peer can save forwarding state: inet-vpn-unicast
NLRI that peer saved forwarding for: inet-vpn-unicast
NLRI that restart is negotiated for: inet-vpn-unicast
NLRI of received end-of-rib markers: inet-vpn-unicast
NLRI of all end-of-rib markers sent: inet-vpn-unicast
Table bgp.l3vpn.0 Bit: 10000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2
  Suppressed due to damping: 0
Table vpn-green.inet.0 Bit: 20001
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2

```

```

    Suppressed due to damping: 0
    Last traffic (seconds): Received 15    Sent 20    Checked 20
    Input messages: Total 40    Updates 2    Refreshes 0    Octets 856
    Output messages: Total 44    Updates 2    Refreshes 0    Octets 1066
    Output Queue[0]: 0 (bgp.13vpn.0, inet-vpn-unicast)
    Output Queue[1]: 0 (vpn-green.inet.0, inet-vpn-unicast)
    Trace options: detail packets
    Trace file: /var/log/bgpr.log size 131072 files 10

```

show bgp neighbor neighbor-address

```

user@host> show bgp neighbor 192.168.1.111
Peer: 10.255.245.12+179 AS 35 Local: 10.255.245.13+2884 AS 35
  Type: Internal    State: Established (route reflector client)Flags: <Sync>
  Last State: OpenConfirm    Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress HoldTime Cluster AddressFamily Rib-group
Refresh>
  Address families configured: inet-vpn-unicast inet-labeled-unicast
  Local Address: 10.255.245.13 Holdtime: 90 Preference: 170
  Flags for NLRI inet-vpn-unicast: AggregateLabel
  Flags for NLRI inet-labeled-unicast: AggregateLabel
  Number of flaps: 0
  Peer ID: 10.255.245.12    Local ID: 10.255.245.13    Active Holdtime: 90
  Keepalive Interval: 30
  BFD: disabled
  NLRI advertised by peer: inet-vpn-unicast inet-labeled-unicast
  NLRI for this session: inet-vpn-unicast inet-labeled-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 300
  Stale routes from peer are kept for: 60
  Restart time requested by this peer: 300
  NLRI that peer supports restart for: inet-unicast inet6-unicast
  NLRI that restart is negotiated for: inet-unicast inet6-unicast
  NLRI of received end-of-rib markers: inet-unicast inet6-unicast
  NLRI of all end-of-rib markers sent: inet-unicast inet6-unicast
  Table inet.0 Bit: 10000
    RIB State: restart is complete
    Send state: in sync
    Active prefixes: 4
    Received prefixes: 6
    Suppressed due to damping: 0
  Table inet6.0 Bit: 20000
    RIB State: restart is complete
    Send state: in sync
    Active prefixes: 0
    Received prefixes: 2
    Suppressed due to damping: 0
  Last traffic (seconds): Received 3    Sent 3    Checked 3
  Input messages: Total 9    Updates 6    Refreshes 0    Octets 403
  Output messages: Total 7    Updates 3    Refreshes 0    Octets 365
  Output Queue[0]: 0 (inet.0, inet-unicast)
  Output Queue[1]: 0 (inet6.0, inet6-unicast)
  Trace options: detail packets
  Trace file: /var/log/bgpr.log size 131072 files 10

```

show bgp neighbor neighbor-address

```

user@host> show bgp neighbor 192.168.4.222
Peer: 192.168.4.222+4902 AS 65501 Local: 192.168.4.221+179 AS 65500
  Type: External    State: Established    Flags: <Sync>

```



```

Last State: OpenConfirm   Last Event: RecvKeepAlive
Last Error: Cease
Export: [ export-policy ] Import: [ import-policy ]
Options: <Preference HoldTime AddressFamily PeerAS PrefixLimit Refresh>
Address families configured: inet-unicast inet-multicast
Holdtime: 60000 Preference: 170
Number of flaps: 4
Last flap event: RecvUpdate
Error: 'Cease' Sent: 5 Recv: 0
Peer ID: 10.255.245.6      Local ID: 10.255.245.5      Active Holdtime: 60000
Keepalive Interval: 20000   Peer index: 0
BFD: disabled, down
Local Interface: fxp0.0
NLRI advertised by peer: inet-unicast inet-multicast
NLRI for this session: inet-unicast inet-multicast
Peer supports Refresh capability (2)
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          8
  Received prefixes:        10
  Accepted prefixes:        10
  Suppressed due to damping: 0
  Advertised prefixes:      3
Table inet.2 Bit: 20000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        0
  Accepted prefixes:        0
  Suppressed due to damping: 0
  Advertised prefixes:      0
Last traffic (seconds): Received 357 Sent 357 Checked 357
Input messages: Total 4 Updates 2 Refreshes 0 Octets 211
Output messages: Total 4 Updates 1 Refreshes 0 Octets 147
Output Queue[0]: 0 (inet.0, inet-unicast)
Output Queue[1]: 0 (inet.2, inet-multicast)
Trace options: all
Trace file: /var/log/bgp size 10485760 files 10

```

show bgp neighbor neighbor-address (BGP Graceful Restart Enabled)

```
user@router> show bgp neighbor 10.255.255.16
```

```

Peer: 10.255.255.16 AS 100      Local: 10.255.255.12 AS 100
Type: Internal   State: Active   Flags: <>
Last State: Idle   Last Event: Start
Last Error: None
Options: <Preference LocalAddress AddressFamily Rib-group Refresh>
Options: <LLGR>
Address families configured: l2vpn
Local Address: 10.255.255.12 Holdtime: 90 Preference: 170
NLRI l2vpn:
Number of flaps: 6
Last flap event: Restart
NLRI we are holding stale routes for: inet-vpn-unicast
Time until stale routes are deleted or become long-lived stale: 00:01:57
Time until end-of-rib is assumed for stale routes: 00:04:43
Table bgp.l3vpn.0
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete

```

```

Send state: not advertising
Active prefixes:          0
Received prefixes:       7
Accepted prefixes:       7
Suppressed due to damping: 0
Table foo.inet.0 Bit: 30000
RIB State: BGP restart is complete
RIB State: VPN restart is complete
Send state: not in sync
Active prefixes:          0
Received prefixes:       7
Accepted prefixes:       7
Suppressed due to damping: 0

```

show bgp neighbor neighbor-address (BGP Long-Lived Graceful Restart)

```

user@router> show bgp neighbor 10.4.12.11

Peer: 10.4.12.11 AS 100          Local: 10.6.128.225 AS 100
Type: Internal      State: Active      Flags: <>
Last State: Idle      Last Event: Start
Last Error: None
Export: [ foo ]
Options: <Preference LocalAddress Refresh GracefulRestart>
Options: <LLGR>
Local Address: 10.6.128.225 Holdtime: 90 Preference: 170
Number of flaps: 3
Last flap event: Restart
Error: 'Cease' Sent: 0 Recv: 1
Time until long-lived stale routes deleted: inet-vpn-unicast 10:00:22
route-target 10:00:22
Table bgp.l3vpn.0
RIB State: BGP restart is complete
RIB State: VPN restart is complete
Send state: not advertising
Active prefixes:          0
Received prefixes:       7
Accepted prefixes:       7
Suppressed due to damping: 0
Table foo.inet.0 Bit: 30000
RIB State: BGP restart is complete
RIB State: VPN restart is complete
Send state: not in sync
Active prefixes:          0
Received prefixes:       7
Accepted prefixes:       7
Suppressed due to damping: 0

```

show bgp neighbor orf neighbor-address detail

```

user@host > show bgp neighbor orf 192.168.165.56 detail
Peer: 192.168.165.56+179 Type: External
Group: ext1

inet-unicast
Filter updates rcv:          1 Immediate:          1
Filter: prefix-based receive
Received filter entries:
seq 1: prefix 2.2.2.2/32: minlen 32: maxlen 32: match deny:

inet6-unicast

```

```
Filter updates rcv:          0 Immediate:          1
Filter: prefix-based receive
Received filter entries:
*:*
```

show bgp summary

List of Syntax	Syntax on page 4124 Syntax (EX Series Switch and QFX Series) on page 4124
Syntax	<pre>show bgp summary <exact-instance <i>instance-name</i>> <group <i>group-name</i>> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switch and QFX Series)	<pre>show bgp summary <exact-instance <i>instance-name</i>> <instance <i>instance-name</i>></pre>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series. exact-instance option introduced in Junos OS Release 11.4. group option introduced in Junos OS Release 13.3.
Description	Display BGP summary information.
Options	<p>none—Display BGP summary information for all routing instances.</p> <p>exact-instance <i>instance-name</i>—(Optional) Display information for the specified instance only.</p> <p>group—Display overview of bgp information for a particular group</p> <p>instance <i>instance-name</i>—(Optional) Display information for all routing instances whose name begins with this string (for example, cust1, cust11, and cust111 are all displayed when you run the show bgp summary instance cust1 command). The instance name can be master for the main instance, or any valid configured instance name or its prefix.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show bgp summary (When a Peer Is Not Established) on page 4127 show bgp summary (When a Peer Is Established) on page 4127 show bgp summary (CLNS) on page 4127 show bgp summary (Layer 2 VPN) on page 4128 show bgp summary (Layer 3 VPN) on page 4128 show bgp summary group on page 4128 show bgp summary (BGP Graceful Restart or Long-Lived Graceful Restart) on page 4129

Output Fields Table 323 describes the output fields for the **show bgp summary** command. Output fields are listed in the approximate order in which they appear.

Table 323: show bgp summary Output Fields

Field Name	Field Description
Groups	Number of BGP groups.
Peers	Number of BGP peers.
Down peers	Number of down BGP peers.
Table	Name of routing table.
Tot Paths	Total number of paths.
Act Paths	Number of active routes.
Suppressed	Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.
History	Number of withdrawn routes stored locally to keep track of damping history.
Damp State	Number of routes with a figure of merit greater than zero, but still active because the value has not reached the threshold at which suppression occurs.
Pending	Routes in process by BGP import policy.
Peer	Address of each BGP peer. Each peer has one line of output.
AS	Peer's AS number.
InPkt	Number of packets received from the peer.
OutPkt	Number of packets sent to the peer.
OutQ	Number of BGP packets that are queued to be transmitted to a particular neighbor. It normally is 0 because the queue usually is emptied quickly.
Flaps	Number of times the BGP session has gone down and then come back up.
Last Up/Down	Last time since the neighbor transitioned to or from the established state.

Table 323: show bgp summary Output Fields (*continued*)

Field Name	Field Description
State #Active /Received/Accepted /Damped	<p>Multipurpose field that displays information about BGP peer sessions. The field's contents depend upon whether a session is established and whether it was established on the main routing device or in a routing instance.</p> <ul style="list-style-type: none"> If a peer is not established, the field shows the state of the peer session: Active, Connect, or Idle. In general, the Idle state is the first stage of a connection. BGP is waiting for a Start event. A session can be idle for other reasons as well. The reason that a session is idle is sometimes displayed. For example: Idle (Removal in progress) or Idle (LicenseFailure). If a BGP session is established on the main routing device, the field shows the number of active, received, accepted, and damped routes that are received from a neighbor and appear in the inet.0 (main) and inet.2 (multicast) routing tables. For example, 8/10/10/2 and 2/4/4/0 indicate the following: <ul style="list-style-type: none"> 8 active routes, 10 received routes, 10 accepted routes, and 2 damped routes from a BGP peer appear in the inet.0 routing table. 2 active routes, 4 received routes, 4 accepted routes, and no damped routes from a BGP peer appear in the inet.2 routing table. If a BGP session is established in a routing instance, the field indicates the established (Establ) state, identifies the specific routing table that receives BGP updates, and shows the number of active, received, and damped routes that are received from a neighbor. For example, Establ VPN-AB.inet.0: 2/4/0 indicates the following: <ul style="list-style-type: none"> The BGP session is established. Routes are received in the VPN-AB.inet.0 routing table. The local routing device has two active routes, four received routes, and no damped routes from a BGP peer. <p>When a BGP session is established, the peers are exchanging update messages.</p> <p>NOTE: When graceful restart or LLGR helper mode is active, the RIB information is now displayed by the show bgp summary command. If a BGP session is established on the main routing device, the field shows the number of active, received, accepted, and damped routes that are received from a neighbor and appear in the inet.0 (main) and inet.2 (multicast) routing tables. For example, 8/10/10/2 and 2/4/4/0 indicate the following:</p> <ul style="list-style-type: none"> 8 active routes, 10 received routes, 10 accepted routes, and 2 damped routes from a BGP peer appear in the inet.0 routing table. 2 active routes, 4 received routes, 4 accepted routes, and no damped routes from a BGP peer appear in the inet.2 routing table.

Sample Output

show bgp summary (When a Peer Is Not Established)

```

user@host> show bgp summary
Groups: 2 Peers: 4 Down peers: 1
Table          Tot Paths  Act Paths Suppressed    History  Damp State   Pending
inet.0          6          4          0          0          0          0          0
Peer           AS      InPkt   OutPkt   OutQ   Flaps  Last Up/Dwn
State|#Active/Received/Damped...
10.0.0.3        65002      86      90      0       2      42:54 0/0/0

0/0/0
10.0.0.4        65002      90      91      0       1      42:54 0/2/0

0/0/0
10.0.0.6        65002      87      90      0       3          3 Active
10.1.12.1       65001      89      89      0       1      42:54 4/4/0

0/0/0

```

show bgp summary (When a Peer Is Established)

```

user@host> show bgp summary
Groups: 1 Peers: 3 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History  Damp State   Pending
inet.0          6          4          0          0          0          0          0
Peer           AS      InPkt   OutPkt   OutQ   Flaps  Last Up/Dwn
State|#Active/Received/Damped...
10.0.0.2        65002    88675    88652      0       2      42:38 2/4/0

0/0/0
10.0.0.3        65002    54528    54532      0       1     2w4d22h 0/0/0

0/0/0
10.0.0.4        65002    51597    51584      0       0     2w3d22h 2/2/0

0/0/0

user@host> show bgp summary logical-system R3
Groups: 2 Peers: 2 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History  Damp State   Pending
bgp.13vpn.0      2          2          0          0          0          0          0
Peer           AS      InPkt   OutPkt   OutQ   Flaps  Last Up/Dwn
State|#Active/Received/Accepted/Damped...
1.1.1.2          2       204     206      0       0      1:30:59
Establ
  bgp.13vpn.0: 2/2/2/0
  red.inet.0: 2/2/2/0
10.1.1.10        3       206     207      0       0      1:31:36
Establ
  red.inet.0: 2/2/2/0

```

show bgp summary (CLNS)

```

user@host> show bgp summary
Groups: 1 Peers: 1 Down peers: 0
Peer           AS      InPkt   OutPkt   OutQ   Flaps  Last Up/Dwn
State|#Active/Received/Damped...
10.245.245.1    200     1735    1737      0       0     14:26:12 Establ

```

```

bgp.isovpn.0: 3/3/0
aaaa.iso.0: 3/3/0

```

show bgp summary (Layer 2 VPN)

```

user@host> show bgp summary
Groups: 1 Peers: 5 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
bgp.l2vpn.0 1 1 0 0 0 0
inet.0 0 0 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last
Up/Dwn State|#Active/Received/Damped...
10.255.245.35 65299 72 74 0 1 19:00 Establ
  bgp.l2vpn.0: 1/1/0
  frame-vpn.l2vpn.0: 1/1/0
10.255.245.36 65299 2164 2423 0 4 19:50 Establ
  bgp.l2vpn.0: 0/0/0
  frame-vpn.l2vpn.0: 0/0/0
10.255.245.37 65299 36 37 0 4 17:07 Establ
  inet.0: 0/0/0
10.255.245.39 65299 138 168 0 6 53:48 Establ
  bgp.l2vpn.0: 0/0/0
  frame-vpn.l2vpn.0: 0/0/0
10.255.245.69 65299 134 140 0 6 53:42 Establ
  inet.0: 0/0/0

```

show bgp summary (Layer 3 VPN)

```

user@host> show bgp summary
Groups: 2 Peers: 2 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
bgp.l3vpn.0 2 2 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Damped...
10.39.1.5 2 21 22 0 0 6:26 Establ
  VPN-AB.inet.0: 1/1/0
10.255.71.15 1 19 21 0 0 6:17 Establ
  bgp.l3vpn.0: 2/2/0
  VPN-A.inet.0: 1/1/0
  VPN-AB.inet.0: 2/2/0
  VPN-B.inet.0: 1/1/0

```

show bgp summary group

```

user@host> show bgp summary group Group2
Groups: 3 Peers: 3 Down peers: 3
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet.0 0 0 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.0.0.1 56 0 0 0 0 51
Idle

user@host> show bgp summary logical-system R3 group toR4
Groups: 2 Peers: 2 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
bgp.l3vpn.0 2 2 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.1.1.10 3 207 207 0 0 1:31:40

```



```
Establ
red.inet.0: 2/2/2/0
```

show bgp summary (BGP Graceful Restart or Long-Lived Graceful Restart)

```
user@router> show route receive-protocol bgp 10.4.12.11 detail
Groups: 2 Peers: 9 Down peers: 1
...
Peer          AS      InPkt    OutPkt    OutQ    Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.255.255.16 100        7         6         0         4         4
Idle
  bgp.13vpn.0: 0/7/7/0
  foo.inet.0: 0/7/7/0
```

show policy damping

List of Syntax	Syntax on page 4130 Syntax (EX Series Switch and QFX Series) on page 4130
Syntax	<pre>show policy damping <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switch and QFX Series)	show policy damping
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Display information about BGP route flap damping parameters.
Options	<p>none—Display information about BGP route flap damping parameters.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Additional Information	In the output from this command, figure-of-merit values correlate with the probability of future instability of a routing device. Routes with higher figure-of-merit values are suppressed for longer periods of time. The figure-of-merit value decays exponentially over time. A figure-of-merit value of zero is assigned to each new route. The value is increased each time the route is withdrawn or readvertised, or when one of its path attributes changes.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear bgp damping on page 4092 • show route damping on page 3430
List of Sample Output	show policy damping on page 4131
Output Fields	Table 324 describes the output fields for the show policy damping command. Output fields are listed in the approximate order in which they appear.

Table 324: show policy damping Output Fields

Field Name	Field Description
Halflife	Decay half-life, in minutes. The value represents the period during which the accumulated figure-of-merit value is reduced by half if the route remains stable. If a route has flapped, but then becomes stable, the figure-of-merit value for the route decays exponentially. For example, for a route with a figure-of-merit value of 1500, if no incidents occur, its figure-of-merit value is reduced to 750 after 15 minutes and to 375 after another 15 minutes.

Table 324: show policy damping Output Fields (*continued*)

Field Name	Field Description
Reuse merit	Figure-of-merit value below which a suppressed route can be used again. A suppressed route becomes reusable when its figure-of-merit value decays to a value below a reuse threshold, and the route once again is considered usable and can be installed in the forwarding table and exported from the routing table.
Suppress/cutoff merit	Figure-of-merit value above which a route is suppressed for use or inclusion in advertisements. When a route's figure-of-merit value reaches a particular level, called the cutoff or suppression threshold, the route is suppressed. When a route is suppressed, the routing table no longer installs the route into the forwarding table and no longer exports this route to any of the routing protocols.
Maximum suppress time	Maximum hold-down time, in minutes. The value represents the maximum time that a route can be suppressed no matter how unstable it has been before this period of stability.
Computed values	<ul style="list-style-type: none"> • Merit ceiling—Maximum merit that a flapping route can collect. • Maximum decay—Maximum decay half-life, in minutes.

Sample Output

show policy damping

```

user@host> show policy damping
Default damping information:
  Halflife: 15 minutes
  Reuse merit: 750 Suppress/cutoff merit: 3000
  Maximum suppress time: 60 minutes
  Computed values:
    Merit ceiling: 12110
    Maximum decay: 6193
Damping information for "standard-damping":
  Halflife: 10 minutes
  Reuse merit: 4000 Suppress/cutoff merit: 8000
  Maximum suppress time: 30 minutes
  Computed values:
    Merit ceiling: 32120
    Maximum decay: 12453

```

show route damping

List of Syntax	Syntax on page 4132 Syntax (EX Series Switch and QFX Series) on page 4132
Syntax	<code>show route damping (decayed history suppressed)</code> <code><brief detail extensive terse></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switch and QFX Series)	<code>show route damping (decayed history suppressed)</code> <code><brief detail extensive terse></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display the BGP routes for which updates might have been reduced because of route flap damping.
Options	brief detail extensive terse —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief. decayed —Display route damping entries that might no longer be valid, but are not suppressed. history —Display entries that have already been withdrawn, but have been logged. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. suppressed —Display entries that have been suppressed and are no longer being installed into the forwarding table or exported by routing protocols.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• clear bgp damping on page 4092• show policy damping on page 4130
List of Sample Output	show route damping decayed detail on page 4135 show route damping history on page 4136 show route damping history detail on page 4136
Output Fields	Table 298 lists the output fields for the show route damping command. Output fields are listed in the approximate order in which they appear.

Table 325: show route damping Output Fields

Field Name	Field Description	Level of Output
<i>routing-table-name</i>	Name of the routing table—for example, inet.0 .	All levels
destinations	Number of destinations for which there are routes in the routing table.	All levels
number routes	Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> • active • holdddown (routes that are in a pending state before being declared inactive) • hidden (the routes are not used because of a routing policy) 	All levels
destination-prefix (entry, announced)	Destination prefix. The entry value is the number of routes for this destination, and the announced value is the number of routes being announced for this destination.	detail extensive
[protocol, preference]	Protocol from which the route was learned and the preference value for the route. <ul style="list-style-type: none"> • +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • -—A hyphen indicates the last active route. • *—An asterisk indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route. <p>In every routing metric except for the BGP LocalPref attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the LocalPref value in the Preference2 field. For example, if the LocalPref value for Route 1 is 100, the Preference2 value is -101. If the LocalPref value for Route 2 is 155, the Preference2 value is -156. Route 2 is preferred because it has a higher LocalPref value and a lower Preference2 value.</p>	All levels
Next-hop reference count	Number of references made to the next hop.	detail extensive
Source	IP address of the route source.	detail extensive
Next hop	Network layer address of the directly reachable neighboring system.	detail extensive
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word Selected .	detail extensive
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to derive a forwarding next hop.	detail extensive
Indirect next hop	Index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.	detail extensive
State	Flags for this route. For a description of possible values for this field, see the output field table for the show route detail command.	detail extensive

Table 325: show route damping Output Fields (*continued*)

Field Name	Field Description	Level of Output
Local AS	AS number of the local routing device.	detail extensive
Peer AS	AS number of the peer routing device.	detail extensive
Age	How long the route has been known.	detail extensive
Metric	Metric for the route.	detail extensive
Task	Name of the protocol that has added the route.	detail extensive
Announcement bits	List of protocols that announce this route. <i>n-Resolve inet</i> indicates that the route is used for route resolution for next hops found in the routing table. <i>n</i> is an index used by Juniper Networks customer support only.	detail extensive
AS path	<p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> • I—IGP. • E—EGP. • ?—Incomplete; typically, the AS path was aggregated. <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> • []—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device or if AS path prepending is configured. • { }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • ()—Parentheses enclose a confederation. • ([])—Parentheses and brackets enclose a confederation set. <p>NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p>	All levels
to	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.	brief none
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word Selected .	brief none
Communities	Community path attribute for the route. See the output field table for the show route detail command.	detail extensive
Localpref	Local preference value included in the route.	All levels
Router ID	BGP router ID as advertised by the neighbor in the open message.	detail extensive

Table 325: show route damping Output Fields (*continued*)

Field Name	Field Description	Level of Output
Merit (last update/now)	Last updated and current figure-of-merit value.	detail extensive
damping-parameters	Name that identifies the damping parameters used, which is defined in the damping statement at the [edit policy-options] hierarchy level.	detail extensive
Last update	Time of most recent change in path attributes.	detail extensive
First update	Time of first change in path attributes, which started the route damping process.	detail extensive
Flaps	Number of times the route has gone up or down or its path attributes have changed.	detail extensive
Suppressed	(suppressed keyword only) This route is currently suppressed. A suppressed route does not appear in the forwarding table and routing protocols do not export it.	All levels
Reusable in	(suppressed keyword only) Time when a suppressed route will again be available.	All levels
Preference will be	(suppressed keyword only) Preference value that will be applied to the route when it is again active.	All levels

Sample Output

show route damping decayed detail

```

user@host> show route damping decayed detail
inet.0: 173319 destinations, 1533668 routes (172625 active, 4 holddown, 108083
hidden)
10.0.111.0/24 (7 entries, 1 announced)
  *BGP      Preference: 170/-101
            Next-hop reference count: 151973
            Source: 172.23.2.129
            Next hop: via so-1/2/0.0
            Next hop: via so-5/1/0.0, selected
            Next hop: via so-6/0/0.0
            Protocol next hop: 172.23.2.129
            Indirect next hop: 89a1a00 264185
            State: <Active Ext>
            Local AS: 65000 Peer AS: 65490
            Age: 3:28      Metric2: 0
            Task: BGP_65490.172.23.2.129+179
            Announcement bits (6): 0-KRT 1-RT 4-KRT 5-BGP.0.0.0.0+179

  6-Resolve tree 2 7-Resolve tree 3
    AS path: 65490 65520 65525 65525 65525 65525 I ()
    Communities: 65501:390 65501:2000 65501:3000 65504:701
    Localpref: 100
    Router ID: 172.23.2.129
    Merit (last update/now): 1934/1790
    damping-parameters: damping-high

```

Last update: 00:03:28 First update: 00:06:40
Flaps: 2

show route damping history

```
user@host> show route damping history
inet.0: 173320 destinations, 1533529 routes (172624 active, 6 holddown, 108122
hidden)
+ = Active Route, - = Last Active, * = Both

10.108.0.0/15      [BGP ] 2d 22:47:58, localpref 100
                  AS path: 65220 65501 65502 I
                  > to 192.168.60.85 via so-3/1/0.0
```

show route damping history detail

```
user@host> show route damping history detail
inet.0: 173319 destinations, 1533435 routes (172627 active, 2 holddown, 108105
hidden)
10.108.0.0/15 (3 entries, 1 announced)
    BGP                /-101
        Next-hop reference count: 69058
        Source: 192.168.60.85
        Next hop: 192.168.60.85 via so-3/1/0.0, selected
        State: <Hidden Ext>
        Inactive reason: Unusable path
        Local AS: 65000 Peer AS: 65220
        Age: 2d 22:48:10
        Task: BGP_65220.192.168.60.85+179
        AS path: 65220 65501 65502 I ()
        Communities: 65501:390 65501:2000 65501:3000 65504:3561
        Localpref: 100
        Router ID: 192.168.80.25
        Merit (last update/now): 1000/932
        damping-parameters: set-normal
        Last update: 00:01:05 First update: 00:01:05
        Flaps: 1
```


show route detail

List of Syntax	Syntax on page 4137 Syntax (EX Series Switches) on page 4137
Syntax	show route detail <destination-prefix> <logical-system (all logical-system-name)>
Syntax (EX Series Switches)	show route detail <destination-prefix>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display detailed information about the active entries in the routing tables.
Options	<p>none—Display all active entries in the routing table on all systems.</p> <p>destination-prefix—(Optional) Display active entries for the specified address or range of addresses.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show route detail on page 4146 show route detail (with BGP Multipath) on page 4152 show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 4153 show route label detail (Multipoint LDP with Multicast-Only Fast Reroute) on page 4153
Output Fields	Table 299 describes the output fields for the show route detail command. Output fields are listed in the approximate order in which they appear.

Table 326: show route detail Output Fields

Field Name	Field Description
<i>routing-table-name</i>	Name of the routing table (for example, inet.0).
<i>number destinations</i>	Number of destinations for which there are routes in the routing table.
<i>number routes</i>	Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> active (routes that are active) holddown (routes that are in the pending state before being declared inactive) hidden (routes that are not used because of a routing policy)

Table 326: show route detail Output Fields (*continued*)

Field Name	Field Description
<i>route-destination</i> (entry, announced)	<p>Route destination (for example:10.0.0.1/24). The entry value is the number of routes for this destination, and the announced value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:</p> <ul style="list-style-type: none"> • MPLS-label (for example, 80001). • interface-name (for example, ge-1/0/2). • neighbor-address:control-word-status:encapsulation type:vc-id:source (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96). <ul style="list-style-type: none"> • neighbor-address—Address of the neighbor. • control-word-status—Whether the use of the control word has been negotiated for this virtual circuit: NoCtrlWord or CtrlWord. • encapsulation type—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport. • vc-id—Virtual circuit identifier. • source—Source of the advertisement: Local or Remote. • source—Source of the advertisement: Local or Remote.
label stacking	<p>(Next-to-the-last-hop routing device for MPLS only) Depth of the MPLS label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels.</p> <ul style="list-style-type: none"> • S=0 route indicates that a packet with an incoming label stack depth of 2 or more exits this routing device with one fewer label (the label-popping operation is performed). • If there is no S= information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).
[<i>protocol, preference</i>]	<p>Protocol from which the route was learned and the preference value for the route.</p> <ul style="list-style-type: none"> • +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • - —A hyphen indicates the last active route. • *—An asterisk indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route. <p>In every routing metric except for the BGP LocalPref attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the LocalPref value in the Preference2 field. For example, if the LocalPref value for Route 1 is 100, the Preference2 value is -101. If the LocalPref value for Route 2 is 155, the Preference2 value is -156. Route 2 is preferred because it has a higher LocalPref value and a lower Preference2 value.</p>
Level	<p>(IS-IS only). In IS-IS, a single AS can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.</p>
Route Distinguisher	IP subnet augmented with a 64-bit prefix.
PMSI	Provider multicast service interface (MVPN routing table).

Table 326: show route detail Output Fields (*continued*)

Field Name	Field Description
Next-hop type	Type of next hop. For a description of possible values for this field, see Table 300 .
Next-hop reference count	Number of references made to the next hop.
Flood nexthop branches exceed maximum message	Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.
Source	IP address of the route source.
Next hop	Network layer address of the directly reachable neighboring system.
via	<p>Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word Selected. This field can also contain the following information:</p> <ul style="list-style-type: none"> • Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.
Label-switched-path lsp-path-name	Name of the LSP used to reach the next hop.
Label operation	MPLS label and operation occurring at this routing device. The operation can be pop (where a label is removed from the top of the stack), push (where another label is added to the label stack), or swap (where a label is replaced by another label).
Interface	(Local only) Local interface name.
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to derive a forwarding next hop.
Indirect next hop	Index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.
State	State of the route (a route can be in more than one state). See Table 301 .
Local AS	AS number of the local routing device.
Age	How long the route has been known.
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.

Table 326: show route detail Output Fields (*continued*)

Field Name	Field Description
Metric	Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
MED-plus-IGP	Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.
TTL-Action	For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signaled and LDP-signaled LSPs or for specific VRF routing instances. For sample output, see show route table .
Task	Name of the protocol that has added the route.
Announcement bits	The number of BGP peers or protocols to which Junos OS has announced this route, followed by the list of the recipients of the announcement. Junos OS can also announce the route to the KRT for installing the route into the Packet Forwarding Engine, to a resolve tree, a L2 VC, or even a VPN. For example, <i>n-Resolve inet</i> indicates that the specified route is used for route resolution for next hops found in the routing table. <ul style="list-style-type: none"> <i>n</i>—An index used by Juniper Networks customer support only.
AS path	<p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> I—IGP. E—EGP. Recorded—The AS path is recorded by the sample process (sampled). ?—Incomplete; typically, the AS path was aggregated. <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> []—Brackets enclose the number that precedes the AS path. This number represents the number of ASs present in the AS path, when calculated as defined in RFC 4271. This value is used in the AS-path merge process, as defined in RFC 4893. []—If more than one AS number is configured on the routing device, or if AS path prepending is configured, brackets enclose the local AS number associated with the AS path. { }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. ()—Parentheses enclose a confederation. ([])—Parentheses and brackets enclose a confederation set. <p>NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p>

Table 326: show route detail Output Fields (*continued*)

Field Name	Field Description
validation-state	<p>(BGP-learned routes) Validation status of the route:</p> <ul style="list-style-type: none"> • Invalid—Indicates that the prefix is found, but either the corresponding AS received from the EBGp peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • Unknown—Indicates that the prefix is not among the prefixes or prefix ranges in the database. • Unverified—Indicates that the origin of the prefix is not verified against the database. This is because the database got populated and the validation is not called for in the BGP import policy, although origin validation is enabled, or the origin validation is not enabled for the BGP peers. • Valid—Indicates that the prefix and autonomous system pair are found in the database.
FECs bound to route	Point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.
Primary Upstream	When multipoint LDP with multicast-only fast reroute (MoFRR) is configured, the primary upstream path. MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path.
RPF Nexthops	When multipoint LDP with MoFRR is configured, the reverse-path forwarding (RPF) next-hop information. Data packets are received from both the primary path and the secondary paths. The redundant packets are discarded at topology merge points due to the RPF checks.
Label	Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.
weight	Value used to distinguish MoFRR primary and backup routes. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.
VC Label	MPLS label assigned to the Layer 2 circuit virtual connection.
MTU	Maximum transmission unit (MTU) of the Layer 2 circuit.
VLAN ID	VLAN identifier of the Layer 2 circuit.
Prefixes bound to route	Forwarding equivalent class (FEC) bound to this route. Applicable only to routes installed by LDP.
Communities	Community path attribute for the route. See Table 302 for all possible values for this field.
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).
control flags	Control flags: none or Site Down .
mtu	Maximum transmission unit (MTU) information.
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.
status vector	Layer 2 VPN and VPLS network layer reachability information (NLRI).

Table 326: show route detail Output Fields (*continued*)

Field Name	Field Description
Accepted Multipath	Current active path when BGP multipath is configured.
Accepted LongLivedStale	The LongLivedStale flag indicates that the route was marked LLGR-stale by this router, as part of the operation of LLGR receiver mode. Either this flag or the LongLivedStaleImport flag may be displayed for a route. Neither of these flags are displayed at the same time as the Stale (ordinary GR stale) flag.
Accepted LongLivedStaleImport	<p>The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy. Either this flag or the LongLivedStale flag may be displayed for a route. Neither of these flags are displayed at the same time as the Stale (ordinary GR stale) flag.</p> <p>Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and import into the inet.0 routing table</p>
ImportAccepted LongLivedStaleImport	<p>Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and imported into the inet.0 routing table</p> <p>The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy.</p>
Accepted MultipathContrib	Path currently contributing to BGP multipath.
Localpref	Local preference value included in the route.
Router ID	BGP router ID as advertised by the neighbor in the open message.
Primary Routing Table	In a routing table group, the name of the primary routing table in which the route resides.
Secondary Tables	In a routing table group, the name of one or more secondary tables in which the route resides.

[Table 300](#) describes all possible values for the Next-hop Types output field.

Table 327: Next-Hop Types Output Field Values

Next-Hop Type	Description
Broadcast (bcast)	Broadcast next hop.
Deny	Deny next hop.
Discard	Discard next hop.
Flood	Flood next hop. Consists of components called branches, up to a maximum of 32 branches. Each flood next-hop branch sends a copy of the traffic to the forwarding interface. Used by P2MP RSVP, P2MP LDP, P2MP CCC, and multicast.
Hold	Next hop is waiting to be resolved into a unicast or multicast type.

Table 327: Next-Hop Types Output Field Values (*continued*)

Next-Hop Type	Description
Indexed (idxd)	Indexed next hop.
Indirect (indr)	Used with applications that have a protocol next hop address that is remote. You are likely to see this next-hop type for internal BGP (IBGP) routes when the BGP next hop is a BGP neighbor that is not directly connected.
Interface	Used for a network address assigned to an interface. Unlike the router next hop, the interface next hop does not reference any specific node on the network.
Local (locl)	Local address on an interface. This next-hop type causes packets with this destination address to be received locally.
Multicast (mcst)	Wire multicast next hop (limited to the LAN).
Multicast discard (mdsc)	Multicast discard.
Multicast group (mgrp)	Multicast group member.
Receive (recv)	Receive.
Reject (rjct)	Discard. An ICMP unreachable message was sent.
Resolve (rslv)	Resolving next hop.
Routed multicast (mcrtr)	Regular multicast next hop.
Router	<p>A specific node or set of nodes to which the routing device forwards packets that match the route prefix.</p> <p>To qualify as next-hop type router, the route must meet the following criteria:</p> <ul style="list-style-type: none"> • Must not be a direct or local subnet for the routing device. • Must have a next hop that is directly connected to the routing device.
Table	Routing table next hop.
Unicast (ucst)	Unicast.
Unilist (ulst)	List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.

Table 301 describes all possible values for the State output field. A route can be in more than one state (for example, <Active NoReadvrt Int Ext>).

Table 328: State Output Field Values

Value	Description
Accounting	Route needs accounting.
Active	Route is active.
Always Compare MED	Path with a lower multiple exit discriminator (MED) is available.
AS path	Shorter AS path is available.
Cisco Non-deterministic MED selection	Route is a clone.
Clone	Cisco nondeterministic MED is enabled and a path with a lower MED is available.
Cluster list length	Length of cluster list sent by the route reflector.
Delete	Route has been deleted.
Ex	Exterior route.
Ext	BGP route received from an external BGP neighbor.
FlashAll	Forces all protocols to be notified of a change to any route, active or inactive, for a prefix. When not set, protocols are informed of a prefix only when the active route changes.
Hidden	Route not used because of routing policy.
IfCheck	Route needs forwarding RPF check.
IGP metric	Path through next hop with lower IGP metric is available.
Inactive reason	Flags for this route, which was not selected as best for a particular destination.
Initial	Route being added.
Int	Interior route.
Int Ext	BGP route received from an internal BGP peer or a BGP confederation peer.
Interior > Exterior > Exterior via Interior	Direct, static, IGP, or EBGp path is available.
Local Preference	Path with a higher local preference value is available.

Table 328: State Output Field Values (*continued*)

Value	Description
Martian	Route is a martian (ignored because it is obviously invalid).
MartianOK	Route exempt from martian filtering.
Next hop address	Path with lower metric next hop is available.
No difference	Path from neighbor with lower IP address is available.
NoReadvrt	Route not to be advertised.
NotBest	Route not chosen because it does not have the lowest MED.
Not Best in its group	Incoming BGP AS is not the best of a group (only one AS can be the best).
NotInstall	Route not to be installed in the forwarding table.
Number of gateways	Path with a greater number of next hops is available.
Origin	Path with a lower origin code is available.
Pending	Route pending because of a hold-down configured on another route.
Release	Route scheduled for release.
RIB preference	Route from a higher-numbered routing table is available.
Route Distinguisher	64-bit prefix added to IP subnets to make them unique.
Route Metric or MED comparison	Route with a lower metric or MED is available.
Route Preference	Route with lower preference value is available
Router ID	Path through a neighbor with lower ID is available.
Secondary	Route not a primary route.
Unusable path	Path is not usable because of one of the following conditions: <ul style="list-style-type: none"> • The route is damped. • The route is rejected by an import policy. • The route is unresolved.
Update source	Last tiebreaker is the lowest IP address value.

Table 302 describes the possible values for the Communities output field.

Table 329: Communities Output Field Values

Value	Description
<i>area-number</i>	4 bytes, encoding a 32-bit area number. For AS-external routes, the value is 0 . A nonzero value identifies the route as internal to the OSPF domain, and as within the identified area. Area numbers are relative to a particular OSPF domain.
bandwidth: local AS number:link-bandwidth-number	Link-bandwidth community value used for unequal-cost load balancing. When BGP has several candidate paths available for multipath purposes, it does not perform unequal-cost load balancing according to the link-bandwidth community unless all candidate paths have this attribute.
domain-id	Unique configurable number that identifies the OSPF domain.
domain-id-vendor	Unique configurable number that further identifies the OSPF domain.
<i>link-bandwidth-number</i>	Link-bandwidth number: from 0 through 4,294,967,295 (bytes per second).
<i>local AS number</i>	Local AS number: from 1 through 65,535 .
<i>options</i>	1 byte. Currently this is only used if the route type is 5 or 7 . Setting the least significant bit in the field indicates that the route carries a type 2 metric.
origin	(Used with VPNs) Identifies where the route came from.
<i>ospf-route-type</i>	1 byte, encoded as 1 or 2 for intra-area routes (depending on whether the route came from a type 1 or a type 2 LSA); 3 for summary routes; 5 for external routes (area number must be 0); 7 for NSSA routes; or 129 for sham link endpoint addresses.
route-type-vendor	Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute 0x8000 . The format is area-number:ospf-route-type:options .
rte-type	Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute 0x0306 . The format is area-number:ospf-route-type:options .
target	Defines which VPN the route participates in; target has the format 32-bit IP address:16-bit number . For example, 10.19.0.0:100.
unknown IANA	Incoming IANA codes with a value between 0x1 and 0x7fff . This code of the BGP extended community attribute is accepted, but it is not recognized.
unknown OSPF vendor community	Incoming IANA codes with a value above 0x8000 . This code of the BGP extended community attribute is accepted, but it is not recognized.

Sample Output

show route detail

```
user@host> show route detail
```

```
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
```

```

10.10.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 29
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 1:31:43
    Task: RT
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I

10.31.1.0/30 (2 entries, 1 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 2
    Next hop: via so-0/3/0.0, selected
    State: <Active Int>
    Local AS: 69
    Age: 1:30:17
    Task: IF
    Announcement bits (1): 3-Resolve tree 2
    AS path: I
  OSPF Preference: 10
    Next-hop reference count: 1
    Next hop: via so-0/3/0.0, selected
    State: <Int>
    Inactive reason: Route Preference
    Local AS: 69
    Age: 1:30:17 Metric: 1
    Area: 0.0.0.0
    Task: OSPF
    AS path: I

10.31.1.1/32 (1 entry, 1 announced)
  *Local Preference: 0
    Next hop type: Local
    Next-hop reference count: 7
    Interface: so-0/3/0.0
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:30:20
    Task: IF
    Announcement bits (1): 3-Resolve tree 2
    AS path: I

...

10.31.2.0/30 (1 entry, 1 announced)
  *OSPF Preference: 10
    Next-hop reference count: 9
    Next hop: via so-0/3/0.0
    Next hop: 10.31.1.6 via ge-3/1/0.0, selected
    State: <Active Int>
    Local AS: 69
    Age: 1:29:56 Metric: 2
    Area: 0.0.0.0
    Task: OSPF
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I

...

```

```
224.0.0.2/32 (1 entry, 1 announced)
  *PIM    Preference: 0
          Next-hop reference count: 18
          State: <Active NoReadvrt Int>
          Local AS:    69
          Age: 1:31:45
          Task: PIM Recv
          Announcement bits (2): 0-KRT 3-Resolve tree 2
          AS path: I

...

224.0.0.22/32 (1 entry, 1 announced)
  *IGMP   Preference: 0
          Next-hop reference count: 18
          State: <Active NoReadvrt Int>
          Local AS:    69
          Age: 1:31:43
          Task: IGMP
          Announcement bits (2): 0-KRT 3-Resolve tree 2
          AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.255.70.103/32 (1 entry, 1 announced)
  State: <FlashAll>
  *RSVP   Preference: 7
          Next-hop reference count: 6
          Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
          Label-switched-path green-r1-r3
          Label operation: Push 100096
          State: <Active Int>
          Local AS:    69
          Age: 1:25:49   Metric: 2
          Task: RSVP
          Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
          AS path: I

10.255.71.238/32 (1 entry, 1 announced)
  State: <FlashAll>
  *RSVP   Preference: 7
          Next-hop reference count: 6
          Next hop: via so-0/3/0.0 weight 0x1, selected
          Label-switched-path green-r1-r2
          State: <Active Int>
          Local AS:    69
          Age: 1:25:49   Metric: 1
          Task: RSVP
          Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
          AS path: I

private__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

47.0005.80ff.f800.0000.0108.0001.0102.5507.1052/152 (1 entry, 0 announced)
  *Direct Preference: 0
          Next hop type: Interface
          Next-hop reference count: 1
          Next hop: via lo0.0, selected
```

```

        State: <Active Int>
        Local AS: 69
        Age: 1:31:44
        Task: IF
        AS path: I

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
0 (1 entry, 1 announced)
    *MPLS Preference: 0
        Next hop type: Receive
        Next-hop reference count: 6
        State: <Active Int>
        Local AS: 69
        Age: 1:31:45 Metric: 1
        Task: MPLS
        Announcement bits (1): 0-KRT
        AS path: I

...

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
299776 (1 entry, 1 announced)
TSI:
KRT in-kernel 299776 /52 -> {Flood}
    *RSVP Preference: 7
        Next hop type: Flood
        Next-hop reference count: 130
        Flood nexthop branches exceed maximum
        Address: 0x8ea65d0

...

299840 (1 entry, 1 announced)
TSI:
KRT in-kernel 299840 /52 -> {indirect(1048575)}
    *RSVP Preference: 7/2
        Next hop type: Flood
        Address: 0x9174a30
        Next-hop reference count: 4
        Next hop type: Router, Next hop index: 798
        Address: 0x9174c28
        Next-hop reference count: 2
        Next hop: 8.0.0.2 via lt-1/2/0.9 weight 0x1
        Label-switched-path R2-to-R4-2p2mp
        Label operation: Pop
        Next hop type: Router, Next hop index: 1048574
        Address: 0x92544f0
        Next-hop reference count: 2
        Next hop: 7.0.0.2 via lt-1/2/0.7 weight 0x1
        Label-switched-path R2-to-R200-p2mp
        Label operation: Pop
        Next hop: 6.0.0.2 via lt-1/2/0.5 weight 0x8001
        Label operation: Pop
        State: <Active Int>
        Age: 1:29 Metric: 1
        Task: RSVP
        Announcement bits (1): 0-KRT
        AS path: I...

800010 (1 entry, 1 announced)
    *VPLS Preference: 7
        Next-hop reference count: 2

```

```
Next hop: via vt-3/2/0.32769, selected
Label operation: Pop
State: <Active Int>
Age: 1:29:30
Task: Common L2 VC
Announcement bits (1): 0-KRT
AS path: I

vt-3/2/0.32769 (1 entry, 1 announced)
  *VPLS Preference: 7
    Next-hop reference count: 2
    Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
    Label-switched-path green-r1-r3
    Label operation: Push 800012, Push 100096(top)
    Protocol next hop: 10.255.70.103
    Push 800012
    Indirect next hop: 87272e4 1048574
    State: <Active Int>
    Age: 1:29:30 Metric2: 2
    Task: Common L2 VC
    Announcement bits (2): 0-KRT 1-Common L2 VC
    AS path: I
    Communities: target:11111:1 Layer2-info: encaps:VPLS,
    control flags:, mtu: 0

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

abcd::10:255:71:52/128 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
    State: <Active Int>
    Local AS: 69
    Age: 1:31:44
    Task: IF
    AS path: I

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:31:44
    Task: IF
    AS path: I

ff02::2/128 (1 entry, 1 announced)
  *PIM Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:31:45
    Task: PIM Recv6
    Announcement bits (1): 0-KRT
    AS path: I

ff02::d/128 (1 entry, 1 announced)
  *PIM Preference: 0
```

```

        Next-hop reference count: 18
        State: <Active NoReadvrt Int>
        Local AS: 69
        Age: 1:31:45
        Task: PIM Recv6
        Announcement bits (1): 0-KRT
        AS path: I

ff02::16/128 (1 entry, 1 announced)
  *MLD Preference: 0
        Next-hop reference count: 18
        State: <Active NoReadvrt Int>
        Local AS: 69
        Age: 1:31:43
        Task: MLD
        Announcement bits (1): 0-KRT
        AS path: I

private.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
  *Direct Preference: 0
        Next hop type: Interface
        Next-hop reference count: 1
        Next hop: via lo0.16385, selected
        State: <Active NoReadvrt Int>
        Age: 1:31:44
        Task: IF
        AS path: I

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

10.255.70.103:1:3:1/96 (1 entry, 1 announced)
  *BGP Preference: 170/-101
        Route Distinguisher: 10.255.70.103:1
        Next-hop reference count: 7
        Source: 10.255.70.103
        Protocol next hop: 10.255.70.103
        Indirect next hop: 2 no-forward
        State: <Secondary Active Int Ext>
        Local AS: 69 Peer AS: 69
        Age: 1:25:49 Metric2: 1
        AIGP 210
        Task: BGP_69.10.255.70.103+179
        Announcement bits (1): 0-green-l2vpn
        AS path: I
        Communities: target:11111:1 Layer2-info: encaps:VPLS,
        control flags:, mtu: 0
        Label-base: 800008, range: 8
        Localpref: 100
        Router ID: 10.255.70.103
        Primary Routing Table bgp.l2vpn.0

10.255.71.52:1:1:1/96 (1 entry, 1 announced)
  *L2VPN Preference: 170/-1
        Next-hop reference count: 5
        Protocol next hop: 10.255.71.52
        Indirect next hop: 0 -
        State: <Active Int Ext>
        Age: 1:31:40 Metric2: 1
        Task: green-l2vpn

```

```

Announcement bits (1): 1-BGP.0.0.0+179
AS path: I
Communities: Layer2-info: encaps:VPLS, control flags:Site-Down,
mtu: 0
Label-base: 800016, range: 8, status-vector: 0x9F

10.255.71.52:1:5:1/96 (1 entry, 1 announced)
  *L2VPN Preference: 170/-101
    Next-hop reference count: 5
    Protocol next hop: 10.255.71.52
    Indirect next hop: 0 -
    State: <Active Int Ext>
    Age: 1:31:40 Metric2: 1
    Task: green-l2vpn
    Announcement bits (1): 1-BGP.0.0.0+179
    AS path: I
    Communities: Layer2-info: encaps:VPLS, control flags:, mtu: 0
    Label-base: 800008, range: 8, status-vector: 0x9F

...

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
  *L2CKT Preference: 7
    Next hop: via so-1/1/2.0 weight 1, selected
    Label-switched-path my-lsp
    Label operation: Push 100000[0]
    Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
    State: <Active Int>
    Local AS: 99
    Age: 10:21
    Task: l2 circuit
    Announcement bits (1): 0-LDP
    AS path: I
    VC Label 100000, MTU 1500, VLAN ID 512

```

show route detail (with BGP Multipath)

```

user@host> show route detail

10.1.1.8/30 (2 entries, 1 announced)
  *BGP Preference: 170/-101
    Next hop type: Router, Next hop index: 262142
    Address: 0x901a010
    Next-hop reference count: 2
    Source: 10.1.1.2
    Next hop: 10.1.1.2 via lt-0/3/0.1, selected
    Next hop: 10.1.1.6 via lt-0/3/0.5
    State: <Active Ext>
    Local AS: 1 Peer AS: 2
    Age: 5:04:43
    Task: BGP_2.10.1.1.2+59955
    Announcement bits (1): 0-KRT
    AS path: 2 I
    Accepted Multipath
    Localpref: 100
    Router ID: 1.1.1.2
  BGP Preference: 170/-101
    Next hop type: Router, Next hop index: 678
    Address: 0x8f97520
    Next-hop reference count: 9

```



```

Source: 10.1.1.6
Next hop: 10.1.1.6 via lt-0/3/0.5, selected
State: <NotBest Ext>
Inactive reason: Not Best in its group - Active preferred
Local AS: 1 Peer AS: 2
Age: 5:04:43
Task: BGP_2.10.1.1.6+58198
AS path: 2 I
Accepted MultipathContrib
Localpref: 100
Router ID: 1.1.1.3

```

show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```

user@host> show route label 299872 detail
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
299872 (1 entry, 1 announced)
    *LDP    Preference: 9
            Next hop type: Flood
            Next-hop reference count: 3
            Address: 0x9097d90
            Next hop: via vt-0/1/0.1
            Next-hop index: 661
            Label operation: Pop
            Address: 0x9172130
            Next hop: via so-0/0/3.0
            Next-hop index: 654
            Label operation: Swap 299872
            State: **Active Int>
            Local AS: 1001
            Age: 8:20      Metric: 1
            Task: LDP
            Announcement bits (1): 0-KRT
            AS path: I
            FECs bound to route: P2MP root-addr 10.255.72.166, grp 232.1.1.1,
src 192.168.142.2

```

show route label detail (Multipoint LDP with Multicast-Only Fast Reroute)

```

user@host> show route label 301568 detail
mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
301568 (1 entry, 1 announced)
    *LDP    Preference: 9
            Next hop type: Flood
            Address: 0x2735208
            Next-hop reference count: 3
            Next hop type: Router, Next hop index: 1397
            Address: 0x2735d2c
            Next-hop reference count: 3
            Next hop: 1.3.8.2 via ge-1/2/22.0
            Label operation: Pop
            Load balance label: None;
            Next hop type: Router, Next hop index: 1395
            Address: 0x2736290
            Next-hop reference count: 3
            Next hop: 1.3.4.2 via ge-1/2/18.0
            Label operation: Pop
            Load balance label: None;
            State: <Active Int AckRequest MulticastRPF>
            Local AS: 10

```

```
Age: 54:05      Metric: 1
Validation State: unverified
Task: LDP
Announcement bits (1): 0-KRT
AS path: I
FECs bound to route: P2MP root-addr 1.1.1.1, grp: 232.1.1.1, src:
192.168.219.11
Primary Upstream : 1.1.1.3:0--1.1.1.2:0
  RPF Nexthops :
    ge-1/2/15.0, 1.2.94.1, Label: 301568, weight: 0x1
    ge-1/2/14.0, 1.2.3.1, Label: 301568, weight: 0x1
Backup Upstream : 1.1.1.3:0--1.1.1.6:0
  RPF Nexthops :
    ge-1/2/20.0, 1.2.96.1, Label: 301584, weight: 0xffffe
    ge-1/2/19.0, 1.3.6.1, Label: 301584, weight: 0xffffe
```

IS-IS Feature Guide for QFX10000 Switches

Overview

- [IS-IS Overview on page 4157](#)

IS-IS Overview

The IS-IS protocol is an interior gateway protocol (IGP) that uses link-state information to make routing decisions.

IS-IS is a link-state IGP that uses the shortest-path-first (SPF) algorithm to determine routes. IS-IS evaluates the topology changes and determines whether to perform a full SPF recalculation or a partial route calculation (PRC). This protocol originally was developed for routing International Organization for Standardization (ISO) Connectionless Network Protocol (CLNP) packets.

Like OSPF routing, IS-IS uses hello packets that allow network convergence to occur quickly when network changes are detected. IS-IS uses the SPF algorithm to determine routes. Using SPF, IS-IS evaluates network topology changes and determines if a full or partial route calculation is required.



NOTE: Because IS-IS uses ISO addresses, the configuration of IP version 6 (IPv6) and IP version 4 (IPv4) implementations of IS-IS is identical.



NOTE: See *Platforms/FPCs That Cannot Forward TCC Encapsulated ISO Traffic* to find a list of those devices and FPC configurations that cannot pass ISO traffic when encapsulated in TCC format.

This section discusses the following topics:

- [IS-IS Terminology on page 4158](#)
- [ISO Network Addresses on page 4158](#)
- [IS-IS Packets on page 4160](#)
- [Persistent Route Reachability on page 4161](#)

- [IS-IS Support for Multipoint Network Clouds on page 4161](#)
- [Installing a Default Route to the Nearest Routing Device That Operates at Both IS-IS Levels on page 4161](#)

IS-IS Terminology

An IS-IS network is a single autonomous system (AS), also called a *routing domain*, that consists of *end systems* and *intermediate systems*. End systems are network entities that send and receive packets. Intermediate systems send and receive packets and relay (forward) packets. (Intermediate system is the Open System Interconnection [OSI] term for a router.) ISO packets are called network PDUs.

In IS-IS, a single AS can be divided into smaller groups called *areas*. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring *Level 1* and *Level 2* intermediate systems. Level 1 systems route within an area; when the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs. No IS-IS area functions strictly as a backbone.

Level 1 routers share intra-area routing information, and Level 2 routers share interarea information about IP addresses available within each area. Uniquely, IS-IS routers can act as both Level 1 and Level 2 routers, sharing intra-area routes with other Level 1 routers and interarea routes with other Level 2 routers.

The propagation of link-state updates is determined by the level boundaries. All routers within a level maintain a complete link-state database of all other routers in the same level. Each router then uses the Dijkstra algorithm to determine the shortest path from the local router to other routers in the link-state database.

ISO Network Addresses

IS-IS uses ISO network addresses. Each address identifies a point of connection to the network, such as a router interface, and is called a *network service access point (NSAP)*.

IS-IS supports multiple NSAP addresses on the loopback lo0 interface.

An end system can have multiple NSAP addresses, in which case the addresses differ only by the last byte (called the *n-selector*). Each NSAP represents a service that is available at that node. In addition to having multiple services, a single node can belong to multiple areas.

Each network entity also has a special network address called a *network entity title (NET)*. Structurally, an NET is identical to an NSAP address but has an n-selector of 00. Most end systems and intermediate systems have one NET. Intermediate systems that participate in multiple areas can have multiple NETs.

The following ISO addresses illustrate the IS-IS address format:

```
49.0001.00a0.c96b.c490.00
49.0001.2081.9716.9018.00
```

NETs take several forms, depending on your network requirements. NET addresses are hexadecimal and range from 8 octets to 20 octets in length. Generally, the format consists

of an authority and format Identifier (AFI), a domain ID, an area ID, a system identifier, and a selector. The simplest format omits the domain ID and is 10 octets long. For example, the NET address 49.0001.1921.6800.1001.00 consists of the following parts:

- 49—AFI
- 0001—Area ID
- 1921.6800.1001—System identifier
- 00—Selector

The system identifier must be unique within the network. For an IP-only network, we recommend using the IP address of an interface on the router. Configuring a loopback NET address with the IP address is helpful when troubleshooting is required on the network.

The first portion of the address is the area number, which is a variable number from 1 through 13 bytes. The first byte of the area number (49) is the authority and format indicator (AFI). The next bytes are the assigned domain (area) identifier, which can be from 0 through 12 bytes. In the examples above, the area identifier is 0001.

The next six bytes form the system identifier. The system identifier can be any six bytes that are unique throughout the entire domain. The system identifier commonly is the media access control (MAC) address (as in the first example, 00a0.c96b.c490) or the IP address expressed in binary-coded decimal (BCD) (as in the second example, 2081.9716.9018, which corresponds to IP address 208.197.169.18). The last byte (00) is the n-selector.



NOTE: The system identifier cannot be 0000.0000.0000. All 0s is an illegal setting, and the adjacency is not formed with this setting.

To provide help with IS-IS debugging, the Junos[®] operating system (Junos OS) supports dynamic mapping of ISO system identifiers to the hostname. Each system can be configured with a hostname, which allows the system identifier-to-hostname mapping to be carried in a dynamic hostname type, length, and value (TLV) tuple in IS-IS link-state PDUs. This enables intermediate systems in the routing domain to learn about the ISO system identifier of a particular intermediate system.

IS-IS Packets

Each IS-IS PDU shares a common header. IS-IS uses the following PDUs to exchange protocol information:

- IS-IS hello (IIH) PDUs—Broadcast to discover the identity of neighboring IS-IS systems and to determine whether the neighbors are Level 1 or Level 2 intermediate systems.

IS-IS hello PDUs establish adjacencies with other routers and have three different formats: one for point-to-point hello packets, one for Level 1 broadcast links, and one for Level 2 broadcast links. Level 1 routers must share the same area address to form an adjacency, while Level 2 routers do not have this limitation. The request for adjacency is encoded in the Circuit type field of the PDU.

Hello PDUs have a preset length assigned to them. The IS-IS router does not resize any PDU to match the maximum transmission unit (MTU) on a router interface. Each interface supports the maximum IS-IS PDU of 1492 bytes, and hello PDUs are padded to meet the maximum value. When the hello is sent to a neighboring router, the connecting interface supports the maximum PDU size.

- Link-state PDUs—Contain information about the state of adjacencies to neighboring IS-IS systems. Link-state PDUs are flooded periodically throughout an area.

Also included is metric and IS-IS neighbor information. Each link-state PDU must be refreshed periodically on the network and is acknowledged by information within a sequence number PDU.

On point-to-point links, each link-state PDU is acknowledged by a partial sequence number PDU (PSNP), but on broadcast links, a complete sequence number PDU (CSNP) is sent out over the network. Any router that finds newer link-state PDU information in the CSNP then purges the out-of-date entry and updates the link-state database.

Link-state PDUs support variable-length subnet mask addressing.

- Complete sequence number PDUs (CSNPs)—Contain a complete list of all link-state PDUs in the IS-IS database. CSNPs are sent periodically on all links, and the receiving systems use the information in the CSNP to update and synchronize their link-state PDU databases. The designated router multicasts CSNPs on broadcast links in place of sending explicit acknowledgments for each link-state PDU.

Contained within the CSNP is a link-state PDU identifier, a lifetime, a sequence number, and a checksum for each entry in the database. Periodically, a CSNP is sent on both broadcast and point-to-point links to maintain a correct database. Also, the advertisement of CSNPs occurs when an adjacency is formed with another router. Like IS-IS hello PDUs, CSNPs come in two types: Level 1 and Level 2.

When a device receives a CSNP, it checks the database entries against its own local link-state database. If it detects missing information, the device requests specific link-state PDU details using a partial sequence number PDU (PSNP).

- Partial sequence number PDUs (PSNPs)—Sent multicast by a receiver when it detects that it is missing a link-state PDU (when its link-state PDU database is out of date). The receiver sends a PSNP to the system that transmitted the CSNP, effectively

requesting that the missing link-state PDU be transmitted. That routing device, in turn, forwards the missing link-state PDU to the requesting routing device.

A PSNP is used by an IS-IS router to request link-state PDU information from a neighboring router. A PSNP can also explicitly acknowledge the receipt of a link-state PDU on a point-to-point link. On a broadcast link, a CSNP is used as implicit knowledge. Like hello PDUs and CSNPs, the PSNP also has two types: Level 1 and Level 2.

When a device compares a CSNP to its local database and determines that a link-state PDU is missing, the router issues a PSNP for the missing link-state PDU, which is returned in a link-state PDU from the router sending the CSNP. The received link-state PDU is then stored in the local database, and an acknowledgment is sent back to the originating router.

Persistent Route Reachability

IPv4 and IPv6 route reachability information in IS-IS link-state PDUs is preserved when you commit a configuration. IP prefixes are preserved with their original packet fragment upon link-state PDU regeneration.


IS-IS Support for Multipoint Network Clouds

IS-IS does not support multipoint configurations. Therefore, when configuring Frame Relay or Asynchronous Transfer Mode (ATM) networks, you must configure them as collections of point-to-point links, not as multipoint clouds.

Installing a Default Route to the Nearest Routing Device That Operates at Both IS-IS Levels

When a routing device that operates as both a Level 1 and Level 2 router (Router B) determines that it can reach at least one area other than its own (for example, in Area Y), it sets the ATTACHED bit in its Level 1 link-state PDU. Thereafter, the Level 1 router (Router A) introduces a default route pointing to the nearest attached routing device that operates as both a Level 1 and Level 2 router (Router B). See [Figure 94](#).

Figure 94: Install Default Route to Nearest Routing Device That Operates at Both Level 1 and Level 2



Related Documentation

- *IS-IS Feature Guide for Routing Devices*

PART 65

Configuring IS-IS Features

- [Configuring IS-IS on page 4165](#)
- [Configuring BFD Authentication for IS-IS on page 4243](#)
- [Configuring Hitless Authentication Key Rollover for IS-IS on page 4255](#)

Configuring IS-IS

- [Understanding Loop-Free Alternate Routes for IS-IS on page 4165](#)
- [Example: Configuring IS-IS on page 4169](#)
- [Example: Configuring a Multi-Level IS-IS Topology to Control Interarea Flooding on page 4175](#)
- [Configuring IS-IS Authentication on page 4183](#)
- [Configuring IS-IS Authentication Without Network-Wide Deployment on page 4184](#)
- [Example: Redistributing OSPF Routes into IS-IS on page 4185](#)
- [Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies on page 4193](#)
- [Example: Configuring IS-IS Multicast Topology on page 4201](#)
- [Example: Configuring Link and Node Protection for IS-IS Routes on page 4217](#)
- [Example: Configuring an IS-IS Default Route Policy on Logical Systems on page 4232](#)
- [Example: Configuring IS-IS for CLNS on page 4237](#)
- [Example: Enabling Packet Checksums on IS-IS Interfaces for Error Checking on page 4239](#)

Understanding Loop-Free Alternate Routes for IS-IS

In Junos OS Release 9.5 and later, support for IS-IS loop-free alternate routes enables IP fast-reroute capability for IS-IS. Junos OS precomputes loop-free backup routes for all IS-IS routes. These backup routes are preinstalled in the Packet Forwarding Engine, which performs a local repair and implements the backup path when the link for a primary next hop for a particular route is no longer available. With local repair, the Packet Forwarding Engine can correct a path failure before it receives recomputed paths from the Routing Engine. Local repair reduces the amount of time needed to reroute traffic to less than 50 milliseconds. In contrast, global repair can take up to 800 milliseconds to compute a new route. Local repair and global repair are thus complementary. Local repair enables traffic to continue to be routed using a backup path until global repair is able to calculate a new route.

A loop-free path is one that does not forward traffic back through the routing device to reach a given destination. That is, a neighbor whose shortest path to the destination traverses the routing device is not used as a backup route to that destination. To determine loop-free alternate paths for IS-IS routes, Junos OS runs shortest-path-first (SPF) calculations on each one-hop neighbor. You can enable support for alternate loop-free

routes on any IS-IS interface. Because it is common practice to enable LDP on an interface for which IS-IS is already enabled, this feature also provides support for LDP label-switched paths (LSPs).



NOTE: If you enable support for alternate loop-free routes on an interface configured for both LDP and IS-IS, you can use the `traceroute` command to trace the active path to the primary next hop.


The level of backup coverage available through IS-IS routes depends on the actual network topology and is typically less than 100 percent for all destinations on any given routing device. You can extend backup coverage to include RSVP LSPs.

Junos OS provides two mechanisms for route redundancy for IS-IS through alternate loop-free routes: link protection and node-link protection. When you enable link protection or node-link protection on an IS-IS interface, Junos OS creates a single alternate path to the primary next hop for all destination routes that traverse a protected interface. Link protection offers per-link traffic protection. Use link protection when you assume that only a single link might become unavailable but that the neighboring node on the primary path would still be available through another interface.

Node-link protection establishes an alternate path through a different routing device altogether. Use node-link protection when you assume that access to a node is lost when a link is no longer available. As a result, Junos OS calculates a backup path that avoids the primary next-hop routing device. In Junos OS Release 9.4 and earlier, only the RSVP protocol supports Packet Forwarding Engine local repair and fast reroute as well as link protection and node protection.

In [Figure 95](#), Case 2 shows how link protection allows source Router A to switch to Link B when the primary next hop Link A to destination Router C fails. However, if Router B fails, Link B also fails, and the protected Link A is lost. If node-link protection is enabled, Router A is able to switch to Link D on Router D and bypass the failed Router B altogether. As shown in Case 1, with node-link protection enabled, Router A has a node-link protection alternate path available through Router D to destination Router C. That means that if Router B fails, Router A can still reach Router C because the path from Router A to Link D remains available as an alternate backup path.

Figure 95: Link Protection and Node-Link Protection Comparison for IS-IS Routes



The Junos OS implementation of support for loop-free alternate paths for IS-IS routes is based on the following standards:

- RFC 5286, *Basic Specification for IP Fast-Reroute: Loop-free Alternates*
- RFC 5714, *IP Fast Reroute Framework*

Configuring Link Protection for IS-IS

You can configure link protection on any interface for which IS-IS is enabled. When you enable link protection, Junos OS creates one alternate path to the primary next hop for all destination routes that traverse a protected interface. Link protection assumes that only a single link becomes unavailable but that the neighboring node would still be available through another interface.



NOTE: You must also configure a per-packet load-balancing routing policy to ensure that the routing protocol process installs all the next hops for a given route in the routing table.

To enable link protection, include the **link-protection** statement at the **[edit protocols isis interface *interface-name*]** hierarchy level:

```
[edit]
protocols {
  isis {
    interface interface-name {
```

```

        link-protection;
    }
}

```

Configuring Node-Link Protection for IS-IS

You can configure node-link protection on any interface for which IS-IS is enabled. Node-link protection establishes an alternate path through a different routing device altogether for all destination routes that traverse a protected interface. Node-link protection assumes that the entire routing device, or node, has failed. Junos OS therefore calculates a backup path that avoids the primary next-hop routing device.



NOTE: You must also configure a per-packet load-balancing routing policy to ensure that the routing protocol process installs all the next hops for a given route in the routing table.

To enable node-link protection, include the **node-link-protection** statement at the **[edit protocols isis interface *interface-name*]** hierarchy level:

```

[edit]
protocols {
  isis {
    interface interface-name {
      node-link-protection;
    }
  }
}

```

Excluding an IS-IS Interface as a Backup for Protected Interfaces

By default, all IS-IS interfaces that belong to the master instance or a specific routing instance are eligible as backup interfaces for protected interfaces. You can specify that any IS-IS interface be excluded from functioning as a backup interface to protected interfaces. To exclude an IS-IS interface as a backup interface, include the **no-eligible-backup** statement at the **[edit protocols isis interface *interface-name*]** hierarchy level:

```

[edit]
protocols {
  isis {
    interface interface-name {
      no-eligible-backup;
    }
  }
}

```

Configuring RSVP Label-Switched Paths as Backup Paths for IS-IS

Relying on the shortest-path-first (SPF) calculation of backup paths for one-hop neighbors might result in less than 100 percent backup coverage for a specific network topology. You can enhance coverage of IS-IS and LDP label-switched paths (LSPs) by

configuring RSVP LSPs as backup paths. To configure a specific RSVP LSP as a backup path, include the **backup** statement at the **[edit protocols mpls label-switched-path *lsp-name*]** hierarchy level:

```
[edit]
protocols {
  mpls {
    label-switched-path lsp-name {
      backup;
      to ip-address;
    }
  }
}
```

When configuring an LSP, you must specify the IP address of the egress routing device with the **to** statement. For detailed information about configuring LSPs and RSVP, see the *Junos OS MPLS Applications Library for Routing Devices*.

Using Operational Mode Commands to Monitor Protected IS-IS Routes

You can issue operational mode commands that provide more details about your link-protected and node-link-protected IS-IS routes. The following guidelines explain the type of information available from the output of each command:

- **show isis backup label-switched-path**—Displays which MPLS LSPs have been designated as backup paths and the current status of those LSPs.
- **show isis backup spf results**—Displays SPF calculations for each neighbor for a given destination. Indicates whether a specific interface or node has been designated as a backup path and why. Use the **no-coverage** option to display only those nodes that do not have backup coverage.
- **show isis backup coverage**—Displays the percentage of nodes and prefixes for each type of address family that is protected.
- **show isis interface detail**—Displays the type of protection (link or node-link) applied to each protected interface.

Related Documentation

- [Example: Configuring Node-Link Protection for IS-IS Routes in a Layer 3 VPN on page 4221](#)

Example: Configuring IS-IS

This example shows how to configure IS-IS.

- [Requirements on page 4170](#)
- [Overview on page 4170](#)
- [Configuration on page 4170](#)
- [Verification on page 4172](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.


Overview

In this example, you configure the two IS-IS routing devices in a single area. The devices have NET addresses 49.0002.0192.0168.0001.00 and 49.0002.0192.0168.0002.00 on the lo0 interfaces. Additionally, you configure the ISO family on the IS-IS interfaces.

For Junos OS security devices only, you configure the **mode packet-based** statement at the **[edit security forwarding-options family iso]** hierarchy level.

Figure 96 shows the topology used in this example.

Figure 96: Simple IS-IS Topology



“CLI Quick Configuration” on page 4170 shows the configuration for both of the devices in Figure 96. The section “Step-by-Step Procedure” on page 4171 describes the steps on Device R1.

Configuration

CLI Quick Configuration	To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.
Device R1	<pre> set security forwarding-options family iso mode packet-based set interfaces ge-1/2/0 unit 0 description to-R2 set interfaces ge-1/2/0 unit 0 family inet address 10.0.0.1/30 set interfaces ge-1/2/0 unit 0 family iso set interfaces lo0 unit 0 family inet address 192.168.0.1/32 set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0001.00 set protocols isis interface ge-1/2/0.0 set protocols isis interface lo0.0 </pre>
Device R2	<pre> set security forwarding-options family iso mode packet-based set interfaces ge-1/2/0 unit 0 description to-R1 set interfaces ge-1/2/0 unit 0 family inet address 10.0.0.2/30 set interfaces ge-1/2/0 unit 0 family iso set interfaces lo0 unit 0 family inet address 192.168.0.2/32 set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0002.00 set protocols isis interface ge-1/2/0.0 set protocols isis interface lo0.0 </pre>

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IS-IS:

1. Enable IS-IS if your router is in secure context.

```
[edit security forwarding-options family iso]
user@R1# set mode packet-based
```
2. Create the interface that connects to Device R2, and configure the ISO family on the interface.

```
[edit interfaces ge-1/2/0 unit 0]
user@R1# set description to-R2
user@R1# set family inet address 10.0.0.1/30
user@R1# set family iso
```
3. Create the loopback interface, set the IP address, and set the NET address.

```
[edit interfaces lo0 unit 0]
user@R1# set family inet address 192.168.0.1/32
user@R1# set family iso address 49.0002.0192.0168.0001.00
```
4. Enable IS-IS on the interfaces.

```
[edit protocols isis]
user@R1# set interface ge-1/2/0.0
user@R1# set interface lo0.0
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show security** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show security
forwarding-options {
  family iso {
    mode packet-based;
  }
}

user@R1# show interfaces
ge-1/2/0 {
  unit 0 {
    description to-R2;
    family inet {
      address 10.0.0.1/30;
    }
    family iso;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.168.0.1/32;
    }
    family iso {
```

```

        address 49.0002.0192.0168.0001.00;
    }
}
}

user@R1# show protocols
isis {
    interface ge-1/2/0.0;
    interface lo0.0;
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying IS-IS Interface Configuration on page 4172](#)
- [Verifying IS-IS Interface Configuration in Detail on page 4172](#)
- [Verifying IS-IS Adjacencies on page 4173](#)
- [Verifying IS-IS Adjacencies in Detail on page 4173](#)

Verifying IS-IS Interface Configuration

Purpose Verify the status of the IS-IS-enabled interfaces.

Action From operational mode, enter the **show isis interface brief** command.

```

user@R1> show isis interface brief
IS-IS interface database:
Interface          L CirID Level 1 DR          Level 2 DR          L1/L2 Metric
lo0.0               3  0x1 Passive                Passive                0/0
ge-1/2/0.0          3  0x1 R2.02                  R2.02                  10/10

```

Meaning Verify that the output shows the intended configuration of the interfaces on which IS-IS is enabled.

Verifying IS-IS Interface Configuration in Detail

Purpose Verify the details of IS-IS-enabled interfaces.

Action From operational mode, enter the **show isis interface detail** command.

```

user@R1> show isis interface detail
IS-IS interface database:
lo0.0
  Index: 75, State: 0x6, Circuit id: 0x1, Circuit type: 0
  LSP interval: 100 ms, CSNP interval: disabled
  Adjacency advertisement: Advertise
  Level Adjacencies Priority Metric Hello (s) Hold (s) Designated Router
    1           0       64      0 Passive
    2           0       64      0 Passive
ge-1/2/0.0
  Index: 77, State: 0x6, Circuit id: 0x1, Circuit type: 3
  LSP interval: 100 ms, CSNP interval: 10 s

```

```

Adjacency advertisement: Advertise
Level Adjacencies Priority Metric Hello (s) Hold (s) Designated Router
1      1            64      10      9.000      27 R2.02 (not us)
2      1            64      10      9.000      27 R2.02 (not us)

```

Meaning Check the following output fields and verify that the output shows the intended configuration of IS-IS-enabled interfaces:

- Interface—Interface configured for IS-IS.
- State—Internal implementation information.
- Circuit id—Circuit identifier.
- Circuit type—Configured level of IS-IS:
 - 1—Level 1 only
 - 2—Level 2 only
 - 3—Level 1 and Level 2
- link-state PDU interval—Time between IS-IS information messages.
- L or Level—Type of adjacency:
 - 1—Level 1 only
 - 2—Level 2 only
 - 3—Level 1 and Level 2
- Adjacencies—Adjacencies established on the interface.
- Priority—Priority value established on the interface.
- Metric—Metric value for the interface.
- Hello(s)—Intervals between hello PDUs.
- Hold(s)—Hold time on the interface.

Verifying IS-IS Adjacencies

Purpose Display brief information about IS-IS neighbors.

Action From operational mode, enter the **show isis adjacency brief** command.

```

user@R1> show isis adjacency brief
Interface      System      L State      Hold (secs) SNPA
ge-1/2/0.0     R2          1 Up          6  0:5:85:8f:c8:bd
ge-1/2/0.0     R2          2 Up          6  0:5:85:8f:c8:bd

```

Meaning Verify the adjacent routers in the IS-IS database.

Verifying IS-IS Adjacencies in Detail

Purpose Display extensive information about IS-IS neighbors.

Action From operational mode, enter the **show isis adjacency extensive** command.

```
user@R1> show isis adjacency extensive
```

```
R2
```

```
Interface: ge-1/2/0.0, Level: 1, State: Up, Expires in 6 secs
Priority: 64, Up/Down transitions: 1, Last transition: 00:40:28 ago
Circuit type: 3, Speaks: IP, IPv6, MAC address: 0:5:85:8f:c8:bd
Topologies: Unicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R2.02, IP addresses: 10.0.0.2
Transition log:
When                State          Event          Down reason
Thu May 31 11:18:48  Up            Seenself
```

```
R2
```

```
Interface: ge-1/2/0.0, Level: 2, State: Up, Expires in 8 secs
Priority: 64, Up/Down transitions: 1, Last transition: 00:40:28 ago
Circuit type: 3, Speaks: IP, IPv6, MAC address: 0:5:85:8f:c8:bd
Topologies: Unicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R2.02, IP addresses: 10.0.0.2
Transition log:
When                State          Event          Down reason
Thu May 31 11:18:48  Up            Seenself
```

Meaning Check the following fields and verify the adjacency information about IS-IS neighbors:

- Interface—Interface through which the neighbor is reachable.
- L or Level—Configured level of IS-IS:
 - 1—Level 1 only
 - 2—Level 2 only
 - 3—Level 1 and Level 2

An exclamation point before the level number indicates that the adjacency is missing an IP address.

- State—Status of the adjacency: **Up**, **Down**, **New**, **One-way**, **Initializing**, or **Rejected**.
- Event—Message that identifies the cause of a state.
- Down reason—Reason the adjacency is down.
- Restart capable—A neighbor is configured for graceful restart.
- Transition log—List of transitions including **When**, **State**, and **Reason**.

**Related
Documentation**

- *Understanding IS-IS Configuration*
- *Example: Configuring IS-IS for GRES with Graceful Restart*
- *Configuring Designated Router Election Priority for IS-IS*

Example: Configuring a Multi-Level IS-IS Topology to Control Interarea Flooding

This example shows how to configure a multi-level IS-IS topology.

- [Requirements on page 4175](#)
- [Overview on page 4175](#)
- [Configuration on page 4176](#)
- [Verification on page 4180](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview


Like OSPF, the IS-IS protocol supports the partitioning of a routing domain into multiple areas with levels that control interarea flooding. The use of multiple levels improves protocol scalability, as Level 2 (backbone) link-state PDUs are normally not flooded into a Level 1 area.

An IS-IS Level 2 area is analogous to the OSPF backbone area (0), while a Level 1 area operates much like an OSPF totally stubby area, in that a default route is normally used to reach both inter-level and AS external routes.

Unlike OSPF, IS-IS area boundaries occur between routers, such that a given routing device is always wholly contained within a particular area. Level 1 adjacencies can be formed between routers that share a common area number, while a Level 2 adjacency can be formed between routers that might or might not share an area number.

[Figure 97](#) shows the topology used in this example.

Figure 97: IS-IS Multi-Level Topology



[“CLI Quick Configuration” on page 4176](#) shows the configuration for all of the devices in [Figure 97](#). The section [“Step-by-Step Procedure” on page 4177](#) describes the steps on Device R5.

This example has the following characteristics:

- Device R5 functions as a Level 1/Level 2 router to interconnect the Level 2 backbone area 49.001 and the Level 1 area 49.002 containing Device R6 and Device R7.
- The system ID is based on the devices' IPv4 lo0 addresses.
- Loss of any individual interface does not totally disrupt the IS-IS operation.
- The IPv4 lo0 addresses of all routers are reachable through IS-IS.
- The link between Device R3 and Device S1 appears in area 49.001 as an intra-area route. No IS-IS adjacencies can be established on this interface. This is accomplished by configuring the [passive](#) statement on Device R3's interface to Device S1.
- The loopback addresses of Level 2 devices do not appear in a Level 1 area.
- There is only one adjacency for each device pairing.

Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device R3

```
set interfaces fe-1/2/0 unit 0 description to-R4
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.17/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/1 unit 0 description to-R5
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.21/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces fe-1/2/2 unit 0 family inet address 10.0.0.41/30
set interfaces fe-1/2/2 unit 0 description to-S1
set interfaces lo0 unit 0 family inet address 192.168.0.3/32
set interfaces lo0 unit 0 family iso address 49.001.0192.0168.0003.00
set protocols isis interface fe-1/2/0.0 level 1 disable
set protocols isis interface fe-1/2/1.0 level 1 disable
set protocols isis interface lo0.0 level 1 disable
set protocols isis interface fe-1/2/2.0 passive
```

Device R4

```
set interfaces fe-1/2/0 unit 0 description to-R3
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.18/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/1 unit 0 description to-R5
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.25/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces lo0 unit 0 family inet address 192.168.0.4/32
set interfaces lo0 unit 0 family iso address 49.001.0192.0168.0004.00
set protocols isis interface fe-1/2/0.0 level 1 disable
set protocols isis interface fe-1/2/1.0 level 1 disable
set protocols isis interface lo0.0 level 1 disable
```


Device R5	<pre> set interfaces fe-1/2/0 unit 0 description to-R3 set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.22/30 set interfaces fe-1/2/0 unit 0 family iso set interfaces fe-1/2/1 unit 0 description to-R4 set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.26/30 set interfaces fe-1/2/1 unit 0 family iso set interfaces fe-1/2/2 unit 0 description to-R6 set interfaces fe-1/2/2 unit 0 family inet address 10.0.0.29/30 set interfaces fe-1/2/2 unit 0 family iso set interfaces fe-1/2/3 unit 0 description to-R7 set interfaces fe-1/2/3 unit 0 family inet address 10.0.0.38/30 set interfaces fe-1/2/3 unit 0 family iso set interfaces lo0 unit 0 family inet address 192.168.0.5/32 set interfaces lo0 unit 0 family iso address 49.002.0192.0168.0005.00 set protocols isis interface fe-1/2/0.0 level 1 disable set protocols isis interface fe-1/2/1.0 level 1 disable set protocols isis interface fe-1/2/2.0 level 2 disable set protocols isis interface fe-1/2/3.0 level 2 disable set protocols isis interface lo0.0 level 1 disable </pre>
Device R6	<pre> set interfaces fe-1/2/0 unit 0 description to-R5 set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.30/30 set interfaces fe-1/2/0 unit 0 family iso set interfaces fe-1/2/1 unit 0 description to-R7 set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.33/30 set interfaces fe-1/2/1 unit 0 family iso set interfaces lo0 unit 0 family inet address 192.168.0.6/32 set interfaces lo0 unit 0 family iso address 49.002.0192.0168.0006.00 set protocols isis interface fe-1/2/0.0 level 2 disable set protocols isis interface fe-1/2/1.0 level 2 disable set protocols isis interface lo0.0 level 2 disable </pre>
Device R7	<pre> set interfaces fe-1/2/0 unit 0 description to-R6 set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.34/30 set interfaces fe-1/2/0 unit 0 family iso set interfaces fe-1/2/1 unit 0 description to-R5 set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.37/30 set interfaces fe-1/2/1 unit 0 family iso set interfaces lo0 unit 0 family inet address 192.168.0.7/32 set interfaces lo0 unit 0 family iso address 49.002.0192.0168.0007.00 set protocols isis interface fe-1/2/0.0 level 2 disable set protocols isis interface fe-1/2/1.0 level 2 disable set protocols isis interface lo0.0 level 2 disable </pre>
Device S1	<pre> set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.42/30 set interfaces fe-1/2/0 unit 0 description to-R3 </pre>
Step-by-Step Procedure	<p>The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see <i>Using the CLI Editor in Configuration Mode</i> in the <i>CLI User Guide</i>.</p> <p>To configure multi-level IS-IS:</p> <ol style="list-style-type: none"> 1. Configure the network interfaces.

Enable IS-IS on the interfaces by including the ISO address family on each interface.

```
[edit interfaces]
user@R5# set fe-1/2/0 unit 0 description to-R3
user@R5# set fe-1/2/0 unit 0 family inet address 10.0.0.22/30
user@R5# set fe-1/2/0 unit 0 family iso
user@R5# set fe-1/2/1 unit 0 description to-R4
user@R5# set fe-1/2/1 unit 0 family inet address 10.0.0.26/30
user@R5# set fe-1/2/1 unit 0 family iso
user@R5# set fe-1/2/2 unit 0 description to-R6
user@R5# set fe-1/2/2 unit 0 family inet address 10.0.0.29/30
user@R5# set fe-1/2/2 unit 0 family iso
user@R5# set fe-1/2/3 unit 0 description to-R7
user@R5# set fe-1/2/3 unit 0 family inet address 10.0.0.38/30
user@R5# set fe-1/2/3 unit 0 family iso
```

2. Configure two loopback interface addresses.

One address is for IPv4.

The other is for the IS-IS area 49.002 so that Device R5 can form adjacencies with the other Level 1 devices in area 49.002. Even though Device R5's NET identifies itself as belonging to the Level 1 area 49.002, its loopback interface is not configured as a Level 1 interface. Doing so would cause the route to Device R5's loopback to be injected into the Level 1 area.

```
[edit interfaces lo0 unit 0]
user@R5# set family inet address 192.168.0.5/32
user@R5# set family iso address 49.002.0192.0168.0005.00
```

3. Specify the IS-IS level on a per-interface basis.

Device R5 becomes adjacent to the other routing devices on the same level on each link.

By default, IS-IS is enabled for IS-IS areas on all interfaces on which the ISO protocol family is enabled (at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level). To disable IS-IS at any particular level on an interface, include the **disable** statement.

Device R5's loopback interface is configured to run Level 2 only. If Level 1 operation were enabled on lo0.0, Device R5 would include its loopback address in its Level 1 link-state PDU, which is incorrect for this example in which the loopback addresses of Level 2 devices must not appear in a Level 1 area.

Unlike OSPF, you must explicitly list the router's lo0 interface at the **[edit protocols isis]** hierarchy level, because this interface is the source of the router's NET, and therefore must be configured as an IS-IS interface. In IS-IS, the lo0 interface operates in the passive mode by default, which is ideal because adjacency formation can never occur on a virtual interface.

```
[edit protocols isis]
user@R5# set interface fe-1/2/0.0 level 1 disable
user@R5# set interface fe-1/2/1.0 level 1 disable
user@R5# set interface fe-1/2/0.0 level 2 disable
user@R5# set interface fe-1/2/3.0 level 2 disable
user@R5# set interface lo0.0 level 1 disable
```

Results From configuration mode, confirm your configuration by entering the **show interfaces** and **show protocols** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@R5# show interfaces
fe-1/2/0 {
  unit 0 {
    description to-R3;
    family inet {
      address 10.0.0.22/30;
    }
    family iso;
  }
}
fe-1/2/1 {
  unit 0 {
    description to-R4;
    family inet {
      address 10.0.0.26/30;
    }
    family iso;
  }
}
fe-1/2/2 {
  unit 0 {
    description to-R6;
    family inet {
      address 10.0.0.29/30;
    }
    family iso;
  }
}
fe-1/2/3 {
  unit 0 {
    description to-R7;
    family inet {
      address 10.0.0.38/30;
    }
    family iso;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.168.0.5/32;
    }
    family iso {
      address 49.002.0192.0168.0005.00;
    }
  }
}
}

user@R5# show protocols
isis {
  interface fe-1/2/0.0 {
    level 1 disable;
  }
}

```

```
}
interface fe-1/2/1.0 {
  level 1 disable;
}
interface fe-1/2/0.0 {
  level 2 disable;
}
interface fe-1/2/3.0 {
  level 2 disable;
}
interface lo0.0 {
  level 1 disable;
}
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Checking Interface-to-Area Associations on page 4180](#)
- [Verifying IS-IS Adjacencies on page 4180](#)
- [Examining the IS-IS Database on page 4181](#)

Checking Interface-to-Area Associations

Purpose Make sure that the interface-to-area associations are configured as expected.

Action From operational mode, enter the **show isis interface** command.

```
user@R5> show isis interface
IS-IS interface database:
Interface          L CirID Level 1 DR      Level 2 DR      L1/L2 Metric
lo0.0              3  0x1 Disabled             Passive         0/0
fe-1/2/0.0         2  0x3 Disabled          R5.03           10/10
fe-1/2/1.0         2  0x2 Disabled          R5.02           10/10
fe-1/2/0.0         1  0x1 R6.02             Disabled        10/10
fe-1/2/3.0         1  0x4 R5.04             Disabled        10/10
```

Meaning The output shows that Device R5's interfaces have been correctly configured with the ISO family, and that the interfaces have been placed into the correct levels.

You can also see that Device R5 has elected itself as the designated intermediate system (DIS) on its broadcast-capable IS-IS interfaces.

Verifying IS-IS Adjacencies

Purpose Verify that the expected adjacencies have formed between Device R5 and its IS-IS neighbors.

Action From operational mode, enter the **show isis adjacency detail** command.

```
user@R5> show isis adjacency detail
```

```

R3
Interface: fe-1/2/0.0, Level: 2, State: Up, Expires in 25 secs
Priority: 64, Up/Down transitions: 1, Last transition: 03:19:31 ago
Circuit type: 2, Speaks: IP, IPv6, MAC address: 0:5:85:8f:c8:bc
Topologies: Unicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R5.03, IP addresses: 10.0.0.21

R4
Interface: fe-1/2/1.0, Level: 2, State: Up, Expires in 24 secs
Priority: 64, Up/Down transitions: 1, Last transition: 03:19:36 ago
Circuit type: 2, Speaks: IP, IPv6, MAC address: 0:5:85:8f:c8:bc
Topologies: Unicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R5.02, IP addresses: 10.0.0.25

R6
Interface: fe-1/2/0.0, Level: 1, State: Up, Expires in 6 secs
Priority: 64, Up/Down transitions: 1, Last transition: 03:20:24 ago
Circuit type: 1, Speaks: IP, IPv6, MAC address: 0:5:85:8f:c8:bd
Topologies: Unicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R6.02, IP addresses: 10.0.0.30

R7
Interface: fe-1/2/3.0, Level: 1, State: Up, Expires in 21 secs
Priority: 64, Up/Down transitions: 1, Last transition: 03:19:29 ago
Circuit type: 1, Speaks: IP, IPv6, MAC address: 0:5:85:8f:c8:bc
Topologies: Unicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R5.04, IP addresses: 10.0.0.37

```

Meaning These results confirm that Device R5 has two Level 2 adjacencies and two Level 1 adjacencies.

Examining the IS-IS Database

Purpose Because Device R5 is a Level 1/Level 2 (L1/L2) attached router, examine the Level 1 link-state database associated with area 49.002 to confirm that loopback addresses from backbone routers are not being advertised into the Level 1 area.

Action From operational mode, enter the **show isis database detail** command.

```

user@R5> show isis database detail
IS-IS level 1 link-state database:

R5.00-00 Sequence: 0x19, Checksum: 0x7488, Lifetime: 727 secs
  IS neighbor: R5.04                      Metric:      10
  IS neighbor: R6.02                      Metric:      10
  IP prefix: 10.0.0.28/30                  Metric:      10 Internal Up
  IP prefix: 10.0.0.36/30                  Metric:      10 Internal Up

R5.04-00 Sequence: 0x14, Checksum: 0x2668, Lifetime: 821 secs
  IS neighbor: R5.00                      Metric:        0
  IS neighbor: R7.00                      Metric:        0

R6.00-00 Sequence: 0x17, Checksum: 0xa65, Lifetime: 774 secs
  IS neighbor: R6.02                      Metric:      10
  IS neighbor: R7.02                      Metric:      10

```

```

IP prefix: 10.0.0.28/30          Metric:      10 Internal Up
IP prefix: 10.0.0.32/30          Metric:      10 Internal Up
IP prefix: 192.168.0.6/32        Metric:      0 Internal Up

R6.02-00 Sequence: 0x13, Checksum: 0xd1c0, Lifetime: 908 secs
IS neighbor: R5.00               Metric:      0
IS neighbor: R6.00               Metric:      0

R7.00-00 Sequence: 0x17, Checksum: 0xe39, Lifetime: 775 secs
IS neighbor: R5.04               Metric:      10
IS neighbor: R7.02               Metric:      10
IP prefix: 10.0.0.32/30          Metric:      10 Internal Up
IP prefix: 10.0.0.36/30          Metric:      10 Internal Up
IP prefix: 192.168.0.7/32        Metric:      0 Internal Up

R7.02-00 Sequence: 0x13, Checksum: 0x404d, Lifetime: 966 secs
IS neighbor: R6.00               Metric:      0
IS neighbor: R7.00               Metric:      0

IS-IS level 2 link-state database:

R3.00-00 Sequence: 0x17, Checksum: 0x5f84, Lifetime: 1085 secs
IS neighbor: R4.02               Metric:      10
IS neighbor: R5.03               Metric:      10
IP prefix: 10.0.0.16/30          Metric:      10 Internal Up
IP prefix: 10.0.0.20/30          Metric:      10 Internal Up
IP prefix: 10.0.0.40/30          Metric:      10 Internal Up
IP prefix: 192.168.0.3/32        Metric:      0 Internal Up

R4.00-00 Sequence: 0x17, Checksum: 0xab3a, Lifetime: 949 secs
IS neighbor: R4.02               Metric:      10
IS neighbor: R5.02               Metric:      10
IP prefix: 10.0.0.16/30          Metric:      10 Internal Up
IP prefix: 10.0.0.24/30          Metric:      10 Internal Up
IP prefix: 192.168.0.4/32        Metric:      0 Internal Up

R4.02-00 Sequence: 0x14, Checksum: 0xf2a8, Lifetime: 1022 secs
IS neighbor: R3.00               Metric:      0
IS neighbor: R4.00               Metric:      0

R5.00-00 Sequence: 0x1f, Checksum: 0x20d7, Lifetime: 821 secs
IS neighbor: R5.02               Metric:      10
IS neighbor: R5.03               Metric:      10
IP prefix: 10.0.0.20/30          Metric:      10 Internal Up
IP prefix: 10.0.0.24/30          Metric:      10 Internal Up
IP prefix: 10.0.0.28/30          Metric:      10 Internal Up
IP prefix: 10.0.0.32/30          Metric:      20 Internal Up
IP prefix: 10.0.0.36/30          Metric:      10 Internal Up
IP prefix: 192.168.0.5/32        Metric:      0 Internal Up
IP prefix: 192.168.0.6/32        Metric:      10 Internal Up
IP prefix: 192.168.0.7/32        Metric:      10 Internal Up

R5.02-00 Sequence: 0x14, Checksum: 0x6135, Lifetime: 977 secs
IS neighbor: R4.00               Metric:      0
IS neighbor: R5.00               Metric:      0

R5.03-00 Sequence: 0x14, Checksum: 0x1483, Lifetime: 1091 secs
IS neighbor: R3.00               Metric:      0
IS neighbor: R5.00               Metric:      0

```

Meaning This display indicates that Device R5's loopback interface is correctly configured to run Level 2 only. Had Level 1 operation been enabled on lo0.0, Device R5 would have then included its loopback address in its Level 1 link-state PDU.

You can also see that Device R5 has Level 2 link-state PDUs, received from its adjacent neighbors.

Like an OSPF totally stubby area, no backbone (Level 2) or external prefixes are leaked into a Level 1 area, by default. Level 1 prefixes are leaked up into the IS-IS backbone, however, as can be seen in Device R5's Level 2 link-state PDU.

Related Documentation

- *Understanding IS-IS Areas to Divide an Autonomous System into Smaller Groups*

Configuring IS-IS Authentication

All IS-IS protocol exchanges can be authenticated to guarantee that only trusted routing devices participate in the autonomous system (AS) routing. By default, IS-IS authentication is disabled on the routing device.

To configure IS-IS authentication, you must define an authentication password and specify the authentication type.

You can configure one of the following authentication methods:

- Simple authentication—Uses a text password that is included in the transmitted packet. The receiving routing device uses an authentication key (password) to verify the packet. Simple authentication is included for compatibility with existing IS-IS implementations. However, we recommend that you do *not* use this authentication method because it is insecure (the text can be “sniffed”).



CAUTION: A simple password that exceeds 254 characters is truncated.

- HMAC-MD5 authentication—Uses an iterated cryptographic hash function. The receiving routing device uses an authentication key (password) to verify the packet.

You can also configure more fine-grained interface-level authentication for hello packets.

To enable authentication and specify an authentication method, include the **authentication-type** statement, specifying the **simple** or **md5** authentication type:

```
authentication-type authentication;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

To configure a password, include the **authentication-key** statement. The authentication password for all routing devices in a domain must be the same.

```
authentication-key key;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

To configure hitless authentication key rollover, include the [authentication-key-chain \(Protocols IS-IS\)](#) statement.

The password can contain up to 255 characters. If you include spaces, enclose all characters in quotation marks (" ").

If you are using the Junos OS IS-IS software with another implementation of IS-IS, the other implementation must be configured to use the same password for the domain, the area, and all interfaces that are shared with a Junos OS implementation.

Authentication of hello packets, partial sequence number PDU (PSNP), and complete sequence number PDU (CSNP) can be suppressed to enable interoperability with the routing software of different vendors. Different vendors handle authentication in various ways, and suppressing authentication for different PDU types might be the simplest way to allow compatibility within the same network.

To configure IS-IS to generate authenticated packets, but not to check the authentication on received packets, include the **no-authentication-check** statement:

no-authentication-check;

To suppress authentication of IS-IS hello packets, include the **no-hello-authentication** statement:

no-hello-authentication;

To suppress authentication of PSNPs, include the **no-psnp-authentication** statement:

no-psnp-authentication;

To suppress authentication of CSNPs, include the **no-csnp-authentication** statement:

no-csnp-authentication;

For a list of hierarchy levels at which you can include these statements, see the statement summary sections for these statements.



NOTE: The **authentication** and the **no-authentication** statements must be configured at the same hierarchy level. Configuring **authentication** at the [edit protocols isis interface *interface-name*] hierarchy level and configuring **no-authentication** at the [edit protocols isis] hierarchy level has no effect.

Related Documentation

- [Configuring IS-IS Authentication Without Network-Wide Deployment on page 4184](#)

Configuring IS-IS Authentication Without Network-Wide Deployment

To allow the use of authentication without requiring network-wide deployment, include the **loose-authentication-check** statement:

[loose-authentication-check;](#)

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

**Related
Documentation**

- [Understanding Hitless Authentication Key Rollover for IS-IS on page 4255](#)
- [Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 4256](#)

Example: Redistributing OSPF Routes into IS-IS

This example shows how to redistribute OSPF routes into an IS-IS network.

- [Requirements on page 4185](#)
- [Overview on page 4185](#)
- [Configuration on page 4186](#)
- [Verification on page 4191](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

Export policy can be applied to IS-IS to facilitate route redistribution.


Junos OS does not support the application of import policy for link-state routing protocols like IS-IS because such policies can lead to inconsistent link-state database (LSDB) entries, which in turn can result in routing inconsistencies.

In this example, OSPF routes 192.168.0/24 through 192.168.3/24 are redistributed into IS-IS area 49.0002 from Device R2.

In addition, policies are configured to ensure that Device R1 can reach destinations on the 10.0.0.44/30 network, and that Device R3 can reach destinations on the 10.0.0.36/30 network. This enables end-to-end reachability.

[Figure 98](#) shows the topology used in this example.

Figure 98: IS-IS Route Redistribution Topology



"CLI Quick Configuration" on page 4186 shows the configuration for all of the devices in Figure 98. The section "Step-by-Step Procedure" on page 4187 describes the steps on Device R2. "Step-by-Step Procedure" on page 4188 describes the steps on Device R3.

Configuration

CLI Quick Configuration	To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.
Device R1	<pre> set interfaces fe-1/2/0 unit 0 description to-R7 set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.38/30 set interfaces fe-1/2/0 unit 0 family iso set interfaces lo0 unit 0 family inet address 172.16.3.5/32 set interfaces lo0 unit 0 family iso address 49.0002.0172.0016.0305.00 set protocols isis interface fe-1/2/0.0 set protocols isis interface lo0.0 </pre>
Device R2	<pre> set interfaces fe-1/2/1 unit 0 description to-R5 set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.37/30 set interfaces fe-1/2/1 unit 0 family iso set interfaces fe-1/2/0 unit 0 description to-OSPF-network set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.45/30 set interfaces lo0 unit 0 family inet address 172.16.9.7/32 set interfaces lo0 unit 0 family iso address 49.0002.0172.0016.0907.00 set protocols isis export ospf-isis set protocols isis export send-direct-to-isis-neighbors set protocols isis interface fe-1/2/1.0 set protocols isis interface lo0.0 set protocols ospf export send-direct-to-ospf-neighbors </pre>

```

set protocols ospf area 0.0.0.1 interface fe-1/2/0.0
set protocols ospf area 0.0.0.1 interface lo0.0 passive
set policy-options policy-statement ospf-isis term 1 from protocol ospf
set policy-options policy-statement ospf-isis term 1 from route-filter 192.168.0.0/22
  longer
set policy-options policy-statement ospf-isis term 1 then accept
set policy-options policy-statement send-direct-to-isis-neighbors from protocol direct
set policy-options policy-statement send-direct-to-isis-neighbors from route-filter
  10.0.0.44/30 exact
set policy-options policy-statement send-direct-to-isis-neighbors then accept
set policy-options policy-statement send-direct-to-ospf-neighbors from protocol direct
set policy-options policy-statement send-direct-to-ospf-neighbors from route-filter
  10.0.0.36/30 exact
set policy-options policy-statement send-direct-to-ospf-neighbors then accept

```

Device R3

```

set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.46/30
set interfaces lo0 unit 0 family inet address 192.168.1.1/32
set interfaces lo0 unit 0 family inet address 192.168.2.1/32
set interfaces lo0 unit 0 family inet address 192.168.3.1/32
set interfaces lo0 unit 0 family inet address 192.168.0.1/32
set protocols ospf export ospf
set protocols ospf area 0.0.0.1 interface fe-1/2/0.0
set protocols ospf area 0.0.0.1 interface lo0.0 passive
set policy-options policy-statement ospf term 1 from protocol static
set policy-options policy-statement ospf term 1 then accept
set routing-options static route 192.168.0.0/24 discard
set routing-options static route 192.168.1.0/24 discard
set routing-options static route 192.168.3.0/24 discard
set routing-options static route 192.168.2.0/24 discard

```

Step-by-Step Procedure

To configure Device R2:

1. Configure the network interfaces.

```

[edit interfaces]
user@R2# set fe-1/2/1 unit 0 description to-R5
user@R2# set fe-1/2/1 unit 0 family inet address 10.0.0.37/30
user@R2# set fe-1/2/1 unit 0 family iso
user@R2# set fe-1/2/0 unit 0 description to-OSPF-network
user@R2# set fe-1/2/0 unit 0 family inet address 10.0.0.45/30
user@R2# set lo0 unit 0 family inet address 172.16.9.7/32
user@R2# set lo0 unit 0 family iso address 49.0002.0172.0016.0907.00

```
2. Configure IS-IS on the interface facing Device R1 and the loopback interface.

```

[edit protocols isis]
user@R2# set interface fe-1/2/1.0
user@R2# set interface lo0.0

```
3. Configure the policy that enables Device R1 to reach the 10.0.0.44/30 network.

```

[edit policy-options policy-statement send-direct-to-isis-neighbors]
user@R2# set from protocol direct
user@R2# set from route-filter 10.0.0.44/30 exact
user@R2# set then accept

```
4. Apply the policy that enables Device R1 to reach the 10.0.0.44/30 network.

```
[edit protocols isis]
user@R2# set export send-direct-to-isis-neighbors
```

5. Configure OSPF on the interfaces.

```
[edit protocols ospf]
user@R2# set area 0.0.0.1 interface fe-1/2/0.0
user@R2# set area 0.0.0.1 interface lo0.0 passive
```

6. Configure the OSPF route redistribution policy.

```
[edit policy-options policy-statement ospf-isis term 1]
user@R2# set from protocol ospf
user@R2# set from route-filter 192.168.0.0/22 longer
user@R2# set then accept
```

7. Apply the OSPF route redistribution policy to the IS-IS instance.

```
[edit protocols isis]
user@R2# set export ospf-isis
```

8. Configure the policy that enables Device R3 to reach the 10.0.0.36/30 network.

```
[edit policy-options policy-statement send-direct-to-ospf-neighbors]
user@R2# set from protocol direct
user@R2# set from route-filter 10.0.0.36/30 exact
user@R2# set then accept
```

9. Apply the policy that enables Device R3 to reach the 10.0.0.36/30 network.

```
[edit protocols ospf]
user@R2# set export send-direct-to-ospf-neighbors
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure multi-level IS-IS:

1. Configure the network interfaces.

Multiple addresses are configured on the loopback interface to simulate multiple route destinations.

```
[edit interfaces]
user@R3# set fe-1/2/0 unit 0 family inet address 10.0.0.46/30
user@R3# set lo0 unit 0 family inet address 192.168.1.1/32
user@R3# set lo0 unit 0 family inet address 192.168.2.1/32
user@R3# set lo0 unit 0 family inet address 192.168.3.1/32
user@R3# set lo0 unit 0 family inet address 192.168.0.1/32
```

2. Configure static routes to the loopback interface addresses.

These are the routes that are redistributed into IS-IS.

```
[edit routing-options static]
user@R3# set route 192.168.0.0/24 discard
user@R3# set route 192.168.1.0/24 discard
user@R3# set route 192.168.3.0/24 discard
user@R3# set route 192.168.2.0/24 discard
```

3. Configure OSPF on the interfaces.

```
[edit protocols ospf area 0.0.0.1]
user@R3# set interface fe-1/2/0.0
user@R3# set interface lo0.0 passive
```
4. Configure the OSPF policy to export the static routes.

```
[edit policy-options policy-statement ospf term 1]
user@R3# set from protocol static
user@R3# set then accept
```
5. Apply the OSPF export policy.

```
[edit protocols ospf]
user@R3# set export ospf
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
Device R2 user@R2# show interfaces
fe-1/2/1 {
  unit 0 {
    description to-R5;
    family inet {
      address 10.0.0.37/30;
    }
    family iso;
  }
}
fe-1/2/0 {
  unit 0 {
    description to-OSPF-network;
    family inet {
      address 10.0.0.45/30;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 172.16.9.7/32;
    }
    family iso {
      address 49.0002.0172.0016.0907.00;
    }
  }
}

user@R2# show protocols
isis {
  export [ ospf-isis send-direct-to-isis-neighbors ];
  interface fe-1/2/1.0;
  interface lo0.0;
}
```

```
ospf {
  export send-direct-to-ospf-neighbors;
  area 0.0.0.1 {
    interface fe-1/2/0.0;
    interface lo0.0 {
      passive;
    }
  }
}

user@R2# show policy-options
policy-statement ospf-isis {
  term 1 {
    from {
      protocol ospf;
      route-filter 192.168.0.0/22 longer;
    }
    then accept;
  }
}
policy-statement send-direct-to-isis-neighbors {
  from {
    protocol direct;
    route-filter 10.0.0.44/30 exact;
  }
  then accept;
}
policy-statement send-direct-to-ospf-neighbors {
  from {
    protocol direct;
    route-filter 10.0.0.36/30 exact;
  }
  then accept;
}
```

Device R3

```
user@R3# show interfaces
fe-1/2/0 {
  unit 0 {
    family inet {
      address 10.0.0.46/30;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.168.1.1/32;
      address 192.168.2.1/32;
      address 192.168.3.1/32;
      address 192.168.0.1/32;
    }
  }
}

user@R3# show protocols
ospf {
```

```

export ospf;
area 0.0.0.1 {
    interface fe-1/2/0.0;
    interface lo0.0 {
        passive;
    }
}
}

user@R3# show policy-options
policy-statement ospf {
    term 1 {
        from protocol static;
        then accept;
    }
}

user@R3# show routing-options
static {
    route 192.168.0.0/24 discard;
    route 192.168.1.0/24 discard;
    route 192.168.3.0/24 discard;
    route 192.168.2.0/24 discard;
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying OSPF Route Advertisement on page 4191](#)
- [Verifying Route Redistribution on page 4192](#)
- [Verifying Connectivity on page 4192](#)

Verifying OSPF Route Advertisement

Purpose Make sure that the expected routes are advertised by OSPF.

Action From operational mode on Device R2, enter the **show route protocol ospf** command.

```
user@R2> show route protocol ospf
```

```
inet.0: 15 destinations, 15 routes (15 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```

192.168.0.0/24    *[OSPF/150] 03:54:21, metric 0, tag 0
                  > to 10.0.0.46 via fe-1/2/0.0
192.168.0.1/32   *[OSPF/10] 03:54:21, metric 1
                  > to 10.0.0.46 via fe-1/2/0.0
192.168.1.0/24   *[OSPF/150] 03:54:21, metric 0, tag 0
                  > to 10.0.0.46 via fe-1/2/0.0
192.168.1.1/32   *[OSPF/10] 03:54:21, metric 1
                  > to 10.0.0.46 via fe-1/2/0.0
192.168.2.0/24   *[OSPF/150] 03:54:21, metric 0, tag 0
                  > to 10.0.0.46 via fe-1/2/0.0

```

```
192.168.2.1/32      *[OSPF/10] 03:54:21, metric 1
                   > to 10.0.0.46 via fe-1/2/0.0
192.168.3.0/24      *[OSPF/150] 03:54:21, metric 0, tag 0
                   > to 10.0.0.46 via fe-1/2/0.0
192.168.3.1/32      *[OSPF/10] 03:54:21, metric 1
                   > to 10.0.0.46 via fe-1/2/0.0
224.0.0.5/32        *[OSPF/10] 03:56:03, metric 1
                   MultiRecv

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Meaning The 192.168/16 routes are advertised by OSPF.

Verifying Route Redistribution

Purpose Make sure that the expected routes are redistributed from OSPF into IS-IS.

Action From operational mode on Device R1, enter the **show route protocol isis** command.

```
user@R1> show route protocol isis
```

```
inet.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
10.0.0.44/30        *[IS-IS/160] 03:45:24, metric 20
                   > to 10.0.0.37 via fe-1/2/0.0
172.16.9.7/32       *[IS-IS/15] 03:49:46, metric 10
                   > to 10.0.0.37 via fe-1/2/0.0
192.168.0.0/24       *[IS-IS/160] 03:49:46, metric 10
                   > to 10.0.0.37 via fe-1/2/0.0
192.168.0.1/32       *[IS-IS/160] 03:49:46, metric 11, tag2 1
                   > to 10.0.0.37 via fe-1/2/0.0
192.168.1.0/24       *[IS-IS/160] 03:49:46, metric 10
                   > to 10.0.0.37 via fe-1/2/0.0
192.168.1.1/32       *[IS-IS/160] 03:49:46, metric 11, tag2 1
                   > to 10.0.0.37 via fe-1/2/0.0
192.168.2.0/24       *[IS-IS/160] 03:49:46, metric 10
                   > to 10.0.0.37 via fe-1/2/0.0
192.168.2.1/32       *[IS-IS/160] 03:49:46, metric 11, tag2 1
                   > to 10.0.0.37 via fe-1/2/0.0
192.168.3.0/24       *[IS-IS/160] 03:49:46, metric 10
                   > to 10.0.0.37 via fe-1/2/0.0
192.168.3.1/32       *[IS-IS/160] 03:49:46, metric 11, tag2 1
                   > to 10.0.0.37 via fe-1/2/0.0
```

```
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Meaning The 192.168/16 routes are redistributed into IS-IS.

Verifying Connectivity

Purpose Check that Device R1 can reach the destinations on Device R3.

Action From operational mode, enter the **ping** command.

```
user@R1> ping 192.168.1.1
```



```

PING 192.168.1.1 (192.168.1.1): 56 data bytes
64 bytes from 192.168.1.1: icmp_seq=0 ttl=63 time=2.089 ms
64 bytes from 192.168.1.1: icmp_seq=1 ttl=63 time=1.270 ms
64 bytes from 192.168.1.1: icmp_seq=2 ttl=63 time=2.135 ms

```

Meaning These results confirm that Device R1 can reach the destinations in the OSPF network.

Related Documentation

- [Understanding Routing Policies](#)

Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies

- [Understanding IS-IS IPv4 and IPv6 Unicast Topologies on page 4193](#)
- [Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies on page 4193](#)

Understanding IS-IS IPv4 and IPv6 Unicast Topologies

You can configure IS-IS to calculate an alternate IPv6 unicast topology, in addition to the normal IPv4 unicast topology, and add the corresponding routes to inet6.0. The IS-IS interface metrics for the IPv4 topology can be configured independently of the IPv6 metrics. You can also selectively disable interfaces from participating in the IPv6 topology while continuing to participate in the IPv4 topology. This enables you to exercise control over the paths that unicast data takes through a network.

A topology is the set of joined nodes. IS-IS evaluates all the paths in a single topology for each IS-IS level and uses the shortest-path-first (SPF) algorithm to determine the best path among all the feasible paths. Topology discovery and SPF calculation is performed in a protocol-neutral fashion because it is done at Layer 2 of the OSI model. If you load the topology with reachability information for a certain protocol (for example, IP), the assumption is that the circuits that are supposed to provide reachability between routing devices can carry the protocol. The SPF algorithm has a per-link orientation, not a per-address family or per-protocol orientation.

Multitopology routing enables you to override this default behavior by enabling a per-address family, per-protocol SPF calculation.

The additional CPU load associated with multiple runs of the SPF algorithm is generally not an issue with the processing power available on today's routing device control planes.

The multitopology extensions alter existing type, length, and value (TLV) tuples by adding a topology ID. Each routing device in a given topology maintains its adjacencies and runs a per-topology SPF calculation.

Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies

This example shows how to configure IS-IS to calculate an alternate IPv6 unicast topology, in addition to the normal IPv4 unicast topology.

- [Requirements on page 4194](#)
- [Overview on page 4194](#)

- [Configuration on page 4195](#)
- [Verification on page 4199](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

This example focuses on IPv4 and IPv6 unicast topologies. The IS-IS interface metrics for the IPv4 topology can be configured independently of the IPv6 metrics. You can also selectively disable interfaces from participating in the IPv6 topology while continuing to participate in the IPv4 topology. This enables you to exercise control over the paths that unicast data takes through a network.

To enable an IPv6 unicast topology for IS-IS, include the **ipv6-unicast** statement:

```
isis {  
  topologies {  
    ipv6-unicast;  
  }  
}
```

To configure a metric for the IPv6 unicast topology, include the **ipv6-unicast-metric** statement:


```
isis {  
  interface interface-name {  
    level level-number {  
      ipv6-unicast-metric number;  
    }  
  }  
}
```

To exclude an interface from the IPv6 unicast topologies for IS-IS, include the **no-ipv6-unicast** statement:

```
isis {  
  interface interface-name {  
    no-ipv6-unicast;  
  }  
}
```

[Figure 99](#) shows the topology used in this example. The black lines indicate link membership in the IPv6 topology. The gray lines indicate membership to the IPv4 topology. Using regular TLVs, it would not be possible to build multiple topologies and run an SPF calculation based on them. The multitopology extensions describe an extension to carry the set of supported protocols in the hello packet. After activating multitopology routing support on a link, the link carries all the topologies that the underlying circuit is able to relay.

Figure 99: IS-IS IPv4 and IPv6 Unicast Topologies



“CLI Quick Configuration” on page 4195 shows the configuration for all of the devices in Figure 99. The section “Step-by-Step Procedure” on page 4197 describes the steps on Device R1.

Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device R1

```

set interfaces so-1/2/0 unit 0 family inet address 10.0.0.1/30
set interfaces so-1/2/0 unit 0 family iso
set interfaces so-1/2/0 unit 0 family inet6 address 2001:db8:0:5::/64 eui-64
set interfaces so-1/2/1 unit 0 family inet address 10.0.0.13/30
set interfaces so-1/2/1 unit 0 family iso
set interfaces so-1/2/2 unit 0 family inet address 10.0.0.17/30
set interfaces so-1/2/2 unit 0 family iso
set interfaces so-1/2/2 unit 0 family inet6 address 2001:db8:0:1::/64 eui-64
set interfaces lo0 unit 0 family inet address 192.168.0.1/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0001.00
set interfaces lo0 unit 0 family inet6 address 2001:db8::1/128
set protocols isis topologies ipv6-unicast
set protocols isis interface so-1/2/0.0
set protocols isis interface so-1/2/1.0 no-ipv6-unicast
set protocols isis interface so-1/2/2.0
set protocols isis interface lo0.0

```

Device R2

```
set interfaces so-1/2/0 unit 0 family inet address 10.0.0.2/30
set interfaces so-1/2/0 unit 0 family iso
set interfaces so-1/2/0 unit 0 family inet6 address 2001:db8:0:5::/64 eui-64
set interfaces so-1/2/1 unit 0 family inet address 10.0.0.9/30
set interfaces so-1/2/1 unit 0 family iso
set interfaces so-1/2/1 unit 0 family inet6 address 2001:db8:0:4::/64 eui-64
set interfaces lo0 unit 0 family inet address 192.168.0.2/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0002.00
set interfaces lo0 unit 0 family inet6 address 2001:db8::2/128
set protocols isis topologies ipv6-unicast
set protocols isis interface so-1/2/0.2
set protocols isis interface so-1/2/1.0
set protocols isis interface lo0.0
```

Device R3

```
set interfaces so-1/2/0 unit 0 family inet address 10.0.0.18/30
set interfaces so-1/2/0 unit 0 family iso
set interfaces so-1/2/0 unit 0 family inet6 address 2001:db8:0:1::/64 eui-64
set interfaces so-1/2/1 unit 0 family inet address 10.0.0.25/30
set interfaces so-1/2/1 unit 0 family iso
set interfaces lo0 unit 0 family inet address 192.168.0.3/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0003.00
set interfaces lo0 unit 0 family inet6 address 2001:db8::3/128
set protocols isis topologies ipv6-unicast
set protocols isis interface so-1/2/0.0
set protocols isis interface so-1/2/1.0 no-ipv6-unicast
set protocols isis interface lo0.0
```

Device R4

```
set interfaces so-1/2/0 unit 0 family inet address 10.0.0.5/30
set interfaces so-1/2/0 unit 0 family iso
set interfaces so-1/2/0 unit 0 family inet6 address 2001:db8:0:2::/64 eui-64
set interfaces so-1/2/1 unit 0 family inet address 10.0.0.10/30
set interfaces so-1/2/1 unit 0 family iso
set interfaces so-1/2/1 unit 0 family inet6 address 2001:db8:0:1::/64 eui-64
set interfaces so-1/2/2 unit 0 family inet address 10.0.0.21/30
set interfaces so-1/2/2 unit 0 family iso
set interfaces lo0 unit 0 family inet address 192.168.0.4/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0004.00
set interfaces lo0 unit 0 family inet6 address 2001:db8::4/128
set protocols isis topologies ipv6-unicast
set protocols isis interface so-1/2/0.0
set protocols isis interface so-1/2/1.0
set protocols isis interface so-1/2/2.0 no-ipv6-unicast
set protocols isis interface lo0.0
```

Device R5

```
set interfaces so-1/2/0 unit 0 family inet address 10.0.0.14/30
set interfaces so-1/2/0 unit 0 family iso
set interfaces so-1/2/1 unit 0 family inet address 10.0.0.22/30
set interfaces so-1/2/1 unit 0 family iso
set interfaces so-1/2/2 unit 0 family inet address 10.0.0.26/30
set interfaces so-1/2/2 unit 0 family iso
set interfaces so-1/2/3 unit 0 family iso
set interfaces so-1/2/3 unit 0 family inet6 address 2001:db8:0:3::/64 eui-64
set interfaces lo0 unit 0 family inet address 192.168.0.5/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0005.00
set interfaces lo0 unit 0 family inet6 address 2001:db8::5/128
```

```

set protocols isis topologies ipv6-unicast
set protocols isis interface so-1/2/0.0 no-ipv6-unicast
set protocols isis interface so-1/2/1.0 no-ipv6-unicast
set protocols isis interface so-1/2/2.0 no-ipv6-unicast
set protocols isis interface so-1/2/3.0
set protocols isis interface lo0.0

```

Device R6

```

set interfaces so-1/2/0 unit 0 family inet address 10.0.0.6/30
set interfaces so-1/2/0 unit 0 family iso
set interfaces so-1/2/0 unit 0 family inet6 address 2001:db8:0:2::/64 eui-64
set interfaces so-1/2/1 unit 0 family iso
set interfaces so-1/2/1 unit 0 family inet6 address 2001:db8:0:3::/64 eui-64
set interfaces lo0 unit 0 family inet address 192.168.0.6/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0006.00
set interfaces lo0 unit 0 family inet6 address 2001:db8::6/128
set protocols isis topologies ipv6-unicast
set protocols isis interface so-1/2/0.0
set protocols isis interface so-1/2/1.0
set protocols isis interface lo0.0

```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure an alternate IPv6 unicast topology:

1. Configure the interfaces.

```

[edit interfaces]
user@R1# set so-1/2/0 unit 0 family inet address 10.0.0.1/30
user@R1# set so-1/2/0 unit 0 family iso
user@R1# set so-1/2/0 unit 0 family inet6 address 2001:db8:0:5::/64 eui-64
user@R1# set so-1/2/1 unit 0 family inet address 10.0.0.13/30
user@R1# set so-1/2/1 unit 0 family iso
user@R1# set so-1/2/2 unit 0 family inet address 10.0.0.17/30
user@R1# set so-1/2/2 unit 0 family iso
user@R1# set so-1/2/2 unit 0 family inet6 address 2001:db8:0:1::/64 eui-64
user@R1# set lo0 unit 0 family inet address 192.168.0.1/32
user@R1# set lo0 unit 0 family iso address 49.0002.0192.0168.0001.00
user@R1# set lo0 unit 0 family inet6 address 2001:db8::1/128

```

2. Enable IS-IS on the interfaces.

```

[edit protocols isis]
user@R1# set interface so-1/2/0.0
user@R1# set interface so-1/2/1.0
user@R1# set interface so-1/2/2.0
user@R1# set interface lo0.0

```

3. Enable multitopology routing on the IS-IS interfaces.

The **ipv6-unicast** statement enables multitopology IS-IS routing on all interfaces that have **family iso** and **family inet6** configured and are listed at the **[edit protocols isis interface]** hierarchy level.

```

[edit protocols isis]
user@R1# set topologies ipv6-unicast

```

4. Disable IPv6 unicast support on a given interface.

If you do not want to run multiprotocol IS-IS routing for IPv6 on a given interface, you can disable multiprotocol routing by including the **no-ipv6-unicast** statement in the IS-IS interface configuration.

```
[edit protocols isis]
user@R1# set interface so-1/2/1.0 no-ipv6-unicast
```

Results From configuration mode, confirm your configuration by entering the **show interfaces** and **show protocols** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show interfaces
so-1/2/0 {
  unit 0 {
    family inet {
      address 10.0.0.1/30;
    }
    family iso;
    family inet6 {
      address 2001:db8:0:5::/64 {
        eui-64;
      }
    }
  }
}
so-1/2/1 {
  unit 0 {
    family inet {
      address 10.0.0.13/30;
    }
    family iso;
  }
}
so-1/2/2 {
  unit 0 {
    family inet {
      address 10.0.0.17/30;
    }
    family iso;
    family inet6 {
      address 2001:db8:0:1::/64 {
        eui-64;
      }
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.168.0.1/32;
    }
    family iso {
      address 49.0002.0192.0168.0001.00;
```

```

    }
    family inet6 {
        address 2001:db8::1/128;
    }
}

user@R1# show protocols
isis {
    topologies ipv6-unicast;
    interface so-1/2/0.0;
    interface so-1/2/1.0 {
        no-ipv6-unicast;
    }
    interface so-1/2/2.0;
    interface lo0.0;
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Checking the Topologies on Neighbors on page 4199](#)
- [Checking the IS-IS SPF Calculations on page 4200](#)
- [Checking the Tcpdump Output on page 4201](#)

Checking the Topologies on Neighbors

Purpose Determine what topologies are supported on neighboring IS-IS devices.

Action From operational mode, enter the **show isis adjacency detail** command.

```
user@R1> show isis adjacency detail
```

R2

```

Interface: so-1/2/0.0, Level: 3, State: Up, Expires in 24 secs
Priority: 0, Up/Down transitions: 1, Last transition: 05:28:16 ago
Circuit type: 3, Speaks: IP, IPv6
Topologies: Unicast, IPV6-Unicast
Restart capable: Yes, Adjacency advertisement: Advertise
IP addresses: 10.0.0.2
IPv6 addresses: fe80::2a0:a514:0:24c

```

R5

```

Interface: so-1/2/1.0, Level: 3, State: Up, Expires in 21 secs
Priority: 0, Up/Down transitions: 1, Last transition: 05:27:47 ago
Circuit type: 3, Speaks: IP, IPv6
Topologies: Unicast
Restart capable: Yes, Adjacency advertisement: Advertise
IP addresses: 10.0.0.14

```

R3

```

Interface: so-1/2/2.0, Level: 3, State: Up, Expires in 22 secs
Priority: 0, Up/Down transitions: 1, Last transition: 05:27:25 ago
Circuit type: 3, Speaks: IP, IPv6
Topologies: Unicast, IPV6-Unicast

```

```
Restart capable: Yes, Adjacency advertisement: Advertise
IP addresses: 10.0.0.18
IPv6 addresses: fe80::2a0:a514:0:124c
```

Meaning As expected, the adjacency with Device R5 only supports the IPv4 unicast topology, while the adjacencies with Device R2 and Device R3 support both the IPv4 and IPv6 topologies.

Checking the IS-IS SPF Calculations

Purpose Verify that separate SPF calculations are being run for IPv4 and IPv6.

Action From operational mode, enter the **show isis spf brief** command.

```
user@R1> show isis spf brief
```

IPv4 Unicast IS-IS level 1 SPF results:

Node	Metric	Interface	NH	Via	SNPA
R6.00	20	so-1/2/1.0	IPV4 R5		
R4.00	20	so-1/2/0.0	IPV4 R2		
R5.00	10	so-1/2/1.0	IPV4 R5		
R3.00	10	so-1/2/2.0	IPV4 R3		
R2.00	10	so-1/2/0.0	IPV4 R2		
R1.00	0				

6 nodes

IPv4 Unicast IS-IS level 2 SPF results:

Node	Metric	Interface	NH	Via	SNPA
R6.00	20	so-1/2/1.0	IPV4 R5		
R4.00	20	so-1/2/0.0	IPV4 R2		
R5.00	10	so-1/2/1.0	IPV4 R5		
R3.00	10	so-1/2/2.0	IPV4 R3		
R2.00	10	so-1/2/0.0	IPV4 R2		
R1.00	0				

6 nodes

IPv6 Unicast IS-IS level 1 SPF results:

Node	Metric	Interface	NH	Via	SNPA
R5.00	40	so-1/2/0.0	IPV6 R2		
R6.00	30	so-1/2/0.0	IPV6 R2		
R4.00	20	so-1/2/0.0	IPV6 R2		
R3.00	10	so-1/2/2.0	IPV6 R3		
R2.00	10	so-1/2/0.0	IPV6 R2		
R1.00	0				

6 nodes

IPv6 Unicast IS-IS level 2 SPF results:

Node	Metric	Interface	NH	Via	SNPA
R5.00	40	so-1/2/0.0	IPV6 R2		
R6.00	30	so-1/2/0.0	IPV6 R2		
R4.00	20	so-1/2/0.0	IPV6 R2		
R3.00	10	so-1/2/2.0	IPV6 R3		
R2.00	10	so-1/2/0.0	IPV6 R2		
R1.00	0				

6 nodes

Meaning As expected, SPF calculations are being performed for IPv4 and IPv6 topologies.

Checking the Tcpcdump Output

Purpose Verify that the link can be a member of both the IPv4 unicast topology and the IPv6 unicast topology.

Action user@R1> **monitor traffic** detail interface so-1/2/0.0
[...]

```
15:52:35.719540 In IS-IS, length 82
p2p IIH, hlen: 20, v: 1, pdu-v: 1, sys-id-len: 6 (0), max-area: 3 (0)
source-id: 0192.0168.0002, holding time: 27s, Flags: [Level 1, Level
2]
circuit-id: 0x01, PDU length: 82
Point-to-point Adjacency State TLV #240, length: 15
Adjacency State: Up (0)
Extended Local circuit-ID: 0x00000054
Neighbor System-ID: 0192.0168.0001
Neighbor Extended Local circuit-ID: 0x00000043
Protocols supported TLV #129, length: 2
NLPID(s): IPv4 (0xcc), IPv6 (0x8e)
IPv4 Interface address(es) TLV #132, length: 4
IPv4 interface address: 10.0.0.2
IPv6 Interface address(es) TLV #232, length: 16
IPv6 interface address: fe80::2a0:a514:0:24c
Area address(es) TLV #1, length: 4
Area address (length: 3): 49.0002
Restart Signaling TLV #211, length: 3
Flags [none], Remaining holding time 0s
Multi Topology TLV #229, length: 4
  IPv4 unicast Topology (0x000), Flags: [none]
  IPv6 unicast Topology (0x002), Flags: [none]
```

Meaning The IS-IS hello (IIH) packet shows that IPv4 and IPv6 are supported. The hello packet lists valid IPv4 and IPv6 addresses, and therefore the routing device can create valid next-hop entries. The supported protocols are listed in the multitopology TLV #229.

Related Documentation

- [Example: Configuring IS-IS Dual Stacking of IPv4 and IPv6 Unicast Addresses](#)

Example: Configuring IS-IS Multicast Topology

- [IS-IS Multicast Topologies Overview on page 4202](#)
- [Example: Configuring IS-IS Multicast Topology on page 4203](#)

IS-IS Multicast Topologies Overview

Most multicast routing protocols perform a reverse-path forwarding (RPF) check on the source of multicast data packets. If a packet comes in on the interface that is used to send data to the source, the packet is accepted and forwarded to one or more downstream interfaces. Otherwise, the packet is discarded and a notification is sent to the multicast routing protocol running on the interface.

In certain instances, the unicast routing table used for the RPF check is also the table used for forwarding unicast data packets. Thus, unicast and multicast routing are congruent. In other cases, where it is preferred that multicast routing be independent of unicast routing, the multicast routing protocols are configured to perform the RPF check using an alternate unicast routing table `inet.2`.

You can configure IS-IS to calculate an alternate IPv4 multicast topology, in addition to the normal IPv4 unicast topology, and add the corresponding routes to `inet.2`. The IS-IS interface metrics for the multicast topology can be configured independently of the unicast metrics. You can also selectively disable interfaces from participating in the multicast topology while continuing to participate in the regular unicast topology. This enables you to exercise control over the paths that multicast data takes through a network so that it is independent of unicast data paths. You can also configure IS-IS to calculate an alternate IPv6 multicast topology, in addition to the normal IPv6 unicast topology.



NOTE: IS-IS only starts advertising the routes when the interface routes are in `inet.2`.

Table 330 lists the various IPv4 statements you can use to configure IS-IS topologies.

Table 330: IPv4 Statements

Statement	Description
<code>ipv4-multicast</code>	Enables an alternate IPv4 multicast topology.
<code>ipv4-multicast-metric <i>number</i></code>	Configures the multicast metric for an alternate IPv4 multicast topology.
<code>no-ipv4-multicast</code>	Excludes an interface from the IPv4 multicast topology.
<code>no-unicast-topology</code>	Excludes an interface from the IPv4 unicast topologies.

Table 331 lists the various IPv6 statements you can use to configure IS-IS topologies.

Table 331: IPv6 Statements

Statement	Description
<code>ipv6-multicast</code>	Enables an alternate IPv6 multicast topology.
<code>ipv6-multicast-metric <i>number</i></code>	Configures the multicast metric for an alternate IPv6 multicast topology.

Table 331: IPv6 Statements (*continued*)

Statement	Description
<code>ipv6-unicast-metric <i>number</i></code>	Configures the unicast metric for an alternate IPv6 multicast topology.
<code>no-ipv6-multicast</code>	Excludes an interface from the IPv6 multicast topology.
<code>no-ipv6-unicast</code>	Excludes an interface from the IPv6 unicast topologies.

For a list of hierarchy levels at which you can include these statements, see the statement summary sections for these statements.

Example: Configuring IS-IS Multicast Topology

This example shows how to configure a multicast topology for an IS-IS network.

- [Requirements on page 4203](#)
- [Overview on page 4203](#)
- [Configuration on page 4204](#)
- [Verification on page 4208](#)

Requirements

Before you begin, configure IS-IS on all routers. See [“Example: Configuring IS-IS” on page 4169](#) for information about the required IS-IS configuration.

This example uses the following hardware and software components:


- Junos OS Release 7.3 or later
- M Series, MX Series, and T Series routers

Overview

This example shows an IS-IS multicast topology configuration. Three routers are connected to each other. A loopback interface is configured on each router.

[Figure 100](#) shows the sample network.

Figure 100: Configuring IS-IS Multicast Topology



g040922

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Router R1

```
set protocols isis traceoptions file isis size 5m world-readable
set protocols isis traceoptions flag error
set protocols isis topologies ipv4-multicast
set protocols isis interface so-0/0/0 level 1 metric 15
set protocols isis interface so-0/0/0 level 1 ipv4-multicast-metric 18
set protocols isis interface so-0/0/0 level 2 metric 20
set protocols isis interface so-0/0/0 level 2 ipv4-multicast-metric 14
set protocols isis interface so-1/0/0 level 1 metric 13
set protocols isis interface so-1/0/0 level 1 ipv4-multicast-metric 12
set protocols isis interface so-1/0/0 level 2 metric 29
set protocols isis interface so-1/0/0 level 2 ipv4-multicast-metric 23
set protocols isis interface fxp0.0 disable
```

Router R2

```
set protocols isis traceoptions file isis size 5m world-readable
set protocols isis traceoptions flag error
set protocols isis topologies ipv4-multicast
set protocols isis interface so-0/0/0 level 1 metric 13
set protocols isis interface so-0/0/0 level 1 ipv4-multicast-metric 12
set protocols isis interface so-0/0/0 level 2 metric 29
```

```

set protocols isis interface so-0/0/0 level 2 ipv4-multicast-metric 23
set protocols isis interface so-1/0/0 level 1 metric 14
set protocols isis interface so-1/0/0 level 1 ipv4-multicast-metric 18
set protocols isis interface so-1/0/0 level 2 metric 32
set protocols isis interface so-1/0/0 level 2 ipv4-multicast-metric 26
set protocols isis interface fxp0.0 disable

```

Router R3

```

set protocols isis traceoptions file isis size 5m world-readable
set protocols isis traceoptions flag error
set protocols isis topologies ipv4-multicast
set protocols isis interface so-0/0/0 level 1 metric 19
set protocols isis interface so-0/0/0 level 1 ipv4-multicast-metric 11
set protocols isis interface so-0/0/0 level 2 metric 27
set protocols isis interface so-0/0/0 level 2 ipv4-multicast-metric 21
set protocols isis interface so-1/0/0 level 1 metric 16
set protocols isis interface so-1/0/0 level 1 ipv4-multicast-metric 26
set protocols isis interface so-1/0/0 level 2 metric 30
set protocols isis interface so-1/0/0 level 2 ipv4-multicast-metric 20
set protocols isis interface fxp0.0 disable

```

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IS-IS multicast topologies:

1. Enable the multicast topology for IS-IS by using the **ipv4-multicast** statement.

Routers R1, R2, and R3

```

[edit protocols isis]
user@host# set traceoptions file isis size 5m world-readable
user@host# set traceoptions flag error
user@host# set topologies ipv4-multicast

```

2. Enable multicast metrics on the first SONET/SDH Interface by using the **ipv4-multicast-metric** statement.

Router R1

```

[edit protocols isis interface so-0/0/0 ]
user@R1# set level 1 metric 15
user@R1# set level 1 ipv4-multicast-metric 18
user@R1# set level 2 metric 20
user@R1# set level 2 ipv4-multicast-metric 14

```

Router R2

```

[edit protocols isis interface so-0/0/0]
user@R2# set level 1 metric 13
user@R2# set level 1 ipv4-multicast-metric 12
user@R2# set level 2 metric 29
user@R2# set level 2 ipv4-multicast-metric 23

```

Router R3

```

[edit protocols isis interface so-0/0/0]

```

```
user@R3# set level 1 metric 19
user@R3# set level 1 ipv4-multicast-metric 11
user@R3# set level 2 metric 27
user@R3# set level 2 ipv4-multicast-metric 21
```

3. Enable multicast metrics on a second sonet Interface by using the **ipv4-multicast-metric** statement.

Router R1

```
[edit protocols isis interface so-1/0/0]
user@R1# set level 1 metric 13
user@R1# set level 1 ipv4-multicast-metric 12
user@R1# set level 2 metric 29
user@R1# set level 2 ipv4-multicast-metric 23
```

Router R2

```
[edit protocols isis interface so-1/0/0]
user@R2# set level 1 metric 14
user@R2# set level 1 ipv4-multicast-metric 18
user@R2# set level 2 metric 32
user@R2# set level 2 ipv4-multicast-metric 26
```

Router R3

```
[edit protocols isis interface so-1/0/0]
user@R3# set level 1 metric 16
user@R3# set level 1 ipv4-multicast-metric 26
user@R3# set level 2 metric 30
user@R3# set level 2 ipv4-multicast-metric 20
```

4. Disable the out-of-band management port, fxp0.

Routers R1, R2, and R3

```
[edit protocols isis]
user@host# set interface fxp0.0 disable
```

5. If you are done configuring the routers, commit the configuration.

Routers R1, R2, and R3

```
[edit]
user@host# commit
```

Results From configuration mode, confirm your configuration by using the **show protocols isis** statement. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Router R1

```
user@R1# show protocols isis

traceoptions {
  file isis size 5m world-readable;
  flag error;
}
topologies ipv4-multicast;
interface so-0/0/0 {
  level 1 {
```

```

        metric 15;
        ipv4-multicast-metric 18;
    }
    level 2 {
        metric 20;
        ipv4-multicast-metric 14;
    }
}
interface so-1/0/0 {
    level 1 {
        metric 13;
        ipv4-multicast-metric 12;
    }
    level 2 {
        metric 29;
        ipv4-multicast-metric 23;
    }
}
interface fxp0.0 {
    disable;
}

```

Router R2

user@R2# show protocols isis

```

traceoptions {
    file isis size 5m world-readable;
    flag error;
}
topologies ipv4-multicast;
interface so-0/0/0 {
    level 1 {
        metric 13;
        ipv4-multicast-metric 12;
    }
    level 2 {
        metric 29;
        ipv4-multicast-metric 23;
    }
}
interface so-1/0/0 {
    level 1 {
        metric 14;
        ipv4-multicast-metric 18;
    }
    level 2 {
        metric 32;
        ipv4-multicast-metric 26;
    }
}
interface fxp0.0 {
    disable;
}

```

Router R3

user@R3# show protocols isis

```

traceoptions {
    file isis size 5m world-readable;
    flag error;
}

```

```
}
topologies ipv4-multicast;
interface so-0/0/0 {
    level 1 {
        metric 19;
        ipv4-multicast-metric 11;
    }
    level 2 {
        metric 27;
        ipv4-multicast-metric 21;
    }
}
interface so-1/0/0 {
    level 1 {
        metric 16;
        ipv4-multicast-metric 26;
    }
    level 2 {
        metric 30;
        ipv4-multicast-metric 20;
    }
}
interface fxp0.0 {
    disable;
}
```

Verification

Confirm that the configuration is working properly.

- [Verifying the Connection Between Routers R1, R2, and R3 on page 4208](#)
- [Verifying That IS-IS Is Configured on page 4210](#)
- [Verifying the Configured Multicast Metric Values on page 4212](#)
- [Verifying the Configuration of the Multicast Topology on page 4213](#)

Verifying the Connection Between Routers R1, R2, and R3

Purpose Make sure that Routers R1, R2, and R3 are connected to each other.

Action Ping the other two routers from any router, to check the connectivity between the three routers as per the network topology.

```
user@R1> ping 10.0.3.9
```

```
PING 10.0.3.9 (10.0.3.9): 56 data bytes
64 bytes from 10.0.3.9: icmp_seq=0 ttl=64 time=1.299 ms
64 bytes from 10.0.3.9: icmp_seq=1 ttl=64 time=52.304 ms
64 bytes from 10.0.3.9: icmp_seq=2 ttl=64 time=1.271 ms
64 bytes from 10.0.3.9: icmp_seq=3 ttl=64 time=1.343 ms
64 bytes from 10.0.3.9: icmp_seq=4 ttl=64 time=1.434 ms
64 bytes from 10.0.3.9: icmp_seq=5 ttl=64 time=1.306 ms
^C
--- 10.0.3.9 ping statistics ---
6 packets transmitted, 6 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.271/9.826/52.304/18.997 ms
```

```
user@R1> ping 10.0.3.10
```



```

PING 10.0.3.10 (10.0.3.10): 56 data bytes
64 bytes from 10.0.3.10: icmp_seq=0 ttl=64 time=1.431 ms
64 bytes from 10.0.3.10: icmp_seq=1 ttl=64 time=1.296 ms
64 bytes from 10.0.3.10: icmp_seq=2 ttl=64 time=1.887 ms
^C
--- 10.0.3.10 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.296/1.538/1.887/0.253 ms

```

```
user@R2> ping 10.0.2.9
```

```

PING 10.0.2.9 (10.0.2.9): 56 data bytes
64 bytes from 10.0.2.9: icmp_seq=0 ttl=64 time=1.365 ms
64 bytes from 10.0.2.9: icmp_seq=1 ttl=64 time=1.813 ms
64 bytes from 10.0.2.9: icmp_seq=2 ttl=64 time=1.290 ms
^C
--- 10.0.2.9 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.290/1.489/1.813/0.231 ms

```

```
user@R2> ping 10.0.2.10
```

```

PING 10.0.2.10 (10.0.2.10): 56 data bytes
64 bytes from 10.0.2.10: icmp_seq=0 ttl=63 time=1.318 ms
64 bytes from 10.0.2.10: icmp_seq=1 ttl=63 time=1.394 ms
64 bytes from 10.0.2.10: icmp_seq=2 ttl=63 time=1.366 ms
64 bytes from 10.0.2.10: icmp_seq=3 ttl=63 time=1.305 ms
^C
--- 10.0.2.10 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.305/1.346/1.394/0.036 ms

```

```
user@R3> ping 10.0.1.10
```

```

PING 10.0.1.10 (10.0.1.10): 56 data bytes
64 bytes from 10.0.1.10: icmp_seq=0 ttl=63 time=1.316 ms
64 bytes from 10.0.1.10: icmp_seq=1 ttl=63 time=1.418 ms
64 bytes from 10.0.1.10: icmp_seq=2 ttl=63 time=1.277 ms
^C
--- 10.0.1.10 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.277/1.337/1.418/0.059 ms

```

```
user@R3> ping 10.0.1.9
```

```

PING 10.0.1.9 (10.0.1.9): 56 data bytes
64 bytes from 10.0.1.9: icmp_seq=0 ttl=64 time=1.381 ms
64 bytes from 10.0.1.9: icmp_seq=1 ttl=64 time=1.499 ms
64 bytes from 10.0.1.9: icmp_seq=2 ttl=64 time=1.300 ms
64 bytes from 10.0.1.9: icmp_seq=3 ttl=64 time=1.397 ms
^C
--- 10.0.1.9 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.300/1.394/1.499/0.071 ms

```

Meaning Routers R1, R2, and R3 have a peer relationship with each other.

Verifying That IS-IS Is Configured

Purpose Make sure that the IS-IS instance is running on Routers R1, R2, and R3, and that they are adjacent to each other.

Action Use the **show isis adjacency detail** command to check the adjacency between the routers.

Router R1

```
user@R1> show isis adjacency detail
```

R2

```
Interface: so-0/0/0, Level: 1, State: Up, Expires in 8 secs
Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:23:59 ago
Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bd
Topologies: IPV4-Multicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R2.02, IP addresses: 10.0.1.10
```

R2

```
Interface: so-0/0/0, Level: 2, State: Up, Expires in 8 secs
Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:23:58 ago
Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bd
Topologies: IPV4-Multicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R2.02, IP addresses: 10.0.1.10
```

R3

```
Interface: so-1/0/0, Level: 1, State: Up, Expires in 7 secs
Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:24:20 ago
Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bd
Topologies: IPV4-Multicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R3.02, IP addresses: 10.0.2.10
```

R3

```
Interface: so-1/0/0, Level: 2, State: Up, Expires in 6 secs
Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:24:20 ago
Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bd
Topologies: IPV4-Multicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R3.02, IP addresses: 10.0.2.10
```

Router R2

```
user@R2> show isis adjacency detail
```

R1

```
Interface: so-0/0/0, Level: 1, State: Up, Expires in 20 secs
Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:27:50 ago
Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bc
Topologies: IPV4-Multicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R2.02, IP addresses: 10.0.1.9
```

R1

```
Interface: so-0/0/0, Level: 2, State: Up, Expires in 26 secs
Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:27:50 ago
Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bc
Topologies: IPV4-Multicast
```

Restart capable: Yes, Adjacency advertisement: Advertise
 LAN id: R2.02, IP addresses: 10.0.1.9

R3

Interface: so-1/0/0, Level: 1, State: Up, Expires in 8 secs
 Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:27:22 ago
 Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bd
 Topologies: IPV4-Multicast
 Restart capable: Yes, Adjacency advertisement: Advertise
 LAN id: R3.03, IP addresses: 10.0.3.10

R3

Interface: so-1/0/0, Level: 2, State: Up, Expires in 8 secs
 Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:27:22 ago
 Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bd
 Topologies: IPV4-Multicast
 Restart capable: Yes, Adjacency advertisement: Advertise
 LAN id: R3.03, IP addresses: 10.0.3.10

Router R3

user@R3> show isis adjacency detail

R2

Interface: so-0/0/0, Level: 1, State: Up, Expires in 18 secs
 Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:33:09 ago
 Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bc
 Topologies: IPV4-Multicast
 Restart capable: Yes, Adjacency advertisement: Advertise
 LAN id: R3.03, IP addresses: 10.0.3.9

R2

Interface: so-0/0/0, Level: 2, State: Up, Expires in 22 secs
 Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:33:09 ago
 Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bc
 Topologies: IPV4-Multicast
 Restart capable: Yes, Adjacency advertisement: Advertise
 LAN id: R3.03, IP addresses: 10.0.3.9

R1

Interface: so-1/0/0, Level: 1, State: Up, Expires in 21 secs
 Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:33:59 ago
 Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bc
 Topologies: IPV4-Multicast
 Restart capable: Yes, Adjacency advertisement: Advertise
 LAN id: R3.02, IP addresses: 10.0.2.9

R1

Interface: so-1/0/0, Level: 2, State: Up, Expires in 19 secs
 Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:33:59 ago
 Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bc
 Topologies: IPV4-Multicast
 Restart capable: Yes, Adjacency advertisement: Advertise
 LAN id: R3.02, IP addresses: 10.0.2.9

Meaning IS-IS is configured on Routers R1, R2, and R3, and they are adjacent to each other.

Verifying the Configured Multicast Metric Values

Purpose Make sure that the SPF calculations are accurate as per the configured multicast metric values on Routers R1, R2, and R3.

Action Use the **show isis spf results** command to check the SPF calculations for the network.

Router R1

```
user@R1> show isis spf results
```

```
...
IPv4 Multicast IS-IS level 1 SPF results:
Node  Metric  Interface  NH  Via  SNPA
R3.03  28         so-1/0/0   IPV4 R3  0:1b:c0:86:54:bd
R2.00  18         so-0/0/0   IPV4 R2  0:1b:c0:86:54:bd
R3.00  17         so-1/0/0   IPV4 R3  0:1b:c0:86:54:bd
R1.00  0
      4 nodes
```

```
IPv4 Multicast IS-IS level 2 SPF results:
Node  Metric  Interface  NH  Via  SNPA
R3.03  40         so-0/0/0   IPV4 R2  0:1b:c0:86:54:bd
R3.00  22         so-1/0/0   IPV4 R3  0:1b:c0:86:54:bd
R2.00  14         so-0/0/0   IPV4 R2  0:1b:c0:86:54:bd
R1.00  0
      4 nodes
```

Router R2

```
user@R2> show isis spf results
```

```
...
IPv4 Multicast IS-IS level 1 SPF results:
Node  Metric  Interface  NH  Via  SNPA
R3.02  29         so-0/0/0   IPV4 R1  0:1b:c0:86:54:bc
R3.00  18         so-1/0/0   IPV4 R3  0:1b:c0:86:54:bd
R1.00  12         so-0/0/0   IPV4 R1  0:1b:c0:86:54:bc
R2.02  12
R2.00  0
      5 nodes
```

```
IPv4 Multicast IS-IS level 2 SPF results:
Node  Metric  Interface  NH  Via  SNPA
R3.02  45         so-0/0/0   IPV4 R1  0:1b:c0:86:54:bc
R3.00  26         so-1/0/0   IPV4 R3  0:1b:c0:86:54:bd
R1.00  23         so-0/0/0   IPV4 R1  0:1b:c0:86:54:bc
R2.02  23
R2.00  0
      5 nodes
```

Router R3

```
user@R3> show isis spf results
```

```
...
IPv4 Multicast IS-IS level 1 SPF results:
Node  Metric  Interface  NH  Via  SNPA
R3.02  26
R1.00  23         so-0/0/0   IPV4 R2  0:1b:c0:86:54:bc
R2.02  23         so-0/0/0   IPV4 R2  0:1b:c0:86:54:bc
R2.00  11         so-0/0/0   IPV4 R2  0:1b:c0:86:54:bc
R3.03  11
```

```

R3.00 0
      6 nodes

IPv4 Multicast IS-IS level 2 SPF results:
Node Metric Interface NH Via SNPA
R2.02 34 so-1/0/0 IPv4 R1 0:1b:c0:86:54:bc
R2.00 21 so-0/0/0 IPv4 R2 0:1b:c0:86:54:bc
R3.03 21
R1.00 20 so-1/0/0 IPv4 R1 0:1b:c0:86:54:bc
R3.02 20
R3.00 0
      6 nodes

```

Meaning The configured multicast metric values are used in SPF calculations for the IS-IS network.

Verifying the Configuration of the Multicast Topology

Purpose Make sure that the multicast topology is configured on Routers R1, R2, and R3.

Action Use the **show isis database detail** command to verify the multicast topology configuration on the routers.

Router R1

```
user@R1> show isis database detail
```

```
IS-IS level 1 link-state database:
```

```

R1.00-00 Sequence: 0x142, Checksum: 0xd07, Lifetime: 663 secs
  IPv4 Unicast IS neighbor: R2.02 Metric: 15
  IPv4 Unicast IS neighbor: R3.02 Metric: 15
  IPv4 Multicast IS neighbor: R2.02 Metric: 18
  IPv4 Multicast IS neighbor: R3.02 Metric: 17
  IP IPv4 Unicast prefix: 10.0.1.8/30 Metric: 15 Internal Up
  IP IPv4 Unicast prefix: 10.0.2.8/30 Metric: 15 Internal Up

```

```

R2.00-00 Sequence: 0x13f, Checksum: 0xf02b, Lifetime: 883 secs
  IPv4 Unicast IS neighbor: R2.02 Metric: 13
  IPv4 Unicast IS neighbor: R3.03 Metric: 14
  IPv4 Multicast IS neighbor: R2.02 Metric: 12
  IPv4 Multicast IS neighbor: R3.03 Metric: 18
  IP IPv4 Unicast prefix: 10.0.1.8/30 Metric: 13 Internal Up
  IP IPv4 Unicast prefix: 10.0.3.8/30 Metric: 14 Internal Up

```

```

R2.02-00 Sequence: 0x13c, Checksum: 0x57e2, Lifetime: 913 secs
  IPv4 Unicast IS neighbor: R1.00 Metric: 0
  IPv4 Unicast IS neighbor: R2.00 Metric: 0

```

```

R3.00-00 Sequence: 0x13c, Checksum: 0xc8de, Lifetime: 488 secs
  IPv4 Unicast IS neighbor: R3.02 Metric: 16
  IPv4 Unicast IS neighbor: R3.03 Metric: 19
  IPv4 Multicast IS neighbor: R3.02 Metric: 26
  IPv4 Multicast IS neighbor: R3.03 Metric: 11
  IP IPv4 Unicast prefix: 10.0.2.8/30 Metric: 16 Internal Up
  IP IPv4 Unicast prefix: 10.0.3.8/30 Metric: 19 Internal Up

```

```

R3.02-00 Sequence: 0x139, Checksum: 0xfb0e, Lifetime: 625 secs
  IPv4 Unicast IS neighbor: R1.00 Metric: 0
  IPv4 Unicast IS neighbor: R3.00 Metric: 0

```

```

R3.03-00 Sequence: 0x138, Checksum: 0xad56, Lifetime: 714 secs
  IPv4 Unicast IS neighbor: R2.00 Metric: 0
  IPv4 Unicast IS neighbor: R3.00 Metric: 0

```

IS-IS level 2 link-state database:

```

R1.00-00 Sequence: 0x142, Checksum: 0x2c7c, Lifetime: 816 secs
  IPv4 Unicast IS neighbor: R2.02 Metric: 20
  IPv4 Unicast IS neighbor: R3.02 Metric: 31
  IPv4 Multicast IS neighbor: R2.02 Metric: 14
  IPv4 Multicast IS neighbor: R3.02 Metric: 22
  IP IPv4 Unicast prefix: 10.0.1.8/30 Metric: 20 Internal Up
  IP IPv4 Unicast prefix: 10.0.2.8/30 Metric: 31 Internal Up
  IP IPv4 Unicast prefix: 10.0.3.8/30 Metric: 29 Internal Up

```

```

R2.00-00 Sequence: 0x13f, Checksum: 0x4826, Lifetime: 966 secs
  IPv4 Unicast IS neighbor: R2.02 Metric: 29
  IPv4 Unicast IS neighbor: R3.03 Metric: 32
  IPv4 Multicast IS neighbor: R2.02 Metric: 23
  IPv4 Multicast IS neighbor: R3.03 Metric: 26
  IP IPv4 Unicast prefix: 10.0.1.8/30 Metric: 29 Internal Up
  IP IPv4 Unicast prefix: 10.0.2.8/30 Metric: 28 Internal Up
  IP IPv4 Unicast prefix: 10.0.3.8/30 Metric: 32 Internal Up

```

```

R2.02-00 Sequence: 0x13c, Checksum: 0x57e2, Lifetime: 966 secs
  IPv4 Unicast IS neighbor: R1.00 Metric: 0
  IPv4 Unicast IS neighbor: R2.00 Metric: 0

```

```

R3.00-00 Sequence: 0x13d, Checksum: 0x1b19, Lifetime: 805 secs
  IPv4 Unicast IS neighbor: R3.02 Metric: 30
  IPv4 Unicast IS neighbor: R3.03 Metric: 27
  IPv4 Multicast IS neighbor: R3.02 Metric: 20
  IPv4 Multicast IS neighbor: R3.03 Metric: 21
  IP IPv4 Unicast prefix: 10.0.1.8/30 Metric: 31 Internal Up
  IP IPv4 Unicast prefix: 10.0.2.8/30 Metric: 30 Internal Up
  IP IPv4 Unicast prefix: 10.0.3.8/30 Metric: 27 Internal Up

```

```

R3.02-00 Sequence: 0x139, Checksum: 0xfb0e, Lifetime: 844 secs
  IPv4 Unicast IS neighbor: R1.00 Metric: 0
  IPv4 Unicast IS neighbor: R3.00 Metric: 0

```

```

R3.03-00 Sequence: 0x139, Checksum: 0xab57, Lifetime: 844 secs
  IPv4 Unicast IS neighbor: R2.00 Metric: 0
  IPv4 Unicast IS neighbor: R3.00 Metric: 0

```

Router R2

```
user@R2> show isis database detail
```

IS-IS level 1 link-state database:

```

R1.00-00 Sequence: 0x142, Checksum: 0xd07, Lifetime: 524 secs
  IPv4 Unicast IS neighbor: R2.02 Metric: 15
  IPv4 Unicast IS neighbor: R3.02 Metric: 15
  IPv4 Multicast IS neighbor: R2.02 Metric: 18
  IPv4 Multicast IS neighbor: R3.02 Metric: 17
  IP IPv4 Unicast prefix: 10.0.1.8/30 Metric: 15 Internal Up
  IP IPv4 Unicast prefix: 10.0.2.8/30 Metric: 15 Internal Up

```

```

R2.00-00 Sequence: 0x13f, Checksum: 0xf02b, Lifetime: 748 secs
  IPv4 Unicast IS neighbor: R2.02 Metric: 13

```

```

IPV4 Unicast IS neighbor: R3.03      Metric:      14
IPV4 Multicast IS neighbor: R2.02     Metric:      12
IPV4 Multicast IS neighbor: R3.03     Metric:      18
IP IPV4 Unicast prefix: 10.0.1.8/30   Metric:      13 Internal Up
IP IPV4 Unicast prefix: 10.0.3.8/30   Metric:      14 Internal Up

R2.02-00 Sequence: 0x13c, Checksum: 0x57e2, Lifetime: 777 secs
IPV4 Unicast IS neighbor: R1.00      Metric:      0
IPV4 Unicast IS neighbor: R2.00      Metric:      0

R3.00-00 Sequence: 0x13d, Checksum: 0xc6df, Lifetime: 1102 secs
IPV4 Unicast IS neighbor: R3.02      Metric:      16
IPV4 Unicast IS neighbor: R3.03      Metric:      19
IPV4 Multicast IS neighbor: R3.02     Metric:      26
IPV4 Multicast IS neighbor: R3.03     Metric:      11
IP IPV4 Unicast prefix: 10.0.2.8/30   Metric:      16 Internal Up
IP IPV4 Unicast prefix: 10.0.3.8/30   Metric:      19 Internal Up

R3.02-00 Sequence: 0x139, Checksum: 0xfb0e, Lifetime: 488 secs
IPV4 Unicast IS neighbor: R1.00      Metric:      0
IPV4 Unicast IS neighbor: R3.00      Metric:      0

R3.03-00 Sequence: 0x138, Checksum: 0xad56, Lifetime: 577 secs
IPV4 Unicast IS neighbor: R2.00      Metric:      0
IPV4 Unicast IS neighbor: R3.00      Metric:      0

IS-IS level 2 link-state database:

R1.00-00 Sequence: 0x142, Checksum: 0x2c7c, Lifetime: 676 secs
IPV4 Unicast IS neighbor: R2.02      Metric:      20
IPV4 Unicast IS neighbor: R3.02      Metric:      31
IPV4 Multicast IS neighbor: R2.02     Metric:      14
IPV4 Multicast IS neighbor: R3.02     Metric:      22
IP IPV4 Unicast prefix: 10.0.1.8/30   Metric:      20 Internal Up
IP IPV4 Unicast prefix: 10.0.2.8/30   Metric:      31 Internal Up
IP IPV4 Unicast prefix: 10.0.3.8/30   Metric:      29 Internal Up

R2.00-00 Sequence: 0x13f, Checksum: 0x4826, Lifetime: 831 secs
IPV4 Unicast IS neighbor: R2.02      Metric:      29
IPV4 Unicast IS neighbor: R3.03      Metric:      32
IPV4 Multicast IS neighbor: R2.02     Metric:      23
IPV4 Multicast IS neighbor: R3.03     Metric:      26
IP IPV4 Unicast prefix: 10.0.1.8/30   Metric:      29 Internal Up
IP IPV4 Unicast prefix: 10.0.2.8/30   Metric:      28 Internal Up
IP IPV4 Unicast prefix: 10.0.3.8/30   Metric:      32 Internal Up

R2.02-00 Sequence: 0x13c, Checksum: 0x57e2, Lifetime: 831 secs
IPV4 Unicast IS neighbor: R1.00      Metric:      0
IPV4 Unicast IS neighbor: R2.00      Metric:      0

R3.00-00 Sequence: 0x13d, Checksum: 0x1b19, Lifetime: 667 secs
IPV4 Unicast IS neighbor: R3.02      Metric:      30
IPV4 Unicast IS neighbor: R3.03      Metric:      27
IPV4 Multicast IS neighbor: R3.02     Metric:      20
IPV4 Multicast IS neighbor: R3.03     Metric:      21
IP IPV4 Unicast prefix: 10.0.1.8/30   Metric:      31 Internal Up
IP IPV4 Unicast prefix: 10.0.2.8/30   Metric:      30 Internal Up
IP IPV4 Unicast prefix: 10.0.3.8/30   Metric:      27 Internal Up

R3.02-00 Sequence: 0x139, Checksum: 0xfb0e, Lifetime: 707 secs
IPV4 Unicast IS neighbor: R1.00      Metric:      0

```

IPv4 Unicast IS neighbor: R3.00 Metric: 0

R3.03-00 Sequence: 0x139, Checksum: 0xab57, Lifetime: 707 secs

IPv4 Unicast IS neighbor: R2.00 Metric: 0

IPv4 Unicast IS neighbor: R3.00 Metric: 0

Router R3

user@R3> show isis database detail

IS-IS level 1 link-state database:

R1.00-00 Sequence: 0x143, Checksum: 0xb08, Lifetime: 1155 secs

IPv4 Unicast IS neighbor: R2.02 Metric: 15

IPv4 Unicast IS neighbor: R3.02 Metric: 15

IPv4 Multicast IS neighbor: R2.02 Metric: 18

IPv4 Multicast IS neighbor: R3.02 Metric: 17

IP IPv4 Unicast prefix: 10.0.1.8/30 Metric: 15 Internal Up

IP IPv4 Unicast prefix: 10.0.2.8/30 Metric: 15 Internal Up

R2.00-00 Sequence: 0x13f, Checksum: 0xf02b, Lifetime: 687 secs

IPv4 Unicast IS neighbor: R2.02 Metric: 13

IPv4 Unicast IS neighbor: R3.03 Metric: 14

IPv4 Multicast IS neighbor: R2.02 Metric: 12

IPv4 Multicast IS neighbor: R3.03 Metric: 18

IP IPv4 Unicast prefix: 10.0.1.8/30 Metric: 13 Internal Up

IP IPv4 Unicast prefix: 10.0.3.8/30 Metric: 14 Internal Up

R2.02-00 Sequence: 0x13c, Checksum: 0x57e2, Lifetime: 716 secs

IPv4 Unicast IS neighbor: R1.00 Metric: 0

IPv4 Unicast IS neighbor: R2.00 Metric: 0

R3.00-00 Sequence: 0x13d, Checksum: 0xc6df, Lifetime: 1044 secs

IPv4 Unicast IS neighbor: R3.02 Metric: 16

IPv4 Unicast IS neighbor: R3.03 Metric: 19

IPv4 Multicast IS neighbor: R3.02 Metric: 26

IPv4 Multicast IS neighbor: R3.03 Metric: 11

IP IPv4 Unicast prefix: 10.0.2.8/30 Metric: 16 Internal Up

IP IPv4 Unicast prefix: 10.0.3.8/30 Metric: 19 Internal Up

R3.02-00 Sequence: 0x139, Checksum: 0xfb0e, Lifetime: 430 secs

IPv4 Unicast IS neighbor: R1.00 Metric: 0

IPv4 Unicast IS neighbor: R3.00 Metric: 0

R3.03-00 Sequence: 0x138, Checksum: 0xad56, Lifetime: 519 secs

IPv4 Unicast IS neighbor: R2.00 Metric: 0

IPv4 Unicast IS neighbor: R3.00 Metric: 0

IS-IS level 2 link-state database:

R1.00-00 Sequence: 0x142, Checksum: 0x2c7c, Lifetime: 617 secs

IPv4 Unicast IS neighbor: R2.02 Metric: 20

IPv4 Unicast IS neighbor: R3.02 Metric: 31

IPv4 Multicast IS neighbor: R2.02 Metric: 14

IPv4 Multicast IS neighbor: R3.02 Metric: 22

IP IPv4 Unicast prefix: 10.0.1.8/30 Metric: 20 Internal Up

IP IPv4 Unicast prefix: 10.0.2.8/30 Metric: 31 Internal Up

IP IPv4 Unicast prefix: 10.0.3.8/30 Metric: 29 Internal Up

R2.00-00 Sequence: 0x13f, Checksum: 0x4826, Lifetime: 769 secs

IPv4 Unicast IS neighbor: R2.02 Metric: 29

IPv4 Unicast IS neighbor: R3.03 Metric: 32


```

IPV4 Multicast IS neighbor: R2.02  Metric:      23
IPV4 Multicast IS neighbor: R3.03  Metric:      26
IP  IPV4 Unicast prefix: 10.0.1.8/30 Metric:      29 Internal Up
IP  IPV4 Unicast prefix: 10.0.2.8/30 Metric:      28 Internal Up
IP  IPV4 Unicast prefix: 10.0.3.8/30 Metric:      32 Internal Up

R2.02-00 Sequence: 0x13c, Checksum: 0x57e2, Lifetime: 769 secs
IPV4 Unicast IS neighbor: R1.00    Metric:      0
IPV4 Unicast IS neighbor: R2.00    Metric:      0

R3.00-00 Sequence: 0x13d, Checksum: 0x1b19, Lifetime: 610 secs
IPV4 Unicast IS neighbor: R3.02    Metric:      30
IPV4 Unicast IS neighbor: R3.03    Metric:      27
IPV4 Multicast IS neighbor: R3.02   Metric:      20
IPV4 Multicast IS neighbor: R3.03   Metric:      21
IP  IPV4 Unicast prefix: 10.0.1.8/30 Metric:      31 Internal Up
IP  IPV4 Unicast prefix: 10.0.2.8/30 Metric:      30 Internal Up
IP  IPV4 Unicast prefix: 10.0.3.8/30 Metric:      27 Internal Up

R3.02-00 Sequence: 0x139, Checksum: 0xfb0e, Lifetime: 649 secs
IPV4 Unicast IS neighbor: R1.00    Metric:      0
IPV4 Unicast IS neighbor: R3.00    Metric:      0

R3.03-00 Sequence: 0x139, Checksum: 0xab57, Lifetime: 649 secs
IPV4 Unicast IS neighbor: R2.00    Metric:      0
IPV4 Unicast IS neighbor: R3.00    Metric:      0

```

Meaning Multicast topology is configured on Routers R1, R2, and R3.

- Related Documentation**
- [Understanding Multitopology Routing in Conjunction with PIM](#)
 - [Example: Configuring Multitopology Routing to Provide Redundancy for Multicast Traffic over Separate Network Paths](#)
 - [Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies on page 4193](#)

Example: Configuring Link and Node Protection for IS-IS Routes

- [Understanding Loop-Free Alternate Routes for IS-IS on page 4218](#)
- [Example: Configuring Node-Link Protection for IS-IS Routes in a Layer 3 VPN on page 4221](#)

Understanding Loop-Free Alternate Routes for IS-IS

In Junos OS Release 9.5 and later, support for IS-IS loop-free alternate routes enables IP fast-reroute capability for IS-IS. Junos OS precomputes loop-free backup routes for all IS-IS routes. These backup routes are preinstalled in the Packet Forwarding Engine, which performs a local repair and implements the backup path when the link for a primary next hop for a particular route is no longer available. With local repair, the Packet Forwarding Engine can correct a path failure before it receives recomputed paths from the Routing Engine. Local repair reduces the amount of time needed to reroute traffic to less than 50 milliseconds. In contrast, global repair can take up to 800 milliseconds to compute a new route. Local repair and global repair are thus complementary. Local repair enables traffic to continue to be routed using a backup path until global repair is able to calculate a new route.

A loop-free path is one that does not forward traffic back through the routing device to reach a given destination. That is, a neighbor whose shortest path to the destination traverses the routing device is not used as a backup route to that destination. To determine loop-free alternate paths for IS-IS routes, Junos OS runs shortest-path-first (SPF) calculations on each one-hop neighbor. You can enable support for alternate loop-free routes on any IS-IS interface. Because it is common practice to enable LDP on an interface for which IS-IS is already enabled, this feature also provides support for LDP label-switched paths (LSPs).



NOTE: If you enable support for alternate loop-free routes on an interface configured for both LDP and IS-IS, you can use the `traceroute` command to trace the active path to the primary next hop.

The level of backup coverage available through IS-IS routes depends on the actual network topology and is typically less than 100 percent for all destinations on any given routing device. You can extend backup coverage to include RSVP LSPs.


Junos OS provides two mechanisms for route redundancy for IS-IS through alternate loop-free routes: link protection and node-link protection. When you enable link protection or node-link protection on an IS-IS interface, Junos OS creates a single alternate path to the primary next hop for all destination routes that traverse a protected interface. Link protection offers per-link traffic protection. Use link protection when you assume that only a single link might become unavailable but that the neighboring node on the primary path would still be available through another interface.

Node-link protection establishes an alternate path through a different routing device altogether. Use node-link protection when you assume that access to a node is lost when a link is no longer available. As a result, Junos OS calculates a backup path that avoids the primary next-hop routing device. In Junos OS Release 9.4 and earlier, only the RSVP protocol supports Packet Forwarding Engine local repair and fast reroute as well as link protection and node protection.

In [Figure 95](#), Case 2 shows how link protection allows source Router A to switch to Link B when the primary next hop Link A to destination Router C fails. However, if Router B

fails, Link B also fails, and the protected Link A is lost. If node-link protection is enabled, Router A is able to switch to Link D on Router D and bypass the failed Router B altogether. As shown in Case 1, with node-link protection enabled, Router A has a node-link protection alternate path available through Router D to destination Router C. That means that if Router B fails, Router A can still reach Router C because the path from Router A to Link D remains available as an alternate backup path.

Figure 101: Link Protection and Node-Link Protection Comparison for IS-IS Routes



The Junos OS implementation of support for loop-free alternate paths for IS-IS routes is based on the following standards:

- RFC 5286, *Basic Specification for IP Fast-Reroute: Loop-free Alternates*
- RFC 5714, *IP Fast Reroute Framework*

Configuring Link Protection for IS-IS

You can configure link protection on any interface for which IS-IS is enabled. When you enable link protection, Junos OS creates one alternate path to the primary next hop for all destination routes that traverse a protected interface. Link protection assumes that only a single link becomes unavailable but that the neighboring node would still be available through another interface.



NOTE: You must also configure a per-packet load-balancing routing policy to ensure that the routing protocol process installs all the next hops for a given route in the routing table.

To enable link protection, include the **link-protection** statement at the **[edit protocols isis interface *interface-name*]** hierarchy level:

```
[edit]
protocols {
  isis {
    interface interface-name {
      link-protection;
    }
  }
}
```

Configuring Node-Link Protection for IS-IS

You can configure node-link protection on any interface for which IS-IS is enabled. Node-link protection establishes an alternate path through a different routing device altogether for all destination routes that traverse a protected interface. Node-link protection assumes that the entire routing device, or node, has failed. Junos OS therefore calculates a backup path that avoids the primary next-hop routing device.



NOTE: You must also configure a per-packet load-balancing routing policy to ensure that the routing protocol process installs all the next hops for a given route in the routing table.

To enable node-link protection, include the **node-link-protection** statement at the **[edit protocols isis interface *interface-name*]** hierarchy level:

```
[edit]
protocols {
  isis {
    interface interface-name {
      node-link-protection;
    }
  }
}
```

Excluding an IS-IS Interface as a Backup for Protected Interfaces

By default, all IS-IS interfaces that belong to the master instance or a specific routing instance are eligible as backup interfaces for protected interfaces. You can specify that any IS-IS interface be excluded from functioning as a backup interface to protected interfaces. To exclude an IS-IS interface as a backup interface, include the **no-eligible-backup** statement at the **[edit protocols isis interface *interface-name*]** hierarchy level:

```
[edit]
protocols {
  isis {
    interface interface-name {
      no-eligible-backup;
    }
  }
}
```

Configuring RSVP Label-Switched Paths as Backup Paths for IS-IS

Relying on the shortest-path-first (SPF) calculation of backup paths for one-hop neighbors might result in less than 100 percent backup coverage for a specific network topology. You can enhance coverage of IS-IS and LDP label-switched paths (LSPs) by configuring RSVP LSPs as backup paths. To configure a specific RSVP LSP as a backup path, include the **backup** statement at the `[edit protocols mpls label-switched-path lsp-name]` hierarchy level:

```
[edit]
protocols {
  mpls {
    label-switched-path lsp-name {
      backup;
      to ip-address;
    }
  }
}
```

When configuring an LSP, you must specify the IP address of the egress routing device with the **to** statement. For detailed information about configuring LSPs and RSVP, see the *Junos OS MPLS Applications Library for Routing Devices*.

Using Operational Mode Commands to Monitor Protected IS-IS Routes

You can issue operational mode commands that provide more details about your link-protected and node-link-protected IS-IS routes. The following guidelines explain the type of information available from the output of each command:

- **show isis backup label-switched-path**—Displays which MPLS LSPs have been designated as backup paths and the current status of those LSPs.
- **show isis backup spf results**—Displays SPF calculations for each neighbor for a given destination. Indicates whether a specific interface or node has been designated as a backup path and why. Use the **no-coverage** option to display only those nodes that do not have backup coverage.
- **show isis backup coverage**—Displays the percentage of nodes and prefixes for each type of address family that is protected.
- **show isis interface detail**—Displays the type of protection (link or node-link) applied to each protected interface.

Example: Configuring Node-Link Protection for IS-IS Routes in a Layer 3 VPN

Node-link protection establishes an alternate path through a different routing device. Use node-link protection when you assume that access to a node is lost when a link is no longer available. Junos OS calculates a backup path that avoids the primary next-hop routing device.

- [Requirements on page 4222](#)
- [Overview on page 4222](#)

- [Configuration on page 4222](#)
- [Verification on page 4229](#)

Requirements

This example requires Junos OS Release 9.5 or later.

No special configuration beyond device initialization is required before configuring this example.


Overview

In this example, core-facing interfaces are enabled for IS-IS Level 2, LDP, and RSVP. Node-link protection is enabled on all the core-facing interfaces, which means that if the primary next hop for any destination that traverses the interfaces becomes unavailable, Junos OS uses a backup link that avoids the next-hop router altogether if necessary.

You also need to configure a routing policy that requires all traffic to use per-packet load balancing in order to enable Packet Forwarding Engine local repair. With local repair, the Packet Forwarding Engine can correct a path failure and implement a backup loop-free alternate route before it receives recomputed paths from the Routing Engine.

[Figure 102](#) shows the topology used in this example.

Figure 102: IS-IS Node-Link Protection Topology



On Device PE1, an RSVP LSP is configured as a backup path for IS-IS. Relying on the shortest-path-first (SPF) calculation of backup paths for one-hop neighbors might result in less than 100 percent backup coverage for a specific network topology. You can enhance coverage of IS-IS and LDP LSPs by configuring RSVP LSPs as backup paths. To configure a specific RSVP LSP as a backup path, include the **backup** statement at the **[edit protocols mpls label-switched-path *lsp-name*]** hierarchy level.

“[CLI Quick Configuration](#)” on [page 4222](#) shows the configuration for all of the devices in [Figure 102](#). The section “[Step-by-Step Procedure](#)” on [page 4226](#) describes the steps on Device P1.

Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network

configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device CE1	set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.1/30 set interfaces lo0 unit 0 family inet address 10.255.1.1/32
Device PE1	set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.2/30 set interfaces fe-1/2/0 unit 0 family iso set interfaces fe-1/2/0 unit 0 family mpls set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.5/30 set interfaces fe-1/2/1 unit 0 family iso set interfaces fe-1/2/1 unit 0 family mpls set interfaces fe-1/2/2 unit 0 family inet address 10.0.0.21/30 set interfaces fe-1/2/2 unit 0 family iso set interfaces fe-1/2/2 unit 0 family mpls set interfaces lo0 unit 0 family inet address 10.255.2.2/32 set interfaces lo0 unit 0 family iso address 49.0001.0010.0000.0202.00 set protocols rsvp interface fe-1/2/2.0 set protocols rsvp interface fe-1/2/1.0 set protocols rsvp interface lo0.0 set protocols rsvp interface fxp0.0 disable set protocols mpls label-switched-path to-p2 backup set protocols mpls label-switched-path to-p2 to 10.255.4.4 set protocols mpls label-switched-path to-p2 ldp-tunneling set protocols mpls interface fe-1/2/2.0 set protocols mpls interface fe-1/2/1.0 set protocols mpls interface lo0.0 set protocols mpls interface fxp0.0 disable set protocols bgp group l3vpn type internal set protocols bgp group l3vpn local-address 10.255.2.2 set protocols bgp group l3vpn family inet-vpn unicast set protocols bgp group l3vpn peer-as 65534 set protocols bgp group l3vpn local-as 65534 set protocols bgp group l3vpn neighbor 10.255.5.5 set protocols isis spf-options delay 1000 set protocols isis interface all node-link-protection set protocols isis interface all level 2 metric 10 set protocols isis interface all level 1 disable set protocols isis interface fxp0.0 disable set protocols isis interface lo0.0 level 2 metric 0 set protocols ldp deaggregate set protocols ldp interface fe-1/2/1.0 set protocols ldp interface fe-1/2/2.0 set protocols ldp interface fxp0.0 disable set protocols ldp interface lo0.0 set policy-options policy-statement ecmp term 1 then load-balance per-packet set routing-instances VPN-A instance-type vrf set routing-instances VPN-A interface fe-1/2/0.0 set routing-instances VPN-A route-distinguisher 65534:1234 set routing-instances VPN-A vrf-target target:65534:1234 set routing-instances VPN-A routing-options static route 10.255.1.1/32 next-hop 10.0.0.1 set routing-options autonomous-system 65534 set routing-options forwarding-table export ecmp
Device P1	set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.6/30

```
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.9/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces fe-1/2/1 unit 0 family mpls
set interfaces fe-1/2/2 unit 0 family inet address 10.0.0.25/30
set interfaces fe-1/2/2 unit 0 family iso
set interfaces fe-1/2/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.3.3/32
set interfaces lo0 unit 0 family iso address 49.0001.0010.0000.0303.00
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols isis spf-options delay 1000
set protocols isis interface all node-link-protection
set protocols isis interface all level 2 metric 10
set protocols isis interface all level 1 disable
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0 level 2 metric 0
set protocols ldp deaggregate
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set policy-options policy-statement ecmp term 1 then load-balance per-packet
set routing-options forwarding-table export ecmp
```

Device P2

```
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.10/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.13/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces fe-1/2/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.4.4/32
set interfaces lo0 unit 0 family iso address 49.0001.0010.0000.0404.00
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols isis spf-options delay 1000
set protocols isis interface all node-link-protection
set protocols isis interface all level 2 metric 10
set protocols isis interface all level 1 disable
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0 level 2 metric 0
set protocols ldp deaggregate
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set policy-options policy-statement ecmp term 1 then load-balance per-packet
set routing-options forwarding-table export ecmp
```

Device P3

```
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.22/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.26/30
set interfaces fe-1/2/1 unit 0 family iso
```



```

set interfaces fe-1/2/1 unit 0 family mpls
set interfaces fe-1/2/2 unit 0 family inet address 10.0.0.30/30
set interfaces fe-1/2/2 unit 0 family iso
set interfaces fe-1/2/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.7.7/32
set interfaces lo0 unit 0 family iso address 49.0001.0010.0000.0707.00
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols isis spf-options delay 1000
set protocols isis interface all node-link-protection
set protocols isis interface all level 2 metric 10
set protocols isis interface all level 1 disable
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0 level 2 metric 0
set protocols ldp deaggregate
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set policy-options policy-statement ecmp term 1 then load-balance per-packet
set routing-options forwarding-table export ecmp

```

Device PE2

```

set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.14/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.17/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces fe-1/2/2 unit 0 family inet address 10.0.0.29/30
set interfaces fe-1/2/2 unit 0 family iso
set interfaces fe-1/2/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.5.5/32
set interfaces lo0 unit 0 family iso address 49.0001.0010.0000.0505.00
set protocols rsvp interface fe-1/2/0.0
set protocols rsvp interface fe-1/2/2.0
set protocols rsvp interface lo0.0
set protocols rsvp interface fxp0.0 disable
set protocols mpls interface fe-1/2/0.0
set protocols mpls interface fe-1/2/2.0
set protocols mpls interface lo0.0
set protocols mpls interface fxp0.0 disable
set protocols bgp group l3vpn type internal
set protocols bgp group l3vpn local-address 10.255.5.5
set protocols bgp group l3vpn family inet-vpn unicast
set protocols bgp group l3vpn peer-as 65534
set protocols bgp group l3vpn local-as 65534
set protocols bgp group l3vpn neighbor 10.255.2.2
set protocols isis spf-options delay 1000
set protocols isis interface all node-link-protection
set protocols isis interface all level 2 metric 10
set protocols isis interface all level 1 disable
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0 level 2 metric 0
set protocols ldp deaggregate
set protocols ldp interface fe-1/2/0.0
set protocols ldp interface fe-1/2/2.0
set protocols ldp interface fxp0.0 disable

```

```
set protocols ldp interface lo0.0
set policy-options policy-statement ecmp term 1 then load-balance per-packet
set routing-instances VPN-A instance-type vrf
set routing-instances VPN-A interface fe-1/2/1.0
set routing-instances VPN-A route-distinguisher 65534:1234
set routing-instances VPN-A vrf-target target:65534:1234
set routing-instances VPN-A routing-options static route 10.255.1.1/32 next-hop 10.0.0.18
set routing-options autonomous-system 65534
set routing-options forwarding-table export ecmp
```

Device CE2 `set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.18/30`
 `set interfaces lo0 unit 0 family inet address 10.255.6.6/32`

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure multi-level IS-IS:

1. Configure the interfaces.

Enable IS-IS and MPLS.

```
[edit interfaces]
user@P1# set fe-1/2/0 unit 0 family inet address 10.0.0.6/30
user@P1# set fe-1/2/0 unit 0 family iso
user@P1# set fe-1/2/0 unit 0 family mpls
user@P1# set fe-1/2/1 unit 0 family inet address 10.0.0.9/30
user@P1# set fe-1/2/1 unit 0 family iso
user@P1# set fe-1/2/1 unit 0 family mpls
user@P1# set fe-1/2/2 unit 0 family inet address 10.0.0.25/30
user@P1# set fe-1/2/2 unit 0 family iso
user@P1# set fe-1/2/2 unit 0 family mpls
user@P1# set lo0 unit 0 family inet address 10.255.3.3/32
user@P1# set lo0 unit 0 family iso address 49.0001.0010.0000.0303.00
```

2. Configure the IS-IS interfaces for Level 2.

```
[edit protocols]
user@P1# set isis interface all level 2 metric 10
user@P1# set isis interface all level 1 disable
user@P1# set isis interface fxp0.0 disable
user@P1# set isis interface lo0.0 level 2 metric 0
```

3. Enable IS-IS node-link protection, which also automatically extends backup coverage to all LDP LSPs.

```
[edit protocols]
user@P1# set isis interface all node-link-protection
```

4. (Optional) Configure a 1000-millisecond time interval between the detection of a topology change and when the SPF algorithm runs.

```
[edit protocols]
user@P1# set isis spf-options delay 1000
```

5. Configure MPLS to use both RSVP and LDP label-switched paths (LSPs).

```
[edit protocols]
```

```

user@P1# set mpls interface all
user@P1# set mpls interface fxp0.0 disable
user@P1# set rsvp interface all
user@P1# set rsvp interface fxp0.0 disable
user@P1# set ldp interface all
user@P1# set ldp interface fxp0.0 disable

```

6. (Optional) For LDP, enable forwarding equivalence class (FEC) deaggregation, which results in faster global convergence.

```

[edit protocols]
user@P1# set ldp deaggregate

```

7. To enable Packet Forwarding Engine local repair, establish a policy that forces the routing protocol process to install all the next hops for a given route.

This policy ensures that the backup route is installed in the forwarding table used by the Packet Forwarding Engine to forward traffic to a given destination.

```

[edit policy-options policy-statement ecmp term 1]
user@P1# set then load-balance per-packet

```

8. Apply the policy to the forwarding table of the local router with the **export** statement.

```

[edit routing-options forwarding-table]
user@P1# set export ecmp

```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@P1# show interfaces
fe-1/2/0 {
  unit 0 {
    family inet {
      address 10.0.0.6/30;
    }
    family iso;
    family mpls;
  }
}
fe-1/2/1 {
  unit 0 {
    family inet {
      address 10.0.0.9/30;
    }
    family iso;
    family mpls;
  }
}
fe-1/2/2 {
  unit 0 {
    family inet {
      address 10.0.0.25/30;
    }
    family iso;
  }
}

```

```
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.255.3.3/32;
        }
        family iso {
            address 49.0001.0010.0000.0303.00;
        }
    }
}

user@P1# show protocols
rsvp {
    interface all;
    interface fxp0.0 {
        disable;
    }
}
mpls {
    interface all;
    interface fxp0.0 {
        disable;
    }
}
isis {
    spf-options delay 1000;
    interface all {
        node-link-protection;
        level 2 metric 10;
        level 1 disable;
    }
    interface fxp0.0 {
        disable;
    }
    interface lo0.0 {
        level 2 metric 0;
    }
}
ldp {
    deaggregate;
    interface all;
    interface fxp0.0 {
        disable;
    }
}

user@P1# show policy-options
policy-statement ecmp {
    term 1 {
        then {
            load-balance per-packet;
        }
    }
}
```

```

user@P1# show routing-options
forwarding-table {
    export ecmp;
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Checking the MPLS LSP Backup Path on page 4229](#)
- [Checking Which Next-Hop Neighbors Are Designated as Backup Paths to the Destination Node on page 4229](#)
- [Checking the Backup Coverage on page 4230](#)
- [Checking the Type of Protection Configured on page 4231](#)

Checking the MPLS LSP Backup Path

Purpose	Display information about the MPLS label-switched-paths (LSPs) designated as the backup route for the IS-IS routes.
Action	On Device PE1, from operational mode, enter the show isis backup label-switched-path command. <pre> user@PE1> show isis backup label-switched-path Backup MPLS LSPs: to-p2, Egress: 10.255.4.4, Status: up, Last change: 01:17:45 TE-metric: 19, Metric: 0, Refcount: 1 </pre>
Meaning	The output shows that the backup path is up and operational.

Checking Which Next-Hop Neighbors Are Designated as Backup Paths to the Destination Node

Purpose	Display SPF calculations for each neighbor for a given destination.
Action	On Device PE1, from operational mode, enter the show isis backup spf results command. <pre> user@PE1> show isis backup spf results IS-IS level 1 SPF results: 0 nodes IS-IS level 2 SPF results: PE2.00 Primary next-hop: fe-1/2/2.0, IPV4, P3, SNPA: 0:5:85:8f:c8:bd Root: P2, Root Metric: 20, Metric: 10, Root Preference: 0x0 track-item: P2.00-00 Eligible, Backup next-hop: fe-1/2/1.0, LSP, to-p2 Root: P3, Root Metric: 10, Metric: 10, Root Preference: 0x0 Not eligible, Reason: Interface is already covered Root: P1, Root Metric: 10, Metric: 20, Root Preference: 0x0 track-item: P3.00-00 Not eligible, Reason: Interface is already covered P2.00 </pre>

```

Primary next-hop: fe-1/2/1.0, IPV4, P1, SNPA: 0:5:85:8f:c8:bd
Root: P2, Root Metric: 20, Metric: 0, Root Preference: 0x0
track-item: P2.00-00
Not eligible, Reason: Primary next-hop link fate sharing
Root: P1, Root Metric: 10, Metric: 10, Root Preference: 0x0
Not eligible, Reason: Primary next-hop link fate sharing
Root: P3, Root Metric: 10, Metric: 20, Root Preference: 0x0
track-item: P1.00-00
Not eligible, Reason: Primary next-hop node fate sharing
P3.00
Primary next-hop: fe-1/2/2.0, IPV4, P3, SNPA: 0:5:85:8f:c8:bd
Root: P2, Root Metric: 20, Metric: 20, Root Preference: 0x0
track-item: P3.00-00
track-item: P2.00-00
track-item: P1.00-00
Eligible, Backup next-hop: fe-1/2/1.0, LSP, to-p2
Root: P3, Root Metric: 10, Metric: 0, Root Preference: 0x0
Not eligible, Reason: Interface is already covered
Root: P1, Root Metric: 10, Metric: 10, Root Preference: 0x0
track-item: P3.00-00
Not eligible, Reason: Interface is already covered
P1.00
Primary next-hop: fe-1/2/1.0, IPV4, P1, SNPA: 0:5:85:8f:c8:bd
Root: P2, Root Metric: 20, Metric: 10, Root Preference: 0x0
track-item: P2.00-00
track-item: P1.00-00
Not eligible, Reason: Primary next-hop link fate sharing
Root: P1, Root Metric: 10, Metric: 0, Root Preference: 0x0
Not eligible, Reason: Primary next-hop link fate sharing
Root: P3, Root Metric: 10, Metric: 10, Root Preference: 0x0
track-item: P1.00-00
Eligible, Backup next-hop: fe-1/2/2.0, IPV4, P3, SNPA: 0:5:85:8f:c8:bd
4 nodes

```

Meaning The output indicates whether a specific interface or node has been designated as a backup path and why.

Checking the Backup Coverage

Purpose Check the percentage of protected nodes and prefixes.

Action From operational mode, enter the **show isis backup coverage** command.

```
user@PE1> show isis backup coverage
```

```
Backup Coverage:
```

Topology	Level	Node	IPv4	IPv6	CLNS
IPV4 Unicast	1	0.00%	0.00%	0.00%	0.00%
IPV4 Unicast	2	75.00%	87.50%	0.00%	0.00%

```
user@P1> show isis backup coverage
```

```
Backup Coverage:
```

Topology	Level	Node	IPv4	IPv6	CLNS
IPV4 Unicast	1	0.00%	0.00%	0.00%	0.00%
IPV4 Unicast	2	75.00%	71.43%	0.00%	0.00%

```
user@P2> show isis backup coverage
```

```
Backup Coverage:
```

Topology	Level	Node	IPv4	IPv6	CLNS
----------	-------	------	------	------	------

```

IPv4 Unicast      1  0.00%  0.00%  0.00%  0.00%
IPv4 Unicast      2  50.00%  37.50%  0.00%  0.00%

```

```
user@P3> show isis backup coverage
```

```
Backup Coverage:
```

```

Topology      Level1  Node   IPv4   IPv6   CLNS
IPv4 Unicast   1    0.00%  0.00%  0.00%  0.00%
IPv4 Unicast   2    75.00%  71.43%  0.00%  0.00%

```

```
user@PE2> show isis backup coverage
```

```
Backup Coverage:
```

```

Topology      Level1  Node   IPv4   IPv6   CLNS
IPv4 Unicast   1    0.00%  0.00%  0.00%  0.00%
IPv4 Unicast   2    50.00%  37.50%  0.00%  0.00%

```

Meaning The level of backup coverage available through IS-IS routes depends on the actual network topology and is typically less than 100 percent for all destinations on any given routing device. You can extend backup coverage to include RSVP LSPs.

Checking the Type of Protection Configured

Purpose On all nodes in the IS-IS domain, check the type and percentage of protected nodes and prefixes.

Action From operational mode, enter the **show isis interface detail** command.

```
user@PE1> show isis interface detail
```

```
IS-IS interface database:
```

```
lo0.0
```

```
Index: 76, State: 0x6, Circuit id: 0x1, Circuit type: 0
```

```
LSP interval: 100 ms, CSNP interval: disabled
```

```
Adjacency advertisement: Advertise
```

```
Level Adjacencies Priority Metric Hello (s) Hold (s) Designated Router
```

```
1          0          64      0 Passive
```

```
2          0          64      0 Passive
```

```
fe-1/2/2.0
```

```
Index: 79, State: 0x6, Circuit id: 0x1, Circuit type: 2
```

```
LSP interval: 100 ms, CSNP interval: 10 s
```

```
Adjacency advertisement: Advertise
```

```
Protection Type: Node Link
```

```
Level Adjacencies Priority Metric Hello (s) Hold (s) Designated Router
```

```
2          1          64      10     9.000     27 P3.03 (not us)
```

```
fe-1/2/1.0
```

```
Index: 77, State: 0x6, Circuit id: 0x1, Circuit type: 2
```

```
LSP interval: 100 ms, CSNP interval: 10 s
```

```
Adjacency advertisement: Advertise
```

```
Protection Type: Node Link
```

```
Level Adjacencies Priority Metric Hello (s) Hold (s) Designated Router
```

```
2          1          64      10     9.000     27 P1.02 (not us)
```

Meaning The output shows that node-link protection is configured on the interfaces.

Related Documentation

- *Example: Configuring BFD for IS-IS*

Example: Configuring an IS-IS Default Route Policy on Logical Systems

This example shows logical systems configured on a single physical router and explains how to configure a default route on one logical system.

- [Requirements on page 4232](#)
- [Overview on page 4232](#)
- [Configuration on page 4233](#)
- [Verification on page 4235](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.


Overview

This example shows a logical system redistributing a default route to other logical systems. All logical systems are running IS-IS. A common reason for a default route is to provide a path for sending traffic destined outside the IS-IS domain.

In this example, the default route is not used for forwarding traffic. The **no-install** statement prevents the route from being installed in the forwarding table of Logical System LS3. If you configure a route so it is not installed in the forwarding table, the route is still eligible to be exported from the routing table to other protocols. The **discard** statement silently drops packets without notice.

Figure 103 shows the sample network.

Figure 103: IS-IS with a Default Route to an ISP



90-40918

Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter commit from configuration mode.

```
set logical-systems LS3 interfaces lt-1/2/0 unit 3 description LS3->LS2
set logical-systems LS3 interfaces lt-1/2/0 unit 3 encapsulation ethernet
set logical-systems LS3 interfaces lt-1/2/0 unit 3 peer-unit 4
set logical-systems LS3 interfaces lt-1/2/0 unit 3 family inet address 10.0.2.1/30
set logical-systems LS3 interfaces lt-1/2/0 unit 3 family iso
set logical-systems LS3 interfaces lt-1/2/0 unit 5 description LS3->LS1
set logical-systems LS3 interfaces lt-1/2/0 unit 5 encapsulation ethernet
set logical-systems LS3 interfaces lt-1/2/0 unit 5 peer-unit 0
set logical-systems LS3 interfaces lt-1/2/0 unit 5 family inet address 10.0.1.1/30
set logical-systems LS3 interfaces lt-1/2/0 unit 5 family iso
set logical-systems LS3 interfaces lo0 unit 3 family iso address 49.0001.1234.1600.2231.00
set logical-systems LS3 protocols isis export isis-default
set logical-systems LS3 protocols isis interface lt-1/2/0.3
set logical-systems LS3 protocols isis interface lt-1/2/0.5
set logical-systems LS3 protocols isis interface lo0.3 passive
set logical-systems LS3 routing-options static route 0.0.0.0/0 discard
set logical-systems LS3 routing-options static route 0.0.0.0/0 no-install
set logical-systems LS3 policy-options policy-statement isis-default from protocol static
set logical-systems LS3 policy-options policy-statement isis-default from route-filter
  0.0.0.0/0 exact
set logical-systems LS3 policy-options policy-statement isis-default then accept
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure an IS-IS default route policy on logical systems:

1. Configure the logical tunnel interfaces.

```
[edit logical-systems LS3 interfaces lt-1/2/0]
user@R1# set unit 3 description LS3->LS2
user@R1# set unit 3 encapsulation ethernet
user@R1# set unit 3 peer-unit 4
user@R1# set unit 3 family inet address 10.0.2.1/30
user@R1# set unit 3 family iso
user@R1# set unit 5 description LS3->LS1
user@R1# set unit 5 encapsulation ethernet
user@R1# set unit 5 peer-unit 0
user@R1# set unit 5 family inet address 10.0.1.1/30
user@R1# set unit 5 family iso
[edit logical-systems LS3 interfaces lo0 unit 3]
user@R1# set family iso address 49.0001.1234.1600.2231.00
```

2. Enable IS-IS on the interfaces.

```
[edit logical-systems LS3 protocols isis]
user@R1# set interface lt-1/2/0.3
```

```
user@R1# set interface lt-1/2/0.5
user@R1# set interface lo0.3 passive
```

3. Configure the default route on Logical System LS3.

```
[edit logical-systems LS3 routing-options]
user@R1# set static route 0.0.0.0/0 discard
user@R1# set static route 0.0.0.0/0 no-install
```

4. Configure the default route policy on Logical System LS3.

```
[edit logical-systems LS3 policy-options]
user@R1# set policy-statement isis-default from protocol static
user@R1# set policy-statement isis-default from route-filter 0.0.0.0/0 exact
user@R1# set policy-statement isis-default then accept
```

5. Apply the export policy to IS-IS on Logical System LS3.

```
[edit logical-systems LS3 protocols isis]
user@R1# set export isis-default
```

6. If you are done configuring the device, commit the configuration.

```
[edit]
user@R1# commit
```

Results

From configuration mode, confirm your configuration by issuing the **show logical-systems LS3** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show logical-systems LS3
interfaces {
  lt-1/2/0 {
    unit 3 {
      description LS3->LS2;
      encapsulation ethernet;
      peer-unit 4;
      family inet {
        address 10.0.2.1/30;
      }
      family iso;
    }
    unit 5 {
      description LS3->LS1;
      encapsulation ethernet;
      peer-unit 0;
      family inet {
        address 10.0.1.1/30;
      }
      family iso;
    }
  }
  lo0 {
    unit 3 {
      family iso {
```

```

        address 49.0001.1234.1600.2231.00;
    }
}
}
}
protocols {
    isis {
        export isis-default;
        interface lt-1/2/0.3;
        interface lt-1/2/0.5;
        interface lo0.3 {
            passive;
        }
    }
}
policy-options {
    policy-statement isis-default {
        from {
            protocol static;
            route-filter 0.0.0.0/0 exact;
        }
        then accept;
    }
}
routing-options {
    static {
        route 0.0.0.0/0 {
            discard;
            no-install;
        }
    }
}
}

```

Verification

Confirm that the configuration is working properly.

Verifying That the Static Route Is Redistributed

Purpose Make sure that the IS-IS policy is working by checking the routing tables.

```

Action user@R1> show route logical-system LS3
inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 00:00:45
                   Discard
10.0.0.0/30        *[IS-IS/15] 1w0d 10:14:14, metric 20
                   to 10.0.2.2 via lt-1/2/0.3
                   > to 10.0.1.2 via lt-1/2/0.5
10.0.1.0/30        *[Direct/0] 1w0d 10:15:18
                   > via lt-1/2/0.5
10.0.1.1/32        *[Local/0] 1w0d 10:15:18
                   Local via lt-1/2/0.5
10.0.2.0/30        *[Direct/0] 1w0d 10:15:18
                   > via lt-1/2/0.3
10.0.2.1/32        *[Local/0] 1w0d 10:15:18
                   Local via lt-1/2/0.3

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

49.0001.1234.1600.2231/72
                   *[Direct/0] 1w0d 10:17:19
                   > via lo0.3

user@R1> show route logical-system LS2
inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[IS-IS/160] 00:01:38, metric 10
                   > to 10.0.2.1 via lt-1/2/0.4
10.0.0.0/30        *[Direct/0] 1w0d 10:16:11
                   > via lt-1/2/0.1
10.0.0.2/32        *[Local/0] 1w0d 10:16:11
                   Local via lt-1/2/0.1
10.0.1.0/30        *[IS-IS/15] 1w0d 10:15:07, metric 20
                   > to 10.0.0.1 via lt-1/2/0.1
                   to 10.0.2.1 via lt-1/2/0.4
10.0.2.0/30        *[Direct/0] 1w0d 10:16:11
                   > via lt-1/2/0.4
10.0.2.2/32        *[Local/0] 1w0d 10:16:11
                   Local via lt-1/2/0.4

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

49.0001.1720.1600.2002/72
                   *[Direct/0] 1w0d 10:18:12
                   > via lo0.2

user@R1> show route logical-system LS1
inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[IS-IS/160] 00:02:01, metric 10
                   > to 10.0.1.1 via lt-1/2/0.0
10.0.0.0/30        *[Direct/0] 1w0d 10:16:34
                   > via lt-1/2/0.2
10.0.0.1/32        *[Local/0] 1w0d 10:16:34
                   Local via lt-1/2/0.2

```

```

10.0.1.0/30      *[Direct/0] 1w0d 10:16:34
                  > via lt-1/2/0.0
10.0.1.2/32      *[Local/0] 1w0d 10:16:34
                  Local via lt-1/2/0.0
10.0.2.0/30      *[IS-IS/15] 1w0d 10:15:55, metric 20
                  to 10.0.1.1 via lt-1/2/0.0
                  > to 10.0.0.2 via lt-1/2/0.2

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

49.0001.1720.1600.1001/72
                  *[Direct/0] 1w0d 10:18:35
                  > via lo0.1

```

Meaning The routing table on Logical System LS3 contains the default 0.0.0.0/0 route from protocol **IS-IS**. The routing tables on Logical System LS1 and Logical System LS2 contain the default 0.0.0.0/0 route from protocol **IS-IS**. If Logical System LS1 and Logical System LS2 receive packets destined for networks not specified in their routing tables, those packets will be sent to Logical System LS3 for further processing. This configuration assumes that Logical System LS3 has a connection to an ISP or another external network.

Related Documentation

- [Example: Creating an Interface on a Logical System](#)

Example: Configuring IS-IS for CLNS

- [Understanding IS-IS for CLNS on page 4237](#)
- [Example: Configuring IS-IS for CLNS on page 4237](#)

Understanding IS-IS for CLNS

IS-IS extensions provide the basic interior gateway protocol (IGP) support for collecting intradomain routing information for Connectionless Network Service (CLNS) destinations within a CLNS network. Routers that learn host addresses through End System-to-Intermediate System (ES-IS) can advertise the addresses to other routers (intermediate systems) by using IS-IS.

For more information about IS-IS, see the ISO 10589 standard.

Example: Configuring IS-IS for CLNS

This example shows how to create a routing instance and enable IS-IS protocol on all interfaces.

- [Requirements on page 4238](#)
- [Overview on page 4238](#)
- [Configuration on page 4238](#)
- [Verification on page 4239](#)

Requirements

Before you begin, configure the network interfaces. See *Interfaces Feature Guide for Security Devices*.

Overview

The configuration instructions in this topic describe how to create a routing-instance called `aaaa`, enable IS-IS on all interfaces, and define BGP export policy name (`dist-bgp`), family (`ISO`), and protocol (`BP`), and apply the export policy to IS-IS.

Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set routing-instances aaaa protocols isis clns-routing
set routing-instances aaaa protocols isis interface all
set routing-instances aaaa protocols isis no-ipv4-routing no-ipv6-routing
set policy-options policy-statement dist-bgp from family iso protocol bgp
set policy-options policy-statement dist-bgp then accept
set routing-instances aaaa protocols isis export dist-bgp
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure IS-IS for CLNS:

1. Configure the routing instance.

```
[edit]
user@host# edit routing-instances aaaa
```
2. Enable CLNS routing.

```
[edit routing-instances aaaa]
user@host# set protocols isis clns-routing
```
3. Enable IS-IS on all interfaces.

```
[edit routing-instances aaaa]
user@host# set protocols isis interface all
```
4. (Optional) Disable IPv4 and IPv6 routing to configure a pure CLNS network .

```
[edit routing-instances aaaa]
user@host# set protocols isis no-ipv4-routing no-ipv6-routing
```
5. Define the BGP export policy name, family, and protocol.

```
[edit policy-options]
user@host# set policy-statement dist-bgp from family iso protocol bgp
```
6. Define the action for the export policy.

```
[edit policy-options]
```

```
user@host# set policy-statement dist-bgp then accept
```

7. Apply the export policy to IS-IS.

```
[edit routing-instances aaaa]
user@host# set protocols isis export dist-bgp
```

Results From configuration mode, confirm your configuration by entering the **show routing-instances** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show routing-instances
aaaa {
  protocols {
    isis {
      export dist-bgp;
      no-ipv4-routing;
      no-ipv6-routing;
      clns-routing;
      interface all;
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying Routing-Instance for CLNS on page 4239](#)
- [Verifying IS-IS for CLNS on page 4239](#)

Verifying Routing-Instance for CLNS

Purpose Verify that the policy options are enabled for the routing instance.

Action From operational mode, enter the **show routing-instances** command.

Verifying IS-IS for CLNS

Purpose Verify that IS-IS is enabled.

Action From operational mode, enter the **show protocols** command.

Example: Enabling Packet Checksums on IS-IS Interfaces for Error Checking

This example shows how to enable packet checksums for IS-IS interfaces.

- [Requirements on page 4240](#)
- [Overview on page 4240](#)
- [Configuration on page 4240](#)
- [Verification on page 4241](#)

Requirements

Before you begin, configure IS-IS on both routers. See [“Example: Configuring IS-IS” on page 4169](#) for information about the sample IS-IS configuration.

Overview


Junos OS supports IS-IS checksums as documented in RFC 3358, *Optional Checksums in Intermediate System to Intermediate System (ISIS)*.

IS-IS protocol data units (PDUs) include link-state PDUs, complete sequence number PDUs (CSNPs), partial sequence number PDUs (PSNPs), and IS-IS hello (IIH) packets. These PDUs can be corrupt due to faulty implementations of Layer 2 hardware or lack of checksums on a specific network technology. Corruption of length or type, length, and value (TLV) fields can lead to the generation of extensive numbers of empty link-state PDUs in the receiving node. Because authentication is not a replacement for a checksum mechanism, you might want to enable the optional checksum TLV on your IS-IS interfaces.

The checksum cannot be enabled with MD5 hello authentication on the same interface.

[Figure 104](#) shows the topology used in this example.

Figure 104: IS-IS Checksum Topology



This example describes the steps on Device R1.

Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device R1

```

set protocols isis traceoptions file isis
set protocols isis traceoptions flag all
set protocols isis interface fe-1/2/0.1 checksum

```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IS-IS checksums:

1. Enable checksums.

```

[edit protocols isis interface fe-1/2/0.1]
user@R1# set checksum

```


2. (Optional) Enable tracing for tracking checksum operations.

```
[edit protocols isis traceoptions]
user@R1# set file isis
user@R1# set flag all
```

Results From configuration mode, confirm your configuration by entering the **show protocols** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show protocols
isis {
  traceoptions {
    file isis;
    flag all;
  }
  interface fe-1/2/0.1 {
    checksum;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

Verifying Checksums

Purpose Verify that checksums are performed.

Action From operational mode, enter the **show log isis | match checksum** command.

```
user@R1> show log isis | match checksum
```

```
May 31 16:47:39.513267      sequence 0x49 checksum 0x8e64
May 31 16:47:39.513394      sequence 0x4e checksum 0x34b3
May 31 16:47:39.513517      sequence 0x50 checksum 0x9dcb
May 31 16:47:46.563781      sequence 0x45 checksum 0x7e1a
May 31 16:47:46.563970      sequence 0x46 checksum 0x226d
May 31 16:47:46.564104      sequence 0x52 checksum 0x99cd
May 31 16:47:46.581087      sequence 0x49 checksum 0x8e64
May 31 16:47:46.581222      sequence 0x4e checksum 0x34b3
May 31 16:47:46.581353      sequence 0x50 checksum 0x9dcb
May 31 16:47:55.799090      sequence 0x45 checksum 0x7e1a
May 31 16:47:55.799223      sequence 0x46 checksum 0x226d
May 31 16:47:55.799347      sequence 0x52 checksum 0x99cd
May 31 16:47:55.818255      sequence 0x49 checksum 0x8e64
May 31 16:47:55.818473      sequence 0x4e checksum 0x34b3
May 31 16:47:55.818606      sequence 0x50 checksum 0x9dcb
May 31 16:48:03.455816      sequence 0x49 checksum 0x8e64
May 31 16:48:03.455973      sequence 0x4e checksum 0x34b3
```

Meaning The output shows that checksum information is captured in the IS-IS trace log file.

- Related Documentation**
- *Understanding Checksums on IS-IS Interfaces for Error Checking*

Configuring BFD Authentication for IS-IS

- [Understanding BFD Authentication for IS-IS on page 4243](#)
- [Example: Configuring BFD for IS-IS on page 4245](#)
- [Example: Configuring BFD Authentication for IS-IS on page 4250](#)

Understanding BFD Authentication for IS-IS

Bidirectional Forwarding Detection (BFD) enables rapid detection of communication failures between adjacent systems. By default, authentication for BFD sessions is disabled. However, when running BFD over Network Layer protocols, the risk of service attacks can be significant. We strongly recommend using authentication if you are running BFD over multiple hops or through insecure tunnels. Beginning with Junos OS Release 9.6, Junos OS supports authentication for BFD sessions running over IS-IS. BFD authentication is only supported in the domestic image and is not available in the export image.

You authenticate BFD sessions by specifying an authentication algorithm and keychain, and then associating that configuration information with a security authentication keychain using the keychain name.

The following sections describe the supported authentication algorithms, security keychains, and level of authentication that can be configured:

- [BFD Authentication Algorithms on page 4243](#)
- [Security Authentication Keychains on page 4244](#)
- [Strict Versus Loose Authentication on page 4244](#)

BFD Authentication Algorithms

Junos OS supports the following algorithms for BFD authentication:

- **simple-password**—Plain-text password. One to 16 bytes of plain text are used to authenticate the BFD session. One or more passwords might be configured. This method is the least secure and should be used only when BFD sessions are not subject to packet interception.
- **keyed-md5**—Keyed Message Digest 5 hash algorithm for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed MD5 uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. With this method, packets are accepted at the receiving

end of the session if one of the keys matches and the sequence number is greater than or equal to the last sequence number received. Although more secure than a simple password, this method is vulnerable to replay attacks. Increasing the rate at which the sequence number is updated can reduce this risk.

- **meticulous-keyed-md5**—Meticulous keyed Message Digest 5 hash algorithm. This method works in the same manner as keyed MD5, but the sequence number is updated with every packet. Although more secure than keyed MD5 and simple passwords, this method might take additional time to authenticate the session.
- **keyed-sha-1**—Keyed Secure Hash Algorithm I for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed SHA uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. The key is not carried within the packets. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than the last sequence number received.
- **meticulous-keyed-sha-1**—Meticulous keyed Secure Hash Algorithm I. This method works in the same manner as keyed SHA, but the sequence number is updated with every packet. Although more secure than keyed SHA and simple passwords, this method might take additional time to authenticate the session.



NOTE: Nonstop active routing (NSR) is not supported with meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

Security Authentication Keychains

The security authentication keychain defines the authentication attributes used for authentication key updates. When the security authentication keychain is configured and associated with a protocol through the keychain name, authentication key updates can occur without interrupting routing and signaling protocols.

The authentication keychain contains one or more keychains. Each keychain contains one or more keys. Each key holds the secret data and the time at which the key becomes valid. The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

BFD allows multiple clients per session, and each client can have its own keychain and algorithm defined. To avoid confusion, we recommend specifying only one security authentication keychain.

Strict Versus Loose Authentication

By default, strict authentication is enabled and authentication is checked at both ends of each BFD session. Optionally, to smooth migration from nonauthenticated sessions to authenticated sessions, you can configure *loose checking*. When loose checking is

configured, packets are accepted without authentication being checked at each end of the session. This feature is intended for transitional periods only.

Related Documentation

- [Example: Configuring BFD Authentication for IS-IS on page 4250](#)

Example: Configuring BFD for IS-IS

This example describes how to configure the Bidirectional Forwarding Detection (BFD) protocol to detect failures in an IS-IS network.



NOTE: BFD is not supported with ISIS for IPV6 on QFX10000 series switches.

- [Requirements on page 4245](#)
- [Overview on page 4245](#)
- [Configuration on page 4245](#)
- [Verification on page 4248](#)

Requirements

Before you begin, configure IS-IS on both routers. See “[Example: Configuring IS-IS](#)” on [page 4169](#) for information about the required IS-IS configuration.

This example uses the following hardware and software components:


- Junos OS Release 7.3 or later
- M Series, MX Series, and T Series routers

Overview

This example shows two routers connected to each other. A loopback interface is configured on each router. IS-IS and BFD protocols are configured on both routers.

[Figure 105](#) shows the sample network.

Figure 105: Configuring BFD for IS-IS



Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Router R1

```
set protocols isis interface so-0/0/0 bfd-liveness-detection detection-time threshold 5
set protocols isis interface so-0/0/0 bfd-liveness-detection minimum-interval 2
set protocols isis interface so-0/0/0 bfd-liveness-detection minimum-receive-interval 1
set protocols isis interface so-0/0/0 bfd-liveness-detection no-adaptation
set protocols isis interface so-0/0/0 bfd-liveness-detection transmit-interval threshold 3
set protocols isis interface so-0/0/0 bfd-liveness-detection transmit-interval
  minimum-interval 1
set protocols isis interface so-0/0/0 bfd-liveness-detection multiplier 2
set protocols isis interface so-0/0/0 bfd-liveness-detection version automatic
```

Router R2

```
set protocols isis interface so-0/0/0 bfd-liveness-detection detection-time threshold 6
set protocols isis interface so-0/0/0 bfd-liveness-detection minimum-interval 3
set protocols isis interface so-0/0/0 bfd-liveness-detection minimum-receive-interval 1
set protocols isis interface so-0/0/0 bfd-liveness-detection no-adaptation
set protocols isis interface so-0/0/0 bfd-liveness-detection transmit-interval threshold 4
set protocols isis interface so-0/0/0 bfd-liveness-detection transmit-interval
  minimum-interval 1
set protocols isis interface so-0/0/0 bfd-liveness-detection multiplier 2
set protocols isis interface so-0/0/0 bfd-liveness-detection version automatic
```

**Step-by-Step
Procedure**

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.



NOTE: To simply configure BFD for IS-IS, only the `minimum-interval` statement is required. The BFD protocol selects default parameters for all the other configuration statements when you use the `bfd-liveness-detection` statement without specifying any parameters.



NOTE: You can change parameters at any time without stopping or restarting the existing session. BFD automatically adjusts to the new parameter value. However, no changes to BFD parameters take place until the values resynchronize with each BFD peer.

To configure BFD for IS-IS on Routers R1 and R2:

1. Enable BFD failure detection for IS-IS.

```
[edit protocols isis]
user@R1# set interface so-0/0/0 bfd-liveness-detection

[edit protocols isis]
user@R2# set interface so-0/0/0 bfd-liveness-detection
```
2. Configure the threshold for the adaptation of the detection time, which must be greater than the multiplier number multiplied by the minimum interval.

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set detection-time threshold 5
```

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set detection-time threshold 6
```

3. Configure the minimum transmit and receive intervals for failure detection.

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set minimum-interval 2
```

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set minimum-interval 3
```

4. Configure only the minimum receive interval for failure detection.

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set minimum-receive-interval 1
```

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set minimum-receive-interval 1
```

5. Disable BFD adaptation.

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set no-adaptation
```

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set no-adaptation
```

6. Configure the threshold for the transmit interval, which must be greater than the minimum transmit interval.

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set transmit-interval threshold 3
```

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set transmit-interval threshold 4
```

7. Configure the minimum transmit interval for failure detection.

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set transmit-interval minimum-interval 1
```

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set transmit-interval minimum-interval 1
```

8. Configure the multiplier number, which is the number of hello packets not received by the neighbor that causes the originating interface to be declared down.

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set multiplier 2
```

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set multiplier 2
```

9. Configure the BFD version used for detection.

The default is to have the version detected automatically.

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set version automatic
```

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set version automatic
```

Results

From configuration mode, confirm your configuration by issuing the **show protocols isis interface** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

user@R1# **show protocols isis interface so-0/0/0**

```
    bfd-liveness-detection {
      version automatic;
      minimum-interval 2;
      minimum-receive-interval 1;
      multiplier 2;
      no-adaptation;
      transmit-interval {
        minimum-interval 1;
        threshold 3;
      }
      detection-time {
        threshold 5;
      }
    }
  ...
```

user@R2# **show protocols isis interface so-0/0/0**

```
    bfd-liveness-detection {
      version automatic;
      minimum-interval 3;
      minimum-receive-interval 1;
      multiplier 2;
      no-adaptation;
      transmit-interval {
        minimum-interval 1;
        threshold 4;
      }
      detection-time {
        threshold 6;
      }
    }
  ...
```

Verification

Confirm that the configuration is working properly.

- [Verifying the Connection Between Routers R1 and R2 on page 4248](#)
- [Verifying That IS-IS Is Configured on page 4249](#)
- [Verifying That BFD Is configured on page 4250](#)

Verifying the Connection Between Routers R1 and R2

Purpose Make sure that Routers R1 and R2 are connected to each other.

Action Ping the other router to check the connectivity between the two routers as per the network topology.

```
user@R1> ping 10.0.0.2
```

```
PING 10.0.0.2 (10.0.0.2): 56 data bytes
64 bytes from 10.0.0.2: icmp_seq=0 ttl=64 time=1.367 ms
64 bytes from 10.0.0.2: icmp_seq=1 ttl=64 time=1.662 ms
64 bytes from 10.0.0.2: icmp_seq=2 ttl=64 time=1.291 ms
^C
--- 10.0.0.2 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.291/1.440/1.662/0.160 ms
```

```
user@R2> ping 10.0.0.1
```

```
PING 10.0.0.1 (10.0.0.1): 56 data bytes
64 bytes from 10.0.0.1: icmp_seq=0 ttl=64 time=1.287 ms
64 bytes from 10.0.0.1: icmp_seq=1 ttl=64 time=1.310 ms
64 bytes from 10.0.0.1: icmp_seq=2 ttl=64 time=1.289 ms
^C
--- 10.0.0.1 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.287/1.295/1.310/0.010 ms
```

Meaning Routers R1 and R2 are connected to each other.

Verifying That IS-IS Is Configured

Purpose Make sure that the IS-IS instance is running on both routers.

Action Use the **show isis database** statement to check if the IS-IS instance is running on both routers, R1 and R2.

```
user@R1> show isis database
```

```
IS-IS level 1 link-state database:
LSP ID          Sequence Checksum Lifetime Attributes
R1.00-00        0x4a571  0x30c5    1195 L1 L2
R2.00-00        0x4a586  0x4b7e    1195 L1 L2
R2.02-00        0x330ca1 0x3492    1196 L1 L2
  3 LSPs
```

```
IS-IS level 2 link-state database:
LSP ID          Sequence Checksum Lifetime Attributes
R1.00-00        0x4a856  0x5db0    1194 L1 L2
R2.00-00        0x4a89d  0x149b    1194 L1 L2
R2.02-00        0x1fb2ff 0xd302    1194 L1 L2
  3 LSPs
```

```
user@R2> show isis database
```

```
IS-IS level 1 link-state database:
LSP ID          Sequence Checksum Lifetime Attributes
R1.00-00        0x4b707  0xcc80    1195 L1 L2
R2.00-00        0x4b71b  0xeb37    1198 L1 L2
R2.02-00        0x33c2ce 0xb52d    1198 L1 L2
  3 LSPs
```

IS-IS level 2 link-state database:

LSP ID	Sequence	Checksum	Lifetime	Attributes
R1.00-00	0x4b9f2	0xee70	1192	L1 L2
R2.00-00	0x4ba41	0x9862	1197	L1 L2
R2.02-00	0x3	0x6242	1198	L1 L2

3 LSPs

Meaning IS-IS is configured on both routers, R1 and R2.

Verifying That BFD Is configured

Purpose Make sure that the BFD instance is running on both routers, R1 and R2.

Action Use the `show bfd session detail` statement to check if BFD instance is running on the routers.

user@R1> show bfd session detail

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.0.0.2	Up	so-0/0/0	2.000	1.000	2

Client ISIS R2, TX interval 0.001, RX interval 0.001
 Client ISIS R1, TX interval 0.001, RX interval 0.001
 Session down time 00:00:00, previous up time 00:00:15
 Local diagnostic NbrSignal, remote diagnostic NbrSignal
 Remote state AdminDown, version 1
 Router 3, routing table index 17

1 sessions, 2 clients

Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps

user@R2> show bfd session detail

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.0.0.1	Up	so-0/0/0	2.000	1.000	2

Client ISIS R2, TX interval 0.001, RX interval 0.001
 Session down time 00:00:00, previous up time 00:00:05
 Local diagnostic NbrSignal, remote diagnostic NbrSignal
 Remote state AdminDown, version 1
 Router 2, routing table index 15

1 sessions, 1 clients

Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps

Meaning BFD is configured on Routers R1 and R2 for detecting failures in the IS-IS network.

Related Documentation

- [Understanding BFD for IS-IS](#)

Example: Configuring BFD Authentication for IS-IS

This example shows how to configure BFD authentication for IS-IS.

- [Requirements on page 4251](#)
- [Overview on page 4251](#)

- [Configuration on page 4251](#)
- [Verification on page 4253](#)

Requirements


Before you begin, configure IS-IS on both routers. See [“Example: Configuring IS-IS” on page 4169](#) for information about the required IS-IS configuration.

Overview

In this example, a BFD authentication keychain is configured with meticulous keyed MD5 authentication.

[Figure 106](#) shows the topology used in this example.

Figure 106: IS-IS BFD Authentication Topology



[“CLI Quick Configuration” on page 4251](#) shows the configuration for both of the devices in [Figure 106](#). The section [“Step-by-Step Procedure” on page 4252](#) describes the steps on Device R1.

Configuration

CLI Quick Configuration	To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.
Device R1	<pre> set security authentication-key-chains key-chain secret123 description for-isis-bfd set security authentication-key-chains key-chain secret123 key 1 secret "\$9\$Cw-yrv" set security authentication-key-chains key-chain secret123 key 1 start-time "2012-5-31.13:00:00 -0700" set security authentication-key-chains key-chain secret123 key 2 secret "\$9\$m5T3" set security authentication-key-chains key-chain secret123 key 2 start-time "2013-5-31.13:00:00 -0700" set security authentication-key-chains key-chain secret123 key 3 secret "\$9\$mTQn" set security authentication-key-chains key-chain secret123 key 3 start-time "2014-5-31.13:00:00 -0700" set protocols isis interface ge-1/2/0.0 bfd-liveness-detection minimum-interval 100 set protocols isis interface ge-1/2/0.0 bfd-liveness-detection authentication key-chain secret123 set protocols isis interface ge-1/2/0.0 bfd-liveness-detection authentication algorithm meticulous-keyed-md5 </pre>
Device R2	<pre> set security authentication-key-chains key-chain secret123 description for-isis-bfd set security authentication-key-chains key-chain secret123 key 1 secret "\$9\$Cw-yrv" set security authentication-key-chains key-chain secret123 key 1 start-time "2012-5-31.13:00:00 -0700" </pre>

```
set security authentication-key-chains key-chain secret123 key 2 secret "$9$m5T3"  
set security authentication-key-chains key-chain secret123 key 2 start-time  
  "2013-5-31.13:00:00 -0700"  
set security authentication-key-chains key-chain secret123 key 3 secret "$9$mTQn"  
set security authentication-key-chains key-chain secret123 key 3 start-time  
  "2014-5-31.13:00:00 -0700"  
set protocols isis interface ge-1/2/0.0 bfd-liveness-detection minimum-interval 100  
set protocols isis interface ge-1/2/0.0 bfd-liveness-detection authentication key-chain  
  secret123  
set protocols isis interface ge-1/2/0.0 bfd-liveness-detection authentication algorithm  
  meticulous-keyed-md5
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IS-IS BFD authentication:

1. Configure the authentication keychain.

```
[edit security authentication-key-chains key-chain secret123]  
user@R1# set description for-isis-bfd  
user@R1# set key 1 secret "$9$cW-yrv"  
user@R1# set key 1 start-time "2012-5-31.13:00:00 -0700"  
user@R1# set key 2 secret "$9$m5T3"  
user@R1# set key 2 start-time "2013-5-31.13:00:00 -0700"  
user@R1# set key 3 secret "$9$mTQn"  
user@R1# set key 3 start-time "2014-5-31.13:00:00 -0700"
```
2. Enable BFD.

```
[edit protocols isis interface ge-1/2/0.0 bfd-liveness-detection]  
user@R1# set minimum-interval 100
```
3. Apply the authentication keychain.

```
[edit protocols isis interface ge-1/2/0.0 bfd-liveness-detection]  
user@R1# set authentication key-chain secret123
```
4. Set the authentication type.

```
[edit protocols isis interface ge-1/2/0.0 bfd-liveness-detection]  
user@R1# set authentication algorithm meticulous-keyed-md5
```

Results From configuration mode, confirm your configuration by entering the **show protocols** and **show security** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show protocols  
isis {  
  interface ge-1/2/0.0 {  
    bfd-liveness-detection {  
      minimum-interval 100;  
      authentication {  
        key-chain secret123;  
        algorithm meticulous-keyed-md5;  
      }  
    }  
  }  
}
```

```

    }
  }
}

user@R1# show security
authentication-key-chains {
  key-chain secret123 {
    description for-isis-bfd;
    key 1 {
      secret "$9$cW-yrv"; ## SECRET-DATA
      start-time "2012-5-31.13:00:00 -0700";
    }
    key 2 {
      secret "$9$m5T3"; ## SECRET-DATA
      start-time "2013-5-31.13:00:00 -0700";
    }
    key 3 {
      secret "$9$mTQn"; ## SECRET-DATA
      start-time "2014-5-31.13:00:00 -0700";
    }
  }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

Verifying IS-IS BFD Authentication

Purpose Verify the status of IS-IS BFD authentication.

Action From operational mode, enter the **show bfd session extensive** command.

```

user@R1> show bfd session extensive

```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.0.0.2	Down	ge-1/2/0.0	0.300	1.000	3

```

Client ISIS L1, TX interval 0.100, RX interval 0.100, Authenticate
  keychain secret123, algo meticulous-keyed-md5, mode strict
Client ISIS L2, TX interval 0.100, RX interval 0.100, Authenticate
  keychain secret123, algo meticulous-keyed-md5, mode strict
Session down time 00:35:13, previous up time 00:12:17
Local diagnostic None, remote diagnostic None
Remote state Up, version 1
Logical system 2, routing table index 85
Min async interval 0.100, min slow interval 1.000
Adaptive async TX interval 0.100, RX interval 0.100
Local min TX interval 1.000, minimum RX interval 0.100, multiplier 3
Remote min TX interval 0.100, min RX interval 0.100, multiplier 3
Local discriminator 2, remote discriminator 1
Echo mode disabled/inactive, no-absorb, no-refresh
Authentication enabled/active, keychain secret123, algo meticulous-keyed-md5,
mode strict
Session ID: 0x100101

```

1 sessions, 2 clients

Cumulative transmit rate 1.0 pps, cumulative receive rate 10.0 pps

Meaning The output shows that BFD authentication is enabled on IS-IS Level 1 and Level 2.

Related Documentation

- *Configuring BFD Authentication for IS-IS*
- [Example: Configuring BFD for IS-IS on page 4245](#)
- *Understanding BFD for IS-IS*

Configuring Hitless Authentication Key Rollover for IS-IS

- [Understanding Hitless Authentication Key Rollover for IS-IS on page 4255](#)
- [Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 4256](#)

Understanding Hitless Authentication Key Rollover for IS-IS

IS-IS protocol exchanges can be authenticated to guarantee that only trusted routing devices participate in routing. By default, authentication is disabled. The authentication algorithm creates an encoded checksum that is included in the transmitted packet. The receiving routing device uses an authentication key (password) to verify the packet's checksum.

If you configure authentication for all peers, each peer in that group inherits the group's authentication.

You can update authentication keys without resetting any IS-IS neighbor sessions. This is referred to as *hitless authentication key rollover*.

Hitless authentication key rollover uses authentication keychains, which consist of the authentication keys that are being updated. The keychain includes multiple keys. Each key in the keychain has a unique start time. At the next key's start time, a rollover occurs from the current key to the next key, and the next key becomes the current key.

You can choose the algorithm through which authentication is established. You can configure MD5 or SHA-1 authentication. You associate a keychain and the authentication algorithm with an IS-IS neighboring session. Each key contains an identifier and a secret password.

The sending peer chooses the active key based on the system time and the start times of the keys in the keychain. The receiving peer determines the key with which it authenticates based on the incoming key identifier.

You can configure either RFC 5304-based encoding or RFC 5310-based encoding for the IS-IS protocol transmission encoding format.

Related Documentation

- [Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 4256](#)

Example: Configuring Hitless Authentication Key Rollover for IS-IS

This example shows how to configure hitless authentication key rollover for IS-IS.

- [Requirements on page 4256](#)
- [Overview on page 4256](#)
- [Configuration on page 4257](#)
- [Verification on page 4260](#)

Requirements

No special configuration beyond device initialization is required before configuring hitless authentication key rollover for IS-IS.

Overview

Authentication guarantees that only trusted routers participate in routing updates. This keychain authentication method is referred to as hitless because the keys roll over from one to the next without resetting any peering sessions or interrupting the routing protocol. Junos OS supports both RFC 5304, *IS-IS Cryptographic Authentication* and RFC 5310, *IS-IS Generic Cryptographic Authentication*.

This example includes the following statements for configuring the keychain:

- **algorithm**—For each key in the keychain, you can specify an encryption algorithm. The algorithm can be SHA-1 or MD-5.
- **key**—A keychain can have multiple keys. Each key within a keychain must be identified by a unique integer value. The range of valid identifier values is from 0 through 63.
- **key-chain**—For each keychain, you must specify a name. This example defines two keychains: **base-key-global** and **base-key-inter**.
- **options**—For each key in the keychain, you can specify the encoding for the message authentication code: **isis-enhanced** or **basic**. The basic (RFC 5304) operation is enabled by default.

When you configure the **isis-enhanced** option, Junos OS sends RFC 5310-encoded routing protocol packets and accepts both RFC 5304-encoded and RFC 5310-encoded routing protocol packets that are received from other devices.

When you configure **basic** (or do not include the **options** statement in the key configuration) Junos OS sends and receives RFC 5304-encoded routing protocols packets, and drops 5310-encoded routing protocol packets that are received from other devices.

Because this setting is for IS-IS only, the TCP and the BFD protocols ignore the encoding option configured in the key.

- **secret**—For each key in the keychain, you must set a secret password. This password can be entered in either encrypted or plain text format in the **secret** statement. It is always displayed in encrypted format.
- **start-time**—Each key must specify a start time in UTC format. Control gets passed from one key to the next. When a configured start time arrives (based on the routing device's clock), the key with that start time becomes active. Start times are specified in the local time zone for a routing device and must be unique within the key chain.


You can apply a keychain globally to all interfaces or more granularly to specific interfaces.

This example includes the following statements for applying the keychain to all interfaces or to particular interfaces:

- **authentication-key-chain**—Enables you to apply a keychain at the global IS-IS level for all Level 1 or all Level 2 interfaces.
- **hello-authentication-key-chain**—Enables you to apply a keychain at the individual IS-IS interface level. The interface configuration overrides the global configuration.

Figure 107 shows the topology used in the example.

Figure 107: Hitless Authentication Key Rollover for IS-IS



This example shows the configuration for Router R0.

Configuration

CLI Quick Configuration

To quickly configure the hitless authentication key rollover for IS-IS, copy the following commands and paste the commands into the CLI.

```
[edit]
set interfaces ge-0/0/0 unit 0 description "interface A"
set interfaces ge-0/0/0 unit 0 family inet address 10.0.0.1/30
set interfaces ge-0/0/0 unit 0 family iso
set interfaces ge-0/0/0 unit 0 family inet6 address fe80::200:f8ff:fe21:67cf/128
set interfaces ge-0/0/1 unit 0 description "interface B"
set interfaces ge-0/0/1 unit 0 family inet address 10.0.0.5/30
set interfaces ge-0/0/1 unit 0 family iso
set interfaces ge-0/0/1 unit 0 family inet6 address 10FB::C:ABC:1F0C:44DA/128
set interfaces ge-0/0/2 unit 0 description "interface C"
set interfaces ge-0/0/2 unit 0 family inet address 10.0.0.9/30
set interfaces ge-0/0/2 unit 0 family iso
set interfaces ge-0/0/2 unit 0 family inet6 address ff06::c3/128
```

```

set security authentication-key-chains key-chain base-key-global key 63 secret
"$9$jfkqfTQnCpBDiCt"
set security authentication-key-chains key-chain base-key-global key 63 start-time
"2011-8-6.06:54:00-0700"
set security authentication-key-chains key-chain base-key-global key 63 algorithm
hmac-sha-1
set security authentication-key-chains key-chain base-key-global key 63 options
isis-enhanced
set security authentication-key-chains key-chain base-key-inter key 0 secret
"$9$8sgx7Vws4ZDkWLGD"
set security authentication-key-chains key-chain base-key-inter key 0 start-time
"2011-8-6.06:54:00-0700"
set security authentication-key-chains key-chain base-key-inter key 0 algorithm md5
set security authentication-key-chains key-chain base-key-inter key 0 options basic
set protocols isis level 1 authentication-key-chain base-key-global
set protocols isis interface ge-0/0/0.0 level 1 hello-authentication-key-chain
base-key-inter

```

Step-by-Step Procedure

To configure hitless authentication key rollover for IS-IS:

1. Configure the Router R0 interfaces.

```

[edit]
user@host# edit interfaces ge-0/0/0 unit 0
[edit interfaces ge-0/0/0 unit 0]
user@host# set description "interface A"
user@host# set family inet address 10.0.0.1/30
user@host# set family iso
user@host# set family inet6 address fe80::200:f8ff:fe21:67cf/128
user@host# exit
[edit]
user@host# edit interfaces ge-0/0/1 unit 0
[edit interfaces ge-0/0/1 unit 0]
user@host# set interfaces ge-0/0/1 unit 0 description "interface B"
user@host# set interfaces ge-0/0/1 unit 0 family inet address 10.0.0.5/30
user@host# set interfaces ge-0/0/1 unit 0 family iso
user@host# set interfaces ge-0/0/1 unit 0 family inet6 address
10fb::c:abc:1f0c:44da/128
user@host# exit
[edit]
user@host# edit interfaces ge-0/0/2 unit 0
[edit interfaces ge-0/0/2 unit 0]
user@host# set description "interface C"
user@host# set family inet address 10.0.0.9/30
user@host# set interfaces ge-0/0/2 unit 0 family iso
user@host# set interfaces ge-0/0/2 unit 0 family inet6 address ff06::c3/128
user@host# exit

```

2. Configure one or more authentication keys.

```

[edit]
user@host# edit security authentication-key-chains key-chain base-key-global
[edit security authentication-key-chains key-chain base-key-global]
user@host# set key 63 secret "$9$jfkqfTQnCpBDiCt"
user@host# set key 63 start-time "2011-8-6.06:54:00-0700"
user@host# set key 63 algorithm hmac-sha-1

```

```

user@host# set key 63 options isis-enhanced
user@host# exit
[edit]
user@host# edit security authentication-key-chains key-chain base-key-inter
[edit security authentication-key-chains key-chain base-key-inter]
user@host# set key 0 secret "$9$8sgx7Vws4ZDkWLGD"
user@host# set key 0 start-time "2011-8-6.06:54:00-0700"
user@host# set key 0 algorithm md5
user@host# set key 0 options basic
user@host# exit

```

3. Apply the base-key-global keychain to all Level 1 IS-IS interfaces on Router R0.

```

[edit]
user@host# edit protocols isis level 1
[edit protocols isis level 1]
set authentication-key-chain base-key-global
user@host# exit

```

4. Apply the base-key-inter keychain to the **ge-0/0/0.0** interface on Router R0.

```

[edit]
user@host# edit protocols isis interface ge-0/0/0.0 level 1
[edit protocols isis interface ge-0/0/0.0 level 1]
set hello-authentication-key-chain base-key-inter
user@host# exit

```

5. If you are done configuring the device, commit the configuration.

```

[edit]
user@host# commit

```

Results

Confirm your configuration by entering the **show interfaces**, **show protocols**, and **show security** commands.

```

user@host# show interfaces
ge-0/0/0 {
  unit 0 {
    description "interface A";
    family inet {
      address 10.0.0.1/30;
    }
    family iso;
    family inet6 {
      address fe80::200:f8ff:fe21:67cf/128;
    }
  }
}
ge-0/0/1 {
  unit 0 {
    description "interface B";
    family inet {
      address 10.0.0.5/30;
    }
    family iso;
  }
}

```

```
        family inet6 {
            address 10FB::C:ABC:1F0C:44DA/128;
        }
    }
}
ge-0/0/2 {
    unit 0 {
        description "interface C";
        family inet {
            address 10.0.0.9/30;
        }
        family iso;
        family inet6 {
            address ff06::c3/128;
        }
    }
}

user@host# show protocols
isis {
    level 1 authentication-key-chain base-key-global;
    interface ge-0/0/0.0 {
        level 1 hello-authentication-key-chain base-key-inter;
    }
}

user@host# show security
authentication-key-chains {
    key-chain base-key-global {
        key 63 {
            secret "$9$jfkqfTQnCpBDiCt"; ## SECRET-DATA
            start-time "2011-8-6.06:54:00-0700";
            algorithm hmac-sha-1;
            options isis-enhanced;
        }
    }
    key-chain base-key-inter {
        key 0 {
            secret "$9$8sgx7Vws4ZDkWLGD"; ## SECRET-DATA
            start-time "2011-8-6.06:54:00-0700";
            algorithm md5;
            options basic;
        }
    }
}
```

Verification

To verify the configuration, run the following commands:

- show isis authentication
- show security keychain

Related Documentation

- [Understanding Hitless Authentication Key Rollover for IS-IS on page 4255](#)

Configuration Statements

- [authentication-key \(Protocols IS-IS\) on page 4263](#)
- [authentication-key-chain \(Protocols IS-IS\) on page 4264](#)
- [authentication-type \(Protocols IS-IS\) on page 4265](#)
- [bfd-liveness-detection \(Protocols IS-IS\) on page 4266](#)
- [checksum \(Protocols IS-IS\) on page 4268](#)
- [csnp-interval on page 4269](#)
- [disable \(Protocols IS-IS\) on page 4270](#)
- [export on page 4271](#)
- [external-preference \(Protocols IS-IS\) on page 4272](#)
- [family \(Protocols IS-IS\) on page 4273](#)
- [graceful-restart \(Protocols IS-IS\) on page 4274](#)
- [hello-authentication-key on page 4275](#)
- [hello-authentication-key-chain on page 4276](#)
- [hello-authentication-type on page 4277](#)
- [hello-interval \(Protocols IS-IS\) on page 4278](#)
- [hello-padding on page 4280](#)
- [hold-time \(Protocols IS-IS\) on page 4282](#)
- [ignore-attached-bit on page 4283](#)
- [interface on page 4284](#)
- [ipv4-multicast on page 4286](#)
- [ipv4-multicast-metric on page 4287](#)
- [ipv6-multicast on page 4288](#)
- [ipv6-multicast-metric on page 4288](#)
- [ipv6-unicast on page 4289](#)
- [ipv6-unicast-metric on page 4290](#)
- [isis on page 4291](#)
- [level \(Global IS-IS\) on page 4292](#)
- [link-protection \(Protocols IS-IS\) on page 4293](#)

- [loose-authentication-check](#) on page 4293
- [lsp-interval](#) on page 4294
- [lsp-lifetime](#) on page 4295
- [max-areas](#) on page 4296
- [mesh-group \(Protocols IS-IS\)](#) on page 4297
- [metric \(Protocols IS-IS\)](#) on page 4298
- [no-adjacency-holddown](#) on page 4299
- [no-authentication-check](#) on page 4300
- [no-csnp-authentication](#) on page 4300
- [no-eligible-backup \(Protocols IS-IS\)](#) on page 4301
- [no-hello-authentication](#) on page 4301
- [no-ipv4-multicast](#) on page 4302
- [no-ipv4-routing](#) on page 4303
- [no-ipv6-multicast](#) on page 4304
- [no-ipv6-routing](#) on page 4305
- [no-ipv6-unicast](#) on page 4306
- [no-psnp-authentication](#) on page 4306
- [no-unicast-topology](#) on page 4307
- [node-link-protection \(Protocols IS-IS\)](#) on page 4308
- [overload \(Protocols IS-IS\)](#) on page 4309
- [passive \(Protocols IS-IS\)](#) on page 4312
- [point-to-point](#) on page 4313
- [preference \(Protocols IS-IS\)](#) on page 4314
- [prefix-export-limit \(Protocols IS-IS\)](#) on page 4315
- [priority \(Protocols IS-IS\)](#) on page 4316
- [reference-bandwidth \(Protocols IS-IS\)](#) on page 4317
- [rib-group \(Protocols IS-IS\)](#) on page 4318
- [spf-options \(Protocols IS-IS\)](#) on page 4319
- [topologies \(Protocols IS-IS\)](#) on page 4320
- [traceoptions \(Protocols IS-IS\)](#) on page 4321
- [traffic-engineering \(Protocols IS-IS\)](#) on page 4324
- [wide-metrics-only](#) on page 4327

authentication-key (Protocols IS-IS)

Syntax	<code>authentication-key key;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis level <i>level-number</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis level <i>level-number</i>], [edit protocols isis level <i>level-number</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis level <i>level-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Authentication key (password). Neighboring routing devices use the password to verify the authenticity of packets sent from this interface. For the key to work, you also must include the authentication-type statement.</p> <p>All routing devices must use the same password. If you are using the Junos OS IS-IS software with another implementation of IS-IS, the other implementation must be configured to use the same password for the domain, the area, and all interfaces adjacent to the Juniper Networks routing device.</p>
Default	If you do not include this statement and the authentication-type statement, IS-IS authentication is disabled.
Options	key —Authentication password. The password can be up to 1024 characters long. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").



CAUTION: A simple password for authentication is truncated if it exceeds 254 characters.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding Hitless Authentication Key Rollover for IS-IS on page 4255 • Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 4256

authentication-key-chain (Protocols IS-IS)

Syntax	authentication-key-chain <i>key-chain-name</i> ;
Hierarchy Level	[edit logical-systems <i>name</i> protocols isis level <i>level-number</i>], [edit logical-systems <i>name</i> routing-instances <i>instance-name</i> protocols isis level <i>level-number</i>], [edit protocols isis level <i>level-number</i>], [edit routing-instances <i>instance-name</i> protocols isis level <i>level-number</i>]
Release Information	Statement introduced in Junos OS Release 11.2. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Apply and enable an authentication keychain to the routing device.
Options	key-chain —Authentication keychain name. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 4256• Example: Configuring Route Authentication for BGP on page 3898• Example: Configuring BFD Authentication for Securing Static Routes on page 3255• Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols• Understanding Hitless Authentication Key Rollover for IS-IS on page 4255

authentication-type (Protocols IS-IS)

Syntax	<code>authentication-type <i>authentication</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis level level-number],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis level level-number],</p> <p>[edit protocols isis level level-number],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis level level-number]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Enable authentication and specify the authentication scheme for IS-IS. If you enable authentication, you must specify a password by including the authentication-key statement.
Default	If you do not include this statement and the authentication-key statement, IS-IS authentication is disabled.
Options	<p><i>authentication</i>—Authentication scheme:</p> <ul style="list-style-type: none"> • md5—Use HMAC authentication in combination with MD5. HMAC-MD5 authentication is defined in RFC 2104, <i>HMAC: Keyed-Hashing for Message Authentication</i>. • simple—Use a simple password for authentication. The password is included in the transmitted packet, making this method of authentication relatively insecure. We recommend that you <i>not</i> use this authentication method.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • authentication-key on page 4263 • no-authentication-check on page 4300 • Understanding Hitless Authentication Key Rollover for IS-IS on page 4255 • Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 4256

bfd-liveness-detection (Protocols IS-IS)

Syntax	<pre>bfd-liveness-detection { authentication { algorithm <i>algorithm-name</i>; key-chain <i>key-chain-name</i>; loose-check; } detection-time { threshold <i>milliseconds</i>; } minimum-interval <i>milliseconds</i>; minimum-receive-interval <i>milliseconds</i>; multiplier <i>number</i>; no-adaptation; transmit-interval { minimum-interval <i>milliseconds</i>; threshold <i>milliseconds</i>; } version (1 automatic); }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis interface interface-name],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface interface-name],</p> <p>[edit protocols isis interface interface-name],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface interface-name]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>detection-time threshold and transmit-interval threshold options added in Junos OS Release 8.2.</p> <p>Support for logical systems introduced in Junos OS Release 8.3.</p> <p>no-adaptation statement introduced in Junos OS Release 9.0.</p> <p>authentication algorithm, authentication key-chain, and authentication loose-check options introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Configure bidirectional failure detection timers and authentication.
Options	<p>authentication algorithm <i>algorithm-name</i>—Configure the algorithm used to authenticate the specified BFD session: simple-password, keyed-md5, keyed-sha-1, meticulous-keyed-md5, meticulous-keyed-sha-1.</p> <p>authentication key-chain <i>key-chain-name</i>—Associate a security key with the specified BFD session using the name of the security keychain. The name you specify must match one of the keychains configured in the authentication-key-chains key-chain statement at the [edit security] hierarchy level.</p>

authentication loose-check—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication might not be configured at both ends of the BFD session.

detection-time threshold *milliseconds*—Configure a threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

minimum-interval *milliseconds*—Configure the minimum interval after which the local routing device transmits a hello packet and then expects to receive a reply from the neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately using the **transmit-interval minimum-interval** and **minimum-receive-interval** statements.

Range: 1 through 255,000

minimum-receive-interval *milliseconds*—Configure the minimum interval after which the local routing device expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the **minimum-interval** statement.

Range: 1 through 255,000

multiplier *number*—Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

Range: 1 through 255

Default: 3

no-adaptation—Specify that BFD sessions not adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

transmit-interval threshold *milliseconds*—Configure the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

transmit-interval minimum-interval *milliseconds*—Configure a minimum interval after which the local routing device transmits hello packets to a neighbor. Optionally, instead of using this statement, you can configure the minimum transmit interval using the **minimum-interval** statement.

Range: 1 through 255,000

version—Configure the BFD version to detect: **1** (BFD version 1) or **automatic** (autodetect the BFD version)

Default: automatic

Required Privilege Level	routing—To view this statement in the configuration.
	routing-control—To add this statement to the configuration.

- Related Documentation**
- [Example: Configuring BFD for IS-IS on page 4245](#)
 - [Understanding BFD for IS-IS](#)
 - [Example: Configuring BFD Authentication for IS-IS on page 4250](#)
 - [Configuring BFD Authentication for IS-IS](#)

checksum (Protocols IS-IS)

Syntax	checksum;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>], [edit protocols isis interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Enable checksums for packets on this interface. Junos OS supports IS-IS checksums as documented in RFC 3358, <i>Optional Checksums in Intermediate System to Intermediate System (ISIS)</i> . The checksum cannot be enabled with MD5 hello authentication on the same interface.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

csnp-interval


Syntax	<code>csnp-interval (seconds disable);</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis interface interface-name],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface interface-name],</p> <p>[edit protocols isis interface interface-name],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface interface-name]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Configure the interval between complete sequence number PDUs (CSNPs) on a LAN interface.</p> <p>If the routing device is the designated router on a LAN, IS-IS sends CSN packets every 10 seconds. If the routing device is on a point-to-point interface, it sends CSN packets every 5 seconds multiplied by the number of IS-IS adjacencies over point-to-point links, which are in UP state.</p> <p>To configure the interface not to send any CSNPs, specify the disable option.</p>
Default	By default, IS-IS sends CSNPs periodically. If the routing device is the designated router on a LAN, IS-IS sends CSNPs every 10 seconds. If the routing device is on a point-to-point interface, it sends CSNPs every 5 seconds multiplied by the number of IS-IS adjacencies over point-to-point links, which are in UP state.
Options	<p>disable—Do not send CSNPs on this interface.</p> <p>seconds—Number of seconds between the sending of CSNPs.</p> <p>Range: 1 through 65,535 seconds</p> <p>Default: 10 seconds on LAN broadcast links. 5 seconds on point-to-point links.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <i>Example: Configuring the Transmission Frequency for CSNP Packets on IS-IS Interfaces</i>

disable (Protocols IS-IS)

Syntax	disable;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols isis traffic-engineering],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis traffic-engineering],</p> <p>[edit protocols isis],</p> <p>[edit protocols isis interface <i>interface-name</i>],</p> <p>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit protocols isis traffic-engineering],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis traffic-engineering]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Disable IS-IS on the routing device, on an interface, or on a level.</p> <p>At the [edit protocols isis traffic-engineering] hierarchy level, disable IS-IS support for traffic engineering.</p> <p>Enabling IS-IS on an interface (by including the interface statement at the [edit protocols isis] or the [edit routing-instances routing-instance-name protocols isis] hierarchy level), disabling it (by including the disable statement), and not actually having IS-IS run on an interface (by including the passive statement) are mutually exclusive states.</p>
Default	<p>IS-IS is enabled for Level 1 and Level 2 routers on all interfaces on which family iso is enabled.</p> <p>IS-IS support for traffic engineering is enabled.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Example: Configuring a Multi-Level IS-IS Topology to Control Interarea Flooding on page 4175

- [IS-IS Overview on page 4157](#)

export

Syntax	<code>export [<i>policy-names</i>];</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis],</p> <p>[edit protocols isis],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Apply one or more policies to routes being exported from the routing table into IS-IS.</p> <p>All routing protocols store the routes that they learn in the routing table. The routing table uses this collected route information to determine the active routes to destinations. The routing table then installs the active routes into its forwarding table and exports them into the routing protocols. It is these exported routes that the protocols advertise.</p> <p>For each protocol, you control which routes the protocol stores in the routing table and which routes the routing table exports into the protocol from the routing table by defining a <i>routing policy</i> for that protocol.</p>
	<p> NOTE: For IS-IS, you cannot apply routing policies that affect how routes are imported into the routing table; doing so with a link-state protocol can easily lead to an inconsistent topology database.</p>
Options	<i>policy-names</i> —Name of one or more policies.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	

external-preference (Protocols IS-IS)

Syntax	<code>external-preference <i>preference</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis level level-number], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis level level-number], [edit protocols isis level level-number], [edit routing-instances <i>routing-instance-name</i> protocols isis level level-number]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the preference of external routes.
Options	<i>preference</i> —Preference value. Range: 0 through 4,294,967,295 ($2^{32} - 1$) Default: 15 (for Level 1 internal routes), 18 (for Level 2 internal routes), 160 (for Level 1 external routes), 165 (for Level 2 external routes)
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Route Preferences Overview</i>• Example: Redistributing OSPF Routes into IS-IS on page 4185• <i>Example: Configuring a Routing Policy to Redistribute BGP Routes with a Specific Community Tag into IS-IS</i>• <i>Understanding BGP Communities and Extended Communities as Routing Policy Match Conditions</i>• <i>Understanding Routing Policies</i>• preference on page 4314

family (Protocols IS-IS)

Syntax	<pre>family inet { shortcuts { multicast-rpf-routes; } } family inet6 { shortcuts; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis traffic-engineering], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis traffic-engineering], [edit protocols isis traffic-engineering], [edit routing-instances <i>routing-instance-name</i> protocols isis traffic-engineering]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.3. Support for IPv6 for IGP shortcuts introduced in Junos OS Release 9.3. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Configure the address family for traffic engineering IS-IS interior gateway protocol (IGP) shortcuts.
Options	<p>inet—IPv4 address family</p> <p>inet6—IPv6 address family</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

graceful-restart (Protocols IS-IS)

Syntax	<pre>graceful-restart { disable; helper-disable; restart-duration <i>seconds</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis], [edit protocols isis]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	<p>Configure graceful restart parameters for IS-IS.</p> <p>Graceful restart allows a routing device to restart with minimal effects to the network, and is enabled for all routing protocols at the [edit routing-options] hierarchy level. When graceful restart is enabled, the restarting routing device is not removed from the network topology during the restart period. The adjacencies are reestablished after restart is complete.</p> <p>On LAN interfaces where IS-IS is configured on a transit router that serves as the designated router (DR), a graceful restart causes:</p> <ul style="list-style-type: none">• The ingress router of the label-switched path (LSP), which passes through the DR, to break the LSP.• The ingress router to re-signal the LSP.
Options	<p>disable—Disable graceful restart for IS-IS.</p> <p>helper-disable—Disable graceful restart helper capability. Helper mode is enabled by default.</p> <p>restart-duration <i>seconds</i>—Time period for the restart to last, in seconds. Range: 30 through 300 seconds Default: 30 seconds</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Routing Protocols Graceful Restart on page 2580


hello-authentication-key

Syntax	<code>hello-authentication-key password;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>number</i>],</p> <p>[edit protocols isis interface <i>interface-name</i> level <i>number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>number</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Configure an authentication key (password) for hello packets. Neighboring routing devices use the password to verify the authenticity of packets sent from an interface. For the key to work, you also must include the hello-authentication-type statement.
Default	By default, hello authentication is not configured on an interface. However, if IS-IS authentication is configured, the hello packets are authenticated using the IS-IS authentication type and password.
Options	<p>password—Authentication password. The password can be up to 255 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • authentication-key on page 4263 • authentication-type on page 4265 • hello-authentication-type on page 4277


hello-authentication-key-chain

Syntax	hello-authentication-key-chain <i>key-chain-name</i> ;
Hierarchy Level	[edit logical-systems <i>name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit logical-systems <i>name</i> routing-instances <i>instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit routing-instances <i>instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]
Release Information	Statement introduced in Junos OS Release 11.2. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Apply an authentication keychain to the IS-IS interface.
Options	<i>key-chain-name</i> —Authentication keychain name. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding Hitless Authentication Key Rollover for IS-IS on page 4255• Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 4256

hello-authentication-type

Syntax	hello-authentication-type (md5 simple);
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>number</i>],</p> <p>[edit protocols isis interface <i>interface-name</i> level <i>number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>number</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Enable authentication on an interface for hello packets. If you enable authentication on hello packets, you must specify a password by including the hello-authentication-key statement.</p> <p>You can configure authentication for a given IS-IS level on an interface. On a point-to-point link, if you enable hello authentication for both IS-IS levels, the password configured for Level 1 is used for both levels.</p>
<div style="display: flex; align-items: center;">  <div style="margin-left: 10px;"> <p>CAUTION: If no authentication is configured for Level 1 on a point-to-point link with both levels enabled, the hello packets are sent without any password, regardless of the Level 2 authentication configurations.</p> </div> </div>	
Default	By default, hello authentication is not configured on an interface. However, if IS-IS authentication is configured, the hello packets are authenticated using the IS-IS authentication type and password.
Options	<p>md5—Specifies Message Digest 5 as the packet verification type.</p> <p>simple—Specifies simple authentication as the packet verification type.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • authentication-key on page 4263 • authentication-type on page 4265 • hello-authentication-key on page 4275

hello-interval (Protocols IS-IS)

Syntax	<code>hello-interval seconds;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Modify the frequency with which the routing device sends hello packets out of an interface, in seconds.</p> <p>Routing devices send hello packets at a fixed interval on all interfaces to establish and maintain neighbor relationships. This interval is advertised in the hello interval field in the hello packet.</p> <p>You can send out hello packets in subsecond intervals. To send out hello packets every 333 milliseconds, set the hold-time value to 1.</p>
Options	<p>seconds—Frequency of transmission for hello packets.</p> <p>Range: 1 through 20,000 seconds</p> <p>Default: 3 seconds (for designated intermediate system [DIS] routers), 9 seconds (for non-DIS routers)</p>
<div>  <p>NOTE: When elected as a designated intermediate system [DIS] router on any LAN adjacency, the hello and hold-timer intervals are scaled down by a factor of 3. This means the default values of 9 and 27 seconds of ISIS hello and hold time intervals are scaled down to 3 and 9 seconds for LAN adjacencies. During switchovers, this hello interval is too short to form LAN adjacencies. Therefore, you can configure one of the following solutions:</p> <ul style="list-style-type: none"> Set the hello and hold time interval for LAN adjacencies to 30 seconds and 90 seconds respectively on both the DIS and a neighboring router. Convert the LAN interfaces to point-to-point IS-IS interfaces. </div>	
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

Related • *hold-time*
Documentation

hello-padding

Syntax	hello-padding (adaptive loose strict);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>], [edit protocols isis interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 8.0. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Configure padding on hello packets to accommodate asymmetrical maximum transfer units (MTUs) from different hosts.</p> <p>This helps to prevent a premature adjacency Up state when one routing device's MTU does not meet the requirements to establish the adjacency.</p> <p>As an OSI Layer 2 protocol, IS-IS does not support data fragmentation. Therefore, maximum packet sizes must be established and supported between two routers. During adjacency establishment, the IS-IS protocol makes sure that the link supports a packet size of 1492 bytes by padding outgoing hello packets up to the maximum packet size of 1492 bytes.</p> <p>This is the default behavior of the Junos OS IS-IS implementation. However, Junos OS provides an option to disable hello padding that can override this behavior.</p> <p>There are four types of hello padding:</p> <ul style="list-style-type: none">• Adaptive padding—On point-to-point connections, the hello packets are padded from the initial detection of a new neighbor until the neighbor verifies the adjacency as Up in the adjacency state type, length, and value (TLV) tuple. If the neighbor does not support the adjacency state TLV, then padding continues. On LAN connections, padding starts from the initial detection of a new neighbor until there is at least one active adjacency on the interface. Adaptive padding has more overhead than loose padding and is able to detect MTU asymmetry from one side of the connection. This one-sided detection can result in generation of extra link-state PDUs that are flooded throughout the network. Specify the adaptive option to configure enough padding to establish an adjacency to neighbors.• Disabled padding—Padding is disabled on all types of interfaces for all adjacency states. Specify the disable option to accommodate interfaces that support less than the default packet size of 1492 bytes.• Loose padding (the default)—The hello packet is padded from the initial detection of a new neighbor until the adjacency transitions to the Up state. Loose padding might not be able to detect certain situations such as asymmetrical MTUs between the routing devices. Specify the loose option to configure enough padding to initialize an adjacency to neighbors.

- **Strict padding**—Padding is done on all interface types and for all adjacency states, and is continuous. Strict padding has the most overhead. The advantage is that strict padding detects MTU issues on both sides of a link. Specify the **strict** option to configure padding to allow all adjacency states with neighbors.

Options **adaptive**—Configure padding until the neighbor adjacency is established and active.

disable—Disable padding on all types of interfaces for all adjacency states.

loose—Configure padding until the state of the adjacency is initialized.

strict—Configure padding for all adjacency states.

Required Privilege routing—To view this statement in the configuration.
Level routing-control—To add this statement to the configuration.

hold-time (Protocols IS-IS)

Syntax	<code>hold-time seconds;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</code> <code>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Set the length of time a neighbor considers this router to be operative (up) after receiving a hello packet. If the neighbor does not receive another hello packet within the specified time, it marks this routing device as inoperative (down). The hold time itself is advertised in the hello packets.</p> <p>The hold time specifies how long a neighbor should consider this routing device to be operative without receiving another hello packet. If the neighbor does not receive a hello packet from this routing device within the hold time, it marks the routing device as being unavailable.</p> <p>For systems configured with graceful routing switchover (GRES) with Graceful Restart, the hold time for Master and Backup Routing Engines should be set to a value higher than 40 seconds. This ensures that adjacencies between the Routing Engine and the neighboring peer 'helper' routers do not time out, stopping graceful restart, and all traffic.</p>
Options	seconds —Hold-time value, in seconds. Range: 3 through 65,535 seconds, or 1 to send out hello packets every 333 milliseconds Default: 9 seconds (for designated intermediate system [DIS] routers), 27 seconds (for non-DIS routers; three times the default hello interval)



NOTE: When elected as a designated intermediate system [DIS] router on any LAN adjacency, the hello and hold-timer intervals are scaled down by a factor of 3. This means the default values of 9 and 27 seconds of ISIS hello and hold time intervals are scaled down to 3 and 9 seconds for LAN adjacencies. During switchovers, this hold time is too short to form LAN adjacencies. Therefore, you can configure one of the following solutions:

- Set the hello and hold time interval for LAN adjacencies to 30 seconds and 90 seconds respectively on both the DIS router and a neighboring router.
 - Convert the LAN interfaces to point-to-point IS-IS interfaces.
-

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Graceful Routing Engine Switchover on page 2627 • <i>Example: Configuring IS-IS</i> • <i>Example: Configuring IS-IS for GRES with Graceful Restart</i> • hello-interval on page 4278

ignore-attached-bit

Syntax	ignore-attached-bit;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis], [edit protocols isis], [edit routing-instances <i>routing-instance-name</i> protocols isis]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Ignore the attached bit on IS-IS Level 1 routers. Configuring this statement enables the routing device to ignore the attached bit on incoming Level 1 link-state PDUs. If the attached bit is ignored, no default route, which points to the routing device which has set the attached bit, is installed.</p> <p>There might be times, such as during a denial-of-service (DoS) attack, that you do not want a Level 1 router to be able to forward traffic based on a default route.</p> <p>To prevent a routing device from being able to reach interarea destinations, you can prevent the routing device from installing the default route without affecting the status of its IS-IS adjacencies. The ignore-attached-bit statement is used to tell the routing device to ignore the presence of the attached bit in Level 1 link-state PDUs, which blocks the installation of the IS-IS default route.</p>
Default	The ignore-attached-bit statement is disabled by default.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

interface

```
Syntax  interface (all | interface-name) {
        disable;
        bfd-liveness-detection {
            authentication {
                algorithm algorithm-name;
                key-chain key-chain-name;
                loose-check;
            }
            detection-time {
                threshold milliseconds;
            }
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            transmit-interval {
                threshold milliseconds;
                minimum-interval milliseconds;
            }
            multiplier number;
        }
        checksum;
        csnp-interval (seconds | disable);
        hello-padding (adaptive | loose | strict);
        ldp-synchronization {
            disable;
            hold-time seconds;
        }
        lsp-interval milliseconds;
        mesh-group (value | blocked);
        no-adjacency-holddown;
        no-ipv4-multicast;
        no-ipv6-multicast;
        no-ipv6-unicast;
        no-unicast-topology;
        passive;
        point-to-point;
        level level-number {
            disable;
            hello-authentication-key key;
            hello-authentication-key-chain key-chain-name;
            hello-authentication-type authentication;
            hello-interval seconds;
            hold-time seconds;
            ipv4-multicast-metric metric;
            ipv6-multicast-metric metric;
            ipv6-unicast-metric metric;
            metric metric;
            passive;
            priority number;
            te-metric metric;
        }
    }
```

Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis],</p> <p>[edit protocols isis],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> <p>no-eligible-remote-backup option introduced in Junos OS Release 14.2 for the MX Series.</p>
Description	<p>Configure interface-specific IS-IS properties. To configure more than one interface, include the interface statement multiple times.</p> <p>Enabling IS-IS on an interface (by including the interface statement at the [edit protocols isis] or the [edit routing-instances <i>routing-instance-name</i> protocols isis] hierarchy level), disabling it (by including the disable statement), and not actually having IS-IS run on an interface (by including the passive statement) are mutually exclusive states.</p>
Options	<p>all—Have Junos OS create IS-IS interfaces automatically. If you include this option, disable IS-IS on the management interface (fxp0).</p> <p>interface-name—Name of an interface. Specify the full interface name, including the physical and logical address components.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring of Interface-Specific IS-IS Properties • Example: Configuring a Multi-Level IS-IS Topology to Control Interarea Flooding on page 4175

ipv4-multicast

Syntax	ipv4-multicast;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis topologies], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis topologies], [edit protocols isis topologies], [edit routing-instances <i>routing-instance-name</i> protocols isis topologies]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure alternate IPv4 multicast topologies.



NOTE: The IS-IS interface metrics for the IPv4 topology can be configured independently of the IPv6 metrics. You can also selectively disable interfaces from participating in the IPv6 topology while continuing to participate in the IPv4 topology. This lets you exercise control over the paths that unicast data takes through a network.

Default	Multicast topologies are disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring IS-IS Multicast Topology on page 4203• IS-IS Multicast Topologies Overview on page 4202

ipv4-multicast-metric

Syntax	<code>ipv4-multicast-metric <i>metric</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Specify the multicast topology metric value for the level.
Options	<p><i>metric</i>—Metric value.</p> <p>Range: 0 through 16,777,215</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring IS-IS Multicast Topology on page 4203 • IS-IS Multicast Topologies Overview on page 4202


ipv6-multicast

Syntax	ipv6-multicast;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis topologies], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis topologies], [edit protocols isis topologies], [edit routing-instances <i>routing-instance-name</i> protocols isis topologies]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure alternate IPv6 multicast topologies.
Default	Multicast topologies are disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring IS-IS Multicast Topology on page 4203• IS-IS Multicast Topologies Overview on page 4202

ipv6-multicast-metric

Syntax	ipv6-multicast-metric <i>metric</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Specify the IPv6 alternate multicast topology metric value for the level.
Options	<i>metric</i> —Metric value. Range: 0 through 16,777,215
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring IS-IS Multicast Topology on page 4203• IS-IS Multicast Topologies Overview on page 4202

ipv6-unicast

Syntax	ipv6-unicast;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis topologies], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis topologies], [edit protocols isis topologies], [edit routing-instances <i>routing-instance-name</i> protocols isis topologies]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure alternate IPv6 unicast topologies. This statement causes IS-IS to calculate an alternate IPv6 unicast topology, in addition to the normal IPv4 unicast topology, and add the corresponding routes to inet6.0.
<div>  <p>NOTE: The IS-IS interface metrics for the IPv4 topology can be configured independently of the IPv6 metrics. You can also selectively disable interfaces from participating in the IPv6 topology while continuing to participate in the IPv4 topology. This lets you exercise control over the paths that unicast data takes through a network.</p> </div>	
Default	IPv6 unicast topologies are disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding IS-IS IPv4 and IPv6 Unicast Topologies on page 4193 • Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies on page 4193

ipv6-unicast-metric

Syntax	<code>ipv6-unicast-metric <i>metric</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Specify the IPv6 unicast topology metric value for the level. The IS-IS interface metrics for the IPv4 topology can be configured independently of the IPv6 metrics.
Options	<i>metric</i> —Metric value. Range: 0 through 16,777,215
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies on page 4193• Understanding IS-IS IPv4 and IPv6 Unicast Topologies on page 4193

isis

Syntax	isis { ... }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols], [edit protocols], [edit routing-instances <i>routing-instance-name</i> protocols]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Enable IS-IS routing on the routing device or for a routing instance. The isis statement is the one statement you must include in the configuration to run IS-IS on the routing device or in a routing instance.
Default	IS-IS is disabled on the routing device.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring IS-IS on page 4169 • Example: Configuring a Multi-Level IS-IS Topology to Control Interarea Flooding on page 4175

level (Global IS-IS)

Syntax	<pre>level <i>level-number</i> { authentication-key <i>key</i>; authentication-key-chain (Protocols IS-IS) <i>key-chain-name</i>; authentication-type <i>type</i>; disable; external-preference <i>preference</i>; no-csnp-authentication; no-hello-authentication; no-psnp-authentication; preference <i>preference</i>; wide-metrics-only; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis], [edit protocols isis], [edit routing-instances <i>routing-instance-name</i> protocols isis]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Configure the global-level properties.</p> <p>You can administratively divide a single AS into smaller groups called areas. You configure each routing device interface to be in an area. Any interface can be in any area. The area address applies to the entire routing device. You cannot specify one interface to be in one area and another interface in a different area. To route between areas, you must have two adjacent Level 2 routers that communicate with each other.</p> <p>Level 1 routers can only route within their IS-IS area. To send traffic outside their area, Level 1 routers must send packets to the nearest intra-area Level 2 router. A routing device can be a Level 1 router, a Level 2 router, or both. You specify the router level on a per-interface basis, and a routing device becomes adjacent to other routing devices on the same level on that link only.</p> <p>You can configure one Level 1 routing process and one Level 2 routing process on each interface, and you can configure the two levels differently.</p>
Options	<p>level-number—IS-IS level number. Values: 1 or 2</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.</p>

link-protection (Protocols IS-IS)

Syntax	link-protection;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>], [edit protocols isis interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 9.5. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Enable link protection on the specified IS-IS interface. Junos OS creates a backup loop-free alternate path to the primary next hop for all destination routes that traverse the protected interface.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding Loop-Free Alternate Routes for IS-IS on page 4165 • Example: Configuring Node-Link Protection for IS-IS Routes in a Layer 3 VPN on page 4221 • node-link-protection on page 4308

loose-authentication-check

Syntax	loose-authentication-check;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols <i>isis</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>isis</i>], [edit protocols <i>isis</i>], [edit routing-instances <i>routing-instance-name</i> protocols <i>isis</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Allow the use of MD5 authentication without requiring network-wide deployment.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding Hitless Authentication Key Rollover for IS-IS on page 4255 • Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 4256

lsp-interval

Syntax	<code>lsp-interval milliseconds;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>], [edit protocols isis interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Configure the link-state PDU interval time.</p> <p>By default, the routing device sends one link-state PDU packet out an interface every 100 milliseconds. To disable the transmission of all link-state PDUs, set the interval to 0.</p> <p>Link-state PDU throttling by use of the lsp-interval statement controls the flooding pace to neighboring routing devices in order to not overload them.</p> <p>Also, consider that control traffic (such as link-state PDUs and related packets) might delay user traffic (information packets) because control traffic always has precedence in terms of scheduling on the routing device interface cards. Unfortunately, the control traffic transmission rate is not decreased on low-bandwidth interfaces, such as DS-0 or fractional T1 and E1 interface. Line control traffic stays the same. On a low-bandwidth circuit that is transmitting 30 full-MTU-sized packets, there is not much bandwidth left over for other types of packets.</p>
Default	By default, the routing device sends one link-state PDU out an interface every 100 milliseconds.
Options	milliseconds —Number of milliseconds between the sending of link-state PDUs. Specifying a value of 0 blocks all link-state PDU transmission. Range: 0 through 1000 milliseconds Default: 100 milliseconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring the Transmission Frequency for Link-State PDUs on IS-IS Interfaces</i>• <i>Understanding Link-State PDU Throttling for IS-IS Interfaces</i>• <i>Example: Configuring the Transmission Frequency for CSNP Packets on IS-IS Interfaces</i>• <i>Understanding the Transmission Frequency for CSNPs on IS-IS Interfaces</i>

lsp-lifetime

Syntax	<code>lsp-lifetime <i>seconds</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis],</p> <p>[edit protocols isis],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Specify how long a link-state PDU originating from the routing device should persist in the network. The routing device sends link-state PDUs often enough so that the link-state PDU lifetime never expires.</p> <p>Because link-state PDUs have a maximum lifetime, they need to be refreshed. Refreshing means that a routing device needs to re-originate its link-state PDUs periodically. The re-origination interval must be less than the link-state PDU's lifetime. For example, if the link-state PDU is valid for 1200 seconds, the routing device needs to refresh the link-state PDU in less than 1200 seconds to avoid removal of the link-state PDU from the link-state database by other routing devices. The recommended maximum link-state PDU origination interval is the lifetime minus 300 seconds. So, in a default environment this would be 900 seconds. In Junos OS, the refresh interval is derived from the lifetime and is equal to the lifetime minus 317 seconds. You can change the lifetime to a higher value to reduce the number of refreshes in the network. (You would rarely want to increase the number of refreshes.) Often these periodic link-state PDU refreshes are referred to as refresh noise, and network administrators want to reduce this noise as much as possible.</p> <p>The show isis overview command displays the link-state PDU lifetime.</p>
Default	By default, link-state PDUs are maintained in network databases for 1200 seconds (20 minutes) before being considered invalid. This length of time, called the <i>LSP lifetime</i> , normally is sufficient to guarantee that link-state PDUs never expire.
Options	<p><i>seconds</i>—link-state PDU lifetime, in seconds.</p> <p>Range: 350 through 65,535 seconds</p> <p>Default: 1200 seconds</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring the Transmission Frequency for Link-State PDUs on IS-IS Interfaces</i> • <i>Understanding Link-State PDU Throttling for IS-IS Interfaces</i> • <i>Example: Configuring the Transmission Frequency for CSNPs on IS-IS Interfaces</i>

- http://www.juniper.net/us/en/training/certification/JNCIP_studyguide.pdf

max-areas

Syntax	<code>max-areas <i>number</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols <i>isis</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code><i>isis</i>]</code> <code>[edit protocols <i>isis</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols <i>isis</i>]</code>
Release Information	Statement introduced in Junos OS Release 8.1. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Modify the maximum number of IS-IS areas advertised.</p> <p>This value is included in the Maximum Address Area field of the IS-IS common PDU header included in all outgoing PDUs.</p> <p>The maximum number of areas you can advertise is restricted to 36 to ensure that the IIH PDUs have enough space to include other type, length, and value (TLV) fields, such as the Authentication and IPv4 and IPv6 Interface Address TLVs.</p>
Options	<p><i>number</i>—Maximum number of areas to include in the IS-IS hello (IIH) PDUs and link-state PDUs.</p> <p>Range: 3 through 36</p> <p>Default: 3</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Understanding IS-IS Areas to Divide an Autonomous System into Smaller Groups• Example: Configuring a Multi-Level IS-IS Topology to Control Interarea Flooding on page 4175

mesh-group (Protocols IS-IS)

Syntax	mesh-group (blocked <i>value</i>);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>], [edit protocols isis interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure an interface to be part of a mesh group, which is a set of fully connected nodes. <i>A mesh group</i> is a set of routing devices that are fully connected. That is, they have a fully meshed topology. When link-state PDUs are being flooded throughout an area, each router within a mesh group receives only a single copy of a link-state PDU instead of receiving one copy from each neighbor, thus minimizing the overhead associated with the flooding of link-state PDUs. To create a mesh group and designate that an interface be part of the group, assign a mesh-group number to all the routing device interfaces in the group. To prevent an interface in the mesh group from flooding link-state PDUs, configure blocking on that interface.
Options	blocked —Configure the interface so that it does not flood link-state PDUs. value —Number that identifies the mesh group. Range: 1 through 4,294,967,295 ($2^{32} - 1$; 32 bits are allocated to identify a mesh group)
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Mesh Groups of IS-IS Interfaces</i> • <i>Understanding IS-IS Mesh Groups</i>

metric (Protocols IS-IS)

Syntax	<code>metric <i>metric</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Specify the metric value for the level.</p> <p>All IS-IS routes have a cost, which is a routing metric that is used in the IS-IS link-state calculation. The cost is an arbitrary, dimensionless integer that can be from 1 through 63, or from 1 through 16,777,215 ($2^{24} - 1$) if you are using wide metrics.</p> <p>Similar to other routing protocols, IS-IS provides a way of exporting routes from the routing table into the IS-IS network. When a route is exported into the IS-IS network without a specified metric, IS-IS uses default metric values for the route, depending on the protocol that was used to learn the route.</p>

Table 332 depicts IS-IS route export metric default values.

Table 332: Default Metric Values for Routes Exported into IS-IS

Protocol Used for Learning the Route	Default Metric Value
Direct	10
Static	Same as reported by the protocol used for exporting the route
Aggregate	10
Generate	10
RIP	Same as reported by the protocol used for exporting the route
OSPF	Same as reported by the protocol used for exporting the route
BGP	10

The default metric values behavior can be customized by using routing policies.

Options `metric`—Metric value.

Range: 1 through 63, or 1 through 16,777,215 (if you have configured wide metrics)

Default: 10 (for all interfaces except lo0), 0 (for the lo0 interface)

Required Privilege	routing—To view this statement in the configuration.
Level	routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Enabling Wide IS-IS Metrics for Traffic Engineering</i> • <i>Understanding Wide IS-IS Metrics for Traffic Engineering</i> • <i>te-metric</i> • wide-metrics-only on page 4327

no-adjacency-holddown

Syntax	no-adjacency-holddown;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis], [edit protocols isis], [edit routing-instances <i>routing-instance-name</i> protocols isis]
Release Information	Statement introduced in Junos OS Release 8.0. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Disable the hold-down timer for IS-IS adjacencies.</p> <p>A hold-down timer delays the advertising of adjacencies by waiting until a time period has elapsed before labeling adjacencies in the up state. You can disable this hold-down timer, which labels adjacencies up faster. However, disabling the hold-down timer creates more frequent link-state PDU updates and SPF computation.</p>
Required Privilege	routing—To view this statement in the configuration.
Level	routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • hold-time on page 4282

no-authentication-check

Syntax	no-authentication-check;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis], [edit protocols isis], [edit routing-instances <i>routing-instance-name</i> protocols isis]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Generate authenticated packets and check the authentication on received packets, but do not reject packets that cannot be authenticated.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• csnp-interval on page 4269• hello-authentication-type on page 4277

no-csnp-authentication

Syntax	no-csnp-authentication;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis level level-number], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis level level-number], [edit protocols isis level level-number], [edit routing-instances <i>routing-instance-name</i> protocols isis level level-number]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Suppress authentication check on complete sequence number PDU (CSNP) packets.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• csnp-interval on page 4269

no-eligible-backup (Protocols IS-IS)

Syntax	no-eligible-backup;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>], [edit protocols isis interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 9.5. Statement introduced in Junos OS Release 9.5 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Exclude the specified interface as a backup interface for IS-IS interfaces on which link protection or node-link protection is enabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding Loop-Free Alternate Routes for IS-IS on page 4165 • Example: Configuring Node-Link Protection for IS-IS Routes in a Layer 3 VPN on page 4221 • link-protection on page 4293 • node-link-protection on page 4308


no-hello-authentication

Syntax	no-hello-authentication;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis level level-number], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis level level-number], [edit protocols isis level level-number], [edit routing-instances <i>routing-instance-name</i> protocols isis level level-number]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Suppress authentication check on complete sequence number hello packets.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • hello-authentication-type on page 4277

no-ipv4-multicast

Syntax	no-ipv4-multicast;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface interface-name], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface interface-name], [edit protocols isis interface interface-name], [edit routing-instances <i>routing-instance-name</i> protocols isis interface interface-name]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Exclude an interface from IPv4 multicast topologies.
Default	Multicast topologies are disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring IS-IS Multicast Topology on page 4203• IS-IS Multicast Topologies Overview on page 4202

no-ipv4-routing

Syntax	no-ipv4-routing;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis],</p> <p>[edit protocols isis],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Disable IP version 4 (IPv4) routing.</p> <p>Disabling IPv4 routing has the following results:</p> <ul style="list-style-type: none"> • The routing device does not advertise the network layer protocol identifier (NLPID) for IPv4 in the Junos OS link-state PDU fragment zero. • The routing device does not advertise any IPv4 prefixes in Junos OS link-state PDUs. • The routing device does not advertise the NLPID for IPv4 in Junos OS hello packets. • The routing device does not advertise any IPv4 addresses in Junos OS hello packets. • The routing device does not calculate any IPv4 routes.
	<div>  <p>NOTE: Note: Even when no-ipv4-routing is configured, an IS-IS traceoptions log can list rejected IPv4 addresses. When a configuration is committed, IS-IS schedules a scan of the routing table to determine whether any routes need to be exported into the IS-IS link state database. The implicit default export policy action is to reject everything. IPv4 addresses from the routing table are examined for export, rejected by the default policy, and the rejections are logged.</p> </div>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies on page 4193 • Understanding IS-IS IPv4 and IPv6 Unicast Topologies on page 4193

no-ipv6-multicast

Syntax	no-ipv6-multicast;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface interface-name], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface interface-name], [edit protocols isis interface interface-name], [edit routing-instances <i>routing-instance-name</i> protocols isis interface interface-name]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Exclude an interface from the IPv6 multicast topologies.
Default	Multicast topologies are disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring IS-IS Multicast Topology on page 4203• IS-IS Multicast Topologies Overview on page 4202

no-ipv6-routing

Syntax	no-ipv6-routing;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis],</p> <p>[edit protocols isis],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	<p>Disable IP version 6 (IPv6) routing.</p> <p>Disabling IPv6 routing has the following results:</p> <ul style="list-style-type: none"> • The routing device does not advertise the network layer protocol identifier (NLPID) for IPv6 in the Junos OS link-state PDU fragment zero. • The routing device does not advertise any IPv6 prefixes in Junos OS link-state PDUs. • The routing device does not advertise the NLPID for IPv6 in Junos OS hello packets. • The routing device does not advertise any IPv6 addresses in Junos OS hello packets. • The routing device does not calculate any IPv6 routes.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies on page 4193 • Understanding IS-IS IPv4 and IPv6 Unicast Topologies on page 4193

no-ipv6-unicast

Syntax	no-ipv6-unicast;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface interface-name], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface interface-name], [edit protocols isis interface interface-name], [edit routing-instances <i>routing-instance-name</i> protocols isis interface interface-name]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Exclude an interface from the IPv6 unicast topologies. This enables you to exercise control over the paths that unicast data takes through a network.
Default	IPv6 unicast topologies are disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies on page 4193• Understanding IS-IS IPv4 and IPv6 Unicast Topologies on page 4193

no-psnp-authentication

Syntax	no-psnp-authentication;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis level level-number], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis level level-number], [edit protocols isis level level-number], [edit routing-instances <i>routing-instance-name</i> protocols isis level level-number]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Suppress authentication check on partial sequence number PDU (PSNP) packets.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring IS-IS Authentication on page 4183

no-unicast-topology

Syntax	no-unicast-topology;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface interface-name], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface interface-name], [edit protocols isis interface interface-name], [edit routing-instances <i>routing-instance-name</i> protocols isis interface interface-name]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Exclude an interface from the IPv4 unicast topologies.
Default	IPv4 unicast topologies are disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring IS-IS Multicast Topology on page 4203 • IS-IS Multicast Topologies Overview on page 4202

node-link-protection (Protocols IS-IS)

Syntax	node-link-protection;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>], [edit protocols isis interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 9.5. Statement introduced in Junos OS Release 9.5 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Enable node-link protection on the specified IS-IS interface. Junos OS creates an alternate loop-free path to the primary next hop for all destination routes that traverse a protected interface. This alternate path avoids the primary next-hop routing device altogether and establishes a path through a different routing device.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding Loop-Free Alternate Routes for IS-IS on page 4165• Example: Configuring Node-Link Protection for IS-IS Routes in a Layer 3 VPN on page 4221• link-protection on page 4293

overload (Protocols IS-IS)

Syntax	<pre> overload { advertise-high-metrics; allow-route-leaking; timeout <i>seconds</i>; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols <i>isis</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>isis</i>],</p> <p>[edit protocols <i>isis</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <i>isis</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Configure the local routing device so that it appears to be overloaded. This statement causes the routing device to continue participating in IS-IS routing, but prevents it from being used for transit traffic. Traffic destined to immediately attached subnets continues to transit the routing device.</p> <p>You can also advertise maximum link metrics in network layer reachability information (NLRI) instead of setting the overload bit.</p> <p>You configure or disable overload mode in IS-IS with or without a timeout. Without a timeout, overload mode is set until it is explicitly deleted from the configuration. With a timeout, overload mode is set if the time elapsed since the IS-IS instance started is less than the specified timeout.</p> <p>A timer is started for the difference between the timeout and the time elapsed since the instance started. If the time elapsed after the IS-IS instance is enabled is less than the specified timeout, overload mode is set. When the timer expires, overload mode is cleared. In overload mode, the routing device IS-IS advertisements are originated with the overload bit set. This causes the transit traffic to take paths around the routing device. However, the overloaded routing device's own links are still accessible.</p> <p>The value of the overload bit depends on these three scenarios:</p> <ol style="list-style-type: none"> 1. When the overload bit has already been set to a given value and the routing process is restarted: Link-state PDUs are regenerated with the overload bit cleared. 2. When the overload bit is reset to a lesser value while the routing process is running: Link-state PDUs are regenerated with the overload bit cleared. 3. When the overload bit is reset to a greater value while the routing process is running: Link-state PDUs are regenerated with the overload bit set to the difference between the old and new value.

In overload mode, the routing device advertisement is originated with all the transit routing device links (except stub) set to a metric of 0xFFFF. The stub routing device links are advertised with the actual cost of the interfaces corresponding to the stub. This causes the transit traffic to avoid the overloaded routing device and take paths around the routing device.

To understand the reason for setting the overload bit, consider that BGP converges slowly. It is not very good at detecting that a neighbor is down because it has slow-paced keepalive timers. Once the BGP neighbor is determined to be down, it can take up to 2 minutes for a BGP router to declare the neighbor down. IS-IS is much quicker. IS-IS only takes 10-30 seconds to detect absent peers. It is the slowness of BGP, more precisely the slowness of internal BGP (IBGP), that necessitates the use of the overload bit. IS-IS and BGP routing are mutually dependent on each other. If both do not converge at the same time, traffic is dropped without notification (black holed).

You might want to configure the routing device so that it appears to be overloaded when you are restarting routing on the device. Setting the overload bit for a fixed amount of time right after a restart of the routing protocol process (rpd) ensures that the router does not receive transit traffic while the routing protocols (especially IBGP) are still converging.

Setting the overload bit is useful when performing hardware or software maintenance work on a routing device. After the maintenance work, clear the overload bit to carry on forwarding transit traffic. Manual clearing of the overload bit is not always possible. What is needed is an automated way of clearing the overload bit after some amount of time. Most networks use a time value of 300 seconds. This 5-minute value provides a good balance, allowing time to bring up even large internal IBGP meshes, while still relatively quick.

Another appropriate application for setting for the overload bit is on dedicated devices such as BGP route reflectors, which are intentionally not meant to carry any transit traffic. In this case, you would not use the timer.

You can verify that the overload bit is set by running the **show isis database** command.

Options **advertise-high-metrics**—Advertise maximum link metrics in NLRI's instead of setting the overload bit.

The **advertise-high-metric** setting is only valid while the routing device is in overload mode. When **advertise-high-metric** is configured, IS-IS does not set the overload bit. Rather, it sets the metric to 63 or 16,777,214, depending whether wide metrics are enabled. This allows the overloaded routing device to be used for transit as a last resort.

An L1-L2 router in overload mode stops leaking route information between L1 and L2 levels and clears its attached bit. This is also true when **advertise-high-metrics** is configured.

allow-route-leaking—Enable leaking of route information into the network even if the overload bit is set.



NOTE: The **allow-route-leaking** option does not work if the routing device is in dynamic overload mode. Dynamic overload can occur if the device has exceeded its resource limits, such as the prefix limit.


timeout seconds—Number of seconds at which the overloading is reset.

Range: 60 through 1800 seconds

Default: 0 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

passive (Protocols IS-IS)

Syntax	<pre> passive { remote-node-id <i>address</i>; remote-node-iso <i>iso-id</i>; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit protocols isis interface <i>interface-name</i>],</p> <p>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>remote-node-id <i>address</i> option introduced in Junos OS Release 14.2.</p> <p>remote-node-iso <i>iso-id</i> option introduced in Junos OS Release 14.2.</p>
Description	<p>Advertise the direct interface addresses on an interface or into a level on the interface without actually running IS-IS on that interface or level.</p> <p>This statement effectively prevents IS-IS from running on the interface. To enable IS-IS on an interface, include the interface statement at the [edit protocols isis] or the [edit routing-instances <i>routing-instance-name</i> protocols isis] hierarchy level. To disable it, include the disable statement at those hierarchy levels. The three states—enabling, disabling, or not running IS-IS on an interface—are mutually exclusive.</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p> NOTE: Configuring IS-IS on a loopback interface automatically renders it as a passive interface, irrespective of whether the passive statement was used in the configuration of the interface.</p> </div> <p>If neither passive mode nor the family iso option is configured on the IS-IS interface, then the routing device treats the interface as not being operational, and no direct IPv4/IPv6 routes are exported into IS-IS. (You configure the family iso option at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] hierarchy level.)</p>
Default	By default, IS-IS must be configured on an interface or a level for direct interface addresses to be advertised into that level.
Options	remote-node-id <i>address</i> —IP address of the remote link.

remote-node-iso *iso-id*—ISO ID of the remote node.



NOTE: The options **remote-node-id** *address* and **remote-node-iso** *iso-id* do not apply under the [edit routing-instances *routing-instance-name* protocols isis] hierarchy level.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [Example: Configuring a Multi-Level IS-IS Topology to Control Interarea Flooding on page 4175](#)
- *disable*

point-to-point

Syntax point-to-point;

Hierarchy Level [edit logical-systems *logical-system-name* protocols isis **interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis **interface** *interface-name*],
[edit protocols isis **interface** *interface-name*],
[edit routing-instances *routing-instance-name* protocols isis **interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure an IS-IS interface to behave like a point-to-point connection.

You can use the **point-to-point** statement to configure a LAN interface to act like a point-to-point interface for IS-IS. You do not need an unnumbered LAN interface, and it has no effect if configured on an interface that is already point-to-point.

The **point-to-point** statement affects only IS-IS protocol procedures on that interface. All other protocols continue to treat the interface as a LAN interface. Only two IS-IS routing devices can be connected to the LAN interface, and both must be configured as point-to-point.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [IS-IS Overview on page 4157](#)
- [Example: Configuring Synchronization Between IS-IS and LDP](#)
- [Understanding LDP-IGP Synchronization](#)
- [Understanding IS-IS Designated Routers](#)

preference (Protocols IS-IS)

Syntax	<code>preference <i>preference</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis level level-number], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis level level-number], [edit protocols isis level level-number], [edit routing-instances <i>routing-instance-name</i> protocols isis level level-number]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Configure the preference of internal routes.</p> <p>Route preferences (also known as administrative distances) are used to select which route is installed in the forwarding table when several protocols calculate routes to the same destination. The route with the lowest preference value is selected.</p> <p>To change the preference values, include the preference statement (for internal routes) or the external-preference statement.</p>
Options	<p>preference—Preference value.</p> <p>Range: 0 through 4,294,967,295 ($2^{32} - 1$)</p> <p>Default: 15 (for Level 1 internal routes), 18 (for Level 2 internal routes), 160 (for Level 1 external routes), 165 (for Level 2 external routes)</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Route Preferences Overview</i>• <i>Example: Configuring a Routing Policy to Redistribute BGP Routes with a Specific Community Tag into IS-IS</i>• Example: Redistributing OSPF Routes into IS-IS on page 4185• <i>Understanding Routing Policies</i>• <i>Understanding BGP Communities and Extended Communities as Routing Policy Match Conditions</i>• external-preference on page 4272

prefix-export-limit (Protocols IS-IS)

Syntax	<code>prefix-export-limit <i>number</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis level level-number],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis level level-number],</p> <p>[edit protocols isis level level-number],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis level level-number]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Configure a limit to the number of prefixes exported into IS-IS.</p> <p>By default, there is no limit to the number of prefixes that can be exported into IS-IS. To configure a limit to the number of prefixes that can be exported into IS-IS, include the prefix-export-limit statement. The prefix-export-limit statement protects the rest of the network from a malicious policy by applying a threshold filter for exported routes.</p> <p>The number of prefixes depends on the size of your network. Good design advice is to set it to double the total number of IS-IS Level 1 and Level 2 routing devices in your network.</p> <p>If the number of prefixes exported into IS-IS exceeds the configured limit, the overload bit is set and the overload state is reached. When other routers detect that this bit is set, they do not use this routing device for transit traffic, but they do use it for packets destined to the overloaded routing device's directly connected networks and IP prefixes. The overload state can be cleared by using the clear isis overload command.</p> <p>The show isis overview command displays the prefix export limit when it is configured.</p>
Options	<p><i>number</i>—Prefix limit.</p> <p>Range: 0 through 4,294,967,295 ($2^{32} - 1$)</p> <p>Default: None</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Redistributing OSPF Routes into IS-IS on page 4185 • Example: Configuring a Routing Policy to Redistribute BGP Routes with a Specific Community Tag into IS-IS • Understanding BGP Communities and Extended Communities as Routing Policy Match Conditions • Understanding Routing Policies

priority (Protocols IS-IS)

Syntax	<code>priority <i>number</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</code> <code>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Configure the interface's priority for becoming the designated router. The interface with the highest priority value becomes that level's designated router.</p> <p>The priority value is meaningful only on a multiaccess network. It has no meaning on a point-to-point interface.</p> <p>A routing device advertises its priority to become a designated router in its hello packets. On all multiaccess networks, IS-IS uses the advertised priorities to elect a designated router for the network. This routing device is responsible for sending network link-state advertisements, which describe all the routing devices attached to the network. These advertisements are flooded throughout a single area.</p> <p>A routing device's priority for becoming the designated router is indicated by an arbitrary number from 0 through 127. Routing devices with a higher value are more likely to become the designated router.</p>
Options	<i>number</i> —Priority value. Range: 0 through 127 Default: 64
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring IS-IS Designated Routers</i>

reference-bandwidth (Protocols IS-IS)

Syntax	<code>reference-bandwidth <i>reference-bandwidth</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis],</p> <p>[edit protocols isis],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Optimize routing based on bandwidth by setting the reference bandwidth used in calculating the default interface cost.</p> <p>All IS-IS interfaces have a cost, which is a routing metric that is used in the IS-IS link-state calculation. Routes with lower total path metrics are preferred over those with higher path metrics. When there are several equal-cost routes to a destination, traffic is distributed equally among them.</p> <p>The cost of a route is described by a single dimensionless metric that is determined using the following formula:</p> $\text{cost} = \text{reference-bandwidth} / \text{bandwidth}$ <p>For example, if you set the reference bandwidth to 1 Gbps (that is, <i>reference-bandwidth</i> is set to 1,000,000,000), a 100-Mbps interface has a routing metric of 10.</p> <p>All IS-IS interfaces have a cost, which is a routing metric that is used in the IS-IS link-state calculation. Routes with lower total path metrics are preferred over those with higher path metrics.</p>
Options	<p><i>reference-bandwidth</i>—Reference bandwidth value in bits per second.</p> <p>Range: 9600 through 1,000,000,000,000 bps</p> <p>Default: None</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Understanding IS-IS Configuration</i> • Example: Configuring IS-IS on page 4169 • http://www.juniper.net/us/en/training/certification/JNCIP_studyguide.pdf

rib-group (Protocols IS-IS)

Syntax	<pre>rib-group { inet <i>group-name</i>; inet6 <i>group-name</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis], [edit protocols isis], [edit routing-instances <i>routing-instance-name</i> protocols isis]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Install routes learned from IS-IS routing instances into routing tables in the IS-IS routing table group. You can install IPv4 routes or IPv6 routes.</p> <p>Support for IPv6 routing table groups in IS-IS enables IPv6 routes that are learned from IS-IS routing instances to be installed into other routing tables defined in an IS-IS routing table group.</p>
Options	<p><i>group-name</i>—Name of the routing table group.</p> <p>inet—Install IPv4 IS-IS routes.</p> <p>inet6—Install IPv6 IS-IS routes.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Example: Exporting Specific Routes from One Routing Table Into Another Routing Table</i>• <i>Example: Importing Direct and Static Routes Into a Routing Instance</i>• <i>Understanding Multiprotocol BGP</i>

spf-options (Protocols IS-IS)

Syntax	<pre>spf-options { delay <i>milliseconds</i>; holddown <i>milliseconds</i>; rapid-runs <i>number</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols <i>isis</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>isis</i>], [edit protocols <i>isis</i>], [edit routing-instances <i>routing-instance-name</i> protocols <i>isis</i>]
Release Information	Statement introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	<p>Configure options for running the shortest-path-first (SPF) algorithm.</p> <p>Running the SPF algorithm is usually the beginning of a series of larger system-wide events. For example, the SPF algorithm can lead to interior gateway protocol (IGP) prefix changes, which then lead to BGP nexthop resolution changes. Consider what happens if there are rapid link changes in the network. The local routing device can become overwhelmed. This is why it sometimes makes sense to throttle the scheduling of the SPF algorithm.</p> <p>You can configure the following SPF options:</p> <ul style="list-style-type: none"> • The delay in the time between the detection of a topology change and when the SPF algorithm actually runs. • The maximum number of times that the SPF algorithm can run in succession before the hold-down timer begins. • The time to hold down, or wait, before running another SPF calculation after the SPF algorithm has run in succession the configured maximum number of times. <p>If the network stabilizes during the hold-down period and the SPF algorithm does not need to run again, the system reverts to the configured values for the delay and rapid-runs statements.</p>
Options	<p>delay <i>milliseconds</i>—Time interval between the detection of a topology change and when the SPF algorithm runs.</p> <p>Range: 50 through 1000 milliseconds</p> <p>Default: 200 milliseconds</p> <p>holddown <i>milliseconds</i>—Time interval to hold down, or wait before a subsequent SPF algorithm runs after the SPF algorithm has run the configured maximum number of times in succession.</p> <p>Range: 2000 through 10,000 milliseconds</p> <p>Default: 5000 milliseconds</p>

rapid-runs *number*—Maximum number of times the SPF algorithm can run in succession.
After the maximum is reached, the holddown interval begins.

Range: 1 through 5

Default: 3

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [Understanding Loop-Free Alternate Routes for IS-IS on page 4165](#)
- [Example: Configuring Node-Link Protection for IS-IS Routes in a Layer 3 VPN on page 4221](#)

topologies (Protocols IS-IS)

Syntax

```
topologies {  
  ipv4-multicast;  
  ipv6-multicast;  
  ipv6-unicast;  
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols *isis*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols *isis*],
[edit protocols *isis*],
[edit routing-instances *routing-instance-name* protocols *isis*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure alternate IS-IS topologies.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies on page 4193](#)
- [Example: Configuring IS-IS Multicast Topology on page 4203](#)
- [IS-IS Multicast Topologies Overview on page 4202](#)
- [Understanding IS-IS IPv4 and IPv6 Unicast Topologies on page 4193](#)

traceoptions (Protocols IS-IS)

Syntax	<pre>traceoptions { file <i>name</i> <size <i>size</i>> <files <i>number</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier> <disable>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis], [edit protocols isis], [edit routing-instances <i>routing-instance-name</i> protocols isis]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure IS-IS protocol-level tracing options. To specify more than one tracing operation, include multiple flag statements.



NOTE: The **traceoptions** statement is not supported on QFabric systems.

Default	The default IS-IS protocol-level tracing options are those inherited from the routing protocols traceoptions statement included at the [edit routing-options] hierarchy level.
Options	<p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>name</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks (" "). All files are placed in the directory /var/log. We recommend that you place IS-IS tracing output in the file isis-log.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>Range: 2 through 1000 files</p> <p>Default: 10 files</p> <p>flag <i>flag</i>—Tracing operation to perform. To specify more than one flag, include multiple flag statements.</p>

IS-IS Protocol-Specific Tracing Flags

- **csn**—Complete sequence number PDU (CSNP) packets
- **error**—Errored IS-IS packets
- **graceful-restart**—Graceful restart operation
- **hello**—Hello packets
- **ldp-synchronization**—Synchronization between IS-IS and LDP
- **lsp**—Link-state PDUs
- **lsp-generation**—Link-state PDU generation packets
- **packets**—All IS-IS protocol packets
- **psn**—Partial sequence number PDU (PSNP) packets
- **spf**—Shortest-path-first calculations

Global Tracing Flags

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations, including adjacency changes

Default: If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information.
- **receive**—Trace the packets being received.
- **send**—Trace the packets being transmitted.

no-world-readable—(Optional) Prevent any user from reading the log file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 128 KB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level	routing and trace—To view this statement in the configuration. routing-control and trace-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring the Transmission Frequency for CSNPs on IS-IS Interfaces</i> • Example: Enabling Packet Checksums on IS-IS Interfaces for Error Checking on page 4239 • <i>Example: Configuring the Transmission Frequency for Link-State PDUs on IS-IS Interfaces</i> • <i>Understanding Link-State PDU Throttling for IS-IS Interfaces</i> • <i>Understanding Checksums on IS-IS Interfaces for Error Checking</i>

traffic-engineering (Protocols IS-IS)

Syntax	<pre> traffic-engineering { disable; credibility-protocol-preference; family inet { shortcuts { multicast-rpf-routes; } } family inet6 { shortcuts; } multipath { lsp-equal-cost; } } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis],</p> <p>[edit protocols isis]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Support for the family statement introduced in Junos OS Release 9.3.</p> <p>Support for the credibility-protocol-preference statement introduced in Junos OS Release 9.4.</p> <p>Support for the multipath statement introduced in Junos OS Release 9.6.</p> <p>Support for the lsp-equal-cost statement introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Configure traffic engineering properties for IS-IS.</p> <p>IS-IS always performs shortest-path-first (SPF) calculations to determine next hops. For prefixes reachable through a particular next hop, IS-IS places that next hop for that prefix in the inet.0 routing table. In addition, for routers running MPLS, IS-IS installs the prefix for IPv4 routes in the inet.3 routing table as well. The inet.3 table, which is present on the ingress router, contains the host address of each MPLS label-switched path (LSP) egress router. BGP uses this routing table to resolve next-hop addresses.</p> <p>If you enable IS-IS traffic engineering shortcuts and if there is a label-switched path to a point along the path to that prefix, IS-IS installs the prefix in the inet.3 routing table and uses the LSP as a next hop. The net result is that for BGP egress routers for which there is no LSP, BGP automatically uses an LSP along the path to reach the egress router.</p> <p>In Junos OS Release 9.3 and later, IS-IS traffic engineering shortcuts support IPv6 routes. LSPs to be used for shortcuts continue to be signaled using IPv4. However, by default, shortcut routes calculated through IPv6 routes are added to the inet6.3 routing table. The default behavior is for only BGP to use LSPs in its calculations. If you configure MPLS so that both BGP and interior gateway protocols use LSPs for forwarding traffic, shortcut routes calculated through IPv6 are added to the inet6.0 routing table. IS-IS ensures that the IPv6 routes running over the IPv4 MPLS LSP are correctly de-encapsulated at the</p>

tunnel egress by pushing an extra IPv6 explicit null label between the IPv6 payload and the IPv4 transport label.

RSVP LSPs with a higher preference than IS-IS routes are not considered during the computation of traffic engineering shortcuts.

To configure IS-IS so that it uses LSPs as shortcuts when installing information in the inet.3 or inet6.3 routing table, include the following statements:

```
family inet {
  shortcuts {
    multicast-rpf-routes;
  }
}
family inet6 {
  shortcuts;
}
```

For IPv4 traffic, include the **inet** statement. For IPv6 traffic, include the **inet6** statement.

To configure load balancing across multiple LSPs, include the **multipath** statement.

When traffic engineering shortcuts are used, RSVP first looks at the **metric2** value, which is derived from the IGP cost. After this, RSVP considers the LSP metric value. So, if a certain path changes for an LSP and the cost changes, not all LSPs are used to load-balance the network.

When a route with an improved metric is added to the IS-IS internal routing table, IS-IS flushes all next-hop information (including LSP next-hop information) for a route. This is undesirable, because certain equal-cost multipath (ECMP) combinations can be lost during route calculation. To override this default behavior for load balancing, include the **lsp-equal-cost** statement to retain the equal cost path information in the routing table.

```
multipath {
  lsp-equal-cost;
}
```

Because the inet.3 routing table is present only on ingress routers, you can configure LSP shortcuts only on these routers.

Default IS-IS traffic engineering support is enabled.

By default, IS-IS supports traffic engineering by exchanging basic information with the traffic engineering database. To disable this support, and to disable IS-IS shortcuts if they are configured, include the **disable** statement.

Options **credibility-protocol-preference**—Specify that IS-IS should use the configured protocol preference for IGP routes to determine the traffic engineering database credibility value. By default, the traffic engineering database prefers IS-IS routes even when the routes of another IGP are configured with a lower, that is, more preferred value. Use this statement to override this default behavior.

The traffic engineering database assigns a credibility value to each IGP and prefers the routes of the IGP with the highest credibility value. In Junos OS Release 9.4 and later, you can configure IS-IS to take protocol preference into account to determine the traffic engineering database credibility value. When protocol preference is used to determine the credibility value, IS-IS routes are not automatically preferred by the traffic engineering database, depending on your configuration. For example, OSPF routes have a default preference value of 10, whereas IS-IS Level 1 routes have a default preference value of 15. When protocol preference is enabled, the credibility value is determined by deducting the protocol preference value from a base value of 512. Using default protocol preference values, OSPF has a credibility value of 502, whereas IS-IS has a credibility value of 497. Because the traffic engineering database prefers IGP routes with the highest credibility value, OSPF routes are now preferred.



NOTE: This feature is also supported for OSPFv2.

lsp-equal-cost—Configure LSPs to be retained as equal cost paths for load balancing when a better path metric is found during the IS-IS internal routing table calculation. When a route with an improved metric is added to the IS-IS internal routing table, IS-IS flushes all next-hop information (including LSP next-hop information) for a route. This is undesirable, because certain equal-cost multipath (ECMP) combinations can be lost during route calculation. To override this default IS-IS behavior, include the **lsp-equal-cost** statement for load balancing, so that the equal cost path information is retained in the routing table.

multipath—Enable load balancing for multiple LSPs.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
 routing-control—To add this statement to the configuration.

Related Documentation

- [Example: Enabling OSPF Traffic Engineering Support on page 4513](#)
- [Example: Enabling IS-IS Traffic Engineering Support](#)
- [traffic-engineering \(OSPF\) on page 4614](#)
- [Using Labeled-Switched Paths to Augment SPF to Compute IGP Shortcuts](#)

wide-metrics-only

Syntax	wide-metrics-only;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis level <i>level-number</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis level <i>level-number</i>], [edit protocols isis level <i>level-number</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis level <i>level-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure IS-IS to generate metric values greater than 63 on a per IS-IS level basis. Normally, IS-IS metrics can have values up to 63, and IS-IS generates two type, length, and value (TLV) tuples, one for an IS-IS adjacency and the second for an IP prefix. To allow IS-IS to support traffic engineering, a second pair of TLVs has been added to IS-IS, one for IP prefixes and the second for IS-IS adjacency and traffic engineering information. With these TLVs, IS-IS metrics can have values up to 16,777,215 ($2^{24} - 1$). To configure IS-IS to generate only the new pair of TLVs and thus to allow the wider range of metric values, include the wide-metrics-only statement.
Default	By default, Junos OS supports the sending and receiving of wide metrics. Junos OS allows a maximum metric value of 63 and generates both pairs of TLVs.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> <i>te-metric</i> <i>Example: Enabling Wide IS-IS Metrics for Traffic Engineering</i> <i>Understanding Wide IS-IS Metrics for Traffic Engineering</i>

CHAPTER 161

Operational Commands

- clear isis adjacency
- clear isis database
- clear isis overload
- clear isis statistics
- show isis adjacency
- show isis authentication
- show isis backup coverage
- show isis backup label-switched-path
- show isis backup spf results
- show isis database
- show isis hostname
- show isis interface
- show isis overview
- show isis route
- show isis spf
- show isis statistics

clear isis adjacency

List of Syntax	Syntax on page 4330 Syntax (EX Series Switches and QFX Series) on page 4330
Syntax	<pre>clear isis adjacency <all> <instance <i>instance-name</i>> <interface <i>interface-name</i>> <logical-system (all <i>logical-system-name</i>)> <neighbor></pre>
Syntax (EX Series Switches and QFX Series)	<pre>clear isis adjacency <all> <instance <i>instance-name</i>> <interface <i>interface-name</i>> <neighbor></pre>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series. all option introduced in Junos OS Release 14.2.
Description	Remove entries from the IS-IS adjacency database.
Options	<p>none all—(Optional) Remove all entries from the adjacency database. Both clear isis adjacency and clear isis adjacency all function identically.</p> <p>instance <i>instance-name</i>—(Optional) Clear all adjacencies for the specified routing instance only.</p> <p>interface <i>interface-name</i>—(Optional) Clear all adjacencies for the specified interface only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>neighbor—(Optional) Clear adjacencies for the specified neighbor only.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show isis adjacency on page 4338
List of Sample Output	clear isis adjacency on page 4331 clear isis adjacency all on page 4331
Output Fields	See show isis adjacency for an explanation of output fields.

Sample Output

clear isis adjacency

The following sample output displays IS-IS adjacency database information before and after the **clear isis adjacency** command is entered:

```
user@host> show isis adjacency
IS-IS adjacency database:
Interface      System          L State          Hold (secs) SNPA
so-1/0/0.0     karaku1         3 Up             26
so-1/1/3.0     1921.6800.5080 3 Up             23
so-5/0/0.0     1921.6800.5080 3 Up             19
```

```
user@host> clear isis adjacency karaku1
```

```
user@host> show isis adjacency
IS-IS adjacency database:
Interface      System          L State          Hold (secs) SNPA
so-1/0/0.0     karaku1         3 Initializing   26
so-1/1/3.0     1921.6800.5080 3 Up             24
so-5/0/0.0     1921.6800.5080 3 Up             21
```

clear isis adjacency all

```
user@host> clear isis adjacency all
IS-IS adjacency database:
Interface      System          L State          Hold (secs) SNPA
so-1/0/0.0     karaku1         3 Initializing   26
so-1/1/3.0     1921.6800.5080 3 Initializing   24
so-5/0/0.0     1921.6800.5080 3 Initializing   21
```

clear isis database

List of Syntax	Syntax on page 4332 Syntax (EX Series Switches and QFX Series) on page 4332
Syntax	<pre>clear isis database <entries> <instance instance-name> <logical-system (all logical-system-name)></pre>
Syntax (EX Series Switches and QFX Series)	<pre>clear isis database <entries> <instance instance-name></pre>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series. Command introduced in 15.1X53-D30 for QFX10002 switch.
Description	Remove the entries from the IS-IS link-state database, which contains prefixes and topology information.
Options	none —Remove all entries from the IS-IS link-state database for all routing instances. entries —(Optional) Name of the database entry. instance instance-name —(Optional) Clear all entries for the specified routing instance. logical-system (all logical-system-name) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show isis database on page 4352
List of Sample Output	clear isis database on page 4332
Output Fields	See show isis database for an explanation of output fields.

Sample Output

clear isis database

The following sample output displays IS-IS link-state database information before and after the **clear isis database** command is entered:

```
user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                Sequence Checksum Lifetime (secs)
crater.00-00          0x12    0x84dd             1139
```

```
1 LSPs
IS-IS level 2 link-state database:
LSP ID          Sequence Checksum Lifetime (secs)
crater.00-00     0x19   0xe92c           1134
badlands.00-00   0x16   0x1454           985
carlsbad.00-00   0x33   0x220b          1015
ranier.00-00     0x2e   0xfc31          1007
1921.6800.5066.00-00 0x11   0x7313           566
1921.6800.5067.00-00 0x14   0xd9d4           939
6 LSPs
```

```
user@host> clear isis database
```

```
user@host> show isis database
IS-IS level 1 link-state database:
LSP ID          Sequence Checksum Lifetime (secs)

IS-IS level 2 link-state database:
LSP ID          Sequence Checksum Lifetime (secs)
```

clear isis overload

List of Syntax	Syntax on page 4334 Syntax (EX Series Switches and QFX Series) on page 4334
Syntax	<code>clear isis overload</code> <code><instance <i>instance-name</i>></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switches and QFX Series)	<code>clear isis overload</code> <code><instance <i>instance-name</i>></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Reset the IS-IS dynamic overload bit. This command can appear to not work, continuing to display overload after execution. The bit is reset only if the root cause is corrected by configuration remotely or locally.</p> <p>When other routers detect that the overload bit is set, they do not use this routing device for transit traffic, but they do use it for packets destined to the overloaded routing device's directly connected networks and IP prefixes.</p>
Options	<p>none—Reset the IS-IS dynamic overload bit.</p> <p>instance <i>instance-name</i>—(Optional) Reset the IS-IS dynamic overload bit for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show isis database on page 4352
List of Sample Output	clear isis overload on page 4334
Output Fields	See show isis database for an explanation of output fields.

Sample Output

clear isis overload

The following sample output displays IS-IS database information before and after the **clear isis overload** command is entered:

```
user@host> show isis database
```

IS-IS level 1 link-state database:

LSP ID	Sequence	Checksum	Lifetime	Attributes
pro3-c.00-00	0x4	0x10db	1185	L1 L2 Overload

1 LSPs

IS-IS level 2 link-state database:

LSP ID	Sequence	Checksum	Lifetime	Attributes
pro3-c.00-00	0x5	0x429f	1185	L1 L2 Overload

pro2-a.00-00	0x91e	0x2589	874	L1 L2
--------------	-------	--------	-----	-------

pro2-a.02-00	0x1	0xcbc	874	L1 L2
--------------	-----	-------	-----	-------

3 LSPs

user@host> clear isis overload

user@host> show isis database

IS-IS level 1 link-state database:

LSP ID	Sequence	Checksum	Lifetime	Attributes
pro3-c.00-00	0xa	0x429e	1183	L1 L2

1 LSPs

IS-IS level 2 link-state database:

LSP ID	Sequence	Checksum	Lifetime	Attributes
pro3-c.00-00	0xc	0x9c39	1183	L1 L2

pro2-a.00-00	0x91e	0x2589	783	L1 L2
--------------	-------	--------	-----	-------

pro2-a.02-00	0x1	0xcbc	783	L1 L2
--------------	-----	-------	-----	-------

3 LSPs

clear isis statistics

List of Syntax	Syntax on page 4336 Syntax (EX Series Switches and QFX Series) on page 4336
Syntax	<code>clear isis statistics</code> <code><instance <i>instance-name</i>></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switches and QFX Series)	<code>clear isis statistics</code> <code><instance <i>instance-name</i>></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Set statistics about IS-IS traffic to zero.
Options	none —Set IS-IS traffic statistics to zero for all routing instances. instance <i>instance-name</i> —(Optional) Set IS-IS traffic statistics to zero for the specified routing instance only. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show isis statistics on page 4378
List of Sample Output	clear isis statistics on page 4336
Output Fields	See show isis statistics for an explanation of output fields.

Sample Output

clear isis statistics

The following sample output displays IS-IS statistics before and after the **clear isis statistics** command is entered:

```
user@host> show isis statistics
IS-IS statistics for merino:
```

PDU type	Received	Processed	Drops	Sent	Rexmit
LSP	12793	12793	0	8666	719
IIH	116751	116751	0	118834	0
CSNP	203956	203956	0	204080	0
PSNP	7356	7350	6	8635	0
Unknown	0	0	0	0	0


```
Totals          340856    340850         6    340215    719
```

```
Total packets received: 340856 Sent: 340934
```

```
SNP queue length:      0 Drops:      0
LSP queue length:      0 Drops:      0
```

```
SPF runs:              1064
Fragments rebuilt:     1087
LSP regenerations:     436
Purges initiated:      0
```

```
user@host> clear isis statistics
```

```
user@host> show isis statistics
IS-IS statistics for merino:
```

PDU type	Received	Processed	Drops	Sent	Rexmit
LSP	0	0	0	0	0
IIH	3	3	0	3	0
CSNP	2	2	0	4	0
PSNP	0	0	0	0	0
Unknown	0	0	0	0	0
Totals	5	5	0	7	0

```
Total packets received: 5 Sent: 7
```

```
SNP queue length:      0 Drops:      0
LSP queue length:      0 Drops:      0
```

```
SPF runs:              0
Fragments rebuilt:     0
LSP regenerations:     0
Purges initiated:      0
```

show isis adjacency

List of Syntax	Syntax on page 4338 Syntax (EX Series Switches and QFX Series) on page 4338
Syntax	<pre>show isis adjacency <system-id> <brief detail extensive> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switches and QFX Series)	<pre>show isis adjacency <system-id> <brief detail extensive> <instance <i>instance-name</i>></pre>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Display information about IS-IS neighbors.
Options	<p>none—Display standard information about IS-IS neighbors for all routing instances.</p> <p><i>system id</i>—(Optional) Display information about IS-IS neighbors for the specified intermediate system.</p> <p><i>brief detail extensive</i>—(Optional) Display standard information about IS-IS neighbors with the specified level of output.</p> <p><i>instance instance-name</i>—(Optional) Display information about IS-IS neighbors for the specified routing instance.</p> <p><i>logical-system (all logical-system-name)</i>—(Optional) Display information about IS-IS neighbors for all logical systems or for a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• clear isis adjacency on page 4330
List of Sample Output	show isis adjacency on page 4340 show isis adjacency brief on page 4340 show isis adjacency detail on page 4341 show isis adjacency extensive on page 4341
Output Fields	Table 333 describes the output fields for the show isis adjacency command. Output fields are listed in the approximate order in which they appear.

Table 333: show isis adjacency Output Fields

Field Name	Field Description	Level of Output
Interface	Interface through which the neighbor is reachable.	All levels
System	System identifier (sysid), displayed as a name, if possible.	brief
L or Level	Level: <ul style="list-style-type: none"> • 1—Level 1 only • 2—Level 2 only • 3—Level 1 and Level 2 An exclamation point (!) preceding the level number indicates that the adjacency is missing an IP address.	All levels
State	State of the adjacency: Up , Down , New , One-way , Initializing , or Rejected .	All levels
Hold (secs)	Remaining hold time of the adjacency.	brief
SNPA	Subnetwork point of attachment (MAC address of the next hop).	brief
Expires in	How long until the adjacency expires, in seconds.	detail
Priority	Priority to become the designated intermediate system.	detail extensive
Up/Down transitions	Count of adjacency status changes from Up to Down or from Down to Up .	detail
Last transition	Time of the last Up/Down transition.	detail
Circuit type	Bit mask of levels on this interface: 1=Level 1 router; 2=Level 2 router; 3=both Level 1 and Level 2 router.	detail
Speaks	Protocols supported by this neighbor.	detail extensive
MAC address	MAC address of the interface.	detail extensive
Topologies	Supported topologies.	detail extensive
Restart capable	Whether a neighbor is capable of graceful restart: Yes or No .	detail extensive
Adjacency advertisement: Advertise	This routing device has signaled to advertise this interface to its neighbors in their link-state PDUs.	detail extensive
Adjacency advertisement: Suppress	This neighbor has signaled not to advertise the interface in the routing device's outbound link-state PDUs.	detail extensive
IP addresses	IP address of this neighbor.	detail extensive

Table 333: show isis adjacency Output Fields (*continued*)

Field Name	Field Description	Level of Output
Transition log	<p>List of recent transitions, including:</p> <ul style="list-style-type: none"> • When—Time at which an IS-IS adjacency transition occurred. • State—Current state of the IS-IS adjacency (up, down, or rejected). <ul style="list-style-type: none"> • Up—Adjacency is up and operational. • Down—Adjacency is down and not available. • Rejected—Adjacency has been rejected. • Event—Type of transition that occurred. <ul style="list-style-type: none"> • Seenself—Possible routing loop has been detected. • Interface down—IS-IS interface has gone down and is no longer available. • Error—Adjacency error. • Down reason—Reason that an IS-IS adjacency is down: <ul style="list-style-type: none"> • 3-Way Handshake Failed—Connection establishment failed. • Address Mismatch—Address mismatch caused link failure. • Aged Out—Link expired. • ISO Area Mismatch—IS-IS area mismatch caused link failure. • Bad Hello—Unacceptable hello message caused link failure. • BFD Session Down—Bidirectional failure detection caused link failure. • Interface Disabled—IS-IS interface is disabled. • Interface Down—IS-IS interface is unavailable. • Interface Level Disabled—IS-IS level is disabled. • Level Changed—IS-IS level has changed on the adjacency. • Level Mismatch—Levels on adjacency are not compatible. • MPLS LSP Down—Label-switched path (LSP) is unavailable. • MT Topology Changed—IS-IS topology has changed. • MT Topology Mismatch—IS-IS topology is mismatched. • Remote System ID Changed—Adjacency peer system ID changed. • Protocol Shutdown—IS-IS protocol is disabled. • CLI Command—Adjacency brought down by user. • Unknown—Unknown. 	extensive

Sample Output

show isis adjacency

```

user@host> show isis adjacency
Interface          System      L State      Hold (secs) SNPA
at-2/3/0.0         ranier      3 Up          23

```

show isis adjacency brief

The output for the **show isis adjacency brief** command is identical to that for the **show isis adjacency** command. For sample output, see [show isis adjacency on page 4340](#).

show isis adjacency detail

```
user@host> show isis adjacency detail
ranier
Interface: at-2/3/0.0, Level: 3, State: Up, Expires in 21 secs
Priority: 0, Up/Down transitions: 1, Last transition: 00:01:09 ago
Circuit type: 3, Speaks: IP, IPv6
Topologies: Unicast
Restart capable: Yes
IP addresses: 11.1.1.2
```

show isis adjacency extensive

```
user@host> show isis adjacency extensive
ranier
Interface: at-2/3/0.0, Level: 3, State: Up, Expires in 22 secs
Priority: 0, Up/Down transitions: 1, Last transition: 00:01:16 ago
Circuit type: 3, Speaks: IP, IPv6
Topologies: Unicast
Restart capable: Yes
IP addresses: 11.1.1.2
Transition log:
When           State      Event      Down reason
Wed Nov  8 21:24:25  Up        Seenself
```

show isis authentication

List of Syntax	Syntax on page 4342 Syntax (EX Series Switches and QFX Series) on page 4342
Syntax	<pre>show isis authentication <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switches and QFX Series)	<pre>show isis authentication <instance <i>instance-name</i>></pre>
Release Information	<p>Command introduced in Junos OS Release 7.5.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for hitless authentication key rollover introduced in Junos OS Release 11.2.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Display information about IS-IS authentication.
Options	<p>none—Display information about IS-IS authentication.</p> <p>instance <i>instance-name</i>—(Optional) Display IS-IS authentication for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> show security keychain on page 5124
List of Sample Output	show isis authentication on page 4343 show isis authentication (With Hitless Authentication Key Rollover Configured) on page 4343
Output Fields	<p>Table 334 describes the output fields for the show isis authentication command. Output fields are listed in the approximate order in which they appear.</p>

Table 334: show isis authentication Output Fields

Field Name	Field Description
Interface	Interface name.
Level	IS-IS level.

Table 334: show isis authentication Output Fields (*continued*)

Field Name	Field Description
IIH Auth	IS-IS Hello (IIH) packet authentication type. Displays the name of the active keychain if hitless authentication key rollover is configured.
CSN Auth	Complete sequence number authentication type.
PSN Auth	Partial sequence number authentication type.
L1 LSP Authentication	Layer 1 link-state PDU authentication type.
L2 LSP Authentication	Layer 2 link-state PDU authentication type.

Sample Output

show isis authentication

```

user@host> show isis authentication
Interface          Level IIH Auth  CSN Auth  PSN Auth
at-2/3/0.0         1    Simple      Simple    Simple
                   2    MD5         MD5       MD5

L1 LSP Authentication: Simple
L2 LSP Authentication: MD5

```

show isis authentication (With Hitless Authentication Key Rollover Configured)

```

user@host> show isis authentication
Interface          Level IIH Auth  CSN Auth  PSN Auth
so-0/1/3.0         2    hakrhello MD5       MD5

L2 LSP Authentication: MD5

```

show isis backup coverage

Syntax	<pre>show isis backup coverage <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switches and QFX Series)	<pre>show isis backup coverage <instance <i>instance-name</i>></pre>
Release Information	<p>Command introduced in Junos OS Release 9.5.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Display information about the level of backup coverage available.
Options	<p>none—Display information about the level of backup coverage available for all the nodes and prefixes in the network.</p> <p>instance <i>instance-name</i>—(Optional) Display information about the level of backup coverage for a specific IS-IS routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Understanding Loop-Free Alternate Routes for IS-IS on page 4165 • Example: Configuring Node-Link Protection for IS-IS Routes in a Layer 3 VPN on page 4221 • show isis backup label-switched-path on page 4346
List of Sample Output	show isis backup coverage on page 4345
Output Fields	Table 335 lists the output fields for the show isis backup coverage command. Output fields are listed in the approximate order in which they appear.

Table 335: show isis backup coverage Output Fields

Field Name	Field Description
Topology	Type of topology or address family: IPv4 Unicast or IPv6 Unicast .
Level	IS-IS level: <ul style="list-style-type: none"> • 1—Level 1 • 2—Level 2
Node	By topology, the percentage of all routes configured on the node that are protected through backup coverage.

Table 335: show isis backup coverage Output Fields (*continued*)

Field Name	Field Description
IPv4	Percentage of IPv4 unicast routes that are protected through backup coverage.
IPv6	Percentage of IPv6 unicast routes that are protected through backup coverage.
CLNS	Percentage of Connectionless Network Service (CLNS) routes that are protected through backup coverage.

Sample Output

show isis backup coverage

```

user@host> show isis backup coverage
Backup Coverage:
  Topology    Level  Node   IPv4   IPv6   CLNS
  IPV4 Unicast    2  28.57%  22.22%  0.00%  0.00%
  IPV6 Unicast    2   0.00%  0.00%  0.00%  0.00%

```

show isis backup label-switched-path

Syntax	show isis backup label-switched-path <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches and QFX Series)	show isis backup label-switched-path
Release Information	Command introduced in Junos OS Release 9.5. Command introduced in Junos OS Release 9.5 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Display information about MPLS label-switched-paths (LSPs) designated as backup routes for IS-IS routes.
Options	none —Display information about MPLS LSPs designated as backup routes for IS-IS routes. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Understanding Loop-Free Alternate Routes for IS-IS on page 4165 • Example: Configuring Node-Link Protection for IS-IS Routes in a Layer 3 VPN on page 4221 • show isis backup coverage on page 4344
List of Sample Output	show isis backup label-switched-path on page 4347
Output Fields	Table 336 lists the output fields for the show isis backup label-switched-path command. Output fields are listed in the approximate order in which they appear.

Table 336: show isis backup label-switched-path Output Fields

Field Name	Field Description
Backup MPLS LSPs	List of MPLS LSPs designated as backup paths for IS-IS routes.
Egress	IP address of the egress routing device for the LSP.

Table 336: show isis backup label-switched-path Output Fields (*continued*)

Field Name	Field Description
Status	<p>State of the LSP:</p> <ul style="list-style-type: none"> • Up—The routing device can detect RSVP hello messages from the neighbor. • Down—The routing device has received one of the following indications: <ul style="list-style-type: none"> • Communication failure from the neighbor. • Communication from IGP that the neighbor is unavailable. • Change in the sequence numbers in the RSVP hello messages sent by the neighbor. • Deleted—LSP is no longer available as a backup path.
Last change	Time elapsed since the neighbor state changed either from up to down or from down to up. The format is <i>hh:mm:ss</i> .
TE-metric	Configured traffic engineering metric.
Metric	Configured metric.

Sample Output

show isis backup label-switched-path

```

user@host> show isis backup label-switched-path
Backup MPLS LSPs:
f-to-g, Egress: 192.168.1.4, Status: up, Last change: 06:12:03
TE-metric: 9, Metric: 0

```

show isis backup spf results

Syntax	<pre>show isis backup spf results <instance <i>instance-name</i>> <level (1 2)> <logical-system (all <i>logical-system-name</i>)> <no-coverage> <topology (ipv4-unicast ipv6-multicast ipv6-unicast unicast)></pre>
Syntax (EX Series Switches)	<pre>show isis backup spf results <instance <i>instance-name</i>> <level (1 2)> <no-coverage> <topology (ipv4-unicast unicast)></pre>
Release Information	Command introduced in Junos OS Release 9.5.
Description	Display information about IS-IS shortest-path-first (SPF) calculations for backup paths.
Options	<p>none—Display information about IS-IS SPF calculations for all backup paths for all destination nodes.</p> <p>instance <i>instance-name</i>—(Optional) Display SPF calculations for backup paths for the specified routing instance.</p> <p>level (1 2)—(Optional) Display SPF calculations for the backup paths for the specified IS-IS level.</p> <p>logical-system <i>logical-system-name</i>—(Optional) Display SPF calculations for the backup paths for all logical systems or on a particular logical system.</p> <p>no-coverage—(Optional) Display SPF calculations only for destinations that do not have backup coverage.</p> <p>topology (ipv4-multicast ipv6-multicast ipv6-unicast unicast)—(Optional) Display SPF calculations for backup paths for the specified topology only.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Link and Node Protection for IS-IS Routes on page 4217• show isis backup coverage on page 4344• Understanding Loop-Free Alternate Routes for IS-IS on page 4165• Example: Configuring Node-Link Protection for IS-IS Routes in a Layer 3 VPN on page 4221
List of Sample Output	<p>show isis backup spf results on page 4349</p> <p>show isis backup spf results no-coverage on page 4350</p>
Output Fields	<p>Table 337 lists the output fields for the show isis backup spf results command. Output fields are listed in the approximate order in which they appear.</p>

Table 337: show isis backup spf results Output Fields

Field Name	Field Description
<i>node-name</i>	Name of the destination node.
Address	Address of the destination node.
Primary next-hop	Interface and name of the node of the primary next hop to reach the destination.
Root	Name of the next-hop neighbor.
Metric	Metric to the node.
Eligible	Indicates that the next-hop neighbor has been designated as a backup path to the destination node.
Backup next-hop	Name of the interface of the backup next hop.
SNPA	Subnetwork point of attachment (MAC address of the next hop).
LSP	Name of the MPLS label-switched path (LSP) designated as a backup path.
Not eligible	Indicates that the next-hop neighbor cannot function as a backup path to the destination.
Reason	Describes why the next-hop neighbor is designated as Not eligible as a backup path.

Sample Output

show isis backup spf results

```

user@host> show isis backup spf results

IS-IS level 1 SPF results:
  0 nodes

IS-IS level 2 SPF results:
banff.00
  Primary next-hop: so-6/0/0.0, IPV4, olympic
  Primary next-hop: ae0.0, IPV4, camaro, SNPA: 0:90:69:f:67:f0
  Primary next-hop: so-6/0/0.0, IPV6, olympic
  Primary next-hop: ae0.0, IPV6, camaro, SNPA: 0:90:69:f:67:f0
  Root: camaro, Root Metric: 10, Metric: 10
  Not eligible, Reason: Primary next-hop multipath
  Root: olympic, Root Metric: 10, Metric: 10
  Not eligible, Reason: Primary next-hop multipath
  Root: glacier, Root Metric: 10, Metric: 25
  Not eligible, Reason: Primary next-hop multipath
crater.00
  Primary next-hop: so-6/0/0.0, IPV4, olympic
  Primary next-hop: so-6/0/0.0, IPV6, olympic

```

```

Root: olympic, Root Metric: 10, Metric: 10
  Not eligible, Reason: Primary next-hop link fate sharing
Root: glacier, Root Metric: 10, Metric: 15
  Eligible, Backup next-hop: as0.0, IPV4, glacier
  Eligible, Backup next-hop: as0.0, IPV6, glacier
Root: camaro, Root Metric: 10, Metric: 20
  Not eligible, Reason: Interface is already covered
olympic.00
Primary next-hop: so-6/0/0.0, IPV4, olympic
Primary next-hop: so-6/0/0.0, IPV6, olympic
Root: olympic, Root Metric: 10, Metric: 0
  Not eligible, Reason: Primary next-hop link fate sharing
Root: camaro, Root Metric: 10, Metric: 20
  track-item: olympic.00-00
  track-item: kobuk.00-00
  Not eligible, Reason: Path loops
Root: glacier, Root Metric: 10, Metric: 20
  track-item: olympic.00-00
  track-item: kobuk.00-00
  Not eligible, Reason: Path loops
camaro.00
Primary next-hop: ae0.0, IPV4, camaro, SNPA: 0:90:69:f:67:f0
Primary next-hop: ae0.0, IPV6, camaro, SNPA: 0:90:69:f:67:f0
Root: camaro, Root Metric: 10, Metric: 0
  Not eligible, Reason: Primary next-hop link fate sharing
Root: glacier, Root Metric: 10, Metric: 20
  track-item: camaro.00-00
  track-item: kobuk.00-00
  Not eligible, Reason: Path loops
Root: olympic, Root Metric: 10, Metric: 20
  track-item: camaro.00-00
  track-item: kobuk.00-00
  Not eligible, Reason: Path loops
glacier.00
Primary next-hop: as0.0, IPV4, glacier
Primary next-hop: as0.0, IPV6, glacier
Root: glacier, Root Metric: 10, Metric: 0
  Not eligible, Reason: Primary next-hop link fate sharing
Root: camaro, Root Metric: 10, Metric: 20
  track-item: glacier.00-00
  track-item: kobuk.00-00
  Not eligible, Reason: Path loops
Root: olympic, Root Metric: 10, Metric: 20
  track-item: glacier.00-00
  track-item: kobuk.00-00
  Not eligible, Reason: Path loops
5 nodes

```

show isis backup spf results no-coverage

```

user@host> show isis backup spf results no-coverage
IS-IS level 1 SPF results:
pro-bng3-k.00
Primary next-hop: fe-1/3/3.0, IPV4, pro-bng3-k, SNPA: b0:c6:9a:2c:f0:de
Primary next-hop: fe-1/3/3.0, IPV6, pro-bng3-k, SNPA: b0:c6:9a:2c:f0:de
Root: pro-bng3-k, Root Metric: 10, Metric: 0, Root Preference: 0x0
Root: pro-bng3-i, Root Metric: 10, Metric: 20, Root Preference: 0x0
  track-item: pro-bng3-k.00-00
  track-item: pro-bng3-j.00-00
pro-bng3-i.00
Primary next-hop: fe-0/1/2.0, IPV4, pro-bng3-i, SNPA: b0:c6:9a:2a:f4:21

```

```

Primary next-hop: fe-0/1/2.0, IPV6, pro-bng3-i, SNPA: b0:c6:9a:2a:f4:21
Root: pro-bng3-i, Root Metric: 10, Metric: 0, Root Preference: 0x0
Root: pro-bng3-k, Root Metric: 10, Metric: 20, Root Preference: 0x0
  track-item: pro-bng3-j.00-00
  track-item: pro-bng3-i.00-00
2 nodes

IS-IS level 2 SPF results:
olympic.00
Primary next-hop: so-6/0/0.0, IPV4, olympic
Primary next-hop: so-6/0/0.0, IPV6, olympic
Root: olympic, Root Metric: 10, Metric: 0
  Not eligible, Reason: Primary next-hop link fate sharing
Root: camaro, Root Metric: 10, Metric: 20
  track-item: olympic.00-00
  track-item: kobuk.00-00
  Not eligible, Reason: Path loops
Root: glacier, Root Metric: 10, Metric: 20
  track-item: olympic.00-00
  track-item: kobuk.00-00
  Not eligible, Reason: Path loops
camaro.00
Primary next-hop: ae0.0, IPV4, camaro, SNPA: 0:90:69:f:67:f0
Primary next-hop: ae0.0, IPV6, camaro, SNPA: 0:90:69:f:67:f0
Root: camaro, Root Metric: 10, Metric: 0
  Not eligible, Reason: Primary next-hop link fate sharing
Root: glacier, Root Metric: 10, Metric: 20
  track-item: camaro.00-00
  track-item: kobuk.00-00
  Not eligible, Reason: Path loops
Root: olympic, Root Metric: 10, Metric: 20
  track-item: camaro.00-00
  track-item: kobuk.00-00
  Not eligible, Reason: Path loops
glacier.00
Primary next-hop: as0.0, IPV4, glacier
Primary next-hop: as0.0, IPV6, glacier
Root: glacier, Root Metric: 10, Metric: 0
  Not eligible, Reason: Primary next-hop link fate sharing
Root: camaro, Root Metric: 10, Metric: 20
  track-item: glacier.00-00
  track-item: kobuk.00-00
  Not eligible, Reason: Path loops
Root: olympic, Root Metric: 10, Metric: 20
  track-item: glacier.00-00
  track-item: kobuk.00-00
  Not eligible, Reason: Path loops
3 nodes

```

show isis database

List of Syntax	Syntax on page 4352 Syntax (EX Series Switch and QFX Series) on page 4352
Syntax	<pre>show isis database <system-id> <brief detail extensive> <instance <i>instance-name</i>> <level (1 2)> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switch and QFX Series)	<pre>show isis database <system-id> <brief detail extensive> <level (1 2)> <instance <i>instance-name</i>></pre>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Display the entries in the Intermediate System-to-Intermediate System (IS-IS) link-state database, which contains data about PDU packets.
Options	<p>none—Display standard information about IS-IS link-state database entries for all routing instances.</p> <p><i>system id</i>—(Optional) Display IS-IS link-state database entries for the specified intermediate system.</p> <p><i>brief detail extensive</i>—(Optional) Display the specified level of output.</p> <p><i>instance instance-name</i>—(Optional) Display IS-IS link-state database entries for the specified routing instance.</p> <p><i>level (1 2)</i>—(Optional) Display IS-IS link-state database entries for the specified IS-IS level.</p> <p><i>logical-system (all logical-system-name)</i>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• clear isis database on page 4332
List of Sample Output	show isis database on page 4354 show isis database brief on page 4355 show isis database detail on page 4355

[show isis database extensive on page 4355](#)

Output Fields Table 338 describes the output fields for the **show isis database** command. Output fields are listed in the approximate order in which they appear. Fields that contain internal IS-IS information useful only in troubleshooting obscure problems are not described in the table. For more details about these fields, contact your customer support representative.

Table 338: show isis database Output Fields

Field Name	Field Description	Level of Output
Interface name	Name of the interface on which the link-state PDU has been received; always IS-IS for this command.	All levels
level	Level of intermediate system: <ul style="list-style-type: none"> • 1—Intermediate system routes within an area; when the destination is outside an area, it routes toward a Level 2 system. • 2—Intermediate system routes between areas and toward other ASs. 	All levels
LSP ID	Link-state PDU identifier.	All levels
Sequence	Sequence number of the link-state PDU.	All levels
Checksum	Checksum value of the link-state PDU.	All levels
Lifetime (secs)	Remaining lifetime of the link-state PDU, in seconds.	All levels
Attributes	Attributes of the specified database: L1 , L2 , Overload , or Attached (L1 only).	none brief
# LSPs	Total number of link-state PDUs in the specified link-state database.	none brief
IP prefix	Prefix advertised by this link-state PDU.	detail extensive
IS neighbor	IS-IS neighbor of the advertising system.	detail extensive
IP prefix	IPv4 prefix advertised by this link-state PDU.	detail extensive
V6 prefix	IPv6 prefix advertised by this link-state PDU.	detail extensive
Metric	Metric of the prefix or neighbor.	detail extensive
Header	<ul style="list-style-type: none"> • LSP ID—Link state PDU identifier of the header. • Length—Header length. • Allocated Length—Amount of length available for the header. • Router ID—Address of the local routing device. • Remaining Lifetime—Remaining lifetime of the link-state PDU, in seconds. 	extensive

Table 338: show isis database Output Fields (*continued*)

Field Name	Field Description	Level of Output
Packet	<ul style="list-style-type: none"> • LSP ID—The identifier for the link-state PDU. • Length—Packet length. • Lifetime—Remaining lifetime, in seconds. • Checksum—The checksum of the link-state PDU. • Sequence—The sequence number of the link-state PDU. Every time the link-state PDU is updated, this number increments. • Attributes—Packet attributes. • NLPID—Network layer protocol identifier. • Fixed length—Specifies the set length for the packet. 	extensive
TLVs	<ul style="list-style-type: none"> • Area Address—Area addresses that the routing device can reach. • Speaks—Supported routing protocols. • IP router id—ID of the routing device (usually the IP address). • IP address—IPv4 address. • Hostname—Assigned name of the routing device. • IP prefix—IP prefix of the routing device. • Metric—IS-IS metric that measures the cost of the adjacency between the originating routing device and the advertised routing device. • IP extended prefix—Extended IP prefix of the routing device. • IS neighbor—Directly attached neighbor's name and metric. • IS extended neighbor—Directly attached neighbor's name, metric, IP address, local interface index, and remote interface index. <p>The interface indexes enable Junos OS to support unnumbered extensions for IS-IS, as described in RFC 4205.</p>	extensive

Sample Output

show isis database

```

user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                Sequence Checksum Lifetime Attributes
kobuk.00-00           0x3     0x3167    1057 L1 L2
camaro.00-00          0x5     0x770e    1091 L1 L2
ranier.00-00          0x4     0xaa95    1091 L1 L2
glacier.00-00         0x4     0x206f    1089 L1 L2
glacier.02-00         0x1     0xd141    1089 L1 L2
badlands.00-00        0x3     0x87a2    1093 L1 L2
  6 LSPs

IS-IS level 2 link-state database:
LSP ID                Sequence Checksum Lifetime Attributes
kobuk.00-00           0x6     0x8d6b    1096 L1 L2
camaro.00-00          0x9     0x877b    1101 L1 L2
ranier.00-00          0x8     0x855d    1103 L1 L2
glacier.00-00         0x7     0xf892    1098 L1 L2
glacier.02-00         0x1     0xd141    1089 L1 L2
badlands.00-00        0x6     0x562     1105 L1 L2
  6 LSPs

```

show isis database brief

The output for the **show isis database brief** command is identical to that for the **show isis database** command. For sample output, see [show isis database on page 4354](#).

show isis database detail

```
user@host> show isis database logical-system CE3 sisira.00-00 detail
```

IS-IS level 1 link-state database:

```
sisira.00-00 Sequence: 0x11, Checksum: 0x10fc, Lifetime: 975 secs
  IS neighbor: hemantha-CE3.02           Metric:      10
  ES neighbor: 0015.0015.0015           Metric:      10 Down
  ES neighbor: 0025.0025.0025           Metric:      10 Down
  ES neighbor: 0030.0030.0030           Metric:      10 Down
  ES neighbor: 0040.0040.0040           Metric:      10 Down
  ES neighbor: sisira                     Metric:       0
  IP prefix: 1.0.0.0/24                  Metric:      10 External Down
  IP prefix: 3.0.0.0/24                  Metric:      10 External Down
  IP prefix: 4.0.0.0/24                  Metric:      10 External Down
  IP prefix: 5.0.0.0/24                  Metric:      10 Internal Up
  IP prefix: 15.15.15.15/32              Metric:      10 External Down
  IP prefix: 25.25.25.25/32              Metric:      10 External Down
  IP prefix: 30.30.30.30/32              Metric:      10 External Down
  IP prefix: 40.40.40.40/32              Metric:      10 External Down
  IP prefix: 60.60.60.60/32              Metric:       0 Internal Up
```

IS-IS level 2 link-state database:

```
sisira.00-00 Sequence: 0x13, Checksum: 0x69ac, Lifetime: 993 secs
  IS neighbor: hemantha-CE3.02           Metric:      10
  IP prefix: 1.0.0.0/24                  Metric:      10 External Down
  IP prefix: 3.0.0.0/24                  Metric:      10 External Down
  IP prefix: 4.0.0.0/24                  Metric:      10 External Down
  IP prefix: 5.0.0.0/24                  Metric:      10 Internal Up
  IP prefix: 15.15.15.15/32              Metric:      10 External Down
  IP prefix: 25.25.25.25/32              Metric:      10 External Down
  IP prefix: 30.30.30.30/32              Metric:      10 External Down
  IP prefix: 40.40.40.40/32              Metric:      10 External Down
  IP prefix: 50.50.50.50/32              Metric:      10 Internal Up
  IP prefix: 60.60.60.60/32              Metric:       0 Internal Up
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0015.0015.0015/152
                                          Metric:      10 External Down
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0025.0025.0025/152
                                          Metric:      10 External Down
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0030.0030.0030/152
                                          Metric:      10 External Down
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0040.0040.0040/152
                                          Metric:      10 External Down
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0060.0060.0060/152
                                          Metric:       0 Internal Up
```

show isis database extensive

```
user@host> show isis database extensive
```

IS-IS level 1 link-state database:

```
sisira.00-00 Sequence: 0x11, Checksum: 0x10fc, Lifetime: 970 secs
```

```

IS neighbor: hemantha-CE3.02                Metric:      10
Two-way fragment: hemantha-CE3.02-00, Two-way first fragment:
hemantha-CE3.02-00
ES neighbor: 0015.0015.0015                Metric:      10 Down
ES neighbor: 0025.0025.0025                Metric:      10 Down
ES neighbor: 0030.0030.0030                Metric:      10 Down
ES neighbor: 0040.0040.0040                Metric:      10 Down
ES neighbor: sisira                        Metric:       0
IP prefix: 1.0.0.0/24                      Metric:      10 External Down
IP prefix: 3.0.0.0/24                      Metric:      10 External Down
IP prefix: 4.0.0.0/24                      Metric:      10 External Down
IP prefix: 5.0.0.0/24                      Metric:      10 Internal Up
IP prefix: 15.15.15.15/32                 Metric:      10 External Down
IP prefix: 25.25.25.25/32                 Metric:      10 External Down
IP prefix: 30.30.30.30/32                 Metric:      10 External Down
IP prefix: 40.40.40.40/32                 Metric:      10 External Down
IP prefix: 60.60.60.60/32                 Metric:       0 Internal Up

```

```

Header: LSP ID: sisira.00-00, Length: 336 bytes
Allocated length: 336 bytes, Router ID: 0.0.0.0
Remaining lifetime: 970 secs, Level: 1, Interface: 333
Estimated free bytes: 144, Actual free bytes: 0
Aging timer expires in: 970 secs
Protocols: IP, IPv6, CLNS

```

```

Packet: LSP ID: sisira.00-00, Length: 336 bytes, Lifetime : 1198 secs
Checksum: 0x10fc, Sequence: 0x11, Attributes: 0xb L1 L2 Attached
NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
Packet type: 18, Packet version: 1, Max area: 0

```

TLVs:

```

Area address: 60.0006.80ff.f800.0000.0108.0001 (13)
Speaks: IP
Speaks: IPV6
Speaks: CLNP
Hostname: sisira
ES neighbor TLV: Internal, Metric: default 0, Up
  ES: sisira
IS neighbor: hemantha-CE3.02, Internal, Metric: default 10
IS extended neighbor: hemantha-CE3.02, Metric: default 10
ES neighbor TLV: External, Metric: default 10, Down
  ES: 0040.0040.0040
ES neighbor TLV: External, Metric: default 10, Down
  ES: 0025.0025.0025
ES neighbor TLV: External, Metric: default 10, Down
  ES: 0015.0015.0015
ES neighbor TLV: External, Metric: default 10, Down
  ES: 0030.0030.0030
IP external prefix: 3.0.0.0/24, Internal, Metric: default 10, Down
IP external prefix: 40.40.40.40/32, Internal, Metric: default 10, Down
IP external prefix: 4.0.0.0/24, Internal, Metric: default 10, Down
IP external prefix: 25.25.25.25/32, Internal, Metric: default 10, Down
IP external prefix: 15.15.15.15/32, Internal, Metric: default 10, Down
IP external prefix: 30.30.30.30/32, Internal, Metric: default 10, Down
IP extended prefix: 3.0.0.0/24 metric 10 down
IP extended prefix: 40.40.40.40/32 metric 10 down
IP extended prefix: 4.0.0.0/24 metric 10 down
IP extended prefix: 25.25.25.25/32 metric 10 down
IP extended prefix: 15.15.15.15/32 metric 10 down
IP extended prefix: 1.0.0.0/24 metric 10 down

```

```

IP extended prefix: 30.30.30.30/32 metric 10 down
IP prefix: 60.60.60.60/32, Internal, Metric: default 0, Up
IP prefix: 5.0.0.0/24, Internal, Metric: default 10, Up
IP extended prefix: 60.60.60.60/32 metric 0 up
IP extended prefix: 5.0.0.0/24 metric 10 up
No queued transmissions

```

IS-IS level 2 link-state database:

```

sisira.00-00 Sequence: 0x13, Checksum: 0x69ac, Lifetime: 988 secs
IS neighbor: hemantha-CE3.02 Metric: 10
Two-way fragment: hemantha-CE3.02-00, Two-way first fragment:
hemantha-CE3.02-00
IP prefix: 1.0.0.0/24 Metric: 10 External Down
IP prefix: 3.0.0.0/24 Metric: 10 External Down
IP prefix: 4.0.0.0/24 Metric: 10 External Down
IP prefix: 5.0.0.0/24 Metric: 10 Internal Up
IP prefix: 15.15.15.15/32 Metric: 10 External Down
IP prefix: 25.25.25.25/32 Metric: 10 External Down
IP prefix: 30.30.30.30/32 Metric: 10 External Down
IP prefix: 40.40.40.40/32 Metric: 10 External Down
IP prefix: 50.50.50.50/32 Metric: 10 Internal Up
IP prefix: 60.60.60.60/32 Metric: 0 Internal Up
ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0015.0015.0015/152
Metric: 10 External Down
ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0025.0025.0025/152
Metric: 10 External Down
ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0030.0030.0030/152
Metric: 10 External Down
ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0040.0040.0040/152
Metric: 10 External Down
ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0060.0060.0060/152
Metric: 0 Internal Up

```

```

Header: LSP ID: sisira.00-00, Length: 427 bytes
Allocated length: 427 bytes, Router ID: 0.0.0.0
Remaining lifetime: 988 secs, Level: 2, Interface: 333
Estimated free bytes: 130, Actual free bytes: 0
Aging timer expires in: 988 secs
Protocols: IP, IPv6, CLNS

```

```

Packet: LSP ID: sisira.00-00, Length: 427 bytes, Lifetime : 1198 secs
Checksum: 0x69ac, Sequence: 0x13, Attributes: 0x3 L1 L2
NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
Packet type: 20, Packet version: 1, Max area: 0

```

TLVs:

```

Area address: 60.0006.80ff.f800.0000.0108.0001 (13)
Speaks: IP
Speaks: IPV6
Speaks: CLNP
Hostname: sisira
IS neighbor: hemantha-CE3.02, Internal, Metric: default 10
IS extended neighbor: hemantha-CE3.02, Metric: default 10
IP external prefix: 3.0.0.0/24, Internal, Metric: default 10, Down
IP external prefix: 40.40.40.40/32, Internal, Metric: default 10, Down
IP external prefix: 4.0.0.0/24, Internal, Metric: default 10, Down
IP external prefix: 25.25.25.25/32, Internal, Metric: default 10, Down
IP external prefix: 15.15.15.15/32, Internal, Metric: default 10, Down
IP external prefix: 1.0.0.0/24, Internal, Metric: default 10, Down
IP external prefix: 30.30.30.30/32, Internal, Metric: default 10, Down

```

```

IP extended prefix: 3.0.0.0/24 metric 10 down
IP extended prefix: 40.40.40.40/32 metric 10 down
IP extended prefix: 4.0.0.0/24 metric 10 down
IP extended prefix: 25.25.25.25/32 metric 10 down
IP extended prefix: 15.15.15.15/32 metric 10 down
IP extended prefix: 1.0.0.0/24 metric 10 down
IP extended prefix: 30.30.30.30/32 metric 10 down
ISO prefix-neighbor TLV: Internal, Metric: default 0, Up
  Prefix : 60.0006.80ff.f800.0000.0108.0001.0060.0060/152
ISO prefix-neighbor TLV: External, Metric: default 10, Down
  Prefix : 60.0006.80ff.f800.0000.0108.0001.0040.0040/152
ISO prefix-neighbor TLV: External, Metric: default 10, Down
  Prefix : 60.0006.80ff.f800.0000.0108.0001.0025.0025/152
ISO prefix-neighbor TLV: External, Metric: default 10, Down
  Prefix : 60.0006.80ff.f800.0000.0108.0001.0015.0015/152
ISO prefix-neighbor TLV: External, Metric: default 10, Down
  Prefix : 60.0006.80ff.f800.0000.0108.0001.0030.0030/152
IP prefix: 60.60.60.60/32, Internal, Metric: default 0, Up
IP prefix: 5.0.0.0/24, Internal, Metric: default 10, Up
IP prefix: 50.50.50.50/32, Internal, Metric: default 10, Up
IP extended prefix: 60.60.60.60/32 metric 0 up
IP extended prefix: 5.0.0.0/24 metric 10 up
IP extended prefix: 50.50.50.50/32 metric 10 up
No queued transmissions

```

```

Router-F.02-00 Sequence: 0x1, Checksum: 0xf5ae, Lifetime: 1153 secs
IS neighbor: Router-E.00 Metric: 0
  Two-way fragment: Router-E.00-00, Two-way first fragment: Router-E.00-00
IS neighbor: Router-F.00 Metric: 0
  Two-way fragment: Router-F.00-00, Two-way first fragment: Router-F.00-00

```

```

Header: LSP ID: Router-F.02-00, Length: 76 bytes
  Allocated length: 284 bytes, Router ID: 0.0.0.0
  Remaining lifetime: 1153 secs, Level: 2, Interface: 101
  Estimated free bytes: 208, Actual free bytes: 208
  Aging timer expires in: 1153 secs

```

```

Packet: LSP ID: Router-F.02-00, Length: 76 bytes, Lifetime : 1183 secs
  Checksum: 0xf5ae, Sequence: 0x1, Attributes: 0x3 <L1 L2>
  NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
  Packet type: 20, Packet version: 1, Max area: 0

```

```

TLVs:
  IS neighbor: Router-F.00, Internal, Metric: default 0
  IS neighbor: Router-E.00, Internal, Metric: default 0
  IS extended neighbor: Router-F.00, Metric: default 0
  IS extended neighbor: Router-E.00, Metric: default 0
No queued transmissions

```

show isis hostname

List of Syntax [Syntax on page 4359](#)
[Syntax \(EX Series Switches and QFX Series\) on page 4359](#)

Syntax show isis hostname
 <logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches and QFX Series) show isis hostname

Release Information Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.0 for EX Series switches.
 Command introduced in Junos OS Release 12.1 for the QFX Series.
 Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Display IS-IS hostname database information.

This command displays the system ID-to-name cache. The output shows if the mapping has been learned by receipt of a Hostname TLV #137 (type dynamic) configured in Junos OS with the **set system host-name** command, or a static mapping defined in Junos OS with the **set system static-host-mapping hostname sysid** command (type static). The local router always has its type set to static even if **static-host-mapping** is not configured.

Options **none**—Display IS-IS hostname database information.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level view

List of Sample Output [show isis hostname on page 4360](#)

Output Fields [Table 339](#) describes the output fields for the **show isis hostname** command. Output fields are listed in the approximate order in which they appear.

Table 339: show isis hostname Output Fields

Field Name	Field Description
System Id	System identifier mapped to the hostname.
Hostname	Hostname mapped to the system identifier.
Type	Type of mapping between system identifier and hostname. <ul style="list-style-type: none"> • Dynamic—Hostname mapping determined as described in RFC 2763, <i>Dynamic Hostname Exchange Mechanism for IS-IS</i>. • Static—Hostname mapping configured by user.

Sample Output

show isis hostname

```
user@host> show isis hostname
IS-IS hostname database:
System Id      Hostname
1921.6800.4201 isis1
1921.6800.4202 isis2
1921.6800.4203 isis3
```

Type
Dynamic
Static
Dynamic

show isis interface

List of Syntax [Syntax on page 4361](#)
[Syntax \(EX Series Switches and QFX Series\) on page 4361](#)

Syntax show isis interface
 <brief | detail | extensive>
 <interface-name>
 <logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches and QFX Series) show isis interface
 <brief | detail | extensive>
 <interface-name>

Release Information Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.0 for EX Series switches.
 Command introduced in Junos OS Release 12.1 for the QFX Series.

Description Display status information about IS-IS-enabled interfaces.



NOTE: If the configured metric for an IS-IS level is above 63, and the wide-metrics-only statement is not configured, the show isis interface detail command and the show isis interface extensive command display 63 as the metric value for that level. Configure the wide-metrics-only statement to generate metric values greater than 63 on a per IS-IS level basis.

The show isis interface command displays the configured metric value for an IS-IS level irrespective of whether is configured or not.

Options none—Display standard information about all IS-IS-enabled interfaces.
 brief | detail | extensive—(Optional) Display the specified level of output.
 interface-name—(Optional) Display information about the specified interface only.
 logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level view

Related Documentation

- [Understanding Wide IS-IS Metrics for Traffic Engineering](#)
- [Example: Enabling Wide IS-IS Metrics for Traffic Engineering](#)

List of Sample Output [show isis interface on page 4364](#)
[show isis interface brief on page 4364](#)
[show isis interface detail on page 4364](#)

[show isis interface extensive on page 4364](#)

Output Fields Table 340 describes the output fields for the **show isis interface** command. Output fields are listed in the approximate order in which they appear.

Table 340: show isis interface Output Fields

Field Name	Field Description	Level of Output
<i>interface-name</i>	Name of the interface.	detail
Designated router	Routing device selected by other routers that is responsible for sending link-state advertisements that describe the network. Used only on broadcast networks.	detail
Index	Interface index assigned by the Junos OS kernel.	detail
State	Internal implementation information.	detail
Circuit id	Circuit identifier. NOTE: Each IS-IS interface is assigned a circuit ID value to identify the interface within the linkstate database. All interfaces (loopback, broadcast, and so on) and all point-to-point links share the locally significant value of 0x01, and this value is not incremented.	detail
Circuit type	Circuit type: <ul style="list-style-type: none"> • 1—Level 1 only • 2—Level 2 only • 3—Level 1 and Level 2 	detail
LSP interval	Interval between link-state PDUs sent from the interface.	detail
CSNP interval	Interval between complete sequence number PDUs sent from the interface.	detail extensive
Sysid	System identifier.	detail
Interface	Interface through which the adjacency is made.	none brief
L or Level	Level: <ul style="list-style-type: none"> • 1—Level 1 only • 2—Level 2 only • 3—Level 1 and Level 2 NOTE: The default IS-IS level on loopback interfaces are always same as the IS-IS level configured on other IS-IS interfaces in a router. You can also configure IS-IS level on loopback interfaces per your requirement.	All levels
CirID	Circuit identifier.	none brief
Level 1 DR	Level 1 designated intermediate system.	none brief
Level 2 DR	Level 2 designated intermediate system.	none brief

Table 340: show isis interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
L1/L2 Metric	Interface's metric for Level 1 and Level 2. If there is no information, the metric is 0.	none brief
Adjacency advertisement: Advertise	This routing device has signaled to advertise this interface to its neighbors in their label-switched paths (LSPs).	detail extensive
Adjacency advertisement: Suppress	This neighbor has signaled not to advertise this interface in the routing device's outbound LSPs.	detail extensive
Adjacencies	Number of adjacencies established on this interface.	detail
Priority	Priority value for this interface.	detail
Metric	Metric value for this interface.	detail
Hello(s) / Hello Interval	Interface's hello interval.	detail extensive
Hold(s) / Hold Time	Interface's hold time.	detail extensive
Designated Router	Router responsible for sending network link-state advertisements, which describe all the routing devices attached to the network.	detail
Hello padding	Type of hello padding: <ul style="list-style-type: none"> • Adaptive—On point-to-point connections, the hello packets are padded from the initial detection of a new neighbor until the neighbor verifies the adjacency as Up in the adjacency state TLV. If the neighbor does not support the adjacency state TLV, then padding continues. On LAN connections, padding starts from the initial detection of a new neighbor until there is at least one active adjacency on the interface. • Loose—(Default) The hello packet is padded from the initial detection of a new neighbor until the adjacency transitions to the Up state. • Strict—Padding is performed on all interface types and for all adjacency states, and is continuous. 	extensive
LDP sync state	Current LDP synchronization state: in sync , in holddown , or not supported .	extensive
reason	Reason for being in the LDP sync state.	extensive
config holdtime	Configured value of the hold timer.	extensive
remaining	If the state is not in sync and the hold time is not infinity, then this field displays the remaining hold time in seconds.	extensive

Sample Output

show isis interface

```
user@host> show isis interface
IS-IS interface database:
Interface          L CirID Level 1 DR      Level 2 DR      L1/L2 Metric
at-2/3/0.0         3   0x1 Point to Point    Point to Point    10/10
lo0.0              3   0x1 Passive          Passive           0/0
```

show isis interface brief

The output for the **show isis interface brief** command is identical to that for the **show isis interface** command. For sample output, see [show isis interface on page 4364](#).

show isis interface detail

```
user@host> show isis interface detail
IS-IS interface database:
at-2/3/0.0
  Index: 66, State: 0x6, Circuit id: 0x1, Circuit type: 3
  LSP interval: 100 ms, CSNP interval: 5 s
  Level Adjacencies Priority Metric Hello (s) Hold (s) Designated Router
    1                1      64    10    9.000    27
    2                1      64    10    9.000    27
lo0.0
  Index: 64, State: 0x6, Circuit id: 0x1, Circuit type: 0
  LSP interval: 100 ms, CSNP interval: disabled
  Adjacency advertisement: Advertise
  Protection Type: Node Link, No eligible Backup
  Level Adjacencies Priority Metric Hello (s) Hold (s) Designated Router
    1                0      64     0 Passive
    2                0      64     0 Passive
```

show isis interface extensive

```
user@host> show isis interface extensive
IS-IS interface database:
xe-6/1/0.0
  Index: 75, State: 0x6, Circuit id: 0x1, Circuit type: 2
  LSP interval: 100 ms, CSNP interval: 10 s, Loose Hello padding
  Adjacency advertisement: Advertise
  Level 1
    Adjacencies: 0, Priority: 64, Metric: 10
    Disabled
  Level 2
    Adjacencies: 1, Priority: 64, Metric: 10
    Hello Interval: 20.000 s, Hold Time: 60 s
    Designated Router: nemean.03
```

show isis overview

Syntax	<pre>show isis overview <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switches and QFX Series)	<pre>show isis overview <instance <i>instance-name</i>></pre>
Release Information	<p>Command introduced in Junos OS Release 8.5.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Display IS-IS overview information.
Options	<p>none—Display standard overview information about IS-IS for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Display overview information for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show isis overview on page 4367
Output Fields	Table 341 lists the output fields for the show isis overview command. Output fields are listed in the approximate order in which they appear.

Table 341: show isis overview Output Fields

Field Name	Field Description
Hostname	Name of the router.
Sysid	Part of the ISO address of the routing device.
Areaid	The area number of the routing device.
Instance	IS-IS routing instance.
Router ID	Router ID of the routing device.
Adjacency holddown	Adjacency holddown capability: enabled or disabled .
Maximum Areas	Maximum number of IS-IS areas advertised by the routing device.
LSP life time	Lifetime of the link-state PDU, in seconds.

Table 341: show isis overview Output Fields (*continued*)

Field Name	Field Description
Attached bit evaluation	Attached bit capability: enabled or disabled .
SPF delay	Delay before performing consecutive shortest-path-first (SPF) calculations.
SPF holddown	Delay before performing additional SPF calculations after the maximum number of consecutive SPF calculations is reached.
SPF rapid runs	Maximum number of SPF calculations that can be performed in succession before the holddown timer begins.
Overload bit at startup is set	Overload bit capability is enabled.
Overload high metrics	Overload high metrics capability: enabled or disabled .
Overload timeout	Time period after which overload is reset and the time that remains before the timer is set to expire.
Traffic engineering	Traffic engineering capability: enabled or disabled .
Restart	Graceful restart capability: enabled or disabled .
Restart duration	Time period for complete reacquisition of IS-IS neighbors.
Helper mode	Graceful restart helper capability: enabled or disabled .
Level	IS-IS level: <ul style="list-style-type: none"> • 1—Level 1 information • 2—Level 2 information
IPv4 is enabled	IP Protocol version 4 capability is enabled.
IPv6 is enabled	IP Protocol version 6 capability is enabled.
Internal route preference	Preference value of internal routes.
External route preference	Preference value of external routes.
Prefix export limit	Number of prefixes allowed to be exported, as configured by the prefix-export-limit statement.
Prefix export count	Number of prefixes exported.
Wide area metrics are enabled	Wide area metrics capability is enabled.

Table 341: show isis overview Output Fields (*continued*)

Field Name	Field Description
Narrow metrics are enabled	Narrow metrics capability is enabled.
Adjacency holddown is active	IS-IS adjacencies come up one after another when adjacency holddown is enabled.

Sample Output

show isis overview

```

user@host> show isis overview
Instance: master
Router ID: 10.255.107.183
Hostname: pro-bng3-a
Sysid: 0192.0168.0001
Areaid: 49.0002
Adjacency holddown: enabled
Maximum Areas: 3
LSP life time: 1200
Attached bit evaluation: enabled
SPF delay: 200 msec, SPF holddown: 5000 msec, SPF rapid runs: 3
IPv4 is enabled, IPv6 is enabled
Traffic engineering: enabled
Restart: Disabled
  Helper mode: Enabled
Level 1
  Internal route preference: 15
  External route preference: 160
  Wide metrics are enabled, Narrow metrics are enabled
  Adjacency holddown is active
Level 2
  Internal route preference: 18
  External route preference: 165
  Prefix export limit: 5, Prefix export count: 5
  Wide metrics are enabled
  Adjacency holddown is active

```

```

user@host> show isis overview logical-system R2
Instance: master
Router ID: 192.168.0.2
Hostname: pro-bng3-a-R2
Sysid: 0192.0168.0002
Areaid: 49.0002
Adjacency holddown: enabled
Maximum Areas: 3
LSP life time: 1200
Attached bit evaluation: enabled
SPF delay: 200 msec, SPF holddown: 5000 msec, SPF rapid runs: 3
IPv4 is enabled, IPv6 is enabled
Traffic engineering: enabled
Restart: Disabled
  Helper mode: Enabled
Level 1
  Internal route preference: 15
  External route preference: 160

```

```
    Prefix export count: 0
    Wide metrics are enabled, Narrow metrics are enabled
Level 2
    Internal route preference: 18
    External route preference: 165
    Prefix export count: 0
    Wide metrics are enabled, Narrow metrics are enabled

user@host> show isis overview logical-system R3
Instance: master
  Router ID: 192.168.0.3
  Hostname: pro-bng3-a-R3
  Sysid: 0192.0168.0003
  Areaid: 49.0002
  Adjacency holddown: enabled
  Maximum Areas: 3
  LSP life time: 1200
  Attached bit evaluation: enabled
  SPF delay: 200 msec, SPF holddown: 5000 msec, SPF rapid runs: 3
  IPv4 is enabled, IPv6 is enabled
  Traffic engineering: enabled
  Restart: Disabled
    Helper mode: Enabled
Level 1
  Internal route preference: 15
  External route preference: 160
  Prefix export count: 0
  Wide metrics are enabled, Narrow metrics are enabled
Level 2
  Internal route preference: 18
  External route preference: 165
  Prefix export count: 0
  Wide metrics are enabled, Narrow metrics are enabled
```


show isis route

List of Syntax	Syntax on page 4369 Syntax (EX Series Switches and QFX Series) on page 4369
Syntax	<pre>show isis route <destination> <inet inet6> <instance instance-name> <logical-system (all logical-system-name)> <topology (ipv4-multicast ipv6-multicast ipv6-unicast unicast)></pre>
Syntax (EX Series Switches and QFX Series)	<pre>show isis route <destination> <inet inet6> <instance instance-name> <topology (ipv4-multicast ipv6-multicast ipv6-unicast unicast)></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Display the routes in the IS-IS routing table.
Options	<p>none—Display all routes in the IS-IS routing table for all supported address families for all routing instances.</p> <p>destination—(Optional) Destination address for the route.</p> <p>inet inet6—(Optional) Display inet (IPv4) or inet6 (IPv6) routes, respectively.</p> <p>instance instance-name—(Optional) Display routes for the specified routing instance only.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>topology (ipv4-multicast ipv6-multicast ipv6-unicast unicast)—(Optional) Display routes for the specified topology only, or use unicast to display information, if available, for both IPv4 and IPv6 unicast topologies.</p>
Required Privilege Level	view
List of Sample Output	show isis route logical-system on page 4370 show isis route (CLNS) on page 4370 show isis route on page 4371
Output Fields	Table 342 describes the output fields for the show isis route command. Output fields are listed in the approximate order in which they appear.

Table 342: show isis route Output Fields

Field Name	Field Description
Current version	Number of the current version of the IS-IS routing table.
L1	Version of Level 1 SPF that was run.
L2	Version of Level 2 SPF that was run.
Prefix	Destination of the route.
L	IS-IS level: <ul style="list-style-type: none"> • 1—Level 1 only • 2—Level 2 only • 3—Level 1 and Level 2
Version	Version of SPF that generated the route.
Metric	Metric value associated with the route.
Type	Metric type: int (internal) or ext (external).
Interface	Interface to the next hop.
Via	System identifier of the next hop, displayed as a name if possible.
ISO Routes	ISO routing table entries.
snpa	MAC address.

Sample Output

show isis route logical-system

```

user@host> show isis route logical-system ls1
IS-IS routing table           Current version: L1: 8 L2: 11
Prefix      L Version Metric Type Interface  Via
10.9.7.0/30  2      11     20 int  gr-0/2/0.0  h
10.9.201.1/32 2      11     60 int  gr-0/2/0.0  h
IPv6 Unicast IS-IS routing table           Current version: L1: 9 L2: 11
Prefix      L Version Metric Type Interface  Via
8009:3::a09:3200/126 2      11     20 int  gr-0/2/0.0  h

```

show isis route (CLNS)

```

user@host> show isis route
IS-IS routing table           Current version: L1: 10 L2: 8
IPv4/IPv6 Routes
Prefix      L Version Metric Type Interface  Via
0.0.0.0/0    1      10     10 int  fe-0/0/1.0  ISIS.0
ISO Routes
Prefix L   Version Metric Type Interface  Via  snpa

```

```

0/0
  1      10      10 int fe-0/0/1.0 isis.0 0:12:0:34:0:56
47.0005.80ff.f800.0000.0108.0001/104
  1      10      0 int
47.0005.80ff.f800.0000.0108.0001.1921.6800.4001/152
  1      10      10 int fe-0/0/1.0 isis.0 0:12:0:34:0:56
47.0005.80ff.f800.0000.0108.0001.1921.6800.4002/152
  1      10      20 int fe-0/0/1.0 isis.0 0:12:0:34:0:56
47.0005.80ff.f800.0000.0108.0002/104
  1      10      0 int
47.0005.80ff.f800.0000.0108.0002.1921.6800.4001/152
  1      10      10 int fe-0/0/1.0 isis.0 0:12:0:34:0:56

```

show isis route

```
user@host> show isis route
```

```

IS-IS routing table          Current version: L1: 4 L2: 13
IPv4/IPv6 Routes
-----
Prefix                      L   Version  Metric Type Interface      NH   Via
10.255.71.52/32             2    13        10   int  ae0.0                 IPV4 camaro
10.255.71.238/32            2    13        20   int  so-6/0/0.0           IPV4 olympic
                             as0.0                 IPV4 glacier
10.255.71.239/32            2    13        20   int  so-6/0/0.0           IPV4 olympic
                             ae0.0                 IPV4 camaro
10.255.71.242/32            2    13        10   int  as0.0                 IPV4 glacier
10.255.71.243/32            2    13        10   int  so-6/0/0.0           IPV4 olympic
12.13.0.0/30                 2    13        20   int  so-6/0/0.0           IPV4 olympic
12.15.0.0/30                 2    13        20   int  so-6/0/0.0           IPV4 olympic
13.15.0.0/30                 2    13        30   int  ae0.0                 IPV4 camaro
                             so-6/0/0.0           IPV4 olympic
                             as0.0                 IPV4 glacier
13.16.0.0/30                 2    13        25   int  as0.0                 IPV4 glacier
14.15.0.0/30                 2    13        20   int  ae0.0                 IPV4 camaro
192.2.1.0/30                 2    13        30   int  so-6/0/0.0           IPV4 olympic
                             as0.0                 IPV4 glacier
1eee::/64                    2    13        30   int  so-6/0/0.0           IPV6 olympic
                             as0.0                 IPV6 glacier
abcd::10:255:71:52/128      2    13        10   int  ae0.0                 IPV6 camaro
abcd::10:255:71:238/128     2    13        20   int  so-6/0/0.0           IPV6 olympic

```

					as0.0	IPv6 glacier
abcd::10:255:71:239/128	2	13	20	int	so-6/0/0.0	IPv6 olympic
					ae0.0	IPv6 camaro
abcd::10:255:71:242/128	2	13	10	int	as0.0	IPv6 glacier
abcd::10:255:71:243/128	2	13	10	int	so-6/0/0.0	IPv6 olympic

show isis spf

List of Syntax [Syntax on page 4373](#)
[Syntax \(EX Series Switches\) on page 4373](#)

Syntax show isis spf (brief | log | results)
 <instance *instance-name*>
 <level (1 | 2)>
 <logical-system (all | *logical-system-name*)>
 <topology (ipv4-multicast | ipv6-multicast | ipv6-unicast | unicast)>

Syntax (EX Series Switches) show isis spf (brief | log | results)
 <instance *instance-name*>
 <level (1 | 2)>
 <topology (ipv4-multicast | ipv6-multicast | ipv6-unicast | unicast)>

Release Information Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.0 for EX Series switches.

Description Display information about IS-IS shortest-path-first (SPF) calculations.

Options **brief**—Display an overview of SPF calculations.

instance *instance instance-name*—(Optional) Display SPF calculations for the specified routing instance.

level (1 | 2)—(Optional) Display SPF calculations for the specified IS-IS level.

log—Display the log of SPF calculations.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

results—Display the results of SPF calculations.

topology (ipv4-multicast | ipv6-multicast | ipv6-unicast | unicast)—(Optional) Display SPF calculations for the specified topology only.

Required Privilege Level view

List of Sample Output [show isis spf log on page 4374](#)
[show isis spf results logical-system on page 4375](#)
[show isis spf results \(CLNS\) on page 4376](#)

Output Fields [Table 343](#) describes the output fields for the **show isis spf** command. Output fields are listed in the approximate order in which they appear.

Table 343: show isis spf Output Fields

Field Name	Field Description
Node	System ID of a node.

Table 343: show isis spf Output Fields (*continued*)

Field Name	Field Description
Metric	Metric to the node.
Interface	Interface of the next hop.
Via	System ID of the next hop.
SNPA	Subnetwork point of attachment (MAC address of the next hop).
Start time	(log option only) Time that the SPF computation started.
Elapsed (secs)	(log option only) Length of time, in seconds, required to complete the SPF computation.
Count	(log option only) Number of times the SPF was triggered.
Reason	(log option only) Reason that the SPF computation was completed.

Sample Output

show isis spf log

```

user@host> show isis spf log logical-system lsl
IS-IS level 1 SPF log:
Start time           Elapsed (secs) Count Reason
Fri Oct 31 12:41:18   0.000069    1 Reconfig
Fri Oct 31 12:41:18   0.000107    3 Updated LSP fix.00-00
Fri Oct 31 12:41:18   0.000050    3 Address change on so-1/2/0.0
Fri Oct 31 12:41:23   0.000033    1 Updated LSP fix.00-00
Fri Oct 31 12:41:28   0.000178    5 New adjacency scat on ge-1/1/0.0
Fri Oct 31 12:41:59   0.000060    1 Updated LSP fix.00-00
Fri Oct 31 12:42:30   0.000161    2 Multi area attachment change
Fri Oct 31 12:56:58   0.000198    1 Periodic SPF
Fri Oct 31 13:10:29   0.000209    1 Periodic SPF
IS-IS level 2 SPF log:
Start time           Elapsed (secs) Count Reason
Fri Oct 31 12:41:18   0.000035    1 Reconfig
Fri Oct 31 12:41:18   0.000047    2 Updated LSP fix.00-00
Fri Oct 31 12:41:18   0.000043    5 Address change on gr-0/2/0.0
Fri Oct 31 12:41:23   0.000022    1 Updated LSP fix.00-00
Fri Oct 31 12:41:59   0.000144    3 New adjacency h on gr-0/2/0.0
Fri Oct 31 12:42:30   0.000257    3 New LSP skag.00-00
Fri Oct 31 12:54:37   0.000195    1 Periodic SPF
Fri Oct 31 12:55:50   0.000178    1 Updated LSP fix.00-00
Fri Oct 31 12:55:55   0.000174    1 Updated LSP h.00-00
Fri Oct 31 12:55:58   0.000176    1 Updated LSP skag.00-00
Fri Oct 31 13:08:14   0.000198    1 Periodic SPF
IPv6 Unicast IS-IS level 1 SPF log:
Start time           Elapsed (secs) Count Reason
Fri Oct 31 12:41:18   0.000028    1 Reconfig
Fri Oct 31 12:41:18   0.000043    3 Updated LSP fix.00-00

```

```

Fri Oct 31 12:41:18      0.000112    4 Updated LSP fix.00-00
Fri Oct 31 12:41:23      0.000059    1 Updated LSP fix.00-00
Fri Oct 31 12:41:25      0.000041    1 Updated LSP fix.00-00
Fri Oct 31 12:41:28      0.000103    5 New adjacency scat on ge-1/1/0.0
Fri Oct 31 12:41:59      0.000040    1 Updated LSP fix.00-00
Fri Oct 31 12:42:30      0.000118    2 Multi area attachment change
Fri Oct 31 12:56:08      0.000289    1 Periodic SPF
Fri Oct 31 13:11:07      0.000214    1 Periodic SPF
IPV6 Unicast IS-IS level 2 SPF log:

```

```

Start time      Elapsed (secs) Count Reason
Fri Oct 31 12:41:18      0.000027    1 Reconfig
Fri Oct 31 12:41:18      0.000039    2 Updated LSP fix.00-00
Fri Oct 31 12:41:18      0.000049    6 Updated LSP fix.00-00
Fri Oct 31 12:41:23      0.000025    1 Updated LSP fix.00-00
Fri Oct 31 12:41:25      0.000023    1 Updated LSP fix.00-00
Fri Oct 31 12:41:59      0.000087    3 New adjacency h on gr-0/2/0.0
Fri Oct 31 12:42:30      0.000123    3 New LSP skag.00-00
Fri Oct 31 12:55:50      0.000121    1 Updated LSP fix.00-00
Fri Oct 31 12:55:55      0.000121    1 Updated LSP h.00-00
Fri Oct 31 12:55:58      0.000121    1 Updated LSP skag.00-00
Fri Oct 31 13:09:46      0.000201    1 Periodic SPF
...

```

show isis spf results logical-system

```
user@host> show isis spf results logical-system ls1
```

```
IS-IS level 1 SPF results:
```

Node	Metric	Interface	Via	SNPA
scat.00	10	ge-1/1/0.0	scat	0:90:69:a6:48:9d
	20	10.9.1.0/30		
fix.02	10			
fix.00	0			
	10	10.9.1.0/30		
	10	10.9.5.0/30		
	10	10.9.6.0/30		
	20	10.9.7.0/30		
	60	10.9.201.1/32		
3 nodes				

```
IS-IS level 2 SPF results:
```

Node	Metric	Interface	Via	SNPA
skag.00	20	gr-0/2/0.0	h	
	30	10.9.7.0/30		
skag.02	20	gr-0/2/0.0	h	
h.00	10	gr-0/2/0.0	h	
	20	10.9.6.0/30		
	20	10.9.7.0/30		
	60	10.9.201.1/32		
fix.00	0			
	10	10.9.1.0/30		
	10	10.9.5.0/30		
	10	10.9.6.0/30		
4 nodes				

```
IPV6 Unicast IS-IS level 1 SPF results:
```

Node	Metric	Interface	Via	SNPA
scat.00	10	ge-1/1/0.0	scat	0:90:69:a6:48:9d
		ge-1/1/0.0	scat	0:90:69:a6:48:9d
	20	8009:1::a09:1400/126		
fix.02	10			

```

fix.00      0
            10      8009:1::a09:1400/126
            10      8009:2::a09:1e00/126
            20      8009:3::a09:3200/126
            10      8009:4::a09:2800/126

  3 nodes

IPv6 Unicast IS-IS level 2 SPF results:
Node      Metric      Interface      Via      SNPA
skag.00    20      gr-0/2/0.0    h
           30      8009:3::a09:3200/126
           20      gr-0/2/0.0    h
skag.02    20      gr-0/2/0.0    h
           10      gr-0/2/0.0    h
h.00       10      gr-0/2/0.0    h
           20      8009:3::a09:3200/126
           20      8009:4::a09:2800/126
fix.00     0
           10      8009:1::a09:1400/126
           10      8009:2::a09:1e00/126
           10      8009:4::a09:2800/126

  4 nodes

Multicast IS-IS level 1 SPF results:
Node      Metric      Interface      Via      SNPA
scat.00    10      ge-1/1/0.0    scat    0:90:69:a6:48:9d
fix.02     10
fix.00     0
  3 nodes

Multicast IS-IS level 2 SPF results:
Node      Metric      Interface      Via      SNPA
skag.00    20      gr-0/2/0.0    h
skag.02    20      gr-0/2/0.0    h
h.00       10      gr-0/2/0.0    h
fix.00     0
  4 nodes
...

```

show isis spf results (CLNS)

```

user@host> show isis spf results
IS-IS level 1 SPF results:
Node      Metric      Interface      Via      SNPA
skag.00 10      fe-0/0/1.0    toothache 0:12:0:34:0:56
           20      fe-0/0/1.0    toothache 0:12:0:34:0:56
           10      192.168.37.64/29
           20      192.168.37.64/29
           20      192.168.37.64/29
pro1-a.02 10
pro1-a.00 0
           0      10.255.245.1/32
           10      192.168.37.64/29
           0      192.168.37.64/29
           0      192.168.37.64/29

  3 nodes

IS-IS level 2 SPF results:
Node      Metric      Interface      Via      SNPA
skag.00 10      fe-0/0/1.0    toothache 0:12:0:34:0:56
           10      fe-0/0/1.0    toothache 0:12:0:34:0:56

```


	20	10.255.245.1/32
	20	192.168.37.64/29
	20	47.0005.80ff.f800.0000.0109.0010/104
pro1-a.02	10	
pro1-a.00	0	
	0	10.255.245.1/32
	10	192.168.37.64/29
3 nodes		

show isis statistics

List of Syntax	Syntax on page 4378 Syntax (EX Series Switches and QFX Series) on page 4378
Syntax	<code>show isis statistics</code> <code><instance <i>instance-name</i>></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switches and QFX Series)	<code>show isis statistics</code> <code><instance <i>instance-name</i>></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Display statistics about IS-IS traffic.
Options	none —Display IS-IS traffic statistics for all routing instances. instance <i>instance-name</i> —(Optional) Display statistics for the specified routing instance. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• clear isis statistics on page 4336
List of Sample Output	show isis statistics on page 4380
Output Fields	Table 344 describes the output fields for the show isis statistics command. Output fields are listed in the approximate order in which they appear.

Table 344: show isis statistics Output Fields

Field Name	Field Description
PDU type	<p>PDU type:</p> <ul style="list-style-type: none"> • CSNP—Complete sequence number PDUs contain a complete list of all link-state PDUs in the IS-IS database. CSNPs are sent periodically on all links, and the receiving systems use the information in the CSNP to update and synchronize their link-state PDU databases. The designated router multicasts CSNPs on broadcast links in place of sending explicit acknowledgments for each link-state PDU. • IIH—IS-IS hello packets are broadcast to discover the identity of neighboring IS-IS systems and to determine whether the neighbors are Level 1 or Level 2 intermediate systems. • LSP—Link-state PDUs contain information about the state of adjacencies to neighboring IS-IS systems. Link-state PDUs are flooded periodically throughout an area. • PSNP—Partial sequence number PDUs are sent multicast by a receiver when it detects that it is missing a link-state PDU (when its link-state PDU database is out of date). The receiver sends a PSNP to the system that transmitted the CSNP, effectively requesting that the missing link-state PDU be transmitted. That routing device, in turn, forwards the missing link-state PDU to the requesting routing device. • Unknown—The PDU type is unknown.
Received	Number of PDUs received since IS-IS started or since the statistics were set to zero.
Processed	Number of PDUs received less the number dropped.
Drops	Number of PDUs dropped.
Sent	Number of PDUs transmitted since IS-IS started or since the statistics were set to zero.
Rexmit	Number of PDUs retransmitted since IS-IS started or since the statistics were set to zero.
Total packets received/sent	Total number of PDUs received and transmitted since IS-IS started or since the statistics were set to zero.
SNP queue length	Number of CSPN and PSNP packets currently waiting in the queue for processing. This value is almost always 0.
LSP queue length	Number of link-state PDUs waiting in the queue for processing. This value is almost always 0.
SPF runs	Number of shortest-path-first (SPF) calculations that have been performed. If this number is incrementing rapidly, it indicates that the network is unstable.
Fragments rebuilt	Number of link-state PDU fragments that the local system has computed.
LSP regenerations	Number of link-state PDUs that have been regenerated. A link-state PDU is regenerated when it is nearing the end of its lifetime and it has not changed.
Purges initiated	Number of purges that the system initiated. A purge is initiated if the software decides that a link-state PDU must be removed from the network.

Sample Output

show isis statistics

```
user@host> show isis statistics
```

```
IS-IS statistics for merino:
```

PDU type	Received	Processed	Drops	Sent	Rexmit
LSP	12227	12227	0	8184	683
IIH	113808	113808	0	115817	0
CSNP	198868	198868	0	198934	0
PSNP	6985	6979	6	8274	0
Unknown	0	0	0	0	0
Totals	331888	331882	6	331209	683

```
Total packets received: 331888 Sent: 331892
```

```
SNP queue length:      0 Drops:      0  
LSP queue length:      0 Drops:      0
```

```
SPF runs:              1014  
Fragments rebuilt:     1038  
LSP regenerations:     425  
Purges initiated:      0
```

OSPF Feature Guide for QFX10000 Switches

PART 66

Overview

- [OSPF Overview on page 4385](#)

CHAPTER 162

OSPF Overview

- [OSPF Overview on page 4386](#)
- [OSPF Packets Overview on page 4391](#)
- [Understanding OSPF External Metrics on page 4393](#)

OSPF Overview

OSPF is an interior gateway protocol (IGP) that routes packets within a single autonomous system (AS). OSPF uses link-state information to make routing decisions, making route calculations using the shortest-path-first (SPF) algorithm (also referred to as the Dijkstra algorithm). Each router running OSPF floods link-state advertisements throughout the AS or area that contain information about that router's attached interfaces and routing metrics. Each router uses the information in these link-state advertisements to calculate the least cost path to each network and create a routing table for the protocol.

Junos OS supports OSPF version 2 (OSPFv2) and OSPF version 3 (OSPFv3), including virtual links, stub areas, and for OSPFv2, authentication. Junos OS does not support type-of-service (ToS) routing.

OSPF was designed for the Transmission Control Protocol/Internet Protocol (TCP/IP) environment and as a result explicitly supports IP subnetting and the tagging of externally derived routing information. OSPF also provides for the authentication of routing updates.

OSPF routes IP packets based solely on the destination IP address contained in the IP packet header. OSPF quickly detects topological changes, such as when router interfaces become unavailable, and calculates new loop-free routes quickly and with a minimum of routing overhead traffic.



NOTE: On SRX Series devices, when only one link-protection is configured under the OSPF interface, the device does not install an alternative route in the forwarding table. When the per-packet load-balancing is enabled as a workaround, the device does not observe both the OSPF metric and sending the traffic through both the interfaces.

An OSPF AS can consist of a single area, or it can be subdivided into multiple areas. In a single-area OSPF network topology, each router maintains a database that describes the topology of the AS. Link-state information for each router is flooded throughout the AS. In a multiarea OSPF topology, each router maintains a database that describes the topology of its area, and link-state information for each router is flooded throughout that area. All routers maintain summarized topologies of other areas within an AS. Within each area, OSPF routers have identical topological databases. When the AS or area topology changes, OSPF ensures that the contents of all routers' topological databases converge quickly.

All OSPFv2 protocol exchanges can be authenticated. OSPFv3 relies on IPsec to provide this functionality. This means that only trusted routers can participate in the AS's routing. A variety of authentication schemes can be used. A single authentication scheme is configured for each area, which enables some areas to use stricter authentication than others.

Externally derived routing data (for example, routes learned from BGP) is passed transparently throughout the AS. This externally derived data is kept separate from the OSPF link-state data. Each external route can be tagged by the advertising router, enabling the passing of additional information between routers on the boundaries of the AS.



NOTE: By default, Junos OS is compatible with RFC 1583, *OSPF Version 2*. In Junos OS Release 8.5 and later, you can disable compatibility with RFC 1583 by including the `no-rfc-1583` statement. For more information, see [“Example: Disabling OSPFv2 Compatibility with RFC 1583”](#) on page 4428.

This topic describes the following information:

- [OSPF Default Route Preference Values on page 4388](#)
- [OSPF Routing Algorithm on page 4388](#)
- [OSPF Three-Way Handshake on page 4389](#)
- [OSPF Version 3 on page 4390](#)

OSPF Default Route Preference Values

The Junos OS routing protocol process assigns a default preference value to each route that the routing table receives. The default value depends on the source of the route. The preference value is from 0 through 4,294,967,295 ($2^{32} - 1$), with a lower value indicating a more preferred route. [Table 345](#) lists the default preference values for OSPF.

Table 345: Default Route Preference Values for OSPF

How Route Is Learned	Default Preference	Statement to Modify Default Preference
OSPF internal route	10	OSPF preference
OSPF AS external routes	150	OSPF external-preference

OSPF Routing Algorithm

OSPF uses the shortest-path-first (SPF) algorithm, also referred to as the Dijkstra algorithm, to determine the route to each destination. All routing devices in an area run this algorithm in parallel, storing the results in their individual topological databases. Routing devices with interfaces to multiple areas run multiple copies of the algorithm. This section provides a brief summary of how the SPF algorithm works.

When a routing device starts, it initializes OSPF and waits for indications from lower-level protocols that the router interfaces are functional. The routing device then uses the OSPF hello protocol to acquire neighbors, by sending hello packets to its neighbors and receiving their hello packets.

On broadcast or nonbroadcast multiaccess networks (physical networks that support the attachment of more than two routing devices), the OSPF hello protocol elects a designated router for the network. This routing device is responsible for sending *link-state advertisements* (LSAs) that describe the network, which reduces the amount of network traffic and the size of the routing devices' topological databases.

The routing device then attempts to form *adjacencies* with some of its newly acquired neighbors. (On multiaccess networks, only the designated router and backup designated router form adjacencies with other routing devices.) Adjacencies determine the distribution

of routing protocol packets. Routing protocol packets are sent and received only on adjacencies, and topological database updates are sent only along adjacencies. When adjacencies have been established, pairs of adjacent routers synchronize their topological databases.

A routing device sends LSA packets to advertise its state periodically and when its state changes. These packets include information about the routing device's adjacencies, which allows detection of nonoperational routing devices.

Using a reliable algorithm, the routing device floods LSAs throughout the area, which ensures that all routing devices in an area have exactly the same topological database. Each routing device uses the information in its topological database to calculate a shortest-path tree, with itself as the root. The routing device then uses this tree to route network traffic.


The description of the SPF algorithm up to this point has explained how the algorithm works within a single area (*intra-area routing*). For internal routers to be able to route to destinations outside the area (*interarea routing*), the area border routers must inject additional routing information into the area. Because the area border routers are connected to the backbone, they have access to complete topological data about the backbone. The area border routers use this information to calculate paths to all destinations outside its area and then advertise these paths to the area's internal routers.

Autonomous system (AS) boundary routers flood information about external autonomous systems throughout the AS, except to stub areas. Area border routers are responsible for advertising the paths to all AS boundary routers.

OSPF Three-Way Handshake

OSPF creates a topology map by flooding LSAs across OSPF-enabled links. LSAs announce the presence of OSPF-enabled interfaces to adjacent OSPF interfaces. The exchange of LSAs establishes bidirectional connectivity between all adjacent OSPF interfaces (neighbors) using a three-way handshake, as shown in [Figure 108](#).

Figure 108: OSPF Three-Way Handshake



In [Figure 108](#), Router A sends hello packets out all its OSPF-enabled interfaces when it comes online. Router B receives the packet, which establishes that Router B can receive traffic from Router A. Router B generates a response to Router A to acknowledge receipt of the hello packet. When Router A receives the response, it establishes that Router B can receive traffic from Router A. Router A then generates a final response packet to inform Router B that Router A can receive traffic from Router B. This three-way handshake ensures bidirectional connectivity.

As new neighbors are added to the network or existing neighbors lose connectivity, the adjacencies in the topology map are modified accordingly through the exchange (or absence) of LSAs. These LSAs advertise only the incremental changes in the network, which helps minimize the amount of OSPF traffic on the network. The adjacencies are shared and used to create the network topology in the topological database.

OSPF Version 3

OSPFv3 is a modified version of OSPF that supports IP version 6 (IPv6) addressing. OSPFv3 differs from OSPFv2 in the following ways:

- All neighbor ID information is based on a 32-bit router ID.
- The protocol runs per link rather than per subnet.
- Router and network link-state advertisements (LSAs) do not carry prefix information.
- Two new LSA types are included: link-LSA and intra-area-prefix-LSA.
- Flooding scopes are as follows:
 - Link-local
 - Area
 - AS
- Link-local addresses are used for all neighbor exchanges except virtual links.
- Authentication is removed. The IPv6 authentication header relies on the IP layer.
- The packet format has changed as follows:
 - Version number 2 is now version number 3.
 - The **db** option field has been expanded to 24 bits.
 - Authentication information has been removed.
 - Hello messages do not have address information.
 - Two new option bits are included: **R** and **V6**.
- Type 3 summary LSAs have been renamed *inter-area-prefix-LSAs*.
- Type 4 summary LSAs have been renamed *inter-area-router-LSAs*.

Related Documentation

- [Understanding OSPF Areas and Backbone Areas on page 4404](#)
- [Understanding OSPF Configurations](#)
- [Example: Disabling OSPFv2 Compatibility with RFC 1583 on page 4428](#)

OSPF Packets Overview

There are several types of link-state advertisement (LSA) packets.

This topic describes the following information:

- [OSPF Packet Header on page 4391](#)
- [Hello Packets on page 4391](#)
- [Database Description Packets on page 4392](#)
- [Link-State Request Packets on page 4392](#)
- [Link-State Update Packets on page 4392](#)
- [Link-State Acknowledgment Packets on page 4392](#)
- [Link-State Advertisement Packet Types on page 4393](#)

OSPF Packet Header

All OSPFv2 packets have a common 24-byte header, and OSPFv3 packets have a common 16-byte header, that contains all information necessary to determine whether OSPF should accept the packet. The header consists of the following fields:

- Version number—The current OSPF version number. This can be either 2 or 3.
- Type—Type of OSPF packet.
- Packet length—Length of the packet, in bytes, including the header.
- Router ID—IP address of the router from which the packet originated.
- Area ID—Identifier of the area in which the packet is traveling. Each OSPF packet is associated with a single area. Packets traveling over a virtual link are labeled with the backbone area ID, 0.0.0.0.
- Checksum—Fletcher checksum.
- Authentication—(OSPFv2 only) Authentication scheme and authentication information.
- Instance ID—(OSPFv3 only) Identifier used when there are multiple OSPFv3 realms configured on a link.

Hello Packets

Routers periodically send hello packets on all interfaces, including virtual links, to establish and maintain neighbor relationships. Hello packets are multicast on physical networks that have a multicast or broadcast capability, which enables dynamic discovery of neighboring routers. (On nonbroadcast networks, dynamic neighbor discovery is not possible, so you must configure all neighbors statically as described in [“Example: Configuring an OSPFv2 Interface on a Nonbroadcast Multiaccess Network” on page 4435.](#))

Hello packets consist of the OSPF header plus the following fields:

- Network mask—(OSPFv2 only) Network mask associated with the interface.
- Hello interval—How often the router sends hello packets. All routers on a shared network must use the same hello interval.
- Options—Optional capabilities of the router.
- Router priority—The router's priority to become the designated router.
- Router dead interval—How long the router waits without receiving any OSPF packets from a router before declaring that router to be down. All routers on a shared network must use the same router dead interval.
- Designated router—IP address of the designated router.
- Backup designated router—IP address of the backup designated router.
- Neighbor—IP addresses of the routers from which valid hello packets have been received within the time specified by the router dead interval.

Database Description Packets

When initializing an adjacency, OSPF exchanges database description packets, which describe the contents of the topological database. These packets consist of the OSPF header, packet sequence number, and the link-state advertisement's header.

Link-State Request Packets

When a router detects that portions of its topological database are out of date, it sends a link-state request packet to a neighbor requesting a precise instance of the database. These packets consist of the OSPF header plus fields that uniquely identify the database information that the router is seeking.

Link-State Update Packets

Link-state update packets carry one or more link-state advertisements one hop farther from their origin. The router multicasts (floods) these packets on physical networks that support multicast or broadcast mode. The router acknowledges all link-state update packets and, if retransmission is necessary, sends the retransmitted advertisements unicast.

Link-state update packets consist of the OSPF header plus the following fields:

- Number of advertisements—Number of link-state advertisements included in this packet.
- Link-state advertisements—The link-state advertisements themselves.

Link-State Acknowledgment Packets

The router sends link-state acknowledgment packets in response to link-state update packets to verify that the update packets have been received successfully. A single acknowledgment packet can include responses to multiple update packets.

Link-state acknowledgment packets consist of the OSPF header plus the link-state advertisement header.

Link-State Advertisement Packet Types

Link-state request, link-state update, and link-state acknowledgment packets are used to reliably flood link-state advertisement packets. OSPF sends the following types of link-state advertisements:

- Router link advertisements—Are sent by all routers to describe the state and cost of the router's links to the area. These link-state advertisements are flooded throughout a single area only.
- Network link advertisements—Are sent by designated routers to describe all the routers attached to the network. These link-state advertisements are flooded throughout a single area only.
- Summary link advertisements—Are sent by area border routers to describe the routes that they know about in other areas. There are two types of summary link advertisements: those used when the destination is an IP network, and those used when the destination is an AS boundary router. Summary link advertisements describe interarea routes, that is, routes to destinations outside the area but within the AS. These link-state advertisements are flooded throughout the advertisement's associated areas.
- AS external link advertisement—Are sent by AS boundary routers to describe external routes that they know about. These link-state advertisements are flooded throughout the AS (except for stub areas).

Each link-state advertisement type describes a portion of the OSPF routing domain. All link-state advertisements are flooded throughout the AS.

Each link-state advertisement packet begins with a common 20-byte header.

Related Documentation

- [OSPF Overview on page 4386](#)
- [Understanding OSPF Areas on page 4397](#)
- [Understanding OSPF Configurations](#)
- [Understanding OSPF Designated Router on page 4400](#)
- [Understanding OSPFv2 Authentication](#)
- [Understanding OSPF Timers on page 4475](#)

Understanding OSPF External Metrics

When OSPF exports route information from external autonomous systems (ASs), it includes a cost, or *external metric*, in the route. OSPF supports two types of external metrics: Type 1 and Type 2. The difference between the two metrics is how OSPF calculates the cost of the route.

- Type 1 external metrics are equivalent to the link-state metric, where the cost is equal to the sum of the internal costs plus the external cost. This means that Type 1 external metrics include the external cost to the destination as well as the cost (metric) to reach the AS boundary router.
- Type 2 external metrics are greater than the cost of any path internal to the AS. Type 2 external metrics use only the external cost to the destination and ignore the cost (metric) to reach the AS boundary router.

By default, OSPF uses the Type 2 external metric.

Both Type 1 and Type 2 external metrics can be present in the AS at the same time. In that event, Type 1 external metrics always takes the precedence.

Type 1 external paths are always preferred over Type 2 external paths. When all paths are Type 2 external paths, the paths with the smallest advertised Type 2 metric are always preferred.

**Related
Documentation**

- [Example: Dynamically Adjusting OSPF Interface Metrics Based on Bandwidth on page 4466](#)

PART 67

Configuring OSPF

- [Basic OSPF Area Configuration on page 4397](#)
- [Advanced OSPF Area Configuration on page 4413](#)
- [OSPF Interface Configuration on page 4431](#)
- [OSPF Route Control Configuration on page 4451](#)
- [OSPF Fault Detection Configuration on page 4475](#)
- [OSPF Redundancy Features Configuration on page 4495](#)
- [OSPF Traffic Engineering Configuration on page 4511](#)
- [OSPF Database Protection Configuration on page 4523](#)
- [OSPF Policy Configuration on page 4527](#)

Basic OSPF Area Configuration

- [Understanding OSPF Areas on page 4397](#)
- [Examples: Configuring OSPF Designated Routers on page 4399](#)
- [Examples: Configuring OSPF Areas on page 4404](#)

Understanding OSPF Areas

In OSPF, a single autonomous system (AS) can be divided into smaller groups called *areas*. This reduces the number of link-state advertisements (LSAs) and other OSPF overhead traffic sent on the network, and it reduces the size of the topology database that each router must maintain. The routing devices that participate in OSPF routing perform one or more functions based on their location in the network.

This topic describes the following OSPF area types and routing device functions:

- [Areas on page 4397](#)
- [Area Border Routers on page 4398](#)
- [Backbone Areas on page 4398](#)
- [AS Boundary Routers on page 4398](#)
- [Backbone Router on page 4398](#)
- [Internal Router on page 4398](#)
- [Stub Areas on page 4398](#)
- [Not-So-Stubby Areas on page 4399](#)
- [Transit Areas on page 4399](#)

Areas

An *area* is a set of networks and hosts within an AS that have been administratively grouped together. We recommend that you configure an area as a collection of contiguous IP subnetted networks. Routing devices that are wholly within an area are called *internal routers*. All interfaces on internal routers are directly connected to networks within the area.

The topology of an area is hidden from the rest of the AS, thus significantly reducing routing traffic in the AS. Also, routing within the area is determined only by the area's topology, providing the area with some protection from bad routing data.

All routing devices within an area have identical topology databases.

Area Border Routers

Routing devices that belong to more than one area and connect one or more OSPF areas to the backbone area are called *area border routers* (ABRs). At least one interface is within the backbone while another interface is in another area. ABRs also maintain a separate topological database for each area to which they are connected.

Backbone Areas

An OSPF *backbone area* consists of all networks in area ID 0.0.0.0, their attached routing devices, and all ABRs. The backbone itself does not have any ABRs. The backbone distributes routing information between areas. The backbone is simply another area, so the terminology and rules of areas apply: a routing device that is directly connected to the backbone is an internal router on the backbone, and the backbone's topology is hidden from the other areas in the AS.

The routing devices that make up the backbone must be physically contiguous. If they are not, you must configure *virtual links* to create the appearance of backbone connectivity. You can create virtual links between any two ABRs that have an interface to a common nonbackbone area. OSPF treats two routing devices joined by a virtual link as if they were connected to an unnumbered point-to-point network.

AS Boundary Routers

Routing devices that exchange routing information with routing devices in non-OSPF networks are called *AS boundary routers*. They advertise externally learned routes throughout the OSPF AS. Depending on the location of the AS boundary router in the network, it can be an ABR, a backbone router, or an internal router (with the exception of stub areas). Internal routers within a stub area cannot be an AS boundary router because stub areas cannot contain any Type 5 LSAs.

Routing devices within the area where the AS boundary router resides know the path to that AS boundary router. Any routing device outside the area only knows the path to the nearest ABR that is in the same area where the AS boundary router resides.

Backbone Router

Backbone routers are routing devices that have one or more interfaces connected to the OSPF backbone area (area ID 0.0.0.0).

Internal Router

Routing devices that connect to only one OSPF area are called *internal routers*. All interfaces on internal routers are directly connected to networks within a single area.

Stub Areas

Stub areas are areas through which or into which AS external advertisements are not flooded. You might want to create stub areas when much of the topological database consists of AS external advertisements. Doing so reduces the size of the topological

databases and therefore the amount of memory required on the internal routers in the stub area.

Routing devices within a stub area rely on the default routes originated by the area's ABR to reach external AS destinations. You must configure the **default-metric** option on the ABR before it advertises a default route. Once configured, the ABR advertises a default route in place of the external routes that are not being advertised within the stub area, so that routing devices in the stub area can reach destinations outside the area.

The following restrictions apply to stub areas: you cannot create a virtual link through a stub area, a stub area cannot contain an AS boundary router, the backbone cannot be a stub area, and you cannot configure an area as both a stub area and a not-so-stubby area.

Not-So-Stubby Areas

An OSPF stub area has no external routes in it, so you cannot redistribute from another protocol into a stub area. A *not-so-stubby area* (NSSA) allows external routes to be flooded within the area. These routes are then leaked into other areas. However, external routes from other areas still do not enter the NSSA.

The following restriction applies to NSSAs: you cannot configure an area as both a stub area and an NSSA.

Transit Areas

Transit areas are used to pass traffic from one adjacent area to the backbone (or to another area if the backbone is more than two hops away from an area). The traffic does not originate in, nor is it destined for, the transit area.

Related Documentation

- [OSPF Overview on page 4386](#)
- [OSPF Packets Overview on page 4391](#)
- [Understanding OSPF Configurations](#)
- [Understanding OSPF Areas and Backbone Areas on page 4404](#)
- [Understanding OSPF Stub Areas, Totally Stubby Areas, and Not-So-Stubby Areas on page 4413](#)

Examples: Configuring OSPF Designated Routers

- [OSPF Designated Router Overview on page 4400](#)
- [Example: Configuring an OSPF Router Identifier on page 4400](#)
- [Example: Controlling OSPF Designated Router Election on page 4402](#)

OSPF Designated Router Overview

Large LANs that have many routing devices and therefore many OSPF adjacencies can produce heavy control-packet traffic as link-state advertisements (LSAs) are flooded across the network. To alleviate the potential traffic problem, OSPF uses designated routers on all multiaccess networks (broadcast and nonbroadcast multiaccess [NBMA] networks types). Rather than broadcasting LSAs to all their OSPF neighbors, the routing devices send their LSAs to the designated router. Each multiaccess network has a designated router, which performs two main functions:

- Originate network link advertisements on behalf of the network.
- Establish adjacencies with all routing devices on the network, thus participating in the synchronizing of the link-state databases.

In LANs, the election of the designated router takes place when the OSPF network is initially established. When the first OSPF links are active, the routing device with the highest router identifier (defined by the **router-id** configuration value, which is typically the IP address of the routing device, or the loopback address) is elected the designated router. The routing device with the second highest router identifier is elected the backup designated router. If the designated router fails or loses connectivity, the backup designated router assumes its role and a new backup designated router election takes place between all the routers in the OSPF network.

OSPF uses the router identifier for two main purposes: to elect a designated router, unless you manually specify a priority value, and to identify the routing device from which a packet is originated. At designated router election, the router priorities are evaluated first, and the routing device with the highest priority is elected designated router. If router priorities tie, the routing device with the highest router identifier, which is typically the routing device's IP address, is chosen as the designated router. If you do not configure a router identifier, the IP address of the first interface to come online is used. This is usually the loopback interface. Otherwise, the first hardware interface with an IP address is used.

At least one routing device on each logical IP network or subnet must be eligible to be the designated router for OSPFv2. At least one routing device on each logical link must be eligible to be the designated router for OSPFv3.

By default, routing devices have a priority of 128. A priority of 0 marks the routing device as ineligible to become the designated router. A priority of 1 means the routing device has the least chance of becoming a designated router. A priority of 255 means the routing device is always the designated router.

Example: Configuring an OSPF Router Identifier

This example shows how to configure an OSPF router identifier.

- [Requirements on page 4401](#)
- [Overview on page 4401](#)
- [Configuration on page 4401](#)
- [Verification on page 4402](#)

Requirements

Before you begin:

- Identify the interfaces on the routing device that will participate in OSPF. You must enable OSPF on all interfaces within the network on which OSPF traffic is to travel.
- Configure the device interfaces. See the *Interfaces Feature Guide for Security Devices*

Overview

The router identifier is used by OSPF to identify the routing device from which a packet originated. Junos OS selects a router identifier according to the following set of rules:

1. By default, Junos OS selects the lowest configured physical IP address of an interface as the router identifier.
2. If a loopback interface is configured, the IP address of the loopback interface becomes the router identifier.
3. If multiple loopback interfaces are configured, the lowest loopback address becomes the router identifier.
4. If a router identifier is explicitly configured using the **router-id address** statement under the **[edit routing-options]** hierarchy level, the above three rules are ignored.



NOTE: 1. The router identifier behavior described here holds good even when configured under **[edit routing-instances routing-instance-name routing-options]** and **[edit logical-systems logical-system-name routing-instances routing-instance-name routing-options]** hierarchy levels.

2. If the router identifier is modified in a network, the link-state advertisements (LSAs) advertised by the previous router identifier are retained in the OSPF database until the LSA retransmit interval has timed out. Hence, it is strongly recommended that you explicitly configure the router identifier under the **[edit routing-options]** hierarchy level to avoid unpredictable behavior if the interface address on a loopback interface changes.

In this example, you configure the OSPF router identifier by setting its router ID value to the IP address of the device, which is 177.162.4.24.

Configuration

CLI Quick Configuration

To quickly configure an OSPF router identifier, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
[edit]
set routing-options router-id 177.162.4.24
```

**Step-by-Step
Procedure**

To configure an OSPF router identifier:

1. Configure the OSPF router identifier by entering the **[router-id]** configuration value.

```
[edit]  
user@host# set routing-options router-id 177.162.4.24
```
2. If you are done configuring the device, commit the configuration.

```
[edit]  
user@host# commit
```

Results

Confirm your configuration by entering the **show routing-options router-id** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show routing-options router-id  
router-id 177.162.4.24;
```

Verification

After you configure the router ID and activate OSPF on the routing device, the router ID is referenced by multiple OSPF operational mode commands that you can use to monitor and troubleshoot the OSPF protocol. The router ID fields are clearly marked in the output.

Example: Controlling OSPF Designated Router Election

This example shows how to control OSPF designated router election.

- [Requirements on page 4402](#)
- [Overview on page 4402](#)
- [Configuration on page 4403](#)
- [Verification on page 4403](#)

Requirements

Before you begin:

- Configure the device interfaces. See the *Interfaces Feature Guide for Security Devices*.
- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 4400.

Overview

This example shows how to control OSPF designated router election. Within the example, you set the OSPF interface to **ge-0/0/1** and the device priority to 200. The higher the priority value, the greater likelihood the routing device will become the designated router.

By default, routing devices have a priority of 128. A priority of 0 marks the routing device as ineligible to become the designated router. A priority of 1 means the routing device has the least chance of becoming a designated router.

Configuration

CLI Quick Configuration To quickly configure an OSPF designated router election, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

```
[edit]
set protocols ospf area 0.0.0.3 interface ge-0/0/1 priority 200
```

Step-by-Step Procedure To control OSPF designated router election:

1. Configure an OSPF interface and specify the device priority.



NOTE: To specify an OSPFv3 interface, include the **ospf3** statement at the [edit protocols] hierarchy level.

```
[edit]
user@host# set protocols ospf area 0.0.0.3 interface ge-0/0/1 priority 200
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Results

Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.3 {
  interface ge-0/0/1.0 {
    priority 200;
  }
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

Verification

Confirm that the configuration is working properly.

- [Verifying the Designated Router Election on page 4403](#)

Verifying the Designated Router Election

Purpose Based on the priority you configured for a specific OSPF interface, you can confirm the address of the area's designated router. The DR ID, DR, or DR-ID field displays the address of the area's designated router. The BDR ID, BDR, or BDR-ID field displays the address of the backup designated router.

Action From operational mode, enter the **show ospf interface** and the **show ospf neighbor** commands for OSPFv2, and enter the **show ospf3 interface** and the **show ospf3 neighbor** commands for OSPFv3.

Related Documentation

- [Understanding OSPF Areas on page 4397](#)
- [Understanding OSPF Configurations](#)

Examples: Configuring OSPF Areas


- [Understanding OSPF Areas and Backbone Areas on page 4404](#)
- [Example: Configuring a Single-Area OSPF Network on page 4406](#)
- [Example: Configuring a Multiarea OSPF Network on page 4408](#)

Understanding OSPF Areas and Backbone Areas

OSPF networks in an autonomous system (AS) are administratively grouped into *areas*. Each area within an AS operates like an independent network and has a unique 32-bit area ID, which functions similar to a network address. Within an area, the topology database contains only information about the area, link-state advertisements (LSAs) are flooded only to nodes within the area, and routes are computed only within the area. The topology of an area is hidden from the rest of the AS, thus significantly reducing routing traffic in the AS. Subnetworks are divided into other areas, which are connected to form the whole of the main network. Routing devices that are wholly within an area are called *internal routers*. All interfaces on internal routers are directly connected to networks within the area.

The central area of an AS, called the *backbone area*, has a special function and is always assigned the area ID 0.0.0.0. (Within a simple, single-area network, this is also the ID of the area.) Area IDs are unique numeric identifiers, in dotted decimal notation, but they are not IP addresses. Area IDs need only be unique within an AS. All other networks or areas in the AS must be directly connected to the backbone area by a routing device that has interfaces in more than one area. These connecting routing devices are called *area border routers* (ABRs). [Figure 109](#) shows an OSPF topology of three areas connected by two ABRs.

Figure 109: Multiarea OSPF Topology




Because all areas are adjacent to the backbone area, OSPF routers send all traffic not destined for their own area through the backbone area. The ABRs in the backbone area are then responsible for transmitting the traffic through the appropriate ABR to the destination area. The ABRs summarize the link-state records of each area and advertise destination address summaries to neighboring areas. The advertisements contain the ID of the area in which each destination lies, so that packets are routed to the appropriate ABR. For example, in the OSPF areas shown in [Figure 109](#), packets sent from Router A to Router C are automatically routed through ABR B.

Junos OS supports active backbone detection. Active backbone detection is implemented to verify that ABRs are connected to the backbone. If the connection to the backbone area is lost, then the routing device's default metric is not advertised, effectively rerouting traffic through another ABR with a valid connection to the backbone. Active backbone detection enables transit through an ABR with no active backbone connection. An ABR advertises to other routing devices that it is an ABR even if the connection to the backbone is down, so that the neighbors can consider it for interarea routes.

An OSPF restriction requires all areas to be directly connected to the backbone area so that packets can be properly routed. All packets are routed first to the backbone area by default. Packets that are destined for an area other than the backbone area are then routed to the appropriate ABR and on to the remote host within the destination area.

In large networks with many areas, in which direct connectivity between all areas and the backbone area is physically difficult or impossible, you can configure virtual links to connect noncontiguous areas. Virtual links use a transit area that contains two or more ABRs to pass network traffic from one adjacent area to another. For example, [Figure 110](#) shows a virtual link between a noncontiguous area and the backbone area through an area connected to both.

Figure 110: OSPF Topology with a Virtual Link



In the topology shown in [Figure 110](#), a virtual link is established between area 0.0.0.3 and the backbone area through area 0.0.0.2. All outbound traffic destined for other areas is routed through area 0.0.0.2 to the backbone area and then to the appropriate ABR. All inbound traffic destined for area 0.0.0.3 is routed to the backbone area and then through area 0.0.0.2.

Example: Configuring a Single-Area OSPF Network

This example shows how to configure a single-area OSPF network.

- [Requirements on page 4406](#)
- [Overview on page 4406](#)
- [Configuration on page 4407](#)
- [Verification on page 4408](#)

Requirements

Before you begin:

- Configure the device interfaces. See the *Interfaces Feature Guide for Security Devices*.
- Configure the router identifiers for the devices in your OSPF network. See [“Example: Configuring an OSPF Router Identifier” on page 4400](#).

Overview


To activate OSPF on a network, you must enable the OSPF protocol on all interfaces within the network on which OSPF traffic is to travel. To enable OSPF, you must configure one or more interfaces on the device within an OSPF area. Once the interfaces are configured, OSPF LSAs are transmitted on all OSPF-enabled interfaces, and the network topology is shared throughout the network.

In an autonomous system (AS), the backbone area is always assigned area ID 0.0.0.0 (within a simple, single-area network, this is also the ID of the area). Area IDs are unique numeric identifiers, in dotted decimal notation. Area IDs need only be unique within an AS. All other networks or areas in the AS must be directly connected to the backbone area by area border routers that have interfaces in more than one area. You must also create a backbone area if your network consists of multiple areas. In this example, you

create the backbone area and add interfaces, such as **ge-0/0/0**, as needed to the OSPF area.

To use OSPF on the device, you must configure at least one OSPF area, such as the one shown in [Figure 111](#).

Figure 111: Typical Single-Area OSPF Network Topology



Configuration

CLI Quick Configuration

To quickly configure a single-area OSPF network, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

```
[edit]
set protocols ospf area 0.0.0.0 interface ge-0/0/0
```

Step-by-Step Procedure

To configure a single-area OSPF network:

1. Configure the single-area OSPF network by specifying the area ID and associated interface.



NOTE: For a single-area OSPFv3 network, include the **ospf3** statement at the [edit protocols] hierarchy level.

```
[edit]
user@host# set protocols ospf area 0.0.0.0 interface ge-0/0/0
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Results

Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.0 {
  interface ge-0/0/0.0;
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

Verification

Confirm that the configuration is working properly.

- [Verifying the Interfaces in the Area on page 4408](#)

Verifying the Interfaces in the Area

Purpose Verify that the interface for OSPF or OSPFv3 has been configured for the appropriate area. Confirm that the Area field displays the value that you configured.

Action From operational mode, enter the **show ospf interface** command for OSPFv2, and enter the **show ospf3 interface** command for OSPFv3.

Example: Configuring a Multiarea OSPF Network

This example shows how to configure a multiarea OSPF network. To reduce traffic and topology maintenance for the devices in an OSPF autonomous system (AS), you can group the OSPF-enabled routing devices into multiple areas.

- [Requirements on page 4408](#)
- [Overview on page 4409](#)
- [Configuration on page 4409](#)
- [Verification on page 4411](#)

Requirements

Before you begin:

- Configure the device interfaces. See the *Interfaces Feature Guide for Security Devices*.
- Configure the router identifiers for the devices in your OSPF network. See "[Example: Configuring an OSPF Router Identifier](#)" on page 4400.
- Control OSPF designated router election. See "[Example: Controlling OSPF Designated Router Election](#)" on page 4402
- Configure a single-area OSPF network. See "[Example: Configuring a Single-Area OSPF Network](#)" on page 4406.


Overview

To activate OSPF on a network, you must enable the OSPF protocol on all interfaces within the network on which OSPF traffic is to travel. To enable OSPF, you must configure one or more interfaces on the device within an OSPF area. Once the interfaces are configured, OSPF LSAs are transmitted on all OSPF-enabled interfaces, and the network topology is shared throughout the network.

Each OSPF area consists of routing devices configured with the same area number. In [Figure 112](#), Router B resides in the backbone area of the AS. The backbone area is always assigned area ID 0.0.0.0. (All area IDs must be unique within an AS.) All other networks or areas in the AS must be directly connected to the backbone area by a router that has interfaces in more than one area. In this example, these area border routers are A, C, D, and E. You create an additional area (area 2) and assign it unique area ID 0.0.0.2, and then add interface `ge-0/0/0` to the OSPF area.

To reduce traffic and topology maintenance for the devices in an OSPF AS, you can group them into multiple areas as shown in [Figure 112](#). In this example, you create the backbone area, create an additional area (area 2) and assign it unique area ID 0.0.0.2, and you configure Device B as the area border router, where interface `ge-0/0/0` participates in OSPF area 0 and interface `ge-0/0/2` participates in OSPF area 2.

Figure 112: Typical Multiarea OSPF Network Topology



Configuration

CLI Quick Configuration

To quickly configure a multiarea OSPF network, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the `[edit]` hierarchy level, and then enter **commit** from configuration mode.

```
Device A [edit]
set protocols ospf area 0.0.0.0 interface ge-0/0/0
set protocols ospf area 0.0.0.0 interface ge-0/0/1

Device C [edit]
set protocols ospf area 0.0.0.0 interface ge-0/0/0
```

Device B [edit]
set protocols ospf area 0.0.0.0 interface ge-0/0/0
set protocols ospf area 0.0.0.2 interface ge-0/0/2

Device D [edit]
set protocols ospf area 0.0.0.2 interface ge-0/0/0
set protocols ospf area 0.0.0.2 interface ge-0/0/2

Device E [edit]
set protocols ospf area 0.0.0.2 interface ge-0/0/2

Step-by-Step Procedure To configure a multiarea OSPF network:

1. Configure the backbone area.



NOTE: For an OSPFv3 network, include the `ospf3` statement at the [edit protocols] hierarchy level.

```
[edit]
user@A# set protocols ospf area 0.0.0.0 interface ge-0/0/0
user@A# set protocols ospf area 0.0.0.0 interface ge-0/0/1
```

```
[edit]
user@C# set protocols ospf area 0.0.0.0 interface ge-0/0/0
```

```
[edit]
user@B# set protocols ospf area 0.0.0.0 interface ge-0/0/0
```

2. Configure an additional area for your OSPF network.



NOTE: For a multiarea OSPFv3 network, include the `ospf3` statement at the [edit protocols] hierarchy level.

```
[edit]
user@host# set protocols ospf area 0.0.0.2 interface ge-0/0/0
user@D# set protocols ospf area 0.0.0.2 interface ge-0/0/2
```

```
[edit]
user@E# set protocols ospf area 0.0.0.2 interface ge-0/0/2
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Results

Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
```

```

area 0.0.0.0 {
    interface ge-0/0/0.0;
    interface ge-0/0/1.0;
}

user@C# show protocols ospf
area 0.0.0.0 {
    interface ge-0/0/0.0;
}

user@B# show protocols ospf
area 0.0.0.0 {
    interface ge-0/0/0.0;
}
area 0.0.0.2 {
    interface ge-0/0/2.0;
}

user@D# show protocols ospf
area 0.0.0.2 {
    interface ge-0/0/0.0;
    interface ge-0/0/2.0;
}

user@E# show protocols ospf
area 0.0.0.2 {
    interface ge-0/0/2.0;
}

```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

Verification

Confirm that the configuration is working properly.

- [Verifying the Interfaces in the Area on page 4411](#)

Verifying the Interfaces in the Area

Purpose	Verify that the interface for OSPF or OSPFv3 has been configured for the appropriate area. Confirm that the Area field displays the value that you configured.
Action	From operational mode, enter the show ospf interface command for OSPFv2, and enter the show ospf3 interface command for OSPFv3.
Related Documentation	<ul style="list-style-type: none"> • Understanding OSPF Areas on page 4397 • Understanding OSPF Configurations

Advanced OSPF Area Configuration

- Examples: Configuring OSPF Stub and Not-So-Stubby Areas on page 4413
- Example: Configuring OSPF Multiarea Adjacency on page 4423
- Example: Disabling OSPFv2 Compatibility with RFC 1583 on page 4428


Examples: Configuring OSPF Stub and Not-So-Stubby Areas

- Understanding OSPF Stub Areas, Totally Stubby Areas, and Not-So-Stubby Areas on page 4413
- Example: Configuring OSPF Stub and Totally Stubby Areas on page 4414
- Example: Configuring OSPF Not-So-Stubby Areas on page 4418

Understanding OSPF Stub Areas, Totally Stubby Areas, and Not-So-Stubby Areas

Figure 113 shows an autonomous system (AS) across which many external routes are advertised. If external routes make up a significant portion of a topology database, you can suppress the advertisements in areas that do not have links outside the network. By doing so, you can reduce the amount of memory the nodes use to maintain the topology database and free it for other uses.

Figure 113: OSPF AS Network with Stub Areas and NSSAs



To control the advertisement of external routes into an area, OSPF uses stub areas. By designating an area border router (ABR) interface to the area as a stub interface, you suppress external route advertisements through the ABR. Instead, the ABR advertises a default route (through itself) in place of the external routes and generates network summary (Type 3) link-state advertisements (LSAs). Packets destined for external routes are automatically sent to the ABR, which acts as a gateway for outbound traffic and routes the traffic appropriately.



NOTE: You must explicitly configure the ABR to generate a default route when attached to a stub or not-so-stubby-area (NSSA). To inject a default route with a specified metric value into the area, you must configure the `default-metric` option and specify a metric value.

For example, area 0.0.0.3 in [Figure 113](#) is not directly connected to the outside network. All outbound traffic is routed through the ABR to the backbone and then to the destination addresses. By designating area 0.0.0.3 as a stub area, you reduce the size of the topology database for that area by limiting the route entries to only those routes internal to the area.

A stub area that only allows routes internal to the area and restricts Type 3 LSAs from entering the stub area is often called a *totally stubby area*. You can convert area 0.0.0.3 to a totally stubby area by configuring the ABR to only advertise and allow the default route to enter into the area. External routes and destinations to other areas are no longer summarized or allowed into a totally stubby area.



NOTE: If you incorrectly configure a totally stubby area, you might encounter network connectivity issues. You should have advanced knowledge of OSPF and understand your network environment before configuring totally stubby areas.

Similar to area 0.0.0.3 in [Figure 113](#), area 0.0.0.4 has no external connections. However, area 0.0.0.4 has static customer routes that are not internal OSPF routes. You can limit the external route advertisements to the area and advertise the static customer routes by designating the area an NSSA. In an NSSA, the AS boundary router generates NSSA external (Type 7) LSAs and floods them into the NSSA, where they are contained. Type 7 LSAs allow an NSSA to support the presence of AS boundary routers and their corresponding external routing information. The ABR converts Type 7 LSAs into AS external (Type 5) LSAs and leaks them to the other areas, but external routes from other areas are not advertised within the NSSA.

Example: Configuring OSPF Stub and Totally Stubby Areas

This example shows how to configure an OSPF stub area and a totally stubby area to control the advertisement of external routes into an area.

- [Requirements on page 4415](#)
- [Overview on page 4415](#)

- [Configuration on page 4416](#)
- [Verification on page 4418](#)

Requirements

Before you begin:

- Configure the device interfaces. See the *Interfaces Feature Guide for Security Devices*.
- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 4400.
- Control OSPF designated router election. See “[Example: Controlling OSPF Designated Router Election](#)” on page 4402
- Configure a multiarea OSPF network. See “[Example: Configuring a Multiarea OSPF Network](#)” on page 4408.

Overview

The backbone area, which is 0 in [Figure 114](#), has a special function and is always assigned the area ID 0.0.0.0. Area IDs are unique numeric identifiers, in dotted decimal notation. Area IDs need only be unique within an autonomous system (AS). All other networks or areas (such as 3, 7, and 9) in the AS must be directly connected to the backbone area by area border routers (ABRs) that have interfaces in more than one area.

Stub areas are areas through which or into which OSPF does not flood AS external link-state advertisements (Type 5 LSAs). You might create stub areas when much of the topology database consists of AS external advertisements and you want to minimize the size of the topology databases on the internal routers in the stub area.

The following restrictions apply to stub areas:

- You cannot create a virtual link through a stub area.
- A stub area cannot contain an AS boundary router.
- You cannot configure the backbone as a stub area.
- You cannot configure an area as both a stub area and an not-so-stubby area (NSSA).

In this example, you configure each routing device in area 7 (area ID 0.0.0.7) as a stub router and some additional settings on the ABR:

- **stub**—Specifies that this area become a stub area and not be flooded with Type 5 LSAs. You must include the **stub** statement on all routing devices that are in area 7 because this area has no external connections.
- **default-metric**—Configures the ABR to generate a default route with a specified metric into the stub area. This default route enables packet forwarding from the stub area to external destinations. You configure this option only on the ABR. The ABR does not automatically generate a default route when attached to a stub. You must explicitly configure this option to generate a default route.

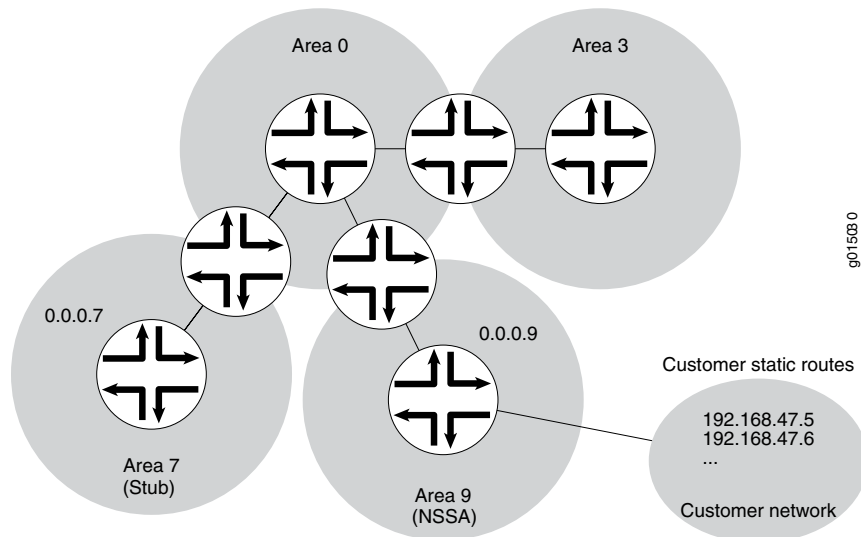
- **no-summaries**—(Optional) Prevents the ABR from advertising summary routes into the stub area by converting the stub area into a totally stubby area. If configured in combination with the **default-metric** statement, a totally stubby area only allows routes internal to the area and advertises the default route into the area. External routes and destinations to other areas are no longer summarized or allowed into a totally stubby area. Only the ABR requires this additional configuration because it is the only routing device within the totally stubby area that creates Type 3 LSAs used to receive and send traffic from outside of the area.

**NOTE:**

In Junos OS Release 8.5 and later, the following applies:

- A router-identifier interface that is not configured to run OSPF is no longer advertised as a stub network in OSPF LSAs.
- OSPF advertises a local route with a prefix length of 32 as a stub link if the loopback interface is configured with a prefix length other than 32. OSPF also advertises the direct route with the configured mask length, as in earlier releases.

Figure 114: OSPF Network Topology with Stub Areas and NSSAs



Configuration

CLI Quick Configuration

- To quickly configure an OSPF stub area, copy the following command and paste it into the CLI. You must configure all routing devices that are part of the stub area.

```
[edit]
set protocols ospf area 0.0.0.7 stub
```

- To quickly configure the ABR to inject a default route into the area, copy the following command and paste it into the CLI. You apply this configuration only on the ABR.

```
[edit]
```



```
set protocols ospf area 0.0.0.7 stub default-metric 10
```

- (Optional) To quickly configure the ABR to restrict all summary advertisements and allow only internal routes and default route advertisements into the area, copy the following command and paste it into the CLI. You apply this configuration only on the ABR.

```
[edit]
set protocols ospf area 0.0.0.7 stub no-summaries
```

Step-by-Step Procedure

To configure OSPF stub areas:

1. On all routing devices in the area, configure an OSPF stub area.



NOTE: To specify an OSPFv3 stub area, include the **ospf3** statement at the **[edit protocols]** hierarchy level.

```
[edit]
user@host# set protocols ospf area 0.0.0.7 stub
```

2. On the ABR, inject a default route into the area.

```
[edit]
user@host# set protocols ospf area 0.0.0.7 stub default-metric 10
```

3. (Optional) On the ABR, restrict summary LSAs from entering the area. This step converts the stub area into a totally stubby area.

```
[edit]
user@host# set protocols ospf area 0.0.0.7 stub no-summaries
```

4. If you are done configuring the devices, commit the configuration.

```
[edit]
user@host# commit
```

Results

Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Configuration on all routing devices:

```
user@host# show protocols ospf
area 0.0.0.7 {
  stub;
}
```

Configuration on the ABR (the output also includes the optional setting):

```
user@host# show protocols ospf
area 0.0.0.7 {
  stub default-metric 10 no-summaries;
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

Verification

Confirm that the configuration is working properly.

- [Verifying the Interfaces in the Area on page 4418](#)
- [Verifying the Type of OSPF Area on page 4418](#)

Verifying the Interfaces in the Area

Purpose Verify that the interface for OSPF has been configured for the appropriate area. Confirm that the output includes Stub as the type of OSPF area.

Action From operational mode, enter the **show ospf interface detail** command for OSPFv2, and enter the **show ospf3 interface detail** command for OSPFv3.

Verifying the Type of OSPF Area

Purpose Verify that the OSPF area is a stub area. Confirm that the output displays Normal Stub as the Stub type.

Action From operational mode, enter the **show ospf overview** command for OSPFv2, and enter the **show ospf3 overview** command for OSPFv3.

Example: Configuring OSPF Not-So-Stubby Areas

This example shows how to configure an OSPF not-so-stubby area (NSSA) to control the advertisement of external routes into an area.

- [Requirements on page 4418](#)
- [Overview on page 4419](#)
- [Configuration on page 4420](#)
- [Verification on page 4423](#)

Requirements

Before you begin:

- Configure the device interfaces. See the *Interfaces Feature Guide for Security Devices*.
- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 4400.
- Control OSPF designated router election. See “[Example: Controlling OSPF Designated Router Election](#)” on page 4402
- Configure a multiarea OSPF network. See “[Example: Configuring a Multiarea OSPF Network](#)” on page 4408.

Overview

The backbone area, which is 0 in [Figure 115](#), has a special function and is always assigned the area ID 0.0.0.0. Area IDs are unique numeric identifiers, in dotted decimal notation. Area IDs need only be unique within an AS. All other networks or areas (such as 3, 7, and 9) in the AS must be directly connected to the backbone area by ABRs that have interfaces in more than one area.

An OSPF stub area has no external routes, so you cannot redistribute routes from another protocol into a stub area. OSPF NSSAs allow external routes to be flooded within the area.

In addition, you might have a situation when exporting Type 7 LSAs into the NSSA is unnecessary. When an AS boundary router is also an ABR with an NSSA attached, Type 7 LSAs are exported into the NSSA by default. If the ABR is attached to multiple NSSAs, a separate Type 7 LSA is exported into each NSSA by default. During route redistribution, this routing device generates both Type 5 LSAs and Type 7 LSAs. You can disable exporting Type 7 LSAs into the NSSA.



NOTE: The following restriction applies to NSSAs: You cannot configure an area as both a stub area and an NSSA.

You configure each routing device in area 9 (area ID 0.0.0.9) with the following setting:

- **nssa**—Specifies an OSPF NSSA. You must include the **nssa** statement on all routing devices in area 9 because this area only has external connections to static routes.

You also configure the ABR in area 9 with the following additional settings:


- **no-summaries**—Prevents the ABR from advertising summary routes into the NSSA. If configured in combination with the **default-metric** statement, the NSSA only allows routes internal to the area and advertises the default route into the area. External routes and destinations to other areas are no longer summarized or allowed into the NSSA. Only the ABR requires this additional configuration because it is the only routing device within the NSSA that creates Type 3 LSAs used to receive and send traffic from outside the area.
- **default-lsa**—Configures the ABR to generate a default route into the NSSA. In this example, you configure the following:
 - **default-metric**—Specifies that the ABR generate a default route with a specified metric into the NSSA. This default route enables packet forwarding from the NSSA to external destinations. You configure this option only on the ABR. The ABR does not automatically generate a default route when attached to an NSSA. You must explicitly configure this option for the ABR to generate a default route.
 - **metric-type**—(Optional) Specifies the external metric type for the default LSA, which can be either Type 1 or Type 2. When OSPF exports route information from external ASs, it includes a cost, or external metric, in the route. The difference between the two metrics is how OSPF calculates the cost of the route. Type 1 external metrics

are equivalent to the link-state metric, where the cost is equal to the sum of the internal costs plus the external cost. Type 2 external metrics use only the external cost assigned by the AS boundary router. By default, OSPF uses the Type 2 external metric.

- **type-7**—(Optional) Floods Type 7 default LSAs into the NSSA if the **no-summaries** statement is configured. By default, when the **no-summaries** statement is configured, a Type 3 LSA is injected into NSSAs for Junos OS release 5.0 and later. To support backward compatibility with earlier Junos OS releases, include the **type-7** statement.

The second example also shows the optional configuration required to disable exporting Type 7 LSAs into the NSSA by including the **no-nssa-abr** statement on the routing device that performs the functions of both an ABR and an AS boundary router.

Figure 115: OSPF Network Topology with Stub Areas and NSSAs



Configuration

- [Configuring Routing Devices to Participate in a Not-So-Stubby-Area on page 4420](#)
- [Disabling the Export of Type 7 Link State Advertisements into Not-So-Stubby Areas on page 4422](#)

Configuring Routing Devices to Participate in a Not-So-Stubby-Area

CLI Quick Configuration

To quickly configure an OSPF NSSA, copy the following command and paste it into the CLI. You must configure all routing devices that are part of the NSSA.

```
[edit]
set protocols ospf area 0.0.0.9 nssa
```

To quickly configure an ABR that participates in an OSPF NSSA, copy the following commands and paste them into the CLI.

```
[edit]
set protocols ospf area 0.0.0.9 nssa default-lsa default-metric 10
set protocols ospf area 0.0.0.9 nssa default-lsa metric-type 1
```

```
set protocols ospf area 0.0.0.9 nssa default-lsa type-7
set protocols ospf area 0.0.0.9 nssa no-summaries
```

Step-by-Step Procedure

To configure OSPF NSSAs:

1. On all routing devices in the area, configure an OSPF NSSA.



NOTE: To specify an OSPFv3 NSSA area, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

[edit]

```
user@host# set protocols ospf area 0.0.0.9 nssa
```

2. On the ABR, enter OSPF configuration mode and specify the NSSA area 0.0.0.9 that you already created.

[edit]

```
user@host# edit protocols ospf area 0.0.0.9 nssa
```

3. On the ABR, inject a default route into the area.

[edit protocols ospf area 0.0.0.9 nssa]

```
user@host# set default-lsa default-metric 10
```

4. (Optional) On the ABR, specify the external metric type for the default route.

[edit protocols ospf area 0.0.0.9 nssa]

```
user@host# set default-lsa metric-type 1
```

5. (Optional) On the ABR, specify the flooding of Type 7 LSAs.

[edit protocols ospf area 0.0.0.9 nssa]

```
user@host# set default-lsa type-7
```

6. On the ABR, restrict summary LSAs from entering the area.

[edit protocols ospf area 0.0.0.9 nssa]

```
user@host# set no-summaries
```

7. If you are done configuring the devices, commit the configuration.

[edit protocols ospf area 0.0.0.9 nssa]

```
user@host# commit
```

Results

Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Configuration on all routing devices in the area:

```
user@host# show protocols ospf
area 0.0.0.9 {
  nssa;
}
```

Configuration on the ABR. The output also includes the optional `metric-type` and `type-7` statements.

```
user@host# show protocols ospf
  area 0.0.0.9 {
    nssa {
      default-lsa {
        default-metric 10;
        metric-type 1;
        type-7;
      }
      no-summaries;
    }
  }
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

Disabling the Export of Type 7 Link State Advertisements into Not-So-Stubby Areas

CLI Quick Configuration To quickly disable exporting Type 7 LSAs into the NSSA, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode. You configure this setting on an AS boundary router that is also an ABR with an NSSA area attached.

```
[edit]
set protocols ospf no-nssa-abr
```

Step-by-Step Procedure You can configure this setting if you have an AS boundary router that is also an ABR with an NSSA area attached.

1. Disable exporting Type 7 LSAs into the NSSA.



NOTE: To specify OSPFv3, include the **ospf3** statement at the [edit protocols] hierarchy level.

```
[edit]
user@host# set protocols ospf no-nssa-abr
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Results Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
  no-nssa-abr;
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

Verification

Confirm that the configuration is working properly.

- [Verifying the Interfaces in the Area on page 4423](#)
- [Verifying the Type of OSPF Area on page 4423](#)
- [Verifying the Type of LSAs on page 4423](#)

Verifying the Interfaces in the Area

Purpose Verify that the interface for OSPF has been configured for the appropriate area. Confirm that the output includes Stub NSSA as the type of OSPF area.

Action From operational mode, enter the **show ospf interface detail** command for OSPFv2, and enter the **show ospf3 interface detail** command for OSPFv3.

Verifying the Type of OSPF Area

Purpose Verify that the OSPF area is a stub area. Confirm that the output displays Not so Stubby Stub as the Stub type.

Action From operational mode, enter the **show ospf overview** command for OSPFv2, and enter the **show ospf3 overview** command for OSPFv3.

Verifying the Type of LSAs

Purpose Verify the type of LSAs that are in the area. If you disabled exporting Type 7 LSAs into an NSSA, confirm that the Type field does not include NSSA as a type of LSA.

Action From operational mode, enter the **show ospf overview** command for OSPFv2, and enter the **show ospf3 overview** command for OSPFv3.

Related Documentation

- *Example: Configuring OSPFv3 Stub and Totally Stubby Areas*
- [Understanding OSPF Areas on page 4397](#)
- *Understanding OSPF Configurations*

Example: Configuring OSPF Multiarea Adjacency

- [Multiarea Adjacency for OSPF on page 4423](#)
- [Example: Configuring Multiarea Adjacency for OSPF on page 4424](#)

Multiarea Adjacency for OSPF

By default, a single interface can belong to only one OSPF area. However, in some situations, you might want to configure an interface to belong to more than one area. Doing so allows the corresponding link to be considered an intra-area link in multiple areas and to be preferred over other higher-cost intra-area paths. For example, you can

configure an interface to belong to multiple areas with a high-speed backbone link between two area border routers (ABRs) so you can create multiarea adjacencies that belong to different areas.

In Junos OS Release 9.2 and later, you can configure a logical interface to belong to more than one OSPFv2 area. Support for OSPFv3 was introduced in Junos OS Release 9.4. As defined in RFC 5185, *OSPF Multi-Area Adjacency*, the ABRs establish multiple adjacencies belonging to different areas over the same logical interface. Each multiarea adjacency is announced as a point-to-point unnumbered link in the configured area by the routers connected to the link. For each area, one of the logical interfaces is treated as primary, and the remaining interfaces that are configured for the area are designated as secondary.

Any logical interface not configured as a secondary interface for an area is treated as the primary interface for that area. A logical interface can be configured as primary interface only for one area. For any other area for which you configure the interface, you must configure it as a secondary interface.

Example: Configuring Multiarea Adjacency for OSPF

This example shows how to configure multiarea adjacency for OSPF.

- [Requirements on page 4424](#)
- [Overview on page 4424](#)
- [Configuration on page 4425](#)
- [Verification on page 4427](#)

Requirements

Before you begin, plan your multiarea OSPF network. See [“Example: Configuring a Multiarea OSPF Network” on page 4408](#).

Overview

By default, a single interface can belong to only one OSPF area. You can configure a single interface to belong in multiple OSPF areas. Doing so allows the corresponding link to be considered an intra-area link in multiple areas and to be preferred over other higher-cost intra-area paths. When configuring a secondary interface, consider the following:

- For OSPFv2, you cannot configure point-to-multipoint and nonbroadcast multiaccess (NBMA) network interfaces as a secondary interface because secondary interfaces are treated as a point-to-point unnumbered link.
- Secondary interfaces are supported for LAN interfaces (the primary interface can be a LAN interface, but any secondary interfaces are treated as point-to-point unnumbered links over the LAN). In this scenario, you must ensure that there are only two routing devices on the LAN or that there are only two routing devices on the LAN that have secondary interfaces configured for a specific OSPF area.
- Since the purpose of a secondary interface is to advertise a topological path through an OSPF area, you cannot configure a secondary interface or a primary interface with one or more secondary interfaces to be passive. Passive interfaces advertise their

address, but do not run the OSPF protocol (adjacencies are not formed and hello packets are not generated).

- Any logical interface not configured as a secondary interface for an area is treated as a primary interface for that area. A logical interface can be configured as the primary interface only for one area. For any other area for which you configure the interface, you must configure it as a secondary interface.
- You cannot configure the **secondary** statement with the **interface all** statement.
- You cannot configure a secondary interface by its IP address.

In this example, you configure an interface to be in two areas, creating a multiarea adjacency with a link between two ABRs: ABR R1 and ABR R2. On each ABR, area 0.0.0.1 contains the primary interface and is the primary link between the ABRs, and area 0.0.0.2 contains the secondary logical interface, which you configure by including the **secondary** statement. You configure interface **so-0/0/0** on ABR R1 and interface **so-1/0/0** on ABR R2.

Configuration

CLI Quick Configuration

To quickly configure a secondary logical interface for an OSPF area, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

Configuration on ABR R1:

```
[edit]
set interfaces so-0/0/0 unit 0 family inet address 192.168.8.45/30
set routing-options router-id 10.255.0.1
set protocols ospf area 0.0.0.1 interface so-0/0/0
set protocols ospf area 0.0.0.2 interface so-0/0/0 secondary
```

Configuration on ABR R2:

```
[edit]
set interfaces so-1/0/0 unit 0 family inet address 192.168.8.37/30
set routing-options router-id 10.255.0.2
set protocols ospf area 0.0.0.1 interface so-1/0/0
set protocols ospf area 0.0.0.2 interface so-1/0/0 secondary
```

Step-by-Step Procedure

To configure a secondary logical interface:

1. Configure the device interfaces.



NOTE: For OSPFv3, on each interface specify the inet6 address family and include the IPv6 address.

```
[edit]
user@R1# set interfaces so-0/0/0 unit 0 family inet address 192.168.8.45/30
[edit]
```

```
user@R2# set interfaces so-1/0/0 unit 0 family inet address 192.168.8.37/30
```

2. Configure the router identifier.

```
[edit]
user@R1# set routing-options router-id 10.255.0.1
```

```
[edit]
user@R2# set routing-options router-id 10.255.0.2
```

3. On each ABR, configure the primary interface for the OSPF area.



NOTE: For OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@R1# set protocols ospf 0.0.0.1 interface so-0/0/0
```

```
[edit ]
user@R2# set protocols ospf 0.0.0.2 interface so-1/0/0
```

4. On each ABR, configure the secondary interface for the OSPF area.

```
[edit ]
user@R1# set protocols ospf area 0.0.0.1 so-0/0/0 secondary
```

```
[edit ]
user@R2# set protocols ospf area 0.0.0.2 so-1/0/0 secondary
```

5. If you are done configuring the devices, commit the configuration.

```
[edit protocols ospf area 0.0.0.1 ]
user@host# commit
```

Results

Confirm your configuration by entering the `show interfaces`, `show routing-options`, and the `show protocols ospf` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Configuration on ABR R1:

```
user@R1# show interfaces
so-0/0/0 {
  unit 0 {
    family inet {
      address 192.168.8.45/30;
    }
  }
}

user@R1# show routing-options
router-id 10.255.0.1;

user@R1# show protocols ospf
area 0.0.0.1 {
  interface so-0/0/0.0;
```

```

}
  area 0.0.0.2 {
    interface so-0/0/0.0 {
      secondary;
    }
  }
}

```

Configuration on ABR R2:

```

user@R2# show interfaces
so-0/0/0 {
  unit 0 {
    family inet {
      address 192.168.8.37/30;
    }
  }
}

user@R2# show routing-options
router-id 10.255.0.2;

user@R2# show protocols ospf
area 0.0.0.1 {
  interface so-1/0/0.0;
}
area 0.0.0.2 {
  interface so-1/0/0.0 {
    secondary;
  }
}

```

Verification

Confirm that the configuration is working properly.

- [Verifying the Secondary Interface on page 4427](#)
- [Verifying the Interfaces in the Area on page 4427](#)
- [Verifying Neighbor Adjacencies on page 4428](#)

Verifying the Secondary Interface

Purpose Verify that the secondary interface appears for the configured area. The Secondary field displays if the interface is configured as a secondary interface. The output might also show the same interface listed in multiple areas.

Action From operational mode, enter the **show ospf interface detail** command for OSPFv2, and enter the **show ospf3 interface detail** command for OSPFv3.

Verifying the Interfaces in the Area

Purpose Verify the interfaces configured for the specified area.

Action From operational mode, enter the **show ospf interface area *area-id*** command for OSPFv2, and enter the **show ospf3 interface area *area-id*** command for OSPFv3..

Verifying Neighbor Adjacencies

Purpose	Verify the primary and secondary neighbor adjacencies. The Secondary field displays if the neighbor is on a secondary interface.
Action	From operational mode, enter the show ospf neighbor detail command for OSPFv2, and enter the show ospf3 neighbor detail command for OSPFv3.
Related Documentation	<ul style="list-style-type: none">• Understanding OSPF Areas on page 4397• Understanding OSPF Areas and Backbone Areas on page 4404• Understanding OSPF Configurations

Example: Disabling OSPFv2 Compatibility with RFC 1583

- [OSPFv2 Compatibility with RFC 1583 Overview on page 4428](#)
- [Example: Disabling OSPFv2 Compatibility with RFC 1583 on page 4428](#)

OSPFv2 Compatibility with RFC 1583 Overview

By default, the Junos OS implementation of OSPFv2 is compatible with RFC 1583, *OSPF Version 2*. This means that Junos OS maintains a single best route to an autonomous system (AS) boundary router in the OSPF routing table, rather than multiple intra-AS paths, if they are available. You can now disable compatibility with RFC 1583. It is preferable to do so when the same external destination is advertised by AS boundary routers that belong to different OSPF areas. When you disable compatibility with RFC 1583, the OSPF routing table maintains the multiple intra-AS paths that are available, which the router uses to calculate AS external routes as defined in RFC 2328, *OSPF Version 2*. Being able to use multiple available paths to calculate an AS external route can prevent routing loops.

Example: Disabling OSPFv2 Compatibility with RFC 1583

This example shows how to disable OSPFv2 compatibility with RFC 1583 on the routing device.

- [Requirements on page 4428](#)
- [Overview on page 4429](#)
- [Configuration on page 4429](#)
- [Verification on page 4429](#)

Requirements

No special configuration beyond device initialization is required before disabling OSPFv2 compatibility with RFC 1583.

Overview

The introduction of RFC 2328 changed the method used to calculate the routes in an OSPF network. By default, the Junos OS implementation of OSPFv2 is compatible with RFC 1583, so OSPF uses the minimum cost to determine the route to any of the networks within the specified range. When you disable RFC 1583 compatibility, OSPF uses the maximum cost to determine the route to any of the networks within the specified range. To minimize the potential for routing loops, configure the same RFC compatibility on all OSPF devices in an OSPF domain.

Configuration

CLI Quick Configuration

To quickly disable OSPFv2 compatibility with RFC 1583, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode. You configure this setting on all devices that are part of the OSPF domain.

```
[edit]
set protocols ospf no-rfc-1583
```

Step-by-Step Procedure

To disable OSPFv2 compatibility with RFC 1583:

1. Disable RFC 1583.

```
[edit]
user@host# set protocols ospf no-rfc-1583
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```



NOTE: Repeat this configuration on each routing device that participates in an OSPF routing domain.

Results

Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
no-rfc-1583;
```

Verification

Confirm that the configuration is working properly.

Verifying the OSPF Routes

Purpose Verify that the OSPF routing table maintains the intra-AS paths with the largest metric, which the router uses to calculate AS external routes.

Action From operational mode, enter the **show ospf route detail** command.

Related Documentation

- [OSPF Overview on page 4386](#)
- *Understanding OSPF Configurations*

OSPF Interface Configuration

- [Examples: Configuring OSPF Interfaces on page 4431](#)
- [Example: Configuring Multiple Address Families for OSPFv3 on page 4446](#)

Examples: Configuring OSPF Interfaces

- [About OSPF Interfaces on page 4431](#)
- [Example: Configuring an Interface on a Broadcast or Point-to-Point Network on page 4432](#)
- [Example: Configuring an OSPFv2 Interface on a Nonbroadcast Multiaccess Network on page 4435](#)
- [Example: Configuring an OSPFv2 Interface on a Point-to-Multipoint Network on page 4438](#)
- [Example: Configuring OSPF Demand Circuits on page 4439](#)
- [Example: Configuring a Passive OSPF Interface on page 4442](#)
- [Example: Configuring OSPFv2 Peer interfaces on page 4444](#)

About OSPF Interfaces

To activate OSPF on a network, you must enable the OSPF protocol on one or more interfaces on each device within the network on which traffic is to travel. How you configure the interface depends on whether the interface is connected to a broadcast or point-to-point network, a point-to-multipoint network, a nonbroadcast multiaccess (NBMA) network, or across a demand circuit.

- A broadcast interface behaves as if the routing device is connected to a LAN.
- A point-to-point interface provides a connection between a single source and a single destination (there is only one OSPF adjacency).
- A point-to-multipoint interface provides a connection between a single source and multiple destinations.
- An NBMA interface behaves in a similar fashion to a point-to-multipoint interface, but you might configure an NBMA interface to interoperate with other equipment.
- A demand circuit is a connection on which you can limit traffic based on user agreements. The demand circuit can limit bandwidth or access time based on agreements between the provider and user.

You can also configure an OSPF interface to be passive, to operate in passive traffic engineering mode, or to be a peer interface.

- A passive interface advertises its address, but does not run the OSPF protocol (adjacencies are not formed and hello packets are not generated).
- An interface operating in OSPF passive traffic engineering mode floods link address information within the autonomous system (AS) and makes it available for traffic engineering calculations.
- A peer interface can be configured for OSPFv2 routing devices. A peer interface is required for Generalized MPLS (GMPLS) to transport traffic engineering information through a link separate from the control channel. You establish this separate link by configuring a peer interface. The peer interface name must match the Link Management Protocol (LMP) peer name. A peer interface is optional for a hierarchy of RSVP label-switched paths (LSPs). After you configure the forwarding adjacency, you can configure OSPFv2 to advertise the traffic engineering properties of a forwarding adjacency to a specific peer.

Point-to-point interfaces differ from multipoint in that only one OSPF adjacency is possible. (A LAN, for instance, can have multiple addresses and can run OSPF on each subnet simultaneously.) As such, when you configure a numbered point-to-point interface to OSPF by name, multiple OSPF interfaces are created. One, which is unnumbered, is the interface on which the protocol is run. An additional OSPF interface is created for each address configured on the interface, if any, which is automatically marked as passive.

For OSPFv3, one OSPF-specific interface must be created per interface name configured under OSPFv3. OSPFv3 does not allow interfaces to be configured by IP address.

Enabling OSPF on an interface (by including the **interface** statement), disabling it (by including the **disable** statement), and not actually having OSPF run on an interface (by including the **passive** statement) are mutually exclusive states.



NOTE: When you configure OSPFv2 on an interface, you must also include the **family inet** statement at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level. When you configure OSPFv3 on an interface, you must also include the **family inet6** statement at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level. In Junos OS Release 9.2 and later, you can configure OSPFv3 to support address families other than unicast IPv6.

Example: Configuring an Interface on a Broadcast or Point-to-Point Network

This example shows how to configure an OSPF interface on a broadcast or point-to-point network.

- [Requirements on page 4433](#)
- [Overview on page 4433](#)

- [Configuration on page 4433](#)
- [Verification on page 4435](#)

Requirements

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See [“Example: Configuring an OSPF Router Identifier” on page 4400](#).
- Control OSPF designated router election. See [“Example: Controlling OSPF Designated Router Election” on page 4402](#)
- Configure a multiarea OSPF network. See [“Example: Configuring a Multiarea OSPF Network” on page 4408](#).

Overview

If the interface on which you are configuring OSPF supports broadcast mode (such as a LAN), or if the interface supports point-to-point mode (such as a PPP interface or a point-to-point logical interface on Frame Relay), you specify the interface by including the IP address or the interface name for OSPFv2, or only the interface name for OSPFv3. In Junos OS Release 9.3 and later, an OSPF point-to-point interface can be an Ethernet interface without a subnet. If you configure an interface on a broadcast network, designated router and backup designated router election is performed.



NOTE: Using both the interface name and the IP address of the same interface produces an invalid configuration.

In this example, you configure interface **ge-0/2/0** as an OSPFv2 interface in OSPF area 0.0.0.1.

Configuration

CLI Quick Configuration

To quickly configure an OSPF interface on a broadcast or point-to-point network and to allow the inbound OSPF into the interfaces that are active, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

```
[edit]
set interfaces ge-0/2/0 unit 0 family inet address 10.0.0.1
set protocols ospf area 0.0.0.1 interface ge-0/2/0
set security zones security-zone Trust host-inbound-traffic protocols all
set security zones security-zone Trust host-inbound-traffic system-services all
set groups global security policies default-policy permit-all
set security zones security-zone Trust interfaces ge-0/2/0
```

Step-by-Step Procedure

To configure an OSPF interface on a broadcast or point-to-point network:

1. Configure the interface.



NOTE: For an OSPFv3 interface, specify an IPv6 address.

```
[edit]
user@host# set interfaces ge-0/2/0 unit 0 family inet address 10.0.0.1
```

2. Create an OSPF area.



NOTE: For an OSPFv3 interface, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# edit protocols ospf area 0.0.0.1
```

3. Assign the interface to the area.

```
[edit protocols ospf area 0.0.0.1 ]
user@host# set interface ge-0/2/0
```

4. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.1 ]
```

5. To allow the inbound OSPF into the interfaces that are active.

```
[edit]
user@host# set security zones security-zone Trust host-inbound-traffic protocols
all
user@host# set security zones security-zone Trust host-inbound-traffic
system-services all
user@host# set groups global security policies default-policy permit-all
user@host# set security zones security-zone Trust interfaces ge-0/2/0
user@host# commit
```

Results

Confirm your configuration by entering the `show interfaces` and the `show protocols ospf` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
ge-0/2/0 {
  unit 0 {
    family inet {
      address 10.0.0.1/32;
    }
  }
}

user@host# show protocols ospf
area 0.0.0.1 {
  interface ge-0/2/0.0;
}
```

To confirm your OSPFv3 configuration, enter the **show interfaces** and the **show protocols ospf3** commands.

Verification

Confirm that the configuration is working properly.

Verifying the OSPF Interface

- | | |
|----------------|---|
| Purpose | Verify the interface configuration. Depending on your deployment, the Type field might display LAN or P2P. |
| Action | From operational mode, enter the show ospf interface detail command for OSPFv2, and enter the show ospf3 interface detail command for OSPFv3. |

Example: Configuring an OSPFv2 Interface on a Nonbroadcast Multiaccess Network

This example shows how to configure an OSPFv2 interface on a nonbroadcast multiaccess (NBMA) network.

- [Requirements on page 4435](#)
- [Overview on page 4435](#)
- [Configuration on page 4436](#)
- [Verification on page 4437](#)

Requirements

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 4400.
- Control OSPF designated router election. See “[Example: Controlling OSPF Designated Router Election](#)” on page 4402.
- Configure a multiarea OSPF network. See “[Example: Configuring a Multiarea OSPF Network](#)” on page 4408.

Overview

When you configure OSPFv2 on an NBMA network, you can use nonbroadcast mode rather than point-to-multipoint mode. Using this mode offers no advantages over point-to-multipoint mode, but it has more disadvantages than point-to-multipoint mode. Nevertheless, you might occasionally find it necessary to configure nonbroadcast mode to interoperate with other equipment. Because there is no autodiscovery mechanism, you must configure each neighbor.

Nonbroadcast mode treats the NBMA network as a partially connected LAN, electing designated and backup designated routers. All routing devices must have a direct connection to both the designated and backup designated routers, or unpredictable results occur.

When you configure the interface, specify either the IP address or the interface name. Using both the IP address and the interface name produces an invalid configuration. For nonbroadcast interfaces, specify the IP address of the nonbroadcast interface as the interface name.

In this example, you configure the Asynchronous Transfer Mode (ATM) interface **at-0/1/0** as an OSPFv2 interface in OSPF area 0.0.0.1, and you specify the following settings:

- **interface-type nbma**—Sets the interface to run in NBMA mode. You must explicitly configure the interface to run in NBMA mode.
- **neighbor address <eligible>**—Specifies the IP address of the neighboring device. OSPF routing devices normally discover their neighbors dynamically by listening to the broadcast or multicast hello packets on the network. Because an NBMA network does not support broadcast (or multicast), the device cannot discover its neighbors dynamically, so you must configure all the neighbors statically. To configure multiple neighbors, include multiple **neighbor** statements. If you want the neighbor to be a designated router, include the **eligible** keyword.
- **poll-interval**—Specifies the length of time, in seconds, before the routing device sends hello packets out of the interface before it establishes adjacency with a neighbor. Routing devices send hello packets for a longer interval on nonbroadcast networks to minimize the bandwidth required on slow WAN links. The range is from 1 through 255 seconds. By default, the device sends hello packets out the interface every 120 seconds before it establishes adjacency with a neighbor.

Once the routing device detects an active neighbor, the hello packet interval changes from the time specified in the **poll-interval** statement to the time specified in the **hello-interval** statement.

Configuration

CLI Quick Configuration

To quickly configure an OSPFv2 interface on an NBMA network, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

```
[edit]
set interfaces at-0/1/0 unit 0 family inet address 192.0.2.1
set protocols ospf area 0.0.0.1 interface at-0/1/0.0 interface-type nbma
set protocols ospf area 0.0.0.1 interface at-0/1/0.0 neighbor 192.0.2.2 eligible
set protocols ospf area 0.0.0.1 interface at-0/1/0.0 poll-interval 130
```

Step-by-Step Procedure

To configure an OSPFv2 interface on an NBMA network:

1. Configure the interface.

```
[edit]
user@host# set interfaces at-0/1/0 unit 0 family inet address 192.0.2.1
```

2. Create an OSPF area.

```
[edit]
user@host# edit protocols ospf area 0.0.0.1
```

3. Assign the interface to the area.
In this example, include the **eligible** keyword to allow the neighbor to be a designated router.


```
[edit protocols ospf area 0.0.0.1 ]
user@host# set interface at-0/1/0 interface-type nbma neighbor 192.0.2.2 eligible
```
4. Configure the poll interval.


```
[edit protocols ospf area 0.0.0.1 ]
user@host# set interface at-0/1/0 poll-interval 130
```
5. If you are done configuring the device, commit the configuration.


```
[edit protocols ospf area 0.0.0.1 ]
user@host# commit
```

Results

Confirm your configuration by entering the **show interfaces** and the **show protocols ospf** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
at-0/1/0 {
  unit 0 {
    family inet {
      address 192.0.2.1/32;
    }
  }
}

user@host# show protocols ospf
area 0.0.0.1 {
  interface at-0/1/0.0 {
    interface-type nbma;
    neighbor 192.0.2.2 eligible;
    poll-interval 130;
  }
}
```

Verification

Confirm that the configuration is working properly.

Verifying the OSPF Interface

Purpose Verify the interface configuration. Confirm that the Type field displays NBMA.

Action From operational mode, enter the **show ospf interface detail** command.

Example: Configuring an OSPFv2 Interface on a Point-to-Multipoint Network

This example shows how to configure an OSPFv2 interface on a point-to-multipoint network.

- [Requirements on page 4438](#)
- [Overview on page 4438](#)
- [Configuration on page 4438](#)
- [Verification on page 4439](#)

Requirements

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See [“Example: Configuring an OSPF Router Identifier” on page 4400](#).
- Control OSPF designated router election. See [“Example: Controlling OSPF Designated Router Election” on page 4402](#)
- Configure a multiarea OSPF network. See [“Example: Configuring a Multiarea OSPF Network” on page 4408](#).

Overview

When you configure OSPFv2 on a nonbroadcast multiaccess (NBMA) network, such as a multipoint Asynchronous Transfer Mode (ATM) or Frame Relay, OSPFv2 operates by default in point-to-multipoint mode. In this mode, OSPFv2 treats the network as a set of point-to-point links. Because there is no autodiscovery mechanism, you must configure each neighbor.

When you configure the interface, specify either the IP address or the interface name. Using both the IP address and the interface name produces an invalid configuration.

In this example, you configure ATM interface **at-0/1/0** as an OSPFv2 interface in OSPF area 0.0.0.1, and you specify 192.0.2.1 as the neighbor's IP address.

Configuration

CLI Quick Configuration

To quickly configure an OSPFv2 interface on a point-to-multipoint network, copy the following commands and paste them into the CLI.

```
[edit]
set interfaces at-0/1/0 unit 0 family inet address 192.0.2.2
set protocols ospf area 0.0.0.1 interface at-0/1/0 neighbor 192.0.2.1
```

Step-by-Step Procedure

To configure an OSPFv2 interface on a point-to-multipoint network:

1. Configure the interface.

```
[edit]
user@host# set interfaces at-0/1/0 unit 0 family inet address 192.0.2.2
```

2. Create an OSPF area.

```
[edit]
user@host# edit protocols ospf area 0.0.0.1
```

3. Assign the interface to the area and specify the neighbor.

```
[edit protocols ospf area 0.0.0.1]
user@host# set interface at-0/1/0 neighbor 192.0.2.1
```

To configure multiple neighbors, include a **neighbor** statement for each neighbor.

4. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.1]
user@host# commit
```

Results

Confirm your configuration by entering the **show interfaces** and the **show protocols ospf** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
at-0/1/0 {
  unit 0 {
    family inet {
      address 192.0.2.2/32;
    }
  }
}

user@host# show protocols ospf
area 0.0.0.1 {
  interface at-0/1/0.0 {
    neighbor 192.0.2.1;
  }
}
```

Verification

Confirm that the configuration is working properly.

Verifying the OSPF Interface

Purpose Verify the interface configuration. Confirm that the Type field displays P2MP.

Action From operational mode, enter the **show ospf interface detail** command.

Example: Configuring OSPF Demand Circuits

This example shows how to configure an OSPF demand circuit interface.

- [Requirements on page 4440](#)
- [Overview on page 4440](#)
- [Configuration on page 4441](#)
- [Verification on page 4442](#)

Requirements

Before you begin:

- Configure the device interfaces. See the *Interfaces Feature Guide for Security Devices*.



NOTE: If you are using OSPF demand circuits over an ISDN link, you must configure an ISDN interface and enable dial-on-demand routing.

- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 4400.
- Configure a single-area OSPF network. See “[Example: Configuring a Single-Area OSPF Network](#)” on page 4406.
- Configure a multiarea OSPF network. See “[Example: Configuring a Multiarea OSPF Network](#)” on page 4408.

Overview

OSPF sends periodic hello packets to establish and maintain neighbor adjacencies and uses link-state advertisements (LSAs) to make routing calculations and decisions. OSPF support for demand circuits is defined in RFC 1793, *Extending OSPF to Support Demand Circuits*, and suppresses the periodic hello packets and LSAs. A demand circuit is a connection on which you can limit traffic based on user agreements. The demand circuit can limit bandwidth or access time based on agreements between the provider and user.

You configure demand circuits on an OSPF interface. When the interface becomes a demand circuit, all hello packets and LSAs are suppressed as soon as OSPF synchronization is achieved. LSAs have a DoNotAge bit that stops the LSA from aging and prevents periodic updates from being sent. Hello packets and LSAs are sent and received on a demand-circuit interface only when there is a change in the network topology. This reduces the amount of traffic through the OSPF interface.

Consider the following when configuring OSPF demand circuits:

- Periodic hellos are only suppressed on point-to-point and point-to-multipoint interfaces. If you configure demand circuits on an OSPF broadcast network or on an OSPF nonbroadcast multiaccess (NBMA) network, periodic hello packets are still sent.
- Demand circuit support on an OSPF point-to-multipoint interface resembles that for point-to-point interfaces. If you configure a point-to-multipoint interface as a demand circuit, the device negotiates hello suppression separately on each interface that is part of the point-to-multipoint network.

This example assumes that you have a point-to-point connection between two devices using SONET/SDH interfaces. A demand-circuit interface automatically negotiates the demand-circuit connection with its OSPF neighbor. If the neighbor does not support demand circuits, then no demand circuit connection is established.

In this example, you configure OSPF interface **so-0/1/0** in OSPF area 0.0.0.1 as a demand circuit.

Configuration

CLI Quick Configuration To quickly configure an OSPF demand circuit interface, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

```
[edit]
set protocols ospf area 0.0.0.1 interface so-0/1/0 demand-circuit
```

Step-by-Step Procedure To configure an OSPF demand circuit interface on one neighboring interface:

1. Create an OSPF area.



NOTE: For OSPFv3, include the **ospf3** statement at the [edit protocols] hierarchy level.

```
[edit ]
user@host# edit protocols ospf area 0.0.0.1
```

2. Configure the neighboring interface as a demand circuit.

```
[edit protocols ospf area 0.0.0.1]
user@host# set interface so-0/1/0 demand-circuit
```

3. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.1]
user@host# commit
```



NOTE: Repeat this entire configuration on the other neighboring interface.

Results

Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols
ospf {
  area 0.0.0.1 {
    interface so-0/1/0.0 {
      demand-circuit;
    }
  }
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

Verification

Confirm that the configuration is working properly.

Verifying the Status of Neighboring Interfaces

- Purpose** Verify information about the neighboring interface. When the neighbor is configured for demand circuits, a DC flag displays.
- Action** From operational mode, enter the **show ospf neighbor detail** command for OSPFv2, and enter the **show ospf3 neighbor detail** command for OSPFv3.

Example: Configuring a Passive OSPF Interface

This example shows how to configure a passive OSPF interface. A passive OSPF interface advertises its address but does not run the OSPF protocol.

- [Requirements on page 4442](#)
- [Overview on page 4442](#)
- [Configuration on page 4443](#)
- [Verification on page 4444](#)

Requirements

Before you begin:

- Configure the device interfaces. See the *Interfaces Feature Guide for Security Devices*.
- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 4400.
- Configure a single-area OSPF network. See “[Example: Configuring a Single-Area OSPF Network](#)” on page 4406.
- Configure a multiarea OSPF network. See “[Example: Configuring a Multiarea OSPF Network](#)” on page 4408.

Overview

By default, OSPF must be configured on an interface for direct interface addresses to be advertised as interior routes. To advertise the direct interface addresses without actually running OSPF on that interface (adjacencies are not formed and hello packets are not generated), you configure that interface as a passive interface.

Enabling OSPF on an interface (by including the **interface** statement), disabling it (by including the **disable** statement), and not actually having OSPF run on an interface (by including the **passive** statement) are mutually exclusive states.



NOTE: If you do not want to see notifications for state changes in a passive OSPF interface, you can disable the OSPF traps for the interface by including the `no-interface-state-traps` statement. The `no-interface-state-traps` statement is supported only for OSPFv2.

In this example, you configure interface `ge-0/2/0` as a passive OSPF interface in area `0.0.0.1` by including the `passive` statement.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the `[edit]` hierarchy level, and then enter `commit` from configuration mode.

```
[edit]
set protocols ospf area 0.0.0.1 interface ge-0/2/0 passive
```

Step-by-Step Procedure

To configure a passive OSPF interface:

1. Create an OSPF area.



NOTE: For an OSPFv3 interface, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# edit protocols ospf area 0.0.0.1
```

2. Configure the passive interface.

```
[edit protocols ospf area 0.0.0.1 ]
user@host# set interface ge-0/2/0 passive
```

3. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.1]
user@host# commit
```

Results

Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
  area 0.0.0.1 {
    interface ge-0/2/0.0 {
      passive;
    }
  }
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

Verification

Confirm that the configuration is working properly.

Verifying the Status of OSPF Interfaces

- Purpose** Verify the status of the OSPF interface. If the interface is passive, the Adj count field is 0 because no adjacencies have been formed. Next to this field, you might also see the word Passive.
- Action** From operational mode, enter the **show ospf interface detail** command for OSPFv2, and enter the **show ospf3 interface detail** command for OSPFv3.

Example: Configuring OSPFv2 Peer interfaces

This example shows how to configure an OSPFv2 peer interface.

- [Requirements on page 4444](#)
- [Overview on page 4444](#)
- [Configuration on page 4445](#)
- [Verification on page 4445](#)

Requirements

Before you begin:

- Configure the device interfaces. See the *Interfaces Feature Guide for Security Devices*.
- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 4400.
- Configure a single-area OSPF network. See “[Example: Configuring a Single-Area OSPF Network](#)” on page 4406.
- Configure a multiarea OSPF network. See “[Example: Configuring a Multiarea OSPF Network](#)” on page 4408.
- Configure Generalized MPLS per your network requirements. .

Overview

You can configure an OSPFv2 peer interface for many reasons, including when you configure Generalized MPLS (GMPLS). This example configures a peer interface for GMPLS. GMPLS requires traffic engineering information to be transported through a link separate from the control channel. You establish this separate link by configuring a peer interface. The OSPFv2 peer interface name must match the Link Management Protocol (LMP) peer name. You configure GMPLS and the LMP settings separately from OSPF.

This example assumes that GMPLS and the LMP peer named **oxc1** are already configured, and you need to configure the OSPFv2 peer interface in area 0.0.0.0.

Configuration

CLI Quick Configuration To quickly configure an OSPFv2 peer interface, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

```
[edit]
set protocols ospf area 0.0.0.0 peer-interface oxc1
```

Step-by-Step Procedure To configure a peer OSPFv2 interface used by the LMP:

1. Create an OSPF area.

```
[edit]
user@host# edit protocols ospf area 0.0.0.0
```

2. Configure the peer interface.

```
[edit protocols ospf area 0.0.0.0]
user@host# set peer-interface oxc1
```

3. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.0]
user@host# commit
```

Results

Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
  area 0.0.0.0 {
    peer-interface oxc1;
  }
```

Verification

Confirm that the configuration is working properly.

Verifying the Configured OSPFv2 Peer

Purpose Verify the status of the OSPFv2 peer. When an OSPFv2 peer is configured for GMPLS, the Peer Name field displays the name of the LMP peer that you created for GMPLS, which is also the configured OSPFv2 peer.

Action From operational mode, enter the **show link-management** command.

Related Documentation

- [OSPF Overview on page 4386](#)
- [Understanding OSPF Configurations](#)

Example: Configuring Multiple Address Families for OSPFv3

- [Understanding Multiple Address Families for OSPFv3 on page 4446](#)
- [Example: Configuring Multiple Address Families for OSPFv3 on page 4446](#)

Understanding Multiple Address Families for OSPFv3

By default, OSPFv3 supports only unicast IPv6 routes. In Junos OS Release 9.2 and later, you can configure OSPFv3 to support multiple address families, including IPv4 unicast, IPv4 multicast, and IPv6 multicast. This multiple address family support allows OSPFv3 to support both IPv6 and IPv4 nodes. Junos OS maps each address family to a separate realm as defined in Internet draft draft-ietf-ospf-af-alt-06.txt, *Support for Address Families in OSPFv3*. Each realm maintains a separate set of neighbors and link-state database.

When you configure multiple address families for OSPFv3, there is a new instance ID field that allows multiple OSPFv3 protocol instances per link. This allows a single link to belong to multiple areas.

You configure each realm independently. We recommend that you configure an area and at least one interface for each realm.

These are the default import and export routing tables for each of the four address families:

- IPv6 unicast: **inet6.0**
- IPv6 multicast: **inet6.2**
- IPv4 unicast: **inet.0**
- IPv4 multicast: **inet.2**

With the exception of virtual links, all configurations supported for the default IPv6 unicast family are supported for the address families that have to be configured as realms.

Example: Configuring Multiple Address Families for OSPFv3

This example shows how to configure multiple address families for OSPFv3.

- [Requirements on page 4446](#)
- [Overview on page 4447](#)
- [Configuration on page 4448](#)
- [Verification on page 4449](#)

Requirements

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See [“Example: Configuring an OSPF Router Identifier” on page 4400](#).

- Configure a single-area OSPF network. See [“Example: Configuring a Single-Area OSPF Network” on page 4406](#).
- Configure a multiarea OSPF network. See [“Example: Configuring a Multiarea OSPF Network” on page 4408](#).

Overview

By default, OSPFv3 supports unicast IPv6 routes, but you can configure OSPFv3 to support multiple address families. To support an address family other than unicast IPv6, you configure a realm that allows OSPFv3 to advertise IPv4 unicast, IPv4 multicast, or IPv6 multicast routes. Junos OS then maps each address family that you configure to a separate realm with its own set of neighbors and link-state database.




NOTE: By default, LDP synchronization is only supported for OSPFv2. If you configure an IPv4 unicast or IPv4 multicast realm, you can also configure LDP synchronization. Since LDP synchronization is only supported for IPv4, this support is only available for OSPFv3 if you configure an IPv4 realm.

When configuring OSPFv3 to support multiple address families, consider the following:

- You configure each realm independently. We recommend that you configure an area and at least one interface for each realm.
- OSPFv3 uses IPv6 link-local addresses as the source of hello packets and next hop calculations. As such, you must enable IPv6 on the link regardless of the additional realm you configure.

[Figure 116](#) shows a connection between Routers R0 and R1. In this example, you configure interface **fe-0/1/0** on Router R0 in area 0 to advertise IPv4 unicast routes, in addition to the default unicast IPv6 routes in area 1, by including the **realm ipv4-unicast** statement. Depending on your network requirements, you can also advertise IPv4 multicast routes by including the **realm-ipv4-multicast** statement, and you can advertise IPv6 multicast routes by including the **realm-ipv6-multicast** statement.

Figure 116: IPv4 Unicast Realm



Configuration

CLI Quick Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Modifying the Junos OS Configuration* in the *CLI User Guide*.

To quickly configure multiple address families for OSPFv3, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

```
[edit]
set interfaces fe-0/1/0 unit 0 family inet address 11.1.2.1/24
set interfaces fe-0/1/0 unit 0 family inet6
set protocols ospf3 area 0.0.0.0 interface fe-0/1/0
set protocols ospf3 realm ipv4-unicast area 0.0.0.0 interface fe-0/1/0
```

Step-by-Step Procedure

To configure multiple address families for OSPFv3:

1. Configure the device interface participating in OSPFv3.


```
[edit]
user@host# set interfaces fe-0/1/0 unit 0 family inet address 11.1.2.1/24
user@host# set interfaces fe-0/1/0 unit 0 family inet6
```
2. Enter OSPFv3 configuration mode.


```
[edit ]
user@host# edit protocols ospf3
```
3. Add the interface you configured to the OSPFv3 area.


```
[edit protocols ospf3 ]
user@host# set area 0.0.0.0 interface fe-0/1/0
```
4. Configure an IPv4 unicast realm. This allows OSPFv3 to support both IPv4 unicast and IPv6 unicast routes.


```
[edit protocols ospf3 ]
user@host# set realm ipv4-unicast area 0.0.0.0 interface fe-0/1/0
```

5. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf3 ]
user@host# commit
```



NOTE: Repeat this entire configuration on the neighboring device that is part of the realm.

Results

Confirm your configuration by entering the **show interfaces** and the **show protocols ospf** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
fe-0/1/0 {
  unit 0 {
    family inet {
      address 11.1.2.1/24;
    }
    family inet6;
  }
}

user@host# show protocols ospf3
realm ipv4-unicast {
  area 0.0.0.0 {
    interface fe-0/1/0.0;
  }
}
area 0.0.0.0 {
  interface fe-0/1/0.0;
}
```

Verification

Confirm that the configuration is working properly.

- [Verifying the Link-State Database on page 4449](#)
- [Verifying the Status of OSPFv3 Interfaces with Multiple Address Families on page 4450](#)

Verifying the Link-State Database

Purpose Verify the status of the link-state database for the configured realm, or address family.

Action From operational mode, enter the **show ospf3 database realm ipv4-unicast** command.

Verifying the Status of OSPFv3 Interfaces with Multiple Address Families

Purpose Verify the status of the interface for the specified OSPFv3 realm, or address family.

Action From operational mode, enter the **show ospf3 interface realm ipv4-unicast** command.

Related Documentation

- [OSPF Overview on page 4386](#)
- *Understanding OSPF Configurations*

OSPF Route Control Configuration

- [Examples: Configuring OSPF Route Summarization on page 4451](#)
- [Examples: Configuring OSPF Traffic Control on page 4460](#)
- [Example: Configuring OSPF Overload Mode on page 4470](#)

Examples: Configuring OSPF Route Summarization

- [Understanding OSPF Route Summarization on page 4451](#)
- [Example: Summarizing Ranges of Routes in OSPF Link-State Advertisements on page 4452](#)
- [Example: Limiting the Number of Prefixes Exported to OSPF on page 4457](#)
- [Configuring OSPF Refresh and Flooding Reduction in Stable Topologies on page 4459](#)

Understanding OSPF Route Summarization

Area border routers (ABRs) send summary link advertisements to describe the routes to other areas. Depending on the number of destinations, an area can get flooded with a large number of link-state records, which can utilize routing device resources. To minimize the number of advertisements that are flooded into an area, you can configure the ABR to coalesce, or summarize, a range of IP addresses and send reachability information about these addresses in a single link-state advertisement (LSA). You can summarize one or more ranges of IP addresses, where all routes that match the specified area range are filtered at the area boundary, and the summary is advertised in their place.

For an OSPF area, you can summarize and filter intra-area prefixes. All routes that match the specified area range are filtered at the area boundary, and the summary is advertised in their place. For an OSPF not-so-stubby area (NSSA), you can only coalesce or filter NSSA external (Type 7) LSAs before they are translated into AS external (Type 5) LSAs and enter the backbone area. All external routes learned within the area that do not fall into the range of one of the prefixes are advertised individually to other areas.

In addition, you can also limit the number of prefixes (routes) that are exported into OSPF. By setting a user-defined maximum number of prefixes, you prevent the routing device from flooding an excessive number of routes into an area.

Example: Summarizing Ranges of Routes in OSPF Link-State Advertisements

This example shows how to summarize routes sent into the backbone area.

- [Requirements on page 4452](#)
- [Overview on page 4452](#)
- [Configuration on page 4453](#)
- [Verification on page 4457](#)

Requirements

Before you begin:

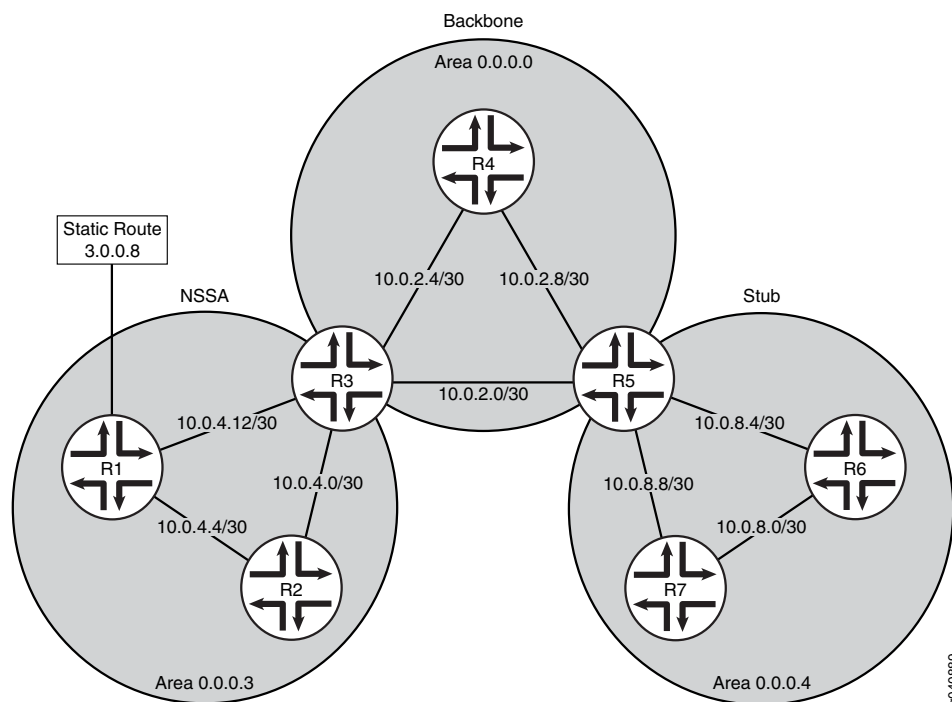
- Configure the router identifiers for the devices in your OSPF network. See [“Example: Configuring an OSPF Router Identifier” on page 4400](#).
- Control OSPF designated router election. See [“Example: Controlling OSPF Designated Router Election” on page 4402](#)
- Configure a static route. See *Examples: Configuring Static Routes* in the *Junos OS Routing Protocols Library for Routing Devices*.

Overview

You can summarize a range of IP addresses to minimize the size of the backbone router's link-state database. All routes that match the specified area range are filtered at the area boundary, and the summary is advertised in their place.

[Figure 117](#) shows the topology used in this example. R5 is the ABR between area 0.0.0.4 and the backbone. The networks in area 0.0.0.4 are 10.0.8.4/30, 10.0.8.0/30, and 10.0.8.8/30, which can be summarized as 10.0.8.0/28. R3 is the ABR between NSSA area 0.0.0.3 and the backbone. The networks in area 0.0.0.3 are 10.0.4.4/30, 10.0.4.0/30, and 10.0.4.12/30, which can be summarized as 10.0.4.0/28. Area 0.0.0.3 also contains external static route 3.0.0.8 that you will prevent from flooding throughout the network.

Figure 117: Summarizing Ranges of Routes in OSPF



In this example, you configure the ABRs for route summarization by including the following settings:

- **area-range**—For an area, summarizes a range of IP addresses when sending summary intra-area link advertisements. For an NSSA, summarizes a range of IP addresses when sending NSSA link-state advertisements (Type 7 LSAs). The specified prefixes are used to aggregate external routes learned within the area when the routes are advertised to other areas.
- **network/mask-length**—Indicates the summarized IP address range and the number of significant bits in the network mask.
- **restrict**—On the NSSA ABR, prevents the configured summary from being advertised. In this example, we do not want to flood the external route outside of area 0.0.0.3.

Configuration

CLI Quick Configuration

- To quickly configure route summarization for an OSPF area, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode. The following is the configuration on ABR R5:

```
[edit]
set interfaces fe-0/0/1 unit 0 family inet address 10.0.8.3/30
set interfaces fe-0/0/2 unit 0 family inet address 10.0.8.4/30
set interfaces fe-0/0/0 unit 0 family inet address 10.0.2.3/30
set interfaces fe-0/0/4 unit 0 family inet address 10.0.2.5/30
set protocols ospf area 0.0.0.4 stub
```

```

set protocols ospf area 0.0.0.4 interface fe-0/0/1
set protocols ospf area 0.0.0.4 interface fe-0/0/2
set protocols ospf area 0.0.0.0 interface fe-0/0/0
set protocols ospf area 0.0.0.0 interface fe-0/0/4
set protocols ospf area 0.0.0.4 area-range 10.0.8.0/28

```

- To quickly configure route summarization for an OSPF NSSA, copy the following commands and paste them into the CLI. The following is the configuration on ABR R3:

```

[edit]
set interfaces fe-0/0/1 unit 0 family inet address 10.0.4.10/30
set interfaces fe-0/0/2 unit 0 family inet address 10.0.4.1/30
set interfaces fe-0/0/0 unit 0 family inet address 10.0.2.1/30
set interfaces fe-0/0/4 unit 0 family inet address 10.0.2.7/30
set protocols ospf area 0.0.0.3 interface fe-0/0/1
set protocols ospf area 0.0.0.3 interface fe-0/0/2
set protocols ospf area 0.0.0.0 interface fe-0/0/0
set protocols ospf area 0.0.0.0 interface fe-0/0/4
set protocols ospf area 0.0.0.3 area-range 10.0.4.0/28
set protocols ospf area 0.0.0.3 nssa
set protocols ospf area 0.0.0.3 nssa area-range 3.0.0.0/8 restrict

```

Step-by-Step Procedure

To summarize routes sent to the backbone area:

1. Configure the interfaces.



NOTE: For OSPFv3, include IPv6 addresses.

```

[edit]
user@R5# set interfaces fe-0/0/1 unit 0 family inet address 10.0.8.3/30
user@R5# set interfaces fe-0/0/2 unit 0 family inet address 10.0.8.4/30
user@R5# set interfaces fe-0/0/0 unit 0 family inet address 10.0.2.3/30
user@R5# set interfaces fe-0/0/4 unit 0 family inet address 10.0.2.5/30

```

```

[edit]
user@R3# set interfaces fe-0/0/1 unit 0 family inet address 10.0.4.10/30
user@R3# set interfaces fe-0/0/2 unit 0 family inet address 10.0.4.1/30
user@R3# set interfaces fe-0/0/0 unit 0 family inet address 10.0.2.1/30
user@R3# set interfaces fe-0/0/4 unit 0 family inet address 10.0.2.7/30

```

2. Configure the type of OSPF area.



NOTE: For OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```

[edit]
user@R5# set protocols ospf area 0.0.0.4 stub

```

```

[edit]
user@R3# set protocols ospf area 0.0.0.3 nssa

```

3. Assign the interfaces to the OSPF areas.

```

user@R5# set protocols ospf area 0.0.0.4 interface fe-0/0/1
user@R5# set protocols ospf area 0.0.0.4 interface fe-0/0/2
user@R5# set protocols ospf area 0.0.0.0 interface fe-0/0/0
user@R5# set protocols ospf area 0.0.0.0 interface fe-0/0/4

user@R3# set protocols ospf area 0.0.0.3 interface fe-0/0/1
user@R3# set protocols ospf area 0.0.0.3 interface fe-0/0/2
user@R3# set protocols ospf area 0.0.0.0 interface fe-0/0/0
user@R3# set protocols ospf area 0.0.0.0 interface fe-0/0/4

```

4. Summarize the routes that are flooded into the backbone.

```

[edit]
user@R5# set protocols ospf area 0.0.0.4 area-range 10.0.8.0/28

[edit]
user@R3# set protocols ospf area 0.0.0.3 area-range 10.0.4.0/28

```

5. On ABR R3, restrict the external static route from leaving area 0.0.0.3.

```

[edit]
user@R3# set protocols ospf area 0.0.0.3 nssa area-range 3.0.0.0/8 restrict

```

6. If you are done configuring the devices, commit the configuration.

```

[edit]
user@host# commit

```

Results

Confirm your configuration by entering the **show interfaces** and the **show protocols ospf** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Configuration on ABR R5:

```

user@R5# show interfaces
fe-0/0/0 {
  unit 0 {
    family inet {
      address 10.0.2.3/32;
    }
  }
}
fe-0/0/1 {
  unit 0 {
    family inet {
      address 10.0.8.3/32;
    }
  }
}
fe-0/0/2 {
  unit 0 {
    family inet {
      address 10.0.8.4/32;
    }
  }
}

```

```
fe-0/0/4 {  
  unit 0 {  
    family inet {  
      address 10.0.2.5/32;  
    }  
  }  
}
```

```
user@R5# show protocols ospf  
area 0.0.0.0 {  
  interface fe-0/0/0.0;  
  interface fe-0/0/4.0;  
}  
area 0.0.0.4 {  
  stub;  
  area-range 10.0.8.0/28;  
  interface fe-0/0/1.0;  
  interface fe-0/0/2.0;  
}
```

Configuration on ABR R3:

```
user@R3# show interfaces  
fe-0/0/0 {  
  unit 0 {  
    family inet {  
      address 10.0.2.1/32;  
    }  
  }  
}  
fe-0/0/1 {  
  unit 0 {  
    family inet {  
      address 10.0.4.10/32;  
    }  
  }  
}  
fe-0/0/2 {  
  unit 0 {  
    family inet {  
      address 10.0.4.1/32;  
    }  
  }  
}  
fe-0/0/4 {  
  unit 0 {  
    family inet {  
      address 10.0.2.7/32;  
    }  
  }  
}  
  
user@R3t# show protocols ospf  
area 0.0.0.0 {  
  interface fe-0/0/0.0;  
  interface fe-0/0/4.0;  
}
```



```

area 0.0.0.3 {
  nssa {
    area-range 3.0.0.0/8 restrict;
  }
  area-range 10.0.4.0/28;
  interface fe-0/0/1.0;
  interface fe-0/0/2.0;
}

```

To confirm your OSPFv3 configuration, enter the **show interfaces** and **show protocols ospf3** commands.

Verification

Confirm that the configuration is working properly.

Verifying the Summarized Route

- | | |
|----------------|--|
| Purpose | Verify that the routes you configured for route summarization are being aggregated by the ABRs before the routes enter the backbone area. Confirm route summarization by checking the entries of the OSPF link-state database for the routing devices in the backbone. |
| Action | From operational mode, enter the show ospf database command for OSPFv2, and enter the show ospf3 database command for OSPFv3. |

Example: Limiting the Number of Prefixes Exported to OSPF

This example shows how to limit the number of prefixes exported to OSPF.

- [Requirements on page 4457](#)
- [Overview on page 4458](#)
- [Configuration on page 4458](#)
- [Verification on page 4459](#)

Requirements

Before you begin:

- Configure the device interfaces. See the *Interfaces Feature Guide for Security Devices*.
- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 4400.
- Control OSPF designated router election. See “[Example: Controlling OSPF Designated Router Election](#)” on page 4402
- Configure a single-area OSPF network. See “[Example: Configuring a Single-Area OSPF Network](#)” on page 4406.
- Configure a multiarea OSPF network. See “[Example: Configuring a Multiarea OSPF Network](#)” on page 4408.

Overview

By default, there is no limit to the number of prefixes (routes) that can be exported into OSPF. By allowing any number of routes to be exported into OSPF, the routing device can become overwhelmed and potentially flood an excessive number of routes into an area.

You can limit the number of routes exported into OSPF to minimize the load on the routing device and prevent this potential problem. If the routing device exceeds the configured prefix export value, the routing device purges the external prefixes and enters into an overload state. This state ensures that the routing device is not overwhelmed as it attempts to process routing information. The prefix export limit number can be a value from 0 through 4,294,967,295.

In this example, you configure a prefix export limit of 100,000 by including the **prefix-export-limit** statement.

Configuration

CLI Quick Configuration

To quickly limit the number of prefixes exported to OSPF, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

```
[edit]
set protocols ospf prefix-export-limit 100000
```

Step-by-Step Procedure

To limit the number of prefixes exported to OSPF:

1. Configure the prefix export limit value.



NOTE: For OSPFv3, include the **ospf3** statement at the [edit protocols] hierarchy level.

```
[edit]
user@host# set protocols ospf prefix-export-limit 100000
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Results

Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
prefix-export-limit 100000;
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

Verification

Confirm that the configuration is working properly.

Verifying the Prefix Export Limit

Purpose Verify the prefix export counter that displays the number of routes exported into OSPF.

Action From operational mode, enter the **show ospf overview** command for OSPFv2, and enter the **show ospf3 overview** command for OSPFv3.

Configuring OSPF Refresh and Flooding Reduction in Stable Topologies

The OSPF standard requires that every link-state advertisement (LSA) be refreshed every 30 minutes. The Juniper Networks implementation refreshes LSAs every 50 minutes. By default, any LSA that is not refreshed expires after 60 minutes. This requirement can result in traffic overhead that makes it difficult to scale OSPF networks. You can override the default behavior by specifying that the DoNotAge bit be set in self-originated LSAs when they are initially sent by the router or switch. Any LSA with the DoNotAge bit set is reflooded only when a change occurs in the LSA. This feature thus reduces protocol traffic overhead while permitting any changed LSAs to be flooded immediately. Routers or switches enabled for flood reduction continue to send hello packets to their neighbors and to age self-originated LSAs in their databases.

The Juniper implementation of OSPF refresh and flooding reduction is based on RFC 4136, *OSPF Refresh and Flooding Reduction in Stable Topologies*. However, the Juniper implementation does not include the forced-flooding interval defined in the RFC. Not implementing the forced-flooding interval ensures that LSAs with the DoNotAge bit set are reflooded only when a change occurs.

This feature is supported for the following:

- OSPFv2 and OSPFv3 interfaces
- OSPFv3 realms
- OSPFv2 and OSPFv3 virtual links
- OSPFv2 sham links
- OSPFv2 peer interfaces
- All routing instances supported by OSPF
- Logical systems

To configure flooding reduction for an OSPF interface, include the **flood-reduction** statement at the **[edit protocols (ospf | ospf3) area area-id interface interface-id]** hierarchy level.



NOTE: If you configure flooding reduction for an interface configured as a demand circuit, the LSAs are not initially flooded, but sent only when their content has changed. Hello packets and LSAs are sent and received on a demand-circuit interface only when a change occurs in the network topology.

In the following example, the OSPF interface **so-0/0/1.0** is configured for flooding reduction. As a result, all the LSAs generated by the routes that traverse the specified interface have the DoNotAge bit set when they are initially flooded, and LSAs are refreshed only when a change occurs.

```
[edit]
protocols ospf {
  area 0.0.0.0 {
    interface so-0/0/1.0 {
      flood-reduction;
    }
    interface lo0.0;
    interface so-0/0/0.0;
  }
}
```



NOTE: Beginning with Junos OS Release 12.2, you can configure a global default link-state advertisement (LSA) flooding interval in OSPF for self-generated LSAs by including the *lsa-refresh-interval minutes* statement at the `[edit protocols (ospf | ospf3)]` hierarchy level. The Juniper Networks implementation refreshes LSAs every 50 minutes. The range is 25 through 50 minutes. By default, any LSA that is not refreshed expires after 60 minutes.

If you have both the global LSA refresh interval configured for OSPF and OSPF flooding reduction configured for a specific interface in an OSPF area, the OSPF flood reduction configuration takes precedence for that specific interface.

- Related Documentation**
- [OSPF Overview on page 4386](#)
 - [Understanding OSPF Configurations](#)

Examples: Configuring OSPF Traffic Control

- [Understanding OSPF Traffic Control on page 4461](#)
- [Example: Controlling the Cost of Individual OSPF Network Segments on page 4462](#)
- [Example: Dynamically Adjusting OSPF Interface Metrics Based on Bandwidth on page 4466](#)
- [Example: Controlling OSPF Route Preferences on page 4468](#)

Understanding OSPF Traffic Control

Once a topology is shared across the network, OSPF uses the topology to route packets between network nodes. Each path between neighbors is assigned a cost based on the throughput, round-trip time, and reliability of the link. The sum of the costs across a particular path between hosts determines the overall cost of the path. Packets are then routed along the shortest path using the shortest-path-first (SPF) algorithm. If multiple equal-cost paths exist between a source and destination address, OSPF routes packets along each path alternately, in round-robin fashion. Routes with lower total path metrics are preferred over those with higher path metrics.

You can use the following methods to control OSPF traffic:

- Control the cost of individual OSPF network segments
- Dynamically adjust OSPF interface metrics based on bandwidth
- Control OSPF route selection

Controlling the Cost of Individual OSPF Network Segments

OSPF uses the following formula to determine the cost of a route:

$$\text{cost} = \text{reference-bandwidth} / \text{interface bandwidth}$$

You can modify the reference-bandwidth value, which is used to calculate the default interface cost. The interface bandwidth value is not user-configurable and refers to the actual bandwidth of the physical interface.

By default, OSPF assigns a default cost metric of 1 to any link faster than 100 Mbps, and a default cost metric of 0 to the loopback interface (**lo0**). No bandwidth is associated with the loopback interface.

To control the flow of packets across the network, OSPF allows you to manually assign a cost (or metric) to a particular path segment. When you specify a metric for a specific OSPF interface, that value is used to determine the cost of routes advertised from that interface. For example, if all routers in the OSPF network use default metric values, and you increase the metric on one interface to 5, all paths through that interface have a calculated metric higher than the default and are not preferred.



NOTE: Any value you configure for the metric overrides the default behavior of using the reference-bandwidth value to calculate the route cost for that interface.

When there are multiple equal-cost routes to the same destination in a routing table, an equal-cost multipath (ECMP) set is formed. If there is an ECMP set for the active route, the Junos OS software uses a hash algorithm to choose one of the next-hop addresses in the ECMP set to install in the forwarding table.

You can configure Junos OS so that multiple next-hop entries in an ECMP set are installed in the forwarding table. Define a load-balancing routing policy by including one or more

policy-statement configuration statements at the **[edit policy-options]** hierarchy level, with the action **load-balance per-packet**. Then apply the routing policy to routes exported from the routing table to the forwarding table.

Dynamically Adjusting OSPF Interface Metrics Based on Bandwidth

You can specify a set of bandwidth threshold values and associated metric values for an OSPF interface or for a topology on an OSPF interface. When the bandwidth of an interface changes, Junos OS automatically sets the interface metric to the value associated with the appropriate bandwidth threshold value. Junos OS uses the smallest configured bandwidth threshold value that is equal to or greater than the actual interface bandwidth to determine the metric value. If the interface bandwidth is greater than any of the configured bandwidth threshold values, the metric value configured for the interface is used instead of any of the bandwidth-based metric values configured. The ability to recalculate the metric for an interface when its bandwidth changes is especially useful for aggregate interfaces.



NOTE: You must also configure a metric for the interface when you enable bandwidth-based metrics.

Controlling OSPF Route Preferences

You can control the flow of packets through the network using route preferences. Route preferences are used to select which route is installed in the forwarding table when several protocols calculate routes to the same destination. The route with the lowest preference value is selected.

By default, internal OSPF routes have a preference value of 10, and external OSPF routes have a preference value of 150. Although the default settings are appropriate for most environments, you might want to modify the default settings if all of the routing devices in your OSPF network use the default preference values, or if you are planning to migrate from OSPF to a different interior gateway protocol (IGP). If all of the devices use the default route preference values, you can change the route preferences to ensure that the path through a particular device is selected for the forwarding table any time multiple equal-cost paths to a destination exist. When migrating from OSPF to a different IGP, modifying the route preferences allows you to perform the migration in a controlled manner.

Example: Controlling the Cost of Individual OSPF Network Segments

This example shows how to control the cost of individual OSPF network segments.

- [Requirements on page 4463](#)
- [Overview on page 4463](#)
- [Configuration on page 4464](#)
- [Verification on page 4466](#)

Requirements

Before you begin:

- Configure the device interfaces. See the *Interfaces Feature Guide for Security Devices*.
- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 4400.
- Control OSPF designated router election. See “[Example: Controlling OSPF Designated Router Election](#)” on page 4402
- Configure a single-area OSPF network. See “[Example: Configuring a Single-Area OSPF Network](#)” on page 4406.

Overview

All OSPF interfaces have a cost, which is a routing metric that is used in the link-state calculation. Routes with lower total path metrics are preferred to those with higher path metrics. In this example, we explore how to control the cost of OSPF network segments.

By default, OSPF assigns a default cost metric of 1 to any link faster than 100 Mbps, and a default cost metric of 0 to the loopback interface (**lo0**). No bandwidth is associated with the loopback interface. This means that all interfaces faster than 100 Mbps have the same default cost metric of 1. If multiple equal-cost paths exist between a source and destination address, OSPF routes packets along each path alternately, in round-robin fashion.


Having the same default metric might not be a problem if all of the interfaces are running at the same speed. If the interfaces operate at different speeds, you might notice that traffic is not routed over the fastest interface because OSPF equally routes packets across the different interfaces. For example, if your routing device has Fast Ethernet and Gigabit Ethernet interfaces running OSPF, each of these interfaces have a default cost metric of 1.

In the first example, you set the reference bandwidth to 10g (10 Gbps, as denoted by 10,000,000,000 bits) by including the **reference-bandwidth** statement. With this configuration, OSPF assigns the Fast Ethernet interface a default metric of 100, and the Gigabit Ethernet interface a metric of 10. Since the Gigabit Ethernet interface has the lowest metric, OSPF selects it when routing packets. The range is 9600 through 1,000,000,000,000 bits.

[Figure 118](#) shows three routing devices in area 0.0.0.0 and assumes that the link between Device R2 and Device R3 is congested with other traffic. You can also control the flow of packets across the network by manually assigning a metric to a particular path segment. Any value you configure for the metric overrides the default behavior of using the reference-bandwidth value to calculate the route cost for that interface. To prevent the traffic from Device R3 going directly to Device R2, you adjust the metric on the interface on Device R3 that connects with Device R1 so that all traffic goes through Device R1.

In the second example, you set the metric to 5 on interface **fe-1/0/1** on Device R3 that connects with Device R1 by including the **metric** statement. The range is 1 through 65,535.

Figure 118: OSPF Metric Configuration



Configuration

- [Configuring the Reference Bandwidth on page 4464](#)
- [Configuring a Metric for a Specific OSPF Interface on page 4465](#)

Configuring the Reference Bandwidth

CLI Quick Configuration

To quickly configure the reference bandwidth, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

```
[edit]
set protocols ospf reference-bandwidth 10g
```

Step-by-Step Procedure

To configure the reference bandwidth:

1. Configure the reference bandwidth to calculate the default interface cost.



NOTE: To specify OSPFv3, include the `ospf3` statement at the [edit protocols] hierarchy level.

```
[edit]
user@host# set protocols ospf reference-bandwidth 10g
```



TIP: As a shortcut in this example, you enter 10g to specify 10 Gbps reference bandwidth. Whether you enter 10g or 10000000000, the output of `show protocols ospf` command displays 10 Gbps as 10g, not 10000000000.

2. If you are done configuring the device, commit the configuration.


```
[edit]
user@host# commit
```



NOTE: Repeat this entire configuration on all routing devices in a shared network.

Results Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
reference-bandwidth 10g;
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

Configuring a Metric for a Specific OSPF Interface

CLI Quick Configuration To quickly configure a metric for a specific OSPF interface, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

```
[edit]
set protocols ospf area 0.0.0.0 interface fe-1/0/1 metric 5
```

Step-by-Step Procedure To configure the metric for a specific OSPF interface:

1. Create an OSPF area.



NOTE: To specify OSPFv3, include the **ospf3** statement at the [edit protocols] hierarchy level.

```
[edit]
user@host# edit protocols ospf area 0.0.0.0
```

2. Configure the metric of the OSPF network segment.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# set interface fe-1/0/1 metric 5
```

3. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# commit
```

Results Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
```

```
area 0.0.0.0 {  
  interface fe-1/0/1.0 {  
    metric 5;  
  }  
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

Verification

Confirm that the configuration is working properly.

- [Verifying the Configured Metric on page 4466](#)
- [Verifying the Route on page 4466](#)

Verifying the Configured Metric

Purpose Verify the metric setting on the interface. Confirm that the Cost field displays the interface's configured metric (cost). When choosing paths to a destination, OSPF uses the path with the lowest cost.

Action From operational mode, enter the **show ospf interface detail** command for OSPFv2, and enter the **show ospf3 interface detail** command for OSPFv3.

Verifying the Route

Purpose When choosing paths to a destination, OSPF uses the path with the lowest total cost. Confirm that OSPF is using the appropriate path.

Action From operational mode, enter the **show route** command.

Example: Dynamically Adjusting OSPF Interface Metrics Based on Bandwidth

This example shows how to dynamically adjust OSPF interface metrics based on bandwidth.

- [Requirements on page 4466](#)
- [Overview on page 4467](#)
- [Configuration on page 4467](#)
- [Verification on page 4468](#)

Requirements

Before you begin:

- Configure the device interfaces. See the *Interfaces Feature Guide for Security Devices*.
- Configure the router identifiers for the devices in your OSPF network. See "[Example: Configuring an OSPF Router Identifier](#)" on page 4400.

- Control OSPF designated router election. See [“Example: Controlling OSPF Designated Router Election” on page 4402](#)
- Configure a single-area OSPF network. See [“Example: Configuring a Single-Area OSPF Network” on page 4406](#).

Overview

You can specify a set of bandwidth threshold values and associated metric values for an OSPF interface. When the bandwidth of an interface changes, Junos OS automatically sets the interface metric to the value associated with the appropriate bandwidth threshold value. When you configure bandwidth-based metric values, you typically configure multiple bandwidth and metric values.

In this example, you configure OSPF interface **ae0** for bandwidth-based metrics by including the **bandwidth-based-metrics** statement and the following settings:

- **bandwidth**—Specifies the bandwidth threshold in bits per second. The range is 9600 through 1,000,000,000,000,000.
- **metric**—Specifies the metric value to associate with a specific bandwidth value. The range is 1 through 65,535.

Configuration

CLI Quick Configuration

To quickly configure bandwidth threshold values and associated metric values for an OSPF interface, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

```
[edit]
set protocols ospf area 0.0.0.0 interface ae0.0 metric 5
set protocols ospf area 0.0.0.0 interface ae0.0 bandwidth-based-metrics bandwidth 1g
metric 60
set protocols ospf area 0.0.0.0 interface ae0.0 bandwidth-based-metrics bandwidth 10g
metric 50
```

To configure the metric for a specific OSPF interface:

1. Create an OSPF area.



NOTE: To specify OSPFv3, include the **ospf3** statement at the [edit protocols] hierarchy level.

```
[edit]
user@host# edit protocols ospf area 0.0.0.0
```

2. Configure the metric of the OSPF network segment.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# set interface ae0 metric 5
```

3. Configure the bandwidth threshold values and associated metric values.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# set interface ae0.0 bandwidth-based-metrics bandwidth 1g metric 60
user@host# set interface ae0.0 bandwidth-based-metrics bandwidth 10g metric 50
```

4. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# commit
```

Results Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.0 {
  interface ae0.0 {
    bandwidth-based-metrics {
      bandwidth 1g metric 60;
      bandwidth 10g metric 50;
    }
    metric 5;
  }
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

Verification

Confirm that the configuration is working properly.

Verifying the Configured Metric

Purpose Verify the metric setting on the interface. Confirm that the Cost field displays the interface's configured metric (cost). When choosing paths to a destination, OSPF uses the path with the lowest cost.

Action From operational mode, enter the **show ospf interface detail** command for OSPFv2, and enter the **show ospf3 interface detail** command for OSPFv3.

Example: Controlling OSPF Route Preferences

This example shows how to control OSPF route selection in the forwarding table. This example also shows how you might control route selection if you are migrating from OSPF to another IGP.

- [Requirements on page 4469](#)
- [Overview on page 4469](#)
- [Configuration on page 4469](#)
- [Verification on page 4470](#)

Requirements

This example assumes that OSPF is properly configured and running in your network, and you want to control route selection because you are planning to migrate from OSPF to a different IGP.

- Configure the device interfaces. See the *Interfaces Feature Guide for Security Devices*.
- Configure the IGP that you want to migrate to.

Overview

Route preferences are used to select which route is installed in the forwarding table when several protocols calculate routes to the same destination. The route with the lowest preference value is selected.

By default, internal OSPF routes have a preference value of 10, and external OSPF routes have a preference value of 150. You might want to modify this setting if you are planning to migrate from OSPF to a different IGP. Modifying the route preferences enables you to perform the migration in a controlled manner.

This example makes the following assumptions:

- OSPF is already running in your network.
- You want to migrate from OSPF to IS-IS.
- You configured IS-IS per your network requirements and confirmed it is working properly.

In this example, you increase the OSPF route preference values to make them less preferred than IS-IS routes by specifying 168 for internal OSPF routes and 169 for external OSPF routes. IS-IS internal routes have a preference of either 15 (for Level 1) or 18 (for Level 2), and external routes have a preference of 160 (for Level 1) or 165 (for Level 2). In general, it is preferred to leave the new protocol at its default settings to minimize complexities and simplify any future addition of routing devices to the network. To modify the OSPF route preference values, configure the following settings:

- **preference**—Specifies the route preference for internal OSPF routes. By default, internal OSPF routes have a value of 10. The range is from 0 through 4,294,967,295 ($2^{32} - 1$).
- **external-preference**—Specifies the route preference for external OSPF routes. By default, external OSPF routes have a value of 150. The range is from 0 through 4,294,967,295 ($2^{32} - 1$).

Configuration

CLI Quick Configuration

To quickly configure the OSPF route preference values, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

```
[edit]
set protocols ospf preference 168 external-preference 169
```

To configure route selection:

1. Enter OSPF configuration mode and set the external and internal routing preferences.



NOTE: To specify OSPFv3, include the **ospf3** statement at the **[edit protocols]** hierarchy level.

```
[edit]
user@host# set protocols ospf preference 168 external-preference 169
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Results Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
preference 168;
external-preference 169;
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

Verification

Confirm that the configuration is working properly.

- [Verifying the Route on page 4470](#)

Verifying the Route

Purpose Verify that the IGP is using the appropriate route. After the new IGP becomes the preferred protocol (in this example, IS-IS), you should monitor the network for any issues. After you confirm that the new IGP is working properly, you can remove the OSPF configuration from the routing device by entering the **delete ospf** command at the **[edit protocols]** hierarchy level.

Action From operational mode, enter the **show route** command.

Related Documentation

- [OSPF Overview on page 4386](#)
- [Understanding OSPF Configurations](#)

Example: Configuring OSPF Overload Mode

- [OSPF Overload Function Overview on page 4471](#)
- [Example: Configuring OSPF to Make Routing Devices Appear Overloaded on page 4472](#)

OSPF Overload Function Overview

If the time elapsed after the OSPF instance is enabled is less than the specified timeout, overload mode is set.

You can configure the local routing device so that it appears to be overloaded. An overloaded routing device determines it is unable to handle any more OSPF transit traffic, which results in sending OSPF transit traffic to other routing devices. OSPF traffic to directly attached interfaces continues to reach the routing device. You might configure overload mode for many reasons, including:

- If you want the routing device to participate in OSPF routing, but do not want it to be used for transit traffic. This could include a routing device that is connected to the network for analysis purposes, but is not considered part of the production network, such as network management routing devices.
- If you are performing maintenance on a routing device in a production network. You can move traffic off that routing device so network services are not interrupted during your maintenance window.

You configure or disable overload mode in OSPF with or without a timeout. Without a timeout, overload mode is set until it is explicitly deleted from the configuration. With a timeout, overload mode is set if the time elapsed since the OSPF instance started is less than the specified timeout.

A timer is started for the difference between the timeout and the time elapsed since the instance started. When the timer expires, overload mode is cleared. In overload mode, the router link-state advertisement (LSA) is originated with all the transit router links (except stub) set to a metric of 0xFFFF. The stub router links are advertised with the actual cost of the interfaces corresponding to the stub. This causes the transit traffic to avoid the overloaded routing device and to take paths around the routing device. However, the overloaded routing device's own links are still accessible.

The routing device can also dynamically enter the overload state, regardless of configuring the device to appear overloaded. For example, if the routing device exceeds the configured OSPF prefix limit, the routing device purges the external prefixes and enters into an overload state.

In cases of incorrect configurations, the huge number of routes might enter OSPF, which can hamper the network performance. To prevent this, **prefix-export-limit** should be configured which will purge externals and prevent the network from the bad impact.

By allowing any number of routes to be exported into OSPF, the routing device can become overwhelmed and potentially flood an excessive number of routes into an area. You can limit the number of routes exported into OSPF to minimize the load on the routing device and prevent this potential problem.

By default, there is no limit to the number of prefixes (routes) that can be exported into OSPF. To prevent this, **prefix-export-limit** should be configured which will purge externals and prevent the network.

To limit the number of prefixes exported to OSPF:

```
[edit]
set protocols ospf prefix-export-limit number
```

The prefix export limit number can be a value from 0 through 4,294,967,295.

Example: Configuring OSPF to Make Routing Devices Appear Overloaded

This example shows how to configure a routing device running OSPF to appear to be overloaded.

- [Requirements on page 4472](#)
- [Overview on page 4472](#)
- [Configuration on page 4473](#)
- [Verification on page 4474](#)

Requirements

Before you begin:

- Configure the device interfaces. See the *Interfaces Feature Guide for Security Devices*.
- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 4400.
- Control OSPF designated router election. See “[Example: Controlling OSPF Designated Router Election](#)” on page 4402.
- Configure a single-area OSPF network. See “[Example: Configuring a Single-Area OSPF Network](#)” on page 4406.
- Configure a multiarea OSPF network. See “[Example: Configuring a Multiarea OSPF Network](#)” on page 4408.

Overview

You can configure a local routing device running OSPF to appear to be overloaded, which allows the local routing device to participate in OSPF routing, but not for transit traffic. When configured, the transit interface metrics are set to the maximum value of 65535.

This example includes the following settings:

- **overload**—Configures the local routing device so it appears to be overloaded. You might configure this if you want the routing device to participate in OSPF routing, but do not want it to be used for transit traffic, or you are performing maintenance on a routing device in a production network.
- **timeout seconds**—(Optional) Specifies the number of seconds at which the overload is reset. If no timeout interval is specified, the routing device remains in the overload state until the overload statement is deleted or a timeout is set. In this example, you configure 60 seconds as the amount of time the routing device remains in the overload state. By default, the timeout interval is 0 seconds (this value is not configured). The range is from 60 through 1800 seconds.

Configuration

CLI Quick Configuration To quickly configure a local routing device to appear as overloaded, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

```
[edit]
set protocols ospf overload timeout 60
```

Step-by-Step Procedure To configure a local routing device to appear overloaded:

1. Enter OSPF configuration mode.



NOTE: To specify OSPFv3, include the **ospf3** statement at the [edit protocols] hierarchy level.

```
[edit]
user@host# edit protocols ospf
```

2. Configure the local routing device to be overloaded.

```
[edit protocols ospf]
user@host# set overload
```

3. (Optional) Configure the number of seconds at which overload is reset.

```
[edit protocols ospf]
user@host# set overload timeout 60
```

4. (Optional) Configure the limit on the number prefixes exported to OSPF, to minimise the load on the routing device and prevent the device from entering the overload mode.

```
[edit protocols ospf]
user@host# set prefix-export-limit 50
```

5. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf]
user@host# commit
```

Results Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration. The output includes the optional **timeout** and **prefix-export-limit** statements.

```
user@host# show protocols ospf
```

```
prefix-export-limit 50;
overload timeout 60;
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

Verification

Confirm that the configuration is working properly.

- [Verifying Traffic Has Moved Off Devices on page 4474](#)
- [Verifying Transit Interface Metrics on page 4474](#)
- [Verifying the Overload Configuration on page 4474](#)
- [Verifying the Viable Next Hop on page 4474](#)

Verifying Traffic Has Moved Off Devices

Purpose Verify that the traffic has moved off the upstream devices.

Action From operational mode, enter the **show interfaces detail** command.

Verifying Transit Interface Metrics

Purpose Verify that the transit interface metrics are set to the maximum value of 65535 on the downstream neighboring device.

Action From operational mode, enter the **show ospf database router detail advertising-router address** command for OSPFv2, and enter the **show ospf3 database router detail advertising-router address** command for OSPFv3.

Verifying the Overload Configuration

Purpose Verify that overload is configured by reviewing the Configured overload field. If the overload timer is also configured, this field also displays the time that remains before it is set to expire.

Action From operational mode, enter the **show ospf overview** command for OSPFv2, and the **show ospf3 overview** command for OSPFv3.

Verifying the Viable Next Hop

Purpose Verify the viable next hop configuration on the upstream neighboring device. If the neighboring device is overloaded, it is not used for transit traffic and is not displayed in the output.

Action From operational mode, enter the **show route address** command.

Related Documentation

- [OSPF Overview on page 4386](#)
- [Understanding OSPF Configurations](#)

OSPF Fault Detection Configuration

- [Example: Configuring OSPF Timers on page 4475](#)
- [Example: Configuring BFD for OSPF on page 4482](#)
- [Example: Configuring BFD Authentication for OSPF on page 4488](#)

Example: Configuring OSPF Timers

- [OSPF Timers Overview on page 4475](#)
- [Example: Configuring OSPF Timers on page 4476](#)

OSPF Timers Overview

OSPF routing devices constantly track the status of their neighbors, sending and receiving hello packets that indicate whether each neighbor still is functioning, and sending and receiving link-state advertisement (LSA) and acknowledgment packets. OSPF sends packets and expects to receive packets at specified intervals.

You configure OSPF timers on the interface of the routing device participating in OSPF. Depending on the timer, the configured interval must be the same on all routing devices on a shared network (area).

You can configure the following OSPF timers:

- Hello interval—Routing devices send hello packets at a fixed interval on all interfaces, including virtual links, to establish and maintain neighbor relationships. The hello interval specifies the length of time, in seconds, before the routing device sends a hello packet out of an interface. This interval must be the same on all routing devices on a shared network. By default, the routing device sends hello packets every 10 seconds (broadcast and point-to-point networks) and 30 seconds (nonbroadcast multiple access (NBMA) networks).
- Poll interval—(OSPFv2, Nonbroadcast networks only) Routing devices send hello packets for a longer interval on nonbroadcast networks to minimize the bandwidth required on slow WAN links. The poll interval specifies the length of time, in seconds, before the routing device sends hello packets out of the interface before establishing adjacency with a neighbor. By default, the routing device sends hello packets every 120 seconds until active neighbors are detected.

Once the routing device detects an active neighbor, the hello packet interval changes from the time specified in the poll interval to the time specified in the hello interval.

- **LSA retransmission interval**—When a routing device sends LSAs to its neighbors, the routing device expects to receive an acknowledgment packet from each neighbor within a certain amount of time. The LSA retransmission interval specifies the length of time, in seconds, that the routing device waits to receive an LSA packet before retransmitting the LSA to an interface's neighbors. By default, the routing device waits 5 seconds for an acknowledgment before retransmitting the LSA.
- **Dead interval**—If a routing device does not receive a hello packet from a neighbor within a fixed amount of time, the routing device modifies its topology database to indicate that the neighbor is nonoperational. The dead interval specifies the length of time, in seconds, that the routing device waits before declaring that a neighboring routing device is unavailable. This is an interval during which the routing device receives no hello packets from the neighbor. This interval must be the same on all routing devices on a shared network. By default, this interval is four times the default hello interval, which is 40 seconds (broadcast and point-to-point networks) and 120 seconds (NBMA networks).
- **Transit delay**—Before a link-state update packet is propagated out of an interface, the routing device must increase the age of the packet. The transit delay sets the estimated time required to transmit a link-state update on the interface. By default, the transit delay is 1 second. You should never have to modify the transit delay time.

Example: Configuring OSPF Timers

This example shows how to configure the OSPF timers.

- [Requirements on page 4476](#)
- [Overview on page 4477](#)
- [Configuration on page 4478](#)
- [Verification on page 4481](#)

Requirements

Before you begin:

- Configure the device interfaces. See the *Interfaces Feature Guide for Security Devices*.
- Configure the router identifiers for the devices in your OSPF network. See "[Example: Configuring an OSPF Router Identifier](#)" on page 4400.
- Control OSPF designated router election. See "[Example: Controlling OSPF Designated Router Election](#)" on page 4402.
- Configure a single-area OSPF network. See "[Example: Configuring a Single-Area OSPF Network](#)" on page 4406.
- Configure a multiarea OSPF network. See "[Example: Configuring a Multiarea OSPF Network](#)" on page 4408.

Overview

The default OSPF timer settings are optimal for most networks. However, depending on your network requirements, you might need to modify the timer settings. This example explains why you might need to modify the following timers:

- Hello interval
- Dead interval
- LSA retransmission interval
- Transit delay

Hello Interval and Dead Interval

The hello interval and the dead interval optimize convergence times by efficiently tracking neighbor status. By lowering the values of the hello interval and the dead interval, you can increase the convergence of OSPF routes if a path fails. These intervals must be the same on all routing devices on a shared network. Otherwise, OSPF cannot establish the appropriate adjacencies.

In the first example, you lower the hello interval to 2 seconds and the dead interval to 8 seconds on point-to-point OSPF interfaces **fe-0/0/1** and **fe-1/0/1** in area 0.0.0.0 by configuring the following settings:

- **hello-interval**—Specifies the length of time, in seconds, before the routing device sends a hello packet out of an interface. By default, the routing device sends hello packets every 10 seconds. The range is from 1 through 255 seconds.
- **dead-interval**—Specifies the length of time, in seconds, that the routing device waits before declaring that a neighboring routing device is unavailable. This is an interval during which the routing device receives no hello packets from the neighbor. By default, the routing device waits 40 seconds (four times the hello interval). The range is 1 through 65,535 seconds.

LSA Retransmission Interval

The link-state advertisement (LSA) retransmission interval optimizes the sending and receiving of LSA and acknowledgement packets. You must configure the LSA retransmission interval to be equal to or greater than 3 seconds to avoid triggering a retransmit trap because the Junos OS delays LSA acknowledgments by up to 2 seconds. If you have a virtual link, you might find increased performance by increasing the value of the LSA retransmission interval.

In the second example, you increase the LSA retransmission timer to 8 seconds on OSPF interface **fe-0/0/1** in area 0.0.0.1 by configuring the following setting:

- **retransmit-interval**—Specifies the length of time, in seconds, that the routing device waits to receive an LSA packet before retransmitting LSA to an interface's neighbors. By default, the routing device retransmits LSAs to its neighbors every 5 seconds. The range is from 1 through 65,535 seconds.

Transit Delay

The transit delay sets the time the routing device uses to age a link-state update packet. If you have a slow link (for example, one with an average propagation delay of multiple seconds), you should increase the age of the packet by a similar amount. Doing this ensures that you do not receive a packet back that is younger than the original copy.

In the final example, you increase the transit delay to 2 seconds on OSPF interface **fe-1/0/1** in area 0.0.0.1. By configuring the following setting, this causes the routing device to age the link-state update packet by 2 seconds:

- **transit-delay**—Sets the estimated time required to transmit a link-state update on the interface. You should never have to modify the transit delay time. By default, the routing device ages the packet by 1 second. The range is from 1 through 65,535 seconds.

Configuration

- [Configuring the Hello Interval and the Dead Interval on page 4478](#)
- [Controlling the LSA Retransmission Interval on page 4479](#)
- [Specifying the Transit Delay on page 4480](#)

Configuring the Hello Interval and the Dead Interval

CLI Quick Configuration

To quickly configure the hello and dead intervals, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

```
[edit]
set protocols ospf area 0.0.0.0 interface fe-0/0/1 hello-interval 2
set protocols ospf area 0.0.0.0 interface fe-0/0/1 dead-interval 8
set protocols ospf area 0.0.0.0 interface fe-1/0/1 hello-interval 2
set protocols ospf area 0.0.0.0 interface fe-1/0/1 dead-interval 8
```

Step-by-Step Procedure

To configure the hello and dead intervals:

1. Create an OSPF area.



NOTE: To specify OSPFv3, include the **ospf3** statement at the [edit protocols] hierarchy level.

[edit]

```
user@host# edit protocols ospf area 0.0.0.0
```

2. Specify the interfaces.

```
[edit protocols ospf area 0.0.0.0]
user@host# set interface fe-0/0/1
user@host# set interface fe-1/0/1
```

3. Configure the hello interval.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# set interface fe-0/0/1 hello-interval 2
user@host# set interface fe-1/0/1 hello-interval 2
```

4. Configure the dead interval.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# set interface fe-0/0/1 dead-interval 8
user@host# set interface fe-1/0/1 dead-interval 8
```

5. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# commit
```



NOTE: Repeat this entire configuration on all routing devices in a shared network.

Results Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.0 {
  interface fe-0/0/1.0 {
    hello-interval 2;
    dead-interval 8;
  }
  interface fe-1/0/1.0 {
    hello-interval 2;
    dead-interval 8;
  }
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

Controlling the LSA Retransmission Interval

CLI Quick Configuration To quickly configure the LSA retransmission interval, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

```
[edit]
set protocols ospf area 0.0.0.1 interface fe-0/0/1 retransmit-interval 8
```

Step-by-Step Procedure To configure the LSA retransmission interval:

1. Create an OSPF area.



NOTE: To specify OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# edit protocols ospf area 0.0.0.1
```

2. Specify the interface.

```
[edit protocols ospf area 0.0.0.1]
user@host# set interface fe-0/0/1
```

3. Configure the LSA retransmission interval.

```
[edit protocols ospf area 0.0.0.1 ]
user@host# set interface fe-0/0/1 retransmit-interval 8
```

4. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.1]
user@host# commit
```

Results Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.1 {
  interface fe-0/0/1.0 {
    retransmit-interval 8;
  }
}
```

To confirm your OSPFv3 configuration, enter the `show protocols ospf3` command.

Specifying the Transit Delay

CLI Quick Configuration To quickly configure the transit delay, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the `[edit]` hierarchy level, and then enter `commit` from configuration mode.

```
[edit]
set protocols ospf area 0.0.0.1 interface fe-1/0/1 transit-delay 2
```

Step-by-Step Procedure To configure the transit delay:

1. Create an OSPF area.



NOTE: To specify OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# edit protocols ospf area 0.0.0.1
```

2. Specify the interface.

```
[edit protocols ospf area 0.0.0.1]
user@host# set interface fe-1/0/1
```

3. Configure the transit delay.

```
[edit protocols ospf area 0.0.0.1 ]
user@host# set interface fe-1/0/1 transit-delay 2
```

4. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.1 ]
user@host# commit
```

Results Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.1 {
  interface fe-1/0/1.0 {
    transit-delay 2;
  }
}
```

To confirm your OSPFv3 configuration, enter the `show protocols ospf3` command.

Verification

Confirm that the configuration is working properly.

Verifying the Timer Configuration

Purpose Verify that the interface for OSPF or OSPFv3 has been configured with the applicable timer values. Confirm that the Hello field, the Dead field, and the ReXmit field display the values that you configured.

Action From operational mode, enter the `show ospf interface detail` for OSPFv2, and enter the `show ospf3 interface detail` command for OSPFv3.

Related Documentation

- [OSPF Overview on page 4386](#)
- [Understanding OSPF Configurations](#)

Example: Configuring BFD for OSPF

- [Understanding BFD for OSPF on page 4482](#)
- [Example: Configuring BFD for OSPF on page 4484](#)

Understanding BFD for OSPF

The Bidirectional Forwarding Detection (BFD) protocol is a simple hello mechanism that detects failures in a network. BFD works with a wide variety of network environments and topologies. A pair of routing devices exchange BFD packets. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the routing device stops receiving a reply after a specified interval. The BFD failure detection timers have shorter time limits than the OSPF failure detection mechanisms, so they provide faster detection.

The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the **clear bfd adaptation** command to return BFD interval timers to their configured values. The **clear bfd adaptation** command is hitless, meaning that the command does not affect traffic flow on the routing device.



NOTE: BFD is supported for OSPFv3 in Junos OS Release 9.3 and later.

You can configure the following BFD protocol settings:

- **detection-time threshold**—Threshold for the adaptation of the detection time. When the BFD session detection time adapts to a value equal to or greater than the configured threshold, a single trap and a single system log message are sent.
- **full-neighbors-only**—Ability to establish BFD sessions only for OSPF neighbors with full neighbor adjacency. The default behavior is to establish BFD sessions for all OSPF neighbors. This setting is available in Junos OS Release 9.5 and later.
- **minimum-interval**—Minimum transmit and receive interval for failure detection. This setting configures both the minimum interval after which the local routing device transmits hello packets and the minimum interval after which the routing device expects to receive a reply from the neighbor with which it has established a BFD session. Both intervals are in milliseconds. You can also specify the minimum transmit and receive intervals separately using the **transmit-interval** **minimum-interval** and **minimum-receive-interval** statements.



NOTE: BFD is an intensive protocol that consumes system resources. Specifying a minimum interval for BFD of less than 100 ms for Routing Engine-based sessions and 10 ms for distributed BFD sessions can cause undesired BFD flapping.

Depending on your network environment, these additional recommendations might apply:

- For large-scale network deployments with a large number of BFD sessions, specify a minimum interval of 300 ms for Routing Engine-based sessions and 100 ms for distributed BFD sessions.
- For very large-scale network deployments with a large number of BFD sessions, contact Juniper Networks customer support for more information.
- For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing (NSR) is configured, specify a minimum interval of 2500 ms for Routing Engine-based sessions. Without NSR, Routing Engine-based sessions can have a minimum interval of 100 ms. In OSPFv3, BFD is always based in the Routing Engine, meaning that BFD is not distributed. For distributed BFD sessions with NSR configured, the minimum interval recommendations are unchanged and depend only on your network deployment.

- **minimum-receive-interval**—Minimum receive interval for failure detection. This setting configures the minimum receive interval, in milliseconds, after which the routing device expects to receive a hello packet from a neighbor with which it has established a BFD session. You can also specify the minimum receive interval using the **minimum-interval** statement.
- **multiplier**—Multiplier for hello packets. This setting configures the number of hello packets that are not received by a neighbor, which causes the originating interface to be declared down. By default, three missed hello packets cause the originating interface to be declared down.
- **no-adaptation**—Disables BFD adaption. This setting disables BFD sessions from adapting to changing network conditions. This setting is available in Junos OS Release 9.0 and later.



NOTE: We recommend that you do not disable BFD adaptation unless it is preferable not to have BFD adaptation in your network.

- **transmit-interval minimum-interval**—Minimum transmit interval for failure detection. This setting configures the minimum transmit interval, in milliseconds, at which the local routing device transmits hello packets to the neighbor with which it has established a BFD session. You can also specify the minimum transmit interval using the **minimum-interval** statement.

- **transmit-interval threshold**—Threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system log message are sent. The threshold value must be greater than the minimum transmit interval. If you attempt to commit a configuration with a threshold value less than the minimum transmit interval, the routing device displays an error and does not accept the configuration.
- **version**—BFD version. This setting configures the BFD version used for detection. You can explicitly configure BFD version 1, or the routing device can automatically detect the BFD version. By default, the routing device automatically detects the BFD version automatically, which is either 0 or 1.

You can also trace BFD operations for troubleshooting purposes.

Example: Configuring BFD for OSPF

This example shows how to configure the Bidirectional Forwarding Detection (BFD) protocol for OSPF.

- [Requirements on page 4484](#)
- [Overview on page 4484](#)
- [Configuration on page 4486](#)
- [Verification on page 4488](#)

Requirements

Before you begin:

- Configure the device interfaces. See the *Interfaces Feature Guide for Security Devices*.
- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 4400.
- Control OSPF designated router election. See “[Example: Controlling OSPF Designated Router Election](#)” on page 4402.
- Configure a single-area OSPF network. See “[Example: Configuring a Single-Area OSPF Network](#)” on page 4406.
- Configure a multiarea OSPF network. See “[Example: Configuring a Multiarea OSPF Network](#)” on page 4408.
- Configure a multiarea OSPF network. See “[Example: Configuring a Multiarea OSPF Network](#)” on page 4408.

Overview

An alternative to adjusting the OSPF hello interval and dead interval settings to increase route convergence is to configure BFD. The BFD protocol is a simple hello mechanism that detects failures in a network. The BFD failure detection timers have shorter timer limits than the OSPF failure detection mechanisms, thereby providing faster detection.

BFD is useful on interfaces that are unable to detect failure quickly, such as Ethernet interfaces. Other interfaces, such as SONET interfaces, already have built-in failure detection. Configuring BFD on those interfaces is unnecessary.

You configure BFD on a pair of neighboring OSPF interfaces. Unlike the OSPF hello interval and dead interval settings, you do not have to enable BFD on all interfaces in an OSPF area.

In this example, you enable failure detection by including the **bfd-liveness-detection** statement on the neighbor OSPF interface **fe-0/1/0** in area 0.0.0.0 and configure the BFD packet exchange interval to 300 milliseconds, configure 4 as the number of missed hello packets that causes the originating interface to be declared down, and configure BFD sessions only for OSPF neighbors with full neighbor adjacency by including the following settings:

- **full-neighbors-only**—In Junos OS Release 9.5 and later, configures the BFD protocol to establish BFD sessions only for OSPF neighbors with full neighbor adjacency. The default behavior is to establish BFD sessions for all OSPF neighbors.
- **minimum-interval**—Configures the minimum interval, in milliseconds, after which the local routing device transmits hello packets as well as the minimum interval after which the routing device expects to receive a reply from the neighbor with which it has established a BFD session. You can configure a number in the range from 1 through 255,000 milliseconds. You can also specify the minimum transmit and receive intervals separately using the **transmit-interval** **minimum-interval** and **minimum-receive-interval** statements.



NOTE: BFD is an intensive protocol that consumes system resources. Specifying a minimum interval for BFD of less than 300 ms for Routing Engine-based sessions and 10 ms for distributed BFD sessions can cause undesired BFD flapping.

Depending on your network environment, these additional recommendations might apply:

- For large-scale network deployments with a large number of BFD sessions, specify a minimum interval of 300 ms for Routing Engine-based sessions and 300 ms for distributed BFD sessions.



NOTE:

- For the `bfdd` process, the detection time interval set is lower than 300 ms. If there is a high priority process such as `ppmd` running on the system, the CPU might spend time on the `ppmd` process rather than the `bfdd` process.
- For branch SRX Series devices, we recommend 1000 ms as the minimum keepalive time interval for BFD packets.

- For very large-scale network deployments with a large number of BFD sessions, contact Juniper Networks customer support for more information.
- For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing (NSR) is configured, specify a minimum interval of 2500 ms for Routing Engine-based sessions. For distributed BFD sessions with NSR configured, the minimum interval recommendations are unchanged and depend only on your network deployment.

- **multiplier**—Configures the number of hello packets not received by a neighbor that causes the originating interface to be declared down. By default, three missed hello packets cause the originating interface to be declared down. You can configure a value in the range from 1 through 255.

Configuration

CLI Quick Configuration

To quickly configure the BFD protocol for OSPF, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the `[edit]` hierarchy level, and then enter **commit** from configuration mode.

```
[edit]
set protocols ospf area 0.0.0.0 interface fe-0/0/1 bfd-liveness-detection minimum-interval 300
set protocols ospf area 0.0.0.0 interface fe-0/0/1 bfd-liveness-detection multiplier 4
set protocols ospf area 0.0.0.0 interface fe-0/0/1 bfd-liveness-detection full-neighbors-only
```

Step-by-Step Procedure To configure the BFD protocol for OSPF on one neighboring interface:

1. Create an OSPF area.



NOTE: To specify OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# edit protocols ospf area 0.0.0.0
```

2. Specify the interface.

```
[edit protocols ospf area 0.0.0.0]
user@host# set interface fe-0/0/1
```

3. Specify the minimum transmit and receive intervals.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# set interface fe-0/0/1 bfd-liveness-detection minimum-interval 300
```

4. Configure the number of missed hello packets that cause the originating interface to be declared down.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# set interface fe-0/0/1 bfd-liveness-detection multiplier 4
```

5. Configure BFD sessions only for OSPF neighbors with full neighbor adjacency.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# set interface fe-0/0/1 bfd-liveness-detection full-neighbors-only
```

6. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# commit
```



NOTE: Repeat this entire configuration on the other neighboring interface.

Results Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.0 {
  interface fe-0/0/1.0 {
    bfd-liveness-detection {
      minimum-interval 300;
      multiplier 4;
      full-neighbors-only;
    }
  }
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

Verification

Confirm that the configuration is working properly.

Verifying the BFD Sessions

Purpose Verify that the OSPF interfaces have active BFD sessions, and that session components have been configured correctly.

Action From operational mode, enter the **show bfd session detail** command.

Meaning The output displays information about the BFD sessions.

- The Address field displays the IP address of the neighbor.
- The Interface field displays the interface you configured for BFD.
- The State field displays the state of the neighbor and should show Full to reflect the full neighbor adjacency that you configured.
- The Transmit Interval field displays the time interval you configured to send BFD packets.
- The Multiplier field displays the multiplier you configured.

Related Documentation

- [Understanding OSPF Configurations](#)
- [Understanding BFD Authentication for OSPF on page 4488](#)

Example: Configuring BFD Authentication for OSPF

- [Understanding BFD Authentication for OSPF on page 4488](#)
- [Configuring BFD Authentication for OSPF on page 4490](#)

Understanding BFD Authentication for OSPF

Bidirectional Forwarding Detection (BFD) enables rapid detection of communication failures between adjacent systems. By default, authentication for BFD sessions is disabled. However, when you run BFD over Network Layer protocols, the risk of service attacks can be significant. We strongly recommend using authentication if you are running BFD over multiple hops or through insecure tunnels. Beginning with Junos OS Release 9.6, Junos OS supports authentication for BFD sessions running over OSPFv2. BFD authentication is not supported on MPLS OAM sessions. BFD authentication is only supported in the Canada and United States version of the Junos OS image and is not available in the export version.

You authenticate BFD sessions by specifying an authentication algorithm and keychain, and then associating that configuration information with a security authentication keychain using the keychain name.

The following sections describe the supported authentication algorithms, security keychains, and level of authentication that can be configured:

- [BFD Authentication Algorithms on page 4489](#)
- [Security Authentication Keychains on page 4490](#)
- [Strict Versus Loose Authentication on page 4490](#)

BFD Authentication Algorithms

Junos OS supports the following algorithms for BFD authentication:

- **simple-password**—Plain-text password. One to 16 bytes of plain text are used to authenticate the BFD session. One or more passwords can be configured. This method is the least secure and should be used only when BFD sessions are not subject to packet interception.
- **keyed-md5**—Keyed Message Digest 5 hash algorithm for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed MD5 uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than or equal to the last sequence number received. Although more secure than a simple password, this method is vulnerable to replay attacks. Increasing the rate at which the sequence number is updated can reduce this risk.
- **meticulous-keyed-md5**—Meticulous keyed Message Digest 5 hash algorithm. This method works in the same manner as keyed MD5, but the sequence number is updated with every packet. Although more secure than keyed MD5 and simple passwords, this method might take additional time to authenticate the session.
- **keyed-sha-1**—Keyed Secure Hash Algorithm I for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed SHA uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. The key is not carried within the packets. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than the last sequence number received.
- **meticulous-keyed-sha-1**—Meticulous keyed Secure Hash Algorithm I. This method works in the same manner as keyed SHA, but the sequence number is updated with every packet. Although more secure than keyed SHA and simple passwords, this method might take additional time to authenticate the session.



NOTE: Nonstop active routing (NSR) is not supported with the meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

Security Authentication Keychains

The security authentication keychain defines the authentication attributes used for authentication key updates. When the security authentication keychain is configured and associated with a protocol through the keychain name, authentication key updates can occur without interrupting routing and signaling protocols.

The authentication keychain contains one or more keychains. Each keychain contains one or more keys. Each key holds the secret data and the time at which the key becomes valid. The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

BFD allows multiple clients per session, and each client can have its own keychain and algorithm defined. To avoid confusion, we recommend specifying only one security authentication keychain.

Strict Versus Loose Authentication

By default, strict authentication is enabled and authentication is checked at both ends of each BFD session. Optionally, to smooth migration from nonauthenticated sessions to authenticated sessions, you can configure *loose checking*. When loose checking is configured, packets are accepted without authentication being checked at each end of the session. This feature is intended for transitional periods only.

Configuring BFD Authentication for OSPF

Beginning with Junos OS Release 9.6, you can configure authentication for BFD sessions running over OSPFv2. Routing instances are also supported.

The following sections provide instructions for configuring and viewing BFD authentication on OSPF:

- [Configuring BFD Authentication Parameters on page 4490](#)
- [Viewing Authentication Information for BFD Sessions on page 4492](#)

Configuring BFD Authentication Parameters

Only three steps are needed to configure authentication on a BFD session:

1. Specify the BFD authentication algorithm for the OSPFv2 protocol.
2. Associate the authentication keychain with the OSPFv2 protocol.
3. Configure the related security authentication keychain.

To configure BFD authentication:

1. Specify the algorithm (**keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**, or **simple-password**) to use for BFD authentication on an OSPF route or routing instance.

[edit]

```
user@host# set protocols ospf area 0.0.0.1 interface if2-ospf bfd-liveness-detection
authentication algorithm keyed-sha-1
```



NOTE: Nonstop active routing (NSR) is not supported with meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

2. Specify the keychain to be used to associate BFD sessions on the specified OSPF route or routing instance with the unique security authentication keychain attributes.

This keychain should match the keychain name configured at the **[edit security authentication key-chains]** hierarchy level.

```
[edit]
```

```
user@host# set protocols ospf area 0.0.0.1 interface if2-ospf bfd-liveness-detection
authentication keychain bfd-ospf
```



NOTE: The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

3. Specify the unique security authentication information for BFD sessions:
 - The matching keychain name as specified in Step 2.
 - At least one key, a unique integer between 0 and 63. Creating multiple keys enables multiple clients to use the BFD session.
 - The secret data used to allow access to the session.
 - The time at which the authentication key becomes active, in the format *yyyy-mm-dd.hh:mm:ss*.

```
[edit security]
```

```
user@host# authentication-key-chains key-chain bfd-ospf key 53 secret
$9$ggaJDmPQ6/tJgF/AtREVsyPsnCtUHm start-time 2009-06-14.10:00:00
```

4. (Optional) Specify loose authentication checking if you are transitioning from nonauthenticated sessions to authenticated sessions.

```
[edit]
```

```
user@host> set protocols ospf interface if2-ospf bfd-liveness-detection authentication
loose-check
```

5. (Optional) View your configuration using the **show bfd session detail** or **show bfd session extensive** command.

6. Repeat the steps in this procedure to configure the other end of the BFD session.



NOTE: BFD authentication is only supported in the Canada and United States version of the Junos OS image and is not available in the export version.

Viewing Authentication Information for BFD Sessions

You can view the existing BFD authentication configuration using the **show bfd session detail** and **show bfd session extensive** commands.

The following example shows BFD authentication configured for the **if2-ospf** BGP group. It specifies the keyed SHA-1 authentication algorithm and a keychain name of **bfd-ospf**. The authentication keychain is configured with two keys. Key 1 contains the secret data “\$9\$ggaJDmPQ6/tJgF/AtREvsyPsnCtUHm” and a start time of June 1, 2009, at 9:46:02 AM PST. Key 2 contains the secret data “\$9\$a5jiKW9l.reP38ny.TszF2/9” and a start time of June 1, 2009, at 3:29:20 PM PST.

```
[edit protocols ospf]
area 0.0.0.1 {
  interface if2-ospf {
    bfd-liveness-detection {
      authentication {
        algorithm keyed-sha-1;
        key-chain bfd-ospf;
      }
    }
  }
}
[edit security]
authentication key-chains {
  key-chain bfd-ospf {
    key 1 {
      secret "$9$ggaJDmPQ6/tJgF/AtREvsyPsnCtUHm";
      start-time "2009-6-1.09:46:02 -0700";
    }
    key 2 {
      secret "$9$a5jiKW9l.reP38ny.TszF2/9";
      start-time "2009-6-1.15:29:20 -0700";
    }
  }
}
```

If you commit these updates to your configuration, you see output similar to the following. In the output for the **show bfd session detail** command, **Authenticate** is displayed to indicate that BFD authentication is configured.

show bfd session detail

```
user@host# show bfd session detail
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.9.1.33	Up	so-7/1/0.0	0.600	0.200	3

Client OSPF, TX interval 0.200, RX interval 0.200, multiplier 3, **Authenticate**
 Session up time 3d 00:34
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 Replicated

1 sessions, 1 clients
 Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps

For more information about the configuration, use the **show bfd session extensive** command. The output for this command provides the keychain name, the authentication algorithm and mode for each client in the session, and the overall BFD authentication configuration status, keychain name, and authentication algorithm and mode.

show bfd session extensive

```
user@host# show bfd session extensive
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.9.1.33	Up	so-7/1/0.0	0.600	0.200	3

Client OSPF, TX interval 0.200, RX interval 0.200, multiplier 3, **Authenticate keychain bfd-ospf, algo keyed-md5, mode loose**

Session up time 3d 00:34
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 Replicated
 Min async interval 0.200, min slow interval 1.000
 Adaptive async tx interval 0.200, rx interval 0.200
 Local min tx interval 0.200, min rx interval 0.200, multiplier 3
 Remote min tx interval 0.100, min rx interval 0.100, multiplier 3
 Threshold transmission interval 0.000, Threshold for detection time 0.000
 Local discriminator 11, remote discriminator 80
 Echo mode disabled/inactive
Authentication enabled/active, keychain bfd-ospf, algo keyed-sha-1, mode strict
 1 sessions, 1 clients
 Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps

- Related Documentation**
- [Understanding OSPF Configurations](#)
 - [Understanding BFD for OSPF on page 4482](#)

OSPF Redundancy Features Configuration

- [Examples: Configuring Graceful Restart for OSPF on page 4495](#)

Examples: Configuring Graceful Restart for OSPF

- [Graceful Restart for OSPF Overview on page 4495](#)
- [Example: Configuring Graceful Restart for OSPF on page 4497](#)
- [Example: Configuring the Helper Capability Mode for OSPFv2 Graceful Restart on page 4501](#)
- [Example: Configuring the Helper Capability Mode for OSPFv3 Graceful Restart on page 4504](#)
- [Example: Disabling Strict LSA Checking for OSPF Graceful Restart on page 4507](#)

Graceful Restart for OSPF Overview

Graceful restart allows a routing device undergoing a restart to inform its adjacent neighbors and peers of its condition. During a graceful restart, the restarting device and its neighbors continue forwarding packets without disrupting network performance. Because neighboring devices assist in the restart (these neighbors are called *helper routers*), the restarting device can quickly resume full operation without recalculating algorithms.



NOTE: On a broadcast link with a single neighbor, when the neighbor initiates an OSPFv3 graceful restart operation, the restart might be terminated at the point when the local routing device assumes the role of a helper. A change in the LSA is considered a topology change, which terminates the neighbor's restart operation.

Graceful restart is disabled by default. You can either globally enable graceful restart for all routing protocols, or you can enable graceful restart specifically for OSPF.

This topic describes the following information:

- [Helper Mode for Graceful Restart on page 4496](#)
- [Planned and Unplanned Graceful Restart on page 4496](#)

Helper Mode for Graceful Restart

When a device enabled for OSPF graceful restart restarts, it retains routes learned before the restart in its forwarding table. The device does not allow new OSPF link-state advertisements (LSAs) to update the routing table. This device continues to forward traffic to other OSPF neighbors (or helper routers), and sends only a limited number of LSAs during the restart period. To reestablish OSPF adjacencies with neighbors, the restarting device must send a grace LSA to all neighbors. In response, the helper routers enter helper mode (the ability to assist a neighboring device attempting a graceful restart) and send an acknowledgment back to the restarting device. If there are no topology changes, the helper routers continue to advertise LSAs as if the restarting device had remained in continuous OSPF operation.



NOTE: Helper mode is enabled by default when you start the routing platform, even if graceful restart is not enabled. You can disable helper mode specifically for OSPF.

When the restarting device receives replies from all the helper routers, the restarting device selects routes, updates the forwarding table, and discards the old routes. At this point, full OSPF adjacencies are reestablished and the restarting device receives and processes OSPF LSAs as usual. When the helper routers no longer receive grace LSAs from the restarting device or when the topology of the network changes, the helper routers also resume normal operation.

Beginning with Junos OS Release 11.4, you can configure restart signaling-based helper mode for OSPFv2 graceful restart configurations. The Junos OS implementation is based on RFC 4811, *OSPF Out-of-Band Link State Database (LSDB) Resynchronization*, RFC 4812, *OSPF Restart Signaling*, and RFC 4813, *OSPF Link-Local Signaling*. In restart signaling-based helper mode implementations, the restarting device informs its restart status to its neighbors only after the restart is complete. When the restart is complete, the restarting device sends hello messages to its helper routers with the restart signal (RS) bit set in the hello packet header. When a helper router receives a hello packet with the RS bit set in the header, the helper router returns a hello message to the restarting device. The reply hello message from the helper router contains the ResyncState flag and the ResyncTimeout timer that enable the restarting device to keep track of the helper routers that are syncing up with it. When all helpers complete the synchronization, the restarting device exits the restart mode.



NOTE: Restart signaling-based graceful restart helper mode is not supported for OSPFv3 configurations.

Planned and Unplanned Graceful Restart

OSPF supports two types of graceful restart: planned and unplanned. During a planned restart, the restarting routing device informs the neighbors before restarting. The neighbors act as if the routing device is still within the network topology, and continue forwarding

traffic to the restarting routing device. A grace period is set to specify when the neighbors should consider the restarting routing device as part of the topology. During an unplanned restart, the routing device restarts without warning.

Example: Configuring Graceful Restart for OSPF

This example shows how to configure graceful restart specifically for OSPF.

- [Requirements on page 4497](#)
- [Overview on page 4497](#)
- [Configuration on page 4498](#)
- [Verification on page 4500](#)

Requirements

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 4400.
- Control OSPF designated router election. See “[Example: Controlling OSPF Designated Router Election](#)” on page 4402.
- Configure a single-area OSPF network. See “[Example: Configuring a Single-Area OSPF Network](#)” on page 4406.
- Configure a multiarea OSPF network. See “[Example: Configuring a Multiarea OSPF Network](#)” on page 4408.

Overview

Graceful restart enables a routing device undergoing a restart to inform its adjacent neighbors and peers of its condition. During a graceful restart, the restarting routing device and its neighbors continue forwarding packets without disrupting network performance. By default, graceful restart is disabled. You can globally enable graceful restart for all routing protocols by including the **graceful-restart** statement at the **[edit routing-options]** hierarchy level, or you can enable graceful restart specifically for OSPF by including the **graceful-restart** statement at the **[edit protocols (ospf|ospf3)]** hierarchy level.

The first example shows how to enable graceful restart and configure the optional settings for the grace period interval. In this example, interfaces **fe-1/1/1** and **fe-1/1/2** are in OSPF area 0.0.0.0, and you configure those interfaces for graceful restart. The grace period interval for OSPF graceful restart is determined as equal to or less than the sum of the **notify-duration** time interval and the **restart-duration** time interval. The grace period is the number of seconds that the routing device’s neighbors continue to advertise the routing device as fully adjacent, regardless of the connection state between the routing device and its neighbors.

The **notify-duration** statement configures how long (in seconds) the routing device notifies helper routers that it has completed graceful restart by sending purged grace link-state advertisements (LSAs) over all interfaces. By default, the routing device sends grace LSAs for 30 seconds. The range is from 1 through 3600 seconds.

The **restart-duration** statement configures the amount of time the routing device waits (in seconds) to complete reacquisition of OSPF neighbors from each area. By default, the routing device allows 180 seconds. The range is from 1 through 3600 seconds.

The second example shows how to disable graceful restart for OSPF by including the **disable** statement.

Configuration

- [Enabling Graceful Restart for OSPF on page 4498](#)
- [Disabling Graceful Restart for OSPF on page 4499](#)

Enabling Graceful Restart for OSPF

CLI Quick Configuration To quickly enable graceful restart for OSPF, copy the following commands and paste them into the CLI.

```
[edit]
set interfaces fe-1/1/1 unit 0 family inet address 10.0.0.4
set interfaces fe-1/1/2 unit 0 family inet address 10.0.0.5
set protocols ospf area 0.0.0.0 interface fe-1/1/1
set protocols ospf area 0.0.0.0 interface fe-1/1/2
set routing-options graceful-restart
set protocols ospf graceful-restart restart-duration 190
set protocols ospf graceful-restart notify-duration 40
```

Step-by-Step Procedure To enable graceful restart for OSPF:

1. Configure the interfaces.



NOTE: For OSPFv3, use IPv6 addresses.

```
[edit]
user@host# set interfaces fe-1/1/1 unit 0 family inet address 10.0.0.4
user@host# set interfaces fe-1/1/2 unit 0 family inet address 10.0.0.5
```

2. Configure OSPF on the interfaces.



NOTE: To specify OSPFv3, include the **ospf3** statement at the **[edit protocols]** hierarchy level.

```
[edit]
user@host# set protocols ospf area 0.0.0.0 interface fe-1/1/1
user@host# set protocols ospf area 0.0.0.0 interface fe-1/1/2
```

3. Configure graceful restart globally

```
[edit]
user@host#edit routing-options graceful-restart
```

4. Configure OSPF graceful restart.

```
[edit]
user@host# edit protocols ospf graceful-restart
```

5. (Optional) Configure the restart duration time.

```
[edit protocols ospf graceful-restart]
user@host# set restart-duration 190
```

6. (Optional) Configure the notify duration time.

```
[edit protocols ospf graceful-restart]
user@host# set notify-duration 40
```

7. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf graceful-restart]
user@host# commit
```

Results Confirm your configuration by entering the **show interfaces** and **show protocols ospf** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
fe-1/1/1 {
  unit 0 {
    family inet {
      address 10.0.0.4/32;
    }
  }
}
fe-1/1/2 {
  unit 0 {
    family inet {
      address 10.0.0.5/32;
    }
  }
}
user@host# show protocols ospf
graceful-restart {
  restart-duration 190;
  notify-duration 40;
}
area 0.0.0.0 {
  interface fe-1/1/1.0;
  interface fe-1/1/2.0;
}
```

To confirm an OSPFv3 configuration, enter the **show interfaces** and the **show protocols ospf3** commands.

Disabling Graceful Restart for OSPF

CLI Quick Configuration To quickly disable graceful restart for OSPF, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

```
[edit]
user@host# set protocols ospf graceful-restart disable
```

**Step-by-Step
Procedure**

To disable graceful restart for OSPF:

1. Disable graceful restart for the OSPF protocol only.

This command does not affect the global graceful restart configuration setting.



NOTE: To specify OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# set protocols ospf graceful-restart disable
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Results Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
graceful-restart disable;
```

To confirm an OSPFv3 configuration, enter the `show protocols ospf3` command.

Verification

Confirm that the configuration is working properly.

- [Verifying the OSPF Graceful Restart Configuration on page 4500](#)
- [Verifying Graceful Restart Status on page 4501](#)

Verifying the OSPF Graceful Restart Configuration

Purpose Verify information about your OSPF graceful restart configuration.

Action From operational mode, enter the `show ospf overview` command for OSPFv2. Enter the `show ospf3 overview` command for OSPFv3.

Meaning The Restart field displays the status of graceful restart as either enabled or disabled. The Restart duration field displays how much time the restarted routing device requires to complete reacquisition of OSPF neighbors. The Restart grace period field displays how much time the neighbors should consider the restarted routing device as part of the topology.

Verifying Graceful Restart Status

Purpose	Verify the status of graceful restart.
Action	From operational mode, enter the show route instance detail command.
Meaning	The Restart State field displays Pending if the restart has not been completed or Complete if the restart has finished. The Path selection timeout field indicates the amount of time remaining until graceful restart is declared complete. There is a more detailed Restart State field that displays a list of protocols that have or have not yet completed graceful restart for the specified routing table.

Example: Configuring the Helper Capability Mode for OSPFv2 Graceful Restart

This example shows how to disable and reenabling the helper mode capability for OSPFv2 graceful restart.

- [Requirements on page 4501](#)
- [Overview on page 4501](#)
- [Configuration on page 4502](#)
- [Verification on page 4504](#)

Requirements

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See [“Example: Configuring an OSPF Router Identifier” on page 4400](#).
- Control OSPF designated router election. See [“Example: Controlling OSPF Designated Router Election” on page 4402](#)
- Configure a single-area OSPF network. See [“Example: Configuring a Single-Area OSPF Network” on page 4406](#).
- Configure a multiarea OSPF network. See [“Example: Configuring a Multiarea OSPF Network” on page 4408](#).

Overview

The OSPF graceful restart helper capability assists a neighboring routing device attempting a graceful restart. By default, the helper capability is globally enabled when you start the routing platform. This means that the helper capability is enabled when you start OSPF, even if graceful restart is not globally enabled or specifically enabled for OSPF. You can further modify your graceful restart configuration to disable the helper capability.

Beginning with Junos OS Release 11.4, you can configure restart signaling-based helper mode for OSPFv2 graceful restart configurations. Both the standard and restart signaling-based helper modes are enabled by default.

In the first example, interfaces **fe-1/1/1** and **fe-1/1/2** are in OSPFv2 area 0.0.0.0, and you configure those interfaces for graceful restart. You then disable the standard OSPFv2 graceful restart helper capability by including the **helper-disable standard** statement. This configuration is useful if you have an environment that contains other vendor equipment that is configured for restart signaling-based graceful restart.



NOTE: The **helper-disable** statement and the **no-strict-lsa-checking** statement cannot be configured at the same time. If you attempt to configure both statements at the same time, the routing device displays a warning message when you enter the **show protocols ospf** command.

The second example shows how to reenable the standard OSPFv2 restart helper capability that you disabled in the first example.

Configuration

- [Disabling Helper Mode for OSPFv2 on page 4502](#)
- [Reenabling Helper Mode for OSPFv2 on page 4503](#)

Disabling Helper Mode for OSPFv2

CLI Quick Configuration

To quickly enable graceful restart for OSPFv2 with helper mode disabled, copy the following commands and paste them into the CLI.

```
[edit]
set interfaces fe-1/1/1 unit 0 family inet address 10.0.0.4
set interfaces fe-1/1/2 unit 0 family inet address 10.0.0.5
set protocols ospf area 0.0.0.0 interface fe-1/1/1
set protocols ospf area 0.0.0.0 interface fe-1/1/2
set protocols ospf graceful-restart helper-disable standard
```

Step-by-Step Procedure

To enable graceful restart for OSPFv2 with helper mode disabled:

1. Configure the interfaces.

```
[edit]
user@host# set interfaces fe-1/1/1 unit 0 family inet address 10.0.0.4
user@host# set interfaces fe-1/1/2 unit 0 family inet address 10.0.0.5
```

2. Configure OSPFv2 on the interfaces

```
[edit]
user@host# set protocols ospf area 0.0.0.0 interface fe-1/1/1
user@host# set protocols ospf area 0.0.0.0 interface fe-1/1/2
```

3. Disable the OSPFv2 graceful restart helper capability.
If you disable the OSPFv2 graceful restart helper capability, you cannot disable strict LSA checking.

```
[edit]
user@host# set protocols ospf graceful-restart helper-disable standard
```

4. If you are done configuring the device, commit the configuration.

```
[edit]
```

```
user@host# commit
```

Results Confirm your configuration by entering the **show interfaces** and the **show protocols ospf** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
fe-1/1/1 {
  unit 0 {
    family inet {
      address 10.0.0.4/32;
    }
  }
}
fe-1/1/2 {
  unit 0 {
    family inet {
      address 10.0.0.5/32;
    }
  }
}
user@host# show protocols ospf
graceful-restart {
  helper-disable {
    standard;
  }
}
area 0.0.0.0 {
  interface fe-1/1/1.0;
  interface fe-1/1/2.0;
}
```

Reenabling Helper Mode for OSPFv2

CLI Quick Configuration To quickly reenabling standard helper-mode for OSPFv2, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

```
[edit]
delete protocols ospf graceful-restart helper-disable standard
```



NOTE: To reenabling restart signaling-based helper mode, include the **restart-signaling** statement. To reenabling both standard and restart signaling-based helper mode, include the **both** statement.

Step-by-Step Procedure To reenabling standard helper mode for OSPFv2:

1. Delete the standard helper-mode statement from the OSPFv2 configuration.

```
[edit]
user@host# delete protocols ospf graceful-restart helper-disable standard
```

2. If you are done configuring the device, commit the configuration.

```
[edit]  
user@host# commit
```

Results After you reenables standard helper mode, the **show protocols ospf** command no longer displays the graceful restart configuration.

Verification

Confirm that the configuration is working properly.

- [Verifying the OSPFv2 Graceful Restart Configuration on page 4504](#)
- [Verifying Graceful Restart Status on page 4504](#)

Verifying the OSPFv2 Graceful Restart Configuration

Purpose Verify information about your OSPFv2 graceful restart configuration. The Restart field displays the status of graceful restart as either enabled or disabled, the Graceful restart helper mode field displays the status of the standard helper mode capability as enabled or disabled, and the Restart-signaling helper mode field displays the status of the restart signaling-based helper mode as enabled or disabled. By default, both standard and restart signaling-based helper modes are enabled.

Action From operational mode, enter the **show ospf overview** command.

Verifying Graceful Restart Status

Purpose Verify the status of graceful restart. The Restart State field displays Pending if the restart has not completed, or Complete if the restart has finished. The Path selection timeout field indicates the amount of time remaining until graceful restart is declared complete. There is a more detailed Restart State field that displays a list of protocols that have completed graceful restart or have not yet completed graceful restart for the specified routing table.

Action From operational mode, enter the **show route instance detail** command.

Example: Configuring the Helper Capability Mode for OSPFv3 Graceful Restart

This example shows how to disable and reenables the helper mode capability for OSPFv3 graceful restart.

- [Requirements on page 4504](#)
- [Overview on page 4505](#)
- [Configuration on page 4505](#)
- [Verification on page 4507](#)

Requirements

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See [“Example: Configuring an OSPF Router Identifier” on page 4400](#).
- Control OSPF designated router election. See [“Example: Controlling OSPF Designated Router Election” on page 4402](#).
- Configure a single-area OSPF network. See [“Example: Configuring a Single-Area OSPF Network” on page 4406](#).
- Configure a multiarea OSPF network. See [“Example: Configuring a Multiarea OSPF Network” on page 4408](#).

Overview

The OSPF graceful restart helper capability assists a neighboring routing device attempting a graceful restart. By default, the helper capability is globally enabled when you start the routing platform. This means that the helper capability is enabled when you start OSPF, even if graceful restart is not globally enabled or specifically enabled for OSPF. You can further modify your graceful restart configuration to disable the helper capability.

In the first example, interfaces **fe-1/1/1** and **fe-1/1/2** are in OSPFv3 area 0.0.0.0, and you configure those interfaces for graceful restart. You then disable the OSPFv3 graceful restart helper capability by including the **helper-disable** statement.



NOTE: The **helper-disable** statement and the **no-strict-lsa-checking** statement cannot be configured at the same time. If you attempt to configure both statements at the same time, the routing device displays a warning message when you enter the **show protocols ospf** command.

The second example shows how to reenabling the OSPFv3 restart helper capability that you disabled in the first example.

Configuration

- [Disabling Helper Mode for OSPFv3 on page 4505](#)
- [Reenabling Helper Mode for OSPFv3 on page 4506](#)

Disabling Helper Mode for OSPFv3

CLI Quick Configuration

To quickly enable graceful restart for OSPFv3 with helper mode disabled, copy the following commands and paste them into the CLI.

```
[edit]
set interfaces fe-1/1/1 unit 0 family inet6 address 2002:0a00:0004::
set interfaces fe-1/1/2 unit 0 family inet6 address 2002:0a00:0005::
set protocols ospf3 area 0.0.0.0 interface fe-1/1/1
set protocols ospf3 area 0.0.0.0 interface fe-1/1/2
set protocols ospf3 graceful-restart helper-disable
```

Step-by-Step Procedure

To enable graceful restart for OSPFv3 with helper mode disabled:

1. Configure the interfaces.

```
[edit]
user@host# set interfaces fe-1/1/1 unit 0 family inet6 address 2002:0a00:0004::
user@host# set interfaces fe-1/1/1 unit 0 family inet address 2002:0a00:0005::
```

2. Configure OSPFv3 on the interfaces

```
[edit]
user@host# set protocols ospf3 area 0.0.0.0 interface fe-1/1/1
user@host# set protocols ospf3 area 0.0.0.0 interface fe-1/1/2
```

3. Disable the OSPFv3 graceful restart helper capability.
If you disable the OSPFv3 graceful restart helper capability, you cannot disable strict LSA checking.

```
[edit]
user@host# set protocols ospf3 graceful-restart helper-disable
```

4. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Results Confirm your configuration by entering the **show interfaces** and the **show protocols ospf3** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
fe-1/1/1 {
  unit 0 {
    family inet6 {
      address 2002:0a00:0004::/128;
    }
  }
}
fe-1/1/2 {
  unit 0 {
    family inet6 {
      address 2002:0a00:0005::/128;
    }
  }
}
user@host# show protocols ospf3
graceful-restart {
  helper-disable;
}
area 0.0.0.0 {
  interface fe-1/1/1.0;
  interface fe-1/1/2.0;
}
```

Reenabling Helper Mode for OSPFv3

CLI Quick Configuration To quickly reenable helper-mode for OSPFv3, copy the following command and paste it into the CLI.

```
[edit]
delete protocols ospf3 graceful-restart helper-disable
```

Step-by-Step Procedure

To reenable helper mode for OSPFv3:

1. Delete the standard helper-mode statement from the OSPFv3 configuration.

```
[edit]
user@host# delete protocols ospf3 graceful-restart helper-disable
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Results After you reenable standard helper mode, the **show protocols ospfs** command no longer displays the graceful restart configuration.

Verification

Confirm that the configuration is working properly.

- [Verifying the OSPFv3 Graceful Restart Configuration on page 4507](#)
- [Verifying Graceful Restart Status on page 4507](#)

Verifying the OSPFv3 Graceful Restart Configuration

Purpose Verify information about your OSPFv3 graceful restart configuration. The Restart field displays the status of graceful restart as either enabled or disabled, and the Helper mode field displays the status of the helper mode capability as either enabled or disabled.

Action From operational mode, enter the **show ospf3 overview** command.

Verifying Graceful Restart Status

Purpose Verify the status of graceful restart. The Restart State field displays Pending if the restart has not completed, or Complete if the restart has finished. The Path selection timeout field indicates the amount of time remaining until graceful restart is declared complete. There is a more detailed Restart State field that displays a list of protocols that have completed graceful restart or have not yet completed graceful restart for the specified routing table.

Action From operational mode, enter the **show route instance detail** command.

Example: Disabling Strict LSA Checking for OSPF Graceful Restart

This example shows how to disable strict link-state advertisement (LSA) checking for OSPF graceful restart.

- [Requirements on page 4508](#)
- [Overview on page 4508](#)
- [Configuration on page 4508](#)
- [Verification on page 4510](#)

Requirements

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See [“Example: Configuring an OSPF Router Identifier” on page 4400](#).
- Control OSPF designated router election. See [“Example: Controlling OSPF Designated Router Election” on page 4402](#).
- Configure a single-area OSPF network. See [“Example: Configuring a Single-Area OSPF Network” on page 4406](#).
- Configure a multiarea OSPF network. See [“Example: Configuring a Multiarea OSPF Network” on page 4408](#).

Overview

You can disable strict LSA checking to prevent the termination of graceful restart by a helping router. You might configure this option for interoperability with other vendor devices. The OSPF graceful restart helper capability must be enabled if you disable strict LSA checking. By default, LSA checking is enabled.

In this example, interfaces **fe-1/1/1** and **fe-1/1/2** are in OSPF area 0.0.0.0, and you configure those interfaces for graceful restart. You then disable strict LSA checking by including the **no-strict-lsa-checking** statement.



NOTE: The **helper-disable** statement and the **no-strict-lsa-checking** statement cannot be configured at the same time. If you attempt to configure both statements at the same time, the routing device displays a warning message when you enter the **show protocols ospf** command.

Configuration

CLI Quick Configuration

To quickly enable graceful restart for OSPF with strict LSA checking disabled, copy the following commands and paste them into the CLI.

```
[edit]
set interfaces fe-1/1/1 unit 0 family inet address 10.0.0.4
set interfaces fe-1/1/2 unit 0 family inet address 10.0.0.5
set protocols ospf area 0.0.0.0 interface fe-1/1/1
set protocols ospf area 0.0.0.0 interface fe-1/1/2
set protocols ospf graceful-restart no-strict-lsa-checking
```

Step-by-Step Procedure

To enable graceful restart for OSPF with strict LSA checking disabled:

1. Configure the interfaces.



NOTE: For OSPFv3, use IPv6 addresses.

```
[edit]
user@host# set interfaces fe-1/1/1 unit 0 family inet address 10.0.0.4
user@host# set interfaces fe-1/1/1 unit 0 family inet address 10.0.0.5
```

2. Configure OSPF on the interfaces



NOTE: To specify OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# set protocols ospf area 0.0.0.0 interface fe-1/1/1
user@host# set protocols ospf area 0.0.0.0 interface fe-1/1/2
```

3. Disable strict LSA checking.
If you disable the strict LSA checking, OSPF graceful restart helper capability must be enabled (which is the default behavior).

```
[edit]
user@host# set protocols ospf graceful-restart no-strict-lsa-checking
```

4. If you are done configuring the device, commit the configuration.

```
[edit ]
user@host# commit
```

Results Confirm your configuration by entering the `show interfaces` and the `show protocols ospf` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
fe-1/1/1 {
  unit 0 {
    family inet {
      address 10.0.0.4/32;
    }
  }
}
fe-1/1/2 {
  unit 0 {
    family inet {
      address 10.0.0.5/32;
    }
  }
}
user@host# show protocols ospf
graceful-restart {
  no-strict-lsa-checking;
}
area 0.0.0.0 {
  interface fe-1/1/1.0;
  interface fe-1/1/2.0;
}
```

To confirm your OSPFv3 configuration, enter the **show interfaces** and the **show protocols ospf3** commands.

Verification

Confirm that the configuration is working properly.

- [Verifying the OSPF Graceful Restart Configuration on page 4510](#)
- [Verifying Graceful Restart Status on page 4510](#)

Verifying the OSPF Graceful Restart Configuration

Purpose Verify information about your OSPF graceful restart configuration. The Restart field displays the status of graceful restart as either enabled or disabled.

Action From operational mode, enter the **show ospf overview** command for OSPFv2, and enter the **show ospf3 overview** command for OSPFv3.

Verifying Graceful Restart Status

Purpose Verify the status of graceful restart. The Restart State field displays Pending if the restart has not completed, or Complete if the restart has finished. The Path selection timeout field indicates the amount of time remaining until graceful restart is declared complete. There is a more detailed Restart State field that displays a list of protocols that have completed graceful restart or have not yet completed graceful restart for the specified routing table.

Action From operational mode, enter the **show route instance detail** command.

Related Documentation

- [OSPF Overview on page 4386](#)
- [Understanding OSPF Configurations](#)
- [Graceful Restart Concepts on page 2579](#) in the *Junos OS High Availability Library for Routing Devices*

OSPF Traffic Engineering Configuration

- [Examples: Configuring OSPF Traffic Engineering on page 4511](#)
- [Example: Configuring OSPF Passive Traffic Engineering Mode on page 4519](#)

Examples: Configuring OSPF Traffic Engineering

- [OSPF Support for Traffic Engineering on page 4511](#)
- [Example: Enabling OSPF Traffic Engineering Support on page 4513](#)
- [Example: Configuring the Traffic Engineering Metric for a Specific OSPF Interface on page 4518](#)

OSPF Support for Traffic Engineering

Traffic engineering allows you to control the path that data packets follow, bypassing the standard routing model, which uses routing tables. Traffic engineering moves flows from congested links to alternate links that would not be selected by the automatically computed destination-based shortest path.

To help provide traffic engineering and MPLS with information about network topology and loading, extensions have been added to the Junos OS implementation of OSPF. When traffic engineering is enabled on the routing device, you can enable OSPF traffic engineering support. When you enable traffic engineering for OSPF, the shortest-path-first (SPF) algorithm takes into account the various label-switched paths (LSPs) configured under MPLS and configures OSPF to generate opaque link-state advertisements (LSAs) that carry traffic engineering parameters. The parameters are used to populate the traffic engineering database. The traffic engineering database is used exclusively for calculating explicit paths for the placement of LSPs across the physical topology. The Constrained Shortest Path First (CSPF) algorithm uses the traffic engineering database to compute the paths that MPLS LSPs take. RSVP uses this path information to set up LSPs and to reserve bandwidth for them.

By default, traffic engineering support is disabled. To enable traffic engineering, include the **traffic-engineering** statement. You can also configure the following OSPF traffic engineering extensions:

- **advertise-unnumbered-interfaces**—(OSPFv2 only) Advertises the link-local identifier in the link-local traffic engineering LSA packet. You do not need to include this statement if RSVP is able to signal unnumbered interfaces as defined in RFC 3477, *Signalling Unnumbered Links in Resource Reservation Protocol - Traffic Engineering (RSVP-TE)*.
- **credibility-protocol-preference**—(OSPFv2 only) Assigns a credibility value to OSPF routes in the traffic engineering database. By default, Junos OS prefers IS-IS routes in the traffic engineering database over other interior gateway protocol (IGP) routes even if the routes of another IGP are configured with a lower, that is, more preferred, preference value. The traffic engineering database assigns a credibility value to each IGP and prefers the routes of the IGP with the highest credibility value. In Junos OS Release 9.4 and later, you can configure OSPF to take protocol preference into account to determine the traffic engineering database credibility value. When protocol preference is used to determine the credibility value, IS-IS routes are not automatically preferred by the traffic engineering database, depending on your configuration.
- **ignore-lsp-metrics**—Ignores RSVP LSP metrics in OSPF traffic engineering shortcut calculations or when you configure LDP over RSVP LSPs. This option avoids mutual dependency between OSPF and RSVP, eliminating the time period when the RSVP metric used for tunneling traffic is not up to date. In addition, If you are using RSVP for traffic engineering, you can run LDP simultaneously to eliminate the distribution of external routes in the core. The LSPs established by LDP are tunneled through the LSPs established by RSVP. LDP effectively treats the traffic-engineered LSPs as single hops.
- **multicast-rpf-routes**—(OSPFv2 only) Installs unicast IPv4 routes (not LSPs) in the multicast routing table (**inet.2**) for multicast reverse-path forwarding (RPF) checks. The **inet.2** routing table consists of unicast routes used for multicast RPF lookup. RPF is an antispoofing mechanism used to check if the packet is coming in on an interface that is also sending data back to the packet source.
- **no-topology**—(OSPFv2 only) To disable the dissemination of link-state topology information. If disabled, traffic engineering topology information is no longer distributed within the OSPF area.
- **shortcuts**—Configures IGP shortcuts, which allows OSPF to use an LSP as the next hop as if it were a logical interface from the ingress routing device to the egress routing device. The address specified in the **to** statement at the **[edit protocols mpls label-switched-path lsp-path-name]** hierarchy level on the ingress routing device must match the router ID of the egress routing device for the LSP to function as a direct link to the egress routing device and to be used as input to the OSPF SPF calculations. When used in this way, LSPs are no different from Asynchronous Transfer Mode (ATM) and Frame Relay virtual circuits (VCs), except that LSPs carry only IPv4 traffic.

OSPFv2 installs the prefix for IPv4 routes in the **inet.0** routing table, and the LSPs are installed by default in the **inet.3** routing table.

OSPFv3 LSPs used for shortcuts continue to be signaled using IPv4. However, by default, shortcut IPv6 routes calculated through OSPFv3 are added to the **inet6.3** routing table. The default behavior is for BGP only to use LSPs in its calculations. If you configure MPLS so that both BGP and IGP use LSPs for forwarding traffic, IPv6 shortcut routes calculated through OSPFv3 are added to the **inet6.0** routing table.



NOTE: Whenever possible, use OSPF IGP shortcuts instead of traffic engineering shortcuts.

- **lsp-metric-info-summary**—Advertises the LSP metric in summary LSAs to treat the LSP as a link. This configuration allows other routing devices in the network to use this LSP. To accomplish this, you need to configure MPLS and OSPF traffic engineering to advertise the LSP metric in summary LSAs.

When you enable traffic engineering on the routing device, you can also configure an OSPF metric that is used exclusively for traffic engineering. The traffic engineering metric is used for information injected into the traffic engineering database. Its value does not affect normal OSPF forwarding.

Example: Enabling OSPF Traffic Engineering Support

This example shows how to enable OSPF traffic engineering support to advertise the label-switched path (LSP) metric in summary link-state advertisements (LSAs).

- [Requirements on page 4513](#)
- [Overview on page 4513](#)
- [Configuration on page 4514](#)
- [Verification on page 4517](#)

Requirements

Before you begin:

- Configure the device interfaces. See the *Interfaces Feature Guide for Security Devices*.
- Configure BGP per your network requirements. See the *BGP Feature Guide for Routing Devices*.
- Configure MPLS per your network requirements. See the *MPLS Feature Guide for Security Devices*.

Overview

You can configure OSPF to treat an LSP as a link and have other routing devices in the network use this LSP. To accomplish this, you configure MPLS and OSPF traffic engineering to advertise the LSP metric in summary LSAs.

In this example, there are four routing devices in area 0.0.0.0, and you want OSPF to treat the LSP named R1-to-R4 that goes from the ingress Device R1 to the egress Device R4 as a link.

For OSPF, you enable traffic engineering on all four routing devices in the area by including the **traffic-engineering** statement. This configuration ensures that the shortest-path-first (SPF) algorithm takes into account the LSPs configured under MPLS and configures OSPF to generate LSAs that carry traffic engineering parameters. You further ensure that OSPF uses the MPLS LSP as the next hop and advertises the LSP metric in summary LSAs, by including the optional **shortcuts lsp-metric-into-summary** statement on the ingress Device R1.

For MPLS, you enable traffic engineering so that MPLS performs traffic engineering on both BGP and IGP destinations by including the **traffic-engineering bgp-igp** statement, and you include the LSP named R1-to-R4 by including the **label-switched-path lsp-path-name to address** statement on the ingress Device R1. The address specified in the **to** statement on the ingress Device R1 must match the router ID of the egress Device R4 for the LSP to function as a direct link to the egress routing device and to be used as input to the OSPF SPF calculations. In this example, the router ID of the egress Device R4 is 10.0.0.4.

Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Modifying the Junos OS Configuration* in the *CLI User Guide*.

CLI Quick Configuration To quickly enable OSPF traffic engineering support to advertise the LSP metric in summary LSAs, copy the following commands and paste them into the CLI.

Configuration on R1:

```
[edit]
set routing-options router-id 10.0.0.1
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf traffic-engineering shortcuts lsp-metric-into-summary
set protocols mpls traffic-engineering bgp-igp
set protocols mpls label-switched-path R1-to-R4 to 10.0.0.4
```

Configuration on R2:

```
[edit]
set routing-options router-id 10.0.0.2
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf traffic-engineering
```

Configuration on R3:

```
[edit]
set routing-options router-id 10.0.0.3
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf traffic-engineering
```

Configuration on R4:

```
[edit]
set routing-options router-id 10.0.0.4
```

```
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf traffic-engineering
```

Step-by-Step Procedure To enable OSPF traffic engineering support to advertise LSP metrics in summary LSAs:

1. Configure the router ID.

```
[edit]
user@R1# set routing-options router-id 10.0.0.1
```

```
[edit]
user@R2# set routing-options router-id 10.0.0.2
```

```
[edit]
user@R3# set routing-options router-id 10.0.0.3
```

```
[edit]
user@R4# set routing-options router-id 10.0.0.4
```

2. Configure the OSPF area and add the interfaces.



NOTE: To specify OSPFv3, include the `ospfv3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@R1# set protocols ospf area 0.0.0.0 interface all
user@R1# set protocols ospf area 0.0.0.0 interface fxp0.0 disable
```

```
[edit]
user@R2# set protocols ospf area 0.0.0.0 interface all
user@R2# set protocols ospf area 0.0.0.0 interface fxp0.0 disable
```

```
[edit]
user@R3# set protocols ospf area 0.0.0.0 interface all
user@R3# set protocols ospf area 0.0.0.0 interface fxp0.0 disable
```

```
[edit]
user@R4# set protocols ospf area 0.0.0.0 interface all
user@R4# set protocols ospf area 0.0.0.0 interface fxp0.0 disable
```

3. Enable OSPF traffic engineering.

```
[edit]
user@R1# set protocols ospf traffic-engineering shortcuts lsp-metric-into-summary
```

```
[edit]
user@R2# set protocols ospf traffic-engineering
```

```
[edit]
user@R3# set protocols ospf traffic-engineering
```

```
[edit]
user@R4# set protocols ospf traffic-engineering
```

4. On Device R1, configure MPLS traffic engineering.

```
[edit ]
user@R1# set protocols mpls traffic-engineering bgp-igp
```

```
user@R1# set protocols mpls label-switched-path R1-to-R4 to 10.0.0.4
```

5. If you are done configuring the devices, commit the configuration.

```
[edit]
user@host# commit
```

Results Confirm your configuration by entering the **show routing-options**, **show protocols ospf**, and **show protocols mpls** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Output for R1:

```
user@host# show routing-options
router-id 10.0.0.1;

user@host# show protocols ospf
  traffic-engineering {
    shortcuts lsp-metric-into-summary;
  }
  area 0.0.0.0 {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }

user@host# show protocols mpls
  traffic-engineering bgp-igp;
  label-switched-path R1-to-R4 {
    to 10.0.0.4;
  }
```

Output for R2:

```
user@host# show routing-options
router-id 10.0.0.2;

user@host# show protocols ospf
  traffic-engineering;
  area 0.0.0.0 {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
```

Output for R3:

```
user@host# show routing-options
router-id 10.0.0.3;

user@host# show protocols ospf
  traffic-engineering;
  area 0.0.0.0 {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
```

```
}

```

Output for R4:

```
user@host# show routing-options
router-id 10.0.0.4;

user@host# show protocols ospf
  traffic-engineering;
  area 0.0.0.0 {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
}
```

To confirm your OSPFv3 configuration, enter the **show routing-options**, **show protocols ospf3**, and **show protocols mpls** commands.

Verification

Confirm that the configuration is working properly.

- [Verifying the Traffic Engineering Capability for OSPF on page 4517](#)
- [Verifying OSPF Entries in the Traffic Engineering Database on page 4517](#)
- [Verifying That the Traffic Engineering Database Is Learning Node Information from OSPF on page 4517](#)

Verifying the Traffic Engineering Capability for OSPF

Purpose Verify that traffic engineering has been enabled for OSPF. By default, traffic engineering is disabled.

Action From operational mode, enter the **show ospf overview** command for OSPFv2, and enter the **show ospf3 overview** for OSPFv3.

Verifying OSPF Entries in the Traffic Engineering Database

Purpose Verify the OSPF information in the traffic engineering database. The Protocol field displays OSPF and the area from which the information was learned.

Action From operational mode, enter the **show ted database** command.

Verifying That the Traffic Engineering Database Is Learning Node Information from OSPF

Purpose Verify that OSPF is reporting node information. The Protocol name field displays OSPF and the area from which the information was learned.

Action From operational mode, enter the **show ted protocol** command.

Example: Configuring the Traffic Engineering Metric for a Specific OSPF Interface

This example shows how to configure the OSPF metric value used for traffic engineering.

- [Requirements on page 4518](#)
- [Overview on page 4518](#)
- [Configuration on page 4518](#)
- [Verification on page 4519](#)

Requirements

Before you begin:

- Configure the device interfaces. See the *Interfaces Feature Guide for Security Devices*.
- Configure OSPF for traffic engineering. See “[Example: Enabling OSPF Traffic Engineering Support](#)” on page 4513

Overview

You can configure an OSPF metric that is used exclusively for traffic engineering. To modify the default value of the traffic engineering metric, include the **te-metric** statement. The OSPF traffic engineering metric does not affect normal OSPF forwarding. By default, the traffic engineering metric is the same value as the OSPF metric. The range is 1 through 65,535.

In this example, you configure the OSPF traffic engineering metric on OSPF interface **fe-0/1/1** in area 0.0.0.0.

Configuration

CLI Quick Configuration

To quickly configure the OSPF traffic engineering metric for a specific interface, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

```
[edit]
set protocols ospf area 0.0.0.0 interface fe-0/1/1 te-metric 10
```

Step-by-Step Procedure

To configure an OSPF traffic engineering metric for a specific interface used only for traffic engineering:

1. Create an OSPF area.



NOTE: To specify OSPFv3, include the **ospf3** statement at the [edit protocols] hierarchy level.

```
[edit]
user@host# edit protocols ospf area 0.0.0.0
```

2. Configure the traffic engineering metric of the OSPF network segments.

```
[edit protocols ospf area 0.0.0.0]
user@host set interface fe-0/1/1 te-metric 10
```

3. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.0]
user@host# commit
```

Results Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.0 {
  interface fe-0/1/1.0 {
    te-metric 10;
  }
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

Verification

Confirm that the configuration is working properly.

Verifying the Configured Traffic Engineering Metric

Purpose Verify the traffic engineering metric value. Confirm that Metric field displays the configured traffic engineering metric.

Action From operational mode, enter the **show ted database extensive** command.

Related Documentation

- [Understanding OSPF Configurations](#)
- [Junos OS MPLS Applications Library for Routing Devices](#)

Example: Configuring OSPF Passive Traffic Engineering Mode

- [OSPF Passive Traffic Engineering Mode on page 4519](#)
- [Example: Configuring OSPF Passive Traffic Engineering Mode on page 4520](#)

OSPF Passive Traffic Engineering Mode

Ordinarily, interior routing protocols such as OSPF are not run on links between autonomous systems. However, for inter-AS traffic engineering to function properly, information about the inter-AS link—in particular, the address on the remote interface—must be made available inside the autonomous system (AS). This information is not normally included either in the external BGP (EBGP) reachability messages or in the OSPF routing advertisements.

To flood this link address information within the AS and make it available for traffic engineering calculations, you must configure OSPF passive mode for traffic engineering on each inter-AS interface. You must also supply the remote address for OSPF to distribute and include it in the traffic engineering database. OSPF traffic engineering mode allows MPLS label-switched paths (LSPs) to dynamically discover OSPF AS boundary routers and to allow routers to establish a traffic engineering LSP across multiple autonomous systems.

Example: Configuring OSPF Passive Traffic Engineering Mode

This example shows how to configure OSPF passive mode for traffic engineering on an inter-AS interface. The AS boundary router link between the EBGP peers must be a directly connected link and must be configured as a passive traffic engineering link.

- [Requirements on page 4520](#)
- [Overview on page 4520](#)
- [Configuration on page 4521](#)
- [Verification on page 4522](#)

Requirements

Before you begin:

- Configure the device interfaces. See the *Interfaces Feature Guide for Security Devices*.
- Configure BGP per your network requirements. See the *BGP Feature Guide for Routing Devices*.
- Configure the LSP per your network requirements. See the *MPLS Feature Guide for Security Devices*.
- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 4400.
- Control OSPF designated router election. See “[Example: Controlling OSPF Designated Router Election](#)” on page 4402.
- Configure a single-area OSPF network. See “[Example: Configuring a Single-Area OSPF Network](#)” on page 4406.
- Configure a multiarea OSPF network. See “[Example: Configuring a Multiarea OSPF Network](#)” on page 4408.

Overview

You can configure OSPF passive mode for traffic engineering on an inter-AS interface. The address used for the remote node of the OSPF passive traffic engineering link must be the same as the address used for the EBGP link. In this example, you configure interface **so-1/1/0** in area 0.0.0.1 as the inter-AS link to distribute traffic engineering information with OSPF within the AS and include the following settings:

- **passive**—Advertises the direct interface addresses on an interface without actually running OSPF on that interface. A passive interface is one for which the address

information is advertised as an internal route in OSPF, but on which the protocol does not run.

- **traffic-engineering**—Configures an interface in OSPF passive traffic-engineering mode to enable dynamic discovery of OSPF AS boundary routers. By default, OSPF passive traffic-engineering mode is disabled.
- **remote-node-id**—Specifies the IP address at the far end of the inter-AS link. In this example, the remote IP address is 192.168.207.2.

Configuration

To quickly configure OSPF passive mode for traffic engineering, copy the following command, remove any line breaks, and paste it into the CLI.

```
[edit]
set protocols ospf area 0.0.0.1 interface so-1/1/0 passive traffic-engineering remote-node-id
192.168.207.2
```

Step-by-Step Procedure

To configure OSPF passive traffic engineering mode:

1. Create an OSPF area.



NOTE: To specify OSPFv3, include the **ospf3** statement at the **[edit protocols]** hierarchy level.

```
[edit]
user@host# set protocols ospf area 0.0.0.1
```

2. Configure interface **so-1/1/0** as a passive interface configured for traffic engineering, and specify the IP address at the far end of the inter-AS link.

```
[edit protocols ospf area 0.0.0.1]
user@host# set interface so-1/1/0 passive traffic-engineering remote-node-id
192.168.207.2
```

3. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf]
user@host# commit
```

Results Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.1 {
  interface so-1/1/0.0 {
    passive {
      traffic-engineering {
        remote-node-id 192.168.207.2;
      }
    }
  }
}
```

```
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

Verification

Confirm that the configuration is working properly.

Verifying the Status of OSPF Interfaces

Purpose Verify the status of OSPF interfaces. If the interface is passive, the Adj count field is 0 because no adjacencies have been formed. Next to this field, you might also see the word Passive.

Action From operational mode, enter the **show ospf interface detail** command for OSPFv2, and enter the **show ospf3 interface detail** command for OSPFv3.

Related Documentation

- *Understanding OSPF Configurations*
- [Understanding OSPF Interfaces on page 4431](#)
- *Junos OS MPLS Applications Library for Routing Devices*

OSPF Database Protection Configuration

- [Example: Configuring OSPF Database Protection on page 4523](#)

Example: Configuring OSPF Database Protection

- [OSPF Database Protection Overview on page 4523](#)
- [Configuring OSPF Database Protection on page 4524](#)

OSPF Database Protection Overview

OSPF database protection allows you to limit the number of link-state advertisements (LSAs) not generated by the local router in a given OSPF routing instance, helping to protect the link-state database from being flooded with excessive LSAs. This feature is particularly useful if VPN routing and forwarding is configured on your provider edge and customer edge routers using OSPF as the routing protocol. An overrun link-state database on the customer edge router can exhaust resources on the provider edge router and impact the rest of the service provider network.

When you enable OSPF database protection, the maximum number of LSAs you specify includes all LSAs whose advertising router ID is not equal to the local router ID (nonself-generated LSAs). These might include external LSAs as well as LSAs with any scope such as the link, area, and autonomous system (AS).

Once the specified maximum LSA count is exceeded, the database typically enters into the ignore state. In this state, all neighbors are brought down, and nonself-generated LSAs are destroyed. In addition, the database sends out hellos but ignores all received packets. As a result, the database does not form any full neighbors, and therefore does not learn about new LSAs. However, if you have configured the **warning-only** option, only a warning is issued and the database does not enter the ignore state but continues to operate as before.

You can also configure one or more of the following options:

- A warning threshold for issuing a warning message before the LSA limit is reached.
- An ignore state time during which the database must remain in the ignore state and after which normal operations can be resumed.
- An ignore state count that limits the number of times the database can enter the ignore state, after which it must enter the isolate state. The isolate state is very similar to the

ignore state, but has one important difference: once the database enters the isolate state, it must remain there until you issue a command to clear database protection before it can return to normal operations.

- A reset time during which the database must stay out of the ignore or isolate state before it is returned to a normal operating state.

Configuring OSPF Database Protection

By configuring OSPF database protection, you can help prevent your OSPF link-state database from being overrun with excessive LSAs that are not generated by the local router. You specify the maximum number of LSAs whose advertising router ID is not the same as the local router ID in an OSPF instance. This feature is particularly useful if your provider edge and customer edge routers are configured with VPN routing and forwarding using OSPF.

OSPF database protection is supported on:

- Logical systems
- All routing instances supported by OSPFv2 and OSPFv3
- OSPFv2 and OSPFv3 topologies
- OSPFv3 realms

To configure OSPF database protection:

1. Include the **database-protection** statement at one of the following hierarchy levels:
 - [edit protocols ospf | ospf3]
 - [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (ospf |ospf3)]
 - [edit routing-instances *routing-instance-name* protocols (ospf |ospf3)]
 - [edit routing-instances *routing-instance-name* protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast)]
2. Include the **maximum-lsa *number*** statement.



NOTE: The **maximum-lsa** statement is mandatory, and there is no default value for it. If you omit this statement, you cannot configure OSPF database protection.

3. (Optional) Include the following statements:
 - **ignore-count *number***—Specify the number of times the database can enter the ignore state before it goes into the isolate state.
 - **ignore-time *seconds***—Specify the time limit the database must remain in the ignore state before it resumes regular operations.

- **reset-time *seconds***—Specify the time during which the database must operate without being in either the ignore or isolate state before it is reset to a normal operating state.
 - **warning-threshold *percent***—Specify the percent of the maximum LSA number that must be exceeded before a warning message is issued.
4. (Optional) Include the **warning-only** statement to prevent the database from entering the ignore state or isolate state when the maximum LSA count is exceeded.



NOTE: If you include the **warning-only** statement, values for the other optional statements at the same hierarchy level are not used when the maximum LSA number is exceeded.

5. Verify your configuration by checking the database protection fields in the output of the **show ospf overview** command.

**Related
Documentation**

- [OSPF Overview on page 4386](#)
- [Understanding OSPF Configurations](#)

OSPF Policy Configuration

- [Examples: Configuring OSPF Routing Policy on page 4527](#)
- [Examples: Configuring Routing Policy for Network Summaries on page 4543](#)

Examples: Configuring OSPF Routing Policy

- [Understanding OSPF Routing Policy on page 4527](#)
- [Example: Injecting OSPF Routes into the BGP Routing Table on page 4529](#)
- [Example: Redistributing Static Routes into OSPF on page 4532](#)
- [Example: Configuring an OSPF Import Policy on page 4535](#)
- [Example: Configuring a Route Filter Policy to Specify Priority for Prefixes Learned Through OSPF on page 4539](#)

Understanding OSPF Routing Policy

Each routing policy is identified by a policy name. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose the entire name in double quotation marks. Each routing policy name must be unique within a configuration. Once a policy is created and named, it must be applied before it is active.

In the **import** statement, you list the name of the routing policy used to filter OSPF external routes from being installed into the routing tables of OSPF neighbors. You can filter the routes, but not link-state address (LSA) flooding. An external route is a route that is outside the OSPF Autonomous System (AS). The import policy does not impact the OSPF database. This means that the import policy has no impact on the link-state advertisements.

In the **export** statement, you list the name of the routing policy to be evaluated when routes are being exported from the routing table into OSPF.

By default, if a routing device has multiple OSPF areas, learned routes from other areas are automatically installed into area 0 of the routing table.

To specify more than one policy and create a policy chain, you list the policies using a space as a separator. If multiple policies are specified, the policies are evaluated in the order in which they are specified. As soon as an accept or reject action is executed, the policy chain evaluation ends.

This topic describes the following information:

- [Routing Policy Terms on page 4528](#)
- [Routing Policy Match Conditions on page 4528](#)
- [Routing Policy Actions on page 4529](#)

Routing Policy Terms

Routing policies are made up of one or more terms. A term is a named structure in which match conditions and actions are defined. You can define one or more terms. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose the entire name in double quotation marks.

Each term contains a set of match conditions and a set of actions:

- Match conditions are criteria that a route must match before the actions can be applied. If a route matches all criteria, one or more actions are applied to the route.
- Actions specify whether to accept or reject the route, control how a series of policies are evaluated, and manipulate the characteristics associated with a route.

Routing Policy Match Conditions

A match condition defines the criteria that a route must match for an action to take place. You can define one or more match conditions for each term. If a route matches all of the match conditions for a particular term, the actions defined for that term are processed.

Each term can include two statements, **from** and **to**, that define the match conditions:

- In the **from** statement, you define the criteria that an incoming route must match. You can specify one or more match conditions. If you specify more than one, they all must match the route for a match to occur.

The **from** statement is optional. If you omit the **from** and the **to** statements, all routes are considered to match.



NOTE: In export policies, omitting the **from** statement from a routing policy term might lead to unexpected results.

- In the **to** statement, you define the criteria that an outgoing route must match. You can specify one or more match conditions. If you specify more than one, they all must match the route for a match to occur.

The order of the match conditions in a term is not important because a route must match all match conditions in a term for an action to be taken.

For a complete list of match conditions, see [Configuring Match Conditions in Routing Policy Terms](#).

Routing Policy Actions

An action defines what the routing device does with the route when the route matches all the match conditions in the **from** and **to** statements for a particular term. If a term does not have **from** and **to** statements, all routes are considered to match and the actions apply to all routes.

Each term can have one or more of the following types of actions. The actions are configured under the **then** statement.

- Flow control actions, which affect whether to accept or reject the route and whether to evaluate the next term or routing policy.
- Actions that manipulate route characteristics.
- Trace action, which logs route matches.

The **then** statement is optional. If you omit it, one of the following occurs:

- The next term in the routing policy, if one exists, is evaluated.
- If the routing policy has no more terms, the next routing policy, if one exists, is evaluated.
- If there are no more terms or routing policies, the **accept** or **reject** action specified by the default policy is executed.

For a complete list of routing policy actions, see *Configuring Actions in Routing Policy Terms*.

Example: Injecting OSPF Routes into the BGP Routing Table

This example shows how to create a policy that injects OSPF routes into the BGP routing table.

- [Requirements on page 4529](#)
- [Overview on page 4530](#)
- [Configuration on page 4530](#)
- [Verification on page 4532](#)
- [Troubleshooting on page 4532](#)

Requirements

Before you begin:

- Configure network interfaces.
- Configure external peer sessions. See [“Example: Configuring External BGP Point-to-Point Peer Sessions” on page 3602](#).
- Configure interior gateway protocol (IGP) sessions between peers.

Overview

In this example, you create a routing policy called **injectpolicy1** and a routing term called **injectterm1**. The policy injects OSPF routes into the BGP routing table.

Configuration

- [Configuring the Routing Policy on page 4530](#)
- [Configuring Tracing for the Routing Policy on page 4531](#)

Configuring the Routing Policy

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

```
set policy-options policy-statement injectpolicy1 term injectterm1 from protocol ospf
set policy-options policy-statement injectpolicy1 term injectterm1 from area 0.0.0.1
set policy-options policy-statement injectpolicy1 term injectterm1 then accept
set protocols bgp export injectpolicy1
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To inject OSPF routes into a BGP routing table:

1. Create the policy term.

```
[edit policy-options policy-statement injectpolicy1]
user@host# set term injectterm1
```

2. Specify OSPF as a match condition.

```
[edit policy-options policy-statement injectpolicy1 term injectterm1]
user@host# set from protocol ospf
```

3. Specify the routes from an OSPF area as a match condition.

```
[edit policy-options policy-statement injectpolicy1 term injectterm1]
user@host# set from area 0.0.0.1
```

4. Specify that the route is to be accepted if the previous conditions are matched.

```
[edit policy-options policy-statement injectpolicy1 term injectterm1]
user@host# set then accept
```

5. Apply the routing policy to BGP.

```
[edit]
user@host# set protocols bgp export injectpolicy1
```

Results Confirm your configuration by entering the **show policy-options** and **show protocols bgp** commands from configuration mode. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@host# show policy-options
policy-statement injectpolicy1 {
  term injectterm1 {
    from {
      protocol ospf;
      area 0.0.0.1;
    }
    then accept;
  }
}

user@host# show protocols bgp
export injectpolicy1;

```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Tracing for the Routing Policy

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

```

set policy-options policy-statement injectpolicy1 term injectterm1 then trace
set routing-options traceoptions file ospf-bgp-policy-log
set routing-options traceoptions file size 5m
set routing-options traceoptions file files 5
set routing-options traceoptions flag policy

```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Include a trace action in the policy.

```

[edit policy-options policy-statement injectpolicy1 term injectterm1]
user@host# then trace

```

2. Configure the tracing file for the output.

```

[edit routing-options traceoptions]
user@host# set file ospf-bgp-policy-log
user@host# set file size 5m
user@host# set file files 5
user@host# set flag policy

```

Results

Confirm your configuration by entering the **show policy-options** and **show routing-options** commands from configuration mode. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@host# show policy-options
policy-statement injectpolicy1 {
  term injectterm1 {
    then {
      trace;
    }
  }
}

```

```
    }  
  }  
  
  user@host# show routing-options  
  traceoptions {  
    file ospf-bgp-policy-log size 5m files 5;  
    flag policy;  
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

Verifying That the Expected BGP Routes Are Present

Purpose Verify the effect of the export policy.

Action From operational mode, enter the **show route** command.

Troubleshooting

- [Using the show log Command to Examine the Actions of the Routing Policy on page 4532](#)

Using the show log Command to Examine the Actions of the Routing Policy

Problem The routing table contains unexpected routes, or routes are missing from the routing table.

Solution If you configure policy tracing as shown in this example, you can run the **show log ospf-bgp-policy-log** command to diagnose problems with the routing policy. The **show log ospf-bgp-policy-log** command displays information about the routes that the **injectpolicy1** policy term analyzes and acts upon.

Example: Redistributing Static Routes into OSPF

This example shows how to create a policy that redistributes static routes into OSPF.

- [Requirements on page 4532](#)
- [Overview on page 4533](#)
- [Configuration on page 4533](#)
- [Verification on page 4534](#)

Requirements

Before you begin:

- Configure the device interfaces. See the *Interfaces Feature Guide for Security Devices*.
- Configure static routes. See *Examples: Configuring Static Routes*.

Overview

In this example, you create a routing policy called `exportstatic1` and a routing term called `exportstatic1`. The policy injects static routes into OSPF. This example includes the following settings:

- **policy-statement**—Defines the routing policy. You specify the name of the policy and further define the elements of the policy. The policy name must be unique and can contain letters, numbers, and hyphens (-) and be up to 255 characters long.
- **term**—Defines the match condition and applicable actions for the routing policy. The term name can contain letters, numbers, and hyphens (-) and be up to 255 characters long. You specify the name of the term and define the criteria that an incoming route must match by including the **from** statement and the action to take if the route matches the conditions by including the **then** statement. In this example you specify the static protocol match condition and the accept action.
- **export**—Applies the export policy you created to be evaluated when routes are being exported from the routing table into OSPF.

Configuration

CLI Quick Configuration

To quickly create a policy that injects static routes into OSPF, copy the following commands and paste them into the CLI.

```
[edit]
set policy-options policy-statement exportstatic1 term exportstatic1 from protocol static
set policy-options policy-statement exportstatic1 term exportstatic1 then accept
set protocols ospf export exportstatic1
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Modifying the Junos OS Configuration* in the *CLI User Guide*.

To inject static routes into OSPF:

1. Create the routing policy.


```
[edit]
user@host# edit policy-options policy-statement exportstatic1
```
2. Create the policy term.


```
[edit policy-options policy-statement exportstatic1]
user@host# set term exportstatic1
```
3. Specify static as a match condition.


```
[edit policy-options policy-statement exportstatic1 term exportstatic1]
user@host# set from protocol static
```
4. Specify that the route is to be accepted if the previous condition is matched.


```
[edit policy-options policy-statement exportstatic1 term exportstatic1]
user@host# set then accept
```
5. Apply the routing policy to OSPF.



NOTE: For OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# set protocols ospf export exportstatic1
```

6. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Results Confirm your configuration by entering the `show policy-options` and `show protocols ospf` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show policy-options
policy-statement exportstatic1 {
  term exportstatic1 {
    from protocol static;
    then accept;
  }
}

user@host# show protocols ospf
export exportstatic1;
```

To confirm your OSPFv3 configuration, enter the `show policy-options` and the `show protocols ospf3` commands.

Verification

Confirm that the configuration is working properly.

- [Verifying That the Expected Static Routes Are Present on page 4534](#)
- [Verifying That AS External LSAs Are Added to the Routing Table on page 4534](#)

Verifying That the Expected Static Routes Are Present

Purpose Verify the effect of the export policy.

Action From operational mode, enter the `show route` command.

Verifying That AS External LSAs Are Added to the Routing Table

Purpose On the routing device where you configured the export policy, verify that the routing device originates an AS external LSA for the static routes that are added to the routing table.

Action From operational mode, enter the `show ospf database` command for OSPFv2, and enter the `show ospf3 database` command for OSPFv3.

Example: Configuring an OSPF Import Policy

This example shows how to create an OSPF import policy. OSPF import policies apply to external routes only. An external route is a route that is outside the OSPF autonomous system (AS).

- [Requirements on page 4535](#)
- [Overview on page 4535](#)
- [Configuration on page 4536](#)
- [Verification on page 4538](#)

Requirements

Before you begin:

- Configure static routes. See *Examples: Configuring Static Routes*.
- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 4400.
- Control OSPF designated router election. See “[Example: Controlling OSPF Designated Router Election](#)” on page 4402.
- Configure a single-area OSPF network. See “[Example: Configuring a Single-Area OSPF Network](#)” on page 4406 .

Overview

External routes are learned by AS boundary routers. External routes can be advertised throughout the OSPF domain if you configure the AS boundary router to redistribute the route into OSPF. An external route might be learned by the AS boundary router from a routing protocol other than OSPF, or the external route might be a static route that you configure on the AS boundary router.

For OSPFv3, the link-state advertisement (LSA) is referred to as the interarea prefix LSA and performs the same function as a network-summary LSA performs for OSPFv2. An area border router (ABR) originates an interarea prefix LSA for each IPv6 prefix that must be advertised into an area.

OSPF import policy allows you to prevent external routes from being added to the routing tables of OSPF neighbors. The import policy does not impact the OSPF database. This means that the import policy has no impact on the link-state advertisements. The filtering is done only on external routes in OSPF. The intra-area and interarea routes are not considered for filtering. The default action is to accept the route when the route does not match the policy.

This example includes the following OSPF policy settings:

- **policy-statement**—Defines the routing policy. You specify the name of the policy and further define the elements of the policy. The policy name must be unique and can contain letters, numbers, and hyphens (-) and be up to 255 characters long.

- **export**—Applies the export policy you created to be evaluated when network summary LSAs are flooded into an area. In this example, the export policy is named `export_static`.
- **import**—Applies the import policy you created to prevent external routes from being added to the routing table. In this example, the import policy is named `filter_routes`.

The devices you configure in this example represent the following functions:

- **R1**—Device R1 is in area 0.0.0.0 and has a direct connection to device R2. R1 has an OSPF export policy configured. The export policy redistributes static routes from R1's routing table into R1's OSPF database. Because the static route is in R1's OSPF database, the route is advertised in an LSA to R1's OSPF neighbor. R1's OSPF neighbor is device R2.
- **R2**—Device R2 is in area 0.0.0.0 and has a direct connection to device R1. R2 has an OSPF import policy configured that matches the static route to the 10.0.16.0/30 network and prevents the static route from being installed in R2's routing table. R2's OSPF neighbor is device R1.

Configuration

CLI Quick Configuration

To quickly configure an OSPF import policy, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

Configuration on Device R1:

```
[edit]
set interfaces so-0/2/0 unit 0 family inet address 10.0.2.1/30
set protocols ospf export export_static
set protocols ospf area 0.0.0.0 interface so-0/2/0
set policy-options policy-statement export_static from protocol static
set policy-options policy-statement export_static then accept
```

Configuration on Device R2:

```
[edit]
set interfaces so-0/2/0 unit 0 family inet address 10.0.2.2/30
set protocols ospf import filter_routes
set protocols ospf area 0.0.0.0 interface so-0/2/0
set policy-options policy-statement filter_routes from route-filter 10.0.16.0/30 exact
set policy-options policy-statement filter_routes then reject
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Modifying the Junos OS Configuration* in the *CLI User Guide*.

To configure an OSPF import policy:

1. Configure the interfaces.

```
[edit]
user@R1# set interfaces so-0/2/0 unit 0 family inet address 10.0.2.1/30
```



```
[edit]
user@R2# set interfaces so-0/2/0 unit 0 family inet address 10.0.2.2/30
```

2. Enable OSPF on the interfaces.



NOTE: For OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@R1# set protocols ospf area 0.0.0.0 interface so-0/2/0
```

```
[edit]
user@R2# set protocols ospf area 0.0.0.0 interface so-0/2/0
```

3. On R1, redistribute the static route into OSPF.

```
[edit]
user@R1# set protocols ospf export export_static
user@R1# set policy-options policy-statement export_static from protocol static
user@R1# set policy-options policy-statement export_static then accept
```

4. On R2, configure the OSPF import policy.

```
[edit]
user@R2# set protocols ospf import filter_routes
user@R2# set policy-options policy-statement filter_routes from route-filter
10.0.16.0/30 exact
user@R2# set policy-options policy-statement filter_routes then reject
```

5. If you are done configuring the devices, commit the configuration.

```
[edit]
user@host# commit
```

Results Confirm your configuration by entering the `show interfaces`, `show policy-options`, and `show protocols ospf` commands on the appropriate device. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Output for R1:

```
user@R1# show interfaces
so-0/2/0 {
  unit 0 {
    family inet {
      address 10.0.2.1/30;
    }
  }
}

user@R1# show policy-options
policy-statement export_static {
  from protocol static;
  then accept;
}
```

```
user@R1# show protocols ospf
export export_static;
area 0.0.0.0 {
    interface so-0/2/0.0;
}
```

Output for R2:

```
user@R2# show interfaces
so-0/2/0 {
    unit 0 {
        family inet {
            address 10.0.2.2/30;
        }
    }
}

user@R2# show policy-options
policy-statement filter_routes {
    from {
        route-filter 10.0.16.0/30 exact;
    }
    then reject;
}

user@R2# show protocols ospf
import filter_routes;
area 0.0.0.0 {
    interface so-0/2/0.0;
}
```

To confirm your OSPFv3 configuration, enter the **show interfaces**, **show policy-options**, **show routing-options**, and **show protocols ospf3** commands on the appropriate device.

Verification

Confirm that the configuration is working properly.

- [Verifying the OSPF Database on page 4538](#)
- [Verifying the Routing Table on page 4538](#)

Verifying the OSPF Database

Purpose Verify that OSPF is advertising the static route in the OSPF database.

Action From operational mode, enter the **show ospf database** for OSPFv2, and enter the **show ospf3 database** command for OSPFv3.

Verifying the Routing Table

Purpose Verify the entries in the routing table.

Action From operational mode, enter the **show route** command.

Example: Configuring a Route Filter Policy to Specify Priority for Prefixes Learned Through OSPF

This example shows how to create an OSPF import policy that prioritizes specific prefixes learned through OSPF.

- [Requirements on page 4539](#)
- [Overview on page 4539](#)
- [Configuration on page 4540](#)
- [Verification on page 4543](#)

Requirements

Before you begin:

- Configure the device interfaces. See the *Interfaces Feature Guide for Security Devices*.
- Configure the router identifiers for the devices in your OSPF network. See [“Example: Configuring an OSPF Router Identifier” on page 4400](#).
- Control OSPF designated router election. See [“Example: Controlling OSPF Designated Router Election” on page 4402](#).
- Configure a single-area OSPF network. See [“Example: Configuring a Single-Area OSPF Network” on page 4406](#).
- Configure a multiarea OSPF network. See [“Example: Configuring a Multiarea OSPF Network” on page 4408](#).

Overview

In a network with a large number of OSPF routes, it can be useful to control the order in which routes are updated in response to a network topology change. In Junos OS Release 9.3 and later, you can specify a priority of high, medium, or low for prefixes included in an OSPF import policy. In the event of an OSPF topology change, high priority prefixes are updated in the routing table first, followed by medium and then low priority prefixes.

OSPF import policy can only be used to set priority or to filter OSPF external routes. If an OSPF import policy is applied that results in a **reject** terminating action for a nonexternal route, then the **reject** action is ignored and the route is accepted anyway. By default, such a route is now installed in the routing table with a priority of low. This behavior prevents traffic black holes, that is, silently discarded traffic, by ensuring consistent routing within the OSPF domain.

In general, OSPF routes that are not explicitly assigned a priority are treated as priority medium, except for the following:

- Summary discard routes have a default priority of low.
- Local routes that are not added to the routing table are assigned a priority of low.

- External routes that are rejected by import policy and thus not added to the routing table are assigned a priority of low.

Any available match criteria applicable to OSPF routes can be used to determine the priority. Two of the most commonly used match criteria for OSPF are the **route-filter** and **tag** statements.

In this example, the routing device is in area 0.0.0.0, with interfaces **fe-0/1/0** and **fe-1/1/0** connecting to neighboring devices. You configure an import routing policy named **ospf-import** to specify a priority for prefixes learned through OSPF. Routes associated with these prefixes are installed in the routing table in the order of the prefixes' specified priority. Routes matching **200.3.0.0/16 orlonger** are installed first because they have a priority of **high**. Routes matching **200.2.0.0/16 orlonger** are installed next because they have a priority of **medium**. Routes matching **200.1.0.0/16 orlonger** are installed last because they have a priority of **low**. You then apply the import policy to OSPF.



NOTE: The priority value takes effect when a new route is installed, or when there is a change to an existing route.

Configuration

CLI Quick Configuration

To quickly configure an OSPF import policy that prioritizes specific prefixes learned through OSPF, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

```
[edit]
set interfaces fe-0/1/0 unit 0 family inet address 192.168.8.4/30
set interfaces fe-0/1/0 unit 0 family inet address 192.168.8.5/30
set policy-options policy-statement ospf-import term t1 from route-filter 200.1.0.0/16
  orlonger
set policy-options policy-statement ospf-import term t1 then priority low
set policy-options policy-statement ospf-import term t1 then accept
set policy-options policy-statement ospf-import term t2 from route-filter 200.2.0.0/16
  orlonger
set policy-options policy-statement ospf-import term t2 then priority medium
set policy-options policy-statement ospf-import term t2 then accept
set policy-options policy-statement ospf-import term t3 from route-filter 200.3.0.0/16
  orlonger
set policy-options policy-statement ospf-import term t3 then priority high
set policy-options policy-statement ospf-import term t3 then accept
set protocols ospf import ospf-import
set protocols ospf area 0.0.0.0 interface fe-0/1/0
set protocols ospf area 0.0.0.0 interface fe-1/1/0
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Modifying the Junos OS Configuration* in the *CLI User Guide*.

To configure an OSPF import policy that prioritizes specific prefixes:

1. Configure the interfaces.

```
[edit]
user@host# set interfaces fe-0/1/0 unit 0 family inet address 192.168.8.4/30
user@host# set interfaces fe-0/2/0 unit 0 family inet address 192.168.8.5/30
```

2. Enable OSPF on the interfaces.



NOTE: For OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# set protocols ospf area 0.0.0.0 interface fe-0/1/0
user@host# set protocols ospf area 0.0.0.0 interface fe-0/2/0
```

3. Configure the policy to specify the priority for prefixes learned through OSPF.

```
[edit ]
user@host# set policy-options policy-statement ospf-import term t1 from route-filter
200.1.0.0/16 orlonger
user@host# set policy-options policy-statement ospf-import term t1 then priority
low
user@host# set policy-options policy-statement ospf-import term t1 then accept
user@host# set policy-options policy-statement ospf-import term t2 from route-filter
200.2.0.0/16 orlonger
user@host# set policy-options policy-statement ospf-import term t2 then priority
medium
user@host# set policy-options policy-statement ospf-import term t2 then accept
user@host# set policy-options policy-statement ospf-import term t3 from route-filter
200.3.0.0/16 orlonger
user@host# set policy-options policy-statement ospf-import term t3 then priority
high
user@host# set policy-options policy-statement ospf-import term t3 then accept
```

4. Apply the policy to OSPF.

```
[edit]
user@host# set protocols ospf import ospf-import
```

5. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Results Confirm your configuration by entering the `show interfaces`, `show policy-options`, and the `show protocols ospf` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
```

```
fe-0/1/0 {
  unit 0 {
    family inet {
      address 192.168.8.4/30;
    }
  }
}
fe-0/2/0 {
  unit 0 {
    family inet {
      address 192.168.8.5/30;
    }
  }
}

user@host# show protocols ospf
import ospf-import;
area 0.0.0.0 {
  interface fe-0/1/0.0;
  interface fe-0/2/0.0;
}

user@host# show policy-options
policy-statement ospf-import {
  term t1 {
    from {
      route-filter 200.1.0.0/16 orlonger;
    }
    then {
      priority low;
      accept;
    }
  }
  term t2 {
    from {
      route-filter 200.2.0.0/16 orlonger;
    }
    then {
      priority medium;
      accept;
    }
  }
  term t3 {
    from {
      route-filter 200.3.0.0/16 orlonger;
    }
    then {
      priority high;
      accept;
    }
  }
}

user@host# show protocols ospf
import ospf-import;
area 0.0.0.0 {
  interface fe-0/1/0.0;
```

```

    interface fe-0/2/0.0;
  }

```

To confirm your OSPFv3 configuration, enter the **show interfaces**, **show policy-options**, and **show protocols ospf3** commands.

Verification

Confirm that the configuration is working properly.

Verifying the Prefix Priority in the OSPF Routing Table

Purpose	Verify the priority assigned to the prefix in the OSPF routing table.
Action	From operational mode, enter the show ospf route detail for OSPFv2, and enter the show ospf3 route detail command for OSPFv3.
Related Documentation	<ul style="list-style-type: none"> • OSPF Overview on page 4386 • Understanding OSPF Configurations • Routing Policy Match Conditions in the Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices • Actions in Routing Policy Terms in the Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices

Examples: Configuring Routing Policy for Network Summaries

- [Import and Export Policies for Network Summaries Overview on page 4543](#)
- [Example: Configuring an OSPF Export Policy for Network Summaries on page 4544](#)
- [Example: Configuring an OSPF Import Policy for Network Summaries on page 4554](#)

Import and Export Policies for Network Summaries Overview

By default, OSPF uses network-summary link-state advertisements (LSAs) to transmit route information across area boundaries. Each area border router (ABR) floods network-summary LSAs to other routing devices in the same area. The ABR also controls which routes from the area are used to generate network-summary LSAs into other areas. Each ABR maintains a separate topological database for each area to which they are connected. In Junos OS Release 9.1 and later, you can configure export and import policies for OSPFv2 and OSPFv3 that enable you to control how network-summary LSAs, which contain information about interarea OSPF prefixes, are distributed and generated. For OSPFv3, the LSA is referred to as the interarea prefix LSA and performs the same function as a network-summary LSA performs for OSPFv2. An ABR originates an interarea prefix LSA for each IPv6 prefix that must be advertised into an area.

The export policy enables you to specify which summary LSAs are flooded into an area. The import policy enables you to control which routes learned from an area are used to generate summary LSAs into other areas. You define a routing policy at the **[edit policy-options policy-statement *policy-name*]** hierarchy level. As with all OSPF export

policies, the default for network-summary LSA export policies is to reject everything. Similarly, as with all OSPF import policies, the default for network-summary LSA import policies is to accept all OSPF routes.

Example: Configuring an OSPF Export Policy for Network Summaries

This example shows how to create an OSPF export policy to control the network-summary (Type 3) LSAs that the ABR floods into an OSPF area.

- [Requirements on page 4544](#)
- [Overview on page 4544](#)
- [Configuration on page 4546](#)
- [Verification on page 4554](#)

Requirements

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See [“Example: Configuring an OSPF Router Identifier” on page 4400](#).
- Control OSPF designated router election. See [“Example: Controlling OSPF Designated Router Election” on page 4402](#)

Overview

OSPF uses network-summary LSAs to transmit route information across area boundaries. Depending on your network environment, you might want to further filter the network-summary LSAs between OSPF areas. For example, if you create OSPF areas to define administrative boundaries, you might not want to advertise internal route information between those areas. To further improve the control of route distribution between multiple OSPF areas, you can configure network summary policies on the ABR for the area that you want to filter the advertisement of network-summary LSAs.



NOTE: For OSPFv3, the LSA is referred to as the interarea prefix LSA and performs the same function as a network-summary LSA performs for OSPFv2. An ABR originates an interarea prefix LSA for each IPv6 prefix that must be advertised into an area. In this topic, the terms network summary policy and network-summary policy are used to describe both OSPFv2 and OSPFv3 functionality.


The following guidelines apply to export network summary policies:

- You should have a thorough understanding of your network before configuring these policies. Incorrect network summary policy configuration might result in an unintended result such as suboptimal routing or dropped traffic.
- We recommend that you use the **route-filter** policy match condition for these types of policies.

- We recommend that you use the **accept** and **reject** routing policy terms for these types of policies.

Figure 119 shows a sample topology with three OSPF areas. R4 generates network summaries for the routes in area 4 and sends them out of area 4 to area 0. R3 generates network summaries for the routes in area 3 and sends them out of area 3 to area 0.

Figure 119: Sample Topology Used for an OSPF Export Network Summary Policy



In this example, you configure R4 with an export network summary policy named `export-policy` that only allows routes that match the 10.0.4.4 prefix from area 3 into area 4. The export policy controls the network-summary LSAs that R4 floods into area 4. This results in only the allowed interarea route to enter area 4, and all other interarea routes to be purged from the OSPF database and the routing table of the devices in area 4. You first define the policy and then apply it to the ABR by including the `network-summary-export` statement for OSPFv2 or the `inter-area-prefix-export` statement for OSPFv3.

The devices operate as follows:

- R1—Device R1 is an internal router in area 3. Interface **fe-0/1/0** has an IP address of 10.0.4.13/30 and connects to R3. Interface **fe-0/0/1** has an IP address of 10.0.4.5/30 and connects to R2.
- R2—Device R2 is an internal router in area 3. Interface **fe-0/0/1** has an IP address of 10.0.4.6/30 and connects to R1. Interface **fe-1/0/0** has an IP address of 10.0.4.1 and connects to R3.
- R3—Device R3 participates in area 3 and area 0. R3 is the ABR between area 3 and area 0, and passes network-summary LSAs between the areas. Interface **fe-1/0/0** has an IP address of 10.0.4.2/30 and connects to R2. Interface **fe-1/1/0** has an IP address of 10.0.4.14/30 and connects to R1. Interface **fe-0/0/1** has an IP address of 10.0.2.1/30 and connects to R4.
- R4—Device R4 participates in area 0 and area 4. R4 is the ABR between area 0 and area 4, and passes network-summary LSAs between the areas. Interface **fe-0/0/1** has

an IP address of 10.0.2.4/30 and connects to R3. Interface **fe-1/1/0** has an IP address of 10.0.8.6/30 and connects to R5. Interface **fe-1/0/0** has an IP address of 10.0.8.9/30 and connects to R6.

- R5—Device R5 is an internal router in area 4. Interface **fe-1/1/0** has an IP address of 10.0.8.5/30 and connects to R4.
- R6—Device R6 is an internal router in area 4. Interface **fe-1/0/0** has an IP address of 10.0.8.10/30 and connects to R4.

Configuration

CLI Quick Configuration

To quickly configure an OSPF export policy for network summaries, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

Configuration on Device R1:

```
[edit]
set interfaces fe-0/1/0 unit 0 family inet address 10.0.4.13/30
set interfaces fe-0/0/1 unit 0 family inet address 10.0.4.5/30
set protocols ospf area 0.0.0.3 interface fe-0/1/0
set protocols ospf area 0.0.0.3 interface fe-0/0/1
```

Configuration on Device R2:

```
[edit]
set interfaces fe-0/1/0 unit 0 family inet address 10.0.4.6/30
set interfaces fe-1/0/0 unit 0 family inet address 10.0.4.1/30
set protocols ospf area 0.0.0.3 interface fe-0/1/0
set protocols ospf area 0.0.0.3 interface fe-1/0/0
```

Configuration on Device R3:

```
[edit]
set interfaces fe-1/0/0 unit 0 family inet address 10.0.4.2/30
set interfaces fe-1/1/0 unit 0 family inet address 10.0.4.14/30
set interfaces fe-0/0/1 unit 0 family inet address 10.0.2.1/30
set protocols ospf area 0.0.0.3 interface fe-1/0/0
set protocols ospf area 0.0.0.3 interface fe-1/1/0
set protocols ospf area 0.0.0.0 interface fe-0/0/1
```

Configuration on Device R4:

```
[edit]
set interfaces fe-0/0/1 unit 0 family inet address 10.0.2.1/30
set interfaces fe-1/1/0 unit 0 family inet address 10.0.8.6/30
set interfaces fe-1/0/0 unit 0 family inet address 10.0.8.9/30
set policy-options policy-statement export-policy term term1 from route-filter 10.0.4.4/30
  prefix-length-range /30-/30
set policy-options policy-statement export-policy term term1 then accept
set protocols ospf area 0.0.0.0 interface fe-0/0/1
set protocols ospf area 0.0.0.4 interface fe-0/1/0
set protocols ospf area 0.0.0.4 interface fe-1/0/0
set protocols ospf area 0.0.0.4 network-summary-export export-policy
```

Configuration on Device R5:

```
[edit]
set interfaces fe-1/1/0 unit 0 family inet address 10.0.8.5/30
set protocols ospf area 0.0.0.4 interface fe-0/1/0
```

Configuration on Device R6:

```
[edit]
set interfaces fe-1/0/0 unit 0 family inet address 10.0.8.10/30
set protocols ospf area 0.0.0.4 interface fe-1/0/0
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Modifying the Junos OS Configuration* in the *CLI User Guide*.

To configure an OSPF export policy for network summaries:

1. Configure the interfaces.



NOTE: For OSPFv3, use IPv6 addresses.

```
[edit]
user@R1# set interfaces fe-0/1/0 unit 0 family inet address 10.0.4.13/30
user@R1# set interfaces fe-0/0/1 unit 0 family inet address 10.0.4.5/30
```

```
[edit]
user@R2# set interfaces fe-0/1/0 unit 0 family inet address 10.0.4.6/30
user@R2# set interfaces fe-1/0/0 unit 0 family inet address 10.0.4.1/30
```

```
[edit]
user@R3# set interfaces fe-1/0/0 unit 0 family inet address 10.0.4.2/30
user@R3# set interfaces fe-1/1/0 unit 0 family inet address 10.0.4.14/30
user@R3# set interfaces fe-0/0/1 unit 0 family inet address 10.0.2.1/30
```

```
[edit]
user@R4# set interfaces fe-0/0/1 unit 0 family inet address 10.0.2.1/30
user@R4# set interfaces fe-1/1/0 unit 0 family inet address 10.0.8.6/30
user@R4# set interfaces fe-1/0/0 unit 0 family inet address 10.0.8.9/30
```

```
[edit]
user@R5# set interfaces fe-1/1/0 unit 0 family inet address 10.0.8.5/30
```

```
[edit]
user@R6# set interfaces fe-1/0/0 unit 0 family inet address 10.0.8.10/30
```

2. Enable OSPF on the interfaces.



NOTE: For OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@R1# set protocols ospf area 0.0.0.3 interface fe-0/1/0
```

```
user@R1# set protocols ospf area 0.0.0.3 interface fe-0/0/1
```

```
[edit]
```

```
user@R2# set protocols ospf area 0.0.0.3 interface fe-0/1/0
```

```
user@R2# set protocols ospf area 0.0.0.3 interface fe-1/0/0
```

```
[edit]
```

```
user@R3# set protocols ospf area 0.0.0.3 interface fe-1/0/0
```

```
user@R3# set protocols ospf area 0.0.0.3 interface fe-1/1/0
```

```
user@R3# set protocols ospf area 0.0.0.0 interface fe-0/0/1
```

```
[edit]
```

```
user@R4# set protocols ospf area 0.0.0.0 interface fe-0/0/1
```

```
user@R4# set protocols ospf area 0.0.0.4 interface fe-1/1/0
```

```
user@R4# set protocols ospf area 0.0.0.4 interface fe-1/0/0
```

```
[edit]
```

```
user@R5# set protocols ospf area 0.0.0.4 interface fe-1/1/0
```

```
[edit]
```

```
user@R6# set protocols ospf area 0.0.0.4 interface fe-1/0/0
```

3. On R4, configure the export network summary policy.

```
[edit]
```

```
user@R4# set policy-options policy-statement export-policy term term1 from  
route-filter 10.0.4.4/30 prefix-length-range /30-/30
```

```
user@R4# set policy-options policy-statement export-policy term term1 then accept
```

4. On R4, apply the export network summary policy to OSPF.



NOTE: For OSPFv3, include the inter-area-prefix-export statement at the [edit protocols ospf3 area *area-id*] hierarchy level.

```
[edit]
```

```
user@R4# set protocols ospf area 0.0.0.4 network-summary-export export-policy
```

5. If you are done configuring the devices, commit the configuration.

```
[edit]
```

```
user@host# commit
```

Results Confirm your configuration by entering the **show interfaces**, **show policy-options**, and **show protocols ospf** commands on the appropriate device. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Output for R1:

```
user@R1# show interfaces  
fe-0/0/1 {  
  unit 0 {  
    family inet {  
      address 10.0.4.5/30;  
    }  
  }  
}
```

```

}
fe-1/1/0 {
  unit 0 {
    family inet {
      address 10.0.4.13/30;
    }
  }
}

user@R1# show protocols ospf
area 0.0.0.3 {
  interface fe-0/1/0.0;
  interface fe-0/0/1.0;
}

```

Output for R2:

```

user@R2# show interfaces
fe-0/1/0 {
  unit 0 {
    family inet {
      address 10.0.4.6/30;
    }
  }
}
fe-1/0/0 {
  unit 0 {
    family inet {
      address 10.0.4.3/30;
    }
  }
}

user@R2# show protocols ospf
area 0.0.0.3 {
  interface fe-0/1/0.0;
  interface fe-1/0/0.0;
}

```

Output for R3:

```

user@R3# show interfaces
fe-0/0/1 {
  unit 0 {
    family inet {
      address 10.0.2.3/30;
    }
  }
}
fe-1/0/0 {
  unit 0 {
    family inet {
      address 10.0.4.2/30;
    }
  }
}
fe-1/1/0 {
  unit 0 {

```

```
        family inet {
            address 10.0.4.14/30;
        }
    }
}
```

```
user@R3# show protocols ospf
area 0.0.0.0 {
    interface fe-0/0/1.0;
}
area 0.0.0.3 {
    interface fe-1/0/0.0;
    interface fe-1/1/0.0;
}
```

Output for R4:

```
user@R4# show interfaces
fe-0/0/1 {
    unit 0 {
        family inet {
            address 10.0.2.4/30;
        }
    }
}
fe-1/0/0 {
    unit 0 {
        family inet {
            address 10.0.8.6/30;
        }
    }
}
fe-1/1/0 {
    unit 0 {
        family inet {
            address 10.0.8.3/30;
        }
    }
}
```

```
user@R4# show protocols ospf
area 0.0.0.0 {
    interface fe-0/0/1.0;
}
area 0.0.0.4 {
    network-summary-export export-policy;
    interface fe-1/0/0.0;
    interface fe-1/1/0.0;
}
```

```
user@R4# show policy-options
policy-statement export-policy {
    term term1 {
        from {
            route-filter 10.0.4.4/30 prefix-length-range /30-/30;
        }
        then accept;
    }
}
```

```

    }
  }

```

Output for R5:

```

user@R5# show interfaces
fe-1/1/0 {
  unit 0 {
    family inet {
      address 10.0.8.5/30;
    }
  }
}

user@R5# show protocols ospf
area 0.0.0.4 {
  interface fe-1/1/0.0;
}

```

Output for R6:

```

user@R6# show interfaces
fe-1/0/0 {
  unit 0 {
    family inet {
      address 10.0.8.7/30;
    }
  }
}

user@R6# show protocols ospf
area 0.0.0.4 {
  interface fe-1/0/0.0;
}

```

```

user@R1# show interfaces
fe-0/0/1 {
  unit 0 {
    family inet {
      address 10.0.4.5/30;
    }
  }
}
fe-1/1/0 {
  unit 0 {
    family inet {
      address 10.0.4.13/30;
    }
  }
}

user@R1# show protocols ospf
area 0.0.0.3 {
  interface fe-0/1/0.0;
  interface fe-0/0/1.0;
}

```

Output for R2:

```
user@R2# show interfaces
fe-0/1/0 {
  unit 0 {
    family inet {
      address 10.0.4.6/30;
    }
  }
}
fe-1/0/0 {
  unit 0 {
    family inet {
      address 10.0.4.1/30;
    }
  }
}

user@R2# show protocols ospf
area 0.0.0.3 {
  interface fe-0/1/0.0;
  interface fe-1/0/0.0;
}
```

Output for R3:

```
user@R3# show interfaces
fe-0/0/1 {
  unit 0 {
    family inet {
      address 10.0.2.1/30;
    }
  }
}
fe-1/0/0 {
  unit 0 {
    family inet {
      address 10.0.4.2/30;
    }
  }
}
fe-1/1/0 {
  unit 0 {
    family inet {
      address 10.0.4.14/30;
    }
  }
}

user@R3# show protocols ospf
area 0.0.0.0 {
  interface fe-0/0/1.0;
}
area 0.0.0.3 {
  interface fe-1/0/0.0;
  interface fe-1/1/0.0;
}
```

Output for R4:


```

user@R4# show interfaces
fe-0/0/1 {
  unit 0 {
    family inet {
      address 10.0.2.1/30;
    }
  }
}
fe-1/0/0 {
  unit 0 {
    family inet {
      address 10.0.8.9/30;
    }
  }
}
fe-1/1/0 {
  unit 0 {
    family inet {
      address 10.0.8.6/30;
    }
  }
}

user@R4# show protocols ospf
area 0.0.0.0 {
  interface fe-0/0/1.0;
}
area 0.0.0.4 {
  network-summary-export export-policy;
  interface fe-1/0/0.0;
  interface fe-1/1/0.0;
}

user@R4# show policy-options
policy-statement export-policy {
  term term1 {
    from {
      route-filter 10.0.4.4/30 prefix-length-range /30-/30;
    }
    then accept;
  }
}

```

Output for R5:

```

user@R5# show interfaces
fe-1/1/0 {
  unit 0 {
    family inet {
      address 10.0.8.5/30;
    }
  }
}

user@R5# show protocols ospf
area 0.0.0.4 {
  interface fe-1/1/0.0;
}

```

```
}
```

Output for R6:

```
user@R6# show interfaces
fe-1/0/0 {
  unit 0 {
    family inet {
      address 10.0.8.10/30;
    }
  }
}

user@R6# show protocols ospf
area 0.0.0.4 {
  interface fe-1/0/0.0;
}
```

To confirm your OSPFv3 configuration, enter the **show interfaces**, **show policy-options**, and **show protocols ospf3** commands on the appropriate device.

Verification

Confirm that the configuration is working properly.

- [Verifying the OSPF Database on page 4554](#)
- [Verifying the Routing Table on page 4554](#)

Verifying the OSPF Database

Purpose Verify that the OSPF database for the devices in area 4 includes the interarea route that we permitted on the ABR R4. The other interarea routes that are not specified should age out or no longer be present in the OSPF database.

Action From operational mode, enter the **show ospf database netsummary area 0.0.0.4** command for OSPFv2, and enter the **show ospf3 database inter-area-prefix area 0.0.0.4** command for OSPFv3.

Verifying the Routing Table

Purpose Verify that the routes corresponding to the rejected network summaries are no longer present in R4's, R5's, or R6's routing table.

Action From operational mode, enter the **show route protocol ospf** command for both OSPFv2 and OSPFv3.

Example: Configuring an OSPF Import Policy for Network Summaries

This example shows how to create an OSPF import policy to control the network-summary (Type 3) LSAs that the ABR advertises out of an OSPF area.

- [Requirements on page 4555](#)
- [Overview on page 4555](#)

- [Configuration on page 4557](#)
- [Verification on page 4562](#)

Requirements

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See [“Example: Configuring an OSPF Router Identifier” on page 4400](#).
- Control OSPF designated router election. See [“Example: Controlling OSPF Designated Router Election” on page 4402](#).

Overview

OSPF uses network-summary LSAs to transmit route information across area boundaries. Depending on your network environment, you might want to further filter the network-summary LSAs between OSPF areas. For example, if you create OSPF areas to define administrative boundaries, you might not want to advertise internal route information between those areas. To further improve the control of route distribution between multiple OSPF areas, you can configure network summary policies on the ABR for the area that you want to filter the advertisement of network-summary LSAs.




NOTE: For OSPFv3, the LSA is referred to as the interarea prefix LSA and performs the same function as a network-summary LSA performs for OSPFv2. An ABR originates an interarea prefix LSA for each IPv6 prefix that must be advertised into an area. In this topic, the terms network summary policy and network-summary policy are used to describe both OSPFv2 and OSPFv3 functionality.

The following guidelines apply to import network summary policies:

- You should have a thorough understanding of your network before configuring these policies. Incorrect network summary policy configuration might result in an unintended result such as suboptimal routing or dropped traffic.
- We recommend that you use the **route-filter** policy match condition for these types of policies.
- We recommend that you use the **accept** and **reject** routing policy terms for these types of policies.

[Figure 120](#) shows a sample topology with three OSPF areas. R4 generates network summaries for the routes in area 4 and sends them out of area 4 to area 0. R3 generates network summaries for the routes in area 3 and sends them out of area 3 to area 0.

Figure 120: Sample Topology Used for an OSPF Import Network Summary Policy



In this example, you configure R3 with an import network summary policy named `import-policy` so R3 only generates network summaries for the route 10.0.4.12/30. The import policy controls the routes and therefore the network summaries that R3 advertises out of area 3, so applying this policy means that R3 only advertises route 10.0.4.12/30 out of area 3. This results in existing network summaries from other interarea routes getting purged from the OSPF database in area 0 and area 4, as well as the routing tables of the devices in areas 0 and area 4. You first define the policy and then apply it to the ABR by including the **network-summary-import** statement for OSPFv2 or the **inter-area-prefix-import** statement for OSPFv3.

The devices operate as follows:

- **R1**—Device R1 is an internal router in area 3. Interface **fe-0/1/0** has an IP address of 10.0.4.13/30 and connects to R3. Interface **fe-0/0/1** has an IP address of 10.0.4.5/30 and connects to R2.
- **R2**—Device R2 is an internal router in area 3. Interface **fe-0/0/1** has an IP address of 10.0.4.6/30 and connects to R1. Interface **fe-1/0/0** has an IP address of 10.0.4.1/30 and connects to R3.
- **R3**—Device R3 participates in area 3 and area 0. R3 is the ABR between area 3 and area 0, and passes network-summary LSAs between the areas. Interface **fe-1/0/0** has an IP address of 10.0.4.2/30 and connects to R2. Interface **fe-1/1/0** has an IP address of 10.0.4.14/30 and connects to R1. Interface **fe-0/0/1** has an IP address of 10.0.2.1/30 and connects to R4.
- **R4**—Device R4 participates in area 0 and area 4. R4 is the ABR between area 0 and area 4, and passes network-summary LSAs between the areas. Interface **fe-0/0/1** has an IP address of 10.0.2.1/30 and connects to R3. Interface **fe-1/1/0** has an IP address of 10.0.8.6/30 and connects to R5. Interface **fe-1/0/0** has an IP address of 10.0.8.9/30 and connects to R6.

- R5—Device R5 is an internal router in area 4. Interface **fe-1/1/0** has an IP address of 10.0.8.5/30 and connects to R4.
- R6—Device R6 is an internal router in area 4. Interface **fe-1/0/0** has an IP address of 10.0.8.10/30 and connects to R4.

Configuration

CLI Quick Configuration

To quickly configure an OSPF import policy for network summaries, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

Configuration on Device R1:

```
[edit]
set interfaces fe-0/1/0 unit 0 family inet address 10.0.4.13/30
set interfaces fe-0/0/1 unit 0 family inet address 10.0.4.5/30
set protocols ospf area 0.0.0.3 interface fe-0/1/0
set protocols ospf area 0.0.0.3 interface fe-0/0/1
```

Configuration on Device R2:

```
[edit]
set interfaces fe-0/1/0 unit 0 family inet address 10.0.4.6/30
set interfaces fe-1/0/0 unit 0 family inet address 10.0.4.1/30
set protocols ospf area 0.0.0.3 interface fe-0/1/0
set protocols ospf area 0.0.0.3 interface fe-1/0/0
```

Configuration on Device R3:

```
[edit]
set interfaces fe-1/0/0 unit 0 family inet address 10.0.4.2/30
set interfaces fe-1/1/0 unit 0 family inet address 10.0.4.14/30
set interfaces fe-0/0/1 unit 0 family inet address 10.0.2.1/30
set policy-options policy-statement import-policy term term1 from route-filter 10.0.4.12/30
  prefix-length-range /30-/30
set policy-options policy-statement import-policy term term1 then accept
set protocols ospf area 0.0.0.3 interface fe-1/0/0
set protocols ospf area 0.0.0.3 interface fe-1/1/0
set protocols ospf area 0.0.0.0 interface fe-0/0/1
set protocols ospf area 0.0.0.3 network-summary-import import-policy
```

Configuration on Device R4:

```
[edit]
set interfaces fe-0/0/1 unit 0 family inet address 10.0.2.1/30
set interfaces fe-1/1/0 unit 0 family inet address 10.0.8.6/30
set interfaces fe-1/0/0 unit 0 family inet address 10.0.8.9/30
set protocols ospf area 0.0.0.0 interface fe-0/0/1
set protocols ospf area 0.0.0.4 interface fe-1/1/0
set protocols ospf area 0.0.0.4 interface fe-1/0/0
```

Configuration on Device R5:

```
[edit]
set interfaces fe-1/1/0 unit 0 family inet address 10.0.8.5/30
```

```
set protocols ospf area 0.0.0.4 interface fe-1/1/0
```

Configuration on Device R6:

```
[edit]
set interfaces fe-1/0/0 unit 0 family inet address 10.0.8.10/30
set protocols ospf area 0.0.0.4 interface fe-1/0/0
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Modifying the Junos OS Configuration* in the *CLI User Guide*.

To configure an OSPF import policy for network summaries:

1. Configure the interfaces.



NOTE: For OSPFv3, use IPv6 addresses.

```
[edit]
user@R1# set interfaces fe-0/1/0 unit 0 family inet address 10.0.4.13/30
user@R1# set interfaces fe-0/0/1 unit 0 family inet address 10.0.4.5/30
```

```
[edit]
user@R2# set interfaces fe-0/1/0 unit 0 family inet address 10.0.4.6/30
user@R2# set interfaces fe-1/0/0 unit 0 family inet address 10.0.4.1/30
```

```
[edit]
user@R3# set interfaces fe-1/0/0 unit 0 family inet address 10.0.4.2/30
user@R3# set interfaces fe-1/1/0 unit 0 family inet address 10.0.4.14/30
user@R3# set interfaces fe-0/0/1 unit 0 family inet address 10.0.2.1/30
```

```
[edit]
user@R4# set interfaces fe-0/0/1 unit 0 family inet address 10.0.2.1/30
user@R4# set interfaces fe-1/1/0 unit 0 family inet address 10.0.8.6/30
user@R4# set interfaces fe-1/0/0 unit 0 family inet address 10.0.8.9/30
```

```
[edit]
user@R5# set interfaces fe-1/1/0 unit 0 family inet address 10.0.8.5/30
```

```
[edit]
user@R6# set interfaces fe-1/0/0 unit 0 family inet address 10.0.8.10/30
```

2. Enable OSPF on the interfaces.



NOTE: For OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@R1# set protocols ospf area 0.0.0.3 interface fe-0/1/0
user@R1# set protocols ospf area 0.0.0.3 interface fe-0/0/1
```

```
[edit]
user@R2# set protocols ospf area 0.0.0.3 interface fe-0/1/0
```

```
user@R2# set protocols ospf area 0.0.0.3 interface fe-1/0/0
```

```
[edit]
```

```
user@R3# set protocols ospf area 0.0.0.3 interface fe-1/0/0
```

```
user@R3# set protocols ospf area 0.0.0.3 interface fe-1/1/0
```

```
user@R3# set protocols ospf area 0.0.0.0 interface fe-0/0/1
```

```
[edit]
```

```
user@R4# set protocols ospf area 0.0.0.0 interface fe-0/0/1
```

```
user@R4# set protocols ospf area 0.0.0.4 interface fe-1/1/0
```

```
user@R4# set protocols ospf area 0.0.0.4 interface fe-1/0/0
```

```
[edit]
```

```
user@R5# set protocols ospf area 0.0.0.4 interface fe-1/1/0
```

```
[edit]
```

```
user@R6# set protocols ospf area 0.0.0.4 interface fe-1/0/0
```

3. On R3, configure the import network summary policy.

```
[edit]
```

```
user@R3# set policy-options policy-statement import-policy term term1 from
route-filter 10.0.4.12/30 prefix-length-range /30-/30
```

```
user@R3# set policy-options policy-statement import-policy term term1 then accept
```

4. On R3, apply the import network summary policy to OSPF.



NOTE: For OSPFv3, include the `inter-area-prefix-export` statement at the `[edit protocols ospf3 area area-id]` hierarchy level.

```
[edit]
```

```
user@R3# set protocols ospf area 0.0.0.3 network-summary-import import-policy
```

5. If you are done configuring the devices, commit the configuration.

```
[edit]
```

```
user@host# commit
```

Results Confirm your configuration by entering the `show interfaces`, `show policy-options`, and `show protocols ospf` commands on the appropriate device. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Output for R1:

```
user@R1# show interfaces
```

```
fe-0/0/1 {
```

```
  unit 0 {
```

```
    family inet {
```

```
      address 10.0.4.5/30;
```

```
    }
```

```
  }
```

```
fe-0/1/0 {
```

```
  unit 0 {
```

```
    family inet {
```

```
        address 10.0.4.13/30;
    }
}

user@R1# show protocols ospf
area 0.0.0.3 {
    interface fe-0/1/0.0;
    interface fe-0/0/1.0;
}
```

Output for R2:

```
user@R2# show interfaces
fe-0/1/0 {
    unit 0 {
        family inet {
            address 10.0.4.6/30;
        }
    }
}
fe-1/0/0 {
    unit 0 {
        family inet {
            address 10.0.4.1/30;
        }
    }
}

user@R2# show protocols ospf
area 0.0.0.3 {
    interface fe-0/1/0.0;
    interface fe-1/0/0.0;
}
```

Output for R3:

```
user@R3# show interfaces
fe-0/0/1 {
    unit 0 {
        family inet {
            address 10.0.2.1/30;
        }
    }
}
fe-1/0/0 {
    unit 0 {
        family inet {
            address 10.0.4.2/30;
        }
    }
}
fe-1/1/0 {
    unit 0 {
        family inet {
            address 10.0.4.14/30;
        }
    }
}
```



```

}

user@R3# show protocols ospf
area 0.0.0.0 {
  interface fe-0/0/1.0;
}
area 0.0.0.3 {
  network-summary-import import policy;
  interface fe-1/0/0.0;
  interface fe-1/1/0.0;
}

user@R3# show policy-options
policy-statement import-policy {
  term term1 {
    from {
      route-filter 10.0.4.12/30 prefix-length-range /30-/30;
    }
    then accept;
  }
}

```

Output for R4:

```

user@R4# show interfaces
fe-0/0/1 {
  unit 0 {
    family inet {
      address 10.0.2.1/30;
    }
  }
}
fe-1/0/0 {
  unit 0 {
    family inet {
      address 10.0.8.9/30;
    }
  }
}
fe-1/1/0 {
  unit 0 {
    family inet {
      address 10.0.8.6/30;
    }
  }
}

user@R4# show protocols ospf
area 0.0.0.0 {
  interface fe-0/0/1.0;
}
area 0.0.0.4 {
  interface fe-0/1/0.0;
  interface fe-1/0/0.0;
}

```

Output for R5:

```
user@R5# show interfaces
fe-1/1/0 {
  unit 0 {
    family inet {
      address 10.0.8.5/30;
    }
  }
}

user@R5# show protocols ospf
area 0.0.0.4 {
  interface fe-1/1/0.0;
}
```

Output for R6:

```
user@R6# show interfaces
fe-1/0/0 {
  unit 0 {
    family inet {
      address 10.0.8.10/30;
    }
  }
}

user@R6# show protocols ospf
area 0.0.0.4 {
  interface fe-1/0/0.0;
}
```

To confirm your OSPFv3 configuration, enter the **show interfaces**, **show policy-options**, and **show protocols ospf3** commands on the appropriate device.

Verification

Confirm that the configuration is working properly.

- [Verifying the OSPF Database on page 4562](#)
- [Verifying the Routing Table on page 4562](#)

Verifying the OSPF Database

Purpose Verify that the OSPF database for the devices in area 4 includes the interarea route that we are advertising from R3. Any other routes from area 3 should not be advertised into area 4, so those entries should age out or no longer be present in the OSPF database.

Action From operational mode, enter the **show ospf database netsummary area 0.0.0.4** command for OSPFv2, and enter the **show ospf3 database inter-area-prefix area 0.0.0.4** command for OSPFv3.

Verifying the Routing Table

Purpose Verify that the specified route is included in R4's, R5's, or R6's routing table. Any other routes from area 3 should not be advertised into area 4.

Action From operational mode, enter the **show route protocol ospf** command for both OSPFv2 and OSPFv3.

- Related Documentation**
- [OSPF Overview on page 4386](#)
 - *Understanding OSPF Configurations*
 - *Routing Policy Match Conditions in the Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*
 - *Actions in Routing Policy Terms in the Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*

PART 68

Monitoring and Troubleshooting

- [Routine Monitoring on page 4567](#)
- [OSPF Monitoring Configuration on page 4569](#)

Routine Monitoring

- [Monitoring OSPF Routing Information on page 4567](#)

Monitoring OSPF Routing Information

Purpose Use the monitoring functionality to monitor OSPF routing information on routing devices.

Action To view OSPF routing information in the CLI, enter the following CLI commands:

- `show ospf neighbor`
- `show ospf interface`
- `show ospf statistics`

Related Documentation

- [show \(ospf | ospf3\) interface on page 4643](#)
- [clear \(ospf | ospf3\) neighbor on page 4623](#)
- [show \(ospf | ospf3\) statistics on page 4671](#)

OSPF Monitoring Configuration

- [Example: Configuring OSPF Trace Options on page 4569](#)

Example: Configuring OSPF Trace Options

- [Tracing OSPF Protocol Traffic on page 4569](#)
- [Example: Tracing OSPF Protocol Traffic on page 4570](#)

Tracing OSPF Protocol Traffic

Tracing operations record detailed messages about the operation of OSPF. You can trace OSPF protocol traffic to help debug OSPF protocol issues. When you trace OSPF protocol traffic, you specify the name of the file and the type of information you want to trace.

You can specify the following OSPF protocol-specific trace options:

- **database-description**—All database description packets, which are used in synchronizing the OSPF topological database
- **error**—OSPF error packets
- **event**—OSPF state transitions
- **flooding**—Link-state flooding packets
- **graceful-restart**—Graceful-restart events
- **hello**—Hello packets, which are used to establish neighbor adjacencies and to determine whether neighbors are reachable
- **ldp-synchronization**—Synchronization events between OSPF and LDP
- **lsa-ack**—Link-state acknowledgment packets, which are used in synchronizing the OSPF topological database
- **lsa-analysis**—Link-state analysis. Specific to the Juniper Networks implementation of OSPF, Junos OS performs LSA analysis before running the shortest-path-first (SPF) algorithm. LSA analysis helps to speed the calculations performed by the SPF algorithm.
- **lsa-request**—Link-state request packets, which are used in synchronizing the OSPF topological database
- **lsa-update**—Link-state updates packets, which are used in synchronizing the OSPF topological database

- **nsr-synchronization**—Nonstop routing synchronization events
- **on-demand**—Trace demand circuit extensions
- **packet-dump**—Dump the contents of selected packet types
- **packets**—All OSPF packets
- **restart-signaling**—(OSPFv2 only) Restart-signaling graceful restart events
- **spf**—Shortest path first (SPF) calculations

You can optionally specify one or more of the following flag modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted



NOTE: Use the **detail** flag modifier with caution as it might cause the CPU to become very busy.

Global tracing options are inherited from the configuration set by the **traceoptions** statement at the **[edit routing-options]** hierarchy level. You can override the following global trace options for the OSPF protocol using the **traceoptions flag** statement included at the **[edit protocols ospf]** hierarchy level:

- **all**—All tracing operations
- **general**—All normal operations and routing table changes (a combination of the normal and route trace operations)
- **normal**—Normal events
- **policy**—Policy processing
- **route**—Routing information
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing



NOTE: Use the trace flag **all** with caution as it might cause the CPU to become very busy.

Example: Tracing OSPF Protocol Traffic

This example shows how to trace OSPF protocol traffic.

- [Requirements on page 4571](#)
- [Overview on page 4571](#)

- [Configuration on page 4572](#)
- [Verification on page 4575](#)

Requirements

This example assumes that OSPF is properly configured and running in your network, and you want to trace OSPF protocol traffic for debugging purposes.

Overview

You can trace OSPF protocol traffic to help debug OSPF protocol issues. When you trace OSPF protocol traffic, you specify the name of the file and the type of information you want to trace. All files are placed in a directory on the routing device's hard disk. On M Series and T Series routers, trace files are stored in the /var/log directory.

This example shows a few configurations that might be useful when debugging OSPF protocol issues. The verification output displayed is specific to each configuration.



TIP: To keep track of your log files, create a meaningful and descriptive name so it is easy to remember the content of the trace file. We recommend that you place global routing protocol tracing output in the file `routing-log`, and OSPF tracing output in the file `ospf-log`.

In the first example, you globally enable tracing operations for all routing protocols that are actively running on your routing device to the file `routing-log`. With this configuration, you keep the default settings for the trace file size and the number of trace files. After enabling global tracing operations, you enable tracing operations to provide detailed information about OSPF packets, including link-state advertisements, requests, and updates, database description packets, and hello packets to the file `ospf-log`, and you configure the following options:

- **size**—Specifies the maximum size of each trace file, in KB, MB, or GB. In this example, you configure 10 KB as the maximum size. When the file reaches its maximum size, it is renamed with a .0 extension. When the file again reaches its maximum size, it is renamed with a .1 extension, and the newly created file is renamed with a .0 extension. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option. You specify **k** for KB, **m** for MB, and **g** for GB. By default, the trace file size is 128 KB. The file size range is 10 KB through the maximum file size supported on your system.
- **files**—Specifies the maximum number of trace files. In this example, you configure a maximum of 5 trace files. When a trace file reaches its maximum size, it is renamed with a .0 extension, then a .1 extension, and so on until the maximum number of trace files is reached. When the maximum number of files is reached, the oldest trace file is overwritten. If you specify a maximum number of files, you must also specify a maximum file size with the **size** option. By default, there are 10 files. The range is 2 through 1000 files.

In the second example, you trace all SPF calculations to the file `ospf-log` by including the `spf` flag. You keep the default settings for the trace file size and the number of trace files.

In the third example, you trace the creation, receipt, and retransmission of all LSAs to the file `ospf-log` by including the `lsa-request`, `lsa-update`, and `lsa-ack` flags. You keep the default settings for the trace file size and the number of trace files.

Configuration

- [Configuring Global Tracing Operations and Tracing OSPF Packet Information on page 4572](#)
- [Tracing SPF Calculations on page 4573](#)
- [Tracing Link-State Advertisements on page 4574](#)

Configuring Global Tracing Operations and Tracing OSPF Packet Information

CLI Quick Configuration

To quickly enable global tracing operations for all routing protocols actively running on your routing device and to trace detailed information about OSPF packets, copy the following commands and paste them into the CLI.

```
[edit]
set routing-options traceoptions file routing-log
set protocols ospf traceoptions file ospf-log
set protocols ospf traceoptions file files 5 size 10k
set protocols ospf traceoptions flag lsa-ack
set protocols ospf traceoptions flag database-description
set protocols ospf traceoptions flag hello
set protocols ospf traceoptions flag lsa-update
set protocols ospf traceoptions flag lsa-request
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Modifying the Junos OS Configuration* in the *CLI User Guide*.

To configure global routing tracing operations and tracing operations for OSPF packets:

1. Configure tracing at the routing options level to collect information about the active routing protocols on your routing device.

```
[edit]
user@host# edit routing-options traceoptions
```

2. Configure the filename for the global trace file.

```
[edit routing-options traceoptions]
user@host# set file routing-log
```

3. Configure the filename for the OSPF trace file.



NOTE: To specify OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
```

- ```

user@host# edit protocols ospf traceoptions
user@host# set file ospf-log

```
4. Configure the maximum number of trace files.

```

[edit protocols ospf traceoptions]
user@host# set file files 5

```
  5. Configure the maximum size of each trace file.

```

[edit protocols ospf traceoptions]
user@host# set file size 10k

```
  6. Configure tracing flags.

```

[edit protocols ospf traceoptions]
user@host# set flag lsa-ack
user@host# set flag database-description
user@host# set flag hello
user@host# set flag lsa-update
user@host# set flag lsa-request

```
  7. If you are done configuring the device, commit the configuration.

```

[edit protocols ospf traceoptions]
user@host# commit

```

**Results** Confirm your configuration by entering the **show routing-options** and the **show protocols ospf** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@host# show routing-options
traceoptions {
 file routing-log;
}

user@host# show protocols ospf
traceoptions {
 file ospf-log size 10k files 5;
 flag lsa-ack;
 flag database-description;
 flag hello;
 flag lsa-update;
 flag lsa-request;
}

```

To confirm your OSPFv3 configuration, enter the **show routing-options** and the **show protocols ospf3** commands.

### *Tracing SPF Calculations*

**CLI Quick Configuration** To quickly trace SPF calculations, copy the following commands and paste them into the CLI.

```

[edit]
set protocols ospf traceoptions file ospf-log
set protocols ospf traceoptions flag spf

```

**Step-by-Step Procedure** To configure SPF tracing operations for OSPF:

1. Configure the filename for the OSPF trace file.



**NOTE:** To specify OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# edit protocols ospf traceoptions
user@host# set file ospf-log
```

2. Configure the SPF tracing flag.

```
[edit protocols ospf traceoptions]
user@host# set flag spf
```

3. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf traceoptions]
user@host# commit
```

**Results** Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
traceoptions {
 file ospf-log ;
 flag spf;
}
```

To confirm your OSPFv3 configuration, enter the `show protocols ospf3` command.

#### *Tracing Link-State Advertisements*

**CLI Quick Configuration** To quickly trace the creation, receipt, and retransmission of all LSAs, copy the following commands and paste them into the CLI.

```
[edit]
set protocols ospf traceoptions file ospf-log
set protocols ospf traceoptions flag lsa-request
set protocols ospf traceoptions flag lsa-update
set protocols ospf traceoptions flag lsa-ack
```

**Step-by-Step Procedure** To configure link-state advertisement tracing operations for OSPF:

1. Configure the filename for the OSPF trace file.



**NOTE:** To specify OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# edit protocols ospf traceoptions
user@host# set file ospf-log
```

2. Configure the link-state advertisement tracing flags.

```
[edit protocols ospf traceoptions]
user@host# set flag lsa-request
user@host# set flag lsa-update
user@host# set flag lsa-ack
```

3. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf traceoptions]
user@host# commit
```

**Results** Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
traceoptions {
 file ospf-log;
 flag lsa-request;
 flag lsa-update;
 flag lsa-ack;
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

### Verification

Confirm that the configuration is working properly.

#### Verifying Trace Operations

**Purpose** Verify that the Trace options field displays the configured trace operations, and verify that the Trace file field displays the location on the routing device where the file is saved, the name of the file to receive the output of the tracing operation, and the size of the file.

**Action** From operational mode, enter the **show ospf overview extensive** command for OSPFv2, and enter the **show ospf3 overview extensive** command for OSPFv3.

**Related Documentation**

- [OSPF Overview on page 4386](#)
- *Understanding OSPF Configurations*
- *Tracing and Logging Junos OS Operations* in the *Junos OS Administration Library for Routing Devices*
- *Example: Tracing Global Routing Protocol Operations* in the *Junos OS Routing Protocols Library for Routing Devices*





## PART 69

# Configuration Statements and Operational Commands

- [OSPF Configuration Statements on page 4579](#)
- [OSPF Operational Commands on page 4617](#)



# OSPF Configuration Statements

- [area](#) on page 4580
- [area-range](#) on page 4582
- [authentication](#) on page 4584
- [context-identifier \(Protocols OSPF\)](#) on page 4585
- [bfd-liveness-detection](#) on page 4586
- [database-protection](#) on page 4590
- [disable \(OSPF\)](#) on page 4592
- [export](#) on page 4594
- [external-preference](#) on page 4595
- [graceful-restart](#) on page 4596
- [import](#) on page 4598
- [interface \(Protocols OSPF\)](#) on page 4599
- [no-nssa-abr](#) on page 4601
- [no-rfc-1583](#) on page 4602
- [ospf](#) on page 4603
- [ospf3](#) on page 4603
- [overload \(Protocols OSPF\)](#) on page 4604
- [preference \(Protocols OSPF\)](#) on page 4606
- [prefix-export-limit \(Protocols OSPF\)](#) on page 4607
- [reference-bandwidth \(Protocols OSPF\)](#) on page 4608
- [rib-group \(Protocols OSPF\)](#) on page 4609
- [topology \(OSPF\)](#) on page 4610
- [traceoptions \(Protocols OSPF\)](#) on page 4611
- [traffic-engineering \(OSPF\)](#) on page 4614

## area

```
Syntax area area-id {
 interface interface-name {
 no-eligible-remote-backup;
 passive;
 topology (ipv4-multicast | name) {
 disable;
 }
 }
 virtual-link neighbor-id router-id transit-area area-id {
 topology (ipv4-multicast | name) {
 disable;
 }
 }
 }
```

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols (**ospf** | ospf3)],  
 [edit logical-systems *logical-system-name* protocols ospf3 realm (ipv4-unicast |  
 ipv4-multicast | ipv6-multicast)],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols  
 (**ospf** | ospf3)],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols  
 ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],  
 [edit protocols (**ospf** | ospf3)],  
 [edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],  
 [edit routing-instances *routing-instance-name* protocols (**ospf** | ospf3)],  
 [edit routing-instances *routing-instance-name* protocols ospf3 realm (ipv4-unicast |  
 ipv4-multicast | ipv6-multicast)]

**Release Information** Statement introduced before Junos OS Release 7.4.  
 Statement introduced in Junos OS Release 9.0 for EX Series switches.  
 Support for the **realm** statement introduced in Junos OS Release 9.2.  
 Support for the **realm** statement introduced in Junos OS Release 9.2 for EX Series switches.  
 Statement introduced in Junos OS Release 11.3 for the QFX Series.  
 Support for the **no-eligible-remote-backup** statement introduced in Junos OS Release 15.1.

**Description** Specify the area identifier for this routing device to use when participating in OSPF routing. All routing devices in an area must use the same area identifier to establish adjacencies.

Specify multiple **area** statements to configure the routing device as an area border router. An area border router does not automatically summarize routes between areas. Use the **area-range** statement to configure route summarization. By definition, an area border router must be connected to the backbone area either through a physical link or through a virtual link. To create a virtual link, include the **virtual-link** statement.

To specify that the routing device is directly connected to the OSPF backbone, include the **area 0.0.0.0** statement.

All routing devices on the backbone must be contiguous. If they are not, use the **virtual-link** statement to create the appearance of connectivity to the backbone.

You can also configure any interface that belongs to one or more topologies to advertise the direct interface addresses without actually running OSPF on that interface. By default, OSPF must be configured on an interface in order for direct interface addresses to be advertised as interior routes.



**NOTE:** If you configure an interface with the **passive** statement, it applies to all the topologies to which the interface belongs. You cannot configure an interface as passive for only one specific topology and have it remain active for any other topologies to which it belongs.

|                                 |                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | <b>area-id</b> —Area identifier. The identifier can be up to 32 bits. It is common to specify the area number as a simple integer or an IP address. Area number <b>0.0.0.0</b> is reserved for the OSPF backbone area.              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding OSPF Areas on page 4397</a></li> <li>• <a href="#">Understanding Multiple Address Families for OSPFv3 on page 4446</a></li> <li>• <i>virtual-link</i></li> </ul> |

## area-range

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <b>area-range</b> <i>network/mask-length</i> <exact> <override-metric <i>metric</i> > <restrict>;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) <b>area</b> <i>area-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) area <i>area-id</i> nssa],</p> <p>[edit logical-systems <i>logical-system-name</i> realm (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) <b>area</b> <i>area-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) area <i>area-id</i> nssa],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> realm (ipv4-unicast   ipv4-multicast   ipv6-multicast) <b>area</b> <i>area-id</i>],</p> <p>[edit protocols (ospf   ospf3) <b>area</b> <i>area-id</i>],</p> <p>[edit protocols (ospf   ospf3) area <i>area-id</i> nssa],</p> <p>[edit protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast) <b>area</b> <i>area-id</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) <b>area</b> <i>area-id</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) area <i>area-id</i> nssa],</p> <p>[edit routing-instances <i>routing-instance-name</i> realm (ipv4-unicast   ipv4-multicast   ipv6-multicast) <b>area</b> <i>area-id</i>]</p> |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>         | <p>(Area border routers only) For an area, summarize a range of IP addresses when sending summary link advertisements (within an area). To summarize multiple ranges, include multiple <b>area-range</b> statements.</p> <p>For a not-so-stubby area (NSSA), summarize a range of IP addresses when sending NSSA link-state advertisements. The specified prefixes are used to aggregate external routes learned within the area when the routes are advertised to other areas. To specify multiple prefixes, include multiple <b>area-range</b> statements. All external routes learned within the area that do not fall into one of the prefixes are advertised individually to other areas.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Default</b>             | By default, area border routing devices do not summarize routes being sent from one area to other areas, but rather send all routes explicitly.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>             | <p><b>exact</b>—(Optional) Summarization of a route is advertised only when an exact match is made with the configured summary range.</p> <p><b>mask-length</b>—Number of significant bits in the network mask.</p> <p><b>network</b>—IP address. You can specify one or more IP addresses.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

**override-metric *metric***—(Optional) Override the metric for the IP address range and configure a specific metric value.

**restrict**—(Optional) Do not advertise the configured summary. This hides all routes that are contained within the summary, effectively creating a route filter.

**Range:** 1 through 16,777,215

|                           |                                                             |
|---------------------------|-------------------------------------------------------------|
| <b>Required Privilege</b> | routing—To view this statement in the configuration.        |
| <b>Level</b>              | routing-control—To add this statement to the configuration. |

|                              |                                                                                                                                                                                    |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Example: Summarizing Ranges of Routes in OSPF Link-State Advertisements Sent into the Backbone Area on page 4452</a></li></ul> |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## authentication

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>authentication {<br/>    md5 key-identifier {<br/>        key key-value;<br/>        start-time YYYY-MM-DD.hh:mm;<br/>    }<br/>    simple-password key;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | <pre>[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> <a href="#">interface</a><br/>  <i>interface-name</i>],<br/>[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> virtual-link],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br/>  ospf area <i>area-id</i> <a href="#">interface</a> <i>interface-name</i>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br/>  ospf area <i>area-id</i> virtual-link],<br/>[edit protocols ospf area <i>area-id</i> <a href="#">interface</a> <i>interface-name</i>],<br/>[edit protocols ospf area <i>area-id</i> virtual-link],<br/>[edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> <a href="#">interface</a><br/>  <i>interface-name</i>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> virtual-link]</pre> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | <p>Configure an authentication key (password). Neighboring routers use the password to verify the authenticity of packets sent from this interface.</p> <p>All routers that are connected to the same IP subnet must use the same authentication scheme and password.</p> <p>The remaining statements are explained separately. See <a href="#">CLI Explorer</a>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding OSPFv2 Authentication</a></li><li>• <a href="#">Example: Configuring MD5 Authentication for OSPFv2 Exchanges</a></li><li>• <a href="#">Example: Configuring a Transition of MD5 Keys on an OSPFv2 Interface</a></li><li>• <a href="#">Example: Configuring Simple Authentication for OSPFv2 Exchanges</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |



---

## context-identifier (Protocols OSPF)

---

|                                 |                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | context-identifier <i>identifier</i>                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) <b>area</b> <i>area-id</i> ],<br>[edit protocols (ospf   ospf3) <b>area</b> <i>area-id</i> ]                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.4.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| <b>Description</b>              | Configure OSPF context-identifier information.                                                                                                                                                  |
| <b>Options</b>                  | <b>identifier</b> —IPv4 address that defines a protection pair. The context identifier is manually configured on both the primary and protector provider edge (PE) devices.                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show ospf context-identifier on page 4633</a></li></ul>                                                                                     |

## bfd-liveness-detection

**Syntax** `bfd-liveness-detection {`  
     `authentication {`  
         `algorithm` *algorithm-name*;  
         `key-chain` *key-chain-name*;  
         `loose-check`;  
     `}`  
     `detection-time {`  
         `threshold` *milliseconds*;  
     `}`  
     `full-neighbors-only`  
     `minimum-interval` *milliseconds*;  
     `minimum-receive-interval` *milliseconds*;  
     `multiplier` *number*;  
     `no-adaptation`;  
     `transmit-interval {`  
         `minimum-interval` *milliseconds*;  
         `threshold` *milliseconds*;  
     `}`  
     `version` (1 | automatic);  
`}`

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols (ospf | ospf3) area *area-id* **interface** *interface-name*],  
 [edit logical-systems *logical-system-name* protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **interface** *interface-name*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (ospf | ospf3) area *area-id* **interface** *interface-name*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **interface** *interface-name*],  
 [edit protocols (ospf | ospf3) area *area-id* **interface** *interface-name*],  
 [edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **interface** *interface-name*],  
 [edit routing-instances *routing-instance-name* protocols (ospf | ospf3) area *area-id* **interface** *interface-name*],  
 [edit routing-instances *routing-instance-name* protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **interface** *interface-name*]

**Release Information** Statement introduced before Junos OS Release 7.4.  
 Statement introduced in Junos OS Release 9.0 for EX Series switches.  
**detection-time threshold** and **transmit-interval threshold** options added in Junos OS Release 8.2.  
 Support for logical systems introduced in Junos OS Release 8.3.  
**no-adaptation** option introduced in Junos OS Release 9.0.  
**no-adaptation** option introduced in Junos OS Release 9.0 for EX Series switches.  
 Support for OSPFv3 introduced in Junos OS Release 9.3.  
 Support for OSPFv3 introduced in Junos OS Release 9.3 for EX Series switches.  
**full-neighbors-only** option introduced in Junos OS Release 9.5.  
**full-neighbors-only** option introduced in Junos OS Release 9.5 for EX Series switches.

**authentication algorithm**, **authentication key-chain**, and **authentication loose-check** options introduced in Junos OS Release 9.6.

Statement introduced in Junos OS Release 12.1 for the QFX Series.

**Description** Configure bidirectional failure detection timers and authentication for OSPF.

The remaining statements are explained separately. See [CLI Explorer](#).

**Options**    **authentication algorithm *algorithm-name***—Configure the algorithm used to authenticate the specified BFD session: **simple-password**, **keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, or **meticulous-keyed-sha-1**.

**authentication key-chain *key-chain-name***—Associate a security key with the specified BFD session using the name of the security keychain. The name you specify must match one of the keychains configured in the **authentication-key-chains key-chain** statement at the **[edit security]** hierarchy level.

**authentication loose-check**—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication may not be configured at both ends of the BFD session.

**detection-time threshold *milliseconds***—Configure a threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

**full-neighbors-only**—Establish BFD sessions only for OSPF neighbors in the full state. The default behavior is to establish BFD sessions for all OSPF neighbors.

**minimum-interval *milliseconds***—Configure the minimum interval after which the local routing device transmits a hello packet and then expects to receive a reply from the neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit and receive intervals separately using the **transmit-interval minimum-interval** and **minimum-receive-interval** statements.

**Range:** 1 through 255,000 milliseconds

**minimum-receive-interval *milliseconds***—Configure the minimum interval after which the routing device expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the **minimum-interval** statement.

**Range:** 1 through 255,000 milliseconds

**multiplier *number***—Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

**Range:** 1 through 255

**Default:** 3

**no-adaptation**—Specify that BFD sessions should not adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

**transmit-interval threshold *milliseconds***—Configure the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

**Range:** 0 through 4,294,967,295 ( $2^{32} - 1$ )

**transmit-interval minimum-interval *milliseconds***—Configure the minimum interval at which the routing device transmits hello packets to a neighbor with which it has established

a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit interval using the **minimum-interval** statement.

**Range:** 1 through 255,000

**version**—Configure the BFD version to detect: **1** (BFD version 1) or **automatic** (autodetect the BFD version).

**Default:** **automatic**

|                                 |                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration. |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------|

|                              |                                                                                                                                                                                                          |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring BFD for OSPF on page 4484</a></li><li>• <a href="#">Example: Configuring BFD Authentication for OSPF on page 4488</a></li></ul> |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## database-protection

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>database-protection {<br/>    ignore-count <i>number</i>;<br/>    ignore-time <i>seconds</i>;<br/>    maximum-lsa <i>number</i>;<br/>    reset-time <i>seconds</i>;<br/>    warning-only;<br/>    warning-threshold <i>percent</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>     | <pre>[edit protocols (<i>ospf</i>   <i>ospf3</i>)],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br/>  (<i>ospf</i>   <i>ospf3</i>)],<br/>[edit routing-instances <i>routing-instance-name</i> protocols (<i>ospf</i>   <i>ospf3</i>)],<br/>[edit routing-instances <i>routing-instance-name</i> protocols <i>ospf3</i> realm (<i>ipv4-unicast</i>  <br/>  <i>ipv4-multicast</i>   <i>ipv6-unicast</i>   <i>ipv6-multicast</i>)]</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 10.2.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>         | Configure the maximum number of link-state advertisements (LSAs) that are not generated by the router or switch in a given OSPF instance.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Default</b>             | By default, OSPF database protection is not enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>             | <p><b>ignore-count <i>number</i></b>—Configure the number of times the database can enter the ignore state. When the ignore count is exceeded, the database enters the isolate state.</p> <p><b>Range:</b> 1 through 32</p> <p><b>Default:</b> 5</p> <p><b>ignore-time <i>seconds</i></b>—Configure the time the database must remain in the ignore state before it resumes regular operations (enters retry state).</p> <p><b>Range:</b> 30 through 3,600 seconds</p> <p><b>Default:</b> 300 seconds</p> <p><b>maximum-lsa <i>number</i></b>—Configure the maximum number of LSAs whose advertising router ID is different from the local router ID in a given OSPF instance. This includes external LSAs as well as LSAs with any scope, such as the link, area, and autonomous system (AS). This value is mandatory.</p> <p><b>Range:</b> 1 through 1,000,000</p> <p><b>Default:</b> None</p> <p><b>reset-time <i>seconds</i></b>—Configure the time period during which the database must operate without being in the ignore or isolate state before it is reset to a normal operating state.</p> <p><b>Range:</b> 60 through 86,400 seconds</p> <p><b>Default:</b> 600 seconds</p> |

**warning-only**—Specify that only a warning should be issued when the maximum LSA number is exceeded. If configured, no other action is taken against the database.

**warning-threshold *percent***—Configure the percentage of the maximum number of LSAs to be exceeded before a warning message is logged.

**Range:** 30 through 100 percent

**Default:** 75 percent

|                           |                                                             |
|---------------------------|-------------------------------------------------------------|
| <b>Required Privilege</b> | routing—To view this statement in the configuration.        |
| <b>Level</b>              | routing-control—To add this statement to the configuration. |

|                              |                                                                                                                                                                                                   |
|------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Understanding OSPF Database Protection on page 4523</a></li><li>• <a href="#">Configuring OSPF Database Protection on page 4524</a></li></ul> |
|------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## disable (OSPF)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | disable;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols (<b>ospf</b>   ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) <b>area</b> <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf <b>area</b> <i>area-id</i> <b>peer-interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) virtual-link],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast) <b>area</b> <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<b>ospf</b>   ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) <b>area</b> <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) virtual-link],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instances</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast) <b>area</b> <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit protocols (<b>ospf</b>   ospf3)],</p> <p>[edit protocols (ospf   ospf3) <b>area</b> <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit protocols (ospf   ospf3) virtual-link],</p> <p>[edit protocols ospf <b>area</b> <i>area-id</i> <b>peer-interface</b> <i>interface-name</i>],</p> <p>[edit protocols ospf <i>area</i> <i>area-id</i> virtual-link neighbor-id <i>router-id</i> transit-area <i>area-id</i>],</p> <p>[edit protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast) <b>area</b> <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (<b>ospf</b>   ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) <b>area</b> <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) virtual-link],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast) <b>area</b> <i>area-id</i> <b>interface</b> <i>interface-name</i>]</p> |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>         | Disable OSPF, an OSPF interface, or an OSPF virtual link.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |



By default, control packets sent to the remote end of a virtual link must be forwarded using the default topology. In addition, the transit area path consists only of links that are in the default topology. You can disable a virtual link for a configured topology, but not for a default topology. Include the **disable** statement at the **[edit protocols ospf area *area-id* virtual-link neighbor-id router-id transit-area *area-id* topology *name*]** hierarchy level.



**NOTE:** If you disable the virtual link by including the **disable** statement at the **[edit protocols ospf area *area-id* virtual-link neighbor-id router-id transit-area *area-id*]** hierarchy level, you disable the virtual link for all topologies, including the default topology. You cannot disable the virtual link only in the default topology.

|                                 |                                                                                                                                                                  |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>                  | The configured object is enabled (operational) unless explicitly disabled.                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Understanding OSPF Configurations</i></li> <li>• <i>Configuring RSVP and OSPF for LMP Peer Interfaces</i></li> </ul> |

## export

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>export [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols (<a href="#">ospf</a>   ospf3)],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast  </code><br><code>  ipv4-multicast   ipv6-multicast)],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  (<a href="#">ospf</a>   ospf3)],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</code><br><code>[edit protocols (<a href="#">ospf</a>   ospf3)],</code><br><code>[edit protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   ospf3)],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast  </code><br><code>  ipv4-multicast   ipv6-multicast)]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.<br>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Apply one or more policies to routes being exported from the routing table into OSPF.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more policies.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding OSPF Routing Policy on page 4527</a></li><li>• <a href="#">Import and Export Policies for Network Summaries Overview on page 4543</a></li><li>• <a href="#">import on page 4598</a></li><li>• <a href="#">import on page 4598</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## external-preference

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>external-preference <i>preference</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols (<a href="#">ospf</a>   ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit protocols (<a href="#">ospf</a>   ospf3)],</p> <p>[edit protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Set the route preference for OSPF external routes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <p><b><i>preference</i></b>—Preference value.</p> <p><b>Range:</b> 0 through 4,294,967,295 (<math>2^{32} - 1</math>)</p> <p><b>Default:</b> 150</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Controlling OSPF Route Preferences on page 4468</a></li> <li>• <a href="#">preference on page 4606</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

## graceful-restart

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre> graceful-restart {   disable;   helper-disable (standard   restart-signaling   both);   no-strict-lsa-checking;   notify-duration <i>seconds</i>;   restart-duration <i>seconds</i>; } </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>     | <pre> [edit logical-systems <i>logical-system-name</i> protocols (<b>ospf</b>   ospf3)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<b>ospf</b>   ospf3)], [edit protocols (<b>ospf</b>   ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols ospf] </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Support for the <b>no-strict-lsa-checking</b> statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the helper mode <b>standard</b>, <b>restart-signaling</b>, and <b>both</b> options introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>         | Configure graceful restart for OSPF.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>             | <p><b>disable</b>—Disable graceful restart for OSPF.</p> <p><b>helper-disable (standard   restart-signaling  both)</b>—Disable helper mode for graceful restart. When helper mode is disabled, a device cannot help a neighboring device that is attempting to restart. Beginning with Junos OS Release 11.4, you can configure restart signaling-based helper mode for OSPFv2 graceful restart configurations. The <b>standard</b>, <b>restart-signaling</b>, and <b>both</b> options are only supported for OSPFv2. Specify <b>standard</b> to disable helper mode for standard graceful restart (based on RFC 3623). Specify <b>restart-signaling</b> to disable helper mode for restart signaling-based graceful restart (based on RFC 4811, RFC 4812, and RFC 4813). Specify <b>both</b> to disable helper mode for both standard and restart signaling-based graceful restart. The last committed statement takes precedence over the previously configured statement.</p> <p><b>Default:</b> Helper mode is enabled by default. For OSPFv2, both standard and restart-signaling based helper modes are enabled by default.</p> <p><b>no-strict-lsa-checking</b>—Disable strict OSPF link-state advertisement (LSA) checking to prevent the termination of graceful restart by a helping router. LSA checking is enabled by default.</p> |



**NOTE:** The **helper-disable** statement and the **no-strict-lsa-checking** statement cannot be configured at the same time. If you attempt to configure both statements at the same time, the routing device displays a warning message when you enter the **show protocols (ospf | ospf3)** command.

**notify-duration *seconds***—Estimated time needed to send out purged grace LSAs over all the interfaces.

**Range:** 1 through 3600 seconds

**Default:** 30 seconds

**restart-duration *seconds***—Estimated time needed to reacquire a full OSPF neighbor from each area.

**Range:** 1 through 3600 seconds

**Default:** 180 seconds

|                           |                                                             |
|---------------------------|-------------------------------------------------------------|
| <b>Required Privilege</b> | routing—To view this statement in the configuration.        |
| <b>Level</b>              | routing-control—To add this statement to the configuration. |

- |                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Graceful Restart for OSPF on page 4497</a></li><li>• <a href="#">Example: Configuring the Helper Capability Mode for OSPFv2 Graceful Restart on page 4501</a></li><li>• <a href="#">Example: Configuring the Helper Capability Mode for OSPFv3 Graceful Restart on page 4504</a></li><li>• <a href="#">Example: Disabling Strict LSA Checking for OSPF Graceful Restart on page 4507</a></li></ul> |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## import

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>import [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols (<a href="#">ospf</a>   ospf3)],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast  </code><br><code>  ipv4-multicast   ipv6-multicast)],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  (<a href="#">ospf</a>   ospf3)],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</code><br><code>[edit protocols (<a href="#">ospf</a>   ospf3)],</code><br><code>[edit protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   ospf3)],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast  </code><br><code>  ipv4-multicast   ipv6-multicast)]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.<br>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Filter OSPF routes from being added to the routing table.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more policies.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding OSPF Routing Policy on page 4527</a></li><li>• <a href="#">Import and Export Policies for Network Summaries Overview on page 4543</a></li><li>• <a href="#">export on page 4594</a></li><li>• <a href="#">export on page 4594</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## interface (Protocols OSPF)

**Syntax** `interface interface-name {`  
     `disable;`  
     `authentication key <key-id identifier>;`  
     `bfd-liveness-detection {`  
         `authentication {`  
             `algorithm algorithm-name;`  
             `key-chain key-chain-name;`  
             `loose-check;`  
         `}`  
         `detection-time {`  
             `threshold milliseconds;`  
         `}`  
         `minimum-interval milliseconds;`  
         `minimum-receive-interval milliseconds;`  
         `transmit-interval {`  
             `threshold milliseconds;`  
             `minimum-interval milliseconds;`  
         `}`  
         `multiplier number;`  
     `}`  
     `dead-interval seconds;`  
     `demand-circuit;`  
     `hello-interval seconds;`  
     `ipsec-sa name;`  
     `interface-type type;`  
     `ldp-synchronization {`  
         `disable;`  
         `hold-time seconds;`  
     `}`  
     `metric metric;`  
     `neighbor address <eligible>;`  
     `no-interface-state-traps;`  
     `passive;`  
     `poll-interval seconds;`  
     `priority number;`  
     `retransmit-interval seconds;`  
     `te-metric metric;`  
     `topology (ipv4-multicast | name) {`  
         `metric metric;`  
     `}`  
     `transit-delay seconds;`  
     `transmit-interval seconds;`  
`}`

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols (ospf | ospf3) *area area-id*],  
 [edit logical-systems *logical-system-name* protocols ospf3 realm (ipv4-unicast |  
 ipv4-multicast | ipv6-multicast) *area area-id*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols  
 (ospf | ospf3) *area area-id*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols  
 ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) *area area-id*],  
 [edit protocols (ospf | ospf3) *area area-id*],

```
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area area-id],
[edit routing-instances routing-instance-name protocols (ospf | ospf3) area area-id],
[edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast |
 ipv4-multicast | ipv6-multicast) area area-id]
```

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Support for the **topology** statement introduced in Junos OS Release 9.0.  
Support for the **topology** statement introduced in Junos OS Release 9.0 for EX Series switches.  
Support for the **realm** statement introduced in Junos OS Release 9.2.  
Support for the **realm** statement introduced in Junos OS Release 9.2 for EX Series switches.  
Support for the **no-interface-state-traps** statement introduced in Junos OS Release 10.3.  
This statement is supported only for OSPFv2.  
Statement introduced in Junos OS Release 11.3 for the QFX Series.

**Description** Enable OSPF routing on a routing device interface.

You must include at least one **interface** statement in the configuration to enable OSPF on the routing device.

**Options** **interface-name**—Name of the interface. Specify the interface by IP address or interface name for OSPFv2, or only the interface name for OSPFv3. Using both the interface name and IP address of the same interface produces an invalid configuration. To configure all interfaces, you can specify **all**. Specifying a particular interface and **all** produces an invalid configuration.



**NOTE:** For nonbroadcast interfaces, specify the IP address of the nonbroadcast interface as **interface-name**.

---

The remaining statements are explained separately.

---



**NOTE:** You cannot run both OSPF and **ethernet-tcc** encapsulation between two Juniper Networks routing devices.

---

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Understanding OSPF Configurations](#)
- [Example: Configuring Multiple Address Families for OSPFv3 on page 4446](#)
- [neighbor](#)



## no-nssa-abr

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-nssa-abr;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols (<a href="#">ospf</a>   ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit protocols (<a href="#">ospf</a>   ospf3)],</p> <p>[edit protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Disable exporting Type 7 link-state advertisements into not-so-stubby-areas (NSSAs) for an autonomous system boundary router (ASBR) or an area border router (ABR).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring OSPF Not-So-Stubby Areas on page 4418</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## no-rfc-1583

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-rfc-1583;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ( <a href="#">ospf</a>   ospf3)],<br>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast  <br>ipv4-multicast   ipv6-multicast)],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>( <a href="#">ospf</a>   ospf3)],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],<br>[edit protocols ( <a href="#">ospf</a>   ospf3)],<br>[edit protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],<br>[edit routing-instances <i>routing-instance-name</i> protocols ( <a href="#">ospf</a>   ospf3)],<br>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast  <br>ipv4-multicast   ipv6-multicast)] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.<br>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Disable compatibility with RFC 1583, <i>OSPF Version 2</i> . If the same external destination is advertised by AS boundary routers that belong to different OSPF areas, disabling compatibility with RFC 1583 can prevent routing loops.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Default</b>                  | Compatibility with RFC 1583 is enabled by default.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control-level—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Disabling OSPFv2 Compatibility with RFC 1583 on page 4428</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |


## ospf

|                                 |                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>ospf { ... }</code>                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols],<br>[edit protocols],<br>[edit routing-instances <i>routing-instance-name</i> protocols]        |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| <b>Description</b>              | Enable OSPF routing on the routing device.<br><br>You must include the <b>ospf</b> statement to enable OSPF on the routing device.                                                                                                                                         |
| <b>Default</b>                  | OSPF is disabled on the routing device.                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Understanding OSPF Configurations</i></li> <li>• <i>[edit protocols ospf] Hierarchy Level</i></li> </ul>                                                                                                                       |

## ospf3

|                                 |                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>ospf3 { ... }</code>                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit protocols],<br>[edit routing-instances <i>routing-instance-name</i> protocols]                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.         |
| <b>Description</b>              | Enable OSPFv3 routing on the routing device.<br><br>You must include the <b>ospf3</b> statement to enable OSPFv3.                                     |
| <b>Default</b>                  | OSPFv3 is disabled on the routing device.                                                                                                             |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Understanding OSPF Configurations</i></li> <li>• <i>[edit protocols ospf3] Hierarchy Level</i></li> </ul> |

## overload (Protocols OSPF)

|                                                                                                                                                                                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                | <pre>overload {     timeout <i>seconds</i>; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>                                                                                                                                                                                       | <p>[edit logical-systems <i>logical-system-name</i> protocols (<b>ospf</b>   ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf <b>topology</b> (default   ipv4-multicast   <i>name</i>)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<b>ospf</b>   ospf3)],</p> <p>[edit logical systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf <b>topology</b> (default   ipv4-multicast   <i>name</i>)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit protocols (<b>ospf</b>   ospf3)],</p> <p>[edit protocols ospf <b>topology</b> (default   ipv4-multicast   <i>name</i>)],</p> <p>[edit protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (<b>ospf</b>   ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf <b>topology</b> (default   ipv4-multicast   <i>name</i>)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> |
| <b>Release Information</b>                                                                                                                                                                                   | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>                                                                                                                                                                                           | <p>Configure the local routing device so that it appears to be overloaded. You might do this when you want the routing device to participate in OSPF routing, but do not want it to be used for transit traffic.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <div>  <p><b>NOTE:</b> Traffic destined to directly attached interfaces continues to reach the routing device.</p> </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                                                                                                                                                                                               | <p><b>timeout <i>seconds</i></b>—(Optional) Number of seconds at which the overloading is reset. If no timeout interval is specified, the routing device remains in overload state until the <b>overload</b> statement is deleted or a timeout is set.</p> <p><b>Range:</b> 60 through 1800 seconds</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

**Default:** 0 seconds



**NOTE:** Multitopology Routing does not support the timeout option.

**Required Privilege Level**  
routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Configuring OSPF to Make Routing Devices Appear Overloaded on page 4472](#)
- *Example: Configuring Multitopology Routing Based on Applications*
- *Example: Configuring Multitopology Routing Based on a Multicast Source*

## preference (Protocols OSPF)

---


|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>preference <i>preference</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols (<a href="#">ospf</a>   ospf3)],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast  </code><br><code>  ipv4-multicast   ipv6-multicast)],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  (<a href="#">ospf</a>   ospf3)],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</code><br><code>[edit protocols (<a href="#">ospf</a>   ospf3)],</code><br><code>[edit protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   ospf3)],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast  </code><br><code>  ipv4-multicast   ipv6-multicast)]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.<br>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Set the route preference for OSPF internal routes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <b><i>preference</i></b> —Preference value.<br><b>Range:</b> 0 through 4,294,967,295 ( $2^{32} - 1$ )<br><b>Default:</b> 10                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Controlling OSPF Route Preferences on page 4468</a></li><li>• <a href="#">external-preference on page 4595</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

## prefix-export-limit (Protocols OSPF)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>prefix-export-limit <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf <a href="#">topology</a> (default   ipv4-multicast   <i>name</i>)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf <a href="#">topology</a> (default   ipv4-multicast   <i>name</i>)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],</p> <p>[edit protocols ospf <a href="#">topology</a> (default   ipv4-multicast   <i>name</i>)],</p> <p>[edit protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf <a href="#">topology</a> (default   ipv4-multicast   <i>name</i>)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Configure a limit to the number of prefixes exported into OSPF.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <p><b><i>number</i></b>—Prefix limit.</p> <p><b>Range:</b> 0 through 4,294,967,295 (<math>2^{32} - 1</math>)</p> <p><b>Default:</b> None</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Limiting the Number of Prefixes Exported to OSPF on page 4457</a></li> <li>• <a href="#">Example: Configuring Multitopology Routing to Provide Redundancy for Multicast Traffic over Separate Network Paths</a></li> <li>• <a href="#">Example: Configuring Multitopology Routing for Class-Based Forwarding of Voice, Video, and Data Traffic</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

- *Understanding Multitopology Routing for Class-Based Forwarding of Voice, Video, and Data Traffic*
- *Understanding Multitopology Routing in Conjunction with PIM*

## reference-bandwidth (Protocols OSPF)

|                                                                                                                                                                                                                                                                                                                                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                      | <code>reference-bandwidth <i>reference-bandwidth</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                             | <pre>[edit logical-systems <i>logical-system-name</i> protocols (<b>ospf</b>   ospf3)], [edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast     ipv4-multicast   ipv6-multicast)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   (<b>ospf</b>   ospf3)], [edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast     ipv4-multicast   ipv6-multicast)], [edit protocols (<b>ospf</b>   ospf3)], [edit protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)], [edit routing-instances <i>routing-instance-name</i> protocols (<b>ospf</b>   ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast     ipv4-multicast   ipv6-multicast)]</pre> |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                         | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                 | <p>Set the reference bandwidth used in calculating the default interface cost. The cost is calculated using the following formula:</p> $\text{cost} = \text{ref-bandwidth} / \text{bandwidth}$                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                                     | <p><b><i>reference-bandwidth</i></b>—Reference bandwidth, in bits per second.</p> <p><b>Range:</b> 9600 through 1,000,000,000,000 bits</p> <p><b>Default:</b> 100 Mbps (100,000,000 bits)</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <div>  <p><b>NOTE:</b> The default behavior is to use the reference-bandwidth value to calculate the cost of OSPF interfaces. You can override this behavior for any OSPF interface by configuring a specific cost with the <b>metric</b> statement.</p> </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                    | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                       | <ul style="list-style-type: none"> <li>• <a href="#">Example: Controlling the Cost of Individual OSPF Network Segments on page 4462</a></li> <li>• <i>metric</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |



## rib-group (Protocols OSPF)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>rib-group group-name;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols (<a href="#">ospf</a>   ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit protocols (<a href="#">ospf</a>   ospf3)],</p> <p>[edit protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Install routes learned from OSPF routing instances into routing tables in the OSPF routing table group.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <b>group-name</b> —Name of the routing table group.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Exporting Specific Routes from One Routing Table Into Another Routing Table</i></li> <li>• <i>Example: Importing Direct and Static Routes Into a Routing Instance</i></li> <li>• <i>Understanding Multiprotocol BGP</i></li> <li>• <a href="#">interface-routes on page 3322</a></li> <li>• <a href="#">rib-group on page 3364</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

## topology (OSPF)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>topology (default   ipv4-multicast   <i>name</i>) {<br/>    spf-options {<br/>        delay <i>milliseconds</i>;<br/>        holddown <i>milliseconds</i>;<br/>        rapid-runs <i>number</i>;<br/>    }<br/>    topology-id <i>number</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | <pre>[edit logical-systems <i>logical-system-name</i> protocols <b>ospf</b>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br/>  <b>ospf</b>],<br/>[edit protocols <b>ospf</b>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols <b>ospf</b>]</pre>                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.0.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Enable a topology for OSPF multitopology routing. You must first configure one or more topologies under the <b>[edit routing-options]</b> hierarchy level.                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <p><b>default</b>—Name of the default topology. This topology is automatically created, and all routes that correspond to it are automatically added to the <b>inet.0</b> routing table. You can modify certain default parameters, such as for the SPF algorithm.</p> <p><b>ipv4-multicast</b>—Name of the topology for IPv4 multicast traffic.</p> <p><b><i>name</i></b>—Name of a topology you configured at the <b>[edit routing-options]</b> hierarchy level to create a topology for a specific type of traffic, such as voice or video.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Multitopology Routing to Provide Redundancy for Multicast Traffic over Separate Network Paths</i></li><li>• <i>Example: Configuring Multitopology Routing for Class-Based Forwarding of Voice, Video, and Data Traffic</i></li><li>• <i>Understanding Multitopology Routing for Class-Based Forwarding of Voice, Video, and Data Traffic</i></li><li>• <i>Understanding Multitopology Routing in Conjunction with PIM</i></li></ul>                                                                                                          |

## traceoptions (Protocols OSPF)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;flag-modifier&gt; &lt;disable&gt;; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>     | <pre>[edit logical-systems <i>logical-system-name</i> protocols (<b>ospf</b>   ospf3)], [edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast     ipv4-multicast   ipv6-multicast)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   (<b>ospf</b>   ospf3)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)], [edit protocols (<b>ospf</b>   ospf3)], [edit protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)], [edit routing-instances <i>routing-instance-name</i> protocols (<b>ospf</b>   ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast     ipv4-multicast   ipv6-multicast)]</pre> |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>         | <p>Configure OSPF protocol-level tracing options.</p> <p>To specify more than one tracing operation, include multiple <b>flag</b> statements.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |



**NOTE:** The **traceoptions** statement is not supported on QFabric systems.

|                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b> | The default OSPF protocol-level tracing options are those inherited from the routing protocols <b>traceoptions</b> statement included at the <b>[edit routing-options]</b> hierarchy level.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b> | <p><b>disable</b>—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>/var/log</b>. We recommend that you place OSPF tracing output in the file <b>ospf-log</b>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and</p> |

so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

**Range:** 2 through 1000 files

**Default:** 10 files

**flag flag**—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

#### OSPF Tracing Flags

- **database-description**—Database description packets, which are used in synchronizing the OSPF and OSPFv3 topological database.
- **error**—OSPF and OSPFv3 error packets.
- **event**—OSPF and OSPFv3 state transitions.
- **flooding**—Link-state flooding packets.
- **graceful-restart**—Graceful-restart events.
- **hello**—Hello packets, which are used to establish neighbor adjacencies and to determine whether neighbors are reachable.
- **ldp-synchronization**—Synchronization events between OSPF and LDP.
- **lsa-ack**—Link-state acknowledgment packets, which are used in synchronizing the OSPF topological database.
- **lsa-analysis**—Link-state analysis. Specific to the Juniper Networks implementation of OSPF, Junos OS performs LSA analysis before running the shortest-path-first (SPF) algorithm. LSA analysis helps to speed the calculations performed by the SPF algorithm.
- **lsa-request**—Link-state request packets, which are used in synchronizing the OSPF topological database.
- **lsa-update**—Link-state updates packets, which are used in synchronizing the OSPF topological database.
- **nsr-synchronization**—Nonstop routing synchronization events.
- **on-demand**—Trace demand circuit extensions.
- **packet-dump**—Content of selected packet types.
- **packets**—All OSPF packets.
- **restart-signaling**—(OSPFv2 only) Restart-signaling graceful restart events.
- **spf**—Shortest-path-first (SPF) calculations.

#### Global Tracing Flags

- **all**—All tracing operations.
- **general**—A combination of the **normal** and **route** trace operations.
- **normal**—All normal operations. If you do not specify this option, only unusual or abnormal operations are traced.
- **policy**—Policy operations and actions.
- **route**—Routing table changes.
- **state**—State transitions.
- **task**—Routing protocol task processing.
- **timer**—Routing protocol timer processing.

**flag-modifier**—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Detailed trace information.
- **receive**—Packets being received.
- **send**—Packets being transmitted.

**no-world-readable**—(Optional) Prevent any user from reading the log file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

**Syntax:** *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 128 KB

**world-readable**—(Optional) Allow any user to read the log file.

|                                 |                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing and trace—To view this statement in the configuration.<br>routing-control and trace-control—To add this statement to the configuration. |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Tracing OSPF Protocol Traffic on page 4570</a></li> </ul>                         |

## traffic-engineering (OSPF)

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>traffic-engineering {<br/>  &lt;advertise-unnumbered-interfaces&gt;;<br/>  &lt;credibility-protocol-preference&gt;;<br/>  ignore-lsp-metrics;<br/>  multicast-rpf-routes;<br/>  no-topology;<br/>  shortcuts {<br/>    lsp-metric-into-summary;<br/>  }<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> protocols ( <a href="#">ospf</a>   ospf3)],<br>[edit protocols ( <a href="#">ospf</a>   ospf3)]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p><b>multicast-rpf-routes</b> option introduced in Junos OS Release 7.5.</p> <p><b>advertise-unnumbered-interfaces</b> option introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for OSPFv3 (<b>ospf3</b>) introduced in Junos OS Release 9.4.</p> <p>Support for OSPFv3 (<b>ospf3</b>) introduced in Junos OS Release 9.4 for EX Series switches.</p> <p><b>credibility-protocol-preference</b> statement introduced in Junos OS Release 9.4.</p> <p><b>credibility-protocol-preference</b> statement introduced in Junos OS Release 9.4 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>         | Enable the OSPF traffic engineering features.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Default</b>             | Traffic engineering support is disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>             | <p><b>advertise-unnumbered-interfaces</b>—(Optional) (OSPFv2 only) Include the link-local identifier in the link-local traffic-engineering link-state advertisement. This statement must be included on both ends of an unnumbered link to allow an ingress LER to update the link in its traffic engineering database and use it for CSPF calculations. The link-local identifier is then used by RSVP to signal unnumbered interfaces as defined in RFC 3477.</p> <p><b>credibility-protocol-preference</b>—(Optional) (OSPFv2 only) Use the configured preference value for OSPF routes to calculate the traffic engineering database credibility value used to select IGP routes. Use this statement to override the default behavior, in which the traffic engineering database prefers IS-IS routes even if OSPF routes are configured with a lower, that is, preferred, preference value. For example, OSPF routes have a default preference value of 10, whereas IS-IS Level 1 routes have a default preference value of 15. When protocol preference is enabled, the credibility value is determined by deducting the protocol preference value from a base value of 512. Using default protocol preference values, OSPF has a credibility value of 502, whereas IS-IS has a credibility value of 497. Because the traffic engineering database prefers IGP routes with the highest credibility value, OSPF routes are now preferred.</p> |

**multicast-rpf-routes**—(Optional) (OSPFv2 only) Install routes for multicast RPF checks into the **inet.2** routing table. The **inet.2** routing table consists of unicast routes used for multicast RPF lookup. RPF is an antispoofing mechanism used to check whether the packet is coming in on an interface that is also sending data back to the packet source.



**NOTE:** You must enable OSPF traffic engineering shortcuts to use the **multicast-rpf-routes** statement. You must not allow LSP advertisements into OSPF when configuring the **multicast-rpf-routes** statement.

**no-topology**—(Optional) (OSPFv2 only) Disable the dissemination of the link-state topology information.

The remaining statements are explained separately.



**CAUTION:** When the OSPF traffic engineering configuration is considerably modified, the routing table entries are deleted and the routing table is recreated. Changes to configuration that can cause this behavior include enabling or disabling:

- Traffic engineering shortcuts
- IGP shortcuts
- LDP tunneling
- Multiprotocol LSP
- Advertise summary metrics
- Multicast RPF routes

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation** • [Example: Enabling OSPF Traffic Engineering Support on page 4513](#)





# OSPF Operational Commands

- `clear (ospf | ospf3) database`
- `clear (ospf | ospf3) database-protection`
- `clear (ospf | ospf3) io-statistics`
- `clear (ospf | ospf3) neighbor`
- `clear (ospf | ospf3) statistics`
- `clear (ospf | ospf3) overload`
- `show (ospf | ospf3) backup coverage`
- `show (ospf | ospf3) backup neighbor`
- `show ospf context-identifier`
- `show ospf database`
- `show (ospf | ospf3) interface`
- `show (ospf | ospf3) io-statistics`
- `show (ospf | ospf3) log`
- `show (ospf | ospf3) neighbor`
- `show (ospf | ospf3) overview`
- `show (ospf | ospf3) route`
- `show (ospf | ospf3) statistics`

```
clear (ospf | ospf3) database
```

- |                                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| List of Syntax                           | <a href="#">Syntax on page 4618</a><br><a href="#">Syntax (EX Series Switch and QFX Series) on page 4618</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Syntax                                   | <pre>clear (ospf   ospf3) database &lt;advertising-router (<i>router-id</i>   self)&gt; &lt;area <i>area-id</i>&gt; &lt;asbrsummary&gt; &lt;external&gt; &lt;instance <i>instance-name</i>&gt; &lt;inter-area-prefix&gt; &lt;inter-area-router&gt; &lt;intra-area-prefix&gt; &lt;link-local&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;lsa-id <i>lsa-id</i>&gt; &lt;netsummary&gt; &lt;network&gt; &lt;nssa&gt; &lt;opaque-area&gt; &lt;purge&gt; &lt;realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)&gt; &lt;router&gt;</pre>                                                                                                                                                                                                                                                                                                                    |
| Syntax (EX Series Switch and QFX Series) | <pre>clear (ospf   ospf3) database &lt;advertising-router (<i>router-id</i>   self)&gt; &lt;area <i>area-id</i>&gt; &lt;asbrsummary&gt; &lt;external&gt; &lt;instance <i>instance-name</i>&gt; &lt;inter-area-prefix&gt; &lt;inter-area-router&gt; &lt;intra-area-prefix&gt; &lt;link-local&gt; &lt;lsa-id <i>lsa-id</i>&gt; &lt;netsummary&gt; &lt;network&gt; &lt;nssa&gt; &lt;opaque-area&gt; &lt;purge&gt; &lt;router&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Release Information                      | <p>Command introduced before Junos OS Release 7.4.</p> <p><b>advertising-router</b> <i>router-id</i>, <b>area</b> <i>area-id</i>, <b>asbrsummary</b>, <b>external</b>, <b>inter-area-prefix</b>, <b>inter-area-router</b>, <b>intra-area-prefix</b>, <b>link-local</b>, <b>lsa-id</b> <i>lsa-id</i>, <b>netsummary</b>, <b>network</b>, <b>nssa</b>, <b>opaque-area</b>, and <b>router</b> options added in Junos OS Release 8.3. You must use the <b>purge</b> command with these options.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>realm</b> option added in Junos OS Release 9.2.</p> <p><b>advertising-router</b> (<i>router-id</i>   <b>self</b>) option added in Junos OS Release 9.5.</p> <p><b>advertising-router</b> (<i>router-id</i>   <b>self</b>) option introduced in Junos OS Release 9.5 for EX Series switches.</p> |

Command introduced in Junos OS Release 11.3 for the QFX Series.

**Description** With the master Routing Engine, delete entries in the Open Shortest Path First (OSPF) link-state advertisement (LSA) database. With the backup Routing Engine, delete the OSPF LSA database and sync the new database with the master Routing Engine. You can also use the **purge** command with any of the options to discard rather than delete the specified LSA entries.



**CAUTION:** This command is useful only for testing. Use it with care, because it causes significant network disruption.

**Options** **none**—Delete all LSAs other than the system's own LSAs, which are regenerated. To resynchronize the database, the system destroys all adjacent neighbors that are in the state **EXSTART** or higher. The neighbors are then reacquired and the databases are synchronized.

**advertising-router** (*router-id* | **self**)—(Optional) Discard entries for the LSA entries advertised by the specified routing device or by this routing device.

**area** *area-id*—(Optional) Discard entries for the LSAs in the specified area.

**asbrsummary**—(Optional) Discard summary AS boundary router LSA entries.

**external**—(Optional) Discard external LSAs.

**instance** *instance-name*—(Optional) Delete or discard entries for the specified routing instance only.

**inter-area-prefix**—(OSPFv3 only) (Optional) Discard interarea prefix LSAs.

**inter-area-router**—(OSPFv3 only) (Optional) Discard interarea router LSAs.

**intra-area-prefix**—(OSPFv3 only) (Optional) Discard intra-area prefix LSAs.

**logical-system** (**all** | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

**link-local**—(Optional) Delete link-local LSAs.

**lsa-id** *lsa-id*—(Optional) Discard the LSA entries with the specified LSA identifier.

**netsummary**—(Optional) Discard summary network LSAs.

**network**—(Optional) Discard network LSAs.

**nssa**—(Optional) Discard not-so-stubby area (NSSA) LSAs.

**opaque-area**—(Optional) Discard opaque area-scope LSAs.

**realm** (**ipv4-multicast** | **ipv4-unicast** | **ipv6-multicast**)—(OSPFv3 only) (Optional) Delete the entries for the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

**router**—(Optional) Discard router LSAs.

**purge**—(Optional) Discard all entries in the link-state advertisement database. All link-state advertisements are set to **MAXAGE** and are flooded. The database is repopulated when the originators of the link-state advertisements receive the **MAXAGE** link-state advertisements and reissue them.

**Required Privilege Level**

clear

**Related Documentation**

- [show ospf database on page 4635](#)
- *show ospf3 database*

**List of Sample Output** [clear ospf database on page 4620](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

[clear ospf database](#)

```
user@host> clear ospf database
```

## clear (ospf | ospf3) database-protection

---

|                                 |                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | clear (ospf   ospf3) database-protection<br><instance <i>instance-name</i> >                                                                                                |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.2.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                             |
| <b>Description</b>              | Clear the Open Shortest Path First (OSPF) link-state database from its isolated state. Reset the ignore count, ignore timer, and reset timer, and resume normal operations. |
| <b>Options</b>                  | <b>instance <i>instance-name</i></b> —(Optional) Clear the OSPF link-state database for the specified routing instance only.                                                |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                       |
| <b>Output Fields</b>            | This command produces no output.                                                                                                                                            |

### Sample Output

#### clear ospf database-protection

```
user@host> clear ospf database-protection
```

## clear (ospf | ospf3) io-statistics

---

|                                                 |                                                                                                                                                                                                                   |
|-------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                           | <a href="#">Syntax on page 4622</a><br><a href="#">Syntax (EX Series Switch and QFX Series) on page 4622</a>                                                                                                      |
| <b>Syntax</b>                                   | clear (ospf   ospf3) io-statistics<br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                        |
| <b>Syntax (EX Series Switch and QFX Series)</b> | clear (ospf   ospf3) io-statistics                                                                                                                                                                                |
| <b>Release Information</b>                      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                          |
| <b>Description</b>                              | Clear Open Shortest Path First (OSPF) input and output statistics.                                                                                                                                                |
| <b>Options</b>                                  | <b>none</b> —Clear OSPF input and output statistics.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>                 | clear                                                                                                                                                                                                             |
| <b>List of Sample Output</b>                    | <a href="#">clear ospf io-statistics on page 4622</a>                                                                                                                                                             |
| <b>Output Fields</b>                            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                             |

### Sample Output

#### clear ospf io-statistics

```
user@host> clear ospf io-statistics
```

## clear (ospf | ospf3) neighbor

|                                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                           | <a href="#">Syntax on page 4623</a><br><a href="#">Syntax (EX Series Switch and QFX Series) on page 4623</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Syntax</b>                                   | <pre>clear (ospf   ospf3) neighbor &lt;area <i>area-id</i>&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;neighbor&gt; &lt;realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Syntax (EX Series Switch and QFX Series)</b> | <pre>clear (ospf   ospf3) neighbor &lt;area <i>area-id</i>&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt; &lt;neighbor&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>                      | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>realm</b> option introduced in Junos OS Release 9.2.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>                              | Tear down Open Shortest Path First (OSPF) neighbor connections.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                                  | <p><b>none</b>—Tear down OSPF connections with all neighbors for all routing instances.</p> <p><b>area <i>area-id</i></b>—(Optional) Tear down neighbor connections for the specified area only.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Tear down neighbor connections for the specified routing instance only.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Tear down neighbor connections for the specified interface only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>neighbor</b>—(Optional) Clear the state of the specified neighbor only.</p> <p><b>realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)</b>—(Optional) (OSPFv3 only) Clear the state of the specified OSPFv3 realm, or address family. Use the <b>realm</b> option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.</p> |
| <b>Required Privilege Level</b>                 | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>                    | <ul style="list-style-type: none"> <li>• <a href="#">show (ospf   ospf3) neighbor on page 4654</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>List of Sample Output</b>                    | <a href="#">clear ospf neighbor on page 4624</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

`clear ospf neighbor`

```
user@host> clear ospf neighbor
```



## clear (ospf | ospf3) statistics

|                                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                           | <a href="#">Syntax on page 4625</a><br><a href="#">Syntax (EX Series Switch and QFX Series) on page 4625</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Syntax</b>                                   | <pre>clear (ospf   ospf3) statistics &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Syntax (EX Series Switch and QFX Series)</b> | <pre>clear (ospf   ospf3) statistics &lt;instance <i>instance-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>                      | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>realm</b> option introduced in Junos OS Release 9.2.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>                              | Clear Open Shortest Path First (OSPF) statistics.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                                  | <p><b>none</b>—Clear OSPF statistics.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Clear statistics for the specified routing instance only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)</b>—(Optional) (OSPFv3 only) Clear statistics for the specified OSPFv3 realm, or address family. Use the <b>realm</b> option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.</p> |
| <b>Required Privilege Level</b>                 | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>                    | <ul style="list-style-type: none"> <li>• <a href="#">show (ospf   ospf3) statistics on page 4671</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>List of Sample Output</b>                    | <a href="#">clear ospf statistics on page 4625</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Output Fields</b>                            | See <a href="#">show (ospf   ospf3) statistics</a> for an explanation of output fields.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## Sample Output

### clear ospf statistics

The following sample output displays OSPF statistics before and after the **clear ospf statistics** command is entered:

```
user@host> show ospf statistics
```

| Packet type | Total | Last 5 seconds |
|-------------|-------|----------------|
|-------------|-------|----------------|

|          | Sent | Received | Sent | Received |
|----------|------|----------|------|----------|
| Hello    | 3254 | 2268     | 3    | 1        |
| DbD      | 41   | 46       | 0    | 0        |
| LSReq    | 8    | 7        | 0    | 0        |
| LSUpdate | 212  | 154      | 0    | 0        |
| LSAck    | 65   | 98       | 0    | 0        |

|                          |   |                    |   |   |
|--------------------------|---|--------------------|---|---|
| DBDs retransmitted       | : | 3, last 5 seconds  | : | 0 |
| LSAs flooded             | : | 12, last 5 seconds | : | 0 |
| LSAs flooded high-prio   | : | 0, last 5 seconds  | : | 0 |
| LSAs retransmitted       | : | 0, last 5 seconds  | : | 0 |
| LSAs transmitted to nbr: | : | 3, last 5 seconds  | : | 0 |
| LSAs requested           | : | 5, last 5 seconds  | : | 0 |
| LSAs acknowledged        | : | 19, last 5 seconds | : | 0 |

|                      |   |   |
|----------------------|---|---|
| Flood queue depth    | : | 0 |
| Total rexmit entries | : | 0 |
| db summaries         | : | 0 |
| lsreq entries        | : | 0 |

Receive errors:  
626 subnet mismatches

user@host> clear ospf statistics

user@host> show ospf statistics

| Packet type | Total |          | Last 5 seconds |          |
|-------------|-------|----------|----------------|----------|
|             | Sent  | Received | Sent           | Received |
| Hello       | 3     | 1        | 3              | 1        |
| DbD         | 0     | 0        | 0              | 0        |
| LSReq       | 0     | 0        | 0              | 0        |
| LSUpdate    | 0     | 0        | 0              | 0        |
| LSAck       | 0     | 0        | 0              | 0        |

|                          |   |                   |   |   |
|--------------------------|---|-------------------|---|---|
| DBDs retransmitted       | : | 0, last 5 seconds | : | 0 |
| LSAs flooded             | : | 0, last 5 seconds | : | 0 |
| LSAs flooded high-prio   | : | 0, last 5 seconds | : | 0 |
| LSAs retransmitted       | : | 0, last 5 seconds | : | 0 |
| LSAs transmitted to nbr: | : | 0, last 5 seconds | : | 0 |
| LSAs requested           | : | 0, last 5 seconds | : | 0 |
| LSAs acknowledged        | : | 0, last 5 seconds | : | 0 |

|                      |   |   |
|----------------------|---|---|
| Flood queue depth    | : | 0 |
| Total rexmit entries | : | 0 |
| db summaries         | : | 0 |
| lsreq entries        | : | 0 |

Receive errors:  
None

## clear (ospf | ospf3) overload

|                                    |                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 4627</a><br><a href="#">Syntax (EX Series Switches) on page 4627</a>                                                                                                                                                                                                                                                                                           |
| <b>Syntax</b>                      | clear (ospf   ospf3) overload<br><instance <i>instance-name</i> ><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                 |
| <b>Syntax (EX Series Switches)</b> | clear (ospf   ospf3) overload<br><instance <i>instance-name</i> >                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                  |
| <b>Description</b>                 | Clear the Open Shortest Path First (OSPF) overload bit and rebuild link-state advertisements (LSAs).                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                     | <p><b>none</b>—Clear the overload bit and rebuild LSAs for all routing instances.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Clear the overload bit and rebuild LSAs for the specified routing instance only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>    | clear                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>List of Sample Output</b>       | <a href="#">clear ospf overload on page 4627</a>                                                                                                                                                                                                                                                                                                                                          |
| <b>Output Fields</b>               | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                     |

## Sample Output

### clear ospf overload

```
user@host> clear ospf overload
```

## show (ospf | ospf3) backup coverage

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show (ospf   ospf3) backup coverage &lt;instance <i>instance-name</i>&gt; &lt; logical-system (all   <i>logical-system-name</i>)&gt; &lt;realm (ipv4-unicast   ipv6-unicast)&gt; &lt;topology <i>topology-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Syntax (QFX Series)</b>      | <pre>show (ospf   ospf3) backup coverage &lt;instance <i>instance-name</i>&gt; &lt;topology <i>topology-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 10.0.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Display information about the level of backup coverage available for all the nodes and prefixes in the network.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <p><b>none</b>—Display information about the level backup coverage for all OSPF routing instances in all logical systems.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Display information about the level of backup coverage for all logical systems or for a specific logical system.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information about the level of backup coverage for a specific OSPF routing instance.</p> <p><b>realm (ipv4-unicast   ipv6-unicast)</b>—(Optional) (OSPFv3 only) Display information about the level of backup coverage for the specific OSPFv3 realm, or address family.</p> <p><b>topology (default   <i>topology-name</i>)</b>—(Optional) (OSPFv2 only) Display information about the level of backup coverage for the specific OSPF topology.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">show (ospf   ospf3) backup lsp</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>List of Sample Output</b>    | <a href="#">show ospf backup coverage on page 4629</a><br><a href="#">show ospf3 backup coverage on page 4629</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Output Fields</b>            | <p><a href="#">Table 346</a> lists the output fields for the <b>show (ospf   ospf3) backup coverage</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

**Table 346: show (ospf | ospf3) backup coverage Output Fields**

| Field Name    | Field Description                                     |
|---------------|-------------------------------------------------------|
| Node Coverage | Information about backup coverage for each OSPF node. |
| Area          | Area number. Area 0.0.0.0 is the backbone.            |

Table 346: show (ospf | ospf3) backup coverage Output Fields (*continued*)

| Field Name             | Field Description                                                                             |
|------------------------|-----------------------------------------------------------------------------------------------|
| <b>Covered Nodes</b>   | Number of nodes for which backup coverage is available.                                       |
| <b>Total Nodes</b>     | Total number of OSPF nodes.                                                                   |
| <b>Route Coverage</b>  | Information about backup coverage for each type of OSPF route.                                |
| <b>Path Type</b>       | Type of OSPF path: <b>Intra</b> , <b>Inter</b> , <b>Ext1</b> , <b>Ext2</b> , and <b>All</b> . |
| <b>Covered Routes</b>  | For each path type, the number of routes for which backup coverage is available.              |
| <b>Total Routes</b>    | For each path type, the total number of configured routes.                                    |
| <b>Percent Covered</b> | For all nodes and for each path type, the percentage for which backup coverage is available.  |

## Sample Output

### show ospf backup coverage

```

user@host> show ospf backup coverage
Topology default coverage:

Node Coverage:

Area Covered Total Percent
 Nodes Nodes Covered
0.0.0.0 4 5 80.00%

Route Coverage:

Path Type Covered Total Percent
 Routes Routes Covered
Intra 8 14 57.14%
Inter 0 0 100.00%
Ext1 0 0 100.00%
Ext2 1 1 100.00%
All 9 15 60.00%

```

### show ospf3 backup coverage

```

user @host > show ospf3 backup coverage
show ospf3 backup coverage
Node Coverage:

Area Covered Total Percent
 Nodes Nodes Covered
0.0.0.0 4 5 80.00%

Route Coverage:

Path Type Covered Total Percent
 Routes Routes Covered

```

|       |   |   |         |
|-------|---|---|---------|
| Intra | 4 | 6 | 66.67%  |
| Inter | 0 | 0 | 100.00% |
| Ext1  | 0 | 0 | 100.00% |
| Ext2  | 1 | 1 | 100.00% |
| All   | 5 | 7 | 71.43%  |

## show (ospf | ospf3) backup neighbor

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--|
| <b>Syntax</b>                   | <pre>show (ospf   ospf3) backup neighbor &lt;area <i>area-id</i>&gt; &lt;instance (default   <i>instance-name</i>)&gt; &lt;logical-system (default   ipv4-multicast   <i>logical-system-name</i>)&gt; &lt;topology (default   ipv4-multicast   <i>topology-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |  |
| <b>Syntax (QFX Series)</b>      | <pre>show (ospf   ospf3) backup neighbor &lt;area <i>area-id</i>&gt; &lt;instance <i>instance-name</i>&gt; &lt;topology (default   ipv4-multicast   <i>topology-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |  |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 10.0.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |  |
| <b>Description</b>              | Display the neighbors through which direct next hops for the backup paths are available.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |  |
| <b>Options</b>                  | <p><b>none</b>—Display all neighbors that have direct next hops for backup paths.</p> <p><b>area <i>area-id</i></b>—(Optional) Display the area information.</p> <p><b>instance (default   <i>instance-name</i>)</b>—(Optional) Display information about the default routing instance or a particular routing instance.</p> <p><b>logical-system (default   ipv4-multicast   <i>logical-system-name</i>)</b>—(Optional) Display information about the default logical system, IPv4 multicast logical system, or a particular logical system.</p> <p><b>topology (default   ipv4-multicast   <i>topology-name</i>)</b>—(OSPFv2 only) (Optional) Display information about the default topology, IPv4 multicast topology, or a particular topology.</p> |  |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">show (ospf   ospf3) backup spf</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |  |
| <b>List of Sample Output</b>    | <a href="#">show ospf backup neighbor on page 4632</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |  |
| <b>Output Fields</b>            | <p><a href="#">Table 347</a> lists the output fields for the <b>show (ospf   ospf3) backup neighbor</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |  |

**Table 347: show (ospf | ospf3) backup neighbor Output Fields**

| Field Name              | Field Description                                 | Level of Output |
|-------------------------|---------------------------------------------------|-----------------|
| Neighbor to Self Metric | Metric from the backup neighbor to the OSPF node. | All levels      |

Table 347: show (ospf |ospf3) backup neighbor Output Fields (*continued*)

| Field Name                     | Field Description                                 | Level of Output |
|--------------------------------|---------------------------------------------------|-----------------|
| <b>Self to Neighbor Metric</b> | Metric from the OSPF node to the backup neighbor. | All levels      |
| <b>Direct next-hop</b>         | Interface and address of the direct next hop.     | All levels      |

## Sample Output

### show ospf backup neighbor

```
user@host> show ospf backup neighbor
Topology default backup neighbors:

Area 0.0.0.5 backup neighbors:

10.0.0.5
 Neighbor to Self Metric: 5
 Self to Neighbor Metric: 5
 Direct next-hop: ge-4/0/0.111 via 10.0.175.5

10.0.0.6
 Neighbor to Self Metric: 5
 Self to Neighbor Metric: 5
 Direct next-hop: ge-4/1/0.110 via 10.0.176.6
```



## show ospf context-identifier

|                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                             | <a href="#">Syntax on page 4633</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 4633</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Syntax</b>                                     | <pre>show ospf context-identifier &lt;brief   detail&gt; &lt;area <i>area-id</i>&gt; &lt;context-id&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Syntax (EX Series Switches and QFX Series)</b> | <pre>show ospf context-identifier &lt;brief   detail&gt; &lt;area <i>area-id</i>&gt; &lt;context-id&gt; &lt;instance <i>instance-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>                        | <p>Command introduced in Junos OS Release 10.4.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>                                | Display the context identifier information processed and advertised by Open Shortest Path First (OSPF) for egress protection.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                                    | <p><b>none</b>—Display information about all context identifiers.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>area <i>area-id</i></b>—(Optional) Display information about the context identifier for the specified area.</p> <p><b>context-id</b>—(Optional) Display information about the specified context identifier.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information about the context identifier for the specified routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                   | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>                      | <ul style="list-style-type: none"> <li><i>egress-protection (Layer 2 circuit)</i> in the <i>Junos OS VPNs Library for Routing Devices</i></li> <li><i>egress-protection (MPLS)</i> in the <i>Junos OS VPNs Library for Routing Devices</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>List of Sample Output</b>                      | <a href="#">show ospf context-identifier on page 4634</a><br><a href="#">show ospf context-identifier detail on page 4634</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Output Fields</b>                              | <a href="#">Table 348</a> lists the output fields for the <b>show ospf context-identifier</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

Table 348: show ospf context-identifier Output Fields

| Field Name                  | Field Description                                                                                                                                                                                                                                                                                                                                                                                  | Level of Output |
|-----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Context</b>              | IPv4 address that defines a protection pair. The context is manually configured on both primary and protector provider edge (PE) devices.                                                                                                                                                                                                                                                          | All levels      |
| <b>Status</b>               | State of the path: <b>active</b> or <b>inactive</b> .                                                                                                                                                                                                                                                                                                                                              | All levels      |
| <b>Metric</b>               | Advertised OSPF metric.                                                                                                                                                                                                                                                                                                                                                                            | All levels      |
| <b>Area</b>                 | OSPF area number.                                                                                                                                                                                                                                                                                                                                                                                  | All levels      |
| <b>Other Advertisements</b> | Other advertisements received by the OSPF node: <ul style="list-style-type: none"> <li>• <b>Advertising router</b>—Address of the device that sent the advertisement.</li> <li>• <b>Type</b>—Type of OSPF path: <b>inter-area</b> and <b>stub</b>.</li> <li>• <b>Metric</b>—Advertised OSPF metric.</li> <li>• <b>None</b>—No additional advertisements were received by the OSPF node.</li> </ul> | <b>detail</b>   |

## Sample Output

### show ospf context-identifier

```
user@host> show ospf context-identifier
Context-id: 2.2.4.3
Status: active, Metric: 65534, PE role: protector, Area: 0.0.0.0
```

### show ospf context-identifier detail

```
user@host> show ospf context-identifier detail
Context-id: 88.24.13.1
Status: inactive, Metric: 0, PE role: protector, Area: 0.0.0.13
Other Advertisements:
Advertising router: 8.8.8.103
Type: stub link
Metric: 65534
```

## show ospf database

**List of Syntax**    [Syntax on page 4635](#)  
                           [Syntax \(EX Series Switch and QFX Series\) on page 4635](#)

**Syntax**    show ospf database  
                   <brief | detail | extensive | summary>  
                   <advertising-router (*address* | self)>  
                   <area *area-id*>  
                   <asbrsummary>  
                   <external>  
                   <instance *instance-name*>  
                   <link-local>  
                   <logical-system (all | *logical-system-name*)>  
                   <lsa-id *lsa-id*>  
                   <netsummary>  
                   <network>  
                   <nssa>  
                   <opaque-area>  
                   <router>

**Syntax (EX Series Switch and QFX Series)**    show ospf database  
                                                           <brief | detail | extensive | summary>  
                                                           <advertising-router (*address* | self)>  
                                                           <area *area-id*>  
                                                           <asbrsummary>  
                                                           <external>  
                                                           <instance *instance-name*>  
                                                           <link-local>  
                                                           <lsa-id *lsa-id*>  
                                                           <netsummary>  
                                                           <network>  
                                                           <nssa>  
                                                           <opaque-area>  
                                                           <router>

**Release Information**    Command introduced before Junos OS Release 7.4.  
                                   Command introduced in Junos OS Release 9.0 for EX Series switches.  
                                   **advertising-router self (*address* | self)** option introduced in Junos OS Release 9.5.  
                                   **advertising-router self (*address* | self)** option introduced in Junos OS Release 9.5 for EX Series switches.  
                                   Command introduced in Junos OS Release 11.3 for the QFX Series.

**Description**    Display the entries in the Open Shortest Path First version 2 (OSPFv2) link-state database, which contains data about link-state advertisement (LSA) packets.

**Options**    **none**—Display standard information about entries in the OSPFv2 link-state database for all routing instances.

**brief | detail | extensive | summary**—(Optional) Display the specified level of output.

**advertising-router (*address* | self)**—(Optional) Display the LSAs advertised either by a particular routing device or by this routing device.

**area** *area-id*—(Optional) Display the LSAs in a particular area.

**asbrsummary**—(Optional) Display summary AS boundary router LSA entries.

**external**—(Optional) Display external LSAs.

**instance** *instance-name*—(Optional) Display all OSPF database information under the named routing instance.

**link-local**—(Optional) Display information about link-local LSAs.

**logical-system** (**all** | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

**lsa-id** *lsa-id*—(Optional) Display the LSA with the specified LSA identifier.

**netsummary**—(Optional) Display summary network LSAs.

**network**—(Optional) Display information about network LSAs.

**nssa**—(Optional) Display information about not-so-stubby area (NSSA) LSAs.

**opaque-area**—(Optional) Display opaque area-scope LSAs.

**router**—(Optional) Display information about router LSAs.

**Required Privilege Level**

view

**Related Documentation**

- [clear \(ospf | ospf3\) database on page 4618](#)

**List of Sample Output**

[show ospf database on page 4638](#)  
[show ospf database brief on page 4638](#)  
[show ospf database detail on page 4638](#)  
[show ospf database extensive on page 4640](#)  
[show ospf database summary on page 4642](#)

**Output Fields**

[Table 349](#) describes the output fields for the **show ospf database** command. Output fields are listed in the approximate order in which they appear.

**Table 349: show ospf database Output Fields**

| Field Name     | Field Description                                                                                                                                        | Level of Output |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>area</b>    | Area number. Area 0.0.0.0 is the backbone area.                                                                                                          | All levels      |
| <b>Type</b>    | Type of link advertisement: <b>ASBRSum</b> , <b>Extern</b> , <b>Network</b> , <b>NSSA</b> , <b>OpaqArea</b> , <b>Router</b> , or <b>Summary</b> .        | All levels      |
| <b>ID</b>      | LSA identifier included in the advertisement. An asterisk preceding the identifier marks database entries that originated from the local routing device. | All levels      |
| <b>Adv Rtr</b> | Address of the routing device that sent the advertisement.                                                                                               | All levels      |

Table 349: show ospf database Output Fields (*continued*)

| Field Name                              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Level of Output         |
|-----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Seq</b>                              | Link sequence number of the advertisement.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | All levels              |
| <b>Age</b>                              | Time elapsed since the LSA was originated, in seconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | All levels              |
| <b>Opt</b>                              | Optional OSPF capabilities associated with the LSA.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | All levels              |
| <b>Cksum</b>                            | Checksum value of the LSA.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | All levels              |
| <b>Len</b>                              | Length of the advertisement, in bytes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | All levels              |
| <b>Router</b>                           | Router link-state advertisement information: <ul style="list-style-type: none"> <li><b>bits</b>—Flags describing the routing device that generated the LSP.</li> <li><b>link count</b>—Number of links in the advertisement.</li> <li><b>id</b>—ID of a routing device or subnet on the link.</li> <li><b>data</b>—For stub networks, the subnet mask; otherwise, the IP address of the routing device that generated the LSP.</li> <li><b>type</b>—Type of link. It can be <b>PointToPoint</b>, <b>Transit</b>, <b>Stub</b>, or <b>Virtual</b>.</li> <li><b>TOS count</b>—Number of type-of-service (ToS) entries in the advertisement.</li> <li><b>TOS 0 metric</b>—Metric for ToS 0.</li> <li><b>TOS</b>—Type-of-service (ToS) value.</li> <li><b>metric</b>—Metric for the ToS.</li> </ul> | <b>detail extensive</b> |
| <b>Network</b>                          | Network link-state advertisement information: <ul style="list-style-type: none"> <li><b>mask</b>—Network mask.</li> <li><b>attached router</b>—ID of the attached neighbor.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>detail extensive</b> |
| <b>Summary</b>                          | Summary link-state advertisement information: <ul style="list-style-type: none"> <li><b>mask</b>—Network mask.</li> <li><b>TOS</b>—Type-of-service (ToS) value.</li> <li><b>metric</b>—Metric for the ToS.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>detail extensive</b> |
| <b>Gen timer</b>                        | How long until the LSA is regenerated.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>extensive</b>        |
| <b>Aging timer</b>                      | How long until the LSA expires.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>extensive</b>        |
| <b>Installed <i>hh:mm:ss</i> ago</b>    | How long ago the route was installed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>extensive</b>        |
| <b>expires in <i>hh:mm:ss</i></b>       | How long until the route expires.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>extensive</b>        |
| <b>sent <i>hh:mm:ss</i> ago</b>         | How long ago the LSA was sent.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>extensive</b>        |
| <b>Last changed <i>hh:mm:ss</i> ago</b> | How long ago the route was changed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>extensive</b>        |

Table 349: show ospf database Output Fields (*continued*)

| Field Name          | Field Description                                                                  | Level of Output  |
|---------------------|------------------------------------------------------------------------------------|------------------|
| <b>Change count</b> | Number of times the route has changed.                                             | <b>extensive</b> |
| <b>Ours</b>         | Indicates that this is a local advertisement.                                      | <b>extensive</b> |
| <b>Router LSAs</b>  | Number of router link-state advertisements in the link-state database.             | <b>summary</b>   |
| <b>Network LSAs</b> | Number of network link-state advertisements in the link-state database.            | <b>summary</b>   |
| <b>Summary LSAs</b> | Number of summary link-state advertisements in the link-state database.            | <b>summary</b>   |
| <b>NSSA LSAs</b>    | Number of not-so-stubby area link-state advertisements in the link-state database. | <b>summary</b>   |

## Sample Output

### show ospf database

```

user@host> show ospf database
OSPF link state database, Area 0.0.0.1
 Type ID Adv Rtr Seq Age Opt Cksum Len
Router 10.255.70.103 10.255.70.103 0x80000002 215 0x20 0x4112 48
Router *10.255.71.242 10.255.71.242 0x80000002 214 0x20 0x11b1 48
Summary *23.1.1.0 10.255.71.242 0x80000002 172 0x20 0x6d72 28
Summary *24.1.1.0 10.255.71.242 0x80000002 177 0x20 0x607e 28
NSSA *33.1.1.1 10.255.71.242 0x80000002 217 0x28 0x73bd 36

 OSPF link state database, Area 0.0.0.2
 Type ID Adv Rtr Seq Age Opt Cksum Len
Router 10.255.71.52 10.255.71.52 0x80000004 174 0x20 0xd021 36
Router *10.255.71.242 10.255.71.242 0x80000003 173 0x20 0xe191 36
Network *23.1.1.1 10.255.71.242 0x80000002 173 0x20 0x9c76 32
Summary *12.1.1.0 10.255.71.242 0x80000001 217 0x20 0xfeec 28
Summary *24.1.1.0 10.255.71.242 0x80000002 177 0x20 0x607e 28
NSSA *33.1.1.1 10.255.71.242 0x80000001 222 0x28 0xe047 36

 OSPF link state database, Area 0.0.0.3
 Type ID Adv Rtr Seq Age Opt Cksum Len
Router 10.255.71.238 10.255.71.238 0x80000003 179 0x20 0x3942 36
Router *10.255.71.242 10.255.71.242 0x80000003 177 0x20 0xf37d 36
Network *24.1.1.1 10.255.71.242 0x80000002 177 0x20 0xc591 32
Summary *12.1.1.0 10.255.71.242 0x80000001 217 0x20 0xfeec 28
Summary *23.1.1.0 10.255.71.242 0x80000002 172 0x20 0x6d72 28
NSSA *33.1.1.1 10.255.71.242 0x80000001 222 0x28 0xeb3b 36

```

### show ospf database brief

The output for the **show ospf database brief** command is identical to that for the **show ospf database** command. For sample output, see [show ospf database on page 4638](#).

### show ospf database detail

```

user@host> show ospf database detail

```

```

OSPF link state database, Area 0.0.0.1
Type ID Adv Rtr Seq Age Opt Cksum Len
Router 10.255.70.103 10.255.70.103 0x80000002 261 0x20 0x4112 48
 bits 0x0, link count 2
 id 10.255.71.242, data 12.1.1.1, Type PointToPoint (1)
 TOS count 0, TOS 0 metric 1
 id 12.1.1.0, data 255.255.255.0, Type Stub (3)
 TOS count 0, TOS 0 metric 1
Router *10.255.71.242 10.255.71.242 0x80000002 260 0x20 0x11b1 48
 bits 0x3, link count 2
 id 10.255.70.103, data 12.1.1.2, Type PointToPoint (1)
 TOS count 0, TOS 0 metric 1
 id 12.1.1.0, data 255.255.255.0, Type Stub (3)
 TOS count 0, TOS 0 metric 1
Summary *23.1.1.0 10.255.71.242 0x80000002 218 0x20 0x6d72 28
 mask 255.255.255.0
 TOS 0x0, metric 1
Summary *24.1.1.0 10.255.71.242 0x80000002 223 0x20 0x607e 28
 mask 255.255.255.0
 TOS 0x0, metric 1
NSSA *33.1.1.1 10.255.71.242 0x80000002 263 0x28 0x73bd 36
 mask 255.255.255.255
 Type 2, TOS 0x0, metric 0, fwd addr 12.1.1.2, tag 0.0.0.0

```

```

OSPF link state database, Area 0.0.0.2
Type ID Adv Rtr Seq Age Opt Cksum Len
Router 10.255.71.52 10.255.71.52 0x80000004 220 0x20 0xd021 36
 bits 0x0, link count 1
 id 23.1.1.1, data 23.1.1.2, Type Transit (2)
 TOS count 0, TOS 0 metric 1
Router *10.255.71.242 10.255.71.242 0x80000003 219 0x20 0xe191 36
 bits 0x3, link count 1
 id 23.1.1.1, data 23.1.1.1, Type Transit (2)
 TOS count 0, TOS 0 metric 1
Network *23.1.1.1 10.255.71.242 0x80000002 219 0x20 0x9c76 32
 mask 255.255.255.0
 attached router 10.255.71.242
 attached router 10.255.71.52
Summary *12.1.1.0 10.255.71.242 0x80000001 263 0x20 0xfeec 28
 mask 255.255.255.0
 TOS 0x0, metric 1
Summary *24.1.1.0 10.255.71.242 0x80000002 223 0x20 0x607e 28
 mask 255.255.255.0
 TOS 0x0, metric 1
NSSA *33.1.1.1 10.255.71.242 0x80000001 268 0x28 0xe047 36
 mask 255.255.255.255
 Type 2, TOS 0x0, metric 0, fwd addr 23.1.1.1, tag 0.0.0.0

```

```

OSPF link state database, Area 0.0.0.3
Type ID Adv Rtr Seq Age Opt Cksum Len
Router 10.255.71.238 10.255.71.238 0x80000003 225 0x20 0x3942 36
 bits 0x0, link count 1
 id 24.1.1.1, data 24.1.1.2, Type Transit (2)
 TOS count 0, TOS 0 metric 1
Router *10.255.71.242 10.255.71.242 0x80000003 223 0x20 0xf37d 36
 bits 0x3, link count 1
 id 24.1.1.1, data 24.1.1.1, Type Transit (2)
 TOS count 0, TOS 0 metric 1
Network *24.1.1.1 10.255.71.242 0x80000002 223 0x20 0xc591 32
 mask 255.255.255.0
 attached router 10.255.71.242

```

```

 attached router 10.255.71.238
Summary *12.1.1.0 10.255.71.242 0x80000001 263 0x20 0xfeec 28
 mask 255.255.255.0
 TOS 0x0, metric 1
Summary *23.1.1.0 10.255.71.242 0x80000002 218 0x20 0x6d72 28
 mask 255.255.255.0
 TOS 0x0, metric 1
NSSA *33.1.1.1 10.255.71.242 0x80000001 268 0x28 0xeb3b 36
 mask 255.255.255.255
 Type 2, TOS 0x0, metric 0, fwd addr 24.1.1.1, tag 0.0.0.0

```

### show ospf database extensive

```

user@host> show ospf database extensive
 OSPF link state database, Area 0.0.0.1
Type ID Adv Rtr Seq Age Opt Cksum Len
Router 10.255.70.103 10.255.70.103 0x80000002 286 0x20 0x4112 48
 bits 0x0, link count 2
 id 10.255.71.242, data 12.1.1.1, Type PointToPoint (1)
 TOS count 0, TOS 0 metric 1
 id 12.1.1.0, data 255.255.255.0, Type Stub (3)
 TOS count 0, TOS 0 metric 1
 Aging timer 00:55:14
 Installed 00:04:43 ago, expires in 00:55:14
 Last changed 00:04:43 ago, Change count: 2
Router *10.255.71.242 10.255.71.242 0x80000002 285 0x20 0x11b1 48
 bits 0x3, link count 2
 id 10.255.70.103, data 12.1.1.2, Type PointToPoint (1)
 TOS count 0, TOS 0 metric 1
 id 12.1.1.0, data 255.255.255.0, Type Stub (3)
 TOS count 0, TOS 0 metric 1
 Gen timer 00:45:15
 Aging timer 00:55:15
 Installed 00:04:45 ago, expires in 00:55:15, sent 00:04:43 ago
 Last changed 00:04:45 ago, Change count: 2, Ours
Summary *23.1.1.0 10.255.71.242 0x80000002 243 0x20 0x6d72 28
 mask 255.255.255.0
 TOS 0x0, metric 1
 Gen timer 00:45:57
 Aging timer 00:55:57
 Installed 00:04:03 ago, expires in 00:55:57, sent 00:04:01 ago
 Last changed 00:04:48 ago, Change count: 1, Ours
Summary *24.1.1.0 10.255.71.242 0x80000002 248 0x20 0x607e 28
 mask 255.255.255.0
 TOS 0x0, metric 1
 Gen timer 00:45:52
 Aging timer 00:55:52
 Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:06 ago
 Last changed 00:04:48 ago, Change count: 1, Ours
NSSA *33.1.1.1 10.255.71.242 0x80000002 288 0x28 0x73bd 36
 mask 255.255.255.255
 Type 2, TOS 0x0, metric 0, fwd addr 12.1.1.2, tag 0.0.0.0
 Gen timer 00:45:12
 Aging timer 00:55:12
 Installed 00:04:48 ago, expires in 00:55:12, sent 00:04:48 ago
 Last changed 00:04:48 ago, Change count: 2, Ours

 OSPF link state database, Area 0.0.0.2
Type ID Adv Rtr Seq Age Opt Cksum Len
Router 10.255.71.52 10.255.71.52 0x80000004 245 0x20 0xd021 36
 bits 0x0, link count 1

```



```

id 23.1.1.1, data 23.1.1.2, Type Transit (2)
TOS count 0, TOS 0 metric 1
Aging timer 00:55:55
Installed 00:04:02 ago, expires in 00:55:55
Last changed 00:04:02 ago, Change count: 2
Router *10.255.71.242 10.255.71.242 0x80000003 244 0x20 0xe191 36
bits 0x3, link count 1
id 23.1.1.1, data 23.1.1.1, Type Transit (2)
TOS count 0, TOS 0 metric 1
Gen timer 00:45:56
Aging timer 00:55:56
Installed 00:04:04 ago, expires in 00:55:56, sent 00:04:02 ago
Last changed 00:04:04 ago, Change count: 2, Ours
Network *23.1.1.1 10.255.71.242 0x80000002 244 0x20 0x9c76 32
mask 255.255.255.0
attached router 10.255.71.242
attached router 10.255.71.52
Gen timer 00:45:56
Aging timer 00:55:56
Installed 00:04:04 ago, expires in 00:55:56, sent 00:04:02 ago
Last changed 00:04:04 ago, Change count: 1, Ours
Summary *12.1.1.0 10.255.71.242 0x80000001 288 0x20 0xfec 28
mask 255.255.255.0
TOS 0x0, metric 1
Gen timer 00:45:12
Aging timer 00:55:12
Installed 00:04:48 ago, expires in 00:55:12, sent 00:04:04 ago
Last changed 00:04:48 ago, Change count: 1, Ours
Summary *24.1.1.0 10.255.71.242 0x80000002 248 0x20 0x607e 28
mask 255.255.255.0
TOS 0x0, metric 1
Gen timer 00:45:52
Aging timer 00:55:52
Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:04 ago
Last changed 00:04:48 ago, Change count: 1, Ours
NSSA *33.1.1.1 10.255.71.242 0x80000001 293 0x28 0xe047 36
mask 255.255.255.255
Type 2, TOS 0x0, metric 0, fwd addr 23.1.1.1, tag 0.0.0.0
Gen timer 00:45:07
Aging timer 00:55:07
Installed 00:04:53 ago, expires in 00:55:07, sent 00:04:04 ago
Last changed 00:04:53 ago, Change count: 1, Ours

OSPF link state database, Area 0.0.0.3
Type ID Adv Rtr Seq Age Opt Cksum Len
Router 10.255.71.238 10.255.71.238 0x80000003 250 0x20 0x3942 36
bits 0x0, link count 1
id 24.1.1.1, data 24.1.1.2, Type Transit (2)
TOS count 0, TOS 0 metric 1
Aging timer 00:55:50
Installed 00:04:07 ago, expires in 00:55:50
Last changed 00:04:07 ago, Change count: 2
Router *10.255.71.242 10.255.71.242 0x80000003 248 0x20 0xf37d 36
bits 0x3, link count 1
id 24.1.1.1, data 24.1.1.1, Type Transit (2)
TOS count 0, TOS 0 metric 1
Gen timer 00:45:52
Aging timer 00:55:52
Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:06 ago
Last changed 00:04:08 ago, Change count: 2, Ours
Network *24.1.1.1 10.255.71.242 0x80000002 248 0x20 0xc591 32

```

```

mask 255.255.255.0
attached router 10.255.71.242
attached router 10.255.71.238
Gen timer 00:45:52
Aging timer 00:55:52
Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:06 ago
Last changed 00:04:08 ago, Change count: 1, Ours
Summary *12.1.1.0 10.255.71.242 0x80000001 288 0x20 0xfeec 28
mask 255.255.255.0
TOS 0x0, metric 1
Gen timer 00:45:12
Aging timer 00:55:12
Installed 00:04:48 ago, expires in 00:55:12, sent 00:04:13 ago
Last changed 00:04:48 ago, Change count: 1, Ours
Summary *23.1.1.0 10.255.71.242 0x80000002 243 0x20 0x6d72 28
mask 255.255.255.0
TOS 0x0, metric 1
Gen timer 00:45:57
Aging timer 00:55:57
Installed 00:04:03 ago, expires in 00:55:57, sent 00:04:01 ago
Last changed 00:04:48 ago, Change count: 1, Ours
NSSA *33.1.1.1 10.255.71.242 0x80000001 293 0x28 0xeb3b 36
mask 255.255.255.255
Type 2, TOS 0x0, metric 0, fwd addr 24.1.1.1, tag 0.0.0.0
Gen timer 00:45:07
Aging timer 00:55:07
Installed 00:04:53 ago, expires in 00:55:07, sent 00:04:13 ago
Last changed 00:04:53 ago, Change count: 1, Ours

```

#### show ospf database summary

```

user@host> show ospf database summary
Area 0.0.0.1:
 2 Router LSAs
 2 Summary LSAs
 1 NSSA LSAs
Area 0.0.0.2:
 2 Router LSAs
 1 Network LSAs
 2 Summary LSAs
 1 NSSA LSAs
Area 0.0.0.3:
 2 Router LSAs
 1 Network LSAs
 2 Summary LSAs
 1 NSSA LSAs
Externals:
Interface fe-2/2/1.0:
Interface ge-0/3/2.0:
Interface so-0/1/2.0:
Interface so-0/1/2.0:

```

## show (ospf | ospf3) interface

|                                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                           | <a href="#">Syntax on page 4643</a><br><a href="#">Syntax (EX Series Switch and QFX Series) on page 4643</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Syntax</b>                                   | <pre>show (ospf   ospf3) interface &lt;brief   detail   extensive&gt; &lt;area <i>area-id</i>&gt; &lt;<i>interface-name</i>&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Syntax (EX Series Switch and QFX Series)</b> | <pre>show (ospf   ospf3) interface &lt;brief   detail   extensive&gt; &lt;area <i>area-id</i>&gt; &lt;<i>interface-name</i>&gt; &lt;instance <i>instance-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>                      | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>area</b> option introduced in Junos OS Release 9.2.</p> <p><b>area</b> option introduced in Junos OS Release 9.2 for EX Series switches.</p> <p><b>realm</b> option introduced in Junos OS Release 9.2.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>                              | Display the status of Open Shortest Path First (OSPF) interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                                  | <p><b>none</b>—Display standard information about the status of all OSPF interfaces for all routing instances</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>area <i>area-id</i></b>—(Optional) Display information about the interfaces that belong to the specified area.</p> <p><b><i>interface-name</i></b>—(Optional) Display information for the specified interface.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display all OSPF interfaces under the named routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)</b>—(Optional) (OSPFv3 only) Display information about the interfaces for the specified OSPFv3 realm, or address family. Use the <b>realm</b> option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.</p> |
| <b>Required Privilege Level</b>                 | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

**List of Sample Output** [show ospf interface brief on page 4646](#)  
[show ospf interface detail on page 4646](#)  
[show ospf3 interface detail on page 4646](#)  
[show ospf interface detail\(When Multiarea Adjacency Is Configured\) on page 4646](#)  
[show ospf interface area area-id on page 4648](#)  
[show ospf interface extensive \(When Flooding Reduction Is Enabled\) on page 4648](#)  
[show ospf interface extensive \(When LDP Synchronization Is Configured\) on page 4648](#)

**Output Fields** Table 350 lists the output fields for the **show (ospf | ospf3) interface** command. Output fields are listed in the approximate order in which they appear.

**Table 350: show (ospf | ospf3) interface Output Fields**

| Field Name              | Field Description                                                                                                                                                                                    | Level of Output         |
|-------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Interface</b>        | Name of the interface running OSPF version 2 or OSPF version 3.                                                                                                                                      | All levels              |
| <b>State</b>            | State of the interface: <b>BDR</b> , <b>Down</b> , <b>DR</b> , <b>DRother</b> , <b>Loop</b> , <b>PtToPt</b> , or <b>Waiting</b> .                                                                    | All levels              |
| <b>Area</b>             | Number of the area that the interface is in.                                                                                                                                                         | All levels              |
| <b>DR ID</b>            | Address of the area's designated router.                                                                                                                                                             | All levels              |
| <b>BDR ID</b>           | Backup designated router for a particular subnet.                                                                                                                                                    | All levels              |
| <b>Nbrs</b>             | Number of neighbors on this interface.                                                                                                                                                               | All levels              |
| <b>Type</b>             | Type of interface: <b>LAN</b> , <b>NBMA</b> , <b>P2MP</b> , <b>P2P</b> , or <b>Virtual</b> .                                                                                                         | <b>detail extensive</b> |
| <b>Address</b>          | IP address of the neighbor.                                                                                                                                                                          | <b>detail extensive</b> |
| <b>Mask</b>             | Netmask of the neighbor.                                                                                                                                                                             | <b>detail extensive</b> |
| <b>Prefix-length</b>    | (OSPFv3) IPv6 prefix length, in bits.                                                                                                                                                                | <b>detail extensive</b> |
| <b>OSPF3-Intf-Index</b> | (OSPFv3) OSPF version 3 interface index.                                                                                                                                                             | <b>detail extensive</b> |
| <b>MTU</b>              | Interface maximum transmission unit (MTU).                                                                                                                                                           | <b>detail extensive</b> |
| <b>Cost</b>             | Interface cost (metric).                                                                                                                                                                             | <b>detail extensive</b> |
| <b>DR addr</b>          | Address of the designated router.                                                                                                                                                                    | <b>detail extensive</b> |
| <b>BDR addr</b>         | Address of the backup designated router.                                                                                                                                                             | <b>detail extensive</b> |
| <b>Adj count</b>        | Number of adjacent neighbors.                                                                                                                                                                        | <b>detail extensive</b> |
| <b>Secondary</b>        | Indicates that this interface is configured as a secondary interface for this area. This interface can belong to more than one area, but can be designated as a primary interface for only one area. | <b>detail extensive</b> |

Table 350: show (ospf | ospf3) interface Output Fields (*continued*)

| Field Name             | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Level of Output         |
|------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Flood Reduction</b> | Indicates that this interface is configured with flooding reduction. All self-originated LSAs from this interface are initially sent with the <b>DoNotAge</b> bit set. As a result, LSAs are refreshed only when a change occurs.                                                                                                                                                                                                                                   | <b>extensive</b>        |
| <b>Priority</b>        | Router priority used in designated router (DR) election on this interface.                                                                                                                                                                                                                                                                                                                                                                                          | <b>detail extensive</b> |
| <b>Flood list</b>      | List of link-state advertisements (LSAs) that might be about to flood this interface.                                                                                                                                                                                                                                                                                                                                                                               | <b>extensive</b>        |
| <b>Ack list</b>        | Acknowledgment list. List of pending acknowledgments on this interface.                                                                                                                                                                                                                                                                                                                                                                                             | <b>extensive</b>        |
| <b>Descriptor list</b> | List of packet descriptors.                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>extensive</b>        |
| <b>Hello</b>           | Configured value for the hello timer.                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive</b> |
| <b>Dead</b>            | Configured value for the dead timer.                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive</b> |
| <b>Auth type</b>       | (OSPFv2) Authentication mechanism for sending and receiving OSPF protocol packets: <ul style="list-style-type: none"> <li>• <b>MD5</b>—The MD5 mechanism is configured in accordance with RFC 2328.</li> <li>• <b>None</b>—No authentication method is configured.</li> <li>• <b>Password</b>—A simple password (RFC 2328) is configured.</li> </ul>                                                                                                                | <b>detail extensive</b> |
| <b>Topology</b>        | (Multiarea adjacency) Name of topology: <b>default</b> or <b>name</b> .                                                                                                                                                                                                                                                                                                                                                                                             |                         |
| <b>LDP sync state</b>  | (OSPFv2 and LDP synchronization) Current state of LDP synchronization: <b>in sync</b> , <b>in holddown</b> , and <b>not supported</b> .                                                                                                                                                                                                                                                                                                                             | <b>extensive</b>        |
| <b>reason</b>          | (OSPFv2 and LDP synchronization) Reason for the current state of LDP synchronization. The LDP session might be up or down, or adjacency might be up or down.                                                                                                                                                                                                                                                                                                        | <b>extensive</b>        |
| <b>config holdtime</b> | (OSPFv2 and LDP synchronization) Configured value of the hold timer.<br><br>If the state is not synchronized, and the hold time is not infinity, the <b>remaining</b> field displays the number of seconds that remain until the configured hold timer expires.                                                                                                                                                                                                     | <b>extensive</b>        |
| <b>IPSec SA name</b>   | (OSPFv2) Name of the IPSec security association name.                                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive</b> |
| <b>Active key ID</b>   | (OSPFv2 and MD5) Number from <b>0</b> to <b>255</b> that uniquely identifies an MD5 key.                                                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b> |
| <b>Start time</b>      | (OSPFv2 and MD5) Time at which the routing device starts using an MD5 key to authenticate OSPF packets transmitted on the interface on which this key is configured. To authenticate received OSPF protocol packets, the key becomes effective immediately after the configuration is committed. If the start time option is not configured, the key is effective immediately for send and receive and is displayed as <b>Start time 1970 Jan 01 00:00:00 PST</b> . | <b>detail extensive</b> |

Table 350: show (ospf | ospf3) interface Output Fields (*continued*)

| Field Name                   | Field Description                          | Level of Output  |
|------------------------------|--------------------------------------------|------------------|
| ReXmit                       | Configured value for the Retransmit timer. | detail extensive |
| Stub, Not Stub, or Stub NSSA | Type of area.                              | detail extensive |

## Sample Output

### show ospf interface brief

```

user@host> show ospf interface brief
Intf State Area DR ID BDR ID Nbrs
at-5/1/0.0 PtToPt 0.0.0.0 0.0.0.0 0.0.0.0 1
ge-2/3/0.0 DR 0.0.0.0 192.168.4.16 192.168.4.15 1
lo0.0 DR 0.0.0.0 192.168.4.16 0.0.0.0 0
so-0/0/0.0 Down 0.0.0.0 0.0.0.0 0.0.0.0 0
so-6/0/1.0 PtToPt 0.0.0.0 0.0.0.0 0.0.0.0 1
so-6/0/2.0 Down 0.0.0.0 0.0.0.0 0.0.0.0 0
so-6/0/3.0 PtToPt 0.0.0.0 0.0.0.0 0.0.0.0 1

```

### show ospf interface detail

```

user@host> show ospf interface detail
Interface State Area DR ID BDR ID Nbrs
fe-0/0/1.0 BDR 0.0.0.0 192.168.37.12 10.255.245.215 1
Type LAN, address 192.168.37.11, Mask 255.255.255.248, MTU 4460, Cost 40
DR addr 192.168.37.12, BDR addr 192.168.37.11, Adj count 1, Priority 128
Hello 10, Dead 40, ReXmit 5, Not Stub
t1-0/2/1.0 PtToPt 0.0.0.0 0.0.0.0 0.0.0.0 0
Type P2P, Address 0.0.0.0, Mask 0.0.0.0, MTU 1500, Cost 2604
Adj count 0
Hello 10, Dead 40, ReXmit 5, Not Stub
Auth type: MD5, Active key ID 3, Start time 2002 Nov 19 10:00:00 PST
IPsec SA Name: sa

```

### show ospf3 interface detail

```

user@host> show ospf3 interface so-0/0/3.0 detail
Interface State Area DR-ID BDR-ID Nbrs
so-0/0/3.0 PtToPt 0.0.0.0 0.0.0.0 0.0.0.0 1
Address fe80::2a0:a5ff:fe28:1dfc, Prefix-length 64
OSPF3-Intf-index 1, Type P2P, MTU 4470, Cost 12, Adj-count 1
Hello 10, Dead 40, ReXmit 5, Not Stub

```

### show ospf interface detail (When Multiarea Adjacency Is Configured)

```

user@host> show ospf interface detail
regress@router> show ospf interface detail
Interface State Area DR ID BDR ID Nbrs
lo0.0 DR 0.0.0.0 10.255.245.2 0.0.0.0 0

Type: LAN, Address: 127.0.0.1, Mask: 255.255.255.255, MTU: 65535, Cost: 0
DR addr: 127.0.0.1, Adj count: 0, Priority: 128
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None

```

```

Topology default (ID 0) -> Cost: 0
100.0 DR 0.0.0.0 10.255.245.2 0.0.0.0 0

Type: LAN, Address: 10.255.245.2, Mask: 255.255.255.255, MTU: 65535, Cost: 0
DR addr: 10.255.245.2, Adj count: 0, Priority: 128
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 0
so-0/0/0.0 PtToPt 0.0.0.0 0.0.0.0 0.0.0.0 1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-0/0/0.0 PtToPt 0.0.0.0 0.0.0.0 0.0.0.0 0

Type: P2P, Address: 192.168.37.46, Mask: 255.255.255.254, MTU: 4470, Cost: 1
Adj count: 0, , Passive
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Passive, Cost: 1
so-1/0/0.0 PtToPt 0.0.0.0 0.0.0.0 0.0.0.0 1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-1/0/0.0 PtToPt 0.0.0.0 0.0.0.0 0.0.0.0 0

Type: P2P, Address: 192.168.37.54, Mask: 255.255.255.254, MTU: 4470, Cost: 1
Adj count: 0, , Passive
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Passive, Cost: 1
so-0/0/0.0 PtToPt 1.1.1.1 0.0.0.0 0.0.0.0 1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1, Secondary
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-1/0/0.0 PtToPt 1.1.1.1 0.0.0.0 0.0.0.0 1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1, Secondary
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-0/0/0.0 PtToPt 2.2.2.2 0.0.0.0 0.0.0.0 1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1, Secondary
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-1/0/0.0 PtToPt 2.2.2.2 0.0.0.0 0.0.0.0 1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1, Secondary

```

```
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
```

#### show ospf interface area area-id

```
user@host> show ospf interface area 1.1.1.1
Interface State Area DR ID BDR ID Nbrs
so-0/0/0.0 PtToPt 1.1.1.1 0.0.0.0 0.0.0.0 1
so-1/0/0.0 PtToPt 1.1.1.1 0.0.0.0 0.0.0.0 1
```

#### show ospf interface extensive (When Flooding Reduction Is Enabled)

```
user@host> show ospf interface extensive
Interface State Area DR ID BDR ID Nbrs
fe-0/0/0.0 PtToPt 0.0.0.0 0.0.0.0 0.0.0.0 0

Type: P2P, Address: 10.10.10.1, Mask: 255.255.255.0, MTU: 1500, Cost: 1
Adj count: 0
Secondary, Flood Reduction
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
```

#### show ospf interface extensive (When LDP Synchronization Is Configured)

```
user@host> show ospf interface extensive
Interface State Area DR ID BDR ID
Nbrs
so-1/0/3.0 Down 0.0.0.0 0.0.0.0 0.0.0.0
0
Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 65535
Adj count: 0
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
LDP sync state: in holddown, for: 00:00:08, reason: LDP down during config
config holddown: 10 seconds, remaining: 1
```



## show (ospf | ospf3) io-statistics

|                                                 |                                                                                                                                                                                                                     |
|-------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                           | <a href="#">Syntax on page 4649</a><br><a href="#">Syntax (EX Series Switch and QFX Series) on page 4649</a>                                                                                                        |
| <b>Syntax</b>                                   | show (ospf   ospf3) io-statistics<br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                           |
| <b>Syntax (EX Series Switch and QFX Series)</b> | show (ospf   ospf3) io-statistics                                                                                                                                                                                   |
| <b>Release Information</b>                      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                            |
| <b>Description</b>                              | Display Open Shortest Path First (OSPF) input and output statistics.                                                                                                                                                |
| <b>Options</b>                                  | <b>none</b> —Display OSPF input and output statistics.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>                 | view                                                                                                                                                                                                                |
| <b>Related Documentation</b>                    | <ul style="list-style-type: none"> <li>• <a href="#">clear (ospf   ospf3) statistics on page 4625</a></li> </ul>                                                                                                    |
| <b>List of Sample Output</b>                    | <a href="#">show ospf io-statistics on page 4650</a>                                                                                                                                                                |
| <b>Output Fields</b>                            | Table 351 lists the output fields for the <b>show ospf io-statistics</b> command. Output fields are listed in the approximate order in which they appear.                                                           |

**Table 351: show (ospf | ospf3) io-statistics Output Fields**

| Field Name      | Field Description                                                                                         |
|-----------------|-----------------------------------------------------------------------------------------------------------|
| Packets read    | Number of OSPF packets read since the last time the routing protocol was started.                         |
| average per run | Total number of packets divided by the total number of times the OSPF read operation is scheduled to run. |
| max run         | Maximum number of packets for a given run among all scheduled runs.                                       |
| Receive errors  | Number of faulty packets received with errors.                                                            |

## Sample Output

### show ospf io-statistics

```
user@host> show ospf io-statistics
```

```
Packets read: 7361, average per run: 1.00, max run: 1
Receive errors:
None
```

## show (ospf | ospf3) log

|                                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                           | <a href="#">Syntax on page 4651</a><br><a href="#">Syntax (EX Series Switch and QFX Series) on page 4651</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Syntax</b>                                   | <pre>show (ospf   ospf3) log &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)&gt; &lt;topology <i>topology-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Syntax (EX Series Switch and QFX Series)</b> | <pre>show (ospf   ospf3) log &lt;instance <i>instance-name</i>&gt; &lt;topology <i>topology-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>                      | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>topology</b> option introduced in Junos OS Release 9.0.</p> <p><b>topology</b> option introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>realm</b> option introduced in Junos OS Release 9.2.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>                              | Display the entries in the Open Shortest Path First (OSPF) log of SPF calculations.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                                  | <p><b>none</b>—Display entries in the OSPF log of SPF calculations for all routing instances.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display entries for the specified routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>topology <i>topology-name</i></b>—(Optional) (OSPFv2 only) Display entries for the specified topology.</p> <p><b>realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)</b>—(OSPFv3 only) (Optional) Display entries for the specified OSPFv3 realm, or address family. Use the <b>realm</b> option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.</p> |
| <b>Required Privilege Level</b>                 | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>List of Sample Output</b>                    | <a href="#">show ospf log on page 4652</a><br><a href="#">show ospf log topology voice on page 4652</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Output Fields</b>                            | Table 352 lists the output fields for the <b>show (ospf   ospf3) log</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

**Table 352: show (ospf | ospf3) log Output Fields**

| Field Name  | Field Description                                                                      |
|-------------|----------------------------------------------------------------------------------------|
| <b>When</b> | Time, in weeks ( <b>w</b> ) and days ( <b>d</b> ), since the SPF calculation was made. |

Table 352: show (ospf | ospf3) log Output Fields (*continued*)

| Field Name | Field Description                                                                                                                                                            |
|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Type       | Type of calculation: Cleanup, External, Interarea, NSSA, Redist, SPF, Stub, Total, or Virtuallink.                                                                           |
| Elapsed    | Amount of time, in seconds, that elapsed during the operation, or the time required to complete the SPF calculation. The start time is the time displayed in the When field. |

## Sample Output

### show ospf log

```

user@host> show ospf log
When Type Elapsed
1w4d 17:25:58 Stub 0.000017
1w4d 17:25:58 SPF 0.000070
1w4d 17:25:58 Stub 0.000019
1w4d 17:25:58 Interarea 0.000054
1w4d 17:25:58 External 0.000005
1w4d 17:25:58 Cleanup 0.000203
1w4d 17:25:58 Total 0.000537
1w4d 17:24:48 SPF 0.000125
1w4d 17:24:48 Stub 0.000017
1w4d 17:24:48 SPF 0.000100
1w4d 17:24:48 Stub 0.000016
1w4d 17:24:48 Interarea 0.000056
1w4d 17:24:48 External 0.000005
1w4d 17:24:48 Cleanup 0.000238
1w4d 17:24:48 Total 0.000600
...

```

### show ospf log topology voice

```

user@host> show ospf log topology voice
Topology voice SPF log:

 Last instance of each event type
When Type Elapsed
00:06:11 SPF 0.000116
00:06:11 Stub 0.000114
00:06:11 Interarea 0.000126
00:06:11 External 0.000067
00:06:11 NSSA 0.000037
00:06:11 Cleanup 0.000186

 Maximum length of each event type
When Type Elapsed
00:13:43 SPF 0.000140
00:13:33 Stub 0.000116
00:13:43 Interarea 0.000128
00:13:33 External 0.000075
00:13:38 NSSA 0.000039
00:13:53 Cleanup 0.000657

```

Last 100 events

| When     | Type      | Elapsed  |
|----------|-----------|----------|
| 00:13:53 | SPF       | 0.000090 |
| 00:13:53 | Stub      | 0.000041 |
| 00:13:53 | Interarea | 0.000123 |
| 00:13:53 | External  | 0.000040 |
| 00:13:53 | NSSA      | 0.000038 |
| 00:13:53 | Cleanup   | 0.000657 |
| 00:13:53 | Total     | 0.001252 |
| .        |           |          |
| .        |           |          |
| 00:06:11 | SPF       | 0.000116 |
| 00:06:11 | Stub      | 0.000114 |
| 00:06:11 | Interarea | 0.000126 |
| 00:06:11 | External  | 0.000067 |
| 00:06:11 | NSSA      | 0.000037 |
| 00:06:11 | Cleanup   | 0.000186 |
| 00:06:11 | Total     | 0.000818 |

## show (ospf | ospf3) neighbor

---

|                                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                           | <a href="#">Syntax on page 4654</a><br><a href="#">Syntax (EX Series Switch and QFX Series) on page 4654</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Syntax</b>                                   | <pre>show (ospf   ospf3) neighbor &lt;brief   detail   extensive&gt; &lt;area <i>area-id</i>&gt; &lt;instance (all   <i>instance-name</i>)&gt; &lt;interface <i>interface-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;neighbor&gt; &lt;realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Syntax (EX Series Switch and QFX Series)</b> | <pre>show (ospf   ospf3) neighbor &lt;brief   detail   extensive&gt; &lt;area <i>area-id</i>&gt; &lt;instance (all   <i>instance-name</i>)&gt; &lt;interface <i>interface-name</i>&gt; &lt;neighbor&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>                      | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>instance all</b> option introduced in Junos OS Release 9.1.</p> <p><b>instance all</b> option introduced in Junos OS Release 9.1 for EX Series switches.</p> <p><b>area</b>, <b>interface</b>, and <b>realm</b> options introduced in Junos OS Release 9.2.</p> <p><b>area</b> and <b>interface</b> options introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>                              | Display information about Open Shortest Path First (OSPF) neighbors.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                                  | <p><b>none</b>—Display standard information about all OSPF neighbors for all routing instances.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>area <i>area-id</i></b>—(Optional) Display information about the OSPF neighbors for the specified area.</p> <p><b>instance (all   <i>instance-name</i>)</b>—(Optional) Display all OSPF interfaces for all routing instances or under the named routing instance.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display information about OSPF neighbors for the specified logical interface.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>neighbor</b>—(Optional) Display information about the specified OSPF neighbor.</p> <p><b>realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)</b>—(Optional) (OSPFv3 only) Display information about the OSPF neighbors for the specified OSPFv3 realm, or address</p> |

family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Required Privilege Level | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Related Documentation    | <ul style="list-style-type: none"><li>• <a href="#">clear (ospf   ospf3) neighbor on page 4623</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                          |
| List of Sample Output    | <a href="#">show ospf neighbor brief on page 4657</a><br><a href="#">show ospf neighbor detail on page 4657</a><br><a href="#">show ospf neighbor extensive on page 4658</a><br><a href="#">show ospf3 neighbor detail on page 4659</a><br><a href="#">show ospf neighbor area area-id on page 4659</a><br><a href="#">show ospf neighbor interface interface-name on page 4659</a><br><a href="#">show ospf3 neighbor instance all (OSPFv3 Multiple Family Address Support Enabled) on page 4659</a> |
| Output Fields            | <a href="#">Table 353</a> lists the output fields for the <b>show (ospf   ospf3) neighbor</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                        |

Table 353: show (ospf | ospf3) neighbor Output Fields

| Field Name | Field Description                                  | Level of Output |
|------------|----------------------------------------------------|-----------------|
| Address    | Address of the neighbor.                           | All levels      |
| Interface  | Interface through which the neighbor is reachable. | All levels      |

Table 353: show (ospf | ospf3) neighbor Output Fields (*continued*)

| Field Name                            | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Level of Output         |
|---------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>State</b>                          | <p>State of the neighbor:</p> <ul style="list-style-type: none"> <li>• <b>Attempt</b>—Valid only for neighbors attached to nonbroadcast networks. It indicates that no recent information has been received from the neighbor, but that a more concerted effort must be made to contact the neighbor.</li> <li>• <b>Down</b>—Initial state of a neighbor conversation. It indicates that no recent information has been received from the neighbor. Hello packets might continue to be sent to neighbors in the <b>Down</b> state, although at a reduced frequency.</li> <li>• <b>Exchange</b>—Routing device is describing its entire link-state database by sending database description packets to the neighbor. Each packet has a sequence number and is explicitly acknowledged.</li> <li>• <b>ExStart</b>—First step in creating an adjacency between the two neighboring routing devices. The goal of this step is to determine which routing device is the master, and to determine the initial sequence number.</li> <li>• <b>Full</b>—Neighboring routing devices are fully adjacent. These adjacencies appear in router link and network link advertisements.</li> <li>• <b>Init</b>—A hello packet has recently been sent by the neighbor. However, bidirectional communication has not yet been established with the neighbor. This state may occur, for example, because the routing device itself did not appear in the neighbor's hello packet.</li> <li>• <b>Loading</b>—Link-state request packets are sent to the neighbor to acquire more recent advertisements that have been discovered (but not yet received) in the <b>Exchange</b> state.</li> <li>• <b>2Way</b>—Communication between the two routing devices is bidirectional. This state has been ensured by the operation of the Hello Protocol. This is the most advanced state short of beginning adjacency establishment. The (backup) designated router is selected from the set of neighbors in state <b>2Way</b> or greater.</li> </ul> | All levels              |
| <b>ID</b>                             | Router ID of the neighbor.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | All levels              |
| <b>Pri</b>                            | Priority of the neighbor to become the designated router.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | All levels              |
| <b>Dead</b>                           | Number of seconds until the neighbor becomes unreachable.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | All levels              |
| <b>Link state acknowledgment list</b> | Number of link-state acknowledgments received.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>extensive</b>        |
| <b>Link state retransmission list</b> | <p>Total number of link-state advertisements retransmitted. For <b>extensive</b> output only, the following information is also displayed:</p> <ul style="list-style-type: none"> <li>• <b>Type</b>—Type of link advertisement: <b>ASBR</b>, <b>Sum</b>, <b>Extern</b>, <b>Network</b>, <b>NSSA</b>, <b>OpaqueArea</b>, <b>Router</b>, or <b>Summary</b>.</li> <li>• <b>LSA ID</b>—LSA identifier included in the advertisement. An asterisk preceding the identifier marks database entries that originated from the local routing device.</li> <li>• <b>Adv rtr</b>—Address of the routing device that sent the advertisement.</li> <li>• <b>Seq</b>—Link sequence number of the advertisement.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive</b> |



Table 353: show (ospf | ospf3) neighbor Output Fields (*continued*)

| Field Name              | Field Description                                                                                                                                                                                                                | Level of Output         |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Neighbor-address</b> | (OSPFv3 only) If the neighbor uses virtual links, the <b>Neighbor-address</b> is the site-local, local, or global address. If the neighbor uses a physical interface, the <b>Neighbor-address</b> is an IPv6 link-local address. | <b>detail extensive</b> |
| <b>area</b>             | Area that the neighbor is in.                                                                                                                                                                                                    | <b>detail extensive</b> |
| <b>OSPF3-Intf-Index</b> | (OSPFv3 only) Displays the OSPFv3 interface index.                                                                                                                                                                               | <b>detail extensive</b> |
| <b>opt</b>              | Option bits received in the hello packets from the neighbor.                                                                                                                                                                     | <b>detail extensive</b> |
| <b>DR or DR-ID</b>      | Address of the designated router.                                                                                                                                                                                                | <b>detail extensive</b> |
| <b>BDR or BDR-ID</b>    | Address of the backup designated router.                                                                                                                                                                                         | <b>detail extensive</b> |
| <b>Up</b>               | Length of time since the neighbor came up.                                                                                                                                                                                       | <b>detail extensive</b> |
| <b>adjacent</b>         | Length of time since the adjacency with the neighbor was established.                                                                                                                                                            | <b>detail extensive</b> |

## Sample Output

### show ospf neighbor brief

```

user@host> show ospf neighbor brief
 Address Intf State ID Pri Dead
192.168.254.225 fxp3.0 2Way 10.250.240.32 128 36
192.168.254.230 fxp3.0 Full 10.250.240.8 128 38
192.168.254.229 fxp3.0 Full 10.250.240.35 128 33
10.1.1.129 fxp2.0 Full 10.250.240.12 128 37
10.1.1.131 fxp2.0 Full 10.250.240.11 128 38
10.1.2.1 fxp1.0 Full 10.250.240.9 128 32
10.1.2.81 fxp0.0 Full 10.250.240.10 128 33

```

### show ospf neighbor detail

```

user@host> show ospf neighbor detail
 Address Interface State ID Pri Dead
10.5.1.2 ge-1/2/0.1 Full 10.5.1.2 128 37
area 0.0.0.1, opt 0x42, DR 10.5.1.2, BDR 10.5.1.1
Up 06:09:28, adjacent 05:17:36
Link state acknowledgment list: 3 entries

Link state retransmission list: 9 entries

10.5.10.2 ge-1/2/0.10 ExStart 10.5.1.38 128 34
area 0.0.0.1, opt 0x42, DR 10.5.10.2, BDR 10.5.10.1
Up 06:09:28
master, seq 0xac1530f8, rexmit DBD in 3 sec
rexmit LSREQ in 0 sec
10.5.11.2 ge-1/2/0.11 Full 10.5.1.42 128 38
area 0.0.0.1, opt 0x42, DR 10.5.11.2, BDR 10.5.11.1
Up 06:09:28, adjacent 05:26:46
Link state retransmission list: 1 entries

```

```

10.5.12.2 ge-1/2/0.12 ExStart 10.5.1.46 128 33
area 0.0.0.1, opt 0x42, DR 10.5.12.2, BDR 10.5.12.1
Up 06:09:28
master, seq 0xac188a68, rexmit DBD in 2 sec
rexmit LSREQ in 0 sec

```

### show ospf neighbor extensive

```

user@host> show ospf neighbor extensive
Address Interface State ID Pri Dead
10.5.1.2 ge-1/2/0.1 Full 10.5.1.2 128 33
area 0.0.0.1, opt 0x42, DR 10.5.1.2, BDR 10.5.1.1
Up 06:09:42, adjacent 05:17:50
Link state retransmission list:

 Type LSA ID Adv rtr Seq
Summary 10.8.56.0 172.25.27.82 0x8000004d
Router 10.5.1.94 10.5.1.94 0x8000005c
Network 10.5.24.2 10.5.1.94 0x80000036
Summary 10.8.57.0 172.25.27.82 0x80000024
Extern 1.10.90.0 10.8.1.2 0x80000041
Extern 1.4.109.0 10.6.1.2 0x80000041
Router 10.5.1.190 10.5.1.190 0x8000005f
Network 10.5.48.2 10.5.1.190 0x8000003d
Summary 10.8.58.0 172.25.27.82 0x8000004d
Extern 1.10.91.0 10.8.1.2 0x80000041
Extern 1.4.110.0 10.6.1.2 0x80000041
Router 10.5.1.18 10.5.1.18 0x8000005f
Network 10.5.5.2 10.5.1.18 0x80000033
Summary 10.8.59.0 172.25.27.82 0x8000003a
Summary 10.8.62.0 172.25.27.82 0x80000025

10.5.10.2 ge-1/2/0.10 ExStart 10.5.1.38 128 38
area 0.0.0.1, opt 0x42, DR 10.5.10.2, BDR 10.5.10.1
Up 06:09:42
master, seq 0xac1530f8, rexmit DBD in 2 sec
rexmit LSREQ in 0 sec

10.5.11.2 ge-1/2/0.11 Full 10.5.1.42 128 33
area 0.0.0.1, opt 0x42, DR 10.5.11.2, BDR 10.5.11.1
Up 06:09:42, adjacent 05:27:00
Link state retransmission list:

 Type LSA ID Adv rtr Seq
Summary 10.8.58.0 172.25.27.82 0x8000004d

```

|         |           |              |            |
|---------|-----------|--------------|------------|
| Extern  | 1.10.91.0 | 10.8.1.2     | 0x80000041 |
| Extern  | 1.1.247.0 | 10.5.1.2     | 0x8000003f |
| Extern  | 1.4.110.0 | 10.6.1.2     | 0x80000041 |
| Router  | 10.5.1.18 | 10.5.1.18    | 0x8000005f |
| Network | 10.5.5.2  | 10.5.1.18    | 0x80000033 |
| Summary | 10.8.59.0 | 172.25.27.82 | 0x8000003a |

### show ospf3 neighbor detail

```
user@host> show ospf3 neighbor detail
ID Interface State Pri Dead
10.255.71.13 fe-0/0/2.0 Full 128 30
Neighbor-address fe80::290:69ff:fe9b:e002
area 0.0.0.0, opt 0x13, OSPF3-Intf-Index 2
DR-ID 10.255.71.13, BDR-ID 10.255.71.12
Up 02:51:43, adjacent 02:51:43
```

### show ospf neighbor area area-id

```
user@host >show ospf neighbor area 1.1.1.1
Address Interface State ID Pri Dead
192.168.37.47 so-0/0/0.0 Full 10.255.245.4 128 33
Area 1.1.1.1
192.168.37.55 so-1/0/0.0 Full 10.255.245.5 128 37
Area 1.1.1.1
```

### show ospf neighbor interface interface-name

```
user@host >show ospf neighbor interface so-0/0/0.0
Address Interface State ID Pri Dead
192.168.37.47 so-0/0/0.0 Full 10.255.245.4 128 37
Area 0.0.0.0
192.168.37.47 so-0/0/0.0 Full 10.255.245.4 128 33
Area 1.1.1.1
192.168.37.47 so-0/0/0.0 Full 10.255.245.4 128 32
Area 2.2.2.2
```

### show ospf3 neighbor instance all (OSPFv3 Multiple Family Address Support Enabled)

```
user @host > show ospf3 neighbor instance all
Instance: ina
Realm: ipv6-unicast
ID Interface State Pri Dead
100.1.1.1 fe-0/0/2.0 Full 128 37
Neighbor-address fe80::217:cb00:c87c:8c03
Instance: inb
Realm: ipv4-unicast
ID Interface State Pri Dead
100.1.2.1 fe-0/0/2.1 Full 128 33
Neighbor-address fe80::217:cb00:c97c:8c03
```

## show (ospf | ospf3) overview

---

|                                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                           | <a href="#">Syntax on page 4660</a><br><a href="#">Syntax (EX Series Switch and QFX Series) on page 4660</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Syntax</b>                                   | <code>show (ospf   ospf3) overview</code><br><code>&lt;brief   extensive&gt;</code><br><code>&lt;instance <i>instance-name</i>&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code><br><code>&lt;realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Syntax (EX Series Switch and QFX Series)</b> | <code>show (ospf   ospf3) overview</code><br><code>&lt;brief   extensive&gt;</code><br><code>&lt;instance <i>instance-name</i>&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>                      | Command introduced in Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br><b>realm</b> option introduced in Junos OS Release 9.2.<br>Database protection introduced in Junos 10.2.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>                              | Display Open Shortest Path First (OSPF) overview information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                                  | <b>none</b> —Display standard information about all OSPF neighbors for all routing instances.<br><br><b>brief   extensive</b> —(Optional) Display the specified level of output.<br><br><b>instance <i>instance-name</i></b> —(Optional) Display all OSPF interfaces under the named routing instance.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.<br><br><b>realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)</b> —(Optional) (OSPFv3 only) Display information about the specified OSPFv3 realm, or address family. Use the <b>realm</b> option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default. |
| <b>Required Privilege Level</b>                 | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>List of Sample Output</b>                    | <a href="#">show ospf overview on page 4662</a><br><a href="#">show ospf overview (With Database Protection) on page 4663</a><br><a href="#">show ospf3 overview (With Database Protection) on page 4663</a><br><a href="#">show ospf overview extensive on page 4663</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Output Fields</b>                            | <a href="#">Table 239</a> lists the output fields for the <b>show ospf overview</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

Table 354: show ospf overview Output Fields

| Field name                       | Field Description                                                                                                                                                                        | Level of Output |
|----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Instance</b>                  | OSPF routing instance.                                                                                                                                                                   | All levels      |
| <b>Router ID</b>                 | Router ID of the routing device.                                                                                                                                                         | All levels      |
| <b>Route table index</b>         | Route table index.                                                                                                                                                                       | All levels      |
| <b>Configured overload</b>       | Overload capability is enabled. If the overload timer is also configured, display the time that remains before it is set to expire. This field is not displayed after the timer expires. | All levels      |
| <b>Topology</b>                  | Topology identifier.                                                                                                                                                                     | All levels      |
| <b>Prefix export count</b>       | Number of prefixes exported into OSPF.                                                                                                                                                   | All levels      |
| <b>Full SPF runs</b>             | Number of complete Shortest Path First calculations.                                                                                                                                     | All levels      |
| <b>SPF delay</b>                 | Delay before performing consecutive Shortest Path First calculations.                                                                                                                    | All levels      |
| <b>SPF holddown</b>              | Delay before performing additional Shortest Path First (SPF) calculations after the maximum number of consecutive SPF calculations is reached.                                           | All levels      |
| <b>SPF rapid runs</b>            | Maximum number of Shortest Path First calculations that can be performed in succession before the hold-down timer begins.                                                                | All levels      |
| <b>LSA refresh time</b>          | Refresh period for link-state advertisement (in minutes).                                                                                                                                | All levels      |
| <b>Database protection state</b> | Current state of database protection.                                                                                                                                                    | All levels      |
| <b>Warning threshold</b>         | Threshold at which a warning message is logged (percentage of maximum LSA count).                                                                                                        | All levels      |
| <b>Non self-generated LSAs</b>   | Number of LSAs whose router ID is not equal to the local router ID: <b>Current</b> , <b>Warning</b> (threshold), and <b>Allowed</b> .                                                    | All levels      |
| <b>Ignore time</b>               | How long the database has been in the ignore state.                                                                                                                                      | All levels      |
| <b>Reset time</b>                | How long the database must stay out of the ignore or isolated state before it returns to normal operations.                                                                              | All levels      |
| <b>Ignore count</b>              | Number of times the database has been in the ignore state: <b>Current</b> and <b>Allowed</b> .                                                                                           | All levels      |
| <b>Restart</b>                   | Graceful restart capability: <b>enabled</b> or <b>disabled</b> .                                                                                                                         | All levels      |
| <b>Restart duration</b>          | Time period for complete reacquisition of OSPF neighbors.                                                                                                                                | All levels      |
| <b>Restart grace period</b>      | Time period for which the neighbors should consider the restarting routing device as part of the topology.                                                                               | All levels      |

Table 354: show ospf overview Output Fields (*continued*)

| Field name                    | Field Description                                                                                                                                    | Level of Output  |
|-------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| Graceful restart helper mode  | (OSPFv2) Standard graceful restart helper capability (based on RFC 3623): <b>enabled</b> or <b>disabled</b> .                                        | All levels       |
| Restart-signaling helper mode | (OSPFv2) Restart signaling-based graceful restart helper capability (based on RFC 4811, RFC 4812, and RFC 4813): <b>enabled</b> or <b>disabled</b> . | All levels       |
| Helper mode                   | (OSPFv3) Graceful restart helper capability: <b>enabled</b> or <b>disabled</b> .                                                                     | All levels       |
| Trace options                 | OSPF-specific trace options.                                                                                                                         | <b>extensive</b> |
| Trace file                    | Name of the file to receive the output of the tracing operation.                                                                                     | <b>extensive</b> |
| Area                          | Area number. Area 0.0.0.0 is the backbone area.                                                                                                      | All levels       |
| Stub type                     | Stub type of area: <b>Normal Stub</b> , <b>Not Stub</b> , or <b>Not so Stubby Stub</b> .                                                             | All levels       |
| Authentication Type           | Type of authentication: <b>None</b> , <b>Password</b> , or <b>MD5</b> .                                                                              | All levels       |
| Area border routers           | Number of area border routers.                                                                                                                       | All levels       |
| Neighbors                     | Number of autonomous system boundary routers.                                                                                                        | All levels       |

## Sample Output

### show ospf overview

```

user@host> show ospf overview
Instance: master
 Router ID: 10.255.245.6
 Route table index: 0
 Configured overload, expires in 118 seconds
 LSA refresh time: 50 minutes
Restart: Enabled
 Restart duration: 20 sec
 Restart grace period: 40 sec
 Helper mode: enabled
Area: 0.0.0.0
 Stub type: Not Stub
 Authentication Type: None
 Area border routers: 0, AS boundary routers: 0
Neighbors
 Up (in full state): 0
Topology: default (ID 0)
 Prefix export count: 0
 Full SPF runs: 1
 SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3

```

**show ospf overview (With Database Protection)**

```

user@host> show ospf overview
Instance: master
Router ID: 10.255.112.218
Route table index: 0
LSA refresh time: 50 minutes
Traffic engineering
Restart: Enabled
 Restart duration: 180 sec
 Restart grace period: 210 sec
 Graceful restart helper mode: Enabled
 Restart-signaling helper mode: Enabled
Database protection state: Normal
Warning threshold: 70 percent
Non self-generated LSAs: Current 582, Warning 700, Allowed 1000
Ignore time: 30, Reset time: 60
Ignore count: Current 0, Allowed 1
Area: 0.0.0.0
 Stub type: Not Stub
 Authentication Type: None
 Area border routers: 0, AS boundary routers: 0
 Neighbors
 Up (in full state): 160
Topology: default (ID 0)
Prefix export count: 0
Full SPF runs: 70
SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
Backup SPF: Not Needed

```

**show ospf3 overview (With Database Protection)**

```

user@host> show ospf3 overview
Instance: master
Router ID: 10.255.112.128
Route table index: 0
LSA refresh time: 50 minutes
Database protection state: Normal
Warning threshold: 80 percent
Non self-generated LSAs: Current 3, Warning 8, Allowed 10
Ignore time: 30, Reset time: 60
Ignore count: Current 0, Allowed 2
Area: 0.0.0.0
 Stub type: Not Stub
 Area border routers: 0, AS boundary routers: 0
 Neighbors
 Up (in full state): 1
Topology: default (ID 0)
Prefix export count: 0
Full SPF runs: 7
SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
Backup SPF: Not Needed

```

**show ospf overview extensive**

```

user@host> show ospf overview extensive
Instance: master
Router ID: 1.1.1.103
Route table index: 0
Full SPF runs: 13, SPF delay: 0.200000 sec
LSA refresh time: 50 minutes

```

```
Restart: Disabled
Trace options: lsa
Trace file: /var/log/ospf size 131072 files 10
Area: 0.0.0.0
 Stub type: Not Stub
 Authentication Type: None
 Area border routers: 0, AS boundary routers: 0
 Neighbors
 Up (in full state): 1
```



## show (ospf | ospf3) route

**List of Syntax**    [Syntax on page 4665](#)  
                          [Syntax \(EX Series Switch and QFX Series\) on page 4665](#)

**Syntax**    show (ospf | ospf3) route  
                  <brief | detail | extensive>  
                  <abr | asbr | extern | inter | intra>  
                  <destination>  
                  <instance (default | ipv4-multicast | *instance-name*)>  
                  <logical-system (default | ipv4-multicast | *logical-system-name*)>  
                  <network>  
                  <no-backup-coverage>  
                  <realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>  
                  <router>  
                  <topology (default | ipv4-multicast | *topology-name*)>  
                  <transit>

**Syntax (EX Series Switch and QFX Series)**    show (ospf | ospf3) route  
                  <brief | detail | extensive>  
                  <abr | asbr | extern | inter | intra>  
                  <destination>  
                  <instance *instance-name*>  
                  <network>  
                  <no-backup-coverage>  
                  <router>  
                  <topology (default | ipv4-multicast | *topology-name*)>  
                  <transit>

**Release Information**    Command introduced before Junos OS Release 7.4.  
                                  Command introduced in Junos OS Release 9.0 for EX Series switches.  
                                  **topology** option introduced in Junos OS Release 9.0.  
                                  **realm** option introduced in Junos OS Release 9.2.  
                                  Command introduced in Junos OS Release 11.3 for the QFX Series.

**Description**    Display the entries in the Open Shortest Path First (OSPF) routing table.

**Options**    **none**—Display standard information about all entries in the OSPF routing table for all routing instances and all topologies.

**destination**—Display routes to the specified IP address (with optional destination prefix length).

**brief | detail | extensive**—(Optional) Display the specified level of output.

**abr**—(Optional) Display routes to area border routers.

**asbr**—(Optional) Display routes to autonomous system border routers.

**extern**—(Optional) Display external routes.

**inter**—(Optional) Display interarea routes.

**intra**—(Optional) Display intra-area routes.

**instance** (**default** | **ipv4-multicast** | *instance-name*)—(Optional) Display entries for the default routing instance, the IPv4 multicast routing instance, or for the specified routing instance.

**logical-system** (**default** | **ipv4-multicast** | *logical-system-name*)—(Optional) Perform this operation on the default logical system, the IPv4 multicast logical system, or on a particular logical system.

**network**—(Optional) Display routes to networks.

**no-backup-coverage**—(Optional) Display routes with no backup coverage.

**realm** (**ipv4-multicast** | **ipv4-unicast** | **ipv6-multicast**)—(OSPFv3 only) (Optional) Display entries in the routing table for the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

**router**—(Optional) Display routes to all routers.

**topology** (**default** | **ipv4-multicast** | *topology-name*)—(OSPFv2 only) (Optional) Display routes for the default OSPF topology, IPv4 multicast topology, or for a particular topology.

**transit**—(Optional) (OSPFv3 only) Display OSPFv3 routes to pseudonodes.

**Required Privilege Level**

view

**List of Sample Output**

[show ospf route on page 4668](#)  
[show ospf route detail on page 4668](#)  
[show ospf3 route on page 4668](#)  
[show ospf3 route detail on page 4669](#)  
[show ospf route topology voice on page 4669](#)

**Output Fields**

[Table 355](#) list the output fields for the **show (ospf | ospf3) route** command. Output fields are listed in the approximate order in which they appear.

**Table 355: show (ospf | ospf3) route Output Fields**

| Field Name      | Field Description         | Output Level |
|-----------------|---------------------------|--------------|
| <b>Topology</b> | Name of the topology.     | All levels   |
| <b>Prefix</b>   | Destination of the route. | All levels   |

Table 355: show (ospf | ospf3) route Output Fields (*continued*)

| Field Name                | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Output Level |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------|
| <b>Path type</b>          | How the route was learned: <ul style="list-style-type: none"> <li><b>Inter</b>—Interarea route</li> <li><b>Ext1</b>—External type 1 route</li> <li><b>Ext2</b>—External type 2 route</li> <li><b>Intra</b>—Intra-area route</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                        | All levels   |
| <b>Route type</b>         | The type of routing device from which the route was learned: <ul style="list-style-type: none"> <li><b>AS BR</b>—Route to AS border router.</li> <li><b>Area BR</b>—Route to area border router.</li> <li><b>Area/AS BR</b>—Route to router that is both an <b>Area BR</b> and <b>AS BR</b>.</li> <li><b>Network</b>—Network router.</li> <li><b>Router</b>—Route to a router that is neither an <b>Area BR</b> nor an <b>AS BR</b>.</li> <li><b>Transit</b>—(OSPFv3 only) Route to a pseudonode representing a transit network, LAN, or nonbroadcast multiaccess (NBMA) link.</li> <li><b>Discard</b>—Route to a summary discard.</li> </ul> | All levels   |
| <b>NH Type</b>            | Next-hop type: <b>LSP</b> or <b>IP</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | All levels   |
| <b>Metric</b>             | Route's metric value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | All levels   |
| <b>NH-interface</b>       | (OSPFv3 only) Interface through which the route's next hop is reachable.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | All levels   |
| <b>NH-addr</b>            | (OSPFv3 only) IPv6 address of the next hop.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | All levels   |
| <b>NextHop Interface</b>  | (OSPFv2 only) Interface through which the route's next hop is reachable.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | All levels   |
| <b>Nexthop addr/label</b> | (OSPFv2 only) If the <b>NH Type</b> is <b>IP</b> , then it is the address of the next hop. If the <b>NH Type</b> is <b>LSP</b> , then it is the name of the label-switched path.                                                                                                                                                                                                                                                                                                                                                                                                                                                              | All levels   |
| <b>Area</b>               | Area ID of the route.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | detail       |
| <b>Origin</b>             | Router from which the route was learned.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | detail       |
| <b>Type 7</b>             | Route was learned through a not-so-stubby area (NSSA) link-state advertisement (LSA).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | detail       |
| <b>P-bit</b>              | Route was learned through NSSA LSA and the propagate bit was set.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | detail       |
| <b>Fwd NZ</b>             | Forwarding address is nonzero. <b>Fwd NZ</b> is only displayed if the route is learned through an NSSA LSA.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | detail       |

Table 355: show (ospf | ospf3) route Output Fields (*continued*)

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Output Level  |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| <b>optional-capability</b> | Optional capabilities propagated in the router LSA. This field is in the output for intra-area router routes only (when <b>Route Type</b> is <b>Area BR</b> , <b>AS BR</b> , <b>Area/AS BR</b> , or <b>Router</b> ), not for interarea router routes or network routes. Three bits in this field are defined as follows: <ul style="list-style-type: none"> <li><b>0x4 (V)</b>—Routing device is at the end of a virtual active link.</li> <li><b>0x2 (E)</b>—Routing device is an autonomous system boundary router.</li> <li><b>0x1 (B)</b>—Routing device is an area border router.</li> </ul> | <b>detail</b> |
| <b>priority</b>            | The priority assigned to the prefix: <ul style="list-style-type: none"> <li><b>high</b></li> <li><b>medium</b></li> <li><b>low</b></li> </ul> <p><b>NOTE:</b> The <b>priority</b> field applies only to routes of type <b>Network</b>.</p>                                                                                                                                                                                                                                                                                                                                                        | <b>detail</b> |

## Sample Output

### show ospf route

```

user@host> show ospf route
Prefix Path Route NH Metric NextHop Nexthop
 Type Type Type Interface addr/label
10.255.71.12 Intra Router IP 1 fe-0/0/2.0 192.16.22.86
10.255.71.13/32 Intra Network IP 0 lo0.0
192.168.222.84/30 Intra Network LSP 1 fe-0/0/2.0 1sp-ab

```

### show ospf route detail

```

user@host> show ospf route detail
Topology default Route Table:

Prefix Path Route NH Metric NextHop Nexthop
 Type Type Type Interface addr/label
10.255.14.174 Inter AS BR IP 210 t1-3/0/1.0
 area 0.0.0.2, origin 10.255.14.185
10.255.14.178 Intra Router IP 200 t3-3/1/3.0
 area 0.0.0.2, origin 10.255.14.178, optional-capability 0x0
10.210.1.0/30 Intra Network IP 10 t3-3/1/2.0
 area 0.0.0.2, origin 10.255.14.172, priority medium
100.1.1.1/32 Inter Network IP 210 t1-3/0/1.0
 area 0.0.0.2, origin 10.255.14.185, priority low
112.3.1.0/24 Ext2 Network IP 0 t1-3/0/1.0
 area 0.0.0.0, origin 10.255.14.174, priority high
200.3.3.0/30 Inter Network IP 220 t1-3/0/1.0
 area 0.0.0.2, origin 10.255.14.185, priority high

```

### show ospf3 route

```

user@host> show ospf3 route
Prefix Path Route NH Metric NextHop Nexthop
 Type Type Type Interface addr/label

```

```

10.255.71.13 Intra Router IP 1
NH-interface fe-0/0/2.0, NH-addr fe80::290:69ff:fe9b:e002
10.255.71.13;0.0.0.2
10.255.245.1 Intra Router IP 40 fxp1.1 192.168.36.17

area 0.0.0.0, origin 10.255.245.1 optional-capability 0x0,
10.255.245.3 Intra AS BR IP 1 fxp2.3 192.168.36.34

area 0.0.0.0, origin 10.255.245.3 optional-capability 0x0,
10.255.245.1/32 Intra Network IP 40 fxp1.1 192.168.36.17

area 0.0.0.0, origin 10.255.245.1, priority high
10.255.245.2/32 Intra Network IP 0 lo0.0
area 0.0.0.0, origin 10.255.245.2, priority medium
10.255.245.3/32 Intra Network IP 1 fxp2.3 192.168.36.34

area 0.0.0.0, origin 10.255.245.3, priority low
Intra Transit IP 1
NH-interface fe-0/0/2.0
192::168:222:84/126 Intra Network IP 1
NH-interface fe-0/0/2.0
abcd::71:12/128 Intra Network IP 0
NH-interface lo0.0
abcd::71:13/128 Intra Network LSP 1
NH-interface fe-0/0/2.0, NH-addr lsp-cd

```

### show ospf3 route detail

```

user@host> show ospf3 route detail
Prefix Path Route NH Metric
 type type
10.255.14.174 Intra Area/AS BR IP 110
NH-interface so-1/2/2.0
Area 0.0.0.0, Origin 10.255.14.174, Optional-capability 0x3
10.255.14.178 Intra Router IP 200
NH-interface t3-3/1/3.0
Area 0.0.0.0, Origin 10.255.14.178, Optional-capability 0x0
10.255.14.185;0.0.0.2 Intra Transit IP 200
NH-interface t1-3/0/1.0
NH-interface so-1/2/2.0
Area 0.0.0.0, Origin 10.255.14.185
1000:1:1::1/128 Inter Network IP 110
NH-interface so-1/2/2.0
Area 0.0.0.0, Origin 10.255.14.174, Priority low
1001:2:1::/48 Ext1 Network IP 110
NH-interface so-1/2/2.0
Area 0.0.0.0, Origin 10.255.14.174, Fwd NZ, Priority medium
1002:1:7::/48 Ext2 Network IP 0
NH-interface so-1/2/2.0
Area 0.0.0.0, Origin 10.255.14.174, Fwd NZ, Priority low
1002:3:4::/48 Ext2 Network IP 0
NH-interface so-1/2/2.0
Area 0.0.0.0, Origin 10.255.14.174, Fwd NZ, Priority high
abcd::10:255:14:172/128 Intra Network IP 0
NH-interface lo0.0
Area 0.0.0.0, Origin 10.255.14.172, Priority low

```

### show ospf route topology voice

```
user@host show ospf route topology voice
```

## Topology voice Route Table:

| Prefix          | Path<br>Type | Route<br>Type | NH<br>Type | Metric | NextHop<br>Interface | Nexthop<br>addr/label |
|-----------------|--------------|---------------|------------|--------|----------------------|-----------------------|
| 10.255.8.2      | Intra        | Router        | IP         | 1      | so-0/2/0.0           |                       |
| 10.255.8.3      | Intra        | Router        | IP         | 2      | so-0/2/0.0           |                       |
| 10.255.8.1/32   | Intra        | Network       | IP         | 0      | lo0.0                |                       |
| 10.255.8.2/32   | Intra        | Network       | IP         | 1      | so-0/2/0.0           |                       |
| 10.255.8.3/32   | Intra        | Network       | IP         | 2      | so-0/2/0.0           |                       |
| 192.168.8.0/29  | Intra        | Network       | IP         | 2      | so-0/2/0.0           |                       |
| 192.168.8.44/30 | Intra        | Network       | IP         | 2      | so-0/2/0.0           |                       |
| 192.168.8.46/32 | Intra        | Network       | IP         | 1      | so-0/2/0.0           |                       |
| 192.168.8.48/30 | Intra        | Network       | IP         | 1      | so-0/2/1.0           |                       |
| 192.168.8.52/30 | Intra        | Network       | IP         | 2      | so-0/2/0.0           |                       |
| 192.168.9.44/30 | Intra        | Network       | IP         | 1      | so-0/2/0.0           |                       |
| 192.168.9.45/32 | Intra        | Network       | IP         | 2      | so-0/2/0.0           |                       |

## show (ospf | ospf3) statistics

|                                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                           | <a href="#">Syntax on page 4671</a><br><a href="#">Syntax (EX Series Switch and QFX Series) on page 4671</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Syntax</b>                                   | <pre>show (ospf   ospf3) statistics &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Syntax (EX Series Switch and QFX Series)</b> | <pre>show (ospf   ospf3) statistics &lt;instance <i>instance-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>                      | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>realm</b> option introduced in Junos OS Release 9.2.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>                              | Display OSPF statistics.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                                  | <p><b>none</b>—Display OSPF statistics for all routing instances.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display all statistics for the specified routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)</b>—(Optional) (OSPFv3 only) Display all statistics for the specified OSPFv3 realm, or address family. Use the <b>realm</b> option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.</p> |
| <b>Required Privilege Level</b>                 | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>                    | <ul style="list-style-type: none"> <li>• <a href="#">clear (ospf   ospf3) statistics on page 4625</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>List of Sample Output</b>                    | <a href="#">show ospf statistics on page 4673</a><br><a href="#">show ospf statistics logical-system all on page 4673</a><br><a href="#">show ospf3 statistics on page 4674</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Output Fields</b>                            | <p><a href="#">Table 356</a> lists the output fields for the <b>show (ospf   ospf3) statistics</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

**Table 356: show (ospf | ospf3) statistics Output Fields**

| Field Name                | Field Description                          |
|---------------------------|--------------------------------------------|
| Packet type               | Type of OSPF packet.                       |
| Total Sent/Total Received | Total number of packets sent and received. |

Table 356: show (ospf | ospf3) statistics Output Fields (*continued*)

| Field Name                                         | Field Description                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Last 5 seconds Sent/Last 5 seconds Received</b> | Total number of packets sent and received in the last 5 seconds.                                                                                                                                                                                                                                                                                                                                                          |
| <b>DBDs retransmitted</b>                          | Total number of database description packets retransmitted, and number retransmitted in the last 5 seconds.                                                                                                                                                                                                                                                                                                               |
| <b>LSAs flooded</b>                                | Total number of link-state advertisements flooded, and number flooded in the last 5 seconds.                                                                                                                                                                                                                                                                                                                              |
| <b>LSAs flooded high-prio</b>                      | <p>Total number of high priority link-state advertisements flooded, and number flooded in the last 5 seconds.</p> <p>A link-state advertisement is deemed a high priority if it has changed since it was last sent.</p>                                                                                                                                                                                                   |
| <b>LSAs retransmitted</b>                          | Total number of link-state advertisements retransmitted, and number retransmitted in the last 5 seconds.                                                                                                                                                                                                                                                                                                                  |
| <b>LSAs transmitted to nbr</b>                     | Total number of link-state advertisements transmitted to a neighbor, and number transmitted in the last 5 seconds.                                                                                                                                                                                                                                                                                                        |
| <b>LSAs requested</b>                              | Total number of link-state advertisements requested by neighboring devices, and number requested in the last 5 seconds.                                                                                                                                                                                                                                                                                                   |
| <b>LSAs acknowledged</b>                           | Total number of link-state advertisements acknowledged, and number acknowledged in the last 5 seconds.                                                                                                                                                                                                                                                                                                                    |
| <b>Flood queue depth</b>                           | Total number of entries in the extended queue.                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Total rexmit entries</b>                        | Total number of retransmission entries waiting to be sent from the OSPF routing instance.                                                                                                                                                                                                                                                                                                                                 |
| <b>db summaries</b>                                | Total number of database description summaries waiting to be sent from the OSPF routing instance.                                                                                                                                                                                                                                                                                                                         |
| <b>lsreq entries</b>                               | Total number of link-state request entries waiting to be sent from the OSPF routing instance.                                                                                                                                                                                                                                                                                                                             |
| <b>Receive errors</b>                              | <p>Number and type of receive errors. Some sample receive errors include:</p> <ul style="list-style-type: none"> <li>• <b>mtu mismatches</b></li> <li>• <b>no interface found</b></li> <li>• <b>no virtual link found</b></li> <li>• <b>nssa mismatches</b></li> <li>• <b>stub area mismatches</b></li> <li>• <b>subnet mismatches</b></li> </ul> <p>If there are no receive errors, the output displays <b>none</b>.</p> |



## Sample Output

### show ospf statistics

```

user@host> show ospf statistics
Packet type Total
 Sent Received
Hello 31 14
DbD 9 10
LSReq 2 2
LSUpdate 8 16
LSAck 9 9

 Last 5 seconds
 Sent Received
Hello 2 2
DbD 0 0
LSReq 0 0
LSUpdate 0 0
LSAck 0 0

DBDs retransmitted : 3, last 5 seconds : 0
LSAs flooded : 12, last 5 seconds : 0
LSAs flooded high-prio : 0, last 5 seconds : 0
LSAs retransmitted : 0, last 5 seconds : 0
LSAs transmitted to nbr: 3, last 5 seconds : 0
LSAs requested : 5, last 5 seconds : 0
LSAs acknowledged : 19, last 5 seconds : 0

Flood queue depth : 0
Total rexmit entries : 0
db summaries : 0
lsreq entries : 0

Receive errors:
 862 no interface found
 115923 no virtual link found

```

### show ospf statistics logical-system all

```

user@host> show ospf statistics logical-system all
logical-system: C
OSPF instance is not running

logical-system: B
Packet type Total
 Sent Received
Hello 313740 313653
DbD 3 2
LSReq 1 1
LSUpdate 2752 1825
LSAck 1821 2747

 Last 5 seconds
 Sent Received
Hello 1 0
DbD 0 0
LSReq 0 0
LSUpdate 0 0
LSAck 0 0

DBDs retransmitted : 0, last 5 seconds : 0
LSAs flooded : 2741, last 5 seconds : 0
LSAs flooded high-prio : 10, last 5 seconds : 0
LSAs retransmitted : 0, last 5 seconds : 0
LSAs transmitted to nbr: 2, last 5 seconds : 0
LSAs requested : 1, last 5 seconds : 0
LSAs acknowledged : 1831, last 5 seconds : 0

Flood queue depth : 0
Total rexmit entries : 0
db summaries : 0
lsreq entries : 0

Receive errors:

```

```

None

logical-system: A

Packet type Total Last 5 seconds
 Sent Received Sent Received
Hello 313698 313695 0 0
 DbD 2 3 0 0
 LSReq 1 1 0 0
LSUpdate 1825 2752 0 0
LSAck 2747 1821 0 0

DBDs retransmitted : 0, last 5 seconds : 0
LSAs flooded : 1825, last 5 seconds : 0
LSAs flooded high-prio : 10, last 5 seconds : 0
LSAs retransmitted : 0, last 5 seconds : 0
LSAs transmitted to nbr: 1, last 5 seconds : 0
LSAs requested : 2, last 5 seconds : 0
LSAs acknowledged : 2748, last 5 seconds : 0

Flood queue depth : 0
Total rexmit entries : 0
db summaries : 0
lsreq entries : 0

Receive errors:
None

```

### show ospf3 statistics

```

user@host> show ospf3 statistics

Packet type Total Last 5 seconds
 Sent Received Sent Received
Hello 0 0 0 0
 DbD 0 0 0 0
 LSReq 0 0 0 0
LSUpdate 0 0 0 0
LSAck 0 0 0 0

DBDs retransmitted : 0, last 5 seconds : 0
LSAs flooded : 0, last 5 seconds : 0
LSAs flooded high-prio : 0, last 5 seconds : 0
LSAs retransmitted : 0, last 5 seconds : 0
LSAs transmitted to nbr: 0, last 5 seconds : 0
LSAs requested : 0, last 5 seconds : 0
LSAs acknowledged : 0, last 5 seconds : 0

Flood queue depth : 0
Total rexmit entries : 0
db summaries : 0
lsreq entries : 0

Receive errors:
None

```

# RIP Feature Guide for QFX10000 Switches



# Overview

- [RIP Overview on page 4677](#)

## RIP Overview

---

RIP is an interior gateway protocol (IGP) that uses a distance-vector algorithm to determine the best route to a destination, using the hop count as the metric.

In a RIP network, each router's forwarding table is distributed among the nodes through the flooding of routing table information. Because topology changes are flooded throughout the network, every node maintains the same list of destinations. Packets are then routed to these destinations based on path-cost calculations done at each node in the network.




**NOTE:** In general, the term *RIP* refers to RIP version 1 and RIP version 2.

This topic contains the following sections:

- [Distance-Vector Routing Protocols on page 4677](#)
- [RIP Protocol Overview on page 4678](#)
- [RIP Packets on page 4679](#)
- [Maximizing Hop Count on page 4680](#)
- [Split Horizon and Poison Reverse Efficiency Techniques on page 4680](#)
- [Limitations of Unidirectional Connectivity on page 4681](#)

## Distance-Vector Routing Protocols

Distance-vector routing protocols transmit routing information that includes a distance vector, typically expressed as the number of hops to the destination. This information is flooded out all protocol-enabled interfaces at regular intervals (every 30 seconds in the case of RIP) to create a network map that is stored in each node's local topology database. [Figure 121](#) shows how distance-vector routing works.

**Figure 121: Distance-Vector Protocol**

In Figure 121, Routers A and B have RIP enabled on adjacent interfaces. Router A has known RIP neighbors Routers C, D, and E, which are 1, 2, and 3 hops away, respectively. Router B has known RIP neighbors Routers X, Y, and Z, which are 1, 2, and 3 hops away, respectively. Every 30 seconds, each router floods its entire routing table information out all RIP-enabled interfaces. In this case, flooding exchanges routing table information across the RIP link.

When Router A receives routing information from Router B, it adds 1 to the hop count to determine the new hop count. For example, Router X has a hop count of 1, but when Router A imports the route to X, the new hop count is 2. The imported route also includes information about where the route was learned, so that the original route is imported as a route to Router X through Router B with a hop count of 2.

When multiple routes to the same host are received, RIP uses the distance-vector algorithm to determine which path to import into the forwarding table. The route with the smallest hop count is imported. If there are multiple routes with the same hop count, all are imported into the forwarding table, and traffic is sent along the paths in round-robin fashion.

## RIP Protocol Overview

The RIP IGP uses the Bellman-Ford, or *distance-vector*, algorithm to determine the best route to a destination. RIP uses the hop count as the metric. RIP enables hosts and routers to exchange information for computing routes through an IP-based network. RIP is intended to be used as an IGP in reasonably homogeneous networks of moderate size.

The Junos<sup>®</sup> operating system (Junos OS) supports RIP versions 1 and 2.



**NOTE:** RIP is not supported for multipoint interfaces.

RIP version 1 packets contain the minimal information necessary to route packets through a network. However, this version of RIP does not support authentication or subnetting.

RIP uses User Datagram Protocol (UDP) port 520.

RIP has the following architectural limitations:

- The longest network path cannot exceed 15 hops (assuming that each network, or hop, has a cost of 1).
- RIP depends on counting to infinity to resolve certain unusual situations—When the network consists of several hundred routers, and when a routing loop has formed, the amount of time and network bandwidth required to resolve a next hop might be great.
- RIP uses only a fixed metric to select a route. Other IGPs use additional parameters, such as measured delay, reliability, and load.

## RIP Packets

RIP packets contain the following fields:

- Command—Indicates whether the packet is a request or response message. Request messages seek information for the router's routing table. Response messages are sent periodically and also when a request message is received. Periodic response messages are called *update messages*. Update messages contain the command and version fields and 25 destinations (by default), each of which includes the destination IP address and the metric to reach that destination.



**NOTE:** Beginning with Junos OS Release 11.1, three additional command field types are available to support RIP demand circuits. When you configure an interface for RIP demand circuits, the command field indicates whether the packet is an update request, update response, or update acknowledge message. Neighbor interfaces send updates on demand, not periodically. These command field types are only valid on interfaces configured for RIP demand circuits. For more detailed information, see *RIP Demand Circuits Overview*.

- Version number—Version of RIP that the originating router is running.
- Address family identifier—Address family used by the originating router. The family is always IP.
- Address—IP address included in the packet.
- Metric—Value of the metric advertised for the address.
- Mask—Mask associated with the IP address (RIP version 2 only).
- Next hop—IP address of the next-hop router (RIP version 2 only).

Routing information is exchanged in a RIP network by RIP request and RIP response packets. A router that has just booted can broadcast a RIP request on all RIP-enabled interfaces. Any routers running RIP on those links receive the request and respond by sending a RIP response packet immediately to the router. The response packet contains the routing table information required to build the local copy of the network topology map.

In the absence of RIP request packets, all RIP routers broadcast a RIP response packet every 30 seconds on all RIP-enabled interfaces. The RIP broadcast is the primary way in which topology information is flooded throughout the network.

Once a router learns about a particular destination through RIP, it starts a timer. Every time it receives a new response packet with information about the destination, the router resets the timer to zero. However, if the router receives no updates about a particular destination for 180 seconds, it removes the destination from its RIP routing table.

In addition to the regular transmission of RIP packets every 30 seconds, if a router detects a new neighbor or detects that an interface is unavailable, it generates a triggered update. The new routing information is immediately broadcast out all RIP-enabled interfaces, and the change is reflected in all subsequent RIP response packets.

## Maximizing Hop Count

The successful routing of traffic across a RIP network requires that every node in the network maintain the same view of the topology. Topology information is broadcast between RIP neighbors every 30 seconds. If Router A is many hops away from a new host, Router B, the route to B might take significant time to propagate through the network and be imported into Router A's routing table. If the two routers are 5 hops away from each other, Router A cannot import the route to Router B until 2.5 minutes after Router B is online (30 seconds per hop). For large numbers of hops, the delay becomes prohibitive. To help prevent this delay from growing arbitrarily large, RIP enforces a maximum hop count of 15 hops. Any prefix that is more than 15 hops away is treated as unreachable and assigned a hop count equal to infinity. This maximum hop count is called the *network diameter*.


## Split Horizon and Poison Reverse Efficiency Techniques

Because RIP functions by periodically flooding the entire routing table out to the network, it generates a lot of traffic. The split horizon and poison reverse techniques can help reduce the amount of network traffic originated by RIP hosts and make the transmission of routing information more efficient.

If a router receives a set of route advertisements on a particular interface, RIP determines that those advertisements do not need to be retransmitted out the same interface. This technique, known as *split horizon*, helps limit the amount of RIP routing traffic by eliminating information that other neighbors on that interface have already learned.

[Figure 122](#) shows an example of the split horizon technique.

Figure 122: Split Horizon Example






In [Figure 122](#), Router A advertises routes to Routers C, D, and E to Router B. In this example, Router A can reach Router C in 2 hops. When Router A advertises the route to Router B, Router B imports it as a route to Router C through Router A in 3 hops. If Router B then readvertised this route to Router A, Router A would import it as a route to Router C through Router B in 4 hops. However, the advertisement from Router B to Router A is unnecessary, because Router A can already reach the route in 2 hops. The split horizon technique helps reduce extra traffic by eliminating this type of route advertisement.

Similarly, the poison reverse technique helps to optimize the transmission of routing information and improve the time to reach network convergence. If Router A learns about unreachable routes through one of its interfaces, it advertises those routes as unreachable (hop count of 16) out the same interface. [Figure 123](#) shows an example of the poison reverse technique.

**Figure 123: Poison Reverse Example**




In [Figure 123](#), Router A learns through one of its interfaces that routes to Routers C, D, and E are unreachable. Router A readvertises those routes out the same interface as unreachable. The advertisement informs Router B that Routers C, D, and E are definitely not reachable through Router A.

## Limitations of Unidirectional Connectivity

Because RIP processes routing information based solely on the receipt of routing table updates, it cannot ensure bidirectional connectivity. As [Figure 124](#) shows, RIP networks are limited by their unidirectional connectivity.

**Figure 124: Limitations of Unidirectional Connectivity**



In [Figure 124](#), Routers A and D flood their routing table information to Router B. Because the path to Router E has the fewest hops when routed through Router A, that route is imported into Router B's forwarding table. However, suppose that Router A can transmit traffic but is not receiving traffic from Router B because of an unavailable link or invalid routing policy. If the only route to Router E is through Router A, any traffic destined for Router A is lost, because bidirectional connectivity was never established.

OSPF establishes bidirectional connectivity with a three-way handshake.

- Related Documentation**
- [RIP Configuration Overview](#)
  - [Example: Configuring RIP on page 4685](#)

## PART 70

# Configuring RIP

- [Example: Configuring RIP on page 4685](#)
- [Example: Configuring Authentication for RIP Routes on page 4693](#)
- [Example: Configuring BFD for RIP on page 4701](#)
- [Example: Configuring BFD Authentication for RIP on page 4709](#)
- [Example: Applying Policies to RIP Routes Imported from Neighbors on page 4719](#)
- [Examples: Controlling Traffic with Metrics in a RIP Network on page 4725](#)
- [Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets on page 4735](#)
- [Example: Redistributing Routes Among RIP Instances on page 4741](#)
- [Example: Configuring RIP Timers on page 4749](#)
- [Example: Tracing RIP Protocol Traffic on page 4757](#)
- [Monitoring RIP Routing Information on page ?](#)



# Example: Configuring RIP

- [Understanding Basic RIP Routing on page 4685](#)
- [Example: Configuring a Basic RIP Network on page 4685](#)

## Understanding Basic RIP Routing

---

RIP is an interior gateway protocol (IGP) that routes packets within a single autonomous system (AS). By default, RIP does not advertise the subnets that are directly connected through the device's interfaces. For traffic to pass through a RIP network, you must create a routing policy to export these routes. Advertising only the direct routes propagates the routes to the immediately adjacent RIP-enabled router only. To propagate all routes through the entire RIP network, you must configure the routing policy to export the routes learned through RIP.

### Related Documentation

- [RIP Overview on page 4677](#)
- [Example: Configuring a Basic RIP Network on page 4685](#)

## Example: Configuring a Basic RIP Network

---

This example shows how to configure a basic RIP network.

- [Requirements on page 4685](#)
- [Overview on page 4685](#)
- [Configuration on page 4686](#)
- [Verification on page 4688](#)

### Requirements

No special configuration beyond device initialization is required before configuring this example.

### Overview


In this example, you configure a basic RIP network, create a RIP group called **rip-group**, and add the directly connected interfaces to the RIP group. Then you configure a routing policy to advertise direct routes using policy statement **advertise-routes-through-rip**.

By default, Junos OS does not advertise RIP routes, not even routes that are learned through RIP. To advertise RIP routes, you must configure and apply an export routing policy that advertises RIP-learned and direct routes.

In Junos OS, you do not need to configure the RIP version. RIP version 2 is used by default.

To use RIP on the device, you must configure RIP on all of the RIP interfaces within the network. [Figure 125](#) shows the topology used in this example.

**Figure 125: Sample RIP Network Topology**



“[CLI Quick Configuration](#)” on [page 4686](#) shows the configuration for all of the devices in [Figure 125](#). The section “[Step-by-Step Procedure](#)” on [page 4687](#) describes the steps on Device R1.

## Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter commit from configuration mode.

**Device R1**

```
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces lo0 unit 1 family inet address 172.16.0.1/32
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.1
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

**Device R2**

```
set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set interfaces lo0 unit 2 family inet address 192.168.2.2/32
set interfaces lo0 unit 2 family inet address 172.16.2.2/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2
set protocols rip group rip-group neighbor fe-1/2/1.5
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

**Device R3**

```

set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a basic RIP network:

1. Configure the network interfaces.

This example shows multiple loopback interface addresses to simulate attached networks.

```

[edit interfaces]
user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30

```

```

user@R1# set lo0 unit 1 family inet address 172.16.0.1/32
user@R1# set lo0 unit 1 family inet address 192.168.1.1/32

```

2. Create the RIP group and add the interface.

To configure RIP in Junos OS, you must configure a group that contains the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

```

[edit protocols rip group rip-group]
user@R1# set neighbor fe-1/2/0.1

```

3. Create the routing policy to advertise both direct and RIP-learned routes.

```

[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R1# set from protocol direct
user@R1# set from protocol rip
user@R1# set then accept

```

4. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```

[edit protocols rip group rip-group]
user@R1# set export advertise-routes-through-rip

```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```

user@R1# show interfaces
fe-1/2/0 {

```

```
 unit 1 {
 family inet {
 address 10.0.0.1/30;
 }
 }
 }
}
lo0 {
 unit 1 {
 family inet {
 address 172.16.0.1/32;
 address 192.168.1.1/32;
 }
 }
}

user@R1# show protocols
rip {
 group rip-group {
 export advertise-routes-through-rip;
 neighbor fe-1/2/0.1;
 }
}

user@R1# show policy-options
policy-statement advertise-routes-through-rip {
 term 1 {
 from protocol [direct rip];
 then accept;
 }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

## Verification

Confirm that the configuration is working properly.

- [Checking the Routing Table on page 4688](#)
- [Looking at the Routes That Device R1 Is Advertising to Device R2 on page 4689](#)
- [Looking at the Routes That Device R1 Is Receiving from Device R2 on page 4689](#)
- [Verifying the RIP-Enabled Interfaces on page 4690](#)
- [Verifying the Exchange of RIP Messages on page 4690](#)
- [Verifying Reachability of All Hosts in the RIP Network on page 4691](#)

---

### Checking the Routing Table

**Purpose** Verify that the routing table is populated with the expected routes..

**Action** From operational mode, enter the **show route protocol rip** command.

```
user@R1> show route protocol rip
inet.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```



```

10.0.0.4/30 *[RIP/100] 00:59:15, metric 2, tag 0
 > to 10.0.0.2 via fe-1/2/0.1
172.16.2.2/32 *[RIP/100] 02:52:48, metric 2, tag 0
 > to 10.0.0.2 via fe-1/2/0.1
172.16.3.3/32 *[RIP/100] 00:45:05, metric 3, tag 0
 > to 10.0.0.2 via fe-1/2/0.1
192.168.2.2/32 *[RIP/100] 02:52:48, metric 2, tag 0
 > to 10.0.0.2 via fe-1/2/0.1
192.168.3.3/32 *[RIP/100] 00:45:05, metric 3, tag 0
 > to 10.0.0.2 via fe-1/2/0.1
224.0.0.9/32 *[RIP/100] 00:45:09, metric 1
 MultiRecv

```

**Meaning** The output shows that the routes have been learned from Device R2 and Device R3.

If you were to delete the **from protocol rip** condition in the routing policy on Device R2, the remote routes from Device R3 would not be learned on Device R1.

### Looking at the Routes That Device R1 Is Advertising to Device R2

**Purpose** Verify that Device R1 is sending the expected routes.

**Action** From operational mode, enter the **show route advertising-protocol rip** command.

```

user@R1> show route advertising-protocol rip 10.0.0.1
inet.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.0.1/32 *[Direct/0] 05:18:26
 > via lo0.1
192.168.1.1/32 *[Direct/0] 05:18:25
 > via lo0.1

```

**Meaning** Device R1 is sending routes to its directly connected networks.

### Looking at the Routes That Device R1 Is Receiving from Device R2

**Purpose** Verify that Device R1 is receiving the expected routes.

**Action** From operational mode, enter the **show route receive-protocol rip** command.

```

user@R1> show route receive-protocol rip 10.0.0.2
inet.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.4/30 *[RIP/100] 02:31:22, metric 2, tag 0
 > to 10.0.0.2 via fe-1/2/0.1
172.16.2.2/32 *[RIP/100] 04:24:55, metric 2, tag 0
 > to 10.0.0.2 via fe-1/2/0.1
172.16.3.3/32 *[RIP/100] 02:17:12, metric 3, tag 0
 > to 10.0.0.2 via fe-1/2/0.1
192.168.2.2/32 *[RIP/100] 04:24:55, metric 2, tag 0
 > to 10.0.0.2 via fe-1/2/0.1
192.168.3.3/32 *[RIP/100] 02:17:12, metric 3, tag 0
 > to 10.0.0.2 via fe-1/2/0.1

```

**Meaning** Device R1 is receiving from Device R2 all of Device R2's directly connected networks. Device R1 is also receiving from Device R2 all of Device R3's directly connected networks, which Device R2 learned from Device R3 through RIP.

### Verifying the RIP-Enabled Interfaces

**Purpose** Verify that all RIP-enabled Interfaces are available and active.

**Action** From operational mode, enter the **show rip neighbor** command.

```
user@R1> show rip neighbor
```

| Neighbor   | Local State | Source Address | Destination Address | Send Mode | Receive Mode | In Met |
|------------|-------------|----------------|---------------------|-----------|--------------|--------|
| fe-1/2/0.1 | Up          | 10.0.0.1       | 224.0.0.9           | mcast     | both         | 1      |

**Meaning** The output shows that the RIP-enabled interface on Device R1 is operational.

In general for this command, the output shows a list of the RIP neighbors that are configured on the device. Verify the following information:

- Each configured interface is present. Interfaces are listed in alphabetical order.
- Each configured interface is up. The state of the interface is listed in the **Local State** column. A state of **Up** indicates that the link is passing RIP traffic. A state of **Dn** indicates that the link is not passing RIP traffic. In a point-to-point link, this state generally means that either the end point is not configured for RIP or the link is unavailable.

### Verifying the Exchange of RIP Messages

**Purpose** Verify that RIP messages are being sent and received on all RIP-enabled interfaces.

**Action** From operational mode, enter the **show rip statistics** command.

```
user@R1> show rip statistics
RIPv2 info: port 520; holddown 120s.
 rts learned rts held down rqsts dropped resps dropped
 5 0 0 0

fe-1/2/0.1: 5 routes learned; 2 routes advertised; timeout 180s; update interval 30s
```

| Counter                 | Total | Last 5 min | Last minute |
|-------------------------|-------|------------|-------------|
| Updates Sent            | 2669  | 10         | 2           |
| Triggered Updates Sent  | 2     | 0          | 0           |
| Responses Sent          | 0     | 0          | 0           |
| Bad Messages            | 0     | 0          | 0           |
| RIPv1 Updates Received  | 0     | 0          | 0           |
| RIPv1 Bad Route Entries | 0     | 0          | 0           |
| RIPv1 Updates Ignored   | 0     | 0          | 0           |
| RIPv2 Updates Received  | 2675  | 11         | 2           |
| RIPv2 Bad Route Entries | 0     | 0          | 0           |
| RIPv2 Updates Ignored   | 0     | 0          | 0           |
| Authentication Failures | 0     | 0          | 0           |
| RIP Requests Received   | 0     | 0          | 0           |

|                      |   |   |   |
|----------------------|---|---|---|
| RIP Requests Ignored | 0 | 0 | 0 |
| none                 | 0 | 0 | 0 |

**Meaning** The output shows the number of RIP routes learned. It also shows the number of RIP updates sent and received on the RIP-enabled interfaces. Verify the following information:

- The number of RIP routes learned matches the number of expected routes learned. Subnets learned by direct connectivity through an outgoing interface are not listed as RIP routes.
- RIP updates are being sent on each RIP-enabled interface. If no updates are being sent, the routing policy might not be configured to export routes.
- RIP updates are being received on each RIP-enabled interface. If no updates are being received, the routing policy might not be configured to export routes on the host connected to that subnet. The lack of updates might also indicate an authentication error.

### Verifying Reachability of All Hosts in the RIP Network

**Purpose** Use the **traceroute** command on each loopback address in the network to verify that all hosts in the RIP network are reachable from each Juniper Networks device.

**Action** From operational mode, enter the **traceroute** command.

```
user@R1> traceroute 192.168.3.3
traceroute to 192.168.3.3 (192.168.3.3), 30 hops max, 40 byte packets
 1 10.0.0.2 (10.0.0.2) 1.094 ms 1.028 ms 0.957 ms
 2 192.168.3.3 (192.168.3.3) 1.344 ms 2.245 ms 2.125 ms
```

**Meaning** Each numbered row in the output indicates a routing hop in the path to the host. The three-time increments indicate the round-trip time (RTT) between the device and the hop for each traceroute packet.

To ensure that the RIP network is healthy, verify the following information:

- The final hop in the list is the host you want to reach.
- The number of expected hops to the host matches the number of hops in the traceroute output. The appearance of more hops than expected in the output indicates that a network segment is probably unreachable. It might also indicate that the incoming or outgoing metric on one or more hosts has been set unexpectedly.

**Related Documentation**

- [Understanding Basic RIP Routing on page 4685](#)
- [RIP Configuration Overview](#)

**Related Documentation**

- [Example: Configuring Point-to-Multipoint RIP Networks](#)



# Example: Configuring Authentication for RIP Routes

- [Understanding RIP Authentication on page 4693](#)
- [Example: Configuring Route Authentication for RIP on page 4693](#)
- [Enabling Authentication with Plain-Text Passwords \(CLI Procedure\) on page 4698](#)
- [Enabling Authentication with MD5 Authentication \(CLI Procedure\) on page 4699](#)

## Understanding RIP Authentication

---

RIPv2 provides authentication support so that RIP links can require authentication keys (passwords) before they become active. Authentication provides an additional layer of security on the network beyond the other security features. By default, this authentication is disabled.

Authentication keys can be specified in either plain-text or MD5 form. Authentication requires all routers within the RIP network or subnetwork to have the same authentication type and key (password) configured.

This type of authentication is not supported on RIPv1 networks.

### Related Documentation

- [RIP Overview on page 4677](#)
- [Enabling Authentication with Plain-Text Passwords \(CLI Procedure\) on page 4698](#)
- [Enabling Authentication with MD5 Authentication \(CLI Procedure\) on page 4699](#)

## Example: Configuring Route Authentication for RIP

---

This example shows how to configure authentication for a RIP network.

- [Requirements on page 4694](#)
- [Overview on page 4694](#)
- [Configuration on page 4694](#)
- [Verification on page 4697](#)

## Requirements

No special configuration beyond device initialization is required before configuring this example.

## Overview


You can configure the router to authenticate RIP route queries. By default, authentication is disabled. You can use one of the following authentication methods:

- Simple authentication—Uses a text password that is included in the transmitted packet. The receiving router uses an authentication key (password) to verify the packet.
- MD5 authentication—Creates an encoded checksum that is included in the transmitted packet. The receiving router uses an authentication key (password) to verify the packet's MD5 checksum.

This example shows MD5 authentication.

Figure 126 shows the topology used in this example.

**Figure 126: RIP Authentication Network Topology**



"CLI Quick Configuration" on page 4694 shows the configuration for all of the devices in Figure 126. The section "Step-by-Step Procedure" on page 4695 describes the steps on Device R1.

## Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Device R1**

```

set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces lo0 unit 1 family inet address 172.16.0.1/32
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.1
set protocols rip authentication-type md5
set protocols rip authentication-key "9ONLRBhreK87dsM8i.5FAAtM8XxNb"
set protocols rip traceoptions file rip-authentication-messages
set protocols rip traceoptions flag auth
set protocols rip traceoptions flag packets

```

```

set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

**Device R2**

```

set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set interfaces lo0 unit 2 family inet address 192.168.2.2/32
set interfaces lo0 unit 2 family inet address 172.16.2.2/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2
set protocols rip group rip-group neighbor fe-1/2/1.5
set protocols rip authentication-type md5
set protocols rip authentication-key "9Lf1Xds2gJDHmoJCu1hKvoJGUjq"
set protocols rip traceoptions file rip-authentication-messages
set protocols rip traceoptions flag auth
set protocols rip traceoptions flag packets
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

**Device R3**

```

set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set protocols rip authentication-type md5
set protocols rip authentication-key "9G.UkP5T39tOz3K87V4oz36/Cu"
set protocols rip traceoptions file rip-authentication-messages
set protocols rip traceoptions flag auth
set protocols rip traceoptions flag packets
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure RIP authentication:

1. Configure the network interfaces.

This example shows multiple loopback interface addresses to simulate attached networks.

```

[edit interfaces]
user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30

user@R1# set lo0 unit 1 family inet address 172.16.0.1/32

```

```
user@R1# set lo0 unit 1 family inet address 192.168.1.1/32
```

2. Create the RIP group and add the interface.

To configure RIP in Junos OS, you must configure a group that contains the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

```
[edit protocols rip group rip-group]
user@R1# set neighbor fe-1/2/0.1
```

3. Create the routing policy to advertise both direct and RIP-learned routes.

```
[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R1# set from protocol direct
user@R1# set from protocol rip
user@R1# set then accept
```

4. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```
[edit protocols rip group rip-group]
user@R1# set export advertise-routes-through-rip
```

5. Require MD5 authentication for RIP route queries received on an interface.

The passwords must match on neighboring RIP routers. If the password does not match, the packet is rejected. The password can be from 1 through 16 contiguous characters long and can include any ASCII strings.

Do not enter the password as shown here. The password shown here is the encrypted password that is displayed in the configuration after the actual password is already configured.

```
[edit protocols rip]
user@R1# set authentication-type md5
user@R1# set authentication-key "9ONLRBhreK87dsM8i.5FAtm8XxNb"
```

6. Configure tracing operations to track authentication.

```
[edit protocols rip traceoptions]
user@R1# set file rip-authentication-messages
user@R1# set flag auth
user@R1# set flag packets
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show interfaces
fe-1/2/0 {
 unit 1 {
 family inet {
 address 10.0.0.1/30;
 }
 }
}
lo0 {
```



```

unit 1 {
 family inet {
 address 172.16.0.1/32;
 address 192.168.1.1/32;
 }
}

user@R1# show protocols
rip {
 traceoptions {
 file rip-authentication-messages;
 flag auth;
 flag packets;
 }
 authentication-type md5;
 authentication-key "9ONLRBhreK87dsM8i.5FAtM8XxNb"; ## SECRET-DATA
 group rip-group {
 export advertise-routes-through-rip;
 neighbor fe-1/2/0.1;
 }
}

user@R1# show policy-options
policy-statement advertise-routes-through-rip {
 term 1 {
 from protocol [direct rip];
 then accept;
 }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

## Verification

Confirm that the configuration is working properly.

- [Checking for Authentication Failures on page 4697](#)
- [Verifying That MD5 Authentication Is Enabled in RIP Update Packets on page 4698](#)

### Checking for Authentication Failures

**Purpose** Verify that there are no authentication failures.

**Action** From operational mode, enter the **show rip statistics** command.

```

user@R1> show rip statistics
RIPv2 info: port 520; holddown 120s.
 rts learned rts held down rqsts dropped resps dropped
 5 0 0 0

fe-1/2/0.1: 5 routes learned; 2 routes advertised; timeout 180s; update interval
30s
Counter Total Last 5 min Last minute

Updates Sent 2669 10 2
Triggered Updates Sent 2 0 0

```

|                         |      |    |   |
|-------------------------|------|----|---|
| Responses Sent          | 0    | 0  | 0 |
| Bad Messages            | 0    | 0  | 0 |
| RIPv1 Updates Received  | 0    | 0  | 0 |
| RIPv1 Bad Route Entries | 0    | 0  | 0 |
| RIPv1 Updates Ignored   | 0    | 0  | 0 |
| RIPv2 Updates Received  | 2675 | 11 | 2 |
| RIPv2 Bad Route Entries | 0    | 0  | 0 |
| RIPv2 Updates Ignored   | 0    | 0  | 0 |
| Authentication Failures | 0    | 0  | 0 |
| RIP Requests Received   | 0    | 0  | 0 |
| RIP Requests Ignored    | 0    | 0  | 0 |
| none                    | 0    | 0  | 0 |

**Meaning** The output shows that there are no authentication failures.

### Verifying That MD5 Authentication Is Enabled in RIP Update Packets

**Purpose** Use tracing operations to verify that MD5 authentication is enabled in RIP updates.

**Action** From operational mode, enter the **show log** command.

```
user@R1> show log rip-authentication-messages | match md5
Feb 15 15:45:13.969462 sending msg 0xb9a8c04, 3 rtes (needs MD5)
Feb 15 15:45:43.229867 sending msg 0xb9a8c04, 3 rtes (needs MD5)
Feb 15 15:46:13.174410 sending msg 0xb9a8c04, 3 rtes (needs MD5)
Feb 15 15:46:42.716566 sending msg 0xb9a8c04, 3 rtes (needs MD5)
Feb 15 15:47:11.425076 sending msg 0xb9a8c04, 3 rtes (needs MD5)
...
```

**Meaning** The **(needs MD5)** output shows that all route updates require MD5 authentication.

**Related Documentation**

- [Understanding Basic RIP Routing on page 4685](#)

## Enabling Authentication with Plain-Text Passwords (CLI Procedure)

To configure authentication that requires a plain-text password to be included in the transmitted packet, enable simple authentication by performing these steps on all RIP devices in the network:

1. Navigate to the top of the configuration hierarchy.
2. Perform the configuration tasks described in [Table 357](#).
3. If you are finished configuring the router, commit the configuration.

**Table 357: Configuring Simple RIP Authentication**

| Task                                                         | CLI Configuration Editor                                                       |
|--------------------------------------------------------------|--------------------------------------------------------------------------------|
| Navigate to <b>Rip</b> level in the configuration hierarchy. | From the <b>[edit]</b> hierarchy level, enter<br><br><b>edit protocols rip</b> |

Table 357: Configuring Simple RIP Authentication (*continued*)

| Task                                                                                                                                                             | CLI Configuration Editor                                                                                   |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------|
| Set the authentication type to <b>simple</b> .                                                                                                                   | Set the authentication type to <b>simple</b> :<br><br><b>set authentication-type simple</b>                |
| Set the authentication key to a simple-text password.<br><br>The password can be from 1 through 16 contiguous characters long and can include any ASCII strings. | Set the authentication key to a simple-text password:<br><br><b>set authentication-key <i>password</i></b> |

- Related Documentation**
- [Understanding RIP Authentication on page 4693](#)
  - [RIP Configuration Overview](#)
  - [Enabling Authentication with MD5 Authentication \(CLI Procedure\) on page 4699](#)

## Enabling Authentication with MD5 Authentication (CLI Procedure)

To configure authentication that requires an MD5 password to be included in the transmitted packet, enable MD5 authentication by performing these steps on all RIP devices in the network:

1. Navigate to the top of the configuration hierarchy.
2. Perform the configuration tasks described in [Table 358](#).
3. If you are finished configuring the router, commit the configuration.

Table 358: Configuring MD5 RIP Authentication

| Task                                                                                                                                             | CLI Configuration Editor                                                              |
|--------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|
| Navigate to <b>Rip</b> level in the configuration hierarchy.                                                                                     | From the <b>[edit]</b> hierarchy level, enter<br><br><b>edit protocols rip</b>        |
| Set the authentication type to <b>MD5</b> .                                                                                                      | Set the authentication type to <b>md5</b> :<br><br><b>set authentication-type md5</b> |
| Set the MD5 authentication key (password).<br><br>The key can be from 1 through 16 contiguous characters long and can include any ASCII strings. | Set the MD5 authentication key:<br><br><b>set authentication-key <i>password</i></b>  |

- Related Documentation**
- [Understanding RIP Authentication on page 4693](#)
  - [RIP Configuration Overview](#)
  - [Enabling Authentication with Plain-Text Passwords \(CLI Procedure\) on page 4698](#)

**Related Documentation**

- [Example: Configuring RIP on page 4685](#)

## Example: Configuring BFD for RIP

- [Understanding BFD for RIP on page 4701](#)
- [Example: Configuring BFD for RIP on page 4702](#)

### Understanding BFD for RIP

---

The Bidirectional Forwarding Detection (BFD) Protocol is a simple hello mechanism that detects failures in a network. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the routing device stops receiving a reply after a specified interval. BFD works with a wide variety of network environments and topologies. BFD failure detection times are shorter than RIP detection times, providing faster reaction times to various kinds of failures in the network. Instead of waiting for the routing protocol neighbor timeout, BFD provides rapid detection of link failures. BFD timers are adaptive and can be adjusted to be more or less aggressive. For example, a timer can adapt to a higher value if the adjacency fails, or a neighbor can negotiate a higher value for a timer than the one configured.

BFD enables quick failover between a primary and a secondary routed path. The protocol tests the operational status of the interface multiple times per second. BFD provides for configuration timers and thresholds for failure detection. For example, if the minimum interval is set for 50 milliseconds and the threshold uses the default value of three missed messages, a failure is detected on an interface within 200 milliseconds of the failure.

Intervening devices (for example, an Ethernet LAN switch) hide link-layer failures from routing protocol peers, such as when two routers are connected by way of a LAN switch, where the local interface status remains up even when a physical fault happens on the remote link. Link-layer failure detection times vary, depending on the physical media and the Layer 2 encapsulation. BFD can provide fast failure detection times for all media types, encapsulations, topologies, and routing protocols.

To enable BFD for RIP, both sides of the connection must receive an update message from the peer. By default, RIP does not export any routes. Therefore, you must enable update messages to be sent by configuring an export policy for routes before a BFD session is triggered.

**Related  
Documentation**

- [Example: Configuring BFD for RIP on page 4702](#)

## Example: Configuring BFD for RIP

---

This example shows how to configure Bidirectional Forwarding Detection (BFD) for a RIP network.

- [Requirements on page 4702](#)
- [Overview on page 4702](#)
- [Configuration on page 4704](#)
- [Verification on page 4706](#)

### Requirements

No special configuration beyond device initialization is required before configuring this example.

### Overview

To enable failure detection, include the **bfd-liveness-detection** statement:

```
bfd-liveness-detection {
 detection-time {
 threshold milliseconds;
 }
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 multiplier number;
 no-adaptation;
 transmit-interval {
 threshold milliseconds;
 minimum-interval milliseconds;
 }
 version (1 | automatic);
}
```

Optionally, you can specify the threshold for the adaptation of the detection time by including the **threshold** statement. When the BFD session detection time adapts to a value equal to or greater than the threshold, a single trap and a system log message are sent.

To specify the minimum transmit and receive interval for failure detection, include the **minimum-interval** statement. This value represents the minimum interval at which the local routing device transmits hello packets as well as the minimum interval at which the routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a value in the range from 1 through 255,000 milliseconds. This examples sets a minimum interval of 600 milliseconds.



**NOTE:** BFD is an intensive protocol that consumes system resources. Specifying a minimum interval for BFD of less than 100 ms for Routing Engine-based sessions and 10 ms for distributed BFD sessions can cause undesired BFD flapping.

Depending on your network environment, these additional recommendations might apply:

- For large-scale network deployments with a large number of BFD sessions, specify a minimum interval of 300 ms for Routing Engine-based sessions and 100 ms for distributed BFD sessions.
- For very large-scale network deployments with a large number of BFD sessions, contact Juniper Networks customer support for more information.
- For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing (NSR) is configured, specify a minimum interval of 2500 ms for Routing Engine-based sessions. For distributed BFD sessions with nonstop active routing configured, the minimum interval recommendations are unchanged and depend only on your network deployment.

You can optionally specify the minimum transmit and receive intervals separately.

To specify only the minimum receive interval for failure detection, include the **minimum-receive-interval** statement. This value represents the minimum interval at which the local routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a value in the range from 1 through 255,000 milliseconds.

To specify only the minimum transmit interval for failure detection, include the **transmit-interval minimum-interval** statement. This value represents the minimum interval at which the local routing device transmits hello packets to the neighbor with which it has established a BFD session. You can configure a value in the range from 1 through 255,000 milliseconds.

To specify the number of hello packets not received by a neighbor that causes the originating interface to be declared down, include the **multiplier** statement. The default is 3, and you can configure a value in the range from 1 through 255.

To specify the threshold for detecting the adaptation of the transmit interval, include the **transmit-interval threshold** statement. The threshold value must be greater than the transmit interval.


To specify the BFD version used for detection, include the **version** statement. The default is to have the version detected automatically.

You can trace BFD operations by including the **traceoptions** statement at the **[edit protocols bfd]** hierarchy level.

In Junos OS Release 9.0 and later, you can configure BFD sessions not to adapt to changing network conditions. To disable BFD adaptation, include the **no-adaptation** statement. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

Figure 127 shows the topology used in this example.

Figure 127: RIP BFD Network Topology



"CLI Quick Configuration" on page 4704 shows the configuration for all of the devices in Figure 127. The section "Step-by-Step Procedure" on page 4705 describes the steps on Device R1.

## Configuration

|                                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>CLI Quick Configuration</b> | To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the <b>[edit]</b> hierarchy level.                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Device R1</b>               | <pre> set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30 set protocols bfd traceoptions file bfd-trace set protocols bfd traceoptions flag all set protocols rip group rip-group export advertise-routes-through-rip set protocols rip group rip-group neighbor fe-1/2/0.1 set protocols rip group rip-group bfd-liveness-detection minimum-interval 600 set policy-options policy-statement advertise-routes-through-rip term 1 from protocol direct set policy-options policy-statement advertise-routes-through-rip term 1 from protocol rip set policy-options policy-statement advertise-routes-through-rip term 1 then accept </pre>                                |
| <b>Device R2</b>               | <pre> set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30 set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30 set protocols rip group rip-group export advertise-routes-through-rip set protocols rip group rip-group neighbor fe-1/2/0.2 set protocols rip group rip-group neighbor fe-1/2/1.5 set protocols rip group rip-group bfd-liveness-detection minimum-interval 600 set policy-options policy-statement advertise-routes-through-rip term 1 from protocol direct set policy-options policy-statement advertise-routes-through-rip term 1 from protocol rip set policy-options policy-statement advertise-routes-through-rip term 1 then accept </pre> |
| <b>Device R3</b>               | <pre> set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30 set protocols rip group rip-group export advertise-routes-through-rip </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |



```

set protocols rip group rip-group neighbor fe-1/2/0.6
set protocols rip group rip-group bfd-liveness-detection minimum-interval 600
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a BFD for a RIP network:

1. Configure the network interfaces.

```

[edit interfaces]
user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30

```

2. Create the RIP group and add the interface.

To configure RIP in Junos OS, you must configure a group that contains the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

```

[edit protocols rip group rip-group]
user@R1# set neighbor fe-1/2/0.1

```

3. Create the routing policy to advertise both direct and RIP-learned routes.

```

[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R1# set from protocol direct
user@R1# set from protocol rip
user@R1# set then accept

```

4. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```

[edit protocols rip group rip-group]
user@R1# set export advertise-routes-through-rip

```

5. Enable BFD.

```

[edit protocols rip group rip-group]
user@R1# set bfd-liveness-detection minimum-interval 600

```

6. Configure tracing operations to track BFD messages.

```

[edit protocols bfd traceoptions]
user@R1# set file bfd-trace
user@R1# set flag all

```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```

user@R1# show interfaces
fe-1/2/0 {

```

```
unit 1 {
 family inet {
 address 10.0.0.1/30;
 }
}

user@R1# show protocols
bfd {
 traceoptions {
 file bfd-trace;
 flag all;
 }
}
rip {
 group rip-group {
 export advertise-routes-through-rip;
 bfd-liveness-detection {
 minimum-interval 600;
 }
 neighbor fe-1/2/0.1;
 }
}

user@R1# show policy-options
policy-statement advertise-routes-through-rip {
 term 1 {
 from protocol [direct rip];
 then accept;
 }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

## Verification

Confirm that the configuration is working properly.

- [Verifying That the BFD Sessions Are Up on page 4706](#)
- [Checking the BFD Trace File on page 4707](#)

---

### Verifying That the BFD Sessions Are Up

---

**Purpose** Make sure that the BFD sessions are operating.

**Action** From operational mode, enter the **show bfd session** command.

```
user@R1> show bfd session
```

| Address  | State | Interface  | Detect Time | Transmit Interval | Multiplier |
|----------|-------|------------|-------------|-------------------|------------|
| 10.0.0.2 | Up    | fe-1/2/0.1 | 1.800       | 0.600             | 3          |

1 sessions, 1 clients  
Cumulative transmit rate 1.7 pps, cumulative receive rate 1.7 pps

**Meaning** The output shows that there are no authentication failures.

### Checking the BFD Trace File

**Purpose** Use tracing operations to verify that BFD packets are being exchanged.

**Action** From operational mode, enter the **show log** command.

```
user@R1> show log bfd-trace
Feb 16 10:26:32 PPM Trace: BFD periodic xmit to 10.0.0.2 (IFL 124, rtbl 53,
single-hop port)
Feb 16 10:26:32 Received Downstream TraceMsg (24) len 86:
Feb 16 10:26:32 IfIndex (3) len 4: 0
Feb 16 10:26:32 Protocol (1) len 1: BFD
Feb 16 10:26:32 Data (9) len 61: (hex) 42 46 44 20 70 61 63 6b 65 74 20 66 72
6f 6d 20 31 30 2e
Feb 16 10:26:32 PPM Trace: BFD packet from 10.0.0.1 (IFL 73, rtbl 56, ttl 255)
absorbed
Feb 16 10:26:32 Received Downstream TraceMsg (24) len 60:
Feb 16 10:26:32 IfIndex (3) len 4: 0
Feb 16 10:26:32 Protocol (1) len 1: BFD
Feb 16 10:26:32 Data (9) len 35: (hex) 42 46 44 20 70 65 72 69 6f 64 69 63 20
78 6d 69 74 20 6f
...
```

**Meaning** The output shows the normal functioning of BFD.

**Related Documentation**

- [Understanding BFD for RIP on page 4701](#)

**Related Documentation**

- [Example: Configuring RIP on page 4685](#)
- [Example: Configuring Authentication for RIP Routes on page 4693](#)
- [Example: Configuring Point-to-Multipoint RIP Networks](#)



# Example: Configuring BFD Authentication for RIP

- [Understanding BFD Authentication for RIP on page 4709](#)
- [Example: Configuring BFD Authentication for RIP on page 4711](#)

## Understanding BFD Authentication for RIP

---

BFD enables rapid detection of communication failures between adjacent systems. By default, authentication for BFD sessions is disabled. However, when running BFD over Network Layer protocols, the risk of service attacks can be significant. We strongly recommend using authentication if you are running BFD over multiple hops or through insecure tunnels. Beginning with Junos OS Release 9.6, Junos OS supports authentication for BFD sessions running over RIP. BFD authentication is only supported in the domestic image and is not available in the export image.

You authenticate BFD sessions by specifying an authentication algorithm and keychain, and then associating that configuration information with a security authentication keychain using the keychain name.

The following sections describe the supported authentication algorithms, security keychains, and the level of authentication that can be configured:

- [BFD Authentication Algorithms on page 4709](#)
- [Security Authentication Keychains on page 4710](#)
- [Strict Versus Loose Authentication on page 4710](#)

## BFD Authentication Algorithms

Junos OS supports the following algorithms for BFD authentication:

- **simple-password**—Plain-text password. One to 16 bytes of plain text are used to authenticate the BFD session. One or more passwords can be configured. This method is the least secure and should be used only when BFD sessions are not subject to packet interception.
- **keyed-md5**—Keyed Message Digest 5 hash algorithm for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed MD5 uses one or more secret keys (generated by the algorithm) and a sequence number

that is updated periodically. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than or equal to the last sequence number received. Although more secure than a simple password, this method is vulnerable to replay attacks. Increasing the rate at which the sequence number is updated can reduce this risk.

- **meticulous-keyed-md5**—Meticulous keyed Message Digest 5 hash algorithm. This method works in the same manner as keyed MD5, but the sequence number is updated with every packet. Although more secure than keyed MD5 and simple passwords, this method might take additional time to authenticate the session.
- **keyed-sha-1**—Keyed Secure Hash Algorithm I for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed SHA uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. The key is not carried within the packets. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than the last sequence number received.
- **meticulous-keyed-sha-1**—Meticulous keyed Secure Hash Algorithm I. This method works in the same manner as keyed SHA, but the sequence number is updated with every packet. Although more secure than keyed SHA and simple passwords, this method might take additional time to authenticate the session.



**NOTE:** Nonstop active routing is not supported with meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

---

## Security Authentication Keychains

The security authentication keychain defines the authentication attributes used for authentication key updates. When the security authentication keychain is configured and associated with a protocol through the keychain name, authentication key updates can occur without interrupting routing and signaling protocols.

The authentication keychain contains one or more keychains. Each keychain contains one or more keys. Each key holds the secret data and the time at which the key becomes valid. The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

BFD allows multiple clients per session, and each client can have its own keychain and algorithm defined. To avoid confusion, we recommend specifying only one security authentication keychain.

## Strict Versus Loose Authentication

By default, strict authentication is enabled and authentication is checked at both ends of each BFD session. Optionally, to smooth migration from nonauthenticated sessions to authenticated sessions, you can configure *loose checking*. When loose checking is

configured, packets are accepted without authentication being checked at each end of the session. This feature is intended for transitional periods only.

#### Related Documentation

- [Example: Configuring BFD Authentication for RIP on page 4711](#)
- [bfd-liveness-detection on page 4767](#)
- **authentication-key-chains** statement in the *Junos OS Administration Library for Routing Devices*
- **show bfd session** command in the [CLI Explorer](#)
- [Example: Configuring BFD for RIP on page 4702](#)

## Example: Configuring BFD Authentication for RIP

This example shows how to configure Bidirectional Forwarding Detection (BFD) authentication for a RIP network.

- [Requirements on page 4711](#)
- [Overview on page 4711](#)
- [Configuration on page 4712](#)
- [Verification on page 4715](#)

### Requirements

No special configuration beyond device initialization is required before configuring this example.

The devices must be running Junos OS Release 9.6 or later.


### Overview

Only three steps are needed to configure authentication on a BFD session:

1. Specify the BFD authentication algorithm for the RIP protocol.
2. Associate the authentication keychain with the RIP protocol.
3. Configure the related security authentication keychain.

[Figure 128](#) shows the topology used in this example.

**Figure 128: RIP BFD Authentication Network Topology**



"CLI Quick Configuration" on page 4712 shows the configuration for all of the devices in Figure 128. The section "Step-by-Step Procedure" on page 4713 describes the steps on Device R1.

## Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

**Device R1**

```
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set protocols bfd traceoptions file bfd-trace
set protocols bfd traceoptions flag all
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.1
set protocols rip group rip-group bfd-liveness-detection minimum-interval 600
set protocols rip group rip-group bfd-liveness-detection authentication key-chain bfd-rip
set protocols rip group rip-group bfd-liveness-detection authentication algorithm
 keyed-md5
set protocols rip group rip-group bfd-liveness-detection authentication loose-check
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
set security authentication-key-chains key-chain bfd-rip key 53 secret
 "9d1V2aZGi.fzDiORSeXxDikqmT"
set security authentication-key-chains key-chain bfd-rip key 53 start-time
 "2012-2-16.12:00:00 -0800"
```

**Device R2**

```
set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2
set protocols rip group rip-group neighbor fe-1/2/1.5
set protocols rip group rip-group bfd-liveness-detection minimum-interval 600
set protocols rip group rip-group bfd-liveness-detection authentication key-chain bfd-rip
set protocols rip group rip-group bfd-liveness-detection authentication algorithm
 keyed-md5
set protocols rip group rip-group bfd-liveness-detection authentication loose-check
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
set security authentication-key-chains key-chain bfd-rip key 53 secret
 "9d1V2aZGi.fzDiORSeXxDikqmT"
set security authentication-key-chains key-chain bfd-rip key 53 start-time
 "2012-2-16.12:00:00 -0800"
```

**Device R3**

```
set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set protocols rip group rip-group bfd-liveness-detection minimum-interval 600
```



```

set protocols rip group rip-group bfd-liveness-detection authentication key-chain bfd-rip
set protocols rip group rip-group bfd-liveness-detection authentication algorithm
 keyed-md5
set protocols rip group rip-group bfd-liveness-detection authentication loose-check
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
set security authentication-key-chains key-chain bfd-rip key 53 secret
 "9d1V2aZGi.fzDiORSeXxDikqmT"
set security authentication-key-chains key-chain bfd-rip key 53 start-time
 "2012-2-16.12:00:00 -0800"

```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a BFD authentication:

1. Configure the network interfaces.

```

[edit interfaces]
user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30

```

2. Create the RIP group and add the interface.

To configure RIP in Junos OS, you must configure a group that contains the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

```

[edit protocols rip group rip-group]
user@R1# set neighbor fe-1/2/0.1

```

3. Create the routing policy to advertise both direct and RIP-learned routes.

```

[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R1# set from protocol direct
user@R1# set from protocol rip
user@R1# set then accept

```

4. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```

[edit protocols rip group rip-group]
user@R1# set export advertise-routes-through-rip

```

5. Enable BFD.

```

[edit protocols rip group rip-group]
user@R1# set bfd-liveness-detection minimum-interval 600

```

6. Specify the algorithm (keyed-md5, keyed-sha-1, meticulous-keyed-md5, meticulous-keyed-sha-1, or simple-password) to use.



**NOTE:** Nonstop active routing is not supported with meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

```
[edit protocols rip group rip-group]
user@R1# set bfd-liveness-detection authentication algorithm keyed-md5
```

7. Specify the keychain to be used to associate BFD sessions on RIP with the unique security authentication keychain attributes.

The keychain you specify must match a keychain name configured at the **[edit security authentication key-chains]** hierarchy level.

The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

```
[edit protocols rip group rip-group]
user@R1# set bfd-liveness-detection authentication key-chain bfd-rip
```

8. (Optional) Specify loose authentication checking if you are transitioning from nonauthenticated sessions to authenticated sessions.

```
[edit protocols rip group rip-group]
user@R1# set bfd-liveness-detection authentication loose-check
```

9. Specify the unique security authentication information for BFD sessions:

- The matching keychain name as specified in Step 7.
- At least one key, a unique integer between 0 and 63. Creating multiple keys allows multiple clients to use the BFD session.
- The secret data used to allow access to the session.
- The time at which the authentication key becomes active, in the format *yyyy-mm-dd.hh:mm:ss*.

```
[edit security authentication-key-chains key-chain bfd-rip]
user@R1# set key 53 secret "9d1V2aZGi.fzDiORSeXxDikqmT"
user@R1# set key 53 start-time "2012-2-16.12:00:00 -0800"
```

10. Configure tracing operations to track BFD authentication.

```
[edit protocols bfd traceoptions]
user@R1# set file bfd-trace
user@R1# set flag all
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show security** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show interfaces
```

```

fe-1/2/0 {
 unit 1 {
 family inet {
 address 10.0.0.1/30;
 }
 }
}

user@R1# show protocols
bfd {
 traceoptions {
 file bfd-trace;
 flag all;
 }
}
rip {
 group rip-group {
 export advertise-routes-through-rip;
 bfd-liveness-detection {
 minimum-interval 600;
 }
 neighbor fe-1/2/0.1;
 }
}

user@R1# show policy-options
policy-statement advertise-routes-through-rip {
 term 1 {
 from protocol [direct rip];
 then accept;
 }
}

user@R1# show security
authentication-key-chains {
 key-chain bfd-rip {
 key 53 {
 secret "9d1V2aZGi.fzDiORSeXxDikqmT"; ## SECRET-DATA
 start-time "2012-2-16.12:00:00 -0800";
 }
 }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

## Verification

Confirm that the configuration is working properly.

- [Verifying That the BFD Sessions Are Authenticated on page 4715](#)
- [Viewing Extensive Information About the BFD Authentication on page 4716](#)
- [Checking the BFD Trace File on page 4716](#)

### Verifying That the BFD Sessions Are Authenticated

**Purpose** Make sure that the BFD sessions are authenticated.

**Action** From operational mode, enter the **show bfd session detail** command.

```
user@R1> show bfd session detail
```

| Address  | State | Interface  | Detect Time | Transmit Interval | Multiplier |
|----------|-------|------------|-------------|-------------------|------------|
| 10.0.0.2 | Up    | fe-1/2/0.1 | 1.800       | 0.600             | 3          |

Client RIP, TX interval 0.600, RX interval 0.600, **Authenticate**  
 Session up time 01:39:34  
 Local diagnostic None, remote diagnostic None  
 Remote state Up, version 1  
 Logical system 6, routing table index 53

1 sessions, 1 clients  
 Cumulative transmit rate 1.7 pps, cumulative receive rate 1.7 pps

**Meaning** **Authenticate** is displayed to indicate that BFD authentication is configured.

### Viewing Extensive Information About the BFD Authentication

**Purpose** View the keychain name, the authentication algorithm and mode for each client in the session, and the BFD authentication configuration status.

**Action** From operational mode, enter the **show bfd session extensive** command.

```
user@R1> show bfd session extensive
```

| Address  | State | Interface  | Detect Time | Transmit Interval | Multiplier |
|----------|-------|------------|-------------|-------------------|------------|
| 10.0.0.2 | Up    | fe-1/2/0.1 | 1.800       | 0.600             | 3          |

Client RIP, TX interval 0.600, RX interval 0.600, **Authenticate**  
**keychain bfd-rip, algo keyed-md5, mode loose**  
 Session up time 01:46:29  
 Local diagnostic None, remote diagnostic None  
 Remote state Up, version 1  
 Logical system 6, routing table index 53  
 Min async interval 0.600, min slow interval 1.000  
 Adaptive async TX interval 0.600, RX interval 0.600  
 Local min TX interval 0.600, minimum RX interval 0.600, multiplier 3  
 Remote min TX interval 0.600, min RX interval 0.600, multiplier 3  
 Local discriminator 225, remote discriminator 226  
 Echo mode disabled/inactive  
**Authentication enabled/active, keychain bfd-rip, algo keyed-md5, mode loose**  
 Session ID: 0x300501

1 sessions, 1 clients  
 Cumulative transmit rate 1.7 pps, cumulative receive rate 1.7 pps

**Meaning** The output shows the keychain name, the authentication algorithm and mode for the client in the session, and the BFD authentication configuration status.

### Checking the BFD Trace File

**Purpose** Use tracing operations to verify that BFD packets are being exchanged.

**Action** From operational mode, enter the **show log** command.

```
user@R1> show log bfd-trace
```

```
Feb 16 10:26:32 PPM Trace: BFD periodic xmit to 10.0.0.2 (IFL 124, rtbl 53,
single-hop port)
Feb 16 10:26:32 Received Downstream TraceMsg (24) len 86:
Feb 16 10:26:32 IfIndex (3) len 4: 0
Feb 16 10:26:32 Protocol (1) len 1: BFD
Feb 16 10:26:32 Data (9) len 61: (hex) 42 46 44 20 70 61 63 6b 65 74 20 66 72
6f 6d 20 31 30 2e
Feb 16 10:26:32 PPM Trace: BFD packet from 10.0.0.1 (IFL 73, rtbl 56, ttl 255)
absorbed
Feb 16 10:26:32 Received Downstream TraceMsg (24) len 60:
Feb 16 10:26:32 IfIndex (3) len 4: 0
Feb 16 10:26:32 Protocol (1) len 1: BFD
Feb 16 10:26:32 Data (9) len 35: (hex) 42 46 44 20 70 65 72 69 6f 64 69 63 20
78 6d 69 74 20 6f
...
```

**Meaning** The output shows the normal functioning of BFD.

**Related Documentation**

- [Understanding BFD Authentication for RIP on page 4709](#)

**Related Documentation**

- [Example: Configuring BFD for RIP on page 4701](#)
- [Example: Configuring Authentication for RIP Routes on page 4693](#)
- [Example: Configuring RIP on page 4685](#)



# Example: Applying Policies to RIP Routes Imported from Neighbors

- [Understanding RIP Import Policy on page 4719](#)
- [Example: Applying Policies to RIP Routes Imported from Neighbors on page 4719](#)

## Understanding RIP Import Policy

---

The default RIP import policy is to accept all received RIP routes that pass a sanity check. To filter routes being imported by the local routing device from its neighbors, include the **import** statement, and list the names of one or more policies to be evaluated. If you specify more than one policy, they are evaluated in order (first to last) and the first matching policy is applied to the route. If no match is found, the local routing device does not import any routes.

### Related Documentation

- [Example: Applying Policies to RIP Routes Imported from Neighbors on page 4719](#)

## Example: Applying Policies to RIP Routes Imported from Neighbors

---

This example shows how to configure an import policy in a RIP network.

- [Requirements on page 4719](#)
- [Overview on page 4719](#)
- [Configuration on page 4720](#)
- [Verification on page 4723](#)

### Requirements

No special configuration beyond device initialization is required before configuring this example.


### Overview

In this example, Device R1 has an import policy that accepts the 10/8 and 192.168/16 RIP routes and rejects all other RIP routes. This means that the 172.16/16 RIP routes are excluded from Device R1's routing table.

An export policy is also shown because an export policy is required as part of the minimum configuration for RIP.

Figure 129 shows the topology used in this example.

**Figure 129: RIP Import Policy Network Topology**



"CLI Quick Configuration" on page 4720 shows the configuration for all of the devices in Figure 129. The section "Step-by-Step Procedure" on page 4721 describes the steps on Device R1.

## Configuration

|                                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>CLI Quick Configuration</b> | To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the <b>[edit]</b> hierarchy level, and then enter <b>commit</b> from configuration mode.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Device R1</b>               | <pre> set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30 set interfaces lo0 unit 1 family inet address 192.168.1.1/32 set protocols rip import rip-import set protocols rip group rip-group export advertise-routes-through-rip set protocols rip group rip-group neighbor fe-1/2/0.1 set policy-options policy-statement advertise-routes-through-rip term 1 from protocol direct set policy-options policy-statement advertise-routes-through-rip term 1 from protocol rip set policy-options policy-statement advertise-routes-through-rip term 1 then accept set policy-options policy-statement rip-import term 1 from protocol rip set policy-options policy-statement rip-import term 1 from route-filter 10.0.0.0/8 orlonger set policy-options policy-statement rip-import term 1 from route-filter 192.168.0.0/16 orlonger set policy-options policy-statement rip-import term 1 then accept set policy-options policy-statement rip-import term 2 then reject </pre> |
| <b>Device R2</b>               | <pre> set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30 set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30 set interfaces lo0 unit 2 family inet address 192.168.2.2/32 set interfaces lo0 unit 2 family inet address 172.16.2.2/32 set protocols rip group rip-group export advertise-routes-through-rip set protocols rip group rip-group neighbor fe-1/2/0.2 set protocols rip group rip-group neighbor fe-1/2/1.5 set policy-options policy-statement advertise-routes-through-rip term 1 from protocol direct </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                          |



```

set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

**Device R3**

```

set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a RIP import policy:

1. Configure the network interfaces.

This example shows multiple loopback interface addresses to simulate attached networks.

```

[edit interfaces]
user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30

user@R1# set lo0 unit 1 family inet address 172.16.0.1/32
user@R1# set lo0 unit 1 family inet address 192.168.1.1/32

```

2. Create the RIP group and add the interface.

To configure RIP in Junos OS, you must configure a group that contains the interfaces on which RIP is enabled.

You do not need to enable RIP on the loopback interface.

```

[edit protocols rip group rip-group]
user@R1# set neighbor fe-1/2/0.1

```

3. Create the routing policy to advertise both direct and RIP-learned routes.

```

[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R1# set from protocol direct
user@R1# set from protocol rip
user@R1# set then accept

```

4. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```

[edit protocols rip group rip-group]
user@R1# set export advertise-routes-through-rip

```

5. Configure the import policy.

```

[edit policy-options policy-statement rip-import]

```

```
user@R1# set term 1 from protocol rip
user@R1# set term 1 from route-filter 10.0.0.0/8 orlonger
user@R1# set term 1 from route-filter 192.168.0.0/16 orlonger
user@R1# set term 1 then accept
user@R1# set term 2 then reject
```

6. Apply the import policy.

```
[edit protocols rip]
user@R1# set import rip-import
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show interfaces
fe-1/2/0 {
 unit 1 {
 family inet {
 address 10.0.0.1/30;
 }
 }
}
lo0 {
 unit 1 {
 family inet {
 address 172.16.0.1/32;
 address 192.168.1.1/32;
 }
 }
}

user@R1# show protocols
rip {
 import rip-import;
 group rip-group {
 export advertise-routes-through-rip;
 neighbor fe-1/2/0.1;
 }
}

user@R1# show policy-options
policy-statement advertise-routes-through-rip {
 term 1 {
 from protocol [direct rip];
 then accept;
 }
}
policy-statement rip-import {
 term 1 {
 from {
 protocol rip;
 route-filter 10.0.0.0/8 orlonger;
 route-filter 192.168.0.0/16 orlonger;
 }
 }
}
```

```

 then accept;
 }
 term 2 {
 then reject;
 }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

## Verification

Confirm that the configuration is working properly.

- [Looking at the Routes That Device R2 Is Advertising to Device R1 on page 4723](#)
- [Looking at the Routes That Device R1 Is Receiving from Device R2 on page 4723](#)
- [Checking the Routing Table on page 4724](#)
- [Testing the Import Policy on page 4724](#)

### [Looking at the Routes That Device R2 Is Advertising to Device R1](#)

**Purpose** Verify that Device R2 is sending the expected routes.

**Action** From operational mode, enter the **show route advertising-protocol rip 10.0.0.2** command.

```
user@R2> show route advertising-protocol rip 10.0.0.2
```

```
inet.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```

10.0.0.4/30 *[Direct/0] 2d 01:17:44
 > via fe-1/2/0.5
172.16.2.2/32 *[Direct/0] 2d 04:09:52
 > via lo0.2
172.16.3.3/32 *[RIP/100] 23:40:02, metric 2, tag 0
 > to 10.0.0.6 via fe-1/2/0.5
192.168.2.2/32 *[Direct/0] 2d 04:09:52
 > via lo0.2
192.168.3.3/32 *[RIP/100] 23:40:02, metric 2, tag 0
 > to 10.0.0.6 via fe-1/2/0.5

```

**Meaning** Device R2 is sending 172.16/16 routes to Device R1.

### [Looking at the Routes That Device R1 Is Receiving from Device R2](#)

**Purpose** Verify that Device R1 is receiving the expected routes.

**Action** From operational mode, enter the **show route receive-protocol rip 10.0.0.2** command.

```
user@R1> show route receive-protocol rip 10.0.0.2
```

```
inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```

10.0.0.4/30 *[RIP/100] 01:06:03, metric 2, tag 0
 > to 10.0.0.2 via fe-1/2/0.1
192.168.2.2/32 *[RIP/100] 01:06:03, metric 2, tag 0

```

```
192.168.3.3/32 > to 10.0.0.2 via fe-1/2/0.1
 *[RIP/100] 01:06:03, metric 3, tag 0
 > to 10.0.0.2 via fe-1/2/0.1
```

**Meaning** The output shows that the 172.16/16 routes are excluded.

---

### Checking the Routing Table

**Purpose** Verify that the routing table is populated with the expected routes.

**Action** From operational mode, enter the **show route protocol rip** command.

```
user@R1> show route protocol rip

inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.4/30 *[RIP/100] 00:54:34, metric 2, tag 0
 > to 10.0.0.2 via fe-1/2/0.1
192.168.2.2/32 *[RIP/100] 00:54:34, metric 2, tag 0
 > to 10.0.0.2 via fe-1/2/0.1
192.168.3.3/32 *[RIP/100] 00:54:34, metric 3, tag 0
 > to 10.0.0.2 via fe-1/2/0.1
224.0.0.9/32 *[RIP/100] 00:49:00, metric 1
 MultiRecv
```

**Meaning** The output shows that the routes have been learned from Device R2 and Device R3.

If you delete or deactivate the import policy, the routing table contains the 172.16/16 routes.

---

### Testing the Import Policy

**Purpose** By using the **test policy** command, monitor the number of rejected prefixes.

**Action** From operational mode, enter the **test policy rip-import 172.16/16** command.

```
user@R1> test policy rip-import 172.16/16
Policy rip-import: 0 prefix accepted, 1 prefix rejected
```

**Meaning** The output shows that the policy rejected one prefix.

**Related Documentation**

- [Example: Configuring RIP on page 4685](#)

## Examples: Controlling Traffic with Metrics in a RIP Network

- [Understanding Traffic Control with Metrics in a RIP Network on page 4725](#)
- [Example: Controlling Traffic in a RIP Network with an Incoming Metric on page 4726](#)
- [Example: Controlling Traffic in a RIP Network with an Outgoing Metric on page 4728](#)
- [Example: Configuring the Metric Value Added to Imported RIP Routes on page 4729](#)

### Understanding Traffic Control with Metrics in a RIP Network

To tune a RIP network and control traffic flowing through the network, you increase or decrease the cost of the paths through the network. RIP provides two ways to modify the path cost: an incoming metric and an outgoing metric, which are each set to 1 by default. These metrics are attributes that manually specify the cost of any route advertised through a host. By increasing or decreasing the metrics—and thus the cost—of links throughout the network, you can control packet transmission across the network.

The incoming metric modifies the cost of an individual segment when a route across the segment is imported into the routing table. For example, if you set the incoming metric on the segment to 3, the individual segment cost along the link is changed from 1 to 3. The increased cost affects all route calculations through that link. Other routes that were previously excluded because of a high hop count might now be selected into the router's forwarding table.

The outgoing metric modifies the path cost for all the routes advertised out a particular interface. Unlike the incoming metric, the outgoing metric modifies the routes that other routers are learning and thereby controls the way they send traffic.

If an exported route was learned from a member of the same RIP group, the metric associated with that route is the normal RIP metric. For example, a RIP route with a metric of 5 learned from a neighbor configured with an incoming metric of 2 is advertised with a combined metric of 7 when advertised to neighbors in the same group. However, if this route was learned from a RIP neighbor in a different group or from a different protocol, the route is advertised with the metric value configured in the outgoing metric for that group.

You might want to increase the metric of routes to decrease the likelihood that a particular route is selected and installed in the routing table. This process is sometimes referred to

as *route poisoning*. Some reasons that you might want to poison a route are that the route is relatively expensive to use, or it has relatively low bandwidth.

A route with a higher metric than another route becomes the active route only when the lower-metric route becomes unavailable. In this way, the higher-metric route serves as a backup path.

One way to increase the metric of imported routes is to configure an import policy. Another way is to include the **metric-in** statement in the RIP neighbor configuration. One way to increase the metric of export routes is to configure an export policy. Another way is to include the **metric-out** statement in the RIP neighbor configuration.

#### Related Documentation

- [RIP Overview on page 4677](#)
- [Example: Controlling Traffic in a RIP Network with an Incoming Metric on page 4726](#)
- [Example: Controlling Traffic in a RIP Network with an Outgoing Metric on page 4728](#)

---

## Example: Controlling Traffic in a RIP Network with an Incoming Metric

---

This example shows how to control traffic with an incoming metric.

- [Requirements on page 4726](#)
- [Overview on page 4726](#)
- [Configuration on page 4727](#)
- [Verification on page 4727](#)


### Requirements

Before you begin, define RIP groups, and add interfaces to the groups. Then configure a routing policy to export directly connected routes and routes learned through the RIP routing exchanges. See “[Example: Configuring a Basic RIP Network](#)” on page 4685.

### Overview

In this example, routes to Router D are received by Router A across both of its RIP-enabled interfaces as shown in [Figure 130](#). Because the route through Router B and the route through Router C have the same number of hops, both routes are imported into the forwarding table. However, because the T3 link from Router B to Router D has a higher bandwidth than the T1 link from Router C to Router D, you want traffic to flow from Router A through Router B to Router D.

Figure 130: Controlling Traffic in a RIP Network with the Incoming Metric



To force this flow, you can modify the route metrics as they are imported into Router A's routing table. By setting the incoming metric on the interface from Router A to Router C, you modify the metric on all routes received through that interface. Setting the incoming route metric on Router A changes only the routes in Router A's routing table, and affects only how Router A sends traffic to Router D. Router D's route selection is based on its own routing table, which, by default, includes no adjusted metric values.

In the example, Router C receives a route advertisement from Router D and readvertises the route to Router A. When Router A receives the route, it applies the incoming metric on the interface. Instead of incrementing the metric by 1 (the default), Router A increments it by 3 (the configured incoming metric), giving the route from Router A to Router D through Router C a total path metric of 4. Because the route through Router B has a metric of 2, it becomes the preferred route for all traffic from Router A to Router D.

This example uses a RIP group called **alpha 1** on interface **g3-0/0/0**.

## Configuration

### Step-by-Step Procedure

To control traffic with an incoming metric:

1. Enable RIP on the interface.  

```
[edit protocols rip]
user@host# set group alpha1 neighbor ge-0/0/0
```
2. Set the incoming metric.  

```
[edit protocols rip]
user@host# set metric-in 3
```
3. If you are done configuring the device, commit the configuration.  

```
[edit]
user@host# commit
```

## Verification

To verify that the configuration is working properly, enter the **show route protocols rip** command.

### Related Documentation

- [Understanding Traffic Control with Metrics in a RIP Network on page 4725](#)

- [RIP Configuration Overview](#)
- [Example: Controlling Traffic in a RIP Network with an Outgoing Metric on page 4728](#)
- [Verifying a RIP Configuration](#)

## Example: Controlling Traffic in a RIP Network with an Outgoing Metric

This example shows how to control traffic with an outgoing metric.

- [Requirements on page 4728](#)
- [Overview on page 4728](#)
- [Configuration on page 4729](#)
- [Verification on page 4729](#)

### Requirements


Before you begin:

- Define RIP groups, and add interfaces to the groups. Then configure a routing policy to export directly connected routes and routes learned through RIP routing exchanges. See [“Example: Configuring a Basic RIP Network” on page 4685](#).
- Control traffic with an incoming metric. See [“Example: Controlling Traffic in a RIP Network with an Incoming Metric” on page 4726](#).

### Overview

In this example, each route from Router A to Router D has two hops as shown in [Figure 131](#). However, because the link from Router A to Router B in the RIP group has a higher bandwidth than the link from Router A to Router C in RIP group Alpha 1, you want traffic from Router D to Router A to flow through Router B. To control the way Router D sends traffic to Router A, you can alter the routes that Router D receives by configuring the outgoing metric on Router A's interfaces in the Alpha 1 RIP group.

**Figure 131: Controlling Traffic in a RIP Network with the Outgoing Metric**





If the outgoing metric for the Alpha 1 RIP group—the A-to-C link—is changed to 3, Router D calculates the total path metric from Router A through Router C as 4. In contrast, the unchanged default total path metric to Router A through Router B in the RIP group is 2. The fact that Router A's interfaces belong to two different RIP groups allows you to configure two different outgoing metrics on its interfaces, because you configure path metrics at the group level.

By configuring the outgoing metric, you control the way Router A sends traffic to Router D. By configuring the outgoing metric on the same router, you control the way Router D sends traffic to Router A.

This example uses an outgoing metric of 3.

## Configuration

### Step-by-Step Procedure

To control traffic with an outgoing metric:

1. Set the outgoing metric.  

```
[edit protocols rip group alpha1]
user@host# set metric-out 3
```
2. If you are done configuring the device, commit the configuration.  

```
[edit]
user@host# commit
```

## Verification

To verify that the configuration is working properly, enter the **show protocols rip** command.

### Related Documentation

- [Understanding Traffic Control with Metrics in a RIP Network on page 4725](#)
- [RIP Configuration Overview](#)
- [Verifying a RIP Configuration](#)

---

## Example: Configuring the Metric Value Added to Imported RIP Routes

This example shows how to change the default metric to be added to incoming routes to control the route selection process.

- [Requirements on page 4729](#)
- [Overview on page 4730](#)
- [Configuration on page 4730](#)
- [Verification on page 4733](#)

## Requirements


No special configuration beyond device initialization is required before configuring this example.

## Overview

Normally, when multiple routes are available, RIP selects the route with the lowest hop count. Changing the default metric enables you to control the route selection process such that a route with a higher hop count can be preferred over of a route with a lower hop count.

Figure 132 shows the topology used in this example.

**Figure 132: RIP Incoming Metrics Network Topology**



Device R1 has two potential paths to reach 172.16.2.2/32. The default behavior is to send traffic out the 0.1/30 interface facing Device R2. Suppose, though, that the path through Device R3 is less expensive to use or has higher bandwidth links. This example shows how to use the **metric-in** statement to ensure that Device R1 uses the path through Device R3 to reach 172.16.2.2/32. “[CLI Quick Configuration](#)” on page 4730 shows the configuration for all of the devices in Figure 132. The section “[Step-by-Step Procedure](#)” on page 4731 describes the steps on Device R1.

## Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Device R1**

```

set interfaces fe-1/2/0 unit 1 description to-R2
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces ge-1/2/1 unit 10 description to-R3
set interfaces ge-1/2/1 unit 10 family inet address 10.0.0.10/30
set interfaces lo0 unit 1 family inet address 172.16.0.1/32
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols rip group primary export advertise-routes-through-rip
set protocols rip group primary neighbor ge-1/2/1.10
set protocols rip group secondary export advertise-routes-through-rip
set protocols rip group secondary neighbor fe-1/2/0.1 metric-in 4
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

**Device R2**

```

set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces ge-1/2/1 unit 5 family inet address 10.0.0.5/30

```

```

set interfaces lo0 unit 2 family inet address 192.168.2.2/32
set interfaces lo0 unit 2 family inet address 172.16.2.2/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2
set protocols rip group rip-group neighbor ge-1/2/1.5
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

**Device R3**

```

set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces ge-1/2/1 unit 9 family inet address 10.0.0.9/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set protocols rip group rip-group neighbor ge-1/2/1.9
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a RIP metrics:

1. Configure the network interfaces.

```

[edit interfaces]
user@R1# set fe-1/2/0 unit 1 description to-R2
user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30

user@R1# set ge-1/2/1 unit 10 description to-R3
user@R1# set ge-1/2/1 unit 10 family inet address 10.0.0.10/30

user@R1# set lo0 unit 1 family inet address 172.16.0.1/32
user@R1# set lo0 unit 1 family inet address 192.168.1.1/32

```

2. Create the RIP groups and add the interfaces.

To configure RIP in Junos OS, you must configure one or more groups that contain the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

For the interface that is facing Device R2, the **metric-in 4** setting causes this route to be less likely to be chosen as the active route.

```

[edit protocols rip]
user@R1# set group primary neighbor ge-1/2/1.10
user@R1# set group secondary neighbor fe-1/2/0.1 metric-in 4

```

3. Create the routing policy to advertise both direct and RIP-learned routes.

```
[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R1# set from protocol direct
user@R1# set from protocol rip
user@R1# set then accept
```

4. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```
[edit protocols rip]
user@R1# set group primary export advertise-routes-through-rip
user@R1# set group secondary export advertise-routes-through-rip
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show interfaces
fe-1/2/0 {
 unit 1 {
 description to-R2;
 family inet {
 address 10.0.0.1/30;
 }
 }
}
ge-1/2/1 {
 unit 10 {
 description to-R3;
 family inet {
 address 10.0.0.10/30;
 }
 }
}
lo0 {
 unit 1 {
 family inet {
 address 172.16.0.1/32;
 address 192.168.1.1/32;
 }
 }
}

user@R1# show protocols
rip {
 group primary {
 export advertise-routes-through-rip;
 neighbor ge-1/2/1.10;
 }
 group secondary {
 export advertise-routes-through-rip;
 neighbor fe-1/2/0.1 {
 metric-in 4;
 }
 }
}
```

```

}
user@R1# show policy-options
policy-statement advertise-routes-through-rip {
 term 1 {
 from protocol [direct rip];
 then accept;
 }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

## Verification

Confirm that the configuration is working properly.

- [Verifying That the Expected Route Is Active on page 4733](#)
- [Removing the metric-in Statement on page 4733](#)

### Verifying That the Expected Route Is Active

**Purpose** Make sure that to reach 172.16.2.2/32, Device R1 uses the path through Device R3.

**Action** From operational mode, enter the **show route 172.16.2.2** command.

```

user@R1> show route 172.16.2.2
inet.0: 12 destinations, 12 routes (12 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.2.2/32 *[RIP/100] 00:15:46, metric 3, tag 0
 > to 10.0.0.9 via ge-1/2/1.10

```

**Meaning** The **to 10.0.0.9 via ge-1/2/1.10** output shows that Device R1 uses the path through Device R3 to reach 172.16.2.2/32. The metric for this route is 3.

### Removing the metric-in Statement

**Purpose** Delete or deactivate the **metric-in** statement to see what happens to the 172.16.2.2/32 route.

**Action** 1. From configuration mode, deactivate the **metric-in** statement.

```

[edit protocols rip group secondary neighbor fe-1/2/0.1]
user@R1# deactivate metric-in
user@R1# commit

```

2. From operational mode, enter the **show route 172.16.2.2** command.

```

user@R1> show route 172.16.2.2
inet.0: 12 destinations, 12 routes (12 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.2.2/32 *[RIP/100] 00:00:06, metric 2, tag 0
 > to 10.0.0.2 via fe-1/2/0.1

```

**Meaning** The **to 10.0.0.2 via fe-1/2/0.1** output shows that Device R1 uses the path through Device R2 to reach 172.16.2.2/32. The metric for this route is 2.

**Related Documentation** • [Understanding Traffic Control with Metrics in a RIP Network on page 4725](#)

**Related Documentation** • [Example: Applying Policies to RIP Routes Imported from Neighbors on page 4719](#)

# Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets

- [Understanding the Sending and Receiving of RIPv1 and RIPv2 Packets on page 4735](#)
- [Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets on page 4735](#)

## Understanding the Sending and Receiving of RIPv1 and RIPv2 Packets

---

RIP version 1 (RIPv1) and RIP version 2 (RIPv2) can run simultaneously. This might make sense when you are migrating a RIPv1 network to a RIPv2 network. This also allows interoperation with a device that supports RIPv1 but not RIPv2.

By default, when RIP is enabled on an interface, Junos OS receives both RIPv1 and RIPv2 packets and sends only RIPv2 packets. You can configure this behavior by including the [send](#) and [receive](#) statements in the RIP configuration.

### Related Documentation

- [Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets on page 4735](#)

## Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets

---

This example shows how to configure whether the RIP update messages conform to RIP version 1 (RIPv1) only, to RIP version 2 (RIPv2) only, or to both versions. You can also disable the sending or receiving of update messages.

- [Requirements on page 4735](#)
- [Overview on page 4736](#)
- [Configuration on page 4736](#)
- [Verification on page 4738](#)

### Requirements


No special configuration beyond device initialization is required before configuring this example.

## Overview

By default, when RIP is enabled on an interface, Junos OS receives both RIPv1 and RIPv2 packets and sends only RIPv2 packets.

Figure 133 shows the topology used in this example.

**Figure 133: Sending and Receiving RIPv1 and RIPv2 Packets Network Topology**



In this example, Device R1 is configured to receive only RIPv2 packets.

“CLI Quick Configuration” on page 4736 shows the configuration for all of the devices in Figure 133. The section “Step-by-Step Procedure” on page 4737 describes the steps on Device R1.

## Configuration

|                                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>CLI Quick Configuration</b> | To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the <b>[edit]</b> hierarchy level, and then enter <b>commit</b> from configuration mode.                                                                                                                                                                                                                                                                                                                  |
| <b>Device R1</b>               | <pre> set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30 set interfaces lo0 unit 1 family inet address 172.16.0.1/32 set interfaces lo0 unit 1 family inet address 192.168.1.1/32 set protocols rip group rip-group export advertise-routes-through-rip set protocols rip group rip-group neighbor fe-1/2/0.1 receive version-2 set policy-options policy-statement advertise-routes-through-rip term 1 from protocol direct set policy-options policy-statement advertise-routes-through-rip term 1 from protocol rip set policy-options policy-statement advertise-routes-through-rip term 1 then accept </pre>                |
| <b>Device R2</b>               | <pre> set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30 set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30 set interfaces lo0 unit 2 family inet address 192.168.2.2/32 set interfaces lo0 unit 2 family inet address 172.16.2.2/32 set protocols rip group rip-group export advertise-routes-through-rip set protocols rip group rip-group neighbor fe-1/2/0.2 set protocols rip group rip-group neighbor fe-1/2/1.5 set policy-options policy-statement advertise-routes-through-rip term 1 from protocol direct set policy-options policy-statement advertise-routes-through-rip term 1 from protocol rip </pre> |



```
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

**Device R3**

```
set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a RIP packet versions that can be received:

1. Configure the network interfaces.

```
[edit interfaces]
user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30
```

```
user@R1# set lo0 unit 1 family inet address 172.16.0.1/32
user@R1# set lo0 unit 1 family inet address 192.168.1.1/32
```

2. Create the RIP groups and add the interfaces.

To configure RIP in Junos OS, you must configure one or more groups that contain the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

For the interface that is facing Device R2, the **receive version-2** setting causes this interface to accept only RIPv2 packets.

```
[edit protocols rip group rip-group]
user@R1# set neighbor fe-1/2/0.1 receive version-2
```

3. Create the routing policy to advertise both direct and RIP-learned routes.

```
[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R1# set from protocol direct
user@R1# set from protocol rip
user@R1# set then accept
```

4. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```
[edit protocols rip group rip-group]
user@R1# set export advertise-routes-through-rip
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```

user@R1# show interfaces
fe-1/2/0 {
 unit 1 {
 family inet {
 address 10.0.0.1/30;
 }
 }
}
lo0 {
 unit 1 {
 family inet {
 address 172.16.0.1/32;
 address 192.168.1.1/32;
 }
 }
}

user@R1# show protocols
rip {
 group rip-group {
 export advertise-routes-through-rip;
 neighbor fe-1/2/0.1 {
 receive version-2;
 }
 }
}

user@R1# show policy-options
policy-statement advertise-routes-through-rip {
 term 1 {
 from protocol [direct rip];
 then accept;
 }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

## Verification

Confirm that the configuration is working properly.

### Verifying That the Receive Mode Is Set to RIPv2 Only

**Purpose** Make sure that the interfacing Device R2 is configured to receive only RIPv2 packets, instead of both RIPv1 and RIPv2 packets.

**Action** From operational mode, enter the **show rip neighbor** command.

```
user@R1> show rip neighbor
```

| Neighbor   | Local State | Source Address | Destination Address | Send Mode | Receive Mode | In Met |
|------------|-------------|----------------|---------------------|-----------|--------------|--------|
| fe-1/2/0.1 | Up          | 10.0.0.1       | 224.0.0.9           | mcast     | v2 only      | 1      |

**Meaning** In the output, the **Receive Mode** field displays **v2 only**. The default **Receive Mode** is **both**.

**Related Documentation**

- [Example: Configuring RIP on page 4685](#)



## Example: Redistributing Routes Among RIP Instances

- [Understanding Route Redistribution Among RIP instances on page 4741](#)
- [Example: Redistributing Routes Between Two RIP Instances on page 4742](#)

### Understanding Route Redistribution Among RIP instances

---

You can redistribute routes among RIP processes. Another way to say this is to export RIP routes from one RIP instance to other RIP instances.

In Junos OS, route redistribution among routing instances is accomplished by using routing table groups, also called RIB groups. Routing table groups allow you to import and export routes from a protocol within one routing table into another routing table.



**NOTE:** In contrast, the policy-based import and export functions allow you import and export routes between different protocols within the same routing table.

Consider the following partial example:

```
protocols {
 rip {
 rib-group inet-to-voice;
 }
}
routing-instances {
 voice {
 protocols {
 rip {
 rib-group voice-to-inet;
 }
 }
 }
}
routing-options {
 rib-groups {
 inet-to-voice {
 import-rib [inet.0 voice.inet.0];
 }
 }
}
```

```
 }
 voice-to-inet {
 import-rib [voice.inet.0 inet.0];
 }
}
```

The way to read the **import-rib** statement is as follows. Take the routes from the protocol (RIP, in this case), and import them into the primary (or local) routing table and also into any other routing tables listed after this. The primary routing table is the routing table where the routing table group is being used. That would be either **inet.0** if used in the main routing instance or **voice.inet.0** if used within the routing instance. In the **inet-to-voice** routing table group, **inet.0** is listed first because this routing table group is used in the main routing instance. In the **voice-to-inet** routing table group, **voice.inet.0** is listed first because this routing table group is used in the voice routing instance.

#### Related Documentation

- [Example: Redistributing Routes Between Two RIP Instances on page 4742](#)

---

## Example: Redistributing Routes Between Two RIP Instances

---

This example shows how to configure a RIP routing instance and control the redistribution of RIP routes between the routing instance and the master instance.

- [Requirements on page 4742](#)
- [Overview on page 4742](#)
- [Configuration on page 4743](#)
- [Verification on page 4746](#)

### Requirements


No special configuration beyond device initialization is required before configuring this example.

### Overview

When you create a routing instance called **voice**, Junos OS creates a routing table called **voice.inet.0**. The example shows how to install routes learned through the master RIP instance into the **voice.inet.0** routing table. The example also shows how to install routes learned through the voice routing instance into **inet.0**. This is done by configuring routing table groups. RIP routes are installed into each routing table that belongs to a routing table group.

[Figure 134](#) shows the topology used in this example.

**Figure 134: Redistributing Routes Between RIP Instances Network Topology**



“CLI Quick Configuration” on page 4743 shows the configuration for all of the devices in Figure 134. The section “Step-by-Step Procedure” on page 4744 describes the steps on Device R2.

## Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

**Device R1**

```

set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces lo0 unit 1 family inet address 172.16.0.1/32
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols rip group to-R2 export advertise-routes-through-rip
set protocols rip group to-R2 neighbor fe-1/2/0.1
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

**Device R2**

```

set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set interfaces lo0 unit 2 family inet address 192.168.2.2/32
set interfaces lo0 unit 2 family inet address 172.16.2.2/32
set protocols rip rib-group inet-to-voice
set protocols rip group to-R3 export advertise-routes-through-rip
set protocols rip group to-R3 neighbor fe-1/2/1.5
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
set routing-instances voice protocols rip group to-R1 export advertise-routes-through-rip
set routing-instances voice interface fe-1/2/0.2
set routing-instances voice protocols rip rib-group voice-to-inet
set routing-instances voice protocols rip group to-R1 neighbor fe-1/2/0.2
set routing-options rib-groups inet-to-voice import-rib inet.0
set routing-options rib-groups inet-to-voice import-rib voice.inet.0
set routing-options rib-groups voice-to-inet import-rib voice.inet.0
set routing-options rib-groups voice-to-inet import-rib inet.0

```

**Device R3**

```
set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group to-R2 export advertise-routes-through-rip
set protocols rip group to-R2 neighbor fe-1/2/0.6
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To redistribute RIP routes between routing instances:

1. Configure the network interfaces.

```
[edit interfaces]
user@R2# set fe-1/2/0 unit 2 family inet address 10.0.0.2/30

user@R2# set fe-1/2/1 unit 5 family inet address 10.0.0.5/30

user@R2# set lo0 unit 2 family inet address 192.168.2.2/32
user@R2# set lo0 unit 2 family inet address 172.16.2.2/32
```

2. Create the routing instance, and add one or more interfaces to the routing instance.

```
[edit routing-instances voice]
user@R2# set interface fe-1/2/0.2
```

3. Create the RIP groups and add the interfaces.

```
[edit protocols rip group to-R3]
user@R2# set neighbor fe-1/2/1.5

[edit routing-instances voice protocols rip group to-R1]
user@R2# set neighbor fe-1/2/0.2
```

4. Create the routing table groups.

```
[edit routing-options rib-groups]
user@R2# set inet-to-voice import-rib inet.0
user@R2# set inet-to-voice import-rib voice.inet.0

user@R2# set voice-to-inet import-rib voice.inet.0
user@R2# set voice-to-inet import-rib inet.0
```

5. Apply the routing table groups.

```
[edit protocols rip]
user@R2# set rib-group inet-to-voice

[edit routing-instances voice protocols rip]
user@R2# set rib-group voice-to-inet
```



6. Create the routing policy to advertise both direct and RIP-learned routes.

```
[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R2# set from protocol direct
user@R2# set from protocol rip
user@R2# set then accept
```

7. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```
[edit protocols rip group to-R3]
user@R2# set export advertise-routes-through-rip
```

```
[edit routing-instances voice protocols rip group to-R1]
user@R2# set export advertise-routes-through-rip
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, **show routing-instances**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R2# show interfaces
fe-1/2/0 {
 unit 2 {
 family inet {
 address 10.0.0.2/30;
 }
 }
}
fe-1/2/1 {
 unit 5 {
 family inet {
 address 10.0.0.5/30;
 }
 }
}
lo0 {
 unit 2 {
 family inet {
 address 192.168.2.2/32;
 address 172.16.2.2/32;
 }
 }
}
}

user@R2# show protocols
rip {
 rib-group inet-to-voice;
 group to-R3 {
 export advertise-routes-through-rip;
 neighbor fe-1/2/1.5;
 }
}

user@R2# show policy-options
policy-statement advertise-routes-through-rip {
```

```
term 1 {
 from protocol [direct rip];
 then accept;
}
}

user@R2# show routing-instances
voice {
 interface fe-1/2/0.2;
 protocols {
 rip {
 rib-group voice-to-inet;
 group to-R1 {
 export advertise-routes-through-rip;
 neighbor fe-1/2/0.2;
 }
 }
 }
}

user@R2# show routing-options
rib-groups {
 inet-to-voice {
 import-rib [inet.0 voice.inet.0];
 }
 voice-to-inet {
 import-rib [voice.inet.0 inet.0];
 }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

## Verification

Confirm that the configuration is working properly.

### Checking the Routing Tables

---

**Purpose** Make sure that the routing tables contain the expected routes.

**Action** From operational mode, enter the **show route protocol rip** command.

```
user@R2> show route protocol rip
inet.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.0.1/32 *[RIP/100] 01:58:14, metric 2, tag 0
 > to 10.0.0.1 via fe-1/2/0.2
172.16.3.3/32 *[RIP/100] 02:06:03, metric 2, tag 0
 > to 10.0.0.6 via fe-1/2/0.5
192.168.1.1/32 *[RIP/100] 01:58:14, metric 2, tag 0
 > to 10.0.0.1 via fe-1/2/0.2
192.168.3.3/32 *[RIP/100] 02:06:03, metric 2, tag 0
 > to 10.0.0.6 via fe-1/2/0.5
224.0.0.9/32 *[RIP/100] 01:44:13, metric 1
 MultiRecv

voice.inet.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
```

+ = Active Route, - = Last Active, \* = Both

```
172.16.0.1/32 *[RIP/100] 02:06:03, metric 2, tag 0
 > to 10.0.0.1 via fe-1/2/0.2
172.16.3.3/32 *[RIP/100] 01:58:14, metric 2, tag 0
 > to 10.0.0.6 via fe-1/2/0.5
192.168.1.1/32 *[RIP/100] 02:06:03, metric 2, tag 0
 > to 10.0.0.1 via fe-1/2/0.2
192.168.3.3/32 *[RIP/100] 01:58:14, metric 2, tag 0
 > to 10.0.0.6 via fe-1/2/0.5
224.0.0.9/32 *[RIP/100] 01:44:13, metric 1
 MultiRecv
```

**Meaning** The output shows that both routing tables contain all of the RIP routes.

- Related Documentation**
- [Example: Configuring RIP on page 4685](#)
  - [Example: Applying Policies to RIP Routes Imported from Neighbors on page 4719](#)



## Example: Configuring RIP Timers

- [Understanding RIP Timers on page 4749](#)
- [Example: Configuring RIP Timers on page 4750](#)

### Understanding RIP Timers

---

RIP uses several timers to regulate its operation.

The update interval is the interval at which routes that are learned by RIP are advertised to neighbors. This timer controls the interval between routing updates. The update interval is set to 30 seconds, by default, with a small random amount of time added when the timer is reset. This added time prevents congestion that can occur if all routing devices update their neighbors simultaneously.

To configure the update time interval, include the **update-interval** statement:

```
update-interval seconds;
```

**seconds** can be a value from 10 through 60.

You can set a route timeout interval. If a route is not refreshed after being installed in the routing table by the specified time interval, the route is marked as invalid and is removed from the routing table after the hold-down period expires.

To configure the route timeout for RIP, include the **route-timeout** statement:

```
route-timeout seconds;
```

**seconds** can be a value from 30 through 360. The default value is 180 seconds.

RIP routes expire when either a route timeout limit is met or a route metric reaches infinity, and the route is no longer valid. However, the expired route is retained in the routing table for a specified period so that neighbors can be notified that the route has been dropped. This time period is set by configuring the hold-down timer. Upon expiration of the hold-down timer, the route is removed from the routing table.

To configure the hold-down timer for RIP, include the **holddown** statement:

```
holddown seconds;
```

**seconds** can be a value from 10 through 180. The default value is 120 seconds.



**NOTE:** In Junos OS Release 11.1 and later, a retransmission timer is available for RIP demand circuits.

Generally, we recommend against changing the RIP timers, unless the effects of a change are well understood. The route timeout should be at least three times the update interval. Normally, the default values are best left in effect for standard operations.

**Related  
Documentation**

- [Example: Configuring RIP Timers on page 4750](#)
- [Example: Configuring RIP Demand Circuits](#)

---

## Example: Configuring RIP Timers

This example shows how to configure the RIP update interval and how to monitor the impact of the change.

- [Requirements on page 4750](#)
- [Overview on page 4750](#)
- [Configuration on page 4751](#)
- [Verification on page 4753](#)

### Requirements

No special configuration beyond device initialization is required before configuring this example.

### Overview


In this example, Device R2 has an update interval of 60 seconds for its neighbor, Device R1, and an update interval of 10 seconds for its neighbor, Device R3.

This example is not necessarily practical, but it is shown for demonstration purposes. Generally, we recommend against changing the RIP timers, unless the effects of a change are well understood. Normally, the default values are best left in effect for standard operations.

An export policy is also shown because an export policy is required as part of the minimum configuration for RIP.

[Figure 135](#) shows the topology used in this example.

Figure 135: RIP Timers Network Topology



"CLI Quick Configuration" on page 4751 shows the configuration for all of the devices in Figure 135. The section "Step-by-Step Procedure" on page 4752 describes the steps on Device R2.

## Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Device R1**

```

set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces lo0 unit 1 family inet address 172.16.0.1/32
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.1
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

**Device R2**

```

set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set interfaces lo0 unit 2 family inet address 192.168.2.2/32
set interfaces lo0 unit 2 family inet address 172.16.2.2/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2 update-interval 60
set protocols rip group rip-group neighbor fe-1/2/1.5 update-interval 10
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

**Device R3**

```

set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 rip

```

```
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the RIP update interval:

1. Configure the network interfaces.

This example shows multiple loopback interface addresses to simulate attached networks.

```
[edit interfaces]
```

```
user@R2# set fe-1/2/0 unit 2 family inet address 10.0.0.2/30
```

```
user@R2# set fe-1/2/1 unit 5 family inet address 10.0.0.5/30
```

```
user@R2# set lo0 unit 2 family inet address 192.168.2.2/32
```

```
user@R2# set lo0 unit 2 family inet address 172.16.2.2/32
```

2. Configure different update intervals for the two RIP neighbors.

To configure RIP in Junos OS, you must configure a group that contains the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

```
[edit protocols rip group rip-group]
```

```
user@R2# set neighbor fe-1/2/0.2 update-interval 60
```

```
user@R2# set neighbor fe-1/2/1.5 update-interval 10
```

3. Create the routing policy to advertise both direct and RIP-learned routes.

```
[edit policy-options policy-statement advertise-routes-through-rip term 1]
```

```
user@R2# set from protocol direct
```

```
user@R2# set from protocol rip
```

```
user@R2# set then accept
```

4. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```
[edit protocols rip group rip-group]
```

```
user@R2# set export advertise-routes-through-rip
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R2# show interfaces
```

```
fe-1/2/0 {
 unit 2 {
 family inet {
 address 10.0.0.2/30;
 }
 }
}
```



```

fe-1/2/1 {
 unit 5 {
 family inet {
 address 10.0.0.5/30;
 }
 }
}
lo0 {
 unit 2 {
 family inet {
 address 192.168.2.2/32;
 address 172.16.2.2/32;
 }
 }
}

user@R2# show protocols
rip {
 group rip-group {
 export advertise-routes-through-rip;
 neighbor fe-1/2/0.2 {
 update-interval 60;
 }
 neighbor fe-1/2/1.5 {
 update-interval 10;
 }
 }
}

user@R2# show policy-options
policy-statement advertise-routes-through-rip {
 term 1 {
 from protocol [direct rip];
 then accept;
 }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

## Verification

Confirm that the configuration is working properly.

- [Checking the RIP Updates Sent by Device R2 on page 4753](#)
- [Checking the RIP Updates Received by Device R2 on page 4754](#)
- [Checking the RIP Updates Received by Device R3 on page 4755](#)

### Checking the RIP Updates Sent by Device R2

**Purpose** Make sure that the RIP update packets are sent at the expected interval.

**Action** From operational mode, enter the **show rip statistics** command.

```

user@R2> show rip statistics
RIPv2 info: port 520; holddown 120s.
 rts learned rts held down rqsts dropped resps dropped

```

```

4 2 0 0

fe-1/2/0.2: 2 routes learned; 5 routes advertised; timeout 180s; update interval
60s

```

| Counter                 | Total | Last 5 min | Last minute |
|-------------------------|-------|------------|-------------|
| -----                   | ----- | -----      | -----       |
| Updates Sent            | 123   | 5          | 1           |
| Triggered Updates Sent  | 0     | 0          | 0           |
| Responses Sent          | 0     | 0          | 0           |
| Bad Messages            | 0     | 0          | 0           |
| RIPv1 Updates Received  | 0     | 0          | 0           |
| RIPv1 Bad Route Entries | 0     | 0          | 0           |
| RIPv1 Updates Ignored   | 0     | 0          | 0           |
| RIPv2 Updates Received  | 244   | 10         | 2           |
| RIPv2 Bad Route Entries | 0     | 0          | 0           |
| RIPv2 Updates Ignored   | 0     | 0          | 0           |
| Authentication Failures | 0     | 0          | 0           |
| RIP Requests Received   | 0     | 0          | 0           |
| RIP Requests Ignored    | 0     | 0          | 0           |
| none                    | 0     | 0          | 0           |

```

fe-1/2/1.5: 2 routes learned; 5 routes advertised; timeout 180s; update interval
10s

```

| Counter                 | Total | Last 5 min | Last minute |
|-------------------------|-------|------------|-------------|
| -----                   | ----- | -----      | -----       |
| Updates Sent            | 734   | 32         | 6           |
| Triggered Updates Sent  | 0     | 0          | 0           |
| Responses Sent          | 0     | 0          | 0           |
| Bad Messages            | 0     | 0          | 0           |
| RIPv1 Updates Received  | 0     | 0          | 0           |
| RIPv1 Bad Route Entries | 0     | 0          | 0           |
| RIPv1 Updates Ignored   | 0     | 0          | 0           |
| RIPv2 Updates Received  | 245   | 11         | 2           |
| RIPv2 Bad Route Entries | 0     | 0          | 0           |
| RIPv2 Updates Ignored   | 0     | 0          | 0           |
| Authentication Failures | 0     | 0          | 0           |
| RIP Requests Received   | 0     | 0          | 0           |
| RIP Requests Ignored    | 0     | 0          | 0           |
| none                    | 0     | 0          | 0           |

**Meaning** The **update interval** field shows that the interval is 60 seconds for Neighbor R1 and 10 seconds for Neighbor R3. The **Updates Sent** field shows that Device R2 is sending updates to Device R1 at roughly 1/6 of the rate that it is sending updates to Device R3.

### Checking the RIP Updates Received by Device R2

**Purpose** Make sure that the RIP update packets are sent at the expected interval.

**Action** From operational mode, enter the **show rip statistics** command.

```

user@R1> show rip statistics
RIPv2 info: port 520; holddown 120s.
 rts learned rts held down rqsts dropped resps dropped
 5 0 0 0

```

```

fe-1/2/0.1: 5 routes learned; 2 routes advertised; timeout 180s; update interval
30s

```

| Counter | Total | Last 5 min | Last minute |
|---------|-------|------------|-------------|
| -----   | ----- | -----      | -----       |

|                               |            |          |          |
|-------------------------------|------------|----------|----------|
| Updates Sent                  | 312        | 10       | 2        |
| Triggered Updates Sent        | 2          | 0        | 0        |
| Responses Sent                | 0          | 0        | 0        |
| Bad Messages                  | 0          | 0        | 0        |
| RIPv1 Updates Received        | 0          | 0        | 0        |
| RIPv1 Bad Route Entries       | 0          | 0        | 0        |
| RIPv1 Updates Ignored         | 0          | 0        | 0        |
| <b>RIPv2 Updates Received</b> | <b>181</b> | <b>5</b> | <b>1</b> |
| RIPv2 Bad Route Entries       | 0          | 0        | 0        |
| RIPv2 Updates Ignored         | 0          | 0        | 0        |
| Authentication Failures       | 0          | 0        | 0        |
| RIP Requests Received         | 1          | 0        | 0        |
| RIP Requests Ignored          | 0          | 0        | 0        |
| none                          | 0          | 0        | 0        |

**Meaning** The **RIPv2 Updates Received** field shows the number of updates received from Device R2.

### Checking the RIP Updates Received by Device R3

**Purpose** Make sure that the RIP update packets are sent at the expected interval.

**Action** From operational mode, enter the **show rip statistics** command.

```
user@R3> show rip statistics
```

```
RIPv2 info: port 520; holddown 120s.
```

```
 rts learned rts held down rqsts dropped resps dropped
 5 0 0 0
```

```
fe-1/2/0.6: 5 routes learned; 2 routes advertised; timeout 180s; update interval
30s
```

| Counter                       | Total      | Last 5 min | Last minute |
|-------------------------------|------------|------------|-------------|
| -----                         | -----      | -----      | -----       |
| Updates Sent                  | 314        | 11         | 2           |
| Triggered Updates Sent        | 1          | 0          | 0           |
| Responses Sent                | 0          | 0          | 0           |
| Bad Messages                  | 0          | 0          | 0           |
| RIPv1 Updates Received        | 0          | 0          | 0           |
| RIPv1 Bad Route Entries       | 0          | 0          | 0           |
| RIPv1 Updates Ignored         | 0          | 0          | 0           |
| <b>RIPv2 Updates Received</b> | <b>827</b> | <b>31</b>  | <b>6</b>    |
| RIPv2 Bad Route Entries       | 0          | 0          | 0           |
| RIPv2 Updates Ignored         | 0          | 0          | 0           |
| Authentication Failures       | 0          | 0          | 0           |
| RIP Requests Received         | 0          | 0          | 0           |
| RIP Requests Ignored          | 0          | 0          | 0           |
| none                          | 0          | 0          | 0           |

**Meaning** The **RIPv2 Updates Received** field shows the number of updates received from Device R2.

**Related Documentation**

- [Example: Configuring RIP on page 4685](#)
- [Example: Configuring RIP Demand Circuits](#)



## Example: Tracing RIP Protocol Traffic

- [Understanding RIP Trace Operations on page 4757](#)
- [Example: Tracing RIP Protocol Traffic on page 4758](#)

### Understanding RIP Trace Operations

---

You can trace various types of RIP protocol traffic to help debug RIP protocol issues.

To trace RIP protocol traffic, include the **traceoptions** statement at the **[edit protocols rip]** hierarchy level:

```
traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
}
```

You can specify the following RIP protocol-specific trace options using the **flag** statement:

- **auth**—RIP authentication
- **error**—RIP error packets
- **expiration**—RIP route expiration processing
- **holddown**—RIP hold-down processing
- **nsr-synchronization**—Nonstop active routing synchronization events
- **packets**—All RIP packets
- **request**—RIP information packets
- **trigger**—RIP triggered updates
- **update**—RIP update packets

You can optionally specify one or more of the following **flag** modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted



**NOTE:** Use the **detail** flag modifier with caution as this may cause the CPU to become very busy.

Global tracing options are inherited from the configuration set by the **traceoptions** statement at the **[edit routing-options]** hierarchy level. You can override the following global trace options for the RIP protocol using the **traceoptions flag** statement included at the **[edit protocols rip]** hierarchy level:

- **all**—All tracing operations
- **general**—All normal operations and routing table changes (a combination of the normal and route trace operations)
- **normal**—Normal events
- **policy**—Policy processing
- **route**—Routing information
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing



**NOTE:** Use the **trace** flag **all** with caution because this may cause the CPU to become very busy.

**Related  
Documentation**

- [Example: Tracing RIP Protocol Traffic on page 4758](#)

---

## Example: Tracing RIP Protocol Traffic

This example shows how to trace RIP protocol operations.

- [Requirements on page 4758](#)
- [Overview on page 4758](#)
- [Configuration on page 4759](#)
- [Verification on page 4761](#)

### Requirements

No special configuration beyond device initialization is required before configuring this example.


### Overview

In this example, Device R1 is set to trace routing information updates.

An export policy is also shown because an export policy is required as part of the minimum configuration for RIP.

Figure 136 shows the topology used in this example.

**Figure 136: RIP Trace Operations Network Topology**



"CLI Quick Configuration" on page 4759 shows the configuration for all of the devices in Figure 136. The section "Step-by-Step Procedure" on page 4760 describes the steps on Device R1.

## Configuration

### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

#### Device R1

```
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces lo0 unit 1 family inet address 172.16.0.1/32
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols rip traceoptions file rip-trace-file
set protocols rip traceoptions flag route
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.1
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

#### Device R2

```
set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set interfaces lo0 unit 2 family inet address 192.168.2.2/32
set interfaces lo0 unit 2 family inet address 172.16.2.2/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2
set protocols rip group rip-group neighbor fe-1/2/1.5
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

#### Device R3

```
set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32
```

```
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the RIP update interval:

1. Configure the network interfaces.

This example shows multiple loopback interface addresses to simulate attached networks.

```
[edit interfaces]
user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30
```

```
user@R1# set lo0 unit 1 family inet address 172.16.0.1/32
user@R1# set lo0 unit 1 family inet address 192.168.1.1/32
```

2. Configure the RIP group, and add the interface to the group.

To configure RIP in Junos OS, you must configure a group that contains the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

```
[edit protocols rip group rip-group]
user@R1# set neighbor fe-1/2/0.1
```

3. Configure RIP tracing operations.

```
[edit protocols rip traceoptions]
user@R1# set file rip-trace-file
user@R1# set flag route
```

4. Create the routing policy to advertise both direct and RIP-learned routes.

```
[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R1# set from protocol direct
user@R1# set from protocol rip
user@R1# set then accept
```

5. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```
[edit protocols rip group rip-group]
user@R1# set export advertise-routes-through-rip
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.



```

user@R1# show interfaces
fe-1/2/0 {
 unit 1 {
 family inet {
 address 10.0.0.1/30;
 }
 }
}
lo0 {
 unit 1 {
 family inet {
 address 172.16.0.1/32;
 address 192.168.1.1/32;
 }
 }
}

user@R1# show protocols
rip {
 traceoptions {
 file rip-trace-file;
 flag route;
 }
 group rip-group {
 export advertise-routes-through-rip;
 neighbor fe-1/2/0.1;
 }
}

user@R1# show policy-options
policy-statement advertise-routes-through-rip {
 term 1 {
 from protocol [direct rip];
 then accept;
 }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

## Verification

Confirm that the configuration is working properly.

### Checking the Log File

**Purpose** Make sure that the RIP route updates are logged in the configured log file.

- Action**
1. Deactivate the extra loopback interface address on Device R3.  

```

[edit interfaces lo0 unit 3 family inet]
user@R3# deactivate address 172.16.3.3/32
user@R3# commit

```
  2. From operational mode on Device R1, enter the **show log rip-trace-file** command with the **| match 172.16.3.3** option.  

```

user@R1> show log rip-trace-file | match 172.16.3.3

```

```
Mar 1 11:39:53.975192 Setting RIPv2 rtbit on route 172.16.3.3/32, tsi =
0xbb69228
Mar 1 11:39:59.847118 172.16.3.3/32: metric-in: 16, change: 3 -> 16; # gw:
1, pkt_upd_src 10.0.0.2, inx: 0, rte_upd_src 10.0.0.2
Mar 1 11:39:59.847568 CHANGE 172.16.3.3/32 nhid 591 gw 10.0.0.2
RIP pref 100/0 metric 3/0 fe-1/2/0.1 <Delete Int>
Mar 1 11:39:59.847629 Best route to 172.16.3.3/32 got deleted. Doing route calculation
on the stored rte-info
```

**Meaning** The output shows that the route to 172.16.3.3/32 was deleted.

**Related  
Documentation**

**Related  
Documentation** • [Example: Configuring RIP on page 4685](#)

# Configuration Statements

- [any-sender](#) on page 4764
- [authentication-key](#) on page 4765
- [authentication-type \(Protocols RIP\)](#) on page 4766
- [bfd-liveness-detection](#) on page 4767
- [check-zero](#) on page 4769
- [export](#) on page 4770
- [group \(Protocols RIP\)](#) on page 4771
- [holddown \(Protocols RIP\)](#) on page 4773
- [import \(Protocols RIP\)](#) on page 4774
- [message-size](#) on page 4775
- [metric-in \(Protocols RIP\)](#) on page 4776
- [metric-out](#) on page 4777
- [neighbor](#) on page 4778
- [preference \(Protocols RIP\)](#) on page 4779
- [receive \(Protocols RIP\)](#) on page 4780
- [rib-group \(Protocols RIP\)](#) on page 4781
- [rip](#) on page 4781
- [route-timeout \(Protocols RIP\)](#) on page 4782
- [send \(Protocols RIP\)](#) on page 4783
- [traceoptions \(Protocols RIP\)](#) on page 4784
- [update-interval \(Protocols RIP\)](#) on page 4787

## any-sender

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | any-sender;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i> ],<br>[edit protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.0.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Disable strict sender address checks.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                         |

## authentication-key

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>authentication-key password;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit protocols <a href="#">rip</a>],</p> <p>[edit protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Require authentication for RIP route queries received on an interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <b><i>password</i></b> —Authentication password. If the password does not match, the packet is rejected. The password can be from 1 through 16 contiguous characters long and can include any ASCII strings.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Route Authentication for RIP on page 4693</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## authentication-type (Protocols RIP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>authentication-type type;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a></code><br><code>  <i>neighbor-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  <a href="#">rip</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</code><br><code>[edit protocols <a href="#">rip</a>],</code><br><code>[edit protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a></code><br><code>  <i>neighbor-name</i>]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Configure the type of authentication for RIP route queries received on an interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Default</b>                  | If you do not include this statement and the <b>authentication-key</b> statement, RIP authentication is disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <b>type</b> —Authentication type: <ul style="list-style-type: none"><li>• <b>md5</b>—Use the MD5 algorithm to create an encoded checksum of the packet. The encoded checksum is included in the transmitted packet. The receiving routing device uses the authentication key to verify the packet, discarding it if the digest does not match. This algorithm provides a more secure authentication scheme.</li><li>• <b>none</b>—Disable authentication. If <b>none</b> is configured, the configured authentication key is ignored.</li><li>• <b>simple</b>—Use a simple password. The password is included in the transmitted packet, which makes this method of authentication relatively insecure. The password can be from 1 through 16 contiguous letters or digits long.</li></ul>                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">authentication-key on page 4765</a></li><li>• <a href="#">Example: Configuring Route Authentication for RIP on page 4693</a></li><li>• <a href="#">authentication-key on page 4765</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

## bfd-liveness-detection

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre> bfd-liveness-detection {     authentication {         algorithm <i>algorithm-name</i>;         key-chain <i>key-chain-name</i>;         loose-check;     }     detection-time {         threshold <i>milliseconds</i>;     }     minimum-interval <i>milliseconds</i>;     minimum-receive-interval <i>milliseconds</i>;     multiplier <i>number</i>;     no-adaptation;     transmit-interval {         minimum-interval <i>milliseconds</i>;         threshold <i>milliseconds</i>;     }     version (1   automatic); } </pre>                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols rip <b>group</b> <i>group-name</i>],<br/> [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br/> rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],<br/> [edit protocols rip <b>group</b> <i>group-name</i>],<br/> [edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <b>neighbor</b><br/> <i>neighbor-name</i>]</p>                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 8.0.</p> <p><b>detection-time threshold</b> and <b>transmit-interval threshold</b> options introduced in Junos OS Release 8.2.</p> <p>Support for logical systems introduced in Junos OS Release 8.3.</p> <p><b>no-adaptation</b> option introduced in Junos OS Release 9.0.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>authentication algorithm</b>, <b>authentication key-chain</b>, and <b>authentication loose-check</b> options introduced in Junos OS Release 9.6.</p> <p><b>authentication algorithm</b>, <b>authentication key-chain</b>, and <b>authentication loose-check</b> options introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> |
| <b>Description</b>         | <p>Configure bidirectional failure detection timers and authentication.</p> <p>The remaining statements are explained separately. See <a href="#">CLI Explorer</a>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>             | <p><b>authentication algorithm</b> <i>algorithm-name</i> —Configure the algorithm used to authenticate the specified BFD session: <b>simple-password</b>, <b>keyed-md5</b>, <b>keyed-sha-1</b>, <b>meticulous-keyed-md5</b>, or <b>meticulous-keyed-sha-1</b>.</p> <p><b>authentication key-chain</b> <i>key-chain-name</i> —Associate a security key with the specified BFD session using the name of the security keychain. The name you specify must</p>                                                                                                                                                                                                                                                                                                                                                                    |

match one of the keychains configured in the **authentication-key-chains key-chain** statement at the **[edit security]** hierarchy level.

**authentication loose-check**—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication may not be configured at both ends of the BFD session.

**detection-time threshold *milliseconds***—Configure a threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

**minimum-interval *milliseconds***—Configure the minimum interval after which the local routing device transmits a hello packet and then expects to receive a reply from the neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately using the **transmit-interval minimum-interval** and **minimum-receive-interval** statements.

**Range:** 1 through 255,000 milliseconds

**minimum-receive-interval *milliseconds***—Configure the minimum interval after which the local routing device expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the **minimum-interval** statement.

**Range:** 1 through 255,000 milliseconds

**multiplier *number***—Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

**Range:** 1 through 255

**Default:** 3

**no-adaptation**—Configure BFD sessions not to adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

**transmit-interval threshold *milliseconds***—Configure the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

**Range:** 0 through 4,294,967,295 ( $2^{32} - 1$ )

**transmit-interval minimum-interval *milliseconds***—Configure a minimum interval after which the local routing device transmits hello packets to a neighbor. Optionally, instead of using this statement, you can configure the minimum transmit interval using the **minimum-interval** statement.

**Range:** 1 through 255,000

**version**—Configure the BFD version to detect: **1** (BFD version 1) or **automatic** (autodetect the BFD version).

**Default:** automatic




|                                 |                                                                                                                                                                                                           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring BFD for RIP on page 4702</a></li> <li>• <a href="#">Example: Configuring BFD Authentication for RIP on page 4711</a></li> </ul> |

## check-zero

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | (check-zero   no-check-zero);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <b>rip</b>],<br/>         [edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],<br/>         [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>rip</b>],<br/>         [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],<br/>         [edit protocols <b>rip</b>],<br/>         [edit protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],<br/>         [edit routing-instances <i>routing-instance-name</i> protocols <b>rip</b>],<br/>         [edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | <p>Check whether the reserved fields in a RIP packet are zero:</p> <ul style="list-style-type: none"> <li>• <b>check-zero</b>—Discard version 1 packets that have nonzero values in the reserved fields and version 2 packets that have nonzero values in the fields that must be zero. This default behavior implements the RIP version 1 and version 2 specifications.</li> <li>• <b>no-check-zero</b>—Receive RIP version 1 packets with nonzero values in the reserved fields or RIP version 2 packets with nonzero values in the fields that must be zero. This is in spite of the fact that they are being sent in violation of the specifications in RFC 1058 and RFC 2453.</li> </ul>                                                                                                                                                                                                                                            |
| <b>Default</b>                  | <b>check-zero</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

## export

---

|                                                                                                                                                                                                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                           | <code>export [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>                                                                                                                                                                                  | [edit logical-systems <i>logical-system-name</i> protocols rip <b>group</b> <i>group-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>rip <b>group</b> <i>group-name</i> ],<br>[edit protocols rip <b>group</b> <i>group-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols rip <b>group</b> <i>group-name</i> ]                                                                                  |
| <b>Release Information</b>                                                                                                                                                                              | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>                                                                                                                                                                                      | <p>Apply a policy to routes being exported to the neighbors.</p> <p>By default, RIP does not export routes it has learned to its neighbors. To enable RIP to export routes, apply one or more export policies.</p> <p>If no routes match the policies, the local routing device does not export any routes to its neighbors. Export policies override any metric values determined through calculations involving the values configured with the <b>metric-in</b> and <b>metric-out</b> statements.</p> |
| <div> <b>NOTE:</b> The export policy on RIP does not support manipulating routing information of the next hop.</div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                                                                                                                                                                                          | <i>policy-names</i> —Name of one or more policies.                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b>                                                                                                                                                                         | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>                                                                                                                                                                            | <ul style="list-style-type: none"><li>• <a href="#">import on page 4774</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                   |

## group (Protocols RIP)

```
Syntax group group-name {
 bfd-liveness-detection {
 authentication {
 algorithm algorithm-name;
 key-chain key-chain-name;
 loose-check;
 }
 detection-time {
 threshold milliseconds;
 }
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 transmit-interval {
 threshold milliseconds;
 minimum-interval milliseconds;
 }
 multiplier number;
 version (0 | 1 | automatic);
 }
 demand-circuit;
 export policy;
 max-retrans-time seconds;
 metric-out metric;
 preference number;
 route-timeout seconds;
 update-interval seconds;
 neighbor neighbor-name {
 authentication-key password;
 authentication-type type;
 bfd-liveness-detection {
 authentication {
 algorithm algorithm-name;
 key-chain key-chain-name;
 loose-check;
 }
 detection-time {
 threshold milliseconds;
 }
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 transmit-interval {
 threshold milliseconds;
 minimum-interval milliseconds;
 }
 multiplier number;
 version (0 | 1 | automatic);
 }
 (check-zero | no-check-zero);
 demand-circuit;
 import policy-name;
 max-retrans-time seconds;
 message-size number;
 }
}
```

```
metric-in metric;
metric-out metric;
receive receive-options;
route-timeout seconds;
send send-options;
update-interval seconds;
}
}
```

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols [rip](#)],  
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols  
[rip](#)],  
[edit protocols [rip](#)],  
[edit routing-instances *routing-instance-name* protocols [rip](#)]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 12.1 for the QFX Series.

**Description** Configure a set of RIP neighbors that share an export policy and metric. The export policy and metric govern what routes to advertise to neighbors in a given group. Each group must contain at least one neighbor. You should create a group for every export policy.

**Options** *group-name*—Name of a group, up to 16 characters long.

The remaining statements are explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Configuring RIP on page 4685](#)

## holddown (Protocols RIP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>holddown seconds;</code>                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit protocols <a href="#">rip</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>]</p>                                   |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                            |
| <b>Description</b>              | <p>Configure how long the expired route is retained in the routing table before being removed.</p> <p>When the hold-down timer runs on RIP demand circuits, routes are advertised as unreachable on other interfaces. When the hold-down timer expires, the route is removed from the routing table if all destinations detect that the route is unreachable or the remaining destinations are down.</p> |
| <b>Options</b>                  | <p><b>seconds</b>—Estimated time to wait before making updates to the routing table.</p> <p><b>Range:</b> 10 through 180 seconds</p> <p><b>Default:</b> 180 seconds</p>                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring RIP Timers on page 4750</a></li> <li>• <a href="#">RIP Demand Circuits Overview</a></li> </ul>                                                                                                                                                                                                                                 |

## import (Protocols RIP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>import [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a></code><br><code>  <i>neighbor-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  <a href="#">rip</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</code><br><code>[edit protocols <a href="#">rip</a>],</code><br><code>[edit protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a></code><br><code>  <i>neighbor-name</i>]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Apply one or more policies to routes being imported by the local routing device from neighbors.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more policies.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Applying Policies to RIP Routes Imported from Neighbors on page 4719</a></li><li>• <a href="#">Junos OS Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</a></li><li>• <a href="#">export on page 4770</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

## message-size

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>message-size <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <i>rip</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <i>neighbor neighbor-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>rip</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <i>neighbor neighbor-name</i>],</p> <p>[edit protocols <i>rip</i>],</p> <p>[edit protocols rip group <i>group-name</i> <i>neighbor neighbor-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <i>rip</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <i>neighbor neighbor-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Specify the number of route entries to be included in every RIP update message. To ensure interoperability with other vendors' equipment, use the standard of 25 route entries per message.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <p><i>number</i>—Number of route entries per update message.</p> <p><b>Range:</b> 25 through 255 entries</p> <p><b>Default:</b> 25 entries</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## metric-in (Protocols RIP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>metric-in <i>metric</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a></code><br><code>  <i>neighbor-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  <a href="#">rip</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</code><br><code>[edit protocols <a href="#">rip</a>],</code><br><code>[edit protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a></code><br><code>  <i>neighbor-name</i>]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Specify the metric to add to incoming routes when the routing device advertises into RIP routes that were learned from other protocols. Use this statement to configure the routing device to prefer RIP routes learned through a specific neighbor.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <i>metric</i> —Metric value.<br><b>Range:</b> 1 through 16<br><b>Default:</b> 1                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring the Metric Value Added to Imported RIP Routes on page 4729</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |



## metric-out

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>metric-out <i>metric</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],</p> <p>[edit protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Specify the metric value to add to routes transmitted to the neighbor. Use this statement to control how other routing devices prefer RIP routes sent from this neighbor.                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <p><b><i>metric</i></b>—Metric value.</p> <p><b>Range:</b> 1 through 16</p> <p><b>Default:</b> 1</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                             |

## neighbor

---

**Syntax**    `neighbor neighbor-name {  
          authentication-key password;  
          authentication-type type;  
          bfd-liveness-detection {  
            authentication {  
              algorithm algorithm-name;  
              key-chain key-chain-name;  
              loose-check;  
            }  
            detection-time {  
              threshold milliseconds;  
            }  
            minimum-interval milliseconds;  
            minimum-receive-interval milliseconds;  
            transmit-interval {  
              threshold milliseconds;  
              minimum-interval milliseconds;  
            }  
            multiplier number;  
            version (0 | 1 | automatic);  
          }  
          (check-zero | no-check-zero);  
          demand-circuit;  
          import policy-name;  
          max-retrans-time seconds;  
          message-size number;  
          metric-in metric;  
          metric-out metric;  
          receive receive-options;  
          route-timeout seconds;  
          send send-options;  
          update-interval seconds;  
          }`

**Hierarchy Level**    [edit logical-systems *logical-system-name* protocols rip *group group-name*],  
                          [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols  
                          rip *group group-name*],  
                          [edit protocols rip *group group-name*],  
                          [edit routing-instances *routing-instance-name* protocols rip *group group-name*]

**Release Information**    Statement introduced before Junos OS Release 7.4.  
                              Statement introduced in Junos OS Release 9.0 for EX Series switches.

**Description**    Configure neighbor-specific RIP parameters, thereby overriding the defaults set for the routing device.

**Options**    *neighbor-name*—Name of an interface over which a routing device communicates to its neighbors.

The remaining statements are explained separately. See [CLI Explorer](#).

|                           |                                                             |
|---------------------------|-------------------------------------------------------------|
| <b>Required Privilege</b> | routing—To view this statement in the configuration.        |
| <b>Level</b>              | routing-control—To add this statement to the configuration. |

## preference (Protocols RIP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>preference <i>preference</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols rip <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip <b>group</b> <i>group-name</i>],</p> <p>[edit protocols rip <b>group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip <b>group</b> <i>group-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | <p>Specify the preference of external routes learned by RIP as compared to those learned from other routing protocols.</p> <p>By default, Junos OS assigns a preference of 100 to routes that originate from RIP. When Junos OS determines a route's preference to become the active route, the software selects the route with the lowest preference and installs this route into the forwarding table.</p>                       |
| <b>Options</b>                  | <p><b>preference</b>—Preference value. A lower value indicates a more preferred route.</p> <p><b>Range:</b> 0 through 4,294,967,295 (<math>2^{32} - 1</math>)</p> <p><b>Default:</b> 100</p>                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Route Preferences Overview</i></li> </ul>                                                                                                                                                                                                                                                                                                                                              |

## receive (Protocols RIP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>receive receive-options;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a></code><br><code>  <i>neighbor-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  <a href="#">rip</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</code><br><code>[edit protocols <a href="#">rip</a>],</code><br><code>[edit protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a></code><br><code>  <i>neighbor-name</i>]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Configure RIP receive options.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <i>receive-options</i> —One of the following: <ul style="list-style-type: none"><li>• <b>both</b>—Accept both RIP version 1 and version 2 packets.</li><li>• <b>none</b>—Do not receive RIP packets.</li><li>• <b>version-1</b>—Accept only RIP version 1 packets.</li><li>• <b>version-2</b>—Accept only RIP version 2 packets.</li></ul> <b>Default:</b> <b>both</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets on page 4735</a></li><li>• <a href="#">send on page 4783</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

## rib-group (Protocols RIP)

|                                 |                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>rib-group <i>group-name</i>;</code>                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a> ],<br>[edit protocols <a href="#">rip</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                              |
| <b>Description</b>              | Install RIP routes into multiple routing tables by configuring a routing table group.                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <i>group-name</i> —Name of the routing table group.                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Redistributing Routes Between Two RIP Instances on page 4742</a></li> </ul>                                                                                                                                                                                                               |

## rip

|                                 |                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>rip {...}</code>                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols],<br>[edit protocols],<br>[edit routing-instances <i>routing-instance-name</i> protocols]        |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| <b>Description</b>              | Enable RIP routing on the routing device.                                                                                                                                                                                                                                  |
| <b>Default</b>                  | RIP is disabled on the routing device.                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring RIP on page 4685</a></li> </ul>                                                                                                                                                                  |

## route-timeout (Protocols RIP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>route-timeout seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rip <a href="#">group</a> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rip <a href="#">group</a> <i>group-name</i> neighbor <i>neighbor-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip <a href="#">group</a> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip <a href="#">group</a> <i>group-name</i> neighbor <i>neighbor-name</i>],</p> <p>[edit protocols <a href="#">rip</a>],</p> <p>[edit protocols rip <a href="#">group</a> <i>group-name</i>],</p> <p>[edit protocols rip <a href="#">group</a> <i>group-name</i> neighbor <i>neighbor-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip <a href="#">group</a> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip <a href="#">group</a> <i>group-name</i> neighbor <i>neighbor-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Configure the route timeout interval for RIP. If a route is not refreshed after being installed in the routing table by the specified timeout interval, the route is marked as invalid and is removed from the routing table after the hold-down period expires.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <p><b>seconds</b>—Estimated time to wait before making updates to the routing table.</p> <p><b>Range:</b> 30 through 360 seconds</p> <p><b>Default:</b> 180 seconds</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring RIP Timers on page 4750</a></li> <li>• <a href="#">RIP Demand Circuits Overview</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

## send (Protocols RIP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>send <i>send-options</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit protocols <a href="#">rip</a>],</p> <p>[edit protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Configure RIP send options.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <p><i>send-options</i>—One of the following:</p> <ul style="list-style-type: none"> <li>• <b>broadcast</b>—Broadcast RIP version 2 packets (RIP version 1 compatible).</li> <li>• <b>multicast</b>—Multicast RIP version 2 packets. This is the default.</li> <li>• <b>none</b>—Do not send RIP updates.</li> <li>• <b>version-1</b>—Broadcast RIP version 1 packets.</li> </ul> <p><b>Default:</b> multicast</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets on page 4735</a></li> <li>• <a href="#">receive on page 4780</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

## traceoptions (Protocols RIP)

|                            |                                                                                                                                                                                                                                                                                                                                                                        |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;flag-modifier&gt; &lt;disable&gt;; }</pre>                                                                                                                                                      |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit protocols <a href="#">rip</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>]</p> |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                          |
| <b>Description</b>         | Set RIP protocol-level tracing options.                                                                                                                                                                                                                                                                                                                                |



**NOTE:** The **traceoptions** statement is not supported on QFabric systems.

**Default** The default RIP protocol-level trace options are inherited from the global **traceoptions** statement.

**Options** **disable**—(Optional) Disable the tracing operation. One use of this option is to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

**file *filename***—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. We recommend that you place RIP tracing output in the file `/var/log/rip-log`.

**files *number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file.0***, then ***trace-file.1***, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum number of files, you must also specify a maximum file size with the **size** option.

**Range:** 2 through 1000 files

**Default:** 10 files

**flag *flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

### RIP Tracing Options

- auth**—RIP authentication



- **error**—RIP error packets
- **expiration**—RIP route expiration processing
- **holddown**—RIP hold-down processing
- **nsr-synchronization**—Nonstop routing synchronization events
- **packets**—All RIP packets
- **request**—RIP information packets such as request, poll, and poll entry packets
- **trigger**—RIP triggered updates
- **update**—RIP update packets

#### Global Tracing Options

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

**Default:** If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing

***flag-modifier***—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information.
- **receive**—Trace the packets being received.
- **receive-detail**—Provide detailed trace information for packets being received.
- **send**—Trace the packets being transmitted.
- **send-detail**—Provide detailed trace information for packets being transmitted.

**no-world-readable**—(Optional) Prevent any user from reading the log file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB) or megabytes (MB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 128 KB

**world-readable**—(Optional) Allow any user to read the log file.

|                                 |                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.                                                                 |
|                                 | routing-control—To add this statement to the configuration.                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Tracing RIP Protocol Traffic on page 4758</a></li></ul> |

## update-interval (Protocols RIP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>update-interval seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a> group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a> group <i>group-name</i> neighbor <i>neighbor-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a> group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a> group <i>group-name</i> neighbor <i>neighbor-name</i>],</p> <p>[edit protocols <a href="#">rip</a>],</p> <p>[edit protocols <a href="#">rip</a> group <i>group-name</i>],</p> <p>[edit protocols <a href="#">rip</a> group <i>group-name</i> neighbor <i>neighbor-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a> group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a> group <i>group-name</i> neighbor <i>neighbor-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Configure the interval at which routes learned by RIP are sent to neighbors. This timer controls the interval between routing updates. This timer is set to 30 seconds, by default, with a small random amount of time added when the timer is reset. This added time prevents congestion that can happen if all routing devices update their neighbors simultaneously.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <p><b>seconds</b>—Estimated time to wait before making updates to the routing table.</p> <p><b>Range:</b> 10 through 60 seconds</p> <p><b>Default:</b> 30 seconds</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">Example: Configuring RIP Timers on page 4750</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |



## CHAPTER 188

# Operational Commands

- `clear rip general-statistics`
- `clear rip statistics`
- `show rip general-statistics`
- `show rip neighbor`
- `show rip statistics`

## clear rip general-statistics

---

|                                                 |                                                                                                                                                                                                         |
|-------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                           | <a href="#">Syntax on page 4790</a><br><a href="#">Syntax (EX Series Switch and QFX Series) on page 4790</a>                                                                                            |
| <b>Syntax</b>                                   | clear rip general-statistics<br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                    |
| <b>Syntax (EX Series Switch and QFX Series)</b> | clear rip general-statistics                                                                                                                                                                            |
| <b>Release Information</b>                      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                |
| <b>Description</b>                              | Clear Routing Information Protocol (RIP) general statistics.                                                                                                                                            |
| <b>Options</b>                                  | <b>none</b> —Clear RIP general statistics.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>                 | clear                                                                                                                                                                                                   |
| <b>Related Documentation</b>                    | <ul style="list-style-type: none"><li>• <a href="#">show rip general-statistics on page 4792</a></li></ul>                                                                                              |
| <b>List of Sample Output</b>                    | <a href="#">clear rip general-statistics on page 4790</a>                                                                                                                                               |
| <b>Output Fields</b>                            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                   |

## Sample Output

### clear rip general-statistics

```
user@host> clear rip general-statistics
```

## clear rip statistics

|                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                             | <a href="#">Syntax on page 4791</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 4791</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Syntax</b>                                     | clear rip statistics<br><instance (all   <i>instance-name</i> )><br><logical-system (all   <i>logical-system-name</i> )><br><neighbor><br><peer (all   <i>address</i> )>                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Syntax (EX Series Switches and QFX Series)</b> | clear rip statistics<br><instance (all   <i>instance-name</i> )><br><neighbor>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>                                | Clear RIP statistics.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                                    | <p><b>none</b>—Reset RIP counters for all neighbors for all routing instances.</p> <p><b>instance (all   <i>instance-name</i>)</b>—(Optional) Clear RIP statistics for all instances or for the specified routing instance only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>neighbor</b>—(Optional) Clear RIP statistics for the specified neighbor only.</p> <p><b>peer (all   <i>address</i>)</b>—(Optional) Clear RIP statistics for a single peer or all peers.</p> |
| <b>Required Privilege Level</b>                   | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>                      | <ul style="list-style-type: none"> <li>• <a href="#">show rip statistics on page 4796</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>List of Sample Output</b>                      | <a href="#">clear rip statistics on page 4791</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Output Fields</b>                              | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## Sample Output

### clear rip statistics

```
user@host> clear rip statistics
```

## show rip general-statistics

|                                                 |                                                                                                                                                                                          |
|-------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                           | <a href="#">Syntax on page 4792</a><br><a href="#">Syntax (EX Series Switch and QFX Series) on page 4792</a>                                                                             |
| <b>Syntax</b>                                   | show rip general-statistics<br><logical-system (all   <i>logical-system-name</i> )>                                                                                                      |
| <b>Syntax (EX Series Switch and QFX Series)</b> | show rip general-statistics                                                                                                                                                              |
| <b>Release Information</b>                      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series. |
| <b>Description</b>                              | Display brief Routing Information Protocol (RIP) statistics.                                                                                                                             |
| <b>Options</b>                                  | none—Display brief RIP statistics.<br><br>logical-system (all   <i>logical-system-name</i> )—(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>                 | view                                                                                                                                                                                     |
| <b>Related Documentation</b>                    | <ul style="list-style-type: none"> <li><a href="#">clear rip general-statistics on page 4790</a></li> </ul>                                                                              |
| <b>List of Sample Output</b>                    | <a href="#">show rip general-statistics on page 4792</a>                                                                                                                                 |
| <b>Output Fields</b>                            | Table 359 lists the output fields for the <b>show rip general-statistics</b> command. Output fields are listed in the approximate order in which they appear.                            |

**Table 359: show rip general-statistics Output Fields**

| Field Name  | Field Description                                      |
|-------------|--------------------------------------------------------|
| bad msgs    | Number of invalid messages received.                   |
| no rcv intf | Number of packets received with no matching interface. |
| curr memory | Amount of memory currently used by RIP.                |
| max memory  | Most memory used by RIP.                               |

## Sample Output

### show rip general-statistics

```
user@host> show rip general-statistics
```



```
RIPv2 I/O info:
 bad msgs : 0
 no recv intf : 0
 curr memory : 0
 max memory : 0
```

## show rip neighbor

---

|                                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| List of Syntax                             | <a href="#">Syntax on page 4794</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 4794</a>                                                                                                                                                                                                                                                                                                                                                                                                |
| Syntax                                     | <pre>show rip neighbor &lt;instance (all   <i>instance-name</i>)&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;name&gt;</pre>                                                                                                                                                                                                                                                                                                                                                             |
| Syntax (EX Series Switches and QFX Series) | <pre>show rip neighbor &lt;instance (all   <i>instance-name</i>)&gt; &lt;name&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Release Information                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                      |
| Description                                | Display information about RIP neighbors.                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Options                                    | <p><b>none</b>—Display information about all RIP neighbors for all instances.</p> <p><b>instance (all   <i>instance-name</i>)</b>—(Optional) Display RIP neighbor information for all instances or for only the specified routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>name</b>—(Optional) Display detailed information about only the specified RIP neighbor.</p> |
| Required Privilege Level                   | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| List of Sample Output                      | <a href="#">show rip neighbor on page 4795</a><br><a href="#">show rip neighbor (With Demand Circuits Configured) on page 4795</a>                                                                                                                                                                                                                                                                                                                                                                            |
| Output Fields                              | <a href="#">Table 360</a> lists the output fields for the <b>show rip neighbor</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                           |

Table 360: show rip neighbor Output Fields

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Neighbor</b>            | Name of the RIP neighbor.<br><br><b>NOTE:</b> Beginning with Junos OS Release 11.1, when you configure demand circuits, the output displays a demand circuit (DC) flag next to neighbor interfaces configured for demand circuits.<br><br>If you configure demand circuits at the <b>[edit protocols rip group group-name neighbor neighbor-name]</b> hierarchy level, the output shows only the neighboring interface that you specifically configured as a demand circuit. If you configure demand circuits at the <b>[edit protocols rip group group-name]</b> hierarchy level, all of the interfaces in the group are configured as demand circuits. Therefore, the output shows all of the interfaces in that group as demand circuits. |
| <b>State</b>               | State of the connection: <b>Up</b> or <b>Dn</b> (Down).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Source Address</b>      | Source address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Destination Address</b> | Destination address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Send Mode</b>           | Send options: <b>broadcast</b> , <b>multicast</b> , <b>none</b> , or <b>version 1</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Receive Mode</b>        | Type of packets to accept: <b>both</b> , <b>none</b> , <b>version 1</b> , or <b>version 2</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>In Met</b>              | Metric added to incoming routes when advertising into RIP routes that were learned from other protocols.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## Sample Output

### show rip neighbor

```

user@host> show rip neighbor
Neighbor Local Source Destination Send Receive In
----- -
ge-2/3/0.0 Up 192.168.9.105 192.168.9.107 bcast both 1
at-5/1/1.42 Dn (null) (null) mcast v2 only 3
at-5/1/0.42 Dn (null) (null) mcast both 3
at-5/1/0.0 Up 20.0.0.1 224.0.0.9 mcast both 3
so-0/0/0.0 Up 192.168.9.97 224.0.0.9 mcast both 3

```

### show rip neighbor (With Demand Circuits Configured)

```

user@host# show rip neighbor
Neighbor Local Source Destination Send Receive In
----- -
so-0/1/0.0(DC) Up 10.10.10.2 224.0.0.9 mcast both 1
so-0/2/0.0(DC) Up 13.13.13.2 224.0.0.9 mcast both 1

```

## show rip statistics

---

|                                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| List of Syntax                             | <a href="#">Syntax on page 4796</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 4796</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Syntax                                     | <pre>show rip statistics &lt;instance (all   <i>instance-name</i>)&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;<i>name</i>&gt; &lt;peer (all   <i>address</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                            |
| Syntax (EX Series Switches and QFX Series) | <pre>show rip statistics &lt;instance (all   <i>instance-name</i>)&gt; &lt;<i>name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Release Information                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Description                                | Display RIP statistics about messages sent and received on an interface, as well as information received from advertisements from other routing devices.                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Options                                    | <p><b>none</b>—Display RIP statistics for all routing instances.</p> <p><b>instance (all   <i>instance-name</i>)</b>—(Optional) Display RIP statistics for all instances or for only the specified routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b><i>name</i></b>—(Optional) Display detailed information about only the specified RIP neighbor.</p> <p><b>peer (all   <i>address</i>)</b>—(Optional) Display RIP statistics for a single peer or all peers.</p> |
| Required Privilege Level                   | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Related Documentation                      | <ul style="list-style-type: none"><li>• <a href="#">clear rip statistics on page 4791</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| List of Sample Output                      | <a href="#">show rip statistics on page 4797</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Output Fields                              | <a href="#">Table 361</a> lists the output fields for the <b>show rip statistics</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                     |

Table 361: show rip statistics Output Fields

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>RIP info</b>          | <p>Information about RIP on the specified interface:</p> <ul style="list-style-type: none"> <li>• <b>port</b>—UDP port number used for RIP.</li> <li>• <b>update interval</b>—Interval between routing table updates, in seconds.</li> <li>• <b>holddown</b>—Hold-down interval, in seconds.</li> <li>• <b>timeout</b>—Timeout interval, in seconds.</li> <li>• <b>restart in progress</b>—Graceful restart status. Displayed when RIP is or has been in the process of graceful restart.</li> <li>• <b>restart time</b>—Estimated time for the graceful restart to finish, in seconds.</li> <li>• <b>restart will complete in</b>—Remaining time for the graceful restart to finish, in seconds.</li> <li>• <b>rts learned</b>—Number of routes learned through RIP.</li> <li>• <b>rts held down</b>—Number of routes held down by RIP.</li> <li>• <b>rqsts dropped</b>—Number of received request packets that were dropped.</li> <li>• <b>resps dropped</b>—Number of received response packets that were dropped.</li> </ul>                                                                                                                                                             |
| <b>logical-interface</b> | <p>Name of the logical interface and its statistics:</p> <ul style="list-style-type: none"> <li>• <b>routes learned</b>—Number of routes learned on the logical interface.</li> <li>• <b>routes advertised</b>—Number of routes advertised by the logical interface.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Counter</b>           | <p>List of counter types:</p> <ul style="list-style-type: none"> <li>• <b>Updates Sent</b>—Number of update messages sent.</li> <li>• <b>Triggered Updates Sent</b>—Number of triggered update messages sent.</li> <li>• <b>Responses Sent</b>—Number of response messages sent.</li> <li>• <b>Bad Messages</b>—Number of invalid messages received.</li> <li>• <b>RIPv1 Updates Received</b>—Number of RIPv1 update messages received.</li> <li>• <b>RIPv1 Bad Route Entries</b>—Number of RIPv1 invalid route entry messages received.</li> <li>• <b>RIPv1 Updates Ignored</b>—Number of RIPv1 update messages ignored.</li> <li>• <b>RIPv2 Updates Received</b>—Number of RIPv2 update messages received.</li> <li>• <b>RIPv2 Bad Route Entries</b>—Number of RIPv2 invalid route entry messages received.</li> <li>• <b>RIPv2 Updates Ignored</b>—Number of RIPv2 update messages that were ignored.</li> <li>• <b>Authentication Failures</b>—Number of received update messages that failed authentication.</li> <li>• <b>RIP Requests Received</b>—Number of RIP request messages received.</li> <li>• <b>RIP Requests Ignored</b>—Number of RIP request messages ignored.</li> </ul> |
| <b>Total</b>             | Total number of packets for the selected counter.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Last 5 min</b>        | Number of packets for the selected counter in the most recent 5-minute period.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Last minute</b>       | Number of packets for the selected counter in the most recent 1-minute period.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## Sample Output

### show rip statistics

```
user@host> show rip statistics so-0/0/0.0
```

RIP info: port 520; update interval: 30s; holddown 180s; timeout 120s  
restart in progress: restart time 60s; restart will complete in 55s  
      rts learned  rts held down  rqsts dropped  resps dropped  
              0              0              0              0

so-0/0/0.0: 0 routes learned; 501 routes advertised

| Counter                 | Total | Last 5 min | Last minute |
|-------------------------|-------|------------|-------------|
| -----                   | ----- | -----      | -----       |
| Updates Sent            | 0     | 0          | 0           |
| Triggered Updates Sent  | 0     | 0          | 0           |
| Responses Sent          | 0     | 0          | 0           |
| Bad Messages            | 0     | 0          | 0           |
| RIPv1 Updates Received  | 0     | 0          | 0           |
| RIPv1 Bad Route Entries | 0     | 0          | 0           |
| RIPv1 Updates Ignored   | 0     | 0          | 0           |
| RIPv2 Updates Received  | 0     | 0          | 0           |
| RIPv2 Bad Route Entries | 0     | 0          | 0           |
| RIPv2 Updates Ignored   | 0     | 0          | 0           |
| Authentication Failures | 0     | 0          | 0           |
| RIP Requests Received   | 0     | 0          | 0           |
| RIP Requests Ignored    | 0     | 0          | 0           |

# MPLS Feature Guide for QFX10000 Switches





## PART 71

# LDP

- [Using LDP on page 4803](#)
- [Configuration Statements and Monitoring Commands for LDP on page 4841](#)



## Using LDP

- [LDP Introduction on page 4804](#)
- [Junos OS LDP Protocol Implementation on page 4804](#)
- [LDP Operation on page 4804](#)
- [Tunneling LDP LSPs in RSVP LSPs on page 4805](#)
- [Tunneling LDP LSPs in RSVP LSPs Overview on page 4805](#)
- [Label Operations on page 4806](#)
- [LDP Message Types on page 4807](#)
- [Discovery Messages on page 4807](#)
- [Session Messages on page 4808](#)
- [Advertisement Messages on page 4808](#)
- [Notification Messages on page 4808](#)
- [LDP Session Protection on page 4808](#)
- [LDP Graceful Restart on page 4809](#)
- [Minimum LDP Configuration on page 4810](#)
- [Enabling and Disabling LDP on page 4810](#)
- [Enabling Strict Targeted Hello Messages for LDP on page 4810](#)
- [Filtering Inbound LDP Label Bindings on page 4811](#)
- [Filtering Outbound LDP Label Bindings on page 4813](#)
- [Specifying the Transport Address Used by LDP on page 4815](#)
- [Collecting LDP Statistics on page 4815](#)
- [Tracing LDP Protocol Traffic on page 4817](#)
- [Example: Configuring LDP Downstream on Demand on page 4820](#)
- [Configuring the LDP Timer for Hello Messages on page 4825](#)
- [Configuring the Delay Before LDP Neighbors Are Considered Down on page 4826](#)
- [Configuring the Interval for LDP Keepalive Messages on page 4827](#)
- [Configuring the LDP Keepalive Timeout on page 4828](#)
- [Configuring LDP Route Preferences on page 4828](#)
- [Configuring LDP Graceful Restart on page 4828](#)

- [Configuring the Prefixes Advertised into LDP from the Routing Table on page 4831](#)
- [Configuring LDP LSP Traceroute on page 4832](#)
- [Configuring Miscellaneous LDP Properties on page 4833](#)

## LDP Introduction

---

The Label Distribution Protocol (LDP) is a protocol for distributing labels in non-traffic-engineered applications. LDP allows routers to establish label-switched paths (LSPs) through a network by mapping network-layer routing information directly to data link layer-switched paths.

These LSPs might have an endpoint at a directly attached neighbor (comparable to IP hop-by-hop forwarding), or at a network egress node, enabling switching through all intermediary nodes. LSPs established by LDP can also traverse traffic-engineered LSPs created by RSVP.

LDP associates a forwarding equivalence class (FEC) with each LSP it creates. The FEC associated with an LSP specifies which packets are mapped to that LSP. LSPs are extended through a network as each router chooses the label advertised by the next hop for the FEC and splices it to the label it advertises to all other routers. This process forms a tree of LSPs that converge on the egress router.

## Junos OS LDP Protocol Implementation

---

The Junos OS implementation of LDP supports LDP version 1. The Junos OS supports a simple mechanism for tunneling between routers in an interior gateway protocol (IGP), to eliminate the required distribution of external routes within the core. The Junos OS allows an MPLS tunnel next hop to all egress routers in the network, with only an IGP running in the core to distribute routes to egress routers. Edge routers run BGP but do not distribute external routes to the core. Instead, the recursive route lookup at the edge resolves to an LSP switched to the egress router. No external routes are necessary on the transit LDP routers.

## LDP Operation

---

You must configure LDP for each interface on which you want LDP to run. LDP creates LSP trees rooted at each egress router for the router ID address that is the subsequent BGP next hop. The ingress point is at every router running LDP. This process provides an inet.3 route to every egress router. If BGP is running, it will attempt to resolve next hops by using the inet.3 table first, which binds most, if not all, of the BGP routes to MPLS tunnel next hops.

Two adjacent routers running LDP become neighbors. If the two routers are connected by more than one interface, they become neighbors on each interface. When LDP routers become neighbors, they establish an LDP session to exchange label information. If per-router labels are in use on both routers, only one LDP session is established between them, even if they are neighbors on multiple interfaces. For this reason, an LDP session is not related to a particular interface.

LDP operates in conjunction with a unicast routing protocol. LDP installs LSPs only when both LDP and the routing protocol are enabled. For this reason, you must enable both LDP and the routing protocol on the same set of interfaces. If this is not done, LSPs might not be established between each egress router and all ingress routers, which might result in loss of BGP-routed traffic.

You can apply policy filters to labels received from and distributed to other routers through LDP. Policy filters provide you with a mechanism to control the establishment of LSPs.

For LDP to run on an interface, MPLS must be enabled on a logical interface on that interface. For more information, see the *Logical Interfaces*.

**Related Documentation**

- [Logical Interfaces](#)

## Tunneling LDP LSPs in RSVP LSPs

You can tunnel LDP LSPs over RSVP LSPs. The following sections describe how tunneling of LDP LSPs in RSVP LSPs works:

- [Tunneling LDP LSPs in RSVP LSPs Overview on page 4805](#)
- [Label Operations on page 4806](#)

## Tunneling LDP LSPs in RSVP LSPs Overview

If you are using RSVP for traffic engineering, you can run LDP simultaneously to eliminate the distribution of external routes in the core. The LSPs established by LDP are tunneled through the LSPs established by RSVP. LDP effectively treats the traffic-engineered LSPs as single hops.

When you configure the router to run LDP across RSVP-established LSPs, LDP automatically establishes sessions with the router at the other end of the LSP. LDP control packets are routed hop-by-hop, rather than carried through the LSP. This routing allows you to use simplex (one-way) traffic-engineered LSPs. Traffic in the opposite direction flows through LDP-established LSPs that follow unicast routing rather than through traffic-engineered tunnels.

If you configure LDP over RSVP LSPs, you can still configure multiple OSPF areas and IS-IS levels in the traffic engineered core and in the surrounding LDP cloud.



**NOTE:** Beginning with Junos OS Release 15.1, multi-instance support is extended to LDP over RSVP tunneling for a virtual router routing instance. This allows splitting of a single routing and MPLS domain into multiple domains so that each domain can be scaled independently. BGP labeled unicast can be used to stitch these domains for service FECs. Each domain uses intra-domain LDP over RSVP LSP for MPLS forwarding.

- Related Documentation**
- [Label Operations on page 4806](#)
  - *Configuring a Hierarchy of RSVP LSPs to Tunnel Multiple RSVP LSPs Over a Single RSVP LSP*

## Label Operations

Figure 137 depicts an LDP LSP being tunneled through an RSVP LSP. (For definitions of label operations, see *MPLS Label Overview*.) The shaded inner oval represents the RSVP domain, whereas the outer oval depicts the LDP domain. RSVP establishes an LSP through routers B, C, D, and E, with the sequence of labels L3, L4. LDP establishes an LSP through Routers A, B, E, F, and G, with the sequence of labels L1, L2, L5. LDP views the RSVP LSP between Routers B and E as a single hop.

When the packet arrives at Router A, it enters the LSP established by LDP, and a label (L1) is pushed onto the packet. When the packet arrives at Router B, the label (L1) is swapped with another label (L2). Because the packet is entering the traffic-engineered LSP established by RSVP, a second label (L3) is pushed onto the packet.

This outer label (L3) is swapped with a new label (L4) at the intermediate router (C) within the RSVP LSP tunnel, and when the penultimate router (D) is reached, the top label is popped. Router E swaps the label (L2) with a new label (L5), and the penultimate router for the LDP-established LSP (F) pops the last label.

Figure 137: Swap and Push When LDP LSPs Are Tunneled Through RSVP LSPs





Figure 138 depicts a double push label operation (L1L2). A double push label operation is used when the ingress router (A) for both the LDP LSP and the RSVP LSP tunneled through it is the same device. Note that Router D is the penultimate hop for the LDP-established LSP, so L2 is popped from the packet by Router D.

Figure 138: Double Push When LDP LSPs Are Tunneled Through RSVP LSPs



## LDP Message Types

LDP uses the message types described in the following sections to establish and remove mappings and to report errors. All LDP messages have a common structure that uses a type, length, and value (TLV) encoding scheme.

- [Discovery Messages on page 4807](#)
- [Session Messages on page 4808](#)
- [Advertisement Messages on page 4808](#)
- [Notification Messages on page 4808](#)

## Discovery Messages

Discovery messages announce and maintain the presence of a router in a network. Routers indicate their presence in a network by sending hello messages periodically. Hello messages are transmitted as UDP packets to the LDP port at the group multicast address for all routers on the subnet.

LDP uses the following discovery procedures:

- **Basic discovery**—A router periodically sends LDP link hello messages through an interface. LDP link hello messages are sent as UDP packets addressed to the LDP discovery port. Receipt of an LDP link hello message on an interface identifies an adjacency with the LDP peer router.
- **Extended discovery**—LDP sessions between routers not directly connected are supported by LDP extended discovery. A router periodically sends LDP targeted hello messages to a specific address. Targeted hello messages are sent as UDP packets addressed to the LDP discovery port at the specific address. The targeted router decides whether to respond to or ignore the targeted hello message. A targeted router that chooses to respond does so by periodically sending targeted hello messages to the initiating router.

## Session Messages

---

Session messages establish, maintain, and terminate sessions between LDP peers. When a router establishes a session with another router learned through the hello message, it uses the LDP initialization procedure over TCP transport. When the initialization procedure is completed successfully, the two routers are LDP peers and can exchange advertisement messages.

## Advertisement Messages

---

Advertisement messages create, change, and delete label mappings for forwarding equivalence classes (FECs). Requesting a label or advertising a label mapping to a peer is a decision made by the local router. In general, the router requests a label mapping from a neighboring router when it needs one and advertises a label mapping to a neighboring router when it wants the neighbor to use a label.

## Notification Messages

---

Notification messages provide advisory information and signal error information. LDP sends notification messages to report errors and other events of interest. There are two kinds of LDP notification messages:

- Error notifications, which signal fatal errors. If a router receives an error notification from a peer for an LDP session, it terminates the LDP session by closing the TCP transport connection for the session and discarding all label mappings learned through the session.
- Advisory notifications, which pass information to a router about the LDP session or the status of some previous message received from the peer.

## LDP Session Protection

---

LDP session protection is based on the LDP targeted hello functionality defined in RFC 5036, *LDP Specification*, and is supported by the Junos OS as well as the LDP implementations of most other vendors. It involves sending unicast User Datagram Protocol (UDP) hello packets to a remote neighbor address and receiving similar packets from the neighbor router.

If you configure LDP session protection on a router, the LDP sessions are maintained as follows:

1. An LDP session is established between a router and a remote neighboring router.
2. If all of the direct links between the routers go down, the LDP session remains up so long as there is IP connectivity between the routers based on another connection over the network.
3. When the direct link between the routers is reestablished, the LDP session is not restarted. The routers simply exchange LDP hellos with each other over the direct link.



They can then begin forwarding LDP-signaled MPLS packets using the original LDP session.

By default, LDP targeted hellos are set to the remote neighbor so long as the LDP session is up, even if there are no more link neighbors to that router. You can also specify the duration you would like to maintain the remote neighbor connection in the absence of link neighbors. When the last link neighbor for a session goes down, the Junos OS starts an LDP session protection timer. If this timer expires before any of the link neighbors come back up, the remote neighbor connection is taken down and the LDP session is terminated. If you configure a different value for the timer while it is currently running, the Junos OS updates the timer to the specified value without disrupting the current state of the LDP session.

## LDP Graceful Restart

---

LDP graceful restart enables a router whose LDP control plane is undergoing a restart to continue to forward traffic while recovering its state from neighboring routers. It also enables a router on which helper mode is enabled to assist a neighboring router that is attempting to restart LDP.

During session initialization, a router advertises its ability to perform LDP graceful restart or to take advantage of a neighbor performing LDP graceful restart by sending the graceful restart TLV. This TLV contains two fields relevant to LDP graceful restart: the reconnect time and the recovery time. The values of the reconnect and recovery times indicate the graceful restart capabilities supported by the router.

When a router discovers that a neighboring router is restarting, it waits until the end of the recovery time before attempting to reconnect. The recovery time is the length of time a router waits for LDP to restart gracefully. The recovery time period begins when an initialization message is sent or received. This time period is also typically the length of time that a neighboring router maintains its information about the restarting router, allowing it to continue to forward traffic.

You can configure LDP graceful restart in both the master instance for the LDP protocol and for a specific routing instance. You can disable graceful restart at the global level for all protocols, at the protocol level for LDP only, and on a specific routing instance. LDP graceful restart is disabled by default, because at the global level, graceful restart is disabled by default. However, helper mode (the ability to assist a neighboring router attempting a graceful restart) is enabled by default.

The following are some of the behaviors associated with LDP graceful restart:

- Outgoing labels are not maintained in restarts. New outgoing labels are allocated.
- When a router is restarting, no label-map messages are sent to neighbors that support graceful restart until the restarting router has stabilized (label-map messages are immediately sent to neighbors that do not support graceful restart). However, all other messages (keepalive, address-message, notification, and release) are sent as usual. Distributing these other messages prevents the router from distributing incomplete information.

- Helper mode and graceful restart are independent. You can disable graceful restart in the configuration, but still allow the router to cooperate with a neighbor attempting to restart gracefully.

## Minimum LDP Configuration

---

To enable LDP on a single interface, include the **ldp** statement and specify the interface using the **interface** statement. This is the minimum LDP configuration. All other LDP configuration statements are optional.

```
ldp {
 interface interface-name;
}
```

To enable LDP on all interfaces, specify **all** for *interface-name*.

For a list of hierarchy levels at which you can include these statements, see the statement summary sections.

## Enabling and Disabling LDP

---

LDP is routing-instance-aware. To enable LDP on a specific interface, include the following statements:

```
ldp {
 interface interface-name;
}
```

For a list of hierarchy levels at which you can include these statements, see the statement summary sections.

To enable LDP on all interfaces, specify **all** for *interface-name*.

If you have configured interface properties on a group of interfaces and want to disable LDP on one of the interfaces, include the **interface** statement with the **disable** option:

```
interface interface-name {
 disable;
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section.

## Enabling Strict Targeted Hello Messages for LDP

---

Use strict targeted hello messages to prevent LDP sessions from being established with remote neighbors that have not been specifically configured. If you configure the **strict-targeted-hellos** statement, an LDP peer does not respond to targeted hello messages coming from a source that is not one of its configured remote neighbors. Configured remote neighbors can include:

- Endpoints of RSVP tunnels for which LDP tunneling is configured
- Layer 2 circuit neighbors

If an unconfigured neighbor sends a hello message, the LDP peer ignores the message and logs an error (with the **error** trace flag) indicating the source. For example, if the LDP peer received a targeted hello from the Internet address 10.0.0.1 and no neighbor with this address is specifically configured, the following message is printed to the LDP log file:

```
LDP: Ignoring targeted hello from 10.0.0.1
```

To enable strict targeted hello messages, include the **strict-targeted-hellos** statement:

```
strict-targeted-hellos;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

## Filtering Inbound LDP Label Bindings

You can filter received LDP label bindings, applying policies to accept or deny bindings advertised by neighboring routers. To configure received-label filtering, include the **import** statement:

```
import [policy-names];
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

The named policy (configured at the **[edit policy-options]** hierarchy level) is applied to all label bindings received from all LDP neighbors. All filtering is done with **from** statements. [Table 362](#) lists the only **from** operators that apply to LDP received-label filtering.

**Table 362: from Operators That Apply to LDP Received-Label Filtering**

| from Operator       | Description                                                                                |
|---------------------|--------------------------------------------------------------------------------------------|
| <b>interface</b>    | Matches on bindings received from a neighbor that is adjacent over the specified interface |
| <b>neighbor</b>     | Matches on bindings received from the specified LDP router ID                              |
| <b>next-hop</b>     | Matches on bindings received from a neighbor advertising the specified interface address   |
| <b>route-filter</b> | Matches on bindings with the specified prefix                                              |

If a binding is filtered, it still appears in the LDP database, but is not considered for installation as part of a label-switched path (LSP).

Generally, applying policies in LDP can be used only to block the establishment of LSPs, not to control their routing. This is because the path that an LSP follows is determined by unicast routing, and not by LDP. However, when there are multiple equal-cost paths to the destination through different neighbors, you can use LDP filtering to exclude some

of the possible next hops from consideration. (Otherwise, LDP chooses one of the possible next hops at random.)

LDP sessions are not bound to interfaces or interface addresses. LDP advertises only per-router (not per-interface) labels; so if multiple parallel links exist between two routers, only one LDP session is established, and it is not bound to a single interface. When a router has multiple adjacencies to the same neighbor, take care to ensure that the filter does what is expected. (Generally, using **next-hop** and **interface** is not appropriate in this case.)

If a label has been filtered (meaning that it has been rejected by the policy and is not used to construct an LSP), it is marked as filtered in the database:

```
user@host> show ldp database
Input label database, 10.10.255.1:0-10.10.255.6:0
Label Prefix
3 10.10.255.6/32 (Filtered)
Output label database, 10.10.255.1:0-10.10.255.6:0
Label Prefix
3 10.10.255.1/32 (Filtered)
```

For more information about how to configure policies for LDP, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

## Examples: Filtering Inbound LDP Label Bindings

Accept only /32 prefixes from all neighbors:

```
[edit]
protocols {
 ldp {
 import only-32;
 ...
 }
}
policy-options {
 policy-statement only-32 {
 term first {
 from {
 route-filter 0.0.0.0/0 upto /31;
 }
 then reject;
 }
 then accept;
 }
}
```

Accept 131.108/16 or longer from router ID 10.10.255.2 and accept all prefixes from all other neighbors:

```
[edit]
protocols {
 ldp {
 import nosy-neighbor;
 ...
 }
}
```

```

}
policy-options {
 policy-statement nosy-neighbor {
 term first {
 from {
 neighbor 10.10.255.2;
 route-filter 131.108.0.0/16 orlonger accept;
 route-filter 0.0.0.0/0 orlonger reject;
 }
 }
 then accept;
 }
}

```

## Filtering Outbound LDP Label Bindings

You can configure export policies to filter LDP outbound labels. You can filter outbound label bindings by applying routing policies to block bindings from being advertised to neighboring routers. To configure outbound label filtering, include the **export** statement:

**export** [*policy-name*];

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

The named export policy (configured at the **[edit policy-options]** hierarchy level) is applied to all label bindings transmitted to all LDP neighbors. The only **from** operator that applies to LDP outbound label filtering is **route-filter**, which matches bindings with the specified prefix. The only **to** operators that apply to outbound label filtering are the operators in [Table 363](#).

**Table 363: to Operators for LDP Outbound-Label Filtering**

| to Operator      | Description                                                                          |
|------------------|--------------------------------------------------------------------------------------|
| <b>interface</b> | Matches on bindings sent to a neighbor that is adjacent over the specified interface |
| <b>neighbor</b>  | Matches on bindings sent to the specified LDP router ID                              |
| <b>next-hop</b>  | Matches on bindings sent to a neighbor advertising the specified interface address   |

If a binding is filtered, the binding is not advertised to the neighboring router, but it can be installed as part of an LSP on the local router. You can apply policies in LDP to block the establishment of LSPs, but not to control their routing. The path an LSP follows is determined by unicast routing, not by LDP.

LDP sessions are not bound to interfaces or interface addresses. LDP advertises only per-router (not per-interface) labels. If multiple parallel links exist between two routers, only one LDP session is established, and it is not bound to a single interface.

Do not use the **next-hop** and **interface** operators when a router has multiple adjacencies to the same neighbor.

Filtered labels are marked in the database:

```
user@host> show ldp database
Input label database, 10.10.255.1:0-10.10.255.3:0
Label Prefix
100007 10.10.255.2/32
3 10.10.255.3/32
Output label database, 10.10.255.1:0-10.10.255.3:0
Label Prefix
3 10.10.255.1/32
100001 10.10.255.6/32 (Filtered)
```

For more information about how to configure policies for LDP, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

## Examples: Filtering Outbound LDP Label Bindings

Block transmission of the route for **10.10.255.6/32** to any neighbors:

```
[edit protocols]
ldp {
 export block-one;
}
policy-options {
 policy-statement block-one {
 term first {
 from {
 route-filter 10.10.255.6/32 exact;
 }
 then reject;
 }
 then accept;
 }
}
```

Send only **131.108/16** or longer to router ID **10.10.255.2**, and send all prefixes to all other routers:

```
[edit protocols]
ldp {
 export limit-lsps;
}
policy-options {
 policy-statement limit-lsps {
 term allow-one {
 from {
 route-filter 131.108.0.0/16 orlonger;
 }
 to {
 neighbor 10.10.255.2;
 }
 then accept;
 }
 term block-the-rest {
```

```

 to {
 neighbor 10.10.255.2;
 }
 then reject;
 }
 then accept;
}
}

```

## Specifying the Transport Address Used by LDP

Routers must first establish a TCP session between each other before they can establish an LDP session. The TCP session enables the routers to exchange the label advertisements needed for the LDP session. To establish the TCP session, each router must learn the other router's transport address. The transport address is an IP address used to identify the TCP session over which the LDP session will run.

To configure the LDP transport address, include the `transport-address` statement:

```
transport-address (router-id | interface);
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

If you specify the **router-id** option, the address of the router identifier is used as the transport address (unless otherwise configured, the router identifier is typically the same as the loopback address). If you specify the **interface** option, the interface address is used as the transport address for any LDP sessions to neighbors that can be reached over that interface. Note that the router identifier is used as the transport address by default.

You cannot specify the **interface** option when there are multiple parallel links to the same LDP neighbor, because the LDP specification requires that the same transport address be advertised on all interfaces to the same neighbor. If LDP detects multiple parallel links to the same neighbor, it disables interfaces to that neighbor one by one until the condition is cleared, either by disconnecting the neighbor on an interface or by specifying the **router-id** option.

**Related Documentation**

- [transport-address on page 4882](#)

## Collecting LDP Statistics

LDP traffic statistics show the volume of traffic that has passed through a particular FEC on a router.

When you configure the **traffic-statistics** statement at the **[edit protocols ldp]** hierarchy level, the LDP traffic statistics are gathered periodically and written to a file. You can configure how often statistics are collected (in seconds) by using the **interval** option. The default collection interval is 5 minutes. You must configure an LDP statistics file; otherwise, LDP traffic statistics are not gathered. If the LSP goes down, the LDP statistics are reset.

To collect LDP traffic statistics, include the **traffic-statistics** statement:

```
traffic-statistics {
```

```

file filename <files number> <size size> <world-readable | no-world-readable>;
interval interval;
no-penultimate-hop;
}

```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

This section includes the following topics:

- [LDP Statistics Output on page 4816](#)
- [Disabling LDP Statistics on the Penultimate-Hop Router on page 4816](#)
- [LDP Statistics Limitations on page 4817](#)

## LDP Statistics Output

The following sample output is from an LDP statistics file:

| FEC               | Type    | Packets | Bytes | Shared |
|-------------------|---------|---------|-------|--------|
| 10.255.350.448/32 | Transit | 0       | 0     | No     |
|                   | Ingress | 0       | 0     | No     |
| 10.255.350.450/32 | Transit | 0       | 0     | Yes    |
|                   | Ingress | 0       | 0     | No     |
| 10.255.350.451/32 | Transit | 0       | 0     | No     |
|                   | Ingress | 0       | 0     | No     |
| 220.220.220.1/32  | Transit | 0       | 0     | Yes    |
|                   | Ingress | 0       | 0     | No     |
| 220.220.220.2/32  | Transit | 0       | 0     | Yes    |
|                   | Ingress | 0       | 0     | No     |
| 220.220.220.3/32  | Transit | 0       | 0     | Yes    |
|                   | Ingress | 0       | 0     | No     |

May 28 15:02:05, read 12 statistics in 00:00:00 seconds

The LDP statistics file includes the following columns of data:

- **read**—Number of bytes of data passed by the FEC since its LSP came up.
- **read**—FEC for which LDP traffic statistics are collected.
- **read**—Number of packets passed by the FEC since its LSP came up.
- **read**—This number (which appears next to the date and time) might differ from the actual number of the statistics displayed. Some of the statistics are summarized before being displayed.
- **Shared**—A **Yes** value indicates that several prefixes are bound to the same label (for example, when several prefixes are advertised with an egress policy). The LDP traffic statistics for this case apply to all the prefixes and should be treated as such.
- **Type**—Type of traffic originating from a router, either **Ingress** (originating from this router) or **Transit** (forwarded through this router).

## Disabling LDP Statistics on the Penultimate-Hop Router

Gathering LDP traffic statistics at the penultimate-hop router can consume excessive system resources, on next-hop routes in particular. This problem is exacerbated if you have configured the **deaggregate** statement in addition to the **traffic-statistics** statement.



For routers reaching their limit of next-hop route usage, we recommend configuring the **no-penultimate-hop** option for the **traffic-statistics** statement:

```
traffic-statistics {
 no-penultimate-hop;
}
```

For a list of hierarchy levels at which you can configure the **traffic-statistics** statement, see the statement summary section for this statement.



**NOTE:** When you configure the **no-penultimate-hop** option, no statistics are available for the FECs that are the penultimate hop for this router.

Whenever you include or remove this option from the configuration, the LDP sessions are taken down and then restarted.

The following sample output is from an LDP statistics file showing routers on which the **no-penultimate-hop** option is configured:

| FEC               | Type    | Packets             | Bytes | Shared |
|-------------------|---------|---------------------|-------|--------|
| 10.255.245.218/32 | Transit | 0                   | 0     | No     |
|                   | Ingress | 4                   | 246   | No     |
| 10.255.245.221/32 | Transit | statistics disabled |       |        |
|                   | Ingress | statistics disabled |       |        |
| 13.1.1.0/24       | Transit | statistics disabled |       |        |
|                   | Ingress | statistics disabled |       |        |
| 13.1.3.0/24       | Transit | statistics disabled |       |        |
|                   | Ingress | statistics disabled |       |        |

## LDP Statistics Limitations

The following are issues related to collecting LDP statistics by configuring the **traffic-statistics** statement:

- You cannot clear the LDP statistics.
- If you shorten the specified interval, a new LDP statistics request is issued only if the statistics timer expires later than the new interval.
- A new LDP statistics collection operation cannot start until the previous one has finished. If the interval is short or if the number of LDP statistics is large, the time gap between the two statistics collections might be longer than the interval.

When an LSP goes down, the LDP statistics are reset.

## Tracing LDP Protocol Traffic

The following sections describe how to configure the trace options to examine LDP protocol traffic:

- [Tracing LDP Protocol Traffic at the Protocol and Routing Instance Levels on page 4818](#)
- [Tracing LDP Protocol Traffic Within FECs on page 4819](#)
- [Examples: Tracing LDP Protocol Traffic on page 4819](#)

## Tracing LDP Protocol Traffic at the Protocol and Routing Instance Levels

To trace LDP protocol traffic, you can specify options in the global **traceoptions** statement at the **[edit routing-options]** hierarchy level, and you can specify LDP-specific options by including the **traceoptions** statement:

```
traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Use the **file** statement to specify the name of the file that receives the output of the tracing operation. All files are placed in the directory `/var/log`. We recommend that you place LDP-tracing output in the file **ldp-log**.

The following trace flags display the operations associated with the sending and receiving of various LDP messages. Each can carry one or more of the following modifiers:

- **address**—Trace the operation of address and address withdrawal messages.
- **binding**—Trace label-binding operations.
- **error**—Trace error conditions.
- **event**—Trace protocol events.
- **initialization**—Trace the operation of initialization messages.
- **label**—Trace the operation of label request, label map, label withdrawal, and label release messages.
- **notification**—Trace the operation of notification messages.
- **packets**—Trace the operation of address, address withdrawal, initialization, label request, label map, label withdrawal, label release, notification, and periodic messages. This modifier is equivalent to setting the **address**, **initialization**, **label**, **notification**, and **periodic** modifiers.

You can also configure the **filter** flag modifier with the **match-on address** sub-option for the **packets** flag. This allows you to trace based on the source and destination addresses of the packets.

- **path**—Trace label-switched path operations.
- **path**—Trace label-switched path operations.
- **periodic**—Trace the operation of hello and keepalive messages.
- **route**—Trace the operation of route messages.
- **state**—Trace protocol state transitions.

## Tracing LDP Protocol Traffic Within FECs

LDP associates a forwarding equivalence class (FEC) with each LSP it creates. The FEC associated with an LSP specifies which packets are mapped to that LSP. LSPs are extended through a network as each router chooses the label advertised by the next hop for the FEC and splices it to the label it advertises to all other routers.

You can trace LDP protocol traffic within a specific FEC and filter LDP trace statements based on an FEC. This is useful when you want to trace or troubleshoot LDP protocol traffic associated with an FEC. The following trace flags are available for this purpose: **route**, **path**, and **binding**.

The following example illustrates how you might configure the LDP **traceoptions** statement to filter LDP trace statements based on an FEC:

```
[edit protocols ldp traceoptions]
set flag route filter match-on fec policy "filter-policy-for-ldp-fec";
```

This feature has the following limitations:

- The filtering capability is only available for FECs composed of IP version 4 (IPv4) prefixes.
- Layer 2 circuit FECs cannot be filtered.
- When you configure both route tracing and filtering, MPLS routes are not displayed (they are blocked by the filter).
- Filtering is determined by the policy and the configured value for the **match-on** option. When configuring the policy, be sure that the default behavior is always **reject**.
- The only **match-on** option is **fec**. Consequently, the only type of policy you should include is a route-filter policy.

## Examples: Tracing LDP Protocol Traffic

Trace LDP path messages in detail:

```
[edit]
protocols {
 ldp {
 traceoptions {
 file ldp size 10m files 5;
 flag path;
 }
 }
}
```

Trace all LDP outgoing messages:

```
[edit]
protocols {
 ldp {
 traceoptions {
 file ldp size 10m files 5;
```

```
 flag packets;
 }
}
```

Trace all LDP error conditions:

```
[edit]
protocols {
 ldp {
 traceoptions {
 file ldp size 10m files 5;
 flag error;
 }
 }
}
```

Trace all LDP incoming messages and all label-binding operations:

```
[edit]
protocols {
 ldp {
 traceoptions {
 file ldp size 10m files 5 world-readable;
 flag packets receive;
 flag binding;
 }
 interface all {
 }
 }
}
```

Trace LDP protocol traffic for an FEC associated with the LSP:

```
[edit]
protocols {
 ldp {
 traceoptions {
 flag route filter match-on fec policy filter-policy-for-ldp-fec;
 }
 }
}
```

---

## Example: Configuring LDP Downstream on Demand

This example shows how to configure LDP downstream on demand. LDP is commonly configured using downstream unsolicited advertisement mode, meaning label advertisements for all routes are received from all LDP peers. As service providers integrate the access and aggregation networks into a single MPLS domain, LDP downstream on demand is needed to distribute the bindings between the access and aggregation networks and to reduce the processing requirements for the control plane.

Downstream nodes could potentially receive tens of thousands of label bindings from upstream aggregation nodes. Instead of learning and storing all label bindings for all possible loopback addresses within the entire MPLS network, the downstream aggregation node can be configured using LDP downstream on demand to only request the label

bindings for the FECs corresponding to the loopback addresses of those egress nodes on which it has services configured.

- [Requirements on page 4821](#)
- [Overview on page 4821](#)
- [Configuration on page 4821](#)
- [Verification on page 4824](#)

## Requirements

This example uses the following hardware and software components:

- M Series router
- Junos OS 12.2

## Overview

You can enable LDP downstream on demand label advertisement for an LDP session by including the **downstream-on-demand** statement at the **[edit protocols ldp session]** hierarchy level. If you have configured downstream on demand, the Juniper Networks router advertises the downstream on demand request to its peer routers. For a downstream on demand session to be established between two routers, both have to advertise downstream on demand mode during LDP session establishment. If one router advertises downstream unsolicited mode and the other advertises downstream on demand, downstream unsolicited mode is used.

## Configuration

### Configuring LDP Downstream on Demand

#### Step-by-Step Procedure

To configure a LDP downstream on demand policy and then configure that policy and enable LDP downstream on demand on the LDP session:

1. Configure the downstream on demand policy (DOD-Request-Loopbacks in this example).

This policy causes the router to forward label request messages only to the FECs that are matched by the DOD-Request-Loopbacks policy.

```
[edit policy-options]
user@host# set prefix-list Request-Loopbacks 10.1.1.1/32
user@host# set prefix-list Request-Loopbacks 10.1.1.2/32
user@host# set prefix-list Request-Loopbacks 10.1.1.3/32
user@host# set prefix-list Request-Loopbacks 10.1.1.4/32
user@host# set policy-statement DOD-Request-Loopbacks term 1 from prefix-list
Request-Loopbacks
user@host# set policy-statement DOD-Request-Loopbacks term 1 then accept
```

2. Specify the DOD-Request-Loopbacks policy using the **dod-request-policy** statement at the **[edit protocols ldp]** hierarchy level.

The policy specified with the **dod-request-policy** statement is used to identify the prefixes to send label request messages. This policy is similar to an egress policy

or an import policy. When processing routes from the inet.0 routing table, the Junos OS software checks for routes matching the **DOD-Request-Loopbacks** policy (in this example). If the route matches the policy and the LDP session is negotiated with DOD advertisement mode, label request messages are sent to the corresponding downstream LDP session.

```
[edit protocols ldp]
user@host# set dod-request-policy DOD-Request-Loopbacks
```

3. Include the **downstream-on-demand** statement in the configuration for the LDP session to enable downstream on demand distribution mode.

```
[edit protocols ldp]
user@host# set session 1.1.1.1 downstream-on-demand
```

---

### Distributing LDP Downstream on Demand Routes into Labeled BGP

---

#### Step-by-Step Procedure

To distribute LDP downstream on demand routes into labeled BGP, use a BGP export policy.

1. Configure the LDP route policy (**redistribute\_ldp** in this example).

```
[edit policy-options]
user@host# set policy-statement redistribute_ldp term 1 from protocol ldp
user@host# set policy-statement redistribute_ldp term 1 from tag 1000
user@host# set policy-statement redistribute_ldp term 1 then accept
```

2. Include the LDP route policy, **redistribute\_ldp** in the BGP configuration (as a part of the BGP group configuration **ebgp-to-abr** in this example).

BGP forwards the LDP routes based on the **redistribute\_ldp** policy to the remote PE router

```
[edit protocols bgp]
user@host# set group ebgp-to-abr type external
user@host# set group ebgp-to-abr local-address 192.168.0.1
user@host# set group ebgp-to-abr peer-as 65319
user@host# set group ebgp-to-abr local-as 65320
user@host# set group ebgp-to-abr neighbor 192.168.6.1 family inet unicast
user@host# set group ebgp-to-abr neighbor 192.168.6.1 family inet labeled-unicast
rib inet.3
user@host# set group ebgp-to-abr neighbor 192.168.6.1 export redistribute_ldp
```

#### Step-by-Step Procedure

To restrict label propagation to other routers configured in downstream unsolicited mode (instead of downstream on demand), configure the following policies:

1. Configure the **dod-routes** policy to accept routes from LDP.

```
user@host# set policy-options policy-statement dod-routes term 1 from protocol ldp
user@host# set policy-options policy-statement dod-routes term 1 from tag 1145307136
user@host# set policy-options policy-statement dod-routes term 1 then accept
```

2. Configure the **do-not-propagate-du-sessions** policy to not forward routes to neighbors 1.1.1.1, 2.2.2.2, and 3.3.3.3.

```

user@host# set policy-options policy-statement do-not-propagate-du-sessions
term 1 to neighbor 1.1.1.1
user@host# set policy-options policy-statement do-not-propagate-du-sessions
term 1 to neighbor 2.2.2.2
user@host# set policy-options policy-statement do-not-propagate-du-sessions
term 1 to neighbor 3.3.3.3
user@host# set policy-options policy-statement do-not-propagate-du-sessions
term 1 then reject

```

3. Configure the **filter-dod-on-du-sessions** policy to prevent the routes examined by the **dod-routes** policy from being forwarded to the neighboring routers defined in the **do-not-propagate-du-sessions** policy.

```

user@host# set policy-options policy-statement filter-dod-routes-on-du-sessions
term 1 from policy dod-routes
user@host# set policy-options policy-statement filter-dod-routes-on-du-sessions
term 1 to policy do-not-propagate-du-sessions

```

4. Specify the **filter-dod-routes-on-du-session** policy as the export policy for BGP group **ebgp-to-abr**.

```

[edit protocols bgp]
user@host# set group ebgp-to-abr neighbor 192.168.6.2 export
filter-dod-routes-on-du-sessions

```

**Results** From configuration mode, confirm your configuration by entering the **show policy-options** and **show protocols ldp** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@host#
show policy-options
prefix-list Request-Loopbacks {
 10.1.1.1/32;
 10.1.1.2/32;
 10.1.1.3/32;
 10.1.1.4/32;
}
policy-statement DOD-Request-Loopbacks {
 term 1 {
 from {
 prefix-list Request-Loopbacks;
 }
 then accept;
 }
}
policy-statement redistribute_ldp {
 term 1 {
 from {
 protocol ldp;
 tag 1000;
 }
 then accept;
 }
}

user@host#
show protocols ldp

```

```
dod-request-policy DOD-Request-Loopbacks;
session 1.1.1.1 {
 downstream-on-demand;
}

user@host#
show protocols bgp
group ebgp-to-abr {
 type external;
 local-address 192.168.0.1;
 peer-as 65319;
 local-as 65320;
 neighbor 192.168.6.1 {
 family inet {
 unicast;
 labeled-unicast {
 rib {
 inet.3;
 }
 }
 }
 }
 export redistribute_ldp;
}
```

## Verification

### Verifying Label Advertisement Mode

---

**Purpose** Confirm that the configuration is working properly.

Use the **show ldp session** command to verify the status of the label advertisement mode for the LDP session.



**Action** Issue the `show ldp session` and `show ldp session detail` commands:

- The following command output for the `show ldp session` command indicates that the **Adv. Mode** (label advertisement mode) is **DOD** (meaning the LDP downstream on demand session is operational):

```
user@host> show ldp session
 Address State Connection Hold time Adv. Mode
 1.1.1.2 Operational Open 22 DOD
```

- The following command output for the `show ldp session detail` command indicates that the **Local Label Advertisement mode** is **Downstream unsolicited**, the default value (meaning downstream on demand is not configured on the local session). Conversely, the **Remote Label Advertisement mode** and the **Negotiated Label Advertisement mode** both indicate that **Downstream on demand** is configured on the remote session

```
user@host> show ldp session detail
Address: 1.1.1.2, State: Operational, Connection: Open, Hold time: 24
Session ID: 1.1.1.1:0--1.1.1.2:0
Next keepalive in 4 seconds
Passive, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
Neighbor types: configured-tunneled
Keepalive interval: 10, Connect retry interval: 1
Local address: 1.1.1.1, Remote address: 1.1.1.2
Up for 17:54:52
Capabilities advertised: none
Capabilities received: none
Protection: disabled
Local - Restart: disabled, Helper mode: enabled,
Remote - Restart: disabled, Helper mode: enabled
Local maximum neighbor reconnect time: 120000 msec
Local maximum neighbor recovery time: 240000 msec
Local Label Advertisement mode: Downstream unsolicited
Remote Label Advertisement mode: Downstream on demand
Negotiated Label Advertisement mode: Downstream on demand
Nonstop routing state: Not in sync
Next-hop addresses received:
 1.1.1.2
```

## Configuring the LDP Timer for Hello Messages

LDP hello messages enable LDP nodes to discover one another and to detect the failure of a neighbor or the link to the neighbor. Hello messages are sent periodically on all interfaces where LDP is enabled.

There are two types of LDP hello messages:

- Link hello messages—Sent through the LDP interface as UDP packets addressed to the LDP discovery port. Receipt of an LDP link hello message on an interface identifies an adjacency with the LDP peer router.
- Targeted hello messages—Sent as UDP packets addressed to the LDP discovery port at a specific address. Targeted hello messages are used to support LDP sessions between routers that are not directly connected. A targeted router determines whether to respond or ignore a targeted hello message. A targeted router that chooses to

respond does so by periodically sending targeted hello messages back to the initiating router.

By default, LDP sends hello messages every 5 seconds for link hello messages and every 15 seconds for targeted hello messages. You can configure the LDP timer to alter how often both types of hello messages are sent. However, you cannot configure a time for the LDP timer that is greater than the LDP hold time. For more information, see [“Configuring the Delay Before LDP Neighbors Are Considered Down” on page 4826](#).

### Configuring the LDP Timer for Link Hello Messages

To modify how often LDP sends link hello messages, specify a new link hello message interval for the LDP timer using the **hello-interval** statement:

```
hello-interval seconds;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

### Configuring the LDP Timer for Targeted Hello Messages

To modify how often LDP sends targeted hello messages, specify a new targeted hello message interval for the LDP timer by configuring the **hello-interval** statement as an option for the **targeted-hello** statement:

```
targeted-hello {
 hello-interval seconds;
}
```

For a list of hierarchy levels at which you can include these statements, see the statement summary sections for these statements.

## Configuring the Delay Before LDP Neighbors Are Considered Down

---

The hold time determines how long an LDP node should wait for a hello message before declaring a neighbor to be down. This value is sent as part of a hello message so that each LDP node tells its neighbors how long to wait. The values sent by each neighbor do not have to match.

The hold time should normally be at least three times the hello interval. The default is 15 seconds for link hello messages and 45 seconds for targeted hello messages. However, it is possible to configure an LDP hold time that is close to the value for the hello interval.



**NOTE:** By configuring an LDP hold time close to the hello interval (less than three times the hello interval), LDP neighbor failures might be detected more quickly. However, this also increases the possibility that the router might declare an LDP neighbor down that is still functioning normally. For more information, see [“Configuring the LDP Timer for Hello Messages” on page 4825](#).

The LDP hold time is also negotiated automatically between LDP peers. When two LDP peers advertise different LDP hold times to one another, the smaller value is used. If an

LDP peer router advertises a shorter hold time than the value you have configured, the peer router's advertised hold time is used. This negotiation can affect the LDP keepalive interval as well.

If the local LDP hold time is not shortened during LDP peer negotiation, the user-configured keepalive interval is left unchanged. However, if the local hold time is reduced during peer negotiation, the keepalive interval is recalculated. If the LDP hold time has been reduced during peer negotiation, the keepalive interval is reduced to one-third of the new hold time value. For example, if the new hold-time value is 45 seconds, the keepalive interval is set to 15 seconds.

This automated keepalive interval calculation can cause different keepalive intervals to be configured on each peer router. This enables the routers to be flexible in how often they send keepalive messages, because the LDP peer negotiation ensures they are sent more frequently than the LDP hold time.

When you reconfigure the hold-time interval, changes do not take effect until after the session is reset. The hold time is negotiated when the LDP peering session is initiated and cannot be renegotiated as long as the session is up (required by RFC 5036, *LDP Specification*). To manually force the LDP session to reset, issue the **clear ldp session** command.

## Configuring the LDP Hold Time for Link Hello Messages

To modify how long an LDP node should wait for a link hello message before declaring the neighbor down, specify a new time in seconds using the **hold-time** statement:

```
hold-time seconds;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

## Configuring the LDP Hold Time for Targeted Hello Messages

To modify how long an LDP node should wait for a targeted hello message before declaring the neighbor down, specify a new time in seconds using the **hold-time** statement as an option for the **targeted-hello** statement:

```
targeted-hello {
 hold-time seconds;
}
```

For a list of hierarchy levels at which you can include these statements, see the statement summary sections for these statements.

## Configuring the Interval for LDP Keepalive Messages

---

The keepalive interval determines how often a message is sent over the session to ensure that the keepalive timeout is not exceeded. If no other LDP traffic is sent over the session in this much time, a keepalive message is sent. The default is 10 seconds. The minimum value is 1 second.

The value configured for the keepalive interval can be altered during LDP session negotiation if the value configured for the LDP hold time on the peer router is lower than the value configured locally. For more information, see [“Configuring the Delay Before LDP Neighbors Are Considered Down” on page 4826](#).

To modify the keepalive interval, include the **keepalive-interval** statement:

```
keepalive-interval seconds;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

---

## Configuring the LDP Keepalive Timeout

After an LDP session is established, messages must be exchanged periodically to ensure that the session is still working. The keepalive timeout defines the amount of time that the neighbor LDP node waits before deciding that the session has failed. This value is usually set to at least three times the keepalive interval. The default is 30 seconds.

To modify the keepalive interval, include the **keepalive-timeout** statement:

```
keepalive-timeout seconds;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

The value configured for the **keepalive-timeout** statement is displayed as the hold time when you issue the **show ldp session detail** command.

---

## Configuring LDP Route Preferences

When several protocols calculate routes to the same destination, route preferences are used to select which route is installed in the forwarding table. The route with the lowest preference value is selected. The preference value can be a number in the range 0 through 255. By default, LDP routes have a preference value of 9.

To modify the route preferences, include the **preference** statement:

```
preference preference;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

---

## Configuring LDP Graceful Restart

When you alter the graceful restart configuration at either the **[edit routing-options graceful-restart]** or **[edit protocols ldp graceful-restart]** hierarchy levels, any running LDP session is automatically restarted to apply the graceful restart configuration. This behavior mirrors the behavior of BGP when you alter its graceful restart configuration.

By default, graceful restart helper mode is enabled, but graceful restart is disabled. Thus, the default behavior of a router is to assist neighboring routers attempting a graceful restart, but not to attempt a graceful restart itself.

To configure LDP graceful restart, see the following sections:

- [Enabling Graceful Restart on page 4829](#)
- [Disabling LDP Graceful Restart or Helper Mode on page 4829](#)
- [Configuring Reconnect Time on page 4830](#)
- [Configuring Recovery Time and Maximum Recovery Time on page 4830](#)

## Enabling Graceful Restart

To enable LDP graceful restart, you also need to enable graceful restart on the router.

To enable graceful restart, include the **graceful-restart** statement:

```
graceful-restart;
```

You can include this statement at the following hierarchy levels:

- **[edit routing-options]**
- **[edit logical-systems *logical-system-name* routing-options]**

The **graceful-restart** statement enables graceful restart for all protocols supporting this feature on the router. For more information about graceful restart, see the *Junos OS Routing Protocols Library for Routing Devices*.

By default, LDP graceful restart is enabled when you enable graceful restart at both the LDP protocol level and on all the routing instances. However, you can disable both LDP graceful restart and LDP graceful restart helper mode.

## Disabling LDP Graceful Restart or Helper Mode

To disable LDP graceful restart and recovery, include the **disable** statement:

```
ldp {
 graceful-restart {
 disable;
 }
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

You can disable helper mode at the LDP protocols level only. You cannot disable helper mode for a specific routing instance. To disable LDP helper mode, include the **helper-disable** statement:

```
ldp {
 graceful-restart {
 helper-disable;
 }
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

The following LDP graceful restart configurations are possible:

- LDP graceful restart and helper mode are both enabled.
- LDP graceful restart is disabled but helper mode is enabled. A router configured in this way cannot restart gracefully but can help a restarting neighbor.
- LDP graceful restart and helper mode are both disabled. The router does not use LDP graceful restart or the graceful restart type, length, and value (TLV) sent in the initialization message. The router behaves as a router that cannot support LDP graceful restart.

A configuration error is issued if you attempt to enable graceful restart and disable helper mode.

## Configuring Reconnect Time

After the LDP connection between neighbors fails, neighbors wait a certain amount of time for the gracefully restarting router to resume sending LDP messages. After the wait period, the LDP session can be reestablished. You can configure the wait period in seconds. This value is included in the fault tolerant session TLV sent in LDP initialization messages when LDP graceful restart is enabled.

Suppose that Router A and Router B are LDP neighbors. Router A is the restarting Router. The reconnect time is the time that Router A tells Router B to wait after Router B detects that Router A restarted.

To configure the reconnect time, include the **reconnect-time** statement:

```
graceful-restart {
 reconnect-time seconds;
}
```

You can set the reconnect time to a value in the range from 30 through 300 seconds. By default, it is 60 seconds.

For a list of hierarchy levels at which you can configure these statements, see the statement summary sections for these statements.

## Configuring Recovery Time and Maximum Recovery Time

The recovery time is the amount of time a router waits for LDP to restart gracefully. The recovery time period begins when an initialization message is sent or received. This period is also typically the amount of time that a neighboring router maintains its information about the restarting router, allowing it to continue to forward traffic.

To prevent a neighboring router from being adversely affected if it receives a false value for the recovery time from the restarting router, you can configure the maximum recovery time on the neighboring router. A neighboring router maintains its state for the shorter of the two times. For example, Router A is performing an LDP graceful restart. It has sent a recovery time of 900 seconds to neighboring Router B. However, Router B has its maximum recovery time configured at 400 seconds. Router B will only wait for 400 seconds before it purges its LDP information from Router A.

To configure recovery time, include the **recovery-time** statement and the **maximum-neighbor-recovery-time** statement:

```

graceful-restart {
 maximum-neighbor-recovery-time seconds;
 recovery-time seconds;
}

```

For a list of hierarchy levels at which you can configure these statements, see the statement summary sections for these statements.

## Configuring the Prefixes Advertised into LDP from the Routing Table

You can control the set of prefixes that are advertised into LDP and cause the router to be the egress router for those prefixes. By default, only the loopback address is advertised into LDP. To configure the set of prefixes from the routing table to be advertised into LDP, include the **egress-policy** statement:

```
egress-policy policy-name;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.



**NOTE:** If you configure an egress policy for LDP that does not include the loopback address, it is no longer advertised in LDP. To continue to advertise the loopback address, you need to explicitly configure it as a part of the LDP egress policy.

The named policy (configured at the **[edit policy-options]** or **[edit logical-systems logical-system-name policy-options]** hierarchy level) is applied to all routes in the routing table. Those routes that match the policy are advertised into LDP. You can control the set of neighbors to which those prefixes are advertised by using the **export** statement. Only **from** operators are considered; you can use any valid **from** operator. For more information, see the *Junos OS Routing Protocols Library for Routing Devices*.

### Example: Configuring the Prefixes Advertised into LDP

Advertise all connected routes into LDP:

```

[edit protocols]
ldp {
 egress-policy connected-only;
}
policy-options {
 policy-statement connected-only {
 from {
 protocol direct;
 }
 then accept;
 }
}

```

## Configuring LDP LSP Traceroute

---

You can trace the route followed by an LDP-signaled LSP. LDP LSP traceroute is based on RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*. This feature allows you to periodically trace all paths in a FEC. The FEC topology information is stored in a database accessible from the CLI.

A topology change does not automatically trigger a trace of an LDP LSP. However, you can manually initiate a traceroute. If the traceroute request is for an FEC that is currently in the database, the contents of the database are updated with the results.

The periodic traceroute feature applies to all FECs specified by the **oam** statement configured at the **[edit protocols ldp]** hierarchy level. To configure periodic LDP LSP traceroute, include the **periodic-traceroute** statement:

```
periodic-traceroute {
 disable;
 exp exp-value;
 fanout fanout-value;
 frequency minutes;
 paths number-of-paths;
 retries retry-attempts;
 source address;
 ttl ttl-value;
 wait seconds;
}
```

You can configure this statement at the following hierarchy levels:

- **[edit protocols ldp oam]**
- **[edit protocols ldp oam fec *address*]**

You can configure the **periodic-traceroute** statement by itself or with any of the following options:

- **exp**—Specify the class of service to use when sending probes.
- **fanout**—Specify the maximum number of next hops to search per node.
- **frequency**—Specify the interval between traceroute attempts.
- **paths**—Specify the maximum number of paths to search.
- **retries**—Specify the number of attempts to send a probe to a specific node before giving up.
- **source**—Specify the IPv4 source address to use when sending probes.
- **ttl**—Specify the maximum time-to-live value. Nodes that are beyond this value are not traced.
- **wait**—Specify the wait interval before resending a probe packet.



---

## Configuring Miscellaneous LDP Properties

---

The following sections describe how to configure a number of miscellaneous LDP properties:

- [Configuring LDP to Use the IGP Route Metric on page 4833](#)
- [Preventing Addition of Ingress Routes to the inet.0 Routing Table on page 4833](#)
- [Multiple-Instance LDP and Carrier-of-Carriers VPNs on page 4834](#)
- [Configuring MPLS and LDP to Pop the Label on the Ultimate-Hop Router on page 4834](#)
- [Enabling LDP over RSVP-Established LSPs on page 4834](#)
- [Enabling LDP over RSVP-Established LSPs in Heterogeneous Networks on page 4835](#)
- [Configuring the TCP MD5 Signature for LDP Sessions on page 4835](#)
- [Configuring LDP Session Protection on page 4836](#)
- [Disabling SNMP Traps for LDP on page 4837](#)
- [Configuring LDP Synchronization with the IGP on LDP Links on page 4837](#)
- [Configuring LDP Synchronization with the IGP on the Router on page 4838](#)
- [Configuring the Label Withdrawal Timer on page 4838](#)
- [Ignoring the LDP Subnet Check on page 4838](#)

### Configuring LDP to Use the IGP Route Metric

Use the **track-igp-metric** statement if you want the interior gateway protocol (IGP) route metric to be used for the LDP routes instead of the default LDP route metric (the default LDP route metric is 1).

To use the IGP route metric, include the **track-igp-metric** statement:

```
track-igp-metric;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

### Preventing Addition of Ingress Routes to the inet.0 Routing Table

By configuring the **no-forwarding** statement, you can prevent ingress routes from being added to the inet.0 routing table instead of the inet.3 routing table even if you enabled the **traffic-engineering bgp-igp** statement at the **[edit protocols mpls]** or the **[edit logical-systems *logical-system-name* protocols mpls]** hierarchy level. By default, the **no-forwarding** statement is disabled.

To omit ingress routes from the inet.0 routing table, include the **no-forwarding** statement:

```
no-forwarding;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

## Multiple-Instance LDP and Carrier-of-Carriers VPNs

By configuring multiple LDP routing instances, you can use LDP to advertise labels in a carrier-of-carriers VPN from a service provider provider edge (PE) router to a customer carrier customer edge (CE) router. This is especially useful when the carrier customer is a basic Internet service provider (ISP) and wants to restrict full Internet routes to its PE routers. By using LDP instead of BGP, the carrier customer shields its other internal routers from the Internet. Multiple-instance LDP is also useful when a carrier customer wants to provide Layer 2 or Layer 3 VPN services to its customers.

For an example of how to configure multiple LDP routing instances for carrier-of-carriers VPNs, see the *Multiple Instances for Label Distribution Protocol Feature Guide*.

## Configuring MPLS and LDP to Pop the Label on the Ultimate-Hop Router

The default advertised label is label 3 (Implicit Null label). If label 3 is advertised, the penultimate-hop router removes the label and sends the packet to the egress router. If ultimate-hop popping is enabled, label 0 (IPv4 Explicit Null label) is advertised. Ultimate-hop popping ensures that any packets traversing an MPLS network include a label.

To configure ultimate-hop popping, include the **explicit-null** statement:

```
explicit-null;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.



**NOTE:** Juniper Networks routers queue packets based on the incoming label. Routers from other vendors might queue packets differently. Keep this in mind when working with networks containing routers from multiple vendors.

For more information about labels, see *MPLS Label Overview* and *MPLS Label Allocation*.

## Enabling LDP over RSVP-Established LSPs

You can run LDP over LSPs established by RSVP, effectively tunneling the LDP-established LSP through the one established by RSVP. To do so, enable LDP on the lo0.0 interface (see “[Enabling and Disabling LDP](#)” on page 4810). You must also configure the LSPs over which you want LDP to operate by including the **ldp-tunneling** statement at the **[edit protocols mpls label-switched-path *lsp-name*]** hierarchy level:

```
[edit]
protocols {
 mpls {
 label-switched-path lsp-name {
 from source;
 to destination;
 ldp-tunneling;
 }
 }
}
```

```
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

**Related  
Documentation**

- [Tunneling LDP LSPs in RSVP LSPs Overview on page 4805](#)

## Enabling LDP over RSVP-Established LSPs in Heterogeneous Networks

Some other vendors use an OSPF metric of 1 for the loopback address. Juniper Networks routers use an OSPF metric of 0 for the loopback address. This might require that you manually configure the RSVP metric when deploying LDP tunneling over RSVP LSPs in heterogeneous networks.

When a Juniper Networks router is linked to another vendor's router through an RSVP tunnel, and LDP tunneling is also enabled, by default the Juniper Networks router might not use the RSVP tunnel to route traffic to the LDP destinations downstream of the other vendor's egress router if the RSVP path has a metric of 1 larger than the physical OSPF path.

To ensure that LDP tunneling functions properly in heterogeneous networks, you can configure OSPF to ignore the RSVP LSP metric by including the **ignore-lsp-metrics** statement:

```
ignore-lsp-metrics;
```

You can configure this statement at the following hierarchy levels:

- **[edit protocols ospf traffic-engineering shortcuts]**
- **[edit logical-systems *logical-system-name* protocols ospf traffic-engineering shortcuts]**

To enable LDP over RSVP LSPs, you also still need to complete the procedure in Section [“Enabling LDP over RSVP-Established LSPs” on page 4834](#).

## Configuring the TCP MD5 Signature for LDP Sessions

You can configure an MD5 signature for an LDP TCP connection to protect against the introduction of spoofed TCP segments into LDP session connection streams.

A router using the MD5 signature option is configured with a password for each peer for which authentication is required. The password is stored encrypted.

LDP hello adjacencies can still be created even when peering interfaces are configured with different security signatures. However, the TCP session cannot be authenticated and is never established.

To configure an MD5 signature for an LDP TCP connection, include the **session** and **authentication-key** statement:

```
session address {
 authentication-key md5-authentication-key;
}
```

For a list of hierarchy levels at which you can include these statements, see the statement summary section for the **session** statement.

Use the **session** statement to configure the address for the remote end of the LDP session.

The **md5-authentication-key** (password) can be up to 69 characters long. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks.

You can also configure an authentication key update mechanism for the LDP routing protocol. This mechanism allows you to update authentication keys without interrupting associated routing and signaling protocols such as Open Shortest Path First (OSPF) and Resource Reservation Setup Protocol (RSVP).

To configure the authentication key update mechanism, include the **key-chain** statement at the **[edit security authentication-key-chains]** hierarchy level, and specify the **key** option to create a keychain consisting of several authentication keys.

```
[edit security authentication-key-chains]
key-chain key-chain-name {
 key key {
 secret secret-data;
 start-time yyyy-mm-dd.hh:mm:ss;
 }
}
```

To configure the authentication key update mechanism for the LDP routing protocol, include the **authentication-key-chain** statement at the **[edit protocols ldp]** hierarchy level to associate the protocol with the **[edit security authentication-key-chains]** authentication keys.

```
[edit protocols ldp]
group group-name {
 neighbor address {
 authentication-key-chain key-chain-name;
 }
}
```

For more information about the authentication key update feature, see *Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols*.

## Configuring LDP Session Protection

An LDP session is normally created between a pair of routers that are connected by one or more links. The routers form one hello adjacency for every link that connects them and associate all the adjacencies with the corresponding LDP session. When the last hello adjacency for an LDP session goes away, the LDP session is terminated. You might want to modify this behavior to prevent an LDP session from being unnecessarily terminated and reestablished.

You can configure the Junos OS to leave the LDP session between two routers up even if there are no hello adjacencies on the links connecting the two routers by configuring the **session-protection** statement. You can optionally specify a time in seconds using the **timeout** option. The session remains up for the duration specified as long as the routers maintain IP network connectivity.

```
session-protection {
```

```

 timeout seconds;
}

```

For a list of hierarchy levels at which you can include this statement, see the statement summary section.

## Disabling SNMP Traps for LDP

Whenever an LDP LSP makes a transition from up to down, or down to up, the router sends an SNMP trap. However, it is possible to disable the LDP SNMP traps on a router, logical system, or routing instance.

For information about the LDP SNMP traps and the proprietary LDP MIB, see the *SNMP MIBs and Traps Reference* and *Interpreting the Enterprise-Specific LDP MIB*.

To disable SNMP traps for LDP, specify the **trap disable** option for the **log-updown** statement:

```

log-updown {
 trap disable;
}

```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

## Configuring LDP Synchronization with the IGP on LDP Links

LDP is a protocol for distributing labels in non-traffic-engineered applications. Labels are distributed along the best path determined by the IGP. If synchronization between LDP and the IGP is not maintained, the LSP goes down. When LDP is not fully operational on a given link (a session is not established and labels are not exchanged), the IGP advertises the link with the maximum cost metric. The link is not preferred but remains in the network topology.

LDP synchronization is supported only on active point-to-point interfaces and LAN interfaces configured as point-to-point under the IGP. LDP synchronization is not supported during graceful restart.

To advertise the maximum cost metric until LDP is operational for synchronization, include the **ldp-synchronization** statement:

```

ldp-synchronization {
 disable;
 hold-time seconds;
}

```

To disable synchronization, include the **disable** statement. To configure the time period to advertise the maximum cost metric for a link that is not fully operational, include the **hold-time** statement.

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

## Configuring LDP Synchronization with the IGP on the Router

You can configure the time the LDP waits before informing the IGP that the LDP neighbor and session for an interface are operational. For large networks with numerous FECs, you might need to configure a longer value to allow enough time for the LDP label databases to be exchanged.

To configure the time the LDP waits before informing the IGP that the LDP neighbor and session are operational, include the **igp-synchronization** statement and specify a time in seconds for the **holddown-interval** option:

```
igp-synchronization holddown-interval seconds;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

## Configuring the Label Withdrawal Timer

The label withdrawal timer delays sending a label withdrawal message for a FEC to a neighbor. When an IGP link to a neighbor fails, the label associated with the FEC has to be withdrawn from all the upstream routers if the neighbor is the next hop for the FEC. After the IGP converges and a label is received from a new next hop, the label is readvertised to all the upstream routers. This is the typical network behavior. By delaying label withdrawal by a small amount of time (for example, until the IGP converges and the router receives a new label for the FEC from the downstream next hop), the label withdrawal and sending a label mapping soon could be avoided. The **label-withdrawal-delay** statement allows you to configure this delay time. By default, the delay is 60 seconds.

If the router receives the new label before the timer runs out, the label withdrawal timer is canceled. However, if the timer runs out, the label for the FEC is withdrawn from all of the upstream routers.

By default, LDP waits for 60 seconds before withdrawing labels to avoid resignaling LSPs multiple times while the IGP is reconverging. To configure the label withdrawal delay time in seconds, include the **label-withdrawal-delay** statement:

```
label-withdrawal-delay seconds;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

## Ignoring the LDP Subnet Check

In Junos OS Release 8.4 and later releases, an LDP source address subnet check is performed during the neighbor establishment procedure. The source address in the LDP link hello packet is matched against the interface address. This causes an interoperability issue with some other vendors' equipment.

To disable the subnet check, include the **allow-subnet-mismatch** statement:

```
allow-subnet-mismatch;
```

This statement can be included at the following hierarchy levels:

- [edit protocols ldp **interface** *interface-name*]
- [edit logical-systems *logical-system-name* protocols ldp **interface** *interface-name*]





# Configuration Statements and Monitoring Commands for LDP

- [Configuration Statements for LDP on page 4841](#)
- [Monitoring Commands for LDP on page 4882](#)

## Configuration Statements for LDP

---

- [allow-subnet-mismatch on page 4842](#)
- [authentication-algorithm on page 4843](#)
- [authentication-key \(Protocols LDP\) on page 4845](#)
- [authentication-key-chain \(Protocols LDP\) on page 4846](#)
- [deaggregate on page 4847](#)
- [disable \(Protocols LDP\) on page 4848](#)
- [dod-request-policy on page 4849](#)
- [downstream-on-demand on page 4849](#)
- [ecmp on page 4850](#)
- [egress-policy on page 4850](#)
- [explicit-null \(Protocols LDP\) on page 4851](#)
- [export \(Protocols LDP\) on page 4851](#)
- [fec on page 4852](#)
- [graceful-restart \(Protocols LDP\) on page 4853](#)
- [hello-interval \(Protocols LDP\) on page 4854](#)
- [helper-disable \(LDP\) on page 4855](#)
- [hold-time \(Protocols LDP\) on page 4856](#)
- [ignore-lsp-metrics on page 4857](#)
- [igp-synchronization on page 4857](#)
- [import \(Protocols LDP\) on page 4858](#)
- [interface \(Protocols LDP\) on page 4859](#)
- [keepalive-interval on page 4860](#)

- [keepalive-timeout](#) on page 4861
- [l2-smart-policy](#) on page 4861
- [label-withdrawal-delay](#) on page 4862
- [ldp](#) on page 4863
- [ldp-synchronization](#) on page 4866
- [ldp-tunneling](#) on page 4866
- [log-updown \(Protocols LDP\)](#) on page 4867
- [maximum-neighbor-recovery-time](#) on page 4868
- [no-forwarding](#) on page 4869
- [policing \(Protocols LDP\)](#) on page 4870
- [preference \(Protocols LDP\)](#) on page 4871
- [reconnect-time](#) on page 4872
- [recovery-time](#) on page 4873
- [session \(ldp\)](#) on page 4874
- [session-protection](#) on page 4875
- [strict-targeted-hellos](#) on page 4875
- [targeted-hello](#) on page 4876
- [traceoptions \(Protocols LDP\)](#) on page 4877
- [track-igp-metric](#) on page 4879
- [traffic-statistics \(Protocols LDP\)](#) on page 4880
- [transport-address](#) on page 4882

---

## allow-subnet-mismatch

---

|                                 |                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>allow-subnet-mismatch;</code>                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols ldp interface <i>interface-name</i>],</code><br><code>[edit protocols ldp interface <i>interface-name</i>]</code>                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.3.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                      |
| <b>Description</b>              | Ignore the LDP subnet check. For Junos OS Release 8.4 and later releases, an LDP source address subnet check was added for the neighbor establishment procedure. The source address in the LDP link hello packet is matched against the interface address. |
| <b>Default</b>                  | The source address in the LDP link hello packet is matched against the interface address.                                                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Ignoring the LDP Subnet Check</a> on page 4838</li></ul>                                                                                                                                               |

## authentication-algorithm

**Syntax** authentication-algorithm *algorithm*;

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols bgp],  
 [edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],  
 [edit logical-systems *logical-system-name* protocols bgp **group** *group-name* neighbor *address*],  
 [edit logical-systems *logical-system-name* protocols ldp session *session-address*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp session *session-address*],  
 [edit logical-systems *logical-system-name* routing-options **bmp**],  
 [edit logical-systems *logical-system-name* routing-options bmp **station** *station-name*],  
 [edit protocols bgp],  
 [edit protocols bgp **group** *group-name*],  
 [edit protocols bgp group *group-name* **neighbor** *address*],  
 [edit protocols ldp session *session-address*],  
 [edit routing-instances *routing-instance-name* protocols bgp],  
 [edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*],  
 [edit routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*],  
 [edit routing-instances *routing-instance-name* protocols ldp session *session-address*],  
 [edit routing-options **bmp**],  
 [edit routing-options bmp **station** *station-name*]

**Release Information** Statement introduced in Junos OS Release 7.6.  
 Statement introduced for BGP in Junos OS Release 8.0.  
 Statement introduced in Junos OS Release 9.0 for EX Series switches.  
 Statement introduced in Junos OS Release 12.3X50 for the QFX Series.  
 Statement introduced for BMP in Junos OS Release 13.2X51-D15 for the QFX Series.  
 Statement introduced for BMP in Junos OS Release 13.3.  
 Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Configure an authentication algorithm type.



**NOTE:** Keep the following points in mind when you configure the authentication algorithm in an IPsec proposal:

- When both ends of an IPsec VPN tunnel contain the same IKE proposal but different IPsec proposals, an error occurs and the tunnel is not established in this scenario. For example, if one end of the tunnel contains router 1 configured with the authentication algorithm as hmac-sha-256-128 and the other end of the tunnel contains router 2 configured with the authentication algorithm as hmac-md5-96, the VPN tunnel is not established.

- When both ends of an IPsec VPN tunnel contain the same IKE proposal but different IPsec proposals, and when one end of the tunnel contains two IPsec proposals to check whether a less secure algorithm is selected or not, an error occurs and the tunnel is not established. For example, if you configure two authentication algorithms for an IPsec proposal as `hmac-sha-256-128` and `hmac-md5-96` on one end of the tunnel, router 1, and if you configure the algorithm for an IPsec proposal as `hmac-md5-96` on the other end of the tunnel, router 2, the tunnel is not established and the number of proposals mismatch.
  - When you configure two IPsec proposals at both ends of a tunnel, such as the `authentication-algorithm hmac-sha-256-128` and `authentication-algorithm hmac-md5-96` statements at the `[edit services ipsec-vpn ipsec proposal proposal-name]` hierarchy level on one of the tunnel, router 1 (with the algorithms in two successive statements to specify the order), and the `authentication-algorithm hmac-md5-96` and `authentication-algorithm hmac-sha-256-128` statements at the `[edit services ipsec-vpn ipsec proposal proposal-name]` hierarchy level on one of the tunnel, router 2 (with the algorithms in two successive statements to specify the order, which is the reverse order of router 1), the tunnel is established in this combination as expected because the number of proposals is the same on both ends and they contain the same set of algorithms. However, the authentication algorithm selected is `hmac-md5-96` and not the stronger algorithm of `hmac-sha-256-128`. This method of selection of the algorithm occurs because the first matching proposal is selected. Also, for a default proposal, regardless of whether the router supports the Advanced Encryption Standard (AES) encryption algorithm, the `3des-cbc` algorithm is chosen and not the `aes-cfb` algorithm, which is because of the first algorithm in the default proposal being selected. In the sample scenario described here, on router 2, if you reverse the order of the algorithm configuration in the proposal so that it is the same order as the one specified on router 1, `hmac-sha-256-128` is selected as the authentication method.
  - You must be aware of the order of proposals in an IPsec policy at the time of configuration if you want the matching of proposals to happen in a certain order of preference, such as the strongest algorithm to be considered first when a match is made when both policies from the two peers have a proposal.
-

**Options** *algorithm*—Specify one of the following types of authentication algorithms:

- **aes-128-cmac-96**—Cipher-based message authentication code (AES128, 96 bits).
- **hmac-sha-1-96**—Hash-based message authentication code (SHA1, 96 bits).
- **md5**—Message digest 5.

**Default:** hmac-sha-1-96



**NOTE:** The default is not displayed in the output of the `show bgp bmp` command unless a key or key-chain is also configured.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Configuring Route Authentication for BGP on page 3898](#)
- [Configuring BGP Monitoring Protocol Version 3 on page 3647](#)

## authentication-key (Protocols LDP)

**Syntax** authentication-key *md5-authentication-key*;

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols ldp session *address*],  
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp session *address*],  
[edit protocols ldp session *address*],  
[edit routing-instances *routing-instance-name* protocols ldp session *address*]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

**Description** Configure the MD5 authentication signature. The maximum length of the authentication signature is 69 characters.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring the TCP MD5 Signature for LDP Sessions on page 4835](#)

## authentication-key-chain (Protocols LDP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | authentication-key-chain <i>key-chain</i> ;                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>name</i> protocols ldp session <i>address</i> ],<br>[edit logical-systems <i>name</i> routing-instances <i>instance-name</i> protocols ldp session <i>address</i> ],<br>[edit protocols ldp session <i>address</i> ],<br>[edit routing-instances <i>instance-name</i> protocols ldp session <i>address</i> ]                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.0.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                           |
| <b>Description</b>              | Apply and enable an authentication keychain to the routing device. Note that the referenced key chain must be defined. When configuring the authentication key update mechanism for LDP, you cannot commit the <b>0.0.0.0/allow</b> statement with authentication keys or key chains. The CLI issues a warning and fails to commit such configurations. |
| <b>Options</b>                  | <i>key-chain</i> —Authentication keychain name. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").                                                                                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols</i></li><li>• <a href="#">Configuring Miscellaneous LDP Properties on page 4833</a></li></ul>                                                                                                                          |

## deaggregate

---

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | deaggregate   no-deaggregate;                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br>[edit protocols ldp],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                           |
| <b>Description</b>              | Control forwarding equivalence class (FEC) deaggregation on the router. The use of the <b>deaggregate</b> statement in LDP is a standard practice that we recommend for LDP deployments.                                                                                            |
| <b>Default</b>                  | Deaggregation is disabled on the router.                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <b>deaggregate</b> —Deaggregate FECs.<br><b>no-deaggregate</b> —Aggregate FECs.                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring FEC Deaggregation</i></li> </ul>                                                                                                                                                                                            |

## disable (Protocols LDP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | disable;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart],<br>[edit logical-systems <i>logical-system-name</i> protocols ldp interface <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options graceful-restart],<br>[edit protocols ldp graceful-restart],<br>[edit protocols ldp interface <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> routing-options graceful-restart] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Explicitly disable LDP on an interface, or explicitly disable LDP graceful restart.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Default</b>                  | LDP is enabled on interfaces configured with the LDP <b>interface</b> statement. LDP graceful restart is automatically enabled when graceful restart is enabled under the <b>[edit routing-options]</b> hierarchy level.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Enabling and Disabling LDP on page 4810</a></li><li>• <a href="#">Configuring LDP Graceful Restart on page 4828</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |



## dod-request-policy

|                                 |                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>dod-request-policy <i>dod-request-policy-name</i>;</code>                                                                                                            |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit protocols ldp]                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                     |
| <b>Description</b>              | Specify the name of the LDP downstream on demand request policy. LDP sends label request messages only for those FECs matching in the downstream on demand request policy. |
| <b>Options</b>                  | <i>dod-request-policy-name</i> —Specify the name of the downstream on demand request policy.                                                                               |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring LDP Downstream on Demand on page 4820</a></li> </ul>                                             |

## downstream-on-demand

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>downstream-on-demand;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit logical systems <i>logical-system-name</i> protocols ldp session <i>session-address</i> ],<br>[edit protocols ldp session <i>session-address</i> ]                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | Enable LDP downstream on demand on the LDP session. LDP is widely deployed in downstream unsolicited advertisement mode. As service providers integrate the access and aggregation networks into a single MPLS domain, LDP downstream on demand is needed to distribute the bindings between access and aggregation networks to minimize the workload for the access node (AN) control plane and to avoid the storage of tens of thousands of label bindings from upstream aggregation nodes. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring LDP Downstream on Demand on page 4820</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                |

## ecmp

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | ecmp;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp oam bfd-liveness-detection],<br>[edit logical-systems <i>logical-system-name</i> protocols ldp oam fec <i>address</i><br>bfd-liveness-detection],<br>[edit protocols ldp oam bfd-liveness-detection],<br>[edit protocols ldp oam fec <i>address</i> bfd-liveness-detection]                                                                                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.<br>Statement introduced in Junos OS Release 15.1X53-D30 for QFX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Allows LDP to establish BFD sessions for all ECMP paths configured for the specified FEC. If you configure the <b>ecmp</b> statement, you must also configure the <b>periodic-traceroute</b> statement for the specified FEC. If you do not do so, the commit operation fails. You can configure the <b>periodic-traceroute</b> statement at the global hierarchy level ([edit protocols ldp oam]) while only configuring the <b>ecmp</b> statement for a specific FEC ([edit protocols ldp oam fec <i>address</i> bfd-liveness-detection]). |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring ECMP-Aware BFD for LDP LSPs</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## egress-policy

---

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | egress-policy [ <i>policy-names</i> ];                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br>[edit protocols ldp],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                           |
| <b>Description</b>              | Control the prefixes advertised into LDP.                                                                                                                                                                                                                                           |
| <b>Default</b>                  | Only the loopback address is advertised.                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more routing policies.                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the Prefixes Advertised into LDP from the Routing Table on page 4831</a></li></ul>                                                                                                                                  |

## explicit-null (Protocols LDP)

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>explicit-null;</code>                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br>[edit protocols ldp],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                           |
| <b>Description</b>              | Advertise label 0 to the egress router of a label-switched path (LSP).                                                                                                                                                                                                              |
| <b>Default</b>                  | If you do not include the <b>explicit-null</b> statement in the MPLS configuration, label 3 (implicit null) is advertised.                                                                                                                                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring MPLS and LDP to Pop the Label on the Ultimate-Hop Router on page 4834</a></li> </ul>                                                                                                                               |

## export (Protocols LDP)

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>export [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br>[edit protocols ldp],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                           |
| <b>Description</b>              | Apply policy filters to outbound LDP label bindings. Filters are applied to all label bindings from all neighbors.                                                                                                                                                                  |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more routing policies.                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Filtering Outbound LDP Label Bindings on page 4813</a></li> </ul>                                                                                                                                                              |

## fec

---

**Syntax**    `fec fec-address {  
              bfd-liveness-detection {  
                  detection-time threshold milliseconds;  
                  ecmp;  
                  failure-action {  
                      remove-nexthop;  
                      remove-route;  
                  }  
                  holddown-interval milliseconds;  
                  ingress-policy ingress-policy-name;  
                  minimum-interval milliseconds;  
                  minimum-receive-interval milliseconds;  
                  minimum-transmit-interval milliseconds;  
                  multiplier detection-time-multiplier;  
                  no-adaptation;  
                  transmit-interval {  
                      minimum-interval milliseconds;  
                      threshold milliseconds;  
                  }  
                  version (0 | 1 | automatic);  
              }  
              no-bfd-liveness-detection;  
              periodic-traceroute {  
                  disable;  
                  exp exp-value;  
                  fanout fanout-value;  
                  frequency minutes;  
                  paths number-of-paths;  
                  retries retry-attempts;  
                  source address;  
                  ttl ttl-value;  
                  wait seconds;  
              }  
          }`

**Hierarchy Level**    [edit logical-systems *logical-systems-name* protocols ldp oam],  
                          [edit protocols ldp oam]

**Release Information**    Statement introduced in Junos OS Release 8.5.  
                              Statement introduced in Junos OS Release 12.2 for EX Series switches.  
                              Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

**Description**    Allows you to configure BFD for a specific LDP forwarding equivalence class (FEC).

**Options**    *fec-address*—Specify the FEC address.  
  
              The other statements are explained separately.

**Required Privilege Level**    routing—To view this statement in the configuration.  
                                  routing-control—To add this statement to the configuration.

**Related Documentation** • [Configuring BFD for LDP LSPs](#)

## graceful-restart (Protocols LDP)

|                            |                                                                                                                                                                                                                                                                                       |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>graceful-restart {   disable;   helper-disable;   maximum-neighbor-recovery-time <i>value</i>;   reconnect-time <i>seconds</i>;   recovery-time <i>value</i>; }</pre>                                                                                                            |
| <b>Hierarchy Level</b>     | <pre>[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]</pre> |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>                                                                                                                                                  |
| <b>Description</b>         | Configure LDP graceful restart on the LDP master protocol instance or for a specific routing instance.                                                                                                                                                                                |



**NOTE:** When you alter the graceful restart configuration at either the [edit routing-options graceful-restart] or [edit protocols ldp graceful-restart] hierarchy levels, any running LDP session is automatically restarted to apply the graceful restart configuration. This behavior mirrors the behavior of BGP when you alter its graceful restart configuration.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation** • [Configuring LDP Graceful Restart on page 4828](#)

## hello-interval (Protocols LDP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>hello-interval <i>seconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols ldp interface <i>interface-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols ldp targeted-hello],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  ldp interface <i>interface-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  ldp targeted-hello],</code><br><code>[edit protocols ldp interface <i>interface-name</i>],</code><br><code>[edit protocols ldp targeted-hello],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols ldp targeted-hello]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Support for LDP targeted hellos added in Junos OS Release 9.5.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Control the LDP timer that regulates how often hello messages are sent. You can control the rate both link hello messages and targeted hello messages are sent depending on the hierarchy level at which you configure the <b>hello-interval</b> statement.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <b><i>seconds</i></b> —Length of time between transmission of hello packets.<br><b>Range:</b> 1 through 65,535 seconds<br><b>Default:</b> 5 seconds for link hello messages, 15 seconds for targeted hello messages                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the LDP Timer for Hello Messages on page 4825</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

---

## helper-disable (LDP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | helper-disable;                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart],<br>[edit protocols ldp graceful-restart],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                               |
| <b>Description</b>              | Disable helper mode for LDP graceful restart. When helper mode is disabled, a router cannot help a neighboring router that is attempting to restart LDP.                                                                                                                                                                                                |
| <b>Default</b>                  | Helper mode is enabled by default on all routing protocols (including LDP) that support graceful restart.                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring LDP Graceful Restart on page 4828</a></li></ul>                                                                                                                                                                                                                                         |

## hold-time (Protocols LDP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>hold-time seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols ldp interface <i>interface-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols ldp targeted-hello],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  ldp interface <i>interface-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  ldp targeted-hello],</code><br><code>[edit protocols ldp interface <i>interface-name</i>],</code><br><code>[edit protocols ldp targeted-hello],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols ldp targeted-hello]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Support for LDP targeted hellos added in Junos OS Release 9.5.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Specify how long an LDP node should wait for a hello message before declaring a neighbor to be down. This value is sent as part of a hello message so that each LDP node tells its neighbors how long to wait. You can specify times for both link hello messages and targeted hello messages depending on the hierarchy level at which you configure the <b>hold-time</b> statement.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <b>seconds</b> —Hold-time value.<br><b>Range:</b> 1 through 65,535 seconds<br><b>Default:</b> 15 seconds for link hello messages, 45 seconds for targeted hello messages                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the Delay Before LDP Neighbors Are Considered Down on page 4826</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |



## ignore-lsp-metrics

|                                 |                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | ignore-lsp-metrics;                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ospf traffic-engineering shortcuts],<br>[edit protocols ospf traffic-engineering shortcuts]                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.5.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                |
| <b>Description</b>              | Cause OSPF to ignore the RSVP LSP metric.<br><br>Some other vendors use an OSPF metric of 1 for the loopback address. Juniper Networks routers use an OSPF metric of 0 for the loopback address. This can cause interoperability problems when you configure LDP tunneling over RSVP LSPs in heterogeneous networks. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Enabling LDP over RSVP-Established LSPs in Heterogeneous Networks on page 4835</a></li> </ul>                                                                                                                                                                   |

## igp-synchronization

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | igp-synchronization holddown-interval <i>seconds</i> ;                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br>[edit protocols ldp],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.5.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                               |
| <b>Description</b>              | Configure the time the LDP waits before informing the IGP that the LDP neighbor and session for an interface are operational. For large networks with numerous FECs, you might need to configure a longer value to allow enough time for the LDP label databases to be exchanged.   |
| <b>Options</b>                  | <b>holddown-interval <i>seconds</i></b> —Time the LDP waits before informing the IGP that the LDP neighbor and session for an interface are operational.<br><b>Default:</b> 10 seconds<br><b>Range:</b> 10 through 60 seconds                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring LDP Synchronization with the IGP on the Router on page 4838</a></li> </ul>                                                                                                                                         |

## import (Protocols LDP)

---

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>import [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br>[edit protocols ldp],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                           |
| <b>Description</b>              | Apply policy filters to received LDP label bindings. Filters are applied to all label bindings from all neighbors.                                                                                                                                                                  |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more routing policies.                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Filtering Inbound LDP Label Bindings on page 4811</a></li></ul>                                                                                                                                                                 |

## interface (Protocols LDP)

|                                 |                                                                                                                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>interface <i>interface-name</i> {     disable;     hello-interval <i>seconds</i>;     hold-time <i>seconds</i>;     transport-address (interface   loopback); }</pre>                                                                                                            |
| <b>Hierarchy Level</b>          | <pre>[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]</pre> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>                                                                                                                                                  |
| <b>Description</b>              | Enable LDP on one or more router interfaces.                                                                                                                                                                                                                                          |
| <b>Default</b>                  | LDP is disabled on all interfaces.                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <p><i>interface-name</i>—Name of an interface. To configure all interfaces, specify <b>all</b>.</p> <p>The remaining statements are explained separately.</p>                                                                                                                         |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Enabling and Disabling LDP on page 4810</a></li> </ul>                                                                                                                                                                           |

## keepalive-interval

---

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>keepalive-interval <i>seconds</i>;</code>                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br>[edit protocols ldp],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                           |
| <b>Description</b>              | Set the keepalive interval value.                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <b><i>seconds</i></b> —Keepalive value.<br><b>Range:</b> 1 through 65,535<br><b>Default:</b> 10 seconds                                                                                                                                                                             |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the Interval for LDP Keepalive Messages on page 4827</a></li></ul>                                                                                                                                                  |

## keepalive-timeout

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>keepalive-timeout <i>seconds</i>;</code>                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br>[edit protocols ldp],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                           |
| <b>Description</b>              | Set the keepalive timeout value. The keepalive timeout defines the amount of time that the neighbor LDP node waits before determining that the session has failed.                                                                                                                  |
| <b>Options</b>                  | <b><i>seconds</i></b> —Keepalive timeout value.<br><b>Range:</b> 1 through 65,535<br><b>Default:</b> 30 seconds                                                                                                                                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring the LDP Keepalive Timeout on page 4828</a></li> </ul>                                                                                                                                                              |

## l2-smart-policy

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>l2-smart-policy;</code>                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br>[edit protocols ldp],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                               |
| <b>Description</b>              | Prevent LDP from exporting IPv4 FECs over sessions with Layer 2 neighbors only. IPv4 FECs received over such sessions are filtered out.                                                                                                                                             |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring LDP IPv4 FEC Filtering</a></li> </ul>                                                                                                                                                                              |

## label-withdrawal-delay

---

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | label-withdrawal-delay <i>seconds</i> ;                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br>[edit protocols ldp],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.1.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                               |
| <b>Description</b>              | Delay the withdrawal of labels to reduce router workload during IGP convergence.                                                                                                                                                                                                    |
| <b>Options</b>                  | <b>seconds</b> —Configure the number of seconds to wait before withdrawing labels for the LDP LSPs.<br><b>Default:</b> 60 seconds<br><b>Range:</b> 0 through 300 seconds                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the Label Withdrawal Timer on page 4838</a></li></ul>                                                                                                                                                               |

## ldp

```
Syntax ldp {
 (deaggregate | no-deaggregate);
 egress-policy [policy-names];
 explicit-null;
 export [policy-names];
 graceful-restart {
 disable;
 helper-disable;
 maximum-neighbor-recovery-time seconds;
 reconnect-time seconds;
 recovery-time seconds;
 }
 import [policy-names];
 interface (interface-name | all) {
 disable;
 hello-interval seconds;
 hold-time seconds;
 transport-address (interface | router-id);
 }
 keepalive-interval seconds;
 keepalive-timeout seconds;
 log-updown {
 trap disable;
 }
 no-forwarding;
 oam {
 bfd-liveness-detection {
 detection-time threshold milliseconds;
 ecmp;
 failure-action {
 remove-nexthop;
 remove-route;
 }
 holddown-interval milliseconds;
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 minimum-transmit-interval milliseconds;
 multiplier detection-time-multiplier;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 }
 }
 fec fec-address {
 bfd-liveness-detection {
 detection-time threshold milliseconds;
 ecmp;
 failure-action {
 remove-nexthop;
 remove-route;
 }
 }
 }
}
```

```
 holddown-interval milliseconds;
 ingress-policy ingress-policy-name;
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 minimum-transmit-interval milliseconds;
 multiplier detection-time-multiplier;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 version (0 | 1 | automatic);
}
no-bfd-liveness-detection;
periodic-traceroute {
 disable;
 exp exp-value;
 fanout fanout-value;
 frequency minutes;
 paths number-of-paths;
 retries retry-attempts;
 source address;
 ttl ttl-value;
 wait seconds;
}
}
ingress-policy ingress-policy-name;
periodic-traceroute {
 disable;
 exp exp-value;
 fanout fanout-value;
 frequency minutes;
 paths number-of-paths;
 retries retry-attempts;
 source address;
 ttl ttl-value;
 wait seconds;
}
}
p2mp;
policing {
 fec fec-address {
 ingress-traffic filter-name;
 transit-traffic filter-name;
 }
}
preference preference;
session address {
 authentication-algorithm algorithm;
 authentication-key authentication-key;
 authentication-key-chain key-chain-name;
}
strict-targeted-hellos;
traceoptions {
 file filename <files number <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
```



```

}
track-igp-metric;
traffic-statistics {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 interval interval;
 no-penultimate-hop;
}
transport-address (address | interface | router-id);
}

```

|                                 |                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols],<br>[edit protocols],<br>[edit routing-instances <i>routing-instance-name</i> protocols] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 11.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                  |
| <b>Description</b>              | Enable LDP routing on the router or switch.<br><br>You must include the <b>ldp</b> statement in the configuration to enable LDP on the router or switch.                                                                                                            |
| <b>Default</b>                  | LDP is disabled on the router.                                                                                                                                                                                                                                      |
| <b>Options</b>                  | The other statements are explained separately.                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Minimum LDP Configuration on page 4810</a></li> <li>• <a href="#">Enabling and Disabling LDP on page 4810</a></li> </ul>                                                                                       |

## ldp-synchronization

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>ldp-synchronization {<br/>    disable;<br/>    hold-time seconds;<br/>}</code>                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ospf interface <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf interface <i>interface-name</i> ],<br>[edit protocols ospf interface <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols ospf interface <i>interface-name</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.5.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Enable synchronization by advertising the maximum cost metric until LDP is operational on the link.                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | The other statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring LDP Synchronization with the IGP on LDP Links on page 4837</a></li></ul>                                                                                                                                                                                                                                                                                    |

## ldp-tunneling

---

|                                 |                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>ldp-tunneling;</code>                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> ],<br>[edit protocols mpls label-switched-path <i>lsp-name</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                            |
| <b>Description</b>              | Enable the LSP to be used for LDP tunneling.                                                                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Enabling LDP over RSVP-Established LSPs on page 4834</a></li></ul>                                               |

## log-updown (Protocols LDP)

---

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | log-updown {<br>trap disable;<br>}                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br>[edit protocols ldp],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                           |
| <b>Description</b>              | Disable LDP traps on the router, logical system, or routing instance.                                                                                                                                                                                                               |
| <b>Options</b>                  | <b>trap disable</b> —Disable LDP traps.<br><b>Default:</b> LDP traps are enabled on the router.                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Disabling SNMP Traps for LDP on page 4837</a></li></ul>                                                                                                                                                                         |

## maximum-neighbor-recovery-time

---

|                                 |                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>maximum-neighbor-recovery-time seconds;</code>                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart],<br>[edit protocols ldp graceful-restart],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4. Statement changed from <b>maximum-recovery-time</b> to <b>maximum-neighbor-recovery-time</b> in Junos OS Release 9.1. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                            |
| <b>Description</b>              | Specify the maximum amount of time to wait before giving up an attempt to gracefully restart.                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <b>seconds</b> —Configure the maximum recovery time, in seconds.<br><b>Range:</b> 120 through 1800 seconds<br><b>Default:</b> 140 seconds                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Recovery Time and Maximum Recovery Time on page 4830</a></li><li>• <a href="#">Configuring Graceful Restart Options for LDP</a></li><li>• <a href="#">no-strict-lsa-checking on page 2667</a></li><li>• <a href="#">recovery-time</a></li></ul>                                         |

## no-forwarding

---

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-forwarding;                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br>[edit protocols ldp],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                           |
| <b>Description</b>              | Do not add ingress routes to the inet.0 routing table even if <a href="#">traffic-engineering bgp-igp</a> (configured at the <a href="#">[edit protocols mpls]</a> hierarchy level) is enabled.                                                                                     |
| <b>Default</b>                  | The <b>no-forwarding</b> statement is disabled. Ingress routes are added to the inet.0 routing table instead of the inet.3 routing table when <a href="#">traffic-engineering bgp-igp</a> is enabled.                                                                               |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Preventing Addition of Ingress Routes to the inet.0 Routing Table on page 4833</a></li> <li>• <a href="#">Configuring Virtual-Router Routing Instances in VPNs</a></li> </ul>                                                  |

## policing (Protocols LDP)

---

|                                 |                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>policing {<br/>    fec <i>fec-address</i> {<br/>        ingress-traffic <i>filter-name</i>;<br/>        transit-traffic <i>filter-name</i>;<br/>    }<br/>}</pre>                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br>[edit protocols ldp],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp]                       |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                 |
| <b>Description</b>              | Enable policing of forwarding equivalence classes (FECs) for LDP.                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <p><b>fec <i>fec-address</i></b>—Specify the address for the FEC.</p> <p><b>ingress-traffic <i>filter-name</i></b>—Specify the name of the filter for policing ingress FEC traffic.</p> <p><b>transit-traffic <i>filter-name</i></b>—Specify the name of the filter for policing transit FEC traffic.</p> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Policers for LDP FECs</i></li></ul>                                                                                                                                                                                                                |

## preference (Protocols LDP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>preference <i>preference</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit protocols ldp interface <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br>[edit protocols ldp],<br>[edit protocols ldp interface <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Set the route preference level for LDP routes.                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <p><i>preference</i>—Preferred value.</p> <p><b>Range:</b> 0 through 255</p> <p><b>Default:</b> 9</p>                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring LDP Route Preferences on page 4828</a></li> </ul>                                                                                                                                                                                                                                                                                                                       |

## reconnect-time

---

|                                 |                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>reconnect-time seconds;</code>                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp <a href="#">graceful-restart</a> ],<br>[edit protocols ldp <a href="#">graceful-restart</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp <a href="#">graceful-restart</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.1.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                  |
| <b>Description</b>              | Specify the length of time required to reestablish a Label Distribution Protocol (LDP) session after graceful restart.                                                                                                                                                 |
| <b>Options</b>                  | <b>seconds</b> —Time required for reconnection.<br><b>Range:</b> 30 through 300<br><b>Default:</b> 60 seconds                                                                                                                                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring LDP Graceful Restart on page 4828</a> on <i>MPLS Applications Feature Guide for Routing Devices</i></li><li>• <i>Configuring Graceful Restart Options for LDP</i></li></ul>                            |



## recovery-time

---

|                                 |                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>recovery-time seconds;</code>                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart],<br>[edit protocols ldp graceful-restart],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                               |
| <b>Description</b>              | Specify the amount of time a router waits for LDP to restart gracefully.                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <b>seconds</b> —Configure the recovery time, in seconds.<br><b>Range:</b> 120 through 1800 seconds<br><b>Default:</b> 140 seconds                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Recovery Time and Maximum Recovery Time on page 4830</a></li> </ul>                                                                                                                                                                                                                    |

## session (ldp)

---

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>session address {<br/>    authentication-algorithm <i>algorithm</i>;<br/>    authentication-key <i>authentication-key</i>;<br/>    authentication-key-chain <i>key-chain-name</i>;<br/>}</pre>                                                                                 |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br>[edit protocols ldp],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br><b>authentication-algorithm</b> statement introduced in Junos OS Release 7.6.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                          |
| <b>Description</b>              | Specify the address for the remote end of the LDP session.<br><br>The remaining statements are explained separately.                                                                                                                                                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the TCP MD5 Signature for LDP Sessions on page 4835</a></li></ul>                                                                                                                                                   |

## session-protection

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | session-protection {<br>timeout <i>seconds</i> ;<br>}                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br>[edit protocols ldp],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp]                                                                                        |
| <b>Description</b>              | Configure when an LDP session is torn down and resigaled after the router stops receiving hello messages from a neighboring router. You might want to modify this behavior to prevent an LDP session from being unnecessarily terminated and reestablished. The LDP session remains up for the duration specified as long as the routers maintain IP network connectivity. |
| <b>Options</b>                  | <b>timeout <i>seconds</i></b> —Time in seconds before the LDP session is torn down and resigaled.<br><b>Range:</b> 1 through 65,535 seconds                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring LDP Session Protection on page 4836</a></li> </ul>                                                                                                                                                                                                                                                        |

## strict-targeted-hellos

---

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | strict-targeted-hellos;                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br>[edit protocols ldp],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                           |
| <b>Description</b>              | Prevent LDP sessions from being established with remote neighbors that have not been specifically configured. LDP peers will not respond to targeted hellos coming from a source that is not one of the configured remote neighbors.                                                |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Enabling Strict Targeted Hello Messages for LDP on page 4810</a></li> </ul>                                                                                                                                                    |

## targeted-hello

---

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | targeted-hello {<br>hello-interval <i>seconds</i> ;<br>hold-time <i>seconds</i> ;<br>}                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br>[edit protocols ldp],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.5.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                               |
| <b>Description</b>              | Specify the LDP timer and LDP hold time for targeted hellos.                                                                                                                                                                                                                        |
| <b>Options</b>                  | The remaining statements are explained separately.                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the LDP Timer for Hello Messages on page 4825</a></li><li>• <a href="#">Configuring the Delay Before LDP Neighbors Are Considered Down on page 4826</a></li></ul>                                                   |

## traceoptions (Protocols LDP)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre> traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;flag-modifier&gt; &lt;disable&gt;; } </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols <i>ldp</i>],<br/> [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>ldp</i>],<br/> [edit protocols <i>ldp</i>],<br/> [edit routing-instances <i>routing-instance-name</i> protocols <i>ldp</i>]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p><b>match-on address</b> option for the <b>filter</b> flag modifier added in Junos OS Release 10.4.</p> <p><b>nsr-synchronization</b> and <b>p2mp-nsr-synchronization</b> operations for <b>flag</b> statement introduced in Junos OS Release 13.3.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>         | Specify LDP protocol-level trace options.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Default</b>             | The default LDP protocol-level trace options are inherited from the routing protocols <b>traceoptions</b> statement included at the [edit routing-options] hierarchy level.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>             | <p><b>disable</b>—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>ldp-log</b>. We recommend that you place LDP tracing output in the file <b>ldp-log</b>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p><b>Range:</b> 2 through 1000</p> <p><b>Default:</b> 2 files</p> <p>If you specify a maximum number of files, you must also include the <b>size</b> statement to specify the maximum file size.</p> <p><b>flag <i>flag</i></b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements.</p> <ul style="list-style-type: none"> <li>• <b>address</b>—Operation of address and address withdrawal messages</li> <li>• <b>binding</b>—Label-binding operations</li> <li>• <b>error</b>—Error conditions</li> <li>• <b>event</b>—Protocol events</li> </ul> |

- **initialization**—Operation of initialization messages
- **label**—Operation of label request, label map, label withdrawal, and label release messages
- **notification**—Operation of notification messages
- **nsr-synchronization**—Nonstop active routing synchronization events
- **p2mp-nsr-synchronization**—Point-to-multipoint nonstop active routing synchronization events
- **packets**—Equivalent to setting **address**, **initialization**, **label**, **notification**, and **periodic** flags (see also the **filter** flag modifier)
- **path**—Label-switched path operations
- **periodic**—Operation of hello and keepalive messages
- **route**—Operation of route messages
- **state**—Protocol state transitions

**flag-modifier**—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information.
- **disable**—Disable this trace flag.
- **filter**—Filter to apply to this flag. The **filter** flag modifier can be applied only to the **route**, **path**, and **binding** flags. This flag modifier has the following options:
  - **match-on**—Match on argument specified. The **match-on** option has the following suboptions:
    - **address**—Filter based on the source and destination addresses of packets. Available for the **packets** flag option only.
    - **fec**—Filter based on the FEC associated with the traced object.
  - **policy *policy-name***—Specify the filter policy.
- **receive**—Packets being received.
- **send**—Packets being transmitted.

**no-world-readable**—(Optional) Prevent all users from reading the log file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches this size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 1 MB

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of files.

**world-readable**—(Optional) Enable any user to read the log file.

**Required Privilege Level** routing and trace—To view this statement in the configuration.  
routing-control and trace-control—To add this statement to the configuration.

**Related Documentation**

- [Tracing LDP Protocol Traffic on page 4817](#)
- *Network Management Administration Guide for Routing Devices*

## track-igp-metric

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | track-igp-metric;                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br>[edit protocols ldp],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                           |
| <b>Description</b>              | Cause the IGP route metric to be used for the LDP routes instead of the default LDP route metric (the default LDP route metric is 1).                                                                                                                                               |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring LDP to Use the IGP Route Metric on page 4833</a></li> </ul>                                                                                                                                                        |

## traffic-statistics (Protocols LDP)

---

|                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax              | <pre>traffic-statistics {<br/>    file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;<br/>    interval <i>seconds</i>;<br/>    no-penultimate-hop;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Hierarchy Level     | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br>[edit protocols ldp],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Release Information | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Description         | LDP traffic statistics display the amount of traffic passed through a router for a particular FEC.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Options             | <p><b>file <i>filename</i></b>—Name of the file to receive the output of the LDP statistics operation.<br/>Enclose the name within quotation marks. All files are placed in the directory <code>/var/log</code>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of LDP statistics files. When a statistics file named <b><i>ldp-stat</i></b> reaches its maximum size, it is renamed <b><i>ldp-stat.0</i></b>, then <b><i>ldp-stat.1</i></b>, and so on, until the maximum number of LDP statistics files is reached. Then the oldest file is overwritten.</p> <p><b>Range:</b> 2 through 1000</p> <p><b>Default:</b> 2 files</p> <p>If you specify a maximum number of files, you also must include the <b>size</b> statement to specify the maximum file size.</p> <p><b>interval <i>seconds</i></b>—(Optional) Specify the interval at which the statistics are polled and written to the file.</p> <p><b>Default:</b> 300 seconds (5 minutes)</p> <p><b>no-penultimate-hop</b>—(Optional) Do not collect traffic statistics on the penultimate hop router.</p> <p><b>no-world-readable</b>—(Optional) Prevent all users from reading the log file.</p> <p><b>size <i>size</i></b>—(Optional) Maximum size of each statistics file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a statistics file named <b><i>ldp-stat</i></b> reaches this size, it is renamed <b><i>ldp-stat.0</i></b>. When <b><i>ldp-stat</i></b> again reaches this size, <b><i>ldp-stat.0</i></b> is renamed <b><i>ldp-stat.1</i></b> and <b><i>ldp-stat</i></b> is renamed <b><i>ldp-stat.0</i></b>. This renaming scheme continues until the maximum number of statistics files is reached. Then the oldest statistics file is overwritten.</p> <p><b>Syntax:</b> <b><i>xk</i></b> to specify KB, <b><i>xm</i></b> to specify MB, or <b><i>xg</i></b> to specify GB</p> <p><b>Range:</b> 10 KB through the maximum file size supported on your system</p> |



**Default:** 1 MB

If you specify a maximum file size, you also must also include the **files** statement to specify the maximum number of files.

**world-readable**—(Optional) Enable log file access for all users.

|                           |                                                             |
|---------------------------|-------------------------------------------------------------|
| <b>Required Privilege</b> | routing—To view this statement in the configuration.        |
| <b>Level</b>              | routing-control—To add this statement to the configuration. |

|                              |                                                                                                          |
|------------------------------|----------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Collecting LDP Statistics on page 4815</a></li></ul> |
|------------------------------|----------------------------------------------------------------------------------------------------------|

## transport-address

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>transport-address (interface   router-id);</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit logical-systems <i>logical-system-name</i> protocols ldp interface <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br>[edit protocols ldp],<br>[edit protocols ldp interface <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i> ]                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | Enables you to configure the IP address used to specify the TCP session for the LDP session. Routers must first establish a TCP session between one another before they can establish an LDP session. The TCP session enables the routers to exchange the label advertisements needed for the LDP session. To establish the TCP session, each router must learn the other router's transport address. The transport address is an IP address used to identify the TCP session over which the LDP session will run.                                                                                                                                                                                                                                                                                                               |
| <b>Default</b>                  | <b>router-id</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <b>interface</b> —The first IP address on the interface is used as the transport address for any LDP sessions to neighbors that can be reached over that interface. You cannot specify the <b>interface</b> option when there are multiple parallel links to the same LDP neighbor, because the LDP specification requires that the same transport address be advertised on all interfaces to the same neighbor. If LDP detects multiple parallel links to the same neighbor, it disables interfaces to that neighbor one by one until the condition is cleared, either by disconnecting the neighbor on an interface or by specifying the <b>router-id</b> option.<br><br><b>router-id</b> —The router identifier is used as the transport address. Unless otherwise configured, the router identifier is the loopback address. |
| <b>Required Privilege Level</b> | <b>interface</b> —To view this statement in the configuration.<br><b>interface-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Specifying the Transport Address Used by LDP on page 4815</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## Monitoring Commands for LDP

---

- [clear ldp neighbor](#)
- [clear ldp session](#)
- [clear ldp statistics](#)
- [ping mpls ldp](#)

- `show ldp database`
- `show ldp fec-filters`
- `show ldp interface`
- `show ldp neighbor`
- `show ldp path`
- `show ldp route`
- `show ldp session`
- `show ldp statistics`
- `show ldp traffic-statistics`
- `traceroute mpls ldp`

## clear ldp neighbor

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>clear ldp neighbor</code><br><code>&lt;instance <i>instance-name</i>&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code><br><code>&lt;neighbor&gt;</code>                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Tear down Label Distribution Protocol (LDP) neighbor connections.                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><b>none</b>—Tear down connections with all LDP neighbors for all routing instances.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Clear the LDP session for the specified routing instance only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>neighbor</b>—(Optional) Clear an LDP session for the specified neighbor (IP address) only.</p> |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show ldp neighbor on page 4902</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>List of Sample Output</b>    | <a href="#">clear ldp neighbor on page 4884</a>                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                                                             |

## Sample Output

### clear ldp neighbor

```
user@host> clear ldp neighbor
```

## clear ldp session

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | clear ldp session<br><destination><br><instance <i>instance-name</i> ><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Clear Label Distribution Protocol (LDP) sessions.                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <p><b>none</b>—Clear LDP sessions for all destinations for all routing instances.</p> <p><b>destination</b>—(Optional) Clear an LDP session for the specified destination (IP address).</p> <p><b>instance <i>instance-name</i></b>—(Optional) Clear the LDP session for the specified routing instance only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show ldp session on page 4910</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                             |
| <b>List of Sample Output</b>    | <a href="#">clear ldp session on page 4885</a>                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                                                         |

## Sample Output

### clear ldp session

```
user@host> clear ldp session
```

## clear ldp statistics

---

|                                 |                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>clear ldp statistics</code><br><code>&lt;instance <i>instance-name</i>&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>                                                                                                                                                                                          |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                    |
| <b>Description</b>              | Set all Label Distribution Protocol (LDP) statistics to zero.                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <b>none</b> —Set all LDP statistics to zero for all routing instances.<br><br><b>instance <i>instance-name</i></b> —(Optional) Clear the LDP session for the specified routing instance only.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show ldp statistics on page 4916</a></li><li>• <a href="#">show ldp traffic-statistics on page 4920</a></li></ul>                                                                                                                                                                                      |
| <b>List of Sample Output</b>    | <a href="#">clear ldp statistics on page 4886</a>                                                                                                                                                                                                                                                                                                          |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                      |

### Sample Output

#### clear ldp statistics

```
user@host> clear ldp statistics
```

## ping mpls ldp

**Syntax**    ping mpls ldp *fec*  
                  <count *count*>  
                  <destination *address*>  
                  <detail>  
                  <exp *forwarding-class*>  
                  <instance *routing-instance-name*>  
                  <logical-system (all | *logical-system-name*)>  
                  <p2mp root-addr *ip-address* lsp-id *identifier*>  
                  <size *bytes*>  
                  <source *source-address*>  
                  <sweep>

**Release Information**    Command introduced before Junos OS Release 7.4.  
                  Command introduced in Junos OS Release 9.0 for EX Series switches.  
                  **size** and **sweep** options introduced in Junos OS Release 9.6.  
                  **instance** option introduced in Junos OS Release 10.0.  
                  **p2mp**, **root-address**, and **lsp-id** options introduced in Junos OS Release 11.2.  
                  Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

**Description**    Check the operability of MPLS LDP-signaled label-switched path (LSP) connections.  
                  Type Ctrl+c to interrupt a **ping mpls** command.

**Options**    **count** *count*—(Optional) Number of ping requests to send. If **count** is not specified, five ping requests are sent. The range of values is 1 through 1,000,000. The default value is 5.

**destination** *address*—(Optional) Specify an address other than the default (127.0.0.1/32) for the ping echo requests. The address can be anything within the 127/8 subnet.

**detail**—(Optional) Display detailed information about the echo requests sent and received.

**exp** *forwarding-class*—(Optional) Value of the forwarding class for the MPLS ping packets.

**fec**—Ping an LDP-signaled LSP using the forwarding equivalence class (FEC) prefix and length.

**instance** *routing-instance-name*—(Optional) Allows you to ping a combination of the routing instance and forwarding equivalence class (FEC) associated with an LSP.

**logical-system** (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on the specified logical system.

**p2mp root-addr** *ip-address* **lsp-id** *identifier*—(Optional) Ping the end points of a point-to-multipoint LSP. Enter the IP address of the point-to-multipoint LSP root and the ID number of the point-to-multipoint LSP.

**size** *bytes*—(Optional) Size of the LSP ping request packet (88 through 65468 bytes). Packets are 4-byte aligned. For example, If you enter a size of 89, 90, 91, or 92, the router or switch uses a size value of 92 bytes. If you enter a packet size that is smaller

than the minimum size, an error message is displayed reminding you of the 88-byte minimum.

**source *source-address***—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (**lo.0**).

**sweep**—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

**Additional Information** If the LSP changes, the label and interface information displayed when you issued the **ping** command continues to be used. You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the remote router or switch to ping an LSP terminating there. You must configure MPLS even if you intend to ping only LDP forwarding equivalence classes (FECs).

You can configure the ping interval for the **ping mpls ldp** command by specifying a new time in seconds using the **lsp-ping-interval** statement at the **[edit protocols ldp oam]** hierarchy level. For more information, see the *MPLS Applications Feature Guide for Routing Devices*.

In asymmetric MTU scenarios, the echo response may be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

**Required Privilege Level** network

**List of Sample Output** [ping mpls ldp fec count on page 4888](#)  
[ping mpls ldp p2mp root-addr lsp-id on page 4888](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. Packets with error codes are not counted in the received packets count. They are accounted for separately.

## Sample Output

### ping mpls ldp fec count

```
user@host> ping mpls ldp 10.255.245.222 count 10
!!!xxx...x--- lsping statistics ---10 packets transmitted, 3 packets received,
70% packet loss 4 packets received with error status, not counted as received.
```

### ping mpls ldp p2mp root-addr lsp-id

```
user@host> ping mpls ldp p2mp root-addr 10.1.1.1/32 lsp-id 1 count 1
Request for seq 1, to interface 71, no label stack.
Request for seq 1, to interface 70, label 299786
Reply for seq 1, egress 10.1.1.3, return code: Egress-ok, time: 18.936 ms
 Local transmit time: 2009-01-12 03:50:03 PST 407.281 ms
 Remote receive time: 2009-01-12 03:50:03 PST 426.217 ms
```



```
Reply for seq 1, egress 10.1.1.4, return code: Egress-ok, time: 18.936 ms
 Local transmit time: 2009-01-12 03:50:03 PST 407.281 ms
 Remote receive time: 2009-01-12 03:50:03 PST 426.217 ms
Reply for seq 1, egress 10.1.1.5, return code: Egress-ok, time: 18.936 ms
 Local transmit time: 2009-01-12 03:50:03 PST 407.281 ms
 Remote receive time: 2009-01-12 03:50:03 PST 426.217 ms
```

## show ldp database

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show ldp database</code><br><code>&lt;brief   detail   extensive&gt;</code><br><code>&lt;inet   l2circuit&gt;</code><br><code>&lt;instance <i>instance-name</i>&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code><br><code>&lt;p2mp&gt;</code><br><code>&lt;session <i>session</i>&gt;</code><br><code>&lt;summary&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br><b>summary</b> option introduced in Junos OS Release 14.2.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Display entries in the LDP database.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <p><b>none</b>—Display standard information about all entries in the LDP database for all routing instances.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>inet   l2circuit</b>—(Optional) Display only IPv4 or Layer 2 circuit bindings.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display routing instance information for the specified instance only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>p2mp</b>—(Optional) Display point-to-multipoint binding information.</p> <p><b>session <i>session</i></b>—(Optional) Display database for the specified session only. <b><i>session</i></b> is the destination address of the LDP session.</p> <p><b>summary</b>—(Optional)—Display summary output. This option displays the number of labels received and advertised for each LDP session.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>List of Sample Output</b>    | <a href="#">show ldp database (master) on page 4893</a><br><a href="#">show ldp database (standby) on page 4894</a><br><a href="#">show ldp database l2circuit detail on page 4894</a><br><a href="#">show ldp database l2circuit extensive on page 4895</a><br><a href="#">show ldp database p2mp (master) on page 4895</a><br><a href="#">show ldp database p2mp (standby) on page 4895</a><br><a href="#">show ldp database p2mp (master) on page 4896</a><br><a href="#">show ldp database p2mp (standby) on page 4896</a><br><a href="#">show ldp database session on page 4896</a><br><a href="#">show ldp database (Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 4897</a>                                                                                                                                                                                                                                                                         |

[show ldp database \(Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs\) on page 4897](#)  
[show ldp database summary on page 4898](#)

**Output Fields** Table 364 describes the output fields for the **show ldp database** command. Output fields are listed in the approximate order in which they appear.

**Table 364: show ldp database Output Fields**

| Field Name                   | Field Description                                                                | Level of Output |
|------------------------------|----------------------------------------------------------------------------------|-----------------|
| <b>Input label database</b>  | Label received from the other router.                                            | All levels      |
| <b>Output label database</b> | Label advertised to the other router.                                            | All levels      |
| <i>session-identifier</i>    | Session identifier, which includes the local and remote label space identifiers. | All levels      |
| <b>Labels received</b>       | Number of labels received from the other router.                                 | All levels      |
| <b>Labels advertised</b>     | Number of labels advertised to the other router.                                 | All levels.     |
| <b>Label</b>                 | Label binding to a route prefix.                                                 | All levels      |

Table 364: show ldp database Output Fields (*continued*)

| Field Name                             | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Level of Output  |
|----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>Prefix</b>                          | <p>Route prefix.</p> <p>It can be one of the following values:</p> <ul style="list-style-type: none"> <li>• IP prefix.</li> <li>• Point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.</li> <li>• Layer 2 encapsulation type.</li> </ul> <p>Layer 2 encapsulation types are displayed in the format <b>L2CKT control word status encapsulation-type vc-number</b>, for example, <b>L2CKT CtlfWord FRAME RELAY VC 2</b></p> <ul style="list-style-type: none"> <li>• <b>control-word-status</b>—Displays whether the use of the control word has been negotiated for this virtual circuit: <ul style="list-style-type: none"> <li>• <b>NoCtrlWord</b></li> <li>• <b>CtrlWord</b></li> </ul> </li> <li>• <b>encapsulation-type</b>—Encapsulation type: <ul style="list-style-type: none"> <li>• <b>FRAME RELAY</b></li> <li>• <b>ATM AAL5</b></li> <li>• <b>ATM CELL</b></li> <li>• <b>VLAN</b></li> <li>• <b>ETHERNET</b></li> <li>• <b>CISCO_HDLC</b></li> <li>• <b>PPP</b></li> </ul> </li> <li>• <b>VC number</b>—Virtual circuit number. It can have any numeric value.</li> <li>• <b>(Stale)</b>—When you display the LDP database for the neighbor of a restarting router, the bindings learned from the restarting neighbor are displayed as (Stale). Stale bindings are deleted if they are not refreshed within the recovery time.</li> </ul> | All levels       |
| <b>MTU</b>                             | MTU of the Layer 2 circuit. MTU is displayed for all encapsulation types except ATM cell encapsulations.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail</b>    |
| <b>VCCV Control Channel types</b>      | <p>Virtual Circuit Connection Verification (VCCV) control channel types.</p> <ul style="list-style-type: none"> <li>• <b>MPLS router alert label</b></li> <li>• <b>MPLS PW label with TTL=1</b></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>extensive</b> |
| <b>VCCV Control Verification types</b> | The only valid VCCV control verification type is <b>LSP ping</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>extensive</b> |
| <b>TDM payload size</b>                | Size of the Time Division Multiplex (TDM) payload.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | All levels       |
| <b>TDM bitrate</b>                     | Bit rate for the TDM traffic.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | All levels       |
| <b>Requested VLAN ID</b>               | (VLANs) VLAN identifier of the Layer 2 circuit.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail</b>    |
| <b>Cell bundle size</b>                | (ATM cell encapsulations) Maximum number of cells that the Layer 2 circuit can receive in a packet.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail</b>    |

Table 364: show ldp database Output Fields (*continued*)

| Field Name   | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Level of Output |
|--------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>State</b> | State of the label binding: <ul style="list-style-type: none"> <li><b>Active</b>—Label binding has been installed and distributed appropriately. A label binding is almost always in this state.</li> <li><b>New</b>—New label that has not yet been distributed. <ul style="list-style-type: none"> <li><b>MapRcv</b>—Waiting to receive a label mapping message.</li> <li><b>MapSend</b>—Waiting to send a label mapping message.</li> <li><b>RelRcv</b>—Waiting to receive a label release message.</li> <li><b>RelRsnd</b>—Waiting to receive a label release message before resending label mapping message.</li> <li><b>RelSend</b>—Waiting to send a label release message.</li> <li><b>ReqSend</b>—Waiting to send a label request message.</li> <li><b>W/dSend</b>—Waiting to send a label withdrawal message.</li> </ul> </li> </ul> | <b>detail</b>   |
| <b>Age</b>   | Time elapsed since the binding was created.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail</b>   |

## Sample Output

### show ldp database (master)

```

user@host> show ldp database extensive
Input label database, 10.255.107.232:0--10.255.107.236:0
 Label Prefix
 299840 10.255.107.232/32
 State: Active
 Age: 9:35
 Entropy Label Capability: No
 3 10.255.107.236/32
 State: Active
 Age: 9:35
 Entropy Label Capability: No
 299776 L2CKT CtrlWord VLAN VC 100
 MTU: 1500 Requested VLAN ID: 600 Flow Label T Bit: 1 Flow Label R
 Bit: 1
 State: Active
 Age: 9:35
 Entropy Label Capability: No
 VCCV Control Channel types:
 PWE3 control word
 MPLS router alert label
 MPLS PW label with TTL=1
 VCCV Control Verification types:
 LSP ping
 BFD with PW-ACH-encapsulation for Fault Detection
 BFD with IP/UDP-encapsulation for Fault Detection

Output label database, 10.255.107.232:0--10.255.107.236:0
 Label Prefix
 3 10.255.107.232/32
 State: Active
 Age: 9:35
 Entropy Label Capability: No
 299776 10.255.107.236/32

```

State: Active  
Age: 9:35  
Entropy Label Capability: No

### show ldp database (standby)

user@host> show ldp database extensive

Input label database, 10.255.107.236:0--10.255.107.234:0

| Label  | Prefix                 |
|--------|------------------------|
| 299808 | 10.255.107.230/32      |
|        | State: Active          |
|        | Age: 1d 2:46:36        |
|        | Standby binding state: |
|        | Map messages: 1        |
|        | Release messages: 0    |
| Label  | Prefix                 |
| 301136 | 10.255.107.232/32      |
|        | State: Active          |
|        | Age: 1d 2:46:36        |
|        | Standby binding state: |
|        | Map messages: 1        |
|        | Release messages: 0    |
| Label  | Prefix                 |
| 3      | 10.255.107.234/32      |
|        | State: Active          |
|        | Age: 1d 2:46:36        |
|        | Standby binding state: |
|        | Map messages: 1        |
|        | Release messages: 0    |
| Label  | Prefix                 |
| 302480 | 10.255.107.236/32      |
|        | State: Active          |
|        | Age: 1d 2:46:36        |
|        | Standby binding state: |
|        | Map messages: 1        |
|        | Release messages: 0    |

Output label database, 10.255.107.236:0--10.255.107.234:0

| Label  | Prefix                                         |
|--------|------------------------------------------------|
| 299904 | 10.255.107.230/32                              |
|        | State: Active                                  |
|        | Age: 1d 2:46:36                                |
| 299936 | 10.255.107.232/32                              |
|        | State: Active                                  |
|        | Age: 1d 2:46:36                                |
| 299872 | 10.255.107.234/32                              |
|        | State: Active                                  |
|        | Age: 1d 2:46:36                                |
| 3      | 10.255.107.236/32                              |
|        | State: Active                                  |
|        | Age: 1d 2:46:36                                |
| 299952 | P2MP root-addr 10.255.107.230, lsp-id 16777217 |
|        | State: Active                                  |
|        | Age: 1d 2:46:36                                |

### show ldp database l2circuit detail

user@host> show ldp database l2circuit detail

Input label database, 10.255.245.44:0--10.255.245.45:0

| Label | Prefix |
|-------|--------|
|-------|--------|

```

100176 L2CKT CtrlWord ATM CELL (VC Mode) VC 100
 Cell bundle size: 80
 State: Active
 Age: 9:48
100256 L2CKT CtrlWord FRAME RELAY VC 101
 MTU: 4470
 State: Active
 Age: 9:48

Output label database, 10.255.245.44:0--10.255.245.45:0
Label Prefix
100048 L2CKT CtrlWord ATM CELL (VC Mode) VC 100
 Cell bundle size: 80
 State: Active
 Age: 9:48
100112 L2CKT CtrlWord FRAME RELAY VC 101
 MTU: 4470
 State: Active
 Age: 9:48

```

#### show ldp database l2circuit extensive

```

user@host> show ldp database l2circuit extensive
Input label database, 10.255.245.198:0--10.255.245.194:0
Label Prefix
299872 L2CKT CtrlWord PPP VC 100
 MTU: 4470
 VCCV Control Channel types:
 MPLS router alert label
 MPLS PW label with TTL=1
 VCCV Control Verification types:
 LSP ping
Label Prefix
 State: Active
 Age: 19:23:08

```

#### show ldp database p2mp (master)

```

user@host> show ldp database p2mp extensive

Input label database, 10.255.107.232:0--10.255.107.236:0
Label Prefix
569649 P2MP root-addr 10.255.107.232, lsp-id 16777217
 State: Active
 Age: 2d 6:41:46

Output label database, 10.255.107.232:0--10.255.107.236:0

Input label database, 10.255.107.232:0--10.255.107.238:0

Output label database, 10.255.107.232:0--10.255.107.238:0
Label Prefix
299888 P2MP root-addr 10.255.107.230, lsp-id 16777217
 State: Active
 Age: 2d 6:41:35

```

#### show ldp database p2mp (standby)

```

user@host> show ldp database p2mp extensive

Input label database, 10.255.107.236:0--10.255.107.232:0

```

```
Label Prefix
299968 P2MP root-addr 10.255.107.230, lsp-id 16777217
 State: Active
 Age: 4d 22:21:57
 Standby binding state:
 Map messages: 1
 Release messages: 0

Output label database, 10.255.107.236:0--10.255.107.232:0
Label Prefix
3 P2MP root-addr 10.255.107.232, lsp-id 1
 State: Active
 Age: 4d 22:21:57
```

### show ldp database p2mp (master)

```
user@host> show ldp database p2mp extensive
```

```
Input label database, 10.255.107.232:0--10.255.107.236:0
Label Prefix
569649 P2MP root-addr 10.255.107.232, lsp-id 16777217
 State: Active
 Age: 2d 6:41:46

Output label database, 10.255.107.232:0--10.255.107.236:0

Input label database, 10.255.107.232:0--10.255.107.238:0

Output label database, 10.255.107.232:0--10.255.107.238:0
Label Prefix
299888 P2MP root-addr 10.255.107.230, lsp-id 16777217
 State: Active
 Age: 2d 6:41:35
```

### show ldp database p2mp (standby)

```
user@host> show ldp database p2mp extensive
```

```
Input label database, 10.255.107.236:0--10.255.107.232:0
Label Prefix
299968 P2MP root-addr 10.255.107.230, lsp-id 16777217
 State: Active
 Age: 4d 22:21:57
 Standby binding state:
 Map messages: 1
 Release messages: 0

Output label database, 10.255.107.236:0--10.255.107.232:0
Label Prefix
3 P2MP root-addr 10.255.107.232, lsp-id 1
 State: Active
 Age: 4d 22:21:57
```

### show ldp database session

```
user@host> show ldp database session 10.1.1.195
Input label database, 10.0.0.194:0--10.1.1.195:0
Label Prefix
100002 10.255.245.197/32
100003 10.255.245.196/32
100004 10.0.0.194/32
```



```

 3 10.1.1.195/32
100000 L2CKT NoCtrlWord FRAME RELAY VC 1
100001 L2CKT CtrlWord FRAME RELAY VC 2
Output label database, 10.0.0.194:0--10.1.1.195:0
 Label Prefix
100003 10.255.245.197/32
100004 10.1.1.195/32
100002 10.255.245.196/32
 3 10.0.0.194/32
100000 L2CKT CtrlWord FRAME RELAY VC 2
100001 L2CKT NoCtrlWord FRAME RELAY VC 1

```

#### show ldp database (Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```

user@host> show ldp database
Input label database, 1.1.1.2:0--1.1.1.3:0
 Label Prefix
299808 1.1.1.2/32
 3 1.1.1.3/32
299792 1.1.1.6/32
299776 10.255.2.227/32
299840 P2MP root-addr 1.1.1.2, grp: 232.2.2.2, src: 1.2.7.7
299824 P2MP root-addr 1.1.1.2, grp: 232.1.1.2, src: 192.168.219.11

Output label database, 1.1.1.2:0--1.1.1.3:0
 Label Prefix
 3 1.1.1.2/32
299776 1.1.1.3/32
299808 1.1.1.6/32
299792 10.255.2.227/32

Input label database, 1.1.1.2:0--1.1.1.6:0
 Label Prefix
299856 1.1.1.2/32
299792 1.1.1.3/32
 3 1.1.1.6/32
299776 10.255.2.227/32
299888 P2MP root-addr 1.1.1.2, grp: 232.2.2.2, src: 1.2.7.7
299808 P2MP root-addr 1.1.1.2, grp: 232.1.1.1, src: 192.168.219.11
299824 P2MP root-addr 1.1.1.2, grp: 232.1.1.2, src: 192.168.219.11
299840 P2MP root-addr 1.1.1.2, grp: 232.1.1.3, src: 192.168.219.11
299872 P2MP root-addr 1.1.1.2, grp: ff3e::1:2, src: abcd::1:2:7:7

Output label database, 1.1.1.2:0--1.1.1.6:0
 Label Prefix
 3 1.1.1.2/32
299776 1.1.1.3/32
299808 1.1.1.6/32
299792 10.255.2.227/32

```

#### show ldp database (Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```

user@host> show ldp database
Input label database, 10.255.2.227:0--1.1.1.3:0
 Label Prefix
299808 1.1.1.2/32
 3 1.1.1.3/32
299792 1.1.1.6/32
299776 10.255.2.227/32

Output label database, 10.255.2.227:0--1.1.1.3:0
 Label Prefix

```

```
299856 1.1.1.2/32
299776 1.1.1.3/32
299792 1.1.1.6/32
3 10.255.2.227/32
```

Input label database, 10.255.2.227:0--1.1.1.6:0

```
Label Prefix
299856 1.1.1.2/32
299776 1.1.1.3/32
3 1.1.1.6/32
299776 10.255.2.227/32
```

Output label database, 10.255.2.227:0--1.1.1.6:0

```
Label Prefix
299856 1.1.1.2/32
299776 1.1.1.3/32
299792 1.1.1.6/32
3 10.255.2.227/32
299888 P2MP root-addr 1.1.1.2, grp: 232.2.2.2, src: 1.2.7.7
299808 P2MP root-addr 1.1.1.2, grp: 232.1.1.1, src: 192.168.219.11
299824 P2MP root-addr 1.1.1.2, grp: 232.1.1.2, src: 192.168.219.11
299840 P2MP root-addr 1.1.1.2, grp: 232.1.1.3, src: 192.168.219.11
299872 P2MP root-addr 1.1.1.2, grp: ff3e::1:2, src: abcd::1:2:7:7
```

#### show ldp database summary

```
user@host> show ldp database summary
```

| Session ID                 | Labels received | Labels advertised |
|----------------------------|-----------------|-------------------|
| 10.255.0.1:0--10.255.0.2:0 | 4               | 4                 |
| 10.255.0.1:0--10.255.0.3:0 | 4               | 4                 |

## show ldp fec-filters

|                                 |                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show ldp fec-filters<br><fec><br><instance <i>instance-name</i> ><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                          |
| <b>Description</b>              | Display information about configured Label Distribution Protocol (LDP) forwarding equivalence class (FEC) filters.                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <p><b>fec</b>—(Optional) Display FEC filter information for the specified FEC.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display FEC filter information for the specified instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                             |
| <b>List of Sample Output</b>    | <a href="#">show ldp fec-filters on page 4899</a>                                                                                                                                                                                                                                                                                                                |
| <b>Output Fields</b>            | Table 365 lists the output fields for the <b>show ldp fec-filters</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                           |

**Table 365: show ldp fec-filters Output Fields**

| Field Name | Field Description                                |
|------------|--------------------------------------------------|
| Ingress    | Names of the FEC filters on the ingress routers. |
| Transit    | Names of the FEC filters on the transit routers. |

## Sample Output

### show ldp fec-filters

```

user@host> show ldp fec-filters 10/8
10.22.1.2/32
 Ingress: f1-10.22.1.2/32 (index: 3)
 Transit: (null) (index: 0)

```

## show ldp interface

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show ldp interface &lt;brief   detail   extensive&gt; &lt;interface-name&gt; &lt;instance instance-name&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | <p>Command introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Display the status of Label Distribution Protocol (LDP)-enabled interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <p><b>none</b>—Display standard status information about all LDP-enabled interface for all routing instances.</p> <p><b>interface-name</b>—(Optional) Display information for the specified interface.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>instance instance-name</b>—(Optional) Display information for the specified routing instance.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>List of Sample Output</b>    | <a href="#">show ldp interface extensive on page 4901</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Output Fields</b>            | Table 366 describes the output fields for the <b>show ldp interface</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                           |

Table 366: show ldp interface Output Fields

| Field Name            | Field Description                                                                                                                                                              | Level of Output                   |
|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------|
| <b>Interface</b>      | Interface name.                                                                                                                                                                | All levels                        |
| <b>Label space ID</b> | Label space identifier that the router is advertising on the interface.                                                                                                        | All levels                        |
| <b>Nbr count</b>      | Number of neighbors on the interface.                                                                                                                                          | All levels                        |
| <b>Next hello</b>     | How long until the next hello packet is sent on this interface, in seconds.                                                                                                    | All levels                        |
| <b>Hello interval</b> | One-third of the negotiated hold time (in seconds). If the user-configured value for the hello interval is smaller than the computed value, the user-configured value is used. | <b>detail</b><br><b>extensive</b> |
| <b>Hold time</b>      | Configured hold time, in seconds.                                                                                                                                              | <b>detail</b><br><b>extensive</b> |

Table 366: show ldp interface Output Fields (*continued*)

| Field Name           | Field Description                                                                 | Level of Output |
|----------------------|-----------------------------------------------------------------------------------|-----------------|
| Transport address    | Address to which the neighbor wants the local route to establish the LDP session. | extensive       |
| Local hello interval | Locally configured hello interval.                                                | extensive       |

## Sample Output

### show ldp interface extensive

```
user@host> show ldp interface extensive
Interface Label space ID Nbr count Next hello
fe-0/0/3.0 10.255.245.6:0 2 0
Hello interval: 1, Hold time: 15, Transport address: 10.255.245.6
Local hello interval: 2, Index: 69
```

## show ldp neighbor

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show ldp neighbor &lt;brief   detail   extensive&gt; &lt;auto-targeted&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;neighbor-address&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | <p>Command introduced before Junos OS Release 7.4.</p> <p><b>neighbor-address</b> option added in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p> <p><b>auto-targeted</b> option added in Junos OS Release 14.2.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | Display Label Distribution Protocol (LDP) neighbor information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <p><b>none</b>—Display standard information about LDP neighbors for all routing instances.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>auto-targeted</b>—(Optional) Display information about LDP neighbors that are automatically targeted using the loopback addresses.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information for the specified routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>neighbor-address</b>—(Optional) Display information about the specified LDP neighbor.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">clear ldp neighbor on page 4884</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>List of Sample Output</b>    | <p><a href="#">show ldp neighbor extensive on page 4903</a></p> <p><a href="#">show ldp neighbor auto-targeted extensive on page 4903</a></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Output Fields</b>            | Table 367 describes the output fields for the <b>show ldp neighbor</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

**Table 367: show ldp neighbor Output Fields**

| Field Name            | Field Description                                  | Level of Output |
|-----------------------|----------------------------------------------------|-----------------|
| <b>Address</b>        | IP address of the neighbor.                        | All levels      |
| <b>Interface</b>      | Interface over which the neighbor was discovered.  | All levels      |
| <b>Label space ID</b> | Label space identifier advertised by the neighbor. | All levels      |

Table 367: show ldp neighbor Output Fields (*continued*)

| Field Name             | Field Description                                                                                           | Level of Output  |
|------------------------|-------------------------------------------------------------------------------------------------------------|------------------|
| Hold time              | Remaining hold time before the neighbor expires, in seconds.                                                | All levels       |
| Transport address      | Address to which the neighbor wants the local route to establish the LDP session.                           | detail           |
| Configuration sequence | Counter that increments whenever the neighbor changes its configuration.                                    | detail           |
| Up for                 | Length of time the LDP neighbor has been in operation.                                                      | detail extensive |
| Reference count        | Reference count for the LDP neighbor.                                                                       | extensive        |
| Hold time              | Displays the neighbor's hold time. The hold time is the proposed hold times for the local and peer routers. | extensive        |
| Proposed local/peer    | Hold time value proposed by the local router and the peer router.                                           | extensive        |

## Sample Output

### show ldp neighbor extensive

```

user@host> show ldp neighbor extensive
Address Interface Label space ID Hold Time
192.168.37.23 so-1/0/0.0 10.255.245.5:0 44
Transport address: 10.255.245.5, Configuration sequence: 6
Up for 00:03:37
Reference count: 1
Hold time: 45, Proposed local/peer: 15/45

```

### show ldp neighbor auto-targeted extensive

```

user@host> show ldp neighbor auto-targeted extensive
Address Interface Label space ID Hold time
10.255.107.236 lo0.0 10.255.107.236:0 41
Transport address: 10.255.107.236, Configuration sequence: 14
Up for 00:10:53
Reference count: 2
Hold time: 45, Proposed local/peer: 45/45
Hello interval: 15
Hello flags: targeted
Neighbor types: Auto-targeted

```

## show ldp path

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show ldp path &lt;brief   detail   extensive&gt; &lt;destination&gt; &lt;instance instance-name&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | <p>Command introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Display Label Distribution Protocol (LDP) label-switched paths (LSPs).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <p><b>none</b>—Display standard information about all LDP LSPs for all routing instances.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>destination</b>—(Optional) Restrict the output to entries that match the specified destination prefix.</p> <p><b>instance instance-name</b>—(Optional) Display information for the specified routing instance only.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>List of Sample Output</b>    | <a href="#">show ldp path extensive on page 4905</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Output Fields</b>            | Table 368 describes the output fields for the <b>show ldp path</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                         |

**Table 368: show ldp path Output Fields**

| Field Name                    | Field Description                                                                                            |
|-------------------------------|--------------------------------------------------------------------------------------------------------------|
| <b>Output Session (label)</b> | Session ID and labels that this system has sent using LDP. These correspond to MPLS packets received.        |
| <b>Input Session (label)</b>  | Session ID and labels that this system has received using LDP. These correspond to MPLS packets transmitted. |
| <b>route</b>                  | MPLS route.                                                                                                  |
| <b>Attached route</b>         | Route corresponding to the LSP.                                                                              |
| <b>Ingress route</b>          | The router acts as the ingress for the LSP.                                                                  |
| <b>Reference count</b>        | Reference count for the LDP neighbor.                                                                        |



Table 368: show ldp path Output Fields (*continued*)

| Field Name    | Field Description                                                               |
|---------------|---------------------------------------------------------------------------------|
| Transit route | Names of the forwarding equivalence class (FEC) filters on the transit routers. |
| Global label  | MPLS label that is used globally.                                               |

## Sample Output

### show ldp path extensive

```
user@host> show ldp path extensive
Output Session (label) Input Session (label)
10.255.14.220:0(3) ()
 Attached route: 10.255.14.221/32
 Reference count: 3, Global label: 3
10.255.14.220:0(100000) 10.255.14.220:0(3)
 Attached route: 10.255.14.220/32, Ingress route
 Reference count: 2, Transit route, Global label: 100000
10.255.14.220:0(100001) 10.255.14.220:0(100001)
 Attached route: 10.255.14.214/32, Ingress route
 Reference count: 2, Transit route, Global label: 100001
```

## show ldp route

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show ldp route &lt;brief   detail   extensive&gt; &lt;destination&gt; &lt;instance instance-name&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | <p>Command introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Display the entries in the Label Distribution Protocol (LDP) internal topology table. The internal topology table contains routes from inet.0 and inet.3 and is used when binding a label to a forwarding equivalence class (FEC).                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <p><b>none</b>—Display standard information about all entries in the LDP internal topology table for all routing instances.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>destination</b>—(Optional) Restrict the output to entries that are longer than the specified destination prefix and prefix length.</p> <p><b>instance instance-name</b>—(Optional) Display entries for the specified routing instance only.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>List of Sample Output</b>    | <p><a href="#">show ldp route detail on page 4908</a></p> <p><a href="#">show ldp route extensive on page 4908</a></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Output Fields</b>            | <a href="#">Table 369</a> describes the output fields for the <b>show ldp route</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

**Table 369: show ldp route Output Fields**

| Field Name                     | Field Description                                         |
|--------------------------------|-----------------------------------------------------------|
| <b>Destination</b>             | Destination prefix.                                       |
| <b>Next-hop intf/lsp/table</b> | Interface that is the next hop to the destination prefix. |
| <b>Next-hop address</b>        | IP address of the next hop.                               |
| <b>Session ID</b>              | LDP session ID.                                           |

Table 369: show ldp route Output Fields (*continued*)

| Field Name              | Field Description                                                                                                                                               |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Route flags             | Information about the route. For example, the <b>Ingress TTL propagate</b> flag indicates that the time-to-live (TTL) value is being propagated with the route. |
| Bound to outgoing label | The route has been bound to LSPs with the label being distributed for that LSP.                                                                                 |
| Topology entry          | The topology that the route is bound to.                                                                                                                        |
| Ingress route status    | Status of the ingress route. For example, it could be <b>Active</b> or <b>Inactive</b> .                                                                        |
| Last modified           | The length of time since the ingress route status last changed.                                                                                                 |

## Sample Output

### show ldap route detail

```

user@host> show ldap route 10.255.8.5 detail
Destination Next-hop intf/lsp Next-hop address
10.255.8.5/32 f1
 Session ID 10.255.170.84:0--10.255.170.92:0
 fe-0/0/0.0 192.168.100.2
 Session ID 10.255.170.84:0--10.255.8.5:0
 so-0/2/1.0
 Session ID 10.255.170.84:0--10.255.8.5:0
 so-0/2/2.0
 Session ID 10.255.170.84:0--10.255.8.3:0
 Bound to outgoing label 299776, Topology entry: 0x8c38a80
 BFD dest addr BFD state LSP-ping Next-hop addr Next-hop intf/lsp
127.0.0.64 up up 192.168.100.2 fe-0/0/0.0
127.0.1.64 up up so-0/2/1.0
127.0.2.64 up up so-0/2/2.0
127.0.3.64 up up f1
.....

```

### show ldap route extensive

```

user@host> show ldap route extensive

Destination Next-hop intf/lsp/table Next-hop address
10.0.0.0/30 ge-1/2/0.18 10.0.0.17
 Session ID 192.168.0.6:0--192.168.0.5:0
 Route flags: None
Destination Next-hop intf/lsp/table Next-hop address
10.0.0.4/30 ge-1/2/0.18 10.0.0.17
 Session ID 192.168.0.6:0--192.168.0.5:0
 Route flags: None
Destination Next-hop intf/lsp/table Next-hop address
10.0.0.8/30 ge-1/2/1.21 10.0.0.22
 Session ID 192.168.0.6:0--192.168.0.4:0
 Route flags: None
Destination Next-hop intf/lsp/table Next-hop address
10.0.0.12/30 ge-1/2/1.21 10.0.0.22
 Session ID 192.168.0.6:0--192.168.0.4:0
 Route flags: None
Destination Next-hop intf/lsp/table Next-hop address
10.0.0.16/30 ge-1/2/0.18
 Route flags: None
Destination Next-hop intf/lsp/table Next-hop address
10.0.0.18/32 ge-1/2/0.18
 Route flags: None
Destination Next-hop intf/lsp/table Next-hop address
10.0.0.20/30 ge-1/2/1.21
 Route flags: None
Destination Next-hop intf/lsp/table Next-hop address
10.0.0.21/32 ge-1/2/1.21
 Route flags: None
Destination Next-hop intf/lsp/table Next-hop address
192.168.0.1/32 ge-1/2/0.18 10.0.0.17
 Session ID 192.168.0.6:0--192.168.0.5:0
 Route flags: None
Destination Next-hop intf/lsp/table Next-hop address
192.168.0.2/32 ge-1/2/1.21 10.0.0.22
 Session ID 192.168.0.6:0--192.168.0.4:0

```

```

 ge-1/2/0.18 10.0.0.17
 Session ID 192.168.0.6:0--192.168.0.5:0
 Route flags: None
Destination Next-hop intf/lsp/table Next-hop address
192.168.0.3/32 ge-1/2/1.21 10.0.0.22
 Session ID 192.168.0.6:0--192.168.0.4:0
 Route flags: None
Destination Next-hop intf/lsp/table Next-hop address
192.168.0.4/32 ge-1/2/1.21 10.0.0.22
 Session ID 192.168.0.6:0--192.168.0.4:0
 Bound to outgoing label 299808, Topology entry: 0x92a483c
 Ingress route status: Active, Last modified: 00:01:19 ago
 Route flags: Ingress TTL propagate, Transit TTL propagate
Destination Next-hop intf/lsp/table Next-hop address
192.168.0.5/32 ge-1/2/0.18 10.0.0.17
 Session ID 192.168.0.6:0--192.168.0.5:0
 Bound to outgoing label 299792, Topology entry: 0x92a47f8
 Ingress route status: Active, Last modified: 00:01:19 ago
 Route flags: Ingress TTL propagate, Transit TTL propagate
Destination Next-hop intf/lsp/table Next-hop address
192.168.0.6/32 lo0.6
 Bound to outgoing label 3, Topology entry: 0x92a4a5c
 Ingress route status: Inactive
 Route type: Egress route
 Route flags: None
Destination Next-hop intf/lsp/table Next-hop address
10.10.20.1/32 fe-1/0/0.0 192.168.199.37
 LSP LDP->10.255.107.230

```

## show ldp session

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show ldp session &lt;brief   detail   extensive&gt; &lt;auto-targeted&gt; &lt;destination&gt; &lt;instance instance-name&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | <p>Command introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p> <p><b>auto-targeted</b> option added in Junos OS Release 14.2.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | Display information about Label Distribution Protocol (LDP) sessions.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <p><b>none</b>—Display standard information about all LDP sessions for all routing instances.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>auto-targeted</b>—(Optional) Display information about LDP sessions that are automatically targeted using loopback addresses.</p> <p><b>destination</b>—(Optional) Restrict LDP session display to the specified address.</p> <p><b>instance instance-name</b>—(Optional) Display routing instance information for the specified instance. If <b>instance-name</b> is omitted, information is displayed for the master instance.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">clear ldp session on page 4885</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>List of Sample Output</b>    | <p><a href="#">show ldp session brief on page 4914</a></p> <p><a href="#">show ldp session detail on page 4914</a></p> <p><a href="#">show ldp session extensive on page 4914</a></p> <p><a href="#">show ldp session auto-targeted detail on page 4915</a></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Output Fields</b>            | Table 370 describes the output fields for the <b>show ldp session</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

Table 370: show ldp session Output Fields

| Field Name | Field Description                                                                                                                                                                                                                                                                           | Level of Output |
|------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Address    | Transport address of the session.                                                                                                                                                                                                                                                           | any             |
| State      | State of the session: <b>Nonexistent</b> , <b>Connecting</b> , <b>Initialized</b> , <b>OpenRec</b> , <b>OpenSent</b> , <b>Operational</b> , or <b>Closing</b> . The states correspond to the state diagram specified in Internet Draft LDP Specification draft-ietf-mpls-rfc3036bis-01.txt. | any             |

Table 370: show ldp session Output Fields (*continued*)

| Field Name                    | Field Description                                                                                                                                                                                                     | Level of Output         |
|-------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Connection</b>             | TCP connection state: <b>Closed</b> , <b>Opening</b> , or <b>Open</b> .                                                                                                                                               | any                     |
| <b>Hold time</b>              | Time remaining until the session will be closed, in seconds.                                                                                                                                                          | any                     |
| <b>Session ID</b>             | LDP identifiers of the peers of this session.                                                                                                                                                                         | detail extensive        |
| <b>Next keepalive</b>         | Time until next keepalive is sent, in seconds.                                                                                                                                                                        | <b>detail extensive</b> |
| <b>Active</b>                 | Whether the local router is playing the active role in the session and during session establishment.                                                                                                                  | <b>detail extensive</b> |
| <b>Passive</b>                | Whether the local router is playing the passive role in the session and during session establishment.                                                                                                                 | <b>detail extensive</b> |
| <b>Maximum PDU</b>            | Maximum protocol data unit (PDU) size (packet size) for the session.                                                                                                                                                  | <b>detail extensive</b> |
| <b>Hold time</b>              | Time remaining until the session will be closed, in seconds. This value corresponds to the one configured using the <b>keepalive-timeout</b> statement configured at the <b>[edit protocols ldp]</b> hierarchy level. | <b>detail extensive</b> |
| <b>Neighbor count</b>         | Number of neighbors that are contributing to the session.                                                                                                                                                             | <b>detail extensive</b> |
| <b>Neighbor types</b>         | Category of LDP session: <b>discovered</b> or <b>auto-targeted</b> .                                                                                                                                                  | any                     |
| <b>Keepalive interval</b>     | Keepalive interval, in seconds.                                                                                                                                                                                       | <b>detail extensive</b> |
| <b>Connect retry interval</b> | TCP connection retry interval, in seconds.                                                                                                                                                                            | <b>detail extensive</b> |
| <b>Local address</b>          | Local transport address.                                                                                                                                                                                              | <b>detail extensive</b> |
| <b>Remote address</b>         | Remote transport address.                                                                                                                                                                                             | <b>detail extensive</b> |
| <b>Up for</b>                 | Time that this session has been up.                                                                                                                                                                                   | <b>detail extensive</b> |
| <b>Last down</b>              | Time since the session last went down.                                                                                                                                                                                | <b>detail extensive</b> |

Table 370: show ldp session Output Fields (*continued*)

| Field Name                                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Level of Output         |
|--------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Reason</b>                              | Reason the session went down: <ul style="list-style-type: none"> <li>• Aborted graceful restart</li> <li>• Authentication key was changed</li> <li>• Bad type length value (TLV)</li> <li>• Bad protocol data unit (PDU) packets</li> <li>• Command-line interface (CLI) command</li> <li>• Connect time expired</li> <li>• Connection error</li> <li>• Connection reset</li> <li>• Error during initialization</li> <li>• Hold time expired</li> <li>• No adjacency or all adjacencies down</li> <li>• Notification received</li> <li>• Received notification from peer</li> <li>• Unexpected End of File (EOF)</li> <li>• Unknown reason</li> </ul>                                                                                                                                                                                                                                                                                                         | <b>detail extensive</b> |
| <b>Number of session flaps</b>             | Number of times the session changes from up to down.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>detail extensive</b> |
| <b>Restarting</b>                          | LDP is in the process of gracefully restarting.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive</b> |
| <b>Capabilities advertised</b>             | LDP capabilities advertised to a peer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail extensive</b> |
| <b>Capabilities received</b>               | LDP capabilities received from a peer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail extensive</b> |
| <b>Protection</b>                          | Information about the status of MPLS LDP session protection.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail extensive</b> |
| <b>restart complete in <i>nnn</i> msec</b> | Amount of time (in milliseconds) remaining until graceful restart is declared complete.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail extensive</b> |
| <b>Local</b>                               | Information about graceful restart for the local end of an LDP session. Graceful restart and helper mode are independent. <ul style="list-style-type: none"> <li>• <b>Restart</b>—Status of the graceful restart feature at the local end of the LDP session: <b>enabled</b> or <b>disabled</b>.</li> <li>• <b>Helper mode</b>—Status of the helper mode feature at the local end of the LDP session: <b>enabled</b> or <b>disabled</b>. When this feature is enabled, the local end of the LDP session can help the restarting router with its LDP restart procedures.</li> <li>• <b>Reconnect time</b>—Amount of time to wait from when a restart is initiated until the router can exchange LDP messages with its neighbors. The default is <b>60000 msec</b> and is not configurable. (<b>Reconnect timeout</b> refers to "FT Reconnect timeout" in draft-ietf-mpls-ldp-restart-06, <i>Internet Draft Graceful Restart Mechanism for LDP</i>.)</li> </ul> | <b>detail extensive</b> |



Table 370: show ldp session Output Fields (*continued*)

| Field Name                         | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Level of Output         |
|------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Remote</b>                      | Information about graceful restart at the remote end of an LDP session. Graceful restart and helper mode are independent. <ul style="list-style-type: none"> <li>• <b>Restart</b>—Status of the graceful restart feature at the remote end of the LDP session: <b>enabled</b> or <b>disabled</b>.</li> <li>• <b>Helper mode</b>—Status of the helper mode feature at the remote end of the LDP session: <b>enabled</b> or <b>disabled</b>. When this feature is enabled, the remote end of the LDP session can help the restarting router with its LDP restart procedures.</li> <li>• <b>Reconnect time</b>—Amount of time in milliseconds from when a restart is initiated until the remote router can exchange LDP messages with its neighbors.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b> |
| <b>Local maximum recovery time</b> | Amount of time during which the restarting node attempts to recover its lost states with help from its neighbors (in milliseconds).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>detail extensive</b> |
| <b>Next-hop addresses received</b> | Next-hop addresses received on the session.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail extensive</b> |
| <b>Queue depth</b>                 | Number of messages that are queued for sending to the peers in the group.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>extensive</b>        |
| <b>Message type</b>                | Type of message being sent: <ul style="list-style-type: none"> <li>• <b>Initialization</b>—Session initialization negotiation messages sent by an LSR to an LDP peer when the transport connection is established.</li> <li>• <b>Keepalive</b>—Keepalive timer messages sent by an LSR to an LDP peer to keep the session active when there is no information or PDU exchanged between them.</li> <li>• <b>Notification</b>—Notification messages (such as state of the LDP session) or error information (such as bad PDU length) sent by an LSR to an LDP peer.</li> <li>• <b>Address</b>—Message sent by an LSR to an LDP peer to advertise interface addresses.</li> <li>• <b>Address withdraw</b>—Message sent by an LSR to an LDP peer to withdraw a previously advertised interface address.</li> <li>• <b>Label mapping</b>—Message sent by an LSR to an LDP peer to advertise label mapping for a forwarding equivalence class (FEC).</li> <li>• <b>Label request</b>—Message sent by an LSR to an LDP peer to request a label mapping for an FEC.</li> <li>• <b>Label withdraw</b>—Message sent by an LSR to an LDP peer to withdraw a previously advertised FEC-label mapping.</li> <li>• <b>Label release</b>—Message sent by an LSR to an LDP peer to notify the peer that a specific FEC-label mapping has been released.</li> <li>• <b>Label abort</b>—Message sent by an LSR to an LDP peer to abort a label request message.</li> <li>• <b>Total</b>—Messages sent and received during the lifetime of the session.</li> <li>• <b>Last 5 seconds</b>—Messages sent and received during the current session.</li> </ul> | <b>extensive</b>        |

## Sample Output

### show ldp session brief

```
user@host> show ldp session brief
 Address State Connection Hold time
10.255.72.160 Operational Open 21
10.255.72.164 Operational Open 20
10.255.72.172 Operational Open 21
```

### show ldp session detail

```
user@host> show ldp session detail
Address: 192.168.0.3, State: Operational, Connection: Open, Hold time: 27
Session ID: 192.168.0.2:0--192.168.0.3:0
Next keepalive in 7 seconds
Passive, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
Neighbor types: discovered
Keepalive interval: 10, Connect retry interval: 1
Local address: 192.168.0.2, Remote address: 192.168.0.3
Up for 00:00:02
Capabilities advertised: none
Capabilities received: none
Protection: disabled
Local - Restart: enabled, Helper mode: enabled, Reconnect time: 60000
Remote - Restart: enabled, Helper mode: enabled, Reconnect time: 60000
Local maximum neighbor reconnect time: 120000 msec
Local maximum neighbor recovery time: 240000 msec
Local Label Advertisement mode: Downstream unsolicited
Remote Label Advertisement mode: Downstream unsolicited
Negotiated Label Advertisement mode: Downstream unsolicited
Nonstop routing state: Not in sync
Next-hop addresses received:
 10.0.0.5
 10.0.0.33
```

### show ldp session extensive

```
user@host> show ldp session extensive
Address: 192.168.0.3, State: Operational, Connection: Open, Hold time: 22
Session ID: 192.168.0.2:0--192.168.0.3:0
Next keepalive in 2 seconds
Passive, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
Neighbor types: discovered
Keepalive interval: 10, Connect retry interval: 1
Local address: 192.168.0.2, Remote address: 192.168.0.3
Up for 00:05:37
Capabilities advertised: none
Capabilities received: none
Protection: disabled
Local - Restart: enabled, Helper mode: enabled, Reconnect time: 60000
Remote - Restart: enabled, Helper mode: enabled, Reconnect time: 60000
Local maximum neighbor reconnect time: 120000 msec
Local maximum neighbor recovery time: 240000 msec
Local Label Advertisement mode: Downstream unsolicited
Remote Label Advertisement mode: Downstream unsolicited
Negotiated Label Advertisement mode: Downstream unsolicited
Nonstop routing state: Not in sync
Next-hop addresses received:
 10.0.0.5
 10.0.0.33
```

Queue depth: 0

| Message type     | Total |          | Last 5 seconds |          |
|------------------|-------|----------|----------------|----------|
|                  | Sent  | Received | Sent           | Received |
| Initialization   | 1     | 1        | 0              | 0        |
| Keepalive        | 33    | 33       | 1              | 1        |
| Notification     | 0     | 0        | 0              | 0        |
| Address          | 1     | 1        | 0              | 0        |
| Address withdraw | 0     | 0        | 0              | 0        |
| Label mapping    | 7     | 5        | 0              | 0        |
| Label request    | 0     | 0        | 0              | 0        |
| Label withdraw   | 3     | 1        | 0              | 0        |
| Label release    | 1     | 3        | 0              | 0        |
| Label abort      | 0     | 0        | 0              | 0        |

### show ldp session auto-targeted detail

```

user@host> show ldp session auto-generated detail
Address: 192.168.1.5, State: Operational, Connection: Open, Hold time: 25
Session ID: 192.168.1.1:0--192.168.1.5:0
Next keepalive in 5 seconds
Passive, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
Neighbor types: discovered, Auto-targeted
 ^^^^^^^^^^^^^^^^^
Keepalive interval: 10, Connect retry interval: 1
Local address: 192.168.1.1, Remote address: 192.168.1.5
Up for 00:00:34
Capabilities advertised: none
Capabilities received: none
Protection: disabled
Local - Restart: disabled, Helper mode: enabled
Remote - Restart: disabled, Helper mode: enabled
Local maximum neighbor reconnect time: 120000 msec
Local maximum neighbor recovery time: 240000 msec
Local Label Advertisement mode: Downstream unsolicited
Remote Label Advertisement mode: Downstream unsolicited
Negotiated Label Advertisement mode: Downstream unsolicited
Nonstop routing state: Not in sync
Next-hop addresses received:
 192.168.1.2
 192.168.1.3

```

## show ldp statistics

|                                 |                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show ldp statistics<br><instance <i>instance-name</i> ><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                              |
| <b>Description</b>              | Display Label Distribution Protocol (LDP) statistics.                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <p><b>none</b>—Display LDP statistics for all routing instances.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information for the specified routing instance only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">clear ldp statistics on page 4886</a></li> </ul>                                                                                                                                                                                                                                                  |
| <b>List of Sample Output</b>    | <a href="#">show ldp statistics on page 4919</a>                                                                                                                                                                                                                                                                                                     |
| <b>Output Fields</b>            | <a href="#">Table 371</a> lists the output fields for the <b>show ldp statistics</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                |

**Table 371: show ldp statistics Output Fields**

| Field Name                    | Field Description                                                    |
|-------------------------------|----------------------------------------------------------------------|
| Total Sent, Received          | Total number of each message type sent and received.                 |
| Last 5 seconds Sent, Received | Number of each message type sent and received in the last 5 seconds. |

Table 371: show ldp statistics Output Fields (*continued*)

| Field Name          | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Message type</b> | <p>LDP message types:</p> <ul style="list-style-type: none"> <li>• <b>Hello</b>—Messages that enable LDP nodes to discover one another and to detect the failure of a neighbor or of the link to the neighbor.</li> <li>• <b>Initialization</b>—Messages that indicate an LDP session has started.</li> <li>• <b>Keepalive</b>—Messages that ensure that the keepalive timeout is not exceeded.</li> <li>• <b>Notification</b>—Advisory information and signal error information.</li> <li>• <b>Address</b>—Messages with address information.</li> <li>• <b>Address withdrawal</b>—Messages regarding address withdrawal.</li> <li>• <b>Label mapping</b>—Messages with label mapping information.</li> <li>• <b>Label request</b>—Request for a label mapping from a neighboring router.</li> <li>• <b>Label withdrawal</b>—Withdrawal message sent by the downstream LSR to recall a label that it previously mapped. If an LSR that has received a label mapping subsequently determines that it no longer needs that label, it can send a label release message that frees the label for use.</li> <li>• <b>Label release</b>—Message sent by the downstream LSR to recall a label that it previously mapped. If an LSR that has received a label mapping subsequently determines that it no longer needs that label, it can send a label release message that frees the label for use.</li> <li>• <b>Label abort</b>—Messages about label interruptions.</li> <li>• <b>All UDP</b>—All hello messages sent by LSRs to the well-known UDP port, 646.</li> <li>• <b>All TCP</b>—All LDP session messages.</li> </ul> |

Table 371: show ldp statistics Output Fields (*continued*)

| Field Name            | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Event type</b>     | <p>LDP events and errors:</p> <ul style="list-style-type: none"> <li>• <b>Sessions opened</b>—Number of LDP sessions that have been opened.</li> <li>• <b>Sessions closed</b>—Number of LDP sessions that have been closed.</li> <li>• <b>Topology changes</b>—Number of changes to the known LDP topology.</li> <li>• <b>No interface</b>—Number of missing interface address messages. When a new LDP session is initialized and before sending label lapping or label request messages, the LSR advertises its interface addresses with one or more address messages.</li> <li>• <b>No session</b>—Number of missing session messages. Session messages are used to establish, maintain, and terminate sessions between LDP peers.</li> <li>• <b>No adjacency</b>—The exchange of hello adjacency messages results in the creation of an adjacency. The LDP identifier, together with the sender's LDP identifier in the PDU header, enables the receiver to match the initialization message with one of its hello adjacencies. If there is no matching hello adjacency, the LSR sends a session the initialization message is rejected.</li> <li>• <b>Unknown version</b>—The LDP protocol version is not supported by the receiver, or it is supported but is not the version negotiated for the session during session establishment.</li> <li>• <b>Malformed PDU</b>—An LDP PDU received on a TCP connection for an LDP session is malformed if the LDP identifier in the PDU header is unknown to the receiver, or if it is known but is not the LDP identifier associated by the receiver with the LDP peer for this LDP session.<br/> An LDP PDU is considered to be malformed if the LDP protocol version is not supported by the receiver, or it is supported but is not the version negotiated for the session during session establishment.<br/> An LDP PDU is considered malformed if the PDU length field is too small (less than 14) or too large (greater than maximum PDU length).</li> <li>• <b>Malformed message</b>—Malformed LDP messages that are part of the LDP discovery mechanism are handled by silently discarding them.<br/> An LDP message is malformed if the message type is unknown. If the message type is less than 0x8000 (high order bit = 0), it is an error signaled by the unknown message type status code.<br/> An LDP message is considered to be malformed if the message length is too large, meaning that the message extends beyond the end of the containing LDP PDU.<br/> The LDP message is considered to be malformed if the message length is too small, meaning that it is smaller than the smallest possible value component.<br/> The LDP message is considered to be malformed if the message is missing one or more mandatory parameters.</li> <li>• <b>Unknown message type</b>—If the message type is less than 0x8000 (high order bit = 0) or greater than or equal to 0x8000 (high order bit = 1) it is considered to be an unknown message.</li> <li>• <b>Inappropriate message</b>—The message is not of the type that the receiver expects to receive.</li> <li>• <b>Malformed TLV</b>—The TLV Length is too large or the receiver cannot decode the TLV value. This can indicate an issue in either the sending or receiving LSR.</li> <li>• <b>Bad TLV value</b>—The TLV Length is too large.</li> <li>• <b>Missing TLV</b>—The TLV is missing one or more mandatory parameters.</li> <li>• <b>PDU too large</b>—The PDF is greater than the maximum PDU length. Section "Initialization Message" in RFC 5036 describes how the maximum PDU length for a session is determined.</li> </ul> |
| <b>Total</b>          | Total number of each event or error.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Last 5 seconds</b> | Number of each event or error in the last 5 seconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## Sample Output

### show ldp statistics

```
user@host> show ldp statistics
```

| Message type     | Total |          | Last 5 seconds |          |
|------------------|-------|----------|----------------|----------|
|                  | Sent  | Received | Sent           | Received |
| Hello            | 265   | 263      | 2              | 2        |
| Initialization   | 2     | 2        | 0              | 0        |
| Keepalive        | 112   | 111      | 1              | 0        |
| Notification     | 0     | 0        | 0              | 0        |
| Address          | 2     | 2        | 0              | 0        |
| Address withdraw | 0     | 0        | 0              | 0        |
| Label mapping    | 7     | 6        | 0              | 0        |
| Label request    | 0     | 0        | 0              | 0        |
| Label withdraw   | 2     | 0        | 0              | 0        |
| Label release    | 0     | 2        | 0              | 0        |
| Label abort      | 0     | 0        | 0              | 0        |
| All UDP          | 265   | 263      | 2              | 2        |
| All TCP          | 123   | 121      | 1              | 0        |

| Event type            | Total | Last 5 seconds |          |
|-----------------------|-------|----------------|----------|
|                       |       | Sent           | Received |
| Sessions opened       | 2     |                | 0        |
| Sessions closed       | 0     |                | 0        |
| Topology changes      | 11    |                | 0        |
| No interface          | 0     |                | 0        |
| No session            | 0     |                | 0        |
| No adjacency          | 0     |                | 0        |
| Unknown version       | 0     |                | 0        |
| Malformed PDU         | 0     |                | 0        |
| Malformed message     | 0     |                | 0        |
| Unknown message type  | 0     |                | 0        |
| Inappropriate message | 0     |                | 0        |
| Malformed TLV         | 0     |                | 0        |
| Bad TLV value         | 0     |                | 0        |
| Missing TLV           | 0     |                | 0        |
| PDU too large         | 0     |                | 0        |

## show ldp traffic-statistics


|                                                                                                                                                                                                                                                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                        | <pre>show ldp traffic-statistics &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;p2mp&gt;</pre>                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>                                                                                                                                                                                                                           | <p>Command introduced before Junos OS Release 7.4.</p> <p><b>p2mp</b> option added in Junos OS Release 11.2.</p> <p>Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.</p>                                                                                                                                                                                                                                                                                |
| <b>Description</b>                                                                                                                                                                                                                                   | Display Label Distribution Protocol (LDP) traffic statistics.                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <div>  <b>NOTE:</b> If nonstop active routing features is configured, <b>show ldp traffic-statistics</b> command is not supported on backup Routing Engines. </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                                                                                                                                                                                                                                       | <p><b>none</b>—Display LDP traffic statistics for all routing instances.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display LDP traffic statistics for the specified routing instance only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>p2mp</b>—(Optional) Display only the data traffic statistics for a point-to-multipoint LSP.</p> |
| <b>Additional Information</b>                                                                                                                                                                                                                        | To collect output from this command on a periodic basis, configure the <a href="#">traffic-statistics</a> statement for the LDP protocol. For more information, see the <i>Junos MPLS Applications Configuration Guide</i> .                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                      | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>                                                                                                                                                                                                                         | <ul style="list-style-type: none"> <li>• <a href="#">clear ldp statistics on page 4886</a></li> <li>• <i>Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain</i></li> <li>• <i>Example: Configuring Multipoint LDP In-Band Signaling for Point-to-Multipoint LSPs</i></li> </ul>                                                                                                                                                                      |
| <b>List of Sample Output</b>                                                                                                                                                                                                                         | <p><a href="#">show ldp traffic-statistics on page 4921</a></p> <p><a href="#">show ldp traffic-statistics p2mp on page 4922</a></p> <p><a href="#">show ldp traffic-statistics p2mp (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 4922</a></p> <p><a href="#">show ldp traffic-statistics p2mp (Multipoint LDP with Multicast-Only Fast Reroute) on page 4922</a></p>                                                                               |
| <b>Output Fields</b>                                                                                                                                                                                                                                 | <a href="#">Table 372</a> lists the output fields for the <b>show ldp traffic-statistics</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                 |



Table 372: show ldp traffic-statistics Output Fields

| Field Name          | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Message type</b> | LDP message types.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>FEC</b>          | Forwarding equivalence class (FEC) for which LDP traffic statistics are collected.<br><br>For P2MP LSPs, FEC appears as a combination of root address and the LSP ID ( <b>root_addr:lsp_id</b> ).<br><br>For M-LDP P2MP LSPs, FEC appears as a combination of root address multicast source address, and multicast group address ( <b>root_addr:lsp_id/grp,src</b> ).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Type</b>         | Type of traffic originating from a router, either <b>Ingress</b> (originating from this router) or <b>Transit</b> (forwarded through this router).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Packets</b>      | Number of packets passed by the FEC since its LSP came up.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Bytes</b>        | Number of bytes of data passed by the FEC since its LSP came up.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Shared</b>       | Whether a label is shared by prefixes: <b>Yes</b> or <b>No</b> . A <b>Yes</b> value indicates that several prefixes are bound to the same label (for example, when several prefixes are advertised with an egress policy). The LDP traffic statistics for this case apply to all the prefixes and should be treated as such.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Nextthop</b>     | The next hop address for P2MP LSPs. (This is the downstream LDP Session ID.)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Label</b>        | For multipoint LDP with multicast-only fast reroute (MoFRR), the multipoint LDP node selects two separate upstream peers and sends two separate labels, one to each upstream peer. The same algorithm described in RFC 6388 is used to select the primary upstream path. The backup upstream path selection again uses the same algorithm but excludes the primary upstream LSR as a candidate. Two streams of MPLS traffic are sent to the egress node from the two different upstream peers. The MPLS traffic from only one of the upstream neighbors is selected as the primary path to accept the traffic, and the other becomes the backup path. The traffic on the backup path is dropped. When the primary upstream path fails, the traffic from the backup path is then accepted. The multipoint LDP node selects the two upstream paths based on the interior gateway protocol (IGP) root node next hop.<br><br>Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label. |
| <b>Backup route</b> | For multipoint LDP with MoFRR, the route that is used if the primary route becomes unavailable.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## Sample Output

### show ldp traffic-statistics

```
user@host> show ldp traffic-statistics
```

| FEC           | Type    | Packets | Bytes | Shared |
|---------------|---------|---------|-------|--------|
| 10.35.3.0/30  | Transit | 0       | 0     | Yes    |
|               | Ingress | 0       | 0     | No     |
| 10.35.10.1/32 | Transit | 0       | 0     | Yes    |

|                        |              |         |          |        |
|------------------------|--------------|---------|----------|--------|
|                        | Ingress      | 0       | 0        | No     |
| 10.255.245.214/32      | Transit      | 0       | 0        | No     |
|                        | Ingress      | 11      | 752      | No     |
| 192.168.37.36/30       | Transit      | 0       | 0        | Yes    |
|                        | Ingress      | 0       | 0        | No     |
| FEC Statistics:        |              |         |          |        |
| FEC(root_addr:lsp_id)  | Nexthop      | Packets | Bytes    | Shared |
| 10.255.72.160:16777217 | 192.168.8.81 | 152056  | 14597376 | No     |
|                        | 192.168.8.1  | 152056  | 14597376 | No     |
|                        | 192.168.8.65 | 152056  | 14597376 | No     |
| NET FEC Statistics:    |              |         |          |        |
| FEC                    | Type         | Packets | Bytes    | Shared |
| 10.255.107.230/32      | Transit      | 30858   | 2022345  | No     |
|                        | Ingress      | 20      | 5120     | No     |

#### show ldp traffic-statistics p2mp

```

user@host> show ldp traffic-statistics p2mp
FEC(root_addr:lsp_id) Nexthop Packets Bytes Shared
10.255.72.160:16777217 192.168.8.81 152056 14597376 No
 192.168.8.1 152056 14597376 No
 192.168.8.65 152056 14597376 No

```

#### show ldp traffic-statistics p2mp (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```

user@host> show ldp traffic-statistics p2mp
P2MP FEC Statistics:

FEC(root_addr:lsp_id/grp,src) Nexthop Packets Bytes
Shared
11.99.0.73:239.10.0.1,11.98.0.10 11.99.0.117 243408 121217184
No
 11.99.0.13 236286 117670428
No
11.99.0.73:239.10.0.2,11.98.0.10 11.99.0.117 248800 123902400
No
 11.99.0.13 240759 119897982
No
11.99.0.73:239.10.0.1,11.98.0.20 11.99.0.117 250286 124642428
No
 11.99.0.13 243741 121383018
No
11.99.0.73:239.10.0.2,11.98.0.20 11.99.0.117 252970 125979060
No
 11.99.0.13 245218 122118564
No

```

#### show ldp traffic-statistics p2mp (Multipoint LDP with Multicast-Only Fast Reroute)

```

user@host> show ldp traffic-statistics p2mp

```

## P2MP FEC Statistics:

| FEC(root_addr:lsp_id/grp,src)                                 | Nexthop | Packets | Bytes |
|---------------------------------------------------------------|---------|---------|-------|
| Shared                                                        |         |         |       |
| 1.1.1.1:232.1.1.1,192.168.219.11, Label: 301568               | 1.3.8.2 | 0       | 0     |
| No                                                            | 1.3.4.2 | 0       | 0     |
| 1.1.1.1:232.1.1.1,192.168.219.11, Label: 301584, Backup route | 1.3.4.2 | 0       | 0     |
| No                                                            | 1.3.8.2 | 0       | 0     |
| 1.1.1.1:232.1.1.2,192.168.219.11, Label: 301600               | 1.3.8.2 | 0       | 0     |
| No                                                            | 1.3.4.2 | 0       | 0     |
| 1.1.1.1:232.1.1.2,192.168.219.11, Label: 301616, Backup route | 1.3.4.2 | 0       | 0     |
| No                                                            | 1.3.8.2 | 0       | 0     |

## traceroute mpls ldp

---

**Syntax** `traceroute mpls <ldp> fec`  
`<destination>`  
`<detail>`  
`<exp>`  
`<fanout>`  
`<logical-system>`  
`<no-resolve>`  
`<paths>`  
`<retries>`  
`<routing-instance>`  
`<source>`  
`<ttl>`  
`<update>`  
`<wait>`

**Release Information** Command introduced in Junos OS Release 8.4.  
Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

**Description** Trace route to a remote host for an MPLS label-switched path signaled by the LDP. Use **traceroute mpls ldp** as a debugging tool to locate MPLS label-switched path forwarding issues in a network. (Currently supported for IPv4 packets only.)

**Options** *fec*—Specify the IP address and optional prefix of the forwarding equivalence class (FEC).  
*destination*—(Optional) Specify the destination address to use when sending probes.  
*detail*—(Optional) Display detailed output.  
*exp*—(Optional) Specify the class-of-service to use when sending probes. The range of values is 0 through 7. The default value is 7.  
*fanout*—(Optional) Specify the maximum number of nexthops to search per node. The range of values is 1 through 16. The default value is 16.  
*logical-system*—(Optional) Specify the name of the logical system for the traceroute attempt.  
*no-resolve*—(Optional) Specify not to resolve the hostname that corresponds to the IP address.  
*paths*—(Optional) Specify the number of paths to search. The range of values is 1 through 255. The default value is 16.  
*retries*—(Optional) Specify the number of times to resend probe. values. The range of values is 1 through 9. The default value is 3.  
*routing-instance routing-instance-name*—(Optional) Specify the name of the routing instance for the traceroute attempt.  
*source source-address*—(Optional) Specify the source address of the outgoing traceroute packets.

**ttl value**—(Optional) Specify the maximum time-to-live value to include in the traceroute request, in seconds. The range of values is **1** through **125** and the default value is **64**.

**update**—(Optional) Update database contents with traceroute results.

**wait seconds**—(Optional) Specify the number of seconds to wait before resending a probe. The range of values is **5** through **15** and the default value is **10** seconds.

**Required Privilege Level**

network

**List of Sample Output**

[traceroute mpls ldp on page 4926](#)  
[traceroute mpls ldp detail on page 4926](#)

**Output Fields**

[Table 373](#) describes the output fields for the **traceroute mpls ldp fec** command and the **traceroute mpls ldp fec detail** commands. Output fields are listed in the approximate order in which they appear.

**Table 373: traceroute mpls ldp Output Fields**

| Field Name    | Field Description                                                                                                        | Level of Output |
|---------------|--------------------------------------------------------------------------------------------------------------------------|-----------------|
| Probe options | Probe options specified in the <b>traceroute mpls ldp fec</b> command.                                                   | all levels      |
| ttl           | Time to live value of the labeled packet.                                                                                | none specified  |
| Label         | Outgoing label used for forwarding the packet along the label-switched paths.                                            | none specified  |
| Protocol      | Signaling protocol used. For this command, it is LDP.                                                                    | none specified  |
| Address       | Address of the next hop.                                                                                                 | none specified  |
| Previous Hop  | Address of the previous hop. Previous hop address of the first hop is <b>null</b> .                                      | none specified  |
| Probe status  | Forwarding status from the first hop to the last-hop label-switching router (egress point in the label-switched paths).  | none specified  |
| Hop           | Address of the hops in the label-switched path from the first hop to the last hop. Depth indicates the level of the hop. | <b>detail</b>   |
| Parent        | Address of the previous hop. Parent value for the first hop is <b>null</b> .                                             | <b>detail</b>   |
| Return Code   | Return code for reporting the result of processing the echo request by the receiver.                                     | <b>detail</b>   |
| Response time | Time for the echo request to reach the receiver.                                                                         | <b>detail</b>   |

Table 373: traceroute mpls ldp Output Fields (*continued*)

| Field Name     | Field Description                                                                                                | Level of Output |
|----------------|------------------------------------------------------------------------------------------------------------------|-----------------|
| Multipath type | Labels or addresses used by the specified multipath type. If multipaths are not used, the value is <b>none</b> . | <b>detail</b>   |
| Label Stack    | Label stack used to forward the packet.                                                                          | <b>detail</b>   |

## Sample Output

### traceroute mpls ldp

```
user@router> traceroute mpls ldp 4.4.4.4
```

```
Probe options: ttl 64, retries 3, wait 10, paths 16, exp 7, fanout 16
ttl Label Protocol Address Previous Hop Probe Status
 1 100016 LDP 24.24.24.1 (null) Success
 2 100000 LDP 20.20.20.2 24.24.24.1 Success
 3 3 LDP 22.22.22.4 20.20.20.2 Egress
```

```
Path 1 via fe-0/3/3.101 destination 127.0.0.64
```

### traceroute mpls ldp detail

```
user@router> traceroute mpls ldp 4.4.4.4 detail
```

```
Probe Options: ttl 64, retries 3, wait 10, paths 3, exp 7
Hop 24.24.24.1 Depth 1
 Parent (null)
 Return code: Label switched at stack-depth 1
 Response time 165.93 msec
 Multipath type: IP bitmask
 Address Range 1: 127.0.0.0 ~ 127.0.3.255
 Label Stack:
 Label 1 Value 100032 Protocol LDP

Hop 20.20.20.2 Depth 2
 Parent 24.24.24.1
 Return code: Upstream interface index unknown label-switched at stack-depth
1
 Response time 19.05 msec
 Multipath type: IP bitmask
 Address Range 1: 127.0.0.0 ~ 127.0.3.255
 Label Stack:
 Label 1 Value 100000 Protocol LDP

Hop 22.22.22.4 Depth 3
 Parent 20.20.20.2
 Return code: Egress-ok at stack-depth 1
 Response time 0.79 msec
 Multipath type: None
 Label Stack:
 Label 1 Value 3 Protocol LDP
```

## PART 72

# MPLS

- [Using MPLS on page 4929](#)
- [Configuration Statements and Command Statements for MPLS on page 5017](#)





# Using MPLS

- [MPLS Overview For QFX Series and EX4600 Switches on page 4930](#)
- [MPLS Feature Support on QFX Series and EX4600 Switches on page 4934](#)
- [MPLS Limitations on QFX Series and EX4600 Switches on page 4943](#)
- [Understanding MPLS Components for QFX Series and EX4600 Switches on page 4946](#)
- [Understanding MPLS Label Operations on page 4951](#)
- [Understanding BGP on page 4955](#)
- [IPv6 Layer 3 VPNs on page 4957](#)
- [Ethernet Pseudowire Overview on page 4958](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4959](#)
- [Understanding Ethernet-over-MPLS \(L2 Circuit\) on page 4962](#)
- [Understanding Using MPLS-Based Layer 3 VPNs on Switches on page 4963](#)
- [Carrier-of-Carriers VPNs on page 4965](#)
- [Interprovider and Carrier-of-Carriers VPNs on page 4966](#)
- [Fast Reroute Overview on page 4967](#)
- [Graceful Restart and MPLS-Related Protocols on page 4969](#)
- [Types of LSPs on page 4971](#)
- [Configuring CoS Bits for an MPLS Network on page 4972](#)
- [Configuring Ethernet over MPLS \(L2 Circuit\) on page 4972](#)
- [Configuring a Global MPLS EXP Classifier on page 4977](#)
- [Configuring MPLS Firewall Filters and Policers on page 4977](#)
- [Configuring MPLS to Gather Statistics on page 4980](#)
- [Configuring MPLS on Provider Edge Switches on page 4981](#)
- [Configuring MPLS on Provider Switches on page 4985](#)
- [Configuring Reporting of Automatic Bandwidth Allocation Statistics for LSPs on page 4986](#)
- [Configuring Static Label Switched Paths for MPLS on page 4989](#)
- [Configuring Rewrite Rules for MPLS EXP Classifiers on page 4992](#)
- [Example: Configuring MPLS-Based Layer 3 VPNs on page 4993](#)
- [Example: Tunneling IPv6 Traffic over MPLS IPv4 Networks on page 5002](#)

- [Verifying That MPLS Is Working Correctly on page 5010](#)
- [MPLS Configuration Guidelines on page 5012](#)
- [Supported MPLS Scaling Values on page 5013](#)
- [MPLS Stitching For Virtual Machine Connection on page 5014](#)

## **MPLS Overview For QFX Series and EX4600 Switches**

---

Multiprotocol Label Switching (MPLS) is a protocol that uses labels to route packets instead of using IP addresses. In a traditional network, each switch performs an IP routing lookup, determines a next-hop based on its routing table, and then forwards a packet to that next-hop. With MPLS, only the first device does a routing lookup, and, instead of finding the next-hop, finds the ultimate destination along with a path to that destination. The path of an MPLS packet is called a label-switched path (LSP).

MPLS applies one or more labels to a packet so it can follow the LSP to the destination. Each switch pops off its label and sends the packet to the next switch label in the sequence.

The Junos OS includes everything you need to configure MPLS. You do not need to install any additional programs or protocols. MPLS is supported on switches with a subset of the commands supported on routers. The Junos MPLS-configured switches can interact with each other and with Junos MPLS-configured routers.

MPLS has the following advantages over conventional packet forwarding:

- Packets arriving on different ports can be assigned different labels.
- A packet arriving at a particular provider edge (PE) switch can be assigned a label that is different from that of the same packet entering the network at a different PE switch. As a result, forwarding decisions that depend on the ingress PE switch can be easily made.
- Sometimes it is desirable to force a packet to follow a particular route that is explicitly chosen at or before the time the packet enters the network, rather than letting it follow the route chosen by the normal dynamic routing algorithm as the packet travels through the network. In MPLS, a label can be used to represent the route so that the packet need not carry the identity of the explicit route.

This topic describes:

- [Why Use MPLS? on page 4931](#)
- [Why Not Use MPLS? on page 4931](#)
- [How Do I Configure MPLS? on page 4931](#)
- [What Does the MPLS Protocol Do? on page 4932](#)
- [How Does MPLS Interface to Other Protocols? on page 4933](#)
- [If I Have Used Cisco MPLS, What Do I Need to Know? on page 4933](#)

## Why Use MPLS?

MPLS reduces the use of the forwarding table by using labels instead of the forwarding table. The size of forwarding tables on a switch are limited by silicon and using exact matching for forwarding to destination devices is cheaper than buying more sophisticated hardware. In addition, MPLS allows you to control where and how traffic is routed on your network – this is called traffic engineering.

Some reasons to use MPLS instead of another switching solution are:

- MPLS can connect different technologies that would not otherwise be compatible---service providers have this compatibility issue when connecting clients with different autonomous systems in their networks. In addition, MPLS has a feature called Fast Reroute that provides alternate backups for paths – this prevents network degradation in case of a switch failure.
- Other IP-based encapsulations such as Generic Route Encapsulation (GRE) or Virtual Extensible Local Area Networks (VXLAN) support only two levels of hierarchy, one for the transport tunnel and one piece of metadata. Using virtual servers means that you need multiple hierarchy levels. For example, one label is needed for top-of-rack (ToR), one label for the egress port that identifies the server, and one for the virtual server.

## Why Not Use MPLS?

There are no protocols to auto-discover MPLS enabled nodes. MPLS protocol just exchanges label values for an LSP. They do not create the LSPs.

You must build the MPLS mesh, switch by switch. We recommend using scripts for this repetitive process.

MPLS hides suboptimal topologies from BGP where multiple exits may exist for the same route.

Large LSPs are limited by the circuits they traverse. You can work around this by creating multiple, parallel LSPs.

## How Do I Configure MPLS?

There are three types of switches you must set up for MPLS:

- Label Edge Router/Switch (LER) or ingress node to the MPLS network. This switch encapsulates the packets.
- Label Switching Routers/Switches (LSR). One or more switches that transfer MPLS packets in the MPLS network.
- Egress router/switch is the final MPLS device that removes the last label before packets leave the MPLS network.

Service providers (SP) use the term provider router (P) for a backbone router/switch doing label switching only. The customer-facing router at the SP is called a provider edge router (PE). Each customer needs a customer edge router (CE) to communicate with

the PE. Customer facing routers typically can terminate IP addresses, L3VPNs, L2VPNs/pseudowires, and VPLS before packets are transferred to the CE.

### Configure the MPLS LER (Ingress) Switch and the Egress Switch

---

To configure MPLS, you must first create one or more named paths on the ingress and egress routers. For each path, you can specify some or all transit routers in the path, or you can leave it empty. See *Configuring the Ingress Router for MPLS-Signaled LSPs* and *Configuring the Intermediate and Egress Routers for MPLS-Signaled LSPs*, *Configuring the Ingress and Egress Router Addresses for LSPs*, and *Configuring the Connection Between Ingress and Egress Routers*.

### Configure LSRs for MPLS

---

Configure one or more MPLS LSRs by following these steps:

1. Configure interfaces on each switch to transmit and receive MPLS packets using the usual interface command with MPLS appended. For example:

```
[edit interfaces ge-0/0/0 unit 0] family mpls;
```

2. Add those same interfaces under [edit protocols mpls]. For example:

```
[edit protocols mpls]
 interface ge-0/0/0;
```

3. Configure the interfaces on each switch to handle MPLS labels with a protocol. For example, for LDP:

```
[edit protocols ldp]
 Interface ge-0/0/0.0;
```

To watch a demo of these configurations, see  
<https://www.youtube.com/watch?v=xegWBCUJ4tE>.

## What Does the MPLS Protocol Do?

Multiprotocol Label Switching (MPLS) is an Internet Engineering Task Force (IETF)-specified framework that provides for the designation, routing, forwarding and switching of traffic flows through the network. In addition, MPLS:

- Specifies mechanisms to manage traffic flows of various granularities, such as flows between different hardware, machines, or even flows between different applications.
- Remains independent of the layer-2 and layer-3 protocols.
- Provides a means to map IP addresses to simple, fixed-length labels used by different packet-forwarding and packet-switching technologies.
- Interfaces to existing routing protocols, such as Resource ReSerVation Protocol (RSVP) and Open Shortest PathFirst (OSPF).
- Supports IP, ATM, and Frame Relay layer-2 protocols.
- Uses these additional technologies:

- FRR: MPLS Fast Reroute improves convergence during a failure by mapping out alternate LSPs in advance.
- Link Protection/ Next-hop backup: A bypass LSP is created for every possible link failure.
- Node Protection/ Next-hop backup: A bypass LSP is created for every possible switch (node) failure.
- VPLS: Creates Ethernet multipoint switching service over MPLS and emulates functions of an L2 switch.
- L3VPN: IP-based VPN customers get individual virtual routing domains.

## How Does MPLS Interface to Other Protocols?

Some of the protocols that work with MPLS are:

- RSVP-TE: Resource Reservation Protocol - Traffic Engineering reserves bandwidth for LSPs.
- LDP: Label Distribution Protocol is the defacto protocol used for distribution of MPLS packets and is usually configured to tunnel inside RSVP-TE.
- IGP: Interior Gateway Protocol is a routing protocol. Edge routers (PE-routers) run BGP between themselves to exchange external (customer) prefixes. Edge and core (P) routers run IGP (usually OSPF or IS-IS) to find optimum path toward BGP next hops. P- and PE-routers use LDP to exchange labels for known IP prefixes (including BGP next hops). LDP indirectly builds end-to-end LSPs across the network core.
- BGP: Border Gateway Protocol (BGP) is allows policy-based routing to take place, using TCP as its transport protocol on port 179 to establish connections. The Junos OS routing protocol software includes BGP version 4. You do not configure BGP---configuring interfaces with MPLS and LDP/RSVP establishes the labels and the ability to transmit packets. BGP automatically determines the routes packets take.
- OSPF and ISIS: These protocols are used for routing between the MPLS PE and CE. Open Shortest Path First (OSPF) is perhaps the most widely used interior gateway protocol (IGP) in large enterprise networks. IS-IS, another link-state dynamic routing protocol, is more common in large service provider networks. Assuming you're running L3VPN to your customers, on the SP edge between the PE and the CE you can run any protocol that your platform supports as a VRF aware instance.

## If I Have Used Cisco MPLS, What Do I Need to Know?

Cisco Networks and Juniper Networks use different MPLS terminology.

| What Cisco Calls:    | Juniper Calls: |
|----------------------|----------------|
| affinities           | admin-groups   |
| autoroute announce   | TE shortcuts   |
| forwarding adjacency | LSP-advertise  |

| What Cisco Calls:       | Juniper Calls:  |
|-------------------------|-----------------|
| tunnel                  | LSP             |
| make-before-break       | adaptive        |
| application-window      | adjust-interval |
| shared risk link groups | fate-sharing    |

**Related Documentation**

- [MPLS Feature Support on QFX Series and EX4600 Switches on page 4934](#)
- [Understanding MPLS Components for QFX Series and EX4600 Switches on page 4946](#)
- [Understanding MPLS Label Operations on page 4951](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4959](#)
- [Junos OS MPLS Applications Library for Routing Devices](#)

## MPLS Feature Support on QFX Series and EX4600 Switches

Multiprotocol Label Switching (MPLS) is a set of procedures for augmenting network layer packets with label stacks, thereby turning them into labeled packets. Service providers frequently use MPLS. Simply put, where traditional networks' routers each perform an IP lookup to determine the next hop, an MPLS network's first device does a routing lookup for the final destination instead of the next hop. A label is then applied to the packet—this is called packet switching. The final destination device removes the label.

A number of Juniper Networks switches are capable of running a subset of MPLS and can therefore communicate, not only with each other, but with Juniper Networks routers running MPLS. This topic describes the major MPLS features that are supported on QFX Series switches and on one EX Series switch, the EX4600. Be sure to check for any exceptions to this support in "[MPLS Limitations on QFX Series and EX4600 Switches](#)" on page 4943.



**NOTE:** EX4600 switches use the same chipset as QFX5100 switches—this is why that EX Series switch is discussed here along with QFX Series switches. Other EX Series switches also support MPLS but with a different feature set.

This topic describes:

- [MPLS Commands Supported by QFX Series and EX4600 Switches on page 4935](#)
- [MPLS Features Supported by QFX Series and EX4600 Switches on page 4935](#)

## MPLS Commands Supported by QFX Series and EX4600 Switches

QFX Series and EX4600 switches support a subset of MPLS features. The command-line interface (CLI) for switches displays all MPLS related configuration statements, even those that are not supported. However, configuring those unsupported statements on a switch has no effect on the operation of the switch. See the following topics for the list of all MPLS related configuration statements:

- “[[edit protocols mpls](#)] Hierarchy Level” on [page 5019](#) for the list of supported configuration statements at the [[edit protocols mpls](#)] hierarchy level
- “[[edit protocols rsvp](#)] Hierarchy Level” on [page 5302](#) for the list of supported configuration statements at the [[edit protocols rsvp](#)] hierarchy level

## MPLS Features Supported by QFX Series and EX4600 Switches

This section lists the major MPLS features supported on QFX Series and EX4600 switches and the Junos OS release in which they were introduced. [Table 374](#) lists the features for the QFX10000 Series switches. [Table 375](#) lists the features for the QFX3500, QFX5100, and QFX5200 switches. [Table 376](#) lists the features for the EX4600 switches.

**Table 374: QFX10000 Switch MPLS Features with Junos OS Release Support**

| Feature                                                                          | QFX10002      | QFX10008      |
|----------------------------------------------------------------------------------|---------------|---------------|
| QFX standalone switch as an MPLS provider edge (PE) switch or provider switch    | 15.1X53-D10   | 15.1X53-D30   |
| Route reflector for BGP labeled routes                                           | 15.1X53-D10   | 15.1X53-D30   |
| Label edge router (LER)                                                          | 15.1X53-D10   | 15.1X53-D30   |
| Label switch router (LSR)                                                        | 15.1X53-D10   | 15.1X53-D30   |
| Border Gateway Protocol (BGP) labeled unicast                                    | 15.1X53-D10   | 15.1X53-D30   |
| Carrier-over-carrier and inter-provider BGP L3 VPN                               | not supported | not supported |
| Class of Service (CoS or QoS) for MPLS traffic                                   | 15.1X53-D10   | 15.1X53-D30   |
| Dynamic LSP count sizing                                                         | 15.1X53-D30   | 15.1X53-D30   |
| Fast Reroute (FRR), one-to-one local protection and many-to-one local protection | 15.1X53-D10   | 15.1X53-D30   |
| Ethernet-over-MPLS pseudowires based on LDP                                      | not supported | not supported |
| FRR using detours and secondary LSP                                              | 15.1X53-D10   | 15.1X53-D30   |

Table 374: QFX10000 Switch MPLS Features with Junos OS Release Support (*continued*)

| Feature                                                                                                                                                                     | QFX10002      | QFX10008      |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|---------------|
| Firewall filters                                                                                                                                                            | 15.1X53-D30   | 15.1X53-D30   |
| Graceful restart for Open Shortest Path First (OSPF) for Resource Reservation Protocol (RSVP)                                                                               | 15.1X53-D10   | 15.1X53-D30   |
| IPv6 tunneling for over an MPLS-based IPv4 network (6PE)                                                                                                                    | 15.1X53-D10   | 15.1X53-D30   |
| Layer 3 VPN 6PE                                                                                                                                                             |               |               |
| Intermediate System to Intermediate System (ISIS) routing protocol as an interior gateway protocol (IGP) for MPLS. IS-IS interior gateway protocol traffic engineering (TE) | 15.1X53-D10   | 15.1X53-D30   |
| L2 circuit (draft Martini/L2Circuit)                                                                                                                                        | not supported | not supported |
| Layer 3 VPNs IPv4                                                                                                                                                           | 15.1X53-D10   | 15.1X53-D30   |
| Label Distribution Protocol (LDP) based signaling over RSVP                                                                                                                 | 15.1X53-D10   | 15.1X53-D30   |
| IP-over-MPLS label-switched paths (LSPs) both static and dynamic links                                                                                                      | 15.1X53-D10   | 15.1X53-D30   |
| MTU signaling in RSVP                                                                                                                                                       | 15.1X53-D10   | 15.1X53-D30   |
| Object access method (OAM) including MPLS ping, traceroute and BFD                                                                                                          | 15.1X53-D10   | 15.1X53-D30   |
| MPLS over integrated bridging and routing (IRB) interfaces                                                                                                                  | 15.1X53-D10   | 15.1X53-D30   |
| Open Shortest Path First traffic engineering (OSPF TE)                                                                                                                      | 15.1X53-D10   | 15.1X53-D30   |
| OSPFv2 as an interior gateway protocol                                                                                                                                      | 15.1X53-D10   | 15.1X53-D30   |
| Per VRF label support                                                                                                                                                       | 15.1X53-D10   | 15.1X53-D30   |
| Ethernet-over-MPLS pseudowires based on LDP                                                                                                                                 | not supported | not supported |
| Pseudowire-over-aggregated Ethernet interfaces (core-facing interface)                                                                                                      | not supported | not supported |
| Pseudowire, static and dynamic                                                                                                                                              | not supported | not supported |
| Resource Reservation Protocol (RSVP)                                                                                                                                        | 15.1X53-D10   | 15.1X53-D30   |
| RSVP bandwidth                                                                                                                                                              | 15.1X53-D10   | 15.1X53-D30   |



Table 374: QFX10000 Switch MPLS Features with Junos OS Release Support (*continued*)

| Feature                                                                                                 | QFX10002      | QFX10008      |
|---------------------------------------------------------------------------------------------------------|---------------|---------------|
| RSVP fast reroute including link-protection, node-link-protection, FRR using detours, and secondary LSP | 15.1X53-D10   | 15.1X53-D30   |
| RSVP Traffic engineering (used to establish LSPs) with IS-IS and OSPF extensions                        | 15.1X53-D10   | 15.1X53-D30   |
| SNMP MIB support                                                                                        | 15.1X53-D10   | 15.1X54-D30   |
| Static and dynamic LSPs                                                                                 | 15.1X53-D10   | 15.1X53-D30   |
| Virtual private wire service (VPWS)                                                                     | not supported | not supported |
| Traffic Engineering (TE)                                                                                | 15.1X53-D10   | 15.1X53-D30   |
| TE auto-bandwidth and RSVP bandwidth                                                                    | 15.1X53-D10   | 15.1X53-D30   |
| Dynamic bandwidth management using ingress LSP splitting and merging                                    |               |               |

Table 375: QFX3500, QFX5100, and QFX5200 MPLS Features with Junos OS Release Support

| Feature                                                                       | QFX3500       | QFX5100                              | QFX5200     |
|-------------------------------------------------------------------------------|---------------|--------------------------------------|-------------|
| QFX standalone switch as an MPLS provider edge (PE) switch or provider switch | 12.2X50-D10   | 13.2X51-D15<br>VC/VCF<br>14.1X53-D30 | 15.1X53-D30 |
| Route reflector for BGP labeled routes                                        | 12.2X50-D10   | 13.2X51-D15<br>VC/VCF<br>14.1X53-D30 | 15.1X53-D30 |
| Label edge router (LER)                                                       | 12.2X50-D10   | 13.2X51-D15<br>VC/VCF<br>14.1X53-D30 | 15.1X53-D30 |
| Label switch router (LSR)                                                     | 12.2X50-D10   | 13.2X51-D15<br>VC/VCF<br>14.1X53-D30 | 15.1X53-D30 |
| Auto Bandwidth                                                                | not supported | 13.2X51-D15<br>VC/VCF<br>14.1X53-D30 | 15.1X53-D30 |

**Table 375: QFX3500, QFX5100, and QFX5200 MPLS Features with Junos OS Release Support (*continued*)**

| Feature                                                                                                                                | QFX3500       | QFX5100                              | QFX5200       |
|----------------------------------------------------------------------------------------------------------------------------------------|---------------|--------------------------------------|---------------|
| Border Gateway Protocol (BGP) labeled unicast                                                                                          | 12.2X50-D10   | 13.2X51-D15<br>VC/VCF<br>14.1X53-D30 | 15.1X53-D30   |
| Carrier-over-carrier and inter-provider BGP L3 VPN                                                                                     | 14.1X53-D15   | 14.1X53-D15                          | 15.1X53-D30   |
| Class of Service (CoS or QoS) for MPLS traffic                                                                                         | 12.3X50-D10   | 13.2X51-D15<br>VC/VCF<br>14.1X53-D30 | 15.1X53-D30   |
| ECMP at LSR routers:<br><ul style="list-style-type: none"> <li>• SWAP</li> <li>• PHP</li> <li>• L3VPN</li> <li>• L2 Circuit</li> </ul> | not supported | not supported                        | 15.1X53-D30   |
| Entropy Labels                                                                                                                         | not supported | 13.2X51-D15<br>VC/VCF<br>14.1X53-D30 | not supported |
| Ethernet-over-MPLS pseudowires based on LDP                                                                                            | 14.1X53-D15   | 14.1X53-D15                          | 15.1X53-D30   |
| Fast Reroute (FRR), one-to-one local protection and many-to-one local protection                                                       | 14.1X53-D10   | 14.1X53-D10<br>VC/VCF not supported  | 15.1X53-D30   |
| FRR using detours and secondary LSP                                                                                                    | not supported | not supported                        | not supported |
| Firewall filters                                                                                                                       | 12.3X50-D10   | 13.2X51-D15<br>VC/VCF<br>14.1X53-D30 | 15.1X53-D30   |
| Flow Labels                                                                                                                            | not supported | 13.2X51-D15<br>VC/VCF<br>14.1X53-D30 | not supported |
| Graceful restart for Open Shortest Path First (OSPF) for Resource Reservation Protocol (RSVP)                                          | 12.2X50-D10   | 13.2X51-D15<br>VC/VCF<br>14.1X53-D30 | 15.1X53-D30   |

**Table 375: QFX3500, QFX5100, and QFX5200 MPLS Features with Junos OS Release Support (*continued*)**

| Feature                                                                                                                                                                     | QFX3500       | QFX5100                                  | QFX5200       |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|------------------------------------------|---------------|
| IPv6 tunneling for over an MPLS-based IPv4 network (6PE)                                                                                                                    | 12.3X50-D10   | 13.2X51-D15                              | 15.1X53-D30   |
| Layer 3 VPN 6PE                                                                                                                                                             |               | VC/VCF<br>14.1X53-D30                    |               |
| Intermediate System to Intermediate System (ISIS) routing protocol as an interior gateway protocol (IGP) for MPLS. IS-IS interior gateway protocol traffic engineering (TE) | 12.2X50-D10   | 13.2X51-D15<br><br>VC/VCF<br>14.1X53-D30 | 15.1X53-D30   |
| L2 circuit (draft Martini/L2Circuit)                                                                                                                                        | 14.1X53-D10   | 14.1X53-D30<br><br>VC/VCF<br>14.1X53-D30 | 15.1X53-D30   |
| Layer 3 VPNs IPv4                                                                                                                                                           | 12.3X50-D10   | 13.2X51-D15<br><br>VC/VCF<br>14.1X53-D30 | 15.1X53-D30   |
| Label Distribution Protocol (LDP) based signaling over RSVP                                                                                                                 | 12.2X50-D10   | 13.2X51-D15<br><br>VC/VCF<br>14.1X53-D30 | 15.1X53-D30   |
| Load balancing at provider routers                                                                                                                                          | not supported | not supported                            | 15.1X53-D30   |
| Loop-free Alternative (LFA)                                                                                                                                                 | not supported | 13.2X51-D15<br><br>VC/VCF<br>14.1X53-D30 | not supported |
| IP-over-MPLS label-switched paths (LSPs) both static and dynamic links                                                                                                      | 12.2X50-D10   | 13.2X51-D15<br><br>VC/VCF<br>14.1X53-D30 | 15.1X53-D30   |
| MTU signaling in RSVP                                                                                                                                                       | 12.3X50-D10   | 13.2X51-D15<br><br>VC/VCF<br>14.1X53-D30 | 15.1X53-D30   |
| Object access method (OAM) including MPLS ping, traceroute and BFD                                                                                                          | 12.3X50-D10   | 13.2X51-D15<br><br>VC/VCF<br>14.1X53-D30 | 15.1X53-D30   |
| MPLS over integrated bridging and routing (IRB) interfaces                                                                                                                  | not supported | not supported                            | not supported |

**Table 375: QFX3500, QFX5100, and QFX5200 MPLS Features with Junos OS Release Support (*continued*)**

| Feature                                                                                                 | QFX3500     | QFX5100                                   | QFX5200     |
|---------------------------------------------------------------------------------------------------------|-------------|-------------------------------------------|-------------|
| Open Shortest Path First traffic engineering (OSPF TE)                                                  | 14.1X53-D10 | 14.1X53-D10<br><br>VC/VCF not supported   | 15.1X53-D30 |
| OSPFv2 as an interior gateway protocol                                                                  | 12.2X50-D10 | 13.2X51-D15<br><br>VC/VCF<br>14.1X53-D30  | 15.1X53-D30 |
| Per VRF label support                                                                                   | 12.2X50-D10 | 13.2X51-D15<br><br>VC/VCF<br>14.1X53-D30  | 15.1X53-D30 |
| Ethernet-over-MPLS pseudowires based on LDP                                                             | 14.1X53-D10 | 14.1X53-D10<br><br>VC/VCF not supported   | 15.1X53-D30 |
| Pseudowire-over-aggregated Ethernet interfaces (core-facing interface)                                  | 14.1X53-D10 | 14.1X53-D15*<br><br>VC/VCF<br>14.1X53-D30 | 15.1X53-D30 |
| Pseudowire, static and dynamic                                                                          | 14.1X53-D10 | 14.1X53-D10<br><br>VC/VCF not supported   | 15.1X53-D30 |
| Resource Reservation Protocol (RSVP)                                                                    | 12.2X50-D10 | 13.2X51-D15<br><br>VC/VCF<br>14.1X53-D30  | 15.1X53-D30 |
| RSVP bandwidth and auto-bandwidth                                                                       | 12.2X50-D10 | 13.2X51-D15<br><br>VC/VCF<br>14.1X53-D30  | 15.1X53-D30 |
| RSVP fast reroute including link-protection, node-link-protection, FRR using detours, and secondary LSP | 14.1X53-D15 | 14.1X53-D15<br><br>VC/VCF not supported   | 15.1X53-D30 |
| RSVP Traffic engineering (used to establish LSPs) with IS-IS and OSPF extensions                        | 12.2X50-D10 | 13.2X51-D15<br><br>VC/VCF<br>14.1X53-D30  | 15.1X53-D30 |

**Table 375: QFX3500, QFX5100, and QFX5200 MPLS Features with Junos OS Release Support (*continued*)**

| Feature                              | QFX3500       | QFX5100                              | QFX5200       |
|--------------------------------------|---------------|--------------------------------------|---------------|
| SNMP MIB support                     | 12.2X50-D10   | 13.2X51-D15<br>VC/VCF<br>14.1X53-D30 | 15.1X53-D30   |
| Static and dynamic LSPs              | 12.2X50-D10   | 13.2X51-D10                          | 15.1X53-D30   |
| Virtual private wire service (VPWS)  | not supported | 14.1-X53-D25                         | not supported |
| Traffic Engineering (TE)             | 13.1X51-D10   | 13.1X51-D10                          | 15.1X53-D30   |
| TE auto-bandwidth and RSVP bandwidth | 13.1X51-D10   | 13.1X51-D10                          | 15.1X53-D30   |

**Table 376: EX4600 Switch MPLS Features with Junos OS Release Support**

| Feature                                                                          | EX4600                                 |
|----------------------------------------------------------------------------------|----------------------------------------|
| QFX standalone switch as an MPLS provider edge (PE) switch or provider switch    | 14.1X53-D15 but not 15.1X53-D10X53-D10 |
| Route reflector for BGP labeled routes                                           | 14.1X53-D15 but not 15.1X53-D10        |
| Label edge router (LER)                                                          | 14.1X53-D15 but not 15.1X53-D10        |
| Label switch router (LSR)                                                        | 14.1X53-D15 but not 15.1X53-D10        |
| Border Gateway Protocol (BGP) labeled unicast                                    | 14.1X53-D15 but not 15.1X53-D10        |
| Carrier-over-carrier and inter-provider BGP L3 VPN                               | not supported                          |
| Class of Service (CoS or QoS) for MPLS traffic                                   | 14.1X53-D15 but not 15.1X53-D10        |
| Fast Reroute (FRR), one-to-one local protection and many-to-one local protection | not supported                          |
| Ethernet-over-MPLS pseudowires based on LDP                                      | 14.1X53-D15 but not 15.1X53-D10        |
| FRR using detours and secondary LSP                                              | not supported                          |
| Firewall filters                                                                 | 14.1X53-D15 but not 15.1X53-D10        |

Table 376: EX4600 Switch MPLS Features with Junos OS Release Support (*continued*)

| Feature                                                                                                                                                                     | EX4600                          |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| Graceful restart for Open Shortest Path First (OSPF) for Resource Reservation Protocol (RSVP)                                                                               | 13.2X51-D25                     |
| IPv6 tunneling for over an MPLS-based IPv4 network (6PE)                                                                                                                    | 14.1X53-D15 but not 15.1X53-D10 |
| Layer 3 VPN 6PE                                                                                                                                                             |                                 |
| Intermediate System to Intermediate System (ISIS) routing protocol as an interior gateway protocol (IGP) for MPLS. IS-IS interior gateway protocol traffic engineering (TE) | 14.1X53-D15 but not 15.1X53-D10 |
| L2 circuit (draft Martini/L2Circuit)                                                                                                                                        | 14.1X53-D15 but not 15.1X53-D10 |
| Layer 3 VPNs IPv4                                                                                                                                                           | 14.1X53-D15 but not 15.1X53-D10 |
| Label Distribution Protocol (LDP) based signaling over RSVP                                                                                                                 | 14.1X53-D15 but not 15.1X53-D10 |
| IP-over-MPLS label-switched paths (LSPs) both static and dynamic links                                                                                                      | 14.1X53-D15 but not 15.1X53-D10 |
| MTU signaling in RSVP                                                                                                                                                       | 14.1X53-D15 but not 15.1X53-D10 |
| Object access method (OAM) including MPLS ping, traceroute and BFD                                                                                                          | 14.1X53-D15 but not 15.1X53-D10 |
| MPLS over integrated bridging and routing (IRB) interfaces                                                                                                                  | not supported                   |
| Open Shortest Path First traffic engineering (OSPF TE)                                                                                                                      | 14.1X53-D15 but not 15.1X53-D10 |
| OSPFv2 as an interior gateway protocol                                                                                                                                      | 13.2X51-D25                     |
| Per VRF label support                                                                                                                                                       | 14.1X53-D15 but not 15.1X53-D10 |
| Ethernet-over-MPLS pseudowires based on LDP                                                                                                                                 | 14.1X53-D15 but not 15.1X53-D10 |
| Pseudowire-over-aggregated Ethernet interfaces (core-facing interface)                                                                                                      | 14.1X53-D15 but not 15.1X53-D10 |
| Pseudowire, static and dynamic                                                                                                                                              | 14.1X53-D15 but not 15.1X53-D10 |
| Resource Reservation Protocol (RSVP)                                                                                                                                        | 14.1X53-D15 but not 15.1X53-D10 |

Table 376: EX4600 Switch MPLS Features with Junos OS Release Support (*continued*)

| Feature                                                                                                 | EX4600                          |
|---------------------------------------------------------------------------------------------------------|---------------------------------|
| RSVP bandwidth and auto-bandwidth                                                                       | 14.1X53-D15 but not 15.1X53-D10 |
| RSVP fast reroute including link-protection, node-link-protection, FRR using detours, and secondary LSP | not supported                   |
| RSVP Traffic engineering (used to establish LSPs) with IS-IS and OSPF extensions                        | 14.1X53-D15 but not 15.1X53-D10 |
| SNMP MIB support                                                                                        | 14.1X53-D15 but not 15.1X53-D10 |
| Static and dynamic LSPs                                                                                 | 14.1X53-D15 but not 15.1X53-D10 |
| Virtual private wire service (VPWS)                                                                     | not supported                   |
| Traffic Engineering (TE)                                                                                | 14.1X53-D15 but not 15.1X53-D10 |
| TE auto-bandwidth and RSVP bandwidth                                                                    | 14.1X53-D15 but not 15.1X53-D10 |

- Related Documentation**
- [MPLS Limitations on QFX Series and EX4600 Switches on page 4943](#)
  - [MPLS Configuration Guidelines on page 5012](#)

## MPLS Limitations on QFX Series and EX4600 Switches

MPLS is fully implemented on routers, while switches support a subset of the MPLS features. The limitations of each switch are listed in a separate section here, even though many of the limitations are duplicates that apply to more than one switch.

- [MPLS Limitations on QFX3500 Switches on page 4943](#)
- [MPLS Limitations on QFX5100 and EX4600 Switches on page 4944](#)
- [MPLS Limitations on QFX5100 Virtual Chassis and Virtual Chassis Fabric on page 4945](#)
- [MPLS Limitations on QFX10000 Switches on page 4946](#)

### MPLS Limitations on QFX3500 Switches

- If you configure the BGP labeled unicast address family (using the **labeled-unicast** statement at the **[edit protocols bgp family inet]** hierarchy level) on a QFX switch or on an EX4600 switch deployed as a route reflector for BGP labeled routes, path selection will occur at the route reflector, and a single best path will be advertised. This will result in loss of BGP multipath information.
- Fast reroute is supported, however the **include-all** and **include-any** options for fast reroute are not supported. For more information, see [“Fast Reroute Overview” on page 4967](#).

- MPLS-based circuit cross-connects (CCC) is not supported—only circuit-based pseudowires are supported.
- MTU signaling in RSVP and discovery is supported in Control Plane. However, this cannot be enforced in data plane.
- With L2 circuit-based pseudowires, if multiple equal-cost RSVP LSP's are available to reach a Layer 2 Circuit neighbor, one LSP is randomly used for forwarding. Use this feature to specify LSPs for specific L2circuit traffic to load-share the traffic in the MPLS core.
- Configuring an MPLS firewall filter on a switch that is deployed as an egress provider edge (PE) switch has no effect.
- Configuring the **revert-timer** statement at the **[edit protocols mpls]** hierarchy level has no effect.

## MPLS Limitations on QFX5100 and EX4600 Switches

- On a QFX5100 switch, you can observe traffic drop after changing your configuration to enable VLAN tagged for MPLS packets. As a result of packet capture, a QFX5100 switch can swap the wrong VLAN ID for MPLS packets.
- Even though both QFX5100 and EX4600 Switches use the same chipset, MPLS support differs. EX4600 switches support only basic MPLS functionality while QFX5100 switches support some of the more advanced features. See [“MPLS Feature Support on QFX Series and EX4600 Switches” on page 4934](#) for details.
- If you configure the BGP labeled unicast address family (using the **labeled-unicast** statement at the **[edit protocols bgp family inet]** hierarchy level) on a QFX switch or on an EX4600 switch deployed as a route reflector for BGP labeled routes, path selection will occur at the route reflector, and a single best path will be advertised. This will result in loss of BGP multipath information.
- Fast reroute is supported, however the **include-all** and **include-any** options for fast reroute are not supported. For more information, see [“Fast Reroute Overview” on page 4967](#).
- MPLS-based circuit cross-connects (CCC) is not supported—only circuit-based pseudowires are supported.
- MTU signaling in RSVP and discovery is supported in the control plane. However, this cannot be enforced in data plane.
- With L2 circuit-based pseudowires, if multiple equal-cost RSVP LSP's are available to reach a Layer 2 Circuit neighbor, one LSP is randomly used for forwarding. Use this feature to specify LSPs for specific L2circuit traffic to load-share the traffic in the MPLS core.
- Configuring an MPLS firewall filter on a switch that is deployed as an egress provider edge (PE) switch has no effect.
- Configuring the **revert-timer** statement at the **[edit protocols mpls]** hierarchy level has no effect.



- ECMP for MPLS is not supported on a QFX5100 switch.
- These are hardware limitations for both EX4600 MPLS and QFX5100 MPLS switches:
  - Push of a maximum of 3 labels is supported in MPLS edge switch if label swap is not done.
  - Push of a maximum of 2 labels is supported in MPLS edge switch if label swap is done.
  - Pop at line rate is supported for a maximum of 2 labels.
  - Global label space is supported but interface-specific label space is not supported.
  - ECMP based on incoming labels at LSR is not supported.
  - QFX switches with Broadcom chips do not support separate next hops for the same label with different S bits (S-0 and S-1). This includes QFX3500, QFX3600, QFX5100, and QFX5200 switches.
  - On QFX5100, the MPLS MTU command can cause unexpected behavior—this is due to SDK chipset limitations on this platform.
- These LDP features are not supported on both QFX5100 switches:
  - LDP multipoint
  - LDP link protection
  - LDP bidirectional forwarding detection (BFD)
  - LDP operation administration and management (OAM)
  - LDP multicast-only fast reroute (MoFRR)
  - LDP equal-cost multipath (ECMP)
- On QFX5100, IRB/L3-sub interfaces are not supported on the NNI port in L2circuit configuration. Because of this, a QFX5100 is unable to ping the neighbor IP of either a direct or local interface.

## MPLS Limitations on QFX5100 Virtual Chassis and Virtual Chassis Fabric

The following MPLS features are not supported by QFX5100 VC and QFX5100 VCF:

- Next-hop LSP Details in section 2.2.6.1
- BFD including BFD triggered FRR
- L2VPN based on BGP (VPWS / draft Kompella)
- VPLS
- Extended-vlan-ccc
- Pseudo-wire protection using Ethernet OAM
- Local switching of pseudo-wire

- Pseudowire fault detection based on VCCV
- QFX switches with Broadcom chips do not support separate next hops for the same label with different S bits (S-0 and S-1). This includes QFX3500, QFX3600, QFX5100, and QFX5200 switches.

## MPLS Limitations on QFX10000 Switches

- Carrier-of-carriers and inter-provider VPNs are not needed on QFX10000 Series switches.
- Configuring an MPLS firewall filter on a switch that is deployed as an egress provider edge (PE) switch has no effect.
- Configuring the **revert-timer** statement at the **[edit protocols mpls]** hierarchy level has no effect.
- These LDP features are not supported on QFX10000 switches:
  - LDP multipoint
  - LDP link protection
  - LDP bidirectional forwarding detection (BFD)
  - LDP operation administration and management (OAM)
  - LDP multicast-only fast reroute (MoFRR)
- On QFX10008, IRB/L3-sub interfaces are not supported on the NNI port in L2circuit configuration. Because of this, a QFX10008 is unable to ping the neighbor IP of either a direct or local interface.

### Related Documentation

- [MPLS Feature Support on QFX Series and EX4600 Switches on page 4934](#)

---

## Understanding MPLS Components for QFX Series and EX4600 Switches

MPLS devices include a number of components. While some components are required for all MPLS applications, others might not be, depending on the specific application.

This topic includes:

- [Provider Edge Switches on page 4946](#)
- [Provider Switch on page 4948](#)
- [Components Required for All Switches in the MPLS Network on page 4948](#)

### Provider Edge Switches

To implement MPLS on a network, you must configure two provider edge (PE) switches—an ingress PE switch and an egress PE switch. In addition, you must configure one or more provider switches as transit switches within the network to support the forwarding of MPLS packets.

The ingress PE switch (the entry point to the MPLS tunnel) receives a packet, analyzes it, and pushes an MPLS label onto it. This label places the packet in a forwarding equivalence class (FEC) and determines its handling and destination through the MPLS tunnel. The egress PE switch (the exit point from the MPLS tunnel) pops the MPLS label off the outgoing packet.

Within an MPLS tunnel, the network traffic is bidirectional. Therefore, each PE switch can be configured to be both an ingress switch and an egress switch, depending on the direction of the traffic.

The following MPLS components are configured on the PE switches but not on the provider switches:

- [MPLS Protocol and Label-Switched Paths on page 4947](#)
- [IP Over MPLS for Customer Edge Interfaces on page 4947](#)
- [BGP Layer 3 VPN Configuration on page 4947](#)
- [Routing Instances for Layer 3 VPN on page 4947](#)
- [Routing Instances for Layer 2 VPN and Layer 3 VPN on page 4948](#)
- [Ethernet Encapsulation for Layer 2 VPN on page 4948](#)

---

### MPLS Protocol and Label-Switched Paths

Each PE switch must be configured to support the MPLS protocol. You must also configure label-switched paths (LSPs) at the `[edit protocols mpls]` hierarchy level.

---

### IP Over MPLS for Customer Edge Interfaces

You can configure the customer edge interfaces of the PE switches for IP over MPLS using a Layer 3 interface and a static route from the ingress PE switch to the egress PE switch. See “[Configuring MPLS on Provider Edge Switches](#)” on page 4981.

---

### BGP Layer 3 VPN Configuration

If you are implementing a Layer 3 virtual private network (VPN), you must configure the BGP routing protocol on the PE switches.

---

### Routing Instances for Layer 3 VPN

If you are implementing a Layer 3 VPN, you must configure a routing instance. A routing instance is a collection of routing tables, interfaces, and routing protocol parameters. The set of interfaces belongs to the routing tables, and the routing protocol parameters control the information in the routing tables.

QFX Series and EX4600 devices support VPN routing and forwarding (VRF) routing instances for Layer 3 VPNs.

Each routing instance has a unique name and a corresponding IP unicast table. For example, if you configure a routing instance with the name **my-instance**, its corresponding IP unicast table will be **my-instance.inet.0**. All routes for **my-instance** are installed in **my-instance.inet.0**.

### Routing Instances for Layer 2 VPN and Layer 3 VPN

---

If you are implementing a Layer 2 VPN or a Layer 3 VPN, you must configure a routing instance. A routing instance is a collection of routing tables, interfaces, and routing protocol parameters. The set of interfaces belongs to the routing tables, and the routing protocol parameters control the information in the routing tables.

QFX Series devices support the following types of routing instances:

- Layer 2 VPN—To support a Layer 2 VPN
- VPN routing and forwarding (VRF)—To support a Layer 3 VPN

Each routing instance has a unique name and a corresponding IP unicast table. For example, if you configure a routing instance with the name **my-instance**, its corresponding IP unicast table will be **my-instance.inet.0**. All routes for **my-instance** are installed in **my-instance.inet.0**.

### Ethernet Encapsulation for Layer 2 VPN

---

If you are implementing a Layer 2 VPN, you must also configure the physical layer encapsulation type on the customer edge interface and within the routing instance.

## Provider Switch

You must configure one or more provider switches as transit switches within the network to support the forwarding of MPLS packets. You can add provider switches without changing the configuration of the PE switches.

A provider switch does not analyze packets. It refers to an MPLS label forwarding table and swaps one label for another. The new label determines the next hop along the MPLS tunnel. A provider switch cannot perform push or pop operations.

## Components Required for All Switches in the MPLS Network

The following MPLS components are configured on both the PE switches and the provider switches:

- [Interior Gateway Protocol on page 4948](#)
- [Traffic Engineering on page 4949](#)
- [MPLS Protocol on page 4949](#)
- [RSVP on page 4949](#)
- [Family mpls on page 4950](#)

### Interior Gateway Protocol

---

MPLS works in coordination with OSPF as the interior gateway protocol (IGP). Therefore, you must configure OSPF as the IGP on the loopback interface and CE-facing interfaces of both the PE switches and the provider switches.

The CE-facing interfaces can be either Gigabit Ethernet or 10-Gigabit Ethernet interfaces, and they can be configured as either individual interfaces or as aggregated Ethernet interfaces.



**NOTE:** The CE-facing interfaces cannot be configured with VLAN tagging or a VLAN ID. When you configure them to belong to family `mpls`, they are removed from the default VLAN if they were members of that VLAN. They operate as an exclusive tunnel for MPLS traffic.

---

## Traffic Engineering

Traffic engineering maps traffic flows onto an existing physical topology and provides the ability to move traffic flow away from the shortest path selected by the IGP and to a potentially less congested physical path across a network.

Traffic engineering enables the selection of specific end-to-end paths to send given types of traffic through your network. You must configure OSPF traffic engineering on the PE switches and the provider switches.

---

## MPLS Protocol

You must enable the MPLS protocol on all switches that participate in the MPLS network and apply it to the core interfaces of both the PE and provider switches. You do not need to apply it to the loopback interface because the MPLS protocol uses the framework established by the RSVP signaling protocol to create LSPs. On the PE switches, the configuration of the MPLS protocol must also include the definition of an LSP.

---

## RSVP

RSVP is a signaling protocol that allocates and distributes labels throughout an MPLS network. RSVP sets up unidirectional paths between the ingress PE switch and the egress PE switch. RSVP makes the LSPs dynamic; it can detect topology changes and outages and establish new LSPs to allow traffic to move around a failure.

You must enable RSVP and apply it to the loopback interface and the core interface of both the PE and provider switches. The path message contains the configured information about the resources required for the LSP to be established.

When the egress PE switch receives the path message, it sends a reservation message back to the ingress PE switch. This reservation message is passed along from switch to switch along the same path as the original path message. Once the ingress PE switch receives this reservation message, an RSVP path is established.

The established LSP stays active as long as the RSVP session remains active. RSVP continues activity through the transmissions and responses to RSVP path and reservation messages. If the messages stop for three minutes, the RSVP session terminates and the LSP is lost.

RSVP runs as a separate software process in Junos OS and is not in the packet-forwarding path.

## Family mpls

---

You must configure the core interfaces used for MPLS traffic to belong to **family mpls**.



**NOTE:** You can enable **family mpls** on either individual interfaces or on aggregated Ethernet interfaces. You cannot enable it on tagged VLAN interfaces.

### Related Documentation

- [MPLS Feature Support on QFX Series and EX4600 Switches on page 4934](#)
- [Understanding Using MPLS-Based Layer 3 VPNs on Switches on page 4963](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4959](#)
- [Configuring MPLS on Provider Edge Switches on page 4981](#)
- [Configuring MPLS on Provider Switches on page 4985](#)
- [Configuring Rewrite Rules for MPLS EXP Classifiers on page 4992](#)
- [Configuring a Global MPLS EXP Classifier on page 4977](#)
- [Configuring Ethernet over MPLS \(L2 Circuit\) on page 4972](#)
- *Junos OS MPLS Applications Library for Routing Devices*
- *Junos OS VPNs Library for Routing Devices*

## Understanding MPLS Label Operations

---

In the traditional packet-forwarding paradigm, as a packet travels from one switch to the next, an independent forwarding decision is made at each hop. The IP network header is analyzed and the next hop is chosen based on this analysis and on the information in the routing table. In an MPLS environment, the analysis of the packet header is made only once, when a packet enters the MPLS tunnel (that is, the path used for MPLS traffic).

When an IP packet enters a label-switched path (LSP), the ingress provider edge (PE) switch examines the packet and assigns it a label based on its destination, placing the label in the packet's header. The label transforms the packet from one that is forwarded based on its IP routing information to one that is forwarded based on information associated with the label. The packet is then forwarded to the next provider switch in the LSP. This switch and all subsequent switches in the LSP do not examine any of the IP routing information in the labeled packet. Rather, they use the label to look up information in their label forwarding table. They then replace the old label with a new label and forward the packet to the next switch in the path. When the packet reaches the egress PE switch, the label is removed, and the packet again becomes a native IP packet and is forwarded based on its IP routing information.

This topic describes:


- [MPLS Label-Switched Paths and MPLS Labels on page 4951](#)
- [Reserved Labels on page 4952](#)
- [MPLS Label Operations on page 4952](#)
- [Penultimate-Hop Popping and Ultimate-Hop Popping on page 4954](#)

### MPLS Label-Switched Paths and MPLS Labels

When a packet enters the MPLS network, it is assigned to an LSP. Each LSP is identified by a label, which is a short (20-bit), fixed-length value at the front of the MPLS label (32 bits). Labels are used as lookup indexes for the label forwarding table. For each label, this table stores forwarding information. Because no additional parsing or lookup is done on the encapsulated packet, MPLS supports the transmission of any other protocols within the packet payload.

[Figure 139](#) shows the encoding of a single label. The encoding appears after data link layer headers, but before any network layer header.

Figure 139: Label Encoding



## Reserved Labels

Labels range from 0 through 1,048,575. Labels 0 through 999,999 are for internal use.

Some of the reserved labels (in the range 0 through 15) have well-defined meanings.

The following reserved labels are used by QFX Series and EX4600 devices:

- 0, IPv4 Explicit Null label—This value is valid only when it is the sole label entry (no label stacking). It indicates that the label must be popped on receipt. Forwarding continues based on the IP version 4 (IPv4) packet.
- 1, Router Alert label—When a packet is received with a top label value of 1, it is delivered to the local software module for processing.
- 3, Implicit Null label—This label is used in the signaling protocol (RSVP) only to request label popping by the downstream switch. It never actually appears in the encapsulation. Labels with a value of 3 must not be used in the data packet as real labels. No payload type (IPv4 or IPv6) is implied with this label.

## MPLS Label Operations

QFX Series and EX4600 devices support the following MPLS label operations:

- Push
- Pop
- Swap





**NOTE:** There is a limit with regard to the number of labels that QFX and EX4600 devices can affix (push operations) to the label stack or remove (pop operations) from the label stack.

- For Push operations—As many as three labels are supported.
- For Pop operations—As many as three labels are supported.

The push operation affixes a new label to the top of the IP packet. For IPv4 packets, the new label is the first label. The time to live (TTL) field value in the packet header is derived from the IP packet header. The push operation cannot be applied to a packet that already has an MPLS label.

The pop operation removes a label from the beginning of the packet. Once the label is removed, the TTL is copied from the label into the IP packet header, and the underlying IP packet is forwarded as a native IP packet.

The swap operation removes an existing MPLS label from an IP packet and replaces it with a new MPLS label, based on the following:

- Incoming interface
- Label
- Label forwarding table

Figure 140 shows an IP packet without a label arriving on the customer edge interface (ge-0/0/1) of the ingress PE switch. The ingress PE switch examines the packet and identifies that packet's destination as the egress PE switch. The ingress PE switch applies label 100 to the packet and sends the MPLS packet to its outgoing MPLS core interface (ge-0/0/5). The MPLS packet is transmitted on the MPLS tunnel through the provider switch, where it arrives at interface ge-0/0/5 with label 100. The provider switch swaps label 100 with label 200 and forwards the MPLS packet through its core interface (ge-0/0/7) to the next hop on the tunnel, which is the egress PE switch. The egress PE switch receives the MPLS packet through its core interface (ge-0/0/7), removes the MPLS label, and sends the IP packet out of its customer edge interface (ge-0/0/1) to a destination that is beyond the tunnel.

**Figure 140: MPLS Label Swapping**




Figure 140 shows the path of a packet as it passes in one direction from the ingress PE switch to the egress PE switch. However, the MPLS configuration also allows traffic to travel in the reverse direction. Thus, each PE switch operates as both an ingress switch and an egress switch.

## Penultimate-Hop Popping and Ultimate-Hop Popping

The switches enable penultimate-hop popping (PHP) by default with IP over MPLS configurations. With PHP, the penultimate provider switch is responsible for popping the MPLS label and forwarding the traffic to the egress PE switch. The egress PE switch then performs an IP route lookup and forwards the traffic. This reduces the processing load on the egress PE switch, because it is not responsible for popping the MPLS label.

- The default advertised label is label 3 (Implicit Null label). If label 3 is advertised, the penultimate-hop switch removes the label and sends the packet to the egress PE switch.
- If ultimate-hop popping is enabled, label 0 (IPv4 Explicit Null label) is advertised and the egress PE switch of the LSP removes the label.

### Related Documentation

- [Understanding MPLS Components for QFX Series and EX4600 Switches on page 4946](#)
- [Configuring MPLS on Provider Edge Switches on page 4981](#)
- [Configuring MPLS on Provider Switches on page 4985](#)
- *Junos OS MPLS Applications Library for Routing Devices*
- *Junos OS VPNs Library for Routing Devices*

---

## Understanding BGP

---

BGP is an exterior gateway protocol (EGP) that is used to exchange routing information among routers in different autonomous systems (ASs). BGP routing information includes the complete route to each destination. BGP uses the routing information to maintain a database of network reachability information, which it exchanges with other BGP systems. BGP uses the network reachability information to construct a graph of AS connectivity, which enables BGP to remove routing loops and enforce policy decisions at the AS level.

Multiprotocol BGP (MBGP) extensions enable BGP to support IP version 6 (IPv6). MBGP defines the attributes `MP_REACH_NLRI` and `MP_UNREACH_NLRI`, which are used to carry IPv6 reachability information. Network layer reachability information (NLRI) update messages carry IPv6 address prefixes of feasible routes.

BGP allows for policy-based routing. You can use routing policies to choose among multiple paths to a destination and to control the redistribution of routing information.

BGP uses TCP as its transport protocol, using port 179 for establishing connections. Running over a reliable transport protocol eliminates the need for BGP to implement update fragmentation, retransmission, acknowledgment, and sequencing.

The Junos OS routing protocol software supports BGP version 4. This version of BGP adds support for Classless Interdomain Routing (CIDR), which eliminates the concept of network classes. Instead of assuming which bits of an address represent the network by looking at the first octet, CIDR allows you to explicitly specify the number of bits in the network address, thus providing a means to decrease the size of the routing tables. BGP version 4 also supports aggregation of routes, including the aggregation of AS paths.

This section discusses the following topics:

- [Autonomous Systems on page 4955](#)
- [AS Paths and Attributes on page 4955](#)
- [External and Internal BGP on page 4956](#)
- [Multiple Instances of BGP on page 4956](#)

### Autonomous Systems

An *autonomous system* (AS) is a set of routers that are under a single technical administration and normally use a single interior gateway protocol and a common set of metrics to propagate routing information within the set of routers. To other ASs, an AS appears to have a single, coherent interior routing plan and presents a consistent picture of what destinations are reachable through it.

### AS Paths and Attributes


The routing information that BGP systems exchange includes the complete route to each destination, as well as additional information about the route. The route to each destination is called the *AS path*, and the additional route information is included in *path attributes*. BGP uses the AS path and the path attributes to completely determine the network topology. Once BGP understands the topology, it can detect and eliminate

routing loops and select among groups of routes to enforce administrative preferences and routing policy decisions.

## External and Internal BGP

BGP supports two types of exchanges of routing information: exchanges among different ASs and exchanges within a single AS. When used among ASs, BGP is called *external BGP* (EBGP) and BGP sessions perform *inter-AS routing*. When used within an AS, BGP is called *internal BGP* (IBGP) and BGP sessions perform *intra-AS routing*. Figure 56 illustrates ASs, IBGP, and EBGP.

Figure 141: ASs, EBGP, and IBGP



A BGP system shares network reachability information with adjacent BGP systems, which are referred to as *neighbors* or *peers*.

BGP systems are arranged into *groups*. In an IBGP group, all peers in the group—called *internal peers*—are in the same AS. Internal peers can be anywhere in the local AS and do not have to be directly connected to one another. Internal groups use routes from an IGP to resolve forwarding addresses. They also propagate external routes among all other internal routers running IBGP, computing the next hop by taking the BGP next hop received with the route and resolving it using information from one of the interior gateway protocols.

In an EBGP group, the peers in the group—called *external peers*—are in different ASs and normally share a subnet. In an external group, the next hop is computed with respect to the interface that is shared between the external peer and the local router.

## Multiple Instances of BGP

You can configure multiple instances of BGP at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols]

Multiple instances of BGP are primarily used for Layer 3 VPN support.

IGP peers and external BGP (EBGP) peers (both nonmultihop and multihop) are all supported for routing instances. BGP peering is established over one of the interfaces configured under the **routing-instances** hierarchy.



**NOTE:** When a BGP neighbor sends BGP messages to the local routing device, the incoming interface on which these messages are received must be configured in the same routing instance that the BGP neighbor configuration exists in. This is true for neighbors that are a single hop away or multiple hops away.

Routes learned from the BGP peer are added to the **instance-name.inet.0** table by default. You can configure import and export policies to control the flow of information into and out of the instance routing table.

For Layer 3 VPN support, configure BGP on the provider edge (PE) router to receive routes from the customer edge (CE) router and to send the instances' routes to the CE router if necessary. You can use multiple instances of BGP to maintain separate per-site forwarding tables for keeping VPN traffic separate on the PE router.

You can configure import and export policies that allow the service provider to control and rate-limit traffic to and from the customer.

You can configure an EBGP multihop session for a VRF routing instance. Also, you can set up the EBGP peer between the PE and CE routers by using the loopback address of the CE router instead of the interface addresses.

**Related  
Documentation**

- [BGP Routes Overview on page 3596](#)
- [BGP Messages Overview on page 3597](#)

## IPv6 Layer 3 VPNs

The interfaces between the PE and CE routers of a Layer 3 VPN can be configured to carry IP version 6 (IPv6) traffic. IP allows numerous nodes on different networks to interoperate seamlessly. IPv4 is currently used in intranets and private networks, as well as the Internet. IPv6 is the successor to IPv4, and is based for the most part on IPv4.

In the Juniper Networks implementation of IPv6, the service provider implements an MPLS-enabled IPv4 backbone to provide VPN service for IPv6 customers. The PE routers have both IPv4 and IPv6 capabilities. They maintain IPv6 VPN routing and forwarding (VRF) tables for their IPv6 sites and encapsulate IPv6 traffic in MPLS frames that are then sent into the MPLS core network.

IPv6 for Layer 3 VPNs is supported for BGP and for static routes.

IPv6 over Layer 3 VPNs is described in RFC 4659, *BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN*.

For more information about IPv6, see the *Junos OS Routing Protocols Library for Routing Devices*.

## Ethernet Pseudowire Overview

---

An Ethernet pseudowire is used to carry Ethernet or 802.3 Protocol Data Units (PDUs) over an MPLS network enabling service providers to offer emulated Ethernet services over existing MPLS networks. Ethernet or 802.3 PDUs are encapsulated within the pseudowire to provide a point-to-point Ethernet service. For the point-to-point Ethernet service, the following fault management features are supported:

- The IEEE 802.3ah standard for Operation, Administration, and Management (OAM). You can configure IEEE 802.3ah OAM link-fault management on Ethernet point-to-point direct links or links across Ethernet repeaters.

Ethernet OAM link-fault management can be used for physical link-level fault detection and management. It uses a new, optional sublayer in the data link layer of the OSI model. Ethernet OAM can be implemented on any full-duplex point-to-point or emulated point-to-point Ethernet link. A system-wide implementation is not required; OAM can be deployed on particular interfaces of a router. Transmitted Ethernet OAM messages or OAM PDUs are of standard length, untagged Ethernet frames within the normal frame length limits in the range 64–1518 bytes.

- Ethernet connectivity fault management (CFM) to monitor the physical link between two routers.
  - Connection protection using the continuity check protocol for fault monitoring . The continuity check protocol is a neighbor discovery and health check protocol that discovers and maintains adjacencies at the VLAN or link level.
  - Path protection using the linktrace protocol for path discovery and fault verification . Similar to IP traceroute, the linktrace protocol maps the path taken to a destination MAC address through one or more bridged networks between the source and destination.

### Related Documentation

- *Configuring IEEE 802.3ah OAM Link-Fault Management*
- *Pseudowire Overview for ACX Series Universal Access Routers*
- *TDM Pseudowires Overview*
- *ATM Pseudowire Overview*

## Understanding CoS MPLS EXP Classifiers and Rewrite Rules

You can use class of service (CoS) within MPLS networks to prioritize certain types of traffic during periods of congestion by applying packet classifiers and rewrite rules to the MPLS traffic. MPLS classifiers are global and apply to all interfaces configured as **family mpls** interfaces.

When a packet enters a customer-edge interface on the ingress provider edge (PE) switch, the switch associates the packet with a particular CoS servicing level before placing the packet onto the label-switched path (LSP). The switches within the LSP utilize the CoS value set at the ingress PE switch to determine the CoS service level. The CoS value embedded in the classifier is translated and encoded in the MPLS header by means of the experimental (EXP) bits.

EXP classifiers map incoming MPLS packets to a forwarding class and a loss priority, and assign MPLS packets to output queues based on the forwarding class mapping. EXP classifiers are behavior aggregate (BA) classifiers.

EXP rewrite rules change (rewrite) the CoS value of the EXP bits in outgoing packets on the egress queues of the switch so that the new (rewritten) value matches the policies of a targeted peer. Policy matching allows the downstream routing platform or switch in a neighboring network to classify each packet into the appropriate service group.



**NOTE:** On QFX5200, QFX5100, QFX3500, QF3600, and EX4600 switches, and on QFabric systems, there is no default EXP classifier. If you want to classify incoming MPLS packets using the EXP bits, you must configure a global EXP classifier. The global EXP classifier applies to all MPLS traffic on interfaces configured as **family mpls**.

On QFX10000 switches, there is a no default EXP classifier. If you want to classify incoming MPLS packets using the EXP bits, you must configure EXP classifiers and apply them to logical interfaces configured as **family mpls**. (You cannot apply classifiers to physical interfaces.). You can configure up to 64 EXP classifiers.

There is no default EXP rewrite rule. If you want to rewrite the EXP bit value at the egress interface, you must configure EXP rewrite rules and apply them to logical interfaces.

EXP classifiers and rewrite rules are applied only to interfaces that are configured as **family mpls** (for example, set interfaces **xe-0/0/35 unit 0 family mpls.**)

This topic includes:

- [EXP Classifiers on page 4960](#)
- [EXP Rewrite Rules on page 4961](#)
- [Schedulers on page 4962](#)

## EXP Classifiers

On QFX5200, QFX5100, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems, unlike DSCP and IEEE 802.1p BA classifiers, EXP classifiers are global to the switch and apply to all switch interfaces that are configured as **family mpls**. On QFX10000 switches, you apply EXP classifiers to individual logical interfaces, and different interfaces can use different EXP classifiers.

When you configure and apply an EXP classifier, MPLS traffic on all **family mpls** interfaces uses the EXP classifier, even on interfaces that also have a fixed classifier. If an interface has both an EXP classifier and a fixed classifier, the EXP classifier is applied to MPLS traffic and the fixed classifier is applied to all other traffic.

Also unlike DSCP and IEEE 802.1p BA classifiers, there is no default EXP classifier. If you want to classify MPLS traffic based on the EXP bits, you must explicitly configure an EXP classifier and apply it to the switch interfaces. Each EXP classifier has eight entries that correspond to the eight EXP CoS values (0 through 7, which correspond to CoS bits 000 through 111).

You can configure up to 64 EXP classifiers.

However, on QFX5200, QFX5100, EX4600, and legacy CLI switches, the switch uses only one MPLS EXP classifier as a global classifier on all interfaces. After you configure an MPLS EXP classifier, you can configure that classifier as the global EXP classifier by including the EXP classifier in the **[edit class-of-service system-defaults classifiers exp]** hierarchy level. All switch interfaces configured as **family mpls** use the global EXP classifier to classify MPLS traffic.

On these switches, only one EXP classifier can be configured as the global EXP classifier at any time. If you want to change the global EXP classifier, delete the global EXP classifier configuration (use the **user@switch# delete class-of-service system-defaults classifiers exp** configuration statement), then configure the new global EXP classifier.

QFX10000 switches do not support global EXP classifiers. You can configure one EXP classifier and apply it to multiple logical interfaces, or configure multiple EXP classifiers and apply different EXP classifiers to different logical interfaces.

If an EXP classifier is not configured, then if a fixed classifier is applied to the interface, the MPLS traffic uses the fixed classifier. (Switches that have a default EXP classifier use the default classifier.) If no EXP classifier and no fixed classifier are applied to the interface, MPLS traffic is treated as best-effort traffic using the 802.1 default untrusted classifier. DSCP classifiers are not applied to MPLS traffic.

On QFX5200, QFX5100, EX4600, and legacy CLI switches, because the EXP classifier is global, you cannot configure some ports to use a fixed IEEE 802.1p classifier for MPLS traffic on some interfaces and the global EXP classifier for MPLS traffic on other interfaces. When you configure a global EXP classifier, all MPLS traffic on all interfaces uses the EXP classifier.





**NOTE:** The switch uses only the outermost label of incoming EXP packets for classification.



**NOTE:** MPLS packets with 802.1Q tags are not supported.

## EXP Rewrite Rules

As MPLS packets enter or exit a network, edge switches might be required to alter the class-of-service (CoS) settings of the packets. EXP rewrite rules set the value of the EXP CoS bits within the header of the outgoing MPLS packet on **family mpls** interfaces. Each rewrite rule reads the current forwarding class and loss priority associated with the packet, locates the chosen CoS value from a table, and writes that CoS value into the packet header, replacing the old CoS value. EXP rewrite rules apply only to MPLS traffic.

EXP rewrite rules apply only to logical interfaces. You cannot apply EXP rewrite rules to physical interfaces.

There are no default EXP rewrite rules. If you want to rewrite the EXP value in MPLS packets, you must configure EXP rewrite rules and apply them to logical interfaces. If no rewrite rules are applied, all MPLS labels that are pushed have a value of zero (0). The EXP value remains unchanged on MPLS labels that are swapped.

You can configure up to 64 EXP rewrite rules, but you can only apply 16 EXP rewrite rules at any time on the switch. On a given logical interface, all pushed MPLS labels have the same EXP rewrite rule applied to them. You can apply different EXP rewrite rules to different logical interfaces on the same physical interface.

You can apply an EXP rewrite rule to an interface that has a DSCP, DSCP IPv6, or IEEE 802.1p rewrite rule. Only MPLS traffic uses the EXP rewrite rule. MPLS traffic does not use DSCP or DSCP IPv6 rewrite rules.

If the switch is performing penultimate hop popping (PHP), EXP rewrite rules do not take effect. If both an EXP classifier and an EXP rewrite rule are configured on the switch, then the EXP value from the last popped label is copied into the inner label. If either an EXP classifier or an EXP rewrite rule (but not both) is configured on the switch, then the inner label EXP value is sent unchanged.



**NOTE:** On each physical interface, either all forwarding classes that are being used on the interface must have rewrite rules configured or no forwarding classes that are being used on the interface can have rewrite rules configured. On any physical port, do not mix forwarding classes with rewrite rules and forwarding classes without rewrite rules.

## Schedulers

The schedulers for using CoS with MPLS are the same as for the other CoS configurations on the switch. Default schedulers are provided only for the best-effort, fcoe, no-loss, and network-control default forwarding classes. If you configure a custom forwarding class for MPLS traffic, you need to configure a scheduler to support that forwarding class and provide bandwidth to that forwarding class.

**Related Documentation**

- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 6672](#)

---

## Understanding Ethernet-over-MPLS (L2 Circuit)

Ethernet-over-MPLS allows sending Layer 2 (L2) Ethernet frames transparently over MPLS. Ethernet-over-MPLS uses a tunneling mechanism for Ethernet traffic through an MPLS-enabled Layer 3 core. It encapsulates Ethernet protocol data units (PDUs) inside MPLS packets and forwards the packets, using label stacking, across the MPLS network. This technology has applications in service provider, enterprise and data center environments. For disaster recovery purposes, data centers are hosted in multiple sites that are geographically distant and interconnected using a WAN network.



**NOTE:** A Layer 2 circuit is similar to a circuit cross-connect (CCC), except that multiple Layer 2 circuits can be transported over a single label-switched path (LSP) tunnel between two provider edge (PE) routers. In contrast, each CCC requires a dedicated LSP.

- 
- [Ethernet-over-MPLS in Data Centers on page 4962](#)

## Ethernet-over-MPLS in Data Centers

For disaster recovery purposes, data centers are hosted in multiple sites that are geographically distant and interconnected using a WAN network. These data centers require L2 connectivity between them for the following reasons:

- To replicate the storage over Fiber Channel IP (FCIP). FCIP works only on the same broadcast domain.
- To run a dynamic routing protocol between the sites.
- To support High Availability clusters that interconnect the nodes hosted in the various data centers.

**Related Documentation**

- [Configuring Ethernet over MPLS \(L2 Circuit\) on page 4972](#)

---

## Understanding Using MPLS-Based Layer 3 VPNs on Switches

---

On the QFX Series switches and on EX4600 switches, you can use MPLS-based Layer 3 virtual private networks (VPNs) to securely connect geographically diverse sites across an MPLS network. MPLS services can be used to connect various sites to a backbone network and to ensure better performance for low-latency applications such as voice over IP (VoIP) and other business-critical functions.

A VPN uses a public telecommunications infrastructure, such as the Internet, to provide remote offices or individual users with secure access to their organization's network. VPNs are designed to provide the same level of performance and security as privately owned or leased networks but without the attendant costs.

This topic describes:

- [MPLS-Based Layer 3 VPNs on page 4963](#)

### MPLS-Based Layer 3 VPNs

In Junos OS, Layer 3 VPNs are based on RFC 4364, *BGP/MPLS IP Virtual Private Networks*. RFC 4364 defines a mechanism by which service providers can use their IP backbones to provide VPN services to their customers. A Layer 3 VPN is a set of sites that share common routing information and whose connectivity is controlled by a collection of policies. The sites that make up a Layer 3 VPN are connected over a provider's existing public Internet backbone.

Customer networks, because they are private, can use either public or private addresses, as defined in RFC 1918, *Address Allocation for Private Internets*. When customer networks that use private addresses connect to the public Internet infrastructure, the private addresses might overlap with the same private addresses used by other network users. BGP/MPLS VPNs solve this problem by adding a VPN identifier prefix to each address from a particular VPN site, thereby creating an address that is unique both within the VPN and on the public Internet. In addition, each VPN has its own VPN-specific routing table that contains the routing information for that VPN only. Two different VPNs can use overlapping addresses. Each route within a VPN is assigned an MPLS label (for example, MPLS-ARCH, MPLS-BGP, or MPLS-ENCAPS). When BGP distributes a VPN route, it also distributes an MPLS label for that route. Before a customer data packet travels across the service provider's backbone, it is encapsulated along with the MPLS label that corresponds to the route within the customer's VPN that is the best match based on the packet's destination address. This MPLS packet is further encapsulated with another MPLS label or with an IP, so that it gets tunneled across the backbone to the egress provider edge (PE) switch. Thus, the backbone core switches do not need to know the VPN routes.

QFX5100 switches also support interprovider VPNs, and carrier-of-carriers VPNs. For more information, see ["Interprovider and Carrier-of-Carriers VPNs" on page 4966](#)

#### Related Documentation

- [Understanding MPLS Label Operations on page 4951](#)
- [Understanding MPLS Components for QFX Series and EX4600 Switches on page 4946](#)

- *Example: Configuring MPLS-Based Layer 2 VPNs*
- *Example: Configuring MPLS-Based Layer 3 VPNs on EX Series Switches*

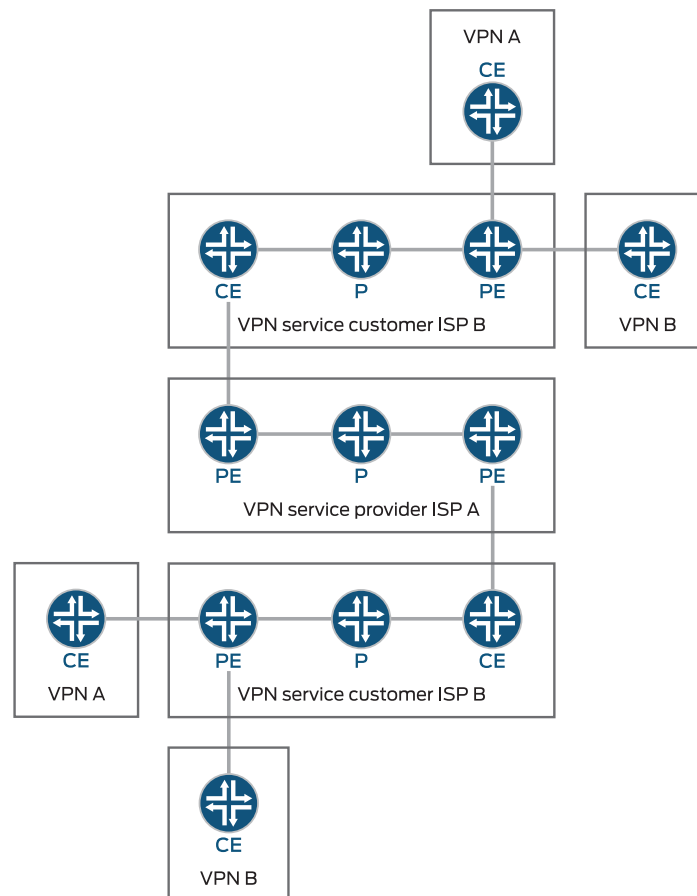
## Carrier-of-Carriers VPNs

The customer of a VPN service provider might be a service provider for the end customer. The following are the two main types of carrier-of-carriers VPNs (as described in RFC 4364):

- “Internet Service Provider as the Customer” on page 4966—The VPN customer is an ISP that uses the VPN service provider’s network to connect its geographically disparate regional networks. The customer does not have to configure MPLS within its regional networks.
- “VPN Service Provider as the Customer” on page 4966—The VPN customer is itself a VPN service provider offering VPN service to its customers. The carrier-of-carriers VPN service customer relies on the backbone VPN service provider for inter-site connectivity. The customer VPN service provider is required to run MPLS within its regional networks.

Figure 142 illustrates the network architecture used for a carrier-of-carriers VPN service.

Figure 142: Carrier-of-Carriers VPN Architecture



This topic covers the following:

- [Internet Service Provider as the Customer on page 4966](#)
- [VPN Service Provider as the Customer on page 4966](#)

## Internet Service Provider as the Customer

In this type of carrier-of-carriers VPN configuration, ISP A configures its network to provide Internet service to ISP B. ISP B provides the connection to the customer wanting Internet service, but the actual Internet service is provided by ISP A.

This type of carrier-of-carriers VPN configuration has the following characteristics:

- The carrier-of-carriers VPN service customer (ISP B) does not need to configure MPLS on its network.
- The carrier-of-carriers VPN service provider (ISP A) must configure MPLS on its network.
- MPLS must also be configured on the CE routers and PE routers connected together in the carrier-of-carriers VPN service customer's and carrier-of-carriers VPN service provider's networks.

## VPN Service Provider as the Customer

A VPN service provider can have customers that are themselves VPN service providers. In this type of configuration, also called a hierarchical or recursive VPN, the customer VPN service provider's VPN-IPv4 routes are considered external routes, and the backbone VPN service provider does not import them into its VRF table. The backbone VPN service provider imports only the customer VPN service provider's internal routes into its VRF table.

The similarities and differences between interprovider and carrier-of-carriers VPNs are shown in [Table 377](#).

**Table 377: Comparison of Interprovider and Carrier-of-Carriers VPNs**

| Feature                                | ISP Customer      | VPN Service Provider Customer                         |
|----------------------------------------|-------------------|-------------------------------------------------------|
| Customer edge device                   | AS border router  | PE router                                             |
| IBGP sessions                          | Carry IPv4 routes | Carry external VPN-IPv4 routes with associated labels |
| Forwarding within the customer network | MPLS is optional  | MPLS is required                                      |

## Interprovider and Carrier-of-Carriers VPNs

All interprovider and carrier-of-carriers VPNs share the following characteristics:

- Each interprovider or carrier-of-carriers VPN customer must distinguish between internal and external customer routes.

- Internal customer routes must be maintained by the VPN service provider in its PE routers.
- External customer routes are carried only by the customer's routing platforms, not by the VPN service provider's routing platforms.

The key difference between interprovider and carrier-of-carriers VPNs is whether the customer sites belong to the same AS or to separate ASs:

- *Understanding Interprovider VPNs*—The customer sites belong to different ASs. You need to configure EBGp to exchange the customer's external routes.
- [“Understanding Carrier-of-Carriers VPNs” on page 4965](#)—The customer sites belong to the same AS. You need to configure IBGP to exchange the customer's external routes.

In general, each service provider in a VPN hierarchy is required to maintain its own internal routes in its P routers, and the internal routes of its customers in its PE routers. By recursively applying this rule, it is possible to create a hierarchy of VPNs.

The following are definitions of the types of PE routers specific to interprovider and carrier-of-carriers VPNs:

- The AS border router is located at the AS border and handles traffic leaving and entering the AS.
- The end PE router is the PE router in the customer VPN; it is connected to the CE router at the end customer's site.


**Related  
Documentation**

- [Understanding Carrier-of-Carriers VPNs on page 4965](#)
- *Understanding Interprovider VPNs*


## Fast Reroute Overview

Fast reroute provides redundancy for an LSP path. When you enable fast reroute, detours are precomputed and preestablished along the LSP. In case of a network failure on the current LSP path, traffic is quickly routed to one of the detours. [Figure 143](#) illustrates an LSP from Router A to Router F, showing the established detours. Each detour is established by an upstream node to avoid the link toward the immediate downstream node and the immediate downstream node itself. Each detour might traverse through one or more label-switched routers (or switches) that are not shown in the figure.

Fast reroute protects traffic against any single point of failure between the ingress and egress routers (or switches). If there are multiple failures along an LSP, fast reroute itself might fail. Also, fast reroute does not protect against failure of the ingress or egress routers.

**Figure 143: Detours Established for an LSP Using Fast Reroute**

If a node detects that a downstream link has failed (using a link-layer-specific liveness detection mechanism) or that a downstream node has failed (for example, using the RSVP neighbor hello protocol), the node quickly switches the traffic to the detour and, at the same time, signals the ingress router about the link or node failure. [Figure 144](#) illustrates the detour taken when the link between Router B and Router C fails.

**Figure 144: Detour After the Link from Router B to Router C Fails**

If the network topology is not rich enough (there are not enough routers with sufficient links to other routers), some of the detours might not succeed. For example, the detour from Router A to Router C in [Figure 143](#) cannot traverse link A-B and Router B. If such a path is not possible, the detour does not occur.

Note that after the node switches traffic to the detour, it might switch the traffic again to a newly calculated detour soon after. This is because the initial detour route might not be the best route. To make rerouting as fast as possible, the node switches traffic onto the initial detour without first verifying that the detour is valid. Once the switch is made, the node recomputes the detour. If the node determines that the initial detour is still valid, traffic continues to flow over this detour. If the node determines that the initial detour is no longer valid, it again switches the traffic to a newly computed detour.



**NOTE:** If you issue `show` commands after the node has switched traffic to the initial detour, the node might indicate that the traffic is still flowing over the original LSP. This situation is temporary and should correct itself quickly.

The time required for a fast-rerouting detour to take effect depends on two independent time intervals:

- Amount of time to detect that there is a link or node failure—This interval depends greatly on the link layer in use and the nature of the failure. For example, failure detection on an SONET/SDH link typically is much faster than on a Gigabit Ethernet link, and both are much faster than detection of a router failure.




- Amount of time required to splice the traffic onto the detour—This operation is performed by the Packet Forwarding Engine, which requires little time to splice traffic onto the detour. The time needed can vary depending on the number of LSPs being switched to detours.

Fast reroute is a short-term patch to reduce packet loss. Because detour computation might not reserve adequate bandwidth, the detours might introduce congestion on the alternate links. The ingress router is the only router that is fully aware of LSP policy constraints and, therefore, is the only router able to come up with adequate long-term alternate paths.

Detours are created by use of RSVP and, like all RSVP sessions, they require extra state and overhead in the network. For this reason, each node establishes at most one detour for each LSP that has fast reroute enabled. Creating more than one detour for each LSP increases the overhead, but serves no practical purpose.

To reduce network overhead further, each detour attempts to merge back into the LSP as soon as possible after the failed node or link. If you can consider an LSP that travels through  $n$  router nodes, it is possible to create  $n - 1$  detours. For instance, in [Figure 145](#), the detour tries to merge back into the LSP at Router D instead of at Router E or Router F. Merging back into the LSP makes the detour scalability problem more manageable. If topology limitations prevent the detour from quickly merging back into the LSP, detours merge with other detours automatically.

Figure 145: Detours Merging into Other Detours



#### Related Documentation

- [fast-reroute on page 5042](#)
- [Configuring Fast Reroute](#)
- [MPLS Feature Support on QFX Series and EX4600 Switches on page 4934](#)
- [Interprovider and Carrier-of-Carriers VPNs on page 4966](#)

## Graceful Restart and MPLS-Related Protocols

This section contains the following topics:

- [LDP on page 4970](#)
- [RSVP on page 4970](#)
- [CCC and TCC on page 4970](#)

## LDP

LDP graceful restart enables a router whose LDP control plane is undergoing a restart to continue to forward traffic while recovering its state from neighboring routers. It also enables a router on which helper mode is enabled to assist a neighboring router that is attempting to restart LDP.

During session initialization, a router advertises its ability to perform LDP graceful restart or to take advantage of a neighbor performing LDP graceful restart by sending the graceful restart TLV. This TLV contains two fields relevant to LDP graceful restart: the reconnect time and the recovery time. The values of the reconnect and recovery times indicate the graceful restart capabilities supported by the router.

The reconnect time is configured in Junos OS as 60 seconds and is not user-configurable. The reconnect time is how long the helper router waits for the restarting router to establish a connection. If the connection is not established within the reconnect interval, graceful restart for the LDP session is terminated. The maximum reconnect time is 120 seconds and is not user-configurable. The maximum reconnect time is the maximum value that a helper router accepts from its restarting neighbor.

When a router discovers that a neighboring router is restarting, it waits until the end of the recovery time before attempting to reconnect. The recovery time is the length of time a router waits for LDP to restart gracefully. The recovery time period begins when an initialization message is sent or received. This time period is also typically the length of time that a neighboring router maintains its information about the restarting router, so it can continue to forward traffic.

You can configure LDP graceful restart both in the master instance for the LDP protocol and for a specific routing instance. You can disable graceful restart at the global level for all protocols, at the protocol level for LDP only, and for a specific routing instance only.

## RSVP

RSVP graceful restart enables a router undergoing a restart to inform its adjacent neighbors of its condition. The restarting router requests a grace period from the neighbor or peer, which can then cooperate with the restarting router. The restarting router can still forward MPLS traffic during the restart period; convergence in the network is not disrupted. The restart is not visible to the rest of the network, and the restarting router is not removed from the network topology. RSVP graceful restart can be enabled on both transit routers and ingress routers. It is available for both point-to-point LSPs and point-to-multipoint LSPs.

## CCC and TCC

CCC and TCC graceful restart enables Layer 2 connections between customer edge (CE) routers to restart gracefully. These Layer 2 connections are configured with the **remote-interface-switch** or **lsp-switch** statements. Because these CCC and TCC connections have an implicit dependency on RSVP LSPs, graceful restart for CCC and TCC uses the RSVP graceful restart capabilities.

RSVP graceful restart must be enabled on the provider edge (PE) routers and provider (P) routers to enable graceful restart for CCC and TCC. Also, because RSVP is used as the signaling protocol for signaling label information, the neighboring router must use helper mode to assist with the RSVP restart procedures.

**Related  
Documentation**

- [Graceful Restart Concepts on page 2579](#)
- *Graceful Restart System Requirements*
- *Configuring Graceful Restart for MPLS-Related Protocols*
- *Configuring Graceful Restart*

---

## Types of LSPs

There are three types of LSPs:

- **Static LSPs**—For static paths, you must manually assign labels on all routers involved (ingress, transit, and egress). No signaling protocol is needed. This procedure is similar to configuring static routes on individual routers. Like static routes, there is no error reporting, liveliness detection, or statistics reporting.
- **LDP-signaled LSPs**—See [“LDP Introduction” on page 4804](#).
- **RSVP-signaled LSPs**—For signaled paths, RSVP is used to set up the path and dynamically assign labels. (RSVP signaling messages are used to set up signaled paths.) You configure only the ingress router. The transit and egress routers accept signaling information from the ingress router, and they set up and maintain the LSP cooperatively. Any errors encountered while establishing an LSP are reported to the ingress router for diagnostics. For signaled LSPs to work, a version of RSVP that supports tunnel extensions must be enabled on all routers.

There are two types of RSVP-signaled LSPs:

- **Explicit-path LSPs**—All intermediate hops of the LSP are manually configured. The intermediate hops can be strict, loose, or any combination of the two. Explicit path LSPs provide you with complete control over how the path is set up. They are similar to static LSPs but require much less configuration.
- **Constrained-path LSPs**—The intermediate hops of the LSP are automatically computed by the software. The computation takes into account information provided by the topology information from the IS-IS or OSPF link-state routing protocol, the current network resource utilization determined by RSVP, and the resource requirements and constraints of the LSP. For signaled constrained-path LSPs to work, either the IS-IS or OSPF protocol and the IS-IS or OSPF traffic engineering extensions must be enabled on all routers.

## Configuring CoS Bits for an MPLS Network

---

When traffic enters a labeled-switch path (LSP) tunnel, the CoS bits in the MPLS header are set in one of two ways:

- The number of the output queue into which the packet was buffered and the packet loss priority (PLP) bit are written into the MPLS header and are used as the packet's CoS value. This behavior is the default, and no configuration is required. The *Class of Service Feature Guide for Routing Devices* explains the IP CoS values, and summarizes how the CoS bits are treated.
- You set a fixed CoS value on all packets entering the LSP tunnel. A fixed CoS value means that all packets entering the LSP receive the same class of service.

To set a fixed CoS value on all packets entering the LSP:

1. Specify a class of service value for the LSP:



**NOTE:** The CoS value set using the `class-of-service` statement at the `[edit protocols mpls]` hierarchy level supersedes the CoS value set at the `[edit class-of-service]` hierarchy level for an interface. Effectively, the CoS value configured for an LSP overrides the CoS value set for an interface.

```
[edit protocols mpls]
user@switch# set class-of-service cos-value
```

### Related Documentation

- [Understanding CoS Classifiers](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4959](#)
- [Configuring a Global MPLS EXP Classifier on page 4977](#)
- [Configuring Rewrite Rules for MPLS EXP Classifiers on page 4992](#)
- [Defining CoS Rewrite Rules on page 6704](#)

## Configuring Ethernet over MPLS (L2 Circuit)

---

To implement Ethernet over MPLS, you must configure a Layer 2 circuit on the provider edge (PE) switches. No special configuration is required on the customer edge (CE) switches. The provider switches require MPLS and LDP to be configured on the interfaces that will be receiving and transmitting MPLS packets.




**NOTE:** A Layer 2 circuit is similar to a circuit cross-connect (CCC), except that multiple Layer 2 circuits can be transported over a single label-switched path (LSP) tunnel between two PE switches. In contrast, each CCC requires a dedicated LSP.

This topic describes how to configure the PE switches to support Ethernet over MPLS. You must configure interfaces and protocols on both the local PE (PE1) and the remote PE (PE2) switches. The interface configuration varies depending upon whether the Layer 2 circuit is port-based or VLAN-based.

Figure 146 shows an example of a Layer 2 circuit configuration.

Figure 146: Ethernet over MPLS Layer 2 Circuit



**NOTE:** This topic refers to the local PE switch as PE1 and the remote PE switch as PE2. It also uses interface names rather than variables to help clarify the connections between the switches. The loopback addresses of the switches are configured as follows:

- PE1: 1.1.1.1
- PE2: 4.4.4.4

- [Configuring the Local PE Switch for Port-Based Layer 2 Circuit \(Pseudo-wire\) on page 4973](#)
- [Configuring the Remote PE Switch for Port-Based Layer 2 Circuit \(Pseudo-wire\) on page 4974](#)
- [Configuring the Local PE Switch for VLAN-Based Layer 2 Circuit on page 4975](#)
- [Configuring the Remote PE Switch for VLAN-Based Layer 2 Circuit on page 4975](#)

### Configuring the Local PE Switch for Port-Based Layer 2 Circuit (Pseudo-wire)



**CAUTION:** Configure MPLS networks with an MTU (maximum transmission unit) that is at least 12 bytes larger than the largest frame size that will be transported by the LSPs. If the size of an encapsulated packet on the ingress LSR exceeds the LSP MTU, that packet is dropped. If an egress LSR receives a packet on a VC LSP with a length (after the label stack and sequencing control word have been popped) that exceeds the MTU of the destination layer 2 interface, that packet is also dropped.

To configure the local PE switch (PE1) for a port-based layer 2 circuit (pseudo-wire):

1. Configure an access CE-facing interface for Ethernet encapsulation:

```
[edit interfaces]
user@switch# set xe-0/0/6 encapsulation ethernet-ccc
```



**NOTE:** On QFX Series switches, the L2 circuit CE facing interface does not support Aggregated Ethernet (AE) interfaces.

2. Configure the Layer 2 circuit from PE1 to PE2:

```
[edit protocols]
user@switch# set l2circuit neighbor 4.4.4.4 interface xe-0/0/6 virtual-circuit-id 1
```

3. Configure the label switched path from PE1 to PE2:

```
[edit protocols]
user@switch# set mpls label-switched-path PE1-to-PE2 to 4.4.4.4
```

4. Configure the protocols on the core and loopback interfaces:

```
[edit protocols]
user@switch# set mpls interface xe-0/0/7
user@switch# set rsdp interface xe-0/0/7
user@switch# set ldp interface lo0.0
```

## Configuring the Remote PE Switch for Port-Based Layer 2 Circuit (Pseudo-wire)

To configure the remote PE switch (PE2) for a port-based layer 2 circuit:

1. Configure an access CE-facing interface for Ethernet encapsulation:

```
[edit interfaces]
user@switch# set xe-0/0/6 encapsulation ethernet-ccc
```



**NOTE:** On QFX Series switches, the L2 circuit CE facing interface does not support AE interfaces.

2. Configure the Layer 2 circuit from PE2 to PE1:

```
[edit protocols]
user@switch# set l2circuit neighbor 1.1.1.1 interface xe-0/0/6 virtual-circuit-id 1
```

3. Configure the label switched path from PE2 to PE1:

```
[edit protocols]
user@switch# set mpls label-switched-path PE2-to-PE1 to 1.1.1.1
```

4. Configure the protocols on the core and loopback interfaces:

```
[edit protocols]
user@switch# set mpls interface xe-0/0/7
user@switch# set rsdp interface xe-0/0/7
user@switch# set ldp interface lo0.0
```

## Configuring the Local PE Switch for VLAN-Based Layer 2 Circuit

To configure the local PE switch (PE1) for a VLAN-based layer 2 circuit:

1. Configure an access CE-facing interface for VLAN encapsulation:

```
[edit interfaces]
user@switch# set xe-0/0/6 encapsulation vlan-ccc
```



**NOTE:** On QFX Series switches, the L2 circuit CE facing interface does not support AE interfaces.

2. Configure the logical unit of the CE-facing interface for VLAN encapsulation:

```
[edit interfaces]
user@switch# set xe-0/0/6 unit 0 encapsulation vlan-ccc
```

3. Configure the logical unit of the CE-facing interface to belong to family ccc:

```
[edit interfaces]
user@switch# set xe-0/0/6 unit 0 family ccc
```

4. Configure the same interface for VLAN tagging:

```
[edit interfaces]
user@switch# set xe-0/0/6 vlan-tagging
```

5. Configure the VLAN ID of the interface:

```
[edit interfaces]
user@switch# set xe-0/0/6 unit 0 vlan-id 600
```

6. Configure the Layer 2 circuit from PE1 to PE2:

```
[edit protocols]
user@switch# set l2circuit neighbor 4.4.4.4 interface xe-0/0/6 virtual-circuit-id 1
```

7. Configure the label switched path from PE1 to PE2:

```
[edit protocols]
user@switch# set mpls label-switched-path PE1-to-PE2 to 4.4.4.4
```

8. Configure the protocols on the core and loopback interfaces:

```
[edit protocols]
user@switch# set mpls interface xe-0/0/7
user@switch# set rsvp interface xe-0/0/7
user@switch# set ldp interface lo0.0
```

## Configuring the Remote PE Switch for VLAN-Based Layer 2 Circuit

To configure the remote PE switch (PE2) for a VLAN-based layer 2 circuit:

1. Configure an access CE-facing interface for VLAN encapsulation:

```
[edit interfaces]
user@switch# set xe-0/0/6 encapsulation vlan-ccc
```



**NOTE:** On QFX Series switches, the L2 circuit CE facing interface does not support AE interfaces.

2. Configure the logical unit of the CE-facing interface for VLAN encapsulation:

```
[edit interfaces]
```

```
user@switch# set xe-0/0/6 unit 0 encapsulation vlan-ccc
```

3. Configure the logical unit of the CE-facing interface to belong to family ccc:

```
[edit interfaces]
```

```
user@switch# set xe-0/0/6 unit 0 family ccc
```

4. Configure the same interface for VLAN tagging:

```
[edit interfaces]
```

```
user@switch# set xe-0/0/6 vlan-tagging
```

5. Configure the VLAN ID of the interface:

```
[edit interfaces]
```

```
user@switch# set xe-0/0/6 unit 0 vlan-id 600
```

6. Configure the Layer 2 circuit from PE2 to PE1:

```
[edit protocols]
```

```
user@switch# set l2circuit neighbor 1.1.1.1 interface xe-0/0/6 virtual-circuit-id 1
```

7. Configure the label switched path from PE2 to PE1:

```
[edit protocols]
```

```
user@switch# set mpls label-switched-path PE2-to-PE1 to 1.1.1.1
```

8. Configure the protocols on the core and loopback interfaces:

```
[edit protocols]
```

```
user@switch# set mpls interface xe-0/0/7
```

```
user@switch# set rsdp interface xe-0/0/7
```

```
user@switch# set ldp interface lo0.0
```

#### Related Documentation

- [Understanding Ethernet-over-MPLS \(L2 Circuit\) on page 4962](#)



## Configuring a Global MPLS EXP Classifier

EXP packet classification associates incoming packets with a particular MPLS CoS servicing level. EXP behavior aggregate (BA) classifiers examine the MPLS EXP value in the packet header to determine the CoS settings applied to the packet. EXP BA classifiers allow you to set the forwarding class and loss priority of an MPLS packet based on the incoming CoS value.

You can configure up to 64 EXP classifiers, however, the switch uses only one MPLS EXP classifier as a global classifier, which is applied only on interfaces configured as **family mpls**. All **family mpls** switch interfaces use the global EXP classifier to classify MPLS traffic.

There is no default EXP classifier. If you want to classify incoming MPLS packets using the EXP bits, you must configure a global EXP classifier. The global classifier applies to all MPLS traffic on all **family mpls** interfaces.

If a global EXP classifier is configured, MPLS traffic on **family mpls** interfaces uses the EXP classifier. If a global EXP classifier is not configured, then if a fixed classifier is applied to the interface, the MPLS traffic uses the fixed classifier. If no EXP classifier and no fixed classifier is applied to the interface, MPLS traffic is treated as best-effort traffic. DSCP classifiers are not applied to MPLS traffic.

To configure an MPLS EXP classifier using the CLI:

1. Create an EXP classifier and associate it with a forwarding class, a loss priority, and a code point:

```
[edit class-of-service classifiers]
user@switch# set (dscp | ieee-802.1 | exp) classifier-name forwarding-class
forwarding-class-name loss-priority level code-points [aliases] [bit-patterns]
```

2. Apply the EXP classifier to the switch interfaces:

```
[edit class-of-service]
user@switch# set system-defaults classifiers exp classifier-name
```

### Related Documentation

- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4959](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 6672](#)
- [Defining CoS Unicast BA Classifiers \(DSCP, DSCP IPv6, IEEE 802.1p\)](#)
- [Defining CoS BA Classifiers \(DSCP, DSCP IPv6, IEEE 802.1p\) on page 6656](#)
- [Configuring Rewrite Rules for MPLS EXP Classifiers on page 4992](#)

## Configuring MPLS Firewall Filters and Policers

You can configure firewall filters to filter MPLS traffic. To use an MPLS firewall filter, you must first configure the filter and then apply it to an interface you have configured for forwarding MPLS traffic. You can also configure a policer for the MPLS filter to police (that is, rate-limit) the traffic on the interface to which the filter is attached.



**NOTE:** You can configure ingress MPLS firewall filters only. Egress MPLS firewall filters are not supported. You cannot apply MPLS firewall filters to loopback interfaces.

When you configure an MPLS firewall filter, you define filtering criteria (terms, with match conditions) for the packets and an action (action, or action modifier) for the switch to take if the packets match the filtering criteria.

- [Table 378](#) describes the match conditions you can configure for MPLS firewall filters at the `[edit firewall family mpls filter filter-name term term-name from]` hierarchy level.



**NOTE:** If a packet has multiple MPLS labels, the filter applies the match conditions to only the bottom label in the label stack.

**Table 378: Supported Match Conditions for MPLS Firewall Filters**

| Match Condition     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>exp number</i>   | <p>Experimental (EXP) bit number or range of bit numbers in the MPLS header of a packet.</p> <p>For <i>number</i>, you can specify one or more values from 0 through 7 in binary, decimal or hexadecimal format, as given below:</p> <ul style="list-style-type: none"> <li>• A single EXP bit—for example, <b>exp 3</b></li> <li>• Several EXP bits—for example, <b>exp 0,4</b></li> <li>• A range of EXP bits—for example, <b>exp [0-5]</b></li> </ul> |
| <i>label number</i> | <p>MPLS label value or range of label values in the MPLS header of a packet.</p> <p>For <i>number</i>, you can specify one or more values from 0 through 1048575 in decimal or hexadecimal format, as given below:</p> <ul style="list-style-type: none"> <li>• A single label—for example, <b>label 3</b></li> <li>• Several labels—for example, <b>label 0,4</b></li> <li>• A range of labels—for example, <b>label [0-5]</b></li> </ul>               |

- [Table 379](#) describes the actions you can configure for MPLS firewall filters at the `[edit firewall family mpls filter filter-name term term-name then]` hierarchy level.

**Table 379: Supported Actions for MPLS Firewall Filters**

| Action                    | Description                                                                                                                                                                                                                                                                |
|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>accept</i>             | Accept a packet                                                                                                                                                                                                                                                            |
| <i>count counter-name</i> | <p>Count the number of packets that pass this filter or term.</p> <p><b>NOTE:</b> We recommend that you configure a counter for each term in a firewall filter, so that you can monitor the number of packets that match the conditions specified in each filter term.</p> |

Table 379: Supported Actions for MPLS Firewall Filters (*continued*)

| Action                     | Description                                                                                                  |
|----------------------------|--------------------------------------------------------------------------------------------------------------|
| <b>discard</b>             | Discard a packet silently without sending an Internet Control Message Protocol (ICMP) message                |
| <b>policer</b>             | Starting with Junos OS 13.2X51-D15, you can send traffic matched by an MPLS filter to a two-color policer.   |
| <b>three-color-policer</b> | Starting with Junos OS 13.2X51-D15, you can send traffic matched by an MPLS filter to a three-color policer. |

- [Configuring an MPLS Firewall Filter on page 4979](#)
- [Applying an MPLS Firewall Filter to an MPLS Interface on page 4979](#)
- [Configuring Policers for LSPs on page 4980](#)

## Configuring an MPLS Firewall Filter

To configure an MPLS firewall filter:

1. Configure the filter name, term name, and at least one match condition—for example, match on MPLS packets with EXP bits set to either 0 or 4:

```
[edit firewall family mpls]
user@switch# set filter ingress-exp-filter term term-one from exp 0,4
```

2. In each firewall filter term, specify the actions to take if the packet matches all the conditions in that term—for example, count MPLS packets with EXP bits set to either 0 or 4:

```
[edit firewall family mpls filter ingress-exp-filter term term-one then]
user@switch# set count counter0
user@switch# set accept
```

## Applying an MPLS Firewall Filter to an MPLS Interface

To apply the MPLS firewall filter to an interface you have configured for forwarding MPLS traffic (using the **family mpls** statement at the **[edit interfaces *interface-name* unit *unit-number*]** hierarchy level):



**NOTE:** You can apply firewall filters only to filter MPLS packets that enter an interface.

1. Apply the firewall filter to an MPLS interface—for example, apply the firewall filter to interface xe-0/0/5:

```
[edit interfaces]
user@switch# set xe-0/0/5 unit 0 family mpls filter input ingress-exp-filter
```

2. Review your configuration and issue the **commit** command:

```
[edit interfaces]
```

```
user@switch# commit
commit complete
```

## Configuring Policers for LSPs

Starting with Junos OS 13.2X51-D15, you can send traffic matched by an MPLS filter to a two-color policer or three-color policer. MPLS LSP policing allows you to control the amount of traffic forwarded through a particular LSP. Policing helps to ensure that the amount of traffic forwarded through an LSP never exceeds the requested bandwidth allocation. LSP policing is supported on regular LSPs, LSPs configured with DiffServ-aware traffic engineering, and multiclass LSPs. You can configure multiple policers for each multiclass LSP. For regular LSPs, each LSP policer is applied to all of the traffic traversing the LSP. The policer's bandwidth limitations become effective as soon as the total sum of traffic traversing the LSP exceeds the configured limit.

You configure the multiclass LSP and DiffServ-aware traffic engineering LSP policers in a filter. The filter can be configured to distinguish between the different class types and apply the relevant policer to each class type. The policers distinguish between class types based on the EXP bits.

You configure LSP policers under the **family any** filter. The **family any** filter is used because the policer is applied to traffic entering the LSP. This traffic might be from different families: IPv6, MPLS, and so on. You do not need to know what sort of traffic is entering the LSP, as long as the match conditions apply to all types of traffic.

When configuring MPLS LSP policers, be aware of the following limitations:

- LSP policers are supported for packet LSPs only.
- LSP policers are supported for unicast next hops only. Multicast next hops are not supported.
- The LSP policer runs before any output filters.
- Traffic sourced from the Routing Engine (for example, ping traffic) does not take the same forwarding path as transit traffic. This type of traffic cannot be policed.

### Related Documentation

- [MPLS Feature Support on QFX Series and EX4600 Switches on page 4934](#)
- [Supported MPLS Scaling Values on page 5013](#)
- [Overview of Policers on page 5999](#)

---

## Configuring MPLS to Gather Statistics

You can configure MPLS so that it periodically gathers traffic statistics about all MPLS sessions, including transit sessions, by configuring the **statistics** statement. You must configure the **statistics** statement if you want to collect MPLS traffic statistics using SNMP polling of MPLS Management Information Bases (MIBs).

To enable or disable MPLS statistics collection, include the **statistics** statement:

```
statistics {
```

```

auto-bandwidth;
file filename <files number> <size size> <world-readable | no-world-readable>;
interval seconds;
no-transit-statistics;
transit-statistics-polling;
}

```

You can configure these statements at the following hierarchy levels:

- **[edit protocols mpls]**
- **[edit logical-systems *logical-system-name* protocols mpls]**

The default interval is 300 seconds.

If you configure the **file** option, the statistics are placed in a file, with one entry per LSP. During the specified interval, the following information is recorded in this file:

- The number of packets, number of bytes, packets per second, and bytes per second transmitted by each LSP. Feature parity for the display of packet and byte statistics for sub-LSPs of a point-to-multipoint LSP on the Junos Trio chipset is supported in Junos OS Releases 11.1R2, 11.2R2, and 11.4.
- The percent of bandwidth transmitted over a given LSP in relation to the bandwidth percentage configured for that LSP. If no bandwidth is configured for an LSP, 0 percent is recorded in the percentage column.

At the end of each periodic report, a summary shows the current time, total number of sessions, number of sessions read, number of sessions ignored, and read errors, if any. Ignored sessions are typically those not in the up state or those with a reserved (0 through 15) incoming label (typically the egress point of an LSP). The reason for a read error appears on the same line as the entry for the LSP on which the error occurred. Gathering statistics is an unreliable process; occasional read errors might affect their accuracy. Sample output follows:

|                                                               |           |              |          |           |     |
|---------------------------------------------------------------|-----------|--------------|----------|-----------|-----|
| lsp6                                                          | 0 pkt     | 0 Byte       | 0 pps    | 0 Bps     | 0   |
| lsp5                                                          | 0 pkt     | 0 Byte       | 0 pps    | 0 Bps     | 0   |
| lsp6.1                                                        | 34845 pkt | 2926980 Byte | 1049 pps | 88179 Bps | 132 |
| lsp5.1                                                        | 0 pkt     | 0 Byte       | 0 pps    | 0 Bps     | 0   |
| lsp4                                                          | 0 pkt     | 0 Byte       | 0 pps    | 0 Bps     | 0   |
| Dec 7 17:28:38 Total 6 sessions: 5 success, 0 fail, 1 ignored |           |              |          |           |     |

#### Related Documentation

- *Configuring Automatic Bandwidth Allocation for LSPs*

## Configuring MPLS on Provider Edge Switches

To implement MPLS, you must configure two provider edge (PE) switches—an ingress PE switch and an egress PE switch—and at least one provider switch. You can configure

the customer edge (CE) interfaces on the PE switches of the MPLS network using IP over MPLS.

This topic describes how to configure an ingress PE switch and an egress PE switch using IP over MPLS:

1. [Configuring the Ingress PE Switch on page 4982](#)
2. [Configuring the Egress PE Switch on page 4983](#)

## Configuring the Ingress PE Switch

To configure the ingress PE switch:

1. Configure an IP address for the loopback interface and the core interfaces:

```
[edit interfaces]
user@switch# set lo0 unit 0 family inet address 192.168.10.1/32
user@switch# set xe-0/0/5 unit 0 family inet address 10.1.5.1/24
user@switch# set xe-0/0/6 unit 0 family inet address 10.1.6.1/24
```



**NOTE:** You cannot use routed VLAN interfaces (RVIs) or Layer 3 subinterfaces as core interfaces.

2. Configure OSPF on the loopback interface and the core interfaces:



**NOTE:** You can use the switch address as an alternative to the loopback interface.

```
[edit protocols ospf]
user@switch# set area 0.0.0.0 interface lo0.0
user@switch# set area 0.0.0.0 interface xe-0/0/5.0
user@switch# set area 0.0.0.0 interface xe-0/0/6.0
```

3. Configure OSPF traffic engineering:

```
[edit protocols ospf]
user@switch# set traffic-engineering
```

4. Configure RSVP on the loopback interface and the core interfaces:

```
[edit protocols rsvp]
user@switch# set interface lo0.0
user@switch# set interface xe-0/0/5.0
user@switch# set interface xe-0/0/6.0
```

5. Configure MPLS traffic engineering.

```
[edit protocols mpls]
user@switch# set traffic-engineering
```

6. Configure MPLS on the core interfaces:

```
[edit protocols mpls]
user@switch# set interface xe-0/0/5.0
user@switch# set interface xe-0/0/6.0
```

7. Configure **family mpls** on the logical units of the core interfaces, thereby identifying the interfaces that will be used for forwarding MPLS packets:

```
[edit interfaces]
```

- ```

user@switch# set xe-0/0/5 unit 0 family mpls
user@switch# set xe-0/0/6 unit 0 family mpls

```
8. Configure a customer edge interface as a Layer 3 routed interface, specifying an IP address:


```

[edit interfaces]
user@switch# set xe-0/0/3 unit 0 family inet address 121.100.10.1/16

```
 9. Configure this Layer 3 customer edge interface for the routing protocol:


```

[edit]
user@switch# set protocols ospf area 0.0.0 interface xe-0/0/3.0

```
 10. Configure an LSP on the ingress PE switch (192.168.10.1) to send IP packets over MPLS to the egress PE switch (192.168.12.1):


```

[edit protocols mpls]
user@switch# set label-switched-path lsp_1 to 192.168.12.1

```
 11. Disable constrained-path LSP computation for this LSP:


```

[edit protocols mpls]
user@switch# set label-switched-path lsp_1 no-cspf

```
 12. Configure a static route from the ingress PE switch to the egress PE switch, thereby indicating to the routing protocol that the packets will be forwarded over the MPLS LSP that has been set up to that destination:


```

[edit routing-options]
user@switch# set static route 2.2.2.0/24 next-hop 192.168.10.1
user@switch# set static route 2.2.2.0/24 resolve

```

Configuring the Egress PE Switch

To configure the egress PE switch:

1. Configure an IP address for the loopback interface and the core interfaces:

```

[edit interfaces]
user@switch# set lo0 unit 0 family inet address 192.168.12.1/32
user@switch# set xe-0/0/5 unit 0 family inet address 10.1.20.1/24
user@switch# set xe-0/0/6 unit 0 family inet address 10.1.21.1/24

```



NOTE: You cannot use routed VLAN interfaces (RVIs) or Layer 3 subinterfaces as core interfaces.

2. Configure OSPF on the loopback interface and the core interfaces:



NOTE: You can use the switch address as an alternative to the loopback interface.

- ```

[edit protocols ospf]
user@switch# set area 0.0.0.0 interface lo0.0
user@switch# set area 0.0.0.0 interface xe-0/0/5.0
user@switch# set area 0.0.0.0 interface xe-0/0/6.0

```
3. Configure RSVP on the loopback interface and the core interfaces:
 

```

[edit protocols rsvp]
user@switch# set rsvp interface lo0.0
user@switch# set rsvp interface xe-0/0/5.0

```

```
user@switch# set rsvp interface xe-0/0/6.0
```

4. Configure MPLS on the core interfaces:

```
[edit protocols mpls]
user@switch# set interface xe-0/0/5.0
user@switch# set interface xe-0/0/6.0
```

5. Configure **family mpls** on the logical units of the core interfaces, thereby identifying the interfaces that will be used for forwarding MPLS packets:

```
[edit interfaces]
user@switch# set xe-0/0/5 unit 0 family mpls
user@switch# set xe-0/0/6 unit 0 family mpls
```

6. Configure a customer edge interface as a Layer 3 routed interface, specifying an IP address:

```
[edit interfaces]
user@switch# set xe-0/0/3 unit 0 family inet address 2.2.2.1/16
```

7. Configure this Layer 3 customer edge interface for the routing protocol:

```
[edit]
user@switch# set protocols ospf area 0.0.0 interface xe-0/0/3
```

8. Configure an LSP on the egress PE switch (192.168.12.1) to send IP packets over MPLS to the ingress PE switch (192.168.10.1):

```
[edit protocols mpls]
user@switch# set label-switched-path lsp_2 to 192.168.10.1
```

9. Disable constrained-path LSP computation for this LSP:

```
[edit protocols mpls]
user@switch# set label-switched-path lsp_2 no-cspf
```

10. Configure a static route from the ingress PE switch to the egress PE switch, thereby indicating to the routing protocol that the packets will be forwarded over the MPLS LSP that has been set up to that destination:

```
[edit routing-options]
user@switch# set static route 121.121.121.0/24 next-hop 192.168.12.1
user@switch# set static route 121.121.121.0/24 resolve
```

#### Related Documentation

- [MPLS Configuration Guidelines on page 5012](#)
- [Configuring MPLS on Provider Switches on page 4985](#)
- [MPLS Feature Support on QFX Series and EX4600 Switches on page 4934](#)
- [Understanding MPLS Components for QFX Series and EX4600 Switches on page 4946](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4959](#)



## Configuring MPLS on Provider Switches

To implement MPLS, you must configure at least one provider switch as a transit switch for the MPLS packets.

MPLS requires the configuration of an interior gateway protocol (OSPF) and a signaling protocol (RSVP) on the core interfaces and the loopback interface of all the switches. This procedure includes the configuration of OSPF on the provider switch.

To configure the provider switch, complete the following tasks:

1. Configure OSPF on the loopback and core interfaces:



**NOTE:** You can use the switch address as an alternative to the loopback interface.

```
[edit protocols ospf]
user@switch# set area 0.0.0.0 interface lo0.0
user@switch# set area 0.0.0.0 interface xe-0/0/5.0
user@switch# set area 0.0.0.0 interface xe-0/0/6.0
user@switch# set area 0.0.0.0 interface ae0
```



**NOTE:** You cannot use routed VLAN interfaces (RVIs) or Layer 3 subinterfaces as core interfaces.

2. Configure MPLS on the core interfaces:

```
[edit protocols mpls]
user@switch# set interface xe-0/0/5.0
user@switch# set interface xe-0/0/6.0
user@switch# set interface ae0
```

3. Configure RSVP on the loopback interface and the core interfaces:

```
[edit protocols rsvp]
user@switch# set interface lo0.0
user@switch# set interface xe-0/0/5.0
user@switch# set interface xe-0/0/6.0
user@switch# set interface ae0
```

4. Configure an IP address for the loopback interface and the core interfaces:

```
[edit interfaces]
user@switch# set lo0 unit 0 family inet address 127.1.1.1/32
user@switch# set xe-0/0/5 unit 0 family inet address 10.1.5.1/24
user@switch# set xe-0/0/6 unit 0 family inet address 10.1.6.1/24
user@switch# set ae0 unit 0 family inet address 10.1.9.2/24
```

5. Configure **family mpls** on the logical units of the core interfaces, thereby identifying the interfaces that will be used for forwarding MPLS packets:

```
[edit interfaces]
user@switch# set xe-0/0/5 unit 0 family mpls
user@switch# set xe-0/0/6 unit 0 family mpls
user@switch# set ae0 unit 0 family mpls
```



**NOTE:** You can configure **family mpls** on either individual interfaces or aggregated Ethernet interfaces. You cannot configure it on tagged VLAN interfaces.

**Related  
Documentation**

- [Configuring MPLS on Provider Edge Switches on page 4981](#)
- [MPLS Configuration Guidelines on page 5012](#)
- [MPLS Feature Support on QFX Series and EX4600 Switches on page 4934](#)
- [Understanding MPLS Components for QFX Series and EX4600 Switches on page 4946](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4959](#)

---

## Configuring Reporting of Automatic Bandwidth Allocation Statistics for LSPs

Automatic bandwidth allocation allows an MPLS tunnel to automatically adjust its bandwidth allocation based on the volume of traffic flowing through the tunnel. You can configure the device to collect statistics related to automatic bandwidth allocation by completing the following steps:

1. To collect statistics related to automatic bandwidth allocation, configure the **auto-bandwidth** option for the **statistics** statement at the **[edit protocols mpls]** hierarchy level. These settings apply to all LSPs configured on the router on which you have also configured the **auto-bandwidth** statement at the **[edit protocols mpls label-switched-path *label-switched-path-name*]** hierarchy level.

```
statistics {
 auto-bandwidth;
 file filename <files number> <size size> <world-readable | no-world-readable>;
 interval seconds;
 no-transit-statistics;
 transit-statistics-polling;
}
```
2. Specify the ***filename*** for the files used to store the MPLS trace operation output using the **file** option. All files are placed in the directory **/var/log**. We recommend that you place MPLS tracing output in the file **mpls-log**.
3. Specify the maximum number of trace files using the **files *number*** option. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file.0***, then ***trace-file.1***, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.
4. Specify the interval for calculating the average bandwidth usage by configuring a time in seconds using the **interval** option. You can also set the adjustment interval on a specific LSP by configuring the **interval** option at the **[edit protocols mpls label-switch-path *label-switched-path-name* statistics]** hierarchy level.



**NOTE:** To prevent unnecessary resignaling of LSPs, it is best to configure an LSP adjustment interval that is at least three times longer than the MPLS automatic bandwidth statistics interval. For example, if you configure a value of 30 seconds for the MPLS automatic bandwidth statistics interval (interval statement at the [edit protocols mpls statistics] hierarchy level), you should configure a value of at least 90 seconds for the LSP adjustment interval (adjust-interval statement at the [edit protocols mpls label-switched-path *label-switched-path-name* auto-bandwidth] hierarchy level).

- To trace automatic bandwidth allocation, include the **autobw-state flag** for the MPLS **traceoptions** statement at the [edit protocols mpls] hierarchy level.

The following configuration enables the MPLS traceoptions for automatic bandwidth allocation. The trace records are stored in a file called **auto-band-trace** (the filename is user configurable):

```
[edit protocols mpls]
traceoptions {
 file auto-band-trace size 10k files 10 world-readable;
 flag autobw-state;
}
```

- Using the **show log** command, you can display the automatic bandwidth allocation statistics file generated when you configure the *auto-bandwidth* statement. The following shows sample log file output taken from an MPLS statistics file named **auto-band-stats** on a router configured with an LSP named **E-D**. The log file shows that LSP **E-D** is operating over its reserved bandwidth limit initially. Before **Oct 30 17:14:57**, the router triggered an automatic bandwidth adjustment (you might see two sessions for an LSP undergoing an automatic bandwidth adjustment). By **Oct 30 17:16:57**, the LSP has been reestablished at a higher bandwidth and is now shown using less than 100 percent of its **Reserved Bw** (reserved bandwidth).

```
user@host> show log auto-band-stats
E-D (LSP ID 5, Tunnel ID 6741) 209 pkt 17094 Byte 1 pps 90 Bps Util
 240.01% Reserved Bw 37 Bps
decr nh 0x952c224, type 4, flags 0x0, n_gw 1, nhid 0 to refcount 10ct 30 17:13:57 Total 1 sessions: 1
success, 0 fail, 0 ignored
E-D (LSP ID 5, Tunnel ID 6741) 241 pkt 19737 Byte 1 pps 88 Bps Util
 234.67% Reserved Bw 37 Bps
decr nh 0x952c224, type 4, flags 0x0, n_gw 1, nhid 0 to refcount 10ct 30 17:14:27 Total 1 sessions: 1
success, 0 fail, 0 ignored
E-D (LSP ID 5, Tunnel ID 6741) 276 pkt 22607 Byte 1 pps 95 Bps Util
 253.34% Reserved Bw 37 Bps
decr nh 0x952c224, type 4, flags 0x0, n_gw 1, nhid 0 to refcount 10ct 30 17:14:57 Total 1 sessions: 1
success, 0 fail, 0 ignored
E-D (LSP ID 5, Tunnel ID 6741) 0 pkt 0 Byte 0 pps 0 Bps Util
 0.00% Reserved Bw 37 Bps
E-D (LSP ID 6, Tunnel ID 6741) 0 pkt 0 Byte 0 pps 0 Bps Util
 0.00% Reserved Bw 101 Bps
decr nh 0x952c224, type 4, flags 0x0, n_gw 1, nhid 0 to refcount 1decr nh 0x952c308, type 4, flags 0x0,
n_gw 1, nhid 0 to refcount 10ct 30 17:15:27 Total 2 sessions: 2 success, 0 fail, 0 ignored
E-D (LSP ID 5, Tunnel ID 6741) 0 pkt 0 Byte 0 pps 0 Bps Util
```

```

0.00% Reserved Bw 37 Bps
E-D (LSP ID 6, Tunnel ID 6741) 33 pkt 2695 Byte 1 pps 89 Bps Util
87.69% Reserved Bw 101 Bps
decr nh 0x952c224, type 4, flags 0x0, n_gw 1, nhid 0 to refcount 1decr nh 0x952c308, type 4, flags 0x0,
n_gw 1, nhid 0 to refcount 10ct 30 17:15:57 Total 2 sessions: 2 success, 0 fail, 0 ignored
E-D (LSP ID 5, Tunnel ID 6741) 0 pkt 0 Byte 0 pps 0 Bps Util
0.00% Reserved Bw 37 Bps
E-D (LSP ID 6, Tunnel ID 6741) 65 pkt 5338 Byte 1 pps 88 Bps Util
86.70% Reserved Bw 101 Bps
decr nh 0x952c224, type 4, flags 0x0, n_gw 1, nhid 0 to refcount 1decr nh 0x952c308, type 4, flags 0x0,
n_gw 1, nhid 0 to refcount 10ct 30 17:16:27 Total 2 sessions: 2 success, 0 fail, 0 ignored
E-D (LSP ID 6, Tunnel ID 6741) 97 pkt 7981 Byte 1 pps 88 Bps Util
86.70% Reserved Bw 101 Bps
decr nh 0x952c308, type 4, flags 0x0, n_gw 1, nhid 0 to refcount 10ct 30 17:16:57 Total 1 sessions: 1
success, 0 fail, 0 ignored

```

7. Issue the `show mpls lsp autobandwidth` command to display current information about automatic bandwidth allocation. The following shows sample output from the `show mpls lsp autobandwidth` command taken at about the same time as the log file shown previously:

```

user@host> show mpls lsp autobandwidth
Lspname Last Requested Reserved Highwater AdjustTime LastAdjust
BW BW BW mark Left (sec)
E-D 300bps 812.005bps 812bps 1.56801kbps 294 sec Wed Oct 30 17:15:26 2013

```

8. Issue the `file show` command to display the MPLS trace file. You need to specify the file location and file name (the file is located in `/var/log/`). The following shows sample trace file output is taken from an MPLS trace file named `auto-band-trace.0.gz` on a router configured with an LSP named `E-D`. The trace file shows that LSP `E-D` is operating over its reserved bandwidth limit initially. At `Oct 30 17:15:26`, the router triggers an automatic bandwidth adjustment (you might see two sessions for an LSP undergoing an automatic bandwidth adjustment). By `Oct 30 17:15:57`, the LSP has been reestablished at a higher bandwidth and is now shown using less than 100 percent of its **Reserved Bw** (reserved bandwidth).

```

user@host> file show /var/log/auto-band-trace.0.gz
Oct 30 17:13:57 trace_on: Tracing to "/var/log/E/auto-band-trace" started
Oct 30 17:13:57.466825 LSP E-D (id 5) new bytes arrived 2714 in 29
sec
Oct 30 17:14:27.466713 E-D (LSP ID 5, Tunnel ID 6741) 241
pkt 19737 Byte 1 pps 88 Bps Util 234.67% Reserved Bw
 37 Bps
Oct 30 17:14:27.466962 LSP E-D (id 5, old id 5); sampled bytes 19737 >
bytes recorded 17094
Oct 30 17:14:27.467035 LSP E-D (id 5) new bytes arrived 2643 in 29
sec
Oct 30 17:14:57.466599 E-D (LSP ID 5, Tunnel ID 6741) 276
pkt 22607 Byte 1 pps 95 Bps Util 253.34% Reserved Bw
 37 Bps
Oct 30 17:14:57.466758 LSP E-D (id 5, old id 5); sampled bytes 22607 >
bytes recorded 19737
Oct 30 17:14:57.466825 LSP E-D (id 5) new bytes arrived 2870 in 29
sec
Oct 30 17:15:26.265816 Adjust Autobw: LSP E-D (id 5) curr adj bw 300bps updated
with 812.005bps
Oct 30 17:15:26.266064 mpls LSP E-D Autobw change 512.005bps >= threshold 75bps
Oct 30 17:15:26.363372 Autobw Success: LSP E-D () (old id 5 new id 6) update
prev active bw 300 bps with 812 bps

```

```

Oct 30 17:15:26.363686 RPD_MPLS_PATH_BANDWIDTH_CHANGE: MPLS path (lsp E-D)
bandwidth changed, path bandwidth 812 bps
Oct 30 17:15:27.364751 RPD_MPLS_LSP_BANDWIDTH_CHANGE: MPLS LSP E-D bandwidth
changed, lsp bandwidth 812 bps
Oct 30 17:15:27.466849 E-D (LSP ID 5, Tunnel ID 6741) 0
pkt 0 Byte 0 pps 0 Bps Util 0.00% Reserved Bw
37 Bps
Oct 30 17:15:27.467050 E-D (LSP ID 6, Tunnel ID 6741) 0
pkt 0 Byte 0 pps 0 Bps Util 0.00% Reserved Bw
101 Bps
Oct 30 17:15:57.466858 E-D (LSP ID 5, Tunnel ID 6741) 0
pkt 0 Byte 0 pps 0 Bps Util 0.00% Reserved Bw
37 Bps
Oct 30 17:15:57.467106 E-D (LSP ID 6, Tunnel ID 6741) 33
pkt 2695 Byte 1 pps 89 Bps Util 87.69% Reserved Bw
101 Bps
Oct 30 17:15:57.467201 LSP E-D (id 6, old id 5); LSP up after autobw adjustment
and active for 30 sec
Oct 30 17:15:57.467398 LSP E-D (id 6) psb bytes 2695 < bytes recorded
22607 total bytes 2695 in 30 sec
Oct 30 17:15:57.467461 First sample of the adjust interval after automatic bw
adjustment
Oct 30 17:15:57.467594 Update curr max avg bw 0bps of LSP E-D with new bw
716.225bps
Oct 30 17:16:27.466830 E-D (LSP ID 5, Tunnel ID 6741) 0
pkt 0 Byte 0 pps 0 Bps Util 0.00% Reserved Bw
37 Bps
Oct 30 17:16:27.467079 E-D (LSP ID 6, Tunnel ID 6741) 65
pkt 5338 Byte 1 pps 88 Bps Util 86.70% Reserved Bw
101 Bps
Oct 30 17:16:27.467171 LSP E-D (id 6, old id 6); sampled bytes 5338 >
bytes recorded 2695
Oct 30 17:16:27.467237 LSP E-D (id 6) new bytes arrived 2643 in 29
sec
Oct 30 17:16:57.466712 E-D (LSP ID 6, Tunnel ID 6741) 97
pkt 7981 Byte 1 pps 88 Bps Util 86.70% Reserved Bw
101 Bps
Oct 30 17:16:57.466870 LSP E-D (id 6, old id 6); sampled bytes 7981 >
bytes recorded 5338

```

- Related Documentation**
- *Configuring Automatic Bandwidth Allocation for LSPs*
  - [show mpls lsp autobandwidth on page 5194](#)

## Configuring Static Label Switched Paths for MPLS

Configuring static label-switched paths (LSPs) for MPLS is similar to configuring static routes on individual switches. As with static routes, there is no error reporting, liveliness detection, or statistics reporting.

To configure static LSPs, configure the ingress PE switch and each provider switch along the path up to and including the egress PE switch.

For the ingress PE switch, configure which packets to tag (based on the packet's destination IP address), configure the next switch in the LSP, and the tag to apply to the packet. Manually assigned labels can have values from 0 through 1,048,575.

For the transit switches in the path, configure the next switch in the path and the tag to apply to the packet. Manually assigned labels can have values from 1,000,000 through 1,048,575.

The egress PE switch removes the label and forwards the packet to the IP destination. However, if the previous switch removed the label, the egress switch examines the packet's IP header and forwards the packet toward its IP destination.

Before you configure a static LSP, you must configure the basic components for an MPLS network:

- Configure two PE switches. See [“Configuring MPLS on Provider Edge Switches” on page 4981](#).



**NOTE:** Do not configure LSPs at the [edit protocols mpls label-switched-path] hierarchy level on the PE switches.

- Configure one or more provider switches. See [“Configuring MPLS on Provider Switches” on page 4985](#).

This topic describes how to configure an ingress PE switch, one or more provider switches, and an egress PE switch for static LSP:

1. [Configuring the Ingress PE Switch on page 4990](#)
2. [Configuring the Provider and the Egress PE Switch on page 4991](#)

## Configuring the Ingress PE Switch

To configure the ingress PE switch:

1. Configure an IP address for every core interface:

```
[edit interfaces]
user@switch# set interface-name unit logical-unit-number family inet address address
```



**NOTE:** You cannot use routed VLAN interfaces (RVIs) or Layer 3 subinterfaces as core interfaces.

2. Configure the name associated with the static LSP:

```
[edit protocols mpls]
user@switch# set static-label-switched-path lsp-name
```

3. Configure the next hop switch for the LSP:

```
[edit protocols mpls]
user@switch# set static-label-switched-path lsp-name ingress next-hop address-of-next-hop
```

4. Specify the address of the egress switch for the LSP:

```
[edit protocols mpls]
user@switch# set static-label-switched-path lsp-name ingress to address-of-egress-switch
```

5. Configure the new label that you want to add to the top of the label stack:

```
[edit protocols mpls]
```

```
user@switch# set static-label-switched-path lsp-name ingress push out-label
```

## Configuring the Provider and the Egress PE Switch

To configure a static LSP for MPLS on the provider and egress PE switch:

1. Configure a transit static LSP:

```
[edit protocols mpls]
```

```
user@switch# set static-label-switched-path lsp-name transit incoming-label
```

2. Configure the next hop switch for the LSP:

```
[edit protocols mpls]
```

```
user@switch# set static-label-switched-path lsp-name transit incoming-label next-hop
address-of-next-hop
```

3. Only for provider switches, remove the label at the top of the label stack and replace it with the specified label:

```
[edit protocols mpls]
```

```
user@switch# set static-label-switched-path lsp-name transit incoming-label swap out-label
```

4. Only for the egress PE switch, remove the label at the top of the label stack:



**NOTE:** If there is another label in the stack, that label becomes the label at the top of the label stack. Otherwise, the packet is forwarded as a native protocol packet (typically, as an IP packet).

```
[edit protocols mpls]
```

```
user@switch# set static-label-switched-path lsp-name transit incoming-label pop
```

### Related Documentation

- [Configuring MPLS on Provider Edge Switches on page 4981](#)
- [Configuring MPLS on Provider Switches on page 4985](#)
- [Understanding MPLS Label Operations on page 4951](#)

## Configuring Rewrite Rules for MPLS EXP Classifiers

You configure EXP rewrite rules to alter CoS values in outgoing MPLS packets on the outbound **family mpls** interfaces of a switch to match the policies of a targeted peer. Policy matching allows the downstream routing platform or switch in a neighboring network to classify each packet into the appropriate service group.

To configure an EXP CoS rewrite rule, create the rule by giving it a name and associating it with a forwarding class, loss priority, and code point. This creates a rewrite table. After the rewrite rule is created, enable it on a logical **family mpls** interface. EXP rewrite rules can only be enabled on logical **family mpls** interfaces, not on physical interfaces or on interfaces of other family types. You can also apply an existing EXP rewrite rule on a logical interface.



**NOTE:** There are no default rewrite rules.

You can configure up to 64 EXP rewrite rules, but you can only use 16 EXP rewrite rules at any time on the switch. On a given **family mpls** logical interface, all pushed MPLS labels have the same EXP rewrite rule applied to them. You can apply different EXP rewrite rules to different logical interfaces on the same physical interface.



**NOTE:** On each physical interface, either all forwarding classes that are being used on the interface must have rewrite rules configured, or no forwarding classes that are being used on the interface can have rewrite rules configured. On any physical port, do not mix forwarding classes with rewrite rules and forwarding classes without rewrite rules.



**NOTE:** To replace an existing rewrite rule on the interface with a new rewrite rule of the same type, first explicitly remove the existing rewrite rule and then apply the new rule.

To create an EXP rewrite rule for MPLS traffic and enable it on a logical interface:

1. Create an EXP rewrite rule:

```
user@switch# set class-of-service rewrite-rules exp rewrite-rule-name forwarding-class forwarding-class-name loss-priority level code-points [aliases] [bit-patterns]
```

For example, to configure an EXP rewrite rule named **exp-rr-1** for a forwarding class named **mpls-1** with a loss priority of **low** that rewrites the EXP code point value to **001**:

```
user@switch# set class-of-service rewrite-rules exp exp-rr-1 forwarding-class mpls-1 loss-priority low code-points 001
```

2. Apply the rewrite rule to a logical interface:

```
user@switch # set class-of-service interfaces interface-name unit logical-unit rewrite-rules exp rewrite-rule-name
```



For example, to apply a rewrite rule named **exp-rr-1** to logical interface **xe-0/0/10.0**:

```
user@switch# set class-of-service interfaces xe-0/0/10 unit 0 rewrite-rules exp exp-rr-1
```



**NOTE:** In this example, all forwarding classes assigned to port **xe-0/0/10** must have rewrite rules. Do not mix forwarding classes that have rewrite rules with forwarding classes that do not have rewrite rules on the same interface.

#### Related Documentation

- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4959](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 6672](#)
- [Monitoring CoS Rewrite Rules on page 7235](#)
- [Defining CoS Rewrite Rules on page 6704](#)

## Example: Configuring MPLS-Based Layer 3 VPNs

You can implement an MPLS-based Layer 3 virtual private network (VPN) on QFX switches to interconnect sites for customers who want the service provider to handle all the Layer 3 routing functions. To support an MPLS-based Layer 3 VPN, you need to add components of the Layer 3 VPN to the configuration of the two provider edge (PE) switches. You do not need to change the configuration of the provider switches.

This example shows how to configure an MPLS-based Layer 3 VPN spanning two corporate sites:

- [Requirements on page 4994](#)
- [Overview and Topology on page 4994](#)
- [Configuring the Local PE Switch on page 4996](#)
- [Configuring the Remote PE Switch on page 4999](#)

## Requirements

This example uses the following software and hardware components:

- Junos OS Release 12.3 or later for the QFX Series
- Three QFX switches

Before you configure the Layer 3 VPN components, you must configure the basic components for an MPLS network:

- Configure two PE switches. See [“Configuring MPLS on Provider Edge Switches” on page 4981](#).
- Configure one or more provider switches. See [“Configuring MPLS on Provider Switches” on page 4985](#).

## Overview and Topology

Layer 3 VPNs allow customers to leverage the service provider’s technical expertise to ensure efficient site-to-site routing. The customer’s customer edge (CE) switch uses a routing protocol such as BGP or OSPF to communicate with the service provider’s provider edge (PE) switch to carry IP prefixes across the network. MPLS-based Layer 3 VPNs use only IP over MPLS; other protocol packets are not supported. This example includes two PE switches, PE1 and PE2.


In the basic MPLS configuration of the PE switches using IP over MPLS, the PE switches were configured to use OSPF as the routing protocol between the MPLS switches and RSVP as the signaling protocol. Traffic engineering was enabled. A label-switched path (LSP) was configured.

The following components must be added to the PE switches for an MPLS-based Layer 3 VPN:

- BGP group with **family inet-vpn unicast**
- Routing instance with instance type **vrf**

[Figure 147](#) illustrates the topology of this MPLS-based Layer 3 VPN.

**Figure 147: MPLS-Based Layer 3 VPN**



[Table 380](#) shows the settings of the customer edge interface on the local CE switch.

Table 380: Local CE Switch in the MPLS-Based Layer 3 VPN Topology

| Property                 | Settings                                                                      | Description                         |
|--------------------------|-------------------------------------------------------------------------------|-------------------------------------|
| Local CE switch hardware | QFX switch                                                                    | CE1                                 |
| Customer edge interface  | <b>ge-0/0/14 unit 0</b><br><b>family inet</b><br><b>address 51.51.0.14/16</b> | Interface that connects CE1 to PE1. |

Table 381 shows the settings of the customer edge interface on the remote CE switch.

Table 381: Remote CE Switch in the MPLS-Based Layer 3 VPN Topology

| Property                  | Settings                                                                      | Description                         |
|---------------------------|-------------------------------------------------------------------------------|-------------------------------------|
| Remote CE switch hardware | QFX switch                                                                    | CE2                                 |
| Customer edge interface   | <b>ge-0/0/14 unit 0</b><br><b>family inet</b><br><b>address 11.22.26.1/16</b> | Interface that connects CE2 to PE2. |

Table 382 shows the Layer 3 VPN components of the local PE switch.

Table 382: Layer 3 VPN Components of the Local PE Switch

| Property                 | Settings                                                                                | Description                                                                                                                                                                                                                                                                       |
|--------------------------|-----------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Local PE switch hardware | QFX switch                                                                              | PE1                                                                                                                                                                                                                                                                               |
| Customer edge interface  | <b>ge-5/0/24 unit 0</b><br><b>family inet</b><br><b>address 51.51.0.1/16</b>            | Connects PE1 to CE1.<br><br><b>NOTE:</b> The <b>family inet</b> configuration should already have been completed as part of the basic MPLS configuration of the PE switch for IP over MPLS. It is included here to show what was specified for that portion of the configuration. |
| Core interface           | <b>xe-0/0/6 unit 0</b><br><b>family inet address 60.0.0.60/16</b><br><b>family mpls</b> | Connects PE1 to P.<br><br><b>NOTE:</b> This portion of the configuration should already have been completed as part of the basic MPLS configuration. It is included here to show what was specified for that portion of the configuration.                                        |
| Loopback interface       | <b>lo0 unit 0</b><br><b>family inet address 21.21.21.21/32</b>                          | <b>NOTE:</b> This portion of the configuration should already have been completed as part of the basic MPLS configuration. It is included here to show what was specified for that portion of the configuration.                                                                  |

Table 382: Layer 3 VPN Components of the Local PE Switch (*continued*)

| Property         | Settings       | Description                              |
|------------------|----------------|------------------------------------------|
| BGP              | <b>bgp</b>     | Added for the Layer 3 VPN configuration. |
| Routing instance | <b>L3VPN-1</b> | Added for the Layer 3 VPN configuration. |

Table 383 shows the Layer 3 VPN components of the remote PE switch.

Table 383: Layer 3 VPN Components of the Remote PE Switch

| Property                  | Settings                                                                                | Description                                                                                                                                                                                                                                                                                                                                            |
|---------------------------|-----------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Remote PE switch hardware | QFX switch                                                                              | PE2                                                                                                                                                                                                                                                                                                                                                    |
| Customer edge interface   | <b>ge-0/0/14 unit 0</b><br><b>family inet</b><br><b>address 11.22.26.14/16</b>          | Connects PE2 to CE2.<br><br>For the Layer 3 VPN configuration, added <b>family mpls</b> .<br><br><b>NOTE:</b> The <b>family inet</b> configuration should already have been completed as part of the basic MPLS configuration of the PE switch for IP over MPLS. It is included here to show what was specified for that portion of the configuration. |
| Core interface            | <b>xe-0/0/6 unit 0</b><br><b>family inet address 60.2.0.60/16</b><br><b>family mpls</b> | Connects PE1 to P.<br><br><b>NOTE:</b> This portion of the configuration should already have been completed as part of the basic MPLS configuration. It is included here to show what was specified for that portion of the configuration.                                                                                                             |
| Loopback interface        | <b>lo0 unit 0</b><br><b>family inet address 22.22.22.22/32</b>                          | <b>NOTE:</b> This portion of the configuration should already have been completed as part of the basic MPLS configuration. It is included here to show what was specified for that portion of the configuration.                                                                                                                                       |
| BGP                       | <b>bgp</b>                                                                              | Added for the Layer 3 VPN configuration.                                                                                                                                                                                                                                                                                                               |
| Routing instances         | <b>L3VPN-1</b>                                                                          | Added for the Layer 3 VPN configuration.                                                                                                                                                                                                                                                                                                               |

## Configuring the Local PE Switch

**CLI Quick Configuration** To quickly configure the Layer 3 VPN components on the local PE switch, copy the following commands and paste them into the switch terminal window of PE1:

```
[edit]
set protocols bgp local-address 21.21.21.21 family inet-vpn unicast
```

```

set protocols bgp group PE1-PE2 type internal
set protocols bgp neighbor 22.22.22.22
set routing-instances L3VPN-1 instance-type vrf
set routing-instances L3VPN-1 description "BETWEEN PE1 AND PE2"
set routing-instances L3VPN-1 interface ge-0/0/14.0
set routing-instances L3VPN-1 route-distinguisher 21:21
set routing-instances L3VPN-1 vrf-target target:21:21
set routing-instances L3VPN-1 vrf-table-label
set routing-options router-id 21.21.21.21
set routing-options autonomous-system 10

```

### Step-by-Step Procedure

To configure the Layer 3 VPN components on the local PE switch:

1. Configure BGP, specifying the loopback address as the local address and specifying family `inet-vpn unicast`:  

```
[edit protocols bgp]
user@switchPE1# set local-address 21.21.21.21 family inet-vpn unicast
```
2. Configure the BGP group, specifying the group name and type:  

```
[edit protocols bgp]
user@switchPE1# set group PE1-PE2 type internal
```
3. Configure the BGP neighbor, specifying the loopback address of the remote PE switch as the neighbor's address:  

```
[edit protocols bgp]
user@switchPE1# set neighbor 22.22.22.22
```
4. Configure the routing instance, specifying the routing-instance name and using `vrf` as the instance type:  

```
[edit routing-instances]
user@switchPE1# set L3VPN-1 instance-type vrf
```
5. Configure a description for this routing instance:  

```
[edit routing-instances]
user@switchPE1# set L3VPN-1 description "BETWEEN PE1 AND PE2"
```
6. Configure the routing instance to use a route distinguisher:  

```
[edit routing-instances]
user@switchPE1# set L3VPN-1 route-distinguisher 21:21
```



**NOTE:** Each routing instance that you configure on a PE switch must have a unique route distinguisher associated with it. VPN routing instances require a route distinguisher to allow BGP to distinguish between potentially identical network layer reachability information (NLRI) messages received from different VPNs. If you configure different VPN routing instances with the same route distinguisher, the commit fails.

7. Configure the VPN routing and forwarding (VRF) target of the routing instance:  

```
[edit routing-instances]
user@switchPE1# set L3VPN-1 vrf-target target:21:21
```



**NOTE:** You can create more complex policies by explicitly configuring VRF import and export policies using the import and export options. See the *Junos OS VPNs Library for Routing Devices*.

8. Configure this routing instance with **vrf-table-label**, which maps the inner label of a packet to a specific VPN routing and forwarding (VRF) table and allows the examination of the encapsulated IP header:

```
[edit routing-instances]
user@switchPE1# set L3VPN-1 vrf-table-label
```

9. Configure the router ID and autonomous system (AS):



**NOTE:** We recommend that you explicitly configure the router identifier under the [edit routing-options] hierarchy level to avoid unpredictable behavior if the interface address on a loopback interface changes.

```
[edit routing-options]
user@switchPE1# set router-id 21.21.21.21 autonomous-system 10
```

**Results** Display the results of the configuration:

```
user@switchPE1> show configuration
```

```
interfaces {
 ge-0/0/14 {
 unit 0 {
 family inet {
 address 51.51.0.1/16;
 }
 }
 }
 lo0 {
 unit 0 {
 family inet {
 address 21.21.21.21/32;
 }
 }
 }
 xe-0/0/6 {
 unit 0 {
 family inet {
 address 60.0.0.60/16;
 }
 family mpls;
 }
 }
}
protocols {
 mpls {
 label-switched-path 21-22 {
 from 21.21.21.21;
 to 22.22.22.22;
 }
 }
}
```

```

 no-cspf;
 }
 interface xe-0/0/6.0;
 interface lo0.0;
 bgp {
 local-address 21.21.21.21;
 family inet-vpn {
 unicast;
 }
 group PE1-PE2 {
 type internal;
 neighbor 22.22.22.22;
 }
 }
 ospf
 traffic-engineering;
 area 0.0.0.0 {
 interface lo0.0;
 interface xe-0/0/6.0;
 }
 }
}
routing-instances {
 L3VPN-1 {
 instance-type vrf;
 description "BETWEEN PE1 AND PE2";
 route-distinguisher 21:21;
 vrf-target target:21:21;
 vrf-table-label;
 }
}
routing-options {
 router-id 21.21.21.21;
 autonomous-system 10;
}

```

## Configuring the Remote PE Switch

**CLI Quick Configuration** To quickly configure the Layer 3 VPN components on the remote PE switch, copy the following commands and paste them into the switch terminal window of PE2:

```

[edit]
set protocols bgp local-address 22.22.22.22 family inet-vpn unicast
set protocols bgp group PE1-PE2 type internal
set protocols bgp neighbor 21.21.21.21
set routing-instances L3VPN-1 instance-type vrf
set routing-instances L3VPN-1 description "BETWEEN PE1 AND PE2"
set routing-instances L3VPN-1 interface ge-11/0/14.0
set routing-instances L3VPN-1 route-distinguisher 21:21
set routing-instances L3VPN-1 vrf-target target:21:21
set routing-instances L3VPN-1 vrf-table-label;
set routing-options router-id 22.22.22.22
set routing-options autonomous-system 10

```

**Step-by-Step  
Procedure**

To configure Layer 3 VPN components on the remote PE switch:

1. Configure BGP, specifying the loopback address as the local address and specifying **family inet-vpn unicast**:  

```
[edit protocols bgp]
user@switchPE2# set local-address 22.22.22.22 family inet-vpn unicast
```
2. Configure the BGP group, specifying the group name and type:  

```
[edit protocols bgp]
user@switchPE2# set group PE1-PE2 type internal
```
3. Configure the BGP neighbor, specifying the loopback address of the remote PE switch as the neighbor's address:  

```
[edit protocols bgp]
user@switchPE2# set neighbor 21.21.21.21
```
4. Configure the routing instance, specifying the routing-instance name and using **vrf** as the instance type:  

```
[edit routing-instances]
user@switchPE2# set L3VPN-1 instance-type vrf
```
5. Configure a description for this routing instance:  

```
[edit routing-instances]
user@switchPE1# set L3VPN-1 description "BETWEEN PE1 AND PE2"
```
6. Configure the routing instance to apply to the customer edge interface:  

```
[edit routing-instances]
user@switchPE2# set L3VPN-1 interface ge-0/0/14.0
```
7. Configure the routing instance to use a route distinguisher, using the format *ip-address:number*:  

```
[edit routing-instances]
user@switchPE2# set L3VPN-1 route-distinguisher 21:21
```
8. Configure the VPN routing and forwarding (VRF) target of the routing instance:  

```
[edit routing-instances]
user@switchPE2# set L3VPN-1 vrf-target target:21:21
```
9. Configure this routing instance with **vrf-table-label**, which maps the inner label of a packet to a specific VPN routing and forwarding (VRF) table and allows the examination of the encapsulated IP header.  

```
[edit routing-instances]
user@switchPE2# set L3VPN-1 vrf-table-label
```
10. Configure the router ID and autonomous system (AS):  

```
[edit routing-options]
user@switchPE2# set router-id 22.22.22.22 autonomous-system 10
```

**Results** Display the results of the configuration:

```
user@switchPE2> show configuration

interfaces {
 ge-0/0/14 {
 unit 0 {
 family inet {
 address 11.22.26.14/16;
 }
 }
 }
}
```



```

}
lo0 {
 unit 0 {
 family inet {
 address 22.22.22.22/32;
 }
 }
}
xe-0/0/6 {
 unit 0 {
 family inet {
 address 60.2.0.60/16;
 }
 family mpls;
 }
}
protocols {
 mpls {
 label-switched-path 22-21 {
 from 22.22.22.22;
 to 21.21.21.21;
 no-cspf;
 }
 interface xe-0/0/6.0;
 interface lo0.0;
 }
 bgp {
 local-address 22.22.22.22;
 family inet-vpn {
 unicast;
 }
 group PE1-PE2 {
 type internal;
 neighbor 21.21.21.21;
 }
 }
 ospf {
 traffic-engineering;
 area 0.0.0.0 {
 interface ge-0/0/14.0;
 interface lo0.0;
 interface xe-0/0/6.0;
 }
 }
}
routing-instances {
 L3VPN-1 {
 instance-type vrf;
 description "BETWEEN PE1 AND PE2";
 route-distinguisher 21:21;
 vrf-target target:21:21;
 vrf-table-label;
 }
}
routing-options {
 router-id 22.22.22.22;
 autonomous-system 10;
}

```

- Related Documentation**
- [Configuring MPLS on Provider Edge Switches on page 4981](#)
  - [Configuring MPLS on Provider Switches on page 4985](#)

---

## Example: Tunneling IPv6 Traffic over MPLS IPv4 Networks

---

This example shows how to configure Junos OS to tunnel IPv6 over an MPLS-based IPv4 network. External BGP (EBGP) is used between the customer edge (CE) and provider edge (PE) devices. The remote CE devices have different AS numbers for loop detection.

- [Requirements on page 5002](#)
- [Overview on page 5002](#)
- [Configuration on page 5005](#)
- [Verification on page 5010](#)

### Requirements

No special configuration beyond device initialization is required before you configure this example.

### Overview

Detailed information about the Juniper Networks implementation of IPv6 over MPLS is described in the following Internet drafts:


- Internet draft draft-ietf-l3vpn-bgp-ipv6-07.txt, *BGP-MPLS IP VPN extension for IPv6 VPN* (expires January 2006)
- Internet draft draft-ooms-v6ops-bgp-tunnel-06.txt, *Connecting IPv6 Islands over IPv4 MPLS using IPv6 Provider Edge Routers* (expires July 2006)

These Internet drafts are available on the IETF website at <http://www.ietf.org/>.

This example shows you how to interconnect a two IPv6 networks over an IPv4-based network core, giving you the ability to provide IPv6 service without having to upgrade the routers in your core network. Multiprotocol Border Gateway Protocol (MP-BGP) is configured to exchange routes between the IPv6 networks, and data is tunneled between these IPv6 networks by means of IPv4-based MPLS.

In [Figure 148](#), PE1 and PE2 are dual-stack BGP routers or switches, meaning they have both IPv4 and IPv6 stacks. The PE devices link the IPv6 networks through the customer edge (CE) routers or switches to the IPv4 core network. The CE devices and the PE devices connect through a link layer that can carry IPv6 traffic. The PE devices use IPv6 on the CE router-facing interfaces and use IPv4 and MPLS on the core-facing interfaces. Note that one of the connected IPv6 networks could be the global IPv6 Internet.

Figure 148: IPv6 Networks Linked by MPLS IPv4 Tunnels



The two PE devices are linked through an MP-BGP session using IPv4 addresses. They use the session to exchange IPv6 routes with an IPv6 (value 2) address family indicator (AFI) and a subsequent AFI (SAFI) (value 4). Each PE router sets the next hop for the IPv6 routes advertised on this session to its own IPv4 address. Because MP-BGP requires the BGP next hop to correspond to the same address family as the network layer reachability information (NLRI), this IPv4 address needs to be embedded within an IPv6 format.

The PE devices can learn the IPv6 routes from the CE devices connected to them using MP-BGP or through static configuration. Note that if BGP is used as the PE-router-to-CE-router protocol, the MP-BGP session between the PE device and CE device could occur over an IPv4 or IPv6 Transmission Control Protocol (TCP) session. Also, the BGP routes exchanged on that session would have SAFI unicast. You must configure an export policy to pass routes between IBGP and EBGp, and between BGP and any other protocol.

The PE routers have MPLS LSPs routed to each others' IPv4 addresses. IPv4 provides signaling for the LSPs by means of RSVP. These LSPs are used to resolve the next-hop addresses of the IPv6 routes learned from MP-BGP. The next hops use IPv4-mapped IPv6 addresses, while the LSPs use IPv4 addresses.

The PE devices always advertise IPv6 routes to each other using a label value of 2, the explicit null label for IPv6 as defined in RFC 3032, *MPLS Label Stack Encoding*. As a consequence, each of the forwarding next hops for the IPv6 routes learned from remote PE routers normally push two labels. The inner label is 2 (this label could be different if the advertising PE device is not a Juniper Networks routing or switching platform), and the outer label is the LSP label. If the LSP is a single-hop LSP, then only Label 2 is pushed.

It is also possible for the PE devices to exchange plain IPv6 routes using SAFI unicast. However, there is one major advantage in exchanging labeled IPv6 routes. The

penultimate-hop router for an MPLS LSP can pop the outer label and then send the packet with the inner label as an MPLS packet. Without the inner label, the penultimate-hop router would need to discover whether the packet is an IPv4 or IPv6 packet to set the protocol field in the Layer 2 header correctly.

When the PE1 device in [Figure 148](#) receives an IPv6 packet from the CE1 device, it performs a lookup in the IPv6 forwarding table. If the destination matches a prefix learned from the CE2 device, then no labels need to be pushed and the packet is simply sent to the CE2 device. If the destination matches a prefix that was learned from the PE2 device, then the PE1 router pushes two labels onto the packet and sends it to the Provider router. The inner label is 2 and the outer label is the LSP label for the PE2 router.

Each provider router in the service provider's network handles the packet as it would any MPLS packet, swapping labels as it passes from provider router to provider router. The penultimate-hop provider router for the LSP pops the outer label and sends the packet to the PE2 router. When the PE2 router receives the packet, it recognizes the IPv6 explicit null label on the packet (Label 2). It pops this label and treats it as an IPv6 packet, performing a lookup in the IPv6 forwarding table and forwarding the packet to the CE3 router.

This example includes the following settings:

- In addition to configuring the **family inet6** statement on all the CE router-facing interfaces, you must also configure the statement on all the core-facing interfaces running MPLS. Both configurations are necessary because the router must be able to process any IPv6 packets it receives on these interfaces. You should not see any regular IPv6 traffic arrive on these interfaces, but you will receive MPLS packets tagged with Label 2. Even though Label 2 MPLS packets are sent in IPv4, these packets are treated as native IPv6 packets.
- You enable IPv6 tunneling by including the **ipv6-tunneling** statement in the configuration for the PE routers. This statement allows IPv6 routes to be resolved over an MPLS network by converting all routes stored in the inet.3 routing table to IPv4-mapped IPv6 addresses and then copying them into the inet6.3 routing table. This routing table can be used to resolve next hops for both inet6 and inet6-vpn routes.



**NOTE:** BGP automatically runs its import policy even when copying routes from a primary routing table group to a secondary routing table group. If IPv4 labeled routes arrive from a BGP session (for example, when you have configured the **labeled-unicast** statement at the `[edit protocols bgp family inet]` hierarchy level on the PE router), the BGP neighbor's import policy also accepts IPv6 routes, since the neighbor's import policy is run while doing the copy operation to the inet6.3 routing table.

---

- When you configure MP-BGP to carry IPv6 traffic, the IPv4 MPLS label is removed at the destination PE router. The remaining IPv6 packet without a label can then be forwarded to the IPv6 network. To enable this, include the **explicit-null** statement in the BGP configuration.

## Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```

Device PE1
set interfaces xe-0/0/5 unit 2 family inet6 address ::10.1.1.2/126
set interfaces xe-0/0/5 unit 2 family mpls
set interfaces xe-0/0/6 unit 5 family inet address 10.1.1.5/30
set interfaces xe-0/0/6 unit 5 family inet6
set interfaces xe-0/0/6 unit 5 family mpls
set interfaces lo0 unit 2 family inet address 1.1.1.2/32
set protocols mpls ipv6-tunneling
set protocols mpls interface xe-0/0/5.2
set protocols mpls interface xe-0/0/6.5
set protocols bgp group toCE1 type external
set protocols bgp group toCE1 local-address ::10.1.1.2
set protocols bgp group toCE1 family inet6 unicast
set protocols bgp group toCE1 export send-bgp6
set protocols bgp group toCE1 peer-as 1
set protocols bgp group toCE1 neighbor ::10.1.1.1
set protocols bgp group toPE2 type internal
set protocols bgp group toPE2 local-address 1.1.1.2
set protocols bgp group toPE2 family inet6 labeled-unicast explicit-null
set protocols bgp group toPE2 export next-hop-self
set protocols bgp group toPE2 export send-v6
set protocols bgp group toPE2 neighbor 1.1.1.4
set protocols ospf area 0.0.0.0 interface xe-0/0/6.5
set protocols ospf area 0.0.0.0 interface lo0.2 passive
set protocols rsvp interface xe-0/0/6.5
set policy-options policy-statement next-hop-self then next-hop self
set policy-options policy-statement send-bgp6 from family inet6
set policy-options policy-statement send-bgp6 from protocol bgp
set policy-options policy-statement send-bgp6 then accept
set policy-options policy-statement send-v6 from family inet6
set policy-options policy-statement send-v6 from protocol bgp
set policy-options policy-statement send-v6 from protocol direct
set policy-options policy-statement send-v6 then accept
set routing-options router-id 1.1.1.2
set routing-options autonomous-system 2

Device PE2
set interfaces xe-0/0/5 unit 10 family inet address 10.1.1.10/30
set interfaces xe-0/0/5 unit 10 family inet6
set interfaces xe-0/0/5 unit 10 family mpls
set interfaces xe-0/0/6 unit 13 family inet6 address ::10.1.1.13/126
set interfaces xe-0/0/6 unit 13 family mpls
set interfaces lo0 unit 4 family inet address 1.1.1.4/32
set protocols mpls ipv6-tunneling
set protocols mpls interface xe-0/0/5.10
set protocols mpls interface xe-0/0/6.13
set protocols bgp group toPE1 type internal
set protocols bgp group toPE1 local-address 1.1.1.4
set protocols bgp group toPE1 family inet6 labeled-unicast explicit-null
set protocols bgp group toPE1 export next-hop-self

```

```
set protocols bgp group toPE1 export send-v6
set protocols bgp group toPE1 neighbor 1.1.1.2
set protocols bgp group toCE3 type external
set protocols bgp group toCE3 local-address ::10.1.1.13
set protocols bgp group toCE3 family inet6 unicast
set protocols bgp group toCE3 export send-bgp6
set protocols bgp group toCE3 peer-as 3
set protocols bgp group toCE3 neighbor ::10.1.1.14
set protocols ospf area 0.0.0.0 interface xe-0/0/5.10
set protocols ospf area 0.0.0.0 interface lo0.4 passive
set protocols rsvp interface xe-0/0/5.10
set policy-options policy-statement next-hop-self then next-hop self
set policy-options policy-statement send-bgp6 from family inet6
set policy-options policy-statement send-bgp6 from protocol bgp
set policy-options policy-statement send-bgp6 then accept
set policy-options policy-statement send-v6 from family inet6
set policy-options policy-statement send-v6 from protocol bgp
set policy-options policy-statement send-v6 from protocol direct
set policy-options policy-statement send-v6 then accept
set routing-options router-id 1.1.1.4
set routing-options autonomous-system 2
```

Device P

```
set interfaces xe-0/0/5 unit 6 family inet address 10.1.1.6/30
set interfaces xe-0/0/5 unit 6 family inet6
set interfaces xe-0/0/5 unit 6 family mpls
set interfaces xe-0/0/6 unit 9 family inet address 10.1.1.9/30
set interfaces xe-0/0/6 unit 9 family inet6
set interfaces xe-0/0/6 unit 9 family mpls
set interfaces lo0 unit 3 family inet address 1.1.1.3/32
set protocols mpls interface xe-0/0/5.6
set protocols mpls interface xe-0/0/6.9
set protocols ospf area 0.0.0.0 interface xe-0/0/5.6
set protocols ospf area 0.0.0.0 interface xe-0/0/6.9
set protocols ospf area 0.0.0.0 interface lo0.3 passive
set protocols rsvp interface xe-0/0/5.6
set protocols rsvp interface xe-0/0/6.9
set routing-options router-id 1.1.1.3
set routing-options autonomous-system 2
```

Device CE1

```
set interfaces xe-0/0/5 unit 1 family inet6 address ::10.1.1.1/126
set interfaces xe-0/0/5 unit 1 family mpls
set interfaces lo0 unit 1 family inet6 address ::1.1.1.1/128
set protocols bgp group toPE1 type external
set protocols bgp group toPE1 local-address ::10.1.1.1
set protocols bgp group toPE1 family inet6 unicast
set protocols bgp group toPE1 export send-v6
set protocols bgp group toPE1 peer-as 2
set protocols bgp group toPE1 neighbor ::10.1.1.2
set policy-options policy-statement send-v6 from family inet6
set policy-options policy-statement send-v6 from protocol direct
set policy-options policy-statement send-v6 then accept
set routing-options router-id 1.1.1.1
set routing-options autonomous-system 1
```

Device CE3

```
set interfaces xe-0/0/5 unit 14 family inet6 address ::10.1.1.14/126
```

```

set interfaces xe-0/0/5 unit 14 family mpls
set interfaces lo0 unit 5 family inet6 address ::1.1.1.5/128
set protocols bgp group toPE2 type external
set protocols bgp group toPE2 local-address ::10.1.1.14
set protocols bgp group toPE2 family inet6 unicast
set protocols bgp group toPE2 export send-v6
set protocols bgp group toPE2 peer-as 2
set protocols bgp group toPE2 neighbor ::10.1.1.13
set policy-options policy-statement send-v6 from family inet6
set policy-options policy-statement send-v6 from protocol direct
set policy-options policy-statement send-v6 then accept
set routing-options router-id 1.1.1.5
set routing-options autonomous-system 3

```

### Configuring Device PE1

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device PE1:

1. Configure the interfaces.

```

[edit interfaces]
user@PE1# set xe-0/0/5 unit 2 family inet6 address ::10.1.1.2/126
user@PE1# set xe-0/0/5 unit 2 family mpls

user@PE1# set xe-0/0/6 unit 5 family inet address 10.1.1.5/30
user@PE1# set xe-0/0/6 unit 5 family inet6
user@PE1# set xe-0/0/6 unit 5 family mpls

user@PE1# set lo0 unit 2 family inet address 1.1.1.2/32

```

2. Configure MPLS on the interfaces.

```

[edit protocols mpls]
user@PE1# set ipv6-tunneling
user@PE1# set interface xe-0/0/5.2
user@PE1# set interface xe-0/0/6.5

```

3. Configure BGP.

```

[edit protocols bgp]
user@PE1# set group toCE1 type external
user@PE1# set group toCE1 local-address ::10.1.1.2
user@PE1# set group toCE1 family inet6 unicast
user@PE1# set group toCE1 export send-bgp6
user@PE1# set group toCE1 peer-as 1
user@PE1# set group toCE1 neighbor ::10.1.1.1

user@PE1# set group toPE2 type internal
user@PE1# set group toPE2 local-address 1.1.1.2
user@PE1# set group toPE2 family inet6 labeled-unicast explicit-null
user@PE1# set group toPE2 export next-hop-self
user@PE1# set group toPE2 export send-v6

```

```
user@PE1# set group toPE2 neighbor 1.1.1.4
```

4. Configure OSPF

```
[edit protocols ospf area 0.0.0.0]
user@PE1# set interface xe-0/0/6.5
user@PE1# set interface lo0.2 passive
```

5. Configure a signaling protocol.

```
[edit protocols]
user@PE1# set rsvp interface xe-0/0/6.5
```

6. Configure the routing policies.

```
[edit policy-options]
user@PE1# set policy-statement next-hop-self then next-hop self
```

```
user@PE1# set policy-statement send-bgp6 from family inet6
user@PE1# set policy-statement send-bgp6 from protocol bgp
user@PE1# set policy-statement send-bgp6 then accept
```

```
user@PE1# set policy-statement send-v6 from family inet6
user@PE1# set policy-statement send-v6 from protocol bgp
user@PE1# set policy-statement send-v6 from protocol direct
user@PE1# set policy-statement send-v6 then accept
```

7. Configure the router ID and the autonomous system (AS) number.

```
[edit routing-options]
user@PE1# set router-id 1.1.1.2
user@PE1# set autonomous-system 2
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show interfaces
xe-0/0/5 {
 unit 2 {
 family inet6 {
 address ::10.1.1.2/126;
 }
 family mpls;
 }
}
xe-0/0/6 {
 unit 5 {
 family inet {
 address 10.1.1.5/30;
 }
 family inet6;
 family mpls;
 }
}
lo0 {
```



```
unit 2 {
 family inet {
 address 1.1.1.2/32;
 }
}

user@R1# show policy-options
policy-statement next-hop-self {
 then {
 next-hop self;
 }
}
policy-statement send-bgp6 {
 from {
 family inet6;
 protocol bgp;
 }
 then accept;
}
policy-statement send-v6 {
 from {
 family inet6;
 protocol [bgp direct];
 }
 then accept;
}

user@R1# show protocols
mpls {
 ipv6-tunneling;
 interface xe-0/0/5.2;
 interface xe-0/0/6.5;
}
bgp {
 group toCE1 {
 type external;
 local-address ::10.1.1.2;
 family inet6 {
 unicast;
 }
 export send-bgp6;
 peer-as 1;
 neighbor ::10.1.1.1;
 }
 group toPE2 {
 type internal;
 local-address 1.1.1.2;
 family inet6 {
 labeled-unicast {
 explicit-null;
 }
 }
 export [next-hop-self send-v6];
 neighbor 1.1.1.4;
 }
}
```

```
ospf {
 area 0.0.0.0 {
 interface xe-0/0/6.5;
 interface lo0.2 {
 passive;
 }
 }
}
rsvp {
 interface xe-0/0/6.5;
}

user@R1# show routing-options
router-id 1.1.1.2;
autonomous-system 2;
```

If you are done configuring the device, enter **commit** from configuration mode. Configure the other devices in the topology, as shown in “[CLI Quick Configuration](#)” on [page 5005](#).

## Verification

Confirm that the configuration is working properly.

### Verifying That the CE Devices Have Connectivity

**Purpose** Make sure that the tunnel is operating.

**Action** From operational mode, enter the **ping** command.

```
user@CE1> ping ::10.1.1.14
PING6(56=40+8+8 bytes) ::10.1.1.1 --> ::10.1.1.14
16 bytes from ::10.1.1.14, icmp_seq=0 hlim=61 time=10.687 ms
16 bytes from ::10.1.1.14, icmp_seq=1 hlim=61 time=9.239 ms
16 bytes from ::10.1.1.14, icmp_seq=2 hlim=61 time=1.842 ms

user@CE3> ping ::10.1.1.1
PING6(56=40+8+8 bytes) ::10.1.1.14 --> ::10.1.1.1
16 bytes from ::10.1.1.1, icmp_seq=0 hlim=61 time=1.484 ms
16 bytes from ::10.1.1.1, icmp_seq=1 hlim=61 time=1.338 ms
16 bytes from ::10.1.1.1, icmp_seq=2 hlim=61 time=1.351 ms
```

**Meaning** The IPv6 CE devices can communicate over the core IPv4 network.

**Related  
Documentation**

## Verifying That MPLS Is Working Correctly

To verify that MPLS is working correctly, perform the following tasks:

1. [Verifying the Physical Layer on the Switches on page 5011](#)
2. [Verifying the Routing Protocol on page 5011](#)

3. [Verifying the Core Interfaces Being Used for the MPLS Traffic on page 5011](#)
4. [Verifying RSVP on page 5012](#)

## Verifying the Physical Layer on the Switches

**Purpose** Verify that the interfaces are up. Perform this verification task on each of the switches.

**Action** user@switch> **show interfaces xe-\* terse**

| Interface  | Admin | Link | Proto        | Local       | Remote |
|------------|-------|------|--------------|-------------|--------|
| xe-0/0/0   | up    | up   |              |             |        |
| xe-0/0/0.0 | up    | up   |              |             |        |
| xe-0/0/1.0 | up    | up   |              |             |        |
| xe-0/0/2.0 | up    | up   |              |             |        |
| xe-0/0/3.0 | up    | up   | inet         | 2.2.2.1/16  |        |
| xe-0/0/4.0 | up    | up   |              |             |        |
| xe-0/0/5.0 | up    | up   | inet<br>mpls | 10.1.5.1/24 |        |
| xe-0/0/6.0 | up    | up   | inet<br>mpls | 10.1.6.1/24 |        |

**Meaning** The **show interfaces terse** command displays status information about the 10-Gigabit Ethernet interfaces on the switch. This output verifies that the interfaces are **up**. The output for the protocol family (Proto column) of the core interfaces (xe-0/0/5.0 and xe-0/0/6.0), shows that these interfaces are configured as both **inet** and **mpls**. The **Local** column for the core interfaces shows the IP address configured for these interfaces.

## Verifying the Routing Protocol

**Purpose** Verify the state of the configured routing protocol. You should perform this verification task on each of the switches. The state should be **Full**. If you have configured OSPF as the routing protocol, use the **show ospf neighbor** command to verify that the routing protocol is communicating with the switch neighbors.

**Action** user@switch> **show ospf neighbor**

| Address   | Interface | State | ID          | Pri | Dead |
|-----------|-----------|-------|-------------|-----|------|
| 127.1.1.1 | xe-0/0/5  | Full  | 10.10.10.10 | 128 | 39   |

**Meaning** The **show ospf neighbor** command displays the status of the routing protocol that has been configured on this switch. The output shows that the state is **Full**, meaning that the routing protocol is operating correctly—that is, hello packets are being exchanged between directly connected neighbors. For additional information on checking and monitoring routing protocols, see the [Junos OS Routing Protocols and Policies Command Reference](#).

## Verifying the Core Interfaces Being Used for the MPLS Traffic

**Purpose** Verify that the state of the MPLS interface is **Up**. You should perform this verification task on each of the switches.

**Action** user@switch> **show mpls interface**

| Interface | State | Administrative groups |
|-----------|-------|-----------------------|
| ge-0/0/5  | Up    | <none>                |
| ge-0/0/6  | Up    | <none>                |

**Meaning** The **show mpls interface** command displays the status of the core interfaces that have been configured to belong to **family mpls**. This output shows that the interface configured to belong to **family mpls** is up.

## Verifying RSVP

**Purpose** Verify the state of the RSVP session. You should perform this verification task on each of the switches.

user@switch> **show rsvp session**

```
Ingress RSVP: 1 sessions
To From State Rt Style Labelin Labelout LSPname
127.1.1.3 127.1.1.1 Up 0 1 FF - 300064 lsp_to_pe2_ge1
Total 1 displayed, Up 1, Down 0

Egress RSVP: 1 sessions
To From State Rt Style Labelin Labelout LSPname
127.1.1.1 127.1.1.3 Up 0 1 FF 299968 - lsp_to_pe1_ge1
Total 1 displayed, Up 1, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

**Meaning** This output confirms that the RSVP sessions are up.

**Related Documentation**

- [Configuring MPLS on Provider Edge Switches on page 4981](#)
- [Configuring MPLS on Provider Switches on page 4985](#)

## MPLS Configuration Guidelines

---

When configuring MPLS on QFX Series devices or on EX4600, note that the number of IP prefixes supported depends on the specific platform being used. See the scale specifications in the data sheet of your device for additional information.

- We recommend the following:
  - If your ingress provider edge (PE) switch needs to support more than 8000 external IP prefixes, use a larger capacity device as an ingress PE switch.
  - If you use a switch as a route reflector for BGP labeled routes, use it as a dedicated route reflector (that is, the switch must not participate in managing data traffic).

- If you use a switch as a PE switch or as a route reflector for BGP labeled routes, configure routing policies on the PE switch and the route reflector to filter external IP routes from the routing table.

The configuration example for a routing policy named `fib_policy` (at the `[edit policy-options` and `[edit routing-options` hierarchy levels) to filter BGP labeled routes from the `inet.0` routing table is given below:

```
user@switch# show policy-options
policy-statement fib_policy {
 from {
 protocol bgp;
 rib inet.0;
 }
 then reject;
}

user@switch# show routing-options
forwarding-table {
 export fib_policy;
}
```

- Packet fragmentation using the `allow-fragmentation` statement at the `[edit protocols mpls path-mtu]` hierarchy level is not supported on QFX Series devices or on the EX4600 switch. Therefore, you must ensure that the maximum transmission unit (MTU) values configured on every MPLS interface is sufficient to handle MPLS packets. The packets whose size exceeds the MTU value of an interface will be dropped.

**Related Documentation**

- [Configuring MPLS on Provider Edge Switches on page 4981](#)
- [Configuring MPLS on Provider Switches on page 4985](#)
- [Configuring a Global MPLS EXP Classifier on page 4977](#)
- [Configuring Rewrite Rules for MPLS EXP Classifiers on page 4992](#)
- [MPLS Feature Support on QFX Series and EX4600 Switches on page 4934](#)

## Supported MPLS Scaling Values

This topic lists the MPLS scaling values supported on QFX Series switches.

[Table 384](#) lists the MPLS scaling values supported on Juniper QFX switches and on the EX4600 switch.

**Table 384: MPLS Scaling Values**

| Feature                                                 | QFX3500 Scaling Value        | QFX5100 and EX4600 Scaling Value | QFX10002 Scaling Value       |
|---------------------------------------------------------|------------------------------|----------------------------------|------------------------------|
| Maximum number of MPLS labels in a packet's label stack | 3 labels for Push operations | 3 labels for Push operations     | 5 labels for Push operations |
|                                                         | 2 labels for Pop operations  | 2 labels for Pop operations      | 8 labels for Pop operations  |
|                                                         | 1 label for Swap operations  | 1 label for Swap operations      | 1 label for Swap operations  |

Table 384: MPLS Scaling Values (*continued*)

| Feature                                                                  | QFX3500 Scaling Value                    | QFX5100 and EX4600 Scaling Value                       | QFX10002 Scaling Value                                                 |
|--------------------------------------------------------------------------|------------------------------------------|--------------------------------------------------------|------------------------------------------------------------------------|
| Maximum number of MPLS labels on provider switches                       | 4096                                     | 16386                                                  | 80000 (Junos limit)                                                    |
| Maximum number of tunnel (combination of routes and LSPs) initiations    | Ingress LSPs: 1024<br>Transit LSPs: 4000 | Ingress LSPs: 1024<br>Transit LSPs: 16386              | Ingress LSPs: 32000 (Junos limit)<br>Transit LSPs: 80000 (Junos limit) |
| Maximum number of unique next-hops on egress provider edge (PE) switches | 512                                      | 512                                                    |                                                                        |
| Maximum number of MPLS firewall filters                                  | 768                                      | 1536                                                   | 8000 ingress<br>8000 egress                                            |
| Virtual Routing and Forwarding (VRF)                                     | 1K                                       | 1K                                                     | 4K                                                                     |
| Layer 3 Host                                                             | IPv4: 8K                                 | <i>See Understanding the Unified Forwarding Table.</i> |                                                                        |
| Layer 3 Longest Prefix Match (LPM)                                       | IPv4: 16K<br>IPv6: 4K                    | <i>See Understanding the Unified Forwarding Table.</i> |                                                                        |

- Related Documentation**
- [MPLS Feature Support on QFX Series and EX4600 Switches on page 4934](#)
  - [MPLS Configuration Guidelines on page 5012](#)

## MPLS Stitching For Virtual Machine Connection

By using MPLS, the stitching feature of Junos OS provides connectivity between virtual machines that reside either on opposite sides of data center routers or in different data centers. An external controller, programmed in the data-plane, assigns MPLS labels to both virtual machines and servers. Then, the signaled MPLS labels are used between the data center routers, generating static link switched paths (LSPs), resolved over BGP labeled unicast, RSVP or LDP, to provide the routes dictated by the labels.


- [When Would I Use Stitching? on page 5015](#)
- [How Does MPLS Stitching Work? on page 5015](#)
- [How Do I Configure Stitching? on page 5016](#)
- [Which Switches Support Stitching? on page 5016](#)
- [Q&A on page 5016](#)

## When Would I Use Stitching?

There are several ways to connect virtual machines. One option when you have virtual machines on opposite sides of a router (or different data centers) is to use MPLS stitching. A typical topology for using MPLS stitching is shown in Figure 149.

**Figure 149: Virtual Machines on Either Side of Routers**

ERROR: Unresolved graphic fileref="" not found in  
 "/cmsxml/default/main/supplemental/STAGING/images/".



The above topology consists of the following MPLS layers: VMs | Servers | ToRs | Router  
 ..... Router | ToRs | Servers | VMs



**NOTE:** The label on the left is the top of the label stack.

## How Does MPLS Stitching Work?

With stitching, the MPLS static allocation of labels demultiplexes incoming traffic onto any device/entity in the next layer in the direction of traffic flow. Essentially, there is a label hierarchy that picks up labels for the correct top-of-rack switch, server, and virtual machine that receives traffic. Static label assignments are done between the top-of-rack switches and the virtual machines.

For example, imagine that traffic is sent from VM1 to VM3 in Figure 149. When traffic exits Server1, its label stack is L1 | L2 | L3 where:

- L1 represents the egress top-of-rack switch ToR1.
- L2 represents the physical server, Server2, towards which the egress-side ToR will forward the traffic.
- L3: represents the virtual machine on Server2 to which the Server2 should deliver the traffic.

When traffic arrives at ToR1, it needs to be sent to ToR2. Since ToR1 and ToR2 are not directly connected, traffic must flow from ToR1 to ToR2 using label-switching starting on the outermost (top) label. Stitching has been added to static-LSP functionality to SWAP L1 to a L-BGP label that ToR2 advertises to ToR1. The label stack now must contain another label at the top to enable forwarding of the labeled packets between ToR1 and ToR2. An L-Top label is added if L-BGP is resolved over RSVP/LDP. If static LSP is resolved over L-BGP, then the top label is swapped with the L-BGP label and there is no L-Top label. When the traffic exits ToR1, the stack is: L-top | L-BGP | L2 | L3.

Traffic from ToR1 to ToR2 is then label switched over any signaled LSP.

When traffic arrives at ToR2, the top label is removed with PHP (popped) and the label stack becomes L-BGP | L2 | L3. Since L-BGP is a implicit null label, ToR2 pops the static LSP label L2 that corresponds to the egress server and then forward the packet to the egress server using the static-LSP configuration on ToR2, which corresponds to a single-hop implicit-NULL LSP.

The outgoing stack becomes L3 and the next-hop is the egress server Server2.

When traffic arrives at the egress server Server2, Server2 pops L3 and delivers the packet to VM3.

## How Do I Configure Stitching?

The new keyword **stitch** for LSPs under the command *transit* has been added to resolve the remote next-hop. The **show mpls static-lsp** command has been extended to show the LSP state as 'InProgress' whenever the LSP is waiting for protocol next-hop resolution by resolver.

## Which Switches Support Stitching?

QFX5100, QFX3500 and EX4600 support the static LSP stitching feature.

## Q&A

Q: Is link and node protection for the next-hop provided by MPLS stitching?

A: link and node protection for the next-hop of transit LSP stitched to L-BGP LSP are not needed. That is provided by L-BGP LSP.

Q: Does stitching work with real L-BGP labels?

A: No, stitching does not work when the L-BGP label is a real label.

### Related Documentation

- [MPLS Feature Support on QFX Series and EX4600 Switches on page 4934](#)



# Configuration Statements and Command Statements for MPLS

- [Configuration Statements for MPLS on page 5017](#)
- [Monitoring Commands for MPLS on page 5105](#)

## Configuration Statements for MPLS

---

- [\[edit protocols mpls\] Hierarchy Level on page 5019](#)
- [adaptive on page 5024](#)
- [adjust-interval on page 5025](#)
- [adjust-threshold on page 5025](#)
- [adjust-threshold-overflow-limit on page 5026](#)
- [adjust-threshold-underflow-limit on page 5026](#)
- [admin-down on page 5027](#)
- [advertisement-hold-time on page 5027](#)
- [associate-backup-pe-groups on page 5028](#)
- [auto-bandwidth \(MPLS Tunnel\) on page 5029](#)
- [backup-pe-group on page 5030](#)
- [bandwidth \(Fast Reroute, Signaled, and Multiclass LSPs\) on page 5031](#)
- [bandwidth-model on page 5032](#)
- [bypass \(Static LSP\) on page 5033](#)
- [class-of-service \(Protocols MPLS\) on page 5034](#)
- [corouted-bidirectional on page 5035](#)
- [corouted-bidirectional-passive on page 5035](#)
- [description \(Protocols MPLS\) on page 5036](#)
- [diffserv-te on page 5037](#)
- [disable \(Protocols MPLS\) on page 5038](#)
- [exclude \(for Fast Reroute\) on page 5039](#)
- [exclude-srlg on page 5040](#)

- [explicit-null \(Protocols MPLS\) on page 5041](#)
- [fast-reroute \(Protocols MPLS\) on page 5042](#)
- [forwarding-table on page 5043](#)
- [from \(Protocols MPLS\) on page 5043](#)
- [gpip on page 5044](#)
- [hop-limit on page 5045](#)
- [include-all \(for Fast Reroute\) on page 5046](#)
- [include-any \(for Fast Reroute\) on page 5046](#)
- [ingress \(LSP\) on page 5047](#)
- [install \(Protocols MPLS\) on page 5048](#)
- [interface \(Protocols MPLS\) on page 5049](#)
- [ipv6-tunneling on page 5050](#)
- [label-switched-path \(Protocols MPLS\) on page 5051](#)
- [ldp-tunneling on page 5053](#)
- [link-protection \(Static LSPs\) on page 5054](#)
- [log-updown \(Protocols MPLS\) on page 5055](#)
- [lsp-attributes on page 5056](#)
- [maximum-bandwidth \(Protocols MPLS\) on page 5057](#)
- [metric \(Protocols MPLS\) on page 5057](#)
- [minimum-bandwidth on page 5058](#)
- [monitor-bandwidth on page 5058](#)
- [mtu-signaling on page 5059](#)
- [no-cspf on page 5060](#)
- [no-decrement-ttl on page 5061](#)
- [no-install-to-address on page 5062](#)
- [no-propagate-ttl on page 5063](#)
- [record on page 5064](#)
- [no-trap on page 5065](#)
- [node-link-protection \(Protocols MPLS\) on page 5066](#)
- [oam \(Protocols MPLS\) on page 5067](#)
- [optimize-aggressive on page 5068](#)
- [optimize-hold-dead-delay on page 5069](#)
- [optimize-switchover-delay on page 5070](#)
- [optimize-timer \(Protocols MPLS\) on page 5071](#)
- [p2mp \(Protocols MPLS\) on page 5072](#)
- [path \(Protocols MPLS\) on page 5073](#)
- [path-mtu on page 5074](#)

- [policing \(Protocols MPLS\) on page 5075](#)
- [policy-statement on page 5076](#)
- [pop on page 5080](#)
- [preference \(Protocols MPLS\) on page 5081](#)
- [primary \(Protocols MPLS\) on page 5082](#)
- [push on page 5083](#)
- [record on page 5084](#)
- [retry-limit on page 5085](#)
- [revert-timer on page 5086](#)
- [rsvp-error-hold-time on page 5087](#)
- [secondary \(Protocols MPLS\) on page 5088](#)
- [select on page 5089](#)
- [signal-bandwidth on page 5089](#)
- [smart-optimize-timer on page 5090](#)
- [standby on page 5091](#)
- [static-label-switched-path on page 5092](#)
- [statistics \(Protocols MPLS\) on page 5094](#)
- [swap on page 5095](#)
- [switching-type on page 5096](#)
- [te-class-matrix on page 5097](#)
- [traffic-engineering \(Protocols MPLS\) on page 5098](#)
- [transit-lsp-association on page 5099](#)
- [te-class-matrix on page 5100](#)
- [to on page 5101](#)
- [traceoptions \(Protocols MPLS\) on page 5102](#)
- [traffic-engineering \(Protocols MPLS\) on page 5104](#)
- [transit-lsp-association on page 5105](#)

### [edit protocols mpls] Hierarchy Level

This topic lists the supported configuration statements at the **[edit protocols mpls]** hierarchy level on the QFX Series and on the EX4600 switch. For more information about these statements, see the *Junos OS MPLS Applications Library for Routing Devices*.



**NOTE:** Both the command-line interface (CLI) on QFX Series devices and the CLI on the EX4600 switch display MPLS configuration statements that are not supported. However, configuring unsupported statements on these devices is ignored and has no effect on the operation of the device.

```
protocols {
```

```
mpls {
 admin-down;
 advertisement-hold-time seconds;
 auto-policing {
 class all (drop | loss-priority-high | loss-priority-low);
 class ctnumber (drop | loss-priority-high | loss-priority-low);
 }
 class-of-service cos-value;
 diffserv-te {
 bandwidth-model {
 extended-mam;
 mam;
 rdm;
 }
 te-class-matrix {
 tnumber {
 priority priority;
 traffic-class {
 ctnumber priority priority;
 }
 }
 }
 }
 disable;
 exclude-srlg;
 explicit-null;
 hop-limit number;
 interface (interface-name | all) {
 disable;
 }
 ipv6-tunneling;
 label-switched-path lsp-name {
 adaptive;
 admin-down;
 associate-backup-pe-groups;
 associate-lsp lsp-name {
 from from-ip-address;
 }
 auto-bandwidth {
 adjust-interval seconds;
 adjust-threshold percentage;
 maximum-bandwidth bps;
 minimum-bandwidth bps;
 monitor-bandwidth;
 }
 backup;
 bandwidth bps {
 ct0 bps;
 ct1 bps;
 ct2 bps;
 ct3 bps;
 }
 class-of-service cos-value;
 corouted-bidirectional;
 corouted-bidirectional-passive;
 description text;
```

```

disable;
exclude-srlg;
fast-reroute {
 (bandwidth bps | bandwidth-percent percentage);
 (exclude [group-names] | no-exclude);
 hop-limit number;
 (include-all [group-names] | no-include-all);
 (include-any [group-names] | no-include-any);
}
from address;
hop-limit number;
install {
 destination-prefix/prefix-length <active>;
}
inter-domain;
ldp-tunneling;
link-protection
lsp-attributes {
 encoding-type (ethernet | packet | pdh | sonet-sdh);
 gpid (ethernet | hdlc | ipv4 | pos-scrambling-crc-16 | pos-no-scrambling-crc-16 |
 pos-scrambling-crc-32 | pos-no-scrambling-crc-32 | ppp);
 signal-bandwidth type;
 switching-type (fiber | lambda | psc-1 | tdm);
}
metric metric;
no-cspf;
no-decrement-ttl;
no-install-to-address;
no-record;
no-record;
node-link-protection
oam{
 lsp-ping-interval seconds;
 mpls-tp-mode seconds;
 traceoptions {
 file filename <files number> <size maximum-file-size> <world-readable |
 no-world-readable>;
 flag flag;
 no-remote-trace;
 }
}
optimize-hold-dead-delay seconds;
optimize-timer seconds;
p2mp lsp-name;
policing {
 filter filter-name;
 no-auto-policing;
}
preference preference;
primary path-name {
 adaptive;
 class-of-service cos-value;
 hop-limit number;
 no-cspf;
 no-decrement-ttl;
 optimize-timer seconds;

```

```
 preference preference;
 (record | no-record);
 select (manual | unconditional);
 standby;
 }
 (record | no-record);
 retry-limit number;
 retry-timer seconds;
 revert-timer seconds;
 secondary path-name {
 adaptive;
 bandwidth bps {
 ct0 bps;
 ct1 bps;
 ct2 bps;
 ct3 bps;
 }
 class-of-service cos-value;
 hop-limit number;
 no-cspf;
 no-decrement-ttl;
 optimize-timer seconds;
 preference preference;
 (record | no-record);
 select (manual | unconditional);
 standby;
 }
 standby;
 jtemplate;
 to address;
 traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
 }
}
log-updown {
 no-trap {
 mpls-lsp-traps;
 rfc3812-traps;
 }
 (syslog | no-syslog);
 trap;
 trap-path-down;
 trap-path-up;
}
mib-mpls-show-p2mp;
no-cspf;
no-decrement-ttl;
no-propagate-ttl;
no-record;
oam {
 lsp-ping-interval seconds;
 mpls-tp-mode seconds;
 traceoptions {
 file filename <files number> <size maximum-file-size> <world-readable |
 no-world-readable>;
```

```

 flag flag;
 no-remote-trace;
 }
}
optimize-aggressive;
optimize-hold-dead-delay;
optimize-switchover-delay;
optimize-timer;
path path-name {
 (address | hostname) <loose | strict>;
}
path-mtu {
 {
 mtu-signaling;
 }
}
preference;
record;
revert-timer;
rsvp-error-hold-time;
smart-optimize-timer;
standby;
static-label-switched-path lsp-name {
 bypass bypass-name {
 description string;
 next-hop (address | interface-name | address/interface-name);
 push out-label;
 to address;
 }
 ingress {
 class-of-service cos-value;
 description string;
 install {
 destination-prefix <active>;
 }
 link-protection
 metric metric;
 next-hop (address | interface-name | address/interface-name);
 node-protection
 no-install-to-address;
 policing {
 filter filter-name;
 no-auto-policing;
 }
 preference preference;
 push out-label;
 to address;
 }
 transit incoming-label {
 description string;
 link-protection
 next-hop (address | interface-name | address/interface-name);
 node-protection
 pop;
 swap out-label;
 }
}

```

```

statistics {
 auto-bandwidth;
 file filename <files number> <size maximum-file-size> <world-readable |
 no-world-readable>;
 interval seconds;
}
traceoptions {
 file filename <files number> <size maximum-file-size> <world-readable |
 no-world-readable>;
 flag flag;
}
traffic-engineering;
transit-lsp-association transit-association-lsp-group-name {
 from-1 address-of-associated-lsp-1;
 from-2 address-of-associated-lsp-2;
 lsp-name-1 name-of-associated-lsp-1;
 lsp-name-2 name-of-associated-lsp-2;
}
}
}

```

**Related Documentation**

- *Junos OS MPLS Applications Library for Routing Devices*

## adaptive

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>adaptive;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i><br>( <a href="#">primary</a>   <a href="#">secondary</a> ) <i>path-name</i> ],<br>[edit protocols mpls label-switched-path <i>lsp-name</i> ],<br>[edit protocols mpls label-switched-path <i>lsp-name</i> ( <a href="#">primary</a>   <a href="#">secondary</a> ) <i>path-name</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | During reroute, do not double-count bandwidth on links shared by the old and new paths. Including this statement causes RSVP to use shared explicit (SE) reservation styles and assists in smooth transition during rerouting.                                                                                                                                                                                                                                                                     |
| <b>Default</b>                  | The configured object is disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Adaptive LSPs</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                               |



## adjust-interval

|                                 |                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>adjust-interval <i>seconds</i>;</code>                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth],<br>[edit protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.                                                                    |
| <b>Description</b>              | Specify the bandwidth reallocation interval.                                                                                                                                                     |
| <b>Options</b>                  | <b><i>seconds</i></b> —Bandwidth reallocation interval, in seconds.<br><b>Range:</b> 300 through 315,360,000 seconds<br><b>Default:</b> 86,400 seconds                                           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Automatic Bandwidth Allocation for LSPs</i></li> </ul>                                                                                   |

## adjust-threshold

|                                 |                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>adjust-threshold <i>percent</i>;</code>                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth],<br>[edit protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth]                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.                                                                                                                                                                                                                  |
| <b>Description</b>              | Specify how sensitive the automatic bandwidth adjustment for a label-switched path (LSP) is to changes in bandwidth utilization.                                                                                                                                                                                                               |
| <b>Options</b>                  | <b><i>percent</i></b> —Bandwidth demand for the current bandwidth adjustment interval is determined and compared to the LSP's current bandwidth allocation. If the percentage difference in bandwidth is greater than or equal to the percentage specified by this statement, the LSP's bandwidth is adjusted to the current bandwidth demand. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Automatic Bandwidth Allocation for LSPs</i></li> </ul>                                                                                                                                                                                                                                 |

## adjust-threshold-overflow-limit

---

|                                 |                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | adjust-threshold-overflow-limit <i>number</i> ;                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth],<br>[edit protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.5.<br>Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.                                                                        |
| <b>Description</b>              | Specify the number of consecutive bandwidth overflow samples before triggering a bandwidth adjustment.                                                                                           |
| <b>Options</b>                  | <b><i>number</i></b> —Number of consecutive bandwidth overflow samples.<br><b>Range:</b> 1 through 65,535<br><b>Default:</b> This feature is disabled by default.                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Automatic Bandwidth Allocation for LSPs</i></li></ul>                                                                                     |

## adjust-threshold-underflow-limit

---

|                                 |                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | adjust-threshold-underflow-limit <i>number</i> ;                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth],<br>[edit protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.3.<br>Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.                                                                       |
| <b>Description</b>              | Specify the number of consecutive bandwidth underflow samples before triggering a bandwidth adjustment.                                                                                          |
| <b>Options</b>                  | <b><i>number</i></b> —Number of consecutive bandwidth underflow samples.<br><b>Range:</b> 1 through 65,535<br><b>Default:</b> This feature is disabled by default.                               |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Automatic Bandwidth Allocation for LSPs</i></li></ul>                                                                                     |

## admin-down

---

|                                 |                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | admin-down;                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> ],<br>[edit protocols mpls label-switched-path <i>lsp-name</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.2.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                |
| <b>Description</b>              | Set a nonpacket GMPLS LSP to the administrative down state. This statement does not affect control path setup or data forwarding for packet LSPs.                    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Allowing Non-Packet GMPLS LSPs to Establish Paths Through Routers Running the Junos OS</i></li> </ul>                    |

## advertisement-hold-time

---

|                                 |                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | advertisement-hold-time <i>seconds</i> ;                                                                                  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls],<br>[edit protocols mpls]                                |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series. |
| <b>Description</b>              | Do not advertise when the LSP goes from up to down, for a certain period of time known as the hold time.                  |
| <b>Options</b>                  | <i>seconds</i> —Hold time, in seconds.<br><b>Range:</b> 0 through 65,535 seconds<br><b>Default:</b> 5 seconds             |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Damping Advertisement of LSP State Changes</i></li> </ul>                     |

## associate-backup-pe-groups

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | associate-backup-pe-groups;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> ],<br>[edit protocols mpls label-switched-path <i>lsp-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Enable an LSP to monitor the status of its destination PE router. You can configure multiple backup PE router groups using the same router's address. Backup PE router groups provide ingress PE router redundancy when point-to-multipoint LSPs are configured for multicast distribution. A failure of this LSP indicates to all of the backup PE router groups that the destination PE router is down. This statement is not tied to a specific backup PE router group. It applies to all groups that are interested in the status of the LSP to the destination address. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Enabling Point-to-Point LSPs to Monitor Egress PE Routers</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## auto-bandwidth (MPLS Tunnel)


|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> auto-bandwidth {   adjust-interval <i>seconds</i>;   adjust-threshold <i>percent</i>;   adjust-threshold-activate-bandwidth <i>bps</i>   adjust-threshold-overflow-limit <i>number</i>;   adjust-threshold-underflow-limit <i>number</i>;   maximum-bandwidth <i>bps</i>;   minimum-bandwidth <i>bps</i>;   minimum-bandwidth-adjust-interval   minimum-bandwidth-adjust-threshold-change   minimum-bandwidth-adjust-threshold-value   monitor-bandwidth; } </pre> |
| <b>Hierarchy Level</b>          | [edit protocols mpls label-switched-path <i>lsp-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Allow an MPLS tunnel to automatically adjust its bandwidth allocation based on the volume of traffic flowing through the tunnel.                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | The statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Automatic Bandwidth Allocation for LSPs</i></li> <li>• <a href="#">request mpls lsp adjust-autobandwidth on page 5122</a></li> </ul>                                                                                                                                                                                                                                                                             |

## backup-pe-group

---

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <pre>backup-pe-group <i>group-name</i> {<br/>    backups [ <i>addresses</i> ];<br/>    local-address <i>address</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                        |
| Hierarchy Level          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],<br>[edit logical-systems <i>logical-system-name</i> routing-options multicast],<br>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],<br>[edit routing-options multicast]                                                                                                                                                                                     |
| Release Information      | Statement introduced in Junos OS Release 9.0.<br>Statement introduced in Junos OS Release 9.5 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                         |
| Description              | Configure a backup provider edge (PE) group for ingress PE redundancy when point-to-multipoint label-switched paths (LSPs) are used for multicast distribution.                                                                                                                                                                                                                                                                                                                                                         |
| Options                  | <p><b>backups <i>addresses</i></b>—Specify the address of backup PE routers for ingress PE redundancy when point-to-multipoint LSPs are used for multicast distribution.</p> <p><b>local-address <i>address</i></b>—Specify the address of the local PE router for ingress PE redundancy when point-to-multipoint LSPs are used for multicast distribution.</p> <p><b><i>pe-group-name</i></b>—Specify the name for the group of PE routers that provide ingress PE router redundancy for point-to-multipoint LSPs.</p> |
| Required Privilege Level | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                     |
| Related Documentation    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Ingress PE Redundancy</i></li><li>• <i>Configuring Ingress PE Router Redundancy for Point-to-Multipoint LSPs</i></li></ul>                                                                                                                                                                                                                                                                                                                              |

## bandwidth (Fast Reroute, Signaled, and Multiclass LSPs)

|                                                                                                                                                                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                             | <pre> bandwidth <i>bps</i> {     ct0 <i>bps</i>;     ct1 <i>bps</i>;     ct2 <i>bps</i>;     ct3 <i>bps</i>; } </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>                                                                                                                                                    | <pre> [edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> <i>fast-reroute</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> (primary   secondary) <i>path-name</i>], [edit protocols mpls], [edit protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i> <i>fast-reroute</i>], [edit protocols mpls label-switched-path <i>lsp-name</i> (primary   secondary) <i>path-name</i>] </pre> |
| <b>Release Information</b>                                                                                                                                                | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>                                                                                                                                                        | <p>When configuring an LSP, specify the traffic rate associated with the LSP.</p> <p>When configuring fast reroute, allocate bandwidth for the reroute path. By default, no bandwidth is reserved for the rerouted path. The fast reroute bandwidth does not need to be identical to that allocated for the LSP itself.</p> <p>When configuring a multiclass LSP, use the <b><i>ctnumber bandwidth</i></b> statements to specify the bandwidth to be allocated for each class type.</p>                                                                                                                                                                                                                                       |
| <b>Options</b>                                                                                                                                                            | <p><b><i>bps</i></b>—Bandwidth, in bits per second. You can specify this as an integer value. You can also use the abbreviations <b>k</b> (for a thousand), <b>m</b> (for a million), or <b>g</b> (for a billion).</p> <p><b>Range:</b> Any positive integer</p> <p><b>Default:</b> 0 (no bandwidth is reserved)</p>                                                                                                                                                                                                                                                                                                                                                                                                          |
| <div>  <b>NOTE:</b> On the ACX Series, <i>bps</i> is the only supported option. </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|                                                                                                                                                                           | <p><b><i>ctnumber bps</i></b>—Bandwidth for the specified class type, in bits per second. You can specify this as an integer value. If you do so, count your zeros carefully, or you can use the abbreviations <b>k</b> (for a thousand), <b>m</b> (for a million), or <b>g</b> (for a billion [also called a thousand million]).</p> <p><b>Range:</b> Any positive integer</p> <p><b>Default:</b> 0 (no bandwidth is reserved)</p>                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b>                                                                                                                                           | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

- Related Documentation**
- *Configuring Fast Reroute*
  - *Configuring the Bandwidth Value for LSPs*
  - *Configuring LSPs for DiffServ-Aware Traffic Engineering*
  - *Configuring Multiclass LSPs*

---

## bandwidth-model

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>bandwidth-model {<br/>    extended-mam;<br/>    mam;<br/>    rdm;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls diffserv-te],<br>[edit protocols mpls diffserv-te]                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Configure the bandwidth model for differentiated services. Note that you cannot configure both bandwidth models at the same time.                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <p><b>extended-mam</b>—The extended maximum allocation model (MAM) is a bandwidth model based on MAM.</p> <p><b>mam</b>—The MAM is defined in RFC 4125, <i>Maximum Allocation Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering</i>.</p> <p><b>rdm</b>—The Russian dolls bandwidth allocation model (RDM) is defined in RFC 4127, <i>Russian Dolls Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering</i>. RDM makes efficient use of bandwidth by allowing the class types to share bandwidth.</p> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Routers for DiffServ-Aware Traffic Engineering</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                            |



## bypass (Static LSP)

|                                 |                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>bypass bypass-name {   bandwidth bps;   description string;   next-hop (address   interface-name   address/interface-name);   push out-label;   to address; }</pre>                                                                                                           |
| <b>Hierarchy Level</b>          | <pre>[edit logical-systems logical-system-name protocols mpls static-label-switched-path   lsp-name], [edit protocols mpls static-label-switched-path lsp-name]</pre>                                                                                                              |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 10.1.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>                                                                                                                                              |
| <b>Description</b>              | <p>Configure specific bandwidth and path constraints for a bypass ingress LSP. It is possible to configure multiple bypass LSPs individually. If you do not, they all share the same path and bandwidth constraints.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Static LSPs</i></li> </ul>                                                                                                                                                                                                 |

## class-of-service (Protocols MPLS)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>class-of-service cos-value;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols mpls],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path</code><br><code>  <i>lsp-name</i> ingress],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols mpls <a href="#">label-switched-path</a> <i>lsp-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols mpls <a href="#">label-switched-path</a> <i>lsp-name</i></code><br><code>  (<a href="#">primary</a>   <a href="#">secondary</a>) <i>path-name</i>],</code><br><code>[edit protocols mpls],</code><br><code>[edit protocols mpls <a href="#">label-switched-path</a> <i>lsp-name</i>],</code><br><code>[edit protocols mpls <a href="#">label-switched-path</a> <i>lsp-name</i> (<a href="#">primary</a>   <a href="#">secondary</a>) <i>path-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path</code><br><code>  <i>lsp-name</i> ingress]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Class-of-service (CoS) value given to all packets in the LSP.<br><br>The CoS value might affect the scheduling or queuing algorithm of traffic traveling along an LSP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <b>cos-value</b> —CoS value. A higher value typically corresponds to a higher level of service.<br><b>Range:</b> 0 through 7<br><b>Default:</b> If you do not specify a CoS value, the IP precedence bits from the packet's IP header are used as the packet's CoS value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Class of Service for MPLS LSPs</i></li><li>• <i>Configuring the Ingress Router for Static LSPs</i></li><li>• <i>Configuring Static LSPs</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

## corouted-bidirectional

|                                 |                                                                                                                                                                                                          |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | corouted-bidirectional;                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> ],<br>[edit protocols mpls label-switched-path <i>lsp-name</i> ]                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                   |
| <b>Description</b>              | Specify that the label-switched path be established as a corouted bidirectional packet LSP. You cannot configure this statement at the same time as the <b>corouted-bidirectional-passive</b> statement. |
| <b>Default</b>                  | This statement is disabled by default.                                                                                                                                                                   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Corouted Bidirectional LSPs</i></li> <li>• <a href="#">corouted-bidirectional-passive on page 5035</a></li> </ul>                                |

## corouted-bidirectional-passive

|                                 |                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | corouted-bidirectional-passive;                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> ],<br>[edit protocols mpls label-switched-path <i>lsp-name</i> ]                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                              |
| <b>Description</b>              | Specify that the label-switched path be a passive LSP associated with a bidirectional LSP when it is signaled at the ingress router. This passive LSP enables the MPLS application to utilize the reverse LSP. You cannot configure this statement at the same time as the <b>corouted-bidirectional</b> statement. |
| <b>Default</b>                  | This statement is disabled by default.                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Corouted Bidirectional LSPs</i></li> <li>• <a href="#">corouted-bidirectional on page 5035</a></li> </ul>                                                                                                                                                   |

## description (Protocols MPLS)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>description text;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols mpls</code><br><code>static-label-switched-path <i>lsp-name</i> bypass],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols mpls</code><br><code>static-label-switched-path <i>lsp-name</i> ingress],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols mpls</code><br><code>static-label-switched-path <i>lsp-name</i> transit <i>incoming-label</i>],</code><br><code>[edit protocols mpls label-switched-path <i>lsp-name</i>],</code><br><code>[edit protocols mpls static-label-switched-path <i>lsp-name</i> bypass],</code><br><code>[edit protocols mpls static-label-switched-path <i>lsp-name</i> ingress],</code><br><code>[edit protocols mpls static-label-switched-path <i>lsp-name</i> transit <i>incoming-label</i>]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Provides a textual description of the LSP. Enclose any descriptive text that includes spaces in quotation marks (" "). Any descriptive text you include is displayed in the output of the <b>show mpls lsp detail</b> command and has no effect on the operation of the LSP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <b>text</b> —Provide a textual description of the LSP. The description text can be no more than 80 characters in length.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring a Text Description for LSPs</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## diffserv-te

|                                 |                                                                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> diffserv-te {   bandwidth-model {     extended-mam;     mam;     rdm;   }   te-class-matrix {     tnumber {       priority <i>priority</i>;       traffic-class {         ctnumber <i>priority priority</i>;       }     }   } } </pre> |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls],<br>[edit protocols mpls]                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                     |
| <b>Description</b>              | Specify properties for differentiated services in traffic engineering.                                                                                                                                                                        |
| <b>Options</b>                  | The statements are explained separately.                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Configuring Routers for DiffServ-Aware Traffic Engineering</i></li> </ul>                                                                                                                           |

## disable (Protocols MPLS)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | disable;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls],<br>[edit logical-systems <i>logical-system-name</i> protocols mpls <a href="#">interface</a> <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> protocols mpls <a href="#">label-switched-path</a> <i>lsp-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> protocols mpls <a href="#">label-switched-path</a> <i>lsp-name</i> auto-bandwidth],<br>[edit protocols mpls],<br>[edit protocols mpls <a href="#">interface</a> <i>interface-name</i> ],<br>[edit protocols mpls <a href="#">label-switched-path</a> <i>lsp-name</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Disable the functionality of the configured object.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Default</b>                  | The configured object is enabled (operational) unless explicitly disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Minimum MPLS Configuration</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

## exclude (for Fast Reroute)

---

|                                 |                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | (exclude [ <i>group-names</i> ]   no-exclude);                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> fast-reroute],<br>[edit protocols mpls label-switched-path <i>lsp-name</i> fast-reroute]                                                   |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series and for EX4600 switches.                                                                                          |
| <b>Description</b>              | Control exclusion of administrative groups: <ul style="list-style-type: none"><li>• <b>exclude</b>—Define the administrative groups to exclude for fast reroute.</li><li>• <b>no-exclude</b>—Disable administrative group exclusion.</li></ul> |
| <b>Options</b>                  | <b><i>group-names</i></b> —Names of one or more groups defined with the <b>admin-groups</b> statement.                                                                                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Fast Reroute</i></li><li>• <i>admin-groups</i></li></ul>                                                                                                                                |


## exclude-srlg

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | exclude-srlg;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit protocols mpls],<br>[edit logical-systems logical-system-name protocols mpls],<br>[edit protocols mpls label-switched-path <i>path-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>path-name</i> ],<br>[edit protocols rsvp interface <i>interface-name</i> link-protection],<br>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection],<br>[edit protocols rsvp interface <i>interface-name</i> link-protection bypass <i>destination</i> ],<br>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection bypass <i>destination</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | <p>Exclude Shared Risk Link Group (SRLG) links for the secondary path for critical links where it is imperative to keep the secondary and primary label-switched paths completely disjoint from any common SRLG.</p> <p>When specified, the Constrained Shortest Path First (CSPF) algorithm excludes any link belonging to the set of SRLGs in the primary path. When not specified and if a link belongs to the set of SRLGs in the primary path, CSPF adds the SRLG cost to the metric, but still accepts the link for computing the path.</p>                                                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Excluding SRLG Links Completely for the Secondary LSP</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |



## explicit-null (Protocols MPLS)

|                                                                                                                                                                                                               |                                                                                                                                                                                                                                                      |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                 | explicit-null;                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>                                                                                                                                                                                        | [edit logical-systems <i>logical-system-name</i> protocols mpls],<br>[edit protocols mpls]                                                                                                                                                           |
| <b>Release Information</b>                                                                                                                                                                                    | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric. |
| <b>Description</b>                                                                                                                                                                                            | Advertise label 0 to the egress router of an LSP.                                                                                                                                                                                                    |
| <b>Default</b>                                                                                                                                                                                                | If you do not include the <b>explicit-null</b> statement in the MPLS configuration, label 3 (implicit null) is advertised.                                                                                                                           |
| <div>  <b>NOTE:</b> Junos OS does not support explicit null routes with next hops to virtual tunnel (vt-) interfaces. </div> |                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b>                                                                                                                                                                               | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                  |
| <b>Related Documentation</b>                                                                                                                                                                                  | <ul style="list-style-type: none"> <li>• <i>Configuring RSVP to Pop the Label on the Ultimate-Hop Router</i></li> </ul>                                                                                                                              |

## fast-reroute (Protocols MPLS)

---

|                                 |                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>fast-reroute {<br/>  (bandwidth <i>bps</i>   bandwidth-percent <i>percentage</i>);<br/>  (exclude [ <i>group-names</i> ]   no-exclude );<br/>  hop-limit <i>number</i>;<br/>  (include-all [ <i>group-names</i> ]   no-include-all);<br/>  (include-any [ <i>group-names</i> ]   no-include-any);<br/>}</pre>                      |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> ],<br>[edit protocols mpls label-switched-path <i>lsp-name</i> ]                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series and for EX4600 switches.<br>Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.                                                            |
| <b>Description</b>              | Establish detours for the LSP so that if a node or link in the LSP fails, the traffic on the LSP can be rerouted with minimal packet loss.                                                                                                                                                                                              |
| <b>Options</b>                  | The statements are explained separately.                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Fast Reroute</a></li><li>• <a href="#">Fast Reroute Overview on page 4967</a></li><li>• <a href="#">MPLS Feature Support on QFX Series and EX4600 Switches on page 4934</a></li><li>• <a href="#">Interprovider and Carrier-of-Carriers VPNs on page 4966</a></li></ul> |

## forwarding-table

|                                 |                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | forwarding-table {<br><b>export</b> [ <i>policy--names</i> ];<br>( <b>indirect-next-hop</b>   no-indirect-next-hop);<br>}                     |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-options],<br>[edit routing-options]                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| <b>Description</b>              | Configure information about the routing device's forwarding table.<br><br>The remaining statements are explained separately.                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Per-Packet Load Balancing on page 3228</a></li> </ul>                        |

## from (Protocols MPLS)

|                                 |                                                                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | from <i>address</i> ;                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> ],<br>[edit protocols mpls label-switched-path <i>lsp-name</i> ]                                                                                 |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric. |
| <b>Description</b>              | Specify the source address to use for the LSP.<br><br>The address you specify does not affect the outgoing interface used by the LSP.                                                                                                                |
| <b>Default</b>                  | If you do not include this statement, the software automatically selects the loopback interface as the address.                                                                                                                                      |
| <b>Options</b>                  | <b>address</b> —IP address.                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring the Ingress and Egress Router Addresses for LSPs</a></li> </ul>                                                                                                                     |

## gp-id

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>gp-id (ethernet   hdlc   ipv4   pos-scrambling-crc-16   pos-no-scrambling-crc-16   pos-scrambling-crc-32   pos-no-scrambling-crc-32   ppp);</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> lsp-attributes],<br>[edit protocols mpls label-switched-path <i>lsp-name</i> lsp-attributes]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br><b>pos-scrambling-crc-16</b> , <b>pos-no-scrambling-crc-16</b> , <b>pos-scrambling-crc-32</b> , and <b>pos-no-scrambling-crc-32</b> options added in Junos OS Release 8.0.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Specify the type of payload carried by the LSP. It can be any of the following: <ul style="list-style-type: none"><li>• <b>ethernet</b>—Ethernet (GPID value: 33)</li><li>• <b>hdlc</b>—High-level Data Link Control (HDLC) (GPID value: 44)</li><li>• <b>ipv4</b>—IP version 4 (GPID value: 0x0800)</li><li>• <b>pos-no-scrambling-crc-16</b>—for interoperability with other vendors' equipment (GPID value: 29)</li><li>• <b>pos-no-scrambling-crc-32</b>—for interoperability with other vendors' equipment (GPID value: 30)</li><li>• <b>pos-scrambling-crc-16</b>—for interoperability with other vendors' equipment (GPID value: 31)</li><li>• <b>pos-scrambling-crc-32</b>—for interoperability with other vendors' equipment (GPID value: 32)</li><li>• <b>ppp</b>—Point-to-Point Protocol (PPP) (GPID value: 50)</li></ul> |
| <b>Default</b>                  | <code>ipv4</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring MPLS LSPs for GMPLS</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

## hop-limit

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>hop-limit <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols mpls],<br/> [edit logical-systems <i>logical-system-name</i> protocols mpls <a href="#">label-switched-path</a> <i>lsp-name</i>],<br/> [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> <a href="#">fast-reroute</a>],<br/> [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> (<a href="#">primary</a>   <a href="#">secondary</a>) <i>path-name</i>],<br/> [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection],<br/> [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>],<br/> [edit protocols mpls],<br/> [edit protocols mpls <a href="#">label-switched-path</a> <i>lsp-name</i>],<br/> [edit protocols mpls label-switched-path <i>lsp-name</i> <a href="#">fast-reroute</a>],<br/> [edit protocols mpls label-switched-path <i>lsp-name</i> (<a href="#">primary</a>   <a href="#">secondary</a>) <i>path-name</i>],<br/> [edit protocols rsvp interface <i>interface-name</i> link-protection],<br/> [edit protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | <p>Specify the maximum number of routers that an LSP can traverse. This limit can be applied to any of the following:</p> <ul style="list-style-type: none"> <li>• LSPs—The configured hop limit includes the ingress and egress routers. You can specify a hop limit for an LSP and for both primary and secondary paths.</li> <li>• Fast reroute detour—Specify the number of additional routers a fast reroute detour can traverse relative to the protected LSP. For example, if an LSP traverses 4 routers, any detour for the LSP can be no more than 10 router hops, including the ingress and egress routers.</li> <li>• Link protection bypass—Specify the maximum number of routers that a link protection bypass can traverse.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <p><b><i>number</i></b>—Maximum number of hops.</p> <p><b>Range:</b> 2 through 255 (for an LSP or for a link protection bypass); 0 through 255 (for fast reroute)</p> <p><b>Default:</b> 255 (for an LSP or for a link protection bypass); 6 (for fast reroute)</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Fast Reroute</i></li> <li>• <i>Limiting the Number of Hops in LSPs</i></li> <li>• <i>Configuring Link Protection on Interfaces Used by LSPs</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## include-all (for Fast Reroute)

---

|                                 |                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | (include-all [ <i>group-names</i> ]   no-include-all);                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> fast-reroute],<br>[edit protocols mpls label-switched-path <i>lsp-name</i> fast-reroute]                                                                          |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series and for EX4600 switches.                                                                                                                 |
| <b>Description</b>              | Control inclusion of administrative groups: <ul style="list-style-type: none"><li>• <b>include-all</b>—Define the administrative groups that must all be included for fast reroute.</li><li>• <b>no-include-all</b>—Disable administrative group inclusion.</li></ul> |
| <b>Options</b>                  | <b>group-names</b> —One or more names of groups defined with the <b>admin-groups</b> statement.                                                                                                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Fast Reroute</i></li></ul>                                                                                                                                                                                     |

## include-any (for Fast Reroute)

---

|                                 |                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | (include-any [ <i>group-names</i> ]   no-include-any);                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> fast-reroute],<br>[edit protocols mpls label-switched-path <i>lsp-name</i> fast-reroute]                                                           |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series and for EX4600 switches.                                                                                                  |
| <b>Description</b>              | Control inclusion of administrative groups: <ul style="list-style-type: none"><li>• <b>include-any</b>—Define the administrative groups to include for fast reroute.</li><li>• <b>no-include-any</b>—Disable administrative group inclusion.</li></ul> |
| <b>Options</b>                  | <b>group-names</b> —One or more names of groups defined with the <b>admin-groups</b> statement.                                                                                                                                                        |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Fast Reroute</i></li></ul>                                                                                                                                                                      |

## ingress (LSP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> ingress {     bandwidth <i>bps</i>;     class-of-service <i>cos-value</i>;     description <i>string</i>;     entropy-label;     install {         destination-prefix &lt;active&gt;;     }     link-protection bypass-name <i>name</i>;     metric <i>metric</i>;     next-hop (<i>address</i>   <i>interface-name</i>   <i>address/interface-name</i>);     node-protection bypass-name <i>name</i> next-next-label <i>label</i>;     no-install-to-address;     policing {         filter <i>filter-name</i>;         no-auto-policing;     }     preference <i>preference</i>;     push <i>out-label</i>;     to <i>address</i>; } </pre> |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i>],</p> <p>[edit protocols mpls static-label-switched-path <i>lsp-name</i>]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 10.1.</p> <p><b>entropy-label</b> option introduced in Junos OS Release 14.1.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | <p>Configure an ingress LSR for a static LSP.</p> <p>The remaining statements are explained separately</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Static LSPs</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## install (Protocols MPLS)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>install {<br/>    <i>destination-prefix</i> &lt;active&gt;;<br/>}</pre>                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i> ingress],<br>[edit protocols mpls label-switched-path <i>lsp-name</i> ],<br>[edit protocols mpls static-label-switched-path <i>lsp-name</i> ingress] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.                                                                                                                      |
| <b>Description</b>              | Associate one or more prefixes with an LSP. When the LSP is up, all the prefixes are installed as entries into the inet.3 or inet6.3 routing table.                                                                                                                                                                                                                       |
| <b>Options</b>                  | <b>active</b> —(Optional) Install the route into the inet.0 or inet6.0 routing table. This allows you to issue a <b>ping</b> or <b>traceroute</b> command on this address.<br><br><b><i>destination-prefix</i></b> —IPv4 or IPv6 address to associate with the LSP.                                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Adding LSP-Related Routes to the inet.3 or inet6.3 Routing Table</i></li></ul>                                                                                                                                                                                                                                                 |



## interface (Protocols MPLS)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>interface (<i>interface-name</i>   all) {   disable;   admin-group [ <i>group-names</i> ];   srlg <i>srlg-name</i>; }</pre>                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls],<br>[edit protocols mpls]                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.</p>                                                                                                                                 |
| <b>Description</b>              | Enable MPLS on one or more interfaces.                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <p><b><i>interface-name</i></b>—Name of the interface on which to configure MPLS. To configure all interfaces, specify <b>all</b>. For details about specifying interfaces, see the <i>Junos OS Network Interfaces Library for Routing Devices</i>.</p> <p><b>srlg <i>srlg-name</i></b>—Name of the SRLG to associate with an interface.</p> <p>The remaining options are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Minimum MPLS Configuration</i></li> <li>• <i>Configuring Static LSPs</i></li> <li>• <i>Example: Configuring SRLG</i></li> </ul>                                                                                                                                                                                                                         |

## ipv6-tunneling

---

|                                 |                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | ipv6-tunneling;                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls],<br>[edit protocols mpls]                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 14.1X53-D30 for QFX Series switches.                                                                                                                                                                      |
| <b>Description</b>              | Allow IPv6 routes to be resolved over an MPLS network by converting LDP and RSVP routes stored in the inet.3 routing table to IPv4-mapped IPv6 addresses and then copying them into the inet6.3 routing table. This routing table can be used to resolve next hops for both inet6 and inet6-vpn routes. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Tunneling IPv6 Traffic over MPLS IPv4 Networks</i></li></ul>                                                                                                                                                                                        |

## label-switched-path (Protocols MPLS)

```

Syntax label-switched-path lsp-name {
 disable;
 adaptive;
 admin-down;
 admin-group {
 exclude [group-names];
 include-all [group-names];
 include-any [group-names];
 }
 auto-bandwidth {
 adjust-interval seconds;
 adjust-threshold percentage;
 maximum-bandwidth bps;
 minimum-bandwidth bps;
 monitor-bandwidth;
 }
 bandwidth bps {
 ct0 bps;
 ct1 bps;
 ct2 bps;
 ct3 bps;
 }
 class-of-service cos-value;
 description text;
 entropy-label;
 fast-reroute {
 (bandwidth bps | bandwidth-percent percentage);
 (exclude [group-names] | no-exclude);
 hop-limit number;
 (include-all [group-names] | no-include-all);
 (include-any [group-names] | no-include-any);
 }
 from address;
 install {
 destination-prefix/prefix-length <active>;
 }
 inter-domain;
 ldp-tunneling;
 link-protection;
 lsp-attributes {
 encoding-type (ethernet | packet | pdh | sonet-sdh);
 gpipid (ethernet | hdlc | ipv4 | pos-scrambling-crc-16 | pos-no-scrambling-crc-16 |
 pos-scrambling-crc-32 | pos-no-scrambling-crc-32 | ppp);
 signal-bandwidth type;
 switching-type (fiber | lambda | psc-1 | tdm);
 }
 metric metric;
 no-cspf;
 no-decrement-ttl;
 node-link-protection;
 optimize-timer seconds;
 p2mp lsp-name;

```

```
policing {
 filter filter-name;
 no-auto-policing;
}
preference preference;
primary path-name {
 adaptive;
 admin-group {
 exclude [group-names];
 include-all [group-names];
 include-any [group-names];
 }
 bandwidth bps {
 ct0 bps;
 ct1 bps;
 ct2 bps;
 ct3 bps;
 }
 class-of-service cos-value;
 hop-limit number;
 no-cspf;
 no-decrement-ttl;
 optimize-timer seconds;
 preference preference;
 priority setup-priority reservation-priority;
 (record | no-record);
 select (manual | unconditional);
 standby;
}
priority setup-priority reservation-priority;
(random | least-fill | most-fill);
(record | no-record);
retry-limit number;
retry-timer seconds;
revert-timer seconds;
secondary path-name {
 adaptive;
 admin-group {
 exclude [group-names];
 include-all [group-names];
 include-any [group-names];
 }
 bandwidth bps {
 ct0 bps;
 ct1 bps;
 ct2 bps;
 ct3 bps;
 }
 class-of-service cos-value;
 hop-limit number;
 no-cspf;
 no-decrement-ttl;
 optimize-timer seconds;
 preference preference;
 priority setup-priority reservation-priority;
 (record | no-record);
```

```

 select (manual | unconditional);
 standby;
 }
 soft-preemption;
 standby;
 to address;
 traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
 }
}

```

|                                 |                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls],<br>[edit protocols mpls]                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Configure an LSP to use in dynamic MPLS. When configuring an LSP, you must specify the address of the egress router in the <b>to</b> statement. All remaining statements are optional.                                                                                                                                         |
| <b>Options</b>                  | <p><b>lsp-name</b>—Name that identifies the LSP. The name can be up to 64 characters and can contain letters, digits, periods, and hyphens. To include other characters, enclose the name in quotation marks. The name must be unique within the ingress router.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Minimum MPLS Configuration</i></li> <li>• <i>Configuring the Ingress and Egress Router Addresses for LSPs</i></li> <li>• <i>Configuring Primary and Secondary LSPs</i></li> </ul>                                                                                                  |

## ldp-tunneling


|                                 |                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | ldp-tunneling;                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> ],<br>[edit protocols mpls label-switched-path <i>lsp-name</i> ] |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>                                 |
| <b>Description</b>              | Enable the LSP to be used for LDP tunneling.                                                                                                                         |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Enabling LDP over RSVP-Established LSPs on page 4834</a></li> </ul>                                             |

## link-protection (Static LSPs)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | link-protection bypass-name <i>name</i> ;                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i> ingress],<br>[edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i> transit <i>incoming-label</i> ],<br>[edit protocols mpls static-label-switched-path <i>lsp-name</i> ingress],<br>[edit protocols mpls static-label-switched-path <i>lsp-name</i> transit <i>incoming-label</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.1.                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Enable link protection on the specified static LSP. Link protection helps to ensure that traffic sent over a specific interface to a neighboring router can continue to reach the router if that interface fails.                                                                                                                                                                                                                                   |
| <b>Default</b>                  | Link protection is disabled.                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | bypass-name <i>name</i> —Bypass LSP name.                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Static LSPs</i></li><li>• <i>Example: Configuring Point-to-Multipoint LSPs with Static Routes</i></li></ul>                                                                                                                                                                                                                                                                                  |

## log-updown (Protocols MPLS)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>log-updown {   no-trap {     mpls-lsp-traps;     rfc3812-traps;   }   (syslog   no-syslog);   trap;   trap-path-down;   trap-path-up; }</pre>                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> protocols mpls],<br>[edit protocols mpls]                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>The <b>mpls-lsp-traps</b> and <b>rfc-3812-traps</b> options added in Junos OS Release 9.0.</p> <p>Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.</p>                                                                 |
| <b>Description</b>         | <p>Log a message or send an SNMP trap whenever an LSP makes a transition from up to down, or vice versa, and whenever an LSP switches from one active path to another. Only the ingress router performs these operations.</p>                                                                                                                                                                                                         |
|                            | <p> <b>NOTE:</b> System log messages for LSPs are generated by default. To disable the default logging of messages for LSPs, configure the <b>no-syslog</b> option under the <b>log-updown</b> statement.</p>                                                                                                                                      |
| <b>Default</b>             | <p>There is no default behavior for this statement. If you do not specify the options, the configuration cannot be committed.</p>                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>             | <p><b>no-syslog</b>—Do not log a message to the system log file.</p> <p><b>no-trap</b>—Do not send an SNMP trap.</p> <p><b>syslog</b>—Log a message to the system log file.</p> <p><b>trap</b>—Send an SNMP trap.</p> <p><b>trap-path-down</b>—Send an SNMP trap when an LSP path goes down.</p> <p><b>trap-path-up</b>—Send an SNMP trap when an LSP path comes up.</p> <p>The <b>no-trap</b> statement is explained separately.</p> |

|                                 |                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring System Log Messages and SNMP Traps for LSPs</i></li><li>• <i>Network Management Administration Guide for Routing Devices</i></li><li>• <a href="#">no-trap on page 5065</a></li><li>• <a href="#">traceoptions (Protocols MPLS) on page 5102</a></li></ul> |

---

## **lsp-attributes**

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>lsp-attributes {<br/>    encoding-type (ethernet   packet   pdh   sonet-sdh);<br/>    <a href="#">gpip</a> (ethernet   hdlc   ipv4   pos-scrambling-crc-16   pos-no-scrambling-crc-16  <br/>        pos-scrambling-crc-32   pos-no-scrambling-crc-32   ppp);<br/>    <a href="#">signal-bandwidth</a> type;<br/>    <a href="#">switching-type</a> (fiber   lambda   psc-1   tdm);<br/>}</pre> |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> ],<br>[edit protocols mpls label-switched-path <i>lsp-name</i> ]                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br><b>pos-scrambling-crc-16</b> , <b>pos-no-scrambling-crc-16</b> , <b>pos-scrambling-crc-32</b> , and <b>pos-no-scrambling-crc-32</b> options added in Junos OS Release 8.0.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                             |
| <b>Description</b>              | Define the parameters signaled during LSP setup. These usually determine the nature of the resource (label) allocated for the LSP.<br><br>The remaining statements are explained separately.                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring MPLS LSPs for GMPLS</i></li></ul>                                                                                                                                                                                                                                                                                                            |



## maximum-bandwidth (Protocols MPLS)

|                                 |                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>maximum-bandwidth <i>bps</i>;</code>                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth],<br>[edit protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.                                                                    |
| <b>Description</b>              | Specify the maximum amount of bandwidth in bits per second (bps).                                                                                                                                |
| <b>Options</b>                  | <i>bps</i> —Maximum amount of bandwidth.                                                                                                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Automatic Bandwidth Allocation for LSPs</i></li> </ul>                                                                                   |

## metric (Protocols MPLS)

|                                 |                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>metric <i>metric</i>;</code>                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i> ingress],<br>[edit protocols mpls label-switched-path <i>lsp-name</i> ],<br>[edit protocols mpls static-label-switched-path <i>lsp-name</i> ingress] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.                                                                                                                      |
| <b>Description</b>              | Compare against another LSP or against an IGP route. To disable dynamic metric tracking, assign a fixed metric value to an LSP. If no metric is assigned, the LSP metric is dynamic and automatically tracks underlying IGP metrics.                                                                                                                                      |
| <b>Options</b>                  | <i>metric</i> —LSP metric value.<br><b>Default:</b> No metric assigned (dynamic)<br><b>Range:</b> 1 through 16,777,215                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring LSP Metrics</i></li> </ul>                                                                                                                                                                                                                                                                                        |

## minimum-bandwidth

---

|                                 |                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>minimum-bandwidth <i>bps</i>;</code>                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth],<br>[edit protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.                                                                    |
| <b>Description</b>              | Set the minimum bandwidth in bps for an LSP with automatic bandwidth allocation enabled.                                                                                                         |
| <b>Options</b>                  | <i>bps</i> —Minimum bandwidth for the LSP.                                                                                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Automatic Bandwidth Allocation for LSPs</i></li></ul>                                                                                     |

## monitor-bandwidth

---

|                                 |                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>monitor-bandwidth;</code>                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth],<br>[edit protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.                                                                    |
| <b>Description</b>              | Do not automatically adjust bandwidth allocation. However, the maximum average bandwidth utilization is monitored on the LSP, and the information is recorded in the MPLS statistics file.       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Automatic Bandwidth Allocation for LSPs</i></li></ul>                                                                                     |

## mtu-signaling

---

|                                 |                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | mtu-signaling;                                                                                                            |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls path-mtu rsvp],<br>[edit protocols mpls path-mtu rsvp]    |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series. |
| <b>Description</b>              | Enable MTU signaling in RSVP.                                                                                             |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring MTU Signaling in RSVP</i></li></ul>                                |

## no-cspf

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-cspf;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls],<br>[edit logical-systems <i>logical-system-name</i> protocols mpls <b>label-switched-path</b> <i>lsp-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> protocols mpls <b>label-switched-path</b> <i>lsp-name</i><br>( <b>primary</b>   <b>secondary</b> ) <i>path-name</i> ],<br>[edit protocols mpls],<br>[edit protocols mpls <b>label-switched-path</b> <i>lsp-name</i> ],<br>[edit protocols mpls <b>label-switched-path</b> <i>lsp-name</i> ( <b>primary</b>   <b>secondary</b> ) <i>path-name</i> ]                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | <p>Disable constrained-path LSP computation.</p> <p>An explicit-path LSP is completely configured through operator action. Once configured, it is initiated only along the explicitly specified path.</p> <p>A constrained-path LSP relies on an ingress router to compute the complete path. The ingress router takes into account the following information during the computation:</p> <ul style="list-style-type: none"><li>• Interior gateway protocol (IGP) topology database</li><li>• Link utilization information from extensions in the IGP link-state database</li><li>• Administrative group information from extensions in the IGP link-state database</li><li>• LSP requirements, including bandwidth, hop count, and administrative group</li></ul> <p>Constrained-path LSPs can generally avoid link failures and congested links. They also permit recomputation (therefore, a new path) during topology changes or unsuccessful setup.</p> |
| <b>Default</b>                  | Constrained-path LSP computation enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Disabling Constrained-Path LSP Computation</i></li><li>• <i>Configuring Explicit-Path LSPs</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

## no-decrement-ttl

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-decrement-ttl;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls],<br>[edit logical-systems <i>logical-system-name</i> protocols mpls <a href="#">label-switched-path</a> <i>lsp-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> protocols mpls <a href="#">label-switched-path</a> <i>lsp-name</i><br>( <a href="#">primary</a>   <a href="#">secondary</a> ) <i>path-name</i> ],<br>[edit protocols mpls],<br>[edit protocols mpls <a href="#">label-switched-path</a> <i>lsp-name</i> ],<br>[edit protocols mpls <a href="#">label-switched-path</a> <i>lsp-name</i> ( <a href="#">primary</a>   <a href="#">secondary</a> ) <i>path-name</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Disable normal time-to-live (TTL) decrementing, which decrements the TTL field in the IP header by 1. This statement decrements the IP TTL by 1 before encapsulating the IP packet within an MPLS packet. When the penultimate router pops off the top label, it does not use the standard write-back procedure of writing the MPLS TTL into the IP TTL field. Therefore, the IP packet is decremented by 1. The ultimate router then decrements the packet by one more for a total cloud appearance of 2, thus hiding the network topology.                                                                                                                      |
| <b>Default</b>                  | Normal TTL decrementing enabled; the TTL field value is decremented by 1 as the packet passes through each label-switched router in the LSP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Disabling Normal TTL Decrementing</a></li> <li>• <a href="#">no-propagate-ttl on page 5063</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## no-install-to-address

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-install-to-address;                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls <a href="#">label-switched-path</a> <i>lsp-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> protocols mpls<br>static-label-switched-path <i>lsp-name</i> ingress],<br>[edit protocols mpls <a href="#">label-switched-path</a> <i>lsp-name</i> ],<br>[edit protocols mpls static-label-switched-path <i>lsp-name</i> ingress] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | Prevent the egress router address configured using the <b>to</b> statement from being installed into the inet.3 and inet.0 routing tables.                                                                                                                                                                                                                                                                   |
| <b>Default</b>                  | The egress router address for an LSP is installed into the inet.3 and inet.0 routing tables.                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Preventing the Addition of Egress Router Addresses to Routing Tables</i></li><li>• <a href="#">to on page 5101</a></li></ul>                                                                                                                                                                                                                                      |

## no-propagate-ttl

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-propagate-ttl;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls],<br>[edit protocols mpls]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | <p>Disable normal time-to-live (TTL) decrementing. You configure this statement once per router, and it affects all RSVP-signaled or LDP-signaled LSPs. When this router acts as an ingress router for an LSP, it pushes an MPLS header with a TTL value of 255, regardless of the IP packet TTL. When the router acts as the penultimate router, it pops the MPLS header without writing the MPLS TTL into the IP packet.</p> <p>When you add the <b>no-propagate-ttl</b> statement to the configuration or delete it from the configuration, the effect takes place immediately. There is no need to clear existing RSVP LSPs or LDP sessions.</p> |
| <b>Default</b>                  | Normal TTL decrementing enabled; the TTL field value is decremented by 1 as the packet passes through each label-switched router in the LSP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Disabling Normal TTL Decrementing</i></li> <li>• <i>Example: Disabling Normal TTL Decrementing in a VRF Routing Instance</i> (on <i>Layer 3 VPNs Feature Guide for Routing Devices</i> or in the <i>Junos VPNs Configuration Guide</i>)</li> <li>• <a href="#">no-decrement-ttl on page 5061</a></li> </ul>                                                                                                                                                                                                                                                                                              |

## record

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | (record   no-record);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls],<br>[edit logical-systems <i>logical-system-name</i> protocols mpls <b>label-switched-path</b> <i>lsp-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> protocols mpls <b>label-switched-path</b> ( <b>primary</b>   <b>secondary</b> ) <i>path-name</i> ],<br>[edit protocols mpls],<br>[edit protocols mpls <b>label-switched-path</b> <i>lsp-name</i> ],<br>[edit protocols mpls <b>label-switched-path</b> <i>lsp-name</i> ( <b>primary</b>   <b>secondary</b> ) <i>path-name</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Specify whether an LSP should actively record the routes in the path. Recording routes requires that all transit routers support the RSVP Record Route object. Recording routes can be useful for diagnostics and loop detection.                                                                                                                                                                                                                                                                                                                                      |
| <b>Default</b>                  | Record routes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Disabling Path Route Recording by LSPs</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |



## no-trap

---

|                                 |                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-trap {<br>mpls-lsp-traps;<br>rfc-3812-traps;<br>}                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls log-updown],<br>[edit protocols mpls log-updown]                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>The <b>mpls-lsp-traps</b> and <b>rfc-3812-traps</b> options added in Junos OS Release 9.0.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                |
| <b>Description</b>              | Prevent the transmission of SNMP traps.                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <p><b>mpls-lsp-traps</b>—Block the MPLS LSP traps defined in the <b>rfc-3812-traps</b>, but allows the <b>rfc3812.mib</b> traps.</p> <p><b>rfc-3812-traps</b>—Block the traps defined in the <b>rfc3812.mib</b>, but allows the MPLS LSP traps defined in the <b>jnx-mpls.mib</b>.</p> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring System Log Messages and SNMP Traps for LSPs</i></li> <li>• <i>Network Management Administration Guide for Routing Devices</i></li> <li>• <a href="#">traceoptions (Protocols MPLS) on page 5102</a></li> </ul>                 |

## node-link-protection (Protocols MPLS)

---

|                                 |                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | node-link-protection;                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> ],<br>[edit protocols mpls label-switched-path <i>lsp-name</i> ]                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series and for EX4600 switches.                                                                                                                                                     |
| <b>Description</b>              | Enable node and link protection on the specified LSP. To fully enable node and link protection, you also need to include the <b>link-protection</b> statement at the [edit protocols <b>rsvp</b> interface <i>interface-name</i> ] hierarchy level.                                                       |
| <b>Default</b>                  | Node and link protection is disabled.                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Node Protection or Link Protection for LSPs</a></li><li>• <a href="#">MPLS Feature Support on QFX Series and EX4600 Switches on page 4934</a></li><li>• <a href="#">Interprovider and Carrier-of-Carriers VPNs on page 4966</a></li></ul> |

## oam (Protocols MPLS)

```
Syntax oam {
 bfd-liveness-detection{
 failure-action teardown;
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 minimum-transmit-interval milliseconds;
 multiplier detection-time-multiplier;
 }
 lsp-ping-interval seconds;
 mpls-tp-mode;
 performance-monitoring {
 querier {
 loss {
 traffic-class tc-value {
 query-interval milliseconds;
 measurement-quantity bytes|packets;
 average-sample-size sample size;
 loss-threshold loss threshold value;
 loss-threshold-window number of samples for loss threshold;
 }
 }
 delay {
 traffic-class tc-value {
 query-interval milliseconds;
 padding-size size;
 average-sample-size sample size;
 rtt-delay-threshold rtt threshold value;
 twcd-delay-threshold twcd threshold value;
 }
 }
 loss-delay {
 traffic-class tc-value {
 query-interval milliseconds;
 measurement-quantity bytes|packets;
 padding-size size;
 average-sample-size sample size;
 loss-threshold loss threshold value;
 loss-threshold-window number of samples for loss threshold;
 rtt-delay-threshold rtt threshold value;
 twcd-delay-threshold twcd threshold value;
 }
 }
 }
 }
 responder {
 loss {
 min-query-interval milliseconds;
 }
 delay {
 min-query-interval milliseconds;
 }
 }
}
```

}

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Hierarchy Level</b>          | [edit protocols mpls],<br>[edit protocols mpls <b>label-switched-path</b> <i>lsp-name</i> ]<br>[edit protocols mpls <b>label-switched-path</b> <i>lsp-name</i> primary <i>path-name</i> ]                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.6.<br><b>lsp-ping-interval</b> option introduced in Junos OS Release 9.4.<br>Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.<br><b>performance-monitoring</b> configuration statement introduced in Junos OS Release 15.1. |
| <b>Description</b>              | Enable Operation, Administration, and Maintenance (OAM) for RSVP-signaled LSPs.                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <b>lsp-ping-interval seconds</b> —Specify the duration of the LSP ping interval in seconds. To issue a ping on an RSVP-signaled LSP, use the <b>ping mpls rsvp</b> command.<br><br>The remaining statements are explained separately.                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring BFD for MPLS IPv4 LSPs</i></li></ul>                                                                                                                                                                                                                                                                                                                           |

---

## optimize-aggressive

---

|                                 |                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | optimize-aggressive;                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls],<br>[edit protocols mpls]                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                  |
| <b>Description</b>              | If enabled, the LSP reoptimization is based solely on the IGP metric. The reoptimization process ignores the available bandwidth ratio calculations, the least-fill 10 percent congestion improvement rule, and the hop-counts rule. This statement makes reoptimization more aggressive than the default. |
| <b>Default</b>                  | Aggressive optimization is disabled.                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Optimizing Signaled LSPs</i></li></ul>                                                                                                                                                                                                                          |

## optimize-hold-dead-delay

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>optimized-hold-dead-delay <i>seconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls],<br>[edit logical-systems <i>logical-system-name</i> protocols mpls label-switch-path <i>lsp-name</i> ],<br>[edit protocols mpls],<br>[edit protocols mpls label-switch-path <i>lsp-name</i> ]                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Allows you to specify the amount of time to delay the tear down of old paths after the router has switched traffic to new optimized paths. You only need to configure this statement on routers acting as the ingress for the affected LSPs (you do not need to configure this statement on transit or egress routers). The specified delay helps to ensure that old paths are not torn down before all routes have been switched over to the new optimized paths. This delay timer starts when the timer specified by the <b>optimize-switchover-dealy</b> statement has elapsed. |
| <b>Options</b>                  | <b><i>seconds</i></b> —Configure the time in seconds to wait before tearing down the old paths that were in use prior to the last LSP optimization.<br><b>Default:</b> 60 seconds<br><b>Range:</b> 0 through 65,535 seconds                                                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Optimizing Signaled LSPs</i></li> <li>• <a href="#">optimize-switchover-delay on page 5070</a></li> <li>• <a href="#">optimize-timer on page 5071</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                       |

## optimize-switchover-delay

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>optimize-switchover-delay seconds;</code>                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls],<br>[edit protocols mpls]                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1R1.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Delays the switch over of LSPs to newly optimized paths. You only need to configure this statement on routers acting as the ingress for the affected LSPs (you do not need to configure this statement on transit or egress routers). The specified delay helps to ensure that the new optimized paths have been established before traffic is switched over from the old paths. |
| <b>Options</b>                  | <b>seconds</b> —Configure the time in seconds to wait before switching LSPs to newly optimized paths.<br><b>Default:</b> 1 second<br><b>Range:</b> 1 through 900 seconds                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Optimizing Signaled LSPs</i></li><li>• <a href="#">optimize-hold-dead-delay on page 5069</a></li><li>• <a href="#">optimize-timer on page 5071</a></li></ul>                                                                                                                                                                          |

## optimize-timer (Protocols MPLS)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>optimize-timer <i>seconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols mpls],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols mpls <i>label-switched-path lsp-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols mpls <i>label-switched-path lsp-name</i></code><br><code>  (<i>primary</i>   <i>secondary</i>) <i>path-name</i>],</code><br><code>[edit protocols mpls],</code><br><code>[edit protocols mpls <i>label-switched-path lsp-name</i>],</code><br><code>[edit protocols mpls <i>label-switched-path lsp-name</i> (<i>primary</i>   <i>secondary</i>) <i>path-name</i>]</code>                                                                                                                                                                                                            |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | <p>Enable periodic reoptimization of an LSP that is already set up. If topology changes occur, an existing path might become suboptimal, and a subsequent recomputation might be able to determine a better path. This feature is useful only on LSPs for which constrained-path computation is enabled; that is, for which the <b>no-cspf</b> statement is not configured. Also, you only need to configure this statement on routers acting as the ingress for the affected LSPs (you do not need to configure this statement on transit or egress routers).</p> <p>To avoid extensive resource consumption that might result because of frequent path recomputations, or to avoid destabilizing the network as a result of constantly changing LSPs, we recommend that you either leave the timer value sufficiently large or disable the timer value.</p> |
| <b>Default</b>                  | The optimize timer is disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <p><b><i>seconds</i></b>—Length of the optimize timer, in seconds.</p> <p><b>Range:</b> 0 through 65,535 seconds</p> <p><b>Default:</b> 0 seconds (the optimize timer is disabled)</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Optimizing Signaled LSPs</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

## p2mp (Protocols MPLS)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>p2mp <i>p2mp-lsp-name</i>;</code>                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> ],<br>[edit protocols mpls label-switched-path <i>lsp-name</i> ]                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.                                                                                                                                         |
| <b>Description</b>              | Specify an LSP as either a point-to-multipoint LSP or as a branch LSP of a point-to-multipoint LSP by specifying the point-to-multipoint LSP path name.                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <b><i>p2mp-lsp-name</i></b> —Name of the point-to-multipoint LSP path that identifies the sequence of nodes that form the point-to-multipoint LSP. The name can contain up to 32 characters and can include letters, digits, periods, and hyphens. To include other characters or use a longer name, enclose the name in quotation marks. The name must be unique within the ingress router. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Primary and Branch LSPs for Point-to-Multipoint LSPs</i></li></ul>                                                                                                                                                                                                                                                                    |



## path (Protocols MPLS)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>path <i>path-name</i> {     (<i>address</i>   <i>hostname</i>) &lt;strict   loose&gt;; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls],<br>[edit protocols mpls]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | <p>Create a named path and optionally specify the sequence of explicit routers that form the path.</p> <p>You must include this statement when configuring explicit LSPs.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><b>address</b>—IP address of each transit router in the LSP. You must specify the address or hostname of each transit router, although you do not need to list each transit router if its type is <b>loose</b>. As an option, you can include the ingress and egress routers in the path. Specify the addresses in order, starting with the ingress router (optional) or the first transit router, and continuing sequentially along the path until reaching the egress router (optional) or the router immediately before the egress router.</p> <p><b>Default:</b> If you do not specify any routers explicitly, no routing limitations are imposed on the LSP.</p> <p><b>hostname</b>—See <b>address</b>.</p> <p><b>Default:</b> If you do not specify any routers explicitly, no routing limitations are imposed on the LSP.</p> <p><b>loose</b>—(Optional) Indicate that the next address in the <b>path</b> statement is a loose link. This means that the LSP can traverse through other routers before reaching this router.</p> <p><b>Default:</b> <b>strict</b></p> <p><b>path-name</b>—Name that identifies the sequence of nodes that form an LSP. The name can contain up to 32 characters and can include letters, digits, periods, and hyphens. To include other characters or use a longer name, enclose the name in quotation marks. The name must be unique within the ingress router.</p> <p><b>strict</b>—(Optional) Indicate that the LSP must go to the next address specified in the <b>path</b> statement without traversing other nodes. This is the default.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Creating Named Paths</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## path-mtu

---

|                                 |                                                                                                                                                   |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>path-mtu {<br/>    allow-fragmentation;<br/>    rsvp {<br/>        mtu-signaling;<br/>    }<br/>}</pre>                                      |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls],<br>[edit protocols mpls]                                                        |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                         |
| <b>Description</b>              | Configure MTU options for MPLS paths, including packet fragmentation and MTU signaling.<br><br>The remaining statements are explained separately. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring MTU Signaling in RSVP</i></li></ul>                                                        |

## policing (Protocols MPLS)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>policing {<br/>    filter <i>filter-name</i>;<br/>    no-auto-policing;<br/>}</pre>                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> protocols mpls<br>static-label-switched-path <i>lsp-name</i> ingress],<br>[edit protocols mpls label-switched-path <i>lsp-name</i> ],<br>[edit protocols mpls static-label-switched-path <i>lsp-name</i> ingress] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Specify the policing filter for the LSP.                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <b>filter <i>filter-name</i></b> —Specify the name of the policing filter.<br><br><b>no-auto-policing</b> —Disable automatic policing on this LSP.                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring MPLS Firewall Filters and Policers</i></li><li>• <i>auto-policing</i></li></ul>                                                                                                                                                                                                                                       |

## policy-statement

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>policy-statement <i>policy-name</i> {<br/>    term <i>term-name</i> {<br/>        from {<br/>            family <i>family-name</i>;<br/>            match-conditions;<br/>            policy <i>subroutine-policy-name</i>;<br/>            prefix-list <i>prefix-list-name</i>;<br/>            prefix-list-filter <i>prefix-list-name</i> <i>match-type</i> &lt;<i>actions</i>&gt;;<br/>            protocol <i>protocol-name</i>;<br/>            route-filter <i>destination-prefix</i> <i>match-type</i> &lt;<i>actions</i>&gt;;<br/>            source-address-filter <i>source-prefix</i> <i>match-type</i> &lt;<i>actions</i>&gt;;<br/>        }<br/>        to {<br/>            match-conditions;<br/>            policy <i>subroutine-policy-name</i>;<br/>        }<br/>        then <i>actions</i>;<br/>    }<br/>    then {<br/>        no-entropy-label-capability;<br/>    }<br/>}</pre> |
| <b>Hierarchy Level</b>     | [edit dynamic policy-options],<br>[edit logical-systems <i>logical-system-name</i> policy-options],<br>[edit policy-options]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Support for configuration in the dynamic database introduced in Junos OS Release 9.5.<br>Support for configuration in the dynamic database introduced in Junos OS Release 9.5 for EX Series switches.<br><b>inet-mdt</b> option introduced in Junos OS Release 10.0R2.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br><b>route-target</b> option introduced in Junos OS Release 12.2.<br>Statement introduced in Junos OS 14.1X53-D20 for the OCX Series.<br><b>protocol</b> and <b>traffic-engineering</b> options introduced in Junos OS Release 14.2.<br><b>no-entropy-label-capability</b> option introduced in Junos OS Release 15.1.                                                                                                              |
| <b>Description</b>         | <p>Define a routing policy, including subroutine policies.</p> <p>A <i>term</i> is a named structure in which match conditions and actions are defined. Routing policies are made up of one or more terms. Each routing policy term is identified by a term name. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose the entire name in double quotation marks.</p> <p>Each term contains a set of match conditions and a set of actions:</p>                                                                                                                                                                                                                                                                                                                                                                                |

- Match conditions are criteria that a route must match before the actions can be applied. If a route matches all criteria, one or more actions are applied to the route.
- Actions specify whether to accept or reject the route, control how a series of policies are evaluated, and manipulate the characteristics associated with a route.

Generally, a router compares a route against the match conditions of each term in a routing policy, starting with the first and moving through the terms in the order in which they are defined, until a match is made and an explicitly configured or default action of **accept** or **reject** is taken. If none of the terms in the policy match the route, the router compares the route against the next policy, and so on, until either an action is taken or the default policy is evaluated.

If none of the match conditions of each term evaluates to true, the final action is executed. The final action is defined in an unnamed term. Additionally, you can define a default action (either **accept** or **reject**) that overrides any action intrinsic to the protocol.

The order of match conditions in a term is not relevant, because a route must match all match conditions in a term for an action to be taken.

To list the routing policies under the **[edit policy-options]** hierarchy level by **policy-statement *policy-name*** in alphabetical order, enter the **show policy-options** configuration command.

**Options** **actions**—(Optional) One or more actions to take if the conditions match. The actions are described in *Configuring Flow Control Actions*.

**family** **family-name**—(Optional) Specify an address family protocol. Specify **inet** for IPv4. Specify **inet6** for 128-bit IPv6, and to enable interpretation of IPv6 router filter addresses. For IS-IS traffic, specify **iso**. For IPv4 multicast VPN traffic, specify **inet-mvpn**. For IPv6 multicast VPN traffic, specify **inet6-mvpn**. For multicast-distribution-tree (MDT) IPv4 traffic, specify **inet-mdt**. For BGP route target VPN traffic, specify **route-target**. For traffic engineering, specify **traffic-engineering**.



**NOTE:** When family is not specified, the routing device or routing instance uses the address family or families carried by BGP. If multiprotocol BGP (MP-BGP) is enabled, the policy defaults to the protocol family or families carried in the network layer reachability information (NLRI) as configured in the **family** statement for BGP. If MP-BGP is not enabled, the policy uses the default BGP address family unicast IPv4.

**from**—(Optional) Match a route based on its source address.

**match-conditions**—(Optional in **from** statement; required in **to** statement) One or more conditions to use to make a match. The qualifiers are described in *Routing Policy Match Conditions*.

**policy subroutine-policy-name**—Use another policy as a match condition within this policy. The name identifying the subroutine policy can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" "). Policy names cannot take the form **\_\_.\*-internal\_\_**, as this form is reserved. For information about how to configure subroutines, see *Understanding Policy Subroutines in Routing Policy Match Conditions*.

**no-entropy-label-capability**—(Optional) Disable the entropy label capability advertisement at egress or transit routes specified in the policy.

**policy subroutine-policy-name**—Use another policy as a match condition within this policy. The name identifying the subroutine policy can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" "). Policy names cannot take the form **\_\_.\*-internal\_\_**, as this form is reserved. For information about how to configure subroutines, see *Understanding Policy Subroutines in Routing Policy Match Conditions*.

**policy-name**—Name that identifies the policy. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" ").

**prefix-list prefix-list-name**—Name of a list of IPv4 or IPv6 prefixes.

**prefix-list-filter prefix-list-name**—Name of a prefix list to evaluate using qualifiers; **match-type** is the type of match (see *Configuring Prefix List Filters*), and **actions** is the action to take if the prefixes match.

**protocol** *protocol-name*—Name of the protocol used to control traffic engineering database import at the originating point.

**route-filter** *destination-prefix match-type <actions>*—(Optional) List of routes on which to perform an immediate match; *destination-prefix* is the IPv4 or IPv6 route prefix to match, *match-type* is the type of match (see *Configuring Route Lists*), and *actions* is the action to take if the *destination-prefix* matches.

**source-address-filter** *source-prefix match-type <actions>*—(Optional) Unicast source addresses in multiprotocol BGP (MBGP) and Multicast Source Discovery Protocol (MSDP) environments on which to perform an immediate match. *source-prefix* is the IPv4 or IPv6 route prefix to match, *match-type* is the type of match (see *Configuring Route Lists*), and *actions* is the action to take if the *source-prefix* matches.

**term** *term-name*—Name that identifies the term. The term name must be unique in the policy. It can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. To include spaces in the name, enclose the entire name in quotation marks (" "). A policy statement can include multiple terms. We recommend that you name all terms. However, you do have the option to include an unnamed term which must be the final term in the policy. To configure an unnamed term, omit the **term** statement when defining match conditions and actions.

**to**—(Optional) Match a route based on its destination address or the protocols into which the route is being advertised.

**then**—(Optional) Actions to take on matching routes. The actions are described in *Configuring Flow Control Actions* and *Configuring Actions That Manipulate Route Characteristics*.

|                                 |                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration. |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>dynamic-db</i></li> </ul>                                                 |

## pop

---

|                                 |                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | pop;                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i> transit <i>incoming-label</i> ],<br>[edit protocols mpls static-label-switched-path <i>lsp-name</i> transit <i>incoming-label</i> ]     |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                          |
| <b>Description</b>              | Remove the label from the top of the label stack. If there is another label in the stack, that label becomes the label at the top of the label stack. Otherwise, the packet is forwarded as a native protocol packet (typically, as an IP packet). |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the Intermediate and Egress Routers for Static LSPs</i></li><li>• <a href="#">swap on page 5095</a></li></ul>                                                                               |



## preference (Protocols MPLS)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>preference <i>preference</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols mpls],<br/> [edit logical-systems <i>logical-system-name</i> protocols mpls <b>label-switched-path</b> <i>lsp-name</i>],<br/> [edit logical-systems <i>logical-system-name</i> protocols mpls <b>label-switched-path</b> <i>lsp-name</i> (<b>primary</b>   <b>secondary</b>) <i>path-name</i>],<br/> [edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i> ingress],<br/> [edit protocols mpls],<br/> [edit protocols mpls <b>label-switched-path</b> <i>lsp-name</i>],<br/> [edit protocols mpls <b>label-switched-path</b> <i>lsp-name</i> (<b>primary</b>   <b>secondary</b>) <i>path-name</i>],<br/> [edit protocols mpls static-label-switched-path <i>lsp-name</i> ingress]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.<br/> Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | <p>Preference for the route.</p> <p>You can optionally configure multiple LSPs between the same pair of ingress and egress routers. This is useful for balancing the load among the LSPs because all LSPs, by default, have the same preference level. To prefer one LSP over another, set different preference levels for individual LSPs. The LSP with the lowest preference value is used. The default preference for LSPs is lower (more preferred) than all learned routes except direct interface routes.</p>                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <p><b><i>preference</i></b>—Preference to assign to the route. A route with a lower preference value is preferred.</p> <p><b>Range:</b> 1 through 255</p> <p><b>Default:</b> 5 for static MPLS LSPs, 7 for RSVP MPLS LSPs, 9 for LDP MPLS LSPs</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.<br/> routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Preference Values for LSPs</i></li> <li>• <i>Configuring Static LSPs</i></li> <li>• <i>Configuring Static LSPs</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## primary (Protocols MPLS)

---

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <pre>primary <i>path-name</i> {<br/>    adaptive;<br/>    admin-group {<br/>        exclude [ <i>group-names</i> ];<br/>        include-all [ <i>group-names</i> ];<br/>        include-any [ <i>group-names</i> ];<br/>    }<br/>    bandwidth <i>bps</i>;<br/>    class-of-service <i>cos-value</i>;<br/>    hop-limit <i>number</i>;<br/>    no-cspf;<br/>    no-decrement-ttl;<br/>    optimize-timer <i>seconds</i>;<br/>    preference <i>preference</i>;<br/>    priority <i>setup-priority reservation-priority</i>;<br/>    (record   no-record);<br/>    select (manual   unconditional);<br/>    standby;<br/>}</pre> |
| Hierarchy Level          | [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> ],<br>[edit protocols mpls label-switched-path <i>lsp-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Release Information      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Description              | <p>Specify the primary path to use for an LSP. You can configure only one primary path.</p> <p>You can optionally specify preference, CoS, and bandwidth values for the primary path, which override any equivalent values that you configure for the LSP (at the [edit mpls label-switched-path <i>lsp-name</i>] hierarchy level).</p>                                                                                                                                                                                                                                                                                          |
| Options                  | <p><b>path-name</b>—Name of a path that you created with the <b>path</b> statement.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Required Privilege Level | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Related Documentation    | <ul style="list-style-type: none"><li>Configuring Primary and Secondary LSPs</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## push

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>push out-label;</code>                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i> bypass],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i> ingress],</p> <p>[edit protocols mpls static-label-switched-path <i>lsp-name</i> bypass],</p> <p>[edit protocols mpls static-label-switched-path <i>lsp-name</i> ingress]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Add a new label to the top of the label stack. This statement is used to configure static LSPs at ingress routers and to configure bypass LSPs for static LSPs.                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <p><b>out-label</b>—Manually assigned outgoing label value.</p> <p><b>Range:</b> 0 through 1,048,575.</p>                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">pop on page 5080</a></li> <li>• <a href="#">swap on page 5095</a></li> <li>• <i>Configuring Static LSPs</i></li> </ul>                                                                                                                                                                                                                                            |

## record

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | (record   no-record);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls],<br>[edit logical-systems <i>logical-system-name</i> protocols mpls <b>label-switched-path</b> <i>lsp-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> protocols mpls <b>label-switched-path</b> (primary   secondary) <i>path-name</i> ],<br>[edit protocols mpls],<br>[edit protocols mpls <b>label-switched-path</b> <i>lsp-name</i> ],<br>[edit protocols mpls <b>label-switched-path</b> <i>lsp-name</i> (primary   secondary) <i>path-name</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Specify whether an LSP should actively record the routes in the path. Recording routes requires that all transit routers support the RSVP Record Route object. Recording routes can be useful for diagnostics and loop detection.                                                                                                                                                                                                                                                                                                      |
| <b>Default</b>                  | Record routes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>Disabling Path Route Recording by LSPs</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## retry-limit

---

|                                 |                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>retry-limit <i>number</i>;</code>                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> ],<br>[edit protocols mpls label-switched-path <i>lsp-name</i> ],                                                                                                              |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                          |
| <b>Description</b>              | Maximum number of times the ingress router tries to establish the primary path. This counter is reset each time a primary path is created successfully. When the limit is exceeded, no more connection attempts are made. Intervention is then required to restart the connection. |
| <b>Options</b>                  | <b><i>number</i></b> —Maximum number of tries to establish the primary path.<br><b>Range:</b> 0 through 10,000<br><b>Default:</b> 0 (The ingress node never stops trying to establish the primary path.)                                                                           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the Connection Between Ingress and Egress Routers</i></li></ul>                                                                                                                                                             |

## revert-timer

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>revert-timer seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls],<br>[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> ],<br>[edit protocols mpls],<br>[edit protocols mpls label-switched-path <i>lsp-name</i> ]                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>BFD behavior modified in Junos OS Release 9.0.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | <p>Specify the amount of time (in seconds) that an LSP must wait before traffic reverts to a primary path. If during this time the primary path experiences any connectivity problem or stability problem, the timer is restarted.</p> <p>If you have configured BFD on the LSP, the Junos OS waits until the BFD session is restored before starting the revert timer counter.</p> <p>If you have configured a value of 0 seconds for the <b>revert-timer</b> statement and traffic is switched to the secondary path, the traffic remains on that path indefinitely. It is never switched back to the primary path unless you intervene.</p> |
| <b>Options</b>                  | <p><b>seconds</b>—Time in seconds.</p> <p><b>Range:</b> 0 through 65,535 seconds</p> <p><b>Default:</b> 60 seconds</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Primary and Secondary LSPs</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

## rsvp-error-hold-time

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>rsvp-error-hold-time <i>seconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls],<br>[edit protocols mpls]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | <p>Amount of time MPLS retains RSVP PathErr messages and considers them for CSPF computations. The more time you configure, the more time a source node (ingress of an RSVP LSP) can have to learn about the failures of its LSP by monitoring PathErr messages transmitted from downstream nodes.</p> <p>Information from the PathErr messages is incorporated into subsequent LSP computations, which can improve the accuracy and speed of LSP setup. Some PathErr messages are also used to update traffic engineering database bandwidth information, reducing inconsistencies between the database and the network.</p> |
| <b>Options</b>                  | <p><b><i>seconds</i></b>—Amount of time MPLS retains RSVP PathErr messages and considers them for CSPF computations.</p> <p><b>Range:</b> 0 through 240 seconds</p> <p><b>Default:</b> 25 seconds</p>                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Improving Traffic Engineering Database Accuracy with RSVP PathErr Messages</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## secondary (Protocols MPLS)

---

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <pre>secondary <i>path-name</i> {<br/>    adaptive;<br/>    admin-group {<br/>        exclude [ <i>group-names</i> ];<br/>        include-all [ <i>group-names</i> ];<br/>        include-any [ <i>group-names</i> ];<br/>    }<br/>    bandwidth <i>bps</i>;<br/>    class-of-service <i>cos-value</i>;<br/>    hop-limit <i>number</i>;<br/>    no-cspf;<br/>    no-decrement-ttl;<br/>    optimize-timer <i>seconds</i>;<br/>    preference <i>preference</i>;<br/>    priority <i>setup-priority reservation-priority</i>;<br/>    (record   no-record);<br/>    retry-limit <i>number</i>;<br/>    retry-timer <i>seconds</i>;<br/>    select (manual   unconditional);<br/>    standby;<br/>}</pre> |
| Hierarchy Level          | [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> ],<br>[edit protocols mpls label-switched-path <i>lsp-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Release Information      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Description              | <p>Specify one or more secondary paths to use for the LSP. You can configure more than one secondary path. All secondary paths are equal, and the first one that is available is chosen.</p> <p>You can specify secondary paths even if you have not specified any primary paths.</p> <p>Optionally, you can specify preference, CoS, and bandwidth values for the secondary path, which override any equivalent values that you configure for the LSP (at the [edit mpls label-switched-path] hierarchy level).</p>                                                                                                                                                                                      |
| Options                  | <p><b><i>path-name</i></b>—Name of a path that you created with the <b>path</b> statement.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Required Privilege Level | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Related Documentation    | <ul style="list-style-type: none"><li>Configuring Primary and Secondary LSPs</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |



## select

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>select (manual   unconditional);</code>                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls <b>label-switched-path</b> <i>lsp-name</i> (primary   secondary) <i>path-name</i> ],<br>[edit protocols mpls <b>label-switched-path</b> <i>lsp-name</i> (primary   secondary) <i>path-name</i> ]                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Specify the conditions under which the path is selected to carry traffic. The <b>manual</b> and <b>unconditional</b> options are mutually exclusive.                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <b>manual</b> —The path is selected for carrying traffic if it is up and stable for at least the revert timer window (potentially before the revert timer has elapsed). Traffic is sent to other working paths if the current path is down or degraded (receiving errors).<br><br><b>unconditional</b> —The path is always selected for carrying traffic, even if it is currently down or degraded (receiving errors). |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Primary and Secondary LSPs</i></li> </ul>                                                                                                                                                                                                                                                                                                                      |

## signal-bandwidth

|                                 |                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>signal-bandwidth type;</code>                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> lsp-attributes],<br>[edit protocols mpls label-switched-path <i>lsp-name</i> lsp-attributes]                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                               |
| <b>Description</b>              | Specify the bandwidth encoding of the signal used for path computation and admission control.                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <b>type</b> —Configure the type of bandwidth encoding used on the LSP. It can be any of the following values: <b>10gigether</b> , <b>ds1</b> , <b>ds3</b> , <b>e1</b> , <b>e3</b> , <b>ethernet</b> , <b>fastether</b> , <b>gigether</b> , <b>stm-1</b> , <b>stm-4</b> , <b>stm-16</b> , <b>stm-64</b> , <b>stm-256</b> , <b>sts-1</b> , <b>vt1-5</b> , or <b>vt2</b> . |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring MPLS LSPs for GMPLS</i></li> </ul>                                                                                                                                                                                                                                                                              |

## smart-optimize-timer

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>smart-optimize-timer seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls],<br>[edit protocols mpls]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | <p>Enable the smart optimization timer. When you enable the smart optimization timer on a router, the Junos OS operates on the assumption that the original LSP path is preferable to any alternate or secondary path. When you enable the smart optimization timer and an LSP fails and its traffic is switched to an alternate path, the smart optimization timer starts and waits 3 minutes (this time is configurable). After 3 minutes have passed, the LSP is switched back to the original path. If the original path fails again and the LSP is switched to an alternate path again, the router waits 1 hour before attempting to switch the LSP back to its original path.</p> <p>If you want to disable the smart optimizer, you can set it to zero. The <b>smart-optimize-timer</b> value in seconds indicates the time before which the LSP is switched back to its primary path in case the primary path becomes available. Otherwise, the time to wait is controlled by the <b>optimize-timer</b>, which is usually set to a high value. Some ISPs have the <b>optimize-timer</b> set to once a day. Sometimes after the smart optimizer causes the LSP to be placed back on its primary path, the primary path goes down again within 60 minutes. When this happens, the <b>smart-optimize-timer</b> is disabled automatically, and the <b>optimize-timer</b> (regular path optimization) goes into effect. This is to protect against a flapping link being used.</p> |
| <b>Default</b>                  | The smart optimization timer is enabled by default.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <p><b>seconds</b>—(Optional) Specify the number of seconds to wait before switching an LSP back to its original path. If you do not specify the number of seconds, the default value is used.</p> <p><b>Range:</b> 0 through 65,535 seconds</p> <p><b>Default:</b> 180 seconds</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the Smart Optimize Timer for LSPs</i></li><li>• <i>Optimizing Signaled LSPs</i></li><li>• <a href="#">optimize-aggressive on page 5068</a></li><li>• <a href="#">optimize-timer on page 5071</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## standby

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | standby;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls],<br>[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i><br>(primary   secondary) <i>path-name</i> ],<br>[edit protocols mpls],<br>[edit protocols mpls label-switched-path <i>lsp-name</i> ],<br>[edit protocols mpls label-switched-path <i>lsp-name</i> (primary   secondary) <i>path-name</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Have the path remain up at all times to provide instant switchover if connectivity problems occur.                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Hot Standby of Secondary Paths for LSPs</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                |

## static-label-switched-path

---

**Syntax**    `static-label-switched-path lsp-name {`  
              `bypass bypass-name {`  
                  `bandwidth bps;`  
                  `description string;`  
                  `next-hop (address | interface-name | address/interface-name);`  
                  `push out-label;`  
                  `to address;`  
                  `}`  
              `ingress {`  
                  `bandwidth bps;`  
                  `class-of-service cos-value;`  
                  `description string;`  
                  `install {`  
                      `destination-prefix <active>;`  
                      `}`  
                  `link-protection bypass-name name;`  
                  `metric metric;`  
                  `next-hop (address | interface-name | address/interface-name);`  
                  `node-protection bypass-name name next-next-label label;`  
                  `no-install-to-address;`  
                  `policing {`  
                      `filter filter-name;`  
                      `no-auto-policing;`  
                      `}`  
                  `preference preference;`  
                  `push out-label;`  
                  `to address;`  
                  `}`  
              `transit incoming-label {`  
                  `bandwidth bps;`  
                  `description string;`  
                  `link-protection bypass-name name;`  
                  `next-hop (address | interface-name | address/interface-name);`  
                  `node-protection bypass-name name next-next-label label;`  
                  `pop;`  
                  `swap out-label;`  
                  `}`  
              `}`

**Hierarchy Level**    `[edit logical-systems logical-system-name protocols mpls],`  
                          `[edit protocols mpls]`

**Release Information**    Statement introduced in Junos OS Release 10.1.  
                              Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

**Description**        Configure a static LSP.

**Options**            *lsp-name*—Name of the path.

                      The remaining statements are explained separately.

|                              |                                                                                  |
|------------------------------|----------------------------------------------------------------------------------|
| <b>Required Privilege</b>    | routing—To view this statement in the configuration.                             |
| <b>Level</b>                 | routing-control—To add this statement to the configuration.                      |
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <i>Configuring Static LSPs</i></li></ul> |

## statistics (Protocols MPLS)

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>statistics {<br/>  auto-bandwidth;<br/>  file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;<br/>  interval <i>seconds</i>;<br/>  no-transit-statistics;<br/>  traffic-class-statistics;<br/>  transit-statistics-polling;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> protocols mpls],<br>[edit protocols mpls]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.<br><b>traffic-class-statistics</b> option introduced in Junos OS Release 14.2.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>         | Enable MPLS statistics collection and reporting.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>             | <p><b>file <i>filename</i></b>—(Optional) Name of the file to receive the output. We recommend that you place MPLS tracing output in the file mpls-stat in the /var/log directory.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <i>file</i> reaches its maximum size, it is renamed <i>file.0</i>, then <i>file.1</i>, and so on, until the maximum number of files is reached. Then, the oldest file is overwritten.</p> <p><b>Range:</b> 2 or more</p> <p><b>Default:</b> 2 files</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the <b>size</b> option.</p> <p><b>interval <i>seconds</i></b>—Interval at which to periodically collect statistics.</p> <p><b>Range:</b> 1 through 65,535</p> <p><b>Default:</b> 300 seconds</p> <p><b>no-world-readable</b>—(Optional) Prevent users from reading the log file.</p> <p><b>size <i>size</i></b>—(Optional) Maximum size of each file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a file named <i>file</i> reaches this size, it is renamed <i>file.0</i>. When the <i>file</i> again reaches its maximum size, <i>file.0</i> is renamed <i>file.1</i> and <i>file</i> is renamed <i>file.0</i>. This renaming scheme continues until the maximum number of files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum file size, you also must specify a maximum number of files with the <b>files</b> option.</p> <p><b>world-readable</b>—(Optional) Enable users to read the log file.</p> <p><b>Syntax:</b> Syntax: <b>xk</b> to specify KB, <b>xm</b> to specify MB, or <b>xg</b> to specify GB</p> <p><b>Range:</b> 10 KB through the maximum file size supported on your system</p> |

**Default:** 1 MB

**traffic-class-statistics**—(Optional) Create counters that maintain data traffic statistics per traffic class at the ingress of all types of LSPs and egress of ultimate hop popping (UHP) point-to-point LSPs. These counters are not created by default and are required to be configured to perform traffic-class-scoped loss measurement.

**transit-statistics-polling**—(Optional) Enable the polling and display of MPLS statistics for LSPs transiting the router. By default, RSVP does not periodically poll for transit LSP statistics. You cannot configure this statement and the **no-transit-statistics** statement at the same time.

The remaining statements are explained separately.

**Required Privilege Level** routing and trace—To view this statement in the configuration.  
routing-control and trace-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring MPLS to Gather Statistics on page 4980](#)
- [Configuring Automatic Bandwidth Allocation for LSPs](#)

## swap

|                                 |                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>swap out-label;</code>                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i> transit <i>incoming-label</i> ],<br>[edit protocols mpls static-label-switched-path <i>lsp-name</i> transit <i>incoming-label</i> ]  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                       |
| <b>Description</b>              | Remove the label at the top of the label stack and replace it with the specified label. Manually assigned incoming labels can have values from 1,000,000 through 1,048,575. This statement is used to configure static LSPs at transit routers. |
| <b>Options</b>                  | <b>out-label</b> —Manually assigned outgoing label value.<br><b>Range:</b> 0 through 1,048,575<br><b>Default:</b> If you do not define the <b>out-label</b> option, the original label value remains unchanged.                                 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">pop on page 5080</a></li> <li>• <a href="#">push on page 5083</a></li> <li>• <a href="#">Configuring Static LSPs</a></li> </ul>                                                            |

## switching-type

---

|                                 |                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | switching-type (fiber   lambda   psc-1   tdm);                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> lsp-attributes],<br>[edit protocols mpls label-switched-path <i>lsp-name</i> lsp-attributes]                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                          |
| <b>Description</b>              | Specify the switching method for the LSP. The switching method can be one of the following values: <ul style="list-style-type: none"><li>• <b>fiber</b>—Fiber switching</li><li>• <b>lambda</b>—Lambda switching</li><li>• <b>psc-1</b>—Packet switching</li><li>• <b>tdm</b>—Time-division multiplexing (TDM) switching</li></ul> |
| <b>Default</b>                  | psc-1                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring MPLS LSPs for GMPLS</i></li></ul>                                                                                                                                                                                                                                           |



## te-class-matrix

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>te-class-matrix {   tenumber {     priority <i>priority</i>;     traffic-class {       <i>ctnumber</i> priority <i>priority</i>;     }   } }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls diffserv-te],<br>[edit protocols mpls diffserv-te]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Specify the traffic engineering class matrix for a multiclass LSP or a DiffServ-aware traffic engineering LSP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Default</b>                  | <p>The default traffic engineering class matrix is:</p> <pre>te-class-matrix {   te0 traffic-class ct0 priority 7;   te1 traffic-class ct1 priority 7;   te2 traffic-class ct2 priority 7;   te3 traffic-class ct3 priority 7;   te4 traffic-class ct0 priority 0;   te5 traffic-class ct1 priority 0;   te6 traffic-class ct2 priority 0;   te7 traffic-class ct3 priority 0; }</pre> <p>If you define any of the traffic engineering classes, all the default values are dropped.</p>                                                                                                                                                                                                         |
| <b>Options</b>                  | <p><b><i>ctnumber</i></b>—Specify the number of the class type. It can be one of four values: <b>ct0</b>, <b>ct1</b>, <b>ct2</b>, or <b>ct3</b>.</p> <p><b><i>priority priority</i></b>—Specify the priority of the class type. It can be one of eight values from 0 through 7.</p> <p><b><i>tenumber</i></b>—Specify the number of the traffic engineering class. It can be one of eight values: <b>te0</b>, <b>te1</b>, <b>te2</b>, <b>te3</b>, <b>te4</b>, <b>te5</b>, <b>te6</b>, or <b>te7</b>. You must configure the traffic engineering classes in order, starting with <b>te0</b>.</p> <p><b><i>traffic-class</i></b>—Specify the traffic class for the traffic engineering class.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Configuring Routers for DiffServ-Aware Traffic Engineering</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

## traffic-engineering (Protocols MPLS)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | traffic-engineering (bgp   bgp-igp   bgp-igp-both-ribs   mpls-forwarding);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls],<br>[edit protocols mpls]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Select whether MPLS performs traffic engineering on BGP destinations only or on both BGP and IGP destinations. Affects only LSPs originating from this routing device, not transit or egress LSPs.                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Default</b>                  | bgp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <p><b>bgp</b>—On BGP destinations only. Ingress routes are installed in the inet.3 routing table.</p> <p><b>bgp-igp</b>—On both BGP and IGP destinations. Ingress routes are installed in the inet.0 routing table. If IGP shortcuts are enabled, the shortcut routes are automatically installed in the inet.0 routing table.</p> <p><b>bgp-igp-both-ribs</b>—On both BGP and IGP destinations. Ingress routes are installed in the inet.0 and inet.3 routing tables. This option is used to support VPNs.</p> <p><b>mpls-forwarding</b>—On both BGP and IGP destinations. Use ingress routes for forwarding only, not for routing.</p> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Traffic Engineering for LSPs</i></li><li>• <i>Configuring MPLS on Provider Edge Switches Using IP Over MPLS (CLI Procedure)</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                         |

## transit-lsp-association

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>transit-lsp-association <i>transit-association-lsp-group-name</i> {     from-1 <i>address-of-associated-lsp-1</i>;     from-2 <i>address-of-associated-lsp-2</i>;     lsp-name-1 <i>name-of-associated-lsp-1</i>;     lsp-name-2 <i>name-of-associated-lsp-2</i>; }</pre>                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit protocols mpls]                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Associate two label-switched paths (LSPs) at a transit node to configure a path for sending and receiving GAL and G-Ach messages for MPLS-TP OAM.                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><i>transit-association-lsp-group-name</i>—Name of the transit association LSP group.</p> <p><i>from-1 address-of-associated-lsp-1</i>—Address of the first associated LSP.</p> <p><i>from-2 address-of-associated-lsp-2</i>—Address of the second associated LSP.</p> <p><i>lsp-name-1 name-of-associated-lsp-1</i>—Name of the first associated LSP.</p> <p><i>lsp-name-2 name-of-associated-lsp-1</i>—Name of the second associated LSP.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring the MPLS Transport Profile for OAM</i></li> </ul>                                                                                                                                                                                                                                                                                                                                |

## te-class-matrix

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>te-class-matrix {<br/>    tenumber {<br/>        priority <i>priority</i>;<br/>        traffic-class {<br/>            ctnumber priority <i>priority</i>;<br/>        }<br/>    }<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls diffserv-te],<br>[edit protocols mpls diffserv-te]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Specify the traffic engineering class matrix for a multiclass LSP or a DiffServ-aware traffic engineering LSP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Default</b>                  | <p>The default traffic engineering class matrix is:</p> <pre>te-class-matrix {<br/>    te0 traffic-class ct0 priority 7;<br/>    te1 traffic-class ct1 priority 7;<br/>    te2 traffic-class ct2 priority 7;<br/>    te3 traffic-class ct3 priority 7;<br/>    te4 traffic-class ct0 priority 0;<br/>    te5 traffic-class ct1 priority 0;<br/>    te6 traffic-class ct2 priority 0;<br/>    te7 traffic-class ct3 priority 0;<br/>}</pre> <p>If you define any of the traffic engineering classes, all the default values are dropped.</p>                                                                                                                                |
| <b>Options</b>                  | <p><b>ctnumber</b>—Specify the number of the class type. It can be one of four values: <b>ct0</b>, <b>ct1</b>, <b>ct2</b>, or <b>ct3</b>.</p> <p><b>priority <i>priority</i></b>—Specify the priority of the class type. It can be one of eight values from 0 through 7.</p> <p><b>tenumber</b>—Specify the number of the traffic engineering class. It can be one of eight values: <b>te0</b>, <b>te1</b>, <b>te2</b>, <b>te3</b>, <b>te4</b>, <b>te5</b>, <b>te6</b>, or <b>te7</b>. You must configure the traffic engineering classes in order, starting with <b>te0</b>.</p> <p><b>traffic-class</b>—Specify the traffic class for the traffic engineering class.</p> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>Configuring Routers for DiffServ-Aware Traffic Engineering</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

to

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | to <i>address</i> ;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i> bypass],<br>[edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i> ingress],<br>[edit protocols mpls label-switched-path <i>lsp-name</i> ],<br>[edit protocols mpls static-label-switched-path <i>lsp-name</i> bypass],<br>[edit protocols mpls static-label-switched-path <i>lsp-name</i> ingress] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Specify the egress router of a dynamic LSP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <i>address</i> —Address of the egress router.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring the Ingress and Egress Router Addresses for LSPs</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## traceoptions (Protocols MPLS)

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>traceoptions {<br/>    file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;<br/>    flag <i>flag</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> protocols mpls],<br>[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> ],<br>[edit protocols mpls],<br>[edit protocols mpls label-switched-path <i>lsp-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.<br><b>ted-export</b> option introduced in Junos OS Release 14.2.<br><b>ted-import</b> option introduced in Junos OS Release 14.2.<br><b>lsp-history</b> option added in Junos OS Release 15.1.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>         | Configure MPLS tracing options at the protocol level or for a label-switched path.<br><br>To specify more than one tracing operation, include multiple <b>flag</b> statements.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Default</b>             | The default MPLS protocol-level tracing options are inherited from the routing protocols <b>traceoptions</b> statement included at the [edit routing-options] hierarchy level.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>             | <p><b>filename</b>—Name of the file to receive the output of the tracing operation. All files are placed in the directory <b>/var/log</b>. We recommend that you place MPLS tracing output in the file <b>mpls-log</b>.</p> <p><b>files number</b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p><b>Range:</b> 2 through 1000</p> <p><b>Default:</b> 2 files</p> <p>If you specify a maximum number of files, you must also include the <b>size</b> statement to specify the maximum file size.</p> <p><b>flag</b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements.</p> <p><b>MPLS Tracing Flags</b></p> <ul style="list-style-type: none"><li>• <b>all</b>—Trace all operations</li><li>• <b>autobw-state</b>—Automatic bandwidth events.</li><li>• <b>connection</b>—All circuit cross-connect (CCC) activity</li><li>• <b>connection-detail</b>—Detailed CCC activity</li></ul> |

- **cspf**—CSPF computations
- **cspf-link**—Links visited during CSPF computations
- **cspf-node**—Nodes visited during CSPF computations
- **error**—MPLS error packets
- **graceful-restart**—Trace MPLS graceful restart events
- **lsp-history**—Trace LSP history events
- **lsping**—Trace lsping packets and return codes
- **nsr-synchronization**—Trace NSR synchronization events
- **nsr-synchronization-detail**—Trace NSR synchronization events in detail
- **state**—All LSP state transitions
- **static**—Trace static label-switched path
- **ted-export**—Trace leaking of entries from **lsdist.0** table into the traffic engineering database
- **ted-import**—Trace leaking traffic engineering database entries into the **lsdist.0** table
- **timer**—Timer usage

**no-world-readable**—(Optional) Allow only certain users to read the log file.

**size** *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches this size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

**Syntax:** *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 1 MB

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of files.

**world-readable**—(Optional) Allow any user to read the log file.

|                                 |                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing and trace—To view this statement in the configuration.<br>routing-control and trace-control—To add this statement to the configuration. |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Tracing MPLS and LSP Packets and Operations</i></li> </ul>                                          |

## traffic-engineering (Protocols MPLS)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | traffic-engineering (bgp   bgp-igp   bgp-igp-both-ribs   mpls-forwarding);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols mpls],<br>[edit protocols mpls]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Select whether MPLS performs traffic engineering on BGP destinations only or on both BGP and IGP destinations. Affects only LSPs originating from this routing device, not transit or egress LSPs.                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Default</b>                  | bgp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <b>bgp</b> —On BGP destinations only. Ingress routes are installed in the inet.3 routing table.<br><br><b>bgp-igp</b> —On both BGP and IGP destinations. Ingress routes are installed in the inet.0 routing table. If IGP shortcuts are enabled, the shortcut routes are automatically installed in the inet.0 routing table.<br><br><b>bgp-igp-both-ribs</b> —On both BGP and IGP destinations. Ingress routes are installed in the inet.0 and inet.3 routing tables. This option is used to support VPNs.<br><br><b>mpls-forwarding</b> —On both BGP and IGP destinations. Use ingress routes for forwarding only, not for routing. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Traffic Engineering for LSPs</i></li><li>• <i>Configuring MPLS on Provider Edge Switches Using IP Over MPLS (CLI Procedure)</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                      |



## transit-lsp-association

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>transit-lsp-association <i>transit-association-lsp-group-name</i> {     from-1 <i>address-of-associated-lsp-1</i>;     from-2 <i>address-of-associated-lsp-2</i>;     lsp-name-1 <i>name-of-associated-lsp-1</i>;     lsp-name-2 <i>name-of-associated-lsp-2</i>; }</pre>                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit protocols mpls]                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Associate two label-switched paths (LSPs) at a transit node to configure a path for sending and receiving GAL and G-Ach messages for MPLS-TP OAM.                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><i>transit-association-lsp-group-name</i>—Name of the transit association LSP group.</p> <p><i>from-1 address-of-associated-lsp-1</i>—Address of the first associated LSP.</p> <p><i>from-2 address-of-associated-lsp-2</i>—Address of the second associated LSP.</p> <p><i>lsp-name-1 name-of-associated-lsp-1</i>—Name of the first associated LSP.</p> <p><i>lsp-name-2 name-of-associated-lsp-1</i>—Name of the second associated LSP.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring the MPLS Transport Profile for OAM</i></li> </ul>                                                                                                                                                                                                                                                                                                                                |

## Monitoring Commands for MPLS

- `clear mpls lsp`
- `monitor label-switched-path`
- `ping mpls bgp`
- `ping mpls l2circuit`
- `ping mpls l3vpn`
- `ping mpls lsp-end-point`
- `request mpls lsp adjust-autobandwidth`
- `show security keychain`
- `show link-management`
- `show link-management peer`
- `show link-management routing`
- `show link-management statistics`

- `show link-management te-link`
- `show mpls call-admission-control`
- `show mpls cspf`
- `show mpls diffserv-te`
- `show route forwarding-table`
- `show mpls interface`
- `show link-management statistics`
- `show link-management te-link`
- `show mpls call-admission-control`
- `show mpls cspf`
- `show mpls diffserv-te`
- `show route forwarding-table`
- `show mpls interface`
- `show mpls lsp`
- `show mpls lsp autobandwidth`
- `show mpls path`
- `show route table`
- `show route forwarding-table`
- `show mpls static-lsp`
- `show ted database`
- `show ted link`
- `show ted protocol`

## clear mpls lsp

**List of Syntax**    [Syntax on page 5107](#)  
                          [Syntax \(EX and QFX Series Switches\) on page 5107](#)

**Syntax**    clear mpls lsp  
                  <autobandwidth>  
                  <logical-system (all | *logical-system-name*)>  
                  <name *name*>  
                  <optimize | optimize-aggressive>  
                  <path *regular-expression*>  
                  <statistics>

**Syntax (EX and QFX Series Switches)**    clear mpls lsp  
                  <autobandwidth>  
                  <name *name*>  
                  <optimize | optimize-aggressive>  
                  <path *regular-expression*>  
                  <statistics>

**Release Information**    Command introduced before Junos OS Release 7.4.  
                                  Command introduced in Junos OS Release 9.5 for EX Series switches.  
                                  Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

**Description**    Release the routes and states associated with MPLS label-switched paths (LSPs), and start new LSPs.



**CAUTION:** This command disconnects existing Resource Reservation Protocol (RSVP) sessions on the ingress routing device. If there is a time lag between the old path being torn down and the new path being set up, this command might impact traffic traveling along the LSPs.

**Options**    **none**—Reset and restart all LSPs that originated from this routing device; that is, all LSPs for which this routing device is the ingress routing device. Depending on the number of LSPs involved, it might take a while to restart all the LSPs.

**autobandwidth**—(Optional) Clear LSP autobandwidth counters.

**logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

**name *name***—(Optional) Reset and restart the specified LSP or group of LSPs. You can include wildcard characters in the interface name, as described in the *Junos Network Interfaces Configuration Guide*.

**optimize | optimize-aggressive**—(Optional) Run nonpreemptive optimization or aggressive optimization computation now.

**path *regular-expression***—(Optional) Clear the specific LSP path matching the specified regular expression.

**statistics**—(Optional) Clear LSP statistics. You cannot clear the MPLS LSP statistics using a regular expression (**name** and **path** options) on transit routers.

**Required Privilege Level**

clear

**Related Documentation**

- [show mpls lsp on page 5176](#)
- [show rsvp session on page 5365](#)

**List of Sample Output** [clear mpls lsp on page 5108](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### clear mpls lsp

```
user@host> clear mpls lsp
```

## monitor label-switched-path

**Syntax** `monitor label-switched-path lsp-name`  
`<logical-system (logical-system-name)>`

**Release Information** Command introduced before Junos OS Release 7.4.  
 Logical system support introduced in Junos OS Release 9.4.  
 Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

**Description** Display the real-time status of the specified RSVP label-switched path (LSP). You can also use this command to monitor LSPs configured within logical systems.

**Options** `logical-system ( logical-system-name )`—(Optional) Perform this operation on all logical systems or on a particular logical system.

*lsp-name*—Name of the LSP.

**Additional Information** You can track the amount of traffic traversing an RSVP LSP and observe its essential parameters, such as uptime, ingress and egress addresses, labels, routes, and ports. Values are typically sampled every second. The display also allows you to scroll to other currently running LSPs. You cannot use this command to display information about static LSPs or LDP-signaled LSPs.

The output of this command shows how much each field has changed since you started the command or since you cleared the counters by using the `c` key. To control the output of the **monitor label-switched-path** command while it is running, use the keys listed in [Table 385](#). The keys are not case-sensitive.

**Table 385: Output Control Keys for the monitor label-switched-path Command**

| Key      | Action                                                                                                                            |
|----------|-----------------------------------------------------------------------------------------------------------------------------------|
| c        | Clears the screen and refreshes the display for this LSP.                                                                         |
| f        | Freezes the display, preventing new information from being displayed.                                                             |
| l        | Monitors a different LSP. After you type l, you can type the new LSP name.                                                        |
| n        | Displays information about the next LSP (whose name is alphabetically higher than the current LSP name) configured on the router. |
| p        | Goes to the previous LSP (whose name is alphabetically lower than the current LSP name) configured on the router.                 |
| q or Esc | Quits the command and returns to the command prompt.                                                                              |
| t        | Thaws, or restarts, the data display for this LSP.                                                                                |

**Required Privilege Level** trace

**List of Sample Output** [monitor label-switched-path on page 511](#)

**Output Fields** [Table 386](#) describes the output fields for the **monitor label-switched-path** command. Output fields are listed in the approximate order in which they appear.

**Table 386: monitor label-switched-path Output Fields**

| Field Name | Field Description                                                                                                                                                                                                                                                                                                                                                                                           |
|------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| (1)        | Displays the following information: <ul style="list-style-type: none"> <li>• <b>hostname</b>—Name of the router.</li> <li>• <b>Seconds</b>—Time elapsed since this display was started.</li> <li>• <b>Time</b>—Current local time.</li> </ul>                                                                                                                                                               |
| (2)        | <b>Delay</b> —Length of the time delay, in milliseconds, required to obtain the information in the monitor display. The first number shows the current sampling delay. The second number shows the shortest delay recorded to date. The third number shows the worst delay recorded to date. This delay can vary substantially depending on the system load.                                                |
| (3)        | Displays the following: <ul style="list-style-type: none"> <li>• <b>To</b>—Destination address of the LSP.</li> <li>• <b>From</b>—Originating address of the LSP.</li> <li>• <b>State</b>—Current state of the LSP: <b>Up</b> or <b>Down</b>.</li> </ul>                                                                                                                                                    |
| (4)        | Displays the following: <ul style="list-style-type: none"> <li>• <b>LSPName</b>—Name of the LSP.</li> <li>• <b>Type</b>—Type of LSP: <b>Ingress</b>, <b>Egress</b>, or <b>Transit</b>.</li> </ul>                                                                                                                                                                                                           |
| (5)        | Displays the following: <ul style="list-style-type: none"> <li>• <b>Label in</b>—Incoming label of the LSP.</li> <li>• <b>Label out</b>—Outgoing label of the LSP.</li> </ul>                                                                                                                                                                                                                               |
| (6)        | <b>Port number</b> —Port number for the sending router, the port number for the receiving router, and the protocol ID. For MPLS traffic engineering applications, the protocol ID is always 0.                                                                                                                                                                                                              |
| (7/8)      | <b>Record route</b> —All intermediate and egress router addresses for this LSP.                                                                                                                                                                                                                                                                                                                             |
| (9/10/11)  | Displays traffic statistics: <ul style="list-style-type: none"> <li>• <b>Output packets</b>—Number of packets that have traversed this LSP, and the change (delta) in the number since the last sample, typically 1 second ago.</li> <li>• <b>Output bytes</b>—Number of bytes that have traversed this LSP, and the change (delta) in the number since the last sample, typically 1 second ago.</li> </ul> |
| (12)       | Displays any errors the router encountered while attempting to retrieve information on the LSP.                                                                                                                                                                                                                                                                                                             |
| (13)       | Lists the keyboard commands you can use to navigate to other LSPs. For a description of the keyboard commands, see <a href="#">Table 385</a> .                                                                                                                                                                                                                                                              |

## Sample Output

### monitor label-switched-path

```
user@host> monitor label-switched-path
(1) host Seconds: 112 Time: 15:32:22
(2) Delay: 0/0/0
(3) To 10.10.10.16, From 10.10.10.17, state: Up
(4) LSPname: k, type: Ingress
(5) Label in: -, Label out: 126000
(6) Port number: sender 1, receiver 45583, protocol 0
(7) Record Route: <self> 192.168.224.196
(8) 192.168.224.202 192.168.224.179
(9) Traffic statistics: Current delta
(10) Output packets: 0 [0]
(11) Output bytes: 0 [0]
(12)
(13)Next='n', Prev='p', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c',
 LSP='l'
```

## ping mpls bgp

---

**Syntax**    ping mpls bgp fec  
             <bottom-label-ttl>  
             <count *count*>  
             <destination *address*>  
             <detail>  
             <exp *forwarding-class*>  
             <instance *routing-instance-name*>  
             <logical-system (all | *logical-system-name*)>  
             <size *bytes*>  
             <source *source-address*>  
             <sweep>

**Release Information**    Command introduced in Junos OS Release 11.1.

**Description**    Check the operability of MPLS BGP-signaled label-switched path (LSP) connections. Press Ctrl+c to interrupt a **ping mpls bgp** command.



**NOTE:** The **ping mpls bgp fec** command only supports single paths.

---

**Options**    **bottom-label-ttl**—(Optional) Time-to-live (TTL) value for the bottom label in the label stack. The range of values is 1 through 255. The default value is **255**.

**count *count***—(Optional) Number of ping requests to send. If **count** is not specified, five ping requests are sent. The range of values is 1 through 1,000,000. The default value is 5.

**destination *address***—(Optional) Specify an address other than the default (127.0.0.1/32) for the ping echo requests. The address can be anything within the 127/8 subnet.

**detail**—(Optional) Display detailed information about the echo requests sent and received.

**exp *forwarding-class***—(Optional) Value of the forwarding class for the MPLS ping packets.

**fec**—Ping a BGP-signaled LSP using the forwarding equivalence class (FEC) prefix and length.

**instance *routing-instance-name***—(Optional) Allows you to ping a combination of the routing instance and forwarding equivalence class (FEC) associated with an LSP.

**logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on the specified logical system.

**size *bytes***—(Optional) Size of the LSP ping request packet (88 through 65468 bytes). Packets are 4-byte aligned. For example, If you enter a size of 89, 90, 91, or 92, the router or switch uses a size value of 92 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 88-byte minimum.



**source *source-address***—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (**lo.0**).

**sweep**—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

**Additional Information** If the LSP changes, the label and interface information displayed when you issued the **ping** command continues to be used. You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the remote router or switch to ping an LSP terminating there. You must configure MPLS even if you intend to ping only BGP forwarding equivalence classes (FECs).

In asymmetric MTU scenarios, the echo response might be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

**Required Privilege Level** network

**List of Sample Output** [ping mpls bgp fec count on page 5113](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. Packets with error codes are not counted in the received packets count. They are accounted for separately. To display the error codes, use the **detail** option (for example, **ping mpls bgp 10.255.245.222 detail**).

## Sample Output

### ping mpls bgp fec count

```
user@host> ping mpls bgp 10.255.245.222 count 10
!!!xxx...x--- 1sping statistics ---10 packets transmitted, 3 packets received,
70% packet loss 4 packets received with error status, not counted as received.
```

## ping mpls l2circuit

---

**Syntax** ping mpls l2circuit (interface *interface-name* | virtual-circuit *virtual-circuit-id* neighbor *address*)  
<count *count*>  
<destination *address*>  
<detail>  
<exp *forwarding-class*>  
<logical-system (all | *logical-system-name*)>  
reply-mode (application-level-control-channel | ip-udp | no-reply)  
<size *bytes*>  
<source *source-address*>  
<sweep>  
<v1>

**Release Information** Command introduced before Junos OS Release 7.4.  
Command introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 12.3X50 for the QFX Series.  
The **size** and **sweep** options were introduced in Junos OS Release 9.6.  
The **reply-mode** option and its suboptions are introduced in Junos OS Release 10.4R1.

**Description** Check the operability of the MPLS Layer 2 circuit connections. Type Ctrl+c to interrupt a ping mpls l2circuit command.

**Options** **count** *count*—(Optional) Number of ping requests to send. If **count** is not specified, five ping requests are sent. The range of values is 1 through **1,000,000**. The default value is 5.

**destination** *address*—(Optional) Specify an address other than the default (**127.0.0.1/32**) for the ping echo requests. The address can be anything within the **127/8** subnet.

**detail**—(Optional) Display detailed information about the echo requests sent and received.

**exp** *forwarding-class*—(Optional) Value of the forwarding class for the MPLS ping packets.

**interface** *interface-name*—Ping an interface configured for the Layer 2 circuit on the egress provider edge (PE) router.

**logical-system** (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on the specified logical system.

**reply-mode**—(Optional) Reply mode for the ping request. This option has the following suboptions:

**application-level-control-channel**—Reply using an application level control channel.

**ip-udp**—Reply using an IPv4 or IPv6 UDP packet.

**no-reply**—Do not reply to the ping request.



**NOTE:** The **reply-mode** option and its suboptions **application-level-control-channel**, **ip-udp**, and **no-reply** are also available in Junos OS Release 10.2R4 and 10.3R2.

**size bytes**—(Optional) Size of the label-switched path (LSP) ping request packet (96 through 65468 bytes). Packets are 4-byte aligned. For example, If you enter a size of 97, 98, 99, or 100, the router or switch uses a size value of 100 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 96-byte minimum.

**source source-address**—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (**lo.0**).

**sweep**—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

**vt**—(Optional) Use the type 9 Layer 2 circuit type, length, and value (TLV).

**virtual-circuit virtual-circuit-id neighbor address**—Ping the virtual circuit identifier on the egress PE router or switch and the specified neighbor, testing the integrity of the Layer 2 circuit between the ingress and egress PE routers or switches.

**Additional Information** You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the egress PE router or switch (the router or switch receiving the MPLS echo packets) to ping a Layer 2 circuit.

In asymmetric MTU scenarios, the echo response might be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

**Required Privilege Level** network

**List of Sample Output** [ping mpls l2circuit interface on page 5116](#)  
[ping mpls l2circuit virtual-circuit detail on page 5116](#)  
[ping mpls l2circuit interface <interface-name> reply-mode on page 5116](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. Packets with an error code are not counted in the received packets count. They are accounted for separately.

## Sample Output

### ping mpls l2circuit interface

```
user@host> ping mpls l2circuit interface so-1/0/0.1
Request for seq 1, to interface 69, labels <100000, 100208>, packet size 100
Reply for seq 1, return code: Egress-ok, time: 0.439 ms
```

### ping mpls l2circuit virtual-circuit detail

```
user@host> ping mpls l2circuit virtual-circuit 200 neighbor 10.255.245.122/32 detail
Request for seq 1, to interface 68, labels <100048, 100128>, packet size 100

Reply for seq 1, return code: Egress-ok time: 0.539 ms
```

### ping mpls l2circuit interface <interface-name> reply-mode

```
user@host> ping mpls l2circuit interface lt-1/2/0.21 reply-mode application-level-control-channel
!!!!
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

## ping mpls l3vpn

**Syntax** ping mpls l3vpn prefix *prefix-name*  
 <*l3vpn-name*>  
 <bottom-label-ttl>  
 <count *count*>  
 <destination *address*>  
 <detail>  
 <exp *forwarding-class*>  
 <logical-system (all | *logical-system-name*)>  
 <size *bytes*>  
 <source *source-address*>  
 <sweep>

**Release Information** Command introduced before Junos OS Release 7.4.  
 Command introduced in Junos OS Release 9.0 for EX Series switches.  
 The **size** and **sweep** options were introduced in Junos OS Release 9.6.  
 Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

**Description** Check the operability of a MPLS Layer 3 virtual private network (VPN) connection. Type Ctrl+c to interrupt a **ping mpls l3vpn** command.

**Options** **bottom-label-ttl**—(Optional) Display the time-to-live value for the bottom label in the label stack.

**count** *count*—(Optional) Number of ping requests to send. If **count** is not specified, five ping requests are sent. The range of values is 1 through **1,000,000**. The default value is **5**.

**destination** *address*—(Optional) Specify an address other than the default (**127.0.0.1/32**) for the ping echo requests. The address can be anything within the **127/8** subnet.

**detail**—(Optional) Display detailed information about the echo requests sent and received.

**exp** *forwarding-class*—(Optional) Value of the forwarding class for the MPLS ping packets.

**l3vpn-name**—(Optional) Layer 3 VPN name.

**logical-system** (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on the specified logical system.

**prefix** *prefix-name*—Ping to test whether a prefix is present in a provider edge (PE) router's or switch's VPN routing and forwarding (VRF) table, by means of a Layer 3 VPN destination prefix. This option does not test the connection between a PE router or switch and a customer edge (CE) router or switch.

**size** *bytes*—(Optional) Size of the label-switched path (LSP) ping request packet (**96** through **65468** bytes). Packets are 4-byte aligned. For example, If you enter a size of 97, 98, 99, or 100, the router or switch uses a size value of 100 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 96-byte minimum.

**source *source-address***—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (**lo.0**).

**sweep**—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

**Additional Information** You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the egress PE router or switch (the router or switch receiving the MPLS echo packets) to ping a Layer 2 circuit.

In asymmetric MTU scenarios, the echo response might be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

If the Layer 3 VPN traffic transits a route reflector within the network, the **ping mpls l3vpn** command does not work.

**Required Privilege Level** network

**List of Sample Output** [ping mpls l3vpn on page 5118](#)  
[ping mpls l3vpn detail on page 5118](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code these packets are not counted in the received packets count. They are accounted for separately.

## Sample Output

### ping mpls l3vpn

```
user@host> ping mpls l3vpn vpn1 prefix 10.255.245.122/32
!!!!
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

### ping mpls l3vpn detail

```
user@host> ping mpls l3vpn vpn1 prefix 10.255.245.122/32 detail
Request for seq 1, to interface 68, labels <100128, 100112>
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 68, labels <100128, 100112>
Reply for seq 2, return code: Egress-ok
Request for seq 3, to interface 68, labels <100128, 100112>
Reply for seq 3, return code: Egress-ok
Request for seq 4, to interface 68, labels <100128, 100112>
Reply for seq 4, return code: Egress-ok
Request for seq 5, to interface 68, labels <100128, 100112>
Reply for seq 5, return code: Egress-ok
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```



## ping mpls lsp-end-point

---

**Syntax**    ping mpls lsp-end-point *prefix-name*  
             <count *count*>  
             <destination *address*>  
             <detail>  
             <exp *forwarding-class*>  
             <instance *routing-instance-name*>  
             <logical-system (all | *logical-system-name*)>  
             <size *bytes*>  
             <source *source-address*>  
             <sweep>

**Release Information**    Command introduced before Junos OS Release 7.4.  
                             Command introduced in Junos OS Release 9.0 for EX Series switches.  
                             The **size** and **sweep** options were introduced in Junos OS Release 9.6.  
                             The **instance** option was introduced in Junos OS Release 10.0.  
                             Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

**Description**    Check the operability of MPLS label-switched path (LSP) endpoint connections. Type Ctrl+c to interrupt a **ping mpls** command.

**Options**    **count** *count*—(Optional) Number of ping requests to send. If **count** is not specified, five ping requests are sent. The range of values is 1 through 1,000,000. The default value is 5.

**destination** *address*—(Optional) Specify an address other than the default (127.0.0.1/32) for the ping echo requests. The address can be anything within the 127/8 subnet.

**detail**—(Optional) Display detailed information about the echo requests sent and received.

**exp** *forwarding-class*—(Optional) Value of the forwarding class for the MPLS ping packets.

**instance** *routing-instance-name*—(Optional) Ping a combination of the routing instance and forwarding equivalence class (FEC) associated with an LSP connection.

**logical-system** (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on the specified logical system.

**prefix-name**—LDP forwarding equivalence class (FEC) prefix or RSVP LSP endpoint address.

**size** *bytes*—(Optional) Size of the LSP ping request packet. If the endpoint is LDP-based, the minimum size of the packet is 88 bytes. If the endpoint is RSVP-based, the minimum size of the packet is 100 bytes. The maximum size in either case is 65468 bytes.

**source** *source-address*—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (lo.0).



**sweep**—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

**Additional Information** If the LSP changes, the label and interface information displayed when you issued the **ping** command continues to be used. You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the remote router or switch to ping an LSP terminating there. You must configure MPLS even if you intend to ping only LDP forwarding equivalence classes (FECs).

In asymmetric MTU scenarios, the echo response might be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

**Required Privilege Level** network

**List of Sample Output** [ping mpls lsp-end-point detail on page 5121](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. Packets with an error code are not counted in the received packets count. They are accounted for separately.

## Sample Output

### [ping mpls lsp-end-point detail](#)

```
user@host> ping mpls lsp-end-point 10.255.245.119 detail
Route to end point address is via LDP FEC
Request for seq 1, to interface 67, label 100032
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 67, label 100032
Reply for seq 2, return code: Egress-ok
Request for seq 3, to interface 67, label 100032
Reply for seq 3, return code: Egress-ok
Request for seq 4, to interface 67, label 100032
Reply for seq 4, return code: Egress-ok
Request for seq 5, to interface 67, label 100032
Reply for seq 5, return code: Egress-ok
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

## request mpls lsp adjust-autobandwidth

---

|                                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                      | <a href="#">Syntax on page 5122</a><br><a href="#">Syntax (EX and QFX Series Switches) on page 5122</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Syntax</b>                              | <pre>request mpls lsp adjust-autobandwidth &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;name <i>lsp-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Syntax (EX and QFX Series Switches)</b> | <pre>request mpls lsp adjust-autobandwidth &lt;name <i>lsp-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>                 | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.5 for EX Series switches.<br>Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>                         | <p>Manually trigger a bandwidth allocation adjustment for active label-switched paths (LSPs).</p> <p>Without running this command, the bandwidth adjustment is recomputed at a configurable interval. The default interval is 5 minutes. If you do not want to wait for the periodic adjustment (for example, during a software demonstration), this command is useful.</p> <p>During bandwidth allocation adjustment, the LSP stays up to enable the bandwidth to be changed without dropping any traffic. This functionality is often referred to as <i>make-before-break</i>.</p> |
| <b>Options</b>                             | <p><b>none</b>—Manually trigger a bandwidth allocation adjustment for all active LSP paths.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>name <i>lsp-name</i></b>—(Optional) Manually trigger a bandwidth allocation adjustment on the specified LSP only.</p>                                                                                                                                                                                              |
| <b>Additional Information</b>              | <p>For this command to work properly, the following conditions must exist:</p> <ul style="list-style-type: none"><li>• Automatic bandwidth allocation must be enabled on the LSP. The parameters for adjustment interval and maximum average bandwidth are not reset after you issue the <b>request mpls lsp adjust-autobandwidth</b> command.</li><li>• The difference between the adjusted bandwidth and the current LSP path bandwidth must be greater than the threshold limit.</li></ul>                                                                                        |
| <b>Required Privilege Level</b>            | clear, maintenance                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>               | <ul style="list-style-type: none"><li>• <a href="#">auto-bandwidth on page 5029</a></li><li>• <i>Configuring Automatic Bandwidth Allocation for LSPs</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                   |

**List of Sample Output** [request mpls lsp adjust-auto-bandwidth on page 5123](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

[request mpls lsp adjust-auto-bandwidth](#)

```
user@host> request mpls lsp adjust-auto-bandwidth
```

## show security keychain

|                                 |                                                                                                                                                                                                                                                                                           |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show security keychain<br><brief   detail>                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.2.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                      |
| <b>Description</b>              | Display information about authentication keychains configured for the Border Gateway Protocol (BGP), the Label Distribution Protocol (LDP) routing protocols, the Bidirectional Forwarding Detection (BFD) protocol, and the Intermediate System-to-Intermediate System (IS-IS) protocol. |
| <b>Options</b>                  | <b>none</b> —Display information about authentication keychains.<br><b>brief   detail</b> —(Optional) Display the specified level of output.                                                                                                                                              |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                      |
| <b>List of Sample Output</b>    | <a href="#">show security keychain brief on page 5125</a><br><a href="#">show security keychain detail on page 5126</a>                                                                                                                                                                   |
| <b>Output Fields</b>            | <a href="#">Table 334</a> describes the output fields for the <b>show security keychain</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                              |

**Table 387: show security keychain Output Fields**

| Field Name               | Field Description                                                                        | Level of Output |
|--------------------------|------------------------------------------------------------------------------------------|-----------------|
| <b>keychain</b>          | The name of the keychain in operation.                                                   | All levels      |
| <b>Active-ID Send</b>    | Number of routing protocols packets sent with the active key.                            | All levels      |
| <b>Active-ID Receive</b> | Number of routing protocols packets received with the active key.                        | All levels      |
| <b>Next-ID Send</b>      | Number of routing protocols packets sent with the next key.                              | All levels      |
| <b>Next-ID Receive</b>   | Number of routing protocols packets received with the next key.                          | All levels      |
| <b>Transition</b>        | Amount of time until the current key will be replaced with the next key in the keychain. | All levels      |
| <b>Tolerance</b>         | Configured clock-skew tolerance, in seconds, for accepting keys for a key chain.         | All levels      |
| <b>Id</b>                | Identification number configured for the current key.                                    | <b>detail</b>   |
| <b>Algorithm</b>         | Authentication algorithm configured for the current key.                                 | <b>detail</b>   |

Table 387: show security keychain Output Fields (*continued*)

| Field Name        | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Level of Output |
|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>State</b>      | <p>State of the current key.</p> <p>The value can be:</p> <ul style="list-style-type: none"> <li>• <b>receive</b></li> <li>• <b>send</b></li> <li>• <b>send-receive</b></li> </ul> <p>For the active key, the <b>State</b> can be <b>send-receive</b>, <b>send</b>, or <b>receive</b>. For keys that have a future start time, the <b>State</b> is <b>inactive</b>. Compare the <b>State</b> field to the <b>Mode</b> field.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>detail</b>   |
| <b>Option</b>     | <p>For IS-IS only, the option determines how Junos OS encodes the message authentication code in routing protocol packets.</p> <p>The values can be:</p> <ul style="list-style-type: none"> <li>• <b>basic</b>—Based on RFC 5304.</li> <li>• <b>isis-enhanced</b>—Based on RFC 5310.</li> </ul> <p>The default value is <b>basic</b>. When you configure the <b>isis-enhanced</b> option, Junos OS sends RFC 5310-encoded routing protocol packets and accepts both RFC 5304-encoded and RFC 5310-encoded routing protocol packets that are received from other devices.</p> <p>When you configure <b>basic</b> (or do not include the <b>options</b> statement in the key configuration) Junos OS sends and receives RFC 5304-encoded routing protocols packets, and drops 5310-encoded routing protocol packets that are received from other devices.</p> <p>Because this setting is for IS-IS only, the TCP and the BFD protocol ignore the encoding option configured in the key.</p> | <b>detail</b>   |
| <b>Start-time</b> | Time that the current key became active.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail</b>   |
| <b>Mode</b>       | <p>Mode of each key (Informational only.)</p> <p>The value can be</p> <ul style="list-style-type: none"> <li>• <b>receive</b></li> <li>• <b>send</b></li> <li>• <b>send-receive</b></li> </ul> <p>The mode of the key is based on the configuration. Suppose you configure two keys, one with a start-time of today and the other with a start-time of next week. For both keys, the <b>Mode</b> can be <b>send-receive</b>, <b>send</b>, or <b>receive</b>, regardless of the configured start-time. Compare the <b>Mode</b> field to the <b>State</b> field.</p>                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail</b>   |

## Sample Output

show security keychain brief

```
user@host> show security keychain brief
```

| keychain | Active-ID |         | Next-ID |         | Transition | Tolerance |
|----------|-----------|---------|---------|---------|------------|-----------|
|          | Send      | Receive | Send    | Receive |            |           |
| hakr     | 3         | 3       | 1       | 1       | 1d 23:58   | 3600      |

#### show security keychain detail

```
user@host> show security keychain detail
keychain Active-ID Next-ID Transition Tolerance
 Send Receive Send Receive
hakr 3 3 1 1 1d 23:58 3600
 Id 3, Algorithm hmac-md5, State send-receive, Option basic
 Start-time Wed Aug 11 16:28:00 2010, Mode send-receive
 Id 1, Algorithm hmac-md5, State inactive, Option basic
 Start-time Fri Aug 20 11:30:57 2010, Mode send-receive
```

## show link-management

|                                 |                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show link-management                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.5 for EX Series switches.                                                                                                                                                                                                        |
| <b>Description</b>              | Display Multiprotocol Label Switching (MPLS) peer and traffic engineering link information.                                                                                                                                                                                                                                  |
| <b>Options</b>                  | This command has no options.                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show link-management peer on page 5131</a></li> <li>• <a href="#">show link-management routing on page 5133</a></li> <li>• <a href="#">show link-management statistics on page 5136</a></li> <li>• <a href="#">show link-management te-link on page 5138</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show link-management on page 5130</a>                                                                                                                                                                                                                                                                            |
| <b>Output Fields</b>            | Table 388 describes the output fields for the <b>show link-management</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                   |

**Table 388: show link-management Output Fields**

| Field Name        | Field Description                                                                                                 |
|-------------------|-------------------------------------------------------------------------------------------------------------------|
| Peer Name         | Name of the peer.                                                                                                 |
| System identifier | Internal identifier for the peer. The range of values is 0 through 64,000.                                        |
| State             | State of the peer: <b>Up</b> or <b>Down</b> .                                                                     |
| Control address   | Address to which a control channel is established.                                                                |
| CC local ID       | Identifier assigned to the control channel by the local peer. The range of values is 1 through 4,294,967,296.     |
| CC remote ID      | Identifier assigned to the control channel by the remote peer. The range of values is 1 through 4,294,967,296.    |
| State             | State of the control channel: <b>Up</b> or <b>Down</b> .                                                          |
| TxSeqNum          | Sequence number of the hello message being sent to the peer. The range of values is 1 through 4,294,967,295.      |
| RcvSeqNum         | Sequence number of the last hello message received from the peer. The range of values is 0 through 4,294,967,295. |

Table 388: show link-management Output Fields (*continued*)

| Field Name                 | Field Description                                                                                                                                                                                                                                                           |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Flags</b>               | Code that provides information about the control channel. Currently supports only code value <b>R</b> , which indicates that the control channel is restarting after a failure in the control plane, as when the Link Management Protocol (LMP) process starts or restarts. |
| <b>TE links</b>            | Traffic-engineered links that are managed by their peer.                                                                                                                                                                                                                    |
| <b>TE link name</b>        | Name of the traffic-engineered link.                                                                                                                                                                                                                                        |
| <b>State</b>               | State of the traffic-engineered link: <b>Up</b> , <b>Down</b> , or <b>Init</b> .                                                                                                                                                                                            |
| <b>Local identifier</b>    | Identifier of the local side of the link.                                                                                                                                                                                                                                   |
| <b>Remote identifier</b>   | Identifier of the remote side of the link.                                                                                                                                                                                                                                  |
| <b>Local address</b>       | Address of the local side of the link.                                                                                                                                                                                                                                      |
| <b>Remote address</b>      | Address of the remote side of the link.                                                                                                                                                                                                                                     |
| <b>Encoding</b>            | Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include <b>SDH/SONET</b> , <b>Ethernet</b> , <b>Packet</b> , and <b>PDH</b> .                                                                               |
| <b>Switching</b>           | Type of switching that can be performed on the traffic-engineered link. Supported values are <b>PSC-1</b> and <b>Packet</b> .                                                                                                                                               |
| <b>Minimum bandwidth</b>   | Smallest single allocation of bandwidth possible on the traffic-engineered link. This number is equal to the smallest bandwidth interface that is a member of the traffic-engineered link (in bps).                                                                         |
| <b>Maximum bandwidth</b>   | Largest single allocation of bandwidth possible on the traffic-engineered link. This number is equal to the largest bandwidth interface that is a member of the link (in bps).                                                                                              |
| <b>Total bandwidth</b>     | Sum of the bandwidth, in bits per second (bps) and megabits per second (Mbps), of all interfaces that are members of the link.                                                                                                                                              |
| <b>Available bandwidth</b> | Sum of the bandwidths of all interfaces that are members of the link and that are not yet allocated (in bps).                                                                                                                                                               |
| <b>Name</b>                | Name of the interface.                                                                                                                                                                                                                                                      |
| <b>State</b>               | State of the interface: <b>Up</b> or <b>Down</b> .                                                                                                                                                                                                                          |
| <b>Local ID</b>            | Identifier of the local side of the interface.                                                                                                                                                                                                                              |
| <b>Remote ID</b>           | Identifier of the remote side of the interface.                                                                                                                                                                                                                             |
| <b>Bandwidth</b>           | Bandwidth, in bps or Mbps, of the member interface.                                                                                                                                                                                                                         |
| <b>Used</b>                | Whether the resource is allocated to an LSP: <b>Yes</b> or <b>No</b> .                                                                                                                                                                                                      |



Table 388: show link-management Output Fields (*continued*)

| Field Name | Field Description |
|------------|-------------------|
| LSP-name   | LSP name.         |

## Sample Output

### show link-management

```
user@host> show link-management
Peer name: PEER-A, System identifier: 11973
State: Up, Control address: 10.255.245.4
 CC local ID CC remote ID State TxSeqNum RcvSeqNum Flags
 24547 24547 Up 1027 1026
TE links:
 pro4-ba

TE link name: pro4-ba, State: Init
Local identifier: 2662, Remote identifier: 0, Encoding: SDH/SONET, Switching:
PSC-1,
Minimum bandwidth: 155.52Mbps, Maximum bandwidth: 155.52Mbps, Total bandwidth:
155.52Mbps,
Available bandwidth: 155.52Mbps
 Name State Local ID Remote ID Bandwidth Used LSP-name
 so-1/0/2 Up 21271 0 155.52Mbps No
```

## show link-management peer

|                                 |                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show link-management peer</code><br><code>&lt;name <i>peer-name</i>&gt;</code>                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.5 for EX Series switches.                                                                                                                                                                                                   |
| <b>Description</b>              | Display Multiprotocol Label Switching (MPLS) peer link information.                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <b>none</b> —Display all peer link information.<br><br><b>name <i>peer-name</i></b> —(Optional) Display information for the specified peer only.                                                                                                                                                                        |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show link-management on page 5127</a></li> <li>• <a href="#">show link-management routing on page 5133</a></li> <li>• <a href="#">show link-management statistics on page 5136</a></li> <li>• <a href="#">show link-management te-link on page 5138</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show link-management peer on page 5132</a>                                                                                                                                                                                                                                                                  |
| <b>Output Fields</b>            | Table 389 describes the output fields for the <b>show link-management peer</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                         |

**Table 389: show link-management peer Output Fields**

| Field Name                 | Field Description                                                                                                                                                                                          |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Peer Name</b>           | Name of the peer.                                                                                                                                                                                          |
| <b>System identifier</b>   | Internal identifier for the peer. The range of values is 0 through 64,000.                                                                                                                                 |
| <b>State</b>               | State of the peer: <b>Up</b> or <b>Down</b> .                                                                                                                                                              |
| <b>Control address</b>     | Address to which a control channel is established.                                                                                                                                                         |
| <b>Hello interval</b>      | How often the routing device sends Link Management Protocol (LMP) hello packets.                                                                                                                           |
| <b>Hello dead interval</b> | How long LMP waits before declaring the control channel to be dead. This is an interval during which the routing device receives no LMP hello packets from the neighbor on a control that is active or up. |
| <b>CC local ID</b>         | Identifier assigned to the control channel by the local peer. The range of values is 1 through 4,294,967,296.                                                                                              |
| <b>CC remote ID</b>        | Identifier assigned to the control channel by the remote peer. The range of values is 1 through 4,294,967,296.                                                                                             |

Table 389: show link-management peer Output Fields (*continued*)

| Field Name       | Field Description                                                                                                                                                                                                                                                           |
|------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>State</b>     | State of the control channel: <b>Up</b> or <b>Down</b> .                                                                                                                                                                                                                    |
| <b>TxSeqNum</b>  | Sequence number of the hello message being sent to the peer. The range of values is <b>1</b> through <b>4,294,967,295</b> .                                                                                                                                                 |
| <b>RcvSeqNum</b> | Sequence number of the last hello message received from the peer. The range of values is <b>0</b> through <b>4,294,967,295</b> .                                                                                                                                            |
| <b>Flags</b>     | Code that provides information about the control channel. Currently supports only code value <b>R</b> , which indicates that the control channel is restarting after a failure in the control plane, as when the Link Management Protocol (LMP) process starts or restarts. |
| <b>TE links</b>  | Traffic-engineered links that are managed by their peer.                                                                                                                                                                                                                    |

## Sample Output

### show link-management peer

```

user@host> show link-management peer
Peer name: sonet, System identifier: 41448
State: Up, Control address: 70.70.70.70
Hello interval: 10000, Hello dead interval: 30000
 CC local ID CC remote ID State TxSeqNum RcvSeqNum Flags
 3265 0 ConfSnd 1 0 R
TE links:
to-sonet

```

## show link-management routing

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show link-management routing<br><peer <name <i>name</i> >   te-link <name <i>name</i> >><br><resource <name <i>name</i> >>                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.5 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Display Multiprotocol Label Switching (MPLS) peer or traffic engineering link information from the routing process.                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <p><b>none</b>—Display all peer and traffic-engineered link information.</p> <p><b>peer &lt;name <i>name</i>&gt;</b>—(Optional) Display information for all peers or for the specified peer only.</p> <p><b>resource &lt;name <i>name</i>&gt;</b>—(Optional) Display information for all resources or for the specified resource only.</p> <p><b>te-link &lt;name <i>name</i>&gt;</b>—(Optional) Display information for all traffic-engineered forwarding paths or for the specified path only.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show link-management on page 5127</a></li> <li>• <a href="#">show link-management peer on page 5131</a></li> <li>• <a href="#">show link-management statistics on page 5136</a></li> <li>• <a href="#">show link-management te-link on page 5138</a></li> </ul>                                                                                                                                                                                 |
| <b>List of Sample Output</b>    | <a href="#">show link-management routing on page 5135</a>                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Output Fields</b>            | Table 390 describes the output fields for the <b>show link-management routing</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                   |

**Table 390: show link-management routing Output Fields**

| Field Name        | Field Description                                                          |
|-------------------|----------------------------------------------------------------------------|
| Peer Name         | Name of the peer.                                                          |
| System identifier | Internal identifier for the peer. The range of values is 0 through 64,000. |
| State             | State of the peer: Up or Down.                                             |
| Control address   | Address to which a control channel is established.                         |
| Control channel   | Interface over which control packets are sent.                             |

Table 390: show link-management routing Output Fields (*continued*)

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>State</b>               | State of the control channel.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>TE link name</b>        | Traffic-engineered link name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>State</b>               | State of the traffic-engineered link: <b>Up</b> or <b>Down</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Local identifier</b>    | Identifier of the local side of the link.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Remote identifier</b>   | Identifier of the remote side of the link.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Local address</b>       | Address of the local side of the link.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Remote address</b>      | Address of the remote side of the link.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Encoding</b>            | Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include <b>SDH/SONET</b> , <b>Ethernet</b> , and <b>Packet</b> .                                                                                                                                                                                                                                                                                                                                                   |
| <b>Minimum bandwidth</b>   | Smallest single allocation of bandwidth, in bits per second (bps) or megabits per second (Mbps), possible on the traffic-engineered link. This number is equal to the smallest bandwidth interface that is a member of the traffic-engineered link.                                                                                                                                                                                                                                                                                |
| <b>Maximum bandwidth</b>   | Largest single allocation of bandwidth, in bps or Mbps, possible on the traffic-engineered link. This number is equal to the largest bandwidth interface that is a member of the link (in bps).                                                                                                                                                                                                                                                                                                                                    |
| <b>Total bandwidth</b>     | Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link.                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Available bandwidth</b> | Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link and that are not yet allocated.                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Resource</b>            | Forwarding adjacency LSP information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Type</b>                | Type of resource. The type is always a forwarding adjacency LSP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>State</b>               | State of the LSP: <b>Up</b> or <b>Down</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>System Identifier</b>   | Internal identifier for the peer. The range of values is <b>0</b> through <b>64,000</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Total bandwidth</b>     | Bandwidth resource, in bps or Mbps, on the TE-link learned from the routing process.                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Traffic parameters</b>  | <ul style="list-style-type: none"> <li>• <b>Encoding</b>—Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include <b>SDH/SONET</b>, <b>Ethernet</b>, and <b>Packet</b>.</li> <li>• <b>Switching</b>—Type of switching that can be performed on the traffic-engineered link: <b>PSC-1</b> and <b>Packet</b>.</li> <li>• <b>Granularity</b>—Layer 2 data for switching Layer 2 LSPs for this resource. Not supported. This value is always <b>unknown</b>.</li> </ul> |

## Sample Output

### show link-management routing

```

user@host> show link-management routing
Peer name: __rpd:fe-0/1/0.0, System identifier: 2147483649
State: Up, Control address: (null)
Control-channel State
fe-0/1/0.0 Active

Peer name: __rpd:fe-0/1/2.0, System identifier: 2147483650
State: Up, Control address: (null)
Control-channel State
fe-0/1/2.0 Active

Peer name: __rpd:so-0/2/0.0, System identifier: 2147483651
State: Down, Control address: (null)
Control-channel State
so-0/2/0.0 State

Peer name: __rpd:so-0/2/1.0, System identifier: 2147483652
State: Down, Control address: (null)
Control-channel State
so-0/2/1.0 State

...

TE link name: __rpd:fe-0/1/0.0, State: Up
Local identifier: 2147483649, Remote identifier: 0,
Local address: 192.168.37.66, Remote address: 192.168.37.66,
Encoding: Ethernet, Minimum bandwidth: 0bps, Maximum bandwidth: 100Mbps,
Total bandwidth: 100Mbps, Available bandwidth: 100Mbps

TE link name: __rpd:fe-0/1/2.0, State: Up
Local identifier: 2147483650, Remote identifier: 0,
Local address: 192.168.37.73, Remote address: 192.168.37.73,
Encoding: Ethernet, Minimum bandwidth: 0bps, Maximum bandwidth: 100Mbps,
Total bandwidth: 100Mbps, Available bandwidth: 100Mbps

TE link name: __rpd:so-0/2/0.0, State: Down
Local identifier: 2147483651, Remote identifier: 0,
Local address: 192.168.37.82, Remote address: 192.168.37.95,
Encoding: Ethernet, Minimum bandwidth: 0bps, Maximum bandwidth: 155.52Mbps,
Total bandwidth: 155.52Mbps, Available bandwidth: 155.52Mbps

...

Resource: falsp-bd, Type: LSP, State: Dn System identifier: 2147483652,
Total bandwidth: 0bps, Traffic parameters: Encoding: Packet, Switching: Packet,
Granularity: Unknown

Resource: falsp-be, Type: LSP, State: Up System identifier: 2147483654,
Total bandwidth: bw[1]=10Mbps, Traffic parameters: Encoding: Packet,
Switching: Packet, Granularity: Unknown

```

## show link-management statistics

|                                 |                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show link-management statistics<br><peer <name <i>name</i> >>                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Command introduced in Junos OS Release 8.0.<br>Command introduced in Junos OS Release 9.5 for EX Series switches.                                                                                                                                                                                                 |
| <b>Description</b>              | Display statistical information for Link Management Protocol (LMP) packets.                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <b>none</b> —Display information for all peers.<br><br><b>peer &lt;name <i>name</i>&gt;</b> —(Optional) Display information for all peers or for the specified peer only.                                                                                                                                         |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show link-management on page 5127</a></li> <li>• <a href="#">show link-management peer on page 5131</a></li> <li>• <a href="#">show link-management routing on page 5133</a></li> <li>• <a href="#">show link-management te-link on page 5138</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show link-management statistics on page 5137</a>                                                                                                                                                                                                                                                      |
| <b>Output Fields</b>            | <a href="#">Table 391</a> describes the output fields for the <b>show link-management statistics</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                             |

**Table 391: show link-management statistics Output Fields**

| Field Name                | Field Description                                                                                                                                                                                     |
|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Received packets          | Number of received packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.     |
| Received bad packets      | Number of received bad packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed. |
| Small packets             | Number of packets that are too small.                                                                                                                                                                 |
| Wrong protocol version    | Number of packets specifying the wrong LMP version.                                                                                                                                                   |
| Messages for unknown peer | Number of packets destined for an unknown peer.                                                                                                                                                       |
| Messages for bad state    | Number of packets indicating a state that does not match the recipient.                                                                                                                               |
| Stale acknowledgments     | Number of <b>configAck</b> and <b>LinkSummaryAck</b> packets received that have a stale message ID.                                                                                                   |



Table 391: show link-management statistics Output Fields (*continued*)

| Field Name                            | Field Description                                                                                                                                                                                                                                                                              |
|---------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Stale negative acknowledgments</b> | Number of <b>configNack</b> and <b>LinkSummaryNack</b> packets received that have a stale message ID.                                                                                                                                                                                          |
| <b>Sent packets</b>                   | Number of sent packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.                                                                                                  |
| <b>Retransmitted packets</b>          | Number of retransmitted packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.                                                                                         |
| <b>Dropped packets</b>                | Number of packets sent, by message type, that have been dropped by the receiver after the LMP retransmission interval has been exceeded. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed. |

## Sample Output

### show link-management statistics

```

user@host> show link-management statistics peer pro4-a
Statistics for peer pro4-a
 Received packets
 Config: 1
 Hello: 2572
 Small packets: 0
 Wrong protocol version: 0
 Messages for unknown peer: 0
 Messages for bad state: 0
 Stale acknowledgments: 0
 Stale negative acknowledgments: 0
 Sent packets
 Config: 2
 ConfigAck: 1
 Hello: 2572
 Retransmitted packets
 Config: 1

```

## show link-management te-link

|                                 |                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show link-management te-link<br><brief   detail><br><name <i>name</i> >                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.5 for EX Series switches.                                                                                                                                                                                                |
| <b>Description</b>              | Display the resources used to set up Multiprotocol Label Switching (MPLS) traffic-engineered forwarding paths.                                                                                                                                                                                                       |
| <b>Options</b>                  | <b>none</b> —Display information for all traffic-engineered links.<br><br><b>brief   detail</b> —(Optional) Display the specified level of output.<br><br><b>name <i>name</i></b> —(Optional) Display information for the specified traffic-engineered link only.                                                    |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show link-management on page 5127</a></li> <li>• <a href="#">show link-management peer on page 5131</a></li> <li>• <a href="#">show link-management routing on page 5133</a></li> <li>• <a href="#">show link-management statistics on page 5136</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show link-management te-link on page 5139</a>                                                                                                                                                                                                                                                            |
| <b>Output Fields</b>            | Table 392 describes the output fields for the <b>show link-management te-link</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                   |

**Table 392: show link-management te-link Output Fields**

| Field Name        | Field Description                                                                                                                                                                             |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| TE link name      | Traffic-engineered link name.                                                                                                                                                                 |
| State             | State of the traffic-engineered link: <b>Up</b> or <b>Down</b> .                                                                                                                              |
| Local identifier  | Identifier of the local side of the link.                                                                                                                                                     |
| Remote identifier | Identifier of the remote side of the link.                                                                                                                                                    |
| Local address     | Address of the local side of the link.                                                                                                                                                        |
| Remote address    | Address of the remote side of the link.                                                                                                                                                       |
| Encoding          | Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include <b>SDH/SONET</b> , <b>Ethernet</b> , <b>Packet</b> , and <b>PDH</b> . |

Table 392: show link-management te-link Output Fields (*continued*)

| Field Name                 | Field Description                                                                                                                                                                                                                                   |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Switching</b>           | Type of switching that can be performed on the traffic-engineered link. Supported values are <b>PSC-1</b> and <b>Packet</b> .                                                                                                                       |
| <b>Minimum bandwidth</b>   | Smallest single allocation of bandwidth, in bits per second (bps) or megabits per second (Mbps), possible on the traffic-engineered link. This number is equal to the smallest bandwidth interface that is a member of the traffic-engineered link. |
| <b>Maximum bandwidth</b>   | Largest single allocation of bandwidth, in bps or Mbps, possible on the traffic-engineered link. This number is equal to the largest bandwidth interface that is a member of the link.                                                              |
| <b>Total bandwidth</b>     | Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link (in bps).                                                                                                                                                      |
| <b>Available Bandwidth</b> | Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link and that are not yet allocated.                                                                                                                                |
| <b>Name</b>                | Name of the interface.                                                                                                                                                                                                                              |
| <b>State</b>               | State of the interface: <b>Up</b> or <b>Down</b> .                                                                                                                                                                                                  |
| <b>Local ID</b>            | Identifier of the local side of the interface.                                                                                                                                                                                                      |
| <b>Remote ID</b>           | Identifier of the remote side of the interface.                                                                                                                                                                                                     |
| <b>Bandwidth</b>           | Bandwidth, in bps or Mbps, of the member interface.                                                                                                                                                                                                 |
| <b>Used</b>                | Whether the resource is allocated to an LSP: <b>Yes</b> or <b>No</b> .                                                                                                                                                                              |
| <b>LSP-name</b>            | LSP name.                                                                                                                                                                                                                                           |

## Sample Output

### show link-management te-link

```

user@host> show link-management te-link
TE link name: FA-bd, State: Up
 Local identifier: 4144, Remote identifier: 0, Local address: 2.2.2.1,
 Remote address: 2.2.2.2, Encoding: Ethernet, Switching: Packet,
 Minimum bandwidth: 0bps, Maximum bandwidth: 0bps, Total bandwidth: 0bps,
 Available bandwidth: 0bps
 Name State Local ID Remote ID Bandwidth Used LSP-name
 falsp-bd Dn 43077 0 0bps No
TE link name: FA-be, State: Up
 Local identifier: 4145, Remote identifier: 0, Local address: 1.1.1.1,
 Remote address: 1.1.1.2, Encoding: Ethernet, Switching: Packet,
 Minimum bandwidth: 0bps, Maximum bandwidth: 10Mbps, Total bandwidth: 10Mbps,
 Available bandwidth: 8Mbps
 Name State Local ID Remote ID Bandwidth Used LSP-name
 falsp-be Up 43076 0 10Mbps Yes e2elasp-bf

```

## show mpls call-admission-control

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 5140</a><br><a href="#">Syntax (EX Series Switches) on page 5140</a>                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Syntax</b>                      | <pre>show mpls call-admission-control &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;lsp-name&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Syntax (EX Series Switches)</b> | <pre>show mpls call-admission-control &lt;lsp-name&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>         | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p><b>instance <i>instance-name</i></b> option added in Junos OS Release 15.1.</p>                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>                 | Display Multiprotocol Label Switching (MPLS) label-switched path (LSP) call admission control (CAC) information.                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                     | <p><b>none</b>—Display CAC information for all LSPs.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display MPLS LSP CAC information for the specified instance. If <i>instance-name</i> is omitted, MPLS LSP CAC information for the master instance is displayed.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b><i>lsp-name</i></b>—(Optional) Display CAC information for the specified LSP only.</p> |
| <b>Additional Information</b>      | The available bandwidth on an LSP path at a particular class type is the total path bandwidth at that class type minus the total bandwidth reserved by any Layer 2 connection at that class type.                                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>List of Sample Output</b>       | <a href="#">show mpls call-admission-control on page 5141</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Output Fields</b>               | <p><a href="#">Table 393</a> describes the output fields for the <b>show mpls call-admission-control</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                                                            |

**Table 393: show mpls call-admission-control Output Fields**

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Available bandwidth</b> | Current available bandwidth on each LSP path. Depending on whether the LSP is an E-LSP or a regular LSP, either per-class bandwidth or a single bandwidth value (corresponding to best-effort bandwidth at <b>ct0</b> ) is displayed. The available bandwidth on an LSP path at a particular class type is the total path bandwidth at that class type minus the total bandwidth reserved by some Layer 2 connections at that class type. |

Table 393: show mpls call-admission-control Output Fields (*continued*)

| Field Name         | Field Description                                                                                     |
|--------------------|-------------------------------------------------------------------------------------------------------|
| Layer2 connections | Different Layer 2 connections that had some bandwidth requirement and were admitted into an LSP path. |
| LSP name           | LSP pathname.                                                                                         |
| Neighbor address   | Neighbor address from which CAC and bandwidth booking are configured for Layer 2 circuits.            |
| Circuit            | Interface name and circuit information.                                                               |
| Primary            | LSP's primary standby path.                                                                           |
| Standby            | LSP's secondary standby path.                                                                         |
| VC bandwidth       | Bandwidth constraints associated with a Layer 2 circuit route.                                        |

## Sample Output

### show mpls call-admission-control

```

user@host# show mpls call-admission-control

LSP name: pro1-be
*Primary
 Available bandwidth: 0bps

LSP name: pro1-be-1
*Primary
 Available bandwidth: 60kbps

LSP name: pro1-be-gold
*Primary
 Available bandwidth: <ct0 50kbps> <ct1 20kbps> <ct2 30kbps> <ct3 0bps>
 Layer2 connections:
 Neighbor address: 10.255.245.215, Circuit: so-0/3/0.0(vc 5)
 VC bandwidth: <ct0 50kbps> <ct1 40kbps> <ct2 40kbps>

LSP name: pro1-be-gold-2
*Primary
 Available bandwidth: <ct0 0bps> <ct1 40kbps> <ct2 40kbps> <ct3 0bps>

LSP name: pro1-be-silver
*Primary prim1
 Available bandwidth: <ct0 10kbps> <ct1 20kbps> <ct2 0bps> <ct3 40kbps>
 Layer2 connections:
 Neighbor address: 10.255.245.215, Circuit: so-0/3/0.1(vc 3)
 VC bandwidth: <ct0 20kbps> <ct1 20kbps>
 Standby sec1
 Available bandwidth: <ct0 10kbps> <ct1 10kbps> <ct2 20kbps> <ct3 0bps>
 Layer2 connections:
 Neighbor address: 10.255.245.215, Circuit: so-0/3/0.1(vc 3)
 VC bandwidth: <ct0 20kbps> <ct1 20kbps>

```

## show mpls cspf

**List of Syntax** [Syntax on page 5142](#)  
[Syntax \(EX Series Switches\) on page 5142](#)

**Syntax** show mpls cspf  
 <instance *instance-name*>  
 <logical-system (all | *logical-system-name*)>

**Syntax (EX Series Switches)** show mpls cspf

**Release Information** Command introduced before Junos OS Release 7.4.  
 Command introduced in Junos OS Release 9.5 for EX Series switches.  
*instance instance-name* option added in Junos OS Release 15.1.

**Description** Display Multiprotocol Label Switching (MPLS) Constrained Shortest Path First (CSPF) statistics.

**Options** none—Display MPLS CSFP statistics.

*instance instance-name*—(Optional) Display MPLS CSPF information for the specified instance. If *instance-name* is omitted, MPLS CSPF information for the master instance is displayed.

*logical-system (all | logical-system-name)*—(Optional) Perform this operation on all logical systems or on a particular logical system.

**Required Privilege Level** view

**List of Sample Output** [show mpls cspf on page 5143](#)

**Output Fields** [Table 394](#) describes the output fields for the **show mpls cspf** command. Output fields are listed in the approximate order in which they appear.

**Table 394: show mpls cspf Output Fields**

| Field Name   | Field Description                                     |
|--------------|-------------------------------------------------------|
| Queue length | Number of LSPs queued for automatic path computation. |
| current      | Current queue length.                                 |
| maximum      | Maximum queue length (high-water mark).               |
| dequeued     | Number of aborted computation attempts.               |
| Paths        | Counters for label-switched path computations.        |
| total        | Sum of the next four fields.                          |

Table 394: show mpls cspf Output Fields (*continued*)

| Field Name          | Field Description                                                                          |
|---------------------|--------------------------------------------------------------------------------------------|
| <b>successful</b>   | Number of path computations that were successfully completed.                              |
| <b>no route</b>     | Number of path computations that failed because the destination is unreachable.            |
| <b>Sys Error</b>    | Number of path computations that failed because of lack of memory.                         |
| <b>CSPFs</b>        | Total number of CSPF computations. A single path might require multiple CSPF computations. |
| <b>Time</b>         | Time, in seconds, required to perform the label-switched path computation.                 |
| <b>Total</b>        | Total amount of time consumed by the CSPF path computation algorithm.                      |
| <b>CSPFs</b>        | Total number of CSPF computations.                                                         |
| <b>Avg per CSPF</b> | Average amount of time required for each CSPF computation.                                 |
| <b>% of rpd</b>     | Percentage of routing process CPU used in the CSPF computation.                            |

## Sample Output

### show mpls cspf

```

user@host> show mpls cspf
CSPF statistics
Queue length current maximum dequeued
 0 0 0
Paths total successful no route sys error CSPFs
 0 0 0 0 0
Time (secs) total CSPFs avg per CSPF % of rpd
 0.000000 0.000000 0.000000 0.0000

```

## show mpls diffserv-te

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 5144</a><br><a href="#">Syntax (EX Series Switches) on page 5144</a>                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Syntax</b>                      | <pre>show mpls diffserve-te &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                            |
| <b>Syntax (EX Series Switches)</b> | show mpls diffserve-te                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>         | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p><b>instance <i>instance-name</i></b> option added in Junos OS Release 15.1.</p>                                                                                                                                                                                                                                                                          |
| <b>Description</b>                 | Display Multiprotocol Label Switching (MPLS) label-switched path (LSP) Differentiated Services (DiffServ) class and preemption priority information.                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                     | <p><b>none</b>—Display DiffServ classes and priorities used by MPLS LSPs.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display DiffServ classes and priorities used by MPLS LSPs for the specified instance. If <b><i>instance-name</i></b> is omitted, DiffServ information for the master instance is displayed.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>List of Sample Output</b>       | <a href="#">show mpls diffserv-te on page 5145</a>                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Output Fields</b>               | Table 395 describes the output fields for the <b>show mpls diffserv-te</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                  |

**Table 395: show mpls diffserv-te Output Fields**

| Field Name             | Field Description                                                                                                                                                                                                                                                                       |
|------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Bandwidth model</b> | Bandwidth constraint model supported. The maximum allocation model (MAM) for EXP-inferred LSPs (E-LSPs) is currently supported.                                                                                                                                                         |
| <b>TE class</b>        | DiffServ traffic engineering class.                                                                                                                                                                                                                                                     |
| <b>Traffic class</b>   | <p>MPLS class type that corresponds to the DiffServ traffic engineering class:</p> <ul style="list-style-type: none"> <li>• <b>ct0</b>—Best effort</li> <li>• <b>ct1</b>—Assured forwarding</li> <li>• <b>ct2</b>—Expedited forwarding</li> <li>• <b>ct3</b>—Network control</li> </ul> |



Table 395: show mpls diffserv-te Output Fields (*continued*)

| Field Name | Field Description                                                                                                                                                                                  |
|------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Priority   | MPLS preemption priority for this class type, a value from 0 through 7. Interior gateway protocols (IGPs) distribute information about the available bandwidth for each traffic engineering class. |

## Sample Output

show mpls diffserv-te

```
user@host> show mpls diffserv-te
Bandwidth model: Maximum Allocation Model with support for E-LSPs.
TE class Traffic class Priority
te0 ct0 3
te1 ct1 2
```

## show route forwarding-table

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show route forwarding-table</code><br><code>&lt;detail   extensive   summary&gt;</code><br><code>&lt;ccc ccc-interface-name&gt;</code><br><code>&lt;destination&gt;</code><br><code>&lt;family family-name&gt;</code><br><code>&lt;label label&gt;</code><br><code>&lt;matching ip_prefix&gt;</code><br><code>&lt;multicast&gt;</code><br><code>&lt;vpn vpn&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.5 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Display the Routing Engine's forwarding table, including the network-layer prefixes and their next hops. This command is used to help verify that the routing protocol process has relayed the correction information to the forwarding table. The Routing Engine constructs and maintains one or more routing tables. From the routing tables, the Routing Engine derives a table of active routes, called the forwarding table.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <p><b>none</b>—Display the routes in the forwarding table.</p> <p><b>detail   extensive   summary</b>—(Optional) Display the specified level of output.</p> <p><b>ccc</b>—(Optional) Display the specified circuit cross-connect interface name for entries to match.</p> <p><b>destination</b>—(Optional) Display the destination prefix.</p> <p><b>family family-name</b>—(Optional) Display routing table entries for the specified family: <b>ethernet-switching</b>, <b>inet</b>, <b>inet6</b>, <b>iso</b>, <b>mpls</b>, <b>vlan classification</b>.</p> <p><b>label label</b>—(Optional) Display route entries for the specified label name.</p> <p><b>matching ip_prefix</b>—(Optional) Display route entries for the specified IP prefix.</p> <p><b>multicast</b>—(Optional) Display route entries for multicast routes.</p> <p><b>vpn vpn</b>—(Optional) Display route entries for the specified VPN.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring MPLS on EX8200 and EX4500 Switches</i></li><li>• <i>Configuring MPLS on EX8200 and EX4500 Provider Switches (CLI Procedure)</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>List of Sample Output</b>    | <p><a href="#">show route forwarding-table on page 5148</a></p> <p><a href="#">show route forwarding-table summary on page 5149</a></p> <p><a href="#">show route forwarding-table extensive on page 5149</a></p> <p><a href="#">show route forwarding-table ccc on page 5151</a></p> <p><a href="#">show route forwarding-table family (MPLS) on page 5151</a></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

[show route forwarding-table family \(IPv6\) on page 5151](#)

[show route forwarding-table label on page 5152](#)

[show route forwarding-table matching on page 5152](#)

[show route forwarding-table multicast on page 5152](#)

**Output Fields** Table 306 lists the output fields for the **show route forwarding-table** command. Output fields are listed in the approximate order in which they appear. Field names might be abbreviated (as shown in parentheses) when no level of output is specified or when the **detail** keyword is used instead of the **extensive** keyword.

**Table 396: show route forwarding-table Output Fields**

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Level of Output                  |
|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------|
| <b>Routing table</b>           | Name of the routing table (for example, <b>inet</b> , <b>inet6</b> , <b>mpls</b> ).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | All levels                       |
| <b>Address family</b>          | Address family (for example, <b>IP</b> , <b>IPv6</b> , <b>ISO</b> , <b>MPLS</b> ).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | All levels                       |
| <b>Destination</b>             | Destination of the route.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail</b> , <b>extensive</b> |
| <b>Route Type (Type)</b>       | How the route was placed into the forwarding table. When the <b>detail</b> keyword is used, the route type might be abbreviated (as shown in parentheses): <ul style="list-style-type: none"> <li><b>cloned (clon)</b>—(TCP or multicast only) Cloned route.</li> <li><b>destination (dest)</b>—Remote addresses directly reachable through an interface.</li> <li><b>destination down (iddn)</b>—Destination route for which the interface is unreachable.</li> <li><b>interface cloned (ifcl)</b>—Cloned route for which the interface is unreachable.</li> <li><b>route down (ifdn)</b>—Interface route for which the interface is unreachable.</li> <li><b>ignore (ignr)</b>—Ignore this route.</li> <li><b>interface (intf)</b>—Installed as a result of configuring an interface.</li> <li><b>permanent (perm)</b>—Routes installed by the kernel when the routing table is initialized.</li> <li><b>user</b>—Routes installed by the routing protocol process or as a result of the configuration.</li> </ul> | All levels                       |
| <b>Route reference (RtRef)</b> | Number of routes to reference.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail</b> , <b>extensive</b> |
| <b>Flags</b>                   | Route type flags: <ul style="list-style-type: none"> <li><b>none</b>—No flags are enabled.</li> <li><b>accounting</b>—Route has accounting enabled.</li> <li><b>cached</b>—Cache route.</li> <li><b>incoming-iface interface-number</b>—Check against incoming interface.</li> <li><b>prefix load balance</b>—Load balancing is enabled for this prefix.</li> <li><b>sent to PFE</b>—Route has been sent to the Packet Forwarding Engine.</li> <li><b>static</b>—Static route.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>extensive</b>                 |
| <b>Nexthop</b>                 | IP address of the next hop to the destination.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail</b> , <b>extensive</b> |

Table 396: show route forwarding-table Output Fields (*continued*)

| Field Name                        | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Level of Output               |
|-----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------|
| <b>Next hop type (Type)</b>       | <p>Next-hop type. When the <b>detail</b> keyword is used, the next-hop type might be abbreviated (as indicated in parentheses):</p> <ul style="list-style-type: none"> <li>• <b>broadcast (bcst)</b>—Broadcast.</li> <li>• <b>deny</b>—Deny.</li> <li>• <b>hold</b>—Next hop is waiting to be resolved into a unicast or multicast type.</li> <li>• <b>indexed (idxd)</b>—Indexed next hop.</li> <li>• <b>indirect (indr)</b>—Indirect next hop.</li> <li>• <b>local (locl)</b>—Local address on an interface.</li> <li>• <b>routed multicast (mcr)</b>—Regular multicast next hop</li> <li>• <b>multicast (mcst)</b>—Wire multicast next hop (limited to the LAN).</li> <li>• <b>multicast discard (mdsc)</b>—Multicast discard.</li> <li>• <b>multicast group (mgrp)</b> —Multicast group member.</li> <li>• <b>receive (rcv)</b>—Receive.</li> <li>• <b>reject (rjct)</b>—Discard. An ICMP unreachable message was sent.</li> <li>• <b>resolve (rslv)</b>—Resolving the next hop.</li> <li>• <b>unicast (ucst)</b>—Unicast.</li> <li>• <b>unilist (ulst)</b>—List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.</li> </ul> | <b>detail, extensive</b>      |
| <b>Index</b>                      | Software index of the next hop that is used to route the traffic for a given prefix.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>detail, extensive none</b> |
| <b>Route interface-index</b>      | Logical interface index from which the route is learned. For example, for interface routes, this is the logical interface index of the route itself. For static routes, this field is zero. For routes learned through routing protocols, this is the logical interface index from which the route is learned.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>extensive</b>              |
| <b>Reference (NhRef)</b>          | Number of routes that refer to this next hop.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>none detail, extensive</b> |
| <b>Next-hop interface (Netif)</b> | Interface used to reach the next hop.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>none detail, extensive</b> |
| <b>Alternate forward nh index</b> | Index number of the alternate next hop interface. Seen with <b>multicast</b> option only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>extensive</b>              |
| <b>Next-hop L3 Interface</b>      | The next hop layer 3 interface. This option can be expressed as a VLAN name and is only seen with the <b>multicast</b> option.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>extensive</b>              |
| <b>Next-hop L2 Interfaces</b>     | The next hop layer 2 interfaces. Seen with <b>multicast</b> option only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>extensive</b>              |

## Sample Output

### show route forwarding-table

```

user@switch> show route forwarding-table

Routing table: default.inet

```

```

Internet:
Destination Type RtRef Next hop Type Index NhRef Netif
default user 2 0:12:f2:21:cf:0 ucst 333 5 me0.0
default perm 0 rjct 36 2
0.0.0.0/32 perm 0 dscd 34 1
2.2.2.0/24 intf 0 rslv 1309 1 ae0.0
2.2.2.0/32 dest 0 2.2.2.0 recv 1307 1 ae0.0
2.2.2.1/32 dest 0 0:21:59:cc:89:c0 ucst 1320 1 ae0.0
2.2.2.2/32 intf 0 2.2.2.2 locl 1308 2
2.2.2.2/32 dest 0 2.2.2.2 locl 1308 2
2.2.2.255/32 dest 0 2.2.2.255 bcst 1306 1 ae0.0
3.3.3.0/24 intf 0 rslv 1313 1 ae1.0
3.3.3.0/32 dest 0 3.3.3.0 recv 1311 1 ae1.0
3.3.3.1/32 intf 0 3.3.3.1 locl 1312 2
3.3.3.1/32 dest 0 3.3.3.1 locl 1312 2
3.3.3.2/32 dest 0 0:21:59:cc:89:c1 ucst 1321 24 ae1.0
3.3.3.255/32 dest 0 3.3.3.255 bcst 1310 1 ae1.0
4.4.4.0/24 user 0 3.3.3.2 ucst 1321 24 ae1.0
8.8.8.8/32 user 0 3.3.3.2 ucst 1321 24 ae1.0
9.9.9.9/32 intf 0 9.9.9.9 locl 1280 1
10.10.10.10/32 user 0 3.3.3.2 ucst 1321 24 ae1.0
10.93.8.0/21 intf 0 rslv 323 1 me0.0
10.93.8.0/32 dest 0 10.93.8.0 recv 321 1 me0.0
10.93.13.238/32 intf 0 10.93.13.238 locl 322 2
10.93.13.238/32 dest 0 10.93.13.238 locl 322 2
10.93.15.254/32 dest 0 0:12:f2:21:cf:0 ucst 333 5 me0.0
10.93.15.255/32 dest 0 10.93.15.255 bcst 320 1 me0.0
14.14.14.0/24 ifdn 0 rslv 1319 1 ge-0/0/25.0
14.14.14.0/32 iddn 0 14.14.14.0 recv 1317 1 ge-0/0/25.0
14.14.14.2/32 user 0 rjct 36 2
14.14.14.2/32 intf 0 14.14.14.2 locl 1318 2
14.14.14.2/32 iddn 0 14.14.14.2 locl 1318 2
14.14.14.255/32 iddn 0 14.14.14.255 bcst 1316 1 ge-0/0/25.0
224.0.0.0/4 perm 1 mdsc 35 1
224.0.0.1/32 perm 0 224.0.0.1 mcst 31 3
224.0.0.5/32 user 1 224.0.0.5 mcst 31 3
255.255.255.255/32 perm 0 bcst 32 1

```

### show route forwarding-table summary

```
user@switch> show route forwarding-table summary
```

```

Routing table: default.inet
Internet:
 user: 6 routes
 perm: 5 routes
 intf: 8 routes
 dest: 12 routes
 ifdn: 1 routes
 iddn: 3 routes

```

### show route forwarding-table extensive

```
user@switch> show route forwarding-table summary
```

```

Routing table: default.inet [Index 0]
Internet:

Destination: default
Route type: user
Route reference: 2
Route interface-index: 0

```

|                                     |                           |              |
|-------------------------------------|---------------------------|--------------|
| Flags: sent to PFE, rt nh decoupled |                           |              |
| Nexthop: 0:12:f2:21:cf:0            |                           |              |
| Next-hop type: unicast              | Index: 333                | Reference: 5 |
| Next-hop interface: me0.0           |                           |              |
| Destination: default                |                           |              |
| Route type: permanent               |                           |              |
| Route reference: 0                  | Route interface-index: 0  |              |
| Flags: none                         |                           |              |
| Next-hop type: reject               | Index: 36                 | Reference: 2 |
| Destination: 0.0.0.0/32             |                           |              |
| Route type: permanent               |                           |              |
| Route reference: 0                  | Route interface-index: 0  |              |
| Flags: sent to PFE                  |                           |              |
| Next-hop type: discard              | Index: 34                 | Reference: 1 |
| Destination: 2.2.2.0/24             |                           |              |
| Route type: interface               |                           |              |
| Route reference: 0                  | Route interface-index: 66 |              |
| Flags: sent to PFE                  |                           |              |
| Next-hop type: resolve              | Index: 1309               | Reference: 1 |
| Next-hop interface: ae0.0           |                           |              |
| Destination: 2.2.2.0/32             |                           |              |
| Route type: destination             |                           |              |
| Route reference: 0                  | Route interface-index: 66 |              |
| Flags: sent to PFE                  |                           |              |
| Nexthop: 2.2.2.0                    |                           |              |
| Next-hop type: receive              | Index: 1307               | Reference: 1 |
| Next-hop interface: ae0.0           |                           |              |
| Destination: 2.2.2.1/32             |                           |              |
| Route type: destination             |                           |              |
| Route reference: 0                  | Route interface-index: 66 |              |
| Flags: sent to PFE                  |                           |              |
| Nexthop: 0:21:59:cc:89:c0           |                           |              |
| Next-hop type: unicast              | Index: 1320               | Reference: 1 |
| Next-hop interface: ae0.0           |                           |              |
| Destination: 2.2.2.2/32             |                           |              |
| Route type: interface               |                           |              |
| Route reference: 0                  | Route interface-index: 0  |              |
| Flags: sent to PFE                  |                           |              |
| Nexthop: 2.2.2.2                    |                           |              |
| Next-hop type: local                | Index: 1308               | Reference: 2 |
| Destination: 2.2.2.2/32             |                           |              |
| Route type: destination             |                           |              |
| Route reference: 0                  | Route interface-index: 66 |              |
| Flags: none                         |                           |              |
| Nexthop: 2.2.2.2                    |                           |              |
| Next-hop type: local                | Index: 1308               | Reference: 2 |
| Destination: 2.2.2.255/32           |                           |              |
| Route type: destination             |                           |              |
| Route reference: 0                  | Route interface-index: 66 |              |
| Flags: sent to PFE                  |                           |              |
| Nexthop: 2.2.2.255                  |                           |              |
| Next-hop type: broadcast            | Index: 1306               | Reference: 1 |
| Next-hop interface: ae0.0           |                           |              |

**show route forwarding-table ccc**

```

user@switch> show route forwarding-table ccc ge-0/0/0.10
Routing table: default.mpls
MPLS:
Destination Type RtRef Next hop Type Index NhRef Netif
ge-0/0/0.10 (CCC) user 0 3.3.3.2 Push 300112 1343 2 ae1.0

```

**show route forwarding-table family (MPLS)**

```

user@switch> show route forwarding-table family mpls

Routing table: default.mpls
MPLS:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0
0 user 0 recv 49 3
1 user 0 recv 49 3
2 user 0 recv 49 3
299776 user 0 Pop 1334 2 ge-0/0/0.10
299792 user 0 Pop 1339 2 ge-0/0/0.14
299808 user 0 Pop 1341 2 ge-0/0/0.2
299824 user 0 Pop 1344 2 ge-0/0/0.11
299840 user 0 Pop 1345 2 ge-0/0/0.13
299856 user 0 Pop 1346 2 ge-0/0/0.18
299872 user 0 Pop 1347 2 ge-0/0/0.16
299888 user 0 Pop 1348 2 ge-0/0/0.7
299904 user 0 Pop 1349 2 ge-0/0/0.20
299920 user 0 Pop 1350 2 ge-0/0/0.19
299936 user 0 Pop 1351 2 ge-0/0/0.17
299952 user 0 Pop 1352 2 ge-0/0/0.9
299968 user 0 Pop 1353 2 ge-0/0/0.1
299984 user 0 Pop 1354 2 ge-0/0/0.12
300000 user 0 Pop 1355 2 ge-0/0/0.8
300016 user 0 Pop 1356 2 ge-0/0/0.4
300032 user 0 Pop 1357 2 ge-0/0/0.5
300048 user 0 Pop 1358 2 ge-0/0/0.3
300064 user 0 Pop 1359 2 ge-0/0/0.15
ge-0/0/0.1 (CCC) user 0 3.3.3.2 Push 300064 1340 2 ae1.0
ge-0/0/0.2 (CCC) user 0 3.3.3.2 Push 299872 1328 2 ae1.0
ge-0/0/0.3 (CCC) user 0 3.3.3.2 Push 299792 1323 2 ae1.0
ge-0/0/0.4 (CCC) user 0 3.3.3.2 Push 300016 1337 2 ae1.0
ge-0/0/0.5 (CCC) user 0 3.3.3.2 Push 299824 1325 2 ae1.0
ge-0/0/0.7 (CCC) user 0 3.3.3.2 Push 299920 1331 2 ae1.0
ge-0/0/0.8 (CCC) user 0 3.3.3.2 Push 299840 1326 2 ae1.0
ge-0/0/0.9 (CCC) user 0 3.3.3.2 Push 299888 1329 2 ae1.0
ge-0/0/0.10 (CCC) user 0 3.3.3.2 Push 300112 1343 2 ae1.0
ge-0/0/0.11 (CCC) user 0 3.3.3.2 Push 299776 1322 2 ae1.0
ge-0/0/0.12 (CCC) user 0 3.3.3.2 Push 299952 1333 2 ae1.0
ge-0/0/0.13 (CCC) user 0 3.3.3.2 Push 300096 1342 2 ae1.0
ge-0/0/0.14 (CCC) user 0 3.3.3.2 Push 299984 1335 2 ae1.0
ge-0/0/0.15 (CCC) user 0 3.3.3.2 Push 299936 1332 2 ae1.0
ge-0/0/0.16 (CCC) user 0 3.3.3.2 Push 299808 1324 2 ae1.0
ge-0/0/0.17 (CCC) user 0 3.3.3.2 Push 300000 1336 2 ae1.0
ge-0/0/0.18 (CCC) user 0 3.3.3.2 Push 300032 1338 2 ae1.0
ge-0/0/0.19 (CCC) user 0 3.3.3.2 Push 299904 1330 2 ae1.0
ge-0/0/0.20 (CCC) user 0 3.3.3.2 Push 299856 1327 2 ae1.0

```

**show route forwarding-table family (IPv6)**

```

user@switch> show route forwarding-table family inet6

```

```

Routing table: default.inet6
Internet6:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0 rjct 44 1
::/128 perm 0 dscd 42 1
ff00::/8 perm 0 mdsc 43 1
ff02::1/128 perm 0 ff02::1 mcst 39 1

```

```

Routing table: default-switch.inet6
Internet6:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0 rjct 530 1
::/128 perm 0 dscd 528 1
2:1::3a00/312 user 0 indr 131070 2
comp 572 1
2:1::3a82/320 user 0 indr 131071 3
comp 573 1
2:1::3af0/320 user 0 indr 131071 3
comp 573 1
2:1:0:ff00::/56 user 0 mdsc 529 2
ff00::/8 perm 0 mdsc 529 2
ff02::1/128 perm 0 ff02::1 mcst 526 1

```

```

Routing table: __master.anon__.inet6
Internet6:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0 rjct 554 1
::/128 perm 0 dscd 552 1
ff00::/8 perm 0 mdsc 553 1
ff02::1/128 perm 0 ff02::1 mcst 550 1

```

### show route forwarding-table label

```
user@switch> show route forwarding-table label 29976
```

```

Routing table: default.mpls
MPLS:
Destination Type RtRef Next hop Type Index NhRef Netif
299776 user 0 Pop 1334 2 ge-0/0/0.10

```

### show route forwarding-table matching

```
user@switch> show route forwarding-table matching 3
```

```

Routing table: default.inet
Internet:

```

### show route forwarding-table multicast

```
user@switch> show route forwarding-table multicast
```

```

Routing table: default.inet
Internet:
Destination Type RtRef Next hop Type Index NhRef Netif
224.0.0.0/4 perm 1 mdsc 35 1
224.0.0.1/32 perm 0 224.0.0.1 mcst 31 3
224.0.0.5/32 user 1 224.0.0.5 mcst 31 3

```

```

Routing table: __master.anon__.inet
Internet:
Destination Type RtRef Next hop Type Index NhRef Netif
224.0.0.0/4 perm 0 mdsc 1289 1

```



```
224.0.0.1/32 perm 0 224.0.0.1 mcst 1285 1
```

```
Routing table: default.inet6
```

```
Internet6:
```

| Destination | Type | RtRef | Next hop | Type | Index | NhRef | Netif |
|-------------|------|-------|----------|------|-------|-------|-------|
| ff00::/8    | perm | 0     |          | mdsc | 43    | 1     |       |
| ff02::1/128 | perm | 0     | ff02::1  | mcst | 39    | 1     |       |

## show mpls interface

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 5154</a><br><a href="#">Syntax (EX Series Switches) on page 5154</a>                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Syntax</b>                      | <pre>show mpls interface &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Syntax (EX Series Switches)</b> | show mpls interface                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>         | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p><b>instance <i>instance-name</i></b> option added in Junos OS Release 15.1.</p>                                                                                                                                                                                                                                                                                |
| <b>Description</b>                 | Display information about Multiprotocol Label Switching (MPLS)-enabled interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                     | <p><b>none</b>—Display information about MPLS-enabled interfaces.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information about MPLS-enabled interfaces for the specified routing instance. If <b>instance-name</b> is omitted, information about MPLS-enabled interfaces is displayed for the master instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Additional Information</b>      | MPLS is enabled on an interface when the interface is configured with both the <b>set protocol mpls interface <i>interface-name</i></b> and <b>set interface <i>interface-name</i> unit 0 family mpls</b> statements.                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>List of Sample Output</b>       | <a href="#">show mpls interface on page 5155</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Output Fields</b>               | <a href="#">Table 397</a> describes the output fields for the <b>show mpls interface</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                          |

**Table 397: show mpls interface Output Fields**

| Field Name                   | Field Description                                      |
|------------------------------|--------------------------------------------------------|
| <b>Interface</b>             | Name of the interface.                                 |
| <b>State</b>                 | State of the interface: <b>Up</b> or <b>Dn</b> (down). |
| <b>Administrative groups</b> | Administratively assigned colors of the link.          |

Table 397: show mpls interface Output Fields (*continued*)

| Field Name                            | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Maximum labels                        | Maximum number of MPLS labels upon which MPLS can operate on a logical interface. This is configured using the <b>maximum-labels</b> statement at the [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family mpls] or the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family mpls] hierarchy levels.                                                                                                                                                                                                            |
| Static protection revert time         | Time (in seconds) that a static LSP must wait before traffic reverts from the bypass path to the original path. This is configured using the <b>protection-revert-time</b> statement at the [edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i> static] or the [edit protocols mpls interface <i>interface-name</i> static] hierarchy levels.                                                                                                                                                                                                                    |
| Always mark connection protection tlv | Enabled or Disabled: Enabled indicates that the <b>always-mark-connection-protection-tlv</b> statement is configured at the [edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i> static] or the [edit protocols mpls interface <i>interface-name</i> static] hierarchy levels. When this statement is configured, it marks all OAM traffic transiting this interface in preparation for switching the traffic to an alternate path based on the OAM functionality. To switch traffic to the bypass LSP, the <b>switch-away-lsps</b> statement must be configured. |
| Switch away lsps                      | Enabled or Disabled: Enabled indicates that the <b>switch-away-lsps</b> statement is configured at the [edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i> static] or the [edit protocols mpls interface <i>interface-name</i> static] hierarchy levels. This enables you to switch an LSP away from a network node using a bypass LSP. This feature can be used in maintenance of active networks when a network device needs to be replaced without interrupting traffic passing through the network. The LSPs can be either static or dynamic.                |

## Sample Output

### show mpls interface

```
user@host> show mpls interface
```

```
Interface: ge-0/2/1.57
State: Up
Administrative group: <none>
Maximum labels: 5
Static protection revert time: 5 seconds
Always mark connection protection tlv: Disabled
Switch away lsps : Disabled
```

## show link-management statistics

|                                 |                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show link-management statistics<br><peer <name <i>name</i> >>                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Command introduced in Junos OS Release 8.0.<br>Command introduced in Junos OS Release 9.5 for EX Series switches.                                                                                                                                                                                                 |
| <b>Description</b>              | Display statistical information for Link Management Protocol (LMP) packets.                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <b>none</b> —Display information for all peers.<br><br><b>peer &lt;name <i>name</i>&gt;</b> —(Optional) Display information for all peers or for the specified peer only.                                                                                                                                         |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show link-management on page 5127</a></li> <li>• <a href="#">show link-management peer on page 5131</a></li> <li>• <a href="#">show link-management routing on page 5133</a></li> <li>• <a href="#">show link-management te-link on page 5138</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show link-management statistics on page 5157</a>                                                                                                                                                                                                                                                      |
| <b>Output Fields</b>            | <a href="#">Table 391</a> describes the output fields for the <b>show link-management statistics</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                             |

**Table 398: show link-management statistics Output Fields**

| Field Name                       | Field Description                                                                                                                                                                                     |
|----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Received packets</b>          | Number of received packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.     |
| <b>Received bad packets</b>      | Number of received bad packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed. |
| <b>Small packets</b>             | Number of packets that are too small.                                                                                                                                                                 |
| <b>Wrong protocol version</b>    | Number of packets specifying the wrong LMP version.                                                                                                                                                   |
| <b>Messages for unknown peer</b> | Number of packets destined for an unknown peer.                                                                                                                                                       |
| <b>Messages for bad state</b>    | Number of packets indicating a state that does not match the recipient.                                                                                                                               |
| <b>Stale acknowledgments</b>     | Number of <b>configAck</b> and <b>LinkSummaryAck</b> packets received that have a stale message ID.                                                                                                   |

Table 398: show link-management statistics Output Fields (*continued*)

| Field Name                            | Field Description                                                                                                                                                                                                                                                                              |
|---------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Stale negative acknowledgments</b> | Number of <b>configNack</b> and <b>LinkSummaryNack</b> packets received that have a stale message ID.                                                                                                                                                                                          |
| <b>Sent packets</b>                   | Number of sent packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.                                                                                                  |
| <b>Retransmitted packets</b>          | Number of retransmitted packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.                                                                                         |
| <b>Dropped packets</b>                | Number of packets sent, by message type, that have been dropped by the receiver after the LMP retransmission interval has been exceeded. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed. |

## Sample Output

### show link-management statistics

```

user@host> show link-management statistics peer pro4-a
Statistics for peer pro4-a
 Received packets
 Config: 1
 Hello: 2572
 Small packets: 0
 Wrong protocol version: 0
 Messages for unknown peer: 0
 Messages for bad state: 0
 Stale acknowledgments: 0
 Stale negative acknowledgments: 0
 Sent packets
 Config: 2
 ConfigAck: 1
 Hello: 2572
 Retransmitted packets
 Config: 1

```

## show link-management te-link

|                                 |                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show link-management te-link<br><brief   detail><br><name <i>name</i> >                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.5 for EX Series switches.                                                                                                                                                                                                |
| <b>Description</b>              | Display the resources used to set up Multiprotocol Label Switching (MPLS) traffic-engineered forwarding paths.                                                                                                                                                                                                       |
| <b>Options</b>                  | <b>none</b> —Display information for all traffic-engineered links.<br><br><b>brief   detail</b> —(Optional) Display the specified level of output.<br><br><b>name <i>name</i></b> —(Optional) Display information for the specified traffic-engineered link only.                                                    |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show link-management on page 5127</a></li> <li>• <a href="#">show link-management peer on page 5131</a></li> <li>• <a href="#">show link-management routing on page 5133</a></li> <li>• <a href="#">show link-management statistics on page 5136</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show link-management te-link on page 5159</a>                                                                                                                                                                                                                                                            |
| <b>Output Fields</b>            | Table 392 describes the output fields for the <b>show link-management te-link</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                   |

Table 399: show link-management te-link Output Fields

| Field Name        | Field Description                                                                                                                                                                             |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| TE link name      | Traffic-engineered link name.                                                                                                                                                                 |
| State             | State of the traffic-engineered link: <b>Up</b> or <b>Down</b> .                                                                                                                              |
| Local identifier  | Identifier of the local side of the link.                                                                                                                                                     |
| Remote identifier | Identifier of the remote side of the link.                                                                                                                                                    |
| Local address     | Address of the local side of the link.                                                                                                                                                        |
| Remote address    | Address of the remote side of the link.                                                                                                                                                       |
| Encoding          | Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include <b>SDH/SONET</b> , <b>Ethernet</b> , <b>Packet</b> , and <b>PDH</b> . |

Table 399: show link-management te-link Output Fields (*continued*)

| Field Name                 | Field Description                                                                                                                                                                                                                                   |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Switching</b>           | Type of switching that can be performed on the traffic-engineered link. Supported values are <b>PSC-1</b> and <b>Packet</b> .                                                                                                                       |
| <b>Minimum bandwidth</b>   | Smallest single allocation of bandwidth, in bits per second (bps) or megabits per second (Mbps), possible on the traffic-engineered link. This number is equal to the smallest bandwidth interface that is a member of the traffic-engineered link. |
| <b>Maximum bandwidth</b>   | Largest single allocation of bandwidth, in bps or Mbps, possible on the traffic-engineered link. This number is equal to the largest bandwidth interface that is a member of the link.                                                              |
| <b>Total bandwidth</b>     | Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link (in bps).                                                                                                                                                      |
| <b>Available Bandwidth</b> | Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link and that are not yet allocated.                                                                                                                                |
| <b>Name</b>                | Name of the interface.                                                                                                                                                                                                                              |
| <b>State</b>               | State of the interface: <b>Up</b> or <b>Down</b> .                                                                                                                                                                                                  |
| <b>Local ID</b>            | Identifier of the local side of the interface.                                                                                                                                                                                                      |
| <b>Remote ID</b>           | Identifier of the remote side of the interface.                                                                                                                                                                                                     |
| <b>Bandwidth</b>           | Bandwidth, in bps or Mbps, of the member interface.                                                                                                                                                                                                 |
| <b>Used</b>                | Whether the resource is allocated to an LSP: <b>Yes</b> or <b>No</b> .                                                                                                                                                                              |
| <b>LSP-name</b>            | LSP name.                                                                                                                                                                                                                                           |

## Sample Output

### show link-management te-link

```

user@host> show link-management te-link
TE link name: FA-bd, State: Up
 Local identifier: 4144, Remote identifier: 0, Local address: 2.2.2.1,
 Remote address: 2.2.2.2, Encoding: Ethernet, Switching: Packet,
 Minimum bandwidth: 0bps, Maximum bandwidth: 0bps, Total bandwidth: 0bps,
 Available bandwidth: 0bps
 Name State Local ID Remote ID Bandwidth Used LSP-name
 falsp-bd Dn 43077 0 0bps No
TE link name: FA-be, State: Up
 Local identifier: 4145, Remote identifier: 0, Local address: 1.1.1.1,
 Remote address: 1.1.1.2, Encoding: Ethernet, Switching: Packet,
 Minimum bandwidth: 0bps, Maximum bandwidth: 10Mbps, Total bandwidth: 10Mbps,
 Available bandwidth: 8Mbps
 Name State Local ID Remote ID Bandwidth Used LSP-name
 falsp-be Up 43076 0 10Mbps Yes e2elasp-bf

```

## show mpls call-admission-control

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 5160</a><br><a href="#">Syntax (EX Series Switches) on page 5160</a>                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Syntax</b>                      | <pre>show mpls call-admission-control &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;lsp-name&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Syntax (EX Series Switches)</b> | <pre>show mpls call-admission-control &lt;lsp-name&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>         | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p><b>instance <i>instance-name</i></b> option added in Junos OS Release 15.1.</p>                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>                 | Display Multiprotocol Label Switching (MPLS) label-switched path (LSP) call admission control (CAC) information.                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                     | <p><b>none</b>—Display CAC information for all LSPs.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display MPLS LSP CAC information for the specified instance. If <b>instance-name</b> is omitted, MPLS LSP CAC information for the master instance is displayed.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>lsp-name</b>—(Optional) Display CAC information for the specified LSP only.</p> |
| <b>Additional Information</b>      | The available bandwidth on an LSP path at a particular class type is the total path bandwidth at that class type minus the total bandwidth reserved by any Layer 2 connection at that class type.                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>List of Sample Output</b>       | <a href="#">show mpls call-admission-control on page 5161</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Output Fields</b>               | <p><a href="#">Table 393</a> describes the output fields for the <b>show mpls call-admission-control</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                                                     |

**Table 400: show mpls call-admission-control Output Fields**

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Available bandwidth</b> | Current available bandwidth on each LSP path. Depending on whether the LSP is an E-LSP or a regular LSP, either per-class bandwidth or a single bandwidth value (corresponding to best-effort bandwidth at <b>ct0</b> ) is displayed. The available bandwidth on an LSP path at a particular class type is the total path bandwidth at that class type minus the total bandwidth reserved by some Layer 2 connections at that class type. |



Table 400: show mpls call-admission-control Output Fields (*continued*)

| Field Name         | Field Description                                                                                     |
|--------------------|-------------------------------------------------------------------------------------------------------|
| Layer2 connections | Different Layer 2 connections that had some bandwidth requirement and were admitted into an LSP path. |
| LSP name           | LSP pathname.                                                                                         |
| Neighbor address   | Neighbor address from which CAC and bandwidth booking are configured for Layer 2 circuits.            |
| Circuit            | Interface name and circuit information.                                                               |
| Primary            | LSP's primary standby path.                                                                           |
| Standby            | LSP's secondary standby path.                                                                         |
| VC bandwidth       | Bandwidth constraints associated with a Layer 2 circuit route.                                        |

## Sample Output

### show mpls call-admission-control

```

user@host# show mpls call-admission-control

LSP name: pro1-be
*Primary
 Available bandwidth: 0bps

LSP name: pro1-be-1
*Primary
 Available bandwidth: 60kbps

LSP name: pro1-be-gold
*Primary
 Available bandwidth: <ct0 50kbps> <ct1 20kbps> <ct2 30kbps> <ct3 0bps>
 Layer2 connections:
 Neighbor address: 10.255.245.215, Circuit: so-0/3/0.0(vc 5)
 VC bandwidth: <ct0 50kbps> <ct1 40kbps> <ct2 40kbps>

LSP name: pro1-be-gold-2
*Primary
 Available bandwidth: <ct0 0bps> <ct1 40kbps> <ct2 40kbps> <ct3 0bps>

LSP name: pro1-be-silver
*Primary prim1
 Available bandwidth: <ct0 10kbps> <ct1 20kbps> <ct2 0bps> <ct3 40kbps>
 Layer2 connections:
 Neighbor address: 10.255.245.215, Circuit: so-0/3/0.1(vc 3)
 VC bandwidth: <ct0 20kbps> <ct1 20kbps>
 Standby sec1
 Available bandwidth: <ct0 10kbps> <ct1 10kbps> <ct2 20kbps> <ct3 0bps>
 Layer2 connections:
 Neighbor address: 10.255.245.215, Circuit: so-0/3/0.1(vc 3)
 VC bandwidth: <ct0 20kbps> <ct1 20kbps>

```

## show mpls cspf

**List of Syntax** [Syntax on page 5162](#)  
[Syntax \(EX Series Switches\) on page 5162](#)

**Syntax** `show mpls cspf`  
`<instance instance-name>`  
`<logical-system (all | logical-system-name)>`

**Syntax (EX Series Switches)** `show mpls cspf`

**Release Information** Command introduced before Junos OS Release 7.4.  
 Command introduced in Junos OS Release 9.5 for EX Series switches.  
**instance *instance-name*** option added in Junos OS Release 15.1.

**Description** Display Multiprotocol Label Switching (MPLS) Constrained Shortest Path First (CSPF) statistics.

**Options** `none`—Display MPLS CSFP statistics.

**instance *instance-name***—(Optional) Display MPLS CSPF information for the specified instance. If *instance-name* is omitted, MPLS CSPF information for the master instance is displayed.

**logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

**Required Privilege Level** view

**List of Sample Output** [show mpls cspf on page 5163](#)

**Output Fields** [Table 394](#) describes the output fields for the **show mpls cspf** command. Output fields are listed in the approximate order in which they appear.

**Table 401: show mpls cspf Output Fields**

| Field Name          | Field Description                                     |
|---------------------|-------------------------------------------------------|
| <b>Queue length</b> | Number of LSPs queued for automatic path computation. |
| <b>current</b>      | Current queue length.                                 |
| <b>maximum</b>      | Maximum queue length (high-water mark).               |
| <b>dequeued</b>     | Number of aborted computation attempts.               |
| <b>Paths</b>        | Counters for label-switched path computations.        |
| <b>total</b>        | Sum of the next four fields.                          |

Table 401: show mpls cspf Output Fields (*continued*)

| Field Name          | Field Description                                                                          |
|---------------------|--------------------------------------------------------------------------------------------|
| <b>successful</b>   | Number of path computations that were successfully completed.                              |
| <b>no route</b>     | Number of path computations that failed because the destination is unreachable.            |
| <b>Sys Error</b>    | Number of path computations that failed because of lack of memory.                         |
| <b>CSPFs</b>        | Total number of CSPF computations. A single path might require multiple CSPF computations. |
| <b>Time</b>         | Time, in seconds, required to perform the label-switched path computation.                 |
| <b>Total</b>        | Total amount of time consumed by the CSPF path computation algorithm.                      |
| <b>CSPFs</b>        | Total number of CSPF computations.                                                         |
| <b>Avg per CSPF</b> | Average amount of time required for each CSPF computation.                                 |
| <b>% of rpd</b>     | Percentage of routing process CPU used in the CSPF computation.                            |

## Sample Output

### show mpls cspf

```

user@host> show mpls cspf
CSPF statistics
Queue length current maximum dequeued
 0 0 0
Paths total successful no route sys error CSPFs
 0 0 0 0 0
Time (secs) total CSPFs avg per CSPF % of rpd
 0.000000 0.000000 0.000000 0.0000

```

## show mpls diffserv-te

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 5164</a><br><a href="#">Syntax (EX Series Switches) on page 5164</a>                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Syntax</b>                      | <pre>show mpls diffserve-te &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                            |
| <b>Syntax (EX Series Switches)</b> | show mpls diffserve-te                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>         | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p><b>instance <i>instance-name</i></b> option added in Junos OS Release 15.1.</p>                                                                                                                                                                                                                                                                          |
| <b>Description</b>                 | Display Multiprotocol Label Switching (MPLS) label-switched path (LSP) Differentiated Services (DiffServ) class and preemption priority information.                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                     | <p><b>none</b>—Display DiffServ classes and priorities used by MPLS LSPs.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display DiffServ classes and priorities used by MPLS LSPs for the specified instance. If <b><i>instance-name</i></b> is omitted, DiffServ information for the master instance is displayed.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>List of Sample Output</b>       | <a href="#">show mpls diffserv-te on page 5165</a>                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Output Fields</b>               | Table 395 describes the output fields for the <b>show mpls diffserv-te</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                  |

**Table 402: show mpls diffserv-te Output Fields**

| Field Name             | Field Description                                                                                                                                                                                                                                                                       |
|------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Bandwidth model</b> | Bandwidth constraint model supported. The maximum allocation model (MAM) for EXP-inferred LSPs (E-LSPs) is currently supported.                                                                                                                                                         |
| <b>TE class</b>        | DiffServ traffic engineering class.                                                                                                                                                                                                                                                     |
| <b>Traffic class</b>   | <p>MPLS class type that corresponds to the DiffServ traffic engineering class:</p> <ul style="list-style-type: none"> <li>• <b>ct0</b>—Best effort</li> <li>• <b>ct1</b>—Assured forwarding</li> <li>• <b>ct2</b>—Expedited forwarding</li> <li>• <b>ct3</b>—Network control</li> </ul> |

Table 402: show mpls diffserv-te Output Fields (*continued*)

| Field Name | Field Description                                                                                                                                                                                  |
|------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Priority   | MPLS preemption priority for this class type, a value from 0 through 7. Interior gateway protocols (IGPs) distribute information about the available bandwidth for each traffic engineering class. |

## Sample Output

### show mpls diffserv-te

```
user@host> show mpls diffserv-te
Bandwidth model: Maximum Allocation Model with support for E-LSPs.
TE class Traffic class Priority
te0 ct0 3
te1 ct1 2
```

## show route forwarding-table

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show route forwarding-table</code><br><code>&lt;detail   extensive   summary&gt;</code><br><code>&lt;ccc ccc-interface-name&gt;</code><br><code>&lt;destination&gt;</code><br><code>&lt;family family-name&gt;</code><br><code>&lt;label label&gt;</code><br><code>&lt;matching ip_prefix&gt;</code><br><code>&lt;multicast&gt;</code><br><code>&lt;vpn vpn&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.5 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Display the Routing Engine's forwarding table, including the network-layer prefixes and their next hops. This command is used to help verify that the routing protocol process has relayed the correction information to the forwarding table. The Routing Engine constructs and maintains one or more routing tables. From the routing tables, the Routing Engine derives a table of active routes, called the forwarding table.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <p><b>none</b>—Display the routes in the forwarding table.</p> <p><b>detail   extensive   summary</b>—(Optional) Display the specified level of output.</p> <p><b>ccc</b>—(Optional) Display the specified circuit cross-connect interface name for entries to match.</p> <p><b>destination</b>—(Optional) Display the destination prefix.</p> <p><b>family family-name</b>—(Optional) Display routing table entries for the specified family: <b>ethernet-switching</b>, <b>inet</b>, <b>inet6</b>, <b>iso</b>, <b>mpls</b>, <b>vlan classification</b>.</p> <p><b>label label</b>—(Optional) Display route entries for the specified label name.</p> <p><b>matching ip_prefix</b>—(Optional) Display route entries for the specified IP prefix.</p> <p><b>multicast</b>—(Optional) Display route entries for multicast routes.</p> <p><b>vpn vpn</b>—(Optional) Display route entries for the specified VPN.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring MPLS on EX8200 and EX4500 Switches</i></li><li>• <i>Configuring MPLS on EX8200 and EX4500 Provider Switches (CLI Procedure)</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>List of Sample Output</b>    | <p><a href="#">show route forwarding-table on page 5168</a></p> <p><a href="#">show route forwarding-table summary on page 5169</a></p> <p><a href="#">show route forwarding-table extensive on page 5169</a></p> <p><a href="#">show route forwarding-table ccc on page 5171</a></p> <p><a href="#">show route forwarding-table family (MPLS) on page 5171</a></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

[show route forwarding-table family \(IPv6\) on page 5171](#)

[show route forwarding-table label on page 5172](#)

[show route forwarding-table matching on page 5172](#)

[show route forwarding-table multicast on page 5172](#)

**Output Fields** Table 306 lists the output fields for the **show route forwarding-table** command. Output fields are listed in the approximate order in which they appear. Field names might be abbreviated (as shown in parentheses) when no level of output is specified or when the **detail** keyword is used instead of the **extensive** keyword.

**Table 403: show route forwarding-table Output Fields**

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Level of Output                  |
|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------|
| <b>Routing table</b>           | Name of the routing table (for example, <b>inet</b> , <b>inet6</b> , <b>mpls</b> ).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | All levels                       |
| <b>Address family</b>          | Address family (for example, <b>IP</b> , <b>IPv6</b> , <b>ISO</b> , <b>MPLS</b> ).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | All levels                       |
| <b>Destination</b>             | Destination of the route.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail</b> , <b>extensive</b> |
| <b>Route Type (Type)</b>       | How the route was placed into the forwarding table. When the <b>detail</b> keyword is used, the route type might be abbreviated (as shown in parentheses): <ul style="list-style-type: none"> <li>• <b>cloned (clon)</b>—(TCP or multicast only) Cloned route.</li> <li>• <b>destination (dest)</b>—Remote addresses directly reachable through an interface.</li> <li>• <b>destination down (iddn)</b>—Destination route for which the interface is unreachable.</li> <li>• <b>interface cloned (ifcl)</b>—Cloned route for which the interface is unreachable.</li> <li>• <b>route down (ifdn)</b>—Interface route for which the interface is unreachable.</li> <li>• <b>ignore (ignr)</b>—Ignore this route.</li> <li>• <b>interface (intf)</b>—Installed as a result of configuring an interface.</li> <li>• <b>permanent (perm)</b>—Routes installed by the kernel when the routing table is initialized.</li> <li>• <b>user</b>—Routes installed by the routing protocol process or as a result of the configuration.</li> </ul> | All levels                       |
| <b>Route reference (RtRef)</b> | Number of routes to reference.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>detail</b> , <b>extensive</b> |
| <b>Flags</b>                   | Route type flags: <ul style="list-style-type: none"> <li>• <b>none</b>—No flags are enabled.</li> <li>• <b>accounting</b>—Route has accounting enabled.</li> <li>• <b>cached</b>—Cache route.</li> <li>• <b>incoming-iface interface-number</b>—Check against incoming interface.</li> <li>• <b>prefix load balance</b>—Load balancing is enabled for this prefix.</li> <li>• <b>sent to PFE</b>—Route has been sent to the Packet Forwarding Engine.</li> <li>• <b>static</b>—Static route.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>extensive</b>                 |
| <b>Nexthop</b>                 | IP address of the next hop to the destination.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>detail</b> , <b>extensive</b> |

Table 403: show route forwarding-table Output Fields (*continued*)

| Field Name                        | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Level of Output               |
|-----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------|
| <b>Next hop type (Type)</b>       | <p>Next-hop type. When the <b>detail</b> keyword is used, the next-hop type might be abbreviated (as indicated in parentheses):</p> <ul style="list-style-type: none"> <li>• <b>broadcast (bcst)</b>—Broadcast.</li> <li>• <b>deny</b>—Deny.</li> <li>• <b>hold</b>—Next hop is waiting to be resolved into a unicast or multicast type.</li> <li>• <b>indexed (idxd)</b>—Indexed next hop.</li> <li>• <b>indirect (indr)</b>—Indirect next hop.</li> <li>• <b>local (locl)</b>—Local address on an interface.</li> <li>• <b>routed multicast (mcr)</b>—Regular multicast next hop</li> <li>• <b>multicast (mcst)</b>—Wire multicast next hop (limited to the LAN).</li> <li>• <b>multicast discard (mdsc)</b>—Multicast discard.</li> <li>• <b>multicast group (mgrp)</b> —Multicast group member.</li> <li>• <b>receive (rcv)</b>—Receive.</li> <li>• <b>reject (rjct)</b>—Discard. An ICMP unreachable message was sent.</li> <li>• <b>resolve (rslv)</b>—Resolving the next hop.</li> <li>• <b>unicast (ucst)</b>—Unicast.</li> <li>• <b>unilist (ulst)</b>—List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.</li> </ul> | <b>detail, extensive</b>      |
| <b>Index</b>                      | Software index of the next hop that is used to route the traffic for a given prefix.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>detail, extensive none</b> |
| <b>Route interface-index</b>      | Logical interface index from which the route is learned. For example, for interface routes, this is the logical interface index of the route itself. For static routes, this field is zero. For routes learned through routing protocols, this is the logical interface index from which the route is learned.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>extensive</b>              |
| <b>Reference (NhRef)</b>          | Number of routes that refer to this next hop.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>none detail, extensive</b> |
| <b>Next-hop interface (Netif)</b> | Interface used to reach the next hop.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>none detail, extensive</b> |
| <b>Alternate forward nh index</b> | Index number of the alternate next hop interface. Seen with <b>multicast</b> option only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>extensive</b>              |
| <b>Next-hop L3 Interface</b>      | The next hop layer 3 interface. This option can be expressed as a VLAN name and is only seen with the <b>multicast</b> option.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>extensive</b>              |
| <b>Next-hop L2 Interfaces</b>     | The next hop layer 2 interfaces. Seen with <b>multicast</b> option only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>extensive</b>              |

## Sample Output

### show route forwarding-table

```

user@switch> show route forwarding-table

Routing table: default.inet

```



| Internet:          |      |       |                  |      |       |       |             |
|--------------------|------|-------|------------------|------|-------|-------|-------------|
| Destination        | Type | RtRef | Next hop         | Type | Index | NhRef | Netif       |
| default            | user | 2     | 0:12:f2:21:cf:0  | ucst | 333   | 5     | me0.0       |
| default            | perm | 0     |                  | rjct | 36    | 2     |             |
| 0.0.0.0/32         | perm | 0     |                  | dscd | 34    | 1     |             |
| 2.2.2.0/24         | intf | 0     |                  | rslv | 1309  | 1     | ae0.0       |
| 2.2.2.0/32         | dest | 0     | 2.2.2.0          | recv | 1307  | 1     | ae0.0       |
| 2.2.2.1/32         | dest | 0     | 0:21:59:cc:89:c0 | ucst | 1320  | 1     | ae0.0       |
| 2.2.2.2/32         | intf | 0     | 2.2.2.2          | loc1 | 1308  | 2     |             |
| 2.2.2.2/32         | dest | 0     | 2.2.2.2          | loc1 | 1308  | 2     |             |
| 2.2.2.255/32       | dest | 0     | 2.2.2.255        | bcst | 1306  | 1     | ae0.0       |
| 3.3.3.0/24         | intf | 0     |                  | rslv | 1313  | 1     | ae1.0       |
| 3.3.3.0/32         | dest | 0     | 3.3.3.0          | recv | 1311  | 1     | ae1.0       |
| 3.3.3.1/32         | intf | 0     | 3.3.3.1          | loc1 | 1312  | 2     |             |
| 3.3.3.1/32         | dest | 0     | 3.3.3.1          | loc1 | 1312  | 2     |             |
| 3.3.3.2/32         | dest | 0     | 0:21:59:cc:89:c1 | ucst | 1321  | 24    | ae1.0       |
| 3.3.3.255/32       | dest | 0     | 3.3.3.255        | bcst | 1310  | 1     | ae1.0       |
| 4.4.4.0/24         | user | 0     | 3.3.3.2          | ucst | 1321  | 24    | ae1.0       |
| 8.8.8.8/32         | user | 0     | 3.3.3.2          | ucst | 1321  | 24    | ae1.0       |
| 9.9.9.9/32         | intf | 0     | 9.9.9.9          | loc1 | 1280  | 1     |             |
| 10.10.10.10/32     | user | 0     | 3.3.3.2          | ucst | 1321  | 24    | ae1.0       |
| 10.93.8.0/21       | intf | 0     |                  | rslv | 323   | 1     | me0.0       |
| 10.93.8.0/32       | dest | 0     | 10.93.8.0        | recv | 321   | 1     | me0.0       |
| 10.93.13.238/32    | intf | 0     | 10.93.13.238     | loc1 | 322   | 2     |             |
| 10.93.13.238/32    | dest | 0     | 10.93.13.238     | loc1 | 322   | 2     |             |
| 10.93.15.254/32    | dest | 0     | 0:12:f2:21:cf:0  | ucst | 333   | 5     | me0.0       |
| 10.93.15.255/32    | dest | 0     | 10.93.15.255     | bcst | 320   | 1     | me0.0       |
| 14.14.14.0/24      | ifdn | 0     |                  | rslv | 1319  | 1     | ge-0/0/25.0 |
| 14.14.14.0/32      | iddn | 0     | 14.14.14.0       | recv | 1317  | 1     | ge-0/0/25.0 |
| 14.14.14.2/32      | user | 0     |                  | rjct | 36    | 2     |             |
| 14.14.14.2/32      | intf | 0     | 14.14.14.2       | loc1 | 1318  | 2     |             |
| 14.14.14.2/32      | iddn | 0     | 14.14.14.2       | loc1 | 1318  | 2     |             |
| 14.14.14.255/32    | iddn | 0     | 14.14.14.255     | bcst | 1316  | 1     | ge-0/0/25.0 |
| 224.0.0.0/4        | perm | 1     |                  | mdsc | 35    | 1     |             |
| 224.0.0.1/32       | perm | 0     | 224.0.0.1        | mcst | 31    | 3     |             |
| 224.0.0.5/32       | user | 1     | 224.0.0.5        | mcst | 31    | 3     |             |
| 255.255.255.255/32 | perm | 0     |                  | bcst | 32    | 1     |             |

### show route forwarding-table summary

```
user@switch> show route forwarding-table summary
```

Routing table: default.inet

Internet:

```

user: 6 routes
perm: 5 routes
intf: 8 routes
dest: 12 routes
ifdn: 1 routes
iddn: 3 routes

```

### show route forwarding-table extensive

```
user@switch> show route forwarding-table summary
```

Routing table: default.inet [Index 0]

Internet:

Destination: default

Route type: user

Route reference: 2

Route interface-index: 0

|                                     |                           |              |
|-------------------------------------|---------------------------|--------------|
| Flags: sent to PFE, rt nh decoupled |                           |              |
| Nexthop: 0:12:f2:21:cf:0            |                           |              |
| Next-hop type: unicast              | Index: 333                | Reference: 5 |
| Next-hop interface: me0.0           |                           |              |
| Destination: default                |                           |              |
| Route type: permanent               |                           |              |
| Route reference: 0                  | Route interface-index: 0  |              |
| Flags: none                         |                           |              |
| Next-hop type: reject               | Index: 36                 | Reference: 2 |
| Destination: 0.0.0.0/32             |                           |              |
| Route type: permanent               |                           |              |
| Route reference: 0                  | Route interface-index: 0  |              |
| Flags: sent to PFE                  |                           |              |
| Next-hop type: discard              | Index: 34                 | Reference: 1 |
| Destination: 2.2.2.0/24             |                           |              |
| Route type: interface               |                           |              |
| Route reference: 0                  | Route interface-index: 66 |              |
| Flags: sent to PFE                  |                           |              |
| Next-hop type: resolve              | Index: 1309               | Reference: 1 |
| Next-hop interface: ae0.0           |                           |              |
| Destination: 2.2.2.0/32             |                           |              |
| Route type: destination             |                           |              |
| Route reference: 0                  | Route interface-index: 66 |              |
| Flags: sent to PFE                  |                           |              |
| Nexthop: 2.2.2.0                    |                           |              |
| Next-hop type: receive              | Index: 1307               | Reference: 1 |
| Next-hop interface: ae0.0           |                           |              |
| Destination: 2.2.2.1/32             |                           |              |
| Route type: destination             |                           |              |
| Route reference: 0                  | Route interface-index: 66 |              |
| Flags: sent to PFE                  |                           |              |
| Nexthop: 0:21:59:cc:89:c0           |                           |              |
| Next-hop type: unicast              | Index: 1320               | Reference: 1 |
| Next-hop interface: ae0.0           |                           |              |
| Destination: 2.2.2.2/32             |                           |              |
| Route type: interface               |                           |              |
| Route reference: 0                  | Route interface-index: 0  |              |
| Flags: sent to PFE                  |                           |              |
| Nexthop: 2.2.2.2                    |                           |              |
| Next-hop type: local                | Index: 1308               | Reference: 2 |
| Destination: 2.2.2.2/32             |                           |              |
| Route type: destination             |                           |              |
| Route reference: 0                  | Route interface-index: 66 |              |
| Flags: none                         |                           |              |
| Nexthop: 2.2.2.2                    |                           |              |
| Next-hop type: local                | Index: 1308               | Reference: 2 |
| Destination: 2.2.2.255/32           |                           |              |
| Route type: destination             |                           |              |
| Route reference: 0                  | Route interface-index: 66 |              |
| Flags: sent to PFE                  |                           |              |
| Nexthop: 2.2.2.255                  |                           |              |
| Next-hop type: broadcast            | Index: 1306               | Reference: 1 |
| Next-hop interface: ae0.0           |                           |              |

**show route forwarding-table ccc**

```

user@switch> show route forwarding-table ccc ge-0/0/0.10
Routing table: default.mpls
MPLS:
Destination Type RtRef Next hop Type Index NhRef Netif
ge-0/0/0.10 (CCC) user 0 3.3.3.2 Push 300112 1343 2 ae1.0

```

**show route forwarding-table family (MPLS)**

```

user@switch> show route forwarding-table family mpls

Routing table: default.mpls
MPLS:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0
0 user 0 recv 49 3
1 user 0 recv 49 3
2 user 0 recv 49 3
299776 user 0 Pop 1334 2 ge-0/0/0.10
299792 user 0 Pop 1339 2 ge-0/0/0.14
299808 user 0 Pop 1341 2 ge-0/0/0.2
299824 user 0 Pop 1344 2 ge-0/0/0.11
299840 user 0 Pop 1345 2 ge-0/0/0.13
299856 user 0 Pop 1346 2 ge-0/0/0.18
299872 user 0 Pop 1347 2 ge-0/0/0.16
299888 user 0 Pop 1348 2 ge-0/0/0.7
299904 user 0 Pop 1349 2 ge-0/0/0.20
299920 user 0 Pop 1350 2 ge-0/0/0.19
299936 user 0 Pop 1351 2 ge-0/0/0.17
299952 user 0 Pop 1352 2 ge-0/0/0.9
299968 user 0 Pop 1353 2 ge-0/0/0.1
299984 user 0 Pop 1354 2 ge-0/0/0.12
300000 user 0 Pop 1355 2 ge-0/0/0.8
300016 user 0 Pop 1356 2 ge-0/0/0.4
300032 user 0 Pop 1357 2 ge-0/0/0.5
300048 user 0 Pop 1358 2 ge-0/0/0.3
300064 user 0 Pop 1359 2 ge-0/0/0.15
ge-0/0/0.1 (CCC) user 0 3.3.3.2 Push 300064 1340 2 ae1.0
ge-0/0/0.2 (CCC) user 0 3.3.3.2 Push 299872 1328 2 ae1.0
ge-0/0/0.3 (CCC) user 0 3.3.3.2 Push 299792 1323 2 ae1.0
ge-0/0/0.4 (CCC) user 0 3.3.3.2 Push 300016 1337 2 ae1.0
ge-0/0/0.5 (CCC) user 0 3.3.3.2 Push 299824 1325 2 ae1.0
ge-0/0/0.7 (CCC) user 0 3.3.3.2 Push 299920 1331 2 ae1.0
ge-0/0/0.8 (CCC) user 0 3.3.3.2 Push 299840 1326 2 ae1.0
ge-0/0/0.9 (CCC) user 0 3.3.3.2 Push 299888 1329 2 ae1.0
ge-0/0/0.10 (CCC) user 0 3.3.3.2 Push 300112 1343 2 ae1.0
ge-0/0/0.11 (CCC) user 0 3.3.3.2 Push 299776 1322 2 ae1.0
ge-0/0/0.12 (CCC) user 0 3.3.3.2 Push 299952 1333 2 ae1.0
ge-0/0/0.13 (CCC) user 0 3.3.3.2 Push 300096 1342 2 ae1.0
ge-0/0/0.14 (CCC) user 0 3.3.3.2 Push 299984 1335 2 ae1.0
ge-0/0/0.15 (CCC) user 0 3.3.3.2 Push 299936 1332 2 ae1.0
ge-0/0/0.16 (CCC) user 0 3.3.3.2 Push 299808 1324 2 ae1.0
ge-0/0/0.17 (CCC) user 0 3.3.3.2 Push 300000 1336 2 ae1.0
ge-0/0/0.18 (CCC) user 0 3.3.3.2 Push 300032 1338 2 ae1.0
ge-0/0/0.19 (CCC) user 0 3.3.3.2 Push 299904 1330 2 ae1.0
ge-0/0/0.20 (CCC) user 0 3.3.3.2 Push 299856 1327 2 ae1.0

```

**show route forwarding-table family (IPv6)**

```

user@switch> show route forwarding-table family inet6

```

```

Routing table: default.inet6
Internet6:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0 rjct 44 1
::/128 perm 0 dscd 42 1
ff00::/8 perm 0 mdsc 43 1
ff02::1/128 perm 0 ff02::1 mcst 39 1

```

```

Routing table: default-switch.inet6
Internet6:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0 rjct 530 1
::/128 perm 0 dscd 528 1
2:1::3a00/312 user 0 indr 131070 2
 comp 572 1
2:1::3a82/320 user 0 indr 131071 3
 comp 573 1
2:1::3af0/320 user 0 indr 131071 3
 comp 573 1
2:1:0:ff00::/56 user 0 mdsc 529 2
ff00::/8 perm 0 mdsc 529 2
ff02::1/128 perm 0 ff02::1 mcst 526 1

```

```

Routing table: __master.anon__.inet6
Internet6:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0 rjct 554 1
::/128 perm 0 dscd 552 1
ff00::/8 perm 0 mdsc 553 1
ff02::1/128 perm 0 ff02::1 mcst 550 1

```

### show route forwarding-table label

```
user@switch> show route forwarding-table label 29976
```

```

Routing table: default.mpls
MPLS:
Destination Type RtRef Next hop Type Index NhRef Netif
299776 user 0 Pop 1334 2 ge-0/0/0.10

```

### show route forwarding-table matching

```
user@switch> show route forwarding-table matching 3
```

```

Routing table: default.inet
Internet:

```

### show route forwarding-table multicast

```
user@switch> show route forwarding-table multicast
```

```

Routing table: default.inet
Internet:
Destination Type RtRef Next hop Type Index NhRef Netif
224.0.0.0/4 perm 1 mdsc 35 1
224.0.0.1/32 perm 0 224.0.0.1 mcst 31 3
224.0.0.5/32 user 1 224.0.0.5 mcst 31 3

```

```

Routing table: __master.anon__.inet
Internet:
Destination Type RtRef Next hop Type Index NhRef Netif
224.0.0.0/4 perm 0 mdsc 1289 1

```

```
224.0.0.1/32 perm 0 224.0.0.1 mcst 1285 1
```

```
Routing table: default.inet6
```

```
Internet6:
```

| Destination | Type | RtRef | Next hop | Type | Index | NhRef | Netif |
|-------------|------|-------|----------|------|-------|-------|-------|
| ff00::/8    | perm | 0     |          | mdsc | 43    | 1     |       |
| ff02::1/128 | perm | 0     | ff02::1  | mcst | 39    | 1     |       |

## show mpls interface

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 5174</a><br><a href="#">Syntax (EX Series Switches) on page 5174</a>                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Syntax</b>                      | <pre>show mpls interface &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Syntax (EX Series Switches)</b> | show mpls interface                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>         | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p><b>instance <i>instance-name</i></b> option added in Junos OS Release 15.1.</p>                                                                                                                                                                                                                                                                                |
| <b>Description</b>                 | Display information about Multiprotocol Label Switching (MPLS)-enabled interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                     | <p><b>none</b>—Display information about MPLS-enabled interfaces.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information about MPLS-enabled interfaces for the specified routing instance. If <b>instance-name</b> is omitted, information about MPLS-enabled interfaces is displayed for the master instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Additional Information</b>      | MPLS is enabled on an interface when the interface is configured with both the <b>set protocol mpls interface <i>interface-name</i></b> and <b>set interface <i>interface-name</i> unit 0 family mpls</b> statements.                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>List of Sample Output</b>       | <a href="#">show mpls interface on page 5175</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Output Fields</b>               | <a href="#">Table 397</a> describes the output fields for the <b>show mpls interface</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                          |

**Table 404: show mpls interface Output Fields**

| Field Name                   | Field Description                                      |
|------------------------------|--------------------------------------------------------|
| <b>Interface</b>             | Name of the interface.                                 |
| <b>State</b>                 | State of the interface: <b>Up</b> or <b>Dn</b> (down). |
| <b>Administrative groups</b> | Administratively assigned colors of the link.          |

Table 404: show mpls interface Output Fields (*continued*)

| Field Name                            | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Maximum labels                        | Maximum number of MPLS labels upon which MPLS can operate on a logical interface. This is configured using the <b>maximum-labels</b> statement at the [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family mpls] or the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family mpls] hierarchy levels.                                                                                                                                                                                                            |
| Static protection revert time         | Time (in seconds) that a static LSP must wait before traffic reverts from the bypass path to the original path. This is configured using the <b>protection-revert-time</b> statement at the [edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i> static] or the [edit protocols mpls interface <i>interface-name</i> static] hierarchy levels.                                                                                                                                                                                                                    |
| Always mark connection protection tlv | Enabled or Disabled: Enabled indicates that the <b>always-mark-connection-protection-tlv</b> statement is configured at the [edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i> static] or the [edit protocols mpls interface <i>interface-name</i> static] hierarchy levels. When this statement is configured, it marks all OAM traffic transiting this interface in preparation for switching the traffic to an alternate path based on the OAM functionality. To switch traffic to the bypass LSP, the <b>switch-away-lsps</b> statement must be configured. |
| Switch away lsps                      | Enabled or Disabled: Enabled indicates that the <b>switch-away-lsps</b> statement is configured at the [edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i> static] or the [edit protocols mpls interface <i>interface-name</i> static] hierarchy levels. This enables you to switch an LSP away from a network node using a bypass LSP. This feature can be used in maintenance of active networks when a network device needs to be replaced without interrupting traffic passing through the network. The LSPs can be either static or dynamic.                |

## Sample Output

### show mpls interface

```
user@host> show mpls interface
```

```
Interface: ge-0/2/1.57
State: Up
Administrative group: <none>
Maximum labels: 5
Static protection revert time: 5 seconds
Always mark connection protection tlv: Disabled
Switch away lsps : Disabled
```

## show mpls lsp

---

**List of Syntax**    [Syntax on page 5176](#)  
                          [Syntax \(EX Series Switches\) on page 5176](#)

**Syntax**    show mpls lsp  
                  <brief | detail | extensive | terse>  
                  <autobandwidth>  
                  <bidirectional | unidirectional>  
                  <bypass>  
                  <count-active-routes>  
                  <defaults>  
                  <descriptions>  
                  <down | up>  
                  <externally-controlled>  
                  <externally-provisioned>  
                  <logical-system (all | *logical-system-name*)>  
                  <lsp-type>  
                  <name *name*>  
                  <p2mp>  
                  <statistics>  
                  <transit>

**Syntax (EX Series Switches)**    show mpls lsp  
                                  <brief | detail | extensive | terse>  
                                  <bidirectional | unidirectional>  
                                  <bypass>  
                                  <descriptions>  
                                  <down | up>  
                                  <externally-controlled>  
                                  <externally-provisioned>  
                                  <lsp-type>  
                                  <name *name*>  
                                  <p2mp>  
                                  <statistics>  
                                  <transit>

**Release Information**    Command introduced before Junos OS Release 7.4.  
                              **defaults** option added in Junos OS Release 8.5.  
                              Command introduced in Junos OS Release 9.5 for EX Series switches.  
                              Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

**Description**    Display information about configured and active dynamic Multiprotocol Label Switching (MPLS) label-switched paths (LSPs).

**Options**    **none**—Display standard information about all configured and active dynamic MPLS LSPs.

**brief | detail | extensive | terse**—(Optional) Display the specified level of output. The extensive option displays the same information as the detail option, but covers the most recent 50 events.

**autobandwidth**—(Optional) Display automatic bandwidth information. This option is explained separately (see [show mpls lsp autobandwidth](#)).



**bidirectional | unidirectional**—(Optional) Display bidirectional or unidirectional LSP information, respectively.

**bypass**—(Optional) Display LSPs used for protecting other LSPs.

**count-active-routes**—(Optional) Display active routes for LSPs.

**defaults**—(Optional) Display the MPLS LSP default settings.

**descriptions**—(Optional) Display the MPLS label-switched path (LSP) descriptions. To view this information, you must configure the description statement at the **[edit protocol mpls lsp]** hierarchy level. Only LSPs with a description are displayed. This command is only valid for the ingress routing device, because the description is not propagated in RSVP messages.

**down | up**—(Optional) Display only LSPs that are inactive or active, respectively.

**externally-controlled**—(Optional) Display the LSPs that are under the control of an external Path Computation Element (PCE).

**externally-provisioned**—(Optional) Display the LSPs that are generated dynamically and provisioned by an external Path Computation Element (PCE).

**instance *instance-name***—(Optional) Display MPLS LSP information for the specified instance. If *instance-name* is omitted, MPLS LSP information is displayed for the master instance.

**logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

***lsp-type***—(Optional) Display information about a particular LSP type:

- **bypass**—Sessions for bypass LSPs.
- **egress**—Sessions that terminate on this routing device.
- **ingress**—Sessions that originate from this routing device.
- **transit**—Sessions that pass through this routing device.

**name *name***—(Optional) Display information about the specified LSP or group of LSPs.

**p2mp**—(Optional) Display information about point-to-multipoint LSPs.

**statistics**—(Optional) (Ingress and transit routers only) Display accounting information about LSPs. Statistics are not available for LSPs on the egress routing device, because the penultimate routing device in the LSP sets the label to 0. Also, as the packet arrives at the egress routing device, the hardware removes its MPLS header and the packet reverts to being an IPv4 packet. Therefore, it is counted as an IPv4 packet, not an MPLS packet.



**NOTE:** If a bypass LSP is configured for the primary static LSP, display cumulative statistics of packets traversing through the protected LSP and bypass LSP when traffic is re-optimized when the protected LSP link is restored. (Bypass LSPs are not supported on QFX Series switches.)

When used with the **bypass** option (**show mpls lsp bypass statistics**), display statistics for the traffic that flows only through the bypass LSP.

**transit**—(Optional) Display LSPs transiting this routing device.

**Required Privilege Level**

view

**Related Documentation**

- [clear mpls lsp on page 5107](#)
- [show mpls lsp autobandwidth on page 5194](#)

**List of Sample Output**

[show mpls lsp defaults on page 5185](#)  
[show mpls lsp descriptions on page 5185](#)  
[show mpls lsp detail on page 5185](#)  
[show mpls lsp extensive on page 5186](#)  
[show mpls lsp detail \(When Egress Protection Is in Effect During a Local Repair\) on page 5187](#)  
[show mpls lsp extensive on page 5187](#)  
[show mpls lsp ingress extensive on page 5189](#)  
[show mpls lsp extensive \(automatic bandwidth adjustment enabled\) on page 5190](#)  
[show mpls lsp bypass extensive on page 5191](#)  
[show mpls lsp p2mp on page 5191](#)  
[show mpls lsp p2mp detail on page 5192](#)  
[show mpls lsp detail count-active-routes on page 5192](#)  
[show mpls lsp statistics extensive on page 5193](#)

**Output Fields**

[Table 405](#) describes the output fields for the **show mpls lsp** command. Output fields are listed in the approximate order in which they appear.

**Table 405: show mpls lsp Output Fields**

| Field Name         | Field Description                                                                                                                                                                                        | Level of Output |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Ingress LSP</b> | Information about LSPs on the ingress routing device. Each session has one line of output.                                                                                                               | All levels      |
| <b>Egress LSP</b>  | Information about the LSPs on the egress routing device. MPLS learns this information by querying RSVP, which holds all the transit and egress session information. Each session has one line of output. | All levels      |
| <b>Transit LSP</b> | Number of LSPs on the transit routing devices and the state of these paths. MPLS learns this information by querying RSVP, which holds all the transit and egress session information.                   | All levels      |

Table 405: show mpls lsp Output Fields (*continued*)

| Field Name                  | Field Description                                                                                                                                                                                                                                                                     | Level of Output         |
|-----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>P2MP name</b>            | Name of the point-to-multipoint LSP. Dynamically generated P2MP LSPs used for VPLS flooding use dynamically generated P2MP LSP names. The name uses the format <i>identifier:vpls:router-id:routing-instance-name</i> . The <i>identifier</i> is automatically generated by Junos OS. | All levels              |
| <b>P2MP branch count</b>    | Number of destination LSPs the point-to-multipoint LSP is transmitting to.                                                                                                                                                                                                            | All levels              |
| <b>P</b>                    | An asterisk (*) under this heading indicates that the LSP is a primary path.                                                                                                                                                                                                          | All levels              |
| <b>address</b>              | ( <b>detail</b> and <b>extensive</b> ) Destination (egress routing device) of the LSP.                                                                                                                                                                                                | <b>detail extensive</b> |
| <b>To</b>                   | Destination (egress routing device) of the session.                                                                                                                                                                                                                                   | <b>brief</b>            |
| <b>From</b>                 | Source (ingress routing device) of the session.                                                                                                                                                                                                                                       | <b>brief detail</b>     |
| <b>State</b>                | State of the LSP handled by this RSVP session: <b>Up</b> , <b>Dn</b> (down), or <b>Restart</b> .                                                                                                                                                                                      | <b>brief detail</b>     |
| <b>Active Route</b>         | Number of active routes (prefixes) installed in the forwarding table. For ingress LSPs, the forwarding table is the primary IPv4 table ( <b>inet.0</b> ). For transit and egress RSVP sessions, the forwarding table is the primary MPLS table ( <b>mpls.0</b> ).                     | <b>detail extensive</b> |
| <b>Rt</b>                   | Number of active routes (prefixes) installed in the routing table. For ingress RSVP sessions, the routing table is the primary IPv4 table ( <b>inet.0</b> ). For transit and egress RSVP sessions, the routing table is the primary MPLS table ( <b>mpls.0</b> ).                     | <b>brief</b>            |
| <b>P</b>                    | Path. An asterisk (*) underneath this column indicates that the LSP is a primary path.                                                                                                                                                                                                | <b>brief</b>            |
| <b>ActivePath</b>           | (Ingress LSP) Name of the active path: <b>Primary</b> or <b>Secondary</b> .                                                                                                                                                                                                           | <b>detail extensive</b> |
| <b>LSPname</b>              | Name of the LSP.                                                                                                                                                                                                                                                                      | <b>brief detail</b>     |
| <b>Statistics</b>           | Displays the number of packets and the number of bytes transmitted over the LSP. These counters are reset to zero whenever the LSP path is optimized (for example, during an automatic bandwidth allocation).                                                                         | <b>extensive</b>        |
| <b>Aggregate statistics</b> | Displays the number of packets and the number of bytes transmitted over the LSP. These counters continue to iterate even if the LSP path is optimized. You can reset these counters to zero using the <b>clear mpls lsp statistics</b> command.                                       | <b>extensive</b>        |
| <b>Packets</b>              | Displays the number of packets transmitted over the LSP.                                                                                                                                                                                                                              | <b>brief extensive</b>  |
| <b>Bytes</b>                | Displays the number of bytes transmitted over the LSP.                                                                                                                                                                                                                                | <b>brief extensive</b>  |
| <b>DiffServInfo</b>         | Type of LSP: multiclass LSP ( <b>multiclass diffServ-TE LSP</b> ) or Differentiated-Services-aware traffic engineering LSP ( <b>diffServ-TE LSP</b> ).                                                                                                                                | <b>detail</b>           |

Table 405: show mpls lsp Output Fields (*continued*)

| Field Name                          | Field Description                                                                                                                                                                                                                                                                                                             | Level of Output         |
|-------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>LSPtype</b>                      | Type of LSP: <ul style="list-style-type: none"> <li>• <b>Static configured</b>—Static</li> <li>• <b>Dynamic configured</b>—Dynamic</li> <li>• <b>Externally controlled</b>—External path computing entity</li> </ul> Also indicates if the LSP is a <b>Penultimate hop popping</b> LSP or an <b>Ultimate hop popping</b> LSP. | <b>detail extensive</b> |
| <b>Bypass</b>                       | (Bypass LSP) Destination address (egress routing device) for the bypass LSP.                                                                                                                                                                                                                                                  | All levels              |
| <b>LSPpath</b>                      | Indicates whether the RSVP session is for the primary or secondary LSP path. <b>LSPpath</b> can be either <b>primary</b> or <b>secondary</b> and can be displayed on the ingress, egress, and transit routing devices.                                                                                                        | <b>detail</b>           |
| <b>Bidir</b>                        | (GMPLS) The LSP allows data to travel in both directions between GMPLS devices.                                                                                                                                                                                                                                               | All levels              |
| <b>Bidirectional</b>                | (GMPLS) The LSP allows data to travel both ways between GMPLS devices.                                                                                                                                                                                                                                                        | All levels              |
| <b>FastReroute desired</b>          | Fast reroute has been requested by the ingress routing device.                                                                                                                                                                                                                                                                | <b>detail</b>           |
| <b>Link protection desired</b>      | <b>detail</b>                                                                                                                                                                                                                                                                                                                 |                         |
| <b>Node/Link protection desired</b> | Link protection has been requested by the ingress routing device.                                                                                                                                                                                                                                                             | <b>detail</b>           |
| <b>LoadBalance</b>                  | (Ingress LSP) CSPF load-balancing rule that was configured to select the LSP's path among equal-cost paths: <b>Most-fill</b> , <b>Least-fill</b> , or <b>Random</b> .                                                                                                                                                         | <b>detail extensive</b> |
| <b>Signal type</b>                  | Signal type for GMPLS LSPs. The signal type determines the peak data rate for the LSP: <b>DS0</b> , <b>DS3</b> , <b>STS-1</b> , <b>STM-1</b> , or <b>STM-4</b> .                                                                                                                                                              | All levels              |
| <b>Encoding type</b>                | LSP encoding type: <b>Packet</b> , <b>Ethernet</b> , <b>PDH</b> , <b>SDH/SONET</b> , <b>Lambda</b> , or <b>Fiber</b> .                                                                                                                                                                                                        | All levels              |
| <b>Switching type</b>               | Type of switching on the links needed for the LSP: <b>Fiber</b> , <b>Lambda</b> , <b>Packet</b> , <b>TDM</b> , or <b>PSC-1</b> .                                                                                                                                                                                              | All levels              |
| <b>GPID</b>                         | Generalized Payload Identifier (identifier of the payload carried by an LSP): <b>HDLC</b> , <b>Ethernet</b> , <b>IPv4</b> , <b>PPP</b> , or <b>Unknown</b> .                                                                                                                                                                  | All levels              |
| <b>Protection</b>                   | Configured protection capability desired for the LSP: <b>Extra</b> , <b>Enhanced</b> , <b>none</b> , <b>One plus one</b> , <b>One to one</b> , or <b>Shared</b> .                                                                                                                                                             | All levels              |
| <b>Upstream label in</b>            | (Bidirectional LSPs) Incoming label for reverse direction traffic for this LSP.                                                                                                                                                                                                                                               | All levels              |
| <b>Upstream label out</b>           | (Bidirectional LSPs) Outgoing label for reverse direction traffic for this LSP.                                                                                                                                                                                                                                               | All levels              |

Table 405: show mpls lsp Output Fields (*continued*)

| Field Name                                          | Field Description                                                                                                                                                                       | Level of Output  |
|-----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>Suggested label received</b>                     | (Bidirectional LSPs) Label the upstream node suggests to use in the Resv message that is sent.                                                                                          | All levels       |
| <b>Suggested label sent</b>                         | (Bidirectional LSPs) Label the downstream node suggests to use in the Resv message that is returned.                                                                                    | All levels       |
| <b>Autobandwidth</b>                                | (Ingress LSP) The LSP is performing autobandwidth allocation.                                                                                                                           | detail extensive |
| <b>MinBW</b>                                        | (Ingress LSP) Configured minimum value of the LSP, in bps.                                                                                                                              | detail extensive |
| <b>MaxBW</b>                                        | (Ingress LSP) Configured maximum value of the LSP, in bps.                                                                                                                              | detail extensive |
| <b>AdjustTimer</b>                                  | (Ingress LSP) Configured value of the bandwidth adjustment timer, indicating the total amount of time allowed before bandwidth adjustment will take place, in seconds.                  | detail extensive |
| <b>Adjustment Threshold</b>                         | (Ingress LSP) Configured value for the <b>adjust-threshold</b> statement. Specifies how sensitive the automatic bandwidth adjustment for an LSP is to changes in bandwidth utilization. | detail extensive |
| <b>Time for Next Adjustment</b>                     | (Ingress LSP) Time in seconds until the next automatic bandwidth adjustment sample is taken.                                                                                            | detail extensive |
| <b>Time of Last Adjustment</b>                      | (Ingress LSP) Date and time since the last automatic bandwidth adjustment was completed.                                                                                                | detail extensive |
| <b>MaxAvgBW util</b>                                | (Ingress LSP) Current value of the actual maximum average bandwidth utilization, in bps.                                                                                                | detail extensive |
| <b>Overflow limit</b>                               | (Ingress LSP) Configured value of the threshold overflow limit.                                                                                                                         | detail extensive |
| <b>Overflow sample count</b>                        | (Ingress LSP) Current value for the overflow sample count.                                                                                                                              | detail extensive |
| <b>Bandwidth Adjustment in <i>nnn</i> second(s)</b> | (Ingress LSP) Current value of the bandwidth adjustment timer, indicating the amount of time remaining until the bandwidth adjustment will take place, in seconds.                      | detail extensive |
| <b>Underflow limit</b>                              | (Ingress LSP) Configured value of the threshold underflow limit.                                                                                                                        | detail extensive |
| <b>Underflow sample count</b>                       | (Ingress LSP) Current value for the underflow sample count.                                                                                                                             | detail extensive |
| <b>Underflow Max AvgBW</b>                          | (Ingress LSP) The highest sample bandwidth among the underflow samples recorded currently. This is the signaling bandwidth if an adjustment occurs because of an underflow.             | detail extensive |

Table 405: show mpls lsp Output Fields (*continued*)

| Field Name                                              | Field Description                                                                                                                                                                                                                                                                                                                           | Level of Output  |
|---------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>Active path indicator</b>                            | (Ingress LSP) A value of * indicates that the path is active. The absence of * indicates that the path is not active. In the following example, "long" is the active path.<br><br>*Primary long<br>Standby short                                                                                                                            | detail extensive |
| <b>Primary</b>                                          | (Ingress LSP) Name of the primary path.                                                                                                                                                                                                                                                                                                     | detail extensive |
| <b>Secondary</b>                                        | (Ingress LSP) Name of the secondary path.                                                                                                                                                                                                                                                                                                   | detail extensive |
| <b>Standby</b>                                          | (Ingress LSP) Name of the path in standby mode.                                                                                                                                                                                                                                                                                             | detail extensive |
| <b>State</b>                                            | (Ingress LSP) State of the path: <b>Up</b> or <b>Dn</b> (down).                                                                                                                                                                                                                                                                             | detail extensive |
| <b>COS</b>                                              | (Ingress LSP) Class-of-service value.                                                                                                                                                                                                                                                                                                       | detail extensive |
| <b>Bandwidth per class</b>                              | (Ingress LSP) Active bandwidth for the LSP path for each MPLS class type, in bps.                                                                                                                                                                                                                                                           | detail extensive |
| <b>Priorities</b>                                       | (Ingress LSP) Configured value of the setup priority and the hold priority respectively (the setup priority is displayed first), where 0 is the highest priority and 7 is the lowest priority. If you have not explicitly configured these values, the default values are displayed (7 for the setup priority and 0 for the hold priority). | detail extensive |
| <b>OptimizeTimer</b>                                    | (Ingress LSP) Configured value of the optimize timer, indicating the total amount of time allowed before path reoptimization, in seconds.                                                                                                                                                                                                   | detail extensive |
| <b>SmartOptimizeTimer</b>                               | (Ingress LSP) Configured value of the smart optimize timer, indicating the total amount of time allowed before path reoptimization, in seconds.                                                                                                                                                                                             | detail extensive |
| <b>Reoptimization in xxx seconds</b>                    | (Ingress LSP) Current value of the optimize timer, indicating the amount of time remaining until the path will be reoptimized, in seconds.                                                                                                                                                                                                  | detail extensive |
| <b>Computed ERO (S [L] denotes strict [loose] hops)</b> | (Ingress LSP) Computed explicit route. A series of hops, each with an address followed by a hop indicator. The value of the hop indicator can be strict (S) or loose (L).                                                                                                                                                                   | detail extensive |
| <b>CSPF metric</b>                                      | (Ingress LSP) Constrained Shortest Path First metric for this path.                                                                                                                                                                                                                                                                         | detail extensive |

Table 405: show mpls lsp Output Fields (*continued*)

| Field Name          | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Level of Output               |
|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------|
| <b>Received RRO</b> | <p>(Ingress LSP) Received record route. A series of hops, each with an address followed by a flag. (In most cases, the received record route is the same as the computed explicit route. If <b>Received RRO</b> is different from <b>Computed ERO</b>, there is a topology change in the network, and the route is taking a detour.) The following flags identify the protection capability and status of the downstream node:</p> <ul style="list-style-type: none"> <li>• <b>0x01</b>—Local protection available. The link downstream from this node is protected by a local repair mechanism. This flag can be set only if the Local protection flag was set in the <b>SESSION_ATTRIBUTE</b> object of the corresponding Path message.</li> <li>• <b>0x02</b>—Local protection in use. A local repair mechanism is in use to maintain this tunnel (usually because of an outage of the link it was routed over previously).</li> <li>• <b>0x03</b>—Combination of <b>0x01</b> and <b>0x02</b>.</li> <li>• <b>0x04</b>—Bandwidth protection. The downstream routing device has a backup path providing the same bandwidth guarantee as the protected LSP for the protected section.</li> <li>• <b>0x08</b>—Node protection. The downstream routing device has a backup path providing protection against link and node failure on the corresponding path section. If the downstream routing device can set up only a link-protection backup path, the <b>Local protection available</b> bit is set but the <b>Node protection</b> bit is cleared.</li> <li>• <b>0x09</b>—Detour is established. Combination of <b>0x01</b> and <b>0x08</b>.</li> <li>• <b>0x10</b>—Preemption pending. The preempting node sets this flag if a pending preemption is in progress for the traffic engine LSP. This flag indicates to the ingress legacy edge router (LER) of this LSP that it should be rerouted.</li> <li>• <b>0x20</b>—Node ID. Indicates that the address specified in the RRO's IPv4 or IPv6 sub-object is a node ID address, which refers to the router address or router ID. Nodes must use the same address consistently.</li> <li>• <b>0xb</b>—Detour is in use. Combination of <b>0x01</b>, <b>0x02</b>, and <b>0x08</b>.</li> </ul> | <b>detail extensive</b>       |
| <b>Index number</b> | (Ingress LSP) Log entry number of each LSP path event. The numbers are in chronological descending order, with a maximum of 50 index numbers displayed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>extensive</b>              |
| <b>Date</b>         | (Ingress LSP) Date of the LSP event.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>extensive</b>              |
| <b>Time</b>         | (Ingress LSP) Time of the LSP event.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>extensive</b>              |
| <b>Event</b>        | (Ingress LSP) Description of the LSP event.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>extensive</b>              |
| <b>Created</b>      | (Ingress LSP) Date and time the LSP was created.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>extensive</b>              |
| <b>Resv style</b>   | (Bypass) RSVP reservation style. This field consists of two parts. The first is the number of active reservations. The second is the reservation style, which can be <b>FF</b> (fixed filter), <b>SE</b> (shared explicit), or <b>WF</b> (wildcard filter).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>brief detail extensive</b> |
| <b>Labelin</b>      | Incoming label for this LSP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>brief detail</b>           |
| <b>Labelout</b>     | Outgoing label for this LSP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>brief detail</b>           |

Table 405: show mpls lsp Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Level of Output     |
|--------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------|
| <b>LSPname</b>                 | Name of the LSP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>brief detail</b> |
| <b>Time left</b>               | Number of seconds remaining in the lifetime of the reservation.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>detail</b>       |
| <b>Since</b>                   | Date and time when the RSVP session was initiated.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail</b>       |
| <b>Tspec</b>                   | Sender's traffic specification, which describes the sender's traffic parameters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail</b>       |
| <b>Port number</b>             | Protocol ID and sender or receiver port used in this RSVP session.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail</b>       |
| <b>PATH rcvfrom</b>            | Address of the previous-hop (upstream) routing device or client, interface the neighbor used to reach this router, and number of packets received from the upstream neighbor.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>detail</b>       |
| <b>PATH sentto</b>             | Address of the next-hop (downstream) routing device or client, interface used to reach this neighbor, and number of packets sent to the downstream routing device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail</b>       |
| <b>RESV rcvfrom</b>            | Address of the previous-hop (upstream) routing device or client, interface the neighbor used to reach this routing device, and number of packets received from the upstream neighbor. The output in this field, which is consistent with that in the <b>PATH rcvfrom</b> field, indicates that the RSVP negotiation is complete.                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail</b>       |
| <b>Record route</b>            | Recorded route for the session, taken from the record route object.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>detail</b>       |
| <b>Soft preempt</b>            | Number of soft preemptions that occurred on a path and when the last soft preemption occurred. Only successful soft preemptions are counted (those that actually resulted in a new path being used).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail</b>       |
| <b>Soft preemption pending</b> | Path is in the process of being soft preempted. This display is removed once the ingress router has calculated a new path.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail</b>       |
| <b>MPLS-TE LSP Defaults</b>    | Default settings for MPLS traffic engineered LSPs: <ul style="list-style-type: none"> <li>• <b>LSP Holding Priority</b>—Determines the degree to which an LSP holds on to its session reservation after the LSP has been set up successfully.</li> <li>• <b>LSP Setup Priority</b>—Determines whether a new LSP that preempts an existing LSP can be established.</li> <li>• <b>Hop Limit</b>—Specifies the maximum number of routers the LSP can traverse (including the ingress and egress).</li> <li>• <b>Bandwidth</b>—Specifies the bandwidth in bits per second for the LSP.</li> <li>• <b>LSP Retry Timer</b>—Length of time in seconds that the ingress router waits between attempts to establish the primary path.</li> </ul> | <b>defaults</b>     |

The XML tag name of the **bandwidth** tag under the **auto-bandwidth** tag has been updated to **maximum-average-bandwidth**. You can see the new tag when you issue the **show mpls lsp extensive** command with the **| display xml** pipe option. If you have any scripts that use the **bandwidth** tag, ensure that they are updated to **maximum-average-bandwidth**.



## Sample Output

### show mpls lsp defaults

```
user@host> show mpls lsp defaults
MPLS-TE LSP Defaults
 LSP Holding Priority 0
 LSP Setup Priority 7
 Hop Limit 255
 Bandwidth 0
 LSP Retry Timer 30 seconds
```

### show mpls lsp descriptions

```
user@host> show mpls lsp descriptions
Ingress LSP: 3 sessions
To LSP name Description
10.0.0.195 to-sanjose to-sanjose-desc
10.0.0.195 to-sanjose-other-desc other-desc
Total 2 displayed, Up 2, Down 0
```

### show mpls lsp detail

```
user@host> show mpls lsp detail
Ingress LSP: 1 sessions

192.168.0.4
 From: 192.168.0.5, State: Up, ActiveRoute: 0, LSPname: E-D
 ActivePath: (primary)
 LSPtype: Static Configured, Penultimate hop popping
 LoadBalance: Random
 Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary State: Up
 Priorities: 7 0
 SmartOptimizeTimer: 180
 Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
 10.0.0.18 S 10.0.0.22 S
 Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
 20=Node-ID):
 10.0.0.18 10.0.0.22
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

192.168.0.5
 From: 192.168.0.4, LSPstate: Up, ActiveRoute: 0
 LSPname: E-D, LSPpath: Primary
 Suggested label received: -, Suggested label sent: -
 Recovery label received: -, Recovery label sent: -
 Resv style: 1 FF, Label in: 3, Label out: -
 Time left: 157, Since: Wed Jul 18 17:55:12 2012
 Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
 Port number: sender 1 receiver 46128 protocol 0
 PATH rcvfrom: 10.0.0.18 (lt-1/2/0.17) 3 pkts
 Adspec: received MTU 1500
 PATH sentto: localclient
 RESV rcvfrom: localclient
 Record route: 10.0.0.22 10.0.0.18 <self>
Total 1 displayed, Up 1, Down 0
```

```
Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

### show mpls lsp extensive

```
user@host> show mpls lsp extensive
Ingress LSP: 1 sessions

192.168.0.4
 From: 192.168.0.5, State: Up, ActiveRoute: 0, LSPname: E-D
 ActivePath: (primary)
 LSPtype: Static Configured, Ultimate hop popping
 LoadBalance: Random
 Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary State: Up
 Priorities: 7 0
 SmartOptimizeTimer: 180
 Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
10.0.0.18 S 10.0.0.22 S
 Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
 10.0.0.18 10.0.0.22
 11 Sep 20 15:54:35.032 Make-before-break: Switched to new instance
 10 Sep 20 15:54:34.029 Record Route: 10.0.0.18 10.0.0.22
 9 Sep 20 15:54:34.029 Up
 8 Sep 20 15:54:20.271 Originate make-before-break call
 7 Sep 20 15:54:20.271 CSPF: computation result accepted 10.0.0.18 10.0.0.22

 6 Sep 20 15:52:10.247 Selected as active path
 5 Sep 20 15:52:10.246 Record Route: 10.0.0.18 10.0.0.22
 4 Sep 20 15:52:10.243 Up
 3 Sep 20 15:52:09.745 Originate Call
 2 Sep 20 15:52:09.745 CSPF: computation result accepted 10.0.0.18 10.0.0.22

 1 Sep 20 15:51:39.903 CSPF failed: no route toward 192.168.0.4
 Created: Thu Sep 20 15:51:08 2012
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

192.168.0.5
 From: 192.168.0.4, LSPstate: Up, ActiveRoute: 0
 LSPname: E-D, LSPpath: Primary
 Suggested label received: -, Suggested label sent: -
 Recovery label received: -, Recovery label sent: -
 Resv style: 1 FF, Label in: 3, Label out: -
 Time left: 148, Since: Thu Sep 20 15:52:10 2012
 Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
 Port number: sender 1 receiver 49601 protocol 0
 PATH rcvfrom: 10.0.0.18 (lt-1/2/0.17) 27 pkts
 Adspec: received MTU 1500
 PATH sentto: localclient
 RESV rcvfrom: localclient
 Record route: 10.0.0.22 10.0.0.18 <self>
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

**show mpls lsp detail (When Egress Protection Is in Effect During a Local Repair)**

```

user@host> show mpls lsp detail
Ingress LSP: 1 sessions

192.168.0.4
 From: 192.168.0.5, State: Up, ActiveRoute: 0, LSPname: E-D
 ActivePath: (primary)
 LSPtype: Static Configured, Penultimate hop popping
 LoadBalance: Random
 Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary State: Up
 Priorities: 7 0
 SmartOptimizeTimer: 180
 Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
10.0.0.18 S 10.0.0.22 S
 Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
 10.0.0.18 10.0.0.22
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

192.168.0.5
 From: 192.168.0.4, LSPstate: Down, ActiveRoute: 0
 LSPname: E-D, LSPpath: Primary
 Suggested label received: -, Suggested label sent: -
 Recovery label received: -, Recovery label sent: -
 Resv style: 1 FF, Label in: 3, Label out: -
 Time left: 157, Since: Wed Jul 18 17:55:12 2012
 Tspecc: rate 0bps size 0bps peak Infbps m 20 M 1500
 Port number: sender 1 receiver 46128 protocol 0
Egress protection PLR as protector: In Use
PATH rcvfrom: 10.0.0.18 (lt-1/2/0.17) 3 pkts
 Adspec: received MTU 1500
 PATH sentto: localclient
 RESV rcvfrom: localclient
 Record route: 10.0.0.22 10.0.0.18 <self>
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

**show mpls lsp extensive**

```

user@host> show mpls lsp extensive
Ingress LSP: 4 sessions

1.1.1.1
 From: 3.3.3.3, State: Up, ActiveRoute: 0, LSPname: m120b-to-mx960
 ActivePath: DEFAULT (primary)
 FastReroute desired
 LSPtype: Static Configured, Penultimate hop popping
 LoadBalance: Random
 Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary DEFAULT State: Up
 Priorities: 7 0
 SmartOptimizeTimer: 180
 Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 310)
10.0.35.5 S 10.0.15.1 S
 Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt

```

```

20=Node-ID):
 10.0.34.4(flag=1) 10.0.14.1
50 Sep 13 16:08:19.712 Record Route: 10.0.35.5(flag=1) 10.0.15.1
49 Sep 13 16:08:16.720 Record Route: 10.0.34.4(flag=1) 10.0.14.1
48 Sep 13 16:08:16.699 Fast-reroute Detour Up
47 Sep 13 16:08:13.702 Record Route: 10.0.34.4 10.0.14.1
46 Sep 13 16:08:13.702 Up
45 Sep 13 16:08:13.672 Originate make-before-break call
44 Sep 13 16:08:13.672 CSPF: computation result accepted 10.0.34.4 10.0.14.1

43 Sep 13 16:08:13.672 Selected as active path
42 Sep 13 16:08:13.672 Make-before-break: Switched to new instance
41 Sep 13 16:08:01.685 Pending path switchover, skip CSPF run[3 times]
40 Sep 13 16:06:33.910 Deselected as active
39 Sep 13 16:06:33.910 Pending path switchover, skip CSPF run

38 Sep 13 16:06:19.521 Record Route: 10.0.35.5 10.0.15.1
37 Sep 13 16:06:19.518 ResvTear received
36 Sep 13 16:06:19.518 Fast-reroute Detour Down
35 Sep 13 16:06:16.676 Record Route: 10.0.35.5(flag=1) 10.0.15.1
34 Sep 13 16:06:13.670 Record Route: 10.0.35.5 10.0.15.1
33 Sep 13 16:06:13.670 Up
32 Sep 13 16:06:13.569 Pending path switchover, skip CSPF run

31 Sep 13 16:06:13.569 CSPF: link down/deleted:
10.0.34.3(3.3.3.3:79)(m120-b-re1.00/3.3.3.3)->0.0.0.0(0.0.0.0:0)(m120-b-re1.04/0.0.0.0)

30 Sep 13 16:06:13.552 Pending path switchover, skip CSPF run

29 Sep 13 16:06:13.552 CSPF: link down/deleted:
0.0.0.0(0.0.0.0:0)(m120-b-re1.04/0.0.0.0)->0.0.0.0(4.4.4.4:0)(m10i-a-re0.00/4.4.4.4)

28 Sep 13 16:06:13.549 Originate make-before-break call
27 Sep 13 16:06:13.549 CSPF: computation result accepted 10.0.35.5 10.0.15.1

26 Sep 13 16:06:13.548 Tunnel local repaired
25 Sep 13 16:06:13.546 Record Route: 10.0.23.2 10.0.12.1
24 Sep 13 16:06:13.546 10.0.34.3: Tunnel local repaired
23 Sep 13 16:06:13.546 10.0.34.3: Down
22 Sep 13 16:03:46.842 Fast-reroute Detour Up
21 Sep 13 16:03:42.730 Record Route: 10.0.34.4(flag=1) 10.0.14.1
20 Sep 13 16:03:39.836 Selected as active path
19 Sep 13 16:03:39.834 Record Route: 10.0.34.4 10.0.14.1
18 Sep 13 16:03:39.834 Up
17 Sep 13 16:03:39.698 Originate Call
16 Sep 13 16:03:39.698 CSPF: computation result accepted 10.0.34.4 10.0.14.1

15 Sep 13 16:03:39.697 Clear Call
14 Sep 13 16:03:39.696 Deselected as active
13 Sep 13 16:03:37.837 Record Route: 10.0.34.4 10.0.14.1
12 Sep 13 16:03:32.829 Fast-reroute Detour Down
11 Sep 13 16:02:15.493 Record Route: 10.0.34.4(flag=1) 10.0.14.1
10 Sep 13 16:02:15.486 Fast-reroute Detour Up
9 Sep 13 16:02:12.468 Record Route: 10.0.34.4 10.0.14.1
8 Sep 13 16:02:07.460 Fast-reroute Detour Down
7 Sep 13 15:57:46.741 Fast-reroute Detour Up
6 Sep 13 15:57:40.768 Record Route: 10.0.34.4(flag=1) 10.0.14.1
5 Sep 13 15:57:37.761 Selected as active path
4 Sep 13 15:57:37.760 Record Route: 10.0.34.4 10.0.14.1
3 Sep 13 15:57:37.760 Up
2 Sep 13 15:57:37.733 Originate Call

```

```

1 Sep 13 15:57:37.733 CSPF: computation result accepted 10.0.34.4 10.0.14.1

Created: Fri Sep 13 15:57:38 2013
Total 1 displayed, Up 1, Down 0

Egress LSP: 4 sessions, 6 detours
Total 0 displayed, Up 0, Down 0

Transit LSP: 6 sessions, 1 detours

1.1.1.1
 From: 3.3.3.3, LSPstate: Up, ActiveRoute: 0
 LSPname: m120b-to-mx960, LSPpath: Primary
 Suggested label received: -, Suggested label sent: -
 Recovery label received: -, Recovery label sent: 302288
 Resv style: 1 FF, Label in: 300416, Label out: 302288
 Time left: 147, Since: Fri Sep 13 16:08:16 2013
 Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
 Port number: sender 4 receiver 13955 protocol 0
 Detour branch from 10.0.34.4, to skip 1.1.1.1, Up
 Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
 Adspec: received MTU 1500
 Path MTU: received 0
 PATH rcvfrom: 10.0.34.4 (ge-4/3/7.0) 7 pkts
 Adspec: received MTU 1500 sent MTU 1500
 PATH sentto: 10.0.35.5 (ge-3/1/0.0) 7 pkts
 RESV rcvfrom: 10.0.35.5 (ge-3/1/0.0) 7 pkts
 Explicit route: 10.0.35.5 10.0.15.1
 Record route: 10.0.34.3 10.0.34.4 <self>10.0.35.5 10.0.15.1
 Label in: 300416, Label out: 302288
Total 1 displayed, Up 1, Down 0

```

### show mpls lsp ingress extensive

```

user@host> show mpls lsp ingress extensive
Ingress LSP: 1 sessions

50.0.0.1
 From: 10.0.0.1, State: Up, ActiveRoute: 0, LSPname: test
 ActivePath: (primary)
 LSPtype: Static Configured
 LoadBalance: Random
 Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary State: Up
 Priorities: 7 0
 OptimizeTimer: 300
 SmartOptimizeTimer: 180
 Reoptimization in 240 second(s).
 Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
 1.1.1.2 S 4.4.4.1 S 5.5.5.2 S
 Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
 20=Node-ID):
 1.1.1.2 4.4.4.1 5.5.5.2
 17 Aug 3 13:17:33.601 CSPF: computation result ignored, new path less avail
 bw[3 times]
 16 Aug 3 13:02:51.283 CSPF: computation result ignored, new path no benefit[2
 times]
 15 Aug 3 12:54:36.678 Selected as active path
 14 Aug 3 12:54:36.676 Record Route: 1.1.1.2 4.4.4.1 5.5.5.2
 13 Aug 3 12:54:36.676 Up
 12 Aug 3 12:54:33.924 Deselected as active

```

```

11 Aug 3 12:54:33.924 Originate Call
10 Aug 3 12:54:33.923 Clear Call
9 Aug 3 12:54:33.923 CSPF: computation result accepted 1.1.1.2 4.4.4.1
5.5.5.2
8 Aug 3 12:54:33.922 2.2.2.2: No Route toward dest
7 Aug 3 12:54:28.177 CSPF: computation result ignored, new path no benefit[4
times]
6 Aug 3 12:35:03.830 Selected as active path
5 Aug 3 12:35:03.828 Record Route: 2.2.2.2 3.3.3.2
4 Aug 3 12:35:03.827 Up
3 Aug 3 12:35:03.814 Originate Call
2 Aug 3 12:35:03.814 CSPF: computation result accepted 2.2.2.2 3.3.3.2
1 Aug 3 12:34:34.921 CSPF failed: no route toward 50.0.0.1
Created: Tue Aug 3 12:34:35 2010
Total 1 displayed, Up 1, Down 0

```

### show mpls lsp extensive (automatic bandwidth adjustment enabled)

```

user@host> show mpls lsp extensive
Ingress LSP: 1 sessions

192.168.0.4
From: 192.168.0.5, State: Up, ActiveRoute: 0, LSPname: E-D
ActivePath: (primary)
Node/Link protection desired
LSPtype: Static Configured, Penultimate hop popping
LoadBalance: Random
Autobandwidth
MinBW: 300bps, MaxBW: 1000bps, Dynamic MinBW: 1000bps
Adjustment Timer: 300 secs AdjustThreshold: 25%
Max AvgBW util: 963.739bps, Bandwidth Adjustment in 0 second(s).
Min BW Adjust Interval: 1000, MinBW Adjust Threshold (in %): 50
Overflow limit: 0, Overflow sample count: 0
Underflow limit: 0, Underflow sample count: 9, Underflow Max AvgBW: 614.421bps

Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary State: Up
Priorities: 7 0
Bandwidth: 1000bps
SmartOptimizeTimer: 180
Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
10.0.0.18 S 10.0.0.22 S
Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
192.168.0.6(flag=0x20) 10.0.0.18(Label=299792) 192.168.0.4(flag=0x20)
10.0.0.22(Label=3)
12 Apr 30 10:25:17.024 Make-before-break: Switched to new instance
11 Apr 30 10:25:16.023 Record Route: 192.168.0.6(flag=0x20)
10.0.0.18(Label=299792) 192.168.0.4(flag=0x20) 10.0.0.22(Label=3)
10 Apr 30 10:25:16.023 Up
9 Apr 30 10:25:16.023 Automatic Autobw adjustment succeeded: BW changes from
300 bps to 1000 bps
8 Apr 30 10:25:15.946 Originate make-before-break call
7 Apr 30 10:25:15.946 CSPF: computation result accepted 10.0.0.18 10.0.0.22

6 Apr 30 10:16:42.891 Selected as active path
5 Apr 30 10:16:42.891 Record Route: 192.168.0.6(flag=0x20)
10.0.0.18(Label=299776) 192.168.0.4(flag=0x20) 10.0.0.22(Label=3)
4 Apr 30 10:16:42.890 Up
3 Apr 30 10:16:42.828 Originate Call
2 Apr 30 10:16:42.828 CSPF: computation result accepted 10.0.0.18 10.0.0.22

```

```

 1 Apr 30 10:16:14.064 CSPF: could not determine self[2 times]
Created: Tue Apr 30 10:15:16 2013
Total 1 displayed, Up 1, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

### show mpls lsp bypass extensive

```

user@host # show mpls lsp bypass extensive

Ingress LSP: 1 sessions

2.2.2.2
 From: 1.1.1.1, LSPstate: Up, ActiveRoute: 0
 LSPname: Bypass->1.1.2.2
 LSPtype: Static Configured
 Suggested label received: -, Suggested label sent: -
 Recovery label received: -, Recovery label sent: 300032
 Resv style: 1 SE, Label in: -, Label out: 300032
 Time left: -, Since: Tue Dec 3 15:19:49 2013
 Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
 Port number: sender 1 receiver 55750 protocol 0
 Type: Bypass LSP
 Number of data route tunnel through: 1
 Number of RSVP session tunnel through: 0
 PATH rcvfrom: localclient
 Adspec: sent MTU 1500
 Path MTU: received 1500
 PATH sentto: 1.1.5.2 (lt-1/2/0.15) 1221 pkts
 RESV rcvfrom: 1.1.5.2 (lt-1/2/0.15) 1221 pkts, Entropy label: No
 Explct route: 1.1.5.2 1.2.5.1
 Record route: <self> 1.1.5.2 1.2.5.1
+ 4 Dec 3 15:19:49 Record Route: 1.1.5.2 1.2.5.1
+ 3 Dec 3 15:19:49 Up
+ 2 Dec 3 15:19:49 CSPF: computation result accepted
+ 1 Dec 3 15:19:47 Originate Call
Total 1 displayed, Up 1, Down 0
Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
Transit LSP: 0 sessions

```

### show mpls lsp p2mp

```

user@host> show mpls lsp p2mp
Ingress LSP: 2 sessions
P2MP name: p2mp-lsp1, P2MP branch count: 1
To From State Rt P ActivePath LSPname
10.255.245.51 10.255.245.50 Up 0 * path1 p2mp-branch-1
P2MP name: p2mp-lsp2, P2MP branch count: 1
To From State Rt P ActivePath LSPname
10.255.245.51 10.255.245.50 Up 0 * path1 p2mp-st-br1
Total 2 displayed, Up 2, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

```
Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

### show mpls lsp p2mp detail

```
user@host> show mpls lsp p2mp detail
Ingress LSP: 2 sessions
P2MP name: p2mp-lsp1, P2MP branch count: 1

10.255.245.51
 From: 10.255.245.50, State: Up, ActiveRoute: 0, LSPname: p2mp-branch-1
 ActivePath: path1 (primary)
 P2MP name: p2mp-lsp1
 LoadBalance: Random
 Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary path1 State: Up
 Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 25)
 192.168.208.17 S
 Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

 192.168.208.17
P2MP name: p2mp-lsp2, P2MP branch count: 1

10.255.245.51
 From: 10.255.245.50, State: Up, ActiveRoute: 0, LSPname: p2mp-st-br1
 ActivePath: path1 (primary)
 P2MP name: p2mp-lsp2
 LoadBalance: Random
 Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary path1 State: Up
 Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 25)
 192.168.208.17 S
 Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

 192.168.208.17
Total 2 displayed, Up 2, Down 0
```

### show mpls lsp detail count-active-routes

```
user@host> show mpls lsp detail count-active-routes
Ingress LSP: 1 sessions

213.119.192.2
 From: 156.154.162.128, State: Up, ActiveRoute: 1, LSPname: to-lahore
 ActivePath: (primary)
 LSPtype: Static Configured
 LoadBalance: Random
 Autobandwidth
 MinBW: 5Mbps MaxBW: 250Mbps
 AdjustTimer: 300 secs
 Max AvgBW util: 0bps, Bandwidth Adjustment in 102 second(s).
 Overflow limit: 0, Overflow sample count: 0
 Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary State: Up
 Priorities: 7 0
 Bandwidth: 5Mbps
 SmartOptimizeTimer: 180
 Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 4)
 10.252.0.177 S
 Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
 20=Node-ID):
```



```

10.252.0.177
Total 1 displayed, Up 1, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

### show mpls lsp statistics extensive

```

user@host> show mpls lsp statistics extensive
Ingress LSP: 1 sessions

192.168.0.4
 From: 192.168.0.5, State: Up, ActiveRoute: 0, LSPName: E-D
 Statistics: Packets 302, Bytes 28992
 Aggregate statistics: Packets 302, Bytes 28992
 ActivePath: (primary)
 LSPType: Static Configured, Penultimate hop popping
 LoadBalance: Random
 Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary State: Up
 Priorities: 7 0
 SmartOptimizeTimer: 180
 Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
10.0.0.18 S 10.0.0.22 S
 Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
 10.0.0.18 10.0.0.22
 6 Oct 3 11:18:28.281 Selected as active path
 5 Oct 3 11:18:28.281 Record Route: 10.0.0.18 10.0.0.22
 4 Oct 3 11:18:28.280 Up
 3 Oct 3 11:18:27.995 Originate Call
 2 Oct 3 11:18:27.995 CSPF: computation result accepted 10.0.0.18 10.0.0.22

 1 Oct 3 11:17:59.118 CSPF failed: no route toward 192.168.0.4[2 times]
 Created: Wed Oct 3 11:17:01 2012
Total 1 displayed, Up 1, Down 0

```

## show mpls lsp autobandwidth

|                                 |                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show mpls lsp autobandwidth</code><br><code>&lt;brief   detail   extensive&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>                                                                                                                                                                                                     |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.4.<br>Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Display automatic bandwidth information for the LSP(s).                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <p><b>brief   detail   extensive</b> — (Optional) Display the specified level of output. The extensive option displays the same information as the detail option, but covers the most recent 50 events.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b> — (Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show mpls lsp on page 5176</a></li> </ul>                                                                                                                                                                                                                                                                            |
| <b>List of Sample Output</b>    | <a href="#">show mpls lsp autobandwidth on page 5195</a>                                                                                                                                                                                                                                                                                                                  |
| <b>Output Fields</b>            | Table 406 describes the output fields for the <b>show mpls lsp autobandwidth</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                         |

Table 406: show mpls lsp autobandwidth Output Fields

| Field Name                   | Field Description                                                                        | Level of Output         |
|------------------------------|------------------------------------------------------------------------------------------|-------------------------|
| <b>To</b>                    | Destination (egress routing device) of the session.                                      | All Levels              |
| <b>From</b>                  | Source (ingress routing device) of the session.                                          | All Levels              |
| <b>LSPname</b>               | Name of the LSP.                                                                         | All Levels              |
| <b>Min BW</b>                | (Ingress LSP) Configured minimum value of the LSP, in bps.                               | <b>detail extensive</b> |
| <b>Max BW</b>                | (Ingress LSP) Configured maximum value of the LSP, in bps.                               | <b>detail extensive</b> |
| <b>Max AvgBW util</b>        | (Ingress LSP) Current value of the actual maximum average bandwidth utilization, in bps. | <b>detail extensive</b> |
| <b>Overflow limit</b>        | (Ingress LSP) Configured value of the threshold overflow limit.                          | <b>detail extensive</b> |
| <b>Overflow sample count</b> | (Ingress LSP) Current value for the overflow sample count.                               | <b>detail extensive</b> |
| <b>Underflow limit</b>       | (Ingress LSP) Configured value of the threshold underflow limit.                         | <b>detail extensive</b> |

Table 406: show mpls lsp autobandwidth Output Fields (*continued*)

| Field Name                      | Field Description                                                                                                                                                                | Level of Output         |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Underflow sample count</b>   | (Ingress LSP) Current value for the underflow sample count.                                                                                                                      | <b>detail extensive</b> |
| <b>Adjustment Timer</b>         | (Ingress LSP) Configured value for the adjust-timer statement, indicating the total amount of time allowed before bandwidth adjustment will take place, in seconds.              | <b>detail extensive</b> |
| <b>Adjustment Threshold</b>     | (Ingress LSP) Configured value for the adjust-threshold statement. Specifies how sensitive the automatic bandwidth adjustment for an LSP is to changes in bandwidth utilization. | <b>detail extensive</b> |
| <b>Time for Next Adjustment</b> | (Ingress LSP) Time in seconds until the next automatic bandwidth adjustment sample is taken.                                                                                     | <b>detail extensive</b> |
| <b>Time of Last Adjustment</b>  | (Ingress LSP) Date and time since the last automatic bandwidth adjustment was completed.                                                                                         | <b>detail extensive</b> |
| <b>Last BW</b>                  | Previous active bandwidth of the LSP.                                                                                                                                            | <b>detail extensive</b> |
| <b>Last Requested BW</b>        | Bandwidth requested in the previous automatic bandwidth adjustment.                                                                                                              | <b>detail extensive</b> |
| <b>Last Signaled BW</b>         | Bandwidth signaled in the previous automatic bandwidth adjustment.                                                                                                               | <b>detail extensive</b> |
| <b>Highest Watermark BW</b>     | Maximum bandwidth used by the LSP.                                                                                                                                               | <b>detail extensive</b> |
| <b>Total AutoBw Adjustments</b> | Total number of attempts to adjust automatic bandwidth including failed and successful adjustments.                                                                              | <b>detail extensive</b> |
| <b>Successful Adjustments</b>   | Number of successful automatic bandwidth adjustments.                                                                                                                            | <b>detail extensive</b> |
| <b>Failed Adjustments</b>       | Number of failed automatic bandwidth adjustments.                                                                                                                                | <b>detail extensive</b> |

## Sample Output

### show mpls lsp autobandwidth

```

user@host> show mpls lsp autobandwidth extensive
To: 10.255.106.133,
From: 10.255.106.135, LSPname: r0-r1
Min BW: 100kbps, Max BW: 0bps, Max AvgBW util: 2.33249Mbps
Overflow limit: 0, Overflow sample count: 0
Underflow limit: 0, Underflow sample count: 0
Adjustment Timer: 300 sec, Adjustment Threshold: 0
Time for Next Adjustment: 23 sec, Time of Last Adjustment: Fri Jun 3 21:05:37
2011
Last BW: 100kbps, Last Requested BW: 2.2169Mbps, Last Signaled BW: 2.2169Mbps,
Highest Watermark BW: 2.33249Mbps
Total AutoBw Adjustments: 1, Successful Adjustments: 1, Failed Adjustments: 0

```



## show mpls path

**List of Syntax** [Syntax on page 5197](#)  
[Syntax \(EX Series Switches\) on page 5197](#)

**Syntax** show mpls path  
 <instance *instance-name*>  
 <logical-system (all | *logical-system-name*)>  
 <path-name>

**Syntax (EX Series Switches)** show mpls path  
 <path-name>

**Release Information** Command introduced before Junos OS Release 7.4.  
 Command introduced in Junos OS Release 9.5 for EX Series switches.  
**instance *instance-name*** option added in Junos OS Release 15.1.

**Description** Display dynamic Multiprotocol Label Switching (MPLS) label-switched paths (LSPs).

**Options** **none**—Display standard information about all MPLS LSPs.

**instance *instance-name***—(Optional) Display the dynamic MPLS LSP for the specified instance. If ***instance-name*** is omitted, dynamic MPLS LSP for the master instance is displayed.

**logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

**path-name**—(Optional) Display information about the specified LSP only.

**Required Privilege Level** view

**List of Sample Output** [show mpls path on page 5198](#)

**Output Fields** [Table 407](#) describes the output fields for the **show mpls path** command. Output fields are listed in the approximate order in which they appear.

**Table 407: show mpls path Output Fields**

| Field Name                  | Field Description                                                 |
|-----------------------------|-------------------------------------------------------------------|
| <b>Path name</b>            | Information about ingress LSPs. Each path has one line of output. |
| <b>Address</b>              | Addresses of the routing devices that form the LSP.               |
| <b>Strict/loose address</b> | Whether the address is configured as a strict or loose address.   |

## Sample Output

show mpls path

```
user@host> show mpls path
Path name Address Strict/loose address
p1 123.456.55.6 Strict
 123.456.1.6 Loose
p2 191.456.1.4 Strict
```

## show route table

|                                                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                                      | <a href="#">Syntax on page 5199</a><br><a href="#">Syntax (EX Series Switches and QFX Series Switches) on page 5199</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Syntax</b>                                              | show route table <i>routing-table-name</i><br><brief   detail   extensive   terse><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Syntax (EX Series Switches and QFX Series Switches)</b> | show route table <i>routing-table-name</i><br><brief   detail   extensive   terse>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>                                 | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 14.1X53-D15 for QFX Series switches.<br>Show route table evpn statement introduced in Junos OS Release 15.1X53-D30 for QFX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>                                         | Display the route entries in a particular routing table.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                                             | <b>brief   detail   extensive   terse</b> —(Optional) Display the specified level of output.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.<br><br><b><i>routing-table-name</i></b> —Display route entries for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the <b>show route table inet</b> command).                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b>                            | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>                               | <ul style="list-style-type: none"> <li>• <a href="#">show route summary on page 3554</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>List of Sample Output</b>                               | <a href="#">show route table bgp.l2.vpn on page 5210</a><br><a href="#">show route table bgp.l3vpn.0 on page 5210</a><br><a href="#">show route table bgp.l3vpn.0 detail on page 5210</a><br><a href="#">show route table bgp.rtarget.0 (When Proxy BGP Route Target Filtering Is Configured) on page 5211</a><br><a href="#">show route table bgp.evpn.0 on page 5212</a><br><a href="#">show route table evpna.evpn.0 on page 5212</a><br><a href="#">show route table inet.0 on page 5212</a><br><a href="#">show route table inet.3 on page 5213</a><br><a href="#">show route table inet6.0 on page 5213</a><br><a href="#">show route table inet6.3 on page 5213</a><br><a href="#">show route table inetflow detail on page 5214</a><br><a href="#">show route table l2circuit.0 on page 5214</a><br><a href="#">show route table mpls on page 5214</a><br><a href="#">show route table mpls extensive on page 5215</a> |

[show route table mpls.0 on page 5215](#)  
[show route table mpls.0 detail \(PTX Series\) on page 5215](#)  
[show route table mpls.0 extensive \(PTX Series\) on page 5216](#)  
[show route table mpls.0 \(RSVP Route—Transit LSP\) on page 5217](#)  
[show route table vpls\\_1 detail on page 5217](#)  
[show route table vpn-a on page 5217](#)  
[show route table vpn-a.mdt.0 on page 5218](#)  
[show route table VPN-A detail on page 5218](#)  
[show route table VPN-AB.inet.0 on page 5219](#)  
[show route table VPN\\_blue.mvpn-inet6.0 on page 5219](#)  
[show route table vrf1.mvpn.0 extensive on page 5219](#)  
[show route table MVPN.mvpn.0 on page 5220](#)  
[show route table inetflow detail on page 5220](#)  
[show route table bgp.evpn.0 extensive |no-more \(EVPN\) on page 5223](#)

**Output Fields** Table 297 describes the output fields for the **show route table** command. Output fields are listed in the approximate order in which they appear.

**Table 408: show route table Output Fields**

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>routing-table-name</i>  | Name of the routing table (for example, inet.0).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Restart complete           | <p>All protocols have restarted for this routing table.</p> <p>Restart state:</p> <ul style="list-style-type: none"> <li>• <b>Pending:</b><i>protocol-name</i>—List of protocols that have not yet completed graceful restart for this routing table.</li> <li>• <b>Complete</b>—All protocols have restarted for this routing table.</li> </ul> <p>For example, if the output shows-</p> <ul style="list-style-type: none"> <li>• LDP.inet.0 : 5 routes (4 active, 1 holddown, 0 hidden)<br/>Restart Pending: OSPF LDP VPN</li> </ul> <p>This indicates that <b>OSPF</b>, <b>LDP</b>, and <b>VPN</b> protocols did not restart for <b>LDP.inet.0</b> routing table.</p> <ul style="list-style-type: none"> <li>• vpls_1.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)<br/>Restart Complete</li> </ul> <p>This indicates that all protocols have restarted for <b>vpls_1.l2vpn.0</b> routing table.</p> |
| <i>number destinations</i> | Number of destinations for which there are routes in the routing table.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <i>number routes</i>       | <p>Number of routes in the routing table and total number of routes in the following states:</p> <ul style="list-style-type: none"> <li>• <b>active</b> (routes that are active)</li> <li>• <b>holddown</b> (routes that are in the pending state before being declared inactive)</li> <li>• <b>hidden</b> (routes that are not used because of a routing policy)</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |



Table 408: show route table Output Fields (*continued*)

| Field Name                                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>route-destination</i><br>(entry, announced) | <p>Route destination (for example:10.0.0.1/24). The <b>entry</b> value is the number of routes for this destination, and the <b>announced</b> value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:</p> <ul style="list-style-type: none"> <li>• <b>MPLS-label</b> (for example, 80001).</li> <li>• <b>interface-name</b> (for example, ge-1/0/2).</li> <li>• <b>neighbor-address:control-word-status:encapsulation type:vc-id:source</b> (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96). <ul style="list-style-type: none"> <li>• <b>neighbor-address</b>—Address of the neighbor.</li> <li>• <b>control-word-status</b>—Whether the use of the control word has been negotiated for this virtual circuit: <b>NoCtrlWord</b> or <b>CtrlWord</b>.</li> <li>• <b>encapsulation type</b>—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport.</li> <li>• <b>vc-id</b>—Virtual circuit identifier.</li> <li>• <b>source</b>—Source of the advertisement: <b>Local</b> or <b>Remote</b>.</li> </ul> </li> <li>• <b>inclusive multicast Ethernet tag route</b>—Type of route destination represented by (for example, 3:100.100.100.10:100::0::10::100.100.100.10/384): <ul style="list-style-type: none"> <li>• <b>route distinguisher</b>—(8 octets) Route distinguisher (RD) must be the RD of the EVPN instance (EVI) that is advertising the NLRI.</li> <li>• <b>Ethernet tag ID</b>—(4 octets) Identifier of the Ethernet tag. Can set to 0 or to a valid Ethernet tag value.</li> <li>• <b>IP address length</b>—(1 octet) Length of IP address in bits.</li> <li>• <b>originating router's IP address</b>—(4 or 16 octets) Must set to the provider edge (PE) device's IP address. This address should be common for all EVIs on the PE device, and may be the PE device's loopback address.</li> </ul> </li> </ul> |
| label stacking                                 | <p>(Next-to-the-last-hop routing device for MPLS only) Depth of the MPLS label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels.</p> <ul style="list-style-type: none"> <li>• <b>S=0 route</b> indicates that a packet with an incoming label stack depth of 2 or more exits this routing device with one fewer label (the label-popping operation is performed).</li> <li>• If there is no <b>S=</b> information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| [ <i>protocol, preference</i> ]                | <p>Protocol from which the route was learned and the preference value for the route.</p> <ul style="list-style-type: none"> <li>• <b>+</b>—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table.</li> <li>• <b>-</b>—A hyphen indicates the last active route.</li> <li>• <b>*</b>—An asterisk indicates that the route is both the active and the last active route. An asterisk before a <b>to</b> line indicates the best subpath to the route.</li> </ul> <p>In every routing metric except for the BGP <b>LocalPref</b> attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the <b>LocalPref</b> value in the <b>Preference2</b> field. For example, if the <b>LocalPref</b> value for Route 1 is 100, the <b>Preference2</b> value is -101. If the <b>LocalPref</b> value for Route 2 is 155, the <b>Preference2</b> value is -156. Route 2 is preferred because it has a higher <b>LocalPref</b> value and a lower <b>Preference2</b> value.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

Table 408: show route table Output Fields (*continued*)

| Field Name                                    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-----------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Level                                         | (IS-IS only). In IS-IS, a single AS can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Route Distinguisher                           | IP subnet augmented with a 64-bit prefix.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| PMSI                                          | Provider multicast service interface (MVPN routing table).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Next-hop type                                 | Type of next hop. For a description of possible values for this field, see <a href="#">Table 300</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Next-hop reference count                      | Number of references made to the next hop.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Flood nexthop branches exceed maximum message | Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Source                                        | IP address of the route source.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Next hop                                      | Network layer address of the directly reachable neighboring system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| via                                           | Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word <b>Selected</b> . This field can also contain the following information: <ul style="list-style-type: none"> <li>Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.</li> <li>Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.</li> </ul> |
| Label-switched-path <i>lsp-path-name</i>      | Name of the LSP used to reach the next hop.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Label operation                               | MPLS label and operation occurring at this routing device. The operation can be <b>pop</b> (where a label is removed from the top of the stack), <b>push</b> (where another label is added to the label stack), or <b>swap</b> (where a label is replaced by another label).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Interface                                     | (Local only) Local interface name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Protocol next hop                             | Network layer address of the remote routing device that advertised the prefix. This address is used to derive a forwarding next hop.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Indirect next hop                             | Index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| State                                         | State of the route (a route can be in more than one state). See <a href="#">Table 301</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

Table 408: show route table Output Fields (*continued*)

| Field Name        | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Local AS          | AS number of the local routing device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Age               | How long the route has been known.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| AIGP              | Accumulated interior gateway protocol (AIGP) BGP attribute.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Metric            | Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| MED-plus-IGP      | Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| TTL-Action        | For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signaled and LDP-signaled LSPs or for specific VRF routing instances.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Task              | Name of the protocol that has added the route.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Announcement bits | <p>The number of BGP peers or protocols to which Junos OS has announced this route, followed by the list of the recipients of the announcement. Junos OS can also announce the route to the KRT for installing the route into the Packet Forwarding Engine, to a resolve tree, a L2 VC, or even a VPN. For example, <i>n-Resolve inet</i> indicates that the specified route is used for route resolution for next hops found in the routing table.</p> <ul style="list-style-type: none"> <li><i>n</i>—An index used by Juniper Networks customer support only.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| AS path           | <p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> <li><b>I</b>—IGP.</li> <li><b>E</b>—EGP.</li> <li><b>Recorded</b>—The AS path is recorded by the sample process (sampled).</li> <li><b>?</b>—Incomplete; typically, the AS path was aggregated.</li> </ul> <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> <li><b>[ ]</b>—Brackets enclose the number that precedes the AS path. This number represents the number of ASs present in the AS path, when calculated as defined in RFC 4271. This value is used in the AS-path merge process, as defined in RFC 4893.</li> <li><b>[ ]</b>—If more than one AS number is configured on the routing device, or if AS path prepending is configured, brackets enclose the local AS number associated with the AS path.</li> <li><b>{ }</b>—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order.</li> <li><b>( )</b>—Parentheses enclose a confederation.</li> <li><b>( [ ] )</b>—Parentheses and brackets enclose a confederation set.</li> </ul> <p><b>NOTE:</b> In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p> |

Table 408: show route table Output Fields (*continued*)

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| validation-state        | <p>(BGP-learned routes) Validation status of the route:</p> <ul style="list-style-type: none"> <li>• <b>Invalid</b>—Indicates that the prefix is found, but either the corresponding AS received from the EBGp peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database.</li> <li>• <b>Unknown</b>—Indicates that the prefix is not among the prefixes or prefix ranges in the database.</li> <li>• <b>Unverified</b>—Indicates that the origin of the prefix is not verified against the database. This is because the database got populated and the validation is not called for in the BGP import policy, although origin validation is enabled, or the origin validation is not enabled for the BGP peers.</li> <li>• <b>Valid</b>—Indicates that the prefix and autonomous system pair are found in the database.</li> </ul> |
| FECs bound to route     | Point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Primary Upstream        | When multipoint LDP with multicast-only fast reroute (MoFRR) is configured, the primary upstream path. MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| RPF Nexthops            | When multipoint LDP with MoFRR is configured, the reverse-path forwarding (RPF) next-hop information. Data packets are received from both the primary path and the secondary paths. The redundant packets are discarded at topology merge points due to the RPF checks.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Label                   | Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| weight                  | Value used to distinguish MoFRR primary and backup routes. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| VC Label                | MPLS label assigned to the Layer 2 circuit virtual connection.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| MTU                     | Maximum transmission unit (MTU) of the Layer 2 circuit.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| VLAN ID                 | VLAN identifier of the Layer 2 circuit.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Prefixes bound to route | Forwarding equivalent class (FEC) bound to this route. Applicable only to routes installed by LDP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Communities             | Community path attribute for the route. See <a href="#">Table 302</a> for all possible values for this field.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Layer2-info: encaps     | Layer 2 encapsulation (for example, VPLS).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| control flags           | Control flags: <b>none</b> or <b>Site Down</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| mtu                     | Maximum transmission unit (MTU) information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Label-Base, range       | First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| status vector           | Layer 2 VPN and VPLS network layer reachability information (NLRI).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

Table 408: show route table Output Fields (*continued*)

| Field Name                          | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Accepted Multipath                  | Current active path when BGP multipath is configured.                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Accepted LongLivedStale             | The LongLivedStale flag indicates that the route was marked LLGR-stale by this router, as part of the operation of LLGR receiver mode. Either this flag or the LongLivedStaleImport flag might be displayed for a route. Neither of these flags is displayed at the same time as the Stale (ordinary GR stale) flag.                                                                                                                                                                  |
| Accepted LongLivedStaleImport       | <p>The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy. Either this flag or the LongLivedStale flag might be displayed for a route. Neither of these flags is displayed at the same time as the Stale (ordinary GR stale) flag.</p> <p>Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and import into the inet.0 routing table</p> |
| ImportAccepted LongLivedStaleImport | <p>Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and imported into the inet.0 routing table</p> <p>The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy.</p>                                                                                                                                                                       |
| Accepted MultipathContrib           | Path currently contributing to BGP multipath.                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Localpref                           | Local preference value included in the route.                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Router ID                           | BGP router ID as advertised by the neighbor in the open message.                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Primary Routing Table               | In a routing table group, the name of the primary routing table in which the route resides.                                                                                                                                                                                                                                                                                                                                                                                           |
| Secondary Tables                    | In a routing table group, the name of one or more secondary tables in which the route resides.                                                                                                                                                                                                                                                                                                                                                                                        |

Table 300 describes all possible values for the Next-hop Types output field.

Table 409: Next-hop Types Output Field Values

| Next-Hop Type     | Description                                                                                                                                                                                                                                                                    |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Broadcast (bcast) | Broadcast next hop.                                                                                                                                                                                                                                                            |
| Deny              | Deny next hop.                                                                                                                                                                                                                                                                 |
| Discard           | Discard next hop.                                                                                                                                                                                                                                                              |
| Flood             | Flood next hop. Consists of components called branches, up to a maximum of 32 branches. Each flood next-hop branch sends a copy of the traffic to the forwarding interface. Used by point-to-multipoint RSVP, point-to-multipoint LDP, point-to-multipoint CCC, and multicast. |

Table 409: Next-hop Types Output Field Values (*continued*)

| Next-Hop Type            | Description                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Hold                     | Next hop is waiting to be resolved into a unicast or multicast type.                                                                                                                                                                                                                                                                                                                                       |
| Indexed (idxd)           | Indexed next hop.                                                                                                                                                                                                                                                                                                                                                                                          |
| Indirect (indr)          | Used with applications that have a protocol next hop address that is remote. You are likely to see this next-hop type for internal BGP (IBGP) routes when the BGP next hop is a BGP neighbor that is not directly connected.                                                                                                                                                                               |
| Interface                | Used for a network address assigned to an interface. Unlike the router next hop, the interface next hop does not reference any specific node on the network.                                                                                                                                                                                                                                               |
| Local (locl)             | Local address on an interface. This next-hop type causes packets with this destination address to be received locally.                                                                                                                                                                                                                                                                                     |
| Multicast (mcst)         | Wire multicast next hop (limited to the LAN).                                                                                                                                                                                                                                                                                                                                                              |
| Multicast discard (mdsc) | Multicast discard.                                                                                                                                                                                                                                                                                                                                                                                         |
| Multicast group (mgrp)   | Multicast group member.                                                                                                                                                                                                                                                                                                                                                                                    |
| Receive (recv)           | Receive.                                                                                                                                                                                                                                                                                                                                                                                                   |
| Reject (rjct)            | Discard. An ICMP unreachable message was sent.                                                                                                                                                                                                                                                                                                                                                             |
| Resolve (rslv)           | Resolving next hop.                                                                                                                                                                                                                                                                                                                                                                                        |
| Routed multicast (mcrtr) | Regular multicast next hop.                                                                                                                                                                                                                                                                                                                                                                                |
| Router                   | <p>A specific node or set of nodes to which the routing device forwards packets that match the route prefix.</p> <p>To qualify as next-hop type router, the route must meet the following criteria:</p> <ul style="list-style-type: none"> <li>• Must not be a direct or local subnet for the routing device.</li> <li>• Must have a next hop that is directly connected to the routing device.</li> </ul> |
| Table                    | Routing table next hop.                                                                                                                                                                                                                                                                                                                                                                                    |
| Unicast (ucst)           | Unicast.                                                                                                                                                                                                                                                                                                                                                                                                   |
| Unilist (ulst)           | List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.                                                                                                                                                                                                                                                                                                                |

Table 301 describes all possible values for the State output field. A route can be in more than one state (for example, <Active NoReadvrt Int Ext>).

**Table 410: State Output Field Values**

| Value                                       | Description                                                                                                                                                                          |
|---------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Accounting                                  | Route needs accounting.                                                                                                                                                              |
| Active                                      | Route is active.                                                                                                                                                                     |
| Always Compare MED                          | Path with a lower multiple exit discriminator (MED) is available.                                                                                                                    |
| AS path                                     | Shorter AS path is available.                                                                                                                                                        |
| Cisco Non-deterministic MED selection       | Cisco nondeterministic MED is enabled, and a path with a lower MED is available.                                                                                                     |
| Clone                                       | Route is a clone.                                                                                                                                                                    |
| Cluster list length                         | Length of cluster list sent by the route reflector.                                                                                                                                  |
| Delete                                      | Route has been deleted.                                                                                                                                                              |
| Ex                                          | Exterior route.                                                                                                                                                                      |
| Ext                                         | BGP route received from an external BGP neighbor.                                                                                                                                    |
| FlashAll                                    | Forces all protocols to be notified of a change to any route, active or inactive, for a prefix. When not set, protocols are informed of a prefix only when the active route changes. |
| Hidden                                      | Route not used because of routing policy.                                                                                                                                            |
| IfCheck                                     | Route needs forwarding RPF check.                                                                                                                                                    |
| IGP metric                                  | Path through next hop with lower IGP metric is available.                                                                                                                            |
| Inactive reason                             | Flags for this route, which was not selected as best for a particular destination.                                                                                                   |
| Initial                                     | Route being added.                                                                                                                                                                   |
| Int                                         | Interior route.                                                                                                                                                                      |
| Int Ext                                     | BGP route received from an internal BGP peer or a BGP confederation peer.                                                                                                            |
| Interior > Exterior > Exterior via Interior | Direct, static, IGP, or EBGp path is available.                                                                                                                                      |

Table 410: State Output Field Values (*continued*)

| Value                          | Description                                                                                                                                                                                                                       |
|--------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Local Preference               | Path with a higher local preference value is available.                                                                                                                                                                           |
| Martian                        | Route is a martian (ignored because it is obviously invalid).                                                                                                                                                                     |
| MartianOK                      | Route exempt from martian filtering.                                                                                                                                                                                              |
| Next hop address               | Path with lower metric next hop is available.                                                                                                                                                                                     |
| No difference                  | Path from neighbor with lower IP address is available.                                                                                                                                                                            |
| NoReadvrt                      | Route not to be advertised.                                                                                                                                                                                                       |
| NotBest                        | Route not chosen because it does not have the lowest MED.                                                                                                                                                                         |
| Not Best in its group          | Incoming BGP AS is not the best of a group (only one AS can be the best).                                                                                                                                                         |
| NotInstall                     | Route not to be installed in the forwarding table.                                                                                                                                                                                |
| Number of gateways             | Path with a greater number of next hops is available.                                                                                                                                                                             |
| Origin                         | Path with a lower origin code is available.                                                                                                                                                                                       |
| Pending                        | Route pending because of a hold-down configured on another route.                                                                                                                                                                 |
| Release                        | Route scheduled for release.                                                                                                                                                                                                      |
| RIB preference                 | Route from a higher-numbered routing table is available.                                                                                                                                                                          |
| Route Distinguisher            | 64-bit prefix added to IP subnets to make them unique.                                                                                                                                                                            |
| Route Metric or MED comparison | Route with a lower metric or MED is available.                                                                                                                                                                                    |
| Route Preference               | Route with lower preference value is available.                                                                                                                                                                                   |
| Router ID                      | Path through a neighbor with lower ID is available.                                                                                                                                                                               |
| Secondary                      | Route not a primary route.                                                                                                                                                                                                        |
| Unusable path                  | Path is not usable because of one of the following conditions: <ul style="list-style-type: none"> <li>• The route is damped.</li> <li>• The route is rejected by an import policy.</li> <li>• The route is unresolved.</li> </ul> |
| Update source                  | Last tiebreaker is the lowest IP address value.                                                                                                                                                                                   |



Table 302 describes the possible values for the Communities output field.

Table 411: Communities Output Field Values

| Value                                                          | Description                                                                                                                                                                                                                                                                           |
|----------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>area-number</i>                                             | 4 bytes, encoding a 32-bit area number. For AS-external routes, the value is 0. A nonzero value identifies the route as internal to the OSPF domain, and as within the identified area. Area numbers are relative to a particular OSPF domain.                                        |
| <b>bandwidth:</b> <i>local AS number:link-bandwidth-number</i> | Link-bandwidth community value used for unequal-cost load balancing. When BGP has several candidate paths available for multipath purposes, it does not perform unequal-cost load balancing according to the link-bandwidth community unless all candidate paths have this attribute. |
| <b>domain-id</b>                                               | Unique configurable number that identifies the OSPF domain.                                                                                                                                                                                                                           |
| <b>domain-id-vendor</b>                                        | Unique configurable number that further identifies the OSPF domain.                                                                                                                                                                                                                   |
| <i>link-bandwidth-number</i>                                   | Link-bandwidth number: from 0 through 4,294,967,295 (bytes per second).                                                                                                                                                                                                               |
| <i>local AS number</i>                                         | Local AS number: from 1 through 65,535.                                                                                                                                                                                                                                               |
| <i>options</i>                                                 | 1 byte. Currently this is only used if the route type is 5 or 7. Setting the least significant bit in the field indicates that the route carries a type 2 metric.                                                                                                                     |
| <b>origin</b>                                                  | (Used with VPNs) Identifies where the route came from.                                                                                                                                                                                                                                |
| <i>ospf-route-type</i>                                         | 1 byte, encoded as 1 or 2 for intra-area routes (depending on whether the route came from a type 1 or a type 2 LSA); 3 for summary routes; 5 for external routes (area number must be 0); 7 for NSSA routes; or 129 for sham link endpoint addresses.                                 |
| <b>route-type-vendor</b>                                       | Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute 0x8000. The format is <b>area-number:ospf-route-type:options</b> .                                                                                  |
| <b>rte-type</b>                                                | Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute 0x0306. The format is <b>area-number:ospf-route-type:options</b> .                                                                                  |
| <b>target</b>                                                  | Defines which VPN the route participates in; <b>target</b> has the format <b>32-bit IP address:16-bit number</b> . For example, 10.19.0.0:100.                                                                                                                                        |
| <b>unknown IANA</b>                                            | Incoming IANA codes with a value between 0x1 and 0x7fff. This code of the BGP extended community attribute is accepted, but it is not recognized.                                                                                                                                     |
| <b>unknown OSPF vendor community</b>                           | Incoming IANA codes with a value above 0x8000. This code of the BGP extended community attribute is accepted, but it is not recognized.                                                                                                                                               |

## Sample Output

### show route table bgp.l2vpn

```
user@host> show route table bgp.l2vpn
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.24.1:1:4:1/96
 *[BGP/170] 01:08:58, localpref 100, from 192.168.24.1
 AS path: I
 > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
```

### show route table bgp.l3vpn.0

```
user@host> show route table bgp.l3vpn.0
bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.71.15:100:10.255.71.17/32
 *[BGP/170] 00:03:59, MED 1, localpref 100, from
10.255.71.15
 AS path: I
 > via so-2/1/0.0, Push 100020, Push 100011(top)
10.255.71.15:200:10.255.71.18/32
 *[BGP/170] 00:03:59, MED 1, localpref 100, from
10.255.71.15
 AS path: I
 > via so-2/1/0.0, Push 100021, Push 100011(top)
```

### show route table bgp.l3vpn.0 detail

```
user@host> show route table bgp.l3vpn.0 detail
bgp.l3vpn.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)

10.255.245.12:1:4.0.0.0/8 (1 entry, 1 announced)
 *BGP Preference: 170/-101
 Route Distinguisher: 10.255.245.12:1
 Source: 10.255.245.12
 Next hop: 192.168.208.66 via fe-0/0/0.0, selected
 Label operation: Push 182449
 Protocol next hop: 10.255.245.12
 Push 182449
 Indirect next hop: 863a630 297
 State: <Active Int Ext>
 Local AS: 35 Peer AS: 35
 Age: 12:19 Metric2: 1
 Task: BGP_35.10.255.245.12+179
 Announcement bits (1): 0-BGP.0.0.0.0+179
 AS path: 30 10458 14203 2914 3356 I (Atomic) Aggregator: 3356 4.68.0.11

 Communities: 2914:420 target:11111:1 origin:56:78
 VPN Label: 182449
 Localpref: 100
 Router ID: 10.255.245.12

10.255.245.12:1:4.17.225.0/24 (1 entry, 1 announced)
 *BGP Preference: 170/-101
 Route Distinguisher: 10.255.245.12:1
 Source: 10.255.245.12
 Next hop: 192.168.208.66 via fe-0/0/0.0, selected
```

```

Label operation: Push 182465
Protocol next hop: 10.255.245.12
Push 182465
Indirect next hop: 863a8f0 305
State: <Active Int Ext>
Local AS: 35 Peer AS: 35
Age: 12:19 Metric2: 1
Task: BGP_35.10.255.245.12+179
Announcement bits (1): 0-BGP.0.0.0.0+179
AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496 6496 I
Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
VPN Label: 182465
Localpref: 100
Router ID: 10.255.245.12

10.255.245.12:1:4.17.226.0/23 (1 entry, 1 announced)
*BGP Preference: 170/-101
Route Distinguisher: 10.255.245.12:1
Source: 10.255.245.12
Next hop: 192.168.208.66 via fe-0/0/0.0, selected
Label operation: Push 182465
Protocol next hop: 10.255.245.12
Push 182465
Indirect next hop: 86bd210 330
State: <Active Int Ext>
Local AS: 35 Peer AS: 35
Age: 12:19 Metric2: 1
Task: BGP_35.10.255.245.12+179
Announcement bits (1): 0-BGP.0.0.0.0+179
AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496
6496 I
Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
VPN Label: 182465
Localpref: 100
Router ID: 10.255.245.12

10.255.245.12:1:4.17.251.0/24 (1 entry, 1 announced)
*BGP Preference: 170/-101
Route Distinguisher: 10.255.245.12:1
Source: 10.255.245.12
Next hop: 192.168.208.66 via fe-0/0/0.0, selected
Label operation: Push 182465
Protocol next hop: 10.255.245.12
Push 182465
Indirect next hop: 86bd210 330
State: <Active Int Ext>
Local AS: 35 Peer AS: 35
Age: 12:19 Metric2: 1
Task: BGP_35.10.255.245.12+179
Announcement bits (1): 0-BGP.0.0.0.0+179
AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496
6496 I
Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
VPN Label: 182465
Localpref: 100

```

#### show route table bgp.rtarget.0 (When Proxy BGP Route Target Filtering Is Configured)

```
user@host> show route table bgp.rtarget.0
```

```

bgp.rtarget.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

```

```

100:100:100/96
 * [RTarget/5] 00:03:14
 Type Proxy
 for 10.255.165.103
 for 10.255.166.124
 Local

```

### show route table bgp.evpn.0

```

user@host> show route table bgp.evpn.0
bgp.evpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2:100.100.100.2:100::0::00:26:88:5f:67:b0/304
 * [BGP/170] 11:00:05, localpref 100, from 100.100.100.2
 AS path: I, validation-state: unverified
 > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
2:100.100.100.2:100::0::00:51:51:51:51:51/304
 * [BGP/170] 11:00:05, localpref 100, from 100.100.100.2
 AS path: I, validation-state: unverified
 > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
2:100.100.100.3:100::0::00:52:52:52:52:52/304
 * [BGP/170] 10:59:58, localpref 100, from 100.100.100.3
 AS path: I, validation-state: unverified
 > to 100.1.13.3 via ge-2/0/8.0, label-switched-path R0toR2
2:100.100.100.3:100::0::a8:d0:e5:5b:01:c8/304
 * [BGP/170] 10:59:58, localpref 100, from 100.100.100.3
 AS path: I, validation-state: unverified
 > to 100.1.13.3 via ge-2/0/8.0, label-switched-path R0toR2
3:100.100.100.2:100::1000::100.100.100.2/304
 * [BGP/170] 11:00:16, localpref 100, from 100.100.100.2
 AS path: I, validation-state: unverified
 > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
3:100.100.100.2:100::2000::100.100.100.2/304
 * [BGP/170] 11:00:16, localpref 100, from 100.100.100.2
 AS path: I, validation-state: unverified
 > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1

```

### show route table evpna.evpn.0

```

user@host> show route table evpna.evpn.0
evpna.evpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

3:100.100.100.10:100::0::10::100.100.100.10/384
 * [EVPN/170] 01:37:09
 Indirect
3:100.100.100.2:100::2000::100.100.100.2/304
 * [EVPN/170] 01:37:12
 Indirect

```

### show route table inet.0

```

user@host> show route table inet.0
inet.0: 12 destinations, 12 routes (11 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0
 * [Static/5] 00:51:57
 > to 111.222.5.254 via fxp0.0

```

```

1.0.0.1/32 *[Direct/0] 00:51:58
 > via at-5/3/0.0
1.0.0.2/32 *[Local/0] 00:51:58
 Local
12.12.12.21/32 *[Local/0] 00:51:57
 Reject
13.13.13.13/32 *[Direct/0] 00:51:58
 > via t3-5/2/1.0
13.13.13.14/32 *[Local/0] 00:51:58
 Local
13.13.13.21/32 *[Local/0] 00:51:58
 Local
13.13.13.22/32 *[Direct/0] 00:33:59
 > via t3-5/2/0.0
127.0.0.1/32 [Direct/0] 00:51:58
 > via lo0.0
111.222.5.0/24 *[Direct/0] 00:51:58
 > via fxp0.0
111.222.5.81/32 *[Local/0] 00:51:58
 Local

```

### show route table inet.3

```

user@host> show route table inet.3
inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

22.0.0.5/32 *[LDP/9] 00:25:43, metric 10, tag 200
 to 1.2.94.2 via lt-1/2/0.49
 > to 1.2.3.2 via lt-1/2/0.23

```

### show route table inet6.0

```

user@host> show route table inet6.0
inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Route, * = Both

fec0:0:0:3::/64 *[Direct/0] 00:01:34
>via fe-0/1/0.0

fec0:0:0:3::/128 *[Local/0] 00:01:34
>Local

fec0:0:0:4::/64 *[Static/5] 00:01:34
>to fec0:0:0:3::ffff via fe-0/1/0.0

```

### show route table inet6.3

```

user@router> show route table inet6.3
inet6.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

::10.255.245.195/128
 *[LDP/9] 00:00:22, metric 1
 > via so-1/0/0.0
::10.255.245.196/128
 *[LDP/9] 00:00:08, metric 1
 > via so-1/0/0.0, Push 100008

```

**show route table inetflow detail**

```

user@host> show route table inetflow detail
inetflow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.12.44.1,*/48 (1 entry, 1 announced)
 *BGP Preference: 170/-101
 Next-hop reference count: 2
 State: <Active Ext>
 Local AS: 65002 Peer AS: 65000
 Age: 4
 Task: BGP_65000.10.12.99.5+3792
 Announcement bits (1): 0-Flow
 AS path: 65000 I
 Communities: traffic-rate:0:0
 Validation state: Accept, Originator: 10.12.99.5
 Via: 10.12.44.0/24, Active
 Localpref: 100
 Router ID: 10.255.71.161

10.12.56.1,*/48 (1 entry, 1 announced)
 *Flow Preference: 5
 Next-hop reference count: 2
 State: <Active>
 Local AS: 65002
 Age: 6:30
 Task: RT Flow
 Announcement bits (2): 0-Flow 1-BGP.0.0.0.0+179
 AS path: I
 Communities: 1:1

```

**show route table l2circuit.0**

```

user@host> show route table l2circuit.0
l2circuit.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.1.1.195:NoCtrlWord:1:1:Local/96
 *[L2CKT/7] 00:50:47
 > via so-0/1/2.0, Push 100049
 via so-0/1/3.0, Push 100049
10.1.1.195:NoCtrlWord:1:1:Remote/96
 *[LDP/9] 00:50:14
 Discard
10.1.1.195:CtrlWord:1:2:Local/96
 *[L2CKT/7] 00:50:47
 > via so-0/1/2.0, Push 100049
 via so-0/1/3.0, Push 100049
10.1.1.195:CtrlWord:1:2:Remote/96
 *[LDP/9] 00:50:14
 Discard

```

**show route table mpls**

```

user@host> show route table mpls
mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0 *[MPLS/0] 00:13:55, metric 1
 Receive
1 *[MPLS/0] 00:13:55, metric 1
 Receive

```

```

2 *[MPLS/0] 00:13:55, metric 1
 Receive
1024 *[VPN/0] 00:04:18
 to table red.inet.0, Pop

```

### show route table mpls extensive

```

user@host> show route table mpls extensive
100000 (1 entry, 1 announced)
TSI:
KRT in-kernel 100000 /36 -> {so-1/0/0.0}
 *LDP Preference: 9
 Next hop: via so-1/0/0.0, selected
 Pop
 State: <Active Int>
 Age: 29:50 Metric: 1
 Task: LDP
 Announcement bits (1): 0-KRT
 AS path: I
 Prefixes bound to route: 10.0.0.194/32

```

### show route table mpls.0

```

user@host> show route table mpls.0
mpls.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0 *[MPLS/0] 00:45:09, metric 1
 Receive
1 *[MPLS/0] 00:45:09, metric 1
 Receive
2 *[MPLS/0] 00:45:09, metric 1
 Receive
100000 *[L2VPN/7] 00:43:04
 > via so-0/1/0.1, Pop
100001 *[L2VPN/7] 00:43:03
 > via so-0/1/0.2, Pop Offset: 4
100002 *[LDP/9] 00:43:22, metric 1
 via so-0/1/2.0, Pop
 > via so-0/1/3.0, Pop
100002(S=0) *[LDP/9] 00:43:22, metric 1
 via so-0/1/2.0, Pop
 > via so-0/1/3.0, Pop
100003 *[LDP/9] 00:43:22, metric 1
 > via so-0/1/2.0, Swap 100002
 via so-0/1/3.0, Swap 100002
100004 *[LDP/9] 00:43:16, metric 1
 via so-0/1/2.0, Swap 100049
 > via so-0/1/3.0, Swap 100049
so-0/1/0.1 *[L2VPN/7] 00:43:04
 > via so-0/1/2.0, Push 100001, Push 100049(top)
 via so-0/1/3.0, Push 100001, Push 100049(top)
so-0/1/0.2 *[L2VPN/7] 00:43:03
 via so-0/1/2.0, Push 100000, Push 100049(top) Offset: -4
 > via so-0/1/3.0, Push 100000, Push 100049(top) Offset: -4

```

### show route table mpls.0 detail (PTX Series)

```

user@host> show route table mpls.0 detail
ge-0/0/2.600 (1 entry, 1 announced)
 *L2VPN Preference: 7
 Next hop type: Indirect

```

```

Address: 0x9438f34
Next-hop reference count: 2
Next hop type: Router, Next hop index: 567
Next hop: 3.0.0.1 via ge-0/0/1.0, selected
Label operation: Push 299808
Label TTL action: prop-ttl
Load balance label: Label 299808:None;
Session Id: 0x1
Protocol next hop: 10.255.255.1
Label operation: Push 299872 Offset: 252
Label TTL action: no-prop-ttl
Load balance label: Label 299872:Flow label PUSH;
Composite next hop: 0x9438ed8 570 INH Session ID: 0x2
Indirect next hop: 0x9448208 262142 INH Session ID: 0x2
State: <Active Int>
Age: 21 Metric2: 1
Validation State: unverified
Task: Common L2 VC
Announcement bits (2): 0-KRT 2-Common L2 VC
AS path: I

```

#### show route table mpls.0 extensive (PTX Series)

```

user@host> show route table mpls.0 extensive
ge-0/0/2.600 (1 entry, 1 announced)
TSI:
KRT in-kernel ge-0/0/2.600.0 /32 -> {composite(570)}
 *L2VPN Preference: 7
 Next hop type: Indirect
 Address: 0x9438f34
 Next-hop reference count: 2
 Next hop type: Router, Next hop index: 567
 Next hop: 3.0.0.1 via ge-0/0/1.0, selected
 Label operation: Push 299808
 Label TTL action: prop-ttl
 Load balance label: Label 299808:None;
 Session Id: 0x1
 Protocol next hop: 10.255.255.1
 Label operation: Push 299872 Offset: 252
 Label TTL action: no-prop-ttl
 Load balance label: Label 299872:Flow label PUSH;
 Composite next hop: 0x9438ed8 570 INH Session ID: 0x2
 Indirect next hop: 0x9448208 262142 INH Session ID: 0x2
 State: <Active Int>
 Age: 47 Metric2: 1
 Validation State: unverified
 Task: Common L2 VC
 Announcement bits (2): 0-KRT 2-Common L2 VC
 AS path: I
 Composite next hops: 1
 Protocol next hop: 10.255.255.1 Metric: 1
 Label operation: Push 299872 Offset: 252
 Label TTL action: no-prop-ttl
 Load balance label: Label 299872:Flow label PUSH;
 Composite next hop: 0x9438ed8 570 INH Session ID: 0x2
 Indirect next hop: 0x9448208 262142 INH Session ID: 0x2
 Indirect path forwarding next hops: 1
 Next hop type: Router
 Next hop: 3.0.0.1 via ge-0/0/1.0
 Session Id: 0x1
 10.255.255.1/32 Originating RIB: inet.3

```



```

Metric: 1
Forwarding nexthops: 1
Node path count: 1
Nexthop: 3.0.0.1 via ge-0/0/1.0

```

### show route table mpls.0 (RSVP Route—Transit LSP)

In the sample output, the 1 in [RSVP/7/1] indicates the secondary preference value. The secondary preference value becomes significant when multiple RSVP LSPs of different types are signaled to the destination. The possible values of RSVP secondary preferences are:

1—Normal Point-to-Point RSVP-TE LSP

2—Point-to-Multipoint (P2MP) RSVP-TE LSP

3—Dynamic RSVP-TE LSP

```
user@host> show route table mpls.0
```

```

mpls.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

```

```

0 *[MPLS/0] 00:37:31, metric 1
 Receive
1 *[MPLS/0] 00:37:31, metric 1
 Receive
2 *[MPLS/0] 00:37:31, metric 1
 Receive
13 *[MPLS/0] 00:37:31, metric 1
 Receive
300352 *[RSVP/7/1] 00:08:00, metric 1
 > to 8.64.0.106 via ge-1/0/1.0, label-switched-path lsp1_p2p
300352(S=0) *[RSVP/7/1] 00:08:00, metric 1
 > to 8.64.0.106 via ge-1/0/1.0, label-switched-path lsp1_p2p
300384 *[RSVP/7/2] 00:05:20, metric 1
 > to 8.64.1.106 via ge-1/0/0.0, Pop
300384(S=0) *[RSVP/7/2] 00:05:20, metric 1
 > to 8.64.1.106 via ge-1/0/0.0, Pop

```

### show route table vpls\_1 detail

```
user@host> show route table vpls_1 detail
```

```

vpls_1.12vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

```

```

1.1.1.11:1000:1:1/96 (1 entry, 1 announced)
*L2VPN Preference: 170/-1
Receive table: vpls_1.12vpn.0
Next-hop reference count: 2
State: <Active Int Ext>
Age: 4:29:47 Metric2: 1
Task: vpls_1-12vpn
Announcement bits (1): 1-BGP.0.0.0.0+179
AS path: I
Communities: Layer2-info: encaps:VPLS, control flags:Site-Down
Label-base: 800000, range: 8, status-vector: 0xFF

```

### show route table vpn-a

```
user@host> show route table vpn-a
```

```

vpn-a.12vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
192.168.16.1:1:1:1/96
 *[VPN/7] 05:48:27
 Discard
192.168.24.1:1:2:1/96
 *[BGP/170] 00:02:53, localpref 100, from 192.168.24.1
 AS path: I
 > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
192.168.24.1:1:3:1/96
 *[BGP/170] 00:02:53, localpref 100, from 192.168.24.1
 AS path: I
 > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am

```

#### show route table vpn-a.mdt.0

```

user@host> show route table vpn-a.mdt.0
vpn-a.mdt.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:1:0:10.255.14.216:232.1.1.1/144
 *[MVPN/70] 01:23:05, metric2 1
 Indirect
1:1:1:10.255.14.218:232.1.1.1/144
 *[BGP/170] 00:57:49, localpref 100, from 10.255.14.218
 AS path: I
 > via so-0/0/0.0, label-switched-path r0e-to-r1
1:1:2:10.255.14.217:232.1.1.1/144
 *[BGP/170] 00:57:49, localpref 100, from 10.255.14.217
 AS path: I
 > via so-0/0/1.0, label-switched-path r0-to-r2

```

#### show route table VPN-A detail

```

user@host> show route table VPN-A detail
VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
10.255.179.9/32 (1 entry, 1 announced)
 *BGP Preference: 170/-101
 Route Distinguisher: 10.255.179.13:200
 Next hop type: Indirect
 Next-hop reference count: 5
 Source: 10.255.179.13
 Next hop type: Router, Next hop index: 732
 Next hop: 10.39.1.14 via fe-0/3/0.0, selected
 Label operation: Push 299824, Push 299824(top)
 Protocol next hop: 10.255.179.13
 Push 299824
 Indirect next hop: 8f275a0 1048574
 State: (Secondary Active Int Ext)
 Local AS: 1 Peer AS: 1
 Age: 3:41:06 Metric: 1 Metric2: 1
 Task: BGP_1.10.255.179.13+64309
 Announcement bits (2): 0-KRT 1-BGP RT Background
 AS path: I
 Communities: target:1:200 rte-type:0.0.0.0:1:0
 Import Accepted
 VPN Label: 299824 TTL Action: vrf-ttl-propagate
 Localpref: 100
 Router ID: 10.255.179.13
 Primary Routing Table bgp.13vpn.0

```

**show route table VPN-AB.inet.0**

```

user@host> show route table VPN-AB.inet.0
VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.0/30 *[OSPF/10] 00:07:24, metric 1
 > via so-7/3/1.0
10.39.1.4/30 *[Direct/0] 00:08:42
 > via so-5/1/0.0
10.39.1.6/32 *[Local/0] 00:08:46
 Local
10.255.71.16/32 *[Static/5] 00:07:24
 > via so-2/0/0.0
10.255.71.17/32 *[BGP/170] 00:07:24, MED 1, localpref 100, from
10.255.71.15
 AS path: I
 > via so-2/1/0.0, Push 100020, Push 100011(top)
10.255.71.18/32 *[BGP/170] 00:07:24, MED 1, localpref 100, from
10.255.71.15
 AS path: I
 > via so-2/1/0.0, Push 100021, Push 100011(top)
10.255.245.245/32 *[BGP/170] 00:08:35, localpref 100
 AS path: 2 I
 > to 10.39.1.5 via so-5/1/0.0
10.255.245.246/32 *[OSPF/10] 00:07:24, metric 1
 > via so-7/3/1.0

```

**show route table VPN\_blue.mvpn-inet6.0**

```

user@host> show route table VPN_blue.mvpn-inet6.0
vpn_blue.mvpn-inet6.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:10.255.2.202:65535:10.255.2.202/432
 *[BGP/170] 00:02:37, localpref 100, from 10.255.2.202
 AS path: I
 > via so-0/1/3.0
1:10.255.2.203:65535:10.255.2.203/432
 *[BGP/170] 00:02:37, localpref 100, from 10.255.2.203
 AS path: I
 > via so-0/1/0.0
1:10.255.2.204:65535:10.255.2.204/432
 *[MVPN/70] 00:57:23, metric2 1
 Indirect
5:10.255.2.202:65535:128::192.168.90.2:128:ffff::1/432
 *[BGP/170] 00:02:37, localpref 100, from 10.255.2.202
 AS path: I
 > via so-0/1/3.0
6:10.255.2.203:65535:65000:128::10.12.53.12:128:ffff::1/432
 *[PIM/105] 00:02:37
 Multicast (IPv6)
7:10.255.2.202:65535:65000:128::192.168.90.2:128:ffff::1/432
 *[MVPN/70] 00:02:37, metric2 1
 Indirect

```

**show route table vrf1.mvpn.0 extensive**

```

user@host> show route table vrf1.mvpn.0 extensive
1:10.255.50.77:1:10.255.50.77/240 (1 entry, 1 announced)
 *MVPN Preference: 70

```

```

PMSI: Flags 0x0: Label 0: RSVP-TE:
Session_13[10.255.50.77:0:25624:10.255.50.77]
 Next hop type: Indirect
 Address: 0xbb2c944
 Next-hop reference count: 360
 Protocol next hop: 10.255.50.77
 Indirect next hop: 0x0 - INH Session ID: 0x0
 State: <Active Int Ext>
 Age: 53:03 Metric2: 1
 Validation State: unverified
 Task: mvpn global task
 Announcement bits (3): 0-PIM.vrf1 1-mvpn global task 2-rt-export

AS path: I

```

### show route table MVPN.mvpn.0

Starting in Junos OS Release 15.1, multicast routes on the locally originated type 7 customer multicast routes are added exclusively by PIM. The functionality of the BGP-MVPN service (which, internally, depends on contributions of state from both the MVPN and PIM protocol components of Junos OS) remains unchanged. MVPN, however, no longer appears as the originator of the locally advertised route. Routes advertised by remote PEs are, as usual, always learned locally from their respective [BGP/...] protocol.

```

user@host> show route table MVPN.mvpn.0
MVPN.mvpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

7:10.255.2.202:65535:65000:128:::192.168.90.2:128:ffff::1/432
 *[PIM/70] 00:02:37, metric2 1
 Indirect
5:100:32:192.168.1.9:32:239.1.1.1/240
 *[PIM/105] 01:51:21
 Multicast (IPv4)
7:100:1:100.32.192.168.5:32:237.1.1.1/240
 *[PIM/105] 01:51:21
 Multicast (IPv4)

```

### show route table inetflow detail

```

user@host> show route table inetflow detail
inetflow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.12.44.1,*/48 (1 entry, 1 announced)
 *BGP Preference: 170/-101
 Next-hop reference count: 2
 State: <Active Ext>
 Local AS: 65002 Peer AS: 65000
 Age: 4
 Task: BGP_65000.10.12.99.5+3792
 Announcement bits (1): 0-Flow
 AS path: 65000 I
 Communities: traffic-rate:0:0
 Validation state: Accept, Originator: 10.12.99.5
 Via: 10.12.44.0/24, Active
 Localpref: 100
 Router ID: 10.255.71.161

10.12.56.1,*/48 (1 entry, 1 announced)
 *Flow Preference: 5
 Next-hop reference count: 2

```

```

State: <Active>
Local AS: 65002
Age: 6:30
Task: RT Flow
Announcement bits (2): 0-Flow 1-BGP.0.0.0+179
AS path: I
Communities: 1:1

user@PE1> show route table green.l2vpn.0 (VPLS Multihoming with FEC 129)
green.l2vpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.2:100:1.1.1.2/96 AD
 *[VPLS/170] 1d 03:11:03, metric2 1
 Indirect
1.1.1.4:100:1.1.1.4/96 AD
 *[BGP/170] 1d 03:11:02, localpref 100, from 1.1.1.4
 AS path: I, validation-state: unverified
 > via ge-1/2/1.5
1.1.1.2:100:1:0/96 MH
 *[VPLS/170] 1d 03:11:03, metric2 1
 Indirect
1.1.1.4:100:1:0/96 MH
 *[BGP/170] 1d 03:11:02, localpref 100, from 1.1.1.4
 AS path: I, validation-state: unverified
 > via ge-1/2/1.5
1.1.1.4:NoCtrlWord:5:100:100:1.1.1.2:1.1.1.4/176
 *[VPLS/7] 1d 03:11:02, metric2 1
 > via ge-1/2/1.5
1.1.1.4:NoCtrlWord:5:100:100:1.1.1.4:1.1.1.2/176
 *[LDP/9] 1d 03:11:02
 Discard

user@host> show route table red extensive
red.inet.0: 364481 destinations, 714087 routes (364480 active, 48448 holddown, 1
hidden)
22.0.0.0/32 (3 entries, 1 announced)
 State: <OnList CalcForwarding>
TSI:
KRT in-kernel 22.0.0.0/32 -> {composite(1048575)} Page 0 idx 1 Type 1 val 0x934342c

 Nexthop: Self
 AS path: [2] I
 Communities: target:2:1
Path 22.0.0.0 from 2.3.0.0 Vector len 4. Val: 1
 @BGP Preference: 170/-1
 Route Distinguisher: 2:1
 Next hop type: Indirect
 Address: 0x258059e4
 Next-hop reference count: 2
 Source: 2.2.0.0
 Next hop type: Router
 Next hop: 10.1.1.1 via ge-1/1/9.0, selected
 Label operation: Push 707633
 Label TTL action: prop-ttl
 Session Id: 0x17d8
 Protocol next hop: 2.2.0.0
 Push 16
 Composite next hop: 0x25805988 - INH Session ID: 0x193c
 Indirect next hop: 0x23eea900 - INH Session ID: 0x193c
 State: <Secondary Active Int Ext ProtectionPath ProtectionCand>

```

```

Local AS: 2 Peer AS: 2
Age: 23 Metric2: 35
Validation State: unverified
Task: BGP_2.2.2.0.0+34549
AS path: I
Communities: target:2:1
Import Accepted
VPN Label: 16
Localpref: 0
Router ID: 2.2.0.0
Primary Routing Table bgp.13vpn.0
Composite next hops: 1
 Protocol next hop: 2.2.0.0 Metric: 35
 Push 16
 Composite next hop: 0x25805988 - INH Session ID: 0x193c
 Indirect next hop: 0x23eea900 - INH Session ID: 0x193c
 Indirect path forwarding next hops: 1
 Next hop type: Router
 Next hop: 10.1.1.1 via ge-1/1/9.0
 Session Id: 0x17d8
 2.2.0.0/32 Originating RIB: inet.3
 Metric: 35 Node path count: 1
 Forwarding nexthops: 1
 Nexthop: 10.1.1.1 via ge-1/1/9.0
BGP Preference: 170/-1
Route Distinguisher: 2:1
Next hop type: Indirect
Address: 0x9347028
Next-hop reference count: 3
Source: 2.3.0.0
Next hop type: Router, Next hop index: 702
Next hop: 10.1.4.2 via ge-1/0/0.0, selected
Label operation: Push 634278
Label TTL action: prop-ttl
Session Id: 0x17d9
Protocol next hop: 2.3.0.0
Push 16
Composite next hop: 0x93463a0 1048575 INH Session ID: 0x17da
Indirect next hop: 0x91e8800 1048574 INH Session ID: 0x17da
State: <Secondary NotBest Int Ext ProtectionPath ProtectionCand>

Inactive reason: Not Best in its group - IGP metric
Local AS: 2 Peer AS: 2
Age: 3:34 Metric2: 70
Validation State: unverified
Task: BGP_2.2.3.0.0+32805
Announcement bits (2): 0-KRT 1-BGP_RT_Background
AS path: I
Communities: target:2:1
Import Accepted
VPN Label: 16
Localpref: 0
Router ID: 2.3.0.0
Primary Routing Table bgp.13vpn.0
Composite next hops: 1
 Protocol next hop: 2.3.0.0 Metric: 70
 Push 16
 Composite next hop: 0x93463a0 1048575 INH Session ID:
0x17da
 Indirect next hop: 0x91e8800 1048574 INH Session ID:
0x17da

```

```

 Indirect path forwarding next hops: 1
 Next hop type: Router
 Next hop: 10.1.4.2 via ge-1/0/0.0
 Session Id: 0x17d9
 2.3.0.0/32 Originating RIB: inet.3
 Metric: 70
 Node path count: 1
 Forwarding nexthops: 1
 Nexthop: 10.1.4.2 via ge-1/0/0.0
#Multipath Preference: 255
 Next hop type: Indirect
 Address: 0x24afca30
 Next-hop reference count: 1
 Next hop type: Router
 Next hop: 10.1.1.1 via ge-1/1/9.0, selected
 Label operation: Push 707633
 Label TTL action: prop-ttl
 Session Id: 0x17d8
 Next hop type: Router, Next hop index: 702
 Next hop: 10.1.4.2 via ge-1/0/0.0
 Label operation: Push 634278
 Label TTL action: prop-ttl
 Session Id: 0x17d9
 Protocol next hop: 2.2.0.0
 Push 16
 Composite next hop: 0x25805988 - INH Session ID: 0x193c
 Indirect next hop: 0x23eea900 - INH Session ID: 0x193c Weight 0x1

 Protocol next hop: 2.3.0.0
 Push 16
 Composite next hop: 0x93463a0 1048575 INH Session ID: 0x17da
 Indirect next hop: 0x91e8800 1048574 INH Session ID: 0x17da Weight

0x4000
 State: <ForwardingOnly Int Ext>
 Inactive reason: Forwarding use only
 Age: 23
 Metric2: 35
 Validation State: unverified
 Task: RT
 AS path: I
 Communities: target:2:1

```

### show route table bgp.evpn.0 extensive |no-more (EVPN)

```

show route table bgp.evpn.0 extensive | no-more
bgp.evpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
2:1000:10::100::00:aa:aa:aa:aa:aa/304 (1 entry, 0 announced)
 *BGP
 Preference: 170/-101
 Route Distinguisher: 1000:10
 Next hop type: Indirect
 Address: 0x9420fd0
 Next-hop reference count: 12
 Source: 1.2.3.4
 Protocol next hop: 1.2.3.4
 Indirect next hop: 0x2 no-forward INH Session ID: 0x0
 State: Local AS: 17 Peer AS:17 Age:21:12 Metric2:1 Validation State:
unverified
 Task: BGP_17.1.2.3.4+50756
 AS path: I
 Communities: target:1111:8388708 encapsulation0:0:0:0:3
 Import Accepted
 Route Label: 100
 ESI: 00:00:00:00:00:00:00:00:00:00

```

```

Localpref: 100
Router ID: 1.2.3.4
Secondary Tables: default-switch.evpn.0
Indirect next hops: 1
 Protocol next hop: 1.2.3.4 Metric: 1
 Indirect next hop: 0x2 no-forward INH Session ID: 0x0
 Indirect path forwarding next hops: 1
 Next hop type: Router
 Next hop: 10.10.10.1 via xe-0/0/1.0
 Session Id: 0x2
 1.2.3.4/32 Originating RIB: inet.0
 Metric: 1 Node path count: 1
 Forwarding nexthops: 2
 Nexthop: 10.92.78.102 via em0.0

2:1000:10::200::00:bb:bb:bb:bb:bb/304 (1 entry, 0 announced)
 *BGP Preference: 170/-101
 Route Distinguisher: 1000:10
 Next hop type: Indirect
 Address: 0x9420fd0
 Next-hop reference count: 12
 Source: 1.2.3.4
 Protocol next hop: 1.2.3.4
 Indirect next hop: 0x2 no-forward INH Session ID: 0x0
 State: Local AS:17 Peer AS:17 Age:19:43 Metric2:1 Validation
State:unverified
 Task: BGP_17.1.2.3.4+50756
 AS path: I
 Communities: target:2222:22 encapsulation0:0:0:0:3
 Import Accepted
 Route Label: 200
 ESI: 00:00:00:00:00:00:00:00:00:00
 Localpref: 100
 Router ID: 1.2.3.4
 Secondary Tables: default-switch.evpn.0
 Indirect next hops: 1
 Protocol next hop: 1.2.3.4 Metric: 1
 Indirect next hop: 0x2 no-forward INH Session ID: 0x0
 Indirect path forwarding next hops: 1
 Next hop type: Router
 Next hop: 10.10.10.1 via xe-0/0/1.0
 Session Id: 0x2
 1.2.3.4/32 Originating RIB: inet.0
 Metric: 1 Node path count: 1
 Forwarding nexthops: 2
 Nexthop: 10.92.78.102 via em0.0

2:1000:10::300::00:cc:cc:cc:cc:cc/304 (1 entry, 0 announced)
 *BGP Preference: 170/-101
 Route Distinguisher: 1000:10
 Next hop type: Indirect
 Address: 0x9420fd0
 Next-hop reference count: 12
 Source: 1.2.3.4
 Protocol next hop: 1.2.3.4
 Indirect next hop: 0x2 no-forward INH Session ID: 0x0
 State: Local AS:17 Peer AS:17 Age:17:21 Metric2:1 Validation State:
unverified Task: BGP 17,1,2,3,4+50756
 AS path: I
 Communities: target:3333:33 encapsulation0:0:0:0:3
 Import Accepted

```



```

Route Label: 300
ESI: 00:00:00:00:00:00:00:00:00
Localpref: 100
Router ID: 1.2.3.4
Secondary Tables: default-switch.evpn.0
Indirect next hops: 1
 Protocol next hop: 1.2.3.4 Metric: 1
 Indirect next hop: 0x2 no-forward INH Session ID: 0x0
 Indirect path forwarding next hops: 1
 Next hop type: Router
 Next hop: 10.10.10.1 via xe-0/0/1.0
 Session Id: 0x2
 1.2.3.4/32 Originating RIB: inet.0
 Metric: 1 Node path count: 1
 Forwarding nexthops: 2
 Nexthop: 10.92.78.102 via em0.0

3:1000:10::100::1.2.3.4/304 (1 entry, 0 announced)
*BGP Preference: 170/-101
Route Distinguisher: 1000:10
PMSI: Flags 0x0: Label 100: Type INGRESS-REPLICATION 1.2.3.4
Next hop type: Indirect
Address: 0x9420fd0
Next-hop reference count: 12
Source: 1.2.3.4
Protocol next hop: 1.2.3.4
Indirect next hop: 0x2 no-forward INH Session ID: 0x0
State: Local AS:17 Peer AS:17 Age:37:01 Metric2:1 Validation State:
unverified Task: BGP 17.1.2.3.4+50756
AS path: I
Communities: target:1111:8388708 encapsulation0:0:0:0:3
Import Accepted
Localpref: 100
Router ID: 1.2.3.4
Secondary Tables: default-switch.evpn.0
Indirect next hops: 1
 Protocol next hop: 1.2.3.4 Metric: 1
 Indirect next hop: 0x2 no-forward INH Session ID: 0x0
 Indirect path forwarding next hops: 1
 Next hop type: Router
 Next hop: 10.10.10.1 via xe-0/0/1.0
 Session Id: 0x2
 1.2.3.4/32 Originating RIB: inet.0
 Metric: 1 Node path count: 1
 Forwarding nexthops: 2
 Nexthop: 10.92.78.102 via em0.0

3:1000:10::200::1.2.3.4/304 (1 entry, 0 announced)
*BGP Preference: 170/-101
Route Distinguisher: 1000:10
PMSI: Flags 0x0: Label 200: Type INGRESS-REPLICATION 1.2.3.4
Next hop type: Indirect
Address: 0x9420fd0
Next-hop reference count: 12
Source: 1.2.3.4
Protocol next hop: 1.2.3.4
Indirect next hop: 0x2 no-forward INH Session ID: 0x0
State: Local AS: 17 Peer AS: 17 Age:35:22 Metric2:1 Validation
State:unverified Task: BGP 17.1.2.3.4+50756
AS path:I Communities: target:2222:22 encapsulation):0:0:0:0:3

```

```
Import Accepted
 Localpref: 100
 Router ID: 1.2.3.4
 Secondary Tables: default-switch.evpn.0
 Indirect next hops: 1
 Protocol next hop: 1.2.3.4 Metric: 1
 Indirect next hop: 0x2 no-forward INH Session ID: 0x0
 Indirect path forwarding next hops: 1
 Next hop type: Router
 Next hop: 10.10.10.1 via xe-0/0/1.0
 Session Id: 0x2
 1.2.3.4/32 Originating RIB: inet.0
 Metric: 1 Node path count: 1
 Forwarding nexthops: 2
 Nexthop: 10.92.78.102 via em0.0

3:1000:10::300::1.2.3.4/304 (1 entry, 0 announced)
 *BGP Preference: 170/-101
 Route Distinguisher: 1000:10
 PMSI: Flags 0x0: Label 300: Type INGRESS-REPLICATION 1.2.3.4
 Next hop type: Indirect
 Address: 0x9420fd0
 Next-hop reference count: 12
 Source: 1.2.3.4
 Protocol next hop: 1.2.3.4
 Indirect next hop: 0x2 no-forward INH Session ID: 0x0
 State: Local AS: 17 Peer AS: 17 Age 35:22 Metric2:1 Validation State:
unverified Task: BGP 17.1.2.3.4+5075
 6 AS path: I Communities: target:3333:33 encapsulation0:0:0:0:3
Import Accepted Localpref:100
 Router ID: 1.2.3.4
 Secondary Tables: default-switch.evpn.0
 Indirect next hops: 1
 Protocol next hop: 1.2.3.4 Metric: 1
 Indirect next hop: 0x2 no-forward INH Session ID: 0x0
 Indirect path forwarding next hops: 1
 Next hop type: Router
 Next hop: 10.10.10.1 via xe-0/0/1.0
 Session Id: 0x2
 1.2.3.4/32 Originating RIB: inet.0
 Metric: 1 Node path count: 1
 Forwarding nexthops: 2
 Nexthop: 10.92.78.102 via em0.0
```

## show route forwarding-table

|                                                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                                | <a href="#">Syntax on page 5227</a><br><a href="#">Syntax (MX Series Routers) on page 5227</a><br><a href="#">Syntax (TX Matrix and TX Matrix Plus Routers) on page 5227</a>                                                                                                                                                                                                                                                                                                                                                            |
| <b>Syntax</b>                                        | <pre>show route forwarding-table &lt;detail   extensive   summary&gt; &lt;all&gt; &lt;ccc interface-name&gt; &lt;destination destination-prefix&gt; &lt;family family   matching matching&gt; &lt;interface-name interface-name&gt; &lt;label name&gt; &lt;matching matching&gt; &lt;multicast&gt; &lt;table (default   logical-system-name/routing-instance-name   routing-instance-name)&gt; &lt;vlan (all   vlan-name)&gt; &lt;vpn vpn&gt;</pre>                                                                                     |
| <b>Syntax (MX Series Routers)</b>                    | <pre>show route forwarding-table &lt;detail   extensive   summary&gt; &lt;all&gt; &lt;bridge-domain (all   domain-name)&gt; &lt;ccc interface-name&gt; &lt;destination destination-prefix&gt; &lt;family family   matching matching&gt; &lt;interface-name interface-name&gt; &lt;label name&gt; &lt;learning-vlan-id learning-vlan-id&gt; &lt;matching matching&gt; &lt;multicast&gt; &lt;table (default   logical-system-name/routing-instance-name   routing-instance-name)&gt; &lt;vlan (all   vlan-name)&gt; &lt;vpn vpn&gt;</pre> |
| <b>Syntax (TX Matrix and TX Matrix Plus Routers)</b> | <pre>show route forwarding-table &lt;detail   extensive   summary&gt; &lt;all&gt; &lt;ccc interface-name&gt; &lt;destination destination-prefix&gt; &lt;family family   matching matching&gt; &lt;interface-name interface-name&gt; &lt;matching matching&gt; &lt;label name&gt; &lt;lcc number&gt; &lt;multicast&gt; &lt;table routing-instance-name&gt; &lt;vpn vpn&gt;</pre>                                                                                                                                                         |
| <b>Release Information</b>                           | <p>Command introduced before Junos OS Release 7.4.</p> <p>Option <b>bridge-domain</b> introduced in Junos OS Release 7.5</p> <p>Option <b>learning-vlan-id</b> introduced in Junos OS Release 8.4</p>                                                                                                                                                                                                                                                                                                                                   |

Options **all** and **vlan** introduced in Junos OS Release 9.6.

Command introduced in Junos OS Release 11.3 for the QFX Series.

Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Display the Routing Engine's forwarding table, including the network-layer prefixes and their next hops. This command is used to help verify that the routing protocol process has relayed the correction information to the forwarding table. The Routing Engine constructs and maintains one or more routing tables. From the routing tables, the Routing Engine derives a table of active routes, called the forwarding table.



**NOTE:** The Routing Engine copies the forwarding table to the Packet Forwarding Engine, the part of the router that is responsible for forwarding packets. To display the entries in the Packet Forwarding Engine's forwarding table, use the **show pfe route** command.

**Options** **none**—Display the routes in the forwarding tables. By default, the **show route forwarding-table** command does not display information about private, or internal, forwarding tables.

**detail | extensive | summary**—(Optional) Display the specified level of output.

**all**—(Optional) Display routing table entries for all forwarding tables, including private, or internal, tables.

**bridge-domain (all | bridge-domain-name)**—(MX Series routers only) (Optional) Display route entries for all bridge domains or the specified bridge domain.

**ccc interface-name**—(Optional) Display route entries for the specified circuit cross-connect interface.

**destination destination-prefix**—(Optional) Destination prefix.

**family family**—(Optional) Display routing table entries for the specified family: **fibre-channel**, **fmembers**, **inet**, **inet6**, **iso**, **mpls**, **tnp**, **unix**, **vpls**, or **vlan-classification**.

**interface-name interface-name**—(Optional) Display routing table entries for the specified interface.

**label name**—(Optional) Display route entries for the specified label.

**lcc number**—(TX Matrix and TX matrix Plus routers only) (Optional) On a routing matrix composed of a TX Matrix router and T640 routers, display information for the specified T640 router (or line-card chassis) connected to the TX Matrix router. On a routing matrix composed of the TX Matrix Plus router and T1600 or T4000 routers, display information for the specified router (line-card chassis) connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**learning-vlan-id** *learning-vlan-id*—(MX Series routers only) (Optional) Display learned information for all VLANs or for the specified VLAN.

**matching** *matching*—(Optional) Display routing table entries matching the specified prefix or prefix length.

**multicast**—(Optional) Display routing table entries for multicast routes.

**table** (*default* | *logical-system-name/routing-instance-name* | *routing-instance-name*)—(Optional) Display route entries for all the routing tables in the main routing instance or for the specified routing instance. If your device supports logical systems, you can also display route entries for the specified logical system and routing instance. To view the routing instances on your device, use the [show route instance](#) command.

**vlan** (*all* | *vlan-name*)—(Optional) Display information for all VLANs or for the specified VLAN.

**vpn** *vpn*—(Optional) Display routing table entries for a specified VPN.

**Required Privilege Level**

view

**List of Sample Output**

[show route forwarding-table on page 5232](#)  
[show route forwarding-table detail on page 5233](#)  
[show route forwarding-table destination extensive \(Weights and Balances\) on page 5233](#)  
[show route forwarding-table extensive on page 5234](#)  
[show route forwarding-table extensive \(RPF\) on page 5235](#)  
[show route forwarding-table family mpls on page 5236](#)  
[show route forwarding-table family vpls on page 5236](#)  
[show route forwarding-table vpls \(Broadcast, unknown unicast, and multicast \(BUM\) hashing is enabled\) on page 5236](#)  
[show route forwarding-table vpls \(Broadcast, unknown unicast, and multicast \(BUM\) hashing is enabled with MAC Statistics\) on page 5237](#)  
[show route forwarding-table family vpls extensive on page 5237](#)  
[show route forwarding-table table default on page 5238](#)  
[show route forwarding-table table logical-system-name/routing-instance-name on page 5239](#)

[show route forwarding-table vpn on page 5240](#)

**Output Fields** Table 306 lists the output fields for the **show route forwarding-table** command. Output fields are listed in the approximate order in which they appear. Field names might be abbreviated (as shown in parentheses) when no level of output is specified, or when the **detail** keyword is used instead of the **extensive** keyword.

Table 412: show route forwarding-table Output Fields

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Level of Output         |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| Logical system          | Name of the logical system. This field is displayed if you specify the <b>table logical-system-name/routing-instance-name</b> option on a device that is configured for and supports logical systems.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | All levels              |
| Routing table           | Name of the routing table (for example, inet, inet6, mpls).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | All levels              |
| Address family          | Address family (for example, IP, IPv6, ISO, MPLS, and VPLS).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | All levels              |
| Destination             | Destination of the route.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b> |
| Route Type (Type)       | How the route was placed into the forwarding table. When the <b>detail</b> keyword is used, the route type might be abbreviated (as shown in parentheses): <ul style="list-style-type: none"> <li><b>cloned (clon)</b>—(TCP or multicast only) Cloned route.</li> <li><b>destination (dest)</b>—Remote addresses directly reachable through an interface.</li> <li><b>destination down (iddn)</b>—Destination route for which the interface is unreachable.</li> <li><b>interface cloned (ifcl)</b>—Cloned route for which the interface is unreachable.</li> <li><b>route down (ifdn)</b>—Interface route for which the interface is unreachable.</li> <li><b>ignore (ignr)</b>—Ignore this route.</li> <li><b>interface (intf)</b>—Installed as a result of configuring an interface.</li> <li><b>permanent (perm)</b>—Routes installed by the kernel when the routing table is initialized.</li> <li><b>user</b>—Routes installed by the routing protocol process or as a result of the configuration.</li> </ul> | All levels              |
| Route Reference (RtRef) | Number of routes to reference.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail extensive</b> |
| Flags                   | Route type flags: <ul style="list-style-type: none"> <li><b>none</b>—No flags are enabled.</li> <li><b>accounting</b>—Route has accounting enabled.</li> <li><b>cached</b>—Cache route.</li> <li><b>incoming-iface interface-number</b>—Check against incoming interface.</li> <li><b>prefix load balance</b>—Load balancing is enabled for this prefix.</li> <li><b>rt nh decoupled</b>—Route has been decoupled from the next hop to the destination.</li> <li><b>sent to PFE</b>—Route has been sent to the Packet Forwarding Engine.</li> <li><b>static</b>—Static route.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>extensive</b>        |
| Next hop                | IP address of the next hop to the destination.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail extensive</b> |

Table 412: show route forwarding-table Output Fields (*continued*)

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Level of Output              |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| Next hop Type (Type)       | <p>Next-hop type. When the <b>detail</b> keyword is used, the next-hop type might be abbreviated (as indicated in parentheses):</p> <ul style="list-style-type: none"> <li>• <b>broadcast (bcst)</b>—Broadcast.</li> <li>• <b>deny</b>—Deny.</li> <li>• <b>discard (dscd)</b> —Discard.</li> <li>• <b>hold</b>—Next hop is waiting to be resolved into a unicast or multicast type.</li> <li>• <b>indexed (idxd)</b>—Indexed next hop.</li> <li>• <b>indirect (indr)</b>—Indirect next hop.</li> <li>• <b>local (locl)</b>—Local address on an interface.</li> <li>• <b>routed multicast (mcrst)</b>—Regular multicast next hop.</li> <li>• <b>multicast (mcst)</b>—Wire multicast next hop (limited to the LAN).</li> <li>• <b>multicast discard (mdsc)</b>—Multicast discard.</li> <li>• <b>multicast group (mgrp)</b>—Multicast group member.</li> <li>• <b>receive (rcv)</b>—Receive.</li> <li>• <b>reject (rjct)</b>—Discard. An ICMP unreachable message was sent.</li> <li>• <b>resolve (rslv)</b>—Resolving the next hop.</li> <li>• <b>unicast (ucst)</b>—Unicast.</li> <li>• <b>unilist (ulst)</b>—List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.</li> </ul> | <b>detail extensive</b>      |
| Index                      | Software index of the next hop that is used to route the traffic for a given prefix.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive none</b> |
| Route interface-index      | Logical interface index from which the route is learned. For example, for interface routes, this is the logical interface index of the route itself. For static routes, this field is zero. For routes learned through routing protocols, this is the logical interface index from which the route is learned.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>extensive</b>             |
| Reference (NhRef)          | Number of routes that refer to this next hop.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail extensive none</b> |
| Next-hop interface (Netif) | Interface used to reach the next hop.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive none</b> |
| Weight                     | Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible (see the <b>Balance</b> field description).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>extensive</b>             |
| Balance                    | Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a router is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>extensive</b>             |
| RPF interface              | List of interfaces from which the prefix can be accepted. Reverse path forwarding (RPF) information is displayed only when <b>rpf-check</b> is configured on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>extensive</b>             |

## Sample Output

### show route forwarding-table

```

user@host> show route forwarding-table
Routing table: default.inet
Internet:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0 Type Index NhRef Netif
0.0.0.0/32 perm 0 dscd 44 1
1.1.1.0/24 ifdn 0 rslv 608 1 ge-2/0/1.0
1.1.1.0/32 iddn 0 1.1.1.0 recv 606 1 ge-2/0/1.0
1.1.1.1/32 user 0 rjct 46 4
1.1.1.1/32 intf 0 1.1.1.1 locl 607 2
1.1.1.1/32 iddn 0 1.1.1.1 locl 607 2
1.1.1.255/32 iddn 0 ff:ff:ff:ff:ff:ff bcst 605 1 ge-2/0/1.0
10.0.0.0/24 intf 0 rslv 616 1 ge-2/0/0.0
10.0.0.0/32 dest 0 10.0.0.0 recv 614 1 ge-2/0/0.0
10.0.0.1/32 intf 0 10.0.0.1 locl 615 2
10.0.0.1/32 dest 0 10.0.0.1 locl 615 2
10.0.0.255/32 dest 0 10.0.0.255 bcst 613 1 ge-2/0/0.0
10.1.1.0/24 ifdn 0 rslv 612 1 ge-2/0/1.0
10.1.1.0/32 iddn 0 10.1.1.0 recv 610 1 ge-2/0/1.0
10.1.1.1/32 user 0 rjct 46 4
10.1.1.1/32 intf 0 10.1.1.1 locl 611 2
10.1.1.1/32 iddn 0 10.1.1.1 locl 611 2
10.1.1.255/32 iddn 0 ff:ff:ff:ff:ff:ff bcst 609 1 ge-2/0/1.0
10.206.0.0/16 user 0 10.209.63.254 ucst 419 20 fxp0.0
10.209.0.0/16 user 1 0:12:1e:ca:98:0 ucst 419 20 fxp0.0
10.209.0.0/18 intf 0 rslv 418 1 fxp0.0
10.209.0.0/32 dest 0 10.209.0.0 recv 416 1 fxp0.0
10.209.2.131/32 intf 0 10.209.2.131 locl 417 2
10.209.2.131/32 dest 0 10.209.2.131 locl 417 2
10.209.17.55/32 dest 0 0:30:48:5b:78:d2 ucst 435 1 fxp0.0
10.209.63.42/32 dest 0 0:23:7d:58:92:ca ucst 434 1 fxp0.0
10.209.63.254/32 dest 0 0:12:1e:ca:98:0 ucst 419 20 fxp0.0
10.209.63.255/32 dest 0 10.209.63.255 bcst 415 1 fxp0.0
10.227.0.0/16 user 0 10.209.63.254 ucst 419 20 fxp0.0

...

Routing table: iso
ISO:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0 rjct 27 1
47.0005.80ff.f800.0000.0108.0003.0102.5524.5220.00
intf 0 locl 28 1

Routing table: inet6
Internet6:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0 rjct 6 1
ff00::/8 perm 0 mdsc 4 1
ff02::1/128 perm 0 ff02::1 mcst 3 1

Routing table: ccc
MPLS:
Interface.Label Type RtRef Next hop Type Index NhRef Netif
default perm 0 rjct 16 1
100004(top)fe-0/0/1.0

```



## show route forwarding-table detail

```

user@host> show route forwarding-table detail
Routing table: inet
Internet:
Destination Type RtRef Next hop Type Index NhRef Netif
default user 2 0:90:69:8e:b1:1b ucst 132 4 fxp0.0
default perm 0 rjct 14 1
10.1.1.0/24 intf 0 ff.3.0.21 ucst 322 1 so-5/3/0.0
10.1.1.0/32 dest 0 10.1.1.0 recv 324 1 so-5/3/0.0
10.1.1.1/32 intf 0 10.1.1.1 locl 321 1
10.1.1.255/32 dest 0 10.1.1.255 bcst 323 1 so-5/3/0.0
10.21.21.0/24 intf 0 ff.3.0.21 ucst 326 1 so-5/3/0.0
10.21.21.0/32 dest 0 10.21.21.0 recv 328 1 so-5/3/0.0
10.21.21.1/32 intf 0 10.21.21.1 locl 325 1
10.21.21.255/32 dest 0 10.21.21.255 bcst 327 1 so-5/3/0.0
127.0.0.1/32 intf 0 127.0.0.1 locl 320 1
172.17.28.19/32 clon 1 192.168.4.254 ucst 132 4 fxp0.0
172.17.28.44/32 clon 1 192.168.4.254 ucst 132 4 fxp0.0

...

Routing table: private1__inet
Internet:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0 rjct 46 1
10.0.0.0/8 intf 0 rslv 136 1 fxp1.0
10.0.0.0/32 dest 0 10.0.0.0 recv 134 1 fxp1.0
10.0.0.4/32 intf 0 10.0.0.4 locl 135 2
10.0.0.4/32 dest 0 10.0.0.4 locl 135 2

...

Routing table: iso
ISO:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0 rjct 38 1

Routing table: inet6
Internet6:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0 rjct 22 1
ff00::/8 perm 0 mdsc 21 1
ff02::1/128 perm 0 ff02::1 mcst 17 1

...

Routing table: mpls
MPLS:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0 rjct 28 1

```

## show route forwarding-table destination extensive (Weights and Balances)

```

user@host> show route forwarding-table destination 3.4.2.1 extensive
Routing table: inet [Index 0]
Internet:

Destination: 3.4.2.1/32
Route type: user
Route reference: 0 Route interface-index: 0

```

|                                |               |              |
|--------------------------------|---------------|--------------|
| Flags: sent to PFE             |               |              |
| Next-hop type: unicast         | Index: 262143 | Reference: 1 |
| Nexthop: 4.4.4.4               |               |              |
| Next-hop type: unicast         | Index: 335    | Reference: 2 |
| Next-hop interface: so-1/1/0.0 | Weight: 22    | Balance: 3   |
| Nexthop: 145.12.1.2            |               |              |
| Next-hop type: unicast         | Index: 337    | Reference: 2 |
| Next-hop interface: so-0/1/2.0 | Weight: 33    | Balance: 33  |

### show route forwarding-table extensive

user@host> show route forwarding-table extensive

Routing table: inet [Index 0]

Internet:

Destination: default

Route type: user

Route reference: 2

Route interface-index: 0

Flags: sent to PFE

Nexthop: 0:90:69:8e:b1:1b

Next-hop type: unicast

Index: 132      Reference: 4

Next-hop interface: fxp0.0

Destination: default

Route type: permanent

Route reference: 0

Route interface-index: 0

Flags: none

Next-hop type: reject

Index: 14      Reference: 1

Destination: 127.0.0.1/32

Route type: interface

Route reference: 0

Route interface-index: 0

Flags: sent to PFE

Nexthop: 127.0.0.1

Next-hop type: local

Index: 320      Reference: 1

...

Routing table: private1\_\_inet [Index 1]

Internet:

Destination: default

Route type: permanent

Route reference: 0

Route interface-index: 0

Flags: sent to PFE

Next-hop type: reject

Index: 46      Reference: 1

Destination: 10.0.0.0/8

Route type: interface

Route reference: 0

Route interface-index: 3

Flags: sent to PFE

Next-hop type: resolve

Index: 136      Reference: 1

Next-hop interface: fxp1.0

...

Routing table: iso [Index 0]

ISO:

Destination: default

Route type: permanent

```

Route reference: 0
Flags: sent to PFE
Next-hop type: reject
Route interface-index: 0
Index: 38 Reference: 1

Routing table: inet6 [Index 0]
Internet6:

Destination: default
Route type: permanent
Route reference: 0
Flags: sent to PFE
Next-hop type: reject
Route interface-index: 0
Index: 22 Reference: 1

Destination: ff00::/8
Route type: permanent
Route reference: 0
Flags: sent to PFE
Next-hop type: multicast discard
Route interface-index: 0
Index: 21 Reference: 1

...

Routing table: private1__inet6 [Index 1]
Internet6:

Destination: default
Route type: permanent
Route reference: 0
Flags: sent to PFE
Next-hop type: reject
Route interface-index: 0
Index: 54 Reference: 1

Destination: fe80::2a0:a5ff:fe3d:375/128
Route type: interface
Route reference: 0
Flags: sent to PFE
Next-hop: fe80::2a0:a5ff:fe3d:375
Next-hop type: local
Route interface-index: 0
Index: 75 Reference: 1

...

```

### show route forwarding-table extensive (RPF)

The next example is based on the following configuration, which enables an RPF check on all routes that are learned from this interface, including the interface route:

```

so-1/1/0 {
 unit 0 {
 family inet {
 rpf-check;
 address 15.95.1.2/30;
 }
 }
}

```

```

user@host> show route forwarding-table extensive
Routing table: inet [Index 0]
Internet:
...
...
Destination: 15.95.1.3/32
Route type: destination
Route reference: 0
Route interface-index: 67

```

```

Flags: sent to PFE
Nexthop: 15.95.1.3
Next-hop type: broadcast Index: 328 Reference: 1
Next-hop interface: so-1/1/0.0
RPF interface: so-1/1/0.0

```

### show route forwarding-table family mpls

```

user@host> show route forwarding-table family mpls
Routing table: mpls
MPLS:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0
0 user 0
1 user 0
2 user 0
100000 user 0 10.31.1.6 swap 100001 fe-1/1/0.0
800002 user 0 Pop vt-0/3/0.32770

vt-0/3/0.32770 (VPLS)
 user 0 indr 351 4
 Push 800000, Push 100002(top)

so-0/0/0.0

```

### show route forwarding-table family vpls

```

user@host> show route forwarding-table family vpls
Routing table: green.vpls
VPLS:
Destination Type RtRef Next hop Type Index NhRef Netif
default dymn 0
default perm 0 rjct 298 1
fe-0/1/0.0 dymn 0 flood 355 1
00:90:69:0c:20:1f/48 <<<<<Remote CE
 dymn 0 indr 351 4
 Push 800000, Push 100002(top)

so-0/0/0.0
00:90:69:85:b0:1f/48 <<<<<Local CE
 dymn 0 ucst 354 2 fe-0/1/0.0

```

### show route forwarding-table vpls (Broadcast, unknown unicast, and multicast (BUM) hashing is enabled)

```

user@host> show route forwarding-table vpls
Routing table: green.vpls
VPLS:
Enabled protocols: BUM hashing
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0 dscd 519 1
lsi.1048832 intf 0 indr 1048574 4
 4.4.3.2 Push 262145 621 2

ge-3/0/0.0
00:19:e2:25:d0:01/48 user 0 ucst 590 5 ge-2/3/9.0
0x30003/51 user 0 comp 627 2
ge-2/3/9.0 intf 0 ucst 590 5 ge-2/3/9.0
ge-3/1/3.0 intf 0 ucst 619 4 ge-3/1/3.0
0x30002/51 user 0 comp 600 2
0x30001/51 user 0 comp 597 2

```

### show route forwarding-table vpls (Broadcast, unknown unicast, and multicast (BUM) hashing is enabled with MAC Statistics)

```

user@host> show route forwarding-table vpls
Routing table: green.vpls
VPLS:
Enabled protocols: BUM hashing, MAC Stats

```

| Destination          | Type | RtRef | Next hop | Type | Index   | NhRef | Netif      |
|----------------------|------|-------|----------|------|---------|-------|------------|
| default              | perm | 0     |          | dscd | 519     | 1     |            |
| 1si.1048834          | intf | 0     |          | indr | 1048574 | 4     |            |
|                      |      |       | 4.4.3.2  | Push | 262145  | 592   | 2          |
| ge-3/0/0.0           |      |       |          |      |         |       |            |
| 00:19:e2:25:d0:01/48 | user | 0     |          | ucst | 590     | 5     | ge-2/3/9.0 |
| 0x30003/51           | user | 0     |          | comp | 630     | 2     |            |
| ge-2/3/9.0           | intf | 0     |          | ucst | 590     | 5     | ge-2/3/9.0 |
| ge-3/1/3.0           | intf | 0     |          | ucst | 591     | 4     | ge-3/1/3.0 |
| 0x30002/51           | user | 0     |          | comp | 627     | 2     |            |
| 0x30001/51           | user | 0     |          | comp | 624     | 2     |            |

### show route forwarding-table family vpls extensive

```

user@host> show route forwarding-table family vpls extensive
Routing table: green.vpls [Index 2]
VPLS:

```

Destination: default

|                                |                              |
|--------------------------------|------------------------------|
| Route type: dynamic            | Route interface-index: 72    |
| Route reference: 0             |                              |
| Flags: sent to PFE             |                              |
| Next-hop type: flood           | Index: 289      Reference: 1 |
| Next-hop type: unicast         | Index: 291      Reference: 3 |
| Next-hop interface: fe-0/1/3.0 |                              |
| Next-hop type: unicast         | Index: 290      Reference: 3 |
| Next-hop interface: fe-0/1/2.0 |                              |

Destination: default

|                        |                              |
|------------------------|------------------------------|
| Route type: permanent  | Route interface-index: 0     |
| Route reference: 0     |                              |
| Flags: none            |                              |
| Next-hop type: discard | Index: 341      Reference: 1 |

Destination: fe-0/1/2.0

|                                |                              |
|--------------------------------|------------------------------|
| Route type: dynamic            | Route interface-index: 69    |
| Route reference: 0             |                              |
| Flags: sent to PFE             |                              |
| Next-hop type: flood           | Index: 293      Reference: 1 |
| Next-hop type: indirect        | Index: 363      Reference: 4 |
| Next-hop type: Push 800016     |                              |
| Next-hop interface: at-1/0/1.0 |                              |
| Next-hop type: indirect        | Index: 301      Reference: 5 |
| Next hop: 10.31.3.2            |                              |
| Next-hop type: Push 800000     |                              |
| Next-hop interface: fe-0/1/1.0 |                              |
| Next-hop type: unicast         | Index: 291      Reference: 3 |
| Next-hop interface: fe-0/1/3.0 |                              |

Destination: fe-0/1/3.0

|                      |                              |
|----------------------|------------------------------|
| Route type: dynamic  | Route interface-index: 70    |
| Route reference: 0   |                              |
| Flags: sent to PFE   |                              |
| Next-hop type: flood | Index: 292      Reference: 1 |

```

Next-hop type: indirect Index: 363 Reference: 4
Next-hop type: Push 800016
Next-hop interface: at-1/0/1.0
Next-hop type: indirect Index: 301 Reference: 5
Next hop: 10.31.3.2
Next-hop type: Push 800000
Next-hop interface: fe-0/1/1.0
Next-hop type: unicast Index: 290 Reference: 3
Next-hop interface: fe-0/1/2.0

Destination: 10:00:00:01:01:01/48
Route type: dynamic
Route reference: 0 Route interface-index: 70
Flags: sent to PFE, prefix load balance
Next-hop type: unicast Index: 291 Reference: 3
Next-hop interface: fe-0/1/3.0
Route used as destination:
 Packet count: 6640 Byte count: 675786
Route used as source
 Packet count: 6894 Byte count: 696424

Destination: 10:00:00:01:01:04/48
Route type: dynamic
Route reference: 0 Route interface-index: 69
Flags: sent to PFE, prefix load balance
Next-hop type: unicast Index: 290 Reference: 3
Next-hop interface: fe-0/1/2.0
Route used as destination:
 Packet count: 96 Byte count: 8079
Route used as source:
 Packet count: 296 Byte count: 24955

Destination: 10:00:00:01:03:05/48
Route type: dynamic
Route reference: 0 Route interface-index: 74
Flags: sent to PFE, prefix load balance
Next-hop type: indirect Index: 301 Reference: 5
Next hop: 10.31.3.2
Next-hop type: Push 800000
Next-hop interface: fe-0/1/1.0

```

### show route forwarding-table table default

```

user@host> show route forwarding-table table default
Routing table: default.inet
Internet:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0
0.0.0.0/32 perm 0
10.0.60.0/30 user 0 10.0.60.13 ucst 713 5 fe-0/1/3.0
10.0.60.12/30 intf 0 rslv 688 1 fe-0/1/3.0
10.0.60.12/32 dest 0 10.0.60.12 recv 686 1 fe-0/1/3.0
10.0.60.13/32 dest 0 0:5:85:8b:bc:22 ucst 713 5 fe-0/1/3.0
10.0.60.14/32 intf 0 10.0.60.14 locl 687 2
10.0.60.14/32 dest 0 10.0.60.14 locl 687 2
10.0.60.15/32 dest 0 10.0.60.15 bcst 685 1 fe-0/1/3.0
10.0.67.12/30 user 0 10.0.60.13 ucst 713 5 fe-0/1/3.0
10.0.80.0/30 ifdn 0 ff.3.0.21 ucst 676 1 so-0/0/1.0
10.0.80.0/32 dest 0 10.0.80.0 recv 678 1 so-0/0/1.0
10.0.80.2/32 user 0 rjct 36 2
10.0.80.2/32 intf 0 10.0.80.2 locl 675 1

```

```

10.0.80.3/32 dest 0 10.0.80.3 bcst 677 1 so-0/0/1.0
10.0.90.12/30 intf 0 rslv 684 1 fe-0/1/0.0
10.0.90.12/32 dest 0 10.0.90.12 recv 682 1 fe-0/1/0.0
10.0.90.14/32 intf 0 10.0.90.14 locl 683 2
10.0.90.14/32 dest 0 10.0.90.14 locl 683 2
10.0.90.15/32 dest 0 10.0.90.15 bcst 681 1 fe-0/1/0.0
10.5.0.0/16 user 0 192.168.187.126 ucst 324 15 fxp0.0
10.10.0.0/16 user 0 192.168.187.126 ucst 324 15 fxp0.0
10.13.10.0/23 user 0 192.168.187.126 ucst 324 15 fxp0.0
10.84.0.0/16 user 0 192.168.187.126 ucst 324 15 fxp0.0
10.150.0.0/16 user 0 192.168.187.126 ucst 324 15 fxp0.0
10.157.64.0/19 user 0 192.168.187.126 ucst 324 15 fxp0.0
10.209.0.0/16 user 0 192.168.187.126 ucst 324 15 fxp0.0

```

...

Routing table: default.iso

ISO:

| Destination | Type | RtRef | Next hop | Type | Index | NhRef | Netif |
|-------------|------|-------|----------|------|-------|-------|-------|
| default     | perm | 0     |          | rjct | 60    | 1     |       |

Routing table: default.inet6

Internet6:

| Destination | Type | RtRef | Next hop | Type | Index | NhRef | Netif |
|-------------|------|-------|----------|------|-------|-------|-------|
| default     | perm | 0     |          | rjct | 44    | 1     |       |
| ::/128      | perm | 0     |          | dscd | 42    | 1     |       |
| ff00::/8    | perm | 0     |          | mdsc | 43    | 1     |       |
| ff02::1/128 | perm | 0     | ff02::1  | mcst | 39    | 1     |       |

Routing table: default.mpls

MPLS:

| Destination | Type | RtRef | Next hop | Type | Index | NhRef | Netif |
|-------------|------|-------|----------|------|-------|-------|-------|
| default     | perm | 0     |          | dscd | 50    | 1     |       |

### show route forwarding-table table logical-system-name/routing-instance-name

```
user@host> show route forwarding-table table R4/vpn-red
```

Logical system: R4

Routing table: vpn-red.inet

Internet:

| Destination        | Type | RtRef | Next hop                                       | Type | Index | NhRef | Netif      |
|--------------------|------|-------|------------------------------------------------|------|-------|-------|------------|
| default            | perm | 0     |                                                | rjct | 563   | 1     |            |
| 0.0.0.0/32         | perm | 0     |                                                | dscd | 561   | 2     |            |
| 1.0.0.1/32         | user | 0     |                                                | dscd | 561   | 2     |            |
| 2.0.2.0/24         | intf | 0     |                                                | rslv | 771   | 1     | ge-1/2/0.3 |
| 2.0.2.0/32         | dest | 0     | 2.0.2.0                                        | recv | 769   | 1     | ge-1/2/0.3 |
| 2.0.2.1/32         | intf | 0     | 2.0.2.1                                        | locl | 770   | 2     |            |
| 2.0.2.1/32         | dest | 0     | 2.0.2.1                                        | locl | 770   | 2     |            |
| 2.0.2.2/32         | dest | 0     | 0.4.80.3.0.1b.c0.d5.e4.bd.0.1b.c0.d5.e4.bc.8.0 | ucst | 789   | 1     | ge-1/2/0.3 |
| 2.0.2.255/32       | dest | 0     | 2.0.2.255                                      | bcst | 768   | 1     | ge-1/2/0.3 |
| 224.0.0.0/4        | perm | 1     |                                                | mdsc | 562   | 1     |            |
| 224.0.0.1/32       | perm | 0     | 224.0.0.1                                      | mcst | 558   | 1     |            |
| 255.255.255.255/32 | perm | 0     |                                                | bcst | 559   | 1     |            |

Logical system: R4

Routing table: vpn-red.iso

ISO:

| Destination | Type | RtRef | Next hop | Type | Index | NhRef | Netif |
|-------------|------|-------|----------|------|-------|-------|-------|
| default     | perm | 0     |          | rjct | 608   | 1     |       |

```

Logical system: R4
Routing table: vpn-red.inet6
Internet6:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0 rjct 708 1
::/128 perm 0 dscd 706 1
ff00::/8 perm 0 mdsc 707 1
ff02::1/128 perm 0 ff02::1 mcst 704 1

```

```

Logical system: R4
Routing table: vpn-red.mpls
MPLS:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0 dscd 638

```

### show route forwarding-table vpn

```

user@host> show route forwarding-table vpn VPN-A
Routing table:: VPN-A.inet
Internet:
Destination Type RtRef Nexthop Type Index NhRef Netif
default perm 0 rjct 4 4
10.39.10.20/30 intf 0 ff.3.0.21 ucst 40 1
so-0/0/0.0
10.39.10.21/32 intf 0 10.39.10.21 locl 36 1
10.255.14.172/32 user 0 ucst 69 2
so-0/0/0.0
10.255.14.175/32 user 0 indr 81 3
Push 100004, Push
100004(top) so-1/0/0.0
224.0.0.0/4 perm 2 mdsc 5 3
224.0.0.1/32 perm 0 224.0.0.1 mcst 1 8
224.0.0.5/32 user 1 224.0.0.5 mcst 1 8
255.255.255.255/32 perm 0 bcst 2 3

```



## show mpls static-lsp

**Syntax** show mpls static-lsp  
 <brief | detail | extensive | terse>  
 <bypass>  
 <descriptions>  
 <down | up>  
 <ingress>  
 <instance *instance-name*>  
 <logical-system (all | *logical-system-name*)>  
 <lsp-type>  
 <name *name*>  
 <statistics>  
 <transit>

**Release Information** Command introduced in Junos OS Release 10.1.  
 Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

**Description** Display information about configured and active static Multiprotocol Label Switching (MPLS) label-switched paths (LSPs).

**Options** **none**—Display standard information about all configured and active static MPLS LSPs.

**brief | detail | extensive | terse**—(Optional) Display the specified level of output. The **extensive** option displays the same information as the **detail** option, but covers the most recent 50 events.

**bypass**—(Optional) Display LSPs used for protecting other static LSPs.

**descriptions**—(Optional) Display the MPLS static LSP descriptions. To view this information, you must configure the description statement at the **[edit protocols mpls static-label-switched-path *path-name* bypass]**, **[edit protocols mpls static-label-switched-path *path-name* ingress]**, or **[edit protocols mpls static-label-switched-path *path-name* transit *incoming-label*]** hierarchy levels. Only static LSPs with a description are displayed.

**down | up**—(Optional) Display only static LSPs that are inactive or active, respectively.

**instance *instance-name***—(Optional) Display information about all configured and active static MPLS LSPs for the specified routing instance. If ***instance-name*** is omitted, information about all configured and active static MPLS LSPs for the master instance is displayed.

**logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

***lsp-type***—(Optional) Display information about a particular LSP type:

- **bypass**—Sessions for bypass LSPs.
- **ingress**—Sessions that originate from this routing device.
- **transit**—Sessions that pass through this routing device.

**name *name***—(Optional) Display information about the specified static LSP or group of LSPs.

**statistics**—(Optional) Display accounting information about static LSPs.

**transit**—(Optional) Display static LSPs transiting this routing device.

**Required Privilege Level** view

**List of Sample Output** [show mpls static-lsp extensive on page 5243](#)  
[show mpls static-lsp statistics ingress on page 5243](#)  
[show mpls static-lsp \(when MPLS stitching is used\) on page 5243](#)

**Output Fields** Table 413 describes the output fields for the **show mpls static-lsp** command. Output fields are listed in the approximate order in which they appear.

**Table 413: show mpls static-lsp Output Fields**

| Field Name                     | Field Description                                                                                                                                                                             | Level of Output          |
|--------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------|
| <b>Ingress LSPs</b>            | Information about the static LSPs on the ingress routing device. Each session has one line of output.                                                                                         | All levels               |
| <b>Transit LSPs</b>            | Number of static LSPs on the transit routing devices and the state of these paths. MPLS learns this information by querying RSVP, which holds all the transit and egress session information. | All levels               |
| <b>Bypass LSPs</b>             | Information about the bypass LSPs configured on the routing device. Each session has one line of output.                                                                                      | All levels               |
| <b>LSPname</b>                 | Name of the static LSP.                                                                                                                                                                       | All levels               |
| <b>To</b>                      | Destination (egress routing device) of the session.                                                                                                                                           | All levels               |
| <b>State</b>                   | State of the static LSP handled by this RSVP session: <b>Up</b> , <b>Dn</b> (down), or <b>Restart</b> .                                                                                       | All levels               |
| <b>Packets</b>                 | Number of packet transiting the static LSP ( <b>statistics</b> option only).                                                                                                                  | All levels               |
| <b>Bytes</b>                   | Number of bytes transiting the static LSP ( <b>statistics</b> option only).                                                                                                                   | All levels               |
| <b>Nexthop</b>                 | IP address for the next-hop router for the static LSP.                                                                                                                                        | <b>detail, extensive</b> |
| <b>Bypass</b>                  | (Bypass LSP) Destination address (egress routing device) for the bypass LSP.                                                                                                                  | All levels               |
| <b>Link protection desired</b> | Link protection has been requested by the ingress routing device.                                                                                                                             | <b>detail, extensive</b> |
| <b>LabelOperation</b>          | Label operation to perform: <b>Push</b> , <b>Pop</b> , <b>Swap</b> .                                                                                                                          | <b>detail, extensive</b> |
| <b>Outgoing-label</b>          | Outgoing label to use for the MPLS packet in either push or swap label operations.                                                                                                            | <b>detail, extensive</b> |

Table 413: show mpls static-lsp Output Fields (*continued*)

| Field Name        | Field Description                                                                                                                                                                                                                   | Level of Output          |
|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------|
| <b>Created</b>    | (Ingress LSP) Date and time the static LSP was created.                                                                                                                                                                             | <b>extensive</b>         |
| <b>Bandwidth</b>  | Bandwidth configured for the static LSP.                                                                                                                                                                                            | <b>detail, extensive</b> |
| <b>Resv style</b> | (Bypass) RSVP reservation style. This field consists of two parts: the number of active reservations and the reservation style, which can be <b>FF</b> (fixed filter), <b>SE</b> (shared explicit), or <b>WF</b> (wildcard filter). | All levels               |

## Sample Output

### show mpls static-lsp extensive

```

user@host> show mpls static-lsp extensive
Ingress LSPs:
LSPName: alpha-to-beta, To: 192.168.14.1
State: Dn
Nexthop: 192.168.10.1
LabelOperation: Push, Outgoing-label: 1000001
Created: Thu Jan 14 16:44:43 2010
Bandwidth: 0 bps
Total 1, displayed 1, Up 0, Down 1

Transit LSPs:
Total 0, displayed 0, Up 0, Down 0

Bypass LSPs:
Total 0, displayed 0, Up 0, Down 0

```

### show mpls static-lsp statistics ingress

```

user@host> show mpls static-lsp statistics ingress
Ingress LSPs:
LSPName To State Packets Bytes
alpha-to-beta 192.168.14.1 Dn NA NA
Total 1, displayed 1, Up 0, Down 1

```

### show mpls static-lsp (when MPLS stitching is used)

The show mpls static-lsp command was extended in Junos release 14.1X53-D25 to accommodate the stitching feature of MPLS. This example shows the LSP state as 'InProgress' because the LSP is waiting for protocol next-hop resolution. For more information, see

```

user@host> show mpls static-lsp
Ingress LSPs:
Total 0, displayed 0, Up 0, Down 0
Transit LSPs: LSPName Incoming-label State
to-165 1000000 InProgress

```

## show ted database

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 5244</a><br><a href="#">Syntax (EX Series Switches) on page 5244</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Syntax</b>                      | <pre>show ted database &lt;brief   detail   extensive&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;<i>system-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Syntax (EX Series Switches)</b> | <pre>show ted database &lt;brief   detail   extensive&gt; &lt;<i>system-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>         | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p><b>instance <i>instance-name</i></b> option added in Junos OS Release 15.1.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>                 | Display the entries in the Multiprotocol Label Switching (MPLS) traffic engineering database.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                     | <p><b>none</b>—Display standard information about all entries in the traffic engineering database.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display routing instance information for the specified instance. If <i>instance-name</i> is omitted, information is displayed for the master instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b><i>system-name</i></b>—(Optional) Display traffic engineering database information for a particular system.</p> |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>List of Sample Output</b>       | <a href="#">show ted database brief on page 5247</a><br><a href="#">show ted database detail on page 5247</a><br><a href="#">show ted database extensive on page 5248</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Output Fields</b>               | <p><a href="#">Table 414</a> describes the output fields for the <b>show ted database</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

**Table 414: show ted database Output Fields**

| Field Name   | Field Description                                                               | Level of Output |
|--------------|---------------------------------------------------------------------------------|-----------------|
| TED database | Number of nodes and pseudonodes participating in IS-IS and OSPF domain routing. | All levels      |

Table 414: show ted database Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                           | Level of Output         |
|--------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| ID                             | Hostname and address of the node that the link is coming from. An address of .00 indicates that the node is the routing device itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode. If the node contains a router ID, it is displayed in parentheses. | <b>brief</b>            |
| NodeID                         | Hostname and address of the node that the link is coming from. An address of .00 indicates that the node is the routing device itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode.                                                                   | <b>extensive</b>        |
| Type                           | Type of node. It can be either <b>Rtr</b> (router) or <b>Net</b> (pseudonode).                                                                                                                                                                                                              | All levels              |
| Age(s)                         | How long since the node was last refreshed, in seconds.                                                                                                                                                                                                                                     | All levels              |
| LnkIn                          | Number of nodes pointing toward this node.                                                                                                                                                                                                                                                  | All levels              |
| LnkOut                         | Number of nodes to which this node points.                                                                                                                                                                                                                                                  | All levels              |
| Protocol                       | Protocol that reported the node information: <ul style="list-style-type: none"> <li>• <b>IS-IS(1)</b>—IS-IS Level 1.</li> <li>• <b>IS-IS(2)</b>—IS-IS Level 2.</li> <li>• <b>OSPF (area-number)</b>—OSPF from the specified area.</li> </ul>                                                | All levels              |
| To                             | Address on the far end of a link.                                                                                                                                                                                                                                                           | <b>detail extensive</b> |
| Local                          | Address of the local interface being used to reach the remote node.                                                                                                                                                                                                                         | <b>detail extensive</b> |
| Remote                         | Address of the interface on the remote node.                                                                                                                                                                                                                                                | <b>detail extensive</b> |
| Local interface index          | The interface indexes enable Junos OS to support unnumbered extensions for IS-IS, as described in RFC 4205.                                                                                                                                                                                 | <b>detail extensive</b> |
| Remote interface index         | The interface indexes enable Junos OS to support unnumbered extensions for IS-IS, as described in RFC 4205.                                                                                                                                                                                 | <b>detail extensive</b> |
| Metric                         | Configured traffic engineering metric.                                                                                                                                                                                                                                                      | <b>extensive</b>        |
| IGP metric                     | Configured interior gateway protocol metric.                                                                                                                                                                                                                                                | <b>extensive</b>        |
| Static BW                      | Total interface bandwidth in bps.                                                                                                                                                                                                                                                           | <b>extensive</b>        |
| Reservable bandwidth           | Subscription factor for the interface, which is the percentage of the link bandwidth that can be used for the RSVP reservation process. You configure this by including the <b>subscription</b> statement when configuring RSVP.                                                            | <b>extensive</b>        |
| <b>Available BW [priority]</b> | (Must include <b>diffserv-te</b> statement when configuring LSPs) Amount of bandwidth actually reserved by RSVP for each priority level. The bandwidth shown is for the entire interface, not for each individual LSP.                                                                      | <b>extensive</b>        |

Table 414: show ted database Output Fields (*continued*)

| Field Name                                             | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Level of Output |
|--------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Diffserv-TE BW Model                                   | Bandwidth constraint model used by the LSPs.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | extensive       |
| Available BW [TE-class]                                | (Must include the <b>diffserv-te</b> statement when configuring LSPs) Amount of bandwidth actually reserved by RSVP for each traffic engineering class.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | extensive       |
| Static BW [CT-class]                                   | Total interface bandwidth used by an MPLS traffic class, in bps.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | extensive       |
| Interface Switching Capability Descriptor ( <i>n</i> ) | <p>Information about the interface switching capability descriptor, which is a subtype length value (TLV) of the link TLV. <i>n</i> is the index number.</p> <ul style="list-style-type: none"> <li>• <b>Switching type</b>—Type of switching to be performed on a particular link: <ul style="list-style-type: none"> <li>• PSC-1—Packet switch-capable 1</li> <li>• PSC-2—Packet switch-capable 2</li> <li>• PSC-3—Packet switch-capable 3</li> <li>• PSC-4—Packet switch-capable 4</li> <li>• L2SC—Layer-2-switch-capable</li> <li>• TDM—Time-division-multiplexing-capable</li> <li>• LSC—Lambda switch-capable</li> <li>• FSC—Fiber switch-capable</li> </ul> </li> <li>• <b>Encoding type</b>—Encoding of the LSP being requested: <ul style="list-style-type: none"> <li>• Packet</li> <li>• Ethernet</li> <li>• ANSI/ETSI PDH</li> <li>• Reserved</li> <li>• SDH /SONET</li> <li>• Digital Wrapper</li> <li>• Lambda (photonic)</li> <li>• Fiber</li> <li>• FiberSDH/SONET</li> </ul> </li> <li>• <b>Maximum LSP BW [priority] bps</b>—Maximum LSP bandwidth information. Amount of bandwidth actually reserved for each priority level. The bandwidth shown is for the entire interface. <ul style="list-style-type: none"> <li>• [<i>n</i>]—Priority level. The range is from 0 (high) through 7 (low).</li> <li>• <i>n</i> Mbps—Amount of the maximum bandwidth.</li> </ul> </li> <li>• <b>Minimum LSP BW</b>—Minimum LSP bandwidth in Mbps. Amount of bandwidth actually reserved for each priority level. The bandwidth shown is for the entire interface. <b>Minimum LSP BW</b> is displayed only when <b>switching type</b> is PSC-1 or TDM.</li> <li>• <b>Interface MTU</b>—Displayed only when <b>switching type</b> is TDM.</li> <li>• <b>Interface supports standard SONET/SDH</b>—Displayed only when <b>switching type</b> is TDM.</li> </ul> | extensive       |

## Sample Output

### show ted database brief

```

user@host> show ted database brief
TED database: 12 ISIS nodes 0 INET nodes
ID Type Age(s) LnkIn LnkOut Protocol
Router-A.00 --- 3178 2 0
Router-B.00 --- 3152 2 0
Router-B.02 Net 802 0 2 IS-IS(2)
 To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
 Local interface index: 0, Remote interface index: 0
 To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
 Local interface index: 0, Remote interface index: 0
ID Type Age(s) LnkIn LnkOut Protocol
Router-C.00 --- 3126 2 0
Router-C.02 Net 38 0 2 IS-IS(2)
 To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
 Local interface index: 0, Remote interface index: 0
 To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
 Local interface index: 0, Remote interface index: 0
ID Type Age(s) LnkIn LnkOut Protocol
Router-D.00 --- 3144 2 0
Router-D.02 Net 723 0 2 IS-IS(2)
 To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
 Local interface index: 0, Remote interface index: 0
 To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
 Local interface index: 0, Remote interface index: 0
ID Type Age(s) LnkIn LnkOut Protocol
Router-D.03 Net 607 0 2 IS-IS(2)
 To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
 Local interface index: 0, Remote interface index: 0
 To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
 Local interface index: 0, Remote interface index: 0
ID Type Age(s) LnkIn LnkOut Protocol
Router-E.00 --- 3178 2 0
Router-E.02 Net 131 0 2 IS-IS(2)
 To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
 Local interface index: 0, Remote interface index: 0
 To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
 Local interface index: 0, Remote interface index: 0
ID Type Age(s) LnkIn LnkOut Protocol
Router-F.00 --- 3153 2 0
Router-F.02 Net 769 0 2 IS-IS(2)
 To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
 Local interface index: 0, Remote interface index: 0
 To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
 Local interface index: 0, Remote interface index: 0

```

### show ted database detail

```

TED database: 12 ISIS nodes 0 INET nodes
ID Type Age(s) LnkIn LnkOut Protocol
Router-A.00 --- 2913 2 0
Router-B.00 --- 2887 2 0
Router-B.02 Net 537 0 2 IS-IS(2)
 To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
 Local interface index: 0, Remote interface index: 0
 To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
 Local interface index: 0, Remote interface index: 0
ID Type Age(s) LnkIn LnkOut Protocol

```

```

Router-C.00 --- 2861 2 0
Router-C.02 Net 597 0 2 IS-IS(2)
 To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
 Local interface index: 0, Remote interface index: 0
 To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
 Local interface index: 0, Remote interface index: 0
ID Type Age(s) LnkIn LnkOut Protocol
Router-D.00 --- 2879 2 0
Router-D.02 Net 458 0 2 IS-IS(2)
 To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
 Local interface index: 0, Remote interface index: 0
 To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
 Local interface index: 0, Remote interface index: 0
ID Type Age(s) LnkIn LnkOut Protocol
Router-D.03 Net 342 0 2 IS-IS(2)
 To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
 Local interface index: 0, Remote interface index: 0
 To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
 Local interface index: 0, Remote interface index: 0
ID Type Age(s) LnkIn LnkOut Protocol
Router-E.00 --- 2913 2 0
Router-E.02 Net 640 0 2 IS-IS(2)
 To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
 Local interface index: 0, Remote interface index: 0
 To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
 Local interface index: 0, Remote interface index: 0
ID Type Age(s) LnkIn LnkOut Protocol
Router-F.00 --- 2888 2 0
Router-F.02 Net 504 0 2 IS-IS(2)
 To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
 Local interface index: 0, Remote interface index: 0
 To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
 Local interface index: 0, Remote interface index: 0

```

### show ted database extensive

```

user@host> show ted database extensive
TED database: 12 ISIS nodes 0 INET nodes
NodeID: Router-A.00
 Type: ---, Age: 3067 secs, LinkIn: 2, LinkOut: 0
NodeID: Router-B.00
 Type: ---, Age: 3041 secs, LinkIn: 2, LinkOut: 0
NodeID: Router-B.02
 Type: Net, Age: 691 secs, LinkIn: 0, LinkOut: 2
 Protocol: IS-IS(2)
 To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
 Local interface index: 0, Remote interface index: 0
 Metric: 0
 IGP metric: 10
 Interface Switching Capability Descriptor(1):
 Switching type: Packet
 Encoding type: Packet
 Maximum LSP BW [priority] bps:
 [0] 0bps [1] 0bps [2] 0bps [3] 0bps
 [4] 0bps [5] 0bps [6] 0bps [7] 0bps
 To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
 Local interface index: 0, Remote interface index: 0
 Metric: 0
 IGP metric: 20
 Interface Switching Capability Descriptor(1):
 Switching type: Packet

```



```

 Encoding type: Packet
 Maximum LSP BW [priority] bps:
 [0] 0bps [1] 0bps [2] 0bps [3] 0bps
 [4] 0bps [5] 0bps [6] 0bps [7] 0bps
NodeID: Router-C.00
Type: ---, Age: 3015 secs, LinkIn: 2, LinkOut: 0
NodeID: Router-C.02
Type: Net, Age: 751 secs, LinkIn: 0, LinkOut: 2
Protocol: IS-IS(2)
To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
Metric: 0
IGP metric: 10
Interface Switching Capability Descriptor(1):
 Switching type: Packet
 Encoding type: Packet
 Maximum LSP BW [priority] bps:
 [0] 0bps [1] 0bps [2] 0bps [3] 0bps
 [4] 0bps [5] 0bps [6] 0bps [7] 0bps
To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
Metric: 0
IGP metric: 10 Interface Switching Capability Descriptor(1):
 Switching type: Packet
 Encoding type: Packet
 Maximum LSP BW [priority] bps:
 [0] 0bps [1] 0bps [2] 0bps [3] 0bps
 [4] 0bps [5] 0bps [6] 0bps [7] 0bps
NodeID: Router-D.00
Type: ---, Age: 3034 secs, LinkIn: 2, LinkOut: 0
NodeID: Router-D.02
Type: Net, Age: 613 secs, LinkIn: 0, LinkOut: 2
Protocol: IS-IS(2)
To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
Metric: 0
IGP metric: 10
Interface Switching Capability Descriptor(1):
 Switching type: Packet
 Encoding type: Packet
 Maximum LSP BW [priority] bps:
 [0] 0bps [1] 0bps [2] 0bps [3] 0bps
 [4] 0bps [5] 0bps [6] 0bps [7] 0bps
To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
Metric: 0
IGP metric: 10
Interface Switching Capability Descriptor(1):
 Switching type: Packet
 Encoding type: Packet
 Maximum LSP BW [priority] bps:
 [0] 0bps [1] 0bps [2] 0bps [3] 0bps
 [4] 0bps [5] 0bps [6] 0bps [7] 0bps
NodeID: Router-D.03
Type: Net, Age: 497 secs, LinkIn: 0, LinkOut: 2
Protocol: IS-IS(2)
To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
Metric: 0
IGP metric: 10
Interface Switching Capability Descriptor(1):

```

```

Switching type: Packet
Encoding type: Packet
Maximum LSP BW [priority] bps:
 [0] 0bps [1] 0bps [2] 0bps [3] 0bps
 [4] 0bps [5] 0bps [6] 0bps [7] 0bps
To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
Metric: 0
IGP metric: 10
Interface Switching Capability Descriptor(1):
 Switching type: Packet
 Encoding type: Packet
 Maximum LSP BW [priority] bps:
 [0] 0bps [1] 0bps [2] 0bps [3] 0bps
 [4] 0bps [5] 0bps [6] 0bps [7] 0bps
NodeID: Router-E.00
Type: ---, Age: 3068 secs, LinkIn: 2, LinkOut: 0
NodeID: Router-E.02
Type: Net, Age: 21 secs, LinkIn: 0, LinkOut: 2
Protocol: IS-IS(2)
To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
Metric: 0
Interface Switching Capability Descriptor(1):
 Switching type: Packet
 Encoding type: Packet
 Maximum LSP BW [priority] bps:
 [0] 0bps [1] 0bps [2] 0bps [3] 0bps
 [4] 0bps [5] 0bps [6] 0bps [7] 0bps
To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
Metric: 0
IGP metric: 10
Interface Switching Capability Descriptor(1):
 Switching type: Packet
 Encoding type: Packet
 Maximum LSP BW [priority] bps:
 [0] 0bps [1] 0bps [2] 0bps [3] 0bps
 [4] 0bps [5] 0bps [6] 0bps [7] 0bps
NodeID: Router-F.00
Type: ---, Age: 3043 secs, LinkIn: 2, LinkOut: 0
NodeID: Router-F.02
Type: Net, Age: 659 secs, LinkIn: 0, LinkOut: 2
Protocol: IS-IS(2)
To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
Metric: 0
IGP metric: 10
Interface Switching Capability Descriptor(1):
 Switching type: Packet
 Encoding type: Packet
 Maximum LSP BW [priority] bps:
 [0] 0bps [1] 0bps [2] 0bps [3] 0bps
 [4] 0bps [5] 0bps [6] 0bps [7] 0bps
To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
Metric: 0
IGP metric: 10
Interface Switching Capability Descriptor(1):
 Switching type: Packet
 Encoding type: Packet
 Maximum LSP BW [priority] bps:

```

|          |          |          |          |
|----------|----------|----------|----------|
| [0] 0bps | [1] 0bps | [2] 0bps | [3] 0bps |
| [4] 0bps | [5] 0bps | [6] 0bps | [7] 0bps |

## show ted link

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 5252</a><br><a href="#">Syntax (EX Series Switches) on page 5252</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Syntax</b>                      | <pre>show ted link &lt;brief   detail&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Syntax (EX Series Switches)</b> | <pre>show ted link &lt;brief   detail&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>         | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p><b>instance <i>instance-name</i></b> option added in Junos OS Release 15.1.</p>                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>                 | Display Multiprotocol Label Switching (MPLS) traffic engineering database link information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                     | <p><b>none</b>—Display standard information about traffic engineering database link information.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display routing instance information for the specified instance. If <b><i>instance-name</i></b> is omitted, information is displayed for the master instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>List of Sample Output</b>       | <a href="#">show ted link brief on page 5253</a><br><a href="#">show ted link detail on page 5253</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Output Fields</b>               | Table 415 describes the output fields for the <b>show ted link</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                  |

Table 415: show ted link Output Fields

| Field Name | Field Description                                                                                                                                                                                                         | Level of Output |
|------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| ID         | Hostname and address of the node that the link is coming from. An address of .00 indicates that the node is the routing device itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode. | <b>brief</b>    |
| -->ID      | Hostname and address of the node that the link is going to. An address of .00 indicates that the node is the routing device itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode.    | <b>brief</b>    |

Table 415: show ted link Output Fields (*continued*)

| Field Name             | Field Description                                                                                                                                                                                                         | Level of Output         |
|------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <i>hostname</i>        | Hostname and address of the node that the link is coming from. An address of .00 indicates that the node is the routing device itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode. | <b>detail</b>           |
| <i>hostname</i>        | Hostname and address of the node that the link is going to. An address of .00 indicates that the node is the routing device itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode.    | <b>detail</b>           |
| Local Path             | Number of paths CSPF on the local routing device has placed on the link.                                                                                                                                                  | All levels              |
| Metric                 | Configured traffic engineering metric.                                                                                                                                                                                    | <b>extensive</b>        |
| IGP metric             | Configured interior gateway protocol metric.                                                                                                                                                                              | <b>detail</b>           |
| Local BW               | Amount of bandwidth the local routing device has placed on the link.                                                                                                                                                      | All levels              |
| Local                  | Address of the local interface being used to reach the remote node.                                                                                                                                                       | <b>detail extensive</b> |
| Remote                 | Address of the interface on the remote node.                                                                                                                                                                              | <b>detail extensive</b> |
| Local interface index  | The interface indexes enable Junos OS to support unnumbered extensions for IS-IS, as described in RFC 4205.                                                                                                               | <b>detail</b>           |
| Remote interface index | The interface indexes enable Junos OS to support unnumbered extensions for IS-IS, as described in RFC 4205.                                                                                                               | <b>detail</b>           |

## Sample Output

### show ted link brief

```

user@host> show ted link brief
ID ->ID LocalPath LocalBW
Router-B.02 Router-A.00 0 0bps
Router-B.02 Router-B.00 0 0bps
Router-C.02 Router-B.00 0 0bps
Router-C.02 Router-C.00 0 0bps
Router-D.02 Router-F.00 0 0bps
Router-D.02 Router-D.00 0 0bps
Router-D.03 Router-D.00 0 0bps
Router-D.03 Router-C.00 0 0bps
Router-E.02 Router-A.00 0 0bps
Router-E.02 Router-E.00 0 0bps
Router-F.02 Router-E.00 0 0bps
Router-F.02 Router-F.00 0 0bps

```

### show ted link detail

```

user@host> show ted link detail
Router-B.02->Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
 Local interface index: 0, Remote interface index: 0
 LocalPath: 0, Metric: 0, IGP metric: 10 AvailBW: 0bps
 localBW [0] 0bps [1] 0bps [2] 0bps [3] 0bps

```

```
 localBW [4] Obps [5] Obps [6] Obps [7] Obps
Router-B.02->Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
 Local interface index: 0, Remote interface index: 0
 LocalPath: 0, Metric: 0, IGP metric: 20 AvailBW: Obps
 localBW [0] Obps [1] Obps [2] Obps [3] Obps
 localBW [4] Obps [5] Obps [6] Obps [7] Obps
Router-C.02->Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
 Local interface index: 0, Remote interface index: 0
 LocalPath: 0, Metric: 0, IGP metric: 40 AvailBW: Obps
 localBW [0] Obps [1] Obps [2] Obps [3] Obps
 localBW [4] Obps [5] Obps [6] Obps [7] Obps
Router-C.02->Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
 Local interface index: 0, Remote interface index: 0
 LocalPath: 0, Metric: 0, IGP metric: 10 AvailBW: Obps
 localBW [0] Obps [1] Obps [2] Obps [3] Obps
 localBW [4] Obps [5] Obps [6] Obps [7] Obps
Router-D.02->Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
 Local interface index: 0, Remote interface index: 0
 LocalPath: 0, Metric: 0, IGP metric: 10 AvailBW: Obps
 localBW [0] Obps [1] Obps [2] Obps [3] Obps
 localBW [4] Obps [5] Obps [6] Obps [7] Obps
Router-D.02->Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
 Local interface index: 0, Remote interface index: 0
 LocalPath: 0, Metric: 0, IGP metric: 60 AvailBW: Obps
 localBW [0] Obps [1] Obps [2] Obps [3] Obps
 localBW [4] Obps [5] Obps [6] Obps [7] Obps
Router-D.03->Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
 Local interface index: 0, Remote interface index: 0
 LocalPath: 0, Metric: 0, IGP metric: 10 AvailBW: Obps
 localBW [0] Obps [1] Obps [2] Obps [3] Obps
 localBW [4] Obps [5] Obps [6] Obps [7] Obps
Router-D.03->Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
 Local interface index: 0, Remote interface index: 0
 LocalPath: 0, Metric: 0, IGP metric: 10 AvailBW: Obps
 localBW [0] Obps [1] Obps [2] Obps [3] Obps
 localBW [4] Obps [5] Obps [6] Obps [7] Obps
Router-E.02->Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
 Local interface index: 0, Remote interface index: 0
 LocalPath: 0, Metric: 0, IGP metric: 60 AvailBW: Obps
 localBW [0] Obps [1] Obps [2] Obps [3] Obps
 localBW [4] Obps [5] Obps [6] Obps [7] Obps
Router-E.02->Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
 Local interface index: 0, Remote interface index: 0
 LocalPath: 0, Metric: 0, IGP metric: 20 AvailBW: Obps
 localBW [0] Obps [1] Obps [2] Obps [3] Obps
 localBW [4] Obps [5] Obps [6] Obps [7] Obps
Router-F.02->Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
 Local interface index: 0, Remote interface index: 0
 LocalPath: 0, Metric: 0, IGP metric: 10 AvailBW: Obps
 localBW [0] Obps [1] Obps [2] Obps [3] Obps
 localBW [4] Obps [5] Obps [6] Obps [7] Obps
Router-F.02->Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
 Local interface index: 0, Remote interface index: 0
 LocalPath: 0, Metric: 0, IGP metric: 40 AvailBW: Obps
 localBW [0] Obps [1] Obps [2] Obps [3] Obps
 localBW [4] Obps [5] Obps [6] Obps [7] Obps
```

## show ted protocol

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 5255</a><br><a href="#">Syntax (EX Series Switches) on page 5255</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Syntax</b>                      | <pre>show ted protocol &lt;brief   detail&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Syntax (EX Series Switches)</b> | <pre>show ted protocol &lt;brief   detail&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>         | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p><b>instance <i>instance-name</i></b> option added in Junos OS Release 15.1.</p>                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>                 | Display information about the protocols from which the Multiprotocol Label Switching (MPLS) traffic engineering database learned about its nodes.                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                     | <p><b>none</b>—Display standard information about the protocols from which the traffic engineering database learned about its nodes.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display routing instance information for the specified instance. If <b><i>instance-name</i></b> is omitted, information is displayed for the master instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>List of Sample Output</b>       | <a href="#">show ted protocol on page 5256</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Output Fields</b>               | Table 416 describes the output fields for the <b>show ted protocol</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

**Table 416: show ted protocol Output Fields**

| Field Name           | Field Description                                                                                                                                                                                                                             |
|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Protocol name</b> | Protocol that reported the node information: <ul style="list-style-type: none"> <li><b>IS-IS(1)</b>—IS-IS Level 1.</li> <li><b>IS-IS(2)</b>—IS-IS Level 2.</li> <li><b>OSPF (<i>area-number</i>)</b>—OSPF from the specified area.</li> </ul> |
| <b>Credibility</b>   | If the protocols provide conflicting information about a node, the protocol with the highest credibility value is the one that the traffic engineering database uses.                                                                         |
| <b>Self node</b>     | Address the protocol uses as the local address.                                                                                                                                                                                               |

## Sample Output

show ted protocol

```
user@host> show ted protocol
Protocol name Credibility Self node
IS-IS(2) 2 (highest) corriedale.00(123.456.1.11)
IS-IS(1) 1 corriedale.00(123.456.1.11)
```



## PART 73

# RSVP

- [Using RSVP on page 5259](#)
- [Configuration Statements and Command Statements for RSVP on page 5301](#)



# Using RSVP

- [Understanding MPLS Components for QFX Series and EX4600 Switches on page 5259](#)
- [RSVP Overview on page 5263](#)
- [MTU Signaling in RSVP on page 5264](#)
- [Tunneling LDP LSPs in RSVP LSPs on page 5264](#)
- [Tunneling LDP LSPs in RSVP LSPs Overview on page 5265](#)
- [Configuring MPLS on Provider Edge Switches on page 5265](#)
- [Configuring MPLS on Provider Switches on page 5269](#)
- [Verifying That MPLS Is Working Correctly on page 5270](#)
- [Dynamic Bandwidth Management Using Container LSP Overview on page 5272](#)

## Understanding MPLS Components for QFX Series and EX4600 Switches

---

MPLS devices include a number of components. While some components are required for all MPLS applications, others might not be, depending on the specific application.

This topic includes:

- [Provider Edge Switches on page 5259](#)
- [Provider Switch on page 5261](#)
- [Components Required for All Switches in the MPLS Network on page 5261](#)

### Provider Edge Switches

To implement MPLS on a network, you must configure two provider edge (PE) switches—an ingress PE switch and an egress PE switch. In addition, you must configure one or more provider switches as transit switches within the network to support the forwarding of MPLS packets.

The ingress PE switch (the entry point to the MPLS tunnel) receives a packet, analyzes it, and pushes an MPLS label onto it. This label places the packet in a forwarding equivalence class (FEC) and determines its handling and destination through the MPLS tunnel. The egress PE switch (the exit point from the MPLS tunnel) pops the MPLS label off the outgoing packet.

Within an MPLS tunnel, the network traffic is bidirectional. Therefore, each PE switch can be configured to be both an ingress switch and an egress switch, depending on the direction of the traffic.

The following MPLS components are configured on the PE switches but not on the provider switches:

- [MPLS Protocol and Label-Switched Paths on page 5260](#)
- [IP Over MPLS for Customer Edge Interfaces on page 5260](#)
- [BGP Layer 3 VPN Configuration on page 5260](#)
- [Routing Instances for Layer 3 VPN on page 5260](#)
- [Routing Instances for Layer 2 VPN and Layer 3 VPN on page 5260](#)
- [Ethernet Encapsulation for Layer 2 VPN on page 5261](#)

---

### MPLS Protocol and Label-Switched Paths

Each PE switch must be configured to support the MPLS protocol. You must also configure label-switched paths (LSPs) at the **[edit protocols mpls]** hierarchy level.

---

### IP Over MPLS for Customer Edge Interfaces

You can configure the customer edge interfaces of the PE switches for IP over MPLS using a Layer 3 interface and a static route from the ingress PE switch to the egress PE switch. See “[Configuring MPLS on Provider Edge Switches](#)” on page 4981.

---

### BGP Layer 3 VPN Configuration

If you are implementing a Layer 3 virtual private network (VPN), you must configure the BGP routing protocol on the PE switches.

---

### Routing Instances for Layer 3 VPN

If you are implementing a Layer 3 VPN, you must configure a routing instance. A routing instance is a collection of routing tables, interfaces, and routing protocol parameters. The set of interfaces belongs to the routing tables, and the routing protocol parameters control the information in the routing tables.

QFX Series and EX4600 devices support VPN routing and forwarding (VRF) routing instances for Layer 3 VPNs.

Each routing instance has a unique name and a corresponding IP unicast table. For example, if you configure a routing instance with the name **my-instance**, its corresponding IP unicast table will be **my-instance.inet.0**. All routes for **my-instance** are installed in **my-instance.inet.0**.

---

### Routing Instances for Layer 2 VPN and Layer 3 VPN

If you are implementing a Layer 2 VPN or a Layer 3 VPN, you must configure a routing instance. A routing instance is a collection of routing tables, interfaces, and routing protocol parameters. The set of interfaces belongs to the routing tables, and the routing protocol parameters control the information in the routing tables.

QFX Series devices support the following types of routing instances:

- Layer 2 VPN—To support a Layer 2 VPN
- VPN routing and forwarding (VRF)—To support a Layer 3 VPN

Each routing instance has a unique name and a corresponding IP unicast table. For example, if you configure a routing instance with the name **my-instance**, its corresponding IP unicast table will be **my-instance.inet.0**. All routes for **my-instance** are installed in **my-instance.inet.0**.

### Ethernet Encapsulation for Layer 2 VPN

If you are implementing a Layer 2 VPN, you must also configure the physical layer encapsulation type on the customer edge interface and within the routing instance.

## Provider Switch

You must configure one or more provider switches as transit switches within the network to support the forwarding of MPLS packets. You can add provider switches without changing the configuration of the PE switches.

A provider switch does not analyze packets. It refers to an MPLS label forwarding table and swaps one label for another. The new label determines the next hop along the MPLS tunnel. A provider switch cannot perform push or pop operations.

## Components Required for All Switches in the MPLS Network

The following MPLS components are configured on both the PE switches and the provider switches:

- [Interior Gateway Protocol on page 5261](#)
- [Traffic Engineering on page 5262](#)
- [MPLS Protocol on page 5262](#)
- [RSVP on page 5262](#)
- [Family mpls on page 5262](#)

### Interior Gateway Protocol

MPLS works in coordination with OSPF as the interior gateway protocol (IGP). Therefore, you must configure OSPF as the IGP on the loopback interface and CE-facing interfaces of both the PE switches and the provider switches.

The CE-facing interfaces can be either Gigabit Ethernet or 10-Gigabit Ethernet interfaces, and they can be configured as either individual interfaces or as aggregated Ethernet interfaces.



**NOTE:** The CE-facing interfaces cannot be configured with VLAN tagging or a VLAN ID. When you configure them to belong to family **mpls**, they are removed from the default VLAN if they were members of that VLAN. They operate as an exclusive tunnel for MPLS traffic.

## Traffic Engineering

---

Traffic engineering maps traffic flows onto an existing physical topology and provides the ability to move traffic flow away from the shortest path selected by the IGP and to a potentially less congested physical path across a network.

Traffic engineering enables the selection of specific end-to-end paths to send given types of traffic through your network. You must configure OSPF traffic engineering on the PE switches and the provider switches.

## MPLS Protocol

---

You must enable the MPLS protocol on all switches that participate in the MPLS network and apply it to the core interfaces of both the PE and provider switches. You do not need to apply it to the loopback interface because the MPLS protocol uses the framework established by the RSVP signaling protocol to create LSPs. On the PE switches, the configuration of the MPLS protocol must also include the definition of an LSP.

## RSVP

---

RSVP is a signaling protocol that allocates and distributes labels throughout an MPLS network. RSVP sets up unidirectional paths between the ingress PE switch and the egress PE switch. RSVP makes the LSPs dynamic; it can detect topology changes and outages and establish new LSPs to allow traffic to move around a failure.

You must enable RSVP and apply it to the loopback interface and the core interface of both the PE and provider switches. The path message contains the configured information about the resources required for the LSP to be established.

When the egress PE switch receives the path message, it sends a reservation message back to the ingress PE switch. This reservation message is passed along from switch to switch along the same path as the original path message. Once the ingress PE switch receives this reservation message, an RSVP path is established.

The established LSP stays active as long as the RSVP session remains active. RSVP continues activity through the transmissions and responses to RSVP path and reservation messages. If the messages stop for three minutes, the RSVP session terminates and the LSP is lost.

RSVP runs as a separate software process in Junos OS and is not in the packet-forwarding path.

## Family mpls

---

You must configure the core interfaces used for MPLS traffic to belong to **family mpls**.



**NOTE:** You can enable **family mpls** on either individual interfaces or on aggregated Ethernet interfaces. You cannot enable it on tagged VLAN interfaces.

---

**Related Documentation**

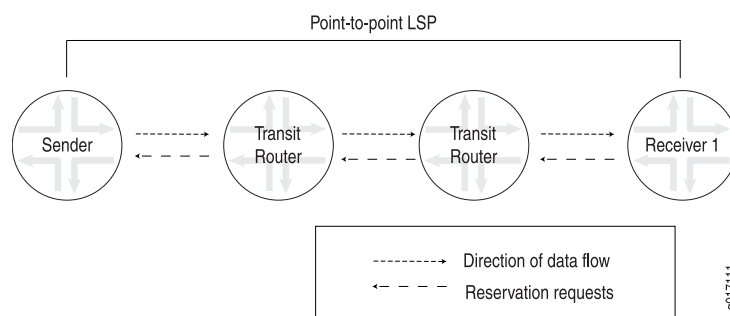
- [MPLS Feature Support on QFX Series and EX4600 Switches on page 4934](#)
- [Understanding Using MPLS-Based Layer 3 VPNs on Switches on page 4963](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4959](#)
- [Configuring MPLS on Provider Edge Switches on page 4981](#)
- [Configuring MPLS on Provider Switches on page 4985](#)
- [Configuring Rewrite Rules for MPLS EXP Classifiers on page 4992](#)
- [Configuring a Global MPLS EXP Classifier on page 4977](#)
- [Configuring Ethernet over MPLS \(L2 Circuit\) on page 4972](#)
- *Junos OS MPLS Applications Library for Routing Devices*
- *Junos OS VPNs Library for Routing Devices*

## RSVP Overview

The RSVP protocol is used by routers to deliver quality-of-service (QoS) requests to all nodes along data flow path(s) and to establish and maintain state for the requested service. RSVP requests generally result in resource reservations in each node along the data path. RSVP has the following attributes:

- Makes resource reservations for unidirectional data flows.
- Allows the receiver of a data flow to initiate and maintain the resource reservation used for that flow, as shown in [Figure 150](#).
- Maintains a soft state in routers and hosts, providing graceful support for dynamic membership changes and automatic adaptation to routing changes.
- Depends upon present and future routing protocols, but is not a routing protocol itself.
- Provides several reservation models or styles to fit a variety of applications.
- Supports both IPv4 and IPv6. Note, you can configure the Junos OS to tunnel IPv6 over an MPLS-based IPv4 network. For more information, see the *MPLS Applications Feature Guide for Routing Devices*.

**Figure 150: RSVP Reservation Request and Data Flow**



## MTU Signaling in RSVP

---

The maximum transmission unit (MTU) is the largest size packet or frame, in bytes, that can be sent in a network. An MTU that is too large might cause retransmissions. Too small an MTU might cause the router to send and handle relatively more header overhead and acknowledgments. There are default values for MTUs associated with various protocols. You can also explicitly configure an MTU on an interface.

When an LSP is created across a set of links with different MTU sizes, the ingress router does not know what the smallest MTU is on the LSP path. By default, the MTU for an LSP is 1,500 bytes.

If this MTU is larger than the MTU of one of the intermediate links, traffic might be dropped, because MPLS packets cannot be fragmented. Also, the ingress router is not aware of this type of traffic loss, because the control plane for the LSP would still function normally.

To prevent this type of packet loss in MPLS LSPs, you can configure MTU signaling in RSVP. This feature is described in RFC 3209. Juniper Networks supports the Integrated Services object for MTU signaling in RSVP. The Integrated Services object is described in RFCs 2210 and 2215. MTU signaling in RSVP is disabled by default.

To avoid packet loss due to MTU mismatches, the ingress router needs to do the following:

- Signal the MTU on the RSVP LSP—To prevent packet loss from an MTU mismatch, the ingress router needs to know what the smallest MTU value is along the path taken by the LSP. Once this MTU value is obtained, the ingress router can assign it to the LSP.
- Fragment packets—Using the assigned MTU value, packets that exceed the size of the MTU can be fragmented into smaller packets on the ingress router before they are encapsulated in MPLS and sent over the RSVP-signaled LSP.

Once both MTU signaling and packet fragmentation have been enabled on an ingress router, any route resolving to an RSVP LSP on this router uses the signaled MTU value. For information about how to configure this feature, see *Configuring MTU Signaling in RSVP*.

The following sections describe how MTU signaling in RSVP works:

- *How the Correct MTU Is Signaled in RSVP*
- *Determining an Outgoing MTU Value*
- *MTU Signaling in RSVP Limitations*

## Tunneling LDP LSPs in RSVP LSPs

---

You can tunnel LDP LSPs over RSVP LSPs. The following sections describe how tunneling of LDP LSPs in RSVP LSPs works:

- [Tunneling LDP LSPs in RSVP LSPs Overview on page 4805](#)
- [Label Operations on page 4806](#)



## Tunneling LDP LSPs in RSVP LSPs Overview

If you are using RSVP for traffic engineering, you can run LDP simultaneously to eliminate the distribution of external routes in the core. The LSPs established by LDP are tunneled through the LSPs established by RSVP. LDP effectively treats the traffic-engineered LSPs as single hops.

When you configure the router to run LDP across RSVP-established LSPs, LDP automatically establishes sessions with the router at the other end of the LSP. LDP control packets are routed hop-by-hop, rather than carried through the LSP. This routing allows you to use simplex (one-way) traffic-engineered LSPs. Traffic in the opposite direction flows through LDP-established LSPs that follow unicast routing rather than through traffic-engineered tunnels.

If you configure LDP over RSVP LSPs, you can still configure multiple OSPF areas and IS-IS levels in the traffic engineered core and in the surrounding LDP cloud.



**NOTE:** Beginning with Junos OS Release 15.1, multi-instance support is extended to LDP over RSVP tunneling for a virtual router routing instance. This allows splitting of a single routing and MPLS domain into multiple domains so that each domain can be scaled independently. BGP labeled unicast can be used to stitch these domains for service FECs. Each domain uses intra-domain LDP over RSVP LSP for MPLS forwarding.

### Related Documentation

- [Label Operations on page 4806](#)
- [Configuring a Hierarchy of RSVP LSPs to Tunnel Multiple RSVP LSPs Over a Single RSVP LSP](#)

## Configuring MPLS on Provider Edge Switches

To implement MPLS, you must configure two provider edge (PE) switches—an ingress PE switch and an egress PE switch—and at least one provider switch. You can configure the customer edge (CE) interfaces on the PE switches of the MPLS network using IP over MPLS.

This topic describes how to configure an ingress PE switch and an egress PE switch using IP over MPLS:

1. [Configuring the Ingress PE Switch on page 5266](#)
2. [Configuring the Egress PE Switch on page 5267](#)

## Configuring the Ingress PE Switch

To configure the ingress PE switch:

1. Configure an IP address for the loopback interface and the core interfaces:

```
[edit interfaces]
user@switch# set lo0 unit 0 family inet address 192.168.10.1/32
user@switch# set xe-0/0/5 unit 0 family inet address 10.1.5.1/24
user@switch# set xe-0/0/6 unit 0 family inet address 10.1.6.1/24
```



**NOTE:** You cannot use routed VLAN interfaces (RVIs) or Layer 3 subinterfaces as core interfaces.

2. Configure OSPF on the loopback interface and the core interfaces:



**NOTE:** You can use the switch address as an alternative to the loopback interface.

```
[edit protocols ospf]
user@switch# set area 0.0.0.0 interface lo0.0
user@switch# set area 0.0.0.0 interface xe-0/0/5.0
user@switch# set area 0.0.0.0 interface xe-0/0/6.0
```

3. Configure OSPF traffic engineering:

```
[edit protocols ospf]
user@switch# set traffic-engineering
```

4. Configure RSVP on the loopback interface and the core interfaces:

```
[edit protocols rsvp]
user@switch# set interface lo0.0
user@switch# set interface xe-0/0/5.0
user@switch# set interface xe-0/0/6.0
```

5. Configure MPLS traffic engineering.

```
[edit protocols mpls]
user@switch# set traffic-engineering
```

6. Configure MPLS on the core interfaces:

```
[edit protocols mpls]
user@switch# set interface xe-0/0/5.0
user@switch# set interface xe-0/0/6.0
```

7. Configure **family mpls** on the logical units of the core interfaces, thereby identifying the interfaces that will be used for forwarding MPLS packets:

```
[edit interfaces]
user@switch# set xe-0/0/5 unit 0 family mpls
user@switch# set xe-0/0/6 unit 0 family mpls
```

8. Configure a customer edge interface as a Layer 3 routed interface, specifying an IP address:

```
[edit interfaces]
user@switch# set xe-0/0/3 unit 0 family inet address 121.100.10.1/16
```

9. Configure this Layer 3 customer edge interface for the routing protocol:

- ```
[edit]
user@switch# set protocols ospf area 0.0.0 interface xe-0/0/3.0
```
10. Configure an LSP on the ingress PE switch (192.168.10.1) to send IP packets over MPLS to the egress PE switch (192.168.12.1):


```
[edit protocols mpls]
user@switch# set label-switched-path lsp_1 to 192.168.12.1
```
 11. Disable constrained-path LSP computation for this LSP:


```
[edit protocols mpls]
user@switch# set label-switched-path lsp_1 no-cspf
```
 12. Configure a static route from the ingress PE switch to the egress PE switch, thereby indicating to the routing protocol that the packets will be forwarded over the MPLS LSP that has been set up to that destination:


```
[edit routing-options]
user@switch# set static route 2.2.2.0/24 next-hop 192.168.10.1
user@switch# set static route 2.2.2.0/24 resolve
```

Configuring the Egress PE Switch

To configure the egress PE switch:

1. Configure an IP address for the loopback interface and the core interfaces:

```
[edit interfaces]
user@switch# set lo0 unit 0 family inet address 192.168.12.1/32
user@switch# set xe-0/0/5 unit 0 family inet address 10.1.20.1/24
user@switch# set xe-0/0/6 unit 0 family inet address 10.1.21.1/24
```



NOTE: You cannot use routed VLAN interfaces (RVIs) or Layer 3 subinterfaces as core interfaces.

2. Configure OSPF on the loopback interface and the core interfaces:



NOTE: You can use the switch address as an alternative to the loopback interface.

- ```
[edit protocols ospf]
user@switch# set area 0.0.0.0 interface lo0.0
user@switch# set area 0.0.0.0 interface xe-0/0/5.0
user@switch# set area 0.0.0.0 interface xe-0/0/6.0
```
3. Configure RSVP on the loopback interface and the core interfaces:
 

```
[edit protocols rsvp]
user@switch# set rsvp interface lo0.0
user@switch# set rsvp interface xe-0/0/5.0
user@switch# set rsvp interface xe-0/0/6.0
```
  4. Configure MPLS on the core interfaces:
 

```
[edit protocols mpls]
user@switch# set interface xe-0/0/5.0
user@switch# set interface xe-0/0/6.0
```
  5. Configure **family mpls** on the logical units of the core interfaces, thereby identifying the interfaces that will be used for forwarding MPLS packets:

```
[edit interfaces]
```

```
user@switch# set xe-0/0/5 unit 0 family mpls
```

```
user@switch# set xe-0/0/6 unit 0 family mpls
```

6. Configure a customer edge interface as a Layer 3 routed interface, specifying an IP address:

```
[edit interfaces]
```

```
user@switch# set xe-0/0/3 unit 0 family inet address 2.2.2.1/16
```

7. Configure this Layer 3 customer edge interface for the routing protocol:

```
[edit]
```

```
user@switch# set protocols ospf area 0.0.0 interface xe-0/0/3
```

8. Configure an LSP on the egress PE switch (192.168.12.1) to send IP packets over MPLS to the ingress PE switch (192.168.10.1):

```
[edit protocols mpls]
```

```
user@switch# set label-switched-path lsp_2 to 192.168.10.1
```

9. Disable constrained-path LSP computation for this LSP:

```
[edit protocols mpls]
```

```
user@switch# set label-switched-path lsp_2 no-cspf
```

10. Configure a static route from the ingress PE switch to the egress PE switch, thereby indicating to the routing protocol that the packets will be forwarded over the MPLS LSP that has been set up to that destination:

```
[edit routing-options]
```

```
user@switch# set static route 121.121.121.0/24 next-hop 192.168.12.1
```

```
user@switch# set static route 121.121.121.0/24 resolve
```

**Related  
Documentation**

- [MPLS Configuration Guidelines on page 5012](#)
- [Configuring MPLS on Provider Switches on page 4985](#)
- [MPLS Feature Support on QFX Series and EX4600 Switches on page 4934](#)
- [Understanding MPLS Components for QFX Series and EX4600 Switches on page 4946](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4959](#)

## Configuring MPLS on Provider Switches

To implement MPLS, you must configure at least one provider switch as a transit switch for the MPLS packets.

MPLS requires the configuration of an interior gateway protocol (OSPF) and a signaling protocol (RSVP) on the core interfaces and the loopback interface of all the switches. This procedure includes the configuration of OSPF on the provider switch.

To configure the provider switch, complete the following tasks:

1. Configure OSPF on the loopback and core interfaces:



**NOTE:** You can use the switch address as an alternative to the loopback interface.

```
[edit protocols ospf]
user@switch# set area 0.0.0.0 interface lo0.0
user@switch# set area 0.0.0.0 interface xe-0/0/5.0
user@switch# set area 0.0.0.0 interface xe-0/0/6.0
user@switch# set area 0.0.0.0 interface ae0
```



**NOTE:** You cannot use routed VLAN interfaces (RVIs) or Layer 3 subinterfaces as core interfaces.

2. Configure MPLS on the core interfaces:

```
[edit protocols mpls]
user@switch# set interface xe-0/0/5.0
user@switch# set interface xe-0/0/6.0
user@switch# set interface ae0
```

3. Configure RSVP on the loopback interface and the core interfaces:

```
[edit protocols rsvp]
user@switch# set interface lo0.0
user@switch# set interface xe-0/0/5.0
user@switch# set interface xe-0/0/6.0
user@switch# set interface ae0
```

4. Configure an IP address for the loopback interface and the core interfaces:

```
[edit interfaces]
user@switch# set lo0 unit 0 family inet address 127.1.1.1/32
user@switch# set xe-0/0/5 unit 0 family inet address 10.1.5.1/24
user@switch# set xe-0/0/6 unit 0 family inet address 10.1.6.1/24
user@switch# set ae0 unit 0 family inet address 10.1.9.2/24
```

5. Configure **family mpls** on the logical units of the core interfaces, thereby identifying the interfaces that will be used for forwarding MPLS packets:

```
[edit interfaces]
user@switch# set xe-0/0/5 unit 0 family mpls
user@switch# set xe-0/0/6 unit 0 family mpls
user@switch# set ae0 unit 0 family mpls
```



**NOTE:** You can configure **family mpls** on either individual interfaces or aggregated Ethernet interfaces. You cannot configure it on tagged VLAN interfaces.

#### Related Documentation

- [Configuring MPLS on Provider Edge Switches on page 4981](#)
- [MPLS Configuration Guidelines on page 5012](#)
- [MPLS Feature Support on QFX Series and EX4600 Switches on page 4934](#)
- [Understanding MPLS Components for QFX Series and EX4600 Switches on page 4946](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4959](#)

## Verifying That MPLS Is Working Correctly

To verify that MPLS is working correctly, perform the following tasks:

1. [Verifying the Physical Layer on the Switches on page 5270](#)
2. [Verifying the Routing Protocol on page 5271](#)
3. [Verifying the Core Interfaces Being Used for the MPLS Traffic on page 5271](#)
4. [Verifying RSVP on page 5271](#)

### Verifying the Physical Layer on the Switches

**Purpose** Verify that the interfaces are up. Perform this verification task on each of the switches.

**Action** `user@switch> show interfaces xe-* terse`

| Interface  | Admin | Link | Proto        | Local       | Remote |
|------------|-------|------|--------------|-------------|--------|
| xe-0/0/0   | up    | up   |              |             |        |
| xe-0/0/0.0 | up    | up   |              |             |        |
| xe-0/0/1.0 | up    | up   |              |             |        |
| xe-0/0/2.0 | up    | up   |              |             |        |
| xe-0/0/3.0 | up    | up   | inet         | 2.2.2.1/16  |        |
| xe-0/0/4.0 | up    | up   |              |             |        |
| xe-0/0/5.0 | up    | up   | inet<br>mpls | 10.1.5.1/24 |        |
| xe-0/0/6.0 | up    | up   | inet<br>mpls | 10.1.6.1/24 |        |

**Meaning** The `show interfaces terse` command displays status information about the 10-Gigabit Ethernet interfaces on the switch. This output verifies that the interfaces are **up**. The output for the protocol family (Proto column) of the core interfaces (xe-0/0/5.0 and xe-0/0/6.0), shows that these interfaces are configured as both **inet** and **mpls**. The **Local** column for the core interfaces shows the IP address configured for these interfaces.

## Verifying the Routing Protocol

**Purpose** Verify the state of the configured routing protocol. You should perform this verification task on each of the switches. The state should be **Full**. If you have configured OSPF as the routing protocol, use the **show ospf neighbor** command to verify that the routing protocol is communicating with the switch neighbors.

**Action** user@switch> **show ospf neighbor**

| Address   | Interface | State | ID          | Pri | Dead |
|-----------|-----------|-------|-------------|-----|------|
| 127.1.1.1 | xe-0/0/5  | Full  | 10.10.10.10 | 128 | 39   |

**Meaning** The **show ospf neighbor** command displays the status of the routing protocol that has been configured on this switch. The output shows that the state is **Full**, meaning that the routing protocol is operating correctly—that is, hello packets are being exchanged between directly connected neighbors. For additional information on checking and monitoring routing protocols, see the [Junos OS Routing Protocols and Policies Command Reference](#).

## Verifying the Core Interfaces Being Used for the MPLS Traffic

**Purpose** Verify that the state of the MPLS interface is **Up**. You should perform this verification task on each of the switches.

**Action** user@switch> **show mpls interface**

| Interface | State | Administrative groups |
|-----------|-------|-----------------------|
| ge-0/0/5  | Up    | <none>                |
| ge-0/0/6  | Up    | <none>                |

**Meaning** The **show mpls interface** command displays the status of the core interfaces that have been configured to belong to **family mpls**. This output shows that the interface configured to belong to **family mpls** is up.

## Verifying RSVP

**Purpose** Verify the state of the RSVP session. You should perform this verification task on each of the switches.

```
user@switch> show mpls session
```

```
Ingress RSVP: 1 sessions
To From State Rt Style Labelin Labelout LSPname
127.1.1.3 127.1.1.1 Up 0 1 FF - 300064 lsp_to_pe2_ge1
Total 1 displayed, Up 1, Down 0
```

```
Egress RSVP: 1 sessions
To From State Rt Style Labelin Labelout LSPname
127.1.1.1 127.1.1.3 Up 0 1 FF 299968 - lsp_to_pe1_ge1
Total 1 displayed, Up 1, Down 0
```

```
Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

**Meaning** This output confirms that the RSVP sessions are up.

**Related Documentation**

- [Configuring MPLS on Provider Edge Switches on page 4981](#)
- [Configuring MPLS on Provider Switches on page 4985](#)

---

## Dynamic Bandwidth Management Using Container LSP Overview

RSVP LSPs with the autobandwidth feature are increasingly deployed in networks to meet traffic engineering needs. However, the current traffic engineering solutions for point-to-point LSPs are inefficient in terms of network bandwidth utilization, mainly because the ingress routers originating the RSVP LSPs either try to fit the LSPs along a particular path without creating parallel LSPs, or do not interact with the other routers in the network and probe for additional available bandwidth.

This feature provides an ingress router with the capability of acquiring as much network bandwidth as possible by creating parallel LSPs dynamically.

- [Understanding RSVP Multipath Extensions on page 5272](#)
- [Junos OS RSVP Multipath Implementation on page 5273](#)
- [Current Traffic Engineering Challenges on page 5274](#)
- [Using Container LSP as a Solution on page 5277](#)
- [Junos OS Container LSP Implementation on page 5279](#)
- [Configuration Statements Supported for Container LSPs on page 5294](#)
- [Impact of Configuring Container LSPs on Network Performance on page 5298](#)
- [Supported and Unsupported Features on page 5298](#)

### Understanding RSVP Multipath Extensions

The RSVP multipath extensions proposed in the IETF [KOMPELLA-MLSP] allow the setup of traffic engineered multipath label-switched paths (container LSPs). The container LSPs, in addition to conforming to traffic engineering constraints, use multiple independent paths from a source to a destination, thereby facilitating load balancing of



traffic. The multipath extensions require changes to the RSVP-TE protocol and allow for merging of labels at the downstream nodes (similar to LDP), which also helps in preserving forwarding resources.

The multipath extensions to RSVP provide the following benefits:

- Ease of configuration. Typically, multiple RSVP LSPs are configured for either load balancing or bin packing. With a container LSP, there is a single entity to provision, manage, and monitor LSPs. Changes in topology are handled easily and autonomously by the ingress LSP, by adding, changing, or removing member LSPs to rebalance traffic, while maintaining the same traffic engineering constraints.
- RSVP equal-cost multipath (ECMP) inherits the standard benefits of ECMP by absorbing traffic surges.
- Multipath traffic engineering allows for better and complete usage of network resources.
- Knowing the relationship among LSPs helps in computing diverse paths with constraint-based routing. It allows adjustment of member LSPs while other member LSPs continue to carry traffic.
- The intermediate routers have an opportunity to merge the labels of member LSPs. This reduces the number of labels that need to get added to the forwarding plane and in turn reduces the convergence time.

If the number of independent ECMP paths is huge, label merging overcomes the platform limitations on maximum (ECMP) next hops. With point-to-point RSVP LSPs that require link or node protection, the next hops are doubled as each LSP is programmed with both primary and backup next hops. RSVP multipath (or ECMP) obviates the need for backup next hops.

- When there is a link failure, the router upstream to the link failure can distribute traffic from the failed link to the remaining ECMP branches, obviating the need for bypass LSPs. The bypass LSP approach not only requires more state when signaling backup LSPs, but also suffers from scaling issues that result in merge-point timing out a protected path state block (PSB) before point of local repair (PLR) gets a chance to signal the backup LSP.

## Junos OS RSVP Multipath Implementation

In order to deploy RSVP multipath (ECMP) in a network, all the nodes through which ECMP LSPs pass must understand RSVP ECMP protocol extensions. This can be a challenge, especially in a multivendor networks.

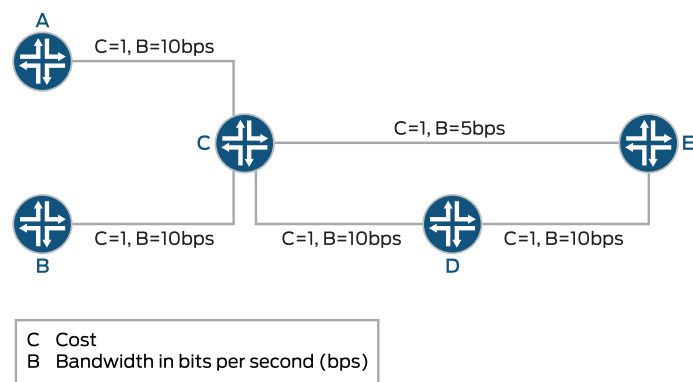
Junos OS implements the RSVP multipath extensions without the need for protocol extensions. A single container LSP, which has the characteristics of ECMP and RSVP TE, is provisioned. A container LSP consists of several member LSPs and is set up between the ingress and egress routing device. Each member LSP takes a different path to the same destination. The ingress routing device is configured with all the required parameters to compute the RSVP ECMP LSP. The parameters configured to compute a set of RSVP point-to-point LSPs can be used by the ingress routing device to compute the container LSP as well.

## Current Traffic Engineering Challenges

The main challenge for traffic engineering is to cope with the dynamics of both topology and traffic demands. Mechanisms are needed that can handle traffic load dynamics in scenarios with sudden changes in traffic demand and dynamically distribute traffic to benefit from available resources.

Figure 151 illustrates a sample network topology with all the LSPs having the same hold and setup priorities, and admission control restricted on the ingress router. All the links are annotated with a tuple (cost and capacity).

Figure 151: Sample Topology



Some of the traffic engineering problems seen in [Figure 151](#) are listed here:

- **Bin Packing**

This problem arises because of a particular order in which LSPs are signaled. The ingress routers might not be able to signal some LSPs with required demands although bandwidth is available in the network, leading to under-utilization of link capacity.

For example, the following LSPs arrive in the sequence mentioned in [Table 417](#).

**Table 417: LSP Sequence Order for Bin Packing**

| Time | Source | Destination | Demand | ERO     |
|------|--------|-------------|--------|---------|
| 1    | A      | E           | 5      | A-C-D-E |
| 2    | B      | E           | 10     | No ERO  |

The LSP originating at Router B is not routable as constraint-based routing fails to find a feasible path. However, if Router B is signaled first, both the LSPs are routable. Bin packing happens because of lack of visibility of individual per-LSP, per-device bandwidth demands at the ingress routing device.

Bin packing can also happen when there is no requirement for ordering of LSPs. For example, if there is an LSP with demand X and there are two different paths to the destination from the ingress router with available bandwidths Y1 and Y2, such that Y1 is less than X, Y2 is less than X, and Y1 plus Y2 is greater than or equal to X.

In this case, even though there are enough network resources in terms of available bandwidth to satisfy the aggregate LSP demand X, the LSP might not be signaled or re-optimized with the new demand. In [Figure 151](#), with container LSP support, the ingress B creates two LSPs each of size 5 when demand 10 is posed. One LSP is routed along B-C-E and another one along B-C-D-E.

- **Deadlock**

Considering [Figure 151](#), the LSPs follow the sequence mentioned in [Table 418](#).

**Table 418: LSP Sequence Order for Deadlock**

| Time | Source | Destination | Demand  | ERO     | Event                                             |
|------|--------|-------------|---------|---------|---------------------------------------------------|
| 1    | A      | E           | 2       | A-C-D-E | Constraint-based routing with RSVP signaling      |
| 2    | B      | E           | 2       | B-C-D-E | Constraint-based routing with RSVP signaling      |
| 3    | A      | E           | 2 to 20 | A-C-D-E | Constraint-based routing fails, no RSVP signaling |

At time 3, the demand on LSP from A to E increases from 2 to 20. If autobandwidth is configured, the change does not get detected until the adjustment timer expires. In the absence of admission control at A, the increased traffic demand might cause traffic to drop on other LSPs that share common links with the mis-behaving LSP.

This happens due to the following reasons:

- Lack of global state at all the ingress routers
- Signaling of mis-behaving demands
- Tearing down of mis-behaving demands

With container LSP configured, ingress A has more chances of splitting the load (even incrementally if not fully) across multiple LSPs. So, LSP from A is less likely to see prolonged traffic loss.

- **Latency Inflation**

Latency inflation is caused by the autobandwidth and other LSPs parameters. Some of the other factors that contribute to latency inflation include:

- LSP priority

LSPs choose longer paths because shorter paths between data centers located in the same city can be congested. The bandwidth on the shorter paths can get exhausted by equal or higher priority LSPs. Due to periodic LSP optimization by autobandwidth, LSP can get rerouted to a higher delay path. When many LSPs undergo less than optimal path selection, they can potentially form a chain of dependencies. Modifying the LSP priorities dynamically is a workaround to the issue; however, dynamically adjusting LSP priorities to find shorter paths is a challenging task.

- All or Nothing policy

When the demand on an LSP increases and at least one of the links along the shorter path is close to its reservation limit, LSP optimization can force the LSP to move to a longer latency path. LSP has to traverse a long path even though the short path is capable of carrying most of the traffic.

- Minimum and maximum bandwidth

Minimum and maximum bandwidth specify the boundaries for LSP sizes. If minimum bandwidth is small, an LSP is more prone to autobandwidth adjustment because a small change in bandwidth is enough to cross the threshold limits. LSPs might reroute although bandwidth is available. On the other hand, if the minimum bandwidth is large, network bandwidth might be wasted. If the maximum bandwidth value is small, a large number of LSPs might be needed at the ingress router to accommodate the application demand. If the maximum bandwidth is large, the LSPs can grow larger in size. Such LSPs can suffer because of an all or nothing policy.

- Autobandwidth adjustment threshold

Bandwidth threshold dictates if LSPs need to be re-optimized and resized. If the value is small, LSPs are frequently re-optimized and rerouted. That might cause CPU spike because applications or protocols, such as BGP resolving over the LSPs, might keep the Routing Engine busy doing next-hop resolution. A large value might make an LSP immobile. With container LSP configured, an LSP is less likely to get subjected

to one or no policy. An ingress router originates multiple LSPs, although not all LSPs potentially traverse high latency paths.

- **Predictability**

Service providers often want predictable behavior in terms of how LSPs get signaled and routed. Currently, without any global coordination, it is difficult to set up the same set of LSPs in a predictable way. Consider the two different orderings in [Table 419](#) and [Table 420](#). The ERO that an LSP uses depends on its signaling time.

**Table 419: LSP Sequence Order for Predictability**

| Time | Source | Destination | Demand | ERO     |
|------|--------|-------------|--------|---------|
| 1    | A      | E           | 5      | A-C-D-E |
| 2    | B      | E           | 5      | B-C-E   |

**Table 420: LSP Sequence Order for Predictability**

| Time | Source | Destination | Demand | ERO     |
|------|--------|-------------|--------|---------|
| 1    | B      | E           | 5      | B-C-E   |
| 2    | A      | E           | 5      | A-C-D-E |

Container LSP does not directly help LSPs find predictable EROs. If LSPs are getting rerouted because of an all or no policy without container LSP configured, such LSPs might see less churn if container LSPs are configured, because smaller LSPs have better chances of finding a shorter or same path.

## Using Container LSP as a Solution

A container LSP can be used as a solution to the challenges faced by the current traffic engineering features. Considering [Figure 151](#), when the demand X on a container LSP increases with the network capacity (max-flow) being more than the demand, the following approaches come into effect with a container LSP:

- [Accommodating the New Demand X on page 5277](#)
- [Creating New LSPs to Meet Demand X on page 5278](#)
- [Assigning Bandwidth to the New LSPs on page 5278](#)
- [Controlling the LSP Paths on page 5278](#)

### Accommodating the New Demand X

In the current implementation, autobandwidth attempts to re-signal an LSP with the new demand X and follows the all or nothing policy as mentioned earlier.

The container LSP approach computes several small (smaller than demand X) bandwidth LSPs such that the aggregate bandwidth is not less than X, and the ingress router performs this adjustment periodically. One of the triggers to create new LSPs or to delete old LSPs

can be changed in aggregate bandwidth. The ingress router then load-balances the incoming traffic across the newly created LSPs.

### Creating New LSPs to Meet Demand X

---

Although the number of new LSPs created can be a maximum of the allowed configurable limit, there is not much benefit from these LSPs once the number of LSPs exceeds the number of possible diverse paths or equal-cost multipaths (ECMPs). The benefit of creating the smaller LSPs is seen when an ingress router uses the newly created LSPs for load-balancing traffic. This, however, depends on the network topology and state.

Creating multiple parallel LSPs by all the ingress routers in the network can lead to scaling issues at the transit routers. Thus, the number of new LSPs to be created depends on the size of the individual LSPs and the given aggregate demand, X in this case.

### Assigning Bandwidth to the New LSPs

---

In general, there can be a number of heuristics to allocate bandwidths to the newly created LSPs. An ingress router can solve an optimization problem in which it can maximize a given utility function. The output of an optimization problem is assigning optimal bandwidth values. However, to solve an optimization problem, the number of newly created LSPs has to be fixed. Therefore, it is complex to optimize the number and size of each LSP. Thus, to simplify the problem, the same amount of bandwidth is assumed for all the newly created LSPs, and then the number of required LSPs is computed.

### Controlling the LSP Paths

---

The flexibility to control the LSP paths is expressed in terms of the configuration for point-to-point LSPs and container LSPs. Controlling the LSP paths using the configuration parameters can be applied under two different aspects:

- **Topology**—There are no topology constraints with this feature. Each member LSP is treated like a point-to-point LSP and is re-optimized individually. An ingress router does not try to compute equal IGP cost paths for all its LSPs, but instead it computes paths for all the LSPs using current traffic engineering database information. While computing a path, constraint-based routing adheres to any constraints specified through the configuration, although there is no change in the constraint-based routing method for path computation.
- **When to create a new LSP**—When to create a new LSP can be explicitly specified. By default, an ingress router periodically computes the aggregate traffic rate by adding up the traffic rate of all the individual LSPs. Looking at the aggregate bandwidth and configuration, the ingress router recomputes the number of LSPs and the bandwidths of the LSPs. The new LSPs are then signaled or the existing LSPs are re-signaled with the updated bandwidth. Instead of looking at the instantaneous aggregate rate, the ingress routers can compute an average (of aggregates) over some duration by removing outlier samples (of aggregates). Managing the LSPs that remain outstanding and active by considering aggregate bandwidth is more scalable than creating the new LSPs based on the usage of a particular LSP. The intervals and thresholds can be configured to track the aggregate traffic and trigger adjustment. These dynamic LSPs co-exist and interoperate with per-LSP autobandwidth configuration.

## Junos OS Container LSP Implementation

A container LSP is an ECMP TE LSP that acts like a container LSP consisting of one or more member LSPs. A point-to-point TE LSP is equivalent to a container LSP with a single member LSP. Member LSPs are added to the container LSP through a process called splitting, and removed from the container LSP through a process called merging.

- [Container LSP Terminology on page 5279](#)
- [LSP Splitting on page 5280](#)
- [LSP Merging on page 5282](#)
- [Node and Link Protection on page 5284](#)
- [Naming Convention on page 5284](#)
- [Normalization on page 5285](#)
- [Constraint-Based Routing Path Computation on page 5290](#)
- [Sampling on page 5291](#)
- [Support for NSR, IPG-FA, and Static Routes on page 5291](#)

### Container LSP Terminology

The following terms are defined in the context of a container LSP:

- **Normalization**—An event occurring periodically when an action is taken to adjust the member LSPs, either to adjust their bandwidths, their number, or both. A normalization process is associated with a sampling process and periodically estimates aggregate utilization of a container LSP.
- **Nominal LSP**—The instance of a container LSP that is always present.
- **Supplementary LSP**—The instances or sub-LSPs of a container LSP, which are dynamically created or removed.

Autobandwidth is run over each of the member LSPs, and each LSP is resized according to the traffic it carries and the autobandwidth configuration parameters. The aggregate demand on a container LSP is tracked by adding up the bandwidth across all the member LSPs.

- **Minimum signaling-bandwidth**—The minimum bandwidth with which a member LSP is signaled at the time of normalization or initialization. This could be different from the minimum-bandwidth defined under autobandwidth.
- **Maximum signaling-bandwidth**—The maximum bandwidth with which a member LSP is signaled at the time of normalization or initialization. This could be different from the maximum-bandwidth defined under autobandwidth.
- **Merging-bandwidth**—Specifies the lower bandwidth threshold on the aggregate bandwidth usage, such that if the aggregate usage falls below this value, the ingress router merges the member LSPs at the time of normalization.
- **Splitting-bandwidth**—Specifies the upper bandwidth threshold on the aggregate bandwidth usage, such that if the aggregate usage exceeds this value, the ingress router splits the member LSPs at the time of normalization.

- **Aggregate minimum-bandwidth**—Sum of merging-bandwidth of the current active member LSPs. This minimum bandwidth is different from the autobandwidth minimum-bandwidth.
- **Aggregate maximum-bandwidth**—Sum of the splitting-bandwidth of the current active member LSPs. This maximum bandwidth is different from the autobandwidth maximum-bandwidth.

---

## LSP Splitting

- [Operational Overview on page 5280](#)
- [Operational Constraints on page 5281](#)
- [Supported Criteria on page 5281](#)
- [Splitting Triggers on page 5282](#)

### *Operational Overview*

The LSP splitting mechanism enables an ingress router to create new member LSPs or to re-signal existing LSPs with different bandwidths within a container LSP when a demand X is placed on the container LSP. With LSP splitting enabled, an ingress router periodically creates a number of LSPs (by signaling new ones or re-signaling existing ones) to accommodate a new aggregate demand X. In the current implementation, an ingress router tries to find an LSP path satisfying a demand X and other constraints. If no path is found, either the LSP is not signaled or it remains up, but with the old reserved bandwidth.

Between two normalization events (splitting or merging), individual LSPs might get re-sigaled with different bandwidths due to the autobandwidth adjustments. If a container LSP is not configured with autobandwidth, then the member LSPs are signaled with the static bandwidth value, if configured. There is no dynamic splitting in this case, as there is no dynamic estimation of aggregate bandwidth. The splitting adjustments with a specific bandwidth value can be manually triggered.



#### NOTE:

Be aware of the following considerations for LSP splitting:

- After LSP splitting, the ingress router continues to inject one forwarding adjacency. Forwarding adjacencies are not supported in IGP for this feature.
  - Between two normalization events, two LSPs might have different bandwidths subjected to autobandwidth constraints.
  - After LSPs are split (or merged), make-before-break uses the fixed filter (FF) style sharing unless the adaptive option is configured. However, two different LSPs do not do the shared explicit (SE) style sharing for this feature.
  - When LSPs are re-sigaled with modified bandwidths, some of the LSPs might not get signaled successfully, leading to failover options.
-



### ***Operational Constraints***

LSP splitting has the following operational constraints:

- LSP bandwidth—Although there are a number of ways to allocate bandwidth values to the LSPs, the Junos OS implementation supports only an equal-bandwidth allocation policy when normalization is done, wherein all the member LSPs are signaled or re-signaled with equal bandwidth.
- Number of LSPs—If an ingress router is configured to have a minimum number of LSPs, it maintains the minimum number of LSPs even if the demand can be satisfied with less than the minimum number of LSPs. In case the ingress router is unable to do constraint-based routing for computations on the sufficient number of LSPs or signal sufficient number of LSPs, the ingress router resorts to a number of fallback options.

By default, an incremental approach is supported as a fallback option (unless configured differently), where an ingress router makes attempts to bring up the sufficient number of LSPs, such that the new aggregate bandwidth exceeds the old aggregate bandwidth (and is as close to the desired demand as possible). The ingress router then load-balances traffic using the LSPs. The LSPs that could not be brought up are removed by the ingress router.

### ***Supported Criteria***

When a container LSP signals a member LSP, the member LSP gets signaled with minimum-signaling-bandwidth. Since each member LSP is configured with autobandwidth, between two normalization events, each LSP can undergo autobandwidth adjustment multiple times. As the traffic demand increases, the ingress router creates additional supplementary LSPs. All member LSPs are used for ECMP, so they should roughly have the same reserved bandwidth after normalization.

For example, if there are K LSPs signaled after normalization, each LSP is signaled with equal bandwidth B. The total aggregate bandwidth reserved is B.K, where B satisfies the following condition:

- Minimum signaling-bandwidth is less than or equal to B, which in turn is less than or equal to the maximum signaling-bandwidth  
(minimum-signaling-bandwidth  $\leq$  B  $\leq$  maximum-signaling-bandwidth)

Until the next normalization event, each member LSP undergoes several autobandwidth adjustments. After any autobandwidth adjustment, if there are N LSPs with reserved bandwidths  $b_i$ , where  $i=1,2,\dots, N$ , each  $b_i$  should satisfy the following condition:

- Minimum bandwidth is less than or equal to  $b_i$ , which in turn is less than or equal to the maximum bandwidth  
(minimum-bandwidth  $\leq b_i \leq$  maximum-bandwidth)

Both the above-mentioned conditions are applicable for per member LSP (nominal and supplementary), and essentially have the reserved bandwidth to exist within a range.

### ***Splitting Triggers***

Every time the normalization timer expires, the ingress router decides if LSP splitting is required. The ingress router works with the aggregate bandwidth instead of the individual LSP bandwidths. The following two variables are defined for aggregate bandwidth:

- **Current-Aggr-Bw**—Sum of reserved bandwidths of all current member LSPs.
- **New-Aggr-Bw**—Sum of traffic rates on all current member LSPs based on sampling.

Taking for example, if there are N member LSPs in the network at the time of normalization, the two approaches to trigger LSP splitting are as follows:

- Absolute trigger—LSP splitting is performed when **New-Aggr-Bw** is greater than **Aggregate-maximum-bandwidth**.

(**New-Aggr-Bw** > **Aggregate-maximum-bandwidth**)

- Relative trigger—The **Current-Aggr-Bw** is compared with **New-Aggr-Bw** at the ingress routing device. LSP splitting is performed when the difference in the bandwidth amount is off by a threshold.

$([1-a] \times \text{Current-Aggr-Bw} < \text{New-Aggr-Bw} < [1+a] \times \text{Current-Aggr-Bw})$ , where  $0 \leq a \leq 1$

When **New-Aggr-Bw** is greater than or equal to  $[1+a]$  multiplied by **Current-Aggr-Bw**, the ingress routing device does not perform normalization, but instead LSP splitting is done. However, when both LSP splitting and LSP merging are configured on the ingress router, LSP splitting is triggered on the ingress router when one of the two conditions is satisfied.

---

### ***LSP Merging***

- [Operational Overview on page 5282](#)
- [Operational Constraints on page 5283](#)
- [Merging Triggers on page 5283](#)

#### ***Operational Overview***

Junos OS supports two kinds of LSPs – CLI-configured LSPs and dynamically created LSPs. The CLI-configured LSPs are created manually and remain in the system until the configuration is modified. The dynamic LSPs are created dynamically by next generation MVPN, BGP virtual private LAN service (VPLS), or LDP, based on a template configuration, and are removed from the system when not used by any application for a certain duration. LSP merging follows a similar approach as dynamic LSPs.

LSP merging enables an ingress routing device to dynamically eliminate some member LSPs of the container LSP so less state information is maintained in the network. If an ingress router provisions several member LSPs between the ingress and egress routers, and there is an overall reduction in aggregate bandwidth (resulting in some LSPs being under-utilized), the ingress router distributes the new traffic load among fewer LSPs.

Although there are a number of ways to merge the member LSPs, Junos OS supports only overall-merge when normalization is being performed. An ingress router considers the aggregate demand and the minimum (or maximum) number of LSPs and revises the number of LSPs that should be active at an ingress routing device. As a result, the following can take place periodically as the normalization timer fires:

- Re-signaling some of the existing LSPs with updated bandwidth
- Creating new LSPs
- Removing some of the existing LSPs

### ***Operational Constraints***

If a container LSP is not configured with autobandwidth, then the member LSPs are signaled with the static bandwidth value, if configured. LSP merging does not happen because there is no dynamic estimation of aggregate bandwidth. However, a manual trigger for splitting and adjusting with a specific bandwidth value can be configured.



#### **NOTE:**

- Nominal LSPs are never deleted as part of LSP merging.
- Before deleting an LSP, the LSP is made inactive, so that traffic shifts to other LSPs before removing the LSP. This is because RSVP sends PathTear before deleting routes and next hops from the Packet Forwarding Engine.
- When member LSPs are re-signaled with modified bandwidth, it might happen that some LSPs do not get signaled successfully.

### ***Merging Triggers***

Every time the normalization timer expires, the ingress router decides if LSP merging is required. The ingress router works with the aggregate bandwidth instead of the individual LSP bandwidths. The following two variables are defined for aggregate bandwidth:

- **Current-Aggr-Bw**—Sum of reserved bandwidths of all current member LSPs.
- **New-Aggr-Bw**—Sum of traffic rates on all current member LSPs based on sampling.

For example, if there are N member LSPs in the network at the time of normalization, the two approaches to trigger LSP merging are as follows:

- Absolute trigger—LSP merging is performed when **New-Aggr-Bw** is less than **Aggregate-minimum-bandwidth**.

(**New-Aggr-Bw** < **Aggregate-maximum-bandwidth**)

- Relative trigger—The **Current-Aggr-Bw** is compared with **New-Aggr-Bw** at the ingress routing device. LSP merging is performed when the difference in the bandwidth amount is off by a threshold.

$$([1-a] \times \text{Current-Aggr-Bw} < \text{New-Aggr-Bw} < [1+a] \times \text{Current-Aggr-Bw}, \text{ where } 0 \leq a \leq 1)$$

When the **New-Aggr-Bw** value is less than or equal to  $[1+a]$  multiplied by the **Current-Aggr-Bw** value, the ingress routing device does not perform normalization, but instead LSP merging is done. However, when both LSP splitting and LSP merging are configured on the ingress router, LSP splitting is triggered on the ingress router when one of the two conditions is satisfied.

### Node and Link Protection

---

Junos OS supports the following mechanisms for node and link protection:

- Fast-reroute
- Link protection
- Node-link protection

Only one of the above-mentioned modes of protection can be configured on an ingress routing device at any given time. All member LSPs (nominal and supplementary) use the same mode of protection that is configured.

### Naming Convention

---

While configuring a container LSP, a name is assigned to the LSP. The name of a nominal and a supplementary LSP is formed by adding the configured-name suffix and an auto-generated suffix to the name of the container LSP. The name of the container LSP is unique and is checked for accuracy during the configuration parsing. The container LSP name should uniquely identify parameters, such as the ingress and egress router names.



**NOTE:** A container LSP member LSP and a point-to-point LSP on an ingress routing device cannot have the same LSP name.

The container LSPs follow a number-based LSP naming convention. For example, if the nominal LSP's configured name is **bob** and the number of member LSPs is  $N$ , the member LSPs are named **bob-*<configured-suffix>-1***, **bob-*<configured-suffix>-2***, ..., and **bob-*<configured-suffix>-N***.

After a normalization event, the number of member LSPs can change. For example, if the number of member LSPs increases from six to eight, then the ingress routing device keeps the first six LSPs named **bob-*<configured-suffix>-1***, **bob-*<configured-suffix>-2***, ..., and **bob-*<configured-suffix>-6***. The two additional LSPs are named **bob-7** and **bob-8**. The original LSPs might need to be re-optimized if their signaled bandwidth changes.

Similarly, if the number of member LSPs reduces from eight to six, the ingress routing device re-signals the member LSPs in such a way that the remaining active LSPs in the system are named **bob-*<configured-suffix>-1***, **bob-*<configured-suffix>-2***, ..., and **bob-*<configured-suffix>-6***.

In the process of creating new LSPs, an RSVP LSP named **bob-*<configured-suffix>-7*** can be configured.

## Normalization

- [Operational Overview on page 5285](#)
- [Operational Constraints on page 5285](#)
- [Inter-Operation with Autobandwidth on page 5286](#)

### Operational Overview

Normalization is an event that happens periodically. When it happens, a decision is made on the number of member LSPs that should remain active and their respective bandwidths in a container LSP. More specifically, the decision is made on whether new supplementary LSPs are to be created, or any existing LSPs are required to be re-signaled or deleted during the normalization event.

Between two normalization events, a member LSP can undergo several autobandwidth adjustments. A normalization timer, similar to re-optimization timer, is configured. The normalization timer interval should be no less than the adjustment interval or optimization timer.



**NOTE:** Normalization is not triggered based on network events, such as topology changes.

### Operational Constraints

Normalization has the following operational constraints:

- Normalization happens only when none of the member LSPs are undergoing re-optimization or make-before-break. Normalization starts when all the member LSPs complete their ongoing make-before-break. If normalization is pending, new optimization should not be attempted until the normalization is complete.
- After normalization, an ingress routing device first computes a set of bandwidth-feasible paths using constraint-based routing computations. If enough constraint-based routing computed paths are not brought up with an aggregate bandwidth value that exceeds the desired bandwidth, several failover actions are taken.
- After a set of bandwidth-feasible paths are available, the ingress routing device signals those paths while keeping the original set of paths up with the old bandwidth values. The make-before-break is done with shared explicit (SE) sharing style, and when some of the LSPs do not get successfully re-signaled, a bounded number of retries is attempted for a specified duration. Only when all the LSPs are successfully signaled does the ingress router switch from the old instance of the container LSP to the newer instance. If all LSPs could not be successfully signaled, the ingress router keeps those instances of members that are up with higher bandwidth values.

For example, if the bandwidth of an old instance of a member LSP (LSP-1) is 1G, the LSP is split into LSP-1 with bandwidth 2G and LSP-2 with bandwidth 2G. If the signaling of LSP-1 with bandwidth 2G fails, the ingress router keeps LSP-1 with bandwidth 1G and LSP-2 with bandwidth 2G.

When there is a signaling failure, the ingress routing device stays in the error state, where some LSPs have updated bandwidth values only if the aggregate bandwidth has increased. The ingress router makes an attempt to bring up those LSPs that could not be successfully signaled, resulting in minimum traffic loss.

- If an LSP goes down between two normalization events, it can increase the load on other LSPs that are up. In order to prevent overuse of other LSPs, premature normalization can be configured in case of LSP failure. LSPs can go down because of pre-emption or lack of node or link protection. It might not be necessary to bring up the LSPs that are down because the normalization process re-runs the constraint-based routing path computations.

### ***Inter-Operation with Autobandwidth***

Taking as an example, there is one nominal LSP named LSP-1 configured with the following parameters:

- Splitting-bandwidth and maximum-signaling-bandwidth of 1G
- Merging-bandwidth and minimum-signaling-bandwidth of 0.8G
- Autobandwidth

Normalization is performed differently in the following scenarios:

- [Changes in Per-LSP Autobandwidth Adjustments on page 5286](#)
- [Changes in Traffic Growth on page 5288](#)
- [Computed Range and Configured Feasible Ranges on page 5288](#)

### ***Changes in Per-LSP Autobandwidth Adjustments***

[Table 421](#) illustrates how normalization splits and merges member LSPs as autobandwidth adjustments change per-LSP bandwidth with unconditional normalization.

**Table 421: Normalization with Per-LSP Autobandwidth Adjustment Changes**

| Normalization Time | Current State                                                               | Events                                                                                                                                                                                                                     | Adjusted State                           |
|--------------------|-----------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------|
| T0                 | No state.                                                                   | Initialization                                                                                                                                                                                                             | LSP-1 is signaled with bandwidth of 0.8G |
| T1                 | LSP-1 usage increases to 1.5G                                               | <ul style="list-style-type: none"> <li>• Multiple autobandwidth adjustments since T0 is possible.</li> <li>• The ingress router decides to split LSP-1 into two LSPs, and creates LSP-2.</li> </ul>                        | LSP-1 = 0.8G<br>LSP-2 = 0.8G             |
| T2                 | LSP-1 usage increase to 2G<br>LSP-2 usage increases to 0.9G (within limits) | <ul style="list-style-type: none"> <li>• Aggregate bandwidth is 2.9G, which exceeds aggregate splitting maximum of 2G.</li> <li>• The ingress router decides to split LSP-1 into three LSPs, and creates LSP-3.</li> </ul> | LSP-1 = 1G<br>LSP-2 = 1G<br>LSP-3 = 1G   |

Table 421: Normalization with Per-LSP Autobandwidth Adjustment Changes (*continued*)

| Normalization Time | Current State                 | Events                                                                                                                                                                                                        | Adjusted State                                       |
|--------------------|-------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------|
| T3                 | LSP-3 usage increases to 1.5G | <ul style="list-style-type: none"> <li>Aggregate bandwidth is 3.5G with a maximum aggregate splitting of 3G.</li> <li>The ingress router decides to split LSP-1 into four LSPs, and creates LSP-4.</li> </ul> | LSP-1 = 1G<br>LSP-2 = 1G<br>LSP-3 = 1G<br>LSP-4 = 1G |
| T4                 | LSP-2 usage drops to 0.5G     | <ul style="list-style-type: none"> <li>Aggregate bandwidth is 3G.</li> <li>The ingress router decides to merge LSP-1 and removes LSP-4.</li> </ul>                                                            | LSP-1 = 1G<br>LSP-2 = 1G<br>LSP-3 = 1G               |

Because autobandwidth is configured on a per-LSP basis, every time there is an autobandwidth adjustment, the ingress router re-signals each LSP with **Max Avg Bw**.

Another approach to handling the changes in per-LSP autobandwidth adjustments is to not allow individual LSPs to run autobandwidth on the ingress router, but to run autobandwidth in passive (monitor) mode. This way, sampling is done at every statistics interval for member LSPs only, and normalization is performed for the container LSP alone instead of acting on individual LSPs adjustment timer expiry.

As a result, the number of re-signaling attempts and bandwidth fluctuations for a given member LSP is reduced. Only the computed bandwidth-values per-member LSP is used by the ingress router to find an aggregate bandwidth to be used during normalization. Configuring autobandwidth adjustment followed by normalization (adjustments and normalization intervals are comparable) can lead to considerable overhead because of re-signaling.

Taking the same example, and applying the second approach, LSP-1 goes from 0.8G to 1.5G and then back to 0.8G. If the normalization timer is of the same order as the adjustment interval, the ingress router leaves LSP-1 alone with its original 0.8G and only signals LSP-2 with 0.8G. This helps achieve the final result of normalization, thus avoiding the extra signaling attempt on LSP-1 with 1.5G at adjustment timer expiry.

Because member LSPs always use equal bandwidth, any adjustment done on member LSPs is undone. The member LSPs are re-signaled with reduced bandwidth when compared to the reserved capacity in adjustment trigger with normalization trigger. Therefore, avoiding adjustment trigger for member LSPs might be useful assuming that normalization and adjustment intervals are of the same order.



**NOTE:** We recommend that the normalization timer be higher than the autobandwidth adjustment interval and regular optimization duration, as the traffic trends are observed at a longer time scale and normalization is performed one-to-three times per day. An LSP can undergo optimization for the following reasons:

- Normal optimization
- Autobandwidth adjustment
- Normalization

### *Changes in Traffic Growth*

Table 422 illustrates how normalization is performed when traffic grows in large factor.

**Table 422: Normalization with Traffic Growth**

| Normalization Time | Current State              | Events                                                                                                                                                                                                | Adjusted State                           |
|--------------------|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------|
| T0                 | No state                   |                                                                                                                                                                                                       | LSP-1 is signaled with bandwidth of 0.8G |
| T1                 | LSP-1 usage increase to 3G | <ul style="list-style-type: none"> <li>• Aggregate usage exceeds maximum splitting bandwidth</li> <li>• The ingress router decides to split LSP-1, and creates two more supplementary LSPs</li> </ul> | LSP-1 = 1G<br>LSP-2 = 1G<br>LSP-3 = 1G   |

Having fewer LSPs is preferred over signaling four LSPs each with 0.8G bandwidth, unless there is a constraint on the minimum number of LSPs.

### *Computed Range and Configured Feasible Ranges*

When an ingress router is configured with the minimum and maximum number of LSPs, and per LSP splitting-bandwidth and merging-bandwidth values, the bandwidth thresholds are used for splitting and merging. For this, the number of LSPs (N) should satisfy the following constraints:

$$\text{minimum-member-lsps} \leq N \leq \text{maximum-member-lsps}$$

At the time of normalization, based on the aggregate demand X:

$$\lceil X/\text{splitting-bandwidth} \rceil \leq N \leq \lfloor X/\text{merging-bandwidth} \rfloor$$

The above-mentioned constraints provide two ranges for N to work from. If the two ranges for N are overlapping, N will be selected from the overlapping interval (lowest possible N) to keep the number of LSPs small in the network.

Otherwise, if maximum-member-lsps is less than  $\lceil X/\text{splitting-bandwidth} \rceil$ , the ingress router keeps (at maximum) the maximum-member-lsps in the system, and the bandwidth of each LSP is  $\lceil X/\text{maximum-member-lsps} \rceil$  or the maximum-signaling-bandwidth, whichever is less. It is possible that some LSPs might not get signaled successfully.



Similarly, if minimum-member-lsps is greater than  $[X/\text{merging-bandwidth}]$ , the ingress router keeps (at minimum) the minimum-member-lsps in the system, and the bandwidth of each LSP is  $[X/\text{minimum-member-lsps}]$  or the minimum-signaling-bandwidth, whichever is less.

Taking as an example, normalization is performed as following in these cases:

- Case 1
  - minimum-member-lsps = 2
  - maximum-member-lsps = 10
  - aggregate demand = 10G
  - merging-bandwidth = 1G
  - splitting-bandwidth = 2.5G

In this case, the ingress routing device signals four member LSPs each with a bandwidth of 2G.

- Case 2
  - minimum-member-lsps = 5
  - maximum-member-lsps = 10
  - aggregate demand = 10G
  - merging-bandwidth = 2.5G
  - splitting-bandwidth = 10G

In this case, the ingress routing device signals five member LSPs each with a bandwidth of 2G. Here, the static configuration on the number of member LSPs takes precedence.

- Case 3
  - minimum-signaling-bandwidth = 5G
  - maximum-signaling-bandwidth = 40G
  - merging-bandwidth = 10G
  - splitting-bandwidth = 50G

When a container LSP comes up, the nominal LSP is signaled with minimum-signaling-bandwidth. At the time of normalization, the new-aggregate-bandwidth is 100G. To find N and the bandwidth of each LSP, N should satisfy the following constraint:

$$100/50 \leq N \leq 100/10, \text{ which gives } 2 \leq N \leq 10$$

Therefore, N is equal to:

- N = 2, bandwidth =  $\min \{100/2G, 40G\} = 40G$

This option does not satisfy the new aggregate of 100G.

- N = 3, bandwidth =  $\min \{100/3G, 40G\} = 33.3G$

This option makes the aggregate bandwidth equal to 100G.

In this case, the ingress routing device signals three LSPs each with a bandwidth of 33.3G.



**NOTE:** The ingress router does not signal an LSP smaller than the minimum-signaling-bandwidth.

---

### Constraint-Based Routing Path Computation

---

Although there are no changes in the general constraint-based routing path computation, with a container LSP, there is a separate module that oversees the normalization process, schedules constraint-based routing events, and schedules switchover from an old instance to a new instance, when appropriate. An ingress routing device has to handle the constraint-based routing path computation periodically. When normalization occurs, an ingress router has to compute constraint-based routing paths, if the number of LSPs or the bandwidth of the LSPs needs to be changed.

For example, there are K LSPs at the ingress router with bandwidth values X-1, X-2, ..., and X-K. The current aggregate bandwidth value is Y, which is the sum of X-1 plus X-2 plus X-K. If there is a new demand of W, the ingress router first computes how many LSPs are required. If the ingress router only needs N LSPs (LSP-1, LSP-2, ..., and LSP-N) each with bandwidth value B, the task of the constraint-based routing module is to provide a set of bandwidth-feasible LSPs that can accommodate the new aggregate demand which is not less than Y.

The ingress router then tries to see if the constraint-based routing paths can be computed successfully for all N LSPs. If the paths for all the LSPs are found successfully, the constraint-based routing module returns the set to the normalization module.

It is possible that the constraint-based routing computation is not successful for some LSPs. In this case, the ingress routing device takes the following action:

- If the configuration allows for incremental-normalization, implying if the ingress router has enough LSPs whose aggregate exceeds Y, the constraint-based routing module returns that set of paths.
- Whether increment-normalization is configured or not, if constraint-based routing paths could not be computed for a sufficient number of LSPs, the ingress router has to repeat the process of finding a new set of LSPs. Initially, the ingress router starts with the lowest value of N from the feasible region. Every time, the ingress router has to revise the number, it linearly increases it by 1. As a result, per LSP bandwidth becomes less and therefore, there is a greater chance of successful signaling. The process is repeated for all feasible values of N (or some bounded number of times or duration as configured).

The ingress router signals the LSPs after successful computations of the constraint-based routing path computation. It might happen that when the LSPs are signaled, signaling of many LSPs fail. In addition to the constraint-based routing path

computations to be successful, the RSVP signaling should also succeed, such that the new aggregate is not less than the old aggregate bandwidth.

### Sampling

Sampling is important for normalization to function. With sampling configured, an ingress routing device is able to make a statistical estimate of the aggregate traffic demands. Every time the sampling timer fires, the ingress routing device can consider traffic rates on different LSPs and compute an aggregate bandwidth sample. This sampling timer is different from the statistics sampling done periodically by RSVP on all LSPs. The aggregate bandwidth is a sample to be used at the time of normalization. An ingress routing device can save past samples to compute an average (or some other statistical measure) and use it the next time normalization happens.

To remove any outlier samples, a sampling token is configured. In other words, from all the aggregate samples collected during the configured time, the bottom and top outliers are ignored before computing a statistical measure from the remaining samples.

The following two methods of computing an aggregate bandwidth value are supported:

- **Average**—All the aggregate bandwidth samples are considered by the ingress routing device, and then all the outlier samples are removed. The average bandwidth value is computed from the remaining samples to be used during normalization.
- **Max**—All the aggregate bandwidth samples are considered by the ingress routing device, and then all the outlier samples are removed. The maximum bandwidth value is picked from the remaining samples to be used during normalization.

The time duration, the number of past aggregate samples to store, the percentile value to determine, and the ignore outliers are user-configurable parameters.

### Support for NSR, IPG-FA, and Static Routes

Starting with Junos OS Release 15.1, container label-switched paths (LSPs) provide support for nonstop active routing (NSR), IGP forwarding adjacency (FA), and static routes to address the requirements of wider business cases.

- [NSR Support on page 5291](#)
- [IPG-FA Support on page 5293](#)
- [Static Route Support on page 5294](#)

#### **NSR Support**

A container LSP has the characteristics of ECMP and RSVP traffic engineering. Because a container LSP consists of several member LSPs between an ingress and an egress router, with each member LSP taking a different path to the same destination, the ingress router is configured with all the parameters necessary to compute an RSVP ECMP LSP. These parameters along with the forwarding state information have to be synchronized between the master and backup Routing Engines to enable the support for nonstop active routing (NSR) for container LSPs. While some of the forwarding state information on the backup Routing Engine is locally built based on the configuration, most of it is built

based on periodic updates from the master Routing Engine. The container LSPs are created dynamically using the replicated states on the backup Routing Engine.

By default, normalization occurs once in every 6 hours and during this time, a number of autobandwidth adjustments happen over each member LSP. A member LSP is resized according to the traffic it carries and the configured autobandwidth configuration parameters. The aggregate demand on a container LSP is tracked by summing up the bandwidth across all the member LSPs.

For RSVP point-to-point LSPs, a Routing Engine switchover can be under any one of the following:

- **Steady state**

In the steady state, the LSP state is up and forwards traffic; however, no other event, such as the make-before-break (MBB), occurs on the LSP. At this stage, the RPD runs on both the Routing Engines, and the switchover event toggles between the master and backup Routing Engine. The backup Routing Engine has the LSP information replicated already. After the switchover, the new master uses the information of the replicated structure to construct the container LSP and en-queues the path (ERO) of LSP in the retrace mode. RSVP signals and checks if the path mentioned in the ERO is reachable. If the RSVP checks fail, then the LSP is restarted. If the RSVP checks succeed, the LSP state remains up.

- **Action leading to make-before-break (MBB)**

A container LSP can be optimized with updated bandwidth, and this change is done in a MBB fashion. During an MBB process, there are two path instances for a given LSP, and the LSP switches from one instance to another. For every Routing Engine switchover, the path is checked to find out where in the MBB process the path is. If the path is in the middle of the MBB process, with the main instance being down and the re-optimized path being up, then MBB can switch over to the new instance. The **show mpls lsp extensive** command output, in this case, is as follows:

```
13 Dec 3 01:33:38.941 Make-before-break: Switched to new instance
12 Dec 3 01:33:37.943 Record Route: 10.1.1.1
11 Dec 3 01:33:37.942 Up
10 Dec 3 01:33:37.942 Automatic Autobw adjustment succeeded: BW changes
from 100 bps to 281669 bps
9 Dec 3 01:33:37.932 Originate make-before-break call
8 Dec 3 01:33:37.931 CSPF: computation result accepted 10.1.1.1
7 Dec 3 01:28:44.228 CSPF: ERO retrace was successful 10.1.1.1
6 Dec 3 01:19:39.931 10.1.1.2 Down: mbb/reopt
5 Dec 3 01:18:29.286 Up: mbb/reopt
4 Dec 3 01:14:47.119 10.1.1.2 Down: mbb/reopt
3 Dec 3 01:13:29.285 Up: mbb/reopt
2 Dec 3 01:10:59.755 Selected as active path: selected by master RE
```

A similar behavior is retained for member LSPs during bandwidth optimization.

A Routing Engine switchover under the steady state (when normalization is not in progress), keeps the container LSPs up and running without any traffic loss. Events, such as an MBB due to autobandwidth adjustments, link status being down, or double failure, in the steady state are similar to a normal RSVP point-to-point LSP.

If the container LSP is in the process of normalization, and the normalization event is triggered either manually or periodically, it goes through the computation and execution phase. In either of the cases, zero percent traffic loss is not guaranteed.

- Normalization in the computation phase

During the computation phase, the master Routing Engine calculates the targeted member LSP count and bandwidth with which each member LSP should be re-signaled. The backup Routing Engine has limited information about the container LSP, such as the LSP name, LSP ID, current bandwidth of its member LSP, member LSP count, and the normalization retry count. If the switchover happens during the computation phase, then the backup Routing Engine is not aware of the targeted member LSP count and the bandwidth to be signaled. Since traffic statistics are not copied to the backup Routing Engine, it cannot compute the targeted member count and bandwidth. In this case, the new master Routing Engine uses the old data stored in the targeted member LSP count and the targeted bandwidth to signal the LSPs.

- Normalization in the execution phase

During the execution phase, RSVP of the master Routing Engine tries to signal the LSPs with the newly calculated bandwidth. If the switchover occurs during the signaling of LSPs with greater bandwidth or during LSP splitting or merging, then the new master Routing Engine uses the information of the targeted member count and bandwidth value to be signaled with, to bring up the LSPs.

### ***IPG-FA Support***

A forwarding adjacency (FA) is a traffic engineering label-switched path (LSP) that is configured between two nodes and used by an interior gateway protocol (IGP) to forward traffic. By default, an IGP does not consider MPLS traffic-engineering tunnels between sites, for traffic forwarding. Forwarding adjacency treats a traffic engineering LSP tunnel as a link in an IGP topology, thus allowing the nodes in the network also to forward the IP traffic to reach the destination over this FA LSP. A forwarding adjacency can be created between routing devices regardless of their location in the network.

To advertise a container LSP as an IGP-FA, the LSP name needs to be configured either under IS-IS or OSPF. For example:

```
IS-IS [edit]
 protocols {
 isis {
 label-switched-path container-lsp-name;
 }
 }

OSPF [edit]
 protocols {
 ospf {
 area 0.0.0.0 {
 label-switched-path container-lsp-name;
 }
 }
 }
```



**NOTE:** The IGP-FA is applied to both container LSPs and regular point-to-point LSPs. If a container LSP and a point-to-point LSP share the same name, the point-to-point LSP is given preference for FA.

### Static Route Support

Static routes often include only one or very few paths to a destination and generally do not change. These routes are used for stitching services when policies and other protocols are not configured.

To advertise a container LSP as a static route, the LSP name needs to be configured under the static route configuration. For example:

```
Static Route [edit]
 routing-options {
 static {
 route destination {
 lsp-next-hop container-lsp-name;
 }
 }
 }
```



**NOTE:** The static route support is applied to both container LSPs and regular point-to-point LSPs. If a container LSP and a point-to-point LSP share the same name, the point-to-point LSP is given preference for static routing.

## Configuration Statements Supported for Container LSPs

Table 423 lists the MPLS LSP configuration statements that apply to RSVP LSP and a container LSP (nominal and supplementary).

The configuration support is defined using the following terms:

- Yes—The configuration statement is supported for this type of LSP.
- No—The configuration statement is not supported for this type of LSP.
- N/A—The configuration statement is not applicable for this type of LSP.

**Table 423: Applicability of RSVP LSPs Configuration to a Container LSP**

| Configuration Statement             | RSVP LSP (Ingress) | Member LSP (Ingress) |
|-------------------------------------|--------------------|----------------------|
| adaptive<br>(Default: non-adaptive) | Yes                | Yes                  |
| admin-down                          | Yes                | Yes                  |
| admin-group                         | Yes                | Yes                  |

Table 423: Applicability of RSVP LSPs Configuration to a Container LSP (*continued*)

| Configuration Statement                                      | RSVP LSP (Ingress) | Member LSP (Ingress) |
|--------------------------------------------------------------|--------------------|----------------------|
| admin-groups-except                                          | Yes                | Yes                  |
| apply-groups                                                 | Yes                | Yes                  |
| apply-groups-except                                          | Yes                | Yes                  |
| associate-backup-pe-groups                                   | Yes                | No                   |
| associate-lsp<br>(No bidirectional support)                  | Yes                | No                   |
| auto-bandwidth                                               | Yes                | Yes                  |
| backup                                                       | Yes                | No                   |
| bandwidth                                                    | Yes                | Yes                  |
| class-of-service                                             | Yes                | Yes                  |
| corouted-bidirectional<br>(No bidirectional support)         | Yes                | No                   |
| corouted-bidirectional-passive<br>(No bidirectional support) | Yes                | No                   |
| description                                                  | Yes                | Yes                  |
| disable                                                      | Yes                | Yes                  |
| egress-protection                                            | Yes                | No                   |
| exclude-srlg                                                 | Yes                | Yes                  |
| fast-reroute<br>(Same fast reroute for all member LSPs)      | Yes                | Yes                  |
| from                                                         | Yes                | Yes                  |
| hop-limit                                                    | Yes                | Yes                  |
| install                                                      | Yes                | Yes                  |

Table 423: Applicability of RSVP LSPs Configuration to a Container LSP (*continued*)

| Configuration Statement                                                      | RSVP LSP (Ingress) | Member LSP (Ingress) |
|------------------------------------------------------------------------------|--------------------|----------------------|
| inter-domain<br>(Same termination router)                                    | Yes                | Yes                  |
| secondary<br>(All LSPs are primary)                                          | Yes                | No                   |
| ldp-tunneling<br>(All LSPs do tunneling)                                     | Yes                | Yes                  |
| least-fill                                                                   | Yes                | Yes                  |
| link-protection<br>(All LSPs share same link protection mechanism)           | Yes                | Yes                  |
| lsp-attributes                                                               | Yes                | Yes                  |
| lsp-external-controller                                                      | Yes                | No                   |
| metric<br>(All LSPs are same)                                                | Yes                | Yes                  |
| most-fill                                                                    | Yes                | Yes                  |
| no-cspf<br>(LSPs use IGP)                                                    | Yes                | Yes                  |
| no-decrement-ttl<br>(All LSPs share same TTL behavior)                       | Yes                | Yes                  |
| no-install-to-address                                                        | Yes                | Yes                  |
| no-record                                                                    | Yes                | Yes                  |
| node-link-protection<br>(All LSPs share same node-link protection mechanism) | Yes                | Yes                  |
| oam                                                                          | Yes                | Yes                  |
| optimize-hold-dead-delay<br>(All LSPs have same value)                       | Yes                | Yes                  |



Table 423: Applicability of RSVP LSPs Configuration to a Container LSP (*continued*)

| Configuration Statement                                 | RSVP LSP<br>(Ingress) | Member LSP<br>(Ingress) |
|---------------------------------------------------------|-----------------------|-------------------------|
| optimize-switchover-delay<br>(All LSPs have same value) | Yes                   | Yes                     |
| optimize-timer<br>(All LSPs have same value)            | Yes                   | Yes                     |
| p2mp                                                    | Yes                   | N/A                     |
| policing<br>(Variable traffic)                          | Yes                   | No                      |
| preference                                              | Yes                   | Yes                     |
| primary<br>(All paths are primary)                      | Yes                   | No                      |
| random                                                  | Yes                   | Yes                     |
| record                                                  | Yes                   | Yes                     |
| retry-limit<br>(Applicable to members)                  | Yes                   | Yes                     |
| retry-timer<br>(Applicable to members)                  | Yes                   | Yes                     |
| revert-timer<br>(No secondary LSP)                      | Yes                   | No                      |
| secondary<br>(All LSPs are primary)                     | Yes                   | No                      |
| soft-preemption                                         | Yes                   | Yes                     |
| standby<br>(All LSPs are standby)                       | Yes                   | No                      |
| template                                                | Yes                   | No                      |
| to                                                      | Yes                   | Yes                     |

Table 423: Applicability of RSVP LSPs Configuration to a Container LSP (*continued*)

| Configuration Statement | RSVP LSP (Ingress) | Member LSP (Ingress) |
|-------------------------|--------------------|----------------------|
| traceoptions            | Yes                | Yes                  |
| ultimate-hop-popping    | Yes                | Yes                  |

## Impact of Configuring Container LSPs on Network Performance

A container LSP is a container LSP that allows multiple member LSPs to co-exist and be managed as a bundle. The member LSPs are similar to independent point-to-point RSVP LSPs. As a result, resource consumption is similar to the sum of resources consumed by each point-to-point RSVP LSP. However, provisioning a container LSP is more efficient, as under-utilized member LSPs are dynamically removed, thus saving memory and CPU resources.

The container LSP features are dependent on the presence of a functional base MPLS RSVP implementation. As a result, a container LSP does not introduce any security considerations beyond the existing considerations for the base MPLS RSVP functionality. The categories of possible attacks and countermeasures are as follows:

- Interaction with processes and router configuration

No new communication mechanisms with external hosts are required for a container LSP. Data arrives at the RSVP module through local software processes and router configuration, other than RSVP neighbor adjacency. Junos OS provides security controls on access to the router and router configuration.

- Communication with external RSVP neighbors

RSVP signaled MPLS LSPs depend on the services of RSVP and IGP to communicate RSVP messages among neighboring routers across the network. Because the RSVP sessions involve communication outside of the local router, they are subject to many forms of attack, such as spoofing of peers, injection of falsified RSVP messages and route updates, and attacks on the underlying TCP/UDP transport for sessions. Junos OS provides countermeasures for such attack vectors.

- Resource limits and denial of service

Junos OS provides several mechanisms through policers and filters to protect against denial-of-service attacks based on injecting higher than the expected traffic demands. At the MPLS LSP level, Junos OS allows operators to configure limits on the LSP bandwidth and the number of LSPs. However, like point-to-point RSVP LSPs, container LSPs do not enforce limits on the volume of traffic forwarded over these LSPs.

## Supported and Unsupported Features

Junos OS supports the following container LSP features:

- Equal-bandwidth-based LSP splitting mechanism
- Aggregate-bandwidth-based LSP splitting and merging in a make-before-break way

- LSP-number-based naming mechanism for dynamically created member LSPs
- Periodic sampling mechanisms to estimate aggregate bandwidth
- Interoperability with auto-bandwidth feature
- ECMP using the dynamically created LSPs
- LDP-tunneling on the dynamically created LSP
- Configuring container LSP using IGP shortcuts
- Aggregated Ethernet links
- Logical systems

Junos OS does **not** support the following container LSP functionality:

- Node and link disjoint paths for different LSPs between an ingress and an egress routing device
- Bandwidth allocation policy different from equal bandwidth policy at the normalization event
- Constraint-based routing path computation to find equal IGP cost paths for different LSPs
- RSVP objects, such as **MLSP\_TUNNEL Sender Template**, and **MLSP\_TUNNEL Filter Specification** defined in [KOMPELLA-MLSP]
- Change in topology as a trigger for LSP splitting and merging
- Change in topology and link failure as a trigger for normalization, unless member LSPs go down
- Egress protection on container LSP
- Container LSP as a backup LSP for IGP interface
- Container LSP configured as IGP interface as forwarding address
- Container LSP as provider tunnel for multicast VPNs
- Dynamic LSPs for normalization
- CCC using container LSP
- Secondary paths for container LSP
- Bidirectional container LSP
- Policing
- Static routes using container LSPs as next hops on a best-effort basis
- External path computing entity, such as PCE
- Graceful Routing Engine switchover
- Nonstop active routing
- Unified ISSU

- Multichassis
- IPv6

**Related  
Documentation**

- *Example: Configuring Dynamic Bandwidth Management Using Container LSP*
- [Maximize Bandwidth Utilization with Juniper Networks TE++](#)

# Configuration Statements and Command Statements for RSVP

- [Configuration Statements for RSVP on page 5301](#)
- [Monitoring Commands for RSVP on page 5341](#)

## Configuration Statements for RSVP

---

- [\[edit protocols rsvp\] Hierarchy Level on page 5302](#)
- [admin-group on page 5304](#)
- [authentication-key \(Protocols RSVP\) on page 5305](#)
- [aggregate \(Protocols RSVP\) on page 5306](#)
- [bandwidth \(Protocols RSVP\) on page 5307](#)
- [bypass \(Signaled LSP\) on page 5308](#)
- [class-of-service \(Protocols RSVP\) on page 5309](#)
- [container-label-switched-path on page 5310](#)
- [disable \(Protocols RSVP\) on page 5311](#)
- [fast-reroute \(Protocols RSVP\) on page 5312](#)
- [graceful-deletion-timeout on page 5312](#)
- [graceful-restart \(Enabling Globally\) on page 5313](#)
- [hello-acknowledgements on page 5314](#)
- [hello-interval \(Protocols RSVP\) on page 5315](#)
- [helper-disable \(Multiple Protocols\) on page 5316](#)
- [hop-limit on page 5317](#)
- [interface \(Protocols RSVP\) on page 5318](#)
- [keep-multiplier on page 5319](#)
- [link-protection \(RSVP\) on page 5320](#)
- [load-balance \(Protocols RSVP\) on page 5321](#)
- [max-bypasses on page 5322](#)
- [maximum-helper-recovery-time on page 5323](#)

- [maximum-helper-restart-time \(RSVP\)](#) on page 5324
- [no-cspf \(Protocols RSVP\)](#) on page 5325
- [no-interface-hello](#) on page 5325
- [no-local-reversion](#) on page 5326
- [no-node-id-subobject](#) on page 5327
- [no-p2mp-sublsp](#) on page 5327
- [node-hello](#) on page 5328
- [optimize-timer \(Protocols RSVP\)](#) on page 5328
- [path \(Protocols RSVP\)](#) on page 5329
- [preemption](#) on page 5330
- [priority \(Protocols RSVP\)](#) on page 5331
- [refresh-time](#) on page 5332
- [reliable](#) on page 5332
- [setup-protection](#) on page 5333
- [subscription](#) on page 5334
- [soft-preemption \(Protocols RSVP\)](#) on page 5335
- [splitting-merging](#) on page 5336
- [traceoptions \(Protocols RSVP\)](#) on page 5338
- [tunnel-services \(RSVP\)](#) on page 5340
- [update-threshold](#) on page 5340

## **[edit protocols rsvp] Hierarchy Level**

This topic lists the supported configuration statements at the **[edit protocols rsvp]** hierarchy level on the QFX Series. For more information about these statements, see the *Junos OS MPLS Applications Library for Routing Devices*.

```
protocols {
 rsvp {
 disable;
 fast-reroute
 graceful-deletion-timeout seconds;
 graceful-restart {
 disable;
 helper-disable;
 maximum-helper-recovery-time seconds;
 maximum-helper-restart-time seconds;
 }
 hello-acknowledgements;
 interface interface-name {
 (aggregate | no-aggregate);
 authentication-key key;
 bandwidth bps;
 disable;
 hello-interval seconds;
 link-protection {
```

```

disable;
admin-group {
 exclude [group-names];
 include-all [group-names];
 include-any [group-names];
}
bandwidth bps;
bypass bypass-name {
 bandwidth bps {
 ct0 bps;
 ct1 bps;
 ct2 bps;
 ct3 bps;
 }
 description text;
 class-of-service cos-value;
 hop-limit number;
 no-cspf;
 path address <strict | loose>;
 priority setup-priority reservation-priority;
 to address;
}
class-of-service cos-value;
hop-limit number;
max-bypasses number;
no-cspf;
no-node-protection;
optimize-timer seconds;
path address <strict | loose>;
priority setup-priority reservation-priority;
subscription percentage;
}
(reliable | no-reliable);
subscription {
 percentage;
 ct0 percentage;
 ct1 percentage;
 ct2 percentage;
 ct3 percentage;
}
update-threshold percentage;
}
keep-multiplier number;
load-balance bandwidth;
no-interface-hello;
no-local-reversion;
no-node-id-subobject;
no-p2mp-sublsp;
node-hello
preemption {
 (aggressive | disabled | normal);
 soft-preemption cleanup-timer seconds;
}
refresh-time seconds;
setup-protection;
traceoptions {

```

```
file filename <files number> <size maximum-file-size> <world-readable |
no-world-readable>;
flag flag <flag-modifier> <disable>;
}
tunnel-services {
 devices device-names;
}
}
}
```

**Related Documentation**

- *Junos OS MPLS Applications Library for Routing Devices*

---

## admin-group

**Syntax**

```
admin-group {
 exclude [group-names];
 include-all [group-names];
 include-any [group-names];
}
```

**Hierarchy Level**

[edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection],  
[edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection bypass *bypass-name*],  
[edit protocols rsvp interface *interface-name* link-protection],  
[edit protocols rsvp interface *interface-name* link-protection bypass *bypass-name*]

**Release Information** Statement introduced in Junos OS Release 9.2.

**Description** Enable you to configure administrative groups for bypass label-switched paths (LSPs). You can configure administrative groups either globally for all bypass LSPs traversing an interface or for just a specific bypass LSP.

**Options**

**exclude *group-names***—Specify the administrative groups to exclude for a bypass LSP.

**include-all *group-names***—Specify the administrative groups whose links the bypass LSP must traverse.

**include-any *group-names***—Specify the administrative groups whose links the bypass LSP can traverse.

**Required Privilege Level**

routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Configuring Link Protection on Interfaces Used by LSPs*



## authentication-key (Protocols RSVP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | authentication-key <i>key</i> ;                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> protocols rsvp peer-interface <i>peer-interface-name</i> ],<br>[edit protocols rsvp interface <i>interface-name</i> ],<br>[edit protocols rsvp peer-interface <i>peer-interface-name</i> ]                                                             |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | <p>Authentication key (password). Neighboring routers use the password to verify the authenticity of packets sent from this interface or peer interface.</p> <p>RSVP uses HMAC-MD5 authentication, which is defined in RFC 2104, <i>HMAC: Keyed-Hashing for Message Authentication</i>.</p> <p>All routers that are connected to the same IP subnet must use the same authentication scheme and password.</p> |
| <b>Options</b>                  | <b>key</b> —Authentication password. It can be 1 through 16 contiguous digits or letters. Separate decimal digits with periods. Separate hexadecimal digits with periods and precede the string with 0x. If you include spaces in the password, enclose the entire password in quotation marks (" ").                                                                                                         |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring RSVP Interfaces</i></li> </ul>                                                                                                                                                                                                                                                                                                                        |

## aggregate (Protocols RSVP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | (aggregate   no-aggregate);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> protocols rsvp peer-interface <i>peer-interface-name</i> ],<br>[edit protocols rsvp interface <i>interface-name</i> ],<br>[edit protocols rsvp peer-interface <i>peer-interface-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | <p>Control the use of RSVP aggregate messages on an interface or peer interface:</p> <ul style="list-style-type: none"><li>• <b>aggregate</b>—Use RSVP aggregate messages.</li><li>• <b>no-aggregate</b>—Do not use RSVP aggregate messages.</li></ul> <p>Aggregate messages can pack multiple RSVP messages into a single transmission, thereby reducing network overhead and enhancing efficiency. The number of supportable sessions and processing overhead are significantly improved when aggregation is enabled.</p> <p>Not all routers connected to a subnet need to support aggregation simultaneously. Each RSVP router negotiates its intention to use aggregate messages on a per-neighbor basis. Only when both routers agree are aggregate messages sent.</p> <p>To have refresh reduction and reliable delivery, you must include the <b>aggregate</b> and <b>reliable</b> statements.</p> |
| <b>Default</b>                  | Aggregation is disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring RSVP Interfaces</i></li><li>• <a href="#">reliable on page 5332</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## bandwidth (Protocols RSVP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>bandwidth <i>bps</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i>],<br/> [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection],<br/> [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>],<br/> [edit protocols rsvp interface <i>interface-name</i>],<br/> [edit protocols rsvp interface <i>interface-name</i> link-protection],<br/> [edit protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | <p>For certain logical interfaces (such as Asynchronous Transfer Mode [ATM], Permanent Virtual Circuit [PVC], or Frame Relay), you cannot determine the correct bandwidth from the hardware. This statement enables you to specify the actual available bandwidth.</p> <p>This statement also enables you to specify the bandwidth for a bypass label switched path (LSP). If you have configured multiple bypasses, this statement is mandatory and is applied to all of the bypass LSPs.</p>                                                                                                                                |
| <b>Default</b>                  | The hardware raw bandwidth is used.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <p><b><i>bps</i></b>—Bandwidth in bits per second. You can specify this as an integer value. If you do so, count your zeros carefully, or you can use the abbreviations <b>k</b> (for a thousand), <b>m</b> (for a million), or <b>g</b> (for a billion [also called a thousand million]).</p> <p><b>Range:</b> Any positive integer</p> <p><b>Default:</b> 0 (no bandwidth is reserved)</p>                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Link Protection on Interfaces Used by LSPs</i></li> <li>• <i>Configuring Link Protection on Interfaces Used by LSPs</i></li> <li>• <i>Configuring Link Protection on Interfaces Used by LSPs</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                           |

## bypass (Signaled LSP)

---

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <pre>bypass <i>bypass-name</i> {<br/>    <b>bandwidth</b> <i>bps</i>;<br/>    <b>description</b> <i>text</i>;<br/>    <b>hop-limit</b> <i>number</i>;<br/>    <b>no-cspf</b>;<br/>    <b>path</b> <i>address</i> &lt;strict   loose&gt;;<br/>    <b>priority</b> <i>setup-priority reservation-priority</i>;<br/>    <b>to address</b>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Hierarchy Level          | [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection],<br>[edit protocols rsvp interface <i>interface-name</i> link-protection]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Release Information      | Statement introduced before Junos OS Release 7.4.<br>The <b>description</b> option was added in Junos OS Release 10.4.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Description              | <p>Enables you to configure specific bandwidth and path constraints for a bypass LSP. It is possible to individually configure multiple bypass LSPs. If you do not configure the bypass LSPs individually, they all share the same path and bandwidth constraints.</p> <p>If you specify the <b>bandwidth</b>, <b>hop-limit</b>, and <b>path</b> statements for the bypass LSP, these values take precedence over the values configured at the [edit protocols rsvp interface <i>interface-name</i> link-protection] hierarchy level. The other attributes (<b>subscription</b>, <b>no-node-protection</b>, and <b>optimize-timer</b>) are inherited from the general constraints.</p>                                                                                                                                                                                                                                                                                                           |
| Options                  | <p><b>bypass-name</b>—(Required) Specify a name for the bypass LSP. The name can be up to 64 characters.</p> <p><b>description</b>—Provides a textual description of the bypass LSP. Enclose any descriptive text that includes spaces in quotation marks (" "). Any descriptive text you include is displayed in the output of the <b>show mpls lsp bypass detail</b> command and has no effect on the operation of the bypass LSP. The description text can be no more than 80 characters in length.</p> <p><b>to address</b>—(Required) Specify the address for the interface of the immediate next-hop node (for link protection) or the next-next-hop node (for node-link protection). The address specified determines whether this is a link protection bypass or a node-link protection bypass. On multiaccess networks (for example, a LAN), this address is also used to specify which next-hop node is being protected.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Related Documentation    | <ul style="list-style-type: none"><li>Configuring Link Protection on Interfaces Used by LSPs</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## class-of-service (Protocols RSVP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>class-of-service <i>cos-value</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>],</p> <p>[edit protocols rsvp interface <i>interface-name</i> link-protection],</p> <p>[edit protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>]</p> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | <p>Class-of-service (CoS) value given to all packets in the bypass LSP. You can specify a single CoS value for all the bypass LSPs traversing an interface. You can also configure CoS values for specific bypass LSPs traversing an interface.</p> <p>The CoS value might affect the scheduling or queuing algorithm of traffic traveling along an LSP.</p>                                                                                                     |
| <b>Options</b>                  | <p><b><i>cos-value</i></b>—CoS value. A higher value typically corresponds to a higher level of service.</p> <p><b>Range:</b> 0 through 7</p> <p><b>Default:</b> If you do not specify a CoS value, the IP precedence bits from the packet's IP header are used as the packet's CoS value.</p>                                                                                                                                                                   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Configuring Link Protection on Interfaces Used by LSPs</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                  |

## container-label-switched-path

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>container-label-switched-path <i>lsp-name</i> {<br/>    disable;<br/>    description <i>description</i>;<br/>    label-switched-path-template;<br/>    splitting-merging;<br/>    suffix <i>string</i>;<br/>    to <i>ip-address</i>;<br/>}</pre>                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit protocols mpls]                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.2.<br>Statement introduced for QFX Switches in Junos OS Release 15.1X53-D30.                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Configure a multi-label-switched path (LSP) tunnel between the ingress and the egress routers. The container LSP consists of several member LSPs to the same destination.                                                                                                                                                                                                        |
| <b>Options</b>                  | <p><b>disable</b>—Disable MPLS container-label-switched path.</p> <p><b>description <i>description</i></b>—Text describing the container LSP.</p> <p><b>suffix <i>string</i></b>—Suffix to generate names of member LSPs of the container LSP.</p> <p><b>to <i>ip-address</i></b>—IP address of the egress router.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                   |

## disable (Protocols RSVP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | disable;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols rsvp],<br>[edit logical-systems <i>logical-system-name</i> protocols rsvp graceful-restart],<br>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection],<br>[edit logical-systems <i>logical-system-name</i> protocols rsvp peer-interface <i>peer-interface-name</i> ],<br>[edit protocols rsvp],<br>[edit protocols rsvp graceful-restart],<br>[edit protocols rsvp interface <i>interface-name</i> ],<br>[edit protocols rsvp interface <i>interface-name</i> link-protection],<br>[edit protocols rsvp peer-interface <i>peer-interface-name</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | Explicitly disable RSVP or RSVP graceful restart. Explicitly disable link protection on the specified interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Default</b>                  | RSVP is enabled on interfaces and peer interfaces configured with the RSVP <b>interface</b> statement. RSVP graceful restart is enabled on the router. Link protection is disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Minimum RSVP Configuration</i></li> <li>• <i>Configuring RSVP Graceful Restart</i></li> <li>• <i>Configuring Link Protection on Interfaces Used by LSPs</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## fast-reroute (Protocols RSVP)

---

|                                 |                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>fast-reroute optimize-timer <i>seconds</i>;</code>                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols rsvp],<br>[edit protocols rsvp]                                                                                                         |
| <b>Release Information</b>      | Statement added in Junos OS Release 7.5.<br>Statement introduced in Junos OS Release 14.1 for the QFX Series.                                                                                      |
| <b>Description</b>              | Configure the optimize timer for fast reroute. The optimize timer triggers a periodic optimization process that recomputes the fast reroute detour LSPs to use network resources more efficiently. |
| <b>Options</b>                  | <b><i>seconds</i></b> —Specify the number of seconds between fast reroute detour LSP optimizations.<br><b>Range:</b> 0 through 65,535 seconds<br><b>Default:</b> 0 (disabled)                      |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the Optimization Interval for Fast Reroute Paths</i></li></ul>                                                                              |


## graceful-deletion-timeout

---

|                                 |                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>graceful-deletion-timeout <i>seconds</i>;</code>                                                                                             |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols rsvp],<br>[edit protocols rsvp]                                                         |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                          |
| <b>Description</b>              | Specify the time, in seconds, before completing graceful deletion of signaling.                                                                    |
| <b>Options</b>                  | <b><i>seconds</i></b> —Time before completing graceful deletion of signaling.<br><b>Range:</b> 1 through 300 seconds<br><b>Default:</b> 30 seconds |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Gracefully Tearing Down GMPLS LSPs</i></li></ul>                                                        |



## graceful-restart (Enabling Globally)

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |                                                                                                                                                                                                                                                                                                      |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <pre> graceful-restart {   disable;   helper-disable;   maximum-helper-recovery-time <i>seconds</i>;   maximum-helper-restart-time <i>seconds</i>;   notify-duration <i>seconds</i>;   recovery-time <i>seconds</i>;   restart-duration <i>seconds</i>;   stale-routes-time <i>seconds</i>; } </pre> |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | [edit logical-systems <i>logical-system-name</i> routing-options],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],<br>[edit routing-options],<br>[edit routing-instances <i>routing-instance-name</i> routing-options]          |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                           |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Configure graceful restart globally to enable the feature. You cannot enable graceful restart for specific protocols unless graceful restart is also enabled globally. You can, optionally, modify the global settings at the individual protocol level.                                             |
| <div>  <b>NOTE:</b> <ul style="list-style-type: none"> <li>For VPNs, the <code>graceful-restart</code> statement allows a router whose VPN control plane is undergoing a restart to continue to forward traffic while recovering its state from neighboring routers.</li> <li>For BGP, if you configure graceful restart after a BGP session has been established, the BGP session restarts and the peers negotiate graceful restart capabilities.</li> <li>LDP sessions flap when <code>graceful-restart</code> configurations change.</li> </ul> </div> |                                                                                                                                                                                                                                                                                                      |
| <b>Default</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Graceful restart is disabled by default.                                                                                                                                                                                                                                                             |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | The remaining statements are explained separately.                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                  |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <ul style="list-style-type: none"> <li><a href="#">Enabling Graceful Restart</a></li> <li><a href="#">Configuring Routing Protocols Graceful Restart on page 2580</a></li> </ul>                                                                                                                     |

- *Configuring Graceful Restart for MPLS-Related Protocols*
- *Configuring VPN Graceful Restart*
- *Configuring Logical System Graceful Restart*
- *Graceful Restart Configuration Statements*
- *Configuring Graceful Restart for QFabric Systems*

---

## hello-acknowledgements

---

|                                 |                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | hello-acknowledgements;                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-systems-name</i> protocols rsvp],<br>[edit protocols rsvp]                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.2.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                           |
| <b>Description</b>              | Enable hello messages from nonsession neighbors to be acknowledged with a hello acknowledgment message. Once hello acknowledgments are enabled, the router continues to acknowledge hello messages from any nonsession RSVP neighbors unless the interface itself goes down or the configuration is changed by an administrator. |
| <b>Default</b>                  | Hello acknowledgments are disabled.                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Hello Acknowledgments for Nonsession RSVP Neighbors</i></li></ul>                                                                                                                                                                                                         |

## hello-interval (Protocols RSVP)

|                                 |                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>hello-interval <i>seconds</i>;</code>                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> protocols rsvp peer-interface <i>peer-interface-name</i> ],<br>[edit protocols rsvp interface <i>interface-name</i> ],<br>[edit protocols rsvp peer-interface <i>peer-interface-name</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                         |
| <b>Description</b>              | Enable the sending of hello packets on the interface.                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <b><i>seconds</i></b> —Length of time between hello packets. A value of 0 disables the sending of hello packets on the interface.<br><b>Range:</b> 1 through 60 seconds<br><b>Default:</b> 9 seconds                                                                                                                                              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring RSVP Interfaces</i></li></ul>                                                                                                                                                                                                                                                              |

## helper-disable (Multiple Protocols)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | helper-disable;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols (isis   ldp   ospf   ospf3   rsvp) <a href="#">graceful-restart</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ldp   ospf   ospf3) <a href="#">graceful-restart</a> ],<br>[edit protocols (isis   ldp   ospf   ospf3   rsvp) <a href="#">graceful-restart</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols (ldp   ospf   ospf3) <a href="#">graceful-restart</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Disable helper mode for graceful restart. When helper mode is disabled, a router or switch cannot help a neighboring router that is attempting to restart.                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Default</b>                  | Helper mode is enabled by default for these supported protocols: IS-IS, LDP, OSPF/OSPFv3, and RSVP.                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Routing Protocols Graceful Restart on page 2580</a></li><li>• <a href="#">Configuring Graceful Restart for MPLS-Related Protocols</a></li></ul>                                                                                                                                                                                                                                                                                                             |

## hop-limit

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>hop-limit <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols mpls],<br/> [edit logical-systems <i>logical-system-name</i> protocols mpls <a href="#">label-switched-path</a> <i>lsp-name</i>],<br/> [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> <a href="#">fast-reroute</a>],<br/> [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> (<a href="#">primary</a>   <a href="#">secondary</a>) <i>path-name</i>],<br/> [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection],<br/> [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>],<br/> [edit protocols mpls],<br/> [edit protocols mpls <a href="#">label-switched-path</a> <i>lsp-name</i>],<br/> [edit protocols mpls label-switched-path <i>lsp-name</i> <a href="#">fast-reroute</a>],<br/> [edit protocols mpls label-switched-path <i>lsp-name</i> (<a href="#">primary</a>   <a href="#">secondary</a>) <i>path-name</i>],<br/> [edit protocols rsvp interface <i>interface-name</i> link-protection],<br/> [edit protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | <p>Specify the maximum number of routers that an LSP can traverse. This limit can be applied to any of the following:</p> <ul style="list-style-type: none"> <li>• LSPs—The configured hop limit includes the ingress and egress routers. You can specify a hop limit for an LSP and for both primary and secondary paths.</li> <li>• Fast reroute detour—Specify the number of additional routers a fast reroute detour can traverse relative to the protected LSP. For example, if an LSP traverses 4 routers, any detour for the LSP can be no more than 10 router hops, including the ingress and egress routers.</li> <li>• Link protection bypass—Specify the maximum number of routers that a link protection bypass can traverse.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <p><b><i>number</i></b>—Maximum number of hops.</p> <p><b>Range:</b> 2 through 255 (for an LSP or for a link protection bypass); 0 through 255 (for fast reroute)</p> <p><b>Default:</b> 255 (for an LSP or for a link protection bypass); 6 (for fast reroute)</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Fast Reroute</i></li> <li>• <i>Limiting the Number of Hops in LSPs</i></li> <li>• <i>Configuring Link Protection on Interfaces Used by LSPs</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## interface (Protocols RSVP)

```
Syntax interface interface-name {
 disable;
 (aggregate | no-aggregate);
 authentication-key key;
 bandwidth bps;
 hello-interval seconds;
 link-protection {
 disable;
 admin-group {
 exclude [group-names];
 include-all [group-names];
 include-any [group-names];
 }
 bandwidth bps;
 bypass bypass-name {
 bandwidth bps {
 ct0 bps;
 ct1 bps;
 ct2 bps;
 ct3 bps;
 }
 description text;
 class-of-service cos-value;
 hop-limit number;
 no-cspf;
 path address <strict | loose>;
 priority setup-priority reservation-priority;
 to address;
 }
 class-of-service cos-value;
 hop-limit number;
 max-bypasses number;
 no-cspf;
 no-node-protection;
 optimize-timer seconds;
 path address <strict | loose>;
 priority setup-priority reservation-priority;
 subscription percentage;
 }
 (reliable | no-reliable);
 subscription percentage {
 ct0 percentage;
 ct1 percentage;
 ct2 percentage;
 ct3 percentage;
 }
 update-threshold threshold;
 }
```

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols rsvp],  
[edit protocols rsvp]

|                                 |                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                  |
| <b>Description</b>              | Enable RSVP on one or more router interfaces.                                                                                                                                                                                                                              |
| <b>Default</b>                  | RSVP is disabled on all interfaces.                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <i>interface-name</i> —Name of an interface. To configure all interfaces, specify <b>all</b> . For details about specifying interfaces, see the <i>Junos OS Network Interfaces Library for Routing Devices</i> .<br><br>The remaining statements are explained separately. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Minimum RSVP Configuration</i></li> </ul>                                                                                                                                                                                      |

## keep-multiplier

---

|                                 |                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | keep-multiplier <i>number</i> ;                                                                                           |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols rsvp],<br>[edit protocols rsvp]                                |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series. |
| <b>Description</b>              | Set the keep multiplier value.                                                                                            |
| <b>Options</b>                  | <i>number</i> —Multiplier value.<br><b>Range:</b> 1 through 255<br><b>Default:</b> 3                                      |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Timers for RSVP Refresh Messages</i></li> </ul>                   |

## link-protection (RSVP)

---

|                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax              | <pre>link-protection {   disable;   admin-group {     exclude [ group-names ];     include-all [ group-names ];     include-any [ group-names ];   }   bandwidth bps;   bypass bypass-name {     bandwidth bps {       ct0 bps;       ct1 bps;       ct2 bps;       ct3 bps;     }     description text;     class-of-service cos-value;     hop-limit number;     no-cspf;     path address &lt;strict   loose&gt;;     priority setup-priority reservation-priority;     to address;   }   class-of-service cos-value;   hop-limit number;   max-bypasses number;   no-cspf;   no-node-protection;   optimize-timer seconds;   path address &lt;strict   loose&gt;;   priority setup-priority reservation-priority;   subscription percentage; }</pre> |
| Hierarchy Level     | [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> ],<br>[edit protocols rsvp interface <i>interface-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Release Information | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series and for EX4600 switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Description         | Enable link protection on the specified interface. Using link protection, you can configure a network to reroute traffic quickly around broken links. To fully enable link protection, you also need to configure the <b>link-protection</b> statement at the [edit protocols mpls label-switched-path <i>lsp-name</i> ] hierarchy level. You can configure single or multiple bypasses for protected interface.                                                                                                                                                                                                                                                                                                                                         |
| Default             | Link protection is disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |



**Options** **no-node-protection**—Disable node-link protection on the RSVP interface. Link protection remains active. When this option is configured, the router can only initiate a next-hop bypass, not a next-next-hop bypass.

The remaining statements are explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Configuring Link Protection on Interfaces Used by LSPs*
- *link-protection (Dynamic LSPs)*

## load-balance (Protocols RSVP)

**Syntax** load-balance {  
    bandwidth;  
}

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols rsvp],  
[edit protocols rsvp]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

**Description** Load-balance traffic between RSVP LSPs.

**Options** **bandwidth**—Load-balance traffic between RSVP LSPs based on the bandwidth configured for each LSP.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Configuring Load Balancing Across RSVP LSPs*

## max-bypasses

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>max-bypasses <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> ],<br>[edit protocols rsvp interface <i>interface-name</i> ]                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Range modified in Junos OS Release 9.3.                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Specify the maximum number of dynamic bypass LSPs permitted for protecting this interface. When this option is configured, multiple bypasses for link protection are enabled. Call admission control (CAC) is also enabled. The limit on bypasses configured applies only to dynamically generated bypass LSPs. By default, this option is disabled and only one dynamic bypass LSP is enabled for each interface. If you configure <b>max-bypasses</b> , you must also configure the <b>bandwidth</b> statement. |
| <b>Options</b>                  | <b>number</b> —Configure the maximum number of bypass LSPs. If you configure a value of 0, no dynamic bypass LSPs are allowed to be established for the interface. Only static bypass LSPs can be configured.<br><b>Range:</b> 0 through 99<br><b>Default:</b> 1                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Link Protection on Interfaces Used by LSPs</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                   |

---

## maximum-helper-recovery-time

---

|                                 |                                                                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>maximum-helper-recovery-time seconds;</code>                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit protocols rsvp <a href="#">graceful-restart</a> ],<br>[edit logical-systems <i>logical-system-name</i> protocols rsvp <a href="#">graceful-restart</a> ]                                                       |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                            |
| <b>Description</b>              | Specify the length of time the router or switch retains the state of its Resource Reservation Protocol (RSVP) neighbors while they undergo a graceful restart.                                                       |
| <b>Options</b>                  | <b>seconds</b> —Length of time that the router retains the state of its Resource Reservation Protocol (RSVP) neighbors while they undergo a graceful restart.<br><b>Range:</b> 1 through 3600<br><b>Default:</b> 180 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Graceful Restart Options for RSVP, CCC, and TCC</i></li><li>• <a href="#">maximum-helper-restart-time (RSVP) on page 5324</a></li></ul>                       |

## maximum-helper-restart-time (RSVP)

---

|                                 |                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>maximum-helper-restart-time seconds;</code>                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit protocols rsvp <a href="#">graceful-restart</a> ],<br>[edit logical-systems <i>logical-system-name</i> protocols rsvp <a href="#">graceful-restart</a> ]                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.3.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                          |
| <b>Description</b>              | Specify the length of time the router or switch waits after it discovers that a neighboring router has gone down before it declares the neighbor down. This value is applied to all RSVP neighbor routers and should be based on the time that the slowest RSVP neighbor requires for restart. |
| <b>Options</b>                  | <b>seconds</b> —The time the router or switch waits after it discovers that a neighboring router has gone down before it declares the neighbor down.<br><b>Range:</b> 1 through 1800<br><b>Default:</b> 60                                                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Graceful Restart Options for RSVP, CCC, and TCC</i></li><li>• <a href="#">maximum-helper-recovery-time on page 5323</a></li></ul>                                                                                                       |

## no-cspf (Protocols RSVP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-cspf;                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection],<br>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i> ],<br>[edit protocols rsvp interface <i>interface-name</i> link-protection],<br>[edit protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.5.                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Disable CSPF computation on all bypass LSPs or on a specific bypass LSP. You need to disable CSPF for link protection to function properly on interarea paths.                                                                                                                                                                                                                                                                                  |
| <b>Default</b>                  | CSPF is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Link Protection on Interfaces Used by LSPs</i></li> </ul>                                                                                                                                                                                                                                                                                                                               |

## no-interface-hello

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-interface-hello;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols rsvp],<br>[edit protocols rsvp]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in JUNOS Release 10.0.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Allows you to explicitly disable RSVP interface hellos globally on the router. This type of configuration might be necessary in networks where the Juniper Networks router has numerous RSVP connections with equipment from other vendors. However, if you disable RSVP interface hellos globally, you can also configure a hello interval on an RSVP interface using the <a href="#">hello-interval (Protocols RSVP)</a> statement. This configuration disables RSVP interface hellos globally but enables RSVP interface hellos on the specified interface. This configuration might be necessary in a heterogeneous network where some devices support RSVP node ID hellos and other devices support RSVP interface hellos. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring RSVP Node ID Hellos</i></li> <li>• <a href="#">hello-interval (Protocols RSVP) on page 5315</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## no-local-reversion

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-local-reversion;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols rsvp],<br>[edit protocols rsvp]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.4.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | <p>Disables RSVP local revertive mode as specified in RFC 4090, <i>Fast Reroute Extensions to RSVP-TE for LSP Tunnels</i>. RSVP local revertive mode is supported on all Juniper Networks routers running the Junos OS. It is the default behavior. If you include this statement, the Juniper Networks router uses global revertive mode instead. You might need to disable RSVP local revertive mode on Juniper Networks routers if your network includes equipment that does not support this mode.</p> <p>The following information can also be found in RFC 4090. Refer to the full RFC for additional information. When an LSP fails, the connection can be repaired locally using a traffic protection mechanism such as fast reroute. To restore the LSP to a full working path, RFC 4090 specifies the following strategies:</p> <ul style="list-style-type: none"><li>• Local revertive mode—Upon detecting that the path is restored, the point of local repair (PLR) resignals each of the LSPs that were formerly routed over the restored path. Every LSP successfully resingaled along the restored path is switched back.</li><li>• Global revertive mode—The ingress router of each tunnel is responsible for reoptimizing the LSPs that used the failed path. There are several potential reoptimization triggers: RSVP error messages, inspection of OSPF LSAs or IS-IS LSPs, and timers. This re-optimization process can proceed as soon as the failure is detected. It is not tied to the restoration of the failed path.</li></ul> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

## no-node-id-subobject

---

|                                 |                                                                                                                                                                                                                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-node-id-subobject;                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols rsvp],<br>[edit protocols rsvp]                                                                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                        |
| <b>Description</b>              | Disable the record route object (RRO) node ID subobject for compatibility with earlier versions of the Junos OS. To interoperate with other vendors' equipment, the Junos OS supports the RRO node ID subobject for use in inter-AS link and node protection configurations. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Inter-AS Node and Link Protection</i></li> </ul>                                                                                                                                                                     |

## no-p2mp-sublsp

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-p2mp-sublsp;                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols rsvp],<br>[edit protocols rsvp]                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.2.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Reject Resv messages that include the S2L_SUB_LSP object. By default, Resv messages that include the S2L_SUB_LSP object are accepted. However, in a network which includes Juniper Networks devices running both Junos OS Release 9.2 and later and Junos OS Release 9.1 and earlier, it is necessary to configure the <b>no-p2mp-sublsp</b> statement on devices running Junos OS Release 9.2 and later to ensure that point-to-multipoint LSPs function properly. |
| <b>Default</b>                  | Resv messages that include the S2L_SUB_LSP object are accepted.                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Preserving Point-to-Multipoint LSP Functioning with Different Junos OS Releases</i></li> </ul>                                                                                                                                                                                                                                                                                                                          |

## node-hello

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | node-hello;                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols rsvp],<br>[edit protocols rsvp]                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced in JUNOS Release 10.0.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Enables node-ID based RSVP hellos globally on all of the RSVP interfaces on the router to allow Juniper Networks routers to interoperate with the equipment of other vendors. By default, the JUNOS Software uses interface-based RSVP hellos and node-ID based RSVP hellos are disabled. If you have not enabled RSVP node IDs on the router, the JUNOS software does not accept any node-ID hello packets. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring RSVP Node ID Hellos</i></li></ul>                                                                                                                                                                                                                                                                                                                     |

## optimize-timer (Protocols RSVP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | optimize-timer <i>seconds</i> ;                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection],<br>[edit protocols rsvp interface <i>interface-name</i> link-protection]                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Configure an optimize timer for a bypass LSP. The optimize timer initiates a periodic optimization process that reshuffles data LSPs among bypass LSPs to achieve the most efficient use of network resources. The optimization process attempts to either minimize the number of bypasses currently in use, minimize the total amount of bandwidth reserved for all bypasses, or both. |
| <b>Options</b>                  | <b><i>seconds</i></b> —Specify the number of seconds between optimizations.<br><b>Range:</b> 0 through 65,535 seconds<br><b>Default:</b> 0 (disabled)                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Link Protection on Interfaces Used by LSPs</i></li></ul>                                                                                                                                                                                                                                                                         |



## path (Protocols RSVP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>path address &lt;strict   loose&gt;;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>],</p> <p>[edit protocols rsvp interface <i>interface-name</i> link-protection],</p> <p>[edit protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Configure an explicit path (a sequence of strict or loose routes) to control where and how a bypass LSP is established. If multiple bypasses are configured, they all will use the same explicit path.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Default</b>                  | No path is configured. CSPF automatically calculates the path the bypass LSP takes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <p><b>address</b>—IP address of each transit router in the LSP. You must specify the address or hostname of each transit router, although you do not need to list each transit router if its type is <b>loose</b>. As an option, you can include the ingress and egress routers in the path. Specify the addresses in order, starting with the ingress router (optional) or the first transit router, and continuing sequentially along the path until reaching the egress router (optional) or the router immediately before the egress router.</p> <p><b>Default:</b> If you do not specify any routers explicitly, no routing limitations are imposed on the bypass LSP.</p> <p><b>loose</b>—(Optional) The next address in the <b>path</b> statement is loose. The LSP can traverse other routers before reaching this router.</p> <p><b>Default:</b> <b>strict</b></p> <p><b>strict</b>—(Optional) The LSP must go to the next address specified in the <b>path</b> statement without traversing other nodes. This is the default.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Link Protection on Interfaces Used by LSPs</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## preemption

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>preemption {<br/>  (aggressive   disabled   normal);<br/>  soft-preemption {<br/>    cleanup-timer <i>seconds</i>;<br/>  }<br/>}</pre>                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols rsvp],<br>[edit protocols rsvp]                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Control RSVP session preemption.                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Default</b>                  | <b>normal</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <p><b>aggressive</b>—Preempt RSVP sessions whenever bandwidth is insufficient to handle all sessions. A session is preempted whenever bandwidth is lowered or a new higher-priority session is established.</p> <p><b>disabled</b>—Do not preempt RSVP sessions.</p> <p><b>normal</b>—Preempt RSVP sessions to accommodate new higher-priority sessions when bandwidth is insufficient to handle all sessions.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Preempting RSVP Sessions</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                            |

## priority (Protocols RSVP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>priority setup-priority reservation-priority;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>],</p> <p>[edit protocols rsvp interface <i>interface-name</i> link-protection],</p> <p>[edit protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>],</p>                                                                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Configure the setup priority and reservation priority for a bypass LSP. If insufficient link bandwidth is available during session establishment, the setup priority is compared with other setup priorities for established sessions on the link to determine whether some of them should be preempted to accommodate the new session. The session with the lower-hold priority is preempted.                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <p><b>reservation-priority</b>—Reservation priority, used to keep a reservation after it has been set up. A smaller number has a higher priority. The priority must be greater than or equal to the setup priority to prevent preemption loops.</p> <p><b>Range:</b> 0 through 7, where 0 is the highest and 7 is the lowest priority.</p> <p><b>Default:</b> 0 (Once the session is set up, no other session can preempt it.)</p> <p><b>setup-priority</b>—Setup priority.</p> <p><b>Range:</b> 0 through 7, where 0 is the highest and 7 is the lowest priority.</p> <p><b>Default:</b> 7 (The session cannot preempt any existing sessions.)</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>Configuring Link Protection on Interfaces Used by LSPs</li> <li>Configuring Priority and Preemption for LSPs</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## refresh-time

---

|                          |                                                                                                                           |
|--------------------------|---------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | refresh-time <i>seconds</i> ;                                                                                             |
| Hierarchy Level          | [edit logical-systems <i>logical-system-name</i> protocols rsvp],<br>[edit protocols rsvp]                                |
| Release Information      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series. |
| Description              | Set the refresh time.                                                                                                     |
| Options                  | <i>seconds</i> —Refresh time.<br><b>Range:</b> 1 through 65,535<br><b>Default:</b> 30 seconds                             |
| Required Privilege Level | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.       |
| Related Documentation    | <ul style="list-style-type: none"><li>• <i>Configuring Timers for RSVP Refresh Messages</i></li></ul>                     |

## reliable

---

|                          |                                                                                                                                                                                                                                                                                                                                                   |
|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | (reliable   no-reliable);                                                                                                                                                                                                                                                                                                                         |
| Hierarchy Level          | [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> protocols rsvp peer-interface <i>peer-interface-name</i> ],<br>[edit protocols rsvp interface <i>interface-name</i> ],<br>[edit protocols rsvp peer-interface <i>peer-interface-name</i> ] |
| Release Information      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                         |
| Description              | Enable reliable message delivery on the interface.<br><br>In order to have refresh reduction and reliable delivery, you must include the <b>aggregate</b> and <b>reliable</b> statements.                                                                                                                                                         |
| Required Privilege Level | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                               |
| Related Documentation    | <ul style="list-style-type: none"><li>• <i>Configuring RSVP Interfaces</i></li><li>• <a href="#">aggregate on page 5306</a></li></ul>                                                                                                                                                                                                             |

---

## setup-protection

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | setup-protection;                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols rsvp],<br>[edit protocols rsvp]                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | The facility-backup fast reroute mechanism can provide setup protection for LSPs which are in the process of being signaled. Both point-to-point LSPs and point-to-multipoint LSPs are supported. You should configure the <b>setup-protection</b> statement on each of the routers along the LSP path on which you want to enable LSP setup protection. You should also configure IGP traffic engineering on all of the routers on the LSP path. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring RSVP Setup Protection</i></li></ul>                                                                                                                                                                                                                                                                                                                                                        |

## subscription

---

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <pre>subscription <i>percentage</i> {<br/>    ct0 <i>percentage</i>;<br/>    ct1 <i>percentage</i>;<br/>    ct2 <i>percentage</i>;<br/>    ct3 <i>percentage</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Hierarchy Level          | [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection],<br>[edit protocols rsvp interface <i>interface-name</i> ],<br>[edit protocols rsvp interface <i>interface-name</i> link-protection]                                                                                                                                                                                                                                                                                                     |
| Release Information      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Description              | Configure the amount of bandwidth subscribed to a class type (when you have enabled Differentiated Services) or bypass LSP (when you have enabled link protection).<br><b>subscription</b> is the percentage of the link bandwidth that can be used for the RSVP reservation process.                                                                                                                                                                                                                                                                                                                                                                           |
| Options                  | <p><b>ctnumber percentage</b>—Percentage of the class-type bandwidth allowed for reservations. If you specify a value greater than 100, you are oversubscribing the class type. You can specify bandwidth subscriptions for class types 0 through 3. This option is not available for bypass LSPs.</p> <p><b>Range:</b> 0 through 65,000</p> <p><b>Default:</b> 100 percent</p> <p><b>percentage</b>—Percentage of the class-type or bypass LSP bandwidth allowed for reservations. If you specify a value greater than 100, you are oversubscribing the class type or bypass LSP.</p> <p><b>Range:</b> 0 through 65,000</p> <p><b>Default:</b> 100 percent</p> |
| Required Privilege Level | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Related Documentation    | <ul style="list-style-type: none"><li>Configuring the Bandwidth Subscription Percentage for LSPs</li><li>Configuring Link Protection on Interfaces Used by LSPs</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

---

## soft-preemption (Protocols RSVP)

---

|                                 |                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>soft-preemption {<br/>    cleanup-timer <i>seconds</i>;<br/>}</code>                                                        |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols RSVP preemption],<br>[edit protocols RSVP preemption]                  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.         |
| <b>Description</b>              | Enable soft preemption to attempt to establish a new path for a preempted LSP before tearing it down.                             |
| <b>Options</b>                  | <b>cleanup-timer</b> —A value of 0 disables soft preemption.<br><b>Range:</b> 0 through 180 seconds<br><b>Default:</b> 30 seconds |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring MPLS Soft Preemption</i></li></ul>                                         |

## splitting-merging

---

|                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax              | <pre>splitting-merging {<br/>    maximum-member-lsps <i>number</i>;<br/>    maximum-signaling-bandwidth <i>bps</i>;<br/>    merging-bandwidth <i>bps</i>;<br/>    minimum-member-lsps <i>number</i>;<br/>    minimum-signaling-bandwidth <i>bps</i>;<br/>    normalization;<br/>    sampling;<br/>    splitting-bandwidth <i>bps</i>;<br/>    splitting-merging-threshold <i>percent</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Hierarchy Level     | [edit protocols mpls container-label-switched-path <i>lsp-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Release Information | Statement introduced in Junos OS Release 14.2.<br>Statement introduced for QFX Switches in Junos OS Release 15.1X53-D30.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Description         | Perform splitting and merging.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Options             | <p><b>maximum-member-lsps <i>number</i></b>—Number of label-switched paths (LSPs) that a container LSP can have as member LSPs at maximum.<br/><b>Default:</b> 1</p> <p><b>maximum-signaling-bandwidth <i>bandwidth</i></b>—Amount of bandwidth in bits per second (bps) that can be signaled for an LSP at maximum after normalization. When <b>maximum-signaling-bandwidth</b> is not configured, the value is derived from the <b>splitting-bandwidth</b>.<br/>When auto-bandwidth adjustment is done between two normalization events, per LSP auto-bandwidth configuration and thresholds are used instead of the <b>splitting-bandwidth</b>.<br/><b>Default:</b> 1 bps</p> <p><b>merging-bandwidth <i>bandwidth</i></b>—Amount of bandwidth in bits per second (bps) that is used for merging during normalization.<br/><b>Default:</b> 1 bps</p> <p><b>minimum-member-lsps <i>number</i></b>—Number of LSPd that a container LSP can have as member LSPs at minimum.<br/><b>Default:</b> 64</p> <p><b>minimum-signaling-bandwidth <i>bandwidth</i></b>—Amount of bandwidth in bits per second (bps) that can be signaled for an LSP at minimum after normalization. When <b>minimum-signaling-bandwidth</b> is not configured, the value is derived from the <b>merging-bandwidth</b>.<br/>When auto-bandwidth adjustment is done between two normalization events, per LSP auto-bandwidth configuration and thresholds are used instead of the <b>merging-bandwidth</b>.<br/><b>Default:</b> 1 bps</p> |



**splitting-bandwidth *bandwidth***—Amount of bandwidth in bits per second (bps) that can be used for splitting during normalization.

**Default:** 1 bps

**splitting-merging-threshold *percent***—Percentage changes in aggregate bandwidth relevant for splitting and merging.

**Default:** 0%

The remaining statements are explained separately.

|                                 |                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration. |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------|

|                              |                                                                                                              |
|------------------------------|--------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">container-label-switched-path on page 5310</a></li></ul> |
|------------------------------|--------------------------------------------------------------------------------------------------------------|

## traceoptions (Protocols RSVP)

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>traceoptions {<br/>    file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;<br/>    flag <i>flag</i> &lt;<i>flag-modifier</i>&gt; &lt;disable&gt;;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> protocols rsvp],<br>[edit protocols rsvp]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>         | Enable RSVP-level trace options.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Default</b>             | The default RSVP-level trace options are those inherited from the routing protocols <b>traceoptions</b> statement included at the [edit routing-options] hierarchy level.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>             | <p><b>disable</b>—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>filename</b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>/var/log</b>. We recommend that you place RSVP tracing output in the file <b>rsvp-log</b>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p><b>Range:</b> 2 through 1000</p> <p><b>Default:</b> 2 files</p> <p>If you specify a maximum number of files, you must also include the <b>size</b> statement to specify the maximum file size.</p> <p><b>flag</b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements.</p> <ul style="list-style-type: none"><li>• <b>all</b>—All tracing operations</li><li>• <b>error</b>—All detected error conditions</li><li>• <b>event</b>—RSVP-related events</li><li>• <b>lmp</b>—RSVP-LMP interactions</li><li>• <b>packets</b>—All RSVP packets</li><li>• <b>path</b>—All path messages</li><li>• <b>pathtear</b>—PathTear messages</li></ul> |

- **resv**—Resv messages
- **resvtear**—ResvTear messages
- **route**—Routing information
- **state**—Session state transitions, including when RSVP-signaled LSPs come up and go down.

**flag-modifier**—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

**no-world-readable**—(Optional) Enable only certain users to read the log file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches this size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 1 MB

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of files.

**world-readable**—(Optional) Enable any user to read the log file.

|                           |                                                                               |
|---------------------------|-------------------------------------------------------------------------------|
| <b>Required Privilege</b> | routing and trace—To view this statement in the configuration.                |
| <b>Level</b>              | routing-control and trace-control—To add this statement to the configuration. |

|                              |                                                                                          |
|------------------------------|------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"> <li>• <i>Tracing RSVP Protocol Traffic</i></li> </ul> |
|------------------------------|------------------------------------------------------------------------------------------|

## tunnel-services (RSVP)

---

|                                 |                                                                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | tunnel-services {<br>devices <i>device-names</i> ;<br>}                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols rsvp],<br>[edit protocols rsvp]                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.1.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                 |
| <b>Description</b>              | Enable ultimate-hop popping on point-to-multipoint LSPs. The Junos OS selects one of the available virtual tunnel (VT) interfaces to de-encapsulate the egress traffic. By default, the selection process is performed automatically. |
| <b>Default</b>                  | Ultimate-hop popping is disabled.                                                                                                                                                                                                     |
| <b>Options</b>                  | <b>devices</b> <i>device-names</i> —Specify which VT interfaces are used to handle the RSVP traffic.<br><b>Range:</b> 0 to 8 devices                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Enabling Ultimate-Hop Popping on Point-to-Multipoint LSPs</i></li></ul>                                                                                                                    |

## update-threshold

---

|                                 |                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | update-threshold <i>threshold</i> ;                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> ],<br>[edit protocols rsvp interface <i>interface-name</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                    |
| <b>Description</b>              | Adjust the threshold at which a change in bandwidth triggers an interior gateway protocol (IGP) update.                                                      |
| <b>Options</b>                  | <b>threshold</b> —Specify the percentage change in bandwidth to trigger an IGP update.<br><b>Range:</b> 1 through 20 percent<br><b>Default:</b> 10 percent   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring RSVP Interfaces</i></li></ul>                                                                         |

## Monitoring Commands for RSVP

---

- `clear mpls container-lsp`
- `clear rsvp session`
- `clear rsvp statistics`
- `ping mpls rsvp`
- `request mpls container-lsp`
- `clear mpls container-lsp`
- `show rsvp interface`
- `show rsvp neighbor`
- `show rsvp session`
- `show rsvp statistics`
- `show rsvp version`
- `traceroute mpls rsvp`

## clear mpls container-lsp

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>clear mpls container-lsp</code><br><code>&lt;autobandwidth&gt;</code><br><code>&lt;history&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code><br><code>&lt;member&gt;</code><br><code>&lt;name <i>name</i>&gt;</code><br><code>&lt;optimize   optimize-aggressive&gt;</code><br><code>&lt;statistics&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.2.<br>Statement introduced for QFX Switches in Junos OS Release 15.1X53-D30.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Release the routes and states associated with MPLS container label-switched paths (LSPs), and start new LSPs.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <p><b>none</b>—Reset and restart all LSPs that originated from this routing device; that is, all LSPs for which this routing device is the ingress routing device. Depending on the number of LSPs involved, it might take a while to restart all the LSPs.</p> <p><b>autobandwidth</b>—(Optional) Clear LSP autobandwidth counters.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>name <i>name</i></b>—(Optional) Reset and restart the specified LSP or group of LSPs. You can include wildcard characters in the interface name, as described in the <i>Junos Network Interfaces Configuration Guide</i>.</p> <p><b>optimize   optimize-aggressive</b>—(Optional) Run nonpreemptive optimization or aggressive optimization computation now.</p> <p><b>statistics</b>—(Optional) Clear LSP statistics. You cannot clear the MPLS LSP statistics using a regular expression (<b>name</b> and <b>path</b> options) on transit routers.</p> |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show mpls container-lsp</a></li><li>• <a href="#">request mpls container-lsp on page 5352</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>List of Sample Output</b>    | <a href="#">clear mpls container-lsp on page 5343</a><br><a href="#">clear mpls container-lsp name on page 5343</a><br><a href="#">clear mpls container-lsp statistics on page 5343</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## Sample Output

clear mpls container-lsp

```
user@host> clear mpls container-lsp
```

clear mpls container-lsp name

```
user@host> clear mpls container-lsp name name
```

clear mpls container-lsp statistics

```
user@host> clear mpls container-lsp statistics
```

## clear rsvp session

---

|                                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                      | <a href="#">Syntax on page 5344</a><br><a href="#">Syntax (EX and QFX Series Switches) on page 5344</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Syntax</b>                              | <pre>clear rsvp session &lt;connection-destination address&gt; &lt;connection-source address&gt; &lt;gracefully&gt; &lt;logical-system (all   logical-system-name)&gt; &lt;lsp-id identifier&gt; &lt;name name&gt; &lt;optimize-fast-reroute&gt; &lt;tunnel-id identifier&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Syntax (EX and QFX Series Switches)</b> | <pre>clear rsvp session &lt;connection-destination address&gt; &lt;connection-source address&gt; &lt;gracefully&gt; &lt;lsp-id identifier&gt; &lt;name name&gt; &lt;optimize-fast-reroute&gt; &lt;tunnel-id identifier&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>                 | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.5 for EX Series switches.<br>Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>                         | Reset and restart Resource Reservation Protocol (RSVP) sessions.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                             | <p><b>none</b>—Reset and restart all RSVP sessions for which this routing device is the ingress, transit, or egress routing device.</p> <p><b>connection-source address</b>—(Optional) Source address for GMPLS and MPLS LSPs from the RSVP sender template.</p> <p><b>connection-destination address</b>—(Optional) Destination address for GMPLS and MPLS LSPs from the RSVP sender template.</p> <p><b>gracefully</b>—(Optional) Gracefully reset an RSVP session for a nonpacket LSP in two passes. In the first pass, the Admin-Status object is signaled along the path to the other endpoint of the RSVP session. In the second pass, the path used by the RSVP session is torn down. This option can only be used on the ingress or egress routing device of the RSVP session and is only valid for nonpacket LSPs.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>lsp-id identifier</b>—(Optional) LSP identifier (source port) for the RSVP sender template.</p> <p><b>name name</b>—(Optional) Reset and restart the specified RSVP session.</p> <p><b>optimize-fast-reroute</b>—(Optional) Begin fast reroute optimization.</p> |



**tunnel-id** *identifier*—(Optional) Tunnel identifier (destination port) for the RSVP session.

**Required Privilege Level** clear

**Related Documentation**

- [clear mpls lsp on page 5107](#)
- [show rsvp session on page 5365](#)

**List of Sample Output** [clear rsvp session on page 5345](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

[clear rsvp session](#)

```
user@host> clear rsvp session
```

## clear rsvp statistics

---

|                                    |                                                                                                                                                                                                                   |
|------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 5346</a><br><a href="#">Syntax (EX Series Switches) on page 5346</a>                                                                                                                   |
| <b>Syntax</b>                      | clear rsvp statistics<br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                     |
| <b>Syntax (EX Series Switches)</b> | clear rsvp statistics                                                                                                                                                                                             |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.5 for EX Series switches.                                                                                             |
| <b>Description</b>                 | Clear Resource Reservation Protocol (RSVP) packet and error statistics.                                                                                                                                           |
| <b>Options</b>                     | <b>none</b> —Clear RSVP packet and error statistics.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>    | clear                                                                                                                                                                                                             |
| <b>Related Documentation</b>       | <ul style="list-style-type: none"><li>• <a href="#">show rsvp statistics on page 5375</a></li></ul>                                                                                                               |
| <b>List of Sample Output</b>       | <a href="#">clear rsvp statistics on page 5346</a>                                                                                                                                                                |
| <b>Output Fields</b>               | When you enter this command, you are provided feedback on the status of your request.                                                                                                                             |

## Sample Output

### clear rsvp statistics

```
user@host> clear rsvp statistics
```

## ping mpls rsvp

**Syntax** ping mpls rsvp  
 <lsp-name>  
 <count count>  
 <destination address>  
 <detail>  
 <dynamic-bypass>  
 <egress egress-address>  
 <exp forwarding-class>  
 <interface interface-name>  
 <logical-system (all | logical-system-name)>  
 <manual-bypass>  
 <multipoint>  
 <size bytes>  
 <source source-address>  
 <standby standby-path-name>  
 <sweep>

**Release Information** Command introduced before Junos OS Release 7.4.  
 The **egress** and **multipoint** options were introduced in Junos OS Release 9.2.  
 The **size** and **sweep** options were introduced in Junos OS Release 9.6.  
 The **dynamic-bypass** and **manual-bypass** options were introduced in Junos OS Release 10.2.  
 Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

**Description** Check the operability of MPLS RSVP-signaled label-switched path (LSP) connections. Type Ctrl+c to interrupt a **ping mpls** command.

**Options** **count count**—(Optional) Number of ping requests to send. If **count** is not specified, five ping requests are sent. The range of values is 1 through 1,000,000. The default value is 5.

**destination address**—(Optional) Specify an address other than the default (127.0.0.1/32) for the ping echo requests. The address can be anything within the 127/8 subnet.

**detail**—(Optional) Display detailed information about the echo requests sent and received.



**NOTE:** When using the **detail** option, the reported time is based on the system time configured on the local and remote routers. Differences in these system times can result in inaccurate one way ping trip times being reported.

In practice, it is difficult to synchronize the system times of independent Juniper Networks routers with sufficient accuracy to provide a meaningful time value for the **detail** option (even when synchronized using NTP).

**dynamic-bypass**—(Optional) Ping dynamically generated bypass LSPs, used for protecting other LSPs.

**egress *egress-address***—(Optional) Only the specified egress router or switch responds to the ping request.

**exp *forwarding-class***—(Optional) Value of the forwarding class for the MPLS ping packets.

**interface**—(Optional) Specify the name of the interface protected by the manual bypass LSP. This option is only available when you have also used the **manual-bypass** option.

**logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on the specified logical system.

***lsp-name***—Ping an RSVP-signaled LSP using an LSP name.

**manual-bypass**—(Optional) Ping manually configured bypass LSPs, used for protecting other LSPs. For this option, you must also specify the interface protected by the manual bypass LSP using the **interface** option.

**multipoint**—(Optional) Send ping requests to each of the egress routers or switches participating in a point-to-multipoint LSP. You can also include the **egress** option to ping a specific egress router or switch participating in a point-to-multipoint LSP.

**size *bytes***—(Optional) Size of the LSP ping request packet (100 through 65468 bytes). Packets are 4-byte aligned. For example, if you enter a size of 101, 102, 103, or 104, the router or switch uses a size value of 104 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 100-byte minimum.

**source *source-address***—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface.

**standby *standby-path-name***—(Optional) Name of the standby path.

**sweep**—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

**Additional Information** If the LSP changes, the label and interface information displayed when you issued the **ping** command continues to be used. You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the remote router or switch to ping an LSP terminating there. You must configure MPLS even if you intend to ping only LDP forwarding equivalence classes (FECs).

In asymmetric MTU scenarios, the echo response might be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

**Required Privilege Level** network

**List of Sample Output** [ping mpls rsvp \(Echo Reply Received\) on page 5349](#)  
[ping mpls rsvp \(Echo Reply with Error Code\) on page 5349](#)

[ping mpls rsvp detail on page 5349](#)

[ping mpls rsvp multipoint egress detail count on page 5349](#)

[ping mpls rsvp multipoint detail count on page 5349](#)

[ping mpls rsvp destination detail count size on page 5350](#)

[ping mpls rsvp destination detail sweep size on page 5350](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. Packets with an error code are not counted in the received packets count. They are accounted for separately.

## Sample Output

### ping mpls rsvp (Echo Reply Received)

```
user@host> ping mpls rsvp test1
!!!!!--- lsping statistics ---5 packets transmitted, 5 packets received, 0% packet
loss
```

### ping mpls rsvp (Echo Reply with Error Code)

```
user@host> ping mpls rsvp test2
!!xxx--- lsping statistics ---5 packets transmitted, 2 packets received, 60%
packet loss3 packets received with error status, not counted as received.
```

### ping mpls rsvp detail

```
user@host> ping mpls rsvp to-green detail
Request for seq 1, to interface 67, labels <100095, 0, 0>
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 67, labels <100095, 0, 0>
Reply for seq 2, return code: Egress-ok
```

### ping mpls rsvp multipoint egress detail count

```
user@host>ping mpls rsvp sample-lsp multipoint egress 192.168.1.3 detail count 1
Request for seq 1, to interface 70, label 299952
Request for seq 1, to interface 70, no label stack.
Request for seq 1, to interface 67, no label stack.

Reply for seq 1, egress 192.168.1.3, return code: Egress-ok, time: 0.242 ms
Local transmit time: 1205310695s 215737us
Remote receive time: 1205310695s 215979us

--- lsping, egress 192.168.1.3 statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
```

### ping mpls rsvp multipoint detail count

```
user@host>ping mpls rsvp sample-lsp multipoint detail count 1
Request for seq 1, to interface 70, label 299952
Request for seq 1, to interface 70, no label stack.
Request for seq 1, to interface 67, no label stack.

Reply for seq 1, return code: Unknown TLV, time: 9.877 m Local transmit time:
1205310615s 347317us
Remote receive time: 1205310615s 357194us
Reply for seq 1, egress 192.168.1.3, return code: Egress-ok, time: 0.351 ms
```

```

Local transmit time: 1205310615s 347262us
Remote receive time: 1205310615s 347613us
Reply for seq 1, egress 192.168.1.13, return code: Egress-ok, time: 0.301 ms
Local transmit time: 1205310615s 347167us
Remote receive time: 1205310615s 347468us
Timeout for seq 1, egress 192.168.1.1
Timeout for seq 1, egress 192.168.1.4
Timeout for seq 1, egress 192.168.1.14

--- lsping, egress 192.168.1.1 statistics ---
1 packets transmitted, 0 packets received, 100% packet loss

--- lsping, egress 192.168.1.3 statistics ---
1 packets transmitted, 1 packets received, 0% packet loss

--- lsping, egress 192.168.1.4 statistics ---
1 packets transmitted, 0 packets received, 100% packet loss

--- lsping, egress 192.168.1.13 statistics ---
1 packets transmitted, 1 packets received, 0% packet loss

--- lsping, egress 192.168.1.14 statistics ---
1 packets transmitted, 0 packets received, 100% packet loss

```

#### ping mpls rsvp destination detail count size

```

user@host> ping mpls rsvp chaser-access destination 192.168.0.1 detail count 1 size 4468

Request for seq 1, to interface 88, label 299984, packet size 4468
Reply for seq 1, return code: Egress-ok, time: 44.804 ms
 Local transmit time: 2009-03-30 22:05:02 CEST 408.629 ms
 Remote receive time: 2009-03-30 22:05:02 CEST 453.433 ms

--- lsping statistics ---
1 packets transmitted, 1 packets received, 0% packet loss

```

#### ping mpls rsvp destination detail sweep size

```

user@router> ping mpls rsvp chaser-access destination 192.168.0.1 detail sweep size 4500
Request for seq 1, to interface 86, no label stack., packet size 100
Reply for seq 1, return code: Egress-ok, time: -39.264 ms
 Local transmit time: 2009-04-24 14:05:40 CEST 541.423 ms
 Remote receive time: 2009-04-24 14:05:40 CEST 502.159 ms
Request for seq 2, to interface 86, no label stack., packet size 2300
Reply for seq 2, return code: Egress-ok, time: -38.179 ms
 Local transmit time: 2009-04-24 14:05:41 CEST 544.240 ms
 Remote receive time: 2009-04-24 14:05:41 CEST 506.061 ms
Request for seq 3, to interface 86, no label stack., packet size 4500
Timeout for seq 3
Request for seq 4, to interface 86, no label stack., packet size 3400
Reply for seq 4, return code: Egress-ok, time: -37.545 ms
 Local transmit time: 2009-04-24 14:05:45 CEST 549.953 ms
 Remote receive time: 2009-04-24 14:05:45 CEST 512.408 ms
Request for seq 5, to interface 86, no label stack., packet size 3952
Reply for seq 5, return code: Egress-ok, time: -37.176 ms
 Local transmit time: 2009-04-24 14:05:46 CEST 555.881 ms
 Remote receive time: 2009-04-24 14:05:46 CEST 518.705 ms
Request for seq 6, to interface 86, no label stack., packet size 4228
Reply for seq 6, return code: Egress-ok, time: -36.962 ms
 Local transmit time: 2009-04-24 14:05:47 CEST 561.809 ms
 Remote receive time: 2009-04-24 14:05:47 CEST 524.847 ms

```

```
Request for seq 7, to interface 86, no label stack., packet size 4368
Reply for seq 7, return code: Egress-ok, time: -36.922 ms
 Local transmit time: 2009-04-24 14:05:48 CEST 568.738 ms
 Remote receive time: 2009-04-24 14:05:48 CEST 531.816 ms
Request for seq 8, to interface 86, no label stack., packet size 4440
Reply for seq 8, return code: Egress-ok, time: -36.855 ms
 Local transmit time: 2009-04-24 14:05:49 CEST 575.669 ms
 Remote receive time: 2009-04-24 14:05:49 CEST 538.814 ms
Request for seq 9, to interface 86, no label stack., packet size 4476
Timeout for seq 9
Request for seq 10, to interface 86, no label stack., packet size 4460
Reply for seq 10, return code: Egress-ok, time: -36.906 ms
 Local transmit time: 2009-04-24 14:05:53 CEST 584.382 ms
 Remote receive time: 2009-04-24 14:05:53 CEST 547.476 ms
Request for seq 11, to interface 86, no label stack., packet size 4480
Timeout for seq 11
Request for seq 12, to interface 86, no label stack., packet size 4472
Timeout for seq 12
Request for seq 13, to interface 86, no label stack., packet size 4468
Reply for seq 13, return code: Egress-ok, time: -36.943 ms
 Local transmit time: 2009-04-24 14:06:00 CEST 594.884 ms
 Remote receive time: 2009-04-24 14:06:00 CEST 557.941 ms
Request for seq 14, to interface 86, no label stack., packet size 4476
Timeout for seq 14
Request for seq 15, to interface 86, no label stack., packet size 4472
Timeout for seq 15

--- lsp ping sweep result---
Maximum Transmission Unit (MTU) is 4468 bytes
```

## request mpls container-lsp

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>request mpls container-lsp</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code><br><code>&lt;name <i>lsp-name</i>&gt;</code><br><code>&lt;adjust-autobandwidth&gt;</code><br><code>&lt;normalization&gt;</code>                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 14.2.<br>Statement introduced for QFX Switches in Junos OS Release 15.1X53-D30.                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Manually trigger a bandwidth allocation adjustment for the container label-switched path (LSP).                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <b>none</b> —Manually trigger a bandwidth allocation adjustment for all active member LSP paths.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.<br><br><b>name <i>lsp-name</i></b> —(Optional) Manually trigger a bandwidth allocation adjustment on the specified member LSP only.<br><br><b>adjust-autobandwidth</b> —(Optional) Request LSP autobandwidth adjustment.<br><br><b>normalization</b> —(Optional) Request container LSP normalization. |
| <b>Required Privilege Level</b> | clear, maintenance                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show mpls container-lsp</a></li><li>• <a href="#">clear mpls container-lsp on page 5342</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>List of Sample Output</b>    | <a href="#">request mpls container-lsp on page 5352</a><br><a href="#">request mpls container-lsp on page 5352</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## Sample Output

### request mpls container-lsp

```
user@host> request mpls container-lsp lsp-name normalize
```

### request mpls container-lsp

```
user@host> request mpls container-lsp normalize bandwidth bps
```



## clear mpls container-lsp

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>clear mpls container-lsp &lt;autobandwidth&gt; &lt;history&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;member&gt; &lt;name <i>name</i>&gt; &lt;optimize   optimize-aggressive&gt; &lt;statistics&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 14.2.</p> <p>Statement introduced for QFX Switches in Junos OS Release 15.1X53-D30.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Release the routes and states associated with MPLS container label-switched paths (LSPs), and start new LSPs.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <p><b>none</b>—Reset and restart all LSPs that originated from this routing device; that is, all LSPs for which this routing device is the ingress routing device. Depending on the number of LSPs involved, it might take a while to restart all the LSPs.</p> <p><b>autobandwidth</b>—(Optional) Clear LSP autobandwidth counters.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>name <i>name</i></b>—(Optional) Reset and restart the specified LSP or group of LSPs. You can include wildcard characters in the interface name, as described in the <i>Junos Network Interfaces Configuration Guide</i>.</p> <p><b>optimize   optimize-aggressive</b>—(Optional) Run nonpreemptive optimization or aggressive optimization computation now.</p> <p><b>statistics</b>—(Optional) Clear LSP statistics. You cannot clear the MPLS LSP statistics using a regular expression (<b>name</b> and <b>path</b> options) on transit routers.</p> |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">show mpls container-lsp</a></li> <li><a href="#">request mpls container-lsp on page 5352</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>List of Sample Output</b>    | <p><a href="#">clear mpls container-lsp on page 5354</a></p> <p><a href="#">clear mpls container-lsp name on page 5354</a></p> <p><a href="#">clear mpls container-lsp statistics on page 5354</a></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## Sample Output

clear mpls container-lsp

```
user@host> clear mpls container-lsp
```

clear mpls container-lsp name

```
user@host> clear mpls container-lsp name name
```

clear mpls container-lsp statistics

```
user@host> clear mpls container-lsp statistics
```

## show rsvp interface

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 5355</a><br><a href="#">Syntax (EX Series Switches) on page 5355</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Syntax</b>                      | <pre>show rsvp interface &lt;brief   detail   extensive&gt; &lt;instance <i>instance-name</i>&gt; &lt;link-management&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Syntax (EX Series Switches)</b> | <pre>show rsvp interface &lt;brief   detail   extensive&gt; &lt;link-management&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>         | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p><b>instance <i>instance-name</i></b> option added in Junos OS Release 15.1.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>                 | Display the status of Resource Reservation Protocol (RSVP)-enabled interfaces and packet statistics.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                     | <p><b>none</b>—Display standard information about the status of RSVP-enabled interfaces and packet statistics.</p> <p><b>brief   detail   extensive   link-management</b>—(Optional) Display the specified level of output.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display RSVP status information for the specified instance. If <i>instance-name</i> is omitted, RSVP status information is displayed for the master instance.</p> <p><b>link-management</b>—(Optional) Use the link-management option to display the control peers and corresponding TE-link information created by the Link Management Protocol (LMP).</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>List of Sample Output</b>       | <a href="#">show rsvp interface brief on page 5358</a><br><a href="#">show rsvp interface detail on page 5358</a><br><a href="#">show rsvp interface extensive on page 5358</a><br><a href="#">show rsvp interface link-management on page 5359</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Output Fields</b>               | Table 424 lists the output fields for the <b>show rsvp interface</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

Table 424: show rsvp interface Output Fields

| Field Name                    | Field Description                                                                                                                                                                                                                                                                                                    | Level of Output  |
|-------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>RSVP interface</b>         | Number of interfaces on which RSVP is active. Each interface has one line of output.                                                                                                                                                                                                                                 | All levels       |
| <b>Interface</b>              | Name of the interface.                                                                                                                                                                                                                                                                                               | All levels       |
| <b>Index</b>                  | Index of the interface.                                                                                                                                                                                                                                                                                              | <b>detail</b>    |
| <b>State</b>                  | State of the interface. <ul style="list-style-type: none"> <li>• <b>Disabled</b>—No traffic engineering information is displayed.</li> <li>• <b>Down</b>—Interface is not operational.</li> <li>• <b>Enabled</b>—Displays traffic engineering information.</li> <li>• <b>Up</b>—Interface is operational.</li> </ul> | All levels       |
| <b>NoAuthentication</b>       | Interface does not support RSVP authentication.                                                                                                                                                                                                                                                                      | <b>detail</b>    |
| <b>NoAggregate</b>            | Interface does not support refresh reduction.                                                                                                                                                                                                                                                                        | <b>detail</b>    |
| <b>NoReliable</b>             | Interface does not support refresh reduction message ID extension.                                                                                                                                                                                                                                                   | <b>detail</b>    |
| <b>NoLinkProtection</b>       | Interface does not support link protection.                                                                                                                                                                                                                                                                          | <b>detail</b>    |
| <b>HelloInterval</b>          | Frequency at which RSVP hellos are sent on this interface (in seconds).                                                                                                                                                                                                                                              | <b>detail</b>    |
| <b>Address</b>                | IP address of the local interface.                                                                                                                                                                                                                                                                                   | <b>detail</b>    |
| <b>Active control channel</b> | Next-hop link address to transmit messages.                                                                                                                                                                                                                                                                          | None specified   |
| <b>TELink</b>                 | Traffic-engineered links that are managed by the peer they are associated with.                                                                                                                                                                                                                                      | None specified   |
| <b>Active resv</b>            | Number of reservations that are actively reserving bandwidth on the interface.                                                                                                                                                                                                                                       | All levels       |
| <b>PreemptionCnt</b>          | Number of times an RSVP session was preempted on this interface.                                                                                                                                                                                                                                                     | <b>detail</b>    |
| <b>Update threshold</b>       | Percentage change in reserved bandwidth to trigger an IGP update.                                                                                                                                                                                                                                                    | <b>detail</b>    |
| <b>Subscription</b>           | User-configured subscription factor.                                                                                                                                                                                                                                                                                 | All levels       |
| <b>bc number</b>              | Bandwidth allocated for the specified bandwidth constraint.                                                                                                                                                                                                                                                          | <b>extensive</b> |
| <b>ct number</b>              | Bandwidth allocated for the specified class type.                                                                                                                                                                                                                                                                    | <b>extensive</b> |
| <b>Static BW</b>              | Total interface bandwidth, in bps.                                                                                                                                                                                                                                                                                   | All levels       |
| <b>Available BW</b>           | Amount of bandwidth that RSVP is allowed to reserve, in bps. It is equal to (static bandwidth * subscription factor).                                                                                                                                                                                                | al levels        |

Table 424: show rsvp interface Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                               | Level of Output |
|--------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Reserved BW</b>             | Currently reserved bandwidth, in bps.                                                                                                                           | All levels      |
| <b>SoftPreemptionCnt</b>       | Number of times a soft preemption occurred on this interface. This number is not included in the <b>PreemptionCnt</b> value.                                    | detail          |
| <b>Overbooked BW</b>           | Currently overbooked bandwidth, in bps, by class type (ct0 through ct3).                                                                                        | detail          |
| <b>Highwater mark</b>          | Highest bandwidth that has ever been reserved on this interface, in bps.                                                                                        | brief           |
| <b>PacketType</b>              | Type of RSVP packet.                                                                                                                                            | detail          |
| <b>Total Sent</b>              | Total number of packets sent.                                                                                                                                   | detail          |
| <b>Total Received</b>          | Total number of packets received since RSVP was enabled.                                                                                                        | detail          |
| <b>Last 5 seconds Sent</b>     | Number of packets sent in the last 5 seconds.                                                                                                                   | detail          |
| <b>Last 5 seconds Received</b> | Number of packets received in the last 5 seconds.                                                                                                               | detail          |
| <b>Path</b>                    | Statistics about Path messages, which are sent from the RSVP sender along the data paths and store path state information in each node along the path.          | detail          |
| <b>PathErr</b>                 | Statistics about PathErr messages, which are advisory messages that are sent upstream to the sender.                                                            | detail          |
| <b>PathTear</b>                | Statistics about PathTear messages, which remove path states and dependent reservation states in any routers along a path.                                      | detail          |
| <b>Resv</b>                    | Statistics about Resv messages, which are sent from the RSVP receiver along the data paths and store reservation state information in each node along the path. | detail          |
| <b>ResvErr</b>                 | Statistics about ResvErr messages, which are advisory messages that are sent when an attempt to establish a reservation fails.                                  | detail          |
| <b>ResvTear</b>                | Statistics about ResvTear messages, which remove reservation states along a path.                                                                               | detail          |
| <b>Hello</b>                   | Number of RSVP hello packets that have been sent to and received from the neighbor.                                                                             | detail          |
| <b>Ack</b>                     | Acknowledge message for refresh reductions.                                                                                                                     | detail          |
| <b>Srefresh</b>                | Summary refresh messages.                                                                                                                                       | detail          |
| <b>EndtoEnd RSVP</b>           | Statistics for the number of end-to-end RSVP messages sent.                                                                                                     | detail          |
| <b>Queue</b>                   | CoS transmit queue number and its associated forwarding class designation.                                                                                      | extensive       |

Table 424: show rsvp interface Output Fields (*continued*)

| Field Name                  | Field Description                                                                                                                                 | Level of Output  |
|-----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>TxRate</b>               | Configured bandwidth in Mbps and configured bandwidth as a percentage of the specified queue.                                                     | <b>extensive</b> |
| <b>Priority</b>             | Weight of the queue relative to other configured queues, in percentage.                                                                           | <b>extensive</b> |
| <i>queue-priority-value</i> | <b>Low, High, None, or Exact.</b> <b>None</b> indicates no rate limiting. <b>Exact</b> indicates the queue transmits at the configured rate only. | <b>extensive</b> |

## Sample Output

### show rsvp interface brief

```

user@host> show rsvp interface brief
RSVP interface: 1 active

```

| Interface | State | Active resv | Subscr-<br>ption | Static<br>BW | Available<br>BW | Reserved<br>BW | Highwater<br>mark |
|-----------|-------|-------------|------------------|--------------|-----------------|----------------|-------------------|
| de0.0     | Up    | 1           | 23%              | 10Mbps       | 989.992kbps     | 1.31Mbps       | 1.31Mbps          |

### show rsvp interface detail

```

user@host> show rsvp interface detail
so-0/1/1.0 Index 6, State: Ena/Up
 NoAuthentication, NoAggregate, NoReliable, NoLinkProtection
 HelloInterval 3(second)
 Address 192.168.207.29, 10.255.245.194
 ActiveResv 0, PreemptionCnt 0, Update threshold 10%
 Subscription 100%, StaticBW 155.52Mbps, AvailableBW 155.52Mbps
 ReservedBW [0] 155Mbps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
 SoftPreemptionCnt1
 OverbookedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 155Mbps[5] 0bps[6] 0bps[7] 0bps
 PacketType
 Total
 Last 5 seconds

```

|               | Total |          | Last 5 seconds |          |
|---------------|-------|----------|----------------|----------|
|               | Sent  | Received | Sent           | Received |
| Path          | 16    | 0        | 1              | 0        |
| PathErr       | 0     | 0        | 0              | 0        |
| PathTear      | 1     | 0        | 0              | 0        |
| Resv          | 0     | 11       | 0              | 1        |
| ResvErr       | 0     | 0        | 0              | 0        |
| ResvTear      | 0     | 0        | 0              | 0        |
| Hello         | 66    | 67       | 1              | 1        |
| Ack           | 0     | 0        | 0              | 0        |
| Srefresh      | 0     | 0        | 0              | 0        |
| EndtoEnd RSVP | 0     | 0        | 0              | 0        |

...

### show rsvp interface extensive

```

user@host> show rsvp interface extensive
so-1/0/0.0 Index 72, State Ena/Up
 NoAuthentication, NoAggregate, NoReliable, NoLinkProtection
 HelloInterval 9(second)
 Address 192.168.213.22, 10.255.240.175
 ActiveResv 1, PreemptionCnt 0, Update threshold 10%
 Subscription 100%,
 bc0 = (ct0+ct1+ct2+ct3), StaticBW 622.08Mbps
 bc1 = (ct1+ct2+ct3), StaticBW 466.56Mbps

```

```

bc2 = (ct2+ct3), StaticBW 311.04Mbps
bc3 = ct3, StaticBW 155.52Mbps
ct0: StaticBW 155.52Mbps, AvailableBW 522.08Mbps
ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
ct1: StaticBW 155.52Mbps, AvailableBW 366.56Mbps
ReservedBW [0] 100Mbps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps

ct2: StaticBW 155.52Mbps, AvailableBW 311.04Mbps
ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
ct3: StaticBW 155.52Mbps, AvailableBW 155.52Mbps
ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps

```

| Queue | TxRate     | Priority | Exact |
|-------|------------|----------|-------|
| 0     | 155.52Mbps | 25%      | Low   |
| 1     | 155.52Mbps | 25%      | Low   |
| 2     | 155.52Mbps | 25%      | Low   |
| 3     | 155.52Mbps | 25%      | Low   |

### show rsvp interface link-management

```

user@host> show rsvp interface link-management
RSVP interface: 2 active
PEER-C State: Up
Active Control Channel: so-0/1/0.0

TElink: TElnk1, Link ID: 37811
ActiveResv 0, PreemptionCnt 0
StaticBW 155.52Mbps, ReservedBW: 0bps, AvailableBW: 155.52Mbps

TElink: TElnk2, Link ID: 37808
ActiveResv 1, PreemptionCnt 0
StaticBW 155.52Mbps, ReservedBW: 0bps, AvailableBW: 155.52Mbps

PEER-B State: Up
Active Control Channel: so-1/0/0.0

TElink: TElnkAB1, Link ID: 1598
ActiveResv 0, PreemptionCnt 0
StaticBW 622.08Mbps, ReservedBW: 0bps, AvailableBW: 622.08Mbps

TElink: TElnkAB2, Link ID: 1597
ActiveResv 0, PreemptionCnt 0
StaticBW 622.08Mbps, ReservedBW: 0bps, AvailableBW: 622.08Mbps

```

## show rsvp neighbor

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 5360</a><br><a href="#">Syntax (EX Series Switches) on page 5360</a>                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Syntax</b>                      | <pre>show rsvp neighbor &lt;brief   detail&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Syntax (EX Series Switches)</b> | <pre>show rsvp neighbor &lt;brief   detail&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>         | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p><b>instance <i>instance-name</i></b> option added in Junos OS Release 15.1.</p>                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>                 | Display Resource Reservation Protocol (RSVP) neighbors that were discovered dynamically during the exchange of RSVP packets.                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                     | <p><b>none</b>—Display standard information about RSVP neighbors.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display the RSVP neighbor information for the specified instance. If <i>instance-name</i> is omitted, RSVP neighbor information is displayed for the master instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>List of Sample Output</b>       | <a href="#">show rsvp neighbor on page 5364</a><br><a href="#">show rsvp neighbor detail on page 5364</a>                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Output Fields</b>               | Table 425 lists the output fields for the <b>show rsvp neighbor</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                          |

**Table 425: show rsvp neighbor Output Fields**

| Field Name           | Field Description                                                                                                                                                    | Level of Output |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>RSVP neighbor</b> | Number of neighbors that the routing device has learned of. Each neighbor has one line of output.                                                                    | All levels      |
| <b>via</b>           | Name of the interface where the neighbor has been detected. In the case of generalized MPLS (GMPLS) LSPs, the name of the peer where the neighbor has been detected. | <b>detail</b>   |
| <b>Address</b>       | Address of a learned neighbor.                                                                                                                                       | All levels      |



Table 425: show rsvp neighbor Output Fields (*continued*)

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Level of Output |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Idle</b>                | Length of time the neighbor has been idle, in seconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | All levels      |
| <b>Up/Dn</b>               | Number of neighbor up or down transitions detected by RSVP hello packets. If the up count is 1 greater than the down count, the neighbor is currently up. Otherwise, the neighbor is down. Neighbors that do not support RSVP hello packets, such as routers running Junos OS Release 3.2 or earlier, are not reported as up or down.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | All levels      |
| <b>Up cnt and Down cnt</b> | Number of neighbor up or down transitions detected by RSVP hello packets. If the up count is 1 greater than the down count, the neighbor is currently up. Otherwise, the neighbor is down. Neighbors that do not support RSVP hello packets, such as routers running Junos OS Release 3.2 or earlier, are not reported as up or down.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>detail</b>   |
| <b>status</b>              | State of the RSVP neighbor: <ul style="list-style-type: none"> <li>• <b>Up</b>—Routing device can detect RSVP Hello messages from the neighbor.</li> <li>• <b>Down</b>—Routing device has received one of the following indications:               <ul style="list-style-type: none"> <li>• Communication failure from the neighbor.</li> <li>• Communication from IGP that the neighbor is unavailable.</li> <li>• Change in the sequence numbers in the RSVP Hello messages sent by the neighbor.</li> </ul> </li> <li>• <b>Restarting</b>—RSVP neighbor is unavailable and might be restarting. The neighbor remains in this state until it has restarted or is declared dead. This state is possible only when graceful restart is enabled.</li> <li>• <b>Restarted</b>—RSVP neighbor has restarted and is undergoing state recovery (graceful restart) procedures.</li> <li>• <b>Dead</b>—Routing device has lost all communication with the RSVP neighbor. Any RSVP sessions with that neighbor are torn down.</li> </ul> | <b>detail</b>   |
| <b>LastChange</b>          | Time elapsed since the neighbor state changed either from up to down or from down to up. The format is <b>hh:mm:ss</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | All levels      |
| <b>Last changed time</b>   | Time elapsed since the neighbor state changed either from up to down or from down to up.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail</b>   |
| <b>HelloInt</b>            | Frequency at which RSVP hellos are sent on this interface (in seconds).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | All levels      |
| <b>HelloTx/Rx</b>          | Number of hello packets sent to and received from the neighbor.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | All levels      |
| <b>Hello</b>               | Number of RSVP hello packets that have been sent to and received from the neighbor.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail</b>   |
| <b>Message received</b>    | Number of Path and Resv messages that this routing device has received from the neighbor.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail</b>   |
| <b>Remote Instance</b>     | Identification provided by the remote routing device during Hello message exchange.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail</b>   |

Table 425: show rsvp neighbor Output Fields (*continued*)

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Level of Output |
|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Local Instance</b>    | Identification sent to the remote routing device during Hello message exchange.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail</b>   |
| <b>Refresh reduction</b> | <p>Measure of processing overhead requests of refresh messages. Refresh reduction extensions improve routing device performance by reducing the process overhead, thus increasing the number of LSPs a routing device can support. <b>Refresh reduction</b> can have the following values:</p> <ul style="list-style-type: none"> <li>• <b>operational</b>—All four RSVP refresh reduction extensions—message ack, bundling, summary refresh, and staged refresh timer—are functional between the two neighboring routing devices. For a detailed explanation of these extensions, see RFC 2961.</li> <li>• <b>incomplete</b>—Some RSVP refresh reduction extensions are functional between the two neighboring routing devices.</li> <li>• <b>no operational</b>—Either the refresh reduction feature has been turned off, or the remote routing device cannot support the refresh reduction extensions.</li> </ul> | <b>detail</b>   |
| <b>Remote end</b>        | <p>Neighboring routing device's status with regard to refresh reduction:</p> <ul style="list-style-type: none"> <li>• <b>enabled</b>—Remote routing device has requested refresh reduction during RSVP message exchanges.</li> <li>• <b>disabled</b>—Remote routing device does not require refresh reduction.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail</b>   |
| <b>Ack-extension</b>     | <p>An RSVP refresh reduction extension:</p> <ul style="list-style-type: none"> <li>• <b>enabled</b>—Both local and remote routing devices support the ack-extension (RFC 2961).</li> <li>• <b>disabled</b>—Remote routing device does not support the ack-extension.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail</b>   |
| <b>Link protection</b>   | <p>Status of the MPLS fast reroute mechanism that protects traffic from link failure:</p> <ul style="list-style-type: none"> <li>• <b>enabled</b>—Link protection feature has been turned on, protecting the neighbor with a bypass LSP.</li> <li>• <b>disabled</b>—No link protection feature has been enabled for this neighbor.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail</b>   |
| <b>LSP name</b>          | Name of the bypass LSP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail</b>   |
| <b>Bypass LSP</b>        | <p>Status of the bypass LSP. It can have the following values:</p> <ul style="list-style-type: none"> <li>• <b>does not exist</b>—Bypass LSP is not available.</li> <li>• <b>connecting</b>—Routing device is in the process of establishing a bypass LSP, and the LSP is not available for link protection at the moment.</li> <li>• <b>operational</b>—Bypass LSP is up and running.</li> <li>• <b>down</b>—Bypass LSP has gone down, with the most probable cause a node or a link failure on the bypass path.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                         | <b>detail</b>   |
| <b>Backup routes</b>     | Number of user LSPs (or routes) that are being protected by a bypass LSP (before link failure).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail</b>   |
| <b>Backup LSPs</b>       | Number of LSPs that have been temporarily established to maintain traffic by refreshing the downstream LSPs during link failure (not a one-to-one correspondence).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>detail</b>   |

Table 425: show rsvp neighbor Output Fields (*continued*)

| Field Name                   | Field Description                                                                                                                                                                                                                                                                                                              | Level of Output |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Bypass explicit route</b> | Explicit route object's (ERO) path that is taken by the bypass LSP.                                                                                                                                                                                                                                                            | <b>detail</b>   |
| <b>Restart time</b>          | Length of time a neighbor waits to receive a Hello from the restarting node before declaring the node dead and deleting the states (in milliseconds).                                                                                                                                                                          | <b>detail</b>   |
| <b>Recovery time</b>         | Length of time during which the restarting node attempts to recover its lost states with help from its neighbors (in milliseconds). Recovery time is advertised by the restarting node to its neighbors, and applies to nodal faults. The restarting node considers its graceful restart complete after this time has elapsed. | <b>detail</b>   |

## Sample Output

### show rsvp neighbor

```
user@host> show rsvp neighbor
RSVP neighbor: 2 learned
Address Idle Up/Dn LastChange HelloInt HelloTx/Rx
192.168.207.203 0 3/2 13:01 3 366/349
192.168.207.207 0 1/0 22:49 3 448/448
```

### show rsvp neighbor detail

```
user@host> show rsvp neighbor detail
RSVP neighbor: 2 learned
Address: 192.168.207.203 via: ecstasy1 status: Up
 Last changed time: 28:47, Idle: 0 sec, Up cnt: 3, Down cnt: 2
 Message received: 632
 Hello: sent 673, received 656, interval 3 sec
 Remote instance: 0x6432838a, Local instance: 0x74b72e36
 Refresh reduction: operational
 Remote end: enabled, Ack-extension: enabled
 Link protection: enabled
 LSP name: Bypass_to_192.168.207.203
 Bypass LSP: operational, Backup routes: 1, Backup LSPs: 0
 Bypass explicit route: 192.168.207.207 192.168.207.224
 Restart time: 60000 msec, Recovery time: 0 msec
```

## show rsvp session

|                                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                      | <a href="#">Syntax on page 5365</a><br><a href="#">Syntax (EX and QFX Series Switches) on page 5365</a>                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Syntax</b>                              | <pre>show rsvp session &lt;brief   detail   extensive   terse&gt; &lt;bidirectional   unidirectional&gt; &lt;bypass&gt; &lt;down   up&gt; &lt;externally-provisioned&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;lsp-type&gt; &lt;name <i>session-name</i>&gt; &lt;p2mp&gt; &lt;session-type&gt; &lt;statistics&gt; &lt;te-link <i>te-link</i>&gt;</pre>                                           |
| <b>Syntax (EX and QFX Series Switches)</b> | <pre>show rsvp session &lt;brief   detail   extensive   terse&gt; &lt;bidirectional   unidirectional&gt; &lt;bypass&gt; &lt;down   up&gt; &lt;externally-provisioned&gt; &lt;interface <i>interface-name</i>&gt; &lt;lsp-type&gt; &lt;name <i>session-name</i>&gt; &lt;p2mp&gt; &lt;session-type&gt; &lt;statistics&gt; &lt;te-link <i>te-link</i>&gt;</pre>                                                                                                                                           |
| <b>Release Information</b>                 | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p><b>externally-provisioned</b> option added in Junos OS Release 13.3.</p> <p>Command introduced in Junos OS Release 13.2X51-D15 for QFX Series.</p> <p><b>instance <i>instance-name</i></b> option added in Junos OS Release 15.1.</p>                                                                                                                              |
| <b>Description</b>                         | Display information about Resource Reservation Protocol (RSVP) sessions.                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                             | <p><b>none</b>—Display standard information about all RSVP sessions.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output.</p> <p><b>bidirectional   unidirectional</b>—(Optional) Display information about bidirectional or unidirectional RSVP sessions only, respectively.</p> <p><b>bypass</b>—(Optional) Display RSVP sessions for bypass LSPs.</p> <p><b>down   up</b>—(Optional) Display only LSPs that are inactive or active, respectively.</p> |

**externally-provisioned**—(Optional) Display the LSPs that are generated dynamically and provisioned by an external Path Computation Element (PCE).

**instance *instance-name***—(Optional) Display RSVP sessions for the specified instance. If *instance-name* is omitted, RSVP session information is displayed for the master instance.

**interface *interface-name***—(Optional) Display RSVP sessions for the specified interface only.

**logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

***lsp-type***—(Optional) Display information about RSVP sessions with regard to LSPs:

- **bypass**—Sessions used for bypass LSPs.
- **lsp**—Sessions used to set up LSPs.
- **nolsp**—Sessions not used to set up LSPs.

**name *session-name***—(Optional) Display information about the named session.

**p2mp**—(Optional) Display point-to-multipoint information.

***session-type***—(Optional) Display information about a particular session type:

- **egress**—Sessions that terminate on this routing device.
- **ingress**—Sessions that originate from this routing device.
- **transit**—Sessions that transit through this routing device.

**statistics**—(Optional) Display packet statistics.

**te-link *te-link***—(Optional) Display sessions with reservations on the specified TE link.

**Required Privilege  
Level**

view

**Related  
Documentation**

- [clear rsvp session on page 5344](#)

**List of Sample Output**

[show rsvp session on page 5370](#)  
[show rsvp session statistics on page 5370](#)  
[show rsvp session detail on page 5371](#)  
[show rsvp session detail \(When Egress Protection is in Standby Mode\) on page 5371](#)  
[show rsvp session detail \(When Egress Protection is in Effect During a Local Repair\) on page 5371](#)  
[show rsvp session detail \(Path MTU Output Field\) on page 5372](#)  
[show rsvp session detail \(GMPLS\) on page 5372](#)  
[show rsvp session extensive on page 5372](#)  
[show rsvp session p2mp \(Ingress Router\) on page 5373](#)  
[show rsvp session p2mp \(Transit Router\) on page 5373](#)

**Output Fields** Table 426 describes the output fields for the **show rsvp session** command. Output fields are listed in the approximate order in which they appear.

**Table 426: show rsvp session Output Fields**

| Field Name        | Field Description                                                                                                                                                                                                                                                                                        | Level of Output |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Ingress RSVP      | Information about ingress RSVP sessions.                                                                                                                                                                                                                                                                 | detail          |
| Ingress RSVP      | Information about ingress RSVP sessions. Each session has one line of output.                                                                                                                                                                                                                            | All levels      |
| Egress RSVP       | Information about egress RSVP sessions.                                                                                                                                                                                                                                                                  | All levels      |
| Transit RSVP      | Information about the transit RSVP sessions.                                                                                                                                                                                                                                                             | All levels      |
| P2MP name         | (Appears only when the <b>p2mp</b> option is specified). Name of the point-to-multipoint LSP path.                                                                                                                                                                                                       | All levels      |
| P2MP branch count | (Appears only when the <b>p2mp</b> option is specified). Number of LSPs receiving packets from the point-to-multipoint LSP.                                                                                                                                                                              | All levels      |
| To                | Destination (egress routing device) of the session.                                                                                                                                                                                                                                                      | All levels      |
| From              | Source (ingress routing device) of the session.                                                                                                                                                                                                                                                          | All levels      |
| State             | State of the path: <b>Up</b> , <b>Down</b> , or <b>AdminDn</b> . <b>AdminDn</b> indicates that the LSP is being taken down gracefully.                                                                                                                                                                   | All levels      |
| Address           | Destination (egress routing device) of the LSP.                                                                                                                                                                                                                                                          | detail          |
| From              | Source (ingress routing device) of the session.                                                                                                                                                                                                                                                          | detail          |
| LSPstate          | State of the LSP that is being handled by this RSVP session. It can be either <b>Up</b> , <b>Dn</b> (down), or <b>AdminDn</b> . <b>AdminDn</b> indicates that the LSP is being taken down gracefully.                                                                                                    | brief detail    |
| Rt                | Number of active routes (prefixes) that have been installed in the routing table. For ingress RSVP sessions, the routing table is the primary IPv4 table ( <b>inet.0</b> ). For transit and egress RSVP sessions, the routing table is the primary MPLS table ( <b>mpls.0</b> ).                         | brief           |
| Active Route      | Number of active routes (prefixes) that have been installed in the forwarding table. For ingress RSVP sessions, the forwarding table is the primary IPv4 table ( <b>inet.0</b> ). For transit and egress RSVP sessions, the forwarding table is the primary MPLS table ( <b>mpls.0</b> ).                | detail          |
| LSPname           | Name of the LSP.                                                                                                                                                                                                                                                                                         | brief detail    |
| LSPpath           | Indicates whether the RSVP session is for the primary or secondary LSP path. <b>LSPpath</b> can be either <b>primary</b> or <b>secondary</b> and can be displayed on the ingress, egress, and transit routing devices. <b>LSPpath</b> can also indicate when a graceful LSP deletion has been triggered. | detail          |

Table 426: show rsvp session Output Fields (*continued*)

| Field Name                      | Field Description                                                                                                                                                                                                                                  | Level of Output     |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------|
| <b>Bypass</b>                   | (Egress routing device) Destination address for the bypass LSP.                                                                                                                                                                                    | <b>detail</b>       |
| <b>Bidir</b>                    | (When LSP is bidirectional) LSP will allow data to travel in both directions between GMPLS devices.                                                                                                                                                | <b>detail</b>       |
| <b>Bidirectional</b>            | (When LSP is bidirectional) LSP will allow data to travel both ways between GMPLS devices.                                                                                                                                                         | <b>detail</b>       |
| <b>Upstream label in</b>        | (When LSP is bidirectional) Incoming label for reverse direction traffic for this LSP.                                                                                                                                                             | <b>detail</b>       |
| <b>Upstream label out</b>       | (When LSP is bidirectional) Outgoing label for reverse direction traffic for this LSP.                                                                                                                                                             | <b>detail</b>       |
| <b>Recovery label received</b>  | (When LSP is bidirectional) Label the upstream node suggests for use in the Resv message that is sent.                                                                                                                                             | <b>detail</b>       |
| <b>Recovery label sent</b>      | (When LSP is bidirectional) Label the downstream node suggests for use in its Resv messages that is returned.                                                                                                                                      | <b>detail</b>       |
| <b>Suggested label received</b> | (When LSP is bidirectional) Label the upstream node suggests for use in the Resv message that is sent.                                                                                                                                             | <b>detail</b>       |
| <b>Suggested label sent</b>     | (When LSP is bidirectional) Label the downstream node suggests for use in its Resv message that is returned.                                                                                                                                       | <b>detail</b>       |
| <b>Resv style or Style</b>      | RSVP reservation style. This field consists of two parts. The first is the number of active reservations. The second is the reservation style, which can be <b>FF</b> (fixed filter), <b>SE</b> (shared explicit), or <b>WF</b> (wildcard filter). | <b>brief detail</b> |
| <b>Label in</b>                 | Incoming label for this LSP.                                                                                                                                                                                                                       | <b>brief detail</b> |
| <b>Label out</b>                | Outgoing label for this LSP.                                                                                                                                                                                                                       | <b>brief detail</b> |
| <b>Time left</b>                | Number of seconds remaining in the lifetime of the reservation.                                                                                                                                                                                    | <b>brief detail</b> |
| <b>Since</b>                    | Date and time when the RSVP session was initiated.                                                                                                                                                                                                 | <b>detail</b>       |
| <b>Tspec</b>                    | Sender's traffic specification, which describes the sender's traffic parameters.                                                                                                                                                                   | <b>detail</b>       |
| <b>DiffServ info</b>            | Indicates whether the LSP is a multiclass LSP ( <b>multiclass diffServ-TE LSP</b> ) or a Differentiated-Services-aware traffic engineering LSP ( <b>diffServ-TE LSP</b> ).                                                                         | <b>detail</b>       |
| <b>bandwidth</b>                | Bandwidth for each class type ( <b>ct0</b> , <b>ct1</b> , <b>ct2</b> , or <b>ct3</b> ).                                                                                                                                                            | <b>detail</b>       |
| <b>Port number</b>              | Protocol ID and sender/receiver port used in this RSVP session.                                                                                                                                                                                    | <b>detail</b>       |
| <b>Attrib flags</b>             | <b>Non-PHP</b> indicates that ultimate hop popping has been requested by the LSP using this RSVP session                                                                                                                                           | <b>extensive</b>    |



Table 426: show rsvp session Output Fields (*continued*)

| Field Name                                          | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Level of Output         |
|-----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>FastReroute desired</b>                          | Fast reroute has been requested by the ingress routing device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail</b>           |
| <b>Soft preemption desired</b>                      | Soft preemption has been requested by the ingress routing device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>detail</b>           |
| <b>FastReroute desired</b>                          | (Data [not a bypass or backup] LSP when the protection scheme has been requested) Fast reroute (one-to-one backup) has been requested by the ingress routing device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail extensive</b> |
| <b>Link protection desired</b>                      | (Data [not a bypass or backup] LSP when the protection scheme has been requested) Link protection (many-to-one backup) has been requested by the ingress routing device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail extensive</b> |
| <b>Node/Link protection desired</b>                 | (Data [not a bypass or backup] LSP when the protection scheme has been requested) Node and link protection (many-to-one backup) has been requested by the ingress routing device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>detail extensive</b> |
| <b>Type</b>                                         | <p>LSP type:</p> <ul style="list-style-type: none"> <li>• <b>Link protected LSP</b>—LSP has been protected by link protection at the outgoing interface. The name of the bypass used is also listed here (<b>extensive</b>).</li> <li>• <b>Node/Link protected LSP</b>—LSP has been protected by node and link protection at the outgoing interface. The name of the bypass used is also listed here (<b>extensive</b>).</li> <li>• <b>Protection down</b>—LSP is not currently protected.</li> <li>• <b>Bypass LSP</b>—LSP that is used to protect one or more user LSPs in case of link failure.</li> <li>• <b>Backup LSP at Point-of-Local-Repair (PLR)</b>—LSP that has been temporarily established to protect a user LSP at the ingress of a failed link.</li> <li>• <b>Backup LSP at Merge Point (MP)</b>—LSP that has been temporarily established to protect a user LSP at the egress of a failed link.</li> </ul> | <b>detail extensive</b> |
| <b>New bypass</b>                                   | New bypass (the bypass name is also displayed) has been activated to protect the LSP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>extensive</b>        |
| <b>Link protection up, using <i>bypass-name</i></b> | Link protection (the bypass name is also displayed) has been activated for the LSP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>extensive</b>        |
| <b>Creating backup LSP, link down</b>               | A <b>link down</b> event occurred, and traffic is being switched over to the bypass LSP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>extensive</b>        |
| <b>Deleting backup LSP, protected LSP restored</b>  | Link has come back up and the LSP has been restored. Because the backup LSP is no longer needed, it is deleted.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>extensive</b>        |
| <b>Path mtu</b>                                     | Displays the value of the path MTU received from the network (through signaling) and the value used for forwarding. This value is only displayed on ingress routing devices with the <b>allow-fragmentation</b> statement configured at the <b>[edit protocols mpls path-mtu]</b> hierarchy level. If there is a detour LSP, the path MTU for the detour is also displayed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>detail</b>           |

Table 426: show rsvp session Output Fields (*continued*)

| Field Name                            | Field Description                                                                                                                                                                                                                                                                   | Level of Output |
|---------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Egress protection<br>PLR as protector | RSVP state on the Protector or the point-of-local-repair (PLR) routing device: <ul style="list-style-type: none"> <li><b>Active</b>— Egress protection is available at the Protector or the PLR routing device.</li> <li><b>In Use</b>— Local repair has been completed.</li> </ul> | detail          |
| PATH rcvfrom                          | Address of the previous-hop (upstream) routing device or client, interface the neighbor used to reach this routing device, and number of packets received from the upstream neighbor.                                                                                               | detail          |
| Adspec                                | MTU signaled from the ingress routing device to the egress routing device by means of the adspec object.                                                                                                                                                                            | detail          |
| PATH sentto                           | Address of the next-hop (downstream) routing device or client, interface used to reach this neighbor (or peer-name in the GMPLS LSP case), and number of packets sent to the downstream routing device.                                                                             | detail          |
| Explct route                          | Explicit route for the session. Normally this value will be the same as that of record route. Differences indicate that path rerouting has occurred, typically during fast reroute.                                                                                                 | detail          |
| Record route                          | Recorded route for the session, taken from the record route object. Normally this value will be the same as that of explct route. Differences indicate that path rerouting has occurred, typically during fast reroute.                                                             | detail          |

## Sample Output

### show rsvp session

```

user@host> show rsvp session
Ingress RSVP: 1 sessions
To From State Rt Style Labelin Labelout LSPName
10.255.245.214 10.255.245.212 AdminDn 0 1 FF - 22293 LSP Bidir
Total 1 displayed, Up 1, Down 0

Egress RSVP: 2 sessions
To From State Rt Style Labelin Labelout LSPName
10.255.245.194 10.255.245.195 Up 0 1 FF 39811 - Gpro3-ba Bidir
10.255.245.194 10.255.245.195 Up 0 1 FF 3 - pro3-ba
Total 2 displayed, Up 2, Down 0

Transit RSVP: 1 sessions
To From State Rt Style Labelin Labelout LSPName
10.255.245.198 10.255.245.197 Up 0 1 SE 100000 3 pro3-de
Total 1 displayed, Up 1, Down 0

```

### show rsvp session statistics

```

user@host> show rsvp session statistics
Ingress RSVP: 2 sessions
To From State Packets Bytes LSPName
10.255.245.24 10.255.245.22 Up 0 0 pro3-bd
10.255.245.24 10.255.245.22 Up 44868 2333136 pro3-bd-2

```

```

Total 2 displayed, Up 2, Down 0
Egress RSVP: 2 sessions
To From State Packets Bytes LSPName
10.255.245.22 10.255.245.24 Up 0 0 pro3-db
10.255.245.22 10.255.245.24 Up 0 0 pro3-db-2
Total 2 displayed, Up 2, Down 0
Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

#### show rsvp session detail

```

user@host> show rsvp session detail
Ingress RSVP: 1 sessions
1.1.1.1
 From: 2.2.2.2, LSPstate: Up, ActiveRoute: 0
 LSPName: to-a, LSPpath: Primary
 Suggested label received: -, Suggested label sent: -
 Recovery label received: -, Recovery label sent: 3
 Resv style: 1 FF, Label in: -, Label out: 3
 Time left: -, Since: Fri Mar 26 18:42:42 2004
 Tspec: rate 300kbps size 300kbps peak Infbps m 20 M 1500
 DiffServ info: diffServ-TE LSP, bandwidth: <ct1 300kbps>
 Port number: sender 1 receiver 15876 protocol 0
 PATH rcvfrom: localclient
 Adspec: sent MTU 1500
 PATH sentto: 192.168.37.16 (t1-0/2/1.0) 1 pkt

```

#### show rsvp session detail (When Egress Protection is in Standby Mode)

```

user@host> show rsvp session detail
Ingress RSVP: 1 sessions
1.1.1.1
 From: 2.2.2.2, LSPstate: Up, ActiveRoute: 0
 LSPName: to-a, LSPpath: Primary
 Suggested label received: -, Suggested label sent: -
 Recovery label received: -, Recovery label sent: 3
 Resv style: 1 FF, Label in: -, Label out: 3
 Time left: -, Since: Fri Mar 26 18:42:42 2004
 Tspec: rate 300kbps size 300kbps peak Infbps m 20 M 1500
 DiffServ info: diffServ-TE LSP, bandwidth: <ct1 300kbps>
 Port number: sender 1 receiver 15876 protocol 0
 Egress protection PLR as protector: Active
 PATH rcvfrom: localclient
 Adspec: sent MTU 1500
 PATH sentto: 192.168.37.16 (t1-0/2/1.0) 1 pkt

```

#### show rsvp session detail (When Egress Protection is in Effect During a Local Repair)

```

user@host> show rsvp session detail
Ingress RSVP: 1 sessions
1.1.1.1
 From: 2.2.2.2, LSPstate: Down, ActiveRoute: 0
 LSPName: to-a, LSPpath: Primary
 Suggested label received: -, Suggested label sent: -
 Recovery label received: -, Recovery label sent: 3
 Resv style: 1 FF, Label in: -, Label out: 3
 Time left: -, Since: Fri Mar 26 18:42:42 2004
 Tspec: rate 300kbps size 300kbps peak Infbps m 20 M 1500
 DiffServ info: diffServ-TE LSP, bandwidth: <ct1 300kbps>
 Port number: sender 1 receiver 15876 protocol 0
 Egress protection PLR as protector: In Use
 PATH rcvfrom: localclient

```

```
Adspec: sent MTU 1500
PATH sentto: 192.168.37.16 (t1-0/2/1.0) 1 pkt
```

#### show rsvp session detail (Path MTU Output Field)

```
user@host> show rsvp session detail
Ingress RSVP: 1 sessions
10.255.245.3
 From: 10.255.245.5, LSPstate: Up, ActiveRoute: 3
 LSPname: to-c, LSPpath: Primary
 Suggested label received: -, Suggested label sent: -
 Recovery label received: -, Recovery label sent: 100432
 Resv style: 1 FF, Label in: -, Label out: 100432
 Time left: -, Since: Mon Aug 16 17:54:40 2006
 Tspec: rate 0bps size 0bps peak Infbps m 20 M 9192
 Port number: sender 1 receiver 57843 protocol 0
 FastReroute desired
 PATH rcvfrom: localclient
 Adspec: sent MTU 4470
 Path mtu: received 4470, using 4458 for forwarding
 PATH sentto: 192.168.37.89 (so-0/2/3.0) 11 pkts
 RESV rcvfrom: 192.168.37.89 (so-0/2/3.0) 10 pkts
 Explct route: 192.168.37.89
 Record route: <self> 192.168.37.89 192.168.37.87
 Detour is Up
 Detour Tspec: rate 0bps size 0bps peak Infbps m 20 M 9192
 Detour adspec: sent MTU 1512
 Path mtu: received 1512, using 1500 for forwarding
```

#### show rsvp session detail (GMPLS)

```
user@host> show rsvp session detail
Ingress RSVP: 1 sessions
192.168.4.1
 From: 192.168.1.1, LSPstate: Dn, ActiveRoute: 0
 LSPname: gmpls-r1-to-r3, LSPpath: Primary
 Bidirectional, Upstream label in: 21253, Upstream label out: -
 Suggested label received: -, Suggested label sent: 21253
 Recovery label received: -, Recovery label sent: -
 Resv style: 0 -, Label in: -, Label out: -
 Time left: -, Since: Mon Aug 16 17:54:40 2006
 Tspec: rate 0bps size 0bps peak 155.52Mbps m 20 M 1500
 Port number: sender 2 receiver 46115 protocol 0
 PATH rcvfrom: localclient
 Adspec: sent MTU 1500
 PATH MTU: received 0
 PATH sentto: 10.35.1.5 (so-0/2/3.0) 11 pkts
 Explct route: 100.100.100.100 93.93.93.93
 Record route: <self> 100.100.100.100 93.93.93.93
 Total 1 displayed, Up 0, Down 1
 Egress RSVP: 0 sessions
 Total 0 displayed, Up 0, Down 0
 Transit RSVP: 0 sessions
 Total 0 displayed, Up 0, Down 0
```

#### show rsvp session extensive

```
user@host> show rsvp session extensive
Ingress RSVP: 1 sessions

192.168.0.4
 From: 192.168.0.5, LSPstate: Up, ActiveRoute: 0
```

```

LSPname: E-D, LSPpath: Primary
LSPtype: Static Configured
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: 299808
Resv style: 1 FF, Label in: -, Label out: 299808
Time left: -, Since: Thu Sep 20 15:54:20 2012
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 2 receiver 61576 protocol 0
Attrib flags: Non-PHP
PATH rcvfrom: localclient
Adspec: sent MTU 1500
Path MTU: received 1500
PATH sentto: 10.0.0.18 (lt-1/2/0.17) 41 pkts
RESV rcvfrom: 10.0.0.18 (lt-1/2/0.17) 40 pkts
Explct route: 10.0.0.18 10.0.0.22
Record route: <self> 10.0.0.18 10.0.0.22
Total 1 displayed, Up 1, Down 0

```

Egress RSVP: 1 sessions

192.168.0.5

```

From: 192.168.0.4, LSPstate: Up, ActiveRoute: 0
LSPname: E-D, LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: -
Resv style: 1 FF, Label in: 3, Label out: -
Time left: 140, Since: Thu Sep 20 15:52:10 2012
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 1 receiver 49601 protocol 0
PATH rcvfrom: 10.0.0.18 (lt-1/2/0.17) 44 pkts
Adspec: received MTU 1500
PATH sentto: localclient
RESV rcvfrom: localclient
Record route: 10.0.0.22 10.0.0.18 <self>
Total 1 displayed, Up 1, Down 0

```

Transit RSVP: 0 sessions

Total 0 displayed, Up 0, Down 0

#### show rsvp session p2mp (Ingress Router)

```

user@host> show rsvp session p2mp
Ingress RSVP: 3 sessions
P2MP name: test, P2MP branch count: 1
 To From State Rt Style Labelin Labelout LSPname
 10.255.10.95 10.255.10.2 Up 0 1 SE - 3 to-pe1
P2MP name: test2, P2MP branch count: 2
 To From State Rt Style Labelin Labelout LSPname
 10.255.10.23 10.255.10.2 Up 0 1 SE - 299776 to-pe3
 10.255.10.16 10.255.10.2 Up 0 1 SE - 299776 to-pe4
Total 3 displayed, Up 3, Down 0

Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

#### show rsvp session p2mp (Transit Router)

```

user@host> show rsvp session p2mp

```

## Ingress RSVP: 1 sessions

P2MP name: test, P2MP branch count: 1

| To           | From         | State | Rt | Style | Labelin | Labelout | LSPname |
|--------------|--------------|-------|----|-------|---------|----------|---------|
| 10.255.10.23 | 10.255.10.95 | Up    | 0  | 1 SE  | -       | 299792   | to-pe2  |

Total 1 displayed, Up 1, Down 0

## Egress RSVP: 1 sessions

P2MP name: test, P2MP branch count: 1

| To           | From        | State | Rt | Style | Labelin | Labelout | LSPname |
|--------------|-------------|-------|----|-------|---------|----------|---------|
| 10.255.10.95 | 10.255.10.2 | Up    | 0  | 1 SE  | 3       | -        | to-pe1  |

Total 1 displayed, Up 1, Down 0

## Transit RSVP: 2 sessions

P2MP name: test2, P2MP branch count: 2

| To           | From        | State | Rt | Style | Labelin | Labelout | LSPname |
|--------------|-------------|-------|----|-------|---------|----------|---------|
| 10.255.10.23 | 10.255.10.2 | Up    | 0  | 1 SE  | 299776  | 299808   | to-pe3  |
| 10.255.10.16 | 10.255.10.2 | Up    | 0  | 1 SE  | 299776  | 299856   | to-pe4  |

Total 2 displayed, Up 2, Down 0

## show rsvp statistics

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 5375</a><br><a href="#">Syntax (EX Series Switches) on page 5375</a>                                                                                                                                                                                                                                                                                                                                                       |
| <b>Syntax</b>                      | <pre>show rsvp statistics &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                       |
| <b>Syntax (EX Series Switches)</b> | show rsvp statistics                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>         | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p><b>instance <i>instance-name</i></b> option added in Junos OS Release 15.1.</p>                                                                                                                                                                                                                                   |
| <b>Description</b>                 | Display Resource Reservation Protocol (RSVP) packet and error statistics.                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                     | <p><b>none</b>—Display RSVP packet and error statistics.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display RSVP packet and error statistics for the specified instance. If <b><i>instance-name</i></b> is omitted, RSVP statistics is displayed for the master instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>       | <ul style="list-style-type: none"> <li><a href="#">clear rsvp statistics on page 5346</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                  |
| <b>List of Sample Output</b>       | <a href="#">show rsvp statistics on page 5378</a>                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Output Fields</b>               | Table 427 describes the output fields for the <b>show rsvp statistics</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                            |

**Table 427: show rsvp statistics Output Fields**

| Field Name                     | Field Description                                        |
|--------------------------------|----------------------------------------------------------|
| <b>Packet Type</b>             | Statistics about different RSVP messages.                |
| <b>Total Sent</b>              | Total number of packets sent since RSVP was enabled.     |
| <b>Total Received</b>          | Total number of packets received since RSVP was enabled. |
| <b>Last 5 seconds Sent</b>     | Total number of packets sent in the last 5 seconds.      |
| <b>Last 5 seconds Received</b> | Number of packets received in the last 5 seconds.        |

Table 427: show rsvp statistics Output Fields (*continued*)

| Field Name                    | Field Description                                                                                                                                            |
|-------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Path</b>                   | Statistics about Path messages, which are sent from the RSVP sender along the data paths and which store path state information in each node along the path. |
| <b>PathErr</b>                | Statistics about PathErr messages, which are advisory messages that are sent upstream to the sender.                                                         |
| <b>PathTear</b>               | Statistics about PathTear messages, which remove path states and dependent reservation states in any routing devices along a path.                           |
| <b>Resv FF</b>                | Statistics about fixed-filter reservation style messages, which consist of distinct reservations among explicit senders.                                     |
| <b>Resv WF</b>                | Statistics about wildcard-filter reservation style messages, which consist of shared reservations among wildcard senders.                                    |
| <b>Res SE</b>                 | Statistics about shared-explicit reservation style messages, which consist of shared reservations among explicit senders.                                    |
| <b>ResvErr</b>                | Statistics about ResvErr messages, which are advisory messages that are sent when an attempt to establish a reservation fails.                               |
| <b>ResvTear</b>               | Statistics about ResvTear messages, which remove reservation states along a path.                                                                            |
| <b>ResvConf</b>               | Statistics about ResvConfirm messages, which are responses to confirm a reservation request.                                                                 |
| <b>Ack</b>                    | Acknowledge message for refresh reductions.                                                                                                                  |
| <b>SRefresh</b>               | Summary refresh messages.                                                                                                                                    |
| <b>Hello</b>                  | Number of RSVP hello packets that have been sent to and received from the neighbor.                                                                          |
| <b>EndtoEnd RSVP</b>          | Statistics for the number of End-to-end RSVP messages.                                                                                                       |
| <b>Errors</b>                 | Statistics about errored RSVP packets.                                                                                                                       |
| <b>Rcv pkt bad length</b>     | The packet was not processed because its length is inappropriate.                                                                                            |
| <b>Rcv pkt unknown type</b>   | The packet is not one of the well-known RSVP types, as defined in RFC 2205, <i>Resource ReSerVation Protocol (RSVP)</i> .                                    |
| <b>Rcv pkt bad version</b>    | The packet is not an RSVP version 1 packet.                                                                                                                  |
| <b>Rcv pkt auth fail</b>      | The packet failed authentication checks.                                                                                                                     |
| <b>Rcv pkt bad checksum</b>   | The RSVP checksum check failed.                                                                                                                              |
| <b>Rcv pkt bad format</b>     | General packet processing failed because the packet was badly formed.                                                                                        |
| <b>Memory allocation fail</b> | An internal resource failure occurred.                                                                                                                       |



Table 427: show rsvp statistics Output Fields (*continued*)

| Field Name                  | Field Description                                                                                                                                                                                                                                                                                                                     |
|-----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| No path information         | A reservation was received, but no sender is active.                                                                                                                                                                                                                                                                                  |
| Resv style conflict         | The same session contains inconsistent reservation styles.                                                                                                                                                                                                                                                                            |
| Port conflict               | There were inconsistent port numbers for the same session.                                                                                                                                                                                                                                                                            |
| Resv no interface           | An interface for the receive reservation packets cannot be located.                                                                                                                                                                                                                                                                   |
| PathErr to client           | Number of PathErr packets delivered to the local client.                                                                                                                                                                                                                                                                              |
| ResvErr to client           | Number of ResvErr packets delivered to the local client.                                                                                                                                                                                                                                                                              |
| Path timeout                | Number of times the sender timed out because the path was removed.                                                                                                                                                                                                                                                                    |
| Resv timeout                | Number of times the receiver timed out because the reservation was removed.                                                                                                                                                                                                                                                           |
| Message out-of-order        | Records the number of RSVP incoming messages that are considered out of order. This is detected from the message ID object's sequence number.                                                                                                                                                                                         |
| Unknown ack msg             | A neighboring routing device replies with an ACK object that contains an unknown message ID. This can indicate a message ID handshake problem. For example, a router receives an ACK for message IDs 1, 2, and 3. However, it only has state for message IDs 1 and 3. The router increments the unknown ack counter by 1.             |
| Recv nack                   | If a neighboring router receives an unknown message ID in an RSVP refresh message, the router sends a Resv nack message back to the sender. This can happen if that neighbor has been rebooted. For this case, the router sends a regular RSVP refresh message to recover the state and start the message-ID handshake process again. |
| Recv duplicated msg-id      | Number of times the same message ID is used by two different RSVP messages. This duplication is usually caused when a neighboring routing device restarts.                                                                                                                                                                            |
| No TE-link to rcv Hop       | Counter of packets discarded because a TE link was not found.                                                                                                                                                                                                                                                                         |
| Rcv pkt disabled interface  | Number of RSVP packets received on an interface that is not enabled for RSVP.                                                                                                                                                                                                                                                         |
| Transmit buffer full        | Number of times the buffer for assembling an outgoing RSVP message was not large enough.                                                                                                                                                                                                                                              |
| Transmit failure            | Number of times the RSVP task failed to send out a packet.                                                                                                                                                                                                                                                                            |
| Receive failure             | Number of times the RSVP task failed to read an incoming packet.                                                                                                                                                                                                                                                                      |
| P2MP RESV discarded by appl | Number of Resv messages discarded because the MPLS label is not valid for the P2MP LSP application.                                                                                                                                                                                                                                   |
| Rate limit                  | Number of RSVP packets dropped due to rate limiting.                                                                                                                                                                                                                                                                                  |

Table 427: show rsvp statistics Output Fields (*continued*)

| Field Name                   | Field Description                                                                                                                                      |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Err msg loop detected</b> | Number of RSVP error messages that have looped back to their originator. This is detected by checking the error node address in the ERROR_SPEC object. |

## Sample Output

### show rsvp statistics

```

user@host> show rsvp statistics
 PacketType Sent Received Last 5 seconds
 Sent Received
 Path 355 408 0 0
 PathErr 2 13 0 0
 PathTear 101 139 0 0
 Resv FF 0 0 0 0
 Resv WF 0 0 0 0
 Resv SE 419 225 0 0
 ResvErr 0 0 0 0
 ResvTear 0 13 0 0
 ResvConf 0 0 0 0
 Ack 682 1414 0 0
 SRefresh 395198 236030 5 2
 Hello 578809 578221 4 4
 EndtoEnd RSVP 0 0 0 0

 Errors Total Last 5 seconds
 Rcv pkt bad length 0 0
 Rcv pkt unknown type 0 0
 Rcv pkt bad version 0 0
 Rcv pkt auth fail 0 0
 Rcv pkt bad checksum 0 0
 Rcv pkt bad format 0 0
 Memory allocation fail 0 0
 No path information 10 0
 Resv style conflict 0 0
 Port conflict 0 0
 Resv no interface 0 0
 PathErr to client 38 0
 ResvErr to client 0 0
 Path timeout 8 0
 Resv timeout 57 0
 Message out-of-order 0 0
 Unknown ack msg 2978 0
 Recv nack 86 0
 Recv duplicated msg-id 5 0
 No TE-link to rcv Hop 0 0
 Rcv pkt disabled interface 0 0
 Transmit buffer full 0 0
 Transmit failure 0 0
 Receive failure 0 0
 P2MP RESV discarded by appl 0 0
 Rate limit 306 0
 Err msg loop detected 0 0

```

## show rsvp version

|                                    |                                                                                                                                                                                                                                          |
|------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 5379</a><br><a href="#">Syntax (EX Series Switches) on page 5379</a>                                                                                                                                          |
| <b>Syntax</b>                      | show rsvp version<br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                |
| <b>Syntax (EX Series Switches)</b> | show rsvp version                                                                                                                                                                                                                        |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.5 for EX Series switches.                                                                                                                    |
| <b>Description</b>                 | Display information about the Resource Reservation Protocol (RSVP) protocol settings, such as the version of the RSVP software, the refresh timer and keep multiplier, and local RSVP graceful restart capabilities on a routing device. |
| <b>Options</b>                     | <b>none</b> —Display RSVP protocol settings.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.                                |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                     |
| <b>List of Sample Output</b>       | <a href="#">show rsvp version on page 5380</a>                                                                                                                                                                                           |
| <b>Output Fields</b>               | <a href="#">Table 428</a> describes the output fields for the <b>show rsvp version</b> command. Output fields are listed in the approximate order in which they appear.                                                                  |

**Table 428: show rsvp version Output Fields**

| Field Name                             | Field Description                                                                                                                                                                                                                                                                                                                               |
|----------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Resource ReSerVation Protocol, version | RSVP software version.                                                                                                                                                                                                                                                                                                                          |
| RSVP protocol                          | Status of RSVP: <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                                                                                                                                                                             |
| R(refresh timer)                       | Configured time interval used to generate periodic RSVP messages.                                                                                                                                                                                                                                                                               |
| K(keep multiplier)                     | Number of RSVP messages that can be lost before an RSVP state is declared stale.                                                                                                                                                                                                                                                                |
| Preemption                             | Currently configured preemption capability: <b>Aggressive</b> , <b>Disabled</b> , or <b>Normal</b> . The default is <b>Normal</b> .                                                                                                                                                                                                             |
| Soft-preemption cleanup                | Time, in seconds, that an LSP is kept after it has been soft preempted. This is a global property of the RSVP protocol.                                                                                                                                                                                                                         |
| Graceful deleting timeout              | Currently configured value for the <b>graceful-deletion-timeout</b> statement. The router that initiates the graceful deletion procedure for an RSVP session waits for the graceful deletion timeout interval to ensure that all routers along the path (especially the ingress and egress routers) have prepared for the LSP to be taken down. |

Table 428: show rsvp version Output Fields (*continued*)

| Field Name                           | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NSR Mode</b>                      | Status of the nonstop active routing feature for RSVP on the restarting device: <b>Disabled</b> , <b>Enabled/Master</b> , or <b>Enabled/Standby</b> .                                                                                                                                                                                                                                                                                |
| <b>NSR State</b>                     | <p>State of the nonstop active routing feature for RSVP on the restarting device.</p> <p>Possible values are:</p> <ul style="list-style-type: none"> <li>• <b>Idle</b></li> <li>• <b>TE-link sync complete</b></li> <li>• <b>Neighbor sync complete</b></li> <li>• <b>Path state sync complete</b></li> <li>• <b>Resv state sync complete</b></li> <li>• <b>Bypass sync complete</b></li> <li>• <b>Init sync complete</b></li> </ul> |
| <b>Setup protection</b>              | Status of point-to-point and point-to-multipoint LSP setup protection configuration on the device: <b>Enabled</b> or <b>Disabled</b>                                                                                                                                                                                                                                                                                                 |
| <b>Graceful restart</b>              | Status of the graceful restart feature for RSVP on the restarting routing device: <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                                                                                                                                                                                                |
| <b>Restart helper mode</b>           | Status of the helper mode feature: <b>Enabled</b> or <b>Disabled</b> . When this feature is enabled, the restarting routing device can help the neighbor with its RSVP restart procedures.                                                                                                                                                                                                                                           |
| <b>Maximum helper restart time</b>   | Number of milliseconds (ms) configured for the maximum helper restart time. The maximum helper restart time is the length of time the routing device waits before declaring that an RSVP neighbor attempting to restart gracefully is down.                                                                                                                                                                                          |
| <b>Maximum helper recovery time</b>  | Number of milliseconds configured for the maximum helper recovery time. The maximum helper recovery time is the amount of time the routing device maintains the state of an RSVP neighbor attempting to restart gracefully.                                                                                                                                                                                                          |
| <b>Restart time</b>                  | Number of milliseconds that a neighbor waits to receive a Hello message from the restarting node before declaring the node dead and deleting the states.                                                                                                                                                                                                                                                                             |
| <b>Recovery time</b>                 | Number of milliseconds during which the restarting node attempts to recover its lost states with help from its neighbors. Recovery time is advertised by the restarting node to its neighbors, and applies to nodal faults. The restarting node considers its graceful restart complete after this time has elapsed.                                                                                                                 |
| <b>P2p transit LSP nexthop mode</b>  | Point-to-point transit LSP nexthop mode on PTX Series devices. The possible values are <b>Chained</b> or <b>Unchained</b>                                                                                                                                                                                                                                                                                                            |
| <b>P2mp transit LSP nexthop mode</b> | Point-to-multipoint transit LSP nexthop mode on PTX Series devices. The possible values are <b>Chained</b> or <b>Unchained</b>                                                                                                                                                                                                                                                                                                       |

## Sample Output

### show rsvp version

```
user@host> show rsvp version
```

```
Resource ReSerVation Protocol, version 1. rfc2205
 RSVP protocol: Enabled
 R(refresh timer): 30 seconds
 K(keep multiplier): 3
 Preemption: Normal
 Soft-preemption cleanup: 30 seconds
 Graceful deletion timeout: 30 seconds
 NSR mode: Enabled/Master
 NSR state: Init sync complete
 Setup protection: Disabled
 Graceful restart: Disabled
 Restart helper mode: Enabled
 Maximum helper restart time: 20000 msec
 Maximum helper recovery time: 180000 msec
 Restart time: 0 msec
 P2p transit LSP nexthop mode: Unchained
 P2mp transit LSP nexthop mode: Unchained
```

## traceroute mpls rsvp

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>traceroute mpls &lt;rsvp&gt; <i>lsp-name</i></code><br><code>&lt;detail&gt;</code><br><code>&lt;egress&gt;</code><br><code>&lt;exp&gt;</code><br><code>&lt;logical-system&gt;</code><br><code>&lt;multipoint&gt;</code><br><code>&lt;no-resolve&gt;</code><br><code>&lt;retries&gt;</code><br><code>&lt;source <i>source-address</i>&gt;</code><br><code>&lt;ttl&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.2.<br><code>egress</code> , <code>multipoint</code> , and <code>ttl</code> options added in Junos OS Release 11.2.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Trace route to a remote host for an MPLS LSP signaled by RSVP. Use <b>traceroute mpls rsvp</b> as a debugging tool to locate MPLS label-switched path (LSP) forwarding issues in a network. (Currently supported for IPv4 packets only.)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <p><b><i>lsp-name</i></b>—Specify the name of the LSP to be traced.</p> <p><b><code>detail</code></b>—(Optional) Display detailed output.</p> <p><b><code>egress</code></b>—(Optional) Request that a specific point-to-multipoint egress node reply to the trace route. The trace route would follow the associated sub-LSP to the egress node.</p> <p><b><code>exp</code></b>—(Optional) Specify the class of service to use when sending probes. The range of values is 0 through 7. The default value is 7.</p> <p><b><code>logical-system</code></b>—(Optional) Specify the name of the logical system for the traceroute attempt.</p> <p><b><code>multipoint</code></b>—(Optional) Perform a trace route on a point-to-multipoint LSP.</p> <p><b><code>no-resolve</code></b>—(Optional) Specify not to resolve the hostname that corresponds to the IP address.</p> <p><b><code>retries</code></b>—(Optional) Specify the number of times to resend probe. The range of values is 1 through 9. The default value is 3.</p> <p><b><code>source <i>source-address</i></code></b>—(Optional) Specify the source address of the outgoing traceroute packets.</p> <p><b><code>ttl</code></b>—(Optional) Specify the number of hops to follow before forcing the trace route to quit.</p> |
| <b>Required Privilege Level</b> | network                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>List of Sample Output</b>    | <a href="#">traceroute mpls rsvp on page 5384</a><br><a href="#">traceroute mpls rsvp detail on page 5384</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

[traceroute mpls rsvp multipoint \(branch node for sub-LSPs\) on page 5385](#)  
[traceroute mpls rsvp multipoint \(single-hop sub-LSPs\) on page 5385](#)

**Output Fields** Table 429 describes the output fields for the **traceroute mpls rsvp *lsp-name*** and **traceroute mpls rsvp *lsp-name* detail** commands. Output fields are listed in the approximate order in which they appear.

**Table 429: traceroute mpls rsvp Output Fields**

| Field Name         | Field Description                                                                                                                                                                                                                                          | Level of Output |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Probe options      | Probe options specified in the <b>traceroute mpls rsvp <i>lsp-name</i></b> command.                                                                                                                                                                        | all levels      |
| ttl                | Time-to-live value of the labeled packet.                                                                                                                                                                                                                  | none specified  |
| Label              | MPLS label used to forward the packets along the LSP.                                                                                                                                                                                                      | none specified  |
| Protocol           | Signaling protocol used. For this command, it is RSVP-TE.                                                                                                                                                                                                  | none specified  |
| Address            | Address of the next hop.                                                                                                                                                                                                                                   | none specified  |
| Previous Hop       | Address of the previous hop. Previous hop address of the first hop is null.                                                                                                                                                                                | none specified  |
| Probe status       | Forwarding status from the first hop to the last-hop label-switching router (egress point in the label-switched paths). Displays <b>Success</b> if the trace to a hop is successful or <b>Egress</b> if the trace has reached the last router on the path. | none specified  |
| Hop                | Address of the hops in the label-switched path from the first hop to the last hop. Depth indicates the level of the hop.                                                                                                                                   | <b>detail</b>   |
| Parent             | Address of the previous hop. Parent value for the first hop is null.                                                                                                                                                                                       | <b>detail</b>   |
| Return Code        | Return code for reporting the result of processing the echo request by the receiver.                                                                                                                                                                       | <b>detail</b>   |
| Sender timestamp   | Displays the timestamp when the MPLS echo request is sent to the next hop.                                                                                                                                                                                 | <b>detail</b>   |
| Receiver timestamp | Timestamp when the echo request from the previous hop is received and acknowledged with an echo response by the next hop.                                                                                                                                  | <b>detail</b>   |
| Response time      | Time for the echo request to reach the receiver.                                                                                                                                                                                                           | <b>detail</b>   |
| MTU                | Size of the largest packet that includes the label stack forwarded to the next hop.                                                                                                                                                                        | <b>detail</b>   |

Table 429: traceroute mpls rsvp Output Fields (*continued*)

| Field Name     | Field Description                                                                                        | Level of Output |
|----------------|----------------------------------------------------------------------------------------------------------|-----------------|
| Multipath type | Labels or addresses used by the specified multipath type. If multipaths are not used, the value is none. | <b>detail</b>   |
| Label stack    | Label stack used to forward the packet.                                                                  | <b>detail</b>   |
| Path           | Displays the sub-lsp path number for this traceroute, the interface used, and the destination address.   | all levels      |

## Sample Output

### traceroute mpls rsvp

```
user@host> traceroute mpls rsvp lsp-chicago-atlanta
```

```
Probe options: retries 3, exp 7
```

| ttl | Label  | Protocol | Address     | Previous Hop | Probe Status |
|-----|--------|----------|-------------|--------------|--------------|
| 1   | 299792 | RSVP-TE  | 192.168.1.2 | (null)       | Success      |
| 2   | 299803 | RSVP-TE  | 192.168.2.3 | 192.168.1.2  | Success      |
| 3   | 3      | RSVP-TE  | 192.168.3.4 | 192.168.2.3  | Egress       |

```
Path 1 via ge-0/0/0.1 destination 127.0.0.64
```

### traceroute mpls rsvp detail

```
user@host> traceroute mpls rsvp lsp-chicago-atlanta detail
```

```
Probe options: retries 3, exp 7
```

```
Hop 192.168.1.2 Depth 1
```

```
Probe status: Success
```

```
Parent: (null)
```

```
Return code: Label-switched at stack-depth 1
```

```
Sender timestamp: 2008-04-17 09:35:27 EDT 400.88 msec
```

```
Receiver timestamp: 2008-04-17 09:35:27 EDT 427.87 msec
```

```
Response time: 26.99 msec
```

```
MTU: Unknown
```

```
Multipath type: IP bitmask
```

```
Address Range 1: 127.0.0.64 ~ 127.0.0.127
```

```
Label Stack:
```

```
Label 1 Value 299792 Protocol RSVP-TE
```

```
Hop 192.168.2.3 Depth 2
```

```
Probe status: Success
```

```
Parent: 192.168.1.2
```

```
Return code: Upstream interface index unknown label-switched at stack-depth
```

```
1
```

```
Sender timestamp: 2008-04-17 09:35:27 EDT 522.13 msec
```

```
Receiver timestamp: 2008-04-17 09:35:27 EDT 548.69 msec
```

```
Response time: 26.55 msec
```

```
MTU: 1518
```

```
Multipath type: IP bitmask
```

```
Address Range 1: 127.0.0.64 ~ 127.0.0.127
```

```
Label Stack:
```

```
Label 1 Value 299803 Protocol RSVP-TE
```



**traceroute mpls rsvp multipoint (branch node for sub-LSPs)**

The following traceroute output is for a point-to-multipoint LSP where the penultimate node is a branch node for the sub-LSPs.

```
user@host> traceroute mpls rsvp multipoint p2mplsp
Probe options: retries 3, exp 7
```

| ttl | Label  | Protocol | Address  | Previous Hop | Probe Status |
|-----|--------|----------|----------|--------------|--------------|
| 1   | 300000 | RSVP-TE  | 81.1.2.2 | (null)       | Success      |
| 2   | 299968 | RSVP-TE  | 81.2.3.3 | 81.1.2.2     | Success      |
| 3   | 299952 | RSVP-TE  | 81.3.4.4 | 81.2.3.3     | Success      |
| 4   | 299920 | RSVP-TE  | 81.4.6.6 | 81.3.4.4     | Egress       |

Path 1 via lt-1/2/0.102 destination 127.0.0.64

| ttl | Label  | Protocol | Address  | Previous Hop | Probe Status |
|-----|--------|----------|----------|--------------|--------------|
| 4   | 299920 | RSVP-TE  | 81.4.5.5 | 81.3.4.4     | Egress       |

Path 2 via lt-1/2/0.102 destination 127.0.0.64

**traceroute mpls rsvp multipoint (single-hop sub-LSPs)**

The following traceroute output is for a point-to-multipoint LSP with multiple single-hop sub-LSPs.

```
user@host> traceroute mpls rsvp multipoint p2mplsp
Probe options: retries 3, exp 7
```

| ttl | Label | Protocol | Address  | Previous Hop | Probe Status |
|-----|-------|----------|----------|--------------|--------------|
| 1   | 0     | RSVP-TE  | 81.1.2.2 | (null)       | Egress       |

Path 1 via lt-1/2/0.102 destination 127.0.0.64

| ttl | Label | Protocol | Address  | Previous Hop | Probe Status |
|-----|-------|----------|----------|--------------|--------------|
| 1   | 0     | RSVP-TE  | 81.1.8.8 | (null)       | Egress       |

Path 2 via lt-1/2/0.108 destination 127.0.0.64

| ttl | Label | Protocol | Address  | Previous Hop | Probe Status |
|-----|-------|----------|----------|--------------|--------------|
| 1   | 0     | RSVP-TE  | 81.1.9.9 | (null)       | Egress       |

Path 3 via lt-1/2/0.109 destination 127.0.0.64



# Multicast Protocols Feature Guide for QFX10000 Switches



# Overview

- [Multicast Overview on page 5389](#)

## Multicast Overview

---

IP has three fundamental types of addresses: unicast, broadcast, and multicast. A *unicast address* is used to send a packet to a single destination. A *broadcast address* is used to send a datagram to an entire subnetwork. A *multicast address* is used to send a datagram to a set of hosts that can be on different subnetworks and that are configured as members of a multicast group.

A multicast datagram is delivered to destination group members with the same best-effort reliability as a standard unicast IP datagram. This means that multicast datagrams are not guaranteed to reach all members of a group or to arrive in the same order in which they were transmitted. The only difference between a multicast IP packet and a unicast IP packet is the presence of a group address in the IP header destination address field. Multicast addresses use the Class D address format.



**NOTE:** On all SRX Series devices, reordering is not supported for multicast fragments. Reordering of unicast fragments is supported.

Individual hosts can join or leave a multicast group at any time. There are no restrictions on the physical location or the number of members in a multicast group. A host can be a member of more than one multicast group at any time. A host does not have to belong to a group to send packets to members of a group.

Routers use a group membership protocol to learn about the presence of group members on directly attached subnetworks. When a host joins a multicast group, it transmits a group membership protocol message for the group or groups that it wants to receive and sets its IP process and network interface card to receive frames addressed to the multicast group.

## Comparing Multicast to Unicast

The Junos<sup>®</sup> operating system (Junos OS) routing protocol process supports a wide variety of routing protocols. These routing protocols carry network information among routing devices not only for *unicast* traffic streams sent between one pair of clients and servers,

but also for *multicast* traffic streams containing video, audio, or both, between a single server source and many client receivers. The routing protocols used for multicast differ in many key ways from unicast routing protocols.

Information is delivered over a network by three basic methods: unicast, broadcast, and multicast.

The differences among unicast, broadcast, and multicast can be summarized as follows:

- Unicast: One-to-one, from one source to one destination.
- Broadcast: One-to-all, from one source to all possible destinations.
- Multicast: One-to-many, from one source to multiple destinations expressing an interest in receiving the traffic.



**NOTE:** This list does not include a special category for many-to-many applications, such as online gaming or videoconferencing, where there are many sources for the same receiver and where receivers often double as sources. Many-to-many is a service model that repeatedly employs one-to-many multicast and therefore requires no unique protocol. The original multicast specification, RFC 1112, supports both the any-source multicast (ASM) many-to-many model and the source-specific multicast (SSM) one-to-many model.

---

With unicast traffic, many streams of IP packets that travel across networks flow from a single source, such as a website server, to a single destination such as a client PC. Unicast traffic is still the most common form of information transfer on networks.

Broadcast traffic flows from a single source to all possible destinations reachable on the network, which is usually a LAN. Broadcasting is the easiest way to make sure traffic reaches its destinations.

Television networks use broadcasting to distribute video and audio. Even if the television network is a cable television (CATV) system, the source signal reaches all possible destinations, which is the main reason that some channels' content is scrambled. Broadcasting is not feasible on the Internet because of the enormous amount of unnecessary information that would constantly arrive at each end user's device, the complexities and impact of scrambling, and related privacy issues.

Multicast traffic lies between the extremes of unicast (one source, one destination) and broadcast (one source, all destinations). Multicast is a "one source, many destinations" method of traffic distribution, meaning only the destinations that explicitly indicate their need to receive the information from a particular source receive the traffic stream.

On an IP network, because destinations (clients) do not often communicate directly with sources (servers), the routing devices between source and destination must be able to determine the topology of the network from the unicast or multicast perspective to avoid routing traffic haphazardly. Multicast routing devices replicate packets received on one input interface and send the copies out on multiple output interfaces.

In IP multicast, the source and destination are almost always hosts and not routing devices. Multicast routing devices distribute the multicast traffic across the network from source to destinations. The multicast routing device must find multicast sources on the network, send out copies of packets on several interfaces, prevent routing loops, connect interested destinations with the proper source, and keep the flow of unwanted packets to a minimum. Standard multicast routing protocols provide most of these capabilities, but some router architectures cannot send multiple copies of packets and so do not support multicasting directly.

## IP Multicast Uses

Multicast allows an IP network to support more than just the unicast model of data delivery that prevailed in the early stages of the Internet. Multicast, originally defined as a host extension in RFC 1112 in 1989, provides an efficient method for delivering traffic flows that can be characterized as one-to-many or many-to-many.

Unicast traffic is not strictly limited to data applications. Telephone conversations, wireless or not, contain digital audio samples and might contain digital photographs or even video and still flow from a single source to a single destination. In the same way, multicast traffic is not strictly limited to multimedia applications. In some data applications, the flow of traffic is from a single source to many destinations that require the packets, as in a news or stock ticker service delivered to many PCs. For this reason, the term *receiver* is preferred to *listener* for multicast destinations, although both terms are common.

Network applications that can function with unicast but are better suited for multicast include collaborative groupware, teleconferencing, periodic or “push” data delivery (stock quotes, sports scores, magazines, newspapers, and advertisements), server or website replication, and distributed interactive simulation (DIS) such as war simulations or virtual reality. Any IP network concerned with reducing network resource overhead for one-to-many or many-to-many data or multimedia applications with multiple receivers benefits from multicast.

If unicast were employed by radio or news ticker services, each radio or PC would have to have a separate traffic session for each listener or viewer at a PC (this is actually the method for some Web-based services). The processing load and bandwidth consumed by the server would increase linearly as more people “tune in” to the server. This is extremely inefficient when dealing with the global scale of the Internet. Unicast places the burden of packet duplication on the server and consumes more and more backbone bandwidth as the number of users grows.

If broadcast were employed instead, the source could generate a single IP packet stream using a broadcast destination address. Although broadcast eliminates the server packet duplication issue, this is not a good solution for IP because IP broadcasts can be sent only to a single subnetwork, and IP routing devices normally isolate IP subnetworks on separate interfaces. Even if an IP packet stream could be addressed to literally go everywhere, and there were no need to “tune” to any source at all, broadcast would be extremely inefficient because of the bandwidth strain and need for uninterested hosts to discard large numbers of packets. Broadcast places the burden of packet rejection on each host and consumes the maximum amount of backbone bandwidth.

For radio station or news ticker traffic, multicast provides the most efficient and effective outcome, with none of the drawbacks and all of the advantages of the other methods. A single source of multicast packets finds its way to every *interested* receiver. As with broadcast, the transmitting host generates only a single stream of IP packets, so the load remains constant whether there is one receiver or one million. The network routing devices replicate the packets and deliver the packets to the proper receivers, but only the replication role is a new one for routing devices. The links leading to subnets consisting of entirely uninterested receivers carry no multicast traffic. Multicast minimizes the burden placed on sender, network, and receiver.


## IP Multicast Terminology

Multicast has its own particular set of terms and acronyms that apply to IP multicast routing devices and networks. [Figure 152](#) depicts some of the terms commonly used in an IP multicast network.

In a multicast network, the key component is the *routing device*, which is able to replicate packets and is therefore multicast-capable. The routing devices in the IP multicast network, which has exactly the same topology as the unicast network it is based on, use a *multicast routing protocol* to build a *distribution tree* that connects receivers (preferred to the multimedia implications of listeners, but listeners is also used) to *sources*. In multicast terminology, the distribution tree is *rooted at the source* (the root of the distribution tree is the source). The interface on the routing device leading toward the source is the *upstream* interface, although the less precise terms *incoming* or *inbound* interface are used as well. To keep bandwidth use to a minimum, it is best for only one upstream interface on the routing device to receive multicast packets. The interface on the routing device leading toward the receivers is the *downstream* interface, although the less precise terms *outgoing* or *outbound* interface are used as well. There can be 0 to  $N-1$  downstream interfaces on a routing device, where  $N$  is the number of logical interfaces on the routing device. To prevent looping, the upstream interface must never receive copies of downstream multicast packets.



Figure 152: Multicast Terminology in an IP Network



Routing loops are disastrous in multicast networks because of the risk of repeatedly replicated packets. One of the complexities of modern multicast routing protocols is the need to avoid routing loops, packet by packet, much more rigorously than in unicast routing protocols.

### Reverse-Path Forwarding for Loop Prevention

The routing device's multicast forwarding state runs more logically based on the reverse path, from the receiver back to the root of the distribution tree. In RPF, every multicast packet received must pass an RPF check before it can be replicated or forwarded on any interface. When it receives a multicast packet on an interface, the routing device verifies that the *source* address in the multicast IP packet is the *destination* address for a unicast IP packet back to the source.

If the outgoing interface found in the unicast routing table is the same interface that the multicast packet was received on, the packet passes the RPF check. Multicast packets that fail the RPF check are dropped, because the incoming interface is not on the shortest path back to the source. routing devices can build and maintain separate tables for RPF purposes.

### Shortest-Path Tree for Loop Prevention

The distribution tree used for multicast is rooted at the source and is the shortest-path tree (SPT), but this path can be long if the source is at the periphery of the network. Providing a *shared tree* on the backbone as the distribution tree locates the multicast source more centrally in the network. Shared distribution trees with roots in the core network are created and maintained by a multicast routing device operating as a rendezvous point (RP), a feature of sparse mode multicast protocols.

## Administrative Scoping for Loop Prevention

Scoping limits the routing devices and interfaces that can forward a multicast packet. Multicast scoping is *administrative* in the sense that a range of multicast addresses is reserved for scoping purposes, as described in RFC 2365, *Administratively Scoped IP Multicast*. routing devices at the boundary must filter multicast packets and ensure that packets do not stray beyond the established limit.

## Multicast Leaf and Branch Terminology

Each subnetwork with hosts on the routing device that has at least one interested receiver is a *leaf* on the distribution tree. routing devices can have multiple leaves on different interfaces and must send a copy of the IP multicast packet out on each interface with a leaf. When a new leaf subnetwork is added to the tree (that is, the interface to the host subnetwork previously received no copies of the multicast packets), a new *branch* is built, the leaf is joined to the tree, and replicated packets are sent out on the interface. The number of leaves on a particular interface does not affect the routing device. The action is the same for one leaf or a hundred.



**NOTE:** On Juniper Networks security devices, if the maximum number of leaves on a multicast distribution tree is exceeded, multicast sessions are created up to the maximum number of leaves, and any multicast sessions that exceed the maximum number of leaves are ignored. The maximum number of leaves on a multicast distribution tree is device specific.

When a branch contains no leaves because there are no interested hosts on the routing device interface leading to that IP subnetwork, the branch is *pruned* from the distribution tree, and no multicast packets are sent out that interface. Packets are replicated and sent out multiple interfaces only where the distribution tree branches at a routing device, and no link ever carries a duplicate flow of packets.

Collections of hosts all receiving the same stream of IP packets, usually from the same multicast source, are called *groups*. In IP multicast networks, traffic is delivered to multicast groups based on an IP multicast address, or *group address*. The groups determine the location of the leaves, and the leaves determine the branches on the multicast network.

## IP Multicast Addressing

Multicast uses the Class D IP address range (224.0.0.0 through 239.255.255.255). Class D addresses are commonly referred to as *multicast addresses* because the entire classful address concept is obsolete. Multicast addresses can never appear as the source address in an IP packet and can only be the destination of a packet.

Multicast addresses usually have a prefix length of /32, although other prefix lengths are allowed. Multicast addresses represent logical groupings of receivers and not physical collections of devices. Blocks of multicast addresses can still be described in terms of prefix length in traditional notation, but only for convenience. For example, the multicast

address range from 232.0.0.0 through 232.255.255.255 can be written as 232.0.0.0/8 or 232/8.

Internet service providers (ISPs) do not typically allocate multicast addresses to their customers because multicast addresses relate to content, not to physical devices. Receivers are not assigned their own multicast addresses, but need to know the multicast address of the content. Sources need to be assigned multicast addresses only to produce the content, not to identify their place in the network. Every source and receiver still needs an ordinary, unicast IP address.

Multicast addressing most often references the receivers, and the source of multicast content is usually not even a member of the multicast group for which it produces content. If the source needs to monitor the packets it produces, monitoring can be done locally, and there is no need to make the packets traverse the network.

Many applications have been assigned a range of multicast addresses for their own use. These applications assign multicast addresses to sessions created by that application. You do not usually need to statically assign a multicast address, but you can do so.

## Multicast Addresses

Multicast host group addresses are defined to be the IP addresses whose high-order four bits are 1110, giving an address range from 224.0.0.0 through 239.255.255.255, or simply 224.0.0.0/4. (These addresses also are referred to as Class D addresses.)

The Internet Assigned Numbers Authority (IANA) maintains a list of registered IP multicast groups. The base address 224.0.0.0 is reserved and cannot be assigned to any group. The block of multicast addresses from 224.0.0.1 through 224.0.0.255 is reserved for local wire use. Groups in this range are assigned for various uses, including routing protocols and local discovery mechanisms.

The range from 239.0.0.0 through 239.255.255.255 is reserved for administratively scoped addresses. Because packets addressed to administratively scoped multicast addresses do not cross configured administrative boundaries, and because administratively scoped multicast addresses are locally assigned, these addresses do not need to be unique across administrative boundaries.

## Layer 2 Frames and IPv4 Multicast Addresses

Multicasting on a LAN is a good place to start an investigation of multicasting at Layer 2. At Layer 2, multicast deals with media access control (MAC) frames and addresses instead of IPv4 or IPv6 packets and addresses. Consider a single LAN, without routing devices, with a multicast source sending to a certain group. The rest of the hosts are receivers interested in the multicast group's content. So the multicast source host generates packets with its unicast IP address as the source, and the multicast group address as the destination.

Which MAC addresses are used on the frame containing this packet? The packet source address—the unicast IP address of the host originating the multicast content—translates easily and directly to the MAC address of the source. But what about the packet's destination address? This is the IP multicast group address. Which destination MAC address for the frame corresponds to the packet's multicast group address?

One option is for LANs simply to use the LAN broadcast MAC address, which guarantees that the frame is processed by every station on the LAN. However, this procedure defeats the whole purpose of multicast, which is to limit the circulation of packets and frames to interested hosts. Also, hosts might have access to many multicast groups, which multiplies the amount of traffic to noninterested destinations. Broadcasting frames at the LAN level to support multicast groups makes no sense.

However, there is an easy way to effectively use Layer 2 frames for multicast purposes. The MAC address has a bit that is set to 0 for unicast (the LAN term is *individual address*) and set to 1 to indicate that this is a multicast address. Some of these addresses are reserved for multicast groups of specific vendors or MAC-level protocols. Internet multicast applications use the range 0x01-00-5E-00-00-00 to 0x01-00-5E-FF-FF-FF. Multicast receivers (hosts running TCP/IP) listen for frames with one of these addresses when the application joins a multicast group. The host stops listening when the application terminates or the host leaves the group at the packet layer (Layer 3).

This means that 3 bytes, or 24 bits, are available to map IPv4 multicast addresses at Layer 3 to MAC multicast addresses at Layer 2. However, all IPv4 addresses, including multicast addresses, are 32 bits long, leaving 8 IP address bits left over. Which method of mapping IPv4 multicast addresses to MAC multicast addresses minimizes the chance of “collisions” (that is, two different IP multicast groups at the packet layer mapping to the same MAC multicast address at the frame layer)?

First, it is important to realize that all IPv4 multicast addresses begin with the same 4 bits (**1110**), so there are really only 4 bits of concern, not 8. A LAN must not drop the last bits of the IPv4 address because these are almost guaranteed to be host bits, depending on the subnet mask. But the high-order bits, the leftmost address bits, are almost always network bits, and there is only one LAN (for now).

One other bit of the remaining 24 MAC address bits is reserved (an initial **0** indicates an Internet multicast address), so the 5 bits following the initial **1110** in the IPv4 address are dropped. The 23 remaining bits are mapped, one for one, into the last 23 bits of the MAC address. An example of this process is shown in [Figure 153](#).

Figure 153: Converting MAC Addresses to Multicast Addresses

|   |                                                                                                       |                             |          |           |           |
|---|-------------------------------------------------------------------------------------------------------|-----------------------------|----------|-----------|-----------|
| 1 | IPv4 header multicast destination address                                                             | 232.                        | 224.     | 202.      | 181       |
|   | Written in hexadecimal                                                                                | E8                          | E0       | CA        | B5        |
|   | Written in binary                                                                                     | 1110 1000 1                 | 110 0000 | 1100 1010 | 1011 0101 |
| 2 | Ignore the first 9 bits and copy the remaining 23 bits                                                | X                           | 110 0000 | 1100 1010 | 1011 0101 |
| 3 | First bit X = 0 for Internet; X = 1 for other                                                         | 0                           | 110 0000 | 1100 1010 | 1011 0101 |
| 4 | Written in hexadecimal                                                                                |                             | 60       | CA        | B5        |
| 5 | MAC address in hexadecimal                                                                            | 01 : 00 : 5E : E0 : CA : B5 |          |           |           |
| 6 | Drop last 24 bits                                                                                     | 01 : 00 : 5E :              |          |           |           |
| 7 | Copy the multicast bits                                                                               | 01 : 00 : 5E : 60 : CA : B5 |          |           |           |
| 8 | MAC frame destination address 01:00:5E:60:CA:B5 corresponds to multicast IPv4 address 232.224.202.181 |                             |          |           |           |

658108

Note that this process means that there are 32 ( $2^5$ ) IPv4 multicast addresses that could map to the same MAC multicast addresses. For example, multicast IPv4 addresses 224.8.7.6 and 229.136.7.6 translate to the same MAC address (0x01-00-5E-08-07-06). This is a real concern, and because the host could be interested in frames sent to both of those multicast groups, the IP software must reject one or the other.



**NOTE:** This “collision” problem does not exist in IPv6 because of the way IPv6 handles multicast groups, but it is always a concern in IPv4. The procedure for placing IPv6 multicast packets inside multicast frames is nearly identical to that for IPv4, except for the MAC destination address 0x3333 prefix (and the lack of “collisions”).

Once the MAC address for the multicast group is determined, the host's operating system essentially orders the LAN interface card to join or leave the multicast group. Once joined to a multicast group, the host accepts frames sent to the multicast address as well as the host's unicast address and ignores other multicast group's frames. It is possible for a host to join and receive multicast content from more than one group at the same time, of course.

## Multicast Interface Lists

To avoid multicast routing loops, every multicast routing device must always be aware of the interface that leads to the source of that multicast group content by the shortest path. This is the upstream (incoming) interface, and packets are never to be forwarded back toward a multicast source. All other interfaces are potential downstream (outgoing) interfaces, depending on the number of branches on the distribution tree.

routing devices closely monitor the status of the incoming and outgoing interfaces, a process that determines the *multicast forwarding state*. A routing device with a multicast forwarding state for a particular multicast group is essentially “turned on” for that group's

content. Interfaces on the routing device's outgoing interface list send copies of the group's packets received on the incoming interface list for that group. The incoming and outgoing interface lists might be different for different multicast groups.

The multicast forwarding state in a routing device is usually written in either (S,G) or (\*,G) notation. These are pronounced “ess comma gee” and “star comma gee,” respectively. In (S,G), the S refers to the unicast IP address of the source for the multicast traffic, and the G refers to the particular multicast group IP address for which S is the source. All multicast packets sent from this source have S as the source address and G as the destination address.

The asterisk (\*) in the (\*,G) notation is a wildcard indicating that the state applies to any multicast application source sending to group G. So, if two sources are originating exactly the same content for multicast group 224.1.1.2, a routing device could use (\*,224.1.1.2) to represent the state of a routing device forwarding traffic from both sources to the group.

## Multicast Routing Protocols

Multicast routing protocols enable a collection of multicast routing devices to build (join) distribution trees when a host on a directly attached subnet, typically a LAN, wants to receive traffic from a certain multicast group, prune branches, locate sources and groups, and prevent routing loops.

There are several multicast routing protocols:

- **Distance Vector Multicast Routing Protocol (DVMRP)**—The first of the multicast routing protocols and hampered by a number of limitations that make this method unattractive for large-scale Internet use. DVMRP is a dense-mode-only protocol, and uses the flood-and-prune or implicit join method to deliver traffic everywhere and then determine where the uninterested receivers are. DVMRP uses source-based distribution trees in the form (S,G), and builds its own multicast routing tables for RPF checks.
- **Multicast OSPF (MOSPF)**—Extends OSPF for multicast use, but only for dense mode. However, MOSPF has an explicit join message, so routing devices do not have to flood their entire domain with multicast traffic from every source. MOSPF uses source-based distribution trees in the form (S,G).
- ***Bidirectional PIM mode***—A variation of PIM. Bidirectional PIM builds bidirectional shared trees that are rooted at a rendezvous point (RP) address. Bidirectional traffic does not switch to shortest path trees as in PIM-SM and is therefore optimized for routing state size instead of path length. This means that the end-to-end latency might be longer compared to PIM sparse mode. Bidirectional PIM routes are always wildcard-source (\*,G) routes. The protocol eliminates the need for (S,G) routes and data-triggered events. The bidirectional (\*,G) group trees carry traffic both upstream from senders toward the RP, and downstream from the RP to receivers. As a consequence, the strict reverse path forwarding (RPF)-based rules found in other PIM modes do not apply to bidirectional PIM. Instead, bidirectional PIM (\*,G) routes forward traffic from all sources and the RP. Bidirectional PIM routing devices must have the ability to accept traffic on many potential incoming interfaces. Bidirectional PIM scales well because it needs no source-specific (S,G) state. Bidirectional PIM is recommended in deployments with many dispersed sources and many dispersed receivers.

- *PIM dense mode*—In this mode of PIM, the assumption is that almost all possible subnets have at least one receiver wanting to receive the multicast traffic from a source, so the network is *flooded* with traffic on all possible branches, then pruned back when branches do not express an interest in receiving the packets, explicitly (by message) or implicitly (time-out silence). This is the *dense mode* of multicast operation. LANs are appropriate networks for dense-mode operation. Some multicast routing protocols, especially older ones, support only dense-mode operation, which makes them inappropriate for use on the Internet. In contrast to DVMRP and MOSPF, PIM dense mode allows a routing device to use any unicast routing protocol and performs RPF checks using the unicast routing table. PIM dense mode has an implicit join message, so routing devices use the flood-and-prune method to deliver traffic everywhere and then determine where the uninterested receivers are. PIM dense mode uses source-based distribution trees in the form (S,G), as do all dense-mode protocols. PIM also supports sparse-dense mode, with mixed sparse and dense groups, but there is no special notation for that operational mode. If *sparse-dense mode* is supported, the multicast routing protocol allows some multicast groups to be sparse and other groups to be dense.
- *PIM sparse mode*—In this mode of PIM, the assumption is that very few of the possible receivers want packets from each source, so the network establishes and sends packets only on branches that have at least one leaf indicating (by message) an interest in the traffic. This multicast protocol allows a routing device to use any unicast routing protocol and performs reverse-path forwarding (RPF) checks using the unicast routing table. PIM sparse mode has an *explicit* join message, so routing devices determine where the interested receivers are and send join messages upstream to their neighbors, building trees from receivers to the rendezvous point (RP). PIM sparse mode uses an RP routing device as the initial source of multicast group traffic and therefore builds distribution trees in the form (\*,G), as do all sparse-mode protocols. PIM sparse mode migrates to an (S,G) source-based tree if that path is shorter than through the RP for a particular multicast group's traffic. WANs are appropriate networks for sparse-mode operation, and indeed a common multicast guideline is not to run dense mode on a WAN under any circumstances.
- *Core Based Trees (CBT)*—Shares all of the characteristics of PIM sparse mode (sparse mode, explicit join, and shared (\*,G) trees), but is said to be more efficient at finding sources than PIM sparse mode. CBT is rarely encountered outside academic discussions. There are no large-scale deployments of CBT, commercial or otherwise.
- *PIM source-specific multicast (SSM)*—Enhancement to PIM sparse mode that allows a client to receive multicast traffic directly from the source, without the help of an RP. Used with IGMPv3 to create a shortest-path tree between receiver and source.
- *IGMPv1*—The original protocol defined in RFC 1112, *Host Extensions for IP Multicasting*. IGMPv1 sends an explicit join message to the routing device, but uses a timeout to determine when hosts leave a group. Three versions of the Internet Group Management Protocol (IGMP) run between receiver hosts and routing devices.
- *IGMPv2*—Defined in RFC 2236, *Internet Group Management Protocol, Version 2*. Among other features, IGMPv2 adds an explicit leave message to the join message.
- *IGMPv3*—Defined in RFC 3376, *Internet Group Management Protocol, Version 3*. Among other features, IGMPv3 optimizes support for a single source of content for a multicast

group, or source-specific multicast (SSM). Used with PIM SSM to create a shortest-path tree between receiver and source.

- Bootstrap Router (BSR) and Auto-Rendezvous Point (RP)—Allow sparse-mode routing protocols to find RPs within the routing domain (autonomous system, or AS). RP addresses can also be statically configured.
- Multicast Source Discovery Protocol (MSDP)—Allows groups located in one multicast routing domain to find RPs in other routing domains. MSDP is not used on an RP if all receivers and sources are located in the same routing domain. Typically runs on the same routing device as PIM sparse mode RP. Not appropriate if all receivers and sources are located in the same routing domain.
- Session Announcement Protocol (SAP) and Session Description Protocol (SDP)—Display multicast session names and correlate the names with multicast traffic. SDP is a session directory protocol that advertises multimedia conference sessions and communicates setup information to participants who want to join the session. A client commonly uses SDP to announce a conference session by periodically multicasting an announcement packet to a well-known multicast address and port using SAP.
- Pragmatic General Multicast (PGM)—Special protocol layer for multicast traffic that can be used between the IP layer and the multicast application to add reliability to multicast traffic. PGM allows a receiver to detect missing information in all cases and request replacement information if the receiver application requires it.

The differences among the multicast routing protocols are summarized in [Table 430](#).

**Table 430: Multicast Routing Protocols Compared**

| Multicast Routing Protocol | Dense Mode | Sparse Mode | Implicit Join | Explicit Join | (S,G) SBT  | (*G) Shared Tree |
|----------------------------|------------|-------------|---------------|---------------|------------|------------------|
| DVMRP                      | Yes        | No          | Yes           | No            | Yes        | No               |
| MOSPF                      | Yes        | No          | No            | Yes           | Yes        | No               |
| PIM dense mode             | Yes        | No          | Yes           | No            | Yes        | No               |
| PIM sparse mode            | No         | Yes         | No            | Yes           | Yes, maybe | Yes, initially   |
| Bidirectional PIM          | No         | No          | No            | Yes           | No         | Yes              |
| CBT                        | No         | Yes         | No            | Yes           | No         | Yes              |
| SSM                        | No         | Yes         | No            | Yes           | Yes, maybe | Yes, initially   |
| IGMPv1                     | No         | Yes         | No            | Yes           | Yes, maybe | Yes, initially   |
| IGMPv2                     | No         | Yes         | No            | Yes           | Yes, maybe | Yes, initially   |
| IGMPv3                     | No         | Yes         | No            | Yes           | Yes, maybe | Yes, initially   |



Table 430: Multicast Routing Protocols Compared (*continued*)

| Multicast Routing Protocol | Dense Mode | Sparse Mode | Implicit Join | Explicit Join | (S,G) SBT  | (*G) Shared Tree |
|----------------------------|------------|-------------|---------------|---------------|------------|------------------|
| BSR and Auto-RP            | No         | Yes         | No            | Yes           | Yes, maybe | Yes, initially   |
| MSDP                       | No         | Yes         | No            | Yes           | Yes, maybe | Yes, initially   |

It is important to realize that retransmissions due to a high bit-error rate on a link or overloaded routing device can make multicast as inefficient as repeated unicast. Therefore, there is a trade-off in many multicast applications regarding the session support provided by the Transmission Control Protocol (TCP) (but TCP always resends missing segments), or the simple drop-and-continue strategy of the User Datagram Protocol (UDP) datagram service (but reordering can become an issue). Modern multicast uses UDP almost exclusively.

### T Series Router Multicast Performance

The Juniper Networks T Series Core Routers handle extreme multicast packet replication requirements with a minimum of router load. Each memory component replicates a multicast packet twice at most. Even in the worst-case scenario involving maximum fan-out, when 1 input port and 63 output ports need a copy of the packet, the T Series routing platform copies a multicast packet only six times. Most multicast distribution trees are much sparser, so in many cases only two or three replications are necessary. In no case does the T Series architecture have an impact on multicast performance, even with the largest multicast fan-out requirements.



## PART 74

# Managing Group Membership

- [Using IGMP on page 5405](#)
- [Using IGMP Snooping on page 5431](#)
- [Using MLD on page 5441](#)
- [Using MLD Snooping on page 5467](#)



## CHAPTER 196

# Using IGMP

- [Understanding Group Membership Protocols on page 5405](#)
- [Understanding IGMP on page 5406](#)
- [Configuring IGMP on page 5408](#)
- [Enabling IGMP on page 5409](#)
- [Changing the IGMP Version on page 5410](#)
- [Modifying the IGMP Host-Query Message Interval on page 5411](#)
- [Modifying the IGMP Last-Member Query Interval on page 5412](#)
- [Specifying Immediate-Leave Host Removal for IGMP on page 5413](#)
- [Filtering Unwanted IGMP Reports at the IGMP Interface Level on page 5414](#)
- [Accepting IGMP Messages from Remote Subnetworks on page 5415](#)
- [Modifying the IGMP Query Response Interval on page 5416](#)
- [Modifying the IGMP Robustness Variable on page 5417](#)
- [Limiting the Maximum IGMP Message Rate on page 5418](#)
- [Enabling IGMP Static Group Membership on page 5418](#)
- [Recording IGMP Join and Leave Events on page 5425](#)
- [Limiting the Number of IGMP Multicast Group Joins on Logical Interfaces on page 5426](#)
- [Tracing IGMP Protocol Traffic on page 5427](#)
- [Disabling IGMP on page 5429](#)

## Understanding Group Membership Protocols

---

There is a big difference between the multicast protocols used between host and routing device and between the multicast routing devices themselves. Hosts on a given subnetwork need to inform their routing device only whether or not they are interested in receiving packets from a certain multicast group. The source host needs to inform its routing devices only that it is the source of traffic for a particular multicast group. In other words, no detailed knowledge of the distribution tree is needed by any hosts; only a group membership protocol is needed to inform routing devices of their participation in a multicast group. Between adjacent routing devices, on the other hand, the multicast routing protocols must avoid loops as they build a detailed sense of the network topology

and distribution tree from source to leaf. So, different multicast protocols are used for the host-router portion and the router-router portion of the multicast network.

Multicast group membership protocols enable a routing device to detect when a host on a directly attached subnet, typically a LAN, wants to receive traffic from a certain multicast group. Even if more than one host on the LAN wants to receive traffic for that multicast group, the routing device sends only one copy of each packet for that multicast group out on that interface, because of the inherent broadcast nature of LANs. When the multicast group membership protocol informs the routing device that there are no interested hosts on the subnet, the packets are withheld and that leaf is pruned from the distribution tree.

The Internet Group Management Protocol (IGMP) and the Multicast Listener Discovery (MLD) Protocol are the standard IP multicast group membership protocols: IGMP and MLD have several versions that are supported by hosts and routing devices:

- IGMPv1—The original protocol defined in RFC 1112. An explicit join message is sent to the routing device, but a timeout is used to determine when hosts leave a group. This process wastes processing cycles on the routing device, especially on older or smaller routing devices.
- IGMPv2—Defined in RFC 2236. Among other features, IGMPv2 adds an explicit leave message to the join message so that routing devices can more easily determine when a group has no interested listeners on a LAN.
- IGMPv3—Defined in RFC 3376. Among other features, IGMPv3 optimizes support for a single source of content for a multicast group, or *source-specific multicast (SSM)*.
- MLDv1—Defined in RFC 2710. MLDv1 is similar to IGMPv2.
- MLDv2—Defined in RFC 3810. MLDv2 similar to IGMPv3.

The various versions of IGMP and MLD are backward compatible. It is common for a routing device to run multiple versions of IGMP and MLD on LAN interfaces. Backward compatibility is achieved by dropping back to the most basic of all versions run on a LAN. For example, if one host is running IGMPv1, any routing device attached to the LAN running IGMPv2 can drop back to IGMPv1 operation, effectively eliminating the IGMPv2 advantages. Running multiple IGMP versions ensures that both IGMPv1 and IGMPv2 hosts find peers for their versions on the routing device.

**Related  
Documentation**

- [Examples: Configuring MLD on page 5444](#)

---

## Understanding IGMP

The Internet Group Management Protocol (IGMP) manages the membership of hosts and routing devices in multicast groups. IP hosts use IGMP to report their multicast group memberships to any immediately neighboring multicast routing devices. Multicast routing devices use IGMP to learn, for each of their attached physical networks, which groups have members.

IGMP is also used as the transport for several related multicast protocols (for example, Distance Vector Multicast Routing Protocol [DVMRP] and Protocol Independent Multicast version 1 [PIMv1]).

A routing device receives explicit join and prune messages from those neighboring routing devices that have downstream group members. When PIM is the multicast protocol in use, IGMP begins the process as follows:

1. To join a multicast group, G, a host conveys its membership information through IGMP.
2. The routing device then forwards data packets addressed to a multicast group G to only those interfaces on which explicit join messages have been received.
3. A designated router (DR) sends periodic join and prune messages toward a group-specific rendezvous point (RP) for each group for which it has active members. One or more routing devices are automatically or statically designated as the RP, and all routing devices must explicitly join through the RP.
4. Each routing device along the path toward the RP builds a wildcard (any-source) state for the group and sends join and prune messages toward the RP.

The term *route entry* is used to refer to the state maintained in a routing device to represent the distribution tree.

A route entry can include such fields as:

- source address
- group address
- incoming interface from which packets are accepted
- list of outgoing interfaces to which packets are sent
- timers
- flag bits

The wildcard route entry's incoming interface points toward the RP.

The outgoing interfaces point to the neighboring downstream routing devices that have sent join and prune messages toward the RP as well as the directly connected hosts that have requested membership to group G.

5. This state creates a shared, RP-centered, distribution tree that reaches all group members.

IGMP is an integral part of IP and must be enabled on all routing devices and hosts that need to receive IP multicast traffic.

For each attached network, a multicast routing device can be either a querier or a nonquerier. The querier routing device periodically sends general query messages to solicit group membership information. Hosts on the network that are members of a multicast group send report messages. When a host leaves a group, it sends a leave group message.

IGMP version 3 (IGMPv3) supports inclusion and exclusion lists. Inclusion lists enable you to specify which sources can send to a multicast group. This type of multicast group is called a source-specific multicast (SSM) group, and its multicast address is 232/8.

IGMPv3 provides support for source filtering. For example, a routing device can specify particular routing devices from which it accepts or rejects traffic. With IGMPv3, a multicast routing device can learn which sources are of interest to neighboring routing devices.

Exclusion mode works the opposite of an inclusion list. It allows any source but the ones listed to send to the SSM group.

IGMPv3 interoperates with versions 1 and 2 of the protocol. However, to remain compatible with older IGMP hosts and routing devices, IGMPv3 routing devices must also implement versions 1 and 2 of the protocol. IGMPv3 supports the following membership-report record types: mode is allowed, allow new sources, and block old sources.

**Related  
Documentation**

- *Supported IP Multicast Protocol Standards*
- *Configuring IGMP*

---

## Configuring IGMP

---

Before you begin:

1. Determine whether the router is directly attached to any multicast sources. Receivers must be able to locate these sources.
2. Determine whether the router is directly attached to any multicast group receivers. If receivers are present, IGMP is needed.
3. Determine whether to configure multicast to use sparse, dense, or sparse-dense mode. Each mode has different configuration considerations.
4. Determine the address of the RP if sparse or sparse-dense mode is used.
5. Determine whether to locate the RP with the static configuration, BSR, or auto-RP method.
6. Determine whether to configure multicast to use its own RPF routing table when configuring PIM in sparse, dense, or sparse-dense mode.
7. Configure the SAP and SDP protocols to listen for multicast session announcements. See *Configuring the Session Announcement Protocol*.

To configure the Internet Group Management Protocol (IGMP), include the **igmp** statement:

```
igmp {
 accounting;
 interface interface-name {
 disable;
 (accounting | no-accounting);
 group-policy [policy-names];
 immediate-leave;
 oif-map map-name;
```



```

promiscuous-mode;
ssm-map ssm-map-name;
static {
 group multicast-group-address {
 exclude;
 group-count number;
 group-increment increment;
 source ip-address {
 source-count number;
 source-increment increment;
 }
 }
}
version version;
}
query-interval seconds;
query-last-member-interval seconds;
query-response-interval seconds;
robust-count number;
traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
}
}

```

You can include this statement at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

By default, IGMP is enabled on all interfaces on which you configure Protocol Independent Multicast (PIM), and on all broadcast interfaces on which you configure the Distance Vector Multicast Routing Protocol (DVMRP).



**NOTE:** You can configure IGMP on an interface without configuring PIM. PIM is generally not needed on IGMP downstream interfaces. Therefore, only one “pseudo PIM interface” is created to represent all IGMP downstream (IGMP-only) interfaces on the router. This reduces the amount of router resources, such as memory, that are consumed. You must configure PIM on upstream IGMP interfaces to enable multicast routing, perform reverse-path forwarding for multicast data packets, populate the multicast forwarding table for upstream interfaces, and in the case of bidirectional PIM and PIM sparse mode, to distribute IGMP group memberships into the multicast routing domain.

## Enabling IGMP

The Internet Group Management Protocol (IGMP) manages multicast groups by establishing, maintaining, and removing groups on a subnet. Multicast routing devices use IGMP to learn which groups have members on each of their attached physical

networks. IGMP must be enabled for the router to receive IPv4 multicast packets. IGMP is only needed for IPv4 networks, because multicast is handled differently in IPv6 networks. IGMP is automatically enabled on all IPv4 interfaces on which you configure PIM and on all IPv4 broadcast interfaces when you configure DVMRP.

If IGMP is not running on an interface—either because PIM and DVMRP are not configured on the interface or because IGMP is explicitly disabled on the interface—you can explicitly enable IGMP.

To explicitly enable IGMP:

1. If PIM and DVMRP are not running on the interface, explicitly enable IGMP by including the interface name.

```
[edit protocols igmp]
user@host# set interface fe-0/0/0.0
```

2. See if IGMP is disabled on any interfaces. In the following example, IGMP is disabled on a Gigabit Ethernet interface.

```
[edit protocols igmp]
user@host# show
interface fe-0/0/0.0;
interface ge-1/0/0.0 {
 disable;
}
```

3. Enable IGMP on the interface by deleting the **disable** statement.

```
[edit protocols igmp]
delete interface ge-1/0/0.0 disable
```

4. Verify the configuration.

```
[edit protocols igmp]
user@host# show
interface fe-0/0/0.0;
interface ge-1/0/0.0;
```

5. Verify the operation of IGMP on the interfaces by checking the output of the **show igmp interface** command.

- Related Documentation**
- [Understanding IGMP on page 5406](#)
  - [Disabling IGMP on page 5429](#)
  - [show igmp interface on page 5801](#)

---

## Changing the IGMP Version

By default, the routing device runs IGMPv2. Routing devices running different versions of IGMP determine the lowest common version of IGMP that is supported by hosts on their subnet and operate in that version.

To enable source-specific multicast (SSM) functionality, you must configure version 3 on the host and the host's directly connected routing device. If a source address is specified in a multicast group that is statically configured, the version must be set to IGMPv3.

If a static multicast group is configured with the source address defined, and the IGMP version is configured to be version 2, the source is ignored and only the group is added. In this case, the join is treated as an IGMPv2 group join.

If you configure the IGMP version setting at the individual interface hierarchy level, it overrides the **interface all** statement.

If you have already configured the routing device to use IGMP version 1 (IGMPv1) and then configure it to use IGMPv2, the routing device continues to use IGMPv1 for up to 6 minutes and then uses IGMPv2.

To change to IGMPv3 for SSM functionality:

1. Configure the IGMP interface.

```
[edit protocols igmp]
user@host# set interface ge-0/0/0 version 3
```

2. Verify the configuration by checking the version field in the output of the **show igmp interfaces** command. The **show igmp statistics** command has version-specific output fields, such as V1 Membership Report, V2 Membership Report, and V3 Membership Report.



**CAUTION:** On MX Series platforms, IGMPv2 and IGMPv3 cannot be configured together on the same interface. Configuring both together causes unexpected behavior in multicast traffic forwarding.

#### Related Documentation

- [Understanding IGMP on page 5406](#)
- [show pim interfaces on page 5883](#)
- [show igmp statistics on page 5805](#)
- RFC 2236, *Internet Group Management Protocol, Version 2*
- RFC 3376, *Internet Group Management Protocol, Version 3*

## Modifying the IGMP Host-Query Message Interval

The objective of IGMP is to keep routers up to date with group membership of the entire subnet. Routers need not know who all the members are, only that members exist. Each host keeps track of which multicast groups are subscribed to. On each link, one router is elected the querier. The IGMP querier router periodically sends general host-query messages on each attached network to solicit membership information. The messages are sent to the all-systems multicast group address, 224.0.0.1.

The query interval, the response interval, and the robustness variable are related in that they are all variables that are used to calculate the group membership timeout. The group membership timeout is the number of seconds that must pass before a multicast router determines that no more members of a host group exist on a subnet. The group membership timeout is calculated as the (robustness variable x query-interval) + (query-response-interval). If no reports are received for a particular group before the group membership timeout has expired, the routing device stops forwarding remotely-originated multicast packets for that group onto the attached network.

By default, host-query messages are sent every 125 seconds. You can change this interval to change the number of IGMP messages sent on the subnet.

To modify the query interval:

1. Configure the interval.

```
[edit protocols igmp]
user@host# set query-interval 200
```

The value can be from 1 through 1024 seconds.

2. Verify the configuration by checking the IGMP Query Interval field in the output of the **show igmp interface** command.
3. Verify the operation of the query interval by checking the Membership Query field in the output of the **show igmp statistics** command.

#### Related Documentation

- [Understanding IGMP on page 5406](#)
- [Modifying the IGMP Query Response Interval on page 5416](#)
- [Modifying the IGMP Robustness Variable on page 5417](#)
- [show igmp interface on page 5801](#)
- [show igmp statistics on page 5805](#)

---

## Modifying the IGMP Last-Member Query Interval

The last-member query interval is the maximum amount of time between group-specific query messages, including those sent in response to leave-group messages. You can configure this interval to change the amount of time it takes a routing device to detect the loss of the last member of a group.

When the routing device that is serving as the querier receives a leave-group message from a host, the routing device sends multiple group-specific queries to the group being left. The querier sends a specific number of these queries at a specific interval. The number of queries sent is called the last-member query count. The interval at which the queries are sent is called the last-member query interval. Because both settings are configurable, you can adjust the leave latency. The IGMP leave latency is the time between a request to leave a multicast group and the receipt of the last byte of data for the multicast group.

The last-member query count x (times) the last-member query interval = (equals) the amount of time it takes a routing device to determine that the last member of a group has left the group and to stop forwarding group traffic.

The default last-member query interval is 1 second. You can configure a subsecond interval up to one digit to the right of the decimal point. The configurable range is 0.1 through 0.9, then in 1-second intervals 1 through 999,999.

To modify this interval:

1. Configure the time (in seconds) that the routing device waits for a report in response to a group-specific query.

```
[edit protocols igmp]
user@host# set query-last-member-interval 0.1
```

2. Verify the configuration by checking the IGMP Last Member Query Interval field in the output of the **show igmp interfaces** command.



**NOTE:** You can configure the last-member query count by configuring the robustness variable. The two are always equal.

**Related  
Documentation**

- [Modifying the IGMP Robustness Variable on page 5417](#)
- [show pim interfaces on page 5883](#)

## Specifying Immediate-Leave Host Removal for IGMP

The immediate leave setting is useful for minimizing the leave latency of IGMP memberships. When this setting is enabled, the routing device leaves the multicast group immediately after the last host leaves the multicast group.

The immediate-leave setting enables host tracking, meaning that the device keeps track of the hosts that send join messages. This allows IGMP to determine when the last host sends a leave message for the multicast group.

When the immediate leave setting is enabled, the device removes an interface from the forwarding-table entry without first sending IGMP group-specific queries to the interface. The interface is pruned from the multicast tree for the multicast group specified in the IGMP leave message. The immediate leave setting ensures optimal bandwidth management for hosts on a switched network, even when multiple multicast groups are being used simultaneously.

When immediate leave is disabled and one host sends a leave group message, the routing device first sends a group query to determine if another receiver responds. If no receiver responds, the routing device removes all hosts on the interface from the multicast group. Immediate leave is disabled by default for both IGMP version 2 and IGMP version 3.



**NOTE:** Although host tracking is enabled for IGMPv2 and MLDv1 when you enable immediate leave, use immediate leave with these versions only when there is one host on the interface. The reason is that IGMPv2 and MLDv1 use a report suppression mechanism whereby only one host on an interface sends a group join report in response to a membership query. The other interested hosts suppress their reports. The purpose of this mechanism is to avoid a flood of reports for the same group. But it also interferes with host tracking, because the router only knows about the one interested host and does not know about the others.

To enable immediate leave on an interface:

1. Configure immediate leave on the IGMP interface.

```
[edit protocols IGMP]
user@host# set interface ge-0/0/0.1 immediate-leave
```

2. Verify the configuration by checking the Immediate Leave field in the output of the `show igmp interface` command.

**Related  
Documentation**

- [Understanding IGMP on page 5406](#)
- [show igmp interface on page 5801](#)

---

## Filtering Unwanted IGMP Reports at the IGMP Interface Level

Suppose you need to limit the subnets that can join a certain multicast group. The **group-policy** statement enables you to filter unwanted IGMP reports at the interface level. When this statement is enabled on a router running IGMP version 2 (IGMPv2) or version 3 (IGMPv3), after the router receives an IGMP report, the router compares the group against the specified group policy and performs the action configured in that policy (for example, rejects the report if the policy matches the defined address or network).

You define the policy to match only IGMP group addresses (for IGMPv2) by using the policy's **route-filter** statement to match the group address. You define the policy to match IGMP (source, group) addresses (for IGMPv3) by using the policy's **route-filter** statement to match the group address and the policy's **source-address-filter** statement to match the source address.



**CAUTION:** On MX Series platforms, IGMPv2 and IGMPv3 cannot be configured together on the same interface. Configuring both together causes unexpected behavior in multicast traffic forwarding.

To filter unwanted IGMP reports:

1. Configure an IGMPv2 policy.

```
[edit policy-statement reject_policy_v2]
```

```

user@host# set from route-filter 224.1.1.1/32 exact
user@host# set from route-filter 239.0.0.0/8 orlonger
user@host# set then reject

```

2. Configure an IGMPv3 policy.

```

[edit policy-statement reject_policy_v3]
user@host# set from route-filter 224.1.1.1/32 exact
user@host# set from route-filter 239.0.0.0/8 orlonger
user@host# set from source-address-filter 10.0.0.0/8 orlonger
user@host# set from source-address-filter 127.0.0.0/8 orlonger
user@host# set then reject

```

3. Apply the policies to the IGMP interfaces on which you prefer not to receive specific group or (source, group) reports. In this example, **ge-0/0/0.1** is running IGMPv2, and **ge-0/1/1.0** is running IGMPv3.

```

[edit protocols igmp]
user@host# set interface ge-0/0/0.1 group-policy reject_policy_v2
user@host# set interface ge-0/1/1.0 group-policy reject_policy_v3

```

4. Verify the operation of the filter by checking the Rejected Report field in the output of the **show igmp statistics** command.

#### Related Documentation

- [Understanding IGMP on page 5406](#)
- [Example: Configuring Policy Chains and Route Filters](#)
- [show igmp statistics on page 5805](#)

## Accepting IGMP Messages from Remote Subnetworks

By default, IGMP interfaces accept IGMP messages only from the same subnet. Including the **promiscuous-mode** statement enables the routing device to accept IGMP messages from indirectly connected subnets.



**NOTE:** When you enable IGMP on an unnumbered Ethernet interface that uses a /32 loopback address as a donor address, you must configure IGMP promiscuous mode to accept the IGMP packets received on this interface.



**NOTE:** When enabling promiscuous-mode, all routers on the ethernet segment must be configured with the promiscuous mode statement. Otherwise, only the interface configured with lowest IPv4 address acts as the querier for IGMP for this Ethernet segment.

To enable IGMP promiscuous mode on an interface:

1. Configure the IGMP interface.

```

[edit protocols igmp]
user@host# set interface ge-0/1/1.0 promiscuous-mode

```

2. Verify the configuration by checking the Promiscuous Mode field in the output of the **show igmp interface** command.
3. Verify the operation of the filter by checking the Rx non-local field in the output of the **show igmp statistics** command.

**Related  
Documentation**

- [Understanding IGMP on page 5406](#)
- *Configuring the Loopback Interface* in the *Junos OS Network Interfaces Library for Routing Devices*
- [show igmp interface on page 5801](#)
- [show igmp statistics on page 5805](#)

---

## Modifying the IGMP Query Response Interval

---

The query response interval is the maximum amount of time that can elapse between when the querier router sends a host-query message and when it receives a response from a host. Configuring this interval allows you to adjust the burst peaks of IGMP messages on the subnet. Set a larger interval to make the traffic less bursty. Bursty traffic refers to an uneven pattern of data transmission: sometimes a very high data transmission rate, whereas at other times a very low data transmission rate.

The query response interval, the host-query interval, and the robustness variable are related in that they are all variables that are used to calculate the group membership timeout. The group membership timeout is the number of seconds that must pass before a multicast router determines that no more members of a host group exist on a subnet. The group membership timeout is calculated as the (robustness variable x query-interval) + (query-response-interval). If no reports are received for a particular group before the group membership timeout has expired, the routing device stops forwarding remotely originated multicast packets for that group onto the attached network.

The default query response interval is 10 seconds. You can configure a subsecond interval up to one digit to the right of the decimal point. The configurable range is 0.1 through 0.9, then in 1-second intervals 1 through 999,999.

To modify the query response interval:

1. Configure the interval.  

```
[edit protocols igmp]
user@host# set query-response-interval 0.4
```
2. Verify the configuration by checking the IGMP Query Response Interval field in the output of the **show igmp interface** command.
3. Verify the operation of the query interval by checking the Membership Query field in the output of the **show igmp statistics** command.

**Related  
Documentation**

- [Understanding IGMP on page 5406](#)
- [Modifying the IGMP Host-Query Message Interval on page 5411](#)



- [Modifying the IGMP Robustness Variable on page 5417](#)
- [show igmp interface on page 5801](#)
- [show igmp statistics on page 5805](#)

## Modifying the IGMP Robustness Variable

Fine-tune the IGMP robustness variable to allow for expected packet loss on a subnet. The robust count automatically changes certain IGMP message intervals for IGMPv2 and IGMPv3. Increasing the robust count allows for more packet loss but increases the leave latency of the subnetwork.

When the query router receives an IGMP leave message on a shared network running IGMPv2, the query router must send an IGMP group query message a specified number of times. The number of IGMP group query messages sent is determined by the robust count.

The value of the robustness variable is also used in calculating the following IGMP message intervals:

- Group member interval—Amount of time that must pass before a multicast router determines that there are no more members of a group on a network. This interval is calculated as follows:  $(\text{robustness variable} \times \text{query-interval}) + (1 \times \text{query-response-interval})$ .
- Other querier present interval—The robust count is used to calculate the amount of time that must pass before a multicast router determines that there is no longer another multicast router that is the querier. This interval is calculated as follows:  $(\text{robustness variable} \times \text{query-interval}) + (0.5 \times \text{query-response-interval})$ .
- Last-member query count—Number of group-specific queries sent before the router assumes there are no local members of a group. The number of queries is equal to the value of the robustness variable.

In IGMPv3, a change of interface state causes the system to immediately transmit a state-change report from that interface. In case the state-change report is missed by one or more multicast routers, it is retransmitted. The number of times it is retransmitted is the robust count minus one. In IGMPv3, the robust count is also a factor in determining the group membership interval, the older version querier interval, and the other querier present interval.

By default, the robustness variable is set to 2. You might want to increase this value if you expect a subnet to lose packets.

The number can be from 2 through 10.

To change the value of the robustness variable:

1. Configure the robust count.

When you set the robust count, you are in effect configuring the number of times the querier retries queries on the connected subnets.

```
[edit protocols igmp]
user@host# set robust-count 5
```

2. Verify the configuration by checking the IGMP Robustness Count field in the output of the **show igmp interfaces** command.

**Related Documentation**

- [Modifying the IGMP Host-Query Message Interval on page 5411](#)
- [Modifying the IGMP Query Response Interval on page 5416](#)
- [Modifying the IGMP Last-Member Query Interval on page 5412](#)
- [show pim interfaces on page 5883](#)
- RFC 2236, *Internet Group Management Protocol, Version 2*
- RFC 3376, *Internet Group Management Protocol, Version 3*

---

## Limiting the Maximum IGMP Message Rate

This section describes how to change the limit for the maximum number of IGMP packets transmitted in 1 second by the router.

Increasing the maximum number of IGMP packets transmitted per second might be useful on a router with a large number of interfaces participating in IGMP.

To change the limit for the maximum number of IGMP packets the router can transmit in 1 second, include the **maximum-transmit-rate** statement and specify the maximum number of packets per second to be transmitted.

**Related Documentation**

- [maximum-transmit-rate \(Protocols IGMP\) on page 5628](#)

---

## Enabling IGMP Static Group Membership

You can create IGMP static group membership to test multicast forwarding without a receiver host. When you enable IGMP static group membership, data is forwarded to an interface without that interface receiving membership reports from downstream hosts. The router on which you enable static IGMP group membership must be the designated router (DR) for the subnet. Otherwise, traffic does not flow downstream.

When enabling IGMP static group membership, you cannot configure multiple groups using the **group-count**, **group-increment**, **source-count**, and **source-increment** statements if the **all** option is specified as the IGMP interface.

Class-of-service (CoS) adjustment is not supported with IGMP static group membership.

In this example, you create static group 225.1.1.1.

1. On the DR, configure the static groups to be created by including the **static** statement and **group** statement and specifying which IP multicast address of the group to be created. When creating groups individually, you must specify a unique address for each group.

```
[edit protocols igmp]
user@host# set interface fe-0/1/2 static group 225.1.1.1
```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```
user@host> show configuration protocol igmp

interface fe-0/1/2.0 {
 static {
 group 225.1.1.1;
 }
}
```

3. After you have committed the configuration and the source is sending traffic, use the **show igmp group** command to verify that static group 225.1.1.1 has been created.

```
user@host> show igmp group
Interface: fe-0/1/2
Group: 225.1.1.1
Source: 10.0.0.2
Last reported by: Local
Timeout: 0 Type: Static
```



**NOTE:** When you configure static IGMP group entries on point-to-point links that connect routing devices to a rendezvous point (RP), the static IGMP group entries do not generate join messages toward the RP.

When you create IGMP static group membership to test multicast forwarding on an interface on which you want to receive multicast traffic, you can specify that a number of static groups be automatically created. This is useful when you want to test forwarding to multiple receivers without having to configure each receiver separately.

In this example, you create three groups.

1. On the DR, configure the number of static groups to be created by including the **group-count** statement and specifying the number of groups to be created.

```
[edit protocols igmp]
user@host# set interface fe-0/1/2 static group 225.1.1.1 group-count 3
```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```
user@host> show configuration protocol igmp

interface fe-0/1/2.0 {
 static {
 group 225.1.1.1 {
 group-count 3;
 }
 }
}
```

3. After you have committed the configuration and after the source is sending traffic, use the **show igmp group** command to verify that static groups 225.1.1.1, 225.1.1.2, and 225.1.1.3 have been created.

```
user@host> show igmp group
Interface: fe-0/1/2
 Group: 225.1.1.1
 Source: 10.0.0.2
 Last reported by: Local
 Timeout: 0 Type: Static
 Group: 225.1.1.2
 Source: 10.0.0.2
 Last reported by: Local
 Timeout: 0 Type: Static
 Group: 225.1.1.3
 Source: 10.0.0.2
 Last reported by: Local
 Timeout: 0 Type: Static
```

When you create IGMP static group membership to test multicast forwarding on an interface on which you want to receive multicast traffic, you can also configure the group address to be automatically incremented for each group created. This is useful when you want to test forwarding to multiple receivers without having to configure each receiver separately and when you do not want the group addresses to be sequential.

In this example, you create three groups and increase the group address by an increment of two for each group.

1. On the DR, configure the group address increment by including the **group-increment** statement and specifying the number by which the address should be incremented for each group. The increment is specified in dotted decimal notation similar to an IPv4 address.

```
[edit protocols igmp]
```

```
user@host# set interface fe-0/1/2 static group 225.1.1.1 group-count 3 group-increment 0.0.0.2
```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```
user@host> show configuration protocol igmp

interface fe-0/1/2.0 {
 version 3;
 static {
 group 225.1.1.1 {
 group-increment 0.0.0.2;
 group-count 3;
 }
 }
}
```

3. After you have committed the configuration and after the source is sending traffic, use the **show igmp group** command to verify that static groups 225.1.1.1, 225.1.1.3, and 225.1.1.5 have been created.

```
user@host> show igmp group
```

```

Interface: fe-0/1/2
 Group: 225.1.1.1
 Source: 10.0.0.2
 Last reported by: Local
 Timeout: 0 Type: Static
 Group: 225.1.1.3
 Source: 10.0.0.2
 Last reported by: Local
 Timeout: 0 Type: Static
 Group: 225.1.1.5
 Source: 10.0.0.2
 Last reported by: Local
 Timeout: 0 Type: Static

```

When you create IGMP static group membership to test multicast forwarding on an interface on which you want to receive multicast traffic, and your network is operating in source-specific multicast (SSM) mode, you can also specify that the multicast source address be accepted. This is useful when you want to test forwarding to multicast receivers from a specific multicast source.

If you specify a group address in the SSM range, you must also specify a source.

If a source address is specified in a multicast group that is statically configured, the IGMP version on the interface must be set to IGMPv3. IGMPv2 is the default value.

In this example, you create group 225.1.1.1 and accept IP address 10.0.0.2 as the only source.

1. On the DR, configure the source address by including the **source** statement and specifying the IPv4 address of the source host.

```

[edit protocols igmp]
user@host# set interface fe-0/1/2 static group 225.1.1.1 source 10.0.0.2

```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```

user@host> show configuration protocol igmp

interface fe-0/1/2.0 {
 version 3;
 static {
 group 225.1.1.1 {
 source 10.0.0.2;
 }
 }
}

```

3. After you have committed the configuration and the source is sending traffic, use the **show igmp group** command to verify that static group 225.1.1.1 has been created and that source 10.0.0.2 has been accepted.

```

user@host> show igmp group
Interface: fe-0/1/2
 Group: 225.1.1.1
 Source: 10.0.0.2
 Last reported by: Local
 Timeout: 0 Type: Static

```

When you create IGMP static group membership to test multicast forwarding on an interface on which you want to receive multicast traffic, you can specify that a number of multicast sources be automatically accepted. This is useful when you want to test forwarding to multicast receivers from more than one specified multicast source.

In this example, you create group 255.1.1.1 and accept addresses 10.0.0.2, 10.0.0.3, and 10.0.0.4 as the sources.

1. On the DR, configure the number of multicast source addresses to be accepted by including the **source-count** statement and specifying the number of sources to be accepted.

```
[edit protocols igmp]
user@host# set interface fe-0/1/2 static group 225.1.1.1 source 10.0.0.2 source-count
3
```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```
user@host> show configuration protocol igmp

interface fe-0/1/2.0 {
 version 3;
 static {
 group 225.1.1.1 {
 source 10.0.0.2 {
 source-count 3;
 }
 }
 }
}
```

3. After you have committed the configuration and the source is sending traffic, use the **show igmp group** command to verify that static group 225.1.1.1 has been created and that sources 10.0.0.2, 10.0.0.3, and 10.0.0.4 have been accepted.

```
user@host> show igmp group
Interface: fe-0/1/2
 Group: 225.1.1.1
 Source: 10.0.0.2
 Last reported by: Local
 Timeout: 0 Type: Static
 Group: 225.1.1.1
 Source: 10.0.0.3
 Last reported by: Local
 Timeout: 0 Type: Static
 Group: 225.1.1.1
 Source: 10.0.0.4
 Last reported by: Local
 Timeout: 0 Type: Static
```

When you configure static groups on an interface on which you want to receive multicast traffic, and specify that a number of multicast sources be automatically accepted, you can also specify the number by which the address should be incremented for each source accepted. This is useful when you want to test forwarding to multiple receivers without having to configure each receiver separately and you do not want the source addresses to be sequential.

In this example, you create group 225.1.1.1 and accept addresses 10.0.0.2, 10.0.0.4, and 10.0.0.6 as the sources.

1. Configure the multicast source address increment by including the **source-increment** statement and specifying the number by which the address should be incremented for each source. The increment is specified in dotted decimal notation similar to an IPv4 address.

```
[edit protocols igmp]
```

```
user@host# set interface fe-0/1/2 static group 225.1.1.1 source 10.0.0.2 source-count 3 source-increment 0.0.0.2
```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```
user@host> show configuration protocol igmp
```

```
interface fe-0/1/2.0 {
 version 3;
 static {
 group 225.1.1.1 {
 source 10.0.0.2 {
 source-count 3;
 source-increment 0.0.0.2;
 }
 }
 }
}
```

3. After you have committed the configuration and after the source is sending traffic, use the **show igmp group** command to verify that static group 225.1.1.1 has been created and that sources 10.0.0.2, 10.0.0.4, and 10.0.0.6 have been accepted.

```
user@host> show igmp group
```

```
Interface: fe-0/1/2
 Group: 225.1.1.1
 Source: 10.0.0.2
 Last reported by: Local
 Timeout: 0 Type: Static
 Group: 225.1.1.1
 Source: 10.0.0.4
 Last reported by: Local
 Timeout: 0 Type: Static
 Group: 225.1.1.1
 Source: 10.0.0.6
 Last reported by: Local
 Timeout: 0 Type: Static
```

When you configure static groups on an interface on which you want to receive multicast traffic and your network is operating in source-specific multicast (SSM) mode, you can specify that certain multicast source addresses be excluded.

By default the multicast source address configured in a static group operates in include mode. In include mode the multicast traffic for the group is accepted from the source address configured. You can also configure the static group to operate in exclude mode. In exclude mode the multicast traffic for the group is accepted from any address other than the source address configured.

If a source address is specified in a multicast group that is statically configured, the IGMP version on the interface must be set to IGMPv3. IGMPv2 is the default value.

In this example, you exclude address 10.0.0.2 as a source for group 225.1.1.1.

1. On the DR, configure a multicast static group to operate in exclude mode by including the **exclude** statement and specifying which IPv4 source address to exclude.

```
[edit protocols igmp]
user@host# set interface fe-0/1/2 static group 225.1.1.1 exclude source 10.0.0.2
```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```
user@host> show configuration protocol igmp

interface fe-0/1/2.0 {
 version 3;
 static {
 group 225.1.1.1 {
 exclude;
 source 10.0.0.2;
 }
 }
}
```

3. After you have committed the configuration and the source is sending traffic, use the **show igmp group detail** command to verify that static group 225.1.1.1 has been created and that the static group is operating in exclude mode.

```
user@host> show igmp group detail
Interface: fe-0/1/2
 Group: 225.1.1.1
 Group mode: Exclude
 Source: 10.0.0.2
 Last reported by: Local
 Timeout: 0 Type: Static
```

#### Related Documentation

- [Enabling MLD Static Group Membership on page 5454](#)
- [group \(Protocols IGMP\) on page 5621](#)
- [group-count \(Protocols IGMP\) on page 5622](#)
- [group-increment \(Protocols IGMP\) on page 5622](#)
- [source-count \(Protocols IGMP\) on page 5636](#)



- [source-increment \(Protocols IGMP\) on page 5637](#)
- [static \(Protocols IGMP\) on page 5638](#)

## Recording IGMP Join and Leave Events

To determine whether IGMP tuning is needed in a network, you can configure the routing device to record IGMP join and leave events. You can record events globally for the routing device or for individual interfaces.

[Table 431](#) describes the recordable IGMP events.

**Table 431: IGMP Event Messages**

| ERRMSG Tag                  | Definition                                                     |
|-----------------------------|----------------------------------------------------------------|
| RPD_IGMP_JOIN               | Records IGMP join events.                                      |
| RPD_IGMP_LEAVE              | Records IGMP leave events.                                     |
| RPD_IGMP_ACCOUNTING_ON      | Records when IGMP accounting is enabled on an IGMP interface.  |
| RPD_IGMP_ACCOUNTING_OFF     | Records when IGMP accounting is disabled on an IGMP interface. |
| RPD_IGMP_MEMBERSHIP_TIMEOUT | Records IGMP membership timeout events.                        |

To enable IGMP accounting:

1. Enable accounting globally or on an IGMP interface. This example shows both options.

```
[edit protocols igmp]
user@host# set accounting
user@host# set interface fe-0/1/0.2 accounting
```

2. Configure the events to be recorded and filter the events to a system log file with a descriptive filename, such as `igmp-events`.

```
[edit system syslog file igmp-events]
user@host# set any info
user@host# set match ".*RPD_IGMP_JOIN.* | .*RPD_IGMP_LEAVE.* |
.*RPD_IGMP_ACCOUNTING.* | .*RPD_IGMP_MEMBERSHIP_TIMEOUT.*"
```

3. Periodically archive the log file.

This example rotates the file size when it reaches 100 KB and keeps three files.

```
[edit system syslog file igmp-events]
user@host# set archive size 100000
user@host# set archive files 3
user@host# set archive archive-sites "ftp://user@host1//var/tmp" password
"anonymous"
user@host# set archive archive-sites "ftp://user@host2//var/tmp" password "test"
user@host# set archive transfer-interval 24
user@host# set archive start-time 2011-01-07:12:30
```

4. You can monitor the system log file as entries are added to the file by running the **monitor start** and **monitor stop** commands.

```
user@host> monitor start igmp-events
```

```
*** igmp-events ***
```

```
Apr 16 13:08:23 host mgd[16416]: UI_CMDLINE_READ_LINE: User 'user', command
'run monitor start igmp-events '
monitor
```

**Related  
Documentation**

- [Understanding IGMP on page 5406](#)
- [Specifying Log File Size, Number, and Archiving Properties on page 1543](#)

---

## Limiting the Number of IGMP Multicast Group Joins on Logical Interfaces

The **group-limit** statement enables you to limit the number of IGMP multicast group joins for logical interfaces. When this statement is enabled on a router running IGMP version 2 (IGMPv2) or version 3 (IGMPv3), the limit is applied upon receipt of the group report. Once the group limit is reached, subsequent join requests are rejected.

When configuring limits for IGMP multicast groups, keep the following in mind:

- Each any-source group (\*G) counts as one group toward the limit.
- Each source-specific group (S,G) counts as one group toward the limit.
- Groups in IGMPv3 exclude mode are counted toward the limit.
- Multiple source-specific groups count individually toward the group limit, even if they are for the same group. For example, (S1, G1) and (S2, G1) would count as two groups toward the configured limit.
- Combinations of any-source groups and source-specific groups count individually toward the group limit, even if they are for the same group. For example, (\*, G1) and (S, G1) would count as two groups toward the configured limit.
- Configuring and committing a group limit on a network that is lower than what already exists on the network results in the removal of all groups from the configuration. The groups must then request to rejoin the network (up to the newly configured group limit).
- You can dynamically limit multicast groups on IGMP logical interfaces using dynamic profiles.

Beginning with Junos OS 12.2, you can optionally configure a system log warning threshold for IGMP multicast group joins received on the logical interface. It is helpful to review the system log messages for troubleshooting purposes and to detect if an excessive amount of IGMP multicast group joins have been received on the interface. These log messages convey when the configured group limit has been exceeded, when the configured threshold has been exceeded, and when the number of groups drop below the configured threshold.

The **group-threshold** statement enables you to configure the threshold at which a warning message is logged. The range is 1 through 100 percent. The warning threshold is a

percentage of the group limit, so you must configure the **group-limit** statement to configure a warning threshold. For instance, when the number of groups exceed the configured warning threshold, but remain below the configured group limit, multicast groups continue to be accepted, and the device logs the warning message. In addition, the device logs a warning message after the number of groups drop below the configured warning threshold. You can further specify the amount of time (in seconds) between the log messages by configuring the **log-interval** statement. The range is 6 through 32,767 seconds.

You might consider throttling log messages because every entry added after the configured threshold and every entry rejected after the configured limit causes a warning message to be logged. By configuring a log interval, you can throttle the amount of system log warning messages generated for IGMP multicast group joins.

To limit multicast group joins on an IGMP logical interface:

1. Access the logical interface at the IGMP protocol hierarchy level.

```
[edit]
user@host# edit protocols igmp interface interface-name
```

2. Specify the group limit for the interface.

```
[edit protocols igmp interface interface-name]
user@host# set group-limit limit
```

3. (Optional) Configure the threshold at which a warning message is logged.

```
[edit protocols igmp interface interface-name]
user@host# set group-threshold value
```

4. (Optional) Configure the amount of time between log messages.

```
[edit protocols igmp interface interface-name]
user@host# set log-interval seconds
```

To confirm your configuration, use the **show protocols igmp** command. To verify the operation of IGMP on the interface, including the configured group limit and the optional warning threshold and interval between log messages, use the **show igmp interface** command.

**Related Documentation**

- [Enabling IGMP Static Group Membership on page 5418](#)

## Tracing IGMP Protocol Traffic

Tracing operations record detailed messages about the operation of routing protocols, such as the various types of routing protocol packets sent and received, and routing policy actions. You can specify which trace operations are logged by including specific tracing flags. The following table describes the flags that you can include.

| Flag                       | Description           |
|----------------------------|-----------------------|
| <b>all</b>                 | Trace all operations. |
| <b>client-notification</b> | Trace notifications.  |

| Flag                     | Description                                                                         |
|--------------------------|-------------------------------------------------------------------------------------|
| <b>general</b>           | Trace general flow.                                                                 |
| <b>group</b>             | Trace group operations.                                                             |
| <b>host-notification</b> | Trace host notifications.                                                           |
| <b>leave</b>             | Trace leave group messages (IGMPv2 only).                                           |
| <b>mtrace</b>            | Trace mtrace packets. Use the <b>mtrace</b> command to troubleshoot the software.   |
| <b>normal</b>            | Trace normal events.                                                                |
| <b>packets</b>           | Trace all IGMP packets.                                                             |
| <b>policy</b>            | Trace policy processing.                                                            |
| <b>query</b>             | Trace IGMP membership query messages, including general and group-specific queries. |
| <b>report</b>            | Trace membership report messages.                                                   |
| <b>route</b>             | Trace routing information.                                                          |
| <b>state</b>             | Trace state transitions.                                                            |
| <b>task</b>              | Trace task processing.                                                              |
| <b>timer</b>             | Trace timer processing.                                                             |

In the following example, tracing is enabled for all routing protocol packets. Then tracing is narrowed to focus only on IGMP packets of a particular type. To configure tracing operations for IGMP:

1. (Optional) Configure tracing at the routing options level to trace all protocol packets.

```
[edit routing-options traceoptions]
user@host# set file all-packets-trace
user@host# set flag all
```

2. Configure the filename for the IGMP trace file.

```
[edit protocols igmp traceoptions]
user@host# set file igmp-trace
```

3. (Optional) Configure the maximum number of trace files.

```
[edit protocols igmp traceoptions]
user@host# set file files 5
```

4. (Optional) Configure the maximum size of each trace file.

```
[edit protocols igmp traceoptions]
user@host# set file size 1m
```

5. (Optional) Enable unrestricted file access.

```
[edit protocols igmp traceoptions]
user@host# set file world-readable
```

6. Configure tracing flags. Suppose you are troubleshooting issues with a particular multicast group. The following example shows how to flag all events for packets associated with the group IP address.

```
[edit protocols igmp traceoptions]
user@host# set flag group | match 232.1.1.2
```

7. View the trace file.

```
user@host> file list /var/log
user@host> file show /var/log/igmp-trace
```

- Related Documentation**
- [Understanding IGMP on page 5406](#)
  - *Tracing and Logging Junos OS Operations*
  - [mtrace on page 5839](#)

---

## Disabling IGMP

To disable IGMP on an interface, include the **disable** statement:

```
disable;
```

You can include this statement at the following hierarchy levels:

- [edit protocols igmp interface *interface-name*]
- [edit logical-systems *logical-system-name* protocols igmp interface *interface-name*]

- Related Documentation**
- [Enabling IGMP on page 5409](#)



# Using IGMP Snooping

- [IGMP Snooping Overview on page 5431](#)
- [Configuring IGMP Snooping on page 5434](#)
- [Configuring VLAN-Specific IGMP Snooping Parameters on page 5435](#)
- [Example: Configuring IGMP Snooping on page 5436](#)
- [Monitoring IGMP Snooping on page 5438](#)
- [Verifying the IGMP Snooping Group Timeout Value on page 5439](#)

## IGMP Snooping Overview

---

With IGMP snooping enabled, a switch monitors the IGMP (Internet Group Management Protocol) traffic between hosts and multicast routers and uses what it learns to forward multicast traffic to only the downstream interfaces that are connected to interested receivers. This conserves bandwidth by allowing the switch to send multicast traffic to only those interfaces that are connected to devices that want to receive the traffic (instead of flooding the traffic to all the downstream VLAN interfaces).

This IGMP snooping topic includes:

- [How IGMP Snooping Works on page 5431](#)
- [How IGMP Snooping Works with Routed VLAN Interfaces on page 5432](#)
- [How Hosts Join and Leave Multicast Groups on page 5432](#)
- [IGMP Snooping and Forwarding Interfaces on page 5433](#)
- [General Forwarding Rules on page 5433](#)

## How IGMP Snooping Works

A switch usually learns unicast MAC addresses by checking the source address field of the frames it receives and then sends any traffic for that unicast address only to the appropriate interface. However, a multicast MAC address can never be the source address for a packet. As a result, when a switch receives traffic for a multicast destination address, it floods the traffic on the relevant VLAN, which can cause a significant amount of traffic to be sent unnecessarily.

IGMP snooping prevents this flooding. When you enable IGMP snooping, the switch monitors IGMP packets between receivers and multicast routers and uses the content

of the packets to build a multicast cache table—a database of multicast groups and the interfaces that are connected to members of the groups. When the switch receives multicast packets, it uses the cache table to selectively forward the traffic to only the interfaces that are connected to members of the appropriate multicast groups.



**NOTE:** IGMP snooping is enabled by default on the default VLAN only. With versions of Junos OS for the QFX Series previous to 13.2, IGMP snooping is enabled by default on all VLANs.



**NOTE:** You cannot configure IGMP snooping on a secondary (private) VLAN.

## How IGMP Snooping Works with Routed VLAN Interfaces

A switch can use a routed VLAN interface (RVI) to forward traffic between VLANs that connect to it. IGMP snooping works with Layer 2 interfaces and RVIs to forward multicast traffic in a switched network.

When a switch receives a multicast packet, its Packet Forwarding Engines perform a multicast lookup on the packet to determine how to forward the packet to its local interfaces. From the results of the lookup, each Packet Forwarding Engine extracts a list of Layer 3 interfaces that have ports local to the Packet Forwarding Engine. If the list includes an RVI, the switch provides a bridge multicast group ID for the RVI to the Packet Forwarding Engine.

For VLANs that include multicast receivers, the bridge multicast ID includes a sub-next-hop ID, which identifies the Layer 2 interfaces in the VLAN that are interested in receiving the multicast stream. The Packet Forwarding Engine then forwards multicast traffic to bridge multicast IDs that have multicast receivers for a given multicast group.

## How Hosts Join and Leave Multicast Groups

Hosts can join multicast groups in two ways:

- By sending an unsolicited IGMP join message to a multicast router that specifies the IP multicast group that the host is attempting to join.
- By sending an IGMP join message in response to a general query from a multicast router.

A multicast router continues to forward multicast traffic to a VLAN provided that at least one host on that VLAN responds to the periodic general IGMP queries. For a host to remain a member of a multicast group, therefore, it must continue to respond to the periodic general IGMP queries.

To leave a multicast group, either a host cannot respond to the periodic general IGMP queries, which results in a “silent leave” (the only leave option for IGMPv1), or a host can send a group-specific IGMPv2 leave message.



## IGMP Snooping and Forwarding Interfaces

To determine how to forward multicast traffic, a switch with IGMP snooping enabled maintains information about the following interfaces in its multicast forwarding table:

- Multicast-router interfaces—These interfaces lead toward multicast routers or IGMP queriers.
- Group-member interfaces—These interfaces lead toward hosts that are members of multicast groups.

The switch learns about these interfaces by monitoring IGMP traffic. If an interface receives IGMP queries or Protocol Independent Multicast (PIM) updates, the switch adds the interface to its multicast forwarding table as a multicast-router interface. If an interface receives membership reports for a multicast group, the switch adds the interface to its multicast forwarding table as a group-member interface.

Table entries for interfaces that the switch learns about are subject to aging. For example, if a learned multicast-router interface does not receive IGMP queries or PIM hellos within a certain interval, the switch removes the entry for that interface from its multicast forwarding table.



**NOTE:** For a switch to learn multicast-router interfaces and group-member interfaces, an IGMP querier must exist in the network. This is often a multicast router, but if there is no multicast router on the local network, you can configure the switch itself to be an IGMP querier.

You can statically configure an interface to be a multicast-router interface or a group-member interface. The switch adds a static interface to its multicast forwarding table without having to learn about the interface, and the entry in the table is not subject to aging. You can have a mix of statically configured and dynamically learned interfaces on a switch.

## General Forwarding Rules

Multicast traffic received on a switch interface in a VLAN on which IGMP snooping is enabled is forwarded according to the following rules.

IGMP traffic is forwarded as follows:

- IGMP general queries received on a multicast-router interface are forwarded to all other interfaces in the VLAN.
- IGMP group-specific queries received on a multicast-router interface are forwarded to only those interfaces in the VLAN that are members of the group.
- IGMP reports received on a host interface are forwarded to multicast-router interfaces in the same VLAN, but not to the other host interfaces in the VLAN.

Multicast traffic that is not IGMP traffic is forwarded as follows:

- A multicast packet with a destination address of 224.0.0.0/24 is flooded to all other interfaces on the VLAN.
- An unregistered multicast packet—that is, a packet for a group that has no current members—is forwarded to all multicast-router interfaces in the VLAN.
- A registered multicast packet is forwarded only to those host interfaces in the VLAN that are members of the multicast group and to all multicast-router interfaces in the VLAN.

**Related  
Documentation**

- [Example: Configuring IGMP Snooping on page 5436](#)
- [Configuring IGMP Snooping on page 5434](#)
- [Monitoring IGMP Snooping on page 5438](#)
- [Configuring IGMP on page 5408](#)
- RFC 3171, *IANA Guidelines for IPv4 Multicast Address Assignments*
- IGMPv1—See RFC 1112, *Host extensions for IP multicasting*.
- IGMPv2—See RFC 2236, *Internet Group Management Protocol, Version 2*.
- IGMPv3—See RFC 3376, *Internet Group Management Protocol, Version 3*.

---

## Configuring IGMP Snooping

With IGMP snooping enabled, a switch monitors the IGMP (Internet Group Management Protocol) traffic between hosts and multicast routers and uses what it learns to forward multicast traffic to only the downstream interfaces that are connected to interested receivers. This conserves bandwidth by allowing the switch to send multicast traffic to only those interfaces that are connected to devices that want to receive the traffic (instead of flooding the traffic to all the downstream VLAN interfaces).



**NOTE:** You cannot configure IGMP snooping on a secondary VLAN.

To enable IGMP snooping and configure individual options as needed for your network by using the CLI:

1. Enable IGMP snooping on a VLAN:

```
[edit protocols]
user@switch# set igmp-snooping vlan employee-vlan
```

2. Configure the switch to immediately remove group membership from interfaces on a VLAN when it receives a leave message through that VLAN, and have it not forward any membership queries for the multicast group to the VLAN (IGMPv2 only):

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan-name immediate-leave
```

3. Configure an interface to belong to a multicast group:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan-name interface interface-name static group
group-address
```

4. Configure an interface to forward IGMP queries received from multicast routers.

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan-name interface interface-name
multicast-router-interface
```

5. Configure the switch to wait for four timeout intervals before timing out a multicast group on a VLAN:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan-name robust-count 4
```

6. If you want a standalone switch to act as an IGMP querier, enter the following:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan-name l2-querier source-address source address
```

The switch uses the address that you configure as the source address in the IGMP queries that it sends. If there are any multicast routers on the same local network, make sure the source address for the IGMP querier is greater (a higher number) than the IP addresses for those routers on the network. This ensures that switch is always the IGMP querier on the network.

#### Related Documentation

- [IGMP Snooping Overview on page 5431](#)
- [Example: Configuring IGMP Snooping on page 5436](#)
- [Monitoring IGMP Snooping on page 5438](#)

## Configuring VLAN-Specific IGMP Snooping Parameters

All of the IGMP snooping statements configured with the **igmp-snooping** statement, with the exception of the **traceoptions** statement, can be qualified with the same statement at the VLAN level. To configure IGMP snooping parameters at the VLAN level, include the **vlan** statement:

```
vlan vlan-id;
 immediate-leave;
 interface interface-name {
 group-limit limit;
 host-only-interface;
 multicast-router-interface;
 static {
 group ip-address {
 source ip-address;
 }
 }
 }
 proxy {
 source-address ip-address;
 }
 query-interval seconds;
 query-last-member-interval seconds;
 query-response-interval seconds;
 robust-count number;
```

```
}
```

You can include this statement at the following hierarchy levels:

- [edit bridge-domains *bridge-domain-name* protocols igmp-snooping]
- [edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping]

#### Related Documentation

- [Layer 2 Frames and IPv4 Multicast Addresses on page 5395](#)
- [Understanding Multicast Snooping](#)

## Example: Configuring IGMP Snooping

With IGMP snooping enabled, a switch monitors the IGMP (Internet Group Management Protocol) traffic between hosts and multicast routers and uses what it learns to forward multicast traffic to only the downstream interfaces that are connected to interested receivers. This conserves bandwidth by allowing the switch to send multicast traffic to only those interfaces that are connected to devices that want to receive the traffic (instead of flooding the traffic to all the downstream VLAN interfaces).

This example describes how to configure IGMP snooping:

- [Requirements on page 5436](#)
- [Overview and Topology on page 5436](#)
- [Configuration on page 5437](#)

### Requirements

This example requires Junos OS Release 11.1 or later on a QFX Series product.

Before you configure IGMP snooping, be sure you have:

- Configured the **employee-vlan** VLAN
- Assigned interfaces **ge-0/0/1**, **ge-0/0/2**, and **ge-0/0/3** to **employee-vlan**

### Overview and Topology

In this example you configure an interface to receive multicast traffic from a source and configure some multicast-related behavior for downstream interfaces. The example assumes that IGMP snooping was previously disabled for the VLAN.

[Table 432](#) shows the components of the topology for this example.

**Table 432: Components of the IGMP Snooping Topology**

| Components                         | Settings                                            |
|------------------------------------|-----------------------------------------------------|
| VLAN name                          | <b>employee-vlan</b> , tag 20                       |
| Interfaces in <b>employee-vlan</b> | <b>ge-0/0/1</b> , <b>ge-0/0/2</b> , <b>ge-0/0/3</b> |

Table 432: Components of the IGMP Snooping Topology (*continued*)

| Components                                    | Settings        |
|-----------------------------------------------|-----------------|
| Multicast IP address for <b>employee-vlan</b> | 225.100.100.100 |

## Configuration

To configure basic IGMP snooping on a switch:

### CLI Quick Configuration

To quickly configure IGMP snooping, copy the following commands and paste them into a terminal window:

```
[edit protocols]
set igmp-snooping vlan employee-vlan
set igmp-snooping vlan employee-vlan interface ge-0/0/3 static group 225.100.100.100
set igmp-snooping vlan employee-vlan interface ge-0/0/2 multicast-router-interface
set igmp-snooping vlan employee-vlan robust-count 4
```

### Step-by-Step Procedure

Configure IGMP snooping:

1. Enable and configure IGMP snooping on the VLAN **employee-vlan**:

```
[edit protocols]
user@switch# set igmp-snooping vlan employee-vlan
```

2. Configure an interface to belong to a multicast group:

```
[edit protocols]
user@switch# set igmp-snooping vlan employee-vlan interface ge-0/0/3 static group
225.100.100.100
```

3. Configure an interface to forward IGMP queries received from multicast routers.

```
[edit protocols]
user@switch# set igmp-snooping vlan employee-vlan interface ge-0/0/2
multicast-router-interface
```

4. Configure the switch to wait for four timeout intervals before timing out a multicast group on a VLAN:

```
[edit protocols]
user@switch# set igmp-snooping vlan employee-vlan robust-count 4
```

### Results

Check the results of the configuration:

```
user@switch# show protocols igmp-snooping
vlan employee-vlan {
 robust-count 4;
}
interface ge-0/0/2 {
 multicast-router-interface;
}
interface ge-0/0/3 {
 static {
 group 225.100.100.100;
 }
}
```

- Related Documentation**
- [IGMP Snooping Overview on page 5431](#)
  - [Configuring IGMP Snooping on page 5434](#)
  - [Changing the IGMP Snooping Group Timeout Value](#)
  - [Monitoring IGMP Snooping on page 5438](#)
  - [Example: Setting Up Bridging with Multiple VLANs.](#)

## Monitoring IGMP Snooping

**Purpose** Use the monitoring feature to view status and information about the IGMP snooping configuration.

**Action** To display IGMP snooping details in the CLI, enter the following commands:

- **show igmp-snooping vlans**
- **show igmp-snooping statistics**
- **show igmp-snooping route**
- **show igmp-snooping membership**

**Meaning** [Table 433](#) summarizes the IGMP snooping details displayed.

**Table 433: Summary of IGMP Snooping Output Fields**

| Field                  | Values                                                             |
|------------------------|--------------------------------------------------------------------|
| IGMP Snooping Monitor  |                                                                    |
| VLAN                   | VLAN for which IGMP snooping is enabled.                           |
| Interfaces             | Interface connected to a multicast router.                         |
| Groups                 | Number of the multicast groups learned by the VLAN.                |
| MRouters               | Multicast router.                                                  |
| Receivers              | Multicast receiver.                                                |
| IGMP Route Information |                                                                    |
| VLAN                   | VLAN for which IGMP snooping is enabled.                           |
| Next-Hop               | Next hop assigned by the switch after performing the route lookup. |
| Group                  | Multicast groups learned by the VLAN.                              |

- Related Documentation**
- [IGMP Snooping Overview on page 5431](#)

- [Example: Configuring IGMP Snooping on page 5436](#)
- [Configuring IGMP Snooping on page 5434](#)
- [Changing the IGMP Snooping Group Timeout Value](#)

---

## Verifying the IGMP Snooping Group Timeout Value

---

**Purpose** Verify that the IGMP snooping group timeout value has been changed correctly from its default value.

**Action** Display the IGMP snooping membership information, which contains the group timeout value that was derived from the IGMP configuration:

```
user@switch> show igmp-snooping membership detail
VLAN: v43 Tag: 43 (Index: 4)
Group: 225.0.0.1
Receiver count: 1, Flags: <v2-hosts>
ge-0/0/15.0 Uptime: 00:00:05 timeout: 510
```

**Meaning** The IGMP snooping group timeout value determines how long a switch waits to receive an IGMP query from a multicast router before removing a multicast group from its multicast cache table. When you enable IGMP snooping, the default IGMP snooping group timeout value of 260 seconds is applied to all VLANs, which means that the switch waits 260 seconds to receive an IGMP query before removing a multicast group from its multicast cache table. You can change the timeout value by using the **robust-count** option.

**Related Documentation**

- [Changing the IGMP Snooping Group Timeout Value](#)





# Using MLD

- [Understanding MLD on page 5441](#)
- [Examples: Configuring MLD on page 5444](#)

## Understanding MLD

---

The Multicast Listener Discovery (MLD) Protocol manages the membership of hosts and routers in multicast groups. IP version 6 (IPv6) multicast routers use MLD to learn, for each of their attached physical networks, which groups have interested listeners. Each routing device maintains a list of host multicast addresses that have listeners for each subnetwork, as well as a timer for each address. However, the routing device does not need to know the address of each listener—just the address of each host. The routing device provides addresses to the multicast routing protocol it uses, which ensures that multicast packets are delivered to all subnetworks where there are interested listeners. In this way, MLD is used as the transport for the Protocol Independent Multicast (PIM) Protocol.


MLD is an integral part of IPv6 and must be enabled on all IPv6 routing devices and hosts that need to receive IP multicast traffic. The Junos OS supports MLD versions 1 and 2. Version 2 is supported for source-specific multicast (SSM) include and exclude modes.

In include mode, the receiver specifies the source or sources it is interested in receiving the multicast group traffic from. Exclude mode works the opposite of include mode. It allows the receiver to specify the source or sources it is not interested in receiving the multicast group traffic from.

For each attached network, a multicast routing device can be either a querier or a nonquerier. A querier routing device, usually one per subnet, solicits group membership information by transmitting MLD queries. When a host reports to the querier routing device that it has interested listeners, the querier routing device forwards the membership information to the rendezvous point (RP) routing device by means of the receiver's (host's) designated router (DR). This builds the rendezvous-point tree (RPT) connecting the host with interested listeners to the RP routing device. The RPT is the initial path used by the sender to transmit information to the interested listeners. Nonquerier routing devices do not transmit MLD queries on a subnet but can do so if the querier routing device fails.

All MLD-configured routing devices start as querier routing devices on each attached subnet (see [Figure 154](#)). The querier routing device on the right is the receiver's DR.

Figure 154: Routing Devices Start Up on a Subnet




To elect the querier routing device, the routing devices exchange query messages containing their IPv6 source addresses. If a routing device hears a query message whose IPv6 source address is numerically lower than its own selected address, it becomes a nonquerier. In Figure 155, the routing device on the left has a source address numerically lower than the one on the right and therefore becomes the querier routing device.




**NOTE:** In the practical application of MLD, several routing devices on a subnet are nonqueriers. If the elected querier routing device fails, query messages are exchanged among the remaining routing devices. The routing device with the lowest IPv6 source address becomes the new querier routing device. The IPv6 Neighbor Discovery Protocol (NDP) implementation drops incoming Neighbor Announcement (NA) messages that have a broadcast or multicast address in the target link-layer address option. This behavior is recommended by RFC 2461.

Figure 155: Querier Routing Device Is Determined




The querier routing device sends general MLD queries on the **link-scope all-nodes** multicast address FF02::1 at short intervals to all attached subnets to solicit group membership information (see Figure 156). Within the query message is the *maximum response delay* value, specifying the maximum allowed delay for the host to respond with a report message.

Figure 156: General Query Message Is Issued




If interested listeners are attached to the host receiving the query, the host sends a report containing the host's IPv6 address to the routing device (see [Figure 157](#)). If the reported address is not yet in the routing device's list of multicast addresses with interested listeners, the address is added to the list and a timer is set for the address. If the address is already on the list, the timer is reset. The host's address is transmitted to the RP in the PIM domain.

**Figure 157: Reports Are Received by the Querier Routing Device**




If the host has no interested multicast listeners, it sends a done message to the querier routing device. On receipt, the querier routing device issues a multicast address-specific query containing the last **listener query interval** value to the multicast address of the host. If the routing device does not receive a report from the multicast address, it removes the multicast address from the list and notifies the RP in the PIM domain of its removal (see [Figure 158](#)).

**Figure 158: Host Has No Interested Receivers and Sends a Done Message to Routing Device**



If a done message is not received by the querier routing device, the querier routing device continues to send multicast address-specific queries. If the timer set for the address on receipt of the last report expires, the querier routing device assumes there are no longer interested listeners on that subnet, removes the multicast address from the list, and notifies the RP in the PIM domain of its removal (see [Figure 159](#)).

**Figure 159: Host Address Timer Expires and Address Is Removed from Multicast Address List**



- Related Documentation**
- [Enabling MLD on page 5448](#)
  - [Example: Recording MLD Join and Leave Events on page 5462](#)

- [Example: Modifying the MLD Robustness Variable on page 5453](#)

## Examples: Configuring MLD

---

- [Understanding MLD on page 5444](#)
- [Configuring MLD on page 5447](#)
- [Enabling MLD on page 5448](#)
- [Modifying the MLD Version on page 5449](#)
- [Modifying the MLD Host-Query Message Interval on page 5449](#)
- [Modifying the MLD Query Response Interval on page 5450](#)
- [Modifying the MLD Last-Member Query Interval on page 5450](#)
- [Specifying Immediate-Leave Host Removal for MLD on page 5451](#)
- [Filtering Unwanted MLD Reports at the MLD Interface Level on page 5452](#)
- [Example: Modifying the MLD Robustness Variable on page 5453](#)
- [Limiting the Maximum MLD Message Rate on page 5454](#)
- [Enabling MLD Static Group Membership on page 5454](#)
- [Example: Recording MLD Join and Leave Events on page 5462](#)
- [Configuring the Number of MLD Multicast Group Joins on Logical Interfaces on page 5464](#)
- [Disabling MLD on page 5466](#)

## Understanding MLD

The Multicast Listener Discovery (MLD) Protocol manages the membership of hosts and routers in multicast groups. IP version 6 (IPv6) multicast routers use MLD to learn, for each of their attached physical networks, which groups have interested listeners. Each routing device maintains a list of host multicast addresses that have listeners for each subnetwork, as well as a timer for each address. However, the routing device does not need to know the address of each listener—just the address of each host. The routing device provides addresses to the multicast routing protocol it uses, which ensures that multicast packets are delivered to all subnetworks where there are interested listeners. In this way, MLD is used as the transport for the Protocol Independent Multicast (PIM) Protocol.

MLD is an integral part of IPv6 and must be enabled on all IPv6 routing devices and hosts that need to receive IP multicast traffic. The Junos OS supports MLD versions 1 and 2. Version 2 is supported for source-specific multicast (SSM) include and exclude modes.


In include mode, the receiver specifies the source or sources it is interested in receiving the multicast group traffic from. Exclude mode works the opposite of include mode. It allows the receiver to specify the source or sources it is not interested in receiving the multicast group traffic from.

For each attached network, a multicast routing device can be either a querier or a nonquerier. A querier routing device, usually one per subnet, solicits group membership information by transmitting MLD queries. When a host reports to the querier routing

device that it has interested listeners, the querier routing device forwards the membership information to the rendezvous point (RP) routing device by means of the receiver's (host's) designated router (DR). This builds the rendezvous-point tree (RPT) connecting the host with interested listeners to the RP routing device. The RPT is the initial path used by the sender to transmit information to the interested listeners. Nonquerier routing devices do not transmit MLD queries on a subnet but can do so if the querier routing device fails.

All MLD-configured routing devices start as querier routing devices on each attached subnet (see Figure 154). The querier routing device on the right is the receiver's DR.

**Figure 160: Routing Devices Start Up on a Subnet**




To elect the querier routing device, the routing devices exchange query messages containing their IPv6 source addresses. If a routing device hears a query message whose IPv6 source address is numerically lower than its own selected address, it becomes a nonquerier. In Figure 155, the routing device on the left has a source address numerically lower than the one on the right and therefore becomes the querier routing device.




**NOTE:** In the practical application of MLD, several routing devices on a subnet are nonqueriers. If the elected querier routing device fails, query messages are exchanged among the remaining routing devices. The routing device with the lowest IPv6 source address becomes the new querier routing device. The IPv6 Neighbor Discovery Protocol (NDP) implementation drops incoming Neighbor Announcement (NA) messages that have a broadcast or multicast address in the target link-layer address option. This behavior is recommended by RFC 2461.

**Figure 161: Querier Routing Device Is Determined**




The querier routing device sends general MLD queries on the **link-scope all-nodes** multicast address FF02::1 at short intervals to all attached subnets to solicit group membership information (see Figure 156). Within the query message is the *maximum response delay* value, specifying the maximum allowed delay for the host to respond with a report message.

Figure 162: General Query Message Is Issued




If interested listeners are attached to the host receiving the query, the host sends a report containing the host's IPv6 address to the routing device (see [Figure 157](#)). If the reported address is not yet in the routing device's list of multicast addresses with interested listeners, the address is added to the list and a timer is set for the address. If the address is already on the list, the timer is reset. The host's address is transmitted to the RP in the PIM domain.

Figure 163: Reports Are Received by the Querier Routing Device




If the host has no interested multicast listeners, it sends a done message to the querier routing device. On receipt, the querier routing device issues a multicast address-specific query containing the last **listener query interval** value to the multicast address of the host. If the routing device does not receive a report from the multicast address, it removes the multicast address from the list and notifies the RP in the PIM domain of its removal (see [Figure 158](#)).

Figure 164: Host Has No Interested Receivers and Sends a Done Message to Routing Device



If a done message is not received by the querier routing device, the querier routing device continues to send multicast address-specific queries. If the timer set for the address on receipt of the last report expires, the querier routing device assumes there are no longer interested listeners on that subnet, removes the multicast address from the list, and notifies the RP in the PIM domain of its removal (see [Figure 159](#)).

Figure 165: Host Address Timer Expires and Address Is Removed from Multicast Address List



## Configuring MLD

To configure the Multicast Listener Discovery (MLD) Protocol, include the **mld** statement:

```
mld {
 accounting;
 interface interface-name {
 disable;
 (accounting | no-accounting);
 group-policy [policy-names];
 immediate-leave;
 oif-map [map-names];
 passive;
 ssm-map ssm-map-name;
 static (Protocols MLD) {
 group mcast-group-address {
 exclude;
 group-count number;
 group-increment increment;
 source ip-address {
 source-count number;
 source-increment increment;
 }
 }
 }
 version version;
 }
 maximum-transmit-rate packets-per-second;
 query-interval seconds;
 query-last-member-interval seconds;
 query-response-interval seconds;
 robust-count number;
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

By default, MLD is enabled on all broadcast interfaces when you configure Protocol Independent Multicast (PIM) or the Distance Vector Multicast Routing Protocol (DVMRP).

## Enabling MLD

The Multicast Listener Discovery (MLD) Protocol manages multicast groups by establishing, maintaining, and removing groups on a subnet. Multicast routing devices use MLD to learn which groups have members on each of their attached physical networks. MLD must be enabled for the router to receive IPv6 multicast packets. MLD is only needed for IPv6 networks, because multicast is handled differently in IPv4 networks. MLD is enabled on all IPv6 interfaces on which you configure PIM and on all IPv6 broadcast interfaces when you configure DVMRP.

MLD specifies different behaviors for multicast listeners and for routers. When a router is also a listener, the router responds to its own messages. If a router has more than one interface to the same link, it needs to perform the router behavior over only one of those interfaces. Listeners, on the other hand, must perform the listener behavior on all interfaces connected to potential receivers of multicast traffic.

If MLD is not running on an interface—either because PIM and DVMRP are not configured on the interface or because MLD is explicitly disabled on the interface—you can explicitly enable MLD.

To explicitly enable MLD:

1. If PIM and DVMRP are not running on the interface, explicitly enable MLD by including the interface name.

```
[edit protocols mld]
user@host# set interface fe-0/0/0.0
```

2. Check to see if MLD is disabled on any interfaces. In the following example, MLD is disabled on a Gigabit Ethernet interface.

```
[edit protocols mld]
user@host# show

interface fe-0/0/0.0;
interface ge-0/0/0.0 {
 disable;
}
```

3. Enable MLD on the interface by deleting the **disable** statement.

```
[edit protocols mld]
delete interface ge-0/0/0.0 disable
```

4. Verify the configuration.

```
[edit protocols mld]
user@host# show

interface fe-0/0/0.0;
interface ge-0/0/0.0;
```

5. Verify the operation of MLD by checking the output of the **show mld interface** command.



## Modifying the MLD Version

By default, the router supports MLD version 1 (MLDv1). To enable the router to use MLD version 2 (MLDv2) for source-specific multicast (SSM) only, include the **version 2** statement.

If you configure the MLD version setting at the individual interface hierarchy level, it overrides configuring the IGMP version using the **interface all** statement.

If a source address is specified in a multicast group that is statically configured, the version must be set to MLDv2.

To change an MLD interface to version 2:

1. Configure the MLD interface.

```
[edit protocols mld]
user@host# set interface fe-0/0/0.0 version 2
```

2. Verify the configuration by checking the **version** field in the output of the **show mld interface** command. The **show mld statistics** command has version-specific output fields, such as the counters in the **MLD Message type** field.

## Modifying the MLD Host-Query Message Interval

The objective of MLD is to keep routers up to date with IPv6 group membership of the entire subnet. Routers need not know who all the members are, only that members exist. Each host keeps track of which multicast groups are subscribed to. On each link, one router is elected the querier. The MLD querier router periodically sends general host-query messages on each attached network to solicit membership information. These messages solicit group membership information and are sent to the **link-scope all-nodes** address **FF02::1**. A general host-query message has a maximum response time that you can set by configuring the query response interval.

The query response timeout, the query interval, and the robustness variable are related in that they are all variables that are used to calculate the multicast listener interval. The multicast listener interval is the number of seconds that must pass before a multicast router determines that no more members of a host group exist on a subnet. The multicast listener interval is calculated as the (robustness variable x query-interval) + (1 x query-response-interval). If no reports are received for a particular group before the multicast listener interval has expired, the routing device stops forwarding remotely-originated multicast packets for that group onto the attached network.

By default, host-query messages are sent every 125 seconds. You can change this interval to change the number of MLD messages sent on the subnet.

To modify the query interval:

1. Configure the interval.

```
[edit protocols mld]
user@host# set query-interval 200
```

The value can be from 1 through 1024 seconds.

2. Verify the configuration by checking the **MLD Query Interval** field in the output of the **show mld interface** command.
3. Verify the operation of the query interval by checking the **Listener Query** field in the output of the **show mld statistics** command.

## Modifying the MLD Query Response Interval

The query response interval is the maximum amount of time that can elapse between when the querier router sends a host-query message and when it receives a response from a host. You can change this interval to adjust the burst peaks of MLD messages on the subnet. Set a larger interval to make the traffic less bursty.

The query response timeout, the query interval, and the robustness variable are related in that they are all variables that are used to calculate the multicast listener interval. The multicast listener interval is the number of seconds that must pass before a multicast router determines that no more members of a host group exist on a subnet. The multicast listener interval is calculated as the (robustness variable x query-interval) + (1 x query-response-interval). If no reports are received for a particular group before the multicast listener interval has expired, the routing device stops forwarding remotely-originated multicast packets for that group onto the attached network.

The default query response interval is 10 seconds. You can configure a subsecond interval up to one digit to the right of the decimal point. The configurable range is 0.1 through 0.9, then in 1-second intervals 1 through 999,999.

To modify the query response interval:

1. Configure the interval.  

```
[edit protocols mld]
user@host# set query-response-interval 0.5
```
2. Verify the configuration by checking the **MLD Query Response Interval** field in the output of the **show mld interface** command.
3. Verify the operation of the query interval by checking the **Listener Query** field in the output of the **show mld statistics** command.

## Modifying the MLD Last-Member Query Interval

The last-member query interval (also called the last-listener query interval) is the maximum amount of time between group-specific query messages, including those sent in response to done messages sent on the **link-scope-all-routers** address FF02::2. You can lower this interval to reduce the amount of time it takes a router to detect the loss of the last member of a group.

When the routing device that is serving as the querier receives a leave-group (done) message from a host, the routing device sends multiple group-specific queries to the group. The querier sends a specific number of these queries, and it sends them at a specific interval. The number of queries sent is called the last-listener query count. The

interval at which the queries are sent is called the last-listener query interval. Both settings are configurable, thus allowing you to adjust the leave latency. The IGMP leave latency is the time between a request to leave a multicast group and the receipt of the last byte of data for the multicast group.

The last-listener query count  $\times$  (times) the last-listener query interval = (equals) the amount of time it takes a routing device to determine that the last member of a group has left the group and to stop forwarding group traffic.

The default last-listener query interval is 1 second. You can configure a subsecond interval up to one digit to the right of the decimal point. The configurable range is 0.1 through 0.9, then in 1-second intervals 1 through 999,999.

To modify this interval:

1. Configure the time (in seconds) that the routing device waits for a report in response to a group-specific query.

```
[edit protocols mld]
```

```
user@host# set query-last-member-interval 0.1
```

2. Verify the configuration by checking the **MLD Last Member Query Interval** field in the output of the **show igmp interfaces** command.



**NOTE:** You can configure the last-member query count by configuring the robustness variable. The two are always equal.

## Specifying Immediate-Leave Host Removal for MLD

The immediate leave setting is useful for minimizing the leave latency of MLD memberships. When this setting is enabled, the routing device leaves the multicast group immediately after the last host leaves the multicast group.

The immediate-leave setting enables host tracking, meaning that the device keeps track of the hosts that send join messages. This allows MLD to determine when the last host sends a leave message for the multicast group.

When the immediate leave setting is enabled, the device removes an interface from the forwarding-table entry without first sending MLD group-specific queries to the interface. The interface is pruned from the multicast tree for the multicast group specified in the MLD leave message. The immediate leave setting ensures optimal bandwidth management for hosts on a switched network, even when multiple multicast groups are being used simultaneously.

When immediate leave is disabled and one host sends a leave group message, the routing device first sends a group query to determine if another receiver responds. If no receiver responds, the routing device removes all hosts on the interface from the multicast group. Immediate leave is disabled by default for both MLD version 1 and MLD version 2.



NOTE: Although host tracking is enabled for IGMPv2 and MLDv1 when you enable immediate leave, use immediate leave with these versions only when there is one host on the interface. The reason is that IGMPv2 and MLDv1 use a report suppression mechanism whereby only one host on an interface sends a group join report in response to a membership query. The other interested hosts suppress their reports. The purpose of this mechanism is to avoid a flood of reports for the same group. But it also interferes with host tracking, because the router only knows about the one interested host and does not know about the others.

To enable immediate leave:

1. Configure immediate leave on the MLD interface.

```
[edit protocols mld]
user@host# set interface ge-0/0/0.1 immediate-leave
```

2. Verify the configuration by checking the **Immediate Leave** field in the output of the `show mld interface` command.

## Filtering Unwanted MLD Reports at the MLD Interface Level

Suppose you need to limit the subnets that can join a certain multicast group. The **group-policy** statement enables you to filter unwanted MLD reports at the interface level.

When the **group-policy** statement is enabled on a router, after the router receives an MLD report, the router compares the group against the specified group policy and performs the action configured in that policy (for example, rejects the report if the policy matches the defined address or network).

You define the policy to match only MLD group addresses (for MLDv1) by using the policy's **route-filter** statement to match the group address. You define the policy to match MLD (source, group) addresses (for MLDv2) by using the policy's **route-filter** statement to match the group address and the policy's **source-address-filter** statement to match the source address.

To filter unwanted MLD reports:

1. Configure an MLDv1 policy.

```
[edit policy-statement reject_policy_v1]
user@host# set from route-filter fec0:1:1:4::/64 exact
user@host# set then reject
```

2. Configure an MLDv2 policy.

```
[edit policy-statement reject_policy_v2]
user@host# set from route-filter fec0:1:1:4::/64 exact
user@host# set from source-address-filter fe80::2e0:81ff:fe05:1a8d/32 orlonger
user@host# set then reject
```

3. Apply the policies to the MLD interfaces where you prefer not to receive specific group or (source, group) reports. In this example, **ge-0/0/0.1** is running MLDv1 and **ge-0/1/1.0** is running MLDv2.

```
[edit protocols mld]
user@host# set interface ge-0/0/0.1 group-policy reject_policy_v1
user@host# set interface ge-0/1/1.0 group-policy reject_policy_v2
```

4. Verify the operation of the filter by checking the **Rejected Report** field in the output of the **show mld statistics** command.

## Example: Modifying the MLD Robustness Variable

This example shows how to configure and verify the MLD robustness variable in a multicast domain.

- [Requirements on page 5453](#)
- [Overview on page 5453](#)
- [Configuration on page 5454](#)
- [Verification on page 5454](#)

### Requirements

Before you begin:

- Configure the router interfaces. See the *Junos OS Network Interfaces Library for Routing Devices*.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Routing Devices*.
- Enable IPv6 unicast routing. See the *Junos OS Routing Protocols Library for Routing Devices*.
- Enable PIM. See [“PIM Overview” on page 5491](#).

### Overview

The MLD robustness variable can be fine-tuned to allow for expected packet loss on a subnet. Increasing the robust count allows for more packet loss but increases the leave latency of the subnetwork.

The value of the robustness variable is used in calculating the following MLD message intervals:

- Group member interval—Amount of time that must pass before a multicast router determines that there are no more members of a group on a network. This interval is calculated as follows: (robustness variable x query-interval) + (1 x query-response-interval).
- Other querier present interval—Amount of time that must pass before a multicast router determines that there is no longer another multicast router that is the querier. This interval is calculated as follows: (robustness variable x query-interval) + (0.5 x query-response-interval).

- Last-member query count—Number of group-specific queries sent before the router assumes there are no local members of a group. The default number is the value of the robustness variable.

By default, the robustness variable is set to 2. The number can be from 2 through 10. You might want to increase this value if you expect a subnet to lose packets.

---

### Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set protocols mld robust-count 5
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To change the value of the robustness variable:

1. Configure the robust count.

```
[edit protocols mld]
user@host# set robust-count 5
```

2. If you are done configuring the device, commit the configuration.

```
[edit protocols mld]
user@host# commit
```

---

### Verification

To verify the configuration is working properly, check the **MLD Robustness Count** field in the output of the **show mld interfaces** command.

## Limiting the Maximum MLD Message Rate

You can change the limit for the maximum number of MLD packets transmitted in 1 second by the router.

Increasing the maximum number of MLD packets transmitted per second might be useful on a router with a large number of interfaces participating in MLD.

To change the limit for the maximum number of MLD packets the router can transmit in 1 second, include the **maximum-transmit-rate** statement and specify the maximum number of packets per second to be transmitted.

## Enabling MLD Static Group Membership

- [Create a MLD Static Group Member on page 5455](#)
- [Automatically create static groups on page 5456](#)

- [Automatically increment group addresses on page 5457](#)
- [Specify multicast source address \(in SSM mode\) on page 5458](#)
- [Automatically specify multicast sources on page 5459](#)
- [Automatically increment source addresses on page 5460](#)
- [Exclude multicast source addresses \(in SSM mode\) on page 5461](#)

### Create a MLD Static Group Member

You can create MLD static group membership to test multicast forwarding without a receiver host. When you enable MLD static group membership, data is forwarded to an interface without that interface receiving membership reports from downstream hosts.

Class-of-service (CoS) adjustment is not supported with MLD static group membership.

When you configure static groups on an interface on which you want to receive multicast traffic, you can specify the number of static groups to be automatically created.

In this example, you create static group ff0e::1:ff05:1a8d.

1. Configure the static groups to be created by including the **static** statement and **group** statement and specifying which IPv6 multicast address of the group to be created.

```
[edit protocols mld]
user@host# set interface fe-0/1/2 static (Protocols MLD) group ff0e::1:ff05:1a8d
```

2. After you commit the configuration, use the **show configuration protocol mld** command to verify the MLD protocol configuration.

```
user@host> show configuration protocol mld

interface fe-0/1/2.0 {
 static {
 group ff0e::1:ff05:1a8d;
 }
}
```

3. After you have committed the configuration and after the source is sending traffic, use the **show mld group** command to verify that static group ff0e::1:ff05:1a8d has been created.

```
user@host> show mld group
Interface: fe-0/1/2
Group: ff0e::1:ff05:1a8d
Group mode: Include
Source: fe80::2e0:81ff:fe05:1a8d
Last reported by: Local
Timeout: 0 Type: Static
```



**NOTE:** You must specify a unique address for each group.

### Automatically create static groups

---

When you create MLD static group membership to test multicast forwarding on an interface on which you want to receive multicast traffic, you can specify that a number of static groups be automatically created. This is useful when you want to test forwarding to multiple receivers without having to configure each receiver separately.

In this example, you create three groups.

1. Configure the number of static groups to be created by including the **group-count** statement and specifying the number of groups to be created.

```
[edit protocols mld]
user@host# set interface fe-0/1/2 static (Protocols MLD) group ff0e::1:ff05:1a8d
group-count 3
```

2. After you commit the configuration, use the **show configuration protocol mld** command to verify the MLD protocol configuration.

```
user@host> show configuration protocol mld

interface fe-0/1/2.0 {
 static {
 group ff0e::1:ff05:1a8d {
 group-count 3;
 }
 }
}
```

3. After you have committed the configuration and the source is sending traffic, use the **show mld group** command to verify that static groups ff0e::1:ff05:1a8d, ff0e::1:ff05:1a8e, and ff0e::1:ff05:1a8f have been created.

```
user@host> show mld group

Interface: fe-0/1/2
 Group: ff0e::1:ff05:1a8d
 Source: fe80::2e0:81ff:fe05:1a8d
 Last reported by: Local
 Timeout: 0 Type: Static
Interface: fe-0/1/2
 Group: ff0e::1:ff05:1a8e
 Source: fe80::2e0:81ff:fe05:1a8d
 Last reported by: Local
 Timeout: 0 Type: Static
Interface: fe-0/1/2
 Group: ff0e::1:ff05:1a8f
 Source: fe80::2e0:81ff:fe05:1a8d
 Last reported by: Local
 Timeout: 0 Type: Static
```



### Automatically increment group addresses

When you configure static groups on an interface on which you want to receive multicast traffic and you specify the number of static groups to be automatically created, you can also configure the group address to be automatically incremented by some number of addresses.

In this example, you create three groups and increase the group address by an increment of two for each group.

1. Configure the group address increment by including the **group-increment** statement and specifying the number by which the address should be incremented for each group. The increment is specified in a format similar to an IPv6 address.

```
[edit protocols mld]
user@host# set interface fe-0/1/2 static (Protocols MLD) group ff0e::1:ff05:1a8d
group-count 3 group-increment ::2
```

2. After you commit the configuration, use the **show configuration protocol mld** command to verify the MLD protocol configuration.

```
user@host> show configuration protocol mld

interface fe-0/1/2.0 {
 static {
 group ff0e::1:ff05:1a8d {
 group-increment ::2;
 group-count 3;
 }
 }
}
```

3. After you have committed the configuration and the source is sending traffic, use the **show mld group** command to verify that static groups ff0e::1:ff05:1a8d, ff0e::1:ff05:1a8f, and ff0e::1:ff05:1a91 have been created.

```
user@host> show mld group

Interface: fe-0/1/2
 Group: ff0e::1:ff05:1a8d
 Source: fe80::2e0:81ff:fe05:1a8d
 Last reported by: Local
 Timeout: 0 Type: Static
Interface: fe-0/1/2
 Group: ff0e::1:ff05:1a8f
 Source: fe80::2e0:81ff:fe05:1a8d
 Last reported by: Local
 Timeout: 0 Type: Static
Interface: fe-0/1/2
 Group: ff0e::1:ff05:1a91
 Source: fe80::2e0:81ff:fe05:1a8d
 Last reported by: Local
 Timeout: 0 Type: Static
```

### Specify multicast source address (in SSM mode)

---

When you configure static groups on an interface on which you want to receive multicast traffic and your network is operating in source-specific multicast (SSM) mode, you can specify the multicast source address to be accepted.

If you specify a group address in the SSM range, you must also specify a source.

If a source address is specified in a multicast group that is statically configured, the MLD version must be set to MLDv2 on the interface. MLDv1 is the default value.

In this example, you create group ff0e::1:ff05:1a8d and accept IPv6 address fe80::2e0:81ff:fe05:1a8d as the only source.

1. Configure the source address by including the **source** statement and specifying the IPv6 address of the source host.

```
[edit protocols mld]
user@host# set interface fe-0/1/2 static (Protocols MLD) group ff0e::1:ff05:1a8d
source fe80::2e0:81ff:fe05:1a8d
```

2. After you commit the configuration, use the **show configuration protocol mld** command to verify the MLD protocol configuration.

```
user@host> show configuration protocol mld

interface fe-0/1/2.0 {
 static {
 group ff0e::1:ff05:1a8d {
 source fe80::2e0:81ff:fe05:1a8d;
 }
 }
}
```

3. After you have committed the configuration and the source is sending traffic, use the **show mld group** command to verify that static group ff0e::1:ff05:1a8d has been created and that source fe80::2e0:81ff:fe05:1a8d has been accepted.

```
user@host> show mld group

Interface: fe-0/1/2
Group: ff0e::1:ff05:1a8d
Source: fe80::2e0:81ff:fe05:1a8d
Last reported by: Local
Timeout: 0 Type: Static
```

### Automatically specify multicast sources

When you configure static groups on an interface on which you want to receive multicast traffic, you can specify a number of multicast sources to be automatically accepted.

In this example, you create static group ff0e::1:ff05:1a8d and accept fe80::2e0:81ff:fe05:1a8d, fe80::2e0:81ff:fe05:1a8e, and fe80::2e0:81ff:fe05:1a8f as the source addresses.

1. Configure the number of multicast source addresses to be accepted by including the **source-count** statement and specifying the number of sources to be accepted.

```
[edit protocols mld]
user@host# set interface fe-0/1/2 static (Protocols MLD) group ff0e::1:ff05:1a8d
source fe80::2e0:81ff:fe05:1a8d source-count 3
```

2. After you commit the configuration, use the **show configuration protocol mld** command to verify the MLD protocol configuration.

```
user@host> show configuration protocol mld

interface fe-0/1/2.0 {
 static {
 group ff0e::1:ff05:1a8d {
 source fe80::2e0:81ff:fe05:1a8d {
 source-count 3;
 }
 }
 }
}
```

3. After you have committed the configuration and the source is sending traffic, use the **show mld group** command to verify that static group ff0e::1:ff05:1a8d has been created and that sources fe80::2e0:81ff:fe05:1a8d, fe80::2e0:81ff:fe05:1a8e, and fe80::2e0:81ff:fe05:1a8f have been accepted.

```
user@host> show mld group

Interface: fe-0/1/2
 Group: ff0e::1:ff05:1a8d
 Source: fe80::2e0:81ff:fe05:1a8d
 Last reported by: Local
 Timeout: 0 Type: Static
Interface: fe-0/1/2
 Group: ff0e::1:ff05:1a8d
 Source: fe80::2e0:81ff:fe05:1a8e
 Last reported by: Local
 Timeout: 0 Type: Static
Interface: fe-0/1/2
 Group: ff0e::1:ff05:1a8d
 Source: fe80::2e0:81ff:fe05:1a8f
 Last reported by: Local
 Timeout: 0 Type: Static
```

### Automatically increment source addresses

---

When you configure static groups on an interface on which you want to receive multicast traffic, and specify a number of multicast sources to be automatically accepted, you can also specify the number by which the address should be incremented for each source accepted.

In this example, you create static group ff0e::1:ff05:1a8d and accept fe80::2e0:81ff:fe05:1a8d, fe80::2e0:81ff:fe05:1a8f, and fe80::2e0:81ff:fe05:1a91 as the sources.

1. Configure the number of multicast source addresses to be accepted by including the **source-increment** statement and specifying the number of sources to be accepted.

```
[edit protocols mld]
```

```
user@host# set interface fe-0/1/2 static (Protocols MLD) group ff0e::1:ff05:1a8d
source fe80::2e0:81ff:fe05:1a8d source-count 3 source-increment ::2
```

2. After you commit the configuration, use the **show configuration protocol mld** command to verify the MLD protocol configuration.

```
user@host> show configuration protocol mld
```

```
interface fe-0/1/2.0 {
 static {
 group ff0e::1:ff05:1a8d {
 source fe80::2e0:81ff:fe05:1a8d {
 source-count 3;
 source-increment ::2;
 }
 }
 }
}
```

3. After you have committed the configuration and the source is sending traffic, use the **show mld group** command to verify that static group ff0e::1:ff05:1a8d has been created and that sources fe80::2e0:81ff:fe05:1a8d, fe80::2e0:81ff:fe05:1a8f, and fe80::2e0:81ff:fe05:1a91 have been accepted.

```
user@host> show mld group
```

```
Interface: fe-0/1/2
 Group: ff0e::1:ff05:1a8d
 Source: fe80::2e0:81ff:fe05:1a8d
 Last reported by: Local
 Timeout: 0 Type: Static
Interface: fe-0/1/2
 Group: ff0e::1:ff05:1a8d
 Source: fe80::2e0:81ff:fe05:1a8f
 Last reported by: Local
 Timeout: 0 Type: Static
Interface: fe-0/1/2
 Group: ff0e2::1:ff05:1a8d
 Source: fe80::2e0:81ff:fe05:1a91
 Last reported by: Local
 Timeout: 0 Type: Static
```

```

Interface: fe-0/1/2
Group: ff0e::1:ff05:1a8d
Group mode: Include
Source: fe80::2e0:81ff:fe05:1a8d
Last reported by: Local
Timeout: 0 Type: Static
Group: ff0e::1:ff05:1a8d
Group mode: Include
Source: fe80::2e0:81ff:fe05:1a8f
Last reported by: Local
Timeout: 0 Type: Static
Group: ff0e::1:ff05:1a8d
Group mode: Include
Source: fe80::2e0:81ff:fe05:1a91
Last reported by: Local
Timeout: 0 Type: Static

```

### Exclude multicast source addresses (in SSM mode)

When you configure static groups on an interface on which you want to receive multicast traffic and your network is operating in source-specific multicast (SSM) mode, you can specify that certain multicast source addresses be excluded.

By default the multicast source address configured in a static group operates in include mode. In include mode the multicast traffic for the group is accepted from the configured source address. You can also configure the static group to operate in exclude mode. In exclude mode the multicast traffic for the group is accepted from any address other than the configured source address.

If a source address is specified in a multicast group that is statically configured, the MLD version must be set to MLDv2 on the interface. MLDv1 is the default value.

In this example, you exclude address fe80::2e0:81ff:fe05:1a8d as a source for group ff0e::1:ff05:1a8d.

1. Configure a multicast static group to operate in exclude mode by including the **exclude** statement and specifying which IPv6 source address to be excluded.

```

[edit protocols mld]
user@host# set interface fe-0/1/2 static (Protocols MLD) group ff0e::1:ff05:1a8d
exclude source fe80::2e0:81ff:fe05:1a8d

```

2. After you commit the configuration, use the **show configuration protocol mld** command to verify the MLD protocol configuration.

```

user@host> show configuration protocol mld

interface fe-0/1/2.0 {
 static {
 group ff0e::1:ff05:1a8d {
 exclude;
 source fe80::2e0:81ff:fe05:1a8d;
 }
 }
}

```

- After you have committed the configuration and the source is sending traffic, use the **show mld group detail** command to verify that static group ff0e::1:ff05:1a8d has been created and that the static group is operating in exclude mode.

```
user@host> show mld group detail
Interface: fe-0/1/2
 Group: ff0e::1:ff05:1a8d
 Group mode: Exclude
 Source: fe80::2e0:81ff:fe05:1a8d
 Last reported by: Local
 Timeout: 0 Type: Static
```

Similar configuration is available for IPv4 multicast traffic using the IGMP protocol.

### Example: Recording MLD Join and Leave Events

This example shows how to determine whether MLD tuning is needed in a network by configuring the routing device to record MLD join and leave events.

- [Requirements on page 5462](#)
- [Overview on page 5462](#)
- [Configuration on page 5463](#)
- [Verification on page 5464](#)

#### Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Routing Devices*.
- Enable IPv6 unicast routing. See the *Junos OS Routing Protocols Library for Routing Devices*.
- Enable PIM. See [“PIM Overview” on page 5491](#).

#### Overview

[Table 434](#) describes the recordable MLD join and leave events.

**Table 434: MLD Event Messages**

| ERRMSG Tag             | Definition                                                   |
|------------------------|--------------------------------------------------------------|
| RPD_MLD_JOIN           | Records MLD join events.                                     |
| RPD_MLD_LEAVE          | Records MLD leave events.                                    |
| RPD_MLD_ACCOUNTING_ON  | Records when MLD accounting is enabled on an MLD interface.  |
| RPD_MLD_ACCOUNTING_OFF | Records when MLD accounting is disabled on an MLD interface. |

Table 434: MLD Event Messages (*continued*)

| ERRMSG Tag                 | Definition                             |
|----------------------------|----------------------------------------|
| RPD_MLD_MEMBERSHIP_TIMEOUT | Records MLD membership timeout events. |

### Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set protocols mld interface fe-0/1/0.2 accounting
set system syslog file mld-events any info
set system syslog file mld-events match ".*RPD_MLD_JOIN.* | .*RPD_MLD_LEAVE.* |
.*RPD_MLD_ACCOUNTING.* | .*RPD_MLD_MEMBERSHIP_TIMEOUT.*"
set system syslog file mld-events archive size 100000
set system syslog file mld-events archive files 3
set system syslog file mld-events archive transfer-interval 1440
set system syslog file mld-events archive archive-sites "ftp://user@host1//var/tmp"
password "anonymous"
set system syslog file mld-events archive archive-sites "ftp://user@host2//var/tmp"
password "test"
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure recording of MLD join and leave events:

1. Enable accounting globally or on an MLD interface. This example shows the interface configuration.

```
[edit protocols mld]
user@host# set interface fe-0/1/0.2 accounting
```

2. Configure the events to be recorded, and filter the events to a system log file with a descriptive filename, such as **mld-events**.

```
[edit system syslog file mld-events]
user@host# set any info
[edit system syslog file mld-events]
user@host# set match ".*RPD_MLD_JOIN.* | .*RPD_MLD_LEAVE.* |
.*RPD_MLD_ACCOUNTING.* | .*RPD_MLD_MEMBERSHIP_TIMEOUT.*"
```

3. Periodically archive the log file.

This example rotates the file every 24 hours (1440 minutes) when it reaches 100 KB and keeps three files.

```
[edit system syslog file mld-events]
user@host# set archive size 100000
[edit system syslog file mld-events]
user@host# set archive files 3
[edit system syslog file mld-events]
```

```
user@host# set archive archive-sites "ftp://user@host1//var/tmp" password
"anonymous"
[edit system syslog file mld-events]
user@host# set archive archive-sites "ftp://user@host2//var/tmp" password "test"
[edit system syslog file mld-events]
user@host# set archive transfer-interval 1440
[edit system syslog file mld-events]
user@host# set archive start-time 2011-01-07:12:30
```

4. If you are done configuring the device, commit the configuration.

```
[edit system syslog file mld-events]]
user@host# commit
```

---

### Verification

You can view the system log file by running the **file show** command.

```
user@host> file show mld-events
```

You can monitor the system log file as entries are added to the file by running the **monitor start** and **monitor stop** commands.

```
user@host> monitor start mld-events
```

```
*** mld-events ***
Apr 16 13:08:23 host mgd[16416]: UI_CMDLINE_READ_LINE: User 'user', command 'run
monitor start mld-events '
monitor
```

## Configuring the Number of MLD Multicast Group Joins on Logical Interfaces

The **group-limit** statement enables you to limit the number of MLD multicast group joins for logical interfaces. When this statement is enabled on a router running MLD version 2, the limit is applied upon receipt of the group report. Once the group limit is reached, subsequent join requests are rejected.

When configuring limits for MLD multicast groups, keep the following in mind:

- Each any-source group (\*G) counts as one group toward the limit.
- Each source-specific group (S,G) counts as one group toward the limit.
- Groups in MLDv2 exclude mode are counted toward the limit.
- Multiple source-specific groups count individually toward the group limit, even if they are for the same group. For example, (S1, G1) and (S2, G1) would count as two groups toward the configured limit.
- Combinations of any-source groups and source-specific groups count individually toward the group limit, even if they are for the same group. For example, (\*, G1) and (S, G1) would count as two groups toward the configured limit.
- Configuring and committing a group limit on a network that is lower than what already exists on the network results in the removal of all groups from the configuration. The



groups must then request to rejoin the network (up to the newly configured group limit).

- You can dynamically limit multicast groups on MLD logical interfaces by using dynamic profiles. For detailed information about creating dynamic profiles, see the *Junos OS Broadband Subscriber Management and Services Library*.

Beginning with Junos OS 12.2, you can optionally configure a system log warning threshold for MLD multicast group joins received on the logical interface. It is helpful to review the system log messages for troubleshooting purposes and to detect if an excessive amount of MLD multicast group joins have been received on the interface. These log messages convey when the configured group limit has been exceeded, when the configured threshold has been exceeded, and when the number of groups drop below the configured threshold.

The **group-threshold** statement enables you to configure the threshold at which a warning message is logged. The range is 1 through 100 percent. The warning threshold is a percentage of the group limit, so you must configure the **group-limit** statement to configure a warning threshold. For instance, when the number of groups exceed the configured warning threshold, but remain below the configured group limit, multicast groups continue to be accepted, and the device logs a warning message. In addition, the device logs a warning message after the number of groups drop below the configured warning threshold. You can further specify the amount of time (in seconds) between the log messages by configuring the **log-interval** statement. The range is 6 through 32,767 seconds.

You might consider throttling log messages because every entry added after the configured threshold and every entry rejected after the configured limit causes a warning message to be logged. By configuring a log interval, you can throttle the amount of system log warning messages generated for MLD multicast group joins.

To limit multicast group joins on an MLD logical interface:

1. Access the logical interface at the MLD protocol hierarchy level.

```
[edit]
user@host# edit protocols mld interface interface-name
```

2. Specify the group limit for the interface.

```
[edit protocols mld interface interface-name]
user@host# set group-limit limit
```

3. (Optional) Configure the threshold at which a warning message is logged.

```
[edit protocols mld interface interface-name]
user@host# set group-threshold value
```

4. (Optional) Configure the amount of time between log messages.

```
[edit protocols mld interface interface-name]
user@host# set log-interval seconds
```

To confirm your configuration, use the **show protocols mld** command. To verify the operation of MLD on the interface, including the configured group limit and the optional warning threshold and interval between log messages, use the **show mld interface** command.

## Disabling MLD

To disable MLD on an interface, include the **disable** statement:

```
interface interface-name {
 disable;
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols mld]
- [edit logical-systems *logical-system-name* protocols mld]

### Related Documentation

- *Configuring IGMP*

# Using MLD Snooping

- [Understanding MLD Snooping on page 5467](#)
- [Configuring MLD Snooping on a VLAN \(CLI Procedure\) on page 5475](#)
- [Example: Configuring MLD Snooping on page 5481](#)
- [Verifying MLD Snooping on page 5484](#)

## Understanding MLD Snooping

---



**NOTE:** This overview uses Junos OS for switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Understanding MLD Snooping*. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 41](#).

Multicast Listener Discovery (MLD) snooping constrains the flooding of IPv6 multicast traffic on VLANs. When MLD snooping is enabled on a VLAN, a Juniper Networks EX switch examines MLD messages between hosts and multicast routers and learns which hosts are interested in receiving traffic for a multicast group. On the basis of what it learns, the switch then forwards multicast traffic only to those interfaces in the VLAN that are connected to interested receivers instead of flooding the traffic to all interfaces.

MLD snooping supports MLD version 1 (MLDv1) and MLDv2. For details on MLDv1 and MLDv2, see the following standards:

- MLDv1—See RFC 2710, *Multicast Listener Discovery (MLD) for IPv6*.
- MLDv2—See RFC 3810, *Multicast Listener Discovery Version 2 (MLDv2) for IPv6*.

This topic covers:

- [How MLD Snooping Works on page 5468](#)
- [MLD Message Types on page 5469](#)
- [How Hosts Join and Leave Multicast Groups on page 5469](#)
- [Support for MLDv2 Multicast Sources on page 5470](#)
- [MLD Snooping and Forwarding Interfaces on page 5470](#)

- [General Forwarding Rules on page 5471](#)
- [Examples of MLD Snooping Multicast Forwarding on page 5471](#)


## How MLD Snooping Works

By default, a switch floods Layer 2 multicast traffic on all of the interfaces belonging to that VLAN on a switch, except for the interface that is the source of the multicast traffic. This behavior can consume significant amounts of bandwidth.

You can enable MLD snooping to avoid this flooding. When you enable MLD snooping, the switch monitors MLD messages between receivers (hosts) and multicast routers and uses the content of the messages to build an IPv6 multicast forwarding table—a database of IPv6 multicast groups and the interfaces that are connected to the interested members of each group. When the switch receives multicast traffic for a multicast group, it uses the forwarding table to forward the traffic only to interfaces that are connected to receivers that belong to the multicast group.

Figure 166 shows an example of multicast traffic flow with MLD snooping enabled.

**Figure 166: Multicast Traffic Flow with MLD Snooping Enabled**



## MLD Message Types

Multicast routers use MLD to learn, for each of their attached physical networks, which groups have interested listeners. In any given subnet, one multicast router is elected to act as an MLD querier. The MLD querier sends out the following types of queries to hosts:

- General query—Asks whether any host is listening to any group.
- Group-specific query—Asks whether any host is listening to a specific multicast group. This query is sent in response to a host leaving the multicast group and allows the router to quickly determine if any remaining hosts are interested in the group.
- Group-and-source-specific query—(MLD version 2 only) Asks whether any host is listening to group multicast traffic from a specific multicast source. This query is sent in response to a host indicating that it is no longer interested in receiving group multicast traffic from the multicast source and allows the router to quickly determine any remaining hosts are interested in receiving group multicast traffic from that source.

Hosts that are multicast listeners send the following kinds of messages:

- Membership report—Indicates that the host wants to join a particular multicast group.
- Leave report—Indicates that the host wants to leave a particular multicast group.

Strictly speaking, only MLDv1 hosts use two different kinds of reports to indicate whether they want to join or leave a group. MLDv2 hosts send only one kind of report, the contents of which indicate whether they want to join or leave a group. However, for simplicity's sake, the MLD snooping documentation uses the term *membership report* for a report that indicates that a host wants to join a group and uses the term *leave report* for a report that indicates a host wants to leave a group.

## How Hosts Join and Leave Multicast Groups

Hosts can join multicast groups in either of two ways:

- By sending an unsolicited membership report that specifies the multicast group that the host is attempting to join.
- By sending a membership report in response to a query from a multicast router.

A multicast router continues to forward multicast traffic to an interface provided that at least one host on that interface responds to the periodic general queries indicating its membership. For a host to remain a member of a multicast group, therefore, it must continue to respond to the periodic general queries.

Hosts can leave multicast groups in either of two ways:

- By not responding to periodic queries within a set interval of time. This results in what is known as a “silent leave.”
- By sending a leave report.



**NOTE:** If a host is connected to the switch through a hub, the host does not automatically leave the multicast group if it disconnects from the hub. The host remains a member of the group until group membership times out and a silent leave occurs. If another host connects to the hub port before the silent leave occurs, the new host might receive the group multicast traffic until the silent leave, even though it never sent an membership report.

---

## Support for MLDv2 Multicast Sources

In MLDv2, a host can send a membership report that includes a list of source addresses. When the host sends a membership report in INCLUDE mode, the host is interested in group multicast traffic only from those sources in the source address list. If host sends a membership report in EXCLUDE mode, the host is interested in group multicast traffic from any source *except* the sources in the source address list. A host can also send an EXCLUDE report in which the source-list parameter is empty, which is known as an EXCLUDE NULL report. An EXCLUDE NULL report indicates that the host wants to join the multicast group and receive packets from all sources.

## MLD Snooping and Forwarding Interfaces

To determine how to forward multicast traffic, a switch with MLD snooping enabled maintains information about the following interfaces in its multicast forwarding table:

- Multicast-router interfaces—These interfaces lead toward multicast routers or MLD queriers.
- Group-member interfaces—These interfaces lead toward hosts that are members of multicast groups.

The switch learns about these interfaces by monitoring MLD traffic. If an interface receives MLD queries, the switch adds the interface to its multicast forwarding table as a multicast-router interface. If an interface receives membership reports for a multicast group, the switch adds the interface to its multicast forwarding table as a group-member interface.

Table entries for interfaces that the switch learns about are subject to aging. For example, if a learned multicast-router interface does not receive MLD queries within a certain interval, the switch removes the entry for that interface from its multicast forwarding table.



**NOTE:** For a switch to learn multicast-router interfaces and group-member interfaces, an MLD querier must exist in the network. For the switch itself to function as an MLD querier, MLD must be enabled on the switch.

You can statically configure an interface to be a multicast-router interface or a group-member interface. The switch adds a static interface to its multicast forwarding table without having to learn about the interface, and the entry in the table is not subject

to aging. You can have a mix of statically configured and dynamically learned interfaces on a switch.

## General Forwarding Rules

Multicast traffic received on a switch interface in a VLAN on which MLD snooping is enabled is forwarded according to the following rules.

MLD protocol traffic is forwarded as follows:

- MLD general queries received on a multicast-router interface are forwarded to all other interfaces in the VLAN.
- MLD group-specific queries received on a multicast-router interface are forwarded to only those interfaces in the VLAN that are members of the group.
- MLD reports received on a host interface are forwarded to multicast-router interfaces in the same VLAN, but not to the other host interfaces in the VLAN.

Multicast traffic that is not MLD protocol traffic is forwarded as follows:

- An unregistered multicast packet—that is, a packet for a group that has no current members—is forwarded to all multicast-router interfaces in the VLAN.
- A registered multicast packet is forwarded only to those host interfaces in the VLAN that are members of the multicast group and to all multicast-router interfaces in the VLAN.

## Examples of MLD Snooping Multicast Forwarding

The following examples are provided to illustrate how MLD snooping forwards multicast traffic in different topologies:

- [Scenario 1: Switch Forwarding Multicast Traffic to a Multicast Router and Hosts on page 5471](#)
- [Scenario 2: Switch Forwarding Multicast Traffic to Another Switch on page 5472](#)
- [Scenario 3: Switch Connected to Hosts Only \(No MLD Querier\) on page 5473](#)
- [Scenario 4: Layer 2/Layer 3 Switch Forwarding Multicast Traffic Between VLANs on page 5474](#)

### Scenario 1: Switch Forwarding Multicast Traffic to a Multicast Router and Hosts

In the topology shown in [Figure 167](#), a switch acting as a Layer 2 device receives multicast traffic belonging to multicast group **ff1e::2010** from Source A, which is connected to the multicast router. It also receives multicast traffic belonging to multicast group **ff15::2** from Source B, which is connected directly to the switch. All interfaces on the switch belong to the same VLAN.


Because the switch receives MLD queries from the multicast router on interface P1, MLD snooping learns that interface P1 is a multicast-router interface and adds the interface to its multicast forwarding table. It forwards any MLD general queries it receives on this

interface to all host interfaces on the switch, and, in turn, forwards membership reports it receives from hosts to the multicast-router interface.

In the example, Hosts A and C have responded to the general queries with membership reports for group **ff1e::2010**. MLD snooping adds interfaces P2 and P4 to its multicast forwarding table as member interfaces for group **ff1e::2010**. It forwards the group multicast traffic received from Source A to Hosts A and C, but not to Hosts B and D.

Host B has responded to the general queries with a membership report for group **ff15::2**. The switch adds interface P3 to its multicast forwarding table as a member interface for group **ff15::2** and forwards multicast traffic it receives from Source B to Host B. The switch also forwards the multicast traffic it receives from Source B to the multicast-router interface P1.

**Figure 167: Scenario 1: Switch Forwarding Multicast Traffic to a Multicast Router and Hosts**



### Scenario 2: Switch Forwarding Multicast Traffic to Another Switch

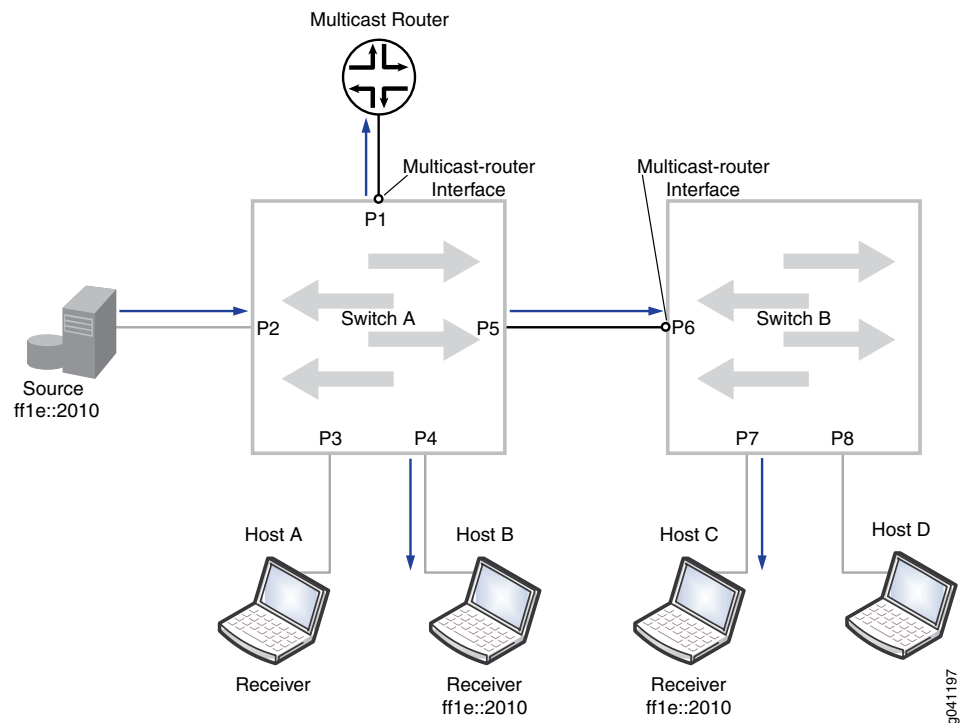
In the topology shown in Figure 168, a multicast source is connected to Switch A. Switch A in turn is connected to another switch, Switch B. Hosts on both Switch A and B are potential members of the multicast group. Both switches are acting as Layer 2 devices, and all interfaces on the switches are members of the same VLAN.

Switch A receives MLD queries from the multicast router on interface P1, making interface P1 a multicast-router interface for Switch A. Switch A forwards all general queries it receives on this interface to the other interfaces on the switch, including the interface connecting Switch B. Because Switch B receives the forwarded MLD queries on interface P6, P6 is the multicast-router interface for Switch B. Switch B forwards the membership



report it receives from Host C to Switch A through its multicast-router interface. Switch A forwards the membership report to its multicast-router interface, includes interface P5 in its multicast forwarding table as a group-member interface, and forwards multicast traffic from the source to Switch B.

**Figure 168: Scenario 2: Switch Forwarding Multicast Traffic to Another Switch**



In certain implementations, you might have to configure P6 on Switch B as a static multicast-router interface to avoid a delay in a host receiving multicast traffic. For example, if Switch B receives unsolicited membership reports from its hosts before it learns which interface is its multicast-router interface, it does not forward those reports to Switch A. If Switch A then receives multicast traffic, it does not forward the traffic to Switch B, because it has not received any membership reports on interface P5. This issue will resolve when the multicast router sends out its next general query; however, it can cause a delay in the host receiving multicast traffic. You can statically configure interface P6 as a multicast-router interface to solve this issue.


### Scenario 3: Switch Connected to Hosts Only (No MLD Querier)

In the topology shown in Figure 169, a switch is connected to a multicast source and to hosts. There is no multicast router in this topology—hence there is no MLD querier. Without an MLD querier to respond to, a host does not send periodic membership reports. As a result, even if the host sends an unsolicited membership report to join a multicast group, its membership in the multicast group will time out.

For MLD snooping to work correctly in this network so that the switch forwards multicast traffic to Hosts A and C only, you can either:

- Configure interfaces P2 and P4 as static group-member interfaces.
- Configure a routed VLAN interface (RVI) on the VLAN and enable MLD on it. In this case, the switch itself acts as an MLD querier, and the hosts can dynamically join the multicast group and refresh their group membership by responding to the queries.

**Figure 169: Scenario 3: Switch Connected to Hosts Only (No MLD Querier)**




#### Scenario 4: Layer 2/Layer 3 Switch Forwarding Multicast Traffic Between VLANs

In the topology shown in [Figure 170](#), a multicast source, Multicast Router A, and Hosts A and B are connected to the switch and are in VLAN 10. Multicast Router B and Hosts C and D are also connected to the switch and are in VLAN 20.

In a pure Layer 2 environment, traffic is not forwarded between VLANs. For Host C to receive the multicast traffic from the source on VLAN 10, RVIs must be created on VLAN 10 and VLAN 20 to permit routing of the multicast traffic between the VLANs.

Figure 170: Scenario 4: Layer 2/Layer 3 Switch Forwarding Multicast Traffic Between VLANs



- Related Documentation**
- *Example: Configuring MLD Snooping*
  - *Configuring MLD Snooping on a VLAN (CLI Procedure)*
  - *Verifying MLD Snooping (CLI Procedure)*

## Configuring MLD Snooping on a VLAN (CLI Procedure)



**NOTE:** This task uses Junos OS with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring MLD Snooping on a VLAN (CLI Procedure)*. For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on [page 41](#).

You can enable MLD snooping on a VLAN to constrain the flooding of IPv6 multicast traffic on the VLAN. When MLD snooping is enabled, a switch examines MLD messages between hosts and multicast routers and learns which hosts are interested in receiving multicast traffic for a multicast group. Based on what it learns, the switch then forwards IPv6 multicast traffic only to those interfaces connected to interested receivers instead of flooding the traffic to all interfaces.

You can perform the following configurations for each VLAN:

- Selectively enable MLD snooping on specific VLANs.
- Specify the MLD version for the general query that the switch sends on an interface when the interface comes up.
- Enable immediate leave to reduce the length of time it takes the switch to stop forwarding multicast traffic when the last member host on the interface leaves the group.
- Configure an interface as a static multicast-router interface so that the switch does not need to dynamically learn that the interface is a multicast-router interface.
- Configure an interface as a static member of a multicast group so that the switch does not need to dynamically learn the interface's membership.
- Change the value for certain timers and counters to match the values configured on the multicast router serving as the MLD querier.

This topic covers:

- [Enabling or Disabling MLD Snooping on VLANs on page 5476](#)
- [Configuring the MLD Version on page 5477](#)
- [Enabling Immediate Leave on page 5477](#)
- [Configuring an Interface as a Multicast-Router Interface on page 5478](#)
- [Configuring Static Group Membership on an Interface on page 5479](#)
- [Changing the Timer and Counter Values on page 5480](#)

## Enabling or Disabling MLD Snooping on VLANs

MLD snooping is not enabled on any VLAN by default. You must explicitly enable MLD snooping on specific interfaces.

- To enable MLD snooping on a specific VLAN:

```
[edit protocols mld-snooping]
user@switch# set vlan vlan-name
```



**NOTE:** You cannot enable MLD snooping on a secondary VLAN.

For example, to enable MLD snooping on VLAN education:

```
[edit protocols mld-snooping]
user@switch# set vlan education
```

- To disable MLD snooping on a specific VLAN:

```
[edit protocols mld-snooping]
user@switch# delete vlan vlan-name
```

You can also deactivate the MLD snooping protocol on the switch without changing the MLD snooping VLAN configurations:

```
[edit]
user@switch# deactivate protocols mld-snooping
```

## Configuring the MLD Version

You can configure the version of MLD queries sent by a switch when MLD snooping is enabled. By default, the switch uses MLD version 1 (MLDv1). If you are using Protocol-Independent Multicast source-specific multicast (PIM-SSM), we recommend that you configure the switch to use MLDv2.

Typically, a switch passively monitors MLD messages sent between multicast routers and hosts and does not send MLD queries. The exception is when a switch detects that an interface has come up. When an interface comes up, the switch sends an immediate general membership query to all hosts on the interface. By doing so, the switch enables the multicast routers to learn group memberships more quickly than they would if they had to wait until the MLD querier sent its next general query.

The MLD version of the general query determines the MLD version of the host membership reports as follows:

- MLD version 1 (MLDv1) general query—Both MLDv1 and MLDv2 hosts respond with an MLDv1 membership report.
- MLDv2 general query—MLDv2 hosts respond with an MLDv2 membership report, while MLDv1 hosts are unable to respond to the query.

By default, the switch sends MLDv1 queries. This ensures compatibility with hosts and multicast routers that support MLDv1 only and cannot process MLDv2 reports. However, if your VLAN contains MLDv2 multicast routers and hosts and the routers are running PIM-SSM, we recommend that you configure MLD snooping for MLDv2. Doing so enables the routers to quickly learn which multicast sources the hosts on the interface want to receive traffic from.



**NOTE:** Configuring the MLD version does not limit the version of MLD messages that the switch can snoop. A switch can snoop both MLDv1 and MLDv2 messages regardless of the MLD version configured.

To configure the MLD version on an interface:

```
[edit protocols]
user@switch# set mld interface interface-name version number
```

For example, to set the MLD version to version 2 on interface ge-0/0/2:

```
[edit protocols]
user@switch# set mld interface ge-0/0/2 version 2
```

## Enabling Immediate Leave

By default, when a switch with MLD snooping enabled receives an MLD leave report on a member interface, it waits for hosts on the interface to respond to MLD group-specific queries to determine whether there still are hosts on the interface interested in receiving

the group multicast traffic. If the switch does not see any membership reports for the group within a set interval of time, it removes the interface's group membership from the multicast forwarding table and stops forwarding multicast traffic for the group to the interface.

You can decrease the leave latency created by this default behavior by enabling immediate leave on a VLAN.

When you enable immediate leave on a VLAN, host tracking is also enabled, allowing the switch to keep track of the hosts on a interface that have joined a multicast group. When the switch receives a leave report from the last member of the group, it immediately stops forwarding traffic to the interface and does not wait for the interface group membership to time out.

Immediate leave is supported for both MLD version 1 (MLDv1) and MLDv2. However, with MLDv1, we recommend that you configure immediate leave only when there is only one MLD host on an interface. In MLDv1, only one host on a interface sends a membership report in response to a group-specific query—any other interested hosts suppress their reports. This report-suppression feature means that the switch only knows about one interested host at any given time.

To enable immediate leave on a VLAN:

```
[edit protocols]
user@switch# set mld-snooping vlan vlan-name immediate-leave
```

## Configuring an Interface as a Multicast-Router Interface

When MLD snooping is enabled on a switch, the switch determines which interfaces face a multicast router by monitoring interfaces for MLD queries or Protocol Independent Multicast (PIM) updates. If the switch receives these messages on an interface, it adds the interface to its multicast forwarding table as a multicast-router interface.

In addition to dynamically learned interfaces, the multicast forwarding table can include interfaces that you explicitly configure to be multicast router interfaces. Unlike the table entries for dynamically learned interfaces, table entries for statically configured interfaces are not subject to aging and deletion from the forwarding table.

Examples of when you might want to configure a static multicast-router interface include:

- You have an unusual network configuration that prevents MLD snooping from reliably learning about a multicast-router interface through monitoring MLD queries or PIM updates.
- Your implementation does not require an MLD querier.
- You have a stable topology and want to avoid the delay the dynamic learning process entails.

To configure an interface as a static multicast-router interface:

```
[edit protocols]
user@switch# set mld-snooping vlan vlan-name interface interface-name
multicast-router-interface
```

For example, to configure ge-0/0/5.0 as a multicast-router interface for VLAN employee:

```
[edit protocols]
user@switch# set mld-snooping vlan employee interface ge-0/0/5.0
multicast-router-interface
```

## Configuring Static Group Membership on an Interface

To determine how to forward multicast packets, a switch with MLD snooping enabled maintains a multicast forwarding table containing a list of host interfaces that have interested listeners for a specific multicast group. The switch learns which host interfaces to add or delete from this table by examining MLD membership reports as they arrive on interfaces on which MLD snooping is enabled.

In addition to such dynamically learned interfaces, the multicast forwarding table can include interfaces that you statically configure to be members of multicast groups. When you configure a static group interface, the switch adds the interface to the forwarding table as a host interface for the group. Unlike an entry for a dynamically learned interface, a static interface entry is not subject to aging and deletion from the forwarding table.

Examples of when you might want to configure static group membership on an interface include:

- You want to simulate an attached multicast receiver for testing purposes.
- The interface has receivers that cannot send MLD membership reports.
- You want the multicast traffic for a specific group to be immediately available to a receiver without any delay imposed by the dynamic join process.

You cannot configure multicast source addresses for a static group interface. The MLD version of a static group interface is always MLD version 1.



**NOTE:** The switch does not simulate MLD membership reports on behalf of a statically configured interface. Thus a multicast router might be unaware that the switch has an interface that is a member of the multicast group. You can configure a static group interface on the router to ensure that the switch receives the group multicast traffic.

To configure a host interface as a static member of a multicast group:

```
[edit protocols]
user@switch# set mld-snooping vlan vlan-name interface interface-name static group
ip-address
```

For example, to configure interface ge-0/0/11.0 in VLAN employee as a static member of multicast group ff1e::1:

```
[edit protocols]
user@switch# set mld-snooping vlan ip-camera-vlan interface ge-0/0/11.0 static group
ff1e::1
```

## Changing the Timer and Counter Values

MLD uses various timers and counters to determine how often an MLD querier sends out membership queries and when group memberships time out. On Juniper Networks switches, the MLD and MLD snooping timers and counters default values are set to the values recommended in RFC 2710, *Multicast Listener Discovery (MLD) for IPv6*. These values work well for most IPv6 multicast deployments.

There might be cases, however, where you might want to adjust the timer and counter values—for example, to reduce burstiness, to reduce leave latency, or to adjust for expected packet loss on a subnet. If you change a timer or counter value for the MLD querier on a VLAN, we recommend that you change the value for all multicast routers and switches on the VLAN so that all devices time out group memberships at approximately the same time.

The following timers and counters are configurable on a switch:

- **query-interval**—The length of time in seconds the MLD querier waits between sending general queries (the default is 125 seconds). You can change this interval to tune the number of MLD messages on the subnet; larger values cause general queries to be sent less often.

To configure the MLD query interval:

```
[edit protocols]
user@switch# set mld-snooping vlan vlan-name query-interval seconds
```

- **query-response-interval**—The maximum length of time in seconds the host waits before it responds (the default is 10 seconds). You can change this interval to accommodate the burst peaks of MLD messages on the subnet. Set a larger interval to make the traffic less bursty.

To configure the MLD query response interval:

```
[edit protocols]
user@switch# set mld-snooping vlan vlan-name query-response-interval seconds
```

- **query-last-member-interval**—The length of time the MLD querier waits between sending group-specific membership queries (the default is 1 second). The MLD querier sends a group-specific query after receiving a leave report from a host. You can decrease this interval to reduce the amount of time it takes for multicast traffic to stop forwarding after the last member leaves a group.

To configure the MLD query last member interval:

```
[edit protocols]
user@switch# set mld-snooping vlan vlan-name query-last-member-interval seconds
```

- **robust-count**—The number of times the querier resends a general membership query or a group-specific membership query (the default is 2 times). You can increase this count to tune for higher anticipated packet loss.

For MLD snooping, you can configure **robust-count** for a specific VLAN. If a VLAN does not have **robust-count** configured, the value is inherited from the value configured for MLD.



To configure **robust-count** for MLD snooping on a VLAN:

[edit protocols]

user@switch# **set mld-snooping vlan *vlan-name* robust-count *number***

The values configured for **query-interval**, **query-response-interval**, and **robust-count** determine the multicast listener interval—the length of time the switch waits for a group membership report after a general query before removing a multicast group from its multicast forwarding table. The switch calculates the multicast listener interval by multiplying **query-interval** value by the **robust-count** value and then adding the **query-response-interval** to the product:

$(\text{query-interval} \times \text{robust-count}) + \text{query-response-interval} = \text{multicast listener interval}$

For example, the multicast listener interval is 260 seconds when the default settings for **query-interval**, **query-response-interval**, and **robust-count** are used:

$(125 \times 2) + 10 = 260$

To display the time remaining in the multicast listener interval before a group times out, use the **show mld-snooping membership** command.

#### Related Documentation

- [Example: Configuring MLD Snooping on page 5481](#)
- [Examples: Configuring MLD on page 5444](#)
- [Verifying MLD Snooping on page 5484](#)

## Example: Configuring MLD Snooping



**NOTE:** This example uses Junos OS with support for the Enhanced Layer 2 Software (ELS) configuration style. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 41](#).

You can enable MLD snooping on a VLAN to constrain the flooding of IPv6 multicast traffic on a VLAN. When MLD snooping is enabled, a switch examines MLD messages between hosts and multicast routers and learns which hosts are interested in receiving multicast traffic for a multicast group. On the basis of what it learns, the switch then forwards IPv6 multicast traffic only to those interfaces connected to interested receivers instead of flooding the traffic to all interfaces.

This example describes how to configure MLD snooping:

- [Requirements on page 5482](#)
- [Overview and Topology on page 5482](#)
- [Configuration on page 5483](#)
- [Verifying MLD Snooping Configuration on page 5484](#)

## Requirements

This example uses the following software and hardware components:

- One switch running Junos OS with ELS
- Junos OS Release 13.3 or later for EX Series switches or Junos OS Release 15.1X53-D10 or later for QFX10000 switches

Before you configure MLD snooping, be sure you have:

- Configured the vlan 100 VLAN on the switch.
- Assigned interfaces ge-0/0/0, ge-0/0/1, ge-0/0/2, and ge-0/0/12 to vlan100.
- Configured ge-0/0/12 as a trunk interface.


See *Configuring VLANs for EX Series Switches (CLI Procedure)* or “[Configuring VLANs](#)” on [page 2120](#).

## Overview and Topology

In this example, interfaces ge-0/0/0, ge-0/0/1, and ge-0/0/2 on the switch are in vlan100 and are connected to hosts that are potential multicast receivers. Interface ge-0/0/12, a trunk interface also in vlan100, is connected to a multicast router. The router acts as the MLD querier and forwards multicast traffic for group ff1e::2010 to the switch from a multicast source.

The topology for this example is illustrated in [Figure 171](#).

**Figure 171: MLD Snooping Topology Example**



In this sample topology, the multicast router forwards multicast traffic to the switch from the source when it receives a membership report for group ff1e::2010 from one of the

hosts—for example, Host B. If MLD snooping is not enabled on vlan100, the switch floods the multicast traffic on all interfaces in vlan100 (except for interface ge-0/0/12). If MLD snooping is enabled on vlan100, the switch monitors the MLD messages between the hosts and router, allowing it to determine that only Host B is interested in receiving the multicast traffic. The switch then forwards the multicast traffic only to interface ge-0/0/1.

This example shows how to enable MLD snooping on vlan100. It also shows how to perform the following optional configurations, which can reduce group join and leave latency:

- Configure immediate leave on the VLAN. When immediate leave is configured, the switch stops forwarding multicast traffic on an interface when it detects that the last member of the multicast group has left the group. If immediate leave is not configured, the switch waits until the group-specific membership queries time out before it stops forwarding traffic.
- Configure ge-0/0/12 as a static multicast-router interface. In this topology, ge-0/0/12 always leads to the multicast router. By statically configuring ge-0/0/12 as a multicast-router interface, you avoid any delay imposed by the switch having to learn that ge-0/0/12 is a multicast-router interface.

## Configuration

To configure MLD snooping on a switch:

### CLI Quick Configuration

To quickly configure MLD snooping, copy the following commands and paste them into the switch terminal window:

```
[edit]
set protocols mld-snooping vlan vlan100
set protocols mld-snooping vlan vlan100 immediate-leave
set protocols mld-snooping vlan vlan100 interface ge-0/0/12 multicast-router-interface
```

### Step-by-Step Procedure

To configure MLD snooping:

1. Enable MLD snooping on the VLAN vlan100:
 

```
[edit protocols]
user@switch# set mld-snooping vlan vlan100
```
2. Configure the switch to immediately remove a group membership from an interface when it receives a leave report from the last member of the group on the interface:
 

```
[edit protocols]
user@switch# set mld-snooping vlan vlan100 immediate-leave
```
3. Statically configure interface ge-0/0/12 as a multicast-router interface:
 

```
[edit protocols]
user@switch# set mld-snooping vlan vlan100 interface ge-0/0/12 multicast-router-interface
```

### Results

Check the results of the configuration:

```
[edit protocols]
user@switch# show mld-snooping
vlan vlan100 {
 immediate-leave;
 interface ge-0/0/12.0 {
 multicast-router-interface;
```

```
 }
}
```

## Verifying MLD Snooping Configuration

To verify that MLD snooping is enabled on the VLAN and the MLD snooping forwarding interfaces are correct, perform the following task:

- [Verifying MLD Snooping Interface Membership on VLAN vlan100 on page 5484](#)

---

### Verifying MLD Snooping Interface Membership on VLAN vlan100

**Purpose** Verify that MLD snooping is enabled on the VLAN vlan 100 and that the multicast-router interface is statically configured:

**Action** Show the MLD snooping information for ge-0/0/12.0:

```
user@switch> show mld snooping interface
Instance: default-switch
```

```
Vlan: vlan100
```

```
Learning-Domain: default
Interface: ge-0/0/12.0
 State: Up Groups: 3
 Immediate leave: On
 Router interface: yes
```

```
Configured Parameters:
MLD Query Interval: 125.0
MLD Query Response Interval: 10.0
MLD Last Member Query Interval: 1.0
MLD Robustness Count: 2
```

**Meaning** MLD snooping is running on **vlan100**, and interface **ge-0/0/12.0** is a statically configured multicast-router interface. Immediate leave is enabled on the interface.

**Related Documentation**

- [Configuring MLD Snooping on a VLAN \(CLI Procedure\) on page 5475](#)
- [Verifying MLD Snooping on page 5484](#)
- [Understanding MLD Snooping on page 5467](#)

---

## Verifying MLD Snooping



**NOTE:** This topic uses Junos OS with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Verifying MLD Snooping (CLI Procedure)*. For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on page 41.

Multicast Listener Discovery (MLD) snooping constrains the flooding of IPv6 multicast traffic on VLANs. This topic describes how to verify MLD snooping operation on a VLAN.

It covers:

- [Verifying MLD Snooping Memberships on page 5485](#)
- [Verifying MLD Snooping Interfaces on page 5485](#)
- [Viewing MLD Snooping Statistics on page 5486](#)
- [Viewing MLD Snooping Routing Information on page 5487](#)

## Verifying MLD Snooping Memberships

**Purpose** Verify that MLD snooping is enabled on a VLAN and determine group memberships.

**Action** Enter the following command:

```
user@switch> show mld snooping membership detail
Instance: default-switch
```

```
Vlan: v1
```

```
Learning-Domain: default
Interface: ge-0/0/1.0, Groups: 1
 Group: ff05::1
 Group mode: Exclude
 Source: ::
 Last reported by: fe80::
 Group timeout: 259 Type: Dynamic
Interface: ge-0/0/2.0, Groups: 0
```

**Meaning** The switch has multicast membership information for one VLAN on the switch, **v1**. MLD snooping might be enabled on other VLANs, but the switch does not have any multicast membership information for them.

- The following information is provided about the group memberships for the VLAN:
  - Currently, the VLAN has membership in only one multicast group, **ff05::1**.
  - The host or hosts that have reported membership in the group are on interface **ge-0/0/1.0**.
  - The last host that reported membership in the group has address **fe80::**.
  - The interface group membership will time out in **259** seconds if no hosts respond to membership queries during this interval.
  - The group membership has been learned by MLD snooping, as indicated by **Dynamic**.

## Verifying MLD Snooping Interfaces

**Purpose** Display MLD snooping information for each interface on which MLD snooping is enabled.

**Action** Enter the following command:

```
user@switch> show mld snooping interface
Instance: default-switch
```

```
Vlan: v100
```

```

Learning-Domain: default
Interface: ge-0/0/1.0
 State: Up Groups: 1
 Immediate leave: Off
 Router interface: no
Interface: ge-0/0/2.0
 State: Up Groups: 0
 Immediate leave: Off
 Router interface: no

Configured Parameters:
MLD Query Interval: 125.0
MLD Query Response Interval: 10.0
MLD Last Member Query Interval: 1.0
MLD Robustness Count: 2

```

**Meaning** MLD snooping is configured on one VLAN on the switch, **v100**. Each interface in each VLAN is listed and the following information is provided:

- How many multicast groups the interface belongs to.
- Whether immediate leave has been configured for the interface.
- Whether the interface is a multicast-router interface.

The output also shows the configured parameters for the MLD querier.

## Viewing MLD Snooping Statistics

**Purpose** Display MLD snooping statistics, such as number of MLD queries, reports, and leaves received and how many of these MLD messages contained errors.

**Action** Enter the following command:

```

user@switch>show mld snooping statistics
Vlan: v1
MLD Message type Received Sent Rx errors
Listener Query (v1/v2) 0 4 0
Listener Report (v1) 447 0 0
Listener Done (v1/v2) 0 0 0
Listener Report (v2) 0 0 0
Other Unknown types 0 0
Vlan: v2
MLD Message type Received Sent Rx errors
Listener Query (v1/v2) 0 4 0
Listener Report (v1) 154 0 0
Listener Done (v1/v2) 0 0 0
Listener Report (v2) 0 0 0
Other Unknown types 0 0
Instance: default-switch
MLD Message type Received Sent Rx errors
Listener Query (v1/v2) 0 8 0
Listener Report (v1) 601 0 0
Listener Done (v1/v2) 0 0 0
Listener Report (v2) 0 0 0
Other Unknown types 0 0

```

```

MLD Global Statistics
Bad Length 0
Bad Checksum 0
Bad Receive If 0
Rx non-local 0
Timed out 0

```

**Meaning** The output shows how many MLD messages of each type—**Queries**, **Done**, **Report**—the switch received or transmitted on interfaces on which MLD snooping is enabled. For each message type, it also shows the number of MLD packets the switch received that had errors—for example, packets that do not conform to the MLDv1 or MLDv2 standards. If the **Rx errors** count increases, verify that the hosts are compliant with MLDv1 or MLDv2 standards. If the switch is unable to recognize the MLD message type for a packet, it counts the packet under **Other Unknown types**.

## Viewing MLD Snooping Routing Information

**Purpose** Display the next-hop information maintained in the multicast snooping forwarding table.

**Action** Enter the following command:

```

user@switch>show multicast snooping route
Nexthop Bulking: OFF

```

```

Family: INET6

```

```

Group: ff00::/8
Source: ::/128
Vlan: v1

Group: ff02::1/128
Source: ::/128
Vlan: v1
Downstream interface list:
ge-1/0/16.0

```

```

Group: ff05::1/128
Source: ::/128
Vlan: v1
Downstream interface list:
ge-1/0/16.0

```

```

Group: ff06::1/128
Source: ::/128
Vlan: v1
Downstream interface list:
ge-1/0/16.0

```

**Meaning** The output shows the next-hop interfaces for a given multicast group on a VLAN. For example, route **ff02::1/128** on VLAN **v1** has the next-hop interface **ge-1/0/16.0**.

**Related Documentation**

- *clear mld snooping membership*
- *clear mld snooping statistics*
- [Example: Configuring MLD Snooping on page 5481](#)

- [Configuring MLD Snooping on a VLAN \(CLI Procedure\) on page 5475](#)



## PART 75

# Configuring PIM

- [Using PIM Basic Features on page 5491](#)
- [Using PIM Sparse Mode on page 5509](#)
- [Using PIM Dense Mode and PIM Sparse-Dense Mode on page 5523](#)
- [Using Source-Specific Multicast on page 5529](#)
- [Using Static RP on page 5545](#)
- [Using Anycast RP on page 5551](#)
- [Using Auto-RP on page 5561](#)
- [Using PIM Bootstrap Router on page 5567](#)
- [Using PIM Filtering on page 5571](#)
- [Using PIM RPT and SPT Cutover on page 5579](#)



# Using PIM Basic Features

- [PIM Overview on page 5491](#)
- [PIM on Aggregated Interfaces on page 5494](#)
- [Changing the PIM Version on page 5494](#)
- [Modifying the PIM Hello Interval on page 5494](#)
- [Preserving Multicast Performance by Disabling Response to the ping Utility on page 5495](#)
- [Configuring PIM Trace Options on page 5496](#)
- [Configuring Interface Priority for PIM Designated Router Selection on page 5498](#)
- [Configuring PIM Designated Router Election on Point-to-Point Links on page 5499](#)
- [Configuring BFD for PIM on page 5500](#)
- [Configuring BFD Authentication for PIM on page 5502](#)
- [Disabling PIM on page 5505](#)

## PIM Overview

---

The predominant multicast routing protocol in use on the Internet today is Protocol Independent Multicast, or PIM. The type of PIM used on the Internet is PIM sparse mode. PIM sparse mode is so accepted that when the simple term “PIM” is used in an Internet context, some form of sparse mode operation is assumed.

PIM emerged as an algorithm to overcome the limitations of dense-mode protocols such as the Distance Vector Multicast Routing Protocol (DVMRP), which was efficient for dense clusters of multicast receivers, but did not scale well for the larger, sparser, groups encountered on the Internet. The Core Based Trees (CBT) Protocol was intended to support sparse mode as well, but CBT, with its all-powerful core approach, made placement of the core critical, and large conference-type applications (many-to-many) resulted in bottlenecks in the core. PIM was designed to avoid the dense-mode scaling issues of DVMRP and the potential performance issues of CBT at the same time.

PIM is one of the most rapidly evolving specifications on the Internet today. Since its introduction in 1995, PIM has already seen two major revisions to its packet structure (PIM version 1 [PIMv1] and PIM version 2 [PIMv2]), two major RFCs (RFC 2362 obsoleted RFC 2117), and numerous drafts describing major components of PIM, such as many-to-many trees and source-specific multicast (SSM). Long-lasting RFCs are not a feature of PIM, and virtually all of PIM must be researched, understood, and implemented

directly from Internet drafts. In fact, no current RFC describes PIMv1 at all. The drafts have all expired, and PIMv1 was never issued as an official RFC.

PIM itself is not nonstandard or unstable, however. PIM has been a promising multicast routing protocol since its inception, especially PIM sparse mode, the first real sparse-mode multicast routing protocol. Work continues on PIM in a number of areas, from bidirectional trees to network management, and the rapid pace of development makes drafts essential for PIM.

PIMv1 and PIMv2 can coexist on the same router and even on the same interface. The main difference between PIMv1 and PIMv2 is the packet format. PIMv1 messages use Internet Group Management Protocol (IGMP) packets, whereas PIMv2 has its own IP protocol number (103) and packet structure. All routers connecting to an IP subnet such as a LAN must use the same PIM version. Some PIM implementations can recognize PIMv1 packets and automatically switch the router interface to PIMv1. Because the difference between PIMv1 and PIMv2 involves the message format, but not the meaning of the message or how the router processes the PIM message, a router can easily mix PIMv1 and PIMv2 interfaces.

PIM is used for efficient routing to multicast groups that might span wide-area and interdomain internetworks. It is called “protocol independent” because it does not depend on a particular unicast routing protocol. Junos OS supports bidirectional mode, sparse mode, dense mode, and sparse-dense mode.

PIM operates in several modes: bidirectional mode, sparse mode, dense mode, and sparse-dense mode. In sparse-dense mode, some multicast groups are configured as dense mode (flood-and-prune, [S,G] state) and others are configured as sparse mode (explicit join to rendezvous point [RP], [\*G] state).

PIM drafts also establish a mode known as PIM source-specific mode, or PIM SSM. In PIM SSM there is only one specific source for the content of a multicast group within a given domain.

Because the PIM mode you choose determines the PIM configuration properties, you first must decide whether PIM operates in bidirectional, sparse, dense, or sparse-dense mode in your network. Each mode has distinct operating advantages in different network environments.

- In sparse mode, routers must join and leave multicast groups explicitly. Upstream routers do not forward multicast traffic to a downstream router unless the downstream router has sent an explicit request (by means of a join message) to the rendezvous point (RP) router to receive this traffic. The RP serves as the root of the shared multicast delivery tree and is responsible for forwarding multicast data from different sources to the receivers.

Sparse mode is well suited to the Internet, where frequent interdomain join messages and prune messages are common.

- Bidirectional PIM is similar to sparse mode, and is especially suited to applications that must scale to support a large number of dispersed sources and receivers. In bidirectional PIM, routers build shared bidirectional trees and do not switch to a source-based tree.

Bidirectional PIM scales well because it needs no source-specific (S,G) state. Instead, it builds only group-specific (\*G) state.

- Unlike sparse mode and bidirectional mode, in which data is forwarded only to routers sending an explicit PIM join request, dense mode implements a *flood-and-prune* mechanism, similar to the Distance Vector Multicast Routing Protocol (DVMRP). In dense mode, a router receives the multicast data on the incoming interface, then forwards the traffic to the outgoing interface list. Flooding occurs periodically and is used to refresh state information, such as the source IP address and multicast group pair. If the router has no interested receivers for the data, and the outgoing interface list becomes empty, the router sends a PIM prune message upstream.

Dense mode works best in networks where few or no prunes occur. In such instances, dense mode is actually more efficient than sparse mode.

- Sparse-dense mode, as the name implies, allows the interface to operate on a per-group basis in either sparse or dense mode. A group specified as “dense” is not mapped to an RP. Instead, data packets destined for that group are forwarded by means of PIM dense mode rules. A group specified as “sparse” is mapped to an RP, and data packets are forwarded by means of PIM sparse-mode rules. Sparse-dense mode is useful in networks implementing auto-RP for PIM sparse mode.



**NOTE:** On SRX Series devices, PIM does not support upstream and downstream interfaces across different virtual routers in flow mode.

## Basic PIM Network Components

PIM dense mode requires only a multicast source and series of multicast-enabled routers running PIM dense mode to allow receivers to obtain multicast content. Dense mode makes sure that all multicast traffic gets everywhere by periodically flooding the network with multicast traffic, and relies on prune messages to make sure that subnets where all receivers are uninterested in that particular multicast group stop receiving packets.

PIM sparse mode is more complicated and requires the establishment of special routers called *rendezvous points (RPs)* in the network core. These routers are where upstream join messages from interested receivers meet downstream traffic from the source of the multicast group content. A network can have many RPs, but PIM sparse mode allows only one RP to be active for any multicast group.

If there is only one RP in a routing domain, the RP and adjacent links might become congested and form a single point of failure for all multicast traffic. Thus, multiple RPs are the rule, but the issue then becomes how other multicast routers find the RP that is the source of the multicast group the receiver is trying to join. This RP-to-group mapping is controlled by a special *bootstrap router (BSR)* running the PIM BSR mechanism. There can be more than one bootstrap router as well, also for single-point-of-failure reasons.

The bootstrap router does not have to be an RP itself, although this is a common implementation. The bootstrap router's main function is to manage the collection of RPs and allow interested receivers to find the source of their group's multicast traffic.

PIM SSM can be seen as a subset of a special case of PIM sparse mode and requires no specialized equipment other than that used for PIM sparse mode (and IGMP version 3).

Bidirectional PIM RPs, unlike RPs for PIM sparse mode, do not need to perform PIM Register tunneling or other specific protocol action. Bidirectional PIM RPs implement no specific functionality. RP addresses are simply a location in the network to rendezvous toward. In fact, for bidirectional PIM, RP addresses need not be loopback interface addresses or even be addresses configured on any router, as long as they are covered by a subnet that is connected to a bidirectional PIM-capable router and advertised to the network.

**Related Documentation**

- *Supported IP Multicast Protocol Standards*

---

## PIM on Aggregated Interfaces

If you configure PIM on an aggregated (**ae-** or **as-**) interface, each of the interfaces in the aggregate is included in the multicast output interface list and carries the single stream of replicated packets in a load-sharing fashion. The multicast aggregate interface is “expanded” into its constituent interfaces in the next-hop database.

**Related Documentation**

- [PIM Overview on page 5491](#)
- [interface on page 5719](#)

---

## Changing the PIM Version

All systems on a subnet must run the same version of PIM.

The default PIM version can be version 1 or version 2, depending on the mode you are configuring. PIMv1 is the default for rendezvous point (RP) mode (at the **[edit protocols pim rp static address address]** hierarchy level). However, PIMv2 is the default for interface mode (at the **[edit protocols pim interface interface-name]** hierarchy level). Explicitly configured versions override the defaults.

To configure the PIM version, include the **version** statement:

```
version (1 | 2);
```

---

## Modifying the PIM Hello Interval

Routing devices send hello messages at a fixed interval on all PIM-enabled interfaces. By using hello messages, routing devices advertise their existence as PIM routing devices on the subnet. With all PIM-enabled routing devices advertised, a single designated router for the subnet is established.

When a routing device is configured for PIM, it sends a hello message at a 30-second default interval. The interval range is from 0 through 255. When the interval counts down to 0, the routing device sends another hello message, and the timer is reset. A routing device that receives no response from a neighbor in 3.5 times the interval value drops

the neighbor. In the case of a 30-second interval, the amount of time a routing device waits for a response is 105 seconds.

If a PIM hello message contains the hold-time option, the neighbor timeout is set to the hold-time sent in the message. If a PIM hello message does not contain the hold-time option, the neighbor timeout is set to the default hello hold time.

To modify how often the routing device sends hello messages out of an interface:

1. This example shows the configuration for the routing instance. Configure the interface globally or in the routing instance.

```
[edit routing-instances PIM.master protocols pim interface fe-3/0/2.0]
user@host# set hello-interval 255
```

2. Verify the configuration by checking the **Hello Option Holdtime** field in the output of the **show pim neighbors detail** command.

```
user@host> show pim neighbors detail
Instance: PIM.master
Interface: fe-3/0/2.0
Address: 192.168.195.37, IPv4, PIM v2, Mode: Sparse
Hello Option Holdtime: 255 seconds
Hello Option DR Priority: 1
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported
Rx Join: Group Source Timeout
225.1.1.1 192.168.195.78 0
225.1.1.1 0

Interface: lo0.0
Address: 10.255.245.91, IPv4, PIM v2, Mode: Sparse
Hello Option Holdtime: 255 seconds
Hello Option DR Priority: 1
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported

Interface: pd-6/0/0.32768
Address: 0.0.0.0, IPv4, PIM v2, Mode: Sparse
Hello Option Holdtime: 255 seconds
Hello Option DR Priority: 0
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported
```

Related Documentation • [show pim neighbors on page 5907](#)

## Preserving Multicast Performance by Disabling Response to the ping Utility

The ping utility uses ICMP Echo messages to verify connectivity to any device with an IP address. However, in the case of multicast applications, a single ping sent to a multicast address can degrade the performance of routers because the stream of packets is replicated multiple times.

You can disable the router's response to ping (ICMP Echo) packets sent to multicast addresses. The system responds normally to unicast ping packets.

To disable the router's response to ping packets sent to multicast addresses:

1. Include the **no-multicast-echo** statement:

```
[edit system]
user@host# set no-multicast-echo
```

2. Verify the configuration by checking the **echo drops with broadcast or multicast destination address** field in the output of the **show system statistics icmp** command.

```
user@host> show system statistics icmp

icmp:
0 drops due to rate limit
0 calls to icmp_error
0 errors not generated because old message was icmp
Output histogram:
echo reply: 21
0 messages with bad code fields
0 messages less than the minimum length
0 messages with bad checksum
0 messages with bad source address
0 messages with bad length
100 echo drops with broadcast or multicast destination address
0 timestamp drops with broadcast or multicast destination address
Input histogram:
echo: 21
21 message responses generated
```

#### Related Documentation

- [Configuring Junos OS to Disable the Routing Engine Response to Multicast Ping Packets on page 86](#)
- *show system statistics icmp*

## Configuring PIM Trace Options

Tracing operations record detailed messages about the operation of routing protocols, such as the various types of routing protocol packets sent and received, and routing policy actions. You can specify which trace operations are logged by including specific tracing flags. The following table describes the flags that you can include.

| Flag                             | Description                                                                                                                                                     |
|----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>all</b>                       | Trace all operations.                                                                                                                                           |
| <b>assert</b>                    | Trace assert messages, which are used to resolve which of the parallel routers connected to a multiaccess LAN is responsible for forwarding packets to the LAN. |
| <b>autorp</b>                    | Trace bootstrap, RP, and auto-RP messages.                                                                                                                      |
| <b>bidirectional-df-election</b> | Trace bidirectional PIM designated-forwarder (DF) election events.                                                                                              |



| Flag                       | Description                                                                                                                                              |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>bootstrap</b>           | Trace bootstrap messages, which are sent periodically by the PIM domain's bootstrap router and are forwarded, hop by hop, to all routers in that domain. |
| <b>general</b>             | Trace general events.                                                                                                                                    |
| <b>graft</b>               | Trace graft and graft acknowledgment messages.                                                                                                           |
| <b>hello</b>               | Trace hello packets, which are sent so that neighboring routers can discover one another.                                                                |
| <b>join</b>                | Trace join messages, which are sent to join a branch onto the multicast distribution tree.                                                               |
| <b>mdt</b>                 | Trace messages related to multicast data tunnels.                                                                                                        |
| <b>normal</b>              | Trace normal events.                                                                                                                                     |
| <b>nsr-synchronization</b> | Trace nonstop routing synchronization events                                                                                                             |
| <b>packets</b>             | Trace all PIM packets.                                                                                                                                   |
| <b>policy</b>              | Trace poison-route-reverse packets.                                                                                                                      |
| <b>prune</b>               | Trace prune messages, which are sent to prune a branch off the multicast distribution tree.                                                              |
| <b>register</b>            | Trace register and register-stop messages. Register messages are sent to the RP when a multicast source first starts sending to a group.                 |
| <b>route</b>               | Trace routing information.                                                                                                                               |
| <b>rp</b>                  | Trace candidate RP advertisements.                                                                                                                       |
| <b>state</b>               | Trace state transitions.                                                                                                                                 |
| <b>task</b>                | Trace task processing.                                                                                                                                   |
| <b>timer</b>               | Trace timer processing.                                                                                                                                  |

In the following example, tracing is enabled for all routing protocol packets. Then tracing is narrowed to focus only on PIM packets of a particular type.

To configure tracing operations for PIM:

1. (Optional) Configure tracing at the [**routing-options** hierarchy level to trace all protocol packets.

[edit routing-options [traceoptions](#)]

```
user@host# set file all-packets-trace
user@host# set flag all
```

2. Configure the filename for the PIM trace file.

```
[edit protocols pim traceoptions]
user@host# set file pim-trace
```

3. (Optional) Configure the maximum number of trace files.

```
[edit protocols pim traceoptions]
user@host# set file files 5
```

4. (Optional) Configure the maximum size of each trace file.

```
[edit protocols pim traceoptions]
user@host# set file size 1m
```

5. (Optional) Enable unrestricted file access.

```
[edit protocols pim traceoptions]
user@host# set file world-readable
```

6. Configure tracing flags.

Suppose you are troubleshooting issues with PIM version 1 control packets that are received on an interface configured for PIM version 2. The following example shows how to trace messages associated with this problem.

```
[edit protocols pim traceoptions]
user@host# set flag packets | match "Rx V1 Require V2"
```

7. View the trace file.

```
user@host> file list /var/log
user@host> file show /var/log/pim-trace
```

- Related Documentation**
- [PIM Overview on page 5491](#)
  - *Tracing and Logging Junos OS Operations*

---

## Configuring Interface Priority for PIM Designated Router Selection

A designated router (DR) sends periodic join messages and prune messages toward a group-specific rendezvous point (RP) for each group for which it has active members. When a Protocol Independent Multicast (PIM) router learns about a source, it originates a Multicast Source Discovery Protocol (MSDP) source-address message if it is the DR on the upstream interface.

By default, every PIM interface has an equal probability (priority 1) of being selected as the DR. Configuring the interface DR priority helps ensure that changing an IP address does not alter your forwarding model.

To configure the interface designated router priority:

1. This example shows the configuration for the routing instance. Configure the interface globally or in the routing instance.

```
[edit routing-instances PIM.master protocols pim interface ge-0/0/0.0 family inet]
```

```
user@host# set priority 5
```

2. Verify the configuration by checking the **Hello Option DR Priority** field in the output of the **show pim neighbors detail** command.

```
user@host> show pim neighbors detail
```

```
Instance: PIM.master
Interface: ge-0/0/0.0
Address: 192.168.195.37, IPv4, PIM v2, Mode: Sparse
Hello Option Holdtime: 65535 seconds
Hello Option DR Priority: 5
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported
Rx Join: Group Source Timeout
225.1.1.1 192.168.195.78 0
225.1.1.1 0
```

```
Interface: lo0.0
Address: 10.255.245.91, IPv4, PIM v2, Mode: Sparse
Hello Option Holdtime: 65535 seconds
Hello Option DR Priority: 1
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported
```

```
Interface: pd-6/0/0.32768
Address: 0.0.0.0, IPv4, PIM v2, Mode: Sparse
Hello Option Holdtime: 65535 seconds
Hello Option DR Priority: 0
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported
```

#### Related Documentation

- [Configuring PIM Designated Router Election on Point-to-Point Links on page 5499](#)
- [Understanding PIM Sparse Mode on page 5509](#)
- [show pim neighbors on page 5907](#)

## Configuring PIM Designated Router Election on Point-to-Point Links

To comply with the latest PIM drafts, enable designated router (DR) election on all PIM interfaces, including point-to-point (P2P) interfaces. (DR election is enabled by default on all other interfaces.) One of the two routers might join a multicast group on its P2P link interface. The DR on that link is responsible for initiating the relevant join messages.

To enable DR election on point-to-point interfaces:

1. On both point-to-point link routers, configure the router globally or in the routing instance. This example shows the configuration for the routing instance.

```
[edit routing-instances PIM.master protocols pim]
user@host# set dr-election-on-p2p
```

2. Verify the configuration by checking the **State** field in the output of the **show pim interfaces** command. The possible values for the **State** field are DR, NotDR, and P2P. When a point-to-point link interface is elected to be the DR, the interface state becomes DR instead of P2P.

3. If the **show pim interfaces** command continues to report the P2P state, consider running the **restart routing** command on both routers on the point-to-point link. Then recheck the state.



**CAUTION:** Do not restart a software process unless specifically asked to do so by your Juniper Networks customer support representative. Restarting a software process during normal operation of a routing platform could cause interruption of packet forwarding and loss of data.

[edit]  
user@host# run restart routing

**Related  
Documentation**

- [Understanding PIM Sparse Mode on page 5509](#)
- [Configuring Interface Priority for PIM Designated Router Selection on page 5498](#)
- [show pim interfaces on page 5883](#)

---

## Configuring BFD for PIM

The Bidirectional Forwarding Detection (BFD) Protocol is a simple hello mechanism that detects failures in a network. BFD works with a wide variety of network environments and topologies. A pair of routing devices exchanges BFD packets. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the routing device stops receiving a reply after a specified interval. The BFD failure detection timers have shorter time limits than the Protocol Independent Multicast (PIM) hello hold time, so they provide faster detection.

The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the **clear bfd adaptation** command to return BFD interval timers to their configured values. The **clear bfd adaptation** command is hitless, meaning that the command does not affect traffic flow on the routing device.

You must specify the minimum transmit and minimum receive intervals to enable BFD on PIM.

To enable failure detection:

1. Configure the interface globally or in a routing instance.

This example shows the global configuration.

[edit protocols pim]

```
user@host# edit interface fe-1/0/0.0 family inet bfd-liveness-detection
```

2. Configure the minimum transmit interval.

This is the minimum interval after which the routing device transmits hello packets to a neighbor with which it has established a BFD session. Specifying an interval smaller than 300 ms can cause undesired BFD flapping.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set transmit-interval 350
```

3. Configure the minimum interval after which the routing device expects to receive a reply from a neighbor with which it has established a BFD session.

Specifying an interval smaller than 300 ms can cause undesired BFD flapping.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set minimum-receive-interval 350
```

4. (Optional) Configure other BFD settings.

As an alternative to setting the receive and transmit intervals separately, configure one interval for both.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set minimum-interval 350
```

5. Configure the threshold for the adaptation of the BFD session detection time.

When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set detection-time threshold 800
```

6. Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set multiplier 50
```

7. Configure the BFD version.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set version 1
```

8. Specify that BFD sessions should not adapt to changing network conditions.

We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set no-adaptation
```

9. Verify the configuration by checking the output of the **show bfd session** command.

**Related Documentation**

- *show bfd session*

## Configuring BFD Authentication for PIM

---

Beginning with Junos OS Release 9.6, you can configure authentication for Bidirectional Forwarding Detection (BFD) sessions running over Protocol Independent Multicast (PIM). Routing instances are also supported. The following steps are needed to configure authentication on a BFD session:

1. Specify the BFD authentication algorithm for the PIM protocol.
2. Associate the authentication keychain with the PIM protocol.
3. Configure the related security authentication keychain.

The following sections provide instructions for configuring and viewing BFD authentication on PIM:

- [Configuring BFD Authentication Parameters on page 5502](#)
- [Viewing Authentication Information for BFD Sessions on page 5503](#)

### Configuring BFD Authentication Parameters

BFD authentication is only supported in the Canada and United States version of the Junos OS image and is not available in the export version.

To configure BFD authentication:

1. Specify the algorithm (**keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**, or **simple-password**) to use for BFD authentication on a PIM route or routing instance.

```
[edit protocols pim]
```

```
user@host# set interface ge-0/1/5 family inet bfd-liveness-detection authentication
algorithm keyed-sha-1
```



**NOTE:** Nonstop active routing (NSR) is not supported with the **meticulous-keyed-md5** and **meticulous-keyed-sha-1** authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

2. Specify the keychain to be used to associate BFD sessions on the specified PIM route or routing instance with the unique security authentication keychain attributes.

The keychain you specify must match the keychain name configured at the **[edit security authentication key-chains]** hierarchy level.

```
[edit protocols pim]
```

```
user@host# set interface ge-0/1/5 family inet bfd-liveness-detection authentication
keychain bfd-pim
```



**NOTE:** The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

3. Specify the unique security authentication information for BFD sessions:

- The matching keychain name as specified in Step 2.
- At least one key, a unique integer between 0 and 63. Creating multiple keys allows multiple clients to use the BFD session.
- The secret data used to allow access to the session.
- The time at which the authentication key becomes active, in the format *yyyy-mm-dd.hh:mm:ss*.

[edit security]

```
user@host# set authentication-key-chains key-chain bfd-pim key 53 secret
9ggaJDmPQ6/tJgF/AtREVsyPsnCtUhm start-time 2009-06-14.10:00:00
```

4. (Optional) Specify loose authentication checking if you are transitioning from nonauthenticated sessions to authenticated sessions.

[edit protocols pim]

```
user@host# set interface ge-0/1/5 family inet bfd-liveness-detection authentication
loose-check
```

5. (Optional) View your configuration by using the **show bfd session detail** or **show bfd session extensive** command.

6. Repeat these steps to configure the other end of the BFD session.

## Viewing Authentication Information for BFD Sessions

You can view the existing BFD authentication configuration by using the **show bfd session detail** and **show bfd session extensive** commands.

The following example shows BFD authentication configured for the **ge-0/1/5** interface. It specifies the keyed SHA-1 authentication algorithm and a keychain name of **bfd-pim**. The authentication keychain is configured with two keys. Key 1 contains the secret data "**\$9\$ggaJDmPQ6/tJgF/AtREVsyPsnCtUhm**" and a start time of June 1, 2009, at 9:46:02 AM PST. Key 2 contains the secret data "**\$9\$a5jiKW9L.reP38ny.TszF2/9**" and a start time of June 1, 2009, at 3:29:20 PM PST.

[edit protocols pim]

```
interface ge-0/1/5 {
 family inet {
 bfd-liveness-detection {
 authentication {
 key-chain bfd-pim;
 algorithm keyed-sha-1;
 }
 }
 }
}
```

```

}
[edit security]
authentication key-chains {
 key-chain bfd-pim {
 key 1 {
 secret "9ggaJDmPQ6/tJgF/AtREVsyPsnCtUHm";
 start-time "2009-6-1.09:46:02 -0700";
 }
 key 2 {
 secret "9a5jiKW9l.reP38ny.TszF2/9";
 start-time "2009-6-1.15:29:20 -0700";
 }
 }
}
}

```

If you commit these updates to your configuration, you see output similar to the following example. In the output for the **show bfd session detail** command, **Authenticate** is displayed to indicate that BFD authentication is configured. For more information about the configuration, use the **show bfd session extensive** command. The output for this command provides the keychain name, the authentication algorithm and mode for each client in the session, and the overall BFD authentication configuration status, keychain name, and authentication algorithm and mode.

#### show bfd session detail

```
user@host# show bfd session detail
```

| Address                                                               | State | Interface  | Detect Time | Transmit Interval | Multiplier |
|-----------------------------------------------------------------------|-------|------------|-------------|-------------------|------------|
| 50.0.0.2                                                              | Up    | ge-0/1/5.0 | 0.900       | 0.300             | 3          |
| Client PIM, TX interval 0.300, RX interval 0.300, <b>Authenticate</b> |       |            |             |                   |            |
| Session up time 3d 00:34                                              |       |            |             |                   |            |
| Local diagnostic None, remote diagnostic NbrSignal                    |       |            |             |                   |            |
| Remote state Up, version 1                                            |       |            |             |                   |            |
| Replicated                                                            |       |            |             |                   |            |

#### show bfd session extensive

```
user@host# show bfd session extensive
```

| Address                                                                               | State | Interface  | Detect Time | Transmit Interval | Multiplier |
|---------------------------------------------------------------------------------------|-------|------------|-------------|-------------------|------------|
| 50.0.0.2                                                                              | Up    | ge-0/1/5.0 | 0.900       | 0.300             | 3          |
| Client PIM, TX interval 0.300, RX interval 0.300, <b>Authenticate</b>                 |       |            |             |                   |            |
| keychain bfd-pim, algo keyed-sha-1, mode strict                                       |       |            |             |                   |            |
| Session up time 00:04:42                                                              |       |            |             |                   |            |
| Local diagnostic None, remote diagnostic NbrSignal                                    |       |            |             |                   |            |
| Remote state Up, version 1                                                            |       |            |             |                   |            |
| Replicated                                                                            |       |            |             |                   |            |
| Min async interval 0.300, min slow interval 1.000                                     |       |            |             |                   |            |
| Adaptive async TX interval 0.300, RX interval 0.300                                   |       |            |             |                   |            |
| Local min TX interval 0.300, minimum RX interval 0.300, multiplier 3                  |       |            |             |                   |            |
| Remote min TX interval 0.300, min RX interval 0.300, multiplier 3                     |       |            |             |                   |            |
| Local discriminator 2, remote discriminator 2                                         |       |            |             |                   |            |
| Echo mode disabled/inactive                                                           |       |            |             |                   |            |
| <b>Authentication enabled/active, keychain bfd-pim, algo keyed-sha-1, mode strict</b> |       |            |             |                   |            |

#### Related Documentation

- *Understanding Bidirectional Forwarding Detection Authentication for PIM*



- [Configuring BFD for PIM on page 5500](#)
- [authentication-key-chains on page 6298](#)
- [bfd-liveness-detection on page 5697](#)
- *show bfd session*

## Disabling PIM

By default, when configured, the PIM protocol is enabled on all interfaces for all families. If desired, you can disable PIM at the protocol, interface, or family hierarchy levels.

The hierarchy in which you configure PIM is critical. In general, the most specific configuration takes precedence. However, if PIM is disabled at the protocol level, then any disable statements with respect to an interface or family are ignored.

For example, the order of precedence for disabling PIM on a particular interface family is:

1. If PIM is disabled at the **[edit protocols pim interface *interface-name* family]** hierarchy level, then PIM is disabled for that interface family.
2. If PIM is not configured at the **[edit protocols pim interface *interface-name* family]** hierarchy level, but is disabled at the **[edit protocols pim interface *interface-name*]** hierarchy level, then PIM is disabled for all families on the specified interface.
3. If PIM is not configured at either the **[edit protocols pim interface *interface-name* family]** hierarchy level or the **[edit protocols pim interface *interface-name*]** hierarchy level, but is disabled at the **[edit protocols pim]** hierarchy level, then the PIM protocol is disabled globally for all interfaces and all families.

The following sections describe how to disable PIM at the various hierarchy levels.

- [Disabling the PIM Protocol on page 5505](#)
- [Disabling PIM on an Interface on page 5506](#)
- [Disabling PIM for a Family on page 5506](#)
- [Disabling PIM for a Rendezvous Point on page 5507](#)

### Disabling the PIM Protocol

You can explicitly disable the PIM protocol. Disabling the PIM protocol disables the protocol for all interfaces and all families. This is accomplished at the **[edit protocols pim]** hierarchy level:

```
[edit protocols]
pim {
 disable;
}
```

To disable the PIM protocol:

1. Include the **disable** statement.

```
user@host# set protocols pim disable
```

2. (Optional) Verify your configuration settings before committing them by using the **show protocols pim** command.

```
user@host# run show protocols pim
```

## Disabling PIM on an Interface

You can disable the PIM protocol on a per-interface basis. This is accomplished at the **[edit protocols pim interface *interface-name*]** hierarchy level:

```
[edit protocols]
pim {
 interface interface-name {
 disable;
 }
}
```

To disable PIM on an interface:

1. Include the **disable** statement.

```
user@host# set protocols pim interface fe-0/1/0 disable
```

2. (Optional) Verify your configuration settings before committing them by using the **show protocols pim** command.

```
user@host# run show protocols pim
```

## Disabling PIM for a Family

You can disable the PIM protocol on a per-family basis. This is accomplished at the **[edit protocols pim family]** hierarchy level:

```
[edit protocols]
pim {
 family inet {
 disable;
 }
 family inet6 {
 disable;
 }
}
```

To disable PIM for a family:

1. Include the **disable** statement.

```
user@host# set protocols pim family inet disable
```

```
user@host# set protocols pim family inet6 disable
```

2. (Optional) Verify your configuration settings before committing them by using the **show protocols pim** command.

```
user@host# run show protocols pim
```

## Disabling PIM for a Rendezvous Point

You can disable the PIM protocol for a rendezvous point (RP) on a per-family basis. This is accomplished at the **[edit protocols pim rp local family]** hierarchy level:

```
[edit protocols]
pim {
 rp {
 local {
 family inet {
 disable;
 }
 family inet6 {
 disable;
 }
 }
 }
}
```

To disable PIM for an RP family:

1. Use the **disable** statement.

```
user@host# set protocols pim rp local family inet disable
user@host# set protocols pim rp local family inet6 disable
```

2. (Optional) Verify your configuration settings before committing them by using the **show protocols pim** command.

```
user@host# run show protocols pim
```



# Using PIM Sparse Mode

- [Understanding PIM Sparse Mode on page 5509](#)
- [Designated Router on page 5512](#)
- [Enabling PIM Sparse Mode on page 5512](#)
- [Configuring PIM Join Load Balancing on page 5513](#)
- [Modifying the Join State Timeout on page 5517](#)
- [Example: Enabling Join Suppression on page 5517](#)

## Understanding PIM Sparse Mode

---

A Protocol Independent Multicast (PIM) sparse-mode domain uses reverse-path forwarding (RPF) to create a path from a data source to the receiver requesting the data. When a receiver issues an explicit join request, an RPF check is triggered. A (\*,G) PIM join message is sent toward the RP from the receiver's designated router (DR). (By definition, this message is actually called a join/prune message, but for clarity in this description, it is called either join or prune, depending on its context.) The join message is multicast hop by hop upstream to the ALL-PIM-ROUTERS group (224.0.0.13) by means of each router's RPF interface until it reaches the RP. The RP router receives the (\*,G) PIM join message and adds the interface on which it was received to the outgoing interface list (OIL) of the rendezvous-point tree (RPT) forwarding state entry. This builds the RPT connecting the receiver with the RP. The RPT remains in effect, even if no active sources generate traffic.



**NOTE:** State—the (\*,G) or (S,G) entries—is the information used for forwarding unicast or multicast packets. S is the source IP address, G is the multicast group address, and \* represents any source sending to group G. Routers keep track of the multicast forwarding state for the incoming and outgoing interfaces for each group.

When a source becomes active, the source DR encapsulates multicast data packets into a PIM register message and sends them by means of unicast to the RP router.

If the RP router has interested receivers in the PIM sparse-mode domain, it sends a PIM join message toward the source to build a shortest-path tree (SPT) back to the source. The source sends multicast packets out on the LAN, and the source DR encapsulates

the packets in a PIM register message and forwards the message toward the RP router by means of unicast. The RP router receives PIM register messages back from the source, and thus adds a new source to the distribution tree, keeping track of sources in a PIM table. Once an RP router receives packets natively (with S,G), it sends a register stop message to stop receiving the register messages by means of unicast.

In actual application, many receivers with multiple SPTs are involved in a multicast traffic flow. To illustrate the process, we track the multicast traffic from the RP router to one receiver. In such a case, the RP router begins sending multicast packets down the RPT toward the receiver's DR for delivery to the interested receivers. When the receiver's DR receives the first packet from the RPT, the DR sends a PIM join message toward the source DR to start building an SPT back to the source. When the source DR receives the PIM join message from the receiver's DR, it starts sending traffic down all SPTs. When the first multicast packet is received by the receiver's DR, the receiver's DR sends a PIM prune message to the RP router to stop duplicate packets from being sent through the RPT. In turn, the RP router stops sending multicast packets to the receiver's DR, and sends a PIM prune message for this source over the RPT toward the source DR to halt multicast packet delivery to the RP router from that particular source.

If the RP router receives a PIM register message from an active source but has no interested receivers in the PIM sparse-mode domain, it still adds the active source into the PIM table. However, after adding the active source into the PIM table, the RP router sends a register stop message. The RP router discovers the active source's existence and no longer needs to receive advertisement of the source (which utilizes resources).



**NOTE:** If the number of PIM join messages exceeds the configured MTU, the messages are fragmented in IPv6 PIM sparse mode. To avoid the fragmentation of PIM join messages, the multicast traffic receives the interface MTU instead of the path MTU.

---

The major characteristics of PIM sparse mode are as follows:

- Routers with downstream receivers join a PIM sparse-mode tree through an explicit join message.
- PIM sparse-mode RPs are the routers where receivers meet sources.
- Senders announce their existence to one or more RPs, and receivers query RPs to find multicast sessions.
- Once receivers get content from sources through the RP, the last-hop router (the router closest to the receiver) can optionally remove the RP from the shared distribution tree (\*,G) if the new source-based tree (S,G) is shorter. Receivers can then get content directly from the source.

The transitional aspect of PIM sparse mode from shared to source-based tree is one of the major features of PIM, because it prevents overloading the RP or surrounding core links.

There are related issues regarding source, RPs, and receivers when sparse mode multicast is used:

- Sources must be able to send to all RPs.
- RPs must all know one another.
- Receivers must send explicit join messages to a known RP.
- Receivers initially need to know only one RP (they later learn about others).
- Receivers can explicitly prune themselves from a tree.
- Receivers that never transition to a source-based tree are effectively running Core Based Trees (CBT).


PIM sparse mode has standard features for all of these issues.

## Rendezvous Point

The RP router serves as the information exchange point for the other routers. All routers in a PIM domain must provide mapping to an RP router. It is the only router that needs to know the active sources for a domain—the other routers just need to know how to reach the RP. In this way, the RP matches receivers with sources.

The RP router is downstream from the source and forms one end of the shortest-path tree. As shown in [Figure 172](#), the RP router is upstream from the receiver and thus forms one end of the rendezvous-point tree.

**Figure 172: Rendezvous Point As Part of the RPT and SPT**



The benefit of using the RP as the information exchange point is that it reduces the amount of state in non-RP routers. No network flooding is required to provide non-RP routers information about active sources.

## RP Mapping Options

RPs can be learned by one of the following mechanisms:

- Static configuration
- Anycast RP
- Auto-RP
- Bootstrap router

We recommend a static RP mapping with anycast RP and a bootstrap router (BSR) with auto-RP configuration, because static mapping provides all the benefits of a bootstrap router and auto-RP without the complexity of the full BSR and auto-RP mechanisms.

- Related Documentation**
- [Understanding Static RP on page 5545](#)
  - [Understanding RP Mapping with Anycast RP on page 5551](#)
  - [Understanding the PIM Bootstrap Router on page 5567](#)
  - [Understanding PIM Auto-RP on page 5561](#)

---

## Designated Router

In a PIM sparse mode (PIM-SM) domain, there are two types of designated routers to consider:

- The receiver DR sends PIM join and PIM prune messages from the receiver network toward the RP.
- The source DR sends PIM register messages from the source network to the RP.

Neighboring PIM routers multicast periodic PIM hello messages to each other every 30 seconds (the default). The PIM hello message usually includes a holdtime value for the neighbor to use, but this is not a requirement. If the PIM hello message does not include a holdtime value, a default timeout value (in Junos OS, 105 seconds) is used. On receipt of a PIM hello message, a router stores the IP address and priority for that neighbor. If the DR priorities match, the router with the highest IP address is selected as the DR.

If a DR fails, a new one is selected using the same process of comparing IP addresses.



**NOTE:** In PIM dense mode (PIM-DM), a DR is elected by the same process that PIM-SM uses. However, the only time that a DR has any effect in PIM-DM is when IGMPv1 is used on the interface. (IGMPv2 is the default.) In this case, the DR also functions as the IGMP Query Router because IGMPv1 does not have a Query Router election mechanism.

---

---

## Enabling PIM Sparse Mode

In PIM sparse mode (PIM-SM), the assumption is that very few of the possible receivers want packets from a source, so the network establishes and sends packets only on branches that have at least one leaf indicating (by message) a desire for the traffic. WANs are appropriate networks for sparse-mode operation.

By default, PIM is disabled. When you enable PIM, it operates in sparse mode by default. You do not need to configure Internet Group Management Protocol (IGMP) version 2 for a sparse mode configuration. After you enable PIM, by default, IGMP version 2 is also enabled.



All systems on a subnet must run the same version of PIM.

The default PIM version can be version 1 or version 2, depending on the mode you are configuring. PIMv1 is the default for rendezvous point (RP) mode (at the **[edit protocols pim rp static address address]** hierarchy level). However, PIMv2 is the default for interface mode (at the **[edit protocols pim interface interface-name]** hierarchy level). Explicitly configured versions override the defaults. The following example explicitly configures PIMv2 on the interfaces.

You can configure PIM sparse mode globally or for a routing instance. This example shows how to configure PIM sparse mode globally on all interfaces. It also shows how to configure a static RP router and how to configure the non-RP routers.

To configure the router properties for PIM sparse mode:

1. Configure the static RP router.

```
[edit protocols pim]
user@host# set rp local family inet address 192.168.3.253
```

2. Configure the RP router interfaces. When configuring all interfaces, exclude the **fxp0.0** management interface by including the **disable** statement for that interface.

```
[edit protocols pim]
user@host# set interface all mode sparse
user@host# set interface all version 2
user@host# set interface fxp0.0 disable
```

3. Configure the non-RP routers. Include the following configuration on all of the non-RP routers.

```
[edit protocols pim]
user@host# set rp static address 192.168.3.253 version 2
user@host# set interface all mode sparse
user@host# set interface all version 2
user@host# set interface fxp0.0 disable
```

4. Monitor the operation of PIM sparse mode.

- [show pim interfaces](#)
- [show pim join](#)
- [show pim neighbors](#)
- [show pim rps](#)

**Related Documentation**

- [Understanding PIM Sparse Mode on page 5509](#)

## Configuring PIM Join Load Balancing

By default, PIM join messages are sent toward a source based on the RPF routing table check. If there is more than one equal-cost path toward the source, then one upstream interface is chosen to send the join message. This interface is also used for all downstream

traffic, so even though there are alternative interfaces available, the multicast load is concentrated on one upstream interface and routing device.

For PIM sparse mode, you can configure PIM join load balancing to spread join messages and traffic across equal-cost upstream paths (interfaces and routing devices) provided by unicast routing toward a source. PIM join load balancing is only supported for PIM sparse mode configurations.

PIM join load balancing is supported on draft-rosen multicast VPNs (also referred to as dual PIM multicast VPNs). PIM join load balancing is not supported on multiprotocol BGP-based multicast VPNs (also referred to as next-generation Layer 3 VPN multicast). When PIM join load balancing is enabled in a draft-rosen Layer 3 VPN scenario, the load balancing is achieved based on the join counts for the far-end PE routing devices, not for any intermediate P routing devices.

If an internal BGP (IBGP) multipath forwarding VPN route is available, the Junos OS uses the multipath forwarding VPN route to send join messages to the remote PE routers to achieve load balancing over the VPN.

By default, when multiple PIM joins are received for different groups, all joins are sent to the same upstream gateway chosen by the unicast routing protocol. Even if there are multiple equal-cost paths available, these alternative paths are not utilized to distribute multicast traffic from the source to the various groups.

When PIM join load balancing is configured, the PIM joins are distributed equally among all equal-cost upstream interfaces and neighbors. Every new join triggers the selection of the least-loaded upstream interface and neighbor. If there are multiple neighbors on the same interface (for example, on a LAN), join load balancing maintains a value for each of the neighbors and distributes multicast joins (and downstream traffic) among these as well.

Join counts for interfaces and neighbors are maintained globally, not on a per-source basis. Therefore, there is no guarantee that joins for a particular source are load-balanced. However, the joins for all sources and all groups known to the routing device are load-balanced. There is also no way to administratively give preference to one neighbor over another: all equal-cost paths are treated the same way.

You can configure message filtering globally or for a routing instance. This example shows the global configuration.

You configure PIM join load balancing on the non-RP routers in the PIM domain.

1. Determine if there are multiple paths available for a source (for example, an RP) with the output of the **show pim join extensive** or **show pim source** commands.

```
user@host> show pim join extensive
Instance: PIM.master Family: INET

Group: 224.1.1.1
 Source: *
 RP: 10.255.245.6
 Flags: sparse,rptree,wildcard
 Upstream interface: t1-0/2/3.0
 Upstream neighbor: 192.168.38.57
 Upstream state: Join to RP
 Downstream neighbors:
 Interface: t1-0/2/1.0
 192.168.38.16 State: JOIN Flags; SRW Timeout: 164
Group: 224.2.127.254
 Source: *
 RP: 10.255.245.6
 Flags: sparse,rptree,wildcard
 Upstream interface: so-0/3/0.0
 Upstream neighbor: 192.168.38.47
 Upstream state: Join to RP
 Downstream neighbors:
 Interface: t1-0/2/3.0
 192.168.38.16 State: JOIN Flags; SRW Timeout: 164
```

Note that for this router, the RP at IP address 10.255.245.6 is the source for two multicast groups: 224.1.1.1 and 224.2.127.254. This router has two equal-cost paths through two different upstream interfaces (**t1-0/2/3.0** and **so-0/3/0.0**) with two different neighbors (192.168.38.57 and 192.168.38.47). This router is a good candidate for PIM join load balancing.

2. On the non-RP router, configure PIM sparse mode and join load balancing.

```
[edit protocols pim]
user@host# set interface all mode sparse version 2
user@host# set join-load-balance
```

3. Then configure the static address of the RP.

```
[edit protocols pim rp]
user@host# set static address 10.10.10.1
```

4. Monitor the operation.

If load balancing is enabled for this router, the number of PIM joins sent on each interface is shown in the output for the **show pim interfaces** command.

```
user@host> show pim interfaces
Instance: PIM.master
```

| Name           | Stat | Mode   | IP V | State | NbrCnt | JoinCnt | DR address         |
|----------------|------|--------|------|-------|--------|---------|--------------------|
| lo0.0          | Up   | Sparse | 4 2  | DR    | 0      | 0       | 10.255.168.58      |
| pe-1/2/0.32769 | Up   | Sparse | 4 2  | P2P   | 0      | 0       |                    |
| so-0/3/0.0     | Up   | Sparse | 4 2  | P2P   | 1      | 1       |                    |
| t1-0/2/1.0     | Up   | Sparse | 4 2  | P2P   | 1      | 0       |                    |
| t1-0/2/3.0     | Up   | Sparse | 4 2  | P2P   | 1      | 1       |                    |
| lo0.0          | Up   | Sparse | 6 2  | DR    | 0      | 0       | fe80::2a0:a5ff:4b7 |

Note that the two equal-cost paths shown by the **show pim interfaces** command now have nonzero join counts. If the counts differ by more than one and were zero (0) when load balancing commenced, an error occurs (joins before load balancing are not redistributed). The join count also appears in the **show pim neighbors detail** output:

```
user@host> show pim neighbors detail
Interface: so-0/3/0.0
```

```
Address: 192.168.38.46, IPv4, PIM v2, Mode: Sparse, Join Count: 0
Hello Option Holdtime: 65535 seconds
Hello Option DR Priority: 1
Hello Option Generation ID: 1689116164
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

```
Address: 192.168.38.47, IPv4, PIM v2, Join Count: 1
BFD: Disabled
Hello Option Holdtime: 105 seconds 102 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 792890329
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

```
Interface: t1-0/2/3.0
```

```
Address: 192.168.38.56, IPv4, PIM v2, Mode: Sparse, Join Count: 0
Hello Option Holdtime: 65535 seconds
Hello Option DR Priority: 1
Hello Option Generation ID: 678582286
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

```
Address: 192.168.38.57, IPv4, PIM v2, Join Count: 1
BFD: Disabled
Hello Option Holdtime: 105 seconds 97 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 1854475503
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

Note that the join count is nonzero on the two load-balanced interfaces toward the upstream neighbors.

PIM join load balancing only takes effect when the feature is configured. Prior joins are not redistributed to achieve perfect load balancing. In addition, if an interface or neighbor fails, the new joins are redistributed among remaining active interfaces and neighbors. However, when the interface or neighbor is restored, prior joins are not redistributed. The **clear pim join-distribution** command redistributes the existing flows to new or restored upstream neighbors. Redistributing the existing flows causes traffic to be disrupted, so we recommend that you perform PIM join redistribution during a maintenance window.

- Related Documentation**
- [clear pim join-distribution](#)
  - [show pim interfaces on page 5883](#)
  - [show pim neighbors on page 5907](#)
  - [show pim source on page 5918](#)

---

## Modifying the Join State Timeout

---

This section describes how to configure the join state timeout.

A downstream router periodically sends join messages to refresh the join state on the upstream router. If the join state is not refreshed before the timeout expires, the join state is removed.

By default, the join state timeout is 210 seconds. You can change this timeout to allow additional time to receive the join messages. Because the messages are called join-prune messages, the name used is the **join-prune-timeout** statement.

To modify the timeout, include the **join-prune-timeout** statement:

```
user@host# set protocols pim join-prune-timeout 230
```

The join timeout value can be from 210 through 420 seconds.

### Related Documentation

- [join-prune-timeout on page 5721](#)

---

## Example: Enabling Join Suppression

---

This example describes how to enable PIM join suppression.

- [Requirements on page 5517](#)
- [Overview on page 5517](#)
- [Configuration on page 5520](#)
- [Verification on page 5521](#)

### Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Routing Devices*.
- Configure PIM Sparse Mode on the interfaces. See [“Enabling PIM Sparse Mode” on page 5512](#).

### Overview

PIM join suppression enables a router on a multiaccess network to defer sending join messages to an upstream router when it sees identical join messages on the same network. Eventually, only one router sends these join messages, and the other routers suppress identical messages. Limiting the number of join messages improves scalability and efficiency by reducing the number of messages sent to the same router.


This example includes the following statements:

- **override-interval**—Sets the maximum time in milliseconds to delay sending override join messages. When a router sees a prune message for a join it is currently suppressing, it waits before it sends an override join message. Waiting helps avoid multiple downstream routers sending override join messages at the same time. The override interval is a random timer with a value of 0 through the maximum override value.
- **propagation-delay**—Sets a value in milliseconds for a prune pending timer, which specifies how long to wait before executing a prune on an upstream router. During this period, the router waits for any prune override join messages that might be currently suppressed. The period for the prune pending timer is the sum of the **override-interval** value and the value specified for **propagation-delay**.
- **reset-tracking-bit**—Enables PIM join suppression on each multiaccess downstream interface. This statement resets a tracking bit field (T-bit) on the LAN prune delay hello option from the default of 1 (join suppression disabled) to 0 (join suppression enabled).

When multiple identical join messages are received, a random join suppression timer is activated, with a range of 66 through 84 milliseconds. The timer is reset each time join suppression is triggered.

[Figure 173](#) shows the topology used in this example.

Figure 173: Join Suppression



The items in [Figure 173](#) represent the following functions:

- Host 0 is the multicast source.
- Host 1, Host 2, Host 3, and Host 4 are receivers.
- Router R0 is the first-hop router and the RP.
- Router R1 is an upstream router.
- Routers R2, R3, R4, and R5 are downstream routers in the multicast LAN.

This example shows the configuration of the downstream devices: Routers R2, R3, R4, and R5.

## Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
[edit]
set protocols pim traceoptions file pim.log
set protocols pim traceoptions file size 5m
set protocols pim traceoptions file world-readable
set protocols pim traceoptions flag join detail
set protocols pim traceoptions flag prune detail
set protocols pim traceoptions flag normal detail
set protocols pim traceoptions flag register detail
set protocols pim rp static address 10.255.112.160
set protocols pim interface all mode sparse
set protocols pim interface all version 2
set protocols pim interface fxp0.0 disable
set protocols pim reset-tracking-bit
set protocols pim propagation-delay 500
set protocols pim override-interval 4000
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure PIM join suppression on a non-RP downstream router in the multicast LAN:

1. Configure PIM sparse mode on the interfaces.

```
[edit]
user@host# edit protocols pim
[edit protocols pim]
user@host# set rp static address 10.255.112.160
[edit protocols pim]
user@host# set interface all mode sparse version 2
[edit protocols pim]
user@host# set interface all version 2
[edit protocols pim]
user@host# set interface fxp0.0 disable
```

2. Enable the join suppression timer.

```
[edit protocols pim]
user@host# set reset-tracking-bit
```

3. Configure the prune override interval value.

```
[edit protocols pim]
user@host# set override-interval 4000
```

4. Configure the propagation delay of the link.

```
[edit protocols pim]
user@host# set propagation-delay 500
```



5. (Optional) Configure PIM tracing operations.

```
[edit protocols pim]
user@host# set traceoptions file pim.log size 5m world-readable
[edit protocols pim]
user@host# set traceoptions flag join detail
[edit protocols pim]
user@host# set traceoptions flag normal detail
[edit protocols pim]
user@host# set traceoptions flag register detail
```

6. If you are done configuring the device, commit the configuration.

```
[edit protocols pim]
user@host# commit
```

## Results

From configuration mode, confirm your configuration by entering the **show protocols** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols
pim {
 traceoptions {
 file pim.log size 5m world-readable;
 flag join detail;
 flag prune detail;
 flag normal detail;
 flag register detail;
 }
 rp {
 static {
 address 10.255.112.160;
 }
 }
 interface all {
 mode sparse;
 version 2;
 }
 interface fxp0.0 {
 disable;
 }
 reset-tracking-bit;
 propagation-delay 500;
 override-interval 4000;
}
```

## Verification

To verify the configuration, run the following commands on the upstream and downstream routers:

- **show pim join extensive**

- [show multicast route extensive](#)

**Related  
Documentation**

- [Example: Configuring the PIM Assert Timeout on page 5588](#)
- [Example: Configuring PIM RPF Selection](#)
- [Example: Configuring the PIM SPT Threshold Policy on page 5590](#)
- [Enabling PIM Sparse Mode on page 5512](#)
- [PIM Overview on page 5491](#)

# Using PIM Dense Mode and PIM Sparse-Dense Mode

- [Understanding PIM Dense Mode on page 5523](#)
- [Understanding PIM Sparse-Dense Mode on page 5525](#)
- [Mixing PIM Sparse and Dense Modes on page 5525](#)
- [Configuring PIM Dense Mode Properties on page 5526](#)
- [Configuring PIM Sparse-Dense Mode Properties on page 5527](#)

## Understanding PIM Dense Mode

---

PIM dense mode is less sophisticated than PIM sparse mode. PIM dense mode is useful for multicast LAN applications, the main environment for all dense mode protocols.

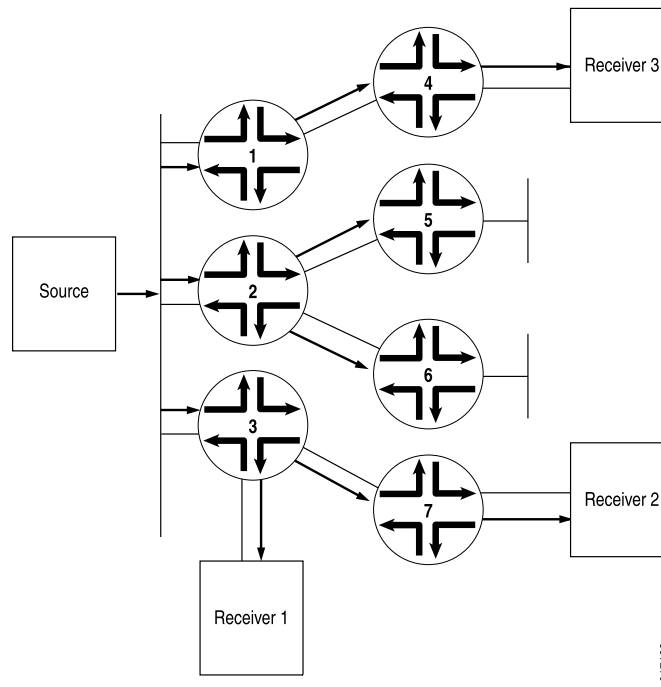
PIM dense mode implements the same flood-and-prune mechanism that DVMRP and other dense mode routing protocols employ. The main difference between DVMRP and PIM dense mode is that PIM dense mode introduces the concept of protocol independence. PIM dense mode can use the routing table populated by any underlying unicast routing protocol to perform reverse-path-forwarding (RPF) checks.

Internet service providers (ISPs) typically appreciate the ability to use any underlying unicast routing protocol with PIM dense mode because they do not need to introduce and manage a separate routing protocol just for RPF checks. While unicast routing protocols extended as multiprotocol BGP (MBGP) and Multitopology Routing in IS-IS (M-IS-IS) were later employed to build special tables to perform RPF checks, PIM dense mode does not require them.

PIM dense mode can use the unicast routing table populated by OSPF, IS-IS, BGP, and so on, or PIM dense mode can be configured to use a special multicast RPF table populated by MBGP or M-IS-IS when performing RPF checks.


Unlike sparse mode, in which data is forwarded only to routers sending an explicit request, dense mode implements a *flood-and-prune* mechanism, similar to DVMRP. In PIM dense mode, there is no RP. A router receives the multicast data on the interface closest to the source, then forwards the traffic to all other interfaces (see [Figure 174](#)).

**Figure 174: Multicast Traffic Flooded from the Source Using PIM Dense Mode**



Flooding occurs periodically. It is used to refresh state information, such as the source IP address and multicast group pair. If the router has no interested receivers for the data, and the OIL becomes empty, the router sends a prune message upstream to stop delivery of multicast traffic (see [Figure 175](#)).

**Figure 175: Prune Messages Sent Back to the Source to Stop Unwanted Multicast Traffic**



## Understanding PIM Sparse-Dense Mode

Sparse-dense mode, as the name implies, allows the interface to operate on a per-group basis in either sparse or dense mode. A group specified as dense is not mapped to an RP. Instead, data packets destined for that group are forwarded by means of PIM dense-mode rules. A group specified as sparse is mapped to an RP, and data packets are forwarded by means of PIM sparse-mode rules.

For information about PIM sparse-mode and PIM dense-mode rules, see [“Understanding PIM Sparse Mode” on page 5509](#) and [“Understanding PIM Dense Mode” on page 5523](#).

- Related Documentation**
- [Understanding PIM Sparse Mode on page 5509](#)
  - [Understanding PIM Dense Mode on page 5523](#)

## Mixing PIM Sparse and Dense Modes

It is possible to mix PIM dense mode, PIM sparse mode, and PIM source-specific multicast (SSM) on the same network, the same router, and even the same interface. This is because modes are effectively tied to multicast groups, an IP multicast group address must be unique for a particular group's traffic, and scoping limits enforce the division between potential or actual overlaps.



**NOTE:** PIM sparse mode was capable of forming shortest-path trees (SPTs) already. Changes to PIM sparse mode to support PIM SSM mainly involved defining behavior in the SSM address range, because shared-tree behavior is prohibited for groups in the SSM address range.

A multicast router employing sparse-dense mode is a good example of mixing PIM modes on the same network or router or interface. Dense modes are easy to support because of the flooding, but scaling issues make dense modes inappropriate for Internet use beyond very restricted uses.

---

## Configuring PIM Dense Mode Properties

---

In PIM dense mode (PIM-DM), the assumption is that almost all possible subnets have at least one receiver wanting to receive the multicast traffic from a source, so the network is flooded with traffic on all possible branches, then pruned back when branches do not express an interest in receiving the packets, explicitly (by message) or implicitly (time-out silence). LANs are appropriate networks for dense-mode operation.

By default, PIM is disabled. When you enable PIM, it operates in sparse mode by default.

You can configure PIM dense mode globally or for a routing instance. This example shows how to configure the routing instance and how to specify that PIM dense mode use **inet.2** as its RPF routing table instead of **inet.0**.

To configure the router properties for PIM dense mode:

1. (Optional) Create an IPv4 routing table group so that interface routes are installed into two routing tables, **inet.0** and **inet.2**.

```
[edit routing-options rib-groups]
user@host# set pim-rg export-rib inet.0
user@host# set pim-rg import-rib [inet.0 inet.2]
```

2. (Optional) Associate the routing table group with a PIM routing instance.

```
[edit routing-instances PIM.dense protocols pim]
user@host# set rib-group inet pim-rg
```

3. Configure the PIM interface. If you do not specify any interfaces, PIM is enabled on all router interfaces. Generally, you specify interface names only if you are disabling PIM on certain interfaces.

```
[edit routing-instances PIM.dense protocols pim]
user@host# set interface fe-0/0/1.0 mode dense
```



**NOTE:** You cannot configure both PIM and Distance Vector Multicast Routing Protocol (DVMRP) in forwarding mode on the same interface. You can configure PIM on the same interface only if you configured DVMRP in unicast-routing mode.

4. Monitor the operation of PIM dense mode by running the **show pim interfaces**, **show pim join**, **show pim neighbors**, and **show pim statistics** commands.

**Related Documentation**

- [Understanding PIM Dense Mode on page 5523](#)
- [Example: Configuring a Dedicated PIM RPF Routing Table](#)

## Configuring PIM Sparse-Dense Mode Properties

Sparse-dense mode allows the interface to operate on a per-group basis in either sparse or dense mode. A group specified as “dense” is not mapped to an RP. Instead, data packets destined for that group are forwarded by means of PIM dense mode rules. A group specified as “sparse” is mapped to an RP, and data packets are forwarded by means of PIM sparse-mode rules. Sparse-dense mode is useful in networks implementing auto-RP for PIM sparse mode.

By default, PIM is disabled. When you enable PIM, it operates in sparse mode by default.

You can configure PIM sparse-dense mode globally or for a routing instance. This example shows how to configure PIM sparse-dense mode globally on all interfaces, specifying that the groups 224.0.1.39 and 224.0.1.40 are using dense mode.

To configure the router properties for PIM sparse-dense mode:

1. Configure the dense-mode groups.

```
[protocols pim]
user@host# set dense-groups 224.0.1.39
user@host# set dense-groups 224.0.1.40
```

2. Configure all interfaces on the routing device to use sparse-dense mode. When configuring all interfaces, exclude the **fxp0.0** management interface by adding the **disable** statement for that interface.

```
[edit protocols pim]
user@host# set interface all mode sparse-dense
user@host# set interface fxp0.0 disable
```

3. Monitor the operation of PIM sparse-dense mode by running the **show pim interfaces**, **show pim join**, **show pim neighbors**, and **show pim statistics** commands.

**Related Documentation**

- [Understanding PIM Sparse-Dense Mode on page 5525](#)





# Using Source-Specific Multicast

- [Source-Specific Multicast Groups Overview on page 5529](#)
- [Understanding PIM SSM on page 5530](#)
- [PIM SSM on page 5531](#)
- [Example: Configuring PIM SSM on a Network on page 5533](#)
- [Example: Configuring an SSM-Only Domain on page 5534](#)
- [Example: Configuring SSM Mapping on page 5535](#)
- [Example: Configuring Source-Specific Multicast Groups with Any-Source Override on page 5537](#)
- [Example: Configuring SSM Maps for Different Groups to Different Sources on page 5541](#)

## Source-Specific Multicast Groups Overview

---

Source-specific multicast (SSM) is a service model that identifies session traffic by both source and group address. SSM implemented in Junos OS has the efficient explicit join procedures of Protocol Independent Multicast (PIM) sparse mode but eliminates the immediate shared tree and rendezvous point (RP) procedures using (\*,G) pairs. The (\*) is a wildcard referring to any source sending to group G, and "G" refers to the IP multicast group. SSM builds shortest-path trees (SPTs) directly represented by (S,G) pairs. The "S" refers to the source's unicast IP address, and the "G" refers to the specific multicast group address. The SSM (S,G) pairs are called channels to differentiate them from any-source multicast (ASM) groups. Although ASM supports both one-to-many and many-to-many communications, ASM's complexity is in its method of source discovery. For example, if you click a link in a browser, the receiver is notified about the group information, but not the source information. With SSM, the client receives both source and group information.

SSM is ideal for one-to-many multicast services such as network entertainment channels. However, many-to-many multicast services might require ASM.

To deploy SSM successfully, you need an end-to-end multicast-enabled network and applications that use an Internet Group Management Protocol version 3 (IGMPv3) or Multicast Listener Discovery version 2 (MLDv2) stack, or you need to configure SSM mapping from IGMPv1 or IGMPv2 to IGMPv3. An IGMPv3 stack provides the capability of a host operating system to use the IGMPv3 protocol. IGMPv3 is available for Windows XP, Windows Vista, and most UNIX operating systems.

SSM mapping allows operators to support an SSM network without requiring all hosts to support IGMPv3. This support exists in static (S,G) configurations, but SSM mapping also supports dynamic per-source group state information, which changes as hosts join and leave the group using IGMP.

SSM is typically supported with a subset of IGMPv3 and PIM sparse mode known as *PIM SSM*. Using SSM, a client can receive multicast traffic directly from the source. PIM SSM uses the PIM sparse-mode functionality to create an SPT between the client and the source, but builds the SPT without the help of an RP.

An SSM-configured network has distinct advantages over a traditionally configured PIM sparse-mode network. There is no need for shared trees or RP mapping (no RP is required), or for RP-to-RP source discovery through the Multicast Source Discovery Protocol (MSDP).

---

## Understanding PIM SSM

RFC 1112, the original multicast RFC, supported both many-to-many and one-to-many models. These came to be known collectively as any-source multicast (ASM) because ASM allowed one or many sources for a multicast group's traffic. However, an ASM network must be able to determine the locations of all sources for a particular multicast group whenever there are interested listeners, no matter where the sources might be located in the network. In ASM, the key function of *source discovery* is a required function of the network itself.

Multicast source discovery appears to be an easy process, but in sparse mode it is not. In dense mode, it is simple enough to flood traffic to every router in the whole network so that every router learns the source address of the content for that multicast group. However, the flooding presents scalability and network resource use issues and is not a viable option in sparse mode.

PIM sparse mode (like any sparse mode protocol) achieves the required source discovery functionality without flooding at the cost of a considerable amount of complexity. The RP routers must be added and must know all multicast sources, and complicated shared distribution trees must be built to the RPs.

In an environment where many sources come and go, such as for a videoconferencing service, ASM is appropriate. However, by ignoring the many-to-many model and focusing attention on the one-to-many source-specific multicast (SSM) model, several commercially promising multicast applications, such as television channel distribution over the Internet, might be brought to the Internet much more quickly and efficiently than if full ASM functionality were required of the network.

PIM SSM is simpler than PIM sparse mode because only the one-to-many model is supported. Initial commercial multicast Internet applications are likely to be available to *subscribers* (that is, receivers that issue join messages) from only a single source (a special case of SSM covers the need for a backup source). PIM SSM therefore forms a subset of PIM sparse mode. PIM SSM builds shortest-path trees (SPTs) rooted at the source immediately because in SSM, the router closest to the interested receiver host is informed of the unicast IP address of the source for the multicast traffic. That is, PIM SSM bypasses the RP connection stage through shared distribution trees, as in PIM sparse mode, and goes directly to the source-based distribution tree.

PIM SSM introduces new terms for many of the concepts in PIM sparse mode. PIM SSM can technically be used in the entire 224/4 multicast address range, although PIM SSM operation is guaranteed only in the 232/8 range (232.0.0/24 is reserved). The new SSM terms are appropriate for Internet video applications and are summarized in [Table 435](#).

**Table 435: ASM and SSM Terminology**

| Term                | Any-Source Multicast  | Source-Specific Multicast         |
|---------------------|-----------------------|-----------------------------------|
| Address identifier  | G                     | S,G                               |
| Address designation | group                 | channel                           |
| Receiver operations | join, leave           | subscribe, unsubscribe            |
| Group address range | 224/4 excluding 232/8 | 224/4 (guaranteed only for 232/8) |

Although PIM SSM describes receiver operations as *subscribe* and *unsubscribe*, the same PIM sparse mode join and leave messages are used by both forms of the protocol. The terminology change distinguishes ASM from SSM even though the receiver messages are identical.

## PIM SSM

PIM source-specific multicast (SSM) uses a subset of PIM sparse mode and IGMP version 3 (IGMPv3) to allow a client to receive multicast traffic directly from the source. PIM SSM uses the PIM sparse-mode functionality to create an SPT between the receiver and the source, but builds the SPT without the help of an RP.

By default, the SSM group multicast address is limited to the IP address range from 232.0.0.0 through 232.255.255.255. However, you can extend SSM operations into another Class D range by including the **ssm-groups** statement at the **[edit routing-options multicast]** hierarchy level. The default SSM address range from 232.0.0.0 through 232.255.255.255 cannot be used in the **ssm-groups** statement. This statement is for adding other multicast addresses to the default SSM group addresses. This statement does not override the default SSM group address range.

You can also configure Junos OS to accept any-source multicast (ASM) join messages (\*G) for group addresses that are within the default or configured range of source-specific multicast (SSM) groups. This allows you to support a mix of any-source and source-specific multicast groups simultaneously.


An SSM-configured network has distinct advantages over a traditionally configured PIM sparse-mode network. There is no need for shared trees or RP mapping (no RP is required), or for RP-to-RP source discovery through MSDP.

Deploying SSM is easy. You need to configure PIM sparse mode on all router interfaces and issue the necessary SSM commands, including specifying IGMPv3 on the receiver's LAN. If PIM sparse mode is not explicitly configured on both the source and group member interfaces, multicast packets are not forwarded. Source lists, supported in IGMPv3, are

used in PIM SSM. As sources become active and start sending multicast packets, interested receivers in the SSM group receive the multicast packets.


In a PIM SSM-configured network, a host subscribes to an SSM channel (by means of IGMPv3), announcing a desire to join group G and source S (see [Figure 176](#)). The directly connected PIM sparse-mode router, the receiver's DR, sends an (S,G) join message to its RPF neighbor for the source. Notice in [Figure 176](#) that the RP is not contacted in this process by the receiver, as would be the case in normal PIM sparse-mode operations.

**Figure 176: Receiver Announces Desire to Join Group G and Source S**




The (S,G) join message initiates the source tree and then builds it out hop by hop until it reaches the source. In [Figure 177](#), the source tree is built across the network to Router 3, the last-hop router connected to the source.

**Figure 177: Router 3 (Last-Hop Router) Joins the Source Tree**



Using the source tree, multicast traffic is delivered to the subscribing host (see [Figure 178](#)).

**Figure 178: (S,G) State Is Built Between the Source and the Receiver**



To configure additional SSM groups, include the **ssm-groups** statement at the **[edit routing-options multicast]** hierarchy level.


#### Related Documentation

- [Source-Specific Multicast Groups Overview on page 5529](#)
- [Example: Configuring Source-Specific Multicast Groups with Any-Source Override on page 5537](#)

## Example: Configuring PIM SSM on a Network

The following example shows how PIM SSM is configured between a receiver and a source in the network illustrated in [Figure 179](#).

Figure 179: Network on Which to Configure PIM SSM



This example shows how to configure the IGMP version to IGMPv3 on all receiving host interfaces.

1. Enable IGMPv3 on all host-facing interfaces, and disable IGMP on the **fxp0.0** interface on Router 1.

```

user@router1# set protocols igmp interface all version 3
user@router1# set protocols igmp interface fxp0.0 disable

```



**NOTE:** When you configure IGMPv3 on a router, hosts on interfaces configured with IGMPv2 cannot join the source tree.

2. After the configuration is committed, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```

user@router1> show configuration protocol igmp

[edit protocols igmp]
interface all {
 version 3;
}
interface fxp0.0 {
 disable;
}

```

3. Use the **show igmp interface** command to verify that IGMP interfaces are configured.

```

user@router1> show igmp interface
Interface State Querier Timeout Version Groups
fe-0/0/0.0 Up 198.58.3.245 213 3 0
fe-0/0/1.0 Up 198.58.3.241 220 3 0
fe-0/0/2.0 Up 198.58.3.237 218 3 0
Configured Parameters:
IGMP Query Interval (1/10 secs): 1250
IGMP Query Response Interval (1/10 secs): 100
IGMP Last Member Query Interval (1/10 secs): 10
IGMP Robustness Count: 2
Derived Parameters:
IGMP Membership Timeout (1/10 secs): 2600
IGMP Other Querier Present Timeout (1/10 secs): 2550

```

4. Use the **show pim join extensive** command to verify the PIM join state on Router 2 and Router 3 (the upstream routers).

```
user@router2> show pim join extensive
232.1.1.1 10.4.1.2 sparse
 Upstream interface: fe-1/1/3.0
 Upstream State: Local Source
 Keepalive timeout: 209
 Downstream Neighbors:
 Interface: so-1/0/2.0
 10.10.71.1 State: Join Flags: S Timeout: 209
```

5. Use the **show pim join extensive** command to verify the PIM join state on Router 1 (the router connected to the receiver).

```
user@router1> show pim join extensive
232.1.1.1 10.4.1.2 sparse
 Upstream interface: so-1/0/2.0
 Upstream State: Join to Source
 Keepalive timeout: 209
 Downstream Neighbors:
 Interface: fe-0/2/3.0
 10.3.1.1 State: Join Flags: S Timeout: Infinity
```



**NOTE:** IP version 6 (IPv6) multicast routers use the Multicast Listener Discovery (MLD) Protocol to manage the membership of hosts and routers in multicast groups and to learn which groups have interested listeners for each attached physical networks. Each routing device maintains a list of host multicast addresses that have listeners for each subnetwork, as well as a timer for each address. However, the routing device does not need to know the address of each listener—just the address of each host. The routing device provides addresses to the multicast routing protocol it uses, which ensures that multicast packets are delivered to all subnetworks where there are interested listeners. In this way, MLD is used as the transport for the Protocol Independent Multicast (PIM) Protocol. MLD is an integral part of IPv6 and must be enabled on all IPv6 routing devices and hosts that need to receive IP multicast traffic. The Junos OS supports MLD versions 1 and 2. Version 2 is supported for source-specific multicast (SSM) include and exclude modes.

**Related Documentation**

- [Example: Configuring SSM Mapping on page 5535](#)
- 

## Example: Configuring an SSM-Only Domain

Deploying an SSM-only domain is much simpler than deploying an ASM domain because it only requires a few configuration steps. Enable PIM sparse mode on all interfaces by adding the **mode** statement at the **[edit protocols pim interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by adding the **disable** statement for that interface. Then configure IGMPv3 on all host-facing interfaces by adding the **version** statement at the **[edit protocols igmp interface interface-name]** hierarchy level.

In the following example, the host-facing interface is **fe-0/1/2**:

```
[edit]
protocols {
 pim {
 interface all {
 mode sparse;
 version 2;
 }
 interface fxp0.0 {
 disable;
 }
 }
 igmp {
 interface fe-0/1/2 {
 version 3;
 }
 }
}
```

## Example: Configuring SSM Mapping

SSM mapping does not require that all hosts support IGMPv3. SSM mapping translates IGMPv1 or IGMPv2 membership reports to an IGMPv3 report. This enables hosts running IGMPv1 or IGMPv2 to participate in SSM until the hosts transition to IGMPv3.

SSM mapping applies to all group addresses that match the policy, not just those that conform to SSM addressing conventions (232.1.1.1/32 for IPv4, ff30::/32 through ff3F::/32 for IPv6).

We recommend separate SSM maps for IPv4 and IPv6 if both address families require SSM support. If you apply an SSM map containing both IPv4 and IPv6 addresses to an interface in an IPv4 context (using IGMP), only the IPv4 addresses in the list are used. If there are no such addresses, no action is taken. Similarly, if you apply an SSM map containing both IPv4 and IPv6 addresses to an interface in an IPv6 context (using MLD), only the IPv6 addresses in the list are used. If there are no such addresses, no action is taken.

In this example, you create a policy to match the group addresses that you want to translate to IGMPv3. Then you define the SSM map that associates the policy with the source addresses where these group addresses are found. Finally, you apply the SSM map to one or more IGMP (for IPv4) or MLD (for IPv6) interfaces.

1. Create an SSM policy named **ssm-policy-example**. The policy terms match the IPv4 SSM group address 232.1.1.1/32 and the IPv6 SSM group address ff35::1/128. All other addresses are rejected.

```
user@router1# set policy-options policy-statement ssm-policy-example term A from
route-filter 232.1.1.1/32 exact
user@router1# set policy-options policy-statement ssm-policy-example term A then
accept
user@router1# set policy-options policy-statement ssm-policy-example term B from
route-filter ff35::1/128 exact
```

```
user@router1# set policy-options policy-statement ssm-policy-example term B then
accept
```

2. After the configuration is committed, use the **show configuration policy-options** command to verify the policy configuration.

```
user@host> show configuration policy-options

[edit policy-options]
policy-statement ssm-policy-example {
 term A {
 from {
 route-filter 232.1.1.1/32 exact;
 }
 then accept;
 }
 term B {
 from {
 route-filter ff35::1/128 exact;
 }
 then accept;
 }
 then reject;
}
```

The group addresses must match the configured policy for SSM mapping to occur.

3. Define two SSM maps, one called **ssm-map-ipv6-example** and one called **ssm-map-ipv4-example**, by applying the policy and configuring the source addresses as a multicast routing option.

```
user@host# set routing-options multicast ssm-map ssm-map-ipv6-example policy
ssm-policy-example
user@host# set routing-options multicast ssm-map ssm-map-ipv6-example source
fec0::1 fec0::12
user@host# set routing-options multicast ssm-map ssm-map-ipv4-example policy
ssm-policy-example
user@host# set routing-options multicast ssm-map ssm-map-ipv4-example source
10.10.10.4
user@host# set routing-options multicast ssm-map ssm-map-ipv4-example source
192.168.43.66
```

4. After the configuration is committed, use the **show configuration routing-options** command to verify the policy configuration.

```
user@host> show configuration routing-options

[edit routing-options]
multicast {
 ssm-map ssm-map-ipv6-example {
 policy ssm-policy-example;
 source [fec0::1 fec0::12];
 }
 ssm-map ssm-map-ipv4-example {
 policy ssm-policy-example;
 source [10.10.10.4 192.168.43.66];
 }
}
```



We recommend separate SSM maps for IPv4 and IPv6.

5. Apply SSM maps for IPv4-to-IGMP interfaces and SSM maps for IPv6-to-MLD interfaces:

```
user@host# set protocols igmp interface fe-0/1/0.0 ssm-map ssm-map-ipv4-example
user@host# set protocols mld interface fe-0/1/1.0 ssm-map ssm-map-ipv6-example
```

6. After the configuration is committed, use the **show configuration protocol** command to verify the IGMP and MLD protocol configuration.

```
user@router1> show configuration protocol

[edit protocols]
igmp {
 interface fe-0/1/0.0 {
 ssm-map ssm-map-ipv4-example;
 }
}
mld {
 interface fe-0/1/1.0 {
 ssm-map ssm-map-ipv6-example;
 }
}
```

7. Use the **show igmp interface** and the **show mld interface** commands to verify that the SSM maps are applied to the interfaces.

```
user@host> show igmp interface fe-0/1/0.0
Interface: fe-0/1/0.0
 Querier: 192.168.224.28
 State: Up Timeout: None Version: 2 Groups: 2
 SSM Map: ssm-map-ipv4-example

user@host> show mld interface fe-0/1/1.0
Interface: fe-0/1/1.0
 Querier: fec0:0:0:0:1::12
 State: Up Timeout: None Version: 2 Groups: 2
 SSM Map: ssm-map-ipv6-example
```

## Example: Configuring Source-Specific Multicast Groups with Any-Source Override

This example shows how to extend source-specific multicast (SSM) group operations beyond the default IP address range of 232.0.0.0 through 232.255.255.255. This example also shows how to accept any-source multicast (ASM) join messages (\*G) for group addresses that are within the default or configured range of SSM groups. This allows you to support a mix of any-source and source-specific multicast groups simultaneously.

- [Requirements on page 5538](#)
- [Overview on page 5538](#)
- [Configuration on page 5539](#)
- [Verification on page 5541](#)

## Requirements


Before you begin, configure the router interfaces. See the *Junos OS Network Interfaces Library for Routing Devices*.

## Overview

To deploy SSM, configure PIM sparse mode on all routing device interfaces and issue the necessary SSM commands, including specifying IGMPv3 or MLDv2 on the receiver's LAN. If PIM sparse mode is not explicitly configured on both the source and group members interfaces, multicast packets are not forwarded. Source lists, supported in IGMPv3 and MLDv2, are used in PIM SSM. Only sources that are specified send traffic to the SSM group.


In a PIM SSM-configured network, a host subscribes to an SSM channel (by means of IGMPv3 or MLDv2) to join group G and source S (see [Figure 180](#)). The directly connected PIM sparse-mode router, the receiver's designated router (DR), sends an (S,G) join message to its reverse-path forwarding (RPF) neighbor for the source. Notice in [Figure 180](#) that the RP is not contacted in this process by the receiver, as would be the case in normal PIM sparse-mode operations.

**Figure 180: Receiver Sends Messages to Join Group G and Source S**




The (S,G) join message initiates the source tree and then builds it out hop by hop until it reaches the source. In [Figure 181](#), the source tree is built across the network to Router 3, the last-hop router connected to the source.

**Figure 181: Router 3 (Last-Hop Router) Joins the Source Tree**



Using the source tree, multicast traffic is delivered to the subscribing host (see [Figure 182](#)).


Figure 182: (S,G) State Is Built Between the Source and the Receiver



SSM can operate in include mode or in exclude mode. In exclude mode the receiver specifies a list of sources that it does not want to receive the multicast group traffic from. The routing device forwards traffic to the receiver from any source except the sources specified in the exclusion list. The receiver accepts traffic from any sources except the sources specified in the exclusion list.

This example works with the simple RPF topology shown in [Figure 183](#).

Figure 183: Simple RPF Topology



## Configuration

### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface all
set protocols pim rp local address 10.255.72.46
set protocols pim rp local group-ranges 239.0.0.0/24
set protocols pim interface fe-1/0/0.0 mode sparse
set protocols pim interface lo0.0 mode sparse
set routing-options multicast ssm-groups 232.0.0.0/8
set routing-options multicast ssm-groups 239.0.0.0/8
set routing-options multicast asm-override-ssm
```

### Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure an RPF policy:

1. Configure OSPF.

```
[edit protocols ospf]
user@host# set area 0.0.0.0 interface fxp0.0 disable
user@host# set area 0.0.0.0 interface all
```

2. Configure PIM sparse mode.

```
[edit protocols pim]
user@host# set rp local address 10.255.72.46
user@host# set rp local group-ranges 239.0.0.0/24
user@host# set interface fe-1/0/0.0 mode sparse
user@host# set interface lo0.0 mode sparse
```

3. Configure additional SSM groups.

```
[edit routing-options]
user@host# set ssm-groups [232.0.0.0/8 239.0.0.0/8]
```

4. Configure the RP to accept ASM join messages for groups within the SSM address range.

```
[edit routing-options]
user@host# set multicast asm-override-ssm
```

5. If you are done configuring the device, commit the configuration.

```
user@host# commit
```

---

## Results

Confirm your configuration by entering the **show protocols** and **show routing-options** commands.

```
user@host# show protocols
ospf {
 area 0.0.0.0 {
 interface fxp0.0 {
 disable;
 }
 interface all;
 }
}
pim {
 rp {
 local {
 address 10.255.72.46;
 group-ranges {
 239.0.0.0/24;
 }
 }
 }
}
interface fe-1/0/0.0 {
 mode sparse;
}
interface lo0.0 {
 mode sparse;
}
}

user@host# show routing-options
multicast {
 ssm-groups [232.0.0.0/8 239.0.0.0/8];
```

```
asm-override-ssm;
}
```

Verification

To verify the configuration, run the following commands:

- `show igmp group`
- `show igmp statistics`
- `show pim join`

Related Documentation

- [Source-Specific Multicast Groups Overview on page 5529](#)

Example: Configuring SSM Maps for Different Groups to Different Sources

- [Multiple SSM Maps and Groups for Interfaces on page 5541](#)
- [Example: Configuring Multiple SSM Maps Per Interface on page 5541](#)

Multiple SSM Maps and Groups for Interfaces

You can configure multiple source-specific multicast (SSM) maps so that different groups map to different sources, which enables a single multicast group to map to different sources for different interfaces.

Example: Configuring Multiple SSM Maps Per Interface

This example shows how to assign more than one SSM map to an IGMP interface.

- [Requirements on page 5541](#)
- [Overview on page 5541](#)
- [Configuration on page 5542](#)
- [Verification on page 5543](#)

Requirements

This example requires Junos OS Release 11.4 or later.

Overview

In this example, you configure a routing policy, POLICY-ipv4-example1, that maps multicast group join messages over an IGMP logical interface to IPv4 multicast source addresses based on destination IP address as follows:

| Routing Policy Name         | Multicast Group Join Messages for a Route Filter at This Destination Address | Multicast Source Addresses   |
|-----------------------------|------------------------------------------------------------------------------|------------------------------|
| POLICY-ipv4-example1 term 1 | 232.1.1.1                                                                    | 10.10.10.4,<br>192.168.43.66 |

| Routing Policy Name         | Multicast Group Join Messages for a Route Filter at This Destination Address | Multicast Source Addresses   |
|-----------------------------|------------------------------------------------------------------------------|------------------------------|
| POLICY-ipv4-example1 term 2 | 232.1.1.2                                                                    | 10.10.10.5,<br>192.168.43.67 |

You apply routing policy POLICY-ipv4-example1 to IGMP logical interface fe-0/1/0.0.

### Configuration

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see the *CLI User Guide*.

To configure this example, perform the following task:

#### CLI Quick Configuration

To quickly configure this example, copy the following configuration commands into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set policy-options policy-statement POLICY-ipv4-example1 term 1 from route-filter
 232.1.1.1/32 exact
set policy-options policy-statement POLICY-ipv4-example1 term 1 then ssm-source
 10.10.10.4
set policy-options policy-statement POLICY-ipv4-example1 term 1 then ssm-source
 192.168.43.66
set policy-options policy-statement POLICY-ipv4-example1 term 1 then accept
set policy-options policy-statement POLICY-ipv4-example1 term 2 from route-filter
 232.1.1.2/32 exact
set policy-options policy-statement POLICY-ipv4-example1 term 2 then ssm-source
 10.10.10.5
set policy-options policy-statement POLICY-ipv4-example1 term 2 then ssm-source
 192.168.43.67
set policy-options policy-statement POLICY-ipv4-example1 term 2 then accept
set protocols igmp interface fe-0/1/0.0 ssm-map-policy POLICY-ipv4-example1
```

#### Step-by-Step Procedure

To configure multiple SSM maps per interface:

1. Configure protocol-independent routing options for route filter 232.1.1.1, and specify the multicast source addresses to which matching multicast groups are to be mapped.

```
[edit policy-options policy-statement POLICY-ipv4-example1 term 1]
user@host# set from route-filter 232.1.1.1/32 exact
user@host# set then ssm-source 10.10.10.4
user@host# set then ssm-source 192.168.43.66
user@host# set then accept
```

2. Configure protocol-independent routing options for route filter 232.1.1.2, and specify the multicast source addresses to which matching multicast groups are to be mapped.

```
[edit policy-options policy-statement POLICY-ipv4-example1 term 2]
user@host# set from route-filter 232.1.1.2/32 exact
user@host# set then ssm-source 10.10.10.5
```

```
user@host# set then ssm-source 192.168.43.67
user@host# set then accept
```

3. Apply the policy map POLICY-ipv4-example1 to IGMP logical interface fe-0/1/1/0.

```
[edit protocols igmp interface fe-0/1/0.0]
user@host# set ssm-map-policy POLICY-ipv4-example1
```

**Results** After the configuration is committed, confirm the configuration by entering the **show policy-options** and **show protocols** configuration mode commands. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
user@host# show policy-options
policy-statement POLICY-ipv4-example1 {
 term 1 {
 from {
 route-filter 232.1.1.1/32 exact;
 }
 then {
 ssm-source [10.10.10.4 192.168.43.66];
 accept;
 }
 }
 term 2 {
 from {
 route-filter 232.1.1.2/32 exact;
 }
 then {
 ssm-source [10.10.10.5 192.168.43.67];
 accept;
 }
 }
}

user@host# show protocols
igmp {
 interface fe-0/1/0.0 {
 ssm-map-policy POLICY-ipv4-example1;
 }
}
```

### Verification

Confirm that the configuration is working properly.

- [Displaying Information About IGMP-Enabled Interfaces on page 5543](#)
- [Displaying the PIM Groups on page 5544](#)
- [Displaying the Entries in the IP Multicast Forwarding Table on page 5544](#)

#### *Displaying Information About IGMP-Enabled Interfaces*

**Purpose** Verify that the SSM map policy POLICY-ipv4-example1 is applied to logical interface fe-0/1/0.0.

**Action** Use the `show igmp interface` operational mode command for the IGMP logical interface to which you applied the SSM map policy.

```
user@host> show igmp interface
Interface: fe-0/1/0.0
 Querier: 10.111.30.1
 State: Up Timeout: None Version: 2 Groups: 2
 SSM Map Policy: POLICY-ipv4-example1;
```

```
Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2
```

```
Derived Parameters:
IGMP Membership Timeout: 260.0
IGMP Other Querier Present Timeout: 255.0
```

The command output displays the name of the IGMP logical interface (fe-0/1/0.0), which is the address of the routing device that has been elected to send membership queries and group information.

#### ***Displaying the PIM Groups***

**Purpose** Verify the Protocol Independent Multicast (PIM) source and group pair (S,G) entries.

**Action** Use the `show pim join extensive 232.1.1.1` operational mode command to display the PIM source and group pair (S,G) entries for the 232.1.1.1 group.

#### ***Displaying the Entries in the IP Multicast Forwarding Table***

**Purpose** Verify that the IP multicast forwarding table displays the multicast route state.

**Action** Use the `show multicast route extensive` operational mode command to display the entries in the IP multicast forwarding table to verify that the **Route state** is active and that the **Forwarding state** is forwarding.

**Related Documentation**

- *Example: Configuring Source-Specific Multicast*
- *Example: Configuring Source-Specific Draft-Rosen 7 Multicast VPNs*



## CHAPTER 204

# Using Static RP

- [Understanding Static RP on page 5545](#)
- [Configuring Local PIM RPs on page 5545](#)
- [Configuring the Static PIM RP Address on the Non-RP Routing Device on page 5547](#)

## Understanding Static RP

---

Protocol Independent Multicast (PIM) sparse mode is the most common multicast protocol used on the Internet. PIM sparse mode is the default mode whenever PIM is configured on any interface of the device. However, because PIM must not be configured on the network management interface, you must disable it on that interface.

Each any-source multicast (ASM) group has a shared tree through which receivers learn about new multicast sources and new receivers learn about all multicast sources. The rendezvous point (RP) router is the root of this shared tree and receives the multicast traffic from the source. To receive multicast traffic from the groups served by the RP, the device must determine the IP address of the RP for the source.

You can configure a static rendezvous point (RP) configuration that is similar to static routes. A static configuration has the benefit of operating in PIM version 1 or version 2. When you configure the static RP, the RP address that you select for a particular group must be consistent across all routers in a multicast domain.

One common way for the device to locate RPs is by static configuration of the IP address of the RP. A static configuration is simple and convenient. However, if the statically defined RP router becomes unreachable, there is no automatic failover to another RP router. To remedy this problem, you can use anycast RP.

### Related Documentation

- [Configuring Local PIM RPs on page 5545](#)
- [Configuring the Static PIM RP Address on the Non-RP Routing Device on page 5547](#)

## Configuring Local PIM RPs

---

Local RP configuration makes the routing device a statically defined RP. Consider statically defining an RP if the network does not have many different RPs defined or if the RP assignment does not change very often. The Junos IPv6 PIM implementation supports

only static RP configuration. Automatic RP announcement and bootstrap routers are not available with IPv6.

You can configure a local RP globally or for a routing instance. This example shows how to configure a local RP in a routing instance for IPv4 or IPv6.

To configure the routing device's RP properties:

1. Configure the routing instance as the local RP.

```
[routing-instances VPN-A protocols pim]
user@host# set rp local
```

2. Configure the IP protocol family and IP address.

IPv6 PIM hello messages are sent to every interface on which you configure **family inet6**, whether at the PIM level of the hierarchy or not. As a result, if you configure an interface with both **family inet** at the **[edit interface *interface-name*]** hierarchy level and **family inet6** at the **[edit protocols pim interface *interface-name*]** hierarchy level, PIM sends both IPv4 and IPv6 hellos to that interface.

By default, PIM operates in sparse mode on an interface. If you explicitly configure sparse mode, PIM uses this setting for all IPv6 multicast groups. However, if you configure sparse-dense mode, PIM does not accept IPv6 multicast groups as dense groups and operates in sparse mode over them.

```
[edit routing-instances VPN-A protocols pim rp local]
user@host# set family inet6 address 2001:db8:85a3::8a2e:370:7334
user@host# set family inet address 10.1.2.254
```

3. (IPv4 only) Configure the routing device's RP priority.



**NOTE:** The priority statement is not supported for IPv6, but is included here for informational purposes. The routing device's priority value for becoming the RP is included in the bootstrap messages that the routing device sends. Use a smaller number to increase the likelihood that the routing device becomes the RP for local multicast groups. Each PIM routing device uses the priority value and other factors to determine the candidate RPs for a particular group range. After the set of candidate RPs is distributed, each routing device determines algorithmically the RP from the candidate RP set using a hash function. By default, the priority value is set to 1. If this value is set to 0, the bootstrap router can override the group range being advertised by the candidate RP.

```
[edit routing-instances VPN-A protocols pim rp local]
user@host# set priority 5
```

4. Configure the groups for which the routing device is the RP.

By default, a routing device running PIM is eligible to be the RP for all IPv4 or IPv6 groups (224.0.0.0/4 or FF70::/12 to FFF0::/12). The following example limits the groups for which this routing device can be the RP.

```
[edit routing-instances VPN-A protocols pim rp local]
```

```
user@host# set group-ranges fec0::/10
user@host# set group-ranges 10.1.2.0/24
```

5. (IPv4 only) Modify the local RP hold time.

If the local routing device is configured as an RP, it is considered a candidate RP for its local multicast groups. For candidate RPs, the hold time is used by the bootstrap router to time out RPs, and applies to the bootstrap RP-set mechanism. The RP hold time is part of the candidate RP advertisement message sent by the local routing device to the bootstrap router. If the bootstrap router does not receive a candidate RP advertisement from an RP within the hold time, it removes that routing device from its list of candidate RPs. The default hold time is 150 seconds.

```
[edit routing-instances VPN-A protocols pim rp local]
user@host# set hold-time 200
```

6. (Optional) Override dynamic RP for the specified group address range.

If you configure both static RP mapping and dynamic RP mapping (such as auto-RP) in a single routing instance, allow the static mapping to take precedence for the given static RP group range, and allow dynamic RP mapping for all other groups.

If you exclude this statement from the configuration and you use both static and dynamic RP mechanisms for different group ranges within the same routing instance, the dynamic RP mapping takes precedence over the static RP mapping, even if static RP is defined for a specific group range.

```
[edit routing-instances VPN-A protocols pim rp local]
user@host# set override
```

7. Monitor the operation of PIM by running the **show pim** commands. Run **show pim ?** to display the supported commands.

- Related Documentation**
- [PIM Overview on page 5491](#)
  - [Understanding MLD on page 5441](#)

## Configuring the Static PIM RP Address on the Non-RP Routing Device

Consider statically defining an RP if the network does not have many different RPs defined or if the RP assignment does not change very often. The Junos IPv6 PIM implementation supports only static RP configuration. Automatic RP announcement and bootstrap routers are not available with IPv6.

You configure a static RP address on the non-RP routing device. This enables the non-RP routing device to recognize the local statically defined RP. For example, if R0 is a non-RP router and R1 is the local RP router, you configure R0 with the static RP address of R1. The static IP address is the routable address assigned to the loopback interface on R1. In the following example, the loopback address of the RP is 2001:db8:85a3::8a2e:370:7334.

You can configure a static RP address globally or for a routing instance. This example shows how to configure a static RP address in a routing instance for IPv6.

To configure the static RP address:

1. On a non-RP routing device, configure the routing instance to point to the routable address assigned to the loopback interface of the RP.

```
[routing-instances VPN-A protocols pim rp]
user@host# set static address 2001:db8:85a3::8a2e:370:7334
```



**NOTE:** Logical systems are also supported. You can configure a static RP address in a logical system only if the logical system is not directly connected to a source.

2. (Optional) Set the PIM sparse mode version.

For each static RP address, you can optionally specify the PIM version. The default PIM version is version 1.

```
[edit routing-instances VPN-A protocols pim rp]
user@host# set static address 2001:db8:85a3::8a2e:370:7334 version 2
```



**NOTE:** The default PIM version can be version 1 or version 2, depending on the mode you are configuring. PIM version 1 is the default for RP mode ([edit pim rp static address *address*]). PIM version 2 is the default for interface mode ([edit pim interface *interface-name*]). Explicitly configured versions override the defaults.

3. (Optional) Set the group address range.

By default, a routing device running PIM is eligible to be the RP for all IPv4 or IPv6 groups (224.0.0.0/4 or FF70::/12 to FFF0::/12). The following example limits the groups for which the 2001:db8:85a3::8a2e:370:7334 address can be the RP.

```
[edit routing-instances VPN-A protocols pim rp]
user@host# set static address 2001:db8:85a3::8a2e:370:7334 group-ranges fec0::/10
```

The RP that you select for a particular group must be consistent across all routers in a multicast domain.

4. (Optional) Override dynamic RP for the specified group address range.

If you configure both static RP mapping and dynamic RP mapping (such as auto-RP) in a single routing instance, allow the static mapping to take precedence for the given static RP group range, and allow dynamic RP mapping for all other groups.

If you exclude this statement from the configuration and you use both static and dynamic RP mechanisms for different group ranges within the same routing instance, the dynamic RP mapping takes precedence over the static RP mapping, even if static RP is defined for a specific group range.

```
[edit routing-instances VPN-A protocols pim rp static address
 2001:db8:85a3::8a2e:370:7334]
user@host# set override
```

5. Monitor the operation of PIM by running the **show pim** commands. Run **show pim ?** to display the supported commands.

- Related Documentation**
- [PIM Overview on page 5491](#)
  - [Understanding MLD on page 5441](#)



## Using Anycast RP

- [Understanding RP Mapping with Anycast RP on page 5551](#)
- [Example: Configuring PIM Anycast With or Without MSDP on page 5552](#)
- [Configuring a PIM Anycast RP Router with MSDP on page 5555](#)
- [Configuring a PIM Anycast RP Router Using Only PIM on page 5556](#)
- [Configuring All PIM Anycast Non-RP Routers on page 5557](#)
- [Example: Configuring Multiple RPs in a Domain with Anycast RP on page 5558](#)

### Understanding RP Mapping with Anycast RP

---

Having a single active rendezvous point (RP) per multicast group is much the same as having a single server providing any service. All traffic converges on this single point, although other servers are sitting idle, and convergence is slow when the resource fails. In multicast specifically, there might be closer RPs on the shared tree, so the use of a single RP is suboptimal.

For the purposes of load balancing and redundancy, you can configure anycast RP. You can use anycast RP within a domain to provide redundancy and RP load sharing. When an RP fails, sources and receivers are taken to a new RP by means of unicast routing. When you configure anycast RP, you bypass the restriction of having one active RP per multicast group, and instead deploy multiple RPs for the same group range. The RP routers share one unicast IP address. Sources from one RP are known to other RPs that use the Multicast Source Discovery Protocol (MSDP). Sources and receivers use the closest RP, as determined by the interior gateway protocol (IGP).

Anycast means that multiple RP routers share the same unicast IP address. Anycast addresses are advertised by the routing protocols. Packets sent to the anycast address are sent to the nearest RP with this address. Anycast addressing is a generic concept and is used in PIM sparse mode to add load balancing and service reliability to RPs.

Anycast RP is defined in RFC3446, *Anycast RP Mechanism Using PIM and MSDP*, and can be found here: <https://www.ietf.org/rfc/rfc3446.txt>.

#### Related Documentation

- [Configuring the Static PIM RP Address on the Non-RP Routing Device on page 5547](#)
- [Example: Configuring Multiple RPs in a Domain with Anycast RP on page 5558](#)
- [Example: Configuring PIM Anycast With or Without MSDP on page 5552](#)

## Example: Configuring PIM Anycast With or Without MSDP

---

When you configure anycast RP, you bypass the restriction of having one active rendezvous point (RP) per multicast group, and instead deploy multiple RPs for the same group range. The RP routers share one unicast IP address. Sources from one RP are known to other RPs that use the Multicast Source Discovery Protocol (MSDP). Sources and receivers use the closest RP, as determined by the interior gateway protocol (IGP).

You can use anycast RP within a domain to provide redundancy and RP load sharing. When an RP stops operating, sources and receivers are taken to a new RP by means of unicast routing.

You can configure anycast RP to use PIM and MSDP for IPv4, or PIM alone for both IPv4 and IPv6 scenarios. Both are discussed in this section.

We recommend a static RP mapping with anycast RP over a bootstrap router and auto-RP configuration because it provides all the benefits of a bootstrap router and auto-RP without the complexity of the BSR and auto-RP mechanisms.

All systems on a subnet must run the same version of PIM.

The default PIM version can be version 1 or version 2, depending on the mode you are configuring. PIMv1 is the default RP mode (at the **[edit protocols pim rp static address address]** hierarchy level). However, PIMv2 is the default for interface mode (at the **[edit protocols pim interface interface-name]** hierarchy level). Explicitly configured versions override the defaults. This example explicitly configures PIMv2 on the interfaces.

The following example shows an anycast RP configuration for the RP routers, first with MSDP and then using PIM alone, and for non-RP routers.

1. For a network using an RP with MSDP, configure the RP using the **lo0** loopback interface, which is always up. Include the **address** statement and specify the unique and routable router ID and the RP address at the **[edit interfaces lo0 unit 0 family inet]** hierarchy level. In this example, the router ID is **198.58.3.254** and the shared RP address is **198.58.3.253**. Include the **primary** statement for the first address. Including the **primary** statement selects the router's primary address from all the preferred addresses on all interfaces.

```
interfaces {
 lo0 {
 description "PIM RP";
 unit 0 {
 family inet {
 address 198.58.3.254/32;
 primary;
 address 198.58.3.253/32;
 }
 }
 }
}
```



2. Specify the RP address. Include the **address** statement at the **[edit protocols pim rp local]** hierarchy level (the same address as the secondary **lo0** interface).

For all interfaces, include the **mode** statement to set the mode to **sparse** and the **version** statement to specify PIM version 2 at the **[edit protocols pim rp local interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by including the **disable** statement for that interface.

```
protocols {
 pim {
 rp {
 local {
 family inet;
 address 198.58.3.253;
 }
 }
 interface all {
 mode sparse;
 version 2;
 }
 interface fxp0.0 {
 disable;
 }
 }
}
```

3. Configure MSDP peering. Include the **peer** statement to configure the address of the MSDP peer at the **[edit protocols msdp]** hierarchy level. For MSDP peering, use the unique, primary addresses instead of the anycast address. To specify the local address for MSDP peering, include the **local-address** statement at the **[edit protocols msdp peer]** hierarchy level.

```
protocols {
 msdp {
 peer 198.58.3.250 {
 local-address 198.58.3.254;
 }
 }
}
```



**NOTE:** If you need to configure a PIM RP for both IPv4 and IPv6 scenarios, perform Step 4 and Step 5. Otherwise, go to Step 6.

4. Configure an RP using the **lo0** loopback interface, which is always up. Include the **address** statement to specify the unique and routable router address and the RP address at the **[edit interfaces lo0 unit 0 family inet]** hierarchy level. In this example, the router ID is **198.58.3.254** and the shared RP address is **198.58.3.253**. Include the **primary** statement on the first address. Including the **primary** statement selects the router's primary address from all the preferred addresses on all interfaces.

```
interfaces {
 lo0 {
 description "PIM RP";
```

```

unit 0 {
 family inet {
 address 198.58.3.254/32 {
 primary;
 }
 address 198.58.3.253/32;
 }
}

```

5. Include the **address** statement at the **[edit protocols pim rp local]** hierarchy level to specify the RP address (the same address as the secondary **lo0** interface).

For all interfaces, include the **mode** statement to set the mode to **sparse**, and the **version** statement to specify PIM version 2 at the **[edit protocols pim rp local interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by including the **disable** statement for that interface.

Include the **anycast-pim** statement to configure anycast RP without MSDP (for example, if IPv6 is used for multicasting). The other RP routers that share the same IP address are configured using the **rp-set** statement. There is one entry for each RP, and the maximum that can be configured is 15. For each RP, specify the routable IP address of the router and whether MSDP source active (SA) messages are forwarded to the RP.

MSDP configuration is not necessary for this type of IPv4 anycast RP configuration.

```

protocols {
 pim {
 rp {
 local {
 family inet {
 address 198.58.3.253;
 anycast-pim {
 rp-set {
 address 198.58.3.240;
 address 198.58.3.241 forward-msdp-sa;
 }
 local-address 198.58.3.254; #If not configured, use lo0 primary
 }
 }
 }
 }
 }
 interface all {
 mode sparse;
 version 2;
 }
 interface fxp0.0 {
 disable;
 }
}

```

6. Configure the non-RP routers. The anycast RP configuration for a non-RP router is the same whether MSDP is used or not. Specify a static RP by adding the address at

the `[edit protocols pim rp static]` hierarchy level. Include the **version** statement at the `[edit protocols pim rp static address]` hierarchy level to specify PIM version 2.

```
protocols {
 pim {
 rp {
 static {
 address 198.58.3.253 {
 version 2;
 }
 }
 }
 }
}
```

7. Include the **mode** statement at the `[edit protocols pim interface all]` hierarchy level to specify sparse mode on all interfaces. Then include the **version** statement at the `[edit protocols pim rp interface all mode]` to configure all interfaces for PIM version 2. When configuring all interfaces, exclude the **fxp0.0** management interface by including the **disable** statement for that interface.

```
protocols {
 pim {
 interface all {
 mode sparse;
 version 2;
 }
 interface fxp0.0 {
 disable;
 }
 }
}
```

## Configuring a PIM Anycast RP Router with MSDP

Add the **address** statement at the `[edit protocols pim rp local]` hierarchy level to specify the RP address (the same address as the secondary **lo0** interface).

For all interfaces, use the **mode** statement to set the mode to **sparse** and the **version** statement to specify PIM version 2 at the `[edit protocols pim rp local interface all]` hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by adding the **disable** statement for that interface.

```
protocols {
 pim {
 rp {
 local {
 family inet;
 address 198.58.3.253;
 }
 interface all {
 mode sparse;
 version 2;
 }
 }
 }
}
```

```
 interface fxp0.0 {
 disable;
 }
 }
}
```

To configure MSDP peering, add the **peer** statement to configure the address of the MSDP peer at the **[edit protocols msdp]** hierarchy level. For MSDP peering, use the unique, primary addresses instead of the anycast address. To specify the local address for MSDP peering, add the **local-address** statement at the **[edit protocols msdp peer]** hierarchy level.

```
protocols {
 msdp {
 peer 198.58.3.250 {
 local-address 198.58.3.254;
 }
 }
}
```

---

## Configuring a PIM Anycast RP Router Using Only PIM

---

In this example, configure an RP using the **lo0** loopback interface, which is always up. Use the **address** statement to specify the unique and routable router address and the RP address at the **[edit interfaces lo0 unit 0 family inet]** hierarchy level. In this case, the router ID is 198.58.3.254/32 and the shared RP address is 198.58.3.253/32. Add the flag statement **primary** to the first address. Using this flag selects the router's primary address from all the preferred addresses on all interfaces.

```
interfaces {
 lo0 {
 description "PIM RP";
 unit 0 {
 family inet {
 address 198.58.3.254/32 {
 primary;
 }
 address 198.58.3.253/32;
 }
 }
 }
}
```

Add the **address** statement at the **[edit protocols pim rp local]** hierarchy level to specify the RP address (the same address as the secondary **lo0** interface).

For all interfaces, use the **mode** statement to set the mode to **sparse**, and include the **version** statement to specify PIM version 2 at the **[edit protocols pim rp local interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by adding the **disable** statement for that interface.

Use the **anycast-pim** statement to configure anycast RP without MSDP (for example, if IPv6 is used for multicasting). The other RP routers that share the same IP address are

configured using the **rp-set** statement. There is one entry for each RP, and the maximum that can be configured is 15. For each RP, specify the routable IP address of the router and whether MSDP source active (SA) messages are forwarded to the RP.

```
protocols {
 pim {
 rp {
 local {
 family inet {
 address 198.58.3.253;
 anycast-pim {
 rp-set {
 address 198.58.3.240;
 address 198.58.3.241 forward-msdp-sa;
 }
 local-address 198.58.3.254; #If not configured, use lo0 primary
 }
 }
 }
 }
 interface all {
 mode sparse;
 version 2;
 }
 interface fxp0.0 {
 disable;
 }
 }
}
```

MSDP configuration is not necessary for this type of IPv4 anycast RP configuration.

## Configuring All PIM Anycast Non-RP Routers

Use the **mode** statement at the **[edit protocols pim rp interface all]** hierarchy level to specify sparse mode on all interfaces. Then add the **version** statement at the **[edit protocols pim rp interface all mode]** to configure all interfaces for PIM version 2. When configuring all interfaces, exclude the **fxp0.0** management interface by adding the **disable** statement for that interface.

```
protocols {
 pim {
 interface all {
 mode sparse;
 version 2;
 }
 interface fxp0.0 {
 disable;
 }
 }
}
```

## Example: Configuring Multiple RPs in a Domain with Anycast RP

---

This example shows how to configure anycast RP on each RP router in the PIM-SM domain. With this configuration you can deploy more than one RP for a single group range. This enables load balancing and redundancy.

- [Requirements on page 5558](#)
- [Overview on page 5558](#)
- [Configuration on page 5558](#)
- [Verification on page 5560](#)

### Requirements

Before you begin:

- Configure the router interfaces. See the *Junos OS Network Interfaces Library for Routing Devices*.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Routing Devices*.
- Configure PIM Sparse Mode on the interfaces. See [“Enabling PIM Sparse Mode” on page 5512](#).

### Overview

When you configure anycast RP, the RP routers in the PIM-SM domain use a shared address. In this example, the shared address is 10.1.1.2/32. Anycast RP uses Multicast Source Discovery Protocol (MSDP) to discover and maintain a consistent view of the active sources. Anycast RP also requires an RP selection method, such as static, auto-RP, or bootstrap RP. This example uses static RP and shows only one RP router configuration.

### Configuration

|                                |                                                                                                                                                                                                                                                                                                                         |
|--------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>CLI Quick Configuration</b> | To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the <b>[edit]</b> hierarchy level.                                          |
| <b>RP Routers</b>              | <pre>set interfaces lo0 unit 0 family inet address 192.168.132.1/32 primary set interfaces lo0 unit 0 family inet address 10.1.1.2/32 set protocols msdp local-address 192.168.132.1 set protocols msdp peer 192.168.12.1 set protocols pim rp local address 10.1.1.2 set routing-options router-id 192.168.132.1</pre> |
| <b>Non-RP Routers</b>          | <pre>set protocols pim rp static address 10.1.1.2</pre>                                                                                                                                                                                                                                                                 |

**Step-by-Step Procedure** The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure anycast RP:

1. On each RP router in the domain, configure the shared anycast address on the router's loopback address.

```
[edit interfaces]
user@host# set lo0 unit 0 family inet address 10.1.1.2/32
```

2. On each RP router in the domain, make sure that the router's regular loopback address is the primary address for the interface, and set the router ID.

```
[edit interfaces]
user@host# set lo0 unit 0 family inet address 192.168.132.1/32 primary
```

```
[edit routing-options]
user@host# set router-id 192.168.132.1
```

3. On each RP router in the domain, configure the local RP address, using the shared address.

```
[edit protocols pim]
user@host# set rp local address 10.1.1.2
```

4. On each RP router in the domain, create MSDP sessions to the other RPs in the domain.

```
[edit protocols msdp]
user@host# set local-address 192.168.132.1
user@host# set peer 192.168.12.1
```

5. On each non-RP router in the domain, configure a static RP address using the shared address.

```
[edit protocols pim]
user@host# set rp static address 10.1.1.2
```

6. If you are done configuring the devices, commit the configuration.

```
user@host# commit
```

## Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
lo0 {
 unit 0 {
 family inet {
 address 192.168.132.1/32 {
 primary;
 }
 }
 }
}
```

```
 address 10.1.1.2/32;
 }
}
```

*On the RP routers:*

```
user@host# show protocols
msdp {
 local-address 192.168.132.1;
 peer 192.168.12.1;
}
pim {
 rp {
 local {
 address 10.1.1.2;
 }
 }
}
```

*On the non-RP routers:*

```
user@host# show protocols
pim {
 rp {
 static {
 address 10.1.1.2;
 }
 }
}

user@host# show routing-options
router-id 192.168.132.1;
```

## Verification

To verify the configuration, run the `show pim rps extensive inet` command.

### Related Documentation

- [Example: Configuring PIM Anycast With or Without MSDP on page 5552](#)
- [Understanding PIM Sparse Mode on page 5509](#)
- [Understanding RP Mapping with Anycast RP on page 5551](#)



# Using Auto-RP

- [Understanding PIM Auto-RP on page 5561](#)
- [Configuring PIM Auto-RP on page 5561](#)

## Understanding PIM Auto-RP

---

You can configure a more dynamic way of assigning rendezvous points (RPs) in a multicast network by means of auto-RP. When you configure auto-RP for a router, the router learns the address of the RP in the network automatically and has the added advantage of operating in PIM version 1 and version 2.

Although auto-RP is a nonstandard (non-RFC-based) function that typically uses dense mode PIM to advertise control traffic, it provides an important failover advantage that simple static RP assignment does not. You can configure multiple routers as RP candidates. If the elected RP fails, one of the other preconfigured routers takes over the RP functions. This capability is controlled by the auto-RP mapping agent.

### Related Documentation

- [Configuring PIM Auto-RP on page 5561](#)

## Configuring PIM Auto-RP

---

For correct operation, every multicast router within a PIM domain must be able to map a particular multicast group address to the same rendezvous point (RP). The auto-RP mechanism is one way that a multicast router can learn the set of group-to-RP mappings. Auto-RP automatically distributes mapping information to routing devices. It simplifies use of multiple RPs for different multicast group ranges, thus allowing multiple RPs to act as backups for each other. Auto-RP relies on a router to act as the RP mapping agent. Potential RPs announce themselves to the mapping agent, and the mapping agent resolves any conflicts.

The mapping agent sends the multicast group-RP mapping information to the other routers using PIM dense mode. The specific groups used are 224.0.1.39 and .40. The first (.39) is used to advertise, the second (.40) is used for discovery. Because PIM dense mode is necessary to enable auto-RP to work, which in turns enables PIM sparse mode to work, you must configure PIM sparse-dense mode in the PIM domains that use auto-RP.

Although auto-RP is a nonstandard (non-RFC-based) function requiring dense mode PIM to advertise control traffic, it provides an important failover advantage that static

RP assignment does not. That is, you can configure multiple routing devices as RP candidates. If the elected RP fails, one of the other preconfigured routing devices takes over the RP functions. This capability is controlled by the auto-RP mapping agent.

Auto-RP operates in PIM version 1 and version 2.

In most cases, how the routing device handles auto-RP discovery, announce, or mapping messages depends on whether the routing device is an RP (configured as local RP) or not. [Table 436](#) shows how the routing device behaves depending on the local RP configuration.

**Table 436: Local RP and Auto-RP Message Types**

| Auto-RP Message Type | Local RP? | Routing Device Behavior                                                                                                                                             |
|----------------------|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| discovery            | No        | Listen for auto-RP mapping messages.                                                                                                                                |
| discovery            | Yes       | Listen for auto-RP mapping messages.                                                                                                                                |
| announce             | No        | Listen for auto-RP mapping messages.                                                                                                                                |
| announce             | Yes       | Listen for auto-RP mapping messages. Send auto-RP announce messages.                                                                                                |
| mapping              | No        | Listen for auto-RP mapping messages. Listen for auto-RP announce messages. If elected mapping agent, send auto-RP mapping messages.                                 |
| mapping              | Yes       | Listen for auto-RP mapping messages. Send auto-RP announce messages. Listen for auto-RP announce messages. If elected mapping agent, send auto-RP mapping messages. |



**NOTE:** If the routing device receives auto-RP announcements split across multiple messages, the routing device loses the information in the previous part of the message as soon as the next part of the message is received.

You can configure auto-RP properties globally or for a routing instance. This example shows the global configuration.

To configure auto-RP properties:

1. Configure PIM in sparse-dense mode on all routing devices in the PIM domain.

```
[edit protocols pim]
user@host# edit
user@host# set interface all mode sparse-dense
```

This configuration allows the routing device to operate in sparse mode for most groups and dense mode for others. The default is to operate in sparse mode unless the routing device is specifically informed of a dense mode group.

2. Configure a routable loopback interface address on all routing devices in the PIM domain.

The routing device joins the auto-RP groups on the configured interfaces and on the loopback interface **lo0.0**. For auto-RP to work correctly, configure a routable IP address on the loopback interface. You cannot use the loopback address 127.0.0.1. Also, you must enable PIM sparse-dense mode on the **lo0.0** interface if you do not specify **interface all**.

```
[edit interfaces lo0.0 unit 0 family inet]
user@host# set address 192.168.0.3 preferred
```

3. Configure the two multicast dense groups on all the routing devices.

Auto-RP requires multicast flooding to announce potential RP candidates and to discover the elected RPs in the network. Multicast flooding occurs through a PIM dense mode model, where group 224.0.1.39 is used for **announce** messages and group 224.0.1.40 is used for **discovery** messages.

```
[edit protocols pim]
user@host# set dense-groups 224.0.1.39/32
user@host# set dense-groups 224.0.1.40/32
```



**TIP:** Step 3 is required. When auto-RP is enabled, the auto-RP announce group (224.0.1.39) and auto-RP-discovery group (224.0.1.40) must be configured explicitly as dense groups. When the auto-RP discovery group is not configured as a dense group, auto-RP is not enabled. When the auto-RP announce group is not configured as a dense group, auto-RP is enabled in the discovery mode only, and mapping and announce modes are disabled.

4. Configure the auto-RP **announce** option.

At least one routing device in the PIM domain must announce auto-RP messages and at least one must map them, or you can configure a routing device to perform both functions.

When a routing device sends announce messages in the network, it is advertising itself as a candidate RP. A routing device configured with this option must also be configured as an RP, or announce messages are not sent.

```
[edit protocols pim rp]
user@host# set local address 192.168.0.1
user@host# set auto-rp announce
```



**NOTE:** You cannot include the `auto-rp announce` option at the `[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim]` hierarchy level.

5. Configure the auto-RP mapping agent.

The mapping agent sends discovery messages to the network, informing all routing devices in a multicast group of which RP to use. If the mapping agent is also an RP, the **mapping** option also allows the routing device to send auto-RP announcements (mapping on an RP allows the routing device to perform both the announcement and mapping functions).

```
[edit protocols pim rp]
user@host# set auto-rp mapping
```

If the mapping agent is also an RP, configure the mapping agent as a local RP.

```
[edit protocols pim rp]
user@host# set local address 192.168.0.2
```

6. Configure mapping agent election.

If you configure the **mapping** option on more than one routing device in the PIM domain, configure mapping agent election on each potential mapping agent.

Auto-RP specifications state that mapping agents do not send mapping messages if they receive messages from a mapping agent with a higher IP address. However, some vendors' mapping agents continue to announce mappings, even in the presence of higher-addressed mapping agents. In other words, some mapping agents will always send mapping messages.

The default auto-RP operation is to perform mapping agent election. To explicitly configure mapping agent election, you can include the **mapping-agent-election** statement. When this option is configured, the mapping agent will stop sending mapping messages if it receives messages from a mapping agent with a higher IP address.

```
[edit protocols pim rp]
user@host# set auto-rp mapping mapping-agent-election
```

Mapping message suppression is disabled with the **no-mapping-agent-election** statement. When this option is configured, the mapping agent will always send mapping messages even in the presence of higher-addressed mapping agents.

To disable mapping agent election for compatibility with other vendors' equipment, include the **no-mapping-agent-election** statement.

```
[edit protocols pim rp]
user@host# set auto-rp mapping no-mapping-agent-election
```

7. Configure the remaining routing devices in the PIM domain to discover the RP.

Discovery enables the routing devices to receive and process discovery messages from the mapping agent. This is the most basic auto-RP option.

```
[edit protocols pim rp]
user@host# set auto-rp discovery
```

8. Monitor the operation of PIM auto-RP routers by running the following commands:

- **show pim interfaces**
- **show pim rps**

- [show pim rps](#)

9. Issue the **show pim rps extensive** command to see information about how an RP is learned, what groups it handles, and the number of groups actively using the RP.

```

user@host> show pim rps extensive
RP: 192.168.5.1
Learned from 192.168.5.1 via: auto-rp
Time Active: 00:34:29
Holdtime: 150 with 108 remaining
Device Index: 6
Subunit: 32769
Interface: pd-0/0/0.32769
Group Ranges:
 224.0.0.0/4
Active groups using RP:
 224.2.2.100
 total 1 groups active
Register State for RP:
Group Source FirstHop RP Address StateRP address Type Holdtime
Timeout

```

In the example, the RP at 192.168.5.1 was learned through auto-RP. The RP is able to support all groups in the 224.0.0.0/4 range (all possible groups). The local router has sent PIM control traffic for the 224.2.2.100 group to the RP.

Additionally, the presence of a Tunnel Physical Interface Card (PIC) in an RP router creates a de-encapsulation interface, which allows the RP to receive multicast traffic from the source. This interface is indicated by **pd-0/0/0.32769**.

#### Related Documentation

- [Understanding PIM Sparse Mode on page 5509](#)
- [show pim interfaces on page 5883](#)
- [show pim rps on page 5911](#)



# Using PIM Bootstrap Router

- [Understanding the PIM Bootstrap Router on page 5567](#)
- [Configuring PIM Bootstrap Properties for IPv4 or IPv6 on page 5567](#)
- [Example: Rejecting PIM Bootstrap Messages at the Boundary of a PIM Domain on page 5569](#)
- [Example: Configuring PIM BSR Filters on page 5569](#)

## Understanding the PIM Bootstrap Router

---

To determine which router is the rendezvous point (RP), all routers within a PIM sparse-mode domain collect bootstrap messages. A PIM sparse-mode domain is a group of routers that all share the same RP router. The domain bootstrap router initiates bootstrap messages, which are sent hop by hop within the domain. The routers use bootstrap messages to distribute RP information dynamically and to elect a bootstrap router when necessary.

**Related  
Documentation**

- [Configuring PIM Bootstrap Properties for IPv4 or IPv6](#)

## Configuring PIM Bootstrap Properties for IPv4 or IPv6

---

For correct operation, every multicast router within a PIM domain must be able to map a particular multicast group address to the same rendezvous point (RP). The bootstrap router mechanism is one way that a multicast router can learn the set of group-to-RP mappings. Bootstrap routers are supported in IPv4 and IPv6.

To determine which routing device is the RP, all routing devices within a PIM domain collect bootstrap messages. A PIM domain is a contiguous set of routing devices that implement PIM. All devices are configured to operate within a common boundary. The domain's bootstrap router initiates bootstrap messages, which are sent hop by hop within the domain. The routing devices use bootstrap messages to distribute RP information dynamically and to elect a bootstrap router when necessary.

You can configure bootstrap properties globally or for a routing instance. This example shows the global configuration.

To configure the bootstrap router properties:

1. Configure the bootstrap priority.

By default, each routing device has a bootstrap priority of 0, which means the routing device can never be the bootstrap router. The routing device with the highest priority value is elected to be the bootstrap router. In the case of a tie, the routing device with the highest IP address is elected to be the bootstrap router. A simple bootstrap configuration assigns a bootstrap priority value to a routing device.



**NOTE:** In the IPv4-only configuration, specifying a bootstrap priority of 0 disables the bootstrap function and does not cause the routing device to send BSR packets with a 0 in the priority field. In the combined IPv4 and IPv6 configuration, specifying a bootstrap priority of 0 does not disable the function, but causes the routing device to send BSR packets with a 0 in the priority field. To disable the bootstrap function in the IPv4 and IPv6 configuration, delete the `bootstrap` statement.

```
user@host# edit protocols pim rp
user@host# set bootstrap family inet priority 3
```

2. (Optional) Create import and export policies to control the flow of bootstrap messages to and from the RP, and apply the policies to PIM. Import and export policies are useful when some of the routers in your PIM domain have interfaces that connect to other PIM domains. Configuring a policy prevents bootstrap messages from crossing domain boundaries. The **import** statement prevents messages from being imported into the RP. The **export** statement prevents messages from being exported from the RP.

```
[edit protocols pim rp]
user@host# set bootstrap family inet import pim-bootstrap-import
user@host# set bootstrap family inet export pim-bootstrap-export
user@host# exit
```

3. Configure the policies.

```
user@host# edit policy-options policy-statement pim-bootstrap-import
[edit policy-options policy-statement pim-bootstrap-import]
user@host# set from interface se-0/0/0
user@host# set then reject
user@host# exit
user@host# edit policy-options policy-statement pim-bootstrap-export
user@host# set from interface se-0/0/0
user@host# set then reject
user@host# exit
```

4. Monitor the operation of PIM bootstrap routers by running the `show pim bootstrap` command.

**Related Documentation**

- [Understanding PIM Sparse Mode on page 5509](#)



- [Example: Rejecting PIM Bootstrap Messages at the Boundary of a PIM Domain on page 5569](#)
- [show pim bootstrap on page 5881](#) in the CLI Explorer

## Example: Rejecting PIM Bootstrap Messages at the Boundary of a PIM Domain

In this example, the **from interface so-0-1/0 then reject** policy statement rejects bootstrap messages from the specified interface (the example is configured for both IPv4 and IPv6 operation):

```
protocols {
 pim {
 rp {
 bootstrap {
 family inet {
 priority 1;
 import pim-import;
 export pim-export;
 }
 family inet6 {
 priority 1;
 import pim-import;
 export pim-export;
 }
 }
 }
 }
}
policy-options {
 policy-statement pim-import {
 from interface so-0/1/0;
 then reject;
 }
 policy-statement pim-export {
 to interface so-0/1/0;
 then reject;
 }
}
```

## Example: Configuring PIM BSR Filters

Configure a filter to prevent BSR messages from entering or leaving your network. Add this configuration to all routers:

```
protocols {
 pim {
 rp {
 bootstrap-import no-bsr;
 bootstrap-export no-bsr;
 }
 }
}
```

```
policy-options {
 policy-statement no-bsr {
 then reject;
 }
}
```

# Using PIM Filtering

- [Understanding Multicast Message Filters on page 5571](#)
- [Filtering MAC Addresses on page 5572](#)
- [Filtering RP and DR Register Messages on page 5572](#)
- [Configuring Interface-Level PIM Neighbor Policies on page 5573](#)
- [Filtering Outgoing PIM Join Messages on page 5574](#)
- [Filtering Incoming PIM Join Messages on page 5575](#)
- [Configuring Register Message Filters on a PIM RP and DR on page 5577](#)

## Understanding Multicast Message Filters

---

Multicast sources and routers generate a considerable number of control messages, especially when using PIM sparse mode. These messages form distribution trees, locate rendezvous points (RPs) and designated routers (DRs), and transition from one type of tree to another. In most cases, this multicast messaging system operates transparently and efficiently. However, in some configurations, more control over the sending and receiving of multicast control messages is necessary.

You can configure multicast filtering to control the sending and receiving of multicast control messages.

To prevent unauthorized groups and sources from registering with an RP router, you can define a routing policy to reject PIM register messages from specific groups and sources and configure the policy on the designated router or the RP router.

- If you configure the reject policy on an RP router, it rejects incoming PIM register messages from the specified groups and sources. The RP router also sends a register stop message by means of unicast to the designated router. On receiving the register stop message, the designated router sends periodic null register messages for the specified groups and sources to the RP router.
- If you configure the reject policy on a designated router, it stops sending PIM register messages for the specified groups and sources to the RP router.



**NOTE:** If you have configured the reject policy on an RP router, we recommend that you configure the same policy on all the RP routers in your multicast network.



**NOTE:** If you delete a group and source address from the reject policy configured on an RP router and commit the configuration, the RP router will register the group and source only when the designated router sends a null register message.

**Related  
Documentation**

- [Filtering MAC Addresses on page 5572](#)
- [Filtering RP and DR Register Messages on page 5572](#)
- [Filtering MSDP SA Messages on page 5600](#)

---

## Filtering MAC Addresses

When a router is exclusively configured with multicast protocols on an interface, multicast sets the interface media access control (MAC) filter to multicast promiscuous mode, and the number of multicast groups is unlimited. However, when the router is not exclusively used for multicasting and other protocols such as OSPF, Routing Information Protocol version 2 (RIPv2), or Network Time Protocol (NTP) are configured on an interface, each of these protocols individually requests that the interface program the MAC filter to pick up its respective multicast group only. In this case, without multicast configured on the interface, the maximum number of multicast MAC filters is limited to 20. For example, the maximum number of interface MAC filters for protocols such as OSPF (multicast group 224.0.0.5) is 20, unless a multicast protocol is also configured on the interface.

No configuration is necessary for MAC filters.

---

## Filtering RP and DR Register Messages

You can filter Protocol Independent Multicast (PIM) register messages sent from the designated router (DR) or to the rendezvous point (RP). The PIM RP keeps track of all active sources in a single PIM sparse mode domain. In some cases, more control over which sources an RP discovers, or which sources a DR notifies other RPs about, is desired. A high degree of control over PIM register messages is provided by RP and DR register message filtering. Message filtering also prevents unauthorized groups and sources from registering with an RP router.

Register messages that are filtered at a DR are not sent to the RP, but the sources are available to local users. Register messages that are filtered at an RP arrive from source DRs, but are ignored by the router. Sources on multicast group traffic can be limited or directed by using RP or DR register message filtering alone or together.

If the action of the register filter policy is to discard the register message, the router needs to send a register-stop message to the DR. Register-stop messages are throttled to prevent malicious users from triggering them on purpose to disrupt the routing process.

Multicast group and source information is encapsulated inside unicast IP packets. This feature allows the router to inspect the multicast group and source information before sending or accepting the PIM register message.

Incoming register messages to an RP are passed through the configured register message filtering policy before any further processing. If the register message is rejected, the RP router sends a register-stop message to the DR. When the DR receives the register-stop message, the DR stops sending register messages for the filtered groups and sources to the RP. Two fields are used for register message filtering:

- Group multicast address
- Source address

The syntax of the existing policy statements is used to configure the filtering on these two fields. The **route-filter** statement is useful for multicast group address filtering, and the **source-address-filter** statement is useful for source address filtering. In most cases, the action is to **reject** the register messages, but more complex filtering policies are possible.

Filtering cannot be performed on other header fields, such as DR address, protocol, or port. In some configurations, an RP might not send register-stop messages when the policy action is to discard the register messages. This has no effect on the operation of the feature, but the router will continue to receive register messages.

When anycast RP is configured, register messages can be sent or received by the RP. All the RPs in the anycast RP set need to be configured with the same RP register message filtering policies. Otherwise, it might be possible to circumvent the filtering policy.

#### Related Documentation

- [Understanding RP Mapping with Anycast RP on page 5551](#)
- [Configuring Register Message Filters on a PIM RP and DR on page 5577](#)

## Configuring Interface-Level PIM Neighbor Policies

You can configure a policy to filter unwanted PIM neighbors. In the following example, the PIM interface compares neighbor IP addresses with the IP address in the policy statement before any hello processing takes place. If any of the neighbor IP addresses (primary or secondary) match the IP address specified in the prefix list, PIM drops the hello packet and rejects the neighbor.

If you configure a PIM neighbor policy after PIM has already established a neighbor adjacency to an unwanted PIM neighbor, the adjacency remains intact until the neighbor hold time expires. When the unwanted neighbor sends another hello message to update its adjacency, the router recognizes the unwanted address and rejects the neighbor.

To configure a policy to filter unwanted PIM neighbors:

1. Configure the policy. The neighbor policy must be a properly structured policy statement that uses a prefix list (or a route filter) containing the neighbor primary address (or any secondary IP addresses) in a prefix list, and the **reject** option to reject the unwanted address.

```
[edit policy-options]
user@host# set prefix-list nbrGroup 1 20.20.20.1/32
user@host# set policy-statement nbr-policy from prefix-list nbrGroup1
user@host# set policy-statement nbr-policy then reject
```

2. Configure the interface globally or in the routing instance. This example shows the configuration for the routing instance.

```
[edit routing-instances PIM.master protocols pim]
user@host# set neighbor-policy nbr-policy
```

3. Verify the configuration by checking the **Hello dropped on neighbor policy** field in the output of the **show pim statistics** command.

- Related Documentation**
- [Understanding PIM Sparse Mode on page 5509](#)
  - [show pim statistics on page 5921](#)

---

## Filtering Outgoing PIM Join Messages

When the core of your network is using MPLS, PIM join and prune messages stop at the customer edge (CE) routers and are not forwarded toward the core, because these routers do not have PIM neighbors on the core-facing interfaces. When the core of your network is using IP, PIM join and prune messages are forwarded to the upstream PIM neighbors in the core of the network.

When the core of your network is using a mix of IP and MPLS, you might want to filter certain PIM join and prune messages at the upstream egress interface of the CE routers.

You can filter PIM sparse mode (PIM-SM) join and prune messages at the egress interfaces for IPv4 and IPv6 in the upstream direction. The messages can be filtered based on the group address, source address, outgoing interface, PIM neighbor, or a combination of these values. If the filter is removed, the join is sent after the PIM periodic join timer expires.

To filter PIM sparse mode join and prune messages at the egress interfaces, create a policy rejecting the group address, source address, outgoing interface, or PIM neighbor, and then apply the policy.

The following example filters PIM join and prune messages for group addresses 224.0.1.2 and 225.1.1.1.

1. In configuration mode, create the policy.

```
user@host# set policy-options policy-statement block-groups term t1 from route-filter
224.0.1.2/32 exact
user@host# set policy-options policy-statement block-groups term t1 from route-filter
225.1.1.1/32 exact
```

```

user@host# set policy-options policy-statement block-groups term t1 then reject
user@host# set policy-options policy-statement block-groups term last then accept

```

2. Verify the policy configuration by running the **show policy-options** command.

```

user@host# show policy-options
policy-statement block-groups {
 term t1 {
 from {
 route-filter 224.0.1.2/32 exact;
 route-filter 225.1.1.1/32 exact;
 then reject;
 }
 term last {
 then accept;
 }
 }
}

```

3. Apply the PIM join and prune message filter.

```

user@host> set protocols pim export block-groups

```

4. After the configuration is committed, use the **show pim statistics** command to verify that outgoing PIM join and prune messages are being filtered.

```

user@host> show pim statistics | grep filtered
RP Filtered Source 0

Rx Joins/Prunes filtered 0

Tx Joins/Prunes filtered 254

```

The egress filter count is shown on the **Tx Joins/Prunes filtered** line.

**Related Documentation**

- [Filtering Incoming PIM Join Messages on page 5575](#)

## Filtering Incoming PIM Join Messages

Multicast scoping controls the propagation of multicast messages. Whereas multicast scoping prevents the actual multicast data packets from flowing in or out of an interface, PIM join filters prevent a state from being created in a router. A state—the (\*,G) or (S,G) entries—is the information used for forwarding unicast or multicast packets. Using PIM join filters prevents the transport of multicast traffic across a network and the dropping of packets at a scope at the edge of the network. Also, PIM join filters reduce the potential for denial-of-service (DoS) attacks and PIM state explosion—large numbers of PIM join messages forwarded to each router on the rendezvous-point tree (RPT), resulting in memory consumption.

To use PIM join filters to efficiently restrict multicast traffic from certain source addresses, create and apply the routing policy across all routers in the network.

See [Table 437](#) for a list of match conditions.

Table 437: PIM Join Filter Match Conditions

| Match Condition              | Matches On                                                                           |
|------------------------------|--------------------------------------------------------------------------------------|
| <b>interface</b>             | Router interface or interfaces specified by name or IP address                       |
| <b>neighbor</b>              | Neighbor address (the source address in the IP header of the join and prune message) |
| <b>route-filter</b>          | Multicast group address embedded in the join and prune message                       |
| <b>source-address-filter</b> | Multicast source address embedded in the join and prune message                      |

The following example shows how to create a PIM join filter. The filter is composed of a route filter and a source address filter—**bad-groups** and **bad-sources**, respectively. the **bad-groups** filter prevents (\*G) or (S,G) join messages from being received for all groups listed. The **bad-sources** filter prevents (S,G) join messages from being received for all sources listed. The **bad-groups** filter and **bad-sources** filter are in two different terms. If route filters and source address filters are in the same term, they are logically ANDed.

To filter incoming PIM join messages:

1. Configure the policy.

```
[edit policy-statement pim-join-filter term bad-groups]
user@host# set from route-filter 224.0.1.2/32 exact
user@host# set from route-filter 239.0.0.0/8 orlonger
user@host# set then reject

[edit policy-statement pim-join-filter term bad-sources]
user@host# set from source-address-filter 10.0.0.0/8 orlonger
user@host# set from source-address-filter 127.0.0.0/8 orlonger
user@host# set then reject

[edit policy-statement pim-join-filter term last]
user@host# set then accept
```

2. Apply one or more policies to routes being imported into the routing table from PIM.

```
[edit protocols pim]
user@host# set import pim-join-filter
```

3. Verify the configuration by checking the output of the **show pim join** and **show policy** commands.

#### Related Documentation

- [Understanding Multicast Administrative Scoping](#)
- [Filtering Outgoing PIM Join Messages on page 5574](#)
- [show pim join on page 5886](#) in the [CLI Explorer](#)
- [show policy](#) in the [CLI Explorer](#)



## Configuring Register Message Filters on a PIM RP and DR

PIM register messages are sent to the rendezvous point (RP) by a designated router (DR). When a source for a group starts transmitting, the DR sends unicast PIM register packets to the RP.

Register messages have the following purposes:

- Notify the RP that a source is sending to a group.
- Deliver the initial multicast packets sent by the source to the RP for delivery down the shortest-path tree (SPT).

The PIM RP keeps track of all active sources in a single PIM sparse mode domain. In some cases, you want more control over which sources an RP discovers, or which sources a DR notifies other RPs about. A high degree of control over PIM register messages is provided by RP or DR register message filtering. Message filtering prevents unauthorized groups and sources from registering with an RP router.

You configure RP or DR register message filtering to control the number and location of multicast sources that an RP discovers. You can apply register message filters on a DR to control outgoing register messages, or apply them on an RP to control incoming register messages.

When anycast RP is configured, all RPs in the anycast RP set need to be configured with the same register message filtering policy.

You can configure message filtering globally or for a routing instance. These examples show the global configuration.

To configure an RP filter to drop the register packets for multicast group range 224.1.1.0/24 from source address 10.10.94.2:

1. On the RP, configure the policy.

```
[edit policy-options policy-statement incoming-policy-for-rp from]
user@host# set route-filter 224.1.1.0/24 orlonger
user@host# set source-address-filter 10.10.94.2/32 exact
user@host# set then reject
user@host# exit
```

2. Apply the policy to the RP.

```
[edit protocols pim rp]
user@host# set rp-register-policy incoming-policy-for-rp
user@host# set local address 10.10.10.5
user@host# exit
```

To configure a DR filter to prevent sending register packets for group range 224.1.1.0/24 and source address 10.10.10.1/32:

1. On the DR, configure the policy.

```
[edit policy-options policy-statement outgoing-policy-for-rp]
```

```
user@host# set from route-filter 224.1.1.0/24 orlonger
user@host# set from source-address-filter 10.10.10.1/32 exact
user@host# set then reject
user@host# exit
```

2. Apply the policy to the DR.

The static address is the address of the RP to which you do not want the DR to send the filtered register messages.

```
[edit protocols pim rp]
user@host# set dr-register-policy outgoing-policy-for-dr
user@host# set static 10.10.10.3
user@host# exit
```

To configure a policy expression to accept register messages for multicast group 224.1.1.5 but reject those for 224.1.1.1:

1. On the RP, configure the policies.

```
[edit policy-options policy-statement reject_224_1_1_1]
user@host# set from route-filter 224.1.1.0/24 orlonger
user@host# set from source-address-filter 10.10.94.2/32 exact
user@host# set then reject
user@host# exit

[edit policy-options policy-statement accept_224_1_1_5]
user@host# set term one from route-filter 224.1.1.5/32 exact
user@host# set term one from source-address-filter 10.10.94.2/32 exact
user@host# set term one then accept
user@host# set term two then reject
user@host# exit
```

2. Apply the policies to the RP.

```
[edit protocols pim rp]
user@host# set rp-register-policy [reject_224_1_1_1 | accept_224_1_1_5]
user@host# set local address 10.10.10.5
```

To monitor the operation of the filters, run the **show pim statistics** command. The command output contains the following fields related to filtering:

- RP Filtered Source
- Rx Joins/Prunes filtered
- Tx Joins/Prunes filtered
- Rx Register msgs filtering drop
- Tx Register msgs filtering drop

#### Related Documentation

- [PIM Sparse Mode Source Registration on page 5581](#)
- [Filtering RP and DR Register Messages on page 5572](#)
- [show pim statistics on page 5921](#)

## Using PIM RPT and SPT Cutover

- [Understanding Multicast Rendezvous Points, Shared Trees, and Rendezvous-Point Trees on page 5579](#)
- [Building an RPT Between the RP and Receivers on page 5580](#)
- [PIM Sparse Mode Source Registration on page 5581](#)
- [Multicast Shortest-Path Tree on page 5584](#)
- [SPT Cutover on page 5585](#)
- [SPT Cutover Control on page 5588](#)
- [Example: Configuring the PIM Assert Timeout on page 5588](#)
- [Example: Configuring the PIM SPT Threshold Policy on page 5590](#)

### Understanding Multicast Rendezvous Points, Shared Trees, and Rendezvous-Point Trees

---

In a shared tree, the root of the distribution tree is a router, not a host, and is located somewhere in the core of the network. In the primary sparse mode multicast routing protocol, Protocol Independent Multicast sparse mode (PIM SM), the core router at the root of the shared tree is the rendezvous point (RP). Packets from the upstream source and join messages from the downstream routers “rendezvous” at this core router.

In the RP model, other routers do not need to know the addresses of the sources for every multicast group. All they need to know is the IP address of the RP router. The RP router discovers the sources for all multicast groups.

The RP model shifts the burden of finding sources of multicast content from each router (the (S,G) notation) to the network (the (\*,G) notation knows only the RP). Exactly how the RP finds the unicast IP address of the source varies, but there must be some method to determine the proper source for multicast content for a particular group.

Consider a set of multicast routers without any active multicast traffic for a certain group. When a router learns that an interested receiver for that group is on one of its directly connected subnets, the router attempts to join the distribution tree for that group back to the RP, not to the actual source of the content.

To join the shared tree, or *rendezvous-point tree (RPT)* as it is called in PIM sparse mode, the router must do the following:

- Determine the IP address of the RP for that group. Determining the address can be as simple as static configuration in the router, or as complex as a set of nested protocols.
- Build the shared tree for that group. The router executes an RPF check on the RP address in its routing table, which produces the interface closest to the RP. The router now detects that multicast packets from this RP for this group need to flow into the router on this RPF interface.
- Send a join message out on this interface using the proper multicast protocol (probably PIM sparse mode) to inform the upstream router that it wants to join the shared tree for that group. This message is a (\*,G) join message because S is not known. Only the RP is known, and the RP is not actually the source of the multicast packets. The router receiving the (\*,G) join message adds the interface on which the message was received to its outgoing interface list (OIL) for the group and also performs an RPF check on the RP address. The upstream router then sends a (\*,G) join message out from the RPF interface toward the source, informing the upstream router that it also wants to join the group.

Each upstream router repeats this process, propagating join messages from the RPF interface, building the shared tree as it goes. The process stops when the join message reaches one of the following:

- The RP for the group that is being joined
- A router along the RPT that already has a multicast forwarding state for the group that is being joined

In either case, the branch is created, and packets can flow from the source to the RP and from the RP to the receiver. Note that there is no guarantee that the shared tree (RPT) is the shortest path tree to the source. Most likely it is not. However, there are ways to “migrate” a shared tree to an SPT once the flow of packets begins. In other words, the forwarding state can transition from (\*,G) to (S,G). The formation of both types of tree depends heavily on the operation of the RPF check and the RPF table.

**Related  
Documentation**

- *Understanding Multicast Reverse Path Forwarding*

---

## Building an RPT Between the RP and Receivers

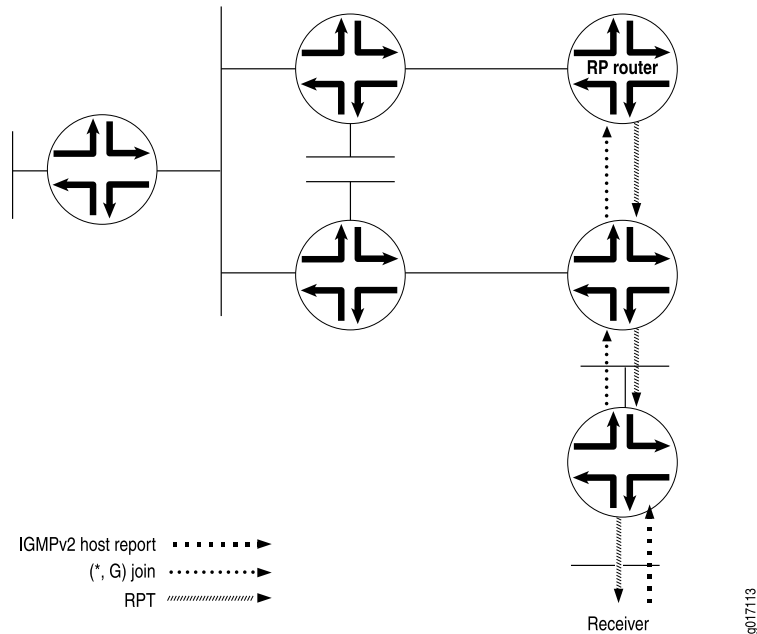
---

The RPT is the path between the RP and receivers (hosts) in a multicast group (see [Figure 184](#)). The RPT is built by means of a PIM join message from a receiver's DR:

1. A receiver sends a request to join group (G) in an Internet Group Management Protocol (IGMP) host membership report. A PIM sparse-mode router, the receiver's DR, receives the report on a directly attached subnet and creates an RPT branch for the multicast group of interest.
2. The receiver's DR sends a PIM join message to its RPF neighbor, the next-hop address in the RPF table, or the unicast routing table.
3. The PIM join message travels up the tree and is multicast to the ALL-PIM-ROUTERS group (224.0.0.13). Each router in the tree finds its RPF neighbor by using either the RPF table or the unicast routing table. This is done until the message reaches the RP

and forms the RPT. Routers along the path set up the multicast forwarding state to forward requested multicast traffic back down the RPT to the receiver.

**Figure 184: Building an RPT Between the RP and the Receiver**




## PIM Sparse Mode Source Registration

The RPT is a unidirectional tree, permitting traffic to flow down from the RP to the receiver in one direction. For multicast traffic to reach the receiver from the source, another branch of the distribution tree, called the shortest-path tree, needs to be built from the source's DR to the RP.

The shortest-path tree is created in the following way:


1. The source becomes active, sending out multicast packets on the LAN to which it is attached. The source's DR receives the packets and encapsulates them in a PIM register message, which it sends to the RP router (see [Figure 185](#)).
2. When the RP router receives the PIM register message from the source, it sends a PIM join message back to the source.

Figure 185: PIM Register Message and PIM Join Message Exchanged




3. The source's DR receives the PIM join message and begins sending traffic down the SPT toward the RP router (see [Figure 186](#)).
4. Once traffic is received by the RP router, it sends a register stop message to the source's DR to stop the register process.

Figure 186: Traffic Sent from the Source to the RP Router



5. The RP router sends the multicast traffic down the RPT toward the receiver (see [Figure 187](#)).

Figure 187: Traffic Sent from the RP Router Toward the Receiver



## Multicast Shortest-Path Tree

---

The distribution tree used for multicast is rooted at the source and is the shortest-path tree (SPT) as well. Consider a set of multicast routers without any active multicast traffic for a certain group (that is, they have no multicast forwarding state for that group). When a router learns that an interested receiver for that group is on one of its directly connected subnets, the router attempts to join the tree for that group.

To join the distribution tree, the router determines the unicast IP address of the source for that group. This address can be a simple static configuration on the router, or as complex as a set of protocols.

To build the SPT for that group, the router executes an a reverse path forwarding (RPF) check on the source address in its routing table. The RPF check produces the interface closest to the source, which is where multicast packets from this source for this group need to flow into the router.

The router next sends a join message out on this interface using the proper multicast protocol to inform the upstream router that it wants to join the distribution tree for that group. This message is an (S,G) join message because both S and G are known. The router receiving the (S,G) join message adds the interface on which the message was received to its output interface list (OIL) for the group and also performs an RPF check on the source address. The upstream router then sends an (S,G) join message out on the RPF interface toward the source, informing the upstream router that it also wants to join the group.

Each upstream router repeats this process, propagating joins out on the RPF interface, building the SPT as it goes. The process stops when the join message does one of two things:

- Reaches the router directly connected to the host that is the source.
- Reaches a router that already has multicast forwarding state for this source-group pair.

In either case, the branch is created, each of the routers has multicast forwarding state for the source-group pair, and packets can flow down the distribution tree from source to receiver. The RPF check at each router makes sure that the tree is an SPT.

SPTs are always the shortest path, but they are not necessarily short. That is, sources and receivers tend to be on the periphery of a router network, not on the backbone, and multicast distribution trees have a tendency to sprawl across almost every router in the network. Because multicast traffic can overwhelm a slow interface, and one packet can easily become a hundred or a thousand on the opposite side of the backbone, it makes sense to provide a shared tree as a distribution tree so that the multicast source can be located more centrally in the network, on the backbone. This sharing of distribution trees with roots in the core network is accomplished by a multicast rendezvous point.

### Related Documentation

- [Understanding Multicast Rendezvous Points, Shared Trees, and Rendezvous-Point Trees on page 5579](#)




## SPT Cutover

Instead of continuing to use the SPT to the RP and the RPT toward the receiver, a direct SPT is created between the source and the receiver in the following way:


1. Once the receiver's DR receives the first multicast packet from the source, the DR sends a PIM join message to its RPF neighbor (see [Figure 188](#)).
2. The source's DR receives the PIM join message, and an additional (S,G) state is created to form the SPT.
3. Multicast packets from that particular source begin coming from the source's DR and flowing down the new SPT to the receiver's DR. The receiver's DR is now receiving two copies of each multicast packet sent by the source—one from the RPT and one from the new SPT.

**Figure 188: Receiver DR Sends a PIM Join Message to the Source**




4. To stop duplicate multicast packets, the receiver's DR sends a PIM prune message toward the RP router, letting it know that the multicast packets from this particular source coming in from the RPT are no longer needed (see [Figure 189](#)).

**Figure 189: PIM Prune Message Is Sent from the Receiver's DR Toward the RP Router**



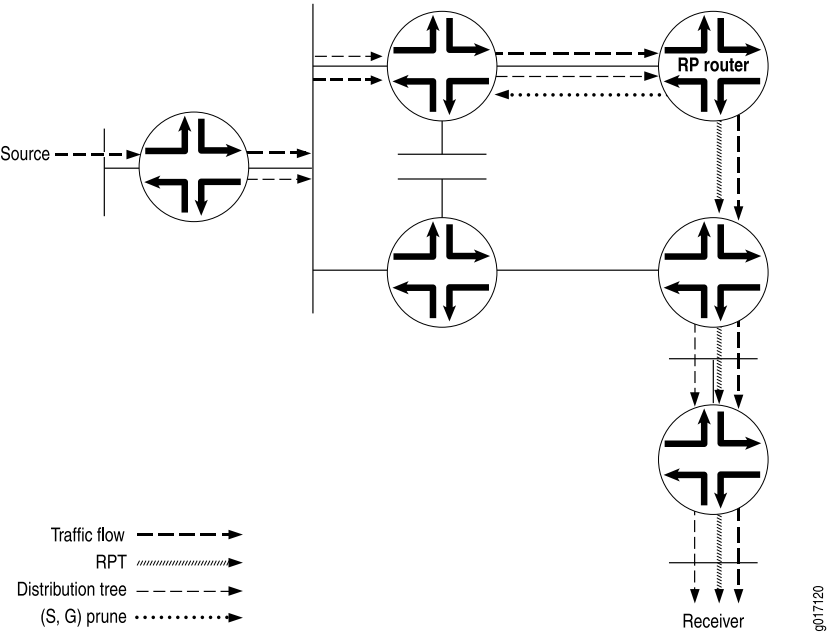
5. The PIM prune message is received by the RP router, and it stops sending multicast packets down to the receiver's DR. The receiver's DR is getting multicast packets only for this particular source over the new SPT. However, multicast packets from the source are still arriving from the source's DR toward the RP router (see [Figure 190](#)).

**Figure 190: RP Router Receives PIM Prune Message**




6. To stop the unneeded multicast packets from this particular source, the RP router sends a PIM prune message to the source's DR (see [Figure 191](#)).

Figure 191: RP Router Sends a PIM Prune Message to the Source DR



7. The receiver's DR now receives multicast packets only for the particular source from the SPT (see Figure 192).

Figure 192: Source's DR Stops Sending Duplicate Multicast Packets Toward the RP Router



## SPT Cutover Control

---

In some cases, the last-hop router needs to stay on the shared tree to the RP and not transition to a direct SPT to the source. You might not want the last-hop router to transition when, for example, a low-bandwidth multicast stream is forwarded from the RP to a last-hop router. All routers between last hop and source must maintain and refresh the SPT state. This can become a resource-intensive activity that does not add much to the network efficiency for a particular pair of source and multicast group addresses.

In these cases, you configure an SPT threshold policy on the last-hop router to control the transition to a direct SPT. An SPT cutover threshold of infinity applied to a source-group address pair means the last-hop router will never transition to a direct SPT. For all other source-group address pairs, the last-hop router transitions immediately to a direct SPT rooted at the source DR.

## Example: Configuring the PIM Assert Timeout

---

This example shows how to configure the timeout period for a PIM assert forwarder.

- [Requirements on page 5588](#)
- [Overview on page 5588](#)
- [Configuration on page 5590](#)

### Requirements

Before you begin:

- Configure the router interfaces. See the *Junos OS Network Interfaces Library for Routing Devices*.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Routing Devices*.
- Configure PIM Sparse Mode on the interfaces. See [“Enabling PIM Sparse Mode” on page 5512](#).

### Overview

The role of PIM assert messages is to determine the forwarder on a network with multiple routers. The forwarder is the router that forwards multicast packets to a network with multicast group members. The forwarder is generally the same as the PIM DR.

A router sends an assert message when it receives a multicast packet on an interface that is listed in the outgoing interface list of the matching routing entry. Receiving a message on an outgoing interface is an indication that more than one router forwards the same multicast packets to a network.

In [Figure 193](#), both routing devices R1 and R2 forward multicast packets for the same (S,G) entry on a network. Both devices detect this situation and both devices send assert messages on the Ethernet network. An assert message contains, in addition to a source

address and group address, a unicast cost metric for sending packets to the source, and a preference metric for the unicast cost. The preference metric expresses a preference between unicast routing protocols. The routing device with the smallest preference metric becomes the forwarder (also called the assert winner). If the preference metrics are equal, the device that sent the lowest unicast cost metric becomes the forwarder. If the unicast metrics are also equal, the routing device with the highest IP address becomes the forwarder. After the transmission of assert messages, only the forwarder continues to forward messages on the network.


When an assert message is received and the RPF neighbor is changed to the assert winner, the assert timer is set to an assert timeout period. The assert timeout period is restarted every time a subsequent assert message for the route entry is received on the incoming interface. When the assert timer expires, the routing device resets its RPF neighbor according to its unicast routing table. Then, if multiple forwarders still exist, the forwarders reenter the assert message cycle. In effect, the assert timeout period determines how often multicast routing devices enter a PIM assert message cycle.

The range is from 5 through 210 seconds. The default is 180 seconds.

Assert messages are useful for LANs that connect multiple routing devices and no hosts.

Figure 193 shows the topology for this example.

**Figure 193: PIM Assert Topology**



## Configuration

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure an assert timeout:

1. Configure the timeout period, in seconds.

```
[edit protocols pim]
user@host# set assert-timeout 60
```

2. (Optional) Trace assert messages.

```
[edit protocols pim]
user@host# set traceoptions file PIM.log
user@host# set traceoptions flag assert detail
```

3. If you are done configuring the device, commit the configuration.

```
user@host# commit
```

4. To verify the configuration, run the following commands:

- `show pim join`
- `show pim statistics`

**Related Documentation**

- [Configuring PIM Trace Options on page 5496](#)
- [SPT Cutover on page 5585](#)
- [SPT Cutover Control on page 5588](#)

---

## Example: Configuring the PIM SPT Threshold Policy

This example shows how to apply a policy that suppresses the transition from the rendezvous-point tree (RPT) rooted at the RP to the shortest-path tree (SPT) rooted at the source.

- [Requirements on page 5590](#)
- [Overview on page 5591](#)
- [Configuration on page 5592](#)
- [Verification on page 5594](#)

## Requirements

Before you begin:

- Configure the router interfaces. See the *Junos OS Network Interfaces Library for Routing Devices*.

- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Routing Devices*.
- Configure PIM Sparse Mode on the interfaces. See [“Enabling PIM Sparse Mode” on page 5512](#).

## Overview

Multicast routing devices running PIM sparse mode can forward the same stream of multicast packets onto the same LAN through an RPT rooted at the RP or through an SPT rooted at the source. In some cases, the last-hop routing device needs to stay on the shared RPT to the RP and not transition to a direct SPT to the source. Receiving the multicast data traffic on SPT is optimal but introduces more state in the network, which might not be desirable in some multicast deployments. Ideally, low-bandwidth multicast streams can be forwarded on the SPT, and high-bandwidth streams can use the SPT. This example shows how to configure such a policy.

This example includes the following settings:

- **spt-threshold**—Enables you to configure an SPT threshold policy on the last-hop routing device to control the transition to a direct SPT. When you include this statement in the main PIM instance, the PE router stays on the RPT for control traffic.
- **infinity**—Applies an SPT cutover threshold of infinity to a source-group address pair, so that the last-hop routing device never transitions to a direct SPT. For all other source-group address pairs, the last-hop routing device transitions immediately to a direct SPT rooted at the source DR. This statement must reference a properly configured policy to set the SPT cutover threshold for a particular source-group pair to infinity. The use of values other than infinity for the SPT threshold is not supported. You can configure more than one policy.
- **policy-statement**—Configures the policy. The simplest type of SPT threshold policy uses a route filter and source address filter to specify the multicast group and source addresses and to set the SPT threshold for that pair of addresses to infinity. The policy is applied to the main PIM instance.

This example sets the SPT transition value for the source-group pair 10.10.10.1 and 224.1.1.1 to infinity. When the policy is applied to the last-hop router, multicast traffic from this source-group pair never transitions to a direct SPT to the source. Traffic will continue to arrive through the RP. However, traffic for any other source-group address combination at this router transitions to a direct SPT to the source.

Note these points when configuring the SPT threshold policy:

- Configuration changes to the SPT threshold policy affect how the routing device handles the SPT transition.

Note these points when configuring the SPT threshold policy:

- Configuration changes to the SPT threshold policy affect how the routing device handles the SPT transition.

Note these points when configuring the SPT threshold policy:

- Configuration changes to the SPT threshold policy affect how the routing device handles the SPT transition.
- When the policy is configured for the first time, the routing device continues to transition to the direct SPT for the source-group address pair until the PIM-join state is cleared with the **clear pim join** command.
- If you do not clear the PIM-join state when you apply the infinity policy configuration for the first time, you must apply it before the PE router is brought up.
- When the policy is deleted for a source-group address pair for the first time, the routing device does not transition to the direct SPT for that source-group address pair until the PIM-join state is cleared with the **clear pim join** command.
- When the policy is changed for a source-group address pair for the first time, the routing device does not use the new policy until the PIM-join state is cleared with the **clear pim join** command.

## Configuration

### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
[edit]
set policy-options policy-statement spt-infinity-policy term one from route-filter
 224.1.1.1/32 exact
set policy-options policy-statement spt-infinity-policy term one from source-address-filter
 10.10.10.1/32 exact
set policy-options policy-statement spt-infinity-policy term one then accept
set policy-options policy-statement spt-infinity-policy term two then reject
set protocols pim spt-threshold infinity spt-infinity-policy
```

### Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see the *CLI User Guide*.

To configure an SPT threshold policy:

1. Apply the policy.

```
[edit]
user@host# edit protocols pim
```



```
[edit protocols pim]
user@host# set spt-threshold infinity spt-infinity-policy
[edit protocols pim]
user@host# exit
```

2. Configure the policy.

```
[edit]
user@host# edit policy-options policy-statement spt-infinity-policy
[edit policy-options policy-statement spt-infinity-policy]
user@host# set term one from route-filter 224.1.1.1/32 exact
[edit policy-options policy-statement spt-infinity-policy]
user@host# set term one from source-address-filter 10.10.10.1/32 exact
[edit policy-options policy-statement spt-infinity-policy]
user@host# set term one then accept
[edit policy-options policy-statement spt-infinity-policy]
user@host# set term two then reject
[edit policy-options policy-statement spt-infinity-policy]
user@host# exit
policy-statement {
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

4. Clear the PIM join cache to force the configuration to take effect.

```
[edit]
user@host# run clear pim join
```

## Results

Confirm your configuration by entering the **show policy-options** command and the **show protocols** command from configuration mode. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show policy-options
policy-statement spt-infinity-policy {
 term one {
 from {
 route-filter 224.1.1.1/32 exact;
 source-address-filter 10.10.10.1/32 exact;
 }
 then accept;
 }
 term two {
 then reject;
 }
}

user@host# show protocols
pim {
 spt-threshold {
 infinity spt-infinity-policy;
 }
}
```

## Verification

To verify the configuration, run the [show pim join](#) command.

**Related Documentation**

- [SPT Cutover Control on page 5588](#)

## PART 76

# Configuring MSDP

- [Using MSDP on page 5597](#)



# Using MSDP

- [Understanding MSDP on page 5597](#)
- [Configuring MSDP on page 5598](#)
- [Filtering MSDP SA Messages on page 5600](#)
- [Tracing MSDP Protocol Traffic on page 5600](#)
- [Configuring the Interface to Accept Traffic from a Remote Source on page 5602](#)
- [Example: Configuring MSDP on page 5603](#)
- [Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 5604](#)
- [Example: Configuring PIM Anycast With or Without MSDP on page 5610](#)
- [Configuring a PIM Anycast RP Router with MSDP on page 5613](#)

## Understanding MSDP

---

The Multicast Source Discovery Protocol (MSDP) is used to connect multicast routing domains. It typically runs on the same router as the Protocol Independent Multicast (PIM) sparse-mode rendezvous point (RP). Each MSDP router establishes adjacencies with internal and external MSDP peers similar to the way BGP establishes peers. These peer routers inform each other about active sources within the domain. When they detect active sources, the routers can send PIM sparse-mode explicit join messages to the active source.

The peer with the higher IP address passively listens to a well-known port number and waits for the side with the lower IP address to establish a Transmission Control Protocol (TCP) connection. When a PIM sparse-mode RP that is running MSDP becomes aware of a new local source, it sends source-active type, length, and values (TLVs) to its MSDP peers. When a source-active TLV is received, a peer-reverse-path-forwarding (peer-RPF) check (not the same as a multicast RPF check) is done to make sure that this peer is in the path that leads back to the originating RP. If not, the source-active TLV is dropped. This TLV is counted as a “rejected” source-active message.

The MSDP peer-RPF check is different from the normal RPF checks done by non-MSDP multicast routers. The goal of the peer-RPF check is to stop source-active messages from looping. Router R accepts source-active messages originated by Router S only from neighbor Router N or an MSDP mesh group member. For more information about configuring MSDP mesh groups, see [“Example: Configuring MSDP with Active Source Limits and Mesh Groups” on page 5604](#).

Router R locates its MSDP peer-RPF neighbor (Router N) deterministically. A series of rules is applied in a particular order to received source-active messages, and the first rule that applies determines the peer-RPF neighbor. All source-active messages from other routers are rejected.

The six rules applied to source-active messages originating at Router S received at Router R from Router X are as follows:

1. If Router X originated the source-active message (Router X is Router S), then Router X is also the peer-RPF neighbor, and its source-active messages are accepted.
2. If Router X is a member of the Router R mesh group, or is the configured peer, then Router X is the peer-RPF neighbor, and its source-active messages are accepted.
3. If Router X is the BGP next hop of the active multicast RPF route toward Router S (Router X installed the route on Router R), then Router X is the peer-RPF neighbor, and its source-active messages are accepted.
4. If Router X is an external BGP (EBGP) or internal BGP (IBGP) peer of Router R, and the last autonomous system (AS) number in the BGP AS-path to Router S is the same as Router X's AS number, then Router X is the peer-RPF neighbor, and its source-active messages are accepted.
5. If Router X uses the same next hop as the next hop to Router S, then Router X is the peer-RPF neighbor, and its source-active messages are accepted.
6. If Router X fits none of these criteria, then Router X is not an MSDP peer-RPF neighbor, and its source-active messages are rejected.

The MSDP peers that receive source-active TLVs can be constrained by BGP reachability information. If the AS path of the network layer reachability information (NLRI) contains the receiving peer's AS number prepended second to last, the sending peer is using the receiving peer as a next hop for this source. If the split horizon information is not being received, the peer can be pruned from the source-active TLV distribution list.

#### Related Documentation

- [Configuring MSDP on page 5598](#)

---

## Configuring MSDP

To configure the Multicast Source Discovery Protocol (MSDP), include the **msdp** statement:

```
msdp {
 disable;
 active-source-limit {
 maximum number;
 threshold number;
 }
 data-encapsulation (disable | enable);
 export [policy-names];
 group group-name {
 ... group-configuration ...
 }
}
```

```

hold-time seconds;
import [policy-names];
local-address address;
keep-alive seconds;
peer address {
 ... peer-configuration ...
}
rib-group group-name;
source ip-prefix </prefix-length> {
 active-source-limit {
 maximum number;
 threshold number;
 }
}
sa-hold-time seconds;
traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
}
group group-name {
 disable;
 export [policy-names];
 import [policy-names];
 local-address address;
 mode (mesh-group | standard);
 peer address {
 ... same statements as at the [edit protocols msdp peer address] hierarchy level shown
 just following ...
 }
 traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
 }
}
peer address {
 disable;
 active-source-limit {
 maximum number;
 threshold number;
 }
 authentication-key peer-key;
 default-peer;
 export [policy-names];
 import [policy-names];
 local-address address;
 traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
 }
}
}

```

You can include this statement at the following hierarchy levels:

- [\[edit protocols\]](#)
- [\[edit routing-instances \*routing-instance-name\* protocols\]](#)
- [\[edit logical-systems \*logical-system-name\* protocols\]](#)
- [\[edit logical-systems \*logical-system-name\* routing-instances \*routing-instance-name\* protocols\]](#)

By default, MSDP is disabled.

- Related Documentation**
- [Example: Configuring MSDP in a Routing Instance](#)
  - [Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 5604](#)

## Filtering MSDP SA Messages

Along with applying MSDP source active (SA) filters on all external MSDP sessions (in and out) to prevent SAs for groups and sources from leaking in and out of the network, you need to apply bootstrap router (BSR) filters. Applying a BSR filter to the boundary of a network prevents foreign BSR messages (which announce RP addresses) from leaking into your network. Since the routers in a PIM sparse-mode domain need to know the address of only one RP router, having more than one in the network can create issues.

If you did not use multicast scoping to create boundary filters for all customer-facing interfaces, you might want to use PIM join filters. Multicast scopes prevent the actual multicast data packets from flowing in or out of an interface. PIM join filters prevent PIM sparse-mode state from being created in the first place. Since PIM join filters apply only to the PIM sparse-mode state, it might be more beneficial to use multicast scoping to filter the actual data.



**NOTE:** When you apply firewall filters, firewall action modifiers, such as **log**, **sample**, and **count**, work only when you apply the filter on an inbound interface. The modifiers do not work on an outbound interface.

- Related Documentation**
- [Filtering Incoming PIM Join Messages on page 5575](#)
  - [Example: Configuring PIM BSR Filters on page 5569](#)

## Tracing MSDP Protocol Traffic

Tracing operations record detailed messages about the operation of routing protocols, such as the various types of routing protocol packets sent and received, and routing policy actions. You can specify which trace operations are logged by including specific tracing flags. The following table describes the flags that you can include.

| Flag | Description           |
|------|-----------------------|
| all  | Trace all operations. |



| Flag                          | Description                              |
|-------------------------------|------------------------------------------|
| <b>general</b>                | Trace general events.                    |
| <b>keepalive</b>              | Trace keepalive messages.                |
| <b>normal</b>                 | Trace normal events.                     |
| <b>packets</b>                | Trace all MSDP packets.                  |
| <b>policy</b>                 | Trace policy processing.                 |
| <b>route</b>                  | Trace MSDP changes to the routing table. |
| <b>source-active</b>          | Trace source-active packets.             |
| <b>source-active-request</b>  | Trace source-active request packets.     |
| <b>source-active-response</b> | Trace source-active response packets.    |
| <b>state</b>                  | Trace state transitions.                 |
| <b>task</b>                   | Trace task processing.                   |
| <b>timer</b>                  | Trace timer processing.                  |

You can configure MSDP tracing for all peers, for all peers in a particular group, or for a particular peer.

In the following example, tracing is enabled for all routing protocol packets. Then tracing is narrowed to focus only on MSDP peers in a particular group. To configure tracing operations for MSDP:

1. (Optional) Configure tracing by including the **traceoptions** statement at the **[edit routing-options]** hierarchy level and set the **all-packets-trace** and **all** flags to trace all protocol packets.

```
[edit routing-options traceoptions]
user@host# set file all-packets-trace
user@host# set flag all
```

2. Configure the filename for the MSDP trace file.

```
[edit protocols msdp group groupa traceoptions]
user@host# set file msdp-trace
```

3. (Optional) Configure the maximum number of trace files.

```
[edit protocols msdp group groupa traceoptions]
user@host# set file files 5
```

4. (Optional) Configure the maximum size of each trace file.

```
[edit protocols msdp group groupa traceoptions]
```

```
user@host# set file size 1m
```

- (Optional) Enable unrestricted file access.

```
[edit protocols msdp group groupa traceoptions]
```

```
user@host# set file world-readable
```

- Configure tracing flags. Suppose you are troubleshooting issues with the source-active cache for **groupa**. The following example shows how to trace messages associated with the group address.

```
[edit protocols msdp group groupa traceoptions]
```

```
user@host# set flag source-active | match 230.0.0.3
```

- View the trace file.

```
user@host> file list /var/log
```

```
user@host> file show /var/log/msdp-trace
```


#### Related Documentation

- [Understanding MSDP on page 5597](#)
- *Tracing and Logging Junos OS Operations in the Junos OS Administration Library for Routing Devices*

## Configuring the Interface to Accept Traffic from a Remote Source

You can configure an incoming interface to accept multicast traffic from a remote source. A remote source is a source that is not on the same subnet as the incoming interface. [Figure 194](#) shows such a topology, where R2 connects to the R1 source on one subnet, and to the incoming interface on R3 (ge-0/3/0.0 in the figure) on another subnet.

**Figure 194: Accepting Multicast Traffic from a Remote Source**



In this topology R2 is a pass-through device not running PIM, so R3 is the first hop router for multicast packets sent from R1. Because R1 and R3 are in different subnets, the default behavior of R3 is to disregard R1 as a remote source. You can have R3 accept multicast traffic from R1, however, by enabling **accept-remote-source** on the target interface.

To accept traffic from a remote source:

- Identify the router and physical interface that you want to receive multicast traffic from the remote source.
- Configure the interface to accept traffic from the remote source.

```
[edit protocols pim interface ge-1/3/0.0]
```

```
user@host# set accept-remote-source
```



**NOTE:** If the interface you identified is not the only path from the remote source, you need to ensure that it is the best path. For example you can configure a static route on the receiver side PE router to the source, or you can prepend the AS path on the other possible routes:

```
[edit policy-options policy-statement as-path-prepend term prepend]
user@host# set from route-filter 192.168.0.0/16 orlonger
user@host# set from route-filter 172.16.0.0/16 orlonger
user@host# set then as-path-prepend "1 1 1"
```

3. Commit the configuration changes.
4. Confirm that the interface you configured accepts traffic from the remote source.

```
user@host# show pim statistics
```

#### Related Documentation

- *Example: Allowing MBGP MVPN Remote Sources*
- *Understanding Prepending AS Numbers to BGP AS Paths*
- [show pim statistics on page 5921](#)

## Example: Configuring MSDP

Configure a router to act as a PIM sparse-mode rendezvous point and an MSDP peer:

```
[edit]
routing-options {
 interface-routes {
 rib-group ifrg;
 }
 rib-groups {
 ifrg {
 import-rib [inet.0 inet.2];
 }
 mcrg {
 export-rib inet.2;
 import-rib inet.2;
 }
 }
}
protocols {
 bgp {
 group lab {
 type internal;
 family any;
 neighbor 192.168.6.18 {
 local-address 192.168.6.17;
 }
 }
 }
}
```

```
pim {
 dense-groups {
 224.0.1.39/32;
 224.0.1.40/32;
 }
 rib-group mcr;
 rp {
 local {
 address 192.168.1.1;
 }
 }
 interface all {
 mode sparse-dense;
 version 1;
 }
}
msdp {
 rib-group mcr;
 group lab {
 peer 192.168.6.18 {
 local-address 192.168.6.17;
 }
 }
}
```

---

## Example: Configuring MSDP with Active Source Limits and Mesh Groups

This example shows how to configure MSDP to filter source-active messages and limit the flooding of source-active messages.

- [Requirements on page 5604](#)
- [Overview on page 5604](#)
- [Configuration on page 5608](#)
- [Verification on page 5609](#)

### Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Routing Devices*.
- Enable PIM sparse mode. See “[PIM Overview](#)” on [page 5491](#).
- Configure the router as a PIM sparse-mode RP. See “[Configuring Local PIM RPs](#)” on [page 5545](#).

### Overview

A router interested in MSDP messages, such as an RP, might have to process a large number of MSDP messages, especially source-active messages, arriving from other

routers. Because of the potential need for a router to examine, process, and create state tables for many MSDP packets, there is a possibility of an MSDP-based denial-of-service (DoS) attack on a router running MSDP. To minimize this possibility, you can configure the router to limit the number of source active messages the router accepts. Also, you can configure a threshold for applying random early detection (RED) to drop some but not all MSDP active source messages.

By default, the router accepts 25,000 source active messages before ignoring the rest. The limit can be from 1 through 1,000,000. The limit is applied to both the number of messages and the number of MSDP peers.

By default, the router accepts 24,000 source-active messages before applying the RED profile to prevent a possible DoS attack. This number can also range from 1 through 1,000,000. The next 1000 messages are screened by the RED profile and the accepted messages processed. If you configure no drop profiles (as this example does not), RED is still in effect and functions as the primary mechanism for managing congestion. In the default RED drop profile, when the packet queue fill-level is 0 percent, the drop probability is 0 percent. When the fill-level is 100 percent, the drop probability is 100 percent.



**NOTE:** The router ignores source-active messages with encapsulated TCP packets. Multicast does not use TCP; segments inside source-active messages are most likely the result of worm activity.

The number configured for the threshold must be less than the number configured for the maximum number of active MSDP sources.

You can configure an active source limit globally, for a group, or for a peer. If active source limits are configured at multiple levels of the hierarchy (as shown in this example), all are applied.

You can configure an active source limit for an address range as well as for a specific peer. A per-source active source limit uses an IP prefix and prefix length instead of a specific address. You can configure more than one per-source active source limit. The longest match determines the limit.

Per-source active source limits can be combined with active source limits at the peer, group, and global (instance) hierarchy level. Per-source limits are applied before any other type of active source limit. Limits are tested in the following order:

- Per-source
- Per-peer or group
- Per-instance

An active source message must “pass” all limits established before being accepted. For example, if a source is configured with an active source limit of 10,000 active multicast groups and the instance is configured with a limit of 5000 (and there are no other sources or limits configured), only 5000 active source messages are accepted from this source.

MSDP mesh groups are groups of peers configured in a full-mesh topology that limits the flooding of source-active messages to neighboring peers. Every mesh group member must have a peer connection with every other mesh group member. When a source-active message is received from a mesh group member, the source-active message is always accepted but is not flooded to other members of the same mesh group. However, the source-active message is flooded to non-mesh group peers or members of other mesh groups. By default, standard flooding rules apply if **mesh-group** is not specified.



**CAUTION:** When configuring MSDP mesh groups, you must configure all members the same way. If you do not configure a full mesh, excessive flooding of source-active messages can occur.

A common application for MSDP mesh groups is peer-reverse-path-forwarding (peer-RPF) check bypass. For example, if there are two MSDP peers inside an autonomous system (AS), and only one of them has an external MSDP session to another AS, the internal MSDP peer often rejects incoming source-active messages relayed by the peer with the external link. Rejection occurs because the external MSDP peer must be reachable by the internal MSDP peer through the next hop toward the source in another AS, and this next-hop condition is not certain. To prevent rejections, configure an MSDP mesh group on the internal MSDP peer so it always accepts source-active messages.



**NOTE:** An alternative way to bypass the peer-RPF check is to configure a default peer. In networks with only one MSDP peer, especially stub networks, the source-active message always needs to be accepted. An MSDP default peer is an MSDP peer from which all source-active messages are accepted without performing the peer-RPF check. You can establish a default peer at the peer or group level by including the **default-peer** statement.


Table 438 explains how flooding is handled by peers in this example. .

**Table 438: Source-Active Message Flooding Explanation**

| Source-Active Message Received From | Source-Active Message Flooded To                     | Source-Active Message Not Flooded To |
|-------------------------------------|------------------------------------------------------|--------------------------------------|
| Peer 21                             | Peer 11, Peer 12, Peer 13, Peer 31, Peer 32          | Peer 22                              |
| Peer 11                             | Peer 21, Peer 22, Peer 31, Peer 32                   | Peer 12, Peer 13                     |
| Peer 31                             | Peer 21, Peer 22, Peer 11, Peer 12, Peer 13, Peer 32 | —                                    |

Figure 195 illustrates source-active message flooding between different mesh groups and peers within the same mesh group.

Figure 195: Source-Active Message Flooding



This example includes the following settings:

- **active-source-limit maximum 10000**—Applies a limit of 10,000 active sources to all other peers.
- **data-encapsulation disable**—On an RP router using MSDP, disables the default encapsulation of multicast data received in MSDP register messages inside MSDP source-active messages.

MSDP data encapsulation mainly concerns bursty sources of multicast traffic. Sources that send only one packet every few minutes have trouble with the timeout of state relationships between sources and their multicast groups (S,G). Routers lose data while they attempt to reestablish (S,G) state tables. As a result, multicast register messages contain data, and this data encapsulation in MSDP source-active messages can be turned on or off through configuration.

By default, MSDP data encapsulation is enabled. An RP running MSDP takes the data packets arriving in the source's register message and encapsulates the data inside an MSDP source-active message.

However, data encapsulation creates both a multicast forwarding cache entry in the **inet.1** table (this is also the forwarding table) and a routing table entry in the **inet.4** table. Without data encapsulation, MSDP creates only a routing table entry in the **inet.4** table. In some circumstances, such as the presence of Internet worms or other forms of DoS attack, the router's forwarding table might fill up with these entries. To prevent the forwarding table from filling up with MSDP entries, you can configure the router not to use MSDP data encapsulation. However, if you disable data encapsulation, the router ignores and discards the encapsulated data. Without data encapsulation, multicast applications with bursty sources having transmit intervals greater than about 3 minutes might not work well.

- **group MSDP-group local-address 10.1.2.3**—Specifies the address of the local router (this router).
- **group MSDP-group mode mesh-group**—Specifies that all peers belonging to the MSDP-group group are mesh group members.

- **group MSDP-group peer 10.10.10.10**—Prevents the sending of source-active messages to neighboring peer 10.10.10.10.
- **group MSDP-group peer 10.10.10.10 active-source-limit maximum 7500**—Applies a limit of 7500 active sources to MSDP peer 10.10.10.10 in group **MSDP-group**.
- **peer 10.0.0.1 active-source-limit maximum 5000 threshold 4000**—Applies a threshold of 4000 active sources and a limit of 5000 active sources to MSDP peer 10.0.0.1.
- **source 10.1.0.0/16 active-source-limit maximum 500**—Applies a limit of 500 active sources to any source on the 10.1.0.0/16 network.

## Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set protocols msdp data-encapsulation disable
set protocols msdp active-source-limit maximum 10000
set protocols msdp peer 10.0.0.1 active-source-limit maximum 5000
set protocols msdp peer 10.0.0.1 active-source-limit threshold 4000
set protocols msdp source 10.1.0.0/16 active-source-limit maximum 500
set protocols msdp group MSDP-group mode mesh-group
set protocols msdp group MSDP-group local-address 10.1.2.3
set protocols msdp group MSDP-group peer 10.10.10.10 active-source-limit maximum
7500
```

**Step-by-Step Procedure** The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure MSDP source active routes and mesh groups:

1. (Optional) Disable data encapsulation.

```
[edit protocols msdp]
user@host# set data-encapsulation disable
```

2. Configure the active source limits.

```
[edit protocols msdp]
user@host# set peer 10.0.0.1 active-source-limit maximum 5000 threshold 4000
user@host# set group MSDP-group peer 10.10.10.10 active-source-limit maximum
7500
user@host# set active-source-limit maximum 10000
user@host# set source 10.1.0.0/16 active-source-limit maximum 500
```

3. Configure the mesh group.

```
[edit protocols msdp]
user@host# set group MSDP-group mode mesh-group
user@host# set group MSDP-group peer 10.10.10.10
user@host# set group MSDP-group local-address 10.1.2.3
```



4. If you are done configuring the device, commit the configuration.

```
[edit routing-instances]
user@host# commit
```

## Results

Confirm your configuration by entering the **show protocols** command.

```
user@host# show protocols
msdp {
 data-encapsulation disable;
 active-source-limit {
 maximum 10000;
 }
 peer 10.0.0.1 {
 active-source-limit {
 maximum 5000;
 threshold 4000;
 }
 }
 source 10.1.0.0/16 {
 active-source-limit {
 maximum 500;
 }
 }
 group MSDP-group {
 mode mesh-group;
 local-address 10.1.2.3;
 peer 10.10.10.10 {
 active-source-limit {
 maximum 7500;
 }
 }
 }
}
```

## Verification

To verify the configuration, run the following commands:

- [show msdp source-active](#)
- [show msdp statistics](#)

### Related Documentation

- [Examples: Configuring MSDP](#)
- [Filtering MSDP SA Messages on page 5600](#)
- [Configuring Local PIM RPs on page 5545](#)

## Example: Configuring PIM Anycast With or Without MSDP

---

When you configure anycast RP, you bypass the restriction of having one active rendezvous point (RP) per multicast group, and instead deploy multiple RPs for the same group range. The RP routers share one unicast IP address. Sources from one RP are known to other RPs that use the Multicast Source Discovery Protocol (MSDP). Sources and receivers use the closest RP, as determined by the interior gateway protocol (IGP).

You can use anycast RP within a domain to provide redundancy and RP load sharing. When an RP stops operating, sources and receivers are taken to a new RP by means of unicast routing.

You can configure anycast RP to use PIM and MSDP for IPv4, or PIM alone for both IPv4 and IPv6 scenarios. Both are discussed in this section.

We recommend a static RP mapping with anycast RP over a bootstrap router and auto-RP configuration because it provides all the benefits of a bootstrap router and auto-RP without the complexity of the BSR and auto-RP mechanisms.

All systems on a subnet must run the same version of PIM.

The default PIM version can be version 1 or version 2, depending on the mode you are configuring. PIMv1 is the default RP mode (at the **[edit protocols pim rp static address address]** hierarchy level). However, PIMv2 is the default for interface mode (at the **[edit protocols pim interface interface-name]** hierarchy level). Explicitly configured versions override the defaults. This example explicitly configures PIMv2 on the interfaces.

The following example shows an anycast RP configuration for the RP routers, first with MSDP and then using PIM alone, and for non-RP routers.

1. For a network using an RP with MSDP, configure the RP using the **lo0** loopback interface, which is always up. Include the **address** statement and specify the unique and routable router ID and the RP address at the **[edit interfaces lo0 unit 0 family inet]** hierarchy level. In this example, the router ID is **198.58.3.254** and the shared RP address is **198.58.3.253**. Include the **primary** statement for the first address. Including the **primary** statement selects the router's primary address from all the preferred addresses on all interfaces.

```
interfaces {
 lo0 {
 description "PIM RP";
 unit 0 {
 family inet {
 address 198.58.3.254/32;
 primary;
 address 198.58.3.253/32;
 }
 }
 }
}
```

2. Specify the RP address. Include the **address** statement at the **[edit protocols pim rp local]** hierarchy level (the same address as the secondary **lo0** interface).

For all interfaces, include the **mode** statement to set the mode to **sparse** and the **version** statement to specify PIM version 2 at the **[edit protocols pim rp local interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by including the **disable** statement for that interface.

```
protocols {
 pim {
 rp {
 local {
 family inet;
 address 198.58.3.253;
 }
 }
 interface all {
 mode sparse;
 version 2;
 }
 interface fxp0.0 {
 disable;
 }
 }
}
```

3. Configure MSDP peering. Include the **peer** statement to configure the address of the MSDP peer at the **[edit protocols msdp]** hierarchy level. For MSDP peering, use the unique, primary addresses instead of the anycast address. To specify the local address for MSDP peering, include the **local-address** statement at the **[edit protocols msdp peer]** hierarchy level.

```
protocols {
 msdp {
 peer 198.58.3.250 {
 local-address 198.58.3.254;
 }
 }
}
```



**NOTE:** If you need to configure a PIM RP for both IPv4 and IPv6 scenarios, perform Step 4 and Step 5. Otherwise, go to Step 6.

4. Configure an RP using the **lo0** loopback interface, which is always up. Include the **address** statement to specify the unique and routable router address and the RP address at the **[edit interfaces lo0 unit 0 family inet]** hierarchy level. In this example, the router ID is **198.58.3.254** and the shared RP address is **198.58.3.253**. Include the **primary** statement on the first address. Including the **primary** statement selects the router's primary address from all the preferred addresses on all interfaces.

```
interfaces {
 lo0 {
 description "PIM RP";
```

```

unit 0 {
 family inet {
 address 198.58.3.254/32 {
 primary;
 }
 address 198.58.3.253/32;
 }
}

```

5. Include the **address** statement at the **[edit protocols pim rp local]** hierarchy level to specify the RP address (the same address as the secondary **lo0** interface).

For all interfaces, include the **mode** statement to set the mode to **sparse**, and the **version** statement to specify PIM version 2 at the **[edit protocols pim rp local interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by including the **disable** statement for that interface.

Include the **anycast-pim** statement to configure anycast RP without MSDP (for example, if IPv6 is used for multicasting). The other RP routers that share the same IP address are configured using the **rp-set** statement. There is one entry for each RP, and the maximum that can be configured is 15. For each RP, specify the routable IP address of the router and whether MSDP source active (SA) messages are forwarded to the RP.

MSDP configuration is not necessary for this type of IPv4 anycast RP configuration.

```

protocols {
 pim {
 rp {
 local {
 family inet {
 address 198.58.3.253;
 anycast-pim {
 rp-set {
 address 198.58.3.240;
 address 198.58.3.241 forward-msdp-sa;
 }
 local-address 198.58.3.254; #If not configured, use lo0 primary
 }
 }
 }
 }
 }
 interface all {
 mode sparse;
 version 2;
 }
 interface fxp0.0 {
 disable;
 }
}

```

6. Configure the non-RP routers. The anycast RP configuration for a non-RP router is the same whether MSDP is used or not. Specify a static RP by adding the address at

the **[edit protocols pim rp static]** hierarchy level. Include the **version** statement at the **[edit protocols pim rp static address]** hierarchy level to specify PIM version 2.

```
protocols {
 pim {
 rp {
 static {
 address 198.58.3.253 {
 version 2;
 }
 }
 }
 }
}
```

7. Include the **mode** statement at the **[edit protocols pim interface all]** hierarchy level to specify sparse mode on all interfaces. Then include the **version** statement at the **[edit protocols pim rp interface all mode]** to configure all interfaces for PIM version 2. When configuring all interfaces, exclude the **fxp0.0** management interface by including the **disable** statement for that interface.

```
protocols {
 pim {
 interface all {
 mode sparse;
 version 2;
 }
 interface fxp0.0 {
 disable;
 }
 }
}
```

## Configuring a PIM Anycast RP Router with MSDP

Add the **address** statement at the **[edit protocols pim rp local]** hierarchy level to specify the RP address (the same address as the secondary **lo0** interface).

For all interfaces, use the **mode** statement to set the mode to **sparse** and the **version** statement to specify PIM version 2 at the **[edit protocols pim rp local interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by adding the **disable** statement for that interface.

```
protocols {
 pim {
 rp {
 local {
 family inet;
 address 198.58.3.253;
 }
 interface all {
 mode sparse;
 version 2;
 }
 }
 }
}
```

```
 interface fxp0.0 {
 disable;
 }
 }
}
```

To configure MSDP peering, add the **peer** statement to configure the address of the MSDP peer at the **[edit protocols msdp]** hierarchy level. For MSDP peering, use the unique, primary addresses instead of the anycast address. To specify the local address for MSDP peering, add the **local-address** statement at the **[edit protocols msdp peer]** hierarchy level.

```
protocols {
 msdp {
 peer 198.58.3.250 {
 local-address 198.58.3.254;
 }
 }
}
```

## PART 77

# Configuration Statements and Operational Commands

- [Configuration Statements \(IGMP\) on page 5617](#)
- [Configuration Statements \(IGMP Snooping\) on page 5643](#)
- [Configuration Statements \(MLD Snooping\) on page 5665](#)
- [Configuration Statements \(PIM\) on page 5687](#)
- [Configuration Statements \(Source-Specific Multicast\) on page 5761](#)
- [Configuration Statements \(MSDP\) on page 5767](#)
- [Operational Commands \(IGMP\) on page 5789](#)
- [Operational Commands \(IGMP Snooping\) on page 5813](#)
- [Operational Commands \(PIM\) on page 5825](#)
- [Operational Commands \(MSDP\) on page 5931](#)





## Configuration Statements (IGMP)

- [accounting \(Protocols IGMP\) on page 5618](#)
- [accounting \(Protocols IGMP Interface\) on page 5618](#)
- [asm-override-ssm on page 5619](#)
- [disable \(Protocols IGMP\) on page 5619](#)
- [exclude \(Protocols IGMP\) on page 5620](#)
- [group \(Protocols IGMP\) on page 5621](#)
- [group-count on page 5622](#)
- [group-increment \(Protocols IGMP\) on page 5622](#)
- [group-limit \(Protocols IGMP\) on page 5623](#)
- [group-policy \(Protocols IGMP\) on page 5623](#)
- [igmp on page 5624](#)
- [immediate-leave \(Protocols IGMP\) on page 5626](#)
- [interface \(Protocols IGMP\) on page 5627](#)
- [maximum-transmit-rate \(Protocols IGMP\) on page 5628](#)
- [oif-map \(IGMP Interface\) on page 5628](#)
- [passive \(IGMP\) on page 5629](#)
- [promiscuous-mode \(Protocols IGMP\) on page 5630](#)
- [query-interval \(Protocols IGMP\) on page 5631](#)
- [query-last-member-interval \(Protocols IGMP\) on page 5632](#)
- [query-response-interval \(Protocols IGMP\) on page 5633](#)
- [robust-count \(Protocols IGMP\) on page 5634](#)
- [source \(Protocols IGMP\) on page 5635](#)
- [source-count \(Protocols IGMP\) on page 5636](#)
- [source-increment \(Protocols IGMP\) on page 5637](#)
- [static \(Protocols IGMP\) on page 5638](#)
- [traceoptions \(Protocols IGMP\) on page 5639](#)
- [version \(Protocols IGMP\) on page 5641](#)

## accounting (Protocols IGMP)

---

|                                 |                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | accounting;                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">igmp</a> ],<br>[edit protocols <a href="#">igmp</a> ]                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series. |
| <b>Description</b>              | Enable the collection of IGMP join and leave event statistics on the system.                                                                                                               |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Recording IGMP Join and Leave Events on page 5425</a></li></ul>                                                                        |

## accounting (Protocols IGMP Interface)

---

|                                 |                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | (accounting   no-accounting);                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">igmp interface</a> <i>interface-name</i> ],<br>[edit protocols <a href="#">igmp interface</a> <i>interface-name</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.   |
| <b>Description</b>              | Enable or disable the collection of IGMP join and leave event statistics for an interface.                                                                                                   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Recording IGMP Join and Leave Events on page 5425</a></li></ul>                                                                          |

## asm-override-ssm

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | asm-override-ssm;                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],<br>[edit logical-systems <i>logical-system-name</i> routing-options multicast],<br>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],<br>[edit routing-options multicast]                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4.<br>Statement introduced in Junos OS Release 9.5 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| <b>Description</b>              | Enable the routing device to accept any-source multicast join messages (*G) for group addresses that are within the default or configured range of source-specific multicast groups.                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Source-Specific Multicast Groups with Any-Source Override on page 5537</a></li> </ul>                                                                                                                                                                                                                                                             |

## disable (Protocols IGMP)


|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | disable;                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp interface</b> <i>interface-name</i> ],<br>[edit protocols <b>igmp interface</b> <i>interface-name</i> ]                     |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series. |
| <b>Description</b>              | Disable IGMP on the system.                                                                                                                                                                    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Disabling IGMP on page 5429</a></li> </ul>                                                                                                |

## exclude (Protocols IGMP)

---

|                                 |                                                                                                                                                                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | exclude;                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> interface <i>interface-name</i> <b>static group</b> <i>multicast-group-address</i> ],<br>[edit protocols <b>igmp</b> interface <i>interface-name</i> <b>static group</b> <i>multicast-group-address</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.3.                                                                                                                                                                                                                                    |
| <b>Description</b>              | Configure the static group to operate in exclude mode. In exclude mode all sources except the address configured are accepted for the group. If this statement is not included, the group operates in include mode.                                                              |
| <b>Required Privilege Level</b> | view-level—To view this statement in the configuration.<br>control-level—To add this statement to the configuration.                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Enabling IGMP Static Group Membership on page 5418</a></li></ul>                                                                                                                                                             |

## group (Protocols IGMP)

|                                                                                                                                                                                                                                                                                                                                   |                                                                                                                                                                                                                                         |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                     | <pre>group <i>multicast-group-address</i> {   exclude;   group-count <i>number</i>;   group-increment <i>increment</i>;   source <i>ip-address</i> {     source-count <i>number</i>;     source-increment <i>increment</i>;   } }</pre> |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                            | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">igmp interface interface-name static</a> ],<br>[edit protocols <a href="#">igmp interface interface-name static</a> ]                                            |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                        | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                          |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                | Specify the IGMP multicast group address and (optionally) the source address for the multicast group being statically configured on an interface.                                                                                       |
| <hr/> <div style="display: flex; align-items: center;">  <div> <p><b>NOTE:</b> You must specify a unique address for each group.</p> </div> </div> <hr/> <p>The remaining statements are explained separately.</p> |                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                   | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                     |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                      | <ul style="list-style-type: none"> <li><a href="#">Enabling IGMP Static Group Membership on page 5418</a></li> </ul>                                                                                                                    |

## group-count

---

|                          |                                                                                                                                                                                                                                                                    |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <code>group-count <i>number</i>;</code>                                                                                                                                                                                                                            |
| Hierarchy Level          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> interface <i>interface-name</i> <b>static group multicast-group-address</b> ],<br>[edit protocols <b>igmp</b> interface <i>interface-name</i> <b>static group multicast-group-address</b> ] |
| Release Information      | Statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                 |
| Description              | Specify the number of static groups to be created.                                                                                                                                                                                                                 |
| Options                  | <i>number</i> —Number of static groups.<br><b>Default:</b><br><b>Range:</b> 1 through 512                                                                                                                                                                          |
| Required Privilege Level | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                |
| Related Documentation    | <ul style="list-style-type: none"><li>• <a href="#">Enabling IGMP Static Group Membership on page 5418</a></li></ul>                                                                                                                                               |

## group-increment (Protocols IGMP)

---

|                          |                                                                                                                                                                                                                                                                    |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <code>group-increment <i>increment</i>;</code>                                                                                                                                                                                                                     |
| Hierarchy Level          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> interface <i>interface-name</i> <b>static group multicast-group-address</b> ],<br>[edit protocols <b>igmp</b> interface <i>interface-name</i> <b>static group multicast-group-address</b> ] |
| Release Information      | Statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                 |
| Description              | Configure the number of times the address should be incremented for each static group created. The increment is specified in dotted decimal notation similar to an IPv4 address.                                                                                   |
| Options                  | <i>increment</i> —Number of times the address should be incremented.<br><b>Default:</b> 0.0.0.1<br><b>Range:</b> 0.0.0.1 through 255.255.255.255                                                                                                                   |
| Required Privilege Level | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                |
| Related Documentation    | <ul style="list-style-type: none"><li>• <a href="#">Enabling IGMP Static Group Membership on page 5418</a></li></ul>                                                                                                                                               |

## group-limit (Protocols IGMP)

|                                 |                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>group-limit <i>limit</i>;</code>                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp interface</b> <i>interface-name</i> ],<br>[edit protocols <b>igmp interface</b> <i>interface-name</i> ]                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.4.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                      |
| <b>Description</b>              | Configure a limit for the number of multicast groups (or [S,G] channels in IGMPv3) allowed on an interface. After this limit is reached, new reports are ignored and all related flows are not flooded on the interface. |
| <b>Default</b>                  | By default, there is no limit to the number of multicast groups that can join the interface.                                                                                                                             |
| <b>Options</b>                  | <i>limit</i> —group limit value for the interface.<br><b>Range:</b> 1 through 32767                                                                                                                                      |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Limiting the Number of IGMP Multicast Group Joins on Logical Interfaces on page 5426</a></li> </ul>                                                                 |

## group-policy (Protocols IGMP)

|                                 |                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>group-policy [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp interface</b> <i>interface-name</i> ],<br>[edit protocols <b>igmp interface</b> <i>interface-name</i> ]                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.1.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                |
| <b>Description</b>              | When this statement is enabled on a router running IGMP version 2 (IGMPv2) or version 3 (IGMPv3), after the routing device receives an IGMP report, the routing device compares the group against the specified group policy and performs the action configured in that policy (for example, rejects the report). |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Filtering Unwanted IGMP Reports at the IGMP Interface Level on page 5414</a></li> </ul>                                                                                                                                                                      |

## igmp

---

**Syntax**    `igmp {  
    accounting;  
    interface interface-name {  
        disable;  
        (accounting | no-accounting);  
        group-limit limit;  
        group-policy [ policy-names ];  
        immediate-leave;  
        oif-map map-name;  
        passive;  
        promiscuous-mode;  
        ssm-map ssm-map-name;  
        ssm-map-policy ssm-map-policy-name;  
        static {  
            group multicast-group-address {  
                exclude;  
                group-count number;  
                group-increment increment;  
                source ip-address {  
                    source-count number;  
                    source-increment increment;  
                }  
            }  
        }  
        version version;  
    }  
    query-interval seconds;  
    query-last-member-interval seconds;  
    query-response-interval seconds;  
    robust-count number;  
    traceoptions {  
        file filename <files number> <size size> <world-readable | no-world-readable>;  
        flag flag <flag-modifier> <disable>;  
    }  
}`

**Hierarchy Level**    [edit logical-systems *logical-system-name* protocols],  
                          [edit protocols]

**Release Information**    Statement introduced before Junos OS Release 7.4.  
                              Statement introduced in Junos OS Release 12.1 for the QFX Series.

**Description**    Enable IGMP on the router. IGMP must be enabled for the router to receive multicast packets.

The remaining statements are explained separately.


**Default**    IGMP is disabled on the router. IGMP is automatically enabled on all broadcast interfaces when you configure Protocol Independent Multicast (PIM) or Distance Vector Multicast Routing Protocol (DVMRP).



**Required Privilege** routing—To view this statement in the configuration.  
**Level** routing-control—To add this statement to the configuration.

**Related Documentation** • [Enabling IGMP on page 5409](#)

## immediate-leave (Protocols IGMP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | immediate-leave;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> interface <i>interface-name</i> ],<br>[edit protocols <b>igmp</b> interface <i>interface-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.3.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | <p>The immediate leave setting is useful for minimizing the leave latency of IGMP memberships. When this setting is enabled, the routing device leaves the multicast group immediately after the last host leaves the multicast group.</p> <p>The immediate leave setting enables host tracking, meaning that the device keeps track of the hosts that send join messages. This allows IGMP to determine when the last host sends a leave message for the multicast group.</p> <p>When the immediate leave setting is enabled, the device removes an interface from the forwarding-table entry without first sending IGMP group-specific queries to the interface. The interface is pruned from the multicast tree for the multicast group specified in the IGMP leave message. The immediate leave setting ensures optimal bandwidth management for hosts on a switched network, even when multiple multicast groups are being used simultaneously.</p> <p>When immediate leave is disabled and one host sends a leave group message, the routing device first sends a group query to determine if another receiver responds. If no receiver responds, the routing device removes all hosts on the interface from the multicast group. Immediate leave is disabled by default for both IGMP version 2 and IGMP version 3.</p> |
|                                 | <p> <b>NOTE:</b> Although host tracking is enabled for IGMPv2 and MLDv1 when you enable immediate leave, use immediate leave with these versions only when there is one host on the interface. The reason is that IGMPv2 and MLDv1 use a report suppression mechanism whereby only one host on an interface sends a group join report in response to a membership query. The other interested hosts suppress their reports. The purpose of this mechanism is to avoid a flood of reports for the same group. But it also interferes with host tracking, because the router only knows about the one interested host and does not know about the others.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">Specifying Immediate-Leave Host Removal for IGMP on page 5413</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

## interface (Protocols IGMP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> interface <i>interface-name</i> {   disable;   (accounting   no-accounting);   group-limit <i>limit</i>;   group-policy [ <i>policy-names</i> ];   immediate-leave;   oif-map <i>map-name</i>;   passive;   promiscuous-mode;   ssm-map <i>ssm-map-name</i>;   ssm-map-policy <i>ssm-map-policy-name</i>;   static {     group <i>multicast-group-address</i> {       exclude;       group-count <i>number</i>;       group-increment <i>increment</i>;       source <i>ip-address</i> {         source-count <i>number</i>;         source-increment <i>increment</i>;       }     }   }   version <i>version</i>; } </pre> |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> ],<br>[edit protocols <b>igmp</b> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Enable IGMP on an interface and configure interface-specific properties.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <p><b><i>interface-name</i></b>—Name of the interface. Specify the full interface name, including the physical and logical address components. To configure all interfaces, you can specify <b>all</b>.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Enabling IGMP on page 5409</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## maximum-transmit-rate (Protocols IGMP)

---


|                                 |                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | maximum-transmit-rate <i>packets-per-second</i> ;                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols igmp],<br>[edit protocols igmp]                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.3.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                 |
| <b>Description</b>              | Limit the transmission rate of IGMP packets                                                                                                                        |
| <b>Options</b>                  | <b>packets-per-second</b> —Maximum number of IGMP packets transmitted in one second by the router.<br><b>Range:</b> 1 through 10000<br><b>Default:</b> 500 packets |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Limiting the Maximum IGMP Message Rate on page 5418</a></li></ul>                                              |

## oif-map (IGMP Interface)

---

|                                 |                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | oif-map <i>map-name</i> ;                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">igmp interface interface-name</a> ],<br>[edit protocols <a href="#">igmp interface interface-name</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                             |
| <b>Description</b>              | Associates an outgoing interface (OIF) map to the IGMP interface. The OIF map is a routing policy statement that can contain multiple terms.                                   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Multicast with Subscriber VLANs</a></li></ul>                                                         |

## passive (IGMP)

|                                                                                                                                                                                                                                                                                                                                                                                 |                                                                                                                                                                                                                                                                                                      |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                   | <code>passive &lt;allow-receive&gt; &lt;send-general-query&gt; &lt;send-group-query&gt;;</code>                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                                          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> interface <i>interface-name</i> ],<br>[edit protocols <b>igmp</b> interface <i>interface-name</i> ]                                                                                                                           |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                      | Statement introduced in Junos OS Release 9.6.<br><b>allow-receive</b> , <b>send-general-query</b> , and <b>send-group-query</b> options were added in Junos OS Release 10.0.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                    |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                              | Specify that IGMP run on the interface and either not send and receive control traffic or selectively send and receive control traffic such as IGMP reports, queries, and leaves.                                                                                                                    |
| <div>  <p><b>NOTE:</b> You can selectively activate up to two out of the three available options for the <b>passive</b> statement while keeping the other functions passive (inactive). Activating all three options would be equivalent to not using the <b>passive</b> statement.</p> </div> |                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                                                                  | <p><b>allow-receive</b>—Enables IGMP to receive control traffic on the interface.</p> <p><b>send-general-query</b>—Enables IGMP to send general queries on the interface.</p> <p><b>send-group-query</b>—Enables IGMP to send group-specific and group-source-specific queries on the interface.</p> |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                                 | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                       |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                                                    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring Multicast with Subscriber VLANs</i></li> <li>• <a href="#">Enabling IGMP on page 5409</a></li> </ul>                                                                                                                                |

## promiscuous-mode (Protocols IGMP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>promiscuous-mode;</code>                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit dynamic-profiles <i>profile-name</i> protocols igmp interface <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> protocols <b>igmp interface</b> <i>interface-name</i> ],<br>[edit protocols <b>igmp interface</b> <i>interface-name</i> ]                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.3.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 9.2 for dynamic profiles.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.          |
| <b>Description</b>              | Specify that the interface accepts IGMP reports from hosts on any subnetwork. Note that when enabling promiscuous-mode, all routing devices on the ethernet segment must be configured with the promiscuous mode statement. Otherwise, only the interface configured with lowest IPv4 address acts as the querier for IGMP for this Ethernet segment. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Dynamic IGMP Configuration Overview</i></li><li>• <i>Configuring Dynamic DHCP Client Access to a Multicast Network</i></li><li>• <a href="#">Accepting IGMP Messages from Remote Subnetworks on page 5415</a></li></ul>                                                                                    |

## query-interval (Protocols IGMP)

|                                 |                                                                                                                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>query-interval seconds;</code>                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">igmp</a> ],<br>[edit protocols <a href="#">igmp</a> ]                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                           |
| <b>Description</b>              | Specify how often the querier routing device sends general host-query messages.                                                                                                                                                                                                                      |
| <b>Options</b>                  | <i>seconds</i> —Time interval.<br><b>Range:</b> 1 through 1024<br><b>Default:</b> 125 seconds                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Modifying the IGMP Host-Query Message Interval on page 5411</a></li> <li>• <a href="#">query-last-member-interval (Protocols IGMP) on page 5632</a></li> <li>• <a href="#">query-response-interval (Protocols IGMP) on page 5633</a></li> </ul> |

## query-last-member-interval (Protocols IGMP)

---

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | query-last-member-interval <i>seconds</i> ;                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">igmp</a> ],<br>[edit protocols <a href="#">igmp</a> ]                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.          |
| <b>Description</b>              | Specify how often the querier routing device sends group-specific query messages.                                                                                                                                                                                                   |
| <b>Options</b>                  | <b>seconds</b> —Time interval, in fractions of a second or seconds.<br><b>Range:</b> 0.1 through 0.9, then in 1-second intervals 1 through 999999<br><b>Default:</b> 1 second                                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Modifying the IGMP Last-Member Query Interval on page 5412</a></li><li>• <a href="#">query-interval (Protocols IGMP) on page 5631</a></li><li>• <a href="#">query-response-interval (Protocols IGMP) on page 5633</a></li></ul> |



## query-response-interval (Protocols IGMP)

|                                 |                                                                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>query-response-interval <i>seconds</i>;</code>                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">igmp</a> ],<br>[edit protocols <a href="#">igmp</a> ]                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.              |
| <b>Description</b>              | Specify how long the querier routing device waits to receive a response to a host-query message from a host.                                                                                                                                                                            |
| <b>Options</b>                  | <b><i>seconds</i></b> —The query response interval must be less than the query interval.<br><b>Range:</b> 1 through 1024<br><b>Default:</b> 10 seconds                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Modifying the IGMP Query Response Interval on page 5416</a></li> <li>• <a href="#">query-interval (Protocols IGMP) on page 5631</a></li> <li>• <a href="#">query-last-member-interval (Protocols IGMP) on page 5632</a></li> </ul> |

## robust-count (Protocols IGMP)

---

|                                 |                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>robust-count <i>number</i>;</code>                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">igmp</a> ],<br>[edit protocols <a href="#">igmp</a> ]                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| <b>Description</b>              | Tune the expected packet loss on a subnet. This factor is used to calculate the group member interval, other querier present interval, and last-member query count.                                                                                                        |
| <b>Options</b>                  | <i>number</i> —Robustness variable.<br><b>Range:</b> 2 through 10<br><b>Default:</b> 2                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Modifying the IGMP Robustness Variable on page 5417</a></li></ul>                                                                                                                                                      |

## source (Protocols IGMP)

---

|                                 |                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>source <i>ip-address</i> {     <i>source-count</i> <i>number</i>;     <i>source-increment</i> <i>increment</i>; }</pre>                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> <b>interface</b> <i>interface-name</i> <b>static group</b> <i>mcast-group-address</i> ],<br>[edit protocols <b>igmp</b> <b>interface</b> <i>interface-name</i> <b>static group</b> <i>mcast-group-address</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.             |
| <b>Description</b>              | Specify the IP version 4 (IPv4) unicast source address for the multicast group being statically configured on an interface.                                                                                                                                                            |
| <b>Options</b>                  | <i>ip-address</i> —IPv4 unicast address.<br><br>The remaining statements are explained separately.                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Enabling IGMP Static Group Membership on page 5418</a></li> </ul>                                                                                                                                                                 |

## source-count (Protocols IGMP)


---

|                                 |                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>source-count <i>number</i>;</code>                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> interface <i>interface-name</i> <b>static group multicast-group-address</b> <b>source</b> ],<br>[edit protocols <b>igmp</b> interface <i>interface-name</i> <b>static group multicast-group-address</b> <b>source</b> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                 |
| <b>Description</b>              | Configure the number of multicast source addresses that should be accepted for each static group created.                                                                                                                                                                                      |
| <b>Options</b>                  | <i>number</i> —Number of source addresses.<br><b>Default:</b> 1<br><b>Range:</b> 1 through 1024                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Enabling IGMP Static Group Membership on page 5418</a></li></ul>                                                                                                                                                                           |

## source-increment (Protocols IGMP)

|                                 |                                                                                                                                                                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | source-increment <i>number</i> ;                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> interface <i>interface-name</i> <b>static group multicast-group-address source</b> ],<br>[edit protocols <b>igmp</b> interface <i>interface-name</i> <b>static group multicast-group-address source</b> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                   |
| <b>Description</b>              | Configure the number of times the multicast source address should be incremented for each static group created. The increment is specified in dotted decimal notation similar to an IPv4 address.                                                                                |
| <b>Options</b>                  | <b>increment</b> —Number of times the source address should be incremented.<br><b>Default:</b> 0.0.0.1<br><b>Range:</b> 0.0.0.1 through 255.255.255.255                                                                                                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Enabling IGMP Static Group Membership on page 5418</a></li> </ul>                                                                                                                                                           |

## static (Protocols IGMP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>static {   group multicast-group-address {     exclude;     group-count number;     group-increment increment;     source ip-address {       source-count number;       source-increment increment;     }   } }</pre>                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp interface</b> <i>interface-name</i> ],<br>[edit protocols <b>igmp interface</b> <i>interface-name</i> ]                                                                                                                                                                                               |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                            |
| <b>Description</b>              | <p>Test multicast forwarding on an interface without a receiver host.</p> <p>The <b>static</b> statement simulates IGMP joins on a routing device statically on an interface without any IGMP hosts. It is supported for both IGMPv2 and IGMPv3 joins. This statement is especially useful for testing multicast forwarding on an interface without a receiver host.</p> |
|                                 | <p> <b>NOTE:</b> To prevent joining too many groups accidentally, the <b>static</b> statement is not supported with the <b>interface all</b> statement.</p>                                                                                                                           |
|                                 | <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | <p>routing and trace—To view this statement in the configuration.</p> <p>routing-control and trace-control—To add this statement to the configuration.</p>                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Enabling IGMP Static Group Membership on page 5418</a></li> </ul>                                                                                                                                                                                                                                                   |

## traceoptions (Protocols IGMP)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre> traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;flag-modifier&gt; &lt;disable&gt;; } </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> ],<br>[edit protocols <b>igmp</b> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>         | <p>Configure IGMP tracing options.</p> <p>To specify more than one tracing operation, include multiple <b>flag</b> statements.</p> <p>To trace the paths of multicast packets, use the <b>mtrace</b> command.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Default</b>             | The default IGMP trace options are those inherited from the routing protocols <b>traceoptions</b> statement included at the [edit routing-options] hierarchy level.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>             | <p><b>disable</b>—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>/var/log</b>. We recommend that you place tracing output in the file <b>igmp-log</b>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you must also include the <b>size</b> statement to specify the maximum file size.</p> <p><b>Range:</b> 2 through 1000 files</p> <p><b>Default:</b> 2 files</p> <p><b>flag</b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements.</p> <p><b>IGMP Tracing Flags</b></p> <ul style="list-style-type: none"> <li><b>leave</b>—Leave group messages (for IGMP version 2 only).</li> <li><b>mtrace</b>—Mtrace packets. Use the <b>mtrace</b> command to troubleshoot the software.</li> </ul> |

- **packets**—All IGMP packets.
- **query**—IGMP membership query messages, including general and group-specific queries.
- **report**—Membership report messages.

#### Global Tracing Flags

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

**Default:** If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

**flag-modifier**—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

**no-stamp**—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

**Default:** If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

**no-world-readable**—(Optional) Do not allow users to read the log file.

**replace**—(Optional) Replace an existing trace file if there is one.

**Default:** If you do not include this option, tracing output is appended to an existing trace file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When **trace-file** again reaches this size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.



If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of trace files.

**Syntax:** *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 1 MB

**world-readable**—(Optional) Allow any user to read the log file.

|                                 |                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing and trace—To view this statement in the configuration.                                                 |
|                                 | routing-control and trace-control—To add this statement to the configuration.                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Tracing IGMP Protocol Traffic on page 5427</a></li> </ul> |

## version (Protocols IGMP)

|                                 |                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>version <i>version</i>;</code>                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp interface</b> <i>interface-name</i> ],<br>[edit protocols <b>igmp interface</b> <i>interface-name</i> ]                                                                                                 |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| <b>Description</b>              | Specify the version of IGMP.                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <p><b>version</b>—IGMP version number.</p> <p><b>Range:</b> 1, 2, or 3</p> <p><b>Default:</b> IGMP version 2</p>                                                                                                                                                           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Changing the IGMP Version on page 5410</a></li> </ul>                                                                                                                                                                 |



## Configuration Statements (IGMP Snooping)

- [all](#) on page 5644
- [data-forwarding](#) on page 5644
- [disable \(IGMP Snooping\)](#) on page 5645
- [group \(IGMP Snooping\)](#) on page 5645
- [group-limit \(IGMP and MLD Snooping\)](#) on page 5646
- [host-only-interface](#) on page 5647
- [igmp-querier](#) on page 5647
- [igmp-snooping](#) on page 5648
- [immediate-leave \(Bridge Domains\)](#) on page 5649
- [interface \(Bridge Domains\)](#) on page 5650
- [interface \(IGMP Snooping\)](#) on page 5651
- [l2-querier](#) on page 5651
- [multicast-router-interface \(IGMP Snooping\)](#) on page 5652
- [query-interval \(Bridge Domains\)](#) on page 5653
- [query-last-member-interval \(Bridge Domains\)](#) on page 5654
- [query-response-interval \(Bridge Domains\)](#) on page 5655
- [receiver](#) on page 5656
- [robust-count \(IGMP Snooping\)](#) on page 5656
- [source-address](#) on page 5657
- [src-address \(IGMP Querier\)](#) on page 5658
- [source-vlans](#) on page 5658
- [static \(IGMP Snooping\)](#) on page 5659
- [traceoptions \(IGMP Snooping\)](#) on page 5660
- [version \(IGMP Snooping\)](#) on page 5662
- [vlan \(IGMP Snooping\)](#) on page 5663

## all

---

|                                 |                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | all;                                                                                                                |
| <b>Hierarchy Level</b>          | [edit protocols igmp-snooping vlan ]                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 15.1 for the QFX series.                                                   |
| <b>Description</b>              | Apply IGMP snooping to all configured VLANs.                                                                        |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration. |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring IGMP Snooping on page 5434</a></li></ul>            |

## data-forwarding

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>data-forwarding {<br/>    receiver {<br/>        source-vlans vlan-list;<br/>        install;<br/>    }<br/>    source {<br/>        groups group-prefix;<br/>    }<br/>}</pre>                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit protocols igmp-snooping vlan (all   <i>vlan-name</i> )]                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6 for EX Series switches.<br>Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | <p>Configure the VLAN to be a multicast source VLAN (MVLAN) or a multicast VLAN registration (MVR) receiver VLAN. Each data-forwarding VLAN, which can be a multicast source VLAN (MVLAN) or a multicast receiver VLAN, must have exactly one source statement or exactly one receiver statement. A data-forwarding VLAN can operate only in IGMP version 2 (IGMPv2) mode.</p> <p>The remaining statements are explained separately.</p> |
| <b>Default</b>                  | Disabled                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Multicast VLAN Registration</i></li><li>• <i>Configuring Multicast VLAN Registration (CLI Procedure)</i></li></ul>                                                                                                                                                                                                                                                       |

## disable (IGMP Snooping)

---

|                                 |                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>disable;</code>                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | <code>[edit protocols igmp-snooping vlan <i>vlan-name</i>]</code>                                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                     |
| <b>Description</b>              | Disable IGMP snooping on all interfaces in a VLAN.                                                                                                                                    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring IGMP Snooping on page 5436</a></li> <li>• <a href="#">Configuring IGMP Snooping on page 5434</a></li> </ul> |

## group (IGMP Snooping)

---

|                                 |                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>group <i>ip-address</i>;</code>                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | <code>[edit protocols igmp-snooping vlan <i>vlan-name</i> <b>interface</b> <i>interface-name</i> <b>static</b>]</code>                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                      |
| <b>Description</b>              | Configure a static multicast group using a valid IP multicast address.                                                                                                                                                                                 |
| <b>Default</b>                  | None.                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <i>ip-address</i> —IP address of the multicast group receiving data on an interface.                                                                                                                                                                   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show igmp-snooping vlans on page 5823</a></li> <li>• <a href="#">Example: Configuring IGMP Snooping on page 5436</a></li> <li>• <a href="#">Configuring IGMP Snooping on page 5434</a></li> </ul> |

## group-limit (IGMP and MLD Snooping)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>group-limit <i>limit</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | <code>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface</a> <i>interface-name</i>],</code><br><code>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> <a href="#">interface</a> <i>interface-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface</a> <i>interface-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols vlan <i>vlan-id</i> igmp-snooping <a href="#">interface</a> <i>interface-name</i>]</code> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Configure a limit for the number of multicast groups (or [S,G] channels in IGMPv3) allowed on an interface. After this limit is reached, new reports are ignored and all related flows are not flooded on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Default</b>                  | By default, there is no limit to the number of multicast groups joining an interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <i>limit</i> —a 32-bit number for the limit on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring IGMP Snooping</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## host-only-interface

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | host-only-interface;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface interface-name</a> ],<br>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> <a href="#">interface interface-name</a> ],<br>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface interface-name</a> ],<br>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols vlan <i>vlan-id</i> igmp-snooping <a href="#">interface interface-name</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Configure an interface as a host-facing interface. IGMP queries received on these interfaces are dropped.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Default</b>                  | The interface can either be a host-side or multicast-router interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring IGMP Snooping</i></li> <li>• <i>multicast-router-interface</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

## igmp-querier

|                                 |                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | igmp-querier <a href="#">source-address</a> <i>source address</i> ;                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit protocols igmp-snooping vlan <i>vlan-name</i> ]                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1X53-D15 for QFabric Systems.                                                                                                                                                                                                                                                        |
| <b>Description</b>              | Configure a QFabric Node device to be an IGMP querier. If there are any multicast routers on the same local network, make sure the source address for the IGMP querier is lower (a smaller number) than the IP addresses for those routers on the network. This ensures that Node is always the IGMP querier on the network.     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring IGMP Snooping on page 5436</a></li> <li>• <a href="#">Configuring IGMP Snooping on page 5434</a></li> <li>• <a href="#">show igmp-snooping vlans on page 5823</a></li> <li>• <a href="#">show configuration protocols igmp on page 5799</a></li> </ul> |

## igmp-snooping

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> igmp-snooping {   vlan <i>vlan-id</i> {     all     immediate-leave;     interface <i>interface-name</i> {       group-limit <i>limit</i>;       host-only-interface;       immediate-leave;       multicast-router-interface;       static {         group <i>ip-address</i> {           source <i>ip-address</i>;         }       }     }   }   l2-querier {     source-address <i>ip-address</i>;   }   proxy {     source-address <i>ip-address</i>;   }   query-interval <i>seconds</i>;   query-last-member-interval <i>seconds</i>;   query-response-interval <i>seconds</i>;   robust-count <i>number</i>;   traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;no-stamp&gt; &lt;replace&gt; &lt;size <i>size</i>&gt; &lt;world-readable         no-world-readable&gt;;     flag <i>flag</i> &lt;flag-modifier&gt;;   } } </pre> |
| <b>Hierarchy Level</b>          | [edit protocols]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Enable IGMP snooping on the router or switch.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Default</b>                  | IGMP snooping is disabled on the router or switch.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | The statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Understanding IGMP Snooping</i></li> <li>• <i>IGMP Snooping in MC-LAG Active-Active Mode</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |



## immediate-leave (Bridge Domains)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>immediate-leave;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>     | <p>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping],</p> <p>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <b>interface</b> <i>interface-name</i>],</p> <p>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <b>interface</b> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols vlan <i>vlan-id</i> igmp-snooping <b>interface</b> <i>interface-name</i>]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b> | Statement introduced in Junos OS Release 8.5.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>         | <p>The immediate leave setting is useful for minimizing the leave latency of IGMP memberships. When this setting is enabled, the routing device leaves the multicast group immediately after the last host leaves the multicast group.</p> <p>The immediate-leave setting enables host tracking, meaning that the device keeps track of the hosts that send join messages. This allows IGMP to determine when the last host sends a leave message for the multicast group.</p> <p>When the immediate leave setting is enabled, the device removes an interface from the forwarding-table entry without first sending IGMP group-specific queries to the interface. The interface is pruned from the multicast tree for the multicast group specified in the IGMP leave message. The immediate leave setting ensures optimal bandwidth management for hosts on a switched network, even when multiple multicast groups are being used simultaneously.</p> <p>When immediate leave is disabled and one host sends a leave group message, the routing device first sends a group query to determine if another receiver responds. If no receiver responds, the routing device removes all hosts on the interface from the multicast group. Immediate leave is disabled by default for both IGMP version 2 and IGMP version 3.</p> |



**NOTE:** Although host tracking is enabled for IGMPv2 and MLDv1 when you enable immediate leave, use immediate leave with these versions only when there is one host on the interface. The reason is that IGMPv2 and MLDv1 use a report suppression mechanism whereby only one host on an interface sends a group join report in response to a membership query. The other interested hosts suppress their reports. The purpose of this mechanism is to avoid a flood of reports for the same group. But it also interferes with host tracking, because the router only knows about the one interested host and does not know about the others.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation** • *Example: Configuring IGMP Snooping*

---

## interface (Bridge Domains)

---

**Syntax** `interface interface-name {  
 group-limit limit;  
 host-only-interface;  
 multicast-router-interface;  
 static {  
 group ip-address {  
 source ip-address;  
 }  
 }  
}`

**Hierarchy Level** [edit bridge-domains *bridge-domain-name* protocols igmp-snooping],  
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id*],  
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping],  
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols vlan *vlan-id* igmp-snooping]

**Release Information** Statement introduced in Junos OS Release 8.5.

**Description** Enable IGMP snooping on an interface and configure interface-specific properties.

**Options** *interface-name*—Name of the interface. Specify the full interface name, including the physical and logical address components. To configure all interfaces, you can specify **all**.

The remaining statements are explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation** • *Example: Configuring IGMP Snooping*

## interface (IGMP Snooping)

|                                 |                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>interface <i>interface-name</i> {     multicast-router-interface;     static {         group <i>ip-address</i>;     } }</pre>                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit protocols igmp-snooping vlan <i>vlan-name</i> ]                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                      |
| <b>Description</b>              | <p>Enable IGMP snooping on an interface and configure interface-specific properties.</p> <p>The remaining statements are explained separately.</p>                                                                                                     |
| <b>Options</b>                  | <i>interface-name</i> —Name of the interface.                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring IGMP Snooping on page 5436</a></li> <li>• <a href="#">Configuring IGMP Snooping on page 5434</a></li> <li>• <a href="#">show igmp-snooping vlans on page 5823</a></li> </ul> |

## l2-querier

|                                 |                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>l2-querier {     source-address <i>ip-address</i>; }</pre>                                                                                       |
| <b>Hierarchy Level</b>          | [edit protocols igmp-snooping vlan],                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2 for the QFX Series.                                                                                     |
| <b>Description</b>              | Configure the switch to be an IGMP querier. Use the <b>source-address</b> statement to configure the source address to use for IGMP snooping queries. |
| <b>Options</b>                  | <p><b>seconds</b>—Time interval.</p> <p><b>Range:</b> 1 through 1024</p> <p><b>Default:</b> 125 seconds</p>                                           |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                        |
| <b>Related Documentation</b>    |                                                                                                                                                       |

## multicast-router-interface (IGMP Snooping)

---

|                                 |                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | multicast-router-interface;                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit protocols igmp-snooping vlan <i>vlan-name</i> <b>interface</b> <i>interface-name</i> ]                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                  |
| <b>Description</b>              | Configure an interface to forward IGMP messages to multicast routers.                                                                                                                                                                              |
| <b>Default</b>                  | Disabled. If this statement is disabled, the interface drops IGMP messages it receives.                                                                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show igmp-snooping vlans on page 5823</a></li><li>• <a href="#">Example: Configuring IGMP Snooping on page 5436</a></li><li>• <a href="#">Configuring IGMP Snooping on page 5434</a></li></ul> |

## query-interval (Bridge Domains)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>query-interval seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | <pre>[edit bridge-domains <i>bridge-domain-name</i> protocols mld-snooping ] , [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <i>interface</i> <i>interface-name</i>], [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> <i>interface</i> <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <i>interface</i> <i>interface-name</i>],[ edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols mld-snooping ] [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> <i>interface</i> <i>interface-name</i>] [edit routing-instances <i>routing-instance-name</i> protocols mld-snooping ] [edit protocols igmp-snooping vlan]</pre> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX series.</p> <p>Statement introduced in Junos OS Release 14.2 for MX series Routers with MPC.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Configure the interval for host-query message timeouts.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <p><i>seconds</i>—Time interval. This value must be greater than the interval set for <code>query-response-interval</code>.</p> <p><b>Range:</b> 1 through 1024</p> <p><b>Default:</b> 125 seconds</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring IGMP Snooping</i></li> <li>• <a href="#">query-last-member-interval (Bridge Domains) on page 5654</a></li> <li>• <a href="#">query-response-interval (Bridge Domains) on page 5655</a></li> <li>• <i>mld-snooping</i></li> <li>• <i>igmp-snooping</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## query-last-member-interval (Bridge Domains)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>query-last-member-interval seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | <code>[edit bridge-domains <i>bridge-domain-name</i> protocols mld-snooping ] ,</code><br><code>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface</a></code><br><code>    <i>interface-name</i>],</code><br><code>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> <a href="#">interface</a></code><br><code>    <i>interface-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols</code><br><code>    igmp-snooping <a href="#">interface</a> <i>interface-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols</code><br><code>    mld-snooping ]</code><br><code>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols</code><br><code>    igmp-snooping vlan <i>vlan-id</i></code><br><code>    [edit routing-instances <i>routing-instance-name</i> protocols mld-snooping ]<a href="#">interface</a></code><br><code>        <i>interface-name</i></code><br><code>    [edit protocols igmp-snooping vlan],</code> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.<br>Statement introduced in Junos OS Release 13.2 for the QFX series.<br>Statement introduced in Junos OS Release 14.2 for MX series Routers with MPC.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Configure the interval for group-specific query timeouts.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <b>seconds</b> —Time interval, in fractions of a second or seconds.<br><b>Range:</b> 0.1 through 0.9, then in 1-second intervals 1 through 1024<br><b>Default:</b> 1 second                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring IGMP Snooping</i></li><li>• <a href="#">query-interval on page 5653</a></li><li>• <a href="#">query-response-interval on page 5655</a></li><li>• <i>mld-snooping</i></li><li>• <i>igmp-snooping</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## query-response-interval (Bridge Domains)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>query-response-interval seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | <pre>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface</a> <i>interface-name</i>], [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snoopingvlan <i>vlan-id</i> <a href="#">interface</a> <i>interface-name</i>], [edit bridge-domains <i>bridge-domain-name</i> protocols mld-snooping ] , [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface</a> <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snoopingvlan <i>vlan-id</i> [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols mld-snooping ]<a href="#">interface</a> <i>interface-name</i>] [edit routing-instances <i>routing-instance-name</i> protocols mld-snooping ] [edit protocols igmp-snooping vlan],</pre> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX series.</p> <p>Statement introduced in Junos OS Release 14.2 for MX series Routers with MPC.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Specify how long to wait to receive a response to a specific query message from a host.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <p><i>seconds</i>—Time interval. This interval should be less than the host-query interval.</p> <p><b>Range:</b> 1 through 1024</p> <p><b>Default:</b> 10 seconds</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring IGMP Snooping</i></li> <li>• <a href="#">query-interval (Bridge Domains) on page 5653</a></li> <li>• <a href="#">query-last-member-interval (Bridge Domains) on page 5654</a></li> <li>• <i>mld-snooping</i></li> <li>• <i>igmp-snooping</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

## receiver

---

|                                 |                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>receiver {<br/>    source-vlans <i>vlan-list</i>;<br/>    install;<br/>}</pre>                                                                                                |
| <b>Hierarchy Level</b>          | [edit protocols igmp-snooping vlan (all   <i>vlan-name</i> ) data-forwarding]                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6 for EX Series switches.<br>Statement introduced in Junos OS Release 12.3 for the QFX Series.                                          |
| <b>Description</b>              | Configure a VLAN as a multicast receiver VLAN of the multicast VLAN (MVLAN).<br><br>The remaining statements are explained separately.                                             |
| <b>Default</b>                  | Disabled                                                                                                                                                                           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Multicast VLAN Registration</i></li><li>• <i>Configuring Multicast VLAN Registration (CLI Procedure)</i></li></ul> |

## robust-count (IGMP Snooping)

---

|                                 |                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>robust-count <i>number</i>;</pre>                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit protocols igmp-snooping vlan <i>vlan-name</i> ]                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                  |
| <b>Description</b>              | Configure the number of intervals the switch waits before removing a multicast group from the multicast forwarding table. Configure the length of each interval using the <b>query-interval</b> statement.                                         |
| <b>Default</b>                  | 2 intervals                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <b><i>number</i></b> —Number of intervals the switch waits before timing out a multicast group.<br><b>Range:</b> 2 through 10                                                                                                                      |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring IGMP Snooping on page 5436</a></li><li>• <a href="#">Configuring IGMP Snooping on page 5434</a></li><li>• <a href="#">show igmp-snooping vlans on page 5823</a></li></ul> |



## source-address

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>source-address <i>ip-address</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | <p>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping proxy],</p> <p>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> proxy],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping proxy],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> proxy]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX series.</p>                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | <p>Specify the IP address to use as the source for IGMP snooping reports in proxy mode. Reports are sent with address 0.0.0.0 as the source address unless there is a source address configured. You can also use this statement to configure the source address to use for IGMP snooping queries.</p>                                                                                                                                                                                       |
| <b>Options</b>                  | <i>ip-address</i> —IP address to use as the source for proxy-mode IGMP snooping reports.                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring IGMP Snooping</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                |

## src-address (IGMP Querier)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>src-address source address;</code>                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit protocols igmp-snooping vlan <i>vlan-name</i> <a href="#">igmp-querier</a> ]<br>[edit protocols igmp-snooping vlan <i>vlan-name</i> <a href="#">l2-querier</a> ]                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.3 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D15 for QFabric Systems.                                                                                                                                                                                                                              |
| <b>Description</b>              | Configure the address that the switch uses as the source address in the IGMP queries that it sends. If there are any multicast routers on the same local network, make sure the source address for the IGMP querier is smaller (a lower number) than the IP addresses for those routers on the network. This ensures that switch is always the IGMP querier on the network. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring IGMP Snooping on page 5436</a></li><li>• <a href="#">Configuring IGMP Snooping on page 5434</a></li><li>• <a href="#">show igmp-snooping vlans on page 5823</a></li><li>• <a href="#">show configuration protocols igmp on page 5799</a></li></ul>                                                 |

## source-vlans

---

|                                 |                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>source-vlans <i>vlan-list</i>;</code>                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit protocols igmp-snooping vlan (all   <i>vlan-name</i> ) data-forwarding receiver]                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6 for EX Series switches.<br>Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                             |
| <b>Description</b>              | Specify a list of multicast VLANs (MVLANS) from which this multicast receiver VLAN receives multicast traffic. Either all of these MVLANS must be in proxy mode or none of them can be in proxy mode. |
| <b>Default</b>                  | Disabled                                                                                                                                                                                              |
| <b>Options</b>                  | <i>vlan-list</i> —Names of the MVLANS.                                                                                                                                                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Multicast VLAN Registration</a></li><li>• <a href="#">Configuring Multicast VLAN Registration (CLI Procedure)</a></li></ul>  |

---

## static (IGMP Snooping)

---

|                                 |                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>static {<br/>    group ip-address;<br/>}</pre>                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit protocols igmp-snooping vlan <i>vlan-name</i> interface <i>interface-name</i> ]                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                  |
| <b>Description</b>              | <p>Statically define multicast groups on an interface.</p> <p>The remaining statement is explained separately.</p>                                                                                                                                 |
| <b>Default</b>                  | No multicast groups are statically defined.                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring IGMP Snooping on page 5436</a></li><li>• <a href="#">Configuring IGMP Snooping on page 5434</a></li><li>• <a href="#">show igmp-snooping vlans on page 5823</a></li></ul> |

## traceoptions (IGMP Snooping)

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>traceoptions {<br/>    file <i>filename</i> &lt;files <i>number</i>&gt; &lt;no-stamp&gt; &lt;size <i>size</i>&gt; &lt;replace&gt; &lt;world-readable  <br/>    no-world-readable&gt;;<br/>    flag <i>flag</i> (detail   disable   receive   send);<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>     | For platforms without ELS:<br><br>[edit protocols igmp-snooping]<br><br>For platforms with ELS:<br><br>[edit protocols igmp-snooping vlan]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b> | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>         | Define tracing operations for IGMP snooping.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Default</b>             | The <b>traceoptions</b> feature is disabled by default.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>             | <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>/var/log</b>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached (<b>xk</b> to specify KB, <b>xm</b> to specify MB, or <b>xg</b> to specify gigabytes), at which point the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the <b>size</b> option.</p> <p><b>Range:</b> 2 through 1000</p> <p><b>Default:</b> 3 files</p> <p><b>flag <i>flag</i></b> —Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:</p> <ul style="list-style-type: none"><li>• <b>all</b>—All tracing operations.</li><li>• <b>general</b>—Trace general IGMP snooping protocol events.</li><li>• <b>krt</b>—Trace communication over routing sockets.</li><li>• <b>nexthop</b>— Trace next-hop related events.</li><li>• <b>normal</b>—Trace normal IGMP snooping protocol events.</li><li>• <b>packets</b>—Trace all IGMP packets.</li><li>• <b>policy</b>—Trace policy processing.</li><li>• <b>query</b>—Trace IGMP membership query messages.</li><li>• <b>report</b>—Trace membership report messages.</li></ul> |

- **route**—Trace routing information.
- **state**—Trace IGMP state transitions.
- **task**—Trace routing protocol task processing.
- **timer**—Trace routing protocol timer processing.
- **vlan**—Trace VLAN related events.

**no-stamp**—(Optional) Do not time stamp trace file.

**no-world-readable**—(Optional) Restrict file access to the user who created the file.

**size size** —(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **files** option. Use **xk** to specify KB, **xm** to specify MB, or **xg** to specify gigabytes.

**Range:** 10 KB through 1 gigabytes

**Default:** 128 KB

**world-readable**—(Optional) Enable unrestricted file access.

|                                 |                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring IGMP Snooping on page 5436</a></li> <li>• <a href="#">Configuring IGMP Snooping on page 5434</a></li> </ul> |

## version (IGMP Snooping)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>version <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit protocols igmp-snooping vlan (all   <i>vlan-name</i> )]                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Specify the IGMP version for the IGMP general query that the switch sends to hosts when an interface comes up. The configured IGMP version affects only the version of the general queries sent by a switch. It does not affect the version of IGMP messages that the switch can snoop. For example, If the switch is configured for IGMP version 1 (IGMPv1), it can snoop IGMPv2 and IGMPv3 messages. |
| <b>Default</b>                  | If you do not configure the <b>version</b> statement, the default is IGMPv2.                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <b>version</b> —IGMP version number.<br><b>Range:</b> 1 and 2.                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring IGMP Snooping (CLI Procedure)</a></li><li>• <a href="#">Configuring IGMP Snooping on page 5434</a></li></ul>                                                                                                                                                                                                                           |

## vlan (IGMP Snooping)

```
Syntax vlan vlan-name {
 immediate-leave;
 interface interface-name {
 group-limit limit;
 host-only-interface;
 multicast-router-interface;
 static {
 group mcast-group-address {
 source ip-address;
 }
 }
 }
 qualified-vlan ;
 proxy {
 source-address ip-address;
 }
 query-interval seconds;
 query-last-member-interval seconds;
 query-response-interval seconds;
 robust-count number;
 }
```

**Hierarchy Level** [edit protocols igmp-snooping ],

**Release Information** Statement introduced in Junos OS Release 8.5.  
Statement introduced in Junos OS Release 13.2 for the QFX series.

**Description** Configure IGMP snooping parameters for a particular VLAN.

**Default** By default, IGMP snooping options apply to all VLANs.

**Options** *vlan-name*—Apply the parameters to this VLAN.  
  
The remaining statements are explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring VLAN-Specific IGMP Snooping Parameters on page 5435](#)
- [igmp-snooping on page 5648](#)





# Configuration Statements (MLD Snooping)

- all on page 5666
- data-forwarding on page 5666
- disable (IGMP Snooping) on page 5667
- group (IGMP Snooping) on page 5667
- group-limit (IGMP and MLD Snooping) on page 5668
- host-only-interface on page 5669
- igmp-querier on page 5669
- igmp-snooping on page 5670
- immediate-leave (Bridge Domains) on page 5671
- interface (Bridge Domains) on page 5672
- interface (IGMP Snooping) on page 5673
- l2-querier on page 5673
- mld-snooping on page 5674
- multicast-router-interface (IGMP Snooping) on page 5675
- query-interval (Bridge Domains) on page 5676
- query-last-member-interval (Bridge Domains) on page 5677
- query-response-interval (Bridge Domains) on page 5678
- receiver on page 5679
- robust-count (IGMP Snooping) on page 5679
- source-address on page 5680
- src-address (IGMP Querier) on page 5681
- source-vlans on page 5681
- static (IGMP Snooping) on page 5682
- traceoptions (IGMP Snooping) on page 5683
- version (IGMP Snooping) on page 5685
- vlan (IGMP Snooping) on page 5686

## all

---

|                                 |                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | all;                                                                                                                |
| <b>Hierarchy Level</b>          | [edit protocols igmp-snooping vlan ]                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 15.1 for the QFX series.                                                   |
| <b>Description</b>              | Apply IGMP snooping to all configured VLANs.                                                                        |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration. |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring IGMP Snooping on page 5434</a></li></ul>            |

## data-forwarding

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>data-forwarding {<br/>  receiver {<br/>    source-vlans <i>vlan-list</i>;<br/>    install;<br/>  }<br/>  source {<br/>    groups <i>group-prefix</i>;<br/>  }<br/>}</pre>                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit protocols igmp-snooping vlan (all   <i>vlan-name</i> )]                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6 for EX Series switches.<br>Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | <p>Configure the VLAN to be a multicast source VLAN (MVLAN) or a multicast VLAN registration (MVR) receiver VLAN. Each data-forwarding VLAN, which can be a multicast source VLAN (MVLAN) or a multicast receiver VLAN, must have exactly one source statement or exactly one receiver statement. A data-forwarding VLAN can operate only in IGMP version 2 (IGMPv2) mode.</p> <p>The remaining statements are explained separately.</p> |
| <b>Default</b>                  | Disabled                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Multicast VLAN Registration</i></li><li>• <i>Configuring Multicast VLAN Registration (CLI Procedure)</i></li></ul>                                                                                                                                                                                                                                                       |

## disable (IGMP Snooping)

---

|                                 |                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>disable;</code>                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | <code>[edit protocols igmp-snooping vlan <i>vlan-name</i>]</code>                                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                     |
| <b>Description</b>              | Disable IGMP snooping on all interfaces in a VLAN.                                                                                                                                    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring IGMP Snooping on page 5436</a></li> <li>• <a href="#">Configuring IGMP Snooping on page 5434</a></li> </ul> |

## group (IGMP Snooping)

---

|                                 |                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>group <i>ip-address</i>;</code>                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | <code>[edit protocols igmp-snooping vlan <i>vlan-name</i> <b>interface</b> <i>interface-name</i> <b>static</b>]</code>                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                      |
| <b>Description</b>              | Configure a static multicast group using a valid IP multicast address.                                                                                                                                                                                 |
| <b>Default</b>                  | None.                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <i>ip-address</i> —IP address of the multicast group receiving data on an interface.                                                                                                                                                                   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show igmp-snooping vlans on page 5823</a></li> <li>• <a href="#">Example: Configuring IGMP Snooping on page 5436</a></li> <li>• <a href="#">Configuring IGMP Snooping on page 5434</a></li> </ul> |

## group-limit (IGMP and MLD Snooping)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>group-limit <i>limit</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | <code>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface</a> <i>interface-name</i>],</code><br><code>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> <a href="#">interface</a> <i>interface-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface</a> <i>interface-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols vlan <i>vlan-id</i> igmp-snooping <a href="#">interface</a> <i>interface-name</i>]</code> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Configure a limit for the number of multicast groups (or [S,G] channels in IGMPv3) allowed on an interface. After this limit is reached, new reports are ignored and all related flows are not flooded on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Default</b>                  | By default, there is no limit to the number of multicast groups joining an interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <i>limit</i> —a 32-bit number for the limit on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring IGMP Snooping</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## host-only-interface

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | host-only-interface;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface interface-name</a> ],<br>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> <a href="#">interface interface-name</a> ],<br>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface interface-name</a> ],<br>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols vlan <i>vlan-id</i> igmp-snooping <a href="#">interface interface-name</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Configure an interface as a host-facing interface. IGMP queries received on these interfaces are dropped.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Default</b>                  | The interface can either be a host-side or multicast-router interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring IGMP Snooping</i></li> <li>• <i>multicast-router-interface</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

## igmp-querier

|                                 |                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | igmp-querier <a href="#">source-address</a> <i>source address</i> ;                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit protocols igmp-snooping vlan <i>vlan-name</i> ]                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1X53-D15 for QFabric Systems.                                                                                                                                                                                                                                                        |
| <b>Description</b>              | Configure a QFabric Node device to be an IGMP querier. If there are any multicast routers on the same local network, make sure the source address for the IGMP querier is lower (a smaller number) than the IP addresses for those routers on the network. This ensures that Node is always the IGMP querier on the network.     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring IGMP Snooping on page 5436</a></li> <li>• <a href="#">Configuring IGMP Snooping on page 5434</a></li> <li>• <a href="#">show igmp-snooping vlans on page 5823</a></li> <li>• <a href="#">show configuration protocols igmp on page 5799</a></li> </ul> |

## igmp-snooping

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> igmp-snooping {   vlan <i>vlan-id</i> {     all     immediate-leave;     interface <i>interface-name</i> {       group-limit <i>limit</i>;       host-only-interface;       immediate-leave;       multicast-router-interface;       static {         group <i>ip-address</i> {           source <i>ip-address</i>;         }       }     }   }   l2-querier {     source-address <i>ip-address</i>;   }   proxy {     source-address <i>ip-address</i>;   }   query-interval <i>seconds</i>;   query-last-member-interval <i>seconds</i>;   query-response-interval <i>seconds</i>;   robust-count <i>number</i>;   traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;no-stamp&gt; &lt;replace&gt; &lt;size <i>size</i>&gt; &lt;world-readable         no-world-readable&gt;;     flag <i>flag</i> &lt;flag-modifier&gt;;   } } </pre> |
| <b>Hierarchy Level</b>          | [edit protocols]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Enable IGMP snooping on the router or switch.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Default</b>                  | IGMP snooping is disabled on the router or switch.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | The statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Understanding IGMP Snooping</i></li> <li>• <i>IGMP Snooping in MC-LAG Active-Active Mode</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## immediate-leave (Bridge Domains)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>immediate-leave;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>     | <pre>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping], [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <b>interface</b> <i>interface-name</i>], [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> <b>interface</b> <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <b>interface</b> <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols vlan <i>vlan-id</i> igmp-snooping <b>interface</b> <i>interface-name</i>]</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b> | Statement introduced in Junos OS Release 8.5.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>         | <p>The immediate leave setting is useful for minimizing the leave latency of IGMP memberships. When this setting is enabled, the routing device leaves the multicast group immediately after the last host leaves the multicast group.</p> <p>The immediate-leave setting enables host tracking, meaning that the device keeps track of the hosts that send join messages. This allows IGMP to determine when the last host sends a leave message for the multicast group.</p> <p>When the immediate leave setting is enabled, the device removes an interface from the forwarding-table entry without first sending IGMP group-specific queries to the interface. The interface is pruned from the multicast tree for the multicast group specified in the IGMP leave message. The immediate leave setting ensures optimal bandwidth management for hosts on a switched network, even when multiple multicast groups are being used simultaneously.</p> <p>When immediate leave is disabled and one host sends a leave group message, the routing device first sends a group query to determine if another receiver responds. If no receiver responds, the routing device removes all hosts on the interface from the multicast group. Immediate leave is disabled by default for both IGMP version 2 and IGMP version 3.</p> |



**NOTE:** Although host tracking is enabled for IGMPv2 and MLDv1 when you enable immediate leave, use immediate leave with these versions only when there is one host on the interface. The reason is that IGMPv2 and MLDv1 use a report suppression mechanism whereby only one host on an interface sends a group join report in response to a membership query. The other interested hosts suppress their reports. The purpose of this mechanism is to avoid a flood of reports for the same group. But it also interferes with host tracking, because the router only knows about the one interested host and does not know about the others.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation** • *Example: Configuring IGMP Snooping*

---

## interface (Bridge Domains)

---

**Syntax** `interface interface-name {  
 group-limit limit;  
 host-only-interface;  
 multicast-router-interface;  
 static {  
 group ip-address {  
 source ip-address;  
 }  
 }  
}`

**Hierarchy Level** [edit bridge-domains *bridge-domain-name* protocols igmp-snooping],  
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id*],  
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols  
igmp-snooping],  
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols  
vlan *vlan-id* igmp-snooping]

**Release Information** Statement introduced in Junos OS Release 8.5.

**Description** Enable IGMP snooping on an interface and configure interface-specific properties.

**Options** *interface-name*—Name of the interface. Specify the full interface name, including the physical and logical address components. To configure all interfaces, you can specify **all**.

The remaining statements are explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation** • *Example: Configuring IGMP Snooping*



## interface (IGMP Snooping)

|                                 |                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>interface <i>interface-name</i> {     multicast-router-interface;     static {         group <i>ip-address</i>;     } }</pre>                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit protocols igmp-snooping vlan <i>vlan-name</i> ]                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                      |
| <b>Description</b>              | <p>Enable IGMP snooping on an interface and configure interface-specific properties.</p> <p>The remaining statements are explained separately.</p>                                                                                                     |
| <b>Options</b>                  | <i>interface-name</i> —Name of the interface.                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring IGMP Snooping on page 5436</a></li> <li>• <a href="#">Configuring IGMP Snooping on page 5434</a></li> <li>• <a href="#">show igmp-snooping vlans on page 5823</a></li> </ul> |

## l2-querier

|                                 |                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>l2-querier {     source-address <i>ip-address</i>; }</pre>                                                                                       |
| <b>Hierarchy Level</b>          | [edit protocols igmp-snooping vlan],                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2 for the QFX Series.                                                                                     |
| <b>Description</b>              | Configure the switch to be an IGMP querier. Use the <b>source-address</b> statement to configure the source address to use for IGMP snooping queries. |
| <b>Options</b>                  | <p><b>seconds</b>—Time interval.</p> <p><b>Range:</b> 1 through 1024</p> <p><b>Default:</b> 125 seconds</p>                                           |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                        |
| <b>Related Documentation</b>    |                                                                                                                                                       |

## mld-snooping

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>mld-snooping {<br/>  vlan (<i>vlan-name</i>) {<br/>    immediate-leave;<br/>    interface (all   <i>interface-name</i>) {<br/>      group-limit <i>limit</i>;<br/>      <a href="#">host-only-interface</a>;<br/>      immediate-leave;<br/>      multicast-router-interface;<br/>      static {<br/>        group <i>ip-address</i> {<br/>          source <i>ip-address</i>;<br/>        }<br/>      }<br/>    }<br/>  }<br/>  qualified-vlan <i>vlan-id</i>;<br/>  query-interval <i>seconds</i>;<br/>  query-last-member-interval <i>seconds</i>;<br/>  query-response-interval <i>seconds</i>;<br/>  robust-count <i>number</i>;<br/>  traceoptions {<br/>    file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;<br/>    flag <i>flag</i> &lt;<i>flag-modifier</i>&gt;;<br/>  }<br/>}</pre> |
| <b>Hierarchy Level</b>          | [edit protocols]<br>[edit routing-instances <i>instance-name</i> protocols]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.3 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Enable and configure MLD snooping.<br><br>The remaining statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring MLD Snooping on page 5481</a></li><li>• <a href="#">Configuring MLD Snooping on a VLAN (CLI Procedure) on page 5475</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

---

## multicast-router-interface (IGMP Snooping)

---

|                                 |                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | multicast-router-interface;                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit protocols igmp-snooping vlan <i>vlan-name</i> <b>interface</b> <i>interface-name</i> ]                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                  |
| <b>Description</b>              | Configure an interface to forward IGMP messages to multicast routers.                                                                                                                                                                              |
| <b>Default</b>                  | Disabled. If this statement is disabled, the interface drops IGMP messages it receives.                                                                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show igmp-snooping vlans on page 5823</a></li><li>• <a href="#">Example: Configuring IGMP Snooping on page 5436</a></li><li>• <a href="#">Configuring IGMP Snooping on page 5434</a></li></ul> |

## query-interval (Bridge Domains)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>query-interval seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | <code>[edit bridge-domains <i>bridge-domain-name</i> protocols mld-snooping ] ,</code><br><code>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface</a></code><br><code><i>interface-name</i>],</code><br><code>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> <a href="#">interface</a></code><br><code><i>interface-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols</code><br><code>igmp-snooping <a href="#">interface</a> <i>interface-name</i>],</code><br><code>edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols</code><br><code>mld-snooping ]</code><br><code>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols</code><br><code>igmp-snooping vlan <i>vlan-id</i> <a href="#">interface</a> <i>interface-name</i>]</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols mld-snooping ]</code><br><code>[edit protocols igmp-snooping vlan]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 8.5.<br>Statement introduced in Junos OS Release 13.2 for the QFX series.<br>Statement introduced in Junos OS Release 14.2 for MX series Routers with MPC.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Configure the interval for host-query message timeouts.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <i>seconds</i> —Time interval. This value must be greater than the interval set for<br>query-response-interval.<br><b>Range:</b> 1 through 1024<br><b>Default:</b> 125 seconds                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring IGMP Snooping</i></li><li>• <a href="#">query-last-member-interval (Bridge Domains) on page 5654</a></li><li>• <a href="#">query-response-interval (Bridge Domains) on page 5655</a></li><li>• <i>mld-snooping</i></li><li>• <i>igmp-snooping</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

## query-last-member-interval (Bridge Domains)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>query-last-member-interval seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <pre>[edit bridge-domains <i>bridge-domain-name</i> protocols mld-snooping ] , [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface</a> <i>interface-name</i>], [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> <a href="#">interface</a> <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface</a> <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols mld-snooping ] [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> [edit routing-instances <i>routing-instance-name</i> protocols mld-snooping ]<a href="#">interface</a> <i>interface-name</i>] [edit protocols igmp-snooping vlan],</pre> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX series.</p> <p>Statement introduced in Junos OS Release 14.2 for MX series Routers with MPC.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Configure the interval for group-specific query timeouts.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <p><b>seconds</b>—Time interval, in fractions of a second or seconds.</p> <p><b>Range:</b> 0.1 through 0.9, then in 1-second intervals 1 through 1024</p> <p><b>Default:</b> 1 second</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring IGMP Snooping</i></li> <li>• <a href="#">query-interval on page 5653</a></li> <li>• <a href="#">query-response-interval on page 5655</a></li> <li>• <i>mld-snooping</i></li> <li>• <i>igmp-snooping</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## query-response-interval (Bridge Domains)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | query-response-interval <i>seconds</i> ;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface interface-name</a> ],<br>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snoopingvlan <i>vlan-id</i> <a href="#">interface interface-name</a> ],<br>[edit bridge-domains <i>bridge-domain-name</i> protocols mld-snooping ] ,<br>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface interface-name</a> ],<br>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snoopingvlan <i>vlan-id</i> ],<br>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols mld-snooping ] <a href="#">interface interface-name</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols mld-snooping ]<br>[edit protocols igmp-snooping vlan ], |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.<br>Statement introduced in Junos OS Release 13.2 for the QFX series.<br>Statement introduced in Junos OS Release 14.2 for MX series Routers with MPC.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Specify how long to wait to receive a response to a specific query message from a host.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <i>seconds</i> —Time interval. This interval should be less than the host-query interval.<br><b>Range:</b> 1 through 1024<br><b>Default:</b> 10 seconds                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring IGMP Snooping</i></li><li>• <a href="#">query-interval (Bridge Domains) on page 5653</a></li><li>• <a href="#">query-last-member-interval (Bridge Domains) on page 5654</a></li><li>• <i>mld-snooping</i></li><li>• <i>igmp-snooping</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## receiver

---

|                                 |                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>receiver {   source-vlans vlan-list;   install; }</pre>                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit protocols igmp-snooping vlan (all   <i>vlan-name</i> ) data-forwarding]                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6 for EX Series switches.<br>Statement introduced in Junos OS Release 12.3 for the QFX Series.                                             |
| <b>Description</b>              | Configure a VLAN as a multicast receiver VLAN of the multicast VLAN (MVLAN).<br><br>The remaining statements are explained separately.                                                |
| <b>Default</b>                  | Disabled                                                                                                                                                                              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring Multicast VLAN Registration</i></li> <li>• <i>Configuring Multicast VLAN Registration (CLI Procedure)</i></li> </ul> |

## robust-count (IGMP Snooping)

---

|                                 |                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | robust-count <i>number</i> ;                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit protocols igmp-snooping vlan <i>vlan-name</i> ]                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                      |
| <b>Description</b>              | Configure the number of intervals the switch waits before removing a multicast group from the multicast forwarding table. Configure the length of each interval using the <b>query-interval</b> statement.                                             |
| <b>Default</b>                  | 2 intervals                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <i>number</i> —Number of intervals the switch waits before timing out a multicast group.<br><b>Range:</b> 2 through 10                                                                                                                                 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring IGMP Snooping on page 5436</a></li> <li>• <a href="#">Configuring IGMP Snooping on page 5434</a></li> <li>• <a href="#">show igmp-snooping vlans on page 5823</a></li> </ul> |

## source-address

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>source-address <i>ip-address</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping proxy],<br>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> proxy],<br>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols<br>igmp-snooping proxy],<br>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols<br>igmp-snooping vlan <i>vlan-id</i> proxy] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.<br>Statement introduced in Junos OS Release 13.2 for the QFX series.                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Specify the IP address to use as the source for IGMP snooping reports in proxy mode. Reports are sent with address 0.0.0.0 as the source address unless there is a source address configured. You can also use this statement to configure the source address to use for IGMP snooping queries.                                                                                                                                                                                 |
| <b>Options</b>                  | <i>ip-address</i> —IP address to use as the source for proxy-mode IGMP snooping reports.                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring IGMP Snooping</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                     |



## src-address (IGMP Querier)

|                                 |                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>src-address source address;</code>                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit protocols igmp-snooping vlan <i>vlan-name</i> <a href="#">igmp-querier</a> ]<br>[edit protocols igmp-snooping vlan <i>vlan-name</i> <a href="#">l2-querier</a> ]                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.3 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D15 for QFabric Systems.                                                                                                                                                                                                                              |
| <b>Description</b>              | Configure the address that the switch uses as the source address in the IGMP queries that it sends. If there are any multicast routers on the same local network, make sure the source address for the IGMP querier is smaller (a lower number) than the IP addresses for those routers on the network. This ensures that switch is always the IGMP querier on the network. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring IGMP Snooping on page 5436</a></li> <li>• <a href="#">Configuring IGMP Snooping on page 5434</a></li> <li>• <a href="#">show igmp-snooping vlans on page 5823</a></li> <li>• <a href="#">show configuration protocols igmp on page 5799</a></li> </ul>                                            |

## source-vlans

|                                 |                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>source-vlans vlan-list;</code>                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit protocols igmp-snooping vlan (all   <i>vlan-name</i> ) data-forwarding receiver]                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6 for EX Series switches.<br>Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                               |
| <b>Description</b>              | Specify a list of multicast VLANs (MVLANS) from which this multicast receiver VLAN receives multicast traffic. Either all of these MVLANS must be in proxy mode or none of them can be in proxy mode.   |
| <b>Default</b>                  | Disabled                                                                                                                                                                                                |
| <b>Options</b>                  | <i>vlan-list</i> —Names of the MVLANS.                                                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Multicast VLAN Registration</a></li> <li>• <a href="#">Configuring Multicast VLAN Registration (CLI Procedure)</a></li> </ul> |

## static (IGMP Snooping)

---

|                                 |                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>static {<br/>    group ip-address;<br/>}</pre>                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit protocols igmp-snooping vlan <i>vlan-name</i> interface <i>interface-name</i> ]                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                  |
| <b>Description</b>              | <p>Statically define multicast groups on an interface.</p> <p>The remaining statement is explained separately.</p>                                                                                                                                 |
| <b>Default</b>                  | No multicast groups are statically defined.                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring IGMP Snooping on page 5436</a></li><li>• <a href="#">Configuring IGMP Snooping on page 5434</a></li><li>• <a href="#">show igmp-snooping vlans on page 5823</a></li></ul> |

## traceoptions (IGMP Snooping)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre> traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;no-stamp&gt; &lt;size <i>size</i>&gt; &lt;replace&gt; &lt;world-readable       no-world-readable&gt;;     flag <i>flag</i> (detail   disable   receive   send); } </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>     | <p>For platforms without ELS:</p> <pre>[edit protocols igmp-snooping]</pre> <p>For platforms with ELS:</p> <pre>[edit protocols igmp-snooping vlan]</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b> | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>         | Define tracing operations for IGMP snooping.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Default</b>             | The <b>traceoptions</b> feature is disabled by default.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>             | <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>/var/log</b>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached (<b>xk</b> to specify KB, <b>xm</b> to specify MB, or <b>xg</b> to specify gigabytes), at which point the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the <b>size</b> option.</p> <p><b>Range:</b> 2 through 1000</p> <p><b>Default:</b> 3 files</p> <p><b>flag <i>flag</i></b>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:</p> <ul style="list-style-type: none"> <li>• <b>all</b>—All tracing operations.</li> <li>• <b>general</b>—Trace general IGMP snooping protocol events.</li> <li>• <b>krt</b>—Trace communication over routing sockets.</li> <li>• <b>nexthop</b>— Trace next-hop related events.</li> <li>• <b>normal</b>—Trace normal IGMP snooping protocol events.</li> <li>• <b>packets</b>—Trace all IGMP packets.</li> <li>• <b>policy</b>—Trace policy processing.</li> <li>• <b>query</b>—Trace IGMP membership query messages.</li> <li>• <b>report</b>—Trace membership report messages.</li> </ul> |

- **route**—Trace routing information.
- **state**—Trace IGMP state transitions.
- **task**—Trace routing protocol task processing.
- **timer**—Trace routing protocol timer processing.
- **vlan**—Trace VLAN related events.

**no-stamp**—(Optional) Do not time stamp trace file.

**no-world-readable**—(Optional) Restrict file access to the user who created the file.

**size size** —(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **files** option. Use **xk** to specify KB, **xm** to specify MB, or **xg** to specify gigabytes.

**Range:** 10 KB through 1 gigabytes

**Default:** 128 KB

**world-readable**—(Optional) Enable unrestricted file access.

|                                 |                                                                   |
|---------------------------------|-------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.              |
|                                 | routing-control—To add this statement to the configuration.       |
| <b>Related Documentation</b>    | • <a href="#">Example: Configuring IGMP Snooping on page 5436</a> |
|                                 | • <a href="#">Configuring IGMP Snooping on page 5434</a>          |

## version (IGMP Snooping)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>version <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit protocols igmp-snooping vlan (all   <i>vlan-name</i> )]                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Specify the IGMP version for the IGMP general query that the switch sends to hosts when an interface comes up. The configured IGMP version affects only the version of the general queries sent by a switch. It does not affect the version of IGMP messages that the switch can snoop. For example, If the switch is configured for IGMP version 1 (IGMPv1), it can snoop IGMPv2 and IGMPv3 messages. |
| <b>Default</b>                  | If you do not configure the <b>version</b> statement, the default is IGMPv2.                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <b>version</b> —IGMP version number.<br><b>Range:</b> 1 and 2.                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring IGMP Snooping (CLI Procedure)</a></li> <li>• <a href="#">Configuring IGMP Snooping on page 5434</a></li> </ul>                                                                                                                                                                                                                        |

## vlan (IGMP Snooping)

---

**Syntax**    `vlan vlan-name {  
                  immediate-leave;  
                  interface interface-name {  
                    group-limit limit;  
                    host-only-interface;  
                    multicast-router-interface;  
                    static {  
                      group multicast-group-address {  
                        source ip-address;  
                      }  
                    }  
                  }  
                  qualified-vlan ;  
                  proxy {  
                    source-address ip-address;  
                  }  
                  query-interval seconds;  
                  query-last-member-interval seconds;  
                  query-response-interval seconds;  
                  robust-count number;  
                  }`

**Hierarchy Level**    [edit protocols igmp-snooping ],

**Release Information**    Statement introduced in Junos OS Release 8.5.  
                              Statement introduced in Junos OS Release 13.2 for the QFX series.

**Description**    Configure IGMP snooping parameters for a particular VLAN.

**Default**    By default, IGMP snooping options apply to all VLANs.

**Options**    *vlan-name*—Apply the parameters to this VLAN.  
  
              The remaining statements are explained separately.

**Required Privilege Level**    routing—To view this statement in the configuration.  
                                  routing-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring VLAN-Specific IGMP Snooping Parameters on page 5435](#)
- [igmp-snooping on page 5648](#)

## Configuration Statements (PIM)

- [address \(Anycast RPs\) on page 5689](#)
- [address \(Local RPs\) on page 5690](#)
- [address \(Static RPs\) on page 5691](#)
- [algorithm on page 5692](#)
- [anycast-pim on page 5693](#)
- [assert-timeout on page 5694](#)
- [authentication on page 5695](#)
- [auto-rp on page 5696](#)
- [bfd-liveness-detection on page 5697](#)
- [bootstrap on page 5698](#)
- [bootstrap-export on page 5699](#)
- [bootstrap-import on page 5700](#)
- [bootstrap-priority on page 5701](#)
- [dense-groups on page 5702](#)
- [detection-time \(BFD for PIM\) on page 5703](#)
- [disable \(PIM\) on page 5704](#)
- [dr-election-on-p2p on page 5705](#)
- [dr-register-policy on page 5705](#)
- [embedded-rp on page 5706](#)
- [export \(Bootstrap\) on page 5707](#)
- [export \(Protocols PIM\) on page 5708](#)
- [family \(Bootstrap\) on page 5709](#)
- [family \(Protocols PIM\) on page 5710](#)
- [family \(Local RP\) on page 5711](#)
- [group \(RPF Selection\) on page 5712](#)
- [group-ranges on page 5713](#)
- [hello-interval on page 5714](#)
- [hold-time \(Protocols PIM\) on page 5715](#)

- [import \(Protocols PIM Bootstrap\) on page 5716](#)
- [import \(Protocols PIM\) on page 5717](#)
- [infinity on page 5718](#)
- [interface on page 5719](#)
- [join-load-balance on page 5720](#)
- [join-prune-timeout on page 5721](#)
- [key-chain on page 5721](#)
- [local on page 5722](#)
- [local-address \(Protocols PIM\) on page 5723](#)
- [loose-check on page 5724](#)
- [mapping-agent-election on page 5725](#)
- [maximum-rps on page 5726](#)
- [minimum-interval \(PIM BFD Liveness Detection\) on page 5727](#)
- [minimum-interval \(PIM BFD Transmit Interval\) on page 5728](#)
- [minimum-receive-interval on page 5729](#)
- [mode \(Protocols PIM\) on page 5729](#)
- [multiplier on page 5730](#)
- [neighbor-policy on page 5730](#)
- [next-hop \(PIM RPF Selection\) on page 5731](#)
- [no-adaptation \(PIM BFD Liveness Detection\) on page 5731](#)
- [override-interval on page 5732](#)
- [pim on page 5733](#)
- [prefix-list \(PIM RPF Selection\) on page 5736](#)
- [priority \(Bootstrap\) on page 5737](#)
- [priority \(PIM Interfaces\) on page 5738](#)
- [priority \(PIM RPs\) on page 5739](#)
- [propagation-delay on page 5740](#)
- [register-probe-time on page 5741](#)
- [reset-tracking-bit on page 5742](#)
- [rib-group \(Protocols PIM\) on page 5743](#)
- [rp on page 5744](#)
- [rp-register-policy on page 5746](#)
- [rp-set on page 5747](#)
- [rpf-selection on page 5748](#)
- [source \(PIM RPF Selection\) on page 5749](#)
- [spt-threshold on page 5750](#)
- [static \(Protocols PIM\) on page 5751](#)



- [threshold \(PIM BFD Detection Time\)](#) on page 5752
- [threshold \(PIM BFD Transmit Interval\)](#) on page 5753
- [transmit-interval \(PIM BFD Liveness Detection\)](#) on page 5754
- [traceoptions \(Protocols PIM\)](#) on page 5755
- [version \(BFD\)](#) on page 5758
- [version \(PIM\)](#) on page 5759
- [wildcard-source \(PIM RPF Selection\)](#) on page 5760

## address (Anycast RPs)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>address <i>address</i> &lt;forward-msdp-sa&gt;;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols pim <b>rp local</b> (inet   inet6) <b>anycast-pim rp-set</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <b>rp local</b> (inet   inet6) <b>anycast-pim rp-set</b>],</p> <p>[edit protocols pim <b>rp local</b> (inet   inet6) <b>anycast-pim rp-set</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp local</b> (inet   inet6) <b>anycast-pim rp-set</b>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Configure the anycast rendezvous point (RP) addresses in the RP set. Multiple addresses can be configured in an RP set. If the RP has peer Multicast Source Discovery Protocol (MSDP) connections, then the RP must forward MSDP source active (SA) messages.                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <p><b><i>address</i></b>—RP address in an RP set.</p> <p><b><i>forward-msdp-sa</i></b>—(Optional) Forward MSDP SAs to this address.</p>                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                             |

## address (Local RPs)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>address <i>address</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim <b>rp local family</b> (inet   inet6)],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>pim <b>rp local family</b> (inet   inet6)],<br>[edit protocols pim <b>rp local family</b> (inet   inet6)],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp local family</b> (inet   inet6)] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                         |
| <b>Description</b>              | Configure the local rendezvous point (RP) address.                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <b><i>address</i></b> —Local RP address.                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Local PIM RPs on page 5545</a></li></ul>                                                                                                                                                                                                                                                                                                                                       |

## address (Static RPs)

|                                 |                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>address address {   group-ranges {     destination-ip-prefix&lt;/prefix-length&gt;;   }   override;   version version; }</pre>                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols pim <b>rp static</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <b>rp static</b>],</p> <p>[edit protocols pim <b>static</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp static</b>]</p>       |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| <b>Description</b>              | <p>Configure static rendezvous point (RP) addresses. You can configure a static RP in a logical system only if the logical system is not directly connected to a source.</p> <p>For each static RP address, you can optionally specify the PIM version and the groups for which this address can be the RP. The default PIM version is version 1.</p>                         |
| <b>Options</b>                  | <p><b>address</b>—Static RP address.</p> <p><b>Default:</b> 224.0.0.0/4</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring the Static PIM RP Address on the Non-RP Routing Device on page 5547</a></li> </ul>                                                                                                                                                                                                                           |

## algorithm

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>algorithm <i>algorithm-name</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit protocols pim interface <i>interface-name</i> bfd-liveness-detection authentication],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection authentication]                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Specify the algorithm to use for BFD authentication.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <p><i>algorithm-name</i>—Name of algorithm to use for BFD authentication:</p> <ul style="list-style-type: none"><li>• <b>simple-password</b>—Plain-text password. One to 16 bytes of plain text. One or more passwords can be configured.</li><li>• <b>keyed-md5</b>—Keyed Message Digest 5 hash algorithm for sessions with transmit and receive rates greater than 100 ms.</li><li>• <b>meticulous-keyed-md5</b>—Meticulous keyed Message Digest 5 hash algorithm.</li><li>• <b>keyed-sha-1</b>—Keyed Secure Hash Algorithm I for sessions with transmit and receive rates greater than 100 ms.</li><li>• <b>meticulous-keyed-sha-1</b>—Meticulous keyed Secure Hash Algorithm I.</li></ul> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Bidirectional Forwarding Detection Authentication for PIM</a></li><li>• <a href="#">Configuring BFD Authentication for PIM on page 5502</a></li><li>• <a href="#">authentication (Protocols PIM) on page 5695</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                 |

## anycast-pim

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>anycast-pim {   rp-set {     address address &lt;forward-msdp-sa&gt;;   } }</pre>                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim <b>rp local family</b> (inet   inet6)],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>pim <b>rp local family</b> (inet   inet6)],<br>[edit protocols pim <b>rp local family</b> (inet   inet6)],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp local family</b> (inet   inet6)] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Configure properties for anycast RP using PIM.<br><br>The remaining statements are explained separately.                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring PIM Anycast With or Without MSDP on page 5552</a></li> </ul>                                                                                                                                                                                                                                                                                                         |

## assert-timeout

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>assert-timeout <i>seconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim],<br>[edit protocols pim],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim]                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Multicast routing devices running PIM sparse mode often forward the same stream of multicast packets onto the same LAN through the rendezvous-point tree (RPT) and shortest-path tree (SPT). PIM assert messages help routing devices determine which routing device forwards the traffic and prunes the RPT for this group. By default, routing devices enter an assert cycle every 180 seconds. You can configure this assert timeout to be between 5 and 210 seconds. |
| <b>Options</b>                  | <b><i>seconds</i></b> —Time for routing device to wait before another assert message cycle.<br><b>Range:</b> 5 through 210 seconds<br><b>Default:</b> 180 seconds                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring the PIM Assert Timeout on page 5588</a></li></ul>                                                                                                                                                                                                                                                                                                                                               |

## authentication

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>authentication {   algorithm <i>algorithm-name</i>;   key-chain <i>key-chain-name</i>;   loose-check; }</pre>                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit protocols pim interface <i>interface-name</i> bfd-liveness-detection],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection]                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Configure the algorithm, security keychain, and level of authentication for BFD sessions running on PIM interfaces.                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | The statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring BFD Authentication for PIM on page 5502</a></li> <li>• <a href="#">Configuring BFD for PIM on page 5500</a></li> <li>• <a href="#">Understanding Bidirectional Forwarding Detection Authentication for PIM</a></li> <li>• <a href="#">bfd-liveness-detection (Protocols PIM) on page 5697</a></li> <li>• <a href="#">key-chain (Protocols PIM) on page 5721</a></li> <li>• <a href="#">loose-check on page 5724</a></li> </ul> |

## auto-rp

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>auto-rp {<br/>    (announce   discovery   mapping);<br/>    (mapping-agent-election   no-mapping-agent-election);<br/>}</pre>                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | <pre>[edit logical-systems <i>logical-system-name</i> protocols pim <i>rp</i>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br/>    pim <i>rp</i>],<br/>[edit protocols pim <i>rp</i>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols pim <i>rp</i>]</pre>                                                                                                            |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Configure automatic RP announcement and discovery.                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <p><b>announce</b>—Configure the routing device to listen only for mapping packets and also to advertise itself if it is an RP.</p> <p><b>discovery</b>—Configure the routing device to listen only for mapping packets.</p> <p><b>mapping</b>—Configures the routing device to announce, listen for and generate mapping packets, and announce that the routing device is eligible to be an RP.</p> <p>The remaining statement is explained separately.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring PIM Auto-RP on page 5561</a></li></ul>                                                                                                                                                                                                                                                                                                                                                       |



## bfd-liveness-detection

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> bfd-liveness-detection {   authentication {     algorithm <i>algorithm-name</i>;     key-chain <i>key-chain-name</i>;     loose-check;   }   detection-time {     threshold <i>milliseconds</i>;   }   minimum-interval <i>milliseconds</i>;   minimum-receive-interval <i>milliseconds</i>;   multiplier <i>number</i>;   no-adaptation;   transmit-interval {     minimum-interval <i>milliseconds</i>;     threshold <i>milliseconds</i>;   }   version (0   1   automatic); } </pre> |
| <b>Hierarchy Level</b>          | <p>[edit protocols pim interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i>]</p>                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.1.</p> <p><b>authentication</b> option introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                              |
| <b>Description</b>              | <p>Configure bidirectional forwarding detection (BFD) timers and authentication for PIM.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring BFD for PIM on page 5500</a></li> <li>• <a href="#">Configuring BFD Authentication for PIM on page 5502</a></li> </ul>                                                                                                                                                                                                                                                                                                        |

## bootstrap

---

|                                 |                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>bootstrap {<br/>    family (inet   inet6) {<br/>        export [ <i>policy-names</i> ];<br/>        import [ <i>policy-names</i> ];<br/>        priority <i>number</i>;<br/>    }<br/>}</pre>                                                                                                                                              |
| <b>Hierarchy Level</b>          | <pre>[edit logical-systems <i>logical-system-name</i> protocols pim <i>rp</i>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br/>  pim <i>rp</i>],<br/>[edit protocols pim <i>rp</i>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols pim <i>rp</i>]</pre> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                       |
| <b>Description</b>              | <p>Configure parameters to control bootstrap routers and messages.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring PIM Bootstrap Properties for IPv4</i></li><li>• <i>Configuring PIM Bootstrap Properties for IPv4 or IPv6</i></li></ul>                                                                                                                                                                   |

## bootstrap-export

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>bootstrap-export [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols pim <a href="#">rp</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <a href="#">rp</a>],</p> <p>[edit protocols pim <a href="#">rp</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <a href="#">rp</a>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                      |
| <b>Description</b>              | Apply one or more export policies to control outgoing PIM bootstrap messages.                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more import policies.                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring PIM Bootstrap Properties for IPv4</i></li> <li>• <i>Configuring PIM Bootstrap Properties for IPv4 or IPv6</i></li> <li>• <a href="#">bootstrap-import on page 5700</a></li> </ul>                                                                                                                                          |

## bootstrap-import

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>bootstrap-import [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols pim <i>rp</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>pim <i>rp</i>],</code><br><code>[edit protocols pim <i>rp</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols pim <i>rp</i>]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                  |
| <b>Description</b>              | Apply one or more import policies to control incoming PIM bootstrap messages.                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more import policies.                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <code>routing</code> —To view this statement in the configuration.<br><code>routing-control</code> —To add this statement to the configuration.                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring PIM Bootstrap Properties for IPv4</i></li><li>• <i>Configuring PIM Bootstrap Properties for IPv4 or IPv6</i></li><li>• <a href="#">bootstrap-export on page 5699</a></li></ul>                                                                                                                                                           |

## bootstrap-priority

---

|                                 |                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>bootstrap-priority <i>number</i>;</code>                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols pim <i>rp</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <i>rp</i>],</p> <p>[edit protocols pim <i>rp</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <i>rp</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                  |
| <b>Description</b>              | Configure whether this routing device is eligible to be a bootstrap router. In the case of a tie, the routing device with the highest IP address is elected to be the bootstrap router.                                                                                                                                                        |
| <b>Options</b>                  | <p><i>number</i>—Priority for becoming the bootstrap router. A value of 0 means that the routing device is not eligible to be the bootstrap router.</p> <p><b>Range:</b> 0 through 255</p> <p><b>Default:</b> 0</p>                                                                                                                            |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring PIM Bootstrap Properties for IPv4</i></li> </ul>                                                                                                                                                                                                                                       |

## dense-groups

---

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>dense-groups {<br/>    addresses;<br/>}</pre>                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim],<br>[edit protocols pim],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                      |
| <b>Description</b>              | Configure which groups are operating in dense mode.                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <b>addresses</b> —Address of groups operating in dense mode.                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring PIM Sparse-Dense Mode Properties on page 5527</a></li></ul>                                                                                                                                                         |

## detection-time (BFD for PIM)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> detection-time {     threshold milliseconds; } </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit protocols pim interface <i>interface-name</i> bfd-liveness-detection],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.2.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for BFD authentication introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | <p>Enable BFD failure detection. The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the <b>clear bfd adaptation</b> command to return BFD interval timers to their configured values. The <b>clear bfd adaptation</b> command is hitless, meaning that the command does not affect traffic flow on the routing device.</p> <p>The remaining statement is explained separately.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring BFD for PIM on page 5500</a></li> <li>• <a href="#">bfd-liveness-detection on page 5697</a></li> <li>• <a href="#">threshold on page 5752</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

## disable (PIM)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | disable;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <pre>[edit logical-systems <i>logical-system-name</i> protocols pim], [edit logical-systems <i>logical-system-name</i> protocols pim <b>family</b> (inet   inet6)], [edit logical-systems <i>logical-system-name</i> protocols pim interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols pim <b>rp local family</b> (inet   inet6)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   pim], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   pim interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   pim <b>rp local family</b> (inet   inet6)], [edit protocols pim], [edit protocols pim <b>family</b> (inet   inet6)], [edit protocols pim interface <i>interface-name</i>], [edit protocols pim <b>rp local family</b> (inet   inet6)], [edit routing-instances <i>routing-instance-name</i> protocols pim], [edit routing-instances <i>routing-instance-name</i> protocols pim <b>family</b> (inet   inet6)], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp local family</b> (inet   inet6)]</pre> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p><b>disable</b> statement extended to the <b>[family]</b> hierarchy level in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Explicitly disable PIM at the protocol, interface or family hierarchy levels.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Disabling PIM on page 5505</a></li><li>• <a href="#">family (Protocols PIM) on page 5710</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |



## dr-election-on-p2p

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | dr-election-on-p2p;                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim],<br>[edit protocols pim],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.1.<br>Statement introduced in Junos OS Release 9.1 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                          |
| <b>Description</b>              | Enable PIM designated router (DR) election on point-to-point (P2P) links.                                                                                                                                                                                                           |
| <b>Default</b>                  | No PIM DR election is performed on point-to-point links.                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring PIM Designated Router Election on Point-to-Point Links on page 5499</a></li> </ul>                                                                                                                                 |

## dr-register-policy

|                                 |                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | dr-register-policy [ <i>policy-names</i> ];                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim <i>rp</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <i>rp</i> ],<br>[edit protocols pim <i>rp</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim <i>rp</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.6.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                      |
| <b>Description</b>              | Apply one or more policies to control outgoing PIM register messages.                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more import policies.                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Register Message Filters on a PIM RP and DR on page 5577</a></li> <li>• <a href="#">rp-register-policy on page 5746</a></li> </ul>                                                                                                                             |

## embedded-rp

---

|                          |                                                                                                                                                                                                                                                                                                                                       |
|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <pre>embedded-rp {<br/>  group-ranges {<br/>    destination-ip-prefix &lt;/prefix-length&gt;;<br/>  }<br/>  maximum-rps limit;<br/>}</pre>                                                                                                                                                                                            |
| Hierarchy Level          | <p>[edit logical-systems <i>logical-system-name</i> protocols pim <b>rp</b>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <b>rp</b>],<br/>[edit protocols pim <b>rp</b>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp</b>]</p> |
| Release Information      | <p>Statement introduced before Junos OS Release 7.4.<br/>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br/>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                               |
| Description              | <p>Configure properties for embedded IP version 6 (IPv6) RPs.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                           |
| Required Privilege Level | <p>routing—To view this statement in the configuration.<br/>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                           |
| Related Documentation    | <ul style="list-style-type: none"><li>• <i>Configuring PIM Embedded RP for IPv6</i></li></ul>                                                                                                                                                                                                                                         |

## export (Bootstrap)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>export [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols pim <a href="#">rp bootstrap family</a> (inet   inet6)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <a href="#">rp bootstrap family</a> (inet   inet6)],</p> <p>[edit protocols pim <a href="#">rp bootstrap family</a> (inet   inet6)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <a href="#">rp bootstrap family</a> (inet   inet6)]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                          |
| <b>Description</b>              | Apply one or more export policies to control outgoing PIM bootstrap messages.                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more import policies.                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring PIM Bootstrap Properties for IPv4</i></li> <li>• <i>Configuring PIM Bootstrap Properties for IPv4 or IPv6</i></li> <li>• <a href="#">import (Protocols PIM Bootstrap) on page 5716</a></li> </ul>                                                                                                                                                                                                                                                          |

## export (Protocols PIM)

---

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | export [ <i>policy-names</i> ];                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim],<br>[edit protocols pim],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                      |
| <b>Description</b>              | Apply one or more export policies to control outgoing PIM join and prune messages. PIM join and prune filters can be applied to PIM-SM and PIM-SSM messages. PIM join and prune filters cannot be applied to PIM-DM messages.                                                       |
| <b>Required Privilege Level</b> | view-level—To view this statement in the configuration.<br>control-level—To add this statement to the configuration.                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Filtering Outgoing PIM Join Messages on page 5574</a></li></ul>                                                                                                                                                                 |

## family (Bootstrap)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>family (inet   inet6) {     export [ <i>policy-names</i> ];     import [ <i>policy-names</i> ];     priority <i>number</i>; }</pre>                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols pim <b>rp bootstrap</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <b>rp bootstrap</b>],</p> <p>[edit protocols pim <b>rp bootstrap</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp bootstrap</b>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                              |
| <b>Description</b>              | Configure which IP protocol type bootstrap properties to apply.                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <p><b>inet</b>—Apply IP version 4 (IPv4) local RP properties.</p> <p><b>inet6</b>—Apply IPv6 local RP properties.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>Configuring PIM Bootstrap Properties for IPv4</li> <li>Configuring PIM Bootstrap Properties for IPv4 or IPv6</li> </ul>                                                                                                                                                                                                                         |

## family (Protocols PIM)

---

|                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                | family (inet   inet6) {<br>disable;<br>}                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>       | [edit logical-systems <i>logical-system-name</i> protocols pim],<br>[edit logical-systems <i>logical-system-name</i> protocols pim interface <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> ],<br>[edit protocols pim],<br>[edit protocols pim interface <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> ] |
| <b>Release Information</b>   | Statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>           | Disable the PIM protocol for the specified family.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>               | <b>inet</b> —Disable the PIM protocol for the IP version 4 (IPv4) address family.<br><br><b>inet6</b> —Disable the PIM protocol for the IP version 6 (IPv6) address family.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Disabling PIM on page 5505</a></li><li>• <i>disable (PIM Graceful Restart)</i></li><li>• <a href="#">disable (PIM) on page 5704</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## family (Local RP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> family (inet   inet6) {     disable;     address address;     anycast-pim {         local-address address;         rp-set {             address address &lt;forward-msdp-sa&gt;;         }     }     group-ranges {         destination-ip-prefix &lt;/prefix-length&gt;;     }     hold-time seconds;     override;     priority number; } </pre>               |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols pim <b>rp local</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <b>rp local</b>],</p> <p>[edit protocols pim <b>rp local</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp local</b>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                          |
| <b>Description</b>              | Configure which IP protocol type local RP properties to apply.                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <p><b>inet</b>—Apply IP version 4 (IPv4) local RP properties.</p> <p><b>inet6</b>—Apply IPv6 local RP properties.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">Configuring Local PIM RPs on page 5545</a></li> </ul>                                                                                                                                                                                                                                                               |

## group (RPF Selection)

---

|                                 |                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>group group-address{   source source-address {     next-hop next-hop-address;   }   wildcard-source {     next-hop next-hop-address;   } }</pre> |
| <b>Hierarchy Level</b>          | [edit routing-instances <i>routing-instance-name</i> edit protocols pim rpf-selection]                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.4.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                   |
| <b>Description</b>              | Configure the PIM group address for which you configure RPF selection <a href="#">group (RPF Selection)</a> .                                         |
| <b>Default</b>                  | By default, PIM RPF selection is not configured.                                                                                                      |
| <b>Options</b>                  | <b>group-address</b> —PIM group address for which you configure RPF selection.                                                                        |
| <b>Required Privilege Level</b> | view-level—To view this statement in the configuration.<br>control-level—To add this statement to the configuration.                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring PIM RPF Selection</i></li></ul>                                                       |



## group-ranges

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>group-ranges {   destination-ip-prefix&lt;/prefix-length&gt;; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols pim <a href="#">rp embedded-rp</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <a href="#">rp embedded-rp</a>],</p> <p>[edit protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit protocols pim <a href="#">rp embedded-rp</a>],</p> <p>[edit protocols pim <a href="#">rp local family</a> (inet   inet6)],</p> <p>[edit protocols pim <a href="#">rp static address</a> <i>address</i>],</p> <p>[edit routing-instances <i>instance-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <a href="#">rp embedded-rp</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <a href="#">rp local family</a> (inet   inet6)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <a href="#">rp static address</a> <i>address</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> <p>Support for bidirectional RP addresses introduced in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 13.3 for the PTX5000 router.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Configure the address ranges of the multicast groups for which this routing device can be a rendezvous point (RP).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Default</b>                  | The routing device is eligible to be the RP for all IPv4 or IPv6 groups (224.0.0.0/4 or FF70::/12 to FFF0::/12).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <i>destination-ip-prefix&lt;/prefix-length&gt;</i> —Addresses or address ranges for which this routing device can be an RP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">Configuring Local PIM RPs on page 5545</a></li> <li><a href="#">Configuring PIM Embedded RP for IPv6</a></li> <li><a href="#">Example: Configuring Bidirectional PIM</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

## hello-interval

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>hello-interval <i>seconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim interface <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>pim interface <i>interface-name</i> ],<br>[edit protocols pim interface <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                             |
| <b>Description</b>              | Specify how often the routing device sends PIM hello packets out of an interface.                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <b><i>seconds</i></b> —Length of time between PIM hello packets.<br><b>Range:</b> 0 through 255<br><b>Default:</b> 30 seconds                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">hold-time (Protocols PIM) on page 5715</a></li><li>• <a href="#">Modifying the PIM Hello Interval on page 5494</a></li></ul>                                                                                                                                                                                                                                           |

## hold-time (Protocols PIM)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>hold-time seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit protocols pim <b>rp local family</b> (inet   inet6)],</p> <p>[edit routing-instances <i>instance-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp local family</b> (inet   inet6)]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional RP addresses introduced in Junos OS Release 12.1.</p>                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Specify the time period for which a neighbor is to consider the sending routing device (this routing device) to be operative (up).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <p><b>seconds</b>—Hold time.</p> <p><b>Range:</b> 0 through 255</p> <p><b>Default:</b> 150 seconds</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Local PIM RPs on page 5545</a></li> <li>• <i>Example: Configuring Bidirectional PIM</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

## import (Protocols PIM Bootstrap)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>import [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim <a href="#">rp bootstrap</a> (inet   inet6)],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>pim <a href="#">rp bootstrap</a> (inet   inet6)],<br>[edit protocols pim <a href="#">rp bootstrap</a> (inet   inet6)],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim <a href="#">rp bootstrap</a> (inet   inet6)] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.6.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                 |
| <b>Description</b>              | Apply one or more import policies to control incoming PIM bootstrap messages.                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more import policies.                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring PIM Bootstrap Properties for IPv4</a></li><li>• <a href="#">Configuring PIM Bootstrap Properties for IPv4 or IPv6</a></li><li>• <a href="#">export (Bootstrap) on page 5707</a></li></ul>                                                                                                                                                                                                              |

## import (Protocols PIM)

---

|                                 |                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>import [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols pim],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim],</p> <p>[edit protocols pim],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>          |
| <b>Description</b>              | Apply one or more policies to routes being imported into the routing table from PIM. Use the <b>import</b> statement to filter PIM join messages and prevent them from entering the network.                                                                                                           |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more policies.                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Filtering Incoming PIM Join Messages on page 5575</a></li> </ul>                                                                                                                                                                                  |

## infinity

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>infinity [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim <a href="#">spt-threshold</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <a href="#">spt-threshold</a> ],<br>[edit protocols pim <a href="#">spt-threshold</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim <a href="#">spt-threshold</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.0.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                      |
| <b>Description</b>              | Apply one or more policies to set the SPT threshold to infinity for a source-group address pair. Use the <b>infinity</b> statement to prevent the last-hop routing device from transitioning from the RPT rooted at the RP to an SPT rooted at the source for that source-group address pair.                                                                                                                   |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more policies.                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring the PIM SPT Threshold Policy on page 5590</a></li></ul>                                                                                                                                                                                                                                                                                |

## interface

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> <b>interface</b> (all   <i>interface-name</i>) {   <b>disable</b>;   family (inet   inet6) {     <b>disable</b>;   }   <b>hello-interval</b> <i>seconds</i>;   mode (dense   sparse   sparse-dense);   <b>neighbor-policy</b> [ <i>policy-names</i> ];   <b>override-interval</b> <i>milliseconds</i>;   <b>priority</b> <i>number</i>;   <b>propagation-delay</b> <i>milliseconds</i>;   <b>reset-tracking-bit</b>;   <b>version</b> <i>version</i>; } </pre> |
| <b>Hierarchy Level</b>          | [edit protocols pim],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim]                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | Enable PIM on an interface and configure interface-specific properties.                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <p><b><i>interface-name</i></b>—Name of the interface. Specify the full interface name, including the physical and logical address components. To configure all interfaces, you can specify <b>all</b>.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">PIM on Aggregated Interfaces on page 5494</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                        |

## join-load-balance

---

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | join-load-balance {<br>automatic;<br>}                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim],<br>[edit protocols pim],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                          |
| <b>Description</b>              | Enable load balancing of PIM join messages across interfaces and routing devices.                                                                                                                                                                                                   |
| <b>Options</b>                  | <b>automatic</b> —Enables automatic load balancing of PIM join messages. When a new interface or neighbor is introduced into the network, ECMP joins are redistributed with minimal disruption to traffic.                                                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring PIM Join Load Balancing on page 5513</a></li><li>• <code>clear pim join-distribution</code> in the <a href="#">CLI Explorer</a></li></ul>                                                                           |



## join-prune-timeout

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | join-prune-timeout <i>seconds</i> ;                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim],<br>[edit protocols pim],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.4.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                  |
| <b>Description</b>              | Configure the timeout for the join state. If the periodic join refresh message is not received before the timeout expires, the join state is removed.                                                                                                                               |
| <b>Options</b>                  | <b>seconds</b> —Number of seconds to wait for the periodic join message to arrive.<br><b>Range:</b> 210 through 240 seconds<br><b>Default:</b> 210 seconds                                                                                                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Modifying the Join State Timeout on page 5517</a></li> </ul>                                                                                                                                                                   |

## key-chain

|                                 |                                                                                                                                                                                                                                                                                                   |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | key-chain <i>key-chain-name</i> ;                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit protocols pim interface <i>interface-name</i> bfd-liveness-detection authentication],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection authentication]                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                |
| <b>Description</b>              | Specify the security keychain to use for BFD authentication.                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <b>key-chain-name</b> —Name of the security keychain to use for BFD authentication. The name is a unique integer between 0 and 63. This must match one of the keychains in the <b>authentication-key-chains</b> statement at the <b>[edit security]</b> hierarchy level.                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring BFD Authentication for PIM on page 5502</a></li> <li>• <a href="#">Understanding Bidirectional Forwarding Detection Authentication for PIM</a></li> <li>• <a href="#">authentication (Protocols PIM) on page 5695</a></li> </ul> |

## local

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> local {   disable;   address address;   family (inet   inet6) {     disable;     address address;     anycast-pim {       local-address address;       rp-set {         address address &lt;forward-msdp-sa&gt;;       }     }     group-ranges {       destination-ip-prefix&lt;/prefix-length&gt;;     }     hold-time seconds;     override;     priority number;   }   group-ranges {     destination-ip-prefix&lt;/prefix-length&gt;;   }   hold-time seconds;   override;   priority number; } </pre> |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols pim <b>rp</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <b>rp</b>],</p> <p>[edit protocols pim <b>rp</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp</b>]</p>                                                                                                                                                                    |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                           |
| <b>Description</b>              | Configure the routing device's RP properties.                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">Configuring Local PIM RPs on page 5545</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                          |

## local-address (Protocols PIM)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>local-address <i>address</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols pim <b>rp local family</b> (inet   inet6) <b>anycast-pim</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <b>rp local family</b> (inet   inet6) <b>anycast-pim</b>],</p> <p>[edit protocols pim <b>rp local family</b> (inet   inet6) <b>anycast-pim</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp local family</b> (inet   inet6) <b>anycast-pim</b>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Configure the routing device local address for the anycast rendezvous point (RP). If this statement is omitted, the router ID is used as this address.                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <b>address</b> —Anycast RP IPv4 or IPv6 address, depending on <b>family</b> configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring PIM Anycast With or Without MSDP on page 5552</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                     |

## loose-check

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | loose-check;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit protocols pim interface <i>interface-name</i> bfd-liveness-detection authentication],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection authentication]                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | <p>Specify loose authentication checking on the BFD session. Use loose authentication for transitional periods only when authentication might not be configured at both ends of the BFD session.</p> <p>By default, strict authentication is enabled and authentication is checked at both ends of each BFD session. Optionally, to smooth migration from nonauthenticated sessions to authenticated sessions, you can configure <i>loose checking</i>. When loose checking is configured, packets are accepted without authentication being checked at each end of the session.</p> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring BFD Authentication for PIM on page 5502</a></li><li>• <a href="#">Understanding Bidirectional Forwarding Detection Authentication for PIM</a></li><li>• <a href="#">authentication (Protocols PIM) on page 5695</a></li></ul>                                                                                                                                                                                                                                                                                        |

## mapping-agent-election

|                                 |                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | (mapping-agent-election   no-mapping-agent-election);                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim <b>rp auto-rp</b> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>pim <b>rp auto-rp</b> ],<br>[edit protocols pim <b>rp auto-rp</b> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp auto-rp</b> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.5.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                         |
| <b>Description</b>              | Configure the routing device mapping announcements as a mapping agent.                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <b>mapping-agent-election</b> —Mapping agents do not announce mappings when receiving mapping messages from a higher-addressed mapping agent.<br><br><b>no-mapping-agent-election</b> —Mapping agents always announce mappings and do not perform mapping agent election.<br><br><b>Default:</b> mapping-agent-election                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring PIM Auto-RP on page 5561</a></li> </ul>                                                                                                                                                                                                                                                           |

## maximum-rps

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>maximum-rps <i>limit</i>;</code>                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim <a href="#">rp embedded-rp</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <a href="#">rp embedded-rp</a> ],<br>[edit protocols pim <a href="#">rp embedded-rp</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim <a href="#">rp embedded-rp</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                      |
| <b>Description</b>              | Limit the number of RPs that the routing device acknowledges.                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <i>limit</i> —Number of RPs.<br><b>Range:</b> 1 through 500<br><b>Default:</b> 100                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring PIM Embedded RP for IPv6</i></li></ul>                                                                                                                                                                                                                                                                                                                       |

## minimum-interval (PIM BFD Liveness Detection)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>minimum-interval <i>milliseconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit protocols pim interface <i>interface-name</i> <b>bfd-liveness-detection</b> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> <b>bfd-liveness-detection</b> ]                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.1.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                 |
| <b>Description</b>              | Configure the minimum interval after which the local routing device transmits hello packets and then expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately using the <b>transmit-interval</b> <b>minimum-interval</b> and <b>minimum-receive-interval</b> statements. |
| <b>Options</b>                  | <b><i>milliseconds</i></b> —Minimum transmit and receive interval.<br><b>Range:</b> 1 through 255,000 milliseconds                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring BFD for PIM on page 5500</a></li> </ul>                                                                                                                                                                                                                                                                                                                   |

## minimum-interval (PIM BFD Transmit Interval)

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>minimum-interval <i>milliseconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>     | [edit protocols pim interface <i>interface-name</i> bfd-liveness-detection transmit-interval],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection transmit-interval]                                                                                                                                                      |
| <b>Release Information</b> | Statement introduced in Junos OS Release 8.2.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Support for BFD authentication introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                    |
| <b>Description</b>         | Configure the minimum interval after which the local routing device transmits hello packets to a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit interval using the <b>minimum-interval</b> statement at the [edit protocols pim interface <i>interface-name</i> bfd-liveness-detection] hierarchy level. |
| <b>Options</b>             | <i>milliseconds</i> —Minimum transmit interval value.<br><b>Range:</b> 1 through 255,000                                                                                                                                                                                                                                                                                                            |



**NOTE:** The threshold value specified in the **threshold** statement must be greater than the value specified in the **minimum-interval** statement for the **transmit-interval** statement.

---

|                                 |                                                                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring BFD for PIM on page 5500</a></li><li>• <a href="#">bfd-liveness-detection on page 5697</a></li><li>• <a href="#">minimum-interval on page 5727</a></li><li>• <a href="#">threshold on page 5753</a></li></ul> |



## minimum-receive-interval

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>minimum-receive-interval <i>milliseconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | <code>[edit protocols pim interface <i>interface-name</i> <b>bfd-liveness-detection</b>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i></code><br><code><b>bfd-liveness-detection</b>]</code>                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.1.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                |
| <b>Description</b>              | Configure the minimum interval after which the local routing device must receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the <b>minimum-interval</b> statement at the <code>[edit protocols pim interface <i>interface-name</i> <b>bfd-liveness-detection</b>]</code> hierarchy level. |
| <b>Options</b>                  | <b><i>milliseconds</i></b> —Minimum receive interval.<br><b>Range:</b> 1 through 255,000 milliseconds                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring BFD for PIM on page 5500</a></li> </ul>                                                                                                                                                                                                                                                                                                              |

## mode (Protocols PIM)

|                                 |                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>mode (dense   sparse   sparse-dense);</code>                                                                                                                                     |
| <b>Hierarchy Level</b>          | <code>[edit protocols pim interface <i>interface-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i>]</code> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                          |
| <b>Description</b>              | Configure PIM to operate in sparse, dense, or sparse-dense mode.                                                                                                                       |
| <b>Options</b>                  | <b>dense</b> —Operate in dense mode.<br><br><b>sparse</b> —Operate in sparse mode.<br><br><b>sparse-dense</b> —Operate in sparse-dense mode.<br><b>Default:</b> sparse                 |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                    |

## multiplier

---

|                                 |                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>multiplier <i>number</i>;</code>                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit protocols pim interface <i>interface-name</i> <b>bfd-liveness-detection</b> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> <b>bfd-liveness-detection</b> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.1.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                 |
| <b>Description</b>              | Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.                                                                                                |
| <b>Options</b>                  | <b><i>number</i></b> —Number of hello packets.<br><b>Range:</b> 1 through 255<br><b>Default:</b> 3                                                                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring BFD for PIM on page 5500</a></li></ul>                                                                                                                     |

## neighbor-policy

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>neighbor-policy [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim interface <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> ],<br>[edit protocols pim interface <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.2.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                              |
| <b>Description</b>              | Apply a PIM interface-level policy to filter neighbor IP addresses.                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <b><i>policy-name</i></b> —Name of the policy that filters neighbor IP addresses.                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Interface-Level PIM Neighbor Policies on page 5573</a></li></ul>                                                                                                                                                                                                                                                                                        |

## next-hop (PIM RPF Selection)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>next-hop <i>next-hop-address</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i> source <i>source-address</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i> wildcard-source],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list <i>prefix-list-addresses</i> source <i>source-address</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list <i>prefix-list-addresses</i> wildcard-source] |
| <b>Release Information</b>      | Statement introduced in JUNOS Release 10.4.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | Configure the specific next-hop address for the PIM group source.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <i>next-hop-address</i> —Specific next-hop address for the PIM group source.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | view-level—To view this statement in the configuration.<br>control-level—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring PIM RPF Selection</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## no-adaptation (PIM BFD Liveness Detection)

|                                 |                                                                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>no-adaptation;</code>                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit protocols pim interface <i>interface-name</i> bfd-liveness-detection],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection]                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Support for BFD authentication introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series. |
| <b>Description</b>              | Configure BFD sessions not to adapt to changing network conditions. We recommend that you <i>do not</i> disable BFD adaptation unless it is preferable to have BFD adaptation disabled in your network.                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring BFD for PIM on page 5500</a></li> <li>• <a href="#">bfd-liveness-detection on page 5697</a></li> </ul>                                                                                         |

## override-interval

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | override-interval <i>milliseconds</i> ;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim],<br>[edit logical-systems <i>logical-system-name</i> protocols pim interface <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>pim interface <i>interface-name</i> ],<br>[edit protocols pim],<br>[edit protocols pim interface <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim]<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.1.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Set the maximum time in milliseconds to delay sending override join messages for a multicast network that has join suppression enabled. When a router or switch sees a prune message for a join it is currently suppressing, it waits for the interval specified by the override timer before it sends an override join message.                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | This is a random timer with a value in milliseconds.<br><b>Range:</b> 0 through maximum override value<br><b>Default:</b> 2000 milliseconds                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Enabling Join Suppression on page 5517</a></li><li>• <a href="#">propagation-delay on page 5740</a></li><li>• <a href="#">reset-tracking-bit on page 5742</a></li></ul>                                                                                                                                                                                                                                                                                                                                                           |

## pim

```

Syntax pim {
 disable;
 assert-timeout seconds;
 dense-groups {
 addresses;
 }
 dr-election-on-p2p;
 export;
 family (inet | inet6) {
 disable;
 }
 graceful-restart {
 disable;
 restart-duration seconds;
 }
 import [policy-names];
 interface interface-name {
 accept-remote-source;
 disable;
 family (inet | inet6) {
 disable;
 }
 hello-interval seconds;
 mode (dense | sparse | sparse-dense);
 neighbor-policy [policy-names];
 override-interval milliseconds;
 priority number;
 propagation-delay milliseconds;
 reset-tracking-bit;
 version version;
 }
 join-load-balance;
 join-prune-timeout;
 nonstop-routing;
 override-interval milliseconds;
 propagation-delay milliseconds;
 reset-tracking-bit;
 rib-group group-name;
 rp {
 auto-rp {
 (announce | discovery | mapping);
 (mapping-agent-election | no-mapping-agent-election);
 }
 bootstrap {
 family (inet | inet6) {
 export [policy-names];
 import [policy-names];
 priority number;
 }
 }
 bootstrap-import [policy-names];
 bootstrap-export [policy-names];
 }
}

```

```
bootstrap-priority number;
dr-register-policy [policy-names];
embedded-rp {
 group-ranges {
 destination-ip-prefix </prefix-length>;
 }
 maximum-rps limit;
}
local {
 family (inet | inet6) {
 address address;
 anycast-pim {
 disable;
 rp-set {
 address address <forward-msdp-sa>;
 }
 local-address address;
 }
 group-ranges {
 destination-ip-prefix </prefix-length>;
 }
 hold-time seconds;
 priority number;
 }
}
rp-register-policy [policy-names];
spt-threshold {
 infinity [policy-names];
}
static {
 address address {
 group-ranges {
 version version;
 destination-ip-prefix </prefix-length>;
 }
 }
}
rpf-selection {
 group group-address {
 source source-address {
 next-hop next-hop-address;
 }
 wildcard-source {
 next-hop next-hop-address;
 }
 }
 prefix-list prefix-list-addresses {
 source source-address {
 next-hop next-hop-address;
 }
 wildcard-source {
 next-hop next-hop-address;
 }
 }
}
traceoptions {
```

```

 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
 }
 tunnel-devices [mt-fpc/pic/port];
}

```

|                                 |                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols],<br>[edit protocols],<br>[edit routing-instances <i>routing-instance-name</i> protocols]                                                                       |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br><b>family</b> statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| <b>Description</b>              | Enable PIM on the routing device.<br><br>The statements are explained separately.                                                                                                                                                                                                                                                         |
| <b>Default</b>                  | PIM is disabled on the routing device.                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                       |

## prefix-list (PIM RPF Selection)

---

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <pre>prefix-list <i>prefix-list-addresses</i> {<br/>    source <i>source-address</i> {<br/>        next-hop <i>next-hop-address</i>;<br/>    }<br/>    wildcard-source {<br/>        next-hop <i>next-hop-address</i>;<br/>    }<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                        |
| Hierarchy Level          | <pre>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i><br/>    source <i>source-address</i>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i><br/>    wildcard-source],<br/>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list<br/>    <i>prefix-list-addresses</i> source <i>source-address</i>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list<br/>    <i>prefix-list-addresses</i> wildcard-source]</pre> |
| Release Information      | Statement introduced in Junos OS Release 10.4.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Description              | (Optional) Configure a list of prefixes (addresses) for multiple PIM groups.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Options                  | <b><i>prefix-list-addresses</i></b> —List of prefixes (addresses) for multiple PIM groups.<br><br>The remaining statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Required Privilege Level | view-level—To view this statement in the configuration.<br>control-level—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Related Documentation    | <ul style="list-style-type: none"><li>• <i>Example: Configuring PIM RPF Selection</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |



## priority (Bootstrap)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>priority <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols pim <b>rp bootstrap</b> (inet   inet6)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <b>rp bootstrap</b> (inet   inet6)],</p> <p>[edit protocols pim <b>rp bootstrap</b> (inet   inet6)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp bootstrap</b> (inet   inet6)]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                          |
| <b>Description</b>              | Configure the routing device's likelihood to be elected as the bootstrap router.                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <p><b>number</b>—Routing device's priority for becoming the bootstrap router. A higher value corresponds to a higher priority.</p> <p><b>Range:</b> 0 through a 32-bit number</p> <p><b>Default:</b> 0 (The routing device has the least likelihood of becoming the bootstrap router and sends packets with a priority of 0.)</p>                                                                                                                  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring PIM Bootstrap Properties for IPv4</i></li> <li>• <i>Configuring PIM Bootstrap Properties for IPv4 or IPv6</i></li> <li>• <a href="#">bootstrap-priority on page 5701</a></li> </ul>                                                                                                                                                                                                        |

## priority (PIM Interfaces)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>priority <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim interface <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>pim interface <i>interface-name</i> ],<br>[edit protocols pim interface <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                 |
| <b>Description</b>              | Configure the routing device's likelihood to be elected as the designated router.                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <b><i>number</i></b> —Routing device's priority for becoming the designated router. A higher value corresponds to a higher priority.<br><b>Range:</b> 0 through 4294967295<br><b>Default:</b> 1 (Each routing device has an equal probability of becoming the DR.)                                                                                                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Interface Priority for PIM Designated Router Selection on page 5498</a></li></ul>                                                                                                                                                                                                                                                                          |

## priority (PIM RPs)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>priority <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit protocols pim <b>rp local family</b> (inet   inet6)],</p> <p>[edit routing-instances <i>instance-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp local family</b> (inet   inet6)]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> <p>Support for bidirectional RP addresses introduced in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 13.3 for the PTX5000 router.</p>                                                                                                                                                                               |
| <b>Description</b>              | <p>For PIM-SM, configure this routing device's priority for becoming an RP.</p> <p>For bidirectional PIM, configure this RP address' priority for becoming an RP.</p> <p>The bootstrap router uses this field when selecting the list of candidate rendezvous points to send in the bootstrap message. A smaller number increases the likelihood that the routing device or RP address becomes the RP. A priority value of 0 means that bootstrap router can override the group range being advertised by the candidate RP.</p>                                                                                                             |
| <b>Options</b>                  | <p><b><i>number</i></b>—Priority for becoming an RP. A lower value corresponds to a higher priority.</p> <p><b>Range:</b> 0 through 255</p> <p><b>Default:</b> 1</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Local PIM RPs on page 5545</a></li> <li>• <a href="#">Example: Configuring Bidirectional PIM</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

## propagation-delay

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>propagation-delay <i>milliseconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit protocols pim],<br>[edit protocols pim interface <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> protocols pim],<br>[edit logical-systems <i>logical-system-name</i> protocols pim interface <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>pim interface <i>interface-name</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.1.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Set a delay for implementing a PIM prune message on the upstream routing device on a multicast network for which join suppression has been enabled. The routing device waits for the prune pending period to detect whether a join message is currently being suppressed by another routing device.                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <b><i>milliseconds</i></b> —Interval for the prune pending timer, which is the sum of the <b>propagation-delay</b> value and the <b>override-interval</b> value.<br><b>Range:</b> 250 through 2000 milliseconds<br><b>Default:</b> 500 milliseconds                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Enabling Join Suppression on page 5517</a></li><li>• <a href="#">override-interval on page 5732</a></li><li>• <a href="#">reset-tracking-bit on page 5742</a></li></ul>                                                                                                                                                                                                                                                                                                                                                            |

---

## register-probe-time

---

|                                 |                                                                                                                                                                  |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>register-probe-time</code> <i>register-probe-time</i> ;                                                                                                    |
| <b>Hierarchy Level</b>          | [edit protocols pim <a href="#">rp</a> ]                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for EX Series switches.<br>Statement introduced in Junos OS Release 14.1X53-D16 for QFX Series switches.           |
| <b>Description</b>              | Specify the amount of time before the register suppression time (RST) expires when a designated switch can send a NULL-Register to the rendezvous point (RP).    |
| <b>Options</b>                  | <i>register-probe-time</i> —Amount of time before the RST expires.<br><b>Default:</b> 5 seconds<br><b>Range:</b> 5 to 60 seconds                                 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">PIM Overview on page 5491</a></li><li>• <a href="#">Understanding PIM Sparse Mode on page 5509</a></li></ul> |

## reset-tracking-bit

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | reset-tracking-bit;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit protocols pim],<br>[edit protocols pim interface <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> protocols pim],<br>[edit logical-systems <i>logical-system-name</i> protocols pim interface <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> ]                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.1.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Change the value of a tracking bit (T-bit) field in the LAN prune delay hello option from the default of 1 to 0, which enables join suppression for a multicast interface. When the network starts receiving multiple identical join messages, join suppression triggers a random timer with a value of 66 through 84 milliseconds ( $1.1 \times \text{periodic}$ through $1.4 \times \text{periodic}$ , where periodic is 60 seconds). This creates an interval during which no identical join messages are sent. Eventually, only one of the identical messages is sent. Join suppression is triggered each time identical messages are sent for the same join. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Enabling Join Suppression on page 5517</a></li><li>• <a href="#">override-interval on page 5732</a></li><li>• <a href="#">propagation-delay on page 5740</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                               |

## rib-group (Protocols PIM)

---

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> rib-group {     inet <i>group-name</i>;     inet6 <i>group-name</i>; } </pre>                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim],<br>[edit protocols pim],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.          |
| <b>Description</b>              | Associate a routing table group with PIM.                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <i>table-name</i> —Name of the routing table. The name must be one that you defined with the <b>rib-groups</b> statement at the [edit routing-options] hierarchy level.                                                                                                             |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>Example: Configuring a Dedicated PIM RPF Routing Table</li> </ul>                                                                                                                                                                            |

## rp

```

Syntax register-probe-time {
 auto-rp {
 (announce | discovery | mapping);
 (mapping-agent-election | no-mapping-agent-election);
 }
 bidirectional {
 address address {
 group-ranges {
 destination-ip-prefix</prefix-length>;
 }
 hold-time seconds;
 priority number;
 }
 }
 bootstrap {
 family (inet | inet6) {
 export [policy-names];
 import [policy-names];
 priority number;
 }
 }
 bootstrap-export [policy-names];
 bootstrap-import [policy-names];
 bootstrap-priority number;
 dr-register-policy [policy-names];
 embedded-rp {
 group-ranges {
 destination-ip-prefix</prefix-length>;
 }
 maximum-rps limit;
 }
 group-rp-mapping {
 family (inet | inet6) {
 log-interval seconds;
 maximum limit;
 threshold value;
 }
 }
 log-interval seconds;
 maximum limit;
 threshold value;
 }
 local {
 family (inet | inet6) {
 disable;
 address address;
 anycast-pim {
 local-address address;
 address address <forward-msdp-sa>;
 rp-set {
 }
 }
 }

```



```

 }
 group-ranges {
 destination-ip-prefix</prefix-length>;
 }
 hold-time seconds;
 override;
 priority number;
}
}
register-limit {
 family (inet | inet6) {
 log-interval seconds;
 maximum limit;
 threshold value;
 }
}
log-interval seconds;
maximum limit;
threshold value;
}
}
register-probe-time register-probe-time;
}
rp-register-policy [policy-names];
static {
 address address {
 override;
 version version;
 group-ranges {
 destination-ip-prefix</prefix-length>;
 }
 }
}
}
}

```

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols pim],  
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols  
pim],  
[edit protocols pim],  
[edit routing-instances *routing-instance-name* protocols pim]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 11.3 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Configure the routing device as an actual or potential RP. A routing device can be an RP  
for more than one group.

The remaining statements are explained separately.

**Default** If you do not include the **rp** statement, the routing device can never become the RP.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Understanding PIM Sparse Mode on page 5509](#)

---

## rp-register-policy

---

**Syntax** `rp-register-policy [ policy-names ];`

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols pim *rp*],  
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim *rp*],  
[edit protocols pim *rp*],  
[edit routing-instances *routing-instance-name* protocols pim *rp*]

**Release Information** Statement introduced in Junos OS Release 7.6.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 11.3 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Apply one or more policies to control incoming PIM register messages.

**Options** *policy-names*—Name of one or more import policies.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring Register Message Filters on a PIM RP and DR on page 5577](#)
- [dr-register-policy on page 5705](#)

## rp-set

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>rp-set {   address address &lt;forward-msdp-sa&gt;; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols pim <b>local family</b> (inet   inet6) <b>anycast-pim</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <b>local family</b> (inet   inet6) <b>anycast-pim</b>],</p> <p>[edit protocols pim <b>local family</b> (inet   inet6) <b>anycast-pim</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>local family</b> (inet   inet6) <b>anycast-pim</b>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                      |
| <b>Description</b>              | <p>Configure a set of rendezvous point (RP) addresses for anycast RP. You can configure up to 15 RPs.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring PIM Anycast With or Without MSDP on page 5552</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                         |

## rpf-selection

---

**Syntax**    `rpf-selection {  
              group group-address {  
                  source source-address {  
                      next-hop next-hop-address;  
                  }  
                  wildcard-source {  
                      next-hop next-hop-address;  
                  }  
              }  
              prefix-list prefix-list-addresses {  
                  source source-address {  
                      next-hop next-hop-address;  
                  }  
                  wildcard-source {  
                      next-hop next-hop-address;  
                  }  
              }  
          }`

**Hierarchy Level**    [edit routing-instances *routing-instance-name* protocols pim]

**Release Information**    Statement introduced in JUNOS Release 10.4.  
Statement introduced in Junos OS Release 11.3 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description**    Configure the PIM RPF next-hop neighbor for a specific group and source for a VRF routing instance.

The remaining statements are explained separately.

**Default**    If you omit the **rpf-selection** statement, PIM RPF checks typically choose the best path determined by the unicast protocol for all multicast flows.

**Options**    *source-address*—Specific source address for the PIM group.

**Required Privilege Level**    view-level—To view this statement in the configuration.  
control-level—To add this statement to the configuration.

**Related Documentation**    • *Example: Configuring PIM RPF Selection*

## source (PIM RPF Selection)

---

|                                 |                                                                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>source source-address {     next-hop next-hop-address; }</pre>                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list <i>prefix-list-addresses</i> ] |
| <b>Release Information</b>      | Statement introduced in JUNOS Release 10.4.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                  |
| <b>Description</b>              | Configure the source address for the PIM group.                                                                                                                                                                                               |
| <b>Options</b>                  | <p><b>source-address</b>—Specific source address for the PIM group.</p> <p>The remaining statements are explained separately.</p>                                                                                                             |
| <b>Required Privilege Level</b> | view-level—To view this statement in the configuration.<br>control-level—To add this statement to the configuration.                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring PIM RPF Selection</i></li> </ul>                                                                                                                                             |

## spt-threshold

---


|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>spt-threshold {<br/>    infinity [ <i>policy-names</i> ];<br/>}</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim],<br>[edit protocols pim],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim]                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.0.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | <p>Set the SPT threshold to infinity for a source-group address pair. Last-hop multicast routing devices running PIM sparse mode can forward the same stream of multicast packets onto the same LAN through an RPT rooted at the RP or an SPT rooted at the source. By default, last-hop routing devices transition to a direct SPT to the source. You can configure this routing device to set the SPT transition value to infinity to prevent this transition for any source-group address pair.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring the PIM SPT Threshold Policy on page 5590</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## static (Protocols PIM)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>static {   address address {     group-ranges {       destination-ip-prefix&lt;/prefix-length&gt;;     }     override;     version version;   } }</pre>                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols pim <b>rp</b>],<br/> [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <b>rp</b>],<br/> [edit protocols pim <b>rp</b>],<br/> [edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp</b>]</p>                                                                                                                                                                                |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.<br/> Statement introduced in Junos OS Release 9.0 for EX Series switches.<br/> Statement introduced in Junos OS Release 11.3 for the QFX Series.<br/> Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                 |
| <b>Description</b>              | <p>Configure static RP addresses. The default static RP address is 224.0.0.0/4. To configure other addresses, include one or more <b>address</b> statements. You can configure a static RP in a logical system only if the logical system is not directly connected to a source.</p> <p>For each static RP address, you can optionally specify the PIM version and the groups for which this address can be the RP. The default PIM version is version 1.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.<br/> routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring the Static PIM RP Address on the Non-RP Routing Device on page 5547</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                     |


## threshold (PIM BFD Detection Time)

---

|                                                                                                                                                                                                                                                                                                                                                                             |                                                                                                                                                                                                                                                                                                                                                    |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                               | <code>threshold <i>milliseconds</i>;</code>                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                                      | [edit protocols pim interface <i>interface-name</i> bfd-liveness-detection detection-time],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection detection-time]                                                                                                           |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                  | Statement introduced in Junos OS Release 8.2.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Support for BFD authentication introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.       |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                          | Specify the threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.                                                                                                                              |
| <div> <b>NOTE:</b> The threshold value must be equal to or greater than the transmit interval.</div> <div>The threshold time must be equal to or greater than the value specified in the <a href="#">minimum-interval</a> or the <a href="#">minimum-receive-interval</a> statement.</div> |                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                                                              | <i>milliseconds</i> —Value for the detection time adaptation threshold.<br><b>Range:</b> 1 through 255,000                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                             | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                                                | <ul style="list-style-type: none"><li>• <a href="#">Configuring BFD for PIM on page 5500</a></li><li>• <a href="#">bfd-liveness-detection on page 5697</a></li><li>• <a href="#">detection-time on page 5703</a></li><li>• <a href="#">minimum-interval on page 5727</a></li><li>• <a href="#">minimum-receive-interval on page 5729</a></li></ul> |



## threshold (PIM BFD Transmit Interval)

|                                                                                                                                                                                                                                                                                                                                     |                                                                                                                                                                                                                                                                        |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                       | <code>threshold <i>milliseconds</i>;</code>                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                              | [edit protocols pim interface <i>interface-name</i> bfd-liveness-detection transmit-interval],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection transmit-interval]                         |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                          | Statement introduced in Junos OS Release 8.2.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                  | Specify the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent.                                                            |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                      | <i>milliseconds</i> —Value for the transmit interval adaptation threshold.<br><b>Range:</b> 0 through 4,294,967,295 ( $2^{32} - 1$ )                                                                                                                                   |
| <div>  <p><b>NOTE:</b> The threshold value specified in the <code>threshold</code> statement must be greater than the value specified in the <code>minimum-interval</code> statement for the <code>transmit-interval</code> statement.</p> </div> |                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                     | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                    |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                        | <ul style="list-style-type: none"> <li>• <a href="#">Configuring BFD for PIM on page 5500</a></li> <li>• <a href="#">bfd-liveness-detection on page 5697</a></li> </ul>                                                                                                |

## transmit-interval (PIM BFD Liveness Detection)

---

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <pre>transmit-interval {<br/>    <i>minimum-interval milliseconds</i>;<br/>    <i>threshold milliseconds</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Hierarchy Level          | [edit protocols pim interface <i>interface-name</i> bfd-liveness-detection],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection]                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Release Information      | Statement introduced in Junos OS Release 8.2.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Support for BFD authentication introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                     |
| Description              | <p>Specify the transmit interval for the <b>bfd-liveness-detection</b> statement. The negotiated transmit interval for a peer is the interval between the sending of BFD packets to peers. The receive interval for a peer is the minimum interval between receiving packets sent from its peer; the receive interval is not negotiated between peers. To determine the transmit interval, each peer compares its configured minimum transmit interval with its peer's minimum receive interval. The larger of the two numbers is accepted as the transmit interval for that peer.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Related Documentation    | <ul style="list-style-type: none"><li>• <a href="#">Configuring BFD for PIM on page 5500</a></li><li>• <a href="#">bfd-liveness-detection on page 5697</a></li><li>• <a href="#">threshold on page 5753</a></li><li>• <a href="#">minimum-interval on page 5728</a></li><li>• <a href="#">minimum-receive-interval on page 5729</a></li></ul>                                                                                                                                                                                                                                                                                                    |

## traceoptions (Protocols PIM)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;flag-modifier&gt; &lt;disable&gt;; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols pim],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim],</p> <p>[edit protocols pim],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>         | <p>Configure PIM tracing options.</p> <p>To specify more than one tracing operation, include multiple <b>flag</b> statements.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Default</b>             | The default PIM trace options are those inherited from the routing protocol's <b>traceoptions</b> statement included at the [edit routing-options] hierarchy level.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>             | <p><b>disable</b>—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>/var/log</b>. We recommend that you place tracing output in the <b>pim-log</b> file.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you must also include the <b>size</b> statement to specify the maximum file size.</p> <p><b>Range:</b> 2 through 1000 files</p> <p><b>Default:</b> 2 files</p> <p><b>flag <i>flag</i></b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements.</p> <p><b>PIM Tracing Flags</b></p> <ul style="list-style-type: none"> <li><b>assert</b>—Assert messages</li> <li><b>bidirectional-df-election</b>—Bidirectional PIM designated-forwarder (DF) election events</li> </ul> |

- **bootstrap**—Bootstrap messages
- **cache**—Packets in the PIM sparse mode routing cache
- **graft**—Graft and graft acknowledgment messages
- **hello**—Hello packets
- **join**—Join messages
- **mt**—Multicast tunnel messages
- **nsr-synchronization**—Nonstop active routing (NSR) synchronization messages
- **packets**—All PIM packets
- **prune**—Prune messages
- **register**—Register and register stop messages
- **rp**—Candidate RP advertisements
- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

**Default:** If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

**flag-modifier**—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

**no-stamp**—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

**Default:** If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

**no-world-readable**—(Optional) Do not allow users to read the log file.

**replace**—(Optional) Replace an existing trace file if there is one.

**Default:** If you do not include this option, tracing output is appended to an existing trace file.

**size** *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When **trace-file** again reaches this size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of trace files.

**Syntax:** *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

**Range:** 0 KB through the maximum file size supported on your system

**Default:** 1 MB

**world-readable**—(Optional) Allow any user to read the log file.

|                                 |                                                                               |
|---------------------------------|-------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing and trace—To view this statement in the configuration.                |
|                                 | routing-control and trace-control—To add this statement to the configuration. |

|                              |                                                                                                                                                                                                                                                                                                                      |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"> <li>• <a href="#">Configuring PIM Trace Options on page 5496</a></li> <li>• <a href="#">Tracing DVMRP Protocol Traffic</a></li> <li>• <a href="#">Tracing MSDP Protocol Traffic on page 5600</a></li> <li>• <a href="#">Configuring PIM Trace Options on page 5496</a></li> </ul> |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## version (BFD)

---

|                                 |                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | version (0   1   automatic);                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit protocols piminterface <i>interface-name</i> <a href="#">bfd-liveness-detection</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> <a href="#">bfd-liveness-detection</a> ]                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.1.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| <b>Description</b>              | Specify the bidirectional forwarding detection (BFD) protocol version that you want to detect.                                                                                                                                                                         |
| <b>Options</b>                  | Configure the BFD version to detect: <b>1</b> (BFD version 1) or <b>automatic</b> (autodetect the BFD version)<br><b>Default:</b> automatic                                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring BFD for PIM on page 5500</a></li></ul>                                                                                                                                                                 |

## version (PIM)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>version <i>version</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols pim interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols pim <b>rp static address</b> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <b>rp static address</b> <i>address</i>],</p> <p>[edit protocols pim interface <i>interface-name</i>],</p> <p>[edit protocols pim <b>rp static address</b> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp static address</b> <i>address</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Specify the version of PIM.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <p><b>version</b>—PIM version number.</p> <p><b>Range:</b> 1 or 2</p> <p><b>Default:</b> PIMv1 for rendezvous point (RP) mode (at the [edit protocols pim rp static address <i>address</i>] hierarchy level). PIMv2 for interface mode (at the [edit protocols pim interface <i>interface-name</i>] hierarchy level).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Enabling PIM Sparse Mode on page 5512</a></li> <li>• <a href="#">Configuring PIM Dense Mode Properties on page 5526</a></li> <li>• <a href="#">Configuring PIM Sparse-Dense Mode Properties on page 5527</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

## wildcard-source (PIM RPF Selection)

---

|                                 |                                                                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | wildcard-source {<br>next-hop next-hop-address;<br>}                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list <i>prefix-list-addresses</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.4.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                               |
| <b>Description</b>              | Use a wildcard for the multicast source instead of (or in addition to) a specific multicast source.<br><br>The remaining statements are explained separately.                                                                                 |
| <b>Required Privilege Level</b> | view-level—To view this statement in the configuration.<br>control-level—To add this statement to the configuration.                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring PIM RPF Selection</i></li></ul>                                                                                                                                               |



# Configuration Statements (Source-Specific Multicast)

- [asm-override-ssm](#) on page 5761
- [policy \(SSM Maps\)](#) on page 5762
- [ssm-groups](#) on page 5763
- [ssm-map \(Protocols IGMP\)](#) on page 5764
- [ssm-map \(Routing Options Multicast\)](#) on page 5765
- [ssm-map-policy \(IGMP\)](#) on page 5766

## [asm-override-ssm](#)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>asm-override-ssm;</code>                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],<br>[edit logical-systems <i>logical-system-name</i> routing-options multicast],<br>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],<br>[edit routing-options multicast]                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4.<br>Statement introduced in Junos OS Release 9.5 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| <b>Description</b>              | Enable the routing device to accept any-source multicast join messages (*G) for group addresses that are within the default or configured range of source-specific multicast groups.                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Source-Specific Multicast Groups with Any-Source Override</a> on page 5537</li> </ul>                                                                                                                                                                                                                                                             |

## policy (SSM Maps)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>policy [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast <a href="#">ssm-map</a> <i>ssm-map-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-options multicast <a href="#">ssm-map</a> <i>ssm-map-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> routing-options multicast <a href="#">ssm-map</a> <i>ssm-map-name</i>],</code><br><code>[edit routing-options multicast <a href="#">ssm-map</a> <i>ssm-map-name</i>]</code> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                         |
| <b>Description</b>              | Apply one or more policies to an SSM map.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more policies for SSM mapping.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | <code>routing</code> —To view this statement in the configuration.<br><code>routing-control</code> —To view this statement in the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring SSM Mapping on page 5535</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

## ssm-groups

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>ssm-groups [ <i>ip-addresses</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit routing-options multicast]</p>                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | <p>Configure source-specific multicast (SSM) groups.</p> <p>By default, the SSM group multicast address is limited to the IP address range from 232.0.0.0 through 232.255.255.255. However, you can extend SSM operations into another Class D range by including the <b>ssm-groups</b> statement in the configuration. The default SSM address range from 232.0.0.0 through 232.255.255.255 cannot be used in the <b>ssm-groups</b> statement. This statement is for adding other multicast addresses to the default SSM group addresses. This statement does not override the default SSM group address range.</p> <p>IGMPv3 supports SSM groups. By utilizing inclusion lists, only sources that are specified send to the SSM group.</p> |
| <b>Options</b>                  | <i>ip-addresses</i> —List of one or more additional SSM group addresses separated by a space.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Source-Specific Multicast Groups with Any-Source Override on page 5537</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## ssm-map (Protocols IGMP)

---

|                                 |                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>ssm-map <i>ssm-map-name</i>;</code>                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> interface <i>interface-name</i> ],<br>[edit protocols <b>igmp</b> interface <i>interface-name</i> ]                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| <b>Description</b>              | Apply an SSM map to an IGMP interface.                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <i>ssm-map-name</i> —Name of SSM map.                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring SSM Mapping on page 5535</a></li></ul>                                                                                                                                                        |

## ssm-map (Routing Options Multicast)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>ssm-map <i>ssm-map-name</i> {     policy [ <i>policy-names</i> ];     source [ <i>addresses</i> ]; }</pre>                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit routing-options multicast]</p>                 |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| <b>Description</b>              | Configure SSM mapping.                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><b><i>ssm-map-name</i></b>—Name of the SSM map.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring SSM Mapping on page 5535</a></li> </ul>                                                                                                                                                                                                                                                      |

## ssm-map-policy (IGMP)

---

|                                 |                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>ssm-map-policy <i>ssm-map-policy-name</i>;</code>                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">igmp interface interface-name</a> ],<br>[edit protocols <a href="#">igmp interface interface-name</a> ]                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.4.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| <b>Description</b>              | Apply an SSM map policy to an IGMP interface.                                                                                                                                                   |
| <b>Options</b>                  | <i>ssm-map-policy-name</i> —Name of SSM map policy.                                                                                                                                             |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring SSM Maps for Different Groups to Different Sources on page 554</a></li></ul>                                           |

## Configuration Statements (MSDP)

- [active-source-limit](#) on page 5768
- [authentication-key](#) on page 5769
- [data-encapsulation](#) on page 5770
- [default-peer](#) on page 5771
- [disable \(Protocols MSDP\)](#) on page 5772
- [export \(Protocols MSDP\)](#) on page 5773
- [group \(Protocols MSDP\)](#) on page 5774
- [import \(Protocols MSDP\)](#) on page 5775
- [local-address \(Protocols MSDP\)](#) on page 5776
- [maximum \(MSDP Active Source Messages\)](#) on page 5777
- [mode \(Protocols MSDP\)](#) on page 5778
- [msdp](#) on page 5779
- [peer \(Protocols MSDP\)](#) on page 5781
- [rib-group \(Protocols MSDP\)](#) on page 5782
- [source \(Protocols MSDP\)](#) on page 5783
- [threshold \(MSDP Active Source Messages\)](#) on page 5784
- [traceoptions \(Protocols MSDP\)](#) on page 5785

## active-source-limit

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>active-source-limit {<br/>    log-interval <i>seconds</i>;<br/>    log-warning <i>value</i>;<br/>    maximum <i>number</i>;<br/>    threshold <i>number</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | <pre>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp</b>],<br/>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp</b> group <i>group-name</i> <b>peer</b> <i>address</i>],<br/>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp</b> <b>peer</b> <i>address</i>],<br/>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp</b> source <i>ip-address/prefix-length</i>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols <b>msdp</b>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br/>    <b>msdp</b> group <i>group-name</i> <b>peer</b> <i>address</i>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br/>    <b>msdp</b> <b>peer</b> <i>address</i>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br/>    <b>msdp</b> source <i>ip-address/prefix-length</i>],<br/>[edit protocols <b>msdp</b>],<br/>[edit protocols <b>msdp</b> group <i>group-name</i> <b>peer</b> <i>address</i>],<br/>[edit protocols <b>msdp</b> <b>peer</b> <i>address</i>],<br/>[edit protocols <b>msdp</b> source <i>ip-address/prefix-length</i>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp</b>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp</b> group <i>group-name</i><br/>    <b>peer</b> <i>address</i>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp</b> <b>peer</b> <i>address</i>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp</b> source<br/>    <i>ip-address/prefix-length</i>]</pre> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Limit the number of active source messages the routing device accepts.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Default</b>                  | If you do not include this statement, the router accepts any number of MSDP active source messages.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | The options are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 5604</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |



## authentication-key

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>authentication-key peer-key;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <code>msdp group group-name peer address</code>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <code>msdp peer address</code>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <code>msdp group group-name peer address</code>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <code>msdp peer address</code>],</p> <p>[edit protocols <code>msdp group group-name peer address</code>],</p> <p>[edit protocols <code>msdp peer address</code>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <code>msdp group group-name peer address</code>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <code>msdp peer address</code>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Associate a Message Digest 5 (MD5) signature option authentication key with an MSDP peering session.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Default</b>                  | If you do not include this statement, the routing device accepts any valid MSDP messages from the peer address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <p><b>peer-key</b>—MD5 authentication key. The peer key can be a text string up to 16 letters and digits long. Strings can include any ASCII characters with the exception of ( , ) , &amp; , and [ . If you include spaces in an MSDP authentication key, enclose all characters in quotation marks ( " " ).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring MSDP in a Routing Instance</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## data-encapsulation

---

|                                 |                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | data-encapsulation (disable   enable);                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp</a> ],<br>[edit protocols <a href="#">msdp</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                          |
| <b>Description</b>              | Configure a rendezvous point (RP) using MSDP to encapsulate multicast data received in MSDP register messages inside forwarded MSDP source-active messages.                                                                                                                                                                                                 |
| <b>Default</b>                  | If you do not include this statement, the RP encapsulates multicast data.                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <b>disable</b> —(Optional) Do not use MSDP data encapsulation.<br><b>enable</b> —Use MSDP data encapsulation.<br><b>Default:</b> enable                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 5604</a></li></ul>                                                                                                                                                                                                          |

## default-peer

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | default-peer;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp peer</b> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp peer</b> <i>address</i>],</p> <p>[edit protocols <b>msdp</b>],</p> <p>[edit protocols <b>msdp group</b> <i>group-name</i>],</p> <p>[edit protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],</p> <p>[edit protocols <b>msdp peer</b> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp peer</b> <i>address</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Establish this peer as the default MSDP peer and accept source-active messages from the peer without the usual peer-reverse-path-forwarding (peer-RPF) check.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 5604</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

## disable (Protocols MSDP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | disable;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | <pre>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp</b>], [edit logical-systems <i>logical-system-name</i> protocols <b>msdp group</b> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>], [edit logical-systems <i>logical-system-name</i> protocols <b>msdp peer</b> <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp</b>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp peer</b> <i>address</i>], [edit protocols <b>msdp</b>], [edit protocols <b>msdp group</b> <i>group-name</i>], [edit protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>], [edit protocols <b>msdp peer</b> <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols <b>msdp</b>], [edit routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols <b>msdp peer</b> <i>address</i>]</pre> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Explicitly disable MSDP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>Disabling MSDP</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

## export (Protocols MSDP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>export [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp group</a> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp group</a> <i>group-name</i> <a href="#">peer</a> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp</a> <a href="#">peer</a> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp group</a> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp group</a> <i>group-name</i> <a href="#">peer</a> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp</a> <a href="#">peer</a> <i>address</i>],</p> <p>[edit protocols <a href="#">msdp</a>],</p> <p>[edit protocols <a href="#">msdp group</a> <i>group-name</i>],</p> <p>[edit protocols <a href="#">msdp group</a> <i>group-name</i> <a href="#">peer</a> <i>address</i>],</p> <p>[edit protocols <a href="#">msdp</a> <a href="#">peer</a> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp group</a> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp group</a> <i>group-name</i> <a href="#">peer</a> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp</a> <a href="#">peer</a> <i>address</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Apply one or more policies to routes being exported from the routing table into MSDP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more policies.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring MSDP in a Routing Instance</i></li> <li>• <a href="#">import on page 5775</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## group (Protocols MSDP)

```
Syntax group group-name {
 disable;
 export [policy-names];
 import [policy-names];
 local-address address;
 mode (mesh-group | standard);
 traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
 }
 peer address; {
 disable;
 active-source-limit {
 maximum number;
 threshold number;
 }
 authentication-key peer-key;
 default-peer;
 export [policy-names];
 import [policy-names];
 local-address address;
 traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
 }
 }
 }
```

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols **msdp**],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **msdp**],  
 [edit protocols **msdp**],  
 [edit routing-instances *routing-instance-name* protocols **msdp**]

**Release Information** Statement introduced before Junos OS Release 7.4.  
 Statement introduced in Junos OS Release 12.1 for the QFX Series.  
 Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Define an MSDP peer group. MSDP peers within groups share common tracing options, if present and not overridden for an individual peer with the **peer** statement. To configure multiple MSDP groups, include multiple **group** statements.

By default, the group's options are identical to the global MSDP options. To override the global options, include group-specific options within the **group** statement.

The group must contain at least one peer.

**Options** **group-name**—Name of the MSDP group.

The remaining statements are explained separately.

|                                 |                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration. |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring MSDP in a Routing Instance</i></li> </ul>          |

## import (Protocols MSDP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>import [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp group</a> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp group</a> <i>group-name</i> <a href="#">peer</a> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp peer</a> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp group</a> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp group</a> <i>group-name</i> <a href="#">peer</a> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp peer</a> <i>address</i>],</p> <p>[edit protocols <a href="#">msdp</a>],</p> <p>[edit protocols <a href="#">msdp group</a> <i>group-name</i>],</p> <p>[edit protocols <a href="#">msdp group</a> <i>group-name</i> <a href="#">peer</a> <i>address</i>],</p> <p>[edit protocols <a href="#">msdp peer</a> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp group</a> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp group</a> <i>group-name</i> <a href="#">peer</a> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp peer</a> <i>address</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Apply one or more policies to routes being imported into the routing table from MSDP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more policies.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring MSDP in a Routing Instance</i></li> <li>• <a href="#">export on page 5773</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## local-address (Protocols MSDP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>local-address address;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp</b>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp group</b> <i>group-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp peer</b> <i>address</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp</b>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp peer</b> <i>address</i>],</code><br><code>[edit protocols <b>msdp</b>],</code><br><code>[edit protocols <b>msdp group</b> <i>group-name</i>],</code><br><code>[edit protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],</code><br><code>[edit protocols <b>msdp peer</b> <i>address</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp</b>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp peer</b> <i>address</i>]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Configure the local end of an MSDP session. You must configure at least one peer for MSDP to function. When configuring a peer, you must include this statement. This address is used to accept incoming connections to the peer and to establish connections to the remote peer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <b>address</b> —IP address of the local end of the connection.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring MSDP in a Routing Instance</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |



## maximum (MSDP Active Source Messages)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>maximum <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp active-source-limit</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp active-source-limit</a>],</p> <p>[edit protocols <a href="#">msdp active-source-limit</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp active-source-limit</a>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                          |
| <b>Description</b>              | Configure the maximum number of MSDP active source messages the router accepts.                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <p><i>number</i>—Maximum number of active source messages.</p> <p><b>Range:</b> 1 through 1,000,000</p> <p><b>Default:</b> 25,000</p>                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 5604</a></li> <li>• <a href="#">threshold (MSDP Active Source Messages) on page 5784</a></li> </ul>                                                                                                                                                                                                                       |

## mode (Protocols MSDP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | mode (mesh-group   standard);                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>msdp group</b> <i>group-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i> ],<br>[edit protocols <b>msdp group</b> <i>group-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                      |
| <b>Description</b>              | Configure groups of peers in a full mesh topology to limit excessive flooding of source-active messages to neighboring peers. The default flooding mode is <b>standard</b> .                                                                                                                                                                                                                                            |
| <b>Default</b>                  | If you do not include this statement, default flooding is applied.                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <b>mesh-group</b> —Group of peers that are mesh group members.<br><br><b>standard</b> —Use standard MSDP source-active flooding rules.<br><b>Default:</b> standard                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 5604</a></li></ul>                                                                                                                                                                                                                                                                      |

## msdp

```

Syntax msdp {
 disable;
 active-source-limit {
 log-interval seconds;
 log-warning value;
 maximum number;
 threshold number;
 }
 data-encapsulation (disable | enable);
 export [policy-names];
 group group-name {
 ...group-configuration ...
 }
 hold-time seconds;
 import [policy-names];
 local-address address;
 keep-alive seconds;
 peer address {
 ...peer-configuration ...
 }
 rib-group group-name;
 source ip-prefix</prefix-length> {
 active-source-limit {
 maximum number;
 threshold number;
 }
 }
 sa-hold-time seconds;
 traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
 }
 group group-name {
 disable;
 export [policy-names];
 import [policy-names];
 local-address address;
 mode (mesh-group | standard);
 peer address {
 ... same statements as at the [edit protocols msdp peer address] hierarchy level shown
 just following ...
 }
 traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
 }
 }
 peer address {
 disable;
 active-source-limit {
 maximum number;
 threshold number;
 }
 }
 }

```

```
 }
 authentication-key peer-key;
 default-peer;
 export [policy-names];
 import [policy-names];
 local-address address;
 traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
 }
}
}
```

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols],  
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols],  
[edit protocols],  
[edit routing-instances *routing-instance-name* protocols]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.4 for EX Series switches.  
Statement introduced in Junos OS Release 12.1 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Enable MSDP on the router or switch. You must also configure at least one peer for MSDP to function.

**Default** MSDP is disabled on the router or switch.

**Options** The statements are explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Configuring MSDP in a Routing Instance*

## peer (Protocols MSDP)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre> peer address {     disable;     active-source-limit {         maximum number;         threshold number;     }     authentication-key peer-key;     default-peer;     export [ policy-names ];     import [ policy-names ];     local-address address;     traceoptions {         file filename &lt;files number&gt; &lt;size size&gt; &lt;world-readable   no-world-readable&gt;;         flag flag &lt;flag-modifier&gt; &lt;disable&gt;;     } } </pre>                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i>],</p> <p>[edit protocols <b>msdp</b>],</p> <p>[edit protocols <b>msdp group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i>]</p>                                               |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>         | <p>Define an MSDP peering relationship. An MSDP routing device must know which routing devices are its peers. You define the peer relationships explicitly by configuring the neighboring routing devices that are the MSDP peers of the local routing device. After peer relationships are established, the MSDP peers exchange messages to advertise active multicast sources. To configure multiple MSDP peers, include multiple <b>peer</b> statements.</p> <p>By default, the peer's options are identical to the global or group-level MSDP options. To override the global or group-level options, include peer-specific options within the <b>peer (Protocols MSDP)</b> statement.</p> <p>At least one peer must be configured for MSDP to function. You must configure <b>address</b> and <b>local-address</b>.</p> |
| <b>Options</b>             | <p><b>address</b>—Name of the MSDP peer.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Configuring MSDP in a Routing Instance*

---

## rib-group (Protocols MSDP)

---

**Syntax** `rib-group group-name;`

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols [msdp](#)],  
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols [msdp](#)],  
[edit protocols [msdp](#)],  
[edit routing-instances *routing-instance-name* protocols [msdp](#)]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 12.1 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Associate a routing table group with MSDP.

**Options** *group-name*—Name of the routing table group. The name must be one that you defined with the **rib-groups** statement at the [edit routing-options] hierarchy level.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Configuring MSDP in a Routing Instance*

## source (Protocols MSDP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>source ip-address &lt;/prefix-length&gt; {     active-source-limit {         maximum number;         threshold number;     } }</pre>                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp</a>],</p> <p>[edit protocols <a href="#">msdp</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp</a>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                          |
| <b>Description</b>              | Limit the number of active source messages the routing device accepts from sources in this address range.                                                                                                                                                                                                                                                                  |
| <b>Default</b>                  | If you do not include this statement, the routing device accepts any number of MSDP active source messages.                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | The other statements are explained separately.                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 5604</a></li> </ul>                                                                                                                                                                                                                       |

## threshold (MSDP Active Source Messages)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>threshold <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp active-source-limit</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br><a href="#">msdp active-source-limit</a> ],<br>[edit protocols <a href="#">msdp active-source-limit</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp active-source-limit</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                             |
| <b>Description</b>              | Configure the random early detection (RED) threshold for MSDP active source messages.<br>This number must be less than the configured or default maximum.                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <b><i>number</i></b> —RED threshold for active source messages.<br><b>Range:</b> 1 through 1,000,000<br><b>Default:</b> 24,000                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 5604</a></li><li>• <a href="#">maximum (MSDP Active Source Messages) on page 5777</a></li></ul>                                                                                                                                                                                                                |



## traceoptions (Protocols MSDP)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre> traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;flag-modifier&gt; &lt;disable&gt;; } </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp group</a> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp group</a> <i>group-name</i> <a href="#">peer</a> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp peer</a> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp group</a> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp group</a> <i>group-name</i> <a href="#">peer</a> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp peer</a> <i>address</i>],</p> <p>[edit protocols <a href="#">msdp</a>],</p> <p>[edit protocols <a href="#">msdp group</a> <i>group-name</i>],</p> <p>[edit protocols <a href="#">msdp group</a> <i>group-name</i> <a href="#">peer</a> <i>address</i>],</p> <p>[edit protocols <a href="#">msdp peer</a> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp group</a> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp group</a> <i>group-name</i> <a href="#">peer</a> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp peer</a> <i>address</i>]</p> |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>         | <p>Configure MSDP tracing options.</p> <p>To specify more than one tracing operation, include multiple <b>flag</b> statements.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Default</b>             | <p>The default MSDP trace options are those inherited from the routing protocol's <b>traceoptions</b> statement included at the <b>[edit routing-options]</b> hierarchy level.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>             | <p><b>disable</b>—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>/var/log</b>. We recommend that you place tracing output in the <b>msdp-log</b> file.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

If you specify a maximum number of files, you must also include the **size** statement to specify the maximum file size.

**Range:** 2 through 1000 files

**Default:** 2 files

**flag *flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

#### MSDP Tracing Flags

- **keepalive**—Keepalive messages
- **packets**—All MSDP packets
- **route**—MSDP changes to the routing table
- **source-active**—Source-active packets
- **source-active-request**—Source-active request packets
- **source-active-response**—Source-active response packets

#### Global Tracing Flags

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

**Default:** If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

***flag-modifier***—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

**no-stamp**—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

**Default:** If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

**no-world-readable**—(Optional) Do not allow any user to read the log file.

**replace**—(Optional) Replace an existing trace file if there is one.

**Default:** If you do not include this option, tracing output is appended to an existing trace file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When **trace-file** again reaches this size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of trace files.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 1 MB

**world-readable**—(Optional) Allow any user to read the log file.

|                                 |                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing and trace—To view this statement in the configuration.<br>routing-control and trace-control—To add this statement to the configuration. |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Tracing MSDP Protocol Traffic on page 5600</a></li> </ul>                                  |



## CHAPTER 217

# Operational Commands (IGMP)

- clear igmp membership
- clear igmp statistics
- show igmp group
- show configuration protocols igmp
- show igmp interface
- show igmp statistics
- show system statistics igmp

## clear igmp membership

---

|                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| List of Syntax                               | <a href="#">Syntax on page 5790</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 5790</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Syntax                                       | <pre>clear igmp membership &lt;group address-range&gt; &lt;interface interface-name&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Syntax (EX Series Switch and the QFX Series) | <pre>clear igmp membership &lt;group address-range&gt; &lt;interface interface-name&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Release Information                          | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.<br>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                       |
| Description                                  | Clear Internet Group Management Protocol (IGMP) group members.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Options                                      | <p><b>none</b>—Clear all IGMP members on all interfaces and for all address ranges.</p> <p><b>group address-range</b>—(Optional) Clear all IGMP members that are in a particular address range. An example of a range is <b>224.2/16</b>. If you omit the destination prefix length, the default is <b>/32</b>.</p> <p><b>interface interface-name</b>—(Optional) Clear all IGMP group members on an interface.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level                     | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Related Documentation                        | <ul style="list-style-type: none"><li>• <a href="#">show igmp group on page 5795</a></li><li>• <a href="#">show igmp interface on page 5801</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                |
| List of Sample Output                        | <a href="#">clear igmp membership on page 5790</a><br><a href="#">clear igmp membership interface on page 5791</a><br><a href="#">clear igmp membership group on page 5792</a>                                                                                                                                                                                                                                                                                                                                                                                           |
| Output Fields                                | See <a href="#">show igmp group</a> for an explanation of output fields.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## Sample Output

### clear igmp membership

The following sample output displays IGMP group information before and after the **clear igmp membership** command is entered:

```

user@host> show igmp group
Interface Group Last Reported Timeout
so-0/0/0 224.2.127.253 10.1.128.1 186
so-0/0/0 224.2.127.254 10.1.128.1 186
so-0/0/0 239.255.255.255 10.1.128.1 187
so-0/0/0 224.1.127.255 10.1.128.1 188
local 224.0.0.6 (null) 0
local 224.0.0.5 (null) 0
local 224.2.127.254 (null) 0
local 239.255.255.255 (null) 0
local 224.0.0.2 (null) 0
local 224.0.0.13 (null) 0

```

```

user@host> clear igmp membership
Clearing Group Membership Info for so-0/0/0
Clearing Group Membership Info for so-1/0/0
Clearing Group Membership Info for so-2/0/0

```

```

user@host> show igmp group
Interface Group Last Reported Timeout
local 224.0.0.6 (null) 0
local 224.0.0.5 (null) 0
local 224.2.127.254 (null) 0
local 239.255.255.255 (null) 0
local 224.0.0.2 (null) 0
local 224.0.0.13 (null) 0

```

### clear igmp membership interface

The following sample output displays IGMP group information before and after the **clear igmp membership interface** command is issued:

```

user@host> show igmp group
Interface Group Last Reported Timeout
so-0/0/0 224.2.127.253 10.1.128.1 210
so-0/0/0 239.255.255.255 10.1.128.1 210
so-0/0/0 224.1.127.255 10.1.128.1 215
so-0/0/0 224.2.127.254 10.1.128.1 216
local 224.0.0.6 (null) 0
local 224.0.0.5 (null) 0
local 224.2.127.254 (null) 0
local 239.255.255.255 (null) 0
local 224.0.0.2 (null) 0
local 224.0.0.13 (null) 0

```

```

user@host> clear igmp membership interface so-0/0/0
Clearing Group Membership Info for so-0/0/0

```

```

user@host> show igmp group
Interface Group Last Reported Timeout
local 224.0.0.6 (null) 0
local 224.0.0.5 (null) 0
local 224.2.127.254 (null) 0
local 239.255.255.255 (null) 0
local 224.0.0.2 (null) 0
local 224.0.0.13 (null) 0

```

## clear igmp membership group

The following sample output displays IGMP group information before and after the **clear igmp membership group** command is entered:

```
user@host> show igmp group
```

| Interface | Group           | Last Reported | Timeout |
|-----------|-----------------|---------------|---------|
| so-0/0/0  | 224.2.127.253   | 10.1.128.1    | 210     |
| so-0/0/0  | 239.255.255.255 | 10.1.128.1    | 210     |
| so-0/0/0  | 224.1.127.255   | 10.1.128.1    | 215     |
| so-0/0/0  | 224.2.127.254   | 10.1.128.1    | 216     |
| local     | 224.0.0.6       | (null)        | 0       |
| local     | 224.0.0.5       | (null)        | 0       |
| local     | 224.2.127.254   | (null)        | 0       |
| local     | 239.255.255.255 | (null)        | 0       |
| local     | 224.0.0.2       | (null)        | 0       |
| local     | 224.0.0.13      | (null)        | 0       |

```
user@host> clear igmp membership group 239.225/16
```

```
Clearing Group Membership Range 239.225.0.0/16 on so-0/0/0
Clearing Group Membership Range 239.225.0.0/16 on so-1/0/0
Clearing Group Membership Range 239.225.0.0/16 on so-2/0/0
```

```
user@host> show igmp group
```

| Interface | Group           | Last Reported | Timeout |
|-----------|-----------------|---------------|---------|
| so-0/0/0  | 224.1.127.255   | 10.1.128.1    | 231     |
| so-0/0/0  | 224.2.127.254   | 10.1.128.1    | 233     |
| so-0/0/0  | 224.2.127.253   | 10.1.128.1    | 236     |
| local     | 224.0.0.6       | (null)        | 0       |
| local     | 224.0.0.5       | (null)        | 0       |
| local     | 224.2.127.254   | (null)        | 0       |
| local     | 239.255.255.255 | (null)        | 0       |
| local     | 224.0.0.2       | (null)        | 0       |
| local     | 224.0.0.13      | (null)        | 0       |



## clear igmp statistics

|                                    |                                                                                                                                                                                                                                                                                                                                      |
|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 5793</a><br><a href="#">Syntax (EX Series Switches) on page 5793</a>                                                                                                                                                                                                                                      |
| <b>Syntax</b>                      | clear igmp statistics<br><interface <i>interface-name</i> ><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                  |
| <b>Syntax (EX Series Switches)</b> | clear igmp statistics<br><interface <i>interface-name</i> >                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                             |
| <b>Description</b>                 | Clear Internet Group Management Protocol (IGMP) statistics.                                                                                                                                                                                                                                                                          |
| <b>Options</b>                     | <b>none</b> —Clear IGMP statistics on all interfaces.<br><br><b>interface <i>interface-name</i></b> —(Optional) Clear IGMP statistics for the specified interface only.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>    | clear                                                                                                                                                                                                                                                                                                                                |
| <b>List of Sample Output</b>       | <a href="#">clear igmp statistics on page 5793</a>                                                                                                                                                                                                                                                                                   |
| <b>Output Fields</b>               | See <a href="#">show igmp statistics</a> for an explanation of output fields.                                                                                                                                                                                                                                                        |

## Sample Output

### clear igmp statistics

The following sample output displays IGMP statistics information before and after the **clear igmp statistics** command is entered:

```
user@host> show igmp statistics
IGMP packet statistics for all interfaces
IGMP Message type Received Sent Rx errors
Membership Query 8883 459 0
V1 Membership Report 0 0 0
DVMRP 19784 35476 0
PIM V1 18310 0 0
Cisco Trace 0 0 0
V2 Membership Report 0 0 0
Group Leave 0 0 0
Mtrace Response 0 0 0
Mtrace Request 0 0 0
Domain Wide Report 0 0 0
V3 Membership Report 0 0 0
Other Unknown types 0 0 0
```

|                                     |   |
|-------------------------------------|---|
| IGMP v3 unsupported type            | 0 |
| IGMP v3 source required for SSM     | 0 |
| IGMP v3 mode not applicable for SSM | 0 |

|                        |      |
|------------------------|------|
| IGMP Global Statistics |      |
| Bad Length             | 0    |
| Bad Checksum           | 0    |
| Bad Receive If         | 0    |
| Rx non-local           | 1227 |

user@host> clear igmp statistics

user@host> show igmp statistics

IGMP packet statistics for all interfaces

| IGMP Message type                   | Received | Sent | Rx errors |
|-------------------------------------|----------|------|-----------|
| Membership Query                    | 0        | 0    | 0         |
| V1 Membership Report                | 0        | 0    | 0         |
| DVMRP                               | 0        | 0    | 0         |
| PIM V1                              | 0        | 0    | 0         |
| Cisco Trace                         | 0        | 0    | 0         |
| V2 Membership Report                | 0        | 0    | 0         |
| Group Leave                         | 0        | 0    | 0         |
| Mtrace Response                     | 0        | 0    | 0         |
| Mtrace Request                      | 0        | 0    | 0         |
| Domain Wide Report                  | 0        | 0    | 0         |
| V3 Membership Report                | 0        | 0    | 0         |
| Other Unknown types                 |          |      | 0         |
| IGMP v3 unsupported type            |          |      | 0         |
| IGMP v3 source required for SSM     |          |      | 0         |
| IGMP v3 mode not applicable for SSM |          |      | 0         |
| IGMP Global Statistics              |          |      |           |
| Bad Length                          | 0        |      |           |
| Bad Checksum                        | 0        |      |           |
| Bad Receive If                      | 0        |      |           |
| Rx non-local                        | 0        |      |           |

## show igmp group

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-----------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 5795</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 5795</a>                                                                                                                                                                                                                                                                                                          |
| <b>Syntax</b>                                       | <pre>show igmp group &lt;brief   detail&gt; &lt;group-name&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>                                                                                                                                                                                                                                                                                                   |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show igmp group &lt;brief   detail&gt; &lt;group-name&gt;</pre>                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>                          | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                   |
| <b>Description</b>                                  | Display Internet Group Management Protocol (IGMP) group membership information.                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                                      | <p><b>none</b>—Display standard information about membership for all IGMP groups.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>group-name</b>—(Optional) Display group membership for the specified IP address only.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>List of Sample Output</b>                        | <a href="#">show igmp group (Include Mode) on page 5796</a><br><a href="#">show igmp group (Exclude Mode) on page 5797</a><br><a href="#">show igmp group brief on page 5797</a><br><a href="#">show igmp group detail on page 5797</a>                                                                                                                                                                                   |
| <b>Output Fields</b>                                | <a href="#">Table 439</a> describes the output fields for the <b>show igmp group</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                     |

**Table 439: show igmp group Output Fields**

| Field Name        | Field Description                                                                                                                                       | Level of Output |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Interface</b>  | Name of the interface that received the IGMP membership report. A name of <b>local</b> indicates that the local routing device joined the group itself. | All levels      |
| <b>Group</b>      | Group address.                                                                                                                                          | All levels      |
| <b>Group Mode</b> | Mode the SSM group is operating in: <b>Include</b> or <b>Exclude</b> .                                                                                  | All levels      |
| <b>Source</b>     | Source address.                                                                                                                                         | All levels      |

Table 439: show igmp group Output Fields (*continued*)

| Field Name       | Field Description                                                                                                                                                                                                         | Level of Output |
|------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Source timeout   | Time remaining until the group traffic is no longer forwarded. The timer is refreshed when a listener in include mode sends a report. A group in exclude mode or configured as a static group displays a zero timer.      | detail          |
| Last reported by | Address of the host that last reported membership in this group.                                                                                                                                                          | All levels      |
| Timeout          | Time remaining until the group membership is removed.                                                                                                                                                                     | brief none      |
| Group timeout    | Time remaining until a group in exclude mode moves to include mode. The timer is refreshed when a listener in exclude mode sends a report. A group in include mode or configured as a static group displays a zero timer. | detail          |
| Type             | Type of group membership: <ul style="list-style-type: none"> <li>• <b>Dynamic</b>—Host reported the membership.</li> <li>• <b>Static</b>—Membership is configured.</li> </ul>                                             | All levels      |

## Sample Output

### show igmp group (Include Mode)

```

user@host> show igmp group
Interface: t1-0/1/0.0
 Group: 232.1.1.1
 Group mode: Include
 Source: 10.0.0.2
 Last reported by: 10.9.5.2
 Timeout: 24 Type: Dynamic
 Group: 232.1.1.1
 Group mode: Include
 Source: 10.0.0.3
 Last reported by: 10.9.5.2
 Timeout: 24 Type: Dynamic
 Group: 232.1.1.1
 Group mode: Include
 Source: 10.0.0.4
 Last reported by: 10.9.5.2
 Timeout: 24 Type: Dynamic
 Group: 232.1.1.2
 Group mode: Include
 Source: 10.0.0.4
 Last reported by: 10.9.5.2
 Timeout: 24 Type: Dynamic
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
 Group: 224.0.0.2
 Source: 0.0.0.0
 Last reported by: Local
 Timeout: 0 Type: Dynamic
 Group: 224.0.0.22
 Source: 0.0.0.0

```

```

Last reported by: Local
Timeout: 0 Type: Dynamic

```

### show igmp group (Exclude Mode)

```

user@host> show igmp group
Interface: t1-0/1/0.0
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
 Group: 224.0.0.2
 Source: 0.0.0.0
 Last reported by: Local
 Timeout: 0 Type: Dynamic
 Group: 224.0.0.22
 Source: 0.0.0.0
 Last reported by: Local
 Timeout: 0 Type: Dynamic

```

### show igmp group brief

The output for the **show igmp group brief** command is identical to that for the **show igmp group** command.

### show igmp group detail

```

user@host> show igmp group detail
Interface: t1-0/1/0.0
 Group: 232.1.1.1
 Group mode: Include
 Source: 10.0.0.2
 Source timeout: 12
 Last reported by: 10.9.5.2
 Group timeout: 0 Type: Dynamic
 Group: 232.1.1.1
 Group mode: Include
 Source: 10.0.0.3
 Source timeout: 12
 Last reported by: 10.9.5.2
 Group timeout: 0 Type: Dynamic
 Group: 232.1.1.1
 Group mode: Include
 Source: 10.0.0.4
 Source timeout: 12
 Last reported by: 10.9.5.2
 Group timeout: 0 Type: Dynamic
 Group: 232.1.1.2
 Group mode: Include
 Source: 10.0.0.4
 Source timeout: 12
 Last reported by: 10.9.5.2
 Group timeout: 0 Type: Dynamic
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
 Group: 224.0.0.2
 Group mode: Exclude
 Source: 0.0.0.0
 Source timeout: 0

```

```
 Last reported by: Local
 Group timeout: 0 Type: Dynamic
Group: 224.0.0.22
 Group mode: Exclude
 Source: 0.0.0.0
 Source timeout: 0
 Last reported by: Local
 Group timeout: 0 Type: Dynamic
```

## show configuration protocols igmp

|                                 |                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show configuration protocols igmp                                                                                                                                         |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.3 for the QFX Series.<br>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                 |
| <b>Description</b>              | Display Internet Group Management Protocol (IGMP) information.                                                                                                            |
| <b>Required Privilege Level</b> | view                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">IGMP Snooping Overview on page 5431</a></li> <li>• <a href="#">Configuring IGMP Snooping on page 5434</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show configuration protocols igmp on page 5799</a>                                                                                                            |
| <b>Output Fields</b>            | <a href="#">Table 439</a> describes the output fields for the <b>show configuration protocols igmp</b> command that relate to IGMP querying.                              |

**Table 440: show igmp group Output Fields**

| Field Name              | Field Description                                                                                        | Level of Output |
|-------------------------|----------------------------------------------------------------------------------------------------------|-----------------|
| accounting              | Enables notification for join and leave events.                                                          | All levels      |
| igmp-querier            | Configured source address for the IGMP querier.                                                          | All levels      |
| interface               | Name of the interface that receives IGMP membership reports.                                             | All levels      |
| query-interval          | Interval at which the IGMP querier sends general host-query messages to solicit membership information.  | All levels      |
| query-response-interval | How long the IGMP querier waits to receive a response from a query message before sending another query. | All levels      |
| src-address             | Source address of IGMP queries.                                                                          |                 |
| version                 | IGMP version.                                                                                            | All levels      |

## Sample Output

### show configuration protocols igmp

```

user@switch> show configuration protocols igmp
query-interval 150;
query-response-interval 50;
accounting;
interface vlan.43 {
 version 2;
}
igmp-querier {

```

```
src-address 10.0.0.2;
}
```



## show igmp interface

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-----------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 5801</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 5801</a>                                                                                                                                                                                                                                                                                                                    |
| <b>Syntax</b>                                       | <pre>show igmp interface &lt;brief   detail&gt; &lt;interface-name&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>                                                                                                                                                                                                                                                                                                     |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show igmp interface &lt;brief   detail&gt; &lt;interface-name&gt;</pre>                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>                          | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                             |
| <b>Description</b>                                  | Display information about Internet Group Management Protocol (IGMP)-enabled interfaces.                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                                      | <p><b>none</b>—Display standard information about all IGMP-enabled interfaces.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>interface-name</b>—(Optional) Display information about the specified IGMP-enabled interface only.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>                        | <ul style="list-style-type: none"> <li>• <a href="#">clear igmp membership on page 5790</a></li> </ul>                                                                                                                                                                                                                                                                                                                              |
| <b>List of Sample Output</b>                        | <a href="#">show igmp interface on page 5803</a><br><a href="#">show igmp interface brief on page 5803</a><br><a href="#">show igmp interface detail on page 5804</a>                                                                                                                                                                                                                                                               |
| <b>Output Fields</b>                                | <a href="#">Table 441</a> describes the output fields for the <b>show igmp interface</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                           |

**Table 441: show igmp interface Output Fields**

| Field Name | Field Description                                                               | Level of Output |
|------------|---------------------------------------------------------------------------------|-----------------|
| Interface  | Name of the interface.                                                          | All levels      |
| Querier    | Address of the routing device that has been elected to send membership queries. | All levels      |

Table 441: show igmp interface Output Fields (*continued*)

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Level of Output |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>State</b>            | State of the interface: <b>Up</b> or <b>Down</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | All levels      |
| <b>SSM Map Policy</b>   | Name of the source-specific multicast (SSM) map policy that has been applied to the IGMP interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | All levels      |
| <b>Timeout</b>          | How long until the IGMP querier is declared to be unreachable, in seconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | All levels      |
| <b>Version</b>          | IGMP version being used on the interface: <b>1</b> , <b>2</b> , or <b>3</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | All levels      |
| <b>Groups</b>           | Number of groups on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | All levels      |
| <b>Immediate Leave</b>  | State of the immediate leave option: <ul style="list-style-type: none"> <li><b>On</b>—Indicates that the router removes a host from the multicast group as soon as the router receives a leave group message from a host associated with the interface.</li> <li><b>Off</b>—Indicates that after receiving a leave group message, instead of removing a host from the multicast group immediately, the router sends a group query to determine if another receiver responds.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                   | All levels      |
| <b>Promiscuous Mode</b> | State of the promiscuous mode option: <ul style="list-style-type: none"> <li><b>On</b>—Indicates that the router can accept IGMP reports from subnetworks that are not associated with its interfaces.</li> <li><b>Off</b>—Indicates that the router can accept IGMP reports only from subnetworks that are associated with its interfaces.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | All levels      |
| <b>Passive</b>          | State of the passive mode option: <ul style="list-style-type: none"> <li><b>On</b>—Indicates that the router can run IGMP on the interface but not send or receive control traffic such as IGMP reports, queries, and leaves.</li> <li><b>Off</b>—Indicates that the router can run IGMP on the interface and send or receive control traffic such as IGMP reports, queries, and leaves.</li> </ul> <p>The <b>passive</b> statement enables you to selectively activate up to two out of a possible three available query or control traffic options. When enabled, the following options appear after the <b>on</b> state declaration:</p> <ul style="list-style-type: none"> <li><b>send-general-query</b>—The interface sends general queries.</li> <li><b>send-group-query</b>—The interface sends group-specific and group-source-specific queries.</li> <li><b>allow-receive</b>—The interface receives control traffic.</li> </ul> | All levels      |
| <b>OIF map</b>          | Name of the OIF map (if configured) associated with the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | All levels      |
| <b>SSM map</b>          | Name of the source-specific multicast (SSM) map (if configured) used on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | All levels      |

Table 441: show igmp interface Output Fields (*continued*)

| Field Name                   | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Level of Output |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Configured Parameters</b> | Information configured by the user: <ul style="list-style-type: none"> <li><b>IGMP Query Interval</b>—Interval (in seconds) at which this router sends membership queries when it is the querier.</li> <li><b>IGMP Query Response Interval</b>—Time (in seconds) that the router waits for a report in response to a general query.</li> <li><b>IGMP Last Member Query Interval</b>—Time (in seconds) that the router waits for a report in response to a group-specific query.</li> <li><b>IGMP Robustness Count</b>—Number of times the router retries a query.</li> </ul> | All levels      |
| <b>Derived Parameters</b>    | Derived information: <ul style="list-style-type: none"> <li><b>IGMP Membership Timeout</b>—Timeout period (in seconds) for group membership. If no report is received for these groups before the timeout expires, the group membership is removed.</li> <li><b>IGMP Other Querier Present Timeout</b>—Time (in seconds) that the router waits for the IGMP querier to send a query.</li> </ul>                                                                                                                                                                              | All levels      |

## Sample Output

### show igmp interface

```

user@host> show igmp interface
Interface: at-0/3/1.0
 Querier: 10.111.30.1
 State: Up Timeout: None Version: 2 Groups: 4
 SSM Map Policy: ssm-policy-A
Interface: so-1/0/0.0
 Querier: 10.111.10.1
 State: Up Timeout: None Version: 2 Groups: 2
 SSM Map Policy: ssm-policy-B
Interface: so-1/0/1.0
 Querier: 10.111.20.1
 State: Up Timeout: None Version: 2 Groups: 4
 SSM Map Policy: ssm-policy-C
Immediate Leave: On
Promiscuous Mode: Off

Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2

Derived Parameters:
IGMP Membership Timeout: 260.0
IGMP Other Querier Present Timeout: 255.0

```

### show igmp interface brief

The output for the **show igmp interface brief** command is identical to that for the **show igmp interface** command. For sample output, see [show igmp interface on page 5803](#).

### [show igmp interface detail](#)

The output for the **show igmp interface detail** command is identical to that for the **show igmp interface** command. For sample output, see [show igmp interface on page 5803](#).

## show igmp statistics

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-----------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 5805</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 5805</a>                                                                                                                                                                                                                                                                                                               |
| <b>Syntax</b>                                       | <pre>show igmp statistics &lt;brief   detail&gt; &lt;interface <i>interface-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                       |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show igmp statistics &lt;brief   detail&gt; &lt;interface <i>interface-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>                          | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                        |
| <b>Description</b>                                  | Display Internet Group Management Protocol (IGMP) statistics.                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                                      | <p><b>none</b>—Display IGMP statistics for all interfaces.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display IGMP statistics about the specified interface only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>                        | <ul style="list-style-type: none"> <li>• <a href="#">clear igmp statistics on page 5793</a></li> </ul>                                                                                                                                                                                                                                                                                                                         |
| <b>List of Sample Output</b>                        | <a href="#">show igmp statistics on page 5806</a><br><a href="#">show igmp statistics interface on page 5807</a>                                                                                                                                                                                                                                                                                                               |
| <b>Output Fields</b>                                | Table 442 describes the output fields for the <b>show igmp statistics</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                     |

**Table 442: show igmp statistics Output Fields**

| Field Name             | Field Description                                                                          |
|------------------------|--------------------------------------------------------------------------------------------|
| IGMP packet statistics | Heading for IGMP packet statistics for all interfaces or for the specified interface name. |

Table 442: show igmp statistics Output Fields (*continued*)

| Field Name                    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>IGMP Message type</b>      | <p>Summary of IGMP statistics:</p> <ul style="list-style-type: none"> <li>• <b>Membership Query</b>—Number of membership queries sent and received.</li> <li>• <b>V1 Membership Report</b>—Number of version 1 membership reports sent and received.</li> <li>• <b>DVMRP</b>—Number of DVMRP messages sent or received.</li> <li>• <b>PIM V1</b>—Number of PIM version 1 messages sent or received.</li> <li>• <b>Cisco Trace</b>—Number of Cisco trace messages sent or received.</li> <li>• <b>V2 Membership Report</b>—Number of version 2 membership reports sent or received.</li> <li>• <b>Group Leave</b>—Number of group leave messages sent or received.</li> <li>• <b>Mtrace Response</b>—Number of Mtrace response messages sent or received.</li> <li>• <b>Mtrace Request</b>—Number of Mtrace request messages sent or received.</li> <li>• <b>Domain Wide Report</b>—Number of domain-wide reports sent or received.</li> <li>• <b>V3 Membership Report</b>—Number of version 3 membership reports sent or received.</li> <li>• <b>Other Unknown types</b>—Number of unknown message types received.</li> <li>• <b>IGMP v3 unsupported type</b>—Number of messages received with unknown and unsupported IGMP version 3 message types.</li> <li>• <b>IGMP v3 source required for SSM</b>—Number of IGMP version 3 messages received that contained no source.</li> <li>• <b>IGMP v3 mode not applicable for SSM</b>—Number of IGMP version 3 messages received that did not contain a mode applicable for source-specific multicast (SSM).</li> </ul> |
| <b>Received</b>               | Number of messages received.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Sent</b>                   | Number of messages sent.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Rx errors</b>              | Number of received packets that contained errors.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>IGMP Global Statistics</b> | <p>Summary of IGMP statistics for all interfaces.</p> <ul style="list-style-type: none"> <li>• <b>Bad Length</b>—Number of messages received with length errors so severe that further classification could not occur.</li> <li>• <b>Bad Checksum</b>—Number of messages received with a bad IP checksum. No further classification was performed.</li> <li>• <b>Bad Receive If</b>—Number of messages received on an interface not enabled for IGMP.</li> <li>• <b>Rx non-local</b>—Number of messages received from senders that are not local.</li> <li>• <b>Timed out</b>—Number of groups that timed out as a result of not receiving an explicit leave message.</li> <li>• <b>Rejected Report</b>—Number of reports dropped because of the IGMP group policy.</li> <li>• <b>Total Interfaces</b>—Number of interfaces configured to support IGMP.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## Sample Output

### show igmp statistics

```

user@host> show igmp statistics
IGMP packet statistics for all interfaces
IGMP Message type Received Sent Rx errors
Membership Query 8883 459 0
V1 Membership Report 0 0 0

```

|                                     |      |   |   |
|-------------------------------------|------|---|---|
| DVMRP                               | 0    | 0 | 0 |
| PIM V1                              | 0    | 0 | 0 |
| Cisco Trace                         | 0    | 0 | 0 |
| V2 Membership Report                | 0    | 0 | 0 |
| Group Leave                         | 0    | 0 | 0 |
| Mtrace Response                     | 0    | 0 | 0 |
| Mtrace Request                      | 0    | 0 | 0 |
| Domain Wide Report                  | 0    | 0 | 0 |
| V3 Membership Report                | 0    | 0 | 0 |
| Other Unknown types                 |      |   | 0 |
| IGMP v3 unsupported type            |      |   | 0 |
| IGMP v3 source required for SSM     |      |   | 0 |
| IGMP v3 mode not applicable for SSM |      |   | 0 |
| IGMP Global Statistics              |      |   |   |
| Bad Length                          | 0    |   |   |
| Bad Checksum                        | 0    |   |   |
| Bad Receive If                      | 0    |   |   |
| Rx non-local                        | 1227 |   |   |
| Timed out                           | 0    |   |   |
| Rejected Report                     | 0    |   |   |
| Total Interfaces                    | 2    |   |   |

#### show igmp statistics interface

```

user@host> show igmp statistics interface fe-1/0/1.0
IGMP interface packet statistics for fe-1/0/1.0
IGMP Message type Received Sent Rx errors
Membership Query 0 230 0
V1 Membership Report 0 0 0

```

## show system statistics igmp

---

|                                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                 | <a href="#">Syntax on page 5808</a><br><a href="#">Syntax (EX Series Switches) on page 5808</a><br><a href="#">Syntax (TX Matrix Router) on page 5808</a><br><a href="#">Syntax (TX Matrix Plus Router) on page 5808</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Syntax</b>                         | show system statistics igmp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Syntax (EX Series Switches)</b>    | show system statistics igmp<br><all-members><br><local><br><member <i>member-id</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Syntax (TX Matrix Router)</b>      | show system statistics igmp<br><all-chassis   all-lcc   lcc <i>number</i>   scc>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Syntax (TX Matrix Plus Router)</b> | show system statistics igmp<br><all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>            | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br><b>sfc</b> option introduced for the TX Matrix Plus router in Junos OS Release 9.6.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.<br>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>                    | Display system-wide Internet Group Management Protocol (IGMP) statistics.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                        | <b>none</b> —Display system statistics for IGMP.<br><br><b>all-chassis</b> —(TX Matrix routers and TX Matrix Plus routers only) (Optional) Display system statistics for IGMP for all the routers in the chassis.<br><br><b>all-lcc</b> —(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system statistics for IGMP for all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for IGMP for all connected T1600 or T4000 LCCs.<br><br><b>all-members</b> —(EX4200 switches only) (Optional) Display IGMP statistics for all members of the Virtual Chassis configuration.<br><br><b>lcc <i>number</i></b> —(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system statistics for IGMP for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for IGMP for a specific router that is connected to the TX Matrix Plus router. |



Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(EX4200 switches only) (Optional) Display IGMP statistics for the local Virtual Chassis member.

**member *member-id***—(EX4200 switches only) (Optional) Display IGMP statistics for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value from 0 through 9.

**scc**—(TX Matrix routers only) (Optional) Display system statistics for IGMP for the TX Matrix router (or switch-card chassis).

**sfc *number***—(TX Matrix Plus routers only) (Optional) Display system statistics for IGMP for the TX Matrix Plus router. Replace *number* with 0.

**Additional Information** By default, when you issue the **show system statistics igmp** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

**Required Privilege Level** view

**Related Documentation** • [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

**List of Sample Output** [show system statistics igmp on page 5809](#)  
[show system statistics igmp \(EX Series Switches\) on page 5810](#)  
[show system statistics igmp \(TX Matrix Plus Router\) on page 5810](#)

## Sample Output

**show system statistics igmp**

```
user@host> show system statistics igmp
igmp:
 17178 messages received
 0 messages received with too few bytes
 0 messages received with bad checksum
 0 membership queries received
```

```
0 membership queries received with invalid field(s)
0 membership reports received
0 membership reports received with invalid field(s)
0 membership reports received for groups to which we belong
0 membership reports sent
```

#### show system statistics igmp (EX Series Switches)

```
user@host> show system statistics igmp
igmp:
 0 messages received
 0 messages received with too few bytes
 0 messages received with bad checksum
 0 membership queries received
 0 membership queries received with invalid fields
 0 membership reports received
 0 membership reports received with invalid fields
 0 membership reports received for groups to which we belong
 0 Membership reports sent
```

#### show system statistics igmp (TX Matrix Plus Router)

```
user@host> show system statistics igmp
sfc0-re0:

igmp:
 0 messages received
 0 messages received with too few bytes
 0 messages received with bad checksum
 0 membership queries received
 0 membership queries received with invalid field(s)
 0 membership reports received
 0 membership reports received with invalid field(s)
 0 membership reports received for groups to which we belong
 0 membership reports sent
```

```
lcc0-re0:

igmp:
 0 messages received
 0 messages received with too few bytes
 0 messages received with bad checksum
 0 membership queries received
 0 membership queries received with invalid field(s)
 0 membership reports received
 0 membership reports received with invalid field(s)
 0 membership reports received for groups to which we belong
 0 membership reports sent
```

```
lcc1-re0:

igmp:
 0 messages received
 0 messages received with too few bytes
 0 messages received with bad checksum
 0 membership queries received
 0 membership queries received with invalid field(s)
 0 membership reports received
 0 membership reports received with invalid field(s)
 0 membership reports received for groups to which we belong
 0 membership reports sent
```

lcc2-re0:

-----  
igmp:

0 messages received  
0 messages received with too few bytes  
0 messages received with bad checksum  
0 membership queries received  
0 membership queries received with invalid field(s)  
0 membership reports received  
0 membership reports received with invalid field(s)  
0 membership reports received for groups to which we belong  
0 membership reports sent

lcc3-re0:

-----  
igmp:

0 messages received  
0 messages received with too few bytes  
0 messages received with bad checksum  
0 membership queries received  
0 membership queries received with invalid field(s)  
0 membership reports received  
0 membership reports received with invalid field(s)  
0 membership reports received for groups to which we belong  
0 membership reports sent



## CHAPTER 218

# Operational Commands (IGMP Snooping)

- `clear igmp-snooping membership`
- `clear igmp-snooping statistics`
- `show igmp-snooping membership`
- `show igmp-snooping route`
- `show igmp-snooping statistics`
- `show igmp-snooping vlans`

## clear igmp-snooping membership

---

|                                 |                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>clear igmp-snooping membership</b><br><b>&lt;vlan <i>vlan-name</i>&gt;</b>                                |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.                                              |
| <b>Description</b>              | Clear IGMP snooping membership information.                                                                  |
| <b>Options</b>                  | <b>vlan <i>vlan-name</i></b> —(Optional) Name of the VLAN.                                                   |
| <b>Required Privilege Level</b> | view                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show igmp-snooping membership on page 5816</a></li></ul> |
| <b>List of Sample Output</b>    | <a href="#">clear igmp-snooping membership on page 5814</a>                                                  |

### Sample Output

#### clear igmp-snooping membership

```
user@switch> clear igmp-snooping membership vlan employee-vlan
```

---

## clear igmp-snooping statistics

---

|                                 |                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>clear igmp-snooping statistics</code>                                                                  |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.                                              |
| <b>Description</b>              | Clear IGMP snooping statistics.                                                                              |
| <b>Required Privilege Level</b> | view                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show igmp-snooping statistics on page 5821</a></li></ul> |
| <b>List of Sample Output</b>    | <a href="#">clear igmp-snooping statistics on page 5815</a>                                                  |

### Sample Output

#### clear igmp-snooping statistics

```
user@switch> clear igmp-snooping statistics
```

## show igmp-snooping membership

|                                 |                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show igmp-snooping membership &lt;brief   detail&gt; &lt;interface <i>interface-name</i>&gt; &lt;vlan <i>vlan-id</i>   <i>vlan-name</i>&gt;</pre>                                                                                                                                                                                                                              |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>IGMPv3 output introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                 |
| <b>Description</b>              | Display IGMP snooping membership information.                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <p><b>none</b>—Display general parameters.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display IGMP snooping information for the specified interface.</p> <p><b>vlan <i>vlan-id</i>   <i>vlan-name</i></b>—(Optional) Display IGMP snooping information for the specified VLAN.</p>     |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Monitoring IGMP Snooping on page 5438</a></li> <li>• <a href="#">Configuring IGMP Snooping on page 5434</a></li> <li>• <a href="#">show igmp-snooping route on page 5819</a></li> <li>• <a href="#">show igmp-snooping statistics on page 5821</a></li> <li>• <a href="#">show igmp-snooping vlans on page 5823</a></li> </ul> |
| <b>List of Sample Output</b>    | <p><a href="#">show igmp-snooping membership on page 5817</a></p> <p><a href="#">show igmp-snooping membership detail on page 5818</a></p>                                                                                                                                                                                                                                          |
| <b>Output Fields</b>            | <p><a href="#">Table 443</a> lists the output fields for the <b>show igmp-snooping membership</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                              |

**Table 443: show igmp-snooping membership Output Fields**

| Field Name | Field Description                 | Level of Output |
|------------|-----------------------------------|-----------------|
| VLAN       | Name of the VLAN.                 | All             |
| Interfaces | Interfaces assigned to the VLAN.  | All             |
| Tag        | Numerical identifier of the VLAN. | <b>detail</b>   |



Table 443: show igmp-snooping membership Output Fields (*continued*)

| Field Name          | Field Description                                                                                                                                                                                     | Level of Output |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Router interfaces   | Names of multicast router interfaces.                                                                                                                                                                 | <b>detail</b>   |
| • static or dynamic | Whether the multicast router interface is <b>static</b> or <b>dynamic</b> .                                                                                                                           | <b>detail</b>   |
| • Uptime            | For static interfaces, length of time since the interface was configured as a multicast router interface; for dynamic interfaces, length of time since the first query was received on the interface. | <b>detail</b>   |
| • timeout           | Query timeout in seconds.                                                                                                                                                                             | <b>detail</b>   |
| Group               | IP multicast address of the multicast group.                                                                                                                                                          | <b>detail</b>   |
| Receiver count      | Number of interfaces that have membership in a multicast group.                                                                                                                                       | <b>detail</b>   |
| Flags               | IGMP version of the host sending a join message.                                                                                                                                                      | <b>detail</b>   |
| Uptime              | Length of time a multicast group has been active on the interface.                                                                                                                                    | <b>detail</b>   |
| timeout             | Time (in seconds) left until the entry for the multicast group is removed.                                                                                                                            | All             |
| Last reporter       | Last host to report membership for the multicast group.                                                                                                                                               | <b>detail</b>   |
| Include source      | Source addresses from which multicast streams are allowed based on IGMPv3 reports.                                                                                                                    | <b>detail</b>   |

## Sample Output

### show igmp-snooping membership

```

user@switch> show igmp-snooping membership
VLAN: v1
 224.1.1.1 * 258 secs
 Interfaces: ge-0/0/0.0
 224.1.1.3 * 258 secs
 Interfaces: ge-0/0/0.0
 224.1.1.5 * 258 secs
 Interfaces: ge-0/0/0.0
 224.1.1.7 * 258 secs

```

```
Interfaces: ge-0/0/0.0
224.1.1.9 * 258 secs
Interfaces: ge-0/0/0.0
224.1.1.11 * 258 secs
Interfaces: ge-0/0/0.0
```

### show igmp-snooping membership detail

```
user@switch> show igmp-snooping membership detail
VLAN: v43 Tag: 43 (Index: 4)
Group: 225.0.0.2
Receiver count: 1, Flags: <V3-hosts>
 ge-0/0/15.0 Uptime: 00:00:11 timeout: 248 Last reporter: 10.2.10.16
 Include source: 1.2.1.1, 1.3.1.1
VLAN: v44 Tag: 44 (Index: 5)
Group: 225.0.0.1
Receiver count: 1, Flags: <V2-hosts>
 ge-0/0/21.0 Uptime: 00:00:02 timeout: 257
VLAN: v110 Tag: 110 (Index: 4)
Router interfaces:
 ge-0/0/3.0 static Uptime: 00:08:45
 ge-0/0/2.0 static Uptime: 00:08:45
 ge-0/0/4.0 dynamic Uptime: 00:16:41 timeout: 254
Group: 225.0.0.3
Receiver count: 1, Flags: <V3-hosts>
 ge-0/0/5.0 Uptime: 00:00:19 timeout: 259
Group: 225.1.1.1
Receiver count: 1, Flags: <V2-hosts>
 ge-0/0/5.0 Uptime: 00:22:43 timeout: 96
Group: 225.2.2.2
Receiver count: 1, Flags: <V2-hosts Static>
 ge-0/0/5.0 Uptime: 00:23:13
```

## show igmp-snooping route

|                                 |                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show igmp-snooping route &lt;brief   detail&gt; &lt;ethernet-switching &lt;brief   detail   vlan (vlan-id   vlan-name )&gt;&gt; &lt;inet &lt;brief   detail   vlan vlan-name&gt;&gt; &lt;vlan vlan-name&gt;</pre>                                                                                                                                                     |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Display IGMP snooping route information.                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <p><b>none</b>—Display general parameters.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>ethernet-switching</b>—(Optional) Display Ethernet switching information.</p> <p><b>inet</b>—(Optional) Display <b>inet</b> information.</p> <p><b>vlan vlan-name</b>—(Optional) Display route information for the specified VLAN.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Monitoring IGMP Snooping on page 5438</a></li> <li>• <a href="#">Configuring IGMP Snooping on page 5434</a></li> <li>• <a href="#">show igmp-snooping statistics on page 5821</a></li> <li>• <a href="#">show igmp-snooping vlans on page 5823</a></li> </ul>                                                         |
| <b>List of Sample Output</b>    | <p><a href="#">show igmp-snooping route on page 5820</a></p> <p><a href="#">show igmp-snooping route vlan v1 on page 5820</a></p>                                                                                                                                                                                                                                          |
| <b>Output Fields</b>            | Table 444 lists the output fields for the <b>show igmp-snooping route</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                 |

**Table 444: show igmp-snooping route Output Fields**

| Field Name | Field Description                              |
|------------|------------------------------------------------|
| Table      | (For internal use only. Value is always 0.)    |
| VLAN       | Name of the VLAN.                              |
| Group      | Multicast group address.                       |
| Interfaces | Interfaces on which IGMP packets were snooped. |
| Next-hop   | ID associated with the next-hop device.        |

## Sample Output

### show igmp-snooping route

```
user@switch> show igmp-snooping route
VLAN Group Next-hop
V11 224.1.1.1, * 533
 Interfaces: ge-0/0/13.0, ge-0/0/1.0
VLAN Group Next-hop
v12 224.1.1.3, * 534
 Interfaces: ge-0/0/13.0, ge-0/0/0.0
```

### show igmp-snooping route vlan v1

```
user@switch> show igmp-snooping route vlan v1
Table: 0
VLAN Group Next-hop
v1 224.1.1.1, * 1266
 Interfaces: ge-0/0/0.0
v1 224.1.1.3, * 1266
 Interfaces: ge-0/0/0.0
v1 224.1.1.5, * 1266
 Interfaces: ge-0/0/0.0
v1 224.1.1.7, * 1266
 Interfaces: ge-0/0/0.0
v1 224.1.1.9, * 1266
 Interfaces: ge-0/0/0.0
v1 224.1.1.11, * 1266
 Interfaces: ge-0/0/0.0
```

## show igmp-snooping statistics

|                                 |                                                                                                                                                                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show igmp-snooping statistics</b>                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                               |
| <b>Description</b>              | Display IGMP snooping statistics.                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Monitoring IGMP Snooping on page 5438</a></li> <li>• <a href="#">Configuring IGMP Snooping on page 5434</a></li> <li>• <a href="#">show igmp-snooping route on page 5819</a></li> <li>• <a href="#">show igmp-snooping vlans on page 5823</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show igmp-snooping statistics on page 5822</a>                                                                                                                                                                                                                                                    |
| <b>Output Fields</b>            | Table 445 lists the output fields for the <b>show igmp-snooping statistics</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                               |

**Table 445: show igmp-snooping statistics Output Fields**

| Field Name        | Field Description                                                                           |
|-------------------|---------------------------------------------------------------------------------------------|
| Bad length        | IGMP packet has illegal or bad length.                                                      |
| Bad checksum      | IGMP or IP checksum is incorrect.                                                           |
| Invalid interface | Packet was received through an invalid interface.                                           |
| Not local         | Number of packets received from senders that are not local.                                 |
| Receive unknown   | Unknown IGMP type.                                                                          |
| Timed out         | Number of timeouts for all multicast groups.                                                |
| IGMP Type         | Type of IGMP message ( <b>Queries</b> , <b>Reports</b> , <b>Leaves</b> , or <b>Other</b> ). |
| Received          | Number of IGMP packets received.                                                            |
| Transmitted       | Number of IGMP packets transmitted.                                                         |
| Recv Errors       | Number of general receive errors.                                                           |

## Sample Output

### show igmp-snooping statistics

```
user@switch> show igmp-snooping statistics
Bad length: 0 Bad checksum: 0 Invalid interface: 0
Not local: 0 Receive unknown: 0 Timed out: 58

IGMP Type Received Transmitted Recv Errors
Queries: 74295 0 0
Reports: 18148423 0 16333523
Leaves: 0 0 0
Other: 0 0 0
```

## show igmp-snooping vlans

|                                 |                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show igmp-snooping vlans</b><br><brief   detail><br><vlan <i>vlan-id</i>   <i>vlan-name</i> >                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Display IGMP snooping VLAN information.                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <p><b>none</b>—Display general parameters.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>vlan <i>vlan-id</i>   vlan <i>vlan-number</i></b>—(Optional) Display VLAN information for the specified VLAN.</p>                                                              |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Monitoring IGMP Snooping on page 5438</a></li> <li>• <a href="#">Configuring IGMP Snooping on page 5434</a></li> <li>• <a href="#">show igmp-snooping route on page 5819</a></li> <li>• <a href="#">show igmp-snooping statistics on page 5821</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show igmp-snooping vlans on page 5824</a><br><a href="#">show igmp-snooping vlans vlan on page 5824</a><br><a href="#">show igmp-snooping vlans vlan detail on page 5824</a>                                                                                                                           |
| <b>Output Fields</b>            | Table 446 lists the output fields for the <b>show igmp-snooping vlans</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                         |

**Table 446: show igmp-snooping vlans Output Fields**

| Field Name             | Field Description                                                       | Level of Output |
|------------------------|-------------------------------------------------------------------------|-----------------|
| <b>VLAN</b>            | Name of the VLAN.                                                       | All levels      |
| <b>IGMP-L2-Querier</b> | Source address for IGMP snooping queries (if switch is an IGMP querier) | All levels      |
| <b>Interfaces</b>      | Number of interfaces in the VLAN.                                       | All levels      |
| <b>Groups</b>          | Number of groups in the VLAN.                                           | All levels      |
| <b>MRouters</b>        | Number of multicast routers associated with the VLAN.                   | All levels      |
| <b>Receivers</b>       | Number of host receivers in the VLAN.                                   | All levels      |

Table 446: show igmp-snooping vlans Output Fields (*continued*)

| Field Name         | Field Description                                                                                                                                                                                          | Level of Output |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Tag                | Numerical identifier of the VLAN.                                                                                                                                                                          | detail          |
| tagged   untagged  | Interface participates in a tagged (802.1Q) or untagged (native) VLAN.                                                                                                                                     | detail          |
| vlan-interface     | Internal VLAN interface identifier.                                                                                                                                                                        | detail          |
| Membership timeout | Membership timeout value.                                                                                                                                                                                  | detail          |
| Querier timeout    | Timeout value for interfaces dynamically marked as router or switch interfaces (interfaces that receive queries). When the querier timeout is reached, the switch marks the interface as a host interface. | detail          |
| Interface          | Name of the interface.                                                                                                                                                                                     | detail          |
| Reporters          | Number of dynamic groups on an interface.                                                                                                                                                                  | detail          |

## Sample Output

### show igmp-snooping vlans

```

user@switch> show igmp-snooping vlans
VLAN Interfaces Groups MRouters Receivers
default 0 0 0 0
v1 11 50 0 0
v10 1 0 0 0
v11 1 0 0 0
v180 3 0 1 0
v181 3 0 0 0
v182 3 0 0 0

```

### show igmp-snooping vlans vlan

```

user@switch> show igmp-snooping vlans vlan v10
user@switch> show igmp-snooping vlans vlan v10
VLAN Interfaces Groups MRouters Receivers
v10 1 0 0 0

```

### show igmp-snooping vlans vlan detail

```

user@switch> show igmp-snooping vlans vlan v10 detail
VLAN: v10, Tag: 10, vlan-interface: vlan.10
 Interface: ge-0/0/10.0, tagged, Groups: 0
IGMP-L2-Querier: Stopped, SourceAddress: 10.10.1.2

```



## Operational Commands (PIM)

- clear multicast bandwidth-admission
- clear multicast scope
- clear multicast sessions
- clear multicast statistics
- clear pim join
- clear pim register
- clear pim statistics
- mtrace
- mtrace from-source
- mtrace monitor
- mtrace to-gateway
- show multicast flow-map
- show multicast interface
- show multicast mrinfo
- show multicast next-hops
- show multicast pim-to-igmp-proxy
- show multicast pim-to-mld-proxy
- show multicast route
- show multicast rpf
- show multicast scope
- show multicast sessions
- show multicast usage
- show pim bootstrap
- show pim interfaces
- show pim join
- show pim neighbors
- show pim rps

- `show pim source`
- `show pim statistics`

## clear multicast bandwidth-admission

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>clear multicast bandwidth-admission &lt;group <i>group-address</i>&gt; &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt; &lt;source <i>source-address</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 8.3.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Reapply IP multicast bandwidth admissions.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <p><b>none</b>—Reapply multicast bandwidth admissions for all IPv4 forwarding entries in the master routing instance.</p> <p><b>group <i>group-address</i></b>—(Optional) Reapply multicast bandwidth admissions for the specified group.</p> <p><b>inet</b>—(Optional) Reapply multicast bandwidth admission settings for IPv4 flows.</p> <p><b>inet6</b>—(Optional) Reapply multicast bandwidth admission settings for IPv6 flows.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Reapply multicast bandwidth admission settings for the specified instance. If you do not specify an instance, the command applies to the master routing instance.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Examines the corresponding outbound interface in the relevant entries and acts as follows:</p> <ul style="list-style-type: none"> <li>• If the interface is congested, and it was admitted previously, it is removed.</li> <li>• If the interface was rejected previously, the <b>clear multicast bandwidth-admission</b> command enables the interface to be admitted as long as enough bandwidth exists on the interface.</li> <li>• If you do not specify an interface, issuing the <b>clear multicast bandwidth-admission</b> command readmits any previously rejected interface for the relevant entries as long as enough bandwidth exists on the interface.</li> </ul> <p>To manually reject previously admitted outbound interfaces, you must specify the interface.</p> <p><b>source <i>source-address</i></b>—(Optional) Use with the <b>group</b> option to reapply multicast bandwidth admission settings for the specified (source, group) entry.</p> |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

**Related Documentation** • [show multicast interface on page 5852](#)

**List of Sample Output** [clear multicast bandwidth-admission on page 5828](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

[clear multicast bandwidth-admission](#)

```
user@host> clear multicast bandwidth-admission
```

## clear multicast scope

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-----------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 5829</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 5829</a>                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Syntax</b>                                       | <pre>clear multicast scope &lt;inet   inet6&gt; &lt;interface <i>interface-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>clear multicast scope &lt;inet   inet6&gt; &lt;interface <i>interface-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>                          | <p>Command introduced in Junos OS Release 7.6.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>inet6</b> option introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                   |
| <b>Description</b>                                  | Clear IP multicast scope statistics.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                                      | <p><b>none</b>—(Same as <b>logical-system all</b>) Clear multicast scope statistics.</p> <p><b>inet</b>—(Optional) Clear multicast scope statistics for IPv4 family addresses.</p> <p><b>inet6</b>—(Optional) Clear multicast scope statistics for IPv6 family addresses.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Clear multicast scope statistics on a specific interface.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>                        | <ul style="list-style-type: none"> <li>• <a href="#">show multicast scope on page 5873</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>List of Sample Output</b>                        | <a href="#">clear multicast scope on page 5829</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Output Fields</b>                                | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

## Sample Output

### clear multicast scope

```
user@host> clear multicast scope
```

## clear multicast sessions

---

|                                              |                                                                                                                                                                                                                                                                                                                                                                                                   |
|----------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| List of Syntax                               | <a href="#">Syntax on page 5830</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 5830</a>                                                                                                                                                                                                                                                                                  |
| Syntax                                       | <code>clear multicast sessions</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code><br><code>&lt;<i>regular-expression</i>&gt;</code>                                                                                                                                                                                                                                 |
| Syntax (EX Series Switch and the QFX Series) | <code>clear multicast sessions</code><br><code>&lt;<i>regular-expression</i>&gt;</code>                                                                                                                                                                                                                                                                                                           |
| Release Information                          | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                          |
| Description                                  | Clear IP multicast sessions.                                                                                                                                                                                                                                                                                                                                                                      |
| Options                                      | <code>none</code> —(Same as <code>logical-system all</code> ) Clear multicast sessions.<br><br><code>logical-system (all   <i>logical-system-name</i>)</code> —(Optional) Perform this operation on all logical systems or on a particular logical system.<br><br><code><i>regular-expression</i></code> —(Optional) Clear only multicast sessions that contain the specified regular expression. |
| Required Privilege Level                     | clear                                                                                                                                                                                                                                                                                                                                                                                             |
| Related Documentation                        | <ul style="list-style-type: none"><li>• <a href="#">show multicast sessions on page 5875</a></li></ul>                                                                                                                                                                                                                                                                                            |
| List of Sample Output                        | <a href="#">clear multicast sessions on page 5830</a>                                                                                                                                                                                                                                                                                                                                             |
| Output Fields                                | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                             |

## Sample Output

### clear multicast sessions

```
user@host> clear multicast sessions
```

## clear multicast statistics

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-----------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 5831</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 5831</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Syntax</b>                                       | clear multicast statistics<br><inet   inet6><br><instance <i>instance-name</i> ><br><interface <i>interface-name</i> ><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | clear multicast statistics<br><inet   inet6><br><instance <i>instance-name</i> ><br><interface <i>interface-name</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>                          | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>                                  | Clear IP multicast statistics.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                                      | <b>none</b> —Clear multicast statistics for all supported address families on all interfaces.<br><b>inet</b> —(Optional) Clear multicast statistics for IPv4 family addresses.<br><b>inet6</b> —(Optional) Clear multicast statistics for IPv6 family addresses.<br><b>instance <i>instance-name</i></b> —(Optional) Clear multicast statistics for the specified instance.<br><b>interface <i>interface-name</i></b> —(Optional) Clear multicast statistics on a specific interface.<br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>                     | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>                        | <ul style="list-style-type: none"> <li><a href="#">show multicast statistics</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>List of Sample Output</b>                        | <a href="#">clear multicast statistics on page 5831</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Output Fields</b>                                | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

## Sample Output

### clear multicast statistics

```
user@host> clear multicast statistics
```

## clear pim join

---

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-----------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 5832</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 5832</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Syntax</b>                                       | <pre>clear pim join &lt;group-address&gt; &lt;inet   inet6&gt; &lt;instance instance-name&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>clear pim join &lt;group-address&gt; &lt;inet   inet6&gt; &lt;instance instance-name&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>                          | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>                                  | Clear the Protocol Independent Multicast (PIM) join and prune states.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                                      | <p><b>none</b>—Clear the PIM join and prune states for all groups, family addresses, and instances.</p> <p><b>group-address</b>—(Optional) Clear the PIM join and prune states for a group address.</p> <p><b>inet   inet6</b>—(Optional) Clear the PIM join and prune states for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance instance-name</b>—(Optional) Clear the join and prune states for a specific PIM-enabled routing instance.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Additional Information</b>                       | The <b>clear pim join</b> command cannot be used to clear the PIM join and prune state on a backup Routing Engine when nonstop active routing is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b>                     | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>                        | <ul style="list-style-type: none"><li>• <a href="#">show pim join on page 5886</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>List of Sample Output</b>                        | <a href="#">clear pim join on page 5833</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Output Fields</b>                                | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |



## Sample Output

clear pim join

```
user@host> clear pim join
```

## clear pim register

---

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-----------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 5834</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 5834</a><br><a href="#">Syntax (PTX Series) on page 5834</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Syntax</b>                                       | <pre>clear pim register &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>clear pim register &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Syntax (PTX Series)</b>                          | <pre>clear pim register &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>                          | Command introduced in Junos OS Release 7.6.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>                                  | Clear Protocol Independent Multicast (PIM) register message counters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                                      | <p><b>none</b>—Clear PIM register message counters for all family addresses, instances, and interfaces.</p> <p><b>inet   inet6</b>—(Optional) Clear PIM register message counters for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Clear register message counters for a specific PIM-enabled routing instance.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Clear PIM register message counters for a specific interface.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Additional Information</b>                       | The <b>clear pim register</b> command cannot be used to clear the PIM register state on a backup Routing Engine when nonstop active routing is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b>                     | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

**Related Documentation** • [show pim statistics on page 5921](#)

**List of Sample Output** [clear pim register on page 5835](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

[clear pim register](#)

```
user@host> clear pim register
```

## clear pim statistics

---

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-----------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 5836</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 5836</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Syntax</b>                                       | <pre>clear pim statistics &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>clear pim statistics &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>                          | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>                                  | Clear Protocol Independent Multicast (PIM) statistics.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                                      | <p><b>none</b>—Clear PIM statistics for all family addresses, instances, and interfaces.</p> <p><b>inet   inet6</b>—(Optional) Clear PIM statistics for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Clear statistics for a specific PIM-enabled routing instance.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Clear PIM statistics for a specific interface.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Additional Information</b>                       | The <b>clear pim statistics</b> command cannot be used to clear the PIM statistics on a backup Routing Engine when nonstop active routing is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b>                     | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>                        | <ul style="list-style-type: none"><li>• <a href="#">show pim statistics on page 5921</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>List of Sample Output</b>                        | <a href="#">clear pim statistics on page 5837</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Output Fields</b>                                | See <a href="#">show pim statistics</a> for an explanation of output fields.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## Sample Output

### clear pim statistics

The following sample output displays PIM statistics before and after the **clear pim statistics** command is entered:

```
user@host> show pim statistics
PIM statistics on all interfaces:
PIM Message type Received Sent Rx errors
Hello 0 0 0
Register 0 0 0
Register Stop 0 0 0
Join Prune 0 0 0
Bootstrap 0 0 0
Assert 0 0 0
Graft 0 0 0
Graft Ack 0 0 0
Candidate RP 0 0 0
V1 Query 2111 4222 0
V1 Register 0 0 0
V1 Register Stop 0 0 0
V1 Join Prune 14200 13115 0
V1 RP Reachability 0 0 0
V1 Assert 0 0 0
V1 Graft 0 0 0
V1 Graft Ack 0 0 0
PIM statistics summary for all interfaces:
Unknown type 0
V1 Unknown type 0
Unknown Version 0
Neighbor unknown 0
Bad Length 0
Bad Checksum 0
Bad Receive If 0
Rx Intf disabled 2007
Rx V1 Require V2 0
Rx Register not RP 0
RP Filtered Source 0
Unknown Reg Stop 0
Rx Join/Prune no state 1040
Rx Graft/Graft Ack no state 0
...
```

```
user@host> clear pim statistics
user@host> show pim statistics
PIM statistics on all interfaces:
PIM Message type Received Sent Rx errors
Hello 0 0 0
Register 0 0 0
Register Stop 0 0 0
Join Prune 0 0 0
Bootstrap 0 0 0
Assert 0 0 0
Graft 0 0 0
Graft Ack 0 0 0
Candidate RP 0 0 0
V1 Query 1 0 0
V1 Register 0 0 0
...
```



## mtrace

|                                 |                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>mtrace source</code><br><logical-system <i>logical-system-name</i> ><br><routing-instance <i>routing-instance-name</i> >                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 9.5 for SRX1400, SRX3400, SRX3600, SRX5600, and SRX5800 devices.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.<br>Command introduced in Junos OS Release 12.3 for the PTX Series. |
| <b>Description</b>              | Display trace information about an IP multicast path.                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <i>source</i> —Source hostname or address.<br><br><i>logical-system (logical-system-name)</i> —(Optional) Perform this operation on a logical system.<br><br><i>routing-instance routing-instance-name</i> —(Optional) Trace a particular routing instance.                                                                                                            |
| <b>Additional Information</b>   | The <b>mtrace</b> command for multicast traffic is similar to the <b>traceroute</b> command used for unicast traffic. Unlike <b>traceroute</b> , <b>mtrace</b> traces traffic backwards, from the receiver to the source.                                                                                                                                              |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                   |
| <b>List of Sample Output</b>    | <a href="#">mtrace source on page 5841</a>                                                                                                                                                                                                                                                                                                                             |
| <b>Output Fields</b>            | <a href="#">Table 447</a> describes the output fields for the <b>mtrace</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                           |

**Table 447: mtrace Output Fields**

| Field Name                        | Field Description                                             |
|-----------------------------------|---------------------------------------------------------------|
| <b>Mtrace from</b>                | IP address of the receiver.                                   |
| <b>to</b>                         | IP address of the source.                                     |
| <b>via group</b>                  | IP address of the multicast group (if any).                   |
| <b>Querying full reverse path</b> | Indicates the full reverse path query has begun.              |
| <i>number-of-hops</i>             | Number of hops from the source to the named router or switch. |
| <i>router-name</i>                | Name of the router or switch for this hop.                    |
| <i>address</i>                    | Address of the router or switch for this hop.                 |

Table 447: mtrace Output Fields (*continued*)

| Field Name      | Field Description                              |
|-----------------|------------------------------------------------|
| <i>protocol</i> | Protocol used (for example, PIM).              |
| Round trip time | Average round-trip time, in milliseconds (ms). |
| total ttl of    | Time-to-live (TTL) threshold.                  |



## Sample Output

### mtrace source

```
user@host> mtrace 192.1.4.2
Mtrace from 192.1.4.2 to 192.1.1.2 via group 0.0.0.0
Querying full reverse path... * *
 0 routerA.lab.mycompany.net (192.1.1.2)
-1 routerB.lab.mycompany.net (192.1.2.2) PIM thresh^ 1
-2 routerC.lab.mycompany.net (192.1.3.2) PIM thresh^ 1
-3 hostA.lab.mycompany.net (192.1.4.2)
Round trip time 2 ms; total ttl of 2 required.
```

## mtrace from-source

---

**Syntax** `mtrace from-source source source`  
`<brief | detail>`  
`<extra-hops extra-hops>`  
`<group group>`  
`<interval interval>`  
`<loop>`  
`<max-hops max-hops>`  
`<max-queries max-queries>`  
`<multicast-response | unicast-response>`  
`<no-resolve>`  
`<no-router-alert>`  
`<response response>`  
`<routing-instance routing-instance-name>`  
`<ttl ttl>`  
`<wait-time wait-time>`

**Release Information** Command introduced before Junos OS Release 7.4.  
Command introduced in Junos OS Release 9.0 for EX Series switches.  
Command introduced in Junos OS Release 11.3 for the QFX Series.

**Description** Display trace information about an IP multicast path from a source to this router or switch. If you specify a group address with this command, Junos OS returns additional information, such as packet rates and losses.

**Options** `brief | detail`—(Optional) Display the specified level of output.

`extra-hops extra-hops`—(Optional) Number of hops to take after reaching a nonresponsive router. You can specify a number between **0** and **255**.

`group group`—(Optional) Group address for which to trace the path. The default group address is **0.0.0.0**.

`interval interval`—(Optional) Number of seconds to wait before gathering statistics again. The default value is **10** seconds.

`loop`—(Optional) Loop indefinitely, displaying rate and loss statistics.

`max-hops max-hops`—(Optional) Maximum hops to trace toward the source. The range of values is **0** through **255**. The default value is **32** hops.

`max-queries max-queries`—(Optional) Maximum number of query attempts for any hop. The range of values is **1** through **32**. The default is **3**.

`multicast-response`—(Optional) Always request the response using multicast.

`no-resolve`—(Optional) Do not attempt to display addresses symbolically.

`no-router-alert`—(Optional) Do not use the router-alert IP option.

`response response`—(Optional) Send trace response to a host or multicast address.

**routing-instance** *routing-instance-name*—(Optional) Trace a particular routing instance.

**source** *source*—Source hostname or address.

**ttl** *tll*—(Optional) IP time-to-live (TTL) value. You can specify a number between 0 and 255. Local queries to the multicast group use a value of 1. Otherwise, the default value is 127.

**unicast-response**—(Optional) Always request the response using unicast.

**wait-time** *wait-time*—(Optional) Number of seconds to wait for a response. The default value is 3.

**Required Privilege Level** view

**List of Sample Output** [mtrace from-source on page 5844](#)

**Output Fields** Table 448 describes the output fields for the **mtrace from-source** command. Output fields are listed in the approximate order in which they appear.

**Table 448: mtrace from-source Output Fields**

| Field Name                        | Field Description                                             |
|-----------------------------------|---------------------------------------------------------------|
| <b>Mtrace from</b>                | IP address of the receiver.                                   |
| <b>to</b>                         | IP address of the source.                                     |
| <b>via group</b>                  | IP address of the multicast group (if any).                   |
| <b>Querying full reverse path</b> | Indicates the full reverse path query has begun.              |
| <b>number-of-hops</b>             | Number of hops from the source to the named router or switch. |
| <b>router-name</b>                | Name of the router or switch for this hop.                    |
| <b>address</b>                    | Address of the router or switch for this hop.                 |
| <b>protocol</b>                   | Protocol used (for example, PIM).                             |
| <b>Round trip time</b>            | Average round-trip time, in milliseconds (ms).                |
| <b>total ttl of</b>               | Time-to-live (TTL) threshold.                                 |
| <b>source</b>                     | Source address.                                               |
| <b>Response Dest</b>              | Response destination address.                                 |
| <b>Overall</b>                    | Average packet rate for all traffic at each hop.              |

Table 448: mtrace from-source Output Fields (*continued*)

| Field Name                                | Field Description                                                                                                |
|-------------------------------------------|------------------------------------------------------------------------------------------------------------------|
| <b>Packet Statistics for Traffic From</b> | Number of packets lost, number of packets sent, percentage of packets lost, and average packet rate at each hop. |
| <b>Receiver</b>                           | IP address receiving the multicast.                                                                              |
| <b>Query source</b>                       | IP address sending the mtrace query.                                                                             |

## Sample Output

### mtrace from-source

```

user@host> mtrace from-source source 192.1.4.2 group 225.1.1.1
Mtrace from 192.1.4.2 to 192.1.1.2 via group 225.1.1.1
Querying full reverse path... * *
 0 routerA.lab.mycompany.net (192.1.1.2)
-1 routerB.lab.mycompany.net (192.1.2.2) PIM thresh^ 1
-2 routerC.lab.mycompany.net (192.1.3.2) PIM thresh^ 1
-3 hostA.lab.mycompany.net (192.1.4.2)
Round trip time 2 ms; total ttl of 2 required.

Waiting to accumulate statistics...Results after 10 seconds:

Source Response Dest Overall Packet Statistics For Traffic From
192.1.4.2 192.1.1.2 Packet 192.1.4.2 To 225.1.1.1
 v ___/ rtt 2 ms Rate Lost/Sent = Pct Rate
192.1.2.1
192.1.3.2 routerC.lab.mycompany.net
 v ^ ttl 2 0/0 = -- 0 pps
192.1.4.1
192.1.2.2 routerB.lab.mycompany.net
 v __ ttl 3 ?/0 0 pps
192.1.1.2 192.1.1.2
Receiver Query Source

```

## mtrace monitor

|                                 |                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | mtrace monitor                                                                                                                                                                           |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series. |
| <b>Description</b>              | Listen passively for IP multicast responses. To exit the <b>mtrace monitor</b> command, type Ctrl+c.                                                                                     |
| <b>Options</b>                  | <b>none</b> —Trace the master instance.                                                                                                                                                  |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                     |
| <b>List of Sample Output</b>    | <a href="#">mtrace monitor on page 5846</a>                                                                                                                                              |
| <b>Output Fields</b>            | <a href="#">Table 449</a> describes the output fields for the <b>mtrace monitor</b> command. Output fields are listed in the approximate order in which they appear.                     |

**Table 449: mtrace monitor Output Fields**

| Field Name              | Field Description                                             |
|-------------------------|---------------------------------------------------------------|
| <b>Mtrace query at</b>  | Date and time of the query.                                   |
| <b>by</b>               | Address of the host issuing the query.                        |
| <b>resp to</b>          | Response destination.                                         |
| <b>qid</b>              | Query ID number.                                              |
| <b>packet from...to</b> | IP address of the query source and default group destination. |
| <b>from...to</b>        | IP address of the multicast source and the response address.  |
| <b>via group</b>        | IP address of the group to trace.                             |
| <b>mxhop</b>            | Maximum hop setting.                                          |

## Sample Output

### mtrace monitor

```
user@host> mtrace monitor
Mtrace query at Oct 22 13:36:14 by 192.1.3.2, resp to 224.0.1.32, qid 74a5b8
packet from 192.1.3.2 to 224.0.0.2
from 192.1.3.2 to 192.1.3.38 via group 224.1.1.1 (mxhop=60)

Mtrace query at Oct 22 13:36:17 by 192.1.3.2, resp to 224.0.1.32, qid 1d07ba
packet from 192.1.3.2 to 224.0.0.2
from 192.1.3.2 to 192.1.3.38 via group 224.1.1.1 (mxhop=60)

Mtrace query at Oct 22 13:36:20 by 192.1.3.2, resp to same, qid 2fea1d
packet from 192.1.3.2 to 224.0.0.2
from 192.1.3.2 to 192.1.3.38 via group 224.1.1.1 (mxhop=60)

Mtrace query at Oct 22 13:36:30 by 192.1.3.2, resp to same, qid 7c88ad
packet from 192.1.3.2 to 224.0.0.2
from 192.1.3.2 to 192.1.3.38 via group 224.1.1.1 (mxhop=60)
```

## mtrace to-gateway

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre> mtrace to-gateway gateway gateway &lt;brief   detail&gt; &lt;extra-hops extra-hops&gt; &lt;group group&gt; &lt;interface interface-name&gt; &lt;interval interval&gt; &lt;loop&gt; &lt;max-hops max-hops&gt; &lt;max-queries max-queries&gt; &lt;multicast-response   unicast-response&gt; &lt;no-resolve&gt; &lt;no-router-alert&gt; &lt;response response&gt; &lt;routing-instance routing-instance-name&gt; &lt;tll ttl&gt; &lt;unicast-response&gt; &lt;wait-time wait-time&gt; </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b> | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>         | Display trace information about a multicast path from this router or switch to a gateway router or switch.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>             | <p><b>gateway gateway</b>—Send the trace query to a gateway multicast address.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>extra-hops extra-hops</b>—(Optional) Number of hops to take after reaching a nonresponsive router or switch. You can specify a number between <b>0</b> and <b>255</b>.</p> <p><b>group group</b>—(Optional) Group address for which to trace the path. The default group address is <b>0.0.0.0</b>.</p> <p><b>interface interface-name</b>—(Optional) Source address for sending the trace query.</p> <p><b>interval interval</b>—(Optional) Number of seconds to wait before gathering statistics again. The default value is <b>10</b>.</p> <p><b>loop</b>—(Optional) Loop indefinitely, displaying rate and loss statistics.</p> <p><b>max-hops max-hops</b>—(Optional) Maximum hops to trace toward the source. You can specify a number between <b>0</b> and <b>255</b>. The default value is <b>32</b>.</p> <p><b>max-queries max-queries</b>—(Optional) Maximum number of query attempts for any hop. You can specify a number between <b>0</b> and <b>255</b>. The default value is <b>3</b>.</p> <p><b>multicast-response</b>—(Optional) Always request the response using multicast.</p> <p><b>no-resolve</b>—(Optional) Do not attempt to display addresses symbolically.</p> |

**no-router-alert**—(Optional) Do not use the router-alert IP option.

**response *response***—(Optional) Send trace response to a host or multicast address.

**routing-instance *routing-instance-name***—(Optional) Trace a particular routing instance.

**ttl *tll***—(Optional) IP time-to-live value. You can specify a number between 0 and 225.

Local queries to the multicast group use TTL 1. Otherwise, the default value is 127.

**unicast-response**—(Optional) Always request the response using unicast.

**wait-time *wait-time***—(Optional) Number of seconds to wait for a response. The default value is 3.

Required Privilege  
Level

view

List of Sample Output [mtrace to-gateway on page 5848](#)

Output Fields [Table 450](#) describes the output fields for the **mtrace to-gateway** command. Output fields are listed in the approximate order in which they appear.

**Table 450: mtrace to-gateway Output Fields**

| Field Name                        | Field Description                                             |
|-----------------------------------|---------------------------------------------------------------|
| <b>Mtrace from</b>                | IP address of the receiver.                                   |
| <b>to</b>                         | IP address of the source.                                     |
| <b>via group</b>                  | IP address of the multicast group (if any).                   |
| <b>Querying full reverse path</b> | Indicates the full reverse path query has begun.              |
| <b><i>number-of-hops</i></b>      | Number of hops from the source to the named router or switch. |
| <b><i>router-name</i></b>         | Name of the router or switch for this hop.                    |
| <b><i>address</i></b>             | Address of the router or switch for this hop.                 |
| <b><i>protocol</i></b>            | Protocol used (for example, PIM).                             |
| <b>Round trip time</b>            | Average round-trip time, in milliseconds (ms).                |
| <b>total ttl of</b>               | Time-to-live (TTL) threshold.                                 |

## Sample Output

### mtrace to-gateway

```
user@host> mtrace to-gateway gateway 192.1.3.2 group 225.1.1.1 interface 192.1.1.73 brief
```

```
Mtrace from 192.1.1.73 to 192.1.1.2 via group 225.1.1.1
```



```
Querying full reverse path... * *
 0 routerA.lab.mycompany.net (192.1.1.2)
-1 routerA.lab.mycompany.net (192.1.1.2) PIM thresh^ 1
-2 routerB.lab.mycompany.net (192.1.2.2) PIM thresh^ 1
-3 routerC.lab.mycompany.net (192.1.3.2) PIM thresh^ 1
Round trip time 2 ms; total ttl of 3 required.
```

## show multicast flow-map

|                                                     |                                                                                                                                                                                                                                                                                                                                              |
|-----------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 5850</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 5850</a>                                                                                                                                                                                                                             |
| <b>Syntax</b>                                       | <pre>show multicast flow-map &lt;brief   detail&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                          |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show multicast flow-map &lt;brief   detail&gt;</pre>                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>                          | <p>Command introduced in Junos OS Release 8.2.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                          |
| <b>Description</b>                                  | Display configuration information about IP multicast flow maps.                                                                                                                                                                                                                                                                              |
| <b>Options</b>                                      | <p><b>none</b>—Display configuration information about IP multicast flow maps on all systems.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                         |
| <b>List of Sample Output</b>                        | <a href="#">show multicast flow-map on page 5851</a><br><a href="#">show multicast flow-map detail on page 5851</a>                                                                                                                                                                                                                          |
| <b>Output Fields</b>                                | Table 451 describes the output fields for the <b>show multicast flow-map</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                |

Table 451: show multicast flow-map Output Fields

| Field Name           | Field Description                                         | Levels of Output |
|----------------------|-----------------------------------------------------------|------------------|
| <b>Name</b>          | Name of the flow map.                                     | All levels       |
| <b>Policy</b>        | Name of the policy associated with the flow map.          | All levels       |
| <b>Cache-timeout</b> | Cache timeout value assigned to the flow map.             | All levels       |
| <b>Bandwidth</b>     | Bandwidth setting associated with the flow map.           | All levels       |
| <b>Adaptive</b>      | Whether or not adaptive mode is enabled for the flow map. | none             |
| <b>Flow-map</b>      | Name of the flow map.                                     | <b>detail</b>    |

Table 451: show multicast flow-map Output Fields (*continued*)

| Field Name                | Field Description                                         | Levels of Output |
|---------------------------|-----------------------------------------------------------|------------------|
| <b>Adaptive Bandwidth</b> | Whether or not adaptive mode is enabled for the flow map. | <b>detail</b>    |
| <b>Redundant Sources</b>  | Redundant sources defined for the same destination group. | <b>detail</b>    |

## Sample Output

### show multicast flow-map

```

user@host> show multicast flow-map
Instance: master
Name Policy Cache timeout Bandwidth Adaptive
map2 policy2 never 2000000 no
map1 policy1 60 seconds 2000000 no

```

## Sample Output

### show multicast flow-map detail

```

user@host> show multicast flow-map detail
Instance: master
Flow-map: map1
 Policy: policy1
 Cache Timeout: 600 seconds
 Bandwidth: 2000000
 Adaptive Bandwidth: yes
 Redundant Sources: 11.11.11.11
 Redundant Sources: 11.11.11.12
 Redundant Sources: 11.11.11.13

```

## show multicast interface

|                                                     |                                                                                                                                                                                                                                  |
|-----------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 5852</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 5852</a>                                                                                                                 |
| <b>Syntax</b>                                       | show multicast interface<br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                 |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | show multicast interface                                                                                                                                                                                                         |
| <b>Release Information</b>                          | Command introduced in Junos OS Release 8.3.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                             |
| <b>Description</b>                                  | Display bandwidth information about IP multicast interfaces.                                                                                                                                                                     |
| <b>Options</b>                                      | <b>none</b> —Display all interfaces that have multicast configured.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                             |
| <b>List of Sample Output</b>                        | <a href="#">show multicast interface on page 5853</a>                                                                                                                                                                            |
| <b>Output Fields</b>                                | <a href="#">Table 452</a> describes the output fields for the <b>show multicast interface</b> command. Output fields are listed in the approximate order in which they appear.                                                   |

**Table 452: show multicast interface Output Fields**

| Field Name                              | Field Description                                                                                                                                                                                                                                                                                                                                           |
|-----------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Interface</b>                        | Name of the multicast interface.                                                                                                                                                                                                                                                                                                                            |
| <b>Maximum bandwidth (bps)</b>          | Maximum bandwidth setting, in bits per second, for this interface.                                                                                                                                                                                                                                                                                          |
| <b>Remaining bandwidth (bps)</b>        | Amount of bandwidth, in bits per second, remaining on the interface.                                                                                                                                                                                                                                                                                        |
| <b>Mapped bandwidth deduction (bps)</b> | Amount of bandwidth, in bits per second, used by any flows that are mapped to the interface.<br><br><b>NOTE:</b> Adding the mapped bandwidth deduction value to the local bandwidth deduction value results in the total deduction value for the interface.<br><br>This field does not appear in the output when the no QoS adjustment feature is disabled. |

Table 452: show multicast interface Output Fields (*continued*)

| Field Name                                   | Field Description                                                                                                                                                                                                                                                                                                                                                          |
|----------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Local bandwidth deduction (bps)</b>       | <p>Amount of bandwidth, in bits per second, used by any mapped flows that are traversing the interface.</p> <p><b>NOTE:</b> Adding the mapped bandwidth deduction value to the local bandwidth deduction value results in the total deduction value for the interface.</p> <p>This field does not appear in the output when the no QoS adjustment feature is disabled.</p> |
| <b>Reverse OIF mapping</b>                   | <p>State of the reverse OIF mapping feature (<b>on</b> or <b>off</b>).</p> <p><b>NOTE:</b> This field does not appear in the output when the no QoS adjustment feature is disabled.</p>                                                                                                                                                                                    |
| <b>Reverse OIF mapping no QoS adjustment</b> | <p>State of the no QoS adjustment feature (<b>on</b> or <b>off</b>) for interfaces that are using reverse OIF mapping.</p> <p><b>NOTE:</b> This field does not appear in the output when the no QoS adjustment feature is disabled.</p>                                                                                                                                    |
| <b>Leave timer</b>                           | <p>Amount of time a mapped interface remains active after the last mapping ends.</p> <p><b>NOTE:</b> This field does not appear in the output when the no QoS adjustment feature is disabled.</p>                                                                                                                                                                          |
| <b>No QoS adjustment</b>                     | <p>State (<b>on</b>) of the no QoS adjustment feature when this feature is enabled.</p> <p><b>NOTE:</b> This field does not appear in the output when the no QoS adjustment feature is disabled.</p>                                                                                                                                                                       |

## Sample Output

### show multicast interface

```

user@host> show multicast interface
Interface Maximum bandwidth (bps) Remaining bandwidth (bps)
fe-0/0/3 10000000 0
fe-0/0/3.210 10000000 -2000000
fe-0/0/3.220 100000000 100000000
fe-0/0/3.230 20000000 18000000
fe-0/0/2.200 100000000 100000000

```

## show multicast mrinfo

|                                 |                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show multicast mrinfo</code><br><code>&lt;host&gt;</code>                                                                                                                                                           |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                  |
| <b>Description</b>              | Display configuration information about IP multicast networks, including neighboring multicast router addresses.                                                                                                          |
| <b>Options</b>                  | <b>none</b> —Display configuration information about all multicast networks.<br><br><b>host</b> —(Optional) Display configuration information about a particular host. Replace <i>host</i> with a hostname or IP address. |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                      |
| <b>List of Sample Output</b>    | <a href="#">show multicast mrinfo on page 5855</a>                                                                                                                                                                        |
| <b>Output Fields</b>            | <a href="#">Table 453</a> describes the output fields for the <b>show multicast mrinfo</b> command. Output fields are listed in the approximate order in which they appear.                                               |

**Table 453: show multicast mrinfo Output Fields**

| Field Name                           | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>source-address</i>                | Query address, hostname (DNS name or IP address of the source address), and multicast protocol version or the software version of another vendor.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <i>ip-address-1—&gt;ip-address-2</i> | Queried router interface address and directly attached neighbor interface address, respectively.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <i>(name or ip-address)</i>          | Name or IP address of neighbor.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <i>[metric/threshold/type/flags]</i> | Neighbor's multicast profile: <ul style="list-style-type: none"> <li><b>metric</b>—Always has a value of 1, because <b>mrinfo</b> queries the directly connected interfaces of a device.</li> <li><b>threshold</b>—Multicast threshold time-to-live (TTL). The range of values is 0 through 255.</li> <li><b>type</b>—Multicast connection type: <b>pim</b> or <b>tunnel</b>.</li> <li><b>flags</b>—Flags for this route: <ul style="list-style-type: none"> <li><b>querier</b>—Queried router is the designated router for the neighboring session.</li> <li><b>leaf</b>—Link is a leaf in the multicast network.</li> <li><b>down</b>—Link status indicator.</li> </ul> </li> </ul> |

## Sample Output

show multicast mrinfo

```
user@host> show multicast mrinfo 10.35.4.1
10.35.4.1 (10.35.4.1) [version 12.0]:
 192.168.195.166 -> 0.0.0.0 (local) [1/0/pim/querier/leaf]
 10.38.20.1 -> 0.0.0.0 (local) [1/0/pim/querier/leaf]
 10.47.1.1 -> 10.47.1.2 (10.47.1.2) [1/5/pim]
 0.0.0.0 -> 0.0.0.0 (local) [1/0/pim/down]
```

## show multicast next-hops

---

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-----------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 5856</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 5856</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Syntax</b>                                       | <pre>show multicast next-hops &lt;brief   detail&gt; &lt;identifier-number&gt; &lt;inet   inet6&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show multicast next-hops &lt;brief   detail&gt; &lt;identifier-number&gt; &lt;inet   inet6&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>                          | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>inet6</b> option introduced in Junos OS Release 10.0 for EX Series switches.</p> <p><b>detail</b> option display of next-hop ID number introduced in Junos OS Release 11.1 for M Series and T Series routers and EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional PIM added in Junos OS Release 12.1.</p>                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>                                  | Display the entries in the IP multicast next-hop table.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                                      | <p><b>none</b>—Display standard information about all entries in the multicast next-hop table for all supported address families.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p>When you include the <b>detail</b> option on M Series and T Series routers and EX Series switches, the downstream interface name includes the next-hop ID number in parentheses, in the form <b>fe-0/1/2.0-(1048574)</b> where <b>1048574</b> is the next-hop ID number.</p> <p><b>identifier-number</b>—(Optional) Show a particular next hop by ID number. The range of values is 1 through <b>65,535</b>.</p> <p><b>inet   inet6</b>—(Optional) Display entries for IPv4 or IPv6 family addresses, respectively.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>List of Sample Output</b>                        | <a href="#">show multicast next-hops on page 5857</a><br><a href="#">show multicast next-hops (Bidirectional PIM on page 5857</a><br><a href="#">show multicast next-hops brief on page 5858</a><br><a href="#">show multicast next-hops detail on page 5858</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |



**Output Fields** Table 454 describes the output fields for the **show multicast next-hops** command. Output fields are listed in the approximate order in which they appear.

**Table 454: show multicast next-hops Output Fields**

| Field Name                     | Field Description                                                                                                                                      |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Family</b>                  | Protocol family (such as <b>INET</b> ).                                                                                                                |
| <b>ID</b>                      | Next-hop identifier of the prefix. The identifier is returned by the routing device's Packet Forwarding Engine.                                        |
| <b>Refcount</b>                | Number of cache entries that are using this next hop.                                                                                                  |
| <b>KRefcount</b>               | Kernel reference count for the next hop.                                                                                                               |
| <b>Downstream interface</b>    | Interface names associated with each multicast next-hop ID.                                                                                            |
| <b>Incoming interface list</b> | List of interfaces that accept incoming traffic. Only shown for routes that do not use strict RPF-based forwarding, for example for bidirectional PIM. |

## Sample Output

### show multicast next-hops

```
user@host> show multicast next-hops
Family: INET
ID Refcount KRefcount Downstream interface
262142 4 2 so-1/0/0.0
262143 2 1 mt-1/1/0.49152
262148 2 1 mt-1/1/0.32769
```

### show multicast next-hops (Bidirectional PIM)

```
user@host> show multicast next-hops
Family: INET
ID Refcount KRefcount Downstream interface
2097151 8 4 ge-0/0/1.0

Family: INET6
ID Refcount KRefcount Downstream interface
2097157 2 1 ge-0/0/1.0

Family: Incoming interface list
ID Refcount KRefcount Downstream interface
513 5 2 lo0.0
 ge-0/0/1.0
514 5 2 lo0.0
 ge-0/0/1.0
 xe-4/1/0.0
515 3 1 lo0.0
 ge-0/0/1.0
 xe-4/1/0.0
544 1 0 lo0.0
 xe-4/1/0.0
```

### show multicast next-hops brief

The output for the **show multicast next-hops brief** command is identical to that for the **show multicast next-hops** command. For sample output, see [show multicast next-hops on page 5857](#).

### show multicast next-hops detail

```
user@host> show multicast next-hops detail
Family: INET
ID Refcount KRefCount Downstream interface
1048577 2 1 fe-0/1/2.0-(1048574)
 ge-0/2/3.0-(1048576)
```

## show multicast pim-to-igmp-proxy

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 5859</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 5859</a>                                                                                                                                                                                                                                                                                                                                    |
| <b>Syntax</b>                                       | <pre>show multicast pim-to-igmp-proxy &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                         |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show multicast pim-to-igmp-proxy &lt;instance <i>instance-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>                          | <p>Command introduced in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 9.6 for EX Series switches.</p> <p><b>instance</b> option introduced in Junos OS Release 10.0.</p> <p><b>instance</b> option introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                    |
| <b>Description</b>                                  | Display configuration information about PIM-to-IGMP message translation, also known as PIM-to-IGMP proxy.                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                                      | <p><b>none</b>—Display configuration information about PIM-to-IGMP message translation for all routing instances.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display configuration information about PIM-to-IGMP message translation for a specific multicast instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>List of Sample Output</b>                        | <a href="#">show multicast pim-to-igmp-proxy on page 5860</a><br><a href="#">show multicast pim-to-igmp-proxy instance on page 5860</a>                                                                                                                                                                                                                                                                                                             |
| <b>Output Fields</b>                                | <p><a href="#">Table 455</a> describes the output fields for the <b>show multicast pim-to-igmp-proxy</b> command. Output fields are listed in the order in which they appear.</p>                                                                                                                                                                                                                                                                   |

**Table 455: show multicast pim-to-igmp-proxy Output Fields**

| Field Name                   | Field Description                                                                                                                                     |
|------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Instance</b>              | Routing instance. Default instance is <b>master</b> (inet.0 routing table).                                                                           |
| <b>Proxy state</b>           | State of PIM-to-IGMP message translation, also known as PIM-to-IGMP proxy, on the configured upstream interfaces: <b>enabled</b> or <b>disabled</b> . |
| <b><i>interface-name</i></b> | Name of upstream interface (no more than two allowed) on which PIM-to-IGMP message translation is configured.                                         |

## Sample Output

### show multicast pim-to-igmp-proxy

```
user@host> show multicast pim-to-igmp-proxy
Instance: master Proxy state: enabled
ge-0/1/0.1
ge-0/1/0.2
```

### show multicast pim-to-igmp-proxy instance

```
user@host> show multicast pim-to-igmp-proxy instance VPN-A
Instance: VPN-A Proxy state: enabled
ge-0/1/0.1
```

## show multicast pim-to-mld-proxy

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-----------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 5861</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 5861</a>                                                                                                                                                                                                                                                                                                                                  |
| <b>Syntax</b>                                       | <pre>show multicast pim-to-mld-proxy &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                        |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show multicast pim-to-mld-proxy &lt;instance <i>instance-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>                          | <p>Command introduced in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 9.6 for EX Series switches.</p> <p><b>instance</b> option introduced in Junos OS Release 10.3.</p> <p><b>instance</b> option introduced in Junos OS Release 10.3 for EX Series switches.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>    |
| <b>Description</b>                                  | Display configuration information about PIM-to-MLD message translation, also known as PIM-to-MLD proxy.                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                                      | <p><b>none</b>—Display configuration information about PIM-to-MLD message translation for all routing instances.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display configuration information about PIM-to-MLD message translation for a specific multicast instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>List of Sample Output</b>                        | <a href="#">show multicast pim-to-mld-proxy on page 5862</a><br><a href="#">show multicast pim-to-mld-proxy instance on page 5862</a>                                                                                                                                                                                                                                                                                                             |
| <b>Output Fields</b>                                | <p><a href="#">Table 456</a> describes the output fields for the <b>show multicast pim-to-mld-proxy</b> command. Output fields are listed in the order in which they appear.</p>                                                                                                                                                                                                                                                                  |

**Table 456: show multicast pim-to-mld-proxy Output Fields**

| Field Name                   | Field Description                                                                                                                                   |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Proxy state</b>           | State of PIM-to-MLD message translation, also known as PIM-to-MLD proxy, on the configured upstream interfaces: <b>enabled</b> or <b>disabled</b> . |
| <b><i>interface-name</i></b> | Name of upstream interface (no more than two allowed) on which PIM-to-MLD message translation is configured.                                        |

## Sample Output

### show multicast pim-to-mld-proxy

```
user@host> show multicast pim-to-mld-proxy
Instance: master Proxy state: enabled
ge-0/5/0.1
ge-0/5/0.2
```

### show multicast pim-to-mld-proxy instance

```
user@host> show multicast pim-to-mld-proxy instance VPN-A
Instance: VPN-A Proxy state: enabled
ge-0/5/0.1
```

## show multicast route

**List of Syntax**    [Syntax on page 5863](#)  
                          [Syntax \(EX Series Switch and the QFX Series\) on page 5863](#)

**Syntax**    show multicast route  
                  <brief | detail | extensive | summary>  
                  <active | all | inactive>  
                  <group *group*>  
                  <inet | inet6>  
                  <instance *instance name*>  
                  <logical-system (all | *logical-system-name*)>  
                  <*regular-expression*>  
                  <source-prefix *source-prefix*>

**Syntax (EX Series Switch and the QFX Series)**    show multicast route  
                  <brief | detail | extensive | summary>  
                  <active | all | inactive>  
                  <group *group*>  
                  <inet | inet6>  
                  <instance *instance name*>  
                  <*regular-expression*>  
                  <source-prefix *source-prefix*>

**Release Information**    Command introduced before Junos OS Release 7.4.  
                                  Command introduced in Junos OS Release 9.0 for EX Series switches.  
                                  inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.  
                                  Command introduced in Junos OS Release 11.3 for the QFX Series.  
                                  Support for bidirectional PIM added in Junos OS Release 12.1.  
                                  Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description**    Display the entries in the IP multicast forwarding table. You can display similar information with the **show route table inet.1** command.

**Options**    **none**—Display standard information about all entries in the multicast forwarding table for all routing instances.

**brief | detail | extensive | summary**—(Optional) Display the specified level of output.

**active | all | inactive**—(Optional) Display all active entries, all entries, or all inactive entries, respectively, in the multicast forwarding table.

**group *group***—(Optional) Display the cache entries for a particular group.

**inet | inet6**—(Optional) Display multicast forwarding table entries for IPv4 or IPv6 family addresses, respectively.

**instance *instance-name***—(Optional) Display entries in the multicast forwarding table for a specific multicast instance.

**logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

**regular-expression**—(Optional) Display information about the multicast forwarding table entries that match a UNIX OS-style regular expression.

**source-prefix source-prefix**—(Optional) Display the cache entries for a particular source prefix.

**Required Privilege Level** view

**Related Documentation** • [Example: Configuring Bidirectional PIM](#)

**List of Sample Output** [show multicast route on page 5865](#)  
[show multicast route \(Bidirectional PIM\) on page 5866](#)  
[show multicast route brief on page 5866](#)  
[show multicast route detail on page 5867](#)  
[show multicast route extensive \(Bidirectional PIM\) on page 5867](#)  
[show multicast route instance <instance-name> extensive on page 5868](#)  
[show multicast route summary on page 5868](#)

**Output Fields** [Table 457](#) describes the output fields for the **show multicast route** command. Output fields are listed in the approximate order in which they appear.

**Table 457: show multicast route Output Fields**

| Field Name                    | Field Description                                                                                                                                                                               | Level of Output         |
|-------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| family                        | IPv4 address family ( <b>INET</b> ) or IPv6 address family ( <b>INET6</b> ).                                                                                                                    | All levels              |
| Group                         | Group address.<br><br>For any-source multicast routes, for example for bidirectional PIM, the group address includes the prefix length.                                                         | All levels              |
| Source                        | Prefix and length of the source as it is in the multicast forwarding table.                                                                                                                     | All levels              |
| Incoming interface list       | List of interfaces that accept incoming traffic. Only shown for routes that do not use strict RPF-based forwarding, for example for bidirectional PIM.                                          | All levels              |
| Upstream interface            | Name of the interface on which the packet with this source prefix is expected to arrive.                                                                                                        | All levels              |
| Upstream rpf interface list   | When multicast-only fast reroute (MoFRR) is enabled, a PIM router propagates join messages on two upstream RPF interfaces to receive multicast traffic on both links for the same join request. | All levels              |
| Downstream interface list     | List of interface names to which the packet with this source prefix is forwarded.                                                                                                               | All levels              |
| Number of outgoing interfaces | Total number of outgoing interfaces for each (S,G) entry.                                                                                                                                       | <b>extensive</b>        |
| Session description           | Name of the multicast session.                                                                                                                                                                  | <b>detail extensive</b> |



Table 457: show multicast route Output Fields (*continued*)

| Field Name                             | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                             | Level of Output   |
|----------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------|
| Statistics                             | Rate at which packets are being forwarded for this source and group entry (in Kbps and pps), and number of packets that have been forwarded to this prefix. If one or more of the kilobits per second packet forwarding statistic queries fails or times out, the statistics field displays <b>Forwarding statistics are not available</b> .<br><br><b>NOTE:</b> On QFX Series switches and OCX Series switches, this field does not report valid statistics. | detail extensive  |
| Next-hop ID                            | Next-hop identifier of the prefix. The identifier is returned by the routing device's Packet Forwarding Engine and is also displayed in the output of the <b>show multicast nexthops</b> command.                                                                                                                                                                                                                                                             | detail extensive  |
| Incoming interface list ID             | For bidirectional PIM, incoming interface list identifier.<br><br>Identifiers for interfaces that accept incoming traffic. Only shown for routes that do not use strict RPF-based forwarding, for example for bidirectional PIM.                                                                                                                                                                                                                              | detail extensive  |
| Upstream protocol                      | The protocol that maintains the active multicast forwarding route for this group or source.<br><br>When the <b>show multicast route extensive</b> command is used with the <b>display-origin-protocol</b> option, the field name is only <b>Protocol</b> and not <b>Upstream Protocol</b> . However, this field also displays the protocol that installed the active route.                                                                                   | detail extensive  |
| Route type                             | Type of multicast route. Values can be (S,G) or (*G).                                                                                                                                                                                                                                                                                                                                                                                                         | summary           |
| Route state                            | Whether the group is <b>Active</b> or <b>Inactive</b> .                                                                                                                                                                                                                                                                                                                                                                                                       | summary extensive |
| Route count                            | Number of multicast routes.                                                                                                                                                                                                                                                                                                                                                                                                                                   | summary           |
| Forwarding state                       | Whether the prefix is pruned or forwarding.                                                                                                                                                                                                                                                                                                                                                                                                                   | extensive         |
| Cache lifetime/timeout                 | Number of seconds until the prefix is removed from the multicast forwarding table. A value of <b>never</b> indicates a permanent forwarding entry. A value of <b>forever</b> indicates routes that do not have keepalive times.                                                                                                                                                                                                                               | extensive         |
| Wrong incoming interface notifications | Number of times that the upstream interface was not available.                                                                                                                                                                                                                                                                                                                                                                                                | extensive         |
| Uptime                                 | Time since the creation of a multicast route.                                                                                                                                                                                                                                                                                                                                                                                                                 | extensive         |

## Sample Output

### show multicast route

```

user@host> show multicast route
Family: INET

Group: 228.0.0.0

```

```
Source: 10.255.14.144/32
Upstream interface: local
Downstream interface list:
 so-1/0/0.0

Group: 239.1.1.1
Source: 10.255.14.144/32
Upstream interface: local
Downstream interface list:
 so-1/0/0.0

Group: 239.1.1.1
Source: 10.255.70.15/32
Upstream interface: so-1/0/0.0
Downstream interface list:
 mt-1/1/0.1081344

Family: INET6
```

### show multicast route (Bidirectional PIM)

```
user@host> show multicast route
Family: INET

Group: 224.1.1.0/24
Source: *
Incoming interface list:
 lo0.0 ge-0/0/1.0
Downstream interface list:
 ge-0/0/1.0

Group: 224.1.3.0/24
Source: *
Incoming interface list:
 lo0.0 ge-0/0/1.0 xe-4/1/0.0
Downstream interface list:
 ge-0/0/1.0

Group: 225.1.1.0/24
Source: *
Incoming interface list:
 lo0.0 ge-0/0/1.0
Downstream interface list:
 ge-0/0/1.0

Group: 225.1.3.0/24
Source: *
Incoming interface list:
 lo0.0 ge-0/0/1.0 xe-4/1/0.0
Downstream interface list:
 ge-0/0/1.0
Family: INET6
```

### show multicast route brief

The output for the **show multicast route brief** command is identical to that for the **show multicast route** command. For sample output, see [show multicast route on page 5865](#) or [show multicast route \(Bidirectional PIM\) on page 5866](#).

**show multicast route detail**

```

user@host> show multicast route detail
Family: INET

Group: 228.0.0.0
 Source: 10.255.14.144/32
 Upstream interface: local
 Downstream interface list:
 so-1/0/0.0
 Session description: Unknown
 Statistics: 8 kbps, 100 pps, 45272 packets
 Next-hop ID: 262142
 Upstream protocol: PIM

Group: 239.1.1.1
 Source: 10.255.14.144/32
 Upstream interface: local
 Downstream interface list:
 so-1/0/0.0
 Session description: Administratively Scoped
 Statistics: 0 kbps, 0 pps, 13404 packets
 Next-hop ID: 262142
 Upstream protocol: PIM

Group: 239.1.1.1
 Source: 10.255.70.15/32
 Upstream interface: so-1/0/0.0
 Downstream interface list:
 mt-1/1/0.1081344
 Session description: Administratively Scoped
 Statistics: 46 kbps, 1000 pps, 921077 packets

 Next-hop ID: 262143
 Upstream protocol: PIM

Family: INET6

```

**show multicast route extensive (Bidirectional PIM)**

```

user@host> show multicast route extensive
Family: INET

Group: 224.1.1.0/24
 Source: *
 Incoming interface list:
 lo0.0 ge-0/0/1.0
 Downstream interface list:
 ge-0/0/1.0
 Number of outgoing interfaces: 1
 Session description: NOB Cross media facilities
 Statistics: 0 kbps, 0 pps, 0 packets
 Next-hop ID: 2097153
 Incoming interface list ID: 585
 Upstream protocol: PIM
 Route state: Active
 Forwarding state: Forwarding
 Cache lifetime/timeout: forever
 Wrong incoming interface notifications: 0

Group: 224.1.3.0/24

```

```
Source: *
Incoming interface list:
 lo0.0 ge-0/0/1.0 xe-4/1/0.0
Downstream interface list:
 ge-0/0/1.0
Number of outgoing interfaces: 1
Session description: NOB Cross media facilities
Statistics: 0 kbps, 0 pps, 0 packets
Next-hop ID: 2097153
Incoming interface list ID: 589
Upstream protocol: PIM
Route state: Active
Forwarding state: Forwarding
Cache lifetime/timeout: forever
Wrong incoming interface notifications: 0
```

Family: INET6

#### show multicast route instance <instance-name> extensive

```
user@host> show multicast route instance mvpn extensive
Family: INET

Group: 239.10.10.10
Source: 2.0.0.2/32
Upstream interface: xe-0/0/0.102
Downstream interface list:
 xe-10/3/0.0 xe-0/3/0.0 xe-0/0/0.106 xe-0/0/0.105
 xe-0/0/0.103 xe-0/0/0.104 xe-0/0/0.107 xe-0/0/0.108
Session description: Administratively Scoped
Statistics: 256 kbps, 3998 pps, 670150 packets
Next-hop ID: 1048579
Upstream protocol: MVPN
Route state: Active
Forwarding state: Forwarding
Cache lifetime/timeout: forever
Wrong incoming interface notifications: 58
Uptime: 00:00:04
```

#### show multicast route summary

```
user@host> show multicast route summary
Instance: master Family: INET

Route type Route state Route count
(S,G) Active 2
(S,G) Inactive 3

Instance: master Family: INET6
```

## show multicast rpf

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 5869</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 5869</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Syntax</b>                                       | <pre>show multicast rpf &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;prefix&gt; &lt;summary&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show multicast rpf &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;prefix&gt; &lt;summary&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>                          | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>                                  | Display information about multicast reverse-path-forwarding (RPF) calculations.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                                      | <p><b>none</b>—Display RPF calculation information for all supported address families.</p> <p><b>inet   inet6</b>—(Optional) Display the RPF calculation information for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information about multicast RPF calculations for a specific multicast instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>prefix</b>—(Optional) Display the RPF calculation information for the specified prefix.</p> <p><b>summary</b>—(Optional) Display a summary of all multicast RPF information.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>List of Sample Output</b>                        | <a href="#">show multicast rpf on page 5870</a><br><a href="#">show multicast rpf inet6 on page 5871</a><br><a href="#">show multicast rpf prefix on page 5872</a><br><a href="#">show multicast rpf summary on page 5872</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

**Output Fields** Table 458 describes the output fields for the **show multicast rpf** command. Output fields are listed in the approximate order in which they appear.

**Table 458: show multicast rpf Output Fields**

| Field Name           | Field Description                                                                                                                                                                                                                                                                                                                                                                                             |
|----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Instance</b>      | Name of the routing instance. (Displayed when multicast is configured within a routing instance.)                                                                                                                                                                                                                                                                                                             |
| <b>Source prefix</b> | Prefix and length of the source as it exists in the multicast forwarding table.                                                                                                                                                                                                                                                                                                                               |
| <b>Protocol</b>      | How the route was learned.                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Interface</b>     | Upstream RPF interface.<br><br><b>NOTE:</b> The displayed interface information does not apply to bidirectional PIM RP addresses. This is because the <b>show multicast rpf</b> command does not take into account equal-cost paths or the designated forwarder. For accurate upstream RPF interface information, always use the <b>show pim join extensive</b> command when bidirectional PIM is configured. |
| <b>Neighbor</b>      | Upstream RPF neighbor.<br><br><b>NOTE:</b> The displayed neighbor information does not apply to bidirectional PIM. This is because the <b>show multicast rpf</b> command does not take into account equal-cost paths or the designated forwarder. For accurate upstream RPF neighbor information, always use the <b>show pim join extensive</b> command when bidirectional PIM is configured.                 |

## Sample Output

### show multicast rpf

```

user@host> show multicast rpf

Multicast RPF table: inet.0, 12 entries

0.0.0.0/0
 Protocol: Static

10.255.14.132/32
 Protocol: Direct
 Interface: lo0.0

10.255.245.91/32
 Protocol: IS-IS
 Interface: so-1/1/1.0
 Neighbor: 192.168.195.21

127.0.0.1/32
Inactive172.16.0.0/12
Protocol: Static
Interface: fxp0.0

```

```

Neighbor: 192.168.14.254

192.168.0.0/16
Protocol: Static
Interface: fxp0.0
Neighbor: 192.168.14.254

192.168.14.0/24
Protocol: Direct
Interface: fxp0.0

192.168.14.132/32
Protocol: Local

192.168.195.20/30
Protocol: Direct
Interface: so-1/1/1.0

192.168.195.22/32
Protocol: Local

192.168.195.36/30
Protocol: IS-IS
Interface: so-1/1/1.0
Neighbor: 192.168.195.21

```

#### show multicast rpf inet6

```

user@host> show multicast rpf inet6

Multicast RPF table: inet6.0, 12 entries

::10.255.14.132/128
 Protocol: Direct
 Interface: lo0.0

::10.255.245.91/128
 Protocol: IS-IS
 Interface: so-1/1/1.0
 Neighbor: fe80::2a0:a5ff:fe28:2e8c

::192.168.195.20/126
 Protocol: Direct
 Interface: so-1/1/1.0

::192.168.195.22/128
 Protocol: Local

::192.168.195.36/126
 Protocol: IS-IS
 Interface: so-1/1/1.0
 Neighbor: fe80::2a0:a5ff:fe28:2e8c

::192.168.195.76/126
 Protocol: Direct
 Interface: fe-2/2/0.0

::192.168.195.77/128
 Protocol: Local

```

```
fe80::/64
Protocol: Direct
Interface: so-1/1/1.0

fe80::290:69ff:fe0c:993a/128
Protocol: Local

fe80::2a0:a5ff:fe12:84f/128
Protocol: Direct
Interface: lo0.0

ff02::2/128
Protocol: PIM

ff02::d/128
Protocol: PIM
```

#### show multicast rpf prefix

```
user@host> show multicast rpf ff02::/16

Multicast RPF table: inet6.0, 13 entries

ff02::2/128
 Protocol: PIM

ff02::d/128
 Protocol: PIM

...
```

#### show multicast rpf summary

```
user@host> show multicast rpf summary

Multicast RPF table: inet.0, 16 entries
Multicast RPF table: inet6.0, 12 entries
```



## show multicast scope

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-----------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 5873</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 5873</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Syntax</b>                                       | <pre>show multicast scope &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show multicast scope &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>                          | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                            |
| <b>Description</b>                                  | Display administratively scoped IP multicast information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                                      | <p><b>none</b>—Display standard information about administratively scoped multicast information for all supported address families in all routing instances.</p> <p><b>inet   inet6</b>—(Optional) Display scoped multicast information for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display administratively scoped information for a specific multicast instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>List of Sample Output</b>                        | <a href="#">show multicast scope on page 5874</a><br><a href="#">show multicast scope inet on page 5874</a><br><a href="#">show multicast scope inet6 on page 5874</a>                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Output Fields</b>                                | <p><a href="#">Table 459</a> describes the output fields for the <b>show multicast scope</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                                                                                                                              |

**Table 459: show multicast scope Output Fields**

| Field Name          | Field Description                                           |
|---------------------|-------------------------------------------------------------|
| <b>Scope name</b>   | Name of the multicast scope.                                |
| <b>Group Prefix</b> | Range of multicast groups that are scoped.                  |
| <b>Interface</b>    | Interface that is the boundary of the administrative scope. |

Table 459: show multicast scope Output Fields (*continued*)

| Field Name      | Field Description                 |
|-----------------|-----------------------------------|
| Resolve Rejects | Number of kernel resolve rejects. |

## Sample Output

### show multicast scope

```
user@host> show multicast scope
```

| Scope name | Group Prefix   | Interface  | Resolve Rejects |
|------------|----------------|------------|-----------------|
| 232-net    | 232.232.0.0/16 | fe-0/0/0.1 | 0               |
| local      | 239.255.0.0/16 | fe-0/0/0.1 | 0               |
| local      | ff05::/16      | fe-0/0/0.1 | 0               |
| larry      | ff05::1234/128 | fe-0/0/0.1 | 0               |

### show multicast scope inet

```
user@host> show multicast scope inet
```

| Scope name | Group Prefix   | Interface  | Resolve Rejects |
|------------|----------------|------------|-----------------|
| 232-net    | 232.232.0.0/16 | fe-0/0/0.1 | 0               |
| local      | 239.255.0.0/16 | fe-0/0/0.1 | 0               |

### show multicast scope inet6

```
user@host> show multicast scope inet6
```

| Scope name | Group Prefix   | Interface  | Resolve Rejects |
|------------|----------------|------------|-----------------|
| local      | ff05::/16      | fe-0/0/0.1 | 0               |
| larry      | ff05::1234/128 | fe-0/0/0.1 | 0               |

## show multicast sessions

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-----------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 5875</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 5875</a>                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Syntax</b>                                       | show multicast sessions<br><brief   detail   extensive><br><logical-system (all   <i>logical-system-name</i> )><br>< <i>regular-expression</i> >                                                                                                                                                                                                                                                                                                                                                        |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | show multicast sessions<br><brief   detail   extensive><br>< <i>regular-expression</i> >                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>                          | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>                                  | Display information about announced IP multicast sessions.                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                                      | <b>none</b> —Display standard information about all multicast sessions for all routing instances.<br><br><b>brief   detail   extensive</b> —(Optional) Display the specified level of output.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.<br><br><b><i>regular-expression</i></b> —(Optional) Display information about announced sessions that match a UNIX-style regular expression. |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>List of Sample Output</b>                        | <a href="#">show multicast sessions on page 5876</a><br><a href="#">show multicast sessions regular-expression detail on page 5876</a>                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Output Fields</b>                                | Table 460 describes the output fields for the <b>show multicast sessions</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                           |

**Table 460: show multicast sessions Output Fields**

| Field Name          | Field Description                               |
|---------------------|-------------------------------------------------|
| <i>session-name</i> | Name of the known announced multicast sessions. |

## Sample Output

### show multicast sessions

```
user@host> show multicast sessions
1-Department of Biological Sciences, LSU
...
Monterey Bay - DockCam
Monterey Bay - JettyCam
Monterey Bay - StandCam
Monterey DockCam
Monterey DockCam / ROV cam
...
NASA TV (MPEG-1)
...
UO Broadcast - NASA Videos - 25 Years of Progress
UO Broadcast - NASA Videos - Journey through the Solar System
UO Broadcast - NASA Videos - Life in the Universe
UO Broadcast - NASA Videos - Nasa and the Airplane
UO Broadcasts OPB's Oregon Story
UO DOD News Clips
UO Medical Management of Biological Casualties (1)
UO Medical Management of Biological Casualties (2)
UO Medical Management of Biological Casualties (3)
...
376 active sessions.
```

### show multicast sessions regular-expression detail

```
user@host> show multicast sessions "NASA TV" detail
SDP Version: 0 Originated by: -@128.223.83.33
Session: NASA TV (MPEG-1)
Description: NASA television in MPEG-1 format, provided by Private University.
Please contact the UO if you have problems with this feed.
Email: Your Name Here <multicast@lists.private.edu>
Phone: Your Name Here <888/555-1212>
Bandwidth: AS:1000
Start time: permanent
Stop time: none
Attribute: type:broadcast
Attribute: tool:IP/TV Content Manager 3.4.14
Attribute: live:capture:1
Attribute: x-iptv-capture:mp1s
Media: video 54302 RTP/AVP 32 31 96 97
Connection Data: 224.2.231.45 ttl 127
Attribute: quality:8
Attribute: framerate:30
Attribute: rtpmap:96 WBIH/90000
Attribute: rtpmap:97 MP4V-ES/90000
Attribute: x-iptv-svr:video 128.223.91.191 live
Attribute: fmtp:32 type=mpeg1
Media: audio 28848 RTP/AVP 14 0 96 3 5 97 98 99 100 101 102 10 11 103 104 105 106
Connection Data: 224.2.145.37 ttl 127
Attribute: rtpmap:96 X-WAVE/8000
Attribute: rtpmap:97 L8/8000/2
Attribute: rtpmap:98 L8/8000
Attribute: rtpmap:99 L8/22050/2
Attribute: rtpmap:100 L8/22050
Attribute: rtpmap:101 L8/11025/2
Attribute: rtpmap:102 L8/11025
Attribute: rtpmap:103 L16/22050/2
```

Attribute: rtpmap:104 L16/22050

1 matching sessions.

## show multicast usage

---

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-----------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 5878</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 5878</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Syntax</b>                                       | <pre>show multicast usage &lt;brief   detail&gt; &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show multicast usage &lt;brief   detail&gt; &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>                          | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>                                  | Display usage information about the 10 most active Distance Vector Multicast Routing Protocol (DVMRP) or Protocol Independent Multicast (PIM) groups.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                                      | <p><b>none</b>—Display multicast usage information for all supported address families for all routing instances.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>inet   inet6</b>—(Optional) Display usage information for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information about the most active DVMRP or PIM groups for a specific multicast instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>List of Sample Output</b>                        | <a href="#">show multicast usage on page 5879</a><br><a href="#">show multicast usage brief on page 5879</a><br><a href="#">show multicast usage instance on page 5879</a><br><a href="#">show multicast usage detail on page 5880</a>                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Output Fields</b>                                | <a href="#">Table 461</a> describes the output fields for the <b>show multicast usage</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

Table 461: show multicast usage Output Fields

| Field Name      | Field Description                                                                                                                                                                        |
|-----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Instance</b> | Name of the routing instance. (Displayed when multicast is configured within a routing instance.)                                                                                        |
| <b>Group</b>    | Group address.                                                                                                                                                                           |
| <b>Sources</b>  | Number of sources.                                                                                                                                                                       |
| <b>Packets</b>  | Number of packets that have been forwarded to this prefix. If one or more of the packets forwarded statistic queries fails or times out, the packets field displays <b>unavailable</b> . |
| <b>Bytes</b>    | Number of bytes that have been forwarded to this prefix. If one or more of the packets forwarded statistic queries fails or times out, the bytes field displays <b>unavailable</b> .     |
| <b>Prefix</b>   | IP address.                                                                                                                                                                              |
| <b>/len</b>     | Prefix length.                                                                                                                                                                           |
| <b>Groups</b>   | Number of multicast groups.                                                                                                                                                              |

## Sample Output

### show multicast usage

```

user@host> show multicast usage
Group Sources Packets Bytes
228.0.0.0 1 52847 4439148
239.1.1.1 2 13450 1125530

Prefix /len Groups Packets Bytes
10.255.14.144 /32 2 66254 5561304
10.255.70.15 /32 1 43 3374...
```

### show multicast usage brief

The output for the **show multicast usage brief** command is identical to that for the **show multicast usage** command. For sample output, see [show multicast usage on page 5879](#).

### show multicast usage instance

```

user@host> show multicast usage instance VPN-A
Group Sources Packets Bytes
224.2.127.254 1 5538 509496
224.0.1.39 1 13 624
224.0.1.40 1 13 624

Prefix /len Groups Packets Bytes
192.168.195.34 /32 1 5538 509496
10.255.14.30 /32 1 13 624
```

```
10.255.245.91 /32 1 13 624
...
```

#### show multicast usage detail

```
user@host> show multicast usage detail
```

| Group                                                   | Sources | Packets | Bytes   |
|---------------------------------------------------------|---------|---------|---------|
| 228.0.0.0                                               | 1       | 53159   | 4465356 |
| Source: 10.255.14.144 /32 Packets: 53159 Bytes: 4465356 |         |         |         |
| 239.1.1.1                                               | 2       | 13450   | 1125530 |
| Source: 10.255.14.144 /32 Packets: 13407 Bytes: 1122156 |         |         |         |
| Source: 10.255.70.15 /32 Packets: 43 Bytes: 3374        |         |         |         |

| Prefix           | /len | Groups         | Packets        | Bytes   |
|------------------|------|----------------|----------------|---------|
| 10.255.14.144    | /32  | 2              | 66566          | 5587512 |
| Group: 228.0.0.0 |      | Packets: 53159 | Bytes: 4465356 |         |
| Group: 239.1.1.1 |      | Packets: 13407 | Bytes: 1122156 |         |
| 10.255.70.15     | /32  | 1              | 43             | 3374    |
| Group: 239.1.1.1 |      | Packets: 43    | Bytes: 3374    |         |



## show pim bootstrap

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                    |
|-----------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 5881</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 5881</a>                                                                                                                                                                                                                                                                                   |
| <b>Syntax</b>                                       | <pre>show pim bootstrap &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                      |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show pim bootstrap &lt;instance <i>instance-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>                          | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>instance</b> option introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                  |
| <b>Description</b>                                  | For sparse mode only, display information about Protocol Independent Multicast (PIM) bootstrap routers.                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                                      | <p><b>none</b>—Display PIM bootstrap router information for all routing instances.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information about bootstrap routers for a specific PIM-enabled routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>List of Sample Output</b>                        | <a href="#">show pim bootstrap on page 5882</a><br><a href="#">show pim bootstrap instance on page 5882</a>                                                                                                                                                                                                                                                                                        |
| <b>Output Fields</b>                                | <a href="#">Table 462</a> describes the output fields for the <b>show pim bootstrap</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                           |

**Table 462: show pim bootstrap Output Fields**

| Field Name           | Field Description                                                            |
|----------------------|------------------------------------------------------------------------------|
| <b>Instance</b>      | Name of the routing instance.                                                |
| <b>BSR</b>           | Bootstrap router.                                                            |
| <b>Pri</b>           | Priority of the routing device as elected to be the bootstrap router.        |
| <b>Local address</b> | Local routing device address.                                                |
| <b>Pri</b>           | Local routing device address priority to be elected as the bootstrap router. |

Table 462: show pim bootstrap Output Fields (*continued*)

| Field Name     | Field Description                                                                                    |
|----------------|------------------------------------------------------------------------------------------------------|
| <b>State</b>   | Local routing device election state: <b>Candidate</b> , <b>Elected</b> , or <b>Ineligible</b> .      |
| <b>Timeout</b> | How long until the local routing device declares the bootstrap router to be unreachable, in seconds. |

## Sample Output

### show pim bootstrap

```
user@host> show pim bootstrap
Instance: PIM.master
```

| BSR                     | Pri | Local address           | Pri | State      | Timeout |
|-------------------------|-----|-------------------------|-----|------------|---------|
| None                    | 0   | 10.255.71.46            | 0   | InEligible | 0       |
| feco:1:1:1:1:0:aff:785c | 34  | feco:1:1:1:1:0:aff:7c12 | 0   | InEligible | 0       |

### show pim bootstrap instance

```
user@host> show pim bootstrap instance VPN-A
Instance: PIM.VPN-A
```

| BSR  | Pri | Local address   | Pri | State      | Timeout |
|------|-----|-----------------|-----|------------|---------|
| None | 0   | 192.168.196.105 | 0   | InEligible | 0       |

## show pim interfaces

**List of Syntax** [Syntax on page 5883](#)

[Syntax \(EX Series Switch and the QFX Series\) on page 5883](#)

**Syntax** show pim interfaces  
 <inet | inet6>  
 <instance *instance-name*>  
 <logical-system (all | *logical-system-name*)>

**Syntax (EX Series Switch and the QFX Series)** show pim interfaces  
 <inet | inet6>  
 <instance *instance-name*>

**Release Information** Command introduced before Junos OS Release 7.4.  
 Command introduced in Junos OS Release 9.0 for EX Series switches.  
**inet6** and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.  
 Command introduced in Junos OS Release 11.3 for the QFX Series.  
 Support for bidirectional PIM added in Junos OS Release 12.1.

**Description** Display information about the interfaces on which Protocol Independent Multicast (PIM) is configured.

**Options** **none**—Display interface information for all family addresses for all routing instances.

**inet | inet6**—(Optional) Display interface information for IPv4 or IPv6 family addresses, respectively.

**instance *instance-name***—(Optional) Display information about interfaces for a specific PIM-enabled routing instance.

**logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

**Required Privilege Level** view

**List of Sample Output** [show pim interfaces on page 5884](#)

**Output Fields** [Table 463](#) describes the output fields for the **show pim interfaces** command. Output fields are listed in the approximate order in which they appear.

**Table 463: show pim interfaces Output Fields**

| Field Name      | Field Description                                                                          |
|-----------------|--------------------------------------------------------------------------------------------|
| <b>Instance</b> | Name of the routing instance.                                                              |
| <b>Name</b>     | Interface name.                                                                            |
| <b>State</b>    | State of the interface. The state also is displayed in the <b>show interfaces</b> command. |

Table 463: show pim interfaces Output Fields (*continued*)

| Field Name          | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Mode</b>         | <p>PIM mode running on the interface:</p> <ul style="list-style-type: none"> <li>• <b>B</b>—In bidirectional mode, multicast groups are carried across the network over bidirectional shared trees. This type of tree minimizes PIM routing state, which is especially important in networks with numerous and dispersed senders and receivers.</li> <li>• <b>S</b>—In sparse mode, routing devices must join and leave multicast groups explicitly. Upstream routing devices do not forward multicast traffic to this routing device unless this device has sent an explicit request (using a join message) to receive multicast traffic.</li> <li>• <b>Dense</b>—Unlike sparse mode, where data is forwarded only to routing devices sending an explicit request, dense mode implements a flood-and-prune mechanism, similar to DVMRP (the first multicast protocol used to support the multicast backbone). (Not supported on QFX Series.)</li> <li>• <b>Sparse-Dense</b>—Sparse-dense mode allows the interface to operate on a per-group basis in either sparse or dense mode. A group specified as <b>dense</b> is not mapped to a rendezvous point (RP). Instead, data packets destined for that group are forwarded using PIM-Dense Mode (PIM-DM) rules. A group specified as <b>sparse</b> is mapped to an RP, and data packets are forwarded using PIM-Sparse Mode (PIM-SM) rules. (Not supported on QFX Series.)</li> </ul> <p>When sparse-dense mode is configured, the output includes both <b>S</b> and <b>D</b>. When bidirectional-sparse mode is configured, the output includes <b>S</b> and <b>B</b>. When bidirectional-sparse-dense mode is configured, the output includes <b>B</b>, <b>S</b>, and <b>D</b>.</p> |
| <b>IP</b>           | Version number of the address family on the interface: <b>4</b> (IPv4) or <b>6</b> (IPv6).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>V</b>            | PIM version running on the interface: 1 or 2.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>State</b>        | <p>State of PIM on the interface:</p> <ul style="list-style-type: none"> <li>• <b>Active</b>—Bidirectional mode is enabled on the interface and on all PIM neighbors.</li> <li>• <b>DR</b>—Designated router.</li> <li>• <b>NotCap</b>—Bidirectional mode is not enabled on the interface. This can happen when bidirectional PIM is not configured locally, when one of the neighbors is not configured for bidirectional PIM, or when one of the neighbors has not implemented the bidirectional PIM protocol.</li> <li>• <b>NotDR</b>—Not the designated router.</li> <li>• <b>P2P</b>—Point to point.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>NbrCnt</b>       | Number of neighbors that have been seen on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>JoinCnt(sg)</b>  | Number of (s,g) join messages that have been seen on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>JointCnt(*g)</b> | Number of (*g) join messages that have been seen on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>DR address</b>   | Address of the designated router.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## Sample Output

### show pim interfaces

```

user@host> show pim interfaces
Stat = Status, V = Version, NbrCnt = Neighbor Count,
S = Sparse, D = Dense, B = Bidirectional,
DR = Designated Router, P2P = Point-to-point link,

```

Active = Bidirectional is active, NotCap = Not Bidirectional Capable

| Name           | Stat | Mode | IP | V | State        | NbrCnt | JoinCnt(sg/*g) | DR address |
|----------------|------|------|----|---|--------------|--------|----------------|------------|
| ge-0/3/0.0     | Up   | S    | 4  | 2 | NotDR,NotCap | 1      | 0/0            | 40.0.0.3   |
| ge-0/3/3.50    | Up   | S    | 4  | 2 | DR,NotCap    | 1      | 9901/100       | 50.0.0.2   |
| ge-0/3/3.51    | Up   | S    | 4  | 2 | DR,NotCap    | 1      | 0/0            | 51.0.0.2   |
| pe-1/2/0.32769 | Up   | S    | 4  | 2 | P2P,NotCap   | 0      | 0/0            |            |

show pim join

|                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| List of Syntax                               | <a href="#">Syntax on page 5886</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 5886</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Syntax                                       | <pre>show pim join &lt;brief   detail   extensive   summary&gt; &lt;bidirectional   dense   sparse&gt; &lt;exact&gt; &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;range&gt; &lt;rp <i>ip-address/prefix</i>   source <i>ip-address/prefix</i>&gt; &lt;sg   star-g&gt;</pre>                                                                                                                                                                                                                                                         |
| Syntax (EX Series Switch and the QFX Series) | <pre>show pim join &lt;brief   detail   extensive   summary&gt; &lt;dense   sparse&gt; &lt;exact&gt; &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;range&gt; &lt;rp <i>ip-address/prefix</i>   source <i>ip-address/prefix</i>&gt; &lt;sg   star-g&gt;</pre>                                                                                                                                                                                                                                                                                                                                   |
| Release Information                          | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>summary</b> option introduced in Junos OS Release 9.6.</p> <p><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Support for bidirectional PIM added in Junos OS Release 12.1.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Multiple new filter options introduced in Junos OS Release 13.2.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description                                  | <p>Display information about Protocol Independent Multicast (PIM) groups for all PIM modes.</p> <p>For bidirectional PIM, display information about PIM group ranges (*G-range) for each active bidirectional RP group range, in addition to each of the joined (*G) routes.</p>                                                                                                                                                                                                                                                                                                                                |
| Options                                      | <p><b>none</b>—Display the standard information about PIM groups for all supported family addresses for all routing instances.</p> <p><b>brief   detail   extensive   summary</b>—(Optional) Display the specified level of output.</p> <p><b>bidirectional   dense   sparse</b>—(Optional) Display information about PIM bidirectional mode, dense mode, or sparse and source-specific multicast (SSM) mode entries.</p> <p><b>exact</b>—(Optional) Display information about only the group that exactly matches the specified group address.</p>                                                             |

**inet | inet6**—(Optional) Display PIM group information for IPv4 or IPv6 family addresses, respectively.

**instance *instance-name***—(Optional) Display information about groups for the specified PIM-enabled routing instance only.

**logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

**range**—(Optional) Address range of the group, specified as *prefix/prefix-length*.

**rp *ip-address/prefix* | source *ip-address/prefix***—(Optional) Display information about the PIM entries with a specified rendezvous point (RP) address and prefix or with a specified source address and prefix. You can omit the prefix.

**sg | star-g**—(Optional) Display information about PIM (S,G) or (\*,G) entries.

**Required Privilege Level**

view

**Related Documentation**

- [clear pim join on page 5832](#)
- [Example: Configuring Bidirectional PIM](#)

**List of Sample Output**

[show pim join summary on page 5891](#)  
[show pim join \(PIM Sparse Mode\) on page 5891](#)  
[show pim join \(Bidirectional PIM\) on page 5892](#)  
[show pim join inet6 on page 5892](#)  
[show pim join inet6 star-g on page 5893](#)  
[show pim join instance <instance-name> on page 5893](#)  
[show pim join detail on page 5893](#)  
[show pim join extensive \(PIM Sparse Mode\) on page 5894](#)  
[show pim join extensive \(Bidirectional PIM\) on page 5895](#)  
[show pim join extensive \(Bidirectional PIM with a Directly Connected Phantom RP\) on page 5896](#)  
[show pim join instance <instance-name> extensive on page 5896](#)  
[show pim join extensive \(Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs\) on page 5897](#)  
[show pim join extensive \(Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs\) on page 5898](#)  
[show pim join summary on page 5899](#)  
[show pim join \(PIM Sparse Mode\) on page 5900](#)  
[show pim join \(Bidirectional PIM\) on page 5900](#)  
[show pim join inet6 on page 5901](#)  
[show pim join inet6 star-g on page 5901](#)  
[show pim join instance <instance-name> on page 5901](#)  
[show pim join detail on page 5902](#)  
[show pim join extensive \(PIM Sparse Mode\) on page 5902](#)  
[show pim join extensive \(Bidirectional PIM\) on page 5903](#)

[show pim join extensive \(Bidirectional PIM with a Directly Connected Phantom RP\) on page 5904](#)

[show pim join instance <instance-name> extensive on page 5904](#)

[show pim join extensive \(Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs\) on page 5905](#)

[show pim join extensive \(Multipoint LDP with Multicast-Only Fast Reroute\) on page 5906](#)

**Output Fields** Table 464 describes the output fields for the **show pim join** command. Output fields are listed in the approximate order in which they appear.

**Table 464: show pim join Output Fields**

| Field Name                               | Field Description                                                                                                                                                                                                                                                                                                                                                                                    | Level of Output                            |
|------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------|
| <b>Instance</b>                          | Name of the routing instance.                                                                                                                                                                                                                                                                                                                                                                        | <b>brief detail extensive summary none</b> |
| <b>Family</b>                            | Name of the address family: <b>inet</b> (IPv4) or <b>inet6</b> (IPv6).                                                                                                                                                                                                                                                                                                                               | <b>brief detail extensive summary none</b> |
| <b>Route type</b>                        | Type of multicast route: (S,G) or (*G).                                                                                                                                                                                                                                                                                                                                                              | <b>summary</b>                             |
| <b>Route count</b>                       | Number of (S,G) routes and number of (*G) routes.                                                                                                                                                                                                                                                                                                                                                    | <b>summary</b>                             |
| <b>R</b>                                 | Rendezvous Point Tree.                                                                                                                                                                                                                                                                                                                                                                               | <b>brief detail extensive none</b>         |
| <b>S</b>                                 | Sparse.                                                                                                                                                                                                                                                                                                                                                                                              | <b>brief detail extensive none</b>         |
| <b>W</b>                                 | Wildcard.                                                                                                                                                                                                                                                                                                                                                                                            | <b>brief detail extensive none</b>         |
| <b>Group</b>                             | Group address.                                                                                                                                                                                                                                                                                                                                                                                       | <b>brief detail extensive none</b>         |
| <b>Bidirectional group prefix length</b> | For bidirectional PIM, length of the IP prefix for RP group ranges.                                                                                                                                                                                                                                                                                                                                  | <b>All levels</b>                          |
| <b>Source</b>                            | Multicast source: <ul style="list-style-type: none"> <li>• * (wildcard value)</li> <li>• <i>ipv4-address</i></li> <li>• <i>ipv6-address</i></li> </ul>                                                                                                                                                                                                                                               | <b>brief detail extensive none</b>         |
| <b>RP</b>                                | Rendezvous point for the PIM group.                                                                                                                                                                                                                                                                                                                                                                  | <b>brief detail extensive none</b>         |
| <b>Flags</b>                             | PIM flags: <ul style="list-style-type: none"> <li>• <b>bidirectional</b>—Bidirectional mode entry.</li> <li>• <b>dense</b>—Dense mode entry.</li> <li>• <b>rptree</b>—Entry is on the rendezvous point tree.</li> <li>• <b>sparse</b>—Sparse mode entry.</li> <li>• <b>spt</b>—Entry is on the shortest-path tree for the source.</li> <li>• <b>wildcard</b>—Entry is on the shared tree.</li> </ul> | <b>brief detail extensive none</b>         |



Table 464: show pim join Output Fields (*continued*)

| Field Name                             | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Level of Output                    |
|----------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------|
| <b>Upstream interface</b>              | <p>RPF interface toward the source address for the source-specific state (S,G) or toward the rendezvous point (RP) address for the non-source-specific state (*G).</p> <p>For bidirectional PIM, <b>RP Link</b> means that the interface is directly connected to a subnet that contains a phantom RP address.</p> <p>A pseudo multipoint LDP (M-LDP) interface appears on egress nodes in M-LDP point-to-multipoint LSPs with inband signaling.</p>                                            | <b>brief detail extensive none</b> |
| <b>Upstream neighbor</b>               | <p>Information about the upstream neighbor: <b>Direct</b>, <b>Local</b>, <b>Unknown</b>, or a specific IP address.</p> <p>For bidirectional PIM, <b>Direct</b> means that the interface is directly connected to a subnet that contains a phantom RP address.</p> <p>The multipoint LDP (M-LDP) root appears on egress nodes in M-LDP point-to-multipoint LSPs with inband signaling.</p>                                                                                                       | <b>extensive</b>                   |
| <b>Upstream state</b>                  | <p>When multicast-only fast reroute (MoFRR) is configured in a PIM domain, the upstream interface for the active path. A PIM router propagates join messages on two upstream RPF interfaces to receive multicast traffic on both links for the same join request. Preference is given to two paths that do not converge to the same immediate upstream router. PIM installs appropriate multicast routes with upstream neighbors as RPF next hops with two (primary and backup) interfaces.</p> | <b>extensive</b>                   |
| <b>Active upstream neighbor</b>        | <p>On the MoFRR primary path, the IP address of the neighbor that is directly connected to the active upstream interface.</p>                                                                                                                                                                                                                                                                                                                                                                   | <b>extensive</b>                   |
| <b>MoFRR Backup upstream interface</b> | <p>The MoFRR upstream interface that is used when the primary path fails.</p> <p>When the primary path fails, the backup path is upgraded to primary, and traffic is forwarded accordingly. If there are alternate paths available, a new backup path is calculated and the appropriate multicast route is updated or installed.</p>                                                                                                                                                            | <b>extensive</b>                   |

Table 464: show pim join Output Fields (*continued*)

| Field Name                  | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Level of Output  |
|-----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>Upstream state</b>       | <p>Information about the upstream interface:</p> <ul style="list-style-type: none"> <li>• <b>Join to RP</b>—Sending a join to the rendezvous point.</li> <li>• <b>Join to Source</b>—Sending a join to the source.</li> <li>• <b>Local RP</b>—Sending neither join messages nor prune messages toward the RP, because this routing device is the rendezvous point.</li> <li>• <b>Local Source</b>—Sending neither join messages nor prune messages toward the source, because the source is locally attached to this routing device.</li> <li>• <b>No Prune to RP</b>—Automatically sent to RP when SPT and RPT are on the same path.</li> <li>• <b>Prune to RP</b>—Sending a prune to the rendezvous point.</li> <li>• <b>Prune to Source</b>—Sending a prune to the source.</li> </ul> <p><b>NOTE:</b> RP group range entries have <b>None</b> in the <b>Upstream state</b> field because RP group ranges do not trigger actual PIM join messages between routing devices.</p>                                   | <b>extensive</b> |
| <b>Downstream neighbors</b> | <p>Information about downstream interfaces:</p> <ul style="list-style-type: none"> <li>• <b>Interface</b>—Interface name for the downstream neighbor.<br/>A pseudo PIM-SM interface appears for all IGMP-only interfaces.<br/>A pseudo multipoint LDP (M-LDP) interface appears on ingress root nodes in M-LDP point-to-multipoint LSPs with inband signaling.</li> <li>• <b>Interface address</b>—Address of the downstream neighbor.</li> <li>• <b>State</b>—Information about the downstream neighbor: <b>join</b> or <b>prune</b>.</li> <li>• <b>Flags</b>—PIM join flags: <b>R</b> (<b>RPtree</b>), <b>S</b> (<b>Sparse</b>), <b>W</b> (<b>Wildcard</b>), or <b>zero</b>.</li> <li>• <b>Uptime</b>—Time since the downstream interface joined the group.</li> <li>• <b>Time since last Join</b>—Time since the last join message was received from the downstream interface.</li> <li>• <b>Time since last Prune</b>—Time since the last prune message was received from the downstream interface.</li> </ul> | <b>extensive</b> |
| <b>Assert Timeout</b>       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |                  |
| <b>Assert Timeout</b>       | Length of time between assert cycles on the downstream interface. Not displayed if the assert timer is null.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>extensive</b> |

Table 464: show pim join Output Fields (*continued*)

| Field Name                                | Field Description                                                                                                                                                                                                                                                                                                             | Level of Output  |
|-------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>Keepalive timeout</b>                  | Time remaining until the downstream join state is updated (in seconds). If the downstream join state is not updated before this keepalive timer reaches zero, the entry is deleted. If there is a directly connected host, <b>Keepalive timeout</b> is <b>Infinity</b> .                                                      | <b>extensive</b> |
| <b>Uptime</b>                             | Time since the creation of (S,G) or (*,G) state. The uptime is not refreshed every time a PIM join message is received for an existing (S,G) or (*,G) state.                                                                                                                                                                  | <b>extensive</b> |
| <b>Bidirectional accepting interfaces</b> | <p>Interfaces on the router that forward bidirectional PIM traffic.</p> <p>The reasons for forwarding bidirectional PIM traffic are that the interface is the winner of the designated forwarder election (<b>DF Winner</b>), or the interface is the reverse path forwarding (RPF) interface toward the RP (<b>RPF</b>).</p> | <b>extensive</b> |

## Sample Output

### show pim join summary

```

user@host> show pim join summary
Instance: PIM.master Family: INET

Route type Route count
(s,g) 2
(*,g) 1

Instance: PIM.master Family: INET6

```

### show pim join (PIM Sparse Mode)

```

user@host> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
 Source: *
 RP: 10.255.14.144
 Flags: sparse,rptree,wildcard
 Upstream interface: Local

Group: 239.1.1.1
 Source: 10.255.14.144
 Flags: sparse,spt
 Upstream interface: Local

Group: 239.1.1.1
 Source: 10.255.70.15
 Flags: sparse,spt
 Upstream interface: so-1/0/0.0

```

```
Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

#### show pim join (Bidirectional PIM)

```
user@host> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.0
 Bidirectional group prefix length: 24
 Source: *
 RP: 10.10.13.2
 Flags: bidirectional,rptree,wildcard
 Upstream interface: ge-0/0/1.0

Group: 224.1.3.0
 Bidirectional group prefix length: 24
 Source: *
 RP: 10.10.1.3
 Flags: bidirectional,rptree,wildcard
 Upstream interface: ge-0/0/1.0 (RP Link)

Group: 225.1.1.0
 Bidirectional group prefix length: 24
 Source: *
 RP: 10.10.13.2
 Flags: bidirectional,rptree,wildcard
 Upstream interface: ge-0/0/1.0

Group: 225.1.3.0
 Bidirectional group prefix length: 24
 Source: *
 RP: 10.10.1.3
 Flags: bidirectional,rptree,wildcard
 Upstream interface: ge-0/0/1.0 (RP Link)

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

#### show pim join inet6

```
user@host> show pim join inet6
Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: ff04::e000:101
 Source: *
 RP: ::46.0.0.13
 Flags: sparse,rptree,wildcard
 Upstream interface: Local

Group: ff04::e000:101
 Source: ::1.1.1.1
 Flags: sparse
 Upstream interface: unknown (no neighbor)

Group: ff04::e800:101
 Source: ::1.1.1.1
 Flags: sparse
 Upstream interface: unknown (no neighbor)
```

```

Group: ff04::e800:101
Source: ::1.1.1.2
Flags: sparse
Upstream interface: unknown (no neighbor)

```

#### show pim join inet6 star-g

```

user@host> show pim join inet6 star-g
Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: ff04::e000:101
Source: *
RP: ::46.0.0.13
Flags: sparse,rptree,wildcard
Upstream interface: Local

```

#### show pim join instance <instance-name>

```

user@host> show pim join instance VPN-A
Instance: PIM.VPN-A Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 235.1.1.2
Source: *
RP: 10.10.47.100
Flags: sparse,rptree,wildcard
Upstream interface: Local

Group: 235.1.1.2
Source: 192.168.195.74
Flags: sparse,spt
Upstream interface: at-0/3/1.0

Group: 235.1.1.2
Source: 192.168.195.169
Flags: sparse
Upstream interface: so-1/0/1.0

Instance: PIM.VPN-A Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```

#### show pim join detail

```

user@host> show pim join detail
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
Source: *
RP: 10.255.14.144
Flags: sparse,rptree,wildcard
Upstream interface: Local

Group: 239.1.1.1
Source: 10.255.14.144
Flags: sparse,spt
Upstream interface: Local

Group: 239.1.1.1
Source: 10.255.70.15

```

```
Flags: sparse,spt
Upstream interface: so-1/0/0.0
```

```
Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

#### show pim join extensive (PIM Sparse Mode)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
 Source: *
 RP: 10.255.14.144
 Flags: sparse,rptree,wildcard
 Upstream interface: Local
 Upstream neighbor: Local
 Upstream state: Local RP
 Uptime: 00:03:49
 Downstream neighbors:
 Interface: so-1/0/0.0
 10.111.10.2 State: Join Flags: SRW Timeout: 174
 Uptime: 00:03:49 Time since last Join: 00:01:49
 Interface: mt-1/1/0.32768
 10.10.47.100 State: Join Flags: SRW Timeout: Infinity
 Uptime: 00:03:49 Time since last Join: 00:01:49
 Number of downstream interfaces: 2

Group: 239.1.1.1
 Source: 10.255.14.144
 Flags: sparse,spt
 Upstream interface: Local
 Upstream neighbor: Local
 Upstream state: Local Source, Local RP
 Keepalive timeout: 344
 Uptime: 00:03:49
 Downstream neighbors:
 Interface: so-1/0/0.0
 10.111.10.2 State: Join Flags: S Timeout: 174
 Uptime: 00:03:49 Time since last Prune: 00:01:49
 Interface: mt-1/1/0.32768
 10.10.47.100 State: Join Flags: S Timeout: Infinity
 Uptime: 00:03:49 Time since last Prune: 00:01:49
 Number of downstream interfaces: 2

Group: 239.1.1.1
 Source: 10.255.70.15
 Flags: sparse,spt
 Upstream interface: so-1/0/0.0
 Upstream neighbor: 10.111.10.2
 Upstream state: Local RP, Join to Source
 Keepalive timeout: 344
 Uptime: 00:03:49
 Downstream neighbors:
 Interface: Pseudo-GMP
 fe-0/0/0.0 fe-0/0/1.0 fe-0/0/3.0
 Interface: so-1/0/0.0 (pruned)
 10.111.10.2 State: Prune Flags: SR Timeout: 174
 Uptime: 00:03:49 Time since last Prune: 00:01:49
 Interface: mt-1/1/0.32768
```

```

10.10.47.100 State: Join Flags: S Timeout: Infinity
Uptime: 00:03:49 Time since last Prune: 00:01:49
Number of downstream interfaces: 3

```

```

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```

### show pim join extensive (Bidirectional PIM)

```

user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```

```

Group: 224.1.1.0
 Bidirectional group prefix length: 24
 Source: *
 RP: 10.10.13.2
 Flags: bidirectional,rptree,wildcard
 Upstream interface: ge-0/0/1.0
 Upstream neighbor: 10.10.1.2
 Upstream state: None
 Uptime: 00:03:49
 Bidirectional accepting interfaces:
 Interface: ge-0/0/1.0 (RPF)
 Interface: lo0.0 (DF Winner)
 Number of downstream interfaces: 0

Group: 225.1.1.0
 Bidirectional group prefix length: 24
 Source: *
 RP: 10.10.13.2
 Flags: bidirectional,rptree,wildcard
 Upstream interface: ge-0/0/1.0
 Upstream neighbor: 10.10.1.2
 Upstream state: None
 Uptime: 00:03:49
 Bidirectional accepting interfaces:
 Interface: ge-0/0/1.0 (RPF)
 Interface: lo0.0 (DF Winner)
 Downstream neighbors:
 Interface: lt-1/0/10.24
 10.0.24.4 State: Join RW Timeout: 185
 Interface: lt-1/0/10.23
 10.0.23.3 State: Join RW Timeout: 184
 Number of downstream interfaces: 2

Group: 225.1.3.0
 Bidirectional group prefix length: 24
 Source: *
 RP: 10.10.1.3
 Flags: bidirectional,rptree,wildcard
 Upstream interface: ge-0/0/1.0 (RP Link)
 Upstream neighbor: Direct
 Upstream state: Local RP
 Uptime: 00:03:49
 Bidirectional accepting interfaces:
 Interface: ge-0/0/1.0 (RPF)
 Interface: lo0.0 (DF Winner)
 Interface: xe-4/1/0.0 (DF Winner)
 Number of downstream interfaces: 0

```

Instance: PIM.master Family: INET6  
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

#### show pim join extensive (Bidirectional PIM with a Directly Connected Phantom RP)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

```
Group: 224.1.3.0
 Bidirectional group prefix length: 24
 Source: *
 RP: 10.10.1.3
 Flags: bidirectional,rptree,wildcard
 Upstream interface: ge-0/0/1.0 (RP Link)
 Upstream neighbor: Direct
 Upstream state: Local RP
 Uptime: 00:03:49
 Bidirectional accepting interfaces:
 Interface: ge-0/0/1.0 (RPF)
 Interface: lo0.0 (DF Winner)
 Interface: xe-4/1/0.0 (DF Winner)
 Number of downstream interfaces: 0
```

#### show pim join instance <instance-name> extensive

```
user@host> show pim join instance VPN-A extensive
Instance: PIM.VPN-A Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

```
Group: 235.1.1.2
 Source: *
 RP: 10.10.47.100
 Flags: sparse,rptree,wildcard
 Upstream interface: Local
 Upstream neighbor: Local
 Upstream state: Local RP
 Uptime: 00:03:49
 Downstream neighbors:
 Interface: mt-1/1/0.32768
 10.10.47.101 State: Join Flags: SRW Timeout: 156
 Uptime: 00:03:49 Time since last Join: 00:01:49
 Number of downstream interfaces: 1
```

```
Group: 235.1.1.2
 Source: 192.168.195.74
 Flags: sparse,spt
 Upstream interface: at-0/3/1.0
 Upstream neighbor: 10.111.30.2
 Upstream state: Local RP, Join to Source
 Keepalive timeout: 156
 Uptime: 00:14:52
```

```
Group: 235.1.1.2
 Source: 192.168.195.169
 Flags: sparse
 Upstream interface: so-1/0/1.0
 Upstream neighbor: 10.111.20.2
 Upstream state: Local RP, Join to Source
 Keepalive timeout: 156
 Uptime: 00:14:52
```



**show pim join extensive (Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)**

```

user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 232.1.1.1
 Source: 192.168.219.11
 Flags: sparse,spt
 Upstream interface: fe-1/3/1.0
 Upstream neighbor: Direct
 Upstream state: Local Source
 Keepalive timeout:
 Uptime: 11:27:55
 Downstream neighbors:
 Interface: Pseudo-MLDP
 Interface: lt-1/2/0.25
 1.2.5.2 State: Join Flags: S Timeout: Infinity
 Uptime: 11:27:55 Time since last Join: 11:27:55

Group: 232.1.1.2
 Source: 192.168.219.11
 Flags: sparse,spt
 Upstream interface: fe-1/3/1.0
 Upstream neighbor: Direct
 Upstream state: Local Source
 Keepalive timeout:
 Uptime: 11:27:41
 Downstream neighbors:
 Interface: Pseudo-MLDP

Group: 232.1.1.3
 Source: 192.168.219.11
 Flags: sparse,spt
 Upstream interface: fe-1/3/1.0
 Upstream neighbor: Direct
 Upstream state: Local Source
 Keepalive timeout:
 Uptime: 11:27:41
 Downstream neighbors:
 Interface: Pseudo-MLDP

Group: 232.2.2.2
 Source: 1.2.7.7
 Flags: sparse,spt
 Upstream interface: lt-1/2/0.27
 Upstream neighbor: Direct
 Upstream state: Local Source
 Keepalive timeout:
 Uptime: 11:27:25
 Downstream neighbors:
 Interface: Pseudo-MLDP

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: ff3e::1:2
 Source: abcd::1:2:7:7
 Flags: sparse,spt
 Upstream interface: lt-1/2/0.27
 Upstream neighbor: Direct

```

```
Upstream state: Local Source
Keepalive timeout:
Uptime: 11:27:26
Downstream neighbors:
 Interface: Pseudo-MLDP
```

#### show pim join extensive (Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 227.1.1.1
 Source: *
 RP: 1.1.1.1
 Flags: sparse,rptree,wildcard
 Upstream interface: Local
 Upstream neighbor: Local
 Upstream state: Local RP
 Uptime: 11:31:33
 Downstream neighbors:
 Interface: fe-1/3/0.0
 192.168.209.9 State: Join Flags: SRW Timeout: Infinity
 Uptime: 11:31:33 Time since last Join: 11:31:32

Group: 232.1.1.1
 Source: 192.168.219.11
 Flags: sparse,spt
 Upstream protocol: MLDP
 Upstream interface: Pseudo MLDP
 Upstream neighbor: MLDP LSP root <1.1.1.2>
 Upstream state: Join to Source
 Keepalive timeout:
 Uptime: 11:31:32
 Downstream neighbors:
 Interface: so-0/1/3.0
 192.168.92.9 State: Join Flags: S Timeout: Infinity
 Uptime: 11:31:30 Time since last Join: 11:31:30
 Downstream neighbors:
 Interface: fe-1/3/0.0
 192.168.209.9 State: Join Flags: S Timeout: Infinity
 Uptime: 11:31:32 Time since last Join: 11:31:32

Group: 232.1.1.2
 Source: 192.168.219.11
 Flags: sparse,spt
 Upstream protocol: MLDP
 Upstream interface: Pseudo MLDP
 Upstream neighbor: MLDP LSP root <1.1.1.2>
 Upstream state: Join to Source
 Keepalive timeout:
 Uptime: 11:31:32
 Downstream neighbors:
 Interface: so-0/1/3.0
 192.168.92.9 State: Join Flags: S Timeout: Infinity
 Uptime: 11:31:30 Time since last Join: 11:31:30
 Downstream neighbors:
 Interface: lt-1/2/0.14
 1.1.4.4 State: Join Flags: S Timeout: 177
 Uptime: 11:30:33 Time since last Join: 00:00:33
 Downstream neighbors:
```

```

Interface: fe-1/3/0.0
 192.168.209.9 State: Join Flags: S Timeout: Infinity
 Uptime: 11:31:32 Time since last Join: 11:31:32

Group: 232.1.1.3
 Source: 192.168.219.11
 Flags: sparse,spt
 Upstream protocol: MLDP
 Upstream interface: Pseudo MLDP
 Upstream neighbor: MLDP LSP root <1.1.1.2>
 Upstream state: Join to Source
 Keepalive timeout:
 Uptime: 11:31:32
 Downstream neighbors:
 Interface: fe-1/3/0.0
 192.168.209.9 State: Join Flags: S Timeout: Infinity
 Uptime: 11:31:32 Time since last Join: 11:31:32

Group: 232.2.2.2
 Source: 1.2.7.7
 Flags: sparse,spt
 Upstream protocol: MLDP
 Upstream interface: Pseudo MLDP
 Upstream neighbor: MLDP LSP root <1.1.1.2>
 Upstream state: Join to Source
 Keepalive timeout:
 Uptime: 11:31:30
 Downstream neighbors:
 Interface: so-0/1/3.0
 192.168.92.9 State: Join Flags: S Timeout: Infinity
 Uptime: 11:31:30 Time since last Join: 11:31:30

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: ff3e::1:2
 Source: abcd::1:2:7:7
 Flags: sparse,spt
 Upstream protocol: MLDP
 Upstream interface: Pseudo MLDP
 Upstream neighbor: MLDP LSP root <1.1.1.2>
 Upstream state: Join to Source
 Keepalive timeout:
 Uptime: 11:31:32
 Downstream neighbors:
 Interface: fe-1/3/0.0
 fe80::21f:12ff:fea5:c4db State: Join Flags: S Timeout: Infinity
 Uptime: 11:31:32 Time since last Join: 11:31:32

```

## Sample Output

### show pim join summary

```

user@host> show pim join summary
Instance: PIM.master Family: INET

Route type Route count
(s,g) 2
(*,g) 1

Instance: PIM.master Family: INET6

```

### show pim join (PIM Sparse Mode)

```
user@host> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
 Source: *
 RP: 10.255.14.144
 Flags: sparse,rptree,wildcard
 Upstream interface: Local

Group: 239.1.1.1
 Source: 10.255.14.144
 Flags: sparse,spt
 Upstream interface: Local

Group: 239.1.1.1
 Source: 10.255.70.15
 Flags: sparse,spt
 Upstream interface: so-1/0/0.0

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

### show pim join (Bidirectional PIM)

```
user@host> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.0
 Bidirectional group prefix length: 24
 Source: *
 RP: 10.10.13.2
 Flags: bidirectional,rptree,wildcard
 Upstream interface: ge-0/0/1.0

Group: 224.1.3.0
 Bidirectional group prefix length: 24
 Source: *
 RP: 10.10.1.3
 Flags: bidirectional,rptree,wildcard
 Upstream interface: ge-0/0/1.0 (RP Link)

Group: 225.1.1.0
 Bidirectional group prefix length: 24
 Source: *
 RP: 10.10.13.2
 Flags: bidirectional,rptree,wildcard
 Upstream interface: ge-0/0/1.0

Group: 225.1.3.0
 Bidirectional group prefix length: 24
 Source: *
 RP: 10.10.1.3
 Flags: bidirectional,rptree,wildcard
 Upstream interface: ge-0/0/1.0 (RP Link)

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

**show pim join inet6**

```

user@host> show pim join inet6
Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: ff04::e000:101
 Source: *
 RP: ::46.0.0.13
 Flags: sparse,rptree,wildcard
 Upstream interface: Local

Group: ff04::e000:101
 Source: ::1.1.1.1
 Flags: sparse
 Upstream interface: unknown (no neighbor)

Group: ff04::e800:101
 Source: ::1.1.1.1
 Flags: sparse
 Upstream interface: unknown (no neighbor)

Group: ff04::e800:101
 Source: ::1.1.1.2
 Flags: sparse
 Upstream interface: unknown (no neighbor)

```

**show pim join inet6 star-g**

```

user@host> show pim join inet6 star-g
Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: ff04::e000:101
 Source: *
 RP: ::46.0.0.13
 Flags: sparse,rptree,wildcard
 Upstream interface: Local

```

**show pim join instance <instance-name>**

```

user@host> show pim join instance VPN-A
Instance: PIM.VPN-A Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 235.1.1.2
 Source: *
 RP: 10.10.47.100
 Flags: sparse,rptree,wildcard
 Upstream interface: Local

Group: 235.1.1.2
 Source: 192.168.195.74
 Flags: sparse,spt
 Upstream interface: at-0/3/1.0

Group: 235.1.1.2
 Source: 192.168.195.169
 Flags: sparse
 Upstream interface: so-1/0/1.0

```

```
Instance: PIM.VPN-A Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

#### show pim join detail

```
user@host> show pim join detail
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
 Source: *
 RP: 10.255.14.144
 Flags: sparse,rptree,wildcard
 Upstream interface: Local

Group: 239.1.1.1
 Source: 10.255.14.144
 Flags: sparse,spt
 Upstream interface: Local

Group: 239.1.1.1
 Source: 10.255.70.15
 Flags: sparse,spt
 Upstream interface: so-1/0/0.0

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

#### show pim join extensive (PIM Sparse Mode)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
 Source: *
 RP: 10.255.14.144
 Flags: sparse,rptree,wildcard
 Upstream interface: Local
 Upstream neighbor: Local
 Upstream state: Local RP
 Uptime: 00:03:49
 Downstream neighbors:
 Interface: so-1/0/0.0
 10.111.10.2 State: Join Flags: SRW Timeout: 174
 Uptime: 00:03:49 Time since last Join: 00:01:49
 Interface: mt-1/1/0.32768
 10.10.47.100 State: Join Flags: SRW Timeout: Infinity
 Uptime: 00:03:49 Time since last Join: 00:01:49
 Number of downstream interfaces: 2

Group: 239.1.1.1
 Source: 10.255.14.144
 Flags: sparse,spt
 Upstream interface: Local
 Upstream neighbor: Local
 Upstream state: Local Source, Local RP
 Keepalive timeout: 344
 Uptime: 00:03:49
 Downstream neighbors:
 Interface: so-1/0/0.0
```

```

 10.111.10.2 State: Join Flags: S Timeout: 174
 Uptime: 00:03:49 Time since last Prune: 00:01:49
 Interface: mt-1/1/0.32768
 10.10.47.100 State: Join Flags: S Timeout: Infinity
 Uptime: 00:03:49 Time since last Prune: 00:01:49
 Number of downstream interfaces: 2

Group: 239.1.1.1
Source: 10.255.70.15
Flags: sparse,spt
Upstream interface: so-1/0/0.0
Upstream neighbor: 10.111.10.2
Upstream state: Local RP, Join to Source
Keepalive timeout: 344
Uptime: 00:03:49
Downstream neighbors:
 Interface: Pseudo-GMP
 fe-0/0/0.0 fe-0/0/1.0 fe-0/0/3.0
 Interface: so-1/0/0.0 (pruned)
 10.111.10.2 State: Prune Flags: SR Timeout: 174
 Uptime: 00:03:49 Time since last Prune: 00:01:49
 Interface: mt-1/1/0.32768
 10.10.47.100 State: Join Flags: S Timeout: Infinity
 Uptime: 00:03:49 Time since last Prune: 00:01:49
 Number of downstream interfaces: 3

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```

#### show pim join extensive (Bidirectional PIM)

```

user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.0
 Bidirectional group prefix length: 24
 Source: *
 RP: 10.10.13.2
 Flags: bidirectional,rptree,wildcard
 Upstream interface: ge-0/0/1.0
 Upstream neighbor: 10.10.1.2
 Upstream state: None
 Uptime: 00:03:49
 Bidirectional accepting interfaces:
 Interface: ge-0/0/1.0 (RPF)
 Interface: lo0.0 (DF Winner)
 Number of downstream interfaces: 0

Group: 225.1.1.0
 Bidirectional group prefix length: 24
 Source: *
 RP: 10.10.13.2
 Flags: bidirectional,rptree,wildcard
 Upstream interface: ge-0/0/1.0
 Upstream neighbor: 10.10.1.2
 Upstream state: None
 Uptime: 00:03:49
 Bidirectional accepting interfaces:
 Interface: ge-0/0/1.0 (RPF)
 Interface: lo0.0 (DF Winner)

```

```
Downstream neighbors:
 Interface: lt-1/0/10.24
 10.0.24.4 State: Join RW Timeout: 185
 Interface: lt-1/0/10.23
 10.0.23.3 State: Join RW Timeout: 184
 Number of downstream interfaces: 2

Group: 225.1.3.0
 Bidirectional group prefix length: 24
 Source: *
 RP: 10.10.1.3
 Flags: bidirectional,rptree,wildcard
 Upstream interface: ge-0/0/1.0 (RP Link)
 Upstream neighbor: Direct
 Upstream state: Local RP
 Uptime: 00:03:49
 Bidirectional accepting interfaces:
 Interface: ge-0/0/1.0 (RPF)
 Interface: lo0.0 (DF Winner)
 Interface: xe-4/1/0.0 (DF Winner)
 Number of downstream interfaces: 0

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

#### show pim join extensive (Bidirectional PIM with a Directly Connected Phantom RP)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.3.0
 Bidirectional group prefix length: 24
 Source: *
 RP: 10.10.1.3
 Flags: bidirectional,rptree,wildcard
 Upstream interface: ge-0/0/1.0 (RP Link)
 Upstream neighbor: Direct
 Upstream state: Local RP
 Uptime: 00:03:49
 Bidirectional accepting interfaces:
 Interface: ge-0/0/1.0 (RPF)
 Interface: lo0.0 (DF Winner)
 Interface: xe-4/1/0.0 (DF Winner)
 Number of downstream interfaces: 0
```

#### show pim join instance <instance-name> extensive

```
user@host> show pim join instance VPN-A extensive
Instance: PIM.VPN-A Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 235.1.1.2
 Source: *
 RP: 10.10.47.100
 Flags: sparse,rptree,wildcard
 Upstream interface: Local
 Upstream neighbor: Local
 Upstream state: Local RP
 Uptime: 00:03:49
 Downstream neighbors:
```



```

Interface: mt-1/1/0.32768
10.10.47.101 State: Join Flags: SRW Timeout: 156
Uptime: 00:03:49 Time since last Join: 00:01:49
Number of downstream interfaces: 1

```

```

Group: 235.1.1.2
Source: 192.168.195.74
Flags: sparse,spt
Upstream interface: at-0/3/1.0
Upstream neighbor: 10.111.30.2
Upstream state: Local RP, Join to Source
Keepalive timeout: 156
Uptime: 00:14:52

```

```

Group: 235.1.1.2
Source: 192.168.195.169
Flags: sparse
Upstream interface: so-1/0/1.0
Upstream neighbor: 10.111.20.2
Upstream state: Local RP, Join to Source
Keepalive timeout: 156
Uptime: 00:14:52

```

#### show pim join extensive (Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```

user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 232.1.1.1
Source: 192.168.219.11
Flags: sparse,spt
Upstream interface: fe-1/3/1.0
Upstream neighbor: Direct
Upstream state: Local Source
Keepalive timeout:
Uptime: 11:27:55
Downstream neighbors:
Interface: Pseudo-MLDP
Interface: lt-1/2/0.25
1.2.5.2 State: Join Flags: S Timeout: Infinity
Uptime: 11:27:55 Time since last Join: 11:27:55

Group: 232.1.1.2
Source: 192.168.219.11
Flags: sparse,spt
Upstream interface: fe-1/3/1.0
Upstream neighbor: Direct
Upstream state: Local Source
Keepalive timeout:
Uptime: 11:27:41
Downstream neighbors:
Interface: Pseudo-MLDP

Group: 232.1.1.3
Source: 192.168.219.11
Flags: sparse,spt
Upstream interface: fe-1/3/1.0
Upstream neighbor: Direct
Upstream state: Local Source
Keepalive timeout:

```

```
Uptime: 11:27:41
Downstream neighbors:
 Interface: Pseudo-MLDP

Group: 232.2.2.2
Source: 1.2.7.7
Flags: sparse,spt
Upstream interface: lt-1/2/0.27
Upstream neighbor: Direct
Upstream state: Local Source
Keepalive timeout:
Uptime: 11:27:25
Downstream neighbors:
 Interface: Pseudo-MLDP

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: ff3e::1:2
Source: abcd::1:2:7:7
Flags: sparse,spt
Upstream interface: lt-1/2/0.27
Upstream neighbor: Direct
Upstream state: Local Source
Keepalive timeout:
Uptime: 11:27:26
Downstream neighbors:
 Interface: Pseudo-MLDP
```

#### show pim join extensive (Multipoint LDP with Multicast-Only Fast Reroute)

```
user@host> show pim join 225.1.1.1 extensive sg
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 225.1.1.1
Source: 10.0.0.1
Flags: sparse,spt
Active upstream interface: fe-1/2/13.0
Active upstream neighbor: 10.0.0.9
MoFRR Backup upstream interface: fe-1/2/14.0
MoFRR Backup upstream neighbor: 10.0.0.21
Upstream state: Join to Source, No Prune to RP
Keepalive timeout: 354
Uptime: 00:00:06
Downstream neighbors:
 Interface: fe-1/2/15.0
 10.0.0.13 State: Join Flags: S Timeout: Infinity
 Uptime: 00:00:06 Time since last Join: 00:00:06
Number of downstream interfaces: 1
```

## show pim neighbors

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-----------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 5907</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 5907</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Syntax</b>                                       | <pre>show pim neighbors &lt;brief   detail&gt; &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show pim neighbors &lt;brief   detail&gt; &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>                          | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional PIM added in Junos OS Release 12.1.</p> <p>Support for the <b>instance all</b> option added in Junos OS Release 12.1.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                       |
| <b>Description</b>                                  | Display information about Protocol Independent Multicast (PIM) neighbors.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                                      | <p><b>none</b>—(Same as <b>brief</b>) Display standard information about PIM neighbors for all supported family addresses for the main instance.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>inet   inet6</b>—(Optional) Display information about PIM neighbors for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance (<i>instance-name</i>   all)</b>—(Optional) Display information about neighbors for the specified PIM-enabled routing instance or for all routing instances.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>List of Sample Output</b>                        | <a href="#">show pim neighbors on page 5909</a><br><a href="#">show pim neighbors brief on page 5909</a><br><a href="#">show pim neighbors instance on page 5909</a><br><a href="#">show pim neighbors detail on page 5909</a><br><a href="#">show pim neighbors detail (With BFD) on page 5910</a>                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Output Fields</b>                                | <p><a href="#">Table 465</a> describes the output fields for the <b>show pim neighbors</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

Table 465: show pim neighbors Output Fields

| Field Name                                       | Field Description                                                                                                                                                                                                                                                                                                                                | Level of Output   |
|--------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------|
| <b>Instance</b>                                  | Name of the routing instance.                                                                                                                                                                                                                                                                                                                    | All levels        |
| <b>Interface</b>                                 | Interface through which the neighbor is reachable.                                                                                                                                                                                                                                                                                               | All levels        |
| <b>Neighbor addr</b>                             | Address of the neighboring PIM routing device.                                                                                                                                                                                                                                                                                                   | All levels        |
| <b>IP</b>                                        | IP version: 4 or 6.                                                                                                                                                                                                                                                                                                                              | All levels        |
| <b>V</b>                                         | PIM version running on the neighbor: 1 or 2.                                                                                                                                                                                                                                                                                                     | All levels        |
| <b>Mode</b>                                      | PIM mode of the neighbor: <b>Sparse</b> , <b>Dense</b> , <b>SparseDense</b> , or <b>Unknown</b> . When the neighbor is running PIM version 2, this mode is always <b>Unknown</b> .                                                                                                                                                               | All levels        |
| <b>Option</b>                                    | Can be one or more of the following: <ul style="list-style-type: none"> <li>• <b>B</b>—Bidirectional Capable.</li> <li>• <b>G</b>—Generation Identifier.</li> <li>• <b>H</b>—Hello Option Holdtime.</li> <li>• <b>L</b>—Hello Option LAN Prune Delay.</li> <li>• <b>P</b>—Hello Option DR Priority.</li> <li>• <b>T</b>—Tracking bit.</li> </ul> | <b>brief</b> none |
| <b>Uptime</b>                                    | Time the neighbor has been operational since the PIM process was last initialized, in the format <b>dd:hh:mm:ss ago</b> for less than a week and <b>nwnd:hh:mm:ss ago</b> for more than a week.                                                                                                                                                  | All levels        |
| <b>Address</b>                                   | Address of the neighboring PIM routing device.                                                                                                                                                                                                                                                                                                   | <b>detail</b>     |
| <b>BFD</b>                                       | Status and operational state of the Bidirectional Forwarding Detection (BFD) protocol on the interface: <b>Enabled</b> , <b>Operational state is up</b> , or <b>Disabled</b> .                                                                                                                                                                   | <b>detail</b>     |
| <b>Hello Option Holdtime</b>                     | Time for which the neighbor is available, in seconds. The range of values is 0 through 65,535.                                                                                                                                                                                                                                                   | <b>detail</b>     |
| <b>Hello Default Holdtime</b>                    | Default holdtime and the time remaining if the <b>holdtime</b> option is not in the received hello message.                                                                                                                                                                                                                                      | <b>detail</b>     |
| <b>Hello Option DR Priority</b>                  | Designated router election priority. The range of values is 0 through 255.                                                                                                                                                                                                                                                                       | <b>detail</b>     |
| <b>Hello Option Generation ID</b>                | 9-digit or 10-digit number used to tag hello messages.                                                                                                                                                                                                                                                                                           | <b>detail</b>     |
| <b>Hello Option Bi-Directional PIM supported</b> | Neighbor can process bidirectional PIM messages.                                                                                                                                                                                                                                                                                                 | <b>detail</b>     |
| <b>Hello Option LAN Prune Delay</b>              | Time to wait before the neighbor receives prune messages, in the format <b>delay nnn ms override nnnn ms</b> .                                                                                                                                                                                                                                   | <b>detail</b>     |

Table 465: show pim neighbors Output Fields (*continued*)

| Field Name                 | Field Description                                                                                                                                                                                                                                                                   | Level of Output |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Join Suppression supported | Neighbor is capable of join suppression.                                                                                                                                                                                                                                            | detail          |
| Rx Join                    | Information about joins received from the neighbor. <ul style="list-style-type: none"> <li><b>Group</b>—Group addresses in the join message.</li> <li><b>Source</b>—Address of the source in the join message.</li> <li><b>Timeout</b>—Time for which the join is valid.</li> </ul> | detail          |

## Sample Output

### show pim neighbors

```

user@host> show pim neighbors
Instance: PIM.master
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority, T = Tracking bit

Interface IP V Mode Option Uptime Neighbor addr
so-1/0/0.0 4 2 HPLG 00:07:10 10.111.10.2

```

### show pim neighbors brief

The output for the **show pim neighbors brief** command is identical to that for the **show pim neighbors** command. For sample output, see [show pim neighbors on page 5909](#).

### show pim neighbors instance

```

user@host> show pim neighbors instance VPN-A
Instance: PIM.VPN-A
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority, T = Tracking bit

Interface IP V Mode Option Uptime Neighbor addr
at-0/3/1.0 4 2 HPLG 00:07:54 10.111.30.2
mt-1/1/0.32768 4 2 HPLG 00:07:22 10.10.47.101
so-1/0/1.0 4 2 HPLG 00:07:50 10.111.20.2

```

### show pim neighbors detail

```

user@host> show pim neighbors detail
Instance: PIM.master
Interface: ge-0/0/1.0

Address: 10.10.1.1, IPv4, PIM v2, Mode: SparseDense, sg Join Count: 0, tsf
Join Count: 2
Hello Option Holdtime: 65535 seconds
Hello Option DR Priority: 1
Hello Option Generation ID: 2053759302
Hello Option Bi-Directional PIM supported
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported

```

```
Address: 10.10.1.2, IPv4, PIM v2, sg Join Count: 0, tsg Join Count: 2
 BFD: Disabled
 Hello Option Holdtime: 105 seconds 93 remaining
 Hello Option DR Priority: 1
 Hello Option Generation ID: 1734018161
 Hello Option Bi-Directional PIM supported
 Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
 Join Suppression supported
```

Interface: lo0.0

```
Address: 10.255.179.246, IPv4, PIM v2, Mode: SparseDense, sg Join Count:
0, tsg Join Count: 0
 Hello Option Holdtime: 65535 seconds
 Hello Option DR Priority: 1
 Hello Option Generation ID: 1997462267
 Hello Option Bi-Directional PIM supported
 Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
 Join Suppression supported
```

#### show pim neighbors detail (With BFD)

```
user@host> show pim neighbors detail
```

Instance: PIM.master

Interface: fe-1/0/0.0

```
Address: 192.168.11.1, IPv4, PIM v2, Mode: Sparse
 Hello Option Holdtime: 65535 seconds
 Hello Option DR Priority: 1
 Hello Option Generation ID: 836607909
 Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

```
Address: 192.168.11.2, IPv4, PIM v2
 BFD: Enabled, Operational state is up
 Hello Default Holdtime: 105 seconds 104 remaining
 Hello Option DR Priority: 1
 Hello Option Generation ID: 1907549685
 Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

Interface: fe-1/0/1.0

```
Address: 192.168.12.1, IPv4, PIM v2
 BFD: Disabled
 Hello Default Holdtime: 105 seconds 80 remaining
 Hello Option DR Priority: 1
 Hello Option Generation ID: 1971554705
 Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

## show pim rps

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-----------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 5911</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 5911</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Syntax</b>                                       | <pre>show pim rps &lt;brief   detail   extensive&gt; &lt;group-address&gt; &lt;inet   inet6&gt; &lt;instance instance-name&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show pim rps &lt;brief   detail   extensive&gt; &lt;group-address&gt; &lt;inet   inet6&gt; &lt;instance instance-name&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>                          | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional PIM added in Junos OS Release 12.1.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>                                  | Display information about Protocol Independent Multicast (PIM) rendezvous points (RPs).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                                      | <p><b>none</b>—Display standard information about PIM RPs for all groups and family addresses for all routing instances.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>group-address</b>—(Optional) Display the RPs for a particular group. If you specify a group address, the output lists the routing device that is the RP for that group.</p> <p><b>inet   inet6</b>—(Optional) Display information for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance instance-name</b>—(Optional) Display information about RPs for a specific PIM-enabled routing instance.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>                        | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Bidirectional PIM</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>List of Sample Output</b>                        | <a href="#">show pim rps on page 5914</a><br><a href="#">show pim rps brief on page 5914</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

[show pim rps <group-address> on page 5914](#)  
[show pim rps <group-address> \(Bidirectional PIM\) on page 5914](#)  
[show pim rps <group-address> \(PIM Dense Mode\) on page 5915](#)  
[show pim rps <group-address> \(SSM Range Without asm-override-ssm Configured\) on page 5915](#)  
[show pim rps <group-address> \(SSM Range With asm-override-ssm Configured and a Sparse-Mode RP\) on page 5915](#)  
[show pim rps <group-address> \(SSM Range With asm-override-ssm Configured and a Bidirectional RP\) on page 5915](#)  
[show pim rps instance on page 5915](#)  
[show pim rps extensive \(PIM Sparse Mode\) on page 5915](#)  
[show pim rps extensive \(Bidirectional PIM\) on page 5916](#)  
[show pim rps extensive \(PIM Anycast RP in Use\) on page 5916](#)

**Output Fields** Table 466 describes the output fields for the **show pim rps** command. Output fields are listed in the approximate order in which they appear.

**Table 466: show pim rps Output Fields**

| Field Name                      | Field Description                                                                                                                                                                                                                                                                                                                                                                          | Level of Output         |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Instance</b>                 | Name of the routing instance.                                                                                                                                                                                                                                                                                                                                                              | All levels              |
| <b>Family or Address family</b> | Name of the address family: <b>inet</b> (IPv4) or <b>inet6</b> (IPv6).                                                                                                                                                                                                                                                                                                                     | All levels              |
| <b>RP address</b>               | Address of the rendezvous point.                                                                                                                                                                                                                                                                                                                                                           | All levels              |
| <b>Type</b>                     | Type of RP: <ul style="list-style-type: none"> <li><b>auto-rp</b>—Address of the RP known through the Auto-RP protocol.</li> <li><b>bootstrap</b>—Address of the RP known through the bootstrap router protocol (BSR).</li> <li><b>embedded</b>—Address of the RP known through an embedded RP (IPv6).</li> <li><b>static</b>—Address of RP known through static configuration.</li> </ul> | <b>brief none</b>       |
| <b>Holdtime</b>                 | How long to keep the RP active, with time remaining, in seconds.                                                                                                                                                                                                                                                                                                                           | All levels              |
| <b>Timeout</b>                  | How long until the local routing device determines the RP to be unreachable, in seconds.                                                                                                                                                                                                                                                                                                   | All levels              |
| <b>Groups</b>                   | Number of groups currently using this RP.                                                                                                                                                                                                                                                                                                                                                  | All levels              |
| <b>Group prefixes</b>           | Addresses of groups that this RP can span.                                                                                                                                                                                                                                                                                                                                                 | <b>brief none</b>       |
| <b>Learned via</b>              | Address and method by which the RP was learned.                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b> |
| <b>Mode</b>                     | The PIM mode of the RP: bidirectional or sparse.<br><br>If a sparse and bidirectional RPs are configured with the same RP address, they appear as separate entries in both formats.                                                                                                                                                                                                        | All levels              |



Table 466: show pim rps Output Fields (*continued*)

| Field Name                    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Level of Output                          |
|-------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------|
| <b>Time Active</b>            | How long the RP has been active, in the format <i>hh:mm:ss</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | detail extensive                         |
| <b>Device Index</b>           | Index value of the order in which Junos OS finds and initializes the interface.<br><br>For bidirectional RPs, the <b>Device Index</b> output field is omitted because bidirectional RPs do not require encapsulation and de-encapsulation interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | detail extensive                         |
| <b>Subunit</b>                | Logical unit number of the interface.<br><br>For bidirectional RPs, the <b>Subunit</b> output field is omitted because bidirectional RPs do not require encapsulation and de-encapsulation interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | detail extensive                         |
| <b>Interface</b>              | Either the encapsulation or the de-encapsulation logical interface, depending on whether this routing device is a designated router (DR) facing an RP router, or is the local RP, respectively.<br><br>For bidirectional RPs, the <b>Interface</b> output field is omitted because bidirectional RPs do not require encapsulation and de-encapsulation interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | detail extensive                         |
| <b>Group Ranges</b>           | Addresses of groups that this RP spans.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | detail extensive<br><i>group-address</i> |
| <b>Active groups using RP</b> | Number of groups currently using this RP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | detail extensive                         |
| <b>total</b>                  | Total number of active groups for this RP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | detail extensive                         |
| <b>Register State for RP</b>  | Current register state for each group: <ul style="list-style-type: none"> <li>• <b>Group</b>—Multicast group address.</li> <li>• <b>Source</b>—Multicast source address for which the PIM register is sent or received, depending on whether this router is a designated router facing an RP router, or is the local RP, respectively:</li> <li>• <b>First Hop</b>—PIM-designated routing device that sent the Register message (the source address in the IP header).</li> <li>• <b>RP Address</b>—RP to which the Register message was sent (the destination address in the IP header).</li> <li>• <b>State</b>: <ul style="list-style-type: none"> <li>On the designated router: <ul style="list-style-type: none"> <li>• <b>Send</b>—Sending Register messages.</li> <li>• <b>Probe</b>—Sent a null register. If a Register-Stop message does not arrive in 5 seconds, the designated router resumes sending Register messages.</li> <li>• <b>Suppress</b>—Received a Register-Stop message. The designated router is waiting for the timer to resume before changing to <b>Probe</b> state.</li> </ul> </li> <li>On the RP: <ul style="list-style-type: none"> <li>• <b>Receive</b>—Receiving Register messages.</li> </ul> </li> </ul> </li> </ul> | extensive                                |
| <b>Anycast-PIM rpset</b>      | If anycast RP is configured, the addresses of the RPs in the set.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | extensive                                |

Table 466: show pim rps Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Level of Output      |
|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------|
| Anycast-PIM local address used | If anycast RP is configured, the local address used by the RP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | extensive            |
| Anycast-PIM Register State     | <p>If anycast RP is configured, the current register state for each group:</p> <ul style="list-style-type: none"> <li>• <b>Group</b>—Multicast group address.</li> <li>• <b>Source</b>—Multicast source address for which the PIM register is sent or received, depending on whether this routing device is a designated router facing an RP router, or is the local RP, respectively.</li> <li>• <b>Origin</b>—How the information was obtained: <ul style="list-style-type: none"> <li>• <b>DIRECT</b>—From a local attachment</li> <li>• <b>MSDP</b>—From the Multicast Source Discovery Protocol (MSDP)</li> <li>• <b>DR</b>—From the designated router</li> </ul> </li> </ul> | extensive            |
| RP selected                    | For sparse mode and bidirectional mode, the identity of the RP for the specified group address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <i>group-address</i> |

## Sample Output

### show pim rps

```

user@host> show pim rps
Instance: PIM.master
Address family INET
RP address Type Mode Holdtime Timeout Groups Group prefixes
10.10.1.3 static bidir 150 None 2 224.1.3.0/24
 225.1.3.0/24
10.10.13.2 static bidir 150 None 2 224.1.1.0/24
 225.1.1.0/24

```

### show pim rps brief

The output for the **show pim rps brief** command is identical to that for the **show pim rps** command. For sample output, see [show pim rps on page 5914](#).

### show pim rps <group-address>

```

user@host> show pim rps 235.100.100.0
Instance: PIM.master
Instance: PIM.master

RP selected: 100.100.100.100

```

### show pim rps <group-address> (Bidirectional PIM)

```

user@host> show pim rps 224.1.1.1
Instance: PIM.master

224.1.0.0/16
 11.4.12.75 (Bidirectional)

RP selected: 11.4.12.75

```

**show pim rps <group-address> (PIM Dense Mode)**

```
user@host> show pim rps 224.1.1.1
Instance: PIM.master

Dense Mode active for group 224.1.1.1
```

**show pim rps <group-address> (SSM Range Without asm-override-ssm Configured)**

```
user@host> show pim rps 224.1.1.1
Instance: PIM.master

Source-specific Mode (SSM) active for group 224.1.1.1
```

**show pim rps <group-address> (SSM Range With asm-override-ssm Configured and a Sparse-Mode RP)**

```
user@host> show pim rps 224.1.1.1
Instance: PIM.master

Source-specific Mode (SSM) active with Sparse Mode ASM override for group 224.1.1.1

224.1.0.0/16
 11.4.12.75

RP selected: 11.4.12.75
```

**show pim rps <group-address> (SSM Range With asm-override-ssm Configured and a Bidirectional RP)**

```
user@host> show pim rps 224.1.1.1
Instance: PIM.master

Source-specific Mode (SSM) active with Sparse Mode ASM override for group 224.1.1.1

224.1.0.0/16
 11.4.12.75 (Bidirectional)

RP selected: (null)
```

**show pim rps instance**

```
user@host> show pim rps instance VPN-A
Instance: PIM.VPN-A
Address family INET
RP address Type Holdtime Timeout Groups Group prefixes
10.10.47.100 static 0 None 1 224.0.0.0/4

Address family INET6
```

**show pim rps extensive (PIM Sparse Mode)**

```
user@host> show pim rps extensive
Instance: PIM.master

Family: INET
RP: 10.255.245.91
Learned via: static configuration
Time Active: 00:05:48
Holdtime: 45 with 36 remaining
Device Index: 122
Subunit: 32768
Interface: pd-6/0/0.32768
Group Ranges:
```

```
224.0.0.0/4, 36s remaining
Active groups using RP:
225.1.1.1
```

```
total 1 groups active
```

```
Register State for RP:
```

| Group     | Source         | FirstHop      | RP Address    | State   | Timeout |
|-----------|----------------|---------------|---------------|---------|---------|
| 225.1.1.1 | 192.168.195.78 | 10.255.14.132 | 10.255.245.91 | Receive | 0       |

#### show pim rps extensive (Bidirectional PIM)

```
user@host> show pim rps extensive
```

```
Instance: PIM.master
```

```
Address family INET
```

```
RP: 10.10.1.3
```

```
Learned via: static configuration
```

```
Mode: Bidirectional
```

```
Time Active: 01:58:07
```

```
Holdtime: 150
```

```
Group Ranges:
```

```
224.1.3.0/24
```

```
225.1.3.0/24
```

```
RP: 10.10.13.2
```

```
Learned via: static configuration
```

```
Mode: Bidirectional
```

```
Time Active: 01:58:07
```

```
Holdtime: 150
```

```
Group Ranges:
```

```
224.1.1.0/24
```

```
225.1.1.0/24
```

#### show pim rps extensive (PIM Anycast RP in Use)

```
user@host> show pim rps extensive
```

```
Instance: PIM.master
```

```
Family: INET
```

```
RP: 10.10.10.2
```

```
Learned via: static configuration
```

```
Time Active: 00:54:52
```

```
Holdtime: 0
```

```
Device Index: 130
```

```
Subunit: 32769
```

```
Interface: pimd.32769
```

```
Group Ranges:
```

```
224.0.0.0/4
```

```
Active groups using RP:
```

```
224.10.10.10
```

```
total 1 groups active
```

```
Anycast-PIM rpset:
```

```
10.100.111.34
```

```
10.100.111.17
```

```
10.100.111.55
```

```
Anycast-PIM local address used: 10.100.111.1
```

```
Anycast-PIM Register State:
```

| Group        | Source     | Origin |
|--------------|------------|--------|
| 224.1.1.1    | 10.10.95.2 | DIRECT |
| 224.1.1.2    | 10.10.95.2 | DIRECT |
| 224.10.10.10 | 10.10.70.1 | MSDP   |
| 224.10.10.11 | 10.10.70.1 | MSDP   |
| 224.20.20.1  | 10.10.71.1 | DR     |

Address family INET6

Anycast-PIM rpset:

ab::1

ab::2

Anycast-PIM local address used: cd::1

Anycast-PIM Register State:

| Group         | Source       | Origin |
|---------------|--------------|--------|
| ::224.1.1.1   | ::10.10.95.2 | DIRECT |
| ::224.1.1.2   | ::10.10.95.2 | DIRECT |
| ::224.20.20.1 | ::10.10.71.1 | DR     |

## show pim source

---

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-----------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 5918</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 5918</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Syntax</b>                                       | <pre>show pim source &lt;brief   detail&gt; &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;source-prefix&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show pim source &lt;brief   detail&gt; &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;source-prefix&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>                          | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>                                  | Display information about the Protocol Independent Multicast (PIM) source reverse path forwarding (RPF) state.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                                      | <p><b>none</b>—Display standard information about the PIM RPF state for all supported family addresses for all routing instances.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>inet   inet6</b>—(Optional) Display information for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information about the RPF state for a specific PIM-enabled routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>source-prefix</b>—(Optional) Display the state for source RPF states in the given range.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>List of Sample Output</b>                        | <a href="#">show pim source on page 5919</a><br><a href="#">show pim source brief on page 5919</a><br><a href="#">show pim source detail on page 5919</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Output Fields</b>                                | <a href="#">Table 467</a> describes the output fields for the <b>show pim source</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

Table 467: show pim source Output Fields

| Field Name         | Field Description                                                     |
|--------------------|-----------------------------------------------------------------------|
| Instance           | Name of the routing instance.                                         |
| Source             | Address of the source or reverse path.                                |
| Prefix/length      | Prefix and prefix length for the route used to reach the RPF address. |
| Upstream interface | RPF interface toward the source address.                              |
| Upstream Neighbor  | Address of the RPF neighbor used to reach the source address.         |

## Sample Output

### show pim source

```

user@host> show pim source
Instance: PIM.master Family: INET

Source 10.255.14.144
 Prefix 10.255.14.144/32
 Upstream interface Local
 Upstream neighbor Local

Source 10.255.70.15
 Prefix 10.255.70.15/32
 Upstream interface so-1/0/0.0
 Upstream neighbor 10.111.10.2

Instance: PIM.master Family: INET6

```

### show pim source brief

The output for the **show pim source brief** command is identical to that for the **show pim source** command. For sample output, see [show pim source on page 5919](#).

### show pim source detail

```

user@host> show pim source detail
Instance: PIM.master Family: INET

Source 10.255.14.144
 Prefix 10.255.14.144/32
 Upstream interface Local
 Upstream neighbor Local
 Active groups:228.0.0.0
 239.1.1.1
 239.1.1.1

Source 10.255.70.15
 Prefix 10.255.70.15/32
 Upstream interface so-1/0/0.0
 Upstream neighbor 10.111.10.2
 Active groups:239.1.1.1

```

Instance: PIM.master Family: INET6



## show pim statistics

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 5921</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 5921</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Syntax</b>                                       | <pre>show pim statistics &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show pim statistics &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>                          | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional PIM added in Junos OS Release 12.1.</p>                                                                                                                                                                              |
| <b>Description</b>                                  | Display Protocol Independent Multicast (PIM) statistics.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                                      | <p><b>none</b>—Display PIM statistics.</p> <p><b>inet   inet6</b>—(Optional) Display IPv4 or IPv6 PIM statistics, respectively.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display statistics for a specific routing instance enabled by Protocol Independent Multicast (PIM).</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display statistics about the specified interface.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>                        | <ul style="list-style-type: none"> <li>• <a href="#">clear pim statistics on page 5836</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>List of Sample Output</b>                        | <a href="#">show pim statistics on page 5927</a><br><a href="#">show pim statistics inet interface &lt;interface-name&gt; on page 5928</a><br><a href="#">show pim statistics inet6 interface &lt;interface-name&gt; on page 5929</a><br><a href="#">show pim statistics interface &lt;interface-name&gt; on page 5929</a>                                                                                                                                                                                                                                            |
| <b>Output Fields</b>                                | <p><a href="#">Table 468</a> describes the output fields for the <b>show pim statistics</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                                                                                                      |

Table 468: show pim statistics Output Fields

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                       |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Instance</b>         | <p>Name of the routing instance.</p> <p>This field only appears if you specify an interface, for example:</p> <ul style="list-style-type: none"> <li>• <b>inet</b> interface <i>interface-name</i></li> <li>• <b>inet6</b> interface <i>interface-name</i></li> <li>• <b>interface</b> <i>interface-name</i></li> </ul>                                                                                                 |
| <b>Family</b>           | <p>Output is for IPv4 or IPv6 PIM statistics. <b>INET</b> indicates IPv4 statistics, and <b>INET6</b> indicates IPv6 statistics.</p> <p>This field only appears if you specify an interface, for example:</p> <ul style="list-style-type: none"> <li>• <b>inet</b> interface <i>interface-name</i></li> <li>• <b>inet6</b> interface <i>interface-name</i></li> <li>• <b>interface</b> <i>interface-name</i></li> </ul> |
| <b>PIM statistics</b>   | PIM statistics for all interfaces or for the specified interface.                                                                                                                                                                                                                                                                                                                                                       |
| <b>PIM message type</b> | Message type for which statistics are displayed.                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Received</b>         | Number of received statistics.                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Sent</b>             | Number of messages sent of a certain type.                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Rx errors</b>        | Number of received packets that contained errors.                                                                                                                                                                                                                                                                                                                                                                       |
| <b>V2 Hello</b>         | PIM version 2 hello packets.                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>V2 Register</b>      | PIM version 2 register packets.                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>V2 Register Stop</b> | PIM version 2 register stop packets.                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>V2 Join Prune</b>    | PIM version 2 join and prune packets.                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>V2 Bootstrap</b>     | PIM version 2 bootstrap packets.                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>V2 Assert</b>        | PIM version 2 assert packets.                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>V2 Graft</b>         | PIM version 2 graft packets.                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>V2 Graft Ack</b>     | PIM version 2 graft acknowledgment packets.                                                                                                                                                                                                                                                                                                                                                                             |
| <b>V2 Candidate RP</b>  | PIM version 2 candidate RP packets.                                                                                                                                                                                                                                                                                                                                                                                     |

Table 468: show pim statistics Output Fields (*continued*)

| Field Name                       | Field Description                                                                                                                                              |
|----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| V2 State Refresh                 | PIM version 2 control messages related to PIM dense mode (PIM-DM) state refresh.<br><br>State refresh is an extension to PIM-DM. It not supported in Junos OS. |
| V2 DF Election                   | PIM version 2 send and receive messages associated with bidirectional PIM designated forwarder election.                                                       |
| V1 Query                         | PIM version 1 query packets.                                                                                                                                   |
| V1 Register                      | PIM version 1 register packets.                                                                                                                                |
| V1 Register Stop                 | PIM version 1 register stop packets.                                                                                                                           |
| V1 Join Prune                    | PIM version 1 join and prune packets.                                                                                                                          |
| V1 RP Reachability               | PIM version 1 RP reachability packets.                                                                                                                         |
| V1 Assert                        | PIM version 1 assert packets.                                                                                                                                  |
| V1 Graft                         | PIM version 1 graft packets.                                                                                                                                   |
| V1 Graft Ack                     | PIM version 1 graft acknowledgment packets.                                                                                                                    |
| AutoRP Announce                  | Auto-RP announce packets.                                                                                                                                      |
| AutoRP Mapping                   | Auto-RP mapping packets.                                                                                                                                       |
| AutoRP Unknown type              | Auto-RP packets with an unknown type.                                                                                                                          |
| Anycast Register                 | Auto-RP announce packets.                                                                                                                                      |
| Anycast Register Stop            | Auto-RP announce packets.                                                                                                                                      |
| Global Statistics                | Summary of PIM statistics for all interfaces.                                                                                                                  |
| Hello dropped on neighbor policy | Number of hello packets dropped because of a configured neighbor policy.                                                                                       |
| Unknown type                     | Number of PIM control packets received with an unknown type.                                                                                                   |
| V1 Unknown type                  | Number of PIM version 1 control packets received with an unknown type.                                                                                         |
| Unknown Version                  | Number of PIM control packets received with an unknown version. The version is not version 1 or version 2.                                                     |

Table 468: show pim statistics Output Fields (*continued*)

| Field Name                             | Field Description                                                                                                         |
|----------------------------------------|---------------------------------------------------------------------------------------------------------------------------|
| <b>Neighbor unknown</b>                | Number of PIM control packets received (excluding PIM hello) without first receiving the hello packet.                    |
| <b>Bad Length</b>                      | Number of PIM control packets received for which the packet size does not match the PIM length field in the packet.       |
| <b>Bad Checksum</b>                    | Number of PIM control packets received for which the calculated checksum does not match the checksum field in the packet. |
| <b>Bad Receive If</b>                  | Number of PIM control packets received on an interface that does not have PIM configured.                                 |
| <b>Rx Bad Data</b>                     | Number of PIM control packets received that contain data for TCP Bad register packets.                                    |
| <b>Rx Intf disabled</b>                | Number of PIM control packets received on an interface that has PIM disabled.                                             |
| <b>Rx V1 Require V2</b>                | Number of PIM version 1 control packets received on an interface configured for PIM version 2.                            |
| <b>Rx V2 Require V1</b>                | Number of PIM version 2 control packets received on an interface configured for PIM version 1.                            |
| <b>Rx Register not RP</b>              | Number of PIM register packets received when the router is not the RP for the group.                                      |
| <b>Rx Register no route</b>            | Number of PIM register packets received when the RP does not have a unicast route back to the source.                     |
| <b>Rx Register no decap if</b>         | Number of PIM register packets received when the RP does not have a de-encapsulation interface.                           |
| <b>Null Register Timeout</b>           | Number of NULL register timeout packets.                                                                                  |
| <b>RP Filtered Source</b>              | Number of PIM packets received when the router has a source address filter configured for the RP.                         |
| <b>Rx Unknown Reg Stop</b>             | Number of register stop messages received with an unknown type.                                                           |
| <b>Rx Join/Prune no state</b>          | Number of join and prune messages received for which the router has no state.                                             |
| <b>Rx Join/Prune on upstream if</b>    | Number of join and prune messages received on the interface used to reach the upstream router, toward the RP.             |
| <b>Rx Join/Prune for invalid group</b> | Number of join or prune messages received for invalid multicast group addresses.                                          |

Table 468: show pim statistics Output Fields (*continued*)

| Field Name                            | Field Description                                                                                                                                           |
|---------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Rx Join/Prune messages dropped</b> | Number of join and prune messages received and dropped.                                                                                                     |
| <b>Rx sparse join for dense group</b> | Number of PIM sparse mode join messages received for a group that is configured for dense mode.                                                             |
| <b>Rx Graft/Graft Ack no state</b>    | Number of graft and graft acknowledgment messages received for which the router or switch has no state.                                                     |
| <b>Rx Graft on upstream if</b>        | Number of graft messages received on the interface used to reach the upstream router, toward the RP.                                                        |
| <b>Rx CRP not BSR</b>                 | Number of BSR messages received in which the PIM message type is Candidate-RP-Advertisement, not Bootstrap.                                                 |
| <b>Rx BSR when BSR</b>                | Number of BSR messages received in which the PIM message type is Bootstrap.                                                                                 |
| <b>Rx BSR not RPF if</b>              | Number of BSR messages received on an interface that is not the RPF interface.                                                                              |
| <b>Rx unknown hello opt</b>           | Number of PIM hello packets received with options that Junos OS does not support.                                                                           |
| <b>Rx data no state</b>               | Number of PIM control packets received for which the router has no state for the data type.                                                                 |
| <b>Rx RP no state</b>                 | Number of PIM control packets received for which the router has no state for the RP.                                                                        |
| <b>Rx aggregate</b>                   | Number of PIM aggregate MDT packets received.                                                                                                               |
| <b>Rx malformed packet</b>            | Number of PIM control packets received with a malformed IP unicast or multicast address family.                                                             |
| <b>No RP</b>                          | Number of PIM control packets received with no RP address.                                                                                                  |
| <b>No register encaps if</b>          | Number of PIM register packets received when the first-hop router does not have an encapsulation interface.                                                 |
| <b>No route upstream</b>              | Number of PIM control packets received when the router does not have a unicast route to the the interface used to reach the upstream router, toward the RP. |
| <b>Nexthop Unusable</b>               | Number of PIM control packets with an unusable nexthop. A path can be unusable if the route is hidden or the link is down.                                  |
| <b>RP mismatch</b>                    | Number of PIM control packets received for which the router has an RP mismatch.                                                                             |

Table 468: show pim statistics Output Fields (*continued*)

| Field Name                                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>RP mode mismatch</b>                    | RP mode (sparse or bidirectional) mismatches encountered when processing join and prune messages.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>RPF neighbor unknown</b>                | Number of PIM control packets received for which the router has an unknown RPF neighbor for the source.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Rx Joins/Prunes filtered</b>            | The number of join and prune messages filtered because of configured route filters and source address filters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Tx Joins/Prunes filtered</b>            | The number of join and prune messages filtered because of configured route filters and source address filters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Embedded-RP invalid addr</b>            | Number of packets received with an invalid embedded RP address in PIM join messages and other types of messages sent between routing domains.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Embedded-RP limit exceed</b>            | Number of times the limit configured with the <b>maximum-rps</b> statement is exceeded. The <b>maximum-rps</b> statement limits the number of embedded RPs created in a specific routing instance. The range is from 1 through 500. The default is 100.                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Embedded-RP added</b>                   | <p>Number of packets in which the embedded RP for IPv6 is added.</p> <p>The following receive events trigger extraction of an IPv6 embedded RP address on the router:</p> <ul style="list-style-type: none"> <li>• Multicast Listener Discovery (MLD) report for an embedded RP multicast group address</li> <li>• PIM join message with an embedded RP multicast group address</li> <li>• Static embedded RP multicast group address associated with an interface</li> <li>• Packets sent to an embedded RP multicast group address received on the DR</li> </ul> <p>An embedded RP node discovered through these receive events is added if it does not already exist on the routing platform.</p> |
| <b>Embedded-RP removed</b>                 | Number of packets in which the embedded RP for IPv6 is removed. The embedded RP is removed whenever all PIM join states using this RP are removed or the configuration changes to remove the embedded RP feature.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Rx Register msgs filtering drop</b>     | Number of received register messages dropped because of a filter configured for PIM register messages.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Tx Register msgs filtering drop</b>     | Number of register messages dropped because of a filter configured for PIM register messages.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Rx Bidir Join/Prune on non-Bidir if</b> | Error counter for join and prune messages received on non-bidirectional PIM interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

Table 468: show pim statistics Output Fields (*continued*)

| Field Name                              | Field Description                                                                          |
|-----------------------------------------|--------------------------------------------------------------------------------------------|
| <b>Rx Bidir Join/Prune on non-DF if</b> | Error counter for join and prune messages received on non-designated forwarder interfaces. |

## Sample Output

### show pim statistics

```

user@host> show pim statistics
PIM Message type Received Sent Rx errors
V2 Hello 15 32 0
V2 Register 0 362 0
V2 Register Stop 483 0 0
V2 Join Prune 18 518 0
V2 Bootstrap 0 0 0
V2 Assert 0 0 0
V2 Graft 0 0 0
V2 Graft Ack 0 0 0
V2 Candidate RP 0 0 0
V2 State Refresh 0 0 0
V2 DF Election 0 0 0
V1 Query 0 0 0
V1 Register 0 0 0
V1 Register Stop 0 0 0
V1 Join Prune 0 0 0
V1 RP Reachability 0 0 0
V1 Assert 0 0 0
V1 Graft 0 0 0
V1 Graft Ack 0 0 0
AutoRP Announce 0 0 0
AutoRP Mapping 0 0 0
AutoRP Unknown type 0 0 0
Anycast Register 0 0 0
Anycast Register Stop 0 0 0

```

#### Global Statistics

```

Hello dropped on neighbor policy 0
Unknown type 0
V1 Unknown type 0
Unknown Version 0
Neighbor unknown 0
Bad Length 0
Bad Checksum 0
Bad Receive If 0
Rx Bad Data 0
Rx Intf disabled 0
Rx V1 Require V2 0
Rx V2 Require V1 0
Rx Register not RP 0
Rx Register no route 0
Rx Register no decap if 0
Null Register Timeout 0
RP Filtered Source 0
Rx Unknown Reg Stop 0
Rx Join/Prune no state 0

```

|                                     |   |
|-------------------------------------|---|
| Rx Join/Prune on upstream if        | 0 |
| Rx Join/Prune for invalid group     | 5 |
| Rx Join/Prune messages dropped      | 0 |
| Rx sparse join for dense group      | 0 |
| Rx Graft/Graft Ack no state         | 0 |
| Rx Graft on upstream if             | 0 |
| Rx CRP not BSR                      | 0 |
| Rx BSR when BSR                     | 0 |
| Rx BSR not RPF if                   | 0 |
| Rx unknown hello opt                | 0 |
| Rx data no state                    | 0 |
| Rx RP no state                      | 0 |
| Rx aggregate                        | 0 |
| Rx malformed packet                 | 0 |
| Rx illegal TTL                      | 0 |
| Rx illegal destination address      | 0 |
| No RP                               | 0 |
| No register encap if                | 0 |
| No route upstream                   | 0 |
| Nexthop Unusable                    | 0 |
| RP mismatch                         | 0 |
| RP mode mismatch                    | 0 |
| RPF neighbor unknown                | 0 |
| Rx Joins/Prunes filtered            | 0 |
| Tx Joins/Prunes filtered            | 0 |
| Embedded-RP invalid addr            | 0 |
| Embedded-RP limit exceed            | 0 |
| Embedded-RP added                   | 0 |
| Embedded-RP removed                 | 0 |
| Rx Register msgs filtering drop     | 0 |
| Tx Register msgs filtering drop     | 0 |
| Rx Bidir Join/Prune on non-Bidir if | 0 |
| Rx Bidir Join/Prune on non-DF if    | 0 |

## Sample Output

show pim statistics inet interface <interface-name>

```
user@host> show pim statistics inet interface ge-0/3/0.0
Instance: PIM.master Family: INET
```

PIM Interface statistics for ge-0/3/0.0

| PIM Message type   | Received | Sent | Rx errors |
|--------------------|----------|------|-----------|
| V2 Hello           | 0        | 4    | 0         |
| V2 Register        | 0        | 0    | 0         |
| V2 Register Stop   | 0        | 0    | 0         |
| V2 Join Prune      | 0        | 0    | 0         |
| V2 Bootstrap       | 0        | 0    | 0         |
| V2 Assert          | 0        | 0    | 0         |
| V2 Graft           | 0        | 0    | 0         |
| V2 Graft Ack       | 0        | 0    | 0         |
| V2 Candidate RP    | 0        | 0    | 0         |
| V1 Query           | 0        | 0    | 0         |
| V1 Register        | 0        | 0    | 0         |
| V1 Register Stop   | 0        | 0    | 0         |
| V1 Join Prune      | 0        | 0    | 0         |
| V1 RP Reachability | 0        | 0    | 0         |
| V1 Assert          | 0        | 0    | 0         |
| V1 Graft           | 0        | 0    | 0         |
| V1 Graft Ack       | 0        | 0    | 0         |



|                       |   |   |   |
|-----------------------|---|---|---|
| AutoRP Announce       | 0 | 0 | 0 |
| AutoRP Mapping        | 0 | 0 | 0 |
| AutoRP Unknown type   | 0 |   |   |
| Anycast Register      | 0 | 0 | 0 |
| Anycast Register Stop | 0 | 0 | 0 |

## Sample Output

**show pim statistics inet6 interface <interface-name>**

```
user@host> show pim statistics inet6 interface ge-0/3/0.0
Instance: PIM.master Family: INET6
```

PIM Interface statistics for ge-0/3/0.0

| PIM Message type      | Received | Sent | Rx errors |
|-----------------------|----------|------|-----------|
| V2 Hello              | 0        | 4    | 0         |
| V2 Register           | 0        | 0    | 0         |
| V2 Register Stop      | 0        | 0    | 0         |
| V2 Join Prune         | 0        | 0    | 0         |
| V2 Bootstrap          | 0        | 0    | 0         |
| V2 Assert             | 0        | 0    | 0         |
| V2 Graft              | 0        | 0    | 0         |
| V2 Graft Ack          | 0        | 0    | 0         |
| V2 Candidate RP       | 0        | 0    | 0         |
| Anycast Register      | 0        | 0    | 0         |
| Anycast Register Stop | 0        | 0    | 0         |

## Sample Output

**show pim statistics interface <interface-name>**

```
user@host> show pim statistics interface ge-0/3/0.0
Instance: PIM.master Family: INET
```

PIM Interface statistics for ge-0/3/0.0

| PIM Message type      | Received | Sent | Rx errors |
|-----------------------|----------|------|-----------|
| V2 Hello              | 0        | 3    | 0         |
| V2 Register           | 0        | 0    | 0         |
| V2 Register Stop      | 0        | 0    | 0         |
| V2 Join Prune         | 0        | 0    | 0         |
| V2 Bootstrap          | 0        | 0    | 0         |
| V2 Assert             | 0        | 0    | 0         |
| V2 Graft              | 0        | 0    | 0         |
| V2 Graft Ack          | 0        | 0    | 0         |
| V2 Candidate RP       | 0        | 0    | 0         |
| V1 Query              | 0        | 0    | 0         |
| V1 Register           | 0        | 0    | 0         |
| V1 Register Stop      | 0        | 0    | 0         |
| V1 Join Prune         | 0        | 0    | 0         |
| V1 RP Reachability    | 0        | 0    | 0         |
| V1 Assert             | 0        | 0    | 0         |
| V1 Graft              | 0        | 0    | 0         |
| V1 Graft Ack          | 0        | 0    | 0         |
| AutoRP Announce       | 0        | 0    | 0         |
| AutoRP Mapping        | 0        | 0    | 0         |
| AutoRP Unknown type   | 0        |      |           |
| Anycast Register      | 0        | 0    | 0         |
| Anycast Register Stop | 0        | 0    | 0         |

Instance: PIM.master Family: INET6

PIM Interface statistics for ge-0/3/0.0

| PIM Message type      | Received | Sent | Rx errors |
|-----------------------|----------|------|-----------|
| V2 Hello              | 0        | 3    | 0         |
| V2 Register           | 0        | 0    | 0         |
| V2 Register Stop      | 0        | 0    | 0         |
| V2 Join Prune         | 0        | 0    | 0         |
| V2 Bootstrap          | 0        | 0    | 0         |
| V2 Assert             | 0        | 0    | 0         |
| V2 Graft              | 0        | 0    | 0         |
| V2 Graft Ack          | 0        | 0    | 0         |
| V2 Candidate RP       | 0        | 0    | 0         |
| Anycast Register      | 0        | 0    | 0         |
| Anycast Register Stop | 0        | 0    | 0         |

## CHAPTER 220

# Operational Commands (MSDP)

- `clear msdp cache`
- `clear msdp statistics`
- `show msdp`
- `show msdp source`
- `show msdp source-active`
- `show msdp statistics`
- `test msdp`

## clear msdp cache

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>clear msdp cache</code><br><code>&lt;all&gt;</code><br><code>&lt;instance <i>instance-name</i>&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code><br><code>&lt;peer <i>peer-address</i>&gt;</code>                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.<br>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Clear the entries in the Multicast Source Discovery Protocol (MSDP) source-active cache.                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <b>all</b> — Clear all MSDP source-active cache entries in the master instance..<br><br><b>instance <i>instance-name</i></b> —(Optional) Clear entries for a specific MSDP instance.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.<br><br><b>peer <i>peer-address</i></b> —(Optional) Clear the MSDP source-active cache entries learned from a specific peer. |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show msdp source-active on page 5938</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                        |
| <b>List of Sample Output</b>    | <a href="#">clear msdp cache all on page 5932</a>                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                                                         |

### Sample Output

#### clear msdp cache all

```
user@host> clear msdp cache all
```

## clear msdp statistics

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | clear msdp statistics<br><instance <i>instance-name</i> ><br><logical-system (all   <i>logical-system-name</i> )><br><peer <i>peer-address</i> >                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Clear Multicast Source Discovery Protocol (MSDP) peer statistics.                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <p><b>none</b>—Clear MSDP statistics for all peers.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Clear statistics for the specified instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>peer <i>peer-address</i></b>—(Optional) Clear the statistics for the specified peer.</p> |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show msdp statistics on page 5941</a></li> </ul>                                                                                                                                                                                                                                                                                                                  |
| <b>List of Sample Output</b>    | <a href="#">clear msdp statistics on page 5933</a>                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                  |

## Sample Output

### clear msdp statistics

```
user@host> clear msdp statistics
```

## show msdp

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show msdp &lt;brief   detail&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;peer <i>peer-address</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Display Multicast Source Discovery Protocol (MSDP) information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <p><b>none</b>—Display standard MSDP information for all routing instances.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information for the specified instance only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>peer <i>peer-address</i></b>—(Optional) Display information about the specified peer only.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show msdp source on page 5936</a></li> <li>• <a href="#">show msdp source-active on page 5938</a></li> <li>• <a href="#">show msdp statistics on page 5941</a></li> </ul>                                                                                                                                                                                                                                                                                                              |
| <b>List of Sample Output</b>    | <p><a href="#">show msdp on page 5935</a></p> <p><a href="#">show msdp brief on page 5935</a></p> <p><a href="#">show msdp detail on page 5935</a></p>                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Output Fields</b>            | Table 469 describes the output fields for the <b>show msdp</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                             |

Table 469: show msdp Output Fields

| Field Name    | Field Description                                                                        | Level of Output |
|---------------|------------------------------------------------------------------------------------------|-----------------|
| Peer address  | IP address of the peer.                                                                  | All levels      |
| Local address | Local address of the peer.                                                               | All levels      |
| State         | Status of the MSDP connection: <b>Listen</b> , <b>Established</b> , or <b>Inactive</b> . | All levels      |
| Last up/down  | Time at which the most recent peer-state change occurred.                                | All levels      |

Table 469: show msdp Output Fields (*continued*)

| Field Name           | Field Description                                                                                                                                                                   | Level of Output |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Peer-Group           | Peer group name.                                                                                                                                                                    | All levels      |
| SA Count             | Number of source-active cache entries advertised by each peer that were accepted, compared to the number that were received, in the format <i>number-accepted/number-received</i> . | All levels      |
| Peer Connect Retries | Number of peer connection retries.                                                                                                                                                  | detail          |
| State timer expires  | Number of seconds before another message is sent to a peer.                                                                                                                         | detail          |
| Peer Times out       | Number of seconds to wait for a response from the peer before the peer is declared unavailable.                                                                                     | detail          |
| SA accepted          | Number of entries in the source-active cache accepted from the peer.                                                                                                                | detail          |
| SA received          | Number of entries in the source-active cache received by the peer.                                                                                                                  | detail          |

## Sample Output

### show msdp

```

user@host> show msdp
Peer address Local address State Last up/down Peer-Group SA Count
198.32.8.193 198.32.8.195 Established 5d 19:25:44 North23 120/150
198.32.8.194 198.32.8.195 Established 3d 19:27:27 North23 300/345
198.32.8.196 198.32.8.195 Established 5d 19:39:36 North23 10/13
198.32.8.197 198.32.8.195 Established 5d 19:32:27 North23 5/6
198.32.8.198 198.32.8.195 Established 3d 19:33:04 North23 2305/3000

```

### show msdp brief

The output for the **show msdp brief** command is identical to that for the **show msdp** command. For sample output, see [show msdp on page 5935](#).

### show msdp detail

```

user@host> show msdp detail
Peer: 10.255.70.15
Local address: 10.255.70.19
State: Established
Peer Connect Retries: 0
State timer expires: 22
Peer Times out: 49
SA accepted: 0
SA received: 0

```

## show msdp source

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show msdp source</code><br><code>&lt;instance <i>instance-name</i>&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code><br><code>&lt;source-address&gt;</code>                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.<br>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Display multicast sources learned from Multicast Source Discovery Protocol (MSDP).                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <b>none</b> —Display standard MSDP source information for all routing instances.<br><br><b>instance <i>instance-name</i></b> —(Optional) Display information for the specified instance only.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.<br><br><b>source-address</b> —(Optional) IP address and optional prefix length. Display information for the specified source address only. |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show msdp on page 5934</a></li><li>• <a href="#">show msdp source-active on page 5938</a></li><li>• <a href="#">show msdp statistics on page 5941</a></li></ul>                                                                                                                                                                                                                                                                                   |
| <b>List of Sample Output</b>    | <a href="#">show msdp source on page 5937</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                         |



**Output Fields** Table 470 describes the output fields for the **show msdp source** command. Output fields are listed in the approximate order in which they appear.

**Table 470: show msdp source Output Fields**

| Field Name     | Field Description                                                                                                                                                                                                                                                       |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Source address | IP address of the source.                                                                                                                                                                                                                                               |
| /Len           | Length of the prefix for this IP address.                                                                                                                                                                                                                               |
| Type           | Discovery method for this multicast source: <ul style="list-style-type: none"> <li>• <b>Configured</b>—Source-active limit explicitly configured for this source.</li> <li>• <b>Dynamic</b>—Source-active limit established when this source was discovered.</li> </ul> |
| Maximum        | Source-active limit applied to this source.                                                                                                                                                                                                                             |
| Threshold      | Source-active threshold applied to this source.                                                                                                                                                                                                                         |
| Exceeded       | Number of source-active messages received from this source exceeding the established maximum.                                                                                                                                                                           |

## Sample Output

**show msdp source**

```

user@host> show msdp source
Source address /Len Type Maximum Threshold Exceeded
0.0.0.0 /0 Configured 5 none 0
10.1.0.0 /16 Configured 500 none 0
10.1.1.1 /32 Configured 10000 none 0
10.1.1.2 /32 Dynamic 6936 none 0
10.1.5.5 /32 Dynamic 500 none 123
10.2.1.1 /32 Dynamic 2 none 0

```

## show msdp source-active

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show msdp source-active</code><br><code>&lt;brief   detail&gt;</code><br><code>&lt;group <i>group</i>&gt;</code><br><code>&lt;instance <i>instance-name</i>&gt;</code><br><code>&lt;local&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code><br><code>&lt;originator <i>originator</i>&gt;</code><br><code>&lt;peer <i>peer-address</i>&gt;</code><br><code>&lt;source <i>source-address</i>&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.<br>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Display the Multicast Source Discovery Protocol (MSDP) source-active cache.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <b>none</b> —Display standard MSDP source-active cache information for all routing instances.<br><br><b>brief   detail</b> —(Optional) Display the specified level of output.<br><br><b>group <i>group</i></b> —(Optional) Display source-active cache information for the specified group.<br><br><b>instance <i>instance-name</i></b> —(Optional) Display information for the specified instance.<br><br><b>local</b> —(Optional) Display all source-active caches originated by this router.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.<br><br><b>originator <i>originator</i></b> —(Optional) Display information about the peer that originated the source-active cache entries.<br><br><b>peer <i>peer-address</i></b> —(Optional) Display the source-active cache of the specified peer.<br><br><b>source <i>source-address</i></b> —(Optional) Display the source-active cache of the specified source. |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show msdp on page 5934</a></li><li>• <a href="#">show msdp source on page 5936</a></li><li>• <a href="#">show msdp statistics on page 5941</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>List of Sample Output</b>    | <a href="#">show msdp source-active on page 5939</a><br><a href="#">show msdp source-active brief on page 5940</a><br><a href="#">show msdp source-active detail on page 5940</a><br><a href="#">show msdp source-active source on page 5940</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

**Output Fields** Table 471 describes the output fields for the **show msdp source-active** command. Output fields are listed in the approximate order in which they appear.

**Table 471: show msdp source-active Output Fields**

| Field Name                              | Field Description                                                                                                                                          |
|-----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Global active source limit exceeded     | Number of times all peers have exceeded configured active source limits.                                                                                   |
| Global active source limit maximum      | Configured number of active source messages accepted by the device.                                                                                        |
| Global active source limit threshold    | Configured threshold for applying random early discard (RED) to drop some but not all MSDP active source messages.                                         |
| Global active source limit log-warning  | Threshold at which a warning message is logged (percentage of the number of active source messages accepted by the device).                                |
| Global active source limit log interval | Time (in seconds) between consecutive log messages.                                                                                                        |
| Group address                           | Multicast address of the group.                                                                                                                            |
| Source address                          | IP address of the source.                                                                                                                                  |
| Peer address                            | IP address of the peer.                                                                                                                                    |
| Originator                              | Router ID configured on the source of the rendezvous point (RP) that originated the message, or the loopback address when the router ID is not configured. |
| Flags                                   | Flags: <b>Accept</b> , <b>Reject</b> , or <b>Filtered</b> .                                                                                                |

## Sample Output

### show msdp source-active

```

user@host> show msdp source-active
Group address Source address Peer address Originator Flags
230.0.0.0 192.168.195.46 local 10.255.14.30 Accept
230.0.0.1 192.168.195.46 local 10.255.14.30 Accept
230.0.0.2 192.168.195.46 local 10.255.14.30 Accept
230.0.0.3 192.168.195.46 local 10.255.14.30 Accept
230.0.0.4 192.168.195.46 local 10.255.14.30 Accept

```

### show msdp source-active brief

The output for the **show msdp source-active brief** command is identical to that for the **show msdp source-active** command. For sample output, see [show msdp source-active on page 5939](#).

### show msdp source-active detail

The output for the **show msdp source-active detail** command is identical to that for the **show msdp source-active** command. For sample output, see [show msdp source-active on page 5939](#).

### show msdp source-active source

```
user@host> show msdp source-active source 192.168.215.246
```

```
Global active source limit exceeded: 0
```

```
Global active source limit maximum: 25000
```

```
Global active source limit threshold: 24000
```

```
Global active source limit log-warning: 100
```

```
Global active source limit log interval: 0
```

| Group address | Source address  | Peer address   | Originator     | Flags  |
|---------------|-----------------|----------------|----------------|--------|
| 226.2.2.1     | 192.168.215.246 | 10.255.182.140 | 10.255.182.140 | Accept |
| 226.2.2.3     | 192.168.215.246 | 10.255.182.140 | 10.255.182.140 | Accept |
| 226.2.2.4     | 192.168.215.246 | 10.255.182.140 | 10.255.182.140 | Accept |
| 226.2.2.5     | 192.168.215.246 | 10.255.182.140 | 10.255.182.140 | Accept |
| 226.2.2.7     | 192.168.215.246 | 10.255.182.140 | 10.255.182.140 | Accept |
| 226.2.2.10    | 192.168.215.246 | 10.255.182.140 | 10.255.182.140 | Accept |
| 226.2.2.11    | 192.168.215.246 | 10.255.182.140 | 10.255.182.140 | Accept |
| 226.2.2.13    | 192.168.215.246 | 10.255.182.140 | 10.255.182.140 | Accept |
| 226.2.2.14    | 192.168.215.246 | 10.255.182.140 | 10.255.182.140 | Accept |
| 226.2.2.15    | 192.168.215.246 | 10.255.182.140 | 10.255.182.140 | Accept |

## show msdp statistics

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show msdp statistics<br><instance <i>instance-name</i> ><br><logical-system (all   <i>logical-system-name</i> )><br><peer <i>peer-address</i> >                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.<br>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Display statistics about Multicast Source Discovery Protocol (MSDP) peers.                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <b>none</b> —Display statistics about all MSDP peers for all routing instances.<br><br><b>instance <i>instance-name</i></b> —(Optional) Display statistics about a specific MSDP instance.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.<br><br><b>peer <i>peer-address</i></b> —(Optional) Display statistics about a particular MSDP peer. |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">clear msdp statistics on page 5933</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                        |
| <b>List of Sample Output</b>    | <a href="#">show msdp statistics on page 5943</a><br><a href="#">show msdp statistics peer on page 5943</a>                                                                                                                                                                                                                                                                                                                                                 |
| <b>Output Fields</b>            | Table 472 describes the output fields for the <b>show msdp statistics</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                  |

**Table 472: show msdp statistics Output Fields**

| Field Name                              | Field Description                                                                                                           |
|-----------------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| Global active source limit exceeded     | Number of times all peers have exceeded configured active source limits.                                                    |
| Global active source limit maximum      | Configured number of active source messages accepted by the device.                                                         |
| Global active source limit threshold    | Configured threshold for applying random early discard (RED) to drop some but not all MSDP active source messages.          |
| Global active source limit log-warning  | Threshold at which a warning message is logged (percentage of the number of active source messages accepted by the device). |
| Global active source limit log interval | Time (in seconds) between consecutive log messages.                                                                         |

Table 472: show msdp statistics Output Fields (*continued*)

| Field Name                                        | Field Description                                                                                                                                                                                     |
|---------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Peer</b>                                       | Address of peer.                                                                                                                                                                                      |
| <b>Last State Change</b>                          | How long ago the peer state changed.                                                                                                                                                                  |
| <b>Last message received from the peer</b>        | How long ago the last message was received from the peer.                                                                                                                                             |
| <b>RPF Failures</b>                               | Number of reverse path forwarding (RPF) failures.                                                                                                                                                     |
| <b>Remote Closes</b>                              | Number of times the remote peer closed.                                                                                                                                                               |
| <b>Peer Timeouts</b>                              | Number of peer timeouts.                                                                                                                                                                              |
| <b>SA messages sent</b>                           | Number of source-active messages sent.                                                                                                                                                                |
| <b>SA messages received</b>                       | Number of source-active messages received.                                                                                                                                                            |
| <b>SA request messages sent</b>                   | Number of source-active request messages sent.                                                                                                                                                        |
| <b>SA request messages received</b>               | Number of source-active request messages received.                                                                                                                                                    |
| <b>SA response messages sent</b>                  | Number of source-active response messages sent.                                                                                                                                                       |
| <b>SA response messages received</b>              | Number of source-active response messages received.                                                                                                                                                   |
| <b>SA messages with zero Entry Count received</b> | Entry Count is a field within SA message that defines how many source/group tuples are present in the SA message. The counter is incremented each time an SA with an Entry Count of zero is received. |
| <b>Active source exceeded</b>                     | Number of times this peer has exceeded configured source-active limits.                                                                                                                               |
| <b>Active source Maximum</b>                      | Configured number of active source messages accepted by this peer.                                                                                                                                    |
| <b>Active source threshold</b>                    | Configured threshold on this peer for applying random early discard (RED) to drop some but not all MSDP active source messages.                                                                       |
| <b>Active source log-warning</b>                  | Configured threshold on this peer at which a warning message is logged (percentage of the number of active source messages accepted by the device).                                                   |
| <b>Active source log-interval</b>                 | Time (in seconds) between consecutive log messages on this peer.                                                                                                                                      |
| <b>Keepalive messages sent</b>                    | Number of keepalive messages sent.                                                                                                                                                                    |

Table 472: show msdp statistics Output Fields (*continued*)

| Field Name                  | Field Description                      |
|-----------------------------|----------------------------------------|
| Keepalive messages received | Number of keepalive messages received. |
| Unknown messages received   | Number of unknown messages received.   |
| Error messages received     | Number of error messages received.     |

## Sample Output

### show msdp statistics

```

user@host> show msdp statistics
Global active source limit exceeded: 0
Global active source limit maximum: 10
Global active source limit threshold: 8
Global active source limit log-warning: 60
Global active source limit log interval: 60

Peer: 10.255.245.39
Last State Change: 11:54:49 (00:24:59)
Last message received from peer: 11:53:32 (00:26:16)
RPF Failures: 0
Remote Closes: 0
Peer Timeouts: 0
SA messages sent: 376
SA messages received: 459
SA messages with zero Entry Count received: 0
SA request messages sent: 0
SA request messages received: 0
SA response messages sent: 0
SA response messages received: 0
Active source exceeded: 0
Active source Maximum: 10
Active source threshold: 8
Active source log-warning: 60
Active source log-interval 120
Keepalive messages sent: 17
Keepalive messages received: 19
Unknown messages received: 0
Error messages received: 0

```

### show msdp statistics peer

```

user@host> show msdp statistics peer 10.255.182.140
Peer: 10.255.182.140
 Last State Change: 8:19:23 (00:01:08)
 Last message received from peer: 8:20:05 (00:00:26)
 RPF Failures: 0
 Remote Closes: 0
 Peer Timeouts: 0
 SA messages sent: 17
 SA messages received: 16
 SA request messages sent: 0
 SA request messages received: 0

```

SA response messages sent: 0  
SA response messages received: 0  
Active source exceeded: 20  
Active source Maximum: 10  
Active source threshold: 8  
Active source log-warning: 60  
Active source log-interval: 120  
Keepalive messages sent: 0  
Keepalive messages received: 0  
Unknown messages received: 0  
Error messages received: 0



## test msdp

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | test msdp (dependent-peers <i>prefix</i>   rpf-peer <i>originator</i> )<br><instance <i>instance-name</i> ><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.<br>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Find Multicast Source Discovery Protocol (MSDP) peers.                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><b>dependent-peers <i>prefix</i></b>—Find downstream dependent MSDP peers.</p> <p><b>rpf-peer <i>originator</i></b>—Find the MSDP reverse-path-forwarding (RPF) peer for the originator.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Find MDSP peers for the specified routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>List of Sample Output</b>    | <a href="#">test msdp dependent-peers on page 5945</a>                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                                                  |

## Sample Output

### test msdp dependent-peers

```
user@host> test msdp dependent-peers 10.0.0.1/24
```



# Security Feature Guide for QFX10000 Switches



## PART 78

# Firewall Filters

- [Configuring Firewall Filters on page 5951](#)



# Configuring Firewall Filters

- [Overview of Firewall Filters on page 5951](#)
- [Understanding How Firewall Filters Are Evaluated on page 5954](#)
- [Understanding How Firewall Filters Control Packet Flows on page 5956](#)
- [Understanding Firewall Filter Match Conditions on page 5957](#)
- [Firewall Filter Match Conditions and Actions for QFX10000 Switches on page 5960](#)
- [Understanding How a Firewall Filter Tests a Protocol on page 5971](#)
- [Understanding Firewall Filter Planning on page 5972](#)
- [Planning the Number of Firewall Filters to Create on page 5973](#)
- [Understanding Firewall Filter Processing Points for Bridged and Routed Packets on page 5978](#)
- [Configuring Firewall Filters on page 5978](#)
- [Applying Firewall Filters to Interfaces on page 5982](#)
- [Understanding Filter-Based Forwarding on page 5983](#)
- [Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on page 5983](#)
- [Monitoring Firewall Filter Traffic on page 5987](#)
- [Verifying That Firewall Filters Are Operational on page 5988](#)
- [Troubleshooting Firewall Filters on page 5989](#)

## Overview of Firewall Filters

---

Firewall filters provide rules that define whether to accept or discard packets that are transiting an interface. If a packet is accepted, you can configure additional actions to perform on the packet, such as class-of-service (CoS) marking (grouping similar types of traffic together and treating each type of traffic as a class with its own level of service priority) and traffic policing (controlling the maximum rate of traffic sent or received). You configure firewall filters to determine whether to accept or discard a packet before it enters or exits any of these:

- Port
- VLAN

- Layer 3 (routed) interface
- Routed VLAN interface (RVI)

An *ingress* firewall filter is applied to packets that are entering an interface or VLAN, and an *egress* firewall filter is applied to packets that are exiting an interface or VLAN.



**NOTE:** Firewall filters are sometimes called *access control lists (ACLs)*.

- [Firewall Filter Types on page 5952](#)
- [Firewall Filter Components on page 5953](#)
- [Firewall Filter Processing on page 5953](#)
- [How Many Filters Are Supported? on page 5953](#)

## Firewall Filter Types

The following firewall filter types are supported:

- Port (Layer 2) firewall filter—Port firewall filters apply to Layer 2 traffic transiting system ports.
- VLAN firewall filter—VLAN firewall filters provide access control for packets that enter a VLAN, are bridged within a VLAN, or leave a VLAN.
- Router (Layer 3) firewall filter—You can apply a router firewall filter in both ingress and egress directions on IPv4 or IPv6 Layer 3 (routed) interfaces, routed VLAN interfaces (RVI) and a loopback interface, which filters traffic sent to the switch itself or generated by the switch. (You apply a filter to a loopback interface in the input direction to protect the switch from unwanted traffic. You also might want to apply a filter to a loopback interface in the output direction so that you can set the forwarding class and DSCP bit value for packets that originate on the switch itself. This feature gives you very fine control over the classification of CPU generated packets. For example, you might want to assign different DSCP values and forwarding classes to traffic generated by different routing protocols so the traffic for those protocols can be treated in a differentiated manner by other devices. You can apply a filter to a loopback interface in the output direction starting with Junos OS 13.2X51-D15.)



**NOTE:** You can apply a firewall filter to a management interface (for example, `me0`) on a QFX and EX4600 standalone switch. You cannot apply a firewall filter to a management interface on a QFX3000-G or QFX3000-M system.

- MPLS filter—You can apply a firewall filter to an MPLS interface

To apply a firewall filter:

1. Configure the firewall filter.
2. Apply the firewall filter to a port, VLAN, or router interface.





**NOTE:** You can apply only one firewall filter to a port, VLAN, or interface for a given direction. For example, for interface `ge-0/0/6.0`, you can apply one filter for the ingress direction and one for the egress direction.

## Firewall Filter Components

In a firewall filter, you first define the family address type (ethernet-switching, inet (for IPv4), inet6 (for IPv6), or mpls) and then define one or more terms that specify the filtering criteria and the action to take if a match occurs.

Each term consists of the following components:

- **Match conditions**—Specify values that a packet must contain to be considered a match. You can specify values for most fields in the IP, TCP, UDP, or ICMP headers. You can also match on interface names.
- **Action**—Specifies what to do if a packet matches the match conditions. A filter can accept, discard, or reject a matching packet and then perform additional actions, such as counting, classifying, and policing. If no action is specified for a term, the default is to accept the matching packet.

## Firewall Filter Processing

If there are multiple terms in a filter, the order of the terms is important. If a packet matches the first term, the switch executes the action defined by that term, and no other terms are evaluated. If the switch does not find a match between the packet and the first term, it compares the packet to the next term. If no match occurs between the packet and the second term, the system continues to compare the packet to each successive term in the filter until a match is found. If the packet does not match any terms in the filter, the switch discards the packet by default.

## How Many Filters Are Supported?

QFX10000 switches support 8K firewall filters and 64K firewall filter terms.

QFX3500, QFX3600, QFX5100, QFX5200, and EX4600 switches, QFabric Node devices, and VCF members support the maximum numbers of firewall filter terms per type of attachment point shown in [Table 473](#).

**Table 473: Supported Firewall Filter Numbers for Specific Switches**

| Filter Type | QFX3500, QFX3600 | QFX5100, EX4600 | QFX5200 |
|-------------|------------------|-----------------|---------|
| Ingress     | 768              | 1536            | 768     |
| Egress      | 1024             | 1024            | 1024    |

These totals are applied in aggregate. For example, on the QFX3500 and QFX3600 you can apply a total of 768 terms in all your port filters, Layer 3 filters, and VLAN filters that are applied in the input direction and 1024 terms in port filters, Layer 3 filters, and VLAN filters that are applied in the output direction. The actual number of filters that these

switches will support depends on how the filters are stored in ternary content addressable memory (TCAM). See [“Planning the Number of Firewall Filters to Create” on page 5973](#) for detailed information about this topic.

**Related  
Documentation**

- [Understanding Firewall Filter Planning on page 5972](#)
- [Understanding Firewall Filter Processing Points for Bridged and Routed Packets on page 5978](#)
- [Understanding How Firewall Filters Are Evaluated on page 5954](#)
- [Understanding Firewall Filter Match Conditions on page 5957](#)
- [Overview of Policers on page 5999](#)
- [Configuring Firewall Filters on page 5978](#)

---

## Understanding How Firewall Filters Are Evaluated

---

A firewall filter consists of one or more terms, and the order of the terms within a filter is important. Before you configure firewall filters, you should understand how switches evaluate the terms within a filter and how packets are evaluated against the terms.

When a firewall filter consists of a single term, the filter is evaluated as follows:

- If the packet matches all the conditions, the action in the **then** statement is taken.
- If the packet matches all the conditions, and no action is specified in the **then** statement, the default action **accept** is taken.
- If the packet does not match all the conditions, the switch discards it.

When a firewall filter consists of more than one term, the filter is evaluated sequentially:

1. The packet is evaluated against the conditions in the **from** statement in the first term.
2. If the packet matches all the conditions in the term, the action in the **then** statement is taken and the evaluation ends. Subsequent terms in the filter are not evaluated.
3. If the packet does not match all the conditions in the term, the packet is evaluated against the conditions in the **from** statement in the second term.

This process continues until the packet matches all the conditions in the **from** statement in one of the subsequent terms or there are no more terms in the filter.

4. If a packet passes through all the terms in the filter without a match, the switch discards it.




**NOTE:** The order of conditions in a **from** statement is not important because a packet must match all the conditions to be considered a match.

---

Figure 196 shows how switches evaluate the terms within a firewall filter.

Figure 196: Evaluation of Terms Within a Firewall Filter



If you do not include a **from** statement in a term, all packets will match the term and be processed by the **then** statement. If a term does not contain a **then** statement or if an action has not been configured in the **then** statement, the term accepts any matching packets.

Every firewall filter contains an implicit **deny** statement at the end of the filter, which is equivalent to the following explicit filter term:

```
term implicit-rule {
 then discard;
}
```

Consequently, a packet that does not match any of the terms in a firewall filter is discarded. If you configure a filter that has no terms, all packets that pass through the filter are discarded.



**NOTE:** Firewall filtering is supported on packets that are at least 64 bytes long.

#### Related Documentation

- [Understanding Firewall Filter Match Conditions on page 5957](#)
- [Overview of Policers on page 5999](#)
- [Configuring Firewall Filters on page 5978](#)

## Understanding How Firewall Filters Control Packet Flows

A switch supports firewall filters that allow you to control flows of data packets and local packets. *Data packets* transit a switch as they are forwarded from a source to a destination. *Local packets* are destined for or sent by a Routing Engine (they do not transit a switch). Local packets usually contain routing protocol data, data for IP services such as Telnet or SSH, or data for administrative protocols such as the Internet Control Message Protocol (ICMP).


Firewall filters affect packet flows entering into or exiting from a switch as follows:

- Ingress firewall filters affect the flow of data packets that are received on switch interfaces. When a switch receives a data packet, the Packet Forwarding Engine in the system that contains the ingress interface determines where to forward the packet by looking in its Layer 2 or Layer 3 forwarding table for the best route to the destination. Data packets are forwarded to an egress interface. Locally destined packets are forwarded to the Routing Engine.
- Egress firewall filters affect data packets that are transiting a switch but do not affect packets sent by the Routing Engine. These filters are applied by the Packet Forwarding Engine in the system that contains the egress interface.

Figure 197 illustrates the application of ingress and egress firewall filters to control the flow of packets through a switch:

1. Ingress firewall filter applied to locally destined packets that are received on switch interfaces and are destined for the Routing Engine.
2. Ingress firewall filter applied to data packets that are received on switch interfaces and will transit the switch.
3. Egress firewall filter applied to data packets that are transiting the switch.

Figure 197: Application of Firewall Filters to Control Packet Flow



**Related Documentation**

- [Understanding Firewall Filter Processing Points for Bridged and Routed Packets on page 5978](#)
- [Understanding How Firewall Filters Are Evaluated on page 5954](#)
- [Configuring Firewall Filters on page 5978](#)

## Understanding Firewall Filter Match Conditions

Before you define terms for firewall filters, you must understand how the conditions in a term are handled and how to specify interface, numeric, address, and bit-field filter match conditions to achieve the desired filter results.

- [Filter Match Conditions on page 5957](#)
- [Numeric Filter Match Conditions on page 5957](#)
- [Interface Filter Match Conditions on page 5958](#)
- [IP Address Filter Match Conditions on page 5958](#)
- [MAC Address Filter Match Conditions on page 5959](#)
- [Bit-Field Filter Match Conditions on page 5959](#)

### Filter Match Conditions

In the **from** statement of a firewall filter term, you specify the conditions that the packet must match for the action in the **then** statement to be taken. All conditions must match for the action to be implemented. The order in which you specify match conditions is not important, because a packet must match all the conditions in a term for a match to occur.

If you specify multiple values for the same condition, a match on any one of those values matches that condition. For example, if you specify multiple IP source addresses using the **source-address** statement, a packet that contains any one of those IP source addresses matches the condition. In some cases you can specify multiple values for the same condition by enclosing the possible values in square brackets, as in:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set protocol (icmp | udp)
```

In other cases you must enter multiple statements, as in:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set source-address 10.1.1.1
user@switch# set source-address 10.1.1.2
```

If you specify no match conditions in a term, that term matches all packets.



**NOTE:** Unlike traditional Junos OS firewall filters, you cannot use **except** in a condition statement to negate the condition.

### Numeric Filter Match Conditions

You can specify numeric filter match conditions that are identified by a numeric value, such as port and protocol numbers. For numeric filter match conditions, you specify the

condition and a single value that a field in a packet must contain to be considered a match.

You can specify the numeric value in one of the following ways:

- Single number—A match occurs if the value of the field matches the number. For example, to match Telnet traffic:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set source-port 23
```

- Text synonym for a single number—A match occurs if the value of the field matches the number that corresponds to the synonym. For example, to match Telnet traffic:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set source-port telnet
```

- To specify multiple values for the same match condition in a filter term, enter each value in its own match statement. For example, a match occurs in the following term if the value of the source port in the packet is 22 or 23.

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set source-port 22
user@switch# set source-port 23
```

## Interface Filter Match Conditions

You can specify an interface filter match condition to match an interface on which a packet is received or transmitted. For example, if you apply a filter to a VLAN you might want the filter to match on some interfaces that participate in the VLAN and not match on other interfaces in the VLAN. When you specify the name of the interface, you must include a logical unit.

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set interface ge-0/0/6.0
```

In this example, the final character (0) specifies the logical unit. You can include the wildcard (\*) as part of the interface name. For example:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set interface ge-0/*/6.0
user@switch# set interface ge-0/1/*0
user@switch# set interface ge-0/0/6.*
```

Note that you must specify a value or a wildcard for the logical unit.

## IP Address Filter Match Conditions

You can specify an address filter match condition to match an IP source or destination address or prefix in a packet. Specify the address or prefix type and the address or prefix itself. For example:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set destination-address 10.2.1.0/24;
```

If you omit the prefix length, it defaults to /32. For example:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set destination-address 10
[edit firewall family family-name filter filter-name term term-name from]
user@switch# show
destination-address {
```

```
10.0.0.0/32;
}
```

To specify more than one IP address or prefix in a filter term, enter each address or prefix in its own match statement. For example, a match occurs in the following term if the source address of a packet matches either of the following prefixes:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set source-address 10.1.0.0/16
user@switch# set source-address 10.2.0.0/16
```

## MAC Address Filter Match Conditions

You can specify a MAC address filter match condition to match a source or destination MAC address. You specify the address type and value that a packet must contain to be considered a match.

You can specify the MAC address as six hexadecimal bytes in any of the following formats:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set destination-mac-address 00:11:22:33:44:55
```

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set destination-mac-address 0011.2233.4455
```

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set destination-mac-address 001122334455
```

Regardless of the formats you use, the system resolves the address to the standard format, in this case 00:11:22:33:44:55.

To specify more than one MAC address in a filter term, enter each MAC address in its own match statement. For example, a match occurs in the following term if the value of the MAC source address matches either of the following addresses:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set source-mac-address 00:11:22:33:44:55
user@switch# set source-mac-address 00:11:22:33:20:15
```

## Bit-Field Filter Match Conditions

You can specify bit-field filter match conditions to match particular bits within certain fields in Ethernet frames and IP, TCP, UDP, and ICMP headers. You usually specify the field and the bit within the field that must be set in a packet to be considered a match.

In most cases you can use a keyword to specify the bit you want to match on. For example, to match on a TCP SYN packet you can enter **syn**, as in:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set tcp-flags syn
```

You can also enter **0x02** because the SYN bit is the third least-significant bit of the 8-bit tcp-flags field:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set tcp-flags 0x02
```

To match multiple bit-field values, use the logical operators, which are described in [Table 474](#). The operators are listed in order from highest precedence to lowest precedence. Operations are evaluated from left to right.

Table 474: Actions for Firewall Filters

| Logical Operators | Description |
|-------------------|-------------|
| !                 | Negation    |
| &                 | Logical AND |
|                   | Logical OR  |

If you use a logical operator, enclose the values in quotation marks and do not include any spaces. For example, the following statement matches the second packet of a TCP handshake:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set tcp-flags "syn&ack"
```

To negate a match, precede the value with an exclamation point. For example, the following statement matches only the initial packet of a TCP handshake:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set tcp-flags "syn!ack"
```

You can use text synonyms to specify some common bit-field matches. For example, the following statement also matches the initial packet of a TCP handshake:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set tcp-initial
```

#### Related Documentation

- [Understanding How a Firewall Filter Tests a Protocol on page 5971](#)
- [Firewall Filter Match Conditions and Actions](#)
- [Configuring Firewall Filters on page 5978](#)

## Firewall Filter Match Conditions and Actions for QFX10000 Switches

Each term in a firewall filter consists of *match conditions* and an *action*. Match conditions are the fields and values that a packet must contain to be considered a match. You can define single or multiple match conditions in *match statements*. You can also include no match statement, in which case the term matches all packets.

When a packet matches a filter, a switch takes the action specified in the term. In addition, you can specify action modifiers to count, mirror, rate-limit, and classify packets. If no match conditions are specified for the term, the switch accepts the packet by default.

This topic describes the various match conditions, actions, and action modifiers that you can define in firewall filters on QFX10000 switches. For similar information about other QFX switches, see *Firewall Filter Match Conditions and Actions*.

- [Table 475](#) describes the match conditions you can specify when configuring a firewall filter. Some of the numeric range and bit-field match conditions allow you to specify a text synonym. To see a list of all the synonyms for a match condition, type ? at the appropriate place in a statement.



- [Table 476](#) shows the actions that you can specify in a term.
- [Table 477](#) shows the action modifiers you can use to count, mirror, rate-limit, and classify packets.



**NOTE:** On QFX10000 switches, do not combine match conditions for Layer 2 and any other layer in a family ethernet-switching filter. (For example, do not include conditions that match MAC addresses and IP addresses in the same filter.) If you do so, the filter will commit successfully but will not work. You will also see the following log message: *L2 filter **filter-name** doesn't support mixed L2 and L3/L4 match conditions. Please re-config.*

Table 475: Supported Match Conditions for Firewall Filters on QFX10000 Switches

| Match Condition                                      | Description                                                                       | Direction and Interface                                                                                                       |
|------------------------------------------------------|-----------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>destination-address</b><br><i>ip-address</i>      | IP destination address field, which is the address of the final destination node. | Ingress Pv4 (inet) interfaces and IPv6 (inet6) interfaces.<br><br>Egress IPv4 (inet) interfaces, and IPv6 (inet6) interfaces. |
| <b>destination-mac-address</b><br><i>mac-address</i> | Destination media access control (MAC) address of the packet.                     | Ingress ports and VLANs.<br><br>Egress ports and VLANs.                                                                       |

Table 475: Supported Match Conditions for Firewall Filters on QFX10000 Switches (*continued*)

| Match Condition                                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Direction and Interface                                                                                                                                            |
|---------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>destination-port <i>value</i></b>              | <p>TCP or UDP destination port field. Typically, you specify this match in conjunction with the <b>protocol</b> match statement. For the following well-known ports you can specify text synonyms (the port numbers are also listed):</p> <p><b>afs</b> (1483), <b>bgp</b> (179), <b>biff</b> (512), <b>bootpc</b> (68), <b>bootps</b> (67),</p> <p><b>cmd</b> (514), <b>cvspserver</b> (2401),</p> <p><b>dhcp</b> (67), <b>domain</b> (53),</p> <p><b>eklogin</b> (2105), <b>ekshell</b> (2106), <b>exec</b> (512),</p> <p><b>finger</b> (79), <b>ftp</b> (21), <b>ftp-data</b> (20),</p> <p><b>http</b> (80), <b>https</b> (443),</p> <p><b>ident</b> (113), <b>imap</b> (143),</p> <p><b>kerberos-sec</b> (88), <b>klogin</b> (543), <b>kpasswd</b> (761), <b>krb-prop</b> (754), <b>krbupdate</b> (760), <b>kshell</b> (544),</p> <p><b>ldap</b> (389), <b>login</b> (513),</p> <p><b>mobileip-agent</b> (434), <b>mobileip-mn</b> (435), <b>msdp</b> (639),</p> <p><b>netbios-dgm</b> (138), <b>netbios-ns</b> (137), <b>netbios-ssn</b> (139), <b>nfsd</b> (2049), <b>nntp</b> (119), <b>ntalk</b> (518), <b>ntp</b> (123),</p> <p><b>pop3</b> (110), <b>pptp</b> (1723), <b>printer</b> (515),</p> <p><b>radacct</b> (1813), <b>radius</b> (1812), <b>rip</b> (520), <b>rkinit</b> (2108),</p> <p><b>smtp</b> (25), <b>snmp</b> (161), <b>snmptrap</b> (162), <b>snpp</b> (444), <b>socks</b> (1080), <b>ssh</b> (22), <b>sunrpc</b> (111), <b>syslog</b> (514),</p> <p><b>tacacs-ds</b> (65), <b>talk</b> (517), <b>telnet</b> (23), <b>tftp</b> (69), <b>timed</b> (525),</p> <p><b>who</b> (513),</p> <p><b>xmcp</b> (177),</p> <p><b>zephyr-clt</b> (2103), <b>zephyr-hm</b> (2104)</p> | <p>Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.</p> <p>Egress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.</p> |
| <b>destination-prefix-list <i>prefix-list</i></b> | <p>IP destination prefix list field. You can define a list of IP address prefixes under a prefix-list alias for frequent use. Define this list at the <b>[edit policy-options]</b> hierarchy level.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <p>Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.</p> <p>Egress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.</p> |

Table 475: Supported Match Conditions for Firewall Filters on QFX10000 Switches (*continued*)

| Match Condition         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Direction and Interface                                                                                          |
|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------|
| <b>dscp value</b>       | <p>Differentiated Services code point (DSCP). The DiffServ protocol uses the type-of-service (ToS) byte in the IP header. The most-significant 6 bits of this byte form the DSCP.</p> <p>You can specify DSCP in hexadecimal, binary, or decimal form.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed):</p> <ul style="list-style-type: none"> <li>• <b>be</b>—best effort (default)</li> <li>• <b>ef (46)</b>—as defined in <a href="#">RFC 3246</a>, <i>An Expedited Forwarding PHB</i>.</li> <li>• <b>af11 (10), af12 (12), af13 (14); af21 (18), af22 (20), af23 (22); af31 (26), af32 (28), af33 (30); af41 (34), af42 (36), af43 (38)</b></li> </ul> <p>These four classes, with three drop precedences in each class, for a total of 12 code points, are defined in <a href="#">RFC 2597</a>, <i>Assured Forwarding PHB</i>.</p> <ul style="list-style-type: none"> <li>• <b>cs0, cs1, cs2, cs3, cs4, cs5, cs6, cs7, cs5</b></li> </ul>                                                                                                                                   | <p>Ingress ports, VLANs, and IPv4 (inet) interfaces.</p> <p>Egress ports, VLANs, and IPv4 (inet) interfaces.</p> |
| <b>ether-type value</b> | <p>Ethernet type field of a packet. The EtherType value specifies what protocol is being transported in the Ethernet frame. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed):</p> <ul style="list-style-type: none"> <li>• <b>aarp (0x80F3)</b>—EtherType value AARP</li> <li>• <b>appletalk (0x809B)</b>—EtherType value AppleTalk</li> <li>• <b>arp (0x0806)</b>—EtherType value ARP</li> <li>• <b>fcoe (0x8906)</b>—EtherType value FCoE</li> <li>• <b>fip (0x8914)</b>—EtherType value FIP</li> <li>• <b>ipv4 (0x0800)</b>—EtherType value IPv4</li> <li>• <b>ipv6 (0x08DD)</b>—EtherType value IPv6</li> <li>• <b>mpls-multicast (0x8848)</b>—EtherType value MPLS multicast</li> <li>• <b>mpls-unicast (0x8847)</b>—EtherType value MPLS unicast</li> <li>• <b>oam (0x88A8)</b>—EtherType value OAM</li> <li>• <b>ppp (0x880B)</b>—EtherType value PPP</li> <li>• <b>pppoe-discovery (0x8863)</b>—EtherType value PPPoE Discovery Stage</li> <li>• <b>pppoe-session (0x8864)</b>—EtherType value PPPoE Session Stage</li> <li>• <b>sna (0x80D5)</b>—EtherType value SNA</li> </ul> | <p>Ingress ports and VLANs.</p> <p>Egress ports and VLANs.</p>                                                   |

Table 475: Supported Match Conditions for Firewall Filters on QFX10000 Switches (*continued*)

| Match Condition             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Direction and Interface                                                                                                                                           |
|-----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>fragment-flags value</b> | <p>IP fragmentation flags. In place of the numeric value, you can specify one of the following text synonyms (the hexadecimal values are also listed):</p> <ul style="list-style-type: none"> <li>• <b>is-fragment</b></li> <li>• <b>dont-fragment (0x4000)</b></li> <li>• <b>more-fragments (0x2000)</b></li> <li>• <b>reserved (0x8000)</b></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Ingress ports, VLANs, and IPv4 (inet) interfaces.                                                                                                                 |
| <b>hop-limit value</b>      | <p>Match the specified hop limit or set of hop limits. Specify a single value or a range of values from 0 through 255.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Ingress and egress IPv6 (inet6) interfaces.                                                                                                                       |
| <b>icmp-code value</b>      | <p>ICMP code field. Because the meaning of the value depends upon the associated <b>icmp-type</b>, you must specify a value for <b>icmp-type</b> along with a value for <b>icmp-code</b>. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed). The keywords are grouped by the ICMP type with which they are associated:</p> <ul style="list-style-type: none"> <li>• <i>IPv4</i>: parameter-problem—<b>ip-header-bad (0)</b>, <b>required-option-missing (1)</b></li> <li>• <i>IPv6</i>: parameter-problem—<b>ip6-header-bad (0)</b>, <b>unrecognized-next-header (1)</b>, <b>unrecognized-option (2)</b></li> <li>• <b>redirect—redirect-for-network (0)</b>, <b>redirect-for-host (1)</b>, <b>redirect-for-tos-and-net (2)</b>, <b>redirect-for-tos-and-host (3)</b></li> <li>• <b>time-exceeded—ttl-eq-zero-during-reassembly (1)</b>, <b>ttl-eq-zero-during-transit (0)</b></li> <li>• <i>IPv4</i>: <b>unreachable—network-unreachable (0)</b>, <b>host-unreachable (1)</b>, <b>protocol-unreachable (2)</b>, <b>port-unreachable (3)</b>, <b>fragmentation-needed (4)</b>, <b>source-route-failed (5)</b>, <b>destination-network-unknown (6)</b>, <b>destination-host-unknown (7)</b>, <b>source-host-isolated (8)</b>, <b>destination-network-prohibited (9)</b>, <b>destination-host-prohibited (10)</b>, <b>network-unreachable-for-TOS (11)</b>, <b>host-unreachable-for-TOS (12)</b>, <b>communication-prohibited-by-filtering (13)</b>, <b>host-precedence-violation (14)</b>, <b>precedence-cutoff-in-effect (15)</b></li> <li>• <i>IPv6</i>: <b>unreachable—address-unreachable (3)</b>, <b>administratively-prohibited (1)</b>, <b>no-route-to-destination (0)</b>, <b>port-unreachable (4)</b></li> </ul> | <p>Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.</p> <p>Egress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces</p> |

Table 475: Supported Match Conditions for Firewall Filters on QFX10000 Switches (*continued*)

| Match Condition                                    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Direction and Interface                                                                                                                                            |
|----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>icmp-type</b> <i>value</i>                      | <p>ICMP message type field. Typically, you specify this match in conjunction with the <b>protocol</b> match statement to determine which protocol is being used on the port. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed):</p> <p><i>IPv4</i>: <b>echo-reply</b> (0), <b>destination unreachable</b> (3), <b>source-quench</b> (4), <b>redirect</b> (5), <b>echo-request</b> (8), <b>IPv4 (inet)-advertisement</b> (9), <b>IPv4 (inet)-solicit</b> (10), <b>time-exceeded</b> (11), <b>parameter-problem</b> (12), <b>timestamp</b> (13), <b>timestamp-reply</b> (14), <b>info-request</b> (15), <b>info-reply</b> (16), <b>mask-request</b> (17), <b>mask-reply</b> (18)</p> <p><i>IPv6</i>: <b>destination-unreachable</b> (1), <b>packet-too-big</b> (2), <b>time-exceeded</b> (3), <b>parameter-problem</b> (4), <b>echo-request</b> (128), <b>echo-reply</b> (129), <b>membership-query</b> (130), <b>membership-report</b> (131), <b>membership-termination</b> (132), <b>router-solicit</b> (133), <b>router-advertisement</b> (134), <b>neighbor-solicit</b> (135), <b>neighbor-advertisement</b> (136), <b>redirect</b> (137), <b>router-renumbering</b> (138), <b>node-information-request</b> (139), <b>node-information-reply</b> (140)</p> <p>See also <b>icmp-code</b> <i>variable</i>.</p> | <p>Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.</p> <p>Egress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.</p> |
| <b>interface</b> <i>interface-name</i>             | <p>Interface on which the packet is received, including the logical unit. You can include the wildcard character (*) as part of an interface name or logical unit.</p> <p><b>NOTE:</b> An interface from which a packet is sent cannot be used as a match condition.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <p>Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.</p> <p>Egress IPv4 (inet) interfaces and IPv6 (inet6) interfaces.</p>                |
| <b>ip-destination-address</b> <i>address</i>       | IPv4 address that is the final destination node address for the packet.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Ingress ports, egress ports, and VLANs.                                                                                                                            |
| <b>ip-options</b>                                  | Specify <b>any</b> to create a match if anything is specified in the options field in the IP header.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Ingress ports, VLANs, and IPv4 (inet) interfaces.                                                                                                                  |
| <b>ip-precedence</b><br><b>ip-precedence-field</b> | <p>IP precedence field. In place of the numeric field value, you can specify one of the following text synonyms (the field values are also listed): <b>critical-ecp</b> (0xa0), <b>flash</b> (0x60), <b>flash-override</b> (0x80), <b>immediate</b> (0x40), <b>internet-control</b> (0xc0), <b>net-control</b> (0xe0), <b>priority</b> (0x20), or <b>routine</b> (0x00).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <p>Ingress ports and VLANs.</p> <p>Egress ports and VLANs.</p>                                                                                                     |
| <b>ip-protocol</b> <i>number</i>                   | IP protocol field.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <p>Ingress ports and VLANs.</p> <p>Egress ports and VLANs.</p>                                                                                                     |

Table 475: Supported Match Conditions for Firewall Filters on QFX10000 Switches (*continued*)

| Match Condition                         | Description                                                                                                                                                                                                                                                                                                                                                                                                                          | Direction and Interface                                                                               |
|-----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------|
| <b>ip-source-address</b> <i>address</i> | IPv4 address of the source node sending the packet.                                                                                                                                                                                                                                                                                                                                                                                  | Ingress ports and VLANs.<br><br>Egress ports and VLANs.                                               |
| <b>ip-version</b> <i>address</i>        | IP version of the packet. Use this condition to match IPv4 or IPv6 header fields in traffic that arrives on a Layer 2 port or VLAN interface.                                                                                                                                                                                                                                                                                        | Ingress ports and VLANs.<br><br>Egress ports and VLANs.                                               |
| <b>is-fragment</b>                      | Using this condition causes a match if the More Fragments flag is enabled in the IP header or if the fragment offset is not zero.                                                                                                                                                                                                                                                                                                    | Ingress ports, VLANs, and IPv4 (inet) interfaces.<br><br>Egress IPv4 (inet) interfaces.               |
| <b>learn-ip-priority</b> <i>number</i>  | Matches the specified IEEE 802.1p VLAN priority bits in the range 0-7.                                                                                                                                                                                                                                                                                                                                                               | Ingress ports and VLANs.<br><br>Egress ports and VLANs.                                               |
| <b>learn-vlan-id</b> <i>number</i>      | Matches the ID of a normal VLAN or the ID of the outer (service) VLAN (for Q-in-Q VLANs). To use filter memory most efficiently and maximize the number of possible filters, use this condition in addition to <b>user-id</b> when you want to match on the inner (customer) VLAN ID. The acceptable values are 1-4095.                                                                                                              | Ingress ports and VLANs.<br><br>Egress ports and VLANs.                                               |
| <b>next-header</b> <i>value</i>         | IPv4 or IPv6 protocol value. In place of the numeric value, you can specify one of the following text synonyms (the numeric values are also listed):<br><br><b>hop-by-hop (0), icmp (1), icmp6 (58), igmp (2), ipip (4), tcp (6), egp (8), udp (17), ipv6 (41), routing (43), fragment (44), rsvp (46), gre (47), esp (50), ah (51), icmp6 (58), no-next-header (59), dstopts (60), ospf (89), pim (103), vrrp (112), sctp (132)</b> | Ingress IPv6 (inet6) interfaces.<br><br>Egress IPv6 (inet6) interfaces.                               |
| <b>packet-length</b> <i>number</i>      | Packet length in bytes. You must enter a number between 0 and 65535.                                                                                                                                                                                                                                                                                                                                                                 | Ingress ports, VLANs, IPv4 (inet), and IPv6 (inet6) interfaces.<br><br>Egress IPv4 (inet) interfaces. |

Table 475: Supported Match Conditions for Firewall Filters on QFX10000 Switches (*continued*)

| Match Condition                                 | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Direction and Interface                                                                                                                                            |
|-------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>precedence value</b>                         | <p>IP precedence bits in the type-of-service (ToS) byte in the IP header. (This byte can also be used for the DiffServ DSCP.) In place of the numeric value, you can specify one of the following text synonyms (the numeric values are also listed):</p> <ul style="list-style-type: none"> <li>• <b>routine (0)</b></li> <li>• <b>priority (1)</b></li> <li>• <b>immediate (2)</b></li> <li>• <b>flash (3)</b></li> <li>• <b>flash-override (4)</b></li> <li>• <b>critical-ecp (5)</b></li> <li>• <b>internet-control (6)</b></li> <li>• <b>net-control (7)</b></li> </ul> | <p>Ingress IPv4 (inet) interfaces.</p> <p>Egress IPv4 (inet) interfaces.</p>                                                                                       |
| <b>protocol type</b>                            | <p>IPv4 or IPv6 protocol value. In place of the numeric value, you can specify one of the following text synonyms (the numeric values are also listed):</p> <p><b>hop-by-hop (0), icmp (1), icmp6, igmp (2), ipip (4), tcp (6), egp (8), udp (17), ipv6 (41), routing (43), fragment (44), rsvp (46), gre (47), esp (50), ah (51), icmp6 (58), no-next-header (59), dstopts (60), ospf (89), pim (103), vrrp (112), sctp (132)</b></p>                                                                                                                                       | <p>Ingress IPv4 (inet) interfaces.</p> <p>Egress IPv4 (inet) interfaces.</p>                                                                                       |
| <b>source-address</b><br><b>ip-address</b>      | IP source address field, which is the address of the node that sent the packet.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <p>Ingress IPv4 (inet) interfaces and IPv6 (inet6) interfaces.</p> <p>Egress IPv4 (inet) interfaces and IPv6 (inet6) interfaces.</p>                               |
| <b>source-mac-address</b><br><b>mac-address</b> | Source media access control (MAC) address of the packet.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <p>Ingress ports and VLANs.</p> <p>Egress ports and VLANs.</p>                                                                                                     |
| <b>source-port value</b>                        | TCP or UDP source port. Typically, you specify this match in conjunction with the <b>protocol</b> match statement. In place of the numeric field, you can specify one of the text synonyms listed under <b>destination-port</b> .                                                                                                                                                                                                                                                                                                                                            | <p>Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.</p> <p>Egress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.</p> |
| <b>source-prefix-list</b> <i>prefix-list</i>    | IP source prefix list. You can define a list of IP address prefixes under a prefix-list alias for frequent use. Define this list at the <b>[edit policy-options]</b> hierarchy level.                                                                                                                                                                                                                                                                                                                                                                                        | <p>Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.</p> <p>Egress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.</p> |

Table 475: Supported Match Conditions for Firewall Filters on QFX10000 Switches (*continued*)

| Match Condition            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Direction and Interface                                                                                                 |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------|
| <b>tcp-established</b>     | <p>Match packets of an established TCP connection. This condition matches packets other than those used to set up a TCP connection—that is, three-way handshake packets are not matched.</p> <p>When you specify <b>tcp-established</b>, a switch does not implicitly verify that the protocol is TCP. You must also specify the <b>protocol tcp</b> match condition.</p>                                                                                                                                                                                                                                                                                                 | <p>Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.</p> <p>Egress IPv4 (inet) interfaces.</p> |
| <b>tcp-flags value</b>     | <p>One or more TCP flags:</p> <ul style="list-style-type: none"> <li>• <b>ack</b> (0x10)</li> <li>• <b>fin</b> (0x01)</li> <li>• <b>push</b> (0x08)</li> <li>• <b>rst</b> (0x04)</li> <li>• <b>syn</b> (0x02)</li> <li>• <b>urgent</b> (0x20)</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                  | <p>Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.</p> <p>Egress IPv4 (inet) interfaces.</p> |
| <b>tcp-initial</b>         | <p>Match the first TCP packet of a connection. A match occurs when the TCP flag <b>SYN</b> is set and the TCP flag <b>ACK</b> is not set.</p> <p>When you specify <b>tcp-initial</b>, a switch does not implicitly verify that the protocol is TCP. You must also specify the <b>protocol tcp</b> match condition.</p>                                                                                                                                                                                                                                                                                                                                                    | <p>Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.</p> <p>Egress IPv4 (inet) interfaces.</p> |
| <b>traffic-class value</b> | <p>8-bit field that specifies the class-of-service (CoS) priority of the packet. The traffic-class field is used to specify a DiffServ code point (DSCP) value. This field was previously used as the type-of-service (ToS) field in IPv4, and, the semantics of this field (for example, DSCP) are identical to those of IPv4.</p> <p>You can specify one of the following text synonyms (the field values are also listed):</p> <p><b>af11 (10), af12 (12), af13 (14), af21 (18), af22 (20), af23 (22), af31 (26), af32 (28), af33 (30), af41 (34), af42 (36), af43 (38), cs0 (0), cs1 (8), cs2 (16), cs3 (24), cs4 (32), cs5 (40), cs6 (48), cs7 (56), ef (46)</b></p> | <p>Ingress IPv6 (inet6) interfaces.</p> <p>Egress IPv6 (inet6) interfaces.</p>                                          |
| <b>ttl value</b>           | <p>IP Time-to-live (TTL) field in decimal. The value can be 1–255.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <p>Ingress IPv4 (inet) interfaces.</p> <p>Egress IPv4 (inet) interfaces.</p>                                            |
| <b>user-vlan-id number</b> | <p>Matches the ID of the inner (customer) VLAN in a Q-in-Q VLAN. To use filter memory most efficiently and maximize the number of possible filters, use in combination with <b>learn-vlan-id</b> to match the outer (service) VLAN ID. The acceptable values are 1–4095.</p>                                                                                                                                                                                                                                                                                                                                                                                              | <p>Ingress ports and VLANs.</p> <p>Egress ports and VLANs.</p>                                                          |



Use **then** statements to define actions that should occur if a packet matches all conditions in a **from** statement. [Table 476](#) shows the actions that you can specify in a term. (If you do not include a **then** statement, the system accepts packets that match the filter.)

**Table 476: Actions for Firewall Filters on QFX10000 Switches**

| Action                                       | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>accept</b>                                | Accept a packet. This is the default action for packets that match a term.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>discard</b>                               | Discard a packet silently without sending an Internet Control Message Protocol (ICMP) message.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>reject <i>message-type</i></b>            | <p>Discard a packet and send a “destination unreachable” ICMPv4 message (type 3). To log rejected packets, configure the <b>syslog</b> action modifier.</p> <p>You can specify one of the following message types: <b>administratively-prohibited</b> (default), <b>bad-host-tos</b>, <b>bad-network-tos</b>, <b>host-prohibited</b>, <b>host-unknown</b>, <b>host-unreachable</b>, <b>network-prohibited</b>, <b>network-unknown</b>, <b>network-unreachable</b>, <b>port-unreachable</b>, <b>precedence-cutoff</b>, <b>precedence-violation</b>, <b>protocol-unreachable</b>, <b>source-host-isolated</b>, <b>source-route-failed</b>, or <b>tcp-reset</b>.</p> <p>If you specify <b>tcp-reset</b>, the system sends a TCP reset if the packet is a TCP packet; otherwise nothing is sent.</p> <p>If you do not specify a message type, the ICMP notification “destination unreachable” is sent with the default message “communication administratively filtered.”</p> <p><b>NOTE:</b> The <b>reject</b> action is supported on ingress interfaces only.</p> |
| <b>routing-instance <i>instance-name</i></b> | Forward matched packets to a virtual routing instance. (The only supported instance type is <b>virtual-router</b> .) Packets can be forwarded to the default instance.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>vlan <i>VLAN-name</i></b>                 | <p>Forward matched packets to a specific VLAN.</p> <p><b>NOTE:</b> The <b>vlan</b> action is supported on ingress interfaces only.</p> <p><b>NOTE:</b> This action is not supported on OCX series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

You can also specify the action modifiers listed in [Table 477](#) to count, mirror, rate-limit, and classify packets.

**Table 477: Action Modifiers for Firewall Filters on QFX10000 Switches**

| Action Modifier                  | Description                                      |
|----------------------------------|--------------------------------------------------|
| <b>count <i>counter-name</i></b> | Count the number of packets that match the term. |

Table 477: Action Modifiers for Firewall Filters on QFX10000 Switches (*continued*)

| Action Modifier                                              | Description                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>forwarding-class</b> <i>class</i>                         | <p>Classify the packet in one of the following default forwarding classes, or in a user-defined forwarding class:</p> <ul style="list-style-type: none"> <li>• <b>best-effort</b></li> <li>• <b>fcoe</b></li> <li>• <b>mcast</b></li> <li>• <b>network-control</b></li> <li>• <b>no-loss</b></li> </ul> <p><b>NOTE:</b> To configure a forwarding class, you must also configure loss priority.</p> |
| <b>log</b>                                                   | <p>Log the packet's header information in the Routing Engine. To view this information, enter the <b>show firewall log</b> operational mode command.</p> <p><b>NOTE:</b> The <b>log</b> action modifier is supported on ingress interfaces only.</p>                                                                                                                                                |
| <b>loss-priority</b> (low   medium-low   medium-high   high) | <p>Set the packet loss priority (PLP).</p> <p><b>NOTE:</b> The <b>loss-priority</b> action modifier is supported on ingress interfaces only.</p> <p><b>NOTE:</b> The <b>loss-priority</b> action modifier is not supported in combination with the <b>policer</b> action.</p>                                                                                                                       |
| <b>policer</b> <i>policer-name</i>                           | <p>Send packets to a policer (for the purpose of applying rate limiting).</p> <p>You can specify a policer for ingress and egress port, VLAN, IPv4 (inet), and IPv6 (inet6) firewall filters.</p> <p><b>NOTE:</b> The <b>policer</b> action modifier is not supported in combination with the <b>loss-priority</b> action.</p>                                                                      |
| <b>port-mirror</b>                                           | <p>(ELS platforms) Mirror traffic (copy packets) to an output interface configured in a port-mirroring instance at the <b>[edit forwarding-options port-mirroring]</b> hierarchy level.</p> <p>You can specify port mirroring for ingress and egress port, VLAN, IPv4 (inet), and IPv6 (inet6) firewall filters.</p>                                                                                |
| <b>port-mirror-instance</b> <i>port-mirror-instance-name</i> | <p>(ELS platforms) Mirror traffic to a port-mirroring instance configured at the <b>[edit forwarding-options port-mirroring]</b> hierarchy level.</p> <p>You can specify port mirroring for ingress and egress port, VLAN, IPv4 (inet), and IPv6 (inet6) firewall filters.</p> <p><b>NOTE:</b></p>                                                                                                  |
| <b>syslog</b>                                                | <p>Log an alert for this packet.</p> <p><b>NOTE:</b> The <b>syslog</b> action modifier is supported on ingress interfaces only.</p>                                                                                                                                                                                                                                                                 |

Table 477: Action Modifiers for Firewall Filters on QFX10000 Switches (*continued*)

| Action Modifier                                            | Description                                                                                                                                                                                                                                                                                                                                   |
|------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>three-color-policer</b> <i>three-color-policer-name</i> | <p>Send packets to a three-color policer (for the purpose of applying rate limiting).</p> <p>You can specify a three-color policer for ingress and egress port, VLAN, IPv4 (inet), and IPv6 (inet6) filters.</p> <p><b>NOTE:</b> The <b>policer</b> action modifier is not supported in combination with the <b>loss-priority</b> action.</p> |

- Related Documentation**
- [Understanding How Firewall Filters Are Evaluated on page 5954](#)
  - [Understanding How a Firewall Filter Tests a Protocol on page 5971](#)
  - [Overview of Policers on page 5999](#)
  - [Understanding Port Mirroring on page 6227](#)
  - [Configuring Firewall Filters on page 5978](#)

## Understanding How a Firewall Filter Tests a Protocol

When examining match conditions in a firewall filter, a switch tests only the fields that you specify. It does not implicitly test any fields that you do not explicitly configure. For example, if you specify a match condition of **source-port ssh**, there is no implied test to determine if the protocol is TCP. In this case, the switch considers any packet that has a value of **22** (decimal) in the 2-byte field that follows a *presumed* IP header to be a match. To ensure that the term matches on TCP packets, you also specify an **ip-protocol tcp** match condition.

For the following match conditions, you should explicitly specify the protocol match condition in the same term:

- **destination-port**—Specify protocol **tcp** or protocol **udp**.
- **icmp-code**—Specify protocol **icmp** and **icmp-type**.
- **icmp-type**—Specify protocol **icmp** or protocol **icmp6**.
- **source-port**—Specify protocol **tcp** or protocol **udp**.
- **tcp-flags**—Specify protocol **tcp**.

- Related Documentation**
- [Understanding Firewall Filter Match Conditions on page 5957](#)
  - [Configuring Firewall Filters on page 5978](#)

## Understanding Firewall Filter Planning

---

Before you create a firewall filter and apply it, determine what you want the filter to accomplish and how to use its match conditions and actions to achieve your goals. It is important that you understand how packets are matched, the default and configured actions of the firewall filter, and where to apply the firewall filter.

You can apply no more than one firewall filter per port, VLAN, or router interface per direction (input and output). For example, for a given port you can apply at most one filter in the input direction and one filter in the output direction. You should try to be conservative in the number of terms (rules) that you include in each firewall filter, because a large number of terms requires longer processing time during a commit operation and can make testing and troubleshooting more difficult.

Before you configure and apply firewall filters, answer the following questions for each of them:

1. What is the purpose of the filter?

For example, the system can drop packets based on header information, rate-limit traffic, classify packets into forwarding classes, log and count packets, or prevent denial-of-service attacks.

2. What are the appropriate match conditions? Determine the packet header fields that the packet must contain for a match. Possible fields include:

- Layer 2 header fields—Source and destination MAC addresses, 802.1Q tag, Ethernet type, or VLAN.
- Layer 3 header fields—Source and destination IP addresses, protocols, and IP options (IP precedence, IP fragmentation flags, or TTL type).
- TCP header fields—Source and destination ports and flags.
- ICMP header fields—Packet type and code.

3. What are the appropriate actions to take if a match occurs?

The system can accept, discard, or reject packets.

4. What additional action modifiers might be required?

For example, you can configure the system to mirror (copy) packets to a specified port, count matching packets, apply traffic management, or police packets.

5. On what port, router interface, or VLAN should the firewall filter be applied?

Start with the following basic guidelines:

- If packets entering or leaving a Layer 2 interface (port) need to be filtered, apply the filter at the **[edit family ethernet switching filter]** hierarchy level. This is a port filter.
- If packets entering or leaving any port in a specific VLAN need to be filtered, use a VLAN filter.

- If packets entering or leaving a Layer 3 (routed) interface or routed VLAN interface (RVI) need to be filtered, use a router firewall filter. Apply the filter to the interface at the **[edit family inet]** hierarchy level. You can also apply a router firewall filter on a loopback interface.

Before you choose the interface or VLAN on which to apply a firewall filter, understand how that placement can affect traffic flow to other interfaces. In general, apply a filter close to the source device if the filter matches on source or destination IP addresses, IP protocols, or protocol information—such as ICMP message types, and TCP or UDP port numbers. However, you should apply a filter close to the destination device if the filter matches *only* on a source IP address. When you apply a filter too close to the source device, the filter could prevent that source device from accessing other services that are available on the network.



**NOTE:** Egress firewall filters do not affect the flow of locally generated control packets from the Routing Engine.

6. In which direction should the firewall filter be applied?

You typically configure different actions for traffic entering an interface than you configure for traffic exiting an interface.

7. How many filters should I create?

See “[Planning the Number of Firewall Filters to Create](#)” on page 5973 for information about how many firewall filters you can apply.

#### Related Documentation

- [Overview of Policers on page 5999](#)
- [Understanding How Firewall Filters Are Evaluated on page 5954](#)
- [Configuring Firewall Filters on page 5978](#)

## Planning the Number of Firewall Filters to Create

- [Understanding How Many Firewall Filters Are Supported on page 5973](#)
- [Egress Filters on page 5975](#)
- [Avoid Configuring too Many Filters on page 5975](#)
- [Configuring TCAM Error Messages on page 5976](#)
- [Policers can Limit Egress Filters on page 5976](#)
- [Planning for Filter-Specific Policers on page 5977](#)
- [Planning for Filter-Based Forwarding on page 5977](#)

## Understanding How Many Firewall Filters Are Supported

QFX3500, QFX3600, QFX5100, and EX4600 switches, QFabric Node devices, and VCF members support the maximum numbers of firewall filter terms per type of attachment point shown in [Table 478](#).

Table 478: Supported Firewall Filter Numbers

| Filter Type | QFX3500, QFX3600 | QFX5100, EX4600 | QFX5200 | QFX10000 |
|-------------|------------------|-----------------|---------|----------|
| Ingress     | 768              | 1536            | 768     | 8192     |
| Egress      | 1024             | 1024            | 1024    | 8192     |

On QFX5100 and QFX5200 switches, you can see how many filters have been programmed of each type by entering **show pfe filter hw summary**.

The totals are applied in aggregate. For example, on the QFX3500 and QFX3600 you can apply a total of 768 terms in all your port filters, Layer 3 filters, and VLAN filters that are applied in the input direction and 1024 terms in port filters, Layer 3 filters, and VLAN filters that are applied in the output direction.



**NOTE:** If you want to create more than 512 egress VLAN filters, your first VLAN ID should be 6 and the subsequent VLAN IDs should increase by 1. For example, to create 1024 egress VLAN filters, the first VLAN ID would be 6, the second ID would be 7, and the sequence would continue through VLAN ID 1029. Similarly, if you want to create fewer than 512 egress VLAN filters but want the total number of terms in those filters to exceed 512, you should number your VLAN IDs in the same manner. If you do not use this approach to create your VLAN IDs, the total number of allowed terms or filters will be less than 1024 and might be 512.

The ternary content addressable memory (TCAM) for firewall filters is divided into slices that accommodate 256 terms, and all the terms in a memory slice must be in filters of the same type and applied in the same direction. A memory slice is reserved as soon as you commit a filter. For example, if you create a port filter and apply it in the input direction, a memory slice is reserved that will only store ingress port filters. If you create and apply only one ingress port filter and that filter has only one term, the rest of this slice is unused and is unavailable for other filter types.

Continuing with the above example for QFX3500 and QFX3600 switches, assume that you create and apply 256 ingress port filters with one term each so that one memory slice is filled. This leaves two more memory slices available for ingress filters. (Remember that the maximum number of ingress terms is 768.) If you then create and apply an ingress Layer 3 filter with one term, another memory slice is reserved for ingress Layer 3 filters. As before, the rest of the slice is unused and is unavailable for different filter types. At this point there is one memory slice available for any ingress filter type.

Now assume that you create and apply a VLAN ingress filter. The final memory slice is reserved for VLAN ingress filters. Memory allocation for ingress filters (once again assuming one term per filter) is as follows:

- Slice 1: Filled with 256 ingress port filters. You cannot commit any more ingress port filters.

- Slice 2: Contains one ingress Layer 3 filter with one term. You can commit 255 more terms in ingress Layer 3 filters.
- Slice 3: Contains one ingress VLAN filter with one term. You can commit 255 more terms in ingress VLAN filters.

Here is another example for QFX3500 and QFX3600 switches. Assume that you create 257 ingress port filters with one term per filter—that is, you create one more term than a single memory slice can accommodate. When you apply the filters and commit the configuration, the filter memory allocation is:

- Slice 1: Filled with 256 ingress port filters. You cannot apply any more ingress port filters.
- Slice 2: Contains one ingress port filter. You can apply 255 more terms in ingress port filters.
- Slice 3: This slice is unassigned. You can create and apply 256 terms in ingress filters of any type (port, Layer 3, or VLAN), but all the filters must be of the same type.

## Egress Filters

All of the preceding principles also apply to egress filters, but four memory slices are used because IPv4 Layer 3 filters and IPv6 Layer 3 filters are stored in separate slices. The memory slices for egress filters are the same size as those for ingress filters, so the maximum number of egress filter terms is therefore 1024.

## Avoid Configuring too Many Filters

If you violate any of these restrictions and commit a configuration that is not in compliance, Junos OS rejects the excessive filters. For example, if you configure 300 ingress port filters and 300 ingress Layer 3 filters and try to commit the configuration, Junos OS does the following (again assuming one term per filter):

- Accepts the 300 ingress port filters (storing them in two memory slices).
- Accepts the first 256 ingress Layer 3 filters it processes (storing them in the third memory slice).
- Rejects the remaining 44 ingress Layer 3 filters.



**NOTE:** In this situation, be sure to delete excessive filters (for example, the remaining 44 ingress Layer 3 filters) from the configuration before you reboot the device. If you reboot a device that has a noncompliant configuration, you cannot predict which filters are installed after the reboot. Using the example above, the 44 ingress Layer 3 filters that were originally rejected might be installed, and 44 of the port filters that were originally accepted might be rejected.

## Configuring TCAM Error Messages

You can configure your switch to display error messages if a filter cannot be installed because there isn't enough TCAM space available. To have TCAM error messages sent to a syslog file, enter

```
set system syslog file filename pfe emergency
```

To have TCAM error messages sent to the console, enter

```
set system syslog console pfe emergency
```

To have TCAM error messages sent to an SSH terminal session, enter

```
set system syslog user user-login pfe emergency
```

## Policers can Limit Egress Filters

On some switches, the number of egress policers that you configure can affect the total number of allowed egress firewall filters. (This issue does not affect QFX10000 switches.) On the affected switches, every policer has two implicit counters that consume two entries in a 1024-entry TCAM that is used for counters, including counters that are configured as action modifiers in firewall filter terms. (Policers consume two entries because one is used for green packets and one is used for nongreen packets regardless of policer type.) If the TCAM becomes full, you cannot commit any more egress firewall filters that have terms with counters. For example, if you configure and commit 512 egress policers (two-color, three-color, or a combination of both policer types), all of the memory entries for counters are used up. If later in your configuration file you insert additional egress firewall filters with terms that also include counters, *none* of the terms in those filters are committed because there is no available memory space for the counters.

Here are some additional examples:

- Assume that you configure egress filters that include a total of 512 policers and no counters. Later in your configuration file you include another egress filter with 10 terms, 1 of which has a counter action modifier. None of the terms in this filter are committed because there is not enough TCAM space for the counter.
- Assume that you configure egress filters that include a total of 500 policers, so 1000 TCAM entries are occupied. Later in your configuration file you include the following two egress filters:
  - Filter A with 20 terms and 20 counters. All the terms in this filter are committed because there is enough TCAM space for all the counters.
  - Filter B comes after Filter A and has five terms and five counters. *None* of the terms in this filter are committed because there is not enough memory space for *all* the counters. (Five TCAM entries are required but only four are available.)

You can prevent this problem from occurring by ensuring that egress firewall filter terms with counter actions are placed earlier in your configuration file than terms that include policers. In this circumstance, Junos OS commits policers even if there is not enough TCAM space for the implicit counters. For example, assume the following:



- You have 1024 egress firewall filter terms with counter actions.
- Later in your configuration file you have an egress filter with 10 terms. None of the terms have counters but one has a policer action modifier.

You can successfully commit the filter with 10 terms even though there is not enough TCAM space for the implicit counters of the policer. The policer is committed without the counters.

## Planning for Filter-Specific Policers

You can configure policers to be filter-specific, which means that Junos OS creates only one policer instance regardless of how many times the policer is referenced. When you do this, rate limiting is applied in aggregate, so if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms, the total bandwidth allowed by the filter is 1 Gbps. However, the behavior of a filter-specific policer is affected by how the firewall filter terms that reference the policer are stored in ternary content addressable memory (TCAM). If you create a filter-specific policer and reference it in multiple firewall filter terms, the policer allows more traffic than expected if the terms are stored in different TCAM slices. For example, if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms that are stored in three separate memory slices, the total bandwidth allowed by the filter is 3 Gbps, not 1 Gbps.

To prevent this unexpected behavior from occurring, use the information about TCAM slices presented above to organize your configuration file so that all the firewall filter terms that reference a given filter-specific policer are stored in the same TCAM slice.

## Planning for Filter-Based Forwarding

You can use firewall filters in conjunction with virtual routing instances to specify different routes for packets to travel in their networks. To set up this feature—called filter-based forwarding—you specify a filter and match criteria and then specify the virtual routing instance to send packets to. Filters used in this way also consume memory in an additional TCAM. See *Understanding FIP Snooping, FBF, and MVR Filter Scalability* for more information. The section *FBF Filter VFP TCAM Consumption* in this topic specifically addresses the number of supported filters when using filter-based forwarding.



**WARNING:** Filter-based forwarding does not work with IPv6 interfaces on some Juniper switches.

### Related Documentation

- [Understanding How Firewall Filters Are Evaluated on page 5954](#)
- [Understanding Firewall Filter Planning on page 5972](#)
- [Configuring Firewall Filters on page 5978](#)
- [Understanding Filter-Based Forwarding on page 5983](#)

## Understanding Firewall Filter Processing Points for Bridged and Routed Packets

---

You apply firewall filters at multiple processing points in the forwarding path. At each processing point, the action to be taken on a packet is determined by the configuration of the filter and the results of the lookup in the forwarding or routing table.

For both bridged (Layer 2) unicast packets and routed (Layer 3) unicast packets, firewall filters are applied in the prescribed order shown below (assuming that each filter is present and a packet is accepted by each one).

Bridged packets:

1. Ingress port filter
2. Ingress VLAN filter
3. Egress VLAN filter
4. Egress port filter

Routed packets:

1. Ingress port firewall filter
2. Ingress VLAN firewall filter (Layer 2 CoS)
3. Ingress router firewall filter (Layer 3 CoS)
4. Egress router firewall filter
5. Egress VLAN firewall filter
6. Egress port filter



**NOTE:** MAC learning occurs before filters are applied, so switches learn the MAC addresses of packets that are dropped by ingress filters.

---

### Related Documentation

- [Overview of Firewall Filters on page 5951](#)
- [Understanding How Firewall Filters Control Packet Flows on page 5956](#)
- [Configuring Firewall Filters on page 5978](#)

## Configuring Firewall Filters

---

You can configure firewall filters in a switch to control traffic that enters switch ports or enters and exits VLANs and Layer 3 (routed) interfaces. To use a firewall filter, you must configure the filter and then apply it to a port, VLAN, or Layer 3 interface.

- [Configuring a Firewall Filter on page 5979](#)
- [Applying a Firewall Filter to a Port on page 5980](#)

- [Applying a Firewall Filter to a VLAN on page 5981](#)
- [Applying a Firewall Filter to a Layer 3 \(Routed\) Interface on page 5981](#)

## Configuring a Firewall Filter

To configure a firewall filter:

1. Configure the family address type, filter name, term name, and at least one match condition—for example, match on packets that contain a specific source address:

```
[edit]
user@switch# set firewall family ethernet-switching filter ingress-port-filter term term-one
from source-address 192.0.2.14
```

For a firewall filter that is applied to a port or VLAN, specify the family address type **ethernet-switching**. For a firewall filter that is applied to a Layer 3 (routed) interface, specify the family address type **inet**.

The filter and term names can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. Each filter name must be unique. A filter can contain one or more terms, and each term name must be unique within a filter.

2. Configure additional match conditions. For example, match on packets that contain a specific source port:

```
[edit firewall family ethernet-switching filter ingress-port-filter term
term-one from]
user@switch# set source-port 80
```

You can specify one or more match conditions in a single **from** statement. For a match to occur, the packet must match all the conditions in the term. The **from** statement is optional, but if included in a term, it cannot be empty. If you omit the **from** statement, all packets are considered to match.

3. If you want to apply a firewall filter to multiple interfaces and be able to see counters specific to each interface, configure the **interface-specific** option:

```
[edit firewall family ethernet-switching filter ingress-port-filter]
user@switch# set interface-specific
```

4. In each firewall filter term, specify the actions to take if the packet matches all the conditions in that term. You can specify an action and action modifiers:

- To specify a filter action, for example, to discard packets that match the conditions of the filter term:

```
[edit firewall family ethernet-switching filter ingress-port-filter term
term-one then]
user@switch# set discard
```

You can specify no more than one action (**accept**, **discard**, **reject**, **routing-instance**, or **vlan**) per term.

- To specify action modifiers, for example, to count and classify packets to a forwarding class:

```
[edit firewall family ethernet-switching filter ingress-port-filter term
term-one then]
user@switch# set count counter-one
user@switch# set forwarding-class expedited-forwarding
user@switch# set loss-priority high
```

You can specify any of the following action modifiers in a **then** statement:

- **analyzer** *analyzer-name*—Mirror port traffic to a specified analyzer, which you must configure at the **[ethernet-switching-options]** level.
- **count** *counter-name*—Count the number of packets that pass this filter term.



**NOTE:** We recommend that you configure a counter for each term in a firewall filter, so that you can monitor the number of packets that match the conditions specified in each filter term.



**NOTE:** On QFX3500 and QFX3600 switches, filters automatically count packets that have been dropped on ingress because of cyclic redundancy check (CRC) errors.

- **forwarding-class** *class*—Assign packets to a forwarding class.
- **log**—Log the packet header information in the Routing Engine.
- **loss-priority** *priority*—Set the priority of dropping a packet.
- **policer** *policer-name*—Apply rate-limiting to the traffic.
- **syslog**—Log an alert for this packet.

If you omit the **then** statement or do not specify an action, packets that match all the conditions in the **from** statement are accepted. However, you should always explicitly configure an action in the **then** statement. You can include no more than one action statement, but you can use any combination of action modifiers. For an action or action modifier to take effect, all conditions in the **from** statement must match.



**NOTE:** Implicit discard is also applicable to a firewall filter applied to the loopback interface, lo0.

## Applying a Firewall Filter to a Port

To apply a firewall filter to an ingress port:

1. Provide a meaningful description of the firewall filter in the configuration of the port to which the filter will be applied:

```
[edit]
user@switch# set interfaces ge-0/0/6 description "filter to limit tcp traffic at trunk port for employee-vlan"
```

2. Apply the filter to the interface, specifying the unit number, family address type, the direction of the filter (for packets entering the port), and the filter name:

```
[edit]
user@switch# set ge-0/0/6 unit 0 family ethernet-switching filter input ingress-port-filter
```

For firewall filters that are applied to ports, the family address type must be **ethernet-switching**.



**NOTE:** You can apply only one filter to a port for a given direction (ingress or egress).

## Applying a Firewall Filter to a VLAN

To apply a firewall filter to a VLAN:

1. Provide a meaningful description of the firewall filter in the configuration of the VLAN to which the filter will be applied:

```
[edit]
user@switch# set vlans employee-vlan vlan-id 20 description "filter to block rogue devices
on employee-vlan"
```

2. Apply firewall filters to filter packets that are entering or exiting the VLAN:

- To apply a filter to match packets that are entering the VLAN:

```
[edit]
user@switch# set vlans employee-vlan vlan-id 20 filter input ingress-vlan-rogue-block
```

- To apply a firewall filter to match packets that are exiting the VLAN:

```
[edit]
user@switch# set vlans employee-vlan vlan-id 20 filter output egress-vlan-filter
```



**NOTE:** You can apply only one filter to a VLAN for a given direction (ingress or egress).

## Applying a Firewall Filter to a Layer 3 (Routed) Interface

To apply a firewall filter to a Layer 3 routed interface:

1. Provide a meaningful description of the firewall filter in the configuration of the interface to which the filter will be applied:

```
[edit]
user@switch# set interfaces ge-0/1/6 description "filter to count and monitor traffic on layer
3 interface"
```

2. You can apply firewall filters to filter packets that enter or exit a Layer 3 routed interface:

- To apply a firewall filter to filter packets that enter a Layer 3 interface:

```
[edit]
user@switch# set interfaces ge-0/1/6 unit 0 family inet filter input ingress-router-filter
```

- To apply a firewall filter to filter packets that exit a Layer 3 interface:

```
[edit]
user@switch# set interfaces ge-0/1/6 unit 0 family inet filter output egress-router-filter
```

For firewall filters applied to Layer 3 routed interfaces, the family address type must be **inet**.



**NOTE:** You can apply only one filter to an interface for a given direction (ingress or egress).

**Related  
Documentation**

- [Overview of Firewall Filters on page 5951](#)
- [Firewall Filter Match Conditions and Actions](#)
- [Verifying That Firewall Filters Are Operational on page 5988](#)
- [Monitoring Firewall Filter Traffic on page 5987](#)
- [Configuring Port Mirroring](#)

---

## Applying Firewall Filters to Interfaces

For a firewall filter to work, you must apply it to at least one interface. To do this, include the **filter** statement when configuring a logical interface at the **[edit interfaces]** hierarchy level:

```
[edit interfaces]
user@switch# set interface-name unit logical-unit-number family family-name filter (input |
output) filter-name
```

In the **input** statement, specify a firewall filter to be evaluated when packets are received on the interface. Input filters applied to a loopback interface affect only traffic destined for the Routing Engine.

In the **output** statement, specify a filter to be evaluated when packets exit the interface.



**NOTE:** When you create a loopback interface, it is important to apply an ingress filter to it so the Routing Engine is protected. We recommend that when you apply a filter to the loopback interface **lo0**, you include the **apply-groups** statement. Doing so ensures that the filter is automatically inherited on every loopback interface, including **lo0** and other loopback interfaces.

**Related  
Documentation**

- [Configuring Firewall Filters on page 5978](#)

## Understanding Filter-Based Forwarding

You can use firewall filters in conjunction with virtual routing instances to specify different routes for packets to travel in their networks. To set up this feature—called filter-based forwarding—you specify a filter and match criteria and then specify the virtual routing instance to send packets to.

You might want to use filter-based forwarding to route specific types of traffic through a firewall or other security device before the traffic continues on its path. You can also use filter-based forwarding to give certain types of traffic preferential treatment. For example, you might want to ensure that the highest-priority traffic is forwarded over a 40-Gigabit Ethernet link. You might also use filter-based forwarding to obtain more control over load balancing than dynamic routing protocols provide.



**NOTE:** You can create as many as 128 filters or terms that direct packets to a given virtual routing instance.

### Related Documentation

- [Understanding Virtual Router Routing Instances on page 3232](#)
- [Overview of Firewall Filters on page 5951](#)
- [Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on page 5983](#)

## Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device

You can configure filter-based forwarding by using a firewall filter to forward matched traffic to a specific virtual routing instance.

This example describes how to set up filter-based forwarding:

- [Requirements on page 5983](#)
- [Overview and Topology on page 5983](#)
- [Configuration on page 5984](#)
- [Verification on page 5986](#)

### Requirements

This example requires Junos OS Release 15.1X53-D10 or later on a QFX10000 switch..

### Overview and Topology

In this example, traffic from one application server that is destined for a different application server is matched by a firewall filter based on the IP address of the source application server. Any matching packets are routed to a virtual routing instance that sends the traffic to a security device. In this case, the security device must be able to

forward the traffic to the destination application server. For this example, assume that the address of the destination application server is 192.168.0.1.



**WARNING:** Filter-based forwarding does not work with IPv6 interfaces on some Juniper switches.

## Configuration

To configure filter-based forwarding:

### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste them into the CLI at the **[edit]** hierarchy level.

```
[edit]
set interfaces xe-0/0/0 unit 0 family inet address 10.1.0.1/24
set interfaces xe-0/0/3 unit 0 family inet address 10.1.3.1/24
set firewall family inet filter f1 term t1 from source-address 10.1.0.50/32
set firewall family inet filter f1 term t1 from protocol tcp
set interfaces xe-0/0/0 unit 0 family inet filter input f1
set routing-instances vrf01 instance-type virtual-router
set routing-instances vrf01 interface xe-0/0/3.0
set routing-instances vrf01 routing-options static route 192.168.0.1/24 next-hop 10.1.3.254
set firewall family inet filter f1 term t1 then routing-instance vrf01
```

### Step-by-Step Procedure

To configure filter-based forwarding:

1. Configure an interface to connect to the application server:  

```
[edit interfaces]
user@switch# set xe-0/0/0 unit 0 family inet address 10.1.0.1/24
```
2. Configure an interface to connect to the security device:  

```
[edit interfaces]
user@switch# set xe-0/0/3 unit 0 family inet address 10.1.3.1/24
```
3. Create a firewall filter that matches packets based on the address of the application server that the traffic will be sent from. Also configure the filter so that it matches only TCP packets:  

```
[edit firewall]
user@switch# set family inet filter f1 term t1 from source-address 10.1.0.50/32
user@switch# set firewall family inet filter f1 term t1 from protocol tcp
```
4. Apply the filter to the interface that connects to the source application server and configure it to match incoming packets:  

```
[edit interfaces]
user@switch# set xe-0/0/0 unit 0 family inet filter input f1
```
5. Create a virtual router:  

```
[edit]
user@switch# set routing-instances vrf01 instance-type virtual-router
```
6. Associate the virtual router with the interface that connects to the security device:  

```
[edit routing-instances]
user@switch# set vrf01 interface xe-0/0/3.0
```
7. Configure the routing information for the virtual routing instance:  

```
[edit routing-instances]
```



```

user@switch# set vrf01 routing-options static route 192.168.0.1/24 next-hop 10.1.3.254
8. Set the filter to forward packets to the virtual router:

[edit firewall]
user@switch# set family inet filter f1 term t1 then routing-instance vrf01

```

## Results

Check the results of the configuration:

```

user@switch> show configuration
interfaces {
 xe-0/0/0 {
 unit 0 {
 family inet {
 filter {
 input f1;
 }
 address 10.1.0.1/24;
 }
 }
 }
 xe-0/0/3 {
 unit 0 {
 family inet {
 address 10.1.3.1/24;
 }
 }
 }
}
firewall {
 family inet {
 filter f1 {
 term t1 {
 from {
 source-address {
 10.1.0.50/32;
 }
 protocol tcp;
 }
 then {
 routing-instance vrf01;
 }
 }
 }
 }
}
routing-instances {
 vrf01 {
 instance-type virtual-router;
 interface xe-0/0/1.0;
 routing-options {
 static {
 route 12.34.56.0/24 next-hop 10.1.3.254;
 }
 }
 }
}

```

```

 }
 }

```

## Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verifying That Filter-Based Forwarding Was Configured on page 5986](#)

### Verifying That Filter-Based Forwarding Was Configured

**Purpose** Verify that filter-based forwarding was properly enabled on the switch.

**Action** 1. Use the `show interfaces filters` command:

```

user@switch> show interfaces filters xe-0/0/0.0
Interface Admin Link Proto Input Filter Output Filter
xe-0/0/0.0 up down inet f1

```

2. Use the `show route forwarding-table` command:

```

user@switch> show route forwarding-table

Routing table: default.inet
Internet:
Destination Type RtRef Next hop Type Index NhRef Netif
default user 1 0:12:f2:21:cf:0 ucst 331 4 me0.0
default perm 0 rjct 36 3
0.0.0.0/32 perm 0 dscd 34 1
10.1.0.0/24 ifdn 0 rslv 613 1 xe-0/0/0.0
10.1.0.0/32 iddn 0 10.1.0.0 recv 611 1 xe-0/0/0.0
10.1.0.1/32 user 0 rjct 36 3
10.1.0.1/32 intf 0 10.1.0.1 locl 612 2
10.1.0.1/32 iddn 0 10.1.0.1 locl 612 2
10.1.0.255/32 iddn 0 10.1.0.255 bcst 610 1 xe-0/0/0.0
10.1.1.0/26 ifdn 0 rslv 583 1 vlan.0
10.1.1.0/32 iddn 0 10.1.1.0 recv 581 1 vlan.0
10.1.1.1/32 user 0 rjct 36 3
10.1.1.1/32 intf 0 10.1.1.1 locl 582 2
10.1.1.1/32 iddn 0 10.1.1.1 locl 582 2
10.1.1.63/32 iddn 0 10.1.1.63 bcst 580 1 vlan.0
255.255.255.255/32 perm 0 bcst 32 1

```

```

Routing table: vrf01.inet
Internet:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0 rjct 559 2
0.0.0.0/32 perm 0 dscd 545 1
10.1.3.0/24 ifdn 0 rslv 617 1 xe-0/0/3.0
10.1.3.0/32 iddn 0 10.1.3.0 recv 615 1 xe-0/0/3.0
10.1.3.1/32 user 0 rjct 559 2
192.168.0.1/24 user 0 10.1.3.254 ucst 616 2 xe-0/0/3.0
192.168.0.1/24 user 0 10.1.3.254 ucst 616 2 xe-0/0/3.0
10.1.3.255/32 iddn 0 10.1.3.255 bcst 614 1 xe-0/0/3.0
224.0.0.0/4 perm 0 mdsc 546 1
224.0.0.1/32 perm 0 224.0.0.1 mcst 529 1
255.255.255.255/32 perm 0 bcst 543 1

```

```

Routing table: default.iso
ISO:

```

| Destination | Type | RtRef | Next hop | Type | Index | NhRef | Netif |
|-------------|------|-------|----------|------|-------|-------|-------|
| default     | perm | 0     |          | rjct | 60    | 1     |       |

Routing table: vrf01.iso

ISO:

| Destination | Type | RtRef | Next hop | Type | Index | NhRef | Netif |
|-------------|------|-------|----------|------|-------|-------|-------|
| default     | perm | 0     |          | rjct | 600   | 1     |       |

**Meaning** The output indicates that the filter was created on the interface and that the virtual routing instance is forwarding matching traffic to the correct IP address.

- Related Documentation**
- [Configuring Firewall Filters on page 5978](#)
  - [Understanding Filter-Based Forwarding on page 5983](#)
  - [Understanding Virtual Router Routing Instances on page 3232](#)

## Monitoring Firewall Filter Traffic

You can use operational mode commands to monitor firewall filter traffic.

- [Monitoring Traffic for All Firewall Filters and Policers That Are Configured on page 5987](#)
- [Monitoring Traffic for a Specific Firewall Filter on page 5988](#)
- [Monitoring Traffic for a Specific Policar on page 5988](#)

### Monitoring Traffic for All Firewall Filters and Policers That Are Configured

**Purpose** Monitor the number of packets and bytes that matched the firewall filters and monitor the number of packets that exceeded policer rate limits:

**Action** Use the **show firewall** operational mode command:

```
user@switch> show firewall
Filter: egress-vlan-watch-employee
Counters:
Name Bytes Packets
counter-employee-web 3348 27
Filter: ingress-port-limit-tcp-icmp
Counters:
Name Bytes Packets
icmp-counter 560 10
Policers:
Name Packets
icmp-connection-policer 10
tcp-connection-policer 0
Filter: ingress-vlan-rogue-block
Filter: ingress-vlan-limit-guest
```

**Meaning** The **show firewall** command displays the names of all firewall filters, counters, and policers that are configured. For each counter that is specified in a filter configuration, the output field shows the byte count and packet count for the term in which the counter

is specified. For each policer that is specified in a filter configuration, the output field shows the packet count for packets that exceed the specified rate limits.

## Monitoring Traffic for a Specific Firewall Filter

**Purpose** Monitor the number of packets and bytes that matched a firewall filter and monitor the number of packets that exceeded policer rate limits.

**Action** Use the **show firewall filter *filter-name*** operational mode command:

```
user@switch> show firewall filter ingress-port-limit-tcp-icmp
Filter: ingress-port-limit-tcp-icmp
Counters:
Name Bytes Packets
icmp-counter 560 10
```

**Meaning** The **show firewall filter *filter-name*** command limits the display information to the counters and policers that are defined for the specified filter.

## Monitoring Traffic for a Specific Policer

**Purpose** Monitor the number of packets that exceeded the rate limits of a policer:

**Action** Use the **show firewall policer *policer-name*** operational mode command:

```
user@switch> show firewall policer icmp-connection-policer
Filter: ingress-port-limit-tcp-icmp
Policers:
Name Packets
icmp-connection-policer 10
```

**Meaning** The **show firewall policer *policer-name*** command displays the number of packets that exceeded the rate limits for the specified policer.

**Related Documentation**

- [Configuring Firewall Filters on page 5978](#)
- [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 6017](#)
- [Verifying That Firewall Filters Are Operational on page 5988](#)

---

## Verifying That Firewall Filters Are Operational

**Purpose** Verify that firewall filters are working properly.

**Action** Use the **show firewall** operational mode command to verify that the firewall filters are working properly:

```
user@switch> show firewall
Filter: egress-vlan-watch-employee
Counters:
Name Bytes Packets
counter-employee-web 0 0
Filter: ingress-port-limit-tcp-icmp
Counters:
```

| Name                             | Bytes   | Packets |
|----------------------------------|---------|---------|
| icmp-counter                     | 560     | 10      |
| Policers:                        |         |         |
| Name                             | Packets |         |
| icmp-connection-policer          | 10      |         |
| tcp-connection-policer           | 0       |         |
| Filter: ingress-vlan-rogue-block |         |         |
| Filter: ingress-vlan-limit-guest |         |         |

**Meaning** The **show firewall** command displays the names of all firewall filters, counters, and policers that are configured. For each counter that is specified in a filter configuration, the output field shows the byte count and packet count for the term in which the counter is specified. In the above example, the **icmp-counter** in the filter **ingress-port-limit-tcp-icmp** shows that the filter matched 10 packets. For each policer that is specified in a filter configuration, the output field shows the packet count for packets that exceed the specified rate limits. The policer **icmp-connection-policer** shows that 10 ICMP packets were policed.

**Related Documentation**

- [Configuring Firewall Filters on page 5978](#)
- [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 6017](#)
- [Monitoring Firewall Filter Traffic on page 5987](#)

## Troubleshooting Firewall Filters

Use the following information to troubleshoot your firewall filter configuration.

- [Troubleshooting QFX10000 Switches on page 5989](#)
- [Troubleshooting Other Switches on page 5990](#)

## Troubleshooting QFX10000 Switches

This section describes issues specific to QFX10000 switches:

- [Do Not Combine Match Conditions for Different Layers on page 5989](#)
- [Layer 2 Packets Cannot be Discarded with Firewall Filters on page 5989](#)
- [Protect-RE \(loopback\) Firewall Filter Does Not Filter Packets Applied to EM0 Interfaces on page 5990](#)

### Do Not Combine Match Conditions for Different Layers

On QFX10000 switches, do not combine match conditions for Layer 2 and any other layer in a **family ethernet-switching** filter. (For example, do not include conditions that match MAC addresses and IP addresses in the same filter.) If you do so, the filter will commit successfully but will not work. You will also see the following log message: **L2 filter *filter-name* doesn't support mixed L2 and L3/L4 match conditions. Please re-config.**

### Layer 2 Packets Cannot be Discarded with Firewall Filters

**Problem Description:** Layer 2 (L2) control packets such as Link Layer Discovery Protocol (LLDP) and bridge protocol data unit (BPDU) cannot be discarded with firewall filters.

**Solution** Configure distributed denial-of-service (DDoS) protection on the L2 control packet and set the aggregate policer bandwidth and burst values to the minimum value of 1. For example,

```
[edit system ddos-protection protocols protocol name]
user@host# set aggregate bandwidth 1
```

```
[edit system ddos-protection protocols protocol name]
user@host# set aggregate burst 1
```

### Protect-RE (loopback) Firewall Filter Does Not Filter Packets Applied to EMO Interfaces

**Problem** **Description:** On QFX10000 switches, the Protect-RE (loopback) firewall filter does not filter packets applied to EMO interfaces including SNMP, Telnet, and other services.

**Solution** This is expected behavior.

## Troubleshooting Other Switches

This section describes issues specific to QFX switches other than QFX10000 switches. This information also applies to OCX1100 switches and EX4600 switches.

- [Firewall Filter Configuration Returns a No Space Available in TCAM Message on page 5990](#)
- [Filter Counts Previously Dropped Packet on page 5992](#)
- [Matching Packets Not Counted on page 5993](#)
- [Counter Reset When Editing Filter on page 5993](#)
- [Cannot Include loss-priority and policer Actions in Same Term on page 5993](#)
- [Cannot Egress Filter Certain Traffic Originating on QFX Switch on page 5994](#)
- [Firewall Filter Match Condition Not Working with Q-in-Q Tunneling on page 5994](#)
- [Egress Firewall Filters with Private VLANs on page 5994](#)
- [Egress Filtering of L2PT Traffic Not Supported on page 5995](#)
- [Cannot Drop BGP Packets in Certain Circumstances on page 5995](#)
- [Invalid Statistics for Policer on page 5995](#)
- [Policers can Limit Egress Filters on page 5995](#)

### Firewall Filter Configuration Returns a No Space Available in TCAM Message

**Problem** **Description:** When a firewall filter configuration exceeds the amount of available Ternary Content Addressable Memory (TCAM) space, the system returns the following **syslogd** message:

```
No space available in tcam.
Rules for filter filter-name will not be installed.
```

A switch returns this message during the commit operation if the firewall filter that has been applied to a port, VLAN, or Layer 3 interface exceeds the amount of space available

in the TCAM table. The filter is not applied, but the commit operation for the firewall filter configuration is completed in the CLI module.

**Solution** When a firewall filter configuration exceeds the amount of available TCAM table space, you must configure a new firewall filter with fewer filter terms so that the space requirements for the filter do not exceed the available space in the TCAM table.

You can perform either of the following procedures to correct the problem:

To delete the filter and its binding and apply the new smaller firewall filter to the same binding:

1. Delete the filter and its binding to ports, VLANs, or Layer 3 interfaces. For example:

```
[edit]
user@switch# delete firewall family ethernet-switching filter ingress-vlan-rogue-block
user@switch# delete vlans employee-vlan description "filter to block rogue devices on
employee-vlan"
user@switch# delete vlans employee-vlan filter input ingress-vlan-rogue-block
```

2. Commit the changes:

```
[edit]
user@switch# commit
```

3. Configure a smaller filter with fewer terms that does not exceed the amount of available TCAM space. For example:

```
[edit]
user@switch# set firewall family ethernet-switching filter new-ingress-vlan-rogue-block ...
```

4. Apply (bind) the new firewall filter to a port, VLAN, or Layer 3 interface. For example:

```
[edit]
user@switch# set vlans employee-vlan description "filter to block rogue devices on
employee-vlan"
user@switch# set vlans employee-vlan filter input new-ingress-vlan-rogue-block
```

5. Commit the changes:

```
[edit]
user@switch# commit
```

To apply a new firewall filter and overwrite the existing binding but not delete the original filter:

1. Configure a firewall filter with fewer terms than the original filter:

```
[edit]
user@switch# set firewall family ethernet-switching filter new-ingress-vlan-rogue-block...
```

2. Apply the firewall filter to the port, VLAN, or Layer 3 interfaces to overwrite the binding of the original filter—for example:

```
[edit]
user@switch# set vlans employee-vlan description "smaller filter to block rogue devices on
employee-vlan"
user@switch# set vlans employee-vlan filter input new-ingress-vlan-rogue-block
```

Because you can apply no more than one firewall filter per VLAN per direction, the binding of the original firewall filter to the VLAN is overwritten with the new firewall filter **new-ingress-vlan-rogue-block**.

3. Commit the changes:

```
[edit]
user@switch# commit
```



**NOTE:** The original filter is not deleted and is still available in the configuration.

---

### Filter Counts Previously Dropped Packet

---

- Problem**    **Description:** If you configure two or more filters in the same direction for a physical interface and one of the filters includes a counter, the counter will be incorrect if the following circumstances apply:
- You configure the filter that is applied to packets first to discard certain packets. For example, imagine that you have a VLAN filter that accepts packets sent to 10.10.1.0/24 addresses and implicitly discards packets sent to any other addresses. You apply the filter to the **admin** VLAN in the output direction, and interface xe-0/0/1 is a member of that VLAN.
  - You configure a subsequent filter to accept and count packets that are dropped by the first filter. In this example, you have a port filter that accepts and counts packets sent to 192.168.1.0/24 addresses that is also applied to xe-0/0/1 in the output direction.

The egress VLAN filter is applied first and correctly discards packets sent to 192.168.1.0/24 addresses. The egress port filter is applied next and counts the discarded packets as matched packets. The packets are not forwarded, but the counter displayed by the egress port filter is incorrect.

Remember that the order in which filters are applied depends on the direction in which they are applied, as indicated here:

Ingress filters:

1. Port (Layer 2) filter
2. VLAN filter
3. Router (Layer 3) filter

Egress filters:

1. Router (Layer 3) filter
2. VLAN filter
3. Port (Layer 2) filter

**Solution**    This is expected behavior.



### Matching Packets Not Counted

---

**Problem** **Description:** If you configure two egress filters with counters for a physical interface and a packet matches both of the filters, only one of the counters includes that packet. For example:

- You configure an egress port filter with a counter for interface xe-0/0/1.
- You configure an egress VLAN filter with a counter for the **adminVLAN**, and interface xe-0/0/1 is a member of that VLAN.
- A packet matches both filters.

In this case, the packet is counted by only one of the counters even though it matched both filters.

**Solution** This is expected behavior.

### Counter Reset When Editing Filter

---

**Problem** **Description:** If you edit a firewall filter term, the value of any counter associated with any term in the same filter is set to 0, including the implicit counter for any policer referenced by the filter. Consider the following examples:

- Assume that your filter has **term1**, **term2**, and **term3**, and each term has a counter that has already counted matching packets. If you edit any of the terms in any way, the counters for all the terms are reset to 0.
- Assume that your filter has **term1** and **term2**. Also assume that **term2** has a **policer** action modifier and the implicit counter of the policer has already counted 1000 matching packets. If you edit **term1** or **term2** in any way, the counter for the policer referenced by **term2** is reset to 0.

**Solution** This is expected behavior.

### Cannot Include loss-priority and policer Actions in Same Term

---

**Problem** **Description:** You cannot include both of the following actions in the same firewall filter term in a QFX Series switch:

- **loss-priority**
- **policer**

If you do so, you see the following error message when you attempt to commit the configuration: "cannot support policer action if loss-priority is configured."

**Solution** This is expected behavior.

### Cannot Egress Filter Certain Traffic Originating on QFX Switch

---

**Problem** **Description:** On a QFX Series switch, you cannot filter certain traffic with a firewall filter applied in the output direction if the traffic originates on the QFX switch. This limitation applies to control traffic for protocols such as ICMP (ping), STP, LACP, and so on.

**Solution** This is expected behavior.

### Firewall Filter Match Condition Not Working with Q-in-Q Tunneling

---

**Problem** **Description:** If you create a firewall filter that includes a match condition of **dot1q-tag** or **dot1q-user-priority** and apply the filter on input to a trunk port that participates in a service VLAN, the match condition does not work if the Q-in-Q EtherType is not 0x8100. (When Q-in-Q tunneling is enabled, trunk interfaces are assumed to be part of the service provider or data center network and therefore participate in service VLANs.)

**Solution** This is expected behavior. To set the Q-in-Q EtherType to 0x8100, enter the **set dot1q-tunneling ethertype 0x8100** statement at the **[edit ethernet-switching-options]** hierarchy level. You must also configure the other end of the link to use the same Ethertype.

### Egress Firewall Filters with Private VLANs

---

**Problem** **Description:** If you apply a firewall filter in the output direction to a primary VLAN, the filter also applies to the secondary VLANs that are members of the primary VLAN when the traffic egresses with the primary VLAN tag or isolated VLAN tag, as listed below:

- Traffic forwarded from a secondary VLAN trunk port to a promiscuous port (trunk or access)
- Traffic forwarded from a secondary VLAN trunk port that carries an isolated VLAN to a PVLAN trunk port.
- Traffic forwarded from a promiscuous port (trunk or access) to a secondary VLAN trunk port
- Traffic forwarded from a PVLAN trunk port. to a secondary VLAN trunk port
- Traffic forwarded from a community port to a promiscuous port (trunk or access)

If you apply a firewall filter in the output direction to a primary VLAN, the filter does *not* apply to traffic that egresses with a community VLAN tag, as listed below:

- Traffic forwarded from a community trunk port to a PVLAN trunk port
- Traffic forwarded from a secondary VLAN trunk port that carries a community VLAN to a PVLAN trunk port
- Traffic forwarded from a promiscuous port (trunk or access) to a community trunk port
- Traffic forwarded from a PVLAN trunk port. to a community trunk port

If you apply a firewall filter in the output direction to a community VLAN, the following behaviors apply:

- The filter is applied to traffic forwarded from a promiscuous port (trunk or access) to a community trunk port (because the traffic egresses with the community VLAN tag).
- The filter is applied to traffic forwarded from a community port to a PVLAN trunk port (because the traffic egresses with the community VLAN tag).
- The filter is *not* applied to traffic forwarded from a community port to a promiscuous port (because the traffic egresses with the primary VLAN tag or untagged).

**Solution** These are expected behaviors. They occur only if you apply a firewall filter to a private VLAN in the output direction and do not occur if you apply a firewall filter to a private VLAN in the input direction.

---

#### Egress Filtering of L2PT Traffic Not Supported

**Problem** **Description:** Egress filtering of L2PT traffic is not supported on the QFX3500 switch. That is, if you configure L2PT to tunnel a protocol on an interface, you cannot also use a firewall filter to filter traffic for that protocol on that interface in the output direction. If you commit a configuration for this purpose, the firewall filter is not applied to the L2PT-tunneled traffic.

**Solution** This is expected behavior.

---

#### Cannot Drop BGP Packets in Certain Circumstances

**Problem** **Description:** BGP packets with a time-to-live (TTL) value greater than 1 cannot be discarded using a firewall filter applied to a loopback interface or applied on input to a Layer 3 interface. BGP packets with TTL value of 1 or 0 can be discarded using a firewall filter applied to a loopback interface or applied on input to a Layer 3 interface.

**Solution** This is expected behavior.

---

#### Invalid Statistics for Policer

**Problem** **Description:** If you apply a single-rate two-color policer in more than 128 terms in a firewall filter, the output of the **show firewall** command displays incorrect data for the policer.

**Solution** This is expected behavior.

---

#### Policers can Limit Egress Filters

**Problem** **Description:** On some switches, the number of egress policers that you configure can affect the total number of allowed egress firewall filters. (This issue does not affect

QFX10000 switches.) On the affected switches, every policer has two implicit counters that consume two entries in a 1024-entry TCAM that is used for counters, including counters that are configured as action modifiers in firewall filter terms. (Policers consume two entries because one is used for green packets and one is used for nongreen packets regardless of policer type.) If the TCAM becomes full, you cannot commit any more egress firewall filters that have terms with counters. For example, if you configure and commit 512 egress policers (two-color, three-color, or a combination of both policer types), all of the memory entries for counters are used up. If later in your configuration file you insert additional egress firewall filters with terms that also include counters, *none* of the terms in those filters are committed because there is no available memory space for the counters.

Here are some additional examples:

- Assume that you configure egress filters that include a total of 512 policers and no counters. Later in your configuration file you include another egress filter with 10 terms, 1 of which has a counter action modifier. None of the terms in this filter are committed because there is not enough TCAM space for the counter.
- Assume that you configure egress filters that include a total of 500 policers, so 1000 TCAM entries are occupied. Later in your configuration file you include the following two egress filters:
  - Filter A with 20 terms and 20 counters. All the terms in this filter are committed because there is enough TCAM space for all the counters.
  - Filter B comes after Filter A and has five terms and five counters. *None* of the terms in this filter are committed because there is not enough memory space for *all* the counters. (Five TCAM entries are required but only four are available.)

**Solution** You can prevent this problem by ensuring that egress firewall filter terms with counter actions are placed earlier in your configuration file than terms that include policers. In this circumstance, Junos OS commits policers even if there is not enough TCAM space for the implicit counters. For example, assume the following:

- You have 1024 egress firewall filter terms with counter actions.
- Later in your configuration file you have an egress filter with 10 terms. None of the terms have counters but one has a policer action modifier.

You can successfully commit the filter with 10 terms even though there is not enough TCAM space for the implicit counters of the policer. The policer is committed without the counters.

## PART 79

# Policers

- [Configuring Policers on page 5999](#)



# Configuring Policers

- [Overview of Policers on page 5999](#)
- [Understanding Policers with Link Aggregation Groups on page 6005](#)
- [Understanding Color-Blind Mode for Single-Rate Tricolor Marking on page 6005](#)
- [Understanding Color-Aware Mode for Single-Rate Tricolor Marking on page 6006](#)
- [Understanding Color-Blind Mode for Two-Rate Tricolor Marking on page 6008](#)
- [Understanding Color-Aware Mode for Two-Rate Tricolor Marking on page 6008](#)
- [Example: Using Two-Color Policers and Prefix Lists on page 6010](#)
- [Example: Using Policers to Manage Oversubscription on page 6013](#)
- [Assigning Forwarding Classes and Loss Priority on page 6015](#)
- [Configuring Color-Blind Egress Policers for Medium-Low PLP on page 6016](#)
- [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 6017](#)
- [Verifying That Two-Color Policers Are Operational on page 6019](#)
- [Verifying That Three-Color Policers Are Operational on page 6020](#)
- [Troubleshooting Policer Configuration on page 6020](#)

## Overview of Policers

---

A switch polices traffic by limiting the input or output transmission rate of a class of traffic according to user-defined criteria. Policing (or rate-limiting) traffic allows you to control the maximum rate of traffic sent or received on an interface and to provide multiple priority levels or classes of service.

- [Policer Overview on page 6000](#)
- [Policer Types on page 6000](#)
- [Policer Actions on page 6001](#)
- [Policer Colors on page 6002](#)
- [Filter-Specific Policers on page 6002](#)
- [Suggested Naming Convention for Policers on page 6003](#)
- [Policer Counters on page 6003](#)
- [Policer Algorithms on page 6003](#)


- [How Many Policers Are Supported? on page 6004](#)
- [Policies Can Limit Egress Firewall Filters on page 6004](#)

## Policer Overview

You use policers to apply limits to traffic flow and set consequences for packets that exceed these limits—usually applying a higher loss priority—so that if packets encounter downstream congestion, they can be discarded first. Policers apply only to unicast packets.

Policers provide two functions: metering and marking. A policer meters (measures) each packet against traffic rates and burst sizes that you configure. It then passes the packet and the metering result to the marker, which assigns a packet loss priority that corresponds to the metering result. [Figure 198](#) illustrates this process.

**Figure 198: Flow of Tricolor Marking Policer Operation**



After you name and configure a policer, you can use it by specifying it as an action in one or more firewall filters.

## Policer Types

A switch supports three types of policers:

- **Single-rate two-color marker**—A two-color policer (or “policer” when used without qualification) meters the traffic stream and classifies packets into two categories of packet loss priority (PLP) according to a configured bandwidth and burst-size limit. You can mark packets that exceed the bandwidth and burst-size limit with a specified PLP or simply discard them.

You can specify this type of policer in an ingress or egress firewall.



**NOTE:** A two-color policer is most useful for metering traffic at the port (physical interface) level.

- **Single-rate three-color marker**—This type of policer is defined in RFC 2697, *A Single Rate Three Color Marker*, as part of an assured forwarding (AF) per-hop-behavior (PHB) classification system for a Differentiated Services (DiffServ) environment. This type of policer meters traffic based on one rate—the configured committed information rate (CIR) as well as the committed burst size (CBS) and the excess burst size (EBS). The CIR specifies the average rate at which bits are admitted to the switch. The CBS specifies the usual burst size in bytes and the EBS specifies the maximum burst size in bytes. The EBS must be greater than or equal to the CBS, and neither can be 0.



You can specify this type of policer in an ingress or egress firewall.



**NOTE:** A single-rate three-color marker (TCM) is most useful when a service is structured according to packet length and not peak arrival rate.

- **Two-rate three-color marker**—This type of policer is defined in RFC 2698, *A Two Rate Three Color Marker*, as part of an assured forwarding per-hop-behavior classification system for a Differentiated Services environment. This type of policer meters traffic based on two rates—the CIR and peak information rate (PIR) along with their associated burst sizes, the CBS and peak burst size (PBS). The PIR specifies the maximum rate at which bits are admitted to the network and must be greater than or equal to the CIR.

You can specify this type of policer in an ingress or egress firewall.



**NOTE:** A two-rate three-color policer is most useful when a service is structured according to arrival rates and not necessarily packet length.

See [Table 479](#) for information about how metering results are applied for each of these policer types.

## Policer Actions

Policer actions are implicit or explicit and vary by policer type. *Implicit* means that Junos OS assigns the loss priority automatically. [Table 479](#) describes the policer actions.

**Table 479: Policer Actions**

| Policer                 | Marking                        | Implicit Action                  | Configurable Action |
|-------------------------|--------------------------------|----------------------------------|---------------------|
| Single-rate two-color   | Green (conforming)             | Assign low loss priority         | None                |
|                         | Red (nonconforming)            | None                             | Discard             |
| Single-rate three-color | Green (conforming)             | Assign low loss priority         | None                |
|                         | Yellow (above the CIR and CBS) | Assign medium-high loss priority | None                |
|                         | Red (above the EBS)            | Assign high loss priority        | Discard             |

Table 479: Policer Actions (*continued*)

| Policer              | Marking                        | Implicit Action                  | Configurable Action |
|----------------------|--------------------------------|----------------------------------|---------------------|
| Two-rate three-color | Green (conforming)             | Assign low loss priority         | None                |
|                      | Yellow (above the CIR and CBS) | Assign medium-high loss priority | None                |
|                      | Red (above the PIR and PBS)    | Assign high loss priority        | Discard             |



**NOTE:** If you specify a policer in an egress firewall filter, the only supported action is **discard**.

## Policer Colors

Single-rate and two-rate three-color policers can operate in two modes:

- **Color-blind**—In color-blind mode, the three-color policer assumes that all packets examined have not been previously marked or metered. In other words, the three-color policer is “blind” to any previous coloring a packet might have had.
- **Color-aware**—In color-aware mode, the three-color policer assumes that all packets examined have been previously marked or metered. In other words, the three-color policer is “aware” of the previous coloring a packet might have had. In color-aware mode, the three-color policer can increase the PLP of a packet but cannot decrease it. For example, if a color-aware three-color policer meters a packet with a medium PLP marking, it can raise the PLP level to high but cannot reduce the PLP level to low.

## Filter-Specific Policers

You can configure policers to be filter-specific, which means that Junos OS creates only one policer instance regardless of how many times the policer is referenced. When you do this on some QFX switches, rate limiting is applied in aggregate, so if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms, the total bandwidth allowed by the filter is 1 Gbps. However, the behavior of a filter-specific policer is affected by how the firewall filter terms that reference the policer are stored in TCAM. If you create a filter-specific policer and reference it in multiple firewall filter terms, the policer allows more traffic than expected if the terms are stored in different TCAM slices. For example, if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms that are stored in three separate memory slices, the total bandwidth allowed by the filter is 3 Gbps, not 1 Gbps. (This behavior does not occur in QFX10000 switches.)

To prevent this unexpected behavior from occurring, use the information about TCAM slices presented in [“Planning the Number of Firewall Filters to Create” on page 5973](#) to

organize your configuration file so that all the firewall filter terms that reference a given filter-specific policer are stored in the same TCAM slice.

## Suggested Naming Convention for Policers

We recommend that you use the naming convention ***policer#-color type*** when configuring three-color policers and ***policer#*** when configuring two-color policers. TCM stands for three-color marker. Because policers can be numerous and must be applied correctly to work, a simple naming convention makes it easier to apply the policers properly. For example, the first single-rate, color-aware three-color policer configured would be named ***srTCM1-ca***. The second two-rate, color-blind three-color configured would be named ***trTCM2-cb***. The elements of this naming convention are explained below:

- sr (single-rate)
- tr (two-rate)
- TCM (tricolor marking)
- 1 or 2 (number of marker)
- ca (color-aware)
- cb (color-blind)

## Policer Counters

On some QFX switches, each policer that you configure includes an implicit counter that counts the number of packets that exceed the rate limits that are specified for the policer. If you use the same policer in multiple terms—either within the same filter or in different filters—the implicit counter counts all the packets that are policed in all of these terms and provides the total amount. (This does not apply to QFX10000 switches.) If you want to obtain separate packet counts for each term on an affected switch, use these options:

- Configure a unique policer for each term.
- Configure only one policer, but use a unique, explicit counter in each term.

## Policer Algorithms

Policing uses the *token-bucket algorithm*, which enforces a limit on average bandwidth while allowing bursts up to a specified maximum value. It offers more flexibility than the *leaky bucket algorithm* in allowing a certain amount of bursty traffic before it starts discarding packets.



**NOTE:** In an environment of light bursty traffic, QFX5200 might not replicate all multicast packets to two or more downstream interfaces. This occurs only at a line rate burst—if traffic is consistent, the issue does not occur. In addition, the issue occurs only when packet size increases beyond 6k in a one gigabit traffic flow.

## How Many Policers Are Supported?

QFX10000 switches support 8K policers (all policer types). QFX5100 and QFX5200 switches support 1535 ingress policers and 1024 egress policers (assuming one policer per firewall filter term).

QFX3500 and QFX3600 standalone switches and QFabric Node devices support the following numbers of policers (assuming one policer per firewall filter term):

- Two-color policers used in ingress firewall filters: 767
- Three-color policers used in ingress firewall filters: 767
- Two-color policers used in egress firewall filters: 1022
- Three-color policers used in egress firewall filters: 512

## Policers Can Limit Egress Firewall Filters

On some switches, the number of egress policers that you configure can affect the total number of allowed egress firewall filters. (This issue does not affect QFX10000 switches.) On the affected switches, every policer has two implicit counters that consume two entries in a 1024-entry TCAM that is used for counters, including counters that are configured as action modifiers in firewall filter terms. (Policers consume two entries because one is used for green packets and one is used for nongreen packets regardless of policer type.) If the TCAM becomes full, you cannot commit any more egress firewall filters that have terms with counters. For example, if you configure and commit 512 egress policers (two-color, three-color, or a combination of both policer types), all of the memory entries for counters are used up. If later in your configuration file you insert additional egress firewall filters with terms that also include counters, *none* of the terms in those filters are committed because there is no available memory space for the counters.

Here are some additional examples:

- Assume that you configure egress filters that include a total of 512 policers and no counters. Later in your configuration file you include another egress filter with 10 terms, 1 of which has a counter action modifier. None of the terms in this filter are committed because there is not enough TCAM space for the counter.
- Assume that you configure egress filters that include a total of 500 policers, so 1000 TCAM entries are occupied. Later in your configuration file you include the following two egress filters:
  - Filter A with 20 terms and 20 counters. All the terms in this filter are committed because there is enough TCAM space for all the counters.
  - Filter B comes after Filter A and has five terms and five counters. *None* of the terms in this filter are committed because there is not enough memory space for *all* the counters. (Five TCAM entries are required but only four are available.)

You can prevent this problem by ensuring that egress firewall filter terms with counter actions are placed earlier in your configuration file than terms that include policers. In

this circumstance, Junos OS commits policers even if there is not enough TCAM space for the implicit counters. For example, assume the following:

- You have 1024 egress firewall filter terms with counter actions.
- Later in your configuration file you have an egress filter with 10 terms. None of the terms have counters but one has a policer action modifier.

You can successfully commit the filter with 10 terms even though there is not enough TCAM space for the implicit counters of the policer. The policer is committed without the counters.

**Related  
Documentation**

- [Understanding Color-Blind Mode for Single-Rate Tricolor Marking on page 6005](#)
- [Understanding Color-Blind Mode for Two-Rate Tricolor Marking on page 6008](#)
- [Understanding Color-Aware Mode for Single-Rate Tricolor Marking on page 6006](#)
- [Understanding Color-Aware Mode for Two-Rate Tricolor Marking on page 6008](#)
- [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 6017](#)

---

## Understanding Policers with Link Aggregation Groups

If you apply a policer to a link aggregation group (LAG) on a standalone switch or QFabric node, the policer applies to all the interfaces in the LAG in aggregate. For example, if you configure a policer to rate-limit at 1 Gbps and apply the policer (by using a firewall filter) to a LAG that has two member interfaces on a single switch or node, the total allowed throughput for both members is 1 Gbps.

If you apply a policer to a LAG that has members on different nodes in a QFabric network Node group or redundant server Node group, the configured rate applies to the interface on each node. For example, if you configure a policer to rate-limit at 1 Gbps and apply the policer to a LAG that has one member on server node A and one member on server node B, the allowed throughput for each member is 1 Gbps, for a total allowed throughput of 2 Gbps.

**Related  
Documentation**

- [Overview of Policers on page 5999](#)
- [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 6017](#)

---

## Understanding Color-Blind Mode for Single-Rate Tricolor Marking

With the color-blind mode of single-rate tricolor marking, all packets are evaluated against the CBS. If a packet exceeds the CBS, it is evaluated against the EBS. In color-blind mode, the policer supports three loss priorities only: low, medium-high, and high.

Packets that exceed the CBS but are below the EBS are marked yellow (medium-high). Packets that exceed the EBS are marked red (high), as shown in [Table 480](#).

Table 480: Color-Blind Mode TCM Color-to-PLP Mapping

| Color  | PLP         | Meaning                                                     |
|--------|-------------|-------------------------------------------------------------|
| Green  | low         | Conforming.                                                 |
| Yellow | medium-high | Packet exceeds the CIR and CBS but does not exceed the EBS. |
| Red    | high        | Packet exceeds the EBS.                                     |

- Related Documentation**
- [Overview of Policers on page 5999](#)
  - [Configuring Color-Blind Egress Policers for Medium-Low PLP on page 6016](#)

## Understanding Color-Aware Mode for Single-Rate Tricolor Marking

In color-aware mode, the treatment the packet receives depends on its classification. Marking can increase a preassigned PLP but cannot decrease it.

### Summary of PLP Changes

Table 481 shows how a packet's incoming priority can be modified with single-rate marking.

Table 481: Color-Aware Mode Single-Rate PLP Mapping

| Incoming PLP | Packet Metered Against      | Possible Cases                                              | Outgoing PLP |
|--------------|-----------------------------|-------------------------------------------------------------|--------------|
| low          | CIR, CBS, and EBS           | Conforming                                                  | low          |
|              |                             | Packet exceeds the CIR and CBS but does not exceed the EBS. | medium-high  |
|              |                             | Packet exceeds the EBS.                                     | high         |
| medium-low   | EBS only                    | Packet does not exceed the EBS.                             | medium-low   |
|              |                             | Packet exceeds the EBS.                                     | high         |
| medium-high  | EBS only                    | Packet does not exceed the EBS.                             | medium-high  |
|              |                             | Packet exceeds the EBS.                                     | high         |
| high         | Not metered by the policer. | All cases.                                                  | high         |

The following sections describe single-rate color-aware PLP mapping in more detail.

### Effect on Green Packets (Low PLP)

---

Packets belonging to the green class have already been marked by a classifier with low PLP. The marking policer can leave the PLP unchanged or increase it to medium-high or high, so these packets are therefore metered against both the CBS and the EBS. For example, if a behavior aggregate or multifield classifier marks a packet with low PLP and the two-rate TCM policer is in color-aware mode, the output loss priority is as follows:

- If the rate of traffic flow is less than the CIR, packets remain marked as low PLP.
- If bursts exceed the CBS but not the EBS, some of the packets are marked as medium-high PLP, and some of the packets remain marked as low PLP.
- If bursts exceed the EBS, some of the packets are marked as high PLP, and some of the packets remain marked as low PLP.

### Effect on Yellow Packets (Medium PLP)

---

Packets belonging to the yellow class have already been marked by a classifier with medium-low or medium-high PLP. The marking policer can leave the PLP unchanged or increase it to high, so these packets are therefore metered against the EBS only. For example, if a behavior aggregate or multifield classifier marks a packet with medium-low PLP and the two-rate TCM policer is in color-aware mode, the output loss priority is as follows:

- If the rate of traffic flow is less than the CBS, the packets remain marked as medium-low PLP.
- If the rate of traffic flow is greater than the CBS but less than the EBS, the packets remain marked as medium-low PLP.
- If the rate of traffic flow is greater than the EBS, some of the packets are marked as high PLP and some remain marked as medium-low PLP.

If a BA or multifield classifier marks a packet with medium-high PLP and the two-rate TCM policer is in color-aware mode, the policer assigns output loss priority as follows:

- If the rate of traffic flow is less than the CBS, the packets remain marked as medium-high PLP.
- If the rate of traffic flow is greater than the CBS but less than the EBS, the packets remain marked as medium-high PLP.
- If the rate of traffic flow is greater than the EBS, some of the packets are marked as high PLP and some remain marked as medium-high PLP.

### Effect on Red Packets (High PLP)

---

Packets belonging to the red class have already been marked by a classifier with high PLP. Because the policer cannot decrease the PLP, it does not change it, and these packets are not metered against the CBS or the EBS.

- Related Documentation**
- [Overview of Policers on page 5999](#)
  - [Configuring Color-Blind Egress Policers for Medium-Low PLP on page 6016](#)

## Understanding Color-Blind Mode for Two-Rate Tricolor Marking

With the color-blind mode of two-rate tricolor marking, all packets are evaluated against the committed information rate (CIR). If a packet exceeds the CIR, it is evaluated against the peak information rate (PIR). Packets that exceed the CIR but are below the PIR are marked yellow (medium-high). Packets that exceed the PIR are marked red (high).

**Table 482: Color-Blind Mode TCM Color-to-PLP Mapping**

| Color  | PLP                | Meaning                                             |
|--------|--------------------|-----------------------------------------------------|
| Green  | <b>low</b>         | Packet does not exceed the CIR.                     |
| Yellow | <b>medium-high</b> | Packet exceeds the CIR but does not exceed the PIR. |
| Red    | <b>high</b>        | Packet exceeds the PIR.                             |

- Related Documentation**
- [Overview of Policers on page 5999](#)
  - [Configuring Color-Blind Egress Policers for Medium-Low PLP on page 6016](#)

## Understanding Color-Aware Mode for Two-Rate Tricolor Marking

In color-aware mode, the treatment the packet receives depends on its classification. Marking can increase the preassigned PLP but cannot decrease it

### Summary of PLP Changes

Table 483 shows how a packet's incoming priority can be modified with two-rate marking.

**Table 483: Color-Aware Mode Two-Rate PLP Mapping**

| Incoming PLP      | Packet Metered Against | Possible Cases                          | Outgoing PLP       |
|-------------------|------------------------|-----------------------------------------|--------------------|
| <b>low</b>        | CIR and PIR            | Packet does not exceed the CIR.         | <b>low</b>         |
|                   |                        | Packet exceeds the CIR but not the PIR. | <b>medium-high</b> |
|                   |                        | Packet exceeds the PIR.                 | <b>high</b>        |
| <b>medium-low</b> | PIR only               | Packet does not exceed the PIR.         | <b>medium-low</b>  |
|                   |                        | Packet exceeds the PIR.                 | <b>high</b>        |



Table 483: Color-Aware Mode Two-Rate PLP Mapping (*continued*)

| Incoming PLP | Packet Metered Against      | Possible Cases                  | Outgoing PLP |
|--------------|-----------------------------|---------------------------------|--------------|
| medium-high  | PIR only                    | Packet does not exceed the PIR. | medium-high  |
|              |                             | Packet exceeds the PIR.         | high         |
| high         | Not metered by the policer. | All cases.                      | high         |

The following sections describe color-aware two-rate PLP mapping in more detail.

### Effect on Green Packets (Low PLP)

Packets belonging to the green class have already been marked by a classifier with low PLP. The marking policer can leave the packet's PLP unchanged or increase the PLP to medium-high or high. These packets are therefore metered against both the CIR and the PIR. For example, if a behavior aggregate or multifield classifier marks a packet with low PLP and the two-rate TCM policer is in color-aware mode, the output loss priority is as follows:

- If the rate of traffic flow is less than the CIR, the packets remain marked as low PLP.
- If the rate of traffic flow is greater than the CIR but less than the PIR, some of the packets are marked as medium-high PLP and some of the packets remain marked as low PLP.
- If the rate of traffic flow is greater than the PIR, some of the packets are marked as high PLP and some of the packets remain marked as low PLP.

### Effect on Yellow Packets (Medium PLP)

Packets belonging to the yellow class have already been marked by a classifier with medium-low or medium-high PLP. The marking policer can leave the PLP unchanged or increase it to high. These packets are therefore metered against the PIR only. For example, if a behavior aggregate (BA) or multifield classifier marks a packet with medium-low PLP and the two-rate TCM policer is in color-aware mode, the policer assigns output loss priority as follows:

- If the rate of traffic flow is less than the CIR, the packets remain marked as medium-low PLP.
- If the rate of traffic flow is greater than the CIR but less than the PIR, the packets remain marked as medium-low PLP.
- If the rate of traffic flow is greater than the PIR, some of the packets are marked as high PLP and some of the packets remain marked as medium-low PLP.

If a BA or multifield classifier marks a packet with medium-high PLP and the two-rate TCM policer is in color-aware mode, the policer assigns output loss priority as follows:

- If the rate of traffic flow is less than the CIR, the packets remain marked as medium-high PLP.
- If the rate of traffic flow is greater than the CIR but less than the PIR, the packets remain marked as medium-high PLP.
- If the rate of traffic flow is greater than the PIR, some of the packets are marked as high PLP and some of the packets remain marked as medium-high PLP.

### Effect on Red Packets (High PLP)

Packets belonging to the red class have already been marked by a classifier with high PLP. Because the policer cannot decrease the PLP, it does not change it, and these packets are not metered against the CIR or the PIR.

#### Related Documentation

- [Overview of Policers on page 5999](#)
- [Configuring Color-Blind Egress Policers for Medium-Low PLP on page 6016](#)

---

## Example: Using Two-Color Policers and Prefix Lists

If you provide specific amounts of bandwidth to internal or external customers, you can use policing to make sure that customers do not consume more bandwidth than they should receive. For example, you might connect many customers to one 10-Gbps interface and want to ensure that none of them congest the interface by using more bandwidth than they have been allotted.

You could accomplish this by creating a two-color policer similar to the following for each customer:

```
firewall {
 policer Limit-Customer-1 {
 if-exceeding {
 bandwidth-limit 100m;
 burst-size-limit 150m;
 }
 then discard;
 }
}
```

Creating a policer for each customer is clearly not a scalable solution, however. As an alternative, you can create prefix lists that group classes of customers and then create policers for each prefix list. For example, you could create prefix lists such as **Class-A-Customer-Prefixes**, **Class-B-Customer-Prefixes**, and **Class-C-Customer-Prefixes** (at the **[edit policy-options]** hierarchy level) and create the following corresponding policers:

```
firewall {
 policer Class-A {
 if-exceeding {
 bandwidth-limit 100m;
 burst-size-limit 150m;
 }
 then discard;
 }
}
```

```

}
policer Class-B {
 if-exceeding {
 bandwidth-limit 75m;
 burst-size-limit 100m;
 }
 then discard;
}
policer Class-C {
 if-exceeding {
 bandwidth-limit 50m;
 burst-size-limit 75m;
 }
 then discard;
}
}

```

You must create filter terms that specify the prefix lists in their **from** statements and the corresponding policers in their **then** statements similar to the following:

```

firewall
family inet {
 filter Class-A-Customers {
 term term-1 {
 from {
 destination-prefix-list {
 Class-A-Customer-Prefixes;
 }
 }
 then policer Class-A;
 }
 }
 filter Class-B-Customers {
 term term-1 {
 from {
 destination-prefix-list {
 Class-B-Customer-Prefixes;
 }
 }
 then policer Class-B;
 }
 }
 filter Class-C-Customers {
 term term-1 {
 from {
 destination-prefix-list {
 Class-C-Customer-Prefixes;
 }
 }
 then policer Class-C;
 }
 }
}
}

```

Here are the steps to create this firewall configuration:

1. Create the first policer:

```
[edit firewall]
user@switch# set policer Class-A if-exceeding bandwidth-limit 100m burst-size-limit 150m
user@switch# set policer Class-A then discard
```

2. Create the second policer:

```
[edit firewall]
user@switch# set policer Class-B if-exceeding bandwidth-limit 75m burst-size-limit 100m
user@switch# set policer Class-B then discard
```

3. Create the third policer:

```
[edit firewall]
user@switch# set policer Class-C if-exceeding bandwidth-limit 50m burst-size-limit 75m
user@switch# set policer Class-C then discard
```

4. Create a filter for class A customers:

```
[edit firewall]
user@switch# edit family inet filter Class-A-Customers
```

5. Configure the filter to send packets matching the Class-A-Customer-Prefixes prefix list to the Class-A policer:

```
[edit firewall family inet filter Class-A-Customers]
user@switch# set term term-1 from source-prefix-list Class-A-Customers
user@switch# set term term-1 then policer Class-A
```

6. Create a filter for class B customers:

```
[edit firewall]
user@switch# edit family inet filter Class-B-Customers
```

7. Configure the filter to send packets matching the Class-B-Customer-Prefixes prefix list to the Class-B policer:

```
[edit firewall family inet filter Class-B-Customers]
user@switch# set term term-1 from source-prefix-list Class-B-Customers
user@switch# set term term-1 then policer Class-B
```

8. Create a filter for class C customers:

```
[edit firewall]
user@switch# edit family inet filter Class-C-Customers
```

9. Configure the filter to send packets matching the Class-C-Customer-Prefixes prefix list to the Class-C policer:

```
[edit firewall family inet filter Class-C-Customers]
user@switch# set term term-1 from source-prefix-list Class-C-Customers
user@switch# set term term-1 then policer Class-C
```

10. Apply the filters you created to the appropriate interfaces in the output direction.



**NOTE:** Note that the implicit deny statement in this filter will block traffic from any source that does not match one of the prefix lists. If you want the filter to allow this traffic, you must include an explicit term for this purpose.

- Related Documentation**
- [Overview of Policers on page 5999](#)
  - [prefix-list](#)

## Example: Using Policers to Manage Oversubscription

You might want to use a policer when an interface is oversubscribed and you want to control what will happen if congestion occurs. For example, you might have servers connected to a switch as listed in [Table 484](#).

**Table 484: Servers Connected to Switch**

| Server Type                | Connection           | IP Address |
|----------------------------|----------------------|------------|
| Network application server | 1-gigabit interface  | 10.0.0.1   |
| Authentication server      | 1-gigabit interface  | 10.0.0.2   |
| Database server            | 10-gigabit interface | 10.0.0.3   |

In this example, users access services provided by the network application server, which requests information from the database server as appropriate. When it receives a request from a user, the network application server first contacts the authentication server to verify the user's credentials. When a user is authenticated and the network application server provides the requested service, all the packets sent from the database server to the application server must transit the 1-Gigabit Ethernet interface connected to the application server twice—once on ingress to the application server and again on egress to the user.

The sequence of events for a user session is as follows:

1. A user connects to the application server and requests a service.
2. The application server requests the user's credentials and relays them to the authentication server.
3. If the authentication server verifies the credentials, the application server initiates the requested service.
4. The application server requests the files necessary to meet the user's request from the database server.
5. The database server sends the requested files to the application server.
6. The application server includes the requested files in its response to the user.

Traffic from the database server to the application server might congest the 1-gigabit interface to which that the application server is connected. This congestion might prevent the server from responding to requests from users and creating new sessions for them. You can use policing to make sure that this does not occur.

To create this firewall configuration, perform the following steps on the database server:

1. Create a policer to drop traffic from the database server to the application server if it exceeds certain limits:

```
[edit firewall]
user@switch# set policer Database-Egress-Policer if-exceeding bandwidth-limit 400
burst-size-limit 500m
user@switch# set policer Database-Egress-Policer then discard
```

2. Create a filter to examine traffic from the database server to the application server:

```
[edit firewall]
user@switch# edit family inet filter Database-Egress-Filter
```

3. Configure the filter to apply the policer to traffic egressing the database server and destined for the application server:

```
[edit firewall family inet filter Database-Egress-Filter]
user@switch# set term term-1 from destination-address 10.0.0.1
user@switch# set term term-1 then policer Database-Egress-Policer
```

4. If required, configure a term to allow traffic from the database server to other destinations (otherwise the traffic will be dropped by the implicit deny statement):

```
[edit firewall family inet filter Database-Egress-Filter]
user@switch# set term term-2 then accept
```

Note that omitting a **from** statement causes the term to match all packets, which is the desired behavior.

5. Install the egress filter as an output filter on the database server interface that is connected the application server:

```
[edit interfaces]
user@switch# set xe-0/0/3 unit 0 family inet filter output Database-Egress-Filter
```

Here is how the final configuration would appear:

```
firewall {
 policer Database-Egress-Policer {
 if-exceeding {
 bandwidth-limit 400;
 burst-size-limit 500m;
 }
 then discard;
 }
 family inet {
 filter Database-Egress-Filter {
 term term-1 {
 from {
 destination-address {
 10.0.0.1/24;
 }
 }
 then policer Database-Egress-Policer;
 }
 term term-2 { # If required, include this term so that traffic from the database server
 # to other destinations is allowed.
 then accept;
 }
 }
 }
}
```

```
}
]
```

#### Related Documentation

- [Overview of Policers on page 5999](#)

## Assigning Forwarding Classes and Loss Priority

You can configure firewall filters to assign packet loss priority (PLP) and forwarding classes so that if congestion occurs, the marked packets can be dropped according to the priority you set. The valid match conditions are one or more of the six packet header fields: destination address, source address, IP protocol, source port, destination port, and DSCP. In other words, you can set the forwarding class and the PLP for each packet entering or an interface with a specific destination address, source address, IP protocol, source port, destination port, or DSCP.



**NOTE:** Junos OS assigns forwarding classes and PLP on ingress only. Do not use a filter that assigns forwarding classes or PLP as an egress filter.

When tricolor marking is enabled, a switch supports four PLP designations: **low**, **medium-low**, **medium-high**, and **high**. You can also specify any of the forwarding classes listed in [Table 485](#)

**Table 485: Unicast Forwarding Classes**

| Unicast Forwarding Class | For CoS Traffic Type                                               |
|--------------------------|--------------------------------------------------------------------|
| <b>be</b>                | Best-effort traffic                                                |
| <b>no-loss</b>           | Guaranteed delivery for TCP traffic                                |
| <b>fcoe</b>              | Guaranteed delivery for Fibre Channel over Ethernet (FCoE) traffic |
| <b>nc</b>                | Network-control traffic                                            |

To assign forwarding classes in firewall filters:

1. Configure the family address type and filter name:

```
[edit]
user@switch# edit firewall family ethernet-switching filter ingress-filter
```

2. Configure the terms of the filter as appropriate, including the **forwarding-class** and **loss-priority** action modifiers. For example, each of the following terms in the filter examines various packet header fields and assigns the appropriate forwarding class and packet loss priority:

- The term **corp-traffic** matches all IPv4 packets with a **10.1.1.0/24** source address and assigns the packets to forwarding class **no-loss** with a loss priority of **low**:

```
[edit firewall family ethernet-switching filter ingress-filter]
user@switch# set term corp-traffic from source-address 10.1.1.0/24;
```

```
user@switch# set term corp-traffic then forwarding-class no-loss
user@switch# set term corp-traffic then loss-priority low
```

- The term **data-traffic** matches all IPv4 packets with a **10.1.2.0/24** source address and assigns the packets to forwarding class **be** (best effort) with a loss priority of **medium-high**:

```
[edit firewall family ethernet-switching filter ingress-filter]
user@switch# set term data-traffic from source-address 10.1.2.0/24;
user@switch# set term data-traffic then forwarding-class be
user@switch# set term data-traffic then loss-priority medium-high
```

- Because the loss of network-generated packets can jeopardize proper network operation, the delay of these packets is preferable to discarding these packets. The term **network-traffic** assigns the packets with an IP precedence of **net-control** to forwarding class **nc** (network control) with a loss priority of **low**:

```
[edit firewall family ethernet-switching filter ingress-filter]
user@switch# set term network-traffic from precedence net-control
user@switch# set term network-traffic then forwarding-class nc
user@switch# set term network-traffic then loss-priority low
```

- The last term **accept-traffic** matches any packets that did not match on any of the preceding terms and assigns the packets to forwarding class **be** with a loss priority of **high**:

```
[edit firewall family ethernet-switching filter ingress-filter]
user@switch# set term accept-traffic then forwarding-class be
user@switch# set term accept-traffic then loss-priority high
```

3. Apply the filter **ingress-filter** to a port, VLAN, or Layer 3 interface. For information about applying the filter, see [“Configuring Firewall Filters” on page 5978](#). (Assigning forwarding classes and PLP is supported only on ingress filters.)

#### Related Documentation

- [Configuring Firewall Filters on page 5978](#)
- [Verifying That Firewall Filters Are Operational on page 5988](#)
- [Monitoring Firewall Filter Traffic on page 5987](#)
- [Overview of Policers on page 5999](#)
- [Understanding CoS Classifiers](#)
- [Understanding CoS Forwarding Classes](#)

---

## Configuring Color-Blind Egress Policers for Medium-Low PLP

If you use color-blind mode and want to configure an egress policer that marks packets to have medium-low PLP, you must configure a single-rate two-color policer at the **[edit firewall policer *policer-name*]** hierarchy level, because color-blind mode does not support medium-low priority. For example:

1. Specify the name of the policer, the bandwidth limit in bits per second (bps) to control the traffic rate on an interface, and the maximum allowed burst size to control the amount of traffic bursting:

```
[edit]
```



```
user@switch# set firewall policer policer-name if-exceeding bandwidth-limit bytes
burst-size-limit bytes
```

2. Specify medium-low loss priority for matching packets:

```
[edit]
```

```
user@switch# set firewall policer policer-name then loss-priority medium-low;
```

3. Apply the filter to a port, VLAN, or Layer 3 interface.

#### Related Documentation

- [Overview of Policers on page 5999](#)
- [Understanding Color-Blind Mode for Single-Rate Tricolor Marking on page 6005](#)
- [Understanding Color-Blind Mode for Two-Rate Tricolor Marking on page 6008](#)
- [Configuring Firewall Filters on page 5978](#)
- [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 6017](#)

## Configuring Two-Color and Three-Color Policers to Control Traffic Rates

You can rate-limit traffic by configuring a policer and specifying it as an action modifier for a term in a firewall filter. By default, if you specify the same policer in multiple terms, Junos OS creates a separate policer instance for each term and applies rate limiting separately for each instance. For example, if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms, each policer instance enforces a 1-Gbps limit. In this case, the total bandwidth allowed by the filter is 3 Gbps.

You can also configure a policer to be filter-specific, which means that Junos OS creates only one policer instance regardless of how many times the policer is referenced. When you do this, rate limiting is applied in aggregate, so if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms, the total bandwidth allowed by the filter is 1 Gbps.



**NOTE:** You can include two-color policer actions on ingress firewall filters only. You can include three-color policer actions on ingress and egress filters.

1. [Configuring Two-Color Policers on page 6017](#)
2. [Configuring Three-Color Policers on page 6018](#)
3. [Specifying Policers in a Firewall Filter Configuration on page 6018](#)
4. [Applying a Firewall Filter That Includes a Policer on page 6019](#)

### Configuring Two-Color Policers

To configure a two-color policer:

1. Specify the name of the policer, the bandwidth limit to control the traffic rate on an interface, and the maximum allowed burst size to control the amount of traffic bursting:

```
[edit firewall]
```

```
user@switch# set policer policer-name <filter-specific> if-exceeding bandwidth-limit bps
burst-size-limit bytes
```

The policer name can contain letters, numbers, and hyphens (-) and can have as many as 64 characters.

The range for the bandwidth limit is 32000 (32k) through 102,300,000,000 (102300m) bps.

To determine the value for the burst-size limit, multiply the bandwidth of the interface on which the filter is applied by the amount of time to allow a burst of traffic at that bandwidth to occur and divide the result by 8:

**maximum burst size = (interface bandwidth) X (allowable time for burst) / (8 bits/byte)**

The range for the burst-size limit is 1 through 2,147,450,880 bytes.

2. Specify the policer action to discard or assign a loss priority to packets that exceed the rate limits:

```
[edit firewall policer policer-name]
user@switch# set then (discard | loss-priority low | loss-priority high)
```

## Configuring Three-Color Policers

To configure a three-color policer:

1. Specify the name of the policer and (optionally) whether to automatically discard packets with high loss priority (PLP):

```
[edit firewall]
user@switch# set three-color-policer policer-name
user@switch# set three-color-policer policer-name action loss-priority high then discard
```

2. Specify whether the three-color policer should be single-rate or two-rate and whether it should be color-aware or color-blind:

```
[edit firewall three-color-policer policer-name]
user@switch# set (single-rate | two-rate) (color-aware | color-blind)
```

3. For single-rate three-color policers, configure the CIR, CBS, and EBS:

```
[edit firewall three-color-policer policer-name single-rate]
user@switch# set committed-information-rate bps
user@switch# set committed-burst-size bytes
user@switch# set excess-burst-size bytes
```

4. For two-rate three-color policers, configure the CIR, CBS, PIR, and PBS:

```
[edit firewall three-color-policer policer-name single-rate]
user@switch# set committed-information-rate bps
user@switch# set committed-burst-size bytes
user@switch# set peak-information-rate bps
user@switch# set peak-burst-size bytes
```

## Specifying Policers in a Firewall Filter Configuration

To use a two-color policer, configure a filter term that includes the action **policer**:

```
[edit firewall family family-name]
user@switch# set filter filter-name term name then name
```

For example, the following commands apply a two-color policer to all packets sent from 192.0.2.0/24.

```
[edit firewall family family-name]
user@switch# set filter limit—hosts term term1 from source-address 192.0.2.0/24
```

```
user@switch# set filter limit—hosts term term1 then policer policer1
```

To use a three-color policer, configure a filter term that includes the action **three-color-policer**:

```
[edit firewall family name]
```

```
user@switch# set filter name term name from match-condition
```

```
user@switch# set filter name term name then three-color-policer (single-rate | two-rate) name
```

For example, the following commands apply a single-rate three-color policer to all packets received or sent by interface **ge-0/0/6** (depending on whether the filter is an ingress or egress filter).

```
[edit firewall family name]
```

```
user@switch# set filter srTCM term term-one from interface ge-0/0/6
```

```
user@switch# set filter srTCM term term-one then three-color-policer single-rate srTCM1-ca
```

You must specify whether the three-color policer is single-rate or two-rate, and this must match the policer itself. Otherwise, the configuration listing includes an error message indicating that the three-color policer you referenced in the filter does not exist.

## Applying a Firewall Filter That Includes a Policer

A firewall filter that includes one or more policer action modifiers must be applied to a port, VLAN, or Layer 3 interface like any other filter. For information about applying firewall filters, see “[Configuring Firewall Filters](#)” on page 5978.



**NOTE:** You can include two-color policer actions on ingress firewall filters only. You can include three-color policer actions on ingress and egress filters.

### Related Documentation

- [Configuring Firewall Filters on page 5978](#)
- [Overview of Policers on page 5999](#)
- [Verifying That Two-Color Policers Are Operational on page 6019](#)
- [Verifying That Three-Color Policers Are Operational on page 6020](#)
- [Configuring Color-Blind Egress Policers for Medium-Low PLP on page 6016](#)

## Verifying That Two-Color Policers Are Operational

**Purpose** Verify that two-color policers in firewall filter configurations are working properly.

**Action** Use the **show firewall policer** operational mode command to verify that the policers are working properly:

```
user@switch> show firewall policer
Filter: egress-vlan-watch-employee
Filter: ingress-port-filter
Filter: ingress-port-limit-tcp-icmp
Policies:
Name Packets
icmp-connection-policer 10
tcp-connection-policer 539
Filter: ingress-vlan-rogue-block
Filter: ingress-vlan-limit-guest
```

**Meaning** The **show firewall policer** command displays the names of all firewall filters and policers that are configured. For each policer that is specified in a filter configuration, the output field shows the current packet count for all packets that exceed the specified rate limits.

**Related Documentation**

- [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 6017](#)
- [Configuring Firewall Filters on page 5978](#)
- [Monitoring Firewall Filter Traffic on page 5987](#)

---

## Verifying That Three-Color Policers Are Operational

**Purpose** Verify that three-color policers in firewall filter configurations are working properly.

**Action** Use the following operational mode commands to verify that a three-color policer is working properly:

- **show class-of-service forwarding-table classifiers**
- **show interfaces *interface-name* extensive**
- **show interfaces queue *interface-name***

**Related Documentation**

- [Overview of Policers on page 5999](#)
- [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 6017](#)

---

## Troubleshooting Policer Configuration

- [Incomplete Count of Packet Drops on page 6020](#)
- [Counter Reset When Editing Filter on page 6021](#)
- [Invalid Statistics for Policer on page 6021](#)
- [Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured on page 6021](#)
- [Filter-Specific Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured on page 6022](#)
- [Policers Can Limit Egress Filters on page 6023](#)

### Incomplete Count of Packet Drops

**Problem Description:** Under certain circumstances, Junos OS might display a misleading number of packets dropped by an ingress policer.

If packets are dropped because of ingress admission control, policer statistics might not show the number of packet drops you would expect by calculating the difference between ingress and egress packet counts. This might happen if you apply an ingress policer to multiple interfaces, and the aggregate ingress rate of those interfaces exceeds the line rate of a common egress interface. In this case, packets might be dropped from the

ingress buffer. These drops are not included in the count of packets dropped by the policer, which causes policer statistics to underreport the total number of drops.

**Solution** This is expected behavior.

### Counter Reset When Editing Filter

**Problem Description:** If you edit a firewall filter term, the value of any counter associated with any term in the same filter is set to 0, including the implicit counter for any policer referenced by the filter. Consider the following examples:

- Assume that your filter has **term1**, **term2**, and **term3**, and each term has a counter that has already counted matching packets. If you edit any of the terms in any way, the counters for all the terms are reset to 0.
- Assume that your filter has **term1** and **term2**. Also assume that **term2** has a **policer** action modifier and the implicit counter of the policer has already counted 1000 matching packets. If you edit **term1** or **term2** in any way, the counter for the policer referenced by **term2** is reset to 0.

**Solution** This is expected behavior.

### Invalid Statistics for Policer

**Problem Description:** If you apply a single-rate two-color policer in more than 128 terms in a firewall filter, the output of the **show firewall** command displays incorrect data for the policer.

**Solution** This is expected behavior.

### Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured

**Problem Description:** If you configure a policer to rate-limit throughput and apply it on egress to multiple interfaces on a QFX3500 switch or Node, the measured aggregate policed rate might be twice the configured rate, depending on which interfaces you apply the policer to. The doubling of the policed rate occurs if you apply a policer to multiple interfaces and *both* of the following are true:

- There is at least one policed interface in the range xe-0/0/0 to xe-0/0/23 or the range xe-0/1/1 to xe-0/1/7.
- There is at least one policed interface in the range xe-0/0/24 to xe-0/0/47 or the range xe-0/1/8 to xe-0/1/15.

For example, if you configure a policer to rate-limit traffic at 1 Gbps and apply the policer (by using a firewall filter) to xe-0/0/0 and xe-0/0/24 in the output direction, each interface is rate-limited at 1 Gbps, for a total allowed throughput of 2 Gbps. The same behavior occurs if you apply the policer to xe-0/1/1 and xe-0/0/24—each interface is rate-limited at 1 Gbps.

If you apply the same policer on egress to multiple interfaces in these groups, each *group* is rate-limited at 1 Gbps. For example, if you apply the policer to xe-0/0/0 through xe-0/0/4 (five interfaces) and xe-0/0/24 through xe-0/0/33 (ten interfaces), each group is rate-limited at 1 Gbps, for a total allowed throughput of 2 Gbps.

Here is another example: If you apply the policer to xe-0/0/0 through xe-0/0/4 and xe-0/1/1 through xe-0/1/5 (a total of ten interfaces), that group is rate-limited at 1 Gbps in aggregate. If you also apply the policer to xe-0/0/24, that one interface is rate-limited at 1 Gbps while the other ten are still rate-limited at 1 Gbps in aggregate.

Interfaces xe-0/1/1 through xe-0/1/15 are physically located on the QSFP+ uplink ports, according to the following scheme:

- xe-0/1/1 through xe-0/1/3 are on Q0.
- xe-0/1/4 through xe-0/1/7 are on Q1.
- xe-0/1/8 through xe-0/1/11 are on Q2.
- xe-0/1/12 through xe-0/1/15 are on Q3.

The doubling of the policed rate occurs only if the policer is applied in the output direction. If you configure a policer as described above but apply it in the input direction, the total allowed throughput for all interfaces is 1 Gbps.

**Solution** This is expected behavior.

## Filter-Specific Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured

**Problem** **Description:** You can configure policers to be filter-specific, which means that Junos OS creates only one policer instance regardless of how many times the policer is referenced. When you do this, rate limiting is applied in aggregate, so if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms, the total bandwidth allowed by the filter is 1 Gbps. However, the behavior of a filter-specific policer is affected by how the firewall filter terms that reference the policer are stored in ternary content addressable memory (TCAM). If you create a filter-specific policer and reference it in multiple firewall filter terms, the policer allows more traffic than expected if the terms are stored in different TCAM slices. For example, if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms that are stored in three separate memory slices, the total bandwidth allowed by the filter is 3 Gbps, not 1 Gbps.

**Solution** To prevent this unexpected behavior, use the information about TCAM slices presented in [“Planning the Number of Firewall Filters to Create” on page 5973](#) to organize your configuration file so that all the firewall filter terms that reference a given filter-specific policer are stored in the same TCAM slice.

## Policers Can Limit Egress Filters

**Problem** **Description:** On some switches, the number of egress policers that you configure can affect the total number of allowed egress firewall filters. (This issue does not affect QFX10000 switches.) On the affected switches, every policer has two implicit counters that consume two entries in a 1024-entry TCAM that is used for counters, including counters that are configured as action modifiers in firewall filter terms. (Policers consume two entries because one is used for green packets and one is used for nongreen packets regardless of policer type.) If the TCAM becomes full, you cannot commit any more egress firewall filters that have terms with counters. For example, if you configure and commit 512 egress policers (two-color, three-color, or a combination of both policer types), all of the memory entries for counters are used up. If later in your configuration file you insert additional egress firewall filters with terms that also include counters, *none* of the terms in those filters are committed because there is no available memory space for the counters.

Here are some additional examples:

- Assume that you configure egress filters that include a total of 512 policers and no counters. Later in your configuration file you include another egress filter with 10 terms, 1 of which has a counter action modifier. None of the terms in this filter are committed because there is not enough TCAM space for the counter.
- Assume that you configure egress filters that include a total of 500 policers, so 1000 TCAM entries are occupied. Later in your configuration file you include the following two egress filters:
  - Filter A with 20 terms and 20 counters. All the terms in this filter are committed because there is enough TCAM space for all the counters.
  - Filter B comes after Filter A and has five terms and five counters. *None* of the terms in this filter are committed because there is not enough memory space for *all* the counters. (Five TCAM entries are required but only four are available.)

**Solution** You can prevent this problem by ensuring that egress firewall filter terms with counter actions are placed earlier in your configuration file than terms that include policers. In this circumstance, Junos OS commits policers even if there is not enough TCAM space for the implicit counters. For example, assume the following:

- You have 1024 egress firewall filter terms with counter actions.
- Later in your configuration file you have an egress filter with 10 terms. None of the terms have counters but one has a policer action modifier.

You can successfully commit the filter with 10 terms even though there is not enough TCAM space for the implicit counters of the policer. The policer is committed without the counters.





## PART 80

# Configuring Port Security

- [Port Security on page 6027](#)



# Port Security

- [Overview of Access Port Protection on page 6027](#)
- [Understanding MAC Limiting and MAC Move Limiting for Port Security on page 6029](#)
- [Verifying That MAC Limiting Is Working Correctly on page 6031](#)
- [Verifying That MAC Move Limiting Is Working Correctly on page 6034](#)
- [Verifying That the Port Error Disable Setting Is Working Correctly on page 6034](#)
- [Configuring Persistent MAC Learning \(CLI Procedure\) on page 6035](#)
- [Understanding Trusted and Untrusted Ports on page 6037](#)
- [Understanding Trusted DHCP Servers for Port Security on page 6037](#)
- [Verifying That a Trusted DHCP Server Is Working Correctly on page 6037](#)
- [Understanding DHCP Option 82 for Port Security on page 6039](#)
- [Understanding Static ARP Entries on page 6041](#)

## Overview of Access Port Protection

---

Port security features on QFX10000 switches can protect a switch against various types of attacks. Protection methods against some common attacks are:

- [Mitigation of Ethernet Switching Table Overflow Attacks on page 6027](#)
- [Mitigation of Rogue DHCP Server Attacks on page 6028](#)
- [Protection Against DHCP Starvation Attacks on page 6028](#)

## Mitigation of Ethernet Switching Table Overflow Attacks

In an overflow attack on an Ethernet switching table, an intruder sends so many requests from new MAC addresses that the table cannot learn all the addresses. The attack forces the switch to send broadcast messages when it needs to send traffic to addresses for which it lacks MAC addresses. In addition to generating unnecessary traffic, the attacker might be able to sniff the broadcast packets.

To mitigate such attacks, you can configure a limit for learned MAC addresses or allow only specific MAC addresses. Use the MAC limit feature to control the total number of MAC addresses that can be added to the Ethernet switching table for the specified interface or interfaces. By setting the MAC addresses that are explicitly allowed, you

ensure that the addresses of network devices whose network access is critical are guaranteed to be included in the Ethernet switching table.

## Mitigation of Rogue DHCP Server Attacks

By default, all access ports are untrusted, and all trunk ports are trusted with regard to DHCP. Trusted ports allow DHCP servers to provide IP addresses and other information to requesting devices. If someone connects an unauthorized DHCP server to a trusted port, the unauthorized server can start issuing IP addresses and configuration information to the network's DHCP clients. The information provided to the clients by this server can disrupt their network access. The unauthorized server might also assign itself as the default gateway device for the network. An attacker can then sniff the network traffic and perpetrate a man-in-the-middle attack—that is, it misdirects traffic intended for a legitimate network device to a device of its choice.

To mitigate this problem, set the interface to which the unauthorized server is connected as untrusted. That action blocks all ingress DHCP server messages from that interface.



**NOTE:** The switch logs all DHCP server packets that are received on untrusted ports. For example:

```
5 untrusted DHCPOFFER received, interface xe-0/0/2.0[65], vlan v1[10] server
ip/mac 12.12.12.1/00:00:00:00:01:12 offer ip/client mac
12.12.12.253/00:AA:BB:CC:DD:01
```

You can use these messages to detect unauthorized DHCP servers on the network.



**NOTE:** If you attach a DHCP server to an access port, you must configure the port as trusted.

## Protection Against DHCP Starvation Attacks

In a DHCP starvation attack, an attacker floods an Ethernet LAN with DHCP requests from spoofed (counterfeit) MAC addresses so that trusted DHCP servers cannot keep up with requests from legitimate DHCP clients. The address space of those servers is completely used up, so they can no longer assign IP addresses and lease times to clients. DHCP requests from those clients are either dropped—that is, the result is a denial of service (DoS)—or directed to a rogue DHCP server set up by the attacker to imitate a legitimate DHCP server.

To protect the switch from DHCP starvation attacks, use the MAC limiting feature. Specify the maximum number of MAC addresses that the switch can learn on the access interfaces to which DHCP clients connect. The DHCP server or servers can then supply only the specified number of IP addresses over each of those interfaces. If a DHCP starvation attack occurs after the maximum number of IP addresses has been assigned, the attack fails.

- Related Documentation**
- [Understanding MAC Limiting and MAC Move Limiting for Port Security on page 6029](#)
  - [Configuring MAC Limiting](#)
  - [Verifying That MAC Limiting Is Working Correctly on page 6031](#)

## Understanding MAC Limiting and MAC Move Limiting for Port Security

MAC limiting protects against flooding of the Ethernet switching table (also known as the MAC forwarding table or Layer 2 forwarding table). You enable this feature on Layer 2 interfaces (ports). MAC move limiting detects MAC movement and MAC spoofing on access interfaces. You enable this feature on VLANs.

- [MAC Limiting on page 6029](#)
- [MAC Move Limiting on page 6030](#)
- [Actions for MAC Limiting on page 6030](#)
- [MAC Addresses That Exceed the MAC Limit or MAC Move Limit on page 6030](#)

### MAC Limiting

MAC limiting sets a limit on the number of MAC addresses that can be learned on a single Layer 2 access interface or on all the Layer 2 access interfaces on the switch. Junos OS provides two MAC limiting methods:

- **Maximum number of MAC addresses**—You configure the maximum number of dynamic MAC addresses allowed per interface. When the limit is exceeded, incoming packets with new MAC addresses can be ignored, dropped, or logged. You can also specify that the interface be shut down or temporarily disabled.
- **Allowed MAC addresses**—You configure specific “allowed” MAC addresses for the access interface. Any MAC address that is not in the list of configured addresses is not learned, and the switch logs an appropriate message. Allowed MAC binds MAC addresses to a VLAN so that the address does not get registered outside the VLAN. If an allowed MAC setting conflicts with a dynamic MAC setting, the allowed MAC setting takes precedence.



**NOTE:** If you do not want the system to log messages about invalid MAC addresses received by an interface that has been configured for allowed MAC addresses, disable the logging by configuring the `no-allowed-mac-log` statement.

You configure MAC limiting per interface, not per VLAN. You can specify the maximum number of dynamic MAC addresses that can be learned on a single Layer 2 access interface (including tagged-access interfaces) or on all Layer 2 access interfaces.

## MAC Move Limiting

MAC move limiting causes the switch to track the number of times a MAC address can move to a new interface (port). It can help to prevent MAC spoofing, and it can also detect and prevent loops.

If a MAC address moves more than the configured number of times within 1 second, the switch performs the configured action. You can configure MAC move limiting to apply to all VLANs or to a specific VLAN.



**CAUTION:** Mac move limiting does not work properly on a QFX5100 switch used as a Node device in a QFabric system. Do not use this feature on a QFX5100 switch in a QFabric system.

## Actions for MAC Limiting

You can choose to have one of the following actions performed when the limit of MAC addresses or the limit of MAC moves is exceeded:

- **drop**—Drop the packet and generate a system log entry. This is the default.
- **log**—Do not drop the packet but generate a system log entry.
- **none**—Take no action.
- **shutdown**—Disable the interface and generate an alarm. If you configure the switch with the **port-error-disable** statement, the disabled interface recovers automatically upon expiration of the specified timeout. If this is not configured, you can bring up the disabled interfaces by running the **clear ethernet-switching port-error** command.

See descriptions of results of these various action settings in “[Verifying That MAC Limiting Is Working Correctly](#)” on page 6031.

If you set a MAC limit to apply to all interfaces on the switch, you can override that setting for a particular interface by specifying action **none**. See *mac-limit* for more information.

## MAC Addresses That Exceed the MAC Limit or MAC Move Limit

If you have configured the **port-error-disable** statement, you can view which interfaces are temporarily disabled because the MAC limit or MAC move limit was exceeded. Use the **show ethernet-switching interfaces** command.

The log messages that indicate the MAC limit or MAC move limit has been exceeded include the offending MAC addresses.

### Related Documentation

- *Understanding Port Security Features to Protect the Access Ports on Your Device Against the Loss of Information and Productivity*
- *Configuring MAC Limiting*
- *Configuring MAC Move Limiting (CLI Procedure)*

- [Verifying That MAC Limiting Is Working Correctly on page 6031](#)
- [Verifying That MAC Move Limiting Is Working Correctly on page 6034](#)
- *Example: Configuring Basic Port Security Features*
- [no-allowed-mac-log on page 6138](#)

## Verifying That MAC Limiting Is Working Correctly

---

MAC limiting protects against flooding of the Ethernet switching table by setting a limit on the number of MAC addresses that can be learned on a single Layer 2 access interface (port).

Junos OS provides two MAC limiting methods:

- **Maximum number of MAC addresses**—You configure the maximum number of dynamic MAC addresses allowed per interface. When the limit is exceeded, incoming packets with new MAC addresses can be ignored, dropped, or logged. You can also specify that the interface be shut down or temporarily disabled.
- **Allowed MAC addresses**—You configure specific “allowed” MAC addresses for the access interface. Any MAC address that is not in the list of configured addresses is not learned, and the switch logs an appropriate message. The allowed MAC method binds MAC addresses to a VLAN so that the address is not registered outside the VLAN. If an allowed MAC setting conflicts with a dynamic MAC setting, the allowed MAC setting takes precedence.

This topic includes the following tasks:

1. [Verifying That MAC Limiting for Dynamic MAC Addresses Is Working Correctly on page 6031](#)
2. [Verifying That Allowed MAC Addresses Are Working Correctly on page 6032](#)
3. [Verifying That Interfaces Are Shut Down on page 6032](#)
4. [Customizing the Ethernet Switching Table Display to View Information for a Specific Interface on page 6033](#)

## Verifying That MAC Limiting for Dynamic MAC Addresses Is Working Correctly

**Purpose**    Verify that MAC limiting for dynamic MAC addresses is working.

**Action** Display the MAC addresses that have been learned. The following sample output shows the results of sending two packets from hosts connected to **xe-1:0/0/1** and five packets from hosts connected to **xe-1:0/0/2**, with both interfaces configured with a MAC limit of **4** and the action **drop**:

```
user@switch> show ethernet-switching table
Ethernet-switching table: 7 entries, 6 learned
```

| VLAN          | MAC address       | Type  | Age | Interfaces   |
|---------------|-------------------|-------|-----|--------------|
| employee-vlan | *                 | Flood | -   | xe-1:0/0/2.0 |
| employee-vlan | 00:05:85:3A:82:77 | Learn | 0   | xe-1:0/0/1.0 |
| employee-vlan | 00:05:85:3A:82:79 | Learn | 0   | xe-1:0/0/1.0 |
| employee-vlan | 00:05:85:3A:82:80 | Learn | 0   | xe-1:0/0/2.0 |
| employee-vlan | 00:05:85:3A:82:81 | Learn | 0   | xe-1:0/0/2.0 |
| employee-vlan | 00:05:85:3A:82:83 | Learn | 0   | xe-1:0/0/2.0 |
| employee-vlan | 00:05:85:3A:82:85 | Learn | 0   | xe-1:0/0/2.0 |

**Meaning** The output shows that the fifth packet received on the **xe-1:0/0/2** interface was dropped because it exceeded the MAC limit for that interface. The address was not learned, and thus an asterisk (\*) rather than an address appears in the MAC address column in the first line of the sample output.

## Verifying That Allowed MAC Addresses Are Working Correctly

**Purpose** Verify that allowed MAC addresses are working.

**Action** Display the MAC cache information after allowed MAC addresses have been configured on an interface. The following sample shows the MAC cache after four allowed MAC addresses had been configured on interface **xe-1:0/0/2** and a fifth MAC address appeared on the interface.

```
user@switch> show ethernet-switching table
Ethernet-switching table: 5 entries, 4 learned
```

| VLAN          | MAC address       | Type  | Age | Interfaces   |
|---------------|-------------------|-------|-----|--------------|
| employee-vlan | 00:05:85:3A:82:80 | Learn | 0   | xe-1:0/0/2.0 |
| employee-vlan | 00:05:85:3A:82:81 | Learn | 0   | xe-1:0/0/2.0 |
| employee-vlan | 00:05:85:3A:82:83 | Learn | 0   | xe-1:0/0/2.0 |
| employee-vlan | 00:05:85:3A:82:85 | Learn | 0   | xe-1:0/0/2.0 |
| employee-vlan | *                 | Flood | -   | xe-1:0/0/2.0 |

**Meaning** Because the fifth address was not allowed it was not learned, and an asterisk (\*) rather than an address appears in the MAC address column in the last line of the sample output.

## Verifying That Interfaces Are Shut Down

**Purpose** Verify that an interface is shut down when the MAC limit is exceeded.



**Action** For more information about interfaces that have been shut down because the MAC limit was exceeded, use the **show ethernet-switching interfaces** command.

```
user@switch> show ethernet-switching interfaces
Interface State VLAN members Tag Tagging Blocking

bme0.32770 down mgmt untagged unblocked
xe-0/0/0.0 down v1 untagged MAC limit exceeded
xe- 0/0/1.0 up v1 untagged unblocked
xe-0/0/2.0 up v1 untagged unblocked
me0.0 up mgmt untagged unblocked
```



**NOTE:** You can configure interfaces to recover automatically when the MAC limit has been exceeded by specifying the **port-error-disable** statement with a **disable timeout** value. The switch automatically restores the disabled interface to service when the disable timeout expires. The **port-error-disable** configuration does not apply to preexisting error conditions—it affects only error conditions that are detected after the **port-error-disable** statement has been enabled and the configuration has been committed. To clear a preexisting error condition and restore the interface to service, use the **clear ethernet-switching port-error** command.

## Customizing the Ethernet Switching Table Display to View Information for a Specific Interface

**Purpose** You can use the **show ethernet-switching table** command to view information for a specific interface.

**Action** For example, to display the MAC addresses that have been learned on the **xe-0/0/2** interface, enter:

```
user@switch> show ethernet-switching table interface xe-0/0/2.0
Ethernet-switching table: 1 unicast entries
```

| VLAN | MAC address       | Type  | Age | Interfaces  |
|------|-------------------|-------|-----|-------------|
| v1   | *                 | Flood | -   | All-members |
| v1   | 00:00:06:00:00:00 | Learn | 0   | xe-0/0/2.0  |

**Meaning** The MAC limit value for the **xe-0/0/2** interface had been set to 1, and the output shows that only one MAC address was learned and added to the MAC cache.

**Related Documentation**

- *Configuring MAC Limiting*
- *Monitoring Port Security*

- *Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure)*
- *Example: Configuring Allowed MAC Addresses to Protect the Switch from DHCP Snooping Database Alteration Attacks*
- *Example: Configuring MAC Limiting to Protect the Switch from DHCP Starvation Attacks*

---

## Verifying That MAC Move Limiting Is Working Correctly

---

**Purpose** Verify that MAC move limiting is working on the switch.

**Action** Display the MAC addresses in the Ethernet switching table when MAC move limiting has been configured for a VLAN. The following sample shows the results after two of the hosts on **ge-0/0/2** sent packets after the MAC addresses for those hosts had moved to other interfaces more than five times in 1 second. The VLAN, **employee-vlan**, was set to a MAC move limit of **5** with the action **drop**:

```
user@switch> show ethernet-switching table
```

```
Ethernet-switching table: 7 entries, 4 learned
```

| VLAN          | MAC address       | Type  | Age | Interfaces |
|---------------|-------------------|-------|-----|------------|
| employee-vlan | 00:05:85:3A:82:77 | Learn | 0   | ge-0/0/1.0 |
| employee-vlan | 00:05:85:3A:82:79 | Learn | 0   | ge-0/0/1.0 |
| employee-vlan | 00:05:85:3A:82:80 | Learn | 0   | ge-0/0/2.0 |
| employee-vlan | 00:05:85:3A:82:81 | Learn | 0   | ge-0/0/2.0 |
| employee-vlan | *                 | Flood | -   | ge-0/0/2.0 |
| employee-vlan | *                 | Flood | -   | ge-0/0/2.0 |

**Meaning** The last two lines of the sample output show that MAC addresses for two hosts on **ge-0/0/2** were not learned, because the hosts had been moved back and forth from the original interfaces more than five times in 1 second.

- Related Documentation**
- *Configuring MAC Move Limiting (CLI Procedure)*
  - *Configuring MAC Move Limiting (J-Web Procedure)*
  - *Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure)*
  - *Example: Configuring Basic Port Security Features*
  - *Monitoring Port Security*

---

## Verifying That the Port Error Disable Setting Is Working Correctly

---

**Purpose** Verify that the port error disable setting is working as expected for MAC limited and storm control interfaces.

**Action** Display information about interfaces:

```
user@switch> show ethernet-switching interfaces
```

| Interface    | State | VLAN members | Blocking                |
|--------------|-------|--------------|-------------------------|
| xe-2:0/0/0.0 | up    | T1122        | unblocked               |
| xe-2:0/0/1.0 | down  | default      | MAC limit exceeded      |
| xe-2:0/0/2.0 | down  | default      | Storm control in effect |
| xe-2:0/0/3.0 | down  | default      | unblocked               |
| xe-2:0/0/4.0 | down  | default      | unblocked               |
| xe-2:0/0/5.0 | down  | default      | unblocked               |
| xe-2:0/0/6.0 | down  | default      | unblocked               |

**Meaning** For interfaces disabled by port security features, the sample output from the **show ethernet-switching interfaces** command specifies the reasons that the interfaces are disabled:

- **MAC limit exceeded**—The interface is temporarily disabled because of a *mac-limit* error. The disabled interface is automatically restored to service when the disable-timeout expires.
- **MAC move limit exceeded**—The interface is temporarily disabled because of a *mac-move-limit* error. The disabled interface is automatically restored to service when the disable-timeout expires.
- **Storm control in effect**—The interface is temporarily disabled because of a storm control error. The disabled interface is automatically restored to service when the disable-timeout expires.

**Related Documentation**

- [Understanding MAC Limiting and MAC Move Limiting for Port Security on page 6029](#)
- [port-error-disable on page 6140](#)

## Configuring Persistent MAC Learning (CLI Procedure)



**NOTE:** This topic uses Junos OS with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring Persistent MAC Learning (CLI Procedure)*. For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on [page 41](#).

Persistent MAC address learning is disabled by default. You can enable it to:

- Help prevent traffic losses for trusted workstations and servers because, if persistent MAC address learning is enabled on an interface, the interface does not have to relearn the addresses from ingress traffic after a restart.
- Protect the switch against security attacks—Use persistent MAC learning in combination with MAC limiting to protect against attacks while still obviating the need to statically configure MAC addresses. When the initial learning of MAC addresses up to the number specified by the MAC limit is done, new addresses are not allowed even after a restart.

The port is secured because after the limit has been reached, additional devices cannot connect to the interface.

To configure persistent MAC learning on an interface and limit the number of allowed MAC addresses:

1. Enable persistent MAC learning on an interface:

```
[edit switch-options]
user@switch# set interface interface-name persistent-learning
```

2. Configure the MAC limit on an interface, and specify the action that the switch takes after the specified limit is exceeded:

```
[edit switch-options]
user@switch# set interface interface-name interface-mac-limit limit packet-action action
```

After you set a new MAC limit for the interface, the system clears existing entries in the MAC address forwarding table associated with the interface.

Values for *action* are:

**drop**—Drop packets with new source MAC addresses, and do not learn the new source MAC addresses.

**drop-and-log**—(EX Series switches only) Drop packets with new source MAC addresses, and generate an alarm, an SNMP trap, or a system log entry.

**log**—(EX Series switches only) Hold packets with new source MAC addresses, and generate an alarm, an SNMP trap, or a system log entry.

**none**—(EX Series switches only) Forward packets with new source MAC addresses, and learn the new source MAC address.

**shutdown**—(EX Series switches only) Disable the specified interface, and generate an alarm, an SNMP trap, or a system log entry.



**TIP:** If you move a device within your network that has a persistent MAC address entry on the switch, use the [clear ethernet-switching table persistent-mac](#) command to clear the persistent MAC address entry from the interface. If you move the device and do not clear the persistent MAC address from the original port it was learned on, then the new port will not learn the MAC address of the device and the device will not be able to connect.

If the original port is down when you move the device, then the new port will learn the MAC address and the device can connect. However, if you do not clear the persistent MAC address on the original port, then when the port restarts, the system reinstalls the persistent MAC address in the forwarding table for that port. If this occurs, the persistent MAC address is removed from the new port and the device loses connectivity.

---

**Related  
Documentation**

- [Configuring Persistent MAC Learning \(CLI Procedure\)](#)

- *Configuring MAC Move Limiting (CLI Procedure)*
- *Understanding Persistent MAC Learning (Sticky MAC)*

## Understanding Trusted and Untrusted Ports

---

By default, all access ports are untrusted and all trunk ports are trusted in regard to DHCP. Trusted ports allow DHCP servers to provide IP addresses and other information to requesting devices. Untrusted ports drop traffic from DHCP servers to prevent unauthorized servers from providing any configuration information to clients.

If you attach a DHCP server to an access port, you must configure the port as trusted. Before you do so, ensure that the server is physically secure—that is, that access to the server is monitored and controlled.

- Related Documentation**
- *Understanding DHCP Snooping for Monitoring DHCP Messages Received from Untrusted Devices*
  - *Example: Configuring Basic Port Security Features*

## Understanding Trusted DHCP Servers for Port Security

---

Any interface on the switching device that connects to a DHCP server can be configured as a trusted port. Configuring a DHCP server on a trusted port protects against rogue DHCP servers sending leases.

Ensure that the DHCP server interface is physically secure—that is, that access to the server is monitored and controlled at the site—before you configure the port as trusted.

- Related Documentation**
- *Example: Configuring a DHCP Server Interface as Untrusted to Protect the Switch from Rogue DHCP Server Attacks*
  - *Enabling a Trusted DHCP Server (CLI Procedure)*

## Verifying That a Trusted DHCP Server Is Working Correctly

---

- Purpose** Verify that a DHCP trusted server is working on the switch. See what happens when the DHCP server is trusted and then untrusted.

**Action** Send some DHCP requests from network devices (here they are DHCP clients) connected to the switch.

Display the DHCP snooping information when the interface on which the DHCP server connects to the switch is trusted. The following output results when requests are sent from the MAC addresses and the server has provided the IP addresses and leases:

```
user@switch> show dhcp snooping binding
```

DHCP Snooping Information:

| MAC Address       | IP Address | Lease | Type    | VLAN          | Interface  |
|-------------------|------------|-------|---------|---------------|------------|
| -----             | -----      | ----- | ----    | ----          | -----      |
| 00:05:85:3A:82:77 | 192.0.2.17 | 600   | dynamic | employee-vlan | ge-0/0/1.0 |
| 00:05:85:3A:82:79 | 192.0.2.18 | 653   | dynamic | employee-vlan | ge-0/0/1.0 |
| 00:05:85:3A:82:80 | 192.0.2.19 | 720   | dynamic | employee-vlan | ge-0/0/2.0 |
| 00:05:85:3A:82:81 | 192.0.2.20 | 932   | dynamic | employee-vlan | ge-0/0/2.0 |
| 00:05:85:3A:82:83 | 192.0.2.21 | 1230  | dynamic | employee-vlan | ge-0/0/2.0 |
| 00:05:85:27:32:88 | 192.0.2.22 | 3200  | dynamic | employee-vlan | ge-0/0/2.0 |

**Meaning** When the interface on which the DHCP server connects to the switch has been set to trusted, the output (see preceding sample) shows, for each MAC address, the assigned IP address and lease time—that is, the time, in seconds, remaining before the lease expires.

If the DHCP server had been configured as untrusted, no entries would be added to the DHCP snooping database and nothing would be shown in the output of the **show dhcp snooping binding** command.

- Related Documentation**
- *Enabling a Trusted DHCP Server (CLI Procedure)*
  - *Enabling a Trusted Port for DHCP*
  - *Example: Configuring Basic Port Security Features*
  - *Example: Configuring a DHCP Server Interface as Untrusted to Protect the Switch from Rogue DHCP Server Attacks*
  - *Monitoring Port Security*
  - *Troubleshooting Port Security*

---

## Understanding DHCP Option 82 for Port Security

---

You can use DHCP option 82, also known as the DHCP relay agent information option, to help protect the switch against attacks such as spoofing (forging) of IP addresses and MAC addresses, and DHCP IP address starvation. Hosts on untrusted access interfaces on Ethernet LAN switches send requests for IP addresses in order to access the Internet. The switch forwards or relays these requests to DHCP servers, and the servers send offers for IP address leases in response. Attackers can use these messages to perpetrate address spoofing and starvation.

Option 82 provides information about the network location of a DHCP client, and the DHCP server uses this information to implement IP addresses or other parameters for the client. The Juniper Networks Junos operating system (Junos OS) implementation of DHCP option 82 supports RFC 3046, *DHCP Relay Agent Information Option*, at <http://tools.ietf.org/html/rfc3046>.

This topic covers:

- [DHCP Option 82 Processing on page 6039](#)
- [Suboption Components of Option 82 on page 6040](#)
- [Configurations That Support Option 82 on page 6040](#)

### DHCP Option 82 Processing

If DHCP option 82 is enabled on the switch, then when a DHCP client that is connected to the switch on an untrusted interface sends a DHCP request, the switch inserts information about the client's network location into the packet header of that request. The switch then sends the request to the DHCP server. The DHCP server reads the option 82 information in the packet header and uses it to implement the IP address or another parameter for the client. See "[Suboption Components of Option 82](#)" on [page 6040](#) for details about option 82 information.

You can enable DHCP option 82 on a single VLAN or on all VLANs on the switch. You can also configure it on Layer 3 interfaces (in routed VLAN interfaces, or RVIs) when the switch is functioning as a relay agent.

When option 82 is enabled on the switch, then this sequence of events occurs when a DHCP client sends a DHCP request:

1. The switch receives the request and inserts the option 82 information in the packet header.
2. The switch forwards or relays the request to the DHCP server.
3. The server uses the DHCP option 82 information to formulate its reply and sends a response back to the switch. It does not alter the option 82 information.
4. The switch strips the option 82 information from the response packet.
5. The switch forwards the response packet to the client.



**NOTE:** To use the DHCP option 82 feature, you must ensure that the DHCP server is configured to accept option 82. If it is not configured to accept option 82, then when it receives requests containing option 82 information, it does not use the information in setting parameters and it does not echo the information in its response message.

---

## Suboption Components of Option 82

When configuring DHCP option 82, you can use the following suboptions:

- **circuit ID**—Identifies the circuit (interface and/or VLAN) on the switch on which the request was received. The circuit ID contains the interface name and/or VLAN name, with the two elements separated by a colon—for example, **xe-0/0/10:vlan1**. If the request packet is received on a Layer 3 interface, the circuit ID is just the interface name—for example, **xe-0/0/10**.

Use the **prefix** option to add an optional prefix to the circuit ID. If you enable the **prefix** option, the hostname for the switch is used as the prefix; for example, **switch1:xe-0/0/10:vlan1**.

You can also specify that the interface description be used rather than the interface name and that the VLAN ID be used rather than the VLAN name.

- **remote ID**—Identifies the host. By default, the remote ID is the MAC address of the switch. You can specify that the remote ID be the hostname of the switch, the interface description, or a character string of your choice. You can also add an optional prefix to the remote ID.
- **vendor ID**—Identifies the vendor of the host. If you specify the **vendor-id** option but do not enter a value, the default value **Juniper** is used. To specify a value, you type a character string.


## Configurations That Support Option 82

You can use option 82 with the following configurations:

- The DHCP client and the DHCP server are on the same VLAN. In this case the switch forwards the requests from the clients on untrusted access interfaces to the server on a trusted interface. For this configuration, you set DHCP option 82 at the **[edit ethernet-switching-options secure-access-port vlan]** hierarchy level.
- The DHCP client or the DHCP server is connected to the switch through a Layer 3 interface and the switch is configured to relay DHCP requests. [Figure 199](#) illustrates a scenario for the switch-as-relay-agent; in this instance, the switch relays requests through a router to the server.



Figure 199: Switch Relays DHCP Requests to Server



For the configuration shown in [Figure 199](#), you set DHCP option 82 at the **[edit forwarding-options helpers bootp]** hierarchy level.

#### Related Documentation

- [Overview of Access Port Protection](#)
- [DHCP and BOOTP Relay Overview on page 6255](#)
- [dhcp-option82](#)
- [Example: Setting Up DHCP Option 82 with a Switch with No Relay Agent Between Clients and a DHCP Server](#)
- [Example: Setting Up DHCP Option 82 with a Switch as a Relay Agent Between Clients and a DHCP Server](#)
- [Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server \(CLI Procedure\)](#)
- [Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server \(CLI Procedure\)](#)

## Understanding Static ARP Entries

You can create explicit mappings between IP addresses and MAC addresses, which are called static ARP table entries. Unlike dynamically learned ARP entries, static entries do not age out. You might want to create static ARP entries in a troubleshooting situation or if your device is unable to learn a MAC address dynamically for any reason.

#### Related Documentation

- [Configuring Static ARP Entries on page 2133](#)

- *arp*

## PART 81

# Configuring Device Security

- [Device Security on page 6045](#)



# Device Security

- [Understanding Storm Control on page 6045](#)
- [Example: Configuring Storm Control to Prevent Network Outages on page 6046](#)
- [Verifying That the Port Error Disable Setting Is Working Correctly on page 6049](#)
- [Understanding Unicast RPF on page 6050](#)
- [Configuring Unicast RPF \(CLI Procedure\) on page 6054](#)
- [Disabling Unicast RPF \(CLI Procedure\) on page 6055](#)
- [Verifying Unicast RPF Status on page 6056](#)
- [Understanding Unknown Unicast Forwarding on page 6059](#)
- [Configuring Unknown Unicast Forwarding \(CLI Procedure\) on page 6059](#)

## Understanding Storm Control

---

A traffic storm occurs when broadcast packets prompt receiving devices to broadcast packets in response. This prompts further responses, creating a snowball effect. The switch is flooded with packets, which creates unnecessary traffic that leads to poor performance or even a complete loss of service by some clients. Storm control causes a device to monitor traffic levels and take a specified action when a specified traffic level—called the *storm control level*—is exceeded, thus preventing packets from proliferating and degrading service. You can configure devices to drop broadcast and unknown unicast packets, shut down interfaces, or temporarily disable interfaces when the storm control level is exceeded.

Storm control is enabled by default on ELS platforms and disabled by default on non-ELS platforms. If storm control is enabled, the default level is 80 percent of the available bandwidth for ingress traffic. You can change the storm control level by configuring it as a specific bandwidth value. (The **level** configuration statement, which allows you to configure the storm control level as a percentage of the combined broadcast and unknown unicast streams, is deprecated and might be removed from future releases. We recommend that you phase out its use and replace it with the **bandwidth** statement.)



**NOTE:** Storm control is not enabled by default on MX platforms.



**NOTE:** When you configure storm control bandwidth, the value you configure is rounded off internally to the closest multiple of 64 Kbps, and the rounded-off value represents the bandwidth that is actually enforced. For example, if you configure a bandwidth limit of 150 Kbps, storm control enforces a bandwidth limit of 128 Kbps.



**NOTE:** On an FCoE-FC gateway, storm control must be disabled on all Ethernet interfaces that belong to an FCoE VLAN to prevent FCoE traffic from being dropped. Configuring storm control on an Ethernet interface that is included in an FCoE-FC gateway may have undesirable effects, including FCoE packet loss. After disabling storm control on all interfaces, enable storm control on any interfaces that are not part of an FCoE-FC gateway on which you want to use storm control. However, on an FCoE transit switch, you can enable storm control on interfaces that carry FCoE traffic.



**CAUTION:** The Junos OS allows you to configure a storm control value that exceeds the bandwidth of the interface. If you configure an interface this way, storm control does not drop broadcast or unknown unicast packets even if they consume all the available bandwidth.

To recognize a storm, you must be able to identify when traffic has reached an abnormal level. Suspect a storm when operations begin timing out and network response times slow down. Users might be unable to access expected services. Monitor the percentage of broadcast and unknown unicast traffic in the network when it is operating normally. This data can then be used as a benchmark to determine when traffic levels are too high. You can then configure storm control to set the level at which you want to drop broadcast and unknown unicast traffic.



**NOTE:** On a QFX10002 switch, if storm control is configured on a VLAN port associated with an IRB interface, unregistered multicast traffic is classified as registered multicast traffic if IGMP snooping is enabled. If IGMP snooping is disabled, the traffic is classified as unknown unicast traffic.

**Related  
Documentation**

- [action-shutdown on page 6120](#)
- [port-error-disable on page 6140](#)
- [storm-control on page 6128](#)

---

## Example: Configuring Storm Control to Prevent Network Outages

Storm control enables you to prevent network outages caused by broadcast storms on the LAN. You can configure storm control on to rate-limit broadcast traffic, multicast

traffic, and unknown unicast traffic at a specified level and to have packets dropped when the specified traffic level is exceeded, thereby preventing packets from proliferating and degrading the LAN.



**NOTE:** This example uses a Junos OS release that supports the Enhanced Layer 2 Software (ELS) configuration style.

- [Requirements on page 6047](#)
- [Overview and Topology on page 6047](#)
- [Configuration on page 6048](#)

## Requirements

This example uses the following hardware and software components:

- One QFX Series switch running Junos OS with ELS
- Junos OS Release 13.2 or later

## Overview and Topology

A storm is generated when messages are broadcast on a network and each message prompts a receiving node to respond by broadcasting its own messages on the network. This, in turn, prompts further responses, creating a snowball effect and resulting in a broadcast storm that can cause network outages.

You can use storm control to prevent broadcast storms by specifying the amount, also known as the *storm control level*, of broadcast traffic, multicast traffic, and unknown unicast traffic to be allowed on an interface. You specify the storm control level as the traffic rate in kilobits per second (Kbps) of the combined applicable traffic streams or as the percentage of available bandwidth used by the combined applicable traffic streams. On ELS systems, storm control is enabled by default on all interfaces at a level of 80 percent of the available bandwidth.

Storm control monitors the level of applicable incoming traffic and compares it with the level that you specify. If the combined level of the applicable traffic exceeds the specified level, the switch drops packets for the controlled traffic types. As an alternative to having the switch drop packets, you can configure storm control to shut down interfaces or temporarily disable interfaces (see the **action-shutdown** statement or the **recovery-timeout** statement) when the storm control level is exceeded.



**NOTE:** If you configure storm control on an aggregated Ethernet interface, the storm-control level is applied to each member interface individually. For example, if the aggregated interface has two members and you configure a storm-control level of 20 kbps, Junos will not detect a storm if one or both of the member interfaces receives traffic at 15 kbps because in neither of these cases does an individual member receive traffic at a rate greater than the configured storm-control level. In this example, Junos detects a storm only if at least one member interface receives traffic at greater than 20 Kbps.

The topology used in this example consists of one switch connected to various network devices. This example shows how to configure the storm control level on interface xe-0/0/0 by setting the level to a traffic rate of 15,000 Kbps, based on the traffic rate of the combined applicable traffic streams. If the combined traffic exceeds this level, the switch drops packets for the controlled traffic types to prevent a network outage.

## Configuration

### CLI Quick Configuration

To quickly configure storm control based on the traffic rate in kilobits per second of the combined traffic streams, copy the following command and paste it into the switch terminal window:

```
[edit]
set forwarding-options storm-control-profiles sc-profile all bandwidth-level 15000
set interfaces xe-0/0/0 unit 0 family ethernet-switching storm-control sc-profile
```

### Step-by-Step Procedure

To configure storm control:

1. Configure a storm control profile, **sc-profile**, and specify the traffic rate in kilobits per second of the combined traffic streams:

```
[edit]
user@switch> set forwarding-options storm-control-profiles sc-profile all bandwidth-level 15000
```

2. Bind the storm control profile, **sc**, to a logical interface:

```
[edit]
user@switch> set interfaces xe-0/0/0 unit 0 family ethernet-switching storm-control sc-profile
```

### Results

Display the results of the configuration:

```
[edit forwarding-options]
user@switch> show storm-control-profiles sc-profile
all {
 bandwidth 15000;
}

[edit]
user@switch> show interfaces xe-0/0/0
unit 0 {
 family ethernet-switching {
 vlan {
 members default;
 }
 }
 storm-control sc-profile;
```



```
}
}
```

- Related Documentation**
- [Understanding Storm Control on page 6045](#)
  - [Configuring Autorecovery for MAC Limited or Storm Control Interfaces \(CLI Procedure\)](#)

## Verifying That the Port Error Disable Setting Is Working Correctly

**Purpose** Verify that the port error disable setting is working as expected for MAC limited and storm control interfaces.

**Action** Display information about interfaces:

```
user@switch> show ethernet-switching interfaces
```

| Interface    | State | VLAN members | Blocking                |
|--------------|-------|--------------|-------------------------|
| xe-2:0/0/0.0 | up    | T1122        | unblocked               |
| xe-2:0/0/1.0 | down  | default      | MAC limit exceeded      |
| xe-2:0/0/2.0 | down  | default      | Storm control in effect |
| xe-2:0/0/3.0 | down  | default      | unblocked               |
| xe-2:0/0/4.0 | down  | default      | unblocked               |
| xe-2:0/0/5.0 | down  | default      | unblocked               |
| xe-2:0/0/6.0 | down  | default      | unblocked               |

**Meaning** For interfaces disabled by port security features, the sample output from the **show ethernet-switching interfaces** command specifies the reasons that the interfaces are disabled:

- **MAC limit exceeded**—The interface is temporarily disabled because of a *mac-limit* error. The disabled interface is automatically restored to service when the disable-timeout expires.
- **MAC move limit exceeded**—The interface is temporarily disabled because of a *mac-move-limit* error. The disabled interface is automatically restored to service when the disable-timeout expires.
- **Storm control in effect**—The interface is temporarily disabled because of a storm control error. The disabled interface is automatically restored to service when the disable-timeout expires.

- Related Documentation**
- [Understanding MAC Limiting and MAC Move Limiting for Port Security on page 6029](#)
  - [port-error-disable on page 6140](#)

## Understanding Unicast RPF

---

Unicast reverse-path forwarding (RPF) helps protect the switch against denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks by verifying the unicast source address of each packet that arrives on an ingress interface where unicast RPF is enabled. It also helps ensure that traffic arriving on ingress interfaces comes from a network source that the receiving interface can reach.

When you enable unicast RPF, by default the switch forwards a packet only if the receiving interface is the best return path to the packet's unicast source address. This is known as strict mode unicast RPF. You can also enable loose mode, which means that the system checks to see if the packet has a source address with a corresponding prefix in the routing table but does not check whether the receiving interface is the best return path to the packet's unicast source address.



.....

**NOTE:** On Juniper Networks EX3200, EX4200, and EX4300 Ethernet Switches, the switch applies unicast RPF *globally* to all interfaces when unicast RPF is configured on any interface. For additional information, see [“Limitations of the Unicast RPF Implementation on EX3200, EX4200, and EX4300 Switches” on page 6053](#).

.....

This topic covers:

- [Unicast RPF for Switches Overview on page 6050](#)
- [Unicast RPF Implementation on page 6051](#)
- [When to Enable Unicast RPF on page 6051](#)
- [When Not to Enable Unicast RPF on page 6052](#)
- [Limitations of the Unicast RPF Implementation on EX3200, EX4200, and EX4300 Switches on page 6053](#)

### Unicast RPF for Switches Overview

Unicast RPF functions as an ingress filter that reduces the forwarding of IP packets that might be spoofing an address. By default, unicast RPF is disabled on the switch interfaces.

The type of unicast RPF provided on the switches—that is, strict mode unicast RPF is especially useful on untrusted interfaces. An untrusted interface is an interface where untrusted users or processes can place packets on the network segment.

The switch supports only the active paths method of determining the best return path back to a unicast source address. The active paths method looks up the best reverse path entry in the forwarding table. It does not consider alternate routes specified using routing-protocol-specific methods when determining the best return path.

If the forwarding table lists the receiving interface as the interface to use to forward the packet back to its unicast source, it is the best return path interface.

Use strict mode unicast RPF only on symmetrically routed interfaces. (For information about symmetrically routed interfaces, see [“When to Enable Unicast RPF” on page 6051](#).)

For more information about strict unicast RPF, see RFC 3704, *Ingress Filtering for Multihomed Networks* at <http://www.ietf.org/rfc/rfc3704.txt>.

## Unicast RPF Implementation

This section includes:

- [Unicast RPF Packet Filtering on page 6051](#)
- [Bootstrap Protocol \(BOOTP\) and DHCP Requests on page 6051](#)
- [Default Route Handling on page 6051](#)

### Unicast RPF Packet Filtering

When you enable unicast RPF on the switch, the switch handles traffic in the following manner:

- If the switch receives a packet on the interface that is the best return path to the unicast source address of that packet, the switch forwards the packet.
- If the best return path from the switch to the packet's unicast source address is not the receiving interface, the switch discards the packet.
- If the switch receives a packet that has a source IP address that does not have a routing entry in the forwarding table, the switch discards the packet.

### Bootstrap Protocol (BOOTP) and DHCP Requests

Bootstrap protocol (BOOTP) and DHCP request packets are sent with a broadcast MAC address and therefore the switch does not perform unicast RPF checks on them. The switch forwards all BOOTP packets and DHCP request packets without performing unicast RPF checks.

### Default Route Handling

If the best return path to the source is the default route (0.0.0.0) and the default route points to **reject**, the switch discards the packets. If the default route points to a valid network interface, the switch performs a normal unicast RPF check on the packets.

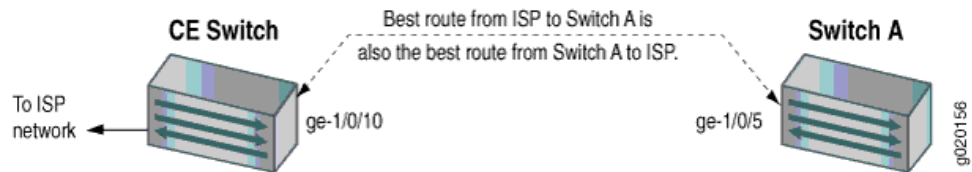
## When to Enable Unicast RPF

Enable unicast RPF when you want to ensure that traffic arriving on a network interface comes from a source that resides on a network that that interface can reach. You can enable unicast RPF on untrusted interfaces to filter spoofed packets. For example, a common application for unicast RPF is to help defend an enterprise network from DoS/DDoS attacks coming from the Internet.

Enable unicast RPF only on symmetrically routed interfaces. A symmetrically routed interface uses the same route in both directions between the source and the destination, as shown in [Figure 200](#). Symmetrical routing means that if an interface receives a packet,

the switch uses the same interface to send a reply to the packet source (the receiving interface matches the forwarding-table entry for the best return path to the source).

**Figure 200: Symmetrically Routed Interfaces**



Enabling unicast RPF on asymmetrically routed interfaces (where different interfaces receive a packet and reply to its source) results in packets from legitimate sources being filtered (discarded) because the best return path is not the same interface that received the packet.

The following switch interfaces are most likely to be symmetrically routed and thus are candidates for unicast RPF enabling:

- The service provider edge to a customer
- The customer edge to a service provider
- A single access point out of the network (usually on the network perimeter)
- A terminal network that has only one link



**NOTE:** Because unicast RPF is enabled globally on EX3200, EX4200, and EX4300 switches, ensure that *all* interfaces are symmetrically routed before you enable unicast RPF on these switches. Enabling unicast RPF on asymmetrically routed interfaces results in packets from legitimate sources being filtered.



**TIP:** Enabling unicast RPF as close as possible to the traffic source stops spoofed traffic before it can proliferate or reach interfaces that do not have unicast RPF enabled.


## When Not to Enable Unicast RPF

Typically, you will not enable unicast RPF if:

- Switch interfaces are multihomed.
- Switch interfaces are trusted interfaces.
- BGP is carrying prefixes and some of those prefixes are not advertised or are not accepted by the ISP under its policy. (The effect in this case is the same as filtering an interface by using an incomplete access list.)
- Switch interfaces face the network core. Core-facing interfaces are usually asymmetrically routed.

An asymmetrically routed interface uses different paths to send and receive packets between the source and the destination, as shown in [Figure 201](#). This means that if an interface receives a packet, that interface does not match the forwarding table entry as the best return path back to the source. If the receiving interface is not the best return path to the source of a packet, unicast RPF causes the switch to discard the packet even though it comes from a valid source.

Figure 201: Asymmetrically Routed Interfaces



**NOTE:** Do not enable unicast RPF on EX3200, EX4200, and EX4300 switches if any switch interfaces are asymmetrically routed, because unicast RPF is enabled globally on all interfaces of these switches. All switch interfaces must be symmetrically routed for you to enable unicast RPF without the risk of the switch discarding traffic that you want to forward.

### Limitations of the Unicast RPF Implementation on EX3200, EX4200, and EX4300 Switches

On EX3200, EX4200, and EX4300 switches, the switch implements unicast RPF on a global basis. You cannot enable unicast RPF on a per-interface basis. Unicast RPF is globally disabled by default.

- When you enable unicast RPF on any interface, it is automatically enabled on all switch interfaces, including link aggregation groups (LAGs), integrated routing and bridging (IRB) interfaces, and routed VLAN interfaces (RVIs).
- When you disable unicast RPF on the interface (or interfaces) on which you enabled unicast RPF, it is automatically disabled on all switch interfaces.



**NOTE:** You must explicitly disable unicast RPF on every interface on which it was explicitly enabled or unicast RPF remains enabled on all switch interfaces.

QFX switches, OCX switches, and EX3200 and EX4200 switches do not perform unicast RPF filtering on equal-cost multipath (ECMP) traffic. The unicast RPF check examines only one best return path to the packet source, but ECMP traffic employs an address block consisting of multiple paths. Using unicast RPF to filter ECMP traffic on these switches can result in the switch discarding packets that you want to forward because the unicast RPF filter does not examine the entire ECMP address block.

- Related Documentation**
- [Example: Configuring Unicast RPF on an EX Series Switch](#)
  - [Configuring Unicast RPF \(CLI Procedure\) on page 6054](#)
  - [Disabling Unicast RPF \(CLI Procedure\) on page 6055](#)

---

## Configuring Unicast RPF (CLI Procedure)

---

Unicast reverse-path forwarding (RPF) can help protect your LAN from denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks on untrusted interfaces. When you enable unicast RPF, by default the switch forwards a packet only if the receiving interface is the best return path to the packet's unicast source address. This is known as strict mode unicast RPF. You can also enable loose mode, which means that the system checks to see if the packet has a source address with a corresponding prefix in the routing table but does not check whether the receiving interface is the best return path to the packet's unicast source address.



**NOTE:** On EX3200, EX4200, and EX4300 switches, you can enable unicast RPF only globally—that is, on all switch interfaces. You cannot enable unicast RPF on a per-interface basis.

Before you begin:

- On an EX8200, EX6200, QFX Series switch, or OCX Series switch, ensure that the selected switch interface is symmetrically routed before you enable unicast RPF. A symmetrically routed interface is an interface that uses the same route in both directions between the source and the destination. Do not enable unicast RPF on asymmetrically routed interfaces. An asymmetrically routed interface uses different paths to send and receive packets between the source and the destination.
- On an EX3200, EX4200, or EX4300 switch, ensure that *all* switch interfaces are symmetrically routed before you enable unicast RPF on an interface. When you enable unicast RPF on any interface, it is enabled globally on all switch interfaces. Do not enable unicast RPF on asymmetrically routed interfaces. An asymmetrically routed interface uses different paths to send and receive packets between the source and the destination.

To enable unicast RPF, configure it explicitly on a selected customer-edge interface:

[edit interfaces]

user@switch# **set interface-name** unit 0 family inet **rpf-check**

To enable unicast RPF loose mode, enter:

[edit interfaces]

user@switch# **set interface-name** unit 0 family inet **rpf-check mode loose**



**BEST PRACTICE:** On EX3200, EX4200, and EX4300 switches, unicast RPF is enabled globally on *all* switch interfaces, regardless of whether you configure it explicitly on only one interface or only on some interfaces.

On EX3200, EX4200, and EX4300 switches, we recommend that you enable unicast RPF explicitly on either all interfaces or only one interface. To avoid possible confusion, do not enable it on only some interfaces:

- Enabling unicast RPF explicitly on only one interface makes it easier if you choose to disable it in the future because you must explicitly disable unicast RPF on every interface on which you explicitly enabled it. If you explicitly enable unicast RPF on two interfaces and you disable it on only one interface, unicast RPF is still implicitly enabled globally on the switch. The drawback of this approach is that the switch displays the flag that indicates that unicast RPF is enabled only on interfaces on which unicast RPF is explicitly enabled, so even though unicast RPF is enabled on all interfaces, this status is not displayed.
- Enabling unicast RPF explicitly on all interfaces makes it easier to know whether unicast RPF is enabled on the switch because every interface shows the correct status. (Only interfaces on which you explicitly enable unicast RPF display the flag that indicates that unicast RPF is enabled.) The drawback of this approach is that if you want to disable unicast RPF, you must explicitly disable it on every interface. If unicast RPF is enabled on any interface, it is implicitly enabled on all interfaces.

#### Related Documentation

- *Example: Configuring Unicast RPF on an EX Series Switch*
- [Verifying Unicast RPF Status on page 6056](#)
- [Disabling Unicast RPF \(CLI Procedure\) on page 6055](#)
- *Troubleshooting Unicast RPF*
- [Understanding Unicast RPF on page 6050](#)

## Disabling Unicast RPF (CLI Procedure)

Unicast reverse-path forwarding (RPF) can help protect your LAN from denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks on untrusted interfaces. Unicast

RPF filters traffic with source addresses that do not use the incoming interface as the best return path back to the source. If the network configuration changes so that an interface that has unicast RPF enabled becomes a trusted interface or becomes asymmetrically routed (the interface that receives a packet is not the best return path to the packet's source), disable unicast RPF.

To disable unicast RPF on an EX3200, EX4200, or EX4300 switch, you must delete it from every interface on which you explicitly configured it. If you do not disable unicast RPF on every interface on which you explicitly enabled it, it remains implicitly enabled on all interfaces. If you attempt to delete unicast RPF from an interface on which it was not explicitly enabled, the **warning: statement not found** message appears. If you do not disable unicast RPF on every interface on which you explicitly enabled it, unicast RPF remains implicitly enabled on all interfaces of the EX3200, EX4200, or EX4300 switch.

On EX8200, EX6200, QFX Series switches, and OCX Series switches, the switch does not apply unicast RPF to an interface unless you explicitly enable that interface for unicast RPF.

To disable unicast RPF, delete its configuration from the interface:

[edit interfaces]

user@switch# **delete** ge-1/0/10 unit 0 family inet **rpf-check**



**NOTE:** On EX3200, EX4200, and EX4300 switches, if you do not disable unicast RPF on every interface on which you explicitly enabled it, unicast RPF remains implicitly enabled on all interfaces.

**Related  
Documentation**

- [Example: Configuring Unicast RPF on an EX Series Switch](#)
- [Verifying Unicast RPF Status on page 6056](#)
- [Configuring Unicast RPF \(CLI Procedure\) on page 6054](#)
- [Understanding Unicast RPF on page 6050](#)

---

## Verifying Unicast RPF Status

---

**Purpose** Verify that unicast reverse-path forwarding (RPF) is enabled and is working on the interface.

**Action** Use one of the **show interfaces *interface-name*** commands with either the **extensive** or **detail** options to verify that unicast RPF is enabled and working on the switch. The following example displays output from the **show interfaces ge- extensive** command.

```
user@switch> show interfaces ge-1/0/10 extensive
Physical interface: ge-1/0/10, Enabled, Physical link is Down
 Interface index: 139, SNMP ifIndex: 58, Generation: 140
 Link-level type: Ethernet, MTU: 1514, Speed: Auto, MAC-REWRITE Error: None,
 Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled,
 Auto-negotiation: Enabled, Remote fault: Online
 Device flags : Present Running
```



```

Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
Link flags : None
CoS queues : 8 supported, 8 maximum usable queues
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:19:e2:50:95:ab, Hardware address: 00:19:e2:50:95:ab
Last flapped : Never
Statistics last cleared: Never
Traffic statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets : 0 0 pps
 Output packets: 0 0 pps
IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets : 0
 Output packets: 0
Input errors:
 Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
 L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
 FIFO errors: 0, Resource errors: 0
Output errors:
 Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

 FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters: Queued packets Transmitted packets Dropped packets

 0 best-effort 0 0 0
 1 assured-forw 0 0 0
 5 expedited-fo 0 0 0
 7 network-cont 0 0 0

Active alarms : LINK
Active defects : LINK
MAC statistics:
 Total octets Receive Transmit
 Total packets 0 0
 Unicast packets 0 0
 Broadcast packets 0 0
 Multicast packets 0 0
 CRC/Align errors 0 0
 FIFO errors 0 0
 MAC control frames 0 0
 MAC pause frames 0 0
 Oversized frames 0
 Jabber frames 0
 Fragment frames 0
 VLAN tagged frames 0
 Code violations 0
Filter statistics:
 Input packet count 0
 Input packet rejects 0
 Input DA rejects 0
 Input SA rejects 0
 Output packet count 0
 Output packet pad count 0
 Output packet error count 0

```

```

 CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
 Negotiation status: Incomplete
Packet Forwarding Engine configuration:
 Destination slot: 1

Logical interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 59) (Generation 135)
Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Local statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Transit statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps
IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
 Protocol inet, Generation: 144, Route table: 0
Flags: uRPF
Addresses, Flags: Is-Preferred Is-Primary

```

**Meaning** The `show interfaces ge-1/0/10 extensive` command (and the `show interfaces ge-1/0/10 detail` command) displays in-depth information about the interface. The **Flags:** output field near the bottom of the display reports the unicast RPF status. If unicast RPF has not been enabled, the **uRPF** flag is not displayed.

On EX3200, EX4200, and EX4300 switches, unicast RPF is implicitly enabled on *all* switch interfaces, including aggregated Ethernet interfaces (also referred to as link aggregation groups or LAGs), integrated routing and bridging (IRB) interfaces, and routed VLAN interfaces (RVIs) when you enable unicast RPF on a single interface. However, the unicast RPF status is shown as enabled only on interfaces for which you have explicitly configured unicast RPF. Thus, the **uRPF** flag is not displayed on interfaces for which you have not explicitly configured unicast RPF even though unicast RPF is implicitly enabled on all interfaces on EX3200 and EX4200 switches.

**Related Documentation**

- *show interfaces xe-*
- *Example: Configuring Unicast RPF on an EX Series Switch*
- [Configuring Unicast RPF \(CLI Procedure\) on page 6054](#)

- [Disabling Unicast RPF \(CLI Procedure\) on page 6055](#)
- *Troubleshooting Unicast RPF*

## Understanding Unknown Unicast Forwarding

Unknown unicast traffic consists of unicast packets with unknown destination MAC addresses. By default, the switch floods these unicast packets that traverse a VLAN to all interfaces that are members of that VLAN. Forwarding this type of traffic can create unnecessary traffic that leads to poor network performance or even a complete loss of network service. This flooding of packets is known as a traffic storm.

To prevent a traffic storm, you can disable the flooding of unknown unicast packets to all VLAN interfaces by configuring specific VLANs or all VLANs to forward all unknown unicast traffic traversing them to a specific interface. You can configure multiple VLANs to forward unknown unicast packets to the same interface or configure different interfaces for different VLANs. This channels the unknown unicast traffic traversing VLANs to specific interfaces instead of flooding all interfaces.

### Related Documentation

- *Configuring Unknown Unicast Forwarding (CLI Procedure)*
- [Configuring Unknown Unicast Forwarding \(CLI Procedure\) on page 6059](#)
- *Understanding Storm Control on EX Series Switches*
- *Understanding Storm Control for Managing Traffic Levels on Switching Devices*
- *Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches*
- *Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches*

## Configuring Unknown Unicast Forwarding (CLI Procedure)



**NOTE:** This task uses Junos OS for EX Series switches or QFX Series with support for the Enhanced Layer 2 Software (ELS) configuration style. If your EX Series switch runs software that does not support ELS, see *Configuring Unknown Unicast Forwarding (CLI Procedure)*. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 41](#)

Unknown unicast traffic consists of packets with unknown destination MAC addresses. By default, the switch floods these packets that traverse a VLAN to all interfaces associated with that VLAN. This flooding of packets is known as a traffic storm and can negatively impact network performance.

To prevent flooding unknown unicast traffic across the switch, configure unknown unicast forwarding to direct all unknown unicast packets within a VLAN to a specific interface.

You can configure each VLAN to divert unknown unicast traffic to a different interface or use the same interface for multiple VLANs.

- [Configuring Unknown Unicast Forwarding on EX4300 Switches on page 6060](#)
- [Configuring Unknown Unicast Forwarding on EX9200 Switches on page 6060](#)

## Configuring Unknown Unicast Forwarding on EX4300 Switches

To configure unknown unicast forwarding options on EX4300 switches:

- Configure unknown unicast forwarding for a specific VLAN and specify the interface to which all unknown unicast traffic will be forwarded:

```
[edit switch-options]
user@switch# set unknown-unicast-forwarding vlan vlan-name interface interface-name
```

- Configure unknown unicast forwarding for all VLANs and specify the interface to which all unknown unicast traffic will be forwarded:

```
[edit switch-options]
user@switch# set unknown-unicast-forwarding vlan all interface interface-name
```

## Configuring Unknown Unicast Forwarding on EX9200 Switches

To configure unknown unicast forwarding on EX9200 switches, you must configure a flood filter and apply it to VLANs for which you want to configure unknown unicast forwarding. Flood filters are firewall filters that are applied only to broadcast, unknown unicast, and multicast (BUM) traffic. If a flood filter is configured, only traffic packets that are of the packet type **unknown-unicast** are forwarded to the interface on which unicast forwarding is configured. A next-hop group redirects the packets according to the action specified in the flood filter.

To configure the next-hop group that receives Layer 2 packets and then configure the interface to which these packets are forwarded:

1. Configure the **next-hop-group** action for the Layer 2 interface expected to receive unknown unicast packets:

```
[edit forwarding-options]
user@switch# set next-hop-group next-hop-group-name group-type layer-2
[edit forwarding-options]
user@switch# set next-hop-group next-hop-group-name interface interface-name
```

For example:

```
[edit forwarding-options]
user@switch# set next-hop-group uuf-nhg group-type layer-2
[edit forwarding-options]
user@switch# set next-hop-group uuf-nhg interface ge-3/1/7.0
```

2. Configure a firewall filter with family address type **ethernet-switching**:

```
[edit firewall]
user@switch# set family ethernet-switching filter filter-name
```

For example:

```
[edit firewall]
user@switch# set family ethernet-switching filter uuf_filter
```

3. Configure a term in the firewall filter for the interface that receives unknown unicast packets (the interface specified in Step 1) to discard unknown unicast packets:

```
[edit firewall family ethernet-switching filter filter-name]
user@switch# set term term-name from interface interface-name
user@switch# set term term-name from traffic-type unknown-unicast
user@switch# set term term-name then discard
```

For example:

```
[edit firewall family ethernet-switching filter uuf_filter]
user@switch# set term source-drop from interface ge-3/1/7.0
user@switch# set term source-drop from traffic-type unknown-unicast
user@switch# set term source-drop then discard
```

4. Configure a term in the firewall filter for unknown unicast packets to be flooded to the interface enabled for unknown unicast forwarding by using **next-hop-group** (in step 1):

```
[edit firewall family ethernet-switching filter filter-name]
user@switch# set term term-name from traffic-type unknown-unicast
user@switch# set term term-name then next-hop-group group-name
```

For example:

```
[edit firewall family ethernet-switching filter uuf_filter]
user@switch# set term uuf-flood from traffic-type unknown-unicast
user@switch# set term uuf-flood then next-hop-group uuf-nhg
```

5. Configure a default term for the firewall filter to forward packets other than unknown unicast packets:

```
[edit firewall family ethernet-switching filter filter-name]
user@switch# set term term-name then accept
```

For example:

```
[edit firewall family ethernet-switching filter uuf_filter]
user@switch# set term fwd-default then accept
```

6. Apply the filter as a flood filter on the VLAN that includes the interface which will receive unknown unicast packets:

```
[edit vlans vlan-name]
user@switch# set forwarding-options flood input filter-name
```

For example:

```
[edit vlans v1]
user@switch# set forwarding-options flood input uuf_filter
```

#### Related Documentation

- [Understanding Unknown Unicast Forwarding on page 6059](#)
- [Verifying That Unknown Unicast Packets Are Forwarded to a Single Interface](#)



## PART 82

# DDoS Protection

- [Overview of DDoS Protection on page 6065](#)
- [Configuring DDoS Protection on page 6069](#)
- [Monitoring DDoS Protection on page 6083](#)





# Overview of DDoS Protection

- [Understanding Distributed Denial-of-Service Protection on QFX Series Switches on page 6065](#)

## Understanding Distributed Denial-of-Service Protection on QFX Series Switches

---

A denial-of-service (DoS) attack is any attempt to deny valid users access to network or server resources by using up all the resources of the network element or server. Distributed denial-of-service attacks (DDoS) involve an attack from multiple sources, enabling a much greater amount of traffic to attack the network. The attacks typically use network protocol control packets to trigger a large number of exceptions to the router or switch control plane. This results in an excessive processing load that disrupts normal network operations.

Junos OS DDoS protection enables both QFX10000 switches and QFX5200 switches to continue functioning while under attack. It identifies and suppresses malicious control packets while enabling legitimate control traffic to be processed. A single point of DDoS protection management enables network administrators to customize profiles for their network control traffic.

To protect against DDoS attacks, you can configure policers for host-bound exception traffic. The policers specify rate limits for all control traffic for a given protocol, or, in some cases, for specific control packet types for a protocol. Control traffic is dropped when it exceeds any configured policer values or, for unconfigured policers, the default policer values. Each violation immediately generates a notification to alert operators about a possible attack. The violation is counted, the time that the violation starts is noted, and the time of the last observed violation is noted. When the traffic rate drops below the bandwidth violation threshold, a recovery timer determines when the traffic flow is considered to have returned to normal. If no further violation occurs before the timer expires, the violation state is cleared and a notification is generated. On QFX Series switches, the timer is set to 300 seconds and cannot be modified.

In addition providing notification of violations through event logging, Junos OS DDoS protection allows you to monitor policers, obtaining information such as the policer configuration, number of violations encountered, date and time of violations, packet arrival rates, and number of packets received or dropped.

## Policer Types and Packet Priorities

DDoS protection includes two types of policers:

- An *aggregate policer* is applied to the complete set of packet types that belong to a protocol group. For example, you can configure an aggregate policer that applies to all RADIUS control packet types or to all DHCP control packet types. You can specify bandwidth and burst limits for aggregate policers. An aggregate policer is available for all protocol groups. Aggregate policers are supported by all protocol groups.
- An *individual policer*, also referred to as a *packet-type policer*, is allocated for a specific control packet type within a protocol group. For example, you can configure a policer for one or more types of RADIUS control packets. You can specify bandwidth and burst limits, prioritize one packet type over another, and enable a packet type to bypass the aggregate policer for the protocol group.

Individual policers are not available for all protocol groups. See [protocols](#) for a list of protocol groups that have individual policers.

A control packet is policed first by its individual policer (if supported) and then by its aggregate policer. A packet dropped by the individual policer never reaches the aggregate policer. A packet that passes the individual policer can subsequently be dropped by the aggregate policer.

Each packet type within a protocol group has a default, configurable priority: low, medium, or high. Each control packet competes with other packets for the bandwidth within the limit imposed by its aggregate policer based on the priority configured for each packet type in the protocol group.

The priority mechanism is absolute. High-priority traffic gets bandwidth in preference to medium-priority and low-priority traffic. Medium-priority traffic gets bandwidth in preference to low-priority traffic. Low-priority traffic can use only the bandwidth left by high-priority and medium-priority traffic. If higher-priority traffic takes all of the bandwidth, then all the lower-priority traffic is dropped.

## Example of Policer Priority Behavior

For example, consider how you might configure packet types within the RADIUS protocol group. Suppose you configure individual policers for accounting and authorization packets, as well as a RADIUS aggregate policer for all RADIUS control packets. You might want to prioritize the RADIUS authorization function over the RADIUS accounting function, and therefore you would assign a high priority to the authorization control packets and a low priority to accounting control packets.

The aggregate policer imposes a total bandwidth limit for the protocol group. Authorization packets passed by their individual policer have access to that bandwidth before accounting packets passed by their individual policer, because the authorization packets have a higher priority. If enough authorization packets are passed that they use all the available bandwidth, then all the accounting packets are dropped, because there is no bandwidth remaining at the aggregate policer.

## Policer Enforcement Points on QFX Series Switches

On both QFX1000 and QFX5200 switches, the DDoS policers operate at the line card level. Control traffic arriving from all ports of a line card converges on the Packet Forwarding Engine, where it is subject to policing. Thus, excess packets are dropped before they reach the Routing Engine, ensuring that the Routing Engine receives only the amount of traffic it can process.

- Related Documentation**
- [Configuring Protection Against DDoS Attacks on page 6069](#)
  - [Verifying and Managing DDoS Protection on page 6083](#)



# Configuring DDoS Protection

- [Configuring Protection Against DDoS Attacks on page 6069](#)
- [Example: Configuring DDoS Protection on QFX Series Switches on page 6070](#)
- [Disabling DDoS Protection Policers and Logging Globally on page 6074](#)
- [Configuring DDoS Protection Policers on QFX Series Switches on page 6074](#)
- [Tracing DDoS Protection Operations on page 6078](#)

## Configuring Protection Against DDoS Attacks

---

DDoS protection is enabled by default for all supported protocol groups and packet types. Default values are present for bandwidth, bandwidth scale, burst, burst scale, priority, and recover time. You can change the DDoS configuration for individual packet types within a protocol group or for the aggregate policer for the protocol group. DDoS logging is enabled by default, but you can disable it globally for all DDoS events or for individual packet types within a protocol group. You can also fine-tune monitoring of DDoS events by configuring tracing operations.

You can disable DDoS protection at the Routing Engine and for all line cards either globally or for individual packet types within a protocol group.



**NOTE:** DDoS protection is supported only on MX Series routers that have only MPCs installed, T4000 routers that have only FPC5s installed, EX9200 switches, QFX5200 switches, and QFX10000 switches. If the router platforms have other line cards in addition to MPCs (MX Series) or FPC5s (T4000), the CLI accepts the configuration but the other line cards are not protected and so the router is not protected. Neither QFX10002 switches nor QFX5200 switches support policers at the Routing Engine.

---

To configure DDoS protection:

1. (Optional) Configure global DDoS settings.  
[See “Disabling DDoS Protection Policers and Logging Globally” on page 6074.](#)
2. (Optional) Configure DDoS settings for individual packet types.

For MX Series routers, T4000 routers, or EX9200 switches, see *Configuring DDoS Protection Policers for Individual Packet Types*. For QFX10000 switches, see [“Configuring DDoS Protection Policers on QFX Series Switches” on page 6074](#).

3. (Optional) Configure tracing for DDoS operations.

See *Tracing DDoS Protection Operations*.

**Related Documentation**

- *Distributed Denial-of-Service (DDoS) Protection Overview*
- [Understanding Distributed Denial-of-Service Protection on QFX Series Switches on page 6065](#)
- *Example: Configuring DDoS Protection*
- [Example: Configuring DDoS Protection on QFX Series Switches on page 6070](#)

---

## Example: Configuring DDoS Protection on QFX Series Switches

This example shows how to configure DDoS protection that enables a switch to quickly identify an attack and prevent a flood of malicious control packets from exhausting system resources.

- [Requirements on page 6070](#)
- [Overview on page 6070](#)
- [Configuration on page 6071](#)
- [Verification on page 6072](#)

### Requirements

DDoS protection requires the following hardware and software:

- QFX Series switch that supports DDoS protection
- Junos OS Release 15.1X53-D10 or later

No special configuration beyond device initialization is required before you can configure this feature.

### Overview

Distributed denial-of-service (DDoS) attacks use multiple sources to flood a network with protocol control packets. This malicious traffic triggers a large number of exceptions in the network and attempts to exhaust the system resources to deny valid users access to the network or server.

DDoS protection is enabled by default on a supported QFX Series switch. This example describes how you can modify the default configuration for the rate-limiting policers that identify excess control traffic and drop the packets before the switch is adversely affected. Sample tasks include configuring an aggregate policer for a protocol group, configuring policers for particular control packet types within a protocol group, and specifying trace options for DDoS operations.

This example does not show all possible configuration choices.

## Configuration

**CLI Quick Configuration** To quickly configure DDoS protection for protocol groups and particular control packet types, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
[edit]
edit system
set ddos-protection protocols radius aggregate bandwidth 150
set ddos-protection protocols radius aggregate burst 2000
set ddos-protection protocols radius accounting bandwidth 100 burst 150
set ddos-protection protocols radius accounting priority low
set ddos-protection protocols radius server bypass-aggregate
set ddos-protection traceoptions file ddos-trace size 10m
set ddos-protection traceoptions flag all
top
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure DDoS protection:

1. Specify a protocol group.  

```
[edit system ddos-protection protocols]
user@host# edit radius
```
2. Configure the maximum traffic rate for the RADIUS aggregate policer; that is, for the combination of all RADIUS packets.  

```
[edit system ddos-protection protocols radius]
user@host# set aggregate bandwidth 150
```
3. Configure the maximum burst rate for the RADIUS aggregate policer.  

```
[edit system ddos-protection protocols radius]
user@host# set aggregate burst 2000
```
4. Configure a different maximum traffic rate and burst size for RADIUS accounting packets.  

```
[edit system ddos-protection protocols radius]
user@host# set accounting bandwidth 100 burst 1500
```
5. Decrease the priority for RADIUS accounting packets.  

```
[edit system ddos-protection protocols radius]
user@host# set accounting priority low
```
6. Prevent RADIUS server control packets from being included in the aggregate bandwidth; that is, server packets do not contribute toward the combined RADIUS traffic to determine whether the aggregate bandwidth is exceeded. However, the server packets are still included in traffic rate statistics.

```
[edit system ddos-protection protocol radius]
user@host# set server bypass-aggregate
```

7. Configure tracing for all DDoS protocol processing events.

```
[edit system ddos-protection traceoptions]
user@host# set file ddos-log
user@host# set file size 10m
user@host# set flag all
```

**Results** From configuration mode, confirm your configuration by entering the **show ddos-protection** command at the **system** hierarchy level.

```
[edit system]

user@host# show ddos-protection

traceoptions {
 file ddos-log size 10m;
 flag all;
}
protocols {

 radius {
 aggregate {
 bandwidth 150;
 burst 2000;
 }
 server {
 bypass-aggregate;
 }
 accounting {
 bandwidth 100;
 burst 1500;
 priority low;
 }
 }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

## Verification

To confirm that the DDoS protection configuration is working properly, perform these tasks:

- [Verifying the DDoS Protection Configuration on page 6072](#)

### Verifying the DDoS Protection Configuration

**Purpose** Verify that the RADIUS policer values have changed from the default.

**Action** From operational mode, enter the **show ddos-protection protocols radius parameters** command.

```
user@host> show ddos-protection protocols radius parameters
Packet types: 5, Modified: 3
* = User configured value
```



## Protocol Group: Radius

```

Packet type: aggregate (Aggregate for all Radius traffic)
Aggregate policer configuration:
 Bandwidth: 150 pps*
 Burst: 2000 packets*
 Recover time: 300 seconds
 Enabled: Yes
Routing Engine information:
 Bandwidth: 150 pps, Burst: 2000 packets, enabled
FPC slot 0 information:
 Bandwidth: 100% (150 pps), Burst: 100% (2000 packets), enabled

Packet type: server (Radius server traffic)
Individual policer configuration:
 Bandwidth: 200 pps
 Burst: 2048 packets
 Priority: High
 Recover time: 300 seconds
 Enabled: Yes
 Bypass aggregate: Yes*
Routing Engine information:
 Bandwidth: 200 pps, Burst: 2048 packets, enabled
FPC slot 0 information:
 Bandwidth: 100% (200 pps), Burst: 100% (2048 packets), enabled

Packet type: accounting (Radius accounting traffic)
Individual policer configuration:
 Bandwidth: 100 pps*
 Burst: 1500 packets*
 Priority: Low*
 Recover time: 300 seconds
 Enabled: Yes
 Bypass aggregate: No
Routing Engine information:
 Bandwidth: 100 pps, Burst: 1500 packets, enabled
FPC slot 0 information:
 Bandwidth: 100% (100 pps), Burst: 100% (1500 packets), enabled

Packet type: authorization (Radius authorization traffic)
Individual policer configuration:
 Bandwidth: 200 pps
 Burst: 2048 packets
 Priority: High
 Recover time: 300 seconds
 Enabled: Yes
 Bypass aggregate: No
Routing Engine information:
 Bandwidth: 200 pps, Burst: 2048 packets, enabled
FPC slot 0 information:
 Bandwidth: 100% (200 pps), Burst: 100% (2048 packets), enabled

```

**Meaning** The command output shows the current configuration of the RADIUS aggregate policer and the RADIUS accounting, server, and authorization control packet policers. Policers values that have been modified from the default values are marked with an asterisk. The output shows that the RADIUS policer configuration has been modified correctly.

- Related Documentation**
- [Understanding Distributed Denial-of-Service Protection on QFX Series Switches on page 6065](#)
  - [Configuring Protection Against DDoS Attacks on page 6069](#)

---

## Disabling DDoS Protection Policers and Logging Globally

---

DDoS policers are enabled by default for all supported protocol groups and packet types.

On MX Series routers, T4000 routers, and EX9200 switches, policers are established at the level of the individual line card and the Routing Engine. You can disable the line card policers globally for all MPCs or FPC5s. You can also disable the Routing Engine policer. When you disable either of these policers, the policers at that level for all protocol groups and packet types are disabled.

On both QFX10002 switches and QFX5200 switches, policers are established at the level of individual line cards only. If you disable line-card policers globally, DDoS protection is disabled on the switch.

DDoS logging is also enabled by default. You can disable all DDoS event logging (including flow detection event logging) for all protocol groups and packet types across the router or switch.



**NOTE:** The global configuration for disabling policers and logging overrides any local configuration for packet types.

---

To configure global DDoS settings:

1. (Optional) Disable line card policers.  

```
[edit system ddos-protection global]
user@host# set disable-fpc
```
2. (Optional) Disable Routing Engine policers (not supported on QFX10002 switches).  

```
[edit system ddos-protection global]
user@host# set disable-routing-engine
```
3. (Optional) Disable event logging.  

```
[edit system ddos-protection global]
user@host# set disable-logging
```

- Related Documentation**
- [Configuring Protection Against DDoS Attacks on page 6069](#)

---

## Configuring DDoS Protection Policers on QFX Series Switches

---

You can modify the DDoS protection configuration as follows:

- Modify the aggregate policer bandwidth and burst values for a protocol group. Default values exist for all protocol groups. See [protocols](#) for the supported protocol groups and their default policer values.
- Modify the policer bandwidth and burst values for individual control packet types within a protocol group, for those groups that support policers for individual packet types. You can specify that packets of a certain type have a higher or lower priority than other types. You can also specify that a packet type bypass the aggregate policer for the protocol group. See [protocols](#) for the supported packet types and their default policer values.
- Scale the bandwidth and burst values for a policer on a line card so that the policer triggers at lower thresholds than the overall protocol thresholds.
- Disable logging for a specific policer.
- Disable a policer on all line cards or on an individual line card. On devices with a single line card, disabling policers on the line card is effectively the same as disabling the policers globally. Note that deleting the configuration for a policer does not disable it—the policer merely reverts to its default settings.



**BEST PRACTICE:** We recommend that you model your network to determine the best values for your situation. Before you configure policers for your network, you can quickly view the default values for all protocol groups and packet types from operational mode by issuing the [show ddos-protection protocols parameters brief](#) command. You can also use the command to specify a single protocol group of interest; for example, issue the [show ddos-protection protocols radius parameters brief](#) command.

This topic describes:

- [Configuring the Aggregate Policer for a Protocol Group on page 6075](#)
- [Configuring Packet-Type Policers for a Protocol Group on page 6076](#)
- [Configuring Policers on Individual Line Cards on page 6077](#)
- [Disabling Policers and Policer Logging on page 6077](#)

## Configuring the Aggregate Policer for a Protocol Group

An aggregate policer exists for each protocol group. The aggregate policer enforces the traffic limits on the control packets for that protocol as a combined group.

To configure the DDoS aggregate policer for a protocol group:

1. Specify the aggregate policer for the protocol group.

```
[edit system ddos-protection protocols]
user@host# edit protocol-group aggregate
```

For example, to specify the DHCPv4v6 aggregate policer:

```
[edit system ddos-protection protocols]
user@host# edit dhcpv4v6 aggregate
```

2. (Optional) Configure the maximum traffic rate the policer allows for the protocol group.

```
[edit system ddos-protection protocols protocol-group aggregate]
user@host# set bandwidth packets-per-second
```

For example, to set a bandwidth of 300 packets per second for DHCPv4 and DHCPv6 packets:

```
[edit system ddos-protection protocols dhcpv4v6 aggregate]
user@host# set bandwidth 300
```

3. (Optional) Configure the maximum number of packets that the policer allows in a burst of traffic.

```
[edit system ddos-protection protocols protocol-group aggregate]
user@host# set burst size
```

For example, to set a maximum of 1500 DHCPv4v6 packets:

```
[edit system ddos-protection protocols dhcpv4v6 aggregate]
user@host# set burst 1500
```

## Configuring Packet-Type Policers for a Protocol Group

Some protocol groups allow you to configure a separate policer for each control packet type. Control traffic is subject first to the packet-type policer and then to the aggregate policer.

To configure a packet-type policer:

1. Specify the protocol group and packet type.

```
[edit system ddos-protection protocols]
user@host# edit protocol-group packet-type
```

For example, to specify the RADIUS protocol group and the authorization packet type:

```
[edit system ddos-protection protocols]
user@host# edit radius authorization
```

2. (Optional) Configure the maximum traffic rate the policer allows for the packet type.

```
[edit system ddos-protection protocols protocol-group packet-type]
user@host# set bandwidth packets-per-second
```

For example, to set a bandwidth of 150 packets per second for RADIUS authorization packets:

```
[edit system ddos-protection protocols radius authorization]
user@host# set bandwidth 150
```

3. (Optional) Configure the maximum number of packets of the packet type that the policer allows in a burst of traffic.

```
[edit system ddos-protection protocols protocol-group packet-type]
user@host# set burst size
```

For example, to set a maximum of 2000 RADIUS authorization packets:

```
[edit system ddos-protection protocols radius authorization]
```

```
user@host# set burst 2000
```

4. (Optional) Set the traffic priority—either high, medium, or low.

```
[edit system ddos-protection protocols protocol-group packet-type]
```

```
user@host# set priority level
```

For example, to specify a low priority for RADIUS accounting packets:

```
[edit system ddos-protection protocols radius accounting]
```

```
user@host# set priority low
```

5. (Optional) Allow packets of the specified type to bypass the aggregate policer.

```
[edit system ddos-protection protocols protocol-group packet-type]
```

```
user@host# set bypass-aggregate
```

For example, to bypass the aggregate policer for RADIUS server packets:

```
[edit system ddos-protection protocols radius server]
```

```
user@host# set bypass-aggregate
```

## Configuring Policers on Individual Line Cards

You can alter a policer behavior on a specific line card by scaling the policer's configured bandwidth and burst values. On switches with a single fixed line card, such as the QFX10002, scaling the policer values affects the entire switch.

- To scale the maximum bandwidth for a policer on a line card:

```
[edit system ddos-protection protocols protocol-group (aggregate | packet-type)]
```

```
user@host# set fpc slot-number bandwidth-scale percentage
```

For example, to scale the maximum bandwidth allowed by the DHCPv4v6 aggregate policer for the line card in slot 3 to 80 percent:

```
[edit system ddos-protection protocols dhcpv4v6 aggregate]
```

```
user@host# set fpc 3 bandwidth-scale 80
```

- To scale the maximum burst size for a policer on the line card:

```
[edit system ddos-protection protocols protocol-group (aggregate | packet-type)]
```

```
user@host# set fpc slot-number burst-scale percentage
```

For example, to scale the maximum burst size to 75 percent for the RADIUS server packets on the line card in slot 1:

```
[edit system ddos-protection protocols radius server]
```

```
user@host# set fpc 1 burst-scale 75
```

## Disabling Policers and Policer Logging

All supported policers are enabled by default. You can disable specific policers on a line card or line cards. Similarly, event logging by policers is enabled by default. You can selectively disable logging by a policer.

- To disable a policer on a specific line card:

```
[edit system ddos-protection protocols]
```

```
user@host# set protocol-group (aggregate | packet-type fpc slot-number) disable-fpc
```

For example, to disable the DDoS policers for the RADIUS authorization packet type on line card 3:

```
[edit system ddos-protection protocols]
user@host# set radius authorization fpc 3 disable-fpc
```

Because both QFX10002 and QFX5200 have a single line card, disabling a policer on that line card effectively disables it for the switch.

- To disable a policer on all line cards:

```
[edit system ddos-protection protocols]
user@host# set protocol-group (aggregate | packet-type) disable-fpc
```

For example, to disable the aggregate policer for the BFD protocol group on all line cards:

```
[edit system ddos-protection protocols]
user@host# set bfd aggregate disable-fpc
```

- To disable event logging by a policer:

```
[edit system ddos-protection protocols]
user@host# set protocol-group (aggregate | packet-type) disable-logging
```

For example, to disable logging by the aggregate BFD policer:

```
[edit system ddos-protection protocols]
user@host# set bfd aggregate disable-logging
```

#### Related Documentation

- [Configuring Protection Against DDoS Attacks on page 6069](#)
- [Example: Configuring DDoS Protection on QFX Series Switches on page 6070](#)

---

## Tracing DDoS Protection Operations

The Junos OS trace feature tracks DDoS protection operations and records events in a log file. The error descriptions captured in the log file provide detailed information to help you solve problems.

By default, nothing is traced. When you enable the tracing operation, the default tracing behavior is as follows:

1. Important events are logged in a file located in the `/var/log` directory. By default, the router uses the filename `jddosd`. You can specify a different filename, but you cannot change the directory in which trace files are located.
2. When the trace log file *filename* reaches 128 kilobytes (KB), it is compressed and renamed *filename.0.gz*. Subsequent events are logged in a new file called *filename*, until it reaches capacity again. At this point, *filename.0.gz* is renamed *filename.1.gz* and *filename* is compressed and renamed *filename.0.gz*. This process repeats until the number of archived files reaches the maximum file number. Then the oldest trace file—the one with the highest number—is overwritten.

You can optionally specify the number of trace files to be from 2 through 1000. You can also configure the maximum file size to be from 10 KB through 1 gigabyte (GB).

(For more information about how log files are created, see the *Junos OS System Log Messages Reference*.)

By default, only the user who configures the tracing operation can access log files. You can optionally configure read-only access for all users.

This topic describes how you can configure all aspects of DDoS tracing operations. It covers:

- [Configuring the DDoS Protection Trace Log Filename on page 6079](#)
- [Configuring the Number and Size of DDoS Protection Log Files on page 6079](#)
- [Configuring Access to the DDoS Protection Log File on page 6080](#)
- [Configuring a Regular Expression for DDoS Protection Messages to Be Logged on page 6080](#)
- [Configuring the DDoS Protection Tracing Flags on page 6080](#)
- [Configuring the Severity Level to Filter Which DDoS Protection Messages Are Logged on page 6081](#)

## Configuring the DDoS Protection Trace Log Filename

By default, the name of the file that records trace output for DDoS protection is **jddosd**. You can specify a different name with the **file** option.

To configure the filename for subscriber management database tracing operations:

- Specify the name of the file used for the trace output.

```
[edit system ddos-protection traceoptions]
user@host# set file ddos_logfile_1
```

## Configuring the Number and Size of DDoS Protection Log Files

You can optionally specify the number of compressed, archived trace log files to be from 2 through 1000. You can also configure the maximum file size to be from 10 KB through 1 gigabyte (GB); the default size is 128 kilobytes (KB).

The archived files are differentiated by a suffix in the format **.number.gz**. The newest archived file is **.0.gz** and the oldest archived file is **.(maximum number)-1.gz**. When the current trace log file reaches the maximum size, it is compressed and renamed, and any existing archived files are renamed. This process repeats until the maximum number of archived files is reached, at which point the oldest file is overwritten.

For example, you can set the maximum file size to 2 MB, and the maximum number of files to 20. When the file that receives the output of the tracing operation, **filename**, reaches 2 MB, **filename** is compressed and renamed **filename.0.gz**, and a new file called **filename** is created. When the new **filename** reaches 2 MB, **filename.0.gz** is renamed **filename.1.gz** and **filename** is compressed and renamed **filename.0.gz**. This process repeats until there are 20 trace files. Then the oldest file, **filename.19.gz**, is simply overwritten when the next oldest file, **filename.18.gz** is compressed and renamed to **filename.19.gz**.

To configure the number and size of trace files:

- Specify the name, number, and size of the file used for the trace output.

```
[edit system ddos-protection traceoptions]
user@host# set file ddos_1_logfile_1 files 20 size 2097152
```

## Configuring Access to the DDoS Protection Log File

By default, only the user who configures the tracing operation can access the log files. You can enable all users to read the log file and you can explicitly set the default behavior of the log file.

To specify that all users can read the log file:

- Configure the log file to be world-readable.

```
[edit system ddos-protection traceoptions]
user@host# set file ddos_1_logfile_1 world-readable
```

To explicitly set the default behavior, only the user who configured tracing can read the log file:

- Configure the log file to be no-world-readable.

```
[edit system ddos-protection traceoptions]
user@host# set file ddos_1_logfile_1 no-world-readable
```

## Configuring a Regular Expression for DDoS Protection Messages to Be Logged

By default, the trace operation output includes all messages relevant to the logged events.

You can refine the output by including regular expressions to be matched.

To configure regular expressions to be matched:

- Configure the regular expression.

```
[edit system ddos-protection traceoptions]
user@host# set file ddos_1_logfile_1 match regex
```

## Configuring the DDoS Protection Tracing Flags

By default, only important events are logged. You can specify which events and operations are logged by specifying one or more tracing flags.

To configure the flags for the events to be logged:

- Configure the flags.

```
[edit system ddos-protection traceoptions]
user@host# set flag flag
```



## Configuring the Severity Level to Filter Which DDoS Protection Messages Are Logged

The messages associated with a logged event are categorized according to severity level. You can use the severity level to determine which messages are logged for the event type. The severity level that you configure depends on the issue that you are trying to resolve. In some cases you might be interested in seeing all messages relevant to the logged event, so you specify **all** or **verbose**. Either choice generates a large amount of output. You can specify a more restrictive severity level, such as **notice** or **info** to filter the messages. By default, the trace operation output includes only messages with a severity level of **error**.

To configure the type of messages to be logged:

- Configure the message severity level.

```
[edit system ddos-protection traceoptions]
user@host# set level severity
```

### Related Documentation

- [Configuring Protection Against DDoS Attacks on page 6069](#)



# Monitoring DDoS Protection

- [Verifying and Managing DDoS Protection on page 6083](#)

## Verifying and Managing DDoS Protection

---

**Purpose** View or clear information about DDoS configurations, states, and statistics.

**Action**

- To display the DDoS policer configuration, violation state, and statistics for all packet types in all protocol groups:

user@host> **show ddos-protection protocols**

If you issue the command before you make any configuration changes, the default policer values are displayed.

- To display the DDoS policer configuration, violation state, and statistics for a particular packet type in a particular protocol group:

user@host> **show ddos-protection protocols protocol-group packet-type**

- To display only the number of DDoS policer violations for all protocol groups:

user@host> **show ddos-protection protocols violations**

- To display a table of the DDoS configuration for all packet types in all protocol groups:

user@host> **show ddos-protection protocols parameters brief**

- To display a complete list of packet statistics and DDoS violation statistics for all packet types in all protocol groups:

user@host> **show ddos-protection protocols statistics detail**

- To display global DDoS violation statistics:

user@host> **show ddos-protection statistics**

- To display the DDoS version number:

user@host> **show ddos-protection version**

- To clear DDoS statistics for all packet types in all protocol groups:

user@host> **clear ddos-protection protocols statistics**

- To clear DDoS statistics for all packet types in a particular protocol group:

user@host> **clear ddos-protection protocols protocol-group statistics**

- To clear DDoS statistics for a particular packet type in a particular protocol group:

user@host> **clear ddos-protection protocols protocol-group statistics packet-type**

- To clear DDoS violation states for all packet types in all protocol groups:

user@host> **clear ddos-protection protocols states**

- To clear DDoS violation states for all packet types in a particular protocol group:

user@host> **clear ddos-protection protocols protocol-group states**

- To clear DDoS violation states for a particular packet type in a particular protocol group:

user@host> **clear ddos-protection protocols protocol-group states packet-type**

**Related  
Documentation**

- *Verifying and Managing Flow Detection*

## PART 83

# Configuration Statements and Operational Commands

- [Configuration Statements \(Firewall Filters\) on page 6087](#)
- [Configuration Statements \(Policers\) on page 6099](#)
- [Configuration Statements \(Device Security\) on page 6119](#)
- [Configuration Statements \(Port Security\) on page 6131](#)
- [Configuration Statements \(DDoS Protection\) on page 6145](#)
- [Operational Commands \(Firewall Filters\) on page 6165](#)
- [Operational Commands \(Port Security\) on page 6177](#)
- [Operational Commands \(DDoS Protection\) on page 6179](#)



## CHAPTER 228

# Configuration Statements (Firewall Filters)

- [family](#) on page 6088
- [filter](#) on page 6089
- [filter \(Layer 2 and Layer 3 Interfaces\)](#) on page 6090
- [filter \(VLANs\)](#) on page 6091
- [firewall](#) on page 6092
- [from](#) on page 6093
- [input \(Forwarding Table\)](#) on page 6094
- [interface-specific](#) on page 6094
- [output \(Forwarding Table\)](#) on page 6095
- [term](#) on page 6096
- [then \(Filters\)](#) on page 6097

## family

---

**Syntax**    family *family-name* {  
              filter *filter-name* {  
                  interface-specific;  
                  term *term-name* {  
                      from {  
                        match-conditions;  
                      }  
                      then {  
                        action;  
                        action-modifiers;  
                      }  
                  }  
              }  
          }

**Hierarchy Level**    [edit [firewall](#)]

**Release Information**    Statement introduced in Junos OS Release 11.1 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.  
**evpn** options introduced in Junos OS Release 15.1 for the MX Series.

**Description**    Configure the fields a firewall filter can match on.

**Options**    *family-name*—Type of addressing protocol:

- **ethernet-switching**—Filter Layer 2 Ethernet packets and Layer 3 (IP) packets (allows some Layer 3 filtering). Not supported on OCX Series switches.
- **evpn**—Filter Ethernet VPN (EVPN) packets.
- **inet**—Filter Layer 3 IPv4 packets (provides additional Layer 3 filter options).
- **inet6**—Filter Layer 3 IPv6 packets (provides additional Layer 3 filter options).
- **mpls**—Filter multiprotocol label switched packets. Not supported on OCX Series switches.

The remaining statements are explained separately.

**Required Privilege Level**    interface—To view this statement in the configuration.  
                                  interface-control—To add this statement to the configuration.

**Related Documentation**

- *Firewall Filter Match Conditions and Actions*
- [Configuring Firewall Filters on page 5978](#)
- [Overview of Firewall Filters on page 5951](#)



## filter

|                                 |                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>filter <i>filter-name</i> {     <i>interface-specific</i>;     term <i>term-name</i> {         from {             <i>match-conditions</i>;         }         then {             <i>action</i>;             <i>action-modifiers</i>;         }     } }</pre>                       |
| <b>Hierarchy Level</b>          | [edit <b>firewall family</b> <i>family-name</i> ]                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                          |
| <b>Description</b>              | Configure firewall filters.                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <p><b><i>filter-name</i></b>—Name that identifies the filter. The name can contain letters, numbers, and hyphens (-), and can be up to 64 characters long. To include spaces in the name, enclose it in quotation marks.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>firewall—To view this statement in the configuration.</p> <p>firewall-control—To add this statement to the configuration.</p>                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Firewall Filter Match Conditions and Actions</i></li> <li><a href="#">Configuring Firewall Filters on page 5978</a></li> <li><a href="#">Overview of Firewall Filters on page 5951</a></li> </ul>                                            |

## filter (Layer 2 and Layer 3 Interfaces)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>filter (input   output) <i>filter-name</i>;</code>                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [ <a href="#">edit</a> <a href="#">interfaces</a> <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> <a href="#">family</a> <i>family-name</i> ]                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Apply a firewall filter to traffic transiting a port or Layer 3 interface.                                                                                                                                                                                                                                                                                                                                |
| <b>Default</b>                  | All incoming traffic is accepted unmodified on the port or Layer 3 interface, and all outgoing traffic is sent unmodified from the port or Layer 3 interface.                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <p><b><i>filter-name</i></b>—Name of a firewall filter defined at the [<a href="#">edit</a> <a href="#">firewall</a> <a href="#">family</a> <i>family-name</i> <a href="#">filter</a>] hierarchy level.</p> <p><b>input</b>—Apply a firewall filter to traffic entering the port or Layer 3 interface.</p> <p><b>output</b>—Apply a firewall filter to traffic exiting the port or Layer 3 interface.</p> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Gigabit Ethernet Interfaces (CLI Procedure)</a></li><li>• <a href="#">Configuring Firewall Filters on page 5978</a></li><li>• <a href="#">Overview of Firewall Filters on page 5951</a></li></ul>                                                                                                                                         |

## filter (VLANs)

---

|                                 |                                                                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>filter (input   output) <i>filter-name</i>;</code>                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | <code>[edit vlans <i>vlan-name</i>],</code><br><code>[edit vlans <i>vlan-name</i> forwarding-options]</code>                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.                                                                                                                                                  |
| <b>Description</b>              | Apply a firewall filter to traffic entering or exiting a VLAN.                                                                                                                                                                                                                                     |
| <b>Default</b>                  | All incoming traffic is accepted unmodified to a VLAN, and all outgoing traffic is sent unmodified from a VLAN.                                                                                                                                                                                    |
| <b>Options</b>                  | <p><b><i>filter-name</i></b>—Name of a firewall filter defined at the <code>[edit firewall family <i>family-name</i> filter]</code> hierarchy level.</p> <p><b>input</b>—Apply a firewall filter to VLAN ingress traffic.</p> <p><b>output</b>—Apply a firewall filter to VLAN egress traffic.</p> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Firewall Filters on page 5978</a></li> <li>• <a href="#">Overview of Firewall Filters on page 5951</a></li> </ul>                                                                                                                 |

## firewall

---

```
Syntax firewall {
 family family-name {
 filter filter-name {
 interface-specific;
 term term-name {
 from {
 match-conditions;
 }
 then {
 action;
 action-modifiers;
 }
 }
 }
 }
 policer policer-name {
 filter-specific;
 if-exceeding {
 bandwidth-limit bps;
 burst-size-limit bytes;
 }
 then {
 policer-action;
 }
 }
 three-color-policer policer-name {
 action {
 loss-priority high then discard;
 }
 single-rate {
 (color-aware | color-blind);
 committed-information-rate bps;
 committed-burst-size bytes;
 excess-burst-size bytes;
 }
 two-rate {
 (color-aware | color-blind);
 committed-information-rate bps;
 committed-burst-size bytes;
 peak-information-rate bps;
 peak-burst-size bytes;
 }
 }
 }
```

Hierarchy Level [\[edit\]](#)

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Configure firewall filters and policers.

The remaining statements are explained separately.

|                                 |                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | firewall—To view this statement in the configuration.<br>firewall-control—To add this statement to the configuration.                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Firewall Filter Match Conditions and Actions</a></li> <li>• <a href="#">Configuring Firewall Filters on page 5978</a></li> <li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 6017</a></li> <li>• <a href="#">Overview of Firewall Filters on page 5951</a></li> </ul> |

## from

|                                 |                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>from {     match-conditions; }</pre>                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit <b>firewall family</b> <i>family-name</i> <b>filter</b> <i>filter-name</i> <b>term</b> <i>term-name</i> ]                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                |
| <b>Description</b>              | Match packet fields to values specified in a match condition. If the <b>from</b> statement is not included in a firewall filter configuration, all packets are considered to match and the actions and action modifiers in the <b>then</b> statement are implemented.                                                        |
| <b>Options</b>                  | <b>match-conditions</b> —Conditions that define the values or fields that the incoming or outgoing packets must contain for a match. You can specify one or more match conditions. If you specify more than one, they all must match for a match to occur and for the action in the <b>then</b> statement to be implemented. |
| <b>Required Privilege Level</b> | firewall—To view this statement in the configuration.<br>firewall-control—To add this statement to the configuration.                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Firewall Filter Match Conditions and Actions</a></li> <li>• <a href="#">Configuring Firewall Filters on page 5978</a></li> <li>• <a href="#">Understanding Firewall Filter Match Conditions on page 5957</a></li> </ul>                                                 |

## input (Forwarding Table)

---

|                                 |                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>input <i>filter-name</i>;</code>                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit forwarding-options family (inet   inet6   mpls   vpls) filter],<br>[edit routing-instances <i>routing-instance-name</i> forwarding-options family (inet   inet6   mpls   vpls) filter] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 15.1X53-D10 for QFX10000 switches..                                                            |
| <b>Description</b>              | Apply a forwarding table filter to ingress traffic of the forwarding table.                                                                                                                  |
| <b>Options</b>                  | <i>filter-name</i> —Name of the applied filter.                                                                                                                                              |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Applying Forwarding Table Filters</a></li></ul>                                                                                          |

## interface-specific

---

|                                 |                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>interface-specific;</code>                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit <a href="#">firewall family</a> <i>family-name</i> <a href="#">filter</a> <i>filter-name</i> ]                                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                          |
| <b>Description</b>              | Configure separate counters for each interface to which a filter is applied.                                                                                                                                                                           |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Firewall Filter Match Conditions and Actions</a></li><li>• <a href="#">Configuring Firewall Filters on page 5978</a></li><li>• <a href="#">Overview of Firewall Filters on page 5951</a></li></ul> |

---

## output (Forwarding Table)

---

|                                 |                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>output <i>filter-name</i>;</code>                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit forwarding-options family (inet   inet6   mpls) filter],<br>[edit routing-instances <i>routing-instance-name</i> forwarding-options family (inet   inet6   mpls) filter] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.5.<br>Statement introduced in Junos OS Release 15.1X53-D10 for QFX10000 switches..                                                  |
| <b>Description</b>              | Configure filtering on the egress traffic of the forwarding table.                                                                                                             |
| <b>Options</b>                  | <i>filter-name</i> —Name of the applied filter.                                                                                                                                |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Applying Forwarding Table Filters</i></li></ul>                                                                                     |

## term

---

|                                 |                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>term <i>term-name</i> {<br/>    from {<br/>        <i>match-conditions</i>;<br/>    }<br/>    then {<br/>        <i>action</i>;<br/>        <i>action-modifiers</i>;<br/>    }<br/>}</pre>                                                                                    |
| <b>Hierarchy Level</b>          | [edit <b>firewall family</b> <i>family-name</i> <b>filter</b> <i>filter-name</i> ]                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                      |
| <b>Description</b>              | Define a firewall filter term.                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <p><b><i>term-name</i></b>—Name that identifies the term. The name can contain letters, numbers, and hyphens (-), and can be up to 64 characters long. To include spaces in the name, enclose it in quotation marks.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>firewall—To view this statement in the configuration.</p> <p>firewall-control—To add this statement to the configuration.</p>                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Firewall Filter Match Conditions and Actions</i></li><li>• <a href="#">Configuring Firewall Filters on page 5978</a></li><li>• <a href="#">Overview of Firewall Filters on page 5951</a></li></ul>                                      |



## then (Filters)

---

|                                 |                                                                                                                                                                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> then {     action;     action-modifiers; } </pre>                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit <b>firewall family</b> <i>family-name</i> <b>filter</b> <i>filter-name</i> <b>term</b> <i>term-name</i> ]                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                    |
| <b>Description</b>              | Configure a firewall filter action.                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <p><b>action</b>—Actions to accept, discard, or forward packets that match all conditions specified in a filter term.</p> <p><b>action-modifiers</b>—Additional actions to analyze, classify, count, or police packets that match all conditions specified in a filter term.</p> |
| <b>Required Privilege Level</b> | <p>firewall—To view this statement in the configuration.</p> <p>firewall-control—To add this statement to the configuration.</p>                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Firewall Filter Match Conditions and Actions</i></li> <li>• <a href="#">Configuring Firewall Filters on page 5978</a></li> <li>• <a href="#">Understanding Firewall Filter Match Conditions on page 5957</a></li> </ul>              |



## Configuration Statements (Policers)

- [action](#) on page 6100
- [bandwidth-limit](#) on page 6100
- [burst-size-limit](#) on page 6101
- [color-aware](#) on page 6102
- [color-blind](#) on page 6103
- [committed-burst-size](#) on page 6104
- [committed-information-rate](#) on page 6105
- [excess-burst-size](#) on page 6106
- [filter-specific](#) on page 6107
- [firewall](#) on page 6108
- [if-exceeding](#) on page 6109
- [loss-priority high then discard \(Three-Color Policer\)](#) on page 6110
- [peak-burst-size](#) on page 6111
- [peak-information-rate](#) on page 6112
- [policer](#) on page 6113
- [single-rate](#) on page 6114
- [then \(Policers\)](#) on page 6115
- [three-color-policer](#) on page 6116
- [two-rate](#) on page 6117

## action

---

|                                 |                                                                                                                                                   |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>action {<br/>    loss-priority high then discard;<br/>}</code>                                                                              |
| <b>Hierarchy Level</b>          | [edit <code>firewall three-color-policer name</code> ]                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.     |
| <b>Description</b>              | Discard traffic on a logical interface using tricolor marking policing.                                                                           |
| <b>Options</b>                  | The statements are explained separately.                                                                                                          |
| <b>Required Privilege Level</b> | <code>firewall</code> —To view this statement in the configuration.<br><code>firewall-control</code> —To add this statement to the configuration. |

## bandwidth-limit

---

|                                 |                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>bandwidth-limit bps;</code>                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit <code>firewall policer policer-name if-exceeding</code> ]                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                             |
| <b>Description</b>              | Specify the traffic rate in bits per second.                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <code>bps</code> —Traffic rate in bits per second. Specify <code>bps</code> as a decimal value or as a decimal number followed by one of the abbreviation <code>k</code> (1000), <code>m</code> (1,000,000), or <code>g</code> (1,000,000,000).<br><b>Range:</b> 32000 bps (32 Kbps) through 10,000,000,000 bps (10 Gbps) |
| <b>Required Privilege Level</b> | <code>firewall</code> —To view this statement in the configuration.<br><code>firewall-control</code> —To add this statement to the configuration.                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 6017</a></li><li>• <a href="#">Overview of Policers on page 5999</a></li></ul>                                                                                                        |

---

## burst-size-limit

---

|                                 |                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>burst-size-limit bytes;</code>                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit <code>firewall policer policer-name if-exceeding</code> ]                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                      |
| <b>Description</b>              | Specify the maximum allowed burst size to control the amount of traffic bursting.                                                                                                                                  |
| <b>Options</b>                  | <b>bytes</b> —Decimal value or a decimal number followed by k (thousand), m (million), or g (giga).<br><b>Range:</b> 1 through 2,147,450,880 bytes (2147 MB)                                                       |
| <b>Required Privilege Level</b> | firewall—To view this statement in the configuration.<br>firewall-control—To add this statement to the configuration.                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 6017</a></li><li>• <a href="#">Overview of Policers on page 5999</a></li></ul> |

## color-aware

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | color-aware;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit <a href="#">firewall three-color-policer</a> <i>policer-name</i> single-rate],<br>[edit <a href="#">firewall three-color-policer</a> <i>policer-name</i> two-rate]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Configure the way preclassified packets are metered. In color-aware mode, the switch can assign a higher packet-loss priority, but cannot assign a lower packet loss priority (PLP). For example, suppose an upstream device assigns medium-high PLP to a packet because the packet exceeded its committed information rate (CIR). The switch cannot change the PLP to low even if the packet conforms to the configured CIR of the appropriate interface. On the other hand, if an upstream device assigns low PLP to a packet but the packet exceeds the CIR and committed burst size (CBS) of the switch interface, the switch can increase the PLP to medium-high. |
| <b>Default</b>                  | If you omit the <b>color-aware</b> statement, the default behavior is color-aware mode.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | firewall—To view this statement in the configuration.<br>firewall-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Overview of Policers on page 5999</a></li><li>• <a href="#">Understanding Color-Aware Mode for Single-Rate Tricolor Marking on page 6006</a></li><li>• <a href="#">Understanding Color-Aware Mode for Two-Rate Tricolor Marking on page 6008</a></li><li>• <a href="#">color-blind on page 6103</a></li></ul>                                                                                                                                                                                                                                                                                                      |


## color-blind

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | color-blind;                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit <b>firewall three-color-policer</b> <i>policer-name</i> single-rate],<br>[edit <b>firewall three-color-policer</b> <i>policer-name</i> two-rate]                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Configure the way preclassified packets are metered. In color-blind mode, the switch ignores any preclassification of packets and can assign a higher or lower packet loss priority (PLP). For example, suppose an upstream device assigns medium-high PLP to a packet because the packet exceeded the CIR on the upstream device. The switch can change the PLP to low if the packet conforms to the CIR of the appropriate interface.                                   |
| <b>Default</b>                  | If you omit the <b>color-blind</b> statement, the default behavior is color-aware mode.                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | firewall—To view this statement in the configuration.<br>firewall-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Overview of Policers on page 5999</a></li> <li>• <a href="#">Understanding Color-Blind Mode for Single-Rate Tricolor Marking on page 6005</a></li> <li>• <a href="#">Understanding Color-Blind Mode for Two-Rate Tricolor Marking on page 6008</a></li> <li>• <a href="#">Configuring Color-Blind Egress Policers for Medium-Low PLP on page 6016</a></li> <li>• <a href="#">color-aware on page 6102</a></li> </ul> |


## committed-burst-size

---

|                                                                                                                                                                                                                                                                                                           |                                                                                                                                                                                                                                                                                         |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                             | <code>committed-burst-size bytes;</code>                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                    | [edit <a href="#">firewall three-color-policer</a> <i>policer-name</i> single-rate],<br>[edit <a href="#">firewall three-color-policer</a> <i>policer-name</i> two-rate]                                                                                                                |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                           |
| <b>Description</b>                                                                                                                                                                                                                                                                                        | Configure the maximum number of bytes allowed for incoming traffic to burst above the committed information rate and still be marked with low packet loss priority (green).                                                                                                             |
| <div> <b>NOTE:</b> When you include the <code>committed-burst-size</code> statement in the configuration, you must also include the <code>committed-information-rate</code> statement at the same hierarchy level.</div> |                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                                                                                                                                                                                                                                                                                            | <b>bytes</b> —Number of bytes. You can specify a value in bytes either as a complete decimal number or as a decimal number followed by the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000).<br><b>Range:</b> 512 bytes through 268435456 bytes (268 MB) |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                           | <b>firewall</b> —To view this statement in the configuration.<br><b>firewall-control</b> —To add this statement to the configuration.                                                                                                                                                   |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                              | <ul style="list-style-type: none"><li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 6017</a></li><li>• <a href="#">Overview of Policers on page 5999</a></li></ul>                                                                      |




## committed-information-rate

|                                                                                                                                                                                                                                                                                                                    |                                                                                                                                                                                                                                                                                                                                             |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                      | <code>committed-information-rate <i>bits-per-second</i>;</code>                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                             | [edit <code>firewall three-color-policer <i>policer-name</i> single-rate</code> ],<br>[edit <code>firewall three-color-policer <i>policer-name</i> two-rate</code> ]                                                                                                                                                                        |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                         | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                               |
| <b>Description</b>                                                                                                                                                                                                                                                                                                 | Configure the guaranteed bandwidth under normal line conditions and the average rate up to which packets are marked with low packet loss priority (green).                                                                                                                                                                                  |
| <div>  <p><b>NOTE:</b> When you include the <code>committed-information-rate</code> statement in the configuration, you must also include the <code>committed-burst-size</code> statement at the same hierarchy level.</p> </div> |                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                                                                                                                                                                                                                                                                                                     | <p><b><i>bits-per-second</i></b>—Number of bits per second. You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000).</p> <p><b>Range:</b> 32,000 bps through 10,000,000,000 bps (10 gbps)</p> |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                    | <p><code>firewall</code>—To view this statement in the configuration.</p> <p><code>firewall-control</code>—To add this statement to the configuration.</p>                                                                                                                                                                                  |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                       | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 6017</a></li> <li>• <a href="#">Overview of Policers on page 5999</a></li> </ul>                                                                                                                       |

## excess-burst-size

---

|                                                                                                                                                                                                                                                                                                                                               |                                                                                                                                                                                                                                                                                         |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                 | <code>excess-burst-size bytes;</code>                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                        | [edit <a href="#">firewall three-color-policer</a> <i>policer-name</i> single-rate]                                                                                                                                                                                                     |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                    | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                           |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                            | Configure the maximum number of bytes allowed for incoming traffic to burst above the committed information rate and still be marked with medium-high packet loss priority (yellow). Packets that exceed the excess burst size (EBS) are marked with high packet loss priority (red).   |
| <div> <b>NOTE:</b> When you include the <code>excess-burst-size</code> statement in the configuration, you must also include the <code>committed-burst-size</code> and <code>committed-information-rate</code> statements at the same hierarchy level.</div> |                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                                | <b>bytes</b> —Number of bytes. You can specify a value in bytes either as a complete decimal number or as a decimal number followed by the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000).<br><b>Range:</b> 512 bytes through 268435456 bytes (268 MB) |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                               | <b>firewall</b> —To view this statement in the configuration.<br><b>firewall-control</b> —To add this statement to the configuration.                                                                                                                                                   |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                  | <ul style="list-style-type: none"><li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 6017</a></li><li>• <a href="#">Overview of Policers on page 5999</a></li></ul>                                                                      |

---

## filter-specific

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | filter-specific;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit <b>firewall policer</b> <i>policer-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | <p>Configure a policer to be filter-specific, which means that Junos OS creates only one policer instance regardless of how many times the policer is referenced. If you use a filter-specific policer in multiple terms, both of the following are true:</p> <ul style="list-style-type: none"><li>• Traffic is policed at the aggregate rate. For example, if you create a policer that has a bandwidth limit of 100 Mbps and use the policer in two terms, the total allowed bandwidth for both terms is 100 Mbps—not 100 Mbps for each term.</li><li>• The implicit counter counts all the packets are that matched by any of the terms. For example, if you reference the same filter-specific policer in term1 and term2, and term1 matches 1000 packets and term2 matches 500 packets, the implicit counter shows 1500 matches for the policer.</li></ul> |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 6017</a></li><li>• <a href="#">Overview of Policers on page 5999</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## firewall

```
Syntax firewall {
 family family-name {
 filter filter-name {
 interface-specific;
 term term-name {
 from {
 match-conditions;
 }
 then {
 action;
 action-modifiers;
 }
 }
 }
 }
 policer policer-name {
 filter-specific;
 if-exceeding {
 bandwidth-limit bps;
 burst-size-limit bytes;
 }
 then {
 policer-action;
 }
 }
 three-color-policer policer-name {
 action {
 loss-priority high then discard;
 }
 single-rate {
 (color-aware | color-blind);
 committed-information-rate bps;
 committed-burst-size bytes;
 excess-burst-size bytes;
 }
 two-rate {
 (color-aware | color-blind);
 committed-information-rate bps;
 committed-burst-size bytes;
 peak-information-rate bps;
 peak-burst-size bytes;
 }
 }
 }
```

Hierarchy Level [\[edit\]](#)

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Configure firewall filters and policers.

The remaining statements are explained separately.

|                                 |                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | firewall—To view this statement in the configuration.<br>firewall-control—To add this statement to the configuration.                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Firewall Filter Match Conditions and Actions</a></li> <li>• <a href="#">Configuring Firewall Filters on page 5978</a></li> <li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 6017</a></li> <li>• <a href="#">Overview of Firewall Filters on page 5951</a></li> </ul> |

## if-exceeding


|                                 |                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>if-exceeding {     bandwidth-limit <i>bps</i>;     burst-size-limit <i>bytes</i>; }</pre>                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit <b>firewall policer</b> <i>policer-name</i> ]                                                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                         |
| <b>Description</b>              | Configure policer rate limits.<br><br>The remaining statements are explained separately.                                                                                                                              |
| <b>Required Privilege Level</b> | firewall—To view this statement in the configuration.<br>firewall-control—To add this statement to the configuration.                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 6017</a></li> <li>• <a href="#">Overview of Policers on page 5999</a></li> </ul> |

## loss-priority high then discard (Three-Color Policer)

---


|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | loss-priority high then discard;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit firewall <b>three-color-policer</b> <i>policer-name</i> <b>action</b> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | <p>For packets with high loss priority, discard the packets. The loss priority setting is not configurable. Include this statement if you do not want the switch to forward packets that have high packet-loss priority.</p> <p>For single-rate three-color policers, Junos OS assigns high loss priority to packets that exceed the committed information rate and the excess burst size.</p> <p>For two-rate three-color policers, Junos OS assigns high loss priority to packets that exceed the peak information rate and the peak burst size.</p> |
| <b>Required Privilege Level</b> | firewall—To view this statement in the configuration.<br>firewall-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 6017</a></li><li>• <a href="#">Overview of Policers on page 5999</a></li></ul>                                                                                                                                                                                                                                                                                                                                     |

## peak-burst-size

|                                                                                                                                                                                                                                                                                                                                          |                                                                                                                                                                                                                                                                                                |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                            | <code>peak-burst-size bytes;</code>                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                   | [edit <a href="#">firewall three-color-policer</a> <i>policer-name</i> two-rate]                                                                                                                                                                                                               |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                               | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                  |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                       | Configure the maximum number of bytes allowed for incoming packets to burst above the peak information rate (PIR) and still be marked with medium-high packet loss priority (yellow). Packets that exceed the peak burst size (PBS) are marked with high packet loss priority (red).           |
| <div>  <b>NOTE:</b> When you include the <code>peak-burst-size</code> statement in the configuration, you must also include the <code>committed-burst-size</code> and <code>peak-information-rate</code> statements at the same hierarchy level. </div> |                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                           | <b>bytes</b> —Number of bytes. You can specify a value in bytes either as a complete decimal number or as a decimal number followed by the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000).<br><b>Range:</b> 1500 bytes through 100,000,000,000 bytes (100 GB) |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                          | <b>firewall</b> —To view this statement in the configuration.<br><b>firewall-control</b> —To add this statement to the configuration.                                                                                                                                                          |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                             | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 6017</a></li> <li>• <a href="#">Overview of Policers on page 5999</a></li> </ul>                                                                          |

## peak-information-rate

---

|                                                                                                                                                                                                                                                                                                                                              |                                                                                                                                                                                                                                                                                                                                                           |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                | <code>peak-information-rate <i>bits-per-second</i>;</code>                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                       | [edit <a href="#">firewall three-color-policer</a> <i>policer-name</i> two-rate]                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                   | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                             |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                           | Configure the maximum achievable rate. Packets that exceed the committed information rate (CIR) but are below the peak information rate (PIR) are marked with medium-high packet loss priority (yellow). Packets that exceed the PIR are marked with high packet loss priority (red). You can configure a discard action for packets that exceed the PIR. |
| <div> <b>NOTE:</b> When you include the <code>peak-information-rate</code> statement in the configuration, you must also include the <code>committed-information-rate</code> and <code>peak-burst-size</code> statements at the same hierarchy level.</div> |                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                               | <b><i>bits-per-second</i></b> —Number of bits per second. You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000).<br><b>Range:</b> 32,000 bps through 10,000,000,000 bps (10 gbps)                         |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                              | <b>firewall</b> —To view this statement in the configuration.<br><b>firewall-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                     |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                 | <ul style="list-style-type: none"><li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 6017</a></li><li>• <a href="#">Overview of Policers on page 5999</a></li></ul>                                                                                                                                        |



## policer

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> policer <i>policer-name</i> {     filter-specific;     if-exceeding {         bandwidth-limit <i>bps</i>;         burst-size-limit <i>bytes</i>;     }     then {         <i>policer-action</i>;     } } </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit <a href="#">firewall</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | <p>Configure policer rate limits and actions. To activate a policer, you must include the <b>policer</b> action modifier in the <b>then</b> statement in a firewall filter term.</p> <p>Each policer that you configure includes an implicit counter that counts the number of packets that exceed the rate limits that are specified for the policer. If you use the same policer in multiple terms—either within the same filter or across filters—the policer’s implicit counter is used to count packets that are policed in all of these terms. If you want to obtain separate packet counts for each term, use these approaches:</p> <ul style="list-style-type: none"> <li>• Configure a unique policer for each term.</li> <li>• Configure only one policer, but use a unique, explicit counter in each term.</li> </ul> |
| <b>Options</b>                  | <p><b><i>policer-name</i></b>—Name that identifies the policer. The name can contain letters, numbers, hyphens (-), and can be up to 64 characters long.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | <p>firewall—To view this statement in the configuration.</p> <p>firewall-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 6017</a></li> <li>• <a href="#">Configuring Firewall Filters on page 5978</a></li> <li>• <a href="#">Overview of Policers on page 5999</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |


## single-rate

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>single-rate {<br/>  (color-aware   color-blind);<br/>  committed-information-rate <i>bps</i>;<br/>  committed-burst-size <i>bytes</i>;<br/>  excess-burst-size <i>bytes</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit <a href="#">firewall three-color-policer</a> <i>policer-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | <p>Configure a single-rate three-color policer in which marking is based on the committed information rate (CIR), committed burst size (CBS), and excess burst size (EBS).</p> <p>Packets that conform to the CIR or the CBS are assigned low loss priority (green). Packets that exceed the CIR and the CBS but are within the EBS are assigned medium-high loss priority (yellow). Packets that exceed the EBS are assigned high loss priority (red).</p> <p>Green and yellow packets are always forwarded; this action is not configurable. You can configure red packets to be discarded. By default, red packets are forwarded.</p> <p>The remaining statements are explained separately.</p> |
| <b>Options</b>                  | <b><i>policer-name</i></b> —Name of the three-color policer. Use this name when you apply the policer to an interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | <b>firewall</b> —To view this statement in the configuration.<br><b>firewall-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 6017</a></li><li>• <a href="#">Overview of Policers on page 5999</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## then (Policers)

---

|                                                                                                                                                                                                                         |                                                                                                                                                                                                                                                                                                                                   |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                           | then {<br><i>policer-action</i> ;<br>}                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>                                                                                                                                                                                                  | [edit <b>firewall</b> <b>policer</b> <i>policer-name</i> ]                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>                                                                                                                                                                                              | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                     |
| <b>Description</b>                                                                                                                                                                                                      | Configure a policer action.                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                                                                                                                                                                                                          | <i>policer-action</i> —Allowed policer actions are <b>discard</b> , <b>loss-priority high</b> , and <b>loss-priority low</b> . <b>discard</b> causes the system to drop traffic that exceeds the rate limits defined by the policer. Use <b>loss-priority high</b> to allow the system to forward matching traffic in some cases. |
| <div>  <b>NOTE:</b> If you specify a policer in an egress firewall filter, the only supported action is <b>discard</b>.         </div> |                                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b>                                                                                                                                                                                         | firewall—To view this statement in the configuration.<br>firewall-control—To add this statement to the configuration.                                                                                                                                                                                                             |
| <b>Related Documentation</b>                                                                                                                                                                                            | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 6017</a></li> <li>• <a href="#">Configuring Firewall Filters on page 5978</a></li> <li>• <a href="#">Overview of Policers on page 5999</a></li> </ul>                                        |

## three-color-policer

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>three-color-policer <i>policer-name</i> {<br/>  action {<br/>    loss-priority high then discard;<br/>  }<br/>  single-rate {<br/>    (color-aware   color-blind);<br/>    committed-information-rate <i>bps</i>;<br/>    committed-burst-size <i>bytes</i>;<br/>    excess-burst-size <i>bytes</i>;<br/>  }<br/>  two-rate {<br/>    (color-aware   color-blind);<br/>    committed-information-rate <i>bps</i>;<br/>    committed-burst-size <i>bytes</i>;<br/>    peak-information-rate <i>bps</i>;<br/>    peak-burst-size <i>bytes</i>;<br/>  }<br/>}</pre> |
| <b>Hierarchy Level</b>          | [edit <a href="#">firewall</a> ],<br>[edit logical-systems <i>logical-system-name</i> firewall]                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Configure a three-color policer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <p><i>policer-name</i>—Name of the three-color policer. Use this name when you apply the policer to an interface.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | firewall—To view this statement in the configuration.<br>firewall-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 6017</a></li><li>• <a href="#">Overview of Policers on page 5999</a></li></ul>                                                                                                                                                                                                                                                                                                                                                    |

## two-rate

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>two-rate {   (color-aware   color-blind);   committed-information-rate <i>bps</i>;   committed-burst-size <i>bytes</i>;   peak-information-rate <i>bps</i>;   peak-burst-size <i>bytes</i>; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit <b>firewall three-color-policer</b> <i>policer-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | <p>Configure a two-rate three-color policer in which marking is based on the committed information rate (CIR), committed burst size (CBS), peak information rate (PIR), and peak burst size (PBS).</p> <p>Packets that conform to the CIR or the CBS are assigned low loss priority (green). Packets that exceed the CIR and the CBS but are within the PIR or the PBS are assigned medium-high loss priority (yellow). Packets that exceed the PIR and the PBS are assigned high loss priority (red).</p> <p>Green and yellow packets are always forwarded; this action is not configurable. You can configure red packets to be discarded. By default, red packets are forwarded.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>firewall—To view this statement in the configuration.</p> <p>firewall-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |



## CHAPTER 230

# Configuration Statements (Device Security)

- [action-shutdown on page 6120](#)
- [bandwidth-level on page 6121](#)
- [bandwidth-percentage on page 6122](#)
- [interface \(Unknown Unicast Forwarding\) on page 6123](#)
- [no-broadcast on page 6124](#)
- [no-multicast on page 6125](#)
- [no-unknown-unicast on page 6126](#)
- [rpf-check on page 6127](#)
- [storm-control on page 6128](#)
- [storm-control-profiles on page 6129](#)
- [unknown-unicast-forwarding on page 6130](#)


## action-shutdown

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | action-shutdown;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | For platforms without ELS:<br><br>[edit ethernet-switching-options <a href="#">storm-control</a> ]<br><br>For platforms with ELS:<br><br>[edit forwarding-options <a href="#">storm-control-profiles</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | <p>Shut down or disable interfaces when the storm control level is exceeded, as follows:</p> <ul style="list-style-type: none"><li>• If you set both the <b>action-shutdown</b> and the <b>port-error-disable</b> statements, the affected interfaces are disabled temporarily and recover automatically when the disable timeout expires.</li><li>• If you set the <b>action-shutdown</b> statement and do not set the <b>port-error-disable</b> statement, the affected interfaces are shut down when the storm control level is exceeded, and they do not recover automatically. You must issue the <b>clear ethernet-switching port-error</b> command to clear the port error and restore the interfaces to service.</li></ul> |
| <b>Default</b>                  | The <b>action-shutdown</b> feature is disabled. If the storm control level is exceeded, the switch drops broadcast and unknown unicast messages on the specified interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Storm Control on page 6045</a></li><li>• <i>Example: Configuring Storm Control to Prevent Network Outages</i></li><li>• <a href="#">port-error-disable on page 6140</a></li><li>• <a href="#">clear ethernet-switching port-error on page 6178</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                     |




## bandwidth-level

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <code>bandwidth-level <i>kbps</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | [edit forwarding-options <a href="#">storm-control-profiles</a> <i>profile-name</i> all]                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1 for MX Series routers.                                                                                                                                                                                                                                        |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Configure the storm control level as the bandwidth in kilobits per second of the available bandwidth used by the combined broadcast, multicast, and unknown unicast traffic streams.                                                                                                                                                                                                                                                                             |
| <div>  <p><b>NOTE:</b> When you configure storm control level on an aggregated Ethernet interface, the storm control level for each member of the aggregated Ethernet interface is set to that bandwidth. For example, if you configure a storm control level of 15,000 Kbps on ae1, and ae1 has two members, ge-0/0/0 and ge-0/0/1, each member has a storm control level of 15,000 Kbps. Thus, the storm control level on ae1 allows a traffic rate of up to 30,000 Kbps of combined broadcast, multicast, and unknown unicast traffic.</p> </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Default</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <p>On EX4300 switches—If you do not specify the storm control level using either the <b>bandwidth-level</b> or the <b>bandwidth-percentage</b> statements, the storm control level defaults to 80 percent of the available bandwidth used by the combined broadcast, unknown unicast, and multicast traffic streams.</p> <p>On EX9200 switches—Storm control is not enabled by default.</p> <p>On MX Series routers—Storm control is not enabled by default.</p> |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <p><b>bandwidth-level <i>kbps</i></b>—Traffic rate in kilobits per second of the combined broadcast, multicast, and unknown unicast traffic streams.</p> <p><b>Range:</b> 100 through 10,000,000</p> <p><b>Range:</b> 100 through 100,000,000 on QFX10000 Series switches</p> <p><b>Default:</b> None</p>                                                                                                                                                        |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <ul style="list-style-type: none"> <li>• <a href="#">bandwidth-percentage on page 6122</a></li> <li>• <i>Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches</i></li> <li>• <i>Example: Configuring Storm Control to Prevent Network Outages on MX Series Routers</i></li> <li>• <i>Configuring or Disabling Storm Control (CLI Procedure)</i></li> </ul>                                                                        |

## bandwidth-percentage

---

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |                                                                                                                                                                                                                                                                                                                                                                                 |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <code>bandwidth-percentage percentage;</code>                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | [edit forwarding-options <a href="#">storm-control-profiles</a> <i>profile-name</i> all]                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for the QFX series.<br>Statement introduced in Junos OS Release 14.1 for MX Series routers.                                                                                                                                                       |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Configure the storm control level as the percentage of available bandwidth used by the combined broadcast, multicast, and unknown unicast traffic streams on an interface. The storm control level is configured as part of the storm control profile.                                                                                                                          |
| <div> <b>NOTE:</b> When you configure storm control level on an aggregated Ethernet interface, the storm control level for each member of the aggregated Ethernet interface is set to that bandwidth. For example, if you configure a storm control level of 15,000 Kbps on ae1, and ae1 has two members, ge-0/0/0 and ge-0/0/1, each member has a storm control level of 15,000 Kbps. Thus, the storm control level on ae1 allows a traffic rate of up to 30,000 Kbps of combined broadcast, multicast, and unknown unicast traffic.</div> |                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Default</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | On EX4300 switches—The storm control level is 80 percent of the available bandwidth used by the combined broadcast, unknown unicast, and multicast traffic streams.<br><br>On EX9200 switches—Storm control is not enabled by default.<br><br>On MX Series routers—Storm control is not enabled by default.                                                                     |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <ul style="list-style-type: none"><li>• <a href="#">bandwidth-level on page 6121</a></li><li>• <i>Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches</i></li><li>• <i>Example: Configuring Storm Control to Prevent Network Outages on MX Series Routers</i></li><li>• <i>Configuring or Disabling Storm Control (CLI Procedure)</i></li></ul> |

## interface (Unknown Unicast Forwarding)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>interface <i>interface-name</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | <ul style="list-style-type: none"> <li>For platforms with ELS:<br/>[edit switch-options <b>unknown-unicast-forwarding</b> vlan <i>vlan-name</i>]</li> <li>For platforms without ELS:<br/>[edit ethernet-switching-options <b>unknown-unicast-forwarding</b> vlan <i>vlan-name</i>]</li> </ul>                                                                                                                                                     |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> <p>Hierarchy level <b>[edit switch-options]</b> introduced in Junos OS Release 13.2X50-D10. (See <a href="#">“Getting Started with Enhanced Layer 2 Software” on page 41</a> for information about ELS.)</p> |
| <b>Description</b>              | Specify the interface to which unknown unicast packets will be forwarded.                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><code>show vlans</code></li> <li><code>show ethernet-switching table</code></li> <li><a href="#">Configuring Unknown Unicast Forwarding (CLI Procedure)</a></li> <li><a href="#">Understanding Unknown Unicast Forwarding on page 6059</a></li> </ul>                                                                                                                                                      |

## no-broadcast

---

|                                 |                                                                                                                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-broadcast;                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | For platforms without ELS:<br><br>[edit ethernet-switching-options <b>storm-control</b> interface (all   <i>interface-name</i> )]<br><br>For platforms with ELS:<br><br>[edit forwarding-options <b>storm-control-profiles</b> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                    |
| <b>Description</b>              | For interfaces configured for storm control, disable broadcast traffic storm control on the interface.                                                                                                                           |
| <b>Default</b>                  | When storm control is enabled on an interface, it is enabled for broadcast traffic (as well as multicast and unknown unicast traffic).                                                                                           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Storm Control on page 6045</a></li><li>• <i>Example: Configuring Storm Control to Prevent Network Outages</i></li></ul>                                        |

---

## no-multicast

---

|                                 |                                                                                                                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-multicast;                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | For platforms without ELS:<br><br>[edit ethernet-switching-options <b>storm-control</b> interface (all   <i>interface-name</i> )]<br><br>For platforms with ELS:<br><br>[edit forwarding-options <b>storm-control-profiles</b> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                    |
| <b>Description</b>              | Disable storm control for all multicast traffic (both registered multicast and unregistered multicast) for the specified interface or for all interfaces.                                                                        |
| <b>Default</b>                  | Storm control is enabled for unknown unicast traffic, multicast traffic, and broadcast traffic.                                                                                                                                  |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Storm Control on page 6045</a></li><li>• <i>Example: Configuring Storm Control to Prevent Network Outages</i></li></ul>                                        |

## no-unknown-unicast

---

|                                 |                                                                                                                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-unknown-unicast;                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | For platforms without ELS:<br><br>[edit ethernet-switching-options <b>storm-control</b> interface (all   <i>interface-name</i> )]<br><br>For platforms with ELS:<br><br>[edit forwarding-options <b>storm-control-profiles</b> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                    |
| <b>Description</b>              | For interfaces configured for storm control, disable unknown unicast traffic storm control on the interface.                                                                                                                     |
| <b>Default</b>                  | When storm control is enabled on an interface, it is enabled for both unknown unicast traffic and broadcast traffic.                                                                                                             |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Storm Control on page 6045</a></li><li>• <i>Example: Configuring Storm Control to Prevent Network Outages</i></li></ul>                                        |


## rpf-check

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>rpf-check;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet],<br>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6]                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.3 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | <p>On EX3200 and EX4200 switches, enable a reverse-path forwarding (RPF) check on unicast traffic (except ECMP packets) on all ingress interfaces.</p> <p>On EX4300 switches, enable a reverse-path forwarding (RPF) check on unicast traffic, including ECMP packets, on all ingress interfaces.</p> <p>On EX8200 and EX6200 switches, enable an RPF check on unicast traffic, including ECMP packets, on the selected ingress interfaces.</p> <p>On QFX Series switches, enable an RPF check on unicast traffic (except ECMP packets) on the selected ingress interfaces.</p> |
| <b>Default</b>                  | Unicast RPF is disabled on all interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring Unicast RPF on an EX Series Switch</i></li> <li>• <a href="#">Configuring Unicast RPF (CLI Procedure) on page 6054</a></li> <li>• <a href="#">Disabling Unicast RPF (CLI Procedure) on page 6055</a></li> <li>• <a href="#">Understanding Unicast RPF on page 6050</a></li> </ul>                                                                                                                                                                                                                              |

## storm-control

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>storm-control storm-control-profile;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> unit <i>number</i> family ethernet-switching],<br>[edit interfaces <i>interface-name</i> unit <i>number</i> family bridge]<br>[edit interfaces <i>interface-name</i> ether-options ethernet-switch-profile]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for the QFX series.<br>Statement introduced in Junos OS Release 14.1 for the MX Series routers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | <p>Bind a storm control profile to a logical interface.</p> <p>On switches running ELS software, storm control is enabled by default on all switch interfaces at a level of 80 percent of the combined broadcast and unknown unicast streams. (For the equivalent statement for platforms running non-ELS software, see <i>storm-control</i>.)</p> <div> <b>NOTE:</b> If you configure storm control on an aggregated Ethernet interface, the storm-control level is applied to each member interface individually. For example, if the aggregated interface has two members and you configure a storm-control level of 20 kbps, Junos will not detect a storm if one or both of the member interfaces receives traffic at 15 kbps because in neither of these cases does an individual member receive traffic at a rate greater than the configured storm-control level. In this example, Junos detects a storm only if at least one member interface receives traffic at greater than 20 Kbps.</div> |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches</i></li><li>• <i>Understanding Storm Control for Managing Traffic Levels on Switching Devices</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |



## storm-control-profiles

**Syntax** `storm-control-profiles profile-name {  
     action-shutdown;  
     all {  
         bandwidth-level;  
         bandwidth-percentage;  
         no-broadcast;  
         no-multicast;  
         no-registered-multicast;  
         no-unknown-unicast;  
         no-unregistered-multicast;  
     }  
 }`

**Hierarchy Level** [edit forwarding-options]

**Release Information** Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.  
 Statement introduced in Junos OS Release 13.2 for the QFX Series.  
 Statement introduced in Junos OS Release 14.1 for MX Series routers.

**Description** Configure a storm control profile on a switch or router. Storm control is used to prevent network outages that are caused by broadcast traffic storms. Storm control enables the switching device to monitor traffic levels and to drop broadcast, multicast, and unknown unicast packets when a specified traffic level—called the storm control level or storm control bandwidth—is exceeded, thus preventing packets from proliferating and degrading the LAN.



**NOTE:** The name of the storm control profile can contain no more than 127 characters.

The remaining statements are explained separately.


**Required Privilege Level** system—To view this statement in the configuration.  
 system-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches*
- *Understanding Storm Control for Managing Traffic Levels on Switching Devices*

## unknown-unicast-forwarding

---


|                                                                                                                                                                                                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                      | <pre>unknown-unicast-forwarding {<br/>  vlan (Unknown Unicast Forwarding) (all   <i>vlan-name</i>){<br/>    interface (Unknown Unicast Forwarding) <i>interface-name</i>;<br/>  }<br/>}</pre>                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>                                                                                                                                                                                             | <ul style="list-style-type: none"><li>For platforms with ELS:<br/>[edit switch-options]</li><li>For platforms without ELS:<br/>[edit ethernet-switching-options]</li></ul>                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>                                                                                                                                                                                         | <p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX Series.</p> <p>Hierarchy level [edit switch-options] introduced in Junos OS Release 13.2X50-D10. (See <a href="#">“Getting Started with Enhanced Layer 2 Software” on page 41</a> for information about ELS.)</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| <b>Description</b>                                                                                                                                                                                                 | Configure the switch to forward all unknown unicast packets in a VLAN or on all VLANs to a particular interface.                                                                                                                                                                                                                                                                                                                           |
| <div> <b>NOTE:</b> Before you can configure unknown unicast forwarding within a VLAN, you must first configure that VLAN.</div> |                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| The remaining statements are explained separately.                                                                                                                                                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Default</b>                                                                                                                                                                                                     | Unknown unicast packets are flooded to all interfaces that belong to the same VLAN.                                                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b>                                                                                                                                                                                    | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>                                                                                                                                                                                       | <ul style="list-style-type: none"><li><i>show vlans</i></li><li><i>show ethernet-switching table</i></li><li><i>Configuring Unknown Unicast Forwarding (CLI Procedure)</i></li><li><a href="#">Understanding Unknown Unicast Forwarding on page 6059</a></li></ul>                                                                                                                                                                         |

## CHAPTER 231

# Configuration Statements (Port Security)

- [circuit-id on page 6132](#)
- [fc-map on page 6134](#)
- [fcoe-trusted on page 6136](#)
- [mac-move-limit on page 6137](#)
- [no-allowed-mac-log on page 6138](#)
- [no-gratuitous-arp-request on page 6139](#)
- [persistent-learning on page 6139](#)
- [port-error-disable on page 6140](#)
- [vendor-id on page 6142](#)
- [write-interval on page 6143](#)

## circuit-id

|                                                                                                                                                                                                                                                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                          | <pre> circuit-id {   prefix {     host-name;     logical-system-name;     routing-instance-name;   }   use-interface-description (device   logical);   use-vlan-id; } </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                 | <ul style="list-style-type: none"> <li>For platforms with Enhanced Layer 2 Software (ELS):<br/>[edit vlans <i>vlan-name</i> forwarding-options dhcp-security option-82 ]</li> <li>For platforms without ELS:<br/>[edit ethernet-switching-options secure-access-port vlan (all   <i>vlan-name</i>) dhcp-option82],<br/>[edit forwarding-options helpers bootp dhcp-option82] ,<br/>[edit forwarding-options helpers bootp interface <i>interface-name</i> dhcp-option82]</li> <li>For MX Series platforms:<br/>[edit bridge-domains <i>bridge-domain-name</i> forwarding-options dhcp-security option-82]</li> </ul> |
| <b>Release Information</b>                                                                                                                                                                                                                             | <p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Hierarchy level [edit vlans <i>vlan-name</i> forwarding-options dhcp-security] introduced in Junos OS Release 13.2X50-D10. (See <a href="#">“Getting Started with Enhanced Layer 2 Software” on page 41</a> for information about ELS.)</p> <p>Statement introduced in Junos OS Release 14.1 for the MX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                          |
| <b>Description</b>                                                                                                                                                                                                                                     | <p>Configure the <b>circuit-id</b> suboption (suboption 1) of DHCP option 82 (the DHCP relay agent information option) in DHCP packets destined for a DHCP server. This suboption identifies the circuit (the interface, the VLAN, or both) on which the DHCP request arrived.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                         |
| <b>Default</b>                                                                                                                                                                                                                                         | <p>If DHCP option 82 is enabled on the switch, the circuit ID is supplied by default in the format <i>interface-name:vlan-name</i> or, on a Layer 3 interface, just <i>interface-name</i>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <div>  <p><b>NOTE:</b> When you configure <b>circuit-id</b>, <b>remote-id</b> is also enabled, even if you do not explicitly configure <b>remote-id</b>.</p> </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                        | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

**Related  
Documentation**

- *Configuring DHCP Option 82 to Help Protect the Switching Devices Against Attacks (CLI Procedure)*
- *Example: Setting Up DHCP Option 82 with a Switch with No Relay Agent Between Clients and a DHCP Server*
- *Example: Setting Up DHCP Option 82 with a Switch as a Relay Agent Between Clients and a DHCP Server*
- *Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)*
- *Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure)*
- *Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)*
- RFC 3046, DHCP Relay Agent Information Option, at <http://tools.ietf.org/html/rfc3046>

## fc-map

---

**Syntax** `fc-map fc-map-value;`

**Hierarchy Level** Original CLI

[edit ethernet-switching options secure-access-port vlan (all | *vlan-name*) examine-fip]

ELS CLI for Platforms that Support FCoE

[edit **v**lans *vlan-name* **f**orwarding-**o**ptions fip-security]



**NOTE:** The `fc-map` configuration statement is in a different hierarchy on the original CLI than on the Enhanced Layer 2 Software (ELS) CLI.

QFX Series that Support FCoE-FC Gateway Configuration

[edit fc-fabrics *fc-fabric-name* protocols fip]

**Release Information** Statement introduced in Junos OS Release 10.4 for EX Series switches.  
Statement introduced in Junos OS Release 11.1 for the QFX Series.  
Statement introduced for the ELS CLI in Junos OS Release 13.2 for the QFX Series.

**Description** Set the FCoE mapped address prefix (FC-MAP) value for the FCoE VLAN to match the FC switch (or FCoE forwarder) FC-MAP value for the FC fabric. The FC-MAP value is a unique MAC address prefix an FC switch uses to identify FCoE traffic for a given FC fabric (traffic on a particular FCoE VLAN).

You can configure the FC-MAP value or use the default value. The default FC-MAP value is different for VN\_Port to VF\_Port (VN2VF\_Port) FIP snooping (0x0EFC00) than for VN\_Port to VN\_Port (VN2VN\_Port) FIP snooping.

The FC switch provides the FC-MAP value to FCoE nodes (ENodes) in the FIP discovery advertisement message. If the EX Series switch or the QFX Series FCoE VLAN FC-MAP value does not match the FC switch FC-MAP value, neither device discovers the FC switch on that VLAN, and the ENodes on that VLAN cannot access the FC switch. The FC switch accepts only FCoE traffic that uses the correct FC-MAP value as part of the VN\_Port MAC address.

When the QFX Series acts as an FCoE-FC gateway, the FC-MAP value for the gateway and the FCoE devices must match the FC switch FC-MAP value in order to communicate with the FC switch.



**NOTE:** Changing the FC-MAP value causes all logins to drop and forces the ENodes to log in again.

**Options** `fc-map-value`—FC-MAP value, hexadecimal value preceded by “0x”.

**Range:** 0x0EFC00 through 0x0EFCFF

**Default:** 0x0EFC00 for VN2VF\_Port FIP snooping 0x0EFD00 for VN2VN\_Port FIP snooping

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *examine-fip*
- *show fip snooping*
- *Example: Configuring an FCoE Transit Switch*
- *Configuring VN2VF\_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch*

## fcoe-trusted

**Syntax** `fcoe-trusted;`

**Hierarchy Level** Original CLI

[edit ethernet-switching-options secure-access-port interface *interface-name*]

ELS CLI for Platforms that Support FCoE

[edit **vlan** *vlan-name* **forwarding-options** fip-security interface *interface-name*]



**NOTE:** The `fcoe-trusted` configuration statement is in a different hierarchy on the original CLI than on the Enhanced Layer 2 Software (ELS) CLI.

QFX Series that Support FCoE-FC Gateway Configuration

[edit fc-fabrics *fc-fabric-name* protocols fip]

**Release Information** Statement introduced in Junos OS Release 10.4 for EX Series switches.  
Statement introduced in Junos OS Release 11.1 for the QFX Series.  
Statement introduced for the FC fabric in Junos OS Release 11.3 for the QFX Series.  
Statement introduced for the ELS CLI in Junos OS Release 13.2 for the QFX Series.

**Description** Configure the specified 10-Gigabit Ethernet interface to trust Fibre Channel over Ethernet (FCoE) traffic. If an interface is connected to another switch such as an FCoE forwarder (FCF) or a transit switch, you can configure the interface as trusted so that the interface forwards FCoE traffic from the switch to the FCoE devices without installing FIP snooping filters.

(QFX Series FCoE-FC gateway) Configure the specified local Fibre Channel fabric to trust FCoE traffic on all ports in the fabric. Changing the fabric ports from untrusted to trusted removes any existing FIP snooping filters from the ports. Changing the fabric ports from trusted to untrusted by removing the **fcoe-trusted** configuration from the fabric forces all of the FCoE sessions on those ports to log out so that when the ENodes and VN\_Ports log in again, the switch can build the appropriate FIP snooping filters.


**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *show fip snooping*
- *Example: Configuring an FCoE Transit Switch*
- *Configuring VN2VF\_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch*
- *Configuring VN2VF\_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch*



## mac-move-limit

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>mac-move-limit <i>limit</i> &lt;fabric-limit <i>limit</i>&gt; action <i>action</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | <p>For platforms without ELS:</p> <pre>[edit ethernet-switching-options secure-access-port (all   <i>vlan-name</i>)]</pre> <p>For platforms with ELS:</p> <pre>[edit vlans <i>vlan-name</i> switch-options],</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Specify the number of times a MAC address can move to a new interface (port) in 1 second and the action to be taken by the switch if the MAC address move limit is exceeded.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|                                 | <div>  <p><b>CAUTION:</b> Mac move limiting does not work properly on a QFX5100 switch used as a Node device in a QFabric system. Do not use this feature on a QFX5100 switch in a QFabric system.</p> </div>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Default</b>                  | The default move limit is unlimited. The default action is <b>drop</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <p><b>fabric-limit</b>—Specify the maximum number of moves in a QFabric system. If you do not specify a fabric limit, the value for <b>mac-move-limit</b> applies to the QFabric system.</p> <p><b>limit</b>—Maximum number of moves to a new interface per second.</p> <p><b>action <i>action</i></b>—(Optional) Action to take when the MAC address move limit is reached:</p> <ul style="list-style-type: none"> <li>• <b>drop</b>—Drop the packet and generate an alarm, an SNMP trap, or a system log entry. This is the default.</li> <li>• <b>log</b>—Do not drop the packet but generate an alarm, an SNMP trap, or a system log entry.</li> <li>• <b>none</b>—No action.</li> <li>• <b>shutdown</b>—Logically disable the interface and generate a system log entry. If you have configured the switch with the <b>port-error-disable</b> statement, the disabled interfaces recover automatically upon expiration of the specified disable timeout. If you have not configured the switch for autorecovery from port error disabled conditions, you can bring up the disabled interfaces by running the <b>clear-ethernet-switch-port</b> command.</li> </ul> |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system—control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

- Related Documentation**
- *Example: Configuring Basic Port Security Features*
  - *Configuring MAC Move Limiting (CLI Procedure)*
  - *Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure)*

---

## no-allowed-mac-log

---

- Syntax** no-allowed-mac-log;
- Hierarchy Level**
- For platforms without ELS:  
[edit ethernet-switching-options secure-access-port interface (all | *interface-name*)]
  - For platforms with ELS:  
[edit switch-options interface *interface-name*]
- Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.
- Description** Specify that the switch should not log messages when it receives packets from invalid MAC addresses on an interface that has been configured for allowed MAC addresses.
- Default** The switch logs messages when it receives packets from invalid MAC addresses on an interface that has been configured for particular allowed (specific) MAC addresses.
- Required Privilege Level**
- routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.
- Related Documentation**
- [Understanding MAC Limiting and MAC Move Limiting for Port Security on page 6029](#)
  - *Configuring MAC Limiting*
  - *mac-limit*

## no-gratuitous-arp-request


|                                 |                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-gratuitous-arp-request;                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces</a> <i>interface-name</i> ],<br>[edit <a href="#">interfaces</a> <i>interface-range</i> <i>interface-name</i> ]                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                     |
| <b>Description</b>              | Configure the switch not to respond to gratuitous ARP requests. You can disable responses to gratuitous ARP requests on both Layer 2 Ethernet switching interfaces and routed VLAN interfaces (RVIs). |
| <b>Default</b>                  | Gratuitous ARP responses are enabled on all Ethernet switching interfaces and RVIs.                                                                                                                   |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring IRB Interfaces on page 2131</a></li> </ul>                                                                                           |

## persistent-learning

|                                 |                                                                                                                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | persistent-learning;                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | <ul style="list-style-type: none"> <li>• For platforms without ELS:<br/>[edit ethernet-switching-options secure-access-port interface (all   <i>interface-name</i>)]</li> <li>• For platforms with ELS:<br/>[edit switch-options interface <i>interface-name</i>]</li> </ul>   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.4 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Hierarchy level [edit switch-options interface <i>interface-name</i> ] introduced in Junos OS Release 13.2X50-D10                |
| <b>Description</b>              | Specify that learned MAC addresses persist on the specified interfaces across restarts of the switch and link-down conditions. This feature is also known as sticky MAC.                                                                                                       |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring Basic Port Security Features</i></li> <li>• <i>Configuring Persistent MAC Learning (CLI Procedure)</i></li> <li>• <a href="#">Configuring Persistent MAC Learning (CLI Procedure) on page 6035</a></li> </ul> |

## port-error-disable

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>port-error-disable {<br/>  (disable-timeout <i>seconds</i>   recovery-timeout <i>seconds</i>);<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | <ul style="list-style-type: none"><li>For platforms without ELS:<br/>[edit ethernet-switching-options]</li><li>For platforms with ELS:<br/>[edit switch-options ]</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 on the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Disable rather than block an interface when enforcing MAC limiting, MAC move limiting, and storm control, and allow the interface to recover automatically from the error condition after a specified period of time:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|                                 | <div> <b>NOTE:</b> The <b>port-error-disable</b> configuration does not apply to preexisting error conditions. It affects only error conditions that are detected after you enable and commit the <b>port-error-disable</b> statement. To clear a preexisting error condition and restore the interface to service, use the <a href="#">clear ethernet-switching port-error</a> command.</div>                                                                                                                                                                                                                                                                                                                                                                                                        |
|                                 | <ul style="list-style-type: none"><li>If you enable the <i>mac-limit</i> statement with the <b>shutdown</b> option and also enable the <b>port-error-disable</b> statement, the switch disables (rather than shuts down) the interface when the MAC address limit is reached.</li><li>If you have enabled the <a href="#">mac-move-limit</a> statement with the <b>shutdown</b> option and you enable the <b>port-error-disable</b> statement, the switch disables (rather than shuts down) the interface when the maximum number of moves to a new interface is reached.</li><li>If you enable the <a href="#">storm-control</a> statement with the <b>action-shutdown</b> option and you also enable <b>port-error-disable</b>, the switch disables (rather than shuts down) the interface when broadcast traffic and unknown unicast traffic exceed the specified levels.</li></ul> |
| <b>Default</b>                  | Not enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li><a href="#">Understanding MAC Limiting and MAC Move Limiting for Port Security on page 6029</a></li><li><a href="#">Understanding Storm Control on page 6045</a></li><li><a href="#">Example: Configuring Storm Control to Prevent Network Outages</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

- *Configuring Autorecovery for MAC Limited or Storm Control Interfaces (CLI Procedure)*
- [action-shutdown on page 6120](#)
- *disable-timeout*
- [clear ethernet-switching port-error on page 6178](#)

## vendor-id

|                                                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                             | <code>vendor-id &lt;string&gt;;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>For Platforms with Enhanced Layer 2 Software (ELS)</b> | <code>[edit vlans <i>vlan-name</i> forwarding-options dhcp-security option-82]</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>For Platforms Without ELS</b>                          | <code>[edit ethernet-switching-options secure-access-port vlan (all   <i>vlan-name</i>) dhcp-option82],</code><br><code>[edit forwarding-options helpers bootp dhcp-option82],</code><br><code>[edit forwarding-options helpers bootp interface <i>interface-name</i> dhcp-option82]</code>                                                                                                                                                                                                                                                                                                                                                           |
| <b>For MX Series Platforms</b>                            | <code>[edit bridge-domains <i>bridge-domain-name</i> forwarding-options dhcp-security option-82]</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>                                | Statement introduced in Junos OS Release 9.3 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Hierarchy level <code>[edit vlans <i>vlan-name</i> forwarding-options dhcp-security]</code> introduced in Junos OS Release 13.2X50-D10. (See <a href="#">“Getting Started with Enhanced Layer 2 Software” on page 41</a> for information about ELS.)<br>Hierarchy level <code>[edit bridge-domains <i>bridge-domain-name</i> forwarding-options dhcp-security]</code> introduced in Junos OS Release 14.1 for the MX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| <b>Description</b>                                        | Insert a vendor ID in the DHCP option 82 information in a DHCP request packet header before forwarding or relaying the request to a DHCP server.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Default</b>                                            | If <b>vendor-id</b> is not explicitly configured for DHCP option 82, then no vendor ID is set.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                                            | <b>string</b> —(Optional) A single string that designates the vendor ID.<br><br><b>Range:</b> 1–255 characters<br><br><b>Default:</b> If you specify <b>vendor-id</b> with no <b>string</b> value, then the default vendor ID <b>Juniper Networks</b> is configured.                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b>                           | <b>system</b> —To view this statement in the configuration.<br><b>system-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>                              | <ul style="list-style-type: none"> <li><i>Configuring DHCP Option 82 to Help Protect the Switching Devices Against Attacks (CLI Procedure)</i></li> <li><i>Example: Setting Up DHCP Option 82 with a Switch with No Relay Agent Between Clients and a DHCP Server</i></li> <li><i>Example: Setting Up DHCP Option 82 with a Switch as a Relay Agent Between Clients and a DHCP Server</i></li> <li><i>Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)</i></li> </ul>                                                                                                                      |

---

## write-interval

---

|                                 |                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>write-interval seconds;</code>                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | For platforms without ELS:<br><br>[edit ethernet-switching-options secure-access-port dhcp-snooping-file]<br><br>For platforms with ELS:<br><br>[edit system processes] <b>dhcp-service</b> dhcp-snooping-file] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                               |
| <b>Description</b>              | Specify how frequently the switch writes the database entries from memory into the specified DHCP snooping database file.                                                                                       |
| <b>Default</b>                  | None                                                                                                                                                                                                            |
| <b>Options</b>                  | <b>seconds</b> —Value in seconds.<br><b>Range:</b> 60 through 86400                                                                                                                                             |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Understanding DHCP Snooping for Monitoring DHCP Messages Received from Untrusted Devices</i></li></ul>                                                               |





## Configuration Statements (DDoS Protection)

- [bandwidth \(DDoS\) on page 6146](#)
- [bandwidth-scale \(DDoS\) on page 6147](#)
- [burst \(DDoS\) on page 6148](#)
- [burst-scale \(DDoS\) on page 6149](#)
- [bypass-aggregate \(DDoS\) on page 6150](#)
- [ddos-protection \(DDoS\) on page 6151](#)
- [disable-fpc \(DDoS\) on page 6152](#)
- [disable-logging \(DDoS\) on page 6153](#)
- [fpc \(DDoS\) on page 6154](#)
- [global \(DDoS\) on page 6155](#)
- [priority \(DDoS\) on page 6156](#)
- [protocols \(DDoS\) on page 6157](#)
- [traceoptions \(DDoS\) on page 6162](#)

## bandwidth (DDoS)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>bandwidth <i>packets-per-second</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | <ul style="list-style-type: none"><li>For MX Series routers, T4000 routers, and EX9200 switches:<br/>[edit system ddos-protection protocols <i>protocol-group</i> (aggregate   <i>packet-type</i>)]</li><li>For QFX10000 switches:<br/>[edit system ddos-protection <b>protocols</b> <i>protocol-group</i> (aggregate   <i>packet-type</i>)]</li></ul>                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 11.2.</p> <p>Statement introduced in Junos OS Release 12.3R2 on EX9200 switches and T4000 routers.</p> <p>Statement introduced in Junos OS Release 15.1X53 on QFX10000 switches.</p> <p>Statement introduced in Junos OS Release 15.1X53-D30 on QFX5200 switches.</p>                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | (MX Series routers with only MPCs, T4000 routers with only FPC5s, EX9200 switches, or QFX10000 switches) Configure the DDoS bandwidth rate limit; that is, the maximum traffic rate (packets per second) allowed by the specified policer. When the value is exceeded, a violation is declared.                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <p><b><i>packets-per-second</i></b>—Number of packets per second that are allowed by the aggregate or packet-type policer.</p> <p><b>Range:</b> 1 through 100,000 packets per second</p> <p><b>Default:</b> The default bandwidth value varies by packet type or protocol. You can view the default values for all packet types or protocols before you begin DDoS protection configuration by entering the <b>show ddos-protection protocols parameters brief</b> command from operational mode. For QFX10000 switches, the default bandwidth limits are also provided in the <b>protocols (DDoS)</b> statement description.</p> |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li><a href="#">Configuring DDoS Protection Policers for Individual Packet Types</a></li><li><a href="#">Configuring DDoS Protection Policers on QFX Series Switches on page 6074</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                 |

## bandwidth-scale (DDoS)

|                                 |                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>bandwidth-scale <i>percentage</i>;</code>                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit system ddos-protection protocols <i>protocol-group</i> (aggregate   <i>packet-type</i> ) <b>fpc</b> <i>slot-number</i> ]                                                                                                                                                                                    |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 11.2.</p> <p>Statement introduced in Junos OS Release 12.3R2 on EX9200 switches and T4000 routers.</p> <p>Statement introduced in Junos OS Release 15.1X53 on QFX10000 switches.</p> <p>Statement introduced in Junos OS Release 15.1X53-D30 on QFX5200 switches.</p> |
| <b>Description</b>              | (MX Series routers with only MPCs, T4000 routers with only FPC5s, EX9200 switches, or QFX10000 switches) Configure the percentage by which the DDoS bandwidth rate limit is scaled down for the aggregate or packet-type policer on the card in the specified slot.                                               |
| <b>Options</b>                  | <p><b><i>percentage</i></b>—Percentage multiplied by the bandwidth rate limit to reduce the number of packets per second allowed for the packet type or protocol.</p> <p><b>Range:</b> 1 through 100 percent</p> <p><b>Default:</b> 100</p>                                                                       |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring DDoS Protection Policers for Individual Packet Types</a></li> <li>• <a href="#">Configuring DDoS Protection Policers on QFX Series Switches on page 6074</a></li> </ul>                                                                          |

## burst (DDoS)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>burst size;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | <ul style="list-style-type: none"><li>For MX Series routers, T4000 routers, and EX9200 switches:<br/>[edit system ddos-protection protocols <i>protocol-group</i> (aggregate   <i>packet-type</i>)]</li><li>For QFX10000 switches:<br/>[edit system ddos-protection <a href="#">protocols</a> <i>protocol-group</i> (aggregate   <i>packet-type</i>)]</li></ul>                                                                                                                                                                                                                             |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 11.2.</p> <p>Statement introduced in Junos OS Release 12.3R2 on EX9200 switches and T4000 routers.</p> <p>Statement introduced in Junos OS Release 15.1X53 on QFX10000 switches.</p> <p>Statement introduced in Junos OS Release 15.1X53-D30 on QFX5200 switches.</p>                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | (MX Series routers with only MPCs, T4000 routers with only FPC5s, EX9200 switches, or QFX10000 switches) Configure the DDoS burst limit; that is, the maximum number of packets that is allowed in a burst of traffic by the specified policer. When this value is exceeded, a violation is declared.                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <p><b>size</b>—Number of packets that are allowed in a burst by the aggregate or packet-type policer.</p> <p><b>Range:</b> 1 through 100,000 packets</p> <p><b>Default:</b> The default burst value varies by packet type or protocol. You can view the default values for all packet types or protocols on an unconfigured router or switch by entering the <b>show ddos-protection protocols parameters brief</b> command from operational mode. For QFX10000 switches, the default bandwidth limits are also provided in the <a href="#">protocols (DDoS)</a> statement description.</p> |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li><a href="#">Configuring DDoS Protection Policers for Individual Packet Types</a></li><li><a href="#">Configuring DDoS Protection Policers on QFX Series Switches on page 6074</a></li></ul>                                                                                                                                                                                                                                                                                                                                                           |

## burst-scale (DDoS)

|                                 |                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>burst-scale <i>percentage</i>;</code>                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit system ddos-protection protocols <i>protocol-group</i> (aggregate   <i>packet-type</i> ) <b>fpc</b> <i>slot-number</i> ]                                                                                                                                                                                    |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 11.2.</p> <p>Statement introduced in Junos OS Release 12.3R2 on EX9200 switches and T4000 routers.</p> <p>Statement introduced in Junos OS Release 15.1X53 on QFX10000 switches.</p> <p>Statement introduced in Junos OS Release 15.1X53-D30 on QFX5200 switches.</p> |
| <b>Description</b>              | (MX Series routers with only MPCs, T4000 routers with only FPC5s, EX9200 switches, or QFX10000 switches) Configure the percentage by which the DDoS burst limit is scaled down for the aggregate or packet-type policer on the card in the specified slot.                                                        |
| <b>Options</b>                  | <p><b><i>percentage</i></b>—Percentage multiplied by the burst limit to reduce the number of packets allowed in a burst for the packet type or protocol.</p> <p><b>Range:</b> 1 through 100 percent</p> <p><b>Default:</b> 100</p>                                                                                |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring DDoS Protection Policers for Individual Packet Types</i></li> <li>• <a href="#">Configuring DDoS Protection Policers on QFX Series Switches on page 6074</a></li> </ul>                                                                                   |

## bypass-aggregate (DDoS)

---

|                                 |                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | bypass-aggregate;                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | <ul style="list-style-type: none"><li>For MX Series routers, T4000 routers, and EX9200 switches:<br/>[edit system ddos-protection protocols <i>protocol-group packet-type</i>]</li><li>For QFX10000 and QFX5200 switches:<br/>[edit system ddos-protection <a href="#">protocols</a> <i>protocol-group packet-type</i>]</li></ul> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 11.2.</p> <p>Statement introduced in Junos OS Release 12.3R2 on EX9200 switches and T4000 routers.</p> <p>Statement introduced in Junos OS Release 15.1X53 on QFX10000 switches.</p> <p>Statement introduced in Junos OS Release 15.1X53-D30 on QFX5200 switches.</p>                 |
| <b>Description</b>              | (MX Series routers with only MPCs, T4000 routers with only FPC5s, EX9200 switches, or QFX10000 switches) Prevent this packet type from being considered by the DDoS aggregate policer. Traffic for the packet type is still included in traffic statistics.                                                                       |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li><i>Configuring DDoS Protection Policers for Individual Packet Types</i></li><li><a href="#">Configuring DDoS Protection Policers on QFX Series Switches on page 6074</a></li></ul>                                                                                                          |

## ddos-protection (DDoS)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> ddos-protection   global {     disable-fpc;     disable-logging;   }   protocols protocol-group (aggregate   packet-type) {     bandwidth packets-per-second;     burst size;     bypass-aggregate;     disable-fpc;     disable-logging;     fpc slot-number {       bandwidth-scale percentage;       burst-scale percentage;       disable-fpc;     }     priority level;   }   traceoptions {     file filename &lt;files number&gt; &lt;match regular-expression &gt; &lt;size maximum-file-size&gt;       &lt;world-readable   no-world-readable&gt;;     flag flag;     level (all   error   info   notice   verbose   warning);     no-remote-trace;   } </pre> |
| <b>Hierarchy Level</b>          | [edit system]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 15.1X53 on QFX10000 switches.</p> <p>Statement introduced in Junos OS Release 15.1X53-D30 on QFX5200 switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | <p>(QFX10000 and QFX5200 switches) Configure DDoS policers.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">Configuring Protection Against DDoS Attacks on page 6069</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## disable-fpc (DDoS)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | disable-fpc;                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit system ddos-protection <a href="#">global</a> ],<br>[edit system ddos-protection protocols <i>protocol-group</i> (aggregate   <i>packet-type</i> )],<br>[edit system ddos-protection protocols <i>protocol-group</i> (aggregate   <i>packet-type</i> ) <a href="#">fpc</a><br><i>slot-number</i> ]                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.2.<br>Support at the [edit system ddos-protection protocols <i>protocol-group</i> (aggregate   <i>packet-type</i> )] hierarchy level introduced in Junos OS Release 12.1.<br>Statement introduced in Junos OS Release 12.3R2 on EX9200 switches and T4000 routers.<br>Statement introduced in Junos OS Release 15.1X53 on QFX10000 switches.<br>Statement introduced in Junos OS Release 15.1X53-D30 on QFX5200 switches.           |
| <b>Description</b>              | (MX Series routers with only MPCs, T4000 routers with only FPC5s, EX9200 switches, QFX5200 switches, or QFX10000 switches) Disable DDoS policers for debugging purposes on the card in the specified slot for a particular packet type within a protocol group, on all cards for a particular packet type within a protocol group, or globally on all cards and for all packet types in all protocols. This statement does not affect the state of the Routing Engine policers. |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Disabling DDoS Protection Policers and Logging Globally on page 6074</a></li><li>• <a href="#">Configuring DDoS Protection Policers for Individual Packet Types</a></li><li>• <a href="#">Configuring DDoS Protection Policers on QFX Series Switches on page 6074</a></li></ul>                                                                                                                                            |



## disable-logging (DDoS)


|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | disable-logging;                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit system ddos-protection <a href="#">global</a> ],<br>[edit system ddos-protection protocols <i>protocol-group</i> (aggregate   <i>packet-type</i> )]                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.2.<br>Support at the [edit system ddos-protection protocols <i>protocol-group</i> (aggregate   <i>packet-type</i> )] hierarchy level introduced in Junos OS Release 12.1.<br>Statement introduced in Junos OS Release 12.3R2 on EX9200 switches and T4000 routers.<br>Statement introduced in Junos OS Release 15.1X53 on QFX10000 switches.<br>Statement introduced in Junos OS Release 15.1X53-D30 on QFX5200 switches. |
| <b>Description</b>              | (MX Series routers with only MPCs, T4000 routers with only FPC5s, EX9200 switches, QFX5200 switches, or QFX10000 switches) Disable device-wide logging of all DDoS violation and flow detection events globally. Disable only logging of events other than flow detection culprit flow events for a particular packet type or for a protocol group. Typically used for debugging purposes.                                                                            |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Disabling DDoS Protection Policers and Logging Globally on page 6074</a></li> <li>• <a href="#">Configuring DDoS Protection Policers for Individual Packet Types</a></li> <li>• <a href="#">Configuring DDoS Protection Policers on QFX Series Switches on page 6074</a></li> <li>• <a href="#">Disabling Automatic Logging of Culprit Flow Events for a Packet Type</a></li> </ul>                              |

## fpc (DDoS)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>fpc slot-number;<br/>  bandwidth-scale percentage;<br/>  burst-scale percentage;<br/>  disable-fpc;<br/>}</pre>                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | <ul style="list-style-type: none"><li>For MX Series routers, T4000 routers, and EX9200 switches:<br/>[edit system ddos-protection protocols <i>protocol-group</i> (aggregate   <i>packet-type</i>)]</li><li>For QFX10000 and QFX5200 switches:<br/>[edit system ddos-protection <a href="#">protocols</a> <i>protocol-group</i> (aggregate   <i>packet-type</i>)]</li></ul> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 11.2.</p> <p>Statement introduced in Junos OS Release 12.3R2 on EX9200 switches and T4000 routers.</p> <p>Statement introduced in Junos OS Release 15.1X53 on QFX10000 switches.</p> <p>Statement introduced in Junos OS Release 15.1X53-D30 on QFX5200 switches.</p>                                                           |
| <b>Description</b>              | (MX Series routers with only MPCs, T4000 routers with only FPC5s, EX9200 switches, or QFX10000 switches) Modify the aggregate or packet-type policer on the specified line card.                                                                                                                                                                                            |
| <b>Options</b>                  | <p><b>slot-number</b>—Slot number of the card.</p> <p><b>Range:</b> Depends on the router or switch model</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li><a href="#">Configuring DDoS Protection Policers for Individual Packet Types</a></li><li><a href="#">Configuring DDoS Protection Policers on QFX Series Switches on page 6074</a></li></ul>                                                                                                                                           |

## global (DDoS)

|                                                                                                                                                                                                                                                                                                                                                                                                                                 |                                                                                                                                                                                                                                                                                                                   |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                                                   | <pre>global {   disable-fpc;   disable-logging;   disable-routing-engine;   flow-detection;   flow-report-rate;   violation-report-rate; }</pre>                                                                                                                                                                  |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                                                                                          | [edit system ddos-protection]                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                                                                      | <p>Statement introduced in Junos OS Release 11.2.</p> <p>Statement introduced in Junos OS Release 12.3R2 on EX9200 switches and T4000 routers.</p> <p>Statement introduced in Junos OS Release 15.1X53 on QFX10000 switches.</p> <p>Statement introduced in Junos OS Release 15.1X53-D30 on QFX5200 switches.</p> |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                              | (MX Series routers with only MPCs, T4000 routers with only FPC5s, EX9200 switches, QFX5200 switches, or QFX10000 switches) Modify DDoS policers, event logging, and flow detection globally for all protocols.                                                                                                    |
| <div style="display: flex; align-items: center;">  <div> <p><b>NOTE:</b> The following statements are not supported on QFX5200 and QFX10000 switches: <code>disable-routing-engine</code>, <code>flow-detection</code>, <code>flow-report-rate</code>, and <code>violation-report-rate</code>.</p> </div> </div> |                                                                                                                                                                                                                                                                                                                   |
| <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                       |                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                                                                                 | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>                                                                                                                                                                                        |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                                                                                                    | <ul style="list-style-type: none"> <li>• <a href="#">Disabling DDoS Protection Policers and Logging Globally on page 6074</a></li> </ul>                                                                                                                                                                          |

## priority (DDoS)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>priority level;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | <ul style="list-style-type: none"><li>For MX Series routers, T4000 routers, and EX9200 switches:<br/>[edit system ddos-protection protocols <i>protocol-group packet-type</i>]</li><li>For QFX10000 and QFX5200 switches:<br/>[edit system ddos-protection <b>protocols</b> <i>protocol-group packet-type</i>]</li></ul>                                                                                                                                                                     |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 11.2.</p> <p>Statement introduced in Junos OS Release 12.3R2 on EX9200 switches and T4000 routers.</p> <p>Statement introduced in Junos OS Release 15.1X53 on QFX10000 switches.</p> <p>Statement introduced in Junos OS Release 15.1X53-D30 on QFX5200 switches.</p>                                                                                                                                                                            |
| <b>Description</b>              | (MX Series routers with only MPCs, T4000 routers with only FPC5s, EX9200 switches, QFX5200 switches, or QFX10000 switches) Configure the priority for the packet type within the parent protocol group. In the event of downstream traffic congestion, high priority packets are provided bandwidth before medium priority packets. In turn, medium priority packets are provided bandwidth before low priority packets. Packets are dropped when there is insufficient available bandwidth. |
| <b>Options</b>                  | <b>level</b> —Priority of the packet type, low, medium, or high.                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li><i>Configuring DDoS Protection Policers for Individual Packet Types</i></li><li><a href="#">Configuring DDoS Protection Policers on QFX Series Switches on page 6074</a></li></ul>                                                                                                                                                                                                                                                                     |

## protocols (DDoS)

**Syntax** `protocols protocol-group (aggregate | packet-type) {  
     bandwidth packets-per-second;  
     burst size;  
     bypass-aggregate;  
     disable-fpc;  
     disable-logging;  
     fpc slot-number {  
         bandwidth-scale percentage;  
         burst-scale percentage;  
         disable-fpc;  
     }  
     priority level;  
}`

**Hierarchy Level** [edit system `ddos-protection`]

**Release Information** Statement introduced in Junos OS Release 11.2.  
 Statement introduced in Junos OS Release 15.1X53 on QFX10000 switches.  
 Statement introduced in Junos OS Release 15.1X53-D30 on QFX5200 switches.

**Description** Configure DDoS policers for all packet types within a protocol group or for a particular packet type within a protocol group.

**Options** **aggregate**—Configure the policer that polices all control packets belonging to the specified protocol as a combined group. An aggregate policer exists for all protocol groups.

**packet-type**—Name of the control packet type to be policed. You can configure a packet-type policer only for the protocol groups listed in [Table 486](#). For all other protocol groups, only aggregate policers are supported. [Table 486](#) lists the packet-type policers and their default configuration.

**Table 486: Packet Types Supported by DDoS Protection on QFX Switches**

| Protocol Group | Packet Type   | Description                        | Default Bandwidth | Default Burst | Default Priority |
|----------------|---------------|------------------------------------|-------------------|---------------|------------------|
| mcast-snoop    | igmp          | Control packets for IGMP snooping  | 500               | 2048          | High             |
|                | mld           | Control packets for MLD snooping   | 500               | 2048          | High             |
|                | pim           | Control packets for PIM snooping   | 500               | 2048          | High             |
|                | unclassified  | Not supported on QFX10000 switches | —                 | —             | —                |
| radius         | accounting    | RADIUS accounting packets          | 200               | 2048          | High             |
|                | authorization | RADIUS authorization packets       | 200               | 2048          | High             |
|                | server        | RADIUS server traffic              | 200               | 2048          | High             |

***protocol-group***—Name of the protocol group for which traffic is policed. You can configure the aggregate policer for any of the following protocol groups listed in [Table 487](#). The table shows the default configuration for the policers.

Table 487: Protocol Groups Supported by DDoS Protection on QFX Switches

| Protocol Group                 | Description                                                  | Default Bandwidth | Default Burst |
|--------------------------------|--------------------------------------------------------------|-------------------|---------------|
| <b>all-fiber-channel-enode</b> | Fiber channel ENode traffic                                  | 10                | 2048          |
| <b>arp</b>                     | ARP traffic                                                  | 500               | 1024          |
| <b>arp-snoop</b>               | ARP snooping traffic                                         | 500               | 2048          |
| <b>bfd</b>                     | Single-hop BFD traffic                                       | 1000              | 2048          |
| <b>bfdv6</b>                   | BFDv6 traffic                                                | 3000              | 10000         |
| <b>bgp</b>                     | BGP traffic                                                  | 1500              | 2048          |
| <b>bridge-control</b>          | Bridge Control traffic                                       | 10                | 2048          |
| <b>dhcpv4v6</b>                | DHCPv4 and DHCPv6 traffic (limits apply to combined traffic) | 500               | 2048          |
| <b>diameter</b>                | Diameter and Gx-Plus traffic                                 | 200               | 2048          |
| <b>dns</b>                     | DNS traffic                                                  | 200               | 2048          |
| <b>dtcp</b>                    | DTCP traffic                                                 | 200               | 2048          |
| <b>egpv6</b>                   | EGPv6 traffic                                                | 10                | 2048          |
| <b>ethernet-tcc</b>            | TCC-encapsulated Ethernet traffic                            | 100               | 2048          |
| <b>ftp</b>                     | FTP traffic                                                  | 500               | 2048          |
| <b>garp-reply</b>              | Gratuitous ARP reply traffic                                 | 100               | 2048          |
| <b>gre</b>                     | GRE traffic                                                  | 500               | 2048          |
| <b>icmp</b>                    | ICMP traffic                                                 | 500               | 2048          |
| <b>igmp</b>                    | IGMPv4 and IGMPv6 traffic                                    | 1000              | 2048          |
| <b>ip-options</b>              | IP traffic with IP packet header options                     | 100               | 2048          |
| <b>isis</b>                    | IS-IS traffic                                                | 1000              | 2048          |
| <b>iso-tcc</b>                 | TCC-encapsulated ISO traffic                                 | 100               | 2048          |
| <b>l2tp</b>                    | Layer 2 protocol tunneling traffic                           | 500               | 2048          |
| <b>lACP</b>                    | LACP traffic                                                 | 300               | 2048          |

Table 487: Protocol Groups Supported by DDoS Protection on QFX Switches (*continued*)

| Protocol Group         | Description                                        | Default Bandwidth | Default Burst |
|------------------------|----------------------------------------------------|-------------------|---------------|
| <b>ldp</b>             | LDP traffic                                        | 1000              | 200           |
| <b>ldp-hello</b>       | LDP hello packets                                  | 1000              | 2048          |
| <b>lldp</b>            | LLDP traffic                                       | 60                | 2048          |
| <b>lmp</b>             | LMP traffic                                        | 100               | 2048          |
| <b>martian-address</b> | Martian address                                    | 200               | 20            |
| <b>mcast-snoop</b>     | Control traffic for multicast snooping             | 500               | 2048          |
| <b>mld</b>             | MLD traffic                                        | 1000              | 2048          |
| <b>msdp</b>            | MSDP traffic                                       | 300               | 2048          |
| <b>multihop-bfd</b>    | Multihop BFD traffic                               | 1500              | 2048          |
| <b>ndpv6</b>           | NDPv6 traffic                                      | 500               | 1024          |
| <b>ntp</b>             | NTP traffic                                        | 200               | 2048          |
| <b>oam-cfm</b>         | OAM CFM traffic                                    | 200               | 2048          |
| <b>oam-lfm</b>         | OAM LFM traffic                                    | 200               | 2048          |
| <b>ospf</b>            | OSPF traffic                                       | 1000              | 200           |
| <b>ospf-hello</b>      | OSPF hello packets                                 | 1500              | 2048          |
| <b>pim-ctrl</b>        | PIM control packets                                | 1000              | 2048          |
| <b>pim-data</b>        | PIM data                                           | 2000              | 2048          |
| <b>proto-802-lx</b>    | 802.1X traffic                                     | 200               | 2048          |
| <b>ptp</b>             | PTP traffic                                        | 100               | 2048          |
| <b>pvstp</b>           | PVSTP traffic                                      | 2000              | 2048          |
| <b>radius</b>          | RADIUS traffic                                     | 200               | 2048          |
| <b>reject</b>          | Packets rejected by a next-hop forwarding decision | 100               | 2048          |
| <b>resolve</b>         |                                                    | 500               | 2048          |



Table 487: Protocol Groups Supported by DDoS Protection on QFX Switches (*continued*)

| Protocol Group | Description                                                                                             | Default Bandwidth | Default Burst |
|----------------|---------------------------------------------------------------------------------------------------------|-------------------|---------------|
|                | Unclassified IPv4 and IPv6 resolve packets sent to the host because of a traffic request resolve action |                   |               |
| <b>rip</b>     | RIP traffic                                                                                             | 100               | 2048          |
| <b>rsvp</b>    | RSVP traffic                                                                                            | 1000              | 2048          |
| <b>snmp</b>    | SNMP traffic                                                                                            | 500               | 2048          |
| <b>ssh</b>     | SSH traffic                                                                                             | 500               | 2048          |
| <b>stp</b>     | STP traffic                                                                                             | 2000              | 2048          |
| <b>tacacs</b>  | TACACS+ traffic                                                                                         | 200               | 2048          |
| <b>telnet</b>  | Telnet traffic                                                                                          | 500               | 2048          |
| <b>tll</b>     | Time to Live packets                                                                                    | 100               | 2048          |
| <b>vrrp</b>    | VRRP traffic                                                                                            | 1000              | 2048          |

The remaining statements are explained separately.

**Required Privilege Level** admin—To view this statement in the configuration.  
admin-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring DDoS Protection Policers on QFX Series Switches on page 6074](#)

## traceoptions (DDoS)

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>traceoptions {<br/>    file <i>filename</i> &lt;files <i>number</i>&gt; &lt;match <i>regular-expression</i> &gt; &lt;size <i>maximum-file-size</i>&gt;<br/>    &lt;world-readable   no-world-readable&gt;;<br/>    flag <i>flag</i>;<br/>    level (all   error   info   notice   verbose   warning);<br/>    no-remote-trace;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>     | [edit system ddos-protection]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 11.2.</p> <p>Statement introduced in Junos OS Release 12.3R2 on EX9200 switches and T4000 routers.</p> <p>Statement introduced in Junos OS Release 15.1X53 on QFX10000 switches.</p> <p>Statement introduced in Junos OS Release 15.1X53-D30 on QFX5200 switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>         | (MX Series routers with only MPCs, T4000 routers with only FPC5s, EX9200 switches, QFX5200 switches, or QFX10000 switches) Define tracing operations for DDoS protection processes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>             | <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the filename within quotation marks. All files are placed in the directory <code>/var/log</code>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the <b>size</b> option.</p> <p><b>Range:</b> 2 through 1000</p> <p><b>Default:</b> 3 files</p> <p><b>flag <i>flag</i></b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements. You can include the following flags:</p> <ul style="list-style-type: none"><li>• <b>all</b>—Trace all operations.</li><li>• <b>config</b>—Trace processing of the DDoS configuration at an extensive level.</li><li>• <b>events</b>—Trace jddosd event processing; currently only exit events are traced.</li><li>• <b>gres</b>—Trace messages exchanged with the kernel and jddosd process that could affect graceful Routing Engine switchover (GRES).</li><li>• <b>init</b>—Trace jddosd initialization.</li><li>• <b>ipc</b>—Trace interface interprocess communication (IPC) messages.</li><li>• <b>memory</b>—Trace memory management code. This flag is not currently supported.</li><li>• <b>protocol</b>—Trace DDoS protocol state processing. Only the violation state is currently traced.</li><li>• <b>rtsock</b>—Trace messages exchanged with the kernel and jddosd process.</li></ul> |

- **signal**—Trace system signals that are passed to jddosd, such as SIGTERM.
- **socket**—Trace socket messages that are passed to jddosd from the Packet Forwarding Engine.
- **state**—Trace state machine events. This flag is not currently supported.
- **timer**—Trace jddosd timer events.
- **ui**—Trace user interface processing. This flag is not currently supported.

**level**—Level of tracing to perform. You can specify any of the following levels:

- **all**—Match all levels.
- **error**—Match error conditions.
- **info**—Match informational messages.
- **notice**—Match notice messages about conditions requiring special handling.
- **verbose**—Match verbose messages.
- **warning**—Match warning messages.

**match *regular-expression***—(Optional) Refine the output to include lines that contain the regular expression.

**no-remote-trace**—Disable remote tracing.

**no-world-readable**—(Optional) Disable unrestricted file access.

**size *maximum-file-size***—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

**Syntax:** *sizek* to specify KB, *sizem* to specify MB, or *sizeg* to specify GB

**Range:** 10,240 through 1,073,741,824

**world-readable**—(Optional) Enable unrestricted file access.

|                                 |                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | trace—To view this statement in the configuration.                                            |
|                                 | trace-control—To add this statement to the configuration.                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Tracing DDoS Protection Operations</i></li> </ul> |



## CHAPTER 233

# Operational Commands (Firewall Filters)

- `clear firewall`
- `show firewall`
- `show firewall policer`
- `show interfaces filters`
- `show pfe filter hw summary`

## clear firewall

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>clear firewall (all   counter <i>counter-name</i>   filter <i>filter-name</i>)</code>                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.<br>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                   |
| <b>Description</b>              | <p>Clear statistics provided by firewall filters.</p> <p>When you clear the counters of a filter, this not only impacts the counters shown by the CLI, but also the ones tracked by SNMP 2.</p>                                                                                                                                                                             |
| <b>Options</b>                  | <p><b>all</b>—Clear the packet and byte counts for all firewall filter counters and clear the packet counts for all policer counters.</p> <p><b>counter <i>counter-name</i></b>—Clear the packet and byte counts for the specified firewall filter counter.</p> <p><b>filter <i>filter-name</i></b>—Clear the packet and byte counts for the specified firewall filter.</p> |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Verifying That Firewall Filters Are Operational on page 5988</a></li><li>• <a href="#">Verifying That Two-Color Policers Are Operational on page 6019</a></li><li>• <a href="#">Overview of Firewall Filters on page 5951</a></li><li>• <a href="#">Overview of Policers on page 5999</a></li></ul>                     |

## Sample Output

### clear firewall all

```
user@switch> clear firewall all
```

### clear firewall counter

```
user@switch> clear firewall counter port-filter-counter
```

### clear firewall filter

```
user@switch> clear firewall filter ingress-port-filter
```

## show firewall

|                                 |                                                                                                                                                                                                                                                                                                                                                                                     |  |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--|
| <b>Syntax</b>                   | <pre>show firewall   &lt;counter <i>counter-name</i>&gt;   &lt;filter <i>filter-name</i>&gt;   &lt;log &lt;detail   interface <i>interface-name</i>&gt;&gt;   &lt;terse&gt;</pre>                                                                                                                                                                                                   |  |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                |  |
| <b>Description</b>              | Display statistics about configured firewall filters.                                                                                                                                                                                                                                                                                                                               |  |
| <b>Options</b>                  | <p><b>counter <i>counter-name</i></b>—(Optional) Display statistics about a particular firewall filter counter.</p> <p><b>filter <i>filter-name</i></b>—(Optional) Display statistics about a particular firewall filter.</p> <p><b>log</b>—(Optional) Display log entries for all firewall filter activity.</p> <p><b>terse</b>—(Optional) Display firewall filter names only.</p> |  |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                |  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Verifying That Firewall Filters Are Operational on page 5988</a></li> <li>• <a href="#">Verifying That Two-Color Policers Are Operational on page 6019</a></li> <li>• <a href="#">Overview of Firewall Filters on page 5951</a></li> <li>• <a href="#">Overview of Policers on page 5999</a></li> </ul>                        |  |
| <b>List of Sample Output</b>    | <p><a href="#">show firewall on page 6168</a></p> <p><a href="#">show firewall filter <i>filter-name</i> on page 6169</a></p> <p><a href="#">show firewall counter <i>counter-name</i> on page 6169</a></p> <p><a href="#">show firewall log on page 6169</a></p> <p><a href="#">show firewall log detail on page 6169</a></p>                                                      |  |
| <b>Output Fields</b>            | <p><a href="#">Table 488</a> lists the output fields for the <b>show firewall</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                              |  |

**Table 488: show firewall Output Fields**

| Field Name | Field Description                                                                                                           | Level of Output |
|------------|-----------------------------------------------------------------------------------------------------------------------------|-----------------|
| Filter     | Name of the filter that is configured at the <code>[edit firewall family <i>family-name</i> filter]</code> hierarchy level. | All levels      |

Table 488: show firewall Output Fields (*continued*)

| Field Name           | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                         | Level of Output |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Counters</b>      | Display filter counter information: <ul style="list-style-type: none"> <li>Name—Name of a filter counter that has been configured with the <b>count</b> firewall filter action modifier.</li> <li>Bytes—Number of bytes that match the filter term where the <b>count</b> action modifier was specified.</li> <li>Packets—Number of packets that matched the filter term where the <b>count</b> action modifier was specified.</li> </ul> | All levels      |
| <b>Policers</b>      | Display policer information: <ul style="list-style-type: none"> <li>Name—Name of the policer that is configured at the <b>[edit firewall policer]</b> hierarchy level.</li> <li>Packets—Number of packets that matched the filter term where the <b>policer</b> action modifier was specified. This is the number of packets that exceeded the rate limits that the policer specifies.</li> </ul>                                         | All levels      |
| <b>Action</b>        | Filter action: <ul style="list-style-type: none"> <li><b>A</b>—Accept</li> <li><b>D</b>—Discard</li> </ul>                                                                                                                                                                                                                                                                                                                                | All levels      |
| <b>Interface</b>     | Interface on which the firewall filter is applied.                                                                                                                                                                                                                                                                                                                                                                                        | All levels      |
| <b>Protocol</b>      | Name of the packet protocol.                                                                                                                                                                                                                                                                                                                                                                                                              | All levels      |
| <b>Packet Length</b> | Length of the packet.                                                                                                                                                                                                                                                                                                                                                                                                                     | All levels      |
| <b>Src Addr</b>      | Source address of the packet.                                                                                                                                                                                                                                                                                                                                                                                                             | All levels      |
| <b>Dest Addr</b>     | Destination address of the packet.                                                                                                                                                                                                                                                                                                                                                                                                        | All levels      |

## Sample Output

### show firewall

```

user@switch> show firewall
Filter: egress-vlan-watch-employee
Counters:
Name Bytes Packets
counter-employee-web 0 0
Filter: ingress-port-limit-tcp-icmp
Counters:
Name Bytes Packets
icmp-counter 560 10
Policers:
Name Packets
icmp-connection-policer 10
tcp-connection-policer 0
Filter: ingress-vlan-rogue-block
Filter: ingress-vlan-limit-guest

```



**show firewall filter filter-name**

```

user@switch> show firewall filter ingress-port-limit-tcp-icmp
Filter: ingress-port-limit-tcp-icmp
Counters:
Name Bytes Packets
icmp-counter 560 10
Policers:
Name Packets
icmp-connection-policer 10
tcp-connection-policer 0

```

**show firewall counter counter-name**

```

user@switch> show firewall counter icmp-counter
Filter: ingress-port-voip-class-filter
Counters:
Name Bytes Packets
icmp-counter 560 10

```

**show firewall log**

```

user@switch> show firewall log
Log :

Time Filter Action Interface Protocol Src Addr
 Dest Addr
08:00:53 pfe R ge-1/0/6.0 ICMP 192.168.3.5
 192.168.3.4
08:00:52 pfe R ge-1/0/6.0 ICMP 192.168.3.5
 192.168.3.4
08:00:51 pfe R ge-1/0/6.0 ICMP 192.168.3.5
 192.168.3.4
08:00:50 pfe R ge-1/0/6.0 ICMP 192.168.3.5
 192.168.3.4
08:00:49 pfe R ge-1/0/6.0 ICMP 192.168.3.5
 192.168.3.4
08:00:48 pfe R ge-1/0/6.0 ICMP 192.168.3.5
 192.168.3.4
08:00:47 pfe R ge-1/0/6.0 ICMP 192.168.3.5
 192.168.3.4

```

**show firewall log detail**

```

user@switch> show firewall log detail
Log :

Time of Log: 2010-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0Name of protocol: TCP, Packet Length: 50824, Source address:
172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2010-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 1020, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2010-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2010-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of

```

```
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2010-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2010-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
```

## show firewall policer

|                                 |                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show firewall policer</code><br><code>&lt;policer-name&gt;</code>                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.<br>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                    |
| <b>Description</b>              | Display statistics about configured policers.                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <b>none</b> —Display the count of policed packets for all configured policers.<br><br><b>policer-name</b> —(Optional) Display the count of policed packets for the specified policer.                                                                                                                                                                        |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Verifying That Firewall Filters Are Operational on page 5988</a></li> <li>• <a href="#">Verifying That Two-Color Policers Are Operational on page 6019</a></li> <li>• <a href="#">Overview of Firewall Filters on page 5951</a></li> <li>• <a href="#">Overview of Policers on page 5999</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show firewall policer on page 6171</a><br><a href="#">show firewall policer policer-name on page 6172</a>                                                                                                                                                                                                                                        |
| <b>Output Fields</b>            | <a href="#">Table 489</a> lists the output fields for the <b>show firewall policer</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                      |

**Table 489: show firewall policer Output Fields**

| Field Name      | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                 | Level of Output |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Filter</b>   | Name of the filter that is configured at the <code>[edit firewall family <i>family-name</i> filter]</code> hierarchy level.                                                                                                                                                                                                                                                                                                       | All levels      |
| <b>Policers</b> | Display policer information: <ul style="list-style-type: none"> <li>• <b>Filter</b>—Name of filter that specifies the <b>policer</b> action modifier.</li> <li>• <b>Name</b>—Name of policer.</li> <li>• <b>Packets</b>—Number of packets that matched the filter term in which the <b>policer</b> action modifier is specified. This is the number of packets that exceed the rate limits that the policer specifies.</li> </ul> | All levels      |

## Sample Output

### show firewall policer

```
user@switch> show firewall policer
Filter: egress-vlan-filter
Filter: ingress-port-filter
```

```
Policers:
Name Packets
icmp-connection-policer 0
tcp-connection-policer 0
Filter: ingress-vlan-rogue-block
```

#### **show firewall policer policer-name**

```
user@switch> show firewall policer tcp-connection-policer
Filter: ingress-port-filter
Policers:
Name Packets
tcp-connection-policer 0
```

## show interfaces filters

|                                 |                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show interfaces filters</code><br><code>&lt;interface-name&gt;</code>                                                                                                       |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.<br>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                         |
| <b>Description</b>              | Display firewall filters that are configured on each interface in a switch.                                                                                                       |
| <b>Options</b>                  | <b>none</b> —Display firewall filter information about all interfaces.<br><br><b>interface-name</b> —(Optional) Display firewall filter information about a particular interface. |
| <b>Required Privilege Level</b> | view                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show firewall on page 6167</a></li> </ul>                                                                                    |
| <b>List of Sample Output</b>    | <a href="#">show interfaces filters on page 6173</a><br><a href="#">show interfaces filters interface-name on page 6174</a>                                                       |
| <b>Output Fields</b>            | <a href="#">Table 490</a> lists the output fields for the <b>show interfaces filters</b> command. Output fields are listed in the approximate order in which they appear.         |

**Table 490: show interfaces filters Output Fields**

| Field Name           | Field Description                                                                          | Level of Output |
|----------------------|--------------------------------------------------------------------------------------------|-----------------|
| <b>Interface</b>     | Name of the physical interface.                                                            | All levels      |
| <b>Admin</b>         | Interface state: <b>up</b> or <b>down</b> .                                                | All levels      |
| <b>Link</b>          | Link state: <b>up</b> or <b>down</b> .                                                     | All levels      |
| <b>Proto</b>         | Protocol that is configured on the interface.                                              | All levels      |
| <b>Input Filter</b>  | Name of the firewall filter to be evaluated when packets are received on the interface.    | All levels      |
| <b>Output Filter</b> | Name of the firewall filter to be evaluated when packets are transmitted on the interface. | All levels      |

## Sample Output

### show interfaces filters

```

user@switch> show interfaces filters
Interface Admin Link Proto Input Filter Output Filter
ge-0/0/6 up up
ge-0/0/6.0 up up inet

```

|             |    |      |
|-------------|----|------|
| ge-0/0/7    | up | down |
| ge-0/0/8    | up | down |
| ge-0/0/9    | up | down |
| ge-0/0/10   | up | down |
| ge-0/0/10.0 | up | down |

#### show interfaces filters interface-name

```
user@switch> show interfaces filters ge-0/0/6
```

| Interface  | Admin | Link | Proto | Input Filter | Output Filter |
|------------|-------|------|-------|--------------|---------------|
| ge-0/0/6   | up    | up   |       |              |               |
| ge-0/0/6.0 | up    | up   | inet  |              |               |

## show pfe filter hw summary

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show pfe filter hw summary                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Command introduced in Junos OS Release 14.1X53-D10 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | <p>Display a summary of the access control list (ACL; also known as firewall filter) ternary content-addressable memory (TCAM) hardware utilization to show the allocated, used, and free TCAM entry space.</p> <p>Command supported on standalone QFX Series switches, QFX5100-only (pure QFX5100) Virtual Chassis Fabric (VCF), QFX5100-only (pure QFX5100) Virtual Chassis (VC), and QFX3500-only (pure QFX3500) VC.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Planning the Number of Firewall Filters to Create on page 5973</a></li> </ul>                                                                                                                                                                                                                                                                                          |
| <b>List of Sample Output</b>    | <a href="#">show pfe filter hw summary on page 6176</a>                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Output Fields</b>            | <a href="#">Table 491</a> lists the output fields for the <b>show pfe filter hw summary</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                |

**Table 491: show pfe filter hw summary Output Fields**

| Field Name       | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Group</b>     | <p>ACL ingress and egress filter groups:</p> <ul style="list-style-type: none"> <li>• iRACL group—ingress routing ACL filter group</li> <li>• iVACL group—ingress VLAN ACL filter group</li> <li>• iPACL group—ingress port ACL filter group</li> <li>• ePACL group—egress port ACL filter group</li> <li>• eVACL group—egress VLAN ACL filter group</li> <li>• eRACL group—egress routing ACL filter group</li> <li>• eRACL IPv6 group—egress IPv6 routing ACL filter group</li> </ul> |
| <b>Group-ID</b>  | Internal identification number of the filter group.                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Allocated</b> | Number of TCAM filter entries allocated to the filter group.                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Used</b>      | Number of TCAM filter entries used by the filter group.                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Free</b>      | Number of TCAM filter entries available for use by the filter group.                                                                                                                                                                                                                                                                                                                                                                                                                    |

## Sample Output

### show pfe filter hw summary

```
user@switch> show pfe filter hw summary
```

| Group                    | Group-ID | Allocated | Used | Free |
|--------------------------|----------|-----------|------|------|
| -----                    |          |           |      |      |
| > Ingress filter groups: |          |           |      |      |
| iRACL group              | 14       | 512       | 4    | 508  |
| iVACL group              | 13       | 512       | 2    | 510  |
| iPACL group              | 12       | 256       | 2    | 254  |
| > Egress filter groups:  |          |           |      |      |
| ePACL group              | 20       | 256       | 3    | 253  |
| eVACL group              | 21       | 256       | 4    | 252  |
| eRACL group              | 22       | 256       | 245  | 11   |
| eRACL IPV6 group         | 24       | 256       | 3    | 253  |



## CHAPTER 234

# Operational Commands (Port Security)

- `clear ethernet-switching port-error`

## clear ethernet-switching port-error

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | clear ethernet-switching port-error<br><interface <i>interface-name</i> >                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Clear all MAC limiting, MAC move limiting, and storm control errors from all the Ethernet switching interfaces on the switch or from the specified interface, and restore the interfaces or the specified interface to service.                                                                                                                                                             |
| <b>Options</b>                  | <b>none</b> —Clear all MAC limiting, MAC move limiting, and storm control errors from all the Ethernet switching interfaces on the switch and restore the interfaces to service.<br><br><b>interface <i>interface-name</i></b> —(Optional) Clear all MAC limiting, MAC move limiting, and storm control errors from the specified interface and restore the interface to service.           |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring MAC Limiting</i></li><li>• <i>Example: Configuring Storm Control to Prevent Network Outages</i></li><li>• <i>Configuring Port Security (CLI Procedure)</i></li><li>• <a href="#">port-error-disable on page 6140</a></li><li>• <i>Configuring Autorecovery for MAC Limited or Storm Control Interfaces (CLI Procedure)</i></li></ul> |
| <b>Output Fields</b>            | This command produces no output.                                                                                                                                                                                                                                                                                                                                                            |

## CHAPTER 235

# Operational Commands (DDos Protection)

- `clear ddos-protection protocols`
- `show ddos-protection protocols`
- `show ddos-protection protocols parameters`
- `show ddos-protection protocols statistics`
- `show ddos-protection statistics`
- `show ddos-protection version`
- `show ddos-protection protocols violations`

## clear ddos-protection protocols

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>clear ddos-protection protocols</b><br><b>&lt;protocol-group &lt;packet-type&gt;&gt; (culprit-flows   states   statistics)</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.2.<br>Option <b>culprit-flows</b> introduced in Junos OS Release 12.3.<br>Command introduced in Junos OS Release 12.3R2 on EX9200 switches and T4000 routers.<br>Command introduced in Junos OS Release 15.1X53 on QFX10000 switches.<br>Command introduced in Junos OS Release 15.1X53-D30 on QFX5200 switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Clear current DDoS protection statistics, violation states, or culprit flows for all packet types in all protocol groups, for all packet types in a particular protocol group, or for a particular packet type in a particular protocol group.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <b>protocol-group</b> —(Optional) Protocol group that is cleared. See <a href="#">show ddos-protection protocols</a> for a list of available groups.<br><br><b>packet-type</b> —(Optional) Packet type in a particular protocol group that is cleared. See <a href="#">show ddos-protection protocols</a> for a list of available packet types.<br><br><b>culprit-flows</b> —Clear culprit flows for a packet type, for a protocol group, or for all protocol groups. This option is not supported on QFX Series switches.<br><br><b>states</b> —Clear DDoS protection violation states for a packet type, for a protocol group, or for all protocol groups.<br><br><b>statistics</b> —Clear DDoS protection statistics such as packet counts and rates for a packet type, for a protocol group, or for all protocol groups. |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show ddos-protection protocols on page 6182</a></li><li>• <a href="#">show ddos-protection statistics on page 6218</a></li><li>• <a href="#">show ddos-protection version on page 6220</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>List of Sample Output</b>    | <a href="#">clear ddos-protection protocols (Clear Statistics for All Protocols) on page 6180</a><br><a href="#">clear ddos-protection protocols (Clear Violation States for Packet Type) on page 6181</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

### Sample Output

#### clear ddos-protection protocols (Clear Statistics for All Protocols)

```
user@host> clear ddos-protection protocols statistics
```

**clear ddos-protection protocols (Clear Violation States for Packet Type)**

```
user@host> clear ddos-protection protocols radius server states
```

## show ddos-protection protocols

---

**Syntax** `show ddos-protection protocols <protocol-group (aggregate | packet-type)>`

**Release Information** Command introduced in Junos OS Release 11.2.  
Command introduced in Junos OS Release 12.3R2 on EX9200 switches and T4000 routers.  
Command introduced in Junos OS Release 15.1X53 on QFX10000 switches.  
Command introduced in Junos OS Release 15.1X53-D30 on QFX5200 switches.

**Description** Display DDoS protection configuration and statistics for protocol groups or individual packet types.

**Options** **none**—Display information for all packet types in all protocol groups.

**aggregate**—(Optional) Display DDoS protection information for the aggregate policer.  
The **aggregate** option is available for all protocol groups.

**packet-type**—(Optional) Display DDoS protection information for the specified packet type in the protocol group. The available packet types vary by protocol group.  
On QFX10000 switches, only aggregate policers are available for protocol groups that are not in the following list:

- **mcast-snoop**—The following packet types are available for the **mcast-snoop** protocol group:
  - **igmp**—Control packets for IGMP snooping.
  - **mld**—Control packets for MLD snooping.
  - **pim**—Control packets for PIM snooping.
- **radius**—The following packet types are available for the **radius** protocol group:
  - **accounting**—RADIUS accounting packets.
  - **authorization**—RADIUS authorization packets.
  - **server**—RADIUS server traffic.

On MX Series routers, T4000 routers, and EX9200 switches, only aggregate policers are available for protocol groups that are not in the following list:

- **dhcpv4**—The following packet types are available for DHCPv4 traffic:
  - **ack**—DHCPACK packets.
  - **bad-packets**—DHCPv4 packets with bad formats.
  - **bootp**—DHCPBOOTP packets.
  - **decline**—DHCPDECLINE packets.
  - **discover**—DHCDISCOVER packets.
  - **force-renew**—DHCPFORCERENEW packets.

- **inform**—DHCPINFORM packets.
- **lease-active**—DHCPLEASEACTIVE packets.
- **lease-query**—DHCPLEASEQUERY packets.
- **lease-unassigned**—DHCPLEASEUNASSIGNED packets.
- **lease-unknown**—DHCPLEASEUNKNOWN packets.
- **nak**—DHCPNAK packets.
- **no-message-type**—DHCP packets that are missing the message type.
- **offer**—DHCOFFER packets.
- **release**—DHCPACK packets.
- **renew**—DHCPRENEW packets.
- **request**—DHCPREQUEST packets.
- **unclassified**— All unclassified packets in the protocol group.
- **dhcpv6**—The following packet types are available for DHCPv6 traffic:
  - **advertise**—ADVERTISE packets.
  - **confirm**—CONFIRM packets.
  - **decline**—DECLINE packets.
  - **information-request**—INFORMATION-REQUEST packets.
  - **leasequery**—LEASEQUERY packets.
  - **leasequery-data**—LEASEQUERY-DATA packets.
  - **leasequery-done**—LEASEQUERY-DONE packets.
  - **leasequery-reply**—LEASEQUERY-REPLY packets.
  - **rebind**—REBIND packets.
  - **reconfigure**—RECONFIGURE packets.
  - **relay-forward**—RELAY-FORWARD packets.
  - **relay-reply**—RELAY-REPLY packets.
  - **release**—RELEASE packets.
  - **renew**—RENEW packets.
  - **reply**—REPLY packets.
  - **request**—REQUEST packets.
  - **solicit**—SOLICIT packets.
  - **unclassified**— All unclassified packets in the protocol group.
- **filter-action**—The following packet types are available for unclassified firewall filter action packets, sent to the host because of reject terms in firewall filters:

- **filter-v4**—Unclassified IPv4 filter action packets.
- **filter-v6**—Unclassified IPv6 filter action packets.
- **other**—All other unclassified filter action packets that are not IPv4 or IPv6.
- **frame-relay**—The following packet types are available for Frame Relay traffic:
  - **frf15**—Multilink frame relay FRF.15 packets.
  - **frf16**—Multilink frame relay FRF.16 packets.
- **ip-fragments**—The following packet types are available for IP fragments:
  - **first-fragment**—First IP fragment.
  - **trail-fragment**—Last IP fragment.
- **ip-options**—The following packet types are available for IP option traffic:
  - **non-v4v6**—Options packets other than IPv4/v6.
  - **router-alert**—Router alert options packets.
  - **unclassified**— All unclassified packets in the protocol group.
- **l2tp**—The following packet types are available for L2TP traffic:
  - **cdn**—Call-Disconnect-Notify message packets.
  - **hello**—Hello message packets.
  - **iccn**—Incoming-Call-Connected message packets.
  - **icrq**—Incoming-Call-Request message packets.
  - **scccn**—Start-Control-Connection-Connected message packets.
  - **sccrq**—Start-Control-Connection-Request message packets.
  - **stopccn**—Stop-Control-Connection-Notification message packets.
  - **unclassified**—All unclassified packets in the protocol group.
- **mcast-snoop**—Control traffic for multicast snooping.
  - **igmp**—Snooped IGMP traffic.
  - **pim**—Snooped PIM control traffic.
- **mlp**—The following MLP packet types are available:
  - **aging-exception**—MLP aging exception packets.
  - **packets**—MLP packets.
  - **unclassified**— All unclassified packets in the protocol group.
- **ppp**—The following PPP packet types are available:
  - **authentication**—PPP authentication protocol packets.
  - **echo-rep**—LCP echo reply packets.



- **echo-req**—LCP echo request packets.
- **ipcp**—IP Control Protocol packets.
- **ipv6cp**—IPv6 Control Protocol packets.
- **isis**—IS-IS packets.
- **lcp**—Link Control Protocol packets.
- **mlppp-lcp**—MLPPP LCP packets.
- **mplscp**—MPLS Control Protocol packets.
- **unclassified**— All unclassified packets in the protocol group.
- **pppoe**—The following PPPoE packet types are available:
  - **padi**—PADI packets.
  - **padm**—PADM packets.
  - **padn**—PADN packets.
  - **pado**—PADO packets.
  - **padr**—PADR packets.
  - **pads**—PADS packets.
  - **padt**—PADT packets.
- **radius**—The following RADIUS packet types are available:
  - **accounting**—RADIUS accounting packets.
  - **authorization**—RADIUS authorization packets.
  - **server**—RADIUS server traffic.
  - **unclassified**— All unclassified packets in the protocol group.
- **re-services**—The following packet type is available for Routing Engine-based HTTP redirect:
  - **captive-portal**—Routing Engine-based captive portal content delivery packets.
- **resolve**—The following packet types are available for unclassified resolve packets, which are sent to the host because of a traffic request resolve action:
  - **mcast-v4**—Unclassified IPv4 multicast resolve packets.
  - **mcast-v6**—Unclassified IPv6 multicast resolve packets.
  - **ucast-v4**—Unclassified IPv4 unicast resolve packets.
  - **ucast-v6**—Unclassified IPv6 unicast resolve packets.
  - **other**—All other unclassified resolve packets.
- **sample**—The following sample packet types are available:

- **host**—Host packets.
- **pfe**—Packet Forwarding Engine packets.
- **syslog**—System log message packets.
- **tap**—TAP packets.
- **tcp-flags**—The following TCP-flagged packet types are available:
  - **established**—TCP ACK and RST connection packets.
  - **initial**—TCP SYN and SYN ACK packets.
- **unclassified**—The following unclassified packet types are available:
  - **control-layer2**—Unclassified layer 2 control packets.
  - **control-v4**—Unclassified IPv4 control packets.
  - **control-v6**—Unclassified IPv6 control packets.
  - **fw-host**—Unclassified send-to-host firewall packets.
  - **host-route-v4**—Unclassified IPv4 routing protocol and host packets in traffic sent to the router local interface address for broadcast and multicast.
  - **host-route-v6**—Unclassified IPv6 routing protocol and host packets in traffic sent to the router local interface address for broadcast and multicast.
  - **mcast-copy**—Unclassified host copy (due to multicast routing) packets.
  - **other**—All unclassified packets that do not belong to another type.
- **virtual-chassis**—The following packet types are available for virtual chassis packets:
  - **control-low**—Low-priority control packets.
  - **control-high**—High-priority control packets.
  - **unclassified**—All unclassified packets in the protocol group.
  - **vc-packets**—All exception packets on the virtual chassis link.
  - **vc-ttl-errors**—Virtual chassis TTL error packets.

**protocol-group**—(Optional) Display DDoS protection information for a protocol group.

[Table 492](#) lists the protocol groups and the platforms they are supported on.

**Table 492: Supported Protocol Groups**

| Protocol Group                 | Description                 | MX Series Routers, T4000 Routers, EX9200 Switches | QFX10000 Switches |
|--------------------------------|-----------------------------|---------------------------------------------------|-------------------|
| <b>all-fiber-channel-enode</b> | Fiber channel ENode traffic | —                                                 | X                 |
| <b>amtv4</b>                   | IPv4 AMT traffic            | X                                                 | —                 |
| <b>amtv6</b>                   | IPv6 AMT traffic            | X                                                 | —                 |

Table 492: Supported Protocol Groups *(continued)*

| Protocol Group         | Description                    | MX Series Routers, T4000 Routers, EX9200 Switches | QFX10000 Switches |
|------------------------|--------------------------------|---------------------------------------------------|-------------------|
| <b>ancp</b>            | ANCP traffic                   | X                                                 | —                 |
| <b>ancpv6</b>          | ANCPv6 traffic                 | X                                                 | —                 |
| <b>arp</b>             | ARP traffic                    | X                                                 | X                 |
| <b>arp-snoop</b>       | ARP snooping traffic           | —                                                 | X                 |
| <b>atm</b>             | ATM traffic                    | X                                                 | —                 |
| <b>bfd</b>             | Single-hop BFD traffic         | X                                                 | X                 |
| <b>bfdv6</b>           | BFDv6 traffic                  | X                                                 | X                 |
| <b>bgp</b>             | BGP traffic                    | X                                                 | X                 |
| <b>bgpv6</b>           | BGPv6 traffic                  | X                                                 | —                 |
| <b>bridge-control</b>  | Bridge Control traffic         | —                                                 | X                 |
| <b>control</b>         | Control traffic                | X                                                 | —                 |
| <b>demux-autosense</b> | Demux autosensing traffic      | X                                                 | —                 |
| <b>dhcpv4</b>          | DHCPv4 traffic                 | X                                                 | —                 |
| <b>dhcpv6</b>          | DHCPv6 traffic                 | X                                                 | —                 |
| <b>dhcpv4v6</b>        | DHCPv4 and DHCPv6 traffic      | —                                                 | X                 |
| <b>diameter</b>        | Diameter and Gx-Plus traffic   | X                                                 | X                 |
| <b>dns</b>             | DNS traffic                    | X                                                 | X                 |
| <b>dtcp</b>            | DTCP traffic                   | X                                                 | X                 |
| <b>dynamic-vlan</b>    | Dynamic VLAN exception traffic | X                                                 | —                 |
| <b>egpv6</b>           | EGPv6 traffic                  | X                                                 | X                 |
| <b>eoam</b>            | EOAM traffic                   | X                                                 | —                 |
| <b>esmc</b>            | ESMC traffic                   | X                                                 | —                 |

Table 492: Supported Protocol Groups *(continued)*

| Protocol Group       | Description                                                                                               | MX Series Routers, T4000 Routers, EX9200 Switches | QFX10000 Switches |
|----------------------|-----------------------------------------------------------------------------------------------------------|---------------------------------------------------|-------------------|
| <b>ethernet-tcc</b>  | TCC-encapsulated Ethernet traffic                                                                         | –                                                 | X                 |
| <b>fab-probe</b>     | Fab out probe packets                                                                                     | X                                                 | –                 |
| <b>filter-action</b> | IPv4 and IPv6 firewall filter action packets sent to the host because of reject terms in firewall filters | X                                                 | –                 |
| <b>firewall-host</b> | Firewall send-to-host traffic                                                                             | X                                                 | –                 |
| <b>frame-relay</b>   | Frame relay traffic                                                                                       | X                                                 | –                 |
| <b>ftp</b>           | FTP traffic                                                                                               | X                                                 | X                 |
| <b>ftpv6</b>         | FTPV6 traffic                                                                                             | X                                                 | –                 |
| <b>garp-reply</b>    | Gratuitous ARP reply traffic                                                                              | –                                                 | X                 |
| <b>gre</b>           | GRE traffic                                                                                               | X                                                 | X                 |
| <b>icmp</b>          | ICMP traffic                                                                                              | X                                                 | X                 |
| <b>igmp</b>          | IGMP traffic                                                                                              | X                                                 | X                 |
| <b>igmpv4v6</b>      | IGMP and MLD traffic                                                                                      | X                                                 | –                 |
| <b>igmpv6</b>        | MLD traffic                                                                                               | X                                                 | –                 |
| <b>inline-ka</b>     | Inline service interfaces keepalive traffic                                                               | X                                                 | –                 |
| <b>inline-svcs</b>   | Inline services traffic                                                                                   | X                                                 | –                 |
| <b>ip-fragments</b>  | IP fragments traffic                                                                                      | X                                                 | –                 |
| <b>ip-options</b>    | IP traffic with IP packet header options                                                                  | X                                                 | X                 |
| <b>isis</b>          | IS-IS traffic                                                                                             | X                                                 | X                 |
| <b>iso-tcc</b>       | TCC-encapsulated ISO traffic                                                                              | –                                                 | X                 |
| <b>jfm</b>           | JFM traffic                                                                                               | X                                                 | –                 |
| <b>keepalive</b>     | Keepalive traffic                                                                                         | X                                                 | –                 |

Table 492: Supported Protocol Groups *(continued)*

| Protocol Group         | Description                                | MX Series Routers, T4000 Routers, EX9200 Switches | QFX10000 Switches |
|------------------------|--------------------------------------------|---------------------------------------------------|-------------------|
| <b>l2tp</b>            | Layer 2 protocol tunneling traffic         | X                                                 | X                 |
| <b>lACP</b>            | LACP traffic                               | X                                                 | X                 |
| <b>ldp</b>             | LDP traffic                                | X                                                 | X                 |
| <b>ldp-hello</b>       | LDP hello packets                          | —                                                 | X                 |
| <b>ldpv6</b>           | LDPv6 traffic                              | X                                                 | —                 |
| <b>lldp</b>            | LLDP traffic                               | X                                                 | X                 |
| <b>lmp</b>             | LMP traffic                                | X                                                 | X                 |
| <b>lmpv6</b>           | LMPv6 traffic                              | X                                                 | —                 |
| <b>mac-host</b>        | Layer 2 MAC send-to-host traffic           | X                                                 | —                 |
| <b>martian-address</b> | Martian address                            | —                                                 | —                 |
| <b>mcast-snoop</b>     | Control traffic for multicast snooping     | X                                                 | X                 |
| <b>mld</b>             | MLD traffic                                | —                                                 | X                 |
| <b>mlp</b>             | MLP traffic                                | X                                                 | —                 |
| <b>msdp</b>            | MSDP traffic                               | X                                                 | X                 |
| <b>multihop-bfd</b>    | Multihop BFD traffic                       | —                                                 | X                 |
| <b>mld</b>             | MLD traffic                                | —                                                 | X                 |
| <b>msdpv6</b>          | MSDPv6 traffic                             | X                                                 | —                 |
| <b>multicast-copy</b>  | Host copy traffic due to multicast routing | X                                                 | —                 |
| <b>mvrp</b>            | MVRP traffic                               | X                                                 | —                 |
| <b>ndpv6</b>           | NDPv6 traffic                              | X                                                 | X                 |
| <b>ntp</b>             | NTP traffic                                | X                                                 | X                 |
| <b>oam-cfm</b>         | OAM CFM traffic                            | —                                                 | X                 |

Table 492: Supported Protocol Groups *(continued)*

| Protocol Group      | Description                                                              | MX Series Routers, T4000 Routers, EX9200 Switches | QFX10000 Switches |
|---------------------|--------------------------------------------------------------------------|---------------------------------------------------|-------------------|
| <b>oam-lfm</b>      | OAM LFM traffic                                                          | X                                                 | X                 |
| <b>ospf</b>         | OSPF traffic                                                             | X                                                 | X                 |
| <b>ospf-hello</b>   | OSPF hello packets                                                       | —                                                 | X                 |
| <b>ospfv3v6</b>     | OSPFv3/IPv6 traffic                                                      | X                                                 | —                 |
| <b>pfe-alive</b>    | Packet Forwarding Engine keepalive traffic                               | X                                                 | —                 |
| <b>pim</b>          | PIM traffic                                                              | X                                                 | —                 |
| <b>pim-ctrl</b>     | PIM control packets                                                      | —                                                 | X                 |
| <b>pim-data</b>     | PIM data                                                                 | —                                                 | X                 |
| <b>pimv6</b>        | PIMv6 traffic                                                            | X                                                 | —                 |
| <b>pmvrp</b>        | PMVRP traffic                                                            | X                                                 | —                 |
| <b>pos</b>          | POS traffic                                                              | X                                                 | —                 |
| <b>ppp</b>          | PPP traffic                                                              | X                                                 | —                 |
| <b>pppoe</b>        | PPPoE traffic                                                            | X                                                 | —                 |
| <b>proto-802-1x</b> | 802.1X traffic                                                           | —                                                 | X                 |
| <b>ptp</b>          | PTP traffic                                                              | X                                                 | X                 |
| <b>pvstp</b>        | PVSTP traffic                                                            | X                                                 | X                 |
| <b>radius</b>       | RADIUS traffic                                                           | X                                                 | X                 |
| <b>re-services</b>  | Captive portal content delivery traffic for Routing Engine HTTP redirect | X                                                 | —                 |
| <b>redirect</b>     | Traffic that triggers ICMP redirects                                     | X                                                 | —                 |
| <b>reject</b>       | Packets rejected by a next-hop forwarding decision                       | X                                                 | X                 |
| <b>rejectv6</b>     | IPv6 packets rejected by a next-hop forwarding decision                  | X                                                 | —                 |

Table 492: Supported Protocol Groups *(continued)*

| Protocol Group         | Description                                                                                             | MX Series Routers, T4000 Routers, EX9200 Switches | QFX10000 Switches |
|------------------------|---------------------------------------------------------------------------------------------------------|---------------------------------------------------|-------------------|
| <b>resolve</b>         | Unclassified IPv4 and IPv6 resolve packets sent to the host because of a traffic request resolve action | X                                                 | X                 |
| <b>rip</b>             | RIP traffic                                                                                             | X                                                 | X                 |
| <b>ripv6</b>           | RIPv6 traffic                                                                                           | X                                                 | —                 |
| <b>rsvp</b>            | RSVP traffic                                                                                            | X                                                 | X                 |
| <b>rsvpv6</b>          | RSVPv6 traffic                                                                                          | X                                                 | —                 |
| <b>snmp</b>            | SNMP traffic                                                                                            | X                                                 | X                 |
| <b>snmpv6</b>          | SNMPv6 traffic                                                                                          | X                                                 | —                 |
| <b>ssh</b>             | SSH traffic                                                                                             | X                                                 | X                 |
| <b>sshv6</b>           | SSHv6 traffic                                                                                           | X                                                 | —                 |
| <b>stp</b>             | STP traffic                                                                                             | X                                                 | X                 |
| <b>syslog</b>          | System log messages UDP traffic on port 6333 for the Routing Engine syslog server                       | X                                                 | —                 |
| <b>tacacs</b>          | TACACS+ traffic                                                                                         | --                                                | X                 |
| <b>tcp-flags</b>       | Traffic with TCP flags                                                                                  | X                                                 | —                 |
| <b>telnet</b>          | Telnet traffic                                                                                          | X                                                 | X                 |
| <b>telnetv6</b>        | Telnetv6 traffic                                                                                        | X                                                 | —                 |
| <b>ttl</b>             | Time to Live packets                                                                                    | X                                                 | X                 |
| <b>tunnel-fragment</b> | Tunnel fragments traffic                                                                                | X                                                 | —                 |
| <b>unclassified</b>    | Unclassified traffic                                                                                    | X                                                 | —                 |
| <b>virtual-chassis</b> | Virtual chassis traffic                                                                                 | X                                                 | —                 |
| <b>vrrp</b>            | VRRP traffic                                                                                            | X                                                 | X                 |
| <b>vrrpv6</b>          | VRRPv6 traffic                                                                                          | X                                                 | —                 |

**Required Privilege Level** view

- Related Documentation**
- [clear ddos-protection protocols on page 6180](#)
  - [show ddos-protection protocols culprit-flows](#)
  - [show ddos-protection protocols flow-detection](#)
  - [show ddos-protection protocols parameters on page 6201](#)
  - [show ddos-protection protocols statistics on page 6208](#)
  - [show ddos-protection protocols violations on page 6221](#)

**List of Sample Output** [show ddos-protection protocols on page 6196](#)  
[show ddos-protection protocols \(Specific Packet Type with Flow Detection Disabled\) on page 6198](#)  
[show ddos-protection protocols \(Specific Packet Type with Flow Detection Enabled and Automatic\) on page 6198](#)  
[show ddos-protection protocols \(Specific Packet Type with Bandwidth Violation\) on page 6199](#)

**Output Fields** [Table 493](#) lists the output fields for the **show ddos-protection protocols** command. Output fields are listed in the approximate order in which they appear.

**Table 493: show ddos-protection protocols Output Fields**

| Field Name                     | Field Description                                                                                             |
|--------------------------------|---------------------------------------------------------------------------------------------------------------|
| <b>Packet types</b>            | Number of packet types                                                                                        |
| <b>Modified</b>                | Number of packets for which policer values have been modified from the default.                               |
| <b>Received traffic</b>        | Number of traffic flows received.                                                                             |
| <b>Currently violated</b>      | Number of flows that are currently violating the flow bandwidth limit.                                        |
| <b>Currently tracked flows</b> | Number of active flows that are being tracked as culprit flows by flow detection.                             |
| <b>Total detected flows</b>    | Total number of culprit flows that have been detected, including those that have recovered or timed out.      |
| <b>Protocol Group</b>          | Name of protocol group.                                                                                       |
| <b>Packet type</b>             | Name of packet type in protocol group.                                                                        |
| <b>Bandwidth</b>               | Bandwidth policer value; number of packets per second that is allowed before a violation is declared.         |
| <b>Burst</b>                   | Burst policer value; the maximum number of packets that is allowed in a burst before a violation is declared. |



Table 493: show ddos-protection protocols Output Fields (*continued*)

| Field Name                   | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Priority                     | Priority of the packet type for individual packet policers that enables more important traffic to pass through in the event of traffic congestion: <b>low</b> , <b>medium</b> , or <b>high</b> . Lower priority packets can be dropped when insufficient bandwidth is available.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Recover time                 | Time that must pass since the last violation before the traffic flow is considered to have recovered from the attack. A notification is generated when the timer expires.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Enabled                      | State of the policer, enabled ( <b>Yes</b> ), disabled ( <b>No</b> ), or partially disabled ( <b>Partial</b> ); <b>Partial</b> indicates that only some of the policer instances are disabled for the policer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Bypass aggregate             | State of the bypass aggregate configuration: <ul style="list-style-type: none"> <li>• Yes—The aggregate policer is bypassed.</li> <li>• No—The aggregate policer is enforced.</li> </ul> This field appears only for individual policers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Flow detection configuration | State of flow detection configured on the router: <ul style="list-style-type: none"> <li>• Detection mode—Mode of operation for suspicious flow detection: automatic, off, or on.</li> <li>• Log flows—State of automatic logging of suspicious traffic flows: on (<b>Yes</b>) or off (<b>No</b>).</li> <li>• Timeout flows—State of culprit flow timeout behavior: flow is suppressed for a configured timeout period (<b>Yes</b>) or flow is suppressed until it is no longer in violation (<b>No</b>).</li> <li>• Detect time—Time in seconds that must pass before a suspicious flow that has exceeded the bandwidth allowed for the packet type is considered to be a culprit flow.</li> <li>• Recover time—Time in seconds that must pass before a culprit flow is considered to have returned to normal. The period starts when the flow drops below the threshold that triggered the last violation.</li> <li>• Timeout time—Time in seconds that a culprit flow is suppressed, if timeouts have been enabled.</li> <li>• Flow aggregation level configuration—Flow detection mode, flow control mode, and flow bandwidth for traffic at each of the traffic flow aggregation levels: subscriber, logical interface, and physical interface. <ul style="list-style-type: none"> <li>• Detection mode—State of flow detection: automatic, off, or on.</li> <li>• Control mode—Mode of controlling culprit traffic: dropped, kept, or policed back to within the allowed bandwidth.</li> <li>• Flow rate—Bandwidth allowed for the control traffic in packets per second.</li> </ul> </li> </ul> |

Table 493: show ddos-protection protocols Output Fields (*continued*)

| Field Name                        | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>System-wide information</b>    | <p>The following information collected for the router:</p> <ul style="list-style-type: none"> <li>• A message indicates whether the policer has been violated.</li> <li>• No. of FPCs currently receiving excess traffic—Number of cards that are currently in violation of a policer.</li> <li>• No. of FPCs that have received excess traffic—Number of cards that have at some point been in violation of a policer.</li> <li>• Violation first detected at—Timestamp of the first violation.</li> <li>• Violation last seen at—Timestamp of the last observed violation.</li> <li>• Duration of violation—Length of the violation.</li> <li>• Number of violations—Number of times the violation has occurred.</li> <li>• Received—Number of packets received at all card slots and the Routing Engine.</li> <li>• Dropped—Number of packets dropped regardless of where they were dropped.</li> <li>• Arrival rate—Current traffic rate for packets arriving from all cards and at the Routing Engine.</li> <li>• Max arrival rate—Highest traffic rate for packets arriving from all cards and at the Routing Engine.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Routing Engine information</b> | <p>The following information collected for the Routing Engine:</p> <ul style="list-style-type: none"> <li>• Bandwidth—Maximum number of packets per second that is allowed.</li> <li>• Burst—Maximum number of packets that is allowed in a burst.</li> <li>• A message indicates the State of the policer, enabled (<b>Yes</b>) or disabled (<b>No</b>).</li> <li>• A message indicates whether the policer has been violated; the policer might be passed at the individual cards, but the combined rate of packets arriving at the Routing Engine can exceed the configured policer value.</li> <li>• Violation first detected at—Timestamp of the first violation.</li> <li>• Violation last seen at—Timestamp of the last observed violation.</li> <li>• Duration of violation—Length of the violation.</li> <li>• Number of violations—Number of times the violation has occurred.</li> <li>• Received—Number of packets received at the Routing Engine from all cards.</li> <li>• Dropped—Number of packets dropped at the Routing Engine; includes packets dropped by the aggregate policer and by individual protocol policers.</li> <li>• Arrival rate—Current traffic rate for packets arriving at the Routing Engine from all cards.</li> <li>• Max arrival rate—Highest traffic rate for packets arriving at the Routing Engine from all cards.</li> <li>• Dropped by aggregate policer—Number of packets dropped by the aggregate policer.</li> <li>• Dropped by individual policers—Number of packets dropped by individual policer.</li> </ul> |

Table 493: show ddos-protection protocols Output Fields (*continued*)

| Field Name                                  | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>FPC slot information</b>                 | <p>The following information collected for the card in the indicated slot:</p> <ul style="list-style-type: none"> <li>• Bandwidth—Bandwidth scaling percentage and the number of packets per second that is allowed before a violation is declared.</li> <li>• Burst—Burst scaling percentage and the maximum number of packets that is allowed in a burst before a violation is declared.</li> <li>• A message indicates whether the policer has been violated.</li> <li>• Violation first detected at—Timestamp of the first violation.</li> <li>• Violation last seen at—Timestamp of the last observed violation.</li> <li>• Duration of violation—Length of the violation.</li> <li>• Number of violations—Number of times the violation has occurred.</li> <li>• Received—Number of packets received on the line card.</li> <li>• Dropped—Number of packets dropped at the line card; includes packets dropped by the aggregate policer and by individual protocol policers.</li> <li>• Arrival rate—Current traffic rate for packets arriving at the line card.</li> <li>• Max arrival rate—Highest traffic rate for packets arriving at the line card.</li> <li>• Dropped by this policer—Number of packets dropped by the individual policer.</li> <li>• Dropped by aggregate policer—Number of packets dropped by the aggregate policer.</li> </ul> <p><b>NOTE:</b> On MX Series routers with built-in MPCs—the MX5, MX10, MX40, MX80, and MX104 routers—this field actually displays information for tfeb0 because these routers have no Flexible PIC Concentrator (FPC) slots. Instead, the Packet Forwarding Engine has two “pseudo” FPCs (FPC 0 and FPC1).</p> |
| <b>Bypass aggr.</b>                         | <p>State of the bypass aggregate configuration:</p> <ul style="list-style-type: none"> <li>• Yes—The aggregate policer configuration is bypassed.</li> <li>• No—The aggregate policer configuration is enforced.</li> </ul> <p>Dashes indicate that the bypass aggregate configuration is not available; this is possible only for aggregate policers.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>FPC Mod</b>                              | <p>Indicates whether configuration has changed from the default for any line cards.</p> <ul style="list-style-type: none"> <li>• No—The default configuration has not changed from the default for the packet type.</li> <li>• Yes—The default configuration has changed from the default for the packet type</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Op mode</b>                              | <p>Mode of operation for suspicious flow detection for the packet type: always-on (<b>on</b>), (<b>auto</b>), or disabled (<b>off</b>).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Policer BW (pps)</b>                     | <p>Bandwidth policer value; number of packets per second that is allowed before a violation is declared.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Aggr level<br/>Op:Fc:Bandwidth (pps)</b> | <p>Flow operation mode, flow control mode, and flow bandwidth for traffic of the packet type at each traffic flow aggregation level: subscriber (<b>sub</b>), logical interface (<b>ifl</b>), and physical interface (<b>ifd</b>).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

Table 493: show ddos-protection protocols Output Fields (*continued*)

| Field Name      | Field Description                                                                                                                                                                                             |
|-----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Log flow</b> | State of automatic logging of suspicious traffic flows for the packet type: on (Yes) or off (No).                                                                                                             |
| <b>Time out</b> | State of culprit flow timeout behavior for the packet type: flow is suppressed or monitored for a configured timeout period (Yes) or flow is suppressed or monitored until it is no longer in violation (No). |

## Sample Output

### show ddos-protection protocols

```
user@host> show ddos-protection protocols
```

```
Packet types: 190, Modified: 0, Received traffic: 12, Currently violated: 3
Currently tracked flows: 0, Total detected flows: 0
* = User configured value
```

```
Protocol Group: IPv4-Unclassified
```

```
Packet type: aggregate (Aggregate for unclassified host-bound IPv4 traff)
```

```
Aggregate policer configuration:
```

```
Bandwidth: 2000 pps
Burst: 10000 packets
Recover time: 300 seconds
Enabled: Yes
```

```
Flow detection configuration:
```

```
Detection mode: Automatic Detect time: 3 seconds
Log flows: No Recover time: 60 seconds
Timeout flows: No Timeout time: 300 seconds
```

```
Flow aggregation level configuration:
```

```
Aggregation level Detection mode Control mode Flow rate
Subscriber Automatic Drop 10 pps
Logical interface Automatic Drop 10 pps
Physical interface Automatic Drop 2000 pps
```

```
System-wide information:
```

```
Aggregate bandwidth is never violated
Received: 0 Arrival rate: 0 pps
Dropped: 0 Max arrival rate: 0 pps
```

```
Routing Engine information:
```

```
Bandwidth: 2000 pps, Burst: 10000 packets, enabled
Aggregate policer is never violated
Received: 0 Arrival rate: 0 pps
Dropped: 0 Max arrival rate: 0 pps
```

```
Dropped by individual policers: 0
```

```
FPC slot 1 information:
```

```
Bandwidth: 100% (2000 pps), Burst: 100% (10000 packets), enabled
Aggregate policer is never violated
Received: 0 Arrival rate: 0 pps
Dropped: 0 Max arrival rate: 0 pps
```

```
Dropped by individual policers: 0
```

```
Dropped by flow suppression: 0
```

```
...
```

```
Protocol Group: PPPoE
```

```

Packet type: aggregate (Aggregate for all PPPoE control traffic)
Aggregate policer configuration:
 Bandwidth: 2000 pps
 Burst: 2000 packets
 Recover time: 300 seconds
 Enabled: Yes
Flow detection configuration:
 Detection mode: Automatic Detect time: 3 seconds
 Log flows: No Recover time: 60 seconds
 Timeout flows: No Timeout time: 300 seconds
Flow aggregation level configuration:
 Aggregation level Detection mode Control mode Flow rate
 Subscriber Automatic Drop 10 pps
 Logical interface Automatic Drop 10 pps
 Physical interface Automatic Drop 2000 pps
System-wide information:
 Aggregate bandwidth is never violated
 Received: 0 Arrival rate: 0 pps
 Dropped: 0 Max arrival rate: 0 pps
Routing Engine information:
 Bandwidth: 2000 pps, Burst: 2000 packets, enabled
 Aggregate policer is never violated
 Received: 0 Arrival rate: 0 pps
 Dropped: 0 Max arrival rate: 0 pps
 Dropped by individual policers: 0
FPC slot 1 information:
 Bandwidth: 100% (2000 pps), Burst: 100% (2000 packets), enabled
 Aggregate policer is never violated
 Received: 0 Arrival rate: 0 pps
 Dropped: 0 Max arrival rate: 0 pps
 Dropped by individual policers: 0
 Dropped by flow suppression: 0

Packet type: padi (PPPoE PADI)
Individual policer configuration:
 Bandwidth: 500 pps
 Burst: 500 packets
 Priority: Low
 Recover time: 300 seconds
 Enabled: Yes
 Bypass aggregate: No
Flow detection configuration:
 Detection mode: Automatic Detect time: 3 seconds
 Log flows: No Recover time: 60 seconds
 Timeout flows: No Timeout time: 300 seconds
Flow aggregation level configuration:
 Aggregation level Detection mode Control mode Flow rate
 Subscriber Automatic Drop 10 pps
 Logical interface Automatic Drop 10 pps
 Physical interface Automatic Drop 500 pps
System-wide information:
 Bandwidth is never violated
 Received: 0 Arrival rate: 0 pps
 Dropped: 0 Max arrival rate: 0 pps
Routing Engine information:
 Bandwidth: 500 pps, Burst: 500 packets, enabled
 Policer is never violated
 Received: 0 Arrival rate: 0 pps
 Dropped: 0 Max arrival rate: 0 pps
 Dropped by aggregate policer: 0

```

```

FPC slot 1 information:
 Bandwidth: 100% (500 pps), Burst: 100% (500 packets), enabled
 Policer is never violated
 Received: 0 Arrival rate: 0 pps
 Dropped: 0 Max arrival rate: 0 pps
 Dropped by aggregate policer: 0
 Dropped by flow suppression: 0
...

```

### show ddos-protection protocols (Specific Packet Type with Flow Detection Disabled)

```

user@host> show ddos-protection protocols pppoe padi
Currently tracked flows: 0, Total detected flows: 0
* = User configured value

Protocol Group: PPPoE

Packet type: padi (PPPoE PADI)
Individual policer configuration:
 Bandwidth: 500 pps
 Burst: 500 packets
 Priority: Low
 Recover time: 300 seconds
 Enabled: Yes
 Bypass aggregate: No
Flow detection configuration:
 Detection mode: Off* Detect time: 3 seconds
 Log flows: No Recover time: 60 seconds
 Timeout flows: No Timeout time: 300 seconds
Flow aggregation level configuration:
 Aggregation level Detection mode Control mode Flow rate
 Subscriber Automatic Drop 10 pps
 Logical interface Automatic Drop 10 pps
 Physical interface Automatic Drop 500 pps
System-wide information:
 Bandwidth is never violated
 Received: 0 Arrival rate: 0 pps
 Dropped: 0 Max arrival rate: 0 pps
Routing Engine information:
 Bandwidth: 500 pps, Burst: 500 packets, enabled
 Policer is never violated
 Received: 0 Arrival rate: 0 pps
 Dropped: 0 Max arrival rate: 0 pps
 Dropped by aggregate policer: 0
FPC slot 1 information:
 Bandwidth: 100% (500 pps), Burst: 100% (500 packets), enabled
 Policer is never violated
 Received: 0 Arrival rate: 0 pps
 Dropped: 0 Max arrival rate: 0 pps
 Dropped by aggregate policer: 0
 Dropped by flow suppression: 0

```

### show ddos-protection protocols (Specific Packet Type with Flow Detection Enabled and Automatic)

```

user@host> show ddos-protection protocols pppoe padi
Currently tracked flows: 0, Total detected flows: 0
* = User configured value

Protocol Group: PPPoE

Packet type: padi (PPPoE PADI)

```

```

Individual policer configuration:
 Bandwidth: 500 pps
 Burst: 500 packets
 Priority: Low
 Recover time: 300 seconds
 Enabled: Yes
 Bypass aggregate: No
Flow detection configuration:
 Detection mode: Automatic Detect time: 3 seconds
 Log flows: No Recover time: 60 seconds
 Timeout flows: No Timeout time: 300 seconds
Flow aggregation level configuration:
 Aggregation level Detection mode Control mode Flow rate
 Subscriber Automatic Drop 10 pps
 Logical interface Automatic Drop 10 pps
 Physical interface Automatic Drop 500 pps
System-wide information:
 Bandwidth is never violated
 Received: 0 Arrival rate: 0 pps
 Dropped: 0 Max arrival rate: 0 pps
Routing Engine information:
 Bandwidth: 500 pps, Burst: 500 packets, enabled
 Policer is never violated
 Received: 0 Arrival rate: 0 pps
 Dropped: 0 Max arrival rate: 0 pps
 Dropped by aggregate policer: 0
FPC slot 1 information:
 Bandwidth: 100% (500 pps), Burst: 100% (500 packets), enabled
 Policer is never violated
 Received: 0 Arrival rate: 0 pps
 Dropped: 0 Max arrival rate: 0 pps
 Dropped by aggregate policer: 0
 Dropped by flow suppression: 0

```

### show ddos-protection protocols (Specific Packet Type with Bandwidth Violation)

```

user@host> show ddos-protection protocols bfd
Packet types: 1, Modified: 0, Received traffic: 1, Currently violated: 1
Currently tracked flows: 1, Total detected flows: 1
* = User configured value

```

Protocol Group: BFD

```

Packet type: aggregate (Aggregate for all bfd traffic)
Aggregate policer configuration:
 Bandwidth: 20000 pps
 Burst: 20000 packets
 Recover time: 300 seconds
 Enabled: Yes
Flow detection configuration:
 Detection mode: Automatic Detect time: 3 seconds
 Log flows: No Recover time: 60 seconds
 Timeout flows: No Timeout time: 300 seconds
Flow aggregation level configuration:
 Aggregation level Detection mode Control mode Flow rate
 Subscriber Automatic Drop 10 pps
 Logical interface Automatic Drop 10 pps
 Physical interface Automatic Drop 20000 pps
System-wide information:
 Aggregate bandwidth is being violated!
 No. of FPCs currently receiving excess traffic: 1

```

**No. of FPCs that have received excess traffic: 1**

Violation first detected at: 2012-10-24 23:40:20 EDT

Violation last seen at: 2012-10-25 10:25:48 EDT

Duration of violation: 10:45:28 Number of violations: 1

Received: 1173471731 Arrival rate: 30304 pps

Dropped: 399135607 Max arrival rate: 30331 pps

**Flow counts:**

| Aggregation level | Current | Total detected |
|-------------------|---------|----------------|
| Subscriber        | 1       | 1              |
| Total             | 1       | 1              |

**Routing Engine information:**

Bandwidth: 20000 pps, Burst: 20000 packets, enabled

Aggregate policer is never violated

Received: 366831604 Arrival rate: 0 pps

Dropped: 0 Max arrival rate: 9522 pps

Dropped by individual policers: 0

**FPC slot 1 information:****Bandwidth: 100% (20000 pps), Burst: 100% (20000 packets), enabled****Aggregate policer is currently being violated!**

Violation first detected at: 2012-10-24 23:40:21 EDT

Violation last seen at: 2012-10-25 10:25:48 EDT

Duration of violation: 10:45:27 Number of violations: 1

Received: 1173471731 Arrival rate: 30304 pps

Dropped: 399135607 Max arrival rate: 30331 pps

Dropped by individual policers: 0

Dropped by aggregate policer: 398854530

Dropped by flow suppression: 281077

**Flow counts:**

| Aggregation level  | Current | Total detected | State  |
|--------------------|---------|----------------|--------|
| Subscriber         | 1       | 1              | Active |
| Logical-interface  | 0       | 0              | Active |
| Physical-interface | 0       | 0              | Active |
| Total              | 1       | 1              |        |



## show ddos-protection protocols parameters

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show ddos-protection protocols &lt;protocol-group&gt; parameters</code><br><code>&lt;brief   detail   terse&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 11.2.</p> <p>Command introduced in Junos OS Release 12.3R2 on EX9200 switches and T4000 routers.</p> <p>Command introduced in Junos OS Release 15.1X53 on QFX10000 switches.</p> <p>Command introduced in Junos OS Release 15.1X53-D30 on QFX5200 switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Display DDoS protection configuration information for all protocol groups or for a particular protocol group.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <p><b>none</b>—Display information for all protocol groups.</p> <p><b>brief   detail   terse</b>—(Optional) Display the specified level of output.</p> <ul style="list-style-type: none"> <li><b>brief</b>—Display basic function information.</li> <li><b>detail</b>—Add information to the <b>brief</b> output; it is identical to the output displayed when you choose no option. The <b>brief</b> and <b>detail</b> options display information for all protocol groups, which can be a long list.</li> <li><b>terse</b>—Display the same level of information as the <b>brief</b> option but only for active protocol groups—groups that show traffic in the <b>Received (packets)</b> column.</li> </ul> <p><b>protocol-group</b>—(Optional) Display information for a particular protocol group. See <a href="#">show ddos-protection protocols</a> for a list of available groups.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">clear ddos-protection protocols on page 6180</a></li> <li><a href="#">show ddos-protection protocols on page 6182</a></li> <li><a href="#">show ddos-protection protocols culprit-flows</a></li> <li><a href="#">show ddos-protection protocols flow-detection</a></li> <li><a href="#">show ddos-protection protocols statistics on page 6208</a></li> <li><a href="#">show ddos-protection protocols violations on page 6221</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>List of Sample Output</b>    | <a href="#">show ddos-protection protocols parameters on page 6203</a><br><a href="#">show ddos-protection protocols parameters brief on page 6204</a><br><a href="#">show ddos-protection protocols dhcpv4 parameters brief on page 6205</a><br><a href="#">show ddos-protection protocols dhcpv4 parameters terse on page 6206</a><br><a href="#">show ddos-protection protocols dhcpv4 parameters on page 6206</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Output Fields</b>            | <a href="#">Table 494</a> lists the output fields for the <b>show ddos-protection protocols parameters</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

Table 494: show ddos-protection protocols parameters Output Fields

| Field Name                         | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                 | Level of Output    |
|------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------|
| <b>Protocol Group</b>              | Name of protocol group.                                                                                                                                                                                                                                                                                                                                                                                                                           | All levels         |
| <b>Packet type</b>                 | Name of packet type in protocol group.                                                                                                                                                                                                                                                                                                                                                                                                            | All levels         |
| <b>Bandwidth</b>                   | Bandwidth policer value; number of packets per second that is allowed before a violation is declared.<br><br>In the <b>brief</b> output, an asterisk indicates the value has been modified from the default.                                                                                                                                                                                                                                      | All levels         |
| <b>Burst</b>                       | Burst policer value; the maximum number of packets that is allowed in a burst before a violation is declared.<br><br>In the <b>brief</b> output, an asterisk indicates the value has been modified from the default.                                                                                                                                                                                                                              | All levels         |
| <b>Priority</b>                    | Priority of the packet type in the event of traffic congestion: <b>low</b> , <b>medium</b> , or <b>high</b> . Lower priority packets can be dropped when insufficient bandwidth is available.<br><br>In the <b>brief</b> output, an asterisk indicates the value has been modified from the default.                                                                                                                                              | All levels         |
| <b>Recover time</b>                | Time that must pass since the last violation before the traffic flow is considered to have recovered from the attack. A notification is generated when the timer expires.<br><br>In the <b>brief</b> output, an asterisk indicates the value has been modified from the default.                                                                                                                                                                  | All levels         |
| <b>Enabled</b>                     | State of the policer, enabled ( <b>Yes</b> ) or disabled ( <b>No</b> ).                                                                                                                                                                                                                                                                                                                                                                           | <b>detail none</b> |
| <b>Bypass aggregate</b>            | State of the bypass aggregate configuration:<br><ul style="list-style-type: none"><li>• Yes—The aggregate policer is bypassed.</li><li>• No—The aggregate policer is enforced.</li></ul> This field appears only for individual policers.                                                                                                                                                                                                         | <b>detail none</b> |
| <b>FPC slot information</b>        | The following configuration information for the card in the indicated slot:<br><ul style="list-style-type: none"><li>• Bandwidth—Bandwidth scale and the number of packets per second that is allowed before a violation is declared</li><li>• Burst—Burst scale and the maximum number of packets that is allowed in a burst before a violation is declared</li><li>• <b>enabled</b> or <b>disabled</b>—State of the line card policer</li></ul> | <b>detail none</b> |
| <b>Number of policers modified</b> | Number of policers that have been changed from the default configuration.<br><br>An asterisk by a particular value indicates that value has been modified.                                                                                                                                                                                                                                                                                        | <b>brief terse</b> |
| <b>Policer Enabled</b>             | State of the policer, enabled ( <b>Yes</b> ), disabled ( <b>No</b> ), or partially disabled ( <b>part.</b> ); <b>part.</b> indicates that only some of the policer instances are disabled for the policer.                                                                                                                                                                                                                                        | <b>brief terse</b> |

Table 494: show ddos-protection protocols parameters Output Fields (*continued*)

| Field Name          | Field Description                                                                                                                                                                                                                                                                                                              | Level of Output    |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------|
| <b>Bypass aggr.</b> | <p>State of the bypass aggregate configuration:</p> <ul style="list-style-type: none"> <li>• Yes—The aggregate policer is bypassed.</li> <li>• No—The aggregate policer is enforced.</li> </ul> <p>Dashes indicate that the bypass aggregate configuration is not available; this is possible only for aggregate policers.</p> | <b>brief terse</b> |
| <b>FPC Mod</b>      | <p>Indicates whether configuration has changed from the default for any line cards.</p> <ul style="list-style-type: none"> <li>• No—The default configuration has not changed from the default for the packet type.</li> <li>• Yes—The default configuration has changed from the default for the packet type</li> </ul>       | <b>brief terse</b> |

## Sample Output

### show ddos-protection protocols parameters

```

user@host> show ddos-protection protocols parameters
Protocol Group: IPv4-Unclassified

Packet type: aggregate (Aggregate for unclassified host-bound IPv4 traffic)
Aggregate policer configuration:
 Bandwidth: 20000 pps
 Burst: 20000 packets
 Priority: medium
 Recover time: 300 seconds
 Enabled: Yes
FPC slot 1 information:
 Bandwidth: 100% (20000 pps), Burst: 100% (20000 packets), enabled

Protocol Group: IPv6-Unclassified

Packet type: aggregate (Aggregate for unclassified host-bound IPv6 traffic)
Aggregate policer configuration:
 Bandwidth: 20000 pps
 Burst: 20000 packets
 Priority: medium
 Recover time: 300 seconds
 Enabled: Yes
FPC slot 1 information:
 Bandwidth: 100% (20000 pps), Burst: 100% (20000 packets), enabled

...

Protocol Group: PPPoE

Packet type: aggregate (Aggregate for all PPPoE control traffic)
Aggregate policer configuration:
 Bandwidth: 800 pps
 Burst: 2000 packets
 Priority: medium
 Recover time: 300 seconds
 Enabled: Yes

```

## FPC slot 1 information:

Bandwidth: 100% (800 pps), Burst: 100% (2000 packets), enabled

## Packet type: padi (PPPoE PADI)

## Individual policer configuration:

Bandwidth: 500 pps  
 Burst: 500 packets  
 Priority: low  
 Recover time: 300 seconds  
 Enabled: Yes  
 Bypass aggregate: No

## FPC slot 1 information:

Bandwidth: 100% (500 pps), Burst: 100% (500 packets), enabled

## Packet type: pado (PPPoE PADO)

## Individual policer configuration:

Bandwidth: 0 pps  
 Burst: 0 packets  
 Priority: low  
 Recover time: 300 seconds  
 Enabled: Yes  
 Bypass aggregate: No

## FPC slot 1 information:

Bandwidth: 100% (0 pps), Burst: 100% (0 packets), enabled

## Packet type: padr (PPPoE PADR)

## Individual policer configuration:

Bandwidth: 500 pps  
 Burst: 500 packets  
 Priority: medium  
 Recover time: 300 seconds  
 Enabled: Yes  
 Bypass aggregate: No

## FPC slot 1 information:

Bandwidth: 100% (500 pps), Burst: 100% (500 packets), enabled

**show ddos-protection protocols parameters brief**

user@host&gt; show ddos-protection protocols parameters brief

Number of policers modified: 3

| Protocol group | Packet type | Bandwidth (pps) | Burst (pkts) | Priority | Recover time(sec) | Policer enabled | Bypass aggr. | FPC mod |
|----------------|-------------|-----------------|--------------|----------|-------------------|-----------------|--------------|---------|
| ipv4-uncls     | aggregate   | 20000           | 20000        | medium   | 300               | yes             | --           | no      |
| ipv6-uncls     | aggregate   | 20000           | 20000        | medium   | 300               | yes             | --           | no      |
| dynvlan        | aggregate   | 1000            | 500          | low      | 300               | yes             | --           | no      |
| ppp            | aggregate   | 16000           | 16000        | medium   | 300               | yes             | --           | no      |
| ppp            | unclass     | 1000            | 500          | low      | 300               | yes             | no           | no      |
| ppp            | lcp         | 12000           | 12000        | low      | 300               | yes             | no           | no      |
| ppp            | auth        | 2000            | 2000         | medium   | 300               | yes             | no           | no      |
| ppp            | ipcp        | 2000            | 2000         | high     | 300               | yes             | no           | no      |
| ppp            | ipv6cp      | 2000            | 2000         | high     | 300               | yes             | no           | no      |
| ppp            | mplscp      | 2000            | 2000         | high     | 300               | yes             | no           | no      |
| ppp            | isis        | 2000            | 2000         | high     | 300               | yes             | no           | no      |
| pppoe          | aggregate   | 800*            | 2000         | medium   | 300               | part.*          | --           | no      |
| pppoe          | padi        | 500             | 500          | low      | 300               | part.           | no           | no      |
| pppoe          | pado        | 0               | 0            | low      | 300               | part.           | no           | no      |
| pppoe          | padr        | 500             | 500          | medium   | 300               | part.           | no           | no      |
| pppoe          | pads        | 0               | 0            | low      | 300               | part.           | no           | no      |
| pppoe          | padt        | 1000            | 1000         | high     | 300               | part.           | no           | no      |
| pppoe          | padm        | 0               | 0            | low      | 300               | part.           | no           | no      |
| pppoe          | padn        | 0               | 0            | low      | 300               | part.           | no           | no      |

```

dhcpv4 aggregate 669* 5000 medium 300 yes -- no
dhcpv4 unclass.. 300 150 low 300 yes no no
dhcpv4 discover 100* 500 low 300 yes no no
dhcpv4 offer 1000 1000 low 300 yes no no
dhcpv4 request 1000 1000 medium 300 yes no no
dhcpv4 decline 500 500 low 300 yes no no
dhcpv4 ack 500 500 medium 300 yes no no
dhcpv4 nak 500 500 low 300 yes no no
dhcpv4 release 2000 2000 high 300 yes no no
dhcpv4 inform 500 500 low 300 yes no no
dhcpv4 renew 2000 2000 high 300 yes no no
dhcpv4 forcerenew 2000 2000 high 300 yes no no
dhcpv4 leasequery 2000 2000 high 300 yes no no
dhcpv4 leaseuna.. 2000 2000 high 300 yes no no
dhcpv4 leaseunk.. 2000 2000 high 300 yes no no
dhcpv4 leaseact.. 2000 2000 high 300 yes no no
dhcpv4 bootp 300 300 low 300 yes no no
dhcpv4 no-msgtype 0 0 low 300 yes no no
dhcpv4 bad-pack.. 0 0 low 300 yes no no

...

icmp aggregate 20000 20000 high 300 yes -- no
igmp aggregate 20000 20000 high 300 yes -- no
ospf aggregate 20000 20000 high 300 yes -- no
rsvp aggregate 20000 20000 high 300 yes -- no
pim aggregate 20000 20000 high 300 yes -- no
rip aggregate 20000 20000 high 300 yes -- no
ptp aggregate 20000 20000 high 300 yes -- no
bfd aggregate 20000 20000 high 300 yes -- no
lmp aggregate 20000 20000 high 300 yes -- no
ldp aggregate 20000 20000 high 300 yes -- no
msdp aggregate 20000 20000 high 300 yes -- no
bgp aggregate 20000 20000 low 300 yes -- no
vrrp aggregate 20000 20000 high 300 yes -- no
telnet aggregate 20000 20000 low 300 yes -- no
ftp aggregate 20000 20000 low 300 yes -- no
ssh aggregate 20000 20000 low 300 yes -- no
snmp aggregate 20000 20000 low 300 yes -- no
anycp aggregate 20000 20000 low 300 yes -- no

...

```

### show ddos-protection protocols dhcpv4 parameters brief

```

user@host> show ddos-protection protocols dhcpv4 parameters brief
Number of policers modified: 2

```

| Protocol | Packet group | Bandwidth (pps) | Burst (pkts) | Priority | Recover time(sec) | Policer enabled | Bypass aggr. | FPC mod |
|----------|--------------|-----------------|--------------|----------|-------------------|-----------------|--------------|---------|
| dhcpv4   | aggregate    | 669*            | 5000         | medium   | 300               | yes             | --           | no      |
| dhcpv4   | unclass..    | 300             | 150          | low      | 300               | yes             | no           | no      |
| dhcpv4   | discover     | 100*            | 500          | low      | 300               | yes             | no           | no      |
| dhcpv4   | offer        | 1000            | 1000         | low      | 300               | yes             | no           | no      |
| dhcpv4   | request      | 1000            | 1000         | medium   | 300               | yes             | no           | no      |
| dhcpv4   | decline      | 500             | 500          | low      | 300               | yes             | no           | no      |
| dhcpv4   | ack          | 500             | 500          | medium   | 300               | yes             | no           | no      |
| dhcpv4   | nak          | 500             | 500          | low      | 300               | yes             | no           | no      |
| dhcpv4   | release      | 2000            | 2000         | high     | 300               | yes             | no           | no      |
| dhcpv4   | inform       | 500             | 500          | low      | 300               | yes             | no           | no      |
| dhcpv4   | renew        | 2000            | 2000         | high     | 300               | yes             | no           | no      |
| dhcpv4   | forcerenew   | 2000            | 2000         | high     | 300               | yes             | no           | no      |

|        |            |      |      |      |     |     |    |    |
|--------|------------|------|------|------|-----|-----|----|----|
| dhcpv4 | leasequery | 2000 | 2000 | high | 300 | yes | no | no |
| dhcpv4 | leaseuna.. | 2000 | 2000 | high | 300 | yes | no | no |
| dhcpv4 | leaseunk.. | 2000 | 2000 | high | 300 | yes | no | no |
| dhcpv4 | leaseact.. | 2000 | 2000 | high | 300 | yes | no | no |
| dhcpv4 | bootp      | 300  | 300  | low  | 300 | yes | no | no |
| dhcpv4 | no-msgtype | 0    | 0    | low  | 300 | yes | no | no |
| dhcpv4 | bad-pack.. | 0    | 0    | low  | 300 | yes | no | no |

### show ddos-protection protocols dhcpv4 parameters terse

```
user@host> show ddos-protection protocols dhcpv4 parameters terse
```

Number of policers modified: 2

| Protocol group | Packet type | Bandwidth (pps) | Burst (pkts) | Priority | Recover time(sec) | Policer enabled | Bypass aggr. | FPC mod |
|----------------|-------------|-----------------|--------------|----------|-------------------|-----------------|--------------|---------|
| dhcpv4         | aggregate   | 669*            | 5000         | medium   | 300               | yes             | --           | no      |
| dhcpv4         | discover    | 100*            | 500          | low      | 300               | yes             | no           | no      |

### show ddos-protection protocols dhcpv4 parameters

```
user@host> show ddos-protection protocols dhcpv4 parameters
```

Protocol Group: DHCPv4

Packet type: aggregate (aggregate for all DHCPv4 traffic)

Aggregate policer configuration:

Bandwidth: 669 pps  
 Burst: 5000 packets  
 Priority: medium  
 Recover time: 300 seconds  
 Enabled: Yes

FPC slot 1 information:

Bandwidth: 100% (669 pps), Burst: 100% (5000 packets), enabled

Packet type: unclassified (Unclassified DHCPv4 traffic)

Individual policer configuration:

Bandwidth: 300 pps  
 Burst: 150 packets  
 Priority: low  
 Recover time: 300 seconds  
 Enabled: Yes  
 Bypass aggregate: No

FPC slot 1 information:

Bandwidth: 100% (300 pps), Burst: 100% (150 packets), enabled

Packet type: discover (DHCPv4 DHCPDISCOVER)

Individual policer configuration:

Bandwidth: 100 pps  
 Burst: 500 packets  
 Priority: low  
 Recover time: 300 seconds  
 Enabled: Yes  
 Bypass aggregate: No

FPC slot 1 information:

Bandwidth: 100% (100 pps), Burst: 100% (500 packets), enabled

Packet type: offer (DHCPv4 DHCPOFFER)

Individual policer configuration:

Bandwidth: 1000 pps  
 Burst: 1000 packets  
 Priority: low  
 Recover time: 300 seconds  
 Enabled: Yes

```
 Bypass aggregate: No
FPC slot 1 information:
 Bandwidth: 100% (1000 pps), Burst: 100% (1000 packets), enabled
```

```
Packet type: request (DHCPv4 DHCPREQUEST)
```

```
Individual policer configuration:
```

```
 Bandwidth: 1000 pps
 Burst: 1000 packets
 Priority: medium
 Recover time: 300 seconds
 Enabled: Yes
 Bypass aggregate: No
```

```
FPC slot 1 information:
 Bandwidth: 100% (1000 pps), Burst: 100% (1000 packets), enabled
```

```
...
```

## show ddos-protection protocols statistics

---

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <b>show ddos-protection protocols</b> < <i>protocol-group</i> > <b>statistics</b><br>< <b>brief</b>   <b>detail</b>   <b>terse</b> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Release Information      | Command introduced in Junos OS Release 11.2.<br>Command introduced in Junos OS Release 12.3R2 on EX9200 switches and T4000 routers.<br>Command introduced in Junos OS Release 15.1X53 on QFX10000 switches.<br>Command introduced in Junos OS Release 15.1X53-D30 on QFX5200 switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Description              | Display traffic statistics and DDoS policer violation statistics for all protocol groups or for a particular protocol group.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Options                  | <b>none</b> —Display information for all protocol groups.<br><br><b>brief</b>   <b>detail</b>   <b>terse</b> —(Optional) Display the specified level of output. <ul style="list-style-type: none"><li>• <b>brief</b>—Display basic function information.</li><li>• <b>detail</b>—Add information to the <b>brief</b> output; it is identical to the output displayed when you choose no option. The <b>brief</b> and <b>detail</b> options display information for all protocol groups, which can be a long list.</li><li>• <b>terse</b>—Display the same level of information as the <b>brief</b> option but only for active protocol groups—groups that show traffic in the <b>Received (packets)</b> column.</li></ul><br><b>protocol-group</b> —(Optional) Display information for a particular protocol group. See <a href="#">show ddos-protection protocols</a> for a list of available groups. |
| Required Privilege Level | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Related Documentation    | <ul style="list-style-type: none"><li>• <a href="#">clear ddos-protection protocols on page 6180</a></li><li>• <a href="#">show ddos-protection protocols on page 6182</a></li><li>• <i>show ddos-protection protocols culprit-flows</i></li><li>• <i>show ddos-protection protocols flow-detection</i></li><li>• <a href="#">show ddos-protection protocols parameters on page 6201</a></li><li>• <a href="#">show ddos-protection protocols violations on page 6221</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                    |
| List of Sample Output    | <a href="#">show ddos-protection protocols statistics on page 6210</a><br><a href="#">show ddos-protection protocols statistics brief on page 6213</a><br><a href="#">show ddos-protection protocols statistics terse on page 6214</a><br><a href="#">show ddos-protection protocols pppoe statistics on page 6215</a><br><a href="#">show ddos-protection protocols pppoe statistics brief on page 6217</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Output Fields            | <a href="#">Table 495</a> lists the output fields for the <b>show ddos-protection protocols statistics</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |



Table 495: show ddos-protection protocols statistics Output Fields

| Field Name                        | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Level of Output |
|-----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Protocol Group</b>             | Name of protocol group.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | All levels      |
| <b>Packet type</b>                | Name of packet type in protocol group.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | All levels      |
| <b>System-wide information</b>    | <p>The following information collected for the router:</p> <ul style="list-style-type: none"> <li>• A message indicates whether the policer has been violated.</li> <li>• No. of FPCs currently receiving excess traffic—Number of cards that are currently in violation of a policer.</li> <li>• No. of FPCs that have received excess traffic—Number of cards that have at some point been in violation of a policer.</li> <li>• Violation first detected at—Timestamp of the first violation.</li> <li>• Violation last seen at—Timestamp of the last observed violation.</li> <li>• Duration of violation—Length of the violation.</li> <li>• Number of violations—Number of times the violation has occurred.</li> <li>• Received—Number of packets received at all card slots and the Routing Engine.</li> <li>• Dropped—Number of packets dropped regardless of where they were dropped.</li> <li>• Arrival rate—Current traffic rate for packets arriving from all cards and at the Routing Engine.</li> <li>• Max arrival rate—Highest traffic rate for packets arriving from all cards and at the Routing Engine.</li> </ul>                                                                                                                                                                | detail none     |
| <b>Routing Engine information</b> | <p>The following information collected for the Routing Engine:</p> <ul style="list-style-type: none"> <li>• A message indicates whether the policer has been violated; the policer might be passed at the individual cards, but the combined rate of packets arriving at the Routing Engine can exceed the configured policer value.</li> <li>• Violation first detected at—Timestamp of the first violation.</li> <li>• Violation last seen at—Timestamp of the last observed violation.</li> <li>• Duration of violation—Length of the violation.</li> <li>• Number of violations—Number of times the violation has occurred.</li> <li>• Received—Number of packets received at the Routing Engine from all cards.</li> <li>• Dropped—Number of packets dropped at the Routing Engine; includes packets dropped by the aggregate policer and by individual protocol policers.</li> <li>• Arrival rate—Current traffic rate for packets arriving at the Routing Engine from all cards.</li> <li>• Max arrival rate—Highest traffic rate for packets arriving at the Routing Engine from all cards.</li> <li>• Dropped by aggregate policer—Number of packets dropped by the aggregate policer.</li> <li>• Dropped by individual policers—Number of packets dropped by individual policer.</li> </ul> | detail none     |

Table 495: show ddos-protection protocols statistics Output Fields (*continued*)

| Field Name                  | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Level of Output    |
|-----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------|
| <b>FPC slot information</b> | <p>The following information collected for the card in the indicated slot:</p> <ul style="list-style-type: none"> <li>• A message indicates whether the policer has been violated</li> <li>• Violation first detected at—Timestamp of the first violation</li> <li>• Violation last seen at—Timestamp of the last observed violation</li> <li>• Duration of violation—Length of the violation</li> <li>• Number of violations—Number of times the violation has occurred</li> <li>• Received—Number of packets received on the line card</li> <li>• Dropped—Number of packets dropped at the line card; includes packets dropped by the aggregate policer and by individual protocol policers</li> <li>• Arrival rate—Current traffic rate for packets arriving at the line card</li> <li>• Max arrival rate—Highest traffic rate for packets arriving at the line card</li> <li>• Dropped by this policer—Number of packets dropped by the individual policer</li> <li>• Dropped by aggregate policer—Number of packets dropped by the aggregate policer</li> </ul> | <b>detail none</b> |
| <b>Received (packets)</b>   | Number of packets of this packet type or protocol group received at all cards and the Routing Engine.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>brief terse</b> |
| <b>Dropped (packets)</b>    | Number of packets dropped for this packet type or protocol group, regardless of where the packets were dropped.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>brief terse</b> |
| <b>Rate (pps)</b>           | Highest observed traffic rate for this packet type or protocol group.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>brief terse</b> |
| <b>Violation counts</b>     | Number of violations of the policer bandwidth.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>brief terse</b> |
| <b>State</b>                | <p>Violation state of the packet type:</p> <ul style="list-style-type: none"> <li>• <b>ok</b>—Policer has not been violated for this packet type</li> <li>• <b>viol</b>—Policer has been violated for this packet type</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>brief terse</b> |

## Sample Output

### show ddos-protection protocols statistics

```

user@host> show ddos-protection protocols statistics
Protocol Group: IPv4-Unclassified

Packet type: aggregate
System-wide information:
 Aggregate bandwidth is never violated
 Received: 0 Arrival rate: 0 pps
 Dropped: 0 Max arrival rate: 0 pps
Routing Engine information:
 Aggregate policer is never violated
 Received: 0 Arrival rate: 0 pps
 Dropped: 0 Max arrival rate: 0 pps
 Dropped by individual policers: 0
FPC slot 1 information:
 Aggregate policer is never violated
 Received: 0 Arrival rate: 0 pps

```

Dropped: 0 Max arrival rate: 0 pps  
 Dropped by individual policers: 0

Protocol Group: IPv6-Unclassified

Packet type: aggregate  
 System-wide information:  
 Aggregate bandwidth is never violated  
 Received: 0 Arrival rate: 0 pps  
 Dropped: 0 Max arrival rate: 0 pps  
 Routing Engine information:  
 Aggregate policer is never violated  
 Received: 0 Arrival rate: 0 pps  
 Dropped: 0 Max arrival rate: 0 pps  
 Dropped by individual policers: 0  
 FPC slot 1 information:  
 Aggregate policer is never violated  
 Received: 0 Arrival rate: 0 pps  
 Dropped: 0 Max arrival rate: 0 pps  
 Dropped by individual policers: 0

Protocol Group: PPPoE

Packet type: aggregate  
 System-wide information:  
 Aggregate bandwidth is never violated  
 Received: 61961244 Arrival rate: 4000 pps  
 Dropped: 0 Max arrival rate: 4002 pps  
 Routing Engine information:  
 Aggregate policer is never violated  
 Received: 15488871 Arrival rate: 1001 pps  
 Dropped: 0 Max arrival rate: 1011 pps  
 Dropped by individual policers: 0  
 FPC slot 1 information:  
 Aggregate policer is never violated  
 Received: 61961244 Arrival rate: 4000 pps  
 Dropped: 46473017 Max arrival rate: 4002 pps  
 Dropped by individual policers: 46473017

Packet type: padi  
 System-wide information:  
 Bandwidth is being violated!  
 No. of FPCs currently receiving excess traffic: 1  
 No. of FPCs that have received excess traffic: 1  
 Violation first detected at: 2011-04-19 08:23:17 PDT  
 Violation last seen at: 2011-04-19 12:41:23 PDT  
 Duration of violation: 04:18:06 Number of violations: 1  
 Received: 30980622 Arrival rate: 2000 pps  
 Dropped: 23236505 Max arrival rate: 2001 pps  
 Routing Engine information:  
 Policar is never violated  
 Received: 7744433 Arrival rate: 500 pps  
 Dropped: 0 Max arrival rate: 505 pps  
 Dropped by aggregate policer: 0  
 FPC slot 1 information:  
 Policar is currently being violated!  
 Violation first detected at: 2011-04-19 08:23:17 PDT  
 Violation last seen at: 2011-04-19 12:41:23 PDT  
 Duration of violation: 04:18:06 Number of violations: 1  
 Received: 30980622 Arrival rate: 2000 pps

Dropped: 23236505 Max arrival rate: 2001 pps  
Dropped by this policer: 23236505  
Dropped by aggregate policer: 0

Packet type: pado

System-wide information:

Bandwidth is never violated

Received: 0 Arrival rate: 0 pps

Dropped: 0 Max arrival rate: 0 pps

Routing Engine information:

Policer is never violated

Received: 0 Arrival rate: 0 pps

Dropped: 0 Max arrival rate: 0 pps

Dropped by aggregate policer: 0

FPC slot 1 information:

Policer is never violated

Received: 0 Arrival rate: 0 pps

Dropped: 0 Max arrival rate: 0 pps

Dropped by aggregate policer: 0

Packet type: padr

System-wide information:

Bandwidth is being violated!

No. of FPCs currently receiving excess traffic: 1

No. of FPCs that have received excess traffic: 1

Violation first detected at: 2011-04-19 08:23:17 PDT

Violation last seen at: 2011-04-19 12:43:23 PDT

Duration of violation: 04:20:06 Number of violations: 1

Received: 31220846 Arrival rate: 2000 pps

Dropped: 23416690 Max arrival rate: 2001 pps

Routing Engine information:

Policer is never violated

Received: 7806417 Arrival rate: 499 pps

Dropped: 0 Max arrival rate: 506 pps

Dropped by aggregate policer: 0

FPC slot 1 information:

Policer is currently being violated!

Violation first detected at: 2011-04-19 08:23:17 PDT

Violation last seen at: 2011-04-19 12:43:23 PDT

Duration of violation: 04:20:06 Number of violations: 1

Received: 31220846 Arrival rate: 2000 pps

Dropped: 23416690 Max arrival rate: 2001 pps

Dropped by this policer: 23416690

Dropped by aggregate policer: 0

Packet type: pads

System-wide information:

Bandwidth is never violated

Received: 0 Arrival rate: 0 pps

Dropped: 0 Max arrival rate: 0 pps

Routing Engine information:

Policer is never violated

Received: 0 Arrival rate: 0 pps

Dropped: 0 Max arrival rate: 0 pps

Dropped by aggregate policer: 0

FPC slot 1 information:

Policer is never violated

Received: 0 Arrival rate: 0 pps

Dropped: 0 Max arrival rate: 0 pps

Dropped by aggregate policer: 0

```

Packet type: padt
System-wide information:
 Bandwidth is never violated
 Received: 0 Arrival rate: 0 pps
 Dropped: 0 Max arrival rate: 0 pps
Routing Engine information:
 Policer is never violated
 Received: 0 Arrival rate: 0 pps
 Dropped: 0 Max arrival rate: 0 pps
 Dropped by aggregate policer: 0
FPC slot 1 information:
 Policer is never violated
 Received: 0 Arrival rate: 0 pps
 Dropped: 0 Max arrival rate: 0 pps
 Dropped by aggregate policer: 0

```

```

Packet type: padm
System-wide information:
 Bandwidth is never violated
 Received: 0 Arrival rate: 0 pps
 Dropped: 0 Max arrival rate: 0 pps
Routing Engine information:
 Policer is never violated
 Received: 0 Arrival rate: 0 pps
 Dropped: 0 Max arrival rate: 0 pps
 Dropped by aggregate policer: 0
FPC slot 1 information:
 Policer is never violated
 Received: 0 Arrival rate: 0 pps
 Dropped: 0 Max arrival rate: 0 pps
 Dropped by aggregate policer: 0

```

```

Packet type: padn
System-wide information:
 Bandwidth is never violated
 Received: 0 Arrival rate: 0 pps
 Dropped: 0 Max arrival rate: 0 pps
Routing Engine information:
 Policer is never violated
 Received: 0 Arrival rate: 0 pps
 Dropped: 0 Max arrival rate: 0 pps
 Dropped by aggregate policer: 0
FPC slot 1 information:
 Policer is never violated
 Received: 0 Arrival rate: 0 pps
 Dropped: 0 Max arrival rate: 0 pps
 Dropped by aggregate policer: 0

```

...

#### show ddos-protection protocols statistics brief

```
user@host> show ddos-protection protocols statistics brief
```

| Protocol group | Packet type | Received (packets) | Dropped (packets) | Rate (pps) | Violation counts | State |
|----------------|-------------|--------------------|-------------------|------------|------------------|-------|
| ipv4-unclass   | aggregate   | 0                  | 0                 | 0          | 0                | ok    |
| ipv6-unclass   | aggregate   | 0                  | 0                 | 0          | 0                | ok    |
| dynvlan        | aggregate   | 0                  | 0                 | 0          | 0                | ok    |
| ppp            | aggregate   | 0                  | 0                 | 0          | 0                | ok    |
| ppp            | unclass     | 0                  | 0                 | 0          | 0                | ok    |

```

ppp lcp 0 0 0 0 ok
ppp auth 0 0 0 0 ok
ppp ipcp 0 0 0 0 ok
ppp ipv6cp 0 0 0 0 ok
ppp mp1scp 0 0 0 0 ok
ppp isis 0 0 0 0 ok
pppoe aggregate 61561238 0 4000 0 ok
pppoe padi 30780619 23086506 2000 1 viol
pppoe pado 0 0 0 0 ok
pppoe padr 30780619 23086499 2000 1 viol
pppoe pads 0 0 0 0 ok
pppoe padt 0 0 0 0 ok
pppoe padm 0 0 0 0 ok
pppoe padn 0 0 0 0 ok
dhcipv4 aggregate 0 0 0 0 ok
dhcipv4 unclass.. 0 0 0 0 ok
dhcipv4 discover 0 0 0 0 ok
dhcipv4 offer 0 0 0 0 ok
dhcipv4 request 0 0 0 0 ok
dhcipv4 decline 0 0 0 0 ok
dhcipv4 ack 0 0 0 0 ok
dhcipv4 nak 0 0 0 0 ok
dhcipv4 release 0 0 0 0 ok
dhcipv4 inform 0 0 0 0 ok
dhcipv4 renew 0 0 0 0 ok
dhcipv4 forcerenew 0 0 0 0 ok
dhcipv4 leasequery 0 0 0 0 ok
dhcipv4 leaseuna.. 0 0 0 0 ok
dhcipv4 leaseunk.. 0 0 0 0 ok
dhcipv4 leaseact.. 0 0 0 0 ok
dhcipv4 bootp 0 0 0 0 ok
dhcipv4 no-msgtype 0 0 0 0 ok
dhcipv4 bad-pack.. 0 0 0 0 ok

...

icmp aggregate 0 0 0 0 ok
igmp aggregate 0 0 0 0 ok
ospf aggregate 0 0 0 0 ok
rsvp aggregate 0 0 0 0 ok
pim aggregate 0 0 0 0 ok
rip aggregate 0 0 0 0 ok
ptp aggregate 0 0 0 0 ok
bfd aggregate 0 0 0 0 ok
lmp aggregate 0 0 0 0 ok
ldp aggregate 0 0 0 0 ok
msdp aggregate 0 0 0 0 ok
bgp aggregate 0 0 0 0 ok
vrrp aggregate 0 0 0 0 ok
telnet aggregate 0 0 0 0 ok

...

```

#### show ddos-protection protocols statistics terse

```

user@host> show ddos-protection protocols statistics terse
Protocol Packet Received Dropped Rate Violation State
group type (packets) (packets) (pps) counts
ipv4-unc1s aggregate 241 0 0 0 ok
icmp aggregate 20 0 0 0 ok

```

|            |           |         |        |   |   |    |
|------------|-----------|---------|--------|---|---|----|
| igmp       | aggregate | 55      | 0      | 0 | 0 | ok |
| ospf       | aggregate | 956     | 0      | 0 | 0 | ok |
| rsvp       | aggregate | 784     | 0      | 0 | 0 | ok |
| ldp        | aggregate | 2984    | 0      | 0 | 0 | ok |
| bgp        | aggregate | 312     | 0      | 0 | 0 | ok |
| lACP       | aggregate | 1744    | 0      | 0 | 0 | ok |
| stp        | aggregate | 9791    | 0      | 0 | 0 | ok |
| arp        | aggregate | 19      | 0      | 0 | 0 | ok |
| pvstp      | aggregate | 393     | 0      | 0 | 0 | ok |
| m1p        | aggregate | 624774  | 0      | 0 | 0 | ok |
| m1p        | packets   | 1714371 | 223937 | 0 | 3 | ok |
| mcast-copy | aggregate | 3018038 | 0      | 0 | 0 | ok |
| igmp-snoop | aggregate | 43      | 0      | 0 | 0 | ok |
| fw-host    | aggregate | 95547   | 0      | 0 | 0 | ok |
| unc1s      | aggregate | 10000   | 0      | 0 | 0 | ok |

### show ddos-protection protocols pppoe statistics

```
user@host> show ddos-protection protocols pppoe statistics
Protocol Group: PPPoE
```

Packet type: aggregate

System-wide information:

Aggregate bandwidth is never violated

Received: 60381200 Arrival rate: 4000 pps

Dropped: 0 Max arrival rate: 4002 pps

Routing Engine information:

Aggregate policer is never violated

Received: 15095242 Arrival rate: 1001 pps

Dropped: 0 Max arrival rate: 1011 pps

Dropped by individual policers: 0

FPC slot 1 information:

Aggregate policer is never violated

Received: 60381200 Arrival rate: 4000 pps

Dropped: 45287921 Max arrival rate: 4002 pps

Dropped by individual policers: 45287921

Packet type: padi

System-wide information:

Bandwidth is being violated!

No. of FPCs currently receiving excess traffic: 1

No. of FPCs that have received excess traffic: 1

Violation first detected at: 2011-04-19 08:23:17 PDT

Violation last seen at: 2011-04-19 12:34:48 PDT

Duration of violation: 04:11:31 Number of violations: 1

Received: 30190600 Arrival rate: 2000 pps

Dropped: 22643960 Max arrival rate: 2001 pps

Routing Engine information:

Policer is never violated

Received: 7547621 Arrival rate: 499 pps

Dropped: 0 Max arrival rate: 505 pps

Dropped by aggregate policer: 0

FPC slot 1 information:

Policer is currently being violated!

Violation first detected at: 2011-04-19 08:23:17 PDT

Violation last seen at: 2011-04-19 12:34:48 PDT

Duration of violation: 04:11:31 Number of violations: 1

Received: 30190600 Arrival rate: 2000 pps

Dropped: 22643960 Max arrival rate: 2001 pps

Dropped by this policer: 22643960

Dropped by aggregate policer: 0

Packet type: pado  
System-wide information:  
Bandwidth is never violated  
Received: 0 Arrival rate: 0 pps  
Dropped: 0 Max arrival rate: 0 pps  
Routing Engine information:  
Policer is never violated  
Received: 0 Arrival rate: 0 pps  
Dropped: 0 Max arrival rate: 0 pps  
Dropped by aggregate policer: 0  
FPC slot 1 information:  
Policer is never violated  
Received: 0 Arrival rate: 0 pps  
Dropped: 0 Max arrival rate: 0 pps  
Dropped by aggregate policer: 0

Packet type: padr  
System-wide information:  
Bandwidth is being violated!  
No. of FPCs currently receiving excess traffic: 1  
No. of FPCs that have received excess traffic: 1  
Violation first detected at: 2011-04-19 08:23:17 PDT  
Violation last seen at: 2011-04-19 12:34:48 PDT  
Duration of violation: 04:11:31 Number of violations: 1  
Received: 30190600 Arrival rate: 2000 pps  
Dropped: 22643961 Max arrival rate: 2001 pps  
Routing Engine information:  
Policer is never violated  
Received: 7547621 Arrival rate: 501 pps  
Dropped: 0 Max arrival rate: 506 pps  
Dropped by aggregate policer: 0  
FPC slot 1 information:  
Policer is currently being violated!  
Violation first detected at: 2011-04-19 08:23:17 PDT  
Violation last seen at: 2011-04-19 12:34:48 PDT  
Duration of violation: 04:11:31 Number of violations: 1  
Received: 30190600 Arrival rate: 2000 pps  
Dropped: 22643961 Max arrival rate: 2001 pps  
Dropped by this policer: 22643961  
Dropped by aggregate policer: 0

Packet type: pads  
System-wide information:  
Bandwidth is never violated  
Received: 0 Arrival rate: 0 pps  
Dropped: 0 Max arrival rate: 0 pps  
Routing Engine information:  
Policer is never violated  
Received: 0 Arrival rate: 0 pps  
Dropped: 0 Max arrival rate: 0 pps  
Dropped by aggregate policer: 0  
FPC slot 1 information:  
Policer is never violated  
Received: 0 Arrival rate: 0 pps  
Dropped: 0 Max arrival rate: 0 pps  
Dropped by aggregate policer: 0

Packet type: padt  
System-wide information:  
Bandwidth is never violated



```

Received: 0 Arrival rate: 0 pps
Dropped: 0 Max arrival rate: 0 pps
Routing Engine information:
 Policer is never violated
Received: 0 Arrival rate: 0 pps
Dropped: 0 Max arrival rate: 0 pps
 Dropped by aggregate policer: 0
FPC slot 1 information:
 Policer is never violated
Received: 0 Arrival rate: 0 pps
Dropped: 0 Max arrival rate: 0 pps
 Dropped by aggregate policer: 0

Packet type: padm
System-wide information:
 Bandwidth is never violated
Received: 0 Arrival rate: 0 pps
Dropped: 0 Max arrival rate: 0 pps
:
 Policer is never violated
Received: 0 Arrival rate: 0 pps
Dropped: 0 Max arrival rate: 0 pps
 Dropped by aggregate policer: 0
FPC slot 1 information:
 Policer is never violated
Received: 0 Arrival rate: 0 pps
Dropped: 0 Max arrival rate: 0 pps
 Dropped by aggregate policer: 0

Packet type: padn
System-wide information:
 Bandwidth is never violated
Received: 0 Arrival rate: 0 pps
Dropped: 0 Max arrival rate: 0 pps
:
 Policer is never violated
Received: 0 Arrival rate: 0 pps
Dropped: 0 Max arrival rate: 0 pps
 Dropped by aggregate policer: 0
FPC slot 1 information:
 Policer is never violated
Received: 0 Arrival rate: 0 pps
Dropped: 0 Max arrival rate: 0 pps
 Dropped by aggregate policer: 0

```

#### show ddos-protection protocols pppoe statistics brief

```
user@host> show ddos-protection protocols pppoe statistics brief
```

| Protocol group | Packet type | Received (packets) | Dropped (packets) | Rate (pps) | Violation counts | State |
|----------------|-------------|--------------------|-------------------|------------|------------------|-------|
| pppoe          | aggregate   | 60901227           | 0                 | 4000       | 0                | ok    |
| pppoe          | padi        | 30450613           | 22838981          | 2000       | 1                | viol  |
| pppoe          | pado        | 0                  | 0                 | 0          | 0                | ok    |
| pppoe          | padr        | 30450614           | 22838977          | 2000       | 1                | viol  |
| pppoe          | pads        | 0                  | 0                 | 0          | 0                | ok    |
| pppoe          | padt        | 0                  | 0                 | 0          | 0                | ok    |
| pppoe          | padm        | 0                  | 0                 | 0          | 0                | ok    |
| pppoe          | padn        | 0                  | 0                 | 0          | 0                | ok    |

## show ddos-protection statistics

|                                 |                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show ddos-protection statistics</b>                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 11.2.</p> <p>Command introduced in Junos OS Release 12.3R2 on EX9200 switches and T4000 routers.</p> <p>Command introduced in Junos OS Release 15.1X53 on QFX10000 switches.</p> <p>Command introduced in Junos OS Release 15.1X53-D30 on QFX5200 switches.</p> |
| <b>Description</b>              | Display DDoS protection global statistics for bandwidth violations.                                                                                                                                                                                                                                       |
| <b>Options</b>                  | This command has no options.                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">clear ddos-protection protocols on page 6180</a></li> <li>• <a href="#">show ddos-protection protocols on page 6182</a></li> <li>• <a href="#">show ddos-protection version on page 6220</a></li> </ul>                                              |
| <b>List of Sample Output</b>    | <a href="#">show ddos-protection statistics on page 6219</a>                                                                                                                                                                                                                                              |
| <b>Output Fields</b>            | Table 496 lists the output fields for the <b>show ddos-protection statistics</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                         |

**Table 496: show ddos-protection statistics Output Fields**

| Field Name                        | Field Description                                                                                 |
|-----------------------------------|---------------------------------------------------------------------------------------------------|
| Policing on routing engine        | Shows whether or not policing is enabled on the Routing Engine.                                   |
| Policing on FPC                   | Shows whether or not policing is enabled on the line card.                                        |
| Flow detection                    | Shows whether or not flow detection is enabled.                                                   |
| Logging                           | Shows whether or not DDoS event logging is enabled.                                               |
| Policer violation report rate     | Shows the violation report rate as a percentage.                                                  |
| Flow report rate                  | Shows the flow report rate as a percentage.                                                       |
| Currently violated packet types   | Number of packet types currently experiencing a bandwidth violation.                              |
| Packet types have seen violations | Number of packet types that have experienced a bandwidth violation since statistics were cleared. |

Table 496: show ddos-protection statistics Output Fields (*continued*)

| Field Name             | Field Description                     |
|------------------------|---------------------------------------|
| Total violation counts | Total number of bandwidth violations. |

## Sample Output

### show ddos-protection statistics

```
user@host> show ddos-protection statistics
DDOS protection global statistics:

 Policing on routing engine: Yes
 Policing on FPC: Yes
 Flow detection: No
 Logging: Yes
 Policer violation report rate: 100
 Flow report rate: 100
 Currently violated packet types: 2
 Packet types have seen violations: 2
 Total violation counts: 2
 Currently tracked flows: 0
 Total detected flows: 0
```

## show ddos-protection version

|                                 |                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show ddos-protection version</b>                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 11.2.</p> <p>Command introduced in Junos OS Release 12.3R2 on EX9200 switches and T4000 routers.</p> <p>Command introduced in Junos OS Release 15.1X53 on QFX10000 switches.</p> <p>Command introduced in Junos OS Release 15.1X53-D30 on QFX5200 switches.</p> |
| <b>Description</b>              | Display the DDoS protection version and the total numbers of protocol groups and packet types that this version can be configured in this version.                                                                                                                                                        |
| <b>Options</b>                  | This command has no options.                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">clear ddos-protection protocols on page 6180</a></li> <li>• <a href="#">show ddos-protection protocols on page 6182</a></li> <li>• <a href="#">show ddos-protection statistics on page 6218</a></li> </ul>                                           |
| <b>List of Sample Output</b>    | <a href="#">show ddos-protection version on page 6220</a>                                                                                                                                                                                                                                                 |
| <b>Output Fields</b>            | Table 497 lists the output fields for the <b>show ddos-protection version</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                            |

**Table 497: show ddos-protection version Output Fields**

| Field Name                        | Field Description                                                |
|-----------------------------------|------------------------------------------------------------------|
| <b>Version</b>                    | Version number of the DDoS protection code.                      |
| <b>Total protocol groups</b>      | Number of protocol groups configured with DDoS protection.       |
| <b>Total tracked packet types</b> | Number of protocol packet types configured with DDoS protection. |

## Sample Output

### show ddos-protection version

```

user@host> show ddos-protection version
DDoS protection, Version 1.0
 Total protocol groups = 83
 Total tracked packet types = 154

```

## show ddos-protection protocols violations

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show ddos-protection protocols &lt;protocol-group&gt; violations</b>                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 11.2.</p> <p>Command introduced in Junos OS Release 12.3R2 on EX9200 switches and T4000 routers.</p> <p>Command introduced in Junos OS Release 15.1X53 on QFX10000 switches.</p> <p>Command introduced in Junos OS Release 15.1X53-D30 on QFX5200 switches.</p>                                                                                                                                                                                                    |
| <b>Description</b>              | Display information about DDoS policer violations for all protocol groups or for a particular protocol group.                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <p><b>none</b>—Display information for all protocol groups.</p> <p><b>protocol-group</b>—(Optional) Name of a particular protocol group. See <a href="#">show ddos-protection protocols</a> for a list of available groups.</p>                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">clear ddos-protection protocols on page 6180</a></li> <li>• <a href="#">show ddos-protection protocols on page 6182</a></li> <li>• <a href="#">show ddos-protection protocols culprit-flows</a></li> <li>• <a href="#">show ddos-protection protocols flow-detection</a></li> <li>• <a href="#">show ddos-protection protocols parameters on page 6201</a></li> <li>• <a href="#">show ddos-protection protocols statistics on page 6208</a></li> </ul> |
| <b>List of Sample Output</b>    | <p><a href="#">show ddos-protection protocols violations on page 6222</a></p> <p><a href="#">show ddos-protection protocols lldp violations on page 6222</a></p> <p><a href="#">show ddos-protection protocols pppoe violations on page 6222</a></p>                                                                                                                                                                                                                                                         |
| <b>Output Fields</b>            | Table 498 lists the output fields for the <b>show ddos-protection protocols violations</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                  |

**Table 498: show ddos-protection protocols violations Output Fields**

| Field Name                                     | Field Description                                                                      |
|------------------------------------------------|----------------------------------------------------------------------------------------|
| Number of packet types that are being violated | Number of individual policers and aggregate policers that are currently being violated |
| Protocol Group                                 | Name of protocol group                                                                 |
| Packet type                                    | Name of packet type in protocol group                                                  |
| Bandwidth (pps)                                | Policer bandwidth                                                                      |

Table 498: show ddos-protection protocols violations Output Fields (*continued*)

| Field Name                              | Field Description                                                                  |
|-----------------------------------------|------------------------------------------------------------------------------------|
| Arrival rate (pps)                      | Current traffic rate for packets arriving from all cards and at the Routing Engine |
| Peak rate (pps)                         | Highest traffic rate for packets arriving from all cards and at the Routing Engine |
| Policer bandwidth violation detected at | Timestamp of the policer violation                                                 |
| Detected on                             | Slot number of the card on which the violation was detected                        |

## Sample Output

### show ddos-protection protocols violations

```

user@host> show ddos-protection protocols violations
Number of packet types that are being violated: 2
Protocol Packet Bandwidth Arrival Peak Policer bandwidth
group type (pps) rate(pps) rate(pps) violation detected at
pppoe padi 500 2000 2001 2011-04-19 08:23:17 PDT
 Detected on: FPC-1
pppoe padr 500 1999 2001 2011-04-19 08:23:17 PDT
 Detected on: FPC-1

```

### show ddos-protection protocols lldp violations

```

user@host> show ddos-protection protocols lldp violations
Number of packet types that are being violated: 0

```

### show ddos-protection protocols pppoe violations

```

user@host> show ddos-protection protocols pppoe violations
Number of packet types that are being violated: 2
Protocol Packet Bandwidth Arrival Peak Policer bandwidth
group type (pps) rate(pps) rate(pps) violation detected at
pppoe padi 500 2000 2001 2011-04-19 08:23:17 PDT
 Detected on: FPC-1
pppoe padr 500 1999 2001 2011-04-19 08:23:17 PDT
 Detected on: FPC-1

```

# System Services Feature Guide for QFX10000 Switches





## PART 84

# Port Mirroring

- [Configuring Port Mirroring on page 6227](#)



# Configuring Port Mirroring

- [Understanding Port Mirroring on page 6227](#)
- [Configuring Port Mirroring on page 6232](#)
- [Examples: Configuring Port Mirroring for Local Analysis on page 6235](#)
- [Example: Configuring Port Mirroring for Remote Analysis on page 6240](#)
- [Example: Mirroring Employee Web Traffic with a Firewall Filter on page 6245](#)
- [Troubleshooting Port Mirroring on page 6248](#)

## Understanding Port Mirroring

---

- [Port Mirroring Overview on page 6227](#)
- [Port Mirroring Instance Types on page 6228](#)
- [Port-Mirroring Terminology on page 6228](#)
- [Port Mirroring and STP on page 6230](#)
- [Port Mirroring Constraints and Limitations on page 6230](#)

### Port Mirroring Overview

Port mirroring copies packets entering or exiting a port or entering a VLAN and sends the copies to a local interface for local monitoring or to a VLAN for remote monitoring. Use port mirroring to send traffic to applications that analyze traffic for purposes such as monitoring compliance, enforcing policies, detecting intrusions, monitoring and predicting traffic patterns, correlating events, and so on.

Port mirroring is needed for traffic analysis on a switch because a switch normally sends packets only to the port to which the destination device is connected. You configure port mirroring on the switch to send copies of unicast traffic to a local interface or a VLAN and run an analyzer application on a device connected to the interface or VLAN. You configure port mirroring by using the **analyzer** statement.

Keep performance in mind when configuring port mirroring. For example, If you mirror traffic from multiple ports, the mirrored traffic may exceed the capacity of the output interface. We recommend that you limit the amount of copied traffic by selecting specific interfaces instead of using the **all** keyword. You can also limit the amount of mirrored traffic by using a firewall filter to send specific traffic to a port mirroring instance. Mirroring only the necessary packets reduces the possibility of a performance impact.

You can use port mirroring to copy any of the following:

- All packets entering or exiting an interface (in any combination)—For example, you can send copies of the packets entering some interfaces and the packets exiting other interfaces to the same local interface or VLAN. If you configure port mirroring to copy packets exiting an interface, traffic that originates on that switch or Node device (in a QFabric system) is not copied when it egresses. Only switched traffic is copied on egress. (See the limitation on egress mirroring below.)
- All packets entering a VLAN—You cannot use port mirroring to copy packets exiting a VLAN.
- Firewall-filtered sample—Sample of packets entering a port or VLAN. Configure a firewall filter to select certain packets for mirroring.



**NOTE:** Firewall filters are not supported on egress ports; therefore, you cannot specify policy-based sampling of packets exiting an interface.

## Port Mirroring Instance Types

To configure port mirroring, you configure an instance of one of the following types:

- Analyzer instance: You must specify the input and output for the instance. This instance type is useful for ensuring that all traffic transiting an interface or VLAN is mirrored and sent to the analyzer device.
- Port-mirroring instance: You do not specify an input for this instance type. Instead, you, create a firewall filter that specifies the required traffic and directs it to the mirror. This instance type is useful for controlling which types of traffic should be mirrored. When you use a port-mirroring instance, you can direct traffic to it in the following ways:
  - Specify the name of the port-mirroring instance in the firewall filter using the **port-mirror-instance *instance-name*** action. You should use this approach if there are multiple port-mirroring instances defined.
  - Configure the filter to send the mirrored packets to the output interface defined in the instance using the **port-mirror** action. You can use this approach if there is only one port-mirroring instance defined.

## Port-Mirroring Terminology

Table 499 lists the terms used in the documentation about port mirroring and provides definitions.

**Table 499: Port Mirroring Terms and Definitions**

| Term              | Description                                                                                                                                                              |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Analyzer instance | Port-mirroring configuration that includes a name, source interfaces or source VLAN, and a destination for mirrored packets (either a local access interface or a VLAN). |

Table 499: Port Mirroring Terms and Definitions (*continued*)

|                                                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-----------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Port mirroring instance                                         | A port-mirroring configuration that does not specify an input.. A firewall filter must be used to send traffic to the port mirror. Use the action <b>port-mirror-instance <i>instance-name</i></b> in the firewall filter configuration to send packets to the port mirror.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Output interface (also known as monitor interface)              | <p>Access interface to which packet copies are sent and to which a device running an analyzer application is connected.</p> <p>The following limitations apply to an output interface:</p> <ul style="list-style-type: none"> <li>• Cannot also be a source port.</li> <li>• Cannot be used for switching.</li> <li>• Cannot be an aggregated Ethernet interface (LAG).</li> <li>• Does not participate in Layer 2 protocols, such as Spanning Tree Protocol (STP).</li> <li>• Loses any existing VLAN associations when you configure it as an analyzer output interface.</li> </ul> <p>If the capacity of the output interface is insufficient to handle the traffic from the source ports, overflow packets are dropped.</p>                                                                                                                                                                                                                                               |
| Output IP address                                               | <p>IP address of the device running an analyzer application. The device can be on a remote network. When you use this feature, the mirrored packets are GRE-encapsulated. The analyzer device must be able to de-encapsulate GRE-encapsulated packets, or the GRE-encapsulated packets must be de-encapsulated before reaching the analyzer device. (You can use a network sniffer to de-encapsulate the packets.)</p> <ul style="list-style-type: none"> <li>• An output IP address cannot be in the same subnetwork as any of the switch's management interfaces.</li> <li>• If you create virtual routing instances and also create an analyzer configuration that includes an output IP address, the output address belongs to the default virtual routing instance (inet.0 routing table).</li> </ul>                                                                                                                                                                    |
| Output VLAN (also known as monitor or analyzer VLAN)            | <p>VLAN to which copies are sent and to which a device running an analyzer application is connected. The analyzer VLAN can span multiple switches.</p> <p>The following limitations apply to an output VLAN:</p> <ul style="list-style-type: none"> <li>• Cannot be a private VLAN or VLAN range.</li> <li>• Cannot be shared by multiple <b>analyzer</b> statements.</li> <li>• An output VLAN interface cannot be a member of any other VLAN.</li> <li>• An output VLAN interface cannot be an aggregated Ethernet interface (LAG).</li> <li>• On some switches, only one interface can be a member of the analyzer VLAN. This limitation does not apply on the QFX10000 switch if traffic is mirrored on ingress. In this case, multiple QFX10000 interfaces can belong to the output VLAN, and traffic is mirrored to all of those interfaces. If traffic is mirrored on egress on a QFX10000 switch, only one interface can be a member of the analyzer VLAN.</li> </ul> |
| Input interface (also known as mirrored or monitored interface) | Interface that provides traffic to be mirrored. This traffic can be entering or exiting the interface. (Ingress or egress traffic can be mirrored.) An input interface cannot also be an output interface for an analyzer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Monitoring station                                              | Computer running an analyzer application.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Local port mirroring                                            | Port-mirroring configuration in which the mirrored packets are sent to an interface on the same switch.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

Table 499: Port Mirroring Terms and Definitions (*continued*)

|                        |                                                                                                                                                                                                                                         |
|------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Remote port mirroring  | Flooding mirrored packets to an output (analyzer) VLAN that you create to receive mirror traffic or sending the mirrored packets to a remote IP address. (You cannot send mirrored packets to a remote IP address on a QFabric system.) |
| Policy-based mirroring | Mirroring of packets that match the match a firewall filter term. The action <b>analyzer</b> <i>analyzer-name</i> is used in the firewall filter to send the packets to the analyzer.                                                   |

## Port Mirroring and STP

The behavior of STP in a port-mirroring configuration depends on the version of Junos OS you are using:

- Junos OS 13.2X50, Junos OS 13.2X51-D25 or earlier, Junos OS 13.2X52: If you enable STP, port mirroring might not work because STP might block the mirrored packets.
- Junos OS 13.2X51-D30, Junos OS 14.1X53: STP is disabled for mirrored traffic. You must ensure that your topology prevents loops for this traffic.

## Port Mirroring Constraints and Limitations

- [Local and Remote Port Mirroring on page 6230](#)
- [Remote Port Mirroring Only on page 6232](#)

### Local and Remote Port Mirroring

The following constraints and limitations apply to local and remote port mirroring:

- You can create a total of four port-mirroring configurations.
- You can create a total of four port-mirroring configurations on each Node group in a QFabric system, subject to the following constraints:
  - As many as four of the configurations can be for local port mirroring.
  - As many as three of the configurations can be for remote port mirroring.
- Regardless of whether you are configuring a standalone switch or a Node group, the following limits apply:
  - There can be no more than two configurations that mirror ingress traffic. (If you configure a firewall filter to send traffic to a port mirror—that is, you use the **analyzer** action modifier in a filter term—this counts as an ingress mirroring configuration for switch or Node group on which the filter is applied.)
  - There can be no more than two configurations that mirror egress traffic.



**NOTE:** On QFabric systems, there is no system-wide limit on the total number of mirror sessions.

- You can configure no more than one type of output in one port-mirroring configuration. That is, you can use no more than one of the following to complete a **set analyzer name output** statement:
  - **interface**
  - **ip-address**
  - **vlan**
- If you configure Junos OS to mirror egress packets, do not configure more than 2000 VLANs on a standalone switch or QFabric system. If you do so, some VLAN packets might contain incorrect VLAN IDs. This applies to any VLAN packets—not only the mirrored copies.
- The **ratio** and **loss-priority** options are not supported.
- Packets with physical layer errors are filtered out and are not sent to the output port or VLAN.
- If you use sFlow monitoring to sample traffic, it does not sample the mirror copies when they exit from the output interface.
- You cannot mirror packets exiting or entering the following ports:
  - Dedicated Virtual Chassis interfaces
  - Management interfaces (me0 or vme0)
  - Fibre Channel interfaces
  - Routed VLAN interfaces
- An aggregated Ethernet interface cannot be an output interface if the input is a VLAN or if traffic is sent to the analyzer by a firewall filter.
- Do not include an 802.1Q subinterface that has a unit number other than 0 in a port mirroring configuration. Port mirroring does not work with subinterfaces if their unit number is not 0. (You configure 802.1Q subinterfaces using the **vlan-tagging** statement.)
- When packet copies are sent out the output interface, they are not modified for any changes that are normally applied on egress, such as CoS rewriting.
- An interface can be the input interface for only one mirroring configuration. Do not use the same interface as the input interface for multiple mirroring configurations.
- CPU-generated packets (such as ARP, ICMP, BPDU, and LACP packets) cannot be mirrored on egress.
- VLAN-based mirroring is not supported for STP traffic.
- (QFabric systems only) If you configure a QFabric analyzer to mirror egress traffic and the input and output interfaces are on different Node devices, the mirrored copies have incorrect VLAN IDs. This limitation does not apply if you configure a QFabric analyzer to mirror egress traffic and the input and output interfaces are on the *same* Node device. In this case the mirrored copies have the correct VLAN IDs (as long as you do not configure more than 2000 VLANs on the QFabric system).

### Remote Port Mirroring Only

---

The following constraints and limitations apply to remote port mirroring:

- If you configure an output IP address, the address cannot be in the same subnetwork as any of the switch's management interfaces.
- If you create virtual routing instances and also create an analyzer configuration that includes an output IP address, the output address belongs to the default virtual routing instance (inet.0 routing table).
- An output VLAN cannot be a private VLAN or VLAN range.
- An output VLAN cannot be shared by multiple **analyzer** statements.
- An output VLAN interface cannot be a member of any other VLAN.
- An output VLAN interface cannot be an aggregated Ethernet interface.

#### Related Documentation

- [Configuring Port Mirroring on page 6232](#)
- [Examples: Configuring Port Mirroring for Local Analysis on page 6235](#)
- [Examples: Configuring Port Mirroring for Local Analysis on page 6235](#)
- [Example: Configuring Port Mirroring for Remote Analysis on page 6240](#)
- [Troubleshooting Port Mirroring on page 6248](#)

## Configuring Port Mirroring

---

You use port mirroring to copy packets and send the copies to a device running an application such as a network analyzer or intrusion detection application so that you can analyze traffic without delaying it. You can mirror traffic entering or exiting a port or entering a VLAN, and you can send the copies to a local access interface or to a VLAN through a trunk interface.

We recommend that you disable port mirroring when you are not using it. To avoid creating a performance issue If you do enable port mirroring, we recommend that you select specific input interfaces instead of using the **all** keyword. You can also limit the amount of mirrored traffic by using a firewall filter.



.....

**NOTE:** This task uses the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring Port Mirroring*. For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on page 41.

.....



.....

**NOTE:** If you want to create additional analyzers without deleting an existing analyzer, first disable the existing analyzer using the **disable analyzer analyzer-name** command.

.....





**NOTE:** You must configure port mirroring output interfaces as family **ethernet-switching**.

- [Configuring Port Mirroring for Local Analysis on page 6233](#)
- [Configuring Port Mirroring for Remote Analysis on page 6233](#)
- [Filtering the Traffic Entering an Analyzer on page 6234](#)

## Configuring Port Mirroring for Local Analysis

To mirror interface traffic to a local interface on the switch:

1. If you want to mirror traffic that is ingressing or egressing specific interfaces, choose a name for the port-mirroring configuration and configure what traffic should be mirrored by specifying the interfaces and direction of traffic:

```
[edit forwarding-options]
user@switch# set analyzer analyzer-name input (ingress | egress) interface interface-name
```



**NOTE:** If you configure Junos OS to mirror egress packets, do not configure more than 2000 VLANs. If you do so, some VLAN packets might contain incorrect VLAN IDs.



**NOTE:** If you configure mirroring for packets that egress an access interface, the original packets lose any VLAN tags when they exit the access interface, but the mirrored (copied) packets retain the VLAN tags when they are sent to the analyzer system.

2. If you want to specify that all traffic entering a VLAN should be mirrored, choose a name for the port-mirroring configuration and specify the VLAN:

```
[edit forwarding-options]
user@switch# set analyzer analyzer-name input ingress vlan vlan-name
```



**NOTE:** You cannot configure port mirroring to copy traffic that egresses a VLAN.

3. Configure the destination interface for the mirrored packets:

```
[edit forwarding-options]
user@switch# set analyzer analyzer-name output interface interface-name
```

## Configuring Port Mirroring for Remote Analysis

To mirror traffic to a VLAN for analysis at a remote location:

1. Configure a VLAN to carry the mirrored traffic:

```
[edit]
```

```
user@switch# set vlans vlan-name vlan-id number
```

2. Configure the interface that connects to another switch (the uplink interface) to trunk mode and associate it with the appropriate VLAN:

```
[edit]
```

```
user@switch# set interfaces interface-name unit 0 family ethernet-switching port-mode
trunk vlan members (vlan-name | vlan-id)
```

3. Configure the analyzer:

- a. Choose a name for the analyzer:

```
[edit forwarding-options]
```

```
user@switch# set analyzer analyzer-name
```

- b. Specify the interface to be mirrored and whether the traffic should be mirrored on ingress or egress:

```
[edit forwarding-options]
```

```
user@switch# set analyzer analyzer-name input (ingress | egress) interface interface-name
```

- c. Specify the appropriate IP address or VLAN as the output (a VLAN is specified in this example:

```
[edit forwarding-options]
```

```
user@switch# set analyzer analyzer-name output vlan (vlan-name | vlan-id)
```

If you specify an IP address as the output, note the following constraints:

- The address cannot be in the same subnet as any of the switch's management interfaces.
- If you create virtual routing instances and also create an analyzer configuration that includes an output IP address, the output address belongs to the default virtual routing instance (*inet.0* routing table).
- The analyzer device must be able to de-encapsulate GRE-encapsulated packets, or the GRE-encapsulated packets must be de-encapsulated before reaching the analyzer device. (You can use a network sniffer to de-encapsulate the packets.)

## Filtering the Traffic Entering an Analyzer

In addition to specifying which traffic to mirror by configuring an analyzer, you can also use a firewall filter to exercise more control over which packets are copied. For example, you might use a filter to specify that only traffic from certain applications be mirrored. The filter can use any of the available match conditions and must have an action of modifier of **port-mirror-instance** *instance-name*. If you use the same analyzer in multiple filters or terms, the output packets are copied only once.

When you use a firewall filter as the input to a port-mirroring instance, you send the copied traffic to a local interface or a VLAN just as you do when a firewall is not involved.

To configure port mirroring with filters:

1. Configure a port-mirroring instance for local or remote analysis. Configure only the output. For example, for local analysis enter:

```
[edit forwarding-options]
```

```
user@switch# set port-mirroring-instance instance-name output interface interface-name
```



**NOTE:** You cannot configure input to this instance.

2. Create a firewall filter using any of the available match conditions. In a **then** term, specify include the action modifier **port-mirror-instance** *instance-name*.
3. Apply the firewall filter to the interfaces or VLAN that should provide the input to the analyzer:

```
[edit]
user@switch# set interfaces interface-name unit 0 family ethernet-switching filter input
filter-name
[edit]
user@switch# set vlan (vlan-name | vlan-id) filter input filter-name
```

#### Related Documentation

- [Understanding Port Mirroring on page 6227](#)
- [Examples: Configuring Port Mirroring for Local Analysis on page 6235](#)
- [Example: Configuring Port Mirroring for Remote Analysis on page 6240](#)
- [Overview of Firewall Filters on page 5951](#)

## Examples: Configuring Port Mirroring for Local Analysis

Use port mirroring to send traffic to applications that analyze traffic for purposes such as monitoring compliance, enforcing policies, detecting intrusions, monitoring and predicting traffic patterns, correlating events, and so on. Port mirroring copies packets entering or exiting an interface or entering a VLAN and sends the copies to a local interface for local monitoring.



**NOTE:** This example uses the Enhanced Layer 2 Software (ELS) configuration style. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 41](#).

This example describes how to configure port mirroring to copy traffic sent by employee computers to a switch to an access interface on the same switch.

- [Requirements on page 6235](#)
- [Overview and Topology on page 6236](#)
- [Example: Mirroring All Employee Traffic for Local Analysis on page 6236](#)
- [Example: Mirroring Employee Web Traffic with a Firewall Filter on page 6237](#)

## Requirements

This example uses the following hardware and software components:

- Junos OS Release 13.2
- A switch

## Overview and Topology

This topic includes two related examples that describe how to mirror traffic entering interfaces on the switch to an access interface on the same switch. The first example shows how to mirror all traffic sent by employee computers to the switch. The second example includes a filter to mirror only the employee traffic going to the Web.

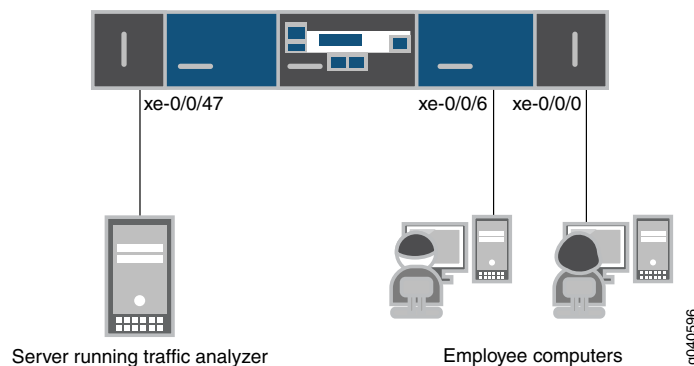
In this example, **xe-0/0/0** and **xe-0/0/6** serve as connections for employee computers. Interface **xe-0/0/47** is connected to a device running an analyzer application.



**NOTE:** Multiple ports mirrored to one interface can cause buffer overflow and dropped packets.

Figure 202 shows the network topology for this example.

Figure 202: Network Topology for Local Port Mirroring Example



### Example: Mirroring All Employee Traffic for Local Analysis

To configure port mirroring for all traffic sent by employee computers for local analysis, perform the tasks explained in this section.

#### CLI Quick Configuration

To quickly configure local port mirroring for ingress traffic to the two ports connected to employee computers, copy the following commands and paste them into a switch terminal window:

```
[edit]
set interfaces xe-0/0/0 unit 0 family ethernet-switching
set interfaces xe-0/0/6 unit 0 family ethernet-switching
set interfaces xe-0/0/47 unit 0 family ethernet-switching
set forwarding-options analyzer employee-monitor input ingress interface xe-0/0/0.0
set forwarding-options analyzer employee-monitor input ingress interface xe-0/0/6.0
set forwarding-options analyzer employee-monitor output interface xe-0/0/47.0
```

#### Step-by-Step Procedure

To configure an analyzer called **employee-monitor** and specify the input (source) interfaces and the output interface:

1. Configure the interfaces connected to employee computers as input interfaces for the port-mirror analyzer **employee-monitor**:

```
[edit forwarding-options]
user@switch# set analyzer employee-monitor input ingress interface xe-0/0/0.0
user@switch# set analyzer employee-monitor input ingress interface xe-0/0/6.0
```

2. Configure the output analyzer interface for the **employee-monitor** analyzer. This will be the destination interface for the mirrored packets:

```
[edit forwarding-options]
user@switch# set analyzer employee-monitor output interface xe-0/0/47.0
```

**Results** Check the results of the configuration:

```
[edit]
user@switch# show forwarding-options analyzer
employee-monitor {
 input {
 ingress {
 interface xe-0/0/0.0;
 interface xe-0/0/6.0;
 }
 }
 output {
 interface {
 xe-0/0/47.0;
 }
 }
}
```

## Example: Mirroring Employee Web Traffic with a Firewall Filter

- [Requirements on page 6237](#)
- [Overview on page 6237](#)
- [Configuring on page 6238](#)
- [Verification on page 6240](#)

### Requirements

This example uses the following hardware and software components:

- One switch
- Junos 13.2X51

### Overview

Rather than mirror all traffic, it is usually desirable to mirror only certain traffic. This is a more-efficient use of your bandwidth and hardware and might be necessary because constraints on these assets. To select specific traffic for mirroring, you use a firewall filter to match the desired traffic and direct it to a port-mirroring instance. The port-mirroring instance then copies the packets and sends them to the output VLAN, interface, or IP address.

## Configuring

To specify that the only traffic that will be mirrored is traffic sent by employees to the Web, perform the tasks explained in this section. To select this traffic for mirroring, you use a firewall filter to specify this traffic and direct it to a port-mirroring instance.

### CLI Quick Configuration

To quickly configure local port mirroring of traffic from employee computers that is destined for the Web, copy the following commands and paste them into a switch terminal window:

```
[edit]
set forwarding-options port-mirroring instance employee-web-monitor output interface
xe-0/0/47.0
set firewall family ethernet-switching filter watch-employee term employee-to-corp from
destination-address 192.0.2.16/28
set firewall family ethernet-switching filter watch-employee term employee-to-corp from
source-address 192.0.2.16/28
set firewall family ethernet-switching filter watch-employee term employee-to-corp then accept
set firewall family ethernet-switching filter watch-employee term employee-to-web from
destination-port 80
set firewall family ethernet-switching filter watch-employee term employee-to-web then
port-mirror-instance employee-web-monitor
set interfaces xe-0/0/0 unit 0 family ethernet-switching filter input watch-employee
set interfaces xe-0/0/6 unit 0 family ethernet-switching filter input watch-employee
```

### Step-by-Step Procedure

To configure local port mirroring of employee-to-web traffic from the two ports connected to employee computers:

1. Configure the output interface:

```
[edit interfaces]
user@switch# set xe-0/0/47 unit 0 family ethernet-switching
```

2. Configure the **employee-web-monitor** output interface. (Configure only the output—the input comes from the filter.)

```
[edit forwarding-options]
user@switch# set port-mirroring instance employee-web-monitor output interface
xe-0/0/47.0
```

3. Configure a firewall filter called **watch-employee** that includes a term to match traffic sent to the Web and send it to the port-mirroring instance **employee-web-monitor**. Traffic to and from the corporate subnet (destination or source address of **192.0.2.16/28**) does not need to be copied, so create another term to accept that traffic before it reaches the term that sends Web traffic to the instance:

```
[edit firewall family ethernet-switching]
user@switch# set filter watch-employee term employee-to-corp from destination-address
192.0.2.16/28
user@switch# set filter watch-employee term employee-to-corp from source-address
192.0.2.16/28
user@switch# set filter watch-employee term employee-to-corp then accept
user@switch# set filter watch-employee term employee-to-web from destination-port 80
user@switch# set filter watch-employee term employee-to-web then port-mirror-instance
employee-web-monitor
```

4. Apply the firewall filter to the appropriate interfaces as an ingress filter (egress filters do not allow analyzers):

```
[edit interfaces]
```

```

user@switch# set xe-0/0/0 unit 0 family ethernet-switching filter input watch-employee
user@switch# set xe-0/0/6 unit 0 family ethernet-switching filter input watch-employee

```

**Results** Check the results of the configuration:

```

[edit]
user@switch# show
forwarding-options {
 port-mirroring {
 instance {
 employee-web-monitor {
 output {
 interface xe-0/0/47.0;
 }
 }
 }
 }
}
...
firewall family ethernet-switching {
 filter watch-employee {
 term employee-to-web {
 from {
 destination-port 80;
 }
 then port-mirror-instance employee-web-monitor;
 }
 }
}
...
interfaces {
 xe-0/0/0 {
 unit 0 {
 family ethernet-switching {
 filter {
 input watch-employee;
 }
 }
 }
 }
 xe-0/0/6 {
 family ethernet-switching {
 filter {
 input watch-employee;
 }
 }
 }
}
}

```

## Verification

---

### *Verifying That the Analyzer Has Been Correctly Created*

**Purpose** Verify that the analyzer named **employee-monitor** or **employee-web-monitor** has been created on the switch with the appropriate input interfaces and appropriate output interface.

**Action** You can verify that the port mirror analyzer has been configured as expected using the **show analyzer** command.

```
user@switch> show forwarding-options analyzer
 Port mirror name : employee-monitor
 Mirror rate : 1
 Maximum packet length : 0
 State : up
 Ingress monitored interfaces : xe-0/0/0.0
 Ingress monitored interfaces : xe-0/0/6.0
 Output interface : xe-0/0/47.0
```

**Meaning** This output shows that the port-mirroring instance **employee-monitor** has a ratio of 1 (mirroring every packet, the default setting), the maximum size of the original packet that was mirrored (0 indicates the entire packet), the state of the configuration (is up indicates that the instance is mirroring the traffic entering the xe-0/0/0, and xe-0/0/6 interfaces, and sending the mirrored traffic to the xe-0/0/47 interface). If the state of the output interface is down or if the output interface is not configured, the value of state will be **down** and the instance will not be programmed for mirroring.

**Related Documentation** [• Understanding Port Mirroring on page 6227](#)

**Related Documentation** [• Understanding Port Mirroring on page 6227](#)  
[• Configuring Port Mirroring on page 6232](#)  
[• Example: Configuring Port Mirroring for Remote Analysis on page 6240](#)

## Example: Configuring Port Mirroring for Remote Analysis

---

Use port mirroring to send traffic to applications that analyze traffic for purposes such as monitoring compliance, enforcing policies, detecting intrusions, monitoring and predicting traffic patterns, correlating events, and so on. Port mirroring copies packets entering or exiting an interface or entering a VLAN and sends the copies either to a local interface for local monitoring or to a VLAN for remote monitoring. This example describes how to configure port mirroring for remote analysis.

- [Requirements on page 6241](#)
- [Overview and Topology on page 6241](#)
- [Mirroring All Employee Traffic for Remote Analysis on page 6241](#)



- [Mirroring Employee-to-Web Traffic for Remote Analysis on page 6242](#)
- [Verification on page 6244](#)

## Requirements

This example uses the following hardware and software components:

- Junos OS Release 13.2 for the QFX Series
- A switch

## Overview and Topology

This topic includes two related examples that describe how to mirror traffic entering ports on the switch to an analyzer VLAN so that you can perform analysis using a remote device. The first example shows how to mirror all traffic sent by employee computers to the switch. The second example includes a filter to mirror only the employee traffic going to the Web.

In this example:

- Interfaces **ge-0/0/0** and **ge-0/0/1** are Layer 2 interfaces that connect to employee computers.
- Interface **ge-0/0/10** is a Layer 2 interface that connects to another switch.
- VLAN **remote-analyzer** is configured on all switches in the topology to carry the mirrored traffic.



**NOTE:** In addition to performing the configuration steps described here, you must also configure the analyzer VLAN (**remote-analyzer** in this example) on the other switches that are used to connect the source switch (the one in this configuration) to the one that the monitoring station is connected to.

## Mirroring All Employee Traffic for Remote Analysis

### CLI Quick Configuration

To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **edit** hierarchy level:

```
[edit]
set vlans remote-analyzer vlan-id 999
set interfaces ge-0/0/10 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/10 unit 0 family ethernet-switching vlan members 999
set forwarding-options analyzer employee-monitor input ingress interface ge-0/0/0.0
set forwarding-options analyzer employee-monitor input ingress interface ge-0/0/1.0
set forwarding-options analyzer employee-monitor output vlan remote-analyzer
```

**Step-by-Step Procedure**

To configure basic remote port mirroring:

1. Configure the analyzer VLAN (called **remote-analyzer** in this example):
 

```
[edit vlans]
user@switch# set vlans remote-analyzer vlan-id 999
```
2. Configure the interface connected to another switch for trunk mode and associate it with the **remote-analyzer** VLAN:
 

```
[edit interfaces]
user@switch# set ge-0/0/10 unit 0 family ethernet-switching port-mode trunk
user@switch# set ge-0/0/10 unit 0 family ethernet-switching vlan members 999
```
3. Configure the **employee-monitor** analyzer:
 

```
[edit forwarding-options]
user@switch# set analyzer employee-monitor
user@switch# set analyzer employee-monitor input ingress interface ge-0/0/0.0
user@switch# set analyzer employee-monitor input ingress interface ge-0/0/1.0
user@switch# set analyzer employee-monitor output vlan remote-analyzer
```
4. Configure the **remote-analyzer** VLAN on the switches that connect this switch to the monitoring workstation.

**Results** Check the results of the configuration:

```
[edit]
user@switch# show
forwarding-options {
 analyzer employee-monitor {
 input {
 ingress {
 interface ge-0/0/0.0;
 interface ge-0/0/1.0;
 }
 }
 output {
 vlan {
 remote-analyzer;
 }
 }
 }
}
```

## Mirroring Employee-to-Web Traffic for Remote Analysis

**CLI Quick Configuration**

To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **edit** hierarchy level:

```
[edit]
set vlans remote-analyzer vlan-id 999
set interfaces ge-0/0/10 unit 0 family ethernet-switching interface-mode trunk
set interfaces ge-0/0/10 unit 0 family ethernet-switching vlan members 999
set forwarding-options port-mirroring instance employee-web-monitor loss-priority high output vlan 999
set firewall family ethernet-switching filter watch-employee term employee-to-web from destination-port 80
```

```

set firewall family ethernet-switching filter watch-employee term employee-to-web then
port-mirror-instance employee-web-monitor employee-web-monitor
set ge-0/0/0 unit 0 family ethernet-switching filter input watch-employee
set interfaces ge-0/0/1 unit 0 family ethernet-switching filter input watch-employee

```

#### Step-by-Step Procedure

1. Configure the analyzer VLAN (called **remote-analyzer** in this example):  

```

[edit vlans]
user@switch# set remote-analyzer vlan-id 999

```
2. Configure an interface to associate it with the **remote-analyzer** VLAN:  

```

[edit interfaces]
user@switch# set interfaces ge-0/0/10 unit 0 family ethernet-switching interface-mode trunk
user@switch# set ge-0/0/10 unit 0 family ethernet-switching vlan members 999

```
3. Configure the **employee-web-monitor** analyzer. (Configure only the output—the input comes from the filter.)  

```

[edit forwarding-options]
user@switch# set forwarding-options port-mirroring instance employee-web-monitor output vlan 999

```
4. Configure a firewall filter called **watch-employee** to match traffic sent to the Web and send it to the analyzer **employee-web-monitor**:  

```

[edit firewall family ethernet-switching]
user@switch# set filter watch-employee term employee-to-web from destination-port 80
user@switch# set filter watch-employee term employee-to-web then port-mirror-instance employee-web-monitor

```
5. Apply the firewall filter to the appropriate interfaces as an ingress filter:  

```

[edit interfaces]
user@switch# set ge-0/0/0 unit 0 family ethernet-switching filter input watch-employee
user@switch# set ge-0/0/1 unit 0 family ethernet-switching filter input watch-employee

```
6. Configure the **remote-analyzer** VLAN on the switches that connect this switch to the monitoring workstation.

**Results** Check the results of the configuration:

```

[edit]
user@switch# show
interfaces {
 ...
 ge-0/0/10 {
 unit 0 {
 family ethernet-switching {
 interface-mode trunk;
 vlan {
 members remote-analyzer;
 }
 }
 }
 }
 ge-0/0/0 {
 unit 0 {
 family ethernet-switching {
 filter {
 input watch-employee;
 }
 }
 }
 }
}

```

```
 }
 }
}
ge-0/0/1 {
 unit 0 {
 family ethernet-switching {
 filter {
 input watch-employee;
 }
 }
 }
}
...
firewall {
 family ethernet-switching {
 ...
 filter watch-employee {
 term employee-to-web {
 from {
 destination-port 80;
 }
 then port-mirror-instance employee-web-monitor;
 }
 }
 }
}
forwarding-options analyzer {
 employee-web-monitor {
 output {
 vlan {
 999;
 }
 }
 }
}
vlands {
 remote-analyzer {
 vlan-id 999;
 }
}
```

## Verification

---

### Verifying That the Analyzer Has Been Correctly Created

**Purpose** Verify that the analyzer named **employee-monitor** or **employee-web-monitor** has been created on the switch with the appropriate input interfaces and appropriate output interface.

**Action** You can verify the port mirror analyzer is configured as expected using the **show analyzer** command.

```
user@switch> show analyzer
Analyzer name : employee-monitor
Output VLAN : remote-analyzer
Ingress monitored interfaces : ge-0/0/0.0
Ingress monitored interfaces : ge-0/0/1.0
```

**Meaning** This output shows that the **employee-monitor** analyzer is mirroring the traffic entering **ge-0/0/0** and **ge-0/0/1** and is sending the mirror traffic to the analyzer **remote-analyzer**.

- Related Documentation**
- [Understanding Port Mirroring on page 6227](#)
  - [Configuring Port Mirroring on page 6232](#)
  - [Examples: Configuring Port Mirroring for Local Analysis on page 6235](#)
  - [Overview of Firewall Filters on page 5951](#)

## Example: Mirroring Employee Web Traffic with a Firewall Filter

- [Requirements on page 6245](#)
- [Overview on page 6245](#)
- [Configuring on page 6245](#)
- [Verification on page 6247](#)

### Requirements

This example uses the following hardware and software components:

- One switch
- Junos 13.2X51

### Overview

Rather than mirror all traffic, it is usually desirable to mirror only certain traffic. This is a more-efficient use of your bandwidth and hardware and might be necessary because constraints on these assets. To select specific traffic for mirroring, you use a firewall filter to match the desired traffic and direct it to a port-mirroring instance. The port-mirroring instance then copies the packets and sends them to the output VLAN, interface, or IP address.

### Configuring

To specify that the only traffic that will be mirrored is traffic sent by employees to the Web, perform the tasks explained in this section. To select this traffic for mirroring, you use a firewall filter to specify this traffic and direct it to a port-mirroring instance.

|                                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>CLI Quick Configuration</b> | <p>To quickly configure local port mirroring of traffic from employee computers that is destined for the Web, copy the following commands and paste them into a switch terminal window:</p> <pre>[edit] set forwarding-options port-mirroring instance employee-web-monitor output interface xe-0/0/47.0 set firewall family ethernet-switching filter watch-employee term employee-to-corp from destination-address 192.0.2.16/28 set firewall family ethernet-switching filter watch-employee term employee-to-corp from source-address 192.0.2.16/28 set firewall family ethernet-switching filter watch-employee term employee-to-corp then accept set firewall family ethernet-switching filter watch-employee term employee-to-web from destination-port 80 set firewall family ethernet-switching filter watch-employee term employee-to-web then port-mirror-instance employee-web-monitor set interfaces xe-0/0/0 unit 0 family ethernet-switching filter input watch-employee set interfaces xe-0/0/6 unit 0 family ethernet-switching filter input watch-employee</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Step-by-Step Procedure</b>  | <p>To configure local port mirroring of employee-to-web traffic from the two ports connected to employee computers:</p> <ol style="list-style-type: none"> <li>1. Configure the output interface: <pre>[edit interfaces] user@switch# set xe-0/0/47 unit 0 family ethernet-switching</pre> </li> <li>2. Configure the <b>employee-web-monitor</b> output interface. (Configure only the output—the input comes from the filter.) <pre>[edit forwarding-options] user@switch# set port-mirroring instance employee-web-monitor output interface xe-0/0/47.0</pre> </li> <li>3. Configure a firewall filter called <b>watch-employee</b> that includes a term to match traffic sent to the Web and send it to the port-mirroring instance <b>employee-web-monitor</b>. Traffic to and from the corporate subnet (destination or source address of <b>192.0.2.16/28</b>) does not need to be copied, so create another term to accept that traffic before it reaches the term that sends Web traffic to the instance: <pre>[edit firewall family ethernet-switching] user@switch# set filter watch-employee term employee-to-corp from destination-address 192.0.2.16/28 user@switch# set filter watch-employee term employee-to-corp from source-address 192.0.2.16/28 user@switch# set filter watch-employee term employee-to-corp then accept user@switch# set filter watch-employee term employee-to-web from destination-port 80 user@switch# set filter watch-employee term employee-to-web then port-mirror-instance employee-web-monitor</pre> </li> <li>4. Apply the firewall filter to the appropriate interfaces as an ingress filter (egress filters do not allow analyzers): <pre>[edit interfaces] user@switch# set xe-0/0/0 unit 0 family ethernet-switching filter input watch-employee user@switch# set xe-0/0/6 unit 0 family ethernet-switching filter input watch-employee</pre> </li> </ol> |
| <b>Results</b>                 | <p>Check the results of the configuration:</p> <pre>[edit] user@switch# show</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

```

forwarding-options {
 port-mirroring {
 instance {
 employee-web-monitor {
 output {
 interface xe-0/0/47.0;
 }
 }
 }
 }
}
...
firewall family ethernet-switching {
 filter watch-employee {
 term employee-to-web {
 from {
 destination-port 80;
 }
 then port-mirror-instance employee-web-monitor;
 }
 }
}
...
interfaces {
 xe-0/0/0 {
 unit 0 {
 family ethernet-switching {
 filter {
 input watch-employee;
 }
 }
 }
 }
 xe-0/0/6 {
 family ethernet-switching {
 filter {
 input watch-employee;
 }
 }
 }
}

```

## Verification

### Verifying That the Analyzer Has Been Correctly Created

**Purpose** Verify that the analyzer named **employee-monitor** or **employee-web-monitor** has been created on the switch with the appropriate input interfaces and appropriate output interface.

**Action** You can verify that the port mirror analyzer has been configured as expected using the **show analyzer** command.

```
user@switch> show forwarding-options analyzer
 Port mirror name : employee-monitor
 Mirror rate : 1
 Maximum packet length : 0
 State : up
 Ingress monitored interfaces : xe-0/0/0.0
 Ingress monitored interfaces : xe-0/0/6.0
 Output interface : xe-0/0/47.0
```

**Meaning** This output shows that the port-mirroring instance **employee-monitor** has a ratio of 1 (mirroring every packet, the default setting), the maximum size of the original packet that was mirrored (**0** indicates the entire packet), the state of the configuration (is up indicates that the instance is mirroring the traffic entering the xe-0/0/0, and xe-0/0/6 interfaces, and sending the mirrored traffic to the xe-0/0/47 interface). If the state of the output interface is down or if the output interface is not configured, the value of state will be **down** and the instance will not be programmed for mirroring.

**Related Documentation**

- [Understanding Port Mirroring on page 6227](#)

---

## Troubleshooting Port Mirroring

- [Port Mirroring Constraints and Limitations on page 6248](#)
- [Egress Port Mirroring with VLAN Translation on page 6250](#)
- [Egress Port Mirroring with Private VLANs on page 6250](#)

### Port Mirroring Constraints and Limitations

- [Local and Remote Port Mirroring on page 6248](#)
- [Remote Port Mirroring Only on page 6250](#)

---

#### Local and Remote Port Mirroring

The following constraints and limitations apply to local and remote port mirroring:

- You can create a total of four port-mirroring configurations.
- You can create a total of four port-mirroring configurations on each Node group in a QFabric system, subject to the following constraints:
  - As many as four of the configurations can be for local port mirroring.
  - As many as three of the configurations can be for remote port mirroring.
- Regardless of whether you are configuring a standalone switch or a Node group, the following limits apply:
  - There can be no more than two configurations that mirror ingress traffic. (If you configure a firewall filter to send traffic to a port mirror—that is, you use the **analyzer**



action modifier in a filter term—this counts as an ingress mirroring configuration for switch or Node group on which the filter is applied.)

- There can be no more than two configurations that mirror egress traffic.



**NOTE:** On QFabric systems, there is no system-wide limit on the total number of mirror sessions.

- You can configure no more than one type of output in one port-mirroring configuration. That is, you can use no more than one of the following to complete a **set analyzer name output** statement:
  - **interface**
  - **ip-address**
  - **vlan**
- If you configure Junos OS to mirror egress packets, do not configure more than 2000 VLANs on a standalone switch or QFabric system. If you do so, some VLAN packets might contain incorrect VLAN IDs. This applies to any VLAN packets—not only the mirrored copies.
- The **ratio** and **loss-priority** options are not supported.
- Packets with physical layer errors are filtered out and are not sent to the output port or VLAN.
- If you use sFlow monitoring to sample traffic, it does not sample the mirror copies when they exit from the output interface.
- You cannot mirror packets exiting or entering the following ports:
  - Dedicated Virtual Chassis interfaces
  - Management interfaces (me0 or vme0)
  - Fibre Channel interfaces
  - Routed VLAN interfaces
- An aggregated Ethernet interface cannot be an output interface if the input is a VLAN or if traffic is sent to the analyzer by a firewall filter.
- Do not include an 802.1Q subinterface that has a unit number other than 0 in a port mirroring configuration. Port mirroring does not work with subinterfaces if their unit number is not 0. (You configure 802.1Q subinterfaces using the **vlan-tagging** statement.)
- When packet copies are sent out the output interface, they are not modified for any changes that are normally applied on egress, such as CoS rewriting.
- An interface can be the input interface for only one mirroring configuration. Do not use the same interface as the input interface for multiple mirroring configurations.
- CPU-generated packets (such as ARP, ICMP, BPDU, and LACP packets) cannot be mirrored on egress.

- VLAN-based mirroring is not supported for STP traffic.
- (QFabric systems only) If you configure a QFabric analyzer to mirror egress traffic and the input and output interfaces are on different Node devices, the mirrored copies have incorrect VLAN IDs. This limitation does not apply if you configure a QFabric analyzer to mirror egress traffic and the input and output interfaces are on the *same* Node device. In this case the mirrored copies have the correct VLAN IDs (as long as you do not configure more than 2000 VLANs on the QFabric system).

### Remote Port Mirroring Only

---

The following constraints and limitations apply to remote port mirroring:

- If you configure an output IP address, the address cannot be in the same subnetwork as any of the switch's management interfaces.
- If you create virtual routing instances and also create an analyzer configuration that includes an output IP address, the output address belongs to the default virtual routing instance (inet.0 routing table).
- An output VLAN cannot be a private VLAN or VLAN range.
- An output VLAN cannot be shared by multiple **analyzer** statements.
- An output VLAN interface cannot be a member of any other VLAN.
- An output VLAN interface cannot be an aggregated Ethernet interface.

### Egress Port Mirroring with VLAN Translation

**Problem Description:** If you create a port-mirroring configuration that mirrors customer VLAN (CVLAN) traffic on egress and the traffic undergoes VLAN translation before being mirrored, the VLAN translation does not apply to the mirrored packets. That is, the mirrored packets retain the service VLAN (SVLAN) tag that should be replaced by the CVLAN tag on egress. The original packets are unaffected—on these packets VLAN translation works properly, and the SVLAN tag is replaced with the CVLAN tag on egress.

**Solution** This is expected behavior.

### Egress Port Mirroring with Private VLANs

**Problem Description:** If you create a port-mirroring configuration that mirrors private VLAN (PVLAN) traffic on egress, the mirrored traffic (the traffic that is sent to the analyzer system) has the VLAN tag of the ingress VLAN instead of the egress VLAN. For example, assume the following PVLAN configuration:

- Promiscuous trunk port that carries primary VLANs pvlan100 and pvlan400.
- Isolated access port that carries secondary VLAN isolated200. This VLAN is a member of primary VLAN pvlan100.

- Community port that carries secondary VLAN comm300. This VLAN is also a member of primary VLAN pvlan100.
- Output interface (monitor interface) that connects to the analyzer system. This interface forwards the mirrored traffic to the analyzer.

If a packet for pvlan100 enters on the promiscuous trunk port and exits on the isolated access port, the original packet is untagged on egress because it is exiting on an access port. However, the mirror copy retains the tag for pvlan100 when it is sent to the analyzer.

Here is another example: If a packet for comm300 ingresses on the community port and egresses on the promiscuous trunk port, the original packet carries the tag for pvlan100 on egress, as expected. However, the mirrored copy retains the tag for comm300 when it is sent to the analyzer.

**Solution** This is expected behavior.

**Related  
Documentation**

- [Understanding Port Mirroring on page 6227](#)
- [Examples: Configuring Port Mirroring for Local Analysis on page 6235](#)
- [Example: Configuring Port Mirroring for Remote Analysis on page 6240](#)
- [Example: Mirroring Employee Web Traffic with a Firewall Filter on page 6237](#)



## PART 85

# DHCP and DHCP Relay

- [Using DHCP and DHCP Relay on page 6255](#)



# Using DHCP and DHCP Relay

- [DHCP and BOOTP Relay Overview on page 6255](#)
- [Configuring DHCP and BOOTP on page 6255](#)
- [Verifying and Managing DHCP Local Server Configuration on page 6256](#)
- [Verifying and Managing DHCPv6 Local Server Configuration on page 6257](#)
- [Verifying and Managing DHCP Relay Configuration on page 6257](#)
- [Verifying and Managing DHCPv6 Relay Configuration on page 6257](#)
- [Understanding Layer 3 Logical Interfaces on page 6258](#)

## DHCP and BOOTP Relay Overview

---

You can configure a Juniper Networks switch to act as a Dynamic Host Configuration Protocol (DHCP) or Bootstrap Protocol (BOOTP) relay agent. This means that if the switch receives a broadcast DHCP or BOOTP request from a locally attached host (client), it relays the message to a specified DHCP or BOOTP server. You should configure the switch to be a DHCP/BOOTP relay agent if you have locally attached hosts and a distant DHCP or BOOTP server.



**NOTE:** Because DHCP and BOOTP messages are broadcast and are not directed to a specific server, switch, or router, Juniper switches cannot function as both a DHCP server and a DHCP/BOOTP relay agent at the same time. The Junos operating system (Junos OS) generates a commit error if both options are configured at the same time, and the commit operation does not succeed until one of the options is removed.

**Related Documentation**

- [Configuring DHCP and BOOTP Relay](#)

## Configuring DHCP and BOOTP

---

You can configure a switch to act as a Dynamic Host Configuration Protocol (DHCP) and Bootstrap Protocol (BOOTP) server or DHCP relay agent. When a switch is a relay agent, if a locally attached host issues a DHCP or BOOTP request as a broadcast message, the switch relays the message to a specified DHCP or BOOTP server. You should configure

a switch to be a DHCP and BOOTP relay agent if you have locally attached hosts and a remote DHCP or BOOTP server.



**NOTE:** This task uses the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring DHCP and BOOTP Relay*. For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on page 41.

To configure a switch to be a server, use the `dhcp-local-server` statement. To configure a switch to be a relay agent, use the `dhcp-relay` statement.

If you want to enable BOOTP support when the switch is configured to be a DHCP server, enter the following statement:

```
[edit system services dhcp-local-server]
user@switch# set overrides bootp-support
```

If you want to enable BOOTP support when the switch is configured to be a DHCP relay agent, enter the following statement:

```
[edit forwarding-options dhcp-relay]
user@switch# set overrides bootp-support
```

---

## Verifying and Managing DHCP Local Server Configuration

---

**Purpose** View or clear information about client address bindings and statistics for the extended DHCP local server.



**NOTE:** If you delete the DHCP server configuration, DHCP server bindings might still remain. To ensure that DHCP bindings are removed, issue the `clear dhcp server binding` command before you delete the DHCP server configuration.

- Action**
- To display the address bindings in the client table on the extended DHCP local server:  
user@host> `show dhcp server binding routing-instance customer routing instance`
  - To display extended DHCP local server statistics:  
user@host> `show dhcp server statistics routing-instance customer routing instance`
  - To clear the binding state of a DHCP client from the client table on the extended DHCP local server:  
user@host> `clear dhcp server binding routing-instance customer routing instance`
  - To clear all extended DHCP local server statistics:  
user@host> `clear dhcp server statistics routing-instance customer routing instance`

**Related Documentation**

- [CLI Explorer](#)



## Verifying and Managing DHCPv6 Local Server Configuration

---

**Purpose** View or clear information about client address bindings and statistics for the DHCPv6 local server.

**Action** • To display the address bindings in the client table on the DHCPv6 local server:

user@host> [show dhcpv6 server binding](#)

• To display DHCPv6 local server statistics:

user@host> [show dhcpv6 server statistics](#)

• To clear all DHCPv6 local server statistics:

user@host> [clear dhcpv6 server binding](#)

• To clear all DHCPv6 local server statistics:

user@host> [clear dhcpv6 server statistics](#)

**Related Documentation** • [CLI Explorer](#)

## Verifying and Managing DHCP Relay Configuration

---

**Purpose** View or clear address bindings or statistics for extended DHCP relay agent clients:

**Action** • To display the address bindings for extended DHCP relay agent clients:

user@host> [show dhcp relay binding](#) routing-instance *customer routing instance*

• To display extended DHCP relay agent statistics:

user@host> [show dhcp relay statistics](#) routing-instance *customer routing instance*

• To clear the binding state of DHCP relay agent clients:

user@host> [clear dhcp relay binding](#) routing-instance *customer routing instance*

• To clear all extended DHCP relay agent statistics:

user@host> [clear dhcp relay statistics](#) routing-instance *customer routing instance*

**Related Documentation** • [CLI Explorer](#)

## Verifying and Managing DHCPv6 Relay Configuration

---

**Purpose** View or clear address bindings or statistics for extended DHCPv6 relay agent clients:

**Action** • To display the address bindings for extended DHCPv6 relay agent clients:

user@host> [show dhcpv6 relay binding](#)

• To display extended DHCPv6 relay agent statistics:

user@host> [show dhcpv6 relay statistics](#)

- To clear the binding state of DHCPv6 relay agent clients:

```
user@host> clear dhcpv6 relay binding
```

- To clear all extended DHCPv6 relay agent statistics:

```
user@host> clear dhcpv6 relay statistics
```

**Related  
Documentation**

- [CLI Explorer](#)

---

## Understanding Layer 3 Logical Interfaces

---

A Layer 3 logical interface is a logical division of a physical interface that operates at the network level and therefore can receive and forward 802.1Q VLAN tags. You can use Layer 3 logical interfaces to route traffic among multiple VLANs along a single trunk line that connects a Juniper Networks switch to a Layer 2 switch. Only one physical connection is required between the switches.

To create Layer 3 logical interfaces on a switch, enable VLAN tagging, partition the physical interface into logical partitions, and bind the VLAN ID to the logical interface.

We recommend that you use the VLAN ID as the logical interface number when you configure the logical interface. QFX Series and EX4600 switches support a maximum of 4089 VLANs, which includes the default VLAN. You can, however, assign a VLAN ID in the range of 1 to 4094, but five of these VLAN IDs are reserved for internal use.

VLAN tagging places the VLAN ID in the frame header, allowing each physical interface to handle multiple VLANs. When you configure multiple VLANs on an interface, you must also enable tagging on that interface. Junos OS on switches supports a subset of the 802.1Q standard for receiving and forwarding routed or bridged Ethernet frames with single VLAN tags and running Virtual Router Redundancy Protocol (VRRP) over 802.1Q-tagged interfaces.

**Related  
Documentation**

- [Interfaces Overview on page 2785](#)
- [Configuring a Layer 3 Logical Interface on page 2860](#)
- *Junos OS Network Interfaces Library for Routing Devices*

## PART 86

# Configuration Statements and Operational Commands

- [Configuration Statements \(Port Mirroring\) on page 6261](#)
- [Configuration Statements \(DHCP and DHCP Relay\) on page 6277](#)
- [Configuration Statements \(Encryption\) on page 6297](#)
- [Operational Command \(Port Mirroring\) on page 6319](#)
- [Operational Commands \(DHCP Local Server\) on page 6323](#)
- [Operational Commands \(DHCP Relay Agent\) on page 6359](#)



## Configuration Statements (Port Mirroring)

- analyzer on page 6262
- egress on page 6263
- ethernet-switching (Port Mirroring) on page 6264
- family (Port Mirroring) on page 6265
- inet (Port Mirroring) on page 6266
- ingress (Port Mirroring) on page 6267
- input on page 6268
- instance (Port Mirroring) on page 6269
- interface (Port Mirroring) on page 6270
- ip-address (Port Mirroring) on page 6271
- no-tag on page 6271
- output on page 6272
- port-mirroring on page 6273
- routing-instance (Port Mirroring) on page 6274
- vlan (Port Mirroring) on page 6275

## analyzer

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre> analyzer {   name {     input {       egress {         interface (all   interface-name);         vlan (vlan-id   vlan-name);       }       ingress {         interface (all   interface-name);         vlan (vlan-id   vlan-name);       }       output {         interface interface-name;         ip-address ip-address;         routing-instance         vlan (vlan-id   vlan-name);       }     }   } } </pre> |
| <b>Hierarchy Level</b>     | <p>For platforms without ELS:</p> <pre>[edit ethernet-switching-options]</pre> <p>For platforms with ELS:</p> <pre>[edit forwarding-options]</pre>                                                                                                                                                                                                                                                                       |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Option <b>output vlan</b> added in Junos OS Release 12.1 for the QFX Series.</p> <p>Option <b>output ip-address</b> added in Junos OS Release 12.3 for the QFX Series for non-ELS platforms and added in 14.1X53-D10 for ELS platforms.</p>                                                                                                  |
| <b>Description</b>         | <p>Configure port mirroring. You can create a total of four port-mirroring configurations on the QFX Series, subject to the following limits:</p> <ul style="list-style-type: none"> <li>• There can be no more than two configurations that mirror ingress traffic.</li> <li>• There can be no more than two configurations that mirror egress traffic.</li> </ul>                                                      |
| <b>Default</b>             | Port mirroring is disabled, and Junos OS creates no default analyzers.                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>             | <b>all</b> —Mirror all the access interfaces. Using this option does not cause the QSFP+ or management interfaces to be mirrored.                                                                                                                                                                                                                                                                                        |



**CAUTION:** Configuring the **all** option in a QFabric system causes all the access interfaces on all the nodes to be mirrored. Be cautious about using this option on a QFabric system.

**name**—Name of the analyzer. The name can include as many as 125 characters; must begin with a letter; and can include uppercase letters, lowercase letters, numbers, dashes, and underscores. No other special characters are allowed.

The remaining statements are explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Understanding Port Mirroring on page 6227](#)
- [Configuring Port Mirroring on page 6232](#)
- [Examples: Configuring Port Mirroring for Local Analysis on page 6235](#)

## egress

**Syntax**

```
egress {
 interface (all | interface-name);
}
vlan (vlan-id | vlan-name);
```

**Hierarchy Level** For platforms without ELS:  
[edit ethernet-switching-options **analyzer name input**]  
For platforms with ELS:  
[edit forwarding-options **analyzer name input**]

**Release Information** Statement introduced in Junos OS Release 11.2 for the QFX Series.

**Description** Specify interface or VLAN for which egressing traffic is mirrored. (The **vlan** statement is not supported on all switches.)

The statement is explained separately.



**NOTE:** If you configure Junos OS to mirror egress packets, do not configure more than 2000 VLANs. If you do so, some of the mirrored packets might contain incorrect VLAN IDs.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Understanding Port Mirroring on page 6227](#)
- [Configuring Port Mirroring on page 6232](#)
- [Examples: Configuring Port Mirroring for Local Analysis on page 6235](#)

## ethernet-switching (Port Mirroring)

---

**Syntax** ethernet-switching;  
    **output** {  
        **interface** *interface-name* {  
        }  
        no-filter-check;  
        }  
    **vlan** *vlan-name* {  
        no-tag;  
    }  
}

**Hierarchy Level** [edit forwarding-options port-mirroring [instance *name*] family]

**Release Information** Statement introduced in Junos OS Release 13.2 for the QFX Series.

**Description** Specify that the output interface for the port mirror will be configured as an **ethernet-switching** interface.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Understanding Port Mirroring on page 6227](#)
- [Configuring Port Mirroring on page 6232](#)
- [Examples: Configuring Port Mirroring for Local Analysis on page 6235](#)



## family (Port Mirroring)

**Syntax** `family (Port Mirroring)`

```

ethernet-switching {
 output {
 interface interface-name {
 }
 no-filter-check;
 }
 vlan vlan-name {
 no-tag;
 }
}
inet
output {
 ip-address address {
 }
 routing-instance instance-name {
 ip-address address {
 }
 }
}

```

**Hierarchy Level** [edit forwarding-options port-mirroring [*instance name*] ]

**Release Information** Statement introduced in Junos OS Release 13.2 for the QFX Series.

**Description** Specify the type of interface that will be used to forward port mirrored packet to an analyzer device..

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Understanding Port Mirroring on page 6227](#)
- [Configuring Port Mirroring on page 6232](#)
- [Examples: Configuring Port Mirroring for Local Analysis on page 6235](#)

## inet (Port Mirroring)

---

**Syntax**

```
inet {
 output {
 ip-address address {
 }
 routing-instance instance-name {
 ip-address address {
 }
 }
 }
}
```

**Hierarchy Level** [edit forwarding-options port-mirroring [instance *name*] family]

**Release Information** Statement introduced in Junos OS Release 14.1X53 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Specify that the output interface will be of type **inet**. Use this statement so that you can send the mirrored packets to the IP address of the device running an analyzer application. The device can be on a remote network. When you use this feature, the mirrored packets are GRE-encapsulated. The analyzer device must be able to de-encapsulate GRE-encapsulated packets, or the GRE-encapsulated packets must be de-encapsulated before reaching the analyzer device. (You can use a network sniffer to de-encapsulate the packets.)



**NOTE:** An output IP address cannot be in the same subnetwork as any of the switch's management interfaces.

---



**NOTE:** If you create virtual routing instances and also create an analyzer configuration that includes an output IP address, the output address belongs to the default virtual routing instance (inet.0 routing table).

---

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

## ingress (Port Mirroring)

---

|                                 |                                                                                                                                                                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>ingress {   interface (all   interface-name);   vlan (vlan-id   vlan-name); }</pre>                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | <p>For platforms without ELS:</p> <pre>[edit ethernet-switching-options analyzer name input]</pre> <p>For platforms with ELS:</p> <pre>[edit forwarding-options analyzer name input]</pre>                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                |
| <b>Description</b>              | <p>Specify the interfaces or VLANs for which incoming traffic is mirrored as part of a port mirroring configuration.</p> <p>The statements are explained separately.</p>                                                                                                         |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.routing-control—To add this statement to the configuration.</p>                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding Port Mirroring on page 6227</a></li> <li>• <a href="#">Configuring Port Mirroring on page 6232</a></li> <li>• <a href="#">Examples: Configuring Port Mirroring for Local Analysis on page 6235</a></li> </ul> |

## input

---

|                                 |                                                                                                                                                                                                                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>input {<br/>  ingress {<br/>    interface (all   <i>interface-name</i>);<br/>    vlan (<i>vlan-id</i>   <i>vlan-name</i>);<br/>  }<br/>  egress {<br/>    interface (all   <i>interface-name</i>);<br/>  }<br/>}</pre>                                                  |
| <b>Hierarchy Level</b>          | <p>For platforms without ELS:</p> <p>[edit ethernet-switching-options <i>analyzer name</i>]</p> <p>For platforms with ELS:</p> <p>[edit forwarding-options <i>analyzer name</i>]</p>                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                            |
| <b>Description</b>              | <p>Define the traffic to be mirrored. The definition can be a combination of traffic entering or exiting specific ports or VLANs.</p> <p>The statements are explained separately.</p>                                                                                        |
| <b>Default</b>                  | No default.                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Port Mirroring on page 6227</a></li><li>• <a href="#">Configuring Port Mirroring on page 6232</a></li><li>• <a href="#">Examples: Configuring Port Mirroring for Local Analysis on page 6235</a></li></ul> |

## instance (Port Mirroring)

```
Syntax instance instance-name{
 family
 ethernet-switching {
 output {
 interface interface-name {
 }
 no-filter-check;
 }
 vlan vlan-name {
 no-tag;
 }
 }
 inet
 output {
 ip-address address {
 }
 routing-instance instance-name {
 ip-address address {
 }
 }
 }
 }
```

**Hierarchy Level** [edit forwarding-options port-mirroring]

**Release Information** Statement introduced in Junos OS Release 13.2 for the QFX Series.

**Description** Specify a port-mirroring configuration (instance). You do not specify an input for this instance. Instead, you, create a firewall filter that specifies the required traffic and directs it to the mirror. This is useful for controlling which types of traffic should be mirrored.


**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**



- [Understanding Port Mirroring on page 6227](#)
- [Configuring Port Mirroring on page 6232](#)
- [Examples: Configuring Port Mirroring for Local Analysis on page 6235](#)
- [Example: Mirroring Employee Web Traffic with a Firewall Filter on page 6237](#)

## interface (Port Mirroring)

---

|                                                                                                                                                                                                                                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                    | <code>interface (all   <i>interface-name</i>);</code>                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                           | <p>For platforms without ELS:</p> <pre>[edit ethernet-switching-options <i>analyzer name</i> input (egress   ingress)],<br/>[edit ethernet-switching-options <i>analyzer name</i> output]</pre> <p>For platforms with ELS:</p> <pre>[edit forwarding-options <i>analyzer name</i> input (egress   ingress)]<br/>[edit forwarding-options <i>analyzer name</i> output]<br/>[edit forwarding-options port-mirroring[instance <i>name</i>] family ethernet-switching output]</pre> |
| <b>Release Information</b>                                                                                                                                                                                                                                                                       | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>                                                                                                                                                                                                                                                                               | Specify the interfaces for which ingressing traffic is mirrored. Specify the interface that mirrored traffic should be copied to (the output interface).                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                                                                                                                                                                                                                                                                                   | <p><code>all</code>—Apply port mirroring to all interfaces on the switch (except the output interface). Mirroring a high volume of traffic can cause performance issues, so you should generally select specific input interfaces.</p>                                                                                                                                                                                                                                          |
| <div><p><b>CAUTION:</b> Configuring <code>all</code> in a QFabric system causes all the access interfaces on all the nodes to be mirrored. Be cautious about using this option on a QFabric system.</p></div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <p><i>interface-name</i>—Apply port mirroring to the specified interface only.</p>                                                                                                                                                                                                               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                  | <p><code>routing</code>—To view this statement in the configuration.</p> <p><code>routing-control</code>—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                     | <ul style="list-style-type: none"><li>• <a href="#">Understanding Port Mirroring on page 6227</a></li><li>• <a href="#">Configuring Port Mirroring on page 6232</a></li><li>• <a href="#">Examples: Configuring Port Mirroring for Local Analysis on page 6235</a></li></ul>                                                                                                                                                                                                    |

## ip-address (Port Mirroring)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>ip-address <i>ip-address</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | <code>[edit forwarding-options] <b>analyzer name</b> <b>output</b>]</code><br><code>[edit forwarding-options port-mirroring [instance <i>name</i>] family ethernet-switching <b>output</b> <b>interface name</b>]</code>                                                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1X53 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | Specify the IP address to which traffic should be mirrored (the IP address of the analyzer system). The device can be on a remote network. The analyzer device must be able to de-encapsulate GRE-encapsulated packets, or the GRE-encapsulated packets must be de-encapsulated before reaching the analyzer device. (You can use a network sniffer to de-encapsulate the packets.) This statement is not supported on QFabric systems. |
|                                 | <div>  <p><b>NOTE:</b> An output IP address cannot be in the same subnetwork as any of the switch's management interfaces.</p> </div>                                                                                                                                                                                                                  |
|                                 | <div>  <p><b>NOTE:</b> If you create virtual routing instances and also create an analyzer configuration that includes an output IP address, the output address belongs to the default virtual routing instance (inet.0 routing table).</p> </div>                                                                                                   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                     |

## no-tag

|                                 |                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>no-tag;</code>                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | <code>[edit forwarding-options <b>analyzer name</b> <b>output</b> <b>vlan</b>]</code><br><code>[edit forwarding-options port-mirroring[instance <i>name</i>] family ethernet-switching <b>output</b> <b>vlan</b>]</code> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 15.1X53-D10 for QFX10000 switches.                                                                                                                                              |
| <b>Description</b>              | Specify that remote mirrored packets are not tagged with the tag of the output (analyzer) VLAN.                                                                                                                          |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Port Mirroring for Remote Analysis on page 6240</a></li> </ul>                                                                                 |

## output

---

|                                 |                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>output {<br/>    interface <i>interface-name</i>;<br/>    ip-address <i>ip-address</i>;<br/>    vlan (<i>vlan-id</i>   <i>vlan-name</i>);<br/>    routing-instance <i>instance-name</i> {<br/>        ip-address <i>address</i> {<br/>        }<br/>    }<br/>}</pre>             |
| <b>Hierarchy Level</b>          | <p>For platforms without ELS:</p> <p>[edit ethernet-switching-options <i>analyzer name</i>]</p> <p>For platforms with ELS:</p> <p>[edit forwarding-options <i>analyzer name</i>]</p> <p>[edit forwarding-options port-mirroring [instance <i>name</i>] family ethernet-switching ]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Option <b>output vlan</b> added in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                           |
| <b>Description</b>              | <p>Configure the destination for mirrored traffic, either an interface on the switch (for local monitoring) or a VLAN (for remote monitoring).</p> <p>The statements are explained separately.</p>                                                                                     |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                         |



## port-mirroring

```

Syntax port-mirroring {
 family {
 ethernet-switching
 output {
 interface interface-name {
 }
 no-filter-check;
 }
 vlan vlan-name {
 no-tag;
 }
 }
 inet
 output {
 ip-address address {
 }
 routing-instance instance-name {
 ip-address address {
 }
 }
 }
 }
 instance instance-name {
 family (Port Mirroring)
 ethernet-switching {
 output {
 interface interface-name {
 }
 no-filter-check;
 }
 vlan vlan-name {
 no-tag;
 }
 }
 inet
 output {
 ip-address address {
 }
 routing-instance instance-name {
 ip-address address {
 }
 }
 }
 }
}

```

**Hierarchy Level** [edit forwarding-options ]

**Release Information** Statement introduced in Junos OS Release 13.2 for the QFX Series.

**Description** Create a port-mirroring configuration.

**Required Privilege** routing—To view this statement in the configuration.  
**Level** routing-control—To add this statement to the configuration.

**Related Documentation**

- [Understanding Port Mirroring on page 6227](#)
- [Configuring Port Mirroring on page 6232](#)
- [Examples: Configuring Port Mirroring for Local Analysis on page 6235](#)
- [Example: Mirroring Employee Web Traffic with a Firewall Filter on page 6237](#)

---

## routing-instance (Port Mirroring)

---

**Syntax** routing-instance *instance-name*;

**Hierarchy Level** [edit forwarding-options] [analyzer name](#) [output](#)]  
[edit forwarding-options port-mirroring [instance *name*] family inet [output interface name](#)]

**Release Information** Statement introduced in Junos OS Release 12.3 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Configure a port mirroring instance. You do not specify an input for this instance. Instead, you, create a firewall filter that specifies the required traffic and directs it to the mirror. This instance type is useful for controlling which types of traffic should be mirrored.

**Required Privilege** routing—To view this statement in the configuration.  
**Level** routing-control—To add this statement to the configuration.

## vlan (Port Mirroring)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>vlan (vlan-id   vlan-name) {<br/>    no-tag;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | <p>For platforms without ELS:</p> <pre>[edit ethernet-switching-options analyzer name input ingress], [edit ethernet-switching-options analyzer name output]</pre> <p>For platforms with ELS:</p> <pre>[edit forwarding-options analyzer name input (egress   ingress)] [edit forwarding-options analyzer name output] [edit forwarding-options port-mirroring[instance name] family ethernet-switching output]</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Option <b>output</b> <b>vlan</b> added in Junos OS Release 12.1 for the QFX Series.</p> <p>Option <b>no-tag</b> added in Junos OS Release 15.1X53-D10 for QFX10000 switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | <p>When used in an <b>input</b> statement, specifies that traffic entering or exiting a VLAN should be mirrored. (You can include this statement in an <b>ingress</b> statement or <b>egress</b> statement within the <b>input</b> statement. It is not supported in an <b>egress</b> statement on all switches)</p> <p>When used in an <b>output</b> statement, specifies that mirrored traffic to be sent to a VLAN for remote monitoring.</p> <p>On some switches, only one interface can be a member of the output (analyzer) VLAN. This limitation does not apply on the QFX10000 switch if traffic is mirrored on ingress. In this case, multiple QFX10000 interfaces can belong to the output VLAN, and traffic is mirrored to all of those interfaces. If traffic is mirrored on egress on a QFX10000 switch, only one interface can be a member of the analyzer VLAN.</p> |
| <b>Options</b>                  | <p><b>vlan-id</b>—Numeric VLAN identifier.</p> <p><b>vlan-name</b>—Name of the VLAN.</p> <p><b>no-tag</b>—Specifies that remote mirrored packets are not tagged with the tag of the output (analyzer) VLAN.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | <p><b>routing</b>—To view this statement in the configuration.</p> <p><b>routing-control</b>—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding Port Mirroring on page 6227</a></li> <li>• <a href="#">Configuring Port Mirroring on page 6232</a></li> <li>• <a href="#">Examples: Configuring Port Mirroring for Local Analysis on page 6235</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |



## CHAPTER 239

# Configuration Statements (DHCP and DHCP Relay)

- [dhcp-local-server](#) on page 6278
- [dhcp-relay](#) on page 6284
- [dhcp-service](#) on page 6289
- [forwarding-options](#) on page 6290
- [forward-only \(DHCP Relay Agent\)](#) on page 6295

## dhcp-local-server

---

```
Syntax dhcp-local-server {
 access-profile profile-name;
 authentication {
 password password-string;
 username-include {
 circuit-type;
 delimiter delimiter-character;
 domain-name domain-name-string;
 interface-name;
 logical-system-name;
 mac-address;
 option-60;
 option-82 <circuit-id> <remote-id>;
 routing-instance-name;
 user-prefix user-prefix-string;
 }
 }
 }
 dhcpv6 {
 access-profile profile-name;
 authentication {
 ...
 }
 group group-name {
 access-profile profile-name;
 authentication {
 ...
 }
 }
 interface interface-name {
 access-profile profile-name;
 exclude;
 liveness-detection {
 failure-action (clear-binding | clear-binding-if-interface-up | log-only);
 method {
 bfd {
 version (0 | 1 | automatic);
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 multiplier number;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 }
 detection-time {
 threshold milliseconds;
 }
 }
 session-mode (automatic | multihop | singlehop);
 holddown-interval milliseconds;
 }
 }
 }
 overrides {
```

```

 interface-client-limit number;
 multi-address-embedded-option-response;
 process-inform {
 pool pool-name;
 }
 rapid-commit;
 }
 service-profile dynamic-profile-name;
 trace;
 upto upto-interface-name;
}
overrides {
 delegated-pool;
 interface-client-limit number;
 multi-address-embedded-option-response;
 process-inform {
 pool pool-name;
 }
 rapid-commit;
}
route-suppression;
service-profile dynamic-profile-name;
}
liveness-detection {
 failure-action (clear-binding | clear-binding-if-interface-up | log-only);
 method {
 bfd {
 version (0 | 1 | automatic);
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 multiplier number;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 detection-time {
 threshold milliseconds;
 }
 session-mode (automatic | multihop | singlehop);
 holddown-interval milliseconds;
 }
 }
}
overrides {
 delegated-pool;
 include-option-82 {
 forcerenew;
 nak;
 }
 interface-client-limit number;
 multi-address-embedded-option-response;
 process-inform {
 pool pool-name;
 }
 rapid-commit;
}

```

```
}
reconfigure {
 attempts attempt-count;
 clear-on-abort;
 strict;
 timeout timeout-value;
 token token-value;
 trigger {
 radius-disconnect;
 }
}
route-suppression;
service-profile dynamic-profile-name;
}
duplicate-clients-in-subnet (incoming-interface | option-82);
dynamic-profile profile-name <aggregate-clients (merge | replace) | use-primary
 primary-profile-name>;
forward-snooped-clients (all-interfaces | configured-interfaces |
 non-configured-interfaces);
group group-name {
 authentication {
 ...
 }
 dynamic-profile profile-name <aggregate-clients (merge | replace) | use-primary
 primary-profile-name>;
 interface interface-name {
 exclude;
 liveness-detection {
 failure-action (clear-binding | clear-binding-if-interface-up | log-only);
 method {
 bfd {
 version (0 | 1 | automatic);
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 multiplier number;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 detection-time {
 threshold milliseconds;
 }
 session-mode (automatic | multihop | singlehop);
 holddown-interval milliseconds;
 }
 }
 }
 }
}
overrides {
 client-discover-match (option60-and-option82 | incoming-interface);
 include-option-82 {
 forcerenew;
 nak;
 }
 interface-client-limit number;
 process-inform {
```



```

 pool pool-name;
 }
}
service-profile dynamic-profile-name;
trace;
upto upto-interface-name;
}
overrides {
 client-discover-match (option60-and-option82 | incoming-interface);
 include-option-82 {
 forcerenew;
 nak;
 }
 interface-client-limit number;
 process-inform {
 pool pool-name;
 }
}
requested-ip-network-match subnet-mask
route-suppression;
service-profile dynamic-profile-name;
}
liveness-detection {
 failure-action (clear-binding | clear-binding-if-interface-up | log-only);
 method {
 bfd {
 version (0 | 1 | automatic);
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 multiplier number;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 detection-time {
 threshold milliseconds;
 }
 session-mode (automatic | multihop | singlehop);
 holddown-interval milliseconds;
 }
 }
}
overrides {
 client-discover-match <option60-and-option82 | incoming-interface>;
 interface-client-limit number;
 process-inform {
 pool pool-name;
 }
}
pool-match-order {
 external-authority;
 ip-address-first;
 option-82;
}
reconfigure {

```

```

 attempts attempt-count;
 clear-on-abort;
 strict;
 timeout timeout-value;
 token token-value;
 trigger {
 radius-disconnect;
 }
}
requested-ip-network-match subnet-mask;
route-suppression;
service-profile dynamic-profile-name;
}

```

**Hierarchy Level** [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services],  
 [edit logical-systems *logical-system-name* system services],  
 [edit routing-instances *routing-instance-name* system services],  
 [edit system services]

**Release Information** Statement introduced in Junos OS Release 9.0.  
 Statement introduced in Junos OS Release 12.1 for EX Series switches.  
 Statement introduced in Junos OS Release 13.2X51 for the QFX Series.  
 Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Configure Dynamic Host Configuration Protocol (DHCP) local server options on the router or switch and enable the router or switch to function as an extended DHCP local server. The DHCP local server receives DHCP request and reply packets from DHCP clients and then responds with an IP address and other optional configuration information to the client.

The extended DHCP local server is incompatible with the DHCP server on J Series routers and so is not supported on J Series routers. Also, the DHCP local server and the DHCP/BOOTP relay server, which are configured under the **[edit forwarding-options helpers]** hierarchy level, cannot both be enabled on the router or switch at the same time. The extended DHCP local server is fully compatible with the extended DHCP relay feature.

The **dhcpv6** stanza configures the router or switch to support Dynamic Host Configuration Protocol for IPv6 (DHCPv6). The DHCPv6 local server is fully compatible with the extended DHCP local server and the extended DHCP relay feature.



**NOTE:** When you configure the **dhcp-local-server** statement at the routing instance hierarchy level, you must use a routing instance type of **virtual-router**.

The remaining statements are explained separately.

**Required Privilege Level** system—To view this statement in the configuration.  
 system-control—To add this statement to the configuration.

- Related Documentation**
- *Extended DHCP Local Server Overview*
  - *DHCPv6 Local Server Overview*

## dhcp-relay

```

Syntax dhcp-relay {
 active-server-group server-group-name;
 authentication {
 password password-string;
 username-include {
 circuit-type;
 delimiter delimiter-character;
 domain-name domain-name-string;
 interface-name;
 logical-system-name;
 mac-address;
 option-60;
 option-82 <circuit-id> <remote-id>;
 routing-instance-name;
 user-prefix user-prefix-string;
 }
 }
 }
 dhcpv6 {
 active-server-group server-group-name;
 authentication {
 password password-string;
 username-include {
 circuit-type;
 client-id;
 delimiter delimiter-character;
 domain-name domain-name-string;
 interface-name;
 logical-system-name;
 relay-agent-interface-id;
 relay-agent-remote-id;
 relay-agent-subscriber-id;
 routing-instance-name;
 user-prefix user-prefix-string;
 }
 }
 dynamic-profile profile-name {
 aggregate-clients (merge | replace);
 use-primary primary-profile-name;
 }
 }
}

group group-name {
 active-server-group server-group-name;
 authentication {
 ...
 }
 dynamic-profile profile-name {
 ...
 }
 interface interface-name {
 exclude;
 liveness-detection {
 failure-action (clear-binding | clear-binding-if-interface-up | log-only);
 }
 }
}

```

```

method {
 bfd {
 version (0 | 1 | automatic);
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 multiplier number;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 detection-time {
 threshold milliseconds;
 }
 session-mode (automatic | multihop | singlehop);
 holddown-interval milliseconds;
 }
}
overrides {
 ...
}
service-profile dynamic-profile-name;
trace;
upto upto-interface-name;
}
service-profile dynamic-profile-name;
}
overrides {
 ...
}
relay-agent-interface-id {
 ...
}
service-profile dynamic-profile-name;
liveness-detection {
 failure-action (clear-binding | clear-binding-if-interface-up | log-only);
 method {
 bfd {
 version (0 | 1 | automatic);
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 multiplier number;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 detection-time {
 threshold milliseconds;
 }
 session-mode (automatic | multihop | singlehop);
 holddown-interval milliseconds;
 }
 }
}
}

```

```
overrides {
 allow-snooped-clients;
 interface-client-limit number;
 no-allow-snooped-clients;
 no-bind-on-request;
 send-release-on-delete;
}
relay-agent-interface-id {
 prefix prefix;
 use-interface-description (logical | device);
}
server-group {
 server-group-name {
 server-ip-address;
 }
}
dynamic-profile profile-name {
 aggregate-clients (merge | replace);
 use-primary primary-profile-name;
}
forward-snooped-clients (all-interfaces | configured-interfaces |
 non-configured-interfaces);
group group-name {
 active-server-group server-group-name;
 authentication {
 ...
 }
}
dynamic-profile profile-name {
 ...
}
interface interface-name {
 exclude;
 liveness-detection {
 failure-action (clear-binding | clear-binding-if-interface-up | log-only);
 method {
 bfd {
 version (0 | 1 | automatic);
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 multiplier number;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 detection-time {
 threshold milliseconds;
 }
 session-mode (automatic | multihop | singlehop);
 holddown-interval milliseconds;
 }
 }
 }
}
overrides {
 ...
}
```

```

 service-profile dynamic-profile-name;
 trace;
 upto upto-interface-name;
 }
 overrides {
 ...
 }
 relay-option-82 {
 ...
 }
 service-profile dynamic-profile-name;
}
liveness-detection {
 failure-action (clear-binding | clear-binding-if-interface-up | log-only);
 method {
 bfd {
 version (0 | 1 | automatic);
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 multiplier number;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 detection-time {
 threshold milliseconds;
 }
 session-mode(automatic | multihop | singlehop);
 holddown-interval milliseconds;
 }
 }
}
overrides {
 allow-snooped-clients;
 always-write-giaddr;
 always-write-option-82;
 client-discover-match <option60-and-option82>;
 disable-relay;
 interface-client-limit number;
 layer2-unicast-replies;
 no-allow-snooped-clients;
 no-bind-on-request;
 proxy-mode;
 replace-ip-source-with;
 send-release-on-delete;
 trust-option-82;
}
relay-option-82 {
 circuit-id {
 prefix prefix;
 use-interface-description (logical | device);
 }
}
server-group {
 server-group-name {

```

```
 server-ip-address;
 }
}
service-profile dynamic-profile-name;
}
```

**Hierarchy Level** [edit forwarding-options],  
[edit vlans forwarding-options]

**Release Information** Statement introduced in Junos OS Release 11.3 for the QFX Series.

**Description** Configure extended Dynamic Host Configuration Protocol (DHCP) relay and DHCPv6 relay options on the switch and enable the switch to function as a DHCP relay agent. A DHCP relay agent forwards DHCP request and reply packets between a DHCP client and a DHCP server.

DHCP relay supports the attachment of dynamic profiles and also interacts with the local AAA Service Framework to use back-end authentication servers, such as RADIUS, to provide subscriber authentication. You can attach dynamic profiles and configure authentication support on a global basis or for a specific group of interfaces.

The extended DHCP and DHCPv6 relay agent options configured with the **dhcp-relay** and **dhcpv6** statements are incompatible with the DHCP/BOOTP relay agent options configured with the **bootp** statement. As a result, the extended DHCP or DHCPv6 relay agent and the DHCP/BOOTP relay agent cannot both be enabled on the router at the same time.

The remaining statements are explained separately.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring DHCP and BOOTP on page 6255](#)



## dhcp-service

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>dhcp-service {     dhcp-snooping-file (<i>local_pathname</i>   <i>remote_URL</i>);     write-interval <i>interval</i>; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit system processes]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1 for MX Series routers.</p>                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | <p>Enable DHCP services on the device. DHCP services automate network-parameter assignment to network devices. The DHCP service process is enabled by default. However, by default, IP-MAC bindings in the DHCP snooping database do not persist through device reboots. You can improve performance after rebooting by configuring the IP-MAC bindings to persist, by configuring a storage location for the DHCP database file. When specifying the location for the DHCP database, you must also specify how frequently the switch writes the database entries into the DHCP snooping database file.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Configuring Persistent Bindings in the DHCP or DHCPv6 Snooping Database to Improve Network Performance (CLI Procedure)</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

## forwarding-options

```
Syntax forwarding-options {
 dhcp-relay {
 active-server-group server-group-name;
 authentication {
 password password-string;
 username-include {
 circuit-type;
 delimiter delimiter-character;
 domain-name domain-name-string;
 interface-name;
 logical-system-name;
 mac-address;
 option-60;
 option-82 <circuit-id> <remote-id>;
 routing-instance-name;
 user-prefix user-prefix-string;
 }
 }
 }
 dhcpv6 {
 active-server-group server-group-name;
 authentication {
 password password-string;
 username-include {
 circuit-type;
 client-id;
 delimiter delimiter-character;
 domain-name domain-name-string;
 interface-name;
 logical-system-name;
 relay-agent-interface-id;
 relay-agent-remote-id;
 relay-agent-subscriber-id;
 routing-instance-name;
 user-prefix user-prefix-string;
 }
 }
 }
 dynamic-profile profile-name {
 aggregate-clients (merge | replace);
 use-primary primary-profile-name;
 }
 group group-name {
 active-server-group server-group-name;
 authentication {
 ...
 }
 dynamic-profile profile-name {
 ...
 }
 }
 interface interface-name {
 exclude;
 liveness-detection {
 failure-action (clear-binding | clear-binding-if-interface-up | log-only);
 }
 }
}
```

```

method {
 bfd {
 version (0 | 1 | automatic);
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 multiplier number;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 detection-time {
 threshold milliseconds;
 }
 session-mode(automatic | multihop | singlehop);
 holddown-interval milliseconds;
 }
}
overrides {
 ...
}
service-profile dynamic-profile-name;
trace;
upto upto-interface-name;
}
service-profile dynamic-profile-name;
}
overrides {
 ...
}
relay-agent-interface-id {
 ...
}
service-profile dynamic-profile-name;
liveness-detection {
 failure-action (clear-binding | clear-binding-if-interface-up | log-only);
 method {
 bfd {
 version (0 | 1 | automatic);
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 multiplier number;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 detection-time {
 threshold milliseconds;
 }
 session-mode(automatic | multihop | singlehop);
 holddown-interval milliseconds;
 }
 }
}
}

```

```
overrides {
 allow-snooped-clients;
 interface-client-limit number;
 no-allow-snooped-clients;
 no-bind-on-request;
 send-release-on-delete;
}
relay-agent-interface-id {
 prefix prefix;
 use-interface-description (logical | device);
}
server-group {
 server-group-name {
 server-ip-address;
 }
}
dynamic-profile profile-name {
 aggregate-clients (merge | replace);
 use-primary primary-profile-name;
}
forward-snooped-clients (all-interfaces | configured-interfaces |
non-configured-interfaces);
group group-name {
 active-server-group server-group-name;
 authentication {
 ...
 }
}
dynamic-profile profile-name {
 ...
}
interface interface-name {
 exclude;
 liveness-detection {
 failure-action (clear-binding | clear-binding-if-interface-up | log-only);
 method {
 bfd {
 version (0 | 1 | automatic);
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 multiplier number;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 detection-time {
 threshold milliseconds;
 }
 }
 session-mode (automatic | multihop | singlehop);
 holddown-interval milliseconds;
 }
 }
}
overrides {
 ...
}
```

```

 service-profile dynamic-profile-name;
 trace;
 upto upto-interface-name;
}
overrides {
 ...
}
relay-option-60 {
 ...
}
relay-option-82 {
 ...
}
service-profile dynamic-profile-name;
}
liveness-detection {
 failure-action (clear-binding | clear-binding-if-interface-up | log-only);
 method {
 bfd {
 version (0 | 1 | automatic);
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 multiplier number;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 detection-time {
 threshold milliseconds;
 }
 session-mode (automatic | multihop | singlehop);
 holddown-interval milliseconds;
 }
 }
}
overrides {
 allow-snooped-clients;
 always-write-giaddr;
 always-write-option-82;
 client-discover-match <option60-and-option82>;
 disable-relay;
 interface-client-limit number;
 layer2-unicast-replies;
 no-allow-snooped-clients;
 no-bind-on-request;
 proxy-mode;
 replace-ip-source-with;
 send-release-on-delete;
 trust-option-82;
}
relay-option-82 {
 circuit-id {
 prefix prefix;
 use-interface-description (logical | device);
 }
}

```

```
 }
 server-group {
 server-group-name {
 server-ip-address;
 }
 }
 service-profile dynamic-profile-name;
}
dhcp-security {
 arp-inspection;
 group group-name {
 interface interface-name {
 static-ip ip-address {
 mac mac-address;
 }
 }
 overrides {
 no-option82;
 trusted;
 untrusted;
 }
 }
}
ip-source-guard;
no-dhcp-snooping;
option-82 {
 circuit-id {
 prefix {
 host-name;
 logical-system-name;
 routing-instance-name;
 }
 use-interface-description (device | logical);
 use-vlan-id;
 }
 remote-id {
 host-name hostname;
 use-interface-description (device | logical);
 use-string string;
 }
 vendor-id {
 use-string string;
 }
}
}
fip-security {
 examine-vn2vf;
 examine-vn2vn {
 beacon-period milliseconds;
 }
 fc-map fc-map-value;
 interface interface-name {
 (fcoe-trusted | no-fcoe-trusted;)
 }
}
}
```

|                                 |                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| <b>Hierarchy Level</b>          | [edit]<br>[edit vlans]                                                                                                      |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 11.3 for QFX Series switches. |
| <b>Description</b>              | Configure traffic forwarding.<br><br>The statements are explained separately.                                               |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.     |

### forward-only (DHCP Relay Agent)

|                                 |                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | forward-only {<br>routing-instance <current   default   <i>routing-instance-name</i> >;<br>}                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit forwarding-options dhcp-relay],<br>[edit forwarding-options dhcp-relay dhcpv6],<br>[edit forwarding-options dhcp-relay group <i>group-name</i> ],<br>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1X53-D25 for QFX Series.                                                                                                                                                                                                                                                        |
| <b>Description</b>              | Specify the location of the DHCP server when configuring secure DHCP traffic between the DHCP server and DHCP client when the two reside in different virtual routing instances.                                                                                                                                            |
| <b>Default</b>                  | Routing instance from where the configuration is applied.                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <b>routing-instance</b> —(Optional) Routing instance in which the DHCP server resides. <ul style="list-style-type: none"> <li><b>current</b>—Routing instance from which the configuration is applied.</li> <li><b>default</b>—Root routing instance.</li> </ul>                                                            |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>DHCP Message Exchange Between DHCP Clients and DHCP Server in Different VRFs</i></li> <li><i>Configuring DHCP Message Exchange Between DHCP Server and Clients in Different Virtual Routing Instances</i></li> </ul>                                                              |





## Configuration Statements (Encryption)

- [authentication-key-chains](#) on page 6298
- [cache-size](#) on page 6299
- [cache-timeout-negative](#) on page 6300
- [ca-name](#) on page 6301
- [certificates](#) on page 6302
- [certification-authority](#) on page 6303
- [crl \(Encryption Interface\)](#) on page 6304
- [encoding](#) on page 6304
- [enrollment-retry](#) on page 6305
- [enrollment-url](#) on page 6305
- [file](#) on page 6306
- [key \(Authentication Keychain\)](#) on page 6307
- [key-chain \(Security\)](#) on page 6308
- [ldap-url](#) on page 6309
- [local](#) on page 6310
- [maximum-certificates](#) on page 6311
- [path-length](#) on page 6311
- [secret](#) on page 6312
- [security](#) on page 6313
- [ssh-known-hosts](#) on page 6314
- [start-time \(Authentication Key Transmission\)](#) on page 6315
- [traceoptions](#) on page 6317

## authentication-key-chains

---

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <pre>authentication-key-chains {<br/>  key-chain <i>key-chain-name</i> {<br/>    description <i>text-string</i>;<br/>    key <i>key</i> {<br/>      algorithm (md5   hmac-sha-1);<br/>      options (basic   isis-enhanced);<br/>      secret <i>secret-data</i>;<br/>      start-time <i>yyyy-mm-dd.hh:mm:ss</i>;<br/>    }<br/>    tolerance <i>seconds</i>;<br/>  }<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Hierarchy Level          | [edit security]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Release Information      | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Support for IS-IS introduced in JUNOS OS Release 11.2.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                          |
| Description              | <p>Configure authentication key updates for the Border Gateway Protocol (BGP), the Label Distribution Protocol (LDP) routing protocols, the Bidirectional Forwarding Detection (BFD) protocol, and the Intermediate System-to-Intermediate System (IS-IS) protocol. When the <b>authentication-key-chains</b> statement is configured at the <b>[edit security]</b> hierarchy level, and is associated with the BGP, LDP, or IS-IS protocols at the <b>[edit protocols]</b> hierarchy level or with the BFD protocol using the <b>bfd-liveness-detection</b> statement, authentication key updates can occur without interrupting routing and signaling protocols such as Open Shortest Path First (OSPF) and Resource Reservation Setup Protocol (RSVP).</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Related Documentation    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols</a></li><li>• <a href="#">Example: Configuring BFD Authentication for Securing Static Routes on page 3255</a></li><li>• <a href="#">Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 4256</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                             |

## cache-size

|                            |                                                                                                                                                                                                                                                                            |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | cache-size <i>bytes</i> ;                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>     | [edit security <a href="#">certificates</a> ]                                                                                                                                                                                                                              |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| <b>Description</b>         | (Encryption interface on M Series and T Series routers and EX Series switches only)<br>Configure the cache size for digital certificates.                                                                                                                                  |
| <b>Options</b>             | <b>bytes</b> —Cache size for digital certificates.<br><b>Range:</b> 64 through 4,294,967,295<br><b>Default:</b> 2 megabytes (MB)                                                                                                                                           |



**NOTE:** We recommend that you limit your cache size to 4 MB.

|                                 |                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Configuring Digital Certificates for an ES PIC</i></li> </ul>        |

## cache-timeout-negative

---

|                            |                                                                                                                                                                                                                                                                            |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | cache-timeout-negative <i>seconds</i> ;                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>     | [edit security <a href="#">certificates</a> ]                                                                                                                                                                                                                              |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| <b>Description</b>         | (Encryption interface on M Series and T Series routers and EX Series switches only)<br>Configure a negative cache for digital certificates.                                                                                                                                |
| <b>Options</b>             | <b>seconds</b> —Negative time to cache digital certificates, in seconds.<br><b>Range:</b> 10 through 4,294,967,295<br><b>Default:</b> 20                                                                                                                                   |



**CAUTION:** Configuring a large negative cache value can lead to a denial-of-service attack.

---

|                                 |                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Digital Certificates for an ES PIC</i></li></ul>        |

---

## ca-name

---

|                                 |                                                                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>ca-name <i>ca-identity</i>;</code>                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit security certificates <a href="#">certification-authority</a> ]                                                                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series |
| <b>Description</b>              | (Encryption interface on M Series and T Series routers and EX Series switches only)<br>Specify the certificate authority (CA) identity to use in the certificate request.                                                                                                 |
| <b>Options</b>                  | <i>ca-identity</i> —CA identity to use in the certificate request.                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Digital Certificates for an ES PIC</i></li></ul>                                                                                                                                                                   |

## certificates

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>certificates {<br/>  cache-size bytes;<br/>  cache-timeout-negative seconds;<br/>  certification-authority ca-profile-name {<br/>    ca-name ca-identity;<br/>    crt file-name;<br/>    encoding (binary   pem);<br/>    enrollment-url url-name;<br/>    file certificate-filename;<br/>    ldap-url url-name;<br/>  }<br/>  enrollment-retry attempts;<br/>  local certificate-name {<br/>    certificate-key-string;<br/>    load-key-file URL filename;<br/>  }<br/>  maximum-certificates number;<br/>  path-length certificate-path-length;<br/>}</pre> |
| <b>Hierarchy Level</b>          | [edit security]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | (Encryption interface on M Series and T Series routers and EX Series switches only)<br>Configure the digital certificates for IPsec.<br><br>The remaining statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>Configuring Digital Certificates for an ES PIC</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## certification-authority

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>certification-authority <i>ca-profile-name</i> {   <i>ca-name</i> <i>ca-identity</i>;   <i>crl</i> <i>file-name</i>;   <i>encoding</i> (binary   pem);   <i>enrollment-url</i> <i>url-name</i>;   <i>file</i> <i>certificate-filename</i>;   <i>ldap-url</i> <i>url-name</i>; }</pre>                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit security <a href="#">certificates</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced before Junos OS Release 12.1 for the SRX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                   |
| <b>Description</b>              | <p>(Encryption interface on M Series and T Series routers and EX Series switches only)</p> <p>Configure a certificate authority profile name.</p> <p>Configure certification authority (CA) for X.509 certificate.</p>                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <ul style="list-style-type: none"> <li>• <i>profile-name</i>—Name of this CA configuration.</li> <li>• <i>ca-name</i> <i>name</i>—Name of the CA.</li> <li>• <i>crl</i> <i>filename</i>—Certificate revocation list (CRL) filename.</li> <li>• <i>encoding</i>—Certificate encoding, either <b>binary</b> or <b>pem</b> (privacy-enhanced mail).</li> <li>• <i>enrollment-url</i> <i>url</i>—Enrollment URL.</li> <li>• <i>file</i> <i>filename</i>—Certificate filename.</li> <li>• <i>ldap-url</i> <i>url</i>—Lightweight Directory Access Protocol (LDAP) URL.</li> </ul> |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Digital Certificates for an ES PIC</i></li> <li>• <i>Network Monitoring and Troubleshooting Guide for Security Devices</i></li> <li>• <i>Security Basics</i></li> <li>• <i>Configuring Digital Certificates for an ES PIC</i></li> </ul>                                                                                                                                                                                                                                                                             |

## crl (Encryption Interface)

---

|                                 |                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>crl file-name;</code>                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit security <a href="#">certificates</a> ]                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                      |
| <b>Description</b>              | (Encryption interface on M Series and T Series routers and EX Series switches only)<br>Configure the certificate revocation list (CRL). A CRL is a time-stamped list identifying revoked certificates, which is signed by a CA and made available to the participating IPsec peers on a regular periodic basis. |
| <b>Options</b>                  | <b>file-name</b> —Specify the file from which to read the CRL.                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Digital Certificates for an ES PIC</i></li></ul>                                                                                                                                                                                                         |

## encoding

---

|                                 |                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>encoding (binary   pem);</code>                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit security ike policy <i>ike-peer-address</i> ],<br>[edit security certificates <a href="#">certification-authority</a> <i>ca-profile-name</i> ]                                                                                                                       |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| <b>Description</b>              | (Encryption interface on M Series and T Series routers and EX Series switches only)<br>Specify the file format used for the <b>local-certificate</b> and <b>local-key-pair</b> statements.                                                                                 |
| <b>Options</b>                  | <b>binary</b> —Binary file format.<br><br><b>pem</b> —Privacy-enhanced mail (PEM), an ASCII base 64 encoded format.<br><b>Default:</b> <b>binary</b>                                                                                                                       |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Digital Certificates for an ES PIC</i></li><li>• <i>Configuring an IKE Policy for Digital Certificates for an ES PIC</i></li></ul>                                                                                  |



## enrollment-retry

---

|                                 |                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>enrollment-retry <i>attempts</i>;</code>                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit security <a href="#">certificates</a> ]                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| <b>Description</b>              | (Encryption interface on M Series and T Series routers and EX Series switches only)<br>Specify how many times a router or switch can resend a digital certificate request.                            |
| <b>Options</b>                  | <i>attempts</i> —Number of enrollment retries.<br><b>Range:</b> 0 through 100<br><b>Default:</b> 0                                                                                                    |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Digital Certificates for an ES PIC</i></li> </ul>                                                                                             |

## enrollment-url

---

|                                 |                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>enrollment-url <i>url-name</i>;</code>                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit security certificates <a href="#">certification-authority</a> <i>ca-profile-name</i> ]                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| <b>Description</b>              | (Encryption interface on M Series and T Series routers and EX Series switches only)<br>Specify where your router or switch sends Simple Certificate Enrollment Protocol-based (SCEP-based) certificate enrollment requests (certificate authority URL).                    |
| <b>Options</b>                  | <i>url-name</i> —Certificate authority URL.                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Digital Certificates for an ES PIC</i></li> </ul>                                                                                                                                                                  |

## file

---

|                                 |                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>file <i>certificate-filename</i>;</code>                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit security certificates <a href="#">certification-authority</a> <i>ca-profile-name</i> ]                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| <b>Description</b>              | (Encryption interface on M Series and T Series routers and EX Series switches only)<br>Specify the file from which to read the digital certificate.                                                                                                                        |
| <b>Options</b>                  | <i>certificate-filename</i> —File from which to read the digital certificate.                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Digital Certificates for an ES PIC</i></li></ul>                                                                                                                                                                    |

## key (Authentication Keychain)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>key key {   algorithm (md5   hmac-sha-1);   options (basic   isis-enhanced);   secret secret-data;   start-time yyyy-mm-dd.hh:mm:ss; }</pre>                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit security authentication-key-chains key-chain <i>key-chain-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Support for IS-IS introduced in Junos OS Release 11.2.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| <b>Description</b>              | Configure the authentication element.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <p><b>key</b>—Each key within a keychain is identified by a unique integer value.</p> <p><b>Range:</b> 0 through 63</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols</a></li> <li>• <a href="#">Example: Configuring BFD Authentication for Securing Static Routes on page 3255</a></li> <li>• <a href="#">Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 4256</a></li> </ul>                                                                                                                                                |

## key-chain (Security)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>keychain <i>key-chain-name</i> {<br/>  description <i>text-string</i>;<br/>  key <i>key</i> {<br/>    algorithm (md5   hmac-sha-1);<br/>    options (basic   isis-enhanced);<br/>    secret <i>secret-data</i>;<br/>    start-time <i>yyyy-mm-dd.hh:mm:ss</i>;<br/>  }<br/>  tolerance <i>seconds</i>;<br/>}</pre>                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit security authentication-key-chains]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Support for IS-IS introduced in Junos OS Release 11.2.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| <b>Description</b>              | Create the key-chain configuration for the Border Gateway Protocol (BGP), the Label Distribution Protocol (LDP) routing protocols, the Bidirectional Forwarding Detection (BFD) protocol, and the Intermediate System-to-Intermediate System (IS-IS) protocol.                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <b><i>key-chain-name</i></b> —Authentication keychain name. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">authentication-key-chains on page 6298</a></li><li>• <a href="#">Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols</a></li><li>• <a href="#">Example: Configuring BFD Authentication for Securing Static Routes on page 3255</a></li><li>• <a href="#">Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 4256</a></li></ul>                                                                                   |


## ldap-url

---

|                                 |                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <ldap-url <i>url-name</i> >;                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit security certificates <a href="#">certification-authority</a> <i>ca-profile-name</i> ]                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series, |
| <b>Description</b>              | (Encryption interface on M Series and T Series routers and EX Series switches only)<br>(Optional) Specify the Lightweight Directory Access Protocol (LDAP) URL for digital certificates.                                                                                   |
| <b>Options</b>                  | <i>url-name</i> —Name of the LDAP URL.                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Digital Certificates for an ES PIC</i></li></ul>                                                                                                                                                                    |

## local

---

|                                                                                                                                                                                                                                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                     | <pre>local <i>certificate-name</i> {<br/>    <i>certificate-key-string</i>;<br/>    load-key-file <i>URL filename</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                            | [edit security <a href="#">certificates</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>                                                                                                                                                                                                                                        | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>                                                                                                                                                                                                                                                | Import a paired X.509 private key and authentication certificate, to enable Junos XML protocol client applications to establish Secure Sockets Layer (SSL) connections to the router or switch.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <div> <b>NOTE:</b> For FIPS mode, the digital security certificates must be compliant with the National Institute of Standards and Technology (NIST) SP 800-131A standard.</div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                                                                                                                                                                                                                                                    | <p><b><i>certificate-key-string</i></b>—String of alphanumeric characters that constitute the private key and certificate.</p> <p><b><i>certificate-name</i></b>—Name that uniquely identifies the certificate.</p> <p><b><i>load-key-file URL filename</i></b>—File that contains the private key and certificate. It can be one of two types of values:</p> <ul style="list-style-type: none"><li>• Pathname of a file on the local disk (assuming you have already used another method to copy the certificate file to the router's or switch's local disk)</li><li>• URL to the certificate file location (for instance, on the computer where the Junos XML protocol client application runs)</li></ul> |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                   | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>                                                                                                                                                                                                                                      | <ul style="list-style-type: none"><li>• <i>Importing SSL Certificates for Junos XML Protocol Support</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## maximum-certificates

|                                 |                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>maximum-certificates <i>number</i>;</code>                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit security <a href="#">certificates</a> ]                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| <b>Description</b>              | (Encryption interface on M Series and T Series routers and EX Series switches only)<br>Configure the maximum number of peer digital certificates to be cached.                                                                                                             |
| <b>Options</b>                  | <b><i>number</i></b> —Maximum number of peer digital certificates to be cached.<br><b>Range:</b> 64 through 4,294,967,295 peer certificates<br><b>Default:</b> 1024 peer certificates                                                                                      |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Digital Certificates for an ES PIC</i></li> </ul>                                                                                                                                                                  |

## path-length

|                                 |                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>path-length <i>certificate-path-length</i>;</code>                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit security <a href="#">certificates</a> ]                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| <b>Description</b>              | (Encryption interface on M Series and T Series routers and EX Series switches only)<br>Configure the digital certificate path length.                                                                                                                                      |
| <b>Options</b>                  | <b><i>certificate-path-length</i></b> —Digital certificate path length.<br><b>Range:</b> 2 through 15 certificates<br><b>Default:</b> 15 certificates                                                                                                                      |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Digital Certificates for an ES PIC</i></li> </ul>                                                                                                                                                                  |

## secret

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>secret <i>secret-data</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit security authentication-key-chains key-chain <i>key-chain-name</i> key <i>key</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Support for IS-IS introduced in JUNOS OS Release 11.2.</p> <p>Statement introduced in Junos OS Release 11.3 for QFX Series switches.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| <b>Description</b>              | Specify a password in encrypted text or plain text format. The secret password always appears in encrypted format.                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <b><i>secret-data</i></b> —Password to use; it can include spaces if the character string is enclosed in quotation marks.                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols</a></li><li>• <a href="#">Example: Configuring BFD Authentication for Securing Static Routes on page 3255</a></li><li>• <a href="#">Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 4256</a></li></ul>                                                                                                                                                         |



## security

```

Syntax security {
 authentication-key-chains {
 key-chain key-chain-name {
 key key {
 secret secret-data;
 start-time yyyy-mm-dd.hh:mm:ss;
 }
 }
 }
 certificates {
 cache-size bytes;
 cache-timeout-negative seconds;
 certification-authority ca-profile-name {
 ca-name ca-identity;
 crl file-name;
 encoding (binary | pem);
 enrollment-url url-name;
 file certificate-filename;
 ldap-url url-name;
 }
 enrollment-retry attempts;
 local certificate-filename {
 certificate-key-string;
 load-key-file key-file-name;
 }
 maximum-certificates number;
 path-length certificate-path-length;
 }
 ssh-known-hosts {
 host {
 fetch-from-server host-name;
 load-key-file file-name;
 }
 }
 traceoptions {
 file filename <files number> <size size>;
 flag flag;
 level level;
 no-remote-trace
 }
 }

```

Hierarchy Level [edit]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Configure security services. Most of the configuration statements do not have default values. If you do not specify an identifier for a statement that does not have a default value, you cannot commit the configuration.

**Required Privilege  
Level**

**Related  
Documentation**

## ssh-known-hosts

---

**Syntax**    ssh-known-hosts {  
              host *host-name* {  
                  fetch-from-server *host-name*;  
                  load-key-file *file-name*;  
              }  
              }

**Hierarchy Level**    [edit security ssh-known-hosts]

**Release Information**    Statement introduced in Junos OS Release 11.1 for the QFX Series.  
                              Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description**    Configure SSH support for known hosts and for administering SSH host key updates.

**Options**    **host *host-name***—Hostname of the SSH known host entry. This option has the following suboptions:

- **fetch-from-server *host-name***—Retrieve SSH public host key information from a specified server.
- **load-key-file *filename***—Import SSH host key information from the `/var/tmp/ssh-known-hosts` file.

**Required Privilege  
Level**    admin—To view this statement in the configuration.  
              admin-control—To add this statement to the configuration.

**Related  
Documentation**    • *Understanding Security Features on the QFabric System*  
                      • *Configuring SSH Host Keys for Secure Copying of Data*


## start-time (Authentication Key Transmission)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>start-time (now   yyyy-mm-dd.hh:mm:ss);</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit security authentication-key-chains key-chain <i>key-chain-name</i> key <i>key</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Support for IS-IS introduced in JUNOS OS Release 11.2.</p> <p>Statement introduced in Junos OS Release 11.3 for QFX Series switches.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | <p>Specify a start time for key transmission. You do not need to specify an end time for the key. If a new key is present with a new start time, the keychain rolls over to the new one. The start time must be unique within the keychain.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <p><b>now</b>—Start time as the current year, month, day, hour, minute, and second.</p> <p><b>daydays</b>—Start time as the specified number of days after the current day. For example, if the current day is the 12th and you configure <b>start-time 2day</b>, the start time will be on the 14th, exactly two days after the configuration is entered.</p> <p><b>hourhours</b>—Start time as the specified number of hours after the current hour. For example, if the current hour is 9:00 and you configure <b>start-time 3hour</b>, the start time will be in 12:00, exactly three hours after the configuration is entered.</p> <p><b>minuteminutes</b>—Start time as the specified number of minutes after the current minute. For example, if the current minute is 27 minutes after the hour and you configure <b>start-time 5min</b>, the start time will be in 32 minutes after the hour, exactly five minutes after the configuration is entered.</p> <p><b>monthmonths</b>—Start time as the specified number of months after the current month. For example, if the current month is March and you configure <b>start-time 4month</b>, the start time will be in July, exactly four months after the configuration is entered.</p> <p><b>secondseconds</b>—Start time as the specified number of seconds after the current second. For example, if the current second is 10:20:40 and you configure <b>start-time 10seconds</b>, the start time will be 10:20:50, exactly 10 seconds after the configuration is entered.</p> <p><b>yearyears</b>—Start time as the specified number of years after the current year. For example, if the current year is 2011 and you configure <b>start-time 1year</b>, the start time will be in 2012, exactly one year after the configuration is entered.</p> <p><b>yyyy-mm-dd.hh:mm:ss</b>—Start time in UTC (Coordinated Universal Time). The start time must be unique within the keychain.</p> |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

**Related  
Documentation**

- [\*Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols\*](#)
- [Example: Configuring BFD Authentication for Securing Static Routes on page 3255](#)
- [Example: Configuring BFD Authentication for Securing Static Routes on page 3255](#)
- [Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 4256](#)

## traceoptions

|                                                                                                                                                                                                                                                                                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                          | <pre> traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt;;     flag all;     flag certificates;     flag database;     flag general;     flag ike;     flag parse;     flag policy-manager;     flag routing-socket;     flag timer;     level     no-remote-trace } </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                 | <p>[edit security],<br/>[edit services ipsec-vpn]</p> <p>Trace options can be configured at either the [edit security] or the [edit services ipsec-vpn] hierarchy level, but not at both levels.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>                                                                                                                                                                                                                                                             | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>                                                                                                                                                                                                                                                                     | <p>Configure security trace options.</p> <p>To specify more than one trace option, include multiple <b>flag</b> statements. Trace option output is recorded in the <code>/var/log/kmd</code> file.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <div style="display: flex; align-items: center;">  <div> <p><b>NOTE:</b> The <code>traceoptions</code> statement is not supported on QFabric systems.</p> </div> </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                                                                                                                                                                                                                                                                         | <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file (for example, <code>kmd</code>) reaches its maximum size, it is renamed <code>kmd.0</code>, then <code>kmd.1</code>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you must also specify a maximum file size with the <b>size</b> option.</p> <p><b>Range:</b> 2 through 1000 files</p> <p><b>Default:</b> 0 files</p> <p><b>size <i>size</i></b>—(Optional) Maximum size of each trace file, in kilobytes (KB). When a trace file (for example, <code>kmd</code>) reaches this size, it is renamed, <code>kmd.0</code>, then <code>kmd.1</code> and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p><b>Default:</b> 1024 KB</p> |

**flag *flag***—Trace operation to perform. To specify more than one trace operation, include multiple **flag** statements.

- **all**—Trace all security events.
- **certificates**—Trace certificate events.
- **database**—Trace database events.
- **general**—Trace general events.
- **ike**—Trace IKE module processing.
- **parse**—Trace configuration processing.
- **policy-manager**—Trace policy manager processing.
- **routing-socket**—Trace routing socket messages.
- **timer**—Trace internal timer events.

**level *level***—(Optional) Set traceoptions level.

- **all**—match all levels.
- **error**—Match error conditions.
- **info**—Match informational messages.
- **notice**—Match conditions that should be handled specially.
- **verbose**—Match verbose messages.
- **warning**—Match warning messages.

**no-remote-trace**—(Optional) Disable remote tracing

|                           |                                                           |
|---------------------------|-----------------------------------------------------------|
| <b>Required Privilege</b> | admin—To view the configuration.                          |
| <b>Level</b>              | admin-control—To add this statement to the configuration. |

|                              |                                                               |
|------------------------------|---------------------------------------------------------------|
| <b>Related Documentation</b> | • <i>Configuring Tracing Operations for Security Services</i> |
|------------------------------|---------------------------------------------------------------|

## CHAPTER 241

# Operational Command (Port Mirroring)

- `show analyzer`

## show analyzer

|                                 |                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show analyzer</b> < <i>analyzer-name</i> >                                                                                                                                  |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                |
| <b>Description</b>              | Display information about port mirroring.                                                                                                                                      |
| <b>Options</b>                  | <i>analyzer-name</i> —(Optional) Displays the status of a specific analyzer (port-mirroring configuration).                                                                    |
| <b>Required Privilege Level</b> | view                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Understanding Layer 2 Port Mirroring</i></li> <li><a href="#">Port Mirroring Constraints and Limitations on page 6230</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show analyzer on page 6320</a>                                                                                                                                     |
| <b>Output Fields</b>            | <a href="#">Table 500</a> describes the output fields for the <b>show analyzer</b> command. Output fields are listed in the approximate order in which they appear.            |

**Table 500: show analyzer Output Fields**

| Field Name                   | Field Description                                                                                                                 |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------|
| Analyzer name                | Name of the analyzer.                                                                                                             |
| Output interface             | Local interface to which mirror packets are sent. If you configure an output interface, you cannot also configure an output VLAN. |
| Output VLAN                  | VLAN to which mirror packets are sent. If you configure an output VLAN, you cannot also configure an output interface.            |
| Egress monitored interfaces  | Interfaces for which egress traffic is mirrored.                                                                                  |
| Egress monitored VLANs       | VLANs for which egress traffic is mirrored.                                                                                       |
| Ingress monitored interfaces | Interfaces for which ingress traffic is mirrored.                                                                                 |
| Ingress monitored VLANs      | VLANs for which ingress traffic is mirrored.                                                                                      |

## Sample Output

### show analyzer

```

user@switch> show analyzer
Analyzer name : employee-monitor
Output interface : ge-0/0/10.0
Output VLAN : remote-analyzer
Egress monitored interfaces : ge-0/0/7.0

```



```
Ingress monitored interfaces : ge-0/0/8.0
Ingress monitored interfaces : ge-0/0/9.0
```



## CHAPTER 242

# Operational Commands (DHCP Local Server)

- clear dhcp server binding
- clear dhcp server statistics
- clear dhcpv6 server binding
- clear dhcpv6 server statistics
- request dhcp server reconfigure
- request dhcpv6 server reconfigure
- request system reboot
- show dhcp server binding
- show dhcp server statistics
- show dhcpv6 server binding
- show dhcpv6 server statistics

## clear dhcp server binding

**Syntax**    `clear dhcp server binding`  
               `<address>`  
               `<all>`  
               `<interface interface-name>`  
               `<interfaces-vlan>`  
               `<interfaces-wildcard>`  
               `<logical-system logical-system-name>`  
               `<routing-instance routing-instance-name>`

**Release Information**    Command introduced in Junos OS Release 9.0.  
                               Options *interfaces-vlan* and *interfaces-wildcard* added in Junos OS Release 12.1.

**Description**    Clear the binding state of a Dynamic Host Configuration Protocol (DHCP) client from the client table on the extended DHCP local server.



**NOTE:** If you delete the DHCP server configuration, DHCP server bindings might still remain. To ensure that DHCP bindings are removed, issue the `clear dhcp server binding` command before you delete the DHCP server configuration.

**Options**    ***address***—(Optional) Clear the binding state for the DHCP client, using one of the following entries:

- *ip-address*—The specified IP address.
- *mac-address*—The specified MAC address.
- *session-id*—The specified session ID.

***all***—(Optional) Clear the binding state for all DHCP clients.

***interface interface-name***—(Optional) Clear the binding state for DHCP clients on the specified interface.



**NOTE:** This option clears all bindings whose initial login requests were received over the specified interface. Dynamic demux login requests are not received over the dynamic demux interface, but rather the underlying interface of the dynamic demux interface. To clear a specific dynamic demux interface, use the *ip-address* or *mac-address* options.

***interfaces-vlan***—(Optional) Clear the binding state on the interface VLAN ID and S-VLAN ID.

***interfaces-wildcard***—(Optional) Clear bindings on a set of interfaces. This option supports the use of the wildcard character (\*).

**logical-system *logical-system-name***—(Optional) Clear the binding state for DHCP clients on the specified logical system.

**routing-instance *routing-instance-name***—(Optional) Clear the binding state for DHCP clients on the specified routing instance.

**Required Privilege Level** view

**Related Documentation**

- [Clearing DHCP Bindings for Subscriber Access](#)
- [show dhcp server binding on page 6341](#)

**List of Sample Output**

- [clear dhcp server binding <ip-address> on page 6325](#)
- [clear dhcp server binding all on page 6325](#)
- [clear dhcp server binding interface on page 6326](#)
- [clear dhcp server binding <interfaces-vlan> on page 6326](#)
- [clear dhcp server binding <interfaces-wildcard> on page 6326](#)

**Output Fields** See [show dhcp server binding](#) for an explanation of output fields.

## Sample Output

### clear dhcp server binding <ip-address>

The following sample output displays the address bindings in the DHCP client table on the extended DHCP local server before and after the **clear dhcp server binding** command is issued.

```
user@host> show dhcp server binding
```

```
2 clients, (0 bound, 0 selecting, 0 renewing, 0 rebinding)
```

| IP address  | Hardware address  | Type   | Lease expires at        |
|-------------|-------------------|--------|-------------------------|
| 100.20.32.1 | 90:00:00:01:00:01 | active | 2007-01-17 11:38:47 PST |
| 100.20.32.3 | 90:00:00:02:00:01 | active | 2007-01-17 11:38:41 PST |

```
user@host> clear dhcp server binding 10.20.32.1
```

```
user@host> show dhcp server binding
```

```
1 clients, (0 bound, 0 selecting, 0 renewing, 0 rebinding)
```

| IP address  | Hardware address  | Type   | Lease expires at        |
|-------------|-------------------|--------|-------------------------|
| 100.20.32.3 | 90:00:00:02:00:01 | active | 2007-01-17 11:38:41 PST |

### clear dhcp server binding all

The following command clears all DHCP local server bindings:

```
user@host> clear dhcp server binding all
```

### clear dhcp server binding interface

The following command clears DHCP local server bindings on a specific interface:

```
user@host> clear dhcp server binding interface fe-0/0/2
```

### clear dhcp server binding <interfaces-vlan>

The following command uses the *interfaces-vlan* option to clear all DHCP local server bindings on top of the underlying interface **ae0**, which clears DHCP bindings on all demux VLANs on top of **ae0**:

```
user@host> clear dhcp server binding ae0
```

### clear dhcp server binding <interfaces-wildcard>

The following command uses the *interfaces-wildcard* option to clear all DHCP local server bindings over a specific interface:

```
user@host> clear dhcp server binding ge-1/0/0.*
```

## clear dhcp server statistics

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>clear dhcp server statistics</code><br><code>&lt;logical-system <i>logical-system-name</i>&gt;</code><br><code>&lt;routing-instance <i>routing-instance-name</i>&gt;</code>                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0.                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Clear all extended Dynamic Host Configuration Protocol (DHCP) local server statistics.                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <p><b>logical-system <i>logical-system-name</i></b>—(Optional) Clear the statistics for DHCP clients on the specified logical system. If you do not specify a logical system, statistics are cleared for the default logical system.</p> <p><b>routing-instance <i>routing-instance-name</i></b>—(Optional) Clear the statistics for DHCP clients on the specified routing instance. If you do not specify a routing instance, statistics are cleared for the default routing instance.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>List of Sample Output</b>    | <a href="#">clear dhcp server statistics on page 6327</a>                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Output Fields</b>            | See <a href="#">show dhcp server statistics</a> for an explanation of output fields.                                                                                                                                                                                                                                                                                                                                                                                                        |

## Sample Output

### clear dhcp server statistics

The following sample output displays the extended DHCP local server statistics before and after the **clear dhcp server statistics** command is issued.

```

user@host> show dhcp server statistics
Packets dropped:
 Total 1
 Lease Time Violation 1

Messages received:
 BOOTREQUEST 89163
 DHCPDECLINE 0
 DHCPDISCOVER 8110
 DHCPINFORM 0
 DHCPRELEASE 0
 DHCPREQUEST 81053

Messages sent:
 BOOTREPLY 32420
 DHCPOFFER 8110
 DHCPACK 8110
 DHCPNAK 8100

user@host> clear dhcp server statistics
user@host> show dhcp server statistics

```

|                    |   |
|--------------------|---|
| Packets dropped:   |   |
| Total              | 0 |
| Messages received: |   |
| BOOTREQUEST        | 0 |
| DHCPDECLINE        | 0 |
| DHCPDISCOVER       | 0 |
| DHCPINFORM         | 0 |
| DHCPRELEASE        | 0 |
| DHCPREQUEST        | 0 |
| Messages sent:     |   |
| BOOTREPLY          | 0 |
| DHCPOFFER          | 0 |
| DHCPACK            | 0 |
| DHCPNAK            | 0 |



## clear dhcpv6 server binding

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>clear dhcpv6 server binding &lt;address&gt; &lt;all&gt; &lt;interface interface-name&gt; &lt;interfaces-vlan&gt; &lt;interfaces-wildcard&gt; &lt;logical-system logical-system-name&gt; &lt;routing-instance routing-instance-name&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 9.6.</p> <p>Options <i>interfaces-vlan</i> and <i>interfaces-wildcard</i> added in Junos OS Release 12.1.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Clear the binding state of a Dynamic Host Configuration Protocol for IPv6 (DHCPv6) client from the client table on the extended DHCPv6 local server.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <p><b>address</b>—(Optional) Clear the binding state for the DHCPv6 client, using one of the following entries:</p> <ul style="list-style-type: none"> <li>• <i>CID</i>—The specified Client ID (CID).</li> <li>• <i>ipv6-prefix</i>—The specified IPv6 prefix.</li> <li>• <i>session-id</i>—The specified session ID.</li> </ul> <p><b>all</b>—(Optional) Clear the binding state for all DHCPv6 clients.</p> <p><b>interface interface-name</b>—(Optional) Clear the binding state for DHCPv6 clients on the specified interface.</p> <p><b>interfaces-vlan</b>—(Optional) Clear the binding state on the interface VLAN ID and S-VLAN ID.</p> <p><b>interfaces-wildcard</b>—(Optional) Clear bindings on a set of interfaces. This option supports the use of the wildcard character (*).</p> <p><b>logical-system logical-system-name</b>—(Optional) Clear the binding state for DHCPv6 clients on the specified logical system.</p> <p><b>routing-instance routing-instance-name</b>—(Optional) Clear the binding state for DHCPv6 clients on the specified routing instance.</p> |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Clearing DHCP Bindings for Subscriber Access</i></li> <li>• <a href="#">show dhcpv6 server binding on page 6350</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>List of Sample Output</b>    | <p><a href="#">clear dhcpv6 server binding all on page 6330</a></p> <p><a href="#">clear dhcpv6 server binding &lt;ipv6-prefix&gt; on page 6330</a></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

[clear dhcpv6 server binding interface on page 6330](#)  
[clear dhcpv6 server binding <interfaces-vlan> on page 6330](#)  
[clear dhcpv6 server binding <interfaces-wildcard> on page 6330](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### clear dhcpv6 server binding all

The following command clears all DHCPv6 local server bindings:

```
user@host> clear dhcpv6 server binding all
```

### clear dhcpv6 server binding <ipv6-prefix>

The following command clears DHCPv6 local server bindings for a specific IPv6 prefix:

```
user@host> clear dhcpv6 server binding 14/0x00010001/0x02b3be8f/0x00109400/0x0005
```

### clear dhcpv6 server binding interface

The following command clears DHCPv6 local server bindings on a specific interface:

```
user@host> clear dhcpv6 server binding interface fe-0/0/2
```

### clear dhcpv6 server binding <interfaces-vlan>

The following command uses the *interfaces-vlan* option to clear all DHCPv6 local server bindings on top of the underlying interface **ae0**, which clears DHCPv6 bindings on all demux VLANs on top of **ae0**:

```
user@host> clear dhcpv6 server binding interface ae0
```

### clear dhcpv6 server binding <interfaces-wildcard>

The following command uses the *interfaces-wildcard* option to clear all DHCPv6 local server bindings over a specific interface:

```
user@host> clear dhcpv6 server binding ge-1/0/0.*
```

## clear dhcpv6 server statistics

---


|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>clear dhcpv6 server statistics &lt;interface <i>interface-name</i>&gt; &lt;logical-system <i>logical-system-name</i>&gt; &lt;routing-instance <i>routing-instance-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.6.                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Clear all extended Dynamic Host Configuration Protocol for IPv6 (DHCPv6) local server statistics.                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <p><b>logical-system <i>logical-system-name</i></b>—(Optional) Clear the statistics for DHCPv6 clients on the specified logical system. If you do not specify a logical system, statistics are cleared for the default logical system.</p> <p><b>routing-instance <i>routing-instance-name</i></b>—(Optional) Clear the statistics for DHCPv6 clients on the specified routing instance. If you do not specify a routing instance, statistics are cleared for the default routing instance.</p> |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show dhcpv6 server statistics on page 6356</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>List of Sample Output</b>    | <a href="#">clear dhcpv6 server statistics on page 6331</a>                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                                                                           |

### Sample Output

#### clear dhcpv6 server statistics

```
user@host> clear dhcpv6 server statistics
```

## request dhcp server reconfigure

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>request dhcp server reconfigure</b> ( <b>all</b>   <i>address</i>   <b>interface</b> <i>interface-name</i>   <b>logical-system</b> <i>logical-system-name</i>   <b>routing-instance</b> <i>routing-instance-name</i> )                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.0.<br>Command introduced in Junos OS Release 12.3R2 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | <p>Initiate reconfiguration processing for the specified DHCP clients if they are in the bound state. If the clients are in the reconfiguring state, this command has no effect. If the clients are in any state other than bound or reconfiguring, this command has the same effect as the <b>clear dhcp server binding</b> command.</p> <p>When the local server state machine starts the reconfiguration process on a bound client, the client transitions to the reconfiguring state and the local server sends a forcerenew message to the client. Because the client was in the bound state before entering the reconfiguring state, all subscriber (or DHCP client) services, such as forwarding and statistics, continue to work. An exponential back-off timer determines the interval at which the forcerenew message is sent. If the final attempt is unsuccessful, the client is returned to its original state by default. You can optionally include the <b>clear-on-abort</b> statement to configure the client to be cleared when reconfiguration fails.</p>                                                                                                                                                    |
| <b>Options</b>                  | <p><b>all</b>—Initiate reconfiguration for all DHCP clients.</p> <p><b>address</b>—Initiate reconfiguration for DHCP client with the specified IP address or MAC address.</p> <p><b>interface <i>interface-name</i></b>—Initiate reconfiguration for all DHCP clients on this logical interface (clients whose initial login requests were received over the specified interface).</p> <div style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;"> <p> <b>NOTE:</b> You cannot use the <b>interface <i>interface-name</i></b> option with the <b>request dhcp server reconfigure</b> command for DHCP passive clients (clients that are added as a result of DHCP snooped packets). For passive clients, the interface is not guaranteed to be the next-hop interface to the client, as is the case for active clients.</p> </div> <p><b>logical-system <i>logical-system-name</i></b>—Initiate reconfiguration for all DHCP clients on the specified logical system.</p> <p><b>routing-instance <i>routing-instance-name</i></b>—Initiate reconfiguration reconfigured for all DHCP clients in the specified routing instance.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

**Related Documentation** • *Configuring Dynamic Client Reconfiguration of Extended Local Server Clients*

**List of Sample Output** [request dhcp server reconfigure on page 6333](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

[request dhcp server reconfigure](#)

```
user@host> request dhcp server reconfigure interface fe-0/0/0.100
```

## request dhcpv6 server reconfigure

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>request dhcpv6 server reconfigure</b> ( <b>all</b>   <b>address</b>   <b>client-id</b>   <b>interface</b> <i>interface-name</i>   <b>logical-system</b> <i>logical-system-name</i>   <b>routing-instance</b> <i>routing-instance-name</i>   <b>session-id</b> )                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.4.<br>Command introduced in Junos OS Release 12.3R2 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | <p>Initiate reconfiguration processing for the specified DHCPv6 clients if they are in the bound state. If the clients are in the reconfiguring state, this command has no effect. If the clients are in any state other than bound or reconfiguring, this command has the same effect as the <b>clear dhcpv6 server binding</b> command.</p> <p>When the local server state machine starts the reconfiguration process on a bound client, the client transitions to the reconfigure state and the local server sends a reconfigure message to the client. Because the client was in the bound state before entering the reconfiguring state, all subscriber (or DHCP client) services, such as forwarding and statistics, continue to work. An exponential back-off timer determines the interval at which the reconfigure message is sent. If the final attempt is unsuccessful, the client is returned to its original state by default. You can optionally include the <b>clear-on-abort</b> statement to configure the client to be cleared when reconfiguration fails.</p> |
| <b>Options</b>                  | <p><b>all</b>—Initiate reconfiguration for all DHCPv6 clients.</p> <p><b>address</b>—Initiate reconfiguration for DHCPv6 client with the specified IPv6 address.</p> <p><b>client-id</b>—Initiate reconfiguration for DHCPv6 client with the specified client ID.</p> <p><b>interface</b> <i>interface-name</i>—Initiate reconfiguration for all DHCPv6 clients on this logical interface (clients whose initial login requests were received over the specified interface).</p> <p><b>logical-system</b> <i>logical-system-name</i>—Initiate reconfiguration for all DHCPv6 clients on the specified logical system.</p> <p><b>routing-instance</b> <i>routing-instance-name</i>—Initiate reconfiguration reconfigured for all DHCPv6 clients in the specified routing instance.</p> <p><b>session-id</b>—Initiate reconfiguration for DHCPv6 client with the specified session ID.</p>                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Dynamic Client Reconfiguration of Extended Local Server Clients</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>List of Sample Output</b>    | <a href="#">request dhcpv6 server reconfigure on page 6335</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## Sample Output

request dhcpv6 server reconfigure

```
user@host> request dhcpv6 server reconfigure 2001::2/16
```

## request system reboot

---

|                                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                 | <a href="#">Syntax on page 6336</a><br><a href="#">Syntax (EX Series Switches) on page 6336</a><br><a href="#">Syntax (TX Matrix Router) on page 6336</a><br><a href="#">Syntax (TX Matrix Plus Router) on page 6336</a><br><a href="#">Syntax (MX Series Router) on page 6336</a>                                                                                                                                                                                           |
| <b>Syntax</b>                         | <code>request system reboot</code><br><code>&lt;at <i>time</i>&gt;</code><br><code>&lt;both-routing-engines&gt;</code><br><code>&lt;in <i>minutes</i>&gt;</code><br><code>&lt;media (compact-flash   disk   removable-compact-flash   usb)&gt;</code><br><code>&lt;message "<i>text</i>"&gt;</code><br><code>&lt;other-routing-engine&gt;</code>                                                                                                                             |
| <b>Syntax (EX Series Switches)</b>    | <code>request system reboot</code><br><code>&lt;all-members&gt;</code><br><code>&lt;at <i>time</i>&gt;</code><br><code>&lt;both-routing-engines&gt;</code><br><code>&lt;in <i>minutes</i>&gt;</code><br><code>&lt;local&gt;</code><br><code>&lt;media (external   internal)&gt;</code><br><code>&lt;member <i>member-id</i>&gt;</code><br><code>&lt;message "<i>text</i>"&gt;</code><br><code>&lt;other-routing-engine&gt;</code><br><code>&lt;slice <i>slice</i>&gt;</code> |
| <b>Syntax (TX Matrix Router)</b>      | <code>request system reboot</code><br><code>&lt;all-chassis   all-lcc   lcc <i>number</i>   scc&gt;</code><br><code>&lt;at <i>time</i>&gt;</code><br><code>&lt;both-routing-engines&gt;</code><br><code>&lt;in <i>minutes</i>&gt;</code><br><code>&lt;media (compact-flash   disk)&gt;</code><br><code>&lt;message "<i>text</i>"&gt;</code><br><code>&lt;other-routing-engine&gt;</code>                                                                                     |
| <b>Syntax (TX Matrix Plus Router)</b> | <code>request system reboot</code><br><code>&lt;all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i>&gt;</code><br><code>&lt;at <i>time</i>&gt;</code><br><code>&lt;both-routing-engines&gt;</code><br><code>&lt;in <i>minutes</i>&gt;</code><br><code>&lt;media (compact-flash   disk)&gt;</code><br><code>&lt;message "<i>text</i>"&gt;</code><br><code>&lt;other-routing-engine&gt;</code><br><code>&lt;partition (1   2   alternate)&gt;</code>                 |
| <b>Syntax (MX Series Router)</b>      | <code>request system reboot</code><br><code>&lt;all-members&gt;</code><br><code>&lt;at <i>time</i>&gt;</code><br><code>&lt;both-routing-engines&gt;</code><br><code>&lt;in <i>minutes</i>&gt;</code><br><code>&lt;local&gt;</code>                                                                                                                                                                                                                                           |



```

<media (external | internal)>
<member member-id>
<message "text">
<other-routing-engine>

```

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Release Information</b> | <p>Command introduced before Junos OS Release 7.4.</p> <p>Option <b>other-routing-engine</b> introduced in Junos OS Release 8.0.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Option <b>sfc</b> introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Option <b>both-routing-engines</b> introduced in Junos OS Release 12.1.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>         | Reboot the software.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>             | <p><b>none</b>—Reboot the software immediately.</p> <p><b>all-chassis</b>—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router or TX Matrix Plus router, reboot all routers connected to the TX Matrix or TX Matrix Plus router, respectively.</p> <p><b>all-lcc</b>—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router or TX Matrix Plus router, reboot all line card chassis connected to the TX Matrix or TX Matrix Plus router, respectively.</p> <p><b>all-members</b>—(EX4200 switches and MX Series routers only) (Optional) Reboot the software on all members of the Virtual Chassis configuration.</p> <p><b>at <i>time</i></b>—(Optional) Time at which to reboot the software, specified in one of the following ways:</p> <ul style="list-style-type: none"> <li>• <b>now</b>—Stop or reboot the software immediately. This is the default.</li> <li>• <b>+<i>minutes</i></b>—Number of minutes from now to reboot the software.</li> <li>• <b><i>yymmddhhmm</i></b>—Absolute time at which to reboot the software, specified as year, month, day, hour, and minute.</li> <li>• <b><i>hh:mm</i></b>—Absolute time on the current day at which to stop the software, specified in 24-hour time.</li> </ul> <p><b>both-routing-engines</b>—(Optional) Reboot both Routing Engines at the same time.</p> <p><b>in <i>minutes</i></b>—(Optional) Number of minutes from now to reboot the software. This option is an alias for the <b>at +<i>minutes</i></b> option.</p> <p><b>lcc <i>number</i></b>—(TX Matrix routers and TX Matrix Plus routers only) (Optional) Line-card chassis number.</p> <p>Replace <i>number</i> with the following values depending on the LCC configuration:</p> <ul style="list-style-type: none"> <li>• 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.</li> <li>• 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.</li> </ul> |

- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(EX4200 switches and MX Series routers only) (Optional) Reboot the software on the local Virtual Chassis member.

**media (compact-flash | disk)**—(Optional) Boot medium for next boot.

**media (external | internal)**—(EX Series switches and MX Series routers only) (Optional) Reboot the boot media:

- **external**—Reboot the external mass storage device.
- **internal**—Reboot the internal flash device.

**member *member-id***—(EX4200 switches and MX Series routers only) (Optional) Reboot the software on the specified member of the Virtual Chassis configuration. For EX4200 switches, replace ***member-id*** with a value from 0 through 9. For an MX Series Virtual Chassis, replace ***member-id*** with a value of 0 or 1.

**message "*text*"**—(Optional) Message to display to all system users before stopping or rebooting the software.

**other-routing-engine**—(Optional) Reboot the other Routing Engine from which the command is issued. For example, if you issue the command from the master Routing Engine, the backup Routing Engine is rebooted. Similarly, if you issue the command from the backup Routing Engine, the master Routing Engine is rebooted.

**partition**—(TX Matrix Plus routers only) (Optional) Reboot using the specified partition on the boot media. This option has the following suboptions:

- **1**—Reboot from partition 1.
- **2**—Reboot from partition 2.
- **alternate**—Reboot from the alternate partition.

**scc**—(TX Matrix routers only) (Optional) Reboot the Routing Engine on the TX Matrix switch-card chassis. If you issue the command from re0, re0 is rebooted. If you issue the command from re1, re1 is rebooted.

**sfc *number***—(TX Matrix Plus routers only) (Optional) Reboot the Routing Engine on the TX Matrix Plus switch-fabric chassis. If you issue the command from re0, re0 is rebooted. If you issue the command from re1, re1 is rebooted. Replace ***number*** with 0.

**slice *slice***—(EX Series switches only) (Optional) Reboot a partition on the boot media. This option has the following suboptions:

- **1**—Power off partition 1.
- **2**—Power off partition 2.

- **alternate**—Reboot from the alternate partition.

**Additional Information** Reboot requests are recorded in the system log files, which you can view with the **show log** command (see [show log](#)). Also, the names of any running processes that are scheduled to be shut down are changed. You can view the process names with the **show system processes** command (see [show system processes](#)).

On a TX Matrix or TX Matrix Plus router, if you issue the **request system reboot** command on the master Routing Engine, all the master Routing Engines connected to the routing matrix are rebooted. If you issue this command on the backup Routing Engine, all the backup Routing Engines connected to the routing matrix are rebooted.



**NOTE:** Before issuing the **request system reboot** command on a TX Matrix Plus router with no options or the **all-chassis**, **all-lcc**, **lcc number**, or **sfc** options, verify that master Routing Engine for all routers in the routing matrix are in the same slot number. If the master Routing Engine for a line-card chassis is in a different slot number than the master Routing Engine for a TX Matrix Plus router, the line-card chassis might become logically disconnected from the routing matrix after the **request system reboot** command.



**NOTE:** To reboot a router that has two Routing Engines, reboot the backup Routing Engine (if you have upgraded it) first, and then reboot the master Routing Engine.

**Required Privilege Level** maintenance

**Related Documentation**

- [clear system reboot on page 337](#)
- [request system halt on page 384](#)
- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

**List of Sample Output**

[request system reboot on page 6340](#)  
[request system reboot \(at 2300\) on page 6340](#)  
[request system reboot \(in 2 Hours\) on page 6340](#)  
[request system reboot \(Immediately\) on page 6340](#)  
[request system reboot \(at 1:20 AM\) on page 6340](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### request system reboot

```
user@host> request system reboot
Reboot the system ? [yes,no] (no)
```

### request system reboot (at 2300)

```
user@host> request system reboot at 2300 message ?Maintenance time!?
Reboot the system ? [yes,no] (no) yes
```

```
shutdown: [pid 186]
*** System shutdown message from root@berry.network.net ***
System going down at 23:00
```

### request system reboot (in 2 Hours)

The following example, which assumes that the time is 5 PM (17:00), illustrates three different ways to request the system to reboot in two hours:

```
user@host> request system reboot at +120
user@host> request system reboot in 120
user@host> request system reboot at 19:00
```

### request system reboot (Immediately)

```
user@host> request system reboot at now
```

### request system reboot (at 1:20 AM)

To reboot the system at 1:20 AM, enter the following command. Because 1:20 AM is the next day, you must specify the absolute time.

```
user@host> request system reboot at 06060120
request system reboot at 120
Reboot the system at 120? [yes,no] (no) yes
```

## show dhcp server binding

**Syntax** `show dhcp server binding`  
`<address>`  
`<interfaces-vlan><brief | detail | summary>`  
`<interface interface-name>`  
`<interfaces-vlan>`  
`<interfaces-wildcard>`  
`<logical-system logical-system-name>`  
`<routing-instance routing-instance-name>`

**Release Information** Command introduced in Junos OS Release 9.0.  
Options *interfaces-vlan* and *interfaces-wildcard* added in Junos OS Release 12.1.

**Description** Display the address bindings in the client table on the extended Dynamic Host Configuration Protocol (DHCP) local server.



**NOTE:** If you delete the DHCP server configuration, DHCP server bindings might still remain. To ensure that DHCP bindings are removed, issue the `clear dhcp server binding` command before you delete the DHCP server configuration.

**Options** *address*—(Optional) Display DHCP binding information for a specific client identified by one of the following entries:

- *ip-address*—The specified IP address.
- *mac-address*—The specified MAC address.
- *session-id*—The specified session ID.

*brief | detail | summary*—(Optional) Display the specified level of output about active client bindings. The default is **brief**, which produces the same output as `show dhcp server binding`.

*interface interface-name*—(Optional) Display information about active client bindings on the specified interface. You can optionally filter on VLAN ID and SVLAN ID.

*interfaces-vlan*—(Optional) Show the binding state information on the interface VLAN ID and S-VLAN ID.

*interfaces-wildcard*—(Optional) The set of interfaces on which to show the binding state information. This option supports the use of the wildcard character (\*).

*logical-system logical-system-name*—(Optional) Display information about active client bindings for DHCP clients on the specified logical system.

*routing-instance routing-instance-name*—(Optional) Display information about active client bindings for DHCP clients on the specified routing instance.

**Required Privilege Level** view

**Related Documentation**

- *Clearing DHCP Bindings for Subscriber Access*
- *Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration*
- [clear dhcp server binding on page 6324](#)

**List of Sample Output**

[show dhcp server binding on page 6344](#)  
[show dhcp server binding detail on page 6345](#)  
[show dhcp server binding detail \(ACI Interface Set Configured\) on page 6345](#)  
[show dhcp server binding interface <vlan-id> on page 6346](#)  
[show dhcp server binding interface <svlan-id> on page 6346](#)  
[show dhcp server binding <ip-address> on page 6346](#)  
[show dhcp server binding <session-id> on page 6346](#)  
[show dhcp server binding summary on page 6346](#)  
[show dhcp server binding <interfaces-vlan> on page 6346](#)  
[show dhcp server binding <interfaces-wildcard> on page 6346](#)

**Output Fields** [Table 501](#) lists the output fields for the **show dhcp server binding** command. Output fields are listed in the approximate order in which they appear.

**Table 501: show dhcp server binding Output Fields**

| Field Name                                                                                                                                                                              | Field Description                                                                                | Level of Output         |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------|-------------------------|
| <i>number</i> clients,<br>( <i>number</i> init,<br><i>number</i> bound,<br><i>number</i> selecting,<br><i>number</i> requesting,<br><i>number</i> renewing,<br><i>number</i> releasing) | Summary counts of the total number of DHCP clients and the number of DHCP clients in each state. | <b>summary</b>          |
| IP address                                                                                                                                                                              | IP address of the DHCP client.                                                                   | <b>brief<br/>detail</b> |
| Session Id                                                                                                                                                                              | Session ID of the subscriber session.                                                            | <b>brief<br/>detail</b> |
| Hardware address                                                                                                                                                                        | Hardware address of the DHCP client.                                                             | <b>brief<br/>detail</b> |
| Expires                                                                                                                                                                                 | Number of seconds in which lease expires.                                                        | <b>brief<br/>detail</b> |

Table 501: show dhcp server binding Output Fields (*continued*)

| Field Name                                    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Level of Output               |
|-----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------|
| <b>State</b>                                  | State of the address binding table on the extended DHCP local server: <ul style="list-style-type: none"> <li>• <b>BOUND</b>—Client has active IP address lease.</li> <li>• <b>FORCERENEW</b>—Client has received forcerenew message from server.</li> <li>• <b>INIT</b>—Initial state.</li> <li>• <b>RELEASE</b>—Client is releasing IP address lease.</li> <li>• <b>RENEWING</b>—Client sending request to renew IP address lease.</li> <li>• <b>REQUESTING</b>—Client requesting a DHCP server.</li> <li>• <b>SELECTING</b>—Client receiving offers from DHCP servers.</li> </ul> | <b>brief</b><br><b>detail</b> |
| <b>Interface</b>                              | Interface on which the request was received.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>brief</b>                  |
| <b>Lease Expires</b>                          | Date and time at which the client's IP address lease expires.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail</b>                 |
| <b>Lease Expires in</b>                       | Number of seconds in which lease expires.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>detail</b>                 |
| <b>Lease Start</b>                            | Date and time at which the client's IP address lease started.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail</b>                 |
| <b>Lease time violated</b>                    | Lease time violation has occurred.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail</b>                 |
| <b>Last Packet Received</b>                   | Date and time at which the router received the last packet.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>detail</b>                 |
| <b>Incoming Client Interface</b>              | Client's incoming interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail</b>                 |
| <b>Client Interface Svlan Id</b>              | S-VLAN ID of the client's incoming interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail</b>                 |
| <b>Client Interface Vlan Id</b>               | VLAN ID of the client's incoming interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>detail</b>                 |
| <b>Demux Interface</b>                        | Name of the IP demultiplexing (demux) interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail</b>                 |
| <b>Server IP Address or Server Identifier</b> | IP address of DHCP server.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>detail</b>                 |
| <b>Server Interface</b>                       | Interface of DHCP server.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>detail</b>                 |
| <b>Client Pool Name</b>                       | Name of address pool used to assign client IP address lease.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail</b>                 |

Table 501: show dhcp server binding Output Fields (*continued*)

| Field Name                   | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Level of Output |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Liveness Detection State     | <p>State of the liveness detection status for a subscriber's Bidirectional Forwarding Detection (BFD) protocol session:</p> <p><b>NOTE:</b> This output field displays status only when liveness detection has been explicitly configured for a subscriber and the liveness detection protocol is actively functioning for that subscriber.</p> <ul style="list-style-type: none"> <li><b>DOWN</b>—Liveness detection has been enabled for a subscriber but the broadband network gateway (BNG) detects that the liveness detection session for the BFD protocol is in the <b>DOWN</b> state.<br/>A liveness detection session that was previously in an <b>UP</b> state has transitioned to a <b>DOWN</b> state, beginning with a liveness detection failure, and ending with the deletion of the client binding. The <b>DOWN</b> state is reported only during this transition period of time.</li> <li><b>UNKNOWN</b>—Liveness detection has been enabled for a subscriber but the actual liveness detection state has not yet been determined.<br/>The <b>UNKNOWN</b> state is reported after a DHCP subscriber initially logs in while the underlying liveness detection protocol handshake, such as BFD, is still processing and the BFD session has not yet reached the <b>UP</b> state.</li> <li><b>UP</b>—Liveness detection has been enabled for a subscriber, and the BNG and the subscriber or client have <i>both</i> determined that the liveness detection session for the BFD protocol is in the <b>UP</b> state.</li> <li><b>WENT_DOWN</b>—State is functionally equivalent to the <b>DOWN</b> state. A liveness detection session that was previously in an <b>UP</b> state has transitioned to a <b>DOWN</b> state implying a liveness detection failure.<br/>The <b>WENT_DOWN</b> state applies to the internal distribution of the liveness detection mechanism between the Junos DHCP Daemon for Subscriber Services (JDHCPd), the BFD plug-in within the Broadband Edge Subscriber Management Daemon (BBE-SMGD), and the Packet Forwarding Engine.</li> </ul> | detail          |
| ACI Interface Set Name       | Internally generated name of the dynamic agent circuit identifier (ACI) interface set.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | detail          |
| ACI Interface Set Index      | Index number of the dynamic ACI interface set.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | detail          |
| ACI Interface Set Session ID | Identifier of the dynamic ACI interface set entry in the session database.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | detail          |

## Sample Output

### show dhcp server binding

```
user@host> show dhcp server binding
```



| IP address   | Session Id | Hardware address  | Expires | State | Interface  |
|--------------|------------|-------------------|---------|-------|------------|
| 100.20.20.15 | 6          | 00:10:94:00:00:01 | 86180   | BOUND | ge-1/0/0.0 |
| 100.20.20.16 | 7          | 00:10:94:00:00:02 | 86180   | BOUND | ge-1/0/0.0 |
| 100.20.20.17 | 8          | 00:10:94:00:00:03 | 86180   | BOUND | ge-1/0/0.0 |
| 100.20.20.18 | 9          | 00:10:94:00:00:04 | 86180   | BOUND | ge-1/0/0.0 |
| 100.20.20.19 | 10         | 00:10:94:00:00:05 | 86180   | BOUND | ge-1/0/0.0 |

### show dhcp server binding detail

```

user@host> show dhcp server binding detail
Client IP Address: 100.20.20.15
 Hardware Address: 00:10:94:00:00:01
 State: BOUND(LOCAL_SERVER_STATE_BOUND_ON_INTF_DELETE)

 Lease Expires: 2009-07-21 10:10:25 PDT
 Lease Expires in: 86151 seconds
 Lease Start: 2009-07-20 10:10:25 PDT
 Incoming Client Interface: ge-1/0/0.0
 Server Ip Address: 100.20.20.9
 Server Interface: none
 Session Id: 6
 Client Pool Name: 6
 Liveness Detection State: UP
Client IP Address: 100.20.20.16
 Hardware Address: 00:10:94:00:00:02
 State: BOUND(LOCAL_SERVER_STATE_BOUND_ON_INTF_DELETE)

 Lease Expires: 2009-07-21 10:10:25 PDT
 Lease Expires in: 86151 seconds
 Lease Start: 2009-07-20 10:10:25 PDT
 Lease time violated: yes
 Incoming Client Interface: ge-1/0/0.0
 Server Ip Address: 100.20.20.9
 Server Interface: none
 Session Id: 7
 Client Pool Name: 7
 Liveness Detection State: UP

```

### show dhcp server binding detail (ACI Interface Set Configured)

```

user@host> show dhcp server binding detail
Client IP Address: 100.20.22.14
 Hardware Address: 00:00:64:34:01:02
 State: BOUND(LOCAL_SERVER_STATE_BOUND)
 Lease Expires: 2012-03-13 09:53:32 PDT
 Lease Expires in: 82660 seconds
 Lease Start: 2012-03-12 10:23:32 PDT
 Last Packet Received: 2012-03-12 10:23:32 PDT
 Incoming Client Interface: demux0.1073741827
 Client Interface Svlan Id: 1802
 Client Interface Vlan Id: 302
 Demux Interface: demux0.1073741832
 Server Identifier: 100.20.200.202
 Session Id: 11
 Client Pool Name: poolA
 Client Profile Name: DEMUXprofile
 Liveness Detection State: UP

```

```

ACI Interface Set Name: aci-1002-demux0.1073741827
ACI Interface Set Index: 2
ACI Interface Set Session ID: 6

```

#### show dhcp server binding interface <vlan-id>

```

user@host> show dhcp server binding interface ge-1/1/0:100
IP address Session Id Hardware address Expires State Interface
200.20.20.15 6 00:10:94:00:00:01 86124 BOUND ge-1/1/0:100

```

#### show dhcp server binding interface <svlan-id>

```

user@host> show dhcp server binding interface ge-1/1/0:10-100
IP address Session Id Hardware address Expires State Interface
200.20.20.16 7 00:10:94:00:00:02 86124 BOUND ge-1/1/0:10-100

```

#### show dhcp server binding <ip-address>

```

user@host> show dhcp server binding 100.20.20.19
IP address Session Id Hardware address Expires State Interface
100.20.20.19 10 00:10:94:00:00:05 86081 BOUND ge-1/0/0.0

```

#### show dhcp server binding <session-id>

```

user@host> show dhcp server binding 6
IP address Session Id Hardware address Expires State Interface
200.20.20.15 6 00:10:94:00:00:01 86124 BOUND ge-1/0/0.0

```

#### show dhcp server binding summary

```

user@host> show dhcp server binding summary
3 clients, (2 init, 1 bound, 0 selecting, 0 requesting, 0 renewing, 0 releasing)

```

#### show dhcp server binding <interfaces-vlan>

```

user@host> show dhcp server binding ge-1/0/0:100-200
IP address Session Id Hardware address Expires State Interface
192.168.0.17 42 00:10:94:00:00:02 86346 BOUND ge-1/0/0.1073741827
192.168.0.16 41 00:10:94:00:00:01 86346 BOUND ge-1/0/0.1073741827

```

#### show dhcp server binding <interfaces-wildcard>

```

user@host> show dhcp server binding ge-1/3/*
IP address Session Id Hardware address Expires State Interface
192.168.0.9 24 00:10:94:00:00:04 86361 BOUND ge-1/3/0.110
192.168.0.8 23 00:10:94:00:00:03 86361 BOUND ge-1/3/0.110
192.168.0.7 22 00:10:94:00:00:02 86361 BOUND ge-1/3/0.110

```

---

## show dhcp server statistics

---

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <b>show dhcp server statistics</b><br><b>&lt;logical-system <i>logical-system-name</i>&gt;</b><br><b>&lt;routing-instance <i>routing-instance-name</i>&gt;</b>                                                                                                                                                                                                                                                                                                                                                                                 |
| Release Information      | Command introduced in Junos OS Release 9.0.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Description              | Display extended Dynamic Host Configuration Protocol (DHCP) local server statistics.                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Options                  | <b>logical-system <i>logical-system-name</i></b> —(Optional) Display information about extended DHCP local server statistics on the specified logical system. If you do not specify a logical system, statistics are displayed for the default logical system.<br><br><b>routing-instance <i>routing-instance-name</i></b> —(Optional) Display information about extended DHCP local server statistics on the specified routing instance. If you do not specify a routing instance, statistics are displayed for the default routing instance. |
| Required Privilege Level | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Related Documentation    | <ul style="list-style-type: none"><li>• <a href="#">clear dhcp server statistics on page 6327</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| List of Sample Output    | <a href="#">show dhcp server statistics on page 6348</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Output Fields            | <a href="#">Table 502</a> lists the output fields for the <b>show dhcp server statistics</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                  |

Table 502: show dhcp server statistics Output Fields

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Packets dropped</b>   | <p>Number of packets discarded by the extended DHCP local server because of errors. Only nonzero statistics appear in the Packets dropped output. When all of the Packets dropped statistics are 0 (zero), only the Total field appears.</p> <ul style="list-style-type: none"> <li>• <b>Total</b>—Total number of packets discarded by the extended DHCP local server</li> <li>• <b>Authentication</b>—Number of packets discarded because they could not be authenticated</li> <li>• <b>Bad hardware address</b>—Number of packets discarded because an invalid hardware address was specified</li> <li>• <b>Bad opcode</b>—Number of packets discarded because an invalid operation code was specified</li> <li>• <b>Bad options</b>—Number of packets discarded because invalid options were specified</li> <li>• <b>Dynamic profile</b>—Number of packets discarded due to dynamic profile information</li> <li>• <b>Invalid server address</b>—Number of packets discarded because an invalid server address was specified</li> <li>• <b>Lease Time Violation</b>—Number of packets discarded because of a lease time violation</li> <li>• <b>No available addresses</b>—Number of packets discarded because there were no addresses available for assignment</li> <li>• <b>No interface match</b>—Number of packets discarded because they did not belong to a configured interface</li> <li>• <b>No routing instance match</b>—Number of packets discarded because they did not belong to a configured routing instance</li> <li>• <b>No valid local address</b>—Number of packets discarded because there was no valid local address</li> <li>• <b>Packet too short</b>—Number of packets discarded because they were too short</li> <li>• <b>Read error</b>—Number of packets discarded because of a system read error</li> <li>• <b>Send error</b>—Number of packets that the extended DHCP local server could not send</li> </ul> |
| <b>Messages received</b> | <p>Number of DHCP messages received.</p> <ul style="list-style-type: none"> <li>• <b>BOOTREQUEST</b>—Number of BOOTP protocol data units (PDUs) received</li> <li>• <b>DHCPDECLINE</b>—Number of DHCP PDUs of type DECLINE received</li> <li>• <b>DHCPDISCOVER</b>—Number of DHCP PDUs of type DISCOVER received</li> <li>• <b>DHCPINFORM</b>—Number of DHCP PDUs of type INFORM received</li> <li>• <b>DHCPRELEASE</b>—Number of DHCP PDUs of type RELEASE received</li> <li>• <b>DHCPREQUEST</b>—Number of DHCP PDUs of type REQUEST received</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Messages sent</b>     | <p>Number of DHCP messages sent.</p> <ul style="list-style-type: none"> <li>• <b>BOOTREPLY</b>—Number of BOOTP PDUs transmitted</li> <li>• <b>DHCPOFFER</b>—Number of DHCP OFFER PDUs transmitted</li> <li>• <b>DHCPACK</b>—Number of DHCP ACK PDUs transmitted</li> <li>• <b>DHCPNACK</b>—Number of DHCP NACK PDUs transmitted</li> <li>• <b>DHCPFORCERENEW</b>—Number of DHCP FORCERENEW PDUs transmitted</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

## Sample Output

### show dhcp server statistics

```

user@host> show dhcp server statistics
Packets dropped:
 Total 1

```

|                      |    |
|----------------------|----|
| Lease Time Violation | 1  |
| Messages received:   |    |
| BOOTREQUEST          | 25 |
| DHCPDECLINE          | 0  |
| DHCPDISCOVER         | 10 |
| DHCPINFORM           | 0  |
| DHCPRELEASE          | 4  |
| DHCPREQUEST          | 10 |
| Messages sent:       |    |
| BOOTREPLY            | 20 |
| DHCPOFFER            | 10 |
| DHCPACK              | 10 |
| DHCPNAK              | 0  |
| DHCPFORCERENEW       | 0  |

## show dhcpv6 server binding

---

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <pre>show dhcpv6 server binding &lt;address&gt; &lt;brief   detail   summary&gt; &lt;interface interface-name&gt; &lt;interfaces-vlan&gt; &lt;interfaces-wildcard&gt; &lt;logical-system logical-system-name&gt; &lt;routing-instance routing-instance-name&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Release Information      | Command introduced in Junos OS Release 9.6.<br>Options <i>interfaces-vlan</i> and <i>interfaces-wildcard</i> added in Junos OS Release 12.1.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Description              | Display the address bindings in the client table on the extended Dynamic Host Configuration Protocol for IPv6 (DHCPv6) local server.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Options                  | <p><b>address</b>—(Optional) One of the following identifiers for the DHCPv6 client whose binding state you want to show:</p> <ul style="list-style-type: none"><li>• <i>CID</i>—The specified Client ID (CID).</li><li>• <i>ipv6-prefix</i>—The specified IPv6 prefix.</li><li>• <i>session-id</i>—The specified session ID.</li></ul> <p><b>brief   detail   summary</b>—(Optional) Display the specified level of output about active client bindings. The default is <b>brief</b>, which produces the same output as <b>show dhcpv6 server binding</b>.</p> <p><b>interface interface-name</b>—(Optional) Display information about active client bindings on the specified interface. You can optionally filter on VLAN ID and SVLAN ID.</p> <p><b>interfaces-vlan</b>—(Optional) Interface VLAN ID or S-VLAN ID interface on which to show binding state information.</p> <p><b>interfaces-wildcard</b>—(Optional) Set of interfaces on which to show binding state information. This option supports the use of the wildcard character (*).</p> <p><b>logical-system logical-system-name</b>—(Optional) Display information about active client bindings for DHCPv6 clients on the specified logical system.</p> <p><b>routing-instance routing-instance-name</b>—(Optional) Display information about active client bindings for DHCPv6 clients on the specified routing instance.</p> |
| Required Privilege Level | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Related Documentation    | <ul style="list-style-type: none"><li>• <i>Clearing DHCP Bindings for Subscriber Access</i></li><li>• <a href="#">clear dhcpv6 server binding on page 6329</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

**List of Sample Output**

- [show dhcpv6 server binding on page 6352](#)
- [show dhcpv6 server binding detail on page 6352](#)
- [show dhcpv6 server binding interface on page 6353](#)
- [show dhcpv6 server binding interface detail on page 6353](#)
- [show dhcpv6 server binding \(IPv6 Prefix\) on page 6354](#)
- [show dhcpv6 server binding \(Session ID\) on page 6354](#)
- [show dhcpv6 server binding \(Interfaces VLAN\) on page 6354](#)
- [show dhcpv6 server binding \(Interfaces Wildcard\) on page 6354](#)
- [show dhcpv6 server binding \(Interfaces Wildcard\) on page 6354](#)
- [show dhcpv6 server binding summary on page 6355](#)

**Output Fields** Table 503 lists the output fields for the **show dhcpv6 server binding** command. Output fields are listed in the approximate order in which they appear.

**Table 503: show dhcpv6 server binding Output Fields**

| Field Name                                                                                                                                                                                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Level of Output     |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------|
| <i>number clients</i> ,<br>( <i>number init</i> ,<br><i>number bound</i> ,<br><i>number selecting</i> ,<br><i>number requesting</i> ,<br><i>number renewing</i> ,<br><i>number releasing</i> ) | Summary counts of the total number of DHCPv6 clients and the number of DHCPv6 clients in each state.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>summary</b>      |
| <b>Prefix</b>                                                                                                                                                                                  | Client's DHCPv6 prefix, or prefix used to support multiple address assignment.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>brief detail</b> |
| <b>Session Id</b>                                                                                                                                                                              | Session ID of the subscriber session.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>brief detail</b> |
| <b>Expires</b>                                                                                                                                                                                 | Number of seconds in which lease expires.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>brief detail</b> |
| <b>State</b>                                                                                                                                                                                   | State of the address binding table on the extended DHCPv6 local server: <ul style="list-style-type: none"> <li>• <b>BOUND</b>—Client has active IP address lease.</li> <li>• <b>INIT</b>—Initial state.</li> <li>• <b>RECONFIGURE</b>—Server has sent reconfigure message to client.</li> <li>• <b>RELEASE</b>—Client is releasing IP address lease.</li> <li>• <b>RENEWING</b>—Client sending request to renew IP address lease.</li> <li>• <b>REQUESTING</b>—Client requesting a DHCPv6 server.</li> <li>• <b>SELECTING</b>—Client receiving offers from DHCPv6 servers.</li> </ul> | <b>brief detail</b> |
| <b>Interface</b>                                                                                                                                                                               | Interface on which the DHCPv6 request was received.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>brief</b>        |
| <b>Client IPv6 Address</b>                                                                                                                                                                     | Client's IPv6 address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail</b>       |
| <b>Client IPv6 Prefix</b>                                                                                                                                                                      | Client's IPv6 prefix.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>detail</b>       |
| <b>Client DUID</b>                                                                                                                                                                             | Client's DHCP Unique Identifier (DUID).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>brief detail</b> |
| <b>Lease expires</b>                                                                                                                                                                           | Date and time at which the client's IP address lease expires.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>detail</b>       |

Table 503: show dhcpv6 server binding Output Fields (*continued*)

| Field Name                 | Field Description                                               | Level of Output |
|----------------------------|-----------------------------------------------------------------|-----------------|
| Lease expires in           | Number of seconds in which lease expires.                       | detail          |
| Preferred Lease Expires    | Date and UTC time at which the client's IPv6 prefix expires.    | detail          |
| Preferred Lease Expires in | Number of seconds at which client's IPv6 prefix expires.        | detail          |
| Lease Start                | Date and time at which the client's address lease was obtained. | detail          |
| Lease time violated        | Lease time violation has occurred.                              | detail          |
| Incoming Client Interface  | Client's incoming interface.                                    | detail          |
| Server IP Address          | IP address of DHCPv6 server.                                    | detail          |
| Server Interface           | Interface of DHCPv6 server.                                     | detail          |
| Client Pool Name           | Address pool used to assign IPv6 address.                       | detail          |
| Client Prefix Pool Name    | Address pool used to assign IPv6 prefix.                        | detail          |
| Client Id length           | Length of the DHCPv6 client ID, in bytes.                       | detail          |
| Client Id                  | ID of the DHCPv6 client.                                        | detail          |

## Sample Output

### show dhcpv6 server binding

```

user@host> show dhcpv6 server binding
Prefix Session Id Expires State Interface Client DUID
2001:bd8:1111:2222::/64 6 86321 BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c0-00:10:94:00:00:01
2001:bd8:1111:2222::/64 7 86321 BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c0-00:10:94:00:00:02
2001:bd8:1111:2222::/64 8 86321 BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c0-00:10:94:00:00:03
2001:bd8:1111:2222::/64 9 86321 BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c1-00:10:94:00:00:04
2001:bd8:1111:2222::/64 10 86321 BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c1-00:10:94:00:00:05
2002::1/74 11 86321 BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c1-00:10:94:00:00:06

```

### show dhcpv6 server binding detail

```

user@host> show dhcpv6 server binding detail

```



```

Session Id: 6
 Client IPv6 Prefix: 2001:bd8:1111:2222::/64
 Client DUID: LL_TIME0x1-0x2e159c0-00:10:94:00:00:01

 State:
 BOUND(LOCAL_SERVER_STATE_BOUND_ON_INTF_DELETE)
 Lease Expires: 2009-07-21 10:41:15 PDT
 Lease Expires in: 86308 seconds
 Preferred Lease Expires: 2012-07-24 00:18:14 UTC
 Preferred Lease Expires in: 600 seconds
 Lease Start: 2009-07-20 10:41:15 PDT
 Lease time violated: yes
 Incoming Client Interface: ge-1/0/0.0
 Server Ip Address: 0.0.0.0
 Server Interface: none
 Client Id Length: 14
 Client Id:
 /0x00010001/0x02e159c0/0x00109400/0x0001

```

```

Session Id: 7
 Client IPv6 Address: 2002::1/128
 Client IPv6 Prefix: 2001:bd8:1111:2222::/64
 Client DUID: LL_TIME0x1-0x2e159c0-00:10:94:00:00:02

 State:
 BOUND(LOCAL_SERVER_STATE_BOUND_ON_INTF_DELETE)
 Lease Expires: 2009-07-21 10:41:15 PDT
 Lease Expires in: 86308 seconds
 Preferred Lease Expires: 2012-07-24 00:18:14 UTC
 Preferred Lease Expires in: 600 seconds
 Lease Start: 2009-07-20 10:41:15 PDT
 Incoming Client Interface: ge-1/0/0.0
 Server Ip Address: 0.0.0.0
 Client Pool Name: bos-v6-pool
 Client Prefix Pool Name: bos-v6-prefix-pool
 Client Id Length: 14
 Client Id:
 /0x00010001/0x02e159c0/0x00109400/0x0002

```

### show dhcpv6 server binding interface

```

user@host> show dhcpv6 server binding interface ge-1/0/0:10-101
Prefix Session Id Expires State Interface Client DUID
2001:bd8:1111:2222::/64 1 86055 BOUND ge-1/0/0.100
LL_TIME0x1-0x4b0a53b9-00:10:94:00:00:01

```

### show dhcpv6 server binding interface detail

```

user@host> show dhcpv6 server binding interface ge-1/0/0:10-101 detail
Session Id: 7
 Client IPv6 Prefix: 2001:bd8:1111:2222::/64
 Client DUID: LL_TIME0x1-0x2e159c0-00:10:94:00:00:02

 State: BOUND(bound)
 Lease Expires: 2009-07-21 10:41:15 PDT
 Lease Expires in: 86136 seconds
 Preferred Lease Expires: 2012-07-24 00:18:14 UTC
 Preferred Lease Expires in: 600 seconds
 Lease Start: 2009-07-20 10:41:15 PDT
 Incoming Client Interface: ge-1/0/0.0
 Server Ip Address: 0.0.0.0

```

```

Server Interface: none
Client Id Length: 14
Client Id:
/0x00010001/0x02e159c0/0x00109400/0x0002

```

### show dhcpv6 server binding (IPv6 Prefix)

```

user@host> show dhcpv6 server binding 14/0x00010001/0x02b3be8f/0x00109400/0x0005
detail
Session Id: 7
Client IPv6 Prefix: 2001:bd8:1111:2222::/64
Client DUID: LL_TIME0x1-0x2e159c0-00:10:94:00:00:02

State: BOUND(bound)
Lease Expires: 2009-07-21 10:41:15 PDT
Lease Expires in: 86136 seconds
Preferred Lease Expires: 2012-07-24 00:18:14 UTC
Preferred Lease Expires in: 600 seconds
Lease Start: 2009-07-20 10:41:15 PDT
Incoming Client Interface: ge-1/0/0.0
Server Ip Address: 0.0.0.0
Server Interface: none
Client Id Length: 14
Client Id:
/0x00010001/0x02e159c0/0x00109400/0x0002

```

### show dhcpv6 server binding (Session ID)

```

user@host> show dhcpv6 server binding 8
Prefix Session Id Expires State Interface Client DUID
2001:DB8::/32 8 86235 BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c0-00:10:94:00:00:03

```

### show dhcpv6 server binding (Interfaces VLAN)

```

user@host> show dhcpv6 server binding ge-1/0/0:100-200
Prefix Session Id Expires State Interface Client DUID
2001:DB8::/32 11 87583 BOUND ge-1/0/0.1073741827
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB9::/32 12 87583 BOUND ge-1/0/0.1073741827
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01

```

### show dhcpv6 server binding (Interfaces Wildcard)

```

user@host> show dhcpv6 server binding demux0
Prefix Session Id Expires State Interface Client DUID
2001:DB8::/32 30 79681 BOUND demux0.1073741824
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB9::/32 31 79681 BOUND demux0.1073741825
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:CB9::/32 32 79681 BOUND demux0.1073741826
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01

```

### show dhcpv6 server binding (Interfaces Wildcard)

```

user@host> show dhcpv6 server binding ge-1/3/*
Prefix Session Id Expires State Interface Client DUID
2001:DB8::/32 22 79681 BOUND ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB9::/32 33 79681 BOUND ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01

```

```
2001:CB9::/32 24 79681 BOUND ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
```

#### show dhcpv6 server binding summary

```
user@host> show dhcpv6 server binding summary
5 clients, (0 init, 5 bound, 0 selecting, 0 requesting, 0 renewing, 0 releasing)
```

## show dhcpv6 server statistics

---

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <b>show dhcpv6 server statistics</b><br><b>&lt;logical-system <i>logical-system-name</i>&gt;</b><br><b>&lt;routing-instance <i>routing-instance-name</i>&gt;</b>                                                                                                                                                                                                                                                                                                                                                                                   |
| Release Information      | Command introduced in Junos OS Release 9.6.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Description              | Display extended Dynamic Host Configuration Protocol for IPv6 (DHCPv6) local server statistics.                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Options                  | <b>logical-system <i>logical-system-name</i></b> —(Optional) Display information about extended DHCPv6 local server statistics on the specified logical system. If you do not specify a logical system, statistics are displayed for the default logical system.<br><br><b>routing-instance <i>routing-instance-name</i></b> —(Optional) Display information about extended DHCPv6 local server statistics on the specified routing instance. If you do not specify a routing instance, statistics are displayed for the default routing instance. |
| Required Privilege Level | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Related Documentation    | <ul style="list-style-type: none"><li>• <a href="#">clear dhcpv6 server statistics on page 6331</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| List of Sample Output    | <a href="#">show dhcpv6 server statistics on page 6357</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Output Fields            | <a href="#">Table 504</a> lists the output fields for the <b>show dhcpv6 server statistics</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                    |

Table 504: show dhcpv6 server statistics Output Fields

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Packets dropped</b>   | <p>Number of packets discarded by the extended DHCPv6 local server because of errors. Only nonzero statistics appear in the Packets dropped output. When all of the Packets dropped statistics are 0 (zero), only the Total field appears.</p> <ul style="list-style-type: none"> <li>• <b>Total</b>—Total number of packets discarded by the extended DHCPv6 local server</li> <li>• <b>Strict Reconfigure</b>—Number of solicit messages discarded because the client does not support reconfiguration</li> <li>• <b>Bad hardware address</b>—Number of packets discarded because an invalid hardware address was specified</li> <li>• <b>Bad opcode</b>—Number of packets discarded because an invalid operation code was specified</li> <li>• <b>Bad options</b>—Number of packets discarded because invalid options were specified</li> <li>• <b>Invalid server address</b>—Number of packets discarded because an invalid server address was specified</li> <li>• <b>Lease Time Violation</b>—Number of packets discarded because of a lease time violation</li> <li>• <b>No available addresses</b>—Number of packets discarded because there were no addresses available for assignment</li> <li>• <b>No interface match</b>—Number of packets discarded because they did not belong to a configured interface</li> <li>• <b>No routing instance match</b>—Number of packets discarded because they did not belong to a configured routing instance</li> <li>• <b>No valid local address</b>—Number of packets discarded because there was no valid local address</li> <li>• <b>Packet too short</b>—Number of packets discarded because they were too short</li> <li>• <b>Read error</b>—Number of packets discarded because of a system read error</li> <li>• <b>Send error</b>—Number of packets that the extended DHCPv6 local server could not send</li> </ul> |
| <b>Messages received</b> | <p>Number of DHCPv6 messages received.</p> <ul style="list-style-type: none"> <li>• <b>DHCPV6_CONFIRM</b>—Number of DHCPv6 CONFIRM PDUs received.</li> <li>• <b>DHCPV6_DECLINE</b>—Number of DHCPv6 DECLINE PDUs received.</li> <li>• <b>DHCPV6_INFORMATION_REQUEST</b>—Number of DHCPv6 INFORMATION-REQUEST PDUs received.</li> <li>• <b>DHCPV6_REBIND</b>—Number of DHCPv6 REBIND PDUs received.</li> <li>• <b>DHCPV6_RELAY_FORW</b>—Number of DHCPv6 RELAY-FORW PDUs received.</li> <li>• <b>DHCPV6_RELAY_REPL</b>—Number of DHCPv6 RELAY-REPL PDUs received.</li> <li>• <b>DHCPV6_RELEASE</b>—Number of DHCPv6 RELEASE PDUs received.</li> <li>• <b>DHCPV6_RENEW</b>—Number of DHCPv6 RENEW PDUs received.</li> <li>• <b>DHCPV6_REQUEST</b>—Number of DHCPv6 REQUEST PDUs received.</li> <li>• <b>DHCPV6_SOLICIT</b>—Number of DHCPv6 SOLICIT PDUs received.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Messages sent</b>     | <p>Number of DHCPv6 messages sent.</p> <ul style="list-style-type: none"> <li>• <b>DHCPV6_ADVERTISE</b>—Number of DHCPv6 ADVERTISE PDUs transmitted.</li> <li>• <b>DHCPV6_REPLY</b>—Number of DHCPv6 ADVERTISE PDUs transmitted.</li> <li>• <b>DHC6_RECONFIGURE</b>—Number of DHCPv6 RECONFIGURE PDUs transmitted.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## Sample Output

### show dhcpv6 server statistics

```
user@host> show dhcpv6 server statistics
```

Dhcpv6 Packets dropped:

|                      |   |
|----------------------|---|
| Total                | 1 |
| Lease Time Violation | 1 |

Messages received:

|                            |   |
|----------------------------|---|
| DHCPV6_DECLINE             | 0 |
| DHCPV6_SOLICIT             | 9 |
| DHCPV6_INFORMATION_REQUEST | 0 |
| DHCPV6_RELEASE             | 0 |
| DHCPV6_REQUEST             | 5 |
| DHCPV6_CONFIRM             | 0 |
| DHCPV6_RENEW               | 0 |
| DHCPV6_REBIND              | 0 |
| DHCPV6_RELAY_FORW          | 0 |
| DHCPV6_RELAY_REPL          | 0 |

Messages sent:

|                    |   |
|--------------------|---|
| DHCPV6_ADVERTISE   | 9 |
| DHCPV6_REPLY       | 5 |
| DHCPV6_RECONFIGURE | 0 |

## CHAPTER 243

# Operational Commands (DHCP Relay Agent)

- clear dhcp relay binding
- clear dhcp relay statistics
- clear dhcpv6 relay binding
- clear dhcpv6 relay statistics
- show dhcp relay binding
- show dhcp relay statistics
- show dhcpv6 relay binding
- show dhcpv6 relay statistics
- show route extensive
- show route protocol

## clear dhcp relay binding

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>clear dhcp relay binding</b><br><b>&lt;address&gt;</b><br><b>&lt;all&gt;</b><br><b>&lt;interface <i>interface-name</i>&gt;</b><br><b>&lt;interfaces-vlan&gt;</b><br><b>&lt;interfaces-wildcard&gt;</b><br><b>&lt;logical-system <i>logical-system-name</i>&gt;</b><br><b>&lt;routing-instance <i>routing-instance-name</i>&gt;</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Command introduced in Junos OS Release 8.3.<br>Options <b>all</b> and <b>interface</b> added in Junos OS Release 8.4.<br>Options <b>interfaces-vlan</b> and <b>interfaces-wildcard</b> added in Junos OS Release 12.1.<br>Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Routers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | Clear the binding state of a Dynamic Host Configuration Protocol (DHCP) client from the client table.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <p><b>address</b>—(Optional) Clear the binding state for the DHCP client, using one of the following entries:</p> <ul style="list-style-type: none"><li>• <b>ip-address</b>—The specified IP address.</li><li>• <b>mac-address</b>—The specified MAC address.</li><li>• <b>session-id</b>—The specified session ID.</li></ul> <p><b>all</b>—(Optional) Clear the binding state for all DHCP clients.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Clear the binding state for DHCP clients on the specified interface.</p> <p><b>interfaces-vlan</b>—(Optional) Clear the binding state on the interface VLAN ID and S-VLAN ID.</p> <p><b>interfaces-wildcard</b>—(Optional) The set of interfaces on which to clear bindings. This option supports the use of the wildcard character (*).</p> <p><b>logical-system <i>logical-system-name</i></b>—(Optional) Clear the binding state for DHCP clients on the specified logical system.</p> <p><b>routing-instance <i>routing-instance-name</i></b>—(Optional) Clear the binding state for DHCP clients on the specified routing instance.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Clearing DHCP Bindings for Subscriber Access</i></li><li>• <a href="#">show dhcp relay binding on page 6370</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |



**List of Sample Output** [clear dhcp relay binding on page 6361](#)  
[clear dhcp relay binding all on page 6361](#)  
[clear dhcp relay binding interface on page 6361](#)  
[clear dhcp relay binding <interfaces-vlan> on page 6361](#)  
[clear dhcp relay binding <interfaces-wildcard> on page 6361](#)

**Output Fields** See [show dhcp relay binding](#) for an explanation of output fields.

## Sample Output

### clear dhcp relay binding

The following sample output displays the address bindings in the DHCP client table before and after the **clear dhcp relay binding** command is issued.

```
user@host> show dhcp relay binding
IP address Hardware address Type Lease expires at
100.20.32.1 90:00:00:01:00:01 active 2007-02-08 16:41:17 EST
192.168.14.8 90:00:01:01:02:01 active 2007-02-10 10:01:06 EST
```

```
user@host> clear dhcp relay binding 100.20.32.1
```

```
user@host> show dhcp relay binding
IP address Hardware address Type Lease expires at
192.168.14.8 90:00:01:01:02:01 active 2007-02-10 10:01:06 EST
```

### clear dhcp relay binding all

The following command clears all DHCP relay agent bindings:

```
user@host> clear dhcp relay binding all
```

### clear dhcp relay binding interface

The following command clears DHCP relay agent bindings on a specific interface:

```
user@host> clear dhcp relay binding interface fe-0/0/3
```

### clear dhcp relay binding <interfaces-vlan>

The following command uses the *interfaces-vlan* option to clear all DHCP relay agent bindings on top of the underlying interface **ae0**, which clears DHCP bindings on all demux VLANs on top of **ae0**:

```
user@host> clear dhcp relay binding interface ae0
```

### clear dhcp relay binding <interfaces-wildcard>

The following command uses the *interfaces-wildcard* option to clear all DHCP relay agent bindings over a specific interface:

```
user@host> clear dhcp relay binding ge-1/0/0.*
```

## clear dhcp relay statistics

---

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <code>clear dhcp relay statistics</code><br><code>&lt;logical-system <i>logical-system-name</i>&gt;</code><br><code>&lt;routing-instance <i>routing-instance-name</i>&gt;</code>                                                                                                                                                                                                                                                                                                                 |
| Syntax                   | Syntax for EX Series switches:<br><br><code>show dhcp relay statistics</code><br><code>&lt;routing-instance <i>routing-instance-name</i>&gt;</code>                                                                                                                                                                                                                                                                                                                                              |
| Release Information      | Command introduced in Junos OS Release 8.3.<br>Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Routers.                                                                                                                                                                                                                                                                                |
| Description              | Clear all Dynamic Host Configuration Protocol (DHCP) relay statistics.                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Options                  | <code>logical-system <i>logical-system-name</i></code> —(On routers only) (Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are cleared for the default logical system.<br><br><code>routing-instance <i>routing-instance-name</i></code> —(Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are cleared for the default routing instance. |
| Required Privilege Level | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Related Documentation    | <ul style="list-style-type: none"><li>• <a href="#">show dhcp relay statistics on page 6375</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                        |
| List of Sample Output    | <a href="#">clear dhcp relay statistics on page 6363</a>                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Output Fields            | <a href="#">Table 505</a> lists the output fields for the <code>clear dhcp relay statistics</code> command.                                                                                                                                                                                                                                                                                                                                                                                      |

Table 505: clear dhcp relay statistics Output Fields

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Packets dropped</b>   | <p>Number of packets discarded by the extended DHCP relay agent application due to errors. Only nonzero statistics appear in the <b>Packets dropped</b> output. When all of the Packets dropped statistics are 0 (zero), only the <b>Total</b> field appears.</p> <ul style="list-style-type: none"> <li>• <b>Total</b>—Total number of packets discarded by the extended DHCP relay agent application.</li> <li>• <b>Bad hardware address</b>—Number of packets discarded because an invalid hardware address was specified.</li> <li>• <b>Bad opcode</b>—Number of packets discarded because an invalid operation code was specified.</li> <li>• <b>Bad options</b>—Number of packets discarded because invalid options were specified.</li> <li>• <b>Invalid server address</b>—Number of packets discarded because an invalid server address was specified.</li> <li>• <b>Lease Time Violation</b>—Number of packets discarded because of a lease time violation</li> <li>• <b>No available addresses</b>—Number of packets discarded because there were no addresses available for assignment.</li> <li>• <b>No interface match</b>—Number of packets discarded because they did not belong to a configured interface.</li> <li>• <b>No routing instance match</b>—Number of packets discarded because they did not belong to a configured routing instance.</li> <li>• <b>No valid local address</b>—Number of packets discarded because there was no valid local address.</li> <li>• <b>Packet too short</b>—Number of packets discarded because they were too short.</li> <li>• <b>Read error</b>—Number of packets discarded because of a system read error.</li> <li>• <b>Send error</b>—Number of packets that the extended DHCP relay application could not send.</li> <li>• <b>Option 60</b>—Number of packets discarded containing DHCP option 60 vendor-specific information.</li> <li>• <b>Option 82</b>—Number of packets discarded because DHCP option 82 information could not be added.</li> </ul> |
| <b>Messages received</b> | <p>Number of DHCP messages received.</p> <ul style="list-style-type: none"> <li>• <b>BOOTREQUEST</b>—Number of BOOTP protocol data units (PDUs) received</li> <li>• <b>DHCPDECLINE</b>—Number of DHCP PDUs of type DECLINE received</li> <li>• <b>DHCPDISCOVER</b>—Number of DHCP PDUs of type DISCOVER received</li> <li>• <b>DHCPINFORM</b>—Number of DHCP PDUs of type INFORM received</li> <li>• <b>DHCPRELEASE</b>—Number of DHCP PDUs of type RELEASE received</li> <li>• <b>DHCPREQUEST</b>—Number of DHCP PDUs of type REQUEST received</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Messages sent</b>     | <p>Number of DHCP messages sent.</p> <ul style="list-style-type: none"> <li>• <b>BOOTREPLY</b>—Number of BOOTP PDUs transmitted</li> <li>• <b>DHCPOFFER</b>—Number of DHCP OFFER PDUs transmitted</li> <li>• <b>DHCPACK</b>—Number of DHCP ACK PDUs transmitted</li> <li>• <b>DHC PNACK</b>—Number of DHCP NACK PDUs transmitted</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## Sample Output

### clear dhcp relay statistics

The following sample output displays the DHCP relay statistics before and after the **clear dhcp relay statistics** command is issued.

```
user@host> show dhcp relay statistics
```

```
Packets dropped:
 Total 1
 Lease Time Violated 1

Messages received:
 BOOTREQUEST 116
 DHCPDECLINE 0
 DHCPDISCOVER 11
 DHCPINFORM 0
 DHCPRELEASE 0
 DHCPREQUEST 105

Messages sent:
 BOOTREPLY 44
 DHCPOFFER 11
 DHCPACK 11
 DHCPNAK 11
```

```
user@host> clear dhcp relay statistics
```

```
user@host> show dhcp relay statistics
```

```
Packets dropped:
 Total 0

Messages received:
 BOOTREQUEST 0
 DHCPDECLINE 0
 DHCPDISCOVER 0
 DHCPINFORM 0
 DHCPRELEASE 0
 DHCPREQUEST 0

Messages sent:
 BOOTREPLY 0
 DHCPOFFER 0
 DHCPACK 0
 DHCPNAK 0
```

## clear dhcpv6 relay binding

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>clear dhcpv6 relay binding &lt;address&gt; &lt;all&gt; &lt;interface interface-name&gt; &lt;interfaces-vlan&gt; &lt;interfaces-wildcard&gt; &lt;logical-system logical-system-name&gt; &lt;routing-instance routing-instance-name&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 11.4.</p> <p>Command introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Options <i>interfaces-vlan</i> and <i>interfaces-wildcard</i> added in Junos OS Release 12.1.</p> <p>Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Routers.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Clear the binding state of Dynamic Host Configuration Protocol for IPv6 (DHCPv6) clients from the client table.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <p><b>address</b>—(Optional) Clear the binding state for the DHCPv6 client, using one of the following entries:</p> <ul style="list-style-type: none"> <li>• <i>CID</i>—The specified Client ID (CID).</li> <li>• <i>ipv6-prefix</i>—The specified IPv6 prefix.</li> <li>• <i>session-id</i>—The specified session ID.</li> </ul> <p><b>all</b>—(Optional) Clear the binding state for all DHCPv6 clients.</p> <p><b>interfaces-vlan</b>—(Optional) Clear the binding state on the interface VLAN ID and S-VLAN ID.</p> <p><b>interfaces-wildcard</b>—(Optional) The set of interfaces on which to clear bindings. This option supports the use of the wildcard character (*).</p> <p><b>interface interface-name</b>—(Optional) Clear the binding state for DHCPv6 clients on the specified interface.</p> <p><b>logical-system logical-system-name</b>—(Optional) Clear the binding state for DHCPv6 clients on the specified logical system.</p> <p><b>routing-instance routing-instance-name</b>—(Optional) Clear the binding state for DHCPv6 clients on the specified routing instance.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Clearing DHCP Bindings for Subscriber Access</i></li> <li>• <a href="#">show dhcpv6 relay binding on page 6378</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

- List of Sample Output**
- [clear dhcpv6 relay binding on page 6366](#)
  - [clear dhcpv6 relay binding <prefix> on page 6366](#)
  - [clear dhcpv6 relay binding all on page 6366](#)
  - [clear dhcv6p relay binding interface on page 6366](#)
  - [clear dhcpv6 relay binding <interfaces-vlan> on page 6367](#)
  - [clear dhcpv6 relay binding <interfaces-wildcard> on page 6367](#)
- Output Fields** See [show dhcpv6 relay binding](#) for an explanation of output fields.

## Sample Output

### clear dhcpv6 relay binding

The following sample output displays the DHCPv6 bindings before and after the **clear dhcpv6 relay binding** command is issued.

```
user@host> show dhcpv6 relay binding
```

| Prefix                                  | Session Id | Expires | State | Interface  | Client DUID |
|-----------------------------------------|------------|---------|-------|------------|-------------|
| 2001:bd8:3c4d:15::/64                   | 1          | 83720   | BOUND | ge-1/0/0.0 |             |
| LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01 |            |         |       |            |             |
| 2001:bd8:3c4d:16::/64                   | 2          | 83720   | BOUND | ge-1/0/0.0 |             |
| LL_TIME0x1-0x4bfa26af-00:10:94:00:00:02 |            |         |       |            |             |
| 2001:bd8:3c4d:17::/64                   | 3          | 83720   | BOUND | ge-1/0/0.0 |             |
| LL_TIME0x1-0x4bfa26af-00:10:94:00:00:03 |            |         |       |            |             |
| 2001:bd8:3c4d:18::/64                   | 4          | 83720   | BOUND | ge-1/0/0.0 |             |
| LL_TIME0x1-0x4bfa26af-00:10:94:00:00:04 |            |         |       |            |             |
| 2001:bd8:3c4d:19::/64                   | 5          | 83720   | BOUND | ge-1/0/0.0 |             |
| LL_TIME0x1-0x4bfa26af-00:10:94:00:00:05 |            |         |       |            |             |
| 2001:bd8:3c4d:20::/64                   | 6          | 83720   | BOUND | ge-1/0/0.0 |             |
| LL_TIME0x1-0x4bfa26af-00:10:94:00:00:06 |            |         |       |            |             |

### clear dhcpv6 relay binding <prefix>

```
user@host> clear dhcpv6 relay binding 2001:bd8:3c4d:15::/64
```

```
user@host> show dhcpv6 relay binding
```

| Prefix                                  | Session Id | Expires | State | Interface  | Client DUID |
|-----------------------------------------|------------|---------|-------|------------|-------------|
| 2001:bd8:3c4d:16::/64                   | 2          | 83720   | BOUND | ge-1/0/0.0 |             |
| LL_TIME0x1-0x4bfa26af-00:10:94:00:00:02 |            |         |       |            |             |
| 2001:bd8:3c4d:17::/64                   | 3          | 83720   | BOUND | ge-1/0/0.0 |             |
| LL_TIME0x1-0x4bfa26af-00:10:94:00:00:03 |            |         |       |            |             |
| 2001:bd8:3c4d:18::/64                   | 4          | 83720   | BOUND | ge-1/0/0.0 |             |
| LL_TIME0x1-0x4bfa26af-00:10:94:00:00:04 |            |         |       |            |             |
| 2001:bd8:3c4d:19::/64                   | 5          | 83720   | BOUND | ge-1/0/0.0 |             |
| LL_TIME0x1-0x4bfa26af-00:10:94:00:00:05 |            |         |       |            |             |
| 2001:bd8:3c4d:20::/64                   | 6          | 83720   | BOUND | ge-1/0/0.0 |             |
| LL_TIME0x1-0x4bfa26af-00:10:94:00:00:06 |            |         |       |            |             |

### clear dhcpv6 relay binding all

The following command clears all DHCP relay agent bindings:

```
user@host> clear dhcpv6 relay binding all
```

### clear dhcv6p relay binding interface

The following command clears DHCPv6 relay agent bindings on a specific interface:

```
user@host> clear dhcpv6 relay binding interface fe-0/0/2
```

#### clear dhcpv6 relay binding <interfaces-vlan>

The following command uses the *interfaces-vlan* option to clear all DHCPv6 relay agent bindings on top of the underlying interface **ae0**, which clears DHCPv6 bindings on all demux VLANs on top of **ae0**:

```
user@host> clear dhcpv6 relay binding interface ae0
```

#### clear dhcpv6 relay binding <interfaces-wildcard>

The following command uses the *interfaces-wildcard* option to clear all DHCPv6 relay agent bindings over a specific interface:

```
user@host> clear dhcpv6 relay binding ge-1/0/0.*
```

## clear dhcpv6 relay statistics

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>clear dhcpv6 relay statistics</b><br><b>&lt;logical-system <i>logical-system-name</i>&gt;</b><br><b>&lt;routing-instance <i>routing-instance-name</i>&gt;</b>                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.4.<br>Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Routers.                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Clear all Dynamic Host Configuration Protocol for IPv6 (DHCPv6) relay statistics.                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <b>logical-system <i>logical-system-name</i></b> —(Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are cleared for the default logical system.<br><br><b>routing-instance <i>routing-instance-name</i></b> —(Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are cleared for the default routing instance. |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>List of Sample Output</b>    | <a href="#">clear dhcpv6 relay statistics on page 6368</a>                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Output Fields</b>            | See <a href="#">show dhcpv6 relay statistics</a> for an explanation of output fields.                                                                                                                                                                                                                                                                                                                                                                              |

## Sample Output

### clear dhcpv6 relay statistics

The following sample output displays the DHCPv6 relay statistics before and after the **clear dhcpv6 relay statistics** command is issued.

```
user@host> show dhcpv6 relay statistics
DHCPv6 Packets dropped:
 Total 0
 Lease Time Violated 1

Messages received:
 DHCPV6_DECLINE 0
 DHCPV6_SOLICIT 10
 DHCPV6_INFORMATION_REQUEST 0
 DHCPV6_RELEASE 0
 DHCPV6_REQUEST 10
 DHCPV6_CONFIRM 0
 DHCPV6_RENEW 0
 DHCPV6_REBIND 0
 DHCPV6_RELAY_REPL 0

Messages sent:
 DHCPV6_ADVERTISE 0
 DHCPV6_REPLY 0
 DHCPV6_RECONFIGURE 0
 DHCPV6_RELAY_FORW 0
```



```
user@host> clear dhcpv6 relay statistics
```

```
user@host> show dhcpv6 relay statistics
```

```
DHCPv6 Packets dropped:
```

```
 Total 0
```

```
Messages received:
```

```
 DHCPV6_DECLINE 0
```

```
 DHCPV6_SOLICIT 0
```

```
 DHCPV6_INFORMATION_REQUEST 0
```

```
 DHCPV6_RELEASE 0
```

```
 DHCPV6_REQUEST 0
```

```
 DHCPV6_CONFIRM 0
```

```
 DHCPV6_RENEW 0
```

```
 DHCPV6_REBIND 0
```

```
 DHCPV6_RELAY_REPL 0
```

```
Messages sent:
```

```
 DHCPV6_ADVERTISE 0
```

```
 DHCPV6_REPLY 0
```

```
 DHCPV6_RECONFIGURE 0
```

```
 DHCPV6_RELAY_FORW 0
```

## show dhcp relay binding

---

**Syntax**    **show dhcp relay binding**  
              <address>  
              <brief>  
              <detail>  
              <interface *interface-name*>  
              <interfaces-vlan>  
              <interfaces-wildcard>  
              <ip-address | mac-address>  
              <logical-system *logical-system-name*>  
              <routing-instance *routing-instance-name*>  
              <summary>

**Release Information**    Command introduced in Junos OS Release 8.3.  
                              Options **interface** and **mac-address** added in Junos OS Release 8.4.  
                              Options **interfaces-vlan** and **interfaces-wildcard** added in Junos OS Release 12.1.  
                              Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Routers.

**Description**    Display the address bindings in the Dynamic Host Configuration Protocol (DHCP) client table.

**Options**    **address**—(Optional) Display DHCP binding information for a specific client identified by one of the following entries:

- *ip-address*—The specified IP address.
- *mac-address*—The specified MAC address.
- *session-id*—The specified session ID.

**brief**—(Optional) Display brief information about the active client bindings. This is the default, and produces the same output as **show dhcp relay binding**.

**detail**—(Optional) Display detailed client binding information.

**interface *interface-name***—(Optional) Perform this operation on the specified interface. You can optionally filter on VLAN ID and SVLAN ID.

**interfaces-vlan**—(Optional) Show the binding state information on the interface VLAN ID and S-VLAN ID.

**interfaces-wildcard**—(Optional) The set of interfaces on which to show binding state information. This option supports the use of the wildcard character (\*).

**logical-system *logical-system-name***—(Optional) Perform this operation on the specified logical system.

**routing-instance *routing-instance-name***—(Optional) Perform this operation on the specified routing instance.

**summary**—(Optional) Display a summary of DHCP client information.

**Required Privilege Level** view

**Related Documentation**

- *Clearing DHCP Bindings for Subscriber Access*
- [clear dhcp relay binding on page 6360](#)

**List of Sample Output**

- [show dhcp relay binding on page 6372](#)
- [show dhcp relay binding detail on page 6373](#)
- [show dhcp relay binding interface on page 6373](#)
- [show dhcp relay binding interface vlan-id on page 6373](#)
- [show dhcp relay binding interface svlan-id on page 6373](#)
- [show dhcp relay binding ip-address on page 6374](#)
- [show dhcp relay binding mac-address on page 6374](#)
- [show dhcp relay binding session-id on page 6374](#)
- [show dhcp relay binding <interfaces-vlan> on page 6374](#)
- [show dhcp relay binding <interfaces-wildcard> on page 6374](#)
- [show dhcp relay binding summary on page 6374](#)

**Output Fields** Table 506 lists the output fields for the **show dhcp relay binding** command. Output fields are listed in the approximate order in which they appear.

**Table 506: show dhcp relay binding Output Fields**

| Field Name                                                                                                                                                                                     | Field Description                                                                                | Level of Output    |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------|--------------------|
| <i>number</i> clients, ( <i>number</i> init, <i>number</i> bound, <i>number</i> selecting, <i>number</i> requesting, <i>number</i> renewing, <i>number</i> rebinding, <i>number</i> releasing) | Summary counts of the total number of DHCP clients and the number of DHCP clients in each state. | <b>summary</b>     |
| IP address                                                                                                                                                                                     | IP address of the DHCP client.                                                                   | <b>briefdetail</b> |
| Session Id                                                                                                                                                                                     | Session ID of the subscriber session.                                                            | <b>briefdetail</b> |
| Generated Remote ID                                                                                                                                                                            | Remote ID generated by the Option 82 Agent Remote ID (suboption 1)                               | <b>detail</b>      |
| Hardware address                                                                                                                                                                               | Hardware address of the DHCP client.                                                             | <b>briefdetail</b> |
| Expires                                                                                                                                                                                        | Number of seconds in which the lease expires.                                                    | <b>briefdetail</b> |

Table 506: show dhcp relay binding Output Fields (*continued*)

| Field Name                       | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Level of Output    |
|----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------|
| <b>State</b>                     | State of the DHCP relay address binding table on the DHCP client: <ul style="list-style-type: none"> <li><b>BOUND</b>—Client has an active IP address lease.</li> <li><b>INIT</b>—Initial state.</li> <li><b>REBINDING</b>—Client is broadcasting a request to renew the IP address lease.</li> <li><b>RELEASE</b>—Client is releasing the IP address lease.</li> <li><b>RENEWING</b>—Client is sending a request to renew the IP address lease.</li> <li><b>REQUESTING</b>—Client is requesting a DHCP server.</li> <li><b>SELECTING</b>—Client is receiving offers from DHCP servers.</li> </ul> | <b>briefdetail</b> |
| <b>Interface</b>                 | Incoming client interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | brief              |
| <b>Lease Expires</b>             | Date and time at which the client's IP address lease expires.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | detail             |
| <b>Lease Expires in</b>          | Number of seconds in which the lease expires.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | detail             |
| <b>Lease Start</b>               | Date and time at which the client's IP address lease started.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | detail             |
| <b>Lease time violated</b>       | Lease time violation has occurred.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | detail             |
| <b>Incoming Client Interface</b> | Client's incoming interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | detail             |
| <b>Server IP Address</b>         | IP address of the DHCP server.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | detail             |
| <b>Server Interface</b>          | Interface of the DHCP server.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | detail             |
| <b>Bootp Relay Address</b>       | IP address of BOOTP relay.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | detail             |
| <b>Type</b>                      | Type of DHCP packet processing performed on the router: <ul style="list-style-type: none"> <li><b>active</b>—Router actively processes and relays DHCP packets.</li> <li><b>passive</b>—Router passively snoops DHCP packets passing through the router.</li> </ul>                                                                                                                                                                                                                                                                                                                                | All levels         |
| <b>Lease expires at</b>          | Date and time at which the client's IP address lease expires.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | All levels         |

## Sample Output

### show dhcp relay binding

```

user@host> show dhcp relay binding
IP address Session Id Hardware address Expires State Interface
100.20.32.11 41 00:10:94:00:00:01 86371 BOUND ge-1/0/0.0
100.20.32.12 42 00:10:94:00:00:02 86371 BOUND ge-1/0/0.0

```

|              |    |                   |       |       |            |
|--------------|----|-------------------|-------|-------|------------|
| 100.20.32.13 | 43 | 00:10:94:00:00:03 | 86371 | BOUND | ge-1/0/0.0 |
| 100.20.32.14 | 44 | 00:10:94:00:00:04 | 86371 | BOUND | ge-1/0/0.0 |
| 100.20.32.15 | 45 | 00:10:94:00:00:05 | 86371 | BOUND | ge-1/0/0.0 |

### show dhcp relay binding detail

```
user@host> show dhcp relay binding detail
```

```
Client IP Address: 100.20.32.11
 Hardware Address: 00:10:94:00:00:01
 State: BOUND(DHCP_RELAY_STATE_BOUND_ON_INTF_DELETE)
 Lease Expires: 2009-07-21 11:00:06 PDT
 Lease Expires in: 86361 seconds
 Lease Start: 2009-07-20 11:00:06 PDT
 Lease time violated: yes
 Last Packet Received: 2009-07-20 11:00:06 PDT
 Incoming Client Interface: ge-1/0/0.0
 Server Ip Address: 100.20.22.2
 Server Interface: none
 Bootp Relay Address: 100.20.32.2
 Session Id: 41
```

```
Client IP Address: 100.20.32.12
 Hardware Address: 00:10:94:00:00:02
 State: BOUND(DHCP_RELAY_STATE_BOUND_ON_INTF_DELETE)
 Lease Expires: 2009-07-21 11:00:06 PDT
 Lease Expires in: 86361 seconds
 Lease Start: 2009-07-20 11:00:06 PDT
 Last Packet Received: 2009-07-20 11:00:06 PDT
 Incoming Client Interface: ge-1/0/0.0
 Server Ip Address: 100.20.22.2
 Server Interface: none
 Bootp Relay Address: 100.20.32.2
 Session Id: 42
 Generated Remote ID: host:ge-1/0/0:100
```

### show dhcp relay binding interface

```
user@host> show dhcp relay binding interface fe-0/0/2
```

| IP address  | Hardware address  | Type   | Lease expires at        |
|-------------|-------------------|--------|-------------------------|
| 100.20.32.1 | 90:00:00:01:00:01 | active | 2007-03-27 15:06:20 EDT |

### show dhcp relay binding interface vlan-id

```
user@host> show dhcp relay binding interface ge-1/1/0:100
```

| IP address   | Session Id | Hardware address  | Expires | State | Interface    |
|--------------|------------|-------------------|---------|-------|--------------|
| 200.20.20.15 | 6          | 00:10:94:00:00:01 | 86124   | BOUND | ge-1/1/0:100 |

### show dhcp relay binding interface svlan-id

```
user@host> show dhcp relay binding interface ge-1/1/0:10-100
```

| IP address | Session Id | Hardware address | Expires | State | Interface |
|------------|------------|------------------|---------|-------|-----------|
|------------|------------|------------------|---------|-------|-----------|

|                 |   |                   |       |       |
|-----------------|---|-------------------|-------|-------|
| 200.20.20.16    | 7 | 00:10:94:00:00:02 | 86124 | BOUND |
| ge-1/1/0:10-100 |   |                   |       |       |

#### show dhcp relay binding ip-address

```
user@host> show dhcp relay binding 100.20.32.13
```

| IP address   | Session Id | Hardware address  | Expires | State | Interface  |
|--------------|------------|-------------------|---------|-------|------------|
| 100.20.32.13 | 43         | 00:10:94:00:00:03 | 86293   | BOUND | ge-1/0/0.0 |

#### show dhcp relay binding mac-address

```
user@host> show dhcp relay binding 00:10:94:00:00:05
```

| IP address   | Session Id | Hardware address  | Expires | State | Interface  |
|--------------|------------|-------------------|---------|-------|------------|
| 100.20.32.15 | 45         | 00:10:94:00:00:05 | 86279   | BOUND | ge-1/0/0.0 |

#### show dhcp relay binding session-id

```
user@host> show dhcp relay binding 41
```

| IP address   | Session Id | Hardware address  | Expires | State | Interface  |
|--------------|------------|-------------------|---------|-------|------------|
| 100.20.32.11 | 41         | 00:10:94:00:00:01 | 86305   | BOUND | ge-1/0/0.0 |

#### show dhcp relay binding <interfaces-vlan>

```
user@host> show dhcp relay binding ge-1/0/0:100-200
```

| IP address   | Session Id | Hardware address  | Expires | State | Interface           |
|--------------|------------|-------------------|---------|-------|---------------------|
| 192.168.0.17 | 42         | 00:10:94:00:00:02 | 86346   | BOUND | ge-1/0/0.1073741827 |
| 192.168.0.16 | 41         | 00:10:94:00:00:01 | 86346   | BOUND |                     |

#### show dhcp relay binding <interfaces-wildcard>

```
user@host> show dhcp relay binding ge-1/3/*
```

| IP address  | Session Id | Hardware address  | Expires | State | Interface    |
|-------------|------------|-------------------|---------|-------|--------------|
| 192.168.0.9 | 24         | 00:10:94:00:00:04 | 86361   | BOUND | ge-1/3/0.110 |
| 192.168.0.8 | 23         | 00:10:94:00:00:03 | 86361   | BOUND |              |
| 192.168.0.7 | 22         | 00:10:94:00:00:02 | 86361   | BOUND | ge-1/3/0.110 |

#### show dhcp relay binding summary

```
user@host> show dhcp relay binding summary
```

3 clients, (2 init, 1 bound, 0 selecting, 0 requesting, 0 renewing, 0 rebinding, 0 releasing)

## show dhcp relay statistics

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show dhcp relay statistics &lt;logical-system <i>logical-system-name</i>&gt; &lt;routing-instance <i>routing-instance-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                 |
| <b>Syntax</b>                   | <p>Syntax for EX Series switches:</p> <pre>show dhcp relay statistics &lt;routing-instance <i>routing-instance-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 8.3.</p> <p>Command introduced in Junos OS Release 12.1 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Routers.</p>                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Display Dynamic Host Configuration Protocol (DHCP) relay statistics.                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <p><b>logical-system <i>logical-system-name</i></b>—(On routers only) (Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are displayed for the default logical system.</p> <p><b>routing-instance <i>routing-instance-name</i></b>—(Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are displayed for the default routing instance.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">clear dhcp relay statistics on page 6362</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>List of Sample Output</b>    | <a href="#">show dhcp relay statistics on page 6377</a>                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Output Fields</b>            | <a href="#">Table 507</a> lists the output fields for the <b>show dhcp relay statistics</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                  |

Table 507: show dhcp relay statistics Output Fields

| Field Name                      | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Packets dropped</b>          | <p>Number of packets discarded by the extended DHCP relay agent application due to errors. Only nonzero statistics appear in the <b>Packets dropped</b> output. When all of the Packets dropped statistics are 0 (zero), only the <b>Total</b> field appears.</p> <ul style="list-style-type: none"> <li>• <b>Total</b>—Total number of packets discarded by the extended DHCP relay agent application.</li> <li>• <b>Bad hardware address</b>—Number of packets discarded because an invalid hardware address was specified.</li> <li>• <b>Bad opcode</b>—Number of packets discarded because an invalid operation code was specified.</li> <li>• <b>Bad options</b>—Number of packets discarded because invalid options were specified.</li> <li>• <b>Invalid server address</b>—Number of packets discarded because an invalid server address was specified.</li> <li>• <b>Lease Time Violation</b>—Number of packets discarded because of a lease time violation</li> <li>• <b>No available addresses</b>—Number of packets discarded because there were no addresses available for assignment.</li> <li>• <b>No interface match</b>—Number of packets discarded because they did not belong to a configured interface.</li> <li>• <b>No routing instance match</b>—Number of packets discarded because they did not belong to a configured routing instance.</li> <li>• <b>No valid local address</b>—Number of packets discarded because there was no valid local address.</li> <li>• <b>Packet too short</b>—Number of packets discarded because they were too short.</li> <li>• <b>Read error</b>—Number of packets discarded because of a system read error.</li> <li>• <b>Send error</b>—Number of packets that the extended DHCP relay application could not send.</li> <li>• <b>Option 60</b>—Number of packets discarded containing DHCP option 60 vendor-specific information.</li> <li>• <b>Option 82</b>—Number of packets discarded because DHCP option 82 information could not be added.</li> </ul> |
| <b>Messages received</b>        | <p>Number of DHCP messages received.</p> <ul style="list-style-type: none"> <li>• <b>BOOTREQUEST</b>—Number of BOOTP protocol data units (PDUs) received</li> <li>• <b>DHCPDECLINE</b>—Number of DHCP PDUs of type DECLINE received</li> <li>• <b>DHCPDISCOVER</b>—Number of DHCP PDUs of type DISCOVER received</li> <li>• <b>DHCPINFORM</b>—Number of DHCP PDUs of type INFORM received</li> <li>• <b>DHCPRELEASE</b>—Number of DHCP PDUs of type RELEASE received</li> <li>• <b>DHCPREQUEST</b>—Number of DHCP PDUs of type REQUEST received</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Messages sent</b>            | <p>Number of DHCP messages sent.</p> <ul style="list-style-type: none"> <li>• <b>BOOTREPLY</b>—Number of BOOTP PDUs transmitted</li> <li>• <b>DHCPOFFER</b>—Number of DHCP OFFER PDUs transmitted</li> <li>• <b>DHCPACK</b>—Number of DHCP ACK PDUs transmitted</li> <li>• <b>DHCPNACK</b>—Number of DHCP NACK PDUs transmitted</li> <li>• <b>DHCPFORCERENEW</b>—Number of DHCP FORCERENEW PDUs transmitted</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>External Server Response</b> | State of the external DHCP server responsiveness.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Packets forwarded</b>        | <p>Number of packets forwarded.</p> <ul style="list-style-type: none"> <li>• <b>BOOTREQUEST</b>—Number of BOOTREQUEST protocol data units (PDUs) forwarded</li> <li>• <b>BOOTREPLY</b>—Number of BOOTREPLY protocol data units (PDUs) forwarded</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |



Table 507: show dhcp relay statistics Output Fields (*continued*)

| Field Name               | Field Description                                 |
|--------------------------|---------------------------------------------------|
| External Server Response | State of the external DHCP server responsiveness. |

## Sample Output

### show dhcp relay statistics

```

user@host> show dhcp relay statistics
Packets dropped:
 Total 34
 Bad hardware address 1
 Bad opcode 1
 Bad options 3
 Invalid server address 5
 Lease Time Violation 1
 No available addresses 1
 No interface match 2
 No routing instance match 9
 No valid local address 4
 Packet too short 2
 Read error 1
 Send error 1
 Option 60 1
 Option 82 2

Messages received:
 BOOTREQUEST 116
 DHCPDECLINE 0
 DHCPDISCOVER 11
 DHCPINFORM 0
 DHCPRELEASE 0
 DHCPREQUEST 105

Messages sent:
 BOOTREPLY 0
 DHCPOFFER 2
 DHCPACK 1
 DHCPNAK 0
 DHCPFORCERENEW 0

Packets forwarded:
 Total 4
 BOOTREQUEST 2
 BOOTREPLY 2

External Server Response:
 State Responding

```

## show dhcpv6 relay binding

---

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <b>show dhcpv6 relay binding</b><br><b>&lt;address&gt;</b><br><b>&lt;brief&gt;</b><br><b>&lt;detail&gt;</b><br><b>&lt;interface <i>interface-name</i>&gt;</b><br><b>&lt;interfaces-vlan&gt;</b><br><b>&lt;interfaces-wildcard&gt;</b><br><b>&lt;logical-system <i>logical-system-name</i>&gt;</b><br><b>&lt;routing-instance <i>routing-instance-name</i>&gt;</b><br><b>&lt;summary&gt;</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Release Information      | Command introduced in Junos OS Release 11.4.<br><i>interfaces-vlan</i> and <i>interfaces-wildcard</i> options introduced in Junos OS Release 12.1.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Description              | Display the DHCPv6 address bindings in the Dynamic Host Configuration Protocol (DHCP) client table.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Options                  | <p><b>address</b>—(Optional) One of the following identifiers for the DHCPv6 client whose binding state you want to show:</p> <ul style="list-style-type: none"><li>• <i>CID</i>—The specified Client ID (CID).</li><li>• <i>ipv6-prefix</i>—The specified IPv6 prefix.</li><li>• <i>session-id</i>—The specified session ID.</li></ul> <p><b>brief</b>—(Optional) Display brief information about the active client bindings. This is the default, and produces the same output as <b>show dhcpv6 relay binding</b>.</p> <p><b>detail</b>—(Optional) Display detailed client binding information.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Perform this operation on the specified interface. You can optionally filter on VLAN ID and S-VLAN ID.</p> <p><b>interfaces-vlan</b>—(Optional) Interface VLAN ID or S-VLAN ID interface on which to show binding state information.</p> <p><b>interfaces-wildcard</b>—(Optional) Set of interfaces on which to show binding state information. This option supports the use of the wildcard character (*).</p> <p><b>logical-system <i>logical-system-name</i></b>—(Optional) Perform this operation on the specified logical system.</p> <p><b>routing-instance <i>routing-instance-name</i></b>—(Optional) Perform this operation on the specified routing instance.</p> <p><b>summary</b>—(Optional) Display a summary of DHCPv6 client information.</p> |
| Required Privilege Level | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

- Related Documentation**
- [Clearing DHCP Bindings for Subscriber Access](#)
  - [clear dhcpv6 relay binding on page 6365](#)

- List of Sample Output**
- [show dhcpv6 relay binding on page 6380](#)
  - [show dhcpv6 relay binding \(Address\) on page 6381](#)
  - [show dhcpv6 relay binding detail \(Client ID\) on page 6381](#)
  - [show dhcpv6 relay binding detail on page 6381](#)
  - [show dhcpv6 relay binding detail \(Multi-Relay Topology\) on page 6382](#)
  - [show dhcpv6 relay binding \(Session ID\) on page 6382](#)
  - [show dhcpv6 relay binding \(Interfaces VLAN\) on page 6382](#)
  - [show dhcpv6 relay binding \(Interfaces Wildcard\) on page 6382](#)
  - [show dhcpv6 relay binding \(Interfaces Wildcard\) on page 6383](#)
  - [show dhcpv6 relay binding summary on page 6383](#)

**Output Fields** Table 508 lists the output fields for the **show dhcpv6 relay binding** command. Output fields are listed in the approximate order in which they appear.

**Table 508: show dhcpv6 relay binding Output Fields**

| Field Name                                                                                                                                   | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Level of Output     |
|----------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------|
| <i>number clients, (number init, number bound, number selecting, number requesting, number renewing, number rebinding, number releasing)</i> | Summary counts of the total number of DHCPv6 clients and the number of DHCPv6 clients in each state.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>summary</b>      |
| <b>Client IPv6 Prefix</b>                                                                                                                    | Prefix of the DHCPv6 client.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>brief detail</b> |
| <b>Client DUID</b>                                                                                                                           | DHCP for IPv6 Unique Identifier (DUID) of the client.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>brief detail</b> |
| <b>Session Id</b>                                                                                                                            | Session ID of the subscriber session.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>brief detail</b> |
| <b>Expires</b>                                                                                                                               | Number of seconds in which the lease expires.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>brief detail</b> |
| <b>State</b>                                                                                                                                 | State of the DHCPv6 relay address binding table on the DHCPv6 client: <ul style="list-style-type: none"> <li>• <b>BOUND</b>—Client has an active IP address lease.</li> <li>• <b>INIT</b>—Initial state.</li> <li>• <b>REBINDING</b>—Client is broadcasting a request to renew the IP address lease.</li> <li>• <b>RELEASE</b>—Client is releasing the IP address lease.</li> <li>• <b>RENEWING</b>—Client is sending a request to renew the IP address lease.</li> <li>• <b>REQUESTING</b>—Client is requesting a DHCPv6 server.</li> <li>• <b>SELECTING</b>—Client is receiving offers from DHCPv6 servers.</li> </ul> | <b>brief detail</b> |
| <b>Interface</b>                                                                                                                             | Incoming client interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>brief</b>        |
| <b>Lease Expires</b>                                                                                                                         | Date and time at which the client's IP address lease expires.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail</b>       |

Table 508: show dhcpv6 relay binding Output Fields (*continued*)

| Field Name                            | Field Description                                                                                                                                                                                                                                                                                                | Level of Output |
|---------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Lease Expires in                      | Number of seconds in which the lease expires.                                                                                                                                                                                                                                                                    | detail          |
| Preferred Lease Expires               | Date and UTC time at which the client's IPv6 prefix expires.                                                                                                                                                                                                                                                     | detail          |
| Preferred Lease Expires in            | Number of seconds at which the client's IPv6 prefix expires.                                                                                                                                                                                                                                                     | detail          |
| Lease Start                           | Date and time at which the client's IP address lease started.                                                                                                                                                                                                                                                    | detail          |
| Lease time violated                   | Lease time violation has occurred.                                                                                                                                                                                                                                                                               | detail          |
| Incoming Client Interface             | Client's incoming interface.                                                                                                                                                                                                                                                                                     | detail          |
| Server Address                        | IP address of the DHCPv6 server.<br><br>Displays <b>unknown</b> for a DHCPv6 relay agent in a multi-relay topology that is not directly adjacent to the DHCPv6 server and does not detect the IP address of the server. In that case, the output instead displays the <b>Next Hop Server Facing Relay</b> field. | detail          |
| Next Hop Server Facing Relay          | Next-hop address in the direction of the DHCPv6 server.                                                                                                                                                                                                                                                          | detail          |
| Server Interface                      | Interface of the DHCPv6 server.                                                                                                                                                                                                                                                                                  | detail          |
| Relay Address                         | IP address of the relay.                                                                                                                                                                                                                                                                                         | detail          |
| Client Pool Name                      | Address pool that granted the client lease.                                                                                                                                                                                                                                                                      | detail          |
| Client ID Length                      | Length of client ID.                                                                                                                                                                                                                                                                                             | All levels      |
| Client Id                             | Client ID.                                                                                                                                                                                                                                                                                                       | All levels      |
| Generated Circuit ID                  | Circuit ID generated by the DHCPv6 Interface-ID option (option 18)                                                                                                                                                                                                                                               | detail          |
| Generated Remote ID Enterprise Number | The Juniper Networks IANA private enterprise number                                                                                                                                                                                                                                                              | detail          |
| Generated Remote ID                   | Remote ID generated by the DHCPv6 Remote-ID option (option 37)                                                                                                                                                                                                                                                   | detail          |

## Sample Output

### show dhcpv6 relay binding

```

user@host> show dhcpv6 relay binding
Prefix Session Id Expires State Interface Client DUID
2001:bd8:3c4d:15::/64 1 83720 BOUND ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01
2001:bd8:3c4d:16::/64 2 83720 BOUND ge-1/0/0.0

```

```

LL_TIME0x1-0x4bfa26af-00:10:94:00:00:02
2001:bd8:3c4d:17::/64 3 83720 BOUND ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:03
2001:bd8:3c4d:18::/64 4 83720 BOUND ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:04
2001:bd8:3c4d:19::/64 5 83720 BOUND ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:05
2001:bd8:3c4d:20::/64 6 83720 BOUND ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:06

```

### show dhcpv6 relay binding (Address)

```

user@host> show dhcpv6 relay binding 2001:bd8:1111:2222::/64 detail
Session Id: 1
 Client IPv6 Prefix: 2001:bd8:3c4d:15::/64
 Client DUID: LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01

 State: BOUND(RELAY_STATE_BOUND)
 Lease Expires: 2011-05-25 07:12:09 PDT
 Lease Expires in: 77115 seconds
 Preferred Lease Expires: 2012-07-24 00:18:14 UTC
 Preferred Lease Expires in: 600 seconds
 Lease Start: 2011-05-24 07:12:09 PDT
 Incoming Client Interface: ge-1/0/0.0
 Server Address: 2008:aaaa:bbbb::1
 Server Interface: none
 Relay Address: 2001:bd8:1111:2222::
 Client Pool Name: pool-25
 Client Id Length: 14
 Client Id:
/0x00010001/0x4bfa26af/0x00109400/0x0001

```

### show dhcpv6 relay binding detail (Client ID)

```

user@host> show dhcpv6 relay binding 14/0x00010001/0x4bfa26af/0x00109400/0x0001
detail
Session Id: 1
 Client IPv6 Prefix: 2001:bd8:3c4d:15::/64
 Client DUID: LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01

 State: BOUND(RELAY_STATE_BOUND)
 Lease Expires: 2011-05-25 07:12:09 PDT
 Lease Expires in: 77115 seconds
 Preferred Lease Expires: 2012-07-24 00:18:14 UTC
 Preferred Lease Expires in: 600 seconds
 Lease Start: 2011-05-24 07:12:09 PDT
 Lease time violated: yes
 Incoming Client Interface: ge-1/0/0.0
 Server Address: 2008:aaaa:bbbb::1
 Server Interface: none
 Relay Address: 2001:bd8:1111:2222::
 Client Pool Name: pool-25
 Client Id Length: 14
 Client Id:
/0x00010001/0x4bfa26af/0x00109400/0x0001

```

### show dhcpv6 relay binding detail

```

user@host> show dhcpv6 relay binding detail
Session Id: 1
 Client IPv6 Prefix: 2001:bd8:3c4d:15::/64

```

```

Client DUID: LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01

State: BOUND(RELAY_STATE_BOUND)
Lease Expires: 2011-05-25 07:12:09 PDT
Lease Expires in: 77115 seconds
Preferred Lease Expires: 2012-07-24 00:18:14 UTC
Preferred Lease Expires in: 600 seconds
Lease Start: 2011-05-24 07:12:09 PDT
Lease time violated: yes
Incoming Client Interface: ge-1/0/0.0
Server Address: 2008:aaaa:bbbb::1
Server Interface: none
Relay Address: 2001:bd8:1111:2222::
Client Pool Name: pool-25
Client Id Length: 14
Client Id: /0x00010001/0x4bfa26af/0x00109400/0x0001
Generated Remote ID Enterprise Number: 1411
Generated Remote ID: host:ge-1/0/0:100

```

#### show dhcpv6 relay binding detail (Multi-Relay Topology)

```

user@host > show dhcpv6 relay binding detail
Session Id: 13
Client IPv6 Prefix: 3000:0:0:8001::5/128
Client DUID: LL0x1-00:00:65:03:01:02
State: BOUND(DHCPV6_RELAY_STATE_BOUND)
Lease Expires: 2011-11-21 06:14:50 PST
Lease Expires in: 293 seconds
Preferred Lease Expires: 2012-07-24 00:18:14 UTC
Preferred Lease Expires in: 600 seconds
Lease Start: 2011-11-21 06:09:50 PST
Incoming Client Interface: ge-1/0/0.0
Server Address: unknown
Next Hop Server Facing Relay: 4000::2
Server Interface: none
Client Id Length: 10
Client Id: /0x00030001/0x00006503/0x0102

```

#### show dhcpv6 relay binding (Session ID)

```

user@host> show dhcpv6 relay binding 41
Prefix Session Id Expires State Interface Client DUID
2001:bd8:3c4d:15::/64 41 78837 BOUND ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01

```

#### show dhcpv6 relay binding (Interfaces VLAN)

```

user@host> show dhcpv6 relay binding ge-1/0/0:100-200
Prefix Session Id Expires State Interface Client DUID
2001:DB8::/32 11 87583 BOUND ge-1/0/0.1073741827
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB9::/32 12 87583 BOUND ge-1/0/0.1073741827
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01

```

#### show dhcpv6 relay binding (Interfaces Wildcard)

```

user@host> show dhcpv6 relay binding demux0
Prefix Session Id Expires State Interface Client DUID
2001:DB8::/32 30 79681 BOUND demux0.1073741824
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB9::/32 31 79681 BOUND demux0.1073741825

```

```

LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:CB9::/32 32 79681 BOUND demux0.1073741826
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01

```

#### show dhcpv6 relay binding (Interfaces Wildcard)

```

user@host> show dhcpv6 relay binding ge-1/3/*
Prefix Session Id Expires State Interface Client DUID
2001:DB8::/32 22 79681 BOUND ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB9::/32 33 79681 BOUND ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:CB9::/32 24 79681 BOUND ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01

```

#### show dhcpv6 relay binding summary

```

user@host> show dhcpv6 relay binding summary
5 clients, (0 init, 5 bound, 0 selecting, 0 requesting, 0 renewing, 0 releasing)

```

## show dhcpv6 relay statistics

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show dhcpv6 relay statistics</b><br><code>&lt;logical-system <i>logical-system-name</i>&gt;</code><br><code>&lt;routing-instance <i>routing-instance-name</i>&gt;</code>                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.4.<br>Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Switches.<br>Command introduced in Junos OS Release 12.3R2 for EX Series switches.                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Display Dynamic Host Configuration Protocol for IPv6 (DHCPv6) relay statistics.                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <p><b>logical-system <i>logical-system-name</i></b>—(Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are displayed for the default logical system.</p> <p><b>routing-instance <i>routing-instance-name</i></b>—(Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are displayed for the default routing instance.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">clear dhcpv6 relay statistics on page 6368</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                |
| <b>List of Sample Output</b>    | <a href="#">show dhcpv6 relay statistics on page 6385</a>                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Output Fields</b>            | <a href="#">Table 509</a> lists the output fields for the <b>show dhcpv6 relay statistics</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                              |

**Table 509: show dhcpv6 relay statistics Output Fields**

| Field Name             | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DHCPv6 Packets dropped | <p>Number of packets discarded by the extended DHCPv6 relay agent application due to errors. Only nonzero statistics appear in the <b>Packets dropped</b> output. When all of the Packets dropped statistics are 0 (zero), only the <b>Total</b> field appears.</p> <ul style="list-style-type: none"> <li><b>Total</b>—Total number of packets discarded by the DHCPv6 relay agent application.</li> <li><b>Bad options</b>—Number of packets discarded because invalid options were specified.</li> <li><b>Bad send</b>—Number of packets that the extended DHCP relay application could not send.</li> <li><b>Bad src address</b>—Number of packets discarded because the family type was not AF_INET6.</li> <li><b>No client id</b>—Number of packets discarded because they could not be matched to a client.</li> <li><b>Lease Time Violation</b>—Number of packets discarded because of a lease time violation</li> <li><b>No safd</b>—Number of packets discarded because they arrived on an unconfigured interface.</li> <li><b>Short packet</b>—Number of packets discarded because they were too short.</li> <li><b>Relay hop count</b>—Number of packets discarded because the hop count in the packet exceeded 32.</li> </ul> |



Table 509: show dhcpv6 relay statistics Output Fields (*continued*)

| Field Name                      | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Messages received</b>        | <p>Number of DHCPv6 messages received.</p> <ul style="list-style-type: none"> <li><b>DHCPV6_DECLINE</b>—Number of DHCPv6 PDUs of type DECLINE received</li> <li><b>DHCPV6_SOLICIT</b>—Number of DHCPv6 PDUs of type SOLICIT received</li> <li><b>DHCPV6_INFORMATION_REQUEST</b>—Number of DHCPv6 PDUs of type INFORMATION-REQUEST received</li> <li><b>DHCPV6_RELEASE</b>—Number of DHCPv6 PDUs of type RELEASE received</li> <li><b>DHCPV6_REQUEST</b>—Number of DHCPv6 PDUs of type REQUEST received</li> <li><b>DHCPV6_CONFIRM</b>—Number of DHCPv6 PDUs of type CONFIRM received</li> <li><b>DHCPV6_RENEW</b>—Number of DHCPv6 PDUs of type RENEW received</li> <li><b>DHCPV6_REBIND</b>—Number of DHCPv6 PDUs of type REBIND received</li> <li><b>DHCPV6_RELAY_REPL</b>—Number of DHCPv6 PDUs of type RELAY-REPL received</li> </ul> |
| <b>Messages sent</b>            | <p>Number of DHCPv6 messages sent.</p> <ul style="list-style-type: none"> <li><b>DHCPV6_ADVERTISE</b>—Number of DHCPv6 ADVERTISE PDUs transmitted</li> <li><b>DHCP_REPLY</b>—Number of DHCPv6 REPLY PDUs transmitted</li> <li><b>DHCP_RECONFIGURE</b>—Number of DHCPv6 RECONFIGURE PDUs transmitted</li> <li><b>DHCP_RELAY_FORW</b>—Number of DHCPv6 RELAY-FORW PDUs transmitted</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Packets forwarded</b>        | <p>Number of packets forwarded by the extended DHCPv6 relay agent application.</p> <ul style="list-style-type: none"> <li><b>FWD REQUEST</b>—Number of DHCPv6 REQUEST packets forwarded</li> <li><b>FWD REPLY</b>—Number of DHCPv6 REPLY packets forwarded</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>External Server Response</b> | State of the external DHCP server responsiveness.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

## Sample Output

### show dhcpv6 relay statistics

```

user@host> show dhcpv6 relay statistics
DHCPv6 Packets dropped:
 Total 1
 Lease Time Violation 1

Messages received:
 DHCPV6_DECLINE 0
 DHCPV6_SOLICIT 10
 DHCPV6_INFORMATION_REQUEST 0
 DHCPV6_RELEASE 0
 DHCPV6_REQUEST 10
 DHCPV6_CONFIRM 0
 DHCPV6_RENEW 0
 DHCPV6_REBIND 0
 DHCPV6_RELAY_REPL 0

Messages sent:
 DHCPV6_ADVERTISE 0
 DHCPV6_REPLY 0

```

|                           |            |
|---------------------------|------------|
| DHCPV6_RECONFIGURE        | 0          |
| DHCPV6_RELAY_FORW         | 0          |
| Packets forwarded:        |            |
| Total                     | 4          |
| FWD REQUEST               | 2          |
| FWD REPLY                 | 2          |
| External Server Response: |            |
| State                     | Responding |

## show route extensive

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 6387</a><br><a href="#">Syntax (EX Series Switches) on page 6387</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Syntax</b>                      | show route extensive<br><destination-prefix><br><logical-system (all   logical-system-name)>                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Syntax (EX Series Switches)</b> | show route extensive<br><destination-prefix>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>                 | Display extensive information about the active entries in the routing tables.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                     | <b>none</b> —Display all active entries in the routing table.<br><br><b>destination-prefix</b> —(Optional) Display active entries for the specified address or range of addresses.<br><br><b>logical-system (all   logical-system-name)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>List of Sample Output</b>       | <a href="#">show route extensive on page 6394</a><br><a href="#">show route extensive (Access Route) on page 6400</a><br><a href="#">show route extensive (BGP PIC Edge) on page 6401</a><br><a href="#">show route extensive (FRR and LFA) on page 6401</a><br><a href="#">show route extensive (Route Reflector) on page 6402</a><br><a href="#">show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 6402</a><br><a href="#">show route label detail (Multipoint LDP with Multicast-Only Fast Reroute) on page 6403</a> |
| <b>Output Fields</b>               | <a href="#">Table 304</a> describes the output fields for the <b>show route extensive</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                           |

Table 510: show route extensive Output Fields

| Field Name                 | Field Description                                                       |
|----------------------------|-------------------------------------------------------------------------|
| <i>routing-table-name</i>  | Name of the routing table (for example, inet.0).                        |
| <i>number destinations</i> | Number of destinations for which there are routes in the routing table. |

Table 510: show route extensive Output Fields (*continued*)

| Field Name                                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>number routes</i>                           | <p>Number of routes in the routing table and total number of routes in the following states:</p> <ul style="list-style-type: none"> <li>• <b>active</b> (routes that are active).</li> <li>• <b>holddown</b> (routes that are in the pending state before being declared inactive).</li> <li>• <b>hidden</b> (routes that are not used because of a routing policy).</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <i>route-destination</i><br>(entry, announced) | <p>Route destination (for example: 10.0.0.1/24). The <b>entry</b> value is the number of route for this destination, and the <b>announced</b> value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:</p> <ul style="list-style-type: none"> <li>• <b>MPLS-label</b> (for example, 80001).</li> <li>• <b>interface-name</b> (for example, ge-1/0/2).</li> <li>• <b>neighbor-address:control-word-status:encapsulation type:vc-id:source</b> (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96). <ul style="list-style-type: none"> <li>• <b>neighbor-address</b>—Address of the neighbor.</li> <li>• <b>control-word-status</b>—Whether the use of the control word has been negotiated for this virtual circuit: <b>NoCtrlWord</b> or <b>CtrlWord</b>.</li> <li>• <b>encapsulation type</b>—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport.</li> <li>• <b>vc-id</b>—Virtual circuit identifier.</li> <li>• <b>source</b>—Source of the advertisement: <b>Local</b> or <b>Remote</b>.</li> </ul> </li> </ul> |
| <b>TSI</b>                                     | Protocol header information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>label stacking</b>                          | <p>(Next-to-the-last-hop routing device for MPLS only) Depth of the Multiprotocol Label Switching (MPLS) label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels.</p> <ul style="list-style-type: none"> <li>• <b>S=0 route</b> indicates that a packet with an incoming label stack depth of two or more exits this router with one fewer label (the label-popping operation is performed).</li> <li>• If there is no <b>S=</b> information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>[protocol, preference]</b>                  | <p>Protocol from which the route was learned and the preference value for the route.</p> <ul style="list-style-type: none"> <li>• <b>+—</b>A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table.</li> <li>• <b>- —</b>A hyphen indicates the last active route.</li> <li>• <b>*—</b>An asterisk indicates that the route is both the active and the last active route. An asterisk before a <b>to</b> line indicates the best subpath to the route.</li> </ul> <p>In every routing metric except for the BGP <b>LocalPref</b> attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the <b>LocalPref</b> value in the <b>Preference2</b> field. For example, if the <b>LocalPref</b> value for Route 1 is 100, the <b>Preference2</b> value is -101. If the <b>LocalPref</b> value for Route 2 is 155, the <b>Preference2</b> value is -156. Route 2 is preferred because it has a higher <b>LocalPref</b> value and a lower <b>Preference2</b> value.</p>                                                                                                                                                                                                                           |

Table 510: show route extensive Output Fields (*continued*)

| Field Name                                           | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Level</b>                                         | (IS-IS only). In IS-IS, a single autonomous system (AS) can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Route Distinguisher</b>                           | IP subnet augmented with a 64-bit prefix.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>PMSI</b>                                          | Provider multicast service interface (MVPN routing table).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Next-hop type</b>                                 | Type of next hop. For a description of possible values for this field, see the Output Field table in the <a href="#">show route detail</a> command.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Next-hop reference count</b>                      | Number of references made to the next hop.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Flood nexthop branches exceed maximum message</b> | Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Source</b>                                        | IP address of the route source.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Next hop</b>                                      | Network layer address of the directly reachable neighboring system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>via</b>                                           | <p>Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word <b>Selected</b>. This field can also contain the following information:</p> <ul style="list-style-type: none"> <li>• <b>Weight</b>—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when Multiprotocol Label Switching (MPLS) label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.</li> <li>• <b>Balance</b>—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable Border Gateway Protocol (BGP) multipath load balancing.</li> </ul> |
| <b>Label-switched-path lsp-path-name</b>             | Name of the label-switched path (LSP) used to reach the next hop.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Label operation</b>                               | MPLS label and operation occurring at this routing device. The operation can be <b>pop</b> (where a label is removed from the top of the stack), <b>push</b> (where another label is added to the label stack), or <b>swap</b> (where a label is replaced by another label).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Offset</b>                                        | Whether the metric has been increased or decreased by an offset value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Interface</b>                                     | (Local only) Local interface name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Protocol next hop</b>                             | Network layer address of the remote routing device that advertised the prefix. This address is used to recursively derive a forwarding next hop.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

Table 510: show route extensive Output Fields (*continued*)

| Field Name                    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b><i>label-operation</i></b> | MPLS label and operation occurring at this routing device. The operation can be <b>pop</b> (where a label is removed from the top of the stack), <b>push</b> (where another label is added to the label stack), or <b>swap</b> (where a label is replaced by another label).                                                                                                                                                                                                                                                 |
| <b>Indirect next hops</b>     | <p>When present, a list of nodes that are used to resolve the path to the next-hop destination, in the order that they are resolved.</p> <p>When BGP PIC Edge is enabled, the output lines that contain <b>Indirect next hop: weight</b> follow next hops that the software can use to repair paths where a link failure occurs. The next-hop weight has one of the following values:</p> <ul style="list-style-type: none"> <li>• 0x1 indicates active next hops.</li> <li>• 0x4000 indicates passive next hops.</li> </ul> |
| <b>State</b>                  | State of the route (a route can be in more than one state). See the Output Field table in the <a href="#">show route detail</a> command.                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Session ID</b>             | The BFD session ID number that represents the protection using MPLS fast reroute (FRR) and loop-free alternate (LFA).                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Weight</b>                 | <p>Weight for the backup path. If the weight of an indirect next hop is larger than zero, the weight value is shown.</p> <p>For sample output, see <a href="#">show route table</a>.</p>                                                                                                                                                                                                                                                                                                                                     |

Table 510: show route extensive Output Fields (*continued*)

| Field Name      | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Inactive reason | <p>If the route is inactive, the reason for its current state is indicated. Typical reasons include:</p> <ul style="list-style-type: none"> <li>• <b>Active preferred</b>—Currently active route was selected over this route.</li> <li>• <b>Always compare MED</b>—Path with a lower multiple exit discriminator (MED) is available.</li> <li>• <b>AS path</b>—Shorter AS path is available.</li> <li>• <b>Cisco Non-deterministic MED selection</b>—Cisco nondeterministic MED is enabled and a path with a lower MED is available.</li> <li>• <b>Cluster list length</b>—Path with a shorter cluster list length is available.</li> <li>• <b>Forwarding use only</b>—Path is only available for forwarding purposes.</li> <li>• <b>IGP metric</b>—Path through the next hop with a lower IGP metric is available.</li> <li>• <b>IGP metric type</b>—Path with a lower OSPF link-state advertisement type is available.</li> <li>• <b>Interior &gt; Exterior &gt; Exterior via Interior</b>—Direct, static, IGP, or EBGp path is available.</li> <li>• <b>Local preference</b>—Path with a higher local preference value is available.</li> <li>• <b>Next hop address</b>—Path with a lower metric next hop is available.</li> <li>• <b>No difference</b>—Path from a neighbor with a lower IP address is available.</li> <li>• <b>Not Best in its group</b>—Occurs when multiple peers of the same external AS advertise the same prefix and are grouped together in the selection process. When this reason is displayed, an additional reason is provided (typically one of the other reasons listed).</li> <li>• <b>Number of gateways</b>—Path with a higher number of next hops is available.</li> <li>• <b>Origin</b>—Path with a lower origin code is available.</li> <li>• <b>OSPF version</b>—Path does not support the indicated OSPF version.</li> <li>• <b>RIB preference</b>—Route from a higher-numbered routing table is available.</li> <li>• <b>Route distinguisher</b>—64-bit prefix added to IP subnets to make them unique.</li> <li>• <b>Route metric or MED comparison</b>—Route with a lower metric or MED is available.</li> <li>• <b>Route preference</b>—Route with a lower preference value is available.</li> <li>• <b>Router ID</b>—Path through a neighbor with a lower ID is available.</li> <li>• <b>Unusable path</b>—Path is not usable because of one of the following conditions: the route is damped, the route is rejected by an import policy, or the route is unresolved.</li> <li>• <b>Update source</b>—Last tiebreaker is the lowest IP address value.</li> </ul> |
| Local AS        | Autonomous system (AS) number of the local routing device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Age             | How long the route has been known.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| AIGP            | Accumulated interior gateway protocol (AIGP) BGP attribute.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Metric          | Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| MED-plus-IGP    | Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| TTL-Action      | <p>For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signalled and LDP-signalled LSPs or for specific VRF routing instances.</p> <p>For sample output, see <a href="#">show route table</a>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

Table 510: show route extensive Output Fields (*continued*)

| Field Name                           | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Task</b>                          | Name of the protocol that has added the route.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Announcement bits</b>             | List of protocols that announce this route. <b>n-Resolve inet</b> indicates that the route is used for route resolution for next hops found in the routing table. <b>n</b> is an index used by Juniper Networks customer support only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>AS path</b>                       | <p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> <li>• <b>I</b>—IGP.</li> <li>• <b>E</b>—EGP.</li> <li>• <b>Recorded</b>—The AS path is recorded by the sample process (sampled).</li> <li>• <b>?</b>—Incomplete; typically, the AS path was aggregated.</li> </ul> <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> <li>• <b>[ ]</b>—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device, or if AS path prepending is configured.</li> <li>• <b>{ }</b>—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order.</li> <li>• <b>( )</b>—Parentheses enclose a confederation.</li> <li>• <b>( [ ] )</b>—Parentheses and brackets enclose a confederation set.</li> </ul> <p><b>NOTE:</b> In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p> |
| <b>validation-state</b>              | <p>(BGP-learned routes) Validation status of the route:</p> <ul style="list-style-type: none"> <li>• <b>Invalid</b>—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database.</li> <li>• <b>Unknown</b>—Indicates that the prefix is not among the prefixes or prefix ranges in the database.</li> <li>• <b>Unverified</b>—Indicates that origin validation is not enabled for the BGP peers.</li> <li>• <b>Valid</b>—Indicates that the prefix and autonomous system pair are found in the database.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>FECs bound to route</b>           | Point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>AS path: I &lt;Originator&gt;</b> | (For route reflected output only) Originator ID attribute set by the route reflector.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |



Table 510: show route extensive Output Fields (*continued*)

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>route status</b>     | <p>Indicates the status of a BGP route:</p> <ul style="list-style-type: none"> <li>• <b>Accepted</b>—The specified BGP route is imported by the default BGP policy.</li> <li>• <b>Import</b>—The route is imported into a Layer 3 VPN routing instance.</li> <li>• <b>Import-Protect</b>—A remote instance egress that is protected.</li> <li>• <b>Multipath</b>—A BGP multipath active route.</li> <li>• <b>MultipathContrib</b>—The route is not active but contributes to the BGP multipath.</li> <li>• <b>Protect</b>—An egress route that is protected.</li> <li>• <b>Stale</b>—A route that is marked stale due to graceful restart.</li> </ul> |
| Primary Upstream        | When multipoint LDP with multicast-only fast reroute (MoFRR) is configured, the primary upstream path. MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path.                                                                                                                                                                                                                                                                                                                                   |
| RPF Nexthops            | When multipoint LDP with MoFRR is configured, the reverse-path forwarding (RPF) next-hop information. Data packets are received from both the primary path and the secondary paths. The redundant packets are discarded at topology merge points due to the RPF checks.                                                                                                                                                                                                                                                                                                                                                                               |
| Label                   | Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.                                                                                                                                                                                                                                                                                                                                                           |
| weight                  | Value used to distinguish MoFRR primary and backup routes. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| VC Label                | MPLS label assigned to the Layer 2 circuit virtual connection.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| MTU                     | Maximum transmission unit (MTU) of the Layer 2 circuit.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| VLAN ID                 | VLAN identifier of the Layer 2 circuit.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Cluster list            | (For route reflected output only) Cluster ID sent by the route reflector.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Originator ID           | (For route reflected output only) Address of router that originally sent the route to the route reflector.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Prefixes bound to route | Forwarding Equivalent Class (FEC) bound to this route. Applicable only to routes installed by LDP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Communities             | Community path attribute for the route. See the Output Field table in the <a href="#">show route detail</a> command for all possible values for this field.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Layer2-info: encaps     | Layer 2 encapsulation (for example, VPLS).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| control flags           | Control flags: <b>none</b> or Site Down.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| mtu                     | Maximum transmission unit (MTU) information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Label-Base, range       | First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

Table 510: show route extensive Output Fields (*continued*)

| Field Name                   | Field Description                                                                                                                                                                                                                                                                                                |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>status vector</b>         | Layer 2 VPN and VPLS network layer reachability information (NLRI).                                                                                                                                                                                                                                              |
| <b>Localpref</b>             | Local preference value included in the route.                                                                                                                                                                                                                                                                    |
| <b>Router ID</b>             | BGP router ID as advertised by the neighbor in the open message.                                                                                                                                                                                                                                                 |
| <b>Primary Routing Table</b> | In a routing table group, the name of the primary routing table in which the route resides.                                                                                                                                                                                                                      |
| <b>Secondary Tables</b>      | In a routing table group, the name of one or more secondary tables in which the route resides.                                                                                                                                                                                                                   |
| <b>Originating RIB</b>       | Name of the routing table whose active route was used to determine the forwarding next-hop entry in the resolution database. For example, in the case of inet.0 resolving through inet.0 and inet.3, this field indicates which routing table, inet.0 or inet.3, provided the best path for a particular prefix. |
| <b>Node path count</b>       | Number of nodes in the path.                                                                                                                                                                                                                                                                                     |
| <b>Forwarding nexthops</b>   | Number of forwarding next hops. The forwarding next hop is the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.                                                                                                                            |

## Sample Output

### show route extensive

```

user@host> show route extensive
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
 *Static Preference: 5
 Next-hop reference count: 29
 Next hop: 192.168.71.254 via fxp0.0, selected
 State: <Active NoReadvrt Int Ext>
 Local AS: 69
 Age: 1:34:06
 Task: RT
 Announcement bits (2): 0-KRT 3-Resolve tree 2
 AS path: I

10.31.1.0/30 (2 entries, 1 announced)
 *Direct Preference: 0
 Next hop type: Interface
 Next-hop reference count: 2
 Next hop: via so-0/3/0.0, selected
 State: <Active Int>
 Local AS: 69
 Age: 1:32:40
 Task: IF
 Announcement bits (1): 3-Resolve tree 2
 AS path: I
 OSPF Preference: 10
 Next-hop reference count: 1
 Next hop: via so-0/3/0.0, selected

```

```

 State: <Int>
 Inactive reason: Route Preference
 Local AS: 69
 Age: 1:32:40 Metric: 1
 Area: 0.0.0.0
 Task: OSPF
 AS path: I

10.31.1.1/32 (1 entry, 1 announced)
 *Local Preference: 0
 Next hop type: Local
 Next-hop reference count: 7
 Interface: so-0/3/0.0
 State: <Active NoReadvrt Int>
 Local AS: 69
 Age: 1:32:43
 Task: IF
 Announcement bits (1): 3-Resolve tree 2
 AS path: I

...

10.31.2.0/30 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.31.2.0/30 -> {10.31.1.6}
 *OSPF Preference: 10
 Next-hop reference count: 9
 Next hop: via so-0/3/0.0
 Next hop: 10.31.1.6 via ge-3/1/0.0, selected
 State: <Active Int>
 Local AS: 69
 Age: 1:32:19 Metric: 2
 Area: 0.0.0.0
 Task: OSPF
 Announcement bits (2): 0-KRT 3-Resolve tree 2
 AS path: I

...

224.0.0.2/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 224.0.0.2/32 -> {}
 *PIM Preference: 0
 Next-hop reference count: 18
 State: <Active NoReadvrt Int>
 Local AS: 69
 Age: 1:34:08
 Task: PIM Recv
 Announcement bits (2): 0-KRT 3-Resolve tree 2
 AS path: I

...

224.0.0.22/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 224.0.0.22/32 -> {}
 *IGMP Preference: 0
 Next-hop reference count: 18
 State: <Active NoReadvrt Int>
 Local AS: 69
 Age: 1:34:06

```

```

Task: IGMP
Announcement bits (2): 0-KRT 3-Resolve tree 2
AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.255.70.103/32 (1 entry, 1 announced)
State: <FlashAll>
*RSVP Preference: 7
Next-hop reference count: 6
Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
Label-switched-path green-r1-r3
Label operation: Push 100096
State: <Active Int>
Local AS: 69
Age: 1:28:12 Metric: 2
Task: RSVP
Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
AS path: I

10.255.71.238/32 (1 entry, 1 announced)
State: <FlashAll>
*RSVP Preference: 7
Next-hop reference count: 6
Next hop: via so-0/3/0.0 weight 0x1, selected
Label-switched-path green-r1-r2
State: <Active Int>
Local AS: 69
Age: 1:28:12 Metric: 1
Task: RSVP
Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
AS path: I

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

...

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

47.0005.80ff.f800.0000.0108.0001.0102.5507.1052/152 (1 entry, 0 announced)
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 1
Next hop: via lo0.0, selected
State: <Active Int>
Local AS: 69
Age: 1:34:07
Task: IF
AS path: I

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

0 (1 entry, 1 announced)
TSI:
KRT in-kernel 0 /36 -> {}
*MPLS Preference: 0
Next hop type: Receive
Next-hop reference count: 6
State: <Active Int>
Local AS: 69
Age: 1:34:08 Metric: 1

```

```

Task: MPLS
Announcement bits (1): 0-KRT
AS path: I

...

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
299776 (1 entry, 1 announced)
TSI:
KRT in-kernel 299776 /52 -> {Flood}
 *RSVP Preference: 7
 Next hop type: Flood
 Next-hop reference count: 130
 Flood nexthop branches exceed maximum
 Address: 0x8ea65d0

...

800010 (1 entry, 1 announced)
TSI:
KRT in-kernel 800010 /36 -> {vt-3/2/0.32769}
 *VPLS Preference: 7
 Next-hop reference count: 2
 Next hop: via vt-3/2/0.32769, selected
 Label operation: Pop
 State: <Active Int>
 Age: 1:31:53
 Task: Common L2 VC
 Announcement bits (1): 0-KRT
 AS path: I

vt-3/2/0.32769 (1 entry, 1 announced)
TSI:
KRT in-kernel vt-3/2/0.32769.0 /16 -> {indirect(1048574)}
 *VPLS Preference: 7
 Next-hop reference count: 2
 Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
 Label-switched-path green-r1-r3
 Label operation: Push 800012, Push 100096(top)
 Protocol next hop: 10.255.70.103
 Push 800012
 Indirect next hop: 87272e4 1048574
 State: <Active Int>
 Age: 1:31:53 Metric2: 2
 Task: Common L2 VC
 Announcement bits (2): 0-KRT 1-Common L2 VC
 AS path: I
 Communities: target:11111:1 Layer2-info: encaps:VPLS,
 control flags:, mtu: 0
 Indirect next hops: 1
 Protocol next hop: 10.255.70.103 Metric: 2
 Push 800012
 Indirect next hop: 87272e4 1048574
 Indirect path forwarding next hops: 1
 Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1
 10.255.70.103/32 Originating RIB: inet.3
 Metric: 2 Node path count: 1
 Forwarding nexthops: 1
 Nexthop: 10.31.1.6 via ge-3/1/0.0

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

```

```
abcd::10:255:71:52/128 (1 entry, 0 announced)
 *Direct Preference: 0
 Next hop type: Interface
 Next-hop reference count: 1
 Next hop: via lo0.0, selected
 State: <Active Int>
 Local AS: 69
 Age: 1:34:07
 Task: IF
 AS path: I

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
 *Direct Preference: 0
 Next hop type: Interface
 Next-hop reference count: 1
 Next hop: via lo0.0, selected
 State: <Active NoReadvrt Int>
 Local AS: 69
 Age: 1:34:07
 Task: IF
 AS path: I

ff02::2/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::2/128 -> {}
 *PIM Preference: 0
 Next-hop reference count: 18
 State: <Active NoReadvrt Int>
 Local AS: 69
 Age: 1:34:08
 Task: PIM Recv6
 Announcement bits (1): 0-KRT
 AS path: I

ff02::d/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::d/128 -> {}
 *PIM Preference: 0
 Next-hop reference count: 18
 State: <Active NoReadvrt Int>
 Local AS: 69
 Age: 1:34:08
 Task: PIM Recv6
 Announcement bits (1): 0-KRT
 AS path: I

ff02::16/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::16/128 -> {}
 *MLD Preference: 0
 Next-hop reference count: 18
 State: <Active NoReadvrt Int>
 Local AS: 69
 Age: 1:34:06
 Task: MLD
 Announcement bits (1): 0-KRT
 AS path: I

private.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

```

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
 *Direct Preference: 0
 Next hop type: Interface
 Next-hop reference count: 1
 Next hop: via lo0.16385, selected
 State: <Active NoReadvrt Int>
 Age: 1:34:07
 Task: IF
 AS path: I

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

10.255.70.103:1:3:1/96 (1 entry, 1 announced)
 *BGP Preference: 170/-101
 Route Distinguisher: 10.255.70.103:1
 Next-hop reference count: 7
 Source: 10.255.70.103
 Protocol next hop: 10.255.70.103
 Indirect next hop: 2 no-forward
 State: <Secondary Active Int Ext>
 Local AS: 69 Peer AS: 69
 Age: 1:28:12 Metric2: 1
 Task: BGP_69.10.255.70.103+179
 Announcement bits (1): 0-green-l2vpn
 AS path: I
 Communities: target:11111:1 Layer2-info: encaps:VPLS,
 control flags:, mtu: 0
 Label-base: 800008, range: 8
 Localpref: 100
 Router ID: 10.255.70.103
 Primary Routing Table bgp.l2vpn.0

10.255.71.52:1:1:1/96 (1 entry, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 8699540
 *L2VPN Preference: 170/-1
 Next-hop reference count: 5
 Protocol next hop: 10.255.71.52
 Indirect next hop: 0 -
 State: <Active Int Ext>
 Age: 1:34:03 Metric2: 1
 Task: green-l2vpn
 Announcement bits (1): 1-BGP.0.0.0.0+179
 AS path: I
 Communities: Layer2-info: encaps:VPLS, control flags:Site-Down,
 mtu: 0
 Label-base: 800016, range: 8, status-vector: 0x9F

10.255.71.52:1:5:1/96 (1 entry, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 8699528
 *L2VPN Preference: 170/-101
 Next-hop reference count: 5
 Protocol next hop: 10.255.71.52
 Indirect next hop: 0 -
 State: <Active Int Ext>
 Age: 1:34:03 Metric2: 1
 Task: green-l2vpn
 Announcement bits (1): 1-BGP.0.0.0.0+179
 AS path: I
 Communities: Layer2-info: encaps:VPLS, control flags:, mtu: 0

```

```
Label-base: 800008, range: 8, status-vector: 0x9F

...

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

TSI:

10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
 *L2CKT Preference: 7
 Next hop: via so-1/1/2.0 weight 1, selected
 Label-switched-path my-lsp
 Label operation: Push 100000[0]
 Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
 State: <Active Int>
 Local AS: 99
 Age: 10:21
 Task: l2 circuit
 Announcement bits (1): 0-LDP
 AS path: I
 VC Label 100000, MTU 1500, VLAN ID 512

55.0.0.0/24 (1 entry, 1 announced)
TSI:
KRT queued (pending) add
 55.0.0.0/24 -> {Push 300112}
 *BGP Preference: 170/-101
 Next hop type: Router
 Address: 0x925c208
 Next-hop reference count: 2
 Source: 10.0.0.9
 Next hop: 10.0.0.9 via lt-1/2/0.15, selected
 Label operation: Push 300112
 Label TTL action: prop-ttl
 State: <Active Ext>
 Local AS: 7019 Peer AS: 13979
 Age: 1w0d 23:06:56
 AIGP: 25
 Task: BGP_13979.10.0.0.9+56732
 Announcement bits (1): 0-KRT
 AS path: 13979 7018 I
 Accepted
 Route Label: 300112
 Localpref: 100
 Router ID: 10.9.9.1
```

#### show route extensive (Access Route)

```
user@host> show route 13.160.0.102 extensive
inet.0: 39256 destinations, 39258 routes (39255 active, 0 holddown, 1 hidden)
13.160.0.102/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 13.160.0.102/32 -> {13.160.0.2}
OSPF area : 0.0.0.0, LSA ID : 13.160.0.102, LSA type : Extern
 *Access Preference: 13
 Next-hop reference count: 78472
 Next hop: 13.160.0.2 via fe-0/0/0.0, selected
 State: <Active Int>
 Age: 12
 Task: RPD Unix Domain Server./var/run/rpd_serv.local
```



```
Announcement bits (2): 0-KRT 1-OSPFv2
AS path: I
```

### show route extensive (BGP PIC Edge)

```
user@host> show route 1.1.1.6 extensive
ed.inet.0: 6 destinations, 9 routes (6 active, 0 holddown, 0 hidden)
 1.1.1.6/32 (3 entries, 2 announced)
 State: <CalcForwarding>
 TSI:
 KRT in-kernel 1.1.1.6/32 -> {indirect(1048574), indirect(1048577)}
 Page 0 idx 0 Type 1 val 9219e30
 Nexthop: Self
 AS path: [2] 3 I
 Communities: target:2:1
 Path 1.1.1.6 from 1.1.1.4 Vector len 4. Val: 0
 ..
 #Multipath Preference: 255
 Next hop type: Indirect
 Address: 0x93f4010
 Next-hop reference count: 2
 ..
 Protocol next hop: 1.1.1.4
 Push 299824
 Indirect next hop: 944c000 1048574 INH Session ID: 0x3
 Indirect next hop: weight 0x1
 Protocol next hop: 1.1.1.5
 Push 299824
 Indirect next hop: 944c1d8 1048577 INH Session ID: 0x4
 Indirect next hop: weight 0x4000
 State: <ForwardingOnly Int Ext>
 Inactive reason: Forwarding use only
 Age: 25 Metric2: 15
 Validation State: unverified
 Task: RT
 Announcement bits (1): 0-KRT
 AS path: 3 I
 Communities: target:2:1
```

### show route extensive (FRR and LFA)

```
user@host> show route 20.31.2.0 extensive
inet.0: 46 destinations, 49 routes (45 active, 0 holddown, 1 hidden)
 20.31.2.0/24 (2 entries, 1 announced)
 State: FlashAll
 TSI:
 KRT in-kernel 20.31.2.0/24 -> {Push 299776, Push 299792}
 *RSVP Preference: 7/1
 Next hop type: Router, Next hop index: 1048574
 Address: 0xbbbc010
 Next-hop reference count: 5
 Next hop: 10.31.1.2 via ge-2/1/8.0 weight 0x1, selected
 Label-switched-path europa-d-to-europa-e
 Label operation: Push 299776
 Label TTL action: prop-ttl
 Session Id: 0x201
 Next hop: 10.31.2.2 via ge-2/1/4.0 weight 0x4001
 Label-switched-path europa-d-to-europa-e
 Label operation: Push 299792
 Label TTL action: prop-ttl
 Session Id: 0x202
```

```

State: Active Int
Local AS: 100
Age: 5:31 Metric: 2
Task: RSVP
Announcement bits (1): 0-KRT
AS path: I
OSPF Preference: 10
Next hop type: Router, Next hop index: 615
Address: 0xb9d78c4
Next-hop reference count: 7
Next hop: 10.31.1.2 via ge-2/1/8.0, selected
Session Id: 0x201
State: Int
Inactive reason: Route Preference
Local AS: 100
Age: 5:35 Metric: 3
Area: 0.0.0.0
Task: OSPF
AS path: I

```

#### show route extensive (Route Reflector)

```

user@host> show route extensive
1.0.0.0/8 (1 entry, 1 announced)

TSI:
KRT in-kernel 1.0.0.0/8 -> {indirect(40)}
*BGP Preference: 170/-101
Source: 192.168.4.214
Protocol next hop: 207.17.136.192 Indirect next hop: 84ac908 40
State: <Active Int Ext>
Local AS: 10458 Peer AS: 10458
Age: 3:09 Metric: 0 Metric2: 0
Task: BGP_10458.192.168.4.214+1033
Announcement bits (2): 0-KRT 4-Resolve inet.0
AS path: 3944 7777 I <Originator>
Cluster list: 1.1.1.1
Originator ID: 10.255.245.88
Communities: 7777:7777
Localpref: 100
Router ID: 4.4.4.4
Indirect next hops: 1
 Protocol next hop: 207.17.136.192 Metric: 0
 Indirect next hop: 84ac908 40
 Indirect path forwarding next hops: 0
 Next hop type: Discard

```

#### show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```

user@host> show route label 299872 detail
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
299872 (1 entry, 1 announced)
*LDP Preference: 9
Next hop type: Flood
Next-hop reference count: 3
Address: 0x9097d90
Next hop: via vt-0/1/0.1
Next-hop index: 661
Label operation: Pop
Address: 0x9172130
Next hop: via so-0/0/3.0

```

```

Next-hop index: 654
Label operation: Swap 299872
State: **Active Int>
Local AS: 1001
Age: 8:20 Metric: 1
Task: LDP
Announcement bits (1): 0-KRT
AS path: I
FECs bound to route: P2MP root-addr 10.255.72.166, grp 232.1.1.1,
src 192.168.142.2

```

#### show route label detail (Multipoint LDP with Multicast-Only Fast Reroute)

```
user@host> show route label 301568 detail
```

```

mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
301568 (1 entry, 1 announced)
 *LDP Preference: 9
 Next hop type: Flood
 Address: 0x2735208
 Next-hop reference count: 3
 Next hop type: Router, Next hop index: 1397
 Address: 0x2735d2c
 Next-hop reference count: 3
 Next hop: 1.3.8.2 via ge-1/2/22.0
 Label operation: Pop
 Load balance label: None;
 Next hop type: Router, Next hop index: 1395
 Address: 0x2736290
 Next-hop reference count: 3
 Next hop: 1.3.4.2 via ge-1/2/18.0
 Label operation: Pop
 Load balance label: None;
 State: <Active Int AckRequest MulticastRPF>
 Local AS: 10
 Age: 54:05 Metric: 1
 Validation State: unverified
 Task: LDP
 Announcement bits (1): 0-KRT
 AS path: I
 FECs bound to route: P2MP root-addr 1.1.1.1, grp: 232.1.1.1, src:
192.168.219.11
 Primary Upstream : 1.1.1.3:0--1.1.1.2:0
 RPF Nexthops :
 ge-1/2/15.0, 1.2.94.1, Label: 301568, weight: 0x1
 ge-1/2/14.0, 1.2.3.1, Label: 301568, weight: 0x1
 Backup Upstream : 1.1.1.3:0--1.1.1.6:0
 RPF Nexthops :
 ge-1/2/20.0, 1.2.96.1, Label: 301584, weight: 0xffffe
 ge-1/2/19.0, 1.3.6.1, Label: 301584, weight: 0xffffe

```

## show route protocol

---

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 6404</a><br><a href="#">Syntax (EX Series Switches) on page 6404</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Syntax</b>                      | <code>show route protocol <i>protocol</i></code><br><brief   detail   extensive   terse><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Syntax (EX Series Switches)</b> | <code>show route protocol <i>protocol</i></code><br><brief   detail   extensive   terse>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Options <b>ospf2</b> and <b>ospf3</b> introduced in Junos OS Release 9.2.<br>Options <b>ospf2</b> and <b>ospf3</b> introduced in Junos OS Release 9.2 for EX Series switches.<br>Option <b>flow</b> introduced in Junos OS Release 10.0.<br>Option <b>flow</b> introduced in Junos OS Release 10.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>                 | Display the route entries in the routing table that were learned from a particular protocol.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                     | <b>brief   detail   extensive   terse</b> —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.<br><br><b><i>protocol</i></b> —Protocol from which the route was learned: <ul style="list-style-type: none"><li>• <b>access</b>—Access route for use by DHCP application</li><li>• <b>access-internal</b>—Access-internal route for use by DHCP application</li><li>• <b>aggregate</b>—Locally generated aggregate route</li><li>• <b>atmvpn</b>—Asynchronous Transfer Mode virtual private network</li><li>• <b>bgp</b>—Border Gateway Protocol</li><li>• <b>ccc</b>—Circuit cross-connect</li><li>• <b>direct</b>—Directly connected route</li><li>• <b>dvmrp</b>—Distance Vector Multicast Routing Protocol</li><li>• <b>esis</b>—End System-to-Intermediate System</li><li>• <b>flow</b>—Locally defined flow-specification route.</li><li>• <b>isis</b>—Intermediate System-to-Intermediate System</li><li>• <b>ldp</b>—Label Distribution Protocol</li><li>• <b>l2circuit</b>—Layer 2 circuit</li><li>• <b>l2vpn</b>—Layer 2 virtual private network</li><li>• <b>local</b>—Local address</li></ul> |

- **mpls**—Multiprotocol Label Switching
- **msdp**—Multicast Source Discovery Protocol
- **ospf**—Open Shortest Path First versions 2 and 3
- **ospf2**—Open Shortest Path First version 2 only
- **ospf3**—Open Shortest Path First version 3 only
- **pim**—Protocol Independent Multicast
- **rip**—Routing Information Protocol
- **ripng**—Routing Information Protocol next generation
- **rsvp**—Resource Reservation Protocol
- **rtarget**—Local route target virtual private network
- **static**—Statically defined route
- **tunnel**—Dynamic tunnel
- **vpn**—Virtual private network



**NOTE:** EX Series switches run a subset of these protocols. See the switch CLI for details.

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>List of Sample Output</b>    | <a href="#">show route protocol access on page 6406</a><br><a href="#">show route protocol access-internal extensive on page 6406</a><br><a href="#">show route protocol bgp on page 6406</a><br><a href="#">show route protocol bgp detail on page 6406</a><br><a href="#">show route protocol bgp extensive on page 6407</a><br><a href="#">show route protocol bgp terse on page 6407</a><br><a href="#">show route protocol direct on page 6407</a><br><a href="#">show route protocol l2circuit detail on page 6408</a><br><a href="#">show route protocol l2vpn extensive on page 6409</a><br><a href="#">show route protocol ldp on page 6409</a><br><a href="#">show route protocol ldp extensive on page 6410</a><br><a href="#">show route protocol ospf (Layer 3 VPN) on page 6411</a><br><a href="#">show route protocol ospf detail on page 6412</a><br><a href="#">show route protocol rip on page 6412</a><br><a href="#">show route protocol rip detail on page 6412</a><br><a href="#">show route protocol ripng table inet6 on page 6412</a> |
| <b>Output Fields</b>            | For information about output fields, see the output field tables for the <a href="#">show route</a> command, the <a href="#">show route detail</a> command, the <a href="#">show route extensive</a> command, or the <a href="#">show route terse</a> command.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## Sample Output

### show route protocol access

```

user@host> show route protocol access
inet.0: 30380 destinations, 30382 routes (30379 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

13.160.0.3/32 *[Access/13] 00:00:09
 > to 13.160.0.2 via fe-0/0/0.0
13.160.0.4/32 *[Access/13] 00:00:09
 > to 13.160.0.2 via fe-0/0/0.0
13.160.0.5/32 *[Access/13] 00:00:09
 > to 13.160.0.2 via fe-0/0/0.0

```

### show route protocol access-internal extensive

```

user@host> show route protocol access-internal 13.160.0.19 extensive
inet.0: 100020 destinations, 100022 routes (100019 active, 0 holddown, 1 hidden)
13.160.0.19/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 13.160.0.19/32 -> {13.160.0.2}
 *Access-internal Preference: 12
 Next-hop reference count: 200000
 Next hop: 13.160.0.2 via fe-0/0/0.0, selected
 State: <Active Int>
 Age: 36
 Task: RPD Unix Domain Server./var/run/rpd_serv.local
 Announcement bits (1): 0-KRT
 AS path: I

```

### show route protocol bgp

```

user@host> show route protocol bgp 192.168.64.0/21
inet.0: 335832 destinations, 335833 routes (335383 active, 0 holddown, 450 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.64.0/21 *[BGP/170] 6d 10:41:16, localpref 100, from 192.168.69.71
 AS path: 10458 14203 2914 4788 4788 I
 > to 192.168.167.254 via fxp0.0

```

### show route protocol bgp detail

```

show route protocol bgp 66.117.63.0/24 exact detail
inet.0: 335805 destinations, 335806 routes (335356 active, 0 holddown, 450 hidden)
66.117.63.0/24 (1 entry, 1 announced)
 *BGP Preference: 170/-101
 Next hop type: Indirect
 Next-hop reference count: 1006436
 Source: 192.168.69.71
 Next hop type: Router, Next hop index: 324
 Next hop: 192.168.167.254 via fxp0.0, selected
 Protocol next hop: 192.168.69.71
 Indirect next hop: 8e166c0 342
 State: <Active Ext>
 Local AS: 69 Peer AS: 10458
 Age: 6d 10:42:42 Metric2: 0
 Task: BGP_10458.192.168.69.71+179
 Announcement bits (3): 0-KRT 2-BGP RT Background 3-Resolve tree
1
 AS path: 10458 14203 2914 4788 4788 I

```

```

Communities: 2914:410 2914:2403 2914:3400
Accepted
Localpref: 100
Router ID: 207.17.136.192

```

### show route protocol bgp extensive

```
user@host> show route protocol bgp 192.168.64.0/21 extensive
```

```
inet.0: 335827 destinations, 335828 routes (335378 active, 0 holddown, 450 hidden)
192.168.64.0/21 (1 entry, 1 announced)
TSI:
```

```

KRT in-kernel 1.9.0.0/16 -> {indirect(342)}
Page 0 idx 1 Type 1 val db31a80
 Nexthop: Self
 AS path: [69] 10458 14203 2914 4788 4788 I
 Communities: 2914:410 2914:2403 2914:3400
Path 1.9.0.0 from 192.168.69.71 Vector len 4. Val: 1
 *BGP Preference: 170/-101
 Next hop type: Indirect
 Next-hop reference count: 1006502
 Source: 192.168.69.71
 Next hop type: Router, Next hop index: 324
 Next hop: 192.168.167.254 via fxp0.0, selected
 Protocol next hop: 192.168.69.71
 Indirect next hop: 8e166c0 342
 State: <Active Ext>
 Local AS: 69 Peer AS: 10458
 Age: 6d 10:44:45 Metric2: 0
 Task: BGP_10458.192.168.69.71+179
 Announcement bits (3): 0-KRT 2-BGP RT Background 3-Resolve tree

```

1

```

AS path: 10458 14203 2914 4788 4788 I
Communities: 2914:410 2914:2403 2914:3400
Accepted
Localpref: 100
Router ID: 207.17.136.192
Indirect next hops: 1
 Protocol next hop: 192.168.69.71
 Indirect next hop: 8e166c0 342
 Indirect path forwarding next hops: 1
 Next hop type: Router
 Next hop: 192.168.167.254 via fxp0.0
 192.168.0.0/16 Originating RIB: inet.0
 Node path count: 1
 Forwarding nexthops: 1
 Nexthop: 192.168.167.254 via fxp0.0

```

### show route protocol bgp terse

```
user@host> show route protocol bgp 192.168.64.0/21 terse
```

```
inet.0: 24 destinations, 32 routes (23 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
```

| A Destination   | P Prf | Metric 1 | Metric 2 | Next hop   | AS path    |
|-----------------|-------|----------|----------|------------|------------|
| 192.168.64.0/21 | B 170 | 100      |          | >100.1.3.2 | 10023 21 I |

### show route protocol direct

```
user@host> show route protocol direct
```

```

inet.0: 335843 destinations, 335844 routes (335394 active, 0 holddown, 450 hidden)
+ = Active Route, - = Last Active, * = Both

8.8.8.0/24 *[Direct/0] 17w0d 10:31:49
 > via fe-1/3/1.0
10.255.165.1/32 *[Direct/0] 25w4d 04:13:18
 > via lo0.0
30.30.30.0/24 *[Direct/0] 17w0d 23:06:26
 > via fe-1/3/2.0
192.168.164.0/22 *[Direct/0] 25w4d 04:13:20
 > via fxp0.0

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

47.0005.80ff.f800.0000.0108.0001.0102.5516.5001/152
 *[Direct/0] 25w4d 04:13:21
 > via lo0.0

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

abcd::10:255:165:1/128
 *[Direct/0] 25w4d 04:13:21
 > via lo0.0
fe80::2a0:a5ff:fe12:ad7/128
 *[Direct/0] 25w4d 04:13:21
 > via lo0.0

```

### show route protocol l2circuit detail

```

user@host> show route protocol l2circuit detail

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
100000 (1 entry, 1 announced)
 *L2CKT Preference: 7
 Next hop: via ge-2/0/0.0, selected
 Label operation: Pop Offset: 4
 State: <Active Int>
 Local AS: 99
 Age: 9:52
 Task: Common L2 VC
 Announcement bits (1): 0-KRT
 AS path: I

ge-2/0/0.0 (1 entry, 1 announced)
 *L2CKT Preference: 7
 Next hop: via so-1/1/2.0 weight 1, selected
 Label-switched-path my-lsp
 Label operation: Push 100000, Push 100000(top)[0] Offset: -4
 Protocol next hop: 10.245.255.63
 Push 100000 Offset: -4
 Indirect next hop: 86af0c0 298
 State: <Active Int>
 Local AS: 99
 Age: 9:52
 Task: Common L2 VC
 Announcement bits (2): 0-KRT 1-Common L2 VC
 AS path: I

```



```

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
 *L2CKT Preference: 7
 Next hop: via so-1/1/2.0 weight 1, selected
 Label-switched-path my-lsp
 Label operation: Push 100000[0]
 Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
 State: <Active Int>
 Local AS: 99
 Age: 10:21
 Task: l2 circuit
 Announcement bits (1): 0-LDP
 AS path: I
 VC Label 100000, MTU 1500, VLAN ID 512

```

### show route protocol l2vpn extensive

```

user@host> show route protocol l2vpn extensive

inet.0: 14 destinations, 15 routes (13 active, 0 holddown, 1 hidden)

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
800001 (1 entry, 1 announced)
TSI:
KRT in-kernel 800001 /36 -> {so-0/0/0.0}
 *L2VPN Preference: 7
 Next hop: via so-0/0/0.0 weight 49087 balance 97%, selected
 Label operation: Pop Offset: 4
 State: <Active Int>
 Local AS: 69
 Age: 7:48
 Task: Common L2 VC
 Announcement bits (1): 0-KRT
 AS path: I

so-0/0/0.0 (1 entry, 1 announced)
TSI:
KRT in-kernel so-0/0/0.0 /16 -> {indirect(288)}
 *L2VPN Preference: 7
 Next hop: via so-0/0/1.0, selected
 Label operation: Push 800000 Offset: -4
 Protocol next hop: 10.255.14.220
 Push 800000 Offset: -4
 Indirect next hop: 85142a0 288
 State: <Active Int>
 Local AS: 69
 Age: 7:48
 Task: Common L2 VC
 Announcement bits (2): 0-KRT 1-Common L2 VC
 AS path: I
 Communities: target:69:1 Layer2-info: encaps:PPP,
 control flags:2, mtu: 0

```

### show route protocol ldp

```

user@host> show route protocol ldp

```

```

inet.0: 12 destinations, 13 routes (12 active, 0 holddown, 0 hidden)

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.16.1/32 *[LDP/9] 1d 23:03:35, metric 1
 > via t1-4/0/0.0, Push 100000
192.168.17.1/32 *[LDP/9] 1d 23:03:35, metric 1
 > via t1-4/0/0.0

private1___.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100064 *[LDP/9] 1d 23:03:35, metric 1
 > via t1-4/0/0.0, Pop
100064(S=0) *[LDP/9] 1d 23:03:35, metric 1
 > via t1-4/0/0.0, Pop
100080 *[LDP/9] 1d 23:03:35, metric 1
 > via t1-4/0/0.0, Swap 100000

```

#### show route protocol ldp extensive

```

user@host> show route protocol ldp extensive
192.168.16.1/32 (1 entry, 1 announced)
 State: <FlashAll>
 *LDP Preference: 9
 Next-hop reference count: 3
 Next hop: via t1-4/0/0.0, selected
 Label operation: Push 100000
 State: <Active Int>
 Local AS: 65500
 Age: 1d 23:03:58 Metric: 1
 Task: LDP
 Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 2
 AS path: I

192.168.17.1/32 (1 entry, 1 announced)
 State: <FlashAll>
 *LDP Preference: 9
 Next-hop reference count: 3
 Next hop: via t1-4/0/0.0, selected
 State: <Active Int>
 Local AS: 65500
 Age: 1d 23:03:58 Metric: 1
 Task: LDP
 Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 2
 AS path: I

private1___.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)

100064 (1 entry, 1 announced)
TSI:
KRT in-kernel 100064 /36 -> {t1-4/0/0.0}
 *LDP Preference: 9
 Next-hop reference count: 2
 Next hop: via t1-4/0/0.0, selected
 State: <Active Int>

```

```

Local AS: 65500
Age: 1d 23:03:58 Metric: 1
Task: LDP
Announcement bits (1): 0-KRT
AS path: I
Prefixes bound to route: 192.168.17.1/32

100064(S=0) (1 entry, 1 announced)
TSI:
KRT in-kernel 100064 /40 -> {t1-4/0/0.0}
 *LDP Preference: 9
 Next-hop reference count: 2
 Next hop: via t1-4/0/0.0, selected
 Label operation: Pop
 State: <Active Int>
 Local AS: 65500
 Age: 1d 23:03:58 Metric: 1
 Task: LDP
 Announcement bits (1): 0-KRT
 AS path: I

100080 (1 entry, 1 announced)
TSI:
KRT in-kernel 100080 /36 -> {t1-4/0/0.0}
 *LDP Preference: 9
 Next-hop reference count: 2
 Next hop: via t1-4/0/0.0, selected
 Label operation: Swap 100000
 State: <Active Int>
 Local AS: 65500
 Age: 1d 23:03:58 Metric: 1
 Task: LDP
 Announcement bits (1): 0-KRT
 AS path: I
 Prefixes bound to route: 192.168.16.1/32

```

### show route protocol ospf (Layer 3 VPN)

```

user@host> show route protocol ospf
inet.0: 40 destinations, 40 routes (39 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.4/30 *[OSPF/10] 00:05:18, metric 4
 > via t3-3/2/0.0
10.39.1.8/30 [OSPF/10] 00:05:18, metric 2
 > via t3-3/2/0.0
10.255.14.171/32 *[OSPF/10] 00:05:18, metric 4
 > via t3-3/2/0.0
10.255.14.179/32 *[OSPF/10] 00:05:18, metric 2
 > via t3-3/2/0.0
224.0.0.5/32 *[OSPF/10] 20:25:55, metric 1

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.16/30 [OSPF/10] 00:05:43, metric 1
 > via so-0/2/2.0
10.255.14.173/32 *[OSPF/10] 00:05:43, metric 1
 > via so-0/2/2.0
224.0.0.5/32 *[OSPF/10] 20:26:20, metric 1

```

### show route protocol ospf detail

```
user@host> show route protocol ospf detail
VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.16/30 (2 entries, 0 announced)
 OSPF Preference: 10
 Nexthop: via so-0/2/2.0, selected
 State: <Int>
 Inactive reason: Route Preference
 Age: 6:25 Metric: 1
 Area: 0.0.0.0
 Task: VPN-AB-OSPF
 AS path: I
 Communities: Route-Type:0.0.0.0:1:0

...
```

### show route protocol rip

```
user@host> show route protocol rip
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32 * [RIP/100] 20:24:34, metric 2
 > to 10.39.1.22 via t3-0/2/2.0
224.0.0.9/32 * [RIP/100] 00:03:59, metric 1
```

### show route protocol rip detail

```
user@host> show route protocol rip detail
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32 (1 entry, 1 announced)
 *RIP Preference: 100
 Nexthop: 10.39.1.22 via t3-0/2/2.0, selected
 State: <Active Int>
 Age: 20:25:02 Metric: 2
 Task: VPN-AB-RIPv2
 Announcement bits (2): 0-KRT 2-BGP.0.0.0.0+179
 AS path: I
 Route learned from 10.39.1.22 expires in 96 seconds
```

### show route protocol ripng table inet6

```
user@host> show route protocol ripng table inet6
inet6.0: 4215 destinations, 4215 routes (4214 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

1111::1/128 * [RIPng/100] 02:13:33, metric 2
 > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::2/128 * [RIPng/100] 02:13:33, metric 2
 > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::3/128 * [RIPng/100] 02:13:33, metric 2
 > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
```

```
1111::4/128 *[RIPng/100] 02:13:33, metric 2
 > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::5/128 *[RIPng/100] 02:13:33, metric 2
 > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::6/128 *[RIPng/100] 02:13:33, metric 2
 > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
```



# Storage Feature Guide for QFX10000 Switches





# Storage Overview

- Overview of FIP on page 6417
- Understanding Fibre Channel Terminology on page 6417

## Overview of FIP

---

Fibre Channel over Ethernet (FCoE) Initialization Protocol (FIP) is a Layer 2 protocol that establishes and maintains Fibre Channel (FC) virtual links between pairs of FCoE devices such as server FCoE Nodes (ENodes) and FC switches. FIP can also establish and maintain virtual links between FCoE devices and an FCoE-FC gateway (such as the QFX3500 switch), where the gateway acts on behalf of the FC switch.

FIP enables FCoE devices to discover one another and to initialize and maintain virtual links over a physical Ethernet network. This allows FCoE devices in the Ethernet network to access storage devices in the FC storage area network (SAN).

FIP solves the problem presented by the FC requirement for point-to-point connections (FC does not permit point-to-multipoint connections) by creating a unique virtual link for each connection between an ENode VN\_Port and an FC switch VF\_Port. Multiple virtual links can use a single physical link and virtual links can traverse Ethernet transit (passthrough) switches while appearing to be direct point-to-point connections to the FC switch.

FIP has its own EtherType (0x8914) to distinguish its traffic from payload-carrying FCoE traffic and other Ethernet traffic. FIP operations occur on a per-VLAN basis.

For more details about FIP, see the Technical Committee T11 organization document *Fibre Channel Backbone - 5 (FC-BB-5) Rev 2.00* available at <http://www.t11.org/ftp/t11/pub/fc/bb-5/09-056v5.pdf>.

## Understanding Fibre Channel Terminology

---

To understand the Fibre Channel (FC) and Fibre Channel over Ethernet (FCoE) capabilities of the QFX Series, you should become familiar with the terms defined in [Table 511](#).

Table 511: Fibre Channel Terms

| Term                            | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| addressing mode                 | <p>Format for the locally unique MAC address the FC switch assigns to FCoE devices for FCoE transactions after FIP establishes a connection between an FCoE device and the FC switch. The two addressing modes are <i>fabric-provided MAC address (FPMA)</i> and <i>server-provided MAC address (SPMA)</i>. The QFX Series supports only FPMA.</p> <p>During FLOGI or FDISC, the ENode advertises the addressing modes it supports. If the FC switch supports an addressing mode that the ENode uses, the virtual link can be established, and the devices can communicate.</p> <p>See also <i>fabric-provided MAC address (FPMA)</i> and <i>server-provided MAC address (SPMA)</i>.</p>                   |
| ALL-ENode-MACs                  | <p>Well-known multicast MAC address to which all FCoE ENodes listen. FCFs send multicast <i>FIP discovery advertisement</i> messages and <i>FIP keepalive</i> messages to the ALL-ENode-MACs address so that ENodes can discover and maintain connections to FCFs. The hexadecimal format of the address is <b>01:10:18:01:00:01</b>.</p> <p>See also <i>well-known address (WKA)</i>.</p>                                                                                                                                                                                                                                                                                                                 |
| ALL-FCF-MACs                    | <p>Well-known multicast MAC address to which all FCFs listen. ENodes send multicast <i>FIP discovery solicitation</i> messages to the ALL-FCF-MACs address to find out which FCFs can accept a login. The hexadecimal format of the address is <b>01:10:18:01:00:02</b>.</p> <p>See also <i>well-known address (WKA)</i>.</p>                                                                                                                                                                                                                                                                                                                                                                              |
| congestion notification         | See <i>quantized congestion notification (QCN)</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| converged network adapter (CNA) | <p>Physical adapter that combines the functions of a Fibre Channel <i>host bus adapter (HBA)</i> to process Fibre Channel frames and a <i>lossless Ethernet network interface card (NIC)</i> to process Ethernet frames. CNAs have one or more Ethernet ports. CNAs encapsulate Fibre Channel frames in Ethernet for FCoE transport and de-encapsulate Fibre Channel frames from FCoE to native Fibre Channel.</p> <p>See also <i>host bus adapter (HBA)</i>.</p>                                                                                                                                                                                                                                          |
| data center bridging (DCB)      | <p>Set of IEEE specifications that enhance the Ethernet standard to allow it to support converged Ethernet (LAN) and Fibre Channel (SAN) traffic on one Ethernet network. DCB features include <i>priority-based flow control (PFC)</i>, <i>enhanced transmission selection (ETS)</i>, <i>Data Center Bridging Capability Exchange protocol (DCBX)</i>, <i>quantized congestion notification (QCN)</i>, and full-duplex 10-Gigabit Ethernet ports.</p> <p>See also <i>priority-based flow control (PFC)</i>, <i>Ethernet PAUSE</i>, <i>enhanced transmission selection (ETS)</i>, <i>Data Center Bridging Capability Exchange protocol (DCBX)</i>, and <i>quantized congestion notification (QCN)</i>.</p> |
| expansion port (E_Port)         | An expansion port in an FC switch/FCF that connects the FC switch/FCF to the E_Port of another FC switch/FCF to form an <i>Interswitch Link (ISL)</i> in a common FC fabric.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

Table 511: Fibre Channel Terms (*continued*)

| Term                                                     | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|----------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Data Center Bridging Capability Exchange protocol (DCBX) | Discovery and exchange protocol for conveying configuration and capabilities among neighbors to ensure consistent configuration across the network. It is an extension of the Link Layer Data Protocol (LLDP, described in IEEE 802.1AB)<br><br>See also <i>data center bridging (DCB)</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| enhanced transmission selection (ETS)                    | Mechanism that provides finer granularity of bandwidth management within a link.<br><br>See also <i>data center bridging (DCB)</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| ENode                                                    | See <i>FCoE Node (ENode)</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| ENode MAC                                                | <i>Lossless Ethernet MAC</i> paired with an <i>FCoE controller</i> in an ENode.<br><br>See also <i>FCoE node (ENode)</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| ENode MAC address                                        | Globally unique address assigned to the CNA by the manufacturer and used to identify the node for FIP transactions.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Ethernet PAUSE                                           | As defined in IEEE 802.3X, a flow control mechanism that temporarily stops the transmission of Ethernet frames on a link for a specified period. A receiving element sends an Ethernet PAUSE frame when a sender transmits data faster than the receiver can accept it. Ethernet PAUSE affects the entire link, not just an individual flow. An Ethernet PAUSE frame temporarily stops all traffic transmission on the link and allows the receiver's input buffer to empty sufficiently to restart traffic on the link. Ethernet PAUSE messages are sent to the previous hop and do not automatically propagate to the source of the congestion.<br><br>See also <i>priority-based flow control (PFC)</i> .                                                                                                                                                                                                                                                                                                                                           |
| fabric                                                   | Interconnection of network nodes using one or more network switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| fabric discovery (FDISC)                                 | Subsequent logins from the same ENode for different users, applications, or virtual machines after an ENode performs an initial FLOGI to log in to a switch.<br><br>FC and FIP FDISC messages serve the same function in FC and FCoE networks, respectively. N_Ports send FC FDISC messages to the FC switch and VN_Ports send FIP FDISC messages to the FCF.<br><br>After an N_Port acquires its initial N_Port ID through the FC FLOGI process, it can acquire additional N_Port IDs by sending an FC FDISC with a new worldwide port name and a source ID of 0x000000. The new port name and blank source ID tell the FC switch to assign a new N_Port ID to the N_Port. The different N_Port IDs allow multiple virtual machines or users on the N_Port to have separate, secure virtual links on the same physical N_Port. These additional ports are also referred to as VN_Ports.<br><br>FIP FDISC works the same way, except the VN_Port logs in using a FIP FLOGI message.<br><br>See also <i>fabric login (FLOGI)</i> and <i>N_Port ID</i> . |

Table 511: Fibre Channel Terms (*continued*)

| Term                               | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| fabric login (FLOGI)               | <p>Creation of a logical connection to the FC switch and establishment of a node's operating environment.</p> <p>For FC devices, an N_Port logs in to the FC network by sending an FC FLOGI message to the F_Port of an FC switch.</p> <p>For FCoE devices, a VN_Port logs in to the FC network by sending a FIP FLOGI message to the VF_Port of an FC switch.</p>                                                                                                                                                                                                                                                                                                                 |
| fabric port (F_Port)               | <p>FC port on an FC switch or an FCF that connects point-to-point to an FC node port (N_Port) on an FC host (server or storage device). An F_Port provides access to fabric services for FC devices.</p> <p>F_Ports are intermediate ports in a connection between FC device end-point N_Ports. For example, a connection between an FC host server and an FC storage device through an FC switch looks like this: FC server N_Port to FC switch ingress F_Port to FC switch egress F_Port to FC storage device N_Port.</p> <p>See also <i>node port (N_Port)</i>.</p>                                                                                                             |
| fabric-provided MAC address (FPMA) | <p>MAC address that an FCF assigns to a single ENode MAC through the FLOGI or FDISC process that is unique to the local fabric. The FPMA uniquely identifies a single VN_Port at that ENode MAC in FCoE transactions with the FCF.</p> <p>Because an ENode can have more than one ENode MAC, an FCF can assign multiple FPMAs to an ENode, one FPMA per ENode MAC.</p> <p>An FPMA is a 48-bit value that consists of two 24-bit values, the N_Port ID and the FC-MAP value. The N_Port ID uniquely identifies the VN_Port and the FC-MAP value identifies the FCF.</p> <p>See also <i>FCoE node (ENode)</i>, <i>N_Port ID</i>, and <i>FCoE mapped address prefix (FC-MAP)</i>.</p> |
| FCF-MAC                            | Lossless Ethernet MAC paired with an FCoE controller in an FCF. The FCF-MAC enables the FCF to handle FCoE traffic.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| FCoE controller                    | <p>Instantiates and terminates VN_Port and VF_Port instances on an ENode. An ENode can have more than one FCoE controller. Each FCoE controller is paired with a lossless Ethernet MAC on the ENode.</p> <p>See also <i>lossless Ethernet MAC</i>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                             |
| FC forwarder (FCF)                 | Alternative term and acronym to refer to an FC switch that has all physical Fibre Channel ports and the necessary set of services as defined in the T11 Organization <i>Fibre Channel Switched Fabric (FC-SW)</i> standards.                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| FCoE forwarder (FCF)               | Defined by the <i>Fibre Channel Backbone - 5 (FC-BB-5) Rev 2.00</i> specification available at <a href="http://www.t11.org/ftp/t11/pub/fc/bb-5/09-056v5.pdf">http://www.t11.org/ftp/t11/pub/fc/bb-5/09-056v5.pdf</a> as a device that has the necessary set of services as defined in FC-SW and the FCoE capabilities to act as an FCoE-based FC switch.                                                                                                                                                                                                                                                                                                                           |

Table 511: Fibre Channel Terms (*continued*)

| Term                                | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| FCoE Initialization Protocol (FIP)  | <p>Layer 2 protocol for endpoint discovery, fabric login, and fabric association. FIP enables FCoE devices and FC switches to discover one another. Through FIP, FCoE nodes can log in to an FC switch, access the SAN FC fabric, and communicate with target FC devices. FIP messages also maintain the connection between the FCoE initiator and the FCF.</p> <p>FIP has its own EtherType (0x8914) to distinguish its traffic from payload-carrying FCoE traffic and other Ethernet traffic.</p>                             |
| FCoE link endpoint (LEP)            | Virtual FC interface mapped onto a physical Ethernet interface to handle FC frame encapsulation and de-encapsulation and transmission and reception of FC frames encapsulated in Ethernet through a single virtual link.                                                                                                                                                                                                                                                                                                        |
| FCoE mapped address prefix (FC-MAP) | <p>24-bit value that identifies the FC switch and is half of the 48-bit FPMA MAC address. The FC-MAP value can be configured on the FC switch and has a default value of 0EFC00h. The FC-MAP value was originally called the Fibre Channel Organizationally Unique Identifier (FC-OUI).</p> <p>See also <i>fabric-provided MAC address (FPMA)</i>.</p>                                                                                                                                                                          |
| FCoE node (ENode)                   | <p>Fibre Channel node that has one or more lossless Ethernet MACs, each paired with an <i>FCoE Controller</i> in order to transmit FCoE frames. An ENode combines FCoE termination functions and the FC stack on a CNA. ENodes present virtual FC interfaces to FC switches or FCFs in the form of VN_Ports, which can establish FCoE virtual links with FC switch/FCF VF_Ports. ENodes perform FCoE related functions in a <i>converged network adapter (CNA)</i>.</p> <p>See also <i>converged network adapter (CNA)</i>.</p> |
| FCoE-FC gateway                     | A form of N_Port virtualizer in which the node-facing ports are FCoE ports and the FC switch-facing ports are FC ports.                                                                                                                                                                                                                                                                                                                                                                                                         |
| FCoE-FCoE gateway                   | A form of N_Port virtualizer in which the node-facing ports are FCoE ports and the FC switch-facing ports are FCoE ports.                                                                                                                                                                                                                                                                                                                                                                                                       |
| FC-FC gateway                       | A form of N_Port virtualizer in which the node-facing ports are FC ports and the FC switch-facing ports are FC ports.                                                                                                                                                                                                                                                                                                                                                                                                           |

Table 511: Fibre Channel Terms (*continued*)

| Term                                                      | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-----------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| FCoE transit switch (also known as a FIP snooping bridge) | <p>Switch with a minimum set of features designed to support FCoE Layer 2 forwarding and FCoE security. The switch can also have optional additional features.</p> <p>Minimum feature support is:</p> <ul style="list-style-type: none"> <li>• Priority-based flow control (PFC)</li> <li>• Enhanced transmission selection (ETS)</li> <li>• Data Center Bridging Capability Exchange Protocol (DCBX), including the FCoE application TLV</li> <li>• FIP snooping (minimum support is FIP automated filter programming at the ENode edge)</li> </ul> <p>Additional FIP snooping capabilities can include learning the virtual FC connection paths (VN2VF, VN2VN, or VE2VE) and monitoring the FIP keepalive mechanisms. Other optional capabilities can also enhance FCoE within the standards. FIP snooping is typically configurable on a per-VLAN basis.</p> <p>A transit switch has an FC stack even though it is not an FC switch or an FCF.</p> |
| FCoE VLAN                                                 | VLAN dedicated to carrying only FCoE traffic. FCoE traffic must travel in a VLAN. Only FCoE interfaces should be members of an FCoE VLAN. Ethernet traffic that is not FCoE traffic must travel in a different VLAN.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Fibre Channel                                             | High-speed network technology used for storage area networks (SANs).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Fibre Channel fabric                                      | <p>Network of Fibre Channel devices that allows communication among devices, device name lookup, security, and redundancy.</p> <p>Also a local fabric on a QFX3500 switch with FCoE interfaces connected to FCoE devices on the Ethernet network and native FC interfaces connected to an FC switch in a SAN.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Fibre Channel ID (FCID)                                   | <p>24-bit value the FC switch assigns to the N_Port or VN_Port as a unique identifier within the local FC network. The FCID consists of an 8-bit domain value, an 8-bit area value, and an 8-bit port value. The FCID is sometimes called an N_Port ID.</p> <p>See also <i>N_Port ID</i>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Fibre Channel over Ethernet (FCoE)                        | <p>Standard for transporting FC frames over Ethernet networks. FCoE encapsulates Fibre Channel frames in Ethernet so that the same high-speed Ethernet physical infrastructure can transport both data and storage traffic while preserving the lossless CoS that FC requires. FCoE has its own EtherType (0x8906) to differentiate it from other Ethernet traffic.</p> <p>FCoE runs on a DCB network. FCoE servers connect to a switch that supports both FCoE and native FC protocols. This allows FCoE servers on the Ethernet network to access FC storage devices in the SAN fabric on one converged network.</p> <p>See also <i>data center bridging (DCB)</i>.</p>                                                                                                                                                                                                                                                                             |

Table 511: Fibre Channel Terms (*continued*)

| Term                        | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Fibre Channel services      | Functions required for establishing FC network connectivity among devices and for managing devices on the FC network, such as login servers, domain managers, name servers, and zone servers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| FC stack                    | <p>FC or FCoE protocol capability implemented on a device to support the FC or FCoE functionality. Having an FC stack does not imply consuming a domain ID.</p> <p>Each FC or FCoE enabled server or storage device has an FC stack. Similarly, an FC or FCoE switch, an FCF, an FCoE-FC gateway, and an FCoE transit switch have FC stacks.</p>                                                                                                                                                                                                                                                                                                                                                                           |
| Fibre Channel switch        | Network switch that implements the Fibre Channel protocol.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| FIP discovery advertisement | <p>Multicast or unicast message that the FC switch (or FCF) transmits to ENodes to advertise the switch's presence on the network so that ENodes can discover the switch and request to log in to the FC fabric.</p> <p>The FC switch periodically sends multicast FIP discovery advertisements to the ALL-ENode-MACs address, a well-known address to which all ENodes listen. The multicast messages advertise the FC switch to all ENodes on the VLAN and serve as keepalive messages to maintain connectivity between the FC switch and ENodes.</p> <p>When an ENode sends a FIP discovery solicitation message to the FC switch, the FC switch responds with a unicast FIP discovery advertisement to that ENode.</p> |
| FIP discovery solicitation  | <p>Multicast or unicast message that an ENode transmits to FC switches (or FCFs) to find compatible switches in the network.</p> <p>When an ENode initializes, it sends a multicast FIP discovery solicitation to the ALL-FCF-MACs address, a well-known address to which all FC switches and FCFs listen. Compatible switches reply with a unicast FIP discovery advertisement.</p> <p>The ENode compiles a list of compatible switches, selects a switch, and logs in to that switch.</p>                                                                                                                                                                                                                                |
| FIP keepalive               | Periodic multicast FIP discovery advertisement sent from the FC switch or FCF to all ENodes to maintain connectivity.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

Table 511: Fibre Channel Terms (*continued*)

| Term                   | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| FIP snooping           | <p>For VN_Port to VF_port (VN2VF) paths, FIP snooping is a security feature enabled for FCoE VLANs on an Ethernet switch that connects ENodes to FC switches or FCFs. FIP snooping inspects data in FIP frames and uses that data to create firewall filters. The filters permit only traffic from sources that perform a successful FLOGI to the FC switch. All other traffic on the VLAN is denied. FIP snooping filters are installed on the ports in the FCoE VLAN.</p> <p>FIP snooping also applies similarly for VN_Port to VN_Port (VN2VN) and VE_Port to VE_Port (VE2VE) paths.</p> <p>FIP snooping can also snoop to provide additional visibility of FCoE Layer 2 operation.</p> <p>See also <i>FCoE node (ENode)</i>.</p> |
| FIP snooping bridge    | See <i>FCoE transit switch</i> and <i>FIP snooping</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| host bus adapter (HBA) | Physical mechanism that connects a host system to other FC network and storage devices. HBAs have a unique worldwide node name (WWNN) for the HBA node, which all of the ports on the HBA share, and each port on an HBA has a unique worldwide port name (WWPN).                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| initiator              | System component that originates an I/O command over an I/O bus or network. An FCoE server sending a request to an FC storage device is an example of an initiator.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| iSCSI transit switch   | <p>Layer 2 Ethernet switch with a minimum set of best-practice Ethernet features to support iSCSI, along with optional enhancements. Minimum feature support is:</p> <ul style="list-style-type: none"> <li>• IEEE 802.3X asymmetric and symmetric flow control on ports not running in DCB mode</li> <li>• Priority-based flow control (PFC)</li> <li>• Enhanced transmission selection (ETS)</li> <li>• Data Center Bridging Capability Exchange Protocol (DCBX), including the iSCSI application TLV</li> </ul> <p>Other capabilities such as Internet storage name service (iSNS) are optional.</p>                                                                                                                              |
| Interswitch link (ISL) | The link between the <i>E_Ports</i> of two FC switches in a common FC fabric. When two FCoE-based FC switches are connected together, there is a virtual ISL through Layer 2.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| logout (LOGO)          | <p>For FC devices, an N_Port logs out from the FC network by sending an FC LOGO message to the F_Port of an FC switch. The switch can also send a LOGO message to an N_Port to terminate its connection.</p> <p>For FCoE devices, a VN_Port logs out from the FC network by sending a FIP LOGO message to the VF_Port of an FC switch. The switch can also send a LOGO message to a VN_Port to terminate its connection.</p>                                                                                                                                                                                                                                                                                                         |



Table 511: Fibre Channel Terms (*continued*)

| Term                            | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| lossless Ethernet MAC           | <p>Full-duplex Ethernet MAC that implements Ethernet extensions to avoid Ethernet frame loss due to congestion and supports at least 2.5-KB jumbo frames. Each lossless Ethernet MAC combines with an FCoE Controller to perform FCoE termination functions on an ENode.</p> <p>See also <i>priority-based flow control (PFC)</i>, <i>quantized congestion notification (QCN)</i>, <i>FCoE controller</i>, and <i>FCoE node (ENode)</i>.</p>                                                                                                                                                                                                                                                                                                                 |
| lossless Ethernet network       | Ethernet network composed of only full-duplex links and lossless Ethernet MACs and with CoS and flow control to prevent dropping of frames.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| lossless transport              | In DCB networks, the ability to switch FCoE frames over an Ethernet network without dropping any frames. Lossless transports uses mechanisms such as priority-based flow control and quantized congestion notification to control traffic flows and avoid congestion.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| N_Port ID                       | See <i>Fibre Channel ID (FCID)</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| N_Port ID virtualizer           | <p>Presents itself as an FC or FCoE switch to external devices, but connects to an actual FC or FCoE switch in the other direction to provide the FC-SW services.</p> <p>An N_Port ID virtualizer logs in to the actual FC or FCoE switch in the same way as a normal node device and uses the NPIV mechanism to proxy incoming FLOGIs to FDISCs on the actual FC or FCoE switch.</p> <p>An N_Port ID virtualizer has an FC stack even though it is not an FC switch or an FCF.</p> <p>The acronym <i>NPV</i> is commonly used for N_Port ID virtualizer even though the acronym is not defined in the standards.</p>                                                                                                                                        |
| N_Port ID Virtualization (NPIV) | <p>NPIV enables a physical N_Port to acquire multiple N_Port IDs. Each N_Port ID maps to a different application (such as a virtual machine) or to a different user. This allows you to associate one F_Port with many N_Port IDs and create multiple discrete, secure virtual links over one physical point-to-point connection.</p> <p>NPIV increases resource and bandwidth utilization and allows the implementation of access control, zoning, and port security on a per-application or per-user basis.</p> <p>After an N_Port performs a FLOGI and receives its first N_Port ID, it can request more N_Port IDs by sending FDISC messages.</p> <p>See also <i>fabric login (FLOGI)</i>, <i>fabric discovery (FDISC)</i>, and <i>virtual link</i>.</p> |

Table 511: Fibre Channel Terms (*continued*)

| Term                              | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| node port (N_Port)                | <p>N_Ports can be in two modes:</p> <ul style="list-style-type: none"> <li>Fabric N_Port—Node port that is an FC host or storage device end port in a point-to-point link between the device and the F_Port of an FC switch. The point-to-point link can be virtual or physical.</li> <li>Point-to-point N_Port—Node port that connects to another N_Port. The QFX3500 switch does not support this configuration.</li> </ul> <p>N_Ports handle creation, detection, and flow of messages to and from the connected devices.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| node worldwide name (NWWN)        | WWN that is unique worldwide and is assigned to an FC node. An NWWN is valid for on multiple ports that are on that node (this identifies the ports as network interfaces of a particular node).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| port mode                         | <p>Role that the port plays in the FC fabric (endpoint device, FC switch connection to endpoint devices, interswitch link).</p> <p>See also <i>node port (N_Port)</i>, <i>virtual node port (VN_Port)</i>, <i>proxy node port (NP_Port)</i>, <i>fabric port (F_Port)</i>, and <i>virtual fabric port (VF_Port)</i>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| port worldwide name (PWWN)        | WWN that is unique worldwide and is assigned to an FC port.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| priority-based flow control (PFC) | <p>Link-level flow control mechanism defined by IEEE 802.1Qbb that allows independent flow control for each class of service (as defined in the 3-bit CoS field of the Ethernet header by IEEE 802.1Q tags) to ensure that no frame loss from congestion occurs in DCB networks.</p> <p>PFC is an enhancement of the Ethernet PAUSE mechanism, but PFC controls classes of flows, whereas Ethernet PAUSE indiscriminately pauses all of the traffic on a link. With PFC, a receiving device can signal a transmitting device to pause transmission based on traffic class.</p> <p>PFC provides application-specific bandwidth reservations so you can ensure that time-critical protocols and applications such as FCoE receive the priority necessary to prevent frame loss. PFC allows the same physical link to carry FCoE traffic and provide lossless service while also carrying loss-tolerant Ethernet traffic.</p> <p>See also <i>Ethernet PAUSE</i>.</p> |
| proxy gateway mode                | Connects FCoE initiators to FC switches in a converged Ethernet and Fibre Channel network and acts as an intermediary for these devices. The FCoE-FC gateway represents and acts for the FCoE initiators in transactions from the FCoE initiators destined for an FC switch, including converting FIP and FCoE frames to FC frames. The gateway represents and acts for an FC switch in transactions from the FC switch destined for an FCoE initiator, including converting FC frames to FIP frames and encapsulating FC frames in Ethernet.                                                                                                                                                                                                                                                                                                                                                                                                                     |
| proxy node port (NP_Port)         | N_Port on the QFX Series that performs proxy functions when it is configured as an FCoE-FC gateway. The NP_Port acts as a proxy for the FCoE device VN_Ports in transactions with the FC switch.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

Table 511: Fibre Channel Terms (*continued*)

| Term                                    | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| quantized congestion notification (QCN) | Mechanism defined by IEEE 802.1Qau that manages network congestion within a Layer 2 domain. When a queue reaches a configured threshold, QCN throttles traffic at the source of the congestion by transmitting messages that propagate back to the source and temporarily stop the source from transmitting. When the queue crosses the threshold that indicates the congestion has dissipated, QCN sends a message to allow the source to resume transmitting frames. |
| session                                 | Fabric login (FLOGI) or fabric discovery (FDISC) login to the FC SAN fabric. Session does not refer to end-to-end server-to-storage sessions.                                                                                                                                                                                                                                                                                                                          |
| server-provided MAC address (SPMA)      | <p>MAC address that an ENode assigns to one of its ENode MACs and is not assigned to any other ENode MAC in the same FCoE VLAN. An SPMA can be associated with more than one VN_Port at that ENode MAC.</p> <p>The QFX Series does not support SPMA.</p> <p>See also <i>ENode MAC</i> and <i>fabric-provided MAC address (FPMA)</i>.</p>                                                                                                                               |
| storage area network (SAN)              | Network whose primary purpose is the transfer of data between computer systems and storage devices. This term is most commonly used in the context of any network that supports block storage, usually iSCSI, FC, and FCoE networks.                                                                                                                                                                                                                                   |
| target                                  | System component that receives an I/O command. An FC storage device that receives a request from a server is an example of a target.                                                                                                                                                                                                                                                                                                                                   |
| VE_Port                                 | Virtual ports created to form a connection (an <i>interswitch link</i> ) between two FCoE-based FC switches as part of a common FC fabric.                                                                                                                                                                                                                                                                                                                             |
| VE2VE (VE_Port to VE_Port)              | The <i>Fibre Channel Backbone - 5 (FC-BB-5) Rev 2.00</i> specification capability of FCFs to connect to each other as a single FCoE FC SAN.                                                                                                                                                                                                                                                                                                                            |
| VN2VF (VN_Port to VF_Port)              | The <i>Fibre Channel Backbone - 5 (FC-BB-5) Rev 2.00</i> specification capability of an ENode to connect to an FCF or to an FCoE-enabled FC SAN.                                                                                                                                                                                                                                                                                                                       |
| VN2VN (VN_Port to VN_Port)              | The <i>Fibre Channel Backbone - 6 (FC-BB-6)</i> specification capability of an ENode to connect directly over Layer 2 to another ENode without the need of any FC-related services. This capability is most often used in small-scale FCoE SANs.                                                                                                                                                                                                                       |
| virtual fabric port (VF_Port)           | <p>Data-forwarding component that emulates an F_Port. A VF_Port is dynamically instantiated on successful completion of a FIP FLOGI exchange and connects to one or more VN_Ports. The term <i>virtual</i> indicates the use of a non-FC link such as an FCoE link.</p> <p>See also <i>fabric port (F_Port)</i>.</p>                                                                                                                                                   |
| virtual link                            | <p>Logical link connecting two FCoE Link End Points (LEPs) over a lossless Ethernet network, for example, the link between a VF_Port and a VN_Port. The MAC addresses of the two LEPs identifies a virtual link.</p> <p>See also <i>FCoE link end point (LEP)</i> and <i>lossless Ethernet network</i>.</p>                                                                                                                                                            |

Table 511: Fibre Channel Terms (*continued*)

| Term                        | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| virtual node port (VN_Port) | <p>Data-forwarding component that emulates an N_Port. With FCoE, a VN_Port is dynamically instantiated on successful completion of a FIP FLOGI exchange and connects to one or more VF_Ports. The term <i>virtual</i> indicates the use of a non-FC link such as an FCoE link.</p> <p>VN_Port is also used for the virtual N_Ports created in both FC and FCoE when additional NPIV-based logins occur over a previously created N_Port-to-VN_Port or N_Port-to-VF_Port connection.</p> <p>See also <i>node port (N_Port)</i>.</p> |
| well-known address (WKA)    | <p>Address identifier used to access a service provided by an FC fabric. The service can be distributed in many elements throughout a fabric, or it can be centralized in one element. A WKA is always accessible, regardless of zoning. An example of a WKA is the <i>ALL-FCF-MACs</i> address to which all FCFs listen.</p>                                                                                                                                                                                                      |
| worldwide name (WWN)        | <p>64-bit identifier that is similar to a MAC address except that it is not used for forwarding. It uniquely identifies an FC device. The WWN is derived from the IEEE organizationally unique identifier (OUI) and vendor-supplied information. A WWN is unique worldwide.</p>                                                                                                                                                                                                                                                    |
| worldwide node name (WWNN)  | <p>See <i>node worldwide name (NWWN)</i>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| worldwide port name (WWPN)  | <p>See <i>port worldwide name (PWWN)</i>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

**Related Documentation**

- [Overview of Fibre Channel](#)

## PART 87

# Configuring FCoE on a Transit Switch

- [Using FCoE on a Transit Switch on page 6431](#)



# Using FCoE on a Transit Switch

- [Understanding FCoE Transit Switch Functionality on page 6431](#)
- [Understanding FCoE on page 6436](#)
- [Configuring VLANs for FCoE Traffic on an FCoE Transit Switch on page 6442](#)
- [Troubleshooting Dropped FCoE Traffic on page 6447](#)
- [Understanding MC-LAGs on an FCoE Transit Switch on page 6451](#)
- [Example: Configuring CoS for FCoE Transit Switch Traffic Across an MC-LAG on page 6454](#)
- [Understanding FCoE and FIP Session High Availability on page 6480](#)
- [Troubleshooting Dropped FIP Traffic on page 6482](#)
- [Understanding OxID Hash Control for FCoE Traffic Load Balancing on Standalone Switches on page 6484](#)
- [Enabling and Disabling CoS OxID Hash Control on Standalone Switches on page 6485](#)

## Understanding FCoE Transit Switch Functionality

---

You can use the switch as a Fibre Channel over Ethernet (FCoE) transit switch. An FCoE transit switch is a Layer 2 data center bridging (DCB) switch that can transport FCoE frames, and when used as an access switch for FCoE devices, implements FCoE Initialization Protocol (FIP) snooping. A DCB switch transports both FCoE and Ethernet LAN traffic over the same network infrastructure while preserving the class of service (CoS) treatment that Fibre Channel (FC) traffic requires.



**NOTE:** Juniper Networks QFX10000 switches do not support FIP snooping. (Aggregation devices do not need to enable FIP snooping because FIP snooping is performed at the FCoE access edge.)

An FCoE transit switch does not encapsulate or de-encapsulate FC frames in Ethernet. It is a switch that transports FC frames that have already been encapsulated in Ethernet between FCoE initiators such as servers and a storage area network (SAN) FC switch that supports both Ethernet and native FC traffic on its interfaces. The transit switch acts as a passthrough switch and is transparent to the FC switch, which detects each connection to an FCoE device as a direct point-to-point link.

When a switch acts as a transit switch, the VLANs you configure for FCoE traffic can use any of the switch ports because the traffic in both directions is standard Ethernet traffic, not native FC traffic.



**NOTE:** The Ethernet interfaces that connect to FCoE devices must include a native VLAN to transport FIP traffic, because FIP VLAN discovery and notification frames are exchanged as untagged packets. It is a good practice to keep the native VLAN separate from the VLANs that carry FCoE traffic. FCoE VLANs should carry only FCoE traffic, but other types of untagged traffic might use the native VLAN.

Switches and QFabric system Node devices that use the original CLI (not the Enhanced Layer 2 (ELS) software) only require that you configure the native VLAN on the FCoE interfaces that belong to the FCoE VLAN by including the `[set interfaces interface-name unit unit family ethernet-switching native-vlan-id native-vlan-id]` statement in the configuration.

Switches that use ELS software require that you include two statements in the configuration to configure a native VLAN on FCoE interfaces. Include the `[set interfaces interface-name native-vlan-id vlan-id]` statement in the configuration to configure the native VLAN on the interface, and also include the `[set interfaces interface-name unit unit family ethernet-switching native-vlan-id vlan-id]` statement in the configuration to configure the port as a member of the native VLAN.

FCoE traffic should use a VLAN dedicated only to FCoE traffic. Do not mix FCoE traffic with standard Ethernet traffic on the same VLAN.





**NOTE:** An FCoE VLAN (any VLAN that carries FCoE traffic) supports only Spanning Tree Protocol (STP) and link aggregation group (LAG) Layer 2 features.

FCoE traffic cannot use a standard LAG because traffic might be hashed to different physical LAG links on different transmissions. This breaks the (virtual) point-to-point link that Fibre Channel traffic requires. If you configure a standard LAG interface for FCoE traffic, FCoE traffic might be rejected by the FC SAN.

QFabric systems support a special LAG called an FCoE LAG, which enables you to transport FCoE traffic and regular Ethernet traffic (traffic that is not FCoE traffic) across the same link aggregation bundle. Standard LAGs use a hashing algorithm to determine which physical link in the LAG is used for a transmission, so communication between two devices might use different physical links in the LAG for different transmissions. An FCoE LAG ensures that FCoE traffic uses the same physical link in the LAG for requests and replies in order to preserve the virtual point-to-point link between the FCoE device converged network adapter (CNA) and the FC SAN switch across the QFabric system Node device. An FCoE LAG does not provide load balancing or link redundancy for FCoE traffic. However, regular Ethernet traffic uses the standard hashing algorithm and receives the usual LAG benefits of load balancing and link redundancy in an FCoE LAG.



**NOTE:** IGMP snooping is enabled by default on all VLANs in all software versions before Junos OS Release 13.2. Disable IGMP snooping on FCoE VLANs if you are using software that is older than Junos OS Release 13.2.



**NOTE:** On a QFX3500 switch or on a QFabric system Node device, the same VLAN cannot be used in both transit switch mode and FCoE-FC gateway mode. (Only QFX3500 switches can be configured in FCoE-FC gateway mode.) If you configure both a transit switch and an FCoE-FC gateway on the same QFX3500 switch or QFabric system Node device, configure different FCoE VLANs for the transit switch and the FCoE-FC gateway.

To support FCoE traffic, transit switches require DCB configuration to implement the lossless transport of FCoE traffic across the Ethernet portion of the network, and transit switches on the access edge require enabling FIP snooping on the FCoE access ports.

With the exception of Virtual Chassis (VC) and mixed-mode Virtual Chassis Fabric (VCF) configurations, switches support the DCB standards for ensuring lossless transport and low latency, and provide 10-Gbps ports for FCoE traffic. VCF configurations that use only QFX5100 switches support DCB standards. For lossless transport to function correctly, you must use priority-based flow control (PFC, described in IEEE 802.1Qbb) to prevent FCoE packet loss during periods of congestion and ensure proper CoS for FCoE traffic.

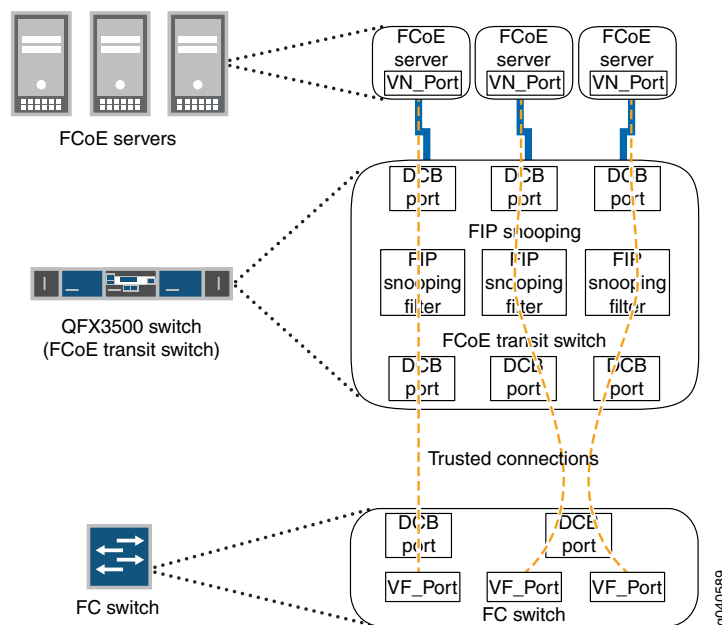
On the FCoE access edge, FIP snooping adds security by filtering access. Only traffic from servers that have successfully logged in to the FC network can pass through the transit switch and reach the FC network. The Technical Committee T11 organization specifications describe two types of FIP snooping:

- The FC-BB-5 specification describes VN\_Port to VF\_Port (VN2VF\_Port) FIP snooping, which provides security for communication between FCoE device VN\_Ports on the Ethernet network and FCF or FC switch VF\_Ports.
- The FC-BB-6 specification describes VN\_Port to VN\_Port (VN2VN\_Port) FIP snooping, which provides security for communication between FCoE device VN\_Ports on the Ethernet network.

To accommodate the larger size of Ethernet-encapsulated frames, FCoE interfaces should be configured with a maximum transmission unit (MTU) size of at least 2180 bytes.

At the access edge, a transit switch transparently connects FCoE-capable devices such as servers in an Ethernet LAN to an FC switch or to a gateway switch (hereafter referred to as the FC switch), as shown in [Figure 203](#). The transit switch acts as a transparent DCB access layer between FCoE servers and the FC switch.

**Figure 203: FCoE Transit Switch Connecting FCoE Devices to an FC Switch**



The transit switch performs FIP snooping at the ports connected to the FCoE devices. For VN2VF\_Port FIP snooping, at the SAN edge, the FC switch must be able to convert the FCoE traffic to native FC traffic. (VN2VN\_Port FIP snooping switches traffic between VN\_Ports directly through the transit switch, without going through the FC switch, so no conversion of FCoE traffic to native FC traffic is needed.)

Encapsulated FCoE traffic flows through the transit switch to the FCoE ports on the FC switch. The FC switch removes the Ethernet encapsulation from the FCoE frames to

restore the native FC frames. Native FC traffic travels out native FC ports to storage devices in the FC SAN.

Native FC traffic from storage devices flows to the FC switch FC ports, and the FC switch encapsulates that traffic in Ethernet as FCoE traffic. The FCoE traffic flows through the transit switch to the appropriate FCoE device.



**NOTE:** The FC switch and FC fabric apply appropriate zoning checks on traffic to and from each ENode and provide FC services (for example, name server, fabric login server, or event server).



**NOTE:** VN\_Port to VN\_Port FIP snooping is supported to allow FCoE initiators and targets to communicate directly through the switch without going through an FCoE forwarder (FCF) or an FC switch. An FCoE VLAN can support either VN2VF\_Port FIP snooping (FC-BB-5) or VN2VN\_Port FIP snooping (FC-BB-6), but not both. The same switch can have multiple FCoE VLANs configured, some FCoE VLANs for VN2VF FIP snooping traffic and others for VN2VN FIP snooping traffic.

Transit switches do not have to be FCoE access edge switches. Transit switches can be intermediate switches between a transit switch at the FCoE access edge and the FC switch. In this case, intermediate transit switches do not need to perform FIP snooping because only the access edge transit switch needs to filter traffic between the FCoE device and the FC network. Once that traffic has been processed by the FIP snooping filters, it does not need to be filtered again. However, intermediate transit switches must support DCB standards to preserve the lossless transport and other CoS characteristics required for FC traffic.

**Related  
Documentation**

- [Understanding DCB Features and Requirements on page 6490](#)
- [Understanding FCoE on page 6436](#)
- [Configuring VLANs for FCoE Traffic on an FCoE Transit Switch on page 6442](#)

## Understanding FCoE

---

Fibre Channel over Ethernet (FCoE) is a method of supporting converged Fibre Channel (FC) and Ethernet traffic on a data center bridging (DCB) network. FCoE encapsulates unmodified FC frames in Ethernet to transport the FC frames over a physical Ethernet network. The T11 Technical Committee, which is the International Committee for Information Technology Standards (INCITS) committee responsible for FC interfaces, developed the FCoE standard to provide a method for transporting FC frames over a DCB network. The T11 document *Fibre Channel Backbone - 5 (FC-BB-5) Rev 2.00* at <http://www.t11.org/ftp/t11/pub/fc/bb-5/09-056v5.pdf> provides details about the FCoE version 1 standard.



**NOTE:** The switch does not support T11 Annex F *FCoE Pre-FIP Virtual Link Instantiation Protocol*.

To the Ethernet network, an FCoE frame is the same as any other Ethernet frame because the Ethernet encapsulation provides the header information needed to forward the frames. However, to achieve the lossless behavior that FC transport requires, the Ethernet network must conform to DCB standards.

DCB standards create an environment over which FCoE can transport native FC traffic encapsulated in Ethernet while preserving the mandatory class of service (CoS) and other characteristics that FC traffic requires.

Supporting FCoE in a DCB network requires that the FCoE devices in the Ethernet network and the FC switches at the edge of the SAN network handle both Ethernet and native FC traffic. To handle Ethernet traffic, an FC switch does one of two things:

- Incorporates FCoE interfaces.
- Uses an FCoE-FC gateway such as a QFX3500 switch to de-encapsulate FCoE traffic from FCoE devices into native FC and to encapsulate native FC traffic from the FC switch into FCoE and forward it to FCoE devices through the Ethernet network.



**NOTE:** Standalone switches support FCoE. Virtual Chassis (VC) and mixed-mode Virtual Chassis Fabric (VCF) configurations do not support FCoE. Pure QFX5100 switch VCFs (consisting of only QFX5100 switches) support FCoE.

FCoE concepts include:

- [FCoE Devices on page 6437](#)
- [FCoE Frames on page 6438](#)
- [Virtual Links on page 6439](#)
- [FCoE VLANs on page 6439](#)

## FCoE Devices

Each FCoE device has a converged network adapter (CNA) that combines the functions of an FC host bus adapter (HBA) and a lossless Ethernet network interface card (NIC) with 10-Gbps Ethernet ports. The portion of the CNA that handles FCoE traffic is called an FCoE Node (ENode). An ENode combines FCoE termination functions and the client part of the FC stack on the CNA.

ENodes present virtual FC interfaces to FC switches in the form of virtual N\_Ports (VN\_Ports). A VN\_Port is an endpoint in a virtual point-to-point connection called a virtual link. The other endpoint of the virtual link is an FC switch (or FCF) port. A VN\_Port emulates a native FC N\_Port and performs similar functions: handling the creation, detection, and flow of messages to and from the FC switch. A single ENode can host multiple VN\_Ports. Each VN\_Port has a separate, unique virtual link with a FC switch.

ENodes contain at least one lossless Ethernet media access controller (MAC). Each Ethernet MAC is paired with an FCoE controller. The lossless Ethernet MAC is a full-duplex Ethernet MAC that implements Ethernet extensions to avoid frame loss due to congestion and supports frames of at least 2500 bytes. The FCoE controller instantiates and terminates VN\_Port instances dynamically as they are needed for FCoE sessions. Each VN\_Port instance has a unique virtual link to an FC switch.



**NOTE:** A *session* is a fabric login (FLOGI) or fabric discovery (FDISC) login to the FC SAN fabric. Session does not refer to end-to-end server-to-storage sessions.


ENodes also contain one FCoE link end point (LEP) for each VN\_Port connection. An FCoE LEP is a virtual FC interface mapped onto the physical Ethernet interface.

An FCoE LEP:

- Transmits and receives FCoE frames on the virtual link.
- Handles FC frame encapsulation for traffic going from the server to the FC switch.
- Performs frame de-encapsulation of traffic received from the FC switch.

Figure 204 shows a block diagram of the major ENode components.

Figure 204: ENode Components



## FCoE Frames

The FCoE protocol specification replaces the FC0 and FC1 layers of the FC stack with Ethernet, but retains the FC frame header. Retaining the FC frame header enables the FC frame to pass directly to a native FC SAN after de-encapsulation. The FCoE header carries the FC start of file (SOF) bits and end of file (EOF) bits in an encoded format. FCoE supports two frame types, control frames and data frames. FCoE Initialization Protocol (FIP) carries all of the discovery and fabric login frames.

FIP control frames handle FCoE device discovery, initializing communication, and maintaining communication. They do not carry a data payload. FIP has its own EtherType (0x8914) to distinguish FIP traffic from FCoE traffic and other Ethernet traffic. To establish communication, the ENode uses the globally unique MAC address assigned to it by the CNA manufacturer.

After FIP establishes a connection between FCoE devices, the FCoE data frames handle the transport of the FC frames encapsulated in Ethernet. FCoE also has its own EtherType (0x8906) to distinguish FCoE frames from other Ethernet traffic and ensure the in-order frame handling that FC requires. FCoE frames include:

- 2112 bytes FC payload
- 24 bytes FC header
- 14 bytes standard Ethernet header
- 14 bytes FCoE header
- 8 bytes cyclic redundancy check (CRC) plus EOF
- 4 bytes VLAN header
- 4 bytes frame check sequence (FCS)

The payload, headers, and checks add up to 2180 bytes. Therefore, interfaces that carry FCoE traffic should have a configured maximum transmission unit (MTU) of 2180 or larger. An MTU size of 2180 bytes is the minimum size; some network administrators prefer an MTU of 2240 or 2500 bytes.

## Virtual Links

Native FC uses point-to-point physical links between FC devices. In FCoE, virtual links replace the physical links. A virtual link emulates a point-to-point link between two FCoE device endpoints, such as a server VN\_Port and an FC switch (or FCF) VF\_Port.

Each FCoE interface can support multiple virtual links. The MAC addresses of the FCoE endpoints (the VN\_Port and the VF\_Port) uniquely identify each virtual link and allow traffic for multiple virtual links to share the same physical link while maintaining data separation and security.

A virtual link exists in one FCoE VLAN and cannot belong to more than one VLAN. Although the FC switch and the FCoE device detect a virtual link as a point-to-point connection, virtual links do not need to be direct connections between a VF\_Port and a VN\_Port. A virtual link can traverse one or more transit switches, also known as passthrough switches. A transit switch can transparently aggregate virtual links while still appearing and functioning as a point-to-point connection to the FCoE devices. However, a virtual link must remain within a single Layer 2 domain.

## FCoE VLANs

All FCoE traffic must travel in a VLAN dedicated to transporting only FCoE traffic. Only FCoE interfaces should be members of an FCoE VLAN. Ethernet traffic that is not FCoE or FIP traffic must travel in a different VLAN.



**NOTE:** On a standalone switch or QFabric system Node device, the same VLAN cannot be used in both transit switch mode and FCoE-FC gateway mode.

---



.....

**NOTE:** FCoE VLANs (any VLAN that carries FCoE traffic) support only Spanning Tree Protocol (STP) and link aggregation group (LAG) Layer 2 features.

FCoE traffic cannot use a standard LAG because traffic might be hashed to different physical LAG links on different transmissions. This breaks the (virtual) point-to-point link that Fibre Channel traffic requires. If you configure a standard LAG interface for FCoE traffic, FCoE traffic might be rejected by the FC SAN.

QFabric systems support a special LAG called an FCoE LAG, which enables you to transport FCoE traffic and regular Ethernet traffic (traffic that is not FCoE traffic) across the same link aggregation bundle. Standard LAGs use a hashing algorithm to determine which physical link in the LAG is used for a transmission, so communication between two devices might use different physical links in the LAG for different transmissions. An FCoE LAG ensures that FCoE traffic uses the same physical link in the LAG for requests and replies in order to preserve the virtual point-to-point link between the FCoE device converged network adapter (CNA) and the FC SAN switch across the QFabric system Node device. An FCoE LAG does not provide load balancing or link redundancy for FCoE traffic. However, regular Ethernet traffic uses the standard hashing algorithm and receives the usual LAG benefits of load balancing and link redundancy in an FCoE LAG.



.....

**NOTE:** IGMP snooping is enabled by default on all VLANs in all software versions before Junos OS R13.2. Disable IGMP snooping on FCoE VLANs if you are using software that is older than 13.2.

.....

You can configure more than one FCoE VLAN, but any given virtual link must be in only one FCoE VLAN.





**NOTE:** All 10-Gigabit Ethernet interfaces that connect to FCoE devices must have a native VLAN configured in order to transport FIP traffic, because FIP VLAN discovery and notification frames are exchanged as untagged packets.

On switches that use the Enhanced Layer 2 Software (ELS) CLI, it is not sufficient only to configure the native VLAN on the interface, the interface must also be configured as a member of the native VLAN. (This is because the ELS CLI does not support tagged-access interface mode, so interfaces that are members of FCoE VLANs must use trunk mode, and trunk port interfaces must be explicitly included as members of a native VLAN.)

In addition, the VLAN ID must match the native VLAN ID that you configure on the physical interface. For example, to configure a native VLAN with an ID of 20 on interface xe-0/0/15 that is a member of an FCoE VLAN, you must include both of the following statements in the configuration:

1. Configure the native VLAN on the interface:

```
user@switch# set interfaces xe-0/0/15 native-vlan-id 20
```

(The equivalent configuration statement on a non-ELS device switch would be `set interfaces xe-0/0/15 unit 0 family ethernet-switching native-vlan-id 20`.)

2. Configure the port as a member of the native VLAN (this step is not required on switches that do not use the ELS software):

```
user@switch# set interfaces xe-0/0/15 unit 0 family ethernet-switching vlan members 20
```



**BEST PRACTICE:** Only FCoE traffic is permitted on the FCoE VLAN. A native VLAN might need to carry untagged traffic of different types and protocols. Therefore, it is a good practice to keep the native VLAN separate from FCoE VLANs.

#### Related Documentation

- [Understanding DCB Features and Requirements on page 6490](#)
- [Understanding FCoE Transit Switch Functionality on page 6431](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 6493](#)
- [Configuring VLANs for FCoE Traffic on an FCoE Transit Switch on page 6442](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 6504](#)

## Configuring VLANs for FCoE Traffic on an FCoE Transit Switch

---

When you configure a switch as a Fibre Channel over Ethernet (FCoE) transit switch, you must configure a VLAN that transports only FCoE traffic. FCoE traffic requires a dedicated VLAN and cannot share a VLAN with any other type of traffic. Because FCoE traffic is tagged traffic, the port (or interface) mode cannot be access mode, it must be either tagged-access port-mode (for switches that run the original CLI) or trunk interface-mode (for switches that run the Enhanced Layer 2 Software (ELS) CLI).

However, each interface that belongs to an FCoE VLAN must not only transport the tagged FCoE traffic, it must also transport the untagged FCoE Initialization Protocol (FIP) traffic. FIP communicates with the storage area network (SAN) Fibre Channel (FC) switch to set up the FCoE session for the FCoE client.

To transport untagged traffic on a tagged-access or trunk mode interface, the interface must have a native VLAN configured on it. Therefore, each interface that belongs to an FCoE VLAN must also have a native VLAN on it.

There are slight differences in the way you configure a native VLAN on an interface, depending on whether the switch uses the ELS CLI or the original CLI. This topic describes both methods.



**NOTE:** FCoE VLANs (any VLAN that carries FCoE traffic) support only Spanning Tree Protocol (STP) and link aggregation group (LAG) Layer 2 features.

FCoE traffic cannot use a standard LAG because traffic might be hashed to different physical LAG links on different transmissions. This breaks the (virtual) point-to-point link that Fibre Channel traffic requires. If you configure a standard LAG interface for FCoE traffic, FCoE traffic might be rejected by the FC SAN.

QFabric systems support a special LAG called an FCoE LAG, which enables you to transport FCoE traffic and regular Ethernet traffic (traffic that is not FCoE traffic) across the same link aggregation bundle. Standard LAGs use a hashing algorithm to determine which physical link in the LAG is used for a transmission, so communication between two devices might use different physical links in the LAG for different transmissions. An FCoE LAG ensures that FCoE traffic uses the same physical link in the LAG for requests and replies in order to preserve the virtual point-to-point link between the FCoE device converged network adapter (CNA) and the FC SAN switch across the QFabric system Node device. An FCoE LAG does not provide load balancing or link redundancy for FCoE traffic. However, regular Ethernet traffic uses the standard hashing algorithm and receives the usual LAG benefits of load balancing and link redundancy in an FCoE LAG.



**NOTE:** To configure an FCoE VLAN on a QFX3500 switch that you are using as an FCoE-FC gateway, you must also configure an FCoE VLAN interface as described in *Configuring an FCoE VLAN Interface on an FCoE-FC Gateway*. (Only the QFX3500 switch supports FCoE-FC gateway configuration.)

FCoE VLAN configuration includes:

- Configuring a VLAN to use as a dedicated FCoE VLAN
- Configuring the interface members of the FCoE VLAN.
- Configuring a native VLAN for FIP traffic.

This topic includes two configuration procedures, one for switches that run the original CLI, and one for switches that run the ELS CLI.

## Original CLI Configuration

To configure an FCoE VLAN on a non-ELS switch:

1. Configure a dedicated FCoE VLAN:

```
[edit vlans]
user@switch# set vlan-name vlan-id vlan-id
```

For example, to configure a VLAN named **fcoe\_vlan** with a VLAN ID of **100** as the FCoE VLAN:

```
[edit vlans]
user@switch# set fcoe_vlan vlan-id 100
```

2. Configure the FCoE VLAN on the interface (use **ethernet-switching** as the family and **tagged-access** as the port mode):

```
[edit interfaces]
user@switch# set interface-name unit unit family family port-mode mode vlan members
vlan-name
```

For example, to configure the interface **xe-0/0/10** as a member of the FCoE VLAN **fcoe\_vlan**:

```
[edit interfaces]
user@switch# set xe-0/0/10 unit 0 family ethernet-switching port-mode tagged-access vlan
members fcoe_vlan
```

3. Configure the Ethernet interface membership in the FCoE VLAN:

```
[edit vlans]
user@switch# set vlan-name interface interface-name
```

For example, to assign the interface **xe-0/0/10.0** to the FCoE VLAN named **fcoe\_vlan**:

```
[edit vlans]
user@switch# set fcoe_vlan interface xe-0/0/10.0
```

4. Configure a native VLAN for the untagged FIP traffic:

```
[edit vlans]
user@switch# set native vlan-id vlan-id
```

For example, to configure the native VLAN with a VLAN ID of 1:

```
[edit vlans]
user@switch# set native vlan-id 1
```

5. Assign member interfaces to the native VLAN:

```
[edit interfaces]
user@switch# set interface-name unit unit family family native-vlan-id vlan-id
```

For example, to configure the interface **xe-0/0/10** as a member of the native VLAN with the native VLAN ID 1:

```
[edit interfaces]
user@switch# set xe-0/0/10 unit 0 family ethernet-switching native-vlan-id 1
```

## ELS CLI Configuration

To configure an FCoE VLAN on a switch running ELS:

1. Configure a dedicated FCoE VLAN:

```
[edit vlans]
user@switch# set vlan-name vlan-id vlan-id
```

For example, to configure a VLAN named **fcoe\_vlan** with a VLAN ID of **100** as the FCoE VLAN:

```
[edit vlans]
user@switch# set fcoe_vlan vlan-id 100
```

2. Configure the FCoE VLAN on the interface (use **ethernet-switching** as the family and **trunk** as the interface mode):

```
[edit interfaces]
user@switch# set interface-name unit unit family family interface-mode mode vlan members
vlan-name
```

For example, to configure the interface **xe-0/0/10** as a member of the FCoE VLAN **fcoe\_vlan**:

```
[edit interfaces]
user@switch# set xe-0/0/10 unit 0 family ethernet-switching interface-mode trunk vlan
members fcoe_vlan
```

3. Configure the Ethernet interface membership in the FCoE VLAN:

```
[edit vlans]
user@switch# set vlan-name interface interface-name
```

For example, to assign the interface **xe-0/0/10.0** to the FCoE VLAN named **fcoe\_vlan**:

```
[edit vlans]
user@switch# set fcoe_vlan interface xe-0/0/10.0
```

4. Configure a native VLAN on the physical Ethernet interface for the untagged FIP traffic:

```
[edit interfaces]
user@switch# set interface-name native-vlan-id vlan-id
```

For example, to configure the native VLAN on interface **xe-0/0/10** with a VLAN ID of **1**:

```
[edit interfaces]
user@switch# set xe-0/0/10 native-vlan-id 1
```

5. Configure the Ethernet interface as a member of the native VLAN:

```
[edit interfaces]
user@switch# set interface-name unit unit family family vlan members native-vlan-id
```



**NOTE:** The *native-vlan-id* number must be the same as the native VLAN ID number that you configured on the physical Ethernet interface (see step 4).

---

For example, to configure the interface **xe-0/0/10** as a member of the native VLAN with the native VLAN ID **1**:

```
[edit interfaces]
```

```
user@switch# set xe-0/0/10 unit 0 family ethernet-switching vlan members 1
```

#### Related Documentation

- [Understanding FCoE on page 6436](#)
- [Understanding FCoE Transit Switch Functionality on page 6431](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 6504](#)

## Troubleshooting Dropped FCoE Traffic

**Problem** **Description:** Fibre Channel over Ethernet (FCoE) traffic for which you want guaranteed delivery is dropped.

**Cause** There are several possible causes of dropped FCoE traffic (the list numbers of the possible causes correspond to the list numbers of the solutions in the *Solution* section.):

1. Priority-based flow control (PFC) is not enabled on the FCoE priority (IEEE 802.1p code point) in both the input and output stanzas of the congestion notification profile.
2. The FCoE traffic is not classified correctly at the ingress interface. FCoE traffic should either use the default **fcoe** forwarding class and classifier configuration (maps the **fcoe** forwarding class to IEEE 802.1p code point 011) or be mapped to a lossless forwarding class and to the code point enabled for PFC on the input and output interfaces.
3. The congestion notification profile that enables PFC on the FCoE priority is not attached to the interface.
4. The forwarding class set (priority group) used for guaranteed delivery traffic does not include the forwarding class used for FCoE traffic.



**NOTE:** This issue can occur only on switches that support enhanced transmission selection (ETS) hierarchical port scheduling. (Direct port scheduling does not use forwarding class sets.)

5. Insufficient bandwidth has been allocated for the FCoE queue or for the forwarding class set to which the FCoE queue belongs.



**NOTE:** This issue can occur for forwarding class sets only on switches that support ETS hierarchical port scheduling. (Direct port scheduling does not use forwarding class sets.)

6. If you are using Junos OS Release 12.2, the **fcoe** forwarding class has been explicitly configured instead of using the default **fcoe** forwarding class configuration (forwarding-class-to-queue mapping).



.....

**NOTE:** If you are using Junos OS Release 12.2, use the default forwarding-class-to-queue mapping for the lossless fcoe and no-loss forwarding classes. If you explicitly configure the lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best effort) traffic and does *not* receive lossless treatment.

.....

7. If you are using Junos OS Release 12.3 or later and you are not using the default **fcoe** forwarding class configuration, the forwarding class used for FCoE is not configured with the **no-loss** packet drop attribute. In Junos OS 12.3 or later, explicit forwarding classes configurations must include the **no-loss** packet drop attribute to be treated as lossless forwarding classes.

**Solution** The list numbers of the possible solutions correspond to the list numbers of the causes in the *Cause* section.



1. Check the congestion notification profile (CNP) to see if PFC is enabled on the FCoE priority (the correct IEEE 802.1p code point) on both input and output interfaces. Use the **show class-of-service congestion-notification** operational command to show the code points that are enabled for PFC in each CNP.

If you are using the default configuration, FCoE traffic is mapped to code point 011 (priority 3). In this case, the input stanza of the CNP should show that PFC is enabled on code point 011, and the output stanza should show that priority 011 is mapped to flow control queue 3.

If you explicitly configured a forwarding class for FCoE traffic, ensure that:

- You specified the **no-loss** packet drop attribute in the forwarding class configuration
- The code point mapped to the FCoE forwarding class in the ingress classifier is the code point enabled for PFC in the CNP input stanza
- The code point and output queue used for FCoE traffic are mapped to each other in the CNP output stanza (if you are not using the default priority and queue, you must explicitly configure each output queue that you want to respond to PFC messages)

For example, if you explicitly configure a forwarding class for FCoE traffic that is mapped to output queue 5 and to code point 101 (priority 5), the output of the **show class-of-service congestion-notification** looks like:

```
Name: fcoe_p5_cnp, Index: 12183
Type: Input
Cable Length: 100 m
 Priority PFC MRU
 000 Disabled
 001 Disabled
 010 Disabled
 011 Disabled
 100 Disabled
 101 Enabled 2500
 110 Disabled
 111 Disabled
Type: Output
 Priority Flow-Control-Queues
 101 5
```

2. Use the **show class-of-service classifier type ieee-802.1p** operational command to check if the classifier maps the forwarding class used for FCoE traffic to the correct IEEE 802.1p code point.
3. Ensure that the congestion notification profile and classifier are attached to the correct ingress interface. Use the operational command **show configuration class-of-service interfaces interface-name**.
4. Check that the forwarding class set includes the forwarding class used for FCoE traffic. Use the operational command **show configuration class-of-service forwarding-class-sets** to show the configured priority groups and their forwarding classes.

5. Verify the amount of bandwidth allocated to the queue mapped to the FCoE forwarding class and to the forwarding class set to which the FCoE traffic queue belongs. Use the **show configuration class-of-service schedulers *scheduler-name*** operational command (specify the scheduler for FCoE traffic as the *scheduler-name*) to see the minimum guaranteed bandwidth (**transmit-rate**) and maximum bandwidth (**shaping-rate**) for the queue.

Use the **show configuration class-of-service traffic-control-profiles *traffic-control-profile*** operational command (specify the traffic control profile used for FCoE traffic as the *traffic-control-profile*) to see the minimum guaranteed bandwidth (**guaranteed-rate**) and maximum bandwidth (**shaping-rate**) for the forwarding class set.

6. Delete the explicit FCoE forwarding-class-to-queue mapping so that the system uses the default FCoE forwarding-class-to-queue mapping. Include the **delete forwarding-classes class fcoe queue-num 3** statement at the **[edit class-of-service]** hierarchy level to remove the explicit configuration. The system then uses the default configuration for the FCoE forwarding class and preserves the lossless treatment of FCoE traffic.
7. Use the **show class-of-service forwarding-class** operational command to display the configured forwarding classes. The *No-Loss* column shows whether lossless transport is enabled or disabled for each forwarding class. If the forwarding class used for FCoE traffic is not enabled for lossless transport, include the **no-loss** packet drop attribute in the forwarding class configuration (**set class-of-service forwarding-classes class *fcoe-forwarding-class-name* queue-num *queue-number* no-loss**).

See “[Example: Configuring CoS PFC for FCoE Traffic](#)” on page 6504 for step-by-step instructions on how to configure PFC for FCoE traffic, including classifier, interface, congestion notification profile, PFC, and bandwidth scheduling configuration.

#### Related Documentation

- [show class-of-service congestion-notification on page 7394](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 6905](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 6504](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 6493](#)

## Understanding MC-LAGs on an FCoE Transit Switch

---

Multichassis link aggregation groups (MC-LAGs) provide redundancy and load balancing between two switches, multihoming support for client devices such as servers, and a loop-free Layer 2 network without running Spanning Tree Protocol (STP).

You can use an MC-LAG to provide a redundant aggregation layer for Fibre Channel over Ethernet (FCoE) traffic. To support lossless transport of FCoE traffic across an MC-LAG, you must configure the appropriate class of service (CoS) on both of the switches with MC-LAG port members. The CoS configuration must be the same on both of the MC-LAG switches because MC-LAGs do not carry forwarding class and IEEE 802.1p priority information.

Ports that are part of an FCoE-FC gateway configuration (a virtual FCoE-FC gateway fabric) do not support MC-LAGs. Ports that are members of an MC-LAG act as pass-through transit switch ports.

Standalone switches support MC-LAGs. QFabric system Node devices do not support MC-LAGs. Virtual Chassis and mixed-mode Virtual Chassis Fabric (VCF) configurations do not support FCoE. Only pure QFX5100 VCFs (consisting of only QFX5100 switches) support FCoE.

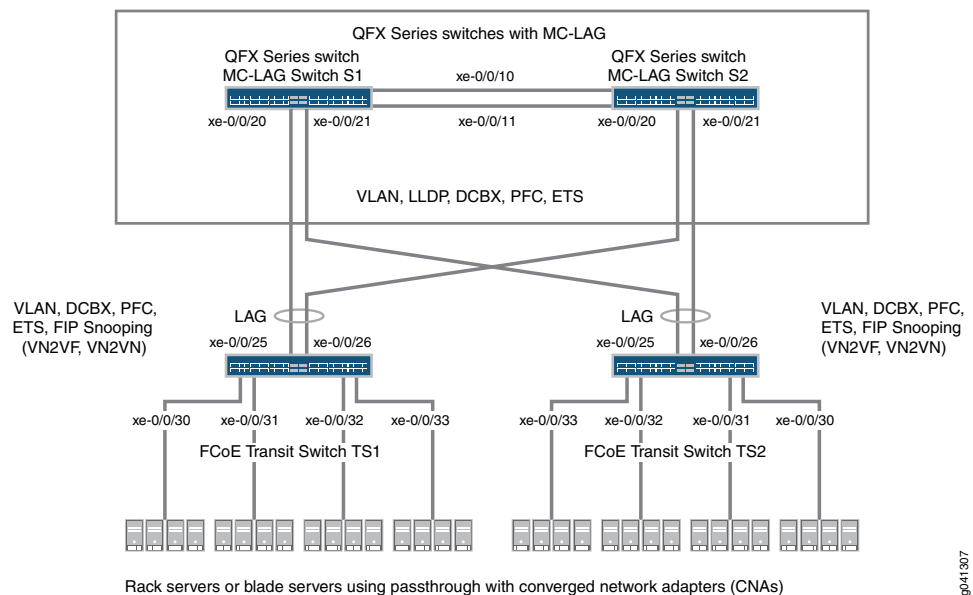
This topic describes:

- [Supported Topology on page 6451](#)
- [FIP Snooping and FCoE Trusted Ports on page 6453](#)
- [CoS and Data Center Bridging \(DCB\) on page 6454](#)

### Supported Topology

Switches that are not directly connected to FCoE hosts and that act as pass-through transit switches support MC-LAGs for FCoE traffic in an *inverted-U* network topology. [Figure 205](#) shows an inverted-U topology using QFX3500 switches.

Figure 205: Supported Topology for an MC-LAG on an FCoE Transit Switch



The following rules and guidelines apply to MC-LAGs when used for FCoE traffic. The rules and guidelines help to ensure the proper handling and lossless transport characteristics required for FCoE traffic.

- The two switches that form the MC-LAG (Switches S1 and S2) cannot use ports that are part of an FCoE-FC gateway fabric. The MC-LAG switch ports must be pass-through transit switch ports (used as part of an intermediate transit switch that is not directly connected to FCoE hosts).
- MC-LAG Switches S1 and S2 cannot be directly connected to the FCoE hosts.
- The two switches that serve as access devices for FCoE hosts (FCoE Transit Switches TS1 and TS2) use standard LAGs to connect to MC-LAG Switches S1 and S2. FCoE Transit Switches TS1 and TS2 can be standalone switches or they can be Node devices in a QFabric system.
- Transit Switches TS1 and TS2 must use transit switch ports for the FCoE hosts and for the standard LAGs to MC-LAG Switches S1 and S2.
- Enable FIP snooping on the FCoE VLAN on Transit Switches TS1 and TS2. You can configure either VN\_Port to VF\_Port (VN2VF\_Port) FIP snooping or VN\_Port to VN\_Port (VN2VN\_Port) FIP snooping, depending on whether the FCoE hosts need to access targets in the FC SAN (VN2VF\_Port FIP snooping) or targets in the Ethernet network (VN2VN\_Port FIP snooping).

FIP snooping should be performed at the access edge and is not supported on MC-LAG switches. Do not enable FIP snooping on MC-LAG Switches S1 and S2. (Do not enable FIP snooping on the MC-LAG ports that connect Switches S1 and S2 to Switches TS1 and TS2 or on the LAG ports that connect Switch S1 to S2.)



**NOTE:** Juniper Networks QFX10000 aggregation switches do not support FIP snooping, so they cannot be used as FIP snooping access switches (Transit Switches TS1 and TS2) in this topology.

- The CoS configuration must be consistent on the MC-LAG switches. Because MC-LAGs carry no forwarding class or priority information, each MC-LAG switch needs to have the same CoS configuration to support lossless transport. (On each MC-LAG switch, the name, egress queue, and CoS provisioning of each forwarding class must be the same, and the priority-based flow control (PFC) configuration must be the same.)

### Transit Switches (Server Access)

The role of FCoE Transit Switches TS1 and TS2 is to connect FCoE hosts in a multihomed fashion to the MC-LAG switches, so Transit Switches TS1 and TS2 act as access switches for the FCoE hosts. (FCoE hosts are directly connected to Transit Switches TS1 and TS2.)

The transit switch configuration depends on whether you want to do VN2VF\_Port FIP snooping or VN2VN\_Port FIP snooping, and whether the transit switches also have ports configured as part of an FCoE-FC gateway virtual fabric. Ports that a QFX3500 switch uses in an FCoE-FC gateway virtual fabric cannot be included in the transit switch LAG connection to the MC-LAG switches. (Ports cannot belong to both a transit switch and an FCoE-FC gateway; you must use different ports for each mode of operation.)

### MC-LAG Switches (FCoE Aggregation)

The role of MC-LAG Switches S1 and S2 is to provide redundant, load-balanced connections between FCoE transit switches. The MC-LAG Switches S1 and S2 act as aggregation switches. FCoE hosts are not directly connected to the MC-LAG switches.

The MC-LAG switch configuration is the same regardless of which type of FIP snooping FCoE Transit Switches TS1 and TS2 perform.

## FIP Snooping and FCoE Trusted Ports

To maintain secure access, enable VN2VF\_Port FIP snooping or VN2VN\_Port FIP snooping at the transit switch access ports connected directly to the FCoE hosts. FIP snooping should be performed at the access edge of the network to prevent unauthorized access. For example, in [Figure 205](#), you enable FIP snooping on the FCoE VLANs on Transit Switches TS1 and TS2 that include the access ports connected to the FCoE hosts.

Do not enable FIP snooping on the switches used to create the MC-LAG. For example, in [Figure 205](#), you would not enable FIP snooping on the FCoE VLANs on Switches S1 and S2.

Configure links between switches as FCoE trusted ports to reduce FIP snooping overhead and ensure that the system performs FIP snooping only at the access edge. In the sample topology, configure the Transit Switch TS1 and TS2 LAG ports connected to the MC-LAG switches as FCoE trusted ports, configure the Switch S1 and S2 MC-LAG ports connected to Switches TS1 and TS2 as FCoE trusted ports, and configure the ports in the LAG that connects Switches S1 to S2 as FCoE trusted ports.

## CoS and Data Center Bridging (DCB)

The MC-LAG links do not carry forwarding class or priority information. The following CoS properties must have the same configuration on each MC-LAG switch or on each MC-LAG interface to support lossless transport:

- FCoE forwarding class name—For example, the forwarding class for FCoE traffic could use the default **fcoe** forwarding class on both MC-LAG switches.
- FCoE output queue—For example, the **fcoe** forwarding class could be mapped to queue 3 on both MC-LAG switches (queue 3 is the default mapping for the **fcoe** forwarding class).
- Classifier—The forwarding class for FCoE traffic must be mapped to the same IEEE 802.1p code point on each member interface of the MC-LAG on both MC-LAG switches. For example, the FCoE forwarding class **fcoe** could be mapped to IEEE 802.1p code point **011** (code point **011** is the default mapping for the **fcoe** forwarding class).
- Priority-based flow control (PFC)—PFC must be enabled on the FCoE code point on each MC-LAG switch and applied to each MC-LAG interface using a congestion notification profile.

You must also configure enhanced transmission selection (ETS) on the MC-LAG interfaces to provide sufficient scheduling resources (bandwidth, priority) for lossless transport. The ETS configuration can be different on each MC-LAG switch, as long as enough resources are scheduled to support lossless transport for the expected FCoE traffic.

Link Layer Discovery Protocol (LLDP) and Data Center Bridging Capability Exchange Protocol (DCBX) must be enabled on each MC-LAG member interface (LLDP and DCBX are enabled by default on all interfaces).



**NOTE:** As with all other FCoE configurations, FCoE traffic requires a dedicated VLAN that carries only FCoE traffic, and IGMP snooping must be disabled on the FCoE VLAN.

---

## Example: Configuring CoS for FCoE Transit Switch Traffic Across an MC-LAG

---

Multichassis link aggregation groups (MC-LAGs) provide redundancy and load balancing between two QFX Series switches, multihoming support for client devices such as servers, and a loop-free Layer 2 network without running Spanning Tree Protocol (STP).



**NOTE:** This example uses the Junos OS Enhanced Layer 2 Software (ELS) configuration style for QFX Series switches. If your switch runs software that does not support ELS, see *Example: Configuring CoS for FCoE Transit Switch Traffic Across an MC-LAG*. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 41](#).

You can use an MC-LAG to provide a redundant aggregation layer for Fibre Channel over Ethernet (FCoE) traffic in an *inverted-U* topology. To support lossless transport of FCoE traffic across an MC-LAG, you must configure the appropriate class of service (CoS) on both of the QFX Series switches with MC-LAG port members. The CoS configuration must be the same on both of the MC-LAG switches because an MC-LAG does not carry forwarding class and IEEE 802.1p priority information.

Ports that are members of an MC-LAG act as FCoE passthrough transit switch ports.



**NOTE:** This example describes how to configure CoS to provide lossless transport for FCoE traffic across an MC-LAG that connects two QFX Series switches. It also describes how to configure CoS on the FCoE transit switches that connect FCoE hosts to the QFX Series switches that form the MC-LAG.

This example does *not* describe how to configure the MC-LAG itself. For a detailed example of MC-LAG configuration, see *Example: Configuring Multichassis Link Aggregation*. However, this example includes a subset of MC-LAG configuration that only shows how to configure interface membership in the MC-LAG.



**NOTE:** Juniper Networks QFX10000 aggregation switches do not support FIP snooping, so they cannot be used as FIP snooping access switches (Transit Switches TS1 and TS2) in this example. However, QFX10000 switches can play the role of the MC-LAG switches (MC-LAG Switch S1 and MC-LAG Switch S2) in this example.

QFX3500 and QFX3600 Virtual Chassis switches do not support FCoE.

This topic describes:

- [Requirements on page 6455](#)
- [Overview on page 6456](#)
- [Configuration on page 6461](#)
- [Verification on page 6471](#)

## Requirements

This example uses the following hardware and software components:

- Two Juniper Networks QFX5100 Switches running the ELS CLI that form an MC-LAG for FCoE traffic.
- Two Juniper Networks QFX5100 Switches running the ELS CLI that provide FCoE server access in transit switch mode and that connect to the MC-LAG switches.
- FCoE servers (or other FCoE hosts) connected to the transit switches.
- Junos OS Release 13.2 or later for the QFX Series.

## Overview

FCoE traffic requires lossless transport. This example shows you how to:

- Configure CoS for FCoE traffic on the two QFX5100 switches that form the MC-LAG, including priority-based flow control (PFC). The example also includes configuration for both enhanced transmission selection (ETS) hierarchical scheduling of resources for the FCoE forwarding class priority and for the forwarding class set priority group, and also direct port scheduling. You can only use one of the scheduling methods on a port. Different switches support different scheduling methods.



**NOTE:** Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

- Configure CoS for FCoE on the two FCoE transit switches that connect FCoE hosts to the MC-LAG switches and enable FIP snooping on the FCoE VLAN at the FCoE transit switch access ports.
- Configure the appropriate port mode, MTU, and FCoE trusted or untrusted state for each interface to support lossless FCoE transport.



**NOTE:** Do not enable IGMP snooping on the FCoE VLAN. (IGMP snooping is enabled on the default VLAN by default, but is disabled by default on all other VLANs.)


---

## Topology

QFX5100 switches that act as transit switches support MC-LAGs for FCoE traffic in an inverted-U network topology, as shown in [Figure 206](#).



Figure 206: Supported Topology for an MC-LAG on an FCoE Transit Switch



**NOTE:** Juniper Networks QFX10000 aggregation switches do not support FIP snooping, so they cannot be used as FIP snooping access switches (Transit Switches TS1 and TS2) in this example. However, QFX10000 switches can play the role of the MC-LAG switches (MC-LAG Switch S1 and MC-LAG Switch S2) in this example.

Table 512 shows the configuration components for this example.

Table 512: Components of the CoS for FCoE Traffic Across an MC-LAG Configuration Topology

| Component                                                                  | Settings                                                                                                                                     |
|----------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| Hardware                                                                   | Four QFX5100 switches running the ELS CLI (two to form the MC-LAG as passthrough transit switches and two transit switches for FCoE access). |
| Forwarding class (all switches)                                            | Default <b>fcoe</b> forwarding class.                                                                                                        |
| Classifier (forwarding class mapping of incoming traffic to IEEE priority) | Default IEEE 802.1p trusted classifier on all FCoE interfaces.                                                                               |

Table 512: Components of the CoS for FCoE Traffic Across an MC-LAG Configuration Topology (*continued*)

| Component                                            | Settings                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| LAGs and MC-LAG                                      | <p>S1—Ports xe-0/0/10 and x-0/0/11 are members of LAG <b>ae0</b>, which connects Switch S1 to Switch S2.<br/>Ports xe-0/0/20 and xe-0/0/21 are members of MC-LAG <b>ae1</b>. All ports are configured in <b>trunk</b> interface mode, as <b>fcoe-trusted</b>, and with an MTU of <b>2180</b>.</p> <p>S2—Ports xe-0/0/10 and x-0/0/11 are members of LAG <b>ae0</b>, which connects Switch S2 to Switch S1.<br/>Ports xe-0/0/20 and xe-0/0/21 are members of MC-LAG <b>ae1</b>. All ports are configured in <b>trunk</b> interface mode, as <b>fcoe-trusted</b>, and with an MTU of <b>2180</b>.</p> <p><b>NOTE:</b> Ports xe-0/0/20 and xe-0/0/21 on Switches S1 and S2 are the members of the MC-LAG.</p> <p>TS1—Ports xe-0/0/25 and x-0/0/26 are members of LAG <b>ae1</b>, configured in <b>trunk</b> interface mode, as <b>fcoe-trusted</b>, and with an MTU of <b>2180</b>.<br/>Ports xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33 are configured in <b>trunk</b> interface mode, with an MTU of <b>2180</b>.</p> <p>TS2—Ports xe-0/0/25 and x-0/0/26 are members of LAG <b>ae1</b>, configured in <b>trunk</b> interface mode, as <b>fcoe-trusted</b>, and with an MTU of <b>2180</b>.<br/>Ports xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33 are configured in <b>trunk</b> interface mode, with an MTU of <b>2180</b>.</p> |
| FCoE queue scheduler (all switches)                  | <p><b>fcoe-sched:</b><br/>Minimum bandwidth <b>3g</b><br/>Maximum bandwidth <b>100%</b><br/>Priority <b>low</b></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Forwarding class-to-scheduler mapping (all switches) | <p>Scheduler map <b>fcoe-map:</b><br/>Forwarding class <b>fcoe</b><br/>Scheduler <b>fcoe-sched</b></p> <p><b>NOTE:</b> If you are using direct port scheduling,</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| PFC congestion notification profile (all switches)   | <p><b>fcoe-cnp:</b><br/>Code point <b>011</b></p> <p>Ingress interfaces:</p> <ul style="list-style-type: none"> <li>• S1—LAG <b>ae0</b> and MC-LAG <b>ae1</b></li> <li>• S2—LAG <b>ae0</b> and MC-LAG <b>ae1</b></li> <li>• TS1—LAG <b>ae1</b>, interfaces <b>xe-0/0/30</b>, <b>xe-0/0/31</b>, <b>xe-0/0/32</b>, and <b>xe-0/0/33</b></li> <li>• TS2—LAG <b>ae1</b>, interfaces <b>xe-0/0/30</b>, <b>xe-0/0/31</b>, <b>xe-0/0/32</b>, and <b>xe-0/0/33</b></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

**Table 512: Components of the CoS for FCoE Traffic Across an MC-LAG Configuration Topology** (*continued*)

| Component                                                         | Settings                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| FCoE VLAN name and tag ID                                         | <p>Name—<b>fcoe_vlan</b><br/>ID—<b>100</b></p> <p>Include the FCoE VLAN on the interfaces that carry FCoE traffic on all four switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| ETS only—forwarding class set (FCoE priority group, all switches) | <p><b>fcoe-pg:</b><br/>Forwarding class <b>fcoe</b></p> <p>Egress interfaces:</p> <ul style="list-style-type: none"> <li>• S1—LAG <b>ae0</b> and MC-LAG <b>ae1</b></li> <li>• S2—LAG <b>ae0</b> and MC-LAG <b>ae1</b></li> <li>• TS1—LAG <b>ae1</b>, interfaces <b>xe-0/0/30</b>, <b>xe-0/0/31</b>, <b>xe-0/0/32</b>, and <b>xe-0/0/33</b></li> <li>• TS2—LAG <b>ae1</b>, interfaces <b>xe-0/0/30</b>, <b>xe-0/0/31</b>, <b>xe-0/0/32</b>, and <b>xe-0/0/33</b></li> </ul>                                                                                                                                                                                                                                                            |
| ETS only—traffic control profile (all switches)                   | <p><b>fcoe-tcp:</b><br/>Scheduler map <b>fcoe-map</b><br/>Minimum bandwidth <b>3g</b><br/>Maximum bandwidth <b>100%</b></p> <p>The traffic control profile is applied to the same interfaces as the forwarding class set, using the same CLI statement. This applies ETS hierarchical scheduling to the interfaces.</p>                                                                                                                                                                                                                                                                                                                                                                                                               |
| Port scheduling only—apply scheduling to interfaces               | <p>On switches that support direct port scheduling, if you use port scheduling, apply scheduling by attaching the scheduler map directly to interfaces:</p> <ul style="list-style-type: none"> <li>• S1—LAG <b>ae0</b> and MC-LAG <b>ae1</b></li> <li>• S2—LAG <b>ae0</b> and MC-LAG <b>ae1</b></li> <li>• TS1—LAG <b>ae1</b>, interfaces <b>xe-0/0/30</b>, <b>xe-0/0/31</b>, <b>xe-0/0/32</b>, and <b>xe-0/0/33</b></li> <li>• TS2—LAG <b>ae1</b>, interfaces <b>xe-0/0/30</b>, <b>xe-0/0/31</b>, <b>xe-0/0/32</b>, and <b>xe-0/0/33</b></li> </ul>                                                                                                                                                                                  |
| FIP snooping                                                      | <p>Enable FIP snooping on Transit Switches TS1 and TS2 on the FCoE VLAN. Configure the LAG interfaces that connect to the MC-LAG switches as FCoE trusted interfaces so that they do not perform FIP snooping.</p> <p>This example enables VN2VN_Port FIP snooping on the FCoE transit switch interfaces connected to the FCoE servers. The example is equally valid with VN2VF_Port FIP snooping enabled on the transit switch access ports. The method of FIP snooping you enable depends on your network configuration.</p> <p><b>NOTE:</b> Juniper Networks QFX10000 aggregation switches do not support FIP snooping, so they cannot be used as FIP snooping access switches (Transit Switches TS1 and TS2) in this example.</p> |



**NOTE:** This example uses the default IEEE 802.1p trusted BA classifier, which is automatically applied to trunk mode interfaces if you do not apply an explicitly configured classifier.

To configure CoS for FCoE traffic across an MC-LAG:

- Use the default FCoE forwarding class and forwarding-class-to-queue mapping (do not explicitly configure the FCoE forwarding class or output queue). The default FCoE forwarding class is **fcoe**, and the default output queue is **queue 3**.
- Use the default trusted BA classifier, which maps incoming packets to forwarding classes by the IEEE 802.1p code point (CoS priority) of the packet. The trusted classifier is the default classifier for interfaces in trunk interface mode. The default trusted classifier maps incoming packets with the IEEE 802.1p code point 3 (**011**) to the FCoE forwarding class. If you choose to configure the BA classifier instead of using the default classifier, you must ensure that FCoE traffic is classified into forwarding classes in exactly the same way on both MC-LAG switches. Using the default classifier ensures consistent classifier configuration on the MC-LAG ports.
- Configure a congestion notification profile that enables PFC on the FCoE code point (code point **011** in this example). The congestion notification profile configuration must be the same on both MC-LAG switches.
- Apply the congestion notification profile to the interfaces.
- Configure the interface mode, MTU, and FCoE trusted or untrusted state for each interface to support lossless FCoE transport.
- For ETS hierarchical port scheduling, configure ETS on the interfaces to provide the bandwidth required for lossless FCoE transport. Configuring ETS includes configuring bandwidth scheduling for the FCoE forwarding class, a forwarding class set (priority group) that includes the FCoE forwarding class, and a traffic control profile to assign bandwidth to the forwarding class set that includes FCoE traffic, and applying the traffic control profile and forwarding class set to interfaces..

On switches that support direct port scheduling, configure CoS properties on interfaces by applying scheduler maps directly to interfaces.

In addition, this example describes how to enable FIP snooping on the Transit Switch TS1 and TS2 ports that are connected to the FCoE servers. To provide secure access, FIP snooping must be enabled on the FCoE access ports.

This example focuses on the CoS configuration to support lossless FCoE transport across an MC-LAG. This example does not describe how to configure the properties of MC-LAGs and LAGs, although it does show you how to configure the port characteristics required to support lossless transport and how to assign interfaces to the MC-LAG and to the LAGs.

Before you configure CoS, configure:

- The MC-LAGs that connect Switches S1 and S2 to Switches TS1 and TS2. (*Example: Configuring Multichassis Link Aggregation* describes how to configure MC-LAGs.)

- The LAGs that connect the Transit Switches TS1 and TS2 to MC-LAG Switches S1 and S2. (“[Configuring Link Aggregation](#)” on [page 2869](#) describes how to configure LAGs.)
- The LAG that connects Switch S1 to Switch S2.

## Configuration

To configure CoS for lossless FCoE transport across an MC-LAG, perform these tasks:

- [MC-LAG Switches S1 and S2 Common Configuration \(Applies to ETS and Port Scheduling\) on page 6463](#)
- [MC-LAG Switches S1 and S2 ETS Hierarchical Scheduling Configuration on page 6464](#)
- [MC-LAG Switches S1 and S2 Port Scheduling Configuration on page 6465](#)
- [FCoE Transit Switches TS1 and TS2 Common Configuration \(Applies to ETS and Port Scheduling\) on page 6465](#)
- [FCoE Transit Switches TS1 and TS2 ETS Hierarchical Scheduling Configuration on page 6467](#)
- [FCoE Transit Switches TS1 and TS2 Port Scheduling Configuration on page 6467](#)
- [Results on page 6468](#)

### CLI Quick Configuration



**NOTE:** The CLI configurations for the MC-LAG switches and for the FCoE transit switches are each separated into three sections:

- Configuration common to all port scheduling methods
- Configuration specific to ETS hierarchical port scheduling
- Configuration specific to direct port scheduling

### MC-LAG Switch S1 and Switch S2

To quickly configure CoS for lossless FCoE transport across an MC-LAG, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI for MC-LAG Switch S1 and MC-LAG Switch S2 at the **[edit]** hierarchy level. The configurations on Switches S1 and S2 are identical because the CoS configuration must be identical, and because this example uses the same ports on both switches.

### MC-LAG Switches Configuration Common to ETS Hierarchical Port Scheduling and to Direct Port Scheduling

```
set class-of-service schedulers fcoe-sched priority low transmit-rate 3g
set class-of-service schedulers fcoe-sched shaping-rate percent 100
set class-of-service scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
set class-of-service congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011 pfc
set class-of-service interfaces ae0 congestion-notification-profile fcoe-cnp
set class-of-service interfaces ae1 congestion-notification-profile fcoe-cnp
set vlans fcoe_vlan vlan-id 100
set interfaces xe-0/0/10 ether-options 802.3ad ae0
set interfaces xe-0/0/11 ether-options 802.3ad ae0
set interfaces xe-0/0/20 ether-options 802.3ad ae1
set interfaces xe-0/0/21 ether-options 802.3ad ae1
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk vlan members fcoe_vlan
set interfaces ae1 unit 0 family ethernet-switching interface-mode trunk vlan members fcoe_vlan
```

```
set interfaces ae0 mtu 2180
set interfaces ae1 mtu 2180
set vlans fcoe_vlan forwarding-options fip-security interface ae0 fcoe-trusted
set vlans fcoe_vlan forwarding-options fip-security interface ae1 fcoe-trusted
```

### MC-LAG Switches Configuration for ETS Hierarchical Port Scheduling

```
set class-of-service forwarding-class-sets fcoe-pg class fcoe
set class-of-service traffic-control-profiles fcoe-tcp scheduler-map fcoe-map guaranteed-rate
3g
set class-of-service traffic-control-profiles fcoe-tcp shaping-rate percent 100
set class-of-service interfaces ae0 forwarding-class-set fcoe-pg output-traffic-control-profile
fcoe-tcp
set class-of-service interfaces ae1 forwarding-class-set fcoe-pg output-traffic-control-profile
fcoe-tcp
```

### MC-LAG Switches Configuration for Direct Port Scheduling

```
set class-of-service interfaces ae0 scheduler-map fcoe-map
set class-of-service interfaces ae1 scheduler-map fcoe-map
```

#### FCoE Transit Switch TS1 and Switch TS2

To quickly configure CoS for lossless FCoE transport across an MC-LAG, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI for Transit Switch TS1 and Transit Switch TS2 at the **[edit]** hierarchy level. The configurations on Switches TS1 and TS2 are identical because the CoS configuration must be identical, and because this example uses the same ports on both switches.

### FCoE Transit Switches Configuration Common to ETS Hierarchical Port Scheduling and to Direct Port Scheduling

```
set class-of-service schedulers fcoe-sched priority low transmit-rate 3g
set class-of-service schedulers fcoe-sched shaping-rate percent 100
set class-of-service scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
set class-of-service congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011 pfc
set class-of-service interfaces ae1 congestion-notification-profile fcoe-cnp
set class-of-service interfaces xe-0/0/30 congestion-notification-profile fcoe-cnp
set class-of-service interfaces xe-0/0/31 congestion-notification-profile fcoe-cnp
set class-of-service interfaces xe-0/0/32 congestion-notification-profile fcoe-cnp
set class-of-service interfaces xe-0/0/33 congestion-notification-profile fcoe-cnp
set vlans fcoe_vlan vlan-id 100
set interfaces xe-0/0/25 ether-options 802.3ad ae1
set interfaces xe-0/0/26 ether-options 802.3ad ae1
set interfaces ae1 unit 0 family ethernet-switching interface-mode trunk vlan members fcoe_vlan
set interfaces xe-0/0/30 unit 0 family ethernet-switching interface-mode trunk vlan members
fcoe_vlan
set interfaces xe-0/0/31 unit 0 family ethernet-switching interface-mode trunk vlan members
fcoe_vlan
set interfaces xe-0/0/32 unit 0 family ethernet-switching interface-mode trunk vlan members
fcoe_vlan
set interfaces xe-0/0/33 unit 0 family ethernet-switching interface-mode trunk vlan members
fcoe_vlan
set interfaces ae1 mtu 2180
set interfaces xe-0/0/30 mtu 2180
set interfaces xe-0/0/31 mtu 2180
set interfaces xe-0/0/32 mtu 2180
set interfaces xe-0/0/33 mtu 2180
set vlans fcoe_vlan forwarding-options fip-security interface ae1 fcoe-trusted
set vlans fcoe_vlan forwarding-options fip-security examine-vn2v2 beacon-period 90000
```

## FCoE Transit Switches Configuration for ETS Hierarchical Port Scheduling

```

set class-of-service forwarding-class-sets fcoe-pg class fcoe
set class-of-service traffic-control-profiles fcoe-tcp scheduler-map fcoe-map guaranteed-rate
3g
set class-of-service traffic-control-profiles fcoe-tcp shaping-rate percent 100
set class-of-service interfaces ae1 forwarding-class-set fcoe-pg output-traffic-control-profile
fcoe-tcp
set class-of-service interfaces xe-0/0/30 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
set class-of-service interfaces xe-0/0/31 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
set class-of-service interfaces xe-0/0/32 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
set class-of-service interfaces xe-0/0/33 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp

```

## FCoE Transit Switches Configuration for Direct Port Scheduling

```

set class-of-service interfaces ae1 scheduler-map fcoe-map
set class-of-service interfaces xe-0/0/30 scheduler-map fcoe-map
set class-of-service interfaces xe-0/0/31 scheduler-map fcoe-map
set class-of-service interfaces xe-0/0/32 scheduler-map fcoe-map
set class-of-service interfaces xe-0/0/33 scheduler-map fcoe-map

```

## MC-LAG Switches S1 and S2 Common Configuration (Applies to ETS and Port Scheduling)

### Step-by-Step Procedure

To configure queue scheduling, PFC, the FCoE VLAN, and LAG and MC-LAG interface membership and characteristics to support lossless FCoE transport across an MC-LAG (this example uses the default **fcoe** forwarding class and the default classifier to map incoming FCoE traffic to the FCoE IEEE 802.1p code point **011**), for both ETS hierarchical port scheduling and port scheduling (common configuration):

1. Configure output scheduling for the FCoE queue:
 

```

[edit class-of-service]
user@switch# set schedulers fcoe-sched priority low transmit-rate 3g
user@switch# set schedulers fcoe-sched shaping-rate percent 100

```
2. Map the FCoE forwarding class to the FCoE scheduler (**fcoe-sched**):
 

```

[edit class-of-service]
user@switch# set scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched

```
3. Enable PFC on the FCoE priority by creating a congestion notification profile (**fcoe-cnp**) that applies FCoE to the IEEE 802.1 code point **011**:
 

```

[edit class-of-service]
user@switch# set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point
011 pfc

```
4. Apply the PFC configuration to the LAG and MC-LAG interfaces:
 

```

[edit class-of-service]
user@switch# set interfaces ae0 congestion-notification-profile fcoe-cnp
user@switch# set interfaces ae1 congestion-notification-profile fcoe-cnp

```
5. Configure the VLAN for FCoE traffic (**fcoe\_vlan**):

```
[edit vlans]
user@switch# set fcoe_vlan vlan-id 100
```

6. Add the member interfaces to the LAG between the two MC-LAG switches:

```
[edit interfaces]
user@switch# set xe-0/0/10 ether-options 802.3ad ae0
user@switch# set xe-0/0/11 ether-options 802.3ad ae0
```

7. Add the member interfaces to the MC-LAG:

```
[edit interfaces]
user@switch# set xe-0/0/20 ether-options 802.3ad ae1
user@switch# set xe-0/0/21 ether-options 802.3ad ae1
```

8. Configure the interface mode as **trunk** and membership in the FCoE VLAN (**fcoe\_vlan**) for the LAG (**ae0**) and for the MC-LAG (**ae1**):

```
[edit interfaces]
user@switch# set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
vlan members fcoe_vlan
user@switch# set interfaces ae1 unit 0 family ethernet-switching interface-mode trunk
vlan members fcoe_vlan
```

9. Set the MTU to **2180** for the LAG and MC-LAG interfaces. 2180 bytes is the minimum size required to handle FCoE packets because of the payload and header sizes; you can configure the MTU to a higher number of bytes if desired, but not less than 2180 bytes:

```
[edit interfaces]
user@switch# set ae0 mtu 2180
user@switch# set ae1 mtu 2180
```

10. Set the LAG and MC-LAG interfaces as FCoE trusted ports. Ports that connect to other switches should be trusted and should not perform FIP snooping:

```
[edit]
user@switch# set vlans fcoe_vlan forwarding-options fip-security interface ae0 fcoe-trusted
user@switch# set vlans fcoe_vlan forwarding-options fip-security interface ae1 fcoe-trusted
```

### MC-LAG Switches S1 and S2 ETS Hierarchical Scheduling Configuration

---

#### Step-by-Step Procedure

To configure the forwarding class set (priority group) and priority group scheduling (in a traffic control profile), and apply the ETS hierarchical scheduling for FCoE traffic to interfaces:

1. Configure the forwarding class set (**fcoe-pg**) for the FCoE traffic:

```
[edit class-of-service]
user@switch# set forwarding-class-sets fcoe-pg class fcoe
```

2. Define the traffic control profile (**fcoe-tcp**) to use on the FCoE forwarding class set:

```
[edit class-of-service]
user@switch# set traffic-control-profiles fcoe-tcp scheduler-map fcoe-map
guaranteed-rate 3g
user@switch# set traffic-control-profiles fcoe-tcp shaping-rate percent 100
```

3. Apply the FCoE forwarding class set and traffic control profile to the LAG and MC-LAG interfaces:

```
[edit class-of-service]
```



```

user@switch# set interfaces ae0 forwarding-class-set fcoe-pg output-traffic-control-profile
fcoe-tcp
user@switch# set interfaces ae1 forwarding-class-set fcoe-pg output-traffic-control-profile
fcoe-tcp

```

### MC-LAG Switches S1 and S2 Port Scheduling Configuration

#### Step-by-Step Procedure

To apply port scheduling for FCoE traffic to interfaces:

1. Apply the scheduler map to the egress ports:

```

set class-of-service interfaces ae0 scheduler-map fcoe-map
set class-of-service interfaces ae1 scheduler-map fcoe-map

```

### FCoE Transit Switches TS1 and TS2 Common Configuration (Applies to ETS and Port Scheduling)

#### Step-by-Step Procedure

The CoS configuration on FCoE Transit Switches TS1 and TS2 is similar to the CoS configuration on MC-LAG Switches S1 and S2. However, the port configurations differ, and you must enable FIP snooping on the Switch TS1 and Switch TS2 FCoE access ports.

To configure queue scheduling, PFC, the FCoE VLAN, and LAG interface membership and characteristics to support lossless FCoE transport across the MC-LAG (this example uses the default **fcoe** forwarding class and the default classifier to map incoming FCoE traffic to the FCoE IEEE 802.1p code point 011, so you do not configure them), or both ETS hierarchical scheduling and port scheduling (common configuration):

1. Configure output scheduling for the FCoE queue:

```

[edit class-of-service]
user@switch# set schedulers fcoe-sched priority low transmit-rate 3g
user@switch# set schedulers fcoe-sched shaping-rate percent 100

```

2. Map the FCoE forwarding class to the FCoE scheduler (**fcoe-sched**):

```

[edit class-of-service]
user@switch# set scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched

```

3. Enable PFC on the FCoE priority by creating a congestion notification profile (**fcoe-cnp**) that applies FCoE to the IEEE 802.1 code point 011:

```

[edit class-of-service]
user@switch# set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point
011 pfc

```

4. Apply the PFC configuration to the LAG interface and to the FCoE access interfaces:

```

[edit class-of-service]
user@switch# set interfaces ae1 congestion-notification-profile fcoe-cnp
user@switch# set class-of-service interfaces xe-0/0/30 congestion-notification-profile
fcoe-cnp
user@switch# set class-of-service interfaces xe-0/0/31 congestion-notification-profile
fcoe-cnp
user@switch# set class-of-service interfaces xe-0/0/32 congestion-notification-profile
fcoe-cnp
user@switch# set class-of-service interfaces xe-0/0/33 congestion-notification-profile
fcoe-cnp

```

5. Configure the VLAN for FCoE traffic (**fcoe\_vlan**):

```
[edit vlans]
user@switch# set fcoe_vlan vlan-id 100
```

6. Add the member interfaces to the LAG:

```
[edit interfaces]
user@switch# set xe-0/0/25 ether-options 802.3ad ae1
user@switch# set xe-0/0/26 ether-options 802.3ad ae1
```

7. On the LAG (**ae1**), configure the interface mode as **trunk** and membership in the FCoE VLAN (**fcoe\_vlan**):

```
[edit interfaces]
user@switch# set interfaces ae1 unit 0 family ethernet-switching interface-mode trunk
vlan members fcoe_vlan
```

8. On the FCoE access interfaces (**xe-0/0/30**, **xe-0/0/31**, **xe-0/0/32**, **xe-0/0/33**), configure the interface mode as **trunk** and membership in the FCoE VLAN (**fcoe\_vlan**):

```
[edit interfaces]
user@switch# set interfaces xe-0/0/30 unit 0 family ethernet-switching interface-mode
trunk vlan members fcoe_vlan
user@switch# set interfaces xe-0/0/31 unit 0 family ethernet-switching interface-mode
trunk vlan members fcoe_vlan
user@switch# set interfaces xe-0/0/32 unit 0 family ethernet-switching interface-mode
trunk vlan members fcoe_vlan
user@switch# set interfaces xe-0/0/33 unit 0 family ethernet-switching interface-mode
trunk vlan members fcoe_vlan
```

9. Set the MTU to **2180** for the LAG and FCoE access interfaces. 2180 bytes is the minimum size required to handle FCoE packets because of the payload and header sizes; you can configure the MTU to a higher number of bytes if desired, but not less than 2180 bytes:

```
[edit interfaces]
user@switch# set ae1 mtu 2180
user@switch# set xe-0/0/30 mtu 2180
user@switch# set xe-0/0/31 mtu 2180
user@switch# set xe-0/0/32 mtu 2180
user@switch# set xe-0/0/33 mtu 2180
```

10. Set the LAG interface as an FCoE trusted port. Ports that connect to other switches should be trusted and should not perform FIP snooping:

```
[edit]
user@switch# set vlans fcoe_vlan forwarding-options fip-security interface ae1 fcoe-trusted
```



**NOTE:** Access ports **xe-0/0/30**, **xe-0/0/31**, **xe-0/0/32**, and **xe-0/0/33** are not configured as FCoE trusted ports. The access ports remain in the default state as untrusted ports because they connect directly to FCoE devices and must perform FIP snooping to ensure network security.

11. Enable FIP snooping on the FCoE VLAN to prevent unauthorized FCoE network access (this example uses **VN2VN\_Port** FIP snooping; the example is equally valid if you use **VN2VF\_Port** FIP snooping):

```
[edit]
```

```
user@switch# set vlans fcoe_vlan forwarding-options fip-security examine-vn2vn
beacon-period 90000
```



**NOTE:** QFX10000 switches do not support FIP snooping and cannot be used as FCoE access transit switches. (QFX10000 switches can be used as FCoE aggregation switches.)

### FCoE Transit Switches TS1 and TS2 ETS Hierarchical Scheduling Configuration

#### Step-by-Step Procedure

To configure the forwarding class set (priority group) and priority group scheduling (in a traffic control profile), and apply the ETS hierarchical scheduling for FCoE traffic to interfaces:

1. Configure the forwarding class set (**fcoe-pg**) for the FCoE traffic:
 

```
[edit class-of-service]
user@switch# set forwarding-class-sets fcoe-pg class fcoe
```
2. Define the traffic control profile (**fcoe-tcp**) to use on the FCoE forwarding class set:
 

```
[edit class-of-service]
user@switch# set traffic-control-profiles fcoe-tcp scheduler-map fcoe-map
guaranteed-rate 3g
user@switch# set traffic-control-profiles fcoe-tcp shaping-rate percent 100
```
3. Apply the FCoE forwarding class set and traffic control profile to the LAG interface and to the FCoE access interfaces:
 

```
[edit class-of-service]
user@switch# set interfaces ae1 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
user@switch# set class-of-service interfaces xe-0/0/30 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
user@switch# set class-of-service interfaces xe-0/0/31 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
user@switch# set class-of-service interfaces xe-0/0/32 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
user@switch# set class-of-service interfaces xe-0/0/33 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
```

### FCoE Transit Switches TS1 and TS2 Port Scheduling Configuration

#### Step-by-Step Procedure

To apply port scheduling for FCoE traffic to interfaces:

1. Apply the scheduler map to the egress ports:
 

```
user@switch# set class-of-service interfaces ae1 scheduler-map fcoe-map
user@switch# set class-of-service interfaces xe-0/0/30 scheduler-map fcoe-map
user@switch# set class-of-service interfaces xe-0/0/31 scheduler-map fcoe-map
user@switch# set class-of-service interfaces xe-0/0/32 scheduler-map fcoe-map
user@switch# set class-of-service interfaces xe-0/0/33 scheduler-map fcoe-map
```

## Results

Display the results of the CoS configuration on MC-LAG Switch S1 and on MC-LAG Switch S2 (the results on both switches are the same). The results are from the ETS hierarchical scheduling configuration, which shows the more complex configuration. Direct port scheduling results would not show the traffic control profile or forwarding class set portions of the configuration, but would display the name of the scheduler map under each interface (instead of the names of the forwarding class set and output traffic control profile). Other than that, they are the same.

```
user@switch> show configuration class-of-service
traffic-control-profiles {
 fcoe-tcp {
 scheduler-map fcoe-map;
 shaping-rate percent 100;
 guaranteed-rate 3000000000;
 }
}
forwarding-class-sets {
 fcoe-pg {
 class fcoe;
 }
}
congestion-notification-profile {
 fcoe-cnp {
 input {
 ieee-802.1 {
 code-point 011 {
 pfc;
 }
 }
 }
 }
}
interfaces {
 ae0 {
 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 congestion-notification-profile fcoe-cnp;
 }
 ae1 {
 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 congestion-notification-profile fcoe-cnp;
 }
}
scheduler-maps {
 fcoe-map {
```

```

 forwarding-class fcoe scheduler fcoe-sched;
 }
}
schedulers {
 fcoe-sched {
 transmit-rate 30000000000;
 shaping-rate percent 100;
 priority low;
 }
}

```



**NOTE:** The forwarding class and classifier configurations are not shown because the show command does not display default portions of the configuration.

For MC-LAG verification commands, see *Example: Configuring Multichassis Link Aggregation*.

Display the results of the CoS configuration on FCoE Transit Switch TS1 and on FCoE Transit Switch TS2 (the results on both transit switches are the same). The results are from the ETS hierarchical port scheduling configuration, which shows the more complex configuration. Direct port scheduling results would not show the traffic control profile or forwarding class set portions of the configuration, but would display the name of the scheduler map under each interface (instead of the names of the forwarding class set and output traffic control profile). Other than that, they are the same.

```

user@switch> show configuration class-of-service
traffic-control-profiles {
 fcoe-tcp {
 scheduler-map fcoe-map;
 shaping-rate percent 100;
 guaranteed-rate 30000000000;
 }
}
forwarding-class-sets {
 fcoe-pg {
 class fcoe;
 }
}
congestion-notification-profile {
 fcoe-cnp {
 input {
 ieee-802.1 {
 code-point 011 {
 pfc;
 }
 }
 }
 }
}
}
interfaces {
 xe-0/0/30 {

```

```
 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 congestion-notification-profile fcoe-cnp;
 }
 xe-0/0/31 {
 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 congestion-notification-profile fcoe-cnp;
 }
 xe-0/0/32 {
 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 congestion-notification-profile fcoe-cnp;
 }
 xe-0/0/33 {
 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 congestion-notification-profile fcoe-cnp;
 }
 ae1 {
 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 congestion-notification-profile fcoe-cnp;
 }
}
scheduler-maps {
 fcoe-map {
 forwarding-class fcoe scheduler fcoe-sched;
 }
}
schedulers {
 fcoe-sched {
 transmit-rate 3000000000;
 shaping-rate percent 100;
 priority low;
 }
}
```



**NOTE:** The forwarding class and classifier configurations are not shown because the `show` command does not display default portions of the configuration.

## Verification

To verify that the CoS components and FIP snooping have been configured and are operating properly, perform these tasks. Because this example uses the default `fcoe` forwarding class and the default IEEE 802.1p trusted classifier, the verification of those configurations is not shown:

- [Verifying That the Output Queue Schedulers Have Been Created on page 6471](#)
- [Verifying That the Priority Group Output Scheduler \(Traffic Control Profile\) Has Been Created \(ETS Configuration Only\) on page 6472](#)
- [Verifying That the Forwarding Class Set \(Priority Group\) Has Been Created \(ETS Configuration Only\) on page 6472](#)
- [Verifying That Priority-Based Flow Control Has Been Enabled on page 6473](#)
- [Verifying That the Interface Class of Service Configuration Has Been Created on page 6474](#)
- [Verifying That the Interfaces Are Correctly Configured on page 6476](#)
- [Verifying That FIP Snooping Is Enabled on the FCoE VLAN on FCoE Transit Switches TS1 and TS2 Access Interfaces on page 6478](#)
- [Verifying That the FIP Snooping Mode Is Correct on FCoE Transit Switches TS1 and TS2 on page 6479](#)

### Verifying That the Output Queue Schedulers Have Been Created

**Purpose** Verify that the output queue scheduler for FCoE traffic has the correct bandwidth parameters and priorities, and is mapped to the correct forwarding class (output queue). Queue scheduler verification is the same on each of the four switches.

**Action** List the scheduler map using the operational mode command `show class-of-service scheduler-map fcoe-map`:

```
user@switch> show class-of-service scheduler-map fcoe-map
Scheduler map: fcoe-map, Index: 9023
```

```
Scheduler: fcoe-sched, Forwarding class: fcoe, Index: 37289
Transmit rate: 3000000000 bps, Rate Limit: none, Buffer size: remainder,
Buffer Limit: none, Priority: low
Excess Priority: unspecified
Shaping rate: 100 percent,
drop-profile-map-set-type: mark
Drop profiles:
 Loss priority Protocol Index Name
 Low any 1 <default-drop-profile>
 Medium high any 1 <default-drop-profile>
 High any 1 <default-drop-profile>
```

**Meaning** The **show class-of-service scheduler-map fcoe-map** command lists the properties of the scheduler map **fcoe-map**. The command output includes:

- The name of the scheduler map (**fcoe-map**)
- The name of the scheduler (**fcoe-sched**)
- The forwarding classes mapped to the scheduler (**fcoe**)
- The minimum guaranteed queue bandwidth (transmit rate **3000000000 bps**)
- The scheduling priority (**low**)
- The maximum bandwidth in the priority group the queue can consume (shaping rate **100 percent**)
- The drop profile loss priority for each drop profile name. This example does not include drop profiles because you do not apply drop profiles to FCoE traffic.

#### Verifying That the Priority Group Output Scheduler (Traffic Control Profile) Has Been Created (ETS Configuration Only)

**Purpose** Verify that the traffic control profile **fcoe-tcp** has been created with the correct bandwidth parameters and scheduler mapping. Priority group scheduler verification is the same on each of the four switches.

**Action** List the FCoE traffic control profile properties using the operational mode command **show class-of-service traffic-control-profile fcoe-tcp**:

```
user@switch> show class-of-service traffic-control-profile fcoe-tcp
Traffic control profile: fcoe-tcp, Index: 18303
 Shaping rate: 100 percent
 Scheduler map: fcoe-map
 Guaranteed rate: 3000000000
```

**Meaning** The **show class-of-service traffic-control-profile fcoe-tcp** command lists all of the configured traffic control profiles. For each traffic control profile, the command output includes:

- The name of the traffic control profile (**fcoe-tcp**)
- The maximum port bandwidth the priority group can consume (shaping rate **100 percent**)
- The scheduler map associated with the traffic control profile (**fcoe-map**)
- The minimum guaranteed priority group port bandwidth (guaranteed rate **3000000000** in bps)

#### Verifying That the Forwarding Class Set (Priority Group) Has Been Created (ETS Configuration Only)

**Purpose** Verify that the FCoE priority group has been created and that the **fcoe** priority (forwarding class) belongs to the FCoE priority group. Forwarding class set verification is the same on each of the four switches.



**Action** List the forwarding class sets using the operational mode command **show class-of-service forwarding-class-set fcoe-pg**:

```
user@switch> show class-of-service forwarding-class-set fcoe-pg
Forwarding class set: fcoe-pg, Type: normal-type, Forwarding class set index:
31420
 Forwarding class Index
 fcoe 1
```

**Meaning** The **show class-of-service forwarding-class-set fcoe-pg** command lists all of the forwarding classes (priorities) that belong to the **fcoe-pg** priority group, and the internal index number of the priority group. The command output shows that the forwarding class set **fcoe-pg** includes the forwarding class **fcoe**.

### Verifying That Priority-Based Flow Control Has Been Enabled

**Purpose** Verify that PFC is enabled on the FCoE code point. PFC verification is the same on each of the four switches.

**Action** List the FCoE congestion notification profile using the operational mode command **show class-of-service congestion-notification fcoe-cnp**:

```
user@switch> show class-of-service congestion-notification fcoe-cnp
Type: Input, Name: fcoe-cnp, Index: 6879
Cable Length: 100 m
 Priority PFC MRU
 000 Disabled
 001 Disabled
 010 Disabled
 011 Enabled 2500
 100 Disabled
 101 Disabled
 110 Disabled
 111 Disabled
Type: Output
 Priority Flow-Control-Queues
 000
 001 0
 010 1
 011 2
 100 3
 101 4
 110 5
 111 6
 111 7
```

**Meaning** The **show class-of-service congestion-notification fcoe-cnp** command lists all of the IEEE 802.1p code points in the congestion notification profile that have PFC enabled. The command output shows that PFC is enabled on code point **011** (**fcoe** queue) for the **fcoe-cnp** congestion notification profile.

The command also shows the default cable length (100 meters), the default maximum receive unit (2500 bytes), and the default mapping of priorities to output queues because this example does not include configuring these options.

### Verifying That the Interface Class of Service Configuration Has Been Created

---

**Purpose** Verify that the CoS properties of the interfaces are correct. The verification output on MC-LAG Switches S1 and S2 differs from the output on FCoE Transit Switches TS1 and TS2.



**NOTE:** The output is from the ETS hierarchical port scheduling configuration to show the more complex configuration. Direct port scheduling results do not show the traffic control profile or forwarding class sets because those elements are configured only for ETS. Instead, the name of the scheduler map is displayed under each interface.

---

**Action** List the interface CoS configuration on MC-LAG Switches S1 and S2 using the operational mode command **show configuration class-of-service interfaces**:

```
user@switch> show configuration class-of-service interfaces
ae0 {
 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 congestion-notification-profile fcoe-cnp;
}

ae1 {
 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 congestion-notification-profile fcoe-cnp;
}
```

List the interface CoS configuration on FCoE Transit Switches TS1 and TS2 using the operational mode command **show configuration class-of-service interfaces**:

```
user@switch> show configuration class-of-service interfaces
xe-0/0/30 {
 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 congestion-notification-profile fcoe-cnp;
}
xe-0/0/31 {
 forwarding-class-set {
 fcoe-pg {
```

```

 output-traffic-control-profile fcoe-tcp;
 }
}
congestion-notification-profile fcoe-cnp;
}
xe-0/0/32 {
 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 congestion-notification-profile fcoe-cnp;
}
xe-0/0/33 {
 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 congestion-notification-profile fcoe-cnp;
}
ae1 {
 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 congestion-notification-profile fcoe-cnp;
}

```

**Meaning** The **show configuration class-of-service interfaces** command lists the class of service configuration for all interfaces. For each interface, the command output includes:

- The name of the interface (for example, **ae0** or **xe-0/0/30**)
- The name of the forwarding class set associated with the interface (**fcoe-pg**)
- The name of the traffic control profile associated with the interface (output traffic control profile, **fcoe-tcp**)
- The name of the congestion notification profile associated with the interface (**fcoe-cnp**)



**NOTE:** Interfaces that are members of a LAG are not shown individually. The LAG or MC-LAG CoS configuration is applied to all interfaces that are members of the LAG or MC-LAG. For example, the interface CoS configuration output on MC-LAG Switches S1 and S2 shows the LAG CoS configuration but does not show the CoS configuration of the member interfaces separately. The interface CoS configuration output on FCoE Transit Switches TS1 and TS2 shows the LAG CoS configuration but also shows the configuration for interfaces xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33, which are not members of a LAG.

### Verifying That the Interfaces Are Correctly Configured

---

**Purpose** Verify that the LAG membership, MTU, VLAN membership, and port mode of the interfaces are correct. The verification output on MC-LAG Switches S1 and S2 differs from the output on FCoE Transit Switches T1 and T2.

**Action** List the interface configuration on MC-LAG Switches S1 and S2 using the operational mode command **show configuration interfaces**:

```
user@switch> show configuration interfaces
```

```
xe-0/0/10 {
 ether-options {
 802.3ad ae0;
 }
}
xe-0/0/11 {
 ether-options {
 802.3ad ae0;
 }
}
xe-0/0/20 {
 ether-options {
 802.3ad ae1;
 }
}
xe-0/0/21 {
 ether-options {
 802.3ad ae1;
 }
}
ae0 {
 mtu 2180;
 unit 0 {
 family ethernet-switching {
 interface-mode trunk;
 vlan {
 members fcoe_vlan;
 }
 }
 }
}
ae1 {
 mtu 2180;
 unit 0 {
 family ethernet-switching {
 interface-mode trunk;
 vlan {
 members fcoe_vlan;
 }
 }
 }
}
```

List the interface configuration on FCoE Transit Switches TS1 and TS2 using the operational mode command **show configuration interfaces**:

```
user@switch> show configuration interfaces
```

```
xe-0/0/25 {
 ether-options {
 802.3ad ae1;
 }
}
xe-0/0/26 {
 ether-options {
 802.3ad ae1;
 }
}
xe-0/0/30 {
 mtu 2180;
 unit 0 {
 family ethernet-switching {
 interface-mode trunk;
 vlan {
 members fcoe_vlan;
 }
 }
 }
}
xe-0/0/31 {
 mtu 2180;
 unit 0 {
 family ethernet-switching {
 interface-mode trunk;
 vlan {
 members fcoe_vlan;
 }
 }
 }
}
xe-0/0/32 {
 mtu 2180;
 unit 0 {
 family ethernet-switching {
 interface-mode trunk;
 vlan {
 members fcoe_vlan;
 }
 }
 }
}
xe-0/0/33 {
 mtu 2180;
 unit 0 {
 family ethernet-switching {
 interface-mode trunk;
 vlan {
 members fcoe_vlan;
 }
 }
 }
}
ae1 {
 mtu 2180;
 unit 0 {
 family ethernet-switching {
 interface-mode trunk;
 vlan {
```

```
 members fcoe_vlan;
 }
}
}
```

**Meaning** The **show configuration interfaces** command lists the configuration of each interface by interface name.

For each interface that is a member of a LAG, the command lists only the name of the LAG to which the interface belongs.

For each LAG interface and for each interface that is not a member of a LAG, the command output includes:

- The MTU (**2180**)
- The unit number of the interface (**0**)
- The interface mode (**trunk** mode both for interfaces that connect two switches and for interfaces that connect to FCoE hosts)
- The name of the VLAN in which the interface is a member (**fcoe\_vlan**)

#### Verifying That FIP Snooping Is Enabled on the FCoE VLAN on FCoE Transit Switches TS1 and TS2 Access Interfaces

---

**Purpose** Verify that FIP snooping is enabled on the FCoE VLAN access interfaces. FIP snooping is enabled only on the FCoE access interfaces, so it is enabled only on FCoE Transit Switches TS1 and TS2. FIP snooping is not enabled on MC-LAG Switches S1 and S2 because FIP snooping is done at the Transit Switch TS1 and TS2 FCoE access ports.

**Action** List the port security configuration on FCoE Transit Switches TS1 and TS2 using the operational mode command **show configuration vlans fcoe\_vlan forwarding-options fip-security**:

```
user@switch> show configuration vlans fcoe_vlan forwarding-options fip-security
interface ae1.0 {
 fcoe-trusted;
}
examine-vn2vn {
 beacon-period 90000;
}
```

**Meaning** The **show configuration vlans fcoe\_vlan forwarding-options fip-security** command lists VLAN FIP security information, including whether a port member of the VLAN is trusted. The command output shows that:

- LAG port **ae1.0**, which connects the FCoE transit switch to the MC-LAG switches, is configured as an FCoE trusted interface. FIP snooping is not performed on the member interfaces of the LAG (**xe-0/0/25** and **xe-0/0/26**).
- VN2VN\_Port FIP snooping is enabled (**examine-vn2vn**) on the FCoE VLAN and the beacon period is set to 90000 milliseconds. On Transit Switches TS1 and TS2, all

interface members of the FCoE VLAN perform FIP snooping unless the interface is configured as FCoE trusted. On Transit Switches TS1 and TS2, interfaces **xe-0/0/30**, **xe-0/0/31**, **xe-0/0/32**, and **xe-0/0/33** perform FIP snooping because they are not configured as FCoE trusted. The interface members of LAG **ae1** (**xe-0/0/25** and **xe-0/0/26**) do not perform FIP snooping because the LAG is configured as FCoE trusted.

### Verifying That the FIP Snooping Mode Is Correct on FCoE Transit Switches TS1 and TS2

**Purpose** Verify that the FIP snooping mode is correct on the FCoE VLAN. FIP snooping is enabled only on the FCoE access interfaces, so it is enabled only on FCoE Transit Switches TS1 and TS2. FIP snooping is not enabled on MC-LAG Switches S1 and S2 because FIP snooping is done at the Transit Switch TS1 and TS2 FCoE access ports.

**Action** List the FIP snooping configuration on FCoE Transit Switches TS1 and TS2 using the operational mode command **show fip snooping brief**:

```
user@switch> show fip snooping brief
VLAN: fcoe_vlan, Mode: VN2VN Snooping
FC-MAP: 0e:fc:00
...
```



**NOTE:** The output has been truncated to show only the relevant information.

**Meaning** The **show fip snooping brief** command lists FIP snooping information, including the FIP snooping VLAN and the FIP snooping mode. The command output shows that:

- The VLAN on which FIP snooping is enabled is **fcoe\_vlan**
- The FIP snooping mode is VN2VN\_Port FIP snooping (**VN2VN Snooping**)

#### Related Documentation

- *Example: Configuring Multichassis Link Aggregation*
- [Configuring Link Aggregation on page 2869](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 6504](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 6771](#)
- [Example: Configuring Queue Schedulers for Port Scheduling on page 6744](#)
- *Understanding Multichassis Link Aggregation*
- [Understanding MC-LAGs on an FCoE Transit Switch on page 6451](#)

## Understanding FCoE and FIP Session High Availability

---

High availability features maintain storage network sessions when a system process is terminated and during certain types of upgrades:

- [High Availability for Fibre Channel Process Termination \(FCoE-FC Gateway Mode, QFX3500 Only\)](#) on page 6480
- [High Availability for FIP Snooping](#) on page 6480
- [Nonstop Software Upgrade \(QFabric Systems\)](#) on page 6481

### High Availability for Fibre Channel Process Termination (FCoE-FC Gateway Mode, QFX3500 Only)

In FCoE-FC gateway mode, the QFX3500 switch provides high availability to restore the FCoE sessions running on the switch in case the Fibre Channel (FC) process is terminated. A session is a fabric login (FLOGI) or fabric discovery (FDISC) login to the FC SAN fabric, not an end-to-end server-to-storage session.

The switch stores FCoE session data in a persistent storage module. If the FC process terminates, the switch restores the existing FCoE sessions on the same interfaces that they were on before the FC process terminated. Data traffic for existing sessions is not affected during session restoration.

For a brief time, the system does not process control traffic because of the FC process restart and session restoration. During this brief time, no new FCoE sessions can be established, and no existing sessions can log out.



**NOTE:** During the restoration process, if the FC process does not receive an *interface up* notification from a particular interface within a certain time, the switch times out the restore operation and discards the data on that interface. The previously existing FCoE sessions on that interface are not restored, and the ENodes must log in again.

---



**NOTE:** An FC process restart and session restoration resets the Fibre Channel statistics.

---

If the FC process terminates repeatedly, the operating system disables the process until you manually restart it. To restart the FC process manually, issue the **restart fibre-channel** command.

### High Availability for FIP Snooping

You can configure the system to perform FIP snooping on Ethernet interfaces that are connected to FCoE devices that have ENodes. The high availability function restores running FIP snooping sessions in case the Ethernet switching process is terminated.





**NOTE:** Juniper Networks QFX10000 switches do not support FIP snooping. (Aggregation devices do not need to enable FIP snooping because FIP snooping is performed at the FCoE access edge.)

The Ethernet switching process stores the FIP snooping state in a persistent storage module. If the Ethernet switching process terminates, the switch restores the existing FIP snooping sessions on the same interfaces that they were on before the Ethernet switching process terminated. The high availability features preserve:

- Logged in ENodes
- Discovered FCFs
- Existing sessions
- Existing FIP snooping filters

The complete restoration process, including reconciling all valid states, takes a maximum of 8 seconds. During the restoration process, the switch can learn a new FCF or a new FC switch, and new ENodes can log in to the FC network. However, FDISC messages from an ENode that is already logged in to the network might be dropped if the ENode has not yet been restored.

When the Ethernet switching process terminates ungracefully, the FIP keepalive timer is reset to the normal initial value, not the value at the time of the Ethernet switching process termination.

In the event of an Ethernet switching process termination, ENodes remain logged in, and existing sessions are not interrupted.



**NOTE:** An Ethernet switching process restart and session restoration resets the FIP snooping statistics.

## Nonstop Software Upgrade (QFabric Systems)

On QFabric system Node groups that have more than one Node device, nonstop software upgrade (NSSU) enables you to upgrade the Node devices with minimal packet loss and maximum uptime. NSSU automates software upgrades on the QFabric system components in an orderly and consistent manner to maximize system uptime.

The system upgrades components with redundant architectures, such as redundant server Node groups and network Node groups that have two or more members, in stages. While the system upgrades one component, the redundant component continues to function.

For example, while one member of a redundant server Node group is upgraded, the other member continues to forward traffic. When the first Node group member completes the upgrade, it comes online while the system upgrades the second member.

NSSU provides high availability for the lossless traffic forwarding required to support storage networks. If your system design includes redundancy (redundant Node devices in Node groups, LAGs, and so on) so that an alternate traffic path is available, when you upgrade a Node device, traffic is not impacted.

In fully redundant topologies, NSSU preserves FIP session, FIP snooping filter, VN2VF\_Port session, and VN2VN\_Port session information and prevents traffic loss in most cases. An exception is that Node devices that are directly connected to ENodes experience momentary traffic loss when the Node device reboots.

**Related Documentation**

- [Understanding FCoE on page 6436](#)

---

## Troubleshooting Dropped FIP Traffic

---

- |                 |                                                                                                                                                                                                                                                                                                      |
|-----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Problem</b>  | <b>Description:</b> Fibre Channel over Ethernet (FCoE) Initialization Protocol (FIP) traffic such as FIP VLAN discovery and notification frames is dropped.                                                                                                                                          |
| <b>Cause</b>    | The interface on which the FIP traffic is dropped does not have a native VLAN configured. FIP VLAN discovery and notification messages are exchanged as untagged packets on the native VLAN. (After the FCoE session with the Fibre Channel switch is established, FCoE traffic uses the FCoE VLAN.) |
| <b>Solution</b> | Check to ensure that every 10-Gigabit Ethernet interface that connects to an FCoE device includes a native VLAN. Configure a native VLAN on all 10-Gigabit Ethernet interfaces that connect to FCoE devices.                                                                                         |



**NOTE:** Make sure that the native VLAN you are using is the same native VLAN that the FCoE devices use for Ethernet traffic.

---

The procedure for configuring a native VLAN on an interface is different on switches that use the original CLI than on switches that use the Enhanced Layer 2 Software (ELS) CLI. This topic provides the configuration procedure for each CLI.

### Configuring a Native VLAN on Switches Using the Original CLI

To configure a native VLAN on an interface:

1. Set the interface port mode to **tagged-access** if you have not already done so:

```
[edit]
user@switch# set interfaces interface unit unit family ethernet-switching port-mode
tagged-access
```

For example, to set the port mode to **tagged-access** for interface **xe-0/0/6.0**:

```
[edit]
user@switch# set interfaces xe-0/0/6 unit 0 family ethernet-switching port-mode
tagged-access
```

2. Configure the native VLAN if it does not already exist:

```
[edit]
user@switch# set vlans vlan-name vlan-id vlan-id
```

For example, to name the native VLAN **native** and use the VLAN ID 1:

```
[edit]
user@switch# set vlans native vlan-id 1
```

3. Configure the native VLAN on the interface:

```
[edit]
user@switch# set interfaces interface unit unit family ethernet-switching native-vlan-id
vlan-id
```

For example, to configure a native VLAN with the VLAN ID 1 on interface **xe-0/0/6.0**:

```
[edit]
user@switch# set interfaces xe-0/0/6 unit 0 family ethernet-switching native-vlan-id 1
```

### Configuring a Native VLAN on Switches Using the ELS CLI

To configure a native VLAN on an interface:

1. Set the interface mode to **trunk** if you have not already done so:

```
[edit]
user@switch# set interfaces interface unit unit family ethernet-switching interface-mode
trunk
```

For example, to set the interface mode to **trunk** for interface **xe-0/0/6.0**:

```
[edit]
user@switch# set interfaces xe-0/0/6 unit 0 family ethernet-switching interface-mode trunk
```

2. Configure the native VLAN if it does not already exist:

```
[edit]
user@switch# set vlans vlan-name vlan-id vlan-id
```

For example, to name the native VLAN **native** and use the VLAN ID 1:

```
[edit]
user@switch# set vlans native vlan-id 1
```

3. Configure the native VLAN on the physical Ethernet interface:

```
[edit]
user@switch# set interfaces interface native-vlan-id vlan-id
```

For example, to configure a native VLAN with the VLAN ID 1 on interface **xe-0/0/6.0**:

```
[edit]
user@switch# set interfaces xe-0/0/6 native-vlan-id 1
```

4. Configure the Ethernet interface as a member of the native VLAN:

```
[edit]
user@switch# set interfaces interface unit unit family ethernet-switching vlan members
vlan-name
```

For example, to configure an Ethernet interface as a member of a native VLAN with the VLAN ID 1 on interface **xe-0/0/6.0**:

```
[edit]
```

```
user@switch# set interfaces xe-0/0/6 unit 0 family ethernet-switching vlan members native
```

**Related  
Documentation**

- [interfaces on page 3010](#)
- [vlans](#)
- [Configuring VLANs for FCoE Traffic on an FCoE Transit Switch on page 6442](#)

---

## Understanding OxID Hash Control for FCoE Traffic Load Balancing on Standalone Switches

---

The originator exchange identifier (OxID) field is one of several fields that the switch can use in its hash function computation for FCoE traffic load balancing over multiple outgoing links in an Ethernet link aggregation group (LAG) on ports that face an FCoE forwarder (FCF). The originator of an exchange between a pair of Fibre Channel (FC) endpoints (such as an FCoE host and an FC storage device) uses the OxID field as an identifier for that exchange. The originator also uses the OxID field to track the progress of the series of sequences that comprise the exchange.

When FCoE traffic traverses a LAG that faces an FCF, it can take multiple different links between the source and destination endpoints. The idea is to distribute the FCoE traffic across the FCF-facing LAG links, thus balancing the link load. The switch creates a hash value from some of the packet header fields, and uses the hash value to assign each packet to one of the LAG links. The switch always uses five packet header fields to compute the hash value:

- Source ID (SID)
- Destination ID (DID)
- Fabric ID (FID)
- Source Port ID (SPID)
- Source Module ID (SMID)

In addition, the OxID field is included by default in the FCoE load-balancing hash computation. However, if you do not want to use the OxID field in the FCoE load-balancing hash computation, you can remove it from the computation by using the **set forwarding-options hash-key family fcoe oxid disable** command.

Including the OxID field in the load-balancing hash computation allows different exchanges between a pair of Fibre Channel (FC) endpoints (such as an FCoE host and an FC storage device) to take different paths across the network, thus improving the aggregate network throughput.

However, if the paths between different sets of FC endpoints have common links, congestion on one set of FC endpoints can affect the other set of endpoints. Such congestion can happen if the FCoE traffic on the two sets of endpoints uses the same priority (IEEE 802.1p code point). It is common for networks to use priority 3 (IEEE 802.1p code point 011) for FCoE traffic. However, you can assign different IEEE priorities to

different lossless FCoE flows as described in [“Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows” on page 6873](#) to further separate the traffic flows.

**Related  
Documentation**

- [Enabling and Disabling CoS OxID Hash Control on Standalone Switches on page 6485](#)

## Enabling and Disabling CoS OxID Hash Control on Standalone Switches

The originator exchange identifier (OxID) field is one of several fields that the switch can use in its hash function computation for FCoE traffic load balancing over multiple outgoing links in an Ethernet link aggregation group (LAG) on ports that face an FCoE forwarder (FCF). You can configure whether or not the switch uses the OxID in the hash computation.

Including the OxID field in the load-balancing hash computation allows different exchanges between a pair of Fibre Channel (FC) endpoints (such as an FCoE host and an FC storage device) to take different paths across the network, thus improving the aggregate network throughput.

However, if the paths between different sets of FC endpoints have common links, congestion on one set of FC endpoints can affect the other set of endpoints. Such congestion can happen if the FCoE traffic on the two sets of endpoints uses the same priority (IEEE 802.1p code point). It is common for networks to use priority 3 (IEEE 802.1p code point 011) for FCoE traffic. However, you can assign different IEEE priorities to different lossless FCoE flows as described in [“Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows” on page 6873](#) to further separate the traffic flows.

OxID hash control is enabled by default.

- To enable OxID hash control field for FCoE traffic load balancing:

```
[edit forwarding-options hash-key]
user@switch# set family fcoe oxid enable
```

- To disable OxID hash control field for FCoE traffic load balancing:

```
[edit forwarding-options hash-key]
user@switch# set family fcoe oxid disable
```

**Related  
Documentation**

- [Understanding OxID Hash Control for FCoE Traffic Load Balancing on Standalone Switches on page 6484](#)



## PART 88

# Configuring DCBX and PFC

- [Using DCBX and PFC on page 6489](#)





## CHAPTER 246

# Using DCBX and PFC

- [Understanding DCB Features and Requirements on page 6490](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 6493](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 6504](#)
- [Understanding DCBX on page 6514](#)
- [Configuring the DCBX Mode on page 6523](#)
- [Configuring DCBX Autonegotiation on page 6524](#)
- [Disabling the ETS Recommendation TLV on page 6527](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 6528](#)
- [Defining an Application for DCBX Application Protocol TLV Exchange on page 6532](#)
- [Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 6533](#)
- [Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 6534](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 6535](#)

## Understanding DCB Features and Requirements

---

Data center bridging (DCB) is a set of enhancements to the IEEE 802.1 bridge specifications. DCB modifies and extends Ethernet behavior to support I/O convergence in the data center. I/O convergence includes but is not limited to the transport of Ethernet LAN traffic and Fibre Channel (FC) storage area network (SAN) traffic on the same physical Ethernet network infrastructure.



Video: [What is Data Center Bridging?](#)

A converged architecture saves cost by reducing the number of networks and switches required to support both types of traffic, reducing the number of interfaces required, reducing cable complexity, and reducing administration activities.

The Juniper Networks QFX Series and EX4600 switches support the DCB features required to transport converged Ethernet and FC traffic while providing the class-of-service (CoS) and other characteristics FC requires for transmitting storage traffic. To accommodate FC traffic, DCB specifications provide:

- A flow control mechanism called priority-based flow control (PFC, described in IEEE 802.1Qbb) to help provide lossless transport.
- A discovery and exchange protocol for conveying configuration and capabilities among neighbors to ensure consistent configuration across the network, called Data Center Bridging Capability Exchange protocol (DCBX), which is an extension of Link Layer Data Protocol (LLDP, described in IEEE 802.1AB).
- A bandwidth management mechanism called enhanced transmission selection (ETS, described in IEEE 802.1Qaz).
- A congestion management mechanism called quantized congestion notification (QCN, described in IEEE 802.1Qau).

The switch supports the PFC, DCBX, and ETS standards but does not support QCN. The switch also provides the high-bandwidth interfaces (10-Gbps minimum) required to support DCB and converged traffic.

This topic describes the DCB standards and requirements the switch supports:

- [Lossless Transport on page 6490](#)
- [ETS on page 6491](#)
- [DCBX on page 6492](#)

### Lossless Transport

FC traffic requires lossless transport (defined as no frames dropped because of congestion). Standard Ethernet does not support lossless transport, but the DCB extensions to Ethernet along with proper buffer management enable an Ethernet network to provide the level of class of service (CoS) necessary to transport FC frames encapsulated in Ethernet over an Ethernet network.

This section describes these factors in creating lossless transport over Ethernet:

- [PFC on page 6491](#)
- [Buffer Management on page 6491](#)
- [Physical Interfaces on page 6491](#)

## PFC

---

PFC is a link-level flow control mechanism similar to Ethernet PAUSE (described in IEEE 802.3x). Ethernet PAUSE stops all traffic on a link for a period of time. PFC enables you to divide traffic on a link into eight priorities and stop the traffic of a selected priority without stopping the traffic assigned to other priorities on the link.

Pausing the traffic of a selected priority enables you to provide lossless transport for traffic assigned that priority and at the same time use standard lossy Ethernet transport for the rest of the link traffic.

## Buffer Management

---

Buffer management is critical to the proper functioning of PFC, because if buffers are allowed to overflow, frames are dropped and transport is not lossless.

For each lossless flow priority, the switch requires sufficient buffer space to:

- Store frames sent during the time it takes to send the PFC pause frame across the cable between devices.
- Store the frames that are already on the wire when the sender receives the PFC pause frame.

The propagation delay due to cable length and speed, as well as processing speed, determines the amount of buffer space needed to prevent frame loss due to congestion.

The switch automatically sets the threshold for sending PFC pause frames to accommodate delay from cables as long as 150 meters (492 feet) and to accommodate large frames that might be on the wire when the switch sends the pause frame. This ensures that the switch sends pause frames early enough to allow the sender to stop transmitting before the receive buffers on the switch overflow.

## Physical Interfaces

---

QFX Series switches support 10-Gbps or faster, full-duplex interfaces. The switch enables DCB capability only on 10-Gbps or faster Ethernet interfaces.

## ETS

PFC divides traffic into up to eight separate streams (priorities, configured on the switch as forwarding classes) on a physical link. ETS enables you to manage the link bandwidth by:

- Grouping the priorities into priority groups (configured on the switch as forwarding class sets).

- Specifying the bandwidth available to each of the priority groups as a percentage of the total available link bandwidth.
- Allocating the bandwidth to the individual priorities in the priority group.

The available link bandwidth is the bandwidth remaining after servicing strict-high priority queues. On QFX5200, QFX5100, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems, we recommend that you always configure a shaping rate to limit the amount of bandwidth a strict-high priority queue can consume by including the [shaping-rate](#) statement in the [\[edit class-of-service schedulers\]](#) hierarchy on the strict-high priority scheduler. This prevents a strict-high priority queue from starving other queues on the port. (On QFX10000 switches, configure a transmit rate on strict-high priority queues to set a maximum amount of bandwidth for strict-high priority traffic.)

Managing link bandwidth with ETS provides several advantages:

- There is uniform management of all types of traffic on the link, both congestion-managed traffic and standard Ethernet traffic.
- When a priority group does not use all of its allocated bandwidth, other priority groups on the link can use that bandwidth as needed.

When a priority in a priority group does not use all of its allocated bandwidth, other priorities in the group can use that bandwidth.

The result is better bandwidth utilization, because priorities that consist of bursty traffic can share bandwidth during periods of low traffic transmission instead of consuming their entire bandwidth allocation when traffic loads are light.

- You can assign traffic types with different service needs to different priorities so that each traffic type receives appropriate treatment.
- Strict priority traffic retains its allocated bandwidth.

## DCBX

DCB devices use DCBX to exchange configuration information with directly connected peers (switches and endpoints such as servers). DCBX is an extension of LLDP. If you disable LLDP on an interface, that interface cannot run DCBX. If you attempt to enable DCBX on an interface on which LLDP is disabled, the configuration commit fails.

DCBX can:

- Discover the DCB capabilities of peers.
- Detect DCB feature misconfiguration or mismatches between peers.
- Configure DCB features on peers.

You can configure DCBX operation for PFC, ETS, and for Layer 2 and Layer 4 applications such as FCoE and iSCSI. DCBX is enabled or disabled on a per-interface basis.

### Related Documentation

- [Understanding FCoE on page 6436](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 6765](#)

- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 6493](#)
- [Understanding DCBX on page 6514](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 6504](#)

## Understanding CoS Flow Control (Ethernet PAUSE and PFC)

Flow control supports lossless transmission by regulating traffic flows to avoid dropping frames during periods of congestion. Flow control stops and resumes the transmission of network traffic between two connected peer nodes on a full-duplex Ethernet physical link. Controlling the flow by pausing and restarting it prevents buffers on the nodes from overflowing and dropping frames. You configure flow control on a per-interface basis.

Two methods of peer-to-peer flow control are supported:

- IEEE 802.3X Ethernet PAUSE



**NOTE:** QFX10000 switches do not support Ethernet PAUSE. Information about Ethernet PAUSE does not apply to QFX10000 switches.

OCX Series switches support symmetric Ethernet PAUSE flow control on Layer 3 tagged interfaces. OCX Series switches do not support asymmetric Ethernet PAUSE flow control. Information about asymmetric flow control does not apply to OCX Series switches.

- IEEE 802.1Qbb priority-based flow control (PFC)



**NOTE:** OCX Series switches do not support PFC or lossless Layer 2 transport. Information about PFC, lossless transport, and congestion notification profiles do not apply to OCX Series switches.



Video: [Why Use PFC in a Data Center Network?](#)

- [General Information about Ethernet PAUSE and PFC and When to Use Them on page 6493](#)
- [Ethernet PAUSE on page 6494](#)
- [PFC on page 6499](#)
- [Lossless Transport Support Summary on page 6502](#)

## General Information about Ethernet PAUSE and PFC and When to Use Them

Ethernet PAUSE and PFC are link-level flow control mechanisms.



**NOTE:** For end-to-end congestion control for best-effort traffic, see [“Understanding CoS Explicit Congestion Notification” on page 6820](#).

Ethernet PAUSE pauses transmission of all traffic on a physical Ethernet link.

PFC decouples the pause function from the physical Ethernet link and enables you to divide traffic on one link into eight priorities. You can think of the eight priorities as eight “lanes” of traffic that are mapped to forwarding classes and output queues. Each priority maps to a 3-bit IEEE 802.1p CoS code point value in the VLAN header. You can enable PFC on one or more priorities (IEEE 802.1p code points) on a link. When PFC-enabled traffic is paused on a link, traffic that is not PFC-enabled continues to flow (or is dropped if congestion is severe enough).

Use Ethernet PAUSE when you want to prevent packet loss on all of the traffic on a link. Use PFC to prevent traffic loss only on a specified type of traffic that require lossless treatment, for example, Fibre Channel over Ethernet (FCoE) traffic.



**NOTE:** Depending on the amount of traffic on a link or assigned to a priority, pausing traffic can cause ingress port congestion and spread congestion through the network.

Ethernet PAUSE and PFC are mutually exclusive configurations on an interface. Attempting to configure both Ethernet PAUSE and PFC on a link causes a commit error.

By default, all forms of flow control are disabled. You must explicitly enable flow control on interfaces to pause traffic.

## Ethernet PAUSE

Ethernet PAUSE is a congestion relief feature that works by providing link-level flow control for all traffic on a full-duplex Ethernet link. Ethernet PAUSE works in both directions on the link. In one direction, an interface generates and sends Ethernet PAUSE messages to stop the connected peer from sending more traffic. In the other direction, the interface responds to Ethernet PAUSE messages it receives from the connected peer to stop sending traffic.



**NOTE:** QFX10000 switches do not support Ethernet PAUSE. Information about Ethernet PAUSE does not apply to QFX10000 switches.

OCX Series switches support symmetric Ethernet PAUSE flow control on Layer 3 tagged interfaces. OCX Series switches do not support asymmetric Ethernet PAUSE flow control. Information about asymmetric flow control does not apply to OCX Series switches.

Ethernet PAUSE also works on aggregated Ethernet interfaces. For example, if the connected peer interfaces are called Node A and Node B:

- When the receive buffers on interface Node A reach a certain level of fullness, the interface generates and sends an Ethernet PAUSE message to the connected peer (interface Node B) to tell the peer to stop sending frames. The Node B buffers store frames until the time period specified in the Ethernet PAUSE frame elapses; then Node B resumes sending frames to Node A.
- When interface Node A receives an Ethernet PAUSE message from interface Node B, interface Node A stops transmitting frames until the time period specified in the Ethernet PAUSE frame elapses; then Node A resumes transmission. (The Node A transmit buffers store frames until Node A resumes sending frames to Node B.)

In this scenario, if Node B sends an Ethernet PAUSE frame with a time value of 0 to Node A, the 0 time value indicates to Node A that it can resume transmission. This happens when the Node B buffer empties to below a certain threshold and the buffer can once again accept traffic.

*Symmetric flow control* means an interface has the same Ethernet PAUSE configuration in both directions. The Ethernet PAUSE generation and Ethernet PAUSE response functions are both configured as enabled, or they are both disabled. You configure symmetric flow control by including the **flow-control** statement at the **[edit interfaces interface-name ether-options]** hierarchy level.

*Asymmetric flow control* allows you to configure the Ethernet PAUSE functionality in each direction independently on an interface. The configuration for generating Ethernet PAUSE messages and for responding to Ethernet PAUSE messages does not have to be the same. It can be enabled in both directions, disabled in both directions, or enabled in one direction and disabled in the other direction. You configure asymmetric flow control by including the **configured-flow-control** statement at the **[edit interfaces interface-name ether-options]** hierarchy level.

On any particular interface, symmetric and asymmetric flow control are mutually exclusive. Asymmetric flow control overrides and disables symmetric flow control. (If PFC is configured on an interface, you cannot commit an Ethernet PAUSE configuration on the interface. Attempting to commit an Ethernet PAUSE configuration on an interface with PFC enabled on one or more queues results in a commit error. To commit the PAUSE configuration, you must first delete the PFC configuration.) Both symmetric and asymmetric flow control are supported.

- [Symmetric Flow Control on page 6495](#)
- [Asymmetric Flow Control on page 6496](#)

### Symmetric Flow Control

Symmetric flow control configures both the receive and transmit buffers in the same state. The interface can both send Ethernet PAUSE messages and respond to them (flow control is enabled), or the interface cannot send Ethernet PAUSE messages or respond to them (flow control is disabled).

When you enable symmetric flow control on an interface, the Ethernet PAUSE behavior depends on the configuration of the connected peer. With symmetric flow control enabled, the interface can perform any Ethernet PAUSE functions that the connected peer can

perform. (When symmetric flow control is disabled, the interface does not send or respond to Ethernet PAUSE messages.)

### Asymmetric Flow Control

Asymmetric flow control enables you to specify independently whether or not the interface receive buffer generates and sends Ethernet PAUSE messages to stop the connected peer from transmitting traffic, and whether or not the interface transmit buffer responds to Ethernet PAUSE messages it receives from the connected peer and stops transmitting traffic. The receive buffer configuration determines if the interface transmits Ethernet PAUSE messages, and the transmit buffer configuration determines if the interface receives and responds to Ethernet PAUSE messages:

- Receive buffers on—Enable Ethernet PAUSE transmission (generate and send Ethernet PAUSE frames)
- Transmit buffers on—Enable Ethernet PAUSE reception (respond to received Ethernet PAUSE frames)

You must explicitly set the flow control for both the receive buffer and the transmit buffer (**on** or **off**) to configure asymmetric Ethernet PAUSE. [Table 513](#) describes the configured flow control state when you set the receive (Rx) and transmit (Tx) buffers on an interface:

**Table 513: Asymmetric Ethernet PAUSE Flow Control Configuration**

| Receive (Rx) Buffer | Transmit (Tx) Buffer | Configured Flow Control State                                                                                                                                                                                         |
|---------------------|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| On                  | Off                  | Interface generates and sends Ethernet PAUSE messages. Interface does not respond to Ethernet PAUSE messages (interface continues to transmit even if peer requests that the interface stop sending traffic).         |
| Off                 | On                   | Interface responds to Ethernet PAUSE messages received from the connected peer, but does not generate or send Ethernet PAUSE messages. (The interface does not request that the connected peer stop sending traffic.) |
| On                  | On                   | Same functionality as symmetric Ethernet PAUSE. Interface generates and sends Ethernet PAUSE messages and responds to received Ethernet PAUSE messages.                                                               |
| Off                 | Off                  | Ethernet PAUSE flow control is disabled.                                                                                                                                                                              |

The configured flow control is the Ethernet PAUSE state configured on the interface.

On 1-Gigabit Ethernet interfaces, autonegotiation of Ethernet PAUSE with the connected peer is supported. (Autonegotiation on 10-Gigabit Ethernet interfaces is not supported.) Autonegotiation enables the interface to exchange state advertisements with the connected peer so that the two devices can agree on the Ethernet PAUSE configuration. Each interface advertises its flow control state to the connected peer using a combination of the Ethernet PAUSE and ASM\_DIR bits, as described in [Table 514](#):



**Table 514: Flow Control State Advertised to the Connected Peer (Autonegotiation)**

| Rx Buffer State | Tx Buffer State | PAUSE Bit | ASM_DIR Bit | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-----------------|-----------------|-----------|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Off             | Off             | 0         | 0           | The interface advertises no Ethernet PAUSE capability. This is equivalent to disabling flow control on an interface.                                                                                                                                                                                                                                                                                                                                                                                       |
| On              | On              | 1         | 0           | The interface advertises symmetric flow control (both the transmission of Ethernet PAUSE messages and the ability to receive and respond to Ethernet PAUSE messages).                                                                                                                                                                                                                                                                                                                                      |
| On              | Off             | 0         | 1           | The interface advertises asymmetric flow control (the transmission of Ethernet PAUSE messages, but not the ability to receive and respond to Ethernet PAUSE messages).                                                                                                                                                                                                                                                                                                                                     |
| Off             | On              | 1         | 1           | The interface advertises both symmetric and asymmetric flow control. Although the interface does not generate and send Ethernet PAUSE requests to the peer, the interface supports both symmetric and asymmetric Ethernet PAUSE configuration on the peer because the peer is not affected if the peer does not receive Ethernet PAUSE requests. (If the interface responds to the peer's Ethernet PAUSE requests, that is sufficient to support either symmetric or asymmetric flow control on the peer.) |

The flow control configuration on each switch interface interacts with the flow control configuration of the connected peer. Each peer advertises its state to the other peer. The interaction of the flow control configuration of the peers determines the flow control

behavior (resolution) between them, as shown in [Table 515](#). The first four columns show the Ethernet PAUSE configuration on the local QFX Series or EX4600 switch and on the connected peer (also known as the *link partner*). The last two columns show the Ethernet PAUSE resolution that results from the local and peer configurations on each interface. This illustrates how the Ethernet PAUSE configuration of each interface affects the Ethernet PAUSE behavior on the other interface.



**NOTE:** In the Resolution columns of the table, disabling Ethernet PAUSE transmit means that the interface receive buffers do not generate and send Ethernet PAUSE messages to the peer. Disabling Ethernet PAUSE receive means that the interface transmit buffers do not respond to Ethernet PAUSE messages received from the peer.

**Table 515: Asymmetric Ethernet PAUSE Behavior on Local and Peer Interfaces**

| Local Interface (QFX Series or EX4600 Switch) |             | Peer Interface |             | Local Resolution                                                  | Peer Resolution                                                   |
|-----------------------------------------------|-------------|----------------|-------------|-------------------------------------------------------------------|-------------------------------------------------------------------|
| PAUSE Bit                                     | ASM_DIR Bit | PAUSE Bit      | ASM_DIR Bit |                                                                   |                                                                   |
| 0                                             | 0           | Don't care     | Don't care  | Disable Ethernet PAUSE transmit and receive                       | Disable Ethernet PAUSE transmit and receive                       |
| 0                                             | 1           | 0              | Don't care  | Disable Ethernet PAUSE transmit and receive                       | Disable Ethernet PAUSE transmit and receive                       |
| 0                                             | 1           | 1              | 0           | Disable Ethernet PAUSE transmit and receive                       | Disable Ethernet PAUSE transmit and receive                       |
| 0                                             | 1           | 1              | 1           | Enable Ethernet PAUSE transmit and disable Ethernet PAUSE receive | Disable Ethernet PAUSE transmit and enable Ethernet PAUSE receive |
| 1                                             | 0           | 0              | Don't care  | Disable Ethernet PAUSE transmit and receive                       | Disable Ethernet PAUSE transmit and receive                       |
| 1                                             | 0           | 1              | Don't care  | Enable Ethernet PAUSE transmit and receive                        | Enable Ethernet PAUSE transmit and receive                        |
| 1                                             | 1           | 0              | 0           | Disable Ethernet PAUSE transmit and receive                       | Disable Ethernet PAUSE transmit and receive                       |
| 1                                             | 1           | 0              | 1           | Enable Ethernet PAUSE receive and disable Ethernet PAUSE transmit | Enable Ethernet PAUSE transmit and disable Ethernet PAUSE receive |
| 1                                             | 1           | Don't care     | Don't care  | Enable Ethernet PAUSE transmit and receive                        | Enable Ethernet PAUSE transmit and receive                        |



**NOTE:** For your convenience, [Table 515](#) replicates Table 28B-3 of Section 2 of the IEEE 802.X specification.

## PFC

PFC is a lossless transport and congestion relief feature that works by providing granular link-level flow control for each IEEE 802.1p code point (priority) on a full-duplex Ethernet link. When the receive buffer on a switch interface fills to a threshold, the switch transmits a pause frame to the sender (the connected peer) to temporarily stop the sender from transmitting more frames. The buffer threshold must be low enough so that the sender has time to stop transmitting frames and the receiver can accept the frames already on the wire before the buffer overflows. The switch automatically sets queue buffer thresholds to prevent frame loss.

When congestion forces one priority on a link to pause, all of the other priorities on the link continue to send frames. Only frames of the paused priority are not transmitted. When the receive buffer empties below another threshold, the switch sends a message that starts the flow again.

You configure PFC using a congestion notification profile (CNP). A CNP has two parts:

- **Input**—Specify the code point (or code points) on which to enable PFC, and optionally specify the maximum receive unit (MRU) and the cable length between the interface and the connected peer interface.
- **Output**—Specify the output queue or output queues that respond to pause messages from the connected peer.

You apply a PFC configuration by configuring a CNP on one or more interfaces. Each interface that uses a particular CNP is enabled to pause traffic identified by the priorities (code points) specified in that CNP. You can configure one CNP on an interface, and you can configure different CNPs on different interfaces. When you configure a CNP on an interface, ingress traffic that is mapped to a priority that the CNP enables for PFC is paused whenever the queue buffer fills to the pause threshold. (The pause threshold is not user-configurable.)

Configure PFC for a priority end to end along the entire data path to create a lossless lane of traffic on the network. You can selectively pause the traffic in any queue without pausing the traffic for other queues on the same link. You can create lossless lanes for traffic such as FCoE, LAN backup, or management, while using standard frame-drop congestion management for IP traffic on the same link.

Potential consequences of flow control are:

- Ingress port congestion (configuring too many lossless flows can cause ingress port congestion)
- A paused priority that causes upstream devices to pause the same priority, thus spreading congestion back through the network

By definition, PFC supports symmetric pause only (as opposed to Ethernet PAUSE, which supports symmetric and asymmetric pause). With symmetric pause, a device can:

- Transmit pause frames to pause incoming traffic. (You configure this using the input stanza of a congestion notification profile.)
- Receive pause frames and stop sending traffic to a device whose buffer is too full to accept more frames. (You configure this using the output stanza of a congestion notification profile.)

Receiving a PFC frame from a connected peer pauses traffic on egress queues based on the IEEE 802.1p priorities that the PFC pause frame identifies. The priorities are 0 through 7. By default, the priorities map to queue numbers 0 through 7, respectively, and to specific forwarding classes, as shown in [Table 516](#):

**Table 516: Default PFC Priority to Queue and Forwarding Class Mapping**

| IEEE 802.1p Priority (Code Point) | Queue | Forwarding Class |
|-----------------------------------|-------|------------------|
| 0 (000)                           | 0     | best-effort      |
| 1 (001)                           | 1     | best-effort      |
| 2 (010)                           | 2     | best-effort      |
| 3 (011)                           | 3     | fcoe             |
| 4 (100)                           | 4     | no-loss          |
| 5 (101)                           | 5     | best-effort      |
| 6 (110)                           | 6     | network-control  |
| 7 (111)                           | 7     | network-control  |

For example, a received PFC pause frame that pauses priority 3 pauses output queue 3. If you do not want to use the default configuration, you can configure customized mapping of priorities to queues and forwarding classes.



**NOTE:** By convention, deployments with converged server access typically use IEEE 802.1p priority 3 for FCoE traffic. The default configuration sets the fcoe forwarding class as a lossless forwarding class that is mapped to queue 3. The default classifier maps incoming priority 3 traffic to the fcoe forwarding class. *However, you must apply PFC to the entire FCoE data path to configure the end-to-end lossless behavior that FCoE traffic requires.*

If your network uses priority 3 for FCoE traffic, we recommend that you use the default configuration. If your network uses a priority other than 3 for FCoE traffic, you can configure lossless FCoE transport on any IEEE 80.21p priority as described in [“Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows” on page 6873](#) and [Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway](#).

To enable PFC on a priority:

1. Specify the IEEE 802.1p code point to pause in the input stanza of a CNP.
2. If you are not using the default lossless forwarding classes, specify the IEEE 802.1p code point to pause and the corresponding output queue in the output stanza of the CNP.
3. Apply the CNP to the ingress interfaces on which you want to pause the traffic.
4. If you are not using the default lossless forwarding classes, apply the CNP to the ingress interfaces on which you want to pause the traffic.



**CAUTION:** Any change to the PFC configuration on a port temporarily blocks the entire port (not just the priorities affected by the PFC change) so that the port can implement the change, then unblocks the port. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

A change to the PFC configuration means any change to a CNP, including changing the input portion of the CNP (enabling or disabling PFC on a priority, or changing the MRU or cable-length values) or changing the output portion of the CNP that enables or disables output flow control on a queue. A PFC configuration change only affects ports that use the changed CNP.

The following actions change the PFC configuration:

- Deleting or disabling a PFC configuration (input or output) in a CNP that is in use on one or more interfaces. For example:
  1. An existing CNP with an input stanza that enables PFC on priorities 3, 5, and 6 is configured on interfaces xe-0/0/20 and xe-0/0/21.
  2. We disable the PFC configuration for priority 6 in the input CNP, and then commit the configuration.

3. The PFC configuration change causes all traffic on interfaces xe-0/0/20 and xe-0/0/21 to stop until the PFC change has been implemented. When the PFC change has been implemented, traffic resumes.

- Configuring a CNP on an interface. (This changes the PFC state by enabling PFC on one or more priorities.)
- Deleting a CNP from an interface. (This changes the PFC state by disabling PFC on one or more priorities.)

---

When you associate the CNP with an interface, the interface uses PFC to send pause requests when the output queue buffer for the lossless traffic fills to the pause threshold.

On switches that use different classifiers for unicast and multdestination traffic, you can map a unicast queue (queue 0 through 7) and a multdestination queue (queue 8, 9, 10, or 11) to the same IEEE 802.1p code point (priority) so that both unicast and multicast traffic use that priority. However, do not map multdestination traffic to lossless output queues. Starting with Junos OS Release 12.3, you can map one priority to multiple output queues.



**NOTE:** You can attach a maximum of one CNP to an interface, but you can create an unlimited number of CNPs that explicitly configure only the input stanza and use the default output stanza.

The output stanza of the CNP maps to a profile that interfaces use to respond to pause messages received from the connected peer. On standalone switches, you can create two CNPs with an explicitly configured output stanza.

When a switch is a Node device in a QFabric system, you can create one CNP with an explicitly configured output stanza. (One fewer profile is available on QFabric systems because the system needs a default profile for fabric interfaces, which are not used as fabric interfaces when the switches are not part of a QFabric system. “[Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows](#)” on page 6873 describes configuring output flow control.

---

## Lossless Transport Support Summary

The switch supports up to six lossless forwarding classes. For lossless transport, you must enable PFC on the IEEE 802.1p priorities (code points) mapped to lossless forwarding classes.



**CAUTION:** Any change to the PFC configuration on a port temporarily blocks the entire port (not just the priorities affected by the PFC change) so that the port can implement the change, then unblocks the port. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

---

The following limitation applies to support lossless transport on QFabric systems only:

- The internal fiber cable length from the QFabric system Node device to the QFabric system Interconnect device cannot exceed 150 meters.

The default CoS configuration provides two lossless forwarding classes, *fcoe* and *no-loss*. If you explicitly configure lossless forwarding classes, you must include the **no-loss** packet drop attribute to enable lossless behavior, or the traffic is not lossless. For both default and explicit lossless forwarding class configuration, you must configure CNP input stanzas to enable PFC on the priority of the lossless traffic and apply the CNPs to ingress interfaces.



**NOTE:** The information in this note applies only to systems that do not run the ELS CLI.

Junos OS Release 12.2 introduced changes to the way the switch handles lossless forwarding classes (including the default *fcoe* and *no-loss* forwarding classes).

In Junos OS Release 12.1, either explicitly configuring the *fcoe* and *no-loss* forwarding classes or using the default configuration for these forwarding classes resulted in the same lossless behavior for traffic mapped to those forwarding classes.

However, in Junos OS Release 12.2, if you explicitly configure the *fcoe* or the *no-loss* forwarding class, that forwarding class is no longer treated as a lossless forwarding class. Traffic mapped to these forwarding classes is treated as lossy (best-effort) traffic. This is true even if the explicit configuration is exactly the same as the default configuration.

If your CoS configuration from Junos OS Release 12.1 or earlier includes the explicit configuration of the *fcoe* or the *no-loss* forwarding class, then when you upgrade to Junos OS Release 12.2, those forwarding classes are not lossless. To preserve the lossless treatment of these forwarding classes, delete the the explicit *fcoe* and *no-loss* forwarding class configuration before you upgrade to Junos OS Release 12.2.

See *Overview of CoS Changes Introduced in Junos OS Release 12.2* for detailed information about this change and how to delete an existing lossless configuration.

In Junos OS Release 12.3, the default behavior of the *fcoe* and *no-loss* forwarding classes is the same as in Junos OS Release 12.2. However, in Junos OS Release 12.3, you can configure up to six lossless forwarding classes. All explicitly configured lossless forwarding classes must include the new *no-loss* packet drop attribute or the forwarding class is lossy.

[“Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows” on page 6873](#) provides detailed information about the explicit configuration of lossless priorities and about the

default configuration of lossless priorities, including the input and output stanzas of the CNP.



**NOTE:** PFC and Ethernet PAUSE are used only on Ethernet interfaces. Fabric (fte) ports on QFabric systems (Node device fabric ports and Interconnect device fabric ports) use link-layer flow control (LLFC) to ensure the appropriate treatment of lossless traffic.

**Related  
Documentation**

- [Understanding DCB Features and Requirements on page 6490](#)
- [Understanding CoS Explicit Congestion Notification on page 6820](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 6905](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 6504](#)

---

## Example: Configuring CoS PFC for FCoE Traffic

---

Priority-based flow control (PFC, described in IEEE 802.1Qbb) is a link-level flow control mechanism that you apply at ingress interfaces. PFC enables you to divide traffic on one physical link into eight priorities. You can think of the eight priorities as eight “lanes” of traffic that correspond to queues (forwarding classes). Each priority is mapped to a 3-bit IEEE 802.1p CoS value in the VLAN header.

You can selectively apply PFC to the traffic in any queue without pausing the traffic in other queues on the same link. You must apply PFC to FCoE traffic to ensure lossless transport.

This example describes how to configure PFC for FCoE traffic:

- [Requirements on page 6504](#)
- [Overview on page 6505](#)
- [Configuration on page 6507](#)
- [Verification on page 6511](#)

### Requirements

This example uses the following hardware and software components:

- One switch
- Junos OS Release 11.1 or later for the QFX Series



## Overview

FCoE traffic requires PFC to ensure lossless packet transport. This example shows you how to configure PFC on FCoE traffic, use the default FCoE forwarding-class-to-queue mapping and:

- Configure a classifier that associates the FCoE forwarding class with FCoE traffic, which is identified by IEEE 802.1p code point 011 (priority 3).
- Configure a congestion notification profile to apply PFC to the FCoE traffic.
- Apply the classifier and the PFC configuration to ingress interfaces.



**NOTE:** Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

- Configure the CoS bandwidth scheduling for the FCoE forwarding class output queue.
- On switches that support enhanced transmission selection (ETS) hierarchical port scheduling, create a forwarding class set (priority group) that includes the FCoE forwarding class; this is required to configure enhanced transmission selection (ETS) and support data center bridging (DCB).
- For ETS, configure the bandwidth scheduling for the FCoE priority group.
- Apply the configuration to ingress and egress interfaces. How this is done differs depending on whether you use ETS or direct port scheduling for the CoS configuration.

For direct port scheduling, you apply a scheduler map directly to the interface. A scheduler map maps schedulers to forwarding classes, and applies the CoS properties of the scheduler to the output queue mapped to the forwarding class.

For ETS hierarchical port scheduling, you apply the scheduler map to a traffic control profile, and then apply the traffic control profile to the interface. The scheduler map maps CoS properties to forwarding classes (and their associated output queues) just as it does for direct port scheduling. The traffic control profile maps CoS properties to the priority group (a group of forwarding classes defined in a forwarding class set) that contains the forwarding class, creating a CoS hierarchy that allocates port bandwidth to a group of forwarding classes (priority group), and then allocates the priority group bandwidth to the individual forwarding classes (see [“Understanding CoS Hierarchical Port Scheduling \(ETS\)”](#) on page 6765).

Each interface in this example acts as both an ingress interface and an egress interface, so the classifier, congestion notification profile, and scheduling are applied to all of the interfaces.

## Topology

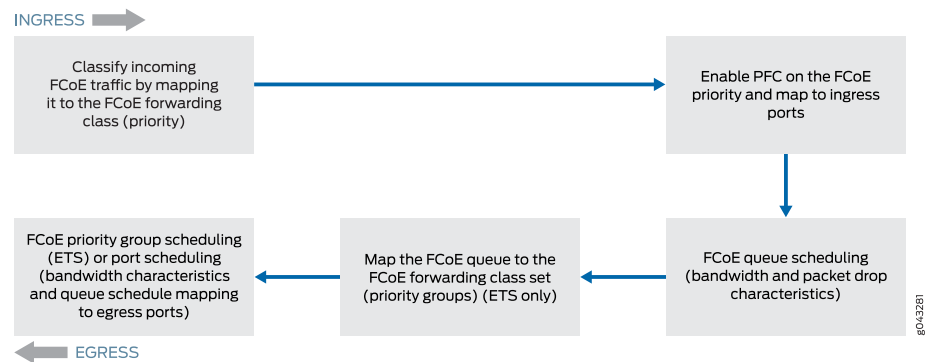
[Table 517](#) shows the configuration components for this example.

Table 517: Components of the PFC for FCoE Traffic Configuration Topology

| Component                                                                                                   | Settings                                                                                                                                                                                                                                                                                                                     |
|-------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Hardware                                                                                                    | One switch                                                                                                                                                                                                                                                                                                                   |
| Behavior aggregate classifier (maps the FCoE forwarding class to incoming packets by IEEE 802.1 code point) | Code point <b>011</b> to forwarding class <b>fcoe</b> and loss priority <b>low</b><br>Ingress interfaces: <b>xe-0/0/31, xe-0/0/32, xe-0/0/33, xe-0/0/34</b>                                                                                                                                                                  |
| PFC congestion notification profile                                                                         | <b>fcoe-cnp:</b><br>Code point <b>011</b><br>Ingress interfaces: <b>xe-0/0/31, xe-0/0/32, xe-0/0/33, xe-0/0/34</b>                                                                                                                                                                                                           |
| FCoE queue scheduler                                                                                        | <b>fcoe-sched:</b><br>Minimum bandwidth <b>3g</b><br>Maximum bandwidth <b>100%</b><br>Priority <b>low</b>                                                                                                                                                                                                                    |
| Forwarding class-to-scheduler mapping                                                                       | Scheduler map <b>fcoe-map:</b><br>Forwarding class <b>fcoe</b><br>Scheduler <b>fcoe-sched</b><br><br>On switches that support direct port scheduling, if you use port scheduling, attach the scheduler map directly to interfaces <b>xe-0/0/31, xe-0/0/32, xe-0/0/33, and xe-0/0/34</b> .                                    |
| ETS only: Forwarding class set (FCoE priority group)                                                        | <b>fcoe-pg:</b><br>Forwarding class <b>fcoe</b><br>Egress interfaces: <b>xe-0/0/31, xe-0/0/32, xe-0/0/33, xe-0/0/34</b>                                                                                                                                                                                                      |
| ETS only: Traffic control profile                                                                           | <b>fcoe-tcp:</b><br>Scheduler map <b>fcoe-map</b><br>Minimum bandwidth <b>3g</b><br>Maximum bandwidth <b>100%</b><br><br>For ETS hierarchical scheduling, attach the traffic control profile (using the <b>output-traffic-control-profile</b> keyword) to interfaces <b>xe-0/0/31, xe-0/0/32, xe-0/0/33, and xe-0/0/34</b> . |

Figure 207 shows a block diagram of the configuration components and the configuration flow of the CLI statements used in the example.

Figure 207: PFC for FCoE Traffic Configuration Components Block Diagram



## Configuration

### CLI Quick Configuration

To quickly configure PFC for FCoE traffic, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level. The configuration is separated into the configuration common to ETS and direct port scheduling, and the portions of the configuration that apply only to ETS and only to port scheduling.

### Common Configuration (Applies to ETS Hierarchical Scheduling and to Port Scheduling)

```
[edit class-of-service]
set classifiers ieee-802.1 fcoe-classifier forwarding-class fcoe loss-priority low code-points 011
set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011 pfc
set interfaces xe-0/0/31 unit 0 classifiers ieee-802.1 fcoe-classifier
set interfaces xe-0/0/32 unit 0 classifiers ieee-802.1 fcoe-classifier
set interfaces xe-0/0/33 unit 0 classifiers ieee-802.1 fcoe-classifier
set interfaces xe-0/0/34 unit 0 classifiers ieee-802.1 fcoe-classifier
set interfaces xe-0/0/31 congestion-notification-profile fcoe-cnp
set interfaces xe-0/0/32 congestion-notification-profile fcoe-cnp
set interfaces xe-0/0/33 congestion-notification-profile fcoe-cnp
set interfaces xe-0/0/34 congestion-notification-profile fcoe-cnp
set schedulers fcoe-sched priority low transmit-rate 3g
set schedulers fcoe-sched shaping-rate percent 100
set scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
```

### Configuration for ETS Hierarchical Scheduling

The ETS-specific portion of this example configures forwarding class set (priority group) membership, priority group CoS settings (traffic control profile), and assigns the priority group and its CoS configuration to the interfaces.

```
[edit class-of-service]
set forwarding-class-sets fcoe-pg class fcoe
set traffic-control-profiles fcoe-tcp scheduler-map fcoe-map guaranteed-rate 3g
```

```
set traffic-control-profiles fcoe-tcp shaping-rate percent 100
set interfaces xe-0/0/31 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
set interfaces xe-0/0/32 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
set interfaces xe-0/0/33 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
set interfaces xe-0/0/34 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
```

## Configuration for Port Scheduling

The port-scheduling-specific portion of this example assigns the scheduler map (which sets the CoS treatment of the forwarding classes in the scheduler map) to the interfaces.

```
[edit class-of-service]
set interfaces xe-0/0/31 scheduler-map fcoe-map
set interfaces xe-0/0/32 scheduler-map fcoe-map
set interfaces xe-0/0/33 scheduler-map fcoe-map
set interfaces xe-0/0/34 scheduler-map fcoe-map
```

### Common Configuration (Applies to ETS Hierarchical Scheduling and to Port Scheduling)

---

#### Step-by-Step Procedure

To configure the ingress classifier for FCoE traffic, PFC on the FCoE traffic, apply the PFC and classifier configurations to interfaces, and configure queue scheduling, for both ETS hierarchical scheduling and port scheduling (common configuration):

1. Configure a classifier to set the loss priority and IEEE 802.1 code point assigned to the FCoE forwarding class at the ingress:  

```
[edit class-of-service]
user@switch# set classifiers ieee-802.1 fcoe-classifier forwarding-class fcoe loss-priority low code-points 011
```
2. Configure PFC on the FCoE queue by applying FCoE to the IEEE 802.1 code point 011:  

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011 pfc
```
3. Apply the PFC configuration to the ingress interfaces:  

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/31 congestion-notification-profile fcoe-cnp
user@switch# set interfaces xe-0/0/32 congestion-notification-profile fcoe-cnp
user@switch# set interfaces xe-0/0/33 congestion-notification-profile fcoe-cnp
user@switch# set interfaces xe-0/0/34 congestion-notification-profile fcoe-cnp
```
4. Assign the classifier to the ingress interfaces:  

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/31 unit 0 classifiers ieee-802.1 fcoe-classifier
user@switch# set interfaces xe-0/0/32 unit 0 classifiers ieee-802.1 fcoe-classifier
user@switch# set interfaces xe-0/0/33 unit 0 classifiers ieee-802.1 fcoe-classifier
user@switch# set interfaces xe-0/0/34 unit 0 classifiers ieee-802.1 fcoe-classifier
```
5. Configure output scheduling for the FCoE queue:  

```
[edit class-of-service]
user@switch# set schedulers fcoe-sched priority low transmit-rate 3g
user@switch# set schedulers fcoe-sched shaping-rate percent 100
```
6. Map the FCoE forwarding class to the FCoE scheduler:

```
[edit class-of-service]
user@switch# set scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
```

### ETS Hierarchical Scheduling Configuration

**Step-by-Step Procedure** To configure the forwarding class set (priority group) and priority group scheduling (in a traffic control profile), and apply the ETS hierarchical scheduling for FCoE traffic to interfaces:

1. Configure the forwarding class set for the FCoE traffic:
 

```
[edit class-of-service]
user@switch# set forwarding-class-sets fcoe-pg class fcoe
```
2. Define the traffic control profile for the FCoE forwarding class set:
 

```
[edit class-of-service]
user@switch# set traffic-control-profiles fcoe-tcp scheduler-map fcoe-map
guaranteed-rate 3g
user@switch# set traffic-control-profiles fcoe-tcp shaping-rate percent 100
```
3. Apply the FCoE forwarding class set and traffic control profile to the egress ports:
 

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/31 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
user@switch# set interfaces xe-0/0/32 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
user@switch# set interfaces xe-0/0/33 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
user@switch# set interfaces xe-0/0/34 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
```

### Port Scheduling Configuration

**Step-by-Step Procedure** To apply port scheduling for FCoE traffic to interfaces:

1. Apply the scheduler map to the egress ports:
 

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/31 scheduler-map fcoe-map
user@switch# set interfaces xe-0/0/32 scheduler-map fcoe-map
user@switch# set interfaces xe-0/0/33 scheduler-map fcoe-map
user@switch# set interfaces xe-0/0/34 scheduler-map fcoe-map
```

### Results

Display the results of the configuration (the system shows only the explicitly configured parameters; it does not show default parameters such as the **fcoe** lossless forwarding class). The results are from the ETS hierarchical scheduling configuration to show the more complex configuration. Direct port scheduling results would not show the traffic control profile or forwarding class set portions of the configuration, and would display the name of the scheduler map under each interface (instead of the names of the forwarding class set and output traffic control profile), but is otherwise the same.

```
user@switch> show configuration class-of-service
classifiers {
```

```
ieee-802.1 fcoe-classifier {
 forwarding-class fcoe {
 loss-priority low code-points 011;
 }
}
traffic-control-profiles {
 fcoe-tcp {
 scheduler-map fcoe-map;
 shaping-rate percent 100;
 guaranteed-rate 3000000000;
 }
}
forwarding-class-sets {
 fcoe-pg {
 class fcoe;
 }
}
congestion-notification-profile {
 fcoe-cnp {
 input {
 ieee-802.1 {
 code-point 011 {
 pfc;
 }
 }
 }
 }
}
}
interfaces {
 xe-0/0/31 {
 congestion-notification-profile fcoe-cnp;
 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 unit 0 {
 classifiers {
 ieee-802.1 fcoe-classifier;
 }
 }
 }
 xe-0/0/32 {
 congestion-notification-profile fcoe-cnp;
 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 unit 0 {
 classifiers {
 ieee-802.1 fcoe-classifier;
 }
 }
 }
 xe-0/0/33 {
```

```

congestion-notification-profile fcoe-cnp;
forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
}
unit 0 {
 classifiers {
 ieee-802.1 fcoe-classifier;
 }
}
}
xe-0/0/34 {
 congestion-notification-profile fcoe-cnp;
 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 unit 0 {
 classifiers {
 ieee-802.1 fcoe-classifier;
 }
 }
}
}
scheduler-maps {
 fcoe-map {
 forwarding-class fcoe scheduler fcoe-sched;
 }
}
schedulers {
 fcoe-sched {
 transmit-rate 3000000000;
 shaping-rate percent 100;
 priority low;
 }
}
}

```



**TIP:** To quickly configure the interfaces, issue the `load merge terminal` command and then copy the hierarchy and paste it into the switch terminal window.

## Verification

To verify that the PFC configuration for FCoE traffic components has been created and is operating properly, perform these tasks:

- [Verifying That Priority-Based Flow Control Has Been Enabled on page 6512](#)
- [Verifying the Ingress Interface PFC Configuration on page 6512](#)

### Verifying That Priority-Based Flow Control Has Been Enabled

---

**Purpose** Verify that PFC is enabled on the FCoE queue to enable lossless transport.

**Action** List the congestion notification profiles using the operational mode command **show class-of-service congestion-notification**:

```
user@switch> show class-of-service congestion-notification
```

```
Type: Input, Name: fcoe-cnp, Index: 51697
```

```
Cable Length: 100 m
```

| Priority | PFC      | MRU  |
|----------|----------|------|
| 000      | Disabled |      |
| 001      | Disabled |      |
| 010      | Disabled |      |
| 011      | Enabled  | 2500 |
| 100      | Disabled |      |
| 101      | Disabled |      |
| 110      | Disabled |      |
| 111      | Disabled |      |

```
Type: Output
```

| Priority | Flow-Control-Queues |
|----------|---------------------|
| 000      | 0                   |
| 001      |                     |
| 010      | 1                   |
| 011      | 2                   |
| 100      | 3                   |
| 101      | 4                   |
| 110      | 5                   |
| 111      | 6                   |
|          | 7                   |

**Meaning** The **show class-of-service congestion-notification** operational command lists all of the congestion notification profiles and which IEEE 802.1p code points have PFC enabled. The command output shows that PFC is enabled on code point 011 for the **fcoe-cnp** congestion notification profile.

The command also shows the default cable length (100 meters), the default maximum receive unit (2500 bytes), and the default mapping of priorities to output queues because this example does not include configuring these options.

### Verifying the Ingress Interface PFC Configuration

---

**Purpose** Verify that the classifier **fcoe-classifier** and the congestion notification profile **fcoe-cnp** are configured on ingress interfaces **xe-0/0/31**, **xe-0/0/32**, **xe-0/0/33**, and **xe-0/0/34**.



**Action** List the ingress interfaces using the operational mode command **show configuration class-of-service interfaces**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/31
congestion-notification-profile fcoe-cnp;
unit 0 {
 classifiers {
 ieee-802.1 fcoe-classifier;
 }
}

user@switch> show configuration class-of-service interfaces xe-0/0/32
congestion-notification-profile fcoe-cnp;
unit 0 {
 classifiers {
 ieee-802.1 fcoe-classifier;
 }
}

user@switch> show configuration class-of-service interfaces xe-0/0/33
congestion-notification-profile fcoe-cnp;
unit 0 {
 classifiers {
 ieee-802.1 fcoe-classifier;
 }
}

user@switch> show configuration class-of-service interfaces xe-0/0/34
congestion-notification-profile fcoe-cnp;
unit 0 {
 classifiers {
 ieee-802.1 fcoe-classifier;
 }
}
```

**Meaning** The **show configuration class-of-service interfaces** commands list the congestion notification profile that is mapped to the interface (**fcoe-cnp**) and the IEEE 802.1p classifier associated with the interface (**fcoe-classifier**).

**Related Documentation**

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 6771](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 6905](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 6493](#)

## Understanding DCBX

---

Data Center Bridging Capability Exchange protocol (DCBX) is an extension of Link Layer Data Protocol (LLDP). If you disable LLDP on an interface, that interface cannot run DCBX. If you attempt to enable DCBX on an interface on which LLDP is disabled, the configuration commit operation fails. Data center bridging (DCB) devices use DCBX to exchange configuration information with directly connected peers.



Video: [What is DCBX Protocol?](#)

This topic describes:

- [DCBX Basics on page 6514](#)
- [DCBX Modes and Support on page 6515](#)
- [DCBX Attribute Types on page 6518](#)
- [DCBX Application Protocol TLV Exchange on page 6519](#)
- [DCBX and PFC on page 6520](#)
- [DCBX and ETS on page 6521](#)

### DCBX Basics

DCBX can:

- Discover the DCB capabilities of peers.
- Detect DCB feature misconfiguration or mismatches between peers.
- Configure DCB features on peers.

You can configure DCBX operation for priority-based flow control (PFC), Layer 2 and Layer 4 applications such as FCoE and iSCSI, and ETS. DCBX is enabled or disabled on a per-interface basis.



**NOTE:** The Juniper Networks QFX10000 does not support enhanced transmission selection (ETS) hierarchical scheduling. Use port scheduling to manage bandwidth on QFX10000 switches.

By default, for PFC and ETS, DCBX automatically negotiates administrative state and configuration with each interface's connected peer. To enable DCBX negotiation for applications, you must configure the applications, map them to IEEE 802.1p code points in an application map, and apply the application map to interfaces.

The FCoE application only needs to be included in an application map when you want an interface to exchange type, length, and values (TLVs) for other applications in addition to FCoE. If FCoE is the only application you want an interface to advertise, then you do not need to use an application map. For ETS, DCBX pushes the switch configuration to

peers if they are set to learn the configuration from the switch (unless you disable sending the ETS recommendation TLV on interfaces in IEEE DCBX mode).

You can override the default behavior for PFC, for ETS, or for all applications mapped to an interface by turning off autonegotiation to force an interface to enable or disable that feature. You can also disable DCBX autonegotiation for applications on an interface by excluding those applications from the application map you apply to that interface or by deleting the application map from the interface.

The default autonegotiation behavior for applications that are mapped to an interface is:

- DCBX is enabled on the interface if the connected peer device also supports DCBX.
- DCBX is disabled on the interface if the connected peer device does not support DCBX.

During negotiation of capabilities, the switch can push the PFC configuration to an attached peer if the peer is configured as “willing” to learn the PFC configuration from other peers. The Juniper Networks switch does not support self autoprovisioning and does not change its configuration during autonegotiation to match the peer configuration. (The Juniper switch is not “willing” to learn the PFC configuration from peers.)



**NOTE:** When a port with DCBX enabled begins to exchange type, length, and value (TLV) entries, optional LLDP TLVs on that port are not advertised to neighbors, so that the switch can interoperate with a wider variety of converged network adapters (CNAs) and Layer 2 switches that support DCBX.

## DCBX Modes and Support

This section describes DCBX support:

- [DCBX Modes \(Versions\) on page 6515](#)
- [Autonegotiation on page 6517](#)
- [CNA Support for DCBX Modes on page 6518](#)
- [Interface Support for DCBX on page 6518](#)

### DCBX Modes (Versions)

The two most common DCBX modes are supported:

- IEEE DCBX—The newest DCBX version. Different TLVs have different subtypes (for example, the subtype for the ETS configuration TLV is 9); the IEEE DCBX Organizationally Unique Identifier (OUI) is 0x0080c2.
- DCBX version 1.01—The Converged Enhanced Ethernet (CEE) version of DCBX. It has a subtype of 2 and an OUI of 0x001b21.

IEEE DCBX and DCBX version 1.01 differ mainly in frame format. DCBX version 1.01 uses one TLV that includes all DCBX attribute information, which is sent as sub-TLVs. IEEE DCBX uses a unique TLV for each DCB attribute.



**NOTE:** The switch does not support pre-CEE (pre-DCB) DCBX versions. Unsupported older versions of DCBX have a subtype of 1 and an OUI of 0x001b21. The switch drops LLDP frames that contain pre-CEE DCBX TLVs.

Table 518 summarizes the differences between IEEE DCBX and DCBX version 1.01, including show command output:

**Table 518: Summary of Differences Between IEEE DCBX and DCBX Version 1.01**

| Characteristic                                                                   | IEEE DCBX                                                                                                                                                                                                                                                                                                                                                                                                                                                                | DCBX Version 1.01                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|----------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| OUI                                                                              | 0x0080c2                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | 0x001b21                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Frame Format                                                                     | Sends a separate, unique TLV for each DCBX attribute. For example, IEEE DCBX uses separate TLVs for ETS, PFC, and each application. Configuration and Recommendation information is sent in different TLVs                                                                                                                                                                                                                                                               | Sends one TLV that includes all DCBX attribute information organized in sub-TLVs. The “willing” bit determines whether or not an interface can change its configuration to match the connected peer.                                                                                                                                                                                                                                                                                                 |
| Symmetric/asymmetric configuration with peer                                     | Asymmetric or symmetric                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Symmetric only                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Differences in the <b>show dcbx interface interface-name</b> operational command | <ul style="list-style-type: none"> <li>Synchronization information is not shown because symmetric configuration is not required.</li> <li>Operational state information is not shown because the operational states do not have to be symmetric.</li> <li>TLV type is shown because unique TLVs are sent for each DCBX attribute.</li> <li>ETS peer Configuration TLV and Recommendation TLV information is shown separately because they are different TLVs.</li> </ul> | <ul style="list-style-type: none"> <li>Synchronization information is shown because symmetric configuration is required.</li> <li>Operational state information is shown because the operational states do have to be symmetric.</li> <li>TLV type is not shown because one TLV is used for all attribute information.</li> <li>Recommendation TLV is not sent (DCBX Version 1.01 uses the “willing” bit to determine whether or not an interface uses the peer interface configuration).</li> </ul> |

For more information about how each DCBX mode exchanges TLVs, see the following specifications:

- For DCBX version 1.01—<http://www.ieee802.org/1/files/private/az-drafts/d2/802-1az-d2-4.pdf>
- For IEEE DCBX—<http://www.ieee802.org/1/files/public/dcs2008/az-wadder-dcbx-capability-exchange-discovery-protocol-108-v101.pdf>



**NOTE:** As of Junos OS Release 12.2, this document is located in a private area of the IEEE website, and access requires a password from the IEEE organization. If you are not an IEEE member, you might not be able to access this document until it moves to the public area of the IEEE website.

You can configure interfaces to use the following DCBX modes:

- IEEE DCBX—The interface uses IEEE DCBX regardless of the configuration on the connected peer.
- DCBX version 1.01—The interface uses DCBX version 1.01 regardless of the configuration on the connected peer.
- Autonegotiation—The interface automatically negotiates with the connected peer to determine the DCBX version the peers use. Autonegotiation is the default DCBX mode.

If you configure a DCBX mode on an interface, the interface ignores DCBX protocol data units (PDUs) it receives from the connected peer if the PDUs do not match the DCBX version configured on the interface. For example, if you configure an interface to use IEEE DCBX and the connected peer sends DCBX version 1.01 LLDP PDUs, the interface ignores the version 1.01 PDUs. If you configure an interface to use DCBX version 1.01 and the peer sends IEEE DCBX LLDP PDUs, the interface ignores the IEEE DCBX PDUs.



**NOTE:** On interfaces that use the IEEE DCBX mode, the `show dcbx neighbors interface interface-name` operational command does not include application, PFC, or ETS operational state in the output.

### Autonegotiation

Autonegotiation is the default DCBX mode. Each interface automatically negotiates with its connected peer to determine the DCBX version that both interfaces use to exchange DCBX information.

When an interface connects to its peer interface, the interface advertises IEEE DCBX TLVs to the peer. If the interface receives one IEEE DCBX PDU from the peer, the interface sets the DCBX mode as IEEE DCBX. If the interface receives three DCBX version 1.01 TLVs from the peer, the interface sets DCBX version 1.01 as the DCBX mode.

Autonegotiation works slightly differently on standalone switches compared to QFabric systems:

- Standalone switches—When an interface connects to its peer interface, the interface advertises IEEE DCBX TLVs to the peer. If the interface receives an IEEE DCBX TLV from the peer, the interface sets IEEE DCBX as the DCBX mode. If the interface receives three consecutive DCBX version 1.01 TLVs from the peer, the interface sets DCBX version 1.01 as the DCBX mode.
- QFabric system—When an interface connects to its peer interface, the interface advertises DCBX version 1.01 TLVs to the peer. If the interface receives an IEEE DCBX

TLVs from the peer, the interface sets IEEE DCBX as the DCBX mode. If the interface receives three consecutive DCBX version 1.01 TLVs from the peer, the interface retains DCBX version 1.01 as the DCBX mode.



**NOTE:** If the link flaps or the LLDP process restarts, the interface starts the autonegotiation process again. The interface does not use the last received DCBX communication mode.

---

### CNA Support for DCBX Modes

Different CNA vendors support different versions and capabilities of DCBX. The DCBX configuration you use on switch interfaces depends on the DCBX features that the CNAs in your network support.

---

### Interface Support for DCBX

You can configure DCBX on 10-Gigabit Ethernet interfaces and on link aggregation group (LAG) interfaces whose member interfaces are all 10-Gigabit Ethernet interfaces.

## DCBX Attribute Types

DCBX has three attribute types:

- **Informational**—These attributes are exchanged using LLDP, but do not affect DCBX state or operation; they only communicate information to the peer. For example, application priority TLVs are informational TLVs.
- **Asymmetric**—The values for these types of attributes do not have to be the same on the connected peer interfaces. Peers exchange asymmetric attributes when the attribute values can differ on each peer interface. The peer interface configurations might match or they might differ. For example, ETS Configuration and Recommendation TLVs are asymmetric TLVs.
- **Symmetric**—The intention is that the values for these types of attributes should be the same on both of the connected peer interfaces. Peer interfaces exchange symmetric attributes to ensure symmetric DCBX configuration for those attributes. For example, PFC Configuration TLVs are symmetric TLVs.

The following sections describe asymmetric and symmetric DCBX attributes:

- [Asymmetric Attributes on page 6518](#)
- [Symmetric Attributes on page 6519](#)

---

### Asymmetric Attributes

DCBX passes asymmetric attributes between connected peer interfaces to communicate parameter information about those attributes (features). The resulting configuration for an attribute might be different on each peer, so the parameters configured on one interface might not match the parameters on the connected peer interface.

There are two types of asymmetric attribute TLVs:

- **Configuration TLV**—Configuration TLVs communicate the current operational state and the state of the “willing” bit. The “willing” bit communicates whether or not the interface is willing to accept and use the configuration from the peer interface. If an interface is “willing,” the interface uses the configuration it receives from the peer interface. (The peer interface configuration can override the configuration on the “willing” interface.) If an interface is “not willing,” the configuration on the interface cannot be overridden by the peer interface configuration.
- **Recommendation TLV**—Recommendation TLVs communicate the parameters the interface recommends that the connected peer interface should use. When an interface sends a Recommendation TLV, if the connected peer is “willing,” the connected peer changes its configuration to match the parameters in the Recommendation TLV.

---

### Symmetric Attributes

DCBX passes symmetric attributes between connected peer interfaces to communicate parameter information about those attributes (features), with the objective that both interfaces should use the same configuration. The intent is that the parameters configured on one interface should match the parameters on the connected peer interface.

There is one type of symmetric attribute TLV, the Configuration TLV. As with asymmetric attributes, symmetric attribute Configuration TLVs communicate the current operational state and the state of the “willing” bit. “Willing” interfaces use the peer interface parameter values for the attribute. (The attribute configuration of the peer overrides the configuration on the “willing” interface.)

## DCBX Application Protocol TLV Exchange

DCBX advertises the switch's capabilities for Layer 2 applications such as FCoE and Layer 4 applications such as iSCSI:

- [Application Protocol TLV Exchange on page 6519](#)
- [FCoE Application Protocol TLV Exchange on page 6519](#)
- [Disabling Application Protocol TLV Exchange on page 6520](#)

---

### Application Protocol TLV Exchange

For all applications, DCBX advertises the application's state and IEEE 802.1p code points on the interfaces to which the application is mapped. If an application is not mapped to an interface, that interface does not advertise the application's TLVs. There is an exception for FCoE application protocol TLV exchange when FCoE is the only application you want DCBX to advertise on an interface.

---

### FCoE Application Protocol TLV Exchange

Protocol TLV exchange for the FCoE application depends on whether FCoE is the only application you want the interface to advertise or whether you want the interface to exchange other application TLVs in addition to FCoE TLVs.

If FCoE is the only application you want DCBX to advertise on an interface, DCBX exchanges FCoE application protocol TLVs by default if the interface:

- Carries FCoE traffic (traffic mapped by CoS configuration to the FCoE forwarding class)
- Has a congestion notification profile with PFC enabled on the FCoE priority (IEEE 802.1p code point)
- Does *not* have an application map



**NOTE:** If no CoS configuration for FCoE is mapped to an interface, that interface does not exchange FCoE application protocol TLVs.

If you want DCBX to advertise FCoE and other applications on an interface, you must specify all of the applications, including FCoE, in an application map, and apply the application map to the desired interfaces.



**NOTE:** If an application map is applied to an interface, the FCoE application must be explicitly configured in the application map, or the interface does not exchange FCoE TLVs.

When DCBX advertises the FCoE application, it advertises the FCoE state and IEEE 802.1p code points. If a peer device connected to a switch interface does not support FCoE, DCBX uses autonegotiation to mark the interface as “FCoE down,” and FCoE is disabled on that interface.

### Disabling Application Protocol TLV Exchange

To disable DCBX application protocol exchange for all applications on an interface, issue the **set protocols dcbx interface *interface-name* applications no-auto-negotiation** command.

You can also disable DCBX application protocol exchange for applications on an interface by deleting the application map from the interface, or by deleting a particular application from the application map. However, when you delete an application from an application map, the application protocol is no longer exchanged on any interface which uses that application map.

## DCBX and PFC

After you enable PFC on a switch interface, DCBX uses autonegotiation to control the operational state of the PFC functionality.

If the peer device connected to the interface supports PFC and is provisioned compatibly with the switch, DCBX sets the PFC operational state to enabled. If the peer device connected to the interface does not support PFC or is not provisioned compatibly with the switch, DCBX sets the operational state to disabled. (PFC must be symmetrical.)

If the peer advertises that it is “willing” to learn its PFC configuration from the switch, DCBX pushes the switch’s PFC configuration to the peer and does not check the peer’s administrative state.



You can manually override DCBX control of the PFC operational state on a per-interface basis by disabling autonegotiation. If you disable autonegotiation on an interface on which you have configured PFC, then PFC is enabled on that interface regardless of the peer configuration. To disable PFC on an interface, do not configure PFC on that interface.

## DCBX and ETS

This section describes:

- [Default DCBX ETS Advertisement on page 6521](#)
- [ETS Advertisement and Peer Configuration on page 6521](#)
- [ETS Recommendation TLV on page 6522](#)

---

### Default DCBX ETS Advertisement

If you do not configure ETS on an interface, the switch automatically creates a default priority group that contains all of the priorities (forwarding classes, which represent output queues) and assigns 100 percent of the port output bandwidth to that priority group. The default priority group is transparent. It does not appear in the configuration and is used for DCBX advertisement. DCBX advertises the default priority group, its priorities, and the assigned bandwidth.

If you configure ETS on an interface, DCBX advertises:

- Each priority group on the interface
- The priorities in each priority group
- The bandwidth properties of each priority group and priority

Any priority on that interface that is not part of an explicitly configured priority group (forwarding class set) is assigned to the automatically generated default priority group and receives no bandwidth. If you configure ETS on an interface, every forwarding class (priority) on that interface for which you want to forward traffic must belong to a forwarding class set (priority group).

---

### ETS Advertisement and Peer Configuration

DCBX does not control the switch's ETS (hierarchical scheduling) operational state. If the connected peer is configured as "willing," DCBX pushes the switch's ETS configuration to the switch's peers if the ETS Recommendation TLV is enabled (it is enabled by default). If the peer does not support ETS or is not consistently provisioned with the switch, DCBX does not change the ETS operational state on the switch. The ETS operational state remains enabled or disabled based only on the switch hierarchical scheduling configuration and is enabled by default.

When ETS is configured, DCBX advertises the priority groups, the priorities in the priority groups, and the bandwidth configuration for the priority groups and priorities. Any priority (essentially a forwarding class or queue) that is not part of a priority group has no scheduling properties and receives no bandwidth.

You can manually override whether DCBX advertises the ETS state to the peer on a per-interface basis by disabling autonegotiation. This does not affect the ETS state on

the switch or on the peer, but it does prevent the switch from sending the Recommendation TLV or the Configuration TLV to the connected peer. To disable ETS on an interface, do not configure priority groups (forwarding class sets) on the interface.

### ETS Recommendation TLV

---

The ETS Recommendation TLV communicates the ETS settings that the switch wants the connected peer interface to use. If the peer interface is “willing,” it changes its configuration to match the configuration in the ETS Recommendation TLV. By default, the switch interfaces send the ETS Recommendation TLV to the peer. The settings communicated are the egress ETS settings defined by configuring hierarchical scheduling on the interface.

We recommend that you use the same ETS settings on the connected peer that you use on the switch interface and that you leave the ETS Recommendation TLV enabled. However, on interfaces that use IEEE DCBX as the DCBX mode, if you want an asymmetric configuration between the switch interface and the connected peer, you can disable the ETS Recommendation TLV by including the **no-recommendation-tlv** statement at the **[edit protocols dcbx interface *interface-name* enhanced-transmission-selection]** hierarchy level.



**NOTE:** You can disable the ETS Recommendation TLV only when the DCBX mode on the interface is IEEE DCBX. Disabling the ETS Recommendation TLV has no effect if the DCBX mode on the interface is DCBX version 1.01. (IEEE DCBX uses separate application attribute TLVs, but DCBX version 1.01 sends all application attributes in the same TLV and uses sub-TLVs to separate the information.)

If you disable the ETS Recommendation TLV, the switch still sends the ETS Configuration TLV to the connected peer. The result is that the connected peer is informed about the switch DCBX ETS configuration, but even if the peer is “willing,” the peer does not change its configuration to match the switch configuration. This is asymmetric configuration—the two interfaces can have different parameter values for the ETS attribute.

For example, if you want a CNA connected to a switch interface to have different bandwidth allocations than the switch ETS configuration, you can disable the ETS Recommendation TLV and configure the CNA for the desired bandwidth. The switch interface and the CNA exchange configuration parameters, but the CNA does not change its configuration to match the switch interface configuration.

#### Related Documentation

- [Understanding DCBX Application Protocol TLV Exchange on page 6528](#)
- [Understanding DCB Features and Requirements on page 6490](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 6493](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 6765](#)
- [Understanding CoS Port Schedulers on QFX Switches on page 6729](#)
- [Understanding FCoE on page 6436](#)

- [Configuring the DCBX Mode on page 6523](#)
- [Configuring DCBX Autonegotiation on page 6524](#)
- [Disabling the ETS Recommendation TLV on page 6527](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 6535](#)

## Configuring the DCBX Mode

You can configure the DCBX mode that an interface uses to communicate with the connected peer. Three DCBX modes are supported:

- **Autonegotiation**—The interface negotiates with the connected peer to determine the DCBX mode. This is the default DCBX mode.
- **IEEE DCBX**—The interface uses IEEE DCBX type, length, and value (TLV) to exchange DCBX information with the connected peer. QFX3500 Node devices come up with IEEE DCBX enabled by default and then autonegotiate with the connected peer to determine the final DCBX mode.
- **DCBX Version 1.01**—The interface uses Converged Enhanced Ethernet (CEE) DCBX version 1.01 TLVs to exchange DCBX information with the connected peer. QFabric system Node devices other than QFX3500 switches come up with DCBX version 1.01 enabled by default and then autonegotiate with the connected peer to determine the final DCBX mode.



**NOTE:** Pre-CEE (pre-DCB) versions of DCBX such as DCBX version 1.00 are not supported. If an interface receives an LLDP frame with pre-CEE DCBX TLVs, the system drops the frame.

Configure the DCBX mode by specifying the mode for one interface or for all interfaces.

- To configure the DCBX mode, specify the interface and the mode:

```
[edit protocols dcbx]
user@switch# set interface interface-name mode (auto-negotiate | ieee-dcbx |
dcbx-version-1.01)
```

For example, to configure DCBX version 1.01 on interface **xe-0/0/21**:

```
user@switch# set protocols dcbx interface xe-0/0/21 mode dcbx-version-1.01
```

To configure IEEE DCBX on all interfaces:

```
user@switch# set protocols dcbx interface all mode ieee-dcbx
```

### Related Documentation

- [Configuring DCBX Autonegotiation on page 6524](#)
- [Disabling the ETS Recommendation TLV on page 6527](#)
- [Understanding DCBX on page 6514](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 6528](#)

- [show dcbx neighbors on page 6581](#)

## Configuring DCBX Autonegotiation

---

Data Center Bridging Capability Exchange protocol (DCBX) discovers the data center bridging (DCB) capabilities of peers by exchanging feature configuration information. DCBX also detects feature misconfiguration and mismatches, and can configure DCB on peers. DCBX is an extension of the Link Layer Discovery Protocol (LLDP), and LLDP must remain enabled on every interface for which you want to use DCBX. If you attempt to enable DCBX on an interface on which LLDP is disabled, the configuration commit operation fails.



**NOTE:** LLDP and DCBX are enabled by default on all interfaces.

The switch supports DCBX autonegotiation for:

- Priority-based flow control (PFC) configuration
- Layer 2 and Layer 4 applications such as Fibre Channel over Ethernet (FCoE) and Internet Small Computer System Interface (iSCSI)
- Enhanced transmission selection (ETS) advertisement

DCBX autonegotiation is configured on a per-interface basis for each supported feature or application. The PFC and application DCBX exchanges use autonegotiation by default. The default autonegotiation behavior is:

- DCBX is enabled on the interface if the connected peer device also supports DCBX.
- DCBX is disabled on the interface if the connected peer device does not support DCBX.

You can override the default behavior for each feature by turning off autonegotiation to force an interface to enable or disable the feature.

Autonegotiation of ETS means that when ETS is enabled on an interface (priority groups are configured), the interface advertises its ETS configuration to the peer device. In this case, priorities (forwarding classes) that are not part of a priority group (forwarding class set) receive no bandwidth and are advertised in an automatically generated default forwarding class. If ETS is not enabled on an interface (no priority groups are configured), all of the priorities are advertised in one automatically generated default priority group that receives 100 percent of the port bandwidth.

Disabling ETS autonegotiation prevents the interface from sending the Recommendation TLV or the Configuration TLV to the connected peer.

On interfaces that use IEEE DCBX mode to exchange DCBX parameters, you can disable autonegotiation of the ETS Recommendation TLV to the peer if you want an asymmetric ETS configuration between the peers. DCBX still exchanges the ETS Configuration TLV if you disable the ETS Recommendation TLV.

Autonegotiation of PFC means that when PFC is enabled on an interface, if the peer device connected to the interface supports PFC and is provisioned compatibly with the switch, DCBX sets the PFC operational state to enabled. If the peer device connected to the interface does not support PFC or is not provisioned compatibly with the switch, DCBX sets the operational state to disabled.

In addition, if the peer advertises that it is “willing” to learn its PFC configuration from the switch, DCBX pushes the switch’s PFC configuration to the peer and does not check the peer’s administrative state. The switch does not learn PFC configuration from peers (the switch does not advertise its state as “willing”).

Disabling PFC autonegotiation prevents the interface from exchanging PFC configuration information with the peer. It forces the interface to enable PFC if PFC is configured on the interface or to disable PFC if PFC is not configured on the interface. If you disable PFC autonegotiation, the assumption is that the peer is also configured manually.

Autonegotiation of applications depends on whether or not you apply an application map to an interface. If you apply an application map to an interface, the interface autonegotiates DCBX for each application in the application map. PFC must be enabled on the FCoE priority (the FCoE IEEE 802.1p code point) for the interface to advertise the FCoE application. The interface only advertises applications that are included in the application map.

For example, if you apply an application map to an interface and the application map does not include the FCoE application, then that interface does not perform DCBX advertisement of FCoE.

If you do not apply an application map to an interface, DCBX does not advertise applications on that interface, with the exception of FCoE, which is handled differently than other applications.



**NOTE:** If you do not apply an application map to an interface, the interface performs autonegotiation of FCoE if the interface carries traffic in the FCoE forwarding class and also has PFC enabled on the FCoE priority. On such interfaces, if DCBX detects that the peer device connected to the interface supports FCoE, the switch advertises its FCoE capability and IEEE 802.1p code point on that interface. If DCBX detects that the peer device connected to the interface does not support FCoE, DCBX marks that interface as “FCoE down” and disables FCoE on the interface.

When DCBX marks an interface as “FCoE down,” the behavior of the switch depends on how you use it in the network:

- When the switch acts as an FCoE transit switch, the interface drops all of the FIP packets it receives. In addition, FIP packets received from an FCoE forwarder (FCF) are not forwarded to interfaces marked as “FCoE down.”
- When the switch acts as an FCoE-FC gateway (only switches that support native Fibre Channel interfaces), it does not send or receive FCoE Initialization Protocol (FIP) packets.

Disabling autonegotiation prevents the interface from exchanging application information with the peer. In this case, the assumption is that the peer is also configured manually.

To disable DCBX autonegotiation of PFC, applications (including FCoE), and ETS using the CLI:

1. Turn off autonegotiation for PFC.

```
[edit]
user@switch# set protocols dcbx interface interface-name priority-flow-control
no-auto-negotiation
```

2. Turn off autonegotiation for applications.

```
[edit]
user@switch# set protocols dcbx interface interface-name applications no-auto-negotiation
```

3. Turn off autonegotiation for ETS.

```
[edit]
user@switch# set protocols dcbx interface interface-name enhanced-transmission-selection
no-auto-negotiation
```

To disable autonegotiation of the ETS Recommendation TLV so that DCBX exchanges only the ETS Configuration TLV:

- [edit protocols dcbx interface *interface-name*]  
user@switch# set enhanced-transmission-selection no-recommendation-tlv

**Related  
Documentation**

- [Example: Configuring DCBX Application Protocol TLV Exchange on page 6535](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 6504](#)
- [Disabling the ETS Recommendation TLV on page 6527](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 6528](#)

## Disabling the ETS Recommendation TLV

The enhanced transmission selection (ETS) Recommendation TLV communicates the ETS settings that the switch wants the connected peer interface to use. If the peer interface is “willing,” the peer interface changes its configuration to match the configuration in the ETS Recommendation TLV. By default, the switch interfaces send the ETS Recommendation TLV to the peer. The settings communicated are the egress ETS settings defined by configuring hierarchical scheduling on the interface.

We recommend that you use the same ETS settings on the connected peer that you use on the switch interface and that you leave the ETS Recommendation TLV enabled. However, on interfaces that use IEEE DCBX as the DCBX mode, if you want an asymmetric configuration between the switch interface and the connected peer, you can disable the ETS Recommendation TLV.



**NOTE:** Disabling the ETS Recommendation TLV on interfaces that use DCBX version 1.01 as the DCBX mode has no effect and does not change DCBX behavior.

If you disable the ETS Recommendation TLV, the switch still sends the ETS Configuration TLV to the connected peer. The result is that the connected peer is informed about the switch DCBX ETS configuration, but even if the peer is “willing,” the peer does not change its configuration to match the switch configuration. This is asymmetric configuration—the two interfaces can have different parameter values for the ETS attribute.

To disable the ETS Recommendation TLV:

- [edit protocols dcbx interface *interface-name*]  
user@switch# **set enhanced-transmission-selection no-recommendation-tlv**

### Related Documentation

- [Configuring the DCBX Mode on page 6523](#)
- [Configuring DCBX Autonegotiation on page 6524](#)
- [Understanding DCBX on page 6514](#)
- [Understanding Data Center Bridging Capability Exchange Protocol for EX Series Switches](#)

## Understanding DCBX Application Protocol TLV Exchange

---

Data Center Bridging Capability Exchange protocol (DCBX) discovers the data center bridging (DCB) capabilities of connected peers. DCBX also advertises the capabilities of applications on interfaces by exchanging application protocol information through application type, length, and value (TLV) elements. DCBX is an extension of Link Layer Discovery Protocol (LLDP). LLDP must remain enabled on every interface on which you want to use DCBX.



**NOTE:** LLDP and DCBX are enabled by default on all interfaces.

Setting up application protocol exchange consists of:

- Defining applications
- Mapping the applications to IEEE 802.1p code points in an *application map*
- Configuring classifiers to prioritize incoming traffic and map the incoming traffic to the application by the traffic code points
- Applying the application maps and classifiers to interfaces

You need to explicitly define the applications that you want an interface to advertise. The FCoE application is a special case (see [“Applications” on page 6528](#)) and only needs to be defined on an interface if you want DCBX to exchange application protocol TLVs for other applications in addition to FCoE on that interface.

You also need to explicitly map all of the defined applications that you want an interface to advertise to IEEE 802.1p code points in an application map. The FCoE application is a special case that only requires inclusion in an application map when you want an interface to use DCBX for other applications in addition to FCoE, as described later in this topic (see [“Application Maps” on page 6529](#)).

This topic describes:

- [Applications on page 6528](#)
- [Application Maps on page 6529](#)
- [Classifying and Prioritizing Application Traffic on page 6530](#)
- [Enabling Interfaces to Exchange Application Protocol Information on page 6531](#)
- [Disabling DCBX Application Protocol Exchange on page 6531](#)

## Applications

Before an interface can exchange application protocol information, you need to define the applications that you want to advertise. The exception is the FCoE application. If FCoE is the only application that you want the interface to advertise, then you do not need to define the FCoE application. You need to define the FCoE application only if you want interfaces to advertise other applications in addition to FCoE.





**NOTE:** If FCoE is the only application that you want DCBX to advertise on an interface, DCBX exchanges FCoE application protocol TLVs by default if the interface:

- Carries FCoE traffic (traffic mapped by CoS configuration to the FCoE forwarding class and applied to the interface)
- Has a congestion notification profile with PFC enabled on the FCoE priority (IEEE 802.1p code point)
- Does *not* have an application map

If you apply an application map to an interface, then all applications that you want DCBX to advertise must be defined and configured in the application map, including the FCoE application.

If no CoS configuration for FCoE is mapped to an interface, that interface does not exchange FCoE application protocol TLVs.

You can define:

- Layer 2 applications by EtherType
- Layer 4 applications by a combination of protocol (TCP or UDP) and destination port number

The EtherType is a two-octet field in the Ethernet frame that denotes the protocol encapsulated in the frame. For a list of common EtherTypes, see <http://standards.ieee.org/develop/regauth/ethertype/eth.txt> on the IEEE standards organization website. For a list of port numbers and protocols, see the *Service Name and Transport Protocol Port Number Registry* at <http://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xml> on the Internet Assigned Numbers Authority (IANA) website.

You must explicitly define each application that you want to advertise, except FCoE. The FCoE application is defined by default (EtherType 0x8906).

## Application Maps

An application map maps defined applications to one or more IEEE 802.1p code points. Each application map contains one or more applications. DCBX includes the configured application code points in the protocol TLVs exchanged with the connected peer.

To exchange protocol TLVs for an application, you must include the application in an application map. The FCoE application is a special case:

- If you want DCBX to exchange application protocol TLVs for more than one application on a particular interface, you must configure the applications, define an application map to map the applications to code points, and apply the application map to the interface. In this case, you must also define the FCoE application and add it to the application map.

This is the same process and treatment required for all other applications. In addition, for DCBX to exchange FCoE application TLVs, you must enable priority-based flow control (PFC) on the FCoE priority (the FCoE IEEE 802.1p code point) on the interface.

- If FCoE is the only application that you want DCBX to advertise on an interface, then you do not need to configure an application map and apply it to the interface. By default, when an interface has no application map, and the interface carries traffic mapped to the FCoE forwarding class, and PFC is enabled on the FCoE priority, the interface advertises FCoE TLVs (autonegotiation mode). DCBX exchanges FCoE application protocol TLVs by default until you apply an application map to the interface, remove the FCoE traffic from the interface (you can do this by removing the or editing the classifier for FCoE traffic), or disable PFC on the FCoE priority.

If you apply an application map to an interface that did not have an application map and was exchanging FCoE application TLVs, and you do not include the FCoE application in the application map, the interface stops exchanging FCoE TLVs. Every interface that has an application map must have FCoE included in the application map (and PFC enabled on the FCoE priority) in order for DCBX to exchange FCoE TLVs.

Mapping an application to code points does two things:

- Maps incoming traffic with the same code points to that application
- Allows you to configure classifiers that map incoming application traffic, by code point, to a forwarding class and a loss priority, in order to apply class of service (CoS) to application traffic and prioritize application traffic

You apply an application map to an interface to enable DCBX application protocol exchange on that interface for each application specified in the application map. All of the applications that you want an interface to advertise must be configured in the application map that you apply to the interface, with the previously noted exception for the FCoE application when FCoE is the only application for which you want DCBX to exchange protocol TLVs on an interface.

## Classifying and Prioritizing Application Traffic

When traffic arrives at an interface, the interface classifies the incoming traffic based on its code points. Classifiers map code points to loss priorities and forwarding classes. The loss priority prioritizes the traffic. The forwarding class determines the traffic output queue and CoS service level.

When you map an application to an IEEE 802.1p code point in an application map and apply the application map to an interface, incoming traffic on the interface that matches the application code points is mapped to the appropriate application. The application receives the loss priority and the CoS associated with the forwarding class for those code points, and is placed in the output queue associated with the forwarding class.

You can use the default classifier or you can configure a classifier to map the application code points defined in the application map to forwarding classes and loss priorities.

## Enabling Interfaces to Exchange Application Protocol Information

Each interface with the **fcoe** forwarding class and PFC enabled on the FCoE code point is enabled for FCoE application protocol exchange by default until you apply an application map to the interface. If you apply an application map to an interface and you want that interface to exchange FCoE application protocol TLVs, you must include the FCoE application in the application map. (In all cases, to achieve lossless transport, you must also enable PFC on the FCoE code point or code points.)

Except when FCoE is the only protocol you want DCBX to advertise on an interface, interfaces on which you want to exchange application protocol TLVs must include the following two items:

- The application map that contains the application(s)
- A classifier



**NOTE:** You must also enable PFC on the code point of any traffic for which you want to achieve lossless transport.

## Disabling DCBX Application Protocol Exchange

To disable DCBX application protocol exchange for all applications on an interface, issue the **set protocols dcbx interface *interface-name* applications no-auto-negotiation** command.

You can also disable DCBX application protocol exchange for applications on an interface by deleting the application map from the interface, or by deleting a particular application from the application map. However, when you delete an application from an application map, the application protocol is no longer exchanged on any interface which uses that application map.

On interfaces that use IEEE DCBX mode to exchange DCBX parameters, you can disable sending the enhanced transmission selection (ETS) Recommendation TLV to the peer if you want an asymmetric ETS configuration between the peers.

### Related Documentation

- [Understanding DCBX on page 6514](#)
- [Configuring DCBX Autonegotiation on page 6524](#)
- [Disabling the ETS Recommendation TLV on page 6527](#)
- [Defining an Application for DCBX Application Protocol TLV Exchange on page 6532](#)
- [Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 6533](#)
- [Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 6534](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 6535](#)

## Defining an Application for DCBX Application Protocol TLV Exchange

Define each application for which you want DCBX to exchange application protocol information. You can define Layer 2 and Layer 4 applications. After you define applications, you map them to IEEE 802.1p code points, and then apply the application map to the interfaces on which you want DCBX to exchange application protocol information with connected peers. (See *Related Documentation* for how to configure application maps and apply them to interfaces, and for an example of the entire procedure that also includes classifier configuration.)



**NOTE:** In Junos OS Release 12.1, the FCoE application was configured by default, so you did not need to configure it in an application map. In Junos OS Release 12.2, if you want DCBX to advertise the FCoE application on an interface and you apply an application map to that interface, you must explicitly configure FCoE in the application map. You also must enable priority-based flow control (PFC) on the FCoE code point on all interfaces that you want to advertise FCoE. If you apply an application map to an interface, the interface sends DCBX TLVs only for the applications configured in the application map.

Define Layer 2 applications by mapping an application name to an EtherType. Define Layer 4 applications by mapping an application name to a protocol (TCP or UDP) and a destination port.

- To define a Layer 2 application, specify the name of the application and its EtherType:

```
[edit applications]
user@switch# set application application-name ether-type ether-type
```

For example, to configure an application named **PTP** (for Precision Time Protocol) that uses the EtherType **0x88F7**:

```
user@switch# set applications application ptp ether-type 0x88F7
```

- To define a Layer 4 application, specify the name of the application, its protocol (TCP or UDP), and its destination port:

```
[edit]
user@switch# set applications application application-name protocol (tcp | udp)
destination-port port-value
```

For example, to configure an application named **iscsi** (for Internet Small Computer System Interface) that uses the protocol **TCP** and the destination port **3260**:

```
user@switch# set applications application iscsi protocol tcp destination-port 3260
```

### Related Documentation

- [Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 6533](#)
- [Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 6534](#)

- [Configuring DCBX Autonegotiation on page 6524](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 6535](#)
- [Example: Configuring DCBX to Support an iSCSI Application](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 6528](#)
- [show dcbx neighbors on page 6581](#)

## Configuring an Application Map for DCBX Application Protocol TLV Exchange

After you define applications for which you want to exchange DCBX application protocol information, map the applications to IEEE 802.1p code points. The IEEE 802.1p code points identify incoming traffic and allow you to map that traffic to the desired application. You then apply the application map to the interfaces on which you want DCBX to exchange application protocol information with connected peers. (See *Related Documentation* for how to define applications and apply the application map to interfaces, and for an example of the entire procedure that also includes classifier configuration.)



**NOTE:** In Junos OS Release 12.1, the FCoE application was configured by default, so you did not need to configure it in an application map. In Junos OS Release 12.2, if you want DCBX to advertise the FCoE application on an interface and you apply an application map to that interface, you must explicitly configure FCoE in the application map. You also must enable priority-based flow control (PFC) on the FCoE code point on all interfaces that you want to advertise FCoE. If you apply an application map to an interface, the interface sends DCBX TLVs only for the applications configured in the application map.

Configure an application map by creating an application map name and mapping an application to one or more IEEE 802.1p code points.

- To define an application map, specify the name of the application map, the name of the application, and the IEEE 802.1p code points of the incoming traffic that you want to associate with the application in the application map:

```
[edit policy-options]
user@switch# set application-maps application-map-name application application-name
code-points [aliases] [bit-patterns]
```

For example, to configure an application map named **ptp-app-map** that includes an application named **PTP** (for Precision Time Protocol) and map the application to IEEE 802.1p code points **001** and **101**:

```
user@switch# set policy-options application-maps ptp-app-map application ptp code points
[001 101]
```

### Related Documentation

- [Defining an Application for DCBX Application Protocol TLV Exchange on page 6532](#)
- [Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 6534](#)

- [Configuring DCBX Autonegotiation on page 6524](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 6535](#)
- [Example: Configuring DCBX to Support an iSCSI Application](#)
- [show dcbx neighbors on page 6581](#)

---

## Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange

---

After you define applications and map them to IEEE 802.1p code points in an application map, apply the application map to the interfaces on which you want DCBX to exchange the application protocol information with connected peers. (See *Related Documentation* for how to define applications and configure application maps to interfaces, and for an example of the entire procedure that also includes classifier configuration.)



**NOTE:** In Junos OS Release 12.1, the FCoE application was configured by default, so you did not need to configure it in an application map. In Junos OS Release 12.2, if you want DCBX to advertise the FCoE application on an interface and you apply an application map to that interface, you must explicitly configure FCoE in the application map. You also must enable priority-based flow control (PFC) on the FCoE code point on all interfaces that you want to advertise FCoE. If you apply an application map to an interface, the interface sends DCBX TLVs only for the applications configured in the application map.

- 
- To apply an application map to a DCBX interface, specify the DCBX interface and the application map name:

[edit protocols]

```
user@switch# set dcbx interface interface-name application-map application-map-name
```

For example, to apply an application map named **ptp-app-map** on interface **xe-0/0/11**:

```
user@switch# set protocols dcbx interface xe-0/0/11 application-map ptp-app-map
```

### Related Documentation

- [Defining an Application for DCBX Application Protocol TLV Exchange on page 6532](#)
- [Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 6533](#)
- [Configuring DCBX Autonegotiation on page 6524](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 6535](#)
- [Example: Configuring DCBX to Support an iSCSI Application](#)
- [show dcbx neighbors on page 6581](#)

## Example: Configuring DCBX Application Protocol TLV Exchange

Data Center Bridging Capability Exchange protocol (DCBX) discovers the data center bridging (DCB) capabilities of connected peers by exchanging application configuration information. DCBX detects feature misconfiguration and mismatches and can configure DCB on peers. DCBX is an extension of the Link Layer Discovery Protocol (LLDP). LLDP must remain enabled on every interface on which you want to use DCBX.



**NOTE:** LLDP and DCBX are enabled by default on all interfaces.

The switch supports DCBX application protocol exchange for Layer 2 and Layer 4 applications such as the Internet Small Computer System Interface (iSCSI). You specify applications by EtherType (for Layer 2 applications) or by the destination port and protocol (for Layer 4 applications; the protocol can be either TCP or UDP).

The switch handles Fibre Channel over Ethernet (FCoE) application protocol exchange differently than other protocols in some cases:

- If FCoE is the only application for which you want to enable DCBX application protocol TLV exchange on an interface, you do not have to explicitly configure the FCoE application or an application map. By default, the switch exchanges FCoE application protocol TLVs on all interfaces that carry FCoE traffic (traffic mapped to the **fcoe** forwarding class) and have priority-based flow control (PFC) enabled on the FCoE priority (the FCoE IEEE 802.1p code point). The default priority mapping for the FCoE application is IEEE 802.1p code point 011 (the default **fcoe** forwarding class code point).
- If you want an interface to use DCBX to exchange application protocol TLVs for any other applications in addition to FCoE, you must configure the applications (including FCoE), define an application map (including FCoE), and apply the application map to the interface. If you apply an application map to an interface, you must explicitly configure the FCoE application, or the interface does not exchange FCoE application protocol TLVs.

This example shows how to configure interfaces to exchange both Layer 2 and Layer 4 applications by configuring one interface to exchange iSCSI and FCoE application protocol information and configuring another interface to exchange iSCSI and Precision Time Protocol (PTP) application protocol information.

- [Requirements on page 6535](#)
- [Overview on page 6536](#)
- [Configuration on page 6539](#)
- [Verification on page 6541](#)

## Requirements

This example uses the following hardware and software components:

- Juniper Networks QFX Series device

- Junos OS Release 12.1 or later for the QFX Series

## Overview

The switch supports DCBX application protocol exchange for:

- Layer 2 applications, defined by EtherType
- Layer 4 applications, defined by destination port and protocol



**NOTE:** DCBX also advertises PFC and enhanced transmission selection (ETS) information. See [“Configuring DCBX Autonegotiation” on page 6524](#) for how DCBX negotiates and advertises configuration information for these features and for the applications.

DCBX is configured on a per-interface basis for each supported feature or application. For applications that you want to enable for DCBX application protocol exchange, you must:

- Define the application name and configure the EtherType or the destination port and protocol (TCP or UDP) of the application. Use the EtherType for Layer 2 applications, and use the destination port and protocol for Layer 4 protocols.
- Map the application to an IEEE 802.1p code point in an application map.
- Add the application map to DCBX interface.

In addition, for all applications (including FCoE, even when you do not use an application map), you either must create an IEEE 802.1p classifier and apply it to the appropriate ingress interfaces or use the default classifier. A classifier maps the code points of incoming traffic to a forwarding class and a loss priority so that ingress traffic is assigned to the correct class of service (CoS). The forwarding class determines the output queue on the egress interface.

If you do not create classifiers, trunk and tagged-access ports use the unicast IEEE 802.1 default trusted classifier. [Table 519](#) shows the default mapping of IEEE 802.1 code-point values to unicast forwarding classes and loss priorities for ports in trunk mode or tagged-access mode. [Table 520](#) shows the default untrusted classifier IEEE 802.1 code-point values to unicast forwarding class mapping for ports in access mode.

**Table 519: Default IEEE 802.1 Classifiers for Trunk Ports and Tagged-Access Ports (Default Trusted Classifier)**

| Code Point | Forwarding Class | Loss Priority |
|------------|------------------|---------------|
| be (000)   | best-effort      | low           |
| be1 (001)  | best-effort      | low           |
| ef (010)   | best-effort      | low           |



**Table 519: Default IEEE 802.1 Classifiers for Trunk Ports and Tagged-Access Ports (Default Trusted Classifier) (continued)**

| Code Point | Forwarding Class | Loss Priority |
|------------|------------------|---------------|
| ef1 (011)  | fcoe             | low           |
| af11 (100) | no-loss          | low           |
| af12 (101) | best-effort      | low           |
| nc1 (110)  | network-control  | low           |
| nc2 (111)  | network-control  | low           |

**Table 520: Default IEEE 802.1 Unicast Classifiers for Access Ports (Default Untrusted Classifier)**

| Code Point | Forwarding Class | Loss Priority |
|------------|------------------|---------------|
| 000        | best-effort      | low           |
| 001        | best-effort      | low           |
| 010        | best-effort      | low           |
| 011        | best-effort      | low           |
| 100        | best-effort      | low           |
| 101        | best-effort      | low           |
| 110        | best-effort      | low           |
| 111        | best-effort      | low           |

### Topology

This example shows how to configure DCBX application protocol exchange for three protocols (iSCSI, PTP, and FCoE) on two interfaces. One interface exchanges iSCSI and FCoE application protocol information, and the other interface exchanges iSCSI and PTP application protocol information.



**NOTE:** You must map FCoE traffic to the interfaces on which you want to forward FCoE traffic. You must also enable PFC on the FCoE interfaces and create an ingress classifier for FCoE traffic, or else use the default classifier.

Table 521 shows the configuration components for this example.

**Table 521: Components of DCBX Application Protocol Exchange Configuration Topology**

| Component                                                         | Settings                                                                                                                                                                                                                                                                                                                                                     |
|-------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Hardware                                                          | QFX Series device                                                                                                                                                                                                                                                                                                                                            |
| LLDP                                                              | Enabled by default on Ethernet interfaces                                                                                                                                                                                                                                                                                                                    |
| DCBX                                                              | Enabled by default on Ethernet interfaces                                                                                                                                                                                                                                                                                                                    |
| iSCSI application (Layer 4)                                       | Application name— <b>iscsi</b><br>protocol— <b>TCP</b><br>destination-port— <b>3260</b><br>code-points— <b>111</b>                                                                                                                                                                                                                                           |
| PTP application (Layer 2)                                         | Application name— <b>ptp</b><br>ether-type— <b>0x88F7</b><br>code-points— <b>001, 101</b>                                                                                                                                                                                                                                                                    |
| FCoE application (Layer 2)                                        | Application name— <b>fcoe</b><br>ether-type— <b>0x8906</b><br>code-points— <b>011</b><br><br><b>NOTE:</b> You explicitly configure the FCoE application because you are applying an application map to the interface. When you apply an application map to an interface, all applications must be explicitly configured and included in the application map. |
| Application maps                                                  | <b>dcbx-iscsi-fcoe-app-map</b> —Maps the iSCSI and FCoE applications to IEEE 802.1p code points<br><br><b>dcbx-iscsi-ptp-app-map</b> —Maps iSCSI and PTP applications to IEEE 802.1p code points                                                                                                                                                             |
| Interfaces                                                        | <b>xe-0/0/10</b> —Configured to exchange FCoE and iSCSI application TLVs (uses application map <b>dcbx-iscsi-fcoe-app-map</b> , carries FCoE traffic, and has PFC enabled on the FCoE priority)<br><br><b>xe-0/0/11</b> —Configured to exchange iSCSI and PTP application TLVs (uses application map <b>dcbx-iscsi-ptp-app-map</b> )                         |
| PFC congestion notification profile for FCoE application exchange | <b>fcoe-cnp:</b> <ul style="list-style-type: none"> <li>Code point—<b>011</b></li> <li>Interface—<b>xe-0/0/10</b></li> </ul>                                                                                                                                                                                                                                 |

**Table 521: Components of DCBX Application Protocol Exchange Configuration Topology (*continued*)**

| Component                                                                                                         | Settings                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Behavior aggregate classifiers (map forwarding classes to incoming packets by the packet's IEEE 802.1 code point) | <p><b>fcoe-iscsi-cl1:</b></p> <ul style="list-style-type: none"> <li>Maps the <b>fcoe</b> forwarding class to the IEEE 802.1p code point used for the FCoE application (011) and a loss priority of <b>high</b></li> <li>Maps the <b>network-control</b> forwarding class to the IEEE 802.1p code point used for the iSCSI application (111) and a loss priority of <b>high</b></li> <li>Applied to interface <b>xe-0/0/10</b></li> </ul> <p><b>iscsi-ntp-cl2:</b></p> <ul style="list-style-type: none"> <li>Maps the <b>network-control</b> forwarding class to the IEEE 802.1p code point used for the iSCSI application (111) and a loss priority of <b>low</b></li> <li>Maps the <b>best-effort</b> forwarding class to the IEEE 802.1p code points used for the PTP application (001 and 101) and a loss priority of <b>low</b></li> <li>Applied to interface <b>xe-0/0/11</b></li> </ul> |



**NOTE:** This example does not include scheduling (bandwidth allocation) configuration or lossless configuration for the iSCSI forwarding class.

## Configuration

### CLI Quick Configuration

To quickly configure DCBX application protocol exchange, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set applications application iSCSI protocol tcp destination-port 3260
set applications application FCoE ether-type 0x8906
set applications application PTP ether-type 0x88F7
set policy-options application-maps dcbx-iscsi-fcoe-app-map application iSCSI code-points 111
set policy-options application-maps dcbx-iscsi-fcoe-app-map application FCoE code-points 011
set policy-options application-maps dcbx-iscsi-ntp-app-map application iSCSI code-points 111
set policy-options application-maps dcbx-iscsi-ntp-app-map application PTP code-points [001 101]
set protocols dcbx interface xe-0/0/10 application-map dcbx-iscsi-fcoe-app-map
set protocols dcbx interface xe-0/0/11 application-map dcbx-iscsi-ntp-app-map
set class-of-service congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011 pfc
set class-of-service interfaces xe-0/0/10 congestion-notification-profile fcoe-cnp
set class-of-service classifiers ieee-802.1 fcoe-iscsi-cl1 import default forwarding-class fcoe
loss-priority high code-points 011
set class-of-service classifiers ieee-802.1 fcoe-iscsi-cl1 import default forwarding-class
network-control loss-priority high code-points 111
set class-of-service classifiers ieee-802.1 iscsi-ntp-cl2 import default forwarding-class
network-control loss-priority low code-points 111
set class-of-service classifiers ieee-802.1 iscsi-ntp-cl2 import default forwarding-class best-effort
loss-priority low code-points [001 101]
set class-of-service interfaces xe-0/0/10 unit 0 classifiers ieee-802.1 fcoe-iscsi-cl1
```

```
set class-of-service interfaces xe-0/0/11 unit 0 classifiers ieee-802.1 iscsi-ptp-cl2
```

### Configuring DCBX Application Protocol TLV Exchange

#### Step-by-Step Procedure

To define the applications, map the applications to IEEE 802.1p code points, apply the applications to interfaces, and create classifiers for DCBX application protocol exchange:

1. Define the iSCSI application by specifying its protocol and destination port, and define the FCoE and PTP applications by specifying their EtherTypes.  

```
[edit applications]
user@switch# set application iSCSI protocol tcp destination-port 3260
user@switch# set application FCoE ether-type 0x8906
user@switch# set application PTP ether-type 0x88F7
```
2. Define an application map that maps the iSCSI and FCoE applications to IEEE 802.1p code points.  

```
[edit policy-options]
user@switch# set application-maps dcbx-iscsi-fcoe-app-map application iSCSI code-points 111
user@switch# set application-maps dcbx-iscsi-fcoe-app-map application FCoE code-points 011
```
3. Define the application map that maps the iSCSI and PTP applications to IEEE 802.1p code points.  

```
[edit policy-options]
user@switch# set application-maps dcbx-iscsi-ptp-app-map application iSCSI code-points 111
user@switch# set application-maps dcbx-iscsi-ptp-app-map application PTP code-points [001 101]
```
4. Apply the iSCSI and FCoE application map to interface **xe-0/0/10**, and apply the iSCSI and PTP application map to interface **xe-0/0/11**.  

```
[edit protocols dcbx]
user@switch# set interface xe-0/0/10 application-map dcbx-iscsi-fcoe-app-map
user@switch# set interface xe-0/0/11 application-map dcbx-iscsi-ptp-app-map
```
5. Create the congestion notification profile to enable PFC on the FCoE code point (**011**), and apply the congestion notification profile to interface **xe-0/0/10**.  

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011 pfc
user@switch# set interfaces xe-0/0/10 congestion-notification-profile fcoe-cnp
```
6. Configure the classifier to apply to the interface that exchanges iSCSI and FCoE application information.  

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 fcoe-iscsi-cl1 import default forwarding-class fcoe loss-priority high code-points 011
user@switch# set ieee-802.1 fcoe-iscsi-cl1 import default forwarding-class network-control loss-priority high code-points 111
```
7. Configure the classifier to apply to the interface that exchanges iSCSI and PTP application information.  

```
[edit class-of-service classifiers]
```

```

user@switch# set ieee-802.1 iscsi-ptp-cl2 import default forwarding-class network-control
loss-priority low code-points 111
user@switch# set ieee-802.1 iscsi-ptp-cl2 import default forwarding-class best-effort
loss-priority low code-points [001101]

```

8. Apply the classifiers to the appropriate interfaces.

```

[edit class-of-service]
user@switch# set interfaces xe-0/0/10 unit 0 classifiers ieee-802.1 fcoe-iscsi-cl1
user@switch# set interfaces xe-0/0/11 unit 0 classifiers ieee-802.1 iscsi-ptp-cl2

```

## Verification

To verify that DCBX application protocol exchange configuration has been created and is operating properly, perform these tasks:

- [Verifying the Application Configuration on page 6541](#)
- [Verifying the Application Map Configuration on page 6541](#)
- [Verifying DCBX Application Protocol Exchange Interface Configuration on page 6542](#)
- [Verifying the PFC Configuration on page 6542](#)
- [Verifying the Classifier Configuration on page 6543](#)

### Verifying the Application Configuration

**Purpose** Verify that DCBX applications have been configured.

**Action** List the applications by using the configuration mode command **show applications**:

```

user@switch# show applications
application iSCSI {
 protocol tcp;
 destination-port 3260;
}

application fcoe {
 ether-type 0x8906;
}

application ptp {
 ether-type 0x88F7;
}

```

**Meaning** The **show applications** configuration mode command lists all of the configured applications and either their protocol and destination port (Layer 4 applications) or their EtherType (Layer 2 applications). The command output shows that the iSCSI application is configured with the **tcp** protocol and destination port **3260**, the FCoE application is configured with the EtherType **0x8906**, and that the PTP application is configured with the EtherType **0x88F7**.

### Verifying the Application Map Configuration

**Purpose** Verify that the application maps have been configured.

**Action** List the application maps by using the configuration mode command **show policy-options application-maps**:

```
user@switch# show policy-options application-maps
dcbx-iscsi-fcoe-app-map {
 application iSCSI code-points 111;
 application FCoE code-points 011;
}

dcbx-iscsi-ntp-app-map {
 application iSCSI code-points 111;
 application PTP code-points [001 101];
}
```

**Meaning** The **show policy-options application-maps** configuration mode command lists all of the configured application maps and the applications that belong to each application map. The command output shows that there are two application maps, **dcbx-iscsi-fcoe-app-map** and **dcbx-iscsi-ntp-app-map**.

The application map **dcbx-iscsi-fcoe-app-map** consists of the iSCSI application, which is mapped to IEEE 802.1p code point 111, and the FCoE application, which is mapped to IEEE 802.1p code point 011.

The application map **dcbx-iscsi-ntp-app-map** consists of the iSCSI application, which is mapped to IEEE 802.1p code point 111, and the PTP application, which is mapped to IEEE 802.1p code points 001 and 101.

---

### Verifying DCBX Application Protocol Exchange Interface Configuration

**Purpose** Verify that the application maps have been applied to the correct interfaces.

**Action** List the application maps by using the configuration mode command **show protocols dcbx**:

```
user@switch# show protocols dcbx
interface xe-0/0/10.0 {
 application-map dcbx-iscsi-fcoe-app-map;
}

interface xe-0/0/11.0 {
 application-map dcbx-iscsi-ntp-app-map;
}
```

**Meaning** The **show protocols dcbx** configuration mode command lists whether the interfaces are enabled for DCBX and lists the application map applied to each interface. The command output shows that interfaces **xe-0/0/10.0** and **xe-0/0/11.0** are enabled for DCBX, and that interface **xe-0/0/10.0** uses application map **dcbx-iscsi-fcoe-app-map**, and interface **xe-0/0/11.0** uses application map **dcbx-iscsi-ntp-app-map**.

---

### Verifying the PFC Configuration

**Purpose** Verify that PFC has been enabled on the FCoE code point and applied to the correct interface.

**Action** Display the PFC configuration to verify that PFC is enabled on the FCoE code point (011) in the congestion notification profile **fcoe-cnp** by using the configuration mode command **show class-of-service congestion-notification-profile**:

```
user@switch# show class-of-service congestion-notification-profile
fcoe-cnp {
 input {
 ieee-802.1 {
 code-point 011 {
 pfc;
 }
 }
 }
}
```

Display the class-of-service (CoS) interface information to verify that the correct interface has PFC enabled for the FCoE application by using the configuration mode command **show class-of-service interfaces**:

```
user@switch# show class-of-service interfaces
xe-0/0/10 {
 congestion-notification-profile fcoe-cnp;
}
```



**NOTE:** The sample output does not include all of the information this command can show. The output is abbreviated to focus on verifying the PFC configuration.

**Meaning** The **show class-of-service congestion-notification-profile** configuration mode command lists the configured congestion notification profiles. The command output shows that the congestion notification profile **fcoe-cnp** has been configured and has enabled PFC on the IEEE 802.1p code point 011 (the default FCoE code point).

The **show class-of-service interfaces** configuration mode command shows the interface CoS configuration. The command output shows that the congestion notification profile **fcoe-cnp**, which enables PFC on the FCoE code point, is applied to interface **xe-0/0/10**.

### Verifying the Classifier Configuration

**Purpose** Verify that the classifiers have been configured and applied to the correct interfaces.

**Action** Display the classifier configuration by using the configuration mode command **show class-of-service**:

```
user@switch# show class-of-service
classifiers {
 ieee-802.1 fcoe-iscsi-cl1 {
 import default;
 forwarding-class network-control {
 loss-priority high code-points 111;
 }
 forwarding-class fcoe {
```

```

 loss-priority high code-points 011;
 }
}
ieee-802.1 iscsi-ntp-cl2 {
 import default;
 forwarding-class network-control {
 loss-priority low code-points 111;
 }
 forwarding-class best-effort {
 loss-priority low code-points [001 101];
 }
}
}
interfaces {
 xe-0/0/10 {
 congestion-notification-profile fcoe-cnp;
 unit 0 {
 classifiers {
 ieee-802.1 fcoe-iscsi-cl1;
 }
 }
 }
 xe-0/0/11 {
 unit 0 {
 classifiers {
 ieee-802.1 iscsi-ntp-cl2;
 }
 }
 }
}
}

```



**NOTE:** The sample output does not include all of the information this command can show. The output is abbreviated to focus on verifying the classifier configuration.

**Meaning** The **show class-of-service** configuration mode command lists the classifier and CoS interface configuration, as well as other information not shown in this example. The command output shows that there are two classifiers configured, **fcoe-iscsi-cl1** and **iscsi-ntp-cl2**.

Classifier **fcoe-iscsi-cl1** uses the **default** classifier as a template and edits the template as follows:

- The forwarding class **network-control** is set to a loss priority of **high** and is mapped to code point **111** (the code point mapped to the iSCSI application).
- The forwarding class **fcoe** is set to a loss priority of **high** and is mapped to code point **011** (the code point mapped by default to the FCoE application).

Classifier **iscsi-ntp-cl2** uses the **default** classifier as a template and edits the template as follows:



- The forwarding class **network-control** is set to a loss priority of **low** and is mapped to IEEE 802.1p code point **111** (the code point mapped to the iSCSI application).
- The forwarding class **best-effort** is set to a loss priority of **low** and is mapped to IEEE 802.1p code points **001** and **101** (the code points mapped by default to the PTP application).

The command output also shows that classifier **fcoe-iscsi-cl1** is mapped to interface **xe-0/0/10.0** and that classifier **iscsi-ntp-cl2** is mapped to interface **xe-0/0/11.0**.

#### Related Documentation

- [Defining an Application for DCBX Application Protocol TLV Exchange on page 6532](#)
- [Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 6533](#)
- [Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 6534](#)
- [Configuring DCBX Autonegotiation on page 6524](#)
- [show dcbx on page 6580](#)
- [show dcbx neighbors on page 6581](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 6528](#)



# Learn About Technology

- [Data Center Technology Overview Videos on page 6547](#)

## Data Center Technology Overview Videos

---

Juniper Information Experience (IX) videos provide brief, high-level overviews of data center technologies and concepts. Each video runs approximately one-and-a-half to two minutes in length. This document contains SDN-related videos and links to conceptual documents that contain other data center technology videos:

- [Learn About Video: Why Do We Need an IP Fabric? on page 6547](#)
- [Learn About Video: What is the Best Control Plane Protocol to Use in a Data Center IP Fabric? on page 6547](#)
- [Learn About Video: Why Use an Overlay Network in a Data Center? on page 6547](#)
- [Conceptual Documents That Contain Technology Overview Videos on page 6548](#)

### Learn About Video: Why Do We Need an IP Fabric?

The video *Why Do We Need an IP Fabric?* presents a brief overview of IP Fabric use cases.



Video: [Why Do We Need an IP Fabric?](#)

### Learn About Video: What is the Best Control Plane Protocol to Use in a Data Center IP Fabric?

The video *What is the Best Control Plane Protocol to Use in a Data Center IP Fabric?* presents a brief overview of the arguments for using Border Gateway Protocol (BGP) as the data center IP fabric control plane protocol.



Video: [What is the Best Control Plane Protocol to Use in a Data Center IP Fabric?](#)

### Learn About Video: Why Use an Overlay Network in a Data Center?

The video *Why Use an Overlay Network in a Data Center?* presents a brief overview of the advantages of data center overlay networks.



Video: [Why Use an Overlay Network in a Data Center?](#)

## Conceptual Documents That Contain Technology Overview Videos

The following conceptual documents include brief video overviews of the technology:

- [Understanding DCB Features and Requirements on page 6490](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 6765](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 6493](#)
- [Understanding DCBX on page 6514](#)
- [Understanding PFC Functionality Across Layer 3 Interfaces on page 6921](#)
- [Virtual Chassis Fabric Overview](#)
- [Understanding In-Service Software Upgrade \(ISSU\) and In-Service Software Upgrade \(ISSU\) System Requirements \(same video\)](#)

## PART 89

# Configuration Statements and Operational Commands

- [Configuration Statements \(FCoE on a Transit Switch\) on page 6551](#)
- [Configuration Statements \(DCBX and PFC\) on page 6555](#)
- [Operational Commands \(FCoE on a Transit Switch\) on page 6575](#)
- [Operational Commands \(DCBX and PFC\) on page 6579](#)



## CHAPTER 248

# Configuration Statements (FCoE on a Transit Switch)

- [family fcoe on page 6552](#)
- [oxid on page 6553](#)

## family fcoe

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | QFX Series Standalone Switches<br><br>family fcoe {<br>oxid (enable   disable);<br>}<br><br>QFabric Systems<br><br>family fcoe {<br>ethernet-interfaces {<br>node-group ( <i>node-group-name</i>   all) {<br>oxid (enable   disable);<br>}<br>}<br>fabric-interfaces {<br>node-group ( <i>node-group-name</i>   all) {<br>oxid (enable   disable);<br>}<br>}<br>} |
| <b>Hierarchy Level</b>          | [edit forwarding-options hash-key]                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.3 for the QFX Series.<br>Ethernet-interfaces and fabric-interfaces statements introduced in Junos OS Release 13.2X52-D10 for the QFX Series.                                                                                                                                                                          |
| <b>Description</b>              | Configure whether or not to use the originator exchange identifier (OxID) field for hash control for FCoE traffic load balancing.                                                                                                                                                                                                                                 |
| <b>Options</b>                  | The statement is explained separately.                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Enabling and Disabling CoS OxID Hash Control on Standalone Switches on page 6485</a></li><li>• <a href="#">Understanding OxID Hash Control for FCoE Traffic Load Balancing on Standalone Switches on page 6484</a></li></ul>                                                                                  |



## oxid

|                                 |                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | oxid (enable   disable)                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | <p>QFX Series Standalone Switches</p> <p>[edit forwarding-options hash-key family fcoe]</p> <p>QFabric Systems</p> <p>[edit forwarding-options hash-key family fcoe ethernet-interfaces node-group (node-group-name   all) {}]</p> <p>[edit forwarding-options hash-key family fcoe fabric-interfaces node-group (node-group-name   all) {}]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 12.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2X52-D10 for the QFabric System.</p>                                                                                                                                                                                     |
| <b>Description</b>              | Enable or disable whether the switch uses the originator exchange identifier (OxID) field for hash control for FCoE traffic load balancing.                                                                                                                                                                                                      |
| <b>Default</b>                  | OxID hash control is enabled by default.                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <b>oxid (enable   disable)</b> —Enable or disable whether the switch uses the OxID hash control field for FCoE traffic load balancing.                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Enabling and Disabling CoS OxID Hash Control on Standalone Switches on page 6485</a></li> <li>• <a href="#">Understanding OxID Hash Control for FCoE Traffic Load Balancing on Standalone Switches on page 6484</a></li> </ul>                                                              |



# Configuration Statements (DCBX and PFC)

- [application \(Application Maps\) on page 6556](#)
- [application \(Applications\) on page 6557](#)
- [application-map on page 6558](#)
- [application-maps on page 6559](#)
- [applications \(Applications\) on page 6560](#)
- [applications \(DCBX\) on page 6561](#)
- [code-points \(Application Maps\) on page 6562](#)
- [dcbx on page 6563](#)
- [dcbx-version on page 6564](#)
- [destination-port \(Applications\) on page 6565](#)
- [disable \(DCBX\) on page 6566](#)
- [enhanced-transmission-selection on page 6567](#)
- [ether-type on page 6568](#)
- [interface \(DCBX\) on page 6569](#)
- [no-recommendation-tlv on page 6570](#)
- [policy-options on page 6571](#)
- [priority-flow-control on page 6572](#)
- [protocol \(Applications\) on page 6573](#)
- [recommendation-tlv on page 6574](#)

## application (Application Maps)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>application <i>application-name</i> {<br/>    <i>code-points</i> [ <i>aliases</i> ] [ <i>bit-patterns</i> ];<br/>}</code>                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit policy-options <b>application-maps</b> <i>application-map-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Add an application to an application map and define the application's code points.                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <i>application-name</i> —Name of the application.<br><br>The remaining statement is explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 6533</a></li><li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 6535</a></li><li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 6528</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li></ul> |

## application (Applications)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> application <i>application-name</i> {     <i>destination-port</i> <i>port-value</i>;     <i>protocol</i> (tcp   udp);     <i>ether-type</i> <i>type</i>; } </pre>                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit applications]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Configure properties to define an application.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <p><i>application-name</i>—Name of the application.</p> <p>The statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Defining an Application for DCBX Application Protocol TLV Exchange on page 6532</a></li> <li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 6535</a></li> <li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 6528</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li> </ul> |

## application-map

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>application-map <i>application-map-name</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit <a href="#">protocols dcbx interface interface-name</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Specify an application map to apply to an interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <i>application-map-name</i> —Name of the application map.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show dcbx neighbors on page 6581</a></li><li>• <a href="#">Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 6534</a></li><li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 6535</a></li><li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 6528</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li></ul> |

## application-maps

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> application-maps <i>application-map-name</i> {   application <i>application-name</i> {     code-points [ <i>aliases</i> ] [ <i>bit-patterns</i> ];   } } </pre>                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit policy-options]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Define an application map by specifying the applications that belong to the application map.                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <p><i>application-map-name</i>—Name of the application map.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 6533</a></li> <li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 6535</a></li> <li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 6528</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li> </ul> |

## applications (Applications)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>applications {<br/>  application application-name {<br/>    destination-port port-value;<br/>    protocol (tcp   udp);<br/>    ether-type type;<br/>  }<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Define applications that DCBX advertises.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | The statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Defining an Application for DCBX Application Protocol TLV Exchange on page 6532</a></li><li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 6535</a></li><li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 6528</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li></ul> |



---

## applications (DCBX)

---

|                                 |                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>applications {<br/>  fcoe {<br/>    no-auto-negotiation;<br/>  }<br/>}</pre>                                                                                                     |
| <b>Hierarchy Level</b>          | [edit <a href="#">protocols dcbx interface</a> <i>interface-name</i> ]                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 12.1 for the EX Series                                                  |
| <b>Description</b>              | Configure Data Center Bridging Capability Exchange protocol (DCBX) applications on an interface.                                                                                      |
| <b>Options</b>                  | The remaining statements are explained separately.                                                                                                                                    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show dcbx neighbors on page 6581</a></li><li>• <a href="#">Understanding DCB Features and Requirements on page 6490</a></li></ul> |

## code-points (Application Maps)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>code-points [ <i>aliases</i> ] [ <i>bit-patterns</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit policy-options <b>application-maps</b> <i>application-map-name</i> <b>application</b> <i>application-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Define one or more code-point aliases or bit sets for an application.                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <i>aliases</i> —Name of the alias or aliases.<br><br><i>bit-patterns</i> —Value of the code-point bits, in decimal form.                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 6533</a></li><li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 6535</a></li><li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 6528</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li></ul> |

## dcbx


|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> dcbx {   disable;   interface (interface-name   all) {     disable;     application-map application-map-name;     applications {       no-auto-negotiation;     }     enhanced-transmission-selection {       no-auto-negotiation;       no-recommendation-tlv;       recommendation-tlv {         no-auto-negotiation;       }     }     dcbx-version (auto-negotiate   ieee-dcbx   dcbx-version-1.01);     priority-flow-control {       no-auto-negotiation;     }   } } </pre> |
| <b>Hierarchy Level</b>          | [edit <a href="#">protocols</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 11.3 for EX Series switches.</p> <p><b>mode</b> and <b>recommendation-tlv</b> statements introduced in Junos OS Release 12.2 for the QFX Series.</p>                                                                                                                                                                                                                |
| <b>Description</b>              | <p>Configure DCBX properties. DCBX is an extension of Link Layer Discovery Protocol (LLDP), and LLDP must remain enabled on every interface for which you want to use DCBX. If you attempt to enable DCBX on an interface on which LLDP is disabled, the configuration commit fails.</p>                                                                                                                                                                                                 |
| <b>Options</b>                  | The statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show dcbx neighbors on page 6581</a></li> <li>• <a href="#">Understanding DCB Features and Requirements on page 6490</a></li> <li>• <a href="#">Configuring DCBX Autonegotiation on page 6524</a></li> <li>• <i>Understanding DCB Features and Requirements on EX Series Switches</i></li> <li>• <i>Disabling DCBX to Disable PFC Autonegotiation on EX Series Switches (CLI Procedure)</i></li> </ul>                              |

## dcbx-version

---

|                                 |                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>dcbx-version (auto-negotiate   ieee-dcbx   dcbx-version-1.01);</code>                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit <a href="#">protocols dcbx interface</a> (all   <i>interface-name</i> )]                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | <p>Set the DCBX version for the specified interface or interfaces.</p> <p>QFX3500 switches come up in IEEE DCBX mode and then autonegotiate with the connected peer to set the DCBX version.</p> <p>QFabric system Node devices come up using DCBX version 1.01, and then autonegotiate with the connected peer to set the DCBX mode.</p> |
| <b>Default</b>                  | The default DCBX mode is autonegotiation.                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><b>auto-negotiate</b>—Automatically negotiate the DCBX version with the connected peer.</p> <p><b>ieee-dcbx</b>—Force the interface to use IEEE DCBX mode, regardless of the peer configuration.</p> <p><b>dcbx-version-1.01</b>—Force the interface to use version 1.01 DCBX mode, regardless of the peer configuration.</p>          |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show dcbx neighbors on page 6581</a></li><li>• <a href="#">Configuring DCBX Autonegotiation on page 6524</a></li><li>• <a href="#">Understanding DCBX on page 6514</a></li></ul>                                                                                                      |

## destination-port (Applications)

|                                                                                                                                                                                                                                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                         | <code>destination-port <i>port-value</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                | [edit applications <b>application</b> <i>application-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>                                                                                                                                                                                                            | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>                                                                                                                                                                                                                    | Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) destination port number, which combines with <b>protocol</b> to identify an application type. The Internet Assigned Numbers Authority (IANA) assigns port numbers. See the IANA <i>Service Name and Transport Protocol Port Number Registry</i> at <a href="http://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xml">http://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xml</a> for a list of assigned port numbers. |
| <div>  <b>NOTE:</b> To create an application for iSCSI, use the protocol <code>tcp</code> with the destination port number <code>3260</code>. </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                                                                                                                                                                                                                        | <i>port-value</i> —Identifier for the port.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b>                                                                                                                                                                                                       | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>                                                                                                                                                                                                          | <ul style="list-style-type: none"> <li>• <a href="#">Defining an Application for DCBX Application Protocol TLV Exchange on page 6532</a></li> <li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 6535</a></li> <li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 6528</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li> </ul>               |

## disable (DCBX)

---


|                                 |                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | disable                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit <a href="#">protocols dcbx</a> ]<br><br>[edit <a href="#">protocols dcbx interface</a> <i>interface-name</i> ]                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 11.3 for EX Series switches.                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Disable Data Center Bridging Capability Exchange protocol (DCBX) on one or more 10-Gigabit Ethernet interfaces.                                                                                                                                                                                                                                                                            |
| <b>Default</b>                  | DCBX is enabled by default on all 10-Gigabit or higher Ethernet interfaces.<br><br>DCBX is enabled by default on all 10-Gigabit Ethernet interfaces on EX4500 CEE-enabled switches.                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring DCBX Autonegotiation on page 6524</a></li><li>• <i>Disabling DCBX to Disable PFC Autonegotiation on EX Series Switches (CLI Procedure)</i></li><li>• <a href="#">Understanding DCB Features and Requirements on page 6490</a></li><li>• <i>Understanding DCB Features and Requirements on EX Series Switches</i></li></ul> |

## enhanced-transmission-selection

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> enhanced-transmission-selection {   no-auto-negotiation;   no-recommendation-tlv;   recommendation-tlv {     no-auto-negotiation;   } } </pre>                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit <a href="#">protocols dcbx interface interface-name</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | <p>Disable advertising the enhanced transmission selection (ETS) state of the interface to the peer. To disable ETS on the interface, do not enable ETS on the interface in the class-of-service (CoS) configuration.</p> <p>Disabling ETS autonegotiation stops the QFX Series from advertising the ETS Configuration TLV and the ETS Recommendation TLV.</p> <p>Disabling the ETS recommendation TLV stops the QFX Series from advertising the ETS Recommendation TLV, but the ETS Configuration TLV is still advertised.</p> |
| <b>Options</b>                  | <p><b>no-auto-negotiation</b>—Disable automatic negotiation of ETS (Configuration TLV and Recommendation TLV)</p> <p><b>no-recommendation-tlv</b>—Disable automatic negotiation of the ETS Recommendation TLV</p> <p><b>recommendation-tlv</b>—Enable automatic negotiation of ETS Recommendation TLV</p>                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show dcbx neighbors on page 6581</a></li> <li>• <a href="#">Configuring DCBX Autonegotiation on page 6524</a></li> <li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 6771</a></li> <li>• <a href="#">Understanding DCB Features and Requirements on page 6490</a></li> </ul>                                                                                                                                                           |

## ether-type

---

|                                                                                                                                                                  |                                                                                                                                                                                                                                                                                                                                                  |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                    | <code>ether-type <i>ether-type</i>;</code>                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>                                                                                                                                           | [edit applications <a href="#">application</a> <i>application-name</i> ]                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>                                                                                                                                       | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                       |
| <b>Description</b>                                                                                                                                               | Two-octet field in an Ethernet frame that defines the protocol encapsulated in the frame payload. See <a href="http://standards.ieee.org/develop/regauth/ethertype/eth.txt">http://standards.ieee.org/develop/regauth/ethertype/eth.txt</a> for a list of Institute of Electrical and Electronics Engineers (IEEE) EtherTypes.                   |
| <div> <b>NOTE:</b> To create a FIP application, use the EtherType 0x8914.</div> |                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                                                                                                                                                   | <i>type</i> —Identifier for the EtherType.                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b>                                                                                                                                  | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                          |
| <b>Related Documentation</b>                                                                                                                                     | <ul style="list-style-type: none"><li>• <a href="#">Defining an Application for DCBX Application Protocol TLV Exchange on page 6532</a></li><li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 6535</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 6528</a></li></ul> |



## interface (DCBX)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> interface (<i>interface-name</i>   all) {   disable;   application-map <i>application-map-name</i>;   applications {     no-auto-negotiation;   }   enhanced-transmission-selection {     no-auto-negotiation;     no-recommendation-tlv;     recommendation-tlv {       no-auto-negotiation;     }   }   dcbx-version (auto-negotiate   ieee-dcbx   dcbx-version-1.01);   priority-flow-control {     no-auto-negotiation;   } } </pre>                                                                                                            |
| <b>Hierarchy Level</b>          | [edit <a href="#">protocols dcbx</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 11.3 for the EX Series switches.</p> <p><b>Mode</b> and <b>recommendation-tlv</b> statements introduced in Junos OS Release 12.2 for the QFX Series.</p>                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Configure DCBX properties on an interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <p><b><i>interface-name</i></b>—Name of the interface.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show dcbx neighbors on page 6581</a></li> <li>• <a href="#">Configuring DCBX Autonegotiation on page 6524</a></li> <li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li> <li>• <a href="#">Understanding DCB Features and Requirements on page 6490</a></li> <li>• <a href="#">Understanding DCB Features and Requirements on EX Series Switches</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li> </ul> |

## no-recommendation-tlv

---

|                                 |                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-recommendation-tlv;                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit <a href="#">protocols dcbx interface interface-name enhanced-transmission-selection</a> ]                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Disable DCBX to send the ETS Recommendation TLV (also known as the Information TLV) on egress. This feature is valid only if the interface DCBX mode is IEEE DCBX. If the interface DCBX mode is DCBX version 1.01, this statement has no effect. (DCBX version 1.01 does not advertise separate TLVs for individual attributes.) |
| <b>Default</b>                  | DCBX-enabled interfaces send the ETS recommendation TLV unless it is disabled.                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show dcbx neighbors on page 6581</a></li><li>• <a href="#">Configuring DCBX Autonegotiation on page 6524</a></li></ul>                                                                                                                                                        |

## policy-options

```
Syntax policy-options
 application-maps application-map-name {
 application application-name {
 code-points [aliases] [bit-patterns];
 }
 }
 policy-statement policy-name {
 term term-name {
 from {
 family family-name;
 match-conditions;
 policy subroutine-policy-name;
 prefix-list prefix-list-name;
 prefix-list-filter prefix-list-name match-type <actions>;
 route-filter destination-prefix match-type <actions>;
 source-address-filter source-prefix match-type <actions>;
 }
 to {
 match-conditions;
 policy subroutine-policy-name;
 }
 then actions;
 }
 }
 }
```

**Hierarchy Level** [edit]

**Release Information** Statement introduced in Junos OS Release 12.1 for the QFX Series.  
Statement introduced in Junos OS Release 12.1 for the EX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Configure options such as application maps for DCBX application protocol exchange and policy statements.

**Required Privilege Level** storage—To view this statement in the configuration.  
storage-control—To add this statement to the configuration.

**Related Documentation**

- [Defining an Application for DCBX Application Protocol TLV Exchange on page 6532](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 6535](#)
- [Example: Configuring DCBX to Support an iSCSI Application](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 6528](#)
- [Understanding DCBX Application Protocol TLV Exchange on EX Series Switches](#)

## priority-flow-control

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>priority-flow-control {<br/>    no-auto-negotiation;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit <a href="#">protocols dcbx interface</a> (all   <i>interface-name</i> )]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 11.3 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Disable autonegotiation of priority-based flow control (PFC) on one or more Ethernet interfaces. Autonegotiation enables PFC on an interface only if the switch and the peer device connected to the switch both support PFC and have the same PFC configuration. Disabling autonegotiation on an interface forces the interface to use the PFC state (enabled or disabled) that is configured on the switch by the configuration and assignment of the congestion notification profile.                                                                                                                                                                                                                             |
| <b>Options</b>                  | <b>no-auto-negotiation</b> —Disable automatic negotiation of PFC.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show dcbx neighbors on page 6581</a></li><li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 6905</a></li><li>• <i>Configuring Priority-Based Flow Control for an EX Series Switch (CLI Procedure)</i></li><li>• <a href="#">Configuring DCBX Autonegotiation on page 6524</a></li><li>• <a href="#">Example: Configuring CoS PFC for FCoE Traffic on page 6504</a></li><li>• <i>Understanding Data Center Bridging Capability Exchange Protocol for EX Series Switches</i></li><li>• <i>Understanding Priority-Based Flow Control</i></li><li>• <a href="#">Understanding DCB Features and Requirements on page 6490</a></li></ul> |

## protocol (Applications)

|                            |                                                                                                                                            |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>protocol (tcp   udp);</code>                                                                                                         |
| <b>Hierarchy Level</b>     | [edit applications <a href="#">application</a> <i>application-name</i> ]                                                                   |
| <b>Release Information</b> | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series. |
| <b>Description</b>         | Networking protocol type, which combines with <b>destination-port</b> to identify an application type.                                     |



**NOTE:** To create an application for iSCSI, use the protocol `tcp` with the destination port number `3260`.

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | <p><code>tcp</code>—Transmission Control Protocol</p> <p><code>udp</code>—User Datagram Protocol</p>                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | <p><code>interface</code>—To view this statement in the configuration.</p> <p><code>interface-control</code>—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Defining an Application for DCBX Application Protocol TLV Exchange on page 6532</a></li> <li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 6535</a></li> <li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 6528</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li> </ul> |

## recommendation-tlv

---

|                                 |                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>recommendation-tlv {<br/>    no-auto-negotiation;<br/>}</pre>                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit <a href="#">protocols dcbx interface interface-name enhanced-transmission-selection</a> ]                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Enable DCBX to send the ETS Recommendation TLV (also known as the Information TLV) on egress. This feature is valid only if the interface DCBX mode is IEEE DCBX. If the interface DCBX mode is DCBX version 1.01, this statement has no effect. (DCBX version 1.01 does not advertise separate TLVs for individual attributes.) |
| <b>Default</b>                  | DCBX-enabled interfaces send the ETS recommendation TLV unless it is disabled.                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <b>no-auto-negotiation</b> —Disable sending of the ETS recommendation TLV.                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show dcbx neighbors on page 6581</a></li><li>• <a href="#">Configuring DCBX Autonegotiation on page 6524</a></li></ul>                                                                                                                                                       |

## CHAPTER 250

# Operational Commands (FCoE on a Transit Switch)

- clear fip vlan-discovery statistics
- show fip vlan-discovery

## clear fip vlan-discovery statistics

---

|                                 |                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | clear fip vlan-discovery statistics                                                                    |
| <b>Release Information</b>      | Command introduced in Junos OS Release 12.1 for the QFX Series.                                        |
| <b>Description</b>              | Clear FIP VLAN discovery statistics.                                                                   |
| <b>Required Privilege Level</b> | view                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show fip vlan-discovery on page 6577</a></li></ul> |
| <b>List of Sample Output</b>    | <a href="#">clear fip vlan-discovery statistics on page 6576</a>                                       |

### Sample Output

#### clear fip vlan-discovery statistics

```
user@switch> clear fip vlan-discovery statistics
```



## show fip vlan-discovery

|                                 |                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show fip vlan-discovery (enodes   statistics)</b>                                                                                                                                                                          |
| <b>Release Information</b>      | Command introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                               |
| <b>Description</b>              | Display FCoE VLAN information from the Fibre Channel switch or FCoE forwarder (FCF).                                                                                                                                          |
| <b>Options</b>                  | <b>enodes</b> —Display VLAN discovery information for each ENode.<br><b>statistics</b> —Display VLAN discovery information statistics.                                                                                        |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">clear fip vlan-discovery statistics on page 6576</a></li> </ul>                                                                                                          |
| <b>List of Sample Output</b>    | <a href="#">show fip vlan-discovery enodes on page 6578</a><br><a href="#">show fip vlan-discovery statistics (QFX3500) on page 6578</a><br><a href="#">show fip vlan-discovery statistics (QFabric Systems) on page 6578</a> |
| <b>Output Fields</b>            | Table 522 lists the output fields for the <b>show fip vlan-discovery</b> command. Output fields are listed in the approximate order in which they appear.                                                                     |

Table 522: show fip vlan-discovery Output Fields

| Field Name                     | Field Description                                                                                                             | Level of Output |
|--------------------------------|-------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Enode-MAC                      | Media access control (MAC) address of the ENode.                                                                              | enodes          |
| Interface                      | Name of the interface.                                                                                                        | enodes          |
| Unsolicited notification count | Number of unsolicited VLAN discovery notifications.                                                                           | All             |
| Solicited notification count   | Number of solicited VLAN discovery notifications.                                                                             | statistics      |
| Node Group Name                | Displays the name of the Node group on QFabric systems.                                                                       | statistics      |
| Request count                  | Number of VLAN discovery requests sent by the ENode. This number should match the <b>Solicited notification count</b> number. | statistics      |
| VLAN tags                      | Tags of the FIP-enabled VLANs.                                                                                                | enodes          |

## Sample Output

### show fip vlan-discovery enodes

```
user@switch> show fip vlan-discovery enodes
```

| Enode-MAC         | Interface  | Unsolicited<br>Notification<br>Count | Vlan Tags |
|-------------------|------------|--------------------------------------|-----------|
| 00:10:94:00:00:02 | xe-0/0/9.0 | 0                                    | 400       |

### show fip vlan-discovery statistics (QFX3500)

```
user@switch> show fip vlan-discovery statistics
```

```
Request count: 0
Solicited notification count: 0
Unsolicited notification count: 1
```

### show fip vlan-discovery statistics (QFabric Systems)

```
user@switch> show fip vlan-discovery statistics
```

```
NW-NG-0:
```

```

Request count: 0
Solicited notification count: 0
Unsolicited notification count: 1
```

```
BBAK0399:
```

```

Request count: 0
Solicited notification count: 0
Unsolicited notification count: 1
```

```
FCC001:
```

```

Request count: 0
Solicited notification count: 0
Unsolicited notification count: 1
```

## CHAPTER 251

# Operational Commands (DCBX and PFC)

- `show dcbx`
- `show dcbx neighbors`

## show dcbx

|                                 |                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show dcbx                                                                                                                                                                     |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                               |
| <b>Description</b>              | List DCBX status (enabled or disabled) and the interfaces on which DCBX is enabled.                                                                                           |
| <b>Required Privilege Level</b> | view                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show dcbx neighbors on page 6581</a></li> <li>• <a href="#">Configuring DCBX Autonegotiation on page 6524</a></li> </ul> |
| <b>Output Fields</b>            | <a href="#">Table 523</a> lists the output fields for the <b>show dcbx</b> command. Output fields are listed in the approximate order in which they appear.                   |

**Table 523: show dcbx output fields**

| Field Name | Field Description                                                                                                                                                                                                                                                     |
|------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DCBX       | Status of DCBX on the switch or for the specified interface: <ul style="list-style-type: none"> <li>• Enabled—DCBX is enabled on the switch or on the specified interface</li> <li>• Disabled—DCBX is disabled on the switch or on the specified interface</li> </ul> |
| Interface  | Name of the interface                                                                                                                                                                                                                                                 |

## Sample Output

### show dcbx

```

user@switch> show dcbx
DCBX : Enabled
Interface DCBX
xe-0/0/9.0 enabled
xe-0/0/32.0 enabled
xe-0/0/36.0 enabled

```

## show dcbx neighbors

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show dcbx neighbors</b><br><b>&lt;interface interface-name&gt;</b><br><b>&lt;terse&gt;</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.<br>Command introduced in Junos OS Release 11.3 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Display information about Data Center Bridging Capability Exchange protocol (DCBX) neighbor interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <b>none</b> —Display information about all DCBX neighbor interfaces.<br><br><b>interface-name</b> —(Optional) Display information for the specified interface.<br><br><b>terse</b> —Display the specified level of output.                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring DCBX Autonegotiation on page 6524</a></li> <li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 6535</a></li> <li>• <a href="#">Example: Configuring an FCoE Transit Switch</a></li> <li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li> <li>• <a href="#">Understanding DCB Features and Requirements on page 6490</a></li> <li>• <a href="#">Understanding Data Center Bridging Capability Exchange Protocol for EX Series Switches</a></li> <li>• <a href="#">dcbx on page 6563</a></li> </ul>                   |
| <b>List of Sample Output</b>    | <a href="#">show dcbx neighbors interface (QFX Series, DCBX Version 1.01 Mode) on page 6594</a><br><a href="#">show dcbx neighbors interface (QFX Series, IEEE DCBX Mode) on page 6596</a><br><a href="#">show dcbx neighbors terse (QFX Series) on page 6598</a><br><a href="#">show dcbx neighbors (EX4500 Switch: FCoE Interfaces on Both Local and Peer with PFC Configured Compatibly) on page 6598</a><br><a href="#">show dcbx neighbors (EX4500 Switch: DCBX Interfaces on Local and Peer Are Configured Compatibly with iSCSI Application) on page 6599</a><br><a href="#">show dcbx neighbors (EX4500 Switch: Includes ETS) on page 6600</a> |
| <b>Output Fields</b>            | <a href="#">Table 524</a> lists the output fields for the <b>show dcbx neighbors</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

**Table 524: show dcbx neighbors Output Fields**

| Field Name | Field Description      |
|------------|------------------------|
| Interface  | Name of the interface. |

Table 524: show dcbx neighbors Output Fields (*continued*)

| Field Name             | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Parent Interface       | Name of the link aggregation group (LAG) interface to which the DCBX interface belongs.                                                                                                                                                                                                                                                                                                                                                                                                                |
| Active-application-map | Name of the application map applied to the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Protocol-Mode          | <p>(QFX Series) DCBX protocol mode the interface uses:</p> <ul style="list-style-type: none"> <li>IEEE DCBX Version—The interface uses IEEE DCBX mode.</li> <li>DCBX Version 1.01—The interface uses DCBX version 1.01.</li> </ul> <p><b>NOTE:</b> On interfaces that use the IEEE DCBX mode, the <b>show dcbx neighbors interface <i>interface-name</i></b> operational command does not include application, PFC, or ETS operational state in the output.</p>                                        |
| Protocol-State         | <p>(DCBX Version 1.01 only) DCBX protocol state synchronization status:</p> <ul style="list-style-type: none"> <li><b>in-sync</b>—The local interface received an acknowledge message from the peer to indicate that the peer received a state change message sent by the local interface.</li> <li><b>ack-pending</b>—The local interface has not yet received an acknowledge message from the peer to indicate that the peer received a state change message sent by the local interface.</li> </ul> |
| Local-Advertisement    | <p>(DCBX Version 1.01 only)</p> <p>Status of advertisements that the local interface sends to the peer.</p>                                                                                                                                                                                                                                                                                                                                                                                            |
| Operational version    | Version of the DCBX standard used.                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| sequence-number        | <p>Number of state change messages sent to the peer.</p> <p>If the interface <b>Protocol-State</b> value is <b>in-sync</b>, this number should match the <b>acknowledge-id</b> number in the <b>Peer-Advertisement</b> section.</p> <p>If the interface <b>Protocol-State</b> value is <b>ack-pending</b>, this number does not match the <b>acknowledge-id</b> number in the <b>Peer-Advertisement</b> section.</p>                                                                                   |
| acknowledge-id         | <p>Number of acknowledge messages received from the peer.</p> <p>If the <b>Protocol-State</b> value is <b>in-sync</b>, this number should match the <b>sequence-number</b> value in the <b>Peer-Advertisement</b> section.</p> <p>If the <b>Protocol-State</b> value is <b>ack-pending</b>, this number does not match the <b>sequence-number</b> value in the <b>Peer-Advertisement</b> section.</p>                                                                                                  |

Table 524: show dcbx neighbors Output Fields (*continued*)

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Peer-Advertisement</b>  | (DCBX Version 1.01 only)<br><br>Status of advertisements that the peer sends to the local interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Operational version</b> | Version of the DCBX standard used.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>sequence-number</b>     | <p>Number of state change messages the peer sent to the local interface.</p> <p>If this number matches the <b>acknowledge-id</b> number in the <b>Local-Advertisement</b> field, this indicates that the local interface has acknowledged all of the peer's state change messages and is synchronized.</p> <p>If this number does not match the <b>acknowledge-id</b> number in the <b>Local-Advertisement</b> field, this indicates that the peer has not yet received an acknowledgment for a state change message from the local interface.</p>            |
| <b>acknowledge-id</b>      | <p>Number of acknowledge messages the peer has received from the local interface.</p> <p>If this number matches the <b>sequence-number</b> value in the <b>Local-Advertisement</b> field, this indicates that the peer has acknowledged all of the local interface's state change messages and is in synchronization.</p> <p>If this number does not match the <b>sequence-number</b> value in the <b>Local-Advertisement</b> field, this indicates that the peer has not yet sent an acknowledgment for a state change message from the local interface.</p> |

Table 524: show dcbx neighbors Output Fields (*continued*)

| Field Name                        | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Feature: PFC</b>               | Priority-based flow control (PFC) feature DCBX state information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Protocol-State</b>             | (DCBX Version 1.01 only)<br><br>DCBX protocol state synchronization status: <ul style="list-style-type: none"> <li>• <b>ack-pending</b>—The local interface has not yet received an acknowledge message from the peer to indicate that the peer received a PFC state change message sent by the local interface.</li> <li>• <b>in-sync</b>—The local interface received an acknowledge message from the peer to indicate that the peer received a PFC state change message sent by the local interface.</li> <li>• <b>not-applicable</b>—PFC autonegotiation is disabled.</li> </ul> |
| <b>Operational State</b>          | (DCBX Version 1.01 only)<br><br>Operational state of the feature: <b>enabled</b> or <b>disabled</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Local-Advertisement</b>        | Status of advertisements that the local interface sends to the peer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Enable</b>                     | (DCBX Version 1.01 only)<br><br>State that the local interface advertises to the peer: <ul style="list-style-type: none"> <li>• <b>Yes</b>—The feature is enabled.</li> <li>• <b>No</b>—The feature is disabled.</li> </ul>                                                                                                                                                                                                                                                                                                                                                          |
| <b>Willing</b>                    | Willingness of the local interface to learn the PFC configuration from the peer using DCBX: <ul style="list-style-type: none"> <li>• <b>Yes</b>—The local interface is willing to learn the PFC configuration from the peer.</li> <li>• <b>No</b>—The local interface is not willing to learn the PFC configuration from the peer.</li> </ul>                                                                                                                                                                                                                                        |
| <b>Mac auth Bypass Capability</b> | (IEEE DCBX only)<br><br>(QFX Series) Media access controller (MAC) authentication bypass provides access to devices based on MAC address authentication. This is not supported, so the only value seen in the local advertisement field is <b>no</b> .                                                                                                                                                                                                                                                                                                                               |
| <b>Error</b>                      | (DCBX Version 1.01 only)<br><br>Configuration compatibility error status: <ul style="list-style-type: none"> <li>• <b>No</b>—No error detected. Local and peer configuration are compatible.</li> <li>• <b>Yes</b>—Error detected. Local and peer configuration are not compatible.</li> </ul>                                                                                                                                                                                                                                                                                       |



Table 524: show dcbx neighbors Output Fields (*continued*)

| Field Name                                            | Field Description                                                                                                                                                                                                                                                                                                                                    |
|-------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Operational State</b>                              | <p>PFC operational state on the interface:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—PFC is enabled on the interface</li> <li>• <b>Disabled</b>—PFC is disabled on the interface</li> </ul>                                                                                                                                        |
| <b>Maximum Traffic Classes capable to support PFC</b> | <p>Largest number of traffic classes the local interface supports for PFC:</p> <ul style="list-style-type: none"> <li>• <b>6</b> (EX Series switches)</li> <li>• <b>6</b> (QFX Series)</li> </ul>                                                                                                                                                    |
| <b>Code Point</b>                                     | <p>PFC code point, which is specified in the 3-bit class-of-service field in the VLAN header.</p>                                                                                                                                                                                                                                                    |
| <b>Admin Mode</b>                                     | <p>PFC administrative state for each code point on the local interface:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—PFC is enabled for the code point.</li> <li>• <b>Disabled</b>—PFC is disabled for the code point.</li> </ul>                                                                                                     |
| <b>Operational Mode</b>                               | <p>(QFX Series) PFC operational mode for each code point:</p> <ul style="list-style-type: none"> <li>• <b>Enable</b>—PFC is enabled on the code point.</li> <li>• <b>Disable</b>—PFC is disabled on the code point.</li> </ul>                                                                                                                       |
| <b>Peer-Advertisement</b>                             | <p>Status of advertisements that the peer sends to the local interface.</p>                                                                                                                                                                                                                                                                          |
| <b>Enable</b>                                         | <p>(DCBX Version 1.01 only)</p> <p>State that the peer advertises to the local interface:</p> <ul style="list-style-type: none"> <li>• <b>Yes</b>—The feature is enabled.</li> <li>• <b>No</b>—The feature is disabled.</li> </ul>                                                                                                                   |
| <b>Willing</b>                                        | <p>Willingness of the peer to learn the PFC configuration from the local interface using DCBX:</p> <ul style="list-style-type: none"> <li>• <b>Yes</b>—The peer is willing to learn the PFC configuration from the local interface.</li> <li>• <b>No</b>—The peer is not willing to learn the PFC configuration from the local interface.</li> </ul> |
| <b>Error</b>                                          | <p>(DCBX Version 1.01 only)</p> <p>Configuration compatibility error status:</p> <ul style="list-style-type: none"> <li>• <b>No</b>—No error detected. Local and peer configuration are compatible.</li> <li>• <b>Yes</b>—Error detected. Local and peer configuration are not compatible.</li> </ul>                                                |

Table 524: show dcbx neighbors Output Fields (*continued*)

| Field Name                                            | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Operational State</b>                              | <p>PFC operational state on the interface:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—PFC is enabled on the interface</li> <li>• <b>Disabled</b>—PFC is disabled on the interface</li> </ul>                                                                                                                                                                                                                                                                                            |
| <b>Mac auth Bypass Capability</b>                     | <p>(IEEE DCBX only)</p> <p>(QFX Series) Media access controller (MAC) authentication bypass provides access to devices based on MAC address authentication. Although the QFX Series does not support this feature, the connected peer might support it. This field reports the peer state:</p> <ul style="list-style-type: none"> <li>• <b>Yes</b>—The connected peer supports MAC authentication bypass.</li> <li>• <b>No</b>—The connected peer does not support MAC authentication bypass.</li> </ul> |
| <b>Maximum Traffic Classes capable to support PFC</b> | <p>Largest number of traffic classes the peer supports for PFC:</p> <ul style="list-style-type: none"> <li>• <b>6</b> (EX Series switches)</li> <li>• <b>8</b> (QFX Series)</li> </ul>                                                                                                                                                                                                                                                                                                                   |
| <b>Code Point</b>                                     | <p>PFC code point, which is specified in the 3-bit class-of-service field in the VLAN header.</p>                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Admin Mode</b>                                     | <p>PFC administrative state for each code point on the peer:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—PFC is enabled for the code point.</li> <li>• <b>Disabled</b>—PFC is disabled for the code point.</li> </ul>                                                                                                                                                                                                                                                                    |

Table 524: show dcbx neighbors Output Fields (*continued*)

| Field Name                  | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Feature: Application</b> | State information for the DCBX application.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Protocol-State</b>       | <p>(DCBX Version 1.01 only)</p> <p>DCBX protocol state synchronization status:</p> <ul style="list-style-type: none"> <li>• <b>in-sync</b>—The local interface received an acknowledge message from the peer to indicate that the peer received an FCoE state change message sent by the local interface.</li> <li>• <b>ack-pending</b>—The local interface has not yet received an acknowledge message from the peer to indicate that the peer received an FCoE state change message sent by the local interface.</li> <li>• <b>not-applicable</b>—The local interface is set to <b>no-auto-negotiation</b> (autonegotiation is disabled). If the interface is associated with an FCoE forwarding class, the interface advertises FCoE capability even if the connected peer does not advertise FCoE capability.</li> </ul> |
| <b>Local-Advertisement</b>  | <p>Status of advertisements that the local interface sends to the peer.</p> <p>If the local interface is set to <b>no-auto-negotiation</b> (autonegotiation is disabled), the local advertisement portion of the output is not shown.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Enable</b>               | <p>(DCBX Version 1.01 only)</p> <p>State that the local interface advertises to the peer:</p> <ul style="list-style-type: none"> <li>• <b>Yes</b>—The feature is enabled.</li> <li>• <b>No</b>—The feature is disabled.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Willing</b>              | <p>(DCBX Version 1.01 only)</p> <p>Willingness of the local interface to learn the FCoE interface state from the peer using DCBX:</p> <ul style="list-style-type: none"> <li>• <b>Yes</b>—The local interface is willing to learn the FCoE interface state from the peer.</li> <li>• <b>No</b>—The local interface is not willing to learn the FCoE interface state from the peer.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Error</b>                | <p>(DCBX Version 1.01 only)</p> <p>Configuration compatibility error status:</p> <ul style="list-style-type: none"> <li>• <b>No</b>—No error detected. The local and peer configuration are compatible.</li> <li>• <b>Yes</b>—Error detected. The local and peer configuration are not compatible.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Appl-Name</b>            | Name of the application:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

Table 524: show dcbx neighbors Output Fields (*continued*)

| Field Name                            | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Ethernet-Type</b>                  | <p>(DCBX Version 1.01 only)</p> <p>Ethernet type (EtherType) of the application. For example, <b>0x8906</b> indicates the EtherType for the FCoE application. Either the EtherType (for Layer 2 applications) or the Socket Number (for Layer 4 applications) of the application is displayed in the output.</p>                                                                                                                                                                                                                                                            |
| <b>Socket-Number</b>                  | <p>Destination port socket number of the application, if applicable. Either the EtherType (for Layer 2 applications) or the Socket Number (for Layer 4 applications) of the application is displayed in the output.</p>                                                                                                                                                                                                                                                                                                                                                     |
| <b>Priority-Field or Priority-Map</b> | <p>Priority assigned to the application.</p> <p>For EX Series switches, the priority of the FCoE application is determined by the PFC congestion notification profile that has been configured and associated with the FCoE interface. For other applications, the priority is based on the application map.</p>                                                                                                                                                                                                                                                            |
| <b>Status</b>                         | <p>(DCBX Version 1.01 only)</p> <p>Local status when autonegotiation is enabled:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—The application feature is enabled on both the local interface and the peer interface. (The local configuration and the peer configuration match.)</li> <li>• <b>Disabled</b>—The local configuration and the peer configuration do not match.</li> </ul> <p><b>NOTE:</b> If there is a configuration mismatch in one application between the switch and the peer, all the other applications including FCoE are disabled.</p> |
| <b>Peer-Advertisement</b>             | <p>Status of advertisements that the peer sends to the local interface.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Enable</b>                         | <p>(DCBX Version 1.01 only)</p> <p>State that the peer advertises to the local interface:</p> <ul style="list-style-type: none"> <li>• <b>Yes</b>—The feature is enabled.</li> <li>• <b>No</b>—The feature is disabled.</li> </ul>                                                                                                                                                                                                                                                                                                                                          |
| <b>Willing</b>                        | <p>(DCBX Version 1.01 only)</p> <p>Willingness of the peer to learn the FCoE interface state from the local interface using DCBX:</p> <ul style="list-style-type: none"> <li>• <b>Yes</b>—The peer is willing to learn the FCoE interface state from the local interface.</li> <li>• <b>No</b>—The peer is not willing to learn the FCoE interface state from the local interface.</li> </ul>                                                                                                                                                                               |

Table 524: show dcbx neighbors Output Fields (*continued*)

| Field Name                            | Field Description                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Error</b>                          | (DCBX Version 1.01 only)<br><br>Configuration compatibility error status: <ul style="list-style-type: none"> <li>• <b>No</b>—No error detected. Local and peer configuration are compatible.</li> <li>• <b>Yes</b>—Error detected. Local and peer configuration are not compatible.</li> </ul>                                                                           |
| <b>Appl-Name</b>                      | Name of the application: <ul style="list-style-type: none"> <li>• <b>FCoE</b>—Fibre Channel over Ethernet</li> </ul>                                                                                                                                                                                                                                                     |
| <b>Ethernet-Type</b>                  | Ethernet type (EtherType) of the application. For example, <b>0x8906</b> indicates the EtherType for the FCoE application. Either the EtherType (for Layer 2 applications) or the Socket-Number (for Layer 4 applications) of the application is displayed in the output.                                                                                                |
| <b>Socket-Number</b>                  | Destination port socket number of the application, if applicable. Either the EtherType (for Layer 2 applications) or the Socket Number (for Layer 4 applications) of the application is displayed in the output.                                                                                                                                                         |
| <b>Priority-Field or Priority-Map</b> | Priority assigned to the application.                                                                                                                                                                                                                                                                                                                                    |
| <b>Status</b>                         | (DCBX Version 1.01 only)<br><br>Peer interface status: <ul style="list-style-type: none"> <li>• <b>Enabled</b>—The application feature is enabled on both the local interface and the peer interface. (The local configuration and the peer configuration match.)</li> <li>• <b>Disabled</b>—The local configuration and the peer configuration do not match.</li> </ul> |

Table 524: show dcbx neighbors Output Fields (*continued*)

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Feature: ETS</b>        | Enhanced Transmission Selection (ETS) DCBX state information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Protocol-State</b>      | (DCBX Version 1.01 only)<br><br>ETS protocol state synchronization status: <ul style="list-style-type: none"> <li>• <b>in-sync</b>—The local interface received an acknowledge message from the peer to indicate that the peer received an ETS state change message sent by the local interface.</li> <li>• <b>ack-pending</b>—The local interface has not yet received an acknowledge message from the peer to indicate that the peer received an ETS state change message sent by the local interface.</li> </ul>                                                          |
| <b>Operational State</b>   | (DCBX Version 1.01 only)<br><br>Operational state of the feature, <b>enabled</b> or <b>disabled</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Local-Advertisement</b> | Status of advertisements that the local interface sends to the peer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Enable</b>              | (DCBX Version 1.01 only)<br><br>State that the local interface advertises to the peer: <ul style="list-style-type: none"> <li>• <b>Yes</b>—The feature is enabled.</li> <li>• <b>No</b>—The feature is disabled.</li> </ul>                                                                                                                                                                                                                                                                                                                                                  |
| <b>TLV Type</b>            | (IEEE DCBX only)<br><br>Type of ETS TLV: <ul style="list-style-type: none"> <li>• <b>Configuration</b>—Advertises the Configuration TLV, which communicates the local ETS configuration to the peer but does not ask the peer to use the configuration.</li> <li>• <b>Recommendation</b>—Advertises the Recommendation TLV, which communicates the local ETS configuration to the peer, and if the peer is “willing,” configures the peer interface to match the local ETS configuration.</li> <li>• <b>Recommendation-or-Configuration</b>—Advertises both TLVs.</li> </ul> |
| <b>Willing</b>             | Willingness of the local interface to learn the ETS state from the peer using DCBX (EX Series switches always advertise <b>No</b> for this field): <ul style="list-style-type: none"> <li>• <b>Yes</b>—Local interface is willing to learn the ETS state from the peer.</li> <li>• <b>No</b>—Local interface is not willing to learn the ETS state from the peer.</li> </ul>                                                                                                                                                                                                 |
| <b>Credit Based Shaper</b> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

Table 524: show dcbx neighbors Output Fields (*continued*)

| Field Name                                            | Field Description                                                                                                                                                                                                                                                                                               |
|-------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                                       | (IEEE DCBX only)<br><br>Alternative method of flow control to buffer-to-buffer credit. The QFX Series does not support a credit-based shaper, so the value of this field is always <b>No</b> .                                                                                                                  |
| <b>Error</b>                                          | (DCBX Version 1.01 only)<br><br>Configuration error status: <ul style="list-style-type: none"> <li>• <b>No</b>—No error. This should always be the switch ETS error state.</li> <li>• <b>Yes</b>—Error detected.</li> </ul>                                                                                     |
| <b>Maximum Traffic Classes capable to support PFC</b> | (DCBX Version 1.01 only)<br><br>Largest number of traffic classes the local interface supports for PFC.                                                                                                                                                                                                         |
| <b>Maximum Traffic Classes supported</b>              | (IEEE DCBX only)<br><br>Largest number of traffic classes the local interface supports for ETS. (EX Series switches support only one traffic class for ETS. However, a different value might be shown for this field.)                                                                                          |
| <b>Code Point</b>                                     | PFC code point, which is specified in the 3-bit class-of-service field in the VLAN header.                                                                                                                                                                                                                      |
| <b>Priority-Group</b>                                 | Class-of-service (CoS) priority group (forwarding class set) identification number.                                                                                                                                                                                                                             |
| <b>Percentage B/W</b>                                 | Configured minimum percentage of link bandwidth allocated to the priority group. Only explicitly configured values appear in this output column. If the link bandwidth is the default percentage, it is not shown. (EX Series switches allocate 100% of link bandwidth to the default priority group, group 7.) |
| <b>Transmission Selection Algorithm</b>               | (IEEE DCBX only)<br><br>The transmission selection algorithm used by the interface. The QFX Series supports ETS but does not support using the credit-based shaper algorithm, so the only value shown in this field is <b>ETS</b> .                                                                             |
| <b>Peer-Advertisement</b>                             | Status of advertisements that the peer sends to the local interface.                                                                                                                                                                                                                                            |
| <b>Enable</b>                                         |                                                                                                                                                                                                                                                                                                                 |

Table 524: show dcbx neighbors Output Fields (*continued*)

| Field Name                                            | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                                       | (DCBX Version 1.01 only)<br><br>State that the peer advertises to the local interface:<br><br><ul style="list-style-type: none"> <li>• <b>Yes</b>—The feature is enabled.</li> <li>• <b>No</b>—The feature is disabled.</li> </ul>                                                                                                                                                                                                                                                                                                                                               |
| <b>TLV Type</b>                                       | (IEEE DCBX only)<br><br>Type of ETS TLV:<br><br><ul style="list-style-type: none"> <li>• <b>Configuration</b>—Advertises the Configuration TLV, which communicates the local ETS configuration to the peer but does not ask the peer to use the configuration.</li> <li>• <b>Recommendation</b>—Advertises the Recommendation TLV, which communicates the local ETS configuration to the peer, and if the peer is “willing,” configures the peer interface to match the local ETS configuration.</li> <li>• <b>Configuration/Recommendation</b>—Advertises both TLVs.</li> </ul> |
| <b>Willing</b>                                        | Willingness of the peer to learn the ETS state from the local interface using DCBX:<br><br><ul style="list-style-type: none"> <li>• <b>Yes</b>—Peer is willing to learn the ETS state from the local interface.</li> <li>• <b>No</b>—Peer is not willing to learn the ETS state from the local interface.</li> </ul>                                                                                                                                                                                                                                                             |
| <b>Credit Based Shaper</b>                            | (IEEE DCBX only)<br><br>Alternative method of flow control to buffer-to-buffer credit. The QFX Series does not support a credit-based shaper, so the value of this field is always <b>No</b> .                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Error</b>                                          | (DCBX Version 1.01 only)<br><br>Configuration error status of the peer:<br><br><ul style="list-style-type: none"> <li>• <b>No</b>—No error in peer ETS TLV.</li> <li>• <b>Yes</b>—Error in peer ETS TLV.</li> </ul>                                                                                                                                                                                                                                                                                                                                                              |
| <b>Maximum Traffic Classes capable to support PFC</b> | (DCBX Version 1.01 only)<br><br>Largest number of traffic classes the local interface supports for PFC.                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Maximum Traffic Classes supported</b>              | (IEEE DCBX only)<br><br>Largest number of traffic classes the local interface supports for ETS. (EX Series switches support only one traffic class for ETS. However, a different value might be shown for this field.)                                                                                                                                                                                                                                                                                                                                                           |
| <b>Code Point</b>                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |



Table 524: show dcbx neighbors Output Fields (*continued*)

| Field Name                              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                  |
|-----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                         | PFC code point, which is specified in the 3-bit class-of-service field in the VLAN header.                                                                                                                                                                                                                                                                                                                         |
| <b>Priority-Group</b>                   | CoS priority group (forwarding class set) identification number.                                                                                                                                                                                                                                                                                                                                                   |
| <b>Percentage B/W</b>                   | Configured minimum percentage of link bandwidth allocated to the priority group. (EX Series switches allocate 100% of link bandwidth to the default priority group, group 7.)                                                                                                                                                                                                                                      |
| <b>Transmission Selection Algorithm</b> | (IEEE DCBX only)<br><br>Transmission selection algorithm used by the interface. The QFX Series supports ETS but does not support using the credit-based shaper algorithm, so the only value shown in this field is <b>ETS</b> .                                                                                                                                                                                    |
| <b>PFC</b>                              | (QFX Series, <b>terse</b> option only) DCBX TLV advertisement state for PFC: <ul style="list-style-type: none"> <li>• Disabled—PFC configuration matches the configuration on the connected peer and PFC is disabled</li> <li>• Enabled—PFC configuration matches the configuration on the connected peer and PFC is enabled</li> <li>• Not Advt—Interface does not advertise PFC to the connected peer</li> </ul> |
| <b>ETS</b>                              | ( <b>terse</b> option only) Local DCBX TLV advertisement state for ETS: <ul style="list-style-type: none"> <li>• Advt—Interface advertises ETS TLVs</li> <li>• Disabled—ETS is disabled on the interface (interface does not advertise ETS)</li> </ul>                                                                                                                                                             |
| <b>ETS Rec</b>                          | ( <b>terse</b> option only) DCBX TLV peer advertisement state for ETS (state received from the connected DCBX peer): <ul style="list-style-type: none"> <li>• Advt—Peer interface advertises ETS TLVs</li> <li>• Not Advt—Peer interface does not advertise ETS</li> </ul> <p><b>NOTE:</b> When the DCBX mode is DCBX version 1.01, no peer information is displayed.</p>                                          |

Table 524: show dcbx neighbors Output Fields (*continued*)

| Field Name | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Version    | <p>(<b>terse</b> option only) The DCBX version used on the interface and whether the DCBX version was autonegotiated or explicitly configured:</p> <ul style="list-style-type: none"> <li>• <b>IEEE</b>—The interface uses IEEE DCBX.</li> <li>• <b>1.01</b>—The interface uses DCBX version 1.01.</li> </ul> <p>When the DCBX version used is the result of autonegotiation, the term (<b>Auto</b>) appears next to the version. For example, <b>IEEE (Auto)</b> indicates that the interface autonegotiated with the connected peer to use IEEE DCBX. Autonegotiation is enabled by default.</p> |

## Sample Output

### show dcbx neighbors interface (QFX Series, DCBX Version 1.01 Mode)

```

user@switch> show dcbx neighbors interface xe-0/0/0
Interface : xe-0/0/0.0 - Parent Interface: ae0.0
Active-application-map: app-map-1
Protocol-State: in-sync
Protocol-Mode: DCBX Version 1.01

Local-Advertisement:
 Operational version: 1
 sequence-number: 130, acknowledge-id: 102

Peer-Advertisement:
 Operational version: 1
 sequence-number: 102, acknowledge-id: 130

Feature: PFC, Protocol-State: in-sync

Operational State: Enabled

Local-Advertisement:
 Enable: Yes, Willing: No, Error: No
 Maximum Traffic Classes capable to support PFC: 8

Code Point Admin Mode Operational Mode
000 Disabled Disable
001 Disabled Disable
010 Disabled Disable
011 Enabled Enable
100 Enabled Enable
101 Disabled Disable
110 Disabled Disable
111 Disabled Disable

Peer-Advertisement:
 Enable: Yes, Willing: No, Error: No
 Maximum Traffic Classes capable to support PFC: 8

Code Point Admin Mode
000 Disabled

```

|     |          |
|-----|----------|
| 001 | Disabled |
| 010 | Disabled |
| 011 | Enabled  |
| 100 | Enabled  |
| 101 | Disabled |
| 110 | Disabled |
| 111 | Disabled |

Feature: Application, Protocol-State: in-sync

Local-Advertisement:

Enable: Yes, Willing: No, Error: No

| App1-Name | Ethernet-Type | Socket-Number | Priority-Map | Status  |
|-----------|---------------|---------------|--------------|---------|
| FCoE      | 0x8906        |               | 00001110     | Enabled |
| iSCSI     |               | 3260          | 10000000     | Enabled |

Peer-Advertisement:

Enable: Yes, Willing: Yes, Error: No

| App1-Name | Ethernet-Type | Socket-Number | Priority-Map | Status  |
|-----------|---------------|---------------|--------------|---------|
| FCoE      | 0x8906        | N/A           | 00001110     | Enabled |

Feature: ETS, Protocol-State: in-sync

Operational State: Enabled

Local-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 8

| Code Point | Priority-Group |
|------------|----------------|
| 000        | 0              |
| 001        | 7              |
| 010        | 7              |
| 011        | 7              |
| 100        | 0              |
| 101        | 1              |
| 110        | 1              |
| 111        | 7              |

| Priority-Group | Percentage B/W |
|----------------|----------------|
| 0              | 40%            |
| 1              | 5%             |

Peer-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 8

| Code Point | Priority-Group |
|------------|----------------|
| 000        | 0              |
| 001        | 7              |
| 010        | 7              |
| 011        | 7              |
| 100        | 0              |
| 101        | 1              |
| 110        | 1              |

|                |                |
|----------------|----------------|
| 111            | 7              |
| Priority-Group | Percentage B/W |
| 0              | 40%            |
| 1              | 5%             |

**show dcbx neighbors interface (QFX Series, IEEE DCBX Mode)**

```
user@switch> show dcbx neighbors interface xe-0/0/0
```

```
Interface : xe-0/0/0.0 - Parent Interface: ae0.0
```

```
Active-application-map: app-map-1
```

```
Protocol-Mode: IEEE-DCBX Version
```

```
Feature: PFC
```

```
Local-Advertisement:
```

```
Willing: No
```

```
Mac auth Bypass Capability: No
```

```
Operational State: Enabled
```

```
Maximum Traffic Classes capable to support PFC: 8
```

| Code Point | Admin Mode |
|------------|------------|
| 000        | Disabled   |
| 001        | Disabled   |
| 010        | Disabled   |
| 011        | Enabled    |
| 100        | Enabled    |
| 101        | Disabled   |
| 110        | Disabled   |
| 111        | Disabled   |

```
Peer-Advertisement:
```

```
Willing: No
```

```
Mac auth Bypass Capability: No
```

```
Operational State: Enabled
```

```
Maximum Traffic Classes capable to support PFC: 8
```

| Code Point | Admin Mode |
|------------|------------|
| 000        | Disabled   |
| 001        | Disabled   |
| 010        | Disabled   |
| 011        | Enabled    |
| 100        | Enabled    |
| 101        | Disabled   |
| 110        | Disabled   |
| 111        | Disabled   |

```
Feature: Application
```

```
Local-Advertisement:
```

| App1-Name | Ethernet-Type | Socket-Number | Priority-field |
|-----------|---------------|---------------|----------------|
| FCoE      | 0x8906        |               | 00001110       |
| iSCSI     |               | 3260          | 10000000       |

```
Peer-Advertisement:
```

| App1-Name | Ethernet-Type | Socket-Number | Priority-field |
|-----------|---------------|---------------|----------------|
|-----------|---------------|---------------|----------------|

|      |        |     |          |
|------|--------|-----|----------|
| FCoE | 0x8906 | N/A | 00001110 |
|------|--------|-----|----------|

Feature: ETS

Local-Advertisement:

TLV Type: Configuration/Recommendation

Willing: No

Credit Based Shaper: No

Maximum Traffic Classes supported: 3

| Code Point | Priority-Group |
|------------|----------------|
| 000        | 0              |
| 001        | 7              |
| 010        | 7              |
| 011        | 7              |
| 100        | 0              |
| 101        | 1              |
| 110        | 1              |
| 111        | 7              |

| Priority-Group | Percentage B/W |
|----------------|----------------|
| 0              | 40%            |
| 1              | 5%             |

| Priority-Group | Transmission Selection Algorithm |
|----------------|----------------------------------|
| 0              | Enhanced Transmission Selection  |
| 1              | Enhanced Transmission Selection  |

Peer-Advertisement:

TLV Type: Configuration

Willing: No

Credit Based Shaper: No

| Code Point | Priority-Group |
|------------|----------------|
| 000        | 0              |
| 001        | 7              |
| 010        | 7              |
| 011        | 7              |
| 100        | 0              |
| 101        | 1              |
| 110        | 1              |
| 111        | 7              |

| Priority-Group | Percentage B/W |
|----------------|----------------|
| 0              | 40%            |
| 1              | 5%             |

| Priority-Group | Transmission Selection Algorithm |
|----------------|----------------------------------|
| 0              | Enhanced Transmission Selection  |
| 1              | Enhanced Transmission Selection  |

Peer-Advertisement:

TLV Type: Recommendation

| Code Point | Priority-Group |
|------------|----------------|
| 000        | 0              |
| 001        | 7              |
| 010        | 7              |
| 011        | 7              |
| 100        | 0              |

|                |                                  |
|----------------|----------------------------------|
| 101            | 1                                |
| 110            | 1                                |
| 111            | 7                                |
| Priority-Group | Percentage B/W                   |
| 0              | 40%                              |
| 1              | 5%                               |
| Priority-Group | Transmission Selection Algorithm |
| 0              | Enhanced Transmission Selection  |
| 1              | Enhanced Transmission Selection  |

**show dcbx neighbors terse (QFX Series)**

```

user@switch> show dcbx neighbors terse
Interface Parent PFC ETS ETS Version
Interface Rec
xe-0/0/8.0 - Enabled Advt Advt IEEE (Auto)
xe-0/0/9.0 - Disabled Disabled 1.01
xe-0/0/11.0 ae0.0 Enabled Advt Advt IEEE (Auto)
xe-0/0/12.0 ae0.0 Enabled Advt Advt IEEE (Auto)
xe-0/0/32.0 - Enabled Advt Not Advt IEEE
xe-0/0/36.0 - Not Advt Advt Advt IEEE

```

**show dcbx neighbors (EX4500 Switch: FCoE Interfaces on Both Local and Peer with PFC Configured Compatibly)**

```

user@switch> show dcbx neighbors interface xe-0/0/14
Interface : xe-0/0/14.0 - Parent Interface: ae0.0
Protocol-State: in-sync

Local-Advertisement:
 Operational version: 0
 sequence-number: 6, acknowledge-id: 6

Peer-Advertisement:
 Operational version: 0
 sequence-number: 6, acknowledge-id: 6

Feature: PFC, Protocol-State: in-sync

Operational State: Enabled

Local-Advertisement:
 Enable: Yes, Willing: No, Error: No
 Maximum Traffic Classes capable to support PFC: 6

Code Point Admin Mode
000 Disabled
001 Disabled
010 Disabled
011 Enabled
100 Disabled
101 Disabled
110 Disabled
111 Disabled

```

## Peer-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 6

| Code Point | Admin Mode |
|------------|------------|
| 000        | Disabled   |
| 001        | Disabled   |
| 010        | Disabled   |
| 011        | Enabled    |
| 100        | Disabled   |
| 101        | Disabled   |
| 110        | Disabled   |
| 111        | Disabled   |

Feature: Application, Protocol-State: in-sync

## Local-Advertisement:

Enable: Yes, Willing: No, Error: No <<< Error bit will not be set as there is no miss configuration between local and peer.

| Appl-Name | Ethernet-Type | Socket-Number | Priority-Map | Status  |
|-----------|---------------|---------------|--------------|---------|
| FCoE      | 0x8906        |               | 00001000     | Enabled |

## Peer-Advertisement:

Enable: Yes, Willing: No, Error: No

| Status  | Appl-Name | Ethernet-Type | Socket-Number | Priority-Map |
|---------|-----------|---------------|---------------|--------------|
| Enabled | FCoE      | 0x8906        |               | 00001000     |

### show dcbx neighbors (EX4500 Switch: DCBX Interfaces on Local and Peer Are Configured Compatibly with iSCSI Application)

user@switch&gt; show dcbx neighbors interface xe-0/0/14

Interface : xe-0/0/14.0 - Parent Interface: ae0.0

Protocol-State: in-sync

Active-application-map: iscsi-map

## Local-Advertisement:

Operational version: 0

sequence-number: 9, acknowledge-id: 12

## Peer-Advertisement:

Operational version: 0

sequence-number: 12, acknowledge-id: 9

Feature: PFC, Protocol-State: in-sync

Operational State: Enabled

## Local-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 6

| Code Point | Admin Mode |
|------------|------------|
| 000        | Disabled   |
| 001        | Disabled   |
| 010        | Disabled   |
| 011        | Enabled    |
| 100        | Disabled   |
| 101        | Disabled   |
| 110        | Disabled   |
| 111        | Disabled   |

## Peer-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 6

| Code Point | Admin Mode |
|------------|------------|
| 000        | Disabled   |
| 001        | Disabled   |
| 010        | Disabled   |
| 011        | Enabled    |
| 100        | Disabled   |
| 101        | Disabled   |
| 110        | Disabled   |
| 111        | Disabled   |

Feature: Application, Protocol-State: in-sync

## Local-Advertisement:

Enable: Yes, Willing: No, Error: No

| Appl-Name | Ethernet-Type | Socket-Number | Priority-Map | Status  |
|-----------|---------------|---------------|--------------|---------|
| FCoE      | 0x8906        |               | 00001000     | Enabled |
| iscsi     |               | 3260          | 00100000     | Enabled |

## Peer-Advertisement:

Enable: Yes, Willing: No, Error: No

| Appl-Name | Ethernet-Type | Socket-Number | Priority-Map | Status  |
|-----------|---------------|---------------|--------------|---------|
| FCoE      | 0x8906        |               | 00001000     | Enabled |
| iscsi     |               | 3260          | 00100000     | Enabled |

**show dcbx neighbors (EX4500 Switch: Includes ETS)**

user@switch&gt; show dcbx neighbors interface xe-0/0/3

Interface : xe-0/0/3.0  
Protocol-State: in-sync  
Active-application-map: map\_iscsi

## Local-Advertisement:

Operational version: 0



sequence-number: 1, acknowledge-id: 5

Peer-Advertisement:

Operational version: 0

sequence-number: 5, acknowledge-id: 1

Feature: PFC, Protocol-State: in-sync

Operational State: Enabled

Local-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 6

| Code Point | Admin Mode |
|------------|------------|
| 000        | Enabled    |
| 001        | Enabled    |
| 010        | Disabled   |
| 011        | Disabled   |
| 100        | Disabled   |
| 101        | Disabled   |
| 110        | Disabled   |
| 111        | Disabled   |

Peer-Advertisement:

Enable: Yes, Willing: Yes, Error: No

Maximum Traffic Classes capable to support PFC: 8

| Code Point | Admin Mode |
|------------|------------|
| 000        | Enabled    |
| 001        | Disabled   |
| 010        | Disabled   |
| 011        | Disabled   |
| 100        | Enabled    |
| 101        | Disabled   |
| 110        | Disabled   |
| 111        | Disabled   |

Feature: Application, Protocol-State: in-sync

Local-Advertisement:

Enable: Yes, Willing: No, Error: No

| App1-Name | Ethernet-Type | Socket-Number | Priority-Map | Status  |
|-----------|---------------|---------------|--------------|---------|
| FCoE      | 0x8906        |               | 00000001     | Enabled |
| iscsi     |               | 3260          | 00000010     | Enabled |

Peer-Advertisement:

Enable: Yes, Willing: Yes, Error: No

| App1-Name | Ethernet-Type | Socket-Number | Priority-Map | Status  |
|-----------|---------------|---------------|--------------|---------|
| FCoE      | 0x8906        |               | 0001000      | Enabled |
| iscsi     |               | 3260          | 00010000     | Enabled |

Feature: ETS, Protocol-State: in-sync

Operational State: Enabled

**Local-Advertisement:**

Enable: Yes, Willing: No, Error: No  
Maximum Traffic Classes supported : 3

| Code Point     | Priority-Group |
|----------------|----------------|
| 000            | 7              |
| 001            | 7              |
| 010            | 7              |
| 011            | 7              |
| 100            | 7              |
| 101            | 7              |
| 110            | 7              |
| 111            | 7              |
| Priority-Group | Percentage B/W |
| 7              | 100%           |

**Peer-Advertisement:**

Enable: Yes, Willing: Yes, Error: No  
Maximum Traffic Classes supported : 8

| Code Point     | Priority-Group |
|----------------|----------------|
| 000            | 0              |
| 001            | 1              |
| 010            | 0              |
| 011            | 0              |
| 100            | 2              |
| 101            | 0              |
| 110            | 0              |
| 111            | 0              |
| Priority-Group | Percentage B/W |
| 0              | 30%            |
| 1              | 40%            |
| 2              | 30%            |

# Traffic Management Feature Guide for QFX Switches



## PART 90

# CoS Overview

- [Basic Concepts on page 6607](#)



## CHAPTER 252

# Basic Concepts

- [Overview of Junos OS CoS on page 6608](#)
- [CoS Support on QFX Series Switches, EX4600 Switches, and QFabric Systems on page 6610](#)
- [Configuring CoS on page 6618](#)
- [Understanding Junos CoS Components on page 6621](#)
- [Assigning CoS Components to Interfaces on page 6627](#)
- [Understanding CoS Packet Flow on page 6628](#)
- [Understanding Default CoS Settings on page 6630](#)
- [CoS Inputs and Outputs Overview on page 6639](#)
- [Overview of Policers on page 6639](#)

## Overview of Junos OS CoS

---

When a network experiences congestion and delay, some packets must be dropped. Junos OS class of service (CoS) enables you to divide traffic into classes and set various levels of throughput and packet loss when congestion occurs. You have greater control over packet loss because you can configure rules tailored to your needs.

You can configure CoS features to provide multiple classes of service for different applications. CoS also allows you to rewrite the Differentiated Services code point (DSCP) or IEEE 802.1p code-point bits of packets leaving an interface, thus allowing you to tailor packets for the network requirements of the remote peers.

CoS provides multiple classes of service for different applications. You can configure multiple forwarding classes for transmitting packets, define which packets are placed into each output queue, schedule the transmission service level for each queue, and manage congestion using a weighted random early detection (WRED) algorithm.

In designing CoS applications, you must carefully consider your service needs, and you must thoroughly plan and design your CoS configuration to ensure consistency and interoperability across all platforms in a CoS domain.

Because CoS is implemented in hardware rather than in software, you can experiment with and deploy CoS features without affecting packet forwarding and switching performance.



**NOTE:** CoS policies can be enabled or disabled on each switch interface. Also, each physical and logical interface on the switch can have associated custom CoS rules.

When you change or when you deactivate and then reactivate the class-of-service configuration, the system experiences packet drops because the system momentarily blocks traffic to change the mapping of incoming traffic to input queues.

---

This topic describes:

- [CoS Standards on page 6608](#)
- [How Junos OS CoS Works on page 6609](#)
- [Default CoS Behavior on page 6610](#)

## CoS Standards

The following RFCs define the standards for CoS capabilities:

- RFC 2474, *Definition of the Differentiated Services Field in the IPv4 and IPv6 Headers*
- RFC 2597, *Assured Forwarding PHB Group*
- RFC 2598, *An Expedited Forwarding PHB*



- RFC 2698, *A Two Rate Three Color Marker*
- RFC 3168, *The Addition of Explicit Congestion Notification (ECN) to IP*

The following data center bridging (DCB) standards are also supported to provide the CoS (and other characteristics) that Fibre Channel over Ethernet (FCoE) requires for transmitting storage traffic over an Ethernet network:

- IEEE 802.1Qbb, priority-based flow control (PFC)
- IEEE 802.1Qaz, enhanced transmission selection (ETS)
- IEEE 802.1AB (LLDP) extension called Data Center Bridging Capability Exchange Protocol (DCBX)



**NOTE:** OCX Series switches and NFX250 Network Services platforms do not support PFC and DCBX.

Juniper Networks QFX10000 switches support both enhanced transmission selection (ETS) hierarchical port scheduling and direct port scheduling.

## How Junos OS CoS Works


Junos OS CoS works by examining traffic entering the edge of your network. The switch classifies traffic into defined service groups to provide the special treatment of traffic across the network. For example, you can send voice traffic across certain links and data traffic across other links. In addition, the data traffic streams can be serviced differently along the network path to ensure that higher-paying customers receive better service. As the traffic leaves the network at the far edge, you can reclassify the traffic to meet the policies of the targeted peer by rewriting the DSCP or IEEE 802.1 code-point bits.

To support CoS, you must configure each switch in the network. Generally, each switch examines the packets that enter it to determine their CoS settings. These settings dictate which packets are transmitted first to the next downstream switch. Switches at the edges of the network might be required to alter the CoS settings of the packets that enter the network to classify the packets into the appropriate service groups.

In [Figure 208](#), Switch A is receiving traffic. As each packet enters, Switch A examines the packet's current CoS settings and classifies the traffic into one of the groupings defined on the switch. This definition allows Switch A to prioritize its resources for servicing the traffic streams it receives. Switch A might alter the CoS settings (forwarding class and loss priority) of the packets to better match the defined traffic groups.

When Switch B receives the packets, it examines the CoS settings, determines the appropriate traffic groups, and processes the packet according to those settings. It then transmits the packets to Switch C, which performs the same actions. Switch D also examines the packets and determines the appropriate groups. Because Switch D sits at the far end of the network, it can reclassify (rewrite) the CoS code-point bits of the packets before transmitting them.

Figure 208: Packet Flow Across the Network



### Default CoS Behavior

If you do not configure CoS settings, the software performs some CoS functions to ensure that the system forwards traffic and protocol packets with minimum delay when the network is experiencing congestion. Some CoS settings, such as classifiers, are automatically applied to each logical interface that you configure. Other settings, such as rewrite rules, are applied only if you explicitly associate them with an interface.

#### Related Documentation

- [CoS Support on QFX Series Switches, EX4600 Switches, and QFabric Systems on page 6610](#)
- [Overview of Policers on page 5999](#)
- [Understanding Junos CoS Components on page 6621](#)
- [Understanding CoS Packet Flow on page 6628](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 6765](#)

### CoS Support on QFX Series Switches, EX4600 Switches, and QFabric Systems

Juniper Networks data center switches differ in some aspects of class-of-service (CoS) support because of differences in the way the switches are used in networks, and because of hardware differences such as different chipsets or different interface capabilities.

This topic provides summaries of CoS support on QFX10000, QFX5100, QFX5200, EX4600, QFX3500, and QFX3600 switches, and QFabric systems for:

- CoS features
- Ethernet interface types
- Classifier and rewrite rule entries on QFX10000 switches

#### CoS Feature Support

The first two tables list CoS feature support for newer ELS-CLI-based platforms ([Table 525](#)) such as the QFX5100, QFX5200, EX4600, and QFX10000 switches, and for legacy-CLI-based platforms ([Table 526](#)) such as QFX3500 switches and QFabric systems. Some legacy-CLI-based platforms can also run the ELS CLI.

**Table 525: QFX10000, QFX5100, QFX5200, and EX4600 Switch CoS Features (As of Software Release 15.1X53-D30)**

| Feature                                                                                       | QFX10000                                                                                      | QFX5100, QFX5200, EX4600                                        |
|-----------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------|-----------------------------------------------------------------|
| Class of service (CoS)—Class-based queuing with prioritization                                | Yes                                                                                           | Yes                                                             |
| CoS—Separate unicast and multi-destination classifiers, forwarding classes, and output queues | No                                                                                            | Yes                                                             |
| CoS—Shared unicast and multideestination classifiers, forwarding classes, and output queues   | Yes                                                                                           | No                                                              |
| CoS support on link aggregation groups (LAGs)                                                 | Yes                                                                                           | Yes                                                             |
| Enhanced transmission selection (ETS) hierarchical port scheduling                            | No                                                                                            | No                                                              |
| Direct port scheduling                                                                        | Yes                                                                                           | Yes                                                             |
| Queue shaping                                                                                 | Yes<br><br><b>NOTE:</b> Uses the <b>transmit-rate</b> statement with the <b>exact</b> option. | Yes<br><br><b>NOTE:</b> Uses the <b>shaping-rate</b> statement. |
| Explicit congestion notification (ECN)                                                        | Yes                                                                                           | Yes                                                             |
| Priority-based flow control (PFC)                                                             | Yes                                                                                           | Yes                                                             |
| Re-marking of bridged packets                                                                 | Yes                                                                                           | Yes                                                             |
| Weighted random early detection (WRED) packet drop profiles and tail drop                     | Yes                                                                                           | Yes                                                             |
| 802.3X Ethernet PAUSE                                                                         | Yes                                                                                           | Yes                                                             |
| Layer 2 ingress packet classification and egress rewrite rules                                | Yes                                                                                           | Yes                                                             |

**Table 525: QFX10000, QFX5100, QFX5200, and EX4600 Switch CoS Features (As of Software Release 15.1X53-D30) (continued)**

| Feature                                                         | QFX10000      | QFX5100, QFX5200, EX4600 |
|-----------------------------------------------------------------|---------------|--------------------------|
| MPLS EXP ingress packet classification and egress rewrite rules | Yes           | Yes                      |
| Layer 3 ingress packet classification and egress rewrite rules  | Yes           | Yes                      |
| Virtual output queue (VOQ) architecture                         | Yes           | No                       |
| Software shared buffer configurability                          | No (uses VOQ) | Yes                      |
| Queue shaping                                                   | Yes           | Yes                      |
| CoS command to detect the source of RED-dropped packets         | Yes           | No                       |

Table 526 shows CoS support for legacy-CLI-based switches.

**Table 526: QFX3500 and QFX3600 Switch, and QFabric System CoS Features (As of Software Release 15.1X53-D30)**

| Feature                                                                                       | QFX3500 | QFX3600 | QFabric System |
|-----------------------------------------------------------------------------------------------|---------|---------|----------------|
| Class of service (CoS)—Class-based queuing with prioritization                                | Yes     | Yes     | Yes            |
| CoS—Separate unicast and multideestination classifiers, forwarding classes, and output queues | Yes     | Yes     | Yes            |
| CoS support on link aggregation groups (LAGs)                                                 | Yes     | Yes     | Yes            |
| Enhanced transmission selection (ETS) hierarchical port scheduling                            | Yes     | Yes     | Yes            |

**Table 526: QFX3500 and QFX3600 Switch, and QFabric System CoS Features (As of Software Release 15.1X53-D30) *(continued)***

| Feature                                                         | QFX3500 | QFX3600 | QFabric System |
|-----------------------------------------------------------------|---------|---------|----------------|
| Direct port scheduling                                          | No      | No      | No             |
| Queue shaping                                                   | Yes     | Yes     | Yes            |
| Explicit congestion notification (ECN)                          | Yes     | Yes     | Yes            |
| Priority-based flow control (PFC)                               | Yes     | Yes     | Yes            |
| Re-marking of bridged packets                                   | Yes     | Yes     | Yes            |
| Priority remapping on native Fibre Channel interfaces           | Yes     | No      | No             |
| Weighted random early detection (WRED) tail-drop profiles       | Yes     | Yes     | Yes            |
| 802.3X Ethernet PAUSE                                           | Yes     | Yes     | Yes            |
| Layer 2 ingress packet classification and egress rewrite rules  | Yes     | Yes     | Yes            |
| MPLS EXP ingress packet classification and egress rewrite rules | Yes     | Yes     | Yes            |
| Layer 3 ingress packet classification and egress rewrite rules  | Yes     | Yes     | Yes            |
| Software buffer configurability                                 | Yes     | Yes     | No             |

**Classifier and Rewrite  
Rule Ethernet Interface  
Type Support**

The next two tables in this topic list CoS Ethernet support for classifiers and rewrite rules on different interface types for QFX10000 switches ([Table 527](#)), and for QFX5100, QFX5200, QFX3500, QFX3600, and EX4600 switches, and QFabric systems ([Table 528](#)).

On QFX10000 switches, you cannot apply classifiers or rewrite rules to Layer 2 or Layer 3 physical interfaces. You can apply classifiers and rewrite rules only to Layer 2 logical interface unit 0. You can apply different classifiers and rewrite rules to different Layer 3 logical interfaces. [Table 527](#) shows on which interfaces you can configure and apply classifiers and rewrite rules.

**Table 527: Ethernet Interface Support for Classifier and Rewrite Rule Configuration (QFX10000 Switches)**

| CoS Classifiers and Rewrite Rules | Layer 2 Physical Interfaces | Layer 2 Logical Interface (Unit 0 Only) | Layer 3 Physical Interfaces | Layer 3 Logical Interfaces |
|-----------------------------------|-----------------------------|-----------------------------------------|-----------------------------|----------------------------|
| Fixed classifier                  | No                          | Yes                                     | No                          | Yes                        |
| DSCP classifier                   | No                          | Yes                                     | No                          | Yes                        |
| DSCP IPv6 classifier              | No                          | Yes                                     | No                          | Yes                        |
| IEEE 802.1p classifier            | No                          | Yes                                     | No                          | Yes                        |
| EXP classifier                    | No                          | Yes                                     | No                          | Yes                        |
| DSCP rewrite rule                 | No                          | Yes                                     | No                          | Yes                        |
| DSCP IPv6 rewrite rule            | No                          | Yes                                     | No                          | Yes                        |
| IEEE 802.1p rewrite rule          | No                          | Yes                                     | No                          | Yes                        |
| EXP rewrite rule                  | No                          | Yes                                     | No                          | Yes                        |

On QFX5100, QFX5200, QFX3500, QFX3600, and EX4600 switches, and QFabric systems, you cannot apply classifiers or rewrite rules to Layer 2 physical interfaces or to Layer 3 logical interfaces. [Table 528](#) shows on which interfaces you can configure and apply classifiers and rewrite rules.

**Table 528: Ethernet Interface Support for Classifier and Rewrite Rule Configuration (QFX5100, QFX5200, EX4600, QFX3500, and QFX3600 Switches, and QFabric Systems)**

| CoS Classifiers and Rewrite Rules | Layer 2 Physical Interfaces | Layer 2 Logical Interface (Unit 0 Only) | Layer 3 Physical Interfaces (If at Least One Logical Layer 3 Interface Is Defined) | Layer 3 Logical Interfaces |
|-----------------------------------|-----------------------------|-----------------------------------------|------------------------------------------------------------------------------------|----------------------------|
| Fixed classifier                  | No                          | Yes                                     | Yes                                                                                | No                         |
| DSCP classifier                   | No                          | Yes                                     | Yes                                                                                | No                         |

**Table 528: Ethernet Interface Support for Classifier and Rewrite Rule Configuration (QFX5100, QFX5200, EX4600, QFX3500, and QFX3600 Switches, and QFabric Systems) (continued)**

| CoS Classifiers and Rewrite Rules | Layer 2 Physical Interfaces                                                                                                                         | Layer 2 Logical Interface (Unit 0 Only) | Layer 3 Physical Interfaces (If at Least One Logical Layer 3 Interface Is Defined) | Layer 3 Logical Interfaces |
|-----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------|------------------------------------------------------------------------------------|----------------------------|
| DSCP IPv6 classifier              | No                                                                                                                                                  | Yes                                     | Yes                                                                                | No                         |
| IEEE 802.1p classifier            | No                                                                                                                                                  | Yes                                     | Yes                                                                                | No                         |
| EXP classifier                    | Global classifier, applies only to all switch interfaces that are configured as <b>family mpls</b> . Cannot be configured on individual interfaces. |                                         |                                                                                    |                            |
| DSCP rewrite rule                 | No                                                                                                                                                  | Yes                                     | Yes                                                                                | No                         |
| DSCP IPv6 rewrite rule            | No                                                                                                                                                  | Yes                                     | Yes                                                                                | No                         |
| IEEE 802.1p rewrite rule          | No                                                                                                                                                  | Yes                                     | Yes                                                                                | No                         |
| EXP rewrite rule                  | No                                                                                                                                                  | Yes                                     | Yes                                                                                | No                         |



**NOTE:** IEEE 802.1p multidestination and DSCP multidestination classifiers are applied to all interfaces and cannot be applied to individual interfaces. No DSCP IPv6 multidestination classifier is supported. IPv6 multidestination traffic uses the DSCP multidestination classifier.

Even though feature support for QFX5100 and QFX5200 switches is the same, they use different chipsets, and therefore have some CoS operational differences. [Table 529](#) details both the similarities and differences for CoS on QFX5100 and QFX5200.

**Table 529: CoS Operational Comparison Between QFX5100 and QFX5200**

| CoS Feature                 | QFX5100                              | QFX5200                                  | Change in Operation                                                                          |
|-----------------------------|--------------------------------------|------------------------------------------|----------------------------------------------------------------------------------------------|
| Memory Management           | Central MMU shared by all ports      | Crosspoint architecture with quad pipe   | No customer visible change.                                                                  |
| Pipes                       | 2                                    | 4                                        | No customer visible change.                                                                  |
| Cell Accounting             | Global access pipes                  | Local to Cross point (4MB / cross point) | No customer visible change.                                                                  |
| Shared Buffer               | 60k Cells (Each cell 208Bytes), 12MB | 80K Cells (Each cell 208 Bytes), 16MB    | No customer visible change. QFX-5200 just supports larger packet buffer space than QFX-5100. |
| Shared buffer pool per pipe | 4 pools per pipe (no change)         |                                          | N/A                                                                                          |

Table 529: CoS Operational Comparison Between QFX5100 and QFX5200 (*continued*)

| CoS Feature             | QFX5100                             | QFX5200                                                                                                                                                                                                                                                                                                                                                                   | Change in Operation                                                                     |
|-------------------------|-------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------|
| Queuing and Scheduling  | LLS and three level hierarchy       | Fixed hierarchical scheduling (FHS) and two-level hierarchy                                                                                                                                                                                                                                                                                                               | ETS and FC-Set are not supported on QFX-5200 due to FHS.                                |
| # Unicast Queues        | 8                                   |                                                                                                                                                                                                                                                                                                                                                                           | N/A                                                                                     |
| # Multicast Queues      | 4                                   | 2                                                                                                                                                                                                                                                                                                                                                                         |                                                                                         |
| CPU Queues              | 44                                  |                                                                                                                                                                                                                                                                                                                                                                           | N/A                                                                                     |
| Host Path Scheduling    | 48 queues directly attached to port | 48 queues attached to LO                                                                                                                                                                                                                                                                                                                                                  | No customer visible change. SDK API change just affects software development effort.    |
| FC2Q                    | 4 profiles (no change)              |                                                                                                                                                                                                                                                                                                                                                                           | N/A                                                                                     |
| DSCP classifier table   | 64 profiles (no change)             |                                                                                                                                                                                                                                                                                                                                                                           | N/A                                                                                     |
| 802.1p classifier table | 64 profiles                         | 64 profiles                                                                                                                                                                                                                                                                                                                                                               | No customer visible change. SDK API change just affects software development effort.    |
| PFC                     | Common headroom buffer              | Per pipe headroom buffer                                                                                                                                                                                                                                                                                                                                                  | Available and used head room buffer is maintained separately for each pipe on QFX-5200. |
| Rewrite                 | 128 profiles                        |                                                                                                                                                                                                                                                                                                                                                                           | No customer visible change. SDK API change just affects software development effort.    |
| WRED                    | 128 profiles per pipe (same)        |                                                                                                                                                                                                                                                                                                                                                                           | N/A                                                                                     |
| Scheduler Levels        |                                     | Eight traffic classes per port. For each traffic class, a pair of unicast and multicast queue is associated per XPE. Another two Queues named Queue Management(QM) and System Control(SC) totalling 10 Unicast and 10 Multicast Queues for each XPE. These physical queues are mapped to logical queues (10 Unicast and 10 Multicast) at each MMU slice controllers level |                                                                                         |
| Queueing Levels         |                                     | Three levels, logical queue level, CoS level, and port level.                                                                                                                                                                                                                                                                                                             |                                                                                         |



Table 529: CoS Operational Comparison Between QFX5100 and QFX5200 (*continued*)

| CoS Feature              | QFX5100                                                                                                                                       | QFX5200                                                                                                                                                            | Change in Operation |
|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------|
| Multidestination Traffic | Default scheduler map reserves 20% bandwidth for multicast and 80% of unicast traffic reserved between BE, FCoE, NoLoss and NC traffic types. | Each level 0 node is receiving both multicast and unicast traffic, so it is not possible to differentiate at the port level to apply shaping on multicast traffic. |                     |

The following limitations on QFX5200 switches do not exist on QFX5100 switches.

- CoS ETS is not supported on QFX5200.
- On QFX5200 switches, it is not possible to support multiple queues with **strict-high** priority because QFX5200 doesn't support flexible hierarchical scheduling. When multiple strict-high priority queues are configured, all of those queues are treated as strict-high priority but the higher number queue among them is given highest priority.
- QFX5200 CoS policers do not support global management counters accessed by all ports. Only management counters local to a pipeline are supported—this means that QFX5200 management counters work only on traffic received on ports that belong to the pipeline in which the counter is created.
- Due to QFX5200 cross-point architecture, all buffer usage counters are maintained separately. When usage counters are displayed with the command **show class-of-service shared-buffer**, various pipe counters are displayed separately.
- On QFX5200, port schedulers are supported instead of FC-SET.
- On QFX5200, it is not possible to group multiple forwarding classes into a forwarding class set and apply output traffic control profile on the fc-set. ETS for an fc-set is not supported. Because each L0 node schedules both the unicast and multicast queue of L1 node, it is not possible to differentiate multicast/unicast traffic at the port level and apply minimum bandwidth between unicast and multicast. It can only be supported at CoS level L0.
- Because QFX5200 does not support flexible hierarchical scheduling, it is not possible to apply a traffic control profile for a group of forwarding classes.

#### QFX10000 Switch Classifier and Rewrite Rule Support (Scaling)

You can configure enough classifiers on QFX10000 switches to handle most, if not all, network scenarios. [Table 530](#) shows how many of each type of classifiers you can configure, and how many entries you can configure per classifier.

**Table 530: Classifier Support by Classifier Type**

| Classifier Type       | Default Classifier Name                                                                         | Maximum Number of Classifiers | Maximum Number of Entries per Classifier |
|-----------------------|-------------------------------------------------------------------------------------------------|-------------------------------|------------------------------------------|
| IEEE 802.1p (Layer 2) | ieee8021p-default (for ports in trunk mode)<br><br>ieee8021p-untrust (for ports in access mode) | 64                            | 16                                       |
| DSCP (Layer 3)        | dscp-default                                                                                    | 64                            | 64                                       |
| DSCP IPv6 (Layer 3)   | dscp-ipv6-default                                                                               | 64                            | 64                                       |
| EXP (MPLS)            | exp-default                                                                                     | 64                            | 8                                        |
| Fixed                 | There is no default fixed classifier                                                            | 8                             | 16                                       |

The number of fixed classifiers supported (8) equals the number of supported forwarding classes (fixed classifiers assign all incoming traffic on an interface to one forwarding class).

There are no default rewrite rules. You can configure enough rewrite rules on QFX10000 switches to handle most, if not all, network scenarios. [Table 531](#) shows how many of each type of rewrite rule you can configure, and how many entries you can configure per rewrite rule.

**Table 531: Rewrite Rule Support by Rewrite Rule Type**

| Rewrite Rule Type     | Maximum Number of Rewrite Rule Sets | Maximum Number of Entries per Rewrite Rule Set |
|-----------------------|-------------------------------------|------------------------------------------------|
| IEEE 802.1p (Layer 2) | 64                                  | 128                                            |
| DSCP (Layer 3)        | 32                                  | 128                                            |
| DSCP IPv6 (Layer 3)   | 32                                  | 128                                            |
| EXP (MPLS)            | 64                                  | 128                                            |

## Configuring CoS

The traffic management class-of-service topics describe how to configure the Junos OS class-of-service (CoS) components. Junos CoS provides a flexible set of tools that enable you to fine tune control over the traffic on your network.

You can define classifiers that classify incoming traffic into forwarding classes to place traffic in groups for transmission. You can map forwarding classes to output queues to define the type of traffic on each output queue. You can configure schedulers for each

output queue to control the service level (priority, bandwidth characteristics) of each type of traffic. You can provide different service levels for the same forwarding classes on different interfaces. Some switches support data center bridging standards so that you can configure lossless transport across the Ethernet network using priority-based flow control (PFC), Data Center Bridging Exchange protocol (DCBX), and enhanced transmission selection (ETS) hierarchical scheduling (OCX Series switches and NFX250 Network Services platform do not support lossless transport, PFC, and DCBX).

You can configure various CoS components individually or in combination to define CoS services.



**NOTE:** When you change the CoS configuration or when you deactivate and then reactivate the CoS configuration, the system experiences packet drops because the system momentarily blocks traffic to change the mapping of incoming traffic to input queues.

The following topics describe how to configure CoS components.



**NOTE:** Links to features that are not supported on the platform for which you are looking up information might not be functional. The platforms that support each feature are listed before each topic link. Feature topics that do not list platforms apply to all platforms.

- (QFX3500, QFX3600, EX4600, NFX250, QFX5100, QFX5200, and OCX1100 switches, and QFabric systems) *Defining CoS Unicast BA Classifiers (DSCP, DSCP IPv6, IEEE 802.1p)*
- (QFX10000 switches) [“Example: Configuring Classifiers” on page 6657](#)
- (QFX3500, QFX3600, EX4600, QFX5100, QFX5200, and OCX1100 switches, and QFabric systems) *Example: Configuring Multidestination (Multicast, Broadcast, DLF) Classifiers*
- (QFX3500, QFX3600, EX4600, QFX5100, QFX5200, and OCX1100 switches, and QFabric systems) *Example: Configuring WRED Drop Profiles*
- (QFX10000 switches) [“Example: Configuring WRED Drop Profiles” on page 6814](#)
- (All switches) [“Example: Configuring Drop Profile Maps” on page 6817](#)
- (QFX3500, QFX3600, EX4600, NFX250, QFX5100, QFX5200, and OCX1100 switches, and QFabric systems) *Example: Configuring Forwarding Classes*
- (QFX10000 switches) [“Example: Configuring Forwarding Classes” on page 6692](#)
- (All switches) [“Example: Configuring Forwarding Class Sets” on page 6697](#)
- (QFX3500, QFX3600, EX4600, NFX250, QFX5100, QFX5200, and OCX1100 switches, and QFabric systems) *Example: Configuring Queue Schedulers*
- (QFX10000 switches) [“Example: Configuring Queue Schedulers for Port Scheduling” on page 6744](#)

- (All switches and NFX250 Network Services platform) [“Example: Configuring Queue Scheduling Priority” on page 6753](#)
- (All switches) [“Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\)” on page 6762](#)
- (All switches and NFX250 Network Services platform) [“Example: Configuring Minimum Guaranteed Output Bandwidth” on page 6798](#)
- (All switches) [“Example: Configuring Maximum Output Bandwidth” on page 6805](#)
- (QFX3500, QFX3600, EX4600, QFX5100, and QFX10000 switches, and QFabric systems) [“Example: Configuring DCBX Application Protocol TLV Exchange” on page 6535](#)
- (All switches) [“Example: Configuring CoS Hierarchical Port Scheduling \(ETS\)” on page 6771](#)
- (QFX3500, QFX3600, EX4600, QFX5100, QFX5200, and QFX10000 switches, and QFabric systems) [“Example: Configuring CoS PFC for FCoE Traffic” on page 6504](#)
- (QFX3500, QFX3600, EX4600, QFX5100, QFX5200, and QFX10000 switches, and QFabric systems) *Example: Configuring CoS for FCoE Transit Switch Traffic Across an MC-LAG*
- (QFX3500 switches and QFabric systems) *Example: Configuring IEEE 802.1p Priority Remapping on an FCoE-FC Gateway*
- (QFX3500, QFX3600, EX4600, QFX5100, QFX5200, and QFX10000 switches, and QFabric systems) [“Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces” on page 6986](#)
- (QFX3500, QFX3600, EX4600, QFX5100, QFX5200, and QFX10000 switches, and QFabric systems) [“Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic \(FCoE Transit Switch\)” on page 6969](#)
- (QFX3500, QFX3600, EX4600, QFX5100, QFX5200, and QFX10000 switches, and QFabric systems) [“Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface” on page 6977](#)
- (QFX3500, QFX3600, EX4600, QFX5100, QFX5200, and QFX10000 switches, and QFabric systems) [“Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications \(FCoE and iSCSI\)” on page 7000](#)
- (All switches and NFX250 Network Services platform) [“Defining CoS Code-Point Aliases” on page 6687](#)
- (All switches and NFX250 Network Services platform) [“Defining CoS Rewrite Rules” on page 6704](#)
- (QFX3500, QFX3600, EX4600, QFX5100, QFX5200, and QFX10000 switches, and QFabric systems) [“Configuring CoS PFC \(Congestion Notification Profiles\)” on page 6905](#)
- (QFX3500 switches, NFX250 Network Services platform, and QFabric systems) *Configuring CoS Fixed Classifier Rewrite Values for Native FC Interfaces (NP\_Ports)*
- (All switches and NFX250 Network Services platform) [“Assigning CoS Components to Interfaces” on page 6627](#)

- (NFX250, QFX3500, QFX3600, EX4600, QFX5100, QFX5200, and OCX1100 switches, and QFabric systems) *Changing the Host Outbound Traffic Default Queue Mapping*
- (All switches) [“Defining CoS Traffic Control Profiles \(Priority Group Scheduling\)” on page 6761](#)
- (QFX3500, QFX3600, EX4600, QFX5100, QFX5200, and OCX1100 switches, and QFabric systems) *Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control*
- (QFX3500, QFX3600, EX4600, QFX5200, and QFX5100 switches, and QFabric systems) *Configuring CoS Asymmetric Ethernet PAUSE Flow Control*
- (QFX3500, QFX3600, EX4600, QFX5100, QFX5200, and QFX10000 switches, and QFabric systems) [“Configuring the DCBX Mode” on page 6523](#)
- (QFX3500, QFX3600, EX4600, QFX5100, QFX5200, and QFX10000 switches, and QFabric systems) [“Configuring DCBX Autonegotiation” on page 6524](#)
- (QFX3500, QFX3600, EX4600, QFX5100, and QFX10000 switches, and QFabric systems) [“Disabling the ETS Recommendation TLV” on page 6527](#)
- (QFX3500, QFX3600, EX4600, QFX5100, QFX5200, and QFX10000 switches, and QFabric systems) [“Defining an Application for DCBX Application Protocol TLV Exchange” on page 6532](#)
- (QFX3500, QFX3600, EX4600, QFX5100, QFX5200, and QFX10000 switches, and QFabric systems) [“Configuring an Application Map for DCBX Application Protocol TLV Exchange” on page 6533](#)
- (QFX3500, QFX3600, EX4600, QFX5100, QFX5200, and QFX10000 switches, and QFabric systems) [“Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange” on page 6534](#)

---

## Understanding Junos CoS Components

This topic describes the Junos OS class-of-service (CoS) components:

- [Code-Point Aliases on page 6622](#)
- [Policers on page 6622](#)
- [Classifiers on page 6622](#)
- [Forwarding Classes on page 6623](#)
- [Forwarding Class Sets on page 6624](#)
- [Flow Control \(Ethernet PAUSE, PFC, and ECN\) on page 6624](#)
- [WRED Profiles and Tail Drop on page 6625](#)
- [Schedulers on page 6626](#)
- [Rewrite Rules on page 6626](#)

## Code-Point Aliases

A *code-point alias* assigns a name to a pattern of code-point bits. You can use this name instead of the bit pattern when you configure other CoS components such as classifiers and rewrite rules.

## Policers

*Policers* limit traffic of a certain class to a specified bandwidth and burst size. Packets exceeding the policer limits can be discarded, or can be assigned to a different forwarding class, a different loss priority, or both. You define policers with filters that you can associate with input interfaces.

## Classifiers

Packet classification associates incoming packets with a particular CoS servicing level. In Junos OS, *classifiers* associate packets with a forwarding class and loss priority and assign packets to output queues based on the associated forwarding class. Junos OS supports two general types of classifiers:

- Behavior aggregate (BA) or CoS value traffic classifiers—Examine the CoS value in the packet header. The value in this single field determines the CoS settings applied to the packet. BA classifiers allow you to set the forwarding class and loss priority of a packet based on the Differentiated Services code point (DSCP) value, IEEE 802.1p value, or MPLS EXP value.



**NOTE:** OCX Series switches and NFX250 Network Services platform do not support MPLS.

---

- Multifield traffic classifiers—Examine multiple fields in the packet, such as source and destination addresses and source and destination port numbers of the packet. With multifield classifiers, you set the forwarding class and loss priority of a packet based on firewall filter rules.

On switches that require the separation of unicast and multideestination (multicast, broadcast, and destination lookup fail) traffic, you create separate unicast classifiers and multideestination classifiers. You cannot assign unicast traffic and multideestination traffic to the same classifier. You can apply unicast classifiers to one or more interfaces. Multideestination classifiers apply to all of the switch interfaces and cannot be applied to individual interfaces. Switches that require the separation of unicast and multideestination traffic have 12 output queues to provide 4 output queues reserved for multideestination traffic.

On switches that do not separate unicast and multideestination traffic, unicast and multideestination traffic use the same classifiers, and you do not create a separate special classifier for multideestination traffic. Switches that do not separate unicast and multideestination traffic have eight output queues because no extra queues are required to separate the traffic.

## Forwarding Classes

*Forwarding classes* group packets for transmission and CoS. You assign each packet to an output queue based on the packet's forwarding class. Forwarding classes affect the forwarding, scheduling, and rewrite marking policies applied to packets as they transit the switch.

Switches provide up to five default forwarding classes:

- best-effort—Best-effort traffic
- fcoe—Fibre Channel over Ethernet traffic
- no-loss—Lossless traffic
- network-control—Network control traffic
- mcast—Multicast traffic



**NOTE:** The default mcast forwarding class applies only to switches that require the separation of unicast and multidestination (multicast, broadcast, and destination lookup fail) traffic. On these switches, you create separate forwarding classes for the two types of traffic. The default mcast forwarding class transports only multidestination traffic, and the default best-effort, fcoe, no-loss, and network-control forwarding classes transport only unicast traffic. Unicast forwarding classes map to unicast output queues, and multidestination forwarding classes map to multidestination output queues. You cannot assign unicast traffic and multidestination traffic to the same forwarding class or to the same output queue. Switches that require the separation of unicast and multidestination traffic have 12 output queues, 8 for unicast traffic and 4 for multidestination traffic.

On switches that do not separate unicast and multidestination traffic, unicast and multidestination traffic use the same forwarding classes and output queues, so the mcast forwarding class is not valid. You do not create separate forwarding classes for multidestination traffic. Switches that do not separate unicast and multidestination traffic have eight output queues because no extra queues are required to separate the traffic.



**NOTE:** On OCX Series switches only, do not map traffic to the default fcoe and no-loss forwarding classes. By default, the DSCP default classifier does not map traffic to the fcoe and no-loss forwarding classes, so by default, OCX Series switches do not classify traffic into those forwarding classes. (On other switches, the fcoe and no-loss forwarding classes provide lossless transport for Layer 2 traffic. OCX Series switches do not support lossless Layer 2 transport.)

Switches support a total of either 12 forwarding classes (8 unicast forwarding classes and 4 multicast forwarding classes), or 8 forwarding classes (unicast and multdestination traffic use the same forwarding classes), which provides flexibility in classifying traffic.

NFX250 Network Services platform provide the following forwarding classes:

- best-effort (be)—Provides no service profile. Loss priority is typically not carried in a CoS value.
- expedited-forwarding (ef)—Provides a low loss, low latency, low jitter, assured bandwidth, end-to-end service.
- assured-forwarding (af)—Provides a group of values you can define and includes four subclasses: AF1, AF2, AF3, and AF4, each with two drop probabilities: low and high.
- network-control (nc)—Supports protocol control and thus is typically high priority.

## Forwarding Class Sets

You can group forwarding classes (output queues) into *forwarding class sets* to apply CoS to groups of traffic that require similar treatment. Forwarding class sets map traffic into priority groups to support enhanced transmission selection (ETS), which is described in IEEE 802.1Qaz.

You can configure up to three unicast forwarding class sets and one multicast forwarding class set. For example, you can configure different forwarding class sets to apply CoS to unicast groups of local area network (LAN) traffic, storage area network (SAN) traffic, and high-performance computing (HPC) traffic, and configure another group for multicast traffic.

Within each forwarding class set, you can configure special CoS treatment for the traffic mapped to each individual queue. This provides the ability to configure CoS in a two-tier hierarchical manner. At the forwarding class set tier, you configure CoS for groups of traffic using a *traffic control profile*. At the queue tier, you configure CoS for individual output queues within a forwarding class set using a *scheduler* that you map to a queue (forwarding class) using a *scheduler map*.

## Flow Control (Ethernet PAUSE, PFC, and ECN)

*Ethernet PAUSE* (described in IEEE 802.3X) is a link-level flow control mechanism. During periods of network congestion, Ethernet PAUSE stops all traffic on a full-duplex Ethernet link for a period of time specified in the PAUSE message.



**NOTE:** QFX10000 switches do not support Ethernet PAUSE.

---

*Priority-based flow control* (PFC) is described in IEEE 802.1Qbb as part of the IEEE data center bridging (DCB) specifications for creating a lossless Ethernet environment to transport loss-sensitive flows such as Fibre Channel over Ethernet (FCoE) traffic.





**NOTE:** OCX Series switches do not support PFC.

PFC is a link-level flow control mechanism similar to Ethernet PAUSE. However, Ethernet PAUSE stops all traffic on a link for a period of time. PFC decouples the pause function from the physical link and divides the traffic on the link into eight priorities (3-bit IEEE 802.1p code points). You can think of the eight priorities as eight “lanes” of traffic. You can apply pause selectively to the traffic on any priority without pausing the traffic on other priorities on the same link.

The granularity that PFC provides allows you to configure different levels of CoS for different types of traffic on the link. You can create lossless lanes for traffic such as FCoE, LAN backup, or management, while using standard frame-drop methods of congestion management for IP traffic on the same link.



**NOTE:** If you transport FCoE traffic, you must enable PFC on the priority assigned to FCoE traffic (usually IEEE 802.1p code point 011 on interfaces that carry FCoE traffic).

*Explicit congestion notification* (ECN) enables end-to-end congestion notification between two endpoints on TCP/IP based networks. ECN must be enabled on both endpoints and on all of the intermediate devices between the endpoints for ECN to work properly. Any device in the transmission path that does not support ECN breaks the end-to-end ECN functionality. ECN notifies networks about congestion with the goal of reducing packet loss and delay by making the sending device decrease the transmission rate until the congestion clears, without dropping packets. RFC 3168, *The Addition of Explicit Congestion Notification (ECN) to IP*, defines ECN.

## WRED Profiles and Tail Drop

A weighted random early detection (WRED) profile (drop profile) defines parameters that enable the network to drop packets during periods of congestion. A *drop profile* defines the conditions under which packets of different loss priorities drop, by determining the probability of dropping a packet for each loss priority when output queues become congested. Drop profiles essentially set a value for a level of queue fullness—when the queue fills to the level of the queue fullness value, packets drop. The combination of queue fill level, the probability of dropping a packet at that fill level, and loss priority of the packet, determine whether a packet is dropped or forwarded. Each pairing of a fill level with a drop probability creates a point on a drop profile curve.

You can associate different drop profiles with different loss priorities to set the probability of dropping packets. You can apply a drop profile for each loss priority to a forwarding class (output queue) by applying a drop profile to a scheduler, and then mapping the scheduler to a forwarding class using a scheduler map. When the queue mapped to the forwarding class experiences congestion, the drop profile determines the level of packet drop for traffic of each loss priority in that queue.

Loss priority affects the scheduling of a packet without affecting the packet's relative ordering. Typically you mark packets exceeding a particular service level with a high loss priority.

Tail drop is a simple drop mechanism that drops all packets indiscriminately during periods of congestion, without differentiating among the packet loss priorities of traffic flows. Tail drop requires only one curve point that corresponds to the maximum depth of the output queue, and drop probability when traffic exceeds the buffer depth is 100 percent (all packets that cannot be stored in the queue are dropped). WRED is superior to tail-drop because WRED enables you to treat traffic of different priorities in a differentiated manner, so that higher priority traffic receives preference, and because of the ability to set multiple points on the drop curve.

## Schedulers

Each switch interface has multiple queues assigned to store packets. The switch determines which queue to service based on a particular method of scheduling. This process often involves determining the sequence in which different types of packets should be transmitted.

You can define the scheduling priority (**priority**), minimum guaranteed bandwidth (**transmit-rate**), maximum bandwidth (**shaping-rate**), and WRED profiles to be applied to a particular queue (forwarding class) for packet transmission. By default, extra bandwidth is shared among queues in proportion to the minimum guaranteed bandwidth of each queue. On switches that support the **excess-rate** statement, you can configure the percentage of shared extra bandwidth an output queue receives independently from the minimum guaranteed bandwidth transmit rate, or you can use default bandwidth sharing based on the transmit rate.

A scheduler map associates a specified forwarding class with a scheduler configuration. You can associate up to four user-defined scheduler maps with the interfaces.

## Rewrite Rules

A *rewrite rule* sets the appropriate CoS bits in the outgoing packet. This allows the next downstream device to classify the packet into the appropriate service group. Rewriting (marking) outbound packets is useful when the switch is at the border of a network and must change the CoS values to meet the policies of the targeted peer.



**NOTE:** Ingress firewall filters can also rewrite forwarding class and loss priority values.

---

### Related Documentation

- [CoS Support on QFX Series Switches, EX4600 Switches, and QFabric Systems on page 6610](#)
- [Understanding CoS Packet Flow on page 6628](#)

## Assigning CoS Components to Interfaces

After you define the following CoS components, you assign them to physical or logical interfaces. Components that you assign to physical interfaces are valid for all of the logical interfaces configured on the physical interface. Components that you assign to a logical interface are valid only for that logical interface.

- Classifiers—Assign only to logical interfaces; on some switches, you can apply classifiers to physical Layer 3 interfaces and the classifiers are applied to all logical interfaces on the physical interface.
- Congestion notification profiles—Assign only to physical interfaces.



**NOTE:** OCX Series switches and NFX250 Network Services platform do not support congestion notification profiles.

- Forwarding classes—Assign to interfaces by mapping to forwarding class sets.
- Forwarding class sets—Assign only to physical interfaces.
- Output traffic control profiles—Assign only to physical interfaces (with a forwarding class set).
- Port schedulers—Assign only to physical interfaces on switches that support port scheduling. Associate the scheduler with a forwarding class in a scheduler map and apply the scheduler map to the physical interface.
- Rewrite rules—Assign only to logical interfaces; on some switches, you can apply classifiers to physical Layer 3 interfaces and the classifiers are applied to all logical interfaces on the physical interface.

You can assign a CoS component to a single interface or to multiple interfaces using wildcards. You can also assign a congestion notification profile or a forwarding class set globally to all interfaces.

To assign CoS components to interfaces:

Assign a CoS component to a physical interface by associating a CoS component (for example, a forwarding class set named **be-priority-group**) with an interface:

```
[edit class-of-service interfaces]
user@switch# set xe-0/0/7 forwarding-class-set be-priority-group
```

Assign a CoS component to a logical interface by associating a CoS component (for example, a classifier named **be\_classifier**) with a logical interface:

```
[edit class-of-service interfaces]
user@switch# set xe-0/0/7 unit 0 classifiers dscp be_classifier
```

Assign a CoS component to multiple interfaces by associating a CoS component (for example, a rewrite rule named **customup-rw**) to all 10-Gigabit Ethernet interfaces on the switch, use wildcard characters for the interface name and logical interface (unit) number:

```
[edit class-of-service interfaces]
user@switch# set xe-* unit * rewrite-rules ieee-802.1 customup-rw
```

Assign a congestion notification profile or a forwarding class set globally to all interfaces using the **set class-of-service interfaces all** statement. For example, to assign a forwarding class set named **be-priority-group** to all interfaces:

```
[edit class-of-service interfaces]
user@switch# set all forwarding-class-set be-priority-group
```



**NOTE:** If there is an existing CoS configuration of any type on an interface, the global configuration is not applied to that particular interface. The global configuration is applied to all interfaces that do not have an existing CoS configuration.

For example, if you configure a rewrite rule, assign it to interfaces **xe-0/0/20.0** and **xe-0/0/22.0**, and then configure a forwarding class set and apply it to all interfaces, the forwarding class set is applied to every interface except **xe-0/0/20** and **xe-0/0/22**.

**Related  
Documentation**

- [Monitoring Interfaces That Have CoS Components on page 7161](#)
- [Understanding Junos CoS Components on page 6621](#)

---

## Understanding CoS Packet Flow

When a packet traverses a switch, the switch provides the appropriate level of service to the packet using either default class-of-service (CoS) settings or CoS settings that you configure. On ingress ports, the switch classifies packets into appropriate forwarding classes and assigns a loss priority to the packets. On egress ports, the switch applies packet scheduling and (if you have configured them) rewrite rules to re-mark packets.

You can configure CoS on Layer 2 logical interfaces, and you can configure CoS on Layer 3 physical interfaces if you have defined at least one logical interface on the Layer 3 physical interface. You cannot configure CoS on Layer 2 physical interfaces and Layer 3 logical interfaces.

For Layer 2 traffic, either use the default CoS settings or configure CoS on each logical interface. You can apply different CoS settings to different Layer 2 logical interfaces.



**NOTE:** OCX Series switches do not support Layer 2 interfaces (family **ethernet-switching**).

For Layer 3 traffic, either use the default CoS settings or configure CoS on the physical interface (not on the logical unit). The switch uses the CoS applied on the physical Layer 3 interface for all logical Layer 3 interfaces configured on the physical Layer 3 interface.

The switch applies CoS to packets as they flow through the system:

- An interface has one or more classifiers of different types applied to it (configure this at the **[edit class-of-service interfaces]** hierarchy level). The classifier types are based on the portion of the incoming packet that the classifier examines (IEEE 802.1p code point bits or DSCP code point bits).
- When a packet enters an ingress port, the classifier assigns the packet to a forwarding class and a loss priority based on the code point bits of the packet (configure this at the **[edit class-of-service classifiers]** hierarchy level).
- The switch assigns each forwarding class to an output queue (configure this at the **[edit class-of-service forwarding-classes]** hierarchy level).
- Input (and output) policers meter traffic and can change the forwarding class and loss priority if a traffic flow exceeds its service level.
- A scheduler map is applied to each interface. When a packet exits an egress port, the scheduler map controls how it is treated (configure this at the **[edit class-of-service interfaces]** hierarchy level). A scheduler map assigns schedulers to forwarding classes (configure this at the **[edit class-of-service scheduler-maps]** hierarchy level).
- A scheduler defines how traffic is treated at the egress interface output queue (configure this at the **[edit class-of-service schedulers]** hierarchy level). You control the transmit rate, shaping rate, priority, and drop profile of each forwarding class by mapping schedulers to forwarding classes in scheduler maps, then applying scheduler maps to interfaces.
- A drop-profile defines how aggressively to drop packets that are mapped to a particular scheduler (configure this at the **[edit class-of-service drop-profiles]** hierarchy level).
- A rewrite rule takes effect as the packet leaves an interface that has a rewrite rule configured (configure this at the **[edit class-of-service rewrite-rules]** hierarchy level). The rewrite rule writes information to the packet (for example, a rewrite rule can re-mark the code point bits of outgoing traffic) according to the forwarding class and loss priority of the packet.

Figure 209 is a high-level flow diagram of how packets from various sources enter switch interfaces, are classified at the ingress, and then scheduled (provided bandwidth) at the egress queues.

**Figure 209: CoS Classifier, Queues, and Scheduler**


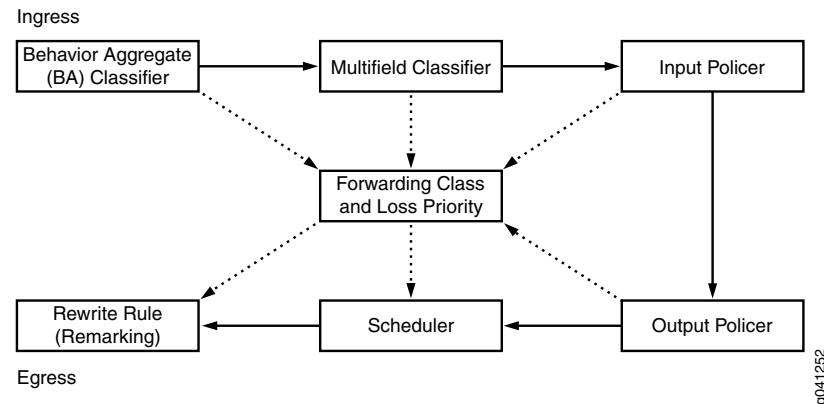


Figure 210 shows the packet flow through the CoS components that you can configure.

**Figure 210: Packet Flow Through Configurable CoS Components**



The middle box (Forwarding Class and Loss Priority) represents two values that you can use on ingress and egress interfaces. The system uses these values for classifying traffic on ingress interfaces and for rewrite rule re-marking on egress interfaces. Each outer box represents a process component. The components in the top row apply to incoming packets. The components in the bottom row apply to outgoing packets.

The solid-line arrows show the direction of packet flow from ingress to egress. The dotted-line arrows that point to the forwarding class and loss priority box indicate processes that configure (set) the forwarding class and loss priority. The dotted-line arrows that point away from the forwarding class and loss priority box indicate processes that use forwarding class and loss priority as input values on which to base actions.

For example, the BA classifier sets the forwarding class and loss priority of incoming packets, so the forwarding class and loss priority are outputs of the classifier and the arrow points away from the classifier. The scheduler receives the forwarding class and loss priority settings, and queues the outgoing packets based on those settings, so the arrow points toward the scheduler.

## Understanding Default CoS Settings

If you do not configure CoS settings, Junos OS performs some CoS functions to ensure that traffic and protocol packets are forwarded with minimum delay when the network experiences congestion. Some default mappings are automatically applied to each logical interface that you configure.

You can display default CoS settings by issuing the **show class-of-service** operational mode command.

This topic describes the default configurations for the following CoS components:

- [Default Forwarding Classes and Queue Mapping on page 6631](#)
- [Default Forwarding Class Sets \(Priority Groups\) on page 6631](#)
- [Default Code-Point Aliases on page 6632](#)

- [Default Classifiers on page 6633](#)
- [Default Rewrite Rules on page 6636](#)
- [Default Drop Profile on page 6636](#)
- [Default Schedulers on page 6636](#)
- [Default Scheduler Maps on page 6638](#)

## Default Forwarding Classes and Queue Mapping

Table 532 shows the default forwarding class to queue mapping along with the packet drop attribute.

**Table 532: Default Forwarding Classes and Queue Mapping**

| Default Forwarding Class | Description                                                                          | Default Queue Mapping | Packet Drop Attribute |
|--------------------------|--------------------------------------------------------------------------------------|-----------------------|-----------------------|
| best-effort (be)         | Best-effort traffic class (priority 0, IEEE 802.1p code point 000)                   | 0                     | drop                  |
| fcoe                     | Guaranteed delivery for FCoE traffic (priority 3, IEEE 802.1p code point 011)        | 3                     | no-loss               |
| no-loss                  | Guaranteed delivery for TCP no-loss traffic (priority 4, IEEE 802.1p code point 100) | 4                     | no-loss               |
| network-control (nc)     | Network control traffic (priority 7, IEEE 802.1p code point 111)                     | 7                     | drop                  |

Unicast and multidestination (multicast, broadcast, and destination lookup fail) traffic use the same forwarding classes and output queues.

## Default Forwarding Class Sets (Priority Groups)

If you do not explicitly configure forwarding class sets, the system automatically creates a default forwarding class set that contains all of the forwarding classes on the switch. The system assigns 100 percent of the port output bandwidth to the default forwarding class set.

Ingress traffic is classified based on the default classifier settings. The forwarding classes (queues) in the default forwarding class set receive bandwidth based on the default scheduler settings. Forwarding classes that are not part of the default scheduler receive no bandwidth.

The default forwarding class set is transparent. It does not appear in the configuration and is used for Data Center Bridging Capability Exchange (DCBX) protocol advertisement.

## Default Code-Point Aliases

Table 533 shows the default mapping of code-point aliases to IEEE code points.

**Table 533: Default IEEE 802.1 Code-Point Aliases**

| CoS Value Types | Mapping |
|-----------------|---------|
| be              | 000     |
| be1             | 001     |
| ef              | 010     |
| ef1             | 011     |
| af11            | 100     |
| af12            | 101     |
| nc1             | 110     |
| nc2             | 111     |

Table 534 shows the default mapping of code-point aliases to DSCP and DSCP IPv6 code points.

**Table 534: Default DSCP and DCSP IPv6 Code-Point Aliases**

| CoS Value Types | Mapping |
|-----------------|---------|
| ef              | 101110  |
| af11            | 001010  |
| af12            | 001100  |
| af13            | 001110  |
| af21            | 010010  |
| af22            | 010100  |
| af23            | 010110  |
| af31            | 011010  |
| af32            | 011100  |
| af33            | 011110  |



Table 534: Default DSCP and DCSP IPv6 Code-Point Aliases (*continued*)

| CoS Value Types | Mapping |
|-----------------|---------|
| af41            | 100010  |
| af42            | 100100  |
| af43            | 100110  |
| be              | 000000  |
| cs1             | 001000  |
| cs2             | 010000  |
| cs3             | 011000  |
| cs4             | 100000  |
| cs5             | 101000  |
| nc1             | 110000  |
| nc2             | 111000  |

## Default Classifiers

The switch applies default IEEE 802.1 and DSCP classifiers to each interface that does not have explicitly configured classifiers. If you explicitly configure one type of classifier but not other types of classifiers, the system uses only the configured classifier and does not use default classifiers for other types of traffic. The switch applies the default MPLS EXP classifier to a logical interface if you enable the MPLS protocol family on that interface.

There are two different default IEEE 802.1 classifiers, a trusted classifier for ports in trunk mode, and an untrusted classifier for ports in access mode. [Table 535](#) shows the default mapping of IEEE 802.1 code-point values to forwarding classes and loss priorities for ports in trunk mode.

Table 535: Default IEEE 802.1 Classifiers for Ports in Trunk Mode (Trusted Classifier)

| Code Point | Forwarding Class | Loss Priority |
|------------|------------------|---------------|
| be (000)   | best-effort      | low           |
| be1 (001)  | best-effort      | low           |
| ef (010)   | best-effort      | low           |

**Table 535: Default IEEE 802.1 Classifiers for Ports in Trunk Mode (Trusted Classifier) (continued)**

| Code Point | Forwarding Class | Loss Priority |
|------------|------------------|---------------|
| ef1 (011)  | fcoe             | low           |
| af11 (100) | no-loss          | low           |
| af12 (101) | best-effort      | low           |
| nc1 (110)  | network-control  | low           |
| nc2 (111)  | network-control  | low           |

Table 536 shows the default mapping of IEEE 802.1p code-point values to forwarding classes and loss priorities for ports in access mode (all incoming traffic is mapped to best-effort forwarding classes).

**Table 536: Default IEEE 802.1 Classifiers for Ports in Access Mode (Untrusted Classifier)**

| Code Point | Forwarding Class | Loss Priority |
|------------|------------------|---------------|
| 000        | best-effort      | low           |
| 001        | best-effort      | low           |
| 010        | best-effort      | low           |
| 011        | best-effort      | low           |
| 100        | best-effort      | low           |
| 101        | best-effort      | low           |
| 110        | best-effort      | low           |
| 111        | best-effort      | low           |

Table 537 shows the default mapping of DSCP code-point values to forwarding classes and loss priorities for DSCP IP and DCSP IPv6.

**Table 537: Default DSCP IP and IPv6 Classifiers**

| Code Point    | Forwarding Class | Loss Priority |
|---------------|------------------|---------------|
| ef (101110)   | best-effort      | low           |
| af11 (001010) | best-effort      | low           |

**Table 537: Default DSCP IP and IPv6 Classifiers (*continued*)**

| Code Point    | Forwarding Class | Loss Priority |
|---------------|------------------|---------------|
| af12 (001100) | best-effort      | low           |
| af13 (001110) | best-effort      | low           |
| af21 (010010) | best-effort      | low           |
| af22 (010100) | best-effort      | low           |
| af23 (010110) | best-effort      | low           |
| af31 (011010) | best-effort      | low           |
| af32 (011100) | best-effort      | low           |
| af33 (011110) | best-effort      | low           |
| af41 (100010) | best-effort      | low           |
| af42 (100100) | best-effort      | low           |
| af43 (100110) | best-effort      | low           |
| be (000000)   | best-effort      | low           |
| cs1 (001000)  | best-effort      | low           |
| cs2 (010000)  | best-effort      | low           |
| cs3 (011000)  | best-effort      | low           |
| cs4 (100000)  | best-effort      | low           |
| cs5 (101000)  | best-effort      | low           |
| nc1 (110000)  | network-control  | low           |
| nc2 (111000)  | network-control  | low           |

Table 538 shows the default mapping of MPLS EXP code-point values to forwarding classes and loss priorities.

**Table 538: Default EXP Classifiers**

| Code Point | Forwarding Class | Loss Priority |
|------------|------------------|---------------|
| 000        | best-effort      | low           |

**Table 538: Default EXP Classifiers (*continued*)**

| Code Point | Forwarding Class     | Loss Priority |
|------------|----------------------|---------------|
| 001        | best-effort          | high          |
| 010        | expedited-forwarding | low           |
| 011        | expedited-forwarding | high          |
| 100        | assured-forwarding   | low           |
| 101        | assured-forwarding   | high          |
| 110        | network-control      | low           |
| 111        | network-control      | high          |

## Default Rewrite Rules

There are no default rewrite rules. If you do not explicitly configure rewrite rules, the switch does not reclassify egress traffic.

## Default Drop Profile

Table 539 shows the default drop profile configuration.

**Table 539: Default Drop Profile**

| Fill Level | Drop Probability |
|------------|------------------|
| 100        | 100              |

## Default Schedulers

Table 540 shows the default scheduler configuration.

**Table 540: Default Schedulers**

| Default Scheduler and Queue Number               | Transmit Rate (Guaranteed Minimum Bandwidth) | Rate Shaping (Maximum Bandwidth) | Excess Bandwidth Sharing | Priority | Buffer Size |
|--------------------------------------------------|----------------------------------------------|----------------------------------|--------------------------|----------|-------------|
| best-effort forwarding class scheduler (queue 0) | 15%                                          | None                             | 15%                      | low      | 15%         |
| fcoe forwarding class scheduler (queue 3)        | 35%                                          | None                             | 35%                      | low      | 35%         |
| no-loss forwarding class scheduler (queue 4)     | 35%                                          | None                             | 35%                      | low      | 35%         |

Table 540: Default Schedulers (*continued*)

| Default Scheduler and Queue Number                   | Transmit Rate (Guaranteed Minimum Bandwidth) | Rate Shaping (Maximum Bandwidth) | Excess Bandwidth Sharing | Priority | Buffer Size |
|------------------------------------------------------|----------------------------------------------|----------------------------------|--------------------------|----------|-------------|
| network-control forwarding class scheduler (queue 7) | 15%                                          | None                             | 15%                      | low      | 15%         |



**NOTE:** By default, the minimum guaranteed bandwidth (transmit rate) also determines the amount of excess (extra) bandwidth that the queue can share. Extra bandwidth is allocated to queues in proportion to the transmit rate of each queue. You can use the `excess-rate` statement to override the default setting and configure the excess bandwidth percentage independently of the transmit rate.

By default, only the four default schedulers shown in Table 540 have traffic mapped to them. Only the forwarding classes and queues associated with the default schedulers receive default bandwidth, based on the default scheduler transmit rate. (You can configure schedulers and forwarding classes to allocate bandwidth to other queues or to change the default bandwidth of a default queue.) If a forwarding class does not transport traffic, the bandwidth allocated to that forwarding class is available to other forwarding classes. Unicast and multdestination (multicast, broadcast, and destination lookup fail) traffic use the same forwarding classes and output queues.

Default scheduling is port scheduling.

Default hierarchical scheduling, known as enhanced transmission selection (ETS, defined in IEEE 802.1Qaz), allocates the total port bandwidth to the four default forwarding classes served by the four default schedulers, as defined by the four default schedulers. The result is the same as direct port scheduling. Configuring hierarchical port scheduling, however, enables you to group forwarding classes that carry similar types of traffic into forwarding class sets (also called priority groups), and to assign port bandwidth to each forwarding class set. The port bandwidth assigned to the forwarding class set is then assigned to the forwarding classes within the forwarding class set. This hierarchy enables you to control port bandwidth allocation with greater granularity, and enables hierarchical sharing of extra bandwidth to better utilize link bandwidth.

Default scheduling uses weighted round-robin (WRR) scheduling. Each queue receives a portion (weight) of the total available interface bandwidth. The scheduling weight is based on the transmit rate of the default scheduler for that queue. For example, queue 7 receives a default scheduling weight of 15 percent of the available bandwidth, and queue 4 receives a default scheduling weight of 35 percent of the available bandwidth. Queues are mapped to forwarding classes (for example, queue 7 is mapped to the network-control forwarding class and queue 4 is mapped to the no-loss forwarding class), so forwarding classes receive the default bandwidth for the queues to which they are mapped. Unused bandwidth is shared with other default queues.

If you want non-default (unconfigured) queues to forward traffic, you should explicitly map traffic to those queues (configure the forwarding classes and queue mapping) and create schedulers to allocate bandwidth to those queues. By default, queues 1, 2, 5, and 6 are unconfigured. Unconfigured queues have a default scheduling weight of 1 so that they can receive a small amount of bandwidth in case they need to forward traffic.

If you map traffic to an unconfigured queue and do not schedule port resources for the queue (configure a scheduler, map it to the forwarding class that is mapped to the queue, and apply the scheduler mapping to the port), the queue receives only the amount of excess bandwidth proportional to its default weight (1). The actual amount of bandwidth an unconfigured queue gets depends on how much bandwidth the other queues on the port are using.

If the other queues use less than their allocated amount of bandwidth, the unconfigured queues can share the unused bandwidth. Configured queues have higher priority for bandwidth than unconfigured queues, so if a configured queue needs more bandwidth, then less bandwidth is available for unconfigured queues. Unconfigured queues always receive a minimum amount of bandwidth based on their scheduling weight (1). If you map traffic to an unconfigured queue, to allocate bandwidth to that queue, configure a scheduler for the forwarding class that is mapped to the queue and apply it to the port.

## Default Scheduler Maps

Table 541 shows the default mapping of forwarding classes to schedulers.

**Table 541: Default Scheduler Maps**

| Forwarding Class | Scheduler                         |
|------------------|-----------------------------------|
| best-effort      | Default BE scheduler              |
| fcoe             | Default FCoE scheduler            |
| no-loss          | No-loss scheduler                 |
| network-control  | Default network-control scheduler |

### Related Documentation

- [Overview of Junos OS CoS on page 6608](#)
- [CoS Support on QFX Series Switches, EX4600 Switches, and QFabric Systems on page 6610](#)
- [Understanding Junos CoS Components on page 6621](#)
- [Understanding Default CoS Scheduling and Classification on page 6666](#)
- [Understanding CoS Classifiers on page 6650](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 6672](#)
- [Understanding CoS Code-Point Aliases on page 6685](#)
- [Understanding CoS Forwarding Classes on page 6688](#)
- [Understanding CoS Rewrite Rules on page 6701](#)

- [Understanding CoS Port Schedulers on QFX Switches on page 6729](#)
- [Understanding CoS WRED Drop Profiles on page 6809](#)

## CoS Inputs and Outputs Overview

Some CoS components map one set of values to another set of values. Each mapping contains one or more inputs and one or more outputs. When you configure a mapping, you set the outputs for a given set of inputs, as shown in [Table 542](#).

**Table 542: CoS Mappings—Inputs and Outputs**

| CoS Mappings                                             | Inputs                                                              | Outputs                                                             | Comments                                                                                                                                                                                                     |
|----------------------------------------------------------|---------------------------------------------------------------------|---------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <a href="#">classifiers</a>                              | <a href="#">code-points</a>                                         | <a href="#">forwarding-class</a> ,<br><a href="#">loss-priority</a> | The map sets the forwarding class and packet loss priority (PLP) for a specific set of code points.                                                                                                          |
| <a href="#">drop-profile-map</a>                         | <a href="#">loss-priority</a> , <a href="#">protocol</a>            | <a href="#">drop-profile</a>                                        | The map sets the drop profile for a specific PLP and protocol type.                                                                                                                                          |
| <a href="#">rewrite-rules</a>                            | <a href="#">loss-priority</a> ,<br><a href="#">forwarding-class</a> | <a href="#">code-points</a>                                         | The map sets the code points for a specific forwarding class and PLP.                                                                                                                                        |
| <a href="#">rewrite-value (Fibre Channel Interfaces)</a> | <a href="#">forwarding-class</a>                                    | <a href="#">code-point</a>                                          | (Systems that support native Fibre Channel interfaces only) The map sets the code point for the forwarding class specified in the fixed classifier attached to the native Fibre Channel (NP_Port) interface. |

**Related Documentation** • [Understanding CoS Packet Flow on page 6628](#)

## Overview of Policers

A switch polices traffic by limiting the input or output transmission rate of a class of traffic according to user-defined criteria. Policing (or rate-limiting) traffic allows you to control the maximum rate of traffic sent or received on an interface and to provide multiple priority levels or classes of service.


- [Policer Overview on page 6640](#)
- [Policer Types on page 6640](#)
- [Policer Actions on page 6641](#)
- [Policer Colors on page 6642](#)
- [Filter-Specific Policers on page 6642](#)
- [Suggested Naming Convention for Policers on page 6642](#)
- [Policer Counters on page 6643](#)
- [Policer Algorithms on page 6643](#)
- [How Many Policers Are Supported? on page 6643](#)
- [Policers Can Limit Egress Firewall Filters on page 6644](#)

## Policer Overview

You use policers to apply limits to traffic flow and set consequences for packets that exceed these limits—usually applying a higher loss priority—so that if packets encounter downstream congestion, they can be discarded first. Policers apply only to unicast packets.

Policers provide two functions: metering and marking. A policer meters (measures) each packet against traffic rates and burst sizes that you configure. It then passes the packet and the metering result to the marker, which assigns a packet loss priority that corresponds to the metering result. [Figure 198](#) illustrates this process.

**Figure 211: Flow of Tricolor Marking Policer Operation**



After you name and configure a policer, you can use it by specifying it as an action in one or more firewall filters.

## Policer Types

A switch supports three types of policers:

- **Single-rate two-color marker**—A two-color policer (or “policer” when used without qualification) meters the traffic stream and classifies packets into two categories of packet loss priority (PLP) according to a configured bandwidth and burst-size limit. You can mark packets that exceed the bandwidth and burst-size limit with a specified PLP or simply discard them.

You can specify this type of policer in an ingress or egress firewall.



**NOTE:** A two-color policer is most useful for metering traffic at the port (physical interface) level.

- **Single-rate three-color marker**—This type of policer is defined in RFC 2697, *A Single Rate Three Color Marker*, as part of an assured forwarding (AF) per-hop-behavior (PHB) classification system for a Differentiated Services (DiffServ) environment. This type of policer meters traffic based on one rate—the configured committed information rate (CIR) as well as the committed burst size (CBS) and the excess burst size (EBS). The CIR specifies the average rate at which bits are admitted to the switch. The CBS specifies the usual burst size in bytes and the EBS specifies the maximum burst size in bytes. The EBS must be greater than or equal to the CBS, and neither can be 0.

You can specify this type of policer in an ingress or egress firewall.





**NOTE:** A single-rate three-color marker (TCM) is most useful when a service is structured according to packet length and not peak arrival rate.

- Two-rate three-color marker—This type of policer is defined in RFC 2698, *A Two Rate Three Color Marker*, as part of an assured forwarding per-hop-behavior classification system for a Differentiated Services environment. This type of policer meters traffic based on two rates—the CIR and peak information rate (PIR) along with their associated burst sizes, the CBS and peak burst size (PBS). The PIR specifies the maximum rate at which bits are admitted to the network and must be greater than or equal to the CIR.

You can specify this type of policer in an ingress or egress firewall.



**NOTE:** A two-rate three-color policer is most useful when a service is structured according to arrival rates and not necessarily packet length.

See [Table 479](#) for information about how metering results are applied for each of these policer types.

## Policer Actions

Policer actions are implicit or explicit and vary by policer type. *Implicit* means that Junos OS assigns the loss priority automatically. [Table 479](#) describes the policer actions.

**Table 543: Policer Actions**

| Policer                 | Marking                        | Implicit Action                  | Configurable Action |
|-------------------------|--------------------------------|----------------------------------|---------------------|
| Single-rate two-color   | Green (conforming)             | Assign low loss priority         | None                |
|                         | Red (nonconforming)            | None                             | Discard             |
| Single-rate three-color | Green (conforming)             | Assign low loss priority         | None                |
|                         | Yellow (above the CIR and CBS) | Assign medium-high loss priority | None                |
|                         | Red (above the EBS)            | Assign high loss priority        | Discard             |
| Two-rate three-color    | Green (conforming)             | Assign low loss priority         | None                |
|                         | Yellow (above the CIR and CBS) | Assign medium-high loss priority | None                |
|                         | Red (above the PIR and PBS)    | Assign high loss priority        | Discard             |



**NOTE:** If you specify a policer in an egress firewall filter, the only supported action is **discard**.

## Policer Colors

Single-rate and two-rate three-color policers can operate in two modes:

- **Color-blind**—In color-blind mode, the three-color policer assumes that all packets examined have not been previously marked or metered. In other words, the three-color policer is “blind” to any previous coloring a packet might have had.
- **Color-aware**—In color-aware mode, the three-color policer assumes that all packets examined have been previously marked or metered. In other words, the three-color policer is “aware” of the previous coloring a packet might have had. In color-aware mode, the three-color policer can increase the PLP of a packet but cannot decrease it. For example, if a color-aware three-color policer meters a packet with a medium PLP marking, it can raise the PLP level to high but cannot reduce the PLP level to low.

## Filter-Specific Policers

You can configure policers to be filter-specific, which means that Junos OS creates only one policer instance regardless of how many times the policer is referenced. When you do this on some QFX switches, rate limiting is applied in aggregate, so if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms, the total bandwidth allowed by the filter is 1 Gbps. However, the behavior of a filter-specific policer is affected by how the firewall filter terms that reference the policer are stored in TCAM. If you create a filter-specific policer and reference it in multiple firewall filter terms, the policer allows more traffic than expected if the terms are stored in different TCAM slices. For example, if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms that are stored in three separate memory slices, the total bandwidth allowed by the filter is 3 Gbps, not 1 Gbps. (This behavior does not occur in QFX10000 switches.)

To prevent this unexpected behavior from occurring, use the information about TCAM slices presented in [“Planning the Number of Firewall Filters to Create” on page 5973](#) to organize your configuration file so that all the firewall filter terms that reference a given filter-specific policer are stored in the same TCAM slice.

## Suggested Naming Convention for Policers

We recommend that you use the naming convention **policertypeTCM#-color type** when configuring three-color policers and **policer#** when configuring two-color policers. TCM stands for three-color marker. Because policers can be numerous and must be applied correctly to work, a simple naming convention makes it easier to apply the policers properly. For example, the first single-rate, color-aware three-color policer configured would be named **srTCM1-ca**. The second two-rate, color-blind three-color configured would be named **trTCM2-cb**. The elements of this naming convention are explained below:

- sr (single-rate)
- tr (two-rate)
- TCM (tricolor marking)
- 1 or 2 (number of marker)
- ca (color-aware)
- cb (color-blind)

## Policer Counters

On some QFX switches, each policer that you configure includes an implicit counter that counts the number of packets that exceed the rate limits that are specified for the policer. If you use the same policer in multiple terms—either within the same filter or in different filters—the implicit counter counts all the packets that are policed in all of these terms and provides the total amount. (This does not apply to QFX10000 switches.) If you want to obtain separate packet counts for each term on an affected switch, use these options:

- Configure a unique policer for each term.
- Configure only one policer, but use a unique, explicit counter in each term.

## Policer Algorithms

Policing uses the *token-bucket algorithm*, which enforces a limit on average bandwidth while allowing bursts up to a specified maximum value. It offers more flexibility than the *leaky bucket algorithm* in allowing a certain amount of bursty traffic before it starts discarding packets.



**NOTE:** In an environment of light bursty traffic, QFX5200 might not replicate all multicast packets to two or more downstream interfaces. This occurs only at a line rate burst—if traffic is consistent, the issue does not occur. In addition, the issue occurs only when packet size increases beyond 6k in a one gigabit traffic flow.

## How Many Policers Are Supported?

QFX10000 switches support 8K policers (all policer types). QFX5100 and QFX5200 switches support 1535 ingress policers and 1024 egress policers (assuming one policer per firewall filter term).

QFX3500 and QFX3600 standalone switches and QFabric Node devices support the following numbers of policers (assuming one policer per firewall filter term):

- Two-color policers used in ingress firewall filters: 767
- Three-color policers used in ingress firewall filters: 767

- Two-color policers used in egress firewall filters: 1022
- Three-color policers used in egress firewall filters: 512

## Policers Can Limit Egress Firewall Filters

On some switches, the number of egress policers that you configure can affect the total number of allowed egress firewall filters. (This issue does not affect QFX10000 switches.) On the affected switches, every policer has two implicit counters that consume two entries in a 1024-entry TCAM that is used for counters, including counters that are configured as action modifiers in firewall filter terms. (Policers consume two entries because one is used for green packets and one is used for nongreen packets regardless of policer type.) If the TCAM becomes full, you cannot commit any more egress firewall filters that have terms with counters. For example, if you configure and commit 512 egress policers (two-color, three-color, or a combination of both policer types), all of the memory entries for counters are used up. If later in your configuration file you insert additional egress firewall filters with terms that also include counters, *none* of the terms in those filters are committed because there is no available memory space for the counters.

Here are some additional examples:

- Assume that you configure egress filters that include a total of 512 policers and no counters. Later in your configuration file you include another egress filter with 10 terms, 1 of which has a counter action modifier. None of the terms in this filter are committed because there is not enough TCAM space for the counter.
- Assume that you configure egress filters that include a total of 500 policers, so 1000 TCAM entries are occupied. Later in your configuration file you include the following two egress filters:
  - Filter A with 20 terms and 20 counters. All the terms in this filter are committed because there is enough TCAM space for all the counters.
  - Filter B comes after Filter A and has five terms and five counters. *None* of the terms in this filter are committed because there is not enough memory space for *all* the counters. (Five TCAM entries are required but only four are available.)

You can prevent this problem by ensuring that egress firewall filter terms with counter actions are placed earlier in your configuration file than terms that include policers. In this circumstance, Junos OS commits policers even if there is not enough TCAM space for the implicit counters. For example, assume the following:

- You have 1024 egress firewall filter terms with counter actions.
- Later in your configuration file you have an egress filter with 10 terms. None of the terms have counters but one has a policer action modifier.

You can successfully commit the filter with 10 terms even though there is not enough TCAM space for the implicit counters of the policer. The policer is committed without the counters.

### Related Documentation

- [Understanding Color-Blind Mode for Single-Rate Tricolor Marking on page 6005](#)

- [Understanding Color-Blind Mode for Two-Rate Tricolor Marking on page 6008](#)
- [Understanding Color-Aware Mode for Single-Rate Tricolor Marking on page 6006](#)
- [Understanding Color-Aware Mode for Two-Rate Tricolor Marking on page 6008](#)
- [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 6017](#)



## PART 91

# Classifying and Rewriting Traffic

- [Using Classifiers, Forwarding Classes, and Rewrite Rules on page 6649](#)





# Using Classifiers, Forwarding Classes, and Rewrite Rules

- [Understanding CoS Classifiers on page 6650](#)
- [Defining CoS BA Classifiers \(DSCP, DSCP IPv6, IEEE 802.1p\) on page 6656](#)
- [Example: Configuring Classifiers on page 6657](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 6660](#)
- [Configuring a Global MPLS EXP Classifier on page 6663](#)
- [Configuring Rewrite Rules for MPLS EXP Classifiers on page 6664](#)
- [Understanding Host Inbound Traffic Classification on page 6665](#)
- [Understanding Default CoS Scheduling and Classification on page 6666](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 6672](#)
- [Understanding CoS Code-Point Aliases on page 6685](#)
- [Defining CoS Code-Point Aliases on page 6687](#)
- [Understanding CoS Forwarding Classes on page 6688](#)
- [Defining CoS Forwarding Classes on page 6691](#)
- [Example: Configuring Forwarding Classes on page 6692](#)
- [Understanding CoS Forwarding Class Sets \(Priority Groups\) on page 6695](#)
- [Defining CoS Forwarding Class Sets on page 6696](#)
- [Example: Configuring Forwarding Class Sets on page 6697](#)
- [Understanding CoS Rewrite Rules on page 6701](#)
- [Defining CoS Rewrite Rules on page 6704](#)
- [Troubleshooting an Unexpected Rewrite Value on page 6705](#)

## Understanding CoS Classifiers

---

Packet classification maps incoming packets to a particular class-of-service (CoS) servicing level. Classifiers map packets to a forwarding class and a loss priority, and they assign packets to output queues based on the forwarding class. There are three general types of classifiers:

- Behavior aggregate (BA) classifiers—DSCP and DSCP IPv6 classify IP and IPv6 traffic, EXP classifies MPLS traffic, and IEEE 802.1p classifies all other traffic. (Although this topic covers EXP classifiers, fsee [“Understanding CoS MPLS EXP Classifiers and Rewrite Rules” on page 4959](#) for more details. EXP classifiers are applied only on **family mpls** interfaces.)
- Fixed classifiers—Fixed classifiers classify all ingress traffic on a physical interface into one forwarding class, regardless of the CoS bits in the packet header.
- Multifield (MF) classifiers—MF classifiers classify traffic based on more than one field in the packet header and take precedence over BA and fixed classifiers.

Classifiers assign incoming unicast and multideestination (multicast, broadcast, and destination lookup fail) traffic to forwarding classes, so that different classes of traffic can receive different treatment. Classification is based on CoS bits, DSCP bits, EXP bits, a forwarding class (fixed classifier), or packet headers (multifield classifiers). Each classifier assigns all incoming traffic that matches the classifier configuration to a particular forwarding class. A classifier can assign both unicast and multideestination traffic to the same forwarding class.

- [Interfaces and Output Queues on page 6650](#)
- [Classifier Support by Type on page 6651](#)
- [Behavior Aggregate Classifiers on page 6651](#)
- [Fixed Classifiers on Ethernet Interfaces on page 6654](#)
- [Multifield Classifiers on page 6654](#)
- [MPLS EXP Classifiers on page 6655](#)
- [Packet Classification for IRB Interfaces on page 6655](#)

## Interfaces and Output Queues

On Gigabit Ethernet interfaces, 10-Gigabit Ethernet interfaces, and link aggregation (LAG) interfaces, you can apply classifiers to Layer 2 logical interface unit 0 (but not to other logical interfaces), and to Layer 3 logical interfaces (you can apply different classifiers to different Layer 3 logical interfaces). You cannot apply classifiers to physical interfaces. [“Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces” on page 6672](#) describes the interaction between classifiers and interfaces in greater detail.

You can configure both a BA classifier and an MF classifier on an interface. If you do this, the BA classification is performed first, and then the MF classification is performed. If the two classification results conflict, the MF classification result overrides the BA classification result.

You cannot configure a fixed classifier and a BA classifier on the same interface.

You can configure either a DSCP or a DSCP IPv6 classifier and also an IEEE 802.1p classifier on the same interface. IP traffic uses the DSCP or DSCP IPv6 classifier. If you configure an interface as **family mpls**, then the interface uses the default MPLS EXP classifier. If you configure an MPLS EXP classifier, then all MPLS traffic on the switch uses the global EXP classifier. All other traffic uses the IEEE classifier. You can configure up to 64 EXP classifiers with up to 8 entries per classifier (one entry for each forwarding class) and apply them to logical interfaces. After you configure an MPLS EXP classifier, you can configure it as the global EXP classifier by including the EXP classifier at the **[edit class-of-service system-defaults classifiers exp]** hierarchy level. All switch interfaces that are configured as **family mpls** use either the default EXP classifier or the global EXP classifier specified in this configuration statement to classify MPLS traffic.

You can apply classifiers to one or more interfaces.

## Classifier Support by Type

You can configure enough classifiers to handle most, if not all, network scenarios.

[Table 544](#) shows how many of each type of classifiers you can configure, and how many entries you can configure per classifier.

**Table 544: Classifier Support by Classifier Type**

| Classifier Type       | Default Classifier Name                      | Maximum Number of Classifiers | Maximum Number of Entries per Classifier |
|-----------------------|----------------------------------------------|-------------------------------|------------------------------------------|
| IEEE 802.1p (Layer 2) | ieee8021p-default (for ports in trunk mode)  | 64                            | 16                                       |
|                       | ieee8021p-untrust (for ports in access mode) |                               |                                          |
| DSCP (Layer 3)        | dscp-default                                 | 64                            | 64                                       |
| DSCP IPv6 (Layer 3)   | dscp-ipv6-default                            | 64                            | 64                                       |
| EXP (MPLS)            | exp-default                                  | 64                            | 8                                        |
| Fixed                 | There is no default fixed classifier         | 8                             | 16                                       |

The number of fixed classifiers supported (8) equals the number of supported forwarding classes (fixed classifiers assign all incoming traffic on an interface to one forwarding class).

## Behavior Aggregate Classifiers

Behavior aggregate classifiers map a class-of-service (CoS) value to a forwarding class and loss priority. The forwarding class determines the output queue. A scheduler uses the loss priority to control packet discard during periods of congestion by associating different drop profiles with different loss priorities.

The switch supports three types of BA classifiers:

- Differentiated Services code point (DSCP) for IP DiffServ (IP and IPv6)
- IEEE 802.1p CoS bits
- MPLS EXP (applies only to interfaces configured as **family mpls**)

BA classifiers are based on fixed-length fields, which makes them computationally more efficient than MF classifiers. Therefore, core devices, which handle high traffic volumes, are normally configured to perform BA classification.

### Default Behavior Aggregate Classification

Juniper Networks Junos OS automatically assigns implicit default classifiers to all logical interfaces based on the type of interface. [Table 545](#) lists different types of interfaces and the corresponding implicit default BA classifiers.

**Table 545: Default BA Classification**

| Type of Interface                | Default BA Classification                       |
|----------------------------------|-------------------------------------------------|
| Layer 2 interface in trunk mode  | <b>ieee8021p-default</b>                        |
| Layer 2 interface in access mode | <b>ieee8021p-untrusted</b>                      |
| Layer 3 interface                | <b>dscp-default</b><br><b>dscp-ipv6-default</b> |
| MPLS interface                   | <b>exp-default</b>                              |



**NOTE:** Default BA classifiers assign traffic only to the **best-effort**, **fcoe**, **no-loss**, and **network-control** forwarding classes.

When you explicitly associate a classifier with a logical interface, you override the default classifier with the explicit classifier.



**NOTE:** You can apply only one DSCP and one IEEE 802.1p classifier to a Layer 2 interface. If both types of classifiers are present, DSCP classifiers take precedence over IEEE 802.1p classifiers. If you configure an EXP classifier and apply it on interfaces configured as **family mpls**, then MPLS traffic uses the EXP classifier on those interfaces.

### Importing a Classifier

You can use any existing classifier, including the default classifiers, as the basis for defining a new classifier. You accomplish this using the **import** statement.

The imported classifier is used as a template and is not modified. The modifications you make become part of a new classifier (and a new template) identified by the name of the new classifier. Whenever you commit a configuration that assigns a new forwarding class-name and loss-priority value to a code-point alias or set of bits, it replaces the old entry in the new classifier template. As a result, you must explicitly specify every CoS value in every packet classification that requires modification.

### PFC Priorities

The eight IEEE 802.1p code points correspond to the eight priorities that priority-based flow control (PFC) uses to differentiate traffic classes for lossless transport. When you map a forwarding class (which maps to an output queue) to an IEEE 802.1p CoS value, the IEEE 802.1p CoS value identifies the PFC priority.

Although you can map a priority to any output queue (by mapping the IEEE 802.1p code point value to a forwarding class), we recommend that the priority and the forwarding class match in a one-to-one correspondence. For example, priority 0 is assigned to queue 0, priority 1 is assigned to queue 1, and so on, as shown in [Table 546](#). A one-to-one correspondence of queue and priority numbers makes it easier to configure and maintain the mapping of forwarding classes to priorities and queues.

**Table 546: Default IEEE 802.1p Code Point to PFC Priority, Output Queue, and Forwarding Class Mapping**

| IEEE 802.1p Code Point | PFC Priority | Output Queue | Forwarding Class and Packet Drop Attribute |
|------------------------|--------------|--------------|--------------------------------------------|
| 000                    | 0            | 0            | best-effort (drop)                         |
| 001                    | 1            | 1            | best-effort (drop)                         |
| 010                    | 2            | 2            | best-effort (drop)                         |
| 011                    | 3            | 3            | fcoe (no-loss)                             |
| 100                    | 4            | 4            | no-loss (no-loss)                          |
| 101                    | 5            | 5            | best-effort (drop)                         |
| 110                    | 6            | 6            | network-control (drop)                     |
| 111                    | 7            | 7            | network-control (drop)                     |



**NOTE:** By convention, deployments with converged server access typically use IEEE 802.1p priority 3 (011) for FCoE traffic. The default mapping of the fcoe forwarding class is to queue 3. Apply priority-based flow control (PFC) to the entire FCoE data path to configure the end-to-end lossless behavior that FCoE requires. We recommend that you use priority 3 for FCoE traffic unless your network architecture requires that you use a different priority.

## Fixed Classifiers on Ethernet Interfaces

Fixed classifiers map all traffic on a physical interface to a forwarding class and a loss priority (as opposed to BA classifiers, which map traffic into multiple different forwarding classes based on the IEEE 802.1p CoS bits field value in the VLAN header or the DSCP field value in the type-of-service bits in the packet IP header). Each forwarding class maps to an output queue. However, when you use a fixed classifier, regardless of the CoS or DSCP bits, all Incoming traffic is classified into the forwarding class specified in the fixed classifier. A scheduler uses the loss priority to control packet discard during periods of congestion by associating different drop profiles with different loss priorities.

You cannot configure a fixed classifier and a DSCP or IEEE 802.1p BA classifier on the same interface. If you configure a fixed classifier on an interface, you cannot configure a DSCP or an IEEE classifier on that interface. If you configure a DSCP classifier, an IEEE classifier, or both classifiers on an interface, you cannot configure a fixed classifier on that interface.



**NOTE:** You can configure both an EXP classifier for MPLS traffic and a fixed classifier on the same interface. When both an EXP classifier and a fixed classifier are applied to an interface, MPLS traffic on interfaces configured as `family mpls` uses the EXP classifier, and all other traffic uses the fixed classifier.

To switch from a fixed classifier to a BA classifier, or to switch from a BA classifier to a fixed classifier, deactivate the existing classifier attachment on the interface, and then attach the new classifier to the interface.



**NOTE:** If you configure a fixed classifier that classifies all incoming traffic into the `fcoe` forwarding class (or any forwarding class designed to handle FCoE traffic), you must ensure that all traffic that enters the interface is FCoE traffic and is tagged with the FCoE IEEE 802.1p code point (priority).

## Multifield Classifiers

Multifield classifiers examine multiple fields in a packet such as source and destination addresses and source and destination port numbers of the packet. With MF classifiers, you set the forwarding class and loss priority of a packet based on firewall filter rules.

MF classification is normally performed at the network edge because of the general lack of DiffServ code point (DSCP) support in end-user applications. On a switch at the edge of a network, an MF classifier provides the filtering functionality that scans through a variety of packet fields to determine the forwarding class for a packet. Typically, a classifier performs matching operations on the selected fields against a configured value.

## MPLS EXP Classifiers

You can configure up to 64 EXP classifiers for MPLS traffic and apply them to **family mpls** interfaces. On **family mpls** interfaces, if a fixed classifier is present on the interface, the EXP classifier overrides the fixed classifier for MPLS traffic only.

You can use the default MPLS EXP classifier or you can configure an EXP classifier and apply it globally to all interfaces that are configured as **family mpls** by including it in the **[edit class-of-service system-defaults classifiers exp]** hierarchy level. On **family mpls** interfaces, if a fixed classifier is present on the interface, the EXP classifier overrides the fixed classifier for MPLS traffic only.

Because the EXP classifier is global, you cannot configure some ports to use a fixed IEEE 802.1p classifier for MPLS traffic on some interfaces and the global EXP classifier for MPLS traffic on other interfaces. When you configure a global EXP classifier, all MPLS traffic on all interfaces uses the EXP classifier, even interfaces that have a fixed classifier.

For details about EXP classifiers, see [“Understanding CoS MPLS EXP Classifiers and Rewrite Rules” on page 4959](#). EXP classifiers are applied only on **family mpls** interfaces.)

## Packet Classification for IRB Interfaces

You cannot apply classifiers directly to integrated routing and bridging (IRB) interfaces because the members of IRBs are VLANs, not ports. However, you can apply classifiers to the VLAN port members of an IRB interface. You can also apply MF classifiers to IRBs.

### Related Documentation

- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4959](#)
- [Understanding CoS Packet Flow on page 6628](#)
- [Understanding Default CoS Settings on page 6630](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 6672](#)
- [Example: Configuring Classifiers on page 6657](#)
- [Configuring a Global MPLS EXP Classifier on page 4977](#)

## Defining CoS BA Classifiers (DSCP, DSCP IPv6, IEEE 802.1p)

Packet classification associates incoming packets with a particular CoS servicing level. Behavior aggregate (BA) classifiers examine the Differentiated Services code point (DSCP or DSCP IPv6) value, the IEEE 802.1p CoS value, or the MPLS EXP value in the packet header to determine the CoS settings applied to the packet. (See [“Configuring a Global MPLS EXP Classifier” on page 4977](#) to learn how to define EXP classifiers for MPLS traffic.) BA classifiers allow you to set the forwarding class and loss priority of a packet based on the incoming CoS value.

Unicast and multideestination (multicast, broadcast, and destination lookup fail) traffic use the same classifiers and forwarding classes.

To configure a DSCP, DSCP IPv6, or IEEE 802.1p BA classifier using the CLI:

1. Create a BA classifier:

- To create a DSCP, DSCP IPv6, or IEEE 802.1p BA classifier based on the default classifier, import the default DSCP, DSCP IPv6, or IEEE 802.1p classifier and associate it with a forwarding class, a loss priority, and a code point:

```
[edit class-of-service classifiers]
user@switch# set (dscp | dscp-ipv6 | ieee-802.1) classifier-name import default
forwarding-class forwarding-class-name loss-priority level code-points [aliases]
[bit-patterns]
```

- To create a BA classifier that is not based on the default classifier, create a DSCP, DSCP IPv6, or IEEE 802.1p classifier and associate it with a forwarding class, a loss priority, and a code point:

```
[edit class-of-service classifiers]
user@switch# set (dscp | dscp-ipv6 | ieee-802.1) classifier-name forwarding-class
forwarding-class-name loss-priority level code-points [aliases] [bit-patterns]
```

2. Apply the classifier to a specific Ethernet interface or to all Ethernet interfaces on the switch.

- To apply the classifier to a specific interface:

```
[edit class-of-service interfaces]
user@switch# set interface-name unit unit classifiers (dscp | dscp-ipv6 | ieee-802.1)
classifier-name
```

- To apply the classifier to all Ethernet interfaces on the switch, use wildcards for the interface name and the logical interface (unit) number:

```
[edit class-of-service interfaces]
user@switch# set xe-* unit * classifiers (dscp | dscp-ipv6 | ieee-802.1) classifier-name
```

### Related Documentation

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 6771](#)
- [Example: Configuring Classifiers on page 6657](#)
- [Configuring a Global MPLS EXP Classifier on page 4977](#)
- [Monitoring CoS Classifiers on page 7231](#)
- [Understanding CoS Classifiers on page 6650](#)



- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4959](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 6672](#)

## Example: Configuring Classifiers

Packet classification associates incoming packets with a particular CoS servicing level. Classifiers associate packets with a forwarding class and loss priority and assign packets to output queues based on the associated forwarding class. You apply classifiers to ingress interfaces.

- [Requirements on page 6657](#)
- [Overview on page 6657](#)
- [Configuring Classifiers on page 6658](#)
- [Verification on page 6658](#)

### Requirements

This example uses the following hardware and software components:

- One switch.
- Junos OS Release 15.1X53-D10 or later for the QFX Series.

### Overview

Junos OS supports three general types of classifiers:

- Behavior aggregate or CoS value traffic classifiers—Examine the CoS value in the packet header. The value in this single field determines the CoS settings applied to the packet. BA classifiers allow you to set the forwarding class and loss priority of a packet based on the Differentiated Services code point (DSCP or DSCP IPv6) value, IEEE 802.1p value, or MPLS EXP value. (EXP classifiers can be applied only to **family mpls** interfaces.)
- Fixed classifiers. Fixed classifiers classify all ingress traffic on a physical interface into one forwarding class, regardless of the CoS bits in the VLAN header or the DSCP bits in the IP packet header.
- Multifield traffic classifiers—Examine multiple fields in the packet, such as source and destination addresses and source and destination port numbers of the packet. With multifield classifiers, you set the forwarding class and loss priority of a packet based on firewall filter rules.

This example describes how to configure a BA classifier called **ba-classifier** as the default IEEE 802.1 mapping of incoming traffic to forwarding classes, and apply it to ingress interface **xe-0/0/10**. The BA classifier assigns loss priorities, as shown in [Table 547](#), to incoming packets in the four default forwarding classes. You can adapt the example to DSCP traffic by specifying a DSCP classifier instead of an IEEE classifier, and by applying DSCP bits instead of CoS bits.

To set multifield classifiers, use firewall filter rules.

Table 547: ba-classifier Loss Priority Assignments

| Forwarding Class | CoS Traffic Type                                                   | ba-classifier Loss Priority to IEEE 802.1p Code Point Mapping | Packet Drop Attribute |
|------------------|--------------------------------------------------------------------|---------------------------------------------------------------|-----------------------|
| <b>be</b>        | Best-effort traffic                                                | Low loss priority code point: <b>000</b>                      | drop                  |
| <b>fcoe</b>      | Guaranteed delivery for Fibre Channel over Ethernet (FCoE) traffic | Low loss priority code point: <b>011</b>                      | no-loss               |
| <b>no-loss</b>   | Guaranteed delivery for TCP traffic                                | Low loss priority code point: <b>100</b>                      | no-loss               |
| <b>nc</b>        | Network-control traffic                                            | Low loss priority code point: <b>110</b>                      | drop                  |

## Configuring Classifiers

To configure an IEEE 802.1 BA classifier named **ba-classifier** as the default IEEE 802.1 classifier:

- Associate code point **000** with forwarding class **be** and loss priority **low**:  

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 ba-classifier import default forwarding-class be loss-priority low code-points 000
```
- Associate code point **011** with forwarding class **fcoe** and loss priority **low**:  

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 ba-classifier forwarding-class fcoe loss-priority low code-points 011
```
- Associate code point **100** with forwarding class **no-loss** and loss priority **low**:  

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 ba-classifier forwarding-class no-loss loss-priority low code-points 100
```
- Associate code point **110** with forwarding class **nc** and loss priority **low**:  

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 ba-classifier forwarding-class nc loss-priority low code-points 110
```
- Apply the classifier to ingress interface **xe-0/0/10**:  

```
[edit class-of-service interfaces]
user@switch# set xe-0/0/10 unit 0 classifiers ieee-802.1 ba-classifier
```

## Verification

To verify the classifier configuration, perform these tasks:

- [Verifying the Classifier Configuration on page 6659](#)
- [Verifying the Ingress Interface Configuration on page 6659](#)

### Verifying the Classifier Configuration

**Purpose** Verify that you configured the classifier with the correct forwarding classes, loss priorities, and code points.

**Action** List the classifier configuration using the operational mode command **show configuration class-of-service classifiers ieee-802.1 ba-classifier**:

```
user@switch> show configuration class-of-service classifiers ieee-802.1 ba-classifier
 forwarding-class be {
 loss-priority low code-points 000;
 }
 forwarding-class fcoe {
 loss-priority low code-points 011;
 }
 forwarding-class no-loss {
 loss-priority low code-points 100;
 }
 forwarding-class nc
 loss-priority low code-points 110;
 }
```

### Verifying the Ingress Interface Configuration

**Purpose** Verify that the classifier **ba-classifier** is attached to ingress interface **xe-0/0/10**.

**Action** List the ingress interface using the operational mode command **show configuration class-of-service interfaces xe-0/0/10**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/10
congestion-notification-profile fcoe-cnp;
unit 0 {
 classifiers {
 ieee-802.1 ba-classifier;
 }
}
```

- Related Documentation**
- [Defining CoS BA Classifiers \(DSCP, DSCP IPv6, IEEE 802.1p\) on page 6656](#)
  - [Configuring a Global MPLS EXP Classifier on page 4977](#)
  - [Configuring Rewrite Rules for MPLS EXP Classifiers on page 4992](#)
  - [Monitoring CoS Classifiers on page 7231](#)
  - [Understanding CoS Classifiers on page 6650](#)
  - [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 6672](#)

## Understanding CoS MPLS EXP Classifiers and Rewrite Rules

---

You can use class of service (CoS) within MPLS networks to prioritize certain types of traffic during periods of congestion by applying packet classifiers and rewrite rules to the MPLS traffic. MPLS classifiers are global and apply to all interfaces configured as **family mpls** interfaces.

When a packet enters a customer-edge interface on the ingress provider edge (PE) switch, the switch associates the packet with a particular CoS servicing level before placing the packet onto the label-switched path (LSP). The switches within the LSP utilize the CoS value set at the ingress PE switch to determine the CoS service level. The CoS value embedded in the classifier is translated and encoded in the MPLS header by means of the experimental (EXP) bits.

EXP classifiers map incoming MPLS packets to a forwarding class and a loss priority, and assign MPLS packets to output queues based on the forwarding class mapping. EXP classifiers are behavior aggregate (BA) classifiers.

EXP rewrite rules change (rewrite) the CoS value of the EXP bits in outgoing packets on the egress queues of the switch so that the new (rewritten) value matches the policies of a targeted peer. Policy matching allows the downstream routing platform or switch in a neighboring network to classify each packet into the appropriate service group.



**NOTE:** On QFX5200, QFX5100, QFX3500, QF3600, and EX4600 switches, and on QFabric systems, there is no default EXP classifier. If you want to classify incoming MPLS packets using the EXP bits, you must configure a global EXP classifier. The global EXP classifier applies to all MPLS traffic on interfaces configured as **family mpls**.

On QFX10000 switches, there is a no default EXP classifier. If you want to classify incoming MPLS packets using the EXP bits, you must configure EXP classifiers and apply them to logical interfaces configured as **family mpls**. (You cannot apply classifiers to physical interfaces.). You can configure up to 64 EXP classifiers.

There is no default EXP rewrite rule. If you want to rewrite the EXP bit value at the egress interface, you must configure EXP rewrite rules and apply them to logical interfaces.

EXP classifiers and rewrite rules are applied only to interfaces that are configured as **family mpls** (for example, set interfaces **xe-0/0/35 unit 0 family mpls.**)

---

This topic includes:

- [EXP Classifiers on page 6661](#)
- [EXP Rewrite Rules on page 6662](#)
- [Schedulers on page 6663](#)

## EXP Classifiers

On QFX5200, QFX5100, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems, unlike DSCP and IEEE 802.1p BA classifiers, EXP classifiers are global to the switch and apply to all switch interfaces that are configured as **family mpls**. On QFX10000 switches, you apply EXP classifiers to individual logical interfaces, and different interfaces can use different EXP classifiers.

When you configure and apply an EXP classifier, MPLS traffic on all **family mpls** interfaces uses the EXP classifier, even on interfaces that also have a fixed classifier. If an interface has both an EXP classifier and a fixed classifier, the EXP classifier is applied to MPLS traffic and the fixed classifier is applied to all other traffic.

Also unlike DSCP and IEEE 802.1p BA classifiers, there is no default EXP classifier. If you want to classify MPLS traffic based on the EXP bits, you must explicitly configure an EXP classifier and apply it to the switch interfaces. Each EXP classifier has eight entries that correspond to the eight EXP CoS values (0 through 7, which correspond to CoS bits 000 through 111).

You can configure up to 64 EXP classifiers.

However, on QFX5200, QFX5100, EX4600, and legacy CLI switches, the switch uses only one MPLS EXP classifier as a global classifier on all interfaces. After you configure an MPLS EXP classifier, you can configure that classifier as the global EXP classifier by including the EXP classifier in the **[edit class-of-service system-defaults classifiers exp]** hierarchy level. All switch interfaces configured as **family mpls** use the global EXP classifier to classify MPLS traffic.

On these switches, only one EXP classifier can be configured as the global EXP classifier at any time. If you want to change the global EXP classifier, delete the global EXP classifier configuration (use the **user@switch# delete class-of-service system-defaults classifiers exp** configuration statement), then configure the new global EXP classifier.

QFX10000 switches do not support global EXP classifiers. You can configure one EXP classifier and apply it to multiple logical interfaces, or configure multiple EXP classifiers and apply different EXP classifiers to different logical interfaces.

If an EXP classifier is not configured, then if a fixed classifier is applied to the interface, the MPLS traffic uses the fixed classifier. (Switches that have a default EXP classifier use the default classifier.) If no EXP classifier and no fixed classifier are applied to the interface, MPLS traffic is treated as best-effort traffic using the 802.1 default untrusted classifier. DSCP classifiers are not applied to MPLS traffic.

On QFX5200, QFX5100, EX4600, and legacy CLI switches, because the EXP classifier is global, you cannot configure some ports to use a fixed IEEE 802.1p classifier for MPLS traffic on some interfaces and the global EXP classifier for MPLS traffic on other interfaces. When you configure a global EXP classifier, all MPLS traffic on all interfaces uses the EXP classifier.



.....  
**NOTE:** The switch uses only the outermost label of incoming EXP packets for classification.  
.....



.....  
**NOTE:** MPLS packets with 802.1Q tags are not supported.  
.....

## EXP Rewrite Rules

As MPLS packets enter or exit a network, edge switches might be required to alter the class-of-service (CoS) settings of the packets. EXP rewrite rules set the value of the EXP CoS bits within the header of the outgoing MPLS packet on **family mpls** interfaces. Each rewrite rule reads the current forwarding class and loss priority associated with the packet, locates the chosen CoS value from a table, and writes that CoS value into the packet header, replacing the old CoS value. EXP rewrite rules apply only to MPLS traffic.

EXP rewrite rules apply only to logical interfaces. You cannot apply EXP rewrite rules to physical interfaces.

There are no default EXP rewrite rules. If you want to rewrite the EXP value in MPLS packets, you must configure EXP rewrite rules and apply them to logical interfaces. If no rewrite rules are applied, all MPLS labels that are pushed have a value of zero (0). The EXP value remains unchanged on MPLS labels that are swapped.

You can configure up to 64 EXP rewrite rules, but you can only apply 16 EXP rewrite rules at any time on the switch. On a given logical interface, all pushed MPLS labels have the same EXP rewrite rule applied to them. You can apply different EXP rewrite rules to different logical interfaces on the same physical interface.

You can apply an EXP rewrite rule to an interface that has a DSCP, DSCP IPv6, or IEEE 802.1p rewrite rule. Only MPLS traffic uses the EXP rewrite rule. MPLS traffic does not use DSCP or DSCP IPv6 rewrite rules.

If the switch is performing penultimate hop popping (PHP), EXP rewrite rules do not take effect. If both an EXP classifier and an EXP rewrite rule are configured on the switch, then the EXP value from the last popped label is copied into the inner label. If either an EXP classifier or an EXP rewrite rule (but not both) is configured on the switch, then the inner label EXP value is sent unchanged.



.....  
**NOTE:** On each physical interface, either all forwarding classes that are being used on the interface must have rewrite rules configured or no forwarding classes that are being used on the interface can have rewrite rules configured. On any physical port, do not mix forwarding classes with rewrite rules and forwarding classes without rewrite rules.  
.....

## Schedulers

The schedulers for using CoS with MPLS are the same as for the other CoS configurations on the switch. Default schedulers are provided only for the best-effort, fcoe, no-loss, and network-control default forwarding classes. If you configure a custom forwarding class for MPLS traffic, you need to configure a scheduler to support that forwarding class and provide bandwidth to that forwarding class.

### Related Documentation

- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 6672](#)

## Configuring a Global MPLS EXP Classifier

EXP packet classification associates incoming packets with a particular MPLS CoS servicing level. EXP behavior aggregate (BA) classifiers examine the MPLS EXP value in the packet header to determine the CoS settings applied to the packet. EXP BA classifiers allow you to set the forwarding class and loss priority of an MPLS packet based on the incoming CoS value.

You can configure up to 64 EXP classifiers, however, the switch uses only one MPLS EXP classifier as a global classifier, which is applied only on interfaces configured as **family mpls**. All **family mpls** switch interfaces use the global EXP classifier to classify MPLS traffic.

There is no default EXP classifier. If you want to classify incoming MPLS packets using the EXP bits, you must configure a global EXP classifier. The global classifier applies to all MPLS traffic on all **family mpls** interfaces.

If a global EXP classifier is configured, MPLS traffic on **family mpls** interfaces uses the EXP classifier. If a global EXP classifier is not configured, then if a fixed classifier is applied to the interface, the MPLS traffic uses the fixed classifier. If no EXP classifier and no fixed classifier is applied to the interface, MPLS traffic is treated as best-effort traffic. DSCP classifiers are not applied to MPLS traffic.

To configure an MPLS EXP classifier using the CLI:

1. Create an EXP classifier and associate it with a forwarding class, a loss priority, and a code point:

```
[edit class-of-service classifiers]
user@switch# set (dscp | ieee-802.1 | exp) classifier-name forwarding-class
forwarding-class-name loss-priority level code-points [aliases] [bit-patterns]
```

2. Apply the EXP classifier to the switch interfaces:

```
[edit class-of-service]
user@switch# set system-defaults classifiers exp classifier-name
```

### Related Documentation

- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4959](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 6672](#)
- [Defining CoS Unicast BA Classifiers \(DSCP, DSCP IPv6, IEEE 802.1p\)](#)

- [Defining CoS BA Classifiers \(DSCP, DSCP IPv6, IEEE 802.1p\) on page 6656](#)
- [Configuring Rewrite Rules for MPLS EXP Classifiers on page 4992](#)

## Configuring Rewrite Rules for MPLS EXP Classifiers

You configure EXP rewrite rules to alter CoS values in outgoing MPLS packets on the outbound **family mpls** interfaces of a switch to match the policies of a targeted peer. Policy matching allows the downstream routing platform or switch in a neighboring network to classify each packet into the appropriate service group.

To configure an EXP CoS rewrite rule, create the rule by giving it a name and associating it with a forwarding class, loss priority, and code point. This creates a rewrite table. After the rewrite rule is created, enable it on a logical **family mpls** interface. EXP rewrite rules can only be enabled on logical **family mpls** interfaces, not on physical interfaces or on interfaces of other family types. You can also apply an existing EXP rewrite rule on a logical interface.



**NOTE:** There are no default rewrite rules.

You can configure up to 64 EXP rewrite rules, but you can only use 16 EXP rewrite rules at any time on the switch. On a given **family mpls** logical interface, all pushed MPLS labels have the same EXP rewrite rule applied to them. You can apply different EXP rewrite rules to different logical interfaces on the same physical interface.



**NOTE:** On each physical interface, either all forwarding classes that are being used on the interface must have rewrite rules configured, or no forwarding classes that are being used on the interface can have rewrite rules configured. On any physical port, do not mix forwarding classes with rewrite rules and forwarding classes without rewrite rules.



**NOTE:** To replace an existing rewrite rule on the interface with a new rewrite rule of the same type, first explicitly remove the existing rewrite rule and then apply the new rule.

To create an EXP rewrite rule for MPLS traffic and enable it on a logical interface:

1. Create an EXP rewrite rule:

```
user@switch# set class-of-service rewrite-rules exp rewrite-rule-name forwarding-class
forwarding-class-name loss-priority level code-points [aliases] [bit-patterns]
```

For example, to configure an EXP rewrite rule named **exp-rr-1** for a forwarding class named **mpls-1** with a loss priority of **low** that rewrites the EXP code point value to **001**:

```
user@switch# set class-of-service rewrite-rules exp exp-rr-1 forwarding-class mpls-1
loss-priority low code-points 001
```



2. Apply the rewrite rule to a logical interface:

```
user@switch # set class-of-service interfaces interface-name unit logical-unit rewrite-rules
exp rewrite-rule-name
```

For example, to apply a rewrite rule named **exp-rr-1** to logical interface **xe-0/0/10.0**:

```
user@switch# set class-of-service interfaces xe-0/0/10 unit 0 rewrite-rules exp exp-rr-1
```



**NOTE:** In this example, all forwarding classes assigned to port xe-0/0/10 must have rewrite rules. Do not mix forwarding classes that have rewrite rules with forwarding classes that do not have rewrite rules on the same interface.

#### Related Documentation

- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4959](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 6672](#)
- [Monitoring CoS Rewrite Rules on page 7235](#)
- [Defining CoS Rewrite Rules on page 6704](#)

## Understanding Host Inbound Traffic Classification

The destination address of traffic that enters the switch can be an external device such as another switch, a router, or a server, or the destination can be the host (the switch Routing Engine or CPU). When the destination is an external device, the DSCP and IEEE 802.1p code-point bits of incoming traffic are preserved as the traffic travels through the switch to the egress port. At the egress port, the code-point bits are either preserved when the packets are sent to the next hop or they are rewritten according to the rewrite rule attached to the egress interface.

When the destination of incoming traffic is the host, DSCP bits are preserved. However, IEEE 802.1p bits are not preserved. The IEEE 802.1p bits of traffic destined for the host are set to zero (0). This does not affect system behavior because the switch prioritizes traffic destined for the host based on the protocol type. For example, the switch gives a higher priority to BPDU traffic than to ping traffic.

## Understanding Default CoS Scheduling and Classification

---

If you do not explicitly configure classifiers and apply them to interfaces, the switch uses the default classifier to group ingress traffic into forwarding classes. If you do not configure scheduling on an interface, the switch uses the default schedulers to provide egress port resources for traffic. Default classification maps all traffic into the default forwarding classes (best-effort, fcoe, no-loss, and network-control). Each default forwarding class has a default scheduler, so traffic mapped to each default forwarding class receives port bandwidth, prioritization, and packet drop characteristics.

The switch supports direct port scheduling and enhanced transmission selection (ETS), also known as hierarchical port scheduling.

Hierarchical scheduling groups IEEE 802.1p priorities (IEEE 802.1p code points, which classifiers map to forwarding classes, which in turn are mapped to output queues) into priority groups (forwarding class sets). If you use only the default traffic scheduling and classification, the switch automatically creates a default priority group that contains all of the priorities (which are mapped to forwarding classes and output queues), and assigns 100 percent of the port output bandwidth to that priority group. The forwarding classes (queues) in the default forwarding class set receive bandwidth based on the default classifier settings. The default priority group is transparent. It does not appear in the configuration and is used for Data Center Bridging Capability Exchange (DCBX) protocol advertisement.



**NOTE:** If you explicitly configure ETS by configuring one or more priority groups on an interface, any forwarding class that is not assigned to a priority group on that interface receives *no bandwidth*. This means that if you configure hierarchical scheduling on an interface, every forwarding class (priority) that you want to forward traffic on that interface must belong to a forwarding class set (priority group). ETS is not supported on QFX5200.

---

This topic describes:

- [Default Classification on page 6666](#)
- [Default Scheduling on page 6669](#)
- [Default DCBX Advertisement on page 6671](#)
- [Default Scheduling and Classification Summary on page 6671](#)

### Default Classification

The default classifiers assign ingress traffic to default forwarding classes and loss priorities. The switch applies default IEEE 802.1, DSCP, and DSCP IPv6 classifiers to each interface that does not have explicitly configured classifiers. If you do not configure and apply EXP classifiers for MPLS traffic to logical interfaces, MPLS traffic on interfaces configured as **family mpls** uses the IEEE classifier.

If you explicitly configure one type of classifier but not other types of classifiers, the system uses only the configured classifier and does not use default classifiers for other types of traffic. There are two default IEEE 802.1 classifiers, a trusted classifier for ports that are in trunk mode, and an untrusted classifier for ports that are in access mode.

Table 548 shows the default mapping of IEEE 802.1 code-point values to forwarding classes and loss priorities for ports in trunk mode.

**Table 548: Default IEEE 802.1 Classifiers for Ports in Trunk Mode (Trusted Classifier)**

| Code Point | Forwarding Class | Loss Priority |
|------------|------------------|---------------|
| be (000)   | best-effort      | low           |
| be1 (001)  | best-effort      | low           |
| ef (010)   | best-effort      | low           |
| ef1 (011)  | fcoe             | low           |
| af11 (100) | no-loss          | low           |
| af12 (101) | best-effort      | low           |
| nc1 (110)  | network-control  | low           |
| nc2 (111)  | network-control  | low           |

Table 549 shows the default mapping of IEEE 802.1p code-point values to forwarding classes and loss priorities for ports in access mode (all incoming traffic is mapped to best-effort forwarding classes).

**Table 549: Default IEEE 802.1 Classifiers for Ports in Access Mode (Untrusted Classifier)**

| Code Point | Forwarding Class | Loss Priority |
|------------|------------------|---------------|
| 000        | best-effort      | low           |
| 001        | best-effort      | low           |
| 010        | best-effort      | low           |
| 011        | best-effort      | low           |
| 100        | best-effort      | low           |
| 101        | best-effort      | low           |
| 110        | best-effort      | low           |

**Table 549: Default IEEE 802.1 Classifiers for Ports in Access Mode (Untrusted Classifier) (*continued*)**

| Code Point | Forwarding Class | Loss Priority |
|------------|------------------|---------------|
| 111        | best-effort      | low           |

Table 550 shows the default mapping of DSCP code-point values to forwarding classes and loss priorities for DSCP IP and DCSP IPv6.

**Table 550: Default DSCP IP and IPv6 Classifiers**

| Code Point    | Forwarding Class | Loss Priority |
|---------------|------------------|---------------|
| ef (101110)   | best-effort      | low           |
| af11 (001010) | best-effort      | low           |
| af12 (001100) | best-effort      | low           |
| af13 (001110) | best-effort      | low           |
| af21 (010010) | best-effort      | low           |
| af22 (010100) | best-effort      | low           |
| af23 (010110) | best-effort      | low           |
| af31 (011010) | best-effort      | low           |
| af32 (011100) | best-effort      | low           |
| af33 (011110) | best-effort      | low           |
| af41 (100010) | best-effort      | low           |
| af42 (100100) | best-effort      | low           |
| af43 (100110) | best-effort      | low           |
| be (000000)   | best-effort      | low           |
| cs1 (001000)  | best-effort      | low           |
| cs2 (010000)  | best-effort      | low           |
| cs3 (011000)  | best-effort      | low           |
| cs4 (100000)  | best-effort      | low           |
| cs5 (101000)  | best-effort      | low           |

**Table 550: Default DSCP IP and IPv6 Classifiers (*continued*)**

| Code Point   | Forwarding Class | Loss Priority |
|--------------|------------------|---------------|
| nc1 (110000) | network-control  | low           |
| nc2 (111000) | network-control  | low           |

[Table 551](#) shows the default mapping of MPLS EXP code-point values to forwarding classes and loss priorities.

**Table 551: Default EXP Classifiers**

| Code Point | Forwarding Class     | Loss Priority |
|------------|----------------------|---------------|
| 000        | best-effort          | low           |
| 001        | best-effort          | high          |
| 010        | expedited-forwarding | low           |
| 011        | expedited-forwarding | high          |
| 100        | assured-forwarding   | low           |
| 101        | assured-forwarding   | high          |
| 110        | network-control      | low           |
| 111        | network-control      | high          |

## Default Scheduling

The default schedulers allocate egress bandwidth resources to egress traffic as shown in [Table 552](#):

**Table 552: Default Scheduler Configuration**

| Default Scheduler and Queue Number               | Transmit Rate (Guaranteed Minimum Bandwidth) | Rate Shaping (Maximum Bandwidth) | Excess Bandwidth Sharing | Priority | Buffer Size |
|--------------------------------------------------|----------------------------------------------|----------------------------------|--------------------------|----------|-------------|
| best-effort forwarding class scheduler (queue 0) | 15%                                          | None                             | 15%                      | low      | 15%         |
| fcoe forwarding class scheduler (queue 3)        | 35%                                          | None                             | 35%                      | low      | 35%         |
| no-loss forwarding class scheduler (queue 4)     | 35%                                          | None                             | 35%                      | low      | 35%         |

Table 552: Default Scheduler Configuration (*continued*)

| Default Scheduler and Queue Number                   | Transmit Rate (Guaranteed Minimum Bandwidth) | Rate Shaping (Maximum Bandwidth) | Excess Bandwidth Sharing | Priority | Buffer Size |
|------------------------------------------------------|----------------------------------------------|----------------------------------|--------------------------|----------|-------------|
| network-control forwarding class scheduler (queue 7) | 15%                                          | None                             | 15%                      | low      | 15%         |



**NOTE:** By default, the minimum guaranteed bandwidth (transmit rate) determines the amount of excess (extra) bandwidth a queue can share. Extra bandwidth is allocated to queues in proportion to the transmit rate of each queue. On switches that support the `excess-rate` statement, you can override the default setting and configure the excess bandwidth percentage independently of the transmit rate on queues that are not strict-high priority queues.

By default, only the four default schedulers shown in Table 552 have traffic mapped to them. Only the forwarding classes and queues associated with the default schedulers receive default bandwidth, based on the default scheduler transmit rate. (You can configure schedulers and forwarding classes to allocate bandwidth to other queues or to change the bandwidth and other scheduling properties of a default queue.) If a forwarding class does not transport traffic, the bandwidth allocated to that forwarding class is available to other forwarding classes. Unicast and multdestination (multicast, broadcast, and destination lookup fail) traffic use the same forwarding classes and output queues.

Default scheduling is port scheduling.

Default hierarchical scheduling, known as enhanced transmission selection (ETS, defined in IEEE 802.1Qaz), allocates the total port bandwidth to the four default forwarding classes served by the four default schedulers, as defined by the four default schedulers. The result is the same as direct port scheduling. Configuring hierarchical port scheduling, however, enables you to group forwarding classes that carry similar types of traffic into forwarding class sets (also called priority groups), and to assign port bandwidth to each forwarding class set. The port bandwidth assigned to the forwarding class set is then assigned to the forwarding classes within the forwarding class set. This hierarchy enables you to control port bandwidth allocation with greater granularity, and enables hierarchical sharing of extra bandwidth to better utilize link bandwidth.

Default scheduling uses weighted round-robin (WRR) scheduling. Each queue receives a portion (weight) of the total available interface bandwidth. The scheduling weight is based on the transmit rate (minimum guaranteed bandwidth) of the default scheduler for that queue. For example, queue 7 receives a default scheduling weight of 15 percent of the available bandwidth, and queue 4 receives a default scheduling weight of 35 percent of the available bandwidth. Queues are mapped to forwarding classes (for example, queue 7 is mapped to the network-control forwarding class and queue 4 is mapped to the no-loss forwarding class). Each forwarding class receives the default

bandwidth for the queue to which it is mapped. Unused bandwidth is shared with other default queues.

If you want non-default (unconfigured) queues to forward traffic, you should explicitly map traffic to those queues (configure the forwarding classes and queue mapping) and create schedulers to allocate bandwidth to those queues. By default, queues 1, 2, 5, and 6 are unconfigured. Unconfigured queues have a default scheduling weight of 1 so that they can receive a small amount of bandwidth in case they need to forward traffic.

If you map traffic to an unconfigured queue and do not schedule port resources for the queue (configure a scheduler, map it to the forwarding class that is mapped to the queue, and apply the scheduler mapping to the port), the queue receives only the amount of excess bandwidth proportional to its default weight (1). The actual amount of bandwidth an unconfigured queue gets depends on how much bandwidth the other queues on the port are using.

If the other queues use less than their allocated amount of bandwidth, the unconfigured queues can share the unused bandwidth. Configured queues have higher priority for bandwidth than unconfigured queues, so if a configured queue needs more bandwidth, then less bandwidth is available for unconfigured queues. Unconfigured queues always receive a minimum amount of bandwidth based on their scheduling weight (1). If you map traffic to an unconfigured queue, to allocate bandwidth to that queue, configure a scheduler for the forwarding class that is mapped to the queue, and apply it to the port.

## Default DCBX Advertisement

When you configure hierarchical scheduling on an interface, DCBX advertises each priority group, the priorities in each priority group, and the bandwidth properties of each priority and priority group.

If you do not configure hierarchical scheduling on an interface, DCBX advertises the automatically created default priority group and its priorities. DCBX also advertises the default bandwidth allocation of the priority group, which is 100 percent of the port bandwidth.

## Default Scheduling and Classification Summary

If you do not configure scheduling on an interface:

- Default classifiers classify ingress traffic.
- Default schedulers schedule egress traffic.
- DCBX advertises a single default priority group with 100 percent of the port bandwidth allocated to that priority group. All priorities (forwarding classes) are assigned to the default priority group and receive bandwidth based on their default schedulers. The default priority group is generated automatically and is not user-configurable.

### Related Documentation

- [Understanding CoS Packet Flow on page 6628](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 6765](#)
- [Understanding Default CoS Settings on page 6630](#)

- [CoS Support on QFX Series Switches, EX4600 Switches, and QFabric Systems on page 6610](#)
- [Understanding CoS Virtual Output Queues \(VOQs\) on QFX10000 Switches on page 6724](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 6672](#)
- [Understanding DCB Features and Requirements on page 6490](#)
- [Example: Configuring Classifiers on page 6657](#)
- [Example: Configuring Queue Schedulers for Port Scheduling on page 6744](#)

## Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces

At ingress interfaces, classifiers group incoming traffic into classes based on the IEEE 802.1p, DSCP, or MPLS EXP class of service (CoS) code point bits in the packet header. At egress interfaces, you can use rewrite rules to change (re-mark) the code point bits before the interface forwards the packets.

You can apply classifiers and rewrite rules to interfaces to control the level of CoS applied to each packet as it traverses the system and the network. This topic describes:

- [Supported Classifier and Rewrite Rule Types on page 6672](#)
- [Ethernet Interfaces Supported for Classifier and Rewrite Rule Configuration on page 6674](#)
- [Default Classifiers on page 6677](#)
- [Default Rewrite Rules on page 6678](#)
- [Classifier Precedence on page 6678](#)
- [Classifier Behavior and Limitations on page 6680](#)
- [Rewrite Rule Precedence and Behavior on page 6681](#)
- [Classifier and Rewrite Rule Configuration Interaction with Ethernet Interface Configuration on page 6681](#)

## Supported Classifier and Rewrite Rule Types

[Table 553](#) shows the supported types of classifiers and rewrite rules supports:

**Table 553: Supported Classifiers and Rewrite Rules**

| Classifier or Rewrite Rule Type        | Description                                                                                                                              |
|----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------|
| Fixed classifier                       | Classifies all ingress traffic on a physical interface into one fixed forwarding class, regardless of the CoS bits in the packet header. |
| DSCP and DSCP IPv6 unicast classifiers | Classifies IP and IPv6 traffic into forwarding classes and assigns loss priorities to the traffic based on DSCP code point bits.         |



Table 553: Supported Classifiers and Rewrite Rules (*continued*)

| Classifier or Rewrite Rule Type                                                                                                                                                                                                                                                 | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IEEE 802.1p unicast classifier                                                                                                                                                                                                                                                  | Classifies Ethernet traffic into forwarding classes and assigns loss priorities to the traffic based on IEEE 802.1p code point bits.                                                                                                                                                                                                                                                                                                                                                             |
| MPLS EXP classifier                                                                                                                                                                                                                                                             | <p>Classifies MPLS traffic into forwarding classes and assigns loss priorities to the traffic on interfaces configured as <b>family mpls</b>.</p> <p>QFX5200, QFX5100, EX4600, QFX3500, and QFX3600 switches, and QFabric systems, use one global EXP classifier on all <b>family mpls</b> switch interfaces.</p> <p>QFX10000 switches do not support global EXP classifiers. You can apply the same EXP classifier or different EXP classifiers to different <b>family mpls</b> interfaces.</p> |
| DSCP multidestination classifier (also used for IPv6 multidestination traffic)                                                                                                                                                                                                  | Classifies IP and IPv6 multicast, broadcast, and destination lookup fail (DLF) traffic into multidestination forwarding classes. Multidestination classifiers are applied to all interfaces and cannot be applied to individual interfaces.                                                                                                                                                                                                                                                      |
| <p><b>NOTE:</b> This applies only to switches that use different classifiers for unicast and multidestination traffic. It does not apply to switches that use the same classifiers for unicast and multidestination traffic.</p> <p>IEEE 802.1p multidestination classifier</p> | <p>Classifies Ethernet multicast, broadcast, and destination lookup fail (DLF) traffic into multidestination forwarding classes. Multidestination classifiers are applied to all interfaces and cannot be applied to individual interfaces.</p> <p><b>NOTE:</b> This applies only to switches that use different classifiers for unicast and multidestination traffic. It does not apply to switches that use the same classifiers for unicast and multidestination traffic.</p>                 |
| DSCP and DSCP IPv6 rewrite rules                                                                                                                                                                                                                                                | Re-marks the DSCP code points of IP and IPv6 packets before forwarding the packets.                                                                                                                                                                                                                                                                                                                                                                                                              |
| IEEE 802.1p rewrite rule                                                                                                                                                                                                                                                        | Re-marks the IEEE 802.1p code points of Ethernet packets before forwarding the packets.                                                                                                                                                                                                                                                                                                                                                                                                          |
| MPLS EXP rewrite rule                                                                                                                                                                                                                                                           | Re-marks the EXP code points of MPLS packets before forwarding the packets on interfaces configured as <b>family mpls</b> .                                                                                                                                                                                                                                                                                                                                                                      |



**NOTE:** On switches that support native Fibre Channel (FC) interfaces, you can specify a rewrite value on native FC interfaces (NP\_Ports) to set the IEEE 802.1p code point of incoming FC traffic when the NP\_Port encapsulates the FC packet in Ethernet before forwarding it to the FCoE network (see *Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway*).

DSCP, IEEE 802.1p, and MPLS EXP classifiers are behavior aggregate (BA) classifiers. On QFX5100, QFX5200, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems, unlike DSCP and IEEE 802.1p classifiers, EXP classifiers are global and apply only to all interfaces that are configured as **family mpls**. On QFX10000 switches, you apply EXP classifiers to individual logical interfaces, and different interfaces can use different EXP classifiers.

Unlike DSCP and IEEE 802.1p BA classifiers, there is no default EXP classifier. Also unlike DSCP and IEEE 802.1p classifiers, for MPLS traffic on **family mpls** interfaces only, EXP classifiers overwrite fixed classifiers. (An interface that has a fixed classifier uses the EXP classifier for MPLS traffic, not the fixed classifier, and the fixed classifier is used for all other traffic.)

On switches that use different classifiers for unicast and multdestination traffic, multdestination classifiers are global and apply to all interfaces; you cannot apply a multdestination classifier to individual interfaces.

Classifying packets into forwarding classes assigns packets to the output queues mapped to those forwarding classes. The traffic classified into a forwarding class receives the CoS scheduling configured for the output queue mapped to that forwarding class.



**NOTE:** In addition to BA classifiers and fixed classifiers, which classify traffic based on the CoS field in the packet header, you can use firewall filters to configure multifield (MF) classifiers. MF classifiers classify traffic based on more than one field in the packet header and take precedence over BA and fixed classifiers.

---

## Ethernet Interfaces Supported for Classifier and Rewrite Rule Configuration

To apply a classifier to incoming traffic or a rewrite rule to outgoing traffic, you need to apply the classifier or rewrite rule to one or more interfaces. When you apply a classifier or rewrite rule to an interface, the interface uses the classifier to group incoming traffic into forwarding classes and uses the rewrite rule to re-mark the CoS code point value of each packet before it leaves the system.

Not all interfaces types support all types of CoS configuration. This section describes:

- [Interface Types That Support Classifier and Rewrite Rule Configuration on page 6674](#)
- [Classifier and Rewrite Rule Physical and Logical Ethernet Interface Support on page 6675](#)
- [Routed VLAN Interfaces \(RVIs\) and Integrated Routing and Bridging \(IRB\) Interfaces on page 6677](#)

---

### Interface Types That Support Classifier and Rewrite Rule Configuration

You can apply classifiers and rewrite rules to Ethernet interfaces. For Layer 3 LAGs, configure BA or fixed classifiers on the LAG (ae) interface. The classifier configured on the LAG is valid on all of the LAG member interfaces.

On switches that support native FC interfaces, you can apply fixed classifiers to native FC interfaces (NP\_Ports). You cannot apply other types of classifiers or rewrite rules to native FC interfaces. You can rewrite the value of the IEEE 802.1p code point of incoming FC traffic when the interface encapsulates it in Ethernet before forwarding it to the FCoE network as described in *Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway*.

### Classifier and Rewrite Rule Physical and Logical Ethernet Interface Support

The Ethernet ports can function as:

- Layer 2 physical interfaces (family ethernet-switching)
- Layer 2 logical interfaces (family ethernet-switching)
- Layer 3 physical interfaces (family inet/inet6)
- Layer 3 logical interfaces (family inet/inet6)
- MPLS interfaces (family mpls)

You can apply CoS classifiers and rewrite rules only to the following interfaces:

- Layer 2 logical interface unit 0



**NOTE:** On a Layer 2 interface, the CoS you configure on logical interface unit 0 applies to all of the logical units on that interface.

- On QFX5100, QFX5200, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems, Layer 3 physical interfaces if at least one logical Layer 3 interface is configured on the physical interface



**NOTE:** The CoS you configure on a Layer 3 physical interface is applied to all of the Layer 3 logical interfaces on that physical interface. This means that each Layer 3 interface uses the same classifiers and rewrite rules for all of the Layer 3 traffic on that interface.

- On QFX10000 switches, Layer 3 logical interfaces. You can apply different classifiers and rewrite rules to different Layer 3 logical interfaces.

### Ethernet Interface Support for Most QFX Series Switches, and QFabric Systems

You cannot apply classifiers or rewrite rules to Layer 2 physical interfaces or to Layer 3 logical interfaces. [Table 554](#) shows on which interfaces you can configure and apply classifiers and rewrite rules.

**Table 554: Ethernet Interface Support for Classifier and Rewrite Rule Configuration (QFX5100, QFX5200, EX4600, QFX3500, and QFX3600 Switches, and QFabric Systems)**

| CoS Classifiers and Rewrite Rules | Layer 2 Physical Interfaces                                                                                                                         | Layer 2 Logical Interface (Unit 0 Only) | Layer 3 Physical Interfaces (If at Least One Logical Layer 3 Interface Is Defined) | Layer 3 Logical Interfaces |
|-----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------|------------------------------------------------------------------------------------|----------------------------|
| Fixed classifier                  | No                                                                                                                                                  | Yes                                     | Yes                                                                                | No                         |
| DSCP classifier                   | No                                                                                                                                                  | Yes                                     | Yes                                                                                | No                         |
| DSCP IPv6 classifier              | No                                                                                                                                                  | Yes                                     | Yes                                                                                | No                         |
| IEEE 802.1p classifier            | No                                                                                                                                                  | Yes                                     | Yes                                                                                | No                         |
| EXP classifier                    | Global classifier, applies only to all switch interfaces that are configured as <b>family mpls</b> . Cannot be configured on individual interfaces. |                                         |                                                                                    |                            |
| DSCP rewrite rule                 | No                                                                                                                                                  | Yes                                     | Yes                                                                                | No                         |
| DSCP IPv6 rewrite rule            | No                                                                                                                                                  | Yes                                     | Yes                                                                                | No                         |
| IEEE 802.1p rewrite rule          | No                                                                                                                                                  | Yes                                     | Yes                                                                                | No                         |
| EXP rewrite rule                  | No                                                                                                                                                  | Yes                                     | Yes                                                                                | No                         |



**NOTE:** IEEE 802.1p multidestination and DSCP multidestination classifiers are applied to all interfaces and cannot be applied to individual interfaces. No DSCP IPv6 multidestination classifier is supported. IPv6 multidestination traffic uses the DSCP multidestination classifier.

#### ***Ethernet Interface Support for QFX10000 Switches***

You cannot apply classifiers or rewrite rules to Layer 2 or Layer 3 physical interfaces. You can apply classifiers and rewrite rules only to Layer 2 logical interface unit 0. You can apply different classifiers and rewrite rules to different Layer 3 logical interfaces. [Table 555](#) shows on which interfaces you can configure and apply classifiers and rewrite rules.

**Table 555: Ethernet Interface Support for Classifier and Rewrite Rule Configuration (QFX10000 Switches)**

| CoS Classifiers and Rewrite Rules | Layer 2 Physical Interfaces | Layer 2 Logical Interface (Unit 0 Only) | Layer 3 Physical Interfaces | Layer 3 Logical Interfaces |
|-----------------------------------|-----------------------------|-----------------------------------------|-----------------------------|----------------------------|
| Fixed classifier                  | No                          | Yes                                     | No                          | Yes                        |
| DSCP classifier                   | No                          | Yes                                     | No                          | Yes                        |

**Table 555: Ethernet Interface Support for Classifier and Rewrite Rule Configuration (QFX10000 Switches) (continued)**

| CoS Classifiers and Rewrite Rules | Layer 2 Physical Interfaces | Layer 2 Logical Interface (Unit 0 Only) | Layer 3 Physical Interfaces | Layer 3 Logical Interfaces |
|-----------------------------------|-----------------------------|-----------------------------------------|-----------------------------|----------------------------|
| DSCP IPv6 classifier              | No                          | Yes                                     | No                          | Yes                        |
| IEEE 802.1p classifier            | No                          | Yes                                     | No                          | Yes                        |
| EXP classifier                    | No                          | Yes                                     | No                          | Yes                        |
| DSCP rewrite rule                 | No                          | Yes                                     | No                          | Yes                        |
| DSCP IPv6 rewrite rule            | No                          | Yes                                     | No                          | Yes                        |
| IEEE 802.1p rewrite rule          | No                          | Yes                                     | No                          | Yes                        |
| EXP rewrite rule                  | No                          | Yes                                     | No                          | Yes                        |

### Routed VLAN Interfaces (RVIs) and Integrated Routing and Bridging (IRB) Interfaces

You cannot apply classifiers and rewrite rules directly to routed VLAN interfaces (RVIs) or integrated routing and bridging (IRB) interfaces because the members of RVIs and IRBs are VLANs, not ports. However, you can apply classifiers and rewrite rules to the VLAN port members of an RVI or an IRB. You can also apply MF classifiers to RVIs and IRBs.

### Default Classifiers

If you do not explicitly configure classifiers on an Ethernet interface, the switch applies default classifiers so that the traffic receives basic CoS treatment. The factors that determine the default classifier applied to the interface include the interface type (Layer 2 or Layer 3), the port mode (trunk, tagged-access, or access), and whether logical interfaces have been configured.

The switch applies default classifiers using the following rules:

- If the physical interface has at least one Layer 3 logical interface configured, the logical interfaces use the default DSCP classifier.
- If the physical interface has a Layer 2 logical interface in trunk mode or tagged-access mode, it uses the default IEEE 802.1p trusted classifier.



**NOTE:** Tagged-access mode is available only on QFX3500 and QFX3600 devices when used as standalone switches or as QFabric system Node devices.

- If the physical interface has a Layer 2 logical interface in access mode, it uses the default IEEE 802.1p untrusted classifier.
- If the physical interface has no logical interface configured, no default classifier is applied.
- On switches that use different classifiers for unicast and multidestination traffic, the default multidestination classifier is the IEEE 802.1p multidestination classifier.
- There is no default MPLS EXP classifier. If you want to classify MPLS traffic using EXP bits on these switches, on QFX10000 switches, configure an EXP classifier and apply it to a logical interface that is configured as **family mpls**. On QFX5100, QFX5200, EX4600, QFX3500 and QFX3600 switches, and on QFabric systems, configure an EXP classifier and configure it as the global system default EXP classifier.

## Default Rewrite Rules

No default rewrite rules are applied to interfaces. If you want to re-mark packets at the egress interface, you must explicitly configure a rewrite rule.

## Classifier Precedence

You can apply multiple classifiers (MF, fixed, IEEE 802.1p, DSCP, or EXP) to an Ethernet interface to handle different types of traffic. (EXP classifiers are global and apply only to all MPLS traffic on all **family mpls** interfaces.) When you apply more than one classifier to an interface, the system uses an order of precedence to determine which classifier to use on interfaces:

- [Classifier Precedence on Physical Ethernet Interfaces \(QFX5200, QFX5100, EX4600, QFX3500, and QFX3600 Switches, and QFabric Systems\) on page 6678](#)
- [Classifier Precedence on Logical Ethernet Interfaces \(All Switches\) on page 6679](#)

### Classifier Precedence on Physical Ethernet Interfaces (QFX5200, QFX5100, EX4600, QFX3500, and QFX3600 Switches, and QFabric Systems)

QFX10000 switches do not support configuring classifiers on physical interfaces. The precedence of classifiers on physical interfaces, from the highest-priority classifier to the lowest-priority classifier, is:

- MF classifier on a logical interface (no classifier has a higher priority than MF classifiers)
- Fixed classifier on the physical interface
- DSCP or DSCP IPv6 classifier on the physical interface
- IEEE 802.1p classifier on the physical interface



**NOTE:** If an EXP classifier is configured, MPLS traffic uses the EXP classifier on all **family mpls** interfaces, even if an MF or fixed classifier is applied to the interface. If an EXP classifier is not configured, then if a fixed classifier is applied to the interface, the MPLS traffic uses the fixed classifier. If no EXP classifier and no fixed classifier is applied to the interface, MPLS traffic is treated as best-effort traffic. DSCP classifiers are not applied to MPLS traffic.

You can apply a DSCP classifier, an IEEE 802.1p classifier, and an EXP classifier on a physical interface. When all three classifiers are on an interface, IP traffic uses the DSCP classifier, MPLS traffic on **family mpls** interfaces uses the EXP classifier, and all other traffic uses the IEEE classifier.



**NOTE:** You cannot apply a fixed classifier and a DSCP or IEEE classifier to the same interface. If a DSCP classifier, an IEEE classifier, or both are on an interface, you cannot apply a fixed classifier to that interface unless you first delete the DSCP and IEEE classifiers. If a fixed classifier is on an interface, you cannot apply a DSCP classifier or an IEEE classifier unless you first delete the fixed classifier.

### Classifier Precedence on Logical Ethernet Interfaces (All Switches)

The precedence of classifiers on logical interfaces, from the highest priority classifier to the lowest priority classifier, is:

- MF classifier on a logical interface (no classifier has a higher priority than MF classifiers).
- Fixed classifier on the logical interface.
- DSCP or DSCP IPv6 classifier on the physical or logical interface..
- IEEE 802.1p classifier on the physical or logical interface.



**NOTE:** If a global EXP classifier is configured, MPLS traffic uses the EXP classifier on all **family mpls** interfaces, even if a fixed classifier is applied to the interface. If a global EXP classifier is not configured, then:

- If a fixed classifier is applied to the interface, the MPLS traffic uses the fixed classifier. If no EXP classifier and no fixed classifier is applied to the interface, MPLS traffic is treated as best-effort traffic.

You can apply both a DSCP classifier and an IEEE 802.1p classifier on a logical interface. When both a DSCP and an IEEE classifier are on an interface, IP traffic uses the DSCP classifier, and all other traffic uses the IEEE classifier. Only MPLS traffic on interfaces configured as **family mpls** uses the EXP classifier.

## Classifier Behavior and Limitations

Consider the following behaviors and constraints when you apply classifiers to Ethernet interfaces. Behaviors for applying classifiers to physical interfaces do not pertain to QFX10000 switches.

- You can configure only one DSCP classifier (IP or IPv6) on a physical interface. You cannot configure both types of DSCP classifier on one physical interface. Both IP and IPv6 traffic use whichever DSCP classifier is configured on the interface.
- When you configure a DSCP or a DSCP IPv6 classifier on a physical interface and the physical interface has at least one logical Layer 3 interface, all packets (IP, IPv6, and non-IP) use that classifier.
- An interface with both a DSCP classifier (IP or IPv6) and an IEEE 802.1p classifier uses the DSCP classifier for IP and IPv6 packets, and uses the IEEE classifier for all other packets.
- Fixed classifiers and BA classifiers (DSCP and IEEE classifiers) are not permitted simultaneously on an interface. If you configure a fixed classifier on an interface, you cannot configure a DSCP or an IEEE classifier on that interface. If you configure a DSCP classifier, an IEEE classifier, or both classifiers on an interface, you cannot configure a fixed classifier on that interface.
- When you configure an IEEE 802.1p classifier on a physical interface and a DSCP classifier is not explicitly configured on that interface, the interface uses the IEEE classifier for all types of packets. No default DSCP classifier is applied to the interface. (In this case, if you want a DSCP classifier on the interface, you must explicitly configure it and apply it to the interface.)
- The system does not apply a default classifier to a physical interface until you create a logical interface on that physical interface. If you configure a Layer 3 logical interface, the system uses the default DSCP classifier. If you configure a Layer 2 logical interface, the system uses the default IEEE 802.1p trusted classifier if the port is in trunk mode or tagged-access mode, or the default IEEE 802.1p untrusted classifier if the port is in access mode.
- MF classifiers configured on logical interfaces take precedence over BA and fixed classifiers, with the exception of the global EXP classifier, which is always used for MPLS traffic on **family mpls** interfaces. (Use firewall filters to configure MF classifiers.) When BA or fixed classifiers are present on an interface, you can still configure an MF classifier on that interface.
- There is no default EXP classifier for MPLS traffic.
- You can configure up to 64 EXP classifiers. On QFX10000 switches, you can apply different EXP classifiers to different interfaces.

However, on On QFX5200, QFX5100, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems, the switch uses only one MPLS EXP classifier as a global classifier on all **family mpls** interfaces. After you configure an MPLS EXP classifier, you can configure it as the global EXP classifier by including the EXP classifier in the **[edit class-of-service system-defaults classifiers exp]** hierarchy level.



All **family mpls** switch interfaces use the EXP classifier specified using this configuration statement to classify MPLS traffic, even on interfaces that have a fixed classifier. No other traffic uses the EXP classifier.

## Rewrite Rule Precedence and Behavior

The following rules apply on Ethernet interfaces for rewrite rules:

- If you configure one DSCP (or DSCP IPv6) rewrite rule and one IEEE 802.1p rewrite rule on an interface, both rewrite rules take effect. Traffic with IP and IPv6 headers use the DSCP rewrite rule, and traffic with a VLAN tag uses the IEEE rewrite rule.
- If you do not explicitly configure a rewrite rule, there is no default rewrite rule, so the system does not apply any rewrite rule to the interface.
- You can apply a DSCP rewrite rule or a DSCP IPv6 rewrite rule to an interface, but you cannot apply both a DSCP and a DSCP IPv6 rewrite rule to the same interface. Both IP and IPv6 packets use the same DSCP rewrite rule, regardless of whether the configured rewrite rule is DSCP or DSCP IPv6.
- MPLS EXP rewrite rules apply only to logical interfaces on **family mpls** interfaces. You cannot apply to an EXP rewrite rule to a physical interface. You can configure up to 64 EXP rewrite rules, but you can only use 16 EXP rewrite rules at any time on the switch.
- A logical interface can use both DSCP (or DSCP IPv6) and EXP rewrite rules.
- DSCP and DSCP IPv6 rewrite rules are not applied to MPLS traffic.
- If the switch is performing penultimate hop popping (PHP), EXP rewrite rules do not take effect. If both an EXP classifier and an EXP rewrite rule are configured on the switch, then the EXP value from the last popped label is copied into the inner label. If either an EXP classifier or an EXP rewrite rule (but not both) is configured on the switch, then the inner label EXP value is sent unchanged.



**NOTE:** On each physical interface, either all forwarding classes that are being used on the interface must have rewrite rules configured or no forwarding classes that are being used on the interface can have rewrite rules configured. On any physical port, do not mix forwarding classes with rewrite rules and forwarding classes without rewrite rules.



**NOTE:** Rewrite rules are applied *before* the egress filter is matched to traffic. Because the code point rewrite occurs before the egress filter is matched to traffic, the egress filter match is based on the rewrite value, not on the original code point value in the packet.

## Classifier and Rewrite Rule Configuration Interaction with Ethernet Interface Configuration

On QFX5100, QFX5200, EX4600, QFX3500, and QFX3600 switches used as standalone switches or as QFabric system Node devices, you can apply classifiers and rewrite rules

only on Layer 2 logical interface unit 0 and Layer 3 physical interfaces (if the Layer 3 physical interface has at least one defined logical interface). On QFX10000 switches, you can apply classifiers and rewrite rules only to Layer 2 logical interface unit 0 and to Layer 3 logical interfaces. This section focuses on BA classifiers, but the interaction between BA classifiers and interfaces described in this section also applies to fixed classifiers and rewrite rules.



**NOTE:** On QFX5100, QFX5200, EX4600, QFX3500, and QFX3600 switches used as standalone switches or as QFabric system Node devices, EXP classifiers, are global and apply to all switch interfaces. See *Defining CoS Multidestination (Multicast, Broadcast, DLF) BA Classifiers* for how to configure multidestination classifiers and see [“Configuring a Global MPLS EXP Classifier” on page 4977](#) for how to configure EXP classifiers.

On switches that use different classifiers for unicast and multidestination traffic, multidestination classifiers are global and apply to all switch interfaces.

There are two components to applying classifiers or rewrite rules to interfaces:

1. Setting the interface family (inet, inet6, or ethernet-switching; ethernet-switching is the default interface family) in the **[edit interfaces]** configuration hierarchy.
2. Applying a classifier or rewrite rule to the interface in the **[edit class-of-service]** hierarchy.

These are separate operations that can be set and committed at different times. Because the type of classifier or rewrite rule you can apply to an interface depends on the interface family configuration, the system performs checks to ensure that the configuration is valid. The method the system uses to notify you of an invalid configuration depends on the **set** operation that causes the invalid configuration.



**NOTE:** QFX10000 switches cannot be misconfigured in the following two ways because you can configure classifiers only on logical interfaces. Only switches that allow classifier configuration on physical and logical interfaces can experience the following misconfigurations.

If applying the classifier or rewrite rule to the interface in the **[edit class-of-service]** hierarchy causes an invalid configuration, the system rejects the configuration and returns a commit check error.

If setting the interface family in the **[edit interfaces]** configuration hierarchy causes an invalid configuration, the system creates a syslog error message. If you receive the error message, you need to remove the classifier or rewrite rule configuration from the logical interface and apply it to the physical interface, or remove the classifier or rewrite rule configuration from the physical interface and apply it to the logical interface. For classifiers, if you do not take action to correct the error, the system programs the default classifier for the interface family on the interface. (There are no default rewrite rules. If the commit check fails, no rewrite rule is applied to the interface.)

Two scenarios illustrate these situations:

- Applying a classifier to an Ethernet interface causes a commit check error
- Configuring the Ethernet interface family causes a syslog error

These scenarios differ on different switches because some switches support classifiers on physical Layer 3 interfaces but not on logical Layer 3 interfaces, while other switches support classifiers on logical Layer 3 interfaces but not on physical Layer 3 interfaces.

Two scenarios illustrate these situations:

- [QFX5100, QFX5200, EX4600, QFX3500, and QFX3600 Switch Scenarios on page 6683](#)



**NOTE:** Both of these scenarios also apply to fixed classifiers and rewrite rules.

### [QFX5100, QFX5200, EX4600, QFX3500, and QFX3600 Switch Scenarios](#)

The following scenarios also apply the QFX5100, QFX5200, EX4600, QFX3500, and QFX3600 switches when they are used as QFabric system Node devices.

#### *Scenario 1: Applying a Classifier to an Ethernet Interface Causes a Commit Check Error*

In Scenario 1, we set the interface family, and then specify an invalid classifier.

1. Set and commit the interface as a Layer 3 (family **inet**) interface:

```
[edit interfaces]
user@switch# set xe-0/0/20 unit 0 family inet
user@switch# commit
```

This commit operation succeeds.

2. Set and commit a DSCP classifier on the logical interface (this example uses a DSCP classifier named **dscp1**):

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 unit 0 classifiers dscp dscp1
user@switch# commit
```

This configuration is not valid, because it attempts to apply a classifier to a Layer 3 logical interface. Because the failure is caused by the class-of-service configuration and not by the interface configuration, the system rejects the commit operation and issues a commit error, not a syslog message.

Note that the commit operation succeeds if you apply the classifier to the physical Layer 3 interface as follows:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 classifiers dscp dscp1
user@switch# commit
```

Because the logical unit is not specified, the classifier is applied to the physical Layer 3 interface in a valid configuration, and the commit check succeeds.

#### *Scenario 2: Configuring the Ethernet Interface Family Causes a Syslog Error*

In Scenario 2, we set the classifier first, and then set an invalid interface type.

1. Set and commit a DSCP classifier on a logical interface that has no existing configuration:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 unit 0 classifiers dscp dscp1
user@switch# commit
```

This commit succeeds. Because no explicit configuration existed on the interface, it is by default a Layer 2 (**family ethernet-switching**) interface. Layer 2 logical interfaces support BA classifiers, so applying the classifier is a valid configuration.

2. Set and commit the interface as a Layer 3 interface (family **inet**) interface:

```
[edit interfaces]
user@switch# set xe-0/0/20 unit 0 family inet
user@switch# commit
```

This configuration is not valid because it attempts to change an interface from Layer 2 (**family ethernet-switching**) to Layer 3 (**family inet**) when a classifier has already been applied to a logical interface. Layer 3 logical interfaces do not support classifiers. Because the failure is caused by the interface configuration and not by the class-of-service configuration, the system does not issue a commit error, but instead issues a syslog message.

When the system issues the syslog message, it programs the default classifier for the interface type on the interface. In this scenario, the interface has been configured as a Layer 3 interface, so the system applies the default DSCP profile to the physical Layer 3 interface.

In this scenario, to install a configured DSCP classifier, remove the misconfigured classifier from the Layer 3 logical interface and apply it to the Layer 3 physical interface. For example:

```
[edit]
user@switch# delete class-of-service interfaces xe-0/0/20 unit 0 classifiers dscp dscp1
user@switch# commit
user@switch# set class-of-service interfaces xe-0/0/20 classifiers dscp dscp1
user@switch# commit
```

- Related Documentation**
- [Understanding CoS Packet Flow on page 6628](#)
  - [Configuring CoS on page 6618](#)

## Understanding CoS Code-Point Aliases

A code-point alias assigns a name to a pattern of code-point bits. You can use this name instead of the bit pattern when you configure other CoS components such as classifiers and rewrite rules.



**NOTE:** This topic applies to all EX Series switches except the EX4600. Because the EX4600 uses a different chipset than other EX Series switches, the code-point aliases on EX4600 match those on QFX Series switches. For EX4600 code-point aliases, see [“Understanding CoS Code-Point Aliases” on page 6685](#).

Behavior aggregate classifiers use class-of-service (CoS) values such as Differentiated Services Code Points (DSCPs) or IEEE 802.1 bits to associate incoming packets with a particular forwarding class and the CoS servicing level associated with that forwarding class. You can assign a meaningful name or alias to the CoS values and use that alias instead of bits when configuring CoS components. These aliases are not part of the specifications but are well known through usage. For example, the alias for DSCP 101110 is widely accepted as ef (expedited forwarding).

When you configure forwarding classes and define classifiers, you can refer to the markers by alias names. You can configure code point alias names for user-defined classifiers. If the value of an alias changes, it alters the behavior of any classifier that references it.

You can configure code-point aliases for the following type of CoS markers:

- dscp or dscp-ipv6—Handles incoming IP and IPv6 packets.
- ieee-802.1—Handles Layer 2 frames.

[Table 556](#) shows the default mapping of code-point aliases to IEEE code points.

**Table 556: Default IEEE 802.1 Code-Point Aliases**

| CoS Value Types | Mapping |
|-----------------|---------|
| be              | 000     |
| be1             | 001     |
| ef              | 010     |
| ef1             | 011     |
| af11            | 100     |
| af12            | 101     |
| nc1             | 110     |

**Table 556: Default IEEE 802.1 Code-Point Aliases (*continued*)**

| CoS Value Types | Mapping |
|-----------------|---------|
| nc2             | 111     |

Table 557 shows the default mapping of code-point aliases to DSCP and DSCP IPv6 code points.

**Table 557: Default DSCP and DSCP IPv6 Code-Point Aliases**

| CoS Value Types | Mapping |
|-----------------|---------|
| ef              | 101110  |
| af11            | 001010  |
| af12            | 001100  |
| af13            | 001110  |
| af21            | 010010  |
| af22            | 010100  |
| af23            | 010110  |
| af31            | 011010  |
| af32            | 011100  |
| af33            | 011110  |
| af41            | 100010  |
| af42            | 100100  |
| af43            | 100110  |
| be              | 000000  |
| cs1             | 001000  |
| cs2             | 010000  |
| cs3             | 011000  |
| cs4             | 100000  |
| cs5             | 101000  |

Table 557: Default DSCP and DSCP IPv6 Code-Point Aliases (*continued*)

| CoS Value Types | Mapping |
|-----------------|---------|
| nc1             | 110000  |
| nc2             | 111000  |

- Related Documentation**
- [Understanding Junos CoS Components on page 6621](#)
  - [Defining CoS Code-Point Aliases on page 6687](#)

## Defining CoS Code-Point Aliases

You can use code-point aliases to streamline the process of configuring CoS features on your switch. A code-point alias assigns a name to a pattern of code-point bits. You can use this name instead of the bit pattern when you configure other CoS components such as classifiers and rewrite rules.

You can configure code-point aliases for the following CoS marker types:

- DSCP or DSCP IPv6—Handles incoming IPv4 or IPv6 packets.
- IEEE 802.1p—Handles Layer 2 frames.

To configure a code-point alias:

1. Specify a CoS marker type (IEEE 802.1 or DSCP).
2. Assign an alias.
3. Specify the code point that corresponds to the alias.

```
[edit class-of-service code-point-aliases]
user@switch# set (dscp | dscp-ipv6 | ieee-802.1) alias-name code-point-bits
```

For example, to configure a code-point alias for an IEEE 802.1 CoS marker type that has the alias name be2 and maps to the code-point bits 001:

```
[edit class-of-service code-point-aliases]
user@switch# set ieee-802.1 be2 001
```

- Related Documentation**
- [Monitoring CoS Code-Point Value Aliases on page 7236](#)
  - [Understanding CoS Code-Point Aliases on page 6685](#)

## Understanding CoS Forwarding Classes

Forwarding classes group traffic and assign the traffic to output queues. Each forwarding class is mapped to an output queue. Classification maps incoming traffic to forwarding classes based on the code point bits in the packet or frame header. Forwarding class to queue mapping defines the output queue used for the traffic classified into a forwarding class.

A classifier must associate each packet with one of the following four default forwarding classes or with a user-configured forwarding class to assign an output queue to the packet:

- **fcoe**—Guaranteed delivery for Fibre Channel over Ethernet (FCoE) traffic.
- **no-loss**—Guaranteed delivery for TCP lossless traffic.
- **best-effort**—Provides best-effort delivery without a service profile. Loss priority is typically not carried in a class-of-service (CoS) value.
- **network-control**—Supports protocol control.

The switch supports up to eight forwarding classes, thus enabling flexible, differentiated, packet classification. For example, you can configure multiple classes of best-effort traffic such as **best-effort**, **best-effort1**, and **best-effort2**.

The switch supports eight output queues (queues 0 through 7). Unicast and multdestination (multicast, broadcast, and destination lookup fail) traffic use the same forwarding classes and output queues.

- [Default Forwarding Classes on page 6688](#)
- [Forwarding Class Configuration Rules on page 6689](#)
- [Lossless Transport Support on page 6691](#)

### Default Forwarding Classes

[Table 558](#) shows the four default forwarding classes. You can rename the forwarding classes. Assigning a new forwarding class name does not alter the default classification or scheduling applied to the queue that is mapped to that forwarding class. CoS configurations can be complex, so unless it is required by your scenario, we recommend that you use the default class names and queue number associations.

**Table 558: Default Forwarding Classes**

| Forwarding Class Name | Default Queue Mapping | Comments                                                                                                                                                                                                                                                                                                                            |
|-----------------------|-----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| best-effort           | 0                     | <p>The software does not apply any special CoS handling to best-effort traffic. This is a backward compatibility feature. Best-effort traffic is usually the first traffic to be dropped during periods of network congestion.</p> <p>By default, this is a lossy forwarding class with a packet drop attribute of <b>drop</b>.</p> |



Table 558: Default Forwarding Classes (*continued*)

| Forwarding Class Name | Default Queue Mapping | Comments                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-----------------------|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| fcoe                  | 3                     | <p>By default, the <b>fcoe</b> forwarding class is a lossless forwarding class designed to handle Fibre Channel over Ethernet (FCoE) traffic. The <b>no-loss</b> packet drop attribute is applied by default.</p> <p><b>NOTE:</b> By convention, deployments with converged server access typically use IEEE 802.1p priority 3 (011) for FCoE traffic. The default mapping of the <b>fcoe</b> forwarding class is to queue 3. Apply priority-based flow control (PFC) to the entire FCoE data path to configure the end-to-end lossless behavior that FCoE requires.</p> <p>We recommend that you use priority 3 for FCoE traffic unless your network architecture requires that you use a different priority.</p> |
| no-loss               | 4                     | <p>By default, this is a lossless forwarding class with a packet drop attribute of <b>no-loss</b>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| network-control       | 7                     | <p>The software delivers packets in this service class with a high priority. (These packets are not delay-sensitive.)</p> <p>Typically, these packets represent routing protocol hello or keepalive messages. Because loss of these packets jeopardizes proper network operation, packet delay is preferable to packet discard.</p> <p>By default, this is a lossy forwarding class with a packet drop attribute of <b>drop</b>.</p>                                                                                                                                                                                                                                                                               |

## Forwarding Class Configuration Rules

Take the following rules into account when you configure forwarding classes:

- [Queue Assignment Rules on page 6689](#)
- [Scheduling Rules on page 6690](#)
- [Rewrite Rules on page 6691](#)

### Queue Assignment Rules

The following rules govern queue assignment:

- CoS configurations that specify more queues than the switch can support are not accepted. The commit operation fails with a detailed message that states the total number of queues available.
- All default CoS configurations are based on queue number. The name of the forwarding class that appears in the default configuration is the forwarding class currently mapped to that queue.
- If you map more than one forwarding class to a queue, all of the forwarding classes mapped to the same queue must have the same packet drop attribute (all of the

forwarding classes must be lossy, or all of the forwarding classes mapped to a queue must be lossless).

You limit the amount of traffic that receives strict-high priority treatment on a strict-high priority queue by configuring a transmit rate. The transmit rate sets the amount of traffic on the queue that receives strict-high priority treatment. The switch treats traffic that exceeds the transmit rate as low priority traffic that receives the queue excess rate bandwidth. Limiting the amount of traffic that receives strict-high priority treatment prevents other queues from being starved while also ensuring that the amount of traffic specified in the transmit rate receives strict-high priority treatment.

If you configure one strict-high priority queue on a port, we strongly recommend that you configure a transmit rate on the queue to prevent it from starving low priority queues on that port. Although it is not mandatory to configure a transmit rate on a strict-high priority queue when there is only one strict-high priority queue on a port, if you do not configure a transmit rate, the strict-high priority queue can consume all of the port bandwidth and starve the other queues.

If you configure more than one strict-high priority queue on a port, you must configure a transmit rate on each of the strict-high priority queues. If you configure more than one strict-high priority queue on a port and you do not configure a transmit rate on the strict-high priority queues, the switch treats only the first queue you configure as a strict-high priority queue. The switch treats the other queues as low priority queues. If you configure a transmit rate on some strict-high priority queues but not on other strict-high priority queues on a port, the switch treats the queues that have a transmit rate as strict-high priority queues, and treats the queues that do not have a transmit rate as low priority queues.

---

### Scheduling Rules

---

When you configure a forwarding class and map traffic to it (that is, you are not using a default classifier and forwarding class), you must also define a scheduling policy for the forwarding class. You can configure either port scheduling or enhanced transmission selection (ETS) hierarchical port scheduling.

Defining a scheduling policy using port scheduling means:

- Mapping a scheduler to the forwarding class in a scheduler map.
- Applying the scheduler map to one or more interfaces.

Defining a scheduling policy using ETS means:

- Mapping a scheduler to the forwarding class in a scheduler map.
- Including the forwarding class in a forwarding class set.
- Associating the scheduler map with a traffic control profile.
- Attaching the traffic control profile to a forwarding class set and applying the traffic control profile to an interface.

## Rewrite Rules

---

On each physical interface, either all forwarding classes that are being used on the interface must have rewrite rules configured, or no forwarding classes that are being used on the interface can have rewrite rules configured. On any physical port, do not mix forwarding classes with rewrite rules and forwarding classes without rewrite rules.

## Lossless Transport Support

The switch supports up to six lossless forwarding classes. For lossless transport, you must enable PFC on the IEEE 802.1p code point of lossless forwarding classes. The following limitations apply to support lossless transport:

- The external cable length from the switch to other devices cannot exceed 300 meters.
- For FCoE traffic, the interface maximum transmission unit (MTU) must be at least 2180 bytes to accommodate the packet payload, headers, and checks.
- Changing any portion of a PFC configuration on a port blocks the entire port until the change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Changing the PFC configuration means any change to a congestion notification profile that is configured on a port (enabling or disabling PFC on a code point, changing the MRU or cable-length value, or specifying an output flow control queue). Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

### Related Documentation

- [Understanding Junos CoS Components on page 6621](#)
- [Understanding CoS Packet Flow on page 6628](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 6493](#)
- [Example: Configuring Forwarding Classes on page 6692](#)

## Defining CoS Forwarding Classes

---

Forwarding classes allow you to group packets for transmission. The switch supports a total of eight forwarding classes. To forward traffic, you map (assign) the forwarding classes to output queues.

The switch has eight output queues, queues 0 through 7. The default forwarding classes are:

- **best-effort**—Best-effort traffic
- **fcoe**—Guaranteed delivery for Fibre Channel over Ethernet traffic
- **no-loss**—Guaranteed delivery for TCP no-loss traffic
- **network-control**—Network control traffic

The NFX250 Network Services platform has the following default forwarding classes:

- best-effort (be)—Provides no service profile. Loss priority is typically not carried in a CoS value.
- expedited-forwarding (ef)—Provides a low loss, low latency, low jitter, assured bandwidth, end-to-end service.
- assured-forwarding (af)—Provides a group of values you can define and includes four subclasses: AF1, AF2, AF3, and AF4, each with two drop probabilities: low and high.
- network-control (nc)—Supports protocol control and thus is typically high priority.

Map forwarding classes to queues using the **class** statement. You can map more than one forwarding class to a single queue, but all forwarding classes mapped to a particular queue must have the same packet drop attribute (all of the forwarding classes must be lossy, or all of the forwarding classes mapped to a queue must be lossless).

```
[edit class-of-service forwarding-classes]
```

```
user@switch# set class class-name queue-num queue-number <no-loss>
```

For example, to create a forwarding class named **be2** and map it to queue 1:

```
[edit class-of-service forwarding-classes]
```

```
user@switch# set class be2 queue-num 1
```

Another example is to create a lossless forwarding class named **fcoe2** and map it to queue 5:

```
[edit class-of-service forwarding-classes]
```

```
user@switch# set class fcoe2 queue-num 5 no-loss
```

#### Related Documentation

- [Example: Configuring Forwarding Classes on page 6692](#)
- [Monitoring CoS Forwarding Classes on page 7232](#)
- [Understanding CoS Forwarding Classes on page 6688](#)
- [Understanding CoS Port Schedulers on QFX Switches on page 6729](#)

---

## Example: Configuring Forwarding Classes

Forwarding classes group packets for transmission. Forwarding classes map to output queues, so the packets assigned to a forwarding class use the output queue mapped to that forwarding class.

- [Requirements on page 6692](#)
- [Overview on page 6693](#)
- [Configuration on page 6694](#)
- [Verification on page 6694](#)

### Requirements

This example uses the following hardware and software components:

- One switch.
- Junos OS Release 15.1X53-D10 or later for the QFX Series.

## Overview

The switch supports a total of eight forwarding classes; you can configure up to six forwarding classes as lossless forwarding classes. To forward traffic, you must map (assign) the forwarding classes to output queues. The switch has eight output queues, queues 0 through 7. The switch provides four default forwarding classes. You can define the remaining four forwarding classes and map them to output queues.

The four default forwarding classes are:

- **be**—Best-effort traffic
- **fcoe**—Guaranteed delivery for Fibre Channel over Ethernet traffic
- **no-loss**—Guaranteed delivery for TCP no-loss traffic
- **nc**—Network control traffic

Map forwarding classes to queues using the **class** statement. You can map more than one forwarding class to a single queue, but all forwarding classes mapped to a particular queue must have the same packet drop attribute (all of the forwarding classes must be lossy, or all of the forwarding classes mapped to a queue must be lossless). The statement format is:

```
[edit class-of-service forwarding-classes]
user@switch# set class class-name queue-num queue-number;
```



**NOTE:** When you configure a forwarding class and map traffic to it (that is, you are not using a default classifier and forwarding class), you must also define a scheduling policy for the forwarding class. You can configure either port scheduling or enhanced transmission selection (ETS) hierarchical port scheduling.

Defining a scheduling policy using port scheduling means:

- Mapping a scheduler to the forwarding class in a scheduler map.
- Applying the scheduler map to one or more interfaces.

Defining a scheduling policy using ETS means:

- Mapping a scheduler to the forwarding class in a scheduler map.
- Including the forwarding class in a forwarding class set.
- Associating the scheduler map with a traffic control profile.
- Attaching the traffic control profile to a forwarding class set and applying the traffic control profile to an interface.

Table 559 shows the configuration forwarding-class-to-queue mapping for this example:

Table 559: Forwarding-Class-to-Queue Example Configuration

| Forwarding Class | Queue |
|------------------|-------|
| best-effort      | 0     |
| be1              | 1     |
| nc               | 7     |

## Configuration

To configure CoS forwarding classes:

1. Map the **best-effort** forwarding class to queue 0:

```
[edit class-of-service forwarding-classes]
user@switch# set class best-effort queue-num 0
```

2. Map the **be1** forwarding class to queue 1:

```
[edit class-of-service forwarding-classes]
user@switch# set class be1 queue-num 1
```

3. Map the **nc** forwarding class to queue 7:

```
[edit class-of-service forwarding-classes]
user@switch# set class nc queue-num 7
```

## Verification

### Verifying the Forwarding-Class-to-Queue Mapping

**Purpose** Verify the forwarding-class-to-queue mapping. (The system shows only the explicitly configured forwarding classes; it does not show default forwarding classes such as **fcoe** and **no-loss**.)

**Action** Verify the results of the forwarding class configuration using the operational mode command **show configuration class-of-service forwarding-classes**:

```
user@switch> show configuration class-of-service forwarding-classes
class best-effort queue-num 0;
class be1 queue-num 1;
class network-control queue-num 7;
```

**Related Documentation**

- [Defining CoS Forwarding Classes on page 6691](#)
- [Monitoring CoS Forwarding Classes on page 7232](#)
- [Understanding CoS Forwarding Classes on page 6688](#)

## Understanding CoS Forwarding Class Sets (Priority Groups)

---

A forwarding class set is the Junos OS configuration construct that equates to a priority group in enhanced transmission selection (ETS, described in IEEE 802.1Qaz). The switch implements ETS using a two-tier hierarchical scheduler.

A priority group is a group of forwarding classes. Each forwarding class is mapped to an output queue and an IEEE 802.1p priority (code points). Classifying traffic into a forwarding class based on its code points, and mapping the forwarding class to a queue, defines the traffic assigned to that queue. The forwarding classes that belong to a priority group share the port bandwidth allocated to that priority group. The traffic mapped to forwarding classes in one priority group usually shares similar traffic-handling requirements.

You can configure up to three unicast forwarding class sets and one multicast forwarding class set. Only unicast forwarding classes can belong to unicast forwarding class sets. Only multicast forwarding classes can belong to the multicast forwarding class set.

If you configure a strict-high priority forwarding class (you can configure only one strict-high priority forwarding class), you must observe the following rules when configuring forwarding class sets:

- You must create a separate forwarding class set for the strict-high priority forwarding class.
- Only one forwarding class set can contain the strict-high priority forwarding class.
- A strict-high priority forwarding class cannot belong to the same forwarding class set as forwarding classes that are not strict-high priority.
- A strict-high priority forwarding class cannot belong to a multidestination forwarding class set.
- You cannot configure a guaranteed minimum bandwidth (guaranteed rate) for a forwarding class set that includes a strict-high priority forwarding class. (You also cannot configure a guaranteed minimum bandwidth for a strict-high forwarding class.)
- We recommend that you always apply a shaping rate to a strict-high priority forwarding class to prevent it from starving the queues mapped to other forwarding classes. If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority forwarding class can use, then the strict-high priority forwarding class can use all of the available port bandwidth and starve other forwarding classes on the port.

You must use hierarchical scheduling if you explicitly configure CoS. The two-tier hierarchical scheduler defines bandwidth resources for the forwarding class set (priority group), and then allocates those resources among the forwarding classes (priorities) that belong to the forwarding class set.

If you do not explicitly configure forwarding class sets, the system automatically creates a default forwarding class set that contains all of the forwarding classes on the switch. The system assigns 100 percent of the port output bandwidth to the default forwarding class set. Ingress traffic is classified based on the default classifier settings. The forwarding classes in the default forwarding class set receive bandwidth based on the default

scheduler settings. Forwarding classes that are not part of the default scheduler receive no bandwidth. The default priority group is transparent. It does not appear in the configuration and is used for Data Center Bridging Capability Exchange Protocol (DCBX) advertisement (except on OCX Series switches, which do not support DCBX).

When you explicitly configure forwarding class sets and apply them to interfaces, on those interfaces, forwarding classes that you do not map to a forwarding class set receive no guaranteed bandwidth. Forwarding classes that belong to the default forwarding class set might receive bandwidth if the other forwarding class sets are not using all of the port bandwidth. However, the amount of bandwidth received by forwarding classes that are not members of a forwarding class set is not guaranteed. In this case, the bandwidth a forwarding class receives if it is not a member of a forwarding class set depends on whether unused port bandwidth is available and therefore is not deterministic.

To guarantee bandwidth for forwarding classes in a predictable manner, be sure to map all forwarding classes that you expect to carry traffic on an interface to a forwarding class set, and apply the forwarding class set to the interface.

**Related  
Documentation**

- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 6765](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 6771](#)
- [Example: Configuring Forwarding Class Sets on page 6697](#)
- [Defining CoS Forwarding Class Sets on page 6696](#)

---

## Defining CoS Forwarding Class Sets

---

A forwarding class set is a priority group for enhanced transmission selection (ETS) traffic control. Each forwarding class set consists of one or more forwarding classes. Classifiers map traffic into forwarding classes based on code points (priority), and forwarding classes are mapped to output queues.

You can configure up to three unicast forwarding class sets and one multicast forwarding class set.

To configure a forwarding class set using the CLI:

1. Assign one or more forwarding classes to the forwarding class set:

```
[edit class-of-service]
user@switch# set forwarding-class-sets forwarding-class-set-name class
forwarding-class-name
```

2. Map the forwarding class set to an interface:

```
[edit class-of-service]
user@switch# set interfaces interface-name forwarding-class-set forwarding-class-set-name
```

**Related  
Documentation**

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 6771](#)
- [Example: Configuring Forwarding Class Sets on page 6697](#)
- [Defining CoS Queue Schedulers](#)



- [Defining CoS Traffic Control Profiles \(Priority Group Scheduling\) on page 6761](#)
- [Understanding CoS Forwarding Class Sets \(Priority Groups\) on page 6695](#)

## Example: Configuring Forwarding Class Sets

A forwarding class set (fc-set) is a priority group for enhanced transmission selection (ETS) traffic control. Each fc-set consists of one or more forwarding classes (priorities). Classifiers map traffic to forwarding classes based on code points, and forwarding classes are mapped to output queues.

ETS enables you to configure link resources (bandwidth and bandwidth sharing characteristics) for an fc-set, and then allocate the fc-set's resources among the forwarding classes that belong to the fc-set. This is called two-tier, or hierarchical, scheduling. Traffic control profiles control the scheduling for the fc-set (priority group), and schedulers control the scheduling for individual forwarding classes (priorities).

- [Requirements on page 6697](#)
- [Overview on page 6697](#)
- [Configuring Forwarding Class Sets on page 6698](#)
- [Verification on page 6699](#)

### Requirements

This example uses the following hardware and software components:

- One switch (this example was tested on a Juniper Networks QFX3500 Switch)
- Junos OS Release 11.1 or later for the QFX Series or Junos OS Release 14.1X53-D20 or later for the OCX Series.

### Overview

You can configure up to three unicast fc-sets and one multicast fc-set. A common way to configure unicast priority groups is to configure separate fc-sets for local area network (LAN) traffic, storage area network (SAN) traffic, and high-performance computing (HPC) traffic, and then assign the appropriate forwarding classes to each fc-set.



**NOTE:** If you configure a strict-high priority forwarding class, you must create an fc-set that is dedicated only to strict-high priority traffic. You can only configure one strict-high priority forwarding class, and only one fc-set can contain a strict-high priority queue. Queues that are not strict-high priority cannot belong to the same fc-set as a strict-high priority queue. The multdestination fc-set cannot contain a strict-high priority queue.

To apply ETS, you use a traffic control profile to map one or more fc-sets to a physical egress port. You can map up to three unicast forwarding class sets and one multdestination forwarding class set to each port. When you map an fc-set to a port,

the port uses hierarchical scheduling to allocate port resources to the priority group (fc-set) and to allocate the priority group resources to the queues (forwarding classes) that belong to the priority group.

This example describes how to:

- Configure three fc-sets called **lan-pg**, **san-pg**, and **hpc-pg**.
- Assign forwarding classes to each of the fc-sets.
- Apply the fc-sets and their output traffic control profiles to an egress interface.

This example does not describe how to configure the forwarding classes assigned to the fc-sets or how to configure traffic control profiles (scheduling). [“Example: Configuring CoS Hierarchical Port Scheduling \(ETS\)” on page 6771](#) provides a complete example of how to configure ETS, including forwarding class and scheduling configuration. [Table 560](#) shows the configuration components for this example:

**Table 560: Components of the Forwarding Class Sets Configuration Example**

| Component                  | Settings                                                                                                                                                                                                                                                                                                                                             |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Hardware                   | QFX3500 switch                                                                                                                                                                                                                                                                                                                                       |
| LAN traffic priority group | Forwarding class set: <b>lan-pg</b><br>Forwarding classes: <b>best-effort-1</b> , <b>best-effort-2</b>                                                                                                                                                                                                                                               |
| SAN traffic priority group | Forwarding class set: <b>san-pg</b><br>Forwarding classes: <b>fcoe</b> , <b>fcoe-2</b><br><br><b>NOTE:</b> OCX Series switches do not support FCoE traffic or lossless Layer 2 transport. If you were configuring this example on an OCX Series switch, you could omit this priority group, or rename it and map different forwarding classes to it. |
| HPC traffic priority group | Forwarding class set: <b>hpc-pg</b><br>Forwarding classes: <b>nc</b> , <b>high-perf</b>                                                                                                                                                                                                                                                              |
| Egress interface           | <b>xe-0/0/7</b>                                                                                                                                                                                                                                                                                                                                      |

## Configuring Forwarding Class Sets

1. Define the **lan-pg** priority group (fc-set) and assign to it the forwarding classes **best-effort-1** and **best-effort-2**:  

```
[edit class-of-service]
user@switch# set forwarding-class-sets lan-pg class best-effort-1
user@switch# set forwarding-class-sets lan-pg class best-effort-2
```
2. Define the **san-pg** priority group and assign to it the forwarding classes **fcoe** and **fcoe-2**:  

```
[edit class-of-service]
user@switch# set forwarding-class-sets san-pg class fcoe
user@switch# set forwarding-class-sets san-pg class fcoe-2
```
3. Define the **hpc-pg** priority group and assign to it the forwarding classes **nc** and **high-perf**:

```
[edit class-of-service]
user@switch# set forwarding-class-sets hpc-pg class nc
user@switch# set forwarding-class-sets hpc-pg class high-perf
```

4. Map the three forwarding class sets to an interface (the output traffic control profiles associated with the forwarding class sets determine the class of service scheduling for the priority groups):

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/7 forwarding-class-set lan-pg
output-traffic-control-profile lan-tcp
user@switch# set interfaces xe-0/0/7 forwarding-class-set san-pg
output-traffic-control-profile san-tcp
user@switch# set interfaces xe-0/0/7 forwarding-class-set hpc-pg
output-traffic-control-profile hpc-tcp
```

## Verification

To verify the priority group configuration, perform these tasks:

- [Verifying Forwarding Class Set Membership on page 6699](#)
- [Verifying the Egress Interface Configuration on page 6699](#)

### Verifying Forwarding Class Set Membership

**Purpose** Verify that you configured the **lan-pg**, **san-pg**, and **hpc-pg** priority groups with the correct forwarding classes.

**Action** List the forwarding class set member configuration using the operational mode command **show configuration class-of-service forwarding-class-sets**:

```
user@switch> show configuration class-of-service forwarding-class-sets
lan-pg {
 class best-effort-1;
 class best-effort-2;
}
san-pg {
 class fcoe;
 class fcoe-2;
}
hpc-pg {
 class high-perf;
 class nc;
}
```

### Verifying the Egress Interface Configuration

**Purpose** Verify that egress interface **xe-0/0/7** is associated with the **lan-pg**, **san-pg**, and **hpc-pg** priority groups and with the correct output traffic control profiles.

**Action** Display the egress interface using the operational mode command **show configuration class-of-service interfaces xe-0/0/7**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/7
forwarding-class-set {
 lan-pg {
```

```
 output-traffic-control-profile lan-tcp;
 }
 san-pg {
 output-traffic-control-profile san-tcp;
 }
 hpc-pg {
 output-traffic-control-profile hpc-tcp;
 }
}
```

**Related  
Documentation**

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 6771](#)
- [Example: Configuring Queue Schedulers](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 6762](#)
- [Defining CoS Forwarding Class Sets on page 6696](#)
- [Understanding CoS Forwarding Class Sets \(Priority Groups\) on page 6695](#)

## Understanding CoS Rewrite Rules

As packets enter or exit a network, edge switches might be required to alter the class-of-service (CoS) settings of the packets. Rewrite rules set the value of the code point bits (Layer 3 DSCP bits, Layer 2 CoS bits, or MPLS EXP bits) within the header of the outgoing packet. Each rewrite rule:

1. Reads the current forwarding class and loss priority associated with the packet.
2. Locates the new (rewrite) code point value from a table.
3. Writes that code point value into the packet header, replacing the old code point value.

Rewrite rules must be assigned to an interface for rewrites to take effect.

You can apply (bind) one DSCP or DSCP IPv6 rewrite rule and one IEEE 802.1p rewrite rule to each interface. You can also bind EXP rewrite rules to **family mpls** logical interfaces to rewrite the CoS bits of MPLS traffic.



**NOTE:** OCX Series switches do not support MPLS and do not support EXP rewrite rules.

You cannot apply both a DSCP and a DSCP IPv6 rewrite rule to the same physical interface. Each physical interface supports only one DSCP rewrite rule. Both IP and IPv6 packets use the same DSCP rewrite rule, regardless if the configured rewrite rule is DSCP or DSCP IPv6. You can apply an EXP rewrite rule on an interface that has DSCP or IEEE rewrite rules. Only MPLS traffic on **family mpls** interfaces uses the EXP rewrite rule.

You *can* apply both a DSCP rewrite rule and a DSCP IPv6 rewrite rule to a logical interface. IPv6 packets are rewritten with DSCP-IPv6 rewrite-rules and IPv4 packets are remarked with DSCP rewrite-rules.



**NOTE:** There are no default rewrite rules. If you want to apply a rewrite rule to outgoing packets, you must explicitly configure the rewrite rule.

You can look at behavior aggregate (BA) classifiers and rewrite rules as two sides of the same coin. A BA classifier reads the code point bits of incoming packets and classifies the packets into forwarding classes, then the system applies the CoS configured for the forwarding class to those packets. Rewrite rules change (rewrite) the code point bits just before the packets leave the system so that the next switch or router can apply the appropriate level of CoS to the packets. When you apply a rewrite rule to an interface, the rewrite rule is the last CoS action performed on the packet before it is forwarded.

Rewrite rules alter CoS values in outgoing packets on the outbound interfaces of an edge switch to accommodate the policies of a targeted peer. This allows the downstream switch in a neighboring network to classify each packet into the appropriate service group.



**NOTE:** On each physical interface, either all forwarding classes that are being used on the interface must have rewrite rules configured or no forwarding classes that are being used on the interface can have rewrite rules configured. On any physical port, do not mix forwarding classes with rewrite rules and forwarding classes without rewrite rules.



**NOTE:** Rewrite rules are applied *before* the egress filter is matched to traffic. Because the code point rewrite occurs before the egress filter is matched to traffic, the egress filter match is based on the rewrite value, not on the original code point value in the packet.

For packets that carry both an inner VLAN tag and an outer VLAN tag, the rewrite rule rewrites only the outer VLAN tag.

MPLS EXP rewrite rules apply only to **family mpls** logical interfaces. You cannot apply to an EXP rewrite rule to a physical interface. You can configure up to 64 EXP rewrite rules, but you can only use 16 EXP rewrite rules at any time on the switch. On a given logical interface, all pushed MPLS labels have the same EXP rewrite rule applied to them. You can apply different EXP rewrite rules to different logical interfaces on the same physical interface.



**NOTE:** If the switch is performing penultimate hop popping (PHP), EXP rewrite rules do not take effect. If both an EXP classifier and an EXP rewrite rule are configured on the switch, then the EXP value from the last popped label is copied into the inner label. If either an EXP classifier or an EXP rewrite rule (but not both) is configured on the switch, then the inner label EXP value is sent unchanged.

You can configure enough rewrite rules to handle most, if not all, network scenarios. [Table 561](#) shows how many of each type of rewrite rules you can configure, and how many entries you can configure per rewrite rule.

**Table 561: Configuring Rewrite Rules**

| Rewrite Rule Type | Maximum Number of Rewrite Rules | Maximum Number of Entries per Rewrite Rule |
|-------------------|---------------------------------|--------------------------------------------|
| IEEE 802.1p       | 64                              | 128                                        |
| DSCP              | 32                              | 128                                        |
| DSCP IPv6         | 32                              | 128                                        |
| MPLS EXP          | 64                              | 128                                        |

You cannot apply rewrite rules directly to integrated routing and bridging (IRB), also known as routed VLAN interfaces (RVIs), because the members of IRBs/RVIs are VLANs, not ports. However, you can apply rewrite rules to the VLAN port members of an IRB/RVI.



**NOTE:** OCX Series switches do not support IRBs/RVIs.

**Related  
Documentation**

- [Understanding Junos CoS Components on page 6621](#)
- [Defining CoS Rewrite Rules on page 6704](#)
- [Configuring Rewrite Rules for MPLS EXP Classifiers on page 4992](#)

## Defining CoS Rewrite Rules

Edge switches might need to change the class-of-service (CoS) settings of the packets. You can configure rewrite rules to alter code point bit values in outgoing packets on the outbound interfaces of a switch so that the CoS treatment matches the policies of a targeted peer. Policy matching allows the downstream routing platform or switch in a neighboring network to classify each packet into the appropriate service group.

To configure a CoS rewrite rule, create the rule by giving it a name and associating it with a forwarding class, loss priority, and code point. This creates a rewrite table. After the rewrite rule is created, enable it on an interface (EXP rewrite rules can only be enabled on **family mpls** logical interfaces, not on physical interfaces). You can also apply an existing rewrite rule on an interface.



**NOTE:** OCX Series switches do not support MPLS, so they do not support EXP rewrite rules.



**NOTE:** On each physical interface, either all forwarding classes that are being used on the interface must have rewrite rules configured, or no forwarding classes that are being used on the interface can have rewrite rules configured. On any physical port, do not mix forwarding classes with rewrite rules and forwarding classes without rewrite rules.



**NOTE:** To replace an existing rewrite rule on the interface with a new rewrite rule of the same type, first explicitly remove the existing rewrite rule and then apply the new rule.



**NOTE:** For packets that carry both an inner VLAN tag and an outer VLAN tag, the rewrite rule rewrites only the outer VLAN tag.

To create rewrite rules and enable them on interfaces:

- To create an 802.1p rewrite rule named **customup-rw** in the rewrite table for all Layer 2 interfaces:

```
[edit class-of-service rewrite-rules]
user@switch# set ieee-802.1 customup-rw forwarding-class be loss-priority low code-point
000
user@switch# set ieee-802.1 customup-rw forwarding-class be loss-priority high code-point
001
user@switch# set ieee-802.1 customup-rw forwarding-class be loss-priority low code-point
010
user@switch# set ieee-802.1 customup-rw forwarding-class fcse loss-priority low code-point
011
```



```

user@switch# set ieee-802.1 customup-rw forwarding-class ef-no-loss loss-priority low
code-point 100
user@switch# set ieee-802.1 customup-rw forwarding-class ef-no-loss loss-priority high
code-point 101
user@switch# set ieee-802.1 customup-rw forwarding-class nc loss-priority low code-point
110
user@switch# set ieee-802.1 customup-rw forwarding-class nc loss-priority high code-point
111

```

- To enable an 802.1p rewrite rule named **customup-rw** on a Layer 2 interface:

```

[edit]
user@switch# set class-of-service interfaces xe-0/0/7 unit 0 rewrite-rules ieee-802.1
customup-rw

```



**NOTE:** All forwarding classes assigned to port xe-0/0/7 must have rewrite rules. Do not mix forwarding classes that have rewrite rules with forwarding classes that do not have rewrite rules on the same physical interface.

- To enable an 802.1p rewrite rule named **customup-rw** on all 10-Gigabit Ethernet interfaces on the switch, use wildcards for the interface name and logical interface (unit) number:

```

[edit]
user@switch# set class-of-service interfaces xe-* unit * rewrite-rules customup-rw

```



**NOTE:** In this case, *all* forwarding classes assigned to *all* 10-Gigabit Ethernet ports must have rewrite rules. Do not mix forwarding classes that have rewrite rules with forwarding classes that do not have rewrite rules on the same physical interface.

#### Related Documentation

- [Monitoring CoS Rewrite Rules on page 7235](#)
- [Configuring Rewrite Rules for MPLS EXP Classifiers on page 4992](#)
- [Understanding CoS Rewrite Rules on page 6701](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4959](#)

## Troubleshooting an Unexpected Rewrite Value

**Problem**    **Description:** Traffic from one or more forwarding classes on an egress port is assigned an unexpected rewrite value.



**NOTE:** For packets that carry both an inner VLAN tag and an outer VLAN tag, the rewrite rules rewrite only the outer VLAN tag.

**Cause** If you configure a rewrite rule for a forwarding class on an egress port, but you do not configure a rewrite rule for every forwarding class on that egress port, then the forwarding classes that do not have a configured rewrite rule are assigned random rewrite values.

For example:

1. Configure forwarding classes **fc1**, **fc2**, and **fc3**.
2. Configure rewrite rules for forwarding classes **fc1** and **fc2**, but not for forwarding class **fc3**.
3. Assign forwarding classes **fc1**, **fc2**, and **fc3** to a port.

When traffic for these forwarding classes flows through the port, traffic for forwarding classes **fc1** and **fc2** is rewritten correctly. However, traffic for forwarding class **fc3** is assigned a random rewrite value.

**Solution** If any forwarding class on an egress port has a configured rewrite rule, then all forwarding classes on that egress port must have a configured rewrite rule. Configuring a rewrite rule for any forwarding class that is assigned a random rewrite value solves the problem.



**TIP:** If you want the forwarding class to use the same code point value assigned to it by the ingress classifier, specify that value as the rewrite rule value. For example, if a forwarding class has the IEEE 802.1 ingress classifier code point value 011, configure a rewrite rule for that forwarding class that uses the IEEE 802.1p code point value 011.



**NOTE:** There are no default rewrite rules. You can bind one rewrite rule for DSCP traffic and one rewrite rule for IEEE 802.1p traffic to an interface. A rewrite rule can contain multiple forwarding-class-to-rewrite-value mappings.

1. To assign a rewrite value to a forwarding class, add the new rewrite value to the same rewrite rule as the other forwarding classes on the port:

```
[edit class-of-service rewrite-rules]
user@switch# set (dscp|ieee-802.1)rewrite-name forwarding-class class-name loss-priority
priority code-point (alias | bits)
```

For example, if the other forwarding classes on the port use rewrite values defined in the rewrite rule **custom-rw**, the forwarding class **be2** is being randomly rewritten, and you want to use IEEE 802.1 code point **002** for the **be2** forwarding class:

```
[edit class-of-service rewrite-rules]
user@switch# set ieee-802.1 custom-rw forwarding-class be2 loss-priority low code-point
002
```

2. Enable the rewrite rule on an interface if it is not already enabled on the desired interface:

```
[edit]
```

```
user@switch# set class-of-service interfaces interface-name unit unit rewrite-rules (dscp |
ieee-802.1) rewrite-rule-name
```

For example, to enable the rewrite rule **custom-rw** on interface **xe-0/0/24.0**:

```
[edit]
user@switch# set class-of-service interfaces xe-0/0/24 unit 0 rewrite-rules ieee-802.1
custom-rw
```

- Related Documentation**
- [interfaces on page 7082](#)
  - [rewrite-rules on page 7088](#)
  - [Defining CoS Rewrite Rules on page 6704](#)
  - [Monitoring CoS Rewrite Rules on page 7235](#)



## PART 92

# Scheduling Traffic

- [Using Schedulers on page 6711](#)



# Using Schedulers

- [Understanding Default CoS Scheduling and Classification on page 6712](#)
- [Understanding CoS Scheduling Behavior and Configuration Considerations on page 6718](#)
- [Understanding CoS Virtual Output Queues \(VOQs\) on QFX10000 Switches on page 6724](#)
- [Understanding CoS Port Schedulers on QFX Switches on page 6729](#)
- [Defining CoS Queue Schedulers for Port Scheduling on page 6742](#)
- [Example: Configuring Queue Schedulers for Port Scheduling on page 6744](#)
- [Troubleshooting Egress Bandwidth That Exceeds the Configured Minimum Bandwidth on page 6749](#)
- [Troubleshooting Egress Bandwidth That Exceeds the Configured Maximum Bandwidth on page 6751](#)
- [Defining CoS Queue Scheduling Priority on page 6752](#)
- [Example: Configuring Queue Scheduling Priority on page 6753](#)
- [Understanding CoS Traffic Control Profiles on page 6757](#)
- [Understanding CoS Priority Group Scheduling on page 6758](#)
- [Defining CoS Traffic Control Profiles \(Priority Group Scheduling\) on page 6761](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 6762](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 6765](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 6771](#)
- [Understanding CoS Priority Group and Queue Guaranteed Minimum Bandwidth on page 6795](#)
- [Example: Configuring Minimum Guaranteed Output Bandwidth on page 6798](#)
- [Understanding CoS Priority Group Shaping and Queue Shaping \(Maximum Bandwidth\) on page 6803](#)
- [Example: Configuring Maximum Output Bandwidth on page 6805](#)
- [Understanding CoS WRED Drop Profiles on page 6809](#)
- [Configuring CoS WRED Drop Profiles on page 6813](#)
- [Example: Configuring WRED Drop Profiles on page 6814](#)
- [Configuring CoS Drop Profile Maps on page 6817](#)
- [Example: Configuring Drop Profile Maps on page 6817](#)

- [Understanding CoS Explicit Congestion Notification on page 6820](#)
- [Example: Configuring ECN on page 6828](#)

## Understanding Default CoS Scheduling and Classification

---

If you do not explicitly configure classifiers and apply them to interfaces, the switch uses the default classifier to group ingress traffic into forwarding classes. If you do not configure scheduling on an interface, the switch uses the default schedulers to provide egress port resources for traffic. Default classification maps all traffic into the default forwarding classes (best-effort, fcoe, no-loss, and network-control). Each default forwarding class has a default scheduler, so traffic mapped to each default forwarding class receives port bandwidth, prioritization, and packet drop characteristics.

The switch supports direct port scheduling and enhanced transmission selection (ETS), also known as hierarchical port scheduling.

Hierarchical scheduling groups IEEE 802.1p priorities (IEEE 802.1p code points, which classifiers map to forwarding classes, which in turn are mapped to output queues) into priority groups (forwarding class sets). If you use only the default traffic scheduling and classification, the switch automatically creates a default priority group that contains all of the priorities (which are mapped to forwarding classes and output queues), and assigns 100 percent of the port output bandwidth to that priority group. The forwarding classes (queues) in the default forwarding class set receive bandwidth based on the default classifier settings. The default priority group is transparent. It does not appear in the configuration and is used for Data Center Bridging Capability Exchange (DCBX) protocol advertisement.



.....

**NOTE:** If you explicitly configure ETS by configuring one or more priority groups on an interface, any forwarding class that is not assigned to a priority group on that interface receives *no bandwidth*. This means that if you configure hierarchical scheduling on an interface, every forwarding class (priority) that you want to forward traffic on that interface must belong to a forwarding class set (priority group). ETS is not supported on QFX5200.

.....

This topic describes:

- [Default Classification on page 6712](#)
- [Default Scheduling on page 6715](#)
- [Default DCBX Advertisement on page 6717](#)
- [Default Scheduling and Classification Summary on page 6717](#)

## Default Classification

The default classifiers assign ingress traffic to default forwarding classes and loss priorities. The switch applies default IEEE 802.1, DSCP, and DSCP IPv6 classifiers to each interface that does not have explicitly configured classifiers. If you do not configure and apply EXP classifiers for MPLS traffic to logical interfaces, MPLS traffic on interfaces configured as **family mpls** uses the IEEE classifier.



If you explicitly configure one type of classifier but not other types of classifiers, the system uses only the configured classifier and does not use default classifiers for other types of traffic. There are two default IEEE 802.1 classifiers, a trusted classifier for ports that are in trunk mode, and an untrusted classifier for ports that are in access mode.

Table 548 shows the default mapping of IEEE 802.1 code-point values to forwarding classes and loss priorities for ports in trunk mode.

**Table 562: Default IEEE 802.1 Classifiers for Ports in Trunk Mode (Trusted Classifier)**

| Code Point | Forwarding Class | Loss Priority |
|------------|------------------|---------------|
| be (000)   | best-effort      | low           |
| be1 (001)  | best-effort      | low           |
| ef (010)   | best-effort      | low           |
| ef1 (011)  | fcoe             | low           |
| af11 (100) | no-loss          | low           |
| af12 (101) | best-effort      | low           |
| nc1 (110)  | network-control  | low           |
| nc2 (111)  | network-control  | low           |

Table 549 shows the default mapping of IEEE 802.1p code-point values to forwarding classes and loss priorities for ports in access mode (all incoming traffic is mapped to best-effort forwarding classes).

**Table 563: Default IEEE 802.1 Classifiers for Ports in Access Mode (Untrusted Classifier)**

| Code Point | Forwarding Class | Loss Priority |
|------------|------------------|---------------|
| 000        | best-effort      | low           |
| 001        | best-effort      | low           |
| 010        | best-effort      | low           |
| 011        | best-effort      | low           |
| 100        | best-effort      | low           |
| 101        | best-effort      | low           |
| 110        | best-effort      | low           |

**Table 563: Default IEEE 802.1 Classifiers for Ports in Access Mode (Untrusted Classifier) (*continued*)**

| Code Point | Forwarding Class | Loss Priority |
|------------|------------------|---------------|
| 111        | best-effort      | low           |

Table 550 shows the default mapping of DSCP code-point values to forwarding classes and loss priorities for DSCP IP and DCSP IPv6.

**Table 564: Default DSCP IP and IPv6 Classifiers**

| Code Point    | Forwarding Class | Loss Priority |
|---------------|------------------|---------------|
| ef (101110)   | best-effort      | low           |
| af11 (001010) | best-effort      | low           |
| af12 (001100) | best-effort      | low           |
| af13 (001110) | best-effort      | low           |
| af21 (010010) | best-effort      | low           |
| af22 (010100) | best-effort      | low           |
| af23 (010110) | best-effort      | low           |
| af31 (011010) | best-effort      | low           |
| af32 (011100) | best-effort      | low           |
| af33 (011110) | best-effort      | low           |
| af41 (100010) | best-effort      | low           |
| af42 (100100) | best-effort      | low           |
| af43 (100110) | best-effort      | low           |
| be (000000)   | best-effort      | low           |
| cs1 (001000)  | best-effort      | low           |
| cs2 (010000)  | best-effort      | low           |
| cs3 (011000)  | best-effort      | low           |
| cs4 (100000)  | best-effort      | low           |
| cs5 (101000)  | best-effort      | low           |

**Table 564: Default DSCP IP and IPv6 Classifiers (*continued*)**

| Code Point   | Forwarding Class | Loss Priority |
|--------------|------------------|---------------|
| nc1 (110000) | network-control  | low           |
| nc2 (111000) | network-control  | low           |

[Table 551](#) shows the default mapping of MPLS EXP code-point values to forwarding classes and loss priorities.

**Table 565: Default EXP Classifiers**

| Code Point | Forwarding Class     | Loss Priority |
|------------|----------------------|---------------|
| 000        | best-effort          | low           |
| 001        | best-effort          | high          |
| 010        | expedited-forwarding | low           |
| 011        | expedited-forwarding | high          |
| 100        | assured-forwarding   | low           |
| 101        | assured-forwarding   | high          |
| 110        | network-control      | low           |
| 111        | network-control      | high          |

## Default Scheduling

The default schedulers allocate egress bandwidth resources to egress traffic as shown in [Table 552](#):

**Table 566: Default Scheduler Configuration**

| Default Scheduler and Queue Number               | Transmit Rate (Guaranteed Minimum Bandwidth) | Rate Shaping (Maximum Bandwidth) | Excess Bandwidth Sharing | Priority | Buffer Size |
|--------------------------------------------------|----------------------------------------------|----------------------------------|--------------------------|----------|-------------|
| best-effort forwarding class scheduler (queue 0) | 15%                                          | None                             | 15%                      | low      | 15%         |
| fcoe forwarding class scheduler (queue 3)        | 35%                                          | None                             | 35%                      | low      | 35%         |
| no-loss forwarding class scheduler (queue 4)     | 35%                                          | None                             | 35%                      | low      | 35%         |

Table 566: Default Scheduler Configuration (*continued*)

| Default Scheduler and Queue Number                   | Transmit Rate (Guaranteed Minimum Bandwidth) | Rate Shaping (Maximum Bandwidth) | Excess Bandwidth Sharing | Priority | Buffer Size |
|------------------------------------------------------|----------------------------------------------|----------------------------------|--------------------------|----------|-------------|
| network-control forwarding class scheduler (queue 7) | 15%                                          | None                             | 15%                      | low      | 15%         |



**NOTE:** By default, the minimum guaranteed bandwidth (transmit rate) determines the amount of excess (extra) bandwidth a queue can share. Extra bandwidth is allocated to queues in proportion to the transmit rate of each queue. On switches that support the `excess-rate` statement, you can override the default setting and configure the excess bandwidth percentage independently of the transmit rate on queues that are not strict-high priority queues.

By default, only the four default schedulers shown in [Table 552](#) have traffic mapped to them. Only the forwarding classes and queues associated with the default schedulers receive default bandwidth, based on the default scheduler transmit rate. (You can configure schedulers and forwarding classes to allocate bandwidth to other queues or to change the bandwidth and other scheduling properties of a default queue.) If a forwarding class does not transport traffic, the bandwidth allocated to that forwarding class is available to other forwarding classes. Unicast and multdestination (multicast, broadcast, and destination lookup fail) traffic use the same forwarding classes and output queues.

Default scheduling is port scheduling.

Default hierarchical scheduling, known as enhanced transmission selection (ETS, defined in IEEE 802.1Qaz), allocates the total port bandwidth to the four default forwarding classes served by the four default schedulers, as defined by the four default schedulers. The result is the same as direct port scheduling. Configuring hierarchical port scheduling, however, enables you to group forwarding classes that carry similar types of traffic into forwarding class sets (also called priority groups), and to assign port bandwidth to each forwarding class set. The port bandwidth assigned to the forwarding class set is then assigned to the forwarding classes within the forwarding class set. This hierarchy enables you to control port bandwidth allocation with greater granularity, and enables hierarchical sharing of extra bandwidth to better utilize link bandwidth.

Default scheduling uses weighted round-robin (WRR) scheduling. Each queue receives a portion (weight) of the total available interface bandwidth. The scheduling weight is based on the transmit rate (minimum guaranteed bandwidth) of the default scheduler for that queue. For example, queue 7 receives a default scheduling weight of 15 percent of the available bandwidth, and queue 4 receives a default scheduling weight of 35 percent of the available bandwidth. Queues are mapped to forwarding classes (for example, queue 7 is mapped to the network-control forwarding class and queue 4 is mapped to the no-loss forwarding class). Each forwarding class receives the default

bandwidth for the queue to which it is mapped. Unused bandwidth is shared with other default queues.

If you want non-default (unconfigured) queues to forward traffic, you should explicitly map traffic to those queues (configure the forwarding classes and queue mapping) and create schedulers to allocate bandwidth to those queues. By default, queues 1, 2, 5, and 6 are unconfigured. Unconfigured queues have a default scheduling weight of 1 so that they can receive a small amount of bandwidth in case they need to forward traffic.

If you map traffic to an unconfigured queue and do not schedule port resources for the queue (configure a scheduler, map it to the forwarding class that is mapped to the queue, and apply the scheduler mapping to the port), the queue receives only the amount of excess bandwidth proportional to its default weight (1). The actual amount of bandwidth an unconfigured queue gets depends on how much bandwidth the other queues on the port are using.

If the other queues use less than their allocated amount of bandwidth, the unconfigured queues can share the unused bandwidth. Configured queues have higher priority for bandwidth than unconfigured queues, so if a configured queue needs more bandwidth, then less bandwidth is available for unconfigured queues. Unconfigured queues always receive a minimum amount of bandwidth based on their scheduling weight (1). If you map traffic to an unconfigured queue, to allocate bandwidth to that queue, configure a scheduler for the forwarding class that is mapped to the queue, and apply it to the port.

## Default DCBX Advertisement

When you configure hierarchical scheduling on an interface, DCBX advertises each priority group, the priorities in each priority group, and the bandwidth properties of each priority and priority group.

If you do not configure hierarchical scheduling on an interface, DCBX advertises the automatically created default priority group and its priorities. DCBX also advertises the default bandwidth allocation of the priority group, which is 100 percent of the port bandwidth.

## Default Scheduling and Classification Summary

If you do not configure scheduling on an interface:

- Default classifiers classify ingress traffic.
- Default schedulers schedule egress traffic.
- DCBX advertises a single default priority group with 100 percent of the port bandwidth allocated to that priority group. All priorities (forwarding classes) are assigned to the default priority group and receive bandwidth based on their default schedulers. The default priority group is generated automatically and is not user-configurable.

### Related Documentation

- [Understanding CoS Packet Flow on page 6628](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 6765](#)
- [Understanding Default CoS Settings on page 6630](#)

- [CoS Support on QFX Series Switches, EX4600 Switches, and QFabric Systems on page 6610](#)
- [Understanding CoS Virtual Output Queues \(VOQs\) on QFX10000 Switches on page 6724](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 6672](#)
- [Understanding DCB Features and Requirements on page 6490](#)
- [Example: Configuring Classifiers on page 6657](#)
- [Example: Configuring Queue Schedulers for Port Scheduling on page 6744](#)

## Understanding CoS Scheduling Behavior and Configuration Considerations

Many factors affect scheduling configuration and bandwidth requirements, including:

- When you configure bandwidth for a forwarding class (each forwarding class is mapped to a queue) or a forwarding class set (priority group), the switch considers only the data as the configured bandwidth. The switch does not account for the bandwidth consumed by the preamble and the interframe gap (IFG). Therefore, when you calculate and configure the bandwidth requirements for a forwarding class or for a forwarding class set, consider the preamble and the IFG as well as the data in the calculations.
- When you configure a forwarding class to carry traffic on the switch (instead of using only default forwarding classes), you must also define a scheduling policy for the user-configured forwarding class. Some switches support enhanced transmission selection (ETS) hierarchical port scheduling, some switches support direct port scheduling, and some switches support both methods of scheduling.

For ETS hierarchical port scheduling, defining a hierarchical scheduling policy using ETS means:

- Mapping a scheduler to the forwarding class in a scheduler map
- Including the forwarding class in a forwarding class set
- Associating the scheduler map with a traffic control profile
- Attaching the traffic control profile to a forwarding class set and an interface

On switches that support port scheduling, defining a scheduling policy means:

- Mapping a scheduler to the forwarding class in a scheduler map.
- Applying the scheduler map to one or more interfaces.
- On each physical interface, either all forwarding classes that are being used on the interface must have rewrite rules configured, or no forwarding classes that are being used on the interface can have rewrite rules configured. On any physical port, do not mix forwarding classes with rewrite rules and forwarding classes without rewrite rules.
- For packets that carry both an inner VLAN tag and an outer VLAN tag, rewrite rules rewrite only the outer VLAN tag.

- For ETS hierarchical port scheduling, configuring the minimum guaranteed bandwidth (**transmit-rate**) for a forwarding class does not work unless you also configure the minimum guaranteed bandwidth (**guaranteed-rate**) for the forwarding class set in the traffic control profile.

Additionally, the sum of the transmit rates of the forwarding classes in a forwarding class set should not exceed the guaranteed rate for the forwarding class set. (You cannot guarantee a minimum bandwidth for the queues that is greater than the minimum bandwidth guaranteed for the entire set of queues.) If you configure transmit rates whose sum exceeds the guaranteed rate of the forwarding class set, the commit check fails and the system rejects the configuration.

- For ETS hierarchical port scheduling, the sum of the forwarding class set guaranteed rates cannot exceed the total port bandwidth. If you configure guaranteed rates whose sum exceeds the port bandwidth, the system sends a syslog message to notify you that the configuration is not valid. However, the system does not perform a commit check. If you commit a configuration in which the sum of the guaranteed rates exceeds the port bandwidth, the hierarchical scheduler behaves unpredictably.
- For ETS hierarchical port scheduling, if you configure the **guaranteed-rate** of a forwarding class set as a percentage, configure all of the transmit rates associated with that forwarding class set as percentages. In this case, if any of the transmit rates are configured as absolute values instead of percentages, the configuration is not valid and the system sends a syslog message.
- There are several factors to consider if you want to configure a strict-high priority queue (forwarding class):
  - You can configure only one strict-high priority queue (forwarding class) on QFX5200, QFX5100, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems.

On QFX10000 switches, there is no limit to the number of strict-high priority queues you can configure.
  - You cannot configure a minimum guaranteed bandwidth (**transmit-rate**) for a strict-high priority queue on QFX5200, QFX5100, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems.

On QFX5200 and QFX10000 switches, you can set the **transmit-rate** on strict-high priority queues to set a limit on the amount of traffic that the queue treats as strict-high priority traffic. Traffic in excess of the **transmit-rate** is treated as best-effort traffic, and receives an excess bandwidth sharing weight of "1", which is the proportion of extra bandwidth the strict-high priority queue can share on the port. Queues that are not strict-high priority queues use the transmit rate (default) or the configured excess rate to determine the proportion (weight) of extra port bandwidth the queue can share. However, you cannot configure an excess rate on a strict-high priority queue, and you cannot change the excess bandwidth sharing weight of "1" on a strict-high priority queue.

For ETS hierarchical port scheduling, you cannot configure a minimum guaranteed bandwidth (**guaranteed-rate**) for a forwarding class set that includes a strict-high priority queue.

- Except on QFX10000 switches, for ETS hierarchical port scheduling only, you must create a separate forwarding class set for a strict-high priority queue. On QFX10000 switches, you can mix strict-high priority and low priority queues in the same forwarding class set.
- Except on QFX10000 switches, for ETS hierarchical port scheduling, only one forwarding class set can contain a strict-high priority queue. On QFX10000 switches, this restriction does not apply.
- Except on QFX10000 switches, for ETS hierarchical port scheduling, a strict-high priority queue cannot belong to the same forwarding class set as queues that are not strict-high priority. (You cannot mix a strict-high priority forwarding class with forwarding classes that are not strict-high priority in one forwarding class set.) On QFX10000 switches, you can mix strict-high priority and low priority queues in the same forwarding class set.
- For ETS hierarchical port scheduling on switches that use different forwarding class sets for unicast and multdestination (multicast, broadcast, and destination lookup fail) traffic, a strict-high priority queue cannot belong to a multdestination forwarding class set.
- On QFX10000 systems, we recommend that you always configure a transmit rate on strict-high priority queues to prevent them from starving other queues. If you do not apply a transmit rate to limit the amount of bandwidth strict-high priority queues can use, then strict-high priority queues can use all of the available port bandwidth and starve other queues on the port.

On QFX5200, QFX5100, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems, we recommend that you always apply a shaping rate to the strict-high priority queue to prevent it from starving other queues. If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.

- On QFabric systems, if any queue that contains outgoing packets does not transmit packets for 12 consecutive seconds, the port automatically resets. Failure of a queue to transmit packets for 12 consecutive seconds might be due to:
  - A strict-high priority queue consuming all of the port bandwidth
  - Several queues consuming all of the port bandwidth
  - Any queue or port receiving continuous priority-based flow control (PFC) or 802.3x Ethernet PAUSE messages (received PFC and PAUSE messages prevent a queue or a port, respectively, from transmitting packets because of network congestion)
  - Other conditions that prevent a queue from obtaining port bandwidth for 12 consecutive seconds

If the cause is a strict-high priority queue consuming all of the port bandwidth, use rate shaping to configure a maximum rate for the strict-high priority queue and prevent it from using all of the port bandwidth. To configure rate shaping, include the **shaping-rate (rate | percent percentage)** statement at the **[edit class-of-service schedulers scheduler-name]** hierarchy level and apply the shaping rate to the strict-high priority



scheduler. We recommend that you always apply a shaping rate to strict-high priority traffic to prevent the strict-high priority queue from starving other queues.

If several queues consume all of the port bandwidth, you can use a scheduler to rate shape those queues and prevent them from using all of the port bandwidth.

- For transmit rates below 1 Gbps, we recommend that you configure the transmit rate as a percentage instead of as a fixed rate. This is because the system converts fixed rates into percentages and might round small fixed rates to a lower percentage. For example, a fixed rate of 350 Mbps is rounded down to 3 percent instead of 3.5 percent.
- When you set the maximum bandwidth for a queue or for a priority group (**shaping-rate**) at 100 Kbps or lower, the traffic shaping behavior is accurate only within +/– 20 percent of the configured **shaping-rate**.
- On QFX10000 switches, configuring rate shaping (**[set class-of-service schedulers scheduler-name transmit-rate (rate | percentage) exact]**) on a LAG interface using the **[edit class-of-service interfaces lag-interface-name scheduler-map scheduler-map-name]** statement can result in scheduled traffic streams receiving more LAG link bandwidth than expected.

You configure rate shaping in a scheduler to set the maximum bandwidth for traffic assigned to a forwarding class on a particular output queue on a port. For example, you can use a scheduler to configure rate shaping on traffic assigned to the best-effort forwarding class mapped to queue 0, and then apply the scheduler to an interface using a scheduler map, to set the maximum bandwidth for best-effort traffic mapped to queue 0 on that port. Traffic in the best-effort forwarding can use no more than the amount of port bandwidth specified by the transmit rate when you use the **exact** option.

LAG interfaces are composed of two or more Ethernet links bundled together to function as a single interface. The switch can hash traffic entering a LAG interface onto any member link in the LAG interface. When you configure rate shaping and apply it to a LAG interface, the way that the switch applies the rate shaping to traffic depends on how the switch hashes the traffic onto the LAG links.

To illustrate how link hashing affects the way the switch applies a shaping rate to LAG traffic, let's look at a LAG interface (**ae0**) that has two member links (**xe-0/0/20** and **xe-0/0/21**). On LAG **ae0**, we configure rate shaping of **2g** for traffic assigned to the **best-effort** forwarding class, which is mapped to output queue 0. When traffic in the **best-effort** forwarding class reaches the LAG interface, the switch hashes the traffic onto one of the two member links.

If the switch hashes all of the **best-effort** traffic onto the same LAG link, the traffic receives a maximum of 2g bandwidth on that link. In this case, the intended cumulative limit of 2g for best-effort traffic on the LAG is enforced.

However, if the switch hashes the **best-effort** traffic onto both of the LAG links, the traffic receives a maximum of 2g bandwidth on *each* LAG link, not 2g as a cumulative total for the entire LAG, so the best-effort traffic receives a maximum of 4g on the LAG, not the 2g set by the rate shaping configuration. When hashing spreads the traffic assigned to an output queue (which is mapped to a forwarding class) across multiple LAG links, the effective rate shaping (cumulative maximum bandwidth) on the LAG is:

(number of LAG member interfaces) x (rate shaping for the output queue) = cumulative LAG rate shaping

- On switches that do not use virtual output queues (VOQs), ingress port congestion can occur during periods of egress port congestion if an ingress port forwards traffic to more than one egress port, and at least one of those egress ports experiences congestion. If this occurs, the congested egress port can cause the ingress port to exceed its fair allocation of ingress buffer resources. When the ingress port exceeds its buffer resource allocation, frames are dropped at the ingress. Ingress port frame drop affects not only the congested egress ports, but also all of the egress ports to which the congested ingress port forwards traffic.

If a congested ingress port drops traffic that is destined for one or more uncongested egress ports, configure a weighted random early detection (WRED) drop profile and apply it to the egress queue that is causing the congestion. The drop profile prevents the congested egress queue from affecting egress queues on other ports by dropping frames at the egress instead of causing congestion at the ingress port.



**NOTE:** On systems that support lossless transport, do not configure drop profiles for lossless forwarding classes such as the default **fcoe** and **no-loss** forwarding classes. FCoE and other lossless traffic queues require lossless behavior. Use priority-based flow control (PFC) to prevent frame drop on lossless priorities.

- On systems that use different classifiers for unicast and multdestination traffic and that support lossless transport, on an ingress port, do not configure classifiers that map the same IEEE 802.1p code point to both a multdestination traffic flow and a lossless unicast traffic flow (such as the default lossless **fcoe** or **no-loss** forwarding classes). Any code point used for multdestination traffic on a port should not be used to classify unicast traffic into a lossless forwarding class on the same port.

If a multdestination traffic flow and a lossless unicast traffic flow use the same code point on a port, the multdestination traffic is treated the same way as the lossless traffic. For example, if priority-based flow control (PFC) is applied to the lossless traffic, the multdestination traffic of the same code point is also paused. During periods of congestion, treating multdestination traffic the same as lossless unicast traffic can create ingress port congestion for the multdestination traffic and affect the multdestination traffic on all of the egress ports the multdestination traffic uses.

For example, the following configuration can cause ingress port congestion for the multdestination flow:

1. For unicast traffic, IEEE 802.1p code point 011 is classified into the **fcoe** forwarding class:  

```
user@switch# set class-of-service classifiers ieee-802.1 ucast_cl forwarding-class fcoe loss-priority low code-points 011
```
2. For multdestination traffic, IEEE 802.1p code point 011 is classified into the **mcast** forwarding class:

```
user@switch# set class-of-service classifiers ieee-802.1 mcast-cl forwarding-class mcast
loss-priority low code-points 011
```

3. The unicast classifier that maps traffic with code point **011** to the **fcoe** forwarding class is mapped to interface **xe-0/0/1**:

```
user@switch# set class-of-service interfaces xe-0/0/1 unit 0 classifiers ieee-802.1 ucast-cl
```

4. The multdestination classifier that maps traffic with code point **011** to the **mcast** forwarding class is mapped to all interfaces (multidestination traffic maps to all interfaces and cannot be mapped to individual interfaces):

```
user@switch# set class-of-service multi-destination classifiers ieee-802.1 mcast-cl
```

Because the same code point (**011**) maps unicast traffic to a lossless traffic flow and also maps multidestination traffic to a multidestination traffic flow, the multidestination traffic flow might experience ingress port congestion during periods of congestion.

To avoid ingress port congestion, do not map the code point used by the multidestination traffic to lossless unicast traffic. For example:

1. Instead of classifying code point **011** into the **fcoe** forwarding class, classify code point **011** into the **best-effort** forwarding class:

```
user@switch# set class-of-service classifiers ieee-802.1 ucast-cl forwarding-class
best-effort loss-priority low code-points 011
```

2. user@switch# set class-of-service classifiers ieee-802.1 mcast-cl forwarding-class mcast loss-priority low code-points 011
  3. user@switch# set class-of-service interfaces xe-0/0/1 unit 0 classifiers ieee-802.1 ucast-cl
  4. user@switch# set class-of-service multi-destination classifiers ieee-802.1 mcast-cl
- Because the code point **011** does not map unicast traffic to a lossless traffic flow, the multidestination traffic flow does not experience ingress port congestion during periods of congestion.

The best practice is to classify unicast traffic with IEEE 802.1p code points that are also used for multidestination traffic into best-effort forwarding classes.

## Understanding CoS Virtual Output Queues (VOQs) on QFX10000 Switches

---

The traditional method of forwarding traffic through a switch is based on buffering ingress traffic in input queues on ingress interfaces, forwarding the traffic across the switch fabric to output queues on egress interfaces, and then buffering traffic again on the output queues before transmitting the traffic to the next hop. The traditional method of queueing packets on an ingress port is storing traffic destined for different egress ports in the same input queue (buffer).

During periods of congestion, the switch might drop packets at the egress port, so the switch might spend resources transporting traffic across the switch fabric to an egress port, only to drop that traffic instead of forwarding it. And because input queues store traffic destined for different egress ports, congestion on one egress port could affect traffic on a different egress port, a condition called *head-of-line blocking (HOLB)*.

*Virtual output queue (VOQ)* architecture takes a different approach:

- Instead of separate physical buffers for input and output queues, the switch uses the physical buffers on the ingress pipeline of each Packet Forwarding Engine (PFE) chip to store traffic for every egress port. Every output queue on an egress port has buffer storage space on every ingress pipeline on all of the PFE chips on the switch. The mapping of ingress pipeline storage space to output queues is 1-to-1, so each output queue receives buffer space on each ingress pipeline.
- Instead of one input queue containing traffic destined for multiple different output queues (a one-to-many mapping), each output queue has a dedicated VOQ comprised of the input buffers on each packet forwarding chip that are dedicated to that output queue (a 1-to-1 mapping). This architecture prevents communication between any two ports from affecting another port.
- Instead of storing traffic on a physical output queue until it can be forwarded, a VOQ does not transmit traffic from the ingress port across the fabric to the egress port until the egress port has the resources to forward the traffic.

A VOQ is a collection of input queues (buffers) that receive and store traffic destined for one output queue on one egress port. Each output queue on each egress port has its own dedicated VOQ, which consists of all of the input queues that are sending traffic to that output queue.

- [VOQ Architecture on page 6724](#)
- [VOQ Advantages on page 6726](#)

### VOQ Architecture

A VOQ represents the ingress buffering for a particular output queue. A unique buffer ID identifies each output queue on a PFE chip. Each of the six PFE chips uses the same unique buffer ID for a particular output queue. The traffic stored using a particular buffer ID on the six PFE chips comprises the traffic destined for one particular output queue on one port, and is the VOQ for that output queue.

A switch that has 72 egress ports with 8 output queues on each port, has 576 VOQs on each PFE chip ( $72 \times 8 = 576$ ). Because the switch has six PFE chips, the switch has a total of 3,456 VOQs ( $576 \times 6 = 3,456$ ).

A VOQ is distributed across all of the PFE chips that are actively sending traffic to that output queue. Each output queue is the sum of the total buffers assigned to that output queue (by its unique buffer ID) across all of the PFE chips. So the output queue itself is virtual, not physical, although the output queue is comprised of physical input queues.

- [Round-Trip Time Buffering on page 6725](#)
- [Requesting and Granting Egress Port Bandwidth on page 6725](#)

### Round-Trip Time Buffering

---

Although there is no output queue buffering during periods of congestion (no long-term storage), there is a small physical output queue buffer on egress line cards to accommodate the round-trip time for traffic to traverse the switch fabric from ingress to egress. The round-trip time consists of the time it takes the ingress port to request egress port resources, receive a grant from the egress port for resources, and transmit the data across the switch fabric.

That means if a packet is not dropped at the switch ingress, and the switch forwards the packet across the fabric to the egress port, the packet will not be dropped and will be forwarded to the next hop. All packet drops take place in the ingress pipeline.

The switch has 4 GB of external DRAM to use as a delay bandwidth buffer (DBB). The DBB provides storage for ingress ports until the ports can forward traffic to egress ports.

### Requesting and Granting Egress Port Bandwidth

---

When packets arrive at an ingress port, the ingress pipeline stores the packet in the ingress queue with the unique buffer ID of the destination output queue. The switch makes the buffering decision after performing the packet lookup. If the packet belongs to a class for which the maximum traffic threshold has been exceeded, the packet might not be buffered and might be dropped. To transport packets across the switch fabric to egress ports:

1. The ingress line card PFE request scheduler sends a request to the egress line card PFE grant scheduler to notify the egress PFE that data is available for transmission.
2. When there is available egress bandwidth, the egress line card grant scheduler responds by sending a bandwidth grant to the ingress line card PFE.
3. The ingress line card PFE receives the grant from the egress line card PFE, and transmits the data to the egress line card.

Ingress packets remain in the VOQ on the ingress port input queues until the output queue is ready to accept and forward more traffic.

Under most conditions, the switch fabric is fast enough to be transparent to egress class-of-service (CoS) policies, so the process of forwarding traffic from the ingress pipeline, across the switch fabric, to egress ports, does not affect the configured CoS

policies for the traffic. The fabric only affects CoS policy if there is a fabric failure or if there is an issue of port fairness.

When a packet ingresses and egresses the same PFE chip (local switching), the packet does not traverse the switch fabric. However, the switch uses the same request and grant mechanism to receive egress bandwidth as packets that cross the fabric, so locally switched packets and packets that arrive at a PFE chip after crossing the switch fabric are treated fairly when the traffic is contending for the same output queue.

## VOQ Advantages

VOQ architecture provides two major advantages:

- [Eliminate Head-of-Line Blocking on page 6726](#)
- [Increase Fabric Efficiency and Utilization on page 6728](#)


### Eliminate Head-of-Line Blocking

VOQ architecture eliminates head-of-line blocking (HOLB) issues. On non-VOQ switches, HOLB occurs when congestion at an egress port affects a different egress port that is not congested. HOLB occurs when the congested port and the uncongested port share the same input queue on an ingress interface.

An example of a HOLB scenario is a switch that has streams of traffic entering one ingress port (IP-1) that are destined for two different egress ports (EP-2 and EP-3):


1. Congestion occurs on egress port EP-2. There is no congestion on egress port EP-3, as shown in [Figure 212](#).

**Figure 212: Congestion Occurs on EP-2**



2. Egress port EP-2 sends a backpressure signal to ingress port IP-1, as shown in [Figure 213](#).


**Figure 213: EP-2 Backpressures IP-1**



3. The backpressure signal causes the ingress port IP-1 to stop sending traffic and to buffer traffic until it receives a signal to resume sending, as shown in [Figure 214](#). Traffic that arrives at ingress port IP-1 destined for uncongested egress port EP-3 is buffered


along with the traffic destined for congested port EP-2, instead of being forwarded to port EP-3.

**Figure 214: Backpressure from EP-2 Causes IP-1 to Buffer Traffic Instead of Sending Traffic, Affecting EP-3**



4. Ingress port IP-1 transmits traffic to uncongested egress port EP-3 only when egress port EP-2 clears enough to allow ingress port IP-1 to resume sending traffic, as shown in Figure 215.


**Figure 215: Congestion on EP-2 Clears, Allowing IP-1 to Resume Sending Traffic to Both Egress Ports**



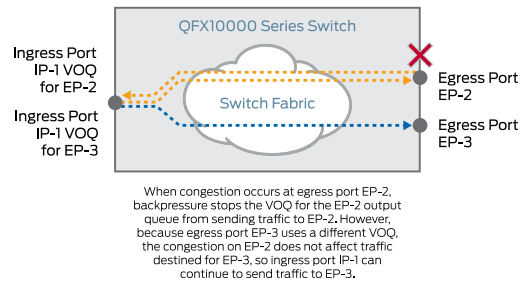
In this way, congested egress port EP-2 negatively affects uncongested egress port EP-3, because both egress ports share the same input queue on ingress port IP-1.

VOQ architecture avoids HOLB by creating a different dedicated virtual queue for each output queue on each interface, as shown in Figure 216.

**Figure 216: Each Egress Port Has a Separate Virtual Output Queue on IP-1**



Because different egress queues do not share the same input queue, a congested egress queue on one port cannot affect an egress queue on a different port, as shown in Figure 217. (For the same reason, a congested egress queue on one port cannot affect another egress queue on the same port—each output queue has its own dedicated virtual output queue composed of ingress interface input queues.)

**Figure 217: Congestion on EP-2 Does Not Affect Uncongested Port EP-3**

Performing queue buffering at the ingress interface ensures that the switch only sends traffic across the fabric to an egress queue if that egress queue is ready to receive that traffic. If the egress queue is not ready to receive traffic, the traffic remains buffered at the ingress interface.

### Increase Fabric Efficiency and Utilization

Traditional output queue architecture has some inherent inefficiencies that VOQ architecture addresses.

- **Packet buffering**—Traditional queueing architecture buffers each packet twice in long-term DRAM storage, once at the ingress interface and once at the egress interface. VOQ architecture buffers each packet only once in long-term DRAM storage, at the ingress interface. The switch fabric is fast enough to be transparent to egress CoS policies, so instead of buffering packets a second time at the egress interface, the switch can forward traffic at a rate that does not require deep egress buffers, without affecting the configured egress CoS policies (scheduling).
- **Consumption of resources**—Traditional queueing architecture sends packets from the ingress interface input queue (buffer), across the switch fabric, to the egress interface output queue (buffer). At the egress interface, packets might be dropped, even though the switch has expended resources transporting the packets across the fabric and storing them in the egress queue. VOQ architecture does not send packets across the fabric to the egress interface until the egress interface is ready to transmit the traffic. This increases system utilization because no resources are wasted transporting and storing packets that are dropped later.

Independent of VOQ architecture, the Juniper Networks switching architecture also provides better fabric utilization because the switch converts packets into cells. Cells have a predictable size, which enables the switch to spray the cells evenly across the fabric links and more fully utilize the fabric links. Packets vary greatly in size, and packet size is not predictable. Packet-based fabrics can deliver no better than 65-70 percent utilization because of the variation and unpredictability of packet sizes. Juniper Networks' cell-based fabrics can deliver a fabric utilization rate of almost 95 percent because of the predictability of and control over cell size.

#### Related Documentation

- [Understanding CoS Port Schedulers on QFX Switches on page 6729](#)
- [Example: Configuring Queue Schedulers for Port Scheduling on page 6744](#)



- [Understanding Default CoS Scheduling and Classification on page 6666](#)

## Understanding CoS Port Schedulers on QFX Switches

---

Port scheduling defines the class-of-service (CoS) properties of output queues. You configure CoS properties in a scheduler, then map the scheduler to a forwarding class. Forwarding classes are in turn mapped to output queues. Classifiers map incoming traffic into forwarding classes based on IEEE 802.1p, DSCP, or EXP code points.

Output queue properties include the amount of interface bandwidth assigned to the queue, the size of the memory buffer allocated for storing packets, the scheduling priority of the queue, and the weighted random early detection (WRED) drop profiles associated with the queue to control packet drop during periods of congestion.

Scheduler maps map schedulers to forwarding classes. The output queue mapped to a forwarding class receives the port resources and properties defined in the scheduler mapped to that forwarding class. You apply a scheduler map to an interface to apply queue scheduling to a port. You can associate different scheduler maps with different interfaces to configure port-specific scheduling for forwarding classes (output queues).



**NOTE:** Port scheduling is simpler to configure than enhanced transmission selection (ETS) two-tier hierarchical port scheduling. Port scheduling allocates port bandwidth to output queues directly, instead of allocating port bandwidth to output queues through a scheduling hierarchy. While port scheduling is simpler, ETS is more flexible.

ETS allocates port bandwidth in a two-tier hierarchy:

- Port bandwidth is first allocated to a priority group using the CoS properties defined in a traffic control profile. A priority group is a group of forwarding classes (which are mapped to output queues) that require similar CoS treatment.
- Priority group bandwidth is allocated to the output queues (which are mapped to forwarding classes) using the properties defined in the output queue scheduler.



**NOTE:** When you configure bandwidth for a queue, the switch considers only the data as the configured bandwidth. The switch does not account for the bandwidth consumed by the preamble and the interframe gap (IFG). Therefore, when you calculate and configure the bandwidth requirements for a queue, consider the preamble and the IFG as well as the data in the calculations.

- [Queue Scheduling Components on page 6730](#)
- [Default Schedulers on page 6732](#)

- [Scheduling Priority on page 6733](#)
- [Bandwidth Scheduling on page 6735](#)
- [Scheduler Drop-Profile Maps on page 6739](#)
- [Buffer Size on page 6739](#)
- [Explicit Congestion Notification on page 6741](#)
- [Scheduler Maps on page 6741](#)

## Queue Scheduling Components

[Table 567](#) provides a quick reference to the scheduler components you can configure to determine the bandwidth properties of output queues (forwarding classes).

**Table 567: Output Queue Scheduler Components**

| Output Queue Scheduler Component | Description                                                                                                                                                                                                                                                                                                          |
|----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Buffer size                      | Sets the size of the queue buffer.                                                                                                                                                                                                                                                                                   |
| Drop profile map                 | Maps a drop profile to a packet loss priority. Drop profile map components include: <ul style="list-style-type: none"> <li>• Drop profile—Sets the probability of dropping packets as the queue fills up.</li> <li>• Loss priority—Sets the traffic packet loss priority to which a drop profile applies.</li> </ul> |
| Excess rate                      | Sets the percentage of extra bandwidth (bandwidth that is not used by other queues) a queue can receive. If not set, the switch uses the transmit rate to determine how much extra bandwidth the queue can use. Extra bandwidth is the bandwidth remaining after all guaranteed bandwidth requirements are met.      |
| Explicit congestion notification | Enables explicit congestion notification (ECN) on the queue.                                                                                                                                                                                                                                                         |
| Priority                         | Sets the scheduling priority applied to the queue.                                                                                                                                                                                                                                                                   |

Table 567: Output Queue Scheduler Components (*continued*)

| Output Queue Scheduler Component | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|----------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Transmit rate                    | <p>Sets the minimum guaranteed bandwidth on low priority queues. By default, if you do not configure an excess rate, extra bandwidth is shared among queues in proportion to the transmit rate of each queue.</p> <p>On strict-high priority queues, sets the amount of bandwidth that receives strict-high priority forwarding treatment. Traffic that exceeds the transmit rate shares in the port excess bandwidth pool based on the strict-high priority excess bandwidth sharing weight of "1", which is not configurable. The actual amount of extra bandwidth that traffic exceeding the transmit rate receives depends on how many other queues consume excess bandwidth and the excess rates of those queues.</p> <p>If you configure two or more strict-high priority queues on a port, you must configure a transmit rate on those queues. However, we strongly recommend that you always configure a transmit rate on strict-high priority queues to prevent them from starving other queues.</p> |

Table 568 provides a quick reference to some related scheduling configuration components.

Table 568: Related Scheduling Components

| Related Scheduling Components       | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Forwarding class                    | Maps traffic classified into the forwarding class at the switch ingress to an output queue. Classifiers map forwarding classes to IEEE 802.1p, DSCP, or EXP code points. A forwarding class, an output queue, and code point bits are mapped to each other and identify the same traffic. (The code point bits identify incoming traffic. Classifiers assign traffic to forwarding classes based on the code point bits. Forwarding classes map to output queues. This mapping determines the output queue each class of traffic uses on the switch egress interfaces.)                     |
| Output queue (virtual output queue) | Output queues are virtual, and are comprised of the physical buffers on the ingress pipeline of each Packet Forwarding Engine (PFE) chip to store traffic for every egress port. Every output queue on an egress port has buffer storage space on every ingress pipeline on all of the PFE chips on the switch. The mapping of ingress pipeline storage space to output queues is 1-to-1, so each output queue receives buffer space on each ingress pipeline. See <a href="#">"Understanding CoS Virtual Output Queues (VOQs) on QFX10000 Switches" on page 6724</a> for more information. |

**Table 568: Related Scheduling Components (*continued*)**

| Related Scheduling Components | Description                                                                                                                                                                                                                                                  |
|-------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Scheduler map                 | Maps schedulers to forwarding classes (forwarding classes are mapped to queues, so a forwarding class represents a queue, and the scheduler mapped to a forwarding class determines the CoS properties of the output queue mapped to that forwarding class). |

## Default Schedulers

If you do not configure CoS, the switch uses its default settings. Each forwarding class requires a scheduler to set the CoS properties of the forwarding class and its output queue. The default configuration has four forwarding classes: best-effort (queue 0), fcoe (queue 3), no-loss (queue 4), and network-control (queue 7). Each default forwarding class is mapped to a default scheduler. You can use the default schedulers or you can define new schedulers for these four forwarding classes. For explicitly configured forwarding classes, you must explicitly configure a queue scheduler to allocate CoS resources to the traffic mapped to each forwarding class.

Table 569 shows the default queue schedulers.

**Table 569: Default Scheduler Configuration**

| Default Scheduler and Queue Number                   | Transmit Rate (Guaranteed Minimum Bandwidth) | Rate Shaping (Maximum Bandwidth) | Excess Bandwidth Sharing | Priority | Buffer Size |
|------------------------------------------------------|----------------------------------------------|----------------------------------|--------------------------|----------|-------------|
| best-effort forwarding class scheduler (queue 0)     | 15%                                          | None                             | 15%                      | low      | 15%         |
| fcoe forwarding class scheduler (queue 3)            | 35%                                          | None                             | 35%                      | low      | 35%         |
| no-loss forwarding class scheduler (queue 4)         | 35%                                          | None                             | 35%                      | low      | 35%         |
| network-control forwarding class scheduler (queue 7) | 15%                                          | None                             | 15%                      | low      | 15%         |



**NOTE:** By default, the minimum guaranteed bandwidth (transmit rate) determines the amount of excess (extra) bandwidth a queue can share. Extra bandwidth is allocated to queues in proportion to the transmit rate of each queue. You can configure bandwidth sharing (excess rate) to override the default setting and configure the excess bandwidth percentage independently of the transmit rate.

By default, only the four default schedulers shown in Table 569 have traffic mapped to them. Only the forwarding classes and queues associated with the default schedulers

receive default bandwidth, based on the default scheduler transmit rate. (You can configure schedulers and forwarding classes to allocate bandwidth to other queues or to change the default bandwidth of a default queue.) If a forwarding class does not transport traffic, the bandwidth allocated to that forwarding class is available to other forwarding classes. Unicast and multdestination (multicast, broadcast, and destination lookup fail) traffic use the same forwarding classes and output queues.

Default scheduling is port scheduling. If you configure scheduling instead of using default scheduling, you can configure port scheduling or enhanced transmission selection (ETS) hierarchical port scheduling.

Default scheduling uses weighted round-robin (WRR) scheduling. Each queue receives a portion (weight) of the total available port bandwidth. The scheduling weight is based on the transmit rate (minimum guaranteed bandwidth) of the default scheduler for that queue. For example, queue 7 receives a default scheduling weight of 15 percent of available port bandwidth, and queue 4 receives a default scheduling weight of 35 percent of available bandwidth. Queues are mapped to forwarding classes (for example, queue 7 is mapped to the network-control forwarding class and queue 4 is mapped to the no-loss forwarding class), so forwarding classes receive the default bandwidth for the queues to which they are mapped. Unused bandwidth is shared with other default queues.

You should explicitly map traffic to non-default (unconfigured) queues and schedule bandwidth resources for those queues if you want to use them to forward traffic. By default, queues 1, 2, 5, and 6 are unconfigured. Unconfigured queues have a default scheduling weight of 1 so that they can receive a small amount of bandwidth in case they need to forward traffic.

If you map traffic to an unconfigured queue and do not schedule bandwidth for the queue, the queue receives only the amount of bandwidth proportional to its default weight (1). The actual amount of bandwidth an unconfigured queue receives depends on how much bandwidth the other queues on the port are using.

If the other queues use less than their allocated amount of bandwidth, the unconfigured queues can share the unused bandwidth. Because of their scheduling weights, configured queues have higher priority for bandwidth than unconfigured queues. If a configured queue needs more bandwidth, then less bandwidth is available for unconfigured queues. However, unconfigured queues always receive a minimum amount of bandwidth based on their scheduling weight (1). If you map traffic to an unconfigured queue, to allocate bandwidth to that queue, configure a scheduler and map it to the forwarding class that is mapped to the queue, and then apply the scheduler map to the port.

## Scheduling Priority

Scheduling priority determines the order in which an interface transmits traffic from its output queues. Priority settings ensure that queues containing important traffic receive prioritized access to the outgoing interface bandwidth. The priority setting in the scheduler determines queue priority (a scheduler map maps the scheduler to a forwarding class, the forwarding class is mapped to an output queue, and the output queue uses the CoS properties defined in the scheduler).

By default, all queues are low priority queues. The switch supports two levels of scheduling priority:

- Low—In the default CoS state, all queues are low priority queues. Low priority queues transmit traffic based on the weighted round-robin (WRR) algorithm. If you configure scheduling priorities higher than low priority on queues, then the higher priority queues are served before the low priority queues.
- Strict-high—You can configure queues as **strict-high** priority. Strict-high priority queues receive preferential treatment over all other queues, and receive all of their configured bandwidth before other queues are serviced. Other queues do not transmit traffic until strict-high priority queues are empty, and they receive the bandwidth that remains after the strict-high priority queues are serviced. Because strict-high priority queues are always serviced first, strict-high priority queues can starve other queues on a port. Carefully consider how much bandwidth you want to allocate to strict-high priority queues to avoid starving other queues.

When you define scheduling priorities for queues instead of using the default priorities (by default all queues are low priority), the switch uses the priorities to determine the order of packet transmission from the queues. The switch services traffic of different scheduling priorities in a strict order, using round-robin (RR) scheduling to arbitrate queue transmission service among queues of the same priority. The switch transmits packets in the following order:

1. Strict-high priority traffic within the configured queue transmit rate (on strict-high priority queues, the transmit rate limits the amount of traffic treated as strict-high priority traffic). When traffic arrives on a strict-high priority queue, the switch forwards it before servicing other queues.
2. Low priority traffic within the configured queue transmit rate (on low priority queues, the transmit rate sets the minimum guaranteed bandwidth)
3. All traffic that exceeds the queue transmit rate using weighted round-robin (WRR) scheduling. Traffic that exceeds the queue transmit rate contends for excess port bandwidth (bandwidth that is not consumed after the port meets all guaranteed bandwidth requirements). The switch allocates and weights excess bandwidth for low priority queues based on the configured queue excess rate, or on the transmit rate if no excess rate is configured. The switch allocates and weights excess bandwidth for strict-high priority queues based on the hard-coded weight “1”, which is not configurable. The actual amount of extra bandwidth that traffic exceeding the transmit rate gets depends on how many other queues consume excess bandwidth and the weighting of those queues.



**NOTE:** If you use the default CoS configuration, all queues are low priority queues and transmit traffic based on the weighted round-robin (WRR) algorithm.

---

## Bandwidth Scheduling

A queue scheduler allocates port bandwidth to a queue (the scheduler is mapped to a forwarding class, and the forwarding class is mapped to a queue). The bandwidth profile, which consists of minimum guaranteed bandwidth, maximum bandwidth (queue shaping), and excess bandwidth sharing properties configured in the scheduler, defines the amount of port bandwidth a queue can consume during normal and congested transmission periods.

The scheduler regularly reevaluates whether each individual queue is within its defined bandwidth profile by comparing the amount of data the queue receives to the amount of bandwidth the scheduler allocates to the queue. When the received amount is less than the guaranteed minimum amount of bandwidth, the queue is considered to be in profile. A queue is out of profile when its received amount is larger than its guaranteed minimum amount. Out of profile queue data is transmitted only if extra (excess) bandwidth is available. Otherwise, it is buffered if buffer space is available. If no buffer space is available, the traffic might be dropped.

The switch provides features that enable you to control the allocation of port bandwidth to queues, so that you can meet the demands of different types of traffic on a port:

- [Minimum Guaranteed Bandwidth on page 6735](#)
- [Maximum Bandwidth \(Rate Shaping on Low Priority Queues and LAGs\) on page 6736](#)
- [Limiting Bandwidth Consumed by Strict-High Priority Queues on page 6737](#)
- [Sharing Extra Bandwidth \(Excess Rate on Low Priority Queues\) on page 6738](#)

---

### Minimum Guaranteed Bandwidth

The transmit rate determines the minimum guaranteed bandwidth for each forwarding class that is mapped to an output queue, and so determines the minimum bandwidth guarantee on that queue.

If you do not want to use the default configuration, you can set the minimum guaranteed bandwidth in several ways, and with several options, using the **[set class-of-service schedulers *scheduler-name* transmit-rate (rate | percent *percentage*) <exact>]** statement:

- **Rate**—Set the minimum guaranteed bandwidth as a fixed amount (rate) in bits-per-second of port bandwidth (for example, 2 Gbps or 800 Mbps).
- **Percent**—Set the minimum guaranteed bandwidth as a percentage of port bandwidth (for example, 25 percent).
- **Exact**—Shape the queue to the transmit rate so that the transmit rate is the maximum amount of bandwidth a queue can use. The queue cannot share extra port bandwidth if you configure the exact option. Configuring a transmit rate as *exact* is how you set a shaping rate to configure the maximum amount of bandwidth the low priority queue can consume, and the maximum is the transmit rate. You cannot use the **exact** option on a strict-high priority queue.
- **Extra bandwidth sharing**—On low-priority queues, if you configure an excess rate, the excess rate determines the amount of extra port bandwidth a queue can use. If you

do not configure an excess rate, the transmit rate determines how much excess (extra) bandwidth a low-priority queue can share. If you do not configure an excess rate, then each queue shares extra bandwidth in proportion to its transmit rate.

You cannot configure an excess rate on strict-high priority queues. Strict-high priority queues share extra bandwidth based on a scheduling weight of “1”, which is not configurable. The actual amount of extra bandwidth that traffic exceeding the transmit rate gets depends on how many other queues consume excess bandwidth and the excess rates of those queues.



**NOTE:** The sum of the transmit rates of the queues on a port should not exceed the total bandwidth of that port. (You cannot guarantee a combined minimum bandwidth for the queues on a port that is greater than the total port bandwidth.)



**NOTE:** For transmit rates below 1 Gbps, we recommend that you configure the transmit rate as a percentage instead of as a fixed rate. This is because the system converts fixed rates into percentages and might round small fixed rates to a lower percentage. For example, a fixed rate of 350 Mbps is rounded down to 3 percent.

The bandwidth a low-priority queue consumes can exceed the configured minimum rate if additional bandwidth is available, and if you do not configure the transmit rate as **exact**. During periods of congestion, the configured transmit rate is the guaranteed minimum bandwidth for the queue. This behavior enables you to ensure that each queue receives the amount of bandwidth appropriate to its required level of service and is also able to share unused bandwidth.

### Maximum Bandwidth (Rate Shaping on Low Priority Queues and LAGs)

---

The optional **exact** keyword in the **[set class-of-service schedulers *scheduler-name* transmit-rate (rate | percent *percentage*) <exact>]** configuration statement shapes the transmission rate of low-priority queues. When you specify the **exact** option, the switch drops traffic that exceeds the configured transmit rate, even if excess bandwidth is available. Rate shaping prevents a queue from using more bandwidth than is appropriate for the planned service level of the traffic on the queue. You cannot use the **exact** option on a strict-high priority queue.

Configuring rate shaping on a LAG interface using the **[edit class-of-service interfaces *lag-interface-name* scheduler-map *scheduler-map-name*]** statement can result in scheduled traffic streams receiving more LAG link bandwidth than expected.

LAG interfaces are composed of two or more Ethernet links bundled together to function as a single interface. The switch can hash traffic entering a LAG interface onto any member link in the LAG interface. When you configure a rate shaping and apply it to a LAG interface, the way that the switch applies the rate shaping to traffic depends on how the switch hashes the traffic onto the LAG links.



To illustrate how link hashing affects the way the switch applies rate shaping to LAG traffic, let's look at a LAG interface named **ae0** that has two member links, **xe-0/0/20** and **xe-0/0/21**. On LAG **ae0**, we configure rate shaping of **2g** by including the **transmit-rate 2g exact** statement in the queue scheduler, and apply the scheduler to traffic assigned to the **best-effort** forwarding class, which is mapped to output queue **0**. When traffic in the **best-effort** forwarding class reaches the LAG interface, the switch hashes the traffic onto one of the two member links.

If the switch hashes all of the **best-effort** traffic onto the same LAG link, the traffic receives a maximum of 2g bandwidth on that link. In this case, the intended cumulative limit of 2g for best effort traffic on the LAG is enforced.

However, if the switch hashes the **best-effort** traffic onto both of the LAG links, the traffic receives a maximum of 2g bandwidth on *each* LAG link, not 2g as a cumulative total for the entire LAG. The result is that best-effort traffic receives a maximum of 4g on the LAG, not the 2g set by the rate shaping statement. When hashing spreads the traffic assigned to an output queue (which is mapped to a forwarding class) across multiple LAG links, the effective shaping rate (cumulative maximum bandwidth) on the LAG is:

(number of LAG member interfaces) x (shaping rate for the output queue) = cumulative LAG shaping rate

### Limiting Bandwidth Consumed by Strict-High Priority Queues

You can limit the amount of traffic that receives strict-high priority treatment on a queue by configuring a transmit rate on the strict-high priority queue. The transmit rate sets the amount of traffic that receives strict-high priority treatment. Traffic that exceeds the transmit rate shares in the port excess bandwidth pool based on the strict-high priority excess bandwidth sharing weight of "1", which is not configurable. The actual amount of extra bandwidth that traffic exceeding the transmit rate gets depends on how many other queues consume excess bandwidth and the excess rates of those queues. Limiting the amount of traffic that receives strict-high priority treatment prevents other queues from being starved, while also ensuring that the amount of traffic specified in the transmit rate receives strict-high priority treatment.



**NOTE:** Configuring a transmit rate on a low-priority queue sets the guaranteed minimum bandwidth of the queue, as described in ["Minimum Guaranteed Bandwidth"](#) on page 6735.



**CAUTION:** If you configure strict-high priority queues, we strongly recommend that you configure a transmit rate on the queues to prevent them from starving low priority queues on that port. This is especially important if you configure more than one strict-high priority queue on a port. Although it is not mandatory to configure a transmit rate on strict-high priority queues, if you do not configure a transmit rate, the strict-high priority queues can consume all of the port bandwidth and starve the other queues.

### Sharing Extra Bandwidth (Excess Rate on Low Priority Queues)

---

Extra bandwidth is essentially the bandwidth remaining after the switch meets all guaranteed bandwidth requirements. Extra bandwidth is available to low-priority traffic when the queues on a port do not use all of the available port bandwidth.

By default, extra port bandwidth is shared among the forwarding classes on a port in proportion to the transmit rate of each queue. You can explicitly configure the amount of extra bandwidth a queue can share by setting an **excess-rate** in the scheduler of a low-priority queue. The configured excess rate overrides the transmit rate and determines the percentage of extra bandwidth the queue can consume.



**NOTE:** You cannot configure an excess rate on a strict-high priority queue. Strict-high priority queues share excess bandwidth based on an excess bandwidth sharing weight of “1”, which is not configurable. The actual amount of extra bandwidth that strict-high priority traffic exceeding the transmit rate receives depends on how many other queues consume excess bandwidth and the excess rates of those queues.

An example of extra bandwidth allocation based on transmit rates is a port that has traffic running on three forwarding classes, **best-effort**, **fcoe**, and **network-control**. In this example, the **best-effort** forwarding class has a transmit rate of 2 Gbps, forwarding class **fcoe** has a transmit rate of 4 Gbps, and **network-control** has a transmit rate of 2 Gbps, for a total of 8 Gbps of the port bandwidth. After servicing the minimum guaranteed bandwidth of these three queues, the port has 2 Gbps of available extra bandwidth.

If all three queues still have packets to forward, the queues receive the extra bandwidth in proportion to their transmit rates, so the **best-effort** queue receives an extra 500 Mbps, the **fcoe** queue receives an extra 1 Gbps, and the **network-control** queue receives an extra 500 Mbps.

If you configure an excess rate for a queue, the excess rate determines the proportion of extra bandwidth that the queue receives in the same way that the default (transmit rate) determines the proportion of extra bandwidth a queue receives. In the previous example, if you configured an excess rate of 20 percent on the **fcoe** forwarding class, and the transmit rates of the **best-effort** and **network-control** forwarding classes remained 2g (with no configured excess rate, so the 2g transmit rate for each queue still determines the excess rate), then the 2 Gbps of extra bandwidth would be allocated evenly among the three queues because all three queues have the same excess rate.

In the previous example, if you configured an excess rate of 10 percent on the **fcoe** forwarding class, and the transmit rates of the **best-effort** and **network-control** forwarding classes remained 2g (again with no configured excess rate, so the 2g transmit rate for each queue still determines the excess rate), the 2 Gbps of extra bandwidth would be allocated 800 Mbps to the **best-effort** queue, 400 Mbps to the **fcoe** queue, and 800 Mbps to the **network-control** queue (again, in proportion to the queue excess rates).

## Scheduler Drop-Profile Maps

Drop-profile maps associate drop profiles with queue schedulers and packet loss priorities (PLPs). Drop profiles set thresholds for dropping packets during periods of congestion, based on the queue fill level and a percentage probability of dropping packets at the specified queue fill level. At different fill levels, a drop profile sets different probabilities of dropping a packet during periods of congestion.

Classifiers assign incoming traffic to forwarding classes (which are mapped to output queues), and also assign a PLP to the incoming traffic. The PLP can be low, medium-high, or high. You can classify traffic with different PLPs into the same forwarding class to differentiate treatment of traffic within the forwarding class.

In a drop profile map, you can configure a different drop profile for each PLP and associate (map) the drop profiles to a queue scheduler. A scheduler map maps the queue scheduler to a forwarding class (output queue). Traffic classified into the forwarding class uses the drop characteristics defined in the drop profiles that the drop profile map associates with the queue scheduler. The drop profile the traffic uses depends on the PLP that the classifier assigns to the traffic. (You can map different drop profiles to the forwarding class for different PLPs.)

In summary:

- Classifiers assign one of three PLPs (low, medium-high, high) to incoming traffic when classifiers assign traffic to a forwarding class.
- Drop profiles set thresholds for packet drop at different queue fill levels.
- Drop profile maps associate a drop profile with each PLP, and then map the drop profiles to schedulers.
- Scheduler maps map schedulers to forwarding classes, and forwarding classes are mapped to output queues. The scheduler mapped to a forwarding class determines the CoS characteristics of the output queue mapped to the forwarding class, including the drop profile mapping.

You associate a scheduler map with an interface to apply the drop profiles and other scheduler elements to traffic in the forwarding class mapped to the scheduler on that interface.

## Buffer Size

On QFX10000 switches, the buffer size is the amount of time in milliseconds of port bandwidth that a queue can use to continue to transmit packets during periods of congestion, before the buffer runs out and packets begin to drop.

The switch can use up to 100 ms total (combined) buffer space for all queues on a port. A buffer-size configured as one percent is equal to 1 ms of buffer usage. A buffer-size of 15 percent (the default value for the best effort and network control queues) is equal to 15 ms of buffer usage.

The total buffer size of the switch is 4 GB. A 40-Gigabit port can use up to 500 MB of buffer space, which is equivalent to 100 ms of port bandwidth on a 40-Gigabit port. A

10-Gigabit port can use up to 125 MB of buffer space, which is equivalent to 100 ms of port bandwidth on a 10-Gigabit port. The total buffer sizes of the eight output queues on a port cannot exceed 100 percent, which is equal to the full 100 ms total buffer available to a port. The maximum amount of buffer space any queue can use is also 100 ms (which equates to a 100 percent buffer-size configuration), but if one queue uses all of the buffer, then no other queue receives buffer space.

There is no minimum buffer allocation, so you can set the buffer-size to zero (0) for a queue. However, we recommend that on queues on which you enable PFC to support lossless transport, you allocate a minimum of 5 ms (a minimum buffer-size of 5 percent). The two default lossless queues, fcoe and no-loss, have default buffer-size values of 35 ms (35 percent).



**NOTE:** If you do not configure buffer-size and you do not explicitly configure a queue scheduler, the default buffer-size is the default transmit rate of the queue. If you explicitly configure a queue scheduler, the default buffer allocations are not used. If you explicitly configure a queue scheduler, configure the buffer-size for each queue in the scheduler, keeping in mind that the total buffer-size of the queues cannot exceed 100 percent (100 ms).

If you do not use the default configuration, you can explicitly configure the queue buffer size in either of two ways:

- As a percentage—The queue receives the specified percentage of dedicated port buffers when the queue is mapped to the scheduler and the scheduler is mapped to a port.
- As a remainder—After the port services the queues that have an explicit percentage buffer size configuration, the remaining port dedicated buffer space is divided equally among the other queues to which a scheduler is attached. (No default or explicit scheduler means no dedicated buffer allocation for the queue.) If you configure a scheduler and you do not specify a buffer size as a percentage, *remainder* is the default setting.

Queue buffer allocation is dynamic, shared among ports as needed. However, a queue cannot use more than its configured amount of buffer space. For example, if you are using the default CoS configuration, the best-effort queue receives a maximum of 15 ms of buffer space because the default transmit rate for the best-effort queue is 15 percent.

If a switch experiences congestion, queues continue to receive their full buffer allocation until 90 percent of the 4 GB buffer space is consumed. When 90 percent of the buffer space is in use, the amount of buffer space per port, per queue, is reduced in proportion to the configured buffer size for each queue. As the percentage of consumed buffer space rises above 90 percent, the amount of buffer space per port, per queue, continues to be reduced.

On 40-Gigabit ports, because the total buffer is 4 GB and the maximum buffer a port can use is 500 MB, up to seven 40-Gigabit ports can consume their full 100 ms allocation of buffer space. However, if an eighth 40-Gigabit port requires the full 500 MB of buffer

space, then the buffer allocations are proportionally reduced because the buffer consumption is above 90 percent.

On 10-Gigabit ports, because the total buffer is 4 GB and the maximum buffer a port can use is 125 MB, up to 28 10-Gigabit ports can consume their full 100 ms allocation of buffer space. However, if a 29th 10-Gigabit port requires the full 125 MB of buffer space, then the buffer allocations are proportionally reduced because the buffer consumption is above 90 percent.

## Explicit Congestion Notification

ECN enables end-to-end congestion notification between two endpoints on TCP/IP based networks. The two endpoints are an ECN-enabled sender and an ECN-enabled receiver. ECN must be enabled on both endpoints and on all of the intermediate devices between the endpoints for ECN to work properly. Any device in the transmission path that does not support ECN breaks the end-to-end ECN functionality. ECN notifies networks about congestion with the goal of reducing packet loss and delay by making the sending device decrease the transmission rate until the congestion clears, without dropping packets.

ECN is disabled by default. Normally, you enable ECN only on queues that handle best-effort traffic because other traffic types use different methods of congestion notification—lossless traffic uses priority-based flow control (PFC) and strict-high priority traffic receives all of the port bandwidth it requires up to the point of a configured rate (see [“Scheduling Priority” on page 6733](#)).

## Scheduler Maps

A scheduler map maps a forwarding class to a queue scheduler. After configuring a scheduler, you must include it in a scheduler map, and apply the scheduler map to an interface to implement the configured queue scheduling.

### Related Documentation

- [Understanding Junos CoS Components on page 6621](#)
- [Understanding CoS Priority Group Scheduling on page 6758](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 6765](#)
- [Understanding CoS Virtual Output Queues \(VOQs\) on QFX10000 Switches on page 6724](#)
- [Understanding CoS Explicit Congestion Notification on page 6820](#)
- [Understanding CoS Scheduling Behavior and Configuration Considerations on page 6718](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 6771](#)
- [Example: Configuring Minimum Guaranteed Output Bandwidth on page 6798](#)
- [Example: Configuring Maximum Output Bandwidth on page 6805](#)
- [Example: Configuring Queue Scheduling Priority on page 6753](#)
- [Example: Configuring Queue Schedulers for Port Scheduling on page 6744](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 6762](#)
- [Example: Configuring WRED Drop Profiles on page 6814](#)

- [Example: Configuring ECN on page 6828](#)

## Defining CoS Queue Schedulers for Port Scheduling

---

Schedulers define the CoS properties of output queues. You configure CoS properties in a scheduler, then map the scheduler to a forwarding class. Forwarding classes are in turn mapped to output queues. Classifiers map incoming traffic into forwarding classes based on IEEE 802.1p, DSCP, or EXP code points. CoS scheduling properties include the amount of interface bandwidth assigned to the queue, the priority of the queue, whether explicit congestion notification (ECN) is enabled on the queue, and the WRED packet drop profiles associated with the queue.

The parameters you configure in a scheduler define the following characteristics for the queues mapped to the scheduler:

- **priority**—One of two bandwidth priorities that queues associated with a scheduler can receive:
  - **low**—The scheduler has low priority.
  - **strict-high**—The scheduler has strict-high priority. Strict-high priority queues receive preferential treatment over low-priority queues and receive all of their configured bandwidth before low-priority queues are serviced. Low-priority queues do not transmit traffic until strict-high priority queues are empty.



**NOTE:** We strongly recommend that you configure a transmit rate on all strict-high priority queues to limit the amount of traffic the switch treats as strict-high priority traffic and prevent strict-high priority queues from starving other queues on the port. This is especially important if you configure more than one strict-high priority queue on a port. If you do not configure a transmit rate to limit the amount of bandwidth strict-high priority queues can use, then the strict-high priority queues can use all of the available port bandwidth and starve other queues on the port.

The switch treats traffic in excess of the transmit rate as best-effort traffic that receives bandwidth from the leftover (excess) port bandwidth pool. On strict-high priority queues, all traffic that exceeds the transmit rate shares in the port excess bandwidth pool based on the strict-high priority excess bandwidth sharing weight of “1”, which is not configurable. The actual amount of extra bandwidth that traffic exceeding the transmit rate receives depends on how many other queues consume excess bandwidth and the excess rates of those queues.

- **transmit-rate**—Minimum guaranteed bandwidth, also known as the *committed information rate (CIR)*, set as a percentage rate or as an absolute value in bits per second. By default, the transmit rate also determines the amount of excess (extra) port bandwidth the queue can share if you do not explicitly configure an excess rate.

Extra bandwidth is allocated among the queues on the port in proportion to the transmit rate of each queue. On queues that are not strict-high priority queues, you can configure a transmit rate as **exact**, which shapes the transmission by setting the transmit rate as the maximum bandwidth the queue can consume on the port.

On strict-high priority queues, the transmit rate sets the amount of bandwidth used for strict-high priority forwarding; traffic in excess of the transmit rate is treated as best-effort traffic that receives the queue excess rate.



**NOTE:** Include the preamble bytes and interframe gap (IFG) bytes as well as the data bytes in your bandwidth calculations.

- **excess-rate**—Percentage of extra bandwidth (bandwidth that is not used by other queues) a low-priority queue can receive. If not set, the switch uses the transmit rate to determine extra bandwidth sharing. You cannot set an excess rate on a strict-high priority queue.
- **drop-profile-map**—Drop profile mapping to a packet loss priority to apply WRED to the scheduler and control packet drop for different packet loss priorities during periods of congestion.
- **buffer-size**—Size of the queue buffer as a percentage of the dedicated buffer space on the port, or as a proportional share of the dedicated buffer space on the port that remains after the explicitly configured queues are served.
- **explicit-congestion-notification**—ECN enable on a best-effort queue. ECN enables end-to-end congestion notification between two ECN-enabled endpoints on TCP/IP based networks. ECN must be enabled on both endpoints and on all of the intermediate devices between the endpoints for ECN to work properly. ECN is disabled by default.



**NOTE:** Do not configure drop profiles for the fcoe and no-loss forwarding classes. FCoE and other lossless traffic queues require lossless behavior. Use priority-based flow control (PFC) to prevent frame drop on lossless priorities.

To apply scheduling properties to traffic, map schedulers to forwarding classes using a scheduler map, and then apply the scheduler map to interfaces. Using different scheduler maps, you can map different schedulers to the same forwarding class on different interfaces, to apply different scheduling to that traffic on different interfaces.

To configure a scheduler using the CLI:

1. Name the scheduler and set the minimum guaranteed bandwidth for the queue; optionally, set a maximum bandwidth limit (shaping rate) on a low priority queue by specifying the **exact** option:

```
[edit class-of-service]
user@switch# set schedulers scheduler-name transmit-rate (rate | percent percentage)
<exact>
```

2. Set the amount of excess bandwidth a low-priority queue can share:

```
[edit class-of-service]
user@switch# set schedulers scheduler-name excess-rate percent percentage
```

3. Set the queue priority:

```
[edit class-of-service schedulers scheduler-name]
user@switch# set priority level
```

4. Specify drop profiles for packet loss priorities using a drop profile map:

```
[edit class-of-service schedulers scheduler-name]
user@switch# set drop-profile-map loss-priority (low | medium-high | high) drop-profile
drop-profile-name
```

5. Configure the size of the buffer space for the queue:

```
[edit class-of-service schedulers scheduler-name]
user@switch# set buffer-size (percent percent | remainder)
```

6. Enable ECN, if desired (on best-effort traffic only):

```
[edit class-of-service schedulers scheduler-name]
user@switch# set explicit-congestion-notification
```

7. Configure a scheduler map to map the scheduler to a forwarding class, which applies the scheduler's properties to the traffic in that forwarding class:

```
[edit class-of-service]
user@switch# set scheduler-maps scheduler-map-name forwarding-class
forwarding-class-name scheduler scheduler-name
```

8. Assign the scheduler map and its associated schedulers to one or more interfaces.

```
[edit class-of-service]
user@switch# set interfaces interface-name scheduler-map scheduler-map-name
```

#### Related Documentation

- [Example: Configuring Queue Schedulers for Port Scheduling on page 6744](#)
- [Example: Configuring ECN on page 6828](#)
- [Defining CoS Queue Scheduling Priority on page 6752](#)
- [Configuring CoS WRED Drop Profiles on page 6813](#)
- [Monitoring CoS Scheduler Maps on page 7291](#)
- [Understanding CoS Port Schedulers on QFX Switches on page 6729](#)
- [Understanding CoS Explicit Congestion Notification on page 6820](#)

---

## Example: Configuring Queue Schedulers for Port Scheduling

Schedulers define the CoS properties of output queues. You configure CoS properties in a scheduler, then map the scheduler to a forwarding class. Forwarding classes are in turn mapped to output queues. Classifiers map incoming traffic into forwarding classes based on IEEE 802.1p, DSCP, or EXP code points. CoS scheduling properties include the amount of interface bandwidth assigned to the queue, the priority of the queue, whether explicit



congestion notification (ECN) is enabled on the queue, and the WRED packet drop profiles associated with the queue.

- [Requirements on page 6745](#)
- [Overview on page 6745](#)
- [Configuring a CoS Scheduler on page 6747](#)
- [Verification on page 6748](#)

## Requirements

This example uses the following hardware and software components:

- One switch.
- Junos OS Release 15.1X53-D10 or later for the QFX Series

## Overview

Scheduler parameters define the following characteristics for the queues mapped to the scheduler:

The parameters you configure in a scheduler define the following characteristics for the queues mapped to the scheduler:

- **priority**—One of three bandwidth priorities that queues associated with a scheduler can receive:
  - **low**—The scheduler has low priority.
  - **strict-high**—The scheduler has strict-high priority. Strict-high priority queues receive preferential treatment over low-priority queues and receive all of their configured bandwidth before low-priority queues are serviced. Low-priority queues do not transmit traffic until strict-high priority queues are empty.



**NOTE:** We strongly recommend that you configure a transmit rate on all strict-high priority queues to limit the amount of traffic the switch treats as strict-high priority traffic and prevent strict-high priority queues from starving other queues on the port. This is especially important if you configure more than one strict-high priority queue on a port. If you do not configure a transmit rate to limit the amount of bandwidth strict-high priority queues can use, then the strict-high priority queues can use all of the available port bandwidth and starve other queues on the port.

The switch treats traffic in excess of the transmit rate as best-effort traffic that receives bandwidth from the leftover (excess) port bandwidth pool. On strict-high priority queues, all traffic that exceeds the transmit rate shares in the port excess bandwidth pool based on the strict-high priority excess bandwidth sharing weight of “1”, which is not configurable. The actual amount of extra bandwidth that traffic exceeding the transmit rate receives depends on how many other queues consume excess bandwidth and the excess rates of those queues.

- **transmit-rate**—Minimum guaranteed bandwidth, also known as the *committed information rate (CIR)*, set as a percentage rate or as an absolute value in bits per second. By default, the transmit rate also determines the amount of excess (extra) port bandwidth the queue can share if you do not explicitly configure an excess rate. Extra bandwidth is allocated among the queues on the port in proportion to the transmit rate of each queue. On queues that are not strict-high priority queues, you can configure a transmit rate as **exact**, which shapes the transmission by setting the transmit rate as the maximum bandwidth the queue can consume on the port.

On strict-high priority queues, the transmit rate sets the amount of bandwidth used for strict-high priority forwarding; traffic in excess of the transmit rate is treated as best-effort traffic that receives the queue excess rate.



**NOTE:** Include the preamble bytes and interframe gap (IFG) bytes as well as the data bytes in your bandwidth calculations.

- **excess-rate**—Percentage of extra bandwidth (bandwidth that is not used by other queues) a low-priority queue can receive. If not set, the switch uses the transmit rate to determine extra bandwidth sharing. You cannot set an excess rate on a strict-high priority queue.
- **drop-profile-map**—Drop profile mapping to a packet loss priority to apply WRED to the scheduler and control packet drop for different packet loss priorities during periods of congestion.

- **buffer-size**—Size of the queue buffer as a percentage of the dedicated buffer space on the port, or as a proportional share of the dedicated buffer space on the port that remains after the explicitly configured queues are served.
- **explicit-congestion-notification**—ECN enable on a best-effort queue. ECN enables end-to-end congestion notification between two ECN-enabled endpoints on TCP/IP based networks. ECN must be enabled on both endpoints and on all of the intermediate devices between the endpoints for ECN to work properly. ECN is disabled by default.



**NOTE:** Do not configure drop profiles for the `fcoe` and `no-loss` forwarding classes. FCoE and other lossless traffic queues require lossless behavior. Use priority-based flow control (PFC) to prevent frame drop on lossless priorities.

Scheduler maps map schedulers to forwarding classes, and forwarding classes are mapped to output queues. After you configure schedulers and map them to forwarding classes in a scheduler map, you attach the scheduler map to an interface to implement the configured scheduling on output queues on that interface.

This process configures the bandwidth properties, scheduling, priority, and WRED characteristics that you map to forwarding classes (and thus to output queues) in a scheduler map.

Table 570 shows the configuration components for this example.

**Table 570: Components of the Port Output Queue Scheduler Configuration Example**

| Component     | Settings                                                                                                                                                                                        |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Hardware      | One switch                                                                                                                                                                                      |
| Scheduler     | Name: <b>be-sched</b><br>Transmit rate: <b>20%</b><br>Buffer size: <b>20%</b><br>Excess rate: <b>20%</b><br>Priority: <b>low</b><br>Drop profile: <b>be-dp</b><br>ECN: <b>disable</b> (default) |
| Scheduler map | Name: <b>be-map</b><br>Forwarding class to associate with the <b>be-sched</b> scheduler: <b>best-effort</b>                                                                                     |

## Configuring a CoS Scheduler

**CLI Quick Configuration** To quickly configure a queue scheduler, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level:

```
[edit class-of-service]
set schedulers be-sched transmit-rate percent 20
```

```
set schedulers be-sched buffer-size percent 20
set schedulers be-sched excess-rate percent 20
set schedulers be-sched priority low
set schedulers be-sched drop-profile-map loss-priority low protocol any drop-profile be-dp
set scheduler-maps be-map forwarding-class best-effort scheduler be-sched
set interfaces xe-0/0/7 scheduler-map be-map
```

To configure a CoS scheduler:

1. Create scheduler (**be-sched**) and map it to the drop profile **be-dp**:

```
[edit class-of-service schedulers]
user@switch# set be-sched transmit-rate percent 20
user@switch# set be-sched buffer-size percent 20
user@switch# set be-sched excess-rate percent 20
user@switch# set be-sched priority low
user@switch# set be-sched drop-profile-map loss-priority low protocol any drop-profile
be-dp
```



**NOTE:** Because ECN is disabled by default, no ECN configuration is shown.

2. Configure scheduler map (**be-map**) to associate the scheduler (**be-sched**) with the forwarding class (**best-effort**):

```
[edit class-of-service scheduler-maps]
user@switch# set be-map forwarding-class best-effort scheduler be-sched
```

3. Associate the scheduler map with an interface to apply scheduling to the best-effort forwarding class output queue:

```
[edit class-of-service]
set interfaces xe-0/0/7 scheduler-map be-map
```

## Verification

To verify that the queue scheduler has been created and is mapped to the correct interfaces, perform these tasks:

- [Verifying the Scheduler Configuration on page 6748](#)
- [Verifying the Scheduler Map Configuration on page 6749](#)
- [Verifying That the Scheduler Is Associated with the Interface on page 6749](#)

---

### Verifying the Scheduler Configuration

**Purpose** Verify that the queue scheduler **be-sched** has been created with a minimum guaranteed bandwidth (**transmit-rate**) of 2 Gbps, an extra bandwidth sharing rate (**excess-rate**) of 20 percent, the priority set to **low**, and the drop profile **be-dp**.

**Action** Display the scheduler using the operational mode command **show configuration class-of-service schedulers be-sched**:

```
user@switch> show configuration class-of-service schedulers be-sched
transmit-rate percent 20;
buffer-size percent 20;
excess-rate percent 20;
```

```
priority low;
drop-profile-map loss-priority low protocol any drop-profile be-dp;
```

### Verifying the Scheduler Map Configuration

---

**Purpose** Verify that the scheduler map **be-map** has been created and associates the forwarding class **best-effort** with the scheduler **be-sched**.

**Action** Display the scheduler map using the operational mode command **show configuration class-of-service scheduler-maps be-map**:

```
user@switch> show configuration class-of-service scheduler-maps be-map
forwarding-class best-effort scheduler be-sched;
```

### Verifying That the Scheduler Is Associated with the Interface

---

**Purpose** Verify that the scheduler map **be-sched** is attached to interface **xe-0/0/7**.

**Action** List the interface using the operational mode command **show configuration class-of-service interfaces xe-0/0/7**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/7
scheduler-map be-map;
```

- Related Documentation**
- [Example: Configuring WRED Drop Profiles on page 6814](#)
  - [Example: Configuring ECN on page 6828](#)
  - [Defining CoS Queue Schedulers for Port Scheduling on page 6742](#)
  - [Monitoring CoS Scheduler Maps on page 7291](#)
  - [Understanding CoS Port Schedulers on QFX Switches on page 6729](#)
  - [Understanding CoS Virtual Output Queues \(VOQs\) on QFX10000 Switches on page 6724](#)

## Troubleshooting Egress Bandwidth That Exceeds the Configured Minimum Bandwidth

---

**Problem Description:** The guaranteed minimum bandwidth of a queue (forwarding class) or a priority group (forwarding class set) when measured at the egress port exceeds the guaranteed minimum bandwidth configured for the queue (transmit-rate) or for the priority group (guaranteed-rate).



**NOTE:** On switches that support enhanced transmission selection (ETS) hierarchical scheduling, the switch allocates guaranteed minimum bandwidth first to a priority group using the guaranteed rate setting in the traffic control profile, and then allocates priority group minimum guaranteed bandwidth to forwarding classes in the priority group using the transmit rate setting in the queue scheduler.

On switches that support direct port scheduling, there is no scheduling hierarchy. The switch allocates port bandwidth to forwarding classes directly, using the transmit rate setting in the queue scheduler.

In this topic, if you are using direct port scheduling on your switch, ignore the references to priority groups and forwarding class sets (priority groups and forwarding class sets are only used for ETS hierarchical port scheduling). For direct port scheduling, only the transmit rate queue scheduler setting can cause the issue described in this topic.

**Cause** When you configure bandwidth for a queue or a priority group, the switch accounts for the configured bandwidth as data only. The switch does not include the preamble and the interframe gap (IFG) associated with frames, so the switch does not account for the bandwidth consumed by the preamble and the IFG in its minimum bandwidth calculations.

The measured egress bandwidth can exceed the configured minimum bandwidth when small packet sizes (64 or 128 bytes) are transmitted because the preamble and the IFG are a larger percentage of the total traffic. For larger packet sizes, the preamble and IFG overhead are a small portion of the total traffic, and the effect on egress bandwidth is minor.



**NOTE:** For ETS, the sum of the queue transmit rates in a priority group should not exceed the guaranteed rate for the priority group. (You cannot guarantee a minimum bandwidth for the queues that is greater than the minimum bandwidth guaranteed for the entire set of queues.)

For port scheduling, the sum of the queue transmit rates should not exceed the port bandwidth.

**Solution** When you calculate the bandwidth requirements for queues and priority groups on which you expect a significant amount of traffic with small packet sizes, consider the transmit rate and the guaranteed rate as the minimum bandwidth for the data only. Add sufficient bandwidth to your calculations to account for the preamble and IFG so that the port bandwidth is sufficient to handle the combined minimum data rate and the preamble and IFG.

If the minimum bandwidth measured at the egress port exceeds the amount of bandwidth that you want to allocate to a queue or to a priority group, reduce the transmit rate for that queue and reduce the guaranteed rate of the priority group that contains the queue.

- Related Documentation**
- [transmit-rate on page 7117](#)
  - [Example: Configuring Minimum Guaranteed Output Bandwidth on page 6798](#)

## **Troubleshooting Egress Bandwidth That Exceeds the Configured Maximum Bandwidth**

**Problem**    **Description:** The maximum bandwidth of a queue when measured at the egress port exceeds the maximum bandwidth rate shaper (**shaping-rate** statement on QFX5200, QFX5100, EX4600, QFX3500, QFX3600, and OCX1100 switches, and on QFabric systems, and **transmit-rate (rate | percentage percent exact** statement on QFX10000 switches) configured for the queue.

**Cause**    When you configure bandwidth for a queue (forwarding class) or a priority group (forwarding class set), the switch accounts for the configured bandwidth as data only. The switch does not rate-shape the preamble and the interframe gap (IFG) associated with frames, so the switch does not account for the bandwidth consumed by the preamble and the IFG in its maximum bandwidth calculations.

The measured egress bandwidth can exceed the configured maximum bandwidth when small packet sizes (64 or 128 bytes) are transmitted because the preamble and the IFG are a larger percentage of the total traffic. For larger packet sizes, the preamble and IFG overhead are a small portion of the total traffic, and the effect on egress bandwidth is minor.

**Solution**    When you calculate the bandwidth requirements for queues on which you expect a significant amount of traffic with small packet sizes, consider the shaping rate as the maximum bandwidth for the data only. Add sufficient bandwidth to your calculations to account for the preamble and IFG so that the port bandwidth is sufficient to handle the combined maximum data rate (shaping rate) and the preamble and IFG.

If the maximum bandwidth measured at the egress port exceeds the amount of bandwidth that you want to allocate to the queue, reduce the shaping rate for that queue.

## Defining CoS Queue Scheduling Priority

---

You can configure the scheduling priority of individual queues by specifying the priority in a scheduler, and then associating the scheduler with a queue by using a scheduler map.



**NOTE:** By default, all queues are low priority queues.

The switch services low priority queues after servicing any queue that has strict-high priority traffic. Strict-high priority queues receive preferential treatment over all other queues and receive all of their configured bandwidth before other queues are serviced. Low-priority queues do not transmit traffic until strict-high priority queues are empty, and receive the bandwidth that remains after the strict-high queues have been serviced.

Different switches handle traffic configured as **strict-high** priority traffic in different ways:

- QFX5100, QFX5200, QFX3500, QFX3600, and EX4600 switches, and QFabric systems—You can configure only one queue as a strict-high priority queue.

On these switches, we recommend that you always apply a shaping rate to strict-high priority queues to prevent them from starving other queues. If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.

- QFX10000 switches—You can configure as many queues as you want as strict-high priority. However, keep in mind that too much strict-high priority traffic can starve low priority queues on the port.



**NOTE:** We strongly recommend that you configure a transmit rate on all strict-high priority queues to limit the amount of traffic the switch treats as strict-high priority traffic and prevent strict-high priority queues from starving other queues on the port. This is especially important if you configure more than one strict-high priority queue on a port. If you do not configure a transmit rate to limit the amount of bandwidth strict-high priority queues can use, then the strict-high priority queues can use all of the available port bandwidth and starve other queues on the port.

The switch treats traffic in excess of the transmit rate as best-effort traffic that receives bandwidth from the leftover (excess) port bandwidth pool. On strict-high priority queues, all traffic that exceeds the transmit rate shares in the port excess bandwidth pool based on the strict-high priority excess bandwidth sharing weight of “1”, which is not configurable. The actual amount of extra bandwidth that traffic exceeding the transmit rate receives depends on how many other queues consume excess bandwidth and the excess rates of those queues.



- To configure queue priority using the CLI:

```
[edit class-of-service]
user@switch# set schedulers scheduler-name priority level
```

**Related  
Documentation**

- [Example: Configuring Queue Scheduling Priority on page 6753](#)
- [Monitoring CoS Scheduler Maps on page 7291](#)

---

## Example: Configuring Queue Scheduling Priority

You can configure the bandwidth scheduling priority of individual queues by specifying the priority in a scheduler, and then using a scheduler map to associate the scheduler with a queue.

- [Requirements on page 6753](#)
- [Overview on page 6753](#)
- [Configuring Queue Scheduling Priority on page 6755](#)
- [Verification on page 6756](#)

### Requirements

This example uses the following hardware and software components:

- One switch.
- Junos OS Release 11.1 or later for the QFX Series or Junos OS Release 14.1X53-D20 or later for the OCX Series.

### Overview

Queues can have one of several bandwidth priorities:

- **strict-high**—Strict-high priority allocates bandwidth to the queue before any other queue receives bandwidth. Other queues receive the bandwidth that remains after the strict-high queue has been serviced. On QFX10000 switches, you can configure as many queues as you want as strict-high priority queues. On QFX5200, QFX5100, QFX3500, QFX3600, and EX4600 switches and on QFabric systems, you can configure only one queue as a strict-high queue.



**NOTE:** On QFX5200 switches, it is not possible to support multiple queues with strict-high priority because QFX5200 doesn't support flexible hierarchical scheduling. When multiple strict-high priority queues are configured, all of those queues are treated as strict-high priority but the higher number queue among them is given highest priority.

On QFX10000 switches, if you configure strict-high priority queues on a port, we strongly recommend that you configure a transmit rate on those queues. The transmit rate sets the amount of traffic that the switch forwards as strict-high priority; traffic in excess of the transmit rate is treated as best-effort traffic that receives the queue excess rate.

Even if you configure only one strict-high priority queue, we strongly recommend that you configure a transmit rate for the queue to prevent it from starving other queues. If you do not configure a transmit rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.

On QFX5200, QFX5100, QFX3500, QFX3600, and EX4600 switches and on QFabric systems, we recommend that you always apply a shaping rate to strict-high priority queues to prevent them from starving other queues. If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.



**NOTE:** On switches that support enhanced transmission selection (ETS) hierarchical scheduling, if you use ETS and you configure a strict-high priority queue, you must create an fc-set that is dedicated only to strict-high priority traffic. Only one fc-set can contain strict-high priority queue. Queues that are not strict-high priority cannot belong to the same fc-set as strict-high priority queues.

On switches that use different output queues for unicast and multidestination traffic, the multidestination fc-set cannot contain strict-high priority queues.

- **low**—Low priority. Traffic with low priority is serviced after any queue that has a strict-high priority.



**NOTE:** By default, all queues are low priority queues.

Table 571 shows the configuration components for this example.

This example describes how to set the queue priority for two forwarding classes (queues) named **fcoe** and **no-loss**. Both queues have a priority of **low**. The scheduler for the **fcoe** queue is named **fcoe-sched** and the scheduler for the **no-loss** queue is named **nl-sched**. One scheduler map, **schedmap1**, associates the schedulers to the queues.

**Table 571: Components of the Queue Scheduler Priority Configuration Example**

| Component  | Settings                                                                  |
|------------|---------------------------------------------------------------------------|
| Hardware   | One switch                                                                |
| Schedulers | <b>fcoe-sched</b> for FCoE traffic<br><b>nl-sched</b> for no-loss traffic |
| Priority   | <b>low</b> for FCoE traffic<br><b>low</b> for no-loss traffic             |

Table 571: Components of the Queue Scheduler Priority Configuration Example (*continued*)

| Component     | Settings                                                                                                                                                                                |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Scheduler map | <b>schedmap1:</b><br><br>FCoE mapping: scheduler <b>fcoe-sched</b> to forwarding class <b>fcoe</b><br><br>No-loss mapping: scheduler <b>nl-sched</b> to forwarding class <b>no-loss</b> |



**NOTE:** OCX Series switches do not support lossless transport. On OCX Series switches, the default DSCP classifier does not map traffic to the default fcoe and no-loss forwarding classes. On an OCX Series switch, you could use this example by substituting other forwarding classes (for example, best-effort or network-control) for the fcoe and no-loss forwarding classes, and naming the schedulers appropriately. The active forwarding classes (**best-effort**, **network-control**, and **mcast**) share the unused bandwidth assigned to the fcoe and no-loss forwarding classes.

## Configuring Queue Scheduling Priority

**CLI Quick Configuration** To quickly configure queue scheduling priority, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level:

```
[edit class-of-service]
set schedulers fcoe-sched priority low
set schedulers nl-sched priority low
set scheduler-maps schedmap1 forwarding-class fcoe scheduler fcoe-sched
set scheduler-maps schedmap1 forwarding-class no-loss scheduler nl-sched
```

To configure queue priority using the CLI:

1. Create the FCoE scheduler with **low** priority:

```
[edit class-of-service]
user@switch# set schedulers fcoe-sched priority low
```

2. Create the no-loss scheduler with **low** priority:

```
[edit class-of-service]
user@switch# set schedulers nl-sched priority low
```

3. Associate the schedulers with the desired queues in the scheduler map:

```
[edit class-of-service]
user@switch# set scheduler-maps schedmap1 forwarding-class fcoe scheduler fcoe-sched
user@switch# set scheduler-maps schedmap1 forwarding-class no-loss scheduler nl-sched
```

## Verification

To verify that you configured the queue scheduling priority for bandwidth and mapped the schedulers to the correct forwarding classes, perform these tasks:

- [Verifying the Queue Scheduling Priority on page 6756](#)
- [Verifying the Scheduler-to-Forwarding-Class Mapping on page 6756](#)

---

### Verifying the Queue Scheduling Priority

**Purpose** Verify that you configured the queue schedulers **fcoe-sched** and **nl-sched** with **low** queue scheduling priority.

**Action** Display the **fcoe-sched** scheduler priority configuration using the operational mode command **show configuration class-of-service schedulers fcoe-sched priority**:

```
user@switch> show configuration class-of-service schedulers fcoe-sched priority
priority low;
```

Display the **nl-sched** scheduler priority configuration using the operational mode command **show configuration class-of-service schedulers nl-sched priority**:

```
user@switch> show configuration class-of-service schedulers nl-sched priority
priority low;
```

---

### Verifying the Scheduler-to-Forwarding-Class Mapping

**Purpose** Verify that you configured the scheduler map **schedmap1** to map scheduler **fcoe-sched** to forwarding class **fcoe** and schedule **nl-sched** to forwarding class **no-loss**.

**Action** Display the scheduler map **schedmap1** using the operational mode command **show configuration class-of-service scheduler-maps schedmap1**:

```
user@switch> show configuration class-of-service scheduler-maps schedmap1
forwarding-class fcoe scheduler fcoe-sched;
forwarding-class no-loss scheduler nl-sched;
```

**Related Documentation**

- [Defining CoS Queue Scheduling Priority on page 6752](#)
- [Monitoring CoS Scheduler Maps on page 7291](#)

## Understanding CoS Traffic Control Profiles

A traffic control profile defines the output bandwidth and scheduling characteristics of forwarding class sets (priority groups). The forwarding classes (which are mapped to output queues) that belong to a forwarding class set (fc-set) share the bandwidth that you assign to the fc-set in the traffic control profile.

This two-tier hierarchical scheduling architecture provides flexibility in allocating resources among forwarding classes, and also:

- Assigns a portion of port bandwidth to an fc-set. You define the port resources for the fc-set in a traffic control profile.
- Allocates fc-set bandwidth among the forwarding classes (queues) that belong to the fc-set. A scheduler map attached to the traffic control profile defines the amount of the fc-set's resources that each forwarding class can use.

Attaching an fc-set and a traffic control profile to a port defines the hierarchical scheduling properties of the group and the forwarding classes that belong to the group.

The ability to create fc-sets supports enhanced transmission selection (ETS), which is described in IEEE 802.1Qaz. When an fc-set does not use its allocated port bandwidth, ETS shares the excess port bandwidth among other fc-sets on the port in proportion to their guaranteed minimum bandwidth (guaranteed rate). This utilizes the port bandwidth better than scheduling schemes that reserve bandwidth for groups even if that bandwidth is not used. ETS shares unused port bandwidth, so traffic groups that need extra bandwidth can use it if the bandwidth is available, while preserving the ability to specify the minimum guaranteed bandwidth for traffic groups.

Traffic control profiles define the following CoS properties for fc-sets:

- Minimum guaranteed bandwidth—Also known as the *committed information rate (CIR)*. This is the minimum amount of port bandwidth the priority group receives. Priorities in the priority group receive their minimum guaranteed bandwidth as a portion of the priority group's minimum guaranteed bandwidth. The **guaranteed-rate** statement defines the minimum guaranteed bandwidth.



**NOTE:** You cannot apply a traffic control profile with a minimum guaranteed bandwidth to a priority group that includes strict-high priority queues.

- Shared excess (extra) bandwidth—When the priority groups on a port do not consume the full amount of bandwidth allocated to them or there is unallocated link bandwidth available, priority groups can contend for that extra bandwidth if they need it. Priorities in the priority group contend for extra bandwidth as a portion of the priority group's extra bandwidth. The amount of extra bandwidth for which a priority group can contend is proportional to the priority group's guaranteed minimum bandwidth (guaranteed rate).

- Maximum bandwidth—Also known as *peak information rate (PIR)*. This is the maximum amount of port bandwidth the priority group receives. Priorities in the priority group receive their maximum bandwidth as a portion of the priority group's maximum bandwidth. The **shaping-rate** statement defines the maximum bandwidth.
- Queue scheduling—Each traffic control profile includes a scheduler map. The scheduler map maps forwarding classes (priorities) to schedulers to define the scheduling characteristics of the individual forwarding classes in the fc-set. The resources scheduled for each forwarding class represent portions of the resources that the traffic control profile schedules for the entire fc-set, not portions of the total link bandwidth. The **scheduler-maps** statement defines the mapping of forwarding classes to schedulers.

**Related  
Documentation**

- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 6765](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 6771](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 6762](#)
- [Defining CoS Traffic Control Profiles \(Priority Group Scheduling\) on page 6761](#)

---

## Understanding CoS Priority Group Scheduling

---

Priority group scheduling defines the class-of-service (CoS) properties of a group of output queues (priorities). Priority group scheduling works with output queue scheduling to create a two-tier hierarchical scheduler. The hierarchical scheduler allocates bandwidth to a group of queues (a priority group, called a forwarding class set in Junos OS configuration). Queue scheduling determines the portion of the priority group bandwidth that the particular queue can use.

You configure priority group scheduling in a traffic control profile and then associate the traffic control profile with a forwarding class set and an interface. You attach a scheduler map to the traffic control profile to specify the queue scheduling characteristics.



**NOTE:** When you configure bandwidth for a queue or a priority group, the switch considers only the data as the configured bandwidth. The switch does not account for the bandwidth consumed by the preamble and the interframe gap (IFG). Therefore, when you calculate and configure the bandwidth requirements for a queue or for a priority group, consider the preamble and the IFG as well as the data in the calculations.

- 
- [Priority Group Scheduling Components on page 6759](#)
  - [Default Traffic Control Profile on page 6759](#)
  - [Guaranteed Rate \(Minimum Guaranteed Bandwidth\) on page 6759](#)
  - [Sharing Extra Bandwidth on page 6760](#)
  - [Shaping Rate \(Maximum Bandwidth\) on page 6760](#)
  - [Scheduler Maps on page 6760](#)

## Priority Group Scheduling Components

[Table 572](#) provides a quick reference to the traffic control profile components you can configure to determine the bandwidth properties of priority groups, and [Table 573](#) provides a quick reference to some related scheduling configuration components.

**Table 572: Priority Group Scheduler Components**

| Traffic Control Profile Component | Description                                                                                                                                                                                      |
|-----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Guaranteed rate                   | Sets the minimum guaranteed port bandwidth for the priority group. Extra port bandwidth is shared among priority groups in proportion to the guaranteed rate of each priority group on the port. |
| Shaping rate                      | Sets the maximum port bandwidth the priority group can consume.                                                                                                                                  |
| Scheduler map                     | Maps schedulers to queues (forwarding classes, also called priorities). This determines the portion of the priority group bandwidth that a queue receives.                                       |

**Table 573: Other Scheduling Components**

| Other Scheduling Components | Description                                                                                                                                 |
|-----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------|
| Forwarding class            | Maps traffic to a queue (priority).                                                                                                         |
| Forwarding class set        | Name of a priority group. You map forwarding classes to priority groups. A forwarding class set consists of one or more forwarding classes. |
| Scheduler                   | Sets the bandwidth and scheduling priority of individual queues (forwarding classes).                                                       |

## Default Traffic Control Profile

There is no default traffic control profile.

### Guaranteed Rate (Minimum Guaranteed Bandwidth)

The guaranteed rate determines the minimum guaranteed bandwidth for each priority group. It also determines how much excess (extra) port bandwidth the priority group can share; each priority group shares extra port bandwidth in proportion to its guaranteed rate. You specify the rate in bits per second as a fixed value such as 3 Mbps or as a percentage of the total port bandwidth.

The minimum transmission bandwidth can exceed the configured rate if additional bandwidth is available from other priority groups on the port. In case of congestion, the configured guaranteed rate is guaranteed for the priority group. This property enables you to ensure that each priority group receives the amount of bandwidth appropriate to its level of service.



**NOTE:** Configuring the minimum guaranteed bandwidth (transmit rate) for a forwarding class does not work unless you also configure the minimum guaranteed bandwidth (guaranteed rate) for the forwarding class set in the traffic control profile.

Additionally, the sum of the transmit rates of the queues in a forwarding class set should not exceed the guaranteed rate for the forwarding class set. (You cannot guarantee a minimum bandwidth for the queues that is greater than the minimum bandwidth guaranteed for the entire set of queues.)

You cannot configure a guaranteed rate for forwarding class sets that include strict-high priority queues.

---

## Sharing Extra Bandwidth

Extra bandwidth is available to priority groups when the priority groups do not use the full amount of available port bandwidth. This extra port bandwidth is shared among the priority groups based on the minimum guaranteed bandwidth of each priority group.

For example, Port A has three priority groups: fc-set-1, fc-set-2, and fc-set-3. Fc-set-1 has a guaranteed rate of 2 Gbps, fc-set-2 has a guaranteed rate of 2 Gbps, and fc-set-3 has a guaranteed rate of 4 Gbps. After servicing the minimum guaranteed bandwidth of these priority groups, the port has an extra 2 Gbps of available bandwidth, and all three priority groups have still have packets to forward. The priority groups receive the extra bandwidth in proportion to their guaranteed rates, so fc-set-1 receives an extra 500 Mbps, fc-set-2 receives an extra 500 Mbps, and fc-set-3 receives an extra 1 Gbps.

## Shaping Rate (Maximum Bandwidth)

The shaping rate determines the maximum bandwidth the priority group can consume. You specify the rate in bits per second as a fixed value such as 5 Mbps or as a percentage of the total port bandwidth.

The maximum bandwidth for a priority group depends on the total bandwidth available on the port and how much bandwidth the other priority groups on the port consume.

## Scheduler Maps

A scheduler map maps schedulers to queues. When you associate a scheduler map with a traffic control profile, then associate the traffic control profile with an interface and a forwarding class set, the scheduling defined by the scheduler map determines the portion of the priority group resources that each individual queue can use.

You can associate up to four user-defined scheduler maps with traffic control profiles.

### Related Documentation

- [Understanding Junos CoS Components on page 6621](#)
- [Understanding CoS Output Queue Schedulers](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 6765](#)
- [Understanding CoS Scheduling Behavior and Configuration Considerations on page 6718](#)



- [Understanding CoS Scheduling on QFabric System Node Device Fabric \(fte\) Ports](#)
- [Understanding Default CoS Scheduling on QFabric System Interconnect Devices \(Junos OS Release 13.1 and Later Releases\)](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 6771](#)
- [Example: Configuring Minimum Guaranteed Output Bandwidth on page 6798](#)
- [Example: Configuring Maximum Output Bandwidth on page 6805](#)
- [Example: Configuring Queue Schedulers](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 6762](#)
- [Example: Configuring WRED Drop Profiles](#)
- [Example: Configuring Drop Profile Maps on page 6817](#)

## Defining CoS Traffic Control Profiles (Priority Group Scheduling)

---

A traffic control profile defines the output bandwidth and scheduling characteristics of forwarding class sets (priority groups). The forwarding classes (which are mapped to output queues) contained in a forwarding class set (fc-set) share the bandwidth resources that you configure in the traffic control profile. A scheduler map associates forwarding classes with schedulers to define how the individual forwarding classes that belong to an fc-set share the bandwidth allocated to that fc-set.

The parameters you configure in a traffic control profile define the following characteristics for the fc-set:

- **guaranteed-rate**—Minimum bandwidth, also known as the *committed information rate (CIR)*. The guaranteed rate also determines the amount of excess (extra) port bandwidth that the fc-set can share. Extra port bandwidth is allocated among the fc-sets on a port in proportion to the guaranteed rate of each fc-set.



**NOTE:** You cannot configure a guaranteed rate for a fc-set that includes strict-high priority queues. If the traffic control profile is for an fc-set that contains strict-high priority queues, do not configure a guaranteed rate.

- **shaping-rate**—Maximum bandwidth, also known as the *peak information rate (PIR)*.
- **scheduler-map**—Bandwidth and scheduling characteristics for the queues, defined by mapping forwarding classes to schedulers. (The queue scheduling characteristics represent amounts or percentages of the fc-set bandwidth, not the amounts or percentages of total link bandwidth.)



**NOTE:** Because a port can have more than one fc-set, when you assign resources to an fc-set, keep in mind that the total port bandwidth must serve all of the queues associated with that port.

To configure a traffic control profile using the CLI:

1. Name the traffic control profile and define the minimum guaranteed bandwidth for the fc-set:

```
[edit class-of-service]
user@switch# set traffic-control-profiles traffic-control-profile-name guaranteed-rate (rate
| percent percentage)
```

2. Define the maximum bandwidth for the fc-set:

```
[edit class-of-service traffic-control-profiles traffic-control-profile-name]
user@switch# set shaping-rate (rate | percent percentage)
```

3. Attach a scheduler map to the traffic control profile:

```
[edit class-of-service traffic-control-profiles traffic-control-profile-name]
user@switch# set scheduler-map scheduler-map-name
```

#### Related Documentation

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 6771](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 6762](#)
- [Example: Configuring Minimum Guaranteed Output Bandwidth on page 6798](#)
- [Example: Configuring Maximum Output Bandwidth on page 6805](#)
- [Defining CoS Queue Schedulers](#)
- [Understanding CoS Traffic Control Profiles on page 6757](#)

---

## Example: Configuring Traffic Control Profiles (Priority Group Scheduling)

A traffic control profile defines the output bandwidth and scheduling characteristics of forwarding class sets (priority groups). The forwarding classes (queues) mapped to a forwarding class set share the bandwidth resources that you configure in the traffic control profile. A scheduler map associates forwarding classes with schedulers to define how the individual queues in a forwarding class set share the bandwidth allocated to that forwarding class set.

- [Requirements on page 6762](#)
- [Overview on page 6763](#)
- [Configuring a Traffic Control Profile on page 6764](#)
- [Verification on page 6764](#)

### Requirements

This example uses the following hardware and software components:

- A Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

## Overview

The parameters you configure in a traffic control profile define the following characteristics for the priority group:

- **guaranteed-rate**—Minimum bandwidth, also known as the *committed information rate (CIR)*. Each fc-set receives a minimum of either the configured amount of absolute bandwidth or the configured percentage of bandwidth. The guaranteed rate also determines the amount of excess (extra) port bandwidth that the fc-set can share. Extra port bandwidth is allocated among the fc-sets on a port in proportion to the guaranteed rate of each fc-set.



**NOTE:** In order for the **transmit-rate** option (minimum bandwidth for a queue that you set using scheduler configuration) to work properly, you must configure the **guaranteed-rate** for the fc-set. If an fc-set does not have a guaranteed minimum bandwidth, the forwarding classes that belong to the fc-set cannot have a guaranteed minimum bandwidth.



**NOTE:** Include the preamble bytes and interframe gap bytes as well as the data bytes in your bandwidth calculations.

- **shaping-rate**—Maximum bandwidth, also known as the *peak information rate (PIR)*. Each fc-set receives a maximum of the configured amount of absolute bandwidth or the configured percentage of bandwidth, even if more bandwidth is available.



**NOTE:** Include the preamble bytes and interframe gap bytes as well as the data bytes in your bandwidth calculations.

- **scheduler-map**—Bandwidth and scheduling characteristics for the queues, defined by mapping forwarding classes to schedulers. (The queue scheduling characteristics represent amounts or percentages of the fc-set bandwidth, not the amounts or percentages of total link bandwidth.)



**NOTE:** Because a port can have more than one fc-set, when you assign resources to an fc-set, keep in mind that the total port bandwidth must serve all of the queues associated with that port.

For example, if you map three fc-sets to a 10-Gigabit Ethernet port, the queues associated with all three of the fc-sets share the 10-Gbps bandwidth as defined by the traffic control profiles. Therefore, the total combined **guaranteed-rate** value of the three fc-sets should not exceed 10 Gbps. If you configure guaranteed rates whose sum exceeds the port bandwidth, the system sends a syslog message to notify you that the configuration is not valid. However, the system does not perform a commit check. If you commit a configuration in which the sum of the guaranteed rates exceeds the port bandwidth, the hierarchical scheduler behaves unpredictably.

The sum of the forwarding class (queue) transmit rates cannot exceed the total **guaranteed-rate** of the fc-set to which the forwarding classes belong. If you configure transmit rates whose sum exceeds the fc-set guaranteed rate, the commit check fails and the system rejects the configuration.

If you configure the **guaranteed-rate** of an fc-set as a percentage, configure all of the transmit rates associated with that fc-set as percentages. In this case, if any of the transmit rates are configured as absolute values instead of percentages, the configuration is not valid and the system sends a syslog message.

---

## Configuring a Traffic Control Profile

This example describes how to configure a traffic control profile named **san-tcp** with a scheduler map named **san-map1** and allocate to it a minimum bandwidth of 4 Gbps and a maximum bandwidth of 8 Gbps:

1. Create the traffic control profile and set the **guaranteed-rate** (minimum guaranteed bandwidth) to **4g**:

```
[edit class-of-service]
user@switch# set traffic-control-profiles san-tcp guaranteed-rate 4g
```

2. Set the **shaping-rate** (maximum guaranteed bandwidth) to **8g**:

```
[edit class-of-service]
user@switch# set traffic-control-profiles san-tcp shaping-rate 8g
```

3. Associate the scheduler map **san-map1** with the traffic control profile:

```
[edit class-of-service]
user@switch# set traffic-control-profiles san-tcp scheduler-map san-map1
```

## Verification

---

### Verifying the Traffic Control Profile Configuration

**Purpose** Verify that you created the traffic control profile **san-tcp** with a minimum guaranteed bandwidth of 4 Gbps, a maximum bandwidth of 8 Gbps, and the scheduler map **san-map1**.

**Action** List the traffic control profile using the operational mode command **show configuration class-of-service traffic-control-profiles san-tcp**:

```
user@switch> show configuration class-of-service traffic-control-profiles san-tcp
scheduler-map san-map1;
shaping-rate percent 8g;
guaranteed-rate 4g;
```

**Related  
Documentation**

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 6771](#)
- [Example: Configuring Minimum Guaranteed Output Bandwidth on page 6798](#)
- [Example: Configuring Maximum Output Bandwidth on page 6805](#)
- [Example: Configuring Queue Schedulers](#)
- [Defining CoS Traffic Control Profiles \(Priority Group Scheduling\) on page 6761](#)
- [Understanding CoS Traffic Control Profiles on page 6757](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 6765](#)

---

## Understanding CoS Hierarchical Port Scheduling (ETS)

Scheduling defines the class-of-service (CoS) properties of output queues. Output queues are mapped to forwarding classes. CoS scheduler properties include the amount of interface bandwidth assigned to the queue, the queue priority, and the drop profiles associated with the queue.

Hierarchical port scheduling is a two-tier process that provides better port bandwidth utilization and greater flexibility to allocate resources to queues (forwarding classes) and to groups of queues (forwarding class sets). Hierarchical scheduling includes the Junos OS implementation of enhanced transmission selection (ETS), as described in IEEE 802.1Qaz.



Video: [What is Enhanced Transmission Selection?](#)

---

This topic describes:

- [Hierarchical Scheduling Tiers on page 6765](#)
- [Hierarchical Scheduling and ETS on page 6766](#)
- [ETS Advertisement in DCBX on page 6768](#)
- [Hierarchical Scheduling Process on page 6768](#)
- [Strict-High Priority Queues and Hierarchical Scheduling on page 6769](#)
- [Default Hierarchical Scheduling on page 6770](#)

### Hierarchical Scheduling Tiers

The two tiers used in hierarchical scheduling are priorities and priority groups, as shown in [Table 574](#).

Table 574: Hierarchical Scheduling Tiers

| Junos OS Configuration Construct | Equivalent ETS Construct | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------------------|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Forwarding class                 | Priority                 | <p>Think about priorities (forwarding classes) as output queues. You map forwarding classes to queues, so each forwarding class represents an output queue.</p> <p>When you use a classifier to map a forwarding class to an IEEE 802.1p code point, the code point identifies that traffic's priority for priority-based flow control (PFC). Thus the forwarding class, the queue mapped to the forwarding class, and the priority (code point) mapped to the forwarding class all identify the same traffic.</p> <p><b>NOTE:</b> OCX Series switches do not support lossless transport or PFC.</p> |
| Forwarding class set             | Priority group           | <p>Priority groups (forwarding class sets) are groups of priorities (forwarding classes). Forwarding class membership in a forwarding class set defines the priority group to which each priority belongs.</p> <p>You can configure up to three unicast priority groups and one multicast priority group.</p>                                                                                                                                                                                                                                                                                        |

You apply scheduling properties to each hierarchical scheduling tier as described in the next section.



**NOTE:** If you explicitly configure one or more priority groups on an interface, any priority (forwarding class) that is not assigned to a priority group (forwarding class set) on that interface is assigned to an automatically generated default priority group and receives *no bandwidth*. This means that if you configure hierarchical scheduling on an interface, every forwarding class that you want to forward traffic on that interface must belong to a forwarding class set.



**NOTE:** On OCX Series switches, by default, classifiers use DSCP code points to map traffic to forwarding classes. However, hierarchical scheduling works in the same manner as when you use IEEE 802.1p code points to classify traffic. The OCX Series classifies traffic into forwarding classes based on DSCP code points, the forwarding classes are mapped to forwarding class sets, and you apply scheduling properties to each of the two tiers.


## Hierarchical Scheduling and ETS

Two-tier hierarchical scheduling manages bandwidth efficiently by enabling you to define the CoS properties for each priority group and for each priority. The first tier of the hierarchical scheduler allocates port bandwidth to a priority group. The second tier of

the hierarchical scheduler determines the portion of the priority group bandwidth that a priority (queue) can use.

The CoS properties of a priority group define the amount of port bandwidth resources available to the queues in that priority group. The CoS properties you configure for each queue specify the amount of the bandwidth available to the queue from the bandwidth allocated to the priority group. [Figure 218](#) shows the relationship of port resource allocation to priority groups, and priority group resource allocation to queues (priorities).

**Figure 218: Hierarchical Scheduling Tiers**



9040722

If a queue (priority) does not use its allocated bandwidth, ETS shares the unused bandwidth among the other queues in the priority group in proportion to the minimum guaranteed rate (transmit rate) scheduled for each queue. If a priority group does not use its allocated bandwidth, ETS shares the unused bandwidth among the priority groups on the port in proportion to the minimum guaranteed rate (guaranteed rate) scheduled for each priority group.

In this way, ETS improves link bandwidth utilization, and it provides each queue and each priority group with the maximum available bandwidth. For example, priorities that consist of bursty traffic can share bandwidth during periods of low traffic transmission, instead of reserving their entire bandwidth allocation when traffic loads are light. All Juniper switches use ETS scheduling, except for QFX5200 and the QFX10000 switches.



**NOTE:** The available link bandwidth is the bandwidth remaining after servicing strict-high priority flows. Strict-high priority takes precedence over all other traffic (we recommend that you configure a shaping rate to limit the maximum amount of bandwidth that a strict-high priority forwarding class can use to prevent starving other queues).

## ETS Advertisement in DCBX

When you configure hierarchical scheduling on a port, Data Center Bridging Capability Exchange protocol (DCBX) advertises:

- Each priority group
- The priorities in each priority group
- The bandwidth properties of each priority group and priority

When you configure hierarchical scheduling on a port, any priority that is not part of an explicitly configured priority group is assigned to the automatically generated default priority group and receives no bandwidth. The default priority group is transparent. It does not appear in the configuration.



**NOTE:** OCX Series switches do support DCBX, so hierarchical scheduling information is not exchanged with connected peers on OCX Series switches.

## Hierarchical Scheduling Process

Hierarchical scheduling consists of multiple configuration steps that create the priorities and the priority groups, schedule their resources, and assign them to interfaces. The steps below correspond to the six blocks in the packet flow diagram shown in [Figure 219](#):

1. Packet classification:
  - Configure classification of incoming traffic into forwarding classes (priorities). This consists of either using the default classifiers or configuring classifiers to map code points and loss priorities to the forwarding classes.
  - Apply the classifiers to ingress interfaces or use the default classifiers. Applying a classifier to an interface groups incoming traffic on the interface into forwarding classes and loss priorities, by applying the classifier code point mapping to the incoming traffic.
2. Configure the output queues for the forwarding classes (priorities). This consists of either using the default forwarding classes and forwarding-class-to-queue mapping, or creating your own forwarding classes and mapping them to output queues.
3. Allocate resources to the forwarding classes:
  - Define resources for the priorities. This consists of configuring schedulers to set minimum guaranteed bandwidth, maximum bandwidth, drop profiles for Weighted Random Early Detection (WRED), and bandwidth priority to apply to a forwarding class. Extra bandwidth is shared among queues in proportion to the minimum guaranteed bandwidth (transmit rate) of each queue.
  - Map resources to priorities. This consists of mapping forwarding classes to schedulers, using a scheduler map.




4. Configure priority groups. This consists of mapping forwarding classes (priorities) to forwarding class sets (priority groups) to define the priorities that belong to each priority group.
5. Define resources for the priority groups. This consists of configuring traffic control profiles to set minimum guaranteed bandwidth (guaranteed rate) and maximum bandwidth (shaping rate) for a priority group. Traffic control profiles also specify a scheduler map, which defines the resources (schedulers) mapped to the priorities in the priority group. Extra port bandwidth is shared among priority groups in proportion to the minimum guaranteed bandwidth of each priority group.

The traffic control profile bandwidth settings determine the port resources available to the priority group. The schedulers specified in the scheduler map determine the amount of priority group resources that each priority receives.

6. Apply hierarchical scheduling to a port. This consists of attaching one or more priority groups (forwarding class sets) to an interface. For each priority group, you also attach a traffic control profile, which contains the scheduling properties of the priority group and the priorities in the priority group. Different priority groups on the same port can use different traffic control profiles, which provides fine tuned control of scheduling for each queue on each interface.

Figure 219: Hierarchical Scheduling Packet Flow



## Strict-High Priority Queues and Hierarchical Scheduling

If you configure a strict-high priority queue, you must observe the following rules:

- You must create a separate forwarding class set (priority group) for the strict-high priority queue.
- Only one forwarding class set can contain strict-high priority queues.
- Strict-high priority queues cannot belong to the same forwarding class set as queues that are not strict-high priority.

- A strict-high priority queue cannot belong to a multidestination forwarding class set.
- We recommend that you always apply a shaping rate to strict-high priority queues to prevent them from starving other queues. If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.



**NOTE:** On a QFabric system, if a fabric (fte) interface handles strict-high priority traffic, you must define a separate forwarding class set (priority group) for strict-high priority traffic. Strict-high priority traffic cannot be mixed with traffic of other priorities in a forwarding class set. For example, you might choose to create different forwarding class sets for best effort, lossless, strict-high priority, and multidestination traffic.

---

## Default Hierarchical Scheduling

If you do not explicitly configure hierarchical scheduling, the switch uses the default settings:

- The switch automatically creates a default forwarding class set that contains all of the forwarding classes on the switch. The switch assigns 100 percent of the port output bandwidth to the default forwarding class set. The default forwarding class set is transparent. It does not appear in the configuration and is used for Data Center Bridging Capability Exchange protocol (DCBX) advertisement.



**NOTE:** OCX Series switches do not support DCBX, so the ETS configuration is not advertised to connected peers.

- Ingress traffic is classified based on the default classifier settings.
- The forwarding classes (queues) in the default forwarding class set receive bandwidth based on the default scheduler settings.

### Related Documentation

- [Understanding CoS Packet Flow on page 6628](#)
- [Understanding CoS Output Queue Schedulers](#)
- [Understanding CoS Priority Group Scheduling on page 6758](#)
- [Benefits of Configuring CoS Hierarchical Port Scheduling](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 6493](#)
- [Understanding CoS Classifiers](#)
- [Understanding CoS Classifiers](#)
- [Understanding Default CoS Scheduling and Classification](#)
- [Understanding Default CoS Scheduling and Classification](#)

- *Understanding CoS Scheduling on QFabric System Node Device Fabric (fte) Ports*
- *Understanding Default CoS Scheduling on QFabric System Interconnect Devices (Junos OS Release 13.1 and Later Releases)*
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 6771](#)
- [Example: Configuring Queue Schedulers](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 6762](#)
- [Example: Configuring Minimum Guaranteed Output Bandwidth on page 6798](#)
- [Example: Configuring Maximum Output Bandwidth on page 6805](#)

---

## Example: Configuring CoS Hierarchical Port Scheduling (ETS)

---

Hierarchical port scheduling defines the class-of-service (CoS) properties of output queues, which are mapped to forwarding classes. Traffic is classified into forwarding classes based on code point (priority), so mapping queues to forwarding classes also maps queues to priorities). Hierarchical port scheduling enables you to group priorities that require similar CoS treatment into priority groups. You define the port bandwidth resources for a priority group, and you define the amount of the priority group's resources that each priority in the group can use.

Hierarchical port scheduling is the Junos OS implementation of enhanced transmission selection (ETS), as described in IEEE 802.1Qaz. One major benefit of hierarchical port scheduling is greater port bandwidth utilization. If a priority group on a port does not use all of its allocated bandwidth, other priority groups on that port can use that bandwidth. Also, if a priority within a priority group does not use its allocated bandwidth, other priorities within that priority group can use that bandwidth.

Configuring hierarchical scheduling is a multistep procedure that includes:

- Mapping forwarding classes to queues
- Defining forwarding class sets (priority groups)
- Defining behavior aggregate classifiers
- Configuring priority-based flow control (PFC) for lossless priorities (queues)
- Applying classifiers and PFC configuration to ingress interfaces
- Defining drop profiles
- Defining schedulers
- Mapping forwarding classes to schedulers
- Defining traffic control profiles
- Assigning priority groups and traffic control profiles to egress ports



NOTE: OCX Series switches do not support lossless transport and do not support PFC. Although this example includes configuring lossless transport with PFC, the portions of the example that do not pertain to lossless transport still apply to OCX Series switches. (You can configure hierarchical scheduling on OCX Series switches, but you cannot configure lossless transport or lossless forwarding classes.)

This example describes how to configure hierarchical scheduling:

- [Requirements on page 6772](#)
- [Overview on page 6772](#)
- [Configuration on page 6776](#)
- [Verification on page 6786](#)

## Requirements

This example uses the following hardware and software components:

- One switch (this example was tested on a Juniper Networks QFX3500 Switch)
- Junos OS Release 11.1 or later for the QFX Series or Junos OS Release 14.1X53-D20 or later for the OCX Series

## Overview

Keep the following considerations in mind when you plan the port bandwidth allocation for priority groups and for individual priorities:

- How much traffic and what types of traffic you expect to traverse the system.
- How you want to divide different types of traffic into priorities (forwarding classes) to apply different CoS treatment to different types of traffic. Dividing traffic into priorities includes:
  - Mapping the code points of ingress traffic to forwarding classes using behavior aggregate (BA) classifiers. This classifies incoming traffic into the appropriate forwarding class based on code point.
  - Mapping forwarding classes to output queues. This defines the output queue for each type of traffic.
  - Attaching the BA classifier to the desired ingress interfaces so that incoming traffic maps to the desired forwarding classes and queues.
- How you want to organize priorities into priority groups (forwarding class sets).

Traffic that requires similar treatment usually belongs in the same priority group. To do this, place forwarding classes that require similar bandwidth, loss, and other

characteristics in the same forwarding class set. For example, you can map all types of best-effort traffic forwarding classes into one forwarding class set.

- How much of the port bandwidth you want to allocate to each priority group and to each of the priorities in each priority group. The following considerations apply to bandwidth allocation:
  - Estimate how much traffic you expect in each forwarding class, and how much traffic you expect in each forwarding class set (the amount of traffic you expect in a forwarding class set is the aggregate amount of traffic in the forwarding classes that belong to the forwarding class set).
  - The combined minimum guaranteed bandwidth of the priorities (forwarding classes) in a priority group should not exceed the minimum guaranteed bandwidth of the priority group (forwarding class set). The transmit rate scheduler parameter defines the minimum guaranteed bandwidth for forwarding classes. Scheduler maps associate schedulers with forwarding classes.
  - The combined minimum guaranteed bandwidth of the priority groups (forwarding class sets) on a port should not exceed the port's total bandwidth. The guaranteed rate parameter in the traffic control profile defines the minimum bandwidth for a forwarding class set. Associating a scheduler map with a traffic control profile sets the scheduling for the individual forwarding classes in the forwarding class set.

This example creates hierarchical port scheduling by defining priority groups for best effort, guaranteed delivery, and high-performance computing (HPC) traffic. Each priority group includes priorities that need to receive similar CoS treatment. Each priority group and each priority within each priority group receive the CoS resources needed to service their flows. Lossless priorities use PFC to prevent packet loss when the network experiences congestion.

### Topology

Table 575 shows the configuration components for this example.



**NOTE:** OCX Series switches do not support lossless transport and do not support PFC. If you eliminate the configuration elements for the default lossless fcoe and no-loss forwarding classes (including classifier, forwarding class set, scheduler, and traffic control profile configuration for those forwarding classes) and for PFC, this example works for OCX Series switches. However, because the default fcoe and no-loss forwarding classes do not carry traffic on OCX Series switches, you can apply the bandwidth allocated to those forwarding classes to other forwarding classes. By default, the active forwarding classes (best-effort, network-control, and mcast) share the unused bandwidth assigned to the fcoe and no-loss forwarding classes.

**Table 575: Components of the Hierarchical Port Scheduling (ETS) Configuration Topology**

| Property | Settings       |
|----------|----------------|
| Hardware | QFX3500 switch |

Table 575: Components of the Hierarchical Port Scheduling (ETS) Configuration Topology (*continued*)

| Property                                                                                                                 | Settings                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Mapping of forwarding classes (priorities) to queues                                                                     | <p><b>best-effort</b> to queue 0</p> <p><b>be2</b> to queue 1</p> <p><b>fcoe</b> (Fibre Channel over Ethernet) to queue 3</p> <p><b>no-loss</b> to queue 4</p> <p><b>hpc</b> (high-performance computing) to queue 5</p> <p><b>network-control</b> to queue 7</p> <p><b>NOTE:</b> On switches that do not support the ELS CLI, if you are using Junos OS Release 12.2 or later, use the default forwarding-class-to-queue mapping for the lossless <b>fcoe</b> and <b>no-loss</b> forwarding classes. If you explicitly configure the default lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (<b>best-effort</b>) traffic and does <i>not</i> receive lossless treatment.</p> <p>On switches that do not support the ELS CLI, in Junos OS Release 12.3 and later, you can include the <i>no-loss</i> packet drop attribute in the explicit forwarding class configuration to configure a lossless forwarding class.</p> |
| Forwarding class sets (priority groups)                                                                                  | <p><b>best-effort-pg:</b> contains forwarding classes <b>best-effort</b>, <b>be2</b>, and <b>network control</b></p> <p><b>guar-delivery-pg:</b> contains forwarding classes <b>fcoe</b> and <b>no-loss</b></p> <p><b>hpc-pg:</b> contains forwarding class <b>hpc</b></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Behavior aggregate classifier (maps forwarding classes and loss priorities to incoming packets by IEEE 802.1 code point) | <p>Name—<b>hsclassifier1</b></p> <p>Code point mapping:</p> <ul style="list-style-type: none"> <li>• <b>000</b> to forwarding class <b>best-effort</b> and loss priority <b>low</b></li> <li>• <b>001</b> to forwarding class <b>be2</b> and loss priority <b>high</b></li> <li>• <b>011</b> to forwarding class <b>fcoe</b> and loss priority <b>low</b></li> <li>• <b>100</b> to forwarding class <b>no-loss</b> and loss priority <b>low</b></li> <li>• <b>101</b> to forwarding class <b>hpc</b> and loss priority <b>low</b></li> <li>• <b>110</b> to forwarding class <b>network-control</b> and loss priority <b>low</b></li> </ul>                                                                                                                                                                                                                                                                                                                               |
| PFC                                                                                                                      | <p>Congestion notification profile name—<b>gd-cnp</b></p> <p>PFC enabled on code points: <b>011</b> (<b>fcoe</b> priority), <b>010</b> (<b>no-loss</b> priority)</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Drop profiles                                                                                                            | <p><b>dp-be-low:</b> drop start point 25, drop end point 50, maximum drop rate 80</p> <p><b>NOTE:</b> The <b>fcoe</b> and <b>no-loss</b> priorities (queues) do not use drop profiles because they are lossless traffic classes.</p> <p><b>dp-be-high:</b> drop start point 10, drop end point 40, maximum drop rate 100</p> <p><b>dp-hpc:</b> drop start point 75, drop end point 90, maximum drop rate 75</p> <p><b>dp-nc:</b> drop start point 80, drop end point 100, maximum drop rate 100</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

**Table 575: Components of the Hierarchical Port Scheduling (ETS) Configuration Topology** (*continued*)

| Property                              | Settings                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Queue schedulers                      | <p><b>be-sched</b>: minimum bandwidth <b>3g</b>, maximum bandwidth <b>100%</b>, priority <b>low</b>, drop profiles <b>dp-be-low</b> and <b>dp-be-high</b></p> <p><b>fcoe-sched</b>: minimum bandwidth <b>2.5g</b>, maximum bandwidth <b>100%</b>, priority <b>low</b></p> <p><b>hpc-sched</b>: minimum bandwidth <b>2g</b>, maximum bandwidth <b>100%</b>, priority <b>low</b>, drop profile <b>dp-hpc</b></p> <p><b>nc-sched</b>: minimum bandwidth <b>500m</b>, maximum bandwidth <b>100%</b>, priority <b>low</b>, drop profile <b>dp-nc</b></p> <p><b>nl-sched</b>: minimum bandwidth <b>2g</b>, maximum bandwidth <b>100%</b>, priority <b>low</b></p> |
| Forwarding class-to-scheduler mapping | <p>Scheduler map <b>be-map</b>:<br/>Forwarding class <b>best-effort</b>, scheduler <b>be-sched</b><br/>Forwarding class <b>be2</b>, scheduler <b>be-sched</b><br/>Forwarding class <b>network-control</b>, scheduler <b>nc-sched</b></p> <p>Scheduler map <b>gd-map</b>:<br/>Forwarding class <b>fcoe</b>, scheduler <b>fcoe-sched</b><br/>Forwarding class <b>no-loss</b>, scheduler <b>nl-sched</b></p> <p>Scheduler map <b>hpc-map</b>:<br/>Forwarding class <b>hpc</b>, scheduler <b>hpc-sched</b></p>                                                                                                                                                  |
| Traffic control profiles              | <p><b>be-tcp</b>: scheduler map <b>be-map</b>, minimum bandwidth <b>3.5g</b>, maximum bandwidth <b>100%</b></p> <p><b>gd-tcp</b>: scheduler map <b>gd-map</b>, minimum bandwidth <b>4.5g</b>, maximum bandwidth <b>100%</b></p> <p><b>hpc-tcp</b>: scheduler map <b>hpc-map</b>, minimum bandwidth <b>2g</b>, maximum bandwidth <b>100%</b></p>                                                                                                                                                                                                                                                                                                             |
| Interfaces                            | <p>This example configures hierarchical port scheduling on interfaces <b>xe-0/0/20</b> and <b>xe-0/0/21</b>. Because traffic is bidirectional, you apply the ingress and egress configuration components to both interfaces:</p> <ul style="list-style-type: none"> <li>• Classifier Name—<b>hsclassifier1</b></li> <li>• Forwarding class sets—<b>best-effort-pg</b>, <b>guar-deliver-pg</b>, <b>hpc-pg</b></li> <li>• Congestion notification profile—<b>gd-cnp</b></li> </ul>                                                                                                                                                                            |

**Figure 220** shows a block diagram of the configuration components and the configuration flow of the CLI statements used in the example. You can perform the configuration steps in a different sequence if you want.

Figure 220: Hierarchical Port Scheduling Components Block Diagram

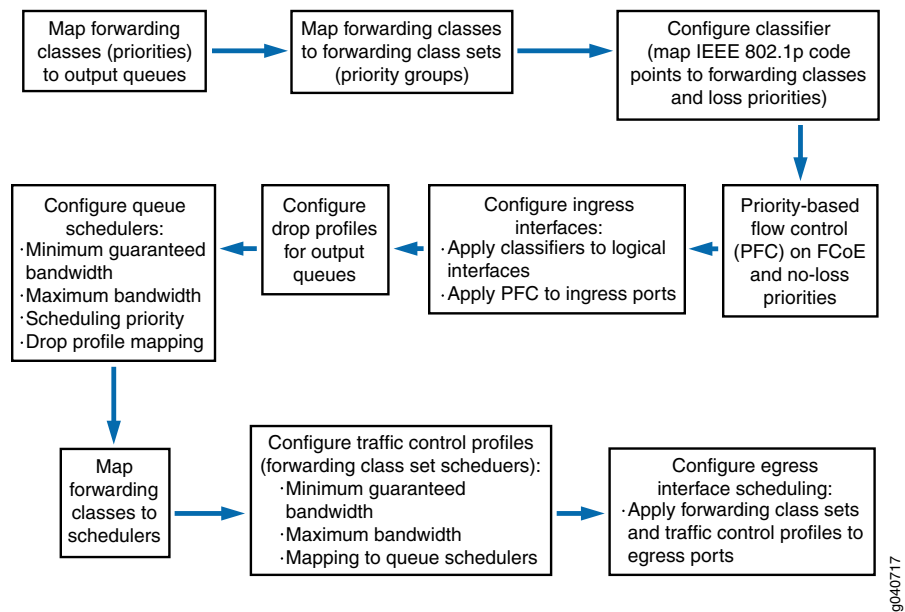



Figure 221 shows a block diagram of the hierarchical scheduling packet flow from ingress to egress.

Figure 221: Hierarchical Port Scheduling Packet Flow Block Diagram



## Configuration

### CLI Quick Configuration

To quickly configure hierarchical port scheduling on systems that support lossless transport, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the [edit class-of-service] hierarchy level:

```
[edit class-of-service]
set forwarding-classes class best-effort queue-num 0
set forwarding-classes class be2 queue-num 1
set forwarding-classes class hpc queue-num 5
set forwarding-classes class network-control queue-num 7
```



```

set forwarding-class-sets best-effort-pg class best-effort
set forwarding-class-sets best-effort-pg class be2
set forwarding-class-sets best-effort-pg class network-control
set forwarding-class-sets guar-delivery-pg class fcoe
set forwarding-class-sets guar-delivery-pg class no-loss
set forwarding-class-sets hpc-pg class hpc
set classifiers ieee-802.1 hsclassifier1 forwarding-class best-effort loss-priority low code-points 000
set classifiers ieee-802.1 hsclassifier1 forwarding-class be2 loss-priority high code-points 001
set classifiers ieee-802.1 hsclassifier1 forwarding-class fcoe loss-priority low code-points 011
set classifiers ieee-802.1 hsclassifier1 forwarding-class no-loss loss-priority low code-points 100
set classifiers ieee-802.1 hsclassifier1 forwarding-class hpc loss-priority low code-points 101
set classifiers ieee-802.1 hsclassifier1 forwarding-class network-control loss-priority low code-points 110
set congestion-notification-profile gd-cnp input ieee-802.1 code-point 011 pfc
set congestion-notification-profile gd-cnp input ieee-802.1 code-point 100 pfc
set interfaces xe-0/0/20 unit 0 classifiers ieee-802.1 hsclassifier1
set interfaces xe-0/0/21 unit 0 classifiers ieee-802.1 hsclassifier1
set interfaces xe-0/0/20 congestion-notification-profile gd-cnp
set interfaces xe-0/0/21 congestion-notification-profile gd-cnp
set drop-profiles dp-be-low interpolate fill-level 25 fill-level 50 drop-probability 0 drop-probability 80
set drop-profiles dp-be-high interpolate fill-level 10 fill-level 40 drop-probability 0 drop-probability 100
set drop-profiles dp-nc interpolate fill-level 80 fill-level 100 drop-probability 0 drop-probability 100
set drop-profiles dp-hpc interpolate fill-level 75 fill-level 90 drop-probability 0 drop-probability 75
set schedulers be-sched priority low transmit-rate 3g
set schedulers be-sched shaping-rate percent 100
set schedulers be-sched drop-profile-map loss-priority low protocol any drop-profile dp-be-low
set schedulers be-sched drop-profile-map loss-priority high protocol any drop-profile dp-be-high
set schedulers fcoe-sched priority low transmit-rate 2500m
set schedulers fcoe-sched shaping-rate percent 100
set schedulers hpc-sched priority low transmit-rate 2g
set schedulers hpc-sched shaping-rate percent 100
set schedulers hpc-sched drop-profile-map loss-priority low protocol any drop-profile dp-hpc
set schedulers nc-sched priority low transmit-rate 500m
set schedulers nc-sched shaping-rate percent 100
set schedulers nc-sched drop-profile-map loss-priority low protocol any drop-profile dp-nc
set schedulers nl-sched priority low transmit-rate 2g
set schedulers nl-sched shaping-rate percent 100
set scheduler-maps be-map forwarding-class best-effort scheduler be-sched
set scheduler-maps be-map forwarding-class be2 scheduler be-sched
set scheduler-maps be-map forwarding-class network-control scheduler nc-sched
set scheduler-maps gd-map forwarding-class fcoe scheduler fcoe-sched
set scheduler-maps gd-map forwarding-class no-loss scheduler nl-sched
set scheduler-maps hpc-map forwarding-class hpc scheduler hpc-sched
set traffic-control-profiles be-tcp scheduler-map be-map guaranteed-rate 3500m
set traffic-control-profiles gd-tcp scheduler-map gd-map guaranteed-rate 4500m
set traffic-control-profiles gd-tcp shaping-rate percent 100
set traffic-control-profiles hpc-tcp scheduler-map hpc-map guaranteed-rate 2g
set traffic-control-profiles hpc-tcp shaping-rate percent 100
set interfaces xe-0/0/20 forwarding-class-set best-effort-pg output-traffic-control-profile be-tcp
set interfaces xe-0/0/20 forwarding-class-set guar-delivery-pg output-traffic-control-profile gd-tcp
set interfaces xe-0/0/20 forwarding-class-set hpc-pg output-traffic-control-profile hpc-tcp
set interfaces xe-0/0/21 forwarding-class-set best-effort-pg output-traffic-control-profile be-tcp

```

```
set interfaces xe-0/0/21 forwarding-class-set guar-delivery-pg output-traffic-control-profile
gd-tcp
set interfaces xe-0/0/21 forwarding-class-set hpc-pg output-traffic-control-profile hpc-tcp
```

## OCX Series Switches

Because OCX Series switches do not support lossless transport, the following subset of the configuration eliminates the lossless configuration elements and provides hierarchical port scheduling for the best-effort, be2, hpc, and network-control forwarding classes. In addition, on OCX Series switches, you would probably use DSCP classifiers and code points instead of IEEE classifiers and code points. To quickly configure hierarchical port scheduling on an OCX Series switch, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the [edit class-of-service] hierarchy level:

```
[edit class-of-service]
set forwarding-classes class best-effort queue-num 0
set forwarding-classes class be2 queue-num 1
set forwarding-classes class hpc queue-num 5
set forwarding-classes class network-control queue-num 7
set forwarding-class-sets best-effort-pg class best-effort
set forwarding-class-sets best-effort-pg class be2
set forwarding-class-sets best-effort-pg class network-control
set forwarding-class-sets hpc-pg class hpc
set classifiers ieee-802.1 hsclassifier1 forwarding-class best-effort loss-priority low code-points
000
set classifiers ieee-802.1 hsclassifier1 forwarding-class be2 loss-priority high code-points 001
set classifiers ieee-802.1 hsclassifier1 forwarding-class hpc loss-priority low code-points 101
set classifiers ieee-802.1 hsclassifier1 forwarding-class network-control loss-priority low
code-points 110
set interfaces xe-0/0/20 unit 0 classifiers ieee-802.1 hsclassifier1
set interfaces xe-0/0/21 unit 0 classifiers ieee-802.1 hsclassifier1
set drop-profiles dp-be-low interpolate fill-level 25 fill-level 50 drop-probability 0 drop-probability
80
set drop-profiles dp-be-high interpolate fill-level 10 fill-level 40 drop-probability 0 drop-probability
100
set drop-profiles dp-nc interpolate fill-level 80 fill-level 100 drop-probability 0 drop-probability
100
set drop-profiles dp-hpc interpolate fill-level 75 fill-level 90 drop-probability 0 drop-probability
75
set schedulers be-sched priority low transmit-rate 3g
set schedulers be-sched shaping-rate percent 100
set schedulers be-sched drop-profile-map loss-priority low protocol any drop-profile dp-be-low
set schedulers be-sched drop-profile-map loss-priority high protocol any drop-profile dp-be-high
set schedulers hpc-sched priority low transmit-rate 2g
set schedulers hpc-sched shaping-rate percent 100
set schedulers hpc-sched drop-profile-map loss-priority low protocol any drop-profile dp-hpc
set schedulers nc-sched priority low transmit-rate 500m
set schedulers nc-sched shaping-rate percent 100
set schedulers nc-sched drop-profile-map loss-priority low protocol any drop-profile dp-nc
set scheduler-maps be-map forwarding-class best-effort scheduler be-sched
set scheduler-maps be-map forwarding-class be2 scheduler be-sched
set scheduler-maps be-map forwarding-class network-control scheduler nc-sched
set scheduler-maps hpc-map forwarding-class hpc scheduler hpc-sched
set traffic-control-profiles be-tcp scheduler-map be-map guaranteed-rate 3500m
set traffic-control-profiles be-tcp shaping-rate percent 100
set traffic-control-profiles hpc-tcp scheduler-map hpc-map guaranteed-rate 2g
set traffic-control-profiles hpc-tcp shaping-rate percent 100
```

```

set interfaces xe-0/0/20 forwarding-class-set best-effort-pg output-traffic-control-profile be-tcp
set interfaces xe-0/0/20 forwarding-class-set hpc-pg output-traffic-control-profile hpc-tcp
set interfaces xe-0/0/21 forwarding-class-set best-effort-pg output-traffic-control-profile be-tcp
set interfaces xe-0/0/21 forwarding-class-set hpc-pg output-traffic-control-profile hpc-tcp

```

### Step-by-Step Procedure

To perform a step-by-step configuration of the forwarding classes (priorities), forwarding class sets (priority groups), classifiers, queue schedulers, PFC, traffic control profiles, and interfaces to set up hierarchical port scheduling (ETS):

1. Configure the forwarding classes (priorities) and map them to unicast output queues (do not explicitly map the **fcoe** and **no-loss** forwarding classes to output queues; use the default configuration):

```

[edit class-of-service]
user@switch# set forwarding-classes class best-effort queue-num 0
user@switch# set forwarding-classes class be2 queue-num 1
user@switch# set forwarding-classes class hpc queue-num 5
user@switch# set forwarding-classes class network-control queue-num 7

```

2. Configure forwarding class sets (priority groups) to group forwarding classes (priorities) that require similar CoS treatment:

```

[edit class-of-service]
user@switch# set forwarding-class-sets best-effort-pg class best-effort
user@switch# set forwarding-class-sets best-effort-pg class be2
user@switch# set forwarding-class-sets best-effort-pg class network-control
user@switch# set forwarding-class-sets guar-delivery-pg class fcoe
user@switch# set forwarding-class-sets guar-delivery-pg class no-loss
user@switch# set forwarding-class-sets hpc-pg class hpc

```



**NOTE:** On OCX Series switches, you would not configure the **guar-delivery-pg** forwarding class set for lossless traffic.

3. Configure a classifier to set the loss priority and IEEE 802.1 code points assigned to each forwarding class at the ingress:

```

[edit class-of-service]
user@switch# set classifiers ieee-802.1 hsclassifier1 forwarding-class best-effort
loss-priority low code-points 000
user@switch# set classifiers ieee-802.1 hsclassifier1 forwarding-class be2 loss-priority
high code-points 001
user@switch# set classifiers ieee-802.1 hsclassifier1 forwarding-class fcoe loss-priority
low code-points 011
user@switch# set classifiers ieee-802.1 hsclassifier1 forwarding-class no-loss loss-priority
low code-points 100
user@switch# set classifiers ieee-802.1 hsclassifier1 forwarding-class hpc loss-priority low
code-points 101
user@switch# set classifiers ieee-802.1 hsclassifier1 forwarding-class network-control
loss-priority low code-points 110

```



**NOTE:** On OCX Series switches, you would not configure the **fcoe** and **no-loss** portions of the classifier.

4. Configure a congestion notification profile to enable PFC on the FCoE and no-loss queue IEEE 802.1 code points:

```
[edit class-of-service]
user@switch# set congestion-notification-profile gd-cnp input ieee-802.1 code-point 011
pfc
user@switch# set congestion-notification-profile gd-cnp input ieee-802.1 code-point 100
pfc
```



**NOTE:** This step does not apply to OCX Series switches, which do not support PFC.

5. Assign the classifier to the interfaces:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 unit 0 classifiers ieee-802.1 hsclassifier1
user@switch# set interfaces xe-0/0/21 unit 0 classifiers ieee-802.1 hsclassifier1
```

6. Apply the PFC configuration to the interfaces:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 congestion-notification-profile gd-cnp
user@switch# set interfaces xe-0/0/21 congestion-notification-profile gd-cnp
```



**NOTE:** This step does not apply to OCX Series switches, which do not support PFC.

7. Configure the drop profile for the best-effort low loss-priority queue:

```
[edit class-of-service]
user@switch# set drop-profiles dp-be-low interpolate fill-level 25 fill-level 50
drop-probability 0 drop-probability 80
```

8. Configure the drop profile for the best-effort high loss-priority queue:

```
[edit class-of-service]
user@switch# set drop-profiles dp-be-high interpolate fill-level 10 fill-level 40
drop-probability 0 drop-probability 100
```

9. Configure the drop profile for the network-control queue:

```
[edit class-of-service]
user@switch# set drop-profiles dp-nc interpolate fill-level 80 fill-level 100 drop-probability
0 drop-probability 100
```

10. Configure the drop profile for the high-performance computing queue:

```
[edit class-of-service]
user@switch# set drop-profiles dp-hpc interpolate fill-level 75 fill-level 90 drop-probability
0 drop-probability 75
```

11. Define the minimum guaranteed bandwidth, priority, maximum bandwidth, and drop profiles for the best-effort queue:

```
[edit class-of-service]
user@switch# set schedulers be-sched priority low transmit-rate 3g
user@switch# set schedulers be-sched shaping-rate percent 100
```

```

user@switch# set schedulers be-sched drop-profile-map loss-priority low protocol any
drop-profile dp-be-low
user@switch# set schedulers be-sched drop-profile-map loss-priority high protocol any
drop-profile dp-be-high

```

12. Define the minimum guaranteed bandwidth, priority, and maximum bandwidth for the FCoE queue:

```

[edit class-of-service]
user@switch# set schedulers fcoe-sched priority low transmit-rate 2500m
user@switch# set schedulers fcoe-sched shaping-rate percent 100

```



**NOTE:** This step does not apply to OCX Series switches, which do not support lossless transport.

13. Define the minimum guaranteed bandwidth, priority, maximum bandwidth, and drop profile for the high-performance computing queue:

```

[edit class-of-service]
user@switch# set schedulers hpc-sched priority low transmit-rate 2g
user@switch# set schedulers hpc-sched shaping-rate percent 100
user@switch# set schedulers hpc-sched drop-profile-map loss-priority low protocol any
drop-profile dp-hpc

```

14. Define the minimum guaranteed bandwidth, priority, maximum bandwidth, and drop profile for the network-control queue:

```

[edit class-of-service]
user@switch# set schedulers nc-sched priority low transmit-rate 500m
user@switch# set schedulers nc-sched shaping-rate percent 100
user@switch# set schedulers nc-sched drop-profile-map loss-priority low protocol any
drop-profile dp-nc

```

15. Define the minimum guaranteed bandwidth, priority, and maximum bandwidth for the no-loss queue:

```

[edit class-of-service]
user@switch# set schedulers nl-sched priority low transmit-rate 2g
user@switch# set schedulers nl-sched shaping-rate percent 100

```



**NOTE:** This step does not apply to OCX Series switches, which do not support lossless transport.

16. Map the schedulers to the appropriate forwarding classes (queues):

```

[edit class-of-service]
user@switch# set scheduler-maps be-map forwarding-class best-effort scheduler be-sched
user@switch# set scheduler-maps be-map forwarding-class be2 scheduler be-sched
user@switch# set scheduler-maps be-map forwarding-class network-control scheduler
nc-sched
user@switch# set scheduler-maps gd-map forwarding-class fcoe scheduler fcoe-sched
user@switch# set scheduler-maps gd-map forwarding-class no-loss scheduler nl-sched
user@switch# set scheduler-maps hpc-map forwarding-class hpc scheduler hpc-sched

```



**NOTE:** On OCX Series switches, because lossless transport is not supported, you would not configure the `gd-map` scheduler map.

17. Define the traffic control profile for the best-effort priority group (queue scheduler to mapping, minimum guaranteed bandwidth, and maximum bandwidth):

```
[edit class-of-service]
user@switch# set traffic-control-profiles be-tcp scheduler-map be-map guaranteed-rate
3500m
user@switch# set traffic-control-profiles be-tcp shaping-rate percent 100
```

18. Define the traffic control profile for the guaranteed delivery priority group (queue to scheduler mapping, minimum guaranteed bandwidth, and maximum bandwidth):

```
[edit class-of-service]
user@switch# set traffic-control-profiles gd-tcp scheduler-map gd-map guaranteed-rate
4500m
user@switch# set traffic-control-profiles gd-tcp shaping-rate percent 100
```



**NOTE:** This step does not apply to OCX Series switches, which do not support lossless transport.

19. Define the traffic control profile for the high-performance computing priority group (queue to scheduler mapping, minimum guaranteed bandwidth, and maximum bandwidth):

```
[edit class-of-service]
user@switch# set traffic-control-profiles hpc-tcp scheduler-map hpc-map guaranteed-rate
2g
user@switch# set traffic-control-profiles hpc-tcp shaping-rate percent 100
```

20. Apply the three priority groups (forwarding class sets) and the appropriate traffic control profiles to the egress ports:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 forwarding-class-set best-effort-pg
output-traffic-control-profile be-tcp
user@switch# set interfaces xe-0/0/20 forwarding-class-set guar-delivery-pg
output-traffic-control-profile gd-tcp
user@switch# set interfaces xe-0/0/20 forwarding-class-set hpc-pg
output-traffic-control-profile hpc-tcp
user@switch# set interfaces xe-0/0/21 forwarding-class-set best-effort-pg
output-traffic-control-profile be-tcp
user@switch# set interfaces xe-0/0/21 forwarding-class-set guar-delivery-pg
output-traffic-control-profile gd-tcp
user@switch# set interfaces xe-0/0/21 forwarding-class-set hpc-pg
output-traffic-control-profile hpc-tcp
```



**NOTE:** Because OCX Series switches do not support lossless transport, on OCX Series switches, you would not apply the `guar-deliver-pg` forwarding class set and the `gd-tcp` traffic control profile to interfaces.

## Results

Display the results of the configuration (the system shows only the explicitly configured parameters; it does not show default parameters such as the **fcoe** and **no-loss** lossless forwarding classes). On OCX Series switches, you would not see the lossless configuration components in the output:

```
user@switch> show configuration class-of-service
classifiers {
 ieee-802.1 hsclassifier1 {
 forwarding-class best-effort {
 loss-priority low code-points 000;
 }
 forwarding-class be2 {
 loss-priority high code-points 001;
 }
 forwarding-class fcoe {
 loss-priority low code-points 011;
 }
 forwarding-class no-loss {
 loss-priority low code-points 100;
 }
 forwarding-class hpc {
 loss-priority low code-points 101;
 }
 forwarding-class network-control {
 loss-priority low code-points 110;
 }
 }
}
drop-profiles {
 dp-be-low {
 interpolate {
 fill-level [25 50];
 drop-probability [0 80];
 }
 }
 dp-be-high {
 interpolate {
 fill-level [10 40];
 drop-probability [0 100];
 }
 }
 dp-hpc {
 interpolate {
 fill-level [75 90];
 drop-probability [0 75];
 }
 }
 dp-nc {
 interpolate {
 fill-level [80 100];
 drop-probability [0 100];
 }
 }
}
```

```
forwarding-classes {
 class best-effort queue-num 0;
 class be2 queue-num 1;
 class hpc queue-num 5;
 class network-control queue-num 7;
}
traffic-control-profiles {
 be-tcp {
 scheduler-map be-map;
 shaping-rate percent 100;
 guaranteed-rate 3500000000;
 }
 gd-tcp {
 scheduler-map gd-map;
 shaping-rate percent 100;
 guaranteed-rate 4500000000;
 }
 hpc-tcp {
 scheduler-map hpc-map;
 shaping-rate percent 100;
 guaranteed-rate 2g;
 }
}
forwarding-class-sets {
 guar-delivery-pg {
 class fcoe;
 class no-loss;
 }
 best-effort-pg {
 class best-effort;
 class be2;
 class network-control;
 }
 hpc-pg {
 class hpc;
 }
}
congestion-notification-profile {
 gd-cnp {
 input {
 ieee-802.1 {
 code-point 011 {
 pfc;
 }
 code-point 100 {
 pfc;
 }
 }
 }
 }
}
interfaces {
 xe-0/0/20 {
 forwarding-class-set {
 best-effort-pg {
 output-traffic-control-profile be-tcp;
 }
 }
 }
}
```



```

 }
 guar-delivery-pg {
 output-traffic-control-profile gd-tcp;
 }
 hpc-pg {
 output-traffic-control-profile hpc-tcp;
 }
}
congestion-notification-profile gd-cnp;
unit 0 {
 classifiers {
 ieee-802.1 hsclassifier1;
 }
}
}
xe-0/0/21 {
 forwarding-class-set {
 best-effort-pg {
 output-traffic-control-profile be-tcp;
 }
 guar-delivery-pg {
 output-traffic-control-profile gd-tcp;
 }
 hpc-pg {
 output-traffic-control-profile hpc-tcp;
 }
 }
 congestion-notification-profile gd-cnp;
 unit 0 {
 classifiers {
 ieee-802.1 hsclassifier1;
 }
 }
}
}
scheduler-maps {
 be-map {
 forwarding-class best-effort scheduler be-sched;
 forwarding-class network-control scheduler nc-sched;
 forwarding-class be2 scheduler be-sched;
 }
 gd-map {
 forwarding-class fcoe scheduler fcoe-sched;
 forwarding-class no-loss scheduler nl-sched;
 }
 hpc-map {
 forwarding-class hpc scheduler hpc-sched;
 }
}
schedulers {
 be-sched {
 transmit-rate 3g;
 shaping-rate percent 100;
 priority low;
 drop-profile-map loss-priority low protocol any drop-profile dp-be-low;
 drop-profile-map loss-priority high protocol any drop-profile dp-be-high;
 }
}

```

```
}
fcoe-sched {
 transmit-rate 2500000000;
 shaping-rate percent 100;
 priority low;
}
hpc-sched {
 transmit-rate 2g;
 shaping-rate percent 100;
 priority low;
 drop-profile-map loss-priority low protocol any drop-profile dp-hpc;
}
nc-sched {
 transmit-rate 500m;
 shaping-rate percent 100;
 priority low;
 drop-profile-map loss-priority low protocol any drop-profile dp-nc;
}
nl-sched {
 transmit-rate 2g;
 shaping-rate percent 100;
 priority low;
}
}
```



**TIP:** To quickly configure the interfaces, issue the `load merge terminal` command, and then copy the hierarchy and paste it into the switch terminal window.

## Verification



**NOTE:** The verification output is based on the full example configuration. On OCX Series switches, you do not see lossless configuration components in the output. Comments about lossless configuration components do not apply to OCX Series switches.

To verify that you created the hierarchical port scheduling components and they are operating properly, perform these tasks:

- [Verifying the Forwarding Classes \(Priorities\) on page 6787](#)
- [Verifying the Forwarding Class Sets \(Priority Groups\) on page 6787](#)
- [Verifying the Classifier on page 6788](#)
- [Verifying Priority-Based Flow Control on page 6788](#)
- [Verifying the Output Queue Schedulers on page 6789](#)
- [Verifying the Drop Profiles on page 6792](#)

- [Verifying the Priority Group Output Schedulers \(Traffic Control Profiles\) on page 6793](#)
- [Verifying the Interface Configuration on page 6794](#)

### Verifying the Forwarding Classes (Priorities)

**Purpose** Verify that you created the forwarding classes and mapped them to the correct queues. (The system shows only the explicitly configured forwarding classes. It does not show default forwarding classes such as **fcoe** and **no-loss**.)

**Action** List the forwarding classes using the operational mode command **show class-of-service forwarding-class**:

```
user@switch> show class-of-service forwarding-class
```

| Forwarding class | ID | Queue | Policing priority | No-Loss  |
|------------------|----|-------|-------------------|----------|
| best-effort      | 0  | 0     | normal            | Disabled |
| be2              | 1  | 3     | normal            | Disabled |
| hpc              | 2  | 4     | normal            | Disabled |
| network-control  | 3  | 7     | normal            | Disabled |
| mcast            | 8  | 8     | normal            | Disabled |

**Meaning** The **show class-of-service forwarding-class** command lists all of the configured forwarding classes, the internal identification number of each forwarding class, the queues that are mapped to the forwarding classes, the policing priority, and whether the forwarding class is lossless (no-loss packet drop attribute enabled) or lossy forwarding class (no-loss packet drop attribute disabled). The command output shows that:

- Forwarding class **best-effort** maps to queue **0** and is lossy
- Forwarding class **be2** maps to queue **1** and is lossy
- Forwarding class **hpc** maps to queue **5** and is lossy
- Forwarding class **network-control** maps to queue **7** and is lossy

In addition, the command lists the default multicast (multidestination) forwarding class and the default queue to which it is mapped.

### Verifying the Forwarding Class Sets (Priority Groups)

**Purpose** Verify that you created the priority groups and that the correct priorities (forwarding classes) belong to the appropriate priority group.

**Action** List the forwarding class sets using the operational mode command **show class-of-service forwarding-class-set**:

```
user@switch> show class-of-service forwarding-class-set
```

```
Forwarding class set: best-effort-pg, Type: normal-type, Forwarding class set
index: 19907
```

| Forwarding class | Index |
|------------------|-------|
| best-effort      | 0     |

|                 |   |
|-----------------|---|
| be2             | 1 |
| network-control | 5 |

Forwarding class set: guar-delivery-pg, Type: normal-type, Forwarding class set index: 43700

| Forwarding class | Index |
|------------------|-------|
| fcoe             | 2     |
| no-loss          | 3     |

Forwarding class set: hpc-pg, Type: normal-type, Forwarding class set index: 60758

| Forwarding class | Index |
|------------------|-------|
| hpc              | 4     |

**Meaning** The `show class-of-service forwarding-class-set` command lists all of the configured forwarding class sets (priority groups), the forwarding classes (priorities) that belong to each priority group, and the internal index number of each priority group. The command output shows that:

- The forwarding class set **best-effort-pg** includes the forwarding classes **best-effort**, **be2**, and **network-control**.
- The forwarding class set **guar-delivery-pg** includes the forwarding classes **fcoe** and **no-loss**.
- The forwarding class set **hpc-pg** includes the forwarding class **hpc**.

### Verifying the Classifier

**Purpose** Verify that the classifier maps forwarding classes to the correct IEEE 802.1p code points and packet loss priorities.

**Action** List the classifier configured for hierarchical port scheduling using the operational mode command `show class-of-service classifier name hsclassifier1`:

```
user@switch> show class-of-service classifier name hsclassifier1
Classifier: hsclassifier1, Code point type: ieee-802.1, Index: 43607
Code point Forwarding class Loss priority
000 best-effort low
001 be2 high
011 fcoe low
100 no-loss low
101 hpc low
110 network-control low
```

**Meaning** The `show class-of-service classifier name hsclassifier1` command lists all of the IEEE 802.1p code points and the loss priorities mapped to all of the forwarding classes in the classifier. The command output shows that the forwarding classes **best-effort**, **be2**, **no-loss**, **fcoe**, **hpc**, and **network-control** have been created and mapped to IEEE 802.1p code points and loss priorities.

### Verifying Priority-Based Flow Control

**Purpose** Verify that PFC is enabled on the correct priorities for lossless transport.

**Action** List the congestion notification profiles using the operational mode command **show class-of-service congestion-notification**:

```
user@switch> show class-of-service congestion-notification
```

```
Type: Input, Name: gd-cnp, Index: 51687
```

```
Cable Length: 100 m
```

| Priority | PFC      | MRU  |
|----------|----------|------|
| 000      | Disabled |      |
| 001      | Disabled |      |
| 010      | Disabled |      |
| 011      | Enabled  | 2500 |
| 100      | Enabled  | 2500 |
| 101      | Disabled |      |
| 110      | Disabled |      |
| 111      | Disabled |      |

```
Type: Output
```

| Priority | Flow-Control-Queues |
|----------|---------------------|
| 000      | 0                   |
| 001      | 0                   |
| 010      | 1                   |
| 011      | 2                   |
| 100      | 3                   |
| 101      | 4                   |
| 110      | 5                   |
| 111      | 6                   |
|          | 7                   |

**Meaning** The **show class-of-service congestion-notification** command lists all of the congestion notification profiles and the IEEE 802.1p code points with PFC enabled. The command output shows that PFC is enabled for code points **011** (**fcoe** priority and queue) and **100** (**no-loss** priority and queue) for the **gd-cnp** congestion notification profile.

The command also shows the default cable length (100 meters), the default maximum receive unit (2500 bytes), and the default mapping of priorities to output queues because this example does not include configuring these options.

### Verifying the Output Queue Schedulers

**Purpose** Verify that you created the output queue schedulers with the correct bandwidth parameters and priorities, mapped to the correct queues, and mapped to the correct drop profiles.

**Action** List the scheduler maps using the operational mode command **show class-of-service scheduler-map**:

```
user@switch> show class-of-service scheduler-map
```

```
Scheduler map: be-map, Index: 64023
```

```
Scheduler: be-sched, Forwarding class: best-effort, Index: 13005
Transmit rate: 3000000000 bps, Rate Limit: none, Buffer size: remainder,
```

Buffer Limit: none, Priority: low  
 Excess Priority: unspecified  
 Shaping rate: 100 percent,  
 drop-profile-map-set-type: mark  
 Drop profiles:

| Loss priority | Protocol | Index | Name                   |
|---------------|----------|-------|------------------------|
| Low           | any      | 55387 | dp-be-low              |
| Medium high   | any      | 1     | <default-drop-profile> |
| High          | any      | 4369  | dp-be-high             |

Scheduler: be-sched, Forwarding class: be2, Index: 13005  
 Transmit rate: 3000000000 bps, Rate Limit: none, Buffer size: remainder,  
 Buffer Limit: none, Priority: low  
 Excess Priority: unspecified  
 Shaping rate: 100 percent,  
 drop-profile-map-set-type: mark  
 Drop profiles:

| Loss priority | Protocol | Index | Name                   |
|---------------|----------|-------|------------------------|
| Low           | any      | 55387 | dp-be-low              |
| Medium high   | any      | 1     | <default-drop-profile> |
| High          | any      | 4369  | dp-be-high             |

Scheduler: nc-sched, Forwarding class: network-control, Index: 45740  
 Transmit rate: 5000000000 bps, Rate Limit: none, Buffer size: remainder,  
 Buffer Limit: none, Priority: low  
 Excess Priority: unspecified  
 Shaping rate: 100 percent,  
 drop-profile-map-set-type: mark  
 Drop profiles:

| Loss priority | Protocol | Index | Name                   |
|---------------|----------|-------|------------------------|
| Low           | any      | 44207 | dp-nc                  |
| Medium high   | any      | 1     | <default-drop-profile> |
| High          | any      | 1     | <default-drop-profile> |

Scheduler map: gd-map, Index: 61447

Scheduler: fcoe-sched, Forwarding class: fcoe, Index: 37289  
 Transmit rate: 25000000000 bps, Rate Limit: none, Buffer size: remainder,  
 Buffer Limit: none, Priority: low  
 Excess Priority: unspecified  
 Shaping rate: 100 percent,  
 drop-profile-map-set-type: mark  
 Drop profiles:

| Loss priority | Protocol | Index | Name                   |
|---------------|----------|-------|------------------------|
| Low           | any      | 44207 | <default-drop-profile> |
| Medium high   | any      | 1     | <default-drop-profile> |
| High          | any      | 1     | <default-drop-profile> |

Scheduler: nl-sched, Forwarding class: no-loss, Index: 29359  
 Transmit rate: 20000000000 bps, Rate Limit: none, Buffer size: remainder,  
 Buffer Limit: none, Priority: low  
 Excess Priority: unspecified  
 Shaping rate: 100 percent,  
 drop-profile-map-set-type: mark  
 Drop profiles:

| Loss priority | Protocol | Index | Name                   |
|---------------|----------|-------|------------------------|
| Low           | any      | 44207 | <default-drop-profile> |
| Medium high   | any      | 1     | <default-drop-profile> |
| High          | any      | 1     | <default-drop-profile> |

Scheduler map: hpc-map, Index: 56941

```
Scheduler: hpc-sched, Forwarding class: hpc, Index: 55900
Transmit rate: 2000000000 bps, Rate Limit: none, Buffer size: remainder,
Buffer Limit: none, Priority: low
Excess Priority: unspecified
Shaping rate: 100 percent,
drop-profile-map-set-type: mark
Drop profiles:
 Loss priority Protocol Index Name
 Low any 57716 dp-hpc
 Medium high any 1 <default-drop-profile>
 High any 1 <default-drop-profile>
```

**Meaning** The **show class-of-service scheduler-map** command lists all of the configured scheduler maps. For each scheduler map, the command output includes:

- The name of the scheduler map (**scheduler-map** field)
- The name of the scheduler (**scheduler** field)
- The forwarding classes mapped to the scheduler (**forwarding-class** field)
- The minimum guaranteed queue bandwidth (**transmit-rate** field)
- The scheduling priority (**priority** field)
- The maximum bandwidth in the priority group the queue can consume (**shaping-rate** field)
- The drop profile loss priority (**loss priority** field) for each drop profile name (**name** field)

The command output shows that:

- The scheduler map **be-map** was created and has these properties:
  - There are two schedulers, **be-sched** and **nc-sched**.
  - The scheduler **be-sched** has two forwarding classes, **best-effort** and **be2**.
  - Scheduler **be-sched** forwarding classes **best-effort** and **be2** share a minimum guaranteed bandwidth of **3,000,000,000 bps**, can consume a maximum of **100 percent** of the priority group bandwidth, and use the drop profile **dp-be-low** for low loss-priority traffic, the default drop profile for medium-high loss-priority traffic, and the drop profile **dp-be-high** for high loss-priority traffic.
  - The scheduler **nc-sched** has one forwarding class, **network-control**.
  - The **network-control** forwarding class has a minimum guaranteed bandwidth of **500,000,000 bps**, can consume a maximum of **100 percent** of the priority group bandwidth, and uses the drop profile **dp-nc** for low loss-priority traffic and the default drop profile for medium-high and high loss priority traffic.
- The scheduler map **gd-map** was created and has these properties:
  - There are two schedulers, **fcoe-sched** and **nl-sched**.
  - The scheduler **fcoe-sched** has one forwarding class, **fcoe**.

- The **fcoe** forwarding class has a minimum guaranteed bandwidth of **2,500,000,000 bps**, and can consume a maximum of **100 percent** of the priority group bandwidth.
- The scheduler **nl-sched** has one forwarding class, **no-loss**.
- The **no-loss** forwarding class has a minimum guaranteed bandwidth of **2,000,000,000 bps**, and can consume a maximum of **100 percent** of the priority group bandwidth.
- The scheduler map **hpc-map** was created and has these properties:
  - There is one scheduler, **hpc-sched**.
  - The scheduler **hpc-sched** has one forwarding class, **hpc**.
  - The **hpc** forwarding class has a minimum guaranteed bandwidth of **2,000,000,000 bps**, can consume a maximum of **100 percent** of the priority group bandwidth, and uses the drop profile **dp-hpc** for low loss-priority traffic and the default drop profile for medium-high and high loss-priority traffic.

---

### Verifying the Drop Profiles

**Purpose** Verify that you created the drop profiles **dp-be-high**, **dp-be-low**, **dp-hpc**, and **dp-nc** with the correct fill levels and drop probabilities.

**Action** List the drop profiles using the operational mode command **show configuration class-of-service drop-profiles**:

```
user@switch> show configuration class-of-service drop-profiles
dp-be-low {
 interpolate {
 fill-level [25 50];
 drop-probability [0 80];
 }
}
dp-be-high {
 interpolate {
 fill-level [10 40];
 drop-probability [0 100];
 }
}
dp-hpc {
 interpolate {
 fill-level [75 90];
 drop-probability [0 75];
 }
}
dp-nc {
 interpolate {
 fill-level [80 100];
 drop-probability [0 100];
 }
}
```

**Meaning** The **show configuration class-of-service drop-profiles** command lists the drop profiles and their properties. The command output shows that there are four drop profiles configured, **dp-be-high**, **dp-be-low**, **dp-hpc**, and **dp-nc**. The output also shows that:



- For **dp-be-low**, the drop start point (the first fill level) is when the queue is 25 percent filled, the drop end point (the second fill level) occurs when the queue is 50 percent filled, and the drop probability at the drop end point is 80 percent.
- For **dp-be-high**, the drop start point (the first fill level) is when the queue is 10 percent filled, the drop end point (the second fill level) occurs when the queue is 40 percent filled, and the drop probability at the drop end point is 100 percent.
- For **dp-hpc**, the drop start point (the first fill level) is when the queue is 75 percent filled, the drop end point (the second fill level) occurs when the queue is 90 percent filled, and the drop probability at the drop end point is 75 percent.
- For **dp-nc**, the drop start point (the first fill level) is when the queue is 80 percent filled, the drop end point (the second fill level) occurs when the queue is 100 percent filled, and the drop probability at the drop end point is 100 percent.

### Verifying the Priority Group Output Schedulers (Traffic Control Profiles)

|                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Purpose</b> | Verify that you created the traffic control profiles <b>be-tcp</b> , <b>gd-tcp</b> , and <b>hpc-tcp</b> with the correct bandwidth parameters and scheduler mapping.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Action</b>  | <p>List the traffic control profiles using the operational mode command <b>show class-of-service traffic-control-profile</b>:</p> <pre> user@switch&gt; show class-of-service traffic-control-profile Traffic control profile: be-tcp, Index: 40535   Shaping rate: 100 percent   Scheduler map: be-map   Guaranteed rate: 3500000000  Traffic control profile: gd-tcp, Index: 37959   Shaping rate: 100 percent   Scheduler map: gd-map   Guaranteed rate: 4500000000  Traffic control profile: hpc-tcp, Index: 47661   Shaping rate: 100 percent   Scheduler map: hpc-map   Guaranteed rate: 2000000000 </pre>                                                                                                                                                                                                                                                                                                                             |
| <b>Meaning</b> | <p>The <b>show class-of-service traffic-control-profile</b> command lists all of the configured traffic control profiles. For each traffic control profile, the command output includes:</p> <ul style="list-style-type: none"> <li>• The name of the traffic control profile (<b>traffic-control-profile</b>)</li> <li>• The maximum port bandwidth the priority group can consume (<b>shaping-rate</b>)</li> <li>• The scheduler map associated with the traffic control profile (<b>scheduler-map</b>)</li> <li>• The minimum guaranteed priority group port bandwidth (<b>guaranteed-rate</b>)</li> </ul> <p>The command output shows that:</p> <ul style="list-style-type: none"> <li>• The traffic control profile <b>be-tcp</b> can consume a maximum of <b>100 percent</b> of the port bandwidth, is associated with the scheduler map <b>be-map</b>, and has a minimum guaranteed bandwidth of <b>3,500,000,000 bps</b>.</li> </ul> |

- The traffic control profile **gd-tcp** can consume a maximum of **100 percent** of the port bandwidth, is associated with the scheduler map **gd-map**, and has a minimum guaranteed bandwidth of **4,500,000,000 bps**.
- The traffic control profile **hpc-tcp** can consume a maximum of **100 percent** of the port bandwidth, is associated with the scheduler map **hpc-map**, and has a minimum guaranteed bandwidth of **2,000,000,000 bps**.

### Verifying the Interface Configuration

---

**Purpose** Verify that the classifier, the congestion notification profile, and the forwarding class sets are configured on interfaces **xe-0/0/20** and **xe-0/0/21**.

**Action** List the interfaces using the operational mode commands **show configuration class-of-service interfaces xe-0/0/20** and **show configuration class-of-service interfaces xe-0/0/21**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/20
forwarding-class-set {
 best-effort-gp {
 output-traffic-control-profile be-tcp;
 }
 guar-delivery-pg {
 output-traffic-control-profile gd-tcp;
 }
 hpc-pg {
 output-traffic-control-profile hpc-tcp;
 }
}
congestion-notification-profile gd_cnp;
unit 0 {
 classifiers {
 ieee-802.1 hsclassifier1;
 }
}
```

```
user@switch> show configuration class-of-service interfaces xe-0/0/21
forwarding-class-set {
 best-effort-gp {
 output-traffic-control-profile be-tcp;
 }
 guar-delivery-pg {
 output-traffic-control-profile gd-tcp;
 }
 hpc-pg {
 output-traffic-control-profile hpc-tcp;
 }
}
congestion-notification-profile gd_cnp;
unit 0 {
 classifiers {
 ieee-802.1 hsclassifier1;
 }
}
```

|                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Meaning</b>               | The <b>show configuration class-of-service interfaces</b> <i>interface-name</i> command shows that each interface includes the forwarding class sets <b>best-effort-pg</b> , <b>guar-delivery-pg</b> , and <b>hpc-pg</b> , congestion notification profile <b>gd-cnp</b> , and the IEEE 802.1p classifier <b>hsclassifier1</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Defining CoS Unicast BA Classifiers (DSCP, DSCP IPv6, IEEE 802.1p)</a></li><li>• <a href="#">Benefits of Configuring CoS Hierarchical Port Scheduling</a></li><li>• <a href="#">Assigning CoS Components to Interfaces on page 6627</a></li><li>• <a href="#">Example: Configuring WRED Drop Profiles</a></li><li>• <a href="#">Example: Configuring Drop Profile Maps on page 6817</a></li><li>• <a href="#">Example: Configuring Forwarding Classes</a></li><li>• <a href="#">Example: Configuring Forwarding Class Sets on page 6697</a></li><li>• <a href="#">Example: Configuring Queue Schedulers</a></li><li>• <a href="#">Example: Configuring Queue Scheduling Priority on page 6753</a></li><li>• <a href="#">Example: Configuring Traffic Control Profiles (Priority Group Scheduling) on page 6762</a></li><li>• <a href="#">Example: Configuring Minimum Guaranteed Output Bandwidth on page 6798</a></li><li>• <a href="#">Example: Configuring Maximum Output Bandwidth on page 6805</a></li><li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 6905</a></li><li>• <a href="#">Overview of CoS Changes Introduced in Junos OS Release 12.2</a></li><li>• <a href="#">Understanding CoS Hierarchical Port Scheduling (ETS) on page 6765</a></li><li>• <a href="#">Understanding CoS Scheduling Behavior and Configuration Considerations on page 6718</a></li><li>• <a href="#">Understanding CoS Scheduling on QFabric System Node Device Fabric (fte) Ports</a></li><li>• <a href="#">Understanding Default CoS Scheduling on QFabric System Interconnect Devices (Junos OS Release 13.1 and Later Releases)</a></li></ul> |

---

## Understanding CoS Priority Group and Queue Guaranteed Minimum Bandwidth

You can set a guaranteed minimum bandwidth for individual forwarding classes (queues) and for groups of forwarding classes called *forwarding class sets* (priority groups). Setting a minimum guaranteed bandwidth ensures that priority groups and queues receive the bandwidth required to support the expected traffic.

This topic covers:

- [Guaranteeing Bandwidth Using Hierarchical Scheduling on page 6796](#)
- [Priority Group Guaranteed Rate \(Guaranteed Minimum Bandwidth\) on page 6797](#)
- [Queue Transmit Rate \(Guaranteed Minimum Bandwidth\) on page 6797](#)

## Guaranteeing Bandwidth Using Hierarchical Scheduling

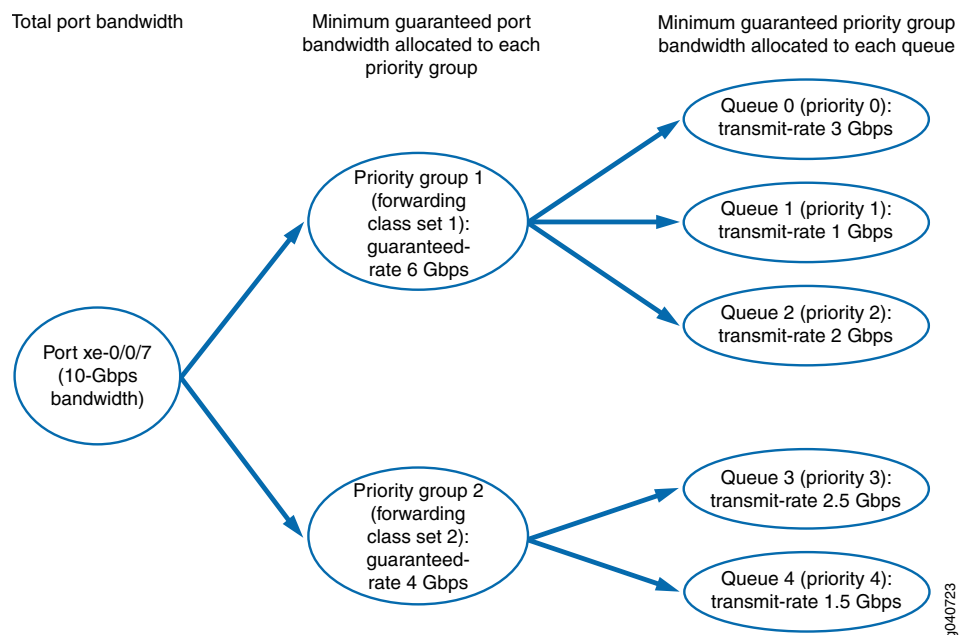
The **guaranteed-rate** value for the priority group (configured in a traffic control profile) defines the minimum amount of bandwidth allocated to a forwarding class set on a port, whereas the **transmit-rate** value of the queue (configured in a scheduler) defines the minimum amount of bandwidth allocated to a particular queue in a priority group. The queue bandwidth is a portion of the priority group bandwidth.



**NOTE:** You cannot configure a minimum guaranteed bandwidth (transmit rate) for a forwarding class that is mapped to a strict-high priority queue, and you cannot configure a minimum guaranteed bandwidth (guaranteed rate) for a priority group that includes strict-high priority queues.

Figure 222 shows how the total port bandwidth is allocated to priority groups (forwarding class sets) based on the guaranteed rate of each priority group. It also shows how the guaranteed bandwidth of each priority group is allocated to the queues in the priority group based on the transmit rate of each queue.

**Figure 222: Allocating Guaranteed Bandwidth Using Hierarchical Scheduling**



The sum of the priority group guaranteed rates cannot exceed the total port bandwidth. If you configure guaranteed rates whose sum exceeds the port bandwidth, the system sends a syslog message to notify you that the configuration is not valid. However, the system does not perform a commit check. If you commit a configuration in which the sum of the guaranteed rates exceeds the port bandwidth, the hierarchical scheduler behaves unpredictably.

The sum of the queue transmit rates cannot exceed the total guaranteed rate of the priority group to which the queues belong. If you configure transmit rates whose sum exceeds the priority group guaranteed rate, the commit check fails and the system rejects the configuration.



**NOTE:** You must set both the priority group **guaranteed-rate** value and the queue **transmit-rate** value in order to configure the minimum bandwidth for individual queues. If you set the **transmit-rate** value but do not set the **guaranteed-rate** value, the configuration fails.

You can set the **guaranteed-rate** value for a priority group without setting the **transmit-rate** value for individual queues in the priority group. However, queues that do not have a configured **transmit-rate** value can become starved for bandwidth if other higher-priority queues need the priority group's bandwidth. To avoid starving a queue, it is a good practice to configure a **transmit-rate** value for most queues.

If you configure the guaranteed rate of a priority group as a percentage, configure all of the transmit rates associated with that priority group as percentages. In this case, if any of the transmit rates are configured as absolute values instead of percentages, the configuration is not valid and the system sends a syslog message.

---

### Priority Group Guaranteed Rate (Guaranteed Minimum Bandwidth)

Setting a priority group (forwarding class set) **guaranteed-rate** enables you to reserve a portion of the port bandwidth for the forwarding classes (queues) in that forwarding class set. The minimum bandwidth (**guaranteed-rate**) that you configure for a priority group sets the minimum bandwidth available to all of the forwarding classes in the forwarding class set.

The combined **guaranteed-rate** value of all of the forwarding class sets associated with an interface cannot exceed the amount of bandwidth available on that interface.

You configure the priority group **guaranteed-rate** in the traffic control profile. You cannot apply a traffic control profile that has a guaranteed rate to a priority group that includes a strict-high priority queue.

### Queue Transmit Rate (Guaranteed Minimum Bandwidth)

Setting a queue (forwarding class) **transmit-rate** enables you to reserve a portion of the priority group bandwidth for the individual queue. For example, a queue that handles Fibre Channel over Ethernet (FCoE) traffic might require a minimum rate of 4 Gbps to ensure the class of service that storage area network (SAN) traffic requires.

The priority group **guaranteed-rate** sets the aggregate minimum amount of bandwidth available to the queues that belong to the priority group. The cumulative total minimum bandwidth the queues consume cannot exceed the minimum bandwidth allocated to

the priority group to which they belong. (The combined transmit rates of the queues in a priority group cannot exceed the priority group's guaranteed rate.)

You must configure the **guaranteed-rate** value of the priority group in order to set a **transmit-rate** value for individual queues that belong to the priority group. The reason is that if there is no guaranteed bandwidth for a priority group, there is no way to guarantee bandwidth for queues in that priority group.

You configure the queue **transmit-rate** in the scheduler configuration. You cannot configure a transmit rate for a strict-high priority queue.

#### Related Documentation

- *Understanding CoS Output Queue Schedulers*
- [Understanding CoS Traffic Control Profiles on page 6757](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 6771](#)
- *Example: Configuring Queue Schedulers*
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 6762](#)
- *Defining CoS Queue Schedulers*
- [Defining CoS Traffic Control Profiles \(Priority Group Scheduling\) on page 6761](#)

---

## Example: Configuring Minimum Guaranteed Output Bandwidth

Scheduling the minimum guaranteed output bandwidth for a queue (forwarding class) requires configuring both tiers of the two-tier hierarchical scheduler. One tier is scheduling the resources for the individual queue. The other tier is scheduling the resources for the priority group (forwarding class set) to which the queue belongs. You set a minimum guaranteed bandwidth to ensure that priority groups and queues receive the bandwidth required to support the expected traffic.

- [Requirements on page 6798](#)
- [Overview on page 6798](#)
- [Configuring Guaranteed Minimum Bandwidth on page 6800](#)
- [Verification on page 6801](#)

### Requirements

This example uses the following hardware and software components:

- A Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series or Junos OS Release 14.1X53-D20 or later for the OCX Series

### Overview

The priority group minimum guaranteed bandwidth defines the minimum total amount of bandwidth available for all of the queues in the priority group to meet their minimum bandwidth requirements.

The **transmit-rate** setting in the scheduler configuration determines the minimum guaranteed bandwidth for an individual queue. The transmit rate also determines the amount of excess (extra) priority group bandwidth that the queue can share. Extra priority group bandwidth is allocated among the queues in the priority group in proportion to the transmit rate of each queue.

The **guaranteed-rate** setting in the traffic control profile configuration determines the minimum guaranteed bandwidth for a priority group. The guaranteed rate also determines the amount of excess (extra) port bandwidth that the priority group can share. Extra port bandwidth is allocated among the priority groups on a port in proportion to the guaranteed rate of each priority group.



**NOTE:** You must configure both the **transmit-rate** value for the queue and the **guaranteed-rate** value for the priority group to set a valid minimum bandwidth guarantee for a queue. (If the priority group does not have a guaranteed minimum bandwidth, there is no guaranteed bandwidth pool from which the queue can take its guaranteed minimum bandwidth.)

The sum of the queue transmit rates in a priority group should not exceed the guaranteed rate for the priority group. (You cannot guarantee a minimum bandwidth for the queues that is greater than the minimum bandwidth guaranteed for the entire set of queues.)



**NOTE:** When you configure bandwidth for a queue or a priority group, the switch considers only the data as the configured bandwidth. The switch does not account for the bandwidth consumed by the preamble and the interframe gap (IFG). Therefore, when you calculate and configure the bandwidth requirements for a queue or for a priority group, consider the preamble and the IFG as well as the data in the calculations.



**NOTE:** You cannot configure minimum guaranteed bandwidth on strict-high priority queues or on a priority group that contains strict-high priority queues.

This example describes how to:

- Configure a transmit rate (minimum guaranteed queue bandwidth) of 2 Gbps for queues in a scheduler named **be-sched**.
- Configure a guaranteed rate (minimum guaranteed priority group bandwidth) of 4 Gbps for a priority group in a traffic control profile named **be-tcp**.
- Assign the scheduler to a queue named **best-effort** by using a scheduler map named **be-map**.
- Associate the scheduler map **be-map** with the traffic control profile **be-tcp**.

- Assign the queue **best-effort** to a priority group named **be-pg**.
- Assign the priority group and the minimum guaranteed bandwidth scheduling to the egress interface **xe-0/0/7**.

Table 576 shows the configuration components for this example:

**Table 576: Components of the Minimum Guaranteed Output Bandwidth Configuration Example**

| Component                                   | Settings                   |
|---------------------------------------------|----------------------------|
| Hardware                                    | QFX3500 switch             |
| Minimum guaranteed queue bandwidth          | Transmit rate: <b>2g</b>   |
| Minimum guaranteed priority group bandwidth | Guaranteed rate: <b>4g</b> |
| Scheduler                                   | <b>be-sched</b>            |
| Scheduler map                               | <b>be-map</b>              |
| Traffic control profile                     | <b>be-tcp</b>              |
| Forwarding class set (priority group)       | <b>be-pg</b>               |
| Queue (forwarding class)                    | <b>best-effort</b>         |
| Egress interface                            | <b>xe-0/0/7</b>            |

## Configuring Guaranteed Minimum Bandwidth

**CLI Quick Configuration** To quickly configure the minimum guaranteed bandwidth for a priority group and a queue, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level:

```
[edit class-of-service]
set schedulers be-sched transmit-rate percent 2g
set traffic-control-profiles be-tcp guaranteed-rate 4g
set scheduler-maps be-map forwarding-class best-effort scheduler be-sched
set traffic-control-profiles be-tcp scheduler-map be-map
set forwarding-class-sets be-pg class best-effort
set interfaces xe-0/0/7 forwarding-class-set be-pg output-traffic-control-profile be-tcp
```

To configure the minimum guaranteed bandwidth hierarchical scheduling for a queue and a priority group:

1. Configure the minimum guaranteed queue bandwidth of 2 Gbps for scheduler **be-sched**:

```
[edit class-of-service schedulers]
user@switch# set be-sched transmit-rate 2g
```

2. Configure the minimum guaranteed priority group bandwidth of 4 Gbps for traffic control profile **be-tcp**:

```
[edit class-of-service traffic-control-profiles]
```



- ```

user@switch# set be-tcp guaranteed-rate 4g

```
3. Associate the scheduler **be-sched** with the **best-effort** queue in the scheduler map **be-map**:


```

[edit class-of-service scheduler-maps]
user@switch# set be-map forwarding-class best-effort scheduler be-sched

```
 4. Associate the scheduler map with the traffic control profile:


```

[edit class-of-service traffic-control-profiles]
user@switch# set be-tcp scheduler-map be-map

```
 5. Assign the **best-effort** queue to the priority group **be-pg**:


```

[edit class-of-service forwarding-class-sets]
user@switch# set be-pg class best-effort

```
 6. Apply the configuration to interface **xe-0/0/7**:


```

[edit class-of-service interfaces]
user@switch# set xe-0/0/7 forwarding-class-set be-pg output-traffic-control-profile be-tcp

```

Verification

To verify the minimum guaranteed output bandwidth configuration, perform these tasks:

- [Verifying the Minimum Guaranteed Queue Bandwidth on page 6801](#)
- [Verifying the Priority Group Minimum Guaranteed Bandwidth and Scheduler Map Association on page 6801](#)
- [Verifying the Scheduler Map Configuration on page 6802](#)
- [Verifying Queue \(Forwarding Class\) Membership in the Priority Group on page 6802](#)
- [Verifying the Egress Interface Configuration on page 6802](#)

Verifying the Minimum Guaranteed Queue Bandwidth

- Purpose** Verify that you configured the minimum guaranteed queue bandwidth as **2g** in the scheduler **be-sched**.
- Action** Display the minimum guaranteed bandwidth in the **be-sched** scheduler configuration using the operational mode command **show configuration class-of-service schedulers be-sched transmit-rate**:

```

user@switch> show configuration class-of-service schedulers be-sched transmit-rate
2g;

```

Verifying the Priority Group Minimum Guaranteed Bandwidth and Scheduler Map Association

- Purpose** Verify that the minimum guaranteed priority group bandwidth is **4g** and the attached scheduler map is **be-map** in the traffic control profile **be-tcp**.
- Action** Display the minimum guaranteed bandwidth in the **be-tcp** traffic control profile configuration using the operational mode command **show configuration class-of-service traffic-control-profiles be-tcp guaranteed-rate**:

```

user@switch> show configuration class-of-service traffic-control-profiles be-tcp guaranteed-rate
4g;

```

Display the scheduler map in the **be-tcp** traffic control profile configuration using the operational mode command **show configuration class-of-service traffic-control-profiles be-tcp scheduler-map**:

```
user@switch> show configuration class-of-service traffic-control-profiles be-tcp scheduler-map scheduler-map be-map;
```

Verifying the Scheduler Map Configuration

Purpose Verify that the scheduler map **be-map** maps the forwarding class **best-effort** to the scheduler **be-sched**.

Action Display the **be-map** scheduler map configuration using the operational mode command **show configuration class-of-service schedulers maps be-map**:

```
user@switch> show configuration class-of-service scheduler-maps be-map forwarding-class best-effort scheduler be-sched;
```

Verifying Queue (Forwarding Class) Membership in the Priority Group

Purpose Verify that the forwarding class set **be-pg** includes the forwarding class **best-effort**.

Action Display the **be-pg** forwarding class set configuration using the operational mode command **show configuration class-of-service forwarding-class-sets be-pg**:

```
user@switch> show configuration class-of-service forwarding-class-sets be-pg class best-effort;
```

Verifying the Egress Interface Configuration

Purpose Verify that the forwarding class set **be-pg** and the traffic control profile **be-tcp** are attached to egress interface **xe-0/0/7**.

Action Display the egress interface using the operational mode command **show configuration class-of-service interfaces xe-0/0/7**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/7 forwarding-class-set {
    be-pg {
        output-traffic-control-profile be-tcp;
    }
}
```

Related Documentation

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 6771](#)
- [Example: Configuring Queue Schedulers](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 6762](#)
- [Example: Configuring Queue Scheduling Priority on page 6753](#)
- [Example: Configuring Forwarding Class Sets on page 6697](#)
- [Understanding CoS Traffic Control Profiles on page 6757](#)

- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 6765](#)

[Understanding CoS Priority Group Shaping and Queue Shaping \(Maximum Bandwidth\)](#)

If the amount of traffic on an interface exceeds the maximum bandwidth available on the interface, it leads to congestion. You can use priority group (forwarding class set) shaping and queue (forwarding class) shaping to manage traffic and avoid congestion.

Configuring a maximum bandwidth sets the most bandwidth a priority group or a queue can use after all of the priority group and queue minimum bandwidth requirements are met, even if more bandwidth is available.

This topic covers:

- [Priority Group Shaping on page 6803](#)
- [Queue Shaping on page 6803](#)
- [Shaping Maximum Bandwidth Using Hierarchical Scheduling on page 6804](#)

Priority Group Shaping

Priority group shaping enables you to shape the aggregate traffic of a forwarding class set on a port to a maximum rate that is less than the line or port rate. The maximum bandwidth ([shaping-rate](#)) that you configure for a priority group sets the maximum bandwidth available to all of the forwarding classes (queues) in the forwarding class set.

If a port has more than one priority group and the combined **shaping-rate** value of the priority groups is greater than the amount of port bandwidth available, the bandwidth is shared proportionally among the priority groups.

You configure the priority group **shaping-rate** in the traffic control profile.

Queue Shaping

Queue shaping throttles the rate at which queues transmit packets. For example, using queue shaping, you can rate-limit a strict-high priority queue so that the strict-priority queue does not lock out (or starve) low-priority queues.



NOTE: We recommend that you always apply a shaping rate to strict-high priority queues to prevent them from starving other queues. If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.

Similarly, for any queue, you can configure queue shaping (**shaping-rate**) to set the maximum bandwidth for a particular queue.

The **shaping-rate** value of the priority group sets the aggregate maximum amount of bandwidth available to the queues that belong to the priority group. On a port, the

cumulative total bandwidth the queues consume cannot exceed the maximum bandwidth of the priority group to which they belong.

If a priority group has more than one queue, and the combined **shaping-rate** of the queues is greater than the amount of bandwidth available to the priority group, the bandwidth is shared proportionally among the queues.

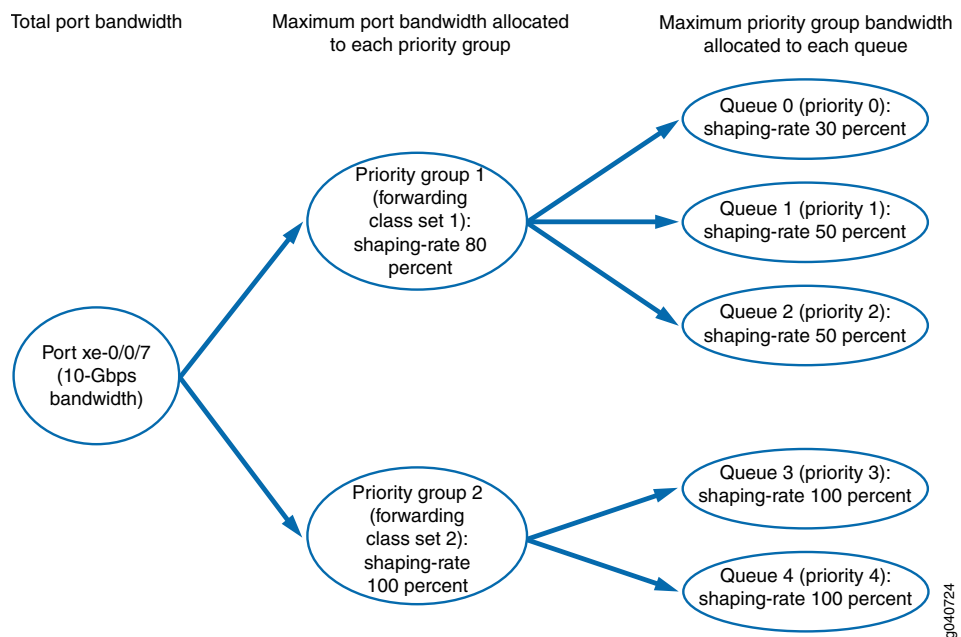
You configure the queue **shaping-rate** in the scheduler configuration, and you set the **shaping-rate** for priority groups in the traffic control profile configuration.

Shaping Maximum Bandwidth Using Hierarchical Scheduling

Priority group shaping defines the maximum bandwidth allocated to a forwarding class set on a port, whereas queue shaping defines a limit on maximum bandwidth usage per queue. The queue bandwidth is a portion of the priority group bandwidth.

Figure 223 shows how the port bandwidth is allocated to priority groups (forwarding class sets) based on the shaping rate of each priority group, and how the bandwidth of each priority group is allocated to the queues in the priority group based on the shaping rate of each queue.

Figure 223: Setting Maximum Bandwidth Using Hierarchical Scheduling



Related Documentation

- [Understanding CoS Output Queue Schedulers](#)
- [Understanding CoS Traffic Control Profiles on page 6757](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 6771](#)
- [Example: Configuring Queue Schedulers](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 6762](#)
- [Defining CoS Queue Schedulers](#)

- [Defining CoS Traffic Control Profiles \(Priority Group Scheduling\) on page 6761](#)

Example: Configuring Maximum Output Bandwidth

Scheduling the maximum output bandwidth for a queue (forwarding class) requires configuring both tiers of the hierarchical scheduler. One tier is scheduling the resources for the individual queue. The other tier is scheduling the resources for the priority group (forwarding class set) to which the queue belongs. You can use priority group and queue shaping to prevent traffic from using more bandwidth than you want the traffic to receive.

- [Requirements on page 6805](#)
- [Overview on page 6805](#)
- [Configuring Maximum Bandwidth on page 6806](#)
- [Verification on page 6807](#)

Requirements

This example uses the following hardware and software components:

- A Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

Overview

The priority group maximum bandwidth defines the maximum total amount of bandwidth available for all of the queues in the priority group.

The **shaping-rate** setting in the scheduler configuration determines the maximum bandwidth for an individual queue.

The **shaping-rate** setting in the traffic control profile configuration determines the maximum bandwidth for a priority group.



NOTE: When you configure bandwidth for a queue or a priority group, the switch considers only the data as the configured bandwidth. The switch does not account for the bandwidth consumed by the preamble and the interframe gap (IFG). Therefore, when you calculate and configure the bandwidth requirements for a queue or for a priority group, consider the preamble and the IFG as well as the data in the calculations.



NOTE: When you set the maximum bandwidth (**shaping-rate**) for a queue or for a priority group at 100 Kbps or less, the traffic shaping behavior is accurate only within +/- 20 percent of the configured **shaping-rate** value.

This example describes how to:

- Configure a maximum rate of 4 Gbps for queues in a scheduler named **be-sched**.
- Configure a maximum rate of 6 Gbps for a priority group in a traffic control profile named **be-tcp**.
- Assign the scheduler to a queue named **best-effort** by using a scheduler map named **be-map**.
- Associate the scheduler map **be-map** with the traffic control profile **be-tcp**.
- Assign the queue **best-effort** to a priority group named **be-pg**.
- Assign the priority group and the bandwidth scheduling to the interface **xe-0/0/7**.

Table 577 shows the configuration components for this example:

Table 577: Components of the Maximum Output Bandwidth Configuration Example

Component	Settings
Hardware	QFX3500 switch
Maximum queue bandwidth	Shaping rate: 4g
Maximum priority group bandwidth	Shaping rate: 6g
Scheduler	be-sched
Scheduler map	be-map
Traffic control profile	be-tcp
Forwarding class set (priority group)	be-pg
Queue (forwarding class)	best-effort
Egress interface	xe-0/0/7

Configuring Maximum Bandwidth

CLI Quick Configuration To quickly configure the maximum bandwidth for a priority group and a queue, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level:

```
[edit class-of-service]
set schedulers be-sched shaping-rate percent 4g
set traffic-control-profiles be-tcp shaping-rate 6g
set scheduler-maps be-map forwarding-class best-effort scheduler be-sched
set traffic-control-profiles be-tcp scheduler-map be-map
set forwarding-class-sets be-pg class best-effort
set interfaces xe-0/0/7 forwarding-class-set be-pg output-traffic-control-profile be-tcp
```

To configure the maximum bandwidth hierarchical scheduling for a queue and a priority group:

1. Configure the maximum queue bandwidth of 4 Gbps for scheduler **be-sched**:

```
[edit class-of-service schedulers]
user@switch# set be-sched shaping-rate 4g
```

2. Configure the maximum priority group bandwidth of 6 Gbps for traffic control profile **be-tcp**:

```
[edit class-of-service traffic-control-profiles]
user@switch# set be-tcp shaping-rate 6g
```

3. Associate the scheduler **be-sched** with the **best-effort** queue in the scheduler map **be-map**:

```
[edit class-of-service scheduler-maps]
user@switch# set be-map forwarding-class best-effort scheduler be-sched
```

4. Associate the scheduler map with the traffic control profile:

```
[edit class-of-service traffic-control-profiles]
user@switch# set be-tcp scheduler-map be-map
```

5. Assign the **best-effort** queue to the priority group **be-pg**:

```
[edit class-of-service forwarding-class-sets]
user@switch# set be-pg class best-effort
```

6. Apply the configuration to interface **xe-0/0/7**:

```
[edit class-of-service interfaces]
user@switch# set xe-0/0/7 forwarding-class-set be-pg output-traffic-control-profile be-tcp
```

Verification

To verify the maximum output bandwidth configuration, perform these tasks:

- [Verifying the Maximum Queue Bandwidth on page 6807](#)
- [Verifying the Priority Group Maximum Bandwidth and Scheduler Map Association on page 6807](#)
- [Verifying the Scheduler Map Configuration on page 6808](#)
- [Verifying Queue \(Forwarding Class\) Membership in the Priority Group on page 6808](#)
- [Verifying the Egress Interface Configuration on page 6808](#)

Verifying the Maximum Queue Bandwidth

Purpose Verify that you configured the maximum queue bandwidth as **4g** in the scheduler **be-sched**.

Action List the maximum bandwidth in the **be-sched** scheduler configuration using the operational mode command **show configuration class-of-service schedulers be-sched shaping-rate**:

```
user@switch> show configuration class-of-service schedulers be-sched shaping-rate
4g;
```

Verifying the Priority Group Maximum Bandwidth and Scheduler Map Association

Purpose Verify that the maximum priority group bandwidth is **6g** and the attached scheduler map is **be-map** in the traffic control profile **be-tcp**.

Action List the maximum bandwidth in the **be-tcp** traffic control profile configuration using the operational mode command **show configuration class-of-service traffic-control-profiles be-tcp shaping-rate**:

```
user@switch> show configuration class-of-service traffic-control-profiles be-tcp shaping-rate
4g;
```

List the scheduler map in the **be-tcp** traffic control profile configuration using the operational mode command **show configuration class-of-service traffic-control-profiles be-tcp scheduler-map**:

```
user@switch> show configuration class-of-service traffic-control-profiles be-tcp scheduler-map
scheduler-map be-map;
```

Verifying the Scheduler Map Configuration

Purpose Verify that the scheduler map **be-map** maps the forwarding class **best-effort** to the scheduler **be-sched**.

Action List the **be-map** scheduler map configuration using the operational mode command **show configuration class-of-service schedulers maps be-map**:

```
user@switch> show configuration class-of-service scheduler-maps be-map
forwarding-class best-effort scheduler be-sched;
```

Verifying Queue (Forwarding Class) Membership in the Priority Group

Purpose Verify that the forwarding class set **be-pg** includes the forwarding class **best-effort**.

Action List the **be-pg** forwarding class set configuration using the operational mode command **show configuration class-of-service forwarding-class-sets be-pg**:

```
user@switch> show configuration class-of-service forwarding-class-sets be-pg
class best-effort;
```

Verifying the Egress Interface Configuration

Purpose Verify that the forwarding class set **be-pg** and the traffic control profile **be-tcp** are attached to egress interface **xe-0/0/7**.

Action List the egress interface using the operational mode command **show configuration class-of-service interfaces xe-0/0/7**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/7
forwarding-class-set {
    be-pg {
        output-traffic-control-profile be-tcp;
    }
}
```

Related Documentation

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 6771](#)
- [Example: Configuring Queue Schedulers](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 6762](#)

- [Example: Configuring Forwarding Class Sets on page 6697](#)
- [Understanding CoS Traffic Control Profiles on page 6757](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 6765](#)

Understanding CoS WRED Drop Profiles

When the number of packets queued is greater than the ability of the switch to empty an output queue, the queue requires a method for determining which packets to drop to relieve the congestion. Weighted random early detection (WRED) drop profiles define the drop probability of packets of different packet loss probabilities (PLPs) as the output queue fills. During periods of congestion, as the output queue fills, the switch drops incoming packets as determined by a drop profile, until the output queue becomes less congested.



NOTE: Do not apply drop profiles to lossless traffic (traffic that belongs to a forwarding class that has the no-loss drop attribute.). Lossless traffic uses priority-based flow control (PFC) to control congestion.

Depending on the drop probabilities, a drop profile can drop many packets long before the buffer becomes full, or it can drop only a few packets even if the buffer is almost full.

You configure drop profiles in the drop profile section of the class-of-service (CoS) configuration hierarchy. You apply drop profiles using a drop profile map in queue scheduler configuration. For each queue scheduler, you can configure separate drop profiles for each PLP (low, medium-high, and high). This enables you to treat traffic of different PLPs in different ways during periods of congestion.

Drop profiles define the meaning (packet drop action) of each PLP by setting different values for when to drop packets and the probability that packets will drop for each PLP.

You can configure a maximum of 32 drop profiles.

- [Drop Profile Parameters on page 6809](#)
- [Default Drop Profile on page 6810](#)
- [Packet Drop Method on page 6811](#)
- [Drop Profile Maps on page 6811](#)
- [Congestion Prevention on page 6811](#)
- [Configuring a WRED Drop Profile and Applying it to an Output Queue on page 6811](#)
- [Drop Profiles on Explicit Congestion Notification Enabled Queues on page 6812](#)

Drop Profile Parameters

Drop profiles specify two values, which work as pairs:

- **Fill level**—The queue fullness value, which represents a percentage of the memory used to store packets in relation to the total amount of memory allocated to the queue.

- Drop probability—The percentage value that corresponds to the likelihood that an individual packet is dropped.

Each queue fill level pairs with a drop probability. As the queue fills to different levels, every time it reaches a fill level configured in a drop profile, the queue applies the drop probability paired with that fill level to the traffic in the queue that exceeds the fill level. You can configure up to 32 pairs of fill levels and drop probabilities to create a customized packet drop probability curve with up to 32 points of differentiation.

Packets are not dropped until they reach the first configured queue fill level. When the queue reaches the first fill level, packets begin to drop at the configured drop probability rate paired with the first fill level. When the queue reaches the second fill level, packets begin to drop at the configured drop probability rate paired with the second fill level. This process continues for the number of fill level/drop probability pairs that you configure in the drop profile.

Drop profiles are interpolated, not segmented. An interpolated drop profile gradually increases the drop probability along a curve between each configured fill level. When the queue reaches the next fill level, the drop probability reaches the drop probability paired with that fill level. A segmented drop profile “jumps” from one fill level/drop probability setting to another in a stepped fashion. The drop probability of traffic does not change as the queue fills until the next fill level is reached.

An example of interpolation is a drop profile with three fill level/drop probability pairs:

- 25 percent queue fill level paired with a 30 percent drop probability
- 50 percent queue fill level paired with a 60 percent drop probability
- 75 percent queue fill level paired with a 100 percent drop probability (all packets that exceed the 75 percent queue fill level are dropped)

The queue drops no packets until its fill level reaches 25 percent. During periods of congestion, when the queue fills above 25 percent full, the queue begins to drop packets at a rate of 30 percent of the packets above the fill level.

However, as the queue continues to fill, it does not continue to drop packets at the 30 percent drop probability. Instead, the drop probability gradually increases as the queue fills to the 50 percent fullness level. When the queue reaches the 50 percent fill level, the drop probability has increased to the configured drop probability pair for the fill level, which is 60 percent.

As the queue continues to fill, the drop probability does not remain at 60 percent, but continues to rise as the queue fills. When the queue reaches the final fill level at 75 percent full, the drop probability has risen to 100 percent and all packets that exceed the 75 percent fill level are dropped.

Default Drop Profile

If you do not configure default profiles and apply them to queue schedulers, the switch uses the default drop profile for lossy traffic classes. In the default drop profile, when the fill level is 0 percent, the drop probability is 0 percent. When the fill level is 100 percent,

the drop probability is 100 percent. During periods of congestion, as soon as packets arrive on a queue, the default profile might begin to drop packets.

Packet Drop Method

When a packet reaches the head of a queue, the switch calculates a random number between 0 and 100. The switch plots the random number against the drop profile using the current fill level of the queue. When the random number falls above the graph line, the queue transmits the packet out the egress interface. When the number falls below graph the line, the switch drops the packet.

Drop Profile Maps

Drop profile maps are part of scheduler configuration. A drop profile map maps drop profiles to packet loss priorities. Specifying the drop profile map in a scheduler associates the drop profile with the forwarding classes (queues) that you map to the scheduler in a scheduler map.

You configure loss priority for a queue in the classifier section of the CoS configuration hierarchy, and the loss priority is applied to the traffic assigned to the forwarding class at the ingress interface.

Each scheduler can have multiple drop profile maps.

Congestion Prevention

Configuring drop profiles on output queues enables you to control how congestion affects other queues on a port. If you do not configure drop profiles and map them to output queues, the switch uses the default drop profile on queues that forward lossy traffic.



NOTE: Do not configure drop profiles for the `fcoe` and `no-loss` forwarding classes. FCoE and other lossless traffic queues (traffic queues that are configured with the `no-loss` packet drop attribute) require lossless behavior. Use priority-based flow control (PFC) to prevent frame drop on lossless priorities.

Configuring a WRED Drop Profile and Applying it to an Output Queue

To configure a WRED packet drop profile and apply it to an output queue:

1. Configure a drop profile using the statement **`set class-of-service drop-profiles profile-name interpolate fill-level level1 level2 ... level32 drop-probability probability1 probability2 ... probability32`**. You can specify as few as two fill level/drop probability pairs or as many as 32 pairs.
2. Map the drop profile to a queue scheduler using the statement **`set class-of-service schedulers scheduler-name drop-profile-map loss-priority (low | medium-high | high) protocol any drop-profile profile-name`**. The name of the drop-profile is the name of the WRED profile configured in Step 1.

3. Map the scheduler, which Step 2 associates with the drop profile, to the output queue using the statement **set class-of-service scheduler-maps *map-name* forwarding-class *forwarding-class-name* scheduler *scheduler-name***. The forwarding class identifies the output queue. Forwarding classes are mapped to output queues by default, and can be remapped to different queues by explicit user configuration. The scheduler name is the scheduler configured in Step 2.
4. Associate the scheduler map with an interface using the statement **set class-of-service interfaces *interface-name* scheduler-map *scheduler-map-name***.

The interface uses the scheduler map to apply the drop profile (and other attributes) to the output queue mapped to the forwarding class on that interface. Because you can use different scheduler maps on different interfaces, the same queue number on different interfaces can handle traffic in different ways.

Drop Profiles on Explicit Congestion Notification Enabled Queues

You must configure a WRED drop profile on queues that you enable for explicit congestion notification (ECN). On ECN-enabled queues, the drop profile sets the threshold for when the queue should mark a packet as experiencing congestion (see [“Understanding CoS Explicit Congestion Notification” on page 6820](#)). When a queue fills to the level at which the WRED drop profile has a packet drop probability greater than zero (0), the switch might mark a packet as experiencing congestion. Whether or not a switch marks a packet as experiencing congestion is the same probability as the drop probability of the queue at that fill level.

On ECN-enabled queues, the switch does not use the drop profile to control dropping packets that are not ECN-capable packets (packets marked non-ECT, ECN code bits 00) during periods of congestion. Instead, the switch uses the tail-drop algorithm to drop non-ECT-capable packets during periods of congestion. When a queue fills to its maximum level of fullness, tail-drop simply drops all subsequently arriving packets until there is space in the queue to buffer more packets. All non-ECT-capable packets are treated the same way.

To apply a WRED drop profile to non-ECT traffic, configure a multifield (MF) classifier to assign non-ECT traffic to a different output queue that is not ECN-enabled, and then apply the WRED drop profile to that queue.

Related Documentation

- [Understanding Junos CoS Components on page 6621](#)
- [Understanding CoS Explicit Congestion Notification on page 6820](#)
- [Example: Configuring WRED Drop Profiles on page 6814](#)
- [Example: Configuring Drop Profile Maps on page 6817](#)
- [Example: Configuring Classifiers on page 6657](#)

Configuring CoS WRED Drop Profiles

You can configure an interpolated weighted random early detection (WRED) profile to control traffic congestion by controlling packet drop characteristics for different packet loss priorities.



NOTE: Do not enable WRED on lossless traffic flows (forwarding classes configured with the no-loss packet drop attribute). Use priority-based flow control (PFC) to prevent packet loss on lossless forwarding classes.

Drop profiles specify two values, which work as pairs:

- Fill level—The queue fullness value, which represents a percentage of the memory used to store packets in relation to the total amount of memory allocated to the queue.
- Drop probability—The percentage value that corresponds to the likelihood that an individual packet is dropped.

Each queue fill level pairs with a drop probability. As the queue fills to different levels, every time it reaches a fill level configured in a drop profile, the queue applies the drop probability paired with that fill level to the traffic in the queue that exceeds the fill level. You can configure up to 32 pairs of fill levels and drop probabilities to create a customized packet drop probability curve with up to 32 points of differentiation.

Packets are not dropped until they reach the first configured queue fill level. When the queue reaches the first fill level, packets begin to drop at the configured drop probability rate paired with the first fill level. When the queue reaches the second fill level, packets begin to drop at the configured drop probability rate paired with the second fill level. This process continues for the number of fill level/drop probability pairs that you configure in the drop profile.

Drop profiles are *interpolated*. An interpolated drop profile gradually increases the drop probability along a curve between each configured fill level. When the queue reaches the next fill level, the drop probability reaches the drop probability paired with that fill level.



NOTE: On ECN-enabled queues, the drop profile sets the threshold for when the queue should mark a packet as experiencing congestion (see [“Understanding CoS Explicit Congestion Notification” on page 6820](#)). On ECN-enabled queues, the switch does not use the drop profile to control dropping packets that are not ECN-capable packets during periods of congestion. Instead, the switch uses the tail-drop algorithm to drop non-ECN-capable packets during periods of congestion. When a queue fills to its maximum level of fullness, tail-drop simply drops all subsequently arriving packets until there is space in the queue to buffer more packets. All non-ECN-capable packets are treated the same way.

To configure a WRED profile using the CLI:

Name the drop profile and set the fill levels and their associated drop probabilities as percentages. For every fill level, there must be a paired drop probability (you must configure the same number of fill levels and drop probabilities).

```
[edit class-of-service]
user@switch# set drop-profile drop-profile-name interpolate fill-level level1 level2 ... level32
drop-probability probability1 probability2 ... probability32
```

**Related
Documentation**

- [Example: Configuring WRED Drop Profiles on page 6814](#)
- [Defining CoS Queue Schedulers for Port Scheduling on page 6742](#)
- [Configuring CoS Drop Profile Maps on page 6817](#)
- [Understanding CoS WRED Drop Profiles on page 6809](#)

Example: Configuring WRED Drop Profiles

You can configure interpolated weighted random early detection (WRED) profiles to control congestion on best-effort traffic by controlling packet drop characteristics for different packet loss priorities.



NOTE: Do not enable WRED on lossless traffic flows. Use priority-based flow control (PFC) to prevent packet loss on lossless forwarding classes.

-
- [Requirements on page 6814](#)
 - [Overview on page 6814](#)
 - [Configuration on page 6816](#)
 - [Verification on page 6816](#)

Requirements

This example uses the following hardware and software components:

- One switch.
- Junos OS Release 15.1X53-D10 or later for the QFX Series.

Overview

You associate WRED drop profiles with loss priorities in a scheduler. When you map the scheduler to a forwarding class (queue), you apply the interpolated drop profile to traffic of the specified loss priority on that queue. Drop profiles specify two values, which work as pairs:

- Fill level—The queue fullness value, which represents a percentage of the memory used to store packets in relation to the total amount of memory allocated to the queue.

- Drop probability—The percentage value that corresponds to the likelihood that an individual packet is dropped.

Each queue fill level pairs with a drop probability. As the queue fills to different levels, every time it reaches a fill level configured in a drop profile, the queue applies the drop probability paired with that fill level to the traffic in the queue that exceeds the fill level. You can configure up to 32 pairs of fill levels and drop probabilities to create a customized packet drop probability curve with up to 32 points of differentiation.

Packets are not dropped until they reach the first configured queue fill level. When the queue reaches the first fill level, packets begin to drop at the configured drop probability rate paired with the first fill level. When the queue reaches the second fill level, packets begin to drop at the configured drop probability rate paired with the second fill level. This process continues for the number of fill level/drop probability pairs that you configure in the drop profile.

Drop profiles are *interpolated*. An interpolated drop profile gradually increases the drop probability along a curve between each configured fill level. When the queue reaches the next fill level, the drop probability reaches the drop probability paired with that fill level.



NOTE: On ECN-enabled queues, the drop profile sets the threshold for when the queue should mark a packet as experiencing congestion (see [“Understanding CoS Explicit Congestion Notification” on page 6820](#)). On ECN-enabled queues, the switch does not use the drop profile to control dropping packets that are not ECN-capable packets during periods of congestion. Instead, the switch uses the tail-drop algorithm to drop non-ECN-capable packets during periods of congestion. When a queue fills to its maximum level of fullness, tail-drop simply drops all subsequently arriving packets until there is space in the queue to buffer more packets. All non-ECN-capable packets are treated the same way.

This example describes how to configure a drop profile with three fill level/drop probability pairs:

- Drop profile name—**be-dp1**
- Queue fill levels—25 percent, 50 percent, 75 percent
- Drop probabilities—30 percent, 60 percent, 100 percent

Each of the three fill levels pairs with a drop probability to program the interpolated drop profile curve.

You apply a drop profile by configuring a drop profile map that maps the drop profile to a packet loss priority, and associate the drop profile and packet loss priority with a scheduler. When you map the scheduler to a forwarding class (queue), the switch applies the drop profile to the packets in the forwarding class that have a matching packet loss priority.

Configuration

To configure a drop profile:

1. Set the drop start point at a **25** percent fill level, an intermediate fill level of **50** percent, and a drop end point of **75** percent. Set the paired drop probabilities to **30** percent, **60** percent, and **100** percent, respectively, for drop profile **be-dp1**:

[edit **class-of-service**]

```
user@switch# set drop-profile be-dp1 interpolate fill-level [ 25 50 75 ] drop-probability [ 30 60 100 ]
```

Verification

Verifying the Drop Profile Configuration

Purpose Verify that you configured the drop profile **be-dp1** with the correct fill levels and drop probabilities.

Action Verify the results of the drop profile configuration using the operational mode command **show configuration class-of-service drop-profiles be-dp1**:

```
user@switch> show configuration class-of-service drop-profiles be-dp1
interpolate {
    fill-level [ 25 50 75 ];
    drop-probability [ 30 60 100 ];
}
```

Related Documentation

- [Example: Configuring Queue Schedulers for Port Scheduling on page 6744](#)
- [Example: Configuring Drop Profile Maps on page 6817](#)
- [Configuring CoS WRED Drop Profiles on page 6813](#)
- [Understanding CoS WRED Drop Profiles on page 6809](#)

Configuring CoS Drop Profile Maps

A drop-profile map associates weighted random early detection (WRED) profiles for traffic of specified packet loss priorities with a scheduler. When you use a scheduler map to map a scheduler to a forwarding class, the drop profile map associated with the scheduler applies the specified WRED drop profile to traffic in the forwarding class that matches the specified packet loss priority.

Drop profile maps enable you to configure different drop profiles for traffic of different packet loss priorities within the same scheduler. You can associate different drop profiles with low-priority, medium-high priority, and high-priority traffic within a single scheduler, and then map that scheduler to a forwarding class. This applies the appropriate drop profile to traffic of each loss priority in a forwarding class. Drop profile maps apply to all traffic protocols.

To configure a drop-profile map:

- For the desired scheduler, configure the traffic loss priority and specify the drop profile you want to use to control the drop characteristics for traffic of that loss priority:

```
[edit class-of-service]
user@switch# set schedulers scheduler-name drop-profile-map loss-priority level protocol
any drop-profile drop-profile-name
```



NOTE: QFX10000 switches do not support the protocol any portion of the configuration. Drop profiles apply to all protocols.

Example: Configuring Drop Profile Maps

A drop-profile map associates a tail-drop profile for traffic of a specified loss priority with a scheduler. When you use a scheduler map to map a scheduler to a forwarding class, the drop profile map associated with the scheduler applies the specified tail-drop profile to traffic in the forwarding class that matches the specified loss priority.

- [Requirements on page 6817](#)
- [Overview on page 6818](#)
- [Configuring a Drop Profile Map on page 6818](#)
- [Verification on page 6818](#)

Requirements

This example uses the following hardware and software components:

- A Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

Overview

Drop profile maps enable you to configure different drop profiles for traffic of different loss priorities within the same scheduler. You can associate different drop profiles with low-priority, medium-high priority, and high-priority traffic within a single scheduler, and then map that scheduler to a forwarding class. This applies the appropriate drop profile to traffic of each loss priority in a forwarding class. Drop profile maps apply to all traffic protocols.

The following example describes how to configure a drop profile map for a scheduler named **mylan** that includes:

- A drop profile called **lp-profile** for low-priority traffic
- A drop profile called **mh-profile** for medium-high priority traffic
- A drop profile called **h-profile** for high-priority traffic

You apply the drop profiles in the drop profile map to a forwarding class by associating the scheduler **mylan** with a forwarding class in a scheduler map.

Configuring a Drop Profile Map

To configure a drop profile map:

1. Configure the drop profile for low-priority traffic:

```
[edit class-of-service]
user@switch# set schedulers mylan drop-profile-map loss-priority low protocol any
drop-profile lp-profile
```

2. Configure the drop profile for medium-high priority traffic:

```
[edit class-of-service]
user@switch# set schedulers mylan drop-profile-map loss-priority medium-high protocol
any drop-profile mh-profile
```

3. Configure the drop profile for high-priority traffic:

```
[edit class-of-service]
user@switch# set schedulers mylan drop-profile-map loss-priority high protocol any
drop-profile h-profile
```

Verification

Verifying the Drop Profile Map Configuration

Purpose Verify that you configured the drop profile map for the scheduler **mylan** with the correct loss priorities and drop profiles.

Action Verify the results of the drop profile map configuration using the operational mode command **show configuration class-of-service schedulers mylan**:

```
user@switch> show configuration class-of-service schedulers mylan
transmit-rate 3g;
shaping-rate percent 100;
priority low;
drop-profile-map loss-priority low protocol any drop-profile lp-profile;
```

```
drop-profile-map loss-priority medium-high protocol any drop-profile mh-profile;  
drop-profile-map loss-priority high protocol any drop-profile h-profile;
```



NOTE: This example does not include configuring scheduler bandwidth and priority. This information (transmit rate, shaping rate, and priority) is shown for completeness.

**Related
Documentation**

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 6771](#)
- [Example: Configuring Queue Schedulers](#)
- [Example: Configuring WRED Drop Profiles](#)
- [Configuring CoS Drop Profile Maps on page 6817](#)
- [Understanding CoS WRED Drop Profiles](#)

Understanding CoS Explicit Congestion Notification

Explicit congestion notification (ECN) enables end-to-end congestion notification between two endpoints on TCP/IP based networks. The two endpoints are an ECN-enabled sender and an ECN-enabled receiver. ECN must be enabled on both endpoints and on all of the intermediate devices between the endpoints for ECN to work properly. Any device in the transmission path that does not support ECN breaks the end-to-end ECN functionality.

ECN notifies networks about congestion with the goal of reducing packet loss and delay by making the sending device decrease the transmission rate until the congestion clears, without dropping packets. RFC 3168, *The Addition of Explicit Congestion Notification (ECN) to IP*, defines ECN.

ECN is disabled by default. Normally, you enable ECN only on queues that handle best-effort traffic because other traffic types use different methods of congestion notification—lossless traffic uses priority-based flow control (PFC) and strict-high priority traffic receives all of the port bandwidth it requires up to the point of a configured maximum rate.



NOTE: OCX Series switches do not support lossless transport and do not support PFC.

You enable ECN on individual output queues (as represented by forwarding classes) by enabling ECN in the queue scheduler configuration, mapping the scheduler to forwarding classes (queues), and then applying the scheduler to interfaces.



NOTE: For ECN to work on a queue, you must also apply a weighted random early detection (WRED) packet drop profile to the queue.

- [How ECN Works on page 6820](#)
- [WRED Drop Profile Control of ECN Thresholds on page 6825](#)
- [Support, Limitations, and Notes on page 6828](#)

How ECN Works

Without ECN, switches respond to network congestion by dropping TCP/IP packets. Dropped packets signal the network that congestion is occurring. Devices on the IP network respond to TCP packet drops by reducing the packet transmission rate to allow the congestion to clear. However, the packet drop method of congestion notification and management has some disadvantages. For example, packets are dropped and must be retransmitted. Also, bursty traffic can cause the network to reduce the transmission rate too much, resulting in inefficient bandwidth utilization.

Instead of dropping packets to signal network congestion, ECN marks packets to signal network congestion, without dropping the packets. For ECN to work, all of the switches

in the path between two ECN-enabled endpoints must have ECN enabled. ECN is negotiated during the establishment of the TCP connection between the endpoints.

ECN-enabled switches determine the queue congestion state based on the WRED packet drop profile configuration applied to the queue, so each ECN-enabled queue must also have a WRED drop profile. If a queue fills to the level at which the WRED drop profile has a packet drop probability greater than zero (0), the switch might mark a packet as experiencing congestion. Whether or not a switch marks a packet as experiencing congestion is the same probability as the drop probability of the queue at that fill level.

ECN communicates whether or not congestion is experienced by marking the two least-significant bits in the differentiated services (DiffServ) field in the IP header. The most significant six bits in the DiffServ field contain the Differentiated Services Code Point (DSCP) bits. The state of the two ECN bits signals whether or not the packet is an ECN-capable packet and whether or not congestion has been experienced.

ECN-capable senders mark packets as ECN-capable. If a sender is not ECN-capable, it marks packets as not not ECN-capable. If an ECN-capable packet experiences congestion at the egress queue of a switch, the switch marks the packet as experiencing congestion. When the packet reaches the ECN-capable receiver (destination endpoint), the receiver echoes the congestion indicator to the sender (source endpoint) by sending a packet marked to indicate congestion.

After receiving the congestion indicator from the receiver, the source endpoint reduces the transmission rate to relieve the congestion. This is similar to the result of TCP congestion notification and management, but instead of dropping the packet to signal network congestion, ECN marks the packet and the receiver echoes the congestion notification to the sender. Because the packet is not dropped, the packet does not need to be retransmitted.

- [ECN Bits in the DiffServ Field on page 6821](#)
- [End-to-End ECN Behavior on page 6822](#)
- [ECN Compared to PFC and Ethernet PAUSE on page 6824](#)

ECN Bits in the DiffServ Field

The two ECN bits in the DiffServ field provide four codes that determine if a packet is marked as an ECN-capable transport (ECT) packet, meaning that both endpoints of the transport protocol are ECN-capable, and if there is congestion experienced (CE), as shown in [Table 578](#):

Table 578: ECN Bit Codes

ECN Bits (Code)	Meaning
00	Non-ECT—Packet is marked as not ECN-capable
01	ECT(1)—Endpoints of the transport protocol are ECN-capable
10	ECT(0)—Endpoints of the transport protocol are ECN-capable

Table 578: ECN Bit Codes (*continued*)

ECN Bits (Code)	Meaning
11	CE—Congestion experienced

Codes 01 and 10 have the same meaning: the sending and receiving endpoints of the transport protocol are ECN-capable. There is no difference between these codes.

End-to-End ECN Behavior

After the sending and receiving endpoints negotiate ECN, the sending endpoint marks packets as ECN-capable by setting the DiffServ ECN field to ECT(1) (01) or ECT(0) (10). Every intermediate switch between the endpoints must have ECN enabled or it does not work.

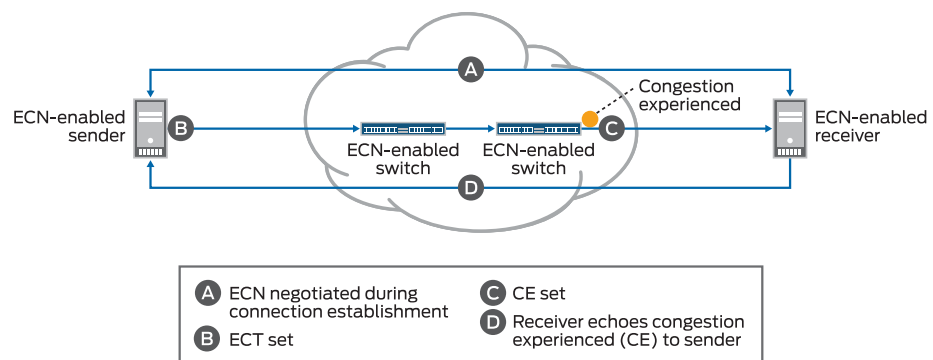
When a packet traverses a switch and experiences congestion at an output queue that uses the WRED packet drop mechanism, the switch marks the packet as experiencing congestion by setting the DiffServ ECN field to CE (11). Instead of dropping the packet (as with TCP congestion notification), the switch forwards the packet.



NOTE: At the egress queue, the WRED algorithm determines whether or not a packet is drop eligible based on the queue fill level (how full the queue is). If a packet is drop eligible and marked as ECN-capable, the packet can be marked CE and forwarded. If a packet is drop eligible and is not marked as ECN-capable, it might be dropped. See “[WRED Drop Profile Control of ECN Thresholds](#)” on page 6825 for more information about the WRED algorithm.

When the packet reaches the receiver endpoint, the CE mark tells the receiver that there is network congestion. The receiver then sends (echoes) a message to the sender that indicates there is congestion on the network. The sender acknowledges the congestion notification message and reduces its transmission rate. [Figure 224](#) summarizes how ECN works to mitigate network congestion:

Figure 224: Explicit Congestion Notification



8042495

End-to-end ECN behavior includes:

1. The ECN-capable sender and receiver negotiate ECN capability during the establishment of their connection.
2. After successful negotiation of ECN capability, the ECN-capable sender sends IP packets with the ECT field set to the receiver.



NOTE: All of the intermediate devices in the path between the sender and the receiver must be ECN-enabled.

3. If the WRED algorithm on a switch egress queue determines that the queue is experiencing congestion and the packet is drop eligible, the switch can mark the packet as “congestion experienced” (CE) to indicate to the receiver that there is congestion on the network. If the packet has already been marked CE (congestion has already been experienced at the egress of another switch), the switch forwards the packet with CE marked.

If there is no congestion at the switch egress queue, the switch forwards the packet and does not change the ECT-enabled marking of the ECN bits, so the packet is still marked as ECN-capable but not as experiencing congestion.

On QFX5200, QFX5100, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems, packets that are not marked as ECN-capable (ECT, 00) are treated according to the WRED drop profile configuration and might be dropped during periods of congestion.

On QFX10000 switches, the switch uses the tail-drop algorithm to drop packets that are marked ECT (00) during periods of congestion. (When a queue fills to its maximum level of fullness, tail-drop simply drops all subsequently arriving packets until there is space in the queue to buffer more packets. All non-ECN-capable packets are treated the same.)

4. The receiver receives a packet marked CE to indicate that congestion was experienced along the congestion path.
5. The receiver echoes (sends) a packet back to the sender with the ECE bit (bit 9) marked in the flag field of the TCP header. The ECE bit is the ECN echo flag bit, which notifies the sender that there is congestion on the network.
6. The sender reduces the data transmission rate and sends a packet to the receiver with the CWR bit (bit 8) marked in the flag field of the TCP header. The CWR bit is the congestion window reduced flag bit, which acknowledges to the receiver that the congestion experienced notification was received.
7. When the receiver receives the CWR flag, the receiver stops setting the ECE bit in replies to the sender.

Table 579 summarizes the behavior of traffic on ECN-enabled queues.

Table 579: Traffic Behavior on ECN-Enabled Queues

Incoming IP Packet Marking of ECN Bits	ECN Configuration on the Output Queue	Action if WRED Algorithm Determines Packet is Drop Eligible	Outgoing Packet Marking of ECN Bits
Non-ECT (00)	Does not matter	Drop (QFX5200, QFX5100, EX4600, QFX3500, QFX3600, QFabric systems). Tail drop occurs when queue reaches maximum fullness because no WRED drop probability is applied (QFX10000 switches).	No ECN bits marked
ECT (10 or 01)	ECN disabled	Drop	Packet dropped—no ECN bits marked
ECT (10 or 01)	ECN enabled	Do not drop. Mark packet as experiencing congestion (CE, bits 11).	Packet marked ECT (11) to indicate congestion
CE (11)	ECN disabled	Drop	Packet dropped—no ECN bits marked
CE (11)	ECN enabled	Do not drop. Packet is already marked as experiencing congestion, forward packet without changing the ECN marking.	Packet marked ECT (11) to indicate congestion

When an output queue is not experiencing congestion as defined by the WRED drop profile mapped to the queue, all packets are forwarded, and no packets are dropped.

ECN Compared to PFC and Ethernet PAUSE

ECN is an end-to-end network congestion notification mechanism for IP traffic. Priority-based flow control (PFC) (IEEE 802.1Qbb) and Ethernet PAUSE (IEEE 802.3X) are different types of congestion management mechanisms.



NOTE: QFX10000 switches do not support Ethernet PAUSE.

OCX Series switches do not support PFC. OCX Series switches support Ethernet PAUSE on tagged Layer 3 interfaces.

ECN requires that an output queue must also have an associated WRED packet drop profile. Output queues used for traffic on which PFC is enabled should not have an associated WRED drop profile. Interfaces on which Ethernet PAUSE is enabled should not have an associated WRED drop profile.

PFC is a peer-to-peer flow control mechanism to support lossless traffic. PFC enables connected peer devices to pause flow transmission during periods of congestion. PFC enables you to pause traffic on a specified type of flow on a link instead of on all traffic

on a link. For example, you can (and should) enable PFC on lossless traffic classes such as the **fcoe** forwarding class. Ethernet PAUSE is also a peer-to-peer flow control mechanism, but instead of pausing only specified traffic flows, Ethernet PAUSE pauses all traffic on a physical link.

With PFC and Ethernet PAUSE, the sending and receiving endpoints of a flow do not communicate congestion information to each other across the intermediate switches. Instead, PFC controls flows between two PFC-enabled peer devices (for example, switches) that support data center bridging (DCB) standards. PFC works by sending a pause message to the connected peer when the flow output queue becomes congested. Ethernet PAUSE simply pauses all traffic on a link during periods of congestion and does not require DCB.

PFC works this way: if a switch output queue fills to a certain threshold, the switch sends a PFC pause message to the connected peer device that is transmitting data. The pause message tells the transmitting switch to pause transmission of the flow. When the congestion clears, the switch sends another PFC message to tell the connected peer to resume transmission. (If the output queue of the transmitting switch also reaches a certain threshold, that switch can in turn send a PFC pause message to the connected peer that is transmitting to it. In this way, PFC can propagate a transmission pause back through the network.)

See [“Understanding CoS Flow Control \(Ethernet PAUSE and PFC\)” on page 6493](#) for more information. For QFX5100 and EX4600 switches only, you can also refer to [“Understanding PFC Functionality Across Layer 3 Interfaces” on page 6921](#).

WRED Drop Profile Control of ECN Thresholds

You apply WRED drop profiles to forwarding classes (which are mapped to output queues) to control how the switch marks ECN-capable packets. A scheduler map associates a drop profile with a scheduler and a forwarding class, and then you apply the scheduler map to interfaces to implement the scheduling properties for the forwarding class on those interfaces.

Drop profiles define queue fill level (the percentage of queue fullness) and drop probability (the percentage probability that a packet is dropped) pairs. When a queue fills to a specified level, traffic that matches the drop profile has the drop probability paired with that fill level. When you configure a drop profile, you configure pairs of fill levels and drop probabilities to control how packets drop at different levels of queue fullness.

The first fill level and drop probability pair is the drop start point. Until the queue reaches the first fill level, packets are not dropped. When the queue reaches the first fill level, packets that exceed the fill level have a probability of being dropped that equals the drop probability paired with the fill level.

The last fill level and drop probability pair is the drop end point. When the queue reaches the last fill level, all packets are dropped unless they are configured for ECN.



NOTE: Lossless queues (forwarding class configured with the no-loss packet drop attribute) and strict-high priority queues do not use drop profiles. Lossless queues use PFC to control the flow of traffic. Strict-high priority queues receive all of the port bandwidth they require up to the configured maximum bandwidth limit (scheduler transmit-rate on QFX10000 switches, and shaping-rate on QFX5200, QFX5100, QFX3500, QFX3600, and EX4600 switches, and QFabric systems).

Different switches support different amounts of fill level/drop probability pairs in drop profiles. For example, QFX10000 switches support 32 fill level/drop probability pairs, so there can be as many as 30 intermediate fill level/drop probability pairs between the drop start and drop endpoints. QFX5200, QFX5100, QFX3500, QFX3600, and EX4600 switches, and QFabric systems support two fill level/drop probability pairs—by definition, the two pairs you configure on these switches are the drop start and drop end points.



NOTE: Do not configure the last fill level as 100 percent.

The drop profile configuration affects ECN packets as follows:

- Drop start point—ECN-capable packets might be marked as congestion experienced (CE).
- Drop end point—ECN-capable packets are always marked CE.

As a queue fills from the drop start point to the drop end point, the probability that an ECN packet is marked CE is the same as the probability that a non-ECN packet is dropped if you apply the drop profile to best-effort traffic. As the queue fills, the probability of an ECN packet being marked CE increases, just as the probability of a non-ECN packet being dropped increases when you apply the drop profile to best-effort traffic.

At the drop end point, all ECN packets are marked CE, but the ECN packets are not dropped. When the queue fill level exceeds the drop end point, all ECN packets are marked CE. (At this point on QFX5200, QFX5100, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems, all non-ECN packets are dropped.) ECN packets (and all other packets) are tail-dropped if the queue fills completely.

To configure a WRED packet drop profile and apply it to an output queue (using hierarchical scheduling on switches that support ETS):

1. Configure a drop profile using the statement **set class-of-service drop-profiles *profile-name* interpolate fill-level *drop-start-point* fill-level *drop-end-point* drop-probability 0 drop-probability *percentage***.
2. Map the drop profile to a queue scheduler using the statement **set class-of-service schedulers *scheduler-name* drop-profile-map loss-priority (low | medium-high | high) protocol any drop-profile *profile-name***. The name of the drop-profile is the name of the WRED profile configured in Step 1.

3. Map the scheduler, which Step 2 associates with the drop profile, to the output queue using the statement **set class-of-service scheduler-maps *map-name* forwarding-class *forwarding-class-name* scheduler *scheduler-name***. The forwarding class identifies the output queue. Forwarding classes are mapped to output queues by default, and can be remapped to different queues by explicit user configuration. The scheduler name is the scheduler configured in Step 2.
4. Associate the scheduler map with a traffic control profile using the statement **set class-of-service traffic-control-profiles *tcp-name* scheduler-map *map-name***. The scheduler map name is the name configured in Step 3.
5. Associate the traffic control profile with an interface using the statement **set class-of-service interface *interface-name* forwarding-class-set *forwarding-class-set-name* output-traffic-control-profile *tcp-name***. The output traffic control profile name is the name of the traffic control profile configured in Step 4.

The interface uses the scheduler map in the traffic control profile to apply the drop profile (and other attributes, including the enable ECN attribute) to the output queue (forwarding class) on that interface. Because you can use different traffic control profiles to map different schedulers to different interfaces, the same queue number on different interfaces can handle traffic in different ways.

To configure a WRED packet drop profile and apply it to an output queue on switches that support port scheduling (ETS hierarchical scheduling is either not supported or not used):

1. Configure a drop profile using the statement **set class-of-service drop-profiles *profile-name* interpolate fill-level *level1* *level2* ... *level32* drop-probability *probability1* *probability2* ... *probability32***. You can specify as few as two fill level/drop probability pairs or as many as 32 pairs.
2. Map the drop profile to a queue scheduler using the statement **set class-of-service schedulers *scheduler-name* drop-profile-map loss-priority (low | medium-high | high) drop-profile *profile-name***. The name of the drop-profile is the name of the WRED profile configured in Step 1.
3. Map the scheduler, which Step 2 associates with the drop profile, to the output queue using the statement **set class-of-service scheduler-maps *map-name* forwarding-class *forwarding-class-name* scheduler *scheduler-name***. The forwarding class identifies the output queue. Forwarding classes are mapped to output queues by default, and can be remapped to different queues by explicit user configuration. The scheduler name is the scheduler configured in Step 2.
4. Associate the scheduler map with an interface using the statement **set class-of-service interfaces *interface-name* scheduler-map *scheduler-map-name***.

The interface uses the scheduler map to apply the drop profile (and other attributes) to the output queue mapped to the forwarding class on that interface. Because you can use different scheduler maps on different interfaces, the same queue number on different interfaces can handle traffic in different ways.

Support, Limitations, and Notes

If the WRED algorithm that is mapped to a queue does not find a packet drop eligible, then the ECN configuration and ECN bits marking does not matter. The packet transport behavior is the same as when ECN is not enabled.

ECN is disabled by default. Normally, you enable ECN only on queues that handle best-effort traffic, and you do not enable ECN on queues that handle lossless traffic or strict-high priority traffic.

ECN supports the following:

- IPv4 and IPv6 packets
- Untagged, single-tagged, and double-tagged packets
- The outer IP header of IP tunneled packets (but not the inner IP header)

ECN does not support the following:

- IP packets with MPLS encapsulation
- The inner IP header of IP tunneled packets (however, ECN works on the outer IP header)
- Multicast, broadcast, and destination lookup fail (DLF) traffic
- Non-IP traffic



NOTE: On QFX10000 switches, when you enable a queue for ECN and apply a WRED drop profile to the queue, the WRED drop profile only sets the thresholds for marking ECN traffic as experiencing congestion (CE, 11). On ECN-enabled queues, the WRED drop profile does not set drop thresholds for non-ECT (00) traffic (traffic that is not ECN-capable). Instead, the switch uses the tail-drop algorithm on traffic that is marked non-ECT on ECN-enabled queues during periods of congestion.

To apply a WRED drop profile to non-ECT traffic, configure a multifield (MF) classifier to assign non-ECT traffic to a different output queue that is not ECN-enabled, and then apply the WRED drop profile to that queue.

Related Documentation

- [Example: Configuring ECN on page 6828](#)

Example: Configuring ECN

This example shows how to enable explicit congestion notification (ECN) on an output queue.

- [Requirements on page 6829](#)
- [Overview on page 6829](#)

- [Configuration on page 6831](#)
- [Verification on page 6833](#)

Requirements

This example uses the following hardware and software components:

- One switch.
- Junos OS Release 13.2X51-D25 or later for the QFX Series or Junos OS Release 14.1X53-D20 for the OCX Series

Overview

ECN enables end-to-end congestion notification between two endpoints on TCP/IP based networks. The two endpoints are an ECN-enabled sender and an ECN-enabled receiver. ECN must be enabled on both endpoints and on all of the intermediate devices between the endpoints for ECN to work properly. Any device in the transmission path that does not support ECN breaks the end-to-end ECN functionality.

A weighted random early detection (WRED) packet drop profile must be applied to the output queues on which ECN is enabled. ECN uses the WRED drop profile thresholds to mark packets when the output queue experiences congestion.

ECN reduces packet loss by forwarding ECN-capable packets during periods of network congestion instead of dropping those packets. (TCP notifies the network about congestion by dropping packets.) During periods of congestion, ECN marks ECN-capable packets that egress from congested queues. When the receiver receives an ECN packet that is marked as experiencing congestion, the receiver echoes the congestion state back to the sender. The sender then reduces its transmission rate to clear the congestion.

ECN is disabled by default. You can enable ECN on best-effort traffic. ECN should not be enabled on lossless traffic queues, which uses priority-based flow control (PFC) for congestion notification, and ECN should not be enabled on strict-high priority traffic queues.

To enable ECN on an output queue, you not only need to enable ECN in the queue scheduler, you also need to:

- Configure a WRED packet drop profile.
- Configure a queue scheduler that includes the WRED drop profile and enables ECN. (This example shows only ECN and drop profile configuration; you can also configure bandwidth, priority, and buffer settings in a scheduler.)
- Map the queue scheduler to a forwarding class (output queue) in a scheduler map.
- If you are using enhanced transmission selection (ETS) hierarchical scheduling, add the forwarding class to a forwarding class set (priority group).

- If you are using ETS, associate the queue scheduler map with a traffic control profile (priority group scheduler for hierarchical scheduling).
- If you are using ETS, apply the traffic control profile and the forwarding class set to an interface. On that interface, the output queue uses the scheduler mapped to the forwarding class, as specified by the scheduler map attached to the traffic control profile. This enables ECN on the queue and applies the WRED drop profile to the queue.

If you are using port scheduling, apply the scheduler map to an interface. On that interface, the output queue uses the scheduler mapped to the forwarding class in the scheduler map, which enables ECN on the queue and applies the WRED drop profile to the queue.

Table 580 shows the configuration components for this example.

Table 580: Components of the ECN Configuration Example

Component	Settings
Hardware	QFX Series switch
Drop profile (with two fill level/drop probability pairs)	Name: be-dp Drop start fill level: 30 percent Drop end fill level: 75 percent Drop probability at drop start (minimum drop rate): 0 percent Drop probability at drop end (maximum drop rate): 80 percent
Scheduler	Name: be-sched ECN: enabled Drop profile: be-dp Transmit rate: 25% Buffer size: 25% Priority: low
Scheduler map	Name: be-map Forwarding class: best-effort Scheduler: be-sched NOTE: By default, the best-effort forwarding class is mapped to output queue 0 .
Forwarding class set (ETS only)	Name: be-pg Forwarding class: best-effort (queue 0)
Traffic control profile (ETS only)	Name: be-tcp Scheduler map: be-map
Interface (ETS only)	Name: xe-0/0/20 Forwarding class set: be-pg (Output) traffic control profile: be-tcp
Interface (port scheduling only)	Name: xe-0/0/20



NOTE: Only switches that support ETS hierarchical scheduling support forwarding class set and traffic control profile configuration. Direct port scheduling does not use the hierarchical scheduling structure.



NOTE: On QFX5100, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems, the WRED drop profile also controls packet drop behavior for traffic that is not ECN-capable (packets marked non-ECT, ECN bit code 00).

On QFX10000 switches, when ECN is enabled on a queue, the WRED drop profile only sets the ECN thresholds, it does not control packet drop on non-ECN packets. On ECN-enabled queues, QFX10000 switches use the tail-drop algorithm on non-ECN packets during periods of congestion. If you do not enable ECN, then the queue uses the WRED packet drop mechanism.

Configuration

CLI Quick Configuration

To quickly configure the drop profile, scheduler with ECN enabled, and to map the scheduler to an output queue on an interface, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

ETS Quick Configuration

```
[edit class-of-service]
set drop-profile be-dp interpolate fill-level 30 fill-level 75 drop-probability 0 drop-probability 80
set schedulers be-sched explicit-congestion-notification
set schedulers be-sched drop-profile-map loss-priority low protocol any drop-profile be-dp
set schedulers be-sched transmit-rate percent 25
set schedulers be-sched buffer-size percent 25
set schedulers be-sched priority low
set scheduler-maps be-map forwarding-class best-effort scheduler be-sched
set forwarding-class-sets be-pg class best-effort
set traffic-control-profiles be-tcp scheduler-map be-map
set interfaces xe-0/0/20 forwarding-class-set be-pg output-traffic-control-profile be-tcp
```

Port Scheduling Quick Configuration (QFX10000 Switches)

```
[edit class-of-service]
set drop-profile be-dp interpolate fill-level 30 fill-level 75 drop-probability 0 drop-probability 80
set schedulers be-sched explicit-congestion-notification
set schedulers be-sched drop-profile-map loss-priority low protocol any drop-profile be-dp
set schedulers be-sched transmit-rate percent 25
set schedulers be-sched buffer-size percent 25
set schedulers be-sched priority low
set scheduler-maps be-map forwarding-class best-effort scheduler be-sched
set interfaces xe-0/0/20 scheduler-map be-map
```

Configuring ECN

Step-by-Step Procedure

To configure ECN:

1. Configure the WRED packet drop profile **be-dp**. This example uses a drop start point of **30** percent, a drop end point of **75** percent, a minimum drop rate of **0** percent, and a maximum drop rate of **80** percent:

```
[edit class-of-service]
user@switch# set drop-profile be-dp interpolate fill-level 30 fill-level 75 drop-probability
0 drop-probability 80
```

2. Create the scheduler **be-sched** with ECN enabled and associate the drop profile **be-dp** with the scheduler:

```
[edit class-of-service]
user@switch# set schedulers be-sched explicit-congestion-notification
user@switch# set schedulers be-sched drop-profile-map loss-priority low protocol any
drop-profile be-dp
user@switch# set be-sched transmit-rate percent 25
user@switch# set be-sched buffer-size percent 25
user@switch# set be-sched priority low
```

3. Map the scheduler **be-sched** to the **best-effort** forwarding class (output queue 0) using scheduler map **be-map**:

```
[edit class-of-service]
user@switch# set scheduler-maps be-map forwarding-class best-effort scheduler be-sched
```

4. If you are using ETS, add the forwarding class **best-effort** to the forwarding class set **be-pg**; if you are using direct port scheduling, skip this step:

```
[edit class-of-service]
user@switch# set forwarding-class-sets be-pg class best-effort
```

5. If you are using ETS, associate the scheduler map **be-map** with the traffic control profile **be-tcp**; if you are using direct port scheduling, skip this step:

```
[edit class-of-service]
user@switch# set traffic-control-profiles be-tcp scheduler-map be-map
```

6. If you are using ETS, associate the traffic control profile **be-tcp** and the forwarding class set **be-pg** with the interface on which you want to enable ECN on the best-effort queue:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 forwarding-class-set be-pg
output-traffic-control-profile be-tcp
```

If you are using direct port scheduling, associate the scheduler map **be-map** with the interface on which you want to enable ECN on the best-effort queue:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 scheduler map be-map
```


Verification

Verifying That ECN Is Enabled

Purpose Verify that ECN is enabled in the scheduler **be-sched** by showing the configuration for the scheduler map **be-map**.

Action Display the scheduler map configuration using the operational mode command **show class-of-service scheduler-map be-map**:

```
user@switch> show class-of-service scheduler-map be-map
Scheduler map: be-map, Index: 12240
```

```
Scheduler:be-sched, Forwarding class: best-effort, Index: 115
Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent,
Buffer Limit: none, Priority: low
Excess Priority: unspecified, Explicit Congestion Notification: enable
Drop profiles:
  Loss priority  Protocol  Index  Name
  Low           any       3312   be-dp
  Medium-high   any       1      <default-drop-profile>
  High          any       1      <default-drop-profile>
```

Meaning The **show class-of-service scheduler-map** operational command shows the configuration of the scheduler associated with the scheduler map and the forwarding class mapped to that scheduler. The output shows that:

- The scheduler associated with the scheduler map is **be-sched**.
- The scheduler map applies to the forwarding class **best-effort** (output queue 0).
- The scheduler **be-sched** has a transmit rate of **25** percent, a queue buffer size of **25** percent, and a drop priority of **low**.
- Explicit congestion notification state is **enable**.
- The WRED drop profile used for low drop priority traffic is **be-dp**.

Related Documentation

- [Understanding CoS Explicit Congestion Notification on page 6820](#)

PART 93

Data Center Bridging (PFC, DCBX) and Flow Control

- [Using Data Center Bridging and Flow Control on page 6837](#)

Using Data Center Bridging and Flow Control

- [Understanding DCB Features and Requirements on page 6838](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 6841](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 6847](#)
- [Disabling the ETS Recommendation TLV on page 6872](#)
- [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 6873](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 6893](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 6905](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 6908](#)
- [Troubleshooting Dropped FCoE Traffic on page 6917](#)
- [Understanding PFC Functionality Across Layer 3 Interfaces on page 6921](#)
- [Example: Configuring PFC Across Layer 3 Interfaces on page 6923](#)
- [Example: Configuring CoS for FCoE Transit Switch Traffic Across an MC-LAG on page 6943](#)
- [Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic \(FCoE Transit Switch\) on page 6969](#)
- [Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 6977](#)
- [Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 6986](#)
- [Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications \(FCoE and iSCSI\) on page 7000](#)
- [Understanding DCBX on page 7016](#)
- [Configuring the DCBX Mode on page 7026](#)
- [Configuring DCBX Autonegotiation on page 7027](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 7030](#)
- [Defining an Application for DCBX Application Protocol TLV Exchange on page 7034](#)
- [Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 7035](#)

- [Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 7036](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 7037](#)

Understanding DCB Features and Requirements

Data center bridging (DCB) is a set of enhancements to the IEEE 802.1 bridge specifications. DCB modifies and extends Ethernet behavior to support I/O convergence in the data center. I/O convergence includes but is not limited to the transport of Ethernet LAN traffic and Fibre Channel (FC) storage area network (SAN) traffic on the same physical Ethernet network infrastructure.



Video: [What is Data Center Bridging?](#)

A converged architecture saves cost by reducing the number of networks and switches required to support both types of traffic, reducing the number of interfaces required, reducing cable complexity, and reducing administration activities.

The Juniper Networks QFX Series and EX4600 switches support the DCB features required to transport converged Ethernet and FC traffic while providing the class-of-service (CoS) and other characteristics FC requires for transmitting storage traffic. To accommodate FC traffic, DCB specifications provide:

- A flow control mechanism called priority-based flow control (PFC, described in IEEE 802.1Qbb) to help provide lossless transport.
- A discovery and exchange protocol for conveying configuration and capabilities among neighbors to ensure consistent configuration across the network, called Data Center Bridging Capability Exchange protocol (DCBX), which is an extension of Link Layer Data Protocol (LLDP, described in IEEE 802.1AB).
- A bandwidth management mechanism called enhanced transmission selection (ETS, described in IEEE 802.1Qaz).
- A congestion management mechanism called quantized congestion notification (QCN, described in IEEE 802.1Qau).

The switch supports the PFC, DCBX, and ETS standards but does not support QCN. The switch also provides the high-bandwidth interfaces (10-Gbps minimum) required to support DCB and converged traffic.

This topic describes the DCB standards and requirements the switch supports:

- [Lossless Transport on page 6839](#)
- [ETS on page 6840](#)
- [DCBX on page 6840](#)

Lossless Transport

FC traffic requires lossless transport (defined as no frames dropped because of congestion). Standard Ethernet does not support lossless transport, but the DCB extensions to Ethernet along with proper buffer management enable an Ethernet network to provide the level of class of service (CoS) necessary to transport FC frames encapsulated in Ethernet over an Ethernet network.

This section describes these factors in creating lossless transport over Ethernet:

- [PFC on page 6839](#)
- [Buffer Management on page 6839](#)
- [Physical Interfaces on page 6839](#)

PFC

PFC is a link-level flow control mechanism similar to Ethernet PAUSE (described in IEEE 802.3x). Ethernet PAUSE stops all traffic on a link for a period of time. PFC enables you to divide traffic on a link into eight priorities and stop the traffic of a selected priority without stopping the traffic assigned to other priorities on the link.

Pausing the traffic of a selected priority enables you to provide lossless transport for traffic assigned that priority and at the same time use standard lossy Ethernet transport for the rest of the link traffic.

Buffer Management

Buffer management is critical to the proper functioning of PFC, because if buffers are allowed to overflow, frames are dropped and transport is not lossless.

For each lossless flow priority, the switch requires sufficient buffer space to:

- Store frames sent during the time it takes to send the PFC pause frame across the cable between devices.
- Store the frames that are already on the wire when the sender receives the PFC pause frame.

The propagation delay due to cable length and speed, as well as processing speed, determines the amount of buffer space needed to prevent frame loss due to congestion.

The switch automatically sets the threshold for sending PFC pause frames to accommodate delay from cables as long as 150 meters (492 feet) and to accommodate large frames that might be on the wire when the switch sends the pause frame. This ensures that the switch sends pause frames early enough to allow the sender to stop transmitting before the receive buffers on the switch overflow.

Physical Interfaces

QFX Series switches support 10-Gbps or faster, full-duplex interfaces. The switch enables DCB capability only on 10-Gbps or faster Ethernet interfaces.

ETS

PFC divides traffic into up to eight separate streams (priorities, configured on the switch as forwarding classes) on a physical link. ETS enables you to manage the link bandwidth by:

- Grouping the priorities into priority groups (configured on the switch as forwarding class sets).
- Specifying the bandwidth available to each of the priority groups as a percentage of the total available link bandwidth.
- Allocating the bandwidth to the individual priorities in the priority group.

The available link bandwidth is the bandwidth remaining after servicing strict-high priority queues. On QFX5200, QFX5100, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems, we recommend that you always configure a shaping rate to limit the amount of bandwidth a strict-high priority queue can consume by including the [shaping-rate](#) statement in the **[edit class-of-service schedulers]** hierarchy on the strict-high priority scheduler. This prevents a strict-high priority queue from starving other queues on the port. (On QFX10000 switches, configure a transmit rate on strict-high priority queues to set a maximum amount of bandwidth for strict-high priority traffic.)

Managing link bandwidth with ETS provides several advantages:

- There is uniform management of all types of traffic on the link, both congestion-managed traffic and standard Ethernet traffic.
- When a priority group does not use all of its allocated bandwidth, other priority groups on the link can use that bandwidth as needed.

When a priority in a priority group does not use all of its allocated bandwidth, other priorities in the group can use that bandwidth.

The result is better bandwidth utilization, because priorities that consist of bursty traffic can share bandwidth during periods of low traffic transmission instead of consuming their entire bandwidth allocation when traffic loads are light.

- You can assign traffic types with different service needs to different priorities so that each traffic type receives appropriate treatment.
- Strict priority traffic retains its allocated bandwidth.

DCBX

DCB devices use DCBX to exchange configuration information with directly connected peers (switches and endpoints such as servers). DCBX is an extension of LLDP. If you disable LLDP on an interface, that interface cannot run DCBX. If you attempt to enable DCBX on an interface on which LLDP is disabled, the configuration commit fails.

DCBX can:

- Discover the DCB capabilities of peers.

- Detect DCB feature misconfiguration or mismatches between peers.
- Configure DCB features on peers.

You can configure DCBX operation for PFC, ETS, and for Layer 2 and Layer 4 applications such as FCoE and iSCSI. DCBX is enabled or disabled on a per-interface basis.

Related Documentation

- [Understanding FCoE on page 6436](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 6765](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 6493](#)
- [Understanding DCBX on page 6514](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 6504](#)

Understanding CoS Hierarchical Port Scheduling (ETS)

Scheduling defines the class-of-service (CoS) properties of output queues. Output queues are mapped to forwarding classes. CoS scheduler properties include the amount of interface bandwidth assigned to the queue, the queue priority, and the drop profiles associated with the queue.

Hierarchical port scheduling is a two-tier process that provides better port bandwidth utilization and greater flexibility to allocate resources to queues (forwarding classes) and to groups of queues (forwarding class sets). Hierarchical scheduling includes the Junos OS implementation of enhanced transmission selection (ETS), as described in IEEE 802.1Qaz.



Video: [What is Enhanced Transmission Selection?](#)

This topic describes:

- [Hierarchical Scheduling Tiers on page 6841](#)
- [Hierarchical Scheduling and ETS on page 6842](#)
- [ETS Advertisement in DCBX on page 6844](#)
- [Hierarchical Scheduling Process on page 6844](#)
- [Strict-High Priority Queues and Hierarchical Scheduling on page 6845](#)
- [Default Hierarchical Scheduling on page 6846](#)

Hierarchical Scheduling Tiers

The two tiers used in hierarchical scheduling are priorities and priority groups, as shown in [Table 574](#).

Table 581: Hierarchical Scheduling Tiers

Junos OS Configuration Construct	Equivalent ETS Construct	Description
Forwarding class	Priority	<p>Think about priorities (forwarding classes) as output queues. You map forwarding classes to queues, so each forwarding class represents an output queue.</p> <p>When you use a classifier to map a forwarding class to an IEEE 802.1p code point, the code point identifies that traffic's priority for priority-based flow control (PFC). Thus the forwarding class, the queue mapped to the forwarding class, and the priority (code point) mapped to the forwarding class all identify the same traffic.</p> <p>NOTE: OCX Series switches do not support lossless transport or PFC.</p>
Forwarding class set	Priority group	<p>Priority groups (forwarding class sets) are groups of priorities (forwarding classes). Forwarding class membership in a forwarding class set defines the priority group to which each priority belongs.</p> <p>You can configure up to three unicast priority groups and one multicast priority group.</p>

You apply scheduling properties to each hierarchical scheduling tier as described in the next section.



NOTE: If you explicitly configure one or more priority groups on an interface, any priority (forwarding class) that is not assigned to a priority group (forwarding class set) on that interface is assigned to an automatically generated default priority group and receives *no bandwidth*. This means that if you configure hierarchical scheduling on an interface, every forwarding class that you want to forward traffic on that interface must belong to a forwarding class set.



NOTE: On OCX Series switches, by default, classifiers use DSCP code points to map traffic to forwarding classes. However, hierarchical scheduling works in the same manner as when you use IEEE 802.1p code points to classify traffic. The OCX Series classifies traffic into forwarding classes based on DSCP code points, the forwarding classes are mapped to forwarding class sets, and you apply scheduling properties to each of the two tiers.

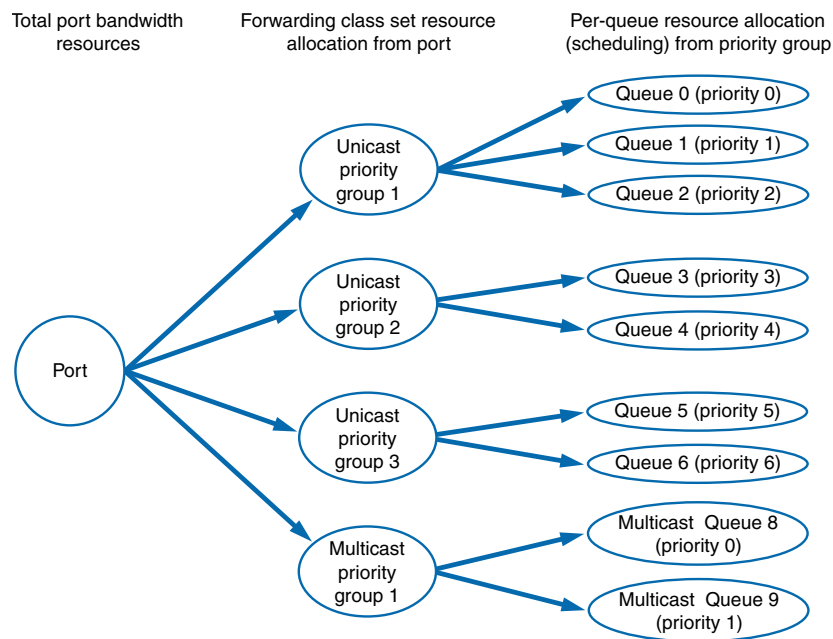
Hierarchical Scheduling and ETS

Two-tier hierarchical scheduling manages bandwidth efficiently by enabling you to define the CoS properties for each priority group and for each priority. The first tier of the hierarchical scheduler allocates port bandwidth to a priority group. The second tier of

the hierarchical scheduler determines the portion of the priority group bandwidth that a priority (queue) can use.

The CoS properties of a priority group define the amount of port bandwidth resources available to the queues in that priority group. The CoS properties you configure for each queue specify the amount of the bandwidth available to the queue from the bandwidth allocated to the priority group. [Figure 218](#) shows the relationship of port resource allocation to priority groups, and priority group resource allocation to queues (priorities).

Figure 225: Hierarchical Scheduling Tiers



9040722

If a queue (priority) does not use its allocated bandwidth, ETS shares the unused bandwidth among the other queues in the priority group in proportion to the minimum guaranteed rate (transmit rate) scheduled for each queue. If a priority group does not use its allocated bandwidth, ETS shares the unused bandwidth among the priority groups on the port in proportion to the minimum guaranteed rate (guaranteed rate) scheduled for each priority group.

In this way, ETS improves link bandwidth utilization, and it provides each queue and each priority group with the maximum available bandwidth. For example, priorities that consist of bursty traffic can share bandwidth during periods of low traffic transmission, instead of reserving their entire bandwidth allocation when traffic loads are light. All Juniper switches use ETS scheduling, except for QFX5200 and the QFX10000 switches.



NOTE: The available link bandwidth is the bandwidth remaining after servicing strict-high priority flows. Strict-high priority takes precedence over all other traffic (we recommend that you configure a shaping rate to limit the maximum amount of bandwidth that a strict-high priority forwarding class can use to prevent starving other queues).

ETS Advertisement in DCBX

When you configure hierarchical scheduling on a port, Data Center Bridging Capability Exchange protocol (DCBX) advertises:

- Each priority group
- The priorities in each priority group
- The bandwidth properties of each priority group and priority

When you configure hierarchical scheduling on a port, any priority that is not part of an explicitly configured priority group is assigned to the automatically generated default priority group and receives no bandwidth. The default priority group is transparent. It does not appear in the configuration.



NOTE: OCX Series switches do support DCBX, so hierarchical scheduling information is not exchanged with connected peers on OCX Series switches.

Hierarchical Scheduling Process

Hierarchical scheduling consists of multiple configuration steps that create the priorities and the priority groups, schedule their resources, and assign them to interfaces. The steps below correspond to the six blocks in the packet flow diagram shown in [Figure 219](#):


1. Packet classification:
 - Configure classification of incoming traffic into forwarding classes (priorities). This consists of either using the default classifiers or configuring classifiers to map code points and loss priorities to the forwarding classes.
 - Apply the classifiers to ingress interfaces or use the default classifiers. Applying a classifier to an interface groups incoming traffic on the interface into forwarding classes and loss priorities, by applying the classifier code point mapping to the incoming traffic.
2. Configure the output queues for the forwarding classes (priorities). This consists of either using the default forwarding classes and forwarding-class-to-queue mapping, or creating your own forwarding classes and mapping them to output queues.
3. Allocate resources to the forwarding classes:
 - Define resources for the priorities. This consists of configuring schedulers to set minimum guaranteed bandwidth, maximum bandwidth, drop profiles for Weighted Random Early Detection (WRED), and bandwidth priority to apply to a forwarding class. Extra bandwidth is shared among queues in proportion to the minimum guaranteed bandwidth (transmit rate) of each queue.
 - Map resources to priorities. This consists of mapping forwarding classes to schedulers, using a scheduler map.

4. Configure priority groups. This consists of mapping forwarding classes (priorities) to forwarding class sets (priority groups) to define the priorities that belong to each priority group.
5. Define resources for the priority groups. This consists of configuring traffic control profiles to set minimum guaranteed bandwidth (guaranteed rate) and maximum bandwidth (shaping rate) for a priority group. Traffic control profiles also specify a scheduler map, which defines the resources (schedulers) mapped to the priorities in the priority group. Extra port bandwidth is shared among priority groups in proportion to the minimum guaranteed bandwidth of each priority group.

The traffic control profile bandwidth settings determine the port resources available to the priority group. The schedulers specified in the scheduler map determine the amount of priority group resources that each priority receives.

6. Apply hierarchical scheduling to a port. This consists of attaching one or more priority groups (forwarding class sets) to an interface. For each priority group, you also attach a traffic control profile, which contains the scheduling properties of the priority group and the priorities in the priority group. Different priority groups on the same port can use different traffic control profiles, which provides fine tuned control of scheduling for each queue on each interface.

Figure 226: Hierarchical Scheduling Packet Flow



Strict-High Priority Queues and Hierarchical Scheduling

If you configure a strict-high priority queue, you must observe the following rules:

- You must create a separate forwarding class set (priority group) for the strict-high priority queue.
- Only one forwarding class set can contain strict-high priority queues.
- Strict-high priority queues cannot belong to the same forwarding class set as queues that are not strict-high priority.

- A strict-high priority queue cannot belong to a multidestination forwarding class set.
- We recommend that you always apply a shaping rate to strict-high priority queues to prevent them from starving other queues. If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.



NOTE: On a QFabric system, if a fabric (fte) interface handles strict-high priority traffic, you must define a separate forwarding class set (priority group) for strict-high priority traffic. Strict-high priority traffic cannot be mixed with traffic of other priorities in a forwarding class set. For example, you might choose to create different forwarding class sets for best effort, lossless, strict-high priority, and multidestination traffic.

Default Hierarchical Scheduling

If you do not explicitly configure hierarchical scheduling, the switch uses the default settings:

- The switch automatically creates a default forwarding class set that contains all of the forwarding classes on the switch. The switch assigns 100 percent of the port output bandwidth to the default forwarding class set. The default forwarding class set is transparent. It does not appear in the configuration and is used for Data Center Bridging Capability Exchange protocol (DCBX) advertisement.



NOTE: OCX Series switches do not support DCBX, so the ETS configuration is not advertised to connected peers.

- Ingress traffic is classified based on the default classifier settings.
- The forwarding classes (queues) in the default forwarding class set receive bandwidth based on the default scheduler settings.

Related Documentation

- [Understanding CoS Packet Flow on page 6628](#)
- [Understanding CoS Output Queue Schedulers](#)
- [Understanding CoS Priority Group Scheduling on page 6758](#)
- [Benefits of Configuring CoS Hierarchical Port Scheduling](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 6493](#)
- [Understanding CoS Classifiers](#)
- [Understanding CoS Classifiers](#)
- [Understanding Default CoS Scheduling and Classification](#)
- [Understanding Default CoS Scheduling and Classification](#)

- *Understanding CoS Scheduling on QFabric System Node Device Fabric (fte) Ports*
- *Understanding Default CoS Scheduling on QFabric System Interconnect Devices (Junos OS Release 13.1 and Later Releases)*
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 6771](#)
- [Example: Configuring Queue Schedulers](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 6762](#)
- [Example: Configuring Minimum Guaranteed Output Bandwidth on page 6798](#)
- [Example: Configuring Maximum Output Bandwidth on page 6805](#)

Example: Configuring CoS Hierarchical Port Scheduling (ETS)

Hierarchical port scheduling defines the class-of-service (CoS) properties of output queues, which are mapped to forwarding classes. Traffic is classified into forwarding classes based on code point (priority), so mapping queues to forwarding classes also maps queues to priorities). Hierarchical port scheduling enables you to group priorities that require similar CoS treatment into priority groups. You define the port bandwidth resources for a priority group, and you define the amount of the priority group's resources that each priority in the group can use.

Hierarchical port scheduling is the Junos OS implementation of enhanced transmission selection (ETS), as described in IEEE 802.1Qaz. One major benefit of hierarchical port scheduling is greater port bandwidth utilization. If a priority group on a port does not use all of its allocated bandwidth, other priority groups on that port can use that bandwidth. Also, if a priority within a priority group does not use its allocated bandwidth, other priorities within that priority group can use that bandwidth.

Configuring hierarchical scheduling is a multistep procedure that includes:

- Mapping forwarding classes to queues
- Defining forwarding class sets (priority groups)
- Defining behavior aggregate classifiers
- Configuring priority-based flow control (PFC) for lossless priorities (queues)
- Applying classifiers and PFC configuration to ingress interfaces
- Defining drop profiles
- Defining schedulers
- Mapping forwarding classes to schedulers
- Defining traffic control profiles
- Assigning priority groups and traffic control profiles to egress ports



NOTE: OCX Series switches do not support lossless transport and do not support PFC. Although this example includes configuring lossless transport with PFC, the portions of the example that do not pertain to lossless transport still apply to OCX Series switches. (You can configure hierarchical scheduling on OCX Series switches, but you cannot configure lossless transport or lossless forwarding classes.)

This example describes how to configure hierarchical scheduling:

- [Requirements on page 6848](#)
- [Overview on page 6848](#)
- [Configuration on page 6852](#)
- [Verification on page 6862](#)

Requirements

This example uses the following hardware and software components:

- One switch (this example was tested on a Juniper Networks QFX3500 Switch)
- Junos OS Release 11.1 or later for the QFX Series or Junos OS Release 14.1X53-D20 or later for the OCX Series

Overview

Keep the following considerations in mind when you plan the port bandwidth allocation for priority groups and for individual priorities:

- How much traffic and what types of traffic you expect to traverse the system.
- How you want to divide different types of traffic into priorities (forwarding classes) to apply different CoS treatment to different types of traffic. Dividing traffic into priorities includes:
 - Mapping the code points of ingress traffic to forwarding classes using behavior aggregate (BA) classifiers. This classifies incoming traffic into the appropriate forwarding class based on code point.
 - Mapping forwarding classes to output queues. This defines the output queue for each type of traffic.
 - Attaching the BA classifier to the desired ingress interfaces so that incoming traffic maps to the desired forwarding classes and queues.
- How you want to organize priorities into priority groups (forwarding class sets).

Traffic that requires similar treatment usually belongs in the same priority group. To do this, place forwarding classes that require similar bandwidth, loss, and other

characteristics in the same forwarding class set. For example, you can map all types of best-effort traffic forwarding classes into one forwarding class set.

- How much of the port bandwidth you want to allocate to each priority group and to each of the priorities in each priority group. The following considerations apply to bandwidth allocation:
 - Estimate how much traffic you expect in each forwarding class, and how much traffic you expect in each forwarding class set (the amount of traffic you expect in a forwarding class set is the aggregate amount of traffic in the forwarding classes that belong to the forwarding class set).
 - The combined minimum guaranteed bandwidth of the priorities (forwarding classes) in a priority group should not exceed the minimum guaranteed bandwidth of the priority group (forwarding class set). The transmit rate scheduler parameter defines the minimum guaranteed bandwidth for forwarding classes. Scheduler maps associate schedulers with forwarding classes.
 - The combined minimum guaranteed bandwidth of the priority groups (forwarding class sets) on a port should not exceed the port's total bandwidth. The guaranteed rate parameter in the traffic control profile defines the minimum bandwidth for a forwarding class set. Associating a scheduler map with a traffic control profile sets the scheduling for the individual forwarding classes in the forwarding class set.

This example creates hierarchical port scheduling by defining priority groups for best effort, guaranteed delivery, and high-performance computing (HPC) traffic. Each priority group includes priorities that need to receive similar CoS treatment. Each priority group and each priority within each priority group receive the CoS resources needed to service their flows. Lossless priorities use PFC to prevent packet loss when the network experiences congestion.

Topology

Table 575 shows the configuration components for this example.



NOTE: OCX Series switches do not support lossless transport and do not support PFC. If you eliminate the configuration elements for the default lossless fcoe and no-loss forwarding classes (including classifier, forwarding class set, scheduler, and traffic control profile configuration for those forwarding classes) and for PFC, this example works for OCX Series switches. However, because the default fcoe and no-loss forwarding classes do not carry traffic on OCX Series switches, you can apply the bandwidth allocated to those forwarding classes to other forwarding classes. By default, the active forwarding classes (best-effort, network-control, and mcast) share the unused bandwidth assigned to the fcoe and no-loss forwarding classes.

Table 582: Components of the Hierarchical Port Scheduling (ETS) Configuration Topology

Property	Settings
Hardware	QFX3500 switch

Table 582: Components of the Hierarchical Port Scheduling (ETS) Configuration Topology (*continued*)

Property	Settings
Mapping of forwarding classes (priorities) to queues	<p>best-effort to queue 0</p> <p>be2 to queue 1</p> <p>fcoe (Fibre Channel over Ethernet) to queue 3</p> <p>no-loss to queue 4</p> <p>hpc (high-performance computing) to queue 5</p> <p>network-control to queue 7</p> <p>NOTE: On switches that do not support the ELS CLI, if you are using Junos OS Release 12.2 or later, use the default forwarding-class-to-queue mapping for the lossless fcoe and no-loss forwarding classes. If you explicitly configure the default lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best-effort) traffic and does <i>not</i> receive lossless treatment.</p> <p>On switches that do not support the ELS CLI, in Junos OS Release 12.3 and later, you can include the <i>no-loss</i> packet drop attribute in the explicit forwarding class configuration to configure a lossless forwarding class.</p>
Forwarding class sets (priority groups)	<p>best-effort-pg: contains forwarding classes best-effort, be2, and network control</p> <p>guar-delivery-pg: contains forwarding classes fcoe and no-loss</p> <p>hpc-pg: contains forwarding class hpc</p>
Behavior aggregate classifier (maps forwarding classes and loss priorities to incoming packets by IEEE 802.1 code point)	<p>Name—hsclassifier1</p> <p>Code point mapping:</p> <ul style="list-style-type: none"> • 000 to forwarding class best-effort and loss priority low • 001 to forwarding class be2 and loss priority high • 011 to forwarding class fcoe and loss priority low • 100 to forwarding class no-loss and loss priority low • 101 to forwarding class hpc and loss priority low • 110 to forwarding class network-control and loss priority low
PFC	<p>Congestion notification profile name—gd-cnp</p> <p>PFC enabled on code points: 011 (fcoe priority), 010 (no-loss priority)</p>
Drop profiles	<p>dp-be-low: drop start point 25, drop end point 50, maximum drop rate 80</p> <p>NOTE: The fcoe and no-loss priorities (queues) do not use drop profiles because they are lossless traffic classes.</p> <p>dp-be-high: drop start point 10, drop end point 40, maximum drop rate 100</p> <p>dp-hpc: drop start point 75, drop end point 90, maximum drop rate 75</p> <p>dp-nc: drop start point 80, drop end point 100, maximum drop rate 100</p>

Table 582: Components of the Hierarchical Port Scheduling (ETS) Configuration Topology (continued)

Property	Settings
Queue schedulers	<p>be-sched: minimum bandwidth 3g, maximum bandwidth 100%, priority low, drop profiles dp-be-low and dp-be-high</p> <p>fcoe-sched: minimum bandwidth 2.5g, maximum bandwidth 100%, priority low</p> <p>hpc-sched: minimum bandwidth 2g, maximum bandwidth 100%, priority low, drop profile dp-hpc</p> <p>nc-sched: minimum bandwidth 500m, maximum bandwidth 100%, priority low, drop profile dp-nc</p> <p>nl-sched: minimum bandwidth 2g, maximum bandwidth 100%, priority low</p>
Forwarding class-to-scheduler mapping	<p>Scheduler map be-map: Forwarding class best-effort, scheduler be-sched Forwarding class be2, scheduler be-sched Forwarding class network-control, scheduler nc-sched</p> <p>Scheduler map gd-map: Forwarding class fcoe, scheduler fcoe-sched Forwarding class no-loss, scheduler nl-sched</p> <p>Scheduler map hpc-map: Forwarding class hpc, scheduler hpc-sched</p>
Traffic control profiles	<p>be-tcp: scheduler map be-map, minimum bandwidth 3.5g, maximum bandwidth 100%</p> <p>gd-tcp: scheduler map gd-map, minimum bandwidth 4.5g, maximum bandwidth 100%</p> <p>hpc-tcp: scheduler map hpc-map, minimum bandwidth 2g, maximum bandwidth 100%</p>
Interfaces	<p>This example configures hierarchical port scheduling on interfaces xe-0/0/20 and xe-0/0/21. Because traffic is bidirectional, you apply the ingress and egress configuration components to both interfaces:</p> <ul style="list-style-type: none"> • Classifier Name—hsclassifier1 • Forwarding class sets—best-effort-pg, guar-deliver-pg, hpc-pg • Congestion notification profile—gd-cnp

Figure 220 shows a block diagram of the configuration components and the configuration flow of the CLI statements used in the example. You can perform the configuration steps in a different sequence if you want.

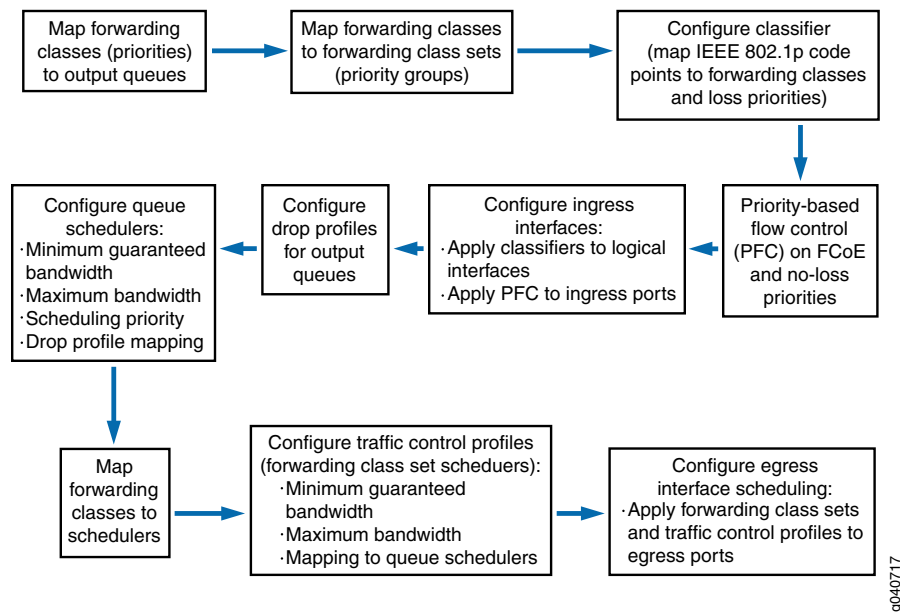
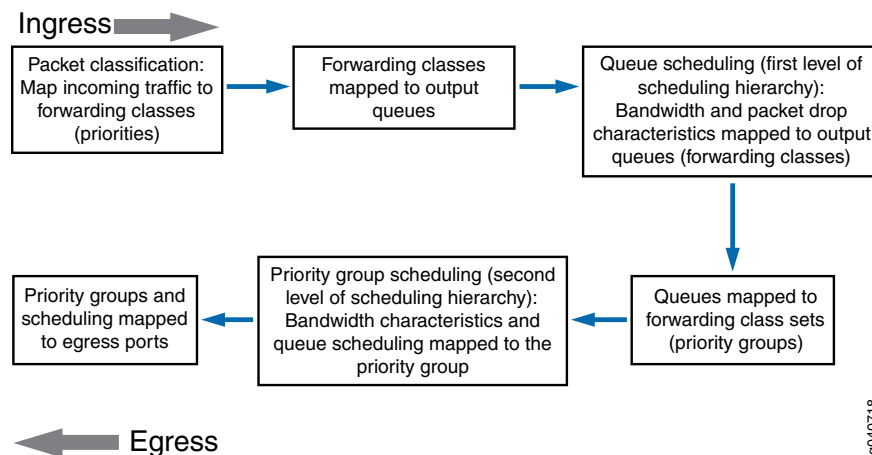
Figure 227: Hierarchical Port Scheduling Components Block Diagram

Figure 221 shows a block diagram of the hierarchical scheduling packet flow from ingress to egress.

Figure 228: Hierarchical Port Scheduling Packet Flow Block Diagram

Configuration

CLI Quick Configuration

To quickly configure hierarchical port scheduling on systems that support lossless transport, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the [edit class-of-service] hierarchy level:

```
[edit class-of-service]
set forwarding-classes class best-effort queue-num 0
set forwarding-classes class be2 queue-num 1
set forwarding-classes class hpc queue-num 5
set forwarding-classes class network-control queue-num 7
```

```

set forwarding-class-sets best-effort-pg class best-effort
set forwarding-class-sets best-effort-pg class be2
set forwarding-class-sets best-effort-pg class network-control
set forwarding-class-sets guar-delivery-pg class fcoe
set forwarding-class-sets guar-delivery-pg class no-loss
set forwarding-class-sets hpc-pg class hpc
set classifiers ieee-802.1 hsclassifier1 forwarding-class best-effort loss-priority low code-points
000
set classifiers ieee-802.1 hsclassifier1 forwarding-class be2 loss-priority high code-points 001
set classifiers ieee-802.1 hsclassifier1 forwarding-class fcoe loss-priority low code-points 011
set classifiers ieee-802.1 hsclassifier1 forwarding-class no-loss loss-priority low code-points 100
set classifiers ieee-802.1 hsclassifier1 forwarding-class hpc loss-priority low code-points 101
set classifiers ieee-802.1 hsclassifier1 forwarding-class network-control loss-priority low
code-points 110
set congestion-notification-profile gd-cnp input ieee-802.1 code-point 011 pfc
set congestion-notification-profile gd-cnp input ieee-802.1 code-point 100 pfc
set interfaces xe-0/0/20 unit 0 classifiers ieee-802.1 hsclassifier1
set interfaces xe-0/0/21 unit 0 classifiers ieee-802.1 hsclassifier1
set interfaces xe-0/0/20 congestion-notification-profile gd-cnp
set interfaces xe-0/0/21 congestion-notification-profile gd-cnp
set drop-profiles dp-be-low interpolate fill-level 25 fill-level 50 drop-probability 0 drop-probability
80
set drop-profiles dp-be-high interpolate fill-level 10 fill-level 40 drop-probability 0 drop-probability
100
set drop-profiles dp-nc interpolate fill-level 80 fill-level 100 drop-probability 0 drop-probability
100
set drop-profiles dp-hpc interpolate fill-level 75 fill-level 90 drop-probability 0 drop-probability
75
set schedulers be-sched priority low transmit-rate 3g
set schedulers be-sched shaping-rate percent 100
set schedulers be-sched drop-profile-map loss-priority low protocol any drop-profile dp-be-low
set schedulers be-sched drop-profile-map loss-priority high protocol any drop-profile dp-be-high
set schedulers fcoe-sched priority low transmit-rate 2500m
set schedulers fcoe-sched shaping-rate percent 100
set schedulers hpc-sched priority low transmit-rate 2g
set schedulers hpc-sched shaping-rate percent 100
set schedulers hpc-sched drop-profile-map loss-priority low protocol any drop-profile dp-hpc
set schedulers nc-sched priority low transmit-rate 500m
set schedulers nc-sched shaping-rate percent 100
set schedulers nc-sched drop-profile-map loss-priority low protocol any drop-profile dp-nc
set schedulers nl-sched priority low transmit-rate 2g
set schedulers nl-sched shaping-rate percent 100
set scheduler-maps be-map forwarding-class best-effort scheduler be-sched
set scheduler-maps be-map forwarding-class be2 scheduler be-sched
set scheduler-maps be-map forwarding-class network-control scheduler nc-sched
set scheduler-maps gd-map forwarding-class fcoe scheduler fcoe-sched
set scheduler-maps gd-map forwarding-class no-loss scheduler nl-sched
set scheduler-maps hpc-map forwarding-class hpc scheduler hpc-sched
set traffic-control-profiles be-tcp scheduler-map be-map guaranteed-rate 3500m
set traffic-control-profiles be-tcp shaping-rate percent 100
set traffic-control-profiles gd-tcp scheduler-map gd-map guaranteed-rate 4500m
set traffic-control-profiles gd-tcp shaping-rate percent 100
set traffic-control-profiles hpc-tcp scheduler-map hpc-map guaranteed-rate 2g
set traffic-control-profiles hpc-tcp shaping-rate percent 100
set interfaces xe-0/0/20 forwarding-class-set best-effort-pg output-traffic-control-profile be-tcp
set interfaces xe-0/0/20 forwarding-class-set guar-delivery-pg output-traffic-control-profile
gd-tcp
set interfaces xe-0/0/20 forwarding-class-set hpc-pg output-traffic-control-profile hpc-tcp
set interfaces xe-0/0/21 forwarding-class-set best-effort-pg output-traffic-control-profile be-tcp

```

```
set interfaces xe-0/0/21 forwarding-class-set guar-delivery-pg output-traffic-control-profile
gd-tcp
set interfaces xe-0/0/21 forwarding-class-set hpc-pg output-traffic-control-profile hpc-tcp
```

OCX Series Switches

Because OCX Series switches do not support lossless transport, the following subset of the configuration eliminates the lossless configuration elements and provides hierarchical port scheduling for the best-effort, be2, hpc, and network-control forwarding classes. In addition, on OCX Series switches, you would probably use DSCP classifiers and code points instead of IEEE classifiers and code points. To quickly configure hierarchical port scheduling on an OCX Series switch, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the [edit class-of-service] hierarchy level:

```
[edit class-of-service]
set forwarding-classes class best-effort queue-num 0
set forwarding-classes class be2 queue-num 1
set forwarding-classes class hpc queue-num 5
set forwarding-classes class network-control queue-num 7
set forwarding-class-sets best-effort-pg class best-effort
set forwarding-class-sets best-effort-pg class be2
set forwarding-class-sets best-effort-pg class network-control
set forwarding-class-sets hpc-pg class hpc
set classifiers ieee-802.1 hsclassifier1 forwarding-class best-effort loss-priority low code-points
000
set classifiers ieee-802.1 hsclassifier1 forwarding-class be2 loss-priority high code-points 001
set classifiers ieee-802.1 hsclassifier1 forwarding-class hpc loss-priority low code-points 101
set classifiers ieee-802.1 hsclassifier1 forwarding-class network-control loss-priority low
code-points 110
set interfaces xe-0/0/20 unit 0 classifiers ieee-802.1 hsclassifier1
set interfaces xe-0/0/21 unit 0 classifiers ieee-802.1 hsclassifier1
set drop-profiles dp-be-low interpolate fill-level 25 fill-level 50 drop-probability 0 drop-probability
80
set drop-profiles dp-be-high interpolate fill-level 10 fill-level 40 drop-probability 0 drop-probability
100
set drop-profiles dp-nc interpolate fill-level 80 fill-level 100 drop-probability 0 drop-probability
100
set drop-profiles dp-hpc interpolate fill-level 75 fill-level 90 drop-probability 0 drop-probability
75
set schedulers be-sched priority low transmit-rate 3g
set schedulers be-sched shaping-rate percent 100
set schedulers be-sched drop-profile-map loss-priority low protocol any drop-profile dp-be-low
set schedulers be-sched drop-profile-map loss-priority high protocol any drop-profile dp-be-high
set schedulers hpc-sched priority low transmit-rate 2g
set schedulers hpc-sched shaping-rate percent 100
set schedulers hpc-sched drop-profile-map loss-priority low protocol any drop-profile dp-hpc
set schedulers nc-sched priority low transmit-rate 500m
set schedulers nc-sched shaping-rate percent 100
set schedulers nc-sched drop-profile-map loss-priority low protocol any drop-profile dp-nc
set scheduler-maps be-map forwarding-class best-effort scheduler be-sched
set scheduler-maps be-map forwarding-class be2 scheduler be-sched
set scheduler-maps be-map forwarding-class network-control scheduler nc-sched
set scheduler-maps hpc-map forwarding-class hpc scheduler hpc-sched
set traffic-control-profiles be-tcp scheduler-map be-map guaranteed-rate 3500m
set traffic-control-profiles be-tcp shaping-rate percent 100
set traffic-control-profiles hpc-tcp scheduler-map hpc-map guaranteed-rate 2g
set traffic-control-profiles hpc-tcp shaping-rate percent 100
```

```

set interfaces xe-0/0/20 forwarding-class-set best-effort-pg output-traffic-control-profile be-tcp
set interfaces xe-0/0/20 forwarding-class-set hpc-pg output-traffic-control-profile hpc-tcp
set interfaces xe-0/0/21 forwarding-class-set best-effort-pg output-traffic-control-profile be-tcp
set interfaces xe-0/0/21 forwarding-class-set hpc-pg output-traffic-control-profile hpc-tcp

```

Step-by-Step Procedure

To perform a step-by-step configuration of the forwarding classes (priorities), forwarding class sets (priority groups), classifiers, queue schedulers, PFC, traffic control profiles, and interfaces to set up hierarchical port scheduling (ETS):

1. Configure the forwarding classes (priorities) and map them to unicast output queues (do not explicitly map the **fcoe** and **no-loss** forwarding classes to output queues; use the default configuration):

```

[edit class-of-service]
user@switch# set forwarding-classes class best-effort queue-num 0
user@switch# set forwarding-classes class be2 queue-num 1
user@switch# set forwarding-classes class hpc queue-num 5
user@switch# set forwarding-classes class network-control queue-num 7

```

2. Configure forwarding class sets (priority groups) to group forwarding classes (priorities) that require similar CoS treatment:

```

[edit class-of-service]
user@switch# set forwarding-class-sets best-effort-pg class best-effort
user@switch# set forwarding-class-sets best-effort-pg class be2
user@switch# set forwarding-class-sets best-effort-pg class network-control
user@switch# set forwarding-class-sets guar-delivery-pg class fcoe
user@switch# set forwarding-class-sets guar-delivery-pg class no-loss
user@switch# set forwarding-class-sets hpc-pg class hpc

```



NOTE: On OCX Series switches, you would not configure the **guar-delivery-pg** forwarding class set for lossless traffic.

3. Configure a classifier to set the loss priority and IEEE 802.1 code points assigned to each forwarding class at the ingress:

```

[edit class-of-service]
user@switch# set classifiers ieee-802.1 hsclassifier1 forwarding-class best-effort
loss-priority low code-points 000
user@switch# set classifiers ieee-802.1 hsclassifier1 forwarding-class be2 loss-priority
high code-points 001
user@switch# set classifiers ieee-802.1 hsclassifier1 forwarding-class fcoe loss-priority
low code-points 011
user@switch# set classifiers ieee-802.1 hsclassifier1 forwarding-class no-loss loss-priority
low code-points 100
user@switch# set classifiers ieee-802.1 hsclassifier1 forwarding-class hpc loss-priority low
code-points 101
user@switch# set classifiers ieee-802.1 hsclassifier1 forwarding-class network-control
loss-priority low code-points 110

```



NOTE: On OCX Series switches, you would not configure the **fcoe** and **no-loss** portions of the classifier.

4. Configure a congestion notification profile to enable PFC on the FCoE and no-loss queue IEEE 802.1 code points:

```
[edit class-of-service]
user@switch# set congestion-notification-profile gd-cnp input ieee-802.1 code-point 011
pfc
user@switch# set congestion-notification-profile gd-cnp input ieee-802.1 code-point 100
pfc
```



NOTE: This step does not apply to OCX Series switches, which do not support PFC.

5. Assign the classifier to the interfaces:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 unit 0 classifiers ieee-802.1 hsclassifier1
user@switch# set interfaces xe-0/0/21 unit 0 classifiers ieee-802.1 hsclassifier1
```

6. Apply the PFC configuration to the interfaces:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 congestion-notification-profile gd-cnp
user@switch# set interfaces xe-0/0/21 congestion-notification-profile gd-cnp
```



NOTE: This step does not apply to OCX Series switches, which do not support PFC.

7. Configure the drop profile for the best-effort low loss-priority queue:

```
[edit class-of-service]
user@switch# set drop-profiles dp-be-low interpolate fill-level 25 fill-level 50
drop-probability 0 drop-probability 80
```

8. Configure the drop profile for the best-effort high loss-priority queue:

```
[edit class-of-service]
user@switch# set drop-profiles dp-be-high interpolate fill-level 10 fill-level 40
drop-probability 0 drop-probability 100
```

9. Configure the drop profile for the network-control queue:

```
[edit class-of-service]
user@switch# set drop-profiles dp-nc interpolate fill-level 80 fill-level 100 drop-probability
0 drop-probability 100
```

10. Configure the drop profile for the high-performance computing queue:

```
[edit class-of-service]
user@switch# set drop-profiles dp-hpc interpolate fill-level 75 fill-level 90 drop-probability
0 drop-probability 75
```

11. Define the minimum guaranteed bandwidth, priority, maximum bandwidth, and drop profiles for the best-effort queue:

```
[edit class-of-service]
user@switch# set schedulers be-sched priority low transmit-rate 3g
user@switch# set schedulers be-sched shaping-rate percent 100
```



```

user@switch# set schedulers be-sched drop-profile-map loss-priority low protocol any
drop-profile dp-be-low
user@switch# set schedulers be-sched drop-profile-map loss-priority high protocol any
drop-profile dp-be-high

```

12. Define the minimum guaranteed bandwidth, priority, and maximum bandwidth for the FCoE queue:

```

[edit class-of-service]
user@switch# set schedulers fcoe-sched priority low transmit-rate 2500m
user@switch# set schedulers fcoe-sched shaping-rate percent 100

```



NOTE: This step does not apply to OCX Series switches, which do not support lossless transport.

13. Define the minimum guaranteed bandwidth, priority, maximum bandwidth, and drop profile for the high-performance computing queue:

```

[edit class-of-service]
user@switch# set schedulers hpc-sched priority low transmit-rate 2g
user@switch# set schedulers hpc-sched shaping-rate percent 100
user@switch# set schedulers hpc-sched drop-profile-map loss-priority low protocol any
drop-profile dp-hpc

```

14. Define the minimum guaranteed bandwidth, priority, maximum bandwidth, and drop profile for the network-control queue:

```

[edit class-of-service]
user@switch# set schedulers nc-sched priority low transmit-rate 500m
user@switch# set schedulers nc-sched shaping-rate percent 100
user@switch# set schedulers nc-sched drop-profile-map loss-priority low protocol any
drop-profile dp-nc

```

15. Define the minimum guaranteed bandwidth, priority, and maximum bandwidth for the no-loss queue:

```

[edit class-of-service]
user@switch# set schedulers nl-sched priority low transmit-rate 2g
user@switch# set schedulers nl-sched shaping-rate percent 100

```



NOTE: This step does not apply to OCX Series switches, which do not support lossless transport.

16. Map the schedulers to the appropriate forwarding classes (queues):

```

[edit class-of-service]
user@switch# set scheduler-maps be-map forwarding-class best-effort scheduler be-sched
user@switch# set scheduler-maps be-map forwarding-class be2 scheduler be-sched
user@switch# set scheduler-maps be-map forwarding-class network-control scheduler
nc-sched
user@switch# set scheduler-maps gd-map forwarding-class fcoe scheduler fcoe-sched
user@switch# set scheduler-maps gd-map forwarding-class no-loss scheduler nl-sched
user@switch# set scheduler-maps hpc-map forwarding-class hpc scheduler hpc-sched

```



NOTE: On OCX Series switches, because lossless transport is not supported, you would not configure the `gd-map` scheduler map.

17. Define the traffic control profile for the best-effort priority group (queue scheduler to mapping, minimum guaranteed bandwidth, and maximum bandwidth):

```
[edit class-of-service]
user@switch# set traffic-control-profiles be-tcp scheduler-map be-map guaranteed-rate
3500m
user@switch# set traffic-control-profiles be-tcp shaping-rate percent 100
```

18. Define the traffic control profile for the guaranteed delivery priority group (queue to scheduler mapping, minimum guaranteed bandwidth, and maximum bandwidth):

```
[edit class-of-service]
user@switch# set traffic-control-profiles gd-tcp scheduler-map gd-map guaranteed-rate
4500m
user@switch# set traffic-control-profiles gd-tcp shaping-rate percent 100
```



NOTE: This step does not apply to OCX Series switches, which do not support lossless transport.

19. Define the traffic control profile for the high-performance computing priority group (queue to scheduler mapping, minimum guaranteed bandwidth, and maximum bandwidth):

```
[edit class-of-service]
user@switch# set traffic-control-profiles hpc-tcp scheduler-map hpc-map guaranteed-rate
2g
user@switch# set traffic-control-profiles hpc-tcp shaping-rate percent 100
```

20. Apply the three priority groups (forwarding class sets) and the appropriate traffic control profiles to the egress ports:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 forwarding-class-set best-effort-pg
output-traffic-control-profile be-tcp
user@switch# set interfaces xe-0/0/20 forwarding-class-set guar-delivery-pg
output-traffic-control-profile gd-tcp
user@switch# set interfaces xe-0/0/20 forwarding-class-set hpc-pg
output-traffic-control-profile hpc-tcp
user@switch# set interfaces xe-0/0/21 forwarding-class-set best-effort-pg
output-traffic-control-profile be-tcp
user@switch# set interfaces xe-0/0/21 forwarding-class-set guar-delivery-pg
output-traffic-control-profile gd-tcp
user@switch# set interfaces xe-0/0/21 forwarding-class-set hpc-pg
output-traffic-control-profile hpc-tcp
```



NOTE: Because OCX Series switches do not support lossless transport, on OCX Series switches, you would not apply the `guar-deliver-pg` forwarding class set and the `gd-tcp` traffic control profile to interfaces.

Results

Display the results of the configuration (the system shows only the explicitly configured parameters; it does not show default parameters such as the **fcoe** and **no-loss** lossless forwarding classes). On OCX Series switches, you would not see the lossless configuration components in the output:

```
user@switch> show configuration class-of-service
classifiers {
  ieee-802.1 hsclassifier1 {
    forwarding-class best-effort {
      loss-priority low code-points 000;
    }
    forwarding-class be2 {
      loss-priority high code-points 001;
    }
    forwarding-class fcoe {
      loss-priority low code-points 011;
    }
    forwarding-class no-loss {
      loss-priority low code-points 100;
    }
    forwarding-class hpc {
      loss-priority low code-points 101;
    }
    forwarding-class network-control {
      loss-priority low code-points 110;
    }
  }
}
drop-profiles {
  dp-be-low {
    interpolate {
      fill-level [ 25 50 ];
      drop-probability [ 0 80 ];
    }
  }
  dp-be-high {
    interpolate {
      fill-level [ 10 40 ];
      drop-probability [ 0 100 ];
    }
  }
  dp-hpc {
    interpolate {
      fill-level [ 75 90 ];
      drop-probability [ 0 75 ];
    }
  }
  dp-nc {
    interpolate {
      fill-level [ 80 100 ];
      drop-probability [ 0 100 ];
    }
  }
}
```

```
forwarding-classes {
  class best-effort queue-num 0;
  class be2 queue-num 1;
  class hpc queue-num 5;
  class network-control queue-num 7;
}
traffic-control-profiles {
  be-tcp {
    scheduler-map be-map;
    shaping-rate percent 100;
    guaranteed-rate 3500000000;
  }
  gd-tcp {
    scheduler-map gd-map;
    shaping-rate percent 100;
    guaranteed-rate 4500000000;
  }
  hpc-tcp {
    scheduler-map hpc-map;
    shaping-rate percent 100;
    guaranteed-rate 2g;
  }
}
forwarding-class-sets {
  guar-delivery-pg {
    class fcoe;
    class no-loss;
  }
  best-effort-pg {
    class best-effort;
    class be2;
    class network-control;
  }
  hpc-pg {
    class hpc;
  }
}
congestion-notification-profile {
  gd-cnp {
    input {
      ieee-802.1 {
        code-point 011 {
          pfc;
        }
        code-point 100 {
          pfc;
        }
      }
    }
  }
}
interfaces {
  xe-0/0/20 {
    forwarding-class-set {
      best-effort-pg {
        output-traffic-control-profile be-tcp;
      }
    }
  }
}
```

```

    }
    guar-delivery-pg {
        output-traffic-control-profile gd-tcp;
    }
    hpc-pg {
        output-traffic-control-profile hpc-tcp;
    }
}
congestion-notification-profile gd-cnp;
unit 0 {
    classifiers {
        ieee-802.1 hsclassifier1;
    }
}
}
xe-0/0/21 {
    forwarding-class-set {
        best-effort-pg {
            output-traffic-control-profile be-tcp;
        }
        guar-delivery-pg {
            output-traffic-control-profile gd-tcp;
        }
        hpc-pg {
            output-traffic-control-profile hpc-tcp;
        }
    }
    congestion-notification-profile gd-cnp;
    unit 0 {
        classifiers {
            ieee-802.1 hsclassifier1;
        }
    }
}
}
scheduler-maps {
    be-map {
        forwarding-class best-effort scheduler be-sched;
        forwarding-class network-control scheduler nc-sched;
        forwarding-class be2 scheduler be-sched;
    }
    gd-map {
        forwarding-class fcoe scheduler fcoe-sched;
        forwarding-class no-loss scheduler nl-sched;
    }
    hpc-map {
        forwarding-class hpc scheduler hpc-sched;
    }
}
schedulers {
    be-sched {
        transmit-rate 3g;
        shaping-rate percent 100;
        priority low;
        drop-profile-map loss-priority low protocol any drop-profile dp-be-low;
        drop-profile-map loss-priority high protocol any drop-profile dp-be-high;
    }
}

```

```
}
fcoe-sched {
    transmit-rate 2500000000;
    shaping-rate percent 100;
    priority low;
}
hpc-sched {
    transmit-rate 2g;
    shaping-rate percent 100;
    priority low;
    drop-profile-map loss-priority low protocol any drop-profile dp-hpc;
}
nc-sched {
    transmit-rate 500m;
    shaping-rate percent 100;
    priority low;
    drop-profile-map loss-priority low protocol any drop-profile dp-nc;
}
nl-sched {
    transmit-rate 2g;
    shaping-rate percent 100;
    priority low;
}
}
```



TIP: To quickly configure the interfaces, issue the `load merge terminal` command, and then copy the hierarchy and paste it into the switch terminal window.

Verification



NOTE: The verification output is based on the full example configuration. On OCX Series switches, you do not see lossless configuration components in the output. Comments about lossless configuration components do not apply to OCX Series switches.

To verify that you created the hierarchical port scheduling components and they are operating properly, perform these tasks:

- [Verifying the Forwarding Classes \(Priorities\) on page 6863](#)
- [Verifying the Forwarding Class Sets \(Priority Groups\) on page 6863](#)
- [Verifying the Classifier on page 6864](#)
- [Verifying Priority-Based Flow Control on page 6864](#)
- [Verifying the Output Queue Schedulers on page 6865](#)
- [Verifying the Drop Profiles on page 6868](#)

- [Verifying the Priority Group Output Schedulers \(Traffic Control Profiles\) on page 6869](#)
- [Verifying the Interface Configuration on page 6870](#)

Verifying the Forwarding Classes (Priorities)

Purpose Verify that you created the forwarding classes and mapped them to the correct queues. (The system shows only the explicitly configured forwarding classes. It does not show default forwarding classes such as **fcoe** and **no-loss**.)

Action List the forwarding classes using the operational mode command **show class-of-service forwarding-class**:

```
user@switch> show class-of-service forwarding-class
```

Forwarding class	ID	Queue	Policing priority	No-Loss
best-effort	0	0	normal	Disabled
be2	1	3	normal	Disabled
hpc	2	4	normal	Disabled
network-control	3	7	normal	Disabled
mcast	8	8	normal	Disabled

Meaning The **show class-of-service forwarding-class** command lists all of the configured forwarding classes, the internal identification number of each forwarding class, the queues that are mapped to the forwarding classes, the policing priority, and whether the forwarding class is lossless (no-loss packet drop attribute enabled) or lossy forwarding class (no-loss packet drop attribute disabled). The command output shows that:

- Forwarding class **best-effort** maps to queue **0** and is lossy
- Forwarding class **be2** maps to queue **1** and is lossy
- Forwarding class **hpc** maps to queue **5** and is lossy
- Forwarding class **network-control** maps to queue **7** and is lossy

In addition, the command lists the default multicast (multidestination) forwarding class and the default queue to which it is mapped.

Verifying the Forwarding Class Sets (Priority Groups)

Purpose Verify that you created the priority groups and that the correct priorities (forwarding classes) belong to the appropriate priority group.

Action List the forwarding class sets using the operational mode command **show class-of-service forwarding-class-set**:

```
user@switch> show class-of-service forwarding-class-set
```

```
Forwarding class set: best-effort-pg, Type: normal-type, Forwarding class set
index: 19907
```

Forwarding class	Index
best-effort	0

be2	1
network-control	5

Forwarding class set: guar-delivery-pg, Type: normal-type, Forwarding class set index: 43700

Forwarding class	Index
fcoe	2
no-loss	3

Forwarding class set: hpc-pg, Type: normal-type, Forwarding class set index: 60758

Forwarding class	Index
hpc	4

Meaning The `show class-of-service forwarding-class-set` command lists all of the configured forwarding class sets (priority groups), the forwarding classes (priorities) that belong to each priority group, and the internal index number of each priority group. The command output shows that:

- The forwarding class set **best-effort-pg** includes the forwarding classes **best-effort**, **be2**, and **network-control**.
- The forwarding class set **guar-delivery-pg** includes the forwarding classes **fcoe** and **no-loss**.
- The forwarding class set **hpc-pg** includes the forwarding class **hpc**.

Verifying the Classifier

Purpose Verify that the classifier maps forwarding classes to the correct IEEE 802.1p code points and packet loss priorities.

Action List the classifier configured for hierarchical port scheduling using the operational mode command `show class-of-service classifier name hsclassifier1`:

```
user@switch> show class-of-service classifier name hsclassifier1
Classifier: hsclassifier1, Code point type: ieee-802.1, Index: 43607
Code point      Forwarding class      Loss priority
000             best-effort           low
001             be2                   high
011             fcoe                  low
100             no-loss               low
101             hpc                   low
110             network-control       low
```

Meaning The `show class-of-service classifier name hsclassifier1` command lists all of the IEEE 802.1p code points and the loss priorities mapped to all of the forwarding classes in the classifier. The command output shows that the forwarding classes **best-effort**, **be2**, **no-loss**, **fcoe**, **hpc**, and **network-control** have been created and mapped to IEEE 802.1p code points and loss priorities.

Verifying Priority-Based Flow Control

Purpose Verify that PFC is enabled on the correct priorities for lossless transport.

Action List the congestion notification profiles using the operational mode command **show class-of-service congestion-notification**:

```
user@switch> show class-of-service congestion-notification
```

```
Type: Input, Name: gd-cnp, Index: 51687
```

```
Cable Length: 100 m
```

Priority	PFC	MRU
000	Disabled	
001	Disabled	
010	Disabled	
011	Enabled	2500
100	Enabled	2500
101	Disabled	
110	Disabled	
111	Disabled	

```
Type: Output
```

Priority	Flow-Control-Queues
000	0
001	0
010	1
011	2
100	3
101	4
110	5
111	6
	7

Meaning The **show class-of-service congestion-notification** command lists all of the congestion notification profiles and the IEEE 802.1p code points with PFC enabled. The command output shows that PFC is enabled for code points **011** (**fcoe** priority and queue) and **100** (**no-loss** priority and queue) for the **gd-cnp** congestion notification profile.

The command also shows the default cable length (100 meters), the default maximum receive unit (2500 bytes), and the default mapping of priorities to output queues because this example does not include configuring these options.

Verifying the Output Queue Schedulers

Purpose Verify that you created the output queue schedulers with the correct bandwidth parameters and priorities, mapped to the correct queues, and mapped to the correct drop profiles.

Action List the scheduler maps using the operational mode command **show class-of-service scheduler-map**:

```
user@switch> show class-of-service scheduler-map
```

```
Scheduler map: be-map, Index: 64023
```

```
Scheduler: be-sched, Forwarding class: best-effort, Index: 13005
Transmit rate: 3000000000 bps, Rate Limit: none, Buffer size: remainder,
```

Buffer Limit: none, Priority: low
 Excess Priority: unspecified
 Shaping rate: 100 percent,
 drop-profile-map-set-type: mark
 Drop profiles:

Loss priority	Protocol	Index	Name
Low	any	55387	dp-be-low
Medium high	any	1	<default-drop-profile>
High	any	4369	dp-be-high

Scheduler: be-sched, Forwarding class: be2, Index: 13005
 Transmit rate: 3000000000 bps, Rate Limit: none, Buffer size: remainder,
 Buffer Limit: none, Priority: low
 Excess Priority: unspecified
 Shaping rate: 100 percent,
 drop-profile-map-set-type: mark
 Drop profiles:

Loss priority	Protocol	Index	Name
Low	any	55387	dp-be-low
Medium high	any	1	<default-drop-profile>
High	any	4369	dp-be-high

Scheduler: nc-sched, Forwarding class: network-control, Index: 45740
 Transmit rate: 5000000000 bps, Rate Limit: none, Buffer size: remainder,
 Buffer Limit: none, Priority: low
 Excess Priority: unspecified
 Shaping rate: 100 percent,
 drop-profile-map-set-type: mark
 Drop profiles:

Loss priority	Protocol	Index	Name
Low	any	44207	dp-nc
Medium high	any	1	<default-drop-profile>
High	any	1	<default-drop-profile>

Scheduler map: gd-map, Index: 61447

Scheduler: fcoe-sched, Forwarding class: fcoe, Index: 37289
 Transmit rate: 25000000000 bps, Rate Limit: none, Buffer size: remainder,
 Buffer Limit: none, Priority: low
 Excess Priority: unspecified
 Shaping rate: 100 percent,
 drop-profile-map-set-type: mark
 Drop profiles:

Loss priority	Protocol	Index	Name
Low	any	44207	<default-drop-profile>
Medium high	any	1	<default-drop-profile>
High	any	1	<default-drop-profile>

Scheduler: nl-sched, Forwarding class: no-loss, Index: 29359
 Transmit rate: 20000000000 bps, Rate Limit: none, Buffer size: remainder,
 Buffer Limit: none, Priority: low
 Excess Priority: unspecified
 Shaping rate: 100 percent,
 drop-profile-map-set-type: mark
 Drop profiles:

Loss priority	Protocol	Index	Name
Low	any	44207	<default-drop-profile>
Medium high	any	1	<default-drop-profile>
High	any	1	<default-drop-profile>

Scheduler map: hpc-map, Index: 56941

```
Scheduler: hpc-sched, Forwarding class: hpc, Index: 55900
Transmit rate: 2000000000 bps, Rate Limit: none, Buffer size: remainder,
Buffer Limit: none, Priority: low
Excess Priority: unspecified
Shaping rate: 100 percent,
drop-profile-map-set-type: mark
Drop profiles:
  Loss priority  Protocol  Index  Name
  Low           any       57716  dp-hpc
  Medium high   any       1      <default-drop-profile>
  High          any       1      <default-drop-profile>
```

Meaning The **show class-of-service scheduler-map** command lists all of the configured scheduler maps. For each scheduler map, the command output includes:

- The name of the scheduler map (**scheduler-map** field)
- The name of the scheduler (**scheduler** field)
- The forwarding classes mapped to the scheduler (**forwarding-class** field)
- The minimum guaranteed queue bandwidth (**transmit-rate** field)
- The scheduling priority (**priority** field)
- The maximum bandwidth in the priority group the queue can consume (**shaping-rate** field)
- The drop profile loss priority (**loss priority** field) for each drop profile name (**name** field)

The command output shows that:

- The scheduler map **be-map** was created and has these properties:
 - There are two schedulers, **be-sched** and **nc-sched**.
 - The scheduler **be-sched** has two forwarding classes, **best-effort** and **be2**.
 - Scheduler **be-sched** forwarding classes **best-effort** and **be2** share a minimum guaranteed bandwidth of **3,000,000,000 bps**, can consume a maximum of **100 percent** of the priority group bandwidth, and use the drop profile **dp-be-low** for low loss-priority traffic, the default drop profile for medium-high loss-priority traffic, and the drop profile **dp-be-high** for high loss-priority traffic.
 - The scheduler **nc-sched** has one forwarding class, **network-control**.
 - The **network-control** forwarding class has a minimum guaranteed bandwidth of **500,000,000 bps**, can consume a maximum of **100 percent** of the priority group bandwidth, and uses the drop profile **dp-nc** for low loss-priority traffic and the default drop profile for medium-high and high loss priority traffic.
- The scheduler map **gd-map** was created and has these properties:
 - There are two schedulers, **fcoe-sched** and **nl-sched**.
 - The scheduler **fcoe-sched** has one forwarding class, **fcoe**.

- The **fcoe** forwarding class has a minimum guaranteed bandwidth of **2,500,000,000 bps**, and can consume a maximum of **100 percent** of the priority group bandwidth.
- The scheduler **nl-sched** has one forwarding class, **no-loss**.
- The **no-loss** forwarding class has a minimum guaranteed bandwidth of **2,000,000,000 bps**, and can consume a maximum of **100 percent** of the priority group bandwidth.
- The scheduler map **hpc-map** was created and has these properties:
 - There is one scheduler, **hpc-sched**.
 - The scheduler **hpc-sched** has one forwarding class, **hpc**.
 - The **hpc** forwarding class has a minimum guaranteed bandwidth of **2,000,000,000 bps**, can consume a maximum of **100 percent** of the priority group bandwidth, and uses the drop profile **dp-hpc** for low loss-priority traffic and the default drop profile for medium-high and high loss-priority traffic.

Verifying the Drop Profiles

Purpose Verify that you created the drop profiles **dp-be-high**, **dp-be-low**, **dp-hpc**, and **dp-nc** with the correct fill levels and drop probabilities.

Action List the drop profiles using the operational mode command **show configuration class-of-service drop-profiles**:

```
user@switch> show configuration class-of-service drop-profiles
dp-be-low {
    interpolate {
        fill-level [ 25 50 ];
        drop-probability [ 0 80 ];
    }
}
dp-be-high {
    interpolate {
        fill-level [ 10 40 ];
        drop-probability [ 0 100 ];
    }
}
dp-hpc {
    interpolate {
        fill-level [ 75 90 ];
        drop-probability [ 0 75 ];
    }
}
dp-nc {
    interpolate {
        fill-level [ 80 100 ];
        drop-probability [ 0 100 ];
    }
}
```

Meaning The **show configuration class-of-service drop-profiles** command lists the drop profiles and their properties. The command output shows that there are four drop profiles configured, **dp-be-high**, **dp-be-low**, **dp-hpc**, and **dp-nc**. The output also shows that:

- For **dp-be-low**, the drop start point (the first fill level) is when the queue is 25 percent filled, the drop end point (the second fill level) occurs when the queue is 50 percent filled, and the drop probability at the drop end point is 80 percent.
- For **dp-be-high**, the drop start point (the first fill level) is when the queue is 10 percent filled, the drop end point (the second fill level) occurs when the queue is 40 percent filled, and the drop probability at the drop end point is 100 percent.
- For **dp-hpc**, the drop start point (the first fill level) is when the queue is 75 percent filled, the drop end point (the second fill level) occurs when the queue is 90 percent filled, and the drop probability at the drop end point is 75 percent.
- For **dp-nc**, the drop start point (the first fill level) is when the queue is 80 percent filled, the drop end point (the second fill level) occurs when the queue is 100 percent filled, and the drop probability at the drop end point is 100 percent.

Verifying the Priority Group Output Schedulers (Traffic Control Profiles)

Purpose	Verify that you created the traffic control profiles be-tcp , gd-tcp , and hpc-tcp with the correct bandwidth parameters and scheduler mapping.
Action	<p>List the traffic control profiles using the operational mode command show class-of-service traffic-control-profile:</p> <pre> user@switch> show class-of-service traffic-control-profile Traffic control profile: be-tcp, Index: 40535 Shaping rate: 100 percent Scheduler map: be-map Guaranteed rate: 3500000000 Traffic control profile: gd-tcp, Index: 37959 Shaping rate: 100 percent Scheduler map: gd-map Guaranteed rate: 4500000000 Traffic control profile: hpc-tcp, Index: 47661 Shaping rate: 100 percent Scheduler map: hpc-map Guaranteed rate: 2000000000 </pre>
Meaning	<p>The show class-of-service traffic-control-profile command lists all of the configured traffic control profiles. For each traffic control profile, the command output includes:</p> <ul style="list-style-type: none"> • The name of the traffic control profile (traffic-control-profile) • The maximum port bandwidth the priority group can consume (shaping-rate) • The scheduler map associated with the traffic control profile (scheduler-map) • The minimum guaranteed priority group port bandwidth (guaranteed-rate) <p>The command output shows that:</p> <ul style="list-style-type: none"> • The traffic control profile be-tcp can consume a maximum of 100 percent of the port bandwidth, is associated with the scheduler map be-map, and has a minimum guaranteed bandwidth of 3,500,000,000 bps.

- The traffic control profile **gd-tcp** can consume a maximum of **100 percent** of the port bandwidth, is associated with the scheduler map **gd-map**, and has a minimum guaranteed bandwidth of **4,500,000,000 bps**.
- The traffic control profile **hpc-tcp** can consume a maximum of **100 percent** of the port bandwidth, is associated with the scheduler map **hpc-map**, and has a minimum guaranteed bandwidth of **2,000,000,000 bps**.

Verifying the Interface Configuration

Purpose Verify that the classifier, the congestion notification profile, and the forwarding class sets are configured on interfaces **xe-0/0/20** and **xe-0/0/21**.

Action List the interfaces using the operational mode commands **show configuration class-of-service interfaces xe-0/0/20** and **show configuration class-of-service interfaces xe-0/0/21**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/20
forwarding-class-set {
    best-effort-gp {
        output-traffic-control-profile be-tcp;
    }
    guar-delivery-pg {
        output-traffic-control-profile gd-tcp;
    }
    hpc-pg {
        output-traffic-control-profile hpc-tcp;
    }
}
congestion-notification-profile gd_cnp;
unit 0 {
    classifiers {
        ieee-802.1 hsclassifier1;
    }
}
```

```
user@switch> show configuration class-of-service interfaces xe-0/0/21
forwarding-class-set {
    best-effort-gp {
        output-traffic-control-profile be-tcp;
    }
    guar-delivery-pg {
        output-traffic-control-profile gd-tcp;
    }
    hpc-pg {
        output-traffic-control-profile hpc-tcp;
    }
}
congestion-notification-profile gd_cnp;
unit 0 {
    classifiers {
        ieee-802.1 hsclassifier1;
    }
}
```

Meaning The `show configuration class-of-service interfaces interface-name` command shows that each interface includes the forwarding class sets **best-effort-pg**, **guar-delivery-pg**, and **hpc-pg**, congestion notification profile **gd-cnp**, and the IEEE 802.1p classifier **hsclassifier1**.

**Related
Documentation**

- *Defining CoS Unicast BA Classifiers (DSCP, DSCP IPv6, IEEE 802.1p)*
- *Benefits of Configuring CoS Hierarchical Port Scheduling*
- *Assigning CoS Components to Interfaces on page 6627*
- *Example: Configuring WRED Drop Profiles*
- *Example: Configuring Drop Profile Maps on page 6817*
- *Example: Configuring Forwarding Classes*
- *Example: Configuring Forwarding Class Sets on page 6697*
- *Example: Configuring Queue Schedulers*
- *Example: Configuring Queue Scheduling Priority on page 6753*
- *Example: Configuring Traffic Control Profiles (Priority Group Scheduling) on page 6762*
- *Example: Configuring Minimum Guaranteed Output Bandwidth on page 6798*
- *Example: Configuring Maximum Output Bandwidth on page 6805*
- *Configuring CoS PFC (Congestion Notification Profiles) on page 6905*
- *Overview of CoS Changes Introduced in Junos OS Release 12.2*
- *Understanding CoS Hierarchical Port Scheduling (ETS) on page 6765*
- *Understanding CoS Scheduling Behavior and Configuration Considerations on page 6718*
- *Understanding CoS Scheduling on QFabric System Node Device Fabric (fte) Ports*
- *Understanding Default CoS Scheduling on QFabric System Interconnect Devices (Junos OS Release 13.1 and Later Releases)*

Disabling the ETS Recommendation TLV

The enhanced transmission selection (ETS) Recommendation TLV communicates the ETS settings that the switch wants the connected peer interface to use. If the peer interface is “willing,” the peer interface changes its configuration to match the configuration in the ETS Recommendation TLV. By default, the switch interfaces send the ETS Recommendation TLV to the peer. The settings communicated are the egress ETS settings defined by configuring hierarchical scheduling on the interface.

We recommend that you use the same ETS settings on the connected peer that you use on the switch interface and that you leave the ETS Recommendation TLV enabled. However, on interfaces that use IEEE DCBX as the DCBX mode, if you want an asymmetric configuration between the switch interface and the connected peer, you can disable the ETS Recommendation TLV.



NOTE: Disabling the ETS Recommendation TLV on interfaces that use DCBX version 1.01 as the DCBX mode has no effect and does not change DCBX behavior.

If you disable the ETS Recommendation TLV, the switch still sends the ETS Configuration TLV to the connected peer. The result is that the connected peer is informed about the switch DCBX ETS configuration, but even if the peer is “willing,” the peer does not change its configuration to match the switch configuration. This is asymmetric configuration—the two interfaces can have different parameter values for the ETS attribute.

To disable the ETS Recommendation TLV:

- [edit protocols dcbx interface *interface-name*]
user@switch# **set enhanced-transmission-selection no-recommendation-tlv**

Related Documentation

- [Configuring the DCBX Mode on page 6523](#)
- [Configuring DCBX Autonegotiation on page 6524](#)
- [Understanding DCBX on page 6514](#)
- [Understanding Data Center Bridging Capability Exchange Protocol for EX Series Switches](#)

Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows

The switch supports up to six lossless forwarding classes. (Junos OS Release 12.3 increased support for lossless priorities from two lossless forwarding classes—the default **fcoe** and **no-loss** forwarding classes—to a maximum of six lossless forwarding classes.) Each forwarding class is mapped to an IEEE 802.1p code point (priority).



NOTE: Junos OS Release 13.1 introduced support for up to six lossless forwarding classes on QFabric systems. Throughout this document, features introduced on standalone switches in Junos OS Release 12.3 are introduced on QFabric systems in Junos OS Release 13.1 unless otherwise noted.

Only switches with native Fibre Channel (FC) interfaces, such as the QFX3500, support native FC traffic and configuration as an FCoE-FC gateway. Throughout this document, features that pertain to native FC traffic and to FCoE-FC gateway configuration apply only to switches that support native FC interfaces.



Video: [Why Use PFC in a Data Center Network?](#)

The default configuration is the same as the default configuration in Junos OS Release 12.2 and is backward-compatible. If you need only two (or fewer) lossless forwarding classes, use the default configuration, in which the **fcoe** and **no-loss** forwarding classes are lossless. If you need more than two lossless forwarding classes, you can use the two default lossless forwarding classes and configure additional lossless forwarding classes. If you do not want to use the default lossless forwarding classes, you can change them, or use only the lossless forwarding classes that you explicitly configure.

- [Default Lossless Priority Configuration on page 6873](#)
- [Configuring Lossless Priorities on page 6876](#)
- [Configuration Rules and Recommendations on page 6889](#)
- [Lossless Transport Features Introduced in Junos OS Release 12.3 \(Legacy Non-ELS CLI\) on page 6890](#)
- [Backward Compatibility with Junos OS Releases Earlier Than Release 12.3 \(Legacy Non-ELS CLI\) on page 6891](#)

Default Lossless Priority Configuration

If you do not explicitly configure forwarding classes, the system uses the default forwarding class configuration, which provides two default lossless forwarding classes (**fcoe** and **no-loss**). (If you change the forwarding class configuration, the changes apply to all traffic on that device because forwarding classes are global to a particular device.)

If you do not explicitly configure classifiers, and you do not explicitly configure flow control to pause output queues (configured in the output stanza of the CNP), the default classifier

and the default output queue pause configurations are applied to all Ethernet interfaces on the switches (or Node devices). You can override the default classifier and the default output queue pause configuration on a per-interface basis by applying an explicit configuration to an Ethernet interface. The default configuration is used on all Ethernet interfaces that do not have an explicit configuration.



NOTE: If you do not configure flow control on output queues, the default configuration uses a one-to-one mapping of IEEE 802.1p code points (priorities) to output queues by number. For example, priority 0 (code point 000) is mapped to queue 0, priority 1 (code point 001) is mapped to queue 1, and so on. If you do not use the default configuration, you must explicitly configure flow control on each output queue that you want to enable for PFC pause in the output stanza of the CNP.

In the default configuration, only queue 3 and queue 4 are enabled to respond to pause messages from the connected peer. For queue 3 to respond to pause messages, priority 3 (code point 011) must be enabled for PFC in the input stanza of the CNP. For queue 4 to respond to pause messages, priority 4 (code point 100) must be enabled for PFC in the input stanza of the CNP.

The default configuration provides the following lossless behavior:

- Two default lossless forwarding classes (the **no-loss** packet drop attribute is applied to these forwarding classes automatically):
fcoe—Mapped to output queue 3
no-loss—Mapped to output queue 4
- A default classifier that maps the fcoe forwarding class to IEEE 802.1p priority 3 (011) and the no-loss forwarding class to IEEE 802.1p priority 4 (100)
- Priority-based flow control (PFC) enabled on Ethernet interface output queues 3 and 4 when those queues carry lossless traffic (traffic that is mapped to the fcoe and no-loss forwarding classes, respectively).

On switches that can be configured as an FCoE-FC gateway, native FC interfaces (NP_Ports), with default flow control enabled on output queue 3 (IEEE 802.1p priority 3) for FCoE/FC traffic.

- DCBX is enabled on all interfaces in autonegotiation mode, and automatically exchanges FCoE application protocol type, length, and values (TLVs) on interfaces that carry FCoE traffic. However, if you explicitly configure DCBX protocol TLV exchange for any application, then you must explicitly configure protocol TLV exchange for every application for which you want DCBX to exchange TLVs, including FCoE.
- On Ethernet ports, PFC buffer calculations use the following default values to determine the headroom buffer size:
Cable length—100 meters (approximately 328 feet)
MRU for priority 3 traffic—2500 bytes
MRU for priority 4 traffic—9216 bytes
Maximum transmission unit (MTU)—1522 (or the configured MTU value for the interface)



NOTE: If you configure flow control on a priority that is not one of the default flow control priorities, the default MRU value is 2500 bytes. For example, if you configure flow control on priority 5 and you do not configure an MRU value, the default MRU value is 2500 bytes.



NOTE: In addition, to support lossless transport, PFC must be enabled explicitly on the lossless IEEE 802.1p priorities (code points) on ingress Ethernet interfaces; no default PFC configuration is applied at ingress interfaces. If you do not enable PFC on lossless priorities, those priorities might experience packet loss during periods of congestion. For example, if you want lossless FCoE traffic and you are using the default fcoe forwarding class, you use a CNP to enable PFC on priority 3 (code point 011), and apply that CNP to all ingress interfaces that carry FCoE traffic.

You can override the default classifier and the default output queue pause configuration on a per-interface basis by applying an explicit configuration to an Ethernet interface.

The default CoS configuration is backward-compatible with the *default* CoS configuration of software releases before Junos OS Release 12.3. If you explicitly configure lossless transport, ensure that the input and output queues corresponding to the lossless forwarding classes are explicitly configured for PFC pause.

Table 583 summarizes the default forwarding classes and their mapping to output queues, IEEE 802.1p priorities, and drop attributes.

Table 583: Mapping of Default Forwarding Class to Queue, IEEE 802.1p Priority, and Drop Attribute

Forwarding Class Name	Output Queue	Priority	Drop Attribute
best-effort	0	0	drop
fcoe	3	3	no-loss
no-loss	4	4	no-loss
network-control	7	7	drop

On switches that use the same forwarding classes and output queues for unicast and multdestination (multicast, broadcast, and destination lookup fail) traffic, these forwarding classes carry both unicast and multdestination traffic. Only unicast traffic is treated as lossless traffic. Multdestination traffic is not treated as lossless traffic, even on lossless output queues.

On switches that use different forwarding classes and output queues for unicast and multdestination traffic, there is one default multdestination forwarding class named *mcast*, which is mapped to output queue 8 with a drop attribute of drop. (Incoming

multidestination traffic on all IEEE 802.1p priorities is mapped to the mcast forwarding class by default.)

Configuring Lossless Priorities

To configure more than two lossless priorities (forwarding classes), or to change the default mapping of lossless forwarding classes to priorities and paused output queues, you must explicitly configure the switch instead of using the default configuration. Configuring lossless priorities includes:

- Configuring forwarding classes with the no-loss packet drop attribute.
- Using a CNP to configure PFC on ingress interfaces and flow control (PFC) on egress interfaces.
- Configuring a classifier to map IEEE 802.1p priorities (code points) to the correct forwarding classes (the forwarding classes for which you want lossless transport).



NOTE: If you expect a large amount of lossless traffic on your network and configure multiple lossless traffic classes, ensure that you reserve enough scheduling resources (bandwidth) and buffer space to support the lossless flows. (For switches that support shared buffer configuration, *Understanding CoS Buffer Configuration* describes how to configure buffers and provides a recommended buffer configuration for networks with larger amounts of lossless traffic. Buffer optimization is automatic on switches that use virtual output queues.)

In addition, on Ethernet interfaces, DCBX must exchange the appropriate application protocol TLVs for the lossless traffic. On switches that can act as an FCoE-FC gateway, you need to remap the FCoE priority on native FC interfaces if your network uses a priority other than 3 (IEEE code point 011) for FCoE traffic. This section describes:

- [Configuring Lossless Forwarding Classes \(Packet Drop Attribute\) on page 6876](#)
- [Congestion Notification Profiles \(PFC Configuration\) on page 6878](#)
- [Configuring DCBX \(Application Protocol TLV Exchange\) on page 6884](#)
- [Fate Sharing Among Traffic Classes on page 6885](#)
- [Transit Switch Configuration Versus FCoE-FC Gateway Configuration on page 6886](#)
- [Configuration Results and Commit Checks on page 6887](#)

Configuring Lossless Forwarding Classes (Packet Drop Attribute)

Junos OS Release 12.3 introduced the *no-loss* parameter for forwarding class configuration. (Although it uses the same name, this is not the no-loss default forwarding class. It is a packet drop attribute you can specify to configure any forwarding class as a lossless forwarding class.)



NOTE: On switches that use different forwarding classes for unicast and multideestination traffic, the forwarding class must be a unicast forwarding class. On switches that use the same forwarding classes for unicast and multideestination traffic, only unicast traffic receives lossless treatment.

You can configure up to six forwarding classes (depending on system architecture and the availability of system resources) as lossless forwarding classes by including the **no-loss** drop attribute at the **[edit class-of-service forwarding-classes class forwarding-class-name queue-num queue-number]** hierarchy level.

If you use the default fcoe or no-loss forwarding classes, they include the no-loss drop attribute by default. If you explicitly configure the fcoe or no-loss forwarding classes and you want to retain their lossless behavior, you *must* include the no-loss drop attribute in the configuration.



NOTE: All forwarding classes mapped to the same output queue must have the same packet drop attribute. (All forwarding classes mapped to the same output queue must be either lossy or lossless. You cannot map both a lossy and a lossless forwarding class to the same queue.)

To avoid fate sharing (a congested flow affecting an uncongested flow), use a one-to-one mapping of lossless forwarding classes to IEEE 802.1p code points (priorities) and queues. Map each lossless forwarding class to a different queue, and classify incoming traffic into forwarding classes so that each forwarding class transports traffic of only one priority (code point).

The fcoe and no-loss forwarding classes are special cases, because in the default configuration, they are configured for lossless behavior (providing that you also enable PFC on the priorities mapped to the fcoe and no-loss forwarding classes in the CNP input stanza).

[Table 584](#) summarizes the possible configurations of the fcoe and no-loss forwarding classes in Junos OS Release 12.3 and later, and the result of those configurations in terms of lossless traffic behavior. It is assumed that PFC, DCBX, and classifiers are properly configured.

Table 584: FCoE and No-Loss Forwarding Class Configuration in Junos OS Release 12.3

Explicit (User-Configured) or Default Forwarding Class Configuration	Packet Drop Attribute	Result and Notes
Default	Default	The fcoe and no-loss forwarding classes are lossless. NOTE: Even if you explicitly configure other forwarding classes (lossy or lossless forwarding classes), the fcoe and no-loss forwarding classes remain lossless because they are not explicitly configured.
Explicit	Not specified in the explicit forwarding class configuration	The fcoe and no-loss forwarding classes are lossy because they do not include the no-loss drop attribute.
Explicit	No-loss	The fcoe and no-loss forwarding classes are lossless.
Explicit, configured in Junos OS Release 12.2 or earlier	Not specified (packet drop attribute was not available before Junos OS Release 12.3)	The fcoe and no-loss forwarding classes are lossy in Junos OS Release 12.3 and later because they do not include the no-loss drop attribute. NOTE: To retain lossless behavior, before you upgrade to Junos OS Release 12.3, delete the explicit configuration so that the system uses the default configuration. Alternatively, you can reconfigure the forwarding classes with the no-loss packet drop attribute after upgrading to Junos OS Release 12.3 or later.

For all other forwarding classes except the **fcoe** and **no-loss** forwarding classes, you must explicitly configure lossless transport by specifying the no-loss packet drop attribute, because the default configuration for all other forwarding classes is lossy (the no-loss packet drop attribute is not applied).

Congestion Notification Profiles (PFC Configuration)

Use CNPs to configure lossless PFC characteristics on input and output interfaces.

The input stanza of a CNP enables PFC on specified IEEE 802.1p priorities (code points) and fine-tunes headroom buffer settings by configuring the maximum receive unit (MRU) value and cable length on ingress interfaces.

The output stanza of a CNP enables PFC (flow control) on output queues for specified IEEE 802.1p priorities so that the queues can respond to PFC pause messages from the connected peer on the priority of your choice. (By default, output queues 3 and 4 respond to received PFC messages when those queues carry lossless traffic in the fcoe and no-loss forwarding classes, respectively.)

To achieve lossless transport, the priority paused at the ingress interfaces must match the priority paused at the egress interfaces for a given traffic flow. For example, if you configure ingress interfaces to pause traffic tagged with IEEE 802.1p priority 5 (code point 101) and priority 5 traffic is mapped to output queue 5, then you must also configure the corresponding output interfaces to pause priority 5 on queue 5. In addition, the forwarding class mapped to queue 5 must be configured as a lossless forwarding class (using the no-loss drop attribute).



CAUTION: Any change to the PFC configuration on a port temporarily blocks the entire port (not just the priorities affected by the PFC change) so that the port can implement the change, then unblocks the port. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

A change to the PFC configuration means any change to a CNP, including changing the input portion of the CNP (enabling or disabling PFC on a priority, or changing the MRU or cable-length values) or changing the output portion the CNP that enables or disables output flow control on a queue. A PFC configuration change only affects ports that use the changed CNP.

The following actions change the PFC configuration:

- Deleting or disabling a PFC configuration (input or output) in a CNP that is in use on one or more interfaces. For example:
 1. An existing CNP with an input stanza that enables PFC on priorities 3, 5, and 6 is configured on interfaces xe-0/0/20 and xe-0/0/21.
 2. We disable the PFC configuration for priority 6 in the input CNP, and then commit the configuration.
 3. The PFC configuration change causes all traffic on interfaces xe-0/0/20 and xe-0/0/21 to stop until the PFC change has been implemented. When the PFC change has been implemented, traffic resumes.
- Configuring a CNP on an interface. (This changes the PFC state by enabling PFC on one or more priorities.)
- Deleting a CNP from an interface. (This changes the PFC state by disabling PFC on one or more priorities.)

Configuring Input Interface Flow Control (PFC and Headroom Buffer Calculation)

On Ethernet interfaces, the input stanza of the CNP enables PFC on specified priorities so that the ingress interface can send a pause message to the connected peer during periods of congestion. Input CNPs also fine-tune the headroom buffers used for PFC support by allowing you to configure the MRU value and cable length (if you do not want to use the default configuration).

Headroom buffers support lossless transport by storing the traffic that arrives at an interface after the interface sends a PFC flow control message to pause incoming traffic.

Until the connected peer receives the flow control message and pauses traffic, the interface continues to receive traffic and must buffer it (and the traffic that is still on the wire after the peer pauses) to prevent packet loss.

The system uses the MRU and the length of the attached physical cable to calculate buffer headroom allocation. The default configuration values are:

- MRU for priority 3 traffic—2500 bytes
- MRU for priority 4 traffic—9216 bytes
- Cable length—100 meters (approximately 328 feet)



NOTE: If you configure flow control on a priority that is not one of the default flow control priorities, the default MRU value is 2500 bytes. For example, if you configure flow control on priority 5 and you do not explicitly configure an MRU value, the default MRU value is 2500 bytes.

You can fine-tune the MRU and the cable length to adjust the size of the headroom buffer on an interface. The switch has a shared global buffer pool and dynamically allocates headroom buffer space to lossless queues as needed.

A lower MRU or a shorter cable length reduces the amount of headroom buffer required on an interface and leaves more headroom buffer space for other interfaces. A higher MRU or a longer cable length increases the amount of headroom buffer space required on an interface and leaves less headroom buffer space for other interfaces.

In many cases, you can better utilize the headroom buffers by reducing the MRU value (for example, an MRU of 2180 is sufficient for most FCoE networks) and by reducing the cable length value if the physical cable is less than 100 meters long.



NOTE: When you configure the headroom buffers by changing the MRU or the cable length, and commit the configuration, the system performs a commit check and rejects the configuration if sufficient headroom buffer space is not available.

However, the system does not perform a commit check but instead returns a syslog error if:

- The buffers are configured on a LAG interface.
- The default classifier is used on the interface (instead of a user-configured classifier).
- The interface has not been created yet.

Configuring Output Interface Flow Control (PFC)

On Ethernet interfaces, you can use the output stanza of the CNP to configure flow control on output queues and enable PFC pause response on specified IEEE 802.1p priorities.



NOTE: On switches that use different output queues for unicast and multdestination traffic, the queue must be a unicast output queue.

By default, output queues 3 and 4 are enabled for PFC pause on priorities 3 (IEEE 802.1p code point 011) and 4 (IEEE 802.1p code point 100). The default PFC pause response supports the default lossless forwarding class configuration, which maps the fcoe forwarding class to queue 3 and priority 3, and maps the no-loss forwarding class to queue 4 and priority 4.

Configuring PFC on output queues enables you to pause any priority on any output queue on any Ethernet interface. Output flow control enables you to use more than two output queues to support lossless traffic flows (you can configure up to six lossless forwarding classes and map them to different output queues that are enabled for PFC pause). Output queue flow control also enables you to support multiple lossless forwarding classes (each mapped to a different priority and output queue) for one class of traffic.



NOTE: Output flow control only works when PFC is enabled in the CNP input stanza on the corresponding priorities on the interface. For example, if you enable output flow control on priority 5 (IEEE 802.1p code point 101), then you must also enable PFC in the CNP on the input stanza on priority 5.

For example, if the converged Ethernet network uses two different priorities for FCoE traffic (for example, priority 3 and priority 5), then you can classify those priorities into different lossless forwarding classes that are mapped to different output queues:

1. Configure two lossless forwarding classes for FCoE traffic, with each forwarding class mapped to a different output queue. For example, you could use the default fcoe forwarding class, which is mapped to queue 3, and you could configure a second lossless forwarding class called fcoe1 and map it to queue 5. The fcoe forwarding class is for priority 3 FCoE traffic (code point 011), and the fcoe1 forwarding class is for priority 5 (code point 101) FCoE traffic.
2. Configure a classifier that maps each forwarding class to the desired IEEE 802.1p code point (priority). If FCoE traffic on both priorities uses one interface, the classifier must classify both forwarding classes to the correct priorities. If FCoE traffic of different priorities uses different interfaces, the classifier configuration on each interface must map the correct priority to the corresponding lossless forwarding class.
3. Apply the classifier to the interfaces that carry FCoE traffic. The classifier determines the mapping of forwarding classes to priorities on each interface.

To configure lossless transport for these forwarding classes, you also need to:

- Enable PFC on the two priorities (3 and 5 in this example) at the ingress interfaces in the CNP input stanza.
- Configure PFC on the output queues and priorities for the forwarding classes in the CNP output stanza so that the interface can respond to pause messages received from the connected peer.



NOTE: When you configure the CNP on an interface, all ingress and egress traffic is blocked until the configuration is implemented, then the interface is unblocked and traffic resumes. During the time the interface is blocked, all queues on the interface experience packet loss.

- Configure DCBX to exchange application protocol TLVs on both FCoE priorities.



NOTE: If you do not configure flow control to pause output queues, the default configuration uses a one-to-one mapping of IEEE 802.1p code points (priorities) to output queues by number. For example, priority 0 (code point 000) is mapped to queue 0, priority 1 (code point 001) is mapped to queue 1, and so on. By default, only queues 3 and 4 are enabled to respond to pause messages from the connected peer, and you must explicitly enable PFC on the corresponding priorities in the CNP input stanza to achieve lossless behavior.

If you do not use the default configuration, you must explicitly configure flow control on each output queue that you want to enable for PFC pause. For example, if you explicitly configure flow control on output queue 5, the default configuration is no longer valid, and only output queue 5 is enabled for PFC pause. Output queues 3 and 4 are no longer enabled for PFC pause, so traffic using those queues no longer responds to PFC pause messages even if the corresponding forwarding class is configured with the no-loss drop attribute. To retain the pause configuration on output queues 3 and 4 and configure flow control on queue 5, you need to explicitly configure flow control on queues 3, 4, and 5.

On switches that use different output queues for unicast and multidestination traffic, you cannot configure flow control to pause a multidestination output queue. You can configure flow control to pause only unicast output queues. On switches that use the same output queues for unicast and multidestination traffic, only unicast traffic receives lossless treatment.

Output Interface Flow Control Profiles

Configuring the CNP output stanza creates an output flow control profile that tells egress ports the queues on which the Ethernet interface should respond to PFC pause messages. Although you can create an unlimited number of CNPs that contain input stanzas only, the number of CNPs that you can configure with output stanzas is limited:

- For standalone switches that are not part of a QFabric system, you can configure up to two output interface flow control profiles. (You can configure up to two CNPs with output stanzas.)
- For QFabric systems, you can configure one output interface flow control profile per Node device. (You can configure one CNP with an output stanza per Node device.)

There are a total of four output flow control profiles.

The system has a default output flow control profile that is applied to all Ethernet interfaces when the CNP attached to the interface has only an input stanza and does not include an output stanza. The default profile responds to PFC pause messages received on queue 3 (for priority 3, for the default fcoe forwarding class) and on queue 4 (for priority 4, for the default no-loss forwarding class), and is effective only if PFC is configured on those priorities in the CNP input stanza.

Additionally, the system has two internal output flow control profiles that it applies automatically to fabric (FTE) ports and to native FC interfaces (NP_Ports). When the switch is not part of a QFabric system, the profile normally used for FTE ports is available for user configuration and provides a second user-configurable profile. (That is why standalone switches have two user-configurable output flow control profiles, but Node devices on a QFabric system have only one user-configurable output flow control profile.)

Because one output CNP can configure PFC pause response on multiple output queues (priorities), one user-configurable output CNP is usually flexible enough to specify the desired PFC response on all programmed interfaces.



NOTE: Each port can use one output flow control profile. You cannot apply more than one profile to one port.

Output flow control profiles can be expressed in table format. For example, [Table 585](#) shows the default output flow control profile that pauses priorities 3 and 4 on queues 3 and 4 (remember that PFC must also be enabled on code points 3 and 4 in the CNP input stanza in order for PFC to work):

Table 585: Default Output Flow Control Profile

IEEE 802.1p Priority Specified in Received PFC Frame	Paused Output Queue
0 (000)	—
1 (001)	—
2 (010)	—
3 (011)	3
4 (100)	4
5 (101)	—
6 (110)	—
7 (111)	—

[Table 586](#) is an example of a user-configured output flow control profile. Using the example from the preceding section, the CNP output stanza configures flow control on output queue 5, and also explicitly configures output flow control on queues 3 and 4 for

the fcoe and no-loss forwarding classes. (If you explicitly configure an output CNP, you must explicitly configure every output queue that you want to respond to PFC messages, because the user-configured profile overrides the default profile. If this example did not include queues 3 and 4, those queues would no longer respond to received PFC messages.)

Table 586: User-Configured Output Flow Control Profile

IEEE 802.1p Priority Specified in Received PFC Frame	Paused Output Queue
0 (000)	—
1 (001)	—
2 (010)	—
3 (011)	3
4 (100)	4
5 (101)	5
6 (110)	—
7 (111)	—

Remember that you must also enable PFC on code points 3, 4, and 5 in the CNP input stanza for this configuration to work. When you configure the CNP on an interface, all ingress and egress traffic is blocked until the configuration is implemented, then the interface is unblocked and traffic resumes. During the time the interface is blocked, all queues on the interface experience packet loss.

Configuring PFC Across Layer 3 Interfaces on QFX5200, QFX5100, EX4600, and QFX10000 Switches

Enabling PFC on traffic flows is based on the IEEE 802.1p code point (priority) in the priority code point (PCP) field of the Ethernet frame header (sometimes known as the CoS bits). To enable PFC on traffic that crosses Layer 3 interfaces, the traffic must be classified by its IEEE 802.1p code point, not by its DSCP (or DSCP IPv6) code point.

See [“Understanding PFC Functionality Across Layer 3 Interfaces” on page 6921](#) for a conceptual overview of how to enable PFC on traffic across Layer 3 interfaces. See [“Example: Configuring PFC Across Layer 3 Interfaces” on page 6923](#) for an example of how to configure PFC on traffic that traverses Layer 3 interfaces.

Configuring DCBX (Application Protocol TLV Exchange)

For applications that require lossless transport, DCBX exchanges application protocol TLVs with the connected peer interface. By default, DCBX advertises FCoE application protocol TLVs on all interfaces that are enabled for DCBX, and by default, DCBX is enabled on all interfaces. DCBX advertises no other applications by default.

For each application (for example, iSCSI) that you want to configure for lossless transport, you must enable the interfaces which carry that application traffic to exchange DCBX protocol TLVs with the connected peer. The TLV exchange allows the peer interfaces to negotiate a compatible configuration to support the application.

If you configure DCBX to advertise any application, the default DCBX advertisement is overridden, and DCBX advertises only the configured applications. If you want an interface to advertise only the FCoE application, you do not have to configure DCBX application protocol TLV exchange; instead, you can use the default configuration.

If you want DCBX to advertise other applications, you must explicitly configure an application map and apply it to the interfaces on which you want to exchange protocol TLVs for those applications. If you want to exchange FCoE application protocol TLVs in addition to other application protocol TLVs, you must also explicitly configure the FCoE application in the application map. [“Understanding DCBX Application Protocol TLV Exchange” on page 6528](#) describes how application mapping works.



NOTE: Lossless transport also requires that you enable PFC on the correct priority (IEEE 802.1p code point) on the ingress interfaces using an input CNP. If the priority you pause at the ingress interfaces is not mapped to queue 3 or queue 4 (the two output queues that are enabled for PFC pause flow control by default), then you must also enable the output queues that correspond to paused input priorities to pause using the output stanza of the CNP.

Fate Sharing Among Traffic Classes

You can configure different lossless (or lossy) traffic flows to share fate—that is, to receive the same CoS treatment.

Fate sharing is not desirable for I/O convergence. Instead of independent control of the fate of each type of flow, different types of flows receive the same treatment. Fate sharing is particularly undesirable for lossless flows. If one lossless flow experiences congestion and must be paused, that affects flows that share fate with the congested flow even if the other flows are not experiencing congestion, and also can cause ingress port congestion. If your network requires that all 802.1p priorities be lossless, you can achieve that by allowing some fate sharing among the eight priorities by spreading them across up to six lossless forwarding classes.

If the number of lossless priorities is less than or equal to the number of configured lossless forwarding classes, then you can avoid fate sharing by configuring a one-to-one mapping of forwarding classes to IEEE 802.1p code points (priorities) and output queues. (Each forwarding class should be mapped to a different output queue and classified to a different priority.)

If you want to configure different traffic flows to share fate, two fate-sharing configurations are supported: mapping one forwarding class to more than one IEEE 802.1p code point (priority), and mapping two forwarding classes to the same output queue:

1. If you map one lossless forwarding class to more than one priority, the traffic tagged with each of the priorities uses the same CoS properties associated (the CoS properties associated with the forwarding class). For example, configuring a forwarding class called fc1, mapping it to queue 1, and mapping it to code points 101 and 110 using a classifier named classify1 results in the traffic tagged with priorities 101 and 110 sharing fate:

```
user@switch# set class-of-service forwarding-classes class fc1 queue-num 1 no-loss
user@switch# set class-of-service classifiers ieee-802.1 classify1 forwarding class fc1
loss-priority low code-points 101
user@switch# set class-of-service classifiers ieee-802.1 classify1 forwarding class fc1
loss-priority low code-points 110
```

In this case, if the traffic mapped to either priority experiences congestion, both priorities are paused because they are mapped to the same forwarding class and are therefore treated similarly.

2. If you map multiple lossless forwarding classes to the same output queue, the traffic mapped to the forwarding classes uses the same output queue. This increases the amount of traffic on the queue, and can create congestion that affects all of the traffic flows that are mapped to the queue. For example, configuring two forwarding classes called fc1 and fc2, mapping both forwarding classes to queue 1, and mapping the forwarding classes to code points 101 and 110 (respectively) using a classifier named classify1, results in the traffic tagged with priorities 101 and 110 sharing fate on the same output queue:

```
user@switch# set class-of-service forwarding-classes class fc1 queue-num 1 no-loss
user@switch# set class-of-service forwarding-classes class fc2 queue-num 1 no-loss
user@switch# set class-of-service classifiers ieee-802.1 classify1 forwarding class fc1
loss-priority low code-points 101
user@switch# set class-of-service classifiers ieee-802.1 classify1 forwarding class fc2
loss-priority low code-points 110
```

In this case, even though the two forwarding classes use different IEEE 802.1p priorities, if one forwarding class experiences congestion, it affects the other forwarding class. The reason is that if the output queue is paused because of congestion on either forwarding class, all traffic that uses that queue is paused. Since both forwarding classes are mapped to the queue, the traffic mapped to both forwarding classes is paused.



NOTE: If you map more than one forwarding class to a queue, all of the forwarding classes mapped to the same queue must have the same packet drop attribute (all of the forwarding classes must be lossy, or all of the forwarding classes mapped to a queue must be lossless).

Transit Switch Configuration Versus FCoE-FC Gateway Configuration

On a transit switch (all Ethernet ports, no native FC ports) that forwards FCoE traffic (or other traffic that requires lossless transport across the Ethernet network), the configuration of classifiers, lossless forwarding classes, DCBX, and PFC on ingress and egress interfaces to support lossless transport is as described in this document.

When a switch acts as an FCoE-FC gateway (if native FC interfaces are supported on your switch), the system uses native FC interfaces (NP_Ports) to connect to the FC switch (or FCoE forwarder) at the FC network edge. You cannot apply CNPs or DCBX to native FC interfaces, only to Ethernet interfaces.

On an FCoE-FC gateway, the Ethernet interface configuration of classifiers, DCBX, and PFC is the same as the Ethernet interface configuration on a transit switch. The configuration of lossless forwarding classes is also the same.

However, supporting lossless transport on native FC interfaces requires that you rewrite the IEEE 802.1p priority value *if* your network uses any priority other than 3 (IEEE code point 011) for FCoE traffic. If your network uses priority 3 for FCoE traffic, you can and should use the default configuration on native FC interfaces.

By default, native FC interfaces tag packets with priority 3 when they encapsulate the incoming FC packets in Ethernet. If your FCoE network uses a different priority than 3 for FCoE traffic, you need to rewrite the priority value to the value that your network uses on the FC interface, classify the FCoE traffic to the correct priority on the Ethernet interfaces, and enable PFC on the correct priority on the Ethernet interfaces, as described in *Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway*.

Configuration Results and Commit Checks

Different configurations of forwarding classes and their drop attributes, classifiers, CNPs (PFC flow control), and Ethernet PAUSE (IEEE 802.3X flow control) result in different system behaviors.

[Table 587](#) describes the results of the possible lossless transport configurations in each case. The assumption in the *Result* column is that the system's buffer headroom calculation resulted in a successful configuration.

However, if the system calculates that there is insufficient buffer space to support the configuration, a commit check prevents you from committing the configuration on an individual Ethernet interface. For LAG interfaces, the system does not issue a commit check error but instead issues a syslog message.



NOTE: After you configure lossless transport for a LAG interface, be sure to check the syslog messages to confirm that the commit was successful.

Table 587: Results of Lossless Priority Configuration

Classifier Configuration	Congestion Notification Profile Configuration	Ethernet PAUSE (IEEE 802.3X) Configuration	Result
None (default classifier)	None	None	System default configuration. No flows are lossless. To achieve lossless behavior for the default fcoe and no-loss forwarding classes, you must configure an input CNP to enable PFC on their IEEE 802.1p code points (011 and 100 respectively).
Classifier with no lossless forwarding classes	None	None	No lossless traffic flows are configured; all traffic is best effort.
Classifier with at least one lossless forwarding class	None	None	Because no CNP is attached to interfaces, PFC is not enabled on the code point of the lossless traffic and no headroom buffer is allocated to the lossless queue, so packets can drop during periods of congestion. This configuration does not achieve lossless behavior.
None (default classifier)	PFC enabled on the fcoe and no-loss forwarding class code points (priorities)	None	The default classifier classifies traffic into two lossless forwarding classes, fcoe and no-loss. The CNP enables PFC on the priorities mapped to both lossless forwarding classes, resulting in lossless behavior for traffic mapped to the fcoe and no-loss forwarding classes.
None (default classifier)	None	Flow control enabled	The system calculates buffer headroom for the physical link based on the interface MTU and the default cable length. The system does not calculate buffer headroom for individual output queues. Because Ethernet PAUSE is enabled on the link instead of PFC being enabled on the lossless priorities, the entire link is paused during periods of congestion. This configuration results in lossless behavior for all of the forwarding classes on the link, but because all traffic is paused, this can cause greater overall network congestion.
Classifier with at least one lossless forwarding class	PFC enabled on the lossless forwarding class code points (priorities)	None	Headroom buffer allocated only to priorities that are mapped to the lossless forwarding classes and on which PFC is enabled. This configuration achieves lossless behavior for the lossless forwarding classes.

Table 587: Results of Lossless Priority Configuration (*continued*)

Classifier Configuration	Congestion Notification Profile Configuration	Ethernet PAUSE (IEEE 802.3X) Configuration	Result
Classifier with no lossless forwarding classes	None	Flow control enabled	The system calculates buffer headroom for the physical link based on the interface MTU and the default cable length, and it pauses all traffic on the link during periods of congestion.
Classifier with at least one lossless forwarding class	None	Flow control enabled	The system calculates buffer headroom for the physical link based on the interface MTU and the default cable length, and it pauses all traffic on the link during periods of congestion.
Classifier with at least one lossless forwarding class	PFC enabled on the lossless forwarding class code points (priorities)	Flow control enabled on a <i>different</i> interface than the interface with the CNP	The system checks the available buffer space for both the PFC-enabled priorities and for the other link. If sufficient buffer space is available, the lossless forwarding classes configured with PFC on one interface and also all of the traffic on the link with Ethernet PAUSE enabled achieve lossless behavior.



NOTE: If you attempt to configure both PFC and Ethernet PAUSE on a link, the system returns a commit error. PFC and Ethernet PAUSE are mutually exclusive configurations on an interface.

Configuration Rules and Recommendations

Keep in mind the following configuration rules and recommendations when you configure lossless traffic flows:

- You can configure a maximum of six lossless forwarding classes (forwarding classes with the no-loss packet drop attribute).
- All forwarding classes that you map to the same queue must have the same packet drop attribute (all of the forwarding classes must be lossy, or all of the forwarding classes must be lossless).
- Do not configure weighted random early detection (WRED) on lossless forwarding classes. (Do not associate a drop profile with a forwarding class that has the no-loss packet drop attribute.)

- On switches that use different forwarding classes and output queues for unicast and multdestination traffic, you cannot configure flow control to pause a multdestination output queue. You can configure PFC flow control only to pause unicast output queues.
- On switches that use different forwarding classes and output queues for unicast and multdestination traffic, forwarding classes mapped to multdestination queues (queues 8 through 11) cannot have the no-loss packet drop attribute. (Multdestination forwarding classes cannot be configured as lossless forwarding classes.)

Lossless Transport Features Introduced in Junos OS Release 12.3 (Legacy Non-ELS CLI)

Support for lossless transport introduced in Junos OS Release 12.3 includes:

- Configuring up to six lossless forwarding classes.
- Configuring PFC pause on output queues to program the output queues that can respond to PFC pause messages received from the connected peer. The priorities you pause on output queues must match the priorities on which you enable PFC on the corresponding ingress interfaces. For example, if you program output queues to pause priorities 3 (011) and 5 (101), then you must also enable pause on priorities 3 and 5 on the corresponding ingress interfaces. Configuring flow control on the output queues and enabling PFC on the corresponding input queues allows you to pause up to six priorities (forwarding classes).
- Controlling the headroom buffer on Ethernet interfaces by configuring the maximum receive unit (MRU) size for the traffic mapped to an IEEE 802.1p priority (configured per priority) and the length of the attached cable (configured per interface). The MRU size can range up to full jumbo packet size (9216 bytes).
- Remapping (rewriting) IEEE 802.1p priorities on native Fibre Channel (FC) interfaces when the system is acting as an FCoE-FC gateway. If the Ethernet (FCoE) network uses a different IEEE 802.1p priority than priority 3 (011) for FCoE traffic, then you can use priority remapping to classify FCoE traffic into a lossless forwarding class mapped to that different priority (see *Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway*).

Lossless transport still requires configuring previously existing features, including enabling PFC on the lossless priorities on ingress interfaces, and configuring classifiers to classify incoming traffic into lossless forwarding classes based on the IEEE 802.1p priority tag of the packet.



NOTE: If you expect a large amount of lossless traffic on your network and configure multiple lossless traffic classes, ensure that you reserve enough scheduling resources (bandwidth) and lossless headroom buffer space to support the lossless flows. (*Understanding CoS Buffer Configuration* describes how to configure buffers and provides a recommended buffer configuration for networks with larger amounts of lossless traffic.)

Backward Compatibility with Junos OS Releases Earlier Than Release 12.3 (Legacy Non-ELS CLI)

The addition of the no-loss packet drop attribute to forwarding class configuration means that when you upgrade from an earlier release to Junos OS Release 12.3, the new software might not preserve the lossless forwarding class configuration of the fcoe and no-loss forwarding classes.

If you used the default forwarding class configuration for the fcoe and no-loss forwarding classes, the CoS configuration is backward-compatible. You do not have to do anything to preserve the lossless behavior of traffic that uses those forwarding classes when you upgrade to Junos OS Release 12.3. (This is because the default configuration of these two forwarding classes includes the no-loss packet drop attribute.)

However, if you explicitly configured the fcoe or the no-loss forwarding class by including the **set forwarding-classes class forwarding-class-name queue-num queue-number** statement at the **[edit class-of-service]** hierarchy level, then those forwarding classes are no longer lossless, they are lossy. (They are lossy because explicit configuration in releases earlier than Junos OS Release 12.3 did not use the no-loss packet drop attribute.) In Junos OS Release 12.3 and later, you must include the no-loss packet drop attribute in explicit forwarding class configurations to configure a lossless forwarding class.

For example, before Junos OS Release 12.3, the following explicit configuration resulted in a lossless forwarding class:

```
user@switch# set class-of-service forwarding-classes class fcoe queue-num 3
```

However, in Junos OS Release 12.3, this configuration is lossy because it does not include the no-loss packet drop attribute. To preserve lossless behavior, after upgrading to Junos OS Release 12.3, you need to add the no-loss drop attribute:

```
user@switch# set class-of-service forwarding-classes class fcoe queue-num 3 no-loss
```

Alternatively, you can delete the explicit configuration before you upgrade to Junos OS Release 12.3 so that the system uses the default forwarding class, which is lossless:

```
user@switch# delete class-of-service forwarding-classes class fcoe queue-num 3
```



NOTE: The explicit configuration of other forwarding classes does not affect the lossless (or lossy) state of the fcoe and no-loss forwarding classes, because only the fcoe and no-loss forwarding classes were lossless forwarding classes before Junos OS Release 12.3. For example, if you explicitly configured the best-effort forwarding class but you used the default fcoe and no-loss forwarding classes in Junos OS Release 12.2, then when you upgrade to Junos OS Release 12.3, the fcoe and no-loss forwarding classes are still lossless (and the best-effort forwarding classes retains its explicit configuration).



NOTE: To achieve lossless behavior for the traffic belonging to any forwarding class, you must also use a CNP to enable PFC on the IEEE 802.1p priority mapped to the forwarding class and apply the CNP to the relevant interfaces, and ensure that DCBX exchanges the protocol TLVs for the application with the connected peer.

**Related
Documentation**

- [Understanding DCBX Application Protocol TLV Exchange on page 6528](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 6493](#)
- [Understanding PFC Functionality Across Layer 3 Interfaces on page 6921](#)
- [Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic \(FCoE Transit Switch\) on page 6969](#)
- [Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 6977](#)
- [Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 6986](#)
- [Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications \(FCoE and iSCSI\) on page 7000](#)
- [Example: Configuring PFC Across Layer 3 Interfaces on page 6923](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 6905](#)

Understanding CoS Flow Control (Ethernet PAUSE and PFC)

Flow control supports lossless transmission by regulating traffic flows to avoid dropping frames during periods of congestion. Flow control stops and resumes the transmission of network traffic between two connected peer nodes on a full-duplex Ethernet physical link. Controlling the flow by pausing and restarting it prevents buffers on the nodes from overflowing and dropping frames. You configure flow control on a per-interface basis.

Two methods of peer-to-peer flow control are supported:

- IEEE 802.3X Ethernet PAUSE



NOTE: QFX10000 switches do not support Ethernet PAUSE. Information about Ethernet PAUSE does not apply to QFX10000 switches.

OCX Series switches support symmetric Ethernet PAUSE flow control on Layer 3 tagged interfaces. OCX Series switches do not support asymmetric Ethernet PAUSE flow control. Information about asymmetric flow control does not apply to OCX Series switches.

- IEEE 802.1Qbb priority-based flow control (PFC)



NOTE: OCX Series switches do not support PFC or lossless Layer 2 transport. Information about PFC, lossless transport, and congestion notification profiles do not apply to OCX Series switches.



Video: [Why Use PFC in a Data Center Network?](#)

- [General Information about Ethernet PAUSE and PFC and When to Use Them on page 6893](#)
- [Ethernet PAUSE on page 6894](#)
- [PFC on page 6899](#)
- [Lossless Transport Support Summary on page 6902](#)

General Information about Ethernet PAUSE and PFC and When to Use Them

Ethernet PAUSE and PFC are link-level flow control mechanisms.



NOTE: For end-to-end congestion control for best-effort traffic, see [“Understanding CoS Explicit Congestion Notification” on page 6820](#).

Ethernet PAUSE pauses transmission of all traffic on a physical Ethernet link.

PFC decouples the pause function from the physical Ethernet link and enables you to divide traffic on one link into eight priorities. You can think of the eight priorities as eight “lanes” of traffic that are mapped to forwarding classes and output queues. Each priority maps to a 3-bit IEEE 802.1p CoS code point value in the VLAN header. You can enable PFC on one or more priorities (IEEE 802.1p code points) on a link. When PFC-enabled traffic is paused on a link, traffic that is not PFC-enabled continues to flow (or is dropped if congestion is severe enough).

Use Ethernet PAUSE when you want to prevent packet loss on all of the traffic on a link. Use PFC to prevent traffic loss only on a specified type of traffic that require lossless treatment, for example, Fibre Channel over Ethernet (FCoE) traffic.



NOTE: Depending on the amount of traffic on a link or assigned to a priority, pausing traffic can cause ingress port congestion and spread congestion through the network.

Ethernet PAUSE and PFC are mutually exclusive configurations on an interface. Attempting to configure both Ethernet PAUSE and PFC on a link causes a commit error.

By default, all forms of flow control are disabled. You must explicitly enable flow control on interfaces to pause traffic.

Ethernet PAUSE

Ethernet PAUSE is a congestion relief feature that works by providing link-level flow control for all traffic on a full-duplex Ethernet link. Ethernet PAUSE works in both directions on the link. In one direction, an interface generates and sends Ethernet PAUSE messages to stop the connected peer from sending more traffic. In the other direction, the interface responds to Ethernet PAUSE messages it receives from the connected peer to stop sending traffic.



NOTE: QFX10000 switches do not support Ethernet PAUSE. Information about Ethernet PAUSE does not apply to QFX10000 switches.

OCX Series switches support symmetric Ethernet PAUSE flow control on Layer 3 tagged interfaces. OCX Series switches do not support asymmetric Ethernet PAUSE flow control. Information about asymmetric flow control does not apply to OCX Series switches.

Ethernet PAUSE also works on aggregated Ethernet interfaces. For example, if the connected peer interfaces are called Node A and Node B:

- When the receive buffers on interface Node A reach a certain level of fullness, the interface generates and sends an Ethernet PAUSE message to the connected peer (interface Node B) to tell the peer to stop sending frames. The Node B buffers store frames until the time period specified in the Ethernet PAUSE frame elapses; then Node B resumes sending frames to Node A.

- When interface Node A receives an Ethernet PAUSE message from interface Node B, interface Node A stops transmitting frames until the time period specified in the Ethernet PAUSE frame elapses; then Node A resumes transmission. (The Node A transmit buffers store frames until Node A resumes sending frames to Node B.)

In this scenario, if Node B sends an Ethernet PAUSE frame with a time value of 0 to Node A, the 0 time value indicates to Node A that it can resume transmission. This happens when the Node B buffer empties to below a certain threshold and the buffer can once again accept traffic.

Symmetric flow control means an interface has the same Ethernet PAUSE configuration in both directions. The Ethernet PAUSE generation and Ethernet PAUSE response functions are both configured as enabled, or they are both disabled. You configure symmetric flow control by including the **flow-control** statement at the **[edit interfaces interface-name ether-options]** hierarchy level.

Asymmetric flow control allows you to configure the Ethernet PAUSE functionality in each direction independently on an interface. The configuration for generating Ethernet PAUSE messages and for responding to Ethernet PAUSE messages does not have to be the same. It can be enabled in both directions, disabled in both directions, or enabled in one direction and disabled in the other direction. You configure asymmetric flow control by including the **configured-flow-control** statement at the **[edit interfaces interface-name ether-options]** hierarchy level.

On any particular interface, symmetric and asymmetric flow control are mutually exclusive. Asymmetric flow control overrides and disables symmetric flow control. (If PFC is configured on an interface, you cannot commit an Ethernet PAUSE configuration on the interface. Attempting to commit an Ethernet PAUSE configuration on an interface with PFC enabled on one or more queues results in a commit error. To commit the PAUSE configuration, you must first delete the PFC configuration.) Both symmetric and asymmetric flow control are supported.

- [Symmetric Flow Control on page 6895](#)
- [Asymmetric Flow Control on page 6895](#)

Symmetric Flow Control

Symmetric flow control configures both the receive and transmit buffers in the same state. The interface can both send Ethernet PAUSE messages and respond to them (flow control is enabled), or the interface cannot send Ethernet PAUSE messages or respond to them (flow control is disabled).

When you enable symmetric flow control on an interface, the Ethernet PAUSE behavior depends on the configuration of the connected peer. With symmetric flow control enabled, the interface can perform any Ethernet PAUSE functions that the connected peer can perform. (When symmetric flow control is disabled, the interface does not send or respond to Ethernet PAUSE messages.)

Asymmetric Flow Control

Asymmetric flow control enables you to specify independently whether or not the interface receive buffer generates and sends Ethernet PAUSE messages to stop the connected

peer from transmitting traffic, and whether or not the interface transmit buffer responds to Ethernet PAUSE messages it receives from the connected peer and stops transmitting traffic. The receive buffer configuration determines if the interface transmits Ethernet PAUSE messages, and the transmit buffer configuration determines if the interface receives and responds to Ethernet PAUSE messages:

- Receive buffers on—Enable Ethernet PAUSE transmission (generate and send Ethernet PAUSE frames)
- Transmit buffers on—Enable Ethernet PAUSE reception (respond to received Ethernet PAUSE frames)

You must explicitly set the flow control for both the receive buffer and the transmit buffer (**on** or **off**) to configure asymmetric Ethernet PAUSE. [Table 513](#) describes the configured flow control state when you set the receive (Rx) and transmit (Tx) buffers on an interface:

Table 588: Asymmetric Ethernet PAUSE Flow Control Configuration

Receive (Rx) Buffer	Transmit (Tx) Buffer	Configured Flow Control State
On	Off	Interface generates and sends Ethernet PAUSE messages. Interface does not respond to Ethernet PAUSE messages (interface continues to transmit even if peer requests that the interface stop sending traffic).
Off	On	Interface responds to Ethernet PAUSE messages received from the connected peer, but does not generate or send Ethernet PAUSE messages. (The interface does not request that the connected peer stop sending traffic.)
On	On	Same functionality as symmetric Ethernet PAUSE. Interface generates and sends Ethernet PAUSE messages and responds to received Ethernet PAUSE messages.
Off	Off	Ethernet PAUSE flow control is disabled.

The configured flow control is the Ethernet PAUSE state configured on the interface.

On 1-Gigabit Ethernet interfaces, autonegotiation of Ethernet PAUSE with the connected peer is supported. (Autonegotiation on 10-Gigabit Ethernet interfaces is not supported.) Autonegotiation enables the interface to exchange state advertisements with the connected peer so that the two devices can agree on the Ethernet PAUSE configuration. Each interface advertises its flow control state to the connected peer using a combination of the Ethernet PAUSE and ASM_DIR bits, as described in [Table 514](#):

Table 589: Flow Control State Advertised to the Connected Peer (Autonegotiation)

Rx Buffer State	Tx Buffer State	PAUSE Bit	ASM_DIR Bit	Description
Off	Off	0	0	The interface advertises no Ethernet PAUSE capability. This is equivalent to disabling flow control on an interface.
On	On	1	0	The interface advertises symmetric flow control (both the transmission of Ethernet PAUSE messages and the ability to receive and respond to Ethernet PAUSE messages).
On	Off	0	1	The interface advertises asymmetric flow control (the transmission of Ethernet PAUSE messages, but not the ability to receive and respond to Ethernet PAUSE messages).
Off	On	1	1	The interface advertises both symmetric and asymmetric flow control. Although the interface does not generate and send Ethernet PAUSE requests to the peer, the interface supports both symmetric and asymmetric Ethernet PAUSE configuration on the peer because the peer is not affected if the peer does not receive Ethernet PAUSE requests. (If the interface responds to the peer's Ethernet PAUSE requests, that is sufficient to support either symmetric or asymmetric flow control on the peer.)

The flow control configuration on each switch interface interacts with the flow control configuration of the connected peer. Each peer advertises its state to the other peer. The interaction of the flow control configuration of the peers determines the flow control

behavior (resolution) between them, as shown in [Table 515](#). The first four columns show the Ethernet PAUSE configuration on the local QFX Series or EX4600 switch and on the connected peer (also known as the *link partner*). The last two columns show the Ethernet PAUSE resolution that results from the local and peer configurations on each interface. This illustrates how the Ethernet PAUSE configuration of each interface affects the Ethernet PAUSE behavior on the other interface.



NOTE: In the Resolution columns of the table, disabling Ethernet PAUSE transmit means that the interface receive buffers do not generate and send Ethernet PAUSE messages to the peer. Disabling Ethernet PAUSE receive means that the interface transmit buffers do not respond to Ethernet PAUSE messages received from the peer.

Table 590: Asymmetric Ethernet PAUSE Behavior on Local and Peer Interfaces

Local Interface (QFX Series or EX4600 Switch)		Peer Interface		Local Resolution	Peer Resolution
PAUSE Bit	ASM_DIR Bit	PAUSE Bit	ASM_DIR Bit		
0	0	Don't care	Don't care	Disable Ethernet PAUSE transmit and receive	Disable Ethernet PAUSE transmit and receive
0	1	0	Don't care	Disable Ethernet PAUSE transmit and receive	Disable Ethernet PAUSE transmit and receive
0	1	1	0	Disable Ethernet PAUSE transmit and receive	Disable Ethernet PAUSE transmit and receive
0	1	1	1	Enable Ethernet PAUSE transmit and disable Ethernet PAUSE receive	Disable Ethernet PAUSE transmit and enable Ethernet PAUSE receive
1	0	0	Don't care	Disable Ethernet PAUSE transmit and receive	Disable Ethernet PAUSE transmit and receive
1	0	1	Don't care	Enable Ethernet PAUSE transmit and receive	Enable Ethernet PAUSE transmit and receive
1	1	0	0	Disable Ethernet PAUSE transmit and receive	Disable Ethernet PAUSE transmit and receive
1	1	0	1	Enable Ethernet PAUSE receive and disable Ethernet PAUSE transmit	Enable Ethernet PAUSE transmit and disable Ethernet PAUSE receive
1	1	Don't care	Don't care	Enable Ethernet PAUSE transmit and receive	Enable Ethernet PAUSE transmit and receive



NOTE: For your convenience, [Table 515](#) replicates Table 28B-3 of Section 2 of the IEEE 802.X specification.

PFC

PFC is a lossless transport and congestion relief feature that works by providing granular link-level flow control for each IEEE 802.1p code point (priority) on a full-duplex Ethernet link. When the receive buffer on a switch interface fills to a threshold, the switch transmits a pause frame to the sender (the connected peer) to temporarily stop the sender from transmitting more frames. The buffer threshold must be low enough so that the sender has time to stop transmitting frames and the receiver can accept the frames already on the wire before the buffer overflows. The switch automatically sets queue buffer thresholds to prevent frame loss.

When congestion forces one priority on a link to pause, all of the other priorities on the link continue to send frames. Only frames of the paused priority are not transmitted. When the receive buffer empties below another threshold, the switch sends a message that starts the flow again.

You configure PFC using a congestion notification profile (CNP). A CNP has two parts:

- **Input**—Specify the code point (or code points) on which to enable PFC, and optionally specify the maximum receive unit (MRU) and the cable length between the interface and the connected peer interface.
- **Output**—Specify the output queue or output queues that respond to pause messages from the connected peer.

You apply a PFC configuration by configuring a CNP on one or more interfaces. Each interface that uses a particular CNP is enabled to pause traffic identified by the priorities (code points) specified in that CNP. You can configure one CNP on an interface, and you can configure different CNPs on different interfaces. When you configure a CNP on an interface, ingress traffic that is mapped to a priority that the CNP enables for PFC is paused whenever the queue buffer fills to the pause threshold. (The pause threshold is not user-configurable.)

Configure PFC for a priority end to end along the entire data path to create a lossless lane of traffic on the network. You can selectively pause the traffic in any queue without pausing the traffic for other queues on the same link. You can create lossless lanes for traffic such as FCoE, LAN backup, or management, while using standard frame-drop congestion management for IP traffic on the same link.

Potential consequences of flow control are:

- Ingress port congestion (configuring too many lossless flows can cause ingress port congestion)
- A paused priority that causes upstream devices to pause the same priority, thus spreading congestion back through the network

By definition, PFC supports symmetric pause only (as opposed to Ethernet PAUSE, which supports symmetric and asymmetric pause). With symmetric pause, a device can:

- Transmit pause frames to pause incoming traffic. (You configure this using the input stanza of a congestion notification profile.)
- Receive pause frames and stop sending traffic to a device whose buffer is too full to accept more frames. (You configure this using the output stanza of a congestion notification profile.)

Receiving a PFC frame from a connected peer pauses traffic on egress queues based on the IEEE 802.1p priorities that the PFC pause frame identifies. The priorities are 0 through 7. By default, the priorities map to queue numbers 0 through 7, respectively, and to specific forwarding classes, as shown in [Table 591](#):

Table 591: Default PFC Priority to Queue and Forwarding Class Mapping

IEEE 802.1p Priority (Code Point)	Queue	Forwarding Class
0 (000)	0	best-effort
1 (001)	1	best-effort
2 (010)	2	best-effort
3 (011)	3	fcoe
4 (100)	4	no-loss
5 (101)	5	best-effort
6 (110)	6	network-control
7 (111)	7	network-control

For example, a received PFC pause frame that pauses priority 3 pauses output queue 3. If you do not want to use the default configuration, you can configure customized mapping of priorities to queues and forwarding classes.



NOTE: By convention, deployments with converged server access typically use IEEE 802.1p priority 3 for FCoE traffic. The default configuration sets the fcoe forwarding class as a lossless forwarding class that is mapped to queue 3. The default classifier maps incoming priority 3 traffic to the fcoe forwarding class. *However, you must apply PFC to the entire FCoE data path to configure the end-to-end lossless behavior that FCoE traffic requires.*

If your network uses priority 3 for FCoE traffic, we recommend that you use the default configuration. If your network uses a priority other than 3 for FCoE traffic, you can configure lossless FCoE transport on any IEEE 802.1p priority as described in [“Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows” on page 6873](#) and [Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway](#).

To enable PFC on a priority:

1. Specify the IEEE 802.1p code point to pause in the input stanza of a CNP.
2. If you are not using the default lossless forwarding classes, specify the IEEE 802.1p code point to pause and the corresponding output queue in the output stanza of the CNP.
3. Apply the CNP to the ingress interfaces on which you want to pause the traffic.
4. If you are not using the default lossless forwarding classes, apply the CNP to the ingress interfaces on which you want to pause the traffic.



CAUTION: Any change to the PFC configuration on a port temporarily blocks the entire port (not just the priorities affected by the PFC change) so that the port can implement the change, then unblocks the port. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

A change to the PFC configuration means any change to a CNP, including changing the input portion of the CNP (enabling or disabling PFC on a priority, or changing the MRU or cable-length values) or changing the output portion of the CNP that enables or disables output flow control on a queue. A PFC configuration change only affects ports that use the changed CNP.

The following actions change the PFC configuration:

- Deleting or disabling a PFC configuration (input or output) in a CNP that is in use on one or more interfaces. For example:
 1. An existing CNP with an input stanza that enables PFC on priorities 3, 5, and 6 is configured on interfaces xe-0/0/20 and xe-0/0/21.
 2. We disable the PFC configuration for priority 6 in the input CNP, and then commit the configuration.

3. The PFC configuration change causes all traffic on interfaces xe-0/0/20 and xe-0/0/21 to stop until the PFC change has been implemented. When the PFC change has been implemented, traffic resumes.

- Configuring a CNP on an interface. (This changes the PFC state by enabling PFC on one or more priorities.)
- Deleting a CNP from an interface. (This changes the PFC state by disabling PFC on one or more priorities.)

When you associate the CNP with an interface, the interface uses PFC to send pause requests when the output queue buffer for the lossless traffic fills to the pause threshold.

On switches that use different classifiers for unicast and multdestination traffic, you can map a unicast queue (queue 0 through 7) and a multdestination queue (queue 8, 9, 10, or 11) to the same IEEE 802.1p code point (priority) so that both unicast and multicast traffic use that priority. However, do not map multdestination traffic to lossless output queues. Starting with Junos OS Release 12.3, you can map one priority to multiple output queues.



NOTE: You can attach a maximum of one CNP to an interface, but you can create an unlimited number of CNPs that explicitly configure only the input stanza and use the default output stanza.

The output stanza of the CNP maps to a profile that interfaces use to respond to pause messages received from the connected peer. On standalone switches, you can create two CNPs with an explicitly configured output stanza.

When a switch is a Node device in a QFabric system, you can create one CNP with an explicitly configured output stanza. (One fewer profile is available on QFabric systems because the system needs a default profile for fabric interfaces, which are not used as fabric interfaces when the switches are not part of a QFabric system. “[Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows](#)” on page 6873 describes configuring output flow control.

Lossless Transport Support Summary

The switch supports up to six lossless forwarding classes. For lossless transport, you must enable PFC on the IEEE 802.1p priorities (code points) mapped to lossless forwarding classes.



CAUTION: Any change to the PFC configuration on a port temporarily blocks the entire port (not just the priorities affected by the PFC change) so that the port can implement the change, then unblocks the port. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

The following limitation applies to support lossless transport on QFabric systems only:

- The internal fiber cable length from the QFabric system Node device to the QFabric system Interconnect device cannot exceed 150 meters.

The default CoS configuration provides two lossless forwarding classes, *fcoe* and *no-loss*. If you explicitly configure lossless forwarding classes, you must include the **no-loss** packet drop attribute to enable lossless behavior, or the traffic is not lossless. For both default and explicit lossless forwarding class configuration, you must configure CNP input stanzas to enable PFC on the priority of the lossless traffic and apply the CNPs to ingress interfaces.



NOTE: The information in this note applies only to systems that do not run the ELS CLI.

Junos OS Release 12.2 introduced changes to the way the switch handles lossless forwarding classes (including the default *fcoe* and *no-loss* forwarding classes).

In Junos OS Release 12.1, either explicitly configuring the *fcoe* and *no-loss* forwarding classes or using the default configuration for these forwarding classes resulted in the same lossless behavior for traffic mapped to those forwarding classes.

However, in Junos OS Release 12.2, if you explicitly configure the *fcoe* or the *no-loss* forwarding class, that forwarding class is no longer treated as a lossless forwarding class. Traffic mapped to these forwarding classes is treated as lossy (best-effort) traffic. This is true even if the explicit configuration is exactly the same as the default configuration.

If your CoS configuration from Junos OS Release 12.1 or earlier includes the explicit configuration of the *fcoe* or the *no-loss* forwarding class, then when you upgrade to Junos OS Release 12.2, those forwarding classes are not lossless. To preserve the lossless treatment of these forwarding classes, delete the the explicit *fcoe* and *no-loss* forwarding class configuration before you upgrade to Junos OS Release 12.2.

See *Overview of CoS Changes Introduced in Junos OS Release 12.2* for detailed information about this change and how to delete an existing lossless configuration.

In Junos OS Release 12.3, the default behavior of the *fcoe* and *no-loss* forwarding classes is the same as in Junos OS Release 12.2. However, in Junos OS Release 12.3, you can configure up to six lossless forwarding classes. All explicitly configured lossless forwarding classes must include the new *no-loss* packet drop attribute or the forwarding class is lossy.

[“Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows” on page 6873](#) provides detailed information about the explicit configuration of lossless priorities and about the

default configuration of lossless priorities, including the input and output stanzas of the CNP.



NOTE: PFC and Ethernet PAUSE are used only on Ethernet interfaces. Fabric (fte) ports on QFabric systems (Node device fabric ports and Interconnect device fabric ports) use link-layer flow control (LLFC) to ensure the appropriate treatment of lossless traffic.

**Related
Documentation**

- [Understanding DCB Features and Requirements on page 6490](#)
- [Understanding CoS Explicit Congestion Notification on page 6820](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 6905](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 6504](#)

Configuring CoS PFC (Congestion Notification Profiles)

A congestion notification profile (CNP) enables priority-based flow control (PFC) on specified IEEE 802.1p priorities (code points). A CNP has two components:

- Input CNP:
 - Enable PFC on a specified priority.
 - Configure the maximum receive unit (MRU) on an interface for traffic that matches the PFC priority (optional).
 - Specify the length of the attached cable on the ingress interface (optional)
- Output CNP (optional): Configure flow control to enable PFC pause on specific output queues for specified priorities.



NOTE: By default, output queues 3 and 4 (which are mapped to default lossless forwarding classes `fcoe` and `no-loss`, respectively) are configured to respond to PFC pause messages received from the connected peer on priorities 3 and 4 (code points 011 and 100, respectively). If you explicitly configure flow control on any output queue, you must configure flow control on every output queue that you want to respond to pause messages. (The explicit configuration overrides the default configuration.)

To achieve lossless behavior, the output queue priorities on which you enable PFC flow control must match the PFC priorities on which you enable PFC on the input interfaces. For example, if you program output queues to pause priorities 3 (011) and 5 (101) in the output component of the CNP, then you must also enable pause on priorities 3 and 5 on the input component of the CNP. (In addition, the forwarding classes mapped to the paused output queues must be lossless forwarding classes.)

Associating a CNP with an interface enables PFC on the ingress traffic that matches the priority specified in the input CNP, and programs the queues listed in the output CNP to pause when the interface receives a PFC pause message from the connected peer. Configure PFC on a priority end to end along the entire data path to create a lossless lane of traffic on the network.



NOTE: You must enable PFC on the priority used by FCoE traffic on ingress interfaces (input CNP). Enable PFC on the FCoE priority on every interface that carries FCoE traffic. By convention, FCoE traffic uses priority 3 (code point 011), which maps to queue 3. If your network uses priority 3 for FCoE traffic, the default forwarding class and classifier configuration support lossless transport, but you must still configure a CNP and apply it to the correct ingress interfaces to enable PFC and achieve lossless transport.

If your network does not use priority 3 for FCoE traffic, you need to configure a classifier that classifies FCoE traffic into a lossless forwarding class, based on the priority your network uses for FCoE traffic. If you are not using the default lossless forwarding class configuration, then you also need to ensure that the output queue mapped to the lossless FCoE forwarding class is programmed to pause.

You can attach only one CNP to an interface. There is no limit to the total number of CNPs you can create.

Configuring a CNP consists of:

- Naming the CNP.
- Specifying the IEEE 802.1 code point (priority) on which you want to enable PFC on ingress interfaces (input CNP).
- Optionally, specifying the MRU and the length of the attached cable on ingress interfaces (input CNP).
- Optionally, configuring flow control (PFC pause) on specified output queues if you want queues other than queues 3 and 4 to respond to pause messages received from the connected peer (output CNP).
- Mapping the CNP to an interface.



NOTE: Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.



NOTE: On QFX5100 and QFX5200, once the headroom buffer is exhausted, any new CNP configuration is not allocated headroom buffer, even if headroom buffer is freed by deletion of an existing CNP. CNP configuration has to be applied again to re-allocate the headroom buffer.

1. Enable PFC on the desired priority in the input CNP and optionally configure the interface MRU for traffic on that priority:

```
[edit class-of-service]
user@switch# set congestion-notification-profile cnp-name input ieee-802.1 code-point
code-point bits pfc mru mru-value
```

For example, to configure a CNP named **fcoe-cnp** that enables PFC on IEEE 802.1 code point **011** and configures an MRU value of **2240**:

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011
pfc mru 2240
```

2. (Optional) Configure the length of the cable attached to the ingress interface:

```
[edit class-of-service]
user@switch# set congestion-notification-profile cnp-name input cable-length
cable-length-value
```

For example, to configure a CNP named **fcoe-cnp** that sets the length of the ingress interface cable to **100** meters:

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe-cnp input cable-length 100
```

3. (Optional) Configure flow control on output queues:

```
[edit class-of-service]
user@switch# set congestion-notification-profile cnp-name output ieee-802.1 code-point
code-point-bits flow-control-queue [queue | list-of-queues]
```

For example, to configure a CNP named **fcoe-cnp** that enables PFC pause flow control on output queues 3 and 5 for FCoE traffic that uses priority 3 (code point **011**) and on output queue 4 for traffic that uses priority 4 (code point **100**):

```
[edit class-of-service]
user@switch# set congestion-notification-profile cnp-name output ieee-802.1 code-point
011 flow-control-queue [3 5]
user@switch# set congestion-notification-profile cnp-name output ieee-802.1 code-point
100 flow-control-queue 4
```

4. Map the CNP to an interface:

```
[edit class-of-service]
user@switch# set interfaces interface congestion-notification-profile cnp-name
```

For example, to map the CNP **fcoe-cnp** to the interface **xe-0/0/7**:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/7 congestion-notification-profile fcoe-cnp
```

Related Documentation

- [Example: Configuring CoS PFC for FCoE Traffic on page 6504](#)
- [Assigning CoS Components to Interfaces on page 6627](#)
- [Monitoring Interfaces That Have CoS Components on page 7161](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 6493](#)
- [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 6873](#)

Example: Configuring CoS PFC for FCoE Traffic

Priority-based flow control (PFC, described in IEEE 802.1Qbb) is a link-level flow control mechanism that you apply at ingress interfaces. PFC enables you to divide traffic on one physical link into eight priorities. You can think of the eight priorities as eight “lanes” of traffic that correspond to queues (forwarding classes). Each priority is mapped to a 3-bit IEEE 802.1p CoS value in the VLAN header.

You can selectively apply PFC to the traffic in any queue without pausing the traffic in other queues on the same link. You must apply PFC to FCoE traffic to ensure lossless transport.

This example describes how to configure PFC for FCoE traffic:

- [Requirements on page 6908](#)
- [Overview on page 6908](#)
- [Configuration on page 6910](#)
- [Verification on page 6915](#)

Requirements

This example uses the following hardware and software components:

- One switch
- Junos OS Release 11.1 or later for the QFX Series

Overview

FCoE traffic requires PFC to ensure lossless packet transport. This example shows you how to configure PFC on FCoE traffic, use the default FCoE forwarding-class-to-queue mapping and:

- Configure a classifier that associates the FCoE forwarding class with FCoE traffic, which is identified by IEEE 802.1p code point 011 (priority 3).
- Configure a congestion notification profile to apply PFC to the FCoE traffic.
- Apply the classifier and the PFC configuration to ingress interfaces.



NOTE: Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

- Configure the CoS bandwidth scheduling for the FCoE forwarding class output queue.
- On switches that support enhanced transmission selection (ETS) hierarchical port scheduling, create a forwarding class set (priority group) that includes the FCoE

forwarding class; this is required to configure enhanced transmission selection (ETS) and support data center bridging (DCB).

- For ETS, configure the bandwidth scheduling for the FCoE priority group.
- Apply the configuration to ingress and egress interfaces. How this is done differs depending on whether you use ETS or direct port scheduling for the CoS configuration.

For direct port scheduling, you apply a scheduler map directly to the interface. A scheduler map maps schedulers to forwarding classes, and applies the CoS properties of the scheduler to the output queue mapped to the forwarding class.

For ETS hierarchical port scheduling, you apply the scheduler map to a traffic control profile, and then apply the traffic control profile to the interface. The scheduler map maps CoS properties to forwarding classes (and their associated output queues) just as it does for direct port scheduling. The traffic control profile maps CoS properties to the priority group (a group of forwarding classes defined in a forwarding class set) that contains the forwarding class, creating a CoS hierarchy that allocates port bandwidth to a group of forwarding classes (priority group), and then allocates the priority group bandwidth to the individual forwarding classes (see [“Understanding CoS Hierarchical Port Scheduling \(ETS\)” on page 6765](#)).

Each interface in this example acts as both an ingress interface and an egress interface, so the classifier, congestion notification profile, and scheduling are applied to all of the interfaces.

Topology

Table 517 shows the configuration components for this example.

Table 592: Components of the PFC for FCoE Traffic Configuration Topology

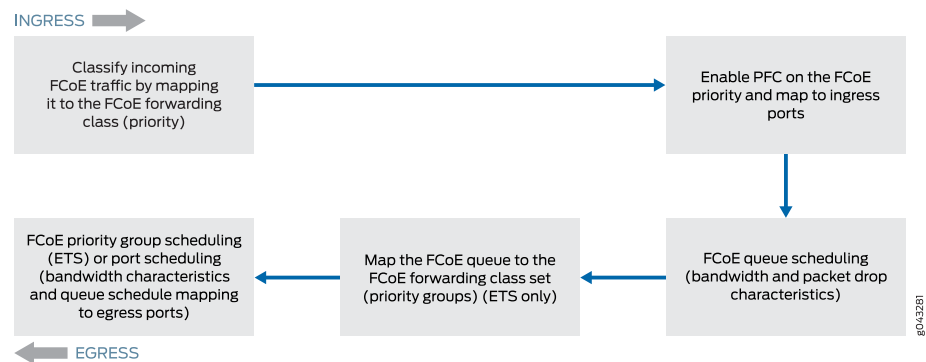
Component	Settings
Hardware	One switch
Behavior aggregate classifier (maps the FCoE forwarding class to incoming packets by IEEE 802.1 code point)	Code point 011 to forwarding class fcoe and loss priority low Ingress interfaces: xe-0/0/31, xe-0/0/32, xe-0/0/33, xe-0/0/34
PFC congestion notification profile	fcoe-cnp: Code point 011 Ingress interfaces: xe-0/0/31, xe-0/0/32, xe-0/0/33, xe-0/0/34
FCoE queue scheduler	fcoe-sched: Minimum bandwidth 3g Maximum bandwidth 100% Priority low

Table 592: Components of the PFC for FCoE Traffic Configuration Topology (*continued*)

Component	Settings
Forwarding class-to-scheduler mapping	<p>Scheduler map fcoe-map:</p> <p>Forwarding class fcoe</p> <p>Scheduler fcoe-sched</p> <p>On switches that support direct port scheduling, if you use port scheduling, attach the scheduler map directly to interfaces xe-0/0/31, xe-0/0/32, xe-0/0/33, and xe-0/0/34.</p>
ETS only: Forwarding class set (FCoE priority group)	<p>fcoe-pg:</p> <p>Forwarding class fcoe</p> <p>Egress interfaces: xe-0/0/31, xe-0/0/32, xe-0/0/33, xe-0/0/34</p>
ETS only: Traffic control profile	<p>fcoe-tcp:</p> <p>Scheduler map fcoe-map</p> <p>Minimum bandwidth 3g</p> <p>Maximum bandwidth 100%</p> <p>For ETS hierarchical scheduling, attach the traffic control profile (using the output-traffic-control-profile keyword) to interfaces xe-0/0/31, xe-0/0/32, xe-0/0/33, and xe-0/0/34.</p>

Figure 207 shows a block diagram of the configuration components and the configuration flow of the CLI statements used in the example.

Figure 229: PFC for FCoE Traffic Configuration Components Block Diagram



Configuration

CLI Quick Configuration

To quickly configure PFC for FCoE traffic, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level. The configuration is separated into the configuration common to ETS and direct port scheduling, and the portions of the configuration that apply only to ETS and only to port scheduling.

Common Configuration (Applies

to ETS Hierarchical Scheduling and to Port Scheduling

```
[edit class-of-service]
set classifiers ieee-802.1 fcoe-classifier forwarding-class fcoe loss-priority low code-points 011
set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011 pfc
set interfaces xe-0/0/31 unit 0 classifiers ieee-802.1 fcoe-classifier
set interfaces xe-0/0/32 unit 0 classifiers ieee-802.1 fcoe-classifier
set interfaces xe-0/0/33 unit 0 classifiers ieee-802.1 fcoe-classifier
set interfaces xe-0/0/34 unit 0 classifiers ieee-802.1 fcoe-classifier
set interfaces xe-0/0/31 congestion-notification-profile fcoe-cnp
set interfaces xe-0/0/32 congestion-notification-profile fcoe-cnp
set interfaces xe-0/0/33 congestion-notification-profile fcoe-cnp
set interfaces xe-0/0/34 congestion-notification-profile fcoe-cnp
set schedulers fcoe-sched priority low transmit-rate 3g
set schedulers fcoe-sched shaping-rate percent 100
set scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
```

Configuration for ETS Hierarchical Scheduling

The ETS-specific portion of this example configures forwarding class set (priority group) membership, priority group CoS settings (traffic control profile), and assigns the priority group and its CoS configuration to the interfaces.

```
[edit class-of-service]
set forwarding-class-sets fcoe-pg class fcoe
set traffic-control-profiles fcoe-tcp scheduler-map fcoe-map guaranteed-rate 3g
set traffic-control-profiles fcoe-tcp shaping-rate percent 100
set interfaces xe-0/0/31 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
set interfaces xe-0/0/32 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
set interfaces xe-0/0/33 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
set interfaces xe-0/0/34 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
```

Configuration for Port Scheduling

The port-scheduling-specific portion of this example assigns the scheduler map (which sets the CoS treatment of the forwarding classes in the scheduler map) to the interfaces.

```
[edit class-of-service]
set interfaces xe-0/0/31 scheduler-map fcoe-map
set interfaces xe-0/0/32 scheduler-map fcoe-map
set interfaces xe-0/0/33 scheduler-map fcoe-map
set interfaces xe-0/0/34 scheduler-map fcoe-map
```

Common Configuration (Applies to ETS Hierarchical Scheduling and to Port Scheduling)

Step-by-Step Procedure

To configure the ingress classifier for FCoE traffic, PFC on the FCoE traffic, apply the PFC and classifier configurations to interfaces, and configure queue scheduling, for both ETS hierarchical scheduling and port scheduling (common configuration):

1. Configure a classifier to set the loss priority and IEEE 802.1 code point assigned to the FCoE forwarding class at the ingress:

```
[edit class-of-service]
user@switch# set classifiers ieee-802.1 fcoe-classifier forwarding-class fcoe loss-priority low code-points 011
```
2. Configure PFC on the FCoE queue by applying FCoE to the IEEE 802.1 code point 011:

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011 pfc
```
3. Apply the PFC configuration to the ingress interfaces:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/31 congestion-notification-profile fcoe-cnp
user@switch# set interfaces xe-0/0/32 congestion-notification-profile fcoe-cnp
user@switch# set interfaces xe-0/0/33 congestion-notification-profile fcoe-cnp
user@switch# set interfaces xe-0/0/34 congestion-notification-profile fcoe-cnp
```
4. Assign the classifier to the ingress interfaces:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/31 unit 0 classifiers ieee-802.1 fcoe-classifier
user@switch# set interfaces xe-0/0/32 unit 0 classifiers ieee-802.1 fcoe-classifier
user@switch# set interfaces xe-0/0/33 unit 0 classifiers ieee-802.1 fcoe-classifier
user@switch# set interfaces xe-0/0/34 unit 0 classifiers ieee-802.1 fcoe-classifier
```
5. Configure output scheduling for the FCoE queue:

```
[edit class-of-service]
user@switch# set schedulers fcoe-sched priority low transmit-rate 3g
user@switch# set schedulers fcoe-sched shaping-rate percent 100
```
6. Map the FCoE forwarding class to the FCoE scheduler:

```
[edit class-of-service]
user@switch# set scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
```

ETS Hierarchical Scheduling Configuration

Step-by-Step Procedure

To configure the forwarding class set (priority group) and priority group scheduling (in a traffic control profile), and apply the ETS hierarchical scheduling for FCoE traffic to interfaces:

1. Configure the forwarding class set for the FCoE traffic:

```
[edit class-of-service]
user@switch# set forwarding-class-sets fcoe-pg class fcoe
```
2. Define the traffic control profile for the FCoE forwarding class set:


```
[edit class-of-service]
user@switch# set traffic-control-profiles fcoe-tcp scheduler-map fcoe-map
guaranteed-rate 3g
user@switch# set traffic-control-profiles fcoe-tcp shaping-rate percent 100
```

3. Apply the FCoE forwarding class set and traffic control profile to the egress ports:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/31 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
user@switch# set interfaces xe-0/0/32 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
user@switch# set interfaces xe-0/0/33 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
user@switch# set interfaces xe-0/0/34 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
```

Port Scheduling Configuration

Step-by-Step Procedure

To apply port scheduling for FCoE traffic to interfaces:

1. Apply the scheduler map to the egress ports:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/31 scheduler-map fcoe-map
user@switch# set interfaces xe-0/0/32 scheduler-map fcoe-map
user@switch# set interfaces xe-0/0/33 scheduler-map fcoe-map
user@switch# set interfaces xe-0/0/34 scheduler-map fcoe-map
```

Results

Display the results of the configuration (the system shows only the explicitly configured parameters; it does not show default parameters such as the **fcoe** lossless forwarding class). The results are from the ETS hierarchical scheduling configuration to show the more complex configuration. Direct port scheduling results would not show the traffic control profile or forwarding class set portions of the configuration, and would display the name of the scheduler map under each interface (instead of the names of the forwarding class set and output traffic control profile), but is otherwise the same.

```
user@switch> show configuration class-of-service
classifiers {
  ieee-802.1 fcoe-classifier {
    forwarding-class fcoe {
      loss-priority low code-points 011;
    }
  }
}
traffic-control-profiles {
  fcoe-tcp {
    scheduler-map fcoe-map;
    shaping-rate percent 100;
    guaranteed-rate 3000000000;
  }
}
forwarding-class-sets {
  fcoe-pg {
    class fcoe;
  }
}
```

```
}
congestion-notification-profile {
  fcoe-cnp {
    input {
      ieee-802.1 {
        code-point 011 {
          pfc;
        }
      }
    }
  }
}
}
interfaces {
  xe-0/0/31 {
    congestion-notification-profile fcoe-cnp;
    forwarding-class-set {
      fcoe-pg {
        output-traffic-control-profile fcoe-tcp;
      }
    }
    unit 0 {
      classifiers {
        ieee-802.1 fcoe-classifier;
      }
    }
  }
  xe-0/0/32 {
    congestion-notification-profile fcoe-cnp;
    forwarding-class-set {
      fcoe-pg {
        output-traffic-control-profile fcoe-tcp;
      }
    }
    unit 0 {
      classifiers {
        ieee-802.1 fcoe-classifier;
      }
    }
  }
  xe-0/0/33 {
    congestion-notification-profile fcoe-cnp;
    forwarding-class-set {
      fcoe-pg {
        output-traffic-control-profile fcoe-tcp;
      }
    }
    unit 0 {
      classifiers {
        ieee-802.1 fcoe-classifier;
      }
    }
  }
  xe-0/0/34 {
    congestion-notification-profile fcoe-cnp;
    forwarding-class-set {
      fcoe-pg {
```

```

        output-traffic-control-profile fcoe-tcp;
    }
}
unit 0 {
    classifiers {
        ieee-802.1 fcoe-classifier;
    }
}
}
}
scheduler-maps {
    fcoe-map {
        forwarding-class fcoe scheduler fcoe-sched;
    }
}
schedulers {
    fcoe-sched {
        transmit-rate 3000000000;
        shaping-rate percent 100;
        priority low;
    }
}
}

```



TIP: To quickly configure the interfaces, issue the **load merge terminal** command and then copy the hierarchy and paste it into the switch terminal window.

Verification

To verify that the PFC configuration for FCoE traffic components has been created and is operating properly, perform these tasks:

- [Verifying That Priority-Based Flow Control Has Been Enabled on page 6915](#)
- [Verifying the Ingress Interface PFC Configuration on page 6916](#)

Verifying That Priority-Based Flow Control Has Been Enabled

Purpose Verify that PFC is enabled on the FCoE queue to enable lossless transport.

Action List the congestion notification profiles using the operational mode command **show class-of-service congestion-notification**:

```

user@switch> show class-of-service congestion-notification
Type: Input, Name: fcoe-cnp, Index: 51697
Cable Length: 100 m

```

Priority	PFC	MRU
000	Disabled	
001	Disabled	
010	Disabled	
011	Enabled	2500
100	Disabled	
101	Disabled	

110	Disabled
111	Disabled
Type: Output	
Priority	Flow-Control-Queues
000	0
001	0
010	1
011	2
100	3
101	4
110	5
111	6
	7

Meaning The **show class-of-service congestion-notification** operational command lists all of the congestion notification profiles and which IEEE 802.1p code points have PFC enabled. The command output shows that PFC is enabled on code point **011** for the **fcoe-cnp** congestion notification profile.

The command also shows the default cable length (100 meters), the default maximum receive unit (2500 bytes), and the default mapping of priorities to output queues because this example does not include configuring these options.

Verifying the Ingress Interface PFC Configuration

Purpose Verify that the classifier **fcoe-classifier** and the congestion notification profile **fcoe-cnp** are configured on ingress interfaces **xe-0/0/31**, **xe-0/0/32**, **xe-0/0/33**, and **xe-0/0/34**.

Action List the ingress interfaces using the operational mode command **show configuration class-of-service interfaces**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/31
congestion-notification-profile fcoe-cnp;
unit 0 {
    classifiers {
        ieee-802.1 fcoe-classifier;
    }
}

user@switch> show configuration class-of-service interfaces xe-0/0/32
congestion-notification-profile fcoe-cnp;
unit 0 {
    classifiers {
        ieee-802.1 fcoe-classifier;
    }
}

user@switch> show configuration class-of-service interfaces xe-0/0/33
congestion-notification-profile fcoe-cnp;
unit 0 {
```

```

        classifiers {
            ieee-802.1 fcoe-classifier;
        }
    }

user@switch> show configuration class-of-service interfaces xe-0/0/34
congestion-notification-profile fcoe-cnp;
unit 0 {
    classifiers {
        ieee-802.1 fcoe-classifier;
    }
}

```

Meaning The **show configuration class-of-service interfaces** commands list the congestion notification profile that is mapped to the interface (**fcoe-cnp**) and the IEEE 802.1p classifier associated with the interface (**fcoe-classifier**).

Related Documentation

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 6771](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 6905](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 6493](#)

Troubleshooting Dropped FCoE Traffic

Problem **Description:** Fibre Channel over Ethernet (FCoE) traffic for which you want guaranteed delivery is dropped.

Cause There are several possible causes of dropped FCoE traffic (the list numbers of the possible causes correspond to the list numbers of the solutions in the *Solution* section.):

1. Priority-based flow control (PFC) is not enabled on the FCoE priority (IEEE 802.1p code point) in both the input and output stanzas of the congestion notification profile.
2. The FCoE traffic is not classified correctly at the ingress interface. FCoE traffic should either use the default **fcoe** forwarding class and classifier configuration (maps the **fcoe** forwarding class to IEEE 802.1p code point 011) or be mapped to a lossless forwarding class and to the code point enabled for PFC on the input and output interfaces.
3. The congestion notification profile that enables PFC on the FCoE priority is not attached to the interface.
4. The forwarding class set (priority group) used for guaranteed delivery traffic does not include the forwarding class used for FCoE traffic.



NOTE: This issue can occur only on switches that support enhanced transmission selection (ETS) hierarchical port scheduling. (Direct port scheduling does not use forwarding class sets.)

5. Insufficient bandwidth has been allocated for the FCoE queue or for the forwarding class set to which the FCoE queue belongs.



NOTE: This issue can occur for forwarding class sets only on switches that support ETS hierarchical port scheduling. (Direct port scheduling does not use forwarding class sets.)

6. If you are using Junos OS Release 12.2, the **fcoe** forwarding class has been explicitly configured instead of using the default **fcoe** forwarding class configuration (forwarding-class-to-queue mapping).



NOTE: If you are using Junos OS Release 12.2, use the default forwarding-class-to-queue mapping for the lossless **fcoe** and **no-loss** forwarding classes. If you explicitly configure the lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best effort) traffic and does *not* receive lossless treatment.

7. If you are using Junos OS Release 12.3 or later and you are not using the default **fcoe** forwarding class configuration, the forwarding class used for FCoE is not configured with the **no-loss** packet drop attribute. In Junos OS 12.3 or later, explicit forwarding classes configurations must include the **no-loss** packet drop attribute to be treated as lossless forwarding classes.

Solution The list numbers of the possible solutions correspond to the list numbers of the causes in the *Cause* section.

1. Check the congestion notification profile (CNP) to see if PFC is enabled on the FCoE priority (the correct IEEE 802.1p code point) on both input and output interfaces. Use the **show class-of-service congestion-notification** operational command to show the code points that are enabled for PFC in each CNP.

If you are using the default configuration, FCoE traffic is mapped to code point 011 (priority 3). In this case, the input stanza of the CNP should show that PFC is enabled on code point 011, and the output stanza should show that priority 011 is mapped to flow control queue 3.

If you explicitly configured a forwarding class for FCoE traffic, ensure that:

- You specified the **no-loss** packet drop attribute in the forwarding class configuration
- The code point mapped to the FCoE forwarding class in the ingress classifier is the code point enabled for PFC in the CNP input stanza
- The code point and output queue used for FCoE traffic are mapped to each other in the CNP output stanza (if you are not using the default priority and queue, you must explicitly configure each output queue that you want to respond to PFC messages)

For example, if you explicitly configure a forwarding class for FCoE traffic that is mapped to output queue 5 and to code point 101 (priority 5), the output of the **show class-of-service congestion-notification** looks like:

```
Name: fcoe_p5_cnp, Index: 12183
Type: Input
Cable Length: 100 m
  Priority  PFC      MRU
  000      Disabled
  001      Disabled
  010      Disabled
  011      Disabled
  100      Disabled
  101      Enabled   2500
  110      Disabled
  111      Disabled
Type: Output
  Priority  Flow-Control-Queues
  101      5
```

2. Use the **show class-of-service classifier type ieee-802.1p** operational command to check if the classifier maps the forwarding class used for FCoE traffic to the correct IEEE 802.1p code point.
3. Ensure that the congestion notification profile and classifier are attached to the correct ingress interface. Use the operational command **show configuration class-of-service interfaces interface-name**.
4. Check that the forwarding class set includes the forwarding class used for FCoE traffic. Use the operational command **show configuration class-of-service forwarding-class-sets** to show the configured priority groups and their forwarding classes.

5. Verify the amount of bandwidth allocated to the queue mapped to the FCoE forwarding class and to the forwarding class set to which the FCoE traffic queue belongs. Use the **show configuration class-of-service schedulers *scheduler-name*** operational command (specify the scheduler for FCoE traffic as the *scheduler-name*) to see the minimum guaranteed bandwidth (**transmit-rate**) and maximum bandwidth (**shaping-rate**) for the queue.

Use the **show configuration class-of-service traffic-control-profiles *traffic-control-profile*** operational command (specify the traffic control profile used for FCoE traffic as the *traffic-control-profile*) to see the minimum guaranteed bandwidth (**guaranteed-rate**) and maximum bandwidth (**shaping-rate**) for the forwarding class set.

6. Delete the explicit FCoE forwarding-class-to-queue mapping so that the system uses the default FCoE forwarding-class-to-queue mapping. Include the **delete forwarding-classes class fcoe queue-num 3** statement at the **[edit class-of-service]** hierarchy level to remove the explicit configuration. The system then uses the default configuration for the FCoE forwarding class and preserves the lossless treatment of FCoE traffic.
7. Use the **show class-of-service forwarding-class** operational command to display the configured forwarding classes. The *No-Loss* column shows whether lossless transport is enabled or disabled for each forwarding class. If the forwarding class used for FCoE traffic is not enabled for lossless transport, include the **no-loss** packet drop attribute in the forwarding class configuration (**set class-of-service forwarding-classes class *fcoe-forwarding-class-name* queue-num *queue-number* no-loss**).

See “[Example: Configuring CoS PFC for FCoE Traffic](#)” on page 6504 for step-by-step instructions on how to configure PFC for FCoE traffic, including classifier, interface, congestion notification profile, PFC, and bandwidth scheduling configuration.

Related Documentation

- [show class-of-service congestion-notification on page 7394](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 6905](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 6504](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 6493](#)

Understanding PFC Functionality Across Layer 3 Interfaces

Priority-based flow control (PFC) allows you to select traffic flows within a link and pause them, so that the output queues associated with the flows do not overflow and drop packets. (PFC is more granular than Ethernet PAUSE, which pauses all traffic on a physical link.) PFC helps you configure lossless transport for traffic flows across a data center bridging network.

However, you might want to create a traffic flow that losslessly traverses the Layer 2 data center bridging network *and* also losslessly traverses a Layer 3 network that connects Ethernet hosts in different Layer 2 networks. On a QFX5200, QFX5100, EX4600, or QFX10000 switch running the Enhanced Layer 2 Software (ELS) CLI, in addition to configuring PFC on Layer 2 (bridging) interfaces, you can configure PFC on traffic that traverses Layer 3 interfaces. This enables you to preserve the lossless characteristics that PFC provides on traffic, even when the traffic crosses Layer 3 interfaces that connect two Layer 2 networks.



Video: [Preserving Lossless Behavior on an SDN or Overlay Network](#)

PFC works the same way across Layer 3 interfaces as it works across Layer 2 interfaces. When an output queue buffer reaches a certain fill level threshold, the switch sends a PFC pause message to the connected peer to pause transmission of the traffic on which PFC is enabled. Pausing the incoming traffic prevents the queue buffer from overflowing and dropping packets, just as on Layer 2 interfaces. When the queue buffer fill level decreases below a certain threshold, the interface sends a message to the connected peer to restart traffic transmission.

Although PFC is a data center bridging technology, PFC also works on Layer 3 interfaces because PFC operates at the queue level. When you use an IEEE 802.1p classifier to classify incoming traffic (map incoming traffic to a forwarding class and a loss priority based on the IEEE 802.1p code point in the Ethernet frame header) and you enable PFC on the appropriate priority (IEEE 802.1p code point), PFC works on Layer 2 and Layer 3 interfaces.



NOTE: Lossless traffic on Layer 3 interfaces *must* use an IEEE 802.1p classifier to classify incoming traffic, because PFC does not use DSCP or DSCP IPv6 code points to identify traffic for flow control. PFC cannot pause traffic flows unless the incoming traffic is classified by an IEEE 802.1p classifier. Do not apply a DSCP (or a DSCP IPv6) classifier to Layer 3 traffic on which you want to enable PFC.

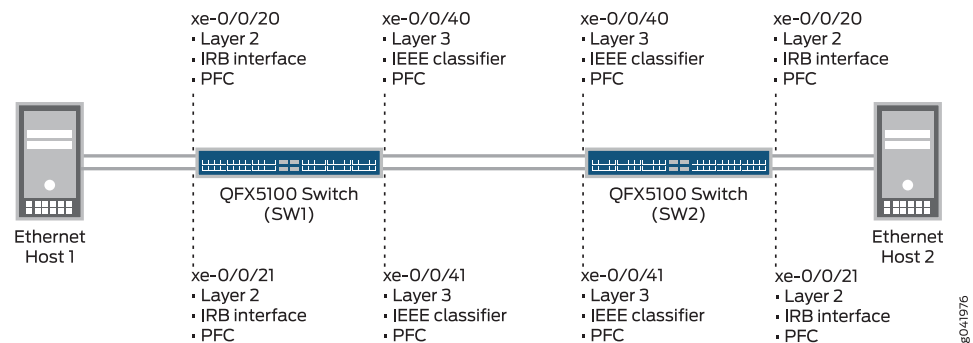
Because PFC functionality relies on the mapping (classifying) of incoming traffic to IEEE 802.1p code points and on enabling PFC on the correct code point(s) at each interface, you must ensure that incoming traffic has the correct 3-bit IEEE 802.1p code point (priority) in the priority code point (PCP) field of the Ethernet frame header (sometimes known as the CoS bits).



NOTE: Layer 3 interfaces do not support FCoE traffic. FCoE traffic must use Layer 2 interfaces and cannot use Layer 3 interfaces. Therefore, you cannot enable PFC on FCoE traffic across Layer 3 interfaces.

Figure 230 shows a topology in which two Ethernet hosts in Layer 2 networks communicate across a Layer 3 network, with PFC enabled on all of the Layer 2 and Layer 3 switch interfaces.

Figure 230: Enabling PFC Across Layer 3 Interface Hops



The Ethernet host-facing interfaces (`xe-0/0/20` and `xe-0/0/21` on both switches) and the Layer 3 network-facing interfaces (interfaces `xe-0/0/40` and `xe-0/0/41` on both switches) require different interface configurations to enable PFC on the Layer 3 interfaces. In addition, the class of service (CoS) for each interface must be configured correctly, including enabling PFC on the traffic that you want to treat as lossless traffic:

Ethernet-host facing interfaces (`xe-0/0/20` and `xe-0/0/21`) require the following configuration:

- Set interfaces as family ethernet-switching
- Set the interface mode as trunk mode
- Create VLANs to carry the traffic
- Create IRB interfaces to place the Layer 2 VLAN traffic on Layer 3 for transport between IP networks
- Create an IEEE 802.1p classifier to classify incoming traffic into the correct forwarding class, based on the IEEE 802.1p code point
- Create a congestion notification profile (CNP) to configure PFC on the IEEE 802.1p code point of the traffic that you want treat as lossless traffic
- Apply the classifier and the CNP to the Layer 2 interfaces
- Configure CoS: lossless forwarding classes, hierarchical port scheduling (also known as enhanced transmission selection), or direct port scheduling, depending on your switch, and apply it to the Layer 2 interfaces

Layer 3 IP network-facing interfaces (xe-0/0/40 and xe-0/0/41) require the following configuration:

- Set interfaces as family inet
- Set VLAN tagging on the interfaces
- Create VLANs to carry the traffic
- Create an IEEE 802.1p classifier to classify incoming traffic into the correct forwarding class, based on the IEEE 802.1p code point (do not use a DSCP or DSCP IPv6 classifier)
- Create a congestion notification profile (CNP) to configure PFC on the IEEE 802.1p code point of the traffic that you want treat as lossless traffic on the Layer 3 interfaces
- Apply the IEEE 802.1p classifier and the CNP to the Layer 3 interfaces
- Configure CoS: lossless forwarding classes, hierarchical port scheduling (enhanced transmission selection), or direct port scheduling, depending on your switch, and apply it to the Layer 3 interfaces



NOTE: Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

When you configure the Layer 2 and Layer 3 interfaces correctly, the switch enables PFC on the traffic between Ethernet Host 1 and Ethernet Host 2 across the entire path between the two hosts. If any output queue in the path on which PFC is enabled experiences congestion, PFC pauses the traffic and prevents packet loss for the flow.

Related Documentation

- [Example: Configuring PFC Across Layer 3 Interfaces on page 6923](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 6493](#)
- [Understanding Integrated Routing and Bridging on page 2124](#) (Topic also applies to IRB interfaces.)

Example: Configuring PFC Across Layer 3 Interfaces

Priority-based flow control (PFC) helps ensure lossless transport across data center bridging interfaces by pausing incoming traffic when output queue buffers fill to a certain threshold. On a QFX5200, QFX5100, EX4600, or QFX10000 switch running the Enhanced Layer 2 Software (ELS) CLI, in addition to configuring PFC on Layer 2 (bridging) interfaces, you can configure PFC on traffic that traverses Layer 3 interfaces. This enables you to preserve the lossless characteristics that PFC provides on traffic, even when the traffic crosses Layer 3 interfaces that connect two Layer 2 networks.

- [Requirements on page 6924](#)
- [Overview on page 6924](#)

- [Configuration on page 6927](#)
- [Verification on page 6936](#)

Requirements

This example uses the following hardware and software components:

- Two switches
- Junos OS Release 13.2 or later for the QFX Series
- Two Ethernet hosts

Overview

On a network that uses two QFX5200, QFX5100, EX4600, or QFX10000 switches to connect hosts on two different Ethernet networks across a Layer 3 network, to configure PFC across the Layer 2 and Layer 3 interfaces, you must:

- Configure the Layer 2 and Layer 3 interfaces on the switches
- Configure VLANs to carry the traffic across the Layer 2 and Layer 3 networks
- Configure integrated routing and bridging (IRB) interfaces on the Layer 2 interfaces to move the Layer 2 VLAN traffic to Layer 3
- Configure and apply the appropriate classifiers to the interfaces
- Configure and apply congestion notification profiles (CNPs) on the interfaces to enable PFC on the traffic that you want to be lossless



NOTE: Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

- Configure lossless forwarding classes and either hierarchical port scheduling (also known as enhanced transmission selection) or direct port scheduling, depending on your switch, on the interfaces



NOTE: PFC operates at the queue level, based on the IEEE 802.1p code point in the priority code point (PCP) field of the Ethernet frame header (sometimes known as the CoS bits). For this reason, traffic on Layer 3 interfaces on which you want to enable PFC must use an IEEE 802.1p classifier to map incoming traffic to forwarding classes (which are in turn mapped to output queues) and loss priorities. You cannot use a DSCP or DSCP IPv6 classifier to classify Layer 3 traffic if you want to enable PFC on traffic flows.

Topology

Figure 231 shows the topology for this example.

Figure 231: Enabling PFC Across Layer 3 Interface Hops

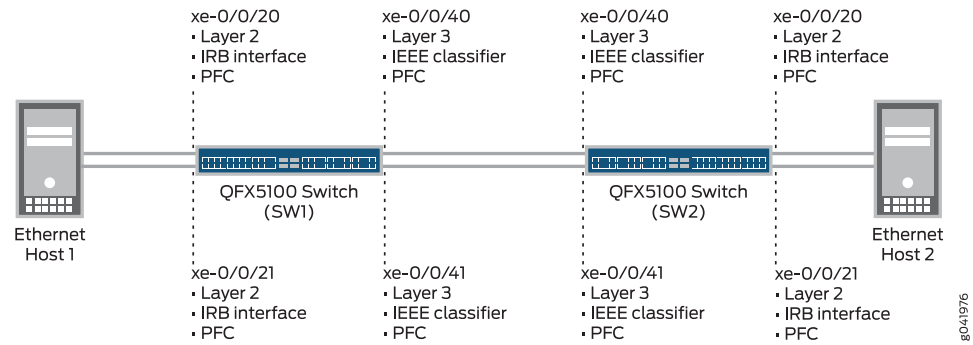


Table 593 shows the configuration components for this example. On the two switches, the Ethernet host-facing interfaces use the same interface names and configuration, and the Layer 3 network-facing interfaces use the same interface names and configuration.

Table 593: Components of the PFC Across Layer 3 Interfaces Topology

Component	Settings
Hardware	Two switches, Switch SW1 and Switch SW2. Two Ethernet hosts
Layer 3 interfaces (xe-0/0/40 and xe-0/0/41) and VLANs	Interface xe-0/0/40: <ul style="list-style-type: none"> Interface family—inet Interface IP address—100.103.1.2/24 VLAN tagging—enabled Interface VLAN ID—103 Interface xe-0/0/41: <ul style="list-style-type: none"> Interface family—inet Interface IP address—100.104.1.2/24 VLAN tagging—enabled Interface VLAN ID—104
Layer 2 interfaces (xe-0/0/20 and xe-0/0/21) and VLAN membership	Family: Ethernet switching Interface mode—trunk Interface xe-0/0/20 VLAN membership—vlan105 Interface xe-0/0/21 VLAN membership—vlan106
VLANs for the IRB interfaces	VLAN unit 105—family inet, IP address 100.105.1.1/24 VLAN unit 106—family inet, IP address 100.106.1.1/24

Table 593: Components of the PFC Across Layer 3 Interfaces Topology (continued)

Component	Settings
Layer 2 IRB interfaces	<p>Interface xe-0/0/20:</p> <ul style="list-style-type: none"> • IRB interface unit—105 • IRB interface family—inert • IRB interface IP address—100.105.1.1/24 • IRB interface VLAN ID—105 • Layer 3 interface name—irb.105 <p>Interface xe-0/0/21:</p> <ul style="list-style-type: none"> • IRB interface unit—106 • IRB interface family—inert • IRB interface IP address—100.106.1.1/24 • IRB interface VLAN ID—106 • Layer 3 interface name—irb.106
Forwarding classes (both switches)	<p>Name—lossless-3 Queue mapping—queue 3 Packet drop attribute—no-loss</p> <p>Name—lossless-4 Queue mapping—queue 4 Packet drop attribute—no-loss</p> <p>NOTE: Matching the forwarding class names (lossless-3 and lossless-4) to the queue number and to the classified IEEE 802.1p code point (priority) creates a configuration that is logical and easy to map because the forwarding class, queue, and priority all use the same number.</p> <p>Name—all-others Queue mapping—queue 0 Packet drop attribute—none</p> <p>NOTE: The forwarding class <i>all-others</i> is for best-effort traffic that traverses the interfaces.</p>
Layer 2 interface behavior aggregate (BA) classifier	<p>Name—lossless-3-4-ieee Forwarding class lossless-3—mapped to code point 011 (IEEE 802.1p priority 3) and a packet loss priority of low Forwarding class lossless-4—mapped to code point 100 (IEEE 802.1p priority 4) and a packet loss priority of low</p> <p>Apply the Layer 2 IEEE 802.1p classifier to both the Layer 2 and the Layer 3 interfaces (xe-0/0/20, xe-0/0/21, xe-0/0/40, and xe-0/0/41).</p>
Congestion notification profile (PFC, both switches)	<p>Name—lossless-cnp PFC enabled on IEEE 802.1p code points—011 (lossless-3 forwarding class and priority), 100 (lossless-4 forwarding class and priority)</p> <p>Apply the CNP to both the Layer 2 and the Layer 3 interfaces (xe-0/0/20, xe-0/0/21, xe-0/0/40, and xe-0/0/41) to enable PFC on IEEE 802.1p code points 011 and 100.</p>

Table 593: Components of the PFC Across Layer 3 Interfaces Topology (*continued*)

Component	Settings
Enhanced transmission selection (ETS) hierarchical port scheduling (only if using ETS)	<p>Hierarchical port scheduling (ETS) includes configuring:</p> <ul style="list-style-type: none"> • Schedulers to assign bandwidth to traffic • Scheduler mapping to forwarding classes • Grouping of the forwarding classes (priorities) in forwarding class sets (priority groups) • A traffic control profile to assign bandwidth to the forwarding class set and to associate the forwarding class set with the scheduler mapping <p>Hierarchical port scheduling also includes applying the hierarchical scheduler (defined in the traffic control profile) to the interfaces.</p> <p>This example focuses on configuring PFC across the Layer 2 and Layer 3 interfaces. To maintain this focus, this example includes the CLI statements needed to configure hierarchical port scheduling, but does not include descriptive explanations of the configuration. The <i>Related Documentation</i> section provides links to example documents that show how to configure hierarchical port scheduling.</p> <p>Apply the scheduling configuration to both the Layer 2 and the Layer 3 interfaces (xe-0/0/20, xe-0/0/21, xe-0/0/40, and xe-0/0/41).</p>
Direct port scheduling (only if using port scheduling instead of ETS)	<p>Direct port scheduling includes configuring:</p> <ul style="list-style-type: none"> • Schedulers to assign bandwidth to traffic • Scheduler mapping to forwarding classes <p>Port scheduling also includes applying the scheduler map to the interfaces.</p> <p>This example focuses on configuring PFC across the Layer 2 and Layer 3 interfaces. To maintain this focus, this example includes the CLI statements needed to configure direct port scheduling, but does not include descriptive explanations of the configuration. The <i>Related Documentation</i> section provides links to example documents that show how to configure port scheduling.</p> <p>Apply the scheduling configuration to both the Layer 2 and the Layer 3 interfaces (xe-0/0/20, xe-0/0/21, xe-0/0/40, and xe-0/0/41).</p>

Configuration

- [Common Configuration \(Applies to ETS Hierarchical Scheduling and to Port Scheduling\) on page 6930](#)
- [ETS Hierarchical Scheduling Configuration on page 6932](#)
- [Port Scheduling Configuration on page 6932](#)
- [Results on page 6932](#)

CLI Quick Configuration To configure PFC across Layer 3 interfaces, copy the following commands, paste them in a text file, remove the line breaks, change variables and details to match your network

configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level. The same configuration applies to both Switch SW1 and Switch SW2. The configuration is separated into the configuration common to ETS and direct port scheduling, and the portions of the configuration that apply only to ETS and only to port scheduling.

**Common
Configuration (Applies
to ETS Hierarchical
Scheduling and to Port
Scheduling)**

```
set interfaces xe-0/0/40 vlan-tagging
set interfaces xe-0/0/40 unit 0 vlan-id 103
set interfaces xe-0/0/40 unit 0 family inet address 100.103.1.2/24
set interfaces xe-0/0/41 vlan-tagging
set interfaces xe-0/0/41 unit 0 vlan-id 104
set interfaces xe-0/0/41 unit 0 family inet address 100.104.1.2/24
set interfaces xe-0/0/20 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/20 unit 0 family ethernet-switching vlan members vlan105
set interfaces xe-0/0/21 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/21 unit 0 family ethernet-switching vlan members vlan106
set interfaces irb unit 105 family inet address 100.105.1.1/24
set interfaces irb unit 106 family inet address 100.106.1.1/24
set vlans vlan105 vlan-id 105
set vlans vlan106 vlan-id 106
set vlans vlan105 l3-interface irb.105
set vlans vlan106 l3-interface irb.106
set class-of-service forwarding-classes class lossless-3 queue-num 3 no-loss
set class-of-service forwarding-classes class lossless-4 queue-num 4 no-loss
set class-of-service forwarding-classes class all-others queue-num 0
set class-of-service classifiers ieee-802.1 lossless-3-4-ieee forwarding-class lossless-3 loss-priority
low code-points 011
set class-of-service classifiers ieee-802.1 lossless-3-4-ieee forwarding-class lossless-4 loss-priority
low code-points 100
set class-of-service congestion-notification-profile lossless-cnp input ieee-802.1 code-point 011
pfc
set class-of-service congestion-notification-profile lossless-cnp input ieee-802.1 code-point 100
pfc
set class-of-service schedulers lossless_sch transmit-rate 6g
set class-of-service schedulers lossless_sch shaping-rate percent 100
set class-of-service schedulers all-others_sch transmit-rate 4g
set class-of-service scheduler-maps lossless_map forwarding-class lossless-3 scheduler
lossless_sch
set class-of-service scheduler-maps lossless_map forwarding-class lossless-4 scheduler
lossless_sch
set class-of-service scheduler-maps all-others_map forwarding-class all-others scheduler
all-others_sch
set class-of-service interfaces xe-0/0/20 congestion-notification-profile lossless-cnp
set class-of-service interfaces xe-0/0/20 unit 0 classifiers ieee-802.1 lossless-3-4-ieee
set class-of-service interfaces xe-0/0/21 congestion-notification-profile lossless-cnp
set class-of-service interfaces xe-0/0/21 unit 0 classifiers ieee-802.1 lossless-3-4-ieee
set class-of-service interfaces xe-0/0/40 congestion-notification-profile lossless-cnp
set class-of-service interfaces xe-0/0/40 classifiers ieee-802.1 lossless-3-4-ieee
set class-of-service interfaces xe-0/0/41 congestion-notification-profile lossless-cnp
set class-of-service interfaces xe-0/0/41 classifiers ieee-802.1 lossless-3-4-ieee
```


Configuration for ETS Hierarchical Scheduling

The ETS-specific portion of this example configures forwarding class set (priority group) membership and priority group CoS settings (traffic control profile), and assigns the priority group and its CoS configuration to the interfaces.

```
set class-of-service forwarding-class-sets lossless_fc_set class lossless-3
set class-of-service forwarding-class-sets lossless_fc_set class lossless-4
set class-of-service forwarding-class-sets all-others_fc_set class all-others
set class-of-service traffic-control-profiles lossless_tcp scheduler-map lossless_map
set class-of-service traffic-control-profiles lossless_tcp guaranteed-rate percent 60
set class-of-service traffic-control-profiles lossless_tcp shaping-rate percent 100
set class-of-service traffic-control-profiles all-others_tcp scheduler-map all-others_map
set class-of-service traffic-control-profiles all-others_tcp guaranteed-rate percent 40
set class-of-service interfaces xe-0/0/20 forwarding-class-set lossless_fc_set
output-traffic-control-profile lossless_tcp
set class-of-service interfaces xe-0/0/20 forwarding-class-set all-others_fc_set
output-traffic-control-profile all-others_tcp
set class-of-service interfaces xe-0/0/21 forwarding-class-set lossless_fc_set
output-traffic-control-profile lossless_tcp
set class-of-service interfaces xe-0/0/21 forwarding-class-set all-others_fc_set
output-traffic-control-profile all-others_tcp
set class-of-service interfaces xe-0/0/40 forwarding-class-set lossless_fc_set
output-traffic-control-profile lossless_tcp
set class-of-service interfaces xe-0/0/40 forwarding-class-set all-others_fc_set
output-traffic-control-profile all-others_tcp
set class-of-service interfaces xe-0/0/41 forwarding-class-set lossless_fc_set
output-traffic-control-profile lossless_tcp
set class-of-service interfaces xe-0/0/41 forwarding-class-set all-others_fc_set
output-traffic-control-profile all-others_tcp
```

Configuration for Port Scheduling

The port-scheduling-specific portion of this example assigns the scheduler maps (which set the CoS treatment of the forwarding classes in the scheduler map) to the interfaces.

```
[edit class-of-service]
set interfaces xe-0/0/20 scheduler-map lossless_map
set interfaces xe-0/0/20 scheduler-map all-others_map
set interfaces xe-0/0/21 scheduler-map lossless_map
set interfaces xe-0/0/21 scheduler-map all-others_map
set interfaces xe-0/0/40 scheduler-map lossless_map
set interfaces xe-0/0/40 scheduler-map all-others_map
set interfaces xe-0/0/41 scheduler-map lossless_map
set interfaces xe-0/0/41 scheduler-map all-others_map
```

Common Configuration (Applies to ETS Hierarchical Scheduling and to Port Scheduling)

Step-by-Step Procedure The following step-by-step procedure shows you how to configure the VLANs, IRB interfaces, lossless forwarding classes, classifiers, PFC settings to enable PFC across Layer 3 interfaces, and the queue scheduling configuration common to ETS and direct port scheduling. For completeness, the ETS hierarchical port scheduling and direct port scheduling configurations are included separately, in the following procedures, but without explanatory text. See the *Related Documentation* links for detailed examples of the scheduling elements of the configuration.

1. Configure the Layer 3 interface VLANs and IP addresses:

```
[edit interfaces]
user@switch# set xe-0/0/40 vlan-tagging
user@switch# set xe-0/0/40 unit 0 vlan-id 103
user@switch# set xe-0/0/40 unit 0 family inet address 100.103.1.2/24
user@switch# set xe-0/0/41 vlan-tagging
user@switch# set xe-0/0/41 unit 0 vlan-id 104
user@switch# set xe-0/0/41 unit 0 family inet address 100.104.1.2/24
```

2. Configure the Layer 2 interface VLAN membership and interface mode:

```
[edit interfaces]
user@switch# set xe-0/0/20 unit 0 family ethernet-switching interface-mode trunk
user@switch# set xe-0/0/20 unit 0 family ethernet-switching vlan members vlan105
user@switch# set xe-0/0/21 unit 0 family ethernet-switching interface-mode trunk
user@switch# set xe-0/0/21 unit 0 family ethernet-switching vlan members vlan106
```

3. Configure the IRB interfaces and VLANs to transport incoming Layer 2 traffic assigned to VLANs vlan105 (of which interface xe-0/0/20 is a member) and vlan106 (of which interface xe-0/0/21 is a member) across Layer 3:

```
[edit]
user@switch# set interfaces irb unit 105 family inet address 100.105.1.1/24
user@switch# set interfaces irb unit 106 family inet address 100.106.1.1/24
user@switch# set vlans vlan105 vlan-id 105
user@switch# set vlans vlan106 vlan-id 106
user@switch# set vlans vlan105 l3-interface irb.105
user@switch# set vlans vlan106 l3-interface irb.106
```

4. Configure the lossless forwarding classes and a best-effort forwarding class for any other traffic that might use the interfaces:

```
[edit class-of-service]
user@switch# set forwarding-classes class lossless-3 queue-num 3 no-loss
user@switch# set forwarding-classes class lossless-4 queue-num 4 no-loss
user@switch# set forwarding-classes class all-others queue-num 0
```

5. Configure the IEEE classifier for the Layer 2 and Layer 3 interfaces to classify incoming traffic into the lossless forwarding classes based on the IEEE 802.1p code point of the traffic:

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 lossless-3-4-ieee forwarding-class lossless-3 loss-priority low code-points 011
user@switch# set ieee-802.1 lossless-3-4-ieee forwarding-class lossless-4 loss-priority low code-points 100
```

6. Configure the CNP to enable PFC on the lossless priorities (the lossless forwarding classes mapped to IEEE 802.1p code points 3 and 4):

```
[edit class-of-service congestion-notification-profile]
user@switch# set lossless-cnp input ieee-802.1 code-point 011 pfc
user@switch# set lossless-cnp input ieee-802.1 code-point 100 pfc
```

7. Apply the Layer 2 IEEE 802.1p classifier and the CNP to the Layer 3 interfaces:

```
[edit class-of-service interfaces]
user@switch# set xe-0/0/40 classifiers ieee-802.1 lossless-3-4-ieee
user@switch# set xe-0/0/40 congestion-notification-profile lossless-cnp
user@switch# set xe-0/0/41 classifiers ieee-802.1 lossless-3-4-ieee
user@switch# set xe-0/0/41 congestion-notification-profile lossless-cnp
```

8. Apply the Layer 2 IEEE 802.1p classifier and the CNP to the Layer 2 interfaces:

```
[edit class-of-service interfaces]
user@switch# xe-0/0/20 unit 0 classifiers ieee-802.1 lossless-3-4-ieee
user@switch# xe-0/0/20 congestion-notification-profile lossless-cnp
user@switch# xe-0/0/21 unit 0 classifiers ieee-802.1 lossless-3-4-ieee
user@switch# xe-0/0/21 congestion-notification-profile lossless-cnp
```

9. Configure queue scheduling to support the lossless configuration and map the schedulers to the forwarding classes (statements included here for completeness; see the *Related Documentation* links for detailed examples of scheduling configuration):

```
[edit class-of-service]
user@switch# set schedulers lossless_sch transmit-rate 6g
user@switch# set schedulers lossless_sch shaping-rate percent 100
user@switch# set schedulers all-others_sch transmit-rate 4g
user@switch# set scheduler-maps lossless_map forwarding-class lossless-3 scheduler
lossless_sch
user@switch# set scheduler-maps lossless_map forwarding-class lossless-4 scheduler
lossless_sch
user@switch# set scheduler-maps all-others_map forwarding-class all-others scheduler
all-others_sch
```

ETS Hierarchical Scheduling Configuration

Step-by-Step Procedure

1. Configure hierarchical scheduling to support the lossless configuration (included here for completeness; see the *Related Documentation* links for detailed examples of scheduling configuration) and apply it to the Layer 2 and Layer 3 interfaces:

```
[edit class-of-service interfaces]
user@switch# set forwarding-class-sets lossless_fc_set class lossless-3
user@switch# set forwarding-class-sets lossless_fc_set class lossless-4
user@switch# set forwarding-class-sets all-others_fc_set class all-others
user@switch# set traffic-control-profiles lossless_tcp scheduler-map lossless_map
user@switch# set traffic-control-profiles lossless_tcp guaranteed-rate percent 60
user@switch# set traffic-control-profiles lossless_tcp shaping-rate percent 100
user@switch# set traffic-control-profiles all-others_tcp scheduler-map all-others_map
user@switch# set traffic-control-profiles all-others_tcp guaranteed-rate percent 40
user@switch# set interfaces xe-0/0/20 forwarding-class-set lossless_fc_set
output-traffic-control-profile lossless_tcp
user@switch# set interfaces xe-0/0/20 forwarding-class-set all-others_fc_set
output-traffic-control-profile all-others_tcp
user@switch# set interfaces xe-0/0/21 forwarding-class-set lossless_fc_set
output-traffic-control-profile lossless_tcp
user@switch# set interfaces xe-0/0/21 forwarding-class-set all-others_fc_set
output-traffic-control-profile all-others_tcp
user@switch# set interfaces xe-0/0/40 forwarding-class-set lossless_fc_set
output-traffic-control-profile lossless_tcp
user@switch# set interfaces xe-0/0/40 forwarding-class-set all-others_fc_set
output-traffic-control-profile all-others_tcp
user@switch# set interfaces xe-0/0/41 forwarding-class-set lossless_fc_set
output-traffic-control-profile lossless_tcp
user@switch# set interfaces xe-0/0/41 forwarding-class-set all-others_fc_set
output-traffic-control-profile all-others_tcp
```

Port Scheduling Configuration

Step-by-Step Procedure

1. Apply port scheduling to support the lossless configuration on interfaces:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 scheduler-map lossless_map
user@switch# set interfaces xe-0/0/20 scheduler-map all-others_map
user@switch# set interfaces xe-0/0/21 scheduler-map lossless_map
user@switch# set interfaces xe-0/0/21 scheduler-map all-others_map
user@switch# set interfaces xe-0/0/40 scheduler-map lossless_map
user@switch# set interfaces xe-0/0/40 scheduler-map all-others_map
user@switch# set interfaces xe-0/0/41 scheduler-map lossless_map
user@switch# set interfaces xe-0/0/41 scheduler-map all-others_map
```

Results

Display the results of the interface, VLAN, and class-of-service configurations (the system shows only the explicitly configured parameters; it does not show default parameters). The results are valid for both Switch SW1 and Switch SW2 because the same configuration is used on both switches. The results are from the ETS hierarchical scheduling configuration, which show the more complex configuration. Direct port scheduling results would not show the traffic control profile or forwarding class set portions of the configuration, but would display the name of the scheduler map under each interface

(instead of the names of the forwarding class set and output traffic control profile). Other than that, the results are the same.

Display the results of the interface configuration:

```
user@switch# show configuration interfaces
```

```
xe-0/0/20 {
  unit 0 {
    family ethernet-switching {
      interface-mode trunk;
      vlan {
        members vlan105;
      }
    }
  }
}
xe-0/0/21 {
  unit 0 {
    family ethernet-switching {
      interface-mode trunk;
      vlan {
        members vlan106;
      }
    }
  }
}
xe-0/0/40 {
  vlan-tagging;
  unit 0 {
    vlan-id 103;
    family inet {
      address 100.103.1.2/24;
    }
  }
}
xe-0/0/41 {
  vlan-tagging;
  unit 0 {
    vlan-id 104;
    family inet {
      address 100.104.1.2/24;
    }
  }
}
irb {
  unit 105 {
    family inet {
      address 100.105.1.1/24;
    }
  }
  unit 106 {
    family inet {
      address 100.106.1.1/24;
    }
  }
}
vlan {
  unit 105 {
    family inet {
      address 100.105.1.1/24;
    }
  }
}
```

```
    }  
    unit 106 {  
        family inet {  
            address 100.106.1.1/24;  
        }  
    }  
}
```

Display the results of the vlan configuration:

```
user@switch# show configuration vlans  
vlan105 {  
    vlan-id 105;  
    l3-interface irb.105;  
}  
vlan106 {  
    vlan-id 106;  
    l3-interface irb.106;  
}
```

Display the results of the class-of-service configuration:

```
user@switch# show configuration class-of-service  
classifiers {  
    ieee-802.1 lossless-3-4-ieee {  
        forwarding-class lossless-3 {  
            loss-priority low code-points 011;  
        }  
        forwarding-class lossless-4 {  
            loss-priority low code-points 100;  
        }  
    }  
}  
forwarding-classes {  
    class lossless-3 queue-num 3 no-loss;  
    class lossless-4 queue-num 4 no-loss;  
    class all-others queue-num 0;  
}  
traffic-control-profiles {  
    lossless_tcp {  
        scheduler-map lossless_map;  
        shaping-rate percent 100;  
        guaranteed-rate percent 60;  
    }  
    all-others_tcp {  
        scheduler-map all-others_map;  
        guaranteed-rate percent 40;  
    }  
}  
forwarding-class-sets {  
    lossless_fc_set {  
        class lossless-3;  
        class lossless-4;  
    }  
    all-others_fc_set {  
        class all-others;  
    }  
}  
congestion-notification-profile {  
    lossless-cnp {
```

```

        input {
            ieee-802.1 {
                code-point 011 {
                    pfc;
                }
                code-point 100 {
                    pfc;
                }
            }
        }
    }
}
interfaces {
    xe-0/0/20 {
        forwarding-class-set {
            lossless_fc_set {
                output-traffic-control-profile lossless_tcp;
            }
            all-others_fc_set {
                output-traffic-control-profile all-others_tcp;
            }
        }
        congestion-notification-profile lossless-cnp;
        unit 0 {
            classifiers {
                ieee-802.1 lossless-3-4-ieee;
            }
        }
    }
    xe-0/0/21 {
        forwarding-class-set {
            all-others_fc_set {
                output-traffic-control-profile all-others_tcp;
            }
            lossless_fc_set {
                output-traffic-control-profile lossless_tcp;
            }
        }
        congestion-notification-profile lossless-cnp;
        unit 0 {
            classifiers {
                ieee-802.1 lossless-3-4-ieee;
            }
        }
    }
    xe-0/0/40 {
        forwarding-class-set {
            lossless_fc_set {
                output-traffic-control-profile lossless_tcp;
            }
            all-others_fc_set {
                output-traffic-control-profile all-others_tcp;
            }
        }
        congestion-notification-profile lossless-cnp;
        classifiers {
            ieee-802.1 lossless-3-4-ieee;
        }
    }
    xe-0/0/41 {
        forwarding-class-set {

```

```
        lossless_fc_set {
            output-traffic-control-profile lossless_tcp;
        }
        all-others_fc_set {
            output-traffic-control-profile all-others_tcp;
        }
    }
    congestion-notification-profile lossless-cnp;
    classifiers {
        ieee-802.1 lossless-3-4-ieee;
    }
}
scheduler-maps {
    lossless_map {
        forwarding-class lossless-3 scheduler lossless_sch;
        forwarding-class lossless-4 scheduler lossless_sch;
    }
    all-others_map {
        forwarding-class all-others scheduler all-others_sch;
    }
}
schedulers {
    lossless_sch {
        transmit-rate 6g;
        shaping-rate percent 100;
    }
    all-others_sch {
        transmit-rate 4g;
    }
}
```



TIP: To quickly configure the switch, issue the `load merge terminal` command, and then copy the hierarchies and paste them into the switch terminal window.

Verification

To verify that the PFC across Layer 3 interfaces configuration has been created and is operating properly, perform these tasks:

- [Verifying the Interface Configuration on page 6937](#)
- [Verifying the VLAN Configuration on page 6939](#)
- [Verifying the PFC Configuration \(Congestion Notification Profile\) on page 6939](#)
- [Verify the Forwarding Class Configuration on page 6940](#)
- [Verifying the Classifier Configuration on page 6940](#)
- [Verifying the Interface CoS Configuration \(Hierarchical Scheduling, PFC, and Classifier Mapping to Interfaces\) on page 6941](#)

Verifying the Interface Configuration

Purpose Verify that the Layer 2 Ethernet interfaces, Layer 3 IP interfaces, IRB interfaces, and VLAN interfaces have been created on the switch and are correctly configured.

Action Display the switch interface configuration using the **show configuration interfaces** command:

```
user@switch> show configuration interfaces
xe-0/0/20 {
  unit 0 {
    family ethernet-switching {
      interface-mode trunk;
      vlan {
        members vlan105;
      }
    }
  }
}
xe-0/0/21 {
  unit 0 {
    family ethernet-switching {
      interface-mode trunk;
      vlan {
        members vlan106;
      }
    }
  }
}
xe-0/0/40 {
  vlan-tagging;
  unit 0 {
    vlan-id 103;
    family inet {
      address 100.103.1.2/24;
    }
  }
}
xe-0/0/41 {
  vlan-tagging;
  unit 0 {
    vlan-id 104;
    family inet {
      address 100.104.1.2/24;
    }
  }
}
irb {
  unit 105 {
    family inet {
      address 100.105.1.1/24;
    }
  }
  unit 106 {
    family inet {
      address 100.106.1.1/24;
    }
  }
}
vlan {
  unit 105 {
    family inet {
      address 100.105.1.1/24;
    }
  }
}
```

```

    unit 106 {
        family inet {
            address 100.106.1.1/24;
        }
    }
}

```

Meaning The **show configuration interfaces** command displays all of the interfaces configured on the switch. The command output shows that:

- Interfaces xe-0/0/20 and xe-0/0/21 are Ethernet interfaces (family ethernet-switching) in trunk interface mode. Interface xe-0/0/20 is a member of VLAN vlan105, and interface xe-0/0/21 is a member of VLAN vlan106.
- Interfaces xe-0/0/40 and xe-0/0/41 are IP interfaces (family inet) with VLAN tagging enabled. Interface xe-0/0/40 has an IP address of 100.103.1.2/24 and a VLAN ID of 103. Interface xe-0/0/41 has an IP address of 100.104.1.2/24 and a VLAN ID of 104.
- Two IRB interfaces are configured, IRB unit 105 with an IP address of 100.105.1.1/24 and IRB unit 106 with an IP address of 100.106.1.1/24.
- Two VLAN interfaces are configured, VLAN unit 105 with an IP address of 100.105.1.1/24 (for IRB interface unit 105) and VLAN unit 106 with an IP address of 100.106.1.1/24 (for IRB interface unit 106).

Verifying the VLAN Configuration

Purpose Verify that VLANs have been created on the switch and are correctly configured.

Action Display the VLAN configuration using the **show configuration vlans** command:

```

user@switch> show configuration vlans
vlan105 {
    vlan-id 105;
    l3-interface irb.105;
}
vlan106 {
    vlan-id 106;
    l3-interface irb.106;
}

```

Meaning The **show configuration vlans** command displays all of the VLANs configured on the switch. The command output shows that:

- VLAN vlan105 has been configured with VLAN ID 105 on IRB interface irb.105.
- VLAN vlan106 has been configured with VLAN ID 106 on IRB interface irb.106.

Verifying the PFC Configuration (Congestion Notification Profile)

Purpose Verify that PFC has been enabled on the correct IEEE 802.1p code points (priorities) in the CNP.

Action Display the PFC configuration using the **show configuration class-of-service congestion-notification-profile** command:

```
user@switch> show configuration class-of-service congestion-notification-profile
lossless-cnp {
    input {
        ieee-802.1 {
            code-point 011 {
                pfc;
            }
            code-point 100 {
                pfc;
            }
        }
    }
}
```

Meaning The **show configuration class-of-service congestion-notification-profile** command displays all of the CNPs configured on the switch. The command output shows that:

- The CNP named **lossless-cnp** is configured on the switch.
- The CNP **lossless-cnp** enables PFC on IEEE 802.1p code points 100 and 100.

Verify the Forwarding Class Configuration

Purpose Verify that the two lossless forwarding classes and the best-effort forwarding class have been configured on the switch.

Action Display the forwarding class configuration using the **show configuration class-of-service forwarding-classes** command:

```
user@switch> show configuration class-of-service forwarding-classes
class lossless-3 queue-num 3 no-loss;
class lossless-4 queue-num 4 no-loss;
class all-others queue-num 0;
```

Meaning The **show configuration class-of-service forwarding-classes** command displays all of the forwarding classes configured on the switch (default forwarding classes are not displayed). The command output shows that:

- Forwarding class **lossless-3** is mapped to queue 3 and is configured as a lossless forwarding class (the **no-loss** attribute is applied)
- Forwarding class **lossless-4** is mapped to queue 4 and is configured as a lossless forwarding class (the **no-loss** attribute is applied)
- Forwarding class **all-others** is mapped to queue 0. It is not a lossless forwarding class (the **no-loss** attribute is not applied).

Verifying the Classifier Configuration

Purpose Verify that the IEEE 802.1p classifier has been configured on the switch.

Action Display the classifier configuration using the **show configuration class-of-service classifiers** command:

```
user@switch> show configuration class-of-service classifiers
ieee-802.1 lossless-3-4-ieee {
    forwarding-class lossless-3 {
        loss-priority low code-points 011;
    }
    forwarding-class lossless-4 {
        loss-priority low code-points 100;
    }
}
```

Meaning The **show configuration class-of-service classifiers** command displays all of the classifiers configured on the switch. The command output shows that the Layer 2 IEEE 802.1p classifier **lossless-3-4-ieee** classifies traffic with the code point 011 into the **lossless-3** forwarding class with a loss priority of **low**, and classifies traffic with the code point 100 into the **lossless-4** forwarding class with a loss priority of **low**.

Verifying the Interface CoS Configuration (Hierarchical Scheduling, PFC, and Classifier Mapping to Interfaces)

Purpose Verify that the interfaces have the correct hierarchical scheduling, PFC, and classifier configurations.



NOTE: The results are from the ETS hierarchical scheduling configuration, which shows the more complex configuration. Direct port scheduling results would not show the traffic control profile or forwarding class set portions of the interface configuration, but would display the name of the scheduler map under each interface instead of the names of the forwarding class set and output traffic control profile. Other than that, they are the same.

Action Display the interface CoS configuration using the **show configuration class-of-service interfaces** command:

```
user@switch> show configuration class-of-service interfaces
xe-0/0/20 {
  forwarding-class-set {
    lossless_fc_set {
      output-traffic-control-profile lossless_tcp;
    }
    all-others_fc_set {
      output-traffic-control-profile all-others_tcp;
    }
  }
  congestion-notification-profile lossless-cnp;
  unit 0 {
    classifiers {
      ieee-802.1 lossless-3-4-ieee;
    }
  }
}
xe-0/0/21 {
  forwarding-class-set {
    all-others_fc_set {
      output-traffic-control-profile all-others_tcp;
    }
    lossless_fc_set {
      output-traffic-control-profile lossless_tcp;
    }
  }
  congestion-notification-profile lossless-cnp;
  unit 0 {
    classifiers {
      ieee-802.1 lossless-3-4-ieee;
    }
  }
}
xe-0/0/40 {
  forwarding-class-set {
    lossless_fc_set {
      output-traffic-control-profile lossless_tcp;
    }
    all-others_fc_set {
      output-traffic-control-profile all-others_tcp;
    }
  }
  congestion-notification-profile lossless-cnp;
  classifiers {
    ieee-802.1 lossless-3-4-ieee;
  }
}
xe-0/0/41 {
  forwarding-class-set {
    lossless_fc_set {
      output-traffic-control-profile lossless_tcp;
    }
    all-others_fc_set {
      output-traffic-control-profile all-others_tcp;
    }
  }
  congestion-notification-profile lossless-cnp;
}
```

```

classifiers {
    ieee-802.1 lossless-3-4-ieee;
}

```

Meaning The `show configuration class-of-service interfaces` command displays all of the CoS components configured on the switch interfaces. The command output shows that:

- The configuration on Layer 2 Ethernet interfaces xe-0/0/20 and xe-0/0/21 includes:
 - Hierarchical scheduling—The forwarding class set **lossless_fc_set** with the traffic control profile **lossless_tcp** for the lossless traffic, and the forwarding class set **all-others_fc_set** with the traffic control profile **all-others_tcp** for the best-effort traffic are applied to both interfaces.
 - PFC—The **lossless-cnp** congestion notification profile is applied to both interfaces.
 - Classifiers—The Layer 2 IEEE 802.1p classifier **lossless-3-4-ieee** is applied to both interfaces.
- The configuration on Layer 3 IP interfaces xe-0/0/40 and xe-0/0/41 includes:
 - Hierarchical scheduling—The forwarding class set **lossless_fc_set** with the traffic control profile **lossless_tcp** for the lossless traffic, and the forwarding class set **all-others_fc_set** with the traffic control profile **all-others_tcp** for the best-effort traffic are applied to both interfaces.
 - PFC—The **lossless-cnp** congestion notification profile is applied to both interfaces.
 - Classifiers—The Layer 2 IEEE 802.1p classifier **lossless-3-4-ieee** is applied to both interfaces. Traffic that would use a DSCP or a DSCP IPv6 classifier if it were configured uses the IEEE 802.1p classifier instead. Using the IEEE 802.1p classifier allows the interface to use PFC to pause traffic during periods of congestion to prevent packet loss.

Related Documentation

- [Understanding PFC Functionality Across Layer 3 Interfaces on page 6921](#)

Example: Configuring CoS for FCoE Transit Switch Traffic Across an MC-LAG

Multichassis link aggregation groups (MC-LAGs) provide redundancy and load balancing between two QFX Series switches, multihoming support for client devices such as servers, and a loop-free Layer 2 network without running Spanning Tree Protocol (STP).



NOTE: This example uses the Junos OS Enhanced Layer 2 Software (ELS) configuration style for QFX Series switches. If your switch runs software that does not support ELS, see *Example: Configuring CoS for FCoE Transit Switch Traffic Across an MC-LAG*. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 41](#).

You can use an MC-LAG to provide a redundant aggregation layer for Fibre Channel over Ethernet (FCoE) traffic in an *inverted-U* topology. To support lossless transport of FCoE traffic across an MC-LAG, you must configure the appropriate class of service (CoS) on both of the QFX Series switches with MC-LAG port members. The CoS configuration must be the same on both of the MC-LAG switches because an MC-LAG does not carry forwarding class and IEEE 802.1p priority information.

Ports that are members of an MC-LAG act as FCoE passthrough transit switch ports.



NOTE: This example describes how to configure CoS to provide lossless transport for FCoE traffic across an MC-LAG that connects two QFX Series switches. It also describes how to configure CoS on the FCoE transit switches that connect FCoE hosts to the QFX Series switches that form the MC-LAG.

This example does *not* describe how to configure the MC-LAG itself. For a detailed example of MC-LAG configuration, see *Example: Configuring Multichassis Link Aggregation*. However, this example includes a subset of MC-LAG configuration that only shows how to configure interface membership in the MC-LAG.



NOTE: Juniper Networks QFX10000 aggregation switches do not support FIP snooping, so they cannot be used as FIP snooping access switches (Transit Switches TS1 and TS2) in this example. However, QFX10000 switches can play the role of the MC-LAG switches (MC-LAG Switch S1 and MC-LAG Switch S2) in this example.

QFX3500 and QFX3600 Virtual Chassis switches do not support FCoE.

This topic describes:

- [Requirements on page 6944](#)
- [Overview on page 6945](#)
- [Configuration on page 6950](#)
- [Verification on page 6960](#)

Requirements

This example uses the following hardware and software components:

- Two Juniper Networks QFX5100 Switches running the ELS CLI that form an MC-LAG for FCoE traffic.
- Two Juniper Networks QFX5100 Switches running the ELS CLI that provide FCoE server access in transit switch mode and that connect to the MC-LAG switches.
- FCoE servers (or other FCoE hosts) connected to the transit switches.
- Junos OS Release 13.2 or later for the QFX Series.

Overview

FCoE traffic requires lossless transport. This example shows you how to:

- Configure CoS for FCoE traffic on the two QFX5100 switches that form the MC-LAG, including priority-based flow control (PFC). The example also includes configuration for both enhanced transmission selection (ETS) hierarchical scheduling of resources for the FCoE forwarding class priority and for the forwarding class set priority group, and also direct port scheduling. You can only use one of the scheduling methods on a port. Different switches support different scheduling methods.



NOTE: Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

- Configure CoS for FCoE on the two FCoE transit switches that connect FCoE hosts to the MC-LAG switches and enable FIP snooping on the FCoE VLAN at the FCoE transit switch access ports.
- Configure the appropriate port mode, MTU, and FCoE trusted or untrusted state for each interface to support lossless FCoE transport.




NOTE: Do not enable IGMP snooping on the FCoE VLAN. (IGMP snooping is enabled on the default VLAN by default, but is disabled by default on all other VLANs.)

Topology

QFX5100 switches that act as transit switches support MC-LAGs for FCoE traffic in an inverted-U network topology, as shown in [Figure 206](#).

Figure 232: Supported Topology for an MC-LAG on an FCoE Transit Switch



NOTE: Juniper Networks QFX10000 aggregation switches do not support FIP snooping, so they cannot be used as FIP snooping access switches (Transit Switches TS1 and TS2) in this example. However, QFX10000 switches can play the role of the MC-LAG switches (MC-LAG Switch S1 and MC-LAG Switch S2) in this example.

Table 512 shows the configuration components for this example.

Table 594: Components of the CoS for FCoE Traffic Across an MC-LAG Configuration Topology

Component	Settings
Hardware	Four QFX5100 switches running the ELS CLI (two to form the MC-LAG as passthrough transit switches and two transit switches for FCoE access).
Forwarding class (all switches)	Default fcoe forwarding class.
Classifier (forwarding class mapping of incoming traffic to IEEE priority)	Default IEEE 802.1p trusted classifier on all FCoE interfaces.

Table 594: Components of the CoS for FCoE Traffic Across an MC-LAG Configuration Topology (continued)

Component	Settings
LAGs and MC-LAG	<p>S1—Ports xe-0/0/10 and x-0/0/11 are members of LAG ae0, which connects Switch S1 to Switch S2. Ports xe-0/0/20 and xe-0/0/21 are members of MC-LAG ae1. All ports are configured in trunk interface mode, as fcoe-trusted, and with an MTU of 2180.</p> <p>S2—Ports xe-0/0/10 and x-0/0/11 are members of LAG ae0, which connects Switch S2 to Switch S1. Ports xe-0/0/20 and xe-0/0/21 are members of MC-LAG ae1. All ports are configured in trunk interface mode, as fcoe-trusted, and with an MTU of 2180.</p> <p>NOTE: Ports xe-0/0/20 and xe-0/0/21 on Switches S1 and S2 are the members of the MC-LAG.</p> <p>TS1—Ports xe-0/0/25 and x-0/0/26 are members of LAG ae1, configured in trunk interface mode, as fcoe-trusted, and with an MTU of 2180. Ports xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33 are configured in trunk interface mode, with an MTU of 2180.</p> <p>TS2—Ports xe-0/0/25 and x-0/0/26 are members of LAG ae1, configured in trunk interface mode, as fcoe-trusted, and with an MTU of 2180. Ports xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33 are configured in trunk interface mode, with an MTU of 2180.</p>
FCoE queue scheduler (all switches)	<p>fcoe-sched: Minimum bandwidth 3g Maximum bandwidth 100% Priority low</p>
Forwarding class-to-scheduler mapping (all switches)	<p>Scheduler map fcoe-map: Forwarding class fcoe Scheduler fcoe-sched</p> <p>NOTE: If you are using direct port scheduling,</p>
PFC congestion notification profile (all switches)	<p>fcoe-cnp: Code point 011</p> <p>Ingress interfaces:</p> <ul style="list-style-type: none"> • S1—LAG ae0 and MC-LAG ae1 • S2—LAG ae0 and MC-LAG ae1 • TS1—LAG ae1, interfaces xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33 • TS2—LAG ae1, interfaces xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33

Table 594: Components of the CoS for FCoE Traffic Across an MC-LAG Configuration Topology (continued)

Component	Settings
FCoE VLAN name and tag ID	<p>Name—fcoe_vlan ID—100</p> <p>Include the FCoE VLAN on the interfaces that carry FCoE traffic on all four switches.</p>
ETS only—forwarding class set (FCoE priority group, all switches)	<p>fcoe-pg: Forwarding class fcoe</p> <p>Egress interfaces:</p> <ul style="list-style-type: none"> • S1—LAG ae0 and MC-LAG ae1 • S2—LAG ae0 and MC-LAG ae1 • TS1—LAG ae1, interfaces xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33 • TS2—LAG ae1, interfaces xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33
ETS only—traffic control profile (all switches)	<p>fcoe-tcp: Scheduler map fcoe-map Minimum bandwidth 3g Maximum bandwidth 100%</p> <p>The traffic control profile is applied to the same interfaces as the forwarding class set, using the same CLI statement. This applies ETS hierarchical scheduling to the interfaces.</p>
Port scheduling only—apply scheduling to interfaces	<p>On switches that support direct port scheduling, if you use port scheduling, apply scheduling by attaching the scheduler map directly to interfaces:</p> <ul style="list-style-type: none"> • S1—LAG ae0 and MC-LAG ae1 • S2—LAG ae0 and MC-LAG ae1 • TS1—LAG ae1, interfaces xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33 • TS2—LAG ae1, interfaces xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33
FIP snooping	<p>Enable FIP snooping on Transit Switches TS1 and TS2 on the FCoE VLAN. Configure the LAG interfaces that connect to the MC-LAG switches as FCoE trusted interfaces so that they do not perform FIP snooping.</p> <p>This example enables VN2VN_Port FIP snooping on the FCoE transit switch interfaces connected to the FCoE servers. The example is equally valid with VN2VF_Port FIP snooping enabled on the transit switch access ports. The method of FIP snooping you enable depends on your network configuration.</p> <p>NOTE: Juniper Networks QFX10000 aggregation switches do not support FIP snooping, so they cannot be used as FIP snooping access switches (Transit Switches TS1 and TS2) in this example.</p>



NOTE: This example uses the default IEEE 802.1p trusted BA classifier, which is automatically applied to trunk mode interfaces if you do not apply an explicitly configured classifier.

To configure CoS for FCoE traffic across an MC-LAG:

- Use the default FCoE forwarding class and forwarding-class-to-queue mapping (do not explicitly configure the FCoE forwarding class or output queue). The default FCoE forwarding class is **fcoe**, and the default output queue is queue **3**.
- Use the default trusted BA classifier, which maps incoming packets to forwarding classes by the IEEE 802.1p code point (CoS priority) of the packet. The trusted classifier is the default classifier for interfaces in trunk interface mode. The default trusted classifier maps incoming packets with the IEEE 802.1p code point 3 (**011**) to the FCoE forwarding class. If you choose to configure the BA classifier instead of using the default classifier, you must ensure that FCoE traffic is classified into forwarding classes in exactly the same way on both MC-LAG switches. Using the default classifier ensures consistent classifier configuration on the MC-LAG ports.
- Configure a congestion notification profile that enables PFC on the FCoE code point (code point **011** in this example). The congestion notification profile configuration must be the same on both MC-LAG switches.
- Apply the congestion notification profile to the interfaces.
- Configure the interface mode, MTU, and FCoE trusted or untrusted state for each interface to support lossless FCoE transport.
- For ETS hierarchical port scheduling, configure ETS on the interfaces to provide the bandwidth required for lossless FCoE transport. Configuring ETS includes configuring bandwidth scheduling for the FCoE forwarding class, a forwarding class set (priority group) that includes the FCoE forwarding class, and a traffic control profile to assign bandwidth to the forwarding class set that includes FCoE traffic, and applying the traffic control profile and forwarding class set to interfaces..

On switches that support direct port scheduling, configure CoS properties on interfaces by applying scheduler maps directly to interfaces.

In addition, this example describes how to enable FIP snooping on the Transit Switch TS1 and TS2 ports that are connected to the FCoE servers. To provide secure access, FIP snooping must be enabled on the FCoE access ports.

This example focuses on the CoS configuration to support lossless FCoE transport across an MC-LAG. This example does not describe how to configure the properties of MC-LAGs and LAGs, although it does show you how to configure the port characteristics required to support lossless transport and how to assign interfaces to the MC-LAG and to the LAGs.

Before you configure CoS, configure:

- The MC-LAGs that connect Switches S1 and S2 to Switches TS1 and TS2. (*Example: Configuring Multichassis Link Aggregation* describes how to configure MC-LAGs.)

- The LAGs that connect the Transit Switches TS1 and TS2 to MC-LAG Switches S1 and S2. (“[Configuring Link Aggregation](#)” on [page 2869](#) describes how to configure LAGs.)
- The LAG that connects Switch S1 to Switch S2.

Configuration

To configure CoS for lossless FCoE transport across an MC-LAG, perform these tasks:

- [MC-LAG Switches S1 and S2 Common Configuration \(Applies to ETS and Port Scheduling\) on page 6952](#)
- [MC-LAG Switches S1 and S2 ETS Hierarchical Scheduling Configuration on page 6953](#)
- [MC-LAG Switches S1 and S2 Port Scheduling Configuration on page 6954](#)
- [FCoE Transit Switches TS1 and TS2 Common Configuration \(Applies to ETS and Port Scheduling\) on page 6954](#)
- [FCoE Transit Switches TS1 and TS2 ETS Hierarchical Scheduling Configuration on page 6956](#)
- [FCoE Transit Switches TS1 and TS2 Port Scheduling Configuration on page 6956](#)
- [Results on page 6957](#)

CLI Quick Configuration



NOTE: The CLI configurations for the MC-LAG switches and for the FCoE transit switches are each separated into three sections:

- Configuration common to all port scheduling methods
- Configuration specific to ETS hierarchical port scheduling
- Configuration specific to direct port scheduling

MC-LAG Switch S1 and Switch S2

To quickly configure CoS for lossless FCoE transport across an MC-LAG, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI for MC-LAG Switch S1 and MC-LAG Switch S2 at the **[edit]** hierarchy level. The configurations on Switches S1 and S2 are identical because the CoS configuration must be identical, and because this example uses the same ports on both switches.

MC-LAG Switches Configuration Common to ETS Hierarchical Port Scheduling and to Direct Port Scheduling

```
set class-of-service schedulers fcoe-sched priority low transmit-rate 3g
set class-of-service schedulers fcoe-sched shaping-rate percent 100
set class-of-service scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
set class-of-service congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011 pfc
set class-of-service interfaces ae0 congestion-notification-profile fcoe-cnp
set class-of-service interfaces ae1 congestion-notification-profile fcoe-cnp
set vlans fcoe_vlan vlan-id 100
set interfaces xe-0/0/10 ether-options 802.3ad ae0
set interfaces xe-0/0/11 ether-options 802.3ad ae0
set interfaces xe-0/0/20 ether-options 802.3ad ae1
set interfaces xe-0/0/21 ether-options 802.3ad ae1
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk vlan members fcoe_vlan
set interfaces ae1 unit 0 family ethernet-switching interface-mode trunk vlan members fcoe_vlan
```

```

set interfaces ae0 mtu 2180
set interfaces ae1 mtu 2180
set vlans fcoe_vlan forwarding-options fip-security interface ae0 fcoe-trusted
set vlans fcoe_vlan forwarding-options fip-security interface ae1 fcoe-trusted

```

MC-LAG Switches Configuration for ETS Hierarchical Port Scheduling

```

set class-of-service forwarding-class-sets fcoe-pg class fcoe
set class-of-service traffic-control-profiles fcoe-tcp scheduler-map fcoe-map guaranteed-rate
3g
set class-of-service traffic-control-profiles fcoe-tcp shaping-rate percent 100
set class-of-service interfaces ae0 forwarding-class-set fcoe-pg output-traffic-control-profile
fcoe-tcp
set class-of-service interfaces ae1 forwarding-class-set fcoe-pg output-traffic-control-profile
fcoe-tcp

```

MC-LAG Switches Configuration for Direct Port Scheduling

```

set class-of-service interfaces ae0 scheduler-map fcoe-map
set class-of-service interfaces ae1 scheduler-map fcoe-map

```

FCoE Transit Switch TS1 and Switch TS2

To quickly configure CoS for lossless FCoE transport across an MC-LAG, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI for Transit Switch TS1 and Transit Switch TS2 at the **[edit]** hierarchy level. The configurations on Switches TS1 and TS2 are identical because the CoS configuration must be identical, and because this example uses the same ports on both switches.

FCoE Transit Switches Configuration Common to ETS Hierarchical Port Scheduling and to Direct Port Scheduling

```

set class-of-service schedulers fcoe-sched priority low transmit-rate 3g
set class-of-service schedulers fcoe-sched shaping-rate percent 100
set class-of-service scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
set class-of-service congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011 pfc
set class-of-service interfaces ae1 congestion-notification-profile fcoe-cnp
set class-of-service interfaces xe-0/0/30 congestion-notification-profile fcoe-cnp
set class-of-service interfaces xe-0/0/31 congestion-notification-profile fcoe-cnp
set class-of-service interfaces xe-0/0/32 congestion-notification-profile fcoe-cnp
set class-of-service interfaces xe-0/0/33 congestion-notification-profile fcoe-cnp
set vlans fcoe_vlan vlan-id 100
set interfaces xe-0/0/25 ether-options 802.3ad ae1
set interfaces xe-0/0/26 ether-options 802.3ad ae1
set interfaces ae1 unit 0 family ethernet-switching interface-mode trunk vlan members fcoe_vlan
set interfaces xe-0/0/30 unit 0 family ethernet-switching interface-mode trunk vlan members
fcoe_vlan
set interfaces xe-0/0/31 unit 0 family ethernet-switching interface-mode trunk vlan members
fcoe_vlan
set interfaces xe-0/0/32 unit 0 family ethernet-switching interface-mode trunk vlan members
fcoe_vlan
set interfaces xe-0/0/33 unit 0 family ethernet-switching interface-mode trunk vlan members
fcoe_vlan
set interfaces ae1 mtu 2180
set interfaces xe-0/0/30 mtu 2180
set interfaces xe-0/0/31 mtu 2180
set interfaces xe-0/0/32 mtu 2180
set interfaces xe-0/0/33 mtu 2180
set vlans fcoe_vlan forwarding-options fip-security interface ae1 fcoe-trusted
set vlans fcoe_vlan forwarding-options fip-security examine-vn2v2 beacon-period 90000

```

FCoE Transit Switches Configuration for ETS Hierarchical Port Scheduling

```
set class-of-service forwarding-class-sets fcoe-pg class fcoe
set class-of-service traffic-control-profiles fcoe-tcp scheduler-map fcoe-map guaranteed-rate
3g
set class-of-service traffic-control-profiles fcoe-tcp shaping-rate percent 100
set class-of-service interfaces ae1 forwarding-class-set fcoe-pg output-traffic-control-profile
fcoe-tcp
set class-of-service interfaces xe-0/0/30 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
set class-of-service interfaces xe-0/0/31 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
set class-of-service interfaces xe-0/0/32 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
set class-of-service interfaces xe-0/0/33 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
```

FCoE Transit Switches Configuration for Direct Port Scheduling

```
set class-of-service interfaces ae1 scheduler-map fcoe-map
set class-of-service interfaces xe-0/0/30 scheduler-map fcoe-map
set class-of-service interfaces xe-0/0/31 scheduler-map fcoe-map
set class-of-service interfaces xe-0/0/32 scheduler-map fcoe-map
set class-of-service interfaces xe-0/0/33 scheduler-map fcoe-map
```

MC-LAG Switches S1 and S2 Common Configuration (Applies to ETS and Port Scheduling)

Step-by-Step Procedure

To configure queue scheduling, PFC, the FCoE VLAN, and LAG and MC-LAG interface membership and characteristics to support lossless FCoE transport across an MC-LAG (this example uses the default **fcoe** forwarding class and the default classifier to map incoming FCoE traffic to the FCoE IEEE 802.1p code point **011**), for both ETS hierarchical port scheduling and port scheduling (common configuration):

1. Configure output scheduling for the FCoE queue:

```
[edit class-of-service]
user@switch# set schedulers fcoe-sched priority low transmit-rate 3g
user@switch# set schedulers fcoe-sched shaping-rate percent 100
```
2. Map the FCoE forwarding class to the FCoE scheduler (**fcoe-sched**):

```
[edit class-of-service]
user@switch# set scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
```
3. Enable PFC on the FCoE priority by creating a congestion notification profile (**fcoe-cnp**) that applies FCoE to the IEEE 802.1 code point **011**:

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point
011 pfc
```
4. Apply the PFC configuration to the LAG and MC-LAG interfaces:

```
[edit class-of-service]
user@switch# set interfaces ae0 congestion-notification-profile fcoe-cnp
user@switch# set interfaces ae1 congestion-notification-profile fcoe-cnp
```
5. Configure the VLAN for FCoE traffic (**fcoe_vlan**):

- ```
[edit vlans]
user@switch# set fcoe_vlan vlan-id 100
```
6. Add the member interfaces to the LAG between the two MC-LAG switches:
 

```
[edit interfaces]
user@switch# set xe-0/0/10 ether-options 802.3ad ae0
user@switch# set xe-0/0/11 ether-options 802.3ad ae0
```
  7. Add the member interfaces to the MC-LAG:
 

```
[edit interfaces]
user@switch# set xe-0/0/20 ether-options 802.3ad ae1
user@switch# set xe-0/0/21 ether-options 802.3ad ae1
```
  8. Configure the interface mode as **trunk** and membership in the FCoE VLAN (**fcoe\_vlan**) for the LAG (**ae0**) and for the MC-LAG (**ae1**):
 

```
[edit interfaces]
user@switch# set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
vlan members fcoe_vlan
user@switch# set interfaces ae1 unit 0 family ethernet-switching interface-mode trunk
vlan members fcoe_vlan
```
  9. Set the MTU to **2180** for the LAG and MC-LAG interfaces. 2180 bytes is the minimum size required to handle FCoE packets because of the payload and header sizes; you can configure the MTU to a higher number of bytes if desired, but not less than 2180 bytes:
 

```
[edit interfaces]
user@switch# set ae0 mtu 2180
user@switch# set ae1 mtu 2180
```
  10. Set the LAG and MC-LAG interfaces as FCoE trusted ports. Ports that connect to other switches should be trusted and should not perform FIP snooping:
 

```
[edit]
user@switch# set vlans fcoe_vlan forwarding-options fip-security interface ae0 fcoe-trusted
user@switch# set vlans fcoe_vlan forwarding-options fip-security interface ae1 fcoe-trusted
```

### MC-LAG Switches S1 and S2 ETS Hierarchical Scheduling Configuration

#### Step-by-Step Procedure

To configure the forwarding class set (priority group) and priority group scheduling (in a traffic control profile), and apply the ETS hierarchical scheduling for FCoE traffic to interfaces:

1. Configure the forwarding class set (**fcoe-pg**) for the FCoE traffic:
 

```
[edit class-of-service]
user@switch# set forwarding-class-sets fcoe-pg class fcoe
```
2. Define the traffic control profile (**fcoe-tcp**) to use on the FCoE forwarding class set:
 

```
[edit class-of-service]
user@switch# set traffic-control-profiles fcoe-tcp scheduler-map fcoe-map
guaranteed-rate 3g
user@switch# set traffic-control-profiles fcoe-tcp shaping-rate percent 100
```
3. Apply the FCoE forwarding class set and traffic control profile to the LAG and MC-LAG interfaces:
 

```
[edit class-of-service]
```

```
user@switch# set interfaces ae0 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
user@switch# set interfaces ae1 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
```

---

### MC-LAG Switches S1 and S2 Port Scheduling Configuration

---

#### Step-by-Step Procedure

To apply port scheduling for FCoE traffic to interfaces:

1. Apply the scheduler map to the egress ports:

```
set class-of-service interfaces ae0 scheduler-map fcoe-map
set class-of-service interfaces ae1 scheduler-map fcoe-map
```

---

### FCoE Transit Switches TS1 and TS2 Common Configuration (Applies to ETS and Port Scheduling)

---

#### Step-by-Step Procedure

The CoS configuration on FCoE Transit Switches TS1 and TS2 is similar to the CoS configuration on MC-LAG Switches S1 and S2. However, the port configurations differ, and you must enable FIP snooping on the Switch TS1 and Switch TS2 FCoE access ports.

To configure queue scheduling, PFC, the FCoE VLAN, and LAG interface membership and characteristics to support lossless FCoE transport across the MC-LAG (this example uses the default **fcoe** forwarding class and the default classifier to map incoming FCoE traffic to the FCoE IEEE 802.1p code point 011, so you do not configure them), or both ETS hierarchical scheduling and port scheduling (common configuration):

1. Configure output scheduling for the FCoE queue:

```
[edit class-of-service]
user@switch# set schedulers fcoe-sched priority low transmit-rate 3g
user@switch# set schedulers fcoe-sched shaping-rate percent 100
```

2. Map the FCoE forwarding class to the FCoE scheduler (**fcoe-sched**):

```
[edit class-of-service]
user@switch# set scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
```

3. Enable PFC on the FCoE priority by creating a congestion notification profile (**fcoe-cnp**) that applies FCoE to the IEEE 802.1 code point 011:

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011 pfc
```

4. Apply the PFC configuration to the LAG interface and to the FCoE access interfaces:

```
[edit class-of-service]
user@switch# set interfaces ae1 congestion-notification-profile fcoe-cnp
user@switch# set class-of-service interfaces xe-0/0/30 congestion-notification-profile fcoe-cnp
user@switch# set class-of-service interfaces xe-0/0/31 congestion-notification-profile fcoe-cnp
user@switch# set class-of-service interfaces xe-0/0/32 congestion-notification-profile fcoe-cnp
user@switch# set class-of-service interfaces xe-0/0/33 congestion-notification-profile fcoe-cnp
```

5. Configure the VLAN for FCoE traffic (**fcoe\_vlan**):

```
[edit vlans]
user@switch# set fcoe_vlan vlan-id 100
```

6. Add the member interfaces to the LAG:

```
[edit interfaces]
user@switch# set xe-0/0/25 ether-options 802.3ad ae1
user@switch# set xe-0/0/26 ether-options 802.3ad ae1
```

7. On the LAG (**ae1**), configure the interface mode as **trunk** and membership in the FCoE VLAN (**fcoe\_vlan**):

```
[edit interfaces]
user@switch# set interfaces ae1 unit 0 family ethernet-switching interface-mode trunk
vlan members fcoe_vlan
```

8. On the FCoE access interfaces (**xe-0/0/30**, **xe-0/0/31**, **xe-0/0/32**, **xe-0/0/33**), configure the interface mode as **trunk** and membership in the FCoE VLAN (**fcoe\_vlan**):

```
[edit interfaces]
user@switch# set interfaces xe-0/0/30 unit 0 family ethernet-switching interface-mode
trunk vlan members fcoe_vlan
user@switch# set interfaces xe-0/0/31 unit 0 family ethernet-switching interface-mode
trunk vlan members fcoe_vlan
user@switch# set interfaces xe-0/0/32 unit 0 family ethernet-switching interface-mode
trunk vlan members fcoe_vlan
user@switch# set interfaces xe-0/0/33 unit 0 family ethernet-switching interface-mode
trunk vlan members fcoe_vlan
```

9. Set the MTU to **2180** for the LAG and FCoE access interfaces. 2180 bytes is the minimum size required to handle FCoE packets because of the payload and header sizes; you can configure the MTU to a higher number of bytes if desired, but not less than 2180 bytes:

```
[edit interfaces]
user@switch# set ae1 mtu 2180
user@switch# set xe-0/0/30 mtu 2180
user@switch# set xe-0/0/31 mtu 2180
user@switch# set xe-0/0/32 mtu 2180
user@switch# set xe-0/0/33 mtu 2180
```

10. Set the LAG interface as an FCoE trusted port. Ports that connect to other switches should be trusted and should not perform FIP snooping:

```
[edit]
user@switch# set vlans fcoe_vlan forwarding-options fip-security interface ae1 fcoe-trusted
```



**NOTE:** Access ports **xe-0/0/30**, **xe-0/0/31**, **xe-0/0/32**, and **xe-0/0/33** are not configured as FCoE trusted ports. The access ports remain in the default state as untrusted ports because they connect directly to FCoE devices and must perform FIP snooping to ensure network security.

11. Enable FIP snooping on the FCoE VLAN to prevent unauthorized FCoE network access (this example uses **VN2VN\_Port** FIP snooping; the example is equally valid if you use **VN2VF\_Port** FIP snooping):

```
[edit]
```

```
user@switch# set vlans fcoe_vlan forwarding-options fip-security examine-vn2vn
beacon-period 90000
```



**NOTE:** QFX10000 switches do not support FIP snooping and cannot be used as FCoE access transit switches. (QFX10000 switches can be used as FCoE aggregation switches.)

### FCoE Transit Switches TS1 and TS2 ETS Hierarchical Scheduling Configuration

#### Step-by-Step Procedure

To configure the forwarding class set (priority group) and priority group scheduling (in a traffic control profile), and apply the ETS hierarchical scheduling for FCoE traffic to interfaces:

1. Configure the forwarding class set (**fcoe-pg**) for the FCoE traffic:
 

```
[edit class-of-service]
user@switch# set forwarding-class-sets fcoe-pg class fcoe
```
2. Define the traffic control profile (**fcoe-tcp**) to use on the FCoE forwarding class set:
 

```
[edit class-of-service]
user@switch# set traffic-control-profiles fcoe-tcp scheduler-map fcoe-map
guaranteed-rate 3g
user@switch# set traffic-control-profiles fcoe-tcp shaping-rate percent 100
```
3. Apply the FCoE forwarding class set and traffic control profile to the LAG interface and to the FCoE access interfaces:
 

```
[edit class-of-service]
user@switch# set interfaces ae1 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
user@switch# set class-of-service interfaces xe-0/0/30 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
user@switch# set class-of-service interfaces xe-0/0/31 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
user@switch# set class-of-service interfaces xe-0/0/32 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
user@switch# set class-of-service interfaces xe-0/0/33 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
```

### FCoE Transit Switches TS1 and TS2 Port Scheduling Configuration

#### Step-by-Step Procedure

To apply port scheduling for FCoE traffic to interfaces:

1. Apply the scheduler map to the egress ports:
 

```
user@switch# set class-of-service interfaces ae1 scheduler-map fcoe-map
user@switch# set class-of-service interfaces xe-0/0/30 scheduler-map fcoe-map
user@switch# set class-of-service interfaces xe-0/0/31 scheduler-map fcoe-map
user@switch# set class-of-service interfaces xe-0/0/32 scheduler-map fcoe-map
user@switch# set class-of-service interfaces xe-0/0/33 scheduler-map fcoe-map
```

## Results

Display the results of the CoS configuration on MC-LAG Switch S1 and on MC-LAG Switch S2 (the results on both switches are the same). The results are from the ETS hierarchical scheduling configuration, which shows the more complex configuration. Direct port scheduling results would not show the traffic control profile or forwarding class set portions of the configuration, but would display the name of the scheduler map under each interface (instead of the names of the forwarding class set and output traffic control profile). Other than that, they are the same.

```
user@switch> show configuration class-of-service
traffic-control-profiles {
 fcoe-tcp {
 scheduler-map fcoe-map;
 shaping-rate percent 100;
 guaranteed-rate 3000000000;
 }
}
forwarding-class-sets {
 fcoe-pg {
 class fcoe;
 }
}
congestion-notification-profile {
 fcoe-cnp {
 input {
 ieee-802.1 {
 code-point 011 {
 pfc;
 }
 }
 }
 }
}
interfaces {
 ae0 {
 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 congestion-notification-profile fcoe-cnp;
 }
 ae1 {
 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 congestion-notification-profile fcoe-cnp;
 }
}
scheduler-maps {
 fcoe-map {
```

```

 forwarding-class fcoe scheduler fcoe-sched;
 }
}
schedulers {
 fcoe-sched {
 transmit-rate 3000000000;
 shaping-rate percent 100;
 priority low;
 }
}

```



**NOTE:** The forwarding class and classifier configurations are not shown because the show command does not display default portions of the configuration.

For MC-LAG verification commands, see *Example: Configuring Multichassis Link Aggregation*.

Display the results of the CoS configuration on FCoE Transit Switch TS1 and on FCoE Transit Switch TS2 (the results on both transit switches are the same). The results are from the ETS hierarchical port scheduling configuration, which shows the more complex configuration. Direct port scheduling results would not show the traffic control profile or forwarding class set portions of the configuration, but would display the name of the scheduler map under each interface (instead of the names of the forwarding class set and output traffic control profile). Other than that, they are the same.

```

user@switch> show configuration class-of-service
traffic-control-profiles {
 fcoe-tcp {
 scheduler-map fcoe-map;
 shaping-rate percent 100;
 guaranteed-rate 3000000000;
 }
}
forwarding-class-sets {
 fcoe-pg {
 class fcoe;
 }
}
congestion-notification-profile {
 fcoe-cnp {
 input {
 ieee-802.1 {
 code-point 011 {
 pfc;
 }
 }
 }
 }
}
}
interfaces {
 xe-0/0/30 {

```

```

 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 congestion-notification-profile fcoe-cnp;
}
xe-0/0/31 {
 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 congestion-notification-profile fcoe-cnp;
}
xe-0/0/32 {
 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 congestion-notification-profile fcoe-cnp;
}
xe-0/0/33 {
 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 congestion-notification-profile fcoe-cnp;
}
ae1 {
 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 congestion-notification-profile fcoe-cnp;
}
}
scheduler-maps {
 fcoe-map {
 forwarding-class fcoe scheduler fcoe-sched;
 }
}
schedulers {
 fcoe-sched {
 transmit-rate 3000000000;
 shaping-rate percent 100;
 priority low;
 }
}
}

```



**NOTE:** The forwarding class and classifier configurations are not shown because the `show` command does not display default portions of the configuration.

## Verification

To verify that the CoS components and FIP snooping have been configured and are operating properly, perform these tasks. Because this example uses the default `fcoe` forwarding class and the default IEEE 802.1p trusted classifier, the verification of those configurations is not shown:

- [Verifying That the Output Queue Schedulers Have Been Created on page 6960](#)
- [Verifying That the Priority Group Output Scheduler \(Traffic Control Profile\) Has Been Created \(ETS Configuration Only\) on page 6961](#)
- [Verifying That the Forwarding Class Set \(Priority Group\) Has Been Created \(ETS Configuration Only\) on page 6961](#)
- [Verifying That Priority-Based Flow Control Has Been Enabled on page 6962](#)
- [Verifying That the Interface Class of Service Configuration Has Been Created on page 6963](#)
- [Verifying That the Interfaces Are Correctly Configured on page 6965](#)
- [Verifying That FIP Snooping Is Enabled on the FCoE VLAN on FCoE Transit Switches TS1 and TS2 Access Interfaces on page 6967](#)
- [Verifying That the FIP Snooping Mode Is Correct on FCoE Transit Switches TS1 and TS2 on page 6968](#)

### Verifying That the Output Queue Schedulers Have Been Created

**Purpose** Verify that the output queue scheduler for FCoE traffic has the correct bandwidth parameters and priorities, and is mapped to the correct forwarding class (output queue). Queue scheduler verification is the same on each of the four switches.

**Action** List the scheduler map using the operational mode command `show class-of-service scheduler-map fcoe-map`:

```
user@switch> show class-of-service scheduler-map fcoe-map
Scheduler map: fcoe-map, Index: 9023
```

```
Scheduler: fcoe-sched, Forwarding class: fcoe, Index: 37289
Transmit rate: 3000000000 bps, Rate Limit: none, Buffer size: remainder,
Buffer Limit: none, Priority: low
Excess Priority: unspecified
Shaping rate: 100 percent,
drop-profile-map-set-type: mark
Drop profiles:
 Loss priority Protocol Index Name
 Low any 1 <default-drop-profile>
 Medium high any 1 <default-drop-profile>
 High any 1 <default-drop-profile>
```



**Meaning** The **show class-of-service scheduler-map fcoe-map** command lists the properties of the scheduler map **fcoe-map**. The command output includes:

- The name of the scheduler map (**fcoe-map**)
- The name of the scheduler (**fcoe-sched**)
- The forwarding classes mapped to the scheduler (**fcoe**)
- The minimum guaranteed queue bandwidth (transmit rate **3000000000 bps**)
- The scheduling priority (**low**)
- The maximum bandwidth in the priority group the queue can consume (shaping rate **100 percent**)
- The drop profile loss priority for each drop profile name. This example does not include drop profiles because you do not apply drop profiles to FCoE traffic.

#### Verifying That the Priority Group Output Scheduler (Traffic Control Profile) Has Been Created (ETS Configuration Only)

**Purpose** Verify that the traffic control profile **fcoe-tcp** has been created with the correct bandwidth parameters and scheduler mapping. Priority group scheduler verification is the same on each of the four switches.

**Action** List the FCoE traffic control profile properties using the operational mode command **show class-of-service traffic-control-profile fcoe-tcp**:

```
user@switch> show class-of-service traffic-control-profile fcoe-tcp
Traffic control profile: fcoe-tcp, Index: 18303
 Shaping rate: 100 percent
 Scheduler map: fcoe-map
 Guaranteed rate: 3000000000
```

**Meaning** The **show class-of-service traffic-control-profile fcoe-tcp** command lists all of the configured traffic control profiles. For each traffic control profile, the command output includes:

- The name of the traffic control profile (**fcoe-tcp**)
- The maximum port bandwidth the priority group can consume (shaping rate **100 percent**)
- The scheduler map associated with the traffic control profile (**fcoe-map**)
- The minimum guaranteed priority group port bandwidth (guaranteed rate **3000000000** in bps)

#### Verifying That the Forwarding Class Set (Priority Group) Has Been Created (ETS Configuration Only)

**Purpose** Verify that the FCoE priority group has been created and that the **fcoe** priority (forwarding class) belongs to the FCoE priority group. Forwarding class set verification is the same on each of the four switches.

**Action** List the forwarding class sets using the operational mode command **show class-of-service forwarding-class-set fcoe-pg**:

```
user@switch> show class-of-service forwarding-class-set fcoe-pg
Forwarding class set: fcoe-pg, Type: normal-type, Forwarding class set index:
31420
 Forwarding class Index
 fcoe 1
```

**Meaning** The **show class-of-service forwarding-class-set fcoe-pg** command lists all of the forwarding classes (priorities) that belong to the **fcoe-pg** priority group, and the internal index number of the priority group. The command output shows that the forwarding class set **fcoe-pg** includes the forwarding class **fcoe**.

---

### Verifying That Priority-Based Flow Control Has Been Enabled

---

**Purpose** Verify that PFC is enabled on the FCoE code point. PFC verification is the same on each of the four switches.

**Action** List the FCoE congestion notification profile using the operational mode command **show class-of-service congestion-notification fcoe-cnp**:

```
user@switch> show class-of-service congestion-notification fcoe-cnp
Type: Input, Name: fcoe-cnp, Index: 6879
Cable Length: 100 m
 Priority PFC MRU
 000 Disabled
 001 Disabled
 010 Disabled
 011 Enabled 2500
 100 Disabled
 101 Disabled
 110 Disabled
 111 Disabled
Type: Output
 Priority Flow-Control-Queues
 000 0
 001 1
 010 2
 011 3
 100 4
 101 5
 110 6
 111 7
```

**Meaning** The **show class-of-service congestion-notification fcoe-cnp** command lists all of the IEEE 802.1p code points in the congestion notification profile that have PFC enabled. The command output shows that PFC is enabled on code point **011** (**fcoe** queue) for the **fcoe-cnp** congestion notification profile.

The command also shows the default cable length (100 meters), the default maximum receive unit (2500 bytes), and the default mapping of priorities to output queues because this example does not include configuring these options.

### Verifying That the Interface Class of Service Configuration Has Been Created

**Purpose** Verify that the CoS properties of the interfaces are correct. The verification output on MC-LAG Switches S1 and S2 differs from the output on FCoE Transit Switches TS1 and TS2.



**NOTE:** The output is from the ETS hierarchical port scheduling configuration to show the more complex configuration. Direct port scheduling results do not show the traffic control profile or forwarding class sets because those elements are configured only for ETS. Instead, the name of the scheduler map is displayed under each interface.

**Action** List the interface CoS configuration on MC-LAG Switches S1 and S2 using the operational mode command **show configuration class-of-service interfaces**:

```
user@switch> show configuration class-of-service interfaces
ae0 {
 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 congestion-notification-profile fcoe-cnp;
}

ae1 {
 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 congestion-notification-profile fcoe-cnp;
}
```

List the interface CoS configuration on FCoE Transit Switches TS1 and TS2 using the operational mode command **show configuration class-of-service interfaces**:

```
user@switch> show configuration class-of-service interfaces
xe-0/0/30 {
 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 congestion-notification-profile fcoe-cnp;
}
xe-0/0/31 {
 forwarding-class-set {
 fcoe-pg {
```

```

 output-traffic-control-profile fcoe-tcp;
 }
}
congestion-notification-profile fcoe-cnp;
}
xe-0/0/32 {
 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 congestion-notification-profile fcoe-cnp;
}
xe-0/0/33 {
 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 congestion-notification-profile fcoe-cnp;
}
ae1 {
 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 congestion-notification-profile fcoe-cnp;
}

```

**Meaning** The **show configuration class-of-service interfaces** command lists the class of service configuration for all interfaces. For each interface, the command output includes:

- The name of the interface (for example, **ae0** or **xe-0/0/30**)
- The name of the forwarding class set associated with the interface (**fcoe-pg**)
- The name of the traffic control profile associated with the interface (output traffic control profile, **fcoe-tcp**)
- The name of the congestion notification profile associated with the interface (**fcoe-cnp**)



**NOTE:** Interfaces that are members of a LAG are not shown individually. The LAG or MC-LAG CoS configuration is applied to all interfaces that are members of the LAG or MC-LAG. For example, the interface CoS configuration output on MC-LAG Switches S1 and S2 shows the LAG CoS configuration but does not show the CoS configuration of the member interfaces separately. The interface CoS configuration output on FCoE Transit Switches TS1 and TS2 shows the LAG CoS configuration but also shows the configuration for interfaces xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33, which are not members of a LAG.

### Verifying That the Interfaces Are Correctly Configured

**Purpose** Verify that the LAG membership, MTU, VLAN membership, and port mode of the interfaces are correct. The verification output on MC-LAG Switches S1 and S2 differs from the output on FCoE Transit Switches T1 and T2.

**Action** List the interface configuration on MC-LAG Switches S1 and S2 using the operational mode command **show configuration interfaces**:

```
user@switch> show configuration interfaces
```

```
xe-0/0/10 {
 ether-options {
 802.3ad ae0;
 }
}
xe-0/0/11 {
 ether-options {
 802.3ad ae0;
 }
}
xe-0/0/20 {
 ether-options {
 802.3ad ae1;
 }
}
xe-0/0/21 {
 ether-options {
 802.3ad ae1;
 }
}
ae0 {
 mtu 2180;
 unit 0 {
 family ethernet-switching {
 interface-mode trunk;
 vlan {
 members fcoe_vlan;
 }
 }
 }
}
ae1 {
 mtu 2180;
 unit 0 {
 family ethernet-switching {
 interface-mode trunk;
 vlan {
 members fcoe_vlan;
 }
 }
 }
}
```

List the interface configuration on FCoE Transit Switches TS1 and TS2 using the operational mode command **show configuration interfaces**:

```
user@switch> show configuration interfaces
```

```
xe-0/0/25 {
 ether-options {
 802.3ad ae1;
 }
}
xe-0/0/26 {
 ether-options {
 802.3ad ae1;
 }
}
xe-0/0/30 {
 mtu 2180;
 unit 0 {
 family ethernet-switching {
 interface-mode trunk;
 vlan {
 members fcoe_vlan;
 }
 }
 }
}
xe-0/0/31 {
 mtu 2180;
 unit 0 {
 family ethernet-switching {
 interface-mode trunk;
 vlan {
 members fcoe_vlan;
 }
 }
 }
}
xe-0/0/32 {
 mtu 2180;
 unit 0 {
 family ethernet-switching {
 interface-mode trunk;
 vlan {
 members fcoe_vlan;
 }
 }
 }
}
xe-0/0/33 {
 mtu 2180;
 unit 0 {
 family ethernet-switching {
 interface-mode trunk;
 vlan {
 members fcoe_vlan;
 }
 }
 }
}
ae1 {
 mtu 2180;
 unit 0 {
 family ethernet-switching {
 interface-mode trunk;
 vlan {
```

```

 members fcoe_vlan;
 }
}
}

```

**Meaning** The **show configuration interfaces** command lists the configuration of each interface by interface name.

For each interface that is a member of a LAG, the command lists only the name of the LAG to which the interface belongs.

For each LAG interface and for each interface that is not a member of a LAG, the command output includes:

- The MTU (**2180**)
- The unit number of the interface (**0**)
- The interface mode (**trunk** mode both for interfaces that connect two switches and for interfaces that connect to FCoE hosts)
- The name of the VLAN in which the interface is a member (**fcoe\_vlan**)

#### Verifying That FIP Snooping Is Enabled on the FCoE VLAN on FCoE Transit Switches TS1 and TS2 Access Interfaces

**Purpose** Verify that FIP snooping is enabled on the FCoE VLAN access interfaces. FIP snooping is enabled only on the FCoE access interfaces, so it is enabled only on FCoE Transit Switches TS1 and TS2. FIP snooping is not enabled on MC-LAG Switches S1 and S2 because FIP snooping is done at the Transit Switch TS1 and TS2 FCoE access ports.

**Action** List the port security configuration on FCoE Transit Switches TS1 and TS2 using the operational mode command **show configuration vlans fcoe\_vlan forwarding-options fip-security**:

```

user@switch> show configuration vlans fcoe_vlan forwarding-options fip-security
interface ae1.0 {
 fcoe-trusted;
}
examine-vn2vn {
 beacon-period 90000;
}

```

**Meaning** The **show configuration vlans fcoe\_vlan forwarding-options fip-security** command lists VLAN FIP security information, including whether a port member of the VLAN is trusted. The command output shows that:

- LAG port **ae1.0**, which connects the FCoE transit switch to the MC-LAG switches, is configured as an FCoE trusted interface. FIP snooping is not performed on the member interfaces of the LAG (**xe-0/0/25** and **xe-0/0/26**).
- VN2VN\_Port FIP snooping is enabled (**examine-vn2vn**) on the FCoE VLAN and the beacon period is set to 90000 milliseconds. On Transit Switches TS1 and TS2, all

interface members of the FCoE VLAN perform FIP snooping unless the interface is configured as FCoE trusted. On Transit Switches TS1 and TS2, interfaces **xe-0/0/30**, **xe-0/0/31**, **xe-0/0/32**, and **xe-0/0/33** perform FIP snooping because they are not configured as FCoE trusted. The interface members of LAG **ae1** (**xe-0/0/25** and **xe-0/0/26**) do not perform FIP snooping because the LAG is configured as FCoE trusted.

### Verifying That the FIP Snooping Mode Is Correct on FCoE Transit Switches TS1 and TS2

---

**Purpose** Verify that the FIP snooping mode is correct on the FCoE VLAN. FIP snooping is enabled only on the FCoE access interfaces, so it is enabled only on FCoE Transit Switches TS1 and TS2. FIP snooping is not enabled on MC-LAG Switches S1 and S2 because FIP snooping is done at the Transit Switch TS1 and TS2 FCoE access ports.

**Action** List the FIP snooping configuration on FCoE Transit Switches TS1 and TS2 using the operational mode command **show fip snooping brief**:

```
user@switch> show fip snooping brief
VLAN: fcoe_vlan, Mode: VN2VN Snooping
FC-MAP: 0e:fc:00
...
```



**NOTE:** The output has been truncated to show only the relevant information.

---

**Meaning** The **show fip snooping brief** command lists FIP snooping information, including the FIP snooping VLAN and the FIP snooping mode. The command output shows that:

- The VLAN on which FIP snooping is enabled is **fcoe\_vlan**
- The FIP snooping mode is VN2VN\_Port FIP snooping (**VN2VN Snooping**)

**Related  
Documentation**

- *Example: Configuring Multichassis Link Aggregation*
- [Configuring Link Aggregation on page 2869](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 6504](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 6771](#)
- [Example: Configuring Queue Schedulers for Port Scheduling on page 6744](#)
- *Understanding Multichassis Link Aggregation*
- [Understanding MC-LAGs on an FCoE Transit Switch on page 6451](#)



## Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch)

The default system configuration supports FCoE traffic on priority 3 (IEEE 802.1p code point 011). If the FCoE traffic on your converged Ethernet network uses priority 3, the only user configuration required for lossless transport is to enable PFC on code point 011 on the FCoE ingress interfaces.

However, if your network uses a different priority than 3 for FCoE traffic, you need to configure lossless FCoE transport on that priority. This example shows you how to configure lossless FCoE transport on a converged Ethernet network that uses priority 5 (IEEE 802.1p code point 101) for FCoE traffic instead of using priority 3.

- [Requirements on page 6969](#)
- [Overview on page 6969](#)
- [Configuration on page 6971](#)
- [Verification on page 6973](#)

### Requirements

This example uses the following hardware and software components:

- One switch used as an FCoE transit switch
- Junos OS Release 12.3 or later for the QFX Series

### Overview

Although FCoE traffic typically uses IEEE 802.1p priority 3 on converged Ethernet networks, some networks use a different priority for FCoE traffic. Regardless of the priority used, FCoE traffic must receive lossless treatment. Supporting lossless behavior for FCoE traffic when your network does not use priority 3 requires configuring:

- A lossless forwarding class for FCoE traffic.
- A behavior aggregate (BA) classifier to map the FCoE forwarding class to the appropriate IEEE 802.1p priority.
- A congestion notification profile (CNP) to enable PFC on the FCoE code point at the interface ingress and to configure flow control on the interface egress. Flow control on the interface egress enables the interface to respond to PFC messages received from the connected peer and pause the correct IEEE 802.1p priority on the correct output queue.



**NOTE:** Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

- A DCBX application and an application map to support DCBX application TLV exchange for the lossless FCoE traffic on the configured FCoE priority. By default, DCBX is enabled on all Ethernet interfaces, but only on priority 3 (IEEE 802.1p code point 011). To support DCBX application TLV exchange when you are not using the default configuration, you must configure all of the applications and map them to interfaces and priorities.

The priorities specified in the BA classifiers, CNP, and DCBX application map must match, or the configuration does not work. You must specify the same lossless FCoE forwarding class in each configuration and use the same IEEE 802.1p code point (priority) so that the FCoE traffic is properly classified into flows and so that those flows receive lossless treatment.

### Topology

This example shows how to configure one lossless FCoE traffic class, map it to a priority other than priority 3, and configure flow control to ensure lossless behavior on the interfaces. This example uses two Ethernet interfaces, xe-0/0/25 and xe-0/0/26. The interfaces connect to a converged Ethernet network that uses IEEE 802.1p priority 5 (code point 101) for FCoE traffic.

The configuration on the two interfaces is the same. Both interfaces use the same explicitly configured lossless FCoE forwarding class and the same ingress classifier. Both interfaces enable PFC on priority 5 and enable flow control on the same output queue (which is mapped to the lossless FCoE forwarding class).

Table 595 shows the configuration components for this example.

**Table 595: Components of the Configuration Topology for FCoE Traffic That Does Not Use Priority 3**

| Component        | Settings                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Hardware         | One switch                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Forwarding class | Name— <b>fcoe1</b><br><br>Queue mapping—queue 5<br><br>Packet drop attribute— <b>no-loss</b><br><br><b>NOTE:</b> A lossless forwarding class can be mapped to any output queue. However, because the <b>fcoe1</b> forwarding class uses priority 5 in this example, matching that traffic to a forwarding class that uses queue 5 creates a configuration that is logical and easy to map because the priority and the queue are identified by the same number. |

**Table 595: Components of the Configuration Topology for FCoE Traffic That Does Not Use Priority 3 (*continued*)**

| Component                | Settings                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| BA classifier            | Name— <b>fcoe_p5</b><br><br>FCoE priority mapping—Forwarding class <b>fcoe1</b> mapped to code point <b>101</b> (IEEE 802.1p priority 5) and a packet loss priority of <b>low</b> .                                                                                                                                                                                                                                                                                                                                                                                                             |
| PFC configuration (CNPs) | CNP name— <b>fcoe_p5_cnp</b><br><br>Input CNP code point— <b>101</b><br><br>MRU— <b>2240</b> bytes<br><br>Cable length— <b>100</b> meters<br><br>Output CNP code point— <b>101</b><br><br>Output CNP flow control queue— <b>5</b><br><br><b>NOTE:</b> When you apply a CNP with an explicit output queue flow control configuration to an interface, the explicit CNP overwrites the default output CNP. The output queues that are enabled for pause in the default configuration (queues 3 and 4) are not enabled for pause unless they are included in the explicitly configured output CNP. |
| DCBX application mapping | Application name— <b>fcoe_p5_app</b><br><br>Application EtherType— <b>0x8906</b><br><br>Application map name— <b>fcoe_p5_app_map</b><br><br>Application map code points— <b>101</b><br><br><b>NOTE:</b> LLDP and DCBX must be enabled on the interface. By default, LLDP and DCBX are enabled on all Ethernet interfaces.                                                                                                                                                                                                                                                                       |



**NOTE:** This example does not include scheduling (bandwidth allocation) configuration or the FIP snooping configuration. This example focuses only on the lossless FCoE priority configuration.

QFX10000 switches do not support FIP snooping. For this reason, QFX10000 switches cannot be used as FCoE access transit switches. QFX10000 switches can be used as intermediate or aggregation transit switches in the FCoE path, between an FCoE access transit switch that performs FIP snooping and an FCF.

## Configuration

### CLI Quick Configuration

To quickly configure a lossless FCoE forwarding class that uses a different priority than IEEE 802.1p priority 3 for FCoE traffic on an FCoE transit switch, copy the following commands, paste them in a text file, remove line breaks, change variables and details

to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set class-of-service forwarding-classes class fcoe1 queue-num 5 no-loss
set class-of-service classifiers ieee-802.1 fcoe_p5 forwarding-class fcoe1 loss-priority low
code-points 101
set class-of-service interfaces xe-0/0/25 unit 0 classifiers ieee-802.1 fcoe_p5
set class-of-service interfaces xe-0/0/26 unit 0 classifiers ieee-802.1 fcoe_p5
set class-of-service congestion-notification-profile fcoe_p5_cnp input ieee-802.1 code-point 101
pfc mru 2240
set class-of-service congestion-notification-profile fcoe_p5_cnp input cable-length 100
set class-of-service congestion-notification-profile fcoe_p5_cnp output ieee-802.1 code-point
101 pfc flow-control-queue 5
set class-of-service interfaces xe-0/0/25 congestion-notification-profile fcoe_p5_cnp
set class-of-service interfaces xe-0/0/26 congestion-notification-profile fcoe_p5_cnp
set applications application fcoe_p5_app ether-type 0x8906
set policy-options application-maps fcoe_p5_app_map application fcoe_p5_app code-points 101
set protocols dcbx interface xe-0/0/25 application-map fcoe_p5_app_map
set protocols dcbx interface xe-0/0/26 application-map fcoe_p5_app_map
```

### Configuring A Lossless FCoE Forwarding Class On IEEE 802.1p Priority 5

#### Step-by-Step Procedure

To configure a lossless forwarding class for FCoE traffic on IEEE 802.1p priority 5 (code point 101), classify FCoE traffic into the lossless forwarding class, configure a congestion notification profile to enable PFC on the FCoE priority and output queue, and configure DCBX application protocol TLV exchange for traffic on the FCoE priority:

1. Configure the lossless forwarding class (named **fcoe1** and mapped to output queue 5) for FCoE traffic on IEEE 802.1p priority 5:

```
[edit class-of-service]
user@switch# set forwarding-classes class fcoe1 queue-num 5 no-loss
```

2. Configure the ingress classifier (**fcoe\_p5**). The classifier maps the FCoE priority (code point 101) to the lossless FCoE forwarding class **fcoe1**:

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 fcoe_p5 forwarding-class fcoe1 loss-priority low code-points
101
```

3. Apply the classifier to interfaces **xe-0/0/25** and **xe-0/0/26**:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/25 unit 0 classifiers ieee-802.1 fcoe_p5
user@switch# set interfaces xe-0/0/26 unit 0 classifiers ieee-802.1 fcoe_p5
```

4. Configure the CNP. The input stanza enables PFC on the FCoE priority (IEEE 802.1p code point 101), sets the MRU value (2240 bytes), and sets the cable length value (100 meters). The output stanza configures flow control on output queue 5 on the FCoE priority:

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe_p5_cnp input ieee-802.1 code-point
101 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_p5_cnp input cable-length 100
user@switch# set congestion-notification-profile fcoe_p5_cnp output ieee-802.1 code-point
101 pfc flow-control-queue 5
```

5. Apply the CNP to the interfaces:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/25 congestion-notification-profile fcoe_p5_cnp
user@switch# set interfaces xe-0/0/26 congestion-notification-profile fcoe_p5_cnp
```

6. Configure the DCBX application for FCoE to map to the Ethernet interfaces, so that DCBX can exchange application protocol TLVs on the IEEE 802.1p priority 5 instead of on the default priority 3:

```
[edit]
user@switch# set applications application fcoe_p5_app ether-type 0x8906
```

7. Configure a DCBX application map to map the FCoE application to the correct IEEE 802.1p FCoE priority:

```
[edit]
user@switch# set policy-options application-maps fcoe_p5_app_map application
fcoe_p5_app code-points 101
```

8. Apply the application map to the Ethernet interfaces so that DCBX exchanges FCoE application TLVs on the correct code point:

```
[edit]
user@switch# set protocols dcbx interface xe-0/0/25 application-map fcoe_p5_app_map
user@switch# set protocols dcbx interface xe-0/0/26 application-map fcoe_p5_app_map
```

## Verification

To verify the configuration and proper operation of the lossless forwarding class and IEEE 802.1p priority, perform these tasks:

- [Verifying the Forwarding Class Configuration on page 6973](#)
- [Verifying the Behavior Aggregate Classifier Configuration on page 6974](#)
- [Verifying the PFC Flow Control Configuration \(CNP\) on page 6974](#)
- [Verifying the Interface Configuration on page 6975](#)
- [Verifying the DCBX Application Configuration on page 6976](#)
- [Verifying the DCBX Application Map Configuration on page 6976](#)
- [Verifying the DCBX Application Protocol Exchange Interface Configuration on page 6976](#)

### Verifying the Forwarding Class Configuration

**Purpose** Verify that the lossless forwarding class **fcoe1** has been created.

**Action** Show the forwarding class configuration by using the operational command **show class-of-service forwarding-class**:

```
user@switch# show class-of-service forwarding-class
```

| Forwarding class | ID | Queue | Policing priority | No-Loss  |
|------------------|----|-------|-------------------|----------|
| best-effort      | 0  | 0     | normal            | Disabled |
| fcoe             | 1  | 3     | normal            | Enabled  |
| no-loss          | 2  | 4     | normal            | Enabled  |
| network-control  | 3  | 7     | normal            | Disabled |
| fcoe1            | 4  | 5     | normal            | Enabled  |
| mcast            | 8  | 8     | normal            | Disabled |

**Meaning** The **show class-of-service forwarding-class** command shows all of the forwarding classes. The command output shows that the **fcoe1** forwarding class is configured on output queue 5 with the no-loss packet drop attribute enabled.

Because we did not explicitly configure the default forwarding classes, they remain in their default state, including the lossless configuration of the **fcoe** and **no-loss** default forwarding classes.

### Verifying the Behavior Aggregate Classifier Configuration

**Purpose** Verify that the classifier maps the forwarding classes to the correct IEEE 802.1p code points (priorities) and packet loss priorities.

**Action** List the classifier configured to support lossless FCoE transport using the operational mode command **show class-of-service classifier**:

```
user@switch> show class-of-service classifier
Classifier: fcoe_p5, Code point type: ieee-802.1, Index: 63065
 Code point Forwarding class Loss priority
 101 fcoe1 low
```

**Meaning** The **show class-of-service classifier** command shows the IEEE 802.1p code points and the loss priorities that are mapped to the forwarding classes in each classifier.

Classifier **fcoe\_p5** maps code point **101** (priority 5) to explicitly configured lossless forwarding class **fcoe1** and a packet loss priority of **low**, and all other priorities to the **best-effort** forwarding class with a packet loss priority of **high**.

### Verifying the PFC Flow Control Configuration (CNP)

**Purpose** Verify that PFC is enabled on the correct input priority and that flow control is configured on the correct output queue in the CNP.

**Action** Display the congestion notification profile using the operational mode command **show class-of-service congestion-notification**:

```
user@switch> show class-of-service congestion-notification
Name: fcoe_p5_cnp, Index: 12137
Type: Input
Cable Length: 100 m
 Priority PFC MRU
 000 Disabled
 001 Disabled
 010 Disabled
 011 Disabled
 100 Disabled
 101 Enabled 2240
 110 Disabled
 111 Disabled
Type: Output
 Priority Flow-Control-Queues
 101
 5
```

**Meaning** The **show class-of-service congestion-notification** command shows the input and output stanzas of the configured CNPs.

The **fcoe\_p5\_cnp** CNP input stanza shows that PFC is enabled on code point **101** (priority 5), the MRU is **2240** bytes, and the cable length is **100** meters. The CNP output stanza shows that output flow control is configured on queue **5** for code point **101** (priority 5).

### Verifying the Interface Configuration

**Purpose** Verify that the correct classifier and congestion notification profile are configured on the interfaces.

**Action** List the ingress interfaces using the operational mode commands **show configuration class-of-service interfaces xe-0/0/25** and **show configuration class-of-service interfaces xe-0/0/26**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/25
congestion-notification-profile fcoe_p5_cnp;
unit 0 {
 classifiers {
 ieee-802.1 fcoe_p5;
 }
}

user@switch> show configuration class-of-service interfaces xe-0/0/26
congestion-notification-profile fcoe_p5_cnp;
unit 0 {
 classifiers {
 ieee-802.1 fcoe_p5;
 }
}
```

**Meaning** Both the **show configuration class-of-service interfaces xe-0/0/25** command and the **show configuration class-of-service interfaces xe-0/0/26** command show that the

congestion notification profile **fcoe\_p5\_cnp** is configured on each interface, and that the IEEE 802.1p classifier associated with each interface is **fcoe\_p5**.

### Verifying the DCBX Application Configuration

- Purpose** Verify that the DCBX application for FCoE is configured.
- Action** List the DCBX applications by using the configuration mode command **show applications**:
- ```
user@switch# show applications
application fcoe_p5_app {
    ether-type 0x8906;
```
- Meaning** The **show applications** configuration mode command shows all of the configured applications. The output shows that the application **fcoe_p5_app** is configured with an EtherType of **0x8906**.

Verifying the DCBX Application Map Configuration

- Purpose** Verify that the application map is configured.
- Action** List the application maps by using the configuration mode command **show policy-options application-maps**:
- ```
user@switch# show policy-options application-maps
fcoe_p5_app_map {
 application fcoe_p5_app code-points 101;
}
```
- Meaning** The **show policy-options application-maps** configuration mode command lists all of the configured application maps and the applications that belong to each application map. The output shows that application map **fcoe\_p5\_app\_map** consists of the application named **fcoe\_p5\_app**, which is mapped to IEEE 802.1p code point 101.

### Verifying the DCBX Application Protocol Exchange Interface Configuration

- Purpose** Verify that the application map is applied to the correct interfaces.
- Action** List the application maps on each interface using the configuration mode command **show protocols dcbx**:
- ```
user@switch# show protocols dcbx
interface xe-0/0/25.0 {
    application-map fcoe_p5_app_map;
}
interface xe-0/0/26.0 {
    application-map fcoe_p5_app_map;
}
```
- Meaning** The **show protocols dcbx** configuration mode command lists the application map association with interfaces. The output shows that interfaces **xe-0/0/25.0** and **xe-0/0/26.0** use application map **fcoe_p5_app_map**.

- Related Documentation**
- [Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 6986](#)
 - [Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 6977](#)
 - [Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications \(FCoE and iSCSI\) on page 7000](#)
 - [Example: Configuring DCBX Application Protocol TLV Exchange on page 6535](#)
 - [Configuring CoS PFC \(Congestion Notification Profiles\) on page 6905](#)
 - [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 6873](#)
 - [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 6493](#)

Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface

The default system configuration supports FCoE traffic on priority 3 (IEEE 802.1p code point 011). If the FCoE traffic on your converged Ethernet network uses priority 3, the only user configuration required for lossless transport is to enable PFC on code point 011 on the FCoE ingress interfaces.

However, if your converged Ethernet network uses more than one priority for FCoE traffic, you need to configure lossless transport for each FCoE priority. This example shows you how to configure lossless FCoE transport on a converged Ethernet network that uses both priority 3 (IEEE 802.1p code point 011) and priority 5 (IEEE 802.1p code point 101) for FCoE traffic.

- [Requirements on page 6977](#)
- [Overview on page 6977](#)
- [Configuration on page 6980](#)
- [Verification on page 6982](#)

Requirements

This example uses the following hardware and software components:

- One switch used as an FCoE transit switch
- Junos OS Release 12.3 or later for the QFX Series

Overview

Some network topologies support FCoE traffic on more than one IEEE 802.1p priority. For example, a converged Ethernet network might include two separate FCoE networks that use different priorities to identify traffic. Interfaces that carry traffic for both FCoE networks need to support lossless FCoE transport on both priorities.

Supporting lossless behavior for two FCoE traffic classes requires configuring:

- At least one lossless forwarding class for FCoE traffic (this example uses the default **fcoe** forwarding class as one of the lossless FCoE forwarding classes, so we need to explicitly configure only one FCoE forwarding class).
- A behavior aggregate (BA) classifier to map the FCoE forwarding classes to the appropriate IEEE 802.1p code points (priorities).
- A congestion notification profile (CNP) to enable PFC on the FCoE code points at the interface ingress and to configure PFC flow control on the interface egress so that the interface can respond to PFC messages received from the connected peer.



NOTE: Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

- DCBX applications and an application map to support DCBX application TLV exchange for the lossless FCoE traffic on the configured FCoE priorities. By default, DCBX is enabled on all Ethernet interfaces, but only on priority 3 (IEEE 802.1p code point 011). To support DCBX application TLV exchange when you are not using the default configuration, you must configure all of the applications and map them to interfaces and priorities.

The priorities specified in the BA classifier, CNP, and DCBX application map must match, or the configuration does not work. You must specify the same lossless FCoE forwarding class in each configuration and use the same IEEE 802.1p code point (priority) so that the FCoE traffic is properly classified into flows and so that those flows receive lossless treatment.

Topology

This example shows how to configure two lossless FCoE traffic classes on an interface, map them to two different priorities, and configure flow control to ensure lossless behavior. This example uses two Ethernet interfaces, xe-0/0/20 and xe-0/0/21, that are connected to the converged Ethernet network. Both interfaces transport FCoE traffic on priorities 3 (011) and 5 (101), and must support lossless transport of that traffic.

Table 596 shows the configuration components for this example.

Table 596: Components of the Two Lossless FCoE Priorities on an Interface Configuration Topology

Component	Settings
Hardware	One switch

Table 596: Components of the Two Lossless FCoE Priorities on an Interface Configuration Topology (*continued*)

Component	Settings
Forwarding classes	<p>Name—fcoe1 Queue mapping—queue 5 Packet drop attribute—no-loss</p> <p>NOTE: A lossless forwarding class can be mapped to any output queue. However, because the fcoe1 forwarding class uses priority 5 in this example, matching that traffic to a forwarding class that uses queue 5 creates a configuration that is logical and easy to map because the priority and the queue are identified by the same number.</p> <p>Name—fcoe This is the default lossless FCoE forwarding class, so no configuration required. The fcoe forwarding class is mapped to priority 3 (IEEE 802.1p code point 011) and to output queue 3 with a packet drop attribute of no-loss.</p>
BA classifier	<p>Name—fcoe_classifier</p> <p>FCoE priority mapping for forwarding class fcoe—mapped to code point 011 (IEEE 802.1p priority 3) and a packet loss priority of low.</p> <p>FCoE priority mapping for forwarding class fcoe1—mapped to code point 101 (IEEE 802.1p priority 5) and a packet loss priority of low.</p>
PFC configuration (CNP)	<p>CNP name—fcoe_cnp</p> <p>Input CNP code points—011 and 101</p> <p>MRU—2240 bytes</p> <p>Cable length—100 meters</p> <p>Output CNP code points—011 and 101</p> <p>Output CNP flow control queues—3 and 5</p> <p>NOTE: When you apply a CNP with an explicit output queue flow control configuration to an interface, the explicit CNP overwrites the default output CNP. The output queues that are enabled for PFC pause in the default configuration (queues 3 and 4) are not enabled for PFC pause unless they are included in the explicitly configured output CNP. In this example, because the explicit output CNP overwrites the default output CNP, we must explicitly configure flow control on queue 3.</p>

Table 596: Components of the Two Lossless FCoE Priorities on an Interface Configuration Topology (*continued*)

Component	Settings
DCBX application mapping	Application name— fcoe_app Application EtherType— 0x8906 Application map name— fcoe_app_map Application map code points— 011 and 101 NOTE: LLDP and DCBX must be enabled on the interface. By default, LLDP and DCBX are enabled on all Ethernet interfaces.
Interfaces	Interfaces xe-0/0/20 and xe-0/0/21 use the same configuration: <ul style="list-style-type: none"> • Classifier—fcoe_classifier • CNP—fcoe_cnp • DCBX application map—fcoe_app_map



NOTE: This example does not include scheduling (bandwidth allocation) configuration or the FIP snooping configuration. This examples focuses only on the lossless FCoE priority configuration.

QFX10000 switches do not support FIP snooping. For this reason, QFX10000 switches cannot be used as FCoE access transit switches. QFX10000 switches can be used as intermediate or aggregation transit switches in the FCoE path, between an FCoE access transit switch that performs FIP snooping and an FCF.

Configuration

CLI Quick Configuration

To quickly configure two lossless FCoE forwarding classes that use different priorities on an FCoE transit switch interface, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set class-of-service forwarding-classes class fcoe1 queue-num 5 no-loss
set class-of-service classifiers ieee-802.1 fcoe_classifier forwarding-class fcoe loss-priority low
code-points 011
set class-of-service classifiers ieee-802.1 fcoe_classifier forwarding-class fcoe1 loss-priority low
code-points 101set class-of-service interfaces xe-0/0/20 unit 0 classifiers ieee-802.1 fcoe_classifier
set class-of-service interfaces xe-0/0/21 unit 0 classifiers ieee-802.1 fcoe_classifier
set class-of-service congestion-notification-profile fcoe_cnp input ieee-802.1 code-point 011 pfc
mru 2240
set class-of-service congestion-notification-profile fcoe_cnp input ieee-802.1 code-point 101 pfc
mru 2240
set class-of-service congestion-notification-profile fcoe_cnp input cable-length 100
set class-of-service congestion-notification-profile fcoe_cnp output ieee-802.1 code-point 011
pfc flow-control-queue 3
set class-of-service congestion-notification-profile fcoe_cnp output ieee-802.1 code-point 101
pfc flow-control-queue 5
```

```

set class-of-service interfaces xe-0/0/20 congestion-notification-profile fcoe_cnp
set class-of-service interfaces xe-0/0/21 congestion-notification-profile fcoe_cnp
set applications application fcoe_app ether-type 0x8906
set policy-options application-maps fcoe_app_map application fcoe_app code-points [011 101]
set protocols dcbx interface xe-0/0/20 application-map fcoe_app_map
set protocols dcbx interface xe-0/0/21 application-map fcoe_app_map

```

Step-by-Step Procedure

To configure two lossless forwarding classes for FCoE traffic on the same interface, classify FCoE traffic into the forwarding classes, configure CNPs to enable PFC on the FCoE priorities and output queues, and configure DCBX application protocol TLV exchange for traffic on both FCoE priorities:

1. Configure lossless forwarding class **fcoe1** and map it to output queue **5** for FCoE traffic that uses IEEE 802.1p priority 5:

```

[edit class-of-service]
user@switch# set forwarding-classes class fcoe1 queue-num 5 no-loss

```



NOTE: This examples uses the default **fcoe** forwarding class as the other lossless FCoE forwarding class.

2. Configure the ingress classifier. The classifier maps the FCoE priorities (IEEE 802.1p code points **011** and **101**) to lossless FCoE forwarding classes **fcoe** and **fcoe1**, respectively:

```

[edit class-of-service classifiers]
user@switch# set ieee-802.1 fcoe_classifier forwarding-class fcoe loss-priority low code-points 011
user@switch# set ieee-802.1 fcoe_classifier forwarding-class fcoe1 loss-priority low code-points 101

```

3. Apply the classifier to the interfaces:

```

[edit class-of-service]
user@switch# set interfaces xe-0/0/20 unit 0 classifiers ieee-802.1 fcoe_classifier
user@switch# set interfaces xe-0/0/21 unit 0 classifiers ieee-802.1 fcoe_classifier

```

4. Configure the CNP. The input stanza enables PFC on the FCoE priorities (IEEE 802.1p code points **011** and **101**), sets the MRU value (2240 bytes), and sets the cable length value (100 meters). The output stanza configures flow control on output queues **3** and **5** on the FCoE priorities:

```

[edit class-of-service]
user@switch# set congestion-notification-profile fcoe_cnp input ieee-802.1 code-point 011 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_cnp input ieee-802.1 code-point 101 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_cnp input cable-length 100
user@switch# set congestion-notification-profile fcoe_cnp output ieee-802.1 code-point 011 pfc flow-control-queue 3
user@switch# set congestion-notification-profile fcoe_cnp output ieee-802.1 code-point 101 pfc flow-control-queue 5

```

5. Apply the CNP to the interfaces:

```

[edit class-of-service]

```

```
user@switch# set interfaces xe-0/0/20 congestion-notification-profile fcoe_cnp
user@switch# set interfaces xe-0/0/21 congestion-notification-profile fcoe_cnp
```

6. Configure a DCBX application for FCoE to map to the Ethernet interfaces, so that DCBX can exchange application protocol TLVs on both of the IEEE 802.1p priorities used for FCoE transport:

```
[edit]
user@switch# set applications application fcoe_app ether-type 0x8906
```

7. Configure a DCBX application map to map the FCoE application to the correct IEEE 802.1p FCoE priorities:

```
[edit]
user@switch# set policy-options application-maps fcoe_app_map application fcoe_app
code-points [011 101]
```

8. Apply the application map to the interfaces so that DCBX exchanges FCoE application TLVs on the correct code points:

```
[edit]
user@switch# set protocols dcbx interface xe-0/0/20 application-map fcoe_app_map
user@switch# set protocols dcbx interface xe-0/0/21 application-map fcoe_app_map
```

Verification

To verify the configuration and proper operation of the lossless forwarding classes and IEEE 802.1p priorities, perform these tasks:

- [Verifying the Forwarding Class Configuration on page 6982](#)
- [Verifying the Behavior Aggregate Classifier Configuration on page 6983](#)
- [Verifying the PFC Flow Control Configuration \(CNP\) on page 6983](#)
- [Verifying the Interface Configuration on page 6984](#)
- [Verifying the DCBX Application Configuration on page 6985](#)
- [Verifying the DCBX Application Map Configuration on page 6985](#)
- [Verifying the DCBX Application Protocol Exchange Interface Configuration on page 6985](#)

Verifying the Forwarding Class Configuration

Purpose Verify that the lossless forwarding class **fcoe1** has been created.

Action Show the forwarding class configuration by using the operational command **show class-of-service forwarding class**:

```
user@switch# show class-of-service forwarding-class
```

Forwarding class	ID	Queue	Policing priority	No-Loss
best-effort	0	0	normal	Disabled
fcoe	1	3	normal	Enabled
no-loss	2	4	normal	Enabled
network-control	3	7	normal	Disabled
fcoe1	4	5	normal	Enabled
mcast	8	8	normal	Disabled

Meaning The **show class-of-service forwarding-class** command shows all of the forwarding classes. The command output shows that the **fcoe1** forwarding class is configured on output queue 5 with the no-loss packet drop attribute enabled.

Because we did not explicitly configure the default forwarding classes, they remain in their default state, including the lossless configuration of the **fcoe** and **no-loss** default forwarding classes.

Verifying the Behavior Aggregate Classifier Configuration

Purpose Verify that the three classifiers map the forwarding classes to the correct IEEE 802.1p code points (priorities) and packet loss priorities.

Action List the classifiers using the operational mode command **show class-of-service classifier**:

```
user@switch> show class-of-service classifier
```

```
Classifier: fcoe_classifier, Code point type: ieee-802.1, Index: 10964
Code point      Forwarding class      Loss priority
011             fcoe                  low
101             fcoe1                 low
```

Meaning The **show class-of-service classifier** command shows the IEEE 802.1p code points and the loss priorities that are mapped to the forwarding classes in each classifier.

Classifier **fcoe_classifier** maps code point **011** to default lossless forwarding class **fcoe** and a packet loss priority of **low**, and maps code point **101** to explicitly configured lossless forwarding class **fcoe1** and a packet loss priority of **low**.

Verifying the PFC Flow Control Configuration (CNP)

Purpose Verify that PFC is enabled on the correct input priorities and that flow control is configured on the correct output queues and priorities.

Action List the CNPs using the operational mode command **show class-of-service congestion-notification**:

```
user@switch> show class-of-service congestion-notification
Name: fcoe_cnp, Index: 46504
Type: Input
Cable Length: 100 m
  Priority  PFC      MRU
  000      Disabled
  001      Disabled
  010      Disabled
  011      Enabled   2240
  100      Disabled
  101      Enabled   2240
  110      Disabled
  111      Disabled
Type: Output
  Priority  Flow-Control-Queues
  011      3
  101      5
```

Meaning The **show class-of-service congestion-notification** command shows the input and output stanzas of the CNP.

The CNP **fcoe_cnp** input stanza shows that PFC is enabled on code points **011** and **101**, the MRU is **2240** bytes on both priorities, and the interface cable length is **100** meters. The CNP output stanza shows that output flow control is configured on queues **3** and **5** for code points **011** and **101**, respectively.

Verifying the Interface Configuration

Purpose Verify that the classifier and congestion notification profile are configured on the interfaces. Both interfaces should show the same configuration.

Action List the ingress interfaces using the operational mode commands **show configuration class-of-service interfaces xe-0/0/20** and **show configuration class-of-service interfaces xe-0/0/21**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/20
congestion-notification-profile fcoe_cnp;
unit 0 {
  classifiers {
    ieee-802.1 fcoe_classifier;
  }
}

user@switch> show configuration class-of-service interfaces xe-0/0/21
congestion-notification-profile fcoe_cnp;
unit 0 {
  classifiers {
    ieee-802.1 fcoe_classifier;
  }
}
```


Meaning The `show configuration class-of-service interfaces xe-0/0/20` command shows that the congestion notification profile `fcoe_cnp` is configured on the interface, and that the IEEE 802.1p classifier associated with the interface is `fcoe_classifier`.

The `show configuration class-of-service interfaces xe-0/0/21` command shows that the congestion notification profile `fcoe_cnp` is configured on the interface, and that the IEEE 802.1p classifier associated with the interface is `fcoe_classifier`.

Verifying the DCBX Application Configuration

Purpose Verify that the DCBX application for FCoE is configured.

Action List the DCBX applications by using the configuration mode command `show applications`:

```
user@switch# show applications
application fcoe_app {
    ether-type 0x8906;
```

Meaning The `show applications` configuration mode command shows all of the configured applications. The output shows that the application `fcoe_app` is configured with an EtherType of `0x8906`.

Verifying the DCBX Application Map Configuration

Purpose Verify that the application map is configured.

Action List the application maps by using the configuration mode command `show policy-options application-maps`:

```
user@switch# show policy-options application-maps
fcoe_app_map {
    application fcoe_app code-points [011 101];
}
```

Meaning The `show policy-options application-maps` configuration mode command lists all of the configured application maps and the applications that belong to each application map. The output shows that application map `fcoe_app_map` consists of the application named `fcoe_app`, which is mapped to IEEE 802.1p code points `011` and `101` (priorities 3 and 5, respectively).

Verifying the DCBX Application Protocol Exchange Interface Configuration

Purpose Verify that the application map is applied to the interfaces.

Action List the application maps on each interface using the configuration mode command `show protocols dcbx`:

```
user@switch# show protocols dcbx
interface xe-0/0/20.0 {
    application-map fcoe_app_map;
}
interface xe-0/0/21.0 {
```

```
    application-map fcoe_app_map;  
}
```

Meaning The **show protocols dcbx** configuration mode command lists the application map association with interfaces. The output shows that interfaces **xe-0/0/20.0** and **xe-0/0/21.0** use application map **fcoe_app_map**.

- Related Documentation**
- [Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 6986](#)
 - [Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic \(FCoE Transit Switch\) on page 6969](#)
 - [Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications \(FCoE and iSCSI\) on page 7000](#)
 - [Example: Configuring DCBX Application Protocol TLV Exchange on page 6535](#)
 - [Configuring CoS PFC \(Congestion Notification Profiles\) on page 6905](#)
 - [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 6873](#)
 - [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 6493](#)

Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces

Although the default configuration provides two lossless forwarding classes mapped to two different IEEE 802.1p priorities (code points), you can explicitly configure up to six lossless forwarding classes and map them to different priorities. You can support up to six different types of lossless traffic, and you can support the same type of traffic if it uses different priorities in different parts of your converged network.

This example shows you how to configure two lossless forwarding classes for FCoE traffic and map them to two different priorities on an FCoE transit switch.

- [Requirements on page 6986](#)
- [Overview on page 6987](#)
- [Configuration on page 6991](#)
- [Verification on page 6994](#)

Requirements

This example uses the following hardware and software components:

- One switch used as an FCoE transit switch
- Junos OS Release 12.3 or later for the QFX Series

Overview

Some network topologies support FCoE traffic on more than one IEEE 802.1p priority. For example, when the switch acts as a transit switch, it could be connected to two QFX3500 switches in FCoE-FC gateway mode. Each of the gateway switches could connect a set of FCoE clients to a different SAN, and each set of FCoE clients could use a different priority for FCoE traffic to avoid fate sharing and maintain separation of the two FCoE networks. In this case, you need to configure two forwarding classes for FCoE traffic, each mapped to a different output queue and a different priority.

Supporting lossless behavior for two FCoE traffic classes requires configuring:

- At least one lossless forwarding class for FCoE traffic (this example uses the default **fcoe** forwarding class as one of the two lossless FCoE forwarding classes, so we need to explicitly configure only one FCoE forwarding class)
- Behavior aggregate (BA) classifiers to map the FCoE forwarding classes to the appropriate IEEE 802.1p code points (priorities) on each interface
- Congestion notification profiles (CNPs) for each interface to enable PFC on the FCoE code points at the interface ingress and to configure PFC flow control on the interface egress so that the interface can respond to PFC messages received from the connected peer



NOTE: Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

- DCBX applications and an application map to support DCBX application TLV exchange for the lossless FCoE traffic on the configured FCoE priorities. By default, DCBX is enabled on all Ethernet interfaces, but only on priority 3 (IEEE 802.1p code point 011). To support DCBX application TLV exchange when you are not using the default configuration, you must configure all of the applications and map them to interfaces and priorities.

The priorities specified in the BA classifiers, CNPs, and DCBX application map must match, or the configuration does not work. You must specify the same lossless FCoE forwarding class in each configuration and use the same IEEE 802.1p code point (priority) so that the FCoE traffic is properly classified into flows and so that those flows receive lossless treatment.

Topology

This example shows how to configure two lossless FCoE traffic classes, map them to two different priorities, and configure flow control to ensure lossless behavior for those priorities on the interfaces. This example uses three Ethernet interfaces, xe-0/0/20, xe-0/0/21, and xe-0/0/22:

- Interface xe-0/0/20 connects to an FCoE-FC gateway that connects to Fibre Channel (FC) SAN 1. FCoE traffic to and from FC SAN 1 uses the default **fcoe** forwarding class and the default mapping to priority 3 (IEEE 802.1p code point 011) and output queue 3.
- Interface xe-0/0/21 connects to another FCoE-FC gateway that connects to Fibre Channel (FC) SAN 2. FCoE traffic to and from FC SAN-2 uses an explicitly configured FCoE forwarding class that is mapped to priority 5 (code point 101) and output queue 5.
- Interface xe-0/0/22 connects to FCoE devices on the converged Ethernet network and handles traffic destined for FC SAN 1 and FC SAN 2. Interface xe-0/0/22 must properly handle lossless FCoE traffic of both priorities (both FCoE forwarding classes), including pausing the traffic on ingress or egress as required.

Figure 233 shows the topology for this example, and Table 597 shows the configuration components for this example.

Figure 233: Topology of the Two Lossless FCoE Priorities Example




Table 597: Components of the Two Lossless FCoE Priorities Configuration Topology

Component	Settings
Hardware	One switch
Forwarding classes	<p>Name—fcoe1 Queue mapping—queue 5 Packet drop attribute—no-loss</p> <p>NOTE: A lossless forwarding class can be mapped to any output queue. However, because the fcoe1 forwarding class uses priority 5 in this example, matching that traffic to a forwarding class that uses queue 5 creates a configuration that is logical and easy to map because the priority and the queue are identified by the same number.</p> <p>Name—fcoe This is the default lossless FCoE forwarding class, so no configuration required. The fcoe forwarding class is mapped to priority 3 (IEEE 802.1p code point 011) and to output queue 3 with a packet drop attribute of no-loss</p>

Table 597: Components of the Two Lossless FCoE Priorities Configuration Topology (*continued*)

Component	Settings
BA classifiers	<p>Each interface requires a different classifier because each interface handles a different subset of FCoE traffic.</p> <ul style="list-style-type: none"> Interface xe-0/0/20 classifier: Name—fcoe_p3 FCoE priority mapping—Forwarding class fcoe mapped to code point 011 (IEEE 802.1p priority 3) and a packet loss priority of low. Interface xe-0/0/21 classifier: Name—fcoe_p5 FCoE priority mapping—Forwarding class fcoe1 mapped to code point 101 (IEEE 802.1p priority 5) and a packet loss priority of low. Interface xe-0/0/22 classifier: Name—fcoe_p3_p5 FCoE priority mapping—Forwarding class fcoe1 mapped to code point 101 and a packet loss priority of low, and forwarding class fcoe mapped to code point 011 and a packet loss priority of low.

Table 597: Components of the Two Lossless FCoE Priorities Configuration Topology (*continued*)

Component	Settings
PFC configuration (CNPs)	<p>Each interface requires a different CNP because each interface handles a different subset of FCoE traffic and must pause that traffic on different priorities.</p> <ul style="list-style-type: none"> Interface xe-0/0/20 CNP: CNP name—fcoe_p3_cnp Input CNP code point—011 MRU—2240 bytes Cable length—100 meters NOTE: Because interface xe-0/0/20 uses the default FCoE configuration, output queue 3 is paused by default and you do not need to configure the output stanza of the CNP. Interface xe-0/0/21 CNP: CNP name—fcoe_p5_cnp Input CNP code point—101 MRU—2240 bytes Cable length—150 meters Output CNP code point—101 Output CNP flow control queue—5 Interface xe-0/0/22 CNP: CNP name—fcoe_p3_p5_cnp Input CNP code points—011 and 101 MRU—2240 bytes (both priorities) Cable length—100 meters Output CNP code points—011 (for queue 3) and 101 (for queue 5) Output CNP flow control queues—3 for priority 3 (code point 011) and 5 for priority 5 (code point 101) <p>NOTE: When you apply a CNP with an explicit output queue flow control configuration to an interface, the explicit CNP overwrites the default output CNP. The output queues that are enabled for pause in the default configuration (queues 3 and 4) are not enabled for pause unless they are included in the explicitly configured output CNP.</p>

Table 597: Components of the Two Lossless FCoE Priorities Configuration Topology (continued)

Component	Settings
DCBX application mapping	<p>Interface xe-0/0/20 does not need an application map because DCBX exchanges application protocol TLVs only on the default FCoE priority (priority 3).</p> <p>Interface xe-0/0/21 requires an application map that enables DCBX application protocol TLV exchange on priority 5 (code point 101) for FCoE traffic. Interface xe-0/0/22 requires an application map that enables DCBX application protocol TLV exchange both on priority 3 (code point 011) and on priority 5 (code point 101) for FCoE traffic.</p> <ul style="list-style-type: none"> Interface xe-0/0/21 DCBX application mapping: <ul style="list-style-type: none"> Application name—fcoe_p5_app Application ether-type—0x8906 Application map name—fcoe_p5_app_map Application map code points—101 Interface xe-0/0/22 DCBX application mapping: <ul style="list-style-type: none"> Application name—fcoe_all_app Application ether-type—0x8906 Application map name—fcoe_all_app_map Application map code points—011 and 101 <p>NOTE: LLDP and DCBX must be enabled on the interface. By default, LLDP and DCBX are enabled on all Ethernet interfaces.</p>



NOTE: This example does not include scheduling (bandwidth allocation) configuration or the FIP snooping configuration. This examples focuses only on the lossless FCoE priority configuration.

QFX10000 switches do not support FIP snooping. For this reason, QFX10000 switches cannot be used as FCoE access transit switches. QFX10000 switches can be used as intermediate or aggregation transit switches in the FCoE path, between an FCoE access transit switch that performs FIP snooping and an FCF.

Configuration

CLI Quick Configuration

To quickly configure two lossless FCoE forwarding classes that use different priorities on an FCoE transit switch, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set class-of-service forwarding-classes class fcoe1 queue-num 5 no-loss
set class-of-service classifiers ieee-802.1 fcoe_p3 forwarding-class fcoe loss-priority low
code-points 011
set class-of-service classifiers ieee-802.1 fcoe_p5 forwarding-class fcoe1 loss-priority low
code-points 101
set class-of-service classifiers ieee-802.1 fcoe_p3_p5 forwarding-class fcoe loss-priority low
code-points 011
```

```

set class-of-service classifiers ieee-802.1p fcoe_p3_p5 forwarding-class fcoe1 loss-priority low
code-points 101
set class-of-service interfaces xe-0/0/20 unit 0 classifiers ieee-802.1p fcoe_p3
set class-of-service interfaces xe-0/0/21 unit 0 classifiers ieee-802.1p fcoe_p5
set class-of-service interfaces xe-0/0/22 unit 0 classifiers ieee-802.1p fcoe_p3_p5
set class-of-service congestion-notification-profile fcoe_p3_cnp input ieee-802.1p code-point 011
pfc mru 2240
set class-of-service congestion-notification-profile fcoe_p3_cnp input cable-length 100
set class-of-service congestion-notification-profile fcoe_p5_cnp input ieee-802.1p code-point 101
pfc mru 2240
set class-of-service congestion-notification-profile fcoe_p5_cnp input cable-length 150
set class-of-service congestion-notification-profile fcoe_p5_cnp output ieee-802.1p code-point
101 pfc flow-control-queue 5
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp input ieee-802.1p code-point
011 pfc mru 2240
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp input ieee-802.1p code-point
101 pfc mru 2240
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp input cable-length 100
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp output ieee-802.1p code-point
011 pfc flow-control-queue 3
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp output ieee-802.1p code-point
101 pfc flow-control-queue 5
set class-of-service interfaces xe-0/0/20 congestion-notification-profile fcoe_p3_cnp
set class-of-service interfaces xe-0/0/21 congestion-notification-profile fcoe_p5_cnp
set class-of-service interfaces xe-0/0/22 congestion-notification-profile fcoe_p3_p5_cnp
set applications application fcoe_p5_app ether-type 0x8906
set applications application fcoe_all_app ether-type 0x8906
set policy-options application-maps fcoe_p5_app_map application fcoe_p5_app code-points 101
set policy-options application-maps fcoe_all_app_map application fcoe_all_app code-points [011
101]
set protocols dcbx interface xe-0/0/21 application-map fcoe_p5_app_map
set protocols dcbx interface xe-0/0/22 application-map fcoe_all_app_map

```

Step-by-Step Procedure

To configure two lossless forwarding classes for FCoE traffic on different interfaces, classify FCoE traffic into the forwarding classes, configure congestion notification profiles to enable PFC on the FCoE priorities and output queues, and configure DCBX application protocol TLV exchange for traffic on both FCoE priorities:

1. Configure lossless forwarding class **fcoe1** and map it to output queue **5** for FCoE traffic that uses IEEE 802.1p priority **5**:

```

[edit class-of-service]
user@switch# set forwarding-classes class fcoe1 queue-num 5 no-loss

```



NOTE: This example uses the default **fcoe** forwarding class as the other lossless FCoE forwarding class.

2. Configure the ingress classifier (**fcoe_p3**) for interface **xe-0/0/20**. The classifier maps the FCoE priority (IEEE 802.1p code point **011**) to lossless FCoE forwarding class **fcoe**:

```

[edit class-of-service classifiers]
user@switch# set ieee-802.1p fcoe_p3 forwarding-class fcoe loss-priority low code-points
011

```


3. Configure the ingress classifier (**fcoe_p5**) for interface **xe-0/0/21**. The classifier maps the FCoE priority (IEEE 802.1p code point **101**) to lossless FCoE forwarding class **fcoe1**:

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 fcoe_p5 forwarding-class fcoe1 loss-priority low code-points
101
```

4. Configure the ingress classifier (**fcoe_p3_p5**) for interface **xe-0/0/22**. The classifier maps the two FCoE priorities (IEEE 802.1p code points **011** and **101**) to the two lossless FCoE forwarding classes **fcoe** and **fcoe1**, respectively:

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 fcoe_p3_p5 forwarding-class fcoe loss-priority low code-points
011
user@switch# set ieee-802.1 fcoe_p3_p5 forwarding-class fcoe1 loss-priority low code-points
101
```

5. Apply each classifier to the appropriate interface:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 unit 0 classifiers ieee-802.1 fcoe_p3
user@switch# set interfaces xe-0/0/21 unit 0 classifiers ieee-802.1 fcoe_p5
user@switch# set interfaces xe-0/0/22 unit 0 classifiers ieee-802.1 fcoe_p3_p5
```

6. Configure the CNP input stanza for interface **xe-0/0/20** to enable PFC on the FCoE priority (IEEE 802.1p code point **011**), set the MRU value (2240 bytes), and set the cable length value (100 meters). No output stanza is needed because queue 3 is paused by default on priority 3, and we are not explicitly configuring output queue flow control for any other queues.

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe_p3_cnp input ieee-802.1 code-point
011 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_p3_cnp input cable-length 100
```

7. Configure the CNP for interface **xe-0/0/21**. The input stanza enables PFC on the FCoE priority (IEEE 802.1p code point **101**), sets the MRU value (2240 bytes), and sets the cable length value (150 meters). The output stanza configures flow control on output queue 5 on the FCoE priority:

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe_p5_cnp input ieee-802.1 code-point
101 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_p5_cnp input cable-length 150
user@switch# set congestion-notification-profile fcoe_p5_cnp output ieee-802.1 code-point
101 pfc flow-control-queue 5
```

8. Configure the CNP for interface **xe-0/0/22**. The input stanza enables PFC on the FCoE priorities (IEEE 802.1p code points **011** and **101**), sets the MRU value (2240 bytes), and sets the cable length value (100 meters). The output stanza configures flow control on output queues 3 and 5 on the FCoE priorities:

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe_p3_p5_cnp input ieee-802.1
code-point 011 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_p3_p5_cnp input ieee-802.1
code-point 101 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_p3_p5_cnp input cable-length 100
```

```

user@switch# set congestion-notification-profile fcoe_p3_p5_cnp output ieee-802.1
code-point 011 pfc flow-control-queue 3
user@switch# set congestion-notification-profile fcoe_p3_p5_cnp output ieee-802.1
code-point 101 pfc flow-control-queue 5

```

9. Apply each CNP to the appropriate interface:

```

[edit class-of-service]
user@switch# set interfaces xe-0/0/20 congestion-notification-profile fcoe_p3_cnp
user@switch# set interfaces xe-0/0/21 congestion-notification-profile fcoe_p5_cnp
user@switch# set interfaces xe-0/0/22 congestion-notification-profile fcoe_p3_p5_cnp

```

10. Configure the DCBX FCoE application and application map to apply to interface xe-0/0/21. Interface xe-0/0/21 uses priority 5 (IEEE 802.1p code point 101) for FCoE traffic, which requires DCBX to exchange FCoE application protocol TLVs on priority 5 on interface xe-0/0/21. Configure an application named **fcoe_p5_app** for FCoE traffic (EtherType 0x8906) and configure an application map named **fcoe_p5_app_map** to map the application to code point 101:

```

[edit]
user@switch# set applications application fcoe_p5_app ether-type 0x8906
user@switch# set policy-options application-maps fcoe_p5_app_map application
fcoe_p5_app code-points 101

```



NOTE: Interface xe-0/0/20 uses the default FCoE configuration (priority 3). DCBX exchanges protocol TLVs for the FCoE application by default, so you do not need to configure DCBX explicitly on interface xe-0/0/20.

11. Configure the DCBX FCoE application and application map to apply to interface xe-0/0/22. Interface xe-0/0/22 uses both priority 3 (IEEE 802.1p code point 011) and priority 5 for FCoE traffic, which requires DCBX to exchange FCoE application protocol TLVs on both priority 3 and priority 5. Configure an application named **fcoe_all_app** for FCoE traffic (EtherType 0x8906) and configure an application map named **fcoe_all_app_map** to map the application to code points 011 and 101:

```

[edit]
user@switch# set applications application fcoe_all_app ether-type 0x8906
user@switch# set policy-options application-maps fcoe_all_app_map application
fcoe_all_app code-points [011 101]

```

12. Apply the application maps to the interfaces xe-0/0/21 and xe-0/0/22 so that DCBX exchanges FCoE application TLVs on the correct code points on each interface:

```

[edit]
user@switch# set protocols dcbx interface xe-0/0/21 application-map fcoe_p5_app_map
user@switch# set protocols dcbx interface xe-0/0/22 application-map fcoe_all_app_map

```

Verification

To verify the configuration and proper operation of the lossless forwarding classes and IEEE 802.1p priorities, perform these tasks:

- [Verifying the Forwarding Class Configuration on page 6995](#)
- [Verifying the Behavior Aggregate Classifier Configuration on page 6995](#)

- [Verifying the PFC Flow Control Configuration \(CNP\) on page 6996](#)
- [Verifying the Interface Configuration on page 6998](#)
- [Verifying the DCBX Application Configuration on page 6999](#)
- [Verifying the DCBX Application Map Configuration on page 6999](#)
- [Verifying the DCBX Application Protocol Exchange Interface Configuration on page 6999](#)

Verifying the Forwarding Class Configuration

Purpose Verify that the lossless forwarding class **fcoe1** has been created.

Action Show the forwarding class configuration by using the operational command **show class-of-service forwarding class**:

```
user@switch# show class-of-service forwarding-class
```

Forwarding class	ID	Queue	Policing priority	No-Loss
best-effort	0	0	normal	Disabled
fcoe	1	3	normal	Enabled
no-loss	2	4	normal	Enabled
network-control	3	7	normal	Disabled
fcoe1	4	5	normal	Enabled
mcast	8	8	normal	Disabled

Meaning The **show class-of-service forwarding-class** command shows all of the forwarding classes. The command output shows that the **fcoe1** forwarding class is configured on output queue **5** with the no-loss packet drop attribute enabled.

Because we did not explicitly configure the default forwarding classes, they remain in their default state, including the lossless configuration of the **fcoe** and **no-loss** default forwarding classes.

Verifying the Behavior Aggregate Classifier Configuration

Purpose Verify that the three classifiers map the forwarding classes to the correct IEEE 802.1p code points (priorities) and packet loss priorities.

Action List the classifiers configured to support lossless FCoE transport using the operational mode command **show class-of-service classifier**:

```
user@switch> show class-of-service classifier
```

Classifier: fcoe_p3, Code point type: ieee-802.1, Index: 13913		
Code point	Forwarding class	Loss priority
011	fcoe	low
Classifier: fcoe_p5, Code point type: ieee-802.1, Index: 63065		
Code point	Forwarding class	Loss priority
101	fcoe1	low

```
Classifier: fcoe_p3_p5, Code point type: ieee-802.1, Index: 10964
Code point      Forwarding class      Loss priority
011             fcoe                  low
101             fcoe1                 low
```

Meaning The **show class-of-service classifier** command shows the IEEE 802.1p code points and the loss priorities that are mapped to the forwarding classes in each classifier. The command output shows that there are three classifiers, **fcoe_p3**, **fcoe_p5**, and **fcoe_p3_p5**.

Classifier **fcoe_p3** maps code point **011** (priority 3) to default lossless forwarding class **fcoe** and a packet loss priority of **low**.

Classifier **fcoe_p5** maps code point **101** (priority 5) to explicitly configured lossless forwarding class **fcoe1** and a packet loss priority of **low**.

Classifier **fcoe_p3_p5** maps code point **011** to default lossless forwarding class **fcoe** and a packet loss priority of **low**, and maps code point **101** to explicitly configured lossless forwarding class **fcoe1** and a packet loss priority of **low**.

Verifying the PFC Flow Control Configuration (CNP)

Purpose Verify that PFC is enabled on the correct input priorities and that flow control is configured on the correct output queues and priorities in each CNP.

Action List the congestion notification profiles using the operational mode command **show class-of-service congestion-notification**:

```
user@switch> show class-of-service congestion-notification
```

```
Name: fcoe_p3_cnp, Index: 12037
```

```
Type: Input
```

```
Cable Length: 100 m
```

Priority	PFC	MRU
000	Disabled	
001	Disabled	
010	Disabled	
011	Enabled	2240
100	Disabled	
101	Disabled	
110	Disabled	
111	Disabled	

```
Type: Output
```

Priority	Flow-Control-Queues
000	
	0
001	
	1
010	
	2
011	
	3
100	
	4
101	
	5
110	
	6
111	

7

Name: fcoe_p3_p5_cnp, Index: 46484

Type: Input

Cable Length: 100 m

Priority	PFC	MRU
000	Disabled	
001	Disabled	
010	Disabled	
011	Enabled	2240
100	Disabled	
101	Enabled	2240
110	Disabled	
111	Disabled	

Type: Output

Priority	Flow-Control-Queues
011	
	3
101	
	5

Name: fcoe_p5_cnp, Index: 12133

Type: Input

Cable Length: 150 m

Priority	PFC	MRU
000	Disabled	
001	Disabled	
010	Disabled	
011	Disabled	
100	Disabled	
101	Enabled	2240
110	Disabled	
111	Disabled	

Type: Output

Priority	Flow-Control-Queues
101	
	5

Meaning The **show class-of-service congestion-notification** command shows the input and output stanzas of the three CNPs. For CNP **fcoe_p3_cnp**, the input stanza shows that PFC is enabled on IEEE 802.1p code point **011** (priority 3), the MRU is **2240** bytes, and the cable length is **100** meters. The CNP output stanza shows the default mapping of priorities to output queues.



NOTE: By default, only queues 3 and 4 are enabled to respond to pause messages from the connected peer. For queue 3 to respond to pause messages, priority 3 (code point 011) must be enabled for PFC in the input stanza. For queue 4 to respond to pause messages, priority 4 (code point 100) must be enabled for PFC in the input stanza. In this example, only queue 3 responds to pause messages from the connected peer on interfaces that use CNP **fcoe_p3_cnp**, because the input stanza enables PFC priority 3 only.

For CNP **fcoe_p3_p5_cnp**, the input stanza shows that PFC is enabled on code points **011** and **101**, the MRU is **2240** bytes on both priorities, and the cable length is **100** meters. The

CNP output stanza shows that output flow control is configured on queues **3** and **5** for code points **011** and **101**, respectively.

For CNP **fcoe_p5_cnp**, the input stanza shows that PFC is enabled on code point **101** (priority 5), the MRU is **2240** bytes, and the cable length is **150** meters. The CNP output stanza shows that output flow control is configured on queue **5** for code point **101** (priority 5).

Verifying the Interface Configuration

Purpose Verify that the correct classifiers and congestion notification profiles are configured on the correct interfaces.

Action List the ingress interfaces using the operational mode commands **show configuration class-of-service interfaces xe-0/0/20**, **show configuration class-of-service interfaces xe-0/0/21**, and **show configuration class-of-service interfaces xe-0/0/22**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/20
congestion-notification-profile fcoe_p3_cnp;
unit 0 {
    classifiers {
        ieee-802.1p fcoe_p3;
    }
}
```

```
user@switch> show configuration class-of-service interfaces xe-0/0/21
congestion-notification-profile fcoe_p5_cnp;
unit 0 {
    classifiers {
        ieee-802.1p fcoe_p5;
    }
}
```

```
user@switch> show configuration class-of-service interfaces xe-0/0/22
congestion-notification-profile fcoe_p3_p5_cnp;
unit 0 {
    classifiers {
        ieee-802.1p fcoe_p3_p5;
    }
}
```

Meaning The **show configuration class-of-service interfaces xe-0/0/20** command shows that the congestion notification profile **fcoe_p3_cnp** is configured on the interface, and that the IEEE 802.1p classifier associated with the interface is **fcoe_p3**.

The **show configuration class-of-service interfaces xe-0/0/21** command shows that the congestion notification profile **fcoe_p5_cnp** is configured on the interface, and that the IEEE 802.1p classifier associated with the interface is **fcoe_p5**.

The **show configuration class-of-service interfaces xe-0/0/22** command shows that the congestion notification profile **fcoe_p3_p5_cnp** is configured on the interface, and that the IEEE 802.1p classifier associated with the interface is **fcoe_p3_p5**.

Verifying the DCBX Application Configuration

- Purpose** Verify that the two DCBX applications for FCoE are configured.
- Action** List the DCBX applications by using the configuration mode command **show applications**:
- ```
user@switch# show applications
application fcoe_all_app {
 ether-type 0x8906;

application fcoe_p5_app {
 ether-type 0x8906;
```
- Meaning** The **show applications** configuration mode command shows all of the configured applications. The output shows that the application **fcoe\_all\_app** is configured with an EtherType of **0x8906** (the correct EtherType for FCoE traffic) and that the application **fcoe\_p5\_app** is also configured with an EtherType of **0x8906**.

### Verifying the DCBX Application Map Configuration

- Purpose** Verify that the application maps are configured.
- Action** List the application maps by using the configuration mode command **show policy-options application-maps**:
- ```
user@switch# show policy-options application-maps
fcoe_all_app_map {
    application fcoe_all_app code-points [011 101];
}
fcoe_p5_app_map {
    application fcoe_p5_app code-points 101;
}
```
- Meaning** The **show policy-options application-maps** configuration mode command lists all of the configured application maps and the applications that belong to each application map. The output shows that there are two application maps.
- Application map **fcoe_all_app_map** consists of the application named **fcoe_all_app** mapped to IEEE 802.1p code points **011** (priority 3) and **101** (priority 5).
- Application map **fcoe_p5_app_map** consists of the application named **fcoe_p5_app** mapped to IEEE 802.1p code point **101** (priority 5).

Verifying the DCBX Application Protocol Exchange Interface Configuration

- Purpose** Verify that the application maps are applied to the correct interfaces.
- Action** List the application maps on each interface using the configuration mode command **show protocols dcbx**:
- ```
user@switch# show protocols dcbx
interface xe-0/0/21.0 {
 application-map fcoe_p5_app_map;
```

```
}
interface xe-0/0/22.0 {
 application-map fcoe_all_app_map;
}
```

**Meaning** The `show protocols dcbx` configuration mode command lists the application map association with interfaces. The output shows that interface `xe-0/0/21.0` uses application map `fcoe_p5_app_map` and interface `xe-0/0/22.0` uses application map `fcoe_all_app_map`.



**NOTE:** Because interface `xe-0/0/20` uses the default lossless FCoE configuration, you do not configure application mapping to interface `xe-0/0/20`. The default configuration automatically exchanges application protocol TLVs for the default FCoE configuration on priority 3 (IEEE 802.1p code point 011).

---

**Related Documentation**

- [Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 6977](#)
- [Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic \(FCoE Transit Switch\) on page 6969](#)
- [Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications \(FCoE and iSCSI\) on page 7000](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 6535](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 6905](#)
- [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 6873](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 6493](#)

## Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI)

---

Although the default configuration provides two lossless forwarding classes mapped to two different IEEE 802.1p priorities (code points), you can explicitly configure up to six lossless forwarding classes and map them to different priorities. You can support up to six different types of lossless traffic, and you can support the same type of traffic on different priorities in different parts of your converged network.

This example shows you how to configure two lossless forwarding classes for FCoE traffic and one lossless forwarding class for iSCSI traffic, and map the forwarding classes to three different priorities. (The converged Ethernet network includes two FCoE networks, each of which uses a different priority to identify FCoE traffic, and an iSCSI network.)

- [Requirements on page 7001](#)
- [Overview on page 7001](#)



- [Configuration on page 7005](#)
- [Verification on page 7009](#)

## Requirements

This example uses the following hardware and software components:

- One switch used as an FCoE transit switch
- Junos OS Release 12.3 or later for the QFX Series

## Overview

Some converged Ethernet networks support FCoE on more than one IEEE 802.1p priority and also require supporting other lossless traffic classes. Interfaces that carry multiple lossless forwarding classes need to support lossless behavior for the priorities mapped to those forwarding classes. To support the two FCoE forwarding classes and the iSCSI forwarding class used in this example, you need to configure:

- At least one lossless forwarding class for FCoE traffic (this example uses the default **fcoe** forwarding class as one of the two lossless FCoE forwarding classes, so we need to explicitly configure only one FCoE forwarding class)
- A lossless forwarding class for iSCSI traffic
- Behavior aggregate (BA) classifiers to map the lossless forwarding classes to the appropriate IEEE 802.1p code points (priorities) on each interface
- Congestion notification profiles (CNPs) for each interface to enable PFC on the FCoE and iSCSI code points at the interface ingress, and to configure PFC flow control on the interface egress so that the interface can respond to PFC messages received from the connected peer



**NOTE:** Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

- DCBX applications and an application map to support DCBX application TLV exchange for the FCoE and iSCSI traffic on the configured lossless priorities. By default, DCBX is enabled on all Ethernet interfaces for FCoE, but only on priority 3 (IEEE 802.1p code point 011). To support DCBX application TLV exchange when you are not using the default configuration, you must configure all of the applications and map them to interfaces and priorities.

The priorities specified in the BA classifiers, CNPs, and DCBX application map must match, or the configuration does not work. You must specify the same lossless FCoE forwarding class in each configuration and use the same IEEE 802.1p code point (priority) so that the FCoE traffic is properly classified into flows and so that those flows receive lossless treatment.

Topology

This example shows how to configure two lossless FCoE traffic classes and one lossless iSCSI traffic class, map them to three different priorities, and configure flow control to ensure lossless behavior for those priorities on the interfaces. This example uses four Ethernet interfaces, xe-0/0/31, xe-0/0/32, xe-0/0/33, and xe-0/0/34:

- Interface xe-0/0/31 handles FCoE traffic on priority 3 (IEEE 802.1p code point 011) and iSCSI traffic on priority 4 (code point 100).
- Interface xe-0/0/32 handles FCoE traffic on priority 5 (code point 101) and iSCSI traffic on priority 4.
- Interface xe-0/0/33 handles FCoE traffic on priority 3 and priority 5.
- Interface xe-0/0/34 handles iSCSI traffic on priority 4.

Figure 234 shows the topology for this example, and Table 598 shows the configuration components for this example.

Figure 234: Topology of the Lossless FCoE and iSCSI Priorities Example




Table 598: Components of the Lossless FCoE and iSCSI Priorities Configuration Topology

| Component | Settings   |
|-----------|------------|
| Hardware  | One switch |

**Table 598: Components of the Lossless FCoE and iSCSI Priorities Configuration Topology (*continued*)**

| Component          | Settings                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Forwarding classes | <p>This example uses one explicitly configured lossless FCoE forwarding class, the default lossless FCoE forwarding class, and one explicitly configured iSCSI forwarding class.</p> <ul style="list-style-type: none"> <li>iSCSI forwarding class:<br/>Name—<b>iscsi</b><br/>Queue mapping—queue 4<br/>Packet drop attribute—<b>no-loss</b></li> <li>FCoE forwarding class (explicitly configured):<br/>Name—<b>fcoe1</b><br/>Queue mapping—queue 5<br/>Packet drop attribute—<b>no-loss</b></li> </ul> <p><b>NOTE:</b> A lossless forwarding class can be mapped to any output queue. However, because the <b>fcoe1</b> forwarding class uses priority 5 in this example, matching that traffic to a forwarding class that uses queue 5 creates a configuration that is logical and easy to map because the priority and the queue are identified by the same number.</p> <ul style="list-style-type: none"> <li>FCoE forwarding class (default)<br/>Name—<b>fcoe</b><br/>The default <b>fcoe</b> forwarding class is mapped to priority 3 (IEEE 802.1p code point 011) and to output queue 3 with a packet drop attribute of <b>no-loss</b>.</li> </ul>                                                                                                                                                                                                                                                                                                                                   |
| BA classifiers     | <p>Each interface requires a different classifier because each interface handles a different subset of FCoE traffic.</p> <ul style="list-style-type: none"> <li>Interface xe-0/0/31 classifier:<br/>Name—<b>fcoe_p3_iscsi</b><br/>FCoE priority mapping—Forwarding class <b>fcoe</b> mapped to code point <b>011</b> (IEEE 802.1p priority 3) and a packet loss priority of <b>low</b>.<br/>iSCSI priority mapping—Forwarding class <b>iscsi</b> mapped to code point <b>100</b> (priority 4) and a packet loss priority of <b>low</b>.</li> <li>Interface xe-0/0/32 classifier:<br/>Name—<b>fcoe_p5_iscsi</b><br/>FCoE priority mapping—Forwarding class <b>fcoe1</b> mapped to code point <b>101</b> (IEEE 802.1p priority 5) and a packet loss priority of <b>low</b>.<br/>iSCSI priority mapping—Forwarding class <b>iscsi</b> mapped to code point <b>100</b> (priority 4) and a packet loss priority of <b>low</b>.</li> <li>Interface xe-0/0/33 classifier:<br/>Name—<b>fcoe_p3_p5</b><br/>FCoE priority mapping—Forwarding class <b>fcoe1</b> mapped to code point <b>101</b> (priority 5) and a packet loss priority of <b>low</b>, and forwarding class <b>fcoe</b> mapped to code point <b>011</b> and a packet loss priority of <b>low</b>.</li> <li>Interface xe-0/0/34 classifier:<br/>Name—<b>iscsi_classifier</b><br/>iSCSI priority mapping—Forwarding class <b>iscsi</b> mapped to code point <b>100</b> (priority 4) and a packet loss priority of <b>low</b>.</li> </ul> |

**Table 598: Components of the Lossless FCoE and iSCSI Priorities Configuration Topology (*continued*)**

| Component                | Settings                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PFC configuration (CNPs) | <p>Each interface requires a different CNP because each interface handles a different subset of FCoE and iSCSI traffic, and must pause that traffic on different priorities.</p> <ul style="list-style-type: none"> <li>Interface xe-0/0/31 CNP:<br/> CNP name—<b>fcoe_p3_cnp</b><br/> Input CNP code points—<b>011</b> and <b>100</b><br/> MRU—2240 bytes for code point <b>011</b>, default value (2500 bytes) for code point <b>100</b><br/> Cable length—100 meters<br/> <br/> <b>NOTE:</b> On interface xe-0/0/31, the FCoE forwarding class is mapped to queue 3 and priority 3 (code point 011), and the iSCSI forwarding class is mapped to queue 4 and priority 4 (code point 100). Therefore, interface xe-0/0/31 does not require an output CNP configuration because queue 3 and queue 4 are enabled for PFC flow control by default on code points 011 and 100, respectively.</li> <li>Interface xe-0/0/32 CNP:<br/> CNP name—<b>fcoe_p5_cnp</b><br/> Input CNP code points—<b>100</b> and <b>101</b><br/> MRU—Default value (2500 bytes) for code point <b>100</b>, <b>2240</b> bytes for code point <b>101</b><br/> Cable length—<b>150</b> meters<br/> Output CNP code points—<b>100</b> and <b>101</b><br/> Output CNP flow control queues—<b>4</b> and <b>5</b></li> <li>Interface xe-0/0/33 CNP:<br/> CNP name—<b>fcoe_p3_p5_cnp</b><br/> Input CNP code points—<b>011</b> and <b>101</b><br/> MRU—<b>2240</b> bytes (both priorities)<br/> Cable length—<b>100</b> meters<br/> Output CNP code points—<b>011</b> and <b>101</b><br/> Output CNP flow control queues—<b>3</b> and <b>5</b></li> <li>Interface xe-0/0/34 CNP:<br/> CNP name—<b>iscsi_cnp</b><br/> Input CNP code point—<b>100</b><br/> MRU—<b>2500</b> bytes (default value)<br/> Cable length—<b>100</b> meters<br/> <br/> <b>NOTE:</b> On interface xe-0/0/34, the iSCSI forwarding class is mapped to queue 4 and priority 4 (code point 100). Interface xe-0/0/34 does not require an output CNP configuration because queue 4 is enabled for PFC flow control by default on code point 100.</li> </ul> <p><b>NOTE:</b> When you apply a CNP with an explicit output queue flow control configuration to an interface, the explicit CNP overwrites the default output CNP. The output queues that are enabled for PFC pause in the default configuration (queues 3 and 4) are not enabled for pause unless they are included in the explicitly configured output CNP.</p> |

**Table 598: Components of the Lossless FCoE and iSCSI Priorities Configuration Topology (continued)**

| Component                | Settings                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DCBX application mapping | <p>This example requires configuring applications for FCoE and iSCSI, including them in the same application map, and applying the application map to all four interfaces.</p> <p>Application map name—<b>dcbx_iscsi_fcoe_app_map</b></p> <ul style="list-style-type: none"> <li>FCoE application name—<b>fcoe_app</b><br/>Application ether-type—<b>0x8906</b><br/>Application map code points—<b>011</b> and <b>101</b></li> <li>iSCSI application name—<b>iscsi_app</b><br/>Application protocol type—<b>tcp</b><br/>Application destination port—<b>3260</b><br/>Application map code point—<b>100</b></li> </ul> <p><b>NOTE:</b> LLDP and DCBX must be enabled on the interface. By default, LLDP and DCBX are enabled on all Ethernet interfaces.</p> |



**NOTE:** This example does not include scheduling (bandwidth allocation) configuration or the FIP snooping configuration. This examples focuses only on the lossless FCoE priority configuration.

QFX10000 switches do not support FIP snooping. For this reason, QFX10000 switches cannot be used as FCoE access transit switches. QFX10000 switches can be used as intermediate or aggregation transit switches in the FCoE path, between an FCoE access transit switch that performs FIP snooping and an FCF.

## Configuration

### CLI Quick Configuration

To quickly configure two lossless FCoE forwarding classes and one lossless iSCSI forwarding class and map them to different priorities, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set class-of-service forwarding-classes class iscsi queue-num 4 no-loss
set class-of-service forwarding-classes class fcoe1 queue-num 5 no-loss
set class-of-service classifiers ieee-802.1 fcoe_p3_iscsi forwarding-class fcoe loss-priority low
code-points 011
set class-of-service classifiers ieee-802.1 fcoe_p3_iscsi forwarding-class iscsi loss-priority low
code-points 100
set class-of-service classifiers ieee-802.1 fcoe_p5_iscsi forwarding-class iscsi loss-priority low
code-points 100
set class-of-service classifiers ieee-802.1 fcoe_p5_iscsi forwarding-class fcoe1 loss-priority low
code-points 101
set class-of-service classifiers ieee-802.1 fcoe_p3_p5 forwarding-class fcoe loss-priority low
code-points 011
set class-of-service classifiers ieee-802.1 fcoe_p3_p5 forwarding-class fcoe1 loss-priority low
code-points 101
```

```

set class-of-service classifiers ieee-802.1 iscsi_classifier forwarding-class iscsi loss-priority low
code-points 100
set class-of-service interfaces xe-0/0/31 unit 0 classifiers ieee-802.1 fcoe_p3_iscsi
set class-of-service interfaces xe-0/0/32 unit 0 classifiers ieee-802.1 fcoe_p5_iscsi
set class-of-service interfaces xe-0/0/33 unit 0 classifiers ieee-802.1 fcoe_p3_p5set
class-of-service interfaces xe-0/0/34 unit 0 classifiers ieee-802.1 iscsi_classifier
set class-of-service congestion-notification-profile fcoe_p3_cnp input ieee-802.1 code-point 011
pfc mru 2240
set class-of-service congestion-notification-profile fcoe_p3_cnp input ieee-802.1 code-point 100
pfc
set class-of-service congestion-notification-profile fcoe_p3_cnp input cable-length 100
set class-of-service congestion-notification-profile fcoe_p5_cnp input ieee-802.1 code-point 100
pfc
set class-of-service congestion-notification-profile fcoe_p5_cnp input ieee-802.1 code-point 101
pfc mru 2240
set class-of-service congestion-notification-profile fcoe_p5_cnp input cable-length 150
set class-of-service congestion-notification-profile fcoe_p5_cnp output ieee-802.1 code-point
100 pfc flow-control-queue 4
set class-of-service congestion-notification-profile fcoe_p5_cnp output ieee-802.1 code-point
101 pfc flow-control-queue 5
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp input ieee-802.1 code-point
011 pfc mru 2240
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp input ieee-802.1 code-point
101 pfc mru 2240
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp input cable-length 100
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp output ieee-802.1 code-point
011 pfc flow-control-queue 3
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp output ieee-802.1 code-point
101 pfc flow-control-queue 5
set class-of-service congestion-notification-profile iscsi_cnp input ieee-802.1 code-point 100 pfc
set class-of-service congestion-notification-profile iscsi_cnp input cable-length 100
set class-of-service interfaces xe-0/0/31 congestion-notification-profile fcoe_p3_cnp
set class-of-service interfaces xe-0/0/32 congestion-notification-profile fcoe_p5_cnp
set class-of-service interfaces xe-0/0/33 congestion-notification-profile fcoe_p3_p5_cnp
set class-of-service interfaces xe-0/0/34 congestion-notification-profile iscsi_cnp
set applications application iscsi_app protocol tcp destination-port 3260
set applications application fcoe_app ether-type 0x8906
set policy-options application-maps dcbx_iscsi_fcoe_app_map application iscsi_app code-points
100
set policy-options application-maps dcbx_iscsi_fcoe_app_map application fcoe_app code-points
[011 101]
set protocols dcbx interface xe-0/0/31 application-map dcbx_iscsi_fcoe_app_map
set protocols dcbx interface xe-0/0/32 application-map dcbx_iscsi_fcoe_app_map
set protocols dcbx interface xe-0/0/33 application-map dcbx_iscsi_fcoe_app_map
set protocols dcbx interface xe-0/0/34 application-map dcbx_iscsi_fcoe_app_map

```

#### Step-by-Step Procedure

To configure two lossless forwarding classes for FCoE traffic and one lossless forwarding class for iSCSI traffic, classify the traffic into the three forwarding classes, configure congestion notification profiles to enable PFC on the FCoE priorities and output queues, and configure DCBX application protocol TLV exchange for traffic on both FCoE priorities:

1. Configure lossless forwarding classes **iscsi** for iSCSI traffic and **fcoe1** for FCoE traffic (this example uses the default **fcoe** forwarding class as the other lossless FCoE forwarding class) and map them to output queues:

```

[edit class-of-service]
user@switch# set forwarding-classes class iscsi queue-num 4 no-loss
user@switch# set forwarding-classes class fcoe1 queue-num 5 no-loss

```

2. Configure the ingress classifier (**fcoe\_p3\_iscsi**) for interface **xe-0/0/31**. The classifier maps the FCoE priority (code point **011**) to lossless FCoE forwarding class **fcoe** and the iSCSI priority (code point **100**) to lossless iSCSI forwarding class **iscsi**:

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 fcoe_p3_iscsi forwarding-class fcoe loss-priority low
code-points 011
user@switch# set ieee-802.1 fcoe_p3_iscsi forwarding-class iscsi loss-priority low
code-points 100
```

3. Configure the ingress classifier (**fcoe\_p5\_iscsi**) for interface **xe-0/0/32**. The classifier maps the FCoE priority (code point **101**) to lossless FCoE forwarding class **fcoe1** and the iSCSI priority (code point **100**) to lossless iSCSI forwarding class **iscsi**:

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 fcoe_p5_iscsi forwarding-class iscsi loss-priority low
code-points 100
user@switch# set ieee-802.1 fcoe_p5_iscsi forwarding-class fcoe1 loss-priority low
code-points 101
```

4. Configure the ingress classifier (**fcoe\_p3\_p5**) for interface **xe-0/0/33**. The classifier maps the two FCoE priorities (code points **011** and **101**) to lossless FCoE forwarding classes **fcoe** and **fcoe1**, respectively:

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 fcoe_p3_p5 forwarding-class fcoe loss-priority low code-points
011
user@switch# set ieee-802.1 fcoe_p3_p5 forwarding-class fcoe1 loss-priority low code-points
101
```

5. Configure the ingress classifier (**iscsi\_classifier**) for interface **xe-0/0/34**. The classifier maps the iSCSI priority (code point **101**) to lossless iSCSI forwarding class **iscsi**:

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 iscsi_classifier forwarding-class iscsi loss-priority low
code-points 100
```

6. Apply each classifier to the appropriate interface:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/31 unit 0 classifiers ieee-802.1 fcoe_p3_iscsi
user@switch# set interfaces xe-0/0/32 unit 0 classifiers ieee-802.1 fcoe_p5_iscsi
user@switch# set interfaces xe-0/0/33 unit 0 classifiers ieee-802.1 fcoe_p3_p5
user@switch# set interfaces xe-0/0/34 unit 0 classifiers ieee-802.1 iscsi_classifier
```

7. Configure the CNP input stanza for interface **xe-0/0/31** to enable PFC on the FCoE and iSCSI priorities that the interface handles (code points **011** and **100**), set the MRU value for the FCoE traffic (2240 bytes), and set the cable length value (100 meters). No output stanza is needed because queues 3 and 4 are paused by default on priorities 3 and 4, respectively, and we are not explicitly configuring output queue flow control for any other queues.

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe_p3_cnp input ieee-802.1 code-point
011 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_p3_cnp input ieee-802.1 code-point
100 pfc
user@switch# set congestion-notification-profile fcoe_p3_cnp input cable-length 100
```

8. Configure the CNP for interface xe-0/0/32. The input stanza enables PFC on the FCoE priority (code point 101), sets the MRU value for FCoE traffic (2240 bytes), enables PFC on the iSCSI priority (code point 100), and sets the cable length value (150 meters). The output stanza configures flow control on output queue 5 on the FCoE priority and on output queue 4 on the iSCSI priority:

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe_p5_cnp input ieee-802.1 code-point
100 pfc
user@switch# set congestion-notification-profile fcoe_p5_cnp input ieee-802.1 code-point
101 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_p5_cnp input cable-length 150
user@switch# set congestion-notification-profile fcoe_p5_cnp output ieee-802.1 code-point
100 pfc flow-control-queue 4
user@switch# set congestion-notification-profile fcoe_p5_cnp output ieee-802.1 code-point
101 pfc flow-control-queue 5
```

9. Configure the CNP for interface xe-0/0/33. The input stanza enables PFC on the FCoE priorities (IEEE 802.1p code points 011 and 101), sets the MRU value (2240 bytes), and sets the cable length value (100 meters). The output stanza configures flow control on output queues 3 and 5 on the FCoE priorities:

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe_p3_p5_cnp input ieee-802.1
code-point 011 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_p3_p5_cnp input ieee-802.1
code-point 101 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_p3_p5_cnp input cable-length 100
user@switch# set congestion-notification-profile fcoe_p3_p5_cnp output ieee-802.1
code-point 011 pfc flow-control-queue 3
user@switch# set congestion-notification-profile fcoe_p3_p5_cnp output ieee-802.1
code-point 101 pfc flow-control-queue 5
```

10. Configure the CNP input stanza for interface xe-0/0/34 to enable PFC on the iSCSI priority (code point 100) and set the cable length value (100 meters). No output stanza is needed because queue 4 is paused by default on priority 4, and we are not explicitly configuring output queue flow control for any other queues.

```
[edit class-of-service]
user@switch# set congestion-notification-profile iscsi_cnp input ieee-802.1 code-point
100 pfc
user@switch# set congestion-notification-profile iscsi_cnp input cable-length 100
```

11. Apply each CNP to the appropriate interface:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/31 congestion-notification-profile fcoe_p3_cnp
user@switch# set interfaces xe-0/0/32 congestion-notification-profile fcoe_p5_cnp
user@switch# set interfaces xe-0/0/33 congestion-notification-profile fcoe_p3_p5_cnp
user@switch# set interfaces xe-0/0/34 congestion-notification-profile iscsi_cnp
```

12. Configure the DCBX applications for FCoE and iSCSI to map to the interfaces so that DCBX can exchange application protocol TLVs on the IEEE 802.1p priorities used for FCoE and iSCSI traffic:

```
[edit]
user@switch# set applications application fcoe_app ether-type 0x8906
user@switch# set applications application iscsi_app protocol tcp destination-port 3260
```



13. Configure a DCBX application map to map the FCoE and iSCSI applications to the correct priorities:

```
[edit]
user@switch# set policy-options application-maps dcbx_iscsi_fcoe_app_map application
fcoe_app code-points [011 101]
user@switch# set policy-options application-maps dcbx_iscsi_fcoe_app_map application
iscsi_app code-points 100
```

14. Apply the application map to the interfaces so that DCBX exchanges FCoE application TLVs on the correct code points:

```
[edit]
user@switch# set protocols dcbx interface xe-0/0/31 application-map
dcbx_iscsi_fcoe_app_map
user@switch# set protocols dcbx interface xe-0/0/32 application-map
dcbx_iscsi_fcoe_app_map
user@switch# set protocols dcbx interface xe-0/0/33 application-map
dcbx_iscsi_fcoe_app_map
user@switch# set protocols dcbx interface xe-0/0/34 application-map
dcbx_iscsi_fcoe_app_map
```

## Verification

To verify the configuration and proper operation of the lossless forwarding classes and IEEE 802.1p priorities, perform these tasks:

- [Verifying the Forwarding Class Configuration on page 7009](#)
- [Verifying the Behavior Aggregate Classifier Configuration on page 7010](#)
- [Verifying the PFC Flow Control Configuration \(CNP\) on page 7011](#)
- [Verifying the Interface Configuration on page 7013](#)
- [Verifying the DCBX Application Configuration on page 7014](#)
- [Verifying the DCBX Application Map Configuration on page 7015](#)
- [Verifying the DCBX Application Protocol Exchange Interface Configuration on page 7015](#)

### Verifying the Forwarding Class Configuration

**Purpose** Verify that the lossless forwarding classes **iscsi** and **fcoe1** have been created and that the default lossless forwarding class **fcoe** is still enabled for lossless transport.

**Action** Show the forwarding class configuration by using the operational command **show class-of-service forwarding-class**:

```
user@switch> show class-of-service forwarding-class
```

| Forwarding class | ID | Queue | Policing priority | No-Loss  |
|------------------|----|-------|-------------------|----------|
| best-effort      | 0  | 0     | normal            | Disabled |
| fcoe             | 1  | 3     | normal            | Enabled  |
| iscsi            | 2  | 4     | normal            | Enabled  |
| network-control  | 3  | 7     | normal            | Disabled |
| fcoe1            | 4  | 5     | normal            | Enabled  |
| mcast            | 8  | 8     | normal            | Disabled |

**Meaning** The **show class-of-service forwarding-class** command shows all of the forwarding classes. The command output shows that the **iscsi** and **fcoe1** forwarding classes are configured on output queues 4 and 5, respectively, with the no-loss packet drop attribute enabled.

Because we did not explicitly configure the default **fcoe** forwarding class, it remains in its default state (lossless configuration).

### Verifying the Behavior Aggregate Classifier Configuration

**Purpose** Verify that the four classifiers map the forwarding classes to the correct IEEE 802.1p code points (priorities) and packet loss priorities.

**Action** List the classifiers configured to support lossless FCoE transport using the operational mode command **show class-of-service classifier**:

```
user@switch> show class-of-service classifier
```

Classifier: fcoe\_p3\_iscsi, Code point type: ieee-802.1, Index: 13915

| Code point | Forwarding class | Loss priority |
|------------|------------------|---------------|
| 011        | fcoe             | low           |
| 100        | iscsi            | low           |

Classifier: fcoe\_p5\_iscsi, Code point type: ieee-802.1, Index: 62035

| Code point | Forwarding class | Loss priority |
|------------|------------------|---------------|
| 100        | iscsi            | low           |
| 101        | fcoe1            | low           |

Classifier: fcoe\_p3\_p5, Code point type: ieee-802.1, Index: 17774

| Code point | Forwarding class | Loss priority |
|------------|------------------|---------------|
| 011        | fcoe             | low           |
| 101        | fcoe1            | low           |

Classifier: iscsi\_classifier, Code point type: ieee-802.1, Index: 31635

| Code point | Forwarding class | Loss priority |
|------------|------------------|---------------|
| 100        | iscsi            | low           |

**Meaning** The **show class-of-service classifier** command shows the IEEE 802.1p code points and the loss priorities that are mapped to the forwarding classes in each classifier. The

command output shows that there are four classifiers, **fcoe\_p3\_iscsi**, **fcoe\_p5\_iscsi**, **fcoe\_p3\_p5**, and **iscsi\_classifier**.

Classifier **fcoe\_p3\_iscsi** maps code point **011** (priority 3) to default lossless forwarding class **fcoe** and a packet loss priority of **low**, and code point **100** (priority 4) to explicitly configured lossless forwarding class **iscsi**.

Classifier **fcoe\_p5\_iscsi** maps code point **100** to explicitly configured forwarding class **iscsi** and a packet loss priority of **low**, and code point **101** (priority 5) to explicitly configured lossless forwarding class **fcoe1** and a packet loss priority of **low**.

Classifier **fcoe\_p3\_p5** maps code point **011** to default lossless forwarding class **fcoe** and a packet loss priority of **low**, and maps code point **101** to explicitly configured lossless forwarding class **fcoe1** and a packet loss priority of **low**.

Classifier **iscsi\_classifier** maps code point **100** to explicitly configured forwarding class **iscsi** and a packet loss priority of **low**.

### Verifying the PFC Flow Control Configuration (CNP)

**Purpose** Verify that PFC is enabled on the correct input priorities and that flow control is configured on the correct output queues and priorities in each CNP.

**Action** List the congestion notification profiles using the operational mode command **show class-of-service congestion-notification**:

```
user@switch> show class-of-service congestion-notification
```

```
Name: fcoe_p3_cnp, Index: 12037
```

```
Type: Input
```

```
Cable Length: 100 m
```

| Priority | PFC      | MRU  |
|----------|----------|------|
| 000      | Disabled |      |
| 001      | Disabled |      |
| 010      | Disabled |      |
| 011      | Enabled  | 2240 |
| 100      | Enabled  | 9216 |
| 101      | Disabled |      |
| 110      | Disabled |      |
| 111      | Disabled |      |

```
Type: Output
```

| Priority | Flow-Control-Queues |
|----------|---------------------|
| 000      | 0                   |
| 001      | 1                   |
| 010      | 2                   |
| 011      | 3                   |
| 100      | 4                   |
| 101      | 5                   |
| 110      | 6                   |
| 111      |                     |

7

Name: fcoe\_p3\_p5\_cnp, Index: 46484

Type: Input

Cable Length: 100 m

| Priority | PFC      | MRU  |
|----------|----------|------|
| 000      | Disabled |      |
| 001      | Disabled |      |
| 010      | Disabled |      |
| 011      | Enabled  | 2240 |
| 100      | Disabled |      |
| 101      | Enabled  | 2240 |
| 110      | Disabled |      |
| 111      | Disabled |      |

Type: Output

| Priority | Flow-Control-Queues |
|----------|---------------------|
| 011      |                     |
|          | 3                   |
| 101      |                     |
|          | 5                   |

Name: fcoe\_p5\_cnp, Index: 12133

Type: Input

Cable Length: 150 m

| Priority | PFC      | MRU  |
|----------|----------|------|
| 000      | Disabled |      |
| 001      | Disabled |      |
| 010      | Disabled |      |
| 011      | Disabled |      |
| 100      | Enabled  | 9216 |
| 101      | Enabled  | 2240 |
| 110      | Disabled |      |
| 111      | Disabled |      |

Type: Output

|     |   |
|-----|---|
| 100 |   |
|     | 4 |
| 101 |   |
|     | 5 |

Name: iscsi\_cnp, Index: 19342

Type: Input

Cable Length: 100 m

| Priority | PFC      | MRU  |
|----------|----------|------|
| 000      | Disabled |      |
| 001      | Disabled |      |
| 010      | Disabled |      |
| 011      | Disabled |      |
| 100      | Enabled  | 9216 |
| 101      | Disabled |      |
| 110      | Disabled |      |
| 111      | Disabled |      |

Type: Output

| Priority | Flow-Control-Queues |
|----------|---------------------|
| 000      |                     |
|          | 0                   |
| 001      |                     |
|          | 1                   |
| 010      |                     |
|          | 2                   |
| 011      |                     |
|          | 3                   |

|     |   |
|-----|---|
| 100 | 4 |
| 101 | 5 |
| 110 | 6 |
| 111 | 7 |

**Meaning** The `show class-of-service congestion-notification` command shows the input and output stanzas of the four CNPs.

For CNP `fcoe_p3_cnp`, the input stanza shows that PFC is enabled on IEEE 802.1p code point `011` (priority 3) with an MRU of `2240` bytes, and cable length of `100` meters. The input stanza also shows that PFC is enabled on code point `100` (priority 4) with the default MRU value of `9216` bytes. The CNP output stanza shows the default mapping of priorities to output queues because no explicit output CNP is configured.



**NOTE:** By default, only queues 3 and 4 are enabled respond to pause messages from the connected peer. For queue 3 to respond to pause messages, priority 3 (code point `011`) must be enabled for PFC in the input stanza. For queue 4 to respond to pause messages, priority 4 (code point `100`) must be enabled for PFC in the input stanza. In this example, only queues 3 and 4 respond to pause messages from the connected peer on interfaces that use CNP `fcoe_p3_cnp` because the input stanza enables PFC only on priorities 3 and 4.

For CNP `fcoe_p3_p5_cnp`, the input stanza shows that PFC is enabled on code points `011` and `101` (priority 5), the MRU is `2240` bytes on both priorities, and the cable length is `100` meters. The CNP output stanza shows that output flow control is configured on queues `3` and `5` for code points `011` and `101`, respectively.

For CNP `fcoe_p5_cnp`, the input stanza shows that PFC is enabled on code points `100` and `101`. The MRU for code point `101` (FCoE traffic) is `2240` bytes and the MRU for code point `100` is `9216`. The interface cable length is `150` meters. The CNP output stanza shows that output flow control is configured on queue `4` for code point `100` and on queue `5` for code point `101`.

For CNP `iscsi_cnp`, the input stanza shows that PFC is enabled on code point `100`, the MRU value is `9216` bytes, and the interface cable length is `100` meters. The CNP output stanza shows the default mapping of priorities to output queues because no explicit output CNP is configured.

### Verifying the Interface Configuration

**Purpose** Verify that the correct classifiers and congestion notification profiles are configured on the correct interfaces.

**Action** List the ingress interfaces using the operational mode commands **show configuration class-of-service interfaces xe-0/0/31**, **show configuration class-of-service interfaces xe-0/0/32**, **show configuration class-of-service interfaces xe-0/0/33**, and **show configuration class-of-service interfaces xe-0/0/34**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/31
congestion-notification-profile fcoe_p3_cnp;
unit 0 {
 classifiers {
 ieee-802.1 fcoe_p3_iscsi;
 }
}
```

```
user@switch> show configuration class-of-service interfaces xe-0/0/32
congestion-notification-profile fcoe_p5_cnp;
unit 0 {
 classifiers {
 ieee-802.1 fcoe_p5_iscsi;
 }
}
```

```
user@switch> show configuration class-of-service interfaces xe-0/0/33
congestion-notification-profile fcoe_p3_p5_cnp;
unit 0 {
 classifiers {
 ieee-802.1 fcoe_p3_p5;
 }
}
```

```
user@switch> show configuration class-of-service interfaces xe-0/0/34
congestion-notification-profile iscsi_cnp;
unit 0 {
 classifiers {
 ieee-802.1 iscsi_classifier;
 }
}
```

**Meaning** The **show configuration class-of-service interfaces xe-0/0/31** command shows that the congestion notification profile **fcoe\_p3\_cnp** is configured on the interface, and that the IEEE 802.1p classifier associated with the interface is **fcoe\_p3\_iscsi**.

The **show configuration class-of-service interfaces xe-0/0/32** command shows that the congestion notification profile **fcoe\_p5\_cnp** is configured on the interface, and that the IEEE 802.1p classifier associated with the interface is **fcoe\_p5\_iscsi**.

The **show configuration class-of-service interfaces xe-0/0/33** command shows that the congestion notification profile **fcoe\_p3\_p5\_cnp** is configured on the interface, and that the IEEE 802.1p classifier associated with the interface is **fcoe\_p3\_p5**.

The **show configuration class-of-service interfaces xe-0/0/34** command shows that the congestion notification profile **iscsi\_cnp** is configured on the interface, and that the IEEE 802.1p classifier associated with the interface is **iscsi\_classifier**.

---

### Verifying the DCBX Application Configuration

**Purpose** Verify that the DCBX applications for FCoE and iSCSI are configured.

**Action** List the DCBX applications by using the configuration mode command **show applications**:

```
user@switch# show applications
application iscsi_app {
 protocol tcp;
 destination-port 3260;
}
application fcoe_app {
 ether-type 0x8906;
```

**Meaning** The **show applications** configuration mode command shows all of the configured applications. The output shows that the application **iscsi\_app** is configured with a protocol value of **tcp** and a destination port value of **3260**, and that the application **fcoe\_app** is configured with an EtherType of **0x8906** (the correct EtherType for FCoE traffic).

---

### Verifying the DCBX Application Map Configuration

**Purpose** Verify that the application map is configured.

**Action** List the application maps by using the configuration mode command **show policy-options application-maps**:

```
user@switch# show policy-options application-maps
dcbx-iscsi-fcoe-app-map {
 application iscsi_app code-points 100;
 application fcoe_app code-points [011 101];
}
```

**Meaning** The **show policy-options application-maps** configuration mode command lists all of the configured application maps and the applications that belong to each application map. The output shows that there is one application map named **dcbx-iscsi-fcoe\_app\_map**. It consists of the application **iscsi\_app** mapped to code point **100** and the application **fcoe\_app** mapped to code points **011** and **101**.

---

### Verifying the DCBX Application Protocol Exchange Interface Configuration

**Purpose** Verify that the application maps are applied to the correct interfaces.

**Action** List the application maps on each interface using the configuration mode command **show protocols dcbx**:

```
user@switch# show protocols dcbx
interface xe-0/0/31.0 {
 application-map dcbx-iscsi-fcoe-app-map;
}
interface xe-0/0/32.0 {
 application-map dcbx-iscsi-fcoe-app-map;
}
interface xe-0/0/33.0 {
 application-map dcbx-iscsi-fcoe-app-map;
}
interface xe-0/0/34.0 {
 application-map dcbx-iscsi-fcoe-app-map;
}
```

**Meaning** The **show protocols dcbx** configuration mode command lists the application map association with interfaces. The output shows that all four interfaces use the application map **dcbx-iscsi-fcoe-app-map**.

- Related Documentation**
- [Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 6977](#)
  - [Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic \(FCoE Transit Switch\) on page 6969](#)
  - [Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 6986](#)
  - [Example: Configuring DCBX Application Protocol TLV Exchange on page 6535](#)
  - [Configuring CoS PFC \(Congestion Notification Profiles\) on page 6905](#)
  - [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 6873](#)
  - [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 6493](#)

---

## Understanding DCBX

Data Center Bridging Capability Exchange protocol (DCBX) is an extension of Link Layer Data Protocol (LLDP). If you disable LLDP on an interface, that interface cannot run DCBX. If you attempt to enable DCBX on an interface on which LLDP is disabled, the configuration commit operation fails. Data center bridging (DCB) devices use DCBX to exchange configuration information with directly connected peers.



Video: [What is DCBX Protocol?](#)

---

This topic describes:

- [DCBX Basics on page 7016](#)
- [DCBX Modes and Support on page 7018](#)
- [DCBX Attribute Types on page 7020](#)
- [DCBX Application Protocol TLV Exchange on page 7022](#)
- [DCBX and PFC on page 7023](#)
- [DCBX and ETS on page 7023](#)

### DCBX Basics

DCBX can:

- Discover the DCB capabilities of peers.
- Detect DCB feature misconfiguration or mismatches between peers.
- Configure DCB features on peers.



You can configure DCBX operation for priority-based flow control (PFC), Layer 2 and Layer 4 applications such as FCoE and iSCSI, and ETS. DCBX is enabled or disabled on a per-interface basis.



**NOTE:** The Juniper Networks QFX10000 does not support enhanced transmission selection (ETS) hierarchical scheduling. Use port scheduling to manage bandwidth on QFX10000 switches.

By default, for PFC and ETS, DCBX automatically negotiates administrative state and configuration with each interface's connected peer. To enable DCBX negotiation for applications, you must configure the applications, map them to IEEE 802.1p code points in an application map, and apply the application map to interfaces.

The FCoE application only needs to be included in an application map when you want an interface to exchange type, length, and values (TLVs) for other applications in addition to FCoE. If FCoE is the only application you want an interface to advertise, then you do not need to use an application map. For ETS, DCBX pushes the switch configuration to peers if they are set to learn the configuration from the switch (unless you disable sending the ETS recommendation TLV on interfaces in IEEE DCBX mode).

You can override the default behavior for PFC, for ETS, or for all applications mapped to an interface by turning off autonegotiation to force an interface to enable or disable that feature. You can also disable DCBX autonegotiation for applications on an interface by excluding those applications from the application map you apply to that interface or by deleting the application map from the interface.

The default autonegotiation behavior for applications that are mapped to an interface is:

- DCBX is enabled on the interface if the connected peer device also supports DCBX.
- DCBX is disabled on the interface if the connected peer device does not support DCBX.

During negotiation of capabilities, the switch can push the PFC configuration to an attached peer if the peer is configured as “willing” to learn the PFC configuration from other peers. The Juniper Networks switch does not support self autoprovisioning and does not change its configuration during autonegotiation to match the peer configuration. (The Juniper switch is not “willing” to learn the PFC configuration from peers.)



**NOTE:** When a port with DCBX enabled begins to exchange type, length, and value (TLV) entries, optional LLDP TLVs on that port are not advertised to neighbors, so that the switch can interoperate with a wider variety of converged network adapters (CNAs) and Layer 2 switches that support DCBX.

## DCBX Modes and Support

This section describes DCBX support:

- [DCBX Modes \(Versions\) on page 7018](#)
- [Autonegotiation on page 7020](#)
- [CNA Support for DCBX Modes on page 7020](#)
- [Interface Support for DCBX on page 7020](#)

### DCBX Modes (Versions)

The two most common DCBX modes are supported:

- IEEE DCBX—The newest DCBX version. Different TLVs have different subtypes (for example, the subtype for the ETS configuration TLV is 9); the IEEE DCBX Organizationally Unique Identifier (OUI) is 0x0080c2.
- DCBX version 1.01—The Converged Enhanced Ethernet (CEE) version of DCBX. It has a subtype of 2 and an OUI of 0x001b21.

IEEE DCBX and DCBX version 1.01 differ mainly in frame format. DCBX version 1.01 uses one TLV that includes all DCBX attribute information, which is sent as sub-TLVs. IEEE DCBX uses a unique TLV for each DCB attribute.



**NOTE:** The switch does not support pre-CEE (pre-DCB) DCBX versions. Unsupported older versions of DCBX have a subtype of 1 and an OUI of 0x001b21. The switch drops LLDP frames that contain pre-CEE DCBX TLVs.

[Table 518](#) summarizes the differences between IEEE DCBX and DCBX version 1.01, including show command output:

**Table 599: Summary of Differences Between IEEE DCBX and DCBX Version 1.01**

| Characteristic                               | IEEE DCBX                                                                                                                                                                                                  | DCBX Version 1.01                                                                                                                                                                                    |
|----------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| OUI                                          | 0x0080c2                                                                                                                                                                                                   | 0x001b21                                                                                                                                                                                             |
| Frame Format                                 | Sends a separate, unique TLV for each DCBX attribute. For example, IEEE DCBX uses separate TLVs for ETS, PFC, and each application. Configuration and Recommendation information is sent in different TLVs | Sends one TLV that includes all DCBX attribute information organized in sub-TLVs. The “willing” bit determines whether or not an interface can change its configuration to match the connected peer. |
| Symmetric/asymmetric configuration with peer | Asymmetric or symmetric                                                                                                                                                                                    | Symmetric only                                                                                                                                                                                       |

**Table 599: Summary of Differences Between IEEE DCBX and DCBX Version 1.01** *(continued)*

| Characteristic                                                                   | IEEE DCBX                                                                                                                                                                                                                                                                                                                                                                                                                                                                | DCBX Version 1.01                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|----------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Differences in the <b>show dcbx interface interface-name</b> operational command | <ul style="list-style-type: none"> <li>Synchronization information is not shown because symmetric configuration is not required.</li> <li>Operational state information is not shown because the operational states do not have to be symmetric.</li> <li>TLV type is shown because unique TLVs are sent for each DCBX attribute.</li> <li>ETS peer Configuration TLV and Recommendation TLV information is shown separately because they are different TLVs.</li> </ul> | <ul style="list-style-type: none"> <li>Synchronization information is shown because symmetric configuration is required.</li> <li>Operational state information is shown because the operational states do have to be symmetric.</li> <li>TLV type is not shown because one TLV is used for all attribute information.</li> <li>Recommendation TLV is not sent (DCBX Version 1.01 uses the "willing" bit to determine whether or not an interface uses the peer interface configuration).</li> </ul> |

For more information about how each DCBX mode exchanges TLVs, see the following specifications:

- For DCBX version 1.01—<http://www.ieee802.org/1/files/public/docs2008/az-wedlar-dcb-capability-exchange-discovery-protocol-108-v101.pdf>
- For IEEE DCBX—<http://www.ieee802.org/1/files/private/az-drafts/d2/802-1az-d2-4.pdf>



**NOTE:** As of Junos OS Release 12.2, this document is located in a private area of the IEEE website, and access requires a password from the IEEE organization. If you are not an IEEE member, you might not be able to access this document until it moves to the public area of the IEEE website.

You can configure interfaces to use the following DCBX modes:

- IEEE DCBX—The interface uses IEEE DCBX regardless of the configuration on the connected peer.
- DCBX version 1.01—The interface uses DCBX version 1.01 regardless of the configuration on the connected peer.
- Autonegotiation—The interface automatically negotiates with the connected peer to determine the DCBX version the peers use. Autonegotiation is the default DCBX mode.

If you configure a DCBX mode on an interface, the interface ignores DCBX protocol data units (PDUs) it receives from the connected peer if the PDUs do not match the DCBX version configured on the interface. For example, if you configure an interface to use IEEE DCBX and the connected peer sends DCBX version 1.01 LLDP PDUs, the interface ignores the version 1.01 PDUs. If you configure an interface to use DCBX version 1.01 and the peer sends IEEE DCBX LLDP PDUs, the interface ignores the IEEE DCBX PDUs.



**NOTE:** On interfaces that use the IEEE DCBX mode, the `show dcbx neighbors interface interface-name` operational command does not include application, PFC, or ETS operational state in the output.

---

### Autonegotiation

Autonegotiation is the default DCBX mode. Each interface automatically negotiates with its connected peer to determine the DCBX version that both interfaces use to exchange DCBX information.

When an interface connects to its peer interface, the interface advertises IEEE DCBX TLVs to the peer. If the interface receives one IEEE DCBX PDU from the peer, the interface sets the DCBX mode as IEEE DCBX. If the interface receives three DCBX version 1.01 TLVs from the peer, the interface sets DCBX version 1.01 as the DCBX mode.

Autonegotiation works slightly differently on standalone switches compared to QFabric systems:

- Standalone switches—When an interface connects to its peer interface, the interface advertises IEEE DCBX TLVs to the peer. If the interface receives an IEEE DCBX TLV from the peer, the interface sets IEEE DCBX as the DCBX mode. If the interface receives three consecutive DCBX version 1.01 TLVs from the peer, the interface sets DCBX version 1.01 as the DCBX mode.
- QFabric system—When an interface connects to its peer interface, the interface advertises DCBX version 1.01 TLVs to the peer. If the interface receives an IEEE DCBX TLVs from the peer, the interface sets IEEE DCBX as the DCBX mode. If the interface receives three consecutive DCBX version 1.01 TLVs from the peer, the interface retains DCBX version 1.01 as the DCBX mode.



**NOTE:** If the link flaps or the LLDP process restarts, the interface starts the autonegotiation process again. The interface does not use the last received DCBX communication mode.

---

### CNA Support for DCBX Modes

Different CNA vendors support different versions and capabilities of DCBX. The DCBX configuration you use on switch interfaces depends on the DCBX features that the CNAs in your network support.

---

### Interface Support for DCBX

You can configure DCBX on 10-Gigabit Ethernet interfaces and on link aggregation group (LAG) interfaces whose member interfaces are all 10-Gigabit Ethernet interfaces.

## DCBX Attribute Types

DCBX has three attribute types:

- **Informational**—These attributes are exchanged using LLDP, but do not affect DCBX state or operation; they only communicate information to the peer. For example, application priority TLVs are informational TLVs.
- **Asymmetric**—The values for these types of attributes do not have to be the same on the connected peer interfaces. Peers exchange asymmetric attributes when the attribute values can differ on each peer interface. The peer interface configurations might match or they might differ. For example, ETS Configuration and Recommendation TLVs are asymmetric TLVs.
- **Symmetric**—The intention is that the values for these types of attributes should be the same on both of the connected peer interfaces. Peer interfaces exchange symmetric attributes to ensure symmetric DCBX configuration for those attributes. For example, PFC Configuration TLVs are symmetric TLVs.

The following sections describe asymmetric and symmetric DCBX attributes:

- [Asymmetric Attributes on page 7021](#)
- [Symmetric Attributes on page 7021](#)

---

### Asymmetric Attributes

DCBX passes asymmetric attributes between connected peer interfaces to communicate parameter information about those attributes (features). The resulting configuration for an attribute might be different on each peer, so the parameters configured on one interface might not match the parameters on the connected peer interface.

There are two types of asymmetric attribute TLVs:

- **Configuration TLV**—Configuration TLVs communicate the current operational state and the state of the “willing” bit. The “willing” bit communicates whether or not the interface is willing to accept and use the configuration from the peer interface. If an interface is “willing,” the interface uses the configuration it receives from the peer interface. (The peer interface configuration can override the configuration on the “willing” interface.) If an interface is “not willing”, the configuration on the interface cannot be overridden by the peer interface configuration.
- **Recommendation TLV**—Recommendation TLVs communicate the parameters the interface recommends that the connected peer interface should use. When an interface sends a Recommendation TLV, if the connected peer is “willing,” the connected peer changes its configuration to match the parameters in the Recommendation TLV.

---

### Symmetric Attributes

DCBX passes symmetric attributes between connected peer interfaces to communicate parameter information about those attributes (features), with the objective that both interfaces should use the same configuration. The intent is that the parameters configured on one interface should match the parameters on the connected peer interface.

There is one type of symmetric attribute TLV, the Configuration TLV. As with asymmetric attributes, symmetric attribute Configuration TLVs communicate the current operational state and the state of the “willing” bit. “Willing” interfaces use the peer interface parameter

values for the attribute. (The attribute configuration of the peer overrides the configuration on the “willing” interface.)

## DCBX Application Protocol TLV Exchange

DCBX advertises the switch's capabilities for Layer 2 applications such as FCoE and Layer 4 applications such as iSCSI:

- [Application Protocol TLV Exchange on page 7022](#)
- [FCoE Application Protocol TLV Exchange on page 7022](#)
- [Disabling Application Protocol TLV Exchange on page 7023](#)

---

### Application Protocol TLV Exchange

For all applications, DCBX advertises the application's state and IEEE 802.1p code points on the interfaces to which the application is mapped. If an application is not mapped to an interface, that interface does not advertise the application's TLVs. There is an exception for FCoE application protocol TLV exchange when FCoE is the only application you want DCBX to advertise on an interface.

---

### FCoE Application Protocol TLV Exchange

Protocol TLV exchange for the FCoE application depends on whether FCoE is the only application you want the interface to advertise or whether you want the interface to exchange other application TLVs in addition to FCoE TLVs.

If FCoE is the only application you want DCBX to advertise on an interface, DCBX exchanges FCoE application protocol TLVs by default if the interface:

- Carries FCoE traffic (traffic mapped by CoS configuration to the FCoE forwarding class)
- Has a congestion notification profile with PFC enabled on the FCoE priority (IEEE 802.1p code point)
- Does *not* have an application map



**NOTE:** If no CoS configuration for FCoE is mapped to an interface, that interface does not exchange FCoE application protocol TLVs.

---

If you want DCBX to advertise FCoE and other applications on an interface, you must specify all of the applications, including FCoE, in an application map, and apply the application map to the desired interfaces.

---



**NOTE:** If an application map is applied to an interface, the FCoE application must be explicitly configured in the application map, or the interface does not exchange FCoE TLVs.

---

When DCBX advertises the FCoE application, it advertises the FCoE state and IEEE 802.1p code points. If a peer device connected to a switch interface does not support FCoE,

DCBX uses autonegotiation to mark the interface as “FCoE down,” and FCoE is disabled on that interface.

### Disabling Application Protocol TLV Exchange

---

To disable DCBX application protocol exchange for all applications on an interface, issue the **set protocols dcbx interface *interface-name* applications no-auto-negotiation** command.

You can also disable DCBX application protocol exchange for applications on an interface by deleting the application map from the interface, or by deleting a particular application from the application map. However, when you delete an application from an application map, the application protocol is no longer exchanged on any interface which uses that application map.

## DCBX and PFC

After you enable PFC on a switch interface, DCBX uses autonegotiation to control the operational state of the PFC functionality.

If the peer device connected to the interface supports PFC and is provisioned compatibly with the switch, DCBX sets the PFC operational state to enabled. If the peer device connected to the interface does not support PFC or is not provisioned compatibly with the switch, DCBX sets the operational state to disabled. (PFC must be symmetrical.)

If the peer advertises that it is “willing” to learn its PFC configuration from the switch, DCBX pushes the switch’s PFC configuration to the peer and does not check the peer’s administrative state.

You can manually override DCBX control of the PFC operational state on a per-interface basis by disabling autonegotiation. If you disable autonegotiation on an interface on which you have configured PFC, then PFC is enabled on that interface regardless of the peer configuration. To disable PFC on an interface, do not configure PFC on that interface.

## DCBX and ETS

This section describes:

- [Default DCBX ETS Advertisement on page 7023](#)
- [ETS Advertisement and Peer Configuration on page 7024](#)
- [ETS Recommendation TLV on page 7024](#)

### Default DCBX ETS Advertisement

---

If you do not configure ETS on an interface, the switch automatically creates a default priority group that contains all of the priorities (forwarding classes, which represent output queues) and assigns 100 percent of the port output bandwidth to that priority group. The default priority group is transparent. It does not appear in the configuration and is used for DCBX advertisement. DCBX advertises the default priority group, its priorities, and the assigned bandwidth.

If you configure ETS on an interface, DCBX advertises:

- Each priority group on the interface

- The priorities in each priority group
- The bandwidth properties of each priority group and priority

Any priority on that interface that is not part of an explicitly configured priority group (forwarding class set) is assigned to the automatically generated default priority group and receives no bandwidth. If you configure ETS on an interface, every forwarding class (priority) on that interface for which you want to forward traffic must belong to a forwarding class set (priority group).

---

### ETS Advertisement and Peer Configuration

DCBX does not control the switch's ETS (hierarchical scheduling) operational state. If the connected peer is configured as "willing," DCBX pushes the switch's ETS configuration to the switch's peers if the ETS Recommendation TLV is enabled (it is enabled by default). If the peer does not support ETS or is not consistently provisioned with the switch, DCBX does not change the ETS operational state on the switch. The ETS operational state remains enabled or disabled based only on the switch hierarchical scheduling configuration and is enabled by default.

When ETS is configured, DCBX advertises the priority groups, the priorities in the priority groups, and the bandwidth configuration for the priority groups and priorities. Any priority (essentially a forwarding class or queue) that is not part of a priority group has no scheduling properties and receives no bandwidth.

You can manually override whether DCBX advertises the ETS state to the peer on a per-interface basis by disabling autonegotiation. This does not affect the ETS state on the switch or on the peer, but it does prevent the switch from sending the Recommendation TLV or the Configuration TLV to the connected peer. To disable ETS on an interface, do not configure priority groups (forwarding class sets) on the interface.

---

### ETS Recommendation TLV

The ETS Recommendation TLV communicates the ETS settings that the switch wants the connected peer interface to use. If the peer interface is "willing," it changes its configuration to match the configuration in the ETS Recommendation TLV. By default, the switch interfaces send the ETS Recommendation TLV to the peer. The settings communicated are the egress ETS settings defined by configuring hierarchical scheduling on the interface.

We recommend that you use the same ETS settings on the connected peer that you use on the switch interface and that you leave the ETS Recommendation TLV enabled. However, on interfaces that use IEEE DCBX as the DCBX mode, if you want an asymmetric configuration between the switch interface and the connected peer, you can disable the ETS Recommendation TLV by including the **no-recommendation-tlv** statement at the **[edit protocols dcbx interface *interface-name* enhanced-transmission-selection]** hierarchy level.





**NOTE:** You can disable the ETS Recommendation TLV only when the DCBX mode on the interface is IEEE DCBX. Disabling the ETS Recommendation TLV has no effect if the DCBX mode on the interface is DCBX version 1.01. (IEEE DCBX uses separate application attribute TLVs, but DCBX version 1.01 sends all application attributes in the same TLV and uses sub-TLVs to separate the information.)

If you disable the ETS Recommendation TLV, the switch still sends the ETS Configuration TLV to the connected peer. The result is that the connected peer is informed about the switch DCBX ETS configuration, but even if the peer is “willing,” the peer does not change its configuration to match the switch configuration. This is asymmetric configuration—the two interfaces can have different parameter values for the ETS attribute.

For example, if you want a CNA connected to a switch interface to have different bandwidth allocations than the switch ETS configuration, you can disable the ETS Recommendation TLV and configure the CNA for the desired bandwidth. The switch interface and the CNA exchange configuration parameters, but the CNA does not change its configuration to match the switch interface configuration.

**Related  
Documentation**

- [Understanding DCBX Application Protocol TLV Exchange on page 6528](#)
- [Understanding DCB Features and Requirements on page 6490](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 6493](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 6765](#)
- [Understanding CoS Port Schedulers on QFX Switches on page 6729](#)
- [Understanding FCoE on page 6436](#)
- [Configuring the DCBX Mode on page 6523](#)
- [Configuring DCBX Autonegotiation on page 6524](#)
- [Disabling the ETS Recommendation TLV on page 6527](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 6535](#)

## Configuring the DCBX Mode

You can configure the DCBX mode that an interface uses to communicate with the connected peer. Three DCBX modes are supported:

- **Autonegotiation**—The interface negotiates with the connected peer to determine the DCBX mode. This is the default DCBX mode.
- **IEEE DCBX**—The interface uses IEEE DCBX type, length, and value (TLV) to exchange DCBX information with the connected peer. QFX3500 Node devices come up with IEEE DCBX enabled by default and then autonegotiate with the connected peer to determine the final DCBX mode.
- **DCBX Version 1.01**—The interface uses Converged Enhanced Ethernet (CEE) DCBX version 1.01 TLVs to exchange DCBX information with the connected peer. QFabric system Node devices other than QFX3500 switches come up with DCBX version 1.01 enabled by default and then autonegotiate with the connected peer to determine the final DCBX mode.



**NOTE:** Pre-CEE (pre-DCB) versions of DCBX such as DCBX version 1.00 are not supported. If an interface receives an LLDP frame with pre-CEE DCBX TLVs, the system drops the frame.

Configure the DCBX mode by specifying the mode for one interface or for all interfaces.

- To configure the DCBX mode, specify the interface and the mode:

```
[edit protocols dcbx]
user@switch# set interface interface-name mode (auto-negotiate | ieee-dcbx |
dcbx-version-1.01)
```

For example, to configure DCBX version 1.01 on interface **xe-0/0/21**:

```
user@switch# set protocols dcbx interface xe-0/0/21 mode dcbx-version-1.01
```

To configure IEEE DCBX on all interfaces:

```
user@switch# set protocols dcbx interface all mode ieee-dcbx
```

### Related Documentation

- [Configuring DCBX Autonegotiation on page 6524](#)
- [Disabling the ETS Recommendation TLV on page 6527](#)
- [Understanding DCBX on page 6514](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 6528](#)
- [show dcbx neighbors on page 6581](#)

## Configuring DCBX Autonegotiation

Data Center Bridging Capability Exchange protocol (DCBX) discovers the data center bridging (DCB) capabilities of peers by exchanging feature configuration information. DCBX also detects feature misconfiguration and mismatches, and can configure DCB on peers. DCBX is an extension of the Link Layer Discovery Protocol (LLDP), and LLDP must remain enabled on every interface for which you want to use DCBX. If you attempt to enable DCBX on an interface on which LLDP is disabled, the configuration commit operation fails.



**NOTE:** LLDP and DCBX are enabled by default on all interfaces.

The switch supports DCBX autonegotiation for:

- Priority-based flow control (PFC) configuration
- Layer 2 and Layer 4 applications such as Fibre Channel over Ethernet (FCoE) and Internet Small Computer System Interface (iSCSI)
- Enhanced transmission selection (ETS) advertisement

DCBX autonegotiation is configured on a per-interface basis for each supported feature or application. The PFC and application DCBX exchanges use autonegotiation by default. The default autonegotiation behavior is:

- DCBX is enabled on the interface if the connected peer device also supports DCBX.
- DCBX is disabled on the interface if the connected peer device does not support DCBX.

You can override the default behavior for each feature by turning off autonegotiation to force an interface to enable or disable the feature.

Autonegotiation of ETS means that when ETS is enabled on an interface (priority groups are configured), the interface advertises its ETS configuration to the peer device. In this case, priorities (forwarding classes) that are not part of a priority group (forwarding class set) receive no bandwidth and are advertised in an automatically generated default forwarding class. If ETS is not enabled on an interface (no priority groups are configured), all of the priorities are advertised in one automatically generated default priority group that receives 100 percent of the port bandwidth.

Disabling ETS autonegotiation prevents the interface from sending the Recommendation TLV or the Configuration TLV to the connected peer.

On interfaces that use IEEE DCBX mode to exchange DCBX parameters, you can disable autonegotiation of the ETS Recommendation TLV to the peer if you want an asymmetric ETS configuration between the peers. DCBX still exchanges the ETS Configuration TLV if you disable the ETS Recommendation TLV.

Autonegotiation of PFC means that when PFC is enabled on an interface, if the peer device connected to the interface supports PFC and is provisioned compatibly with the

switch, DCBX sets the PFC operational state to enabled. If the peer device connected to the interface does not support PFC or is not provisioned compatibly with the switch, DCBX sets the operational state to disabled.

In addition, if the peer advertises that it is “willing” to learn its PFC configuration from the switch, DCBX pushes the switch’s PFC configuration to the peer and does not check the peer’s administrative state. The switch does not learn PFC configuration from peers (the switch does not advertise its state as “willing”).

Disabling PFC autonegotiation prevents the interface from exchanging PFC configuration information with the peer. It forces the interface to enable PFC if PFC is configured on the interface or to disable PFC if PFC is not configured on the interface. If you disable PFC autonegotiation, the assumption is that the peer is also configured manually.

Autonegotiation of applications depends on whether or not you apply an application map to an interface. If you apply an application map to an interface, the interface autonegotiates DCBX for each application in the application map. PFC must be enabled on the FCoE priority (the FCoE IEEE 802.1p code point) for the interface to advertise the FCoE application. The interface only advertises applications that are included in the application map.

For example, if you apply an application map to an interface and the application map does not include the FCoE application, then that interface does not perform DCBX advertisement of FCoE.

If you do not apply an application map to an interface, DCBX does not advertise applications on that interface, with the exception of FCoE, which is handled differently than other applications.



**NOTE:** If you do not apply an application map to an interface, the interface performs autonegotiation of FCoE if the interface carries traffic in the FCoE forwarding class and also has PFC enabled on the FCoE priority. On such interfaces, if DCBX detects that the peer device connected to the interface supports FCoE, the switch advertises its FCoE capability and IEEE 802.1p code point on that interface. If DCBX detects that the peer device connected to the interface does not support FCoE, DCBX marks that interface as “FCoE down” and disables FCoE on the interface.

---

When DCBX marks an interface as “FCoE down,” the behavior of the switch depends on how you use it in the network:

- When the switch acts as an FCoE transit switch, the interface drops all of the FIP packets it receives. In addition, FIP packets received from an FCoE forwarder (FCF) are not forwarded to interfaces marked as “FCoE down.”
- When the switch acts as an FCoE-FC gateway (only switches that support native Fibre Channel interfaces), it does not send or receive FCoE Initialization Protocol (FIP) packets.

Disabling autonegotiation prevents the interface from exchanging application information with the peer. In this case, the assumption is that the peer is also configured manually.

To disable DCBX autonegotiation of PFC, applications (including FCoE), and ETS using the CLI:

1. Turn off autonegotiation for PFC.

```
[edit]
user@switch# set protocols dcbx interface interface-name priority-flow-control
no-auto-negotiation
```

2. Turn off autonegotiation for applications.

```
[edit]
user@switch# set protocols dcbx interface interface-name applications no-auto-negotiation
```

3. Turn off autonegotiation for ETS.

```
[edit]
user@switch# set protocols dcbx interface interface-name enhanced-transmission-selection
no-auto-negotiation
```

To disable autonegotiation of the ETS Recommendation TLV so that DCBX exchanges only the ETS Configuration TLV:

- [edit protocols dcbx interface *interface-name*]  
user@switch# set enhanced-transmission-selection no-recommendation-tlv

#### Related Documentation

- [Example: Configuring DCBX Application Protocol TLV Exchange on page 6535](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 6504](#)
- [Disabling the ETS Recommendation TLV on page 6527](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 6528](#)

## Understanding DCBX Application Protocol TLV Exchange

---

Data Center Bridging Capability Exchange protocol (DCBX) discovers the data center bridging (DCB) capabilities of connected peers. DCBX also advertises the capabilities of applications on interfaces by exchanging application protocol information through application type, length, and value (TLV) elements. DCBX is an extension of Link Layer Discovery Protocol (LLDP). LLDP must remain enabled on every interface on which you want to use DCBX.



**NOTE:** LLDP and DCBX are enabled by default on all interfaces.

Setting up application protocol exchange consists of:

- Defining applications
- Mapping the applications to IEEE 802.1p code points in an *application map*
- Configuring classifiers to prioritize incoming traffic and map the incoming traffic to the application by the traffic code points
- Applying the application maps and classifiers to interfaces

You need to explicitly define the applications that you want an interface to advertise. The FCoE application is a special case (see [“Applications” on page 6528](#)) and only needs to be defined on an interface if you want DCBX to exchange application protocol TLVs for other applications in addition to FCoE on that interface.

You also need to explicitly map all of the defined applications that you want an interface to advertise to IEEE 802.1p code points in an application map. The FCoE application is a special case that only requires inclusion in an application map when you want an interface to use DCBX for other applications in addition to FCoE, as described later in this topic (see [“Application Maps” on page 6529](#)).

This topic describes:

- [Applications on page 7030](#)
- [Application Maps on page 7031](#)
- [Classifying and Prioritizing Application Traffic on page 7032](#)
- [Enabling Interfaces to Exchange Application Protocol Information on page 7033](#)
- [Disabling DCBX Application Protocol Exchange on page 7033](#)

### Applications

Before an interface can exchange application protocol information, you need to define the applications that you want to advertise. The exception is the FCoE application. If FCoE is the only application that you want the interface to advertise, then you do not need to define the FCoE application. You need to define the FCoE application only if you want interfaces to advertise other applications in addition to FCoE.



**NOTE:** If FCoE is the only application that you want DCBX to advertise on an interface, DCBX exchanges FCoE application protocol TLVs by default if the interface:

- Carries FCoE traffic (traffic mapped by CoS configuration to the FCoE forwarding class and applied to the interface)
- Has a congestion notification profile with PFC enabled on the FCoE priority (IEEE 802.1p code point)
- Does *not* have an application map

If you apply an application map to an interface, then all applications that you want DCBX to advertise must be defined and configured in the application map, including the FCoE application.

If no CoS configuration for FCoE is mapped to an interface, that interface does not exchange FCoE application protocol TLVs.

You can define:

- Layer 2 applications by EtherType
- Layer 4 applications by a combination of protocol (TCP or UDP) and destination port number

The EtherType is a two-octet field in the Ethernet frame that denotes the protocol encapsulated in the frame. For a list of common EtherTypes, see <http://standards.ieee.org/develop/regauth/ethertype/eth.txt> on the IEEE standards organization website. For a list of port numbers and protocols, see the *Service Name and Transport Protocol Port Number Registry* at <http://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xml> on the Internet Assigned Numbers Authority (IANA) website.

You must explicitly define each application that you want to advertise, except FCoE. The FCoE application is defined by default (EtherType 0x8906).

## Application Maps

An application map maps defined applications to one or more IEEE 802.1p code points. Each application map contains one or more applications. DCBX includes the configured application code points in the protocol TLVs exchanged with the connected peer.

To exchange protocol TLVs for an application, you must include the application in an application map. The FCoE application is a special case:

- If you want DCBX to exchange application protocol TLVs for more than one application on a particular interface, you must configure the applications, define an application map to map the applications to code points, and apply the application map to the interface. In this case, you must also define the FCoE application and add it to the application map.

This is the same process and treatment required for all other applications. In addition, for DCBX to exchange FCoE application TLVs, you must enable priority-based flow control (PFC) on the FCoE priority (the FCoE IEEE 802.1p code point) on the interface.

- If FCoE is the only application that you want DCBX to advertise on an interface, then you do not need to configure an application map and apply it to the interface. By default, when an interface has no application map, and the interface carries traffic mapped to the FCoE forwarding class, and PFC is enabled on the FCoE priority, the interface advertises FCoE TLVs (autonegotiation mode). DCBX exchanges FCoE application protocol TLVs by default until you apply an application map to the interface, remove the FCoE traffic from the interface (you can do this by removing the or editing the classifier for FCoE traffic), or disable PFC on the FCoE priority.

If you apply an application map to an interface that did not have an application map and was exchanging FCoE application TLVs, and you do not include the FCoE application in the application map, the interface stops exchanging FCoE TLVs. Every interface that has an application map must have FCoE included in the application map (and PFC enabled on the FCoE priority) in order for DCBX to exchange FCoE TLVs.

Mapping an application to code points does two things:

- Maps incoming traffic with the same code points to that application
- Allows you to configure classifiers that map incoming application traffic, by code point, to a forwarding class and a loss priority, in order to apply class of service (CoS) to application traffic and prioritize application traffic

You apply an application map to an interface to enable DCBX application protocol exchange on that interface for each application specified in the application map. All of the applications that you want an interface to advertise must be configured in the application map that you apply to the interface, with the previously noted exception for the FCoE application when FCoE is the only application for which you want DCBX to exchange protocol TLVs on an interface.

## Classifying and Prioritizing Application Traffic

When traffic arrives at an interface, the interface classifies the incoming traffic based on its code points. Classifiers map code points to loss priorities and forwarding classes. The loss priority prioritizes the traffic. The forwarding class determines the traffic output queue and CoS service level.

When you map an application to an IEEE 802.1p code point in an application map and apply the application map to an interface, incoming traffic on the interface that matches the application code points is mapped to the appropriate application. The application receives the loss priority and the CoS associated with the forwarding class for those code points, and is placed in the output queue associated with the forwarding class.

You can use the default classifier or you can configure a classifier to map the application code points defined in the application map to forwarding classes and loss priorities.



## Enabling Interfaces to Exchange Application Protocol Information

Each interface with the **fcoe** forwarding class and PFC enabled on the FCoE code point is enabled for FCoE application protocol exchange by default until you apply an application map to the interface. If you apply an application map to an interface and you want that interface to exchange FCoE application protocol TLVs, you must include the FCoE application in the application map. (In all cases, to achieve lossless transport, you must also enable PFC on the FCoE code point or code points.)

Except when FCoE is the only protocol you want DCBX to advertise on an interface, interfaces on which you want to exchange application protocol TLVs must include the following two items:

- The application map that contains the application(s)
- A classifier



**NOTE:** You must also enable PFC on the code point of any traffic for which you want to achieve lossless transport.

## Disabling DCBX Application Protocol Exchange

To disable DCBX application protocol exchange for all applications on an interface, issue the **set protocols dcbx interface *interface-name* applications no-auto-negotiation** command.

You can also disable DCBX application protocol exchange for applications on an interface by deleting the application map from the interface, or by deleting a particular application from the application map. However, when you delete an application from an application map, the application protocol is no longer exchanged on any interface which uses that application map.

On interfaces that use IEEE DCBX mode to exchange DCBX parameters, you can disable sending the enhanced transmission selection (ETS) Recommendation TLV to the peer if you want an asymmetric ETS configuration between the peers.

### Related Documentation

- [Understanding DCBX on page 6514](#)
- [Configuring DCBX Autonegotiation on page 6524](#)
- [Disabling the ETS Recommendation TLV on page 6527](#)
- [Defining an Application for DCBX Application Protocol TLV Exchange on page 6532](#)
- [Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 6533](#)
- [Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 6534](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 6535](#)

## Defining an Application for DCBX Application Protocol TLV Exchange

Define each application for which you want DCBX to exchange application protocol information. You can define Layer 2 and Layer 4 applications. After you define applications, you map them to IEEE 802.1p code points, and then apply the application map to the interfaces on which you want DCBX to exchange application protocol information with connected peers. (See *Related Documentation* for how to configure application maps and apply them to interfaces, and for an example of the entire procedure that also includes classifier configuration.)



**NOTE:** In Junos OS Release 12.1, the FCoE application was configured by default, so you did not need to configure it in an application map. In Junos OS Release 12.2, if you want DCBX to advertise the FCoE application on an interface and you apply an application map to that interface, you must explicitly configure FCoE in the application map. You also must enable priority-based flow control (PFC) on the FCoE code point on all interfaces that you want to advertise FCoE. If you apply an application map to an interface, the interface sends DCBX TLVs only for the applications configured in the application map.

Define Layer 2 applications by mapping an application name to an EtherType. Define Layer 4 applications by mapping an application name to a protocol (TCP or UDP) and a destination port.

- To define a Layer 2 application, specify the name of the application and its EtherType:

```
[edit applications]
user@switch# set application application-name ether-type ether-type
```

For example, to configure an application named **PTP** (for Precision Time Protocol) that uses the EtherType **0x88F7**:

```
user@switch# set applications application ptp ether-type 0x88F7
```

- To define a Layer 4 application, specify the name of the application, its protocol (TCP or UDP), and its destination port:

```
[edit]
user@switch# set applications application application-name protocol (tcp | udp)
destination-port port-value
```

For example, to configure an application named **iscsi** (for Internet Small Computer System Interface) that uses the protocol **TCP** and the destination port **3260**:

```
user@switch# set applications application iscsi protocol tcp destination-port 3260
```

### Related Documentation

- [Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 6533](#)
- [Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 6534](#)

- [Configuring DCBX Autonegotiation on page 6524](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 6535](#)
- [Example: Configuring DCBX to Support an iSCSI Application](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 6528](#)
- [show dcbx neighbors on page 6581](#)

## Configuring an Application Map for DCBX Application Protocol TLV Exchange

After you define applications for which you want to exchange DCBX application protocol information, map the applications to IEEE 802.1p code points. The IEEE 802.1p code points identify incoming traffic and allow you to map that traffic to the desired application. You then apply the application map to the interfaces on which you want DCBX to exchange application protocol information with connected peers. (See *Related Documentation* for how to define applications and apply the application map to interfaces, and for an example of the entire procedure that also includes classifier configuration.)



**NOTE:** In Junos OS Release 12.1, the FCoE application was configured by default, so you did not need to configure it in an application map. In Junos OS Release 12.2, if you want DCBX to advertise the FCoE application on an interface and you apply an application map to that interface, you must explicitly configure FCoE in the application map. You also must enable priority-based flow control (PFC) on the FCoE code point on all interfaces that you want to advertise FCoE. If you apply an application map to an interface, the interface sends DCBX TLVs only for the applications configured in the application map.

Configure an application map by creating an application map name and mapping an application to one or more IEEE 802.1p code points.

- To define an application map, specify the name of the application map, the name of the application, and the IEEE 802.1p code points of the incoming traffic that you want to associate with the application in the application map:

```
[edit policy-options]
user@switch# set application-maps application-map-name application application-name
code-points [aliases] [bit-patterns]
```

For example, to configure an application map named **ptp-app-map** that includes an application named **PTP** (for Precision Time Protocol) and map the application to IEEE 802.1p code points **001** and **101**:

```
user@switch# set policy-options application-maps ptp-app-map application ptp code points
[001 101]
```

### Related Documentation

- [Defining an Application for DCBX Application Protocol TLV Exchange on page 6532](#)
- [Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 6534](#)

- [Configuring DCBX Autonegotiation on page 6524](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 6535](#)
- [Example: Configuring DCBX to Support an iSCSI Application](#)
- [show dcbx neighbors on page 6581](#)

## Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange

---

After you define applications and map them to IEEE 802.1p code points in an application map, apply the application map to the interfaces on which you want DCBX to exchange the application protocol information with connected peers. (See *Related Documentation* for how to define applications and configure application maps to interfaces, and for an example of the entire procedure that also includes classifier configuration.)



**NOTE:** In Junos OS Release 12.1, the FCoE application was configured by default, so you did not need to configure it in an application map. In Junos OS Release 12.2, if you want DCBX to advertise the FCoE application on an interface and you apply an application map to that interface, you must explicitly configure FCoE in the application map. You also must enable priority-based flow control (PFC) on the FCoE code point on all interfaces that you want to advertise FCoE. If you apply an application map to an interface, the interface sends DCBX TLVs only for the applications configured in the application map.

- To apply an application map to a DCBX interface, specify the DCBX interface and the application map name:

[edit protocols]

```
user@switch# set dcbx interface interface-name application-map application-map-name
```

For example, to apply an application map named **ptp-app-map** on interface **xe-0/0/11**:

```
user@switch# set protocols dcbx interface xe-0/0/11 application-map ptp-app-map
```

### Related Documentation

- [Defining an Application for DCBX Application Protocol TLV Exchange on page 6532](#)
- [Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 6533](#)
- [Configuring DCBX Autonegotiation on page 6524](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 6535](#)
- [Example: Configuring DCBX to Support an iSCSI Application](#)
- [show dcbx neighbors on page 6581](#)

## Example: Configuring DCBX Application Protocol TLV Exchange

Data Center Bridging Capability Exchange protocol (DCBX) discovers the data center bridging (DCB) capabilities of connected peers by exchanging application configuration information. DCBX detects feature misconfiguration and mismatches and can configure DCB on peers. DCBX is an extension of the Link Layer Discovery Protocol (LLDP). LLDP must remain enabled on every interface on which you want to use DCBX.



**NOTE:** LLDP and DCBX are enabled by default on all interfaces.

The switch supports DCBX application protocol exchange for Layer 2 and Layer 4 applications such as the Internet Small Computer System Interface (iSCSI). You specify applications by EtherType (for Layer 2 applications) or by the destination port and protocol (for Layer 4 applications; the protocol can be either TCP or UDP).

The switch handles Fibre Channel over Ethernet (FCoE) application protocol exchange differently than other protocols in some cases:

- If FCoE is the only application for which you want to enable DCBX application protocol TLV exchange on an interface, you do not have to explicitly configure the FCoE application or an application map. By default, the switch exchanges FCoE application protocol TLVs on all interfaces that carry FCoE traffic (traffic mapped to the **fcoe** forwarding class) and have priority-based flow control (PFC) enabled on the FCoE priority (the FCoE IEEE 802.1p code point). The default priority mapping for the FCoE application is IEEE 802.1p code point 011 (the default **fcoe** forwarding class code point).
- If you want an interface to use DCBX to exchange application protocol TLVs for any other applications in addition to FCoE, you must configure the applications (including FCoE), define an application map (including FCoE), and apply the application map to the interface. If you apply an application map to an interface, you must explicitly configure the FCoE application, or the interface does not exchange FCoE application protocol TLVs.

This example shows how to configure interfaces to exchange both Layer 2 and Layer 4 applications by configuring one interface to exchange iSCSI and FCoE application protocol information and configuring another interface to exchange iSCSI and Precision Time Protocol (PTP) application protocol information.

- [Requirements on page 7037](#)
- [Overview on page 7038](#)
- [Configuration on page 7041](#)
- [Verification on page 7043](#)

## Requirements

This example uses the following hardware and software components:

- Juniper Networks QFX Series device

- Junos OS Release 12.1 or later for the QFX Series

## Overview

The switch supports DCBX application protocol exchange for:

- Layer 2 applications, defined by EtherType
- Layer 4 applications, defined by destination port and protocol



**NOTE:** DCBX also advertises PFC and enhanced transmission selection (ETS) information. See [“Configuring DCBX Autonegotiation” on page 6524](#) for how DCBX negotiates and advertises configuration information for these features and for the applications.

DCBX is configured on a per-interface basis for each supported feature or application. For applications that you want to enable for DCBX application protocol exchange, you must:

- Define the application name and configure the EtherType or the destination port and protocol (TCP or UDP) of the application. Use the EtherType for Layer 2 applications, and use the destination port and protocol for Layer 4 protocols.
- Map the application to an IEEE 802.1p code point in an application map.
- Add the application map to DCBX interface.

In addition, for all applications (including FCoE, even when you do not use an application map), you either must create an IEEE 802.1p classifier and apply it to the appropriate ingress interfaces or use the default classifier. A classifier maps the code points of incoming traffic to a forwarding class and a loss priority so that ingress traffic is assigned to the correct class of service (CoS). The forwarding class determines the output queue on the egress interface.

If you do not create classifiers, trunk and tagged-access ports use the unicast IEEE 802.1 default trusted classifier. [Table 519](#) shows the default mapping of IEEE 802.1 code-point values to unicast forwarding classes and loss priorities for ports in trunk mode or tagged-access mode. [Table 520](#) shows the default untrusted classifier IEEE 802.1 code-point values to unicast forwarding class mapping for ports in access mode.

**Table 600: Default IEEE 802.1 Classifiers for Trunk Ports and Tagged-Access Ports (Default Trusted Classifier)**

| Code Point | Forwarding Class | Loss Priority |
|------------|------------------|---------------|
| be (000)   | best-effort      | low           |
| be1 (001)  | best-effort      | low           |
| ef (010)   | best-effort      | low           |

**Table 600: Default IEEE 802.1 Classifiers for Trunk Ports and Tagged-Access Ports (Default Trusted Classifier) (*continued*)**

| Code Point | Forwarding Class | Loss Priority |
|------------|------------------|---------------|
| ef1 (011)  | fcoe             | low           |
| af11 (100) | no-loss          | low           |
| af12 (101) | best-effort      | low           |
| nc1 (110)  | network-control  | low           |
| nc2 (111)  | network-control  | low           |

**Table 601: Default IEEE 802.1 Unicast Classifiers for Access Ports (Default Untrusted Classifier)**

| Code Point | Forwarding Class | Loss Priority |
|------------|------------------|---------------|
| 000        | best-effort      | low           |
| 001        | best-effort      | low           |
| 010        | best-effort      | low           |
| 011        | best-effort      | low           |
| 100        | best-effort      | low           |
| 101        | best-effort      | low           |
| 110        | best-effort      | low           |
| 111        | best-effort      | low           |

### Topology

This example shows how to configure DCBX application protocol exchange for three protocols (iSCSI, PTP, and FCoE) on two interfaces. One interface exchanges iSCSI and FCoE application protocol information, and the other interface exchanges iSCSI and PTP application protocol information.



**NOTE:** You must map FCoE traffic to the interfaces on which you want to forward FCoE traffic. You must also enable PFC on the FCoE interfaces and create an ingress classifier for FCoE traffic, or else use the default classifier.

Table 521 shows the configuration components for this example.

**Table 602: Components of DCBX Application Protocol Exchange Configuration Topology**

| Component                                                         | Settings                                                                                                                                                                                                                                                                                                                                                     |
|-------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Hardware                                                          | QFX Series device                                                                                                                                                                                                                                                                                                                                            |
| LLDP                                                              | Enabled by default on Ethernet interfaces                                                                                                                                                                                                                                                                                                                    |
| DCBX                                                              | Enabled by default on Ethernet interfaces                                                                                                                                                                                                                                                                                                                    |
| iSCSI application (Layer 4)                                       | Application name— <b>iscsi</b><br>protocol— <b>TCP</b><br>destination-port— <b>3260</b><br>code-points— <b>111</b>                                                                                                                                                                                                                                           |
| PTP application (Layer 2)                                         | Application name— <b>ptp</b><br>ether-type— <b>0x88F7</b><br>code-points— <b>001, 101</b>                                                                                                                                                                                                                                                                    |
| FCoE application (Layer 2)                                        | Application name— <b>fcoe</b><br>ether-type— <b>0x8906</b><br>code-points— <b>011</b><br><br><b>NOTE:</b> You explicitly configure the FCoE application because you are applying an application map to the interface. When you apply an application map to an interface, all applications must be explicitly configured and included in the application map. |
| Application maps                                                  | <b>dcbx-iscsi-fcoe-app-map</b> —Maps the iSCSI and FCoE applications to IEEE 802.1p code points<br><br><b>dcbx-iscsi-ptp-app-map</b> —Maps iSCSI and PTP applications to IEEE 802.1p code points                                                                                                                                                             |
| Interfaces                                                        | <b>xe-0/0/10</b> —Configured to exchange FCoE and iSCSI application TLVs (uses application map <b>dcbx-iscsi-fcoe-app-map</b> , carries FCoE traffic, and has PFC enabled on the FCoE priority)<br><br><b>xe-0/0/11</b> —Configured to exchange iSCSI and PTP application TLVs (uses application map <b>dcbx-iscsi-ptp-app-map</b> )                         |
| PFC congestion notification profile for FCoE application exchange | <b>fcoe-cnp:</b> <ul style="list-style-type: none"> <li>Code point—<b>011</b></li> <li>Interface—<b>xe-0/0/10</b></li> </ul>                                                                                                                                                                                                                                 |



Table 602: Components of DCBX Application Protocol Exchange Configuration Topology (*continued*)

| Component                                                                                                         | Settings                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Behavior aggregate classifiers (map forwarding classes to incoming packets by the packet's IEEE 802.1 code point) | <b>fcoe-iscsi-cl1:</b> <ul style="list-style-type: none"> <li>Maps the <b>fcoe</b> forwarding class to the IEEE 802.1p code point used for the FCoE application (011) and a loss priority of <b>high</b></li> <li>Maps the <b>network-control</b> forwarding class to the IEEE 802.1p code point used for the iSCSI application (111) and a loss priority of <b>high</b></li> <li>Applied to interface <b>xe-0/0/10</b></li> </ul> <b>iscsi-ntp-cl2:</b> <ul style="list-style-type: none"> <li>Maps the <b>network-control</b> forwarding class to the IEEE 802.1p code point used for the iSCSI application (111) and a loss priority of <b>low</b></li> <li>Maps the <b>best-effort</b> forwarding class to the IEEE 802.1p code points used for the PTP application (001 and 101) and a loss priority of <b>low</b></li> <li>Applied to interface <b>xe-0/0/11</b></li> </ul> |



**NOTE:** This example does not include scheduling (bandwidth allocation) configuration or lossless configuration for the iSCSI forwarding class.

## Configuration

### CLI Quick Configuration

To quickly configure DCBX application protocol exchange, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```

set applications application iSCSI protocol tcp destination-port 3260
set applications application FCoE ether-type 0x8906
set applications application PTP ether-type 0x88F7
set policy-options application-maps dcbx-iscsi-fcoe-app-map application iSCSI code-points 111
set policy-options application-maps dcbx-iscsi-fcoe-app-map application FCoE code-points 011
set policy-options application-maps dcbx-iscsi-ntp-app-map application iSCSI code-points 111
set policy-options application-maps dcbx-iscsi-ntp-app-map application PTP code-points [001 101]
set protocols dcbx interface xe-0/0/10 application-map dcbx-iscsi-fcoe-app-map
set protocols dcbx interface xe-0/0/11 application-map dcbx-iscsi-ntp-app-map
set class-of-service congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011 pfc
set class-of-service interfaces xe-0/0/10 congestion-notification-profile fcoe-cnp
set class-of-service classifiers ieee-802.1 fcoe-iscsi-cl1 import default forwarding-class fcoe
loss-priority high code-points 011
set class-of-service classifiers ieee-802.1 fcoe-iscsi-cl1 import default forwarding-class
network-control loss-priority high code-points 111
set class-of-service classifiers ieee-802.1 iscsi-ntp-cl2 import default forwarding-class
network-control loss-priority low code-points 111
set class-of-service classifiers ieee-802.1 iscsi-ntp-cl2 import default forwarding-class best-effort
loss-priority low code-points [001 101]
set class-of-service interfaces xe-0/0/10 unit 0 classifiers ieee-802.1 fcoe-iscsi-cl1

```

```
set class-of-service interfaces xe-0/0/11 unit 0 classifiers ieee-802.1 iscsi-ptp-cl2
```

### Configuring DCBX Application Protocol TLV Exchange

#### Step-by-Step Procedure

To define the applications, map the applications to IEEE 802.1p code points, apply the applications to interfaces, and create classifiers for DCBX application protocol exchange:

1. Define the iSCSI application by specifying its protocol and destination port, and define the FCoE and PTP applications by specifying their EtherTypes.  

```
[edit applications]
user@switch# set application iSCSI protocol tcp destination-port 3260
user@switch# set application FCoE ether-type 0x8906
user@switch# set application PTP ether-type 0x88F7
```
2. Define an application map that maps the iSCSI and FCoE applications to IEEE 802.1p code points.  

```
[edit policy-options]
user@switch# set application-maps dcbx-iscsi-fcoe-app-map application iSCSI code-points 111
user@switch# set application-maps dcbx-iscsi-fcoe-app-map application FCoE code-points 011
```
3. Define the application map that maps the iSCSI and PTP applications to IEEE 802.1p code points.  

```
[edit policy-options]
user@switch# set application-maps dcbx-iscsi-ptp-app-map application iSCSI code-points 111
user@switch# set application-maps dcbx-iscsi-ptp-app-map application PTP code-points [001 101]
```
4. Apply the iSCSI and FCoE application map to interface **xe-0/0/10**, and apply the iSCSI and PTP application map to interface **xe-0/0/11**.  

```
[edit protocols dcbx]
user@switch# set interface xe-0/0/10 application-map dcbx-iscsi-fcoe-app-map
user@switch# set interface xe-0/0/11 application-map dcbx-iscsi-ptp-app-map
```
5. Create the congestion notification profile to enable PFC on the FCoE code point (011), and apply the congestion notification profile to interface **xe-0/0/10**.  

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011 pfc
user@switch# set interfaces xe-0/0/10 congestion-notification-profile fcoe-cnp
```
6. Configure the classifier to apply to the interface that exchanges iSCSI and FCoE application information.  

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 fcoe-iscsi-cl1 import default forwarding-class fcoe loss-priority high code-points 011
user@switch# set ieee-802.1 fcoe-iscsi-cl1 import default forwarding-class network-control loss-priority high code-points 111
```
7. Configure the classifier to apply to the interface that exchanges iSCSI and PTP application information.  

```
[edit class-of-service classifiers]
```

```

user@switch# set ieee-802.1 iscsi-ptp-cl2 import default forwarding-class network-control
loss-priority low code-points 111
user@switch# set ieee-802.1 iscsi-ptp-cl2 import default forwarding-class best-effort
loss-priority low code-points [001 101]

```

8. Apply the classifiers to the appropriate interfaces.

```

[edit class-of-service]
user@switch# set interfaces xe-0/0/10 unit 0 classifiers ieee-802.1 fcoe-iscsi-cl1
user@switch# set interfaces xe-0/0/11 unit 0 classifiers ieee-802.1 iscsi-ptp-cl2

```

## Verification

To verify that DCBX application protocol exchange configuration has been created and is operating properly, perform these tasks:

- [Verifying the Application Configuration on page 7043](#)
- [Verifying the Application Map Configuration on page 7043](#)
- [Verifying DCBX Application Protocol Exchange Interface Configuration on page 7044](#)
- [Verifying the PFC Configuration on page 7044](#)
- [Verifying the Classifier Configuration on page 7045](#)

### Verifying the Application Configuration

**Purpose** Verify that DCBX applications have been configured.

**Action** List the applications by using the configuration mode command **show applications**:

```

user@switch# show applications
application iSCSI {
 protocol tcp;
 destination-port 3260;
}

application fcoe {
 ether-type 0x8906;
}

application ptp {
 ether-type 0x88F7;
}

```

**Meaning** The **show applications** configuration mode command lists all of the configured applications and either their protocol and destination port (Layer 4 applications) or their EtherType (Layer 2 applications). The command output shows that the iSCSI application is configured with the **tcp** protocol and destination port **3260**, the FCoE application is configured with the EtherType **0x8906**, and that the PTP application is configured with the EtherType **0x88F7**.

### Verifying the Application Map Configuration

**Purpose** Verify that the application maps have been configured.

**Action** List the application maps by using the configuration mode command **show policy-options application-maps**:

```
user@switch# show policy-options application-maps
dcbx-iscsi-fcoe-app-map {
 application iSCSI code-points 111;
 application FCoE code-points 011;
}

dcbx-iscsi-ptp-app-map {
 application iSCSI code-points 111;
 application PTP code-points [001 101];
}
```

**Meaning** The **show policy-options application-maps** configuration mode command lists all of the configured application maps and the applications that belong to each application map. The command output shows that there are two application maps, **dcbx-iscsi-fcoe-app-map** and **dcbx-iscsi-ptp-app-map**.

The application map **dcbx-iscsi-fcoe-app-map** consists of the iSCSI application, which is mapped to IEEE 802.1p code point 111, and the FCoE application, which is mapped to IEEE 802.1p code point 011.

The application map **dcbx-iscsi-ptp-app-map** consists of the iSCSI application, which is mapped to IEEE 802.1p code point 111, and the PTP application, which is mapped to IEEE 802.1p code points 001 and 101.

---

### Verifying DCBX Application Protocol Exchange Interface Configuration

**Purpose** Verify that the application maps have been applied to the correct interfaces.

**Action** List the application maps by using the configuration mode command **show protocols dcbx**:

```
user@switch# show protocols dcbx
interface xe-0/0/10.0 {
 application-map dcbx-iscsi-fcoe-app-map;
}

interface xe-0/0/11.0 {
 application-map dcbx-iscsi-ptp-app-map;
}
```

**Meaning** The **show protocols dcbx** configuration mode command lists whether the interfaces are enabled for DCBX and lists the application map applied to each interface. The command output shows that interfaces **xe-0/0/10.0** and **xe-0/0/11.0** are enabled for DCBX, and that interface **xe-0/0/10.0** uses application map **dcbx-iscsi-fcoe-app-map**, and interface **xe-0/0/11.0** uses application map **dcbx-iscsi-ptp-app-map**.

---

### Verifying the PFC Configuration

**Purpose** Verify that PFC has been enabled on the FCoE code point and applied to the correct interface.

**Action** Display the PFC configuration to verify that PFC is enabled on the FCoE code point (011) in the congestion notification profile **fcoe-cnp** by using the configuration mode command **show class-of-service congestion-notification-profile**:

```
user@switch# show class-of-service congestion-notification-profile
fcoe-cnp {
 input {
 ieee-802.1 {
 code-point 011 {
 pfc;
 }
 }
 }
}
```

Display the class-of-service (CoS) interface information to verify that the correct interface has PFC enabled for the FCoE application by using the configuration mode command **show class-of-service interfaces**:

```
user@switch# show class-of-service interfaces
xe-0/0/10 {
 congestion-notification-profile fcoe-cnp;
}
```



**NOTE:** The sample output does not include all of the information this command can show. The output is abbreviated to focus on verifying the PFC configuration.

**Meaning** The **show class-of-service congestion-notification-profile** configuration mode command lists the configured congestion notification profiles. The command output shows that the congestion notification profile **fcoe-cnp** has been configured and has enabled PFC on the IEEE 802.1p code point **011** (the default FCoE code point).

The **show class-of-service interfaces** configuration mode command shows the interface CoS configuration. The command output shows that the congestion notification profile **fcoe-cnp**, which enables PFC on the FCoE code point, is applied to interface **xe-0/0/10**.

### Verifying the Classifier Configuration

**Purpose** Verify that the classifiers have been configured and applied to the correct interfaces.

**Action** Display the classifier configuration by using the configuration mode command **show class-of-service**:

```
user@switch# show class-of-service
classifiers {
 ieee-802.1 fcoe-iscsi-cl1 {
 import default;
 forwarding-class network-control {
 loss-priority high code-points 111;
 }
 forwarding-class fcoe {
```

```

 loss-priority high code-points 011;
 }
}
ieee-802.1 iscsi-ntp-cl2 {
 import default;
 forwarding-class network-control {
 loss-priority low code-points 111;
 }
 forwarding-class best-effort {
 loss-priority low code-points [001 101];
 }
}
}
interfaces {
 xe-0/0/10 {
 congestion-notification-profile fcoe-cnp;
 unit 0 {
 classifiers {
 ieee-802.1 fcoe-iscsi-cl1;
 }
 }
 }
 xe-0/0/11 {
 unit 0 {
 classifiers {
 ieee-802.1 iscsi-ntp-cl2;
 }
 }
 }
}
}

```



**NOTE:** The sample output does not include all of the information this command can show. The output is abbreviated to focus on verifying the classifier configuration.

**Meaning** The **show class-of-service** configuration mode command lists the classifier and CoS interface configuration, as well as other information not shown in this example. The command output shows that there are two classifiers configured, **fcoe-iscsi-cl1** and **iscsi-ntp-cl2**.

Classifier **fcoe-iscsi-cl1** uses the **default** classifier as a template and edits the template as follows:

- The forwarding class **network-control** is set to a loss priority of **high** and is mapped to code point **111** (the code point mapped to the iSCSI application).
- The forwarding class **fcoe** is set to a loss priority of **high** and is mapped to code point **011** (the code point mapped by default to the FCoE application).

Classifier **iscsi-ntp-cl2** uses the **default** classifier as a template and edits the template as follows:

- The forwarding class **network-control** is set to a loss priority of **low** and is mapped to IEEE 802.1p code point **111** (the code point mapped to the iSCSI application).
- The forwarding class **best-effort** is set to a loss priority of **low** and is mapped to IEEE 802.1p code points **001** and **101** (the code points mapped by default to the PTP application).

The command output also shows that classifier **fcoe-iscsi-cl1** is mapped to interface **xe-0/0/10.0** and that classifier **iscsi-ptp-cl2** is mapped to interface **xe-0/0/11.0**.

#### Related Documentation

- [Defining an Application for DCBX Application Protocol TLV Exchange on page 6532](#)
- [Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 6533](#)
- [Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 6534](#)
- [Configuring DCBX Autonegotiation on page 6524](#)
- [show dcbx on page 6580](#)
- [show dcbx neighbors on page 6581](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 6528](#)





# Learn About Technology

- [Data Center Technology Overview Videos on page 7049](#)

## Data Center Technology Overview Videos

---

Juniper Information Experience (IX) videos provide brief, high-level overviews of data center technologies and concepts. Each video runs approximately one-and-a-half to two minutes in length. This document contains SDN-related videos and links to conceptual documents that contain other data center technology videos:

- [Learn About Video: Why Do We Need an IP Fabric? on page 7049](#)
- [Learn About Video: What is the Best Control Plane Protocol to Use in a Data Center IP Fabric? on page 7049](#)
- [Learn About Video: Why Use an Overlay Network in a Data Center? on page 7049](#)
- [Conceptual Documents That Contain Technology Overview Videos on page 7050](#)

### Learn About Video: Why Do We Need an IP Fabric?

The video *Why Do We Need an IP Fabric?* presents a brief overview of IP Fabric use cases.



Video: [Why Do We Need an IP Fabric?](#)

### Learn About Video: What is the Best Control Plane Protocol to Use in a Data Center IP Fabric?

The video *What is the Best Control Plane Protocol to Use in a Data Center IP Fabric?* presents a brief overview of the arguments for using Border Gateway Protocol (BGP) as the data center IP fabric control plane protocol.



Video: [What is the Best Control Plane Protocol to Use in a Data Center IP Fabric?](#)

### Learn About Video: Why Use an Overlay Network in a Data Center?

The video *Why Use an Overlay Network in a Data Center?* presents a brief overview of the advantages of data center overlay networks.



Video: [Why Use an Overlay Network in a Data Center?](#)

## Conceptual Documents That Contain Technology Overview Videos

The following conceptual documents include brief video overviews of the technology:

- [Understanding DCB Features and Requirements on page 6490](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 6765](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 6493](#)
- [Understanding DCBX on page 6514](#)
- [Understanding PFC Functionality Across Layer 3 Interfaces on page 6921](#)
- [Virtual Chassis Fabric Overview](#)
- [Understanding In-Service Software Upgrade \(ISSU\) and In-Service Software Upgrade \(ISSU\) System Requirements \(same video\)](#)

# Configuration Statements and Operational Commands

- [Configuration Statements \(Basic Concepts\) on page 7053](#)
- [Configuration Statements \(Classifiers and Rewrite Rules\) on page 7061](#)
- [Configuration Statements \(Scheduling\) on page 7091](#)
- [Configuration Statements \(Data Center Bridging and Flow Control\) on page 7121](#)
- [Operational Commands \(Basic Concepts\) on page 7161](#)
- [Operational Commands \(Classifiers and Rewrite Rules\) on page 7231](#)
- [Operational Commands \(Scheduling\) on page 7291](#)
- [Operational Commands \(Data Center Bridging and Flow Control\) on page 7393](#)



## CHAPTER 257

# Configuration Statements (Basic Concepts)

- [class-of-service](#) on page 7054
- [traceoptions \(Class of Service\)](#) on page 7058

## class-of-service

```

Syntax class-of-service {
 classifiers {
 (dscp | dscp-ipv6 | ieee-802.1 | exp) classifier-name {
 import (classifier-name | default);
 forwarding-class class-name {
 loss-priority level {
 code-points [aliases] [bit-patterns];
 }
 }
 }
 }
 code-point-aliases {
 (dscp | dscp-ipv6 | ieee-802.1) {
 alias-name bits;
 }
 }
 congestion-notification-profile profile-name {
 input {
 ieee-802.1 {
 code-point [code-point-bits] {
 pfc {
 mru mru-value;
 }
 }
 }
 cable-length cable-length-value;
 }
 output {
 ieee-802.1 {
 code-point [code-point-bits] {
 flow-control-queue [queue | list-of-queues];
 }
 }
 }
 }
 drop-profiles {
 profile-name {
 interpolate {
 fill-level low-value fill-level high-value drop-probability 0 drop-probability high-value;
 }
 }
 }
 forwarding-class class-name {
 loss-priority level {
 code-points [aliases] [bit-patterns];
 }
 }
 forwarding-class class-name {
 scheduler scheduler-name;
 }
 forwarding-class-sets forwarding-class-set-name {
 class class-name;
 }
 }

```

```

}
forwarding-classes {
 class {
 class-name {
 queue-num queue-number <no-loss>;
 }
 }
}
host-outbound-traffic {
 forwarding-class class-name;
 dscp-code-point code-point;
}
interfaces {
 interface-name {
 congestion-notification-profile profile-name {
 }
 forwarding-class lossless-forwarding-class-name;
 forwarding-class-set forwarding-class-set-name {
 output-traffic-control-profile profile-name;
 }
 rewrite-value {
 input {
 ieee-802.1 {
 code-point code-point-bits;
 }
 }
 }
 scheduler-map scheduler-map-name
 unit logical-unit-number {
 classifiers {
 (dscp | dscp-ipv6 | ieee-802.1 | exp) (classifier-name | default);
 }
 forwarding-class class-name;
 rewrite-rules {
 (dscp | dscp-ipv6 | ieee-802.1 | exp) (classifier-name | default);
 }
 }
 }
}
multi-destination {
 classifiers {
 (dscp | ieee-802.1) classifier-name;
 }
}
rewrite-rules {
 (dscp | dscp-ipv6 | ieee-802.1 | exp) classifier-name {
 import (rewrite-name | default);
 forwarding-class class-name {
 loss-priority priority code-point (alias | bits);
 }
 }
}
scheduler-map-forwarding-class-sets {
 fabric-scheduler-map-name {
 forwarding-class-set fabric-forwarding-class-set-name scheduler scheduler-name;
 }
}

```

```
}
scheduler-maps {
 map-name {
 forwarding-class class-name scheduler scheduler-name;
 }
}
schedulers {
 scheduler-name {
 buffer-size (percent percentage | remainder);
 drop-profile-map loss-priority (low | medium-high | high) protocol protocol drop-profile
 drop-profile-name;
 excess-rate percent percentage;
 explicit-congestion-notification;
 priority priority;
 shaping-rate (rate | percent percentage);
 transmit-rate (percent percentage) <exact>;
 }
}
shared-buffer {
 egress {
 percent percent;
 buffer-partition (lossless | lossy | multicast) {
 percent percent
 }
 }
 ingress {
 percent percent;
 buffer-partition (lossless | lossless-headroom | lossy) {
 percent percent
 }
 }
}
system-defaults {
 classifiers exp classifier-name;
}
traffic-control-profiles profile-name {
 guaranteed-rate(rate| percent percentage);
 scheduler-map map-name;
 shaping-rate (rate| percent percentage);
}
}
```

Hierarchy Level [\[edit\]](#)



**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.



**NOTE:** Not all switches support all portions of the class of service hierarchy. For example, some switches use the same classifiers for unicast and multidestination traffic, and those switches do not support the multi-destination classifier hierarchy, and some switches do not support shared buffer configuration, and those switches do not support the shared-buffer hierarchy.



**NOTE:** OCX Series switches do not support MPLS exp classifiers and rewrite rules (including MPLS system defaults), and they do not support congestion notification profiles.

**Description** Configure class-of-service parameters on the switch.

The remaining statements are explained separately.

**Default** If you do not configure any CoS features, the default CoS settings are used.

**Required Privilege Level** interfaces—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Assigning CoS Components to Interfaces on page 6627](#)
- [Overview of Junos OS CoS on page 6608](#)

## traceoptions (Class of Service)

---

|                            |                                                                                                                                                                                                                                       |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>traceoptions {<br/>  file <i>filename</i> &lt;size <i>size</i>&gt; &lt;files <i>number</i>&gt;<br/>  &lt;world-readable   no-world-readable&gt;;<br/>  flag <i>flag</i> &lt;flag-modifier&gt;;<br/>  no-remote-trace<br/>}</pre> |
| <b>Hierarchy Level</b>     | [edit <a href="#">class-of-service</a> ]                                                                                                                                                                                              |
| <b>Release Information</b> | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                         |
| <b>Description</b>         | Set class-of-service (CoS) tracing options.                                                                                                                                                                                           |



**NOTE:** The `traceoptions` statement is not supported on the QFabric system.

**Default** Traceoptions is disabled.

**Options** **file *filename***—Name of the file to receive the tracing operation output. Enclose the name in quotation marks. Traceoption output files are located in the `/var/log/` directory.

**files *number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file.0***. The traceoption output continues in a second trace file named ***trace-file.1***. When ***trace-file.1*** reaches its maximum size, output continues in a third file named ***trace-file.2***, and so on. When the maximum number of trace files is reached, the oldest trace file is overwritten.

If you specify a maximum number of files, you must also specify a maximum file size with the size option.

**Range:** 2 through 1000 files

**Default:** 1 trace file

**flag**—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements:

- **all**—Trace all operations.
- **asynch**—Trace asynchronous configuration processing.
- **chassis-scheduler**—Trace chassis stream scheduler processing.
- **cos-adjustment**—Trace CoS rate adjustments.
- **dynamic**—Trace dynamic CoS functions.
- **hardware-database**—Trace the chassis hardware database related processing.

- **init**—Trace initialization events.
- **performance-monitor**—Trace performance monitor counters.
- **process**—Trace configuration processing.
- **restart**—Trace restart processing.
- **route-socket**—Trace route-socket events.
- **show**—Trace show command servicing.
- **snmp**—Trace SNMP-related processing.
- **util**—Trace utilities.

The following are the global tracing options:

- **all**—Perform all tracing operations
- **parse**—Trace parser processing.

**no-remote-trace**—(Optional) Disable remote tracing.

**no-world-readable**—(Optional) Prevent any user from reading the log file.

**size *size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**. Incoming tracefile data is logged in the now empty **trace-file**. When **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option.

**Syntax:** *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

**Range:** 10 KB through the maximum file size of 4 GB (maximum is lower if 4 GB is not supported on your system)

**Default:** 1 MB

**world-readable**—(Optional) Allow any user to read the log file.



|                                 |                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration. |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------|



# Configuration Statements (Classifiers and Rewrite Rules)

- [class \(Forwarding Classes\) on page 7062](#)
- [class \(Forwarding Class Sets\) on page 7063](#)
- [classifiers on page 7064](#)
- [code-point \(Rewrite Rules\) on page 7066](#)
- [code-point-aliases on page 7067](#)
- [code-points \(CoS\) on page 7068](#)
- [dscp on page 7069](#)
- [dscp-ipv6 on page 7071](#)
- [exp on page 7073](#)
- [forwarding-class on page 7075](#)
- [forwarding-class-sets on page 7076](#)
- [forwarding-classes on page 7077](#)
- [ieee-802.1 on page 7079](#)
- [import on page 7081](#)
- [interfaces \(Class of Service\) on page 7082](#)
- [loss-priority \(Classifiers\) on page 7084](#)
- [loss-priority \(Rewrite Rules\) on page 7085](#)
- [queue-num on page 7086](#)
- [rewrite-rules on page 7088](#)
- [unit on page 7089](#)

## class (Forwarding Classes)

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |                                                                                                                                                                                                                                                                                                                                               |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <pre>class {   class-name {     queue-num queue-number &lt;no-loss&gt;;   } }</pre>                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | [edit <a href="#">class-of-service forwarding-classes</a> ]                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>No-loss option introduced in Junos OS Release 12.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                        |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <p>Map one or more forwarding classes to a single queue. You can map unicast forwarding classes to a unicast queue (0 through 7) and multdestination forwarding classes to a multicast queue (8 through 11). The queue to which you map a forwarding class determines if the forwarding class is a unicast or multicast forwarding class.</p> |
| <p> <b>NOTE:</b> On systems that do not use the ELS CLI, if you are using Junos OS Release 12.2, use the default forwarding-class-to-queue mapping for the lossless fcoe and no-loss forwarding classes. If you explicitly configure the lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best effort) traffic and does <i>not</i> receive lossless treatment.</p>                                                                                                                                                                                                                                                                                                                                                          |                                                                                                                                                                                                                                                                                                                                               |
| <p> <b>NOTE:</b> On systems that do not use the ELS CLI, if you are using Junos OS Release 12.3 or later, the default configuration is the same as the default configuration for Junos OS Release 12.2, and the default behavior is the same (the fcoe and no-loss forwarding classes receive lossless treatment). However, if you explicitly configure lossless forwarding classes, you can configure up to six lossless forwarding classes by specifying the no-loss option. If you do not specify the no-loss option in an explicit forwarding class configuration, the forwarding class is lossy. For example, if you explicitly configure the fcoe forwarding class and you do not include the no-loss option, the fcoe forwarding class is lossy, not lossless.</p> |                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <p><b>class-name</b> —Name of the forwarding class.</p> <p>The remaining statement is explained separately.</p>                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                           |

- Related Documentation**
- [Example: Configuring Forwarding Classes](#)
  - [Example: Configuring Forwarding Classes on page 6692](#)
  - [Understanding CoS Forwarding Classes](#)
  - [Understanding CoS Forwarding Classes on page 6688](#)
  - [Understanding CoS Forwarding Classes](#)

## class (Forwarding Class Sets)

---

|                                 |                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>class <i>class-name</i>;</code>                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service forwarding-class-sets</a> <i>forwarding-class-set-name</i> ]                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                     |
| <b>Description</b>              | Group forwarding classes into sets of forwarding classes (priority groups). You can group some or all of the configured forwarding classes into up to three unicast forwarding class sets and one multidestination forwarding class set.                                                                                          |
| <b>Options</b>                  | <i>class-name</i> —Name of the forwarding class.                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | interfaces—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 6771</a></li> <li>• <a href="#">Example: Configuring Forwarding Class Sets on page 6697</a></li> <li>• <a href="#">Understanding CoS Forwarding Class Sets (Priority Groups) on page 6695</a></li> </ul> |

## classifiers

|                                                                                  |                                                                                                                                                                                                                                                                                                                                     |
|----------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                                                            | <a href="#">Syntax (BA Classifiers) on page 7064</a><br><a href="#">Syntax (Multidestination BA Classifiers) on page 7064</a><br><a href="#">Syntax (Interface Classifier Association: DSCP, DSCP IPv6, IEEE) on page 7064</a><br><a href="#">Syntax (Global EXP Interface Classifier Association with Interfaces) on page 7064</a> |
| <b>Syntax (BA Classifiers)</b>                                                   | <pre> classifiers {   (dscp   dscp-ipv6   ieee-802.1) classifier-name {     import (classifier-name   default);     forwarding-class class-name {       loss-priority level {         code-points [ aliases ] [ bit-patterns ];       }     }   } } </pre>                                                                          |
| <b>Syntax (Multidestination BA Classifiers)</b>                                  | <pre> classifiers {   (dscp   ieee-802.1) classifier-name; } </pre>                                                                                                                                                                                                                                                                 |
| <b>Syntax (Interface Classifier Association: DSCP, DSCP IPv6, IEEE)</b>          | <pre> classifiers {   (dscp   dscp-ipv6   ieee-802.1) (default   classifier-name); } </pre>                                                                                                                                                                                                                                         |
| <b>Syntax (Global EXP Interface Classifier Association with Interfaces)</b>      | <pre> classifiers {   exp classifier-name; } </pre>                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level (BA Classifiers)</b>                                          | [edit <a href="#">class-of-service</a> ],                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level (Multidestination BA Classifiers)</b>                         | [edit <a href="#">class-of-service</a> multi-destination],                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level (Interface Classifier Association: DSCP, DSCP IPv6, IEEE)</b> | [edit <a href="#">class-of-service</a> <a href="#">interfaces</a> <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> ]                                                                                                                                                                                           |
| <b>Hierarchy Level (Global EXP Classifier)</b>                                   | [edit <a href="#">class-of-service</a> system-defaults]                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>                                                       | <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>EXP statement introduced in Junos OS Release 12.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                               |
| <b>Description</b>                                                               | Define a unicast or multidestination CoS behavior aggregate (BA) classifier.                                                                                                                                                                                                                                                        |





**NOTE:** OCX Series switches do not support MPLS, so they do not support EXP classifier configuration.

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | The statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Defining CoS Unicast BA Classifiers (DSCP, DSCP IPv6, IEEE 802.1p)</i></li> <li>• <a href="#">Defining CoS BA Classifiers (DSCP, DSCP IPv6, IEEE 802.1p) on page 6656</a></li> <li>• <a href="#">Configuring a Global MPLS EXP Classifier on page 4977</a></li> <li>• <i>Example: Configuring Unicast Classifiers</i></li> <li>• <i>Example: Configuring Multidestination (Multicast, Broadcast, DLF) Classifiers</i></li> <li>• <i>Understanding CoS Classifiers</i></li> <li>• <a href="#">Understanding CoS Classifiers on page 6650</a></li> <li>• <i>Understanding CoS Classifiers</i></li> <li>• <a href="#">Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4959</a></li> </ul> |

## code-point (Rewrite Rules)

---

|                            |                                                                                                                                                                                                            |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>code-point [ <i>alias</i> ] [ <i>bit-pattern</i> ];</code>                                                                                                                                           |
| <b>Hierarchy Level</b>     | [edit <a href="#">class-of-service rewrite-rules</a> ( <a href="#">dscp</a>   <a href="#">ieee-802.1</a> ) <a href="#">forwarding-class</a> <i>class-name</i> <a href="#">loss-priority</a> <i>level</i> ] |
| <b>Release Information</b> | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                              |
| <b>Description</b>         | Configure a code-point alias or bit set to apply to a forwarding class for a rewrite rule.                                                                                                                 |




**NOTE:** OCX Series switches do not support MPLS, so they do not support EXP rewrite rules.

---

|                                 |                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | <i>alias</i> —Name of the alias.<br><br><i>bit-pattern</i> —Value of the code-point bits, in decimal form.                                                                                                                |
| <b>Required Privilege Level</b> | <i>interfaces</i> —To view this statement in the configuration.<br><i>interface-control</i> —To add this statement to the configuration.                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Defining CoS Rewrite Rules on page 6704</a></li><li>• <a href="#">Understanding CoS Classifiers</a></li><li>• <a href="#">Understanding CoS Classifiers</a></li></ul> |

## code-point-aliases

|                                                                                                                                                                                                            |                                                                                                                                                                                                                                               |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                              | <pre>code-point-aliases {   (dscp  dscp-ipv6   ieee-802.1   exp) {     alias-name bits;   } }</pre>                                                                                                                                           |
| <b>Hierarchy Level</b>                                                                                                                                                                                     | [edit <a href="#">class-of-service</a> ]                                                                                                                                                                                                      |
| <b>Release Information</b>                                                                                                                                                                                 | <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                      |
| <b>Description</b>                                                                                                                                                                                         | Define an alias for a CoS marker. You can use the alias instead of the bit pattern when you specify the code point during configuration.                                                                                                      |
| <div>  <p><b>NOTE:</b> OCX Series switches do not support MPLS, so they do not support EXP code-point aliases.</p> </div> |                                                                                                                                                                                                                                               |
| <b>Options</b>                                                                                                                                                                                             | <p><b>(dscp   dscp-ipv6   ieee-802.1   exp)</b>—Set the type of classifier for which you are creating an alias.</p> <p><b>alias-name</b>—Name of the code-point alias.</p> <p><b>bits</b> —Value of the code-point bits, in decimal form.</p> |
| <b>Required Privilege Level</b>                                                                                                                                                                            | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                           |
| <b>Related Documentation</b>                                                                                                                                                                               | <ul style="list-style-type: none"> <li>• <a href="#">Defining CoS Code-Point Aliases on page 6687</a></li> <li>• <a href="#">Understanding CoS Code-Point Aliases on page 6685</a></li> </ul>                                                 |

## code-points (CoS)

---

|                            |                                                                                                                                               |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>code-points [ <i>aliases</i> ] [ <i>bit-patterns</i> ];</code>                                                                          |
| <b>Hierarchy Level</b>     | [edit class-of-service classifiers (dscp   ieee-802.1) <i>classifier-name</i> forwarding-class <i>class-name</i> loss-priority <i>level</i> ] |
| <b>Release Information</b> | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| <b>Description</b>         | Configure one or more code-point aliases or bit sets to apply to a forwarding class.                                                          |



**NOTE:** OCX Series switches do not support MPLS, so they do not support EXP code points or code point aliases.

---

|                                 |                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | <i>aliases</i> —Name of the alias or aliases.<br><br><i>bit-patterns</i> —Value of the code-point bits, in decimal form.                                           |
| <b>Required Privilege Level</b> | interfaces—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Understanding Interfaces</i></li><li>• <i>Example: Configuring BA Classifiers on Transparent Mode Devices</i></li></ul> |

## dscp

|                                                    |                                                                                                                                                                                                                                                                                                                 |
|----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| List of Syntax                                     | <a href="#">Syntax (Classifier) on page 7069</a><br><a href="#">Syntax (Code-Point Alias) on page 7069</a><br><a href="#">Syntax (Multidestination Classifier) on page 7069</a><br><a href="#">Syntax (Interface Classifier Association) on page 7069</a><br><a href="#">Syntax (Rewrite Rule) on page 7069</a> |
| Syntax (Classifier)                                | <pre>dscp classifier-name {   import (classifier-name   default);   forwarding-class class-name {     loss-priority level {       code-points [ aliases ] [ bit-patterns ];     }   } }</pre>                                                                                                                   |
| Syntax (Code-Point Alias Configuration)            | <pre>dscp alias-name bit-pattern;</pre>                                                                                                                                                                                                                                                                         |
| Syntax (Multidestination Classifier Configuration) | <pre>dscp classifier-name;</pre>                                                                                                                                                                                                                                                                                |
| Syntax (Interface Classifier Association)          | <pre>dscp (classifier-name   default);</pre>                                                                                                                                                                                                                                                                    |
| Syntax (Rewrite Rule Configuration)                | <pre>dscp rewrite-name {   import (rewrite-name   default);   forwarding-class class-name {     loss-priority level {       code-point [ aliases ] [ bit-patterns ];     }   } }</pre>                                                                                                                          |
| Hierarchy Level (Classifier)                       | [edit <a href="#">class-of-service classifiers</a> ],                                                                                                                                                                                                                                                           |
| Hierarchy Level (Code-Point Aliases)               | [edit <a href="#">class-of-service code-point-aliases</a> ],                                                                                                                                                                                                                                                    |
| Hierarchy Level (Multidestination Classifier)      | [edit <a href="#">class-of-service multi-destination classifiers</a> ],                                                                                                                                                                                                                                         |
| Hierarchy Level (Interface Classifier Association) | [edit <a href="#">class-of-service interfaces interface-name unit logical-unit-number classifiers</a> ],<br>[edit <a href="#">class-of-service interfaces interface-name unit logical-unit-number rewrite-rules</a> ],                                                                                          |
| Hierarchy Level (Rewrite Rule)                     | [edit <a href="#">class-of-service rewrite-rules</a> ]                                                                                                                                                                                                                                                          |
| Release Information                                | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                   |

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b>              | Define the Differentiated Services code point (DSCP) mapping that is applied to the packets.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <p><i>classifier-name</i>—Name of the classifier.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Unicast Classifiers</i></li><li>• <a href="#">Example: Configuring Classifiers on page 6657</a></li><li>• <a href="#">Defining CoS Code-Point Aliases on page 6687</a></li><li>• <a href="#">Defining CoS Rewrite Rules on page 6704</a></li><li>• <a href="#">Assigning CoS Components to Interfaces on page 6627</a></li><li>• <i>Understanding CoS Classifiers</i></li><li>• <a href="#">Understanding CoS Classifiers on page 6650</a></li><li>• <i>Understanding CoS Classifiers</i></li><li>• <a href="#">Understanding CoS Rewrite Rules on page 6701</a></li><li>• <a href="#">Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 6672</a></li><li>• <i>Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces</i></li></ul> |

## dscp-ipv6

|                                                     |                                                                                                                                                                                                                                            |
|-----------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax (Classifier) on page 7071</a><br><a href="#">Syntax (Code-Point Alias) on page 7071</a><br><a href="#">Syntax (Interface Classifier Association) on page 7071</a><br><a href="#">Syntax (Rewrite Rule) on page 7071</a> |
| <b>Syntax (Classifier)</b>                          | <pre>dscp-ipv6 classifier-name {   import (classifier-name   default);   forwarding-class class-name {     loss-priority level {       code-points [ aliases ] [ bit-patterns ];     }   } }</pre>                                         |
| <b>Syntax (Code-Point Alias)</b>                    | <pre>dscp-ipv6 alias-name bit-pattern;</pre>                                                                                                                                                                                               |
| <b>Syntax (Interface Classifier Association)</b>    | <pre>dscp-ipv6 (classifier-name   default);</pre>                                                                                                                                                                                          |
| <b>Syntax (Rewrite Rule)</b>                        | <pre>dscp-ipv6 rewrite-name {   import (rewrite-name   default);   forwarding-class class-name {     loss-priority level {       code-point [ aliases ] [ bit-patterns ];     }   } }</pre>                                                |
| <b>Hierarchy (Classifier)</b>                       | [edit <a href="#">class-of-service classifiers</a> ],                                                                                                                                                                                      |
| <b>Hierarchy (Code-Point Alias)</b>                 | [edit <a href="#">class-of-service code-point-aliases</a> ],                                                                                                                                                                               |
| <b>Hierarchy (Interface Classifier Association)</b> | [edit <a href="#">class-of-service interfaces interface-name unit logical-unit-number classifiers</a> ],<br>[edit <a href="#">class-of-service interfaces interface-name unit logical-unit-number rewrite-rules</a> ],                     |
| <b>Hierarchy (Rewrite Rule)</b>                     | [edit <a href="#">class-of-service rewrite-rules</a> ]                                                                                                                                                                                     |
| <b>Release Information</b>                          | Statement introduced in Junos OS Release 12.2 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series                                                                                               |
| <b>Description</b>                                  | Define the Differentiated Services code point (DSCP) IPv6 mapping that is applied to the packets.                                                                                                                                          |



**NOTE:** On switches that use different classifiers for unicast and multideestination (multicast, broadcast, and destination lookup fail) traffic, there is no DSCP IPv6 classifier for multideestination (multicast, broadcast,

and destination lookup fail) traffic. Multidestination IPv6 traffic uses the multidestination DSCP classifier.

.....

|                                 |                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | The statements are explained separately.                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | interfaces—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Defining CoS Code-Point Aliases on page 6687</a></li><li>• <a href="#">Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 6672</a></li><li>• <i>Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces</i></li></ul> |



## exp

|                                               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-----------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                                        | <pre>exp classifier-name {     import (classifier-name   default);     forwarding-class class-name {         loss-priority level {             code-points [ aliases ] [ bit-patterns ];         }     } }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Rewrite Rule Configuration                    | <pre>exp rewrite-name {     import (rewrite-name   default);     forwarding-class class-name {         loss-priority level {             code-point [ aliases ] [ bit-patterns ];         }     } }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Global Classifier Association with Interfaces | exp classifier-name;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Hierarchy Level                               | <pre>[edit class-of-service classifiers] [edit class-of-service rewrite-rules] [edit class-of-service system-defaults classifiers]</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Release Information                           | Statement introduced in Junos OS Release 12.3X50 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Description                                   | <p>Define the EXP code point mapping that is applied to MPLS packets. EXP classifiers are not applied to any traffic except MPLS traffic. EXP classifiers are applied only to interfaces that are configured as <b>family mpls</b> (for example, <b>set interfaces xe-0/0/35 unit 0 family mpls</b>.)</p> <p>There are no default EXP classifiers. You can configure up to 64 EXP classifiers.</p> <p>On QFX10000 switches, you can configure and apply EXP classifiers to interfaces in the same way that you configure and apply DSCP, DSCP IPv6, and IEEE classifiers to interfaces. Different interfaces can have different EXP classifiers. QFX10000 switches do not support global EXP classifiers.</p> <p>However, QFX5100, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems, the switch uses only one EXP classifier as a global MPLS classifier on all interfaces. You specify the global EXP classifier in the <b>[edit class-of-service system-defaults]</b> hierarchy.</p> |
| Options                                       | <b>classifier-name</b> —Name of the EXP classifier.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Required Privilege Level                      | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

**Related  
Documentation**

- [Configuring Rewrite Rules for MPLS EXP Classifiers on page 4992](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4959](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 6672](#)

## forwarding-class

|                                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                | <a href="#">Classifier on page 7075</a><br><a href="#">Rewrite Rule on page 7075</a><br><a href="#">Scheduler Map on page 7075</a><br><a href="#">Interface on page 7075</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Classifier</b>                    | <pre>forwarding-class class-name {   loss-priority level {     code-points [ aliases ] [ bit-patterns ];   } }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Rewrite Rule</b>                  | <pre>forwarding-class class-name {   loss-priority level {     code-point [ aliases ] [ bit-patterns ];   } }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Scheduler Map</b>                 | <pre>forwarding-class class-name {   scheduler scheduler-name; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Interface</b>                     | <pre>forwarding-class class-name;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Classifier Hierarchy Level</b>    | [edit <a href="#">class-of-service classifiers</a> (dscp   dscp-ipv6   ieee-802.1   exp) classifier-name],                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Rewrite Rule Hierarchy Level</b>  | [edit <a href="#">class-of-service rewrite-rules</a> ] (dscp   dscp-ipv6   ieee-802.1) rewrite-name   exp],                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Scheduler Map Hierarchy Level</b> | [edit <a href="#">class-of-service scheduler-maps</a> map-name],                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Interface Hierarchy Level</b>     | [edit <a href="#">class-of-service interfaces</a> interface-name unit logical-unit-number]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>           | <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>                   | <ul style="list-style-type: none"> <li>Classifiers—Assign incoming traffic to the specified forwarding class based on the specified code point values and assign that traffic the specified loss priority</li> <li>Rewrite rules—At the egress interface, change (rewrite) the value of the code point bits and the loss priority to specified new values for traffic assigned to the specified forwarding class, before forwarding the traffic to the next hop.</li> <li>Scheduler maps—Apply the specified scheduler to the specified forwarding class.</li> <li>Interfaces—Assign the specified forwarding class to the interface to use as a fixed classifier (all incoming traffic on the interface is classified into that forwarding class).</li> </ul> |



**NOTE:** OCX Series switches do not support MPLS, so they do not support EXP classifiers or rewrite rules.

**Options** *class-name*—Name of the forwarding class.

The remaining statements are explained separately.

**Required Privilege** interfaces—To view this statement in the configuration.

**Level** interface-control—To add this statement to the configuration.

---

## forwarding-class-sets

---

**Syntax** forwarding-class-sets *forwarding-class-set-name* {  
    *class class-name*;  
}

**Hierarchy Level** [edit [class-of-service](#)]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series

**Description** Assign forwarding classes to forwarding class sets (priority groups).

**Options** *forwarding-class-set-name*—Name of the forwarding class set.

The remaining statement is explained separately.

**Required Privilege** interfaces—To view this statement in the configuration.

**Level** interface-control—To add this statement to the configuration.

- Related Documentation**
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 6771](#)
  - [Example: Configuring Forwarding Class Sets on page 6697](#)
  - [Understanding CoS Forwarding Class Sets \(Priority Groups\) on page 6695](#)

## forwarding-classes

**Syntax**

```
forwarding-classes {
 class {
 class-name {
 queue-num queue-number
 no-loss {
 }
 }
 }
}
```

**Hierarchy Level** [edit [class-of-service](#)]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.  
No-loss option introduced in Junos OS Release 12.3 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series

**Description** Map one or more forwarding classes to a single output queue.

Switches that use different forwarding classes for unicast and multdestination (multicast, broadcast, and destination lookup fail) traffic support 12 forwarding classes and 12 output queues (0 through 11). You map unicast forwarding classes to a unicast queue (0 through 7) and multdestination forwarding classes to a multdestination queue (8 through 11). The queue to which you map a forwarding class determines if the forwarding class is a unicast or multdestination forwarding class.

Switches that use the same forwarding classes for unicast and multdestination traffic support eight forwarding classes and eight output queues (0 through 7). You map forwarding classes to output queues. All traffic classified into one forwarding class (unicast and multdestination) uses the same output queue.

You cannot configure weighted random early detection (WRED) packet drop on forwarding classes configured with the no-loss packet drop attribute. Do not associate a drop profile with lossless forwarding classes.



**NOTE:** If you map more than one forwarding class to a queue, all of the forwarding classes mapped to the queue must have the same packet drop attribute (all of the forwarding classes must be lossy, or all of the forwarding classes mapped to a queue must be lossless).

OCX Series switches do not support the no-loss packet drop attribute and do not support lossless forwarding classes. On OCX Series switches, do not configure the no-loss packet drop attribute on forwarding classes, and do not map traffic to the default fcoe and no-loss forwarding classes (both of these default forwarding classes carry the no-loss packet drop attribute).



NOTE: On switches that do not use the ELS CLI, if you are using Junos OS Release 12.2, use the default forwarding-class-to-queue mapping for the lossless fcoe and no-loss forwarding classes. If you explicitly configure the lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best effort) traffic and does *not* receive lossless treatment.



NOTE: On switches that do not use the ELS CLI, if you are using Junos OS Release 12.3 or later, the default configuration is the same as the default configuration for Junos OS Release 12.2, and the default behavior is the same (the fcoe and no-loss forwarding classes receive lossless treatment). However, if you explicitly configure lossless forwarding classes, you can configure up to six lossless forwarding classes by specifying the no-loss option. If you do not specify the no-loss option in an explicit forwarding class configuration, the forwarding class is lossy. For example, if you explicitly configure the fcoe forwarding class and you do not include the no-loss option, the fcoe forwarding class is lossy, not lossless.

**Options**     The statements are explained separately.

**Required Privilege**     interfaces—To view this statement in the configuration.  
**Level**                     interface-control—To add this statement to the configuration.


## ieee-802.1

|                                                    |                                                                                                                                                                                                                                                                                                                 |
|----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| List of Syntax                                     | <a href="#">Syntax (Classifier) on page 7079</a><br><a href="#">Syntax (Code-Point Alias) on page 7079</a><br><a href="#">Syntax (Multidestination Classifier) on page 7079</a><br><a href="#">Syntax (Interface Classifier Association) on page 7079</a><br><a href="#">Syntax (Rewrite Rule) on page 7079</a> |
| Syntax (Classifier)                                | <pre> ieee-802.1 classifier-name {   import (classifier-name   default);   forwarding-class class-name {     loss-priority level {       code-points [ aliases ] [ bit-patterns ];     }   } }</pre>                                                                                                            |
| Syntax (Code-Point Alias Configuration)            | <pre> ieee-802.1 alias-name bit-pattern;</pre>                                                                                                                                                                                                                                                                  |
| Syntax (Multidestination Classifier Configuration) | <pre> ieee-802.1 classifier-name;</pre>                                                                                                                                                                                                                                                                         |
| Syntax (Interface Classifier Association)          | <pre> ieee-802.1 (classifier-name   default);</pre>                                                                                                                                                                                                                                                             |
| Syntax (Rewrite Rule Configuration)                | <pre> ieee-802.1 rewrite-name {   import (rewrite-name   default);   forwarding-class class-name {     loss-priority level {       code-point [ aliases ] [ bit-patterns ];     }   } }</pre>                                                                                                                   |
| Hierarchy Level (Classifier)                       | [edit <a href="#">class-of-service classifiers</a> ],                                                                                                                                                                                                                                                           |
| Hierarchy Level (Code-Point Alias)                 | [edit <a href="#">class-of-service code-point-aliases</a> ],                                                                                                                                                                                                                                                    |
| Hierarchy Level (Multidestination Classifier)      | [edit <a href="#">class-of-service multi-destination classifiers</a> ],                                                                                                                                                                                                                                         |
| Hierarchy Level (Interface Classifier Association) | [edit <a href="#">class-of-service interfaces interface-name unit logical-unit-number classifiers</a> ],<br>[edit <a href="#">class-of-service interfaces interface-name unit logical-unit-number rewrite-rules</a> ],                                                                                          |
| Hierarchy Level (Rewrite Rule)                     | [edit <a href="#">class-of-service rewrite-rules</a> ]                                                                                                                                                                                                                                                          |
| Release Information                                | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                   |

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b>              | Configure an IEEE 802.1 classifier, configure an IEEE 802.1 code-point alias, apply a fixed IEEE 802.1 classifier to an interface, or apply an IEEE-802.1 rewrite rule.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <p><i>classifier-name</i>—Name of the classifier.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Unicast Classifiers</i></li><li>• <a href="#">Defining CoS Code-Point Aliases on page 6687</a></li><li>• <a href="#">Defining CoS Rewrite Rules on page 6704</a></li><li>• <a href="#">Assigning CoS Components to Interfaces on page 6627</a></li><li>• <a href="#">Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 6672</a></li><li>• <i>Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces</i></li><li>• <i>Understanding CoS Classifiers</i></li><li>• <i>Understanding CoS Classifiers</i></li><li>• <a href="#">Understanding CoS Rewrite Rules on page 6701</a></li></ul> |



## import

|                                                                                                                                                                                                                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                  | <code>import (<i>import</i>   default);</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>                                                                                                                                                                                         | [edit <a href="#">class-of-service classifiers</a> ( <a href="#">dscp</a>   <a href="#">dscp-ipv6</a>   <a href="#">ieee-802.1</a>   <a href="#">exp</a> ) <i>classifier-name</i> ],<br>[edit <a href="#">class-of-service rewrite-rules</a> ( <a href="#">dscp</a>   <a href="#">dscp-ipv6</a>   <a href="#">ieee-802.1</a>   <a href="#">exp</a> ) <i>classifier-name</i> ]                                                                                                                                                                      |
| <b>Release Information</b>                                                                                                                                                                                     | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>                                                                                                                                                                                             | Specify a default or previously defined classifier.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <div>  <b>NOTE:</b> OCX Series switches do not support MPLS, so they do not support EXP classifiers and rewrite rules. </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                                                                                                                                                                                                 | <p><b><i>import</i></b>—Name of the classifier mapping configured at the [edit <a href="#">class-of-service classifiers</a>] hierarchy level.</p> <p><b><i>default</i></b>—Default classifier mapping.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b>                                                                                                                                                                                | <p><b>interfaces</b>—To view this statement in the configuration.</p> <p><b>interface-control</b>—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>                                                                                                                                                                                   | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Unicast Classifiers</a></li> <li>• <a href="#">Defining CoS BA Classifiers (DSCP, DSCP IPv6, IEEE 802.1p) on page 6656</a></li> <li>• <a href="#">Defining CoS Rewrite Rules on page 6704</a></li> <li>• <a href="#">Understanding CoS Classifiers</a></li> <li>• <a href="#">Understanding CoS Classifiers</a></li> <li>• <a href="#">Understanding CoS Classifiers on page 6650</a></li> <li>• <a href="#">Understanding CoS Rewrite Rules on page 6701</a></li> </ul> |

## interfaces (Class of Service)

```
Syntax interfaces {
 interface-name {
 congestion-notification-profile profile-name {
 }
 forwarding-class forwarding-class-name;
 forwarding-class-set forwarding-class-set-name {
 output-traffic-control-profile profile-name;
 }
 rewrite-value {
 input {
 ieee-802.1p {
 code-point code-point-bits;
 }
 }
 }
 scheduler-map scheduler-map-name
 unit logical-unit-number {
 classifiers {
 (dscp | dscp-ipv6 | ieee-802.1p | exp) (classifier-name | default);
 }
 forwarding-class class-name;
 rewrite-rules {
 (dscp | dscp-ipv6 | ieee-802.1p | exp) (classifier-name | default);
 }
 }
 }
}
```

**Hierarchy Level** [edit [class-of-service](#)]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series

**Description** Configure interface-specific CoS properties for incoming packets.



**NOTE:** Only switches that support direct port scheduling also support applying a scheduler map directly to an interface. When using enhanced transmission selection (ETS) hierarchical port scheduling, you cannot apply a scheduler map directly to an interface; instead, you associate the scheduler map with a traffic control profile and apply the traffic control profile to the interface.



**NOTE:** Only switches that support native Fibre Channel interfaces support the rewrite-value statement, which enables you to rewrite the IEEE 802.1p code points on native Fibre Channel interfaces.



**NOTE:** OCX Series switches do not support MPLS, so they do not support EXP classifiers or rewrite rules. OCX Series switches do not support the congestion-notification-profile configuration statement, which applies priority-based flow control (PFC) to interface output queues.

**Options** *interface-name*—Name of the interface.

The statements are explained separately.

**Required Privilege Level** interfaces—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Assigning CoS Components to Interfaces on page 6627](#)
- [Interfaces Overview on page 2785](#)

## loss-priority (Classifiers)

---

|                            |                                                                                                                                                                                                                     |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>loss-priority level {<br/>    code-points [ aliases ] [ bit-patterns ];<br/>}</code>                                                                                                                          |
| <b>Hierarchy Level</b>     | [edit <a href="#">class-of-service classifiers</a> ( <a href="#">dscp</a>   <a href="#">dscp-ipv6</a>   <a href="#">ieee-802.1</a> ) <i>classifier-name</i><br><a href="#">forwarding-class</a> <i>class-name</i> ] |
| <b>Release Information</b> | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                       |
| <b>Description</b>         | Configure packet loss priority value for a specific set of code-point aliases and bit patterns.                                                                                                                     |



**NOTE:** OCX Series switches do not support MPLS, so they do not support EXP classifiers.

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | <p><i>level</i>—Can be one of the following:</p> <ul style="list-style-type: none"><li>• <b>low</b>—Packet has low loss priority.</li><li>• <b>medium-high</b>—Packet has medium-high loss priority.</li><li>• <b>high</b>—Packet has high loss priority.</li></ul> <p>The remaining statement is explained separately.</p>                                            |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Unicast Classifiers</i></li><li>• <a href="#">Defining CoS BA Classifiers (DSCP, DSCP IPv6, IEEE 802.1p)</a> on page 6656</li><li>• <i>Understanding CoS Classifiers</i></li><li>• <a href="#">Understanding CoS Classifiers on page 6650</a></li><li>• <i>Understanding CoS Classifiers</i></li></ul> |

## loss-priority (Rewrite Rules)

|                            |                                                                                                                                                                                                                                                                                                      |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>loss-priority <i>level</i> {<br/>    code-point (<i>alias</i>   <i>bit-pattern</i>);<br/>}</code>                                                                                                                                                                                              |
| <b>Hierarchy Level</b>     | [edit <code>class-of-service rewrite-rules (dscp   ieee-802.1) rewrite-name forwarding-class class-name</code> ]                                                                                                                                                                                     |
| <b>Release Information</b> | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                        |
| <b>Description</b>         | Specify a loss priority to which to apply a rewrite rule. The rewrite rule sets the code-point aliases and bit patterns for a specific forwarding class and loss priority. Packets that match the forwarding class and loss priority are rewritten with the rewrite code-point alias or bit pattern. |



**NOTE:** OCX Series switches do not support MPLS, so they do not support EXP rewrite rules.

|                                 |                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | <p><i>level</i>—Can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>low</b>—Packet has low loss priority.</li> <li>• <b>medium-high</b>—Packet has medium-high loss priority.</li> <li>• <b>high</b>—Packet has high loss priority.</li> </ul> <p>The remaining statement is explained separately.</p> |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Defining CoS Rewrite Rules on page 6704</a></li> <li>• <a href="#">Understanding CoS Rewrite Rules on page 6701</a></li> </ul>                                                                                                                                             |

## queue-num

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>queue-num <i>queue-number</i> &lt;no-loss&gt;;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>     | [edit <code>class-of-service forwarding-classes class <i>class-name</i></code> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b> | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>No-loss option introduced in Junos OS Release 12.3 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>         | <p>Map a forwarding class to an output queue number. Optionally, configure the forwarding class as a lossless forwarding class. Each switch provides enough output queues so that you can map forwarding classes to queues on a one-to-one basis, so each forwarding class can have a dedicated output queue.</p> <p>On switches that use different forwarding classes and output queues for unicast and multdestination (multicast, broadcast, destination lookup fail) traffic, the switch supports 12 forwarding classes and 12 output queues, eight of each for unicast traffic and four of each for multdestination traffic. You can map some or all of the eight unicast forwarding classes to a unicast queue (0 through 7) and some or all of the four multdestination forwarding classes to the a multdestination queue (8 through 11). You cannot map a forwarding class to more than one queue (each forwarding class maps to one and only one queue), but you can map multiple forwarding classes to one queue. The queue to which you map a forwarding class determines if the forwarding class is a unicast or multdestination forwarding class.</p> <p>On switches that use the same forwarding classes and output queues for unicast and multdestination traffic, the switch supports eight forwarding classes and eight output queues. You can map some or all of the eight of the forwarding classes to queues (0 through 7). You cannot map a forwarding class to more than one queue (each forwarding class maps to one and only one queue), but you can map multiple forwarding classes to one queue.</p> <p>You cannot configure weighted random early detection (WRED) packet drop on forwarding classes configured with the no-loss packet drop attribute. Do not associate a drop profile with lossless forwarding classes. Instead, use priority-based flow control (PFC) to prevent frame drop on lossless forwarding classes.</p> |



**NOTE:** If you map more than one forwarding class to a queue, all of the forwarding classes mapped to the same queue must have the same packet drop attribute (all of the forwarding classes must be lossy, or all of the forwarding classes mapped to a queue must be lossless).

OCX Series switches do not support the no-loss packet drop attribute and do not support lossless forwarding classes. On OCX Series switches, do not configure the no-loss packet drop attribute on forwarding classes, and do not map traffic to the default fcoe and no-loss forwarding classes (both of these default forwarding classes carry the no-loss packet drop attribute).



**NOTE:** On systems that do not use the ELS CLI, if you are using Junos OS Release 12.2, use the default forwarding-class-to-queue mapping for the lossless fcoe and no-loss forwarding classes. If you explicitly configure lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best effort) traffic and does *not* receive lossless treatment.



**NOTE:** On systems that do not use the ELS CLI, if you are using Junos OS Release 12.3 or later, the default configuration is the same as the default configuration for Junos OS Release 12.2, and the default behavior is the same (the fcoe and no-loss forwarding classes receive lossless treatment). However, if you explicitly configure lossless forwarding classes, you can configure up to six lossless forwarding classes by specifying the no-loss option. If you do not specify the no-loss option in an explicit forwarding class configuration, the forwarding class is lossy. For example, if you explicitly configure the fcoe forwarding class and you do not include the no-loss option, the fcoe forwarding class is lossy, not lossless.


**Options** *queue-number*—(Switches that use different output queues for unicast and multidestination traffic) Number of the CoS unicast queue (0 through 7) or the CoS multidestination queue (8 through 11).

*queue-number*—(Switches that use the same output queues for unicast and multidestination traffic) Number of the CoS queue (0 through 7).

**no-loss**—Optional packet drop attribute keyword to configure the forwarding class as lossless.

**Required Privilege** interfaces—To view this statement in the configuration.  
**Level** interface-control—To add this statement to the configuration.

## rewrite-rules

|                                                                  |                                                                                                                                                                                                                                                                                                                                                                                       |
|------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                                            | <a href="#">Syntax (Rewrite Rule Configuration) on page 7088</a><br><a href="#">Syntax (Rewrite Rule Association with Interface) on page 7088</a>                                                                                                                                                                                                                                     |
| <b>Syntax (Rewrite Rule Configuration)</b>                       | <pre>rewrite-rules {   (dscp   dscp-ipv6   ieee-802.1   exp) rewrite-name {     import (rewrite-name   default);     forwarding-class class-name {       loss-priority priority code-point (alias   bits);     }   } }</pre>                                                                                                                                                          |
| <b>Syntax (Rewrite Rule Association with Interface)</b>          | <pre>rewrite-rules {   (dscp   dscp-ipv6   ieee-802.1   exp) rewrite-name; }</pre>                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level (Rewrite Rule Configuration)</b>              | [edit <a href="#">class-of-service</a> ],                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level (Rewrite Rule Association with Interface)</b> | [edit <a href="#">class-of-service interfaces</a> <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> ]                                                                                                                                                                                                                                                             |
| <b>Release Information</b>                                       | <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>EXP statement introduced in Junos OS Release 12.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                 |
| <b>Description</b>                                               | <p>Configure rewrite rules that map traffic to code points when traffic exits the system, and apply the rewrite rules to a specific interface.</p> <p>MPLS EXP rewrite rules can only be bound to logical interfaces, not to physical interfaces. You can configure up to 64 EXP rewrite rules, but you can use only 16 EXP rewrite rules on switch interfaces at any given time.</p> |
|                                                                  | <div>  <p><b>NOTE:</b> OCX Series switches do not support MPLS, so they do not support EXP rewrite rules.</p> </div>                                                                                                                                                                               |
| <b>Options</b>                                                   | The statements are explained separately.                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b>                                  | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>                                     | <ul style="list-style-type: none"> <li>• <a href="#">Defining CoS Rewrite Rules on page 6704</a></li> <li>• <a href="#">Configuring Rewrite Rules for MPLS EXP Classifiers on page 4992</a></li> </ul>                                                                                                                                                                                |



- [Understanding CoS Rewrite Rules on page 6701](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4959](#)

## unit

**Syntax** `unit logical-unit-number {  
     classifiers {  
         (dscp | dscp-ipv6 | ieee-802.1 | exp) (classifier-name | default);  
     }  
     forwarding-class class-name;  
     rewrite-rules {  
         (dscp | dscp-ipv6 | ieee-802.1 | exp) (classifier-name | default);  
     }  
 }`

**Hierarchy Level** [edit [class-of-service interfaces](#) *interface-name*]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.  
 Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Configure a logical interface on the physical device. You must configure a logical interface to use the physical device.



**NOTE:** OCX Series switches do not support MPLS, so they do not support EXP classifiers and rewrite rules.

**Options** *logical-unit-number*—Number of the logical unit.

**Range:** 0 through 16,385

The remaining statements are explained separately.

**Required Privilege Level** *interfaces*—To view this statement in the configuration.  
*interface-control*—To add this statement to the configuration.

**Related Documentation** • [Assigning CoS Components to Interfaces on page 6627](#)



## Configuration Statements (Scheduling)

- [buffer-size on page 7092](#)
- [drop-probability on page 7097](#)
- [drop-profile on page 7098](#)
- [drop-profile-map on page 7098](#)
- [drop-profiles on page 7099](#)
- [excess-rate on page 7100](#)
- [explicit-congestion-notification on page 7101](#)
- [fill-level on page 7102](#)
- [forwarding-class on page 7104](#)
- [forwarding-class-set on page 7105](#)
- [guaranteed-rate on page 7106](#)
- [interpolate on page 7107](#)
- [loss-priority \(Drop Profiles\) on page 7108](#)
- [output-traffic-control-profile on page 7108](#)
- [priority \(Schedulers\) on page 7109](#)
- [protocol \(Drop Profile Map\) on page 7110](#)
- [scheduler on page 7111](#)
- [scheduler-map on page 7111](#)
- [scheduler-maps on page 7112](#)
- [schedulers on page 7113](#)
- [shaping-rate on page 7114](#)
- [traffic-control-profiles on page 7116](#)
- [transmit-rate on page 7117](#)

## buffer-size

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>buffer-size (percent <i>percent</i>   remainder);</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>     | [edit <code>class-of-service schedulers <i>scheduler-name</i></code> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b> | Statement introduced in Junos OS Release 12.3 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>         | <p>On all switches, you configure the proportion of port buffers allocated to a particular output queue using the following process:</p> <ol style="list-style-type: none"><li>1. Configure a scheduler and set the <b>buffer-size</b> option.</li><li>2. Use a scheduler map to map the scheduler to the forwarding class that is mapped to the queue to which you want to apply the buffer size.</li></ol> <p>For example, suppose that you want to change the dedicated buffer allocation for FCoE traffic. FCoE traffic is mapped to the <code>fcoe</code> forwarding class, and the <code>fcoe</code> forwarding class is mapped to queue 3 (this is the default configuration). To use default FCoE traffic mapping, in the scheduler map configuration, map the scheduler to the <b>fcoe</b> forwarding class.</p> <ol style="list-style-type: none"><li>3. If you are using enhanced transmission selection (ETS) hierarchical scheduling, associate the scheduler map with the traffic control profile you want to use on the egress ports that carry FCoE traffic. If you are using direct port scheduling, skip this step.</li><li>4. If you are using ETS, associate the traffic control profile that includes the scheduler map with the desired egress ports. For this example, you associate the traffic control profile with the ports that carry FCoE traffic. If you are using port scheduling, associate the scheduler map with the desired egress ports.</li></ol> <p>Queue 3, which is mapped to the <code>fcoe</code> forwarding class and therefore to the FCoE traffic, receives the dedicated buffer allocation specified in the <b>buffer-size</b> statement.</p> |



**NOTE:** The total of all of the explicitly configured buffer size percentages for all of the queues on a port cannot exceed 100 percent.

---

### QFX10000 Switches

On QFX10000 switches, the buffer size is the amount of time in milliseconds of port bandwidth that a queue can use to continue to transmit packets during periods of congestion, before the buffer runs out and packets begin to drop.

The switch can use up to 100 ms total (combined) buffer space for all queues on a port. A buffer-size configured as one percent is equal to 1 ms of buffer usage. A buffer-size of 15 percent (the default value for the best effort and network control queues) is equal to 15 ms of buffer usage.

The total buffer size of the switch is 4 GB. A 40-Gigabit port can use up to 500 MB of buffer space, which is equivalent to 100 ms of port bandwidth on a 40-Gigabit port. A 10-Gigabit port can use up to 125 MB of buffer space, which is equivalent to 100 ms of port bandwidth on a 10-Gigabit port. The total buffer sizes of the eight output queues on a port cannot exceed 100 percent, which is equal to the full 100 ms total buffer available to a port. The maximum amount of buffer space any queue can use is also 100 ms (which equates to a 100 percent buffer-size configuration), but if one queue uses all of the buffer, then no other queue receives buffer space.

There is no minimum buffer allocation, so you can set the buffer-size to zero (0) for a queue. However, we recommend that on queues on which you enable PFC to support lossless transport, you allocate a minimum of 5 ms (a minimum buffer-size of 5 percent). The two default lossless queues, fcoe and no-loss, have buffer-size default values of 35 ms (35 percent).

Queue buffer allocation is dynamic, shared among ports as needed. However, a queue cannot use more than its configured amount of buffer space. For example, if you are using the default CoS configuration, the best-effort queue receives a maximum of 15 ms of buffer space because the default transmit rate for the best-effort queue is 15 percent.

If a switch experiences congestion, queues continue to receive their full buffer allocation until 90 percent of the 4 GB buffer space is consumed. When 90 percent of the buffer space is in use, the amount of buffer space per port, per queue, is reduced in proportion to the configured buffer size for each queue. As the percentage of consumed buffer space rises above 90 percent, the amount of buffer space per port, per queue, continues to be reduced.

On 40-Gigabit ports, because the total buffer is 4 GB and the maximum buffer a port can use is 500 MB, up to seven 40-Gigabit ports can consume their full 100 ms allocation of buffer space. However, if an eighth 40-Gigabit port requires the full 500 MB of buffer space, then the buffer allocations are proportionally reduced because the buffer consumption is above 90 percent.

On 10-Gigabit ports, because the total buffer is 4 GB and the maximum buffer a port can use is 125 MB, up to 28 10-Gigabit ports can consume their full 100 ms allocation of buffer space. However, if a 29th 10-Gigabit port requires the full 125 MB of buffer space, then the buffer allocations are proportionally reduced because the buffer consumption is above 90 percent.

**QFX5100, EX4600,  
QFX3500, and  
QFX3600 Switches,  
and QFabric Systems**

Set the dedicated buffer size of the egress queue that you bind the scheduler to in the scheduler map configuration. The switch allocates space from the global dedicated buffer pool to ports and queues in a hierarchical manner. The switch allocates an equal number of dedicated buffers to each egress port, so each egress port receives the same amount of dedicated buffer space. The amount of dedicated buffer space per port is not configurable.

However, the **buffer-size** statement allows you to control the way each port allocates its share of dedicated buffers to its queues. For example, if a port only uses two queues to forward traffic, you can configure the port to allocate all of its dedicated buffer space to those two ports and avoid wasting buffer space on queues that are not in use. We recommend that the buffer size should be the same size as the minimum guaranteed transmission rate (the **transmit-rate**).

**Default** The default behavior of the differs on different switches.

#### QFX10000 Switches

If you do not configure buffer-size and you do not explicitly configure a queue scheduler, the default buffer-size is the default transmit rate of the queue. If you explicitly configure a queue scheduler, the default buffer allocations are not used. If you explicitly configure a queue scheduler, configure the buffer-size for each queue in the scheduler, keeping in mind that the total buffer-size of the queues cannot exceed 100 percent (100 ms).

[Table 603](#) shows the default queue buffer sizes on QFX10000 switches. The default buffer size is the same as the default transmit rate for each default queue:

**Table 603: Default Output Queue Buffer Sizes (QFX10000 Switches)**

| Queue Number | Forwarding Class | Transmit Rate | Buffer Size |
|--------------|------------------|---------------|-------------|
| 0            | best-effort      | 15%           | 15%         |
| 3            | fcoe             | 35%           | 35%         |
| 4            | no-loss          | 35%           | 35%         |
| 7            | network-control  | 15%           | 15%         |

By default, only the queues mapped to the default forwarding classes receive buffer space from the port buffer pool. (Buffers are not wasted on queues that do not carry traffic.)

#### QFX5100, EX4600, QFX3500, and QFX3600 Switches, and QFabric Systems

The port allocates dedicated buffers to queues that have an explicitly configured scheduler buffer size. If you do not explicitly configure a scheduler buffer size for a queue, the port serves the explicitly configured queues first. Then the port divides the remaining dedicated buffers equally among the queues that have an explicitly attached scheduler *without* an explicitly configured buffer size configuration. (If you configure a scheduler, but you do not configure the buffer size parameter, the default is equivalent to configuring the buffer size with the **remainder** option.)

If you use the default scheduler and scheduler map on a port (no explicit scheduler configuration), then the port allocates its dedicated buffer pool to queues based on the default scheduling. [Table 604](#) shows the default queue buffer sizes. The default buffer size is the same as the default transmit rate for each default queue:

**Table 604: Default Output Queue Buffer Sizes (QFX5100, EX4600, QFX3500, and QFX3600 Switches, and QFabric Systems)**

| Queue Number | Forwarding Class | Transmit Rate | Buffer Size |
|--------------|------------------|---------------|-------------|
| 0            | best-effort      | 5%            | 5%          |
| 3            | fcoe             | 35%           | 35%         |
| 4            | no-loss          | 35%           | 35%         |
| 7            | network-control  | 5%            | 5%          |
| 8            | mcast            | 20%           | 20%         |

By default, only the queues mapped to the default forwarding classes receive buffer space from the port buffer pool. (Buffers are not wasted on queues that do not carry traffic.)



**NOTE:** OCX Series switches do not support lossless transport. On OCX Series switches, do not map traffic to the lossless default fcoe and no-loss forwarding classes. OCX Series default DSCP classification does not map traffic to the fcoe and no-loss forwarding classes, so by default, the OCX system does not classify traffic into those forwarding classes. (On other switches, the fcoe and no-loss forwarding classes provide lossless transport for Layer 2 traffic. OCX Series switches do not support lossless Layer 2 transport.) The active forwarding classes (**best-effort**, **network-control**, and **mcast**) share the unused bandwidth assigned to the fcoe and no-loss forwarding classes.

**Options**    **percent percent**—Percentage of the port dedicated buffer pool allocated to the queue (or queues) mapped to the scheduler.

**remainder**—Remaining dedicated buffer pool after the port satisfies the needs of the explicitly configured buffers. The port divides the remaining buffers equally among the queues that are explicitly attached to a scheduler but that do not have an explicit buffer size configuration (or are configured with **remainder** as the buffer size).

**Required Privilege Level**    interfaces—To view this statement in the configuration.  
                                          interface-control—To add this statement to the configuration.



## drop-probability

|                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                                            | QFX5100, EX4600, QFX3500, and QFX3600, Switches, QFabric Systems on page 7097<br>QFX10000 Switches on page 7097                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| QFX5100, EX4600, QFX3500, and QFX3600, Switches, QFabric Systems | drop-probability 0 drop-probability <i>high-value</i> ;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| QFX10000 Switches                                                | drop-probability <i>percentage1 percentage2 ... percentage32</i> ;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>                                           | [edit <a href="#">class-of-service drop-profiles profile-name interpolate</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>                                       | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>                                               | <p>When configuring WRED, map the packet <b>drop-probability</b> to the fullness of a queue (<b>fill-level</b>). You configure the <b>fill-level</b> and <b>drop-probability</b> statements in related pairs. The pairs of fill level and drop probability values set a probability of dropping packets at a specified queue fullness value.</p> <p>On switches that support only two fill level/drop probability pairs, the first drop probability is always zero. The first fill level/drop probability pair sets the drop start point, and the second fill level/drop probability pair sets the drop end point.</p> <p>On switches that support 32 fill level/drop probability pairs, the first fill level/drop probability pair sets the drop start point, and the last fill level/drop probability pair sets the drop end point.</p> <p>As the queue fills from the drop start point to the drop end point, the rate of packet drop increases in a curve pattern. The higher the queue fill level, the higher the probability of dropping packets.</p>                       |
| <b>Options</b>                                                   | <p>0 (switches that support only two fill level/drop probability pairs)—Probability that packets will drop at the lowest <b>fill-level</b> value. This is always zero, because until the queue reaches the specified low <b>fill-level</b> value, no packets are scheduled to drop.</p> <p><b>Range:</b> 0</p> <p><b>high-value</b> (switches that support only two fill level/drop probability pairs)—The maximum probability that packets will drop before queue fullness exceeds the high value of the queue <b>fill-level</b>, expressed as a percentage. If the queue fills beyond the high <b>fill-level</b> value, all packets drop.</p> <p><b>Range:</b> 0 through 100 percent</p> <p><i>percentage1 percentage2 ... percentage32</i> (switches that support 32 fill level/drop probability pairs)—The probability that packets will drop before the queue fullness exceeds the <b>fill-level</b> value, expressed as a percentage. Each drop probability pairs with a queue fill level to define the probability of a packet dropping at a specified queue fullness.</p> |

**Range:** 0 through 100 percent

**Required Privilege** interface—To view this statement in the configuration.  
**Level** interface-control—To add this statement to the configuration.

---

## drop-profile

---

|                                 |                                                                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>drop-profile <i>profile-name</i>;</code>                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service schedulers scheduler-name</a> <b>drop-profile-map</b> <i>loss-priority</i> (low   medium-high   high) <a href="#">protocol protocol</a> ]                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                           |
| <b>Description</b>              | Define drop profiles for weighted random early detection (WRED). When a packet arrives, WRED checks the queue fill level specified in the drop profile. If the fill level corresponds to a nonzero drop probability, the WRED algorithm determines whether to drop the arriving packet. |
| <b>Options</b>                  | <i>profile-name</i> —Name of the drop profile.                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | interfaces—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Drop Profile Maps on page 6817</a></li></ul>                                                                                                                                                                   |

---

## drop-profile-map

---

|                                 |                                                                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>drop-profile-map <i>loss-priority</i> (low   medium-high   high) <a href="#">protocol protocol</a> <b>drop-profile</b> <i>drop-profile-name</i>;</code>                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service schedulers scheduler-name</a> ]                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                           |
| <b>Description</b>              | Map a drop profile to a loss priority and protocol for weighted random early detection (WRED). When a packet arrives, WRED checks the queue fill level. If the fill level corresponds to a nonzero drop probability, the WRED algorithm determines whether to drop the arriving packet. |
| <b>Options</b>                  | The statements are explained separately.                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | interfaces—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Drop Profile Maps on page 6817</a></li></ul>                                                                                                                                                                   |

## drop-profiles

|                                                                  |                                                                                                                                                                                                                                                                                                                                     |
|------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                                            | QFX5100, EX4600, QFX3500, and QFX3600, Switches, QFabric Systems on page 7099<br>QFX10000 Switches on page 7099                                                                                                                                                                                                                     |
| QFX5100, EX4600, QFX3500, and QFX3600, Switches, QFabric Systems | <pre>drop-profiles {   profile-name {     interpolate {       fill-level low-value fill-level high-value drop-probability 0 drop-probability high-value;     }   } }</pre>                                                                                                                                                          |
| QFX10000 Switches                                                | <pre>drop-profiles {   profile-name {     interpolate {       fill-level level1 level2 ... level32 drop-probability percent1 percent2 ... percent32;     }   } }</pre>                                                                                                                                                              |
| <b>Hierarchy Level</b>                                           | [edit <a href="#">class-of-service</a> ]                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>                                       | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                       |
| <b>Description</b>                                               | <p>Define drop profiles for weighted random early detection (WRED).</p> <p>For a packet to be dropped, it must match the drop profile. When a packet arrives, WRED checks the queue fill level. If the fill level corresponds to a nonzero drop probability, the WRED algorithm determines whether to drop the arriving packet.</p> |
| <b>Options</b>                                                   | <p><i>profile-name</i>—Name of the drop profile.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b>                                  | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                  |

## excess-rate

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>excess-rate percent <i>percentage</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service traffic-control-profiles <i>profile-name</i></a> ],<br>[edit <a href="#">class-of-service schedulers <i>scheduler-name</i></a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 15.1X53-D10 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | <p>Determine the percentage of excess port bandwidth for which a queue (forwarding class) that is not a strict-high priority queue or forwarding class set (priority group) can contend. Excess bandwidth is the extra port bandwidth left after strict-high priority queues and the guaranteed minimum bandwidth requirements of other queues (as determined by each queue's transmit rate) are satisfied. With the exception of strict-high priority queues, the switch allocates extra port bandwidth to queues or to priority groups based on the configured excess rate. If you do not configure an excess rate for a queue, the default excess rate is the same as the transmit rate.</p> <p>You cannot configure an excess rate on strict-high priority queues. Strict-high priority queues receive extra bandwidth based on an extra bandwidth sharing weight of "1", which is not configurable. However, the switch serves traffic on strict-high priority queues up to the configured transmit rate before it serves any other queues, so by configuring an appropriate transmit rate on a strict-high priority queue, you can guarantee strict-high priority traffic on that queue is treated in the manner you want.</p> |
| <b>Options</b>                  | <p><b>percent <i>percentage</i></b>—Percentage of the excess bandwidth to share.</p> <p><b>Range:</b> 0 through 100 percent</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Defining CoS Queue Schedulers for Port Scheduling on page 6742</a></li><li>• <a href="#">Example: Configuring Traffic Control Profiles (Priority Group Scheduling) on page 6762</a></li><li>• <a href="#">Understanding CoS Port Schedulers on QFX Switches on page 6729</a></li><li>• <a href="#">Understanding CoS Traffic Control Profiles on page 6757</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

## explicit-congestion-notification

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>explicit-congestion-notification;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service schedulers</a> <i>scheduler-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X51 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | <p>Enable explicit congestion notification (ECN) on the output queue (forwarding class) or output queues (forwarding classes) mapped to the scheduler. ECN enables end-to-end congestion notification between two endpoints on TCP/IP based networks. The two endpoints are an ECN-enabled sender and an ECN-enabled receiver. ECN must be enabled on both endpoints and on all of the intermediate devices between the endpoints for ECN to work properly. Any device in the transmission path that does not support ECN breaks the end-to-end ECN functionality.</p> <p>A weighted random early detection (WRED) packet drop profile must be applied to the output queues on which ECN is enabled. ECN uses the WRED drop profile thresholds to mark packets when the output queue experiences congestion.</p> <p>ECN reduces packet loss by forwarding ECN-capable packets during periods of network congestion instead of dropping those packets. (TCP notifies the network about congestion by dropping packets.) During periods of congestion, ECN marks ECN-capable packets that egress from congested queues. When the receiver receives an ECN packet that is marked as experiencing congestion, the receiver echoes the congestion state back to the sender. The sender then reduces its transmission rate to clear the congestion.</p> |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring ECN on page 6828</a></li> <li>• <a href="#">Understanding CoS Explicit Congestion Notification on page 6820</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## fill-level

---

|                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| List of Syntax                                                   | QFX5100, EX4600, QFX3500, and QFX3600, Switches, QFabric Systems on page 7102<br>QFX10000 Switches on page 7102                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| QFX5100, EX4600, QFX3500, and QFX3600, Switches, QFabric Systems | fill-level <i>low-value</i> fill-level <i>high-value</i> ;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| QFX10000 Switches                                                | fill-level <i>level1 level2 ... level32</i> ;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Hierarchy Level                                                  | [edit <a href="#">class-of-service drop-profiles profile-name interpolate</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Release Information                                              | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Description                                                      | <p>When configuring weighted random early detection (WRED), map the fullness of a queue to a packet <a href="#">drop-probability</a> value. You configure the <b>fill-level</b> and <b>drop-probability</b> statements in related pairs. The pairs of fill level and drop probability values set a probability of dropping packets at a specified queue fullness value.</p> <p>The first fill level is the packet drop start point. Packets do not drop until the queue fullness reaches the first fill level. The last fill level is the packet drop end point. After the queue exceeds the fullness set by the drop end point, all non-ECN packets are dropped. As the queue fills from the drop start point to the drop end point, the rate of packet drop increases in a curve pattern. The higher the queue fill level, the higher the probability of dropping packets.</p> <p>On switches that support only two fill level/drop probability pairs, the two pairs are the drop start point and the drop end point. On switches that support up to 32 fill level/drop probability pairs, you can configure intermediate interpolations between the drop start point and the drop end point, which provides greater flexibility in controlling the packet drop curve.</p> |



**NOTE:** Do not configure the last fill level as 100 percent.

---

|         |                                                                                                                                                                                                                                                                                                                         |
|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Options | <b>low-value</b> (switches that support only two fill level/drop probability pairs)—Fullness of the queue before packets begin to drop, expressed as a percentage. The low value must be less than the high value.<br><br><b>Range:</b> 0 through 100                                                                   |
|         | <b>high-value</b> (switches that support only two fill level/drop probability pairs)—Fullness of the queue before it reaches the maximum drop probability. If the queue fills beyond the fill level high value, all packets drop. The high value must be greater than the low value.<br><br><b>Range:</b> 0 through 100 |

*level1 level2 ... level32* (switches that support 32 fill level/drop probability pairs)—The queue fullness level, expressed as a percentage. Each fill level pairs with a drop probability to define the probability of a packet dropping at a specified queue fullness.

**Range:** 0 through 100

|                           |                                                               |
|---------------------------|---------------------------------------------------------------|
| <b>Required Privilege</b> | interface—To view this statement in the configuration.        |
| <b>Level</b>              | interface-control—To add this statement to the configuration. |

## forwarding-class

---

|                                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                | <a href="#">Classifier on page 7104</a><br><a href="#">Rewrite Rule on page 7104</a><br><a href="#">Scheduler Map on page 7104</a><br><a href="#">Interface on page 7104</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Classifier</b>                    | <pre>forwarding-class class-name {<br/>    loss-priority level {<br/>        code-points [ aliases ] [ bit-patterns ];<br/>    }<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Rewrite Rule</b>                  | <pre>forwarding-class class-name {<br/>    loss-priority level {<br/>        code-point [ aliases ] [ bit-patterns ];<br/>    }<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Scheduler Map</b>                 | <pre>forwarding-class class-name {<br/>    scheduler scheduler-name;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Interface</b>                     | <pre>forwarding-class class-name;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Classifier Hierarchy Level</b>    | [edit <a href="#">class-of-service classifiers</a> (dscp   dscp-ipv6   ieee-802.1   exp) classifier-name],                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Rewrite Rule Hierarchy Level</b>  | [edit <a href="#">class-of-service rewrite-rules</a> ] (dscp   dscp-ipv6   ieee-802.1) rewrite-name   exp],                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Scheduler Map Hierarchy Level</b> | [edit <a href="#">class-of-service scheduler-maps</a> map-name],                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Interface Hierarchy Level</b>     | [edit <a href="#">class-of-service interfaces</a> interface-name unit logical-unit-number]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>           | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>                   | <ul style="list-style-type: none"><li>• Classifiers—Assign incoming traffic to the specified forwarding class based on the specified code point values and assign that traffic the specified loss priority</li><li>• Rewrite rules—At the egress interface, change (rewrite) the value of the code point bits and the loss priority to specified new values for traffic assigned to the specified forwarding class, before forwarding the traffic to the next hop.</li><li>• Scheduler maps—Apply the specified scheduler to the specified forwarding class.</li><li>• Interfaces—Assign the specified forwarding class to the interface to use as a fixed classifier (all incoming traffic on the interface is classified into that forwarding class).</li></ul> |





**NOTE:** OCX Series switches do not support MPLS, so they do not support EXP classifiers or rewrite rules.

**Options** *class-name*—Name of the forwarding class.

The remaining statements are explained separately.

**Required Privilege** interfaces—To view this statement in the configuration.

**Level** interface-control—To add this statement to the configuration.

## forwarding-class-set

**Syntax** `forwarding-class-set forwarding-class-set-name {  
    output-traffic-control-profile profile-name;  
}`

**Hierarchy Level** [edit [class-of-service interfaces](#) *interface-name*]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Apply a previously defined forwarding class set to an output traffic control profile.

**Options** *forwarding-class-set-name*—Name of the forwarding class set.



The remaining statement is explained separately.

**Required Privilege** interfaces—To view this statement in the configuration.

**Level** interface-control—To add this statement to the configuration.

- Related Documentation**
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 6771](#)
  - [Assigning CoS Components to Interfaces on page 6627](#)
  - [Understanding CoS Forwarding Class Sets \(Priority Groups\) on page 6695](#)

## guaranteed-rate

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>guaranteed-rate (rate  percent <i>percentage</i>);</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit <code>class-of-service traffic-control-profiles traffic-control-profile-name</code> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Configure a guaranteed minimum rate of transmission for a traffic control profile. The sum of the guaranteed rates of all of the forwarding class sets (priority groups) on a port should not exceed the total port bandwidth. The guaranteed rate also determines the amount of excess (extra) port bandwidth that the priority group (forwarding class set) can share. Extra port bandwidth is allocated among the priority groups on a port in proportion to the guaranteed rate of each priority group.                                                                                |
|                                 | <p> <b>NOTE:</b> You cannot configure a guaranteed rate for a forwarding class set (priority group) that includes strict-high priority queues. If the traffic control profile is for a forwarding class set that contains strict-high priority queues, do not configure a guaranteed rate.</p>                                                                                                                                                                                                            |
| <b>Default</b>                  | If you do not specify a guaranteed rate, the guaranteed rate is zero (0) and there is no minimum guaranteed bandwidth.                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|                                 | <p> <b>NOTE:</b> If you do not configure a guaranteed rate for a traffic control profile, the queues that belong to any forwarding class set (priority group) that uses that traffic control profile cannot have a configured transmit rate. The result is that there is no minimum guaranteed bandwidth for those queues and that those queues can be starved during periods of congestion.</p>                                                                                                        |
| <b>Options</b>                  | <p><b>percent <i>percentage</i></b>—Minimum percentage of transmission capacity allocated to the forwarding class set or logical interface.<br/><b>Range:</b> 1 through 100 percent</p> <p><b><i>rate</i></b>—Minimum transmission rate allocated to the forwarding class set or logical interface, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000).<br/><b>Range:</b> 1000 through 10,000,000,000 bps</p> |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

- Related Documentation**
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 6771](#)
  - [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 6762](#)
  - [Example: Configuring Minimum Guaranteed Output Bandwidth on page 6798](#)
  - [Understanding CoS Traffic Control Profiles on page 6757](#)
  - [output-traffic-control-profile on page 7108](#)

## interpolate

|                                                                  |                                                                                                                                                                                                           |
|------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                                            | QFX5100, EX4600, QFX3500, and QFX3600, Switches, QFabric Systems on page 7107<br>QFX10000 Switches on page 7107                                                                                           |
| QFX5100, EX4600, QFX3500, and QFX3600, Switches, QFabric Systems | <pre>interpolate {   fill-level low-value fill-level high-value;   drop-probability 0 drop-probability high-value; }</pre>                                                                                |
| QFX10000 Switches                                                | <pre>interpolate {   fill-level level1 level2 ... level32 drop-probability percent1 percent2 ... percent32; }</pre>                                                                                       |
| <b>Hierarchy Level</b>                                           | [edit class-of-service <a href="#">drop-profiles</a> <i>profile-name</i> ]                                                                                                                                |
| <b>Release Information</b>                                       | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                             |
| <b>Description</b>                                               | Specify values for interpolating the relationship between queue fill level and drop probability for weighted random early detection (WRED) drop profiles.<br><br>The statements are explained separately. |
| <b>Required Privilege Level</b>                                  | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                   |

## loss-priority (Drop Profiles)

---

|                                 |                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>loss-priority level protocol protocol drop-profile profile-name;</code>                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service schedulers scheduler-name drop-profile-map</a> ]                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                 |
| <b>Description</b>              | Configure packet loss priority value for a weighted random early detection (WRED) drop profile mapped to a system drop profile.                                                                                                                                                                                               |
| <b>Options</b>                  | <p><i>level</i>—Can be one of the following:</p> <ul style="list-style-type: none"><li>• <b>low</b>—Packet has low loss priority.</li><li>• <b>medium-high</b>—Packet has medium-high loss priority.</li><li>• <b>high</b>—Packet has high loss priority.</li></ul> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                           |

## output-traffic-control-profile

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>output-traffic-control-profile profile-name;</code>                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service interfaces interface-name forwarding-class-set forwarding-class-set-name</a> ]                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Apply an output traffic scheduling and shaping profile to a forwarding class set (priority group).                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <p><i>profile-name</i>—Name of the traffic-control profile to apply to the specified forwarding class set.</p>                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 6771</a></li><li>• <a href="#">Example: Configuring Traffic Control Profiles (Priority Group Scheduling) on page 6762</a></li><li>• <a href="#">Assigning CoS Components to Interfaces on page 6627</a></li><li>• <a href="#">Understanding CoS Traffic Control Profiles on page 6757</a></li></ul> |

## priority (Schedulers)

|                            |                                                                                                                                               |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>priority <i>priority</i>;</code>                                                                                                        |
| <b>Hierarchy Level</b>     | [edit <code>class-of-service schedulers scheduler-name</code> ]                                                                               |
| <b>Release Information</b> | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| <b>Description</b>         | Specify the packet bandwidth-scheduling priority value.                                                                                       |



**NOTE:** On QFabric systems, the `priority` statement is valid only for Node device queue scheduling. The `priority` statement is not allowed for Interconnect device queue scheduling. If you map a scheduler that includes a `priority` configuration to a fabric forwarding class at the [edit `class-of-service scheduler-map-fcset`] hierarchy level, the system generates a commit error. (On the Interconnect device, fabric `fc-sets` are not user-definable. Only the `fabric_fcset_strict_high` fabric `fc-set` is configured with high priority, and this configuration cannot be changed.)

**Options** `priority`—It can be one of the following:

- **low**—Scheduler has low priority.
- **strict-high**—Scheduler has strict high priority. On QFX5100, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems, you can configure only one queue as a strict-high priority queue. On QFX10000 switches, you can configure as many strict-high priority queues as you want. However, because strict-high priority traffic takes precedence over all other traffic, too much strict-high priority traffic can starve the other output queues.

Strict-high priority allocates the scheduled bandwidth to the packets on the queue before any other queue receives bandwidth. Other queues receive the bandwidth that remains after the strict-high queue has been serviced.



**NOTE:** On QFX10000 switches, we strongly recommend that you apply a transmit rate to strict-high priority queues to prevent them from starving other queues. A transmit rate configured on a strict-high priority queue limits the amount of traffic that receives strict-high priority treatment to the amount or percentage set by the transmit rate. The switch treats traffic in excess of the transmit rate as best-effort traffic that receives bandwidth from the leftover (excess) port bandwidth pool. On strict-high priority queues, all traffic that exceeds the transmit rate shares in the port excess bandwidth pool based on the strict-high priority excess bandwidth sharing weight of “1”, which is not configurable. The actual amount of extra bandwidth that traffic exceeding the transmit rate receives depends on

how many other queues consume excess bandwidth and the excess rates of those queues.

On QFX5100, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems, we recommend that you always apply a shaping rate to strict-high priority queues to prevent them from starving other queues. A shaping rate (shaper) sets the maximum amount of bandwidth a queue can consume. (Unlike using the transmit rate on a QFX10000 switch to limit traffic that receives strict-high priority treatment, traffic that exceeds the shaping rate is dropped, and is not treated as best-effort traffic that shares in excess bandwidth.) If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.

---

|                           |                                                               |
|---------------------------|---------------------------------------------------------------|
| <b>Required Privilege</b> | interfaces—To view this statement in the configuration.       |
| <b>Level</b>              | interface-control—To add this statement to the configuration. |

---

## protocol (Drop Profile Map)

---

|                            |                                                                                                                                                                                                        |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>protocol <i>protocol</i> drop-profile <i>profile-name</i>;</code>                                                                                                                                |
| <b>Hierarchy Level</b>     | [edit <code>class-of-service schedulers <i>scheduler-name</i> drop-profile-map loss-priority (low   medium-high   high)</code> ]                                                                       |
| <b>Release Information</b> | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                          |
| <b>Description</b>         | Configure the protocol type for the specified weighted random early detection (WRED) drop profile.                                                                                                     |
| <b>Options</b>             | <b><i>protocol</i></b> —Type of protocol. The protocol can be: <ul style="list-style-type: none"><li>• <b>any</b>—Accept any protocol type.</li></ul> The remaining statement is explained separately. |
| <b>Required Privilege</b>  | interfaces—To view this statement in the configuration.                                                                                                                                                |
| <b>Level</b>               | interface-control—To add this statement to the configuration.                                                                                                                                          |

## scheduler

|                            |                                                                                                                                               |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>scheduler <i>scheduler-name</i>;</code>                                                                                                 |
| <b>Hierarchy Level</b>     | [edit <code>class-of-service scheduler-maps <i>map-name</i> forwarding-class <i>class-name</i></code> ]                                       |
| <b>Release Information</b> | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| <b>Description</b>         | Map a scheduler to a forwarding class using a scheduler map.                                                                                  |



**NOTE:** On QFX5200 only, absolute CoS rate limits for transmit rate and shaping rate do not reflect 50g and 100g interfaces. Therefore this statement does not affect those interfaces for QFX5200 in release 15.1X53-D30.

|                                 |                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | <code><i>scheduler-name</i></code> —Name of the scheduler to map to the forwarding class.                                |
| <b>Required Privilege Level</b> | interfaces—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration. |

## scheduler-map

|                                                                      |                                                                                                                                               |
|----------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                        | <code>scheduler-map <i>map-name</i>;</code>                                                                                                   |
| <b>Enhanced Transmission Selection (ETS) Hierarchical Scheduling</b> | [edit <code>class-of-service traffic-control-profiles <i>traffic-control-profile-name</i></code> ]                                            |
| <b>Port Scheduling</b>                                               | [edit <code>class-of-service interfaces <i>interface-name</i></code> ]                                                                        |
| <b>Release Information</b>                                           | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| <b>Description</b>                                                   | Associate a scheduler map with a traffic control profile.                                                                                     |
| <b>Options</b>                                                       | <code><i>map-name</i></code> —Name of the scheduler map.                                                                                      |
| <b>Required Privilege Level</b>                                      | interfaces—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                      |

## scheduler-maps

---

|                                 |                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>scheduler-maps {<br/>  map-name {<br/>    forwarding-class class-name scheduler scheduler-name;<br/>  }<br/>}</pre>                      |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service</a> ]                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| <b>Description</b>              | Specify a scheduler map name to map a scheduler configuration to a forwarding class.                                                          |
| <b>Options</b>                  | <p><i>map-name</i>—Name of the scheduler map.</p> <p>The remaining statements are explained separately.</p>                                   |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>           |



## schedulers

|                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| List of Syntax                                                   | QFX5100, EX4600, QFX3500, and QFX3600, Switches, QFabric Systems on page 7113<br>QFX10000 Switches on page 7113                                                                                                                                                                                                                                                                                                                                 |
| QFX5100, EX4600, QFX3500, and QFX3600, Switches, QFabric Systems | <pre> schedulers {   scheduler-name {     buffer-size (percent <i>percentage</i>   remainder);     drop-profile-map loss-priority (low   medium-high   high) protocol <i>protocol</i> drop-profile       drop-profile-name;     explicit-congestion-notification;     priority <i>priority</i>;     shaping-rate (rate   percent <i>percentage</i>);     transmit-rate (percent <i>percentage</i>);   } } </pre>                                |
| QFX10000 Switches                                                | <pre> schedulers {   scheduler-name {     buffer-size (percent <i>percentage</i>   remainder);     drop-profile-map loss-priority (low   medium-high   high) protocol <i>protocol</i> drop-profile       drop-profile-name;     excess-rate;     explicit-congestion-notification;     priority <i>priority</i>;     shaping-rate (rate   percent <i>percentage</i>);     transmit-rate (percent <i>percentage</i>) &lt;exact&gt;;   } } </pre> |
| Hierarchy Level                                                  | [edit <a href="#">class-of-service</a> ]                                                                                                                                                                                                                                                                                                                                                                                                        |
| Release Information                                              | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                   |
| Description                                                      | Specify scheduler name and parameter values such as minimum bandwidth ( <a href="#">transmit-rate</a> ), maximum bandwidth ( <a href="#">shaping-rate</a> ), and priority ( <a href="#">priority</a> ).                                                                                                                                                                                                                                         |
| Options                                                          | <p><b><i>scheduler-name</i></b> —Name of the scheduler.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                           |
| Required Privilege Level                                         | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                             |

## shaping-rate

**Syntax** `shaping-rate (rate | percent percentage);`

**Hierarchy Level** `[edit class-of-service schedulers scheduler-name],`  
`[edit class-of-service traffic-control-profiles profile-name]`



**NOTE:** Only switches that support enhanced transmission selection (ETS) hierarchical scheduling support the **traffic-control-profiles** hierarchy.

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Configure the shaping rate. The shaping rate throttles the rate of packet transmission by setting a maximum bandwidth (rate in bits per second) or a maximum percentage of bandwidth for a queue or a forwarding class set. You specify the maximum bandwidth for a queue by using a scheduler map to associate a forwarding class (queue) with a scheduler that has a configured shaping rate.

For ETS configuration, you specify the maximum bandwidth for a forwarding class set by setting the shaping rate for a traffic control profile, then you associate the scheduler map with the traffic control profile, and then you apply the traffic control profile and a forwarding class set to an interface.

For simple port scheduling configuration, you apply the scheduler map directly to an interface (instead of indirectly through the traffic control profile as in ETS).

We recommend that you configure the shaping rate as an absolute maximum usage and not as additional usage beyond the configured transmit rate (the minimum guaranteed bandwidth for a queue) or the configured guaranteed rate (the minimum guaranteed bandwidth for a forwarding class set).



**NOTE:** When you set the maximum bandwidth (**shaping-rate** value) for a queue or for a priority group at 100 Kbps or less, the traffic shaping behavior is accurate only within +/- 20 percent of the configured **shaping-rate** value.



**NOTE:** On QFX5200, QFX5100, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems, we recommend that you always apply a shaping rate to strict-high priority queues to prevent them from starving other queues. If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.



**NOTE:** On QFX5200 Series switches, a granularity of 64kbps is supported for the shaping rate. Therefore, the shaping rate on queues for 100g interfaces might not be applied correctly.



**NOTE:** QFX10000 Series switches do not support the shaping-rate statement. However, you can configure the transmit-rate exact option to prevent a queue from consuming more bandwidth than you want the queue to consume.

On QFX10000 Series switches, we recommend that you use the transmit rate to set a limit on the amount of bandwidth that receives strict-high priority treatment on a strict-high priority queue. Traffic up to the transmit rate receives strict-high priority treatment. Traffic in excess of the transmit rate is treated as best-effort traffic that receives the strict-high priority queue excess rate weight of “1”. Do not use a shaping rate to set a maximum bandwidth limit on strict-high priority queues on QFX10000 Series switches.

**Default** If you do not configure a shaping rate, the default shaping rate is 100 percent (all of the available bandwidth), which is the equivalent of no rate shaping.

**Options** **percent *percentage***—Shaping rate as a percentage of the available interface bandwidth.  
**Range:** 1 through 100 percent

**rate**—Peak (maximum) rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000).

**Range:** 1000 through 10,000,000,000 bps

**Required Privilege Level** interfaces—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

- Related Documentation**
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 6771](#)
  - [Example: Configuring Queue Schedulers](#)
  - [Example: Configuring Queue Schedulers for Port Scheduling on page 6744](#)
  - [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 6762](#)
  - [Understanding CoS Output Queue Schedulers](#)
  - [Understanding CoS Port Schedulers on QFX Switches on page 6729](#)
  - [Understanding CoS Traffic Control Profiles on page 6757](#)

## traffic-control-profiles

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>traffic-control-profiles <i>profile-name</i> {<br/>    <b>guaranteed-rate</b> (<i>rate</i>  percent <i>percentage</i>);<br/>    <b>scheduler-map</b> <i>map-name</i>;<br/>    <b>shaping-rate</b> (<i>rate</i>  percent <i>percentage</i>);<br/>}</pre>                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit <b>class-of-service</b> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Configure traffic shaping and scheduling profiles for forwarding class sets (priority groups) to implement enhanced transmission selection (ETS) or for logical interfaces.                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <p><b>profile-name</b>—Name of the traffic-control profile. This name is also used to specify an output traffic control profile.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | interfaces—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 6771</a></li><li>• <a href="#">Example: Configuring Traffic Control Profiles (Priority Group Scheduling) on page 6762</a></li><li>• <a href="#">Example: Configuring Forwarding Class Sets on page 6697</a></li><li>• <a href="#">Assigning CoS Components to Interfaces on page 6627</a></li><li>• <a href="#">output-traffic-control-profile on page 7108</a></li><li>• <a href="#">Understanding CoS Traffic Control Profiles on page 6757</a></li></ul> |

## transmit-rate

|                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                                            | QFX5100, EX4600, QFX3500, and QFX3600, Switches, QFabric Systems on page 7117<br>QFX10000 Switches on page 7117                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| QFX5100, EX4600, QFX3500, and QFX3600, Switches, QFabric Systems | <code>transmit-rate (rate   percent <i>percentage</i>);</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| QFX10000 Switches                                                | <code>transmit-rate (rate   percent <i>percentage</i>) &lt;exact&gt;;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>                                           | [edit <a href="#">class-of-service schedulers scheduler-name</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>                                       | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.<br>Exact option introduced in Junos OS Release 15.1X53-D10 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>                                               | <p>On QFX5100, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems, the transmit rate specifies the minimum guaranteed transmission rate or percentage for a queue (forwarding class) scheduler. The queue transmit rate also determines the amount of excess (extra) priority group bandwidth that the queue can share on switches that support enhanced transmission selection (ETS) hierarchical scheduling.</p> <p>On QFX10000 switches, the transmit rate specifies the minimum guaranteed transmission rate or percentage for a queue (forwarding class) scheduler. The queue transmit rate also determines the amount of excess (extra) port bandwidth the queue can share if you do not explicitly configure an excess rate in the scheduler. The transmit rate also determines the amount of excess (extra) priority group bandwidth that the queue can share on switches that support enhanced transmission selection (ETS) hierarchical scheduling.</p> <p>On QFX10000 switch strict-high priority queues, the transmit rate limits the amount of traffic the switch treats as strict-high priority traffic. Traffic up to the transmit rate receives strict-high priority treatment. The switch treats traffic that exceeds the transmit rate as best-effort traffic that receives an excess bandwidth sharing weight of “1”; you cannot configure an excess rate on a strict-high priority queue, and unlike queues with other scheduling priorities, the switch does not use the transmit rate to determine extra bandwidth sharing for strict-high priority queues.</p> |



**CAUTION:** We strongly recommend that you configure a transmit rate on strict-high priority queues to limit the amount of traffic the switch treats as strict-high priority traffic on those queues. This is especially important if you configure more than one strict-high priority queue on a port. To prevent a strict-high priority queue from starving the other queues on a port, we recommend that you always configure a transmit rate, even if you only configure one strict-high priority queue.



NOTE: For ETS, the transmit-rate setting works only if you also configure the **guaranteed-rate** in the traffic control profile that is attached to the forwarding class set to which the queue belongs. If you do not configure the guaranteed rate, the minimum guaranteed rate for individual queues that you set using the transmit-rate statement does not work. The sum of all queue transmit rates in a forwarding class set should not exceed the traffic control profile guaranteed rate.



NOTE: On QFX5100, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems, you cannot configure a transmit rate for a strict-high priority queue. Queues (forwarding classes) with a configured transmit rate cannot be included in a forwarding class set that has a strict-high priority queue. To prevent strict-high priority queues from consuming all of the available bandwidth on these switches, we recommend that you configure a shaping rate to set a maximum amount of bandwidth for strict-high priority queues.



NOTE: For transmit rates below 1 Gbps, we recommend that you configure the transmit rate as a percentage instead of as a fixed rate. This is because the system converts fixed rates into percentages and may round small fixed rates to a lower percentage. For example, a fixed rate of 350 Mbps is rounded down to 3 percent instead of 3.5 percent.

**Default** On QFX5100, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems, if you do not configure the transmit rate, the default scheduler transmission rate and buffer size percentages for queues 0 through 11 are:

**Table 605: Default Transmit Rates for QFX5100, EX4600, QFX3500, and QFX3600 Switches, and QFabric Systems**

| Queue Number        | Default Minimum Guaranteed Bandwidth (Transmit Rate) |
|---------------------|------------------------------------------------------|
| 0 (best-effort)     | 5 %                                                  |
| 1                   | 0                                                    |
| 2                   | 0                                                    |
| 3 (fcoe)            | 35 %                                                 |
| 4 (no-loss)         | 35 %                                                 |
| 5                   | 0                                                    |
| 6                   | 0                                                    |
| 7 (network control) | 5 %                                                  |
| 8 (mcast)           | 20 %                                                 |
| 9                   | 0                                                    |
| 10                  | 0                                                    |
| 11                  | 0                                                    |



**NOTE:** OCX Series switches do not support lossless transport. The OCX Series default DSCP classifier does not classify traffic into the default lossless fcoe and no-loss forwarding classes. The bandwidth that the default scheduler allocates to the default fcoe and no-loss forwarding classes on other switches is allocated to the default best-effort, network-control, and mcast forwarding classes on OCX Series switches.

On QFX10000 switches, if you do not configure the transmit rate, the default scheduler transmission rate and buffer size percentages for queues 0 through 7 are:

Table 606: Default Transmit Rates for QFX10000 Switches

| Queue Number        | Default Minimum Guaranteed Bandwidth (Transmit Rate) |
|---------------------|------------------------------------------------------|
| 0 (best-effort)     | 15 %                                                 |
| 1                   | 0                                                    |
| 2                   | 0                                                    |
| 3 (fcoe)            | 35 %                                                 |
| 4 (no-loss)         | 35 %                                                 |
| 5                   | 0                                                    |
| 6                   | 0                                                    |
| 7 (network control) | 15 %                                                 |

Configure schedulers if you want to change the minimum guaranteed bandwidth and other queue characteristics.

**Options** **rate**—Minimum transmission rate for the queue, in bps. You can specify a value in bits-per-second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

**Range:** 1000 through 10,000,000,000 bps on 10-Gigabit interfaces, 1000 through 40,000,000,000 bps on 40-Gigabit interfaces.

**percent** **percentage**—Minimum percentage of transmission capacity allocated to the queue. A percentage of zero means that there is no minimum bandwidth guarantee for the queue.

**Range:** 0 through 100 percent

**exact**—(QFX10000 switches only) Shape queues that are not strict-high priority queues to the transmit rate so that the transmit rate is the maximum bandwidth limit. Traffic that exceeds the exact transmit rate is dropped. You cannot set an excess rate on queues configured as **transmit-rate (rate | percentage) exact** because the purpose of setting an exact transmit rate is to set a maximum bandwidth (shaping rate) on the traffic.

**Required Privilege Level** interfaces—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.



# Configuration Statements (Data Center Bridging and Flow Control)

- [application \(Application Maps\) on page 7122](#)
- [application \(Applications\) on page 7123](#)
- [application-map on page 7124](#)
- [application-maps on page 7125](#)
- [applications \(Applications\) on page 7126](#)
- [applications \(DCBX\) on page 7127](#)
- [cable-length \(Congestion Notification\) on page 7128](#)
- [code-point \(Input Congestion Notification\) on page 7129](#)
- [code-point \(Output Congestion Notification\) on page 7130](#)
- [code-points \(Application Maps\) on page 7131](#)
- [congestion-notification-profile on page 7132](#)
- [dcbx on page 7134](#)
- [dcbx-version on page 7135](#)
- [destination-port \(Applications\) on page 7136](#)
- [disable \(DCBX\) on page 7137](#)
- [enhanced-transmission-selection on page 7138](#)
- [ether-type on page 7139](#)
- [flow-control on page 7140](#)
- [flow-control-queue \(Output Congestion Notification\) on page 7142](#)
- [ieee-802.1 \(Input Congestion Notification\) on page 7143](#)
- [ieee-802.1 \(Output Congestion Notification\) on page 7144](#)
- [input \(Congestion Notification\) on page 7145](#)
- [interface \(DCBX\) on page 7146](#)
- [interfaces \(Class of Service\) on page 7147](#)
- [mru on page 7149](#)
- [output \(Congestion Notification\) on page 7150](#)

- [pfc \(Input Congestion Notification\) on page 7151](#)
- [policy-options on page 7152](#)
- [priority-flow-control on page 7153](#)
- [protocol \(Applications\) on page 7154](#)
- [recommendation-tlv on page 7155](#)
- [rx-buffers on page 7156](#)
- [tx-buffers on page 7158](#)

---

## application (Application Maps)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>application <i>application-name</i> {<br/>    <a href="#">code-points</a> [ <i>aliases</i> ] [ <i>bit-patterns</i> ];<br/>}</code>                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit policy-options <a href="#">application-maps</a> <i>application-map-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Add an application to an application map and define the application's code points.                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <i>application-name</i> —Name of the application.<br><br>The remaining statement is explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 6533</a></li><li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 6535</a></li><li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 6528</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li></ul> |

## application (Applications)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> application <i>application-name</i> {     <i>destination-port</i> <i>port-value</i>;     <i>protocol</i> (tcp   udp);     <i>ether-type</i> <i>type</i>; } </pre>                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit applications]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Configure properties to define an application.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <p><i>application-name</i>—Name of the application.</p> <p>The statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Defining an Application for DCBX Application Protocol TLV Exchange on page 6532</a></li> <li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 6535</a></li> <li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 6528</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li> </ul> |

## application-map

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>application-map <i>application-map-name</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit <a href="#">protocols dcbx interface interface-name</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Specify an application map to apply to an interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <i>application-map-name</i> —Name of the application map.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show dcbx neighbors on page 6581</a></li><li>• <a href="#">Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 6534</a></li><li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 6535</a></li><li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 6528</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li></ul> |

## application-maps

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> application-maps <i>application-map-name</i> {   application <i>application-name</i> {     code-points [ <i>aliases</i> ] [ <i>bit-patterns</i> ];   } } </pre>                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit policy-options]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Define an application map by specifying the applications that belong to the application map.                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <p><b><i>application-map-name</i></b>—Name of the application map.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 6533</a></li> <li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 6535</a></li> <li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 6528</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li> </ul> |

## applications (Applications)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>applications {<br/>  application application-name {<br/>    destination-port port-value;<br/>    protocol (tcp   udp);<br/>    ether-type type;<br/>  }<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Define applications that DCBX advertises.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | The statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Defining an Application for DCBX Application Protocol TLV Exchange on page 6532</a></li><li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 6535</a></li><li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 6528</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li></ul> |


## applications (DCBX)

---

|                                 |                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>applications {<br/>    fcoe {<br/>        no-auto-negotiation;<br/>    }<br/>}</pre>                                                                                             |
| <b>Hierarchy Level</b>          | [edit <a href="#">protocols dcbx interface interface-name</a> ]                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 12.1 for the EX Series                                                  |
| <b>Description</b>              | Configure Data Center Bridging Capability Exchange protocol (DCBX) applications on an interface.                                                                                      |
| <b>Options</b>                  | The remaining statements are explained separately.                                                                                                                                    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show dcbx neighbors on page 6581</a></li><li>• <a href="#">Understanding DCB Features and Requirements on page 6490</a></li></ul> |

## cable-length (Congestion Notification)

---

|                                                                                                                                                                                                                                                                                                                                             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                               | <code>cable-length <i>cable-length-value</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                      | [edit <a href="#">class-of-service congestion-notification-profile <i>profile-name</i> input</a> ]                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                  | Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                          | <p>Specify the length of the cable between the interface and its peer interface in meters. The system uses the cable length and the maximum receive unit (MRU) to calculate the amount of buffer headroom reserved to support priority-based flow control (PFC). The shorter the cable length and lower the MRU, the less headroom buffer space is required for PFC.</p>                                                                                                                                 |
| <div> <b>NOTE:</b> You can also set a maximum transmission unit (MTU) value (the largest packet size the interface sends) for interfaces by including the <code>mtu</code> statement at the [edit interfaces <i>interface-name</i>] hierarchy level.</div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Default</b>                                                                                                                                                                                                                                                                                                                              | The default cable length value is 100 meters (approximately 328 feet).                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                              | <code><i>cable-length-value</i></code> —Length of the cable in meters. (Generally from 1 to 300 meters, but there is no configuration restriction.)                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                             | <code>interfaces</code> —To view this statement in the configuration.<br><code>interface-control</code> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                | <ul style="list-style-type: none"><li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 6905</a></li><li>• <a href="#">Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 6986</a></li><li>• <a href="#">Understanding CoS Flow Control (Ethernet PAUSE and PFC) on page 6493</a></li><li>• <a href="#">Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 6873</a></li></ul> |



## code-point (Input Congestion Notification)

---

|                                 |                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>code-point [<i>code-point-bits</i>] {<br/>    pfc {<br/>        mru <i>mru-value</i>;<br/>    }<br/>}</pre>                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit <b>class-of-service congestion-notification-profile</b> <i>profile-name</i> <b>input ieee-802.1</b> ]                                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Enable priority-based flow control (PFC) on an IEEE 802.1p code point (priority).                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <p><b><i>code-point-bits</i></b>—3-bit value in decimal form.</p> <p>The remaining statements are described separately.</p>                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring CoS PFC for FCoE Traffic on page 6504</a></li><li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 6905</a></li><li>• <a href="#">Understanding CoS Flow Control (Ethernet PAUSE and PFC) on page 6493</a></li></ul> |

## code-point (Output Congestion Notification)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>code-point [ <i>code-point-bits</i> ] {<br/>    <i>flow-control-queue</i> [ <i>queue</i>   <i>list-of-queues</i> ];<br/>}</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit <i>class-of-service congestion-notification-profile profile-name output ieee-802.1</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Specify the IEEE 802.1p code point bits that identify the traffic you want to enable for priority-based flow control (PFC) pause.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Default</b>                  | <p>By default, IEEE 802.1p priorities 3 and 4 (code points 011 and 100, respectively) are enabled for PFC pause on all Ethernet interfaces. If you explicitly configure priorities to pause and the output queues on which to enable pause, the explicit configuration overrides the default configuration. When you apply an explicit output congestion notification profile to an interface, only the priorities and queues specified in the output congestion notification profile are enabled for pause on that interface.</p> <p>For example, if you configure an output congestion notification profile that specifies priority 2 (code point 010), then traffic with IEEE 802.1p priority 2 is paused on the configured output queue during periods of congestion. However, traffic with priority 3 and priority 4 is not programmed to pause, because the explicit configuration overwrites the default configuration, and the explicit configuration does not pause priority 3 and priority 4. If you configure an explicit output congestion notification profile, all of the priorities you want to enable for PFC and all of the output queues you want to pause must be explicitly configured.</p> |
| <b>Options</b>                  | <p><i>code-point-bits</i>—3-bit value in decimal form.</p> <p>The remaining statements are described separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 6905</a></li><li>• <a href="#">Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 6986</a></li><li>• <a href="#">Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 6977</a></li><li>• <a href="#">Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch) on page 6969</a></li><li>• <a href="#">Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI) on page 7000</a></li><li>• <a href="#">Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 6873</a></li></ul>                                                                                                                                                                                                                                                                                             |

## code-points (Application Maps)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>code-points [ <i>aliases</i> ] [ <i>bit-patterns</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit policy-options <b>application-maps</b> <i>application-map-name</i> <b>application</b> <i>application-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Define one or more code-point aliases or bit sets for an application.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <i>aliases</i> —Name of the alias or aliases.<br><br><i>bit-patterns</i> —Value of the code-point bits, in decimal form.                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 6533</a></li> <li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 6535</a></li> <li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 6528</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li> </ul> |

## congestion-notification-profile

|                                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                | <pre> congestion-notification-profile <i>profile-name</i> {   input {     ieee-802.1 {       code-point [<i>code-point-bits</i>] {         pfc {           mru <i>mru-value</i>;         }       }     }     cable-length <i>cable-length-value</i>;   }   output {     ieee-802.1 {       code-point [<i>code-point-bits</i>] {         flow-control-queue [<i>queue</i>   <i>list-of-queues</i>];       }     }   } } </pre> |
| <b>Interface Congestion Notification Profile Association</b> | <pre> congestion-notification-profile <i>profile-name</i> { </pre>                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>                                       | <pre> [edit <i>class-of-service</i>], [edit <i>class-of-service interfaces interface-name</i>] </pre>                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>                                   | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>                                           | Configure a congestion notification profile to enable priority-based flow control (PFC) on traffic specified by an IEEE 802.1 code point, and apply the profile to an interface.                                                                                                                                                                                                                                               |



**NOTE:** You must configure PFC for FCoE traffic. Each interface that carries FCoE traffic should be configured for PFC on the FCoE code point (usually 011).

You can attach a maximum of one congestion notification profile to an interface. There is no limit to the total number of congestion notification profiles you can create.



**NOTE:** Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

**Options**    *profile-name*—Name of the congestion notification profile.

The remaining statements are explained separately.

**Required Privilege**    interface—To view this statement in the configuration.  
**Level**    interface-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 6905](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 6504](#)
- [Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 6986](#)
- [Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 6977](#)
- [Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic \(FCoE Transit Switch\) on page 6969](#)
- [Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications \(FCoE and iSCSI\) on page 7000](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 6493](#)
- [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 6873](#)

## dcbx

---

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <pre>dcbx {<br/>  disable;<br/>  interface (interface-name   all) {<br/>    disable;<br/>    application-map application-map-name;<br/>    applications {<br/>      no-auto-negotiation;<br/>    }<br/>    enhanced-transmission-selection {<br/>      no-auto-negotiation;<br/>      no-recommendation-tlv;<br/>      recommendation-tlv {<br/>        no-auto-negotiation;<br/>      }<br/>    }<br/>    dcbx-version (auto-negotiate   ieee-dcbx   dcbx-version-1.01);<br/>    priority-flow-control {<br/>      no-auto-negotiation;<br/>    }<br/>  }<br/>}</pre> |
| Hierarchy Level          | [edit <a href="#">protocols</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Release Information      | <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 11.3 for EX Series switches.</p> <p><b>mode</b> and <b>recommendation-tlv</b> statements introduced in Junos OS Release 12.2 for the QFX Series.</p>                                                                                                                                                                                                                                                                                              |
| Description              | <p>Configure DCBX properties. DCBX is an extension of Link Layer Discovery Protocol (LLDP), and LLDP must remain enabled on every interface for which you want to use DCBX. If you attempt to enable DCBX on an interface on which LLDP is disabled, the configuration commit fails.</p>                                                                                                                                                                                                                                                                               |
| Options                  | <p>The statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Related Documentation    | <ul style="list-style-type: none"><li>• <a href="#">show dcbx neighbors on page 6581</a></li><li>• <a href="#">Understanding DCB Features and Requirements on page 6490</a></li><li>• <a href="#">Configuring DCBX Autonegotiation on page 6524</a></li><li>• <i>Understanding DCB Features and Requirements on EX Series Switches</i></li><li>• <i>Disabling DCBX to Disable PFC Autonegotiation on EX Series Switches (CLI Procedure)</i></li></ul>                                                                                                                  |


## dcbx-version

---

|                                 |                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>dcbx-version (auto-negotiate   ieee-dcbx   dcbx-version-1.01);</code>                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit <a href="#">protocols dcbx interface</a> (all   <i>interface-name</i> )]                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | <p>Set the DCBX version for the specified interface or interfaces.</p> <p>QFX3500 switches come up in IEEE DCBX mode and then autonegotiate with the connected peer to set the DCBX version.</p> <p>QFabric system Node devices come up using DCBX version 1.01, and then autonegotiate with the connected peer to set the DCBX mode.</p> |
| <b>Default</b>                  | The default DCBX mode is autonegotiation.                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><b>auto-negotiate</b>—Automatically negotiate the DCBX version with the connected peer.</p> <p><b>ieee-dcbx</b>—Force the interface to use IEEE DCBX mode, regardless of the peer configuration.</p> <p><b>dcbx-version-1.01</b>—Force the interface to use version 1.01 DCBX mode, regardless of the peer configuration.</p>          |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show dcbx neighbors on page 6581</a></li> <li>• <a href="#">Configuring DCBX Autonegotiation on page 6524</a></li> <li>• <a href="#">Understanding DCBX on page 6514</a></li> </ul>                                                                                                  |

## destination-port (Applications)

---

|                                                                                                                                                                                                                                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                             | <code>destination-port <i>port-value</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                    | [edit applications <b>application</b> <i>application-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>                                                                                                                                                                                                                | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>                                                                                                                                                                                                                        | <p>Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) destination port number, which combines with <b>protocol</b> to identify an application type. The Internet Assigned Numbers Authority (IANA) assigns port numbers. See the IANA <i>Service Name and Transport Protocol Port Number Registry</i> at <a href="http://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xml">http://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xml</a> for a list of assigned port numbers.</p> |
| <hr/>                                                                                                                                                                                                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <div> <b>NOTE:</b> To create an application for iSCSI, use the protocol <code>tcp</code> with the destination port number <code>3260</code>.</div> <hr/> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                                                                                                                                                                                                                            | <i>port-value</i> —Identifier for the port.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b>                                                                                                                                                                                                           | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>                                                                                                                                                                                                              | <ul style="list-style-type: none"><li>• <a href="#">Defining an Application for DCBX Application Protocol TLV Exchange on page 6532</a></li><li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 6535</a></li><li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 6528</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li></ul>                            |



## disable (DCBX)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | disable                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit <a href="#">protocols dcbx</a> ]<br><br>[edit <a href="#">protocols dcbx interface</a> <i>interface-name</i> ]                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 11.3 for EX Series switches.                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Disable Data Center Bridging Capability Exchange protocol (DCBX) on one or more 10-Gigabit Ethernet interfaces.                                                                                                                                                                                                                                                                                 |
| <b>Default</b>                  | DCBX is enabled by default on all 10-Gigabit or higher Ethernet interfaces.<br><br>DCBX is enabled by default on all 10-Gigabit Ethernet interfaces on EX4500 CEE-enabled switches.                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring DCBX Autonegotiation on page 6524</a></li> <li>• <i>Disabling DCBX to Disable PFC Autonegotiation on EX Series Switches (CLI Procedure)</i></li> <li>• <a href="#">Understanding DCB Features and Requirements on page 6490</a></li> <li>• <i>Understanding DCB Features and Requirements on EX Series Switches</i></li> </ul> |

## enhanced-transmission-selection

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>enhanced-transmission-selection {<br/>  no-auto-negotiation;<br/>  no-recommendation-tlv;<br/>  recommendation-tlv {<br/>    no-auto-negotiation;<br/>  }<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit <a href="#">protocols dcbx interface interface-name</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | <p>Disable advertising the enhanced transmission selection (ETS) state of the interface to the peer. To disable ETS on the interface, do not enable ETS on the interface in the class-of-service (CoS) configuration.</p> <p>Disabling ETS autonegotiation stops the QFX Series from advertising the ETS Configuration TLV and the ETS Recommendation TLV.</p> <p>Disabling the ETS recommendation TLV stops the QFX Series from advertising the ETS Recommendation TLV, but the ETS Configuration TLV is still advertised.</p> |
| <b>Options</b>                  | <p><b>no-auto-negotiation</b>—Disable automatic negotiation of ETS (Configuration TLV and Recommendation TLV)</p> <p><b>no-recommendation-tlv</b>—Disable automatic negotiation of the ETS Recommendation TLV</p> <p><b>recommendation-tlv</b>—Enable automatic negotiation of ETS Recommendation TLV</p>                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show dcbx neighbors on page 6581</a></li><li>• <a href="#">Configuring DCBX Autonegotiation on page 6524</a></li><li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 6771</a></li><li>• <a href="#">Understanding DCB Features and Requirements on page 6490</a></li></ul>                                                                                                                                                                |

## ether-type

---

|                            |                                                                                                                                                                                                                                                                                                                                |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>ether-type <i>ether-type</i>;</code>                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>     | [edit applications <a href="#">application</a> <i>application-name</i> ]                                                                                                                                                                                                                                                       |
| <b>Release Information</b> | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                     |
| <b>Description</b>         | Two-octet field in an Ethernet frame that defines the protocol encapsulated in the frame payload. See <a href="http://standards.ieee.org/develop/regauth/ethertype/eth.txt">http://standards.ieee.org/develop/regauth/ethertype/eth.txt</a> for a list of Institute of Electrical and Electronics Engineers (IEEE) EtherTypes. |




**NOTE:** To create a FIP application, use the EtherType 0x8914.

---

|                                 |                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | <i>type</i> —Identifier for the EtherType.                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Defining an Application for DCBX Application Protocol TLV Exchange on page 6532</a></li> <li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 6535</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 6528</a></li> </ul> |

## flow-control

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | (flow-control   no-flow-control);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">ether-options</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | <p>Explicitly enable or disable symmetric Ethernet PAUSE flow control, which regulates the flow of packets from the switch to the remote side of the connection by pausing all traffic flows on a link during periods of network congestion. Symmetric flow control means that Ethernet PAUSE is enabled in both directions. The interface generates and sends Ethernet PAUSE messages when the receive buffers fill to a certain threshold and the interface responds to PAUSE messages received from the connected peer. By default, flow control is disabled.</p> <p>You can configure asymmetric flow control by including the <b>configured-flow-control</b> statement at the [edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">ether-options</a> hierarchy level. Symmetric flow control and asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.</p> <div><div></div><div><p><b>NOTE:</b> Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual “lanes” of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC).</p><p>Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.</p><p>OCX Series switches do not support PFC.</p></div></div> <div><ul style="list-style-type: none"><li>• <b>flow-control</b>—Enable flow control; flow control is useful when the remote device is a Gigabit Ethernet switch.</li><li>• <b>no-flow-control</b>—Disable flow control.</li></ul></div> |
| <b>Default</b>                  | Flow control is disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">configured-flow-control on page 2980</a></li><li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2826</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

- *Configuring Gigabit and 10-Gigabit Ethernet Interfaces*
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 6493](#)
- *Junos OS Network Interfaces Library for Routing Devices*

## flow-control-queue (Output Congestion Notification)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>flow-control-queue [ <i>queue</i>   <i>list-of-queues</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service congestion-notification-profile <i>profile-name</i> output ieee-802.1 code-point <i>code-point-bits</i></a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | <p>Specify one or more output queues to pause, to support priority-based flow control (PFC). The specified queues pause when the interface receives a PFC frame with a matching IEEE 802.1p code point.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Default</b>                  | <p>Queue 3 (mapped to the fcoe forwarding class) and queue 4 (mapped to the no-loss forwarding class) are programmed as flow control queues to pause. No other output queues are programmed to pause by default.</p> <p>If you configure flow control queues explicitly, only the queues that you specify are programmed to pause. The explicit flow control queue to pause configuration overrides the default setting, so the queues paused in the default configuration are no longer paused by default.</p> <p>For example, if you configure queue 2 as a flow control queue, then queue 2 pauses when congestion occurs, but queues 3 and 4 do not pause because they were not explicitly specified. To enable pause on output queues 2, 3, and 4, you must explicitly configure all three of the queues as flow control queues.</p> <p>The same behavior applies to the IEEE 802.1p code points (priorities) on which PFC is enabled. By default, priorities 3 (011) and 4 (100) are enabled for PFC pause. If you explicitly configure flow control queues to pause, you must also explicitly configure pause for each priority (code point) that you want to pause, because the explicit configuration overrides the default configuration.</p> |
| <b>Options</b>                  | <code>[ <i>queue</i>   <i>list-of-queues</i> ]</code> —The output queue or a list of output queues to pause.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 6905</a></li><li>• <a href="#">Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 6986</a></li><li>• <a href="#">Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 6977</a></li><li>• <a href="#">Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch) on page 6969</a></li><li>• <a href="#">Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI) on page 7000</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                |

- [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 6873](#)

## ieee-802.1 (Input Congestion Notification)

|                                 |                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> ieee-802.1 {   code-point [code-point-bits] {     pfc {       mru mru-value;     }   } } </pre>                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit <b>class-of-service congestion-notification-profile</b> <i>profile-name</i> <b>input</b> ]                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Configure an IEEE 802.1 code point and apply priority-based flow control (PFC) to packets with that code point.                                                                                                                                                                                                               |
| <b>Options</b>                  | The statements are described separately.                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring CoS PFC for FCoE Traffic on page 6504</a></li> <li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 6905</a></li> <li>• <a href="#">Understanding CoS Flow Control (Ethernet PAUSE and PFC) on page 6493</a></li> </ul> |

## ieee-802.1 (Output Congestion Notification)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>ieee-802.1 {<br/>    code-point [ code-point-bits ] {<br/>        flow-control-queue [ queue   list-of-queues ];<br/>    }<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit <b>class-of-service congestion-notification-profile</b> <i>profile-name</i> <b>output</b> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Configure an IEEE 802.1 code point and apply priority-based flow control (PFC) to packets with that code point on output queues.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | The statements are described separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 6905</a></li><li>• <a href="#">Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 6986</a></li><li>• <a href="#">Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 6977</a></li><li>• <a href="#">Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch) on page 6969</a></li><li>• <a href="#">Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI) on page 7000</a></li><li>• <a href="#">Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 6873</a></li></ul> |



## input (Congestion Notification)

```
Syntax input {
 ieee-802.1 {
 code-point [code-point-bits] {
 pfc {
 mru mru-value;
 }
 }
 }
 cable-length cable-length-value;
 }
```

**Hierarchy Level** [edit [class-of-service congestion-notification-profile](#) *profile-name*]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure priority-based flow control (PFC) on incoming traffic.

**Options** The remaining statements are explained separately.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Configuring CoS PFC for FCoE Traffic on page 6504](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 6905](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 6493](#)

## interface (DCBX)

---

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <pre>interface (<i>interface-name</i>   all) {<br/>  disable;<br/>  application-map <i>application-map-name</i>;<br/>  applications {<br/>    no-auto-negotiation;<br/>  }<br/>  enhanced-transmission-selection {<br/>    no-auto-negotiation;<br/>    no-recommendation-tlv;<br/>    recommendation-tlv {<br/>      no-auto-negotiation;<br/>    }<br/>  }<br/>  dcbx-version (auto-negotiate   ieee-dcbx   dcbx-version-1.01);<br/>  priority-flow-control {<br/>    no-auto-negotiation;<br/>  }<br/>}</pre>                                   |
| Hierarchy Level          | [edit <a href="#">protocols dcbx</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Release Information      | <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 11.3 for the EX Series switches.</p> <p><b>Mode</b> and <b>recommendation-tlv</b> statements introduced in Junos OS Release 12.2 for the QFX Series.</p>                                                                                                                                                                                                                                                                      |
| Description              | Configure DCBX properties on an interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Options                  | <p><i>interface-name</i>—Name of the interface.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Related Documentation    | <ul style="list-style-type: none"><li>• <a href="#">show dcbx neighbors on page 6581</a></li><li>• <a href="#">Configuring DCBX Autonegotiation on page 6524</a></li><li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li><li>• <a href="#">Understanding DCB Features and Requirements on page 6490</a></li><li>• <a href="#">Understanding DCB Features and Requirements on EX Series Switches</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li></ul> |

## interfaces (Class of Service)

```
Syntax interfaces {
 interface-name {
 congestion-notification-profile profile-name {
 }
 forwarding-class forwarding-class-name;
 forwarding-class-set forwarding-class-set-name {
 output-traffic-control-profile profile-name;
 }
 rewrite-value {
 input {
 ieee-802.1{
 code-point code-point-bits;
 }
 }
 }
 scheduler-map scheduler-map-name
 unit logical-unit-number {
 classifiers {
 (dscp | dscp-ipv6 | ieee-802.1 | exp) (classifier-name | default);
 }
 forwarding-class class-name;
 rewrite-rules {
 (dscp | dscp-ipv6 | ieee-802.1 | exp) (classifier-name | default);
 }
 }
 }
}
```

**Hierarchy Level** [edit [class-of-service](#)]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series

**Description** Configure interface-specific CoS properties for incoming packets.



**NOTE:** Only switches that support direct port scheduling also support applying a scheduler map directly to an interface. When using enhanced transmission selection (ETS) hierarchical port scheduling, you cannot apply a scheduler map directly to an interface; instead, you associate the scheduler map with a traffic control profile and apply the traffic control profile to the interface.



**NOTE:** Only switches that support native Fibre Channel interfaces support the rewrite-value statement, which enables you to rewrite the IEEE 802.1p code points on native Fibre Channel interfaces.



**NOTE:** OCX Series switches do not support MPLS, so they do not support EXP classifiers or rewrite rules. OCX Series switches do not support the congestion-notification-profile configuration statement, which applies priority-based flow control (PFC) to interface output queues.


**Options**    *interface-name*—Name of the interface.

The statements are explained separately.

**Required Privilege Level**    interfaces—To view this statement in the configuration.  
                                         interface-control—To add this statement to the configuration.

**Related Documentation**    • [Assigning CoS Components to Interfaces on page 6627](#)  
                                         • [Interfaces Overview on page 2785](#)

## mru

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>mru mru-value;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service congestion-notification-profile profile-name input ieee-802.1 code-point code-point-bits pfc</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Configure the maximum receive unit (MRU) of the interface in bytes (incoming packet sizes must be less than or equal to the MRU, or the packets are dropped). The system uses the MRU and the cable length to calculate the amount of buffer headroom reserved to support priority-based flow control (PFC). The lower the MRU and the shorter the cable length, the less headroom buffer space is required for PFC.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|                                 | <div>  <p><b>NOTE:</b> You can also set a maximum transmission unit (MTU) value (the largest packet size the interface sends) for interfaces by including the <code>mtu</code> statement at the [edit <a href="#">interfaces interface-name</a>] hierarchy level.</p> </div>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Default</b>                  | <p>For priority 3 traffic, the default MRU value is 2500 bytes.</p> <p>For priority 4 traffic, the default MRU value is 9216 bytes.</p> <p>For user-configured priorities, the default MRU value is 2500 bytes.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <b>mru-value</b> —Value of the maximum packet receive unit size in bytes (generally from 1500 to 9216 bytes, but there is no configuration restriction).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | <p><code>interfaces</code>—To view this statement in the configuration.</p> <p><code>interface-control</code>—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 6905</a></li> <li>• <a href="#">Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 6986</a></li> <li>• <a href="#">Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 6977</a></li> <li>• <a href="#">Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch) on page 6969</a></li> <li>• <a href="#">Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI) on page 7000</a></li> <li>• <a href="#">Understanding CoS Flow Control (Ethernet PAUSE and PFC) on page 6493</a></li> <li>• <a href="#">Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 6873</a></li> </ul> |

## output (Congestion Notification)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>output {<br/>  ieee-802.1 {<br/>    code-point [code-point-bits] {<br/>      flow-control-queue [queue   list-of-queues];<br/>    }<br/>  }<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit <b>class-of-service congestion-notification-profile</b> <i>profile-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Configure priority-based flow control (PFC) on output queues.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | The remaining statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 6905</a></li><li>• <a href="#">Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 6986</a></li><li>• <a href="#">Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 6977</a></li><li>• <a href="#">Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch) on page 6969</a></li><li>• <a href="#">Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI) on page 7000</a></li><li>• <a href="#">Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 6873</a></li></ul> |

## pfc (Input Congestion Notification)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>pfc {<br/>    <b>mru</b> <i>mru-value</i>;<br/>}</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit <b>class-of-service</b> <b>congestion-notification-profile</b> <i>profile-name</i> <b>input</b> <b>ieee-802.1</b> <b>code-point</b> <i>code-point-bits</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Enable and configure ingress interface priority-based flow control (PFC).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | The remaining statement is explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | interfaces—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 6905</a></li> <li>• <a href="#">Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 6986</a></li> <li>• <a href="#">Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 6977</a></li> <li>• <a href="#">Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch) on page 6969</a></li> <li>• <a href="#">Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI) on page 7000</a></li> <li>• <a href="#">Understanding CoS Flow Control (Ethernet PAUSE and PFC) on page 6493</a></li> <li>• <a href="#">Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 6873</a></li> </ul> |

## policy-options

---

**Syntax**    `policy-options`  
             `application-maps` *application-map-name* {  
               `application` *application-name* {  
                 `code-points` [ *aliases* ] [ *bit-patterns* ];  
               }  
             }  
             `policy-statement` *policy-name* {  
               `term` *term-name* {  
                 `from` {  
                   *family* *family-name*;  
                   *match-conditions*;  
                   *policy* *subroutine-policy-name*;  
                   *prefix-list* *prefix-list-name*;  
                   *prefix-list-filter* *prefix-list-name* *match-type* <*actions*>;  
                   *route-filter* *destination-prefix* *match-type* <*actions*>;  
                   *source-address-filter* *source-prefix* *match-type* <*actions*>;  
                 }  
                 `to` {  
                   *match-conditions*;  
                   *policy* *subroutine-policy-name*;  
                 }  
                 `then` *actions*;  
               }  
             }

**Hierarchy Level**    [edit]

**Release Information**    Statement introduced in Junos OS Release 12.1 for the QFX Series.  
                             Statement introduced in Junos OS Release 12.1 for the EX Series.  
                             Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description**    Configure options such as application maps for DCBX application protocol exchange and policy statements.

**Required Privilege Level**    `storage`—To view this statement in the configuration.  
                                     `storage-control`—To add this statement to the configuration.

**Related Documentation**

- [Defining an Application for DCBX Application Protocol TLV Exchange on page 6532](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 6535](#)
- [Example: Configuring DCBX to Support an iSCSI Application](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 6528](#)
- [Understanding DCBX Application Protocol TLV Exchange on EX Series Switches](#)




## priority-flow-control

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>priority-flow-control {<br/>    no-auto-negotiation;<br/>}</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit <a href="#">protocols dcbx interface</a> (all   <i>interface-name</i> )]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 11.3 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Disable autonegotiation of priority-based flow control (PFC) on one or more Ethernet interfaces. Autonegotiation enables PFC on an interface only if the switch and the peer device connected to the switch both support PFC and have the same PFC configuration. Disabling autonegotiation on an interface forces the interface to use the PFC state (enabled or disabled) that is configured on the switch by the configuration and assignment of the congestion notification profile.                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <b>no-auto-negotiation</b> —Disable automatic negotiation of PFC.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show dcbx neighbors on page 6581</a></li> <li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 6905</a></li> <li>• <i>Configuring Priority-Based Flow Control for an EX Series Switch (CLI Procedure)</i></li> <li>• <a href="#">Configuring DCBX Autonegotiation on page 6524</a></li> <li>• <a href="#">Example: Configuring CoS PFC for FCoE Traffic on page 6504</a></li> <li>• <i>Understanding Data Center Bridging Capability Exchange Protocol for EX Series Switches</i></li> <li>• <i>Understanding Priority-Based Flow Control</i></li> <li>• <a href="#">Understanding DCB Features and Requirements on page 6490</a></li> </ul> |

## protocol (Applications)

---


|                                                                                                                                                                                                                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                          | <code>protocol (tcp   udp);</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>                                                                                                                                                                                                 | [edit applications <a href="#">application</a> <i>application-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>                                                                                                                                                                                             | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>                                                                                                                                                                                                     | Networking protocol type, which combines with <b>destination-port</b> to identify an application type.                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <div> <b>NOTE:</b> To create an application for iSCSI, use the protocol <code>tcp</code> with the destination port number 3260.</div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                                                                                                                                                                                                         | <code>tcp</code> —Transmission Control Protocol<br><br><code>udp</code> —User Datagram Protocol                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b>                                                                                                                                                                                        | <code>interface</code> —To view this statement in the configuration.<br><code>interface-control</code> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>                                                                                                                                                                                           | <ul style="list-style-type: none"><li>• <a href="#">Defining an Application for DCBX Application Protocol TLV Exchange on page 6532</a></li><li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 6535</a></li><li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 6528</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li></ul> |

## recommendation-tlv

---


|                                 |                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>recommendation-tlv {<br/>    no-auto-negotiation;<br/>}</pre>                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit <a href="#">protocols dcbx interface interface-name enhanced-transmission-selection</a> ]                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Enable DCBX to send the ETS Recommendation TLV (also known as the Information TLV) on egress. This feature is valid only if the interface DCBX mode is IEEE DCBX. If the interface DCBX mode is DCBX version 1.01, this statement has no effect. (DCBX version 1.01 does not advertise separate TLVs for individual attributes.) |
| <b>Default</b>                  | DCBX-enabled interfaces send the ETS recommendation TLV unless it is disabled.                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <b>no-auto-negotiation</b> —Disable sending of the ETS recommendation TLV.                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show dcbx neighbors on page 6581</a></li><li>• <a href="#">Configuring DCBX Autonegotiation on page 6524</a></li></ul>                                                                                                                                                       |

## rx-buffers

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | rx-buffers (on   off);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | [edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">ether-options</a> <a href="#">configured-flow-control</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <p>Enable or disable an interface to generate and send Ethernet PAUSE messages. If you enable the receive buffers to generate and send PAUSE messages, when the receive buffers reach a certain level of fullness, the interface sends a PAUSE message to the connected peer. If the connected peer is properly configured, it stops transmitting frames to the interface on the entire link. When the interface receive buffer empties below a certain threshold, the interface sends a message to the connected peer to resume sending frames.</p> <p>Ethernet PAUSE prevents buffers from overflowing and dropping packets during periods of network congestion. If the other devices in the network are also configured to support PAUSE, PAUSE supports lossless operation. Use the <b>rx-buffers</b> statement with the <b>tx-buffers</b> statement to configure asymmetric Ethernet PAUSE on an interface. (Use the <b>flow-control</b> statement to enable symmetric PAUSE and the <b>no-flow-control</b> statement to disable symmetric PAUSE on an interface. Symmetric flow control and asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.)</p> |
| <div>  <p><b>NOTE:</b> Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual “lanes” of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC).</p> <p>Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.</p> </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Default</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Flow control is disabled. You must explicitly configure Ethernet PAUSE flow control on interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>on   off</b> —Enable or disable an interface to generate and send Ethernet PAUSE messages.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <ul style="list-style-type: none"> <li>• <a href="#">flow-control on page 2996</a></li> <li>• <a href="#">tx-buffers on page 3033</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

- *Configuring CoS Asymmetric Ethernet PAUSE Flow Control*
- *Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control*
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 6493](#)

## tx-buffers

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | tx-buffers (on   off);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">ether-options</a> <a href="#">configured-flow-control</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | <p>Enable or disable an interface to respond to received Ethernet PAUSE messages. If you enable the transmit buffers to respond to PAUSE messages, when the interface receives a PAUSE message from the connected peer, the interface stops transmitting frames on the entire link. When the receive buffer on the connected peer empties below a certain threshold, the peer interface sends a message to the paused interface to resume sending frames.</p> <p>Ethernet PAUSE prevents buffers from overflowing and dropping packets during periods of network congestion. If the other devices in the network are also configured to support PAUSE, PAUSE supports lossless operation. Use the <b>tx-buffers</b> statement with the <b>rx-buffers</b> statement to configure asymmetric Ethernet PAUSE on an interface. (Use the <b>flow-control</b> statement to enable symmetric PAUSE and the <b>no-flow-control</b> statement to disable symmetric PAUSE on an interface. Symmetric flow control and asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.)</p> |
|                                 | <p> <b>NOTE:</b> Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual “lanes” of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC).</p> <p>Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Default</b>                  | Flow control is disabled. You must explicitly configure Ethernet PAUSE flow control on interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | on   off—Enable or disable an interface to respond to an Ethernet PAUSE message.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">flow-control on page 2996</a></li> <li>• <a href="#">rx-buffers on page 3027</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

- *Configuring CoS Asymmetric Ethernet PAUSE Flow Control*
- *Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control*
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 6493](#)





# Operational Commands (Basic Concepts)

- [Monitoring Interfaces That Have CoS Components on page 7161](#)
- [show class-of-service](#)
- [show class-of-service interface](#)
- [show class-of-service shared-buffer](#)
- [show pfe filter hw summary](#)
- [show pfe next-hop](#)
- [show pfe route](#)
- [show pfe terse](#)
- [show pfe version](#)
- [show interfaces voq](#)

## Monitoring Interfaces That Have CoS Components

**Purpose** Use the monitoring functionality to display details about the physical and logical interfaces and the CoS components assigned to them.

**Action** To monitor interfaces that have CoS components in the CLI, enter the command:

```
user@switch> show class-of-service interface
```

To monitor a specific interface in the CLI, enter the command:

```
user@switch> show class-of-service interface interface-name
```

**Meaning** [Table 607](#) summarizes key output fields for CoS interfaces.

**Table 607: Summary of Key CoS Interfaces Output Fields**

| Field              | Values                                                              |
|--------------------|---------------------------------------------------------------------|
| Physical interface | Name of a physical interface to which CoS components are assigned.  |
| Index              | Index of this interface or the internal index of a specific object. |
| Queues supported   | Number of queues you can configure on the interface.                |

Table 607: Summary of Key CoS Interfaces Output Fields (*continued*)

| Field                         | Values                                                                                                                                                                                                                                        |
|-------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Queues in use                 | Number of queues currently configured.                                                                                                                                                                                                        |
| Scheduler map                 | Name of the scheduler map associated with this interface.                                                                                                                                                                                     |
| Congestion-notification       | Status of congestion notification (enabled or disabled).<br><br><b>NOTE:</b> OCX Series switches do not support congestion notification profiles.                                                                                             |
| Rewrite Input IEEE Code-point | (Fibre Channel NP_Port interfaces only) IEEE 802.1p code point (priority) the interface assigns to incoming Fibre Channel (FC) traffic when the interface encapsulates the FC traffic in Ethernet before forwarding it onto the FCoE network. |
| Logical Interface             | Name of a logical interface on the physical interface to which CoS components are assigned.                                                                                                                                                   |
| Object                        | Category of an object—for example, <b>classifier</b> , <b>scheduler-map</b> , or <b>rewrite</b> .                                                                                                                                             |
| Name                          | Name of the object—for example, <b>ba-classifier</b> .                                                                                                                                                                                        |
| Type                          | Type of the object—for example, <b>ieee8021p</b> for a classifier.                                                                                                                                                                            |

**Related Documentation** • [Assigning CoS Components to Interfaces on page 6627](#)

## show class-of-service

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show class-of-service</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.<br>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Display the class-of-service (CoS) information.                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Monitoring CoS Code-Point Value Aliases on page 7236</a></li> <li>• <a href="#">Monitoring CoS Classifiers on page 7231</a></li> <li>• <a href="#">Monitoring CoS Forwarding Classes on page 7232</a></li> <li>• <a href="#">Monitoring Interfaces That Have CoS Components on page 7161</a></li> <li>• <a href="#">Monitoring CoS Scheduler Maps on page 7291</a></li> <li>• <a href="#">Monitoring CoS Rewrite Rules on page 7235</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show class-of- service on page 7164</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Output Fields</b>            | <a href="#">Table 608</a> lists the output fields for the <b>show class-of-service</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                 |

**Table 608: show class-of-service Output Fields**

| Field Name              | Field Description                                                                                                                                                                                                                                            | Level of Output |
|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Forwarding class</b> | The forwarding class configuration: <ul style="list-style-type: none"> <li>• <b>Forwarding class</b>—Name of the forwarding class.</li> <li>• <b>ID</b>—Forwarding class ID.</li> <li>• <b>Queue</b>—Queue number.</li> </ul>                                | All levels      |
| <b>Code point type</b>  | The type of code-point alias: <ul style="list-style-type: none"> <li>• <b>dscp</b>—Aliases for DiffServ code point (DSCP) values.</li> <li>• <b>ieee-802.1</b>—Aliases for IEEE 802.1p values.</li> <li>• <b>exp</b>—Aliases for MPLS EXP values.</li> </ul> | All levels      |
| <b>Alias</b>            | Names given to CoS values.                                                                                                                                                                                                                                   | All levels      |
| <b>Bit pattern</b>      | Set of bits associated with an alias.                                                                                                                                                                                                                        | All levels      |
| <b>Classifier</b>       | Name of the classifier.                                                                                                                                                                                                                                      | All levels      |
| <b>Code point</b>       | Code-point values.                                                                                                                                                                                                                                           | All levels      |
| <b>Loss priority</b>    | Loss priority assigned to specific CoS values and aliases of the classifier.                                                                                                                                                                                 | All levels      |

Table 608: show class-of-service Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                              | Level of Output |
|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Rewrite rule</b>            | Name of the rewrite rule if one has been configured.                                                                                                                                                                                           | All levels      |
| <b>Drop profile</b>            | Name of the drop profile.                                                                                                                                                                                                                      | All levels      |
| <b>Type</b>                    | Type of drop profile. QFX Series supports only the <b>discrete</b> type of drop-profile.                                                                                                                                                       | All levels      |
| <b>Fill level</b>              | Percentage of queue buffer fullness in a drop profile at which packets begin to drop during periods of congestion.                                                                                                                             | All levels      |
| <b>Scheduler map</b>           | Name of the scheduler map.                                                                                                                                                                                                                     | All levels      |
| <b>Scheduler</b>               | Name of the scheduler.                                                                                                                                                                                                                         | All levels      |
| <b>Transmit rate</b>           | Transmission rate of the scheduler.                                                                                                                                                                                                            | All levels      |
| <b>Buffer size</b>             | Delay buffer size in the queue.                                                                                                                                                                                                                | All levels      |
| <b>Drop profiles</b>           | Drop profiles configured for the specified scheduler.                                                                                                                                                                                          | All levels      |
| <b>Protocol</b>                | Transport protocol corresponding to the drop profile.                                                                                                                                                                                          | All levels      |
| <b>Name</b>                    | Name of the drop profile.                                                                                                                                                                                                                      | All levels      |
| <b>Queues supported</b>        | Number of queues that can be configured on the interface.                                                                                                                                                                                      | All levels      |
| <b>Queues in use</b>           | Number of queues currently configured.                                                                                                                                                                                                         | All levels      |
| <b>Physical interface</b>      | Name of the physical interface.                                                                                                                                                                                                                | All levels      |
| <b>Scheduler map</b>           | Name of the scheduler map.                                                                                                                                                                                                                     | All levels      |
| <b>Congestion-notification</b> | Enabled if a congestion notification profile is applied to the interface; disabled if no congestion notification profile is applied to the interface.<br><br><b>NOTE:</b> OCX Series switches do not support congestion notification profiles. | All levels      |
| <b>Forwarding class set</b>    | Name of the forwarding class set (priority group).<br><br><b>NOTE:</b> Only on systems that support enhanced transmission selection (ETS) hierarchical port scheduling.                                                                        |                 |
| <b>Index</b>                   | Internal index of an object.                                                                                                                                                                                                                   | All levels      |

## Sample Output

### show class-of- service

```
user@switch> show class-of-service
```

| Forwarding class | ID | Queue |
|------------------|----|-------|
| best-effort      | 0  | 0     |
| fcoe             | 1  | 3     |
| no-loss          | 2  | 4     |
| network-control  | 3  | 7     |
| mcast            | 8  | 8     |

Code point type: dscp

| Alias | Bit pattern |
|-------|-------------|
| af11  | 001010      |
| af12  | 001100      |
| ...   | ...         |

Code point type: ieee-802.1

| Alias | Bit pattern |
|-------|-------------|
| af11  | 100         |
| ...   | ...         |

Classifier: dscp-default, Code point type: dscp, Index: 7

| Code point | Forwarding class | Loss priority |
|------------|------------------|---------------|
| 000000     | best-effort      | low           |
| 000001     | best-effort      | low           |
| ...        | ...              | ...           |

Classifier: ieee8021p-default, Code point type: ieee-802.1, Index: 11

| Code point | Forwarding class | Loss priority |
|------------|------------------|---------------|
| 000        | best-effort      | low           |
| 001        | best-effort      | low           |
| 010        | best-effort      | low           |
| 011        | fcoe             | low           |
| 100        | no-loss          | low           |
| 101        | best-effort      | low           |
| 110        | network-control  | low           |
| 111        | network-control  | low           |

Drop profile:<default-drop-profile>, Type: discrete, Index: 1

Fill level  
100

Scheduler map: <default>, Index: 2

Scheduler: <default-be>, Forwarding class: best-effort, Index: 21

Transmit rate: 5 percent, Rate Limit: none, Buffer size: 5 percent, Buffer Limit: none,

Priority: low

Excess Priority: low

drop-profile-map-set-type: mark

Drop profiles:

| Loss priority | Protocol | Index | Name                   |
|---------------|----------|-------|------------------------|
| Low           | any      | 1     | <default-drop-profile> |
| Medium high   | any      | 1     | <default-drop-profile> |
| High          | any      | 1     | <default-drop-profile> |

Scheduler: <default-fcoe>, Forwarding class: fcoe, Index: 50

Transmit rate: 35 percent, Rate Limit: none, Buffer size: 35 percent, Buffer Limit: none,

Priority: low

Excess Priority: low

drop-profile-map-set-type: mark

Drop profiles:

| Loss priority | Protocol | Index | Name |
|---------------|----------|-------|------|
|---------------|----------|-------|------|

|             |     |   |                        |
|-------------|-----|---|------------------------|
| Low         | any | 1 | <default-drop-profile> |
| Medium high | any | 1 | <default-drop-profile> |
| High        | any | 1 | <default-drop-profile> |

Scheduler: <default-noloss>, Forwarding class: no-loss, Index: 51  
 Transmit rate: 35 percent, Rate Limit: none, Buffer size: 35 percent, Buffer  
 Limit: none,

Priority: low  
 Excess Priority: low  
 drop-profile-map-set-type: mark

Drop profiles:

| Loss priority | Protocol | Index | Name                   |
|---------------|----------|-------|------------------------|
| Low           | any      | 1     | <default-drop-profile> |
| Medium high   | any      | 1     | <default-drop-profile> |
| High          | any      | 1     | <default-drop-profile> |

Scheduler: <default-nc>, Forwarding class: network-control, Index: 23  
 Transmit rate: 5 percent, Rate Limit: none, Buffer size: 5 percent, Buffer  
 Limit: none,

Priority: low  
 Excess Priority: low  
 drop-profile-map-set-type: mark

Drop profiles:

| Loss priority | Protocol | Index | Name                   |
|---------------|----------|-------|------------------------|
| Low           | any      | 1     | <default-drop-profile> |
| Medium high   | any      | 1     | <default-drop-profile> |
| High          | any      | 1     | <default-drop-profile> |

Scheduler: <default-mcast>, Forwarding class: mcast, Index: 49  
 Transmit rate: 20 percent, Rate Limit: none, Buffer size: 20 percent, Buffer  
 Limit: none,

Priority: low  
 Excess Priority: low  
 drop-profile-map-set-type: mark

Drop profiles:

| Loss priority | Protocol | Index | Name                   |
|---------------|----------|-------|------------------------|
| Low           | any      | 1     | <default-drop-profile> |
| Medium high   | any      | 1     | <default-drop-profile> |
| High          | any      | 1     | <default-drop-profile> |

Physical interface: xe-0/0/0, Index: 129

Queues supported: 12, Queues in use: 12

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

Physical interface: xe-0/0/1, Index: 130

Queues supported: 12, Queues in use: 12

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

... ..

Forwarding class set: lan-fcset, Type: normal-type, Forwarding class set index:  
 7

| Forwarding class | Index |
|------------------|-------|
| best-effort      | 0     |

## show class-of-service interface

**Syntax** `show class-of-service interface`  
`<comprehensive | detail> <interface-name>`

**Release Information** Command introduced before Junos OS Release 7.4.  
 Command introduced in Junos OS Release 9.0 for EX Series switches.  
 Forwarding class map information added in Junos OS Release 9.4.  
 Command introduced in Junos OS Release 11.1 for the QFX Series.  
 Command introduced in Junos OS Release 12.1 for the PTX Series Packet Transport Switches.  
 Command introduced in Junos OS Release 12.2 for the ACX Series Universal Access routers.  
 Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.  
 Options **detail** and **comprehensive** introduced in Junos OS Release 11.4.  
 Command introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.

**Description** Display the logical and physical interface associations for the classifier, rewrite rules, and scheduler map objects.



**NOTE:** On routing platforms with dual Routing Engines, running this command on the backup Routing Engine, with or without any of the available options, is not supported and produces the following error message:

**error: the class-of-service subsystem is not running**

**Options** **none**—Display CoS associations for all physical and logical interfaces.

**comprehensive**—(M Series, MX Series, and T Series routers) (Optional) Display comprehensive quality-of-service (QoS) information about all physical and logical interfaces.

**detail**—(M Series, MX Series, and T Series routers) (Optional) Display QoS and CoS information based on the interface.

If the **interface** *interface-name* is a physical interface, the output includes:

- Brief QoS information about the physical interface
- Brief QoS information about the logical interface
- CoS information about the physical interface
- Brief information about filters or policers of the logical interface
- Brief CoS information about the logical interface

If the **interface** *interface-name* is a logical interface, the output includes:

- Brief QoS information about the logical interface

- Information about filters or policers for the logical interface
- CoS information about the logical interface

**interface-name**—(Optional) Display class-of-service (CoS) associations for the specified interface.

**none**—Display CoS associations for all physical and logical interfaces.

**Required Privilege Level**

view

**Related Documentation**

- *Verifying and Managing Junos OS Enhanced Subscriber Management*

**List of Sample Output**

[show class-of-service interface \(Physical\) on page 7179](#)  
[show class-of-service interface \(Logical\) on page 7180](#)  
[show class-of-service interface \(Gigabit Ethernet\) on page 7180](#)  
[show class-of-service interface \(ANCP\) on page 7180](#)  
[show class-of-service interface \(PPPoE Interface\) on page 7180](#)  
[show class-of-service interface \(T4000 Routers with Type 5 FPCs\) on page 7180](#)  
[show class-of-service interface detail on page 7181](#)  
[show class-of-service interface comprehensive on page 7181](#)  
[show class-of-service interface \(ACX Series Routers\) on page 7192](#)  
[show class-of-service interface \(PPPoE Subscriber Interface for Enhanced Subscriber Management\) on page 7195](#)

**Output Fields**

[Table 609](#) describes the output fields for the **show class-of-service interface** command. Output fields are listed in the approximate order in which they appear.

**Table 609: show class-of-service interface Output Fields**

| Field Name         | Field Description                                                                                                                                                                                                                                                                                          |                                                      |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------|
| Physical interface | Name of a physical interface.                                                                                                                                                                                                                                                                              |                                                      |
| Index              | Index of this interface or the internal index of this object.<br><br>(Enhanced subscriber management for MX Series routers) Index values for dynamic CoS traffic control profiles and dynamic scheduler maps are larger for enhanced subscriber management than they are for legacy subscriber management. |                                                      |
| Dedicated Queues   | Status of dedicated queues configured on an interface. Supported only on Trio MPC/MIC interfaces on MX Series routers.                                                                                                                                                                                     | Number of queues you can configure on the interface. |
| Queues supported   | Number of queues you can configure on the interface.                                                                                                                                                                                                                                                       |                                                      |
| Queues in use      | Number of queues currently configured.                                                                                                                                                                                                                                                                     |                                                      |



Table 609: show class-of-service interface Output Fields (*continued*)

| Field Name                                 | Field Description                                                                                                                                                                                                                                                                                     |
|--------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Total non-default queues created</b>    | Number of queues created in addition to the default queues. Supported only on Trio MPC/MIC interfaces on MX Series routers.<br><br>(Enhanced subscriber management for MX Series routers) This field is not displayed for enhanced subscriber management.                                             |
| <b>Rewrite Input IEEE Code-point</b>       | (QFX Series only) IEEE 802.1p code point (priority) rewrite value. Incoming traffic from the Fibre Channel (FC) SAN is classified into the forwarding class specified in the native FC interface (NP_Port) fixed classifier and uses the priority specified as the IEEE 802.1p rewrite value.         |
| <b>Shaping rate</b>                        | Maximum transmission rate on the physical interface. You can configure the shaping rate on the physical interface, or on the logical interface, but not on both. Therefore, the <b>Shaping rate</b> field is displayed for either the physical interface or the logical interface.                    |
| <b>Scheduler map</b>                       | Name of the output scheduler map associated with this interface.<br><br>(Enhanced subscriber management for MX Series routers) The name of the dynamic scheduler map object is associated with a generated UID (for example, <b>SMAP-1_UID1002</b> ) instead of with a subscriber interface.          |
| <b>Scheduler map forwarding class sets</b> | (QFX Series only) Name of the output fabric scheduler map associated with a QFabric system Interconnect device interface.                                                                                                                                                                             |
| <b>Input shaping rate</b>                  | For Gigabit Ethernet IQ2 PICs, maximum transmission rate on the input interface.                                                                                                                                                                                                                      |
| <b>Input scheduler map</b>                 | For Gigabit Ethernet IQ2 PICs, name of the input scheduler map associated with this interface.                                                                                                                                                                                                        |
| <b>Chassis scheduler map</b>               | Name of the scheduler map associated with the packet forwarding component queues.                                                                                                                                                                                                                     |
| <b>Rewrite</b>                             | Name and type of the rewrite rules associated with this interface.                                                                                                                                                                                                                                    |
| <b>Traffic-control-profile</b>             | Name of the associated traffic control profile.<br><br>(Enhanced subscriber management for MX Series routers) The name of the dynamic traffic control profile object is associated with a generated UID (for example, <b>TC_PROF_100_199_SERIES_UID1006</b> ) instead of with a subscriber interface. |
| <b>Classifier</b>                          | Name and type of classifiers associated with this interface.                                                                                                                                                                                                                                          |
| <b>Forwarding-class-map</b>                | Name of the forwarding map associated with this interface.                                                                                                                                                                                                                                            |
| <b>Congestion-notification</b>             | (QFX Series and EX4600 switches only) Congestion notification state, <b>enabled</b> or <b>disabled</b> .                                                                                                                                                                                              |
| <b>Logical interface</b>                   | Name of a logical interface.                                                                                                                                                                                                                                                                          |
| <b>Object</b>                              | Category of an object: <b>Classifier</b> , <b>Fragmentation-map</b> (for LSQ interfaces only), <b>Scheduler-map</b> , <b>Rewrite</b> , <b>Translation Table</b> (for IQE PICs only), or <b>traffic-class-map</b> (for T4000 routers with Type 5 FPCs).                                                |
| <b>Name</b>                                | Name of an object.                                                                                                                                                                                                                                                                                    |

Table 609: show class-of-service interface Output Fields (*continued*)

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Type</b>             | Type of an object: <b>dscp</b> , <b>dscp-ipv6</b> , <b>exp</b> , <b>ieee-802.1</b> , <b>ip</b> , <b>inet-precedence</b> , or <b>ieee-802.1ad</b> (for traffic class map on T4000 routers with Type 5 FPCs)..                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Link-level type</b>  | Encapsulation on the physical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>MTU</b>              | MTU size on the physical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Speed</b>            | Speed at which the interface is running.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Loopback</b>         | Whether loopback is enabled and the type of loopback.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Source filtering</b> | Whether source filtering is enabled or disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Flow control</b>     | Whether flow control is enabled or disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Auto-negotiation</b> | (Gigabit Ethernet interfaces) Whether autonegotiation is enabled or disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Remote-fault</b>     | (Gigabit Ethernet interfaces) Remote fault status. <ul style="list-style-type: none"> <li>• <b>Online</b>—Autonegotiation is manually configured as online.</li> <li>• <b>Offline</b>—Autonegotiation is manually configured as offline.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Device flags</b>     | The <b>Device flags</b> field provides information about the physical device and displays one or more of the following values: <ul style="list-style-type: none"> <li>• <b>Down</b>—Device has been administratively disabled.</li> <li>• <b>Hear-Own-Xmit</b>—Device receives its own transmissions.</li> <li>• <b>Link-Layer-Down</b>—The link-layer protocol has failed to connect with the remote endpoint.</li> <li>• <b>Loopback</b>—Device is in physical loopback.</li> <li>• <b>Loop-Detected</b>—The link layer has received frames that it sent, thereby detecting a physical loopback.</li> <li>• <b>No-Carrier</b>—On media that support carrier recognition, no carrier is currently detected.</li> <li>• <b>No-Multicast</b>—Device does not support multicast traffic.</li> <li>• <b>Present</b>—Device is physically present and recognized.</li> <li>• <b>Promiscuous</b>—Device is in promiscuous mode and recognizes frames addressed to all physical addresses on the media.</li> <li>• <b>Quench</b>—Transmission on the device is quenched because the output buffer is overflowing.</li> <li>• <b>Recv-All-Multicasts</b>—Device is in multicast promiscuous mode and therefore provides no multicast filtering.</li> <li>• <b>Running</b>—Device is active and enabled.</li> </ul> |

Table 609: show class-of-service interface Output Fields (*continued*)

| Field Name             | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Interface flags</b> | <p>The <b>Interface flags</b> field provides information about the physical interface and displays one or more of the following values:</p> <ul style="list-style-type: none"> <li>• <b>Admin-Test</b>—Interface is in test mode and some sanity checking, such as loop detection, is disabled.</li> <li>• <b>Disabled</b>—Interface is administratively disabled.</li> <li>• <b>Down</b>—A hardware failure has occurred.</li> <li>• <b>Hardware-Down</b>—Interface is nonfunctional or incorrectly connected.</li> <li>• <b>Link-Layer-Down</b>—Interface keepalives have indicated that the link is incomplete.</li> <li>• <b>No-Multicast</b>—Interface does not support multicast traffic.</li> <li>• <b>No-receive No-transmit</b>—Passive monitor mode is configured on the interface.</li> <li>• <b>Point-To-Point</b>—Interface is point-to-point.</li> <li>• <b>Pop all MPLS labels from packets of depth</b>—MPLS labels are removed as packets arrive on an interface that has the <b>pop-all-labels</b> statement configured. The depth value can be one of the following: <ul style="list-style-type: none"> <li>• <b>1</b>—Takes effect for incoming packets with one label only.</li> <li>• <b>2</b>—Takes effect for incoming packets with two labels only.</li> <li>• <b>[ 1 2 ]</b>—Takes effect for incoming packets with either one or two labels.</li> </ul> </li> <li>• <b>Promiscuous</b>—Interface is in promiscuous mode and recognizes frames addressed to all physical addresses.</li> <li>• <b>Recv-All-Multicasts</b>—Interface is in multicast promiscuous mode and provides no multicast filtering.</li> <li>• <b>SNMP-Traps</b>—SNMP trap notifications are enabled.</li> <li>• <b>Up</b>—Interface is enabled and operational.</li> </ul> |
| <b>Flags</b>           | <p>The <b>Logical interface flags</b> field provides information about the logical interface and displays one or more of the following values:</p> <ul style="list-style-type: none"> <li>• <b>ACFC Encapsulation</b>—Address control field Compression (ACFC) encapsulation is enabled (negotiated successfully with a peer).</li> <li>• <b>Device-down</b>—Device has been administratively disabled.</li> <li>• <b>Disabled</b>—Interface is administratively disabled.</li> <li>• <b>Down</b>—A hardware failure has occurred.</li> <li>• <b>Clear-DF-Bit</b>—GRE tunnel or IPsec tunnel is configured to clear the Don't Fragment (DF) bit.</li> <li>• <b>Hardware-Down</b>—Interface protocol initialization failed to complete successfully.</li> <li>• <b>PFC</b>—Protocol field compression is enabled for the PPP session.</li> <li>• <b>Point-To-Point</b>—Interface is point-to-point.</li> <li>• <b>SNMP-Traps</b>—SNMP trap notifications are enabled.</li> <li>• <b>Up</b>—Interface is enabled and operational.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Encapsulation</b>   | Encapsulation on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Admin</b>           | Administrative state of the interface ( <b>Up</b> or <b>Down</b> ).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Link</b>            | Status of physical link ( <b>Up</b> or <b>Down</b> ).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Proto</b>           | Protocol configured on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

Table 609: show class-of-service interface Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Input Filter</b>            | Names of any firewall filters to be evaluated when packets are received on the interface, including any filters attached through activation of dynamic service.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Output Filter</b>           | Names of any firewall filters to be evaluated when packets are transmitted on the interface, including any filters attached through activation of dynamic service.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Link flags</b>              | Provides information about the physical link and displays one or more of the following values: <ul style="list-style-type: none"> <li>• <b>ACFC</b>—Address control field compression is configured. The Point-to-Point Protocol (PPP) session negotiates the ACFC option.</li> <li>• <b>Give-Up</b>—Link protocol does not continue connection attempts after repeated failures.</li> <li>• <b>Loose-LCP</b>—PPP does not use the Link Control Protocol (LCP) to indicate whether the link protocol is operational.</li> <li>• <b>Loose-LMI</b>—Frame Relay does not use the Local Management Interface (LMI) to indicate whether the link protocol is operational.</li> <li>• <b>Loose-NCP</b>—PPP does not use the Network Control Protocol (NCP) to indicate whether the device is operational.</li> <li>• <b>Keepalives</b>—Link protocol keepalives are enabled.</li> <li>• <b>No-Keepalives</b>—Link protocol keepalives are disabled.</li> <li>• <b>PFC</b>—Protocol field compression is configured. The PPP session negotiates the PFC option.</li> </ul> |
| <b>Hold-times</b>              | Current interface hold-time up and hold-time down, in milliseconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>CoS queues</b>              | Number of CoS queues configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Last flapped</b>            | Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago)</b> . For example, <b>Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Statistics last cleared</b> | Number and rate of bytes and packets received and transmitted on the physical interface. <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>IPv6 transit statistics</b> | Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

Table 609: show class-of-service interface Output Fields (*continued*)

| Field Name           | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Input errors</b>  | <p>Input errors on the interface. The labels are explained in the following list:</p> <ul style="list-style-type: none"> <li>• <b>Errors</b>—Sum of the incoming frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li>• <b>Runts</b>—Number of frames received that are smaller than the runt threshold.</li> <li>• <b>Giants</b>—Number of frames received that are larger than the giant threshold.</li> <li>• <b>Bucket Drops</b>—Drops resulting from the traffic load exceeding the interface transmit or receive leaky bucket configuration.</li> <li>• <b>Policed discards</b>—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle.</li> <li>• <b>L3 incompletes</b>—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. Layer 3 incomplete errors can be ignored by configuring the <b>ignore-l3-incompletes</b> statement.</li> <li>• <b>L2 channel errors</b>—Number of times the software did not find a valid logical interface for an incoming frame.</li> <li>• <b>L2 mismatch timeouts</b>—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable.</li> <li>• <b>HS link CRC errors</b>—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces.</li> <li>• <b>HS link FIFO overflows</b>—Number of FIFO overflows on the high-speed links between the ASICs responsible for handling the router interfaces.</li> </ul> |
| <b>Output errors</b> | <p>Output errors on the interface. The labels are explained in the following list:</p> <ul style="list-style-type: none"> <li>• <b>Carrier transitions</b>—Number of times the interface has gone from <b>down</b> to <b>up</b>. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC is malfunctioning.</li> <li>• <b>Errors</b>—Sum of the outgoing frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> </ul> <p><b>NOTE:</b> Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the <b>Drops</b> field does not always use the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p> <ul style="list-style-type: none"> <li>• <b>Aged packets</b>—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware.</li> <li>• <b>HS link FIFO underflows</b>—Number of FIFO underflows on the high-speed links between the ASICs responsible for handling the router interfaces.</li> <li>• <b>MTU errors</b>—Number of packets whose size exceeds the MTU of the interface.</li> </ul>                                                                                                                                                                                                                                                                                                                |
| <b>Egress queues</b> | Total number of egress queues supported on the specified interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

Table 609: show class-of-service interface Output Fields (*continued*)

| Field Name                                  | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Queue counters</b>                       | <p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul> <p><b>NOTE:</b> Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the <b>Dropped packets</b> field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p>                                                                                                                                                                                                                                   |
| <b>SONET alarms</b><br><b>SONET defects</b> | <p>(SONET) SONET media-specific alarms and defects that prevent the interface from passing packets. When a defect persists for a certain period, it is promoted to an alarm. Based on the router configuration, an alarm can ring the red or yellow alarm bell on the router or light the red or yellow alarm LED on the craft interface. See these fields for possible alarms and defects: <b>SONET PHY</b>, <b>SONET section</b>, <b>SONET line</b>, and <b>SONET path</b>.</p>                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>SONET PHY</b>                            | <p>Counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> <li>• <b>Seconds</b>—Number of seconds the defect has been active.</li> <li>• <b>Count</b>—Number of times that the defect has gone from inactive to active.</li> <li>• <b>State</b>—State of the error. A state other than <b>OK</b> indicates a problem.</li> </ul> <p>The <b>SONET PHY</b> field has the following subfields:</p> <ul style="list-style-type: none"> <li>• <b>PLL Lock</b>—Phase-locked loop</li> <li>• <b>PHY Light</b>—Loss of optical signal</li> </ul>                                                                                                                                                                                                                                                                                        |
| <b>SONET section</b>                        | <p>Counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> <li>• <b>Seconds</b>—Number of seconds the defect has been active.</li> <li>• <b>Count</b>—Number of times that the defect has gone from inactive to active.</li> <li>• <b>State</b>—State of the error. A state other than <b>OK</b> indicates a problem.</li> </ul> <p>The <b>SONET section</b> field has the following subfields:</p> <ul style="list-style-type: none"> <li>• <b>BIP-BI</b>—Bit interleaved parity for SONET section overhead</li> <li>• <b>SEF</b>—Severely errored framing</li> <li>• <b>LOS</b>—Loss of signal</li> <li>• <b>LOF</b>—Loss of frame</li> <li>• <b>ES-S</b>—Errored seconds (section)</li> <li>• <b>SES-S</b>—Severely errored seconds (section)</li> <li>• <b>SEFS-S</b>—Severely errored framing seconds (section)</li> </ul> |

Table 609: show class-of-service interface Output Fields (*continued*)

| Field Name        | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>SONET line</b> | <p>Active alarms and defects, plus counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> <li>• <b>Seconds</b>—Number of seconds the defect has been active.</li> <li>• <b>Count</b>—Number of times that the defect has gone from inactive to active.</li> <li>• <b>State</b>—State of the error. A state other than <b>OK</b> indicates a problem.</li> </ul> <p>The <b>SONET line</b> field has the following subfields:</p> <ul style="list-style-type: none"> <li>• <b>BIP-B2</b>—Bit interleaved parity for SONET line overhead</li> <li>• <b>REI-L</b>—Remote error indication (near-end line)</li> <li>• <b>RDI-L</b>—Remote defect indication (near-end line)</li> <li>• <b>AIS-L</b>—Alarm indication signal (near-end line)</li> <li>• <b>BERR-SF</b>—Bit error rate fault (signal failure)</li> <li>• <b>BERR-SD</b>—Bit error rate defect (signal degradation)</li> <li>• <b>ES-L</b>—Errored seconds (near-end line)</li> <li>• <b>SES-L</b>—Severely errored seconds (near-end line)</li> <li>• <b>UAS-L</b>—Unavailable seconds (near-end line)</li> <li>• <b>ES-LFE</b>—Errored seconds (far-end line)</li> <li>• <b>SES-LFE</b>—Severely errored seconds (far-end line)</li> <li>• <b>UAS-LFE</b>—Unavailable seconds (far-end line)</li> </ul>      |
| <b>SONET path</b> | <p>Active alarms and defects, plus counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> <li>• <b>Seconds</b>—Number of seconds the defect has been active.</li> <li>• <b>Count</b>—Number of times that the defect has gone from inactive to active.</li> <li>• <b>State</b>—State of the error. A state other than <b>OK</b> indicates a problem.</li> </ul> <p>The <b>SONET path</b> field has the following subfields:</p> <ul style="list-style-type: none"> <li>• <b>BIP-B3</b>—Bit interleaved parity for SONET section overhead</li> <li>• <b>REI-P</b>—Remote error indication</li> <li>• <b>LOP-P</b>—Loss of pointer (path)</li> <li>• <b>AIS-P</b>—Path alarm indication signal</li> <li>• <b>RDI-P</b>—Path remote defect indication</li> <li>• <b>UNEQ-P</b>—Path unequipped</li> <li>• <b>PLM-P</b>—Path payload (signal) label mismatch</li> <li>• <b>ES-P</b>—Errored seconds (near-end STS path)</li> <li>• <b>SES-P</b>—Severely errored seconds (near-end STS path)</li> <li>• <b>UAS-P</b>—Unavailable seconds (near-end STS path)</li> <li>• <b>ES-PFE</b>—Errored seconds (far-end STS path)</li> <li>• <b>SES-PFE</b>—Severely errored seconds (far-end STS path)</li> <li>• <b>UAS-PFE</b>—Unavailable seconds (far-end STS path)</li> </ul> |

Table 609: show class-of-service interface Output Fields (*continued*)

| Field Name                             | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Received SONET overhead                | <p>Values of the received and transmitted SONET overhead:</p> <ul style="list-style-type: none"> <li><b>C2</b>—Signal label. Allocated to identify the construction and content of the STS-level SPE and for PDI-P.</li> <li><b>F1</b>—Section user channel byte. This byte is set aside for the purposes of users.</li> <li><b>K1 and K2</b>—These bytes are allocated for APS signaling for the protection of the multiplex section.</li> <li><b>J0</b>—Section trace. This byte is defined for STS-1 number 1 of an STS-N signal. Used to transmit a 1-byte fixed-length string or a 16-byte message so that a receiving terminal in a section can verify its continued connection to the intended transmitter.</li> <li><b>S1</b>—Synchronization status. The S1 byte is located in the first STS-1 number of an STS-N signal.</li> <li><b>Z3 and Z4</b>—Allocated for future use.</li> </ul>                                                                                                                                                                                |
| Transmitted SONET overhead             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Received path trace                    | <p>SONET/SDH interfaces allow path trace bytes to be sent inband across the SONET/SDH link. Juniper Networks and other router manufacturers use these bytes to help diagnose misconfigurations and network errors by setting the transmitted path trace message so that it contains the system hostname and name of the physical interface. The received path trace value is the message received from the router at the other end of the fiber. The transmitted path trace value is the message that this router transmits.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Transmitted path trace                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| HDLC configuration                     | <p>Information about the HDLC configuration.</p> <ul style="list-style-type: none"> <li><b>Policing bucket</b>—Configured state of the receiving policer.</li> <li><b>Shaping bucket</b>—Configured state of the transmitting shaper.</li> <li><b>Giant threshold</b>—Giant threshold programmed into the hardware.</li> <li><b>Runt threshold</b>—Runt threshold programmed into the hardware.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Packet Forwarding Engine configuration | <p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> <li><b>Destination slot</b>—FPC slot number.</li> <li><b>PLP byte</b>—Packet Level Protocol byte.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| CoS information                        | <p>Information about the CoS queue for the physical interface.</p> <ul style="list-style-type: none"> <li><b>CoS transmit queue</b>—Queue number and its associated user-configured forwarding class name.</li> <li><b>Bandwidth %</b>—Percentage of bandwidth allocated to the queue.</li> <li><b>Bandwidth bps</b>—Bandwidth allocated to the queue (in bps).</li> <li><b>Buffer %</b>—Percentage of buffer space allocated to the queue.</li> <li><b>Buffer usec</b>—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time.</li> <li><b>Priority</b>—Queue priority: <b>low</b> or <b>high</b>.</li> <li><b>Limit</b>—Displayed if rate limiting is configured for the queue. Possible values are <b>none</b> and <b>exact</b>. If <b>exact</b> is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If <b>none</b> is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.</li> </ul> |
| Forwarding classes                     | Total number of forwarding classes supported on the specified interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Egress queues                          | Total number of egress queues supported on the specified interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |



Table 609: show class-of-service interface Output Fields (*continued*)

| Field Name           | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Queue                | Queue number.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Forwarding classes   | Forwarding class name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Queued Packets       | Number of packets queued to this queue.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Queued Bytes         | Number of bytes queued to this queue. The byte counts vary by PIC type.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Transmitted Packets  | Number of packets transmitted by this queue. When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (displayed under the <b>Packet Forwarding Engine Chassis Queues</b> field) shows the prefragmentation values.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Transmitted Bytes    | Number of bytes transmitted by this queue. The byte counts vary by PIC type.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Tail-dropped packets | Number of packets dropped because of tail drop.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| RED-dropped packets  | <p>Number of packets dropped because of random early detection (RED).</p> <ul style="list-style-type: none"> <li>• (M Series and T Series routers only) On M320 and M120 routers and the T Series routers, the total number of dropped packets is displayed. On all other M Series routers, the output classifies dropped packets into the following categories: <ul style="list-style-type: none"> <li>• <b>Low, non-TCP</b>—Number of low-loss priority non-TCP packets dropped because of RED.</li> <li>• <b>Low, TCP</b>—Number of low-loss priority TCP packets dropped because of RED.</li> <li>• <b>High, non-TCP</b>—Number of high-loss priority non-TCP packets dropped because of RED.</li> <li>• <b>High, TCP</b>—Number of high-loss priority TCP packets dropped because of RED.</li> </ul> </li> <li>• (MX Series routers with enhanced DPCs, and T Series routers with enhanced FPCs only) The output classifies dropped packets into the following categories: <ul style="list-style-type: none"> <li>• <b>Low</b>—Number of low-loss priority packets dropped because of RED.</li> <li>• <b>Medium-low</b>—Number of medium-low loss priority packets dropped because of RED.</li> <li>• <b>Medium-high</b>—Number of medium-high loss priority packets dropped because of RED.</li> <li>• <b>High</b>—Number of high-loss priority packets dropped because of RED.</li> </ul> </li> </ul> <p><b>NOTE:</b> Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), this field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p> |

Table 609: show class-of-service interface Output Fields (*continued*)

| Field Name        | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| RED-dropped bytes | <p>Number of bytes dropped because of RED. The byte counts vary by PIC type.</p> <ul style="list-style-type: none"> <li>(M Series and T Series routers only) On M320 and M120 routers and the T Series routers, only the total number of dropped bytes is displayed. On all other M Series routers, the output classifies dropped bytes into the following categories: <ul style="list-style-type: none"> <li><b>Low, non-TCP</b>—Number of low-loss priority non-TCP bytes dropped because of RED.</li> <li><b>Low, TCP</b>—Number of low-loss priority TCP bytes dropped because of RED.</li> <li><b>High, non-TCP</b>—Number of high-loss priority non-TCP bytes dropped because of RED.</li> <li><b>High, TCP</b>—Number of high-loss priority TCP bytes dropped because of RED.</li> </ul> </li> </ul> <p><b>NOTE:</b> Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), this field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p> |
| Transmit rate     | Configured transmit rate of the scheduler. The rate is a percentage of the total interface bandwidth.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Rate Limit        | <p>Rate limiting configuration of the queue. Possible values are :</p> <ul style="list-style-type: none"> <li><b>None</b>—No rate limit.</li> <li><b>exact</b>—Queue transmits at the configured rate.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Buffer size       | Delay buffer size in the queue.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Priority          | Scheduling priority configured as <b>low</b> or <b>high</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Excess Priority   | Priority of the excess bandwidth traffic on a scheduler: <b>low</b> , <b>medium-low</b> , <b>medium-high</b> , <b>high</b> , or <b>none</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Drop profiles     | <p>Display the assignment of drop profiles.</p> <ul style="list-style-type: none"> <li><b>Loss priority</b>—Packet loss priority for drop profile assignment.</li> <li><b>Protocol</b>—Transport protocol for drop profile assignment.</li> <li><b>Index</b>—Index of the indicated object. Objects that have indexes in this output include schedulers and drop profiles.</li> <li><b>Name</b>—Name of the drop profile.</li> <li><b>Type</b>—Type of the drop profile: <b>discrete</b> or <b>interpolated</b>.</li> <li><b>Fill Level</b>—Percentage fullness of a queue.</li> <li><b>Drop probability</b>—Drop probability at this fill level.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                              |
| Excess Priority   | Priority of the excess bandwidth traffic on a scheduler.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

Table 609: show class-of-service interface Output Fields (*continued*)

| Field Name                    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Drop profiles</b>          | <p>Display the assignment of drop profiles.</p> <ul style="list-style-type: none"> <li>• <b>Loss priority</b>—Packet loss priority for drop profile assignment.</li> <li>• <b>Protocol</b>—Transport protocol for drop profile assignment.</li> <li>• <b>Index</b>—Index of the indicated object. Objects that have indexes in this output include schedulers and drop profiles.</li> <li>• <b>Name</b>—Name of the drop profile.</li> <li>• <b>Type</b>—Type of the drop profile: <b>discrete</b> or <b>interpolated</b>.</li> <li>• <b>Fill Level</b>—Percentage fullness of a queue.</li> <li>• <b>Drop probability</b>—Drop probability at this fill level.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Adjustment information</b> | <p>Display the assignment of shaping-rate adjustments on a scheduler node or queue.</p> <ul style="list-style-type: none"> <li>• <b>Adjusting application</b>—Application that is performing the shaping-rate adjustment. <ul style="list-style-type: none"> <li>• The adjusting application can appear as <b>ancp LS-0</b>, which is the Junos OS Access Node Control Profile process (<b>ancpd</b>) that performs shaping-rate adjustments on schedule nodes.</li> <li>• The adjusting application can also appear as <b>pppoe</b>, which adjusts the shaping-rate and overhead-accounting class-of-service attributes on dynamic subscriber interfaces in a broadband access network based on access line parameters in Point-to-Point Protocol over Ethernet (PPPoE) Tags [TR-101]. This feature is supported on MPC/MIC interfaces on MX Series routers. The shaping rate is based on the actual-data-rate-downstream attribute. The overhead accounting value is based on the access-loop-encapsulation attribute and specifies whether the access loop uses Ethernet (frame mode) or ATM (cell mode).</li> </ul> </li> <li>• <b>Adjustment type</b>—Type of adjustment: <b>absolute</b> or <b>delta</b>.</li> <li>• <b>Configured shaping rate</b>—Shaping rate configured for the scheduler node or queue.</li> <li>• <b>Adjustment value</b>—Value of adjusted shaping rate.</li> <li>• <b>Adjustment target</b>—Level of shaping-rate adjustment performed: <b>node</b> or <b>queue</b>.</li> <li>• <b>Adjustment overhead-accounting mode</b>—Configured shaping mode: <b>frame</b> or <b>cell</b>.</li> <li>• <b>Adjustment overhead bytes</b>—Number of bytes that the ANCP agent adds to or subtracts from the actual downstream frame overhead before reporting the adjusted values to CoS.</li> <li>• <b>Adjustment target</b>—Level of shaping-rate adjustment performed: <b>node</b> or <b>queue</b>.</li> <li>• <b>Adjustment multicast index</b>—</li> </ul> |

## Sample Output

### show class-of-service interface (Physical)

```

user@host> show class-of-service interface so-0/2/3
Physical interface: so-0/2/3, Index: 135
Queues supported: 8, Queues in use: 4
Total non-default queues created: 4
Scheduler map: <default>, Index: 2032638653

Logical interface: fe-0/0/1.0, Index: 68, Dedicated Queues: no
Shaping rate: 32000
Object Name Type Index
Scheduler-map <default> exp 27
Rewrite exp-default exp 21
Classifier exp-default exp 5

```

|                      |                      |     |   |
|----------------------|----------------------|-----|---|
| Classifier           | ipprec-compatibility | ip  | 8 |
| Forwarding-class-map | exp-default          | exp | 5 |

### show class-of-service interface (Logical)

```
user@host> show class-of-service interface so-0/2/3.0
Logical interface: so-0/2/3.0, Index: 68, Dedicated Queues: no
Shaping rate: 32000
Object Name Type Index
Scheduler-map <default> 27
Rewrite exp-default exp 21
Classifier exp-default exp 5
Classifier ipprec-compatibility ip 8
Forwarding-class-map exp-default exp 5
```

### show class-of-service interface (Gigabit Ethernet)

```
user@host> show class-of-service interface ge-6/2/0
Physical interface: ge-6/2/0, Index: 175
Queues supported: 4, Queues in use: 4
Scheduler map: <default>, Index: 2
Input scheduler map: <default>, Index: 3
Chassis scheduler map: <default-chassis>, Index: 4
```

### show class-of-service interface (ANCP)

```
user@host> show class-of-service interface pp0.1073741842
Logical interface: pp0.1073741842, Index: 341
Object Name Type Index
Traffic-control-profile TCP-CVLAN Output 12408
Classifier dscp-ipv6-compatibility dscp-ipv6 9
Classifier ipprec-compatibility ip 13

Adjusting application: ancp LS-0
Adjustment type: absolute
Configured shaping rate: 4000000
Adjustment value: 11228000
Adjustment overhead-accounting mode: Frame Mode
Adjustment overhead bytes: 50
Adjustment target: node
```

### show class-of-service interface (PPPoE Interface)

```
user@host> show class-of-service interface pp0.1
Logical interface: pp0.1, Index: 85
Object Name Type Index
Traffic-control-profile tcp-pppoe.o.pp0.1 Output 2726446535
Classifier ipprec-compatibility ip 13

Adjusting application: PPPoE
Adjustment type: absolute
Adjustment value: 5000000
Adjustment overhead-accounting mode: cell
Adjustment target: node
```

### show class-of-service interface (T4000 Routers with Type 5 FPCs)

```
user@host> show class-of-service interface xe-4/0/0
Physical interface: xe-4/0/0, Index: 153
Queues supported: 8, Queues in use: 4
Shaping rate: 5000000000 bps
Scheduler map: <default>, Index: 2
```

|       |                                          |                      |      |
|-------|------------------------------------------|----------------------|------|
|       | Logical interface: xe-4/0/0.0, Index: 77 |                      |      |
|       | Object                                   | Name                 | Type |
| Index | Classifier                               | ipprec-compatibility | ip   |
| 13    |                                          |                      |      |

[show class-of-service interface detail](#)

```
user@host> show class-of-service interface ge-0/3/0 detail
```

```
Physical interface: ge-0/3/0, Enabled, Physical link is Up
 Link-level type: Ethernet, MTU: 1518, Speed: 1000mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x4000
```

```
Physical interface: ge-0/3/0, Index: 138
Queues supported: 4, Queues in use: 5
Shaping rate: 50000 bps
Scheduler map: interface-scheduler-map, Index: 58414
Input shaping rate: 10000 bps
Input scheduler map: scheduler-map, Index: 15103
Chassis scheduler map: <default-chassis>, Index: 4
Congestion-notification: Disabled
```

```
Logical interface ge-0/3/0.0
 Flags: SNMP-Traps 0x4000 VLAN-Tag [0x8100.1] Encapsulation: ENET2
 inet
 mpls
```

|                         |             |            |                       |               |                |
|-------------------------|-------------|------------|-----------------------|---------------|----------------|
| Interface<br>ge-0/3/0.0 | Admin<br>up | Link<br>up | Proto<br>inet<br>mpls | Input Filter  | Output Filter  |
| Interface<br>ge-0/3/0.0 | Admin<br>up | Link<br>up | Proto<br>inet<br>mpls | Input Policer | Output Policer |

| Logical interface: ge-0/3/0.0, Index: 68 |                      |                |       |
|------------------------------------------|----------------------|----------------|-------|
| Object                                   | Name                 | Type           | Index |
| Rewrite                                  | exp-default          | exp (mpls-any) | 33    |
| Classifier                               | exp-default          | exp            | 10    |
| Classifier                               | ipprec-compatibility | ip             | 13    |

```
Logical interface ge-0/3/0.1
 Flags: SNMP-Traps 0x4000 VLAN-Tag [0x8100.2] Encapsulation: ENET2
 inet
```

|            |       |      |       |       |         |        |         |
|------------|-------|------|-------|-------|---------|--------|---------|
| Interface  | Admin | Link | Proto | Input | Filter  | Output | Filter  |
| ge-0/3/0.1 | up    | up   | inet  |       |         |        |         |
| Interface  | Admin | Link | Proto | Input | Policer | Output | Policer |
| ge-0/3/0.1 | up    | up   | inet  |       |         |        |         |

```
Logical interface: ge-0/3/0.1, Index: 69
 Object Name Type Index
 Classifier ipprec-compatibility ip 13
```

show class-of-service interface comprehensive

```
user@host> show class-of-service interface ge-0/3/0 comprehensive
```

```

Physical interface: ge-0/3/0, Enabled, Physical link is Up
 Interface index: 138, SNMP ifIndex: 601, Generation: 141
 Link-level type: Ethernet, MTU: 1518, Speed: 1000mbps, BPDU Error: None,
 MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled, Flow
 control: Enabled,
 Auto-negotiation: Enabled, Remote fault: Online
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x4000
 CoS queues : 4 supported, 4 maximum usable queues
 Schedulers : 256
 Hold-times : Up 0 ms, Down 0 ms
 Current address: 00:14:f6:f4:b4:5d, Hardware address: 00:14:f6:f4:b4:5d
 Last flapped : 2010-09-07 06:35:22 PDT (15:14:42 ago)
 Statistics last cleared: Never
Traffic statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps
IPv6 total statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Ingress traffic statistics at Packet Forwarding Engine:
 Input bytes : 0 0 bps
 Input packets: 0 0 pps
 Drop bytes : 0 0 bps
 Drop packets: 0 0 pps
Label-switched interface (LSI) traffic statistics:
 Input bytes : 0 0 bps
 Input packets: 0 0 pps
Input errors:
 Errors: 0, Drops: 0, Framing errors: 0, Runt: 0, Policed discards: 0, L3
 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
 Resource errors: 0
Output errors:
 Carrier transitions: 5, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
 FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Ingress queues: 4 supported, 5 in use
Queue counters: Queued packets Transmitted packets Dropped packets

 0 af3 0 0 0
 1 af2 0 0 0
 2 ef2 0 0 0
 3 ef1 0 0 0

Egress queues: 4 supported, 5 in use
Queue counters: Queued packets Transmitted packets Dropped packets

 0 af3 0 0 0
 1 af2 0 0 0
 2 ef2 0 0 0
 3 ef1 0 0 0

```

```

Active alarms : None
Active defects : None
MAC statistics:
 Total octets Receive Transmit
 Total packets 0 0
 Unicast packets 0 0
 Broadcast packets 0 0
 Multicast packets 0 0
 CRC/Align errors 0 0
 FIFO errors 0 0
 MAC control frames 0 0
 MAC pause frames 0 0
 Oversized frames 0
 Jabber frames 0
 Fragment frames 0
 VLAN tagged frames 0
 Code violations 0
Filter statistics:
 Input packet count 0
 Input packet rejects 0
 Input DA rejects 0
 Input SA rejects 0
 Output packet count 0
 Output packet pad count 0
 Output packet error count 0
 CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
 Negotiation status: Complete
 Link partner:
 Link mode: Full-duplex, Flow control: Symmetric/Asymmetric, Remote fault:
OK
 Local resolution:
 Flow control: Symmetric, Remote fault: Link OK
Packet Forwarding Engine configuration:
 Destination slot: 0
CoS information:
 Direction : Output
 CoS transmit queue Bandwidth Buffer Priority
Limit
 % bps % usec
 2 ef2 39 19500 0 120 high
none
 Direction : Input
 CoS transmit queue Bandwidth Buffer Priority
Limit
 % bps % usec
 0 af3 30 3000 45 0 low
none

Physical interface: ge-0/3/0, Enabled, Physical link is Up
 Interface index: 138, SNMP ifIndex: 601
Forwarding classes: 16 supported, 5 in use
Ingress queues: 4 supported, 5 in use
Queue: 0, Forwarding classes: af3
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available

```

```

 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: af2
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: ef2
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: ef1
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Forwarding classes: 16 supported, 5 in use
Egress queues: 4 supported, 5 in use
Queue: 0, Forwarding classes: af3
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RL-dropped packets : 0 0 pps
 RL-dropped bytes : 0 0 bps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: af2
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RL-dropped packets : 0 0 pps
 RL-dropped bytes : 0 0 bps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: ef2
 Queued:
 Packets : 0 0 pps

```



```

 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RL-dropped packets : 0 0 pps
 RL-dropped bytes : 0 0 bps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: ef1
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RL-dropped packets : 0 0 pps
 RL-dropped bytes : 0 0 bps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 5 in use
Queue: 0, Forwarding classes: af3
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : Not Available
 RED-dropped bytes : Not Available
Queue: 1, Forwarding classes: af2
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : Not Available
 RED-dropped bytes : Not Available
Queue: 2, Forwarding classes: ef2
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : Not Available
 RED-dropped bytes : Not Available
Queue: 3, Forwarding classes: ef1
Queued:
 Packets : 108546 0 pps
 Bytes : 12754752 376 bps
Transmitted:
 Packets : 108546 0 pps
 Bytes : 12754752 376 bps

```

```

Tail-dropped packets : 0 0 pps
RED-dropped packets : Not Available
RED-dropped bytes : Not Available

```

```

Physical interface: ge-0/3/0, Index: 138
Queues supported: 4, Queues in use: 5
Shaping rate: 50000 bps

```

```
Scheduler map: interface-scheduler-map, Index: 58414
```

```

Scheduler: ef2, Forwarding class: ef2, Index: 39155
 Transmit rate: 39 percent, Rate Limit: none, Buffer size: 120 us, Buffer
 Limit: none, Priority: high
 Excess Priority: unspecified
 Drop profiles:
 Loss priority Protocol Index Name
 Low any 1 < default-drop-profile>
 Medium low any 1 < default-drop-profile>
 Medium high any 1 < default-drop-profile>
 High any 1 < default-drop-profile>
 Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
 Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
 Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
 Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
 Input shaping rate: 10000 bps
 Input scheduler map: scheduler-map

```

```
Scheduler map: scheduler-map, Index: 15103
```

```

Scheduler: af3, Forwarding class: af3, Index: 35058
 Transmit rate: 30 percent, Rate Limit: none, Buffer size: 45 percent, Buffer
 Limit: none, Priority: low
 Excess Priority: unspecified
 Drop profiles:
 Loss priority Protocol Index Name
 Low any 40582 green
 Medium low any 1 < default-drop-profile>
 Medium high any 1 < default-drop-profile>
 High any 18928 yellow
 Drop profile: green, Type: discrete, Index: 40582
 Fill level Drop probability
 50 0
 100 100
 Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
 Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
 Drop profile: yellow, Type: discrete, Index: 18928
 Fill level Drop probability
 50 0
 100 100

```

Chassis scheduler map: < default-drop-profile>  
 Scheduler map: < default-drop-profile>, Index: 4

Scheduler: < default-drop-profile>, Forwarding class: af3, Index: 25  
 Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer  
 Limit: none, Priority: low  
 Excess Priority: low  
 Drop profiles:

| Loss priority | Protocol | Index | Name                    |
|---------------|----------|-------|-------------------------|
| Low           | any      | 1     | < default-drop-profile> |
| Medium low    | any      | 1     | < default-drop-profile> |
| Medium high   | any      | 1     | < default-drop-profile> |
| High          | any      | 1     | < default-drop-profile> |

Drop profile: < default-drop-profile>, Type: discrete, Index: 1  
 Fill level      Drop probability  
                   100                   100

Drop profile: < default-drop-profile>, Type: discrete, Index: 1  
 Fill level      Drop probability  
                   100                   100

Drop profile: < default-drop-profile>, Type: discrete, Index: 1  
 Fill level      Drop probability  
                   100                   100

Drop profile: < default-drop-profile>, Type: discrete, Index: 1  
 Fill level      Drop probability  
                   100                   100

Scheduler: < default-drop-profile>, Forwarding class: af2, Index: 25  
 Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer  
 Limit: none, Priority: low  
 Excess Priority: low  
 Drop profiles:

| Loss priority | Protocol | Index | Name                    |
|---------------|----------|-------|-------------------------|
| Low           | any      | 1     | < default-drop-profile> |
| Medium low    | any      | 1     | < default-drop-profile> |
| Medium high   | any      | 1     | < default-drop-profile> |
| High          | any      | 1     | < default-drop-profile> |

Drop profile: < default-drop-profile>, Type: discrete, Index: 1  
 Fill level      Drop probability  
                   100                   100

Drop profile: < default-drop-profile>, Type: discrete, Index: 1  
 Fill level      Drop probability  
                   100                   100

Drop profile: < default-drop-profile>, Type: discrete, Index: 1  
 Fill level      Drop probability  
                   100                   100

Drop profile: < default-drop-profile>, Type: discrete, Index: 1  
 Fill level      Drop probability  
                   100                   100

Scheduler: < default-drop-profile>, Forwarding class: ef2, Index: 25  
 Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer  
 Limit: none, Priority: low  
 Excess Priority: low  
 Drop profiles:

| Loss priority | Protocol | Index | Name                    |
|---------------|----------|-------|-------------------------|
| Low           | any      | 1     | < default-drop-profile> |
| Medium low    | any      | 1     | < default-drop-profile> |
| Medium high   | any      | 1     | < default-drop-profile> |
| High          | any      | 1     | < default-drop-profile> |

Drop profile: < default-drop-profile>, Type: discrete, Index: 1  
 Fill level      Drop probability

```

 100 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100

Scheduler: < default-drop-profile>, Forwarding class: ef1, Index: 25
 Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer
Limit: none, Priority: low
 Excess Priority: low
 Drop profiles:
 Loss priority Protocol Index Name
 Low any 1 < default-drop-profile>
 Medium low any 1 < default-drop-profile>
 Medium high any 1 < default-drop-profile>
 High any 1 < default-drop-profile>
Drop profile: , Type: discrete, Index: 1
 Fill level Drop probability
 100 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
 Congestion-notification: Disabled
Forwarding class
priority Policing priority ID Queue Restricted queue Fabric
af3 normal 0 0 0 low
af2 normal 1 1 1 low
ef2 normal 2 2 2 high
ef1 normal 3 3 3 high
af1 normal 4 4 0 low
 normal

Logical interface ge-0/3/0.0 (Index 68) (SNMP ifIndex 152) (Generation 159)
 Flags: SNMP-Traps 0x4000 VLAN-Tag [0x8100.1] Encapsulation: ENET2
 Traffic statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
 Local statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
 Transit statistics:
 Input bytes : 0 0 bps

```

```

Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
Protocol inet, MTU: 1500, Generation: 172, Route table: 0
 Flags: Sendbcst-pkt-to-re
 Input Filters: filter-in-ge-0/3/0.0-i,
 Policer: Input: p1-ge-0/3/0.0-inet-i
Protocol mpls, MTU: 1488, Maximum labels: 3, Generation: 173, Route table: 0
 Flags: Is-Primary
 Output Filters: exp-filter,,,,,

Logical interface ge-1/2/0.0 (Index 347) (SNMP ifIndex 638) (Generation 156)

Forwarding class ID Queue Restricted queue Fabric priority Policing priority
 SPU priority
best-effort 0 0 0 low normal
 low

Aggregate Forwarding-class statistics per forwarding-class
Aggregate Forwarding-class statistics:
Forwarding-class statistics:

Forwarding-class best-effort statistics:
 Input unicast bytes: 0
 Output unicast bytes: 0
 Input unicast packets: 0
 Output unicast packets: 0

 Input multicast bytes: 0
 Output multicast bytes: 0
 Input multicast packets: 0
 Output multicast packets: 0

Forwarding-class expedited-forwarding statistics:
 Input unicast bytes: 0
 Output unicast bytes: 0
 Input unicast packets: 0
 Output unicast packets: 0

 Input multicast bytes: 0
 Output multicast bytes: 0
 Input multicast packets: 0
 Output multicast packets: 0

IPv4 protocol forwarding-class statistics:
Forwarding-class statistics:
Forwarding-class best-effort statistics:

 Input unicast bytes: 0
 Output unicast bytes: 0
 Input unicast packets: 0
 Output unicast packets: 0

 Input multicast bytes: 0
 Output multicast bytes: 0
 Input multicast packets: 0
 Output multicast packets: 0

Forwarding-class expedited-forwarding statistics:
 Input unicast bytes: 0

```

```

Output unicast bytes: 0
Input unicast packets: 0
Output unicast packets: 0

```

```

Input multicast bytes: 0
Output multicast bytes: 0
Input multicast packets: 0
Output multicast packets: 0

```

```

IPv6 protocol forwarding-class statistics:
Forwarding-class statistics:
 Forwarding-class best-effort statistics:

```

```

Input unicast bytes: 0
Output unicast bytes: 0
Input unicast packets: 0
Output unicast packets: 0

```

```

Input multicast bytes: 0
Output multicast bytes: 0
Input multicast packets: 0
Output multicast packets: 0

```

```
Forwarding-class expedited-forwarding statistics:
```

```

Input unicast bytes: 0
Output unicast bytes: 0
Input unicast packets: 0
Output unicast packets: 0

```

```
Logical interface ge-0/3/0.0 (Index 68) (SNMP ifIndex 152)
```

```

Flags: SNMP-Traps 0x4000 VLAN-Tag [0x8100.1] Encapsulation: ENET2
Input packets : 0
Output packets: 0

```

|            |       |      |       |                        |               |
|------------|-------|------|-------|------------------------|---------------|
| Interface  | Admin | Link | Proto | Input Filter           | Output Filter |
| ge-0/3/0.0 | up    | up   | inet  | filter-in-ge-0/3/0.0-i |               |
|            |       |      | mpls  |                        | exp-filter    |

|            |       |      |       |                      |                |
|------------|-------|------|-------|----------------------|----------------|
| Interface  | Admin | Link | Proto | Input Policer        | Output Policer |
| ge-0/3/0.0 | up    | up   |       |                      |                |
|            |       |      | inet  | p1-ge-0/3/0.0-inet-i |                |
|            |       |      | mpls  |                      |                |

```
Filter: filter-in-ge-0/3/0.0-i
```

```
Counters:
```

| Name                         | Bytes | Packets |
|------------------------------|-------|---------|
| count-filter-in-ge-0/3/0.0-i | 0     | 0       |

```
Filter: exp-filter
```

```
Counters:
```

| Name                  | Bytes | Packets |
|-----------------------|-------|---------|
| count-exp-seven-match | 0     | 0       |
| count-exp-zero-match  | 0     | 0       |

```
Policers:
```

| Name                 | Packets |
|----------------------|---------|
| p1-ge-0/3/0.0-inet-i | 0       |

```
Logical interface: ge-0/3/0.0, Index: 68
```

| Object  | Name        | Type           | Index |
|---------|-------------|----------------|-------|
| Rewrite | exp-default | exp (mpls-any) | 33    |

Rewrite rule: exp-default, Code point type: exp, Index: 33

| Forwarding class | Loss priority | Code point |       |
|------------------|---------------|------------|-------|
| af3              | low           | 000        |       |
| af3              | high          | 001        |       |
| af2              | low           | 010        |       |
| af2              | high          | 011        |       |
| ef2              | low           | 100        |       |
| ef2              | high          | 101        |       |
| ef1              | low           | 110        |       |
| ef1              | high          | 111        |       |
| Object           | Name          | Type       | Index |
| Classifier       | exp-default   | exp        | 10    |

Classifier: exp-default, Code point type: exp, Index: 10

| Code point | Forwarding class     | Loss priority |       |
|------------|----------------------|---------------|-------|
| 000        | af3                  | low           |       |
| 001        | af3                  | high          |       |
| 010        | af2                  | low           |       |
| 011        | af2                  | high          |       |
| 100        | ef2                  | low           |       |
| 101        | ef2                  | high          |       |
| 110        | ef1                  | low           |       |
| 111        | ef1                  | high          |       |
| Object     | Name                 | Type          | Index |
| Classifier | ipprec-compatibility | ip            | 13    |

Classifier: ipprec-compatibility, Code point type: inet-precedence, Index: 13

| Code point       | Forwarding class | Loss priority |                  |        |
|------------------|------------------|---------------|------------------|--------|
| 000              | af3              | low           |                  |        |
| 001              | af3              | high          |                  |        |
| 010              | af3              | low           |                  |        |
| 011              | af3              | high          |                  |        |
| 100              | af3              | low           |                  |        |
| 101              | af3              | high          |                  |        |
| 110              | ef1              | low           |                  |        |
| 111              | ef1              | high          |                  |        |
| Forwarding class | ID               | Queue         | Restricted queue | Fabric |
| priority         |                  |               |                  |        |
| af3              | 0                | 0             | 0                | low    |
| normal           |                  |               |                  |        |
| af2              | 1                | 1             | 1                | low    |
| normal           |                  |               |                  |        |
| ef2              | 2                | 2             | 2                | high   |
| normal           |                  |               |                  |        |
| ef1              | 3                | 3             | 3                | high   |
| normal           |                  |               |                  |        |
| af1              | 4                | 4             | 0                | low    |
| normal           |                  |               |                  |        |

Logical interface ge-0/3/0.1 (Index 69) (SNMP ifIndex 154) (Generation 160)

Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.2 ] Encapsulation: ENET2

Traffic statistics:

|                   |   |
|-------------------|---|
| Input bytes :     | 0 |
| Output bytes :    | 0 |
| Input packets:    | 0 |
| Output packets:   | 0 |
| Local statistics: |   |
| Input bytes :     | 0 |

```

Output bytes : 0
Input packets: 0
Output packets: 0
Transit statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
Protocol inet, MTU: 1500, Generation: 174, Route table: 0
Flags: Sendbroadcast-pkt-to-re

```

```

Logical interface ge-0/3/0.1 (Index 69) (SNMP ifIndex 154)
Flags: SNMP-Traps 0x4000 VLAN-Tag [0x8100.2] Encapsulation: ENET2
Input packets : 0
Output packets: 0

```

```

Interface Admin Link Proto Input Filter Output Filter
ge-0/3/0.1 up up mpls
Interface Admin Link Proto Input Policer Output Policer
ge-0/3/0.1 up up mpls

```

```
Logical interface: ge-0/3/0.1, Index: 69
```

| Object     | Name                 | Type | Index |
|------------|----------------------|------|-------|
| Classifier | ipprec-compatibility | ip   | 13    |

```
Classifier: ipprec-compatibility, Code point type: inet-precedence, Index: 13
```

| Code point | Forwarding class | Loss priority |
|------------|------------------|---------------|
| 000        | af3              | low           |
| 001        | af3              | high          |
| 010        | af3              | low           |
| 011        | af3              | high          |
| 100        | af3              | low           |
| 101        | af3              | high          |
| 110        | ef1              | low           |
| 111        | ef1              | high          |

| Forwarding class | ID | Queue | Restricted queue | Fabric |
|------------------|----|-------|------------------|--------|
| priority         |    |       |                  |        |
| af3              | 0  | 0     | 0                | low    |
| af2              | 1  | 1     | 1                | low    |
| ef2              | 2  | 2     | 2                | high   |
| ef1              | 3  | 3     | 3                | high   |
| af1              | 4  | 4     | 0                | low    |

### show class-of-service interface (ACX Series Routers)

```

user@host-g11# show class-of-service interface
Physical interface: at-0/0/0, Index: 130
Queues supported: 4, Queues in use: 4
Scheduler map: <default>, Index: 2
Congestion-notification: Disabled

```



Logical interface: at-0/0/0.0, Index: 69

Logical interface: at-0/0/0.32767, Index: 70

Physical interface: at-0/0/1, Index: 133

Queues supported: 4, Queues in use: 4

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

Logical interface: at-0/0/1.0, Index: 71

Logical interface: at-0/0/1.32767, Index: 72

Physical interface: ge-0/1/0, Index: 146

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

| Object     | Name         | Type      | Index |
|------------|--------------|-----------|-------|
| Rewrite    | dscp-default | dscp      | 31    |
| Classifier | d1           | dscp      | 11331 |
| Classifier | ci           | ieee8021p | 583   |

Logical interface: ge-0/1/0.0, Index: 73

| Object  | Name       | Type           | Index |
|---------|------------|----------------|-------|
| Rewrite | custom-exp | exp (mpls-any) | 46413 |

Logical interface: ge-0/1/0.1, Index: 74

Logical interface: ge-0/1/0.32767, Index: 75

Physical interface: ge-0/1/1, Index: 147

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

| Object     | Name                 | Type | Index |
|------------|----------------------|------|-------|
| Classifier | ipprec-compatibility | ip   | 13    |

Logical interface: ge-0/1/1.0, Index: 76

Physical interface: ge-0/1/2, Index: 148

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

| Object     | Name | Type              | Index |
|------------|------|-------------------|-------|
| Rewrite    | ri   | ieee8021p (outer) | 35392 |
| Classifier | ci   | ieee8021p         | 583   |

Physical interface: ge-0/1/3, Index: 149

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

| Object     | Name                 | Type | Index |
|------------|----------------------|------|-------|
| Classifier | ipprec-compatibility | ip   | 13    |

Logical interface: ge-0/1/3.0, Index: 77

| Object  | Name        | Type           | Index |
|---------|-------------|----------------|-------|
| Rewrite | custom-exp2 | exp (mpls-any) | 53581 |

Physical interface: ge-0/1/4, Index: 150

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

|                                              |                      |      |       |
|----------------------------------------------|----------------------|------|-------|
| Congestion-notification: Disabled            |                      |      |       |
| Object                                       | Name                 | Type | Index |
| Classifier                                   | ipprec-compatibility | ip   | 13    |
| Physical interface: ge-0/1/5, Index: 151     |                      |      |       |
| Queues supported: 8, Queues in use: 5        |                      |      |       |
| Scheduler map: <default>, Index: 2           |                      |      |       |
| Congestion-notification: Disabled            |                      |      |       |
| Object                                       | Name                 | Type | Index |
| Classifier                                   | ipprec-compatibility | ip   | 13    |
| Physical interface: ge-0/1/6, Index: 152     |                      |      |       |
| Queues supported: 8, Queues in use: 5        |                      |      |       |
| Scheduler map: <default>, Index: 2           |                      |      |       |
| Congestion-notification: Disabled            |                      |      |       |
| Object                                       | Name                 | Type | Index |
| Classifier                                   | ipprec-compatibility | ip   | 13    |
| Physical interface: ge-0/1/7, Index: 153     |                      |      |       |
| Queues supported: 8, Queues in use: 5        |                      |      |       |
| Scheduler map: <default>, Index: 2           |                      |      |       |
| Congestion-notification: Disabled            |                      |      |       |
| Object                                       | Name                 | Type | Index |
| Classifier                                   | d1                   | dscp | 11331 |
| Physical interface: ge-0/2/0, Index: 154     |                      |      |       |
| Queues supported: 8, Queues in use: 5        |                      |      |       |
| Scheduler map: <default>, Index: 2           |                      |      |       |
| Congestion-notification: Disabled            |                      |      |       |
| Object                                       | Name                 | Type | Index |
| Classifier                                   | ipprec-compatibility | ip   | 13    |
| Physical interface: ge-0/2/1, Index: 155     |                      |      |       |
| Queues supported: 8, Queues in use: 5        |                      |      |       |
| Scheduler map: <default>, Index: 2           |                      |      |       |
| Congestion-notification: Disabled            |                      |      |       |
| Object                                       | Name                 | Type | Index |
| Classifier                                   | ipprec-compatibility | ip   | 13    |
| Logical interface: ge-0/2/1.0, Index: 78     |                      |      |       |
| Logical interface: ge-0/2/1.32767, Index: 79 |                      |      |       |
| Physical interface: xe-0/3/0, Index: 156     |                      |      |       |
| Queues supported: 8, Queues in use: 5        |                      |      |       |
| Scheduler map: <default>, Index: 2           |                      |      |       |
| Congestion-notification: Disabled            |                      |      |       |
| Object                                       | Name                 | Type | Index |
| Classifier                                   | ipprec-compatibility | ip   | 13    |
| Logical interface: xe-0/3/0.0, Index: 80     |                      |      |       |
| Physical interface: xe-0/3/1, Index: 157     |                      |      |       |
| Queues supported: 8, Queues in use: 5        |                      |      |       |
| Scheduler map: <default>, Index: 2           |                      |      |       |
| Congestion-notification: Disabled            |                      |      |       |
| Object                                       | Name                 | Type | Index |
| Classifier                                   | ipprec-compatibility | ip   | 13    |
| Logical interface: xe-0/3/1.0, Index: 81     |                      |      |       |

```
[edit]
user@host-g11#
```

### show class-of-service interface (PPPoE Subscriber Interface for Enhanced Subscriber Management)

```
user@host> show class-of-service interface pp0.3221225474
 Logical interface: pp0.3221225475, Index: 3221225475
Object Name Type Index
Traffic-control-profile TC_PROF_100_199_SERIES_UID1006 Output 4294967312
Scheduler-map SMAP-1_UID1002 Output 4294967327
Rewrite-Output ieee-rewrite ieee8021p 60432
Rewrite-Output rule1 ip 50463

 Adjusting application: PPPoE IA tags
 Adjustment type: absolute
 Configured shaping rate: 11000000
 Adjustment value: 5000000
 Adjustment target: node

 Adjusting application: ucac
 Adjustment type: delta
 Configured shaping rate: 5000000
 Adjustment value: 100000
 Adjustment target: node
```

## show class-of-service shared-buffer

**Syntax** `show class-of-service shared-buffer`  
`<egress | ingress>`

**Release Information** Command introduced in Junos OS Release 12.3 for the QFX Series.  
 Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Display the shared buffer allocation and partitioning configuration.



**NOTE:** Due to QFX5200 cross-point architecture, all buffer usage counters are maintained separately. When usage counters are displayed with the command `show class-of-service shared-buffer` on QFX5200, various pipe counters are displayed separately.

**Options** **none**—Display ingress and egress shared buffer settings.  
**egress**—(Optional) Display the egress shared buffer settings.  
**ingress**—(Optional) Display the ingress shared buffer settings.

**Required Privilege Level** view

**Related Documentation**

- *Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic*
- *Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic*
- *Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic*
- *Configuring Global Ingress and Egress Shared Buffers*
- *Understanding CoS Buffer Configuration*

**List of Sample Output** [show class-of-service shared-buffer on page 7197](#)

**Output Fields** [Table 610](#) describes the output fields for the `show class-of-service shared-buffer` command. Output fields are listed in the approximate order in which they appear.

**Table 610: show class-of-service shared-buffer Output Fields**

| Field Name          | Field Description                                                                                                   |
|---------------------|---------------------------------------------------------------------------------------------------------------------|
| <b>Ingress</b>      | Ingress shared buffer configuration.                                                                                |
| <b>Total Buffer</b> | Total buffer space available to the ports in KB. This is the combined dedicated buffer pool and shared buffer pool. |

Table 610: show class-of-service shared-buffer Output Fields (*continued*)

| Field Name                           | Field Description                                                                                                                                                                                                                                                       |
|--------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Dedicated Buffer</b>              | Buffer space allocated to the dedicated buffer pool in KB.                                                                                                                                                                                                              |
| <b>Shared Buffer</b>                 | Buffer space allocated to the shared buffer pool in KB.                                                                                                                                                                                                                 |
| <b>Lossless</b>                      | Buffer space allocated to the lossless traffic buffer pool in KB.                                                                                                                                                                                                       |
| <b>Lossless Headroom</b>             | Buffer space allocated to the lossless headroom traffic buffer pool to support priority-based flow control (PFC) and Ethernet PAUSE in KB. (Ingress ports only.)<br><br><i>NOTE:</i> OCX Series switches do not support PFC.                                            |
| <b>Lossy</b>                         | Buffer space allocated to the lossy (best-effort) traffic buffer pool in KB.                                                                                                                                                                                            |
| <b>Lossless Headroom Utilization</b> | Utilization of the ingress lossless headroom buffer pool. (These fields can help you to determine how much headroom buffer space you need to reserve to support PFC and Ethernet PAUSE for lossless flows.)<br><br><i>NOTE:</i> OCX Series switches do not support PFC. |
| <b>Node Device</b>                   | Index number that identifies the switch. On a QFX3500 switch, this field always has a value of zero (0).                                                                                                                                                                |
| <b>Total</b>                         | Size of the lossless headroom ingress buffer pool in KB.                                                                                                                                                                                                                |
| <b>Used</b>                          | Amount in KB of lossless headroom ingress buffer used.                                                                                                                                                                                                                  |
| <b>Free</b>                          | Amount in KB of lossless headroom ingress buffer free (unused).                                                                                                                                                                                                         |
| <b>Egress</b>                        | Egress shared buffer configuration.                                                                                                                                                                                                                                     |
| <b>Multicast</b>                     | Buffer space allocated to the multicast traffic buffer pool in KB. (Egress ports only.)                                                                                                                                                                                 |

## Sample Output

### show class-of-service shared-buffer

```
user@switch> show class-of-service shared-buffer
```

```
Ingress:
```

```
Total Buffer : 9360.00 KB
Dedicated Buffer : 2158.00 KB
Shared Buffer : 7202.00 KB
 Lossless : 648.18 KB
 Lossless Headroom : 3240.90 KB
 Lossy : 3312.92 KB
```

```
Lossless Headroom Utilization:
```

```
Node Device Total Used Free
0 3240.90 KB 0.00 KB 3240.90 KB
```

Egress:

|                  |   |            |
|------------------|---|------------|
| Total Buffer     | : | 9360.00 KB |
| Dedicated Buffer | : | 2704.00 KB |
| Shared Buffer    | : | 6656.00 KB |
| Lossless         | : | 3328.00 KB |
| Multicast        | : | 1264.64 KB |
| Lossy            | : | 2063.36 KB |

## show pfe filter hw summary

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show pfe filter hw summary                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Command introduced in Junos OS Release 14.1X53-D10 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | <p>Display a summary of the access control list (ACL; also known as firewall filter) ternary content-addressable memory (TCAM) hardware utilization to show the allocated, used, and free TCAM entry space.</p> <p>Command supported on standalone QFX Series switches, QFX5100-only (pure QFX5100) Virtual Chassis Fabric (VCF), QFX5100-only (pure QFX5100) Virtual Chassis (VC), and QFX3500-only (pure QFX3500) VC.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Planning the Number of Firewall Filters to Create on page 5973</a></li> </ul>                                                                                                                                                                                                                                                                                          |
| <b>List of Sample Output</b>    | <a href="#">show pfe filter hw summary on page 7200</a>                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Output Fields</b>            | <a href="#">Table 491</a> lists the output fields for the <b>show pfe filter hw summary</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                |

**Table 611: show pfe filter hw summary Output Fields**

| Field Name       | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Group</b>     | <p>ACL ingress and egress filter groups:</p> <ul style="list-style-type: none"> <li>• iRACL group—ingress routing ACL filter group</li> <li>• iVACL group—ingress VLAN ACL filter group</li> <li>• iPACL group—ingress port ACL filter group</li> <li>• ePACL group—egress port ACL filter group</li> <li>• eVACL group—egress VLAN ACL filter group</li> <li>• eRACL group—egress routing ACL filter group</li> <li>• eRACL IPv6 group—egress IPv6 routing ACL filter group</li> </ul> |
| <b>Group-ID</b>  | Internal identification number of the filter group.                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Allocated</b> | Number of TCAM filter entries allocated to the filter group.                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Used</b>      | Number of TCAM filter entries used by the filter group.                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Free</b>      | Number of TCAM filter entries available for use by the filter group.                                                                                                                                                                                                                                                                                                                                                                                                                    |

## Sample Output

### show pfe filter hw summary

```
user@switch> show pfe filter hw summary
```

| Group                    | Group-ID | Allocated | Used | Free |
|--------------------------|----------|-----------|------|------|
| -----                    |          |           |      |      |
| > Ingress filter groups: |          |           |      |      |
| iRACL group              | 14       | 512       | 4    | 508  |
| iVACL group              | 13       | 512       | 2    | 510  |
| iPACL group              | 12       | 256       | 2    | 254  |
| > Egress filter groups:  |          |           |      |      |
| ePACL group              | 20       | 256       | 3    | 253  |
| eVACL group              | 21       | 256       | 4    | 252  |
| eRACL group              | 22       | 256       | 245  | 11   |
| eRACL IPV6 group         | 24       | 256       | 3    | 253  |



## show pfe next-hop

|                                                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                                | <a href="#">Syntax on page 7201</a><br><a href="#">Syntax (TX Matrix and TX Matrix Plus Routers) on page 7201</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Syntax</b>                                        | <pre>show pfe next-hop &lt;interface <i>interface-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Syntax (TX Matrix and TX Matrix Plus Routers)</b> | <pre>show pfe next-hop &lt;fpc <i>slot</i>&gt; &lt;interface <i>interface-name</i>&gt; &lt;lcc <i>number</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>                           | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>                                   | Display Packet Forwarding Engine next-hop information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                                       | <p><b>none</b>—Display all Packet Forwarding Engine next-hop information.</p> <p><b>fpc <i>slot</i></b>—(TX Matrix and TX Matrix Plus routers only) (Optional) Show the next hops for a Flexible PIC Concentrator (FPC) slot.</p> <p>On a TX Matrix router, if you specify the number of a T640 router by using the <b>lcc <i>number</i></b> option (the recommended method), replace <b><i>slot</i></b> with a value from 0 through 7. Otherwise, replace <b><i>slot</i></b> with a value from 0 through 31. On a TX Matrix Plus router, if you specify the number of a T1600 router by using the <b>lcc <i>number</i></b> option (the recommended method), replace <b><i>slot</i></b> with a value from 0 through 7. Otherwise, replace <b><i>slot</i></b> with a value from 0 through 31. For example, the following commands have the same result:</p> <pre>user@host&gt; show pfe next-hop fpc 1 lcc 1 user@host&gt; show pfe next-hop fpc 9</pre> <p><b>interface <i>interface-name</i></b>—(Optional) Display the Packet Forwarding Engine next-hop interface.</p> <p><b>lcc <i>number</i></b>—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, the slot number of the T640 router (or line-card chassis) that houses the FPC. On a TX Matrix Plus router, the slot number of the T1600 router (or line-card chassis) that houses the FPC. Replace <b><i>number</i></b> with a value from 0 through 3.</p> |
| <b>Required Privilege Level</b>                      | admin                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>                         | <ul style="list-style-type: none"> <li>• <i>Routing Matrix with TXP-T1600 Configuration</i></li> <li>• <i>Routing Matrix with TXP-T1600-3D Configuration</i></li> <li>• <i>Routing Matrix with TXP-T4000-3D Configuration</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

- *Routing Matrix with a TXP-Mixed-LCC-3D Configuration*

**List of Sample Output**    [show pfe next-hop on page 7203](#)  
                                  [show pfe next-hop fpc \(TX Matrix Router\) on page 7203](#)  
                                  [show pfe next-hop fpc \(TX Matrix Plus Router\) on page 7204](#)

**Output Fields** Table 612 lists the output fields for the **show pfe next-hop** command. Output fields are listed in the approximate order in which they appear.

**Table 612: show pfe next-hop Output Fields**

| Field Name    | Field Description                                      |
|---------------|--------------------------------------------------------|
| ID            | The next-hop ID for the entry.                         |
| Type          | The next-hop type for the entry.                       |
| Interface     | The interface to which the next-hop entry is assigned. |
| Protocol      | The protocol type for the next-hop entry.              |
| Encap         | Encapsulation type for the next-hop entry.             |
| Next Hop Addr | Next-hop address for the next-hop entry.               |
| MTU           | MTU value for the nexthop entry.                       |

## Sample Output

### show pfe next-hop

```

user@host> show pfe next-hop
Nexthop Info:
 ID Type Interface Protocol Encap Next Hop Addr MTU
 ---- - - - - - -
 4 Mcast - IPv4 - 0.0.0.0 0
 5 Bcast - IPv4 - - 0
 7 Discard - IPv4 - - 0
 8 MDiscard - IPv4 - - 0
 9 Reject - IPv4 - - 0
 13 Local - IPv4 - 192.168.4.60 0
 14 Resolve fxp0.0 IPv4 Unspecified - 0
 17 Local - IPv4 - 127.0.0.1 0
 18 Unicast fxp0.0 IPv4 Unspecified 192.168.4.254 0
 21 Local - IPv4 - 11.1.0.1 0
 22 Unicast at-0/1/0.0 IPv4 ATM SNAP 11.1.0.2 4482
 ...

```

### show pfe next-hop fpc (TX Matrix Router)

```

user@host> show pfe next-hop fpc 1
Slot 1
Nexthop Info:
 ID Type Interface Next Hop Addr Protocol Encap MTU
 ---- - - - - - -
 5 Mcast - default IPv4 - 0
 6 Bcast - - IPv4 - 0
 8 Discard - - IPv4 - 0
 9 MDiscard - - IPv4 - 0
 13 Mcast - default IPV6 - 0
 17 MDiscard - - IPV6 - 0
 18 Reject - - IPV6 - 0
 24 Discard - - None - 0

```

|     |                   |         |                |      |             |   |
|-----|-------------------|---------|----------------|------|-------------|---|
| 68  | Local             | -       | 192.168.66.113 | IPv4 | -           | 0 |
| 69  | Resolve           | fxp0.0  | -              | IPv4 | Unspecified | 0 |
| 70  | Unicast           | fxp0.0  | 192.168.71.254 | IPv4 | Unspecified | 0 |
| 256 | Local             | -       | 10.71.71.1     | IPv4 | -           | 0 |
| 257 | Local             | -       | 127.0.0.1      | IPv4 | -           | 0 |
| 258 | Mcast.local..1    | default | -              | IPv4 | Unspecified | 0 |
| 259 | Bcast.local..1    | -       | -              | IPv4 | Unspecified | 0 |
| 261 | Discard.local..1  | -       | -              | IPv4 | Unspecified | 0 |
| 262 | MDiscard.local..1 | -       | -              | IPv4 | Unspecified | 0 |
| 269 | Mcast.local..1    | default | -              | IPv6 | Unspecified | 0 |
| 271 | Discard.local..1  | -       | -              | IPv6 | Unspecified | 0 |
| ... |                   |         |                |      |             |   |

**show pfe next-hop fpc (TX Matrix Plus Router)**

user@host&gt; show pfe next-hop fpc 0

Slot 0

| ID    | Type     | Interface | Next Hop Addr       | Protocol | Encap       | MTU |
|-------|----------|-----------|---------------------|----------|-------------|-----|
| ----- |          |           |                     |          |             |     |
| 31    | Mcast    | -         | default             | IPv4     | -           | 0   |
| 32    | Bcast    | -         | -                   | IPv4     | -           | 0   |
| 34    | Discard  | -         | -                   | IPv4     | -           | 0   |
| 35    | MDiscard | -         | -                   | IPv4     | -           | 0   |
| 36    | Reject   | -         | -                   | IPv4     | -           | 0   |
| 39    | Mcast    | -         | default             | IPv6     | -           | 0   |
| 42    | Discard  | -         | -                   | IPv6     | -           | 0   |
| 43    | MDiscard | -         | -                   | IPv6     | -           | 0   |
| 44    | Reject   | -         | -                   | IPv6     | -           | 0   |
| 49    | Receive  | -         | -                   | MPLS     | -           | 0   |
| 50    | Discard  | -         | -                   | MPLS     | -           | 0   |
| 111   | Mcast    | .local..1 | default             | IPv4     | Unspecified | 0   |
| 112   | Bcast    | .local..1 | -                   | IPv4     | Unspecified | 0   |
| 114   | Discard  | .local..1 | -                   | IPv4     | Unspecified | 0   |
| 115   | MDiscard | .local..1 | -                   | IPv4     | Unspecified | 0   |
| 116   | Reject   | .local..1 | -                   | IPv4     | Unspecified | 0   |
| 119   | Mcast    | .local..1 | default             | IPv6     | Unspecified | 0   |
| 122   | Discard  | .local..1 | -                   | IPv6     | Unspecified | 0   |
| 123   | MDiscard | .local..1 | -                   | IPv6     | Unspecified | 0   |
| 124   | Reject   | .local..1 | -                   | IPv6     | Unspecified | 0   |
| 191   | Mcast    | .local..2 | default             | IPv4     | Unspecified | 0   |
| 192   | Bcast    | .local..2 | -                   | IPv4     | Unspecified | 0   |
| 194   | Discard  | .local..2 | -                   | IPv4     | Unspecified | 0   |
| 195   | MDiscard | .local..2 | -                   | IPv4     | Unspecified | 0   |
| 196   | Reject   | .local..2 | -                   | IPv4     | Unspecified | 0   |
| 322   | Local    | -         | 10.1.0.5            | IPv4     | -           | 0   |
| 323   | Resolve  | bcm0.0    | -                   | IPv4     | Unspecified | 0   |
| 326   | Local    | -         | 129.0.0.5           | IPv4     | -           | 0   |
| 327   | Resolve  | bcm0.0    | -                   | IPv4     | Unspecified | 0   |
| 328   | Local    | -         | fe80::201:ff:fe01:5 | IPv6     | -           | 0   |
| 329   | Receive  | bcm0.0    | ff02::1:ff01:5      | IPv6     | Unspecified | 0   |
| 330   | Receive  | bcm0.0    | fe80::              | IPv6     | Unspecified | 0   |
| 331   | Resolve  | bcm0.0    | -                   | IPv6     | Unspecified | 0   |
| 332   | Local    | -         | fec0::a:1:0:5       | IPv6     | -           | 0   |
| 333   | Receive  | bcm0.0    | ff02::1:ff00:5      | IPv6     | Unspecified | 0   |
| 334   | Receive  | bcm0.0    | fec0::              | IPv6     | Unspecified | 0   |
| 335   | Resolve  | bcm0.0    | -                   | IPv6     | Unspecified | 0   |
| 348   | Local    | -         | 192.168.178.4       | IPv4     | -           | 0   |
| 349   | Resolve  | em0.0     | -                   | IPv4     | Unspecified | 0   |

|     |          |             |                      |      |             |   |
|-----|----------|-------------|----------------------|------|-------------|---|
| 350 | Unicast  | em0.0       | 192.168.178.126      | IPv4 | Unspecified | 0 |
| 357 | Local    | -           | fe80::201:1ff:fe01:5 | IPv6 | -           | 0 |
| 512 | Local    | -           | 10.255.178.11        | IPv4 | -           | 0 |
| 513 | Local    | -           | 127.0.0.1            | IPv4 | -           | 0 |
| 515 | Local    | -           | abcd::10:255:178:11  | IPv6 | -           | 0 |
| 516 | Local    | -           | fe80::200:ff:fe00:0  | IPv6 | -           | 0 |
| 517 | Local    | -           | 127.0.0.1            | IPv4 | -           | 0 |
| 518 | Mcast    | .local..3   | default              | IPv4 | Unspecified | 0 |
| 519 | Bcast    | .local..3   | -                    | IPv4 | Unspecified | 0 |
| 521 | Discard  | .local..3   | -                    | IPv4 | Unspecified | 0 |
| 522 | MDiscard | .local..3   | -                    | IPv4 | Unspecified | 0 |
| 523 | Reject   | .local..3   | -                    | IPv4 | Unspecified | 0 |
| 531 | Mcast    | .local..3   | default              | IPv6 | Unspecified | 0 |
| 533 | Discard  | .local..3   | -                    | IPv6 | Unspecified | 0 |
| 534 | MDiscard | .local..3   | -                    | IPv6 | Unspecified | 0 |
| 535 | Reject   | .local..3   | -                    | IPv6 | Unspecified | 0 |
| 539 | Mgroup   | -           | -                    | IPv4 | -           | 0 |
| 540 | Bcast    | ge-15/0/3.0 | -                    | IPv4 | Ethernet    | 0 |
| 541 | Receive  | ge-15/0/3.0 | 14.2.1.0             | IPv4 | Ethernet    | 0 |
| 542 | Local    | -           | 14.2.1.1             | IPv4 | -           | 0 |
| 543 | Resolve  | ge-15/0/3.0 | -                    | IPv4 | Ethernet    | 0 |
| 544 | Bcast    | ge-31/0/4.0 | -                    | IPv4 | Ethernet    | 0 |
| 545 | Receive  | ge-31/0/4.0 | 14.1.1.0             | IPv4 | Ethernet    | 0 |
| 546 | Local    | -           | 14.1.1.1             | IPv4 | -           | 0 |
| 547 | Resolve  | ge-31/0/4.0 | -                    | IPv4 | Ethernet    | 0 |
| 548 | Unicast  | ge-31/0/4.0 | 14.1.1.2             | IPv4 | Ethernet    | 0 |
| 549 | Unicast  | ge-15/0/3.0 | 14.2.1.2             | IPv4 | Ethernet    | 0 |
| 550 | Bcast    | ae1.0       | -                    | IPv4 | Ethernet    | 0 |
| 551 | Receive  | ae1.0       | 11.1.1.0             | IPv4 | Ethernet    | 0 |
| 552 | Local    | -           | 11.1.1.1             | IPv4 | -           | 0 |
| 553 | Resolve  | ae1.0       | -                    | IPv4 | Ethernet    | 0 |
| 554 | Aggreg.  | ae1.0       | -                    | IPv4 | Ethernet    | 0 |
| 555 | Unicast  | ge-23/0/8.0 | 11.1.1.2             | IPv4 | Ethernet    | 0 |
| 556 | Unicast  | ge-7/0/9.0  | 11.1.1.2             | IPv4 | Ethernet    | 0 |
| 557 | Aggreg.  | ae1.0       | -                    | MPLS | Ethernet    | 0 |
| 558 | Unicast  | ge-23/0/8.0 | -                    | MPLS | Ethernet    | 0 |
| 559 | Unicast  | ge-7/0/9.0  | -                    | MPLS | Ethernet    | 0 |
| 560 | Aggreg.  | ae1.0       | -                    | MPLS | Ethernet    | 0 |
| 561 | Unicast  | ge-23/0/8.0 | -                    | MPLS | Ethernet    | 0 |
| 562 | Unicast  | ge-7/0/9.0  | -                    | MPLS | Ethernet    | 0 |

## show pfe route

---

|                                                      |                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                                | <a href="#">Syntax on page 7206</a><br><a href="#">Syntax (EX Series Switch and QFX Series) on page 7206</a><br><a href="#">Syntax (QFX Series) on page 7206</a><br><a href="#">Syntax (MX Series) on page 7206</a><br><a href="#">Syntax (TX Matrix and TX Matrix Plus Routers) on page 7206</a>                                                                                                                               |
| <b>Syntax</b>                                        | <pre>show pfe route &lt;&lt;inet6   ip   iso&gt; &lt;prefix prefix&gt;   &lt;table &lt;table-name&gt; &lt;index index&gt; &lt;prefix prefix&gt;&gt;&gt; &lt;mpls&gt; &lt;summary&gt;</pre>                                                                                                                                                                                                                                      |
| <b>Syntax (EX Series Switch and QFX Series)</b>      | <pre>show pfe route &lt;&lt;inet6   ip&gt; &lt;prefix prefix&gt;   &lt;table &lt;table-name&gt; &lt;index index&gt; &lt;prefix prefix&gt;&gt;&gt; &lt;mpls&gt; &lt;summary&gt;</pre>                                                                                                                                                                                                                                            |
| <b>Syntax (QFX Series)</b>                           | <pre>show pfe route &lt;&lt;inet6   ip&gt; &lt;prefix prefix&gt;   &lt;table &lt;table-name&gt; &lt;index index&gt; &lt;prefix prefix&gt;&gt; &lt;hw (host   lpm   multicast)&gt;&gt; &lt;&lt;clnp&gt; &lt;prefix prefix&gt;   &lt;table &lt;table-name&gt; &lt;index index&gt; &lt;prefix prefix&gt;&gt;&gt; &lt;mpls&gt; &lt;summary&gt; &lt;hw&gt;</pre>                                                                     |
| <b>Syntax (MX Series)</b>                            | <pre>show pfe route &lt;&lt;inet6   ip&gt; &lt;prefix prefix&gt;   &lt;table &lt;table-name&gt; &lt;index index&gt; &lt;prefix prefix&gt;&gt;&gt; &lt;dhcp&gt; &lt;mpls&gt; &lt;summary&gt;</pre>                                                                                                                                                                                                                               |
| <b>Syntax (TX Matrix and TX Matrix Plus Routers)</b> | <pre>show pfe route &lt;fpc slot&gt; &lt;&lt;inet6   ip   iso&gt; &lt;prefix prefix&gt;   &lt;table &lt;table-name&gt; &lt;index index&gt; &lt;prefix prefix&gt;&gt;&gt; &lt;lcc number&gt; &lt;mpls&gt; &lt;summary&gt;</pre>                                                                                                                                                                                                  |
| <b>Release Information</b>                           | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.<br>Command introduced in Junos OS Release 13.3 for the MX Series.<br>Command option <b>hw</b> introduced in Junos OS Release 14.1X53-D10 for the QFX Series.<br>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| <b>Description</b>                                   | Display the routes in the Packet Forwarding Engine forwarding table. The Packet Forwarding Engine forwards packets between input and output interfaces.                                                                                                                                                                                                                                                                         |



**NOTE:** The Routing Engine maintains a master copy of the forwarding table. It copies the forwarding table to the Packet Forwarding Engine, which is the part of the router or switch responsible for forwarding packets. To display the routes in the Routing Engine forwarding table, use the **show route forwarding table** command. For more information, see the [CLI Explorer](#).

**Options** **none**—Display all Packet Forwarding Engine forwarding table information.

**clnp**—(Optional) Show International Standards Organization (ISO) connectionless-mode network protocol (CLNP) route table information.

**dhcp**—(Optional) Display Packet Forwarding Engine DHCP-Snooping route table information.

**fpc slot**—(TX Matrix and TX Matrix Plus routers only) (Optional) Show the next hops for a Flexible PIC Concentrator (FPC) slot.

- On a TX Matrix router, if you specify the number of a T640 router by using the **lcc number** option (the recommended method), replace **slot** with a value from **0** through **7**. Otherwise, replace **slot** with a value from **0** through **31**.
- On a TX Matrix Plus router, if you specify the number of a T1600 router by using the **lcc number** option (the recommended method), replace **slot** with a value from **0** through **7**. Otherwise, replace **slot** with a value from **0** through **31**.
- On a TX Matrix Plus router in the TXP-T1600-3D, TXP-T4000-3D, or TXP-Mixed-LCC-3D configuration, if you specify the number of a T1600 or T4000 router by using the **lcc number** option (the recommended method), replace **slot** with a value from **0** through **7**. Otherwise, replace **slot** with a value from **0** through **63**.

For example, the following commands have the same result:

```
user@host> show pfe route fpc 1 lcc 1
user@host> show pfe route fpc 9
```

**host**—(QFX standalone switches, pure mode QFX5100-only VCF and VC, and pure mode QFX3500-only VC) (Optional) Display host routes installed in the on-chip hardware table.

**hw**—(QFX standalone switches, pure mode QFX5100-only VCF and VC, and pure mode QFX3500-only VC) (Optional) Display routes installed in the on-chip hardware table (as opposed to displaying routes from the routing table and the PFE forwarding table before they are installed in the hardware).

**index index**—(Optional) Display table index.

**inet6**—(Optional) Display Packet Forwarding Engine IPv6 routes.

**ip**—(Optional) Display Packet Forwarding Engine IPv4 routes.

**iso**—(Optional) Display ISO version routing tables.

**lcc *number***—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, the slot number of the T640 router (or line-card chassis) that houses the FPC. On a TX Matrix Plus router, the slot number of the T1600 router (or line-card chassis) that houses the FPC. Replace *number* with a value from 0 through 3.

**mpls**—(Optional) Display Packet Forwarding Engine MPLS information.

**multicast**—(QFX standalone switches, pure mode QFX5100-only VCF and VC, and pure mode QFX3500-only VC) (Optional) Display multicast routes installed in the on-chip hardware table.

**prefix *prefix***—(Optional) IPv4 or IPv6 prefix for which to show table entries.

**summary**—(Optional) Display summary of Packet Forwarding Engine information.

**table <*table-name*>**—(Optional) Display table information.

**Required Privilege Level**

admin

**Related Documentation**

- *Routing Matrix with TXP-T1600 Configuration*
- *Routing Matrix with TXP-T1600-3D Configuration*
- *Routing Matrix with TXP-T4000-3D Configuration*
- *Routing Matrix with a TXP-Mixed-LCC-3D Configuration*

**List of Sample Output**

[show pfe route ip on page 7210](#)  
[show pfe route iso on page 7210](#)  
[show pfe route lcc summary \(TX Matrix Router\) on page 7210](#)  
[show pfe route lcc summary \(TX Matrix Plus Router\) on page 7212](#)  
[show pfe route summary \(MX Series Router\) on page 7213](#)  
[show pfe route summary hw \(QFX Series, EX4600 Switches, OCX Series\) on page 7213](#)  
[show pfe route ip hw host \(QFX Series\) on page 7214](#)

**Output Fields**

Table 613 lists the output fields for the **show pfe route** command. Output fields are listed in the approximate order in which they appear.

**Table 613: show pfe route Output Fields**

| Field Name  | Field Description                  |
|-------------|------------------------------------|
| Destination | Destination address for the entry. |
| NH IP Addr  | Next-hop IP address for the entry. |
| Type        | Next-hop type for the entry        |
| NH ID       | Next-hop ID for the entry          |



Table 613: show pfe route Output Fields (*continued*)

| Field Name       | Field Description                                  |
|------------------|----------------------------------------------------|
| <b>Encap</b>     | Encapsulation type for the next-hop entry.         |
| <b>Interface</b> | Interface to which the next-hop entry is assigned. |

Table 614 lists the output fields for the QFX Series **show pfe route** hardware table (**hw**) commands. Output fields are listed in the approximate order in which they appear.

Table 614: QFX Series, EX4600 switches, and OCX Series show pfe route Hardware Table Output Fields

| Field Name             | Field Description                                                                                                                                                                                   |
|------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Max</b>             | Maximum routing entries per route type.                                                                                                                                                             |
| <b>Used</b>            | Number of routing entries consumed per route type.                                                                                                                                                  |
| <b>Free</b>            | Number of unused routing entries per route type.                                                                                                                                                    |
| <b>% Free</b>          | Percentage of unused routing entries per route type.                                                                                                                                                |
| <b>Rtt</b>             | Internal routing engine index number of the route table.                                                                                                                                            |
| <b>VRF</b>             | Internal hardware index number for the corresponding route table.                                                                                                                                   |
| <b>Destination</b>     | Destination address for the entry.                                                                                                                                                                  |
| <b>Type</b>            | ( <b>show pfe route summary hw</b> )—Route type for the entry: IPv4 or IPv6 route, and host, LPM, or multicast route.<br><br>( <b>show pfe route (ip   inet6) hw</b> )—Next-hop type for the entry. |
| <b>NH ID</b>           | Next-hop ID for the entry                                                                                                                                                                           |
| <b>Interface</b>       | Interface to which the next-hop entry is assigned.                                                                                                                                                  |
| <b>HW NH-ID</b>        | Internal hardware index number of the next-hop.                                                                                                                                                     |
| <b>Src-MAC-Address</b> | Source MAC address.                                                                                                                                                                                 |
| <b>Port</b>            | Port number.                                                                                                                                                                                        |
| <b>Dst-MAC-Address</b> | Destination MAC address.                                                                                                                                                                            |
| <b>VLAN</b>            | ID of the multicast group VLAN.                                                                                                                                                                     |
| <b>GROUP</b>           | Internal hardware index number of the multicast group next-hop.                                                                                                                                     |

Table 614: QFX Series, EX4600 switches, and OCX Series show pfe route Hardware Table Output Fields (*continued*)

| Field Name | Field Description                             |
|------------|-----------------------------------------------|
| CLASS      | Internal class number of the multicast group. |

## Sample Output

### show pfe route ip

```
user@host> show pfe route ip
```

```
IPv4 Route Table 0, default.0, 0x0:
Destination NH IP Addr Type NH ID Interface

default 127.0.0.1 Discard 8
127.0.0.1 127.0.0.1 Local 256
172.16/12 192.168.71.254 Unicast 68 fxp0.0
192.168.0/18 192.168.71.254 Unicast 68 fxp0.0
192.168.40/22 192.168.71.254 Unicast 68 fxp0.0
192.168.64/18 192.168.71.254 Unicast 68 fxp0.0
192.168.64/21 192.168.71.254 Resolve 67 fxp0.0
192.168.71.249 192.168.71.249 Local 66
192.168.220.0/30 192.168.220.0 Resolve 303 fe-0/0/0.0
192.168.220.0 192.168.220.0 Receive 301 fe-0/0/0.0
224.0.0.1 Mcast 5
255.255.255.255 Bcast 6
```

```
...
```

### show pfe route iso

```
user@host# show pfe route iso
```

```
CLNS Route Table 0, CLNP.0, 0x0:
Destination Type NH ID Interface

default Reject 60
47.0005.80ff.f800.0000.0108.0001.0102.5508.2159/152 Local 514
49.0001.00a0.c96b.c491/72 Local 536
```

### show pfe route lcc summary (TX Matrix Router)

```
user@host> show pfe route lcc 2 summary
```

```
Slot 0
```

```
IPv4 Route Tables:
Index Routes Size(b)

Default 43 3081
1 4 281
```

```
MPLS Route Tables:
Index Routes Size(b)

Default 1 68
```

## IPv6 Route Tables:

| Index   | Routes | Size(b) |
|---------|--------|---------|
| -----   | -----  | -----   |
| Default | 9      | 717     |
| 1       | 5      | 389     |

Slot 1

## IPv4 Route Tables:

| Index   | Routes | Size(b) |
|---------|--------|---------|
| -----   | -----  | -----   |
| Default | 43     | 3081    |
| 1       | 4      | 281     |

## MPLS Route Tables:

| Index   | Routes | Size(b) |
|---------|--------|---------|
| -----   | -----  | -----   |
| Default | 1      | 68      |

## IPv6 Route Tables:

| Index   | Routes | Size(b) |
|---------|--------|---------|
| -----   | -----  | -----   |
| Default | 9      | 717     |
| 1       | 5      | 389     |

Slot 16

## IPv4 Route Tables:

| Index   | Routes | Size(b) |
|---------|--------|---------|
| -----   | -----  | -----   |
| Default | 41     | 2938    |
| 1       | 4      | 281     |

## MPLS Route Tables:

| Index   | Routes | Size(b) |
|---------|--------|---------|
| -----   | -----  | -----   |
| Default | 1      | 68      |

## IPv6 Route Tables:

| Index   | Routes | Size(b) |
|---------|--------|---------|
| -----   | -----  | -----   |
| Default | 9      | 717     |
| 1       | 5      | 389     |

Slot 17

## IPv4 Route Tables:

| Index   | Routes | Size(b) |
|---------|--------|---------|
| -----   | -----  | -----   |
| Default | 41     | 2938    |
| 1       | 4      | 281     |

## MPLS Route Tables:

| Index   | Routes | Size(b) |
|---------|--------|---------|
| -----   | -----  | -----   |
| Default | 1      | 68      |

## IPv6 Route Tables:

| Index   | Routes | Size(b) |
|---------|--------|---------|
| -----   | -----  | -----   |
| Default | 9      | 717     |
| 1       | 5      | 389     |

### show pfe route lcc summary (TX Matrix Plus Router)

user@host> show pfe route lcc 2 summary

Slot 0

#### IPv4 Route Tables:

| Index   | Routes | Size(b) |
|---------|--------|---------|
| -----   | -----  | -----   |
| Default | 25     | 2266    |
| 1       | 9      | 815     |
| 2       | 6      | 545     |
| 3       | 5      | 453     |
| 4       | 15     | 1371    |
| 5       | 5      | 453     |
| 6       | 13     | 1187    |

#### MPLS Route Tables:

| Index   | Routes | Size(b) |
|---------|--------|---------|
| -----   | -----  | -----   |
| Default | 1      | 88      |
| 4       | 5      | 452     |

#### IPv6 Route Tables:

| Index   | Routes | Size(b) |
|---------|--------|---------|
| -----   | -----  | -----   |
| Default | 7      | 697     |
| 1       | 13     | 1305    |
| 3       | 4      | 385     |
| 4       | 4      | 385     |
| 5       | 4      | 385     |
| 6       | 18     | 1833    |

Slot 6

#### IPv4 Route Tables:

| Index   | Routes | Size(b) |
|---------|--------|---------|
| -----   | -----  | -----   |
| Default | 25     | 2266    |
| 1       | 9      | 815     |
| 2       | 6      | 545     |
| 3       | 5      | 453     |
| 4       | 15     | 1371    |
| 5       | 5      | 453     |
| 6       | 13     | 1187    |

#### MPLS Route Tables:

| Index   | Routes | Size(b) |
|---------|--------|---------|
| -----   | -----  | -----   |
| Default | 1      | 88      |
| 4       | 5      | 452     |

#### IPv6 Route Tables:

| Index | Routes | Size(b) |
|-------|--------|---------|
|-------|--------|---------|

```

Default 7 697
1 13 1305
3 4 385
4 4 385
5 4 385
6 18 1833
...

```

### show pfe route summary (MX Series Router)

```
user@host> show pfe route summary
```

```
Slot 0
```

```

DHCP-Snooping Route Tables:
Index Routes Size(b)

Default 1 144

```

```

IPv4 Route Tables:
Index Routes Size(b)

Default 25 2266
1 9 815
2 6 545
3 5 453
4 15 1371
5 5 453
6 13 1187

```

```

MPLS Route Tables:
Index Routes Size(b)

Default 1 88
4 5 452

```

```

IPv6 Route Tables:
Index Routes Size(b)

Default 7 697
1 13 1305
3 4 385
4 4 385
5 4 385
6 18 1833

```

```
...
```

### show pfe route summary hw (QFX Series, EX4600 Switches, OCX Series)

```
user@switch> show pfe route summary hw
```

```
Slot 0
```

```
Unit: 0
```

```
Profile active: l2-profile-three
```

```

Type Max Used Free % free

IPv4 Host 8192 103 8073 98.55
IPv4 LPM 16384 9 16369 99.91
IPv4 Mcast 4096 2 4037 98.56

```

|                |      |   |      |       |
|----------------|------|---|------|-------|
| IPv6 Host      | 4096 | 6 | 4037 | 98.56 |
| IPv6 LPM(< 64) | 8192 | 3 | 8185 | 99.91 |
| IPv6 LPM(> 64) | 256  | 1 | 255  | 99.61 |
| IPv6 Mcast     | 2048 | 0 | 2019 | 98.58 |

**show pfe route ip hw host (QFX Series)**

```
user@switch> show pfe route ip host hw
```

```
Slot 0
```

```
Unit: 0
```

```
IPv4 Host entries present: 103
```

| Rtt     | VRF    | Destination                                   | Type                   | NH-ID | Interface   |
|---------|--------|-----------------------------------------------|------------------------|-------|-------------|
|         |        | HW NH-ID Src-MAC-Address Port Dst-MAC-Address |                        |       |             |
| 4       | 3      | 255.255.255.255                               | Bcast                  | 1695  | .local. .4  |
| ifl 550 | 100003 | 00:00:00:01:02:03 127                         | 00:00:00:01:02:03      |       |             |
| 0       | 1      | 200.1.1.42                                    | Unicast                | 1743  | et-0/1/1 .0 |
| ifl 559 | 100268 | 84:18:88:de:96:fd 53                          | 00:00:00:21:12:23      |       |             |
| 0       | 1      | 200.1.1.56                                    | Unicast                | 1743  | et-0/1/1 .0 |
| ifl 559 | 100268 | 84:18:88:de:96:fd 53                          | 00:00:00:21:12:23      |       |             |
| 0       | 1      | 200.1.1.61                                    | Unicast                | 1743  | et-0/1/1 .0 |
| ifl 559 | 100268 | 84:18:88:de:96:fd 53                          | 00:00:00:21:12:23      |       |             |
| 0       | 1      | 11.1.1.2                                      | Unicast                | 1743  | et-0/1/1 .0 |
| ifl 559 | 100268 | 84:18:88:de:96:fd 53                          | 00:00:00:21:12:23      |       |             |
| 0       | 1      | 200.1.1.73                                    | Unicast                | 1743  | et-0/1/1 .0 |
| ifl 559 | 100268 | 84:18:88:de:96:fd 53                          | 00:00:00:21:12:23      |       |             |
| 0       | 1      | 200.1.1.76                                    | Unicast                | 1743  | et-0/1/1 .0 |
| ifl 559 | 100268 | 84:18:88:de:96:fd 53                          | 00:00:00:21:12:23      |       |             |
| 0       | 1      | 200.1.1.18                                    | Unicast                | 1743  | et-0/1/1 .0 |
| ifl 559 | 100268 | 84:18:88:de:96:fd 53                          | 00:00:00:21:12:23      |       |             |
| 0       | 1      | 200.1.1.5                                     | Unicast                | 1743  | et-0/1/1 .0 |
| ifl 559 | 100268 | 84:18:88:de:96:fd 53                          | 00:00:00:21:12:23      |       |             |
| 0       | 1      | 200.1.1.23                                    | Unicast                | 1743  | et-0/1/1 .0 |
| ifl 559 | 100268 | 84:18:88:de:96:fd 53                          | 00:00:00:21:12:23      |       |             |
| 0       | 1      | 101.1.1.255                                   | Bcast                  | 1664  | ae0 .0      |
| ifl 544 | 100003 | 00:00:00:01:02:03 127                         | 00:00:00:01:02:03      |       |             |
| 0       | 1      | 200.1.1.40                                    | Unicast                | 1743  | et-0/1/1 .0 |
| ifl 559 | 100268 | 84:18:88:de:96:fd 53                          | 00:00:00:21:12:23      |       |             |
| 0       | 1      | 200.1.1.58                                    | Unicast                | 1743  | et-0/1/1 .0 |
| ifl 559 | 100268 | 84:18:88:de:96:fd 53                          | 00:00:00:21:12:23. . . |       |             |
| . . .   |        |                                               |                        |       |             |

## show pfe terse

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 7215</a><br><a href="#">Syntax (TX Matrix and TX Matrix Plus Router) on page 7215</a><br><a href="#">Syntax (MX Series Router) on page 7215</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Syntax</b>                                       | show pfe terse                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Syntax (TX Matrix and TX Matrix Plus Router)</b> | show pfe terse<br><lcc <i>number</i>   scc><br><sfc <i>number</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Syntax (MX Series Router)</b>                    | show pfe terse<br><all-members><br><local><br><member <i>member-id</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>                          | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.<br>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>                                  | Display Packet Forwarding Engine status information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                                      | <p><b>none</b>—Display brief information about the Packet Forwarding Engine.</p> <p><b>all-members</b>—(MX Series routers only) (Optional) Display Packet Forwarding Engine status information for all members in the Virtual Chassis configuration.</p> <p><b>lcc <i>number</i></b>—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display Packet Forwarding Engine information for a specific T640 router (or line-card chassis) that is connected to a TX Matrix router. On a TX Matrix Plus router, display Packet Forwarding Engine information for a specific T1600 router (or line-card chassis) that is connected to a TX Matrix Plus router. Replace <i>number</i> with a value from 0 through 3.</p> <p><b>local</b>—(MX Series routers only) (Optional) Display Packet Forwarding Engine status information for the local Virtual Chassis member.</p> <p><b>member <i>member-id</i></b>—(MX Series routers only) (Optional) Display Packet Forwarding Engine status information for the specified member of the Virtual Chassis configuration. Replace <i>member-id</i> with a value of 0 or 1.</p> <p><b>scc</b>—(TX Matrix routers only) (Optional) Display Packet Forwarding Engine information for the TX Matrix router (or switch-card chassis).</p> <p><b>sfc</b>—(TX Matrix Plus routers only) (Optional) Display Packet Forwarding Engine information for the TX Matrix Plus router (or switch-fabric chassis).</p> |
| <b>Required Privilege Level</b>                     | admin                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

- List of Sample Output    [show pfe terse \(TX Matrix Router\) on page 7216](#)  
                              [show pfe terse \(TX Matrix Plus Router\) on page 7216](#)  
                              [show pfe terse sfc \(TX Matrix Plus Router\) on page 7216](#)

## Sample Output

### show pfe terse (TX Matrix Router)

```
user@host> show pfe terse
Slot Type Slot State Flags Uptime
 0 SFM Present Online 0x0bf 01:25:42
 2 SFM Present Online 0x0bf 01:25:40
 0 FPC Present Online 0x102 01:25:57
 1 FPC Present Online 0x102 01:25:55
 2 FPC Present Online 0x102 01:25:53
```

### show pfe terse (TX Matrix Plus Router)

```
user@host> show pfe terse
sfc0-re0:

Slot Type Slot State Uptime
 0 LCC Present Online 2d 05:26

lcc0-re0:

Slot Type Slot State Uptime
 0 GFPC Present Online 2d 05:25
 1 GFPC Present Online 2d 05:25
```

### show pfe terse sfc (TX Matrix Plus Router)

```
user@host> show pfe terse sfc 0
sfc0-re0:

Slot Type Slot State Uptime
 0 LCC Present Online 2d 05:25
```



## show pfe version

---

|                                 |                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show pfe version <brief   detail>                                                                                                                                                            |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.<br>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| <b>Description</b>              | Display Packet Forwarding Engine version information.                                                                                                                                        |
| <b>Options</b>                  | brief   detail—Display the specified level of output.                                                                                                                                        |
| <b>Required Privilege Level</b> | admin                                                                                                                                                                                        |
| <b>List of Sample Output</b>    | <a href="#">show pfe version brief on page 7217</a><br><a href="#">show pfe version detail on page 7217</a>                                                                                  |

### Sample Output

#### show pfe version brief

```
user@host> show pfe version brief
PFED release 11.1D0 built by builder on 2010-11-11 05:16:11 UTC
```

#### show pfe version detail

```
user@host> show pfe version detail
PFED release 11.1D0 built by builder on 2010-11-11 05:16:11 UTC

junos-core01.juniper.net:/volume/build/junos/rpd_feb11/11.1/development/20101111.0/obj-i386/
junos/usr.sbin/pfed
```

## show interfaces voq

---

**Syntax**    `show interfaces voq interface-name`  
              `<forwarding-class forwarding-class-name>`  
              `<non-zero>`  
              `<source-fpc source-fpc-number>`

**Release Information**    Command introduced in Junos OS Release 14.1 for the PTX Series Routers  
                              Command introduced in Junos OS Release 15.1X53-D20 for QFX10000 switches.

**Description**    Display the random early detection (RED) drop statistics from all ingress Packet Forwarding Engines associated with the specified physical egress interface. In the VOQ architecture, egress output queues (shallow buffers) buffer data in virtual queues on ingress Packet Forwarding Engines. In cases of congestion, you can use this command to identify which ingress Packet Forwarding Engine is the source of RED-dropped packets contributing to congestion.



**NOTE:** On the PTX Series routers and QFX10000 switches, these statistics include tail-dropped packets.

---

**Options**    **interface *interface-name***—Display the ingress VOQ RED drop statistics for the specified egress interface.

**forwarding-class *forwarding-class-name***—Display VOQ RED drop statistics for a specified forwarding class.

**non-zero**—Display only non-zero VOQ RED drop statistics counters.

**source-fpc *source-fpc-number***—Display VOQ RED drop statistics for the specified source FPC.

**Additional Information**

- On PTX Series routers, you can display VOQ statistics for only the WAN physical interface.
- VOQ statistics for aggregated physical interfaces are not supported. Statistics for an aggregated interface are the summation of the queue statistics of the child links of that aggregated interface. You can use the **show interfaces queue** command to identify the child link which is experiencing congestion and then view the VOQ statistics on the respective child link using the **show interfaces voq** command.

For information on virtual output queuing on PTX routers, see *Understanding Virtual Output Queues on PTX Series Packet Transport Routers*. For information on virtual output queueing on QFX10000 switches, see [“Understanding CoS Virtual Output Queues \(VOQs\) on QFX10000 Switches” on page 6724](#).

**Required Privilege Level**    view

- Related Documentation**
- [Understanding Virtual Output Queues on PTX Series Packet Transport Routers](#)
  - [Understanding CoS Virtual Output Queues \(VOQs\) on QFX10000 Switches on page 6724](#)

**List of Sample Output**

[show interfaces voq \(For a Specific Physical Interface\) \(PTX Series Routers\) on page 7220](#)  
[show interfaces voq \(For a Specific Physical Interface\) \(QFX10000 Switches\) on page 7224](#)  
[show interfaces voq et-7/0/0 \(For a Specific Forwarding Class\) on page 7226](#)  
[show interfaces voq et-5/0/12 \(For a Specific Source FPC\) on page 7227](#)  
[show interfaces voq et-5/0/12 \(For a Specific Forwarding Class and Source FPC\) on page 7229](#)  
[show interfaces voq et-7/0/0 \(Non-Zero\) on page 7229](#)  
[show interfaces voq et-7/0/0 \(For a Specific Forwarding Class and Non-Zero\) on page 7230](#)

**Output Fields** [Table 615](#) lists the output fields for the show interfaces queue command. Output fields are listed in the approximate order in which they appear.

**Table 615: show interfaces voq Output Fields**

| Field Name          | Field Description                                                                                                                                                                            |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Physical interface  | Name of the physical interface.                                                                                                                                                              |
| Enabled             | State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .                                                         |
| Interface index     | Physical interface's index number, which reflects its initialization sequence.                                                                                                               |
| SNMP ifIndex        | SNMP index number for the interface.                                                                                                                                                         |
| Queue               | Egress queue number.                                                                                                                                                                         |
| Forwarding classes  | Forwarding class name.                                                                                                                                                                       |
| FPC number          | Number of the Flexible PIC Concentrator (FPC) located on ingress.                                                                                                                            |
| PFE                 | Number of the Packet Forwarding Engine providing virtual output queues on the ingress.                                                                                                       |
| RED-dropped packets | Number of packets per second (pps) dropped because of random early detection (RED).<br><br><b>NOTE:</b> On the PTX Series routers, these statistics include tail-dropped packets.            |
| RED-dropped bytes   | Number of bytes per second dropped because of RED. The byte counts vary by interface hardware.<br><br><b>NOTE:</b> On the PTX Series routers, these statistics include tail-dropped packets. |

## Sample Output

### show interfaces voq (For a Specific Physical Interface) (PTX Series Routers)

The following example shows ingress RED-dropped statistics for the egress Ethernet interface configured on port 0 of Physical Interface Card (PIC) 0, located on the FPC in slot 7.

The sample output below shows that the cause of the congestion is ingress Packet Forwarding Engine PFE 0, which resides on FPC number 4, as denoted by the count of RED-dropped packets and RED-dropped bytes for egress queue 0, forwarding classes best-effort and egress queue 3, forwarding class network control.

```

user@host> show interfaces voq et-7/0/0
Physical interface: et-7/0/0, Enabled, Physical link is Up
 Interface index: 155, SNMP ifIndex: 699

Queue: 0, Forwarding classes: best-effort

FPC number: 1
 PFE: 0
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
 PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
 PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
 PFE: 3
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

FPC number: 4
 PFE: 0
 RED-dropped packets : 19969426 2323178 pps
 RED-dropped bytes : 2196636860 2044397464 bps
 PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
 PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
 PFE: 3
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

FPC number: 6
 PFE: 0
 RED-dropped packets : 19969424 2321205 pps
 RED-dropped bytes : 2196636640 2042660808 bps
 PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
 PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
 PFE: 3
 RED-dropped packets : 0 0 pps

```

```

 RED-dropped bytes : 0 0 bps
PFE: 4
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 5
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 6
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 7
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

FPC number: 7
PFE: 0
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 3
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

Queue: 1, Forwarding classes: expedited-forwarding

FPC number: 1
PFE: 0
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 3
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

FPC number: 4
PFE: 0
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 3
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

FPC number: 6
PFE: 0
 RED-dropped packets : 0 0 pps

```

```

 RED-dropped bytes : 0 0 bps
PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 3
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 4
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 5
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 6
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 7
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

FPC number: 7
PFE: 0
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 3
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

Queue: 2, Forwarding classes: assured-forwarding

FPC number: 1
PFE: 0
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 3
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

FPC number: 4
PFE: 0
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 2
```

```

 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
 PFE: 3
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

FPC number: 6
 PFE: 0
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
 PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
 PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
 PFE: 3
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
 PFE: 4
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
 PFE: 5
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
 PFE: 6
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
 PFE: 7
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

FPC number: 7
 PFE: 0
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
 PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
 PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
 PFE: 3
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

Queue: 3, Forwarding classes: network-control

FPC number: 1
 PFE: 0
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
 PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
 PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
 PFE: 3
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

```

```
FPC number: 4
PFE: 0
 RED-dropped packets : 16338670 1900314 pps
 RED-dropped bytes : 1797253700 1672276976 bps
PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 3
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

FPC number: 6
PFE: 0
 RED-dropped packets : 16338698 1899163 pps
 RED-dropped bytes : 1797256780 1671263512 bps
PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 3
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 4
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 5
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 6
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 7
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

FPC number: 7
PFE: 0
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 3
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
```

#### show interfaces voq (For a Specific Physical Interface) (QFX10000 Switches)

The sample output below shows congestion on ingress PFE 1 on FPC number 0, and on ingress PFE 2 on FPC number 1, as denoted by the count of RED-dropped packets and RED-dropped bytes for best-effort egress queue 0.



```
user@host> show interfaces voq et-1/0/0
```

```
Physical interface: et-1/0/0, Enabled, Physical link is Up
```

```
Interface index: 659, SNMP ifIndex: 539
```

```
Queue: 0, Forwarding classes: best-effort
```

```
FPC number: 0
```

```
PFE: 0
```

```
RED-dropped packets : 0 0 pps
```

```
RED-dropped bytes : 0 0 bps
```

```
PFE: 1
```

```
RED-dropped packets : 411063248 16891870 pps
```

```
RED-dropped bytes : 52616095744 17297275600 bps
```

```
PFE: 2
```

```
RED-dropped packets : 0 0 pps
```

```
RED-dropped bytes : 0 0 bps
```

```
FPC number: 1
```

```
PFE: 0
```

```
RED-dropped packets : 0 0 pps
```

```
RED-dropped bytes : 0 0 bps
```

```
PFE: 1
```

```
RED-dropped packets : 0 0 pps
```

```
RED-dropped bytes : 0 0 bps
```

```
PFE: 2
```

```
RED-dropped packets : 411063012 16891870 pps
```

```
RED-dropped bytes : 52616065536 17297275376 bps
```

```
Queue: 3, Forwarding classes: fcoe
```

```
FPC number: 0
```

```
PFE: 0
```

```
RED-dropped packets : 0 0 pps
```

```
RED-dropped bytes : 0 0 bps
```

```
PFE: 1
```

```
RED-dropped packets : 0 0 pps
```

```
RED-dropped bytes : 0 0 bps
```

```
PFE: 2
```

```
RED-dropped packets : 0 0 pps
```

```
RED-dropped bytes : 0 0 bps
```

```
FPC number: 1
```

```
PFE: 0
```

```
RED-dropped packets : 0 0 pps
```

```
RED-dropped bytes : 0 0 bps
```

```
PFE: 1
```

```
RED-dropped packets : 0 0 pps
```

```
RED-dropped bytes : 0 0 bps
```

```
PFE: 2
```

```
RED-dropped packets : 0 0 pps
```

```
RED-dropped bytes : 0 0 bps
```

```
Queue: 4, Forwarding classes: no-loss
```

```
FPC number: 0
```

```
PFE: 0
```

```
RED-dropped packets : 0 0 pps
```

```
RED-dropped bytes : 0 0 bps
```

```
PFE: 1
```

```
RED-dropped packets : 0 0 pps
```

```
RED-dropped bytes : 0 0 bps
```

```

PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

FPC number: 1
PFE: 0
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

Queue: 7, Forwarding classes: network-control

FPC number: 0
PFE: 0
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

FPC number: 1
PFE: 0
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

```

#### show interfaces voq et-7/0/0 (For a Specific Forwarding Class)

```

user@host> show interfaces voq et-7/0/0 forwarding-class best-effort
Physical interface: et-7/0/0, Enabled, Physical link is Up
Interface index: 155, SNMP ifIndex: 699

```

Queue: 0, Forwarding classes: best-effort

```

FPC number: 1
PFE: 0
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 3
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

```

```

FPC number: 4
PFE: 0
 RED-dropped packets : 66604786 2321519 pps
 RED-dropped bytes : 7326526460 2042936776 bps
PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 3
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

FPC number: 6
PFE: 0
 RED-dropped packets : 66604794 371200 pps
 RED-dropped bytes : 7326527340 326656000 bps
PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 3
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 4
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 5
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 6
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 7
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

FPC number: 7
PFE: 0
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 3
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

```

#### show interfaces voq et-5/0/12 (For a Specific Source FPC)

```

user@host> show interfaces voq et-5/0/12 source-fpc 0
Physical interface: et-5/0/12, Enabled, Physical link is Up
 Interface index: 166, SNMP ifIndex: 1104

Queue: 0, Forwarding classes: best-effort

```

FPC number: 0

PFE: 0

|                       |   |       |
|-----------------------|---|-------|
| RED-dropped packets : | 0 | 0 pps |
| RED-dropped bytes :   | 0 | 0 bps |

PFE: 1

|                       |   |       |
|-----------------------|---|-------|
| RED-dropped packets : | 0 | 0 pps |
| RED-dropped bytes :   | 0 | 0 bps |

PFE: 2

|                       |   |       |
|-----------------------|---|-------|
| RED-dropped packets : | 0 | 0 pps |
| RED-dropped bytes :   | 0 | 0 bps |

PFE: 3

|                       |   |       |
|-----------------------|---|-------|
| RED-dropped packets : | 0 | 0 pps |
| RED-dropped bytes :   | 0 | 0 bps |

Queue: 1, Forwarding classes: expedited-forwarding

FPC number: 0

PFE: 0

|                       |   |       |
|-----------------------|---|-------|
| RED-dropped packets : | 0 | 0 pps |
| RED-dropped bytes :   | 0 | 0 bps |

PFE: 1

|                       |   |       |
|-----------------------|---|-------|
| RED-dropped packets : | 0 | 0 pps |
| RED-dropped bytes :   | 0 | 0 bps |

PFE: 2

|                       |   |       |
|-----------------------|---|-------|
| RED-dropped packets : | 0 | 0 pps |
| RED-dropped bytes :   | 0 | 0 bps |

PFE: 3

|                       |   |       |
|-----------------------|---|-------|
| RED-dropped packets : | 0 | 0 pps |
| RED-dropped bytes :   | 0 | 0 bps |

Queue: 2, Forwarding classes: assured-forwarding

FPC number: 0

PFE: 0

|                       |   |       |
|-----------------------|---|-------|
| RED-dropped packets : | 0 | 0 pps |
| RED-dropped bytes :   | 0 | 0 bps |

PFE: 1

|                       |   |       |
|-----------------------|---|-------|
| RED-dropped packets : | 0 | 0 pps |
| RED-dropped bytes :   | 0 | 0 bps |

PFE: 2

|                       |   |       |
|-----------------------|---|-------|
| RED-dropped packets : | 0 | 0 pps |
| RED-dropped bytes :   | 0 | 0 bps |

PFE: 3

|                       |   |       |
|-----------------------|---|-------|
| RED-dropped packets : | 0 | 0 pps |
| RED-dropped bytes :   | 0 | 0 bps |

Queue: 3, Forwarding classes: network-control

FPC number: 0

PFE: 0

|                       |   |       |
|-----------------------|---|-------|
| RED-dropped packets : | 0 | 0 pps |
| RED-dropped bytes :   | 0 | 0 bps |

PFE: 1

|                       |   |       |
|-----------------------|---|-------|
| RED-dropped packets : | 0 | 0 pps |
| RED-dropped bytes :   | 0 | 0 bps |

PFE: 2

|                       |   |       |
|-----------------------|---|-------|
| RED-dropped packets : | 0 | 0 pps |
| RED-dropped bytes :   | 0 | 0 bps |

PFE: 3

```

RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

#### show interfaces voq et-5/0/12 (For a Specific Forwarding Class and Source FPC)

```

user@host> show interfaces voq et-5/0/12 forwarding-class best-effort source-fpc 5
Physical interface: et-5/0/12, Enabled, Physical link is Up
Interface index: 166, SNMP ifIndex: 1104

```

Queue: 0, Forwarding classes: best-effort

FPC number: 5

PFE: 0

```

RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

PFE: 1

```

RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

PFE: 2

```

RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

PFE: 3

```

RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

PFE: 4

```

RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

PFE: 5

```

RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

PFE: 6

```

RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

PFE: 7

```

RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

#### show interfaces voq et-7/0/0 (Non-Zero)

```

user@host> show interfaces voq et-7/0/0 non-zero

```

```

Physical interface: et-7/0/0, Enabled, Physical link is Up
Interface index: 155, SNMP ifIndex: 699

```

Queue: 0, Forwarding classes: best-effort

FPC number: 4

PFE: 0

```

RED-dropped packets : 95862238 2301586 pps
RED-dropped bytes : 10544846180 2025396264 bps

```

FPC number: 6

PFE: 0

```

RED-dropped packets : 95866639 2322569 pps
RED-dropped bytes : 10545330290 2043860728 bps

```

Queue: 3, Forwarding classes: network-control

FPC number: 4

PFE: 0

```
RED-dropped packets : 78433066 1899727 pps
RED-dropped bytes : 8627637260 1671760384 bps

FPC number: 6
PFE: 0
RED-dropped packets : 78436704 1900628 pps
RED-dropped bytes : 8628037440 1672553432 bps
```

#### show interfaces voq et-7/0/0 (For a Specific Forwarding Class and Non-Zero)

```
user@host show interfaces voq et-7/0/0 forwarding-class best-effort non-zero
```

```
Physical interface: et-7/0/0, Enabled, Physical link is Up
```

```
Interface index: 155, SNMP ifIndex: 699
```

```
Queue: 0, Forwarding classes: best-effort
```

```
FPC number: 4
PFE: 0
RED-dropped packets : 119540012 2322319 pps
RED-dropped bytes : 13149401320 2043640784 bps
```

```
FPC number: 6
PFE: 0
RED-dropped packets : 119540049 2322988 pps
RED-dropped bytes : 13149405390 2044229744 bps
```

## Operational Commands (Classifiers and Rewrite Rules)

- [Monitoring CoS Classifiers on page 7231](#)
- [Monitoring CoS Forwarding Classes on page 7232](#)
- [Monitoring CoS Rewrite Rules on page 7235](#)
- [Monitoring CoS Code-Point Value Aliases on page 7236](#)
- [show class-of-service classifier](#)
- [show class-of-service code-point-aliases](#)
- [show class-of-service forwarding-class](#)
- [show class-of-service forwarding-class-set](#)
- [show class-of-service forwarding-table](#)
- [show class-of-service forwarding-table classifier](#)
- [show class-of-service forwarding-table classifier mapping](#)
- [show class-of-service forwarding-table rewrite-rule](#)
- [show class-of-service forwarding-table rewrite-rule mapping](#)
- [show class-of-service interface](#)
- [show class-of-service multi-destination](#)
- [show class-of-service rewrite-rule](#)

---

### Monitoring CoS Classifiers

**Purpose** Display the mapping of incoming CoS values to forwarding class and loss priority for each classifier.

**Action** To monitor CoS classifiers in the CLI, enter the CLI command:

```
user@switch> show class-of-service classifier
```

To monitor a particular classifier in the CLI, enter the CLI command:

```
user@switch> show class-of-service classifier name classifier-name
```

To monitor a particular type of classifier in the CLI, enter the CLI command:

```
user@switch> show class-of-service classifier type classifier-type
```

**Meaning** Table 616 summarizes key output fields for CoS classifiers.

**Table 616: Summary of Key CoS Classifier Output Fields**

| Field            | Values                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Classifier       | Name of a classifier.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Code point type  | <p>Type of classifier:</p> <ul style="list-style-type: none"> <li>• <b>dscp</b>—All classifiers of the DSCP type.</li> <li>• <b>ieee-802.1</b>—All classifiers of the IEEE 802.1 type.</li> <li>• <b>ieee-mcast</b>—All classifiers of the IEEE 802.1 multicast type.</li> </ul> <p><b>NOTE:</b> QFX10000 switches do not use different classifiers for unicast and multdestination (multicast, broadcast, destination lookup fail) traffic, so multicast-specific classifiers are not supported.</p> <ul style="list-style-type: none"> <li>• <b>exp</b>—All classifiers of the MPLS exp type.</li> </ul> <p><b>NOTE:</b> OCX Series switches do not support MPLS.</p> |
| Index            | Internal index of the classifier.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Code point       | DSCP or IEEE 802.1 code point value of the incoming packets, in bits. These values are used for classification.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Forwarding Class | Name of the forwarding class that the classifier assigns to an incoming packet. This class affects the forwarding and scheduling policies that are applied to the packet as it transits the switch.                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Loss Priority    | Loss priority value that the classifier assigns to the incoming packet based on its code point value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

## Monitoring CoS Forwarding Classes

**Purpose** Use the monitoring functionality to view the current assignment of CoS forwarding classes to queue numbers on the system.

**Action** To monitor CoS forwarding classes in the CLI, enter the following CLI command:

```
user@switch> show class-of-service forwarding-class
```

**Meaning** Some switches use different forwarding classes, output queues, and classifiers for unicast and multdestination (multicast, broadcast, destination lookup fail) traffic. These switches support 12 forwarding classes and output queues, eight for unicast traffic and four for multdestination traffic.



Some switches use the same forwarding classes, output queues, and classifiers for unicast and multdestination traffic. These switches support eight forwarding classes and eight output queues.

[Table 617](#) summarizes key output fields on switches that use different forwarding classes and output queues for unicast and multdestination traffic.

**Table 617: Summary of Key CoS Forwarding Class Output Fields on Switches that Separate Unicast and Multidestination Traffic**

| Field            | Values                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Forwarding Class | <p>Names of forwarding classes assigned to queue numbers. By default, the following unicast forwarding classes are assigned to queues 0, 3, 4, and 7, respectively:</p> <ul style="list-style-type: none"> <li>• <b>best-effort</b>—Provides no special CoS handling of packets. Loss priority is typically not carried in a CoS value.</li> <li>• <b>fcoe</b>—Provides guaranteed delivery for Fibre Channel over Ethernet (FCoE) traffic.</li> <li>• <b>no-loss</b>—Provides guaranteed delivery for TCP lossless traffic</li> <li>• <b>network-control</b>—Packets can be delayed but not dropped.</li> </ul> <p>By default, the following multdestination forwarding class is assigned to queue 8:</p> <ul style="list-style-type: none"> <li>• <b>mcast</b>—Provides no special CoS handling of packets.</li> </ul> |
| Queue            | <p>Queue number corresponding to (mapped to) the forwarding class name.</p> <p>By default, four queues (0, 3, 4, and 7) are assigned to unicast forwarding classes and one queue (8) is assigned to a multdestination forwarding class:</p> <ul style="list-style-type: none"> <li>• Queue 0—<b>best-effort</b></li> <li>• Queue 3—<b>fcoe</b></li> <li>• Queue 4—<b>no-loss</b></li> <li>• Queue 7—<b>network-control</b></li> <li>• Queue 8—<b>mcast</b></li> </ul>                                                                                                                                                                                                                                                                                                                                                    |

**Table 617: Summary of Key CoS Forwarding Class Output Fields on Switches that Separate Unicast and Multidestination Traffic (*continued*)**

| Field   | Values                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| No-Loss | <p>Packet drop attribute associated with each forwarding class:</p> <ul style="list-style-type: none"> <li>Disabled—The forwarding class is configured for lossy transport (packets might drop during periods of congestion)</li> <li>Enabled—The forwarding class is configured for lossless transport</li> </ul> <p><b>NOTE:</b> To achieve lossless transport, you must ensure that priority-based flow control (PFC) and DCBX are properly configured on the lossless priority (IEEE 802.1p code point), and that sufficient port bandwidth is reserved for the lossless traffic flows.</p> <p>OCX Series switches do not support lossless transport.</p> |



**NOTE:** OCX Series switches do not support the default lossless forwarding classes `fcoe` and `no-loss`, and do not support the no-loss packet drop attribute used to configure lossless forwarding classes. On OCX Series switches, do not map traffic to the default `fcoe` and `no-loss` forwarding classes (both of these default forwarding classes carry the no-loss packet drop attribute), and do not configure the no-loss packet drop attribute on forwarding classes.

Table 618 summarizes key output fields on switches that use the same forwarding classes and output queues for unicast and multidestination traffic.

**Table 618: Summary of Key CoS Forwarding Class Output Fields on Switches That Do Not Separate Unicast and Multidestination Traffic**

| Field            | Values                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Forwarding Class | <p>Names of forwarding classes assigned to queue numbers. By default, the following forwarding classes are assigned to queues 0, 3, 4, and 7, respectively:</p> <ul style="list-style-type: none"> <li><b>best-effort</b>—Provides no special CoS handling of packets. Loss priority is typically not carried in a CoS value.</li> <li><b>fcoe</b>—Provides guaranteed delivery for Fibre Channel over Ethernet (FCoE) traffic.</li> <li><b>no-loss</b>—Provides guaranteed delivery for TCP lossless traffic</li> <li><b>network-control</b>—Packets can be delayed but not dropped.</li> </ul> |

**Table 618: Summary of Key CoS Forwarding Class Output Fields on Switches That Do Not Separate Unicast and Multidestination Traffic (*continued*)**

| Field   | Values                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Queue   | <p>Queue number corresponding to (mapped to) the forwarding class name.</p> <p>By default, four queues (0, 3, 4, and 7) are assigned to forwarding classes:</p> <ul style="list-style-type: none"> <li>• Queue 0—<b>best-effort</b></li> <li>• Queue 3—<b>fcoe</b></li> <li>• Queue 4—<b>no-loss</b></li> <li>• Queue 7—<b>network-control</b></li> </ul>                                                                                                                                                                                                                                                                                                           |
| No-Loss | <p>Packet drop attribute associated with each forwarding class:</p> <ul style="list-style-type: none"> <li>• Disabled—The forwarding class is configured for lossy transport (packets might drop during periods of congestion).</li> <li>• Enabled—The forwarding class is configured for lossless transport.</li> </ul> <p><b>NOTE:</b> To achieve lossless transport, you must ensure that priority-based flow control (PFC) and DCBX are properly configured on the lossless priority (IEEE 802.1p code point), and that sufficient port bandwidth is reserved for the lossless traffic flows.</p> <p>OCX Series switches do not support lossless transport.</p> |

## Monitoring CoS Rewrite Rules

**Purpose** Use the monitoring functionality to display information about CoS value rewrite rules, which are based on the forwarding class and loss priority.

**Action** To monitor CoS rewrite rules in the CLI, enter the CLI command:

```
user@switch> show class-of-service rewrite-rule
```

To monitor a particular rewrite rule in the CLI, enter the CLI command:

```
user@switch> show class-of-service rewrite-rule name rewrite-rule-name
```

To monitor a particular type of rewrite rule (for example, DSCP, DSCP IPv6, IEEE-802.1, or MPLS EXP) in the CLI, enter the CLI command:

```
user@switch> show class-of-service rewrite-rule type rewrite-rule-type
```

**Meaning** [Table 619](#) summarizes key output fields for CoS rewrite rules.

**Table 619: Summary of Key CoS Rewrite Rule Output Fields**

| Field            | Values                                                                                                                                                                                                                                                                                                                      |
|------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Rewrite rule     | Name of the rewrite rule.                                                                                                                                                                                                                                                                                                   |
| Code point type  | Rewrite rule type: <ul style="list-style-type: none"> <li>• <b>dscp</b>—For IPv4 DiffServ traffic.</li> <li>• <b>dscp-ipv6</b>—For IPv6 Diffserv traffic.</li> <li>• <b>ieee-802.1</b>—For Layer 2 traffic.</li> <li>• <b>exp</b>—For MPLS traffic.</li> </ul> <p><b>NOTE:</b> OCX Series switches do not support MPLS.</p> |
| Index            | Internal index for the rewrite rule.                                                                                                                                                                                                                                                                                        |
| Forwarding class | Name of the forwarding class that is used to determine CoS values for rewriting in combination with loss priority.<br><br>Rewrite rules are applied to CoS values in outgoing packets based on forwarding class and loss priority setting.                                                                                  |
| Loss priority    | Level of loss priority that is used to determine CoS values for rewriting in combination with forwarding class.                                                                                                                                                                                                             |
| Code point       | Rewrite code point value.                                                                                                                                                                                                                                                                                                   |

**Related Documentation** • [Defining CoS Rewrite Rules on page 6704](#)

## Monitoring CoS Code-Point Value Aliases

**Purpose** Use the monitoring functionality to display information about the CoS code-point value aliases that the system is currently using to represent DSCP and IEEE 802.1p code point bits.

**Action** To monitor CoS value aliases in the CLI, enter the CLI command:

```
user@switch> show class-of-service code-point-aliases
```

To monitor a specific type of code-point alias (DSCP, DSCP IPv6, IEEE 802.1, or MPLS EXP) in the CLI, enter the CLI command:

```
user@switch> show class-of-service code-point-aliases ieee-802.1
```

**Meaning** [Table 620](#) summarizes key output fields for CoS value aliases.

Table 620: Summary of Key CoS Value Alias Output Fields

| Field           | Values                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Code point type | Type of the CoS value: <ul style="list-style-type: none"><li>• <b>dscp</b>—Examines Layer 3 packet headers for IP packet classification.</li><li>• <b>dscp-ipv6</b>—Examines Layer 3 packet headers for IPv6 packet classification.</li><li>• <b>ieee-802.1</b>—Examines Layer 2 packet headers for packet classification.</li><li>• <b>exp</b>—Examines MPLS packet headers for packet classification.</li></ul> <p><b>NOTE:</b> OCX Series switches do not support MPLS.</p> |
| Alias           | Name given to a set of bits—for example, <b>af11</b> is a name for bits <b>001010</b> .                                                                                                                                                                                                                                                                                                                                                                                        |
| Bit pattern     | Set of bits associated with the alias.                                                                                                                                                                                                                                                                                                                                                                                                                                         |

**Related Documentation** • [Defining CoS Code-Point Aliases on page 6687](#)

## show class-of-service classifier

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show class-of-service classifier &lt;name <i>name</i>&gt; &lt;type dscp   type dscp-ipv6   type exp   type ieee-802.1   type inet-precedence&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | For each class-of-service (CoS) classifier, display the mapping of code point value to forwarding class and loss priority.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <p><b>none</b>—Display all classifiers.</p> <p><b>name <i>name</i></b>—(Optional) Display named classifier.</p> <p><b>type dscp</b>—(Optional) Display all classifiers of the Differentiated Services code point (DSCP) type.</p> <p><b>type dscp-ipv6</b>—(Optional) Display all classifiers of the DSCP for IPv6 type.</p> <p><b>type exp</b>—(Optional) Display all classifiers of the MPLS experimental (EXP) type.</p> <p><b>type ieee-802.1</b>—(Optional) Display all classifiers of the ieee-802.1 type.</p> <p><b>type inet-precedence</b>—(Optional) Display all classifiers of the inet-precedence type.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>List of Sample Output</b>    | <p><a href="#">show class-of-service classifier type ieee-802.1 on page 7239</a></p> <p><a href="#">show class-of-service classifier type ieee-802.1 (QFX Series) on page 7239</a></p>                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Output Fields</b>            | <p><a href="#">Table 621</a> describes the output fields for the <b>show class-of-service classifier</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                                                                                                                                           |

**Table 621: show class-of-service classifier Output Fields**

| Field Name      | Field Description                                                                                                                                                       |
|-----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Classifier      | Name of the classifier.                                                                                                                                                 |
| Code point type | Type of the classifier: <b>exp</b> (not on EX Series switch), <b>dscp</b> , <b>dscp-ipv6</b> (not on EX Series switch), <b>ieee-802.1</b> , or <b>inet-precedence</b> . |
| Index           | Internal index of the classifier.                                                                                                                                       |
| Code point      | Code point value used for classification                                                                                                                                |

Table 621: show class-of-service classifier Output Fields (*continued*)

| Field Name              | Field Description                                                                                                                                                                                                 |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Forwarding class</b> | Classification of a packet affecting the forwarding, scheduling, and marking policies applied as the packet transits the router.                                                                                  |
| <b>Loss priority</b>    | Loss priority value used for classification. For most platforms, the value is <b>high</b> or <b>low</b> . For some platforms, the value is <b>high</b> , <b>medium-high</b> , <b>medium-low</b> , or <b>low</b> . |

## Sample Output

### show class-of-service classifier type ieee-802.1

```

user@host> show class-of-service classifier type ieee-802.1
Classifier: ieee802.1-default, Code point type: ieee-802.1, Index: 3
Code Point Forwarding Class Loss priority
000 best-effort low
001 best-effort high
010 expedited-forwarding low
011 expedited-forwarding high
100 assured-forwarding low
101 assured-forwarding medium-high
110 network-control low
111 network-control high

Classifier: users-ieee802.1, Code point type: ieee-802.1
Code point Forwarding class Loss priority
100 expedited-forwarding low

```

### show class-of-service classifier type ieee-802.1 (QFX Series)

```

user@switch> show class-of-service classifier type ieee-802.1
Classifier: ieee8021p-default, Code point type: ieee-802.1, Index: 11
Code point Forwarding class Loss priority
000 best-effort low
001 best-effort low
010 best-effort low
011 fcoe low
100 no-loss low
101 best-effort low
110 network-control low
111 network-control low

Classifier: ieee8021p-untrust, Code point type: ieee-802.1, Index: 16
Code point Forwarding class Loss priority
000 best-effort low
001 best-effort low
010 best-effort low
011 best-effort low
100 best-effort low
101 best-effort low
110 best-effort low
111 best-effort low

Classifier: ieee-mcast, Code point type: ieee-802.1, Index: 46
Code point Forwarding class Loss priority
000 mcast low

```

|     |       |     |
|-----|-------|-----|
| 001 | mcast | low |
| 010 | mcast | low |
| 011 | mcast | low |
| 100 | mcast | low |
| 101 | mcast | low |
| 110 | mcast | low |
| 111 | mcast | low |



## show class-of-service code-point-aliases

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show class-of-service code-point-aliases</code><br><code>&lt;dscp   dscp-ipv6   exp   ieee-802.1   inet-precedence&gt;</code>                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.<br>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                              |
| <b>Description</b>              | Display the mapping of class-of-service (CoS) code point aliases to corresponding bit patterns.                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><b>none</b>—Display code point aliases of all code point types.</p> <p><b>dscp</b>—(Optional) Display Differentiated Services code point (DSCP) aliases.</p> <p><b>dscp-ipv6</b>—(Optional) Display IPv6 DSCP aliases.</p> <p><b>exp</b>—(Optional) Display MPLS EXP code point aliases.</p> <p><b>ieee-802.1</b>—(Optional) Display IEEE-802.1 code point aliases.</p> <p><b>inet-precedence</b>—(Optional) Display IPv4 precedence code point aliases.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>List of Sample Output</b>    | <a href="#">show class-of-service code-point-aliases exp on page 7242</a>                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Output Fields</b>            | <a href="#">Table 622</a> describes the output fields for the <b>show class-of-service code-point-aliases</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                  |

**Table 622: show class-of-service code-point-aliases Output Fields**

| Field Name             | Field Description                                                                                                                                                                                                           |
|------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Code point type</b> | Type of the code points displayed: <b>dscp</b> , <b>dscp-ipv6</b> (not on EX Series switch), <b>exp</b> (not on EX Series switch or the QFX Series), <b>ieee-802.1</b> , or <b>inet-precedence</b> (not on the QFX Series). |
| <b>Alias</b>           | Alias for a bit pattern.                                                                                                                                                                                                    |
| <b>Bit pattern</b>     | Bit pattern for which the alias is displayed.                                                                                                                                                                               |

## Sample Output

`show class-of-service code-point-aliases exp`

```
user@host> show class-of-service code-point-aliases exp
Code point type: exp
Alias Bit pattern
af11 100
af12 101
be 000
be1 001
cs6 110
cs7 111
ef 010
ef1 011
nc1 110
nc2 111
```

## show class-of-service forwarding-class

|                                 |                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show class-of-service forwarding-class                                                                                                                                                                                                         |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.<br>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                |
| <b>Description</b>              | Display information about forwarding classes, including the mapping of forwarding classes to queue numbers.                                                                                                                                    |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Monitoring CoS Forwarding Classes</i></li> <li>• <a href="#">Monitoring CoS Forwarding Classes on page 7232</a></li> </ul>                                                                         |
| <b>List of Sample Output</b>    | <a href="#">show class-of-service forwarding-class on page 7244</a><br><a href="#">show class-of-service forwarding-class (EX8200 Switch) on page 7244</a><br><a href="#">show class-of-service forwarding-class (QFX Series) on page 7244</a> |
| <b>Output Fields</b>            | <a href="#">Table 623</a> describes the output fields for the <b>show class-of-service forwarding-class</b> command. Output fields are listed in the approximate order in which they appear.                                                   |

**Table 623: show class-of-service forwarding-class Output Fields**

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Forwarding class</b>  | Name of the forwarding class.                                                                                                                                                                                                                                                                                                                                                                                |
| <b>ID</b>                | Forwarding class identifier.                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Queue</b>             | CoS output queue mapped to the forwarding class.                                                                                                                                                                                                                                                                                                                                                             |
| <b>Policing priority</b> | Not supported on EX Series switches or the QFX Series and can be ignored.                                                                                                                                                                                                                                                                                                                                    |
| <b>Fabric priority</b>   | (EX8200 switches only) Fabric priority for the forwarding class, either <b>high</b> or <b>low</b> . Determines the priority of packets entering the switch fabric.                                                                                                                                                                                                                                           |
| <b>No-Loss</b>           | (QFX Series only) Packet loss attribute to differentiate lossless forwarding classes from lossy forwarding classes: <ul style="list-style-type: none"> <li>• Disabled—Lossless transport is not configured on the forwarding class (packet drop attribute is <b>drop</b>).</li> <li>• Enabled—Lossless transport is configured on the forwarding class (packet drop attribute is <b>no-loss</b>).</li> </ul> |

## Sample Output

### show class-of-service forwarding-class

```

user@switch> show class-of-service forwarding-class
Forwarding class ID Queue Policing priority
best-effort 0 0 normal
expedited-forwarding 1 5 normal
assured-forwarding 2 1 normal
network-control 3 7 normal

```

## Sample Output

### show class-of-service forwarding-class (EX8200 Switch)

```

user@switch> show class-of-service forwarding-class
Forwarding class ID Queue Fabric priority
best-effort 0 0 low
expedited-forwarding 1 5 low
assured-forwarding 2 1 low
network-control 3 7 low
mcast-be 4 2 low
mcast-ef 5 4 low
mcast-af 6 6 low

```

## Sample Output

### show class-of-service forwarding-class (QFX Series)

```

user@switch> show class-of-service forwarding-class
Forwarding class ID Queue Policing priority No-Loss
best-effort 0 0 normal Disabled
fcoe 1 3 normal Enabled
no-loss 2 4 normal Enabled
network-control 3 7 normal Disabled
mcast 8 8 normal Disabled

```

On switches that do not use different forwarding classes and output queues for unicast and multdestination (multicast, broadcast, destination lookup fail) traffic, there is no **mcast** forwarding class and there is no queue 8. (Switches that use different forwarding classes and output queues for unicast and multdestination traffic support 12 forwarding classes and output queues, of which four of each are dedicated to multdestination traffic. Switches that use the same forwarding classes and output queues for unicast and multdestination traffic support eight forwarding classes and eight output queues.)

## show class-of-service forwarding-class-set

|                                 |                                                                                                                                                                                                                                                                                                          |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show class-of-service forwarding-class-set</code><br><code>&lt;forwarding-class-set-name&gt;</code>                                                                                                                                                                                                |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.<br>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                |
| <b>Description</b>              | Display the forwarding classes associated with each forwarding class set.                                                                                                                                                                                                                                |
| <b>Options</b>                  | <b>none</b> —Display all forwarding class sets.<br><br><b>forwarding-class-set-name</b> —(Optional) Display the forwarding classes associated with the specified forwarding class set.                                                                                                                   |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding CoS Forwarding Class Sets (Priority Groups) on page 6695</a></li> <li>• <a href="#">Defining CoS Forwarding Class Sets on page 6696</a></li> <li>• <a href="#">Example: Configuring Forwarding Class Sets on page 6697</a></li> </ul> |
| <b>Output Fields</b>            | <a href="#">Table 624</a> describes the output fields for the <b>show class-of-service forwarding-class-set</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                         |

**Table 624: show class-of-service forwarding-class-set Output Fields**

| Field Name                        | Field Description                   |
|-----------------------------------|-------------------------------------|
| <b>Forwarding class set</b>       | Name of the forwarding class set.   |
| <b>Type</b>                       | Internal Junos OS type.             |
| <b>Forwarding class set index</b> | Index of this forwarding class set. |
| <b>Forwarding class</b>           | Name of a forwarding class.         |
| <b>Index</b>                      | Index of this forwarding class.     |

## Sample Output

### show class-of-service forwarding-class-set

```

user@switch> show class-of-service forwarding-class-set
Forwarding class set: san_fcset, Type: normal-type, Forwarding class set index:
37839
 Forwarding class Index
 fcoe 1

Forwarding class set: lan_fcset, Type: normal-type, Forwarding class set index:

```

37840

| Forwarding class | Index |
|------------------|-------|
| best-effort      | 0     |

Forwarding class set: multicast\_fcset, Type: normal-type, Forwarding class set index: 37841

| Forwarding class | Index |
|------------------|-------|
| mcast            | 8     |

## show class-of-service forwarding-table

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-----------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 7247</a><br><a href="#">Syntax (TX Matrix and TX Matrix Plus Router) on page 7247</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Syntax</b>                                       | show class-of-service forwarding-table                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Syntax (TX Matrix and TX Matrix Plus Router)</b> | show class-of-service forwarding-table<br><lcc number>   <sfc number>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>                          | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>                                  | <p>Display the entire class-of-service (CoS) configuration as it exists in the forwarding table. Executing this command is equivalent to executing all <b>show class-of-service forwarding-table</b> commands in succession.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                                      | <p><b>lcc number</b>—(TX Matrix and TX Matrix Plus router only) (Optional) On a TX Matrix router, display the forwarding table configuration for a specific T640 router (or line-card chassis) configured in a routing matrix. On a TX Matrix Plus router, display the forwarding table configuration for a specific router (or line-card chassis) configured in the routing matrix.</p> <p>Replace <i>number</i> with the following values depending on the LCC configuration:</p> <ul style="list-style-type: none"> <li>• 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.</li> <li>• 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.</li> <li>• 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.</li> <li>• 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.</li> </ul> <p><b>sfc number</b>—(TX Matrix Plus routers only) (Optional) Display the forwarding table configuration for the TX Matrix Plus router. Replace <i>number</i> with 0.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>List of Sample Output</b>                        | <a href="#">show class-of-service forwarding-table on page 7248</a><br><a href="#">show class-of-service forwarding-table lcc (TX Matrix Plus Router) on page 7249</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Output Fields</b>                                | <p>See the output field descriptions for <b>show class-of-service forwarding-table</b> commands:</p> <ul style="list-style-type: none"> <li>• <a href="#">show class-of-service forwarding-table classifier</a></li> <li>• <a href="#">show class-of-service forwarding-table classifier mapping</a></li> <li>• <a href="#">show class-of-service forwarding-table drop-profile</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

- *show class-of-service forwarding-table fabric scheduler-map*
- *show class-of-service forwarding-table rewrite-rule*
- *show class-of-service forwarding-table rewrite-rule mapping*
- *show class-of-service forwarding-table scheduler-map*

## Sample Output

### show class-of-service forwarding-table

```

user@host> show class-of-service forwarding-table
Classifier table index: 9, # entries: 8, Table type: EXP
Entry # Code point Forwarding-class # PLP
 0 000 0 0
 1 001 0 1
 2 010 1 0
 3 011 1 1
 4 100 2 0
 5 101 2 1
 6 110 3 0
 7 111 3 1

Interface Index Table Index/ Q num Table type
sp-0/0/0.1001 66 11 11 IPv4 precedence
sp-0/0/0.2001 67 11 11 IPv4 precedence
sp-0/0/0.16383 68 11 11 IPv4 precedence
fe-0/0/0.0 69 11 11 IPv4 precedence

Interface: sp-0/0/0 (Index: 129, Map index: 2, Map type: FINAL,
Num of queues: 2):
 Entry 0 (Scheduler index: 16, Forwarding-class #: 0):
 Tx rate: 0 Kb (95%), Buffer size: 95 percent
 Priority low
 PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1
 Entry 1 (Scheduler index: 18, Forwarding-class #: 3):
 Tx rate: 0 Kb (5%), Buffer size: 5 percent
 Priority low
 PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1

Interface: fe-0/0/0 (Index: 137, Map index: 2, Map type: FINAL,
Num of queues: 2):
 Entry 0 (Scheduler index: 16, Forwarding-class #: 0):
 Tx rate: 0 Kb (95%), Buffer size: 95 percent
 Priority low
 PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1
 Entry 1 (Scheduler index: 18, Forwarding-class #: 3):
 Tx rate: 0 Kb (5%), Buffer size: 5 percent
 Priority low
 PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1

Interface: fe-0/0/1 (Index: 138, Map index: 2, Map type: FINAL,
Num of queues: 2):
 Entry 0 (Scheduler index: 16, Forwarding-class #: 0):
 Tx rate: 0 Kb (95%), Buffer size: 95 percent
 Priority low
 PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1
 Entry 1 (Scheduler index: 18, Forwarding-class #: 3):
 Tx rate: 0 Kb (5%), Buffer size: 5 percent
 Priority low

```



PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1

...

RED drop profile index: 1, # entries: 1

| Entry | Fullness(%) | Drop<br>Probability(%) |
|-------|-------------|------------------------|
| 0     | 100         | 100                    |

### show class-of-service forwarding-table lcc (TX Matrix Plus Router)

user@host> show class-of-service forwarding-table lcc 0  
lcc0-re0:

-----

Classifier table index: 9, # entries: 64, Table type: IPv6 DSCP

| Entry # | Code point | Forwarding-class # | PLP |
|---------|------------|--------------------|-----|
| 0       | 000000     | 0                  | 0   |
| 1       | 000001     | 0                  | 0   |
| 2       | 000010     | 0                  | 0   |
| 3       | 000011     | 0                  | 0   |
| 4       | 000100     | 0                  | 0   |
| 5       | 000101     | 0                  | 0   |
| 6       | 000110     | 0                  | 0   |
| 7       | 000111     | 0                  | 0   |
| 8       | 001000     | 0                  | 0   |
| 9       | 001001     | 0                  | 0   |
| 10      | 001010     | 0                  | 0   |
| 11      | 001011     | 0                  | 0   |
| 12      | 001100     | 0                  | 0   |
| 13      | 001101     | 0                  | 0   |
| 14      | 001110     | 0                  | 0   |
| 15      | 001111     | 0                  | 0   |
| 16      | 010000     | 0                  | 0   |
| 17      | 010001     | 0                  | 0   |
| 18      | 010010     | 0                  | 0   |
| 19      | 010011     | 0                  | 0   |
| 20      | 010100     | 0                  | 0   |
| 21      | 010101     | 0                  | 0   |
| 22      | 010110     | 0                  | 0   |
| 23      | 010111     | 0                  | 0   |
| 24      | 011000     | 0                  | 0   |
| 25      | 011001     | 0                  | 0   |
| 26      | 011010     | 0                  | 0   |
| 27      | 011011     | 0                  | 0   |
| 28      | 011100     | 0                  | 0   |
| 29      | 011101     | 0                  | 0   |
| 30      | 011110     | 0                  | 0   |
| 31      | 011111     | 0                  | 0   |
| 32      | 100000     | 0                  | 0   |
| 33      | 100001     | 0                  | 0   |
| 34      | 100010     | 0                  | 0   |
| 35      | 100011     | 0                  | 0   |
| 36      | 100100     | 0                  | 0   |
| 37      | 100101     | 0                  | 0   |
| 38      | 100110     | 0                  | 0   |
| 39      | 100111     | 0                  | 0   |
| 40      | 101000     | 0                  | 0   |
| 41      | 101001     | 0                  | 0   |
| 42      | 101010     | 0                  | 0   |
| 43      | 101011     | 0                  | 0   |

|     |        |   |   |
|-----|--------|---|---|
| 44  | 101100 | 0 | 0 |
| 45  | 101101 | 0 | 0 |
| 46  | 101110 | 0 | 0 |
| ... |        |   |   |

## show class-of-service forwarding-table classifier

|                                 |                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show class-of-service forwarding-table classifier                                                                                                                                                       |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.<br>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.            |
| <b>Description</b>              | Display the mapping of code point value to queue number and loss priority for each classifier as it exists in the forwarding table.                                                                     |
| <b>Options</b>                  | This command has no options.                                                                                                                                                                            |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                    |
| <b>List of Sample Output</b>    | <a href="#">show class-of-service forwarding-table classifier on page 7251</a>                                                                                                                          |
| <b>Output Fields</b>            | <a href="#">Table 625</a> describes the output fields for the <b>show class-of-service forwarding-table classifier</b> command. Output fields are listed in the approximate order in which they appear. |

**Table 625: show class-of-service forwarding-table classifier Output Fields**

| Field Name                    | Field Description                                                                                                                                                                                                                                                                                                                                                                |
|-------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Classifier table index</b> | Index of the classifier table.                                                                                                                                                                                                                                                                                                                                                   |
| <b>entries</b>                | Total number of entries.                                                                                                                                                                                                                                                                                                                                                         |
| <b>Table type</b>             | Type of code points in the table: <b>DSCP</b> , <b>EXP</b> (not on the QFX Series), <b>IEEE 802.1</b> , <b>IPv4 precedence</b> (not on the QFX Series), or <b>IPv6 DSCP</b> .                                                                                                                                                                                                    |
| <b>Entry #</b>                | Entry number.                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Code point</b>             | Code point value used for classification.                                                                                                                                                                                                                                                                                                                                        |
| <b>Forwarding-class #</b>     | Forwarding class to which the code point is assigned.                                                                                                                                                                                                                                                                                                                            |
| <b>PLP</b>                    | Packet loss priority value set by classification. For most platforms, the value can be <b>0</b> or <b>1</b> . For some platforms, the value is <b>0</b> , <b>1</b> , <b>2</b> , or <b>3</b> . The value <b>0</b> represents low PLP. The value <b>1</b> represents <b>high</b> PLP. The value <b>2</b> represents medium-low PLP. The value <b>3</b> represents medium-high PLP. |

## Sample Output

### show class-of-service forwarding-table classifier

```

user@host> show class-of-service forwarding-table classifier
Classifier table index: 62436, # entries: 64, Table type: DSCP

Entry # Code point Forwarding-class # PLP

```

|     |        |   |   |
|-----|--------|---|---|
| 0   | 000000 | 0 | 0 |
| 1   | 000001 | 0 | 0 |
| 2   | 000010 | 0 | 0 |
| 3   | 000011 | 0 | 0 |
| 4   | 000100 | 0 | 0 |
| 5   | 000101 | 0 | 0 |
| 6   | 000110 | 0 | 0 |
| 7   | 000111 | 0 | 0 |
| 8   | 001000 | 0 | 0 |
| 9   | 001001 | 0 | 0 |
| 10  | 001010 | 1 | 1 |
| 11  | 001011 | 0 | 0 |
| ... |        |   |   |
| 60  | 111100 | 0 | 0 |
| 61  | 111101 | 0 | 0 |
| 62  | 111110 | 0 | 0 |
| 63  | 111111 | 0 | 0 |

## show class-of-service forwarding-table classifier mapping

|                                 |                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show class-of-service forwarding-table classifier mapping                                                                                                                                                       |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.<br>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                    |
| <b>Description</b>              | For each logical interface, display either the table index of the classifier for a given code point type or the queue number (if it is a fixed classification) in the forwarding table.                         |
| <b>Options</b>                  | This command has no options.                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                            |
| <b>List of Sample Output</b>    | <a href="#">show class-of-service forwarding-table classifier mapping on page 7253</a>                                                                                                                          |
| <b>Output Fields</b>            | <a href="#">Table 626</a> describes the output fields for the <b>show class-of-service forwarding-table classifier mapping</b> command. Output fields are listed in the approximate order in which they appear. |

**Table 626: show class-of-service forwarding-table classifier mapping Output Fields**

| Field Name                | Field Description                                                                                                                                                                            |
|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Table index/ Q num</b> | If the table type is <b>Fixed</b> , the number of the queue to which the interface is mapped. For all other types, this value is the classifier index number.                                |
| <b>Interface</b>          | Name of the logical interface. This field can also show the physical interface (QFX Series).                                                                                                 |
| <b>Index</b>              | Logical interface index.                                                                                                                                                                     |
| <b>Table type</b>         | Type of code points in the table: <b>DSCP</b> , <b>EXP</b> (not on the QFX Series), <b>Fixed</b> , <b>IEEE 802.1</b> , <b>IPv4 precedence</b> (not on the QFX Series), or <b>IPv6 DSCP</b> . |

## Sample Output

### show class-of-service forwarding-table classifier mapping

```

user@host> show class-of-service forwarding-table classifier mapping
Table index/
Interface Index Q num Table type
so-5/0/0.0 10 62436 DSCP
so-0/1/0.0 11 62436 DSCP
so-0/2/0.0 12 1 Fixed
so-0/2/1.0 13 62436 DSCP
so-0/2/1.0 13 62437 IEEE 802.1

```

|            |    |       |                 |
|------------|----|-------|-----------------|
| so-0/2/2.0 | 14 | 62436 | DSCP            |
| so-0/2/2.0 | 14 | 62438 | IPv4 precedence |

## show class-of-service forwarding-table rewrite-rule

|                                 |                                                                                                                                                                                                           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show class-of-service forwarding-table rewrite-rule                                                                                                                                                       |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.<br>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.              |
| <b>Description</b>              | Display mapping of queue number and loss priority to code point value for each rewrite rule as it exists in the forwarding table.                                                                         |
| <b>Options</b>                  | This command has no options.                                                                                                                                                                              |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                      |
| <b>List of Sample Output</b>    | <a href="#">show class-of-service forwarding-table rewrite-rule on page 7255</a>                                                                                                                          |
| <b>Output Fields</b>            | <a href="#">Table 627</a> describes the output fields for the <b>show class-of-service forwarding-table rewrite-rule</b> command. Output fields are listed in the approximate order in which they appear. |

**Table 627: show class-of-service forwarding-table rewrite-rule Output Fields**

| Field Name                 | Field Description                                                                                                                                                                                           |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Rewrite table index</b> | Index for this rewrite rule.                                                                                                                                                                                |
| <b># entries</b>           | Number of entries in this rewrite rule.                                                                                                                                                                     |
| <b>Table type</b>          | Type of table: <b>DSCP</b> , <b>EXP</b> (not on the QFX Series), <b>EXP-PUSH-3</b> (not on the QFX Series), <b>IEEE 802.1,IPv4 precedence</b> (not on the QFX Series), <b>IPv6 DSCP</b> , or <b>Fixed</b> . |
| <b>Q#</b>                  | Queue number to which this entry is assigned.                                                                                                                                                               |
| <b>Low bits</b>            | Code point value for low-priority loss profile.                                                                                                                                                             |
| <b>State</b>               | State of this code point: <b>enabled</b> , <b>rewritten</b> , or <b>disabled</b> .                                                                                                                          |
| <b>High bits</b>           | Code point value for high-priority loss profile.                                                                                                                                                            |

## Sample Output

### show class-of-service forwarding-table rewrite-rule

```

user@host> show class-of-service forwarding-table rewrite-rule
Rewrite table index: 3753, # entries: 4, Table type: DSCP
Q# Low bits State High bits State
0 000111 Enabled 001010 Enabled
2 000000 Disabled 001100 Enabled

```

|   |        |         |        |         |
|---|--------|---------|--------|---------|
| 1 | 101110 | Enabled | 110111 | Enabled |
| 3 | 110000 | Enabled | 111000 | Enabled |



## show class-of-service forwarding-table rewrite-rule mapping

|                                 |                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show class-of-service forwarding-table rewrite-rule mapping                                                                                                                                                       |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.<br>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                      |
| <b>Description</b>              | For each logical interface, display the table identifier of the rewrite rule map for each code point type.                                                                                                        |
| <b>Options</b>                  | This command has no options.                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                              |
| <b>List of Sample Output</b>    | <a href="#">show class-of-service forwarding-table rewrite-rule mapping on page 7257</a>                                                                                                                          |
| <b>Output Fields</b>            | <a href="#">Table 628</a> describes the output fields for the <b>show class-of-service forwarding-table rewrite-rule mapping</b> command. Output fields are listed in the approximate order in which they appear. |

**Table 628: show class-of-service forwarding-table rewrite-rule mapping Output Fields**

| Field Name         | Field Description                                                                                                                                                                                                                                                         |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Interface</b>   | Name of the logical interface. This field can also show the physical interface (QFX Series).                                                                                                                                                                              |
| <b>Index</b>       | Logical interface index.                                                                                                                                                                                                                                                  |
| <b>Table index</b> | Rewrite table index.                                                                                                                                                                                                                                                      |
| <b>Type</b>        | Type of classifier: <b>DSCP</b> , <b>EXP</b> (not on the QFX Series), <b>EXP-PUSH-3</b> (not on the QFX Series), <b>EXP-SWAP-PUSH-2</b> (not on the QFX Series), <b>IEEE 802.1</b> , <b>IPv4 precedence</b> (not on the QFX Series), <b>IPv6 DSCP</b> , or <b>Fixed</b> . |

## Sample Output

### show class-of-service forwarding-table rewrite-rule mapping

```

user@host> show class-of-service forwarding-table rewrite-rule mapping
Interface Index Table index Type
so-5/0/0.0 10 3753 DSCP
so-0/1/0.0 11 3753 DSCP
so-0/2/0.0 12 3753 DSCP
so-0/2/1.0 13 3753 DSCP
so-0/2/2.0 14 3753 DSCP
so-0/2/3.0 15 3753 DSCP

```

## show class-of-service interface

---

**Syntax**    `show class-of-service interface`  
              `<comprehensive | detail> <interface-name>`

**Release Information**    Command introduced before Junos OS Release 7.4.  
                              Command introduced in Junos OS Release 9.0 for EX Series switches.  
                              Forwarding class map information added in Junos OS Release 9.4.  
                              Command introduced in Junos OS Release 11.1 for the QFX Series.  
                              Command introduced in Junos OS Release 12.1 for the PTX Series Packet Transport Switches.  
                              Command introduced in Junos OS Release 12.2 for the ACX Series Universal Access routers.  
                              Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.  
                              Options **detail** and **comprehensive** introduced in Junos OS Release 11.4.  
                              Command introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.

**Description**    Display the logical and physical interface associations for the classifier, rewrite rules, and scheduler map objects.



**NOTE:** On routing platforms with dual Routing Engines, running this command on the backup Routing Engine, with or without any of the available options, is not supported and produces the following error message:

**error: the class-of-service subsystem is not running**

**Options**    **none**—Display CoS associations for all physical and logical interfaces.

**comprehensive**—(M Series, MX Series, and T Series routers) (Optional) Display comprehensive quality-of-service (QoS) information about all physical and logical interfaces.

**detail**—(M Series, MX Series, and T Series routers) (Optional) Display QoS and CoS information based on the interface.

If the **interface** *interface-name* is a physical interface, the output includes:

- Brief QoS information about the physical interface
- Brief QoS information about the logical interface
- CoS information about the physical interface
- Brief information about filters or policers of the logical interface
- Brief CoS information about the logical interface

If the **interface** *interface-name* is a logical interface, the output includes:

- Brief QoS information about the logical interface

- Information about filters or policers for the logical interface
- CoS information about the logical interface

**interface-name**—(Optional) Display class-of-service (CoS) associations for the specified interface.

**none**—Display CoS associations for all physical and logical interfaces.

**Required Privilege Level** view

**Related Documentation**

- *Verifying and Managing Junos OS Enhanced Subscriber Management*

**List of Sample Output** [show class-of-service interface \(Physical\) on page 7270](#)  
[show class-of-service interface \(Logical\) on page 7271](#)  
[show class-of-service interface \(Gigabit Ethernet\) on page 7271](#)  
[show class-of-service interface \(ANCP\) on page 7271](#)  
[show class-of-service interface \(PPPoE Interface\) on page 7271](#)  
[show class-of-service interface \(T4000 Routers with Type 5 FPCs\) on page 7271](#)  
[show class-of-service interface detail on page 7272](#)  
[show class-of-service interface comprehensive on page 7272](#)  
[show class-of-service interface \(ACX Series Routers\) on page 7283](#)  
[show class-of-service interface \(PPPoE Subscriber Interface for Enhanced Subscriber Management\) on page 7286](#)

**Output Fields** [Table 609](#) describes the output fields for the **show class-of-service interface** command. Output fields are listed in the approximate order in which they appear.

**Table 629: show class-of-service interface Output Fields**

| Field Name         | Field Description                                                                                                                                                                                                                                                                                          |                                                      |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------|
| Physical interface | Name of a physical interface.                                                                                                                                                                                                                                                                              |                                                      |
| Index              | Index of this interface or the internal index of this object.<br><br>(Enhanced subscriber management for MX Series routers) Index values for dynamic CoS traffic control profiles and dynamic scheduler maps are larger for enhanced subscriber management than they are for legacy subscriber management. |                                                      |
| Dedicated Queues   | Status of dedicated queues configured on an interface. Supported only on Trio MPC/MIC interfaces on MX Series routers.                                                                                                                                                                                     | Number of queues you can configure on the interface. |
| Queues supported   | Number of queues you can configure on the interface.                                                                                                                                                                                                                                                       |                                                      |
| Queues in use      | Number of queues currently configured.                                                                                                                                                                                                                                                                     |                                                      |

Table 629: show class-of-service interface Output Fields (*continued*)

| Field Name                                 | Field Description                                                                                                                                                                                                                                                                                     |
|--------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Total non-default queues created</b>    | Number of queues created in addition to the default queues. Supported only on Trio MPC/MIC interfaces on MX Series routers.<br><br>(Enhanced subscriber management for MX Series routers) This field is not displayed for enhanced subscriber management.                                             |
| <b>Rewrite Input IEEE Code-point</b>       | (QFX Series only) IEEE 802.1p code point (priority) rewrite value. Incoming traffic from the Fibre Channel (FC) SAN is classified into the forwarding class specified in the native FC interface (NP_Port) fixed classifier and uses the priority specified as the IEEE 802.1p rewrite value.         |
| <b>Shaping rate</b>                        | Maximum transmission rate on the physical interface. You can configure the shaping rate on the physical interface, or on the logical interface, but not on both. Therefore, the <b>Shaping rate</b> field is displayed for either the physical interface or the logical interface.                    |
| <b>Scheduler map</b>                       | Name of the output scheduler map associated with this interface.<br><br>(Enhanced subscriber management for MX Series routers) The name of the dynamic scheduler map object is associated with a generated UID (for example, <b>SMAP-1_UID1002</b> ) instead of with a subscriber interface.          |
| <b>Scheduler map forwarding class sets</b> | (QFX Series only) Name of the output fabric scheduler map associated with a QFabric system Interconnect device interface.                                                                                                                                                                             |
| <b>Input shaping rate</b>                  | For Gigabit Ethernet IQ2 PICs, maximum transmission rate on the input interface.                                                                                                                                                                                                                      |
| <b>Input scheduler map</b>                 | For Gigabit Ethernet IQ2 PICs, name of the input scheduler map associated with this interface.                                                                                                                                                                                                        |
| <b>Chassis scheduler map</b>               | Name of the scheduler map associated with the packet forwarding component queues.                                                                                                                                                                                                                     |
| <b>Rewrite</b>                             | Name and type of the rewrite rules associated with this interface.                                                                                                                                                                                                                                    |
| <b>Traffic-control-profile</b>             | Name of the associated traffic control profile.<br><br>(Enhanced subscriber management for MX Series routers) The name of the dynamic traffic control profile object is associated with a generated UID (for example, <b>TC_PROF_100_199_SERIES_UID1006</b> ) instead of with a subscriber interface. |
| <b>Classifier</b>                          | Name and type of classifiers associated with this interface.                                                                                                                                                                                                                                          |
| <b>Forwarding-class-map</b>                | Name of the forwarding map associated with this interface.                                                                                                                                                                                                                                            |
| <b>Congestion-notification</b>             | (QFX Series and EX4600 switches only) Congestion notification state, <b>enabled</b> or <b>disabled</b> .                                                                                                                                                                                              |
| <b>Logical interface</b>                   | Name of a logical interface.                                                                                                                                                                                                                                                                          |
| <b>Object</b>                              | Category of an object: <b>Classifier</b> , <b>Fragmentation-map</b> (for LSQ interfaces only), <b>Scheduler-map</b> , <b>Rewrite</b> , <b>Translation Table</b> (for IQE PICs only), or <b>traffic-class-map</b> (for T4000 routers with Type 5 FPCs).                                                |
| <b>Name</b>                                | Name of an object.                                                                                                                                                                                                                                                                                    |

Table 629: show class-of-service interface Output Fields (*continued*)

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Type</b>             | Type of an object: <b>dscp</b> , <b>dscp-ipv6</b> , <b>exp</b> , <b>ieee-802.1</b> , <b>ip</b> , <b>inet-precedence</b> , or <b>ieee-802.1ad</b> (for traffic class map on T4000 routers with Type 5 FPCs)..                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Link-level type</b>  | Encapsulation on the physical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>MTU</b>              | MTU size on the physical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Speed</b>            | Speed at which the interface is running.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Loopback</b>         | Whether loopback is enabled and the type of loopback.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Source filtering</b> | Whether source filtering is enabled or disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Flow control</b>     | Whether flow control is enabled or disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Auto-negotiation</b> | (Gigabit Ethernet interfaces) Whether autonegotiation is enabled or disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Remote-fault</b>     | (Gigabit Ethernet interfaces) Remote fault status. <ul style="list-style-type: none"> <li>• <b>Online</b>—Autonegotiation is manually configured as online.</li> <li>• <b>Offline</b>—Autonegotiation is manually configured as offline.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Device flags</b>     | The <b>Device flags</b> field provides information about the physical device and displays one or more of the following values: <ul style="list-style-type: none"> <li>• <b>Down</b>—Device has been administratively disabled.</li> <li>• <b>Hear-Own-Xmit</b>—Device receives its own transmissions.</li> <li>• <b>Link-Layer-Down</b>—The link-layer protocol has failed to connect with the remote endpoint.</li> <li>• <b>Loopback</b>—Device is in physical loopback.</li> <li>• <b>Loop-Detected</b>—The link layer has received frames that it sent, thereby detecting a physical loopback.</li> <li>• <b>No-Carrier</b>—On media that support carrier recognition, no carrier is currently detected.</li> <li>• <b>No-Multicast</b>—Device does not support multicast traffic.</li> <li>• <b>Present</b>—Device is physically present and recognized.</li> <li>• <b>Promiscuous</b>—Device is in promiscuous mode and recognizes frames addressed to all physical addresses on the media.</li> <li>• <b>Quench</b>—Transmission on the device is quenched because the output buffer is overflowing.</li> <li>• <b>Recv-All-Multicasts</b>—Device is in multicast promiscuous mode and therefore provides no multicast filtering.</li> <li>• <b>Running</b>—Device is active and enabled.</li> </ul> |

Table 629: show class-of-service interface Output Fields (*continued*)

| Field Name             | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Interface flags</b> | <p>The <b>Interface flags</b> field provides information about the physical interface and displays one or more of the following values:</p> <ul style="list-style-type: none"> <li>• <b>Admin-Test</b>—Interface is in test mode and some sanity checking, such as loop detection, is disabled.</li> <li>• <b>Disabled</b>—Interface is administratively disabled.</li> <li>• <b>Down</b>—A hardware failure has occurred.</li> <li>• <b>Hardware-Down</b>—Interface is nonfunctional or incorrectly connected.</li> <li>• <b>Link-Layer-Down</b>—Interface keepalives have indicated that the link is incomplete.</li> <li>• <b>No-Multicast</b>—Interface does not support multicast traffic.</li> <li>• <b>No-receive No-transmit</b>—Passive monitor mode is configured on the interface.</li> <li>• <b>Point-To-Point</b>—Interface is point-to-point.</li> <li>• <b>Pop all MPLS labels from packets of depth</b>—MPLS labels are removed as packets arrive on an interface that has the <b>pop-all-labels</b> statement configured. The depth value can be one of the following: <ul style="list-style-type: none"> <li>• <b>1</b>—Takes effect for incoming packets with one label only.</li> <li>• <b>2</b>—Takes effect for incoming packets with two labels only.</li> <li>• <b>[ 1 2 ]</b>—Takes effect for incoming packets with either one or two labels.</li> </ul> </li> <li>• <b>Promiscuous</b>—Interface is in promiscuous mode and recognizes frames addressed to all physical addresses.</li> <li>• <b>Recv-All-Multicasts</b>—Interface is in multicast promiscuous mode and provides no multicast filtering.</li> <li>• <b>SNMP-Traps</b>—SNMP trap notifications are enabled.</li> <li>• <b>Up</b>—Interface is enabled and operational.</li> </ul> |
| <b>Flags</b>           | <p>The <b>Logical interface flags</b> field provides information about the logical interface and displays one or more of the following values:</p> <ul style="list-style-type: none"> <li>• <b>ACFC Encapsulation</b>—Address control field Compression (ACFC) encapsulation is enabled (negotiated successfully with a peer).</li> <li>• <b>Device-down</b>—Device has been administratively disabled.</li> <li>• <b>Disabled</b>—Interface is administratively disabled.</li> <li>• <b>Down</b>—A hardware failure has occurred.</li> <li>• <b>Clear-DF-Bit</b>—GRE tunnel or IPsec tunnel is configured to clear the Don't Fragment (DF) bit.</li> <li>• <b>Hardware-Down</b>—Interface protocol initialization failed to complete successfully.</li> <li>• <b>PFC</b>—Protocol field compression is enabled for the PPP session.</li> <li>• <b>Point-To-Point</b>—Interface is point-to-point.</li> <li>• <b>SNMP-Traps</b>—SNMP trap notifications are enabled.</li> <li>• <b>Up</b>—Interface is enabled and operational.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Encapsulation</b>   | Encapsulation on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Admin</b>           | Administrative state of the interface ( <b>Up</b> or <b>Down</b> ).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Link</b>            | Status of physical link ( <b>Up</b> or <b>Down</b> ).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Proto</b>           | Protocol configured on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

Table 629: show class-of-service interface Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Input Filter</b>            | Names of any firewall filters to be evaluated when packets are received on the interface, including any filters attached through activation of dynamic service.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Output Filter</b>           | Names of any firewall filters to be evaluated when packets are transmitted on the interface, including any filters attached through activation of dynamic service.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Link flags</b>              | Provides information about the physical link and displays one or more of the following values: <ul style="list-style-type: none"> <li>• <b>ACFC</b>—Address control field compression is configured. The Point-to-Point Protocol (PPP) session negotiates the ACFC option.</li> <li>• <b>Give-Up</b>—Link protocol does not continue connection attempts after repeated failures.</li> <li>• <b>Loose-LCP</b>—PPP does not use the Link Control Protocol (LCP) to indicate whether the link protocol is operational.</li> <li>• <b>Loose-LMI</b>—Frame Relay does not use the Local Management Interface (LMI) to indicate whether the link protocol is operational.</li> <li>• <b>Loose-NCP</b>—PPP does not use the Network Control Protocol (NCP) to indicate whether the device is operational.</li> <li>• <b>Keepalives</b>—Link protocol keepalives are enabled.</li> <li>• <b>No-Keepalives</b>—Link protocol keepalives are disabled.</li> <li>• <b>PFC</b>—Protocol field compression is configured. The PPP session negotiates the PFC option.</li> </ul> |
| <b>Hold-times</b>              | Current interface hold-time up and hold-time down, in milliseconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>CoS queues</b>              | Number of CoS queues configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Last flapped</b>            | Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago)</b> . For example, <b>Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Statistics last cleared</b> | Number and rate of bytes and packets received and transmitted on the physical interface. <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>IPv6 transit statistics</b> | Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

Table 629: show class-of-service interface Output Fields (*continued*)

| Field Name           | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Input errors</b>  | <p>Input errors on the interface. The labels are explained in the following list:</p> <ul style="list-style-type: none"> <li>• <b>Errors</b>—Sum of the incoming frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li>• <b>Runts</b>—Number of frames received that are smaller than the runt threshold.</li> <li>• <b>Giants</b>—Number of frames received that are larger than the giant threshold.</li> <li>• <b>Bucket Drops</b>—Drops resulting from the traffic load exceeding the interface transmit or receive leaky bucket configuration.</li> <li>• <b>Policed discards</b>—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle.</li> <li>• <b>L3 incompletes</b>—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. Layer 3 incomplete errors can be ignored by configuring the <b>ignore-l3-incompletes</b> statement.</li> <li>• <b>L2 channel errors</b>—Number of times the software did not find a valid logical interface for an incoming frame.</li> <li>• <b>L2 mismatch timeouts</b>—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable.</li> <li>• <b>HS link CRC errors</b>—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces.</li> <li>• <b>HS link FIFO overflows</b>—Number of FIFO overflows on the high-speed links between the ASICs responsible for handling the router interfaces.</li> </ul> |
| <b>Output errors</b> | <p>Output errors on the interface. The labels are explained in the following list:</p> <ul style="list-style-type: none"> <li>• <b>Carrier transitions</b>—Number of times the interface has gone from <b>down</b> to <b>up</b>. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC is malfunctioning.</li> <li>• <b>Errors</b>—Sum of the outgoing frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> </ul> <p><b>NOTE:</b> Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the <b>Drops</b> field does not always use the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p> <ul style="list-style-type: none"> <li>• <b>Aged packets</b>—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware.</li> <li>• <b>HS link FIFO underflows</b>—Number of FIFO underflows on the high-speed links between the ASICs responsible for handling the router interfaces.</li> <li>• <b>MTU errors</b>—Number of packets whose size exceeds the MTU of the interface.</li> </ul>                                                                                                                                                                                                                                                                                                                |
| <b>Egress queues</b> | Total number of egress queues supported on the specified interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |



Table 629: show class-of-service interface Output Fields (*continued*)

| Field Name                                  | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Queue counters</b>                       | <p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul> <p><b>NOTE:</b> Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the <b>Dropped packets</b> field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p>                                                                                                                                                                                                                                   |
| <b>SONET alarms</b><br><b>SONET defects</b> | <p>(SONET) SONET media-specific alarms and defects that prevent the interface from passing packets. When a defect persists for a certain period, it is promoted to an alarm. Based on the router configuration, an alarm can ring the red or yellow alarm bell on the router or light the red or yellow alarm LED on the craft interface. See these fields for possible alarms and defects: <b>SONET PHY</b>, <b>SONET section</b>, <b>SONET line</b>, and <b>SONET path</b>.</p>                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>SONET PHY</b>                            | <p>Counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> <li>• <b>Seconds</b>—Number of seconds the defect has been active.</li> <li>• <b>Count</b>—Number of times that the defect has gone from inactive to active.</li> <li>• <b>State</b>—State of the error. A state other than <b>OK</b> indicates a problem.</li> </ul> <p>The <b>SONET PHY</b> field has the following subfields:</p> <ul style="list-style-type: none"> <li>• <b>PLL Lock</b>—Phase-locked loop</li> <li>• <b>PHY Light</b>—Loss of optical signal</li> </ul>                                                                                                                                                                                                                                                                                        |
| <b>SONET section</b>                        | <p>Counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> <li>• <b>Seconds</b>—Number of seconds the defect has been active.</li> <li>• <b>Count</b>—Number of times that the defect has gone from inactive to active.</li> <li>• <b>State</b>—State of the error. A state other than <b>OK</b> indicates a problem.</li> </ul> <p>The <b>SONET section</b> field has the following subfields:</p> <ul style="list-style-type: none"> <li>• <b>BIP-BI</b>—Bit interleaved parity for SONET section overhead</li> <li>• <b>SEF</b>—Severely errored framing</li> <li>• <b>LOS</b>—Loss of signal</li> <li>• <b>LOF</b>—Loss of frame</li> <li>• <b>ES-S</b>—Errored seconds (section)</li> <li>• <b>SES-S</b>—Severely errored seconds (section)</li> <li>• <b>SEFS-S</b>—Severely errored framing seconds (section)</li> </ul> |

Table 629: show class-of-service interface Output Fields (*continued*)

| Field Name        | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>SONET line</b> | <p>Active alarms and defects, plus counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> <li>• <b>Seconds</b>—Number of seconds the defect has been active.</li> <li>• <b>Count</b>—Number of times that the defect has gone from inactive to active.</li> <li>• <b>State</b>—State of the error. A state other than <b>OK</b> indicates a problem.</li> </ul> <p>The <b>SONET line</b> field has the following subfields:</p> <ul style="list-style-type: none"> <li>• <b>BIP-B2</b>—Bit interleaved parity for SONET line overhead</li> <li>• <b>REI-L</b>—Remote error indication (near-end line)</li> <li>• <b>RDI-L</b>—Remote defect indication (near-end line)</li> <li>• <b>AIS-L</b>—Alarm indication signal (near-end line)</li> <li>• <b>BERR-SF</b>—Bit error rate fault (signal failure)</li> <li>• <b>BERR-SD</b>—Bit error rate defect (signal degradation)</li> <li>• <b>ES-L</b>—Errored seconds (near-end line)</li> <li>• <b>SES-L</b>—Severely errored seconds (near-end line)</li> <li>• <b>UAS-L</b>—Unavailable seconds (near-end line)</li> <li>• <b>ES-LFE</b>—Errored seconds (far-end line)</li> <li>• <b>SES-LFE</b>—Severely errored seconds (far-end line)</li> <li>• <b>UAS-LFE</b>—Unavailable seconds (far-end line)</li> </ul>      |
| <b>SONET path</b> | <p>Active alarms and defects, plus counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> <li>• <b>Seconds</b>—Number of seconds the defect has been active.</li> <li>• <b>Count</b>—Number of times that the defect has gone from inactive to active.</li> <li>• <b>State</b>—State of the error. A state other than <b>OK</b> indicates a problem.</li> </ul> <p>The <b>SONET path</b> field has the following subfields:</p> <ul style="list-style-type: none"> <li>• <b>BIP-B3</b>—Bit interleaved parity for SONET section overhead</li> <li>• <b>REI-P</b>—Remote error indication</li> <li>• <b>LOP-P</b>—Loss of pointer (path)</li> <li>• <b>AIS-P</b>—Path alarm indication signal</li> <li>• <b>RDI-P</b>—Path remote defect indication</li> <li>• <b>UNEQ-P</b>—Path unequipped</li> <li>• <b>PLM-P</b>—Path payload (signal) label mismatch</li> <li>• <b>ES-P</b>—Errored seconds (near-end STS path)</li> <li>• <b>SES-P</b>—Severely errored seconds (near-end STS path)</li> <li>• <b>UAS-P</b>—Unavailable seconds (near-end STS path)</li> <li>• <b>ES-PFE</b>—Errored seconds (far-end STS path)</li> <li>• <b>SES-PFE</b>—Severely errored seconds (far-end STS path)</li> <li>• <b>UAS-PFE</b>—Unavailable seconds (far-end STS path)</li> </ul> |

Table 629: show class-of-service interface Output Fields (*continued*)

| Field Name                             | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Received SONET overhead                | Values of the received and transmitted SONET overhead: <ul style="list-style-type: none"> <li>• <b>C2</b>—Signal label. Allocated to identify the construction and content of the STS-level SPE and for PDI-P.</li> <li>• <b>F1</b>—Section user channel byte. This byte is set aside for the purposes of users.</li> <li>• <b>K1 and K2</b>—These bytes are allocated for APS signaling for the protection of the multiplex section.</li> <li>• <b>J0</b>—Section trace. This byte is defined for STS-1 number 1 of an STS-N signal. Used to transmit a 1-byte fixed-length string or a 16-byte message so that a receiving terminal in a section can verify its continued connection to the intended transmitter.</li> <li>• <b>S1</b>—Synchronization status. The S1 byte is located in the first STS-1 number of an STS-N signal.</li> <li>• <b>Z3 and Z4</b>—Allocated for future use.</li> </ul>                                                                                                                                                                                  |
| Transmitted SONET overhead             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Received path trace                    | SONET/SDH interfaces allow path trace bytes to be sent inband across the SONET/SDH link. Juniper Networks and other router manufacturers use these bytes to help diagnose misconfigurations and network errors by setting the transmitted path trace message so that it contains the system hostname and name of the physical interface. The received path trace value is the message received from the router at the other end of the fiber. The transmitted path trace value is the message that this router transmits.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Transmitted path trace                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| HDLC configuration                     | Information about the HDLC configuration. <ul style="list-style-type: none"> <li>• <b>Policing bucket</b>—Configured state of the receiving policer.</li> <li>• <b>Shaping bucket</b>—Configured state of the transmitting shaper.</li> <li>• <b>Giant threshold</b>—Giant threshold programmed into the hardware.</li> <li>• <b>Runt threshold</b>—Runt threshold programmed into the hardware.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Packet Forwarding Engine configuration | Information about the configuration of the Packet Forwarding Engine: <ul style="list-style-type: none"> <li>• <b>Destination slot</b>—FPC slot number.</li> <li>• <b>PLP byte</b>—Packet Level Protocol byte.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| CoS information                        | Information about the CoS queue for the physical interface. <ul style="list-style-type: none"> <li>• <b>CoS transmit queue</b>—Queue number and its associated user-configured forwarding class name.</li> <li>• <b>Bandwidth %</b>—Percentage of bandwidth allocated to the queue.</li> <li>• <b>Bandwidth bps</b>—Bandwidth allocated to the queue (in bps).</li> <li>• <b>Buffer %</b>—Percentage of buffer space allocated to the queue.</li> <li>• <b>Buffer usec</b>—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time.</li> <li>• <b>Priority</b>—Queue priority: <b>low</b> or <b>high</b>.</li> <li>• <b>Limit</b>—Displayed if rate limiting is configured for the queue. Possible values are <b>none</b> and <b>exact</b>. If <b>exact</b> is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If <b>none</b> is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.</li> </ul> |
| Forwarding classes                     | Total number of forwarding classes supported on the specified interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Egress queues                          | Total number of egress queues supported on the specified interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

Table 629: show class-of-service interface Output Fields (*continued*)

| Field Name           | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Queue                | Queue number.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Forwarding classes   | Forwarding class name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Queued Packets       | Number of packets queued to this queue.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Queued Bytes         | Number of bytes queued to this queue. The byte counts vary by PIC type.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Transmitted Packets  | Number of packets transmitted by this queue. When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (displayed under the <b>Packet Forwarding Engine Chassis Queues</b> field) shows the prefragmentation values.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Transmitted Bytes    | Number of bytes transmitted by this queue. The byte counts vary by PIC type.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Tail-dropped packets | Number of packets dropped because of tail drop.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| RED-dropped packets  | <p>Number of packets dropped because of random early detection (RED).</p> <ul style="list-style-type: none"> <li>• (M Series and T Series routers only) On M320 and M120 routers and the T Series routers, the total number of dropped packets is displayed. On all other M Series routers, the output classifies dropped packets into the following categories: <ul style="list-style-type: none"> <li>• <b>Low, non-TCP</b>—Number of low-loss priority non-TCP packets dropped because of RED.</li> <li>• <b>Low, TCP</b>—Number of low-loss priority TCP packets dropped because of RED.</li> <li>• <b>High, non-TCP</b>—Number of high-loss priority non-TCP packets dropped because of RED.</li> <li>• <b>High, TCP</b>—Number of high-loss priority TCP packets dropped because of RED.</li> </ul> </li> <li>• (MX Series routers with enhanced DPCs, and T Series routers with enhanced FPCs only) The output classifies dropped packets into the following categories: <ul style="list-style-type: none"> <li>• <b>Low</b>—Number of low-loss priority packets dropped because of RED.</li> <li>• <b>Medium-low</b>—Number of medium-low loss priority packets dropped because of RED.</li> <li>• <b>Medium-high</b>—Number of medium-high loss priority packets dropped because of RED.</li> <li>• <b>High</b>—Number of high-loss priority packets dropped because of RED.</li> </ul> </li> </ul> <p><b>NOTE:</b> Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), this field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p> |

Table 629: show class-of-service interface Output Fields (*continued*)

| Field Name        | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| RED-dropped bytes | <p>Number of bytes dropped because of RED. The byte counts vary by PIC type.</p> <ul style="list-style-type: none"> <li>(M Series and T Series routers only) On M320 and M120 routers and the T Series routers, only the total number of dropped bytes is displayed. On all other M Series routers, the output classifies dropped bytes into the following categories: <ul style="list-style-type: none"> <li><b>Low, non-TCP</b>—Number of low-loss priority non-TCP bytes dropped because of RED.</li> <li><b>Low, TCP</b>—Number of low-loss priority TCP bytes dropped because of RED.</li> <li><b>High, non-TCP</b>—Number of high-loss priority non-TCP bytes dropped because of RED.</li> <li><b>High, TCP</b>—Number of high-loss priority TCP bytes dropped because of RED.</li> </ul> </li> </ul> <p><b>NOTE:</b> Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), this field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p> |
| Transmit rate     | Configured transmit rate of the scheduler. The rate is a percentage of the total interface bandwidth.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Rate Limit        | <p>Rate limiting configuration of the queue. Possible values are :</p> <ul style="list-style-type: none"> <li><b>None</b>—No rate limit.</li> <li><b>exact</b>—Queue transmits at the configured rate.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Buffer size       | Delay buffer size in the queue.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Priority          | Scheduling priority configured as <b>low</b> or <b>high</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Excess Priority   | Priority of the excess bandwidth traffic on a scheduler: <b>low</b> , <b>medium-low</b> , <b>medium-high</b> , <b>high</b> , or <b>none</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Drop profiles     | <p>Display the assignment of drop profiles.</p> <ul style="list-style-type: none"> <li><b>Loss priority</b>—Packet loss priority for drop profile assignment.</li> <li><b>Protocol</b>—Transport protocol for drop profile assignment.</li> <li><b>Index</b>—Index of the indicated object. Objects that have indexes in this output include schedulers and drop profiles.</li> <li><b>Name</b>—Name of the drop profile.</li> <li><b>Type</b>—Type of the drop profile: <b>discrete</b> or <b>interpolated</b>.</li> <li><b>Fill Level</b>—Percentage fullness of a queue.</li> <li><b>Drop probability</b>—Drop probability at this fill level.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                              |
| Excess Priority   | Priority of the excess bandwidth traffic on a scheduler.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

Table 629: show class-of-service interface Output Fields (*continued*)

| Field Name                    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Drop profiles</b>          | <p>Display the assignment of drop profiles.</p> <ul style="list-style-type: none"> <li>• <b>Loss priority</b>—Packet loss priority for drop profile assignment.</li> <li>• <b>Protocol</b>—Transport protocol for drop profile assignment.</li> <li>• <b>Index</b>—Index of the indicated object. Objects that have indexes in this output include schedulers and drop profiles.</li> <li>• <b>Name</b>—Name of the drop profile.</li> <li>• <b>Type</b>—Type of the drop profile: <b>discrete</b> or <b>interpolated</b>.</li> <li>• <b>Fill Level</b>—Percentage fullness of a queue.</li> <li>• <b>Drop probability</b>—Drop probability at this fill level.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Adjustment information</b> | <p>Display the assignment of shaping-rate adjustments on a scheduler node or queue.</p> <ul style="list-style-type: none"> <li>• <b>Adjusting application</b>—Application that is performing the shaping-rate adjustment. <ul style="list-style-type: none"> <li>• The adjusting application can appear as <b>ancp LS-0</b>, which is the Junos OS Access Node Control Profile process (<b>ancpd</b>) that performs shaping-rate adjustments on schedule nodes.</li> <li>• The adjusting application can also appear as <b>pppoe</b>, which adjusts the shaping-rate and overhead-accounting class-of-service attributes on dynamic subscriber interfaces in a broadband access network based on access line parameters in Point-to-Point Protocol over Ethernet (PPPoE) Tags [TR-101]. This feature is supported on MPC/MIC interfaces on MX Series routers. The shaping rate is based on the actual-data-rate-downstream attribute. The overhead accounting value is based on the access-loop-encapsulation attribute and specifies whether the access loop uses Ethernet (frame mode) or ATM (cell mode).</li> </ul> </li> <li>• <b>Adjustment type</b>—Type of adjustment: <b>absolute</b> or <b>delta</b>.</li> <li>• <b>Configured shaping rate</b>—Shaping rate configured for the scheduler node or queue.</li> <li>• <b>Adjustment value</b>—Value of adjusted shaping rate.</li> <li>• <b>Adjustment target</b>—Level of shaping-rate adjustment performed: <b>node</b> or <b>queue</b>.</li> <li>• <b>Adjustment overhead-accounting mode</b>—Configured shaping mode: <b>frame</b> or <b>cell</b>.</li> <li>• <b>Adjustment overhead bytes</b>—Number of bytes that the ANCP agent adds to or subtracts from the actual downstream frame overhead before reporting the adjusted values to CoS.</li> <li>• <b>Adjustment target</b>—Level of shaping-rate adjustment performed: <b>node</b> or <b>queue</b>.</li> <li>• <b>Adjustment multicast index</b>—</li> </ul> |

## Sample Output

### show class-of-service interface (Physical)

```

user@host> show class-of-service interface so-0/2/3
Physical interface: so-0/2/3, Index: 135
Queues supported: 8, Queues in use: 4
Total non-default queues created: 4
Scheduler map: <default>, Index: 2032638653

Logical interface: fe-0/0/1.0, Index: 68, Dedicated Queues: no
Shaping rate: 32000

```

| Object        | Name        | Type | Index |
|---------------|-------------|------|-------|
| Scheduler-map | <default>   |      | 27    |
| Rewrite       | exp-default | exp  | 21    |
| Classifier    | exp-default | exp  | 5     |

|                      |                      |     |   |
|----------------------|----------------------|-----|---|
| Classifier           | ipprec-compatibility | ip  | 8 |
| Forwarding-class-map | exp-default          | exp | 5 |

### show class-of-service interface (Logical)

```

user@host> show class-of-service interface so-0/2/3.0
Logical interface: so-0/2/3.0, Index: 68, Dedicated Queues: no
Shaping rate: 32000
Object Name Type Index
Scheduler-map <default>
Rewrite exp-default exp 21
Classifier exp-default exp 5
Classifier ipprec-compatibility ip 8
Forwarding-class-map exp-default exp 5

```

### show class-of-service interface (Gigabit Ethernet)

```

user@host> show class-of-service interface ge-6/2/0
Physical interface: ge-6/2/0, Index: 175
Queues supported: 4, Queues in use: 4
Scheduler map: <default>, Index: 2
Input scheduler map: <default>, Index: 3
Chassis scheduler map: <default-chassis>, Index: 4

```

### show class-of-service interface (ANCP)

```

user@host> show class-of-service interface pp0.1073741842
Logical interface: pp0.1073741842, Index: 341
Object Name Type Index
Traffic-control-profile TCP-CVLAN Output 12408
Classifier dscp-ipv6-compatibility dscp-ipv6 9
Classifier ipprec-compatibility ip 13

Adjusting application: ancp LS-0
Adjustment type: absolute
Configured shaping rate: 4000000
Adjustment value: 11228000
Adjustment overhead-accounting mode: Frame Mode
Adjustment overhead bytes: 50
Adjustment target: node

```

### show class-of-service interface (PPPoE Interface)

```

user@host> show class-of-service interface pp0.1
Logical interface: pp0.1, Index: 85
Object Name Type Index
Traffic-control-profile tcp-pppoe.o.pp0.1 Output 2726446535
Classifier ipprec-compatibility ip 13

Adjusting application: PPPoE
Adjustment type: absolute
Adjustment value: 5000000
Adjustment overhead-accounting mode: cell
Adjustment target: node

```

### show class-of-service interface (T4000 Routers with Type 5 FPCs)

```

user@host> show class-of-service interface xe-4/0/0
Physical interface: xe-4/0/0, Index: 153
Queues supported: 8, Queues in use: 4
Shaping rate: 5000000000 bps
Scheduler map: <default>, Index: 2

```

```

Congestion-notification: Disabled

Logical interface: xe-4/0/0.0, Index: 77
 Object Name Type
Index
 Classifier ipprec-compatibility ip
13

```

### show class-of-service interface detail

```
user@host> show class-of-service interface ge-0/3/0 detail
```

```

Physical interface: ge-0/3/0, Enabled, Physical link is Up
 Link-level type: Ethernet, MTU: 1518, Speed: 1000mbps, Loopback: Disabled,
 Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
 Remote fault: Online
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x4000

```

```

Physical interface: ge-0/3/0, Index: 138
Queues supported: 4, Queues in use: 5
Shaping rate: 50000 bps
Scheduler map: interface-scheduler-map, Index: 58414
Input shaping rate: 10000 bps
Input scheduler map: scheduler-map, Index: 15103
Chassis scheduler map: <default-chassis>, Index: 4
Congestion-notification: Disabled

```

```

Logical interface ge-0/3/0.0
 Flags: SNMP-Traps 0x4000 VLAN-Tag [0x8100.1] Encapsulation: ENET2
 inet
 mpls

```

| Interface  | Admin | Link | Proto | Input Filter  | Output Filter  |
|------------|-------|------|-------|---------------|----------------|
| ge-0/3/0.0 | up    | up   | inet  |               |                |
|            |       |      | mpls  |               |                |
| Interface  | Admin | Link | Proto | Input Policer | Output Policer |
| ge-0/3/0.0 | up    | up   | inet  |               |                |
|            |       |      | mpls  |               |                |

```

Logical interface: ge-0/3/0.0, Index: 68
 Object Name Type Index
 Rewrite exp-default exp (mpls-any) 33
 Classifier exp-default exp 10
 Classifier ipprec-compatibility ip 13

```

```

Logical interface ge-0/3/0.1
 Flags: SNMP-Traps 0x4000 VLAN-Tag [0x8100.2] Encapsulation: ENET2
 inet

```

| Interface  | Admin | Link | Proto | Input Filter  | Output Filter  |
|------------|-------|------|-------|---------------|----------------|
| ge-0/3/0.1 | up    | up   | inet  |               |                |
| Interface  | Admin | Link | Proto | Input Policer | Output Policer |
| ge-0/3/0.1 | up    | up   | inet  |               |                |

```

Logical interface: ge-0/3/0.1, Index: 69
 Object Name Type Index
 Classifier ipprec-compatibility ip 13

```

### show class-of-service interface comprehensive

```
user@host> show class-of-service interface ge-0/3/0 comprehensive
```



```

Physical interface: ge-0/3/0, Enabled, Physical link is Up
 Interface index: 138, SNMP ifIndex: 601, Generation: 141
 Link-level type: Ethernet, MTU: 1518, Speed: 1000Mbps, BPDU Error: None,
 MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled, Flow
 control: Enabled,
 Auto-negotiation: Enabled, Remote fault: Online
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x4000
 CoS queues : 4 supported, 4 maximum usable queues
 Schedulers : 256
 Hold-times : Up 0 ms, Down 0 ms
 Current address: 00:14:f6:f4:b4:5d, Hardware address: 00:14:f6:f4:b4:5d
 Last flapped : 2010-09-07 06:35:22 PDT (15:14:42 ago)
 Statistics last cleared: Never
Traffic statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps
IPv6 total statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Ingress traffic statistics at Packet Forwarding Engine:
 Input bytes : 0 0 bps
 Input packets: 0 0 pps
 Drop bytes : 0 0 bps
 Drop packets: 0 0 pps
Label-switched interface (LSI) traffic statistics:
 Input bytes : 0 0 bps
 Input packets: 0 0 pps
Input errors:
 Errors: 0, Drops: 0, Framing errors: 0, Runt: 0, Policed discards: 0, L3
 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
 Resource errors: 0
Output errors:
 Carrier transitions: 5, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
 FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Ingress queues: 4 supported, 5 in use
Queue counters: Queued packets Transmitted packets Dropped packets

 0 af3 0 0 0
 1 af2 0 0 0
 2 ef2 0 0 0
 3 ef1 0 0 0

Egress queues: 4 supported, 5 in use
Queue counters: Queued packets Transmitted packets Dropped packets

 0 af3 0 0 0
 1 af2 0 0 0
 2 ef2 0 0 0
 3 ef1 0 0 0

```

```

Active alarms : None
Active defects : None
MAC statistics:
 Receive Transmit
 Total octets 0 0
 Total packets 0 0
 Unicast packets 0 0
 Broadcast packets 0 0
 Multicast packets 0 0
 CRC/Align errors 0 0
 FIFO errors 0 0
 MAC control frames 0 0
 MAC pause frames 0 0
 Oversized frames 0
 Jabber frames 0
 Fragment frames 0
 VLAN tagged frames 0
 Code violations 0
Filter statistics:
 Input packet count 0
 Input packet rejects 0
 Input DA rejects 0
 Input SA rejects 0
 Output packet count 0
 Output packet pad count 0
 Output packet error count 0
 CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
 Negotiation status: Complete
 Link partner:
 Link mode: Full-duplex, Flow control: Symmetric/Asymmetric, Remote fault:
OK
 Local resolution:
 Flow control: Symmetric, Remote fault: Link OK
Packet Forwarding Engine configuration:
 Destination slot: 0
CoS information:
 Direction : Output
 CoS transmit queue Bandwidth Buffer Priority
Limit
 % bps % usec
 2 ef2 39 19500 0 120 high
none
 Direction : Input
 CoS transmit queue Bandwidth Buffer Priority
Limit
 % bps % usec
 0 af3 30 3000 45 0 low
none

Physical interface: ge-0/3/0, Enabled, Physical link is Up
 Interface index: 138, SNMP ifIndex: 601
Forwarding classes: 16 supported, 5 in use
Ingress queues: 4 supported, 5 in use
Queue: 0, Forwarding classes: af3
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : Not Available

```

```

RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: af2
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: ef2
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: ef1
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Forwarding classes: 16 supported, 5 in use
Egress queues: 4 supported, 5 in use
Queue: 0, Forwarding classes: af3
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RL-dropped packets : 0 0 pps
 RL-dropped bytes : 0 0 bps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: af2
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RL-dropped packets : 0 0 pps
 RL-dropped bytes : 0 0 bps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: ef2
 Queued:
 Packets : 0 0 pps

```

```

 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RL-dropped packets : 0 0 pps
 RL-dropped bytes : 0 0 bps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: ef1
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RL-dropped packets : 0 0 pps
 RL-dropped bytes : 0 0 bps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 5 in use
Queue: 0, Forwarding classes: af3
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : Not Available
 RED-dropped bytes : Not Available
Queue: 1, Forwarding classes: af2
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : Not Available
 RED-dropped bytes : Not Available
Queue: 2, Forwarding classes: ef2
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : Not Available
 RED-dropped bytes : Not Available
Queue: 3, Forwarding classes: ef1
 Queued:
 Packets : 108546 0 pps
 Bytes : 12754752 376 bps
 Transmitted:
 Packets : 108546 0 pps
 Bytes : 12754752 376 bps

```

```

Tail-dropped packets : 0 0 pps
RED-dropped packets : Not Available
RED-dropped bytes : Not Available

```

```

Physical interface: ge-0/3/0, Index: 138
Queues supported: 4, Queues in use: 5
Shaping rate: 50000 bps

```

```
Scheduler map: interface-scheduler-map, Index: 58414
```

```

Scheduler: ef2, Forwarding class: ef2, Index: 39155
 Transmit rate: 39 percent, Rate Limit: none, Buffer size: 120 us, Buffer
 Limit: none, Priority: high
 Excess Priority: unspecified
 Drop profiles:
 Loss priority Protocol Index Name
 Low any 1 < default-drop-profile>
 Medium low any 1 < default-drop-profile>
 Medium high any 1 < default-drop-profile>
 High any 1 < default-drop-profile>
 Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
 Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
 Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
 Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
 Input shaping rate: 10000 bps
 Input scheduler map: scheduler-map

```

```
Scheduler map: scheduler-map, Index: 15103
```

```

Scheduler: af3, Forwarding class: af3, Index: 35058
 Transmit rate: 30 percent, Rate Limit: none, Buffer size: 45 percent, Buffer
 Limit: none, Priority: low
 Excess Priority: unspecified
 Drop profiles:
 Loss priority Protocol Index Name
 Low any 40582 green
 Medium low any 1 < default-drop-profile>
 Medium high any 1 < default-drop-profile>
 High any 18928 yellow
 Drop profile: green, Type: discrete, Index: 40582
 Fill level Drop probability
 50 0
 100 100
 Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
 Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
 Drop profile: yellow, Type: discrete, Index: 18928
 Fill level Drop probability
 50 0
 100 100

```

```

Chassis scheduler map: < default-drop-profile>
Scheduler map: < default-drop-profile>, Index: 4

Scheduler: < default-drop-profile>, Forwarding class: af3, Index: 25
 Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer
 Limit: none, Priority: low
 Excess Priority: low
 Drop profiles:
 Loss priority Protocol Index Name
 Low any 1 < default-drop-profile>
 Medium low any 1 < default-drop-profile>
 Medium high any 1 < default-drop-profile>
 High any 1 < default-drop-profile>
 Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
 Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
 Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
 Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100

Scheduler: < default-drop-profile>, Forwarding class: af2, Index: 25
 Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer
 Limit: none, Priority: low
 Excess Priority: low
 Drop profiles:
 Loss priority Protocol Index Name
 Low any 1 < default-drop-profile>
 Medium low any 1 < default-drop-profile>
 Medium high any 1 < default-drop-profile>
 High any 1 < default-drop-profile>
 Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
 Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
 Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
 Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100

Scheduler: < default-drop-profile>, Forwarding class: ef2, Index: 25
 Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer
 Limit: none, Priority: low
 Excess Priority: low
 Drop profiles:
 Loss priority Protocol Index Name
 Low any 1 < default-drop-profile>
 Medium low any 1 < default-drop-profile>
 Medium high any 1 < default-drop-profile>
 High any 1 < default-drop-profile>
 Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100

```

```

 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100

Scheduler: < default-drop-profile>, Forwarding class: ef1, Index: 25
 Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer
Limit: none, Priority: low
 Excess Priority: low
 Drop profiles:
 Loss priority Protocol Index Name
 Low any 1 < default-drop-profile>
 Medium low any 1 < default-drop-profile>
 Medium high any 1 < default-drop-profile>
 High any 1 < default-drop-profile>
Drop profile: , Type: discrete, Index: 1
 Fill level Drop probability
 100 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
 Congestion-notification: Disabled
Forwarding class
priority Policing priority ID Queue Restricted queue Fabric
af3 normal 0 0 0 low
af2 normal 1 1 1 low
ef2 normal 2 2 2 high
ef1 normal 3 3 3 high
af1 normal 4 4 0 low

Logical interface ge-0/3/0.0 (Index 68) (SNMP ifIndex 152) (Generation 159)
 Flags: SNMP-Traps 0x4000 VLAN-Tag [0x8100.1] Encapsulation: ENET2
 Traffic statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
 Local statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
 Transit statistics:
 Input bytes : 0 0 bps

```

```

Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
Protocol inet, MTU: 1500, Generation: 172, Route table: 0
 Flags: Sendbcst-pkt-to-re
 Input Filters: filter-in-ge-0/3/0.0-i,
 Policer: Input: p1-ge-0/3/0.0-inet-i
Protocol mpls, MTU: 1488, Maximum labels: 3, Generation: 173, Route table: 0
 Flags: Is-Primary
 Output Filters: exp-filter,,,,,

Logical interface ge-1/2/0.0 (Index 347) (SNMP ifIndex 638) (Generation 156)

Forwarding class ID Queue Restricted queue Fabric priority Policing priority
 SPU priority
best-effort 0 0 0 low normal
 low

Aggregate Forwarding-class statistics per forwarding-class
Aggregate Forwarding-class statistics:
Forwarding-class statistics:

Forwarding-class best-effort statistics:
 Input unicast bytes: 0
 Output unicast bytes: 0
 Input unicast packets: 0
 Output unicast packets: 0

 Input multicast bytes: 0
 Output multicast bytes: 0
 Input multicast packets: 0
 Output multicast packets: 0

Forwarding-class expedited-forwarding statistics:
 Input unicast bytes: 0
 Output unicast bytes: 0
 Input unicast packets: 0
 Output unicast packets: 0

 Input multicast bytes: 0
 Output multicast bytes: 0
 Input multicast packets: 0
 Output multicast packets: 0

IPv4 protocol forwarding-class statistics:
Forwarding-class statistics:
Forwarding-class best-effort statistics:

 Input unicast bytes: 0
 Output unicast bytes: 0
 Input unicast packets: 0
 Output unicast packets: 0

 Input multicast bytes: 0
 Output multicast bytes: 0
 Input multicast packets: 0
 Output multicast packets: 0

Forwarding-class expedited-forwarding statistics:
 Input unicast bytes: 0

```



```

Output unicast bytes: 0
Input unicast packets: 0
Output unicast packets: 0

```

```

Input multicast bytes: 0
Output multicast bytes: 0
Input multicast packets: 0
Output multicast packets: 0

```

```

IPv6 protocol forwarding-class statistics:
Forwarding-class statistics:
 Forwarding-class best-effort statistics:

```

```

Input unicast bytes: 0
Output unicast bytes: 0
Input unicast packets: 0
Output unicast packets: 0

```

```

Input multicast bytes: 0
Output multicast bytes: 0
Input multicast packets: 0
Output multicast packets: 0

```

```
Forwarding-class expedited-forwarding statistics:
```

```

Input unicast bytes: 0
Output unicast bytes: 0
Input unicast packets: 0
Output unicast packets: 0

```

```
Logical interface ge-0/3/0.0 (Index 68) (SNMP ifIndex 152)
```

```

Flags: SNMP-Traps 0x4000 VLAN-Tag [0x8100.1] Encapsulation: ENET2
Input packets : 0
Output packets: 0

```

|            |       |      |       |                        |               |
|------------|-------|------|-------|------------------------|---------------|
| Interface  | Admin | Link | Proto | Input Filter           | Output Filter |
| ge-0/3/0.0 | up    | up   | inet  | filter-in-ge-0/3/0.0-i |               |
|            |       |      | mpls  |                        | exp-filter    |

|            |       |      |       |                      |                |
|------------|-------|------|-------|----------------------|----------------|
| Interface  | Admin | Link | Proto | Input Policer        | Output Policer |
| ge-0/3/0.0 | up    | up   |       |                      |                |
|            |       |      | inet  | p1-ge-0/3/0.0-inet-i |                |
|            |       |      | mpls  |                      |                |

```
Filter: filter-in-ge-0/3/0.0-i
```

```
Counters:
```

| Name                         | Bytes | Packets |
|------------------------------|-------|---------|
| count-filter-in-ge-0/3/0.0-i | 0     | 0       |

```
Filter: exp-filter
```

```
Counters:
```

| Name                  | Bytes | Packets |
|-----------------------|-------|---------|
| count-exp-seven-match | 0     | 0       |
| count-exp-zero-match  | 0     | 0       |

```
Policers:
```

| Name                 | Packets |
|----------------------|---------|
| p1-ge-0/3/0.0-inet-i | 0       |

```
Logical interface: ge-0/3/0.0, Index: 68
```

| Object  | Name        | Type           | Index |
|---------|-------------|----------------|-------|
| Rewrite | exp-default | exp (mpls-any) | 33    |

Rewrite rule: exp-default, Code point type: exp, Index: 33

| Forwarding class | Loss priority | Code point |       |
|------------------|---------------|------------|-------|
| af3              | low           | 000        |       |
| af3              | high          | 001        |       |
| af2              | low           | 010        |       |
| af2              | high          | 011        |       |
| ef2              | low           | 100        |       |
| ef2              | high          | 101        |       |
| ef1              | low           | 110        |       |
| ef1              | high          | 111        |       |
| Object           | Name          | Type       | Index |
| Classifier       | exp-default   | exp        | 10    |

Classifier: exp-default, Code point type: exp, Index: 10

| Code point | Forwarding class     | Loss priority |       |
|------------|----------------------|---------------|-------|
| 000        | af3                  | low           |       |
| 001        | af3                  | high          |       |
| 010        | af2                  | low           |       |
| 011        | af2                  | high          |       |
| 100        | ef2                  | low           |       |
| 101        | ef2                  | high          |       |
| 110        | ef1                  | low           |       |
| 111        | ef1                  | high          |       |
| Object     | Name                 | Type          | Index |
| Classifier | ipprec-compatibility | ip            | 13    |

Classifier: ipprec-compatibility, Code point type: inet-precedence, Index: 13

| Code point       | Forwarding class | Loss priority |                  |        |
|------------------|------------------|---------------|------------------|--------|
| 000              | af3              | low           |                  |        |
| 001              | af3              | high          |                  |        |
| 010              | af3              | low           |                  |        |
| 011              | af3              | high          |                  |        |
| 100              | af3              | low           |                  |        |
| 101              | af3              | high          |                  |        |
| 110              | ef1              | low           |                  |        |
| 111              | ef1              | high          |                  |        |
| Forwarding class | ID               | Queue         | Restricted queue | Fabric |
| priority         |                  |               |                  |        |
| af3              | 0                | 0             | 0                | low    |
| normal           |                  |               |                  |        |
| af2              | 1                | 1             | 1                | low    |
| normal           |                  |               |                  |        |
| ef2              | 2                | 2             | 2                | high   |
| normal           |                  |               |                  |        |
| ef1              | 3                | 3             | 3                | high   |
| normal           |                  |               |                  |        |
| af1              | 4                | 4             | 0                | low    |
| normal           |                  |               |                  |        |

Logical interface ge-0/3/0.1 (Index 69) (SNMP ifIndex 154) (Generation 160)

Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.2 ] Encapsulation: ENET2

Traffic statistics:

|                   |   |
|-------------------|---|
| Input bytes :     | 0 |
| Output bytes :    | 0 |
| Input packets:    | 0 |
| Output packets:   | 0 |
| Local statistics: |   |
| Input bytes :     | 0 |

```

Output bytes : 0
Input packets: 0
Output packets: 0
Transit statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
Protocol inet, MTU: 1500, Generation: 174, Route table: 0
Flags: Sendbcst-pkt-to-re

```

```

Logical interface ge-0/3/0.1 (Index 69) (SNMP ifIndex 154)
Flags: SNMP-Traps 0x4000 VLAN-Tag [0x8100.2] Encapsulation: ENET2
Input packets : 0
Output packets: 0

```

```

Interface Admin Link Proto Input Filter Output Filter
ge-0/3/0.1 up up mpls
Interface Admin Link Proto Input Policer Output Policer
ge-0/3/0.1 up up mpls

```

```
Logical interface: ge-0/3/0.1, Index: 69
```

| Object     | Name                 | Type | Index |
|------------|----------------------|------|-------|
| Classifier | ipprec-compatibility | ip   | 13    |

```
Classifier: ipprec-compatibility, Code point type: inet-precedence, Index: 13
```

| Code point | Forwarding class | Loss priority |
|------------|------------------|---------------|
| 000        | af3              | low           |
| 001        | af3              | high          |
| 010        | af3              | low           |
| 011        | af3              | high          |
| 100        | af3              | low           |
| 101        | af3              | high          |
| 110        | ef1              | low           |
| 111        | ef1              | high          |

| Forwarding class | ID | Queue | Restricted queue | Fabric |
|------------------|----|-------|------------------|--------|
| af3              | 0  | 0     | 0                | low    |
| af2              | 1  | 1     | 1                | low    |
| ef2              | 2  | 2     | 2                | high   |
| ef1              | 3  | 3     | 3                | high   |
| af1              | 4  | 4     | 0                | low    |

### show class-of-service interface (ACX Series Routers)

```

user@host-g11# show class-of-service interface
Physical interface: at-0/0/0, Index: 130
Queues supported: 4, Queues in use: 4
Scheduler map: <default>, Index: 2
Congestion-notification: Disabled

```

Logical interface: at-0/0/0.0, Index: 69

Logical interface: at-0/0/0.32767, Index: 70

Physical interface: at-0/0/1, Index: 133

Queues supported: 4, Queues in use: 4

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

Logical interface: at-0/0/1.0, Index: 71

Logical interface: at-0/0/1.32767, Index: 72

Physical interface: ge-0/1/0, Index: 146

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

| Object     | Name         | Type      | Index |
|------------|--------------|-----------|-------|
| Rewrite    | dscp-default | dscp      | 31    |
| Classifier | d1           | dscp      | 11331 |
| Classifier | ci           | ieee8021p | 583   |

Logical interface: ge-0/1/0.0, Index: 73

| Object  | Name       | Type           | Index |
|---------|------------|----------------|-------|
| Rewrite | custom-exp | exp (mpls-any) | 46413 |

Logical interface: ge-0/1/0.1, Index: 74

Logical interface: ge-0/1/0.32767, Index: 75

Physical interface: ge-0/1/1, Index: 147

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

| Object     | Name                 | Type | Index |
|------------|----------------------|------|-------|
| Classifier | ipprec-compatibility | ip   | 13    |

Logical interface: ge-0/1/1.0, Index: 76

Physical interface: ge-0/1/2, Index: 148

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

| Object     | Name | Type              | Index |
|------------|------|-------------------|-------|
| Rewrite    | ri   | ieee8021p (outer) | 35392 |
| Classifier | ci   | ieee8021p         | 583   |

Physical interface: ge-0/1/3, Index: 149

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

| Object     | Name                 | Type | Index |
|------------|----------------------|------|-------|
| Classifier | ipprec-compatibility | ip   | 13    |

Logical interface: ge-0/1/3.0, Index: 77

| Object  | Name        | Type           | Index |
|---------|-------------|----------------|-------|
| Rewrite | custom-exp2 | exp (mpls-any) | 53581 |

Physical interface: ge-0/1/4, Index: 150

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

```

 Congestion-notification: Disabled
Object Name Type Index
Classifier ipprec-compatibility ip 13

Physical interface: ge-0/1/5, Index: 151
Queues supported: 8, Queues in use: 5
Scheduler map: <default>, Index: 2
Congestion-notification: Disabled
Object Name Type Index
Classifier ipprec-compatibility ip 13

Physical interface: ge-0/1/6, Index: 152
Queues supported: 8, Queues in use: 5
Scheduler map: <default>, Index: 2
Congestion-notification: Disabled
Object Name Type Index
Classifier ipprec-compatibility ip 13

Physical interface: ge-0/1/7, Index: 153
Queues supported: 8, Queues in use: 5
Scheduler map: <default>, Index: 2
Congestion-notification: Disabled
Object Name Type Index
Classifier d1 dscp 11331

Physical interface: ge-0/2/0, Index: 154
Queues supported: 8, Queues in use: 5
Scheduler map: <default>, Index: 2
Congestion-notification: Disabled
Object Name Type Index
Classifier ipprec-compatibility ip 13

Physical interface: ge-0/2/1, Index: 155
Queues supported: 8, Queues in use: 5
Scheduler map: <default>, Index: 2
Congestion-notification: Disabled
Object Name Type Index
Classifier ipprec-compatibility ip 13

 Logical interface: ge-0/2/1.0, Index: 78

 Logical interface: ge-0/2/1.32767, Index: 79

Physical interface: xe-0/3/0, Index: 156
Queues supported: 8, Queues in use: 5
Scheduler map: <default>, Index: 2
Congestion-notification: Disabled
Object Name Type Index
Classifier ipprec-compatibility ip 13

 Logical interface: xe-0/3/0.0, Index: 80

Physical interface: xe-0/3/1, Index: 157
Queues supported: 8, Queues in use: 5
Scheduler map: <default>, Index: 2
Congestion-notification: Disabled
Object Name Type Index
Classifier ipprec-compatibility ip 13

 Logical interface: xe-0/3/1.0, Index: 81

```

```
[edit]
user@host-g11#
```

#### show class-of-service interface (PPPoE Subscriber Interface for Enhanced Subscriber Management)

```
user@host> show class-of-service interface pp0.3221225474
 Logical interface: pp0.3221225475, Index: 3221225475
Object Name Type Index
Traffic-control-profile TC_PROF_100_199_SERIES_UID1006 Output 4294967312
Scheduler-map SMAP-1_UID1002 Output 4294967327
Rewrite-Output ieee-rewrite 60432
Rewrite-Output rule1 ip 50463

 Adjusting application: PPPoE IA tags
 Adjustment type: absolute
 Configured shaping rate: 11000000
 Adjustment value: 5000000
 Adjustment target: node

 Adjusting application: ucac
 Adjustment type: delta
 Configured shaping rate: 5000000
 Adjustment value: 100000
 Adjustment target: node
```

## show class-of-service multi-destination

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show class-of-service multi-destination                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.<br>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | For each class-of-service (CoS) multideestination classifier, display the classifier type.                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <b>none</b> —Display all multideestination classifiers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Defining CoS Multideestination (Multicast, Broadcast, DLF) BA Classifiers</i></li> <li>• <i>Example: Configuring Multideestination (Multicast, Broadcast, DLF) Classifiers</i></li> <li>• <i>Understanding CoS Classifiers</i></li> <li>• <i>Understanding CoS Classifiers</i></li> <li>• <a href="#">Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 6672</a></li> <li>• <i>Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces</i></li> </ul> |
| <b>Output Fields</b>            | <a href="#">Table 630</a> describes the output fields for the <b>show class-of-service multi-destination</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                 |

**Table 630: show class-of-service multi-destination Output Fields**

| Field Name              | Field Description                                          |
|-------------------------|------------------------------------------------------------|
| <b>Family ethernet</b>  | Family to which the classifier belongs.                    |
| <b>Classifier Name</b>  | Name of the classifier.                                    |
| <b>Classifier Type</b>  | Type of the classifier: <b>dscp</b> or <b>ieee-802.1</b> . |
| <b>Classifier Index</b> | Internal index of the classifier.                          |

## Sample Output

### show class-of-service multi-destination

```
user@switch> show class-of-service multi-destination
```

```

Family ethernet:
Classifier Name Classifier Type Classifier Index
ba-mcast-classifier ieee-802.1 62376

```

## show class-of-service rewrite-rule

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show class-of-service rewrite-rule<br><name <i>name</i> ><br><type <i>type</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Display the mapping of forwarding classes and loss priority to code point values.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <p><b>none</b>—Display all rewrite rules.</p> <p><b>name <i>name</i></b>—(Optional) Display the specified rewrite rule.</p> <p><b>type <i>type</i></b>—(Optional) Display the rewrite rule of the specified type. The rewrite rule type can be one of the following:</p> <ul style="list-style-type: none"> <li><b>dscp</b>—For IPv4 traffic.</li> <li><b>dscp-ipv6</b>—For IPv6 traffic.</li> <li><b>exp</b>—For MPLS traffic.</li> <li><b>frame-relay-de</b>— For Frame Relay traffic.</li> <li><b>ieee-802.1</b>—For Layer 2 traffic.</li> <li><b>inet-precedence</b>—For IPv4 traffic.</li> </ul> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">Rewrite Rules Overview</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>List of Sample Output</b>    | <a href="#">show class-of-service rewrite-rule type dscp on page 7289</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Output Fields</b>            | Table 631 describes the output fields for the <b>show class-of-service rewrite-rule</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                              |

**Table 631: show class-of-service rewrite-rule Output Fields**

| Field Name              | Field Description                                                                                                                          |
|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Rewrite rule</b>     | Name of the rewrite rule.                                                                                                                  |
| <b>Code point type</b>  | Type of rewrite rule: <b>dscp</b> , <b>dscp-ipv6</b> , <b>exp</b> , <b>frame-relay-de</b> , or <b>inet-precedence</b> .                    |
| <b>Forwarding class</b> | Classification of a packet affecting the forwarding, scheduling, and marking policies applied as the packet transits the router or switch. |
| <b>Index</b>            | Internal index for this particular rewrite rule.                                                                                           |



Table 631: show class-of-service rewrite-rule Output Fields (*continued*)

| Field Name    | Field Description            |
|---------------|------------------------------|
| Loss priority | Loss priority for rewriting. |
| Code point    | Code point value to rewrite. |

## Sample Output

### show class-of-service rewrite-rule type dscp

```

user@host> show class-of-service rewrite-rule type dscp
Rewrite rule: dscp-default, Code point type: dscp
 Forwarding class Loss priority Code point
 gold high 000000
 silver low 110000
 silver high 111000
 bronze low 001010
 bronze high 001100
 lead high 101110

Rewrite rule: abc-dscp-rewrite, Code point type: dscp, Index: 3245
 Forwarding class Loss priority Code point
 gold low 000111
 gold high 001010
 silver low 110000
 silver high 111000
 bronze high 001100
 lead low 101110
 lead high 110111

```



# Operational Commands (Scheduling)

- [Monitoring CoS Scheduler Maps on page 7291](#)
- [show class-of-service drop-profile](#)
- [show class-of-service forwarding-table](#)
- [show class-of-service forwarding-table drop-profile](#)
- [show class-of-service forwarding-table scheduler-map](#)
- [show class-of-service interface](#)
- [show class-of-service scheduler-map](#)
- [show class-of-service traffic-control-profile](#)
- [show interfaces queue](#)
- [show interfaces voq](#)

## Monitoring CoS Scheduler Maps

**Purpose** Use the monitoring functionality to display assignments of CoS forwarding classes to schedulers.

**Action** To monitor CoS scheduler maps in the CLI, enter the CLI command:

```
user@switch> show class-of-service scheduler-map
```

To monitor a specific scheduler map in the CLI, enter the CLI command:

```
user@switch> show class-of-service scheduler-map scheduler-map-name
```

**Meaning** [Table 632](#) summarizes key output fields for CoS scheduler maps.

**Table 632: Summary of Key CoS Scheduler Maps Output Fields**

| Field         | Values                                                                   |
|---------------|--------------------------------------------------------------------------|
| Scheduler map | Name of a scheduler map that maps forwarding classes to schedulers.      |
| Index         | Index of a specific object—scheduler maps, schedulers, or drop profiles. |

**Table 632: Summary of Key CoS Scheduler Maps Output Fields** (*continued*)

| Field            | Values                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Scheduler        | Name of a scheduler that controls queue properties such as bandwidth and scheduling priority.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Forwarding class | Name(s) of the forwarding class(es) to which the scheduler is mapped.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Transmit rate    | Guaranteed minimum bandwidth configured on the queue mapped to the scheduler. On strict-high priority queues on QFX10000 switches, defines the maximum amount of traffic on the queue that is treated as strict-high priority traffic.                                                                                                                                                                                                                                                                                                                                                                              |
| Priority         | <p>Scheduling priority of traffic on a queue:</p> <ul style="list-style-type: none"> <li>• <b>strict-high</b>—Packets on a strict-high priority queue are transmitted first, before all other traffic, up to the configured maximum bandwidth (shaping rate). On QFX3500, QFX3600, EX4600, and OCX series switches, and on QFabric system, only one queue can be configured as <b>strict-high</b> priority. On QFX10000 switches, you can configure more than one strict-high priority queue.</li> <li>• <b>low</b>—Packets in this queue are transmitted after packets in the <b>strict-high</b> queue.</li> </ul> |
| Drop Profiles    | Name and index of a drop profile that is mapped to a specific loss priority and protocol pair. The drop profile determines the way best effort queues drop packets during periods of congestion.                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Loss Priority    | Packet loss priority mapped to the drop profile. You can configure different drop profiles for <b>low</b> , <b>medium-high</b> , and <b>high</b> loss priority traffic.                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Protocol         | Transport protocol of the drop profile for the particular priority.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Name             | Name of the drop profile.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## show class-of-service drop-profile

|                                 |                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show class-of-service drop-profile<br><profile-name <i>profile-name</i> >                                                                                                                                                                                          |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.<br>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| <b>Description</b>              | Display data points for each class-of-service (CoS) random early detection (RED) drop profile.                                                                                                                                                                     |
| <b>Options</b>                  | <b>none</b> —Display all drop profiles.<br><br><b>profile-name <i>profile-name</i></b> —(Optional) Display the specified profile only.                                                                                                                             |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                               |
| <b>List of Sample Output</b>    | <a href="#">show class-of-service drop-profile on page 7294</a><br><a href="#">show class-of-service drop-profile (EX4200 Switch) on page 7294</a><br><a href="#">show class-of-service drop-profile (EX8200 Switch) on page 7294</a>                              |
| <b>Output Fields</b>            | <a href="#">Table 633</a> describes the output fields for the <b>show class-of-service drop-profile</b> command. Output fields are listed in the approximate order in which they appear.                                                                           |

**Table 633: show class-of-service drop-profile Output Fields**

| Field Name              | Field Description                                                                                                                                                                                                                     |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Drop profile</b>     | Name of a drop profile.                                                                                                                                                                                                               |
| <b>Type</b>             | Type of drop profile: <ul style="list-style-type: none"> <li>• <b>discrete</b> (default)</li> <li>• <b>interpolated</b> (EX8200 switches, QFX Series switches, QFabric systems, EX4600 switches, OCX Series switches only)</li> </ul> |
| <b>Index</b>            | Internal index of this drop profile.                                                                                                                                                                                                  |
| <b>Fill Level</b>       | Percentage fullness of a queue.                                                                                                                                                                                                       |
| <b>Drop probability</b> | Drop probability at this fill level.                                                                                                                                                                                                  |

## Sample Output

### show class-of-service drop-profile

```
user@host> show class-of-service drop-profile
Drop profile: <default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
Drop profile: user-drop-profile, Type: interpolated, Index: 2989
 Fill level Drop probability
 0 0
 1 1
 2 2
 4 4
 5 5
 6 6
 8 8
 10 10
 12 15
 14 20
 15 23
... 64 entries total
 90 96
 92 96
 94 97
 95 98
 96 98
 98 99
 99 99
 100 100
```

### show class-of-service drop-profile (EX4200 Switch)

```
user@switch> show class-of-service drop-profile
Drop profile: <default-drop-profile>, Type: discrete, Index: 1
 Fill level
 100
Drop profile: dp1, Type: discrete, Index: 40496
 Fill level
 10
```

### show class-of-service drop-profile (EX8200 Switch)

```
user@switch> show class-of-service drop-profile
Drop profile: <default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
Drop profile: dp1, Type: interpolated, Index: 40496
 Fill level Drop probability
 0 0
 1 80
 2 90
 4 90
 5 90
 6 90
 8 90
 10 90
 12 91
 14 91
 15 91
 16 91
```

|                                                 |                  |
|-------------------------------------------------|------------------|
| 18                                              | 91               |
| 20                                              | 91               |
| 22                                              | 92               |
| 24                                              | 92               |
| 25                                              | 92               |
| 26                                              | 92               |
| 28                                              | 92               |
| 30                                              | 92               |
| 32                                              | 93               |
| 34                                              | 93               |
| 35                                              | 93               |
| 36                                              | 93               |
| 38                                              | 93               |
| 40                                              | 93               |
| 42                                              | 94               |
| 44                                              | 94               |
| 45                                              | 94               |
| 46                                              | 94               |
| 48                                              | 94               |
| 49                                              | 94               |
| 51                                              | 95               |
| 52                                              | 95               |
| 54                                              | 95               |
| 55                                              | 95               |
| 56                                              | 95               |
| 58                                              | 95               |
| 60                                              | 95               |
| 62                                              | 96               |
| 64                                              | 96               |
| 65                                              | 96               |
| 66                                              | 96               |
| 68                                              | 96               |
| 70                                              | 96               |
| 72                                              | 97               |
| 74                                              | 97               |
| 75                                              | 97               |
| 76                                              | 97               |
| 78                                              | 97               |
| 80                                              | 97               |
| 82                                              | 98               |
| 84                                              | 98               |
| 85                                              | 98               |
| 86                                              | 98               |
| 88                                              | 98               |
| 90                                              | 98               |
| 92                                              | 99               |
| 94                                              | 99               |
| 95                                              | 99               |
| 96                                              | 99               |
| 98                                              | 99               |
| 99                                              | 99               |
| 100                                             | 100              |
| Drop profile: dp2, Type: discrete, Index: 40499 |                  |
| Fill level                                      | Drop probability |
| 10                                              | 5                |
| 50                                              | 50               |

## show class-of-service forwarding-table

---

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-----------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 7296</a><br><a href="#">Syntax (TX Matrix and TX Matrix Plus Router) on page 7296</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Syntax</b>                                       | show class-of-service forwarding-table                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Syntax (TX Matrix and TX Matrix Plus Router)</b> | show class-of-service forwarding-table<br><lcc <i>number</i> >   <sfc <i>number</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>                          | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.<br>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>                                  | Display the entire class-of-service (CoS) configuration as it exists in the forwarding table. Executing this command is equivalent to executing all <b>show class-of-service forwarding-table</b> commands in succession.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                                      | <p><b>lcc <i>number</i></b>—(TX Matrix and TX Matrix Plus router only) (Optional) On a TX Matrix router, display the forwarding table configuration for a specific T640 router (or line-card chassis) configured in a routing matrix. On a TX Matrix Plus router, display the forwarding table configuration for a specific router (or line-card chassis) configured in the routing matrix.</p> <p>Replace <i>number</i> with the following values depending on the LCC configuration:</p> <ul style="list-style-type: none"><li>• 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.</li><li>• 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.</li><li>• 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.</li><li>• 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.</li></ul> <p><b>sfc <i>number</i></b>—(TX Matrix Plus routers only) (Optional) Display the forwarding table configuration for the TX Matrix Plus router. Replace <i>number</i> with 0.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>List of Sample Output</b>                        | <a href="#">show class-of-service forwarding-table on page 7297</a><br><a href="#">show class-of-service forwarding-table lcc (TX Matrix Plus Router) on page 7298</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Output Fields</b>                                | See the output field descriptions for <b>show class-of-service forwarding-table</b> commands: <ul style="list-style-type: none"><li>• <a href="#">show class-of-service forwarding-table classifier</a></li><li>• <a href="#">show class-of-service forwarding-table classifier mapping</a></li><li>• <a href="#">show class-of-service forwarding-table drop-profile</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |



- *show class-of-service forwarding-table fabric scheduler-map*
- *show class-of-service forwarding-table rewrite-rule*
- *show class-of-service forwarding-table rewrite-rule mapping*
- *show class-of-service forwarding-table scheduler-map*

## Sample Output

### show class-of-service forwarding-table

```

user@host> show class-of-service forwarding-table
Classifier table index: 9, # entries: 8, Table type: EXP
Entry # Code point Forwarding-class # PLP
0 000 0 0
1 001 0 1
2 010 1 0
3 011 1 1
4 100 2 0
5 101 2 1
6 110 3 0
7 111 3 1

Interface Index Table Index/ Q num Table type
sp-0/0/0.1001 66 11 11 IPv4 precedence
sp-0/0/0.2001 67 11 11 IPv4 precedence
sp-0/0/0.16383 68 11 11 IPv4 precedence
fe-0/0/0.0 69 11 11 IPv4 precedence

Interface: sp-0/0/0 (Index: 129, Map index: 2, Map type: FINAL,
Num of queues: 2):
 Entry 0 (Scheduler index: 16, Forwarding-class #: 0):
 Tx rate: 0 Kb (95%), Buffer size: 95 percent
 Priority low
 PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1
 Entry 1 (Scheduler index: 18, Forwarding-class #: 3):
 Tx rate: 0 Kb (5%), Buffer size: 5 percent
 Priority low
 PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1

Interface: fe-0/0/0 (Index: 137, Map index: 2, Map type: FINAL,
Num of queues: 2):
 Entry 0 (Scheduler index: 16, Forwarding-class #: 0):
 Tx rate: 0 Kb (95%), Buffer size: 95 percent
 Priority low
 PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1
 Entry 1 (Scheduler index: 18, Forwarding-class #: 3):
 Tx rate: 0 Kb (5%), Buffer size: 5 percent
 Priority low
 PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1

Interface: fe-0/0/1 (Index: 138, Map index: 2, Map type: FINAL,
Num of queues: 2):
 Entry 0 (Scheduler index: 16, Forwarding-class #: 0):
 Tx rate: 0 Kb (95%), Buffer size: 95 percent
 Priority low
 PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1
 Entry 1 (Scheduler index: 18, Forwarding-class #: 3):
 Tx rate: 0 Kb (5%), Buffer size: 5 percent
 Priority low

```

PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1

...

RED drop profile index: 1, # entries: 1

| Entry | Fullness(%) | Drop<br>Probability(%) |
|-------|-------------|------------------------|
| 0     | 100         | 100                    |

### show class-of-service forwarding-table lcc (TX Matrix Plus Router)

user@host> show class-of-service forwarding-table lcc 0  
lcc0-re0:

-----

Classifier table index: 9, # entries: 64, Table type: IPv6 DSCP

| Entry # | Code point | Forwarding-class # | PLP |
|---------|------------|--------------------|-----|
| 0       | 000000     | 0                  | 0   |
| 1       | 000001     | 0                  | 0   |
| 2       | 000010     | 0                  | 0   |
| 3       | 000011     | 0                  | 0   |
| 4       | 000100     | 0                  | 0   |
| 5       | 000101     | 0                  | 0   |
| 6       | 000110     | 0                  | 0   |
| 7       | 000111     | 0                  | 0   |
| 8       | 001000     | 0                  | 0   |
| 9       | 001001     | 0                  | 0   |
| 10      | 001010     | 0                  | 0   |
| 11      | 001011     | 0                  | 0   |
| 12      | 001100     | 0                  | 0   |
| 13      | 001101     | 0                  | 0   |
| 14      | 001110     | 0                  | 0   |
| 15      | 001111     | 0                  | 0   |
| 16      | 010000     | 0                  | 0   |
| 17      | 010001     | 0                  | 0   |
| 18      | 010010     | 0                  | 0   |
| 19      | 010011     | 0                  | 0   |
| 20      | 010100     | 0                  | 0   |
| 21      | 010101     | 0                  | 0   |
| 22      | 010110     | 0                  | 0   |
| 23      | 010111     | 0                  | 0   |
| 24      | 011000     | 0                  | 0   |
| 25      | 011001     | 0                  | 0   |
| 26      | 011010     | 0                  | 0   |
| 27      | 011011     | 0                  | 0   |
| 28      | 011100     | 0                  | 0   |
| 29      | 011101     | 0                  | 0   |
| 30      | 011110     | 0                  | 0   |
| 31      | 011111     | 0                  | 0   |
| 32      | 100000     | 0                  | 0   |
| 33      | 100001     | 0                  | 0   |
| 34      | 100010     | 0                  | 0   |
| 35      | 100011     | 0                  | 0   |
| 36      | 100100     | 0                  | 0   |
| 37      | 100101     | 0                  | 0   |
| 38      | 100110     | 0                  | 0   |
| 39      | 100111     | 0                  | 0   |
| 40      | 101000     | 0                  | 0   |
| 41      | 101001     | 0                  | 0   |
| 42      | 101010     | 0                  | 0   |
| 43      | 101011     | 0                  | 0   |

|     |        |   |   |
|-----|--------|---|---|
| 44  | 101100 | 0 | 0 |
| 45  | 101101 | 0 | 0 |
| 46  | 101110 | 0 | 0 |
| ... |        |   |   |

## show class-of-service forwarding-table drop-profile

|                                 |                                                                                                                                                                                                           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show class-of-service forwarding-table drop-profile                                                                                                                                                       |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.<br>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.              |
| <b>Description</b>              | Display the data points of all random early detection (RED) drop profiles as they exist in the forwarding table.                                                                                          |
| <b>Options</b>                  | This command has no options.                                                                                                                                                                              |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                      |
| <b>List of Sample Output</b>    | <a href="#">show class-of-service forwarding-table drop-profile on page 7300</a>                                                                                                                          |
| <b>Output Fields</b>            | <a href="#">Table 634</a> describes the output fields for the <b>show class-of-service forwarding-table drop-profile</b> command. Output fields are listed in the approximate order in which they appear. |

**Table 634: show class-of-service forwarding-table drop-profile Output Fields**

| Field Name             | Field Description                                         |
|------------------------|-----------------------------------------------------------|
| RED drop profile index | Index of this drop profile.                               |
| # entries              | Number of entries in a particular RED drop profile index. |
| Entry                  | Drop profile entry number.                                |
| Fullness(%)            | Percentage fullness of a queue.                           |
| Drop probability(%)    | Drop probability at this fill level.                      |

## Sample Output

### show class-of-service forwarding-table drop-profile

```

user@host> show class-of-service forwarding-table drop-profile
RED drop profile index: 4, # entries: 1
 Drop
Entry Fullness(%) Probability(%)
 0 100 100

RED drop profile index: 8742, # entries: 3
 Drop
Entry Fullness(%) Probability(%)
 0 10 10
 1 20 20
 2 30 30

```

RED drop profile index: 24627, # entries: 64

| Drop  |             |                |
|-------|-------------|----------------|
| Entry | Fullness(%) | Probability(%) |
| 0     | 0           | 0              |
| 1     | 1           | 1              |
| 2     | 2           | 2              |
| 3     | 4           | 4              |
| ...   |             |                |
| 61    | 98          | 99             |
| 62    | 99          | 99             |
| 63    | 100         | 100            |

RED drop profile index: 25393, # entries: 64

| Drop  |             |                |
|-------|-------------|----------------|
| Entry | Fullness(%) | Probability(%) |
| 0     | 0           | 0              |
| 1     | 1           | 1              |
| 2     | 2           | 2              |
| 3     | 4           | 4              |
| ...   |             |                |
| 61    | 98          | 98             |
| 62    | 99          | 99             |
| 63    | 100         | 100            |

## show class-of-service forwarding-table scheduler-map

|                                 |                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show class-of-service forwarding-table scheduler-map                                                                                                                                                       |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.<br>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.               |
| <b>Description</b>              | For each physical interface, display the scheduler map information as it exists in the forwarding table.                                                                                                   |
| <b>Options</b>                  | This command has no options.                                                                                                                                                                               |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                       |
| <b>List of Sample Output</b>    | <a href="#">show class-of-service forwarding-table scheduler-map on page 7303</a>                                                                                                                          |
| <b>Output Fields</b>            | <a href="#">Table 635</a> describes the output fields for the <b>show class-of-service forwarding-table scheduler-map</b> command. Output fields are listed in the approximate order in which they appear. |

**Table 635: show class-of-service forwarding-table scheduler-map Output Fields**

| Field Name         | Field Description                                                                                                                                                                                                                                                                           |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Interface          | Name of the physical interface.                                                                                                                                                                                                                                                             |
| Index              | Physical interface index.                                                                                                                                                                                                                                                                   |
| Map index          | Scheduler map index.                                                                                                                                                                                                                                                                        |
| Num of queues      | Number of queues defined in this scheduler map.                                                                                                                                                                                                                                             |
| Entry              | Number of this entry in the scheduler map.                                                                                                                                                                                                                                                  |
| Scheduler index    | Scheduler policy index.                                                                                                                                                                                                                                                                     |
| Forwarding-class # | Forwarding class number to which this entry is applied.                                                                                                                                                                                                                                     |
| Tx rate            | Configured transmit rate of the scheduler (in bps). The rate is a percentage of the total interface bandwidth, or the keyword <b>remainder</b> , which indicates that the scheduler receives the remaining bandwidth of the interface.                                                      |
| Max buffer delay   | Amount of transmit delay (in milliseconds) or buffer size of the queue. This amount is a percentage of the total interface buffer allocation or the keyword <b>remainder</b> , which indicates that the buffer is sized according to what remains after other scheduler buffer allocations. |
| Priority           | <ul style="list-style-type: none"> <li><b>high</b>—Queue priority is high.</li> <li><b>low</b>—Queue priority is low.</li> </ul>                                                                                                                                                            |

Table 635: show class-of-service forwarding-table scheduler-map Output Fields (*continued*)

| Field Name      | Field Description                                                                                           |
|-----------------|-------------------------------------------------------------------------------------------------------------|
| PLP high        | Drop profile index for a high packet loss priority profile.                                                 |
| PLP low         | Drop profile index for a low packet loss priority profile.                                                  |
| PLP medium-high | Drop profile index for a medium-high packet loss priority profile.                                          |
| PLP medium-low  | Drop profile index for a medium-low packet loss priority profile.                                           |
| TCP PLP high    | Drop profile index for a high TCP packet loss priority profile.                                             |
| TCP PLP low     | Drop profile index for a low TCP packet loss priority profile.                                              |
| Policy is exact | If this line appears in the output, exact rate limiting is enabled. Otherwise, no rate limiting is enabled. |

## Sample Output

### show class-of-service forwarding-table scheduler-map

```

user@host> show class-of-service forwarding-table scheduler-map
Interface: so-5/0/0 (Index: 9, Map index: 17638, Num of queues: 2):
 Entry 0 (Scheduler index: 6090, Forwarding-class #: 0):
 Tx rate: 0 Kb (30%), Max buffer delay: 39 bytes (0%)
 Priority low
 PLP high: 25393, PLP low: 24627, TCP PLP high: 25393, TCP PLP low: 8742
 Policy is exact
 Entry 1 (Scheduler index: 38372, Forwarding-class #: 1):
 Traffic chunk: Max = 0 bytes, Min = 0 bytes
 Tx rate: 0 Kb (40%), Max buffer delay: 68 bytes (0%)
 Priority high
 PLP high: 25393, PLP low: 24627, TCP PLP high: 25393, TCP PLP low: 8742

Interface: at-6/1/0 (Index: 10, Map index: 17638, Num of queues: 2):
 Entry 0 (Scheduler index: 6090, Forwarding-class #: 0):
 Traffic chunk: Max = 0 bytes, Min = 0 bytes
 Tx rate: 0 Kb (30%), Max buffer delay: 39 bytes (0%)
 Priority high
 PLP high: 25393, PLP low: 24627, TCP PLP high: 25393, TCP PLP low: 8742
 Entry 1 (Scheduler index: 38372, Forwarding-class #: 1):
 Traffic chunk: Max = 0 bytes, Min = 0 bytes
 Tx rate: 0 Kb (40%), Max buffer delay: 68 bytes (0%)
 Priority low
 PLP high: 25393, PLP low: 24627, TCP PLP high: 25393, TCP PLP low: 8742

```

## show class-of-service interface

---

**Syntax**    `show class-of-service interface`  
              `<comprehensive | detail> <interface-name>`

**Release Information**    Command introduced before Junos OS Release 7.4.  
                              Command introduced in Junos OS Release 9.0 for EX Series switches.  
                              Forwarding class map information added in Junos OS Release 9.4.  
                              Command introduced in Junos OS Release 11.1 for the QFX Series.  
                              Command introduced in Junos OS Release 12.1 for the PTX Series Packet Transport Switches.  
                              Command introduced in Junos OS Release 12.2 for the ACX Series Universal Access routers.  
                              Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.  
                              Options **detail** and **comprehensive** introduced in Junos OS Release 11.4.  
                              Command introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.

**Description**    Display the logical and physical interface associations for the classifier, rewrite rules, and scheduler map objects.



**NOTE:** On routing platforms with dual Routing Engines, running this command on the backup Routing Engine, with or without any of the available options, is not supported and produces the following error message:

**error: the class-of-service subsystem is not running**

**Options**    **none**—Display CoS associations for all physical and logical interfaces.

**comprehensive**—(M Series, MX Series, and T Series routers) (Optional) Display comprehensive quality-of-service (QoS) information about all physical and logical interfaces.

**detail**—(M Series, MX Series, and T Series routers) (Optional) Display QoS and CoS information based on the interface.

If the **interface** *interface-name* is a physical interface, the output includes:

- Brief QoS information about the physical interface
- Brief QoS information about the logical interface
- CoS information about the physical interface
- Brief information about filters or policers of the logical interface
- Brief CoS information about the logical interface

If the **interface** *interface-name* is a logical interface, the output includes:

- Brief QoS information about the logical interface



- Information about filters or policers for the logical interface
- CoS information about the logical interface

**interface-name**—(Optional) Display class-of-service (CoS) associations for the specified interface.

**none**—Display CoS associations for all physical and logical interfaces.

**Required Privilege Level** view

**Related Documentation**

- *Verifying and Managing Junos OS Enhanced Subscriber Management*

**List of Sample Output** [show class-of-service interface \(Physical\) on page 7316](#)  
[show class-of-service interface \(Logical\) on page 7317](#)  
[show class-of-service interface \(Gigabit Ethernet\) on page 7317](#)  
[show class-of-service interface \(ANCP\) on page 7317](#)  
[show class-of-service interface \(PPPoE Interface\) on page 7317](#)  
[show class-of-service interface \(T4000 Routers with Type 5 FPCs\) on page 7317](#)  
[show class-of-service interface detail on page 7318](#)  
[show class-of-service interface comprehensive on page 7318](#)  
[show class-of-service interface \(ACX Series Routers\) on page 7329](#)  
[show class-of-service interface \(PPPoE Subscriber Interface for Enhanced Subscriber Management\) on page 7332](#)

**Output Fields** [Table 609](#) describes the output fields for the **show class-of-service interface** command. Output fields are listed in the approximate order in which they appear.

**Table 636: show class-of-service interface Output Fields**

| Field Name         | Field Description                                                                                                                                                                                                                                                                                          |                                                      |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------|
| Physical interface | Name of a physical interface.                                                                                                                                                                                                                                                                              |                                                      |
| Index              | Index of this interface or the internal index of this object.<br><br>(Enhanced subscriber management for MX Series routers) Index values for dynamic CoS traffic control profiles and dynamic scheduler maps are larger for enhanced subscriber management than they are for legacy subscriber management. |                                                      |
| Dedicated Queues   | Status of dedicated queues configured on an interface. Supported only on Trio MPC/MIC interfaces on MX Series routers.                                                                                                                                                                                     | Number of queues you can configure on the interface. |
| Queues supported   | Number of queues you can configure on the interface.                                                                                                                                                                                                                                                       |                                                      |
| Queues in use      | Number of queues currently configured.                                                                                                                                                                                                                                                                     |                                                      |

Table 636: show class-of-service interface Output Fields (*continued*)

| Field Name                                 | Field Description                                                                                                                                                                                                                                                                                     |
|--------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Total non-default queues created</b>    | Number of queues created in addition to the default queues. Supported only on Trio MPC/MIC interfaces on MX Series routers.<br><br>(Enhanced subscriber management for MX Series routers) This field is not displayed for enhanced subscriber management.                                             |
| <b>Rewrite Input IEEE Code-point</b>       | (QFX Series only) IEEE 802.1p code point (priority) rewrite value. Incoming traffic from the Fibre Channel (FC) SAN is classified into the forwarding class specified in the native FC interface (NP_Port) fixed classifier and uses the priority specified as the IEEE 802.1p rewrite value.         |
| <b>Shaping rate</b>                        | Maximum transmission rate on the physical interface. You can configure the shaping rate on the physical interface, or on the logical interface, but not on both. Therefore, the <b>Shaping rate</b> field is displayed for either the physical interface or the logical interface.                    |
| <b>Scheduler map</b>                       | Name of the output scheduler map associated with this interface.<br><br>(Enhanced subscriber management for MX Series routers) The name of the dynamic scheduler map object is associated with a generated UID (for example, <b>SMAP-1_UID1002</b> ) instead of with a subscriber interface.          |
| <b>Scheduler map forwarding class sets</b> | (QFX Series only) Name of the output fabric scheduler map associated with a QFabric system Interconnect device interface.                                                                                                                                                                             |
| <b>Input shaping rate</b>                  | For Gigabit Ethernet IQ2 PICs, maximum transmission rate on the input interface.                                                                                                                                                                                                                      |
| <b>Input scheduler map</b>                 | For Gigabit Ethernet IQ2 PICs, name of the input scheduler map associated with this interface.                                                                                                                                                                                                        |
| <b>Chassis scheduler map</b>               | Name of the scheduler map associated with the packet forwarding component queues.                                                                                                                                                                                                                     |
| <b>Rewrite</b>                             | Name and type of the rewrite rules associated with this interface.                                                                                                                                                                                                                                    |
| <b>Traffic-control-profile</b>             | Name of the associated traffic control profile.<br><br>(Enhanced subscriber management for MX Series routers) The name of the dynamic traffic control profile object is associated with a generated UID (for example, <b>TC_PROF_100_199_SERIES_UID1006</b> ) instead of with a subscriber interface. |
| <b>Classifier</b>                          | Name and type of classifiers associated with this interface.                                                                                                                                                                                                                                          |
| <b>Forwarding-class-map</b>                | Name of the forwarding map associated with this interface.                                                                                                                                                                                                                                            |
| <b>Congestion-notification</b>             | (QFX Series and EX4600 switches only) Congestion notification state, <b>enabled</b> or <b>disabled</b> .                                                                                                                                                                                              |
| <b>Logical interface</b>                   | Name of a logical interface.                                                                                                                                                                                                                                                                          |
| <b>Object</b>                              | Category of an object: <b>Classifier</b> , <b>Fragmentation-map</b> (for LSQ interfaces only), <b>Scheduler-map</b> , <b>Rewrite</b> , <b>Translation Table</b> (for IQE PICs only), or <b>traffic-class-map</b> (for T4000 routers with Type 5 FPCs).                                                |
| <b>Name</b>                                | Name of an object.                                                                                                                                                                                                                                                                                    |

Table 636: show class-of-service interface Output Fields (*continued*)

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Type</b>             | Type of an object: <b>dscp</b> , <b>dscp-ipv6</b> , <b>exp</b> , <b>ieee-802.1</b> , <b>ip</b> , <b>inet-precedence</b> , or <b>ieee-802.1ad</b> (for traffic class map on T4000 routers with Type 5 FPCs)..                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Link-level type</b>  | Encapsulation on the physical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>MTU</b>              | MTU size on the physical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Speed</b>            | Speed at which the interface is running.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Loopback</b>         | Whether loopback is enabled and the type of loopback.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Source filtering</b> | Whether source filtering is enabled or disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Flow control</b>     | Whether flow control is enabled or disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Auto-negotiation</b> | (Gigabit Ethernet interfaces) Whether autonegotiation is enabled or disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Remote-fault</b>     | (Gigabit Ethernet interfaces) Remote fault status. <ul style="list-style-type: none"> <li>• <b>Online</b>—Autonegotiation is manually configured as online.</li> <li>• <b>Offline</b>—Autonegotiation is manually configured as offline.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Device flags</b>     | The <b>Device flags</b> field provides information about the physical device and displays one or more of the following values: <ul style="list-style-type: none"> <li>• <b>Down</b>—Device has been administratively disabled.</li> <li>• <b>Hear-Own-Xmit</b>—Device receives its own transmissions.</li> <li>• <b>Link-Layer-Down</b>—The link-layer protocol has failed to connect with the remote endpoint.</li> <li>• <b>Loopback</b>—Device is in physical loopback.</li> <li>• <b>Loop-Detected</b>—The link layer has received frames that it sent, thereby detecting a physical loopback.</li> <li>• <b>No-Carrier</b>—On media that support carrier recognition, no carrier is currently detected.</li> <li>• <b>No-Multicast</b>—Device does not support multicast traffic.</li> <li>• <b>Present</b>—Device is physically present and recognized.</li> <li>• <b>Promiscuous</b>—Device is in promiscuous mode and recognizes frames addressed to all physical addresses on the media.</li> <li>• <b>Quench</b>—Transmission on the device is quenched because the output buffer is overflowing.</li> <li>• <b>Recv-All-Multicasts</b>—Device is in multicast promiscuous mode and therefore provides no multicast filtering.</li> <li>• <b>Running</b>—Device is active and enabled.</li> </ul> |

Table 636: show class-of-service interface Output Fields (*continued*)

| Field Name             | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Interface flags</b> | <p>The <b>Interface flags</b> field provides information about the physical interface and displays one or more of the following values:</p> <ul style="list-style-type: none"> <li>• <b>Admin-Test</b>—Interface is in test mode and some sanity checking, such as loop detection, is disabled.</li> <li>• <b>Disabled</b>—Interface is administratively disabled.</li> <li>• <b>Down</b>—A hardware failure has occurred.</li> <li>• <b>Hardware-Down</b>—Interface is nonfunctional or incorrectly connected.</li> <li>• <b>Link-Layer-Down</b>—Interface keepalives have indicated that the link is incomplete.</li> <li>• <b>No-Multicast</b>—Interface does not support multicast traffic.</li> <li>• <b>No-receive No-transmit</b>—Passive monitor mode is configured on the interface.</li> <li>• <b>Point-To-Point</b>—Interface is point-to-point.</li> <li>• <b>Pop all MPLS labels from packets of depth</b>—MPLS labels are removed as packets arrive on an interface that has the <b>pop-all-labels</b> statement configured. The depth value can be one of the following: <ul style="list-style-type: none"> <li>• <b>1</b>—Takes effect for incoming packets with one label only.</li> <li>• <b>2</b>—Takes effect for incoming packets with two labels only.</li> <li>• <b>[ 1 2 ]</b>—Takes effect for incoming packets with either one or two labels.</li> </ul> </li> <li>• <b>Promiscuous</b>—Interface is in promiscuous mode and recognizes frames addressed to all physical addresses.</li> <li>• <b>Recv-All-Multicasts</b>—Interface is in multicast promiscuous mode and provides no multicast filtering.</li> <li>• <b>SNMP-Traps</b>—SNMP trap notifications are enabled.</li> <li>• <b>Up</b>—Interface is enabled and operational.</li> </ul> |
| <b>Flags</b>           | <p>The <b>Logical interface flags</b> field provides information about the logical interface and displays one or more of the following values:</p> <ul style="list-style-type: none"> <li>• <b>ACFC Encapsulation</b>—Address control field Compression (ACFC) encapsulation is enabled (negotiated successfully with a peer).</li> <li>• <b>Device-down</b>—Device has been administratively disabled.</li> <li>• <b>Disabled</b>—Interface is administratively disabled.</li> <li>• <b>Down</b>—A hardware failure has occurred.</li> <li>• <b>Clear-DF-Bit</b>—GRE tunnel or IPsec tunnel is configured to clear the Don't Fragment (DF) bit.</li> <li>• <b>Hardware-Down</b>—Interface protocol initialization failed to complete successfully.</li> <li>• <b>PFC</b>—Protocol field compression is enabled for the PPP session.</li> <li>• <b>Point-To-Point</b>—Interface is point-to-point.</li> <li>• <b>SNMP-Traps</b>—SNMP trap notifications are enabled.</li> <li>• <b>Up</b>—Interface is enabled and operational.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Encapsulation</b>   | Encapsulation on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Admin</b>           | Administrative state of the interface ( <b>Up</b> or <b>Down</b> ).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Link</b>            | Status of physical link ( <b>Up</b> or <b>Down</b> ).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Proto</b>           | Protocol configured on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

Table 636: show class-of-service interface Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Input Filter</b>            | Names of any firewall filters to be evaluated when packets are received on the interface, including any filters attached through activation of dynamic service.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Output Filter</b>           | Names of any firewall filters to be evaluated when packets are transmitted on the interface, including any filters attached through activation of dynamic service.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Link flags</b>              | Provides information about the physical link and displays one or more of the following values: <ul style="list-style-type: none"> <li>• <b>ACFC</b>—Address control field compression is configured. The Point-to-Point Protocol (PPP) session negotiates the ACFC option.</li> <li>• <b>Give-Up</b>—Link protocol does not continue connection attempts after repeated failures.</li> <li>• <b>Loose-LCP</b>—PPP does not use the Link Control Protocol (LCP) to indicate whether the link protocol is operational.</li> <li>• <b>Loose-LMI</b>—Frame Relay does not use the Local Management Interface (LMI) to indicate whether the link protocol is operational.</li> <li>• <b>Loose-NCP</b>—PPP does not use the Network Control Protocol (NCP) to indicate whether the device is operational.</li> <li>• <b>Keepalives</b>—Link protocol keepalives are enabled.</li> <li>• <b>No-Keepalives</b>—Link protocol keepalives are disabled.</li> <li>• <b>PFC</b>—Protocol field compression is configured. The PPP session negotiates the PFC option.</li> </ul> |
| <b>Hold-times</b>              | Current interface hold-time up and hold-time down, in milliseconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>CoS queues</b>              | Number of CoS queues configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Last flapped</b>            | Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago)</b> . For example, <b>Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Statistics last cleared</b> | Number and rate of bytes and packets received and transmitted on the physical interface. <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>IPv6 transit statistics</b> | Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

Table 636: show class-of-service interface Output Fields (*continued*)

| Field Name           | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Input errors</b>  | <p>Input errors on the interface. The labels are explained in the following list:</p> <ul style="list-style-type: none"> <li>• <b>Errors</b>—Sum of the incoming frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li>• <b>Runts</b>—Number of frames received that are smaller than the runt threshold.</li> <li>• <b>Giants</b>—Number of frames received that are larger than the giant threshold.</li> <li>• <b>Bucket Drops</b>—Drops resulting from the traffic load exceeding the interface transmit or receive leaky bucket configuration.</li> <li>• <b>Policed discards</b>—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle.</li> <li>• <b>L3 incompletes</b>—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. Layer 3 incomplete errors can be ignored by configuring the <b>ignore-l3-incompletes</b> statement.</li> <li>• <b>L2 channel errors</b>—Number of times the software did not find a valid logical interface for an incoming frame.</li> <li>• <b>L2 mismatch timeouts</b>—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable.</li> <li>• <b>HS link CRC errors</b>—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces.</li> <li>• <b>HS link FIFO overflows</b>—Number of FIFO overflows on the high-speed links between the ASICs responsible for handling the router interfaces.</li> </ul> |
| <b>Output errors</b> | <p>Output errors on the interface. The labels are explained in the following list:</p> <ul style="list-style-type: none"> <li>• <b>Carrier transitions</b>—Number of times the interface has gone from <b>down</b> to <b>up</b>. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC is malfunctioning.</li> <li>• <b>Errors</b>—Sum of the outgoing frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> </ul> <p><b>NOTE:</b> Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the <b>Drops</b> field does not always use the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p> <ul style="list-style-type: none"> <li>• <b>Aged packets</b>—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware.</li> <li>• <b>HS link FIFO underflows</b>—Number of FIFO underflows on the high-speed links between the ASICs responsible for handling the router interfaces.</li> <li>• <b>MTU errors</b>—Number of packets whose size exceeds the MTU of the interface.</li> </ul>                                                                                                                                                                                                                                                                                                                |
| <b>Egress queues</b> | Total number of egress queues supported on the specified interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

Table 636: show class-of-service interface Output Fields (*continued*)

| Field Name                                  | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Queue counters</b>                       | <p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul> <p><b>NOTE:</b> Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the <b>Dropped packets</b> field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p>                                                                                                                                                                                                                                   |
| <b>SONET alarms</b><br><b>SONET defects</b> | <p>(SONET) SONET media-specific alarms and defects that prevent the interface from passing packets. When a defect persists for a certain period, it is promoted to an alarm. Based on the router configuration, an alarm can ring the red or yellow alarm bell on the router or light the red or yellow alarm LED on the craft interface. See these fields for possible alarms and defects: <b>SONET PHY</b>, <b>SONET section</b>, <b>SONET line</b>, and <b>SONET path</b>.</p>                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>SONET PHY</b>                            | <p>Counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> <li>• <b>Seconds</b>—Number of seconds the defect has been active.</li> <li>• <b>Count</b>—Number of times that the defect has gone from inactive to active.</li> <li>• <b>State</b>—State of the error. A state other than <b>OK</b> indicates a problem.</li> </ul> <p>The <b>SONET PHY</b> field has the following subfields:</p> <ul style="list-style-type: none"> <li>• <b>PLL Lock</b>—Phase-locked loop</li> <li>• <b>PHY Light</b>—Loss of optical signal</li> </ul>                                                                                                                                                                                                                                                                                        |
| <b>SONET section</b>                        | <p>Counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> <li>• <b>Seconds</b>—Number of seconds the defect has been active.</li> <li>• <b>Count</b>—Number of times that the defect has gone from inactive to active.</li> <li>• <b>State</b>—State of the error. A state other than <b>OK</b> indicates a problem.</li> </ul> <p>The <b>SONET section</b> field has the following subfields:</p> <ul style="list-style-type: none"> <li>• <b>BIP-BI</b>—Bit interleaved parity for SONET section overhead</li> <li>• <b>SEF</b>—Severely errored framing</li> <li>• <b>LOS</b>—Loss of signal</li> <li>• <b>LOF</b>—Loss of frame</li> <li>• <b>ES-S</b>—Errored seconds (section)</li> <li>• <b>SES-S</b>—Severely errored seconds (section)</li> <li>• <b>SEFS-S</b>—Severely errored framing seconds (section)</li> </ul> |

Table 636: show class-of-service interface Output Fields (*continued*)

| Field Name        | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>SONET line</b> | <p>Active alarms and defects, plus counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> <li>• <b>Seconds</b>—Number of seconds the defect has been active.</li> <li>• <b>Count</b>—Number of times that the defect has gone from inactive to active.</li> <li>• <b>State</b>—State of the error. A state other than <b>OK</b> indicates a problem.</li> </ul> <p>The <b>SONET line</b> field has the following subfields:</p> <ul style="list-style-type: none"> <li>• <b>BIP-B2</b>—Bit interleaved parity for SONET line overhead</li> <li>• <b>REI-L</b>—Remote error indication (near-end line)</li> <li>• <b>RDI-L</b>—Remote defect indication (near-end line)</li> <li>• <b>AIS-L</b>—Alarm indication signal (near-end line)</li> <li>• <b>BERR-SF</b>—Bit error rate fault (signal failure)</li> <li>• <b>BERR-SD</b>—Bit error rate defect (signal degradation)</li> <li>• <b>ES-L</b>—Errored seconds (near-end line)</li> <li>• <b>SES-L</b>—Severely errored seconds (near-end line)</li> <li>• <b>UAS-L</b>—Unavailable seconds (near-end line)</li> <li>• <b>ES-LFE</b>—Errored seconds (far-end line)</li> <li>• <b>SES-LFE</b>—Severely errored seconds (far-end line)</li> <li>• <b>UAS-LFE</b>—Unavailable seconds (far-end line)</li> </ul>      |
| <b>SONET path</b> | <p>Active alarms and defects, plus counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> <li>• <b>Seconds</b>—Number of seconds the defect has been active.</li> <li>• <b>Count</b>—Number of times that the defect has gone from inactive to active.</li> <li>• <b>State</b>—State of the error. A state other than <b>OK</b> indicates a problem.</li> </ul> <p>The <b>SONET path</b> field has the following subfields:</p> <ul style="list-style-type: none"> <li>• <b>BIP-B3</b>—Bit interleaved parity for SONET section overhead</li> <li>• <b>REI-P</b>—Remote error indication</li> <li>• <b>LOP-P</b>—Loss of pointer (path)</li> <li>• <b>AIS-P</b>—Path alarm indication signal</li> <li>• <b>RDI-P</b>—Path remote defect indication</li> <li>• <b>UNEQ-P</b>—Path unequipped</li> <li>• <b>PLM-P</b>—Path payload (signal) label mismatch</li> <li>• <b>ES-P</b>—Errored seconds (near-end STS path)</li> <li>• <b>SES-P</b>—Severely errored seconds (near-end STS path)</li> <li>• <b>UAS-P</b>—Unavailable seconds (near-end STS path)</li> <li>• <b>ES-PFE</b>—Errored seconds (far-end STS path)</li> <li>• <b>SES-PFE</b>—Severely errored seconds (far-end STS path)</li> <li>• <b>UAS-PFE</b>—Unavailable seconds (far-end STS path)</li> </ul> |



Table 636: show class-of-service interface Output Fields (*continued*)

| Field Name                             | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Received SONET overhead                | Values of the received and transmitted SONET overhead: <ul style="list-style-type: none"> <li>• <b>C2</b>—Signal label. Allocated to identify the construction and content of the STS-level SPE and for PDI-P.</li> <li>• <b>F1</b>—Section user channel byte. This byte is set aside for the purposes of users.</li> <li>• <b>K1 and K2</b>—These bytes are allocated for APS signaling for the protection of the multiplex section.</li> <li>• <b>J0</b>—Section trace. This byte is defined for STS-1 number 1 of an STS-N signal. Used to transmit a 1-byte fixed-length string or a 16-byte message so that a receiving terminal in a section can verify its continued connection to the intended transmitter.</li> <li>• <b>S1</b>—Synchronization status. The S1 byte is located in the first STS-1 number of an STS-N signal.</li> <li>• <b>Z3 and Z4</b>—Allocated for future use.</li> </ul>                                                                                                                                                                                  |
| Transmitted SONET overhead             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Received path trace                    | SONET/SDH interfaces allow path trace bytes to be sent inband across the SONET/SDH link. Juniper Networks and other router manufacturers use these bytes to help diagnose misconfigurations and network errors by setting the transmitted path trace message so that it contains the system hostname and name of the physical interface. The received path trace value is the message received from the router at the other end of the fiber. The transmitted path trace value is the message that this router transmits.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Transmitted path trace                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| HDLC configuration                     | Information about the HDLC configuration. <ul style="list-style-type: none"> <li>• <b>Policing bucket</b>—Configured state of the receiving policer.</li> <li>• <b>Shaping bucket</b>—Configured state of the transmitting shaper.</li> <li>• <b>Giant threshold</b>—Giant threshold programmed into the hardware.</li> <li>• <b>Runt threshold</b>—Runt threshold programmed into the hardware.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Packet Forwarding Engine configuration | Information about the configuration of the Packet Forwarding Engine: <ul style="list-style-type: none"> <li>• <b>Destination slot</b>—FPC slot number.</li> <li>• <b>PLP byte</b>—Packet Level Protocol byte.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| CoS information                        | Information about the CoS queue for the physical interface. <ul style="list-style-type: none"> <li>• <b>CoS transmit queue</b>—Queue number and its associated user-configured forwarding class name.</li> <li>• <b>Bandwidth %</b>—Percentage of bandwidth allocated to the queue.</li> <li>• <b>Bandwidth bps</b>—Bandwidth allocated to the queue (in bps).</li> <li>• <b>Buffer %</b>—Percentage of buffer space allocated to the queue.</li> <li>• <b>Buffer usec</b>—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time.</li> <li>• <b>Priority</b>—Queue priority: <b>low</b> or <b>high</b>.</li> <li>• <b>Limit</b>—Displayed if rate limiting is configured for the queue. Possible values are <b>none</b> and <b>exact</b>. If <b>exact</b> is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If <b>none</b> is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.</li> </ul> |
| Forwarding classes                     | Total number of forwarding classes supported on the specified interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Egress queues                          | Total number of egress queues supported on the specified interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

Table 636: show class-of-service interface Output Fields (*continued*)

| Field Name           | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Queue                | Queue number.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Forwarding classes   | Forwarding class name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Queued Packets       | Number of packets queued to this queue.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Queued Bytes         | Number of bytes queued to this queue. The byte counts vary by PIC type.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Transmitted Packets  | Number of packets transmitted by this queue. When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (displayed under the <b>Packet Forwarding Engine Chassis Queues</b> field) shows the prefragmentation values.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Transmitted Bytes    | Number of bytes transmitted by this queue. The byte counts vary by PIC type.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Tail-dropped packets | Number of packets dropped because of tail drop.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| RED-dropped packets  | <p>Number of packets dropped because of random early detection (RED).</p> <ul style="list-style-type: none"> <li>• (M Series and T Series routers only) On M320 and M120 routers and the T Series routers, the total number of dropped packets is displayed. On all other M Series routers, the output classifies dropped packets into the following categories: <ul style="list-style-type: none"> <li>• <b>Low, non-TCP</b>—Number of low-loss priority non-TCP packets dropped because of RED.</li> <li>• <b>Low, TCP</b>—Number of low-loss priority TCP packets dropped because of RED.</li> <li>• <b>High, non-TCP</b>—Number of high-loss priority non-TCP packets dropped because of RED.</li> <li>• <b>High, TCP</b>—Number of high-loss priority TCP packets dropped because of RED.</li> </ul> </li> <li>• (MX Series routers with enhanced DPCs, and T Series routers with enhanced FPCs only) The output classifies dropped packets into the following categories: <ul style="list-style-type: none"> <li>• <b>Low</b>—Number of low-loss priority packets dropped because of RED.</li> <li>• <b>Medium-low</b>—Number of medium-low loss priority packets dropped because of RED.</li> <li>• <b>Medium-high</b>—Number of medium-high loss priority packets dropped because of RED.</li> <li>• <b>High</b>—Number of high-loss priority packets dropped because of RED.</li> </ul> </li> </ul> <p><b>NOTE:</b> Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), this field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p> |

Table 636: show class-of-service interface Output Fields (*continued*)

| Field Name        | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| RED-dropped bytes | <p>Number of bytes dropped because of RED. The byte counts vary by PIC type.</p> <ul style="list-style-type: none"> <li>(M Series and T Series routers only) On M320 and M120 routers and the T Series routers, only the total number of dropped bytes is displayed. On all other M Series routers, the output classifies dropped bytes into the following categories: <ul style="list-style-type: none"> <li><b>Low, non-TCP</b>—Number of low-loss priority non-TCP bytes dropped because of RED.</li> <li><b>Low, TCP</b>—Number of low-loss priority TCP bytes dropped because of RED.</li> <li><b>High, non-TCP</b>—Number of high-loss priority non-TCP bytes dropped because of RED.</li> <li><b>High, TCP</b>—Number of high-loss priority TCP bytes dropped because of RED.</li> </ul> </li> </ul> <p><b>NOTE:</b> Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), this field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p> |
| Transmit rate     | Configured transmit rate of the scheduler. The rate is a percentage of the total interface bandwidth.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Rate Limit        | <p>Rate limiting configuration of the queue. Possible values are :</p> <ul style="list-style-type: none"> <li><b>None</b>—No rate limit.</li> <li><b>exact</b>—Queue transmits at the configured rate.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Buffer size       | Delay buffer size in the queue.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Priority          | Scheduling priority configured as <b>low</b> or <b>high</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Excess Priority   | Priority of the excess bandwidth traffic on a scheduler: <b>low</b> , <b>medium-low</b> , <b>medium-high</b> , <b>high</b> , or <b>none</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Drop profiles     | <p>Display the assignment of drop profiles.</p> <ul style="list-style-type: none"> <li><b>Loss priority</b>—Packet loss priority for drop profile assignment.</li> <li><b>Protocol</b>—Transport protocol for drop profile assignment.</li> <li><b>Index</b>—Index of the indicated object. Objects that have indexes in this output include schedulers and drop profiles.</li> <li><b>Name</b>—Name of the drop profile.</li> <li><b>Type</b>—Type of the drop profile: <b>discrete</b> or <b>interpolated</b>.</li> <li><b>Fill Level</b>—Percentage fullness of a queue.</li> <li><b>Drop probability</b>—Drop probability at this fill level.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                              |
| Excess Priority   | Priority of the excess bandwidth traffic on a scheduler.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

Table 636: show class-of-service interface Output Fields (*continued*)

| Field Name                    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Drop profiles</b>          | <p>Display the assignment of drop profiles.</p> <ul style="list-style-type: none"> <li>• <b>Loss priority</b>—Packet loss priority for drop profile assignment.</li> <li>• <b>Protocol</b>—Transport protocol for drop profile assignment.</li> <li>• <b>Index</b>—Index of the indicated object. Objects that have indexes in this output include schedulers and drop profiles.</li> <li>• <b>Name</b>—Name of the drop profile.</li> <li>• <b>Type</b>—Type of the drop profile: <b>discrete</b> or <b>interpolated</b>.</li> <li>• <b>Fill Level</b>—Percentage fullness of a queue.</li> <li>• <b>Drop probability</b>—Drop probability at this fill level.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Adjustment information</b> | <p>Display the assignment of shaping-rate adjustments on a scheduler node or queue.</p> <ul style="list-style-type: none"> <li>• <b>Adjusting application</b>—Application that is performing the shaping-rate adjustment. <ul style="list-style-type: none"> <li>• The adjusting application can appear as <b>ancp LS-0</b>, which is the Junos OS Access Node Control Profile process (<b>ancpd</b>) that performs shaping-rate adjustments on schedule nodes.</li> <li>• The adjusting application can also appear as <b>pppoe</b>, which adjusts the shaping-rate and overhead-accounting class-of-service attributes on dynamic subscriber interfaces in a broadband access network based on access line parameters in Point-to-Point Protocol over Ethernet (PPPoE) Tags [TR-101]. This feature is supported on MPC/MIC interfaces on MX Series routers. The shaping rate is based on the actual-data-rate-downstream attribute. The overhead accounting value is based on the access-loop-encapsulation attribute and specifies whether the access loop uses Ethernet (frame mode) or ATM (cell mode).</li> </ul> </li> <li>• <b>Adjustment type</b>—Type of adjustment: <b>absolute</b> or <b>delta</b>.</li> <li>• <b>Configured shaping rate</b>—Shaping rate configured for the scheduler node or queue.</li> <li>• <b>Adjustment value</b>—Value of adjusted shaping rate.</li> <li>• <b>Adjustment target</b>—Level of shaping-rate adjustment performed: <b>node</b> or <b>queue</b>.</li> <li>• <b>Adjustment overhead-accounting mode</b>—Configured shaping mode: <b>frame</b> or <b>cell</b>.</li> <li>• <b>Adjustment overhead bytes</b>—Number of bytes that the ANCP agent adds to or subtracts from the actual downstream frame overhead before reporting the adjusted values to CoS.</li> <li>• <b>Adjustment target</b>—Level of shaping-rate adjustment performed: <b>node</b> or <b>queue</b>.</li> <li>• <b>Adjustment multicast index</b>—</li> </ul> |

## Sample Output

### show class-of-service interface (Physical)

```

user@host> show class-of-service interface so-0/2/3
Physical interface: so-0/2/3, Index: 135
Queues supported: 8, Queues in use: 4
Total non-default queues created: 4
Scheduler map: <default>, Index: 2032638653

Logical interface: fe-0/0/1.0, Index: 68, Dedicated Queues: no
Shaping rate: 32000

```

| Object        | Name        | Type | Index |
|---------------|-------------|------|-------|
| Scheduler-map | <default>   |      | 27    |
| Rewrite       | exp-default | exp  | 21    |
| Classifier    | exp-default | exp  | 5     |

|                      |                      |     |   |
|----------------------|----------------------|-----|---|
| Classifier           | ipprec-compatibility | ip  | 8 |
| Forwarding-class-map | exp-default          | exp | 5 |

#### show class-of-service interface (Logical)

```
user@host> show class-of-service interface so-0/2/3.0
Logical interface: so-0/2/3.0, Index: 68, Dedicated Queues: no
Shaping rate: 32000
Object Name Type Index
Scheduler-map <default> 27
Rewrite exp-default exp 21
Classifier exp-default exp 5
Classifier ipprec-compatibility ip 8
Forwarding-class-map exp-default exp 5
```

#### show class-of-service interface (Gigabit Ethernet)

```
user@host> show class-of-service interface ge-6/2/0
Physical interface: ge-6/2/0, Index: 175
Queues supported: 4, Queues in use: 4
Scheduler map: <default>, Index: 2
Input scheduler map: <default>, Index: 3
Chassis scheduler map: <default-chassis>, Index: 4
```

#### show class-of-service interface (ANCP)

```
user@host> show class-of-service interface pp0.1073741842
Logical interface: pp0.1073741842, Index: 341
Object Name Type Index
Traffic-control-profile TCP-CVLAN Output 12408
Classifier dscp-ipv6-compatibility dscp-ipv6 9
Classifier ipprec-compatibility ip 13

Adjusting application: ancp LS-0
Adjustment type: absolute
Configured shaping rate: 4000000
Adjustment value: 11228000
Adjustment overhead-accounting mode: Frame Mode
Adjustment overhead bytes: 50
Adjustment target: node
```

#### show class-of-service interface (PPPoE Interface)

```
user@host> show class-of-service interface pp0.1
Logical interface: pp0.1, Index: 85
Object Name Type Index
Traffic-control-profile tcp-pppoe.o.pp0.1 Output 2726446535
Classifier ipprec-compatibility ip 13

Adjusting application: PPPoE
Adjustment type: absolute
Adjustment value: 5000000
Adjustment overhead-accounting mode: cell
Adjustment target: node
```

#### show class-of-service interface (T4000 Routers with Type 5 FPCs)

```
user@host> show class-of-service interface xe-4/0/0
Physical interface: xe-4/0/0, Index: 153
Queues supported: 8, Queues in use: 4
Shaping rate: 5000000000 bps
Scheduler map: <default>, Index: 2
```

```

Congestion-notification: Disabled

Logical interface: xe-4/0/0.0, Index: 77
 Object Name Type
Index
 Classifier ipprec-compatibility ip
13

```

### show class-of-service interface detail

```
user@host> show class-of-service interface ge-0/3/0 detail
```

```

Physical interface: ge-0/3/0, Enabled, Physical link is Up
 Link-level type: Ethernet, MTU: 1518, Speed: 1000mbps, Loopback: Disabled,
 Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
 Remote fault: Online
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x4000

```

```

Physical interface: ge-0/3/0, Index: 138
Queues supported: 4, Queues in use: 5
Shaping rate: 50000 bps
Scheduler map: interface-scheduler-map, Index: 58414
Input shaping rate: 10000 bps
Input scheduler map: scheduler-map, Index: 15103
Chassis scheduler map: <default-chassis>, Index: 4
Congestion-notification: Disabled

```

```

Logical interface ge-0/3/0.0
 Flags: SNMP-Traps 0x4000 VLAN-Tag [0x8100.1] Encapsulation: ENET2
 inet
 mpls

```

| Interface  | Admin | Link | Proto | Input Filter  | Output Filter  |
|------------|-------|------|-------|---------------|----------------|
| ge-0/3/0.0 | up    | up   | inet  |               |                |
|            |       |      | mpls  |               |                |
| Interface  | Admin | Link | Proto | Input Policer | Output Policer |
| ge-0/3/0.0 | up    | up   | inet  |               |                |
|            |       |      | mpls  |               |                |

```

Logical interface: ge-0/3/0.0, Index: 68
 Object Name Type Index
 Rewrite exp-default exp (mpls-any) 33
 Classifier exp-default exp 10
 Classifier ipprec-compatibility ip 13

```

```

Logical interface ge-0/3/0.1
 Flags: SNMP-Traps 0x4000 VLAN-Tag [0x8100.2] Encapsulation: ENET2
 inet

```

| Interface  | Admin | Link | Proto | Input Filter  | Output Filter  |
|------------|-------|------|-------|---------------|----------------|
| ge-0/3/0.1 | up    | up   | inet  |               |                |
| Interface  | Admin | Link | Proto | Input Policer | Output Policer |
| ge-0/3/0.1 | up    | up   | inet  |               |                |

```

Logical interface: ge-0/3/0.1, Index: 69
 Object Name Type Index
 Classifier ipprec-compatibility ip 13

```

### show class-of-service interface comprehensive

```
user@host> show class-of-service interface ge-0/3/0 comprehensive
```

```

Physical interface: ge-0/3/0, Enabled, Physical link is Up
 Interface index: 138, SNMP ifIndex: 601, Generation: 141
 Link-level type: Ethernet, MTU: 1518, Speed: 1000Mbps, BPDU Error: None,
 MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled, Flow
 control: Enabled,
 Auto-negotiation: Enabled, Remote fault: Online
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x4000
 CoS queues : 4 supported, 4 maximum usable queues
 Schedulers : 256
 Hold-times : Up 0 ms, Down 0 ms
 Current address: 00:14:f6:f4:b4:5d, Hardware address: 00:14:f6:f4:b4:5d
 Last flapped : 2010-09-07 06:35:22 PDT (15:14:42 ago)
 Statistics last cleared: Never
Traffic statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps
IPv6 total statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Ingress traffic statistics at Packet Forwarding Engine:
 Input bytes : 0 0 bps
 Input packets: 0 0 pps
 Drop bytes : 0 0 bps
 Drop packets: 0 0 pps
Label-switched interface (LSI) traffic statistics:
 Input bytes : 0 0 bps
 Input packets: 0 0 pps
Input errors:
 Errors: 0, Drops: 0, Framing errors: 0, Runt: 0, Policed discards: 0, L3
 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
 Resource errors: 0
Output errors:
 Carrier transitions: 5, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
 FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Ingress queues: 4 supported, 5 in use
Queue counters: Queued packets Transmitted packets Dropped packets

 0 af3 0 0 0
 1 af2 0 0 0
 2 ef2 0 0 0
 3 ef1 0 0 0

Egress queues: 4 supported, 5 in use
Queue counters: Queued packets Transmitted packets Dropped packets

 0 af3 0 0 0
 1 af2 0 0 0
 2 ef2 0 0 0
 3 ef1 0 0 0

```

```

Active alarms : None
Active defects : None
MAC statistics:
 Receive Transmit
 Total octets 0 0
 Total packets 0 0
 Unicast packets 0 0
 Broadcast packets 0 0
 Multicast packets 0 0
 CRC/Align errors 0 0
 FIFO errors 0 0
 MAC control frames 0 0
 MAC pause frames 0 0
 Oversized frames 0
 Jabber frames 0
 Fragment frames 0
 VLAN tagged frames 0
 Code violations 0
Filter statistics:
 Input packet count 0
 Input packet rejects 0
 Input DA rejects 0
 Input SA rejects 0
 Output packet count 0
 Output packet pad count 0
 Output packet error count 0
 CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
 Negotiation status: Complete
 Link partner:
 Link mode: Full-duplex, Flow control: Symmetric/Asymmetric, Remote fault:
OK
 Local resolution:
 Flow control: Symmetric, Remote fault: Link OK
Packet Forwarding Engine configuration:
 Destination slot: 0
CoS information:
 Direction : Output
 CoS transmit queue Bandwidth Buffer Priority
Limit
 % bps % usec
 2 ef2 39 19500 0 120 high
none
 Direction : Input
 CoS transmit queue Bandwidth Buffer Priority
Limit
 % bps % usec
 0 af3 30 3000 45 0 low
none

Physical interface: ge-0/3/0, Enabled, Physical link is Up
 Interface index: 138, SNMP ifIndex: 601
Forwarding classes: 16 supported, 5 in use
Ingress queues: 4 supported, 5 in use
Queue: 0, Forwarding classes: af3
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : Not Available

```



```

RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: af2
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: ef2
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: ef1
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Forwarding classes: 16 supported, 5 in use
Egress queues: 4 supported, 5 in use
Queue: 0, Forwarding classes: af3
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RL-dropped packets : 0 0 pps
 RL-dropped bytes : 0 0 bps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: af2
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RL-dropped packets : 0 0 pps
 RL-dropped bytes : 0 0 bps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: ef2
 Queued:
 Packets : 0 0 pps

```

```

 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RL-dropped packets : 0 0 pps
 RL-dropped bytes : 0 0 bps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: ef1
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RL-dropped packets : 0 0 pps
 RL-dropped bytes : 0 0 bps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 5 in use
Queue: 0, Forwarding classes: af3
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : Not Available
 RED-dropped bytes : Not Available
Queue: 1, Forwarding classes: af2
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : Not Available
 RED-dropped bytes : Not Available
Queue: 2, Forwarding classes: ef2
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : Not Available
 RED-dropped bytes : Not Available
Queue: 3, Forwarding classes: ef1
 Queued:
 Packets : 108546 0 pps
 Bytes : 12754752 376 bps
 Transmitted:
 Packets : 108546 0 pps
 Bytes : 12754752 376 bps

```

```

Tail-dropped packets : 0 0 pps
RED-dropped packets : Not Available
RED-dropped bytes : Not Available

```

```

Physical interface: ge-0/3/0, Index: 138
Queues supported: 4, Queues in use: 5
Shaping rate: 50000 bps

```

```
Scheduler map: interface-scheduler-map, Index: 58414
```

```

Scheduler: ef2, Forwarding class: ef2, Index: 39155
 Transmit rate: 39 percent, Rate Limit: none, Buffer size: 120 us, Buffer
 Limit: none, Priority: high
 Excess Priority: unspecified
 Drop profiles:
 Loss priority Protocol Index Name
 Low any 1 < default-drop-profile>
 Medium low any 1 < default-drop-profile>
 Medium high any 1 < default-drop-profile>
 High any 1 < default-drop-profile>
 Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
 Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
 Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
 Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
 Input shaping rate: 10000 bps
 Input scheduler map: scheduler-map

```

```
Scheduler map: scheduler-map, Index: 15103
```

```

Scheduler: af3, Forwarding class: af3, Index: 35058
 Transmit rate: 30 percent, Rate Limit: none, Buffer size: 45 percent, Buffer
 Limit: none, Priority: low
 Excess Priority: unspecified
 Drop profiles:
 Loss priority Protocol Index Name
 Low any 40582 green
 Medium low any 1 < default-drop-profile>
 Medium high any 1 < default-drop-profile>
 High any 18928 yellow
 Drop profile: green, Type: discrete, Index: 40582
 Fill level Drop probability
 50 0
 100 100
 Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
 Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
 Drop profile: yellow, Type: discrete, Index: 18928
 Fill level Drop probability
 50 0
 100 100

```

```

Chassis scheduler map: < default-drop-profile>
Scheduler map: < default-drop-profile>, Index: 4

Scheduler: < default-drop-profile>, Forwarding class: af3, Index: 25
 Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer
 Limit: none, Priority: low
 Excess Priority: low
 Drop profiles:
 Loss priority Protocol Index Name
 Low any 1 < default-drop-profile>
 Medium low any 1 < default-drop-profile>
 Medium high any 1 < default-drop-profile>
 High any 1 < default-drop-profile>
 Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
 Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
 Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
 Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100

Scheduler: < default-drop-profile>, Forwarding class: af2, Index: 25
 Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer
 Limit: none, Priority: low
 Excess Priority: low
 Drop profiles:
 Loss priority Protocol Index Name
 Low any 1 < default-drop-profile>
 Medium low any 1 < default-drop-profile>
 Medium high any 1 < default-drop-profile>
 High any 1 < default-drop-profile>
 Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
 Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
 Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
 Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100

Scheduler: < default-drop-profile>, Forwarding class: ef2, Index: 25
 Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer
 Limit: none, Priority: low
 Excess Priority: low
 Drop profiles:
 Loss priority Protocol Index Name
 Low any 1 < default-drop-profile>
 Medium low any 1 < default-drop-profile>
 Medium high any 1 < default-drop-profile>
 High any 1 < default-drop-profile>
 Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100

```

```

 100 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100

Scheduler: < default-drop-profile>, Forwarding class: ef1, Index: 25
 Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer
Limit: none, Priority: low
 Excess Priority: low
 Drop profiles:
 Loss priority Protocol Index Name
 Low any 1 < default-drop-profile>
 Medium low any 1 < default-drop-profile>
 Medium high any 1 < default-drop-profile>
 High any 1 < default-drop-profile>
Drop profile: , Type: discrete, Index: 1
 Fill level Drop probability
 100 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
 Congestion-notification: Disabled
Forwarding class
priority Policing priority ID Queue Restricted queue Fabric
af3 normal 0 0 0 low
af2 normal 1 1 1 low
ef2 normal 2 2 2 high
ef1 normal 3 3 3 high
af1 normal 4 4 0 low
 normal

Logical interface ge-0/3/0.0 (Index 68) (SNMP ifIndex 152) (Generation 159)
 Flags: SNMP-Traps 0x4000 VLAN-Tag [0x8100.1] Encapsulation: ENET2
 Traffic statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
 Local statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
 Transit statistics:
 Input bytes : 0 0 bps

```

```

Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
Protocol inet, MTU: 1500, Generation: 172, Route table: 0
 Flags: Sendbcst-pkt-to-re
 Input Filters: filter-in-ge-0/3/0.0-i,
 Policer: Input: p1-ge-0/3/0.0-inet-i
Protocol mpls, MTU: 1488, Maximum labels: 3, Generation: 173, Route table: 0
 Flags: Is-Primary
 Output Filters: exp-filter,,,,,

Logical interface ge-1/2/0.0 (Index 347) (SNMP ifIndex 638) (Generation 156)

Forwarding class ID Queue Restricted queue Fabric priority Policing priority
 SPU priority
best-effort 0 0 0 low normal
 low

Aggregate Forwarding-class statistics per forwarding-class
Aggregate Forwarding-class statistics:
Forwarding-class statistics:

Forwarding-class best-effort statistics:
 Input unicast bytes: 0
 Output unicast bytes: 0
 Input unicast packets: 0
 Output unicast packets: 0

 Input multicast bytes: 0
 Output multicast bytes: 0
 Input multicast packets: 0
 Output multicast packets: 0

Forwarding-class expedited-forwarding statistics:
 Input unicast bytes: 0
 Output unicast bytes: 0
 Input unicast packets: 0
 Output unicast packets: 0

 Input multicast bytes: 0
 Output multicast bytes: 0
 Input multicast packets: 0
 Output multicast packets: 0

IPv4 protocol forwarding-class statistics:
Forwarding-class statistics:
Forwarding-class best-effort statistics:

 Input unicast bytes: 0
 Output unicast bytes: 0
 Input unicast packets: 0
 Output unicast packets: 0

 Input multicast bytes: 0
 Output multicast bytes: 0
 Input multicast packets: 0
 Output multicast packets: 0

Forwarding-class expedited-forwarding statistics:
 Input unicast bytes: 0

```

```

Output unicast bytes: 0
Input unicast packets: 0
Output unicast packets: 0

```

```

Input multicast bytes: 0
Output multicast bytes: 0
Input multicast packets: 0
Output multicast packets: 0

```

```

IPv6 protocol forwarding-class statistics:
Forwarding-class statistics:
 Forwarding-class best-effort statistics:

```

```

Input unicast bytes: 0
Output unicast bytes: 0
Input unicast packets: 0
Output unicast packets: 0

```

```

Input multicast bytes: 0
Output multicast bytes: 0
Input multicast packets: 0
Output multicast packets: 0

```

```
Forwarding-class expedited-forwarding statistics:
```

```

Input unicast bytes: 0
Output unicast bytes: 0
Input unicast packets: 0
Output unicast packets: 0

```

```
Logical interface ge-0/3/0.0 (Index 68) (SNMP ifIndex 152)
```

```

Flags: SNMP-Traps 0x4000 VLAN-Tag [0x8100.1] Encapsulation: ENET2
Input packets : 0
Output packets: 0

```

|            |       |      |       |                        |               |
|------------|-------|------|-------|------------------------|---------------|
| Interface  | Admin | Link | Proto | Input Filter           | Output Filter |
| ge-0/3/0.0 | up    | up   | inet  | filter-in-ge-0/3/0.0-i |               |
|            |       |      | mpls  |                        | exp-filter    |

|            |       |      |       |                      |                |
|------------|-------|------|-------|----------------------|----------------|
| Interface  | Admin | Link | Proto | Input Policer        | Output Policer |
| ge-0/3/0.0 | up    | up   |       |                      |                |
|            |       |      | inet  | p1-ge-0/3/0.0-inet-i |                |
|            |       |      | mpls  |                      |                |

```
Filter: filter-in-ge-0/3/0.0-i
```

```
Counters:
```

| Name                         | Bytes | Packets |
|------------------------------|-------|---------|
| count-filter-in-ge-0/3/0.0-i | 0     | 0       |

```
Filter: exp-filter
```

```
Counters:
```

| Name                  | Bytes | Packets |
|-----------------------|-------|---------|
| count-exp-seven-match | 0     | 0       |
| count-exp-zero-match  | 0     | 0       |

```
Policers:
```

| Name                 | Packets |
|----------------------|---------|
| p1-ge-0/3/0.0-inet-i | 0       |

```
Logical interface: ge-0/3/0.0, Index: 68
```

| Object  | Name        | Type           | Index |
|---------|-------------|----------------|-------|
| Rewrite | exp-default | exp (mpls-any) | 33    |

Rewrite rule: exp-default, Code point type: exp, Index: 33

| Forwarding class | Loss priority | Code point |       |
|------------------|---------------|------------|-------|
| af3              | low           | 000        |       |
| af3              | high          | 001        |       |
| af2              | low           | 010        |       |
| af2              | high          | 011        |       |
| ef2              | low           | 100        |       |
| ef2              | high          | 101        |       |
| ef1              | low           | 110        |       |
| ef1              | high          | 111        |       |
| Object           | Name          | Type       | Index |
| Classifier       | exp-default   | exp        | 10    |

Classifier: exp-default, Code point type: exp, Index: 10

| Code point | Forwarding class     | Loss priority |       |
|------------|----------------------|---------------|-------|
| 000        | af3                  | low           |       |
| 001        | af3                  | high          |       |
| 010        | af2                  | low           |       |
| 011        | af2                  | high          |       |
| 100        | ef2                  | low           |       |
| 101        | ef2                  | high          |       |
| 110        | ef1                  | low           |       |
| 111        | ef1                  | high          |       |
| Object     | Name                 | Type          | Index |
| Classifier | ipprec-compatibility | ip            | 13    |

Classifier: ipprec-compatibility, Code point type: inet-precedence, Index: 13

| Code point       | Forwarding class | Loss priority |                  |        |
|------------------|------------------|---------------|------------------|--------|
| 000              | af3              | low           |                  |        |
| 001              | af3              | high          |                  |        |
| 010              | af3              | low           |                  |        |
| 011              | af3              | high          |                  |        |
| 100              | af3              | low           |                  |        |
| 101              | af3              | high          |                  |        |
| 110              | ef1              | low           |                  |        |
| 111              | ef1              | high          |                  |        |
| Forwarding class | ID               | Queue         | Restricted queue | Fabric |
| priority         |                  |               |                  |        |
| af3              | 0                | 0             | 0                | low    |
| normal           |                  |               |                  |        |
| af2              | 1                | 1             | 1                | low    |
| normal           |                  |               |                  |        |
| ef2              | 2                | 2             | 2                | high   |
| normal           |                  |               |                  |        |
| ef1              | 3                | 3             | 3                | high   |
| normal           |                  |               |                  |        |
| af1              | 4                | 4             | 0                | low    |
| normal           |                  |               |                  |        |

Logical interface ge-0/3/0.1 (Index 69) (SNMP ifIndex 154) (Generation 160)

Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.2 ] Encapsulation: ENET2

Traffic statistics:

|                   |   |
|-------------------|---|
| Input bytes :     | 0 |
| Output bytes :    | 0 |
| Input packets:    | 0 |
| Output packets:   | 0 |
| Local statistics: |   |
| Input bytes :     | 0 |



```

Output bytes : 0
Input packets: 0
Output packets: 0
Transit statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
Protocol inet, MTU: 1500, Generation: 174, Route table: 0
Flags: Sendbroadcast-pkt-to-re

```

```

Logical interface ge-0/3/0.1 (Index 69) (SNMP ifIndex 154)
Flags: SNMP-Traps 0x4000 VLAN-Tag [0x8100.2] Encapsulation: ENET2
Input packets : 0
Output packets: 0

```

```

Interface Admin Link Proto Input Filter Output Filter
ge-0/3/0.1 up up mpls
Interface Admin Link Proto Input Policer Output Policer
ge-0/3/0.1 up up mpls

```

```
Logical interface: ge-0/3/0.1, Index: 69
```

| Object     | Name                 | Type | Index |
|------------|----------------------|------|-------|
| Classifier | ipprec-compatibility | ip   | 13    |

```
Classifier: ipprec-compatibility, Code point type: inet-precedence, Index: 13
```

| Code point | Forwarding class | Loss priority |
|------------|------------------|---------------|
| 000        | af3              | low           |
| 001        | af3              | high          |
| 010        | af3              | low           |
| 011        | af3              | high          |
| 100        | af3              | low           |
| 101        | af3              | high          |
| 110        | ef1              | low           |
| 111        | ef1              | high          |

| Forwarding class | ID | Queue | Restricted queue | Fabric |
|------------------|----|-------|------------------|--------|
| priority         |    |       |                  |        |
| af3              | 0  | 0     | 0                | low    |
| af2              | 1  | 1     | 1                | low    |
| ef2              | 2  | 2     | 2                | high   |
| ef1              | 3  | 3     | 3                | high   |
| af1              | 4  | 4     | 0                | low    |

### show class-of-service interface (ACX Series Routers)

```

user@host-g11# show class-of-service interface
Physical interface: at-0/0/0, Index: 130
Queues supported: 4, Queues in use: 4
Scheduler map: <default>, Index: 2
Congestion-notification: Disabled

```

Logical interface: at-0/0/0.0, Index: 69

Logical interface: at-0/0/0.32767, Index: 70

Physical interface: at-0/0/1, Index: 133

Queues supported: 4, Queues in use: 4

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

Logical interface: at-0/0/1.0, Index: 71

Logical interface: at-0/0/1.32767, Index: 72

Physical interface: ge-0/1/0, Index: 146

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

| Object     | Name         | Type      | Index |
|------------|--------------|-----------|-------|
| Rewrite    | dscp-default | dscp      | 31    |
| Classifier | d1           | dscp      | 11331 |
| Classifier | ci           | ieee8021p | 583   |

Logical interface: ge-0/1/0.0, Index: 73

| Object  | Name       | Type           | Index |
|---------|------------|----------------|-------|
| Rewrite | custom-exp | exp (mpls-any) | 46413 |

Logical interface: ge-0/1/0.1, Index: 74

Logical interface: ge-0/1/0.32767, Index: 75

Physical interface: ge-0/1/1, Index: 147

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

| Object     | Name                 | Type | Index |
|------------|----------------------|------|-------|
| Classifier | ipprec-compatibility | ip   | 13    |

Logical interface: ge-0/1/1.0, Index: 76

Physical interface: ge-0/1/2, Index: 148

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

| Object     | Name | Type              | Index |
|------------|------|-------------------|-------|
| Rewrite    | ri   | ieee8021p (outer) | 35392 |
| Classifier | ci   | ieee8021p         | 583   |

Physical interface: ge-0/1/3, Index: 149

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

| Object     | Name                 | Type | Index |
|------------|----------------------|------|-------|
| Classifier | ipprec-compatibility | ip   | 13    |

Logical interface: ge-0/1/3.0, Index: 77

| Object  | Name        | Type           | Index |
|---------|-------------|----------------|-------|
| Rewrite | custom-exp2 | exp (mpls-any) | 53581 |

Physical interface: ge-0/1/4, Index: 150

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

```

Congestion-notification: Disabled
Object Name Type Index
Classifier ipprec-compatibility ip 13

Physical interface: ge-0/1/5, Index: 151
Queues supported: 8, Queues in use: 5
Scheduler map: <default>, Index: 2
Congestion-notification: Disabled
Object Name Type Index
Classifier ipprec-compatibility ip 13

Physical interface: ge-0/1/6, Index: 152
Queues supported: 8, Queues in use: 5
Scheduler map: <default>, Index: 2
Congestion-notification: Disabled
Object Name Type Index
Classifier ipprec-compatibility ip 13

Physical interface: ge-0/1/7, Index: 153
Queues supported: 8, Queues in use: 5
Scheduler map: <default>, Index: 2
Congestion-notification: Disabled
Object Name Type Index
Classifier d1 dscp 11331

Physical interface: ge-0/2/0, Index: 154
Queues supported: 8, Queues in use: 5
Scheduler map: <default>, Index: 2
Congestion-notification: Disabled
Object Name Type Index
Classifier ipprec-compatibility ip 13

Physical interface: ge-0/2/1, Index: 155
Queues supported: 8, Queues in use: 5
Scheduler map: <default>, Index: 2
Congestion-notification: Disabled
Object Name Type Index
Classifier ipprec-compatibility ip 13

Logical interface: ge-0/2/1.0, Index: 78

Logical interface: ge-0/2/1.32767, Index: 79

Physical interface: xe-0/3/0, Index: 156
Queues supported: 8, Queues in use: 5
Scheduler map: <default>, Index: 2
Congestion-notification: Disabled
Object Name Type Index
Classifier ipprec-compatibility ip 13

Logical interface: xe-0/3/0.0, Index: 80

Physical interface: xe-0/3/1, Index: 157
Queues supported: 8, Queues in use: 5
Scheduler map: <default>, Index: 2
Congestion-notification: Disabled
Object Name Type Index
Classifier ipprec-compatibility ip 13

Logical interface: xe-0/3/1.0, Index: 81

```

```
[edit]
user@host-g11#
```

#### show class-of-service interface (PPPoE Subscriber Interface for Enhanced Subscriber Management)

```
user@host> show class-of-service interface pp0.3221225474
 Logical interface: pp0.3221225475, Index: 3221225475
Object Name Type Index
Traffic-control-profile TC_PROF_100_199_SERIES_UID1006 Output 4294967312
Scheduler-map SMAP-1_UID1002 Output 4294967327
Rewrite-Output ieee-rewrite ieee8021p 60432
Rewrite-Output rule1 ip 50463

 Adjusting application: PPPoE IA tags
 Adjustment type: absolute
 Configured shaping rate: 11000000
 Adjustment value: 5000000
 Adjustment target: node

 Adjusting application: ucac
 Adjustment type: delta
 Configured shaping rate: 5000000
 Adjustment value: 100000
 Adjustment target: node
```

## show class-of-service scheduler-map

|                                 |                                                                                                                                                                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show class-of-service scheduler-map</code><br><code>&lt;name&gt;</code>                                                                                                                                                |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.<br>Command introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management. |
| <b>Description</b>              | Display the mapping of schedulers to forwarding classes and a summary of scheduler parameters for each entry.                                                                                                                |
| <b>Options</b>                  | <b>none</b> —Display all scheduler maps.<br><br><b>name</b> —(Optional) Display a summary of scheduler parameters for each forwarding class to which the named scheduler is assigned.                                        |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Verifying and Managing Junos OS Enhanced Subscriber Management</i></li> </ul>                                                                                                      |
| <b>List of Sample Output</b>    | <a href="#">show class-of-service scheduler-map on page 7334</a>                                                                                                                                                             |
| <b>Output Fields</b>            | <a href="#">Table 637</a> describes the output fields for the <b>show class-of-service scheduler-map</b> command. Output fields are listed in the approximate order in which they appear.                                    |

**Table 637: show class-of-service scheduler-map Output Fields**

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                                             |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Scheduler map</b>    | Name of the scheduler map.<br><br>(Enhanced subscriber management for MX Series routers) The name of the dynamic scheduler map object is associated with a generated UID (for example, <b>SMAP-1_UID1002</b> ) instead of with a subscriber interface.                                                                                        |
| <b>Index</b>            | Index of the indicated object. Objects having indexes in this output include scheduler maps, schedulers, and drop profiles.<br><br>(Enhanced subscriber management for MX Series routers) Index values for dynamic CoS traffic control profiles are larger for enhanced subscriber management than they are for legacy subscriber management. |
| <b>Scheduler</b>        | Name of the scheduler.                                                                                                                                                                                                                                                                                                                        |
| <b>Forwarding class</b> | Classification of a packet affecting the forwarding, scheduling, and marking policies applied as the packet transits the router.                                                                                                                                                                                                              |

Table 637: show class-of-service scheduler-map Output Fields (*continued*)

| Field Name                              | Field Description                                                                                                                                                                                                                                                                                          |
|-----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Transmit rate</b>                    | Configured transmit rate of the scheduler (in bps). The rate is a percentage of the total interface bandwidth, or the keyword <b>remainder</b> , which indicates that the scheduler receives the remaining bandwidth of the interface.                                                                     |
| <b>Rate Limit</b>                       | Rate limiting configuration of the queue. Possible values are <b>none</b> , meaning no rate limiting, and <b>exact</b> , meaning the queue only transmits at the configured rate.                                                                                                                          |
| <b>Maximum buffer delay</b>             | Amount of transmit delay (in milliseconds) or the buffer size of the queue. The buffer size is shown as a percentage of the total interface buffer allocation, or by the keyword <b>remainder</b> to indicate that the buffer is sized according to what remains after other scheduler buffer allocations. |
| <b>Priority</b>                         | Scheduling priority: <b>low</b> or <b>high</b> .                                                                                                                                                                                                                                                           |
| <b>Excess priority</b>                  | Priority of excess bandwidth: <b>low</b> , <b>medium-low</b> , <b>medium-high</b> , <b>high</b> , or <b>none</b> .                                                                                                                                                                                         |
| <b>Explicit Congestion Notification</b> | (QFX Series, OCX Series, and EX4600 switches only) Explicit congestion notification (ECN) state: <ul style="list-style-type: none"> <li>• Disable—ECN is disabled on the specified scheduler</li> <li>• Enable—ECN is enabled on the specified scheduler</li> </ul> ECN is disabled by default.            |
| <b>Adjust minimum</b>                   | Minimum shaping rate for an adjusted queue, in bps.                                                                                                                                                                                                                                                        |
| <b>Adjust percent</b>                   | Bandwidth adjustment applied to a queue, in percent.                                                                                                                                                                                                                                                       |
| <b>Drop profiles</b>                    | Table displaying the assignment of drop profiles by name and index to a given loss priority and protocol pair.                                                                                                                                                                                             |
| <b>Loss priority</b>                    | Packet loss priority for drop profile assignment.                                                                                                                                                                                                                                                          |
| <b>Protocol</b>                         | Transport protocol for drop profile assignment.                                                                                                                                                                                                                                                            |
| <b>Name</b>                             | Name of the drop profile.                                                                                                                                                                                                                                                                                  |

## Sample Output

### show class-of-service scheduler-map

```
user@host> show class-of-service scheduler-map
Scheduler map: dd-scheduler-map, Index: 84
```

```
Scheduler: aa-scheduler, Index: 8721, Forwarding class: aa-forwarding-class
Transmit rate: 30 percent, Rate Limit: none, Maximum buffer delay: 39 ms,
Priority: high
Drop profiles:
 Loss priority Protocol Index Name
 Low non-TCP 8724 aa-drop-profile
 Low TCP 9874 bb-drop-profile
 High non-TCP 8833 cc-drop-profile
 High TCP 8484 dd-drop-profile
```

Scheduler: bb-scheduler, Forwarding class: aa-forwarding-class  
Transmit rate: 40 percent, Rate limit: none, Maximum buffer delay: 68 ms,  
Priority: high  
Drop profiles:

| Loss priority | Protocol | Index | Name            |
|---------------|----------|-------|-----------------|
| Low           | non-TCP  | 8724  | aa-drop-profile |
| Low           | TCP      | 9874  | bb-drop-profile |
| High          | non-TCP  | 8833  | cc-drop-profile |
| High          | TCP      | 8484  | dd-drop-profile |

## show class-of-service traffic-control-profile

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show class-of-service traffic-control-profile</code><br><code>&lt;profile-name&gt;</code>                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.<br>Command introduced in Junos OS Release 12.2 for ACX Series Routers.<br>Command introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.                                                                                                                                                                    |
| <b>Description</b>              | For Gigabit Ethernet IQ PICs, Channelized IQ PICs, EQ DPCs, and Trio MPC/MIC interfaces only, display traffic shaping and scheduling profiles.<br><br>(ACX Series routers) For ATM IMA pseudowire interfaces, display traffic shaping and scheduling profiles.                                                                                                                                                                                                         |
| <b>Options</b>                  | <b>none</b> —Display all profiles.<br><br><b>profile-name</b> —(Optional) Display information about a single profile.                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Verifying and Managing Junos OS Enhanced Subscriber Management</i></li> </ul>                                                                                                                                                                                                                                                                                                                                              |
| <b>List of Sample Output</b>    | <a href="#">show class-of-service traffic-control-profile on page 7338</a><br><a href="#">show class-of-service traffic-control-profile (MX Series routers with Clear Channel Multi-Rate CE MIC) on page 7339</a><br><a href="#">show class-of-service traffic-control-profile (ACX Series routers with ATM IMA pseudowire interfaces) on page 7339</a><br><a href="#">show class-of-service traffic-control-profile (Enhanced Subscriber Management) on page 7339</a> |
| <b>Output Fields</b>            | Table 638 describes the output fields for the <b>show class-of-service traffic-control-profile</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                    |

**Table 638: show class-of-service traffic-control-profile Output Fields**

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                      |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Traffic control profile</b> | <p>Name of the traffic control profile.</p> <p>(Enhanced subscriber management for MX Series routers) The name of the dynamic traffic control profile object is associated with a generated UID (for example, <code>TC_PROF_100_199_SERIES_UID1000</code>) instead of with a subscriber interface.</p> |



Table 638: show class-of-service traffic-control-profile Output Fields (*continued*)

| Field Name                          | Field Description                                                                                                                                                                                                                                                                                                    |
|-------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Index</b>                        | Index number of the traffic control profile.<br><br>(Enhanced subscriber management for MX Series routers) Index values for dynamic CoS traffic control profiles are larger for enhanced subscriber management than they are for legacy subscriber management.                                                       |
| <b>ATM Service</b>                  | (MX Series routers with ATM Multi-Rate CE MIC) Configured category of ATM service. Possible values: <ul style="list-style-type: none"> <li>• cbr—Constant bit rate.</li> <li>• rtvbr—Real time variable bit rate.</li> <li>• nrtvbr—Non real time variable bit rate.</li> <li>• ubr—Unspecified bit rate.</li> </ul> |
| <b>Maximum Burst Size</b>           | Configured maximum burst size, in cells.                                                                                                                                                                                                                                                                             |
| <b>Peak rate</b>                    | Configured peak rate, in cps.                                                                                                                                                                                                                                                                                        |
| <b>Sustained rate</b>               | Configured sustained rate, in cps.                                                                                                                                                                                                                                                                                   |
| <b>Shaping rate</b>                 | Configured shaping rate, in bps.<br><br><b>NOTE:</b> (MX Series routers with ATM Multi-Rate CE MIC) Configured peak rate, in cps.                                                                                                                                                                                    |
| <b>Shaping rate burst</b>           | Configured burst size for the shaping rate, in bytes.<br><br><b>NOTE:</b> (MX Series routers with ATM Multi-Rate CE MIC) Configured maximum burst rate, in cells.                                                                                                                                                    |
| <b>Shaping rate priority high</b>   | Configured shaping rate for high-priority traffic, in bps.                                                                                                                                                                                                                                                           |
| <b>Shaping rate priority medium</b> | Configured shaping rate for medium-priority traffic, in bps.                                                                                                                                                                                                                                                         |
| <b>Shaping rate priority low</b>    | Configured shaping rate for low-priority traffic, in bps.                                                                                                                                                                                                                                                            |
| <b>Shaping rate excess high</b>     | Configured shaping rate for high-priority excess traffic, in bps.                                                                                                                                                                                                                                                    |
| <b>Shaping rate excess low</b>      | Configured shaping rate for low-priority excess traffic, in bps.                                                                                                                                                                                                                                                     |
| <b>Scheduler map</b>                | Name of the associated scheduler map.<br><br>(Enhanced subscriber management for MX Series routers) The name of the dynamic scheduler map object is associated with a generated UID (for example, <b>SMAP-1_UID1002</b> ) instead of with a subscriber interface.                                                    |
| <b>Delay Buffer rate</b>            | Configured delay buffer rate, in bps.                                                                                                                                                                                                                                                                                |

**Table 638: show class-of-service traffic-control-profile Output Fields (*continued*)**

| Field Name                      | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Excess rate</b>              | Configured excess rate, in percent or proportion.                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Excess rate high</b>         | Configured excess rate for high priority traffic, in percent or proportion.                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Excess rate low</b>          | Configured excess rate for low priority traffic, in percent or proportion.                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Guaranteed rate</b>          | <p>Configured guaranteed rate, in bps or cps.</p> <p><b>NOTE:</b> (MX Series routers with ATM Multi-Rate CE MIC) This value depends on the ATM service category chosen. Possible values:</p> <ul style="list-style-type: none"> <li>• <b>cbr</b>—Guaranteed rate is equal to the configured peak rate in cps.</li> <li>• <b>rtvbr</b>—Guaranteed rate is equal to the configured sustained rate in cps.</li> <li>• <b>nrtvbr</b>—Guaranteed rate is equal to the configured sustained rate in cps.</li> </ul> |
| <b>Guaranteed rate burst</b>    | Configured burst size for the guaranteed rate, in bytes.                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>adjust-minimum</b>           | Configured minimum shaping rate for an adjusted queue, in bps.                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>overhead accounting mode</b> | Configured shaping mode: <b>Frame Mode</b> or <b>Cell Mode</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Overhead bytes</b>           | Configured byte adjustment value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Adjust parent</b>            | <p>Configured shaping-rate adjustment for parent scheduler nodes. If enabled, this field appears.</p> <p><b>flow-aware</b> indicates that the parent scheduler node is adjusted only once per multicast channel.</p>                                                                                                                                                                                                                                                                                          |

## Sample Output

### show class-of-service traffic-control-profile

```

user@host> show class-of-service traffic-control-profile
Traffic control profile: Profile1, Index: 57625
 Scheduler map: m1
 Delay Buffer rate: 500000
 Guaranteed rate: 1000000

Traffic control profile: Profile2, Index: 57624
 Scheduler map: m2
 Delay Buffer rate: 600000
 Guaranteed rate: 2000000

Traffic control profile: Profile3, Index: 57627
 Scheduler map: m3

```

```

Delay Buffer rate: 800000
Guaranteed rate: 3000000
.Excess rate high: proportion 4

Traffic control profile: Profile4, Index: 57626
Scheduler map: m4
Delay Buffer rate: 750000
Guaranteed rate: 4000000
..adjust-minimum 20000000

Traffic control profile: foo, Index: 57626
Shaping rate: 100000000
Scheduler map: <default>
Overhead accounting mode: Frame Mode
Frame mode overhead accounting bytes: -12
Adjust parent: flow-aware

```

#### show class-of-service traffic-control-profile (MX Series routers with Clear Channel Multi-Rate CE MIC)

```

user@host> show class-of-service traffic-control-profile
Traffic control profile: at-vbr1, Index: 11395
ATM Service: RTVBR
Scheduler map: m3
Shaping rate: 1000 cps
Shaping rate burst: 500 cells
Delay Buffer rate: 2000 cps
Guaranteed rate: 1000 cps

Traffic control profile: foo, Index: 38286
ATM Service: UBR
Scheduler map: m3
overhead accounting mode: Frame Mode

```

#### show class-of-service traffic-control-profile (ACX Series routers with ATM IMA pseudowire interfaces)

```

user@host> show class-of-service traffic-control-profile
Traffic control profile: foo, Index: 38286
ATM Service: RTVBR
Shaping rate: 2000 cps
Shaping rate burst: 200 cells
Scheduler map: <default>
Delay Buffer rate: 1000 cps
Guaranteed rate: 1700 cps

```

#### show class-of-service traffic-control-profile (Enhanced Subscriber Management)

```

user@host> show class-of-service traffic-control-profile
Traffic control profile: TC_PROF_100_199_SERIES_UID1000, Index: 4294967313
Shaping rate: 11000000
Shaping rate burst: 1 bytes
Scheduler map: SMAP-1_UID1002
Delay Buffer rate: 5000000
Overhead accounting mode: Cell Mode
Frame mode overhead accounting bytes: -4
Cell mode overhead accounting bytes: 20

```

## show interfaces queue

---

**Syntax**    show interfaces queue  
              <aggregate | remaining-traffic>  
              <both-ingress-egress>  
              <egress>  
              <forwarding-class *forwarding-class*>  
              <ingress>  
              <interface-name *interface-name*>  
              <l2-statistics>

**Release Information**    Command introduced before Junos OS Release 7.4.  
                              **both-ingress-egress**, **egress**, and **ingress** options introduced in Junos OS Release 7.6.  
                              Command introduced in Junos OS Release 11.1 for the QFX Series.  
                              **l2-statistics** option introduced in Junos OS Release 12.1.  
                              Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description**    Display class-of-service (CoS) queue information for physical interfaces.

**Options**    **none**—Show detailed CoS queue statistics for all physical interfaces.

**aggregate**—(Optional) Display the aggregated queuing statistics of all logical interfaces that have traffic-control profiles configured. (Not on the QFX Series.)

**both-ingress-egress**—(Optional) On Gigabit Ethernet Intelligent Queuing 2 (IQ2) PICs, display both ingress and egress queue statistics. (Not on the QFX Series.)

**egress**—(Optional) Display egress queue statistics.

**forwarding-class *forwarding-class***—(Optional) Forwarding class name for this queue.  
                              Shows detailed CoS statistics for the queue associated with the specified forwarding class.

**ingress**—(Optional) On Gigabit Ethernet IQ2 PICs, display ingress queue statistics. (Not on the QFX Series.)

**interface-name *interface-name***—(Optional) Show detailed CoS queue statistics for the specified interface.

**l2-statistics**—(Optional) Display Layer 2 statistics for MLPPP, FRF.15, and FRF.16 bundles

**remaining-traffic**—(Optional) Display the remaining-traffic queue statistics of all logical interfaces that have traffic-control profiles configured.

### Overhead for Layer 2 Statistics

Transmitted packets and transmitted byte counts are displayed for the Layer 2 level with the addition of encapsulation overheads applied for fragmentation, as shown in [Table 280](#). Others counters, such as packets and bytes queued (input) and drop counters, are displayed at the Layer 3 level. In the case of link fragmentation and interleaving (LFI) for which fragmentation is not applied, corresponding Layer 2 overheads are added, as shown in [Table 280](#).

Table 639: Layer 2 Overhead and Transmitted Packets or Byte Counts

| Protocol       | Fragmentation       |                                   | LFI |
|----------------|---------------------|-----------------------------------|-----|
|                | First fragmentation | Second to <i>n</i> fragmentations |     |
|                | Bytes               | Bytes                             |     |
| MLPPP (Long)   | 13                  | 12                                | 8   |
| MLPPP (short)  | 11                  | 10                                | 8   |
| MLFR (FRF15)   | 12                  | 10                                | 8   |
| MFR (FRF16)    | 10                  | 8                                 | -   |
| MCMLPPP(Long)  | 13                  | 12                                | -   |
| MCMLPPP(Short) | 11                  | 10                                | -   |

## Layer 2 Statistics—Fragmentation Overhead Calculation

## MLPPP/MC-MLPPP Overhead details:

=====

## Fragment 1:

```

Outer PPP header : 4 bytes
Long or short sequence MLPPP header : 4 bytes or 2 bytes
Inner PPP header : 1 byte
HDLC flag and FCS bytes : 4 bytes

```

## Fragments 2 .. n :

```

Outer PPP header : 4 bytes
Long or short sequence MLPPP header : 4 bytes or 2 bytes
HDLC flag and FCS bytes : 4 bytes

```

## MLFR (FRF15) Overhead details:

=====

## Fragment 1:

```

Framereelay header : 2 bytes
Control,NLPID : 2 bytes
Fragmentaion header : 2 bytes
Inner proto : 2 bytes
HDLC flag and FCS : 4 bytes

```

## Fragments 2 ...n :

```

Framereelay header : 2 bytes
Control,NLPID : 2 bytes
Fragmentaion header : 2 bytes
HDLC flag and FCS : 4 bytes

```

## MFR (FRF16) Overhead details:

=====

Fragment 1:

|                     |   |         |
|---------------------|---|---------|
| Fragmentaion header | : | 2 bytes |
| Framereelay header  | : | 2 bytes |
| Inner proto         | : | 2 bytes |
| HDLC flag and FCS   | : | 4 bytes |

Fragments 2 ...n :

|                     |   |         |
|---------------------|---|---------|
| Fragmentaion header | : | 2 bytes |
| Framereelay header  | : | 2 bytes |
| HDLC flag and FCS   | : | 4 bytes |

## Overhead with LFI

MLPPP(Long & short sequence):

=====

|                   |   |         |
|-------------------|---|---------|
| Outer PPP header  | : | 4 bytes |
| HDLC flag and FCS | : | 4 bytes |

MLFR (FRF15):

=====

|                    |   |         |
|--------------------|---|---------|
| Framereelay header | : | 2 bytes |
| Control,NLPID      | : | 2 bytes |
| HDLC flag and FCS  | : | 4 bytes |

The following examples show overhead for different cases:

- A 1000-byte packet is sent to a mlppp bundle without any fragmentation. At the Layer 2 level, bytes transmitted is 1013 in 1 packet. This overhead is for MLPPP long sequence encap.
- A 1000-byte packet is sent to a mlppp bundle with a fragment threshold of 250byte. At the Layer 2 level, bytes transmitted is 1061 bytes in 5 packets.
- A 1000-byte LFI packet is sent to an mlppp bundle. At the Layer 2 level, bytes transmitted is 1008 in 1 packet.

**remaining-traffic**—(Optional) Display the queuing statistics of all logical interfaces that do not have traffic-control profiles configured. (Not on the QFX Series.)

## Additional Information

For rate-limited interfaces hosted on Modular Interface Cards (MICs), Modular Port Concentrators (MPCs), or Enhanced Queuing DPCs, rate-limit packet-drop operations occur *before* packets are queued for transmission scheduling. For such interfaces, the statistics for queued traffic do not include the packets that have already been dropped due to rate limiting, and consequently the displayed statistics for queued traffic are the same as the displayed statistics for transmitted traffic.



**NOTE:** For rate-limited interfaces hosted on other types of hardware, rate-limit packet-drop operations occur *after* packets are queued for transmission scheduling. For these other interface types, the statistics for queued traffic include the packets that are later dropped due to rate limiting, and consequently the displayed statistics for queued traffic equals the sum of the statistics for transmitted and rate-limited traffic.

On M Series routers (except for the M320 and M120 routers), this command is valid only for a PIC installed on an enhanced Flexible PIC Concentrator (FPC).

Queue statistics for aggregated interfaces are supported on the M Series and T Series routers only. Statistics for an aggregated interface are the summation of the queue statistics of the child links of that aggregated interface. You can view the statistics for a child interface by using the **show interfaces statistics** command for that child interface.

When you configure tricolor marking on a 10-port 1-Gigabit Ethernet PIC, for queues 6 and 7 only, the output does not display the number of queued bytes and packets, or the number of bytes and packets dropped because of RED. If you do not configure tricolor marking on the interface, these statistics are available for all queues.

For the 4-port Channelized OC12 IQE PIC and 1-port Channelized OC48 IQE PIC, the **Packet Forwarding Engine Chassis Queues** field represents traffic bound for a particular physical interface on the PIC. For all other PICs, the **Packet Forwarding Engine Chassis Queues** field represents the total traffic bound for the PIC.

For Gigabit Ethernet IQ2 PICs, the **show interfaces queue** command output does not display the number of tail-dropped packets. This limitation does not apply to Packet Forwarding Engine chassis queues.

When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (under the **Packet Forwarding Engine Chassis Queues** field) shows the prefragmentation values.

The behavior of the **egress** queues for the **Routing Engine-Generated Traffic** is not same as the configured queue for MLPPP and MFR configurations.

For information about how to configure CoS, see the *Junos OS Network Interfaces Library for Routing Devices*. For related CoS operational mode commands, see the [CLI Explorer](#).

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>List of Sample Output</b>    | <a href="#">show interfaces queue (Rate-Limited Interface on a Gigabit Ethernet MIC in an MPC) on page 7348</a><br><a href="#">show interfaces queue (Aggregated Ethernet on a T320 Router) on page 7349</a><br><a href="#">show interfaces queue (Gigabit Ethernet on a T640 Router) on page 7351</a><br><a href="#">show interfaces queue aggregate (Gigabit Ethernet Enhanced DPC) on page 7351</a><br><a href="#">show interfaces queue (Gigabit Ethernet IQ2 PIC) on page 7355</a><br><a href="#">show interfaces queue both-ingress-egress (Gigabit Ethernet IQ2 PIC) on page 7358</a><br><a href="#">show interfaces queue ingress (Gigabit Ethernet IQ2 PIC) on page 7360</a><br><a href="#">show interfaces queue egress (Gigabit Ethernet IQ2 PIC) on page 7361</a><br><a href="#">show interfaces queue remaining-traffic (Gigabit Ethernet Enhanced DPC) on page 7363</a><br><a href="#">show interfaces queue (Channelized OC12 IQE Type 3 PIC in SONET Mode) on page 7365</a><br><a href="#">show interfaces queue (QFX Series) on page 7375</a><br><a href="#">show interfaces queue l2-statistics (lsq interface) on page 7376</a><br><a href="#">show interfaces queue lsq (lsq-ifd) on page 7377</a> |
| <b>Output Fields</b>            | Table 281 lists the output fields for the <b>show interfaces queue</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

Table 640: show interfaces queue Output Fields

| Field Name                                                                                                                                          | Field Description                                                                                                                                                                                                                                                                                                                                                         |
|-----------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Physical interface</b>                                                                                                                           | Name of the physical interface.                                                                                                                                                                                                                                                                                                                                           |
| <b>Enabled</b>                                                                                                                                      | State of the interface. Possible values are described in the "Enabled Field" section under <i>Common Output Fields Description</i> .                                                                                                                                                                                                                                      |
| <b>Interface index</b>                                                                                                                              | Physical interface's index number, which reflects its initialization sequence.                                                                                                                                                                                                                                                                                            |
| <b>SNMP ifIndex</b>                                                                                                                                 | SNMP index number for the interface.                                                                                                                                                                                                                                                                                                                                      |
| <b>Forwarding classes supported</b>                                                                                                                 | Total number of forwarding classes supported on the specified interface.                                                                                                                                                                                                                                                                                                  |
| <b>Forwarding classes in use</b>                                                                                                                    | Total number of forwarding classes in use on the specified interface.                                                                                                                                                                                                                                                                                                     |
| <b>Ingress queues supported</b>                                                                                                                     | On Gigabit Ethernet IQ2 PICs only, total number of ingress queues supported on the specified interface.                                                                                                                                                                                                                                                                   |
| <b>Ingress queues in use</b>                                                                                                                        | On Gigabit Ethernet IQ2 PICs only, total number of ingress queues in use on the specified interface.                                                                                                                                                                                                                                                                      |
| <b>Output queues supported</b>                                                                                                                      | Total number of output queues supported on the specified interface.                                                                                                                                                                                                                                                                                                       |
| <b>Output queues in use</b>                                                                                                                         | Total number of output queues in use on the specified interface.                                                                                                                                                                                                                                                                                                          |
| <b>Egress queues supported</b>                                                                                                                      | Total number of egress queues supported on the specified interface.                                                                                                                                                                                                                                                                                                       |
| <b>Egress queues in use</b>                                                                                                                         | Total number of egress queues in use on the specified interface.                                                                                                                                                                                                                                                                                                          |
| <b>Queue counters (Ingress)</b>                                                                                                                     | CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces. <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul> |
| <b>Burst size</b>                                                                                                                                   | (Logical interfaces on IQ PICs only) Maximum number of bytes up to which the logical interface can burst. The burst size is based on the shaping rate applied to the interface.                                                                                                                                                                                           |
| The following output fields are applicable to both interface component and Packet Forwarding component in the <b>show interfaces queue</b> command: |                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Queue</b>                                                                                                                                        | Queue number.                                                                                                                                                                                                                                                                                                                                                             |
| <b>Forwarding classes</b>                                                                                                                           | Forwarding class name.                                                                                                                                                                                                                                                                                                                                                    |



Table 640: show interfaces queue Output Fields (*continued*)

| Field Name                  | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Queued Packets</b>       | <p>Number of packets queued to this queue.</p> <p><b>NOTE:</b> For Gigabit Ethernet IQ2 interfaces, the Queued Packets count is calculated by the Junos OS interpreting one frame buffer as one packet. If the queued packets are very large or very small, the calculation might not be completely accurate for transit traffic. The count is completely accurate for traffic terminated on the router.</p> <p>For rate-limited interfaces hosted on MICs or MPCs only, this statistic does not include traffic dropped due to rate limiting. For more information, see <a href="#">“Additional Information” on page 3148</a>.</p> |
| <b>Queued Bytes</b>         | <p>Number of bytes queued to this queue. The byte counts vary by interface hardware. For more information, see <a href="#">Table 282</a>.</p> <p>For rate-limited interfaces hosted on MICs or MPCs only, this statistic does not include traffic dropped due to rate limiting. For more information, see <a href="#">“Additional Information” on page 3148</a>.</p>                                                                                                                                                                                                                                                                |
| <b>Transmitted Packets</b>  | <p>Number of packets transmitted by this queue. When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (displayed under the <b>Packet Forwarding Engine Chassis Queues</b> field) shows the prefragmentation values.</p> <p><b>NOTE:</b> For Layer 2 statistics, see <a href="#">“Overhead for Layer 2 Statistics” on page 3146</a></p>                                                                                                                                                                                          |
| <b>Transmitted Bytes</b>    | <p>Number of bytes transmitted by this queue. The byte counts vary by interface hardware. For more information, see <a href="#">Table 282</a>.</p> <p><b>NOTE:</b> On MX Series routers, this number can be inaccurate when you issue the command for a physical interface repeatedly and in quick succession, because the statistics for the child nodes are collected infrequently. Wait ten seconds between successive iterations to avoid this situation.</p> <p><b>NOTE:</b> For Layer 2 statistics, see <a href="#">“Overhead for Layer 2 Statistics” on page 3146</a></p>                                                    |
| <b>Tail-dropped packets</b> | <p>Number of packets dropped because of tail drop.</p> <p><b>NOTE:</b> The <b>Tail-dropped packets</b> counter is not supported on the PTX Series Packet Transport Routers.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>RL-dropped packets</b>   | <p>Number of packets dropped due to rate limiting.</p> <p>For rate-limited interfaces hosted on MICs, MPCs, and Enhanced Queuing DPCs only, this statistic is not included in the queued traffic statistics. For more information, see <a href="#">“Additional Information” on page 3148</a>.</p> <p><b>NOTE:</b> The <b>RL-dropped packets</b> counter is not supported on the PTX Series Packet Transport Routers, and is omitted from the output.</p>                                                                                                                                                                            |
| <b>RL-dropped bytes</b>     | <p>Number of bytes dropped due to rate limiting.</p> <p>For rate-limited interfaces hosted on MICs, MPCs, and Enhanced Queuing DPCs only, this statistic is not included in the queued traffic statistics. For more information, see <a href="#">“Additional Information” on page 3148</a>.</p>                                                                                                                                                                                                                                                                                                                                     |

Table 640: show interfaces queue Output Fields (*continued*)

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>RED-dropped packets</b> | <p>Number of packets dropped because of random early detection (RED).</p> <ul style="list-style-type: none"> <li>(M Series and T Series routers only) On M320 and M120 routers and the T Series routers, the total number of dropped packets is displayed. On all other M Series routers, the output classifies dropped packets into the following categories: <ul style="list-style-type: none"> <li><b>Low, non-TCP</b>—Number of low-loss priority non-TCP packets dropped because of RED.</li> <li><b>Low, TCP</b>—Number of low-loss priority TCP packets dropped because of RED.</li> <li><b>High, non-TCP</b>—Number of high-loss priority non-TCP packets dropped because of RED.</li> <li><b>High, TCP</b>—Number of high-loss priority TCP packets dropped because of RED.</li> </ul> </li> <li>(MX Series routers with enhanced DPCs, and T Series routers with enhanced FPCs only) The output classifies dropped packets into the following categories: <ul style="list-style-type: none"> <li><b>Low</b>—Number of low-loss priority packets dropped because of RED.</li> <li><b>Medium-low</b>—Number of medium-low loss priority packets dropped because of RED.</li> <li><b>Medium-high</b>—Number of medium-high loss priority packets dropped because of RED.</li> <li><b>High</b>—Number of high-loss priority packets dropped because of RED.</li> </ul> </li> </ul> <p><b>NOTE:</b> Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), this field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p> |
| <b>RED-dropped bytes</b>   | <p>Number of bytes dropped because of RED. The byte counts vary by interface hardware. For more information, see <a href="#">Table 282</a>.</p> <ul style="list-style-type: none"> <li>(M Series and T Series routers only) On M320 and M120 routers and the T Series routers, only the total number of dropped bytes is displayed. On all other M Series routers, the output classifies dropped bytes into the following categories: <ul style="list-style-type: none"> <li><b>Low, non-TCP</b>—Number of low-loss priority non-TCP bytes dropped because of RED.</li> <li><b>Low, TCP</b>—Number of low-loss priority TCP bytes dropped because of RED.</li> <li><b>High, non-TCP</b>—Number of high-loss priority non-TCP bytes dropped because of RED.</li> <li><b>High, TCP</b>—Number of high-loss priority TCP bytes dropped because of RED.</li> </ul> </li> </ul> <p><b>NOTE:</b> Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), this field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

Byte counts vary by interface hardware. [Table 282](#) shows how the byte counts on the outbound interfaces vary depending on the interface hardware. [Table 282](#) is based on the assumption that outbound interfaces are sending IP traffic with 478 bytes per packet.

Table 641: Byte Count by Interface Hardware

| Interface Hardware               | Output Level                | Byte Count Includes                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Comments                                                                                                                                                                                                     |
|----------------------------------|-----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Gigabit Ethernet IQ and IQE PICs | Interface                   | <p>Queued: 490 bytes per packet, representing 478 bytes of Layer 3 packet + 12 bytes</p> <p>Transmitted: 490 bytes per packet, representing 478 bytes of Layer 3 packet + 12 bytes</p> <p>RED dropped: 496 bytes per packet representing 478 bytes of Layer 3 packet + 18 bytes</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <p>The 12 additional bytes include 6 bytes for the destination MAC address + 4 bytes for the VLAN + 2 bytes for the Ethernet type.</p> <p>For RED dropped, 6 bytes are added for the source MAC address.</p> |
|                                  | Packet forwarding component | <p>Queued: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p> <p>Transmitted: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | —                                                                                                                                                                                                            |
| Non-IQ PIC                       | Interface                   | <p>T Series, TX Series, T1600, and MX Series routers:</p> <ul style="list-style-type: none"> <li>• Queued: 478 bytes of Layer 3 packet.</li> <li>• Transmitted: 478 bytes of Layer 3 packet.</li> </ul> <p>T4000 routers with Type 5 FPCs :</p> <ul style="list-style-type: none"> <li>• Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead 8 bytes preamble + 12 bytes Inter frame Gap.</li> <li>• Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead 8 bytes preamble + 12 bytes Interframe Gap.</li> </ul> <p>M Series routers:</p> <ul style="list-style-type: none"> <li>• Queued: 478 bytes of Layer 3 packet.</li> <li>• Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead.</li> </ul> <p>PTX Series Packet Transport Routers:</p> <ul style="list-style-type: none"> <li>• Queued: The sum of the transmitted bytes and the RED dropped bytes.</li> <li>• Transmitted: Full Layer 2 overhead (including all L2 encapsulation and CRC) + 12 inter-packet gap + 8 for the preamble.</li> <li>• RED dropped: Full Layer 2 overhead (including all L2 encapsulation and CRC) + 12 inter-packet gap + 8 for the preamble (does not include the VLAN header or MPLS pushed bytes).</li> </ul> | <p>The Layer 2 overhead is 14 bytes for non-VLAN traffic and 18 bytes for VLAN traffic.</p>                                                                                                                  |

Table 641: Byte Count by Interface Hardware (*continued*)

| Interface Hardware                                   | Output Level                | Byte Count Includes                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Comments                                                                                                                           |
|------------------------------------------------------|-----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------|
| IQ and IQE PICs with a SONET/SDH interface           | Interface                   | <p>Queued: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes</p> <p>Transmitted: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes</p> <p>RED dropped: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes</p>                                                                                                                                                                                                                       | The additional 4 bytes are for the Layer 2 Point-to-Point Protocol (PPP) header.                                                   |
|                                                      | Packet forwarding component | <p>Queued: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p> <p>Transmitted: 486 bytes per packet, representing 478 bytes of Layer 3 packet + 8 bytes</p>                                                                                                                                                                                                                                                                                                                              | For transmitted packets, the additional 8 bytes includes 4 bytes for the PPP header and 4 bytes for a cookie.                      |
| Non-IQ PIC with a SONET/SDH interface                | Interface                   | <p>T Series, TX Series, T1600, and MX Series routers:</p> <ul style="list-style-type: none"> <li>Queued: 478 bytes of Layer 3 packet.</li> <li>Transmitted: 478 bytes of Layer 3 packet.</li> </ul> <p>M Series routers:</p> <ul style="list-style-type: none"> <li>Queued: 478 bytes of Layer 3 packet.</li> <li>Transmitted: 483 bytes per packet, representing 478 bytes of Layer 3 packet + 5 bytes</li> <li>RED dropped: 478 bytes per packet, representing 478 bytes of Layer 3 packet</li> </ul> | For transmitted packets, the additional 5 bytes includes 4 bytes for the PPP header and 1 byte for the packet loss priority (PLP). |
| Interfaces configured with Frame Relay Encapsulation | Interface                   | The default Frame Relay overhead is 7 bytes. If you configure the Frame Check Sequence (FCS) to 4 bytes, then the overhead increases to 10 bytes.                                                                                                                                                                                                                                                                                                                                                       |                                                                                                                                    |
| 1-port 10-Gigabit Ethernet IQ2 and IQ2-E PICs        | Interface                   | <p>Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including CRC.</p> <p>Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including CRC.</p>                                                                                                                                                                                                                                                                                                                        | The Layer 2 overhead is 18 bytes for non-VLAN traffic and 22 bytes for VLAN traffic.                                               |
| 4-port 1G IQ2 and IQ2-E PICs                         | Packet forwarding component | Queued: 478 bytes of Layer 3 packet.                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | —                                                                                                                                  |
| 8-port 1G IQ2 and IQ2-E PICs                         |                             | Transmitted: 478 bytes of Layer 3 packet.                                                                                                                                                                                                                                                                                                                                                                                                                                                               |                                                                                                                                    |

## Sample Output

### show interfaces queue (Rate-Limited Interface on a Gigabit Ethernet MIC in an MPC)

The following example shows queue information for the rate-limited interface ge-4/2/0 on a Gigabit Ethernet MIC in an MPC. For rate-limited queues for interfaces hosted on MICs or MPCs, rate-limit packet drops occur prior to packet output queuing. In the

command output, the nonzero statistics displayed in the **RL-dropped packets** and **RL-dropped bytes** fields quantify the traffic dropped to rate-limit queue 0 output to 10 percent of 1 gigabyte (100 megabits) per second. Because the RL-dropped traffic is not included in the **Queued** statistics, the statistics displayed for queued traffic are the same as the statistics for transmitted traffic.

```
user@host> show interfaces queue ge-4/2/0
Physical interface: ge-4/2/0, Enabled, Physical link is Up
 Interface index: 203, SNMP ifIndex: 1054
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
 Queued:
 Packets : 131300649 141751 pps
 Bytes : 11287964840 99793248 bps
 Transmitted:
 Packets : 131300649 141751 pps
 Bytes : 11287964840 99793248 bps
 Tail-dropped packets : 0 0 pps
 RL-dropped packets : 205050862 602295 pps
 RL-dropped bytes : 13595326612 327648832 bps
 RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
```

### show interfaces queue (Aggregated Ethernet on a T320 Router)

The following example shows that the aggregated Ethernet interface, **ae1**, has traffic on queues **af1** and **af12**:

```
user@host> show interfaces queue ae1
Physical interface: ae1, Enabled, Physical link is Up
 Interface index: 158, SNMP ifIndex: 33 Forwarding classes: 8 supported, 8 in use
Output queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: be
 Queued:
 Packets : 5 0 pps
 Bytes : 242 0 bps
 Transmitted:
 Packets : 5 0 pps
 Bytes : 242 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: af1
 Queued:
 Packets : 42603765 595484 pps
```

```

Bytes : 5453281920 609776496 bps
Transmitted:
Packets : 42603765 595484 pps
Bytes : 5453281920 609776496 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: ef1
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: nc
Queued:
Packets : 45 0 pps
Bytes : 3930 0 bps
Transmitted:
Packets : 45 0 pps
Bytes : 3930 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 4, Forwarding classes: af11
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 5, Forwarding classes: ef11
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 6, Forwarding classes: af12
Queued:
Packets : 31296413 437436 pps
Bytes : 4005940864 447935200 bps
Transmitted:
Packets : 31296413 437436 pps
Bytes : 4005940864 447935200 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 7, Forwarding classes: nc2
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps

```

```

Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

#### show interfaces queue (Gigabit Ethernet on a T640 Router)

```

user@host> show interfaces queue
Physical interface: ge-7/0/1, Enabled, Physical link is Up
 Interface index: 150, SNMP ifIndex: 42
 Forwarding classes: 8 supported, 8 in use
 Output queues: 8 supported, 8 in use
 Queue: 0, Forwarding classes: be
 Queued:
 Packets : 13 0 pps
 Bytes : 622 0 bps
 Transmitted:
 Packets : 13 0 pps
 Bytes : 622 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
 Queue: 1, Forwarding classes: af1
 Queued:
 Packets : 1725947945 372178 pps
 Bytes : 220921336960 381110432 bps
 Transmitted:
 Packets : 1725947945 372178 pps
 Bytes : 220921336960 381110432 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
 Queue: 2, Forwarding classes: ef1
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
 Queue: 3, Forwarding classes: nc
 Queued:
 Packets : 571 0 pps
 Bytes : 49318 336 bps
 Transmitted:
 Packets : 571 0 pps
 Bytes : 49318 336 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

```

#### show interfaces queue aggregate (Gigabit Ethernet Enhanced DPC)

```

user@host> show interfaces queue ge-2/2/9 aggregate

```

```

Physical interface: ge-2/2/9, Enabled, Physical link is Up
 Interface index: 238, SNMP ifIndex: 71
 Forwarding classes: 16 supported, 4 in use
 Ingress queues: 4 supported, 4 in use
 Queue: 0, Forwarding classes: best-effort
 Queued:
 Packets : 148450735 947295 pps
 Bytes : 8016344944 409228848 bps
 Transmitted:
 Packets : 76397439 487512 pps
 Bytes : 4125461868 210602376 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 72053285 459783 pps
 Low : 72053285 459783 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes : 3890877444 198626472 bps
 Low : 3890877444 198626472 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
 Queue: 1, Forwarding classes: expedited-forwarding
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
 Queue: 2, Forwarding classes: assured-forwarding
 Queued:
 Packets : 410278257 473940 pps
 Bytes : 22156199518 204742296 bps
 Transmitted:
 Packets : 4850003 4033 pps
 Bytes : 261900162 1742256 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 405425693 469907 pps
 Low : 405425693 469907 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes : 21892988124 203000040 bps
 Low : 21892988124 203000040 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
 Queue: 3, Forwarding classes: network-control
 Queued:
 Packets : 0 0 pps

```



```

Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
Low : 0 0 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 0 0 pps
RED-dropped bytes : 0 0 bps
Low : 0 0 bps
Medium-low : 0 0 bps
Medium-high : 0 0 bps
High : 0 0 bps
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets : 76605230 485376 pps
Bytes : 5209211400 264044560 bps
Transmitted:
Packets : 76444631 484336 pps
Bytes : 5198235612 263478800 bps
Tail-dropped packets : Not Available
RED-dropped packets : 160475 1040 pps
Low : 160475 1040 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 0 0 pps
RED-dropped bytes : 10912300 565760 bps
Low : 10912300 565760 bps
Medium-low : 0 0 bps
Medium-high : 0 0 bps
High : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
Low : 0 0 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 0 0 pps
RED-dropped bytes : 0 0 bps
Low : 0 0 bps
Medium-low : 0 0 bps
Medium-high : 0 0 bps
High : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets : 4836136 3912 pps
Bytes : 333402032 2139056 bps
Transmitted:
Packets : 3600866 1459 pps
Bytes : 244858888 793696 bps
Tail-dropped packets : Not Available

```

```

RED-dropped packets : 1225034 2450 pps
 Low : 1225034 2450 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 83302312 1333072 bps
 Low : 83302312 1333072 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps

```

#### Packet Forwarding Engine Chassis Queues:

Queues: 4 supported, 4 in use

Queue: 0, Forwarding classes: best-effort

```

Queued:
 Packets : 77059796 486384 pps
 Bytes : 3544750624 178989576 bps
Transmitted:
 Packets : 77059797 486381 pps
 Bytes : 3544750670 178988248 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps

```

Queue: 1, Forwarding classes: expedited-forwarding

```

Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps

```

```

 High : 0 0 pps
 RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
 Queued:
 Packets : 4846580 3934 pps
 Bytes : 222942680 1447768 bps
 Transmitted:
 Packets : 4846580 3934 pps
 Bytes : 222942680 1447768 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 3, Forwarding classes: network-control
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps

```

#### show interfaces queue (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-7/1/3
Physical interface: ge-7/1/3, Enabled, Physical link is Up
 Interface index: 170, SNMP ifIndex: 70 Forwarding classes: 16 supported, 4 in
 use Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
 Queued:
 Packets : 418390039 10 pps
 Bytes : 38910269752 7440 bps
 Transmitted:
 Packets : 418390039 10 pps
 Bytes : 38910269752 7440 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding

```

```

Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
 Packets : 7055 1 pps
 Bytes : 451552 512 bps
Transmitted:
 Packets : 7055 1 pps
 Bytes : 451552 512 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Forwarding classes: 16 supported, 4 in use Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
 Packets : 1031 0 pps
 Bytes : 143292 0 bps
Transmitted:
 Packets : 1031 0 pps
 Bytes : 143292 0 bps
 Tail-dropped packets : Not Available
 RL-dropped packets : 0 0 pps
 RL-dropped bytes : 0 0 bps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RL-dropped packets : 0 0 pps
 RL-dropped bytes : 0 0 bps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps

```

```

Tail-dropped packets : Not Available
RL-dropped packets : 0 0 pps
RL-dropped bytes : 0 0 bps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets : 77009 11 pps
Bytes : 6894286 7888 bps
Transmitted:
Packets : 77009 11 pps
Bytes : 6894286 7888 bps
Tail-dropped packets : Not Available
RL-dropped packets : 0 0 pps
RL-dropped bytes : 0 0 bps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets : 1031 0 pps
Bytes : 147328 0 bps
Transmitted:
Packets : 1031 0 pps
Bytes : 147328 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
Low, non-TCP : 0 0 pps
Low, TCP : 0 0 pps
High, non-TCP : 0 0 pps
High, TCP : 0 0 pps
RED-dropped bytes : 0 0 bps
Low, non-TCP : 0 0 bps
Low, TCP : 0 0 bps
High, non-TCP : 0 0 bps
High, TCP : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
Low, non-TCP : 0 0 pps
Low, TCP : 0 0 pps
High, non-TCP : 0 0 pps
High, TCP : 0 0 pps
RED-dropped bytes : 0 0 bps
Low, non-TCP : 0 0 bps
Low, TCP : 0 0 bps
High, non-TCP : 0 0 bps
High, TCP : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:

```

```

Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
 Low, non-TCP : 0 0 pps
 Low, TCP : 0 0 pps
 High, non-TCP : 0 0 pps
 High, TCP : 0 0 pps
RED-dropped bytes : 0 0 bps
 Low, non-TCP : 0 0 bps
 Low, TCP : 0 0 bps
 High, non-TCP : 0 0 bps
 High, TCP : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
 Packets : 94386 12 pps
 Bytes : 13756799 9568 bps
Transmitted:
 Packets : 94386 12 pps
 Bytes : 13756799 9568 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
 Low, non-TCP : 0 0 pps
 Low, TCP : 0 0 pps
 High, non-TCP : 0 0 pps
 High, TCP : 0 0 pps
RED-dropped bytes : 0 0 bps
 Low, non-TCP : 0 0 bps
 Low, TCP : 0 0 bps
 High, non-TCP : 0 0 bps
 High, TCP : 0 0 bps

```

#### show interfaces queue both-ingress-egress (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-6/2/0 both-ingress-egress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
 Packets : Not Available
 Bytes : 0 0 bps
Transmitted:
 Packets : 254 0 pps
 Bytes : 16274 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
 Packets : Not Available
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding

```

```

Queued:
 Packets : Not Available
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
 Packets : Not Available
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
 Packets : Not Available
 Bytes : 0 0 bps
Transmitted:
 Packets : 3 0 pps
 Bytes : 126 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
 Packets : Not Available
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
 Packets : Not Available
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
 Packets : Not Available
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

```

```

Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
 Queued:
 Packets : 80564692 0 pps
 Bytes : 3383717100 0 bps
 Transmitted:
 Packets : 80564692 0 pps
 Bytes : 3383717100 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
 Queued:
 Packets : 80564685 0 pps
 Bytes : 3383716770 0 bps
 Transmitted:
 Packets : 80564685 0 pps
 Bytes : 3383716770 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
 Queued:
 Packets : 9397 0 pps
 Bytes : 3809052 232 bps
 Transmitted:
 Packets : 9397 0 pps
 Bytes : 3809052 232 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

```

### show interfaces queue ingress (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-6/2/0 ingress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
 Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
 Queued:
 Packets : Not Available
 Bytes : 0 0 bps
 Transmitted:
 Packets : 288 0 pps
 Bytes : 18450 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

```



```

Queue: 1, Forwarding classes: expedited-forwarding
 Queued:
 Packets : Not Available
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
 Queued:
 Packets : Not Available
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
 Queued:
 Packets : Not Available
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

```

#### show interfaces queue egress (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-6/2/0 egress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
 Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
 Queued:
 Packets : Not Available
 Bytes : 0 0 bps
 Transmitted:
 Packets : 3 0 pps
 Bytes : 126 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
 Queued:
 Packets : Not Available
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
 Queued:

```

```

Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets : 80564692 0 pps
Bytes : 3383717100 0 bps
Transmitted:
Packets : 80564692 0 pps
Bytes : 3383717100 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
Packets : 80564685 0 pps
Bytes : 3383716770 0 bps
Transmitted:
Packets : 80564685 0 pps
Bytes : 3383716770 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets : 9538 0 pps
Bytes : 3819840 0 bps
Transmitted:
Packets : 9538 0 pps
Bytes : 3819840 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

## show interfaces queue remaining-traffic (Gigabit Ethernet Enhanced DPC)

```

user@host> show interfaces queue ge-2/2/9 remaining-traffic
Physical interface: ge-2/2/9, Enabled, Physical link is Up
 Interface index: 238, SNMP ifIndex: 71
Forwarding classes: 16 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
 Queued:
 Packets : 110208969 472875 pps
 Bytes : 5951284434 204282000 bps
 Transmitted:
 Packets : 110208969 472875 pps
 Bytes : 5951284434 204282000 bps
 Tail-dropped packets : Not Available
 RED-dropped packets :
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes :
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets :
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes :
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets :
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes :
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps

```

```

 High : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
 Packets : 109355853 471736 pps
 Bytes : 7436199152 256627968 bps
Transmitted:
 Packets : 109355852 471736 pps
 Bytes : 7436198640 256627968 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps

```

```

Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps

```

#### show interfaces queue (Channelized OC12 IQE Type 3 PIC in SONET Mode)

```

user@host> show interfaces queue t3-1/1/0:7
Physical interface: t3-1/1/0:7, Enabled, Physical link is Up

Interface index: 192, SNMP ifIndex: 1948

Description: full T3 interface connect to 6ce13 t3-3/1/0:7 for FR testing -
Lam

Forwarding classes: 16 supported, 9 in use

Egress queues: 8 supported, 8 in use

Queue: 0, Forwarding classes: DEFAULT

Queued:

Packets : 214886 13449 pps
Bytes : 9884756 5164536 bps

Transmitted:

Packets : 214886 13449 pps
Bytes : 9884756 5164536 bps

```

|                        |   |       |
|------------------------|---|-------|
| Tail-dropped packets : | 0 | 0 pps |
| RED-dropped packets :  | 0 | 0 pps |
| Low :                  | 0 | 0 pps |
| Medium-low :           | 0 | 0 pps |
| Medium-high :          | 0 | 0 pps |
| High :                 | 0 | 0 pps |
| RED-dropped bytes :    | 0 | 0 bps |
| Low :                  | 0 | 0 bps |
| Medium-low :           | 0 | 0 bps |
| Medium-high :          | 0 | 0 bps |
| High :                 | 0 | 0 bps |

Queue: 1, Forwarding classes: REALTIME

Queued:

|           |   |       |
|-----------|---|-------|
| Packets : | 0 | 0 pps |
| Bytes :   | 0 | 0 bps |

Transmitted:

|                        |   |       |
|------------------------|---|-------|
| Packets :              | 0 | 0 pps |
| Bytes :                | 0 | 0 bps |
| Tail-dropped packets : | 0 | 0 pps |
| RED-dropped packets :  | 0 | 0 pps |
| Low :                  | 0 | 0 pps |
| Medium-low :           | 0 | 0 pps |
| Medium-high :          | 0 | 0 pps |
| High :                 | 0 | 0 pps |
| RED-dropped bytes :    | 0 | 0 bps |
| Low :                  | 0 | 0 bps |
| Medium-low :           | 0 | 0 bps |
| Medium-high :          | 0 | 0 bps |
| High :                 | 0 | 0 bps |

Queue: 2, Forwarding classes: PRIVATE

## Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
|---------|---|---|-------|

|       |   |   |       |
|-------|---|---|-------|
| Bytes | : | 0 | 0 bps |
|-------|---|---|-------|

## Transmitted:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
|---------|---|---|-------|

|       |   |   |       |
|-------|---|---|-------|
| Bytes | : | 0 | 0 bps |
|-------|---|---|-------|

|                      |   |   |       |
|----------------------|---|---|-------|
| Tail-dropped packets | : | 0 | 0 pps |
|----------------------|---|---|-------|

|                     |   |   |       |
|---------------------|---|---|-------|
| RED-dropped packets | : | 0 | 0 pps |
|---------------------|---|---|-------|

|     |   |   |       |
|-----|---|---|-------|
| Low | : | 0 | 0 pps |
|-----|---|---|-------|

|            |   |   |       |
|------------|---|---|-------|
| Medium-low | : | 0 | 0 pps |
|------------|---|---|-------|

|             |   |   |       |
|-------------|---|---|-------|
| Medium-high | : | 0 | 0 pps |
|-------------|---|---|-------|

|      |   |   |       |
|------|---|---|-------|
| High | : | 0 | 0 pps |
|------|---|---|-------|

|                   |   |   |       |
|-------------------|---|---|-------|
| RED-dropped bytes | : | 0 | 0 bps |
|-------------------|---|---|-------|

|     |   |   |       |
|-----|---|---|-------|
| Low | : | 0 | 0 bps |
|-----|---|---|-------|

|            |   |   |       |
|------------|---|---|-------|
| Medium-low | : | 0 | 0 bps |
|------------|---|---|-------|

|             |   |   |       |
|-------------|---|---|-------|
| Medium-high | : | 0 | 0 bps |
|-------------|---|---|-------|

|      |   |   |       |
|------|---|---|-------|
| High | : | 0 | 0 bps |
|------|---|---|-------|

## Queue: 3, Forwarding classes: CONTROL

## Queued:

|         |   |    |       |
|---------|---|----|-------|
| Packets | : | 60 | 0 pps |
|---------|---|----|-------|

|       |   |      |       |
|-------|---|------|-------|
| Bytes | : | 4560 | 0 bps |
|-------|---|------|-------|

## Transmitted:

|         |   |    |       |
|---------|---|----|-------|
| Packets | : | 60 | 0 pps |
|---------|---|----|-------|

|       |   |      |       |
|-------|---|------|-------|
| Bytes | : | 4560 | 0 bps |
|-------|---|------|-------|

|                      |   |   |       |
|----------------------|---|---|-------|
| Tail-dropped packets | : | 0 | 0 pps |
|----------------------|---|---|-------|

|                     |   |   |       |
|---------------------|---|---|-------|
| RED-dropped packets | : | 0 | 0 pps |
|---------------------|---|---|-------|

|     |   |   |       |
|-----|---|---|-------|
| Low | : | 0 | 0 pps |
|-----|---|---|-------|

|            |   |   |       |
|------------|---|---|-------|
| Medium-low | : | 0 | 0 pps |
|------------|---|---|-------|

|             |   |   |       |
|-------------|---|---|-------|
| Medium-high | : | 0 | 0 pps |
|-------------|---|---|-------|

|      |   |   |       |
|------|---|---|-------|
| High | : | 0 | 0 pps |
|------|---|---|-------|

|                   |   |   |       |
|-------------------|---|---|-------|
| RED-dropped bytes | : | 0 | 0 bps |
|-------------------|---|---|-------|

|             |   |   |       |
|-------------|---|---|-------|
| Low         | : | 0 | 0 bps |
| Medium-low  | : | 0 | 0 bps |
| Medium-high | : | 0 | 0 bps |
| High        | : | 0 | 0 bps |

Queue: 4, Forwarding classes: CLASS\_B\_OUTPUT

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |
| Low                  | : | 0 | 0 bps |
| Medium-low           | : | 0 | 0 bps |
| Medium-high          | : | 0 | 0 bps |
| High                 | : | 0 | 0 bps |

Queue: 5, Forwarding classes: CLASS\_C\_OUTPUT

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |



|                     |   |   |       |
|---------------------|---|---|-------|
| RED-dropped packets | : | 0 | 0 pps |
| Low                 | : | 0 | 0 pps |
| Medium-low          | : | 0 | 0 pps |
| Medium-high         | : | 0 | 0 pps |
| High                | : | 0 | 0 pps |
| RED-dropped bytes   | : | 0 | 0 bps |
| Low                 | : | 0 | 0 bps |
| Medium-low          | : | 0 | 0 bps |
| Medium-high         | : | 0 | 0 bps |
| High                | : | 0 | 0 bps |

Queue: 6, Forwarding classes: CLASS\_V\_OUTPUT

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |
| Low                  | : | 0 | 0 bps |
| Medium-low           | : | 0 | 0 bps |
| Medium-high          | : | 0 | 0 bps |
| High                 | : | 0 | 0 bps |

Queue: 7, Forwarding classes: CLASS\_S\_OUTPUT, GETS

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
|---------|---|---|-------|

|                      |   |   |       |
|----------------------|---|---|-------|
| Bytes                | : | 0 | 0 bps |
| Transmitted:         |   |   |       |
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |
| Low                  | : | 0 | 0 bps |
| Medium-low           | : | 0 | 0 bps |
| Medium-high          | : | 0 | 0 bps |
| High                 | : | 0 | 0 bps |

Packet Forwarding Engine Chassis Queues:

Queues: 8 supported, 8 in use

Queue: 0, Forwarding classes: DEFAULT

Queued:

|         |   |          |             |
|---------|---|----------|-------------|
| Packets | : | 371365   | 23620 pps   |
| Bytes   | : | 15597330 | 7936368 bps |

Transmitted:

|                      |   |          |             |
|----------------------|---|----------|-------------|
| Packets              | : | 371365   | 23620 pps   |
| Bytes                | : | 15597330 | 7936368 bps |
| Tail-dropped packets | : | 0        | 0 pps       |
| RED-dropped packets  | : | 0        | 0 pps       |
| Low                  | : | 0        | 0 pps       |
| Medium-low           | : | 0        | 0 pps       |
| Medium-high          | : | 0        | 0 pps       |

|                   |   |   |       |
|-------------------|---|---|-------|
| High              | : | 0 | 0 pps |
| RED-dropped bytes | : | 0 | 0 bps |
| Low               | : | 0 | 0 bps |
| Medium-low        | : | 0 | 0 bps |
| Medium-high       | : | 0 | 0 bps |
| High              | : | 0 | 0 bps |

Queue: 1, Forwarding classes: REALTIME

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |
| Low                  | : | 0 | 0 bps |
| Medium-low           | : | 0 | 0 bps |
| Medium-high          | : | 0 | 0 bps |
| High                 | : | 0 | 0 bps |

Queue: 2, Forwarding classes: PRIVATE

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Transmitted:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

|                        |   |       |
|------------------------|---|-------|
| Tail-dropped packets : | 0 | 0 pps |
| RED-dropped packets :  | 0 | 0 pps |
| Low :                  | 0 | 0 pps |
| Medium-low :           | 0 | 0 pps |
| Medium-high :          | 0 | 0 pps |
| High :                 | 0 | 0 pps |
| RED-dropped bytes :    | 0 | 0 bps |
| Low :                  | 0 | 0 bps |
| Medium-low :           | 0 | 0 bps |
| Medium-high :          | 0 | 0 bps |
| High :                 | 0 | 0 bps |

Queue: 3, Forwarding classes: CONTROL

Queued:

|           |         |        |
|-----------|---------|--------|
| Packets : | 32843   | 0 pps  |
| Bytes :   | 2641754 | 56 bps |

Transmitted:

|                        |         |        |
|------------------------|---------|--------|
| Packets :              | 32843   | 0 pps  |
| Bytes :                | 2641754 | 56 bps |
| Tail-dropped packets : | 0       | 0 pps  |
| RED-dropped packets :  | 0       | 0 pps  |
| Low :                  | 0       | 0 pps  |
| Medium-low :           | 0       | 0 pps  |
| Medium-high :          | 0       | 0 pps  |
| High :                 | 0       | 0 pps  |
| RED-dropped bytes :    | 0       | 0 bps  |
| Low :                  | 0       | 0 bps  |
| Medium-low :           | 0       | 0 bps  |
| Medium-high :          | 0       | 0 bps  |
| High :                 | 0       | 0 bps  |

Queue: 4, Forwarding classes: CLASS\_B\_OUTPUT

## Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
|---------|---|---|-------|

|       |   |   |       |
|-------|---|---|-------|
| Bytes | : | 0 | 0 bps |
|-------|---|---|-------|

## Transmitted:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
|---------|---|---|-------|

|       |   |   |       |
|-------|---|---|-------|
| Bytes | : | 0 | 0 bps |
|-------|---|---|-------|

|                      |   |   |       |
|----------------------|---|---|-------|
| Tail-dropped packets | : | 0 | 0 pps |
|----------------------|---|---|-------|

|                     |   |   |       |
|---------------------|---|---|-------|
| RED-dropped packets | : | 0 | 0 pps |
|---------------------|---|---|-------|

|     |   |   |       |
|-----|---|---|-------|
| Low | : | 0 | 0 pps |
|-----|---|---|-------|

|            |   |   |       |
|------------|---|---|-------|
| Medium-low | : | 0 | 0 pps |
|------------|---|---|-------|

|             |   |   |       |
|-------------|---|---|-------|
| Medium-high | : | 0 | 0 pps |
|-------------|---|---|-------|

|      |   |   |       |
|------|---|---|-------|
| High | : | 0 | 0 pps |
|------|---|---|-------|

|                   |   |   |       |
|-------------------|---|---|-------|
| RED-dropped bytes | : | 0 | 0 bps |
|-------------------|---|---|-------|

|     |   |   |       |
|-----|---|---|-------|
| Low | : | 0 | 0 bps |
|-----|---|---|-------|

|            |   |   |       |
|------------|---|---|-------|
| Medium-low | : | 0 | 0 bps |
|------------|---|---|-------|

|             |   |   |       |
|-------------|---|---|-------|
| Medium-high | : | 0 | 0 bps |
|-------------|---|---|-------|

|      |   |   |       |
|------|---|---|-------|
| High | : | 0 | 0 bps |
|------|---|---|-------|

## Queue: 5, Forwarding classes: CLASS\_C\_OUTPUT

## Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
|---------|---|---|-------|

|       |   |   |       |
|-------|---|---|-------|
| Bytes | : | 0 | 0 bps |
|-------|---|---|-------|

## Transmitted:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
|---------|---|---|-------|

|       |   |   |       |
|-------|---|---|-------|
| Bytes | : | 0 | 0 bps |
|-------|---|---|-------|

|                      |   |   |       |
|----------------------|---|---|-------|
| Tail-dropped packets | : | 0 | 0 pps |
|----------------------|---|---|-------|

|                     |   |   |       |
|---------------------|---|---|-------|
| RED-dropped packets | : | 0 | 0 pps |
|---------------------|---|---|-------|

|     |   |   |       |
|-----|---|---|-------|
| Low | : | 0 | 0 pps |
|-----|---|---|-------|

|            |   |   |       |
|------------|---|---|-------|
| Medium-low | : | 0 | 0 pps |
|------------|---|---|-------|

|             |   |   |       |
|-------------|---|---|-------|
| Medium-high | : | 0 | 0 pps |
|-------------|---|---|-------|

|      |   |   |       |
|------|---|---|-------|
| High | : | 0 | 0 pps |
|------|---|---|-------|

|                   |   |   |       |
|-------------------|---|---|-------|
| RED-dropped bytes | : | 0 | 0 bps |
|-------------------|---|---|-------|

|             |   |   |       |
|-------------|---|---|-------|
| Low         | : | 0 | 0 bps |
| Medium-low  | : | 0 | 0 bps |
| Medium-high | : | 0 | 0 bps |
| High        | : | 0 | 0 bps |

Queue: 6, Forwarding classes: CLASS\_V\_OUTPUT

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |
| Low                  | : | 0 | 0 bps |
| Medium-low           | : | 0 | 0 bps |
| Medium-high          | : | 0 | 0 bps |
| High                 | : | 0 | 0 bps |

Queue: 7, Forwarding classes: CLASS\_S\_OUTPUT, GETS

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |

|                     |   |   |       |
|---------------------|---|---|-------|
| RED-dropped packets | : | 0 | 0 pps |
| Low                 | : | 0 | 0 pps |
| Medium-low          | : | 0 | 0 pps |
| Medium-high         | : | 0 | 0 pps |
| High                | : | 0 | 0 pps |
| RED-dropped bytes   | : | 0 | 0 bps |
| Low                 | : | 0 | 0 bps |
| Medium-low          | : | 0 | 0 bps |
| Medium-high         | : | 0 | 0 bps |
| High                | : | 0 | 0 bps |

### show interfaces queue (QFX Series)

```

user@switch> show interfaces queue xe-0/0/15
Physical interface: xe-0/0/15, Enabled, Physical link is Up
Interface index: 49165, SNMP ifIndex: 539
Forwarding classes: 12 supported, 8 in use
Egress queues: 12 supported, 8 in use
Queue: 0, Forwarding classes: best-effort
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped packets: 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: fcoe
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped packets: 0 0 pps
 Total-dropped bytes : 0 0 bps
0 bps
Queue: 4, Forwarding classes: no-loss
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped packets: 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 7, Forwarding classes: network-control

```

```

Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped packets: 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 8, Forwarding classes: mcast
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped packets: 0 0 pps
 Total-dropped bytes : 0 0 bps

```

#### show interfaces queue l2-statistics (lsq interface)

```

user@switch> show interfaces queue lsq-2/2/0.2 l2-statistics
Logical interface lsq-2/2/0.2 (Index 69) (SNMP ifIndex 1598)
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Burst size: 0
Queue: 0, Forwarding classes: be
Queued:
 Packets : 1 0 pps
 Bytes : 1001 0 bps
Transmitted:
 Packets : 5 0 pps
 Bytes : 1062 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: ef
Queued:
 Packets : 1 0 pps
 Bytes : 1500 0 bps
Transmitted:
 Packets : 6 0 pps
 Bytes : 1573 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: af
Queued:
 Packets : 1 0 pps
 Bytes : 512 0 bps
Transmitted:
 Packets : 3 0 pps
 Bytes : 549 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: nc
Queued:
 Packets : 0 0 pps

```



```

Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
=====

```

### show interfaces queue lsq (lsq-afd)

```

user@switch> show interfaces queue lsq-1/0/0
Logical interface lsq-1/0/0 (Index 348) (SNMP ifIndex 660)
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Burst size: 0
Queue: 0, Forwarding classes: be
 Queued:
 Packets : 55576 1206 pps
 Bytes : 29622008 5145472 bps
 Transmitted:
 Packets : 55576 1206 pps
 Bytes : 29622008 5145472 bps
 Tail-dropped packets : 0 0 pps
 RL-dropped packets : 0 0 pps
 RL-dropped bytes : 0 0 bps
 RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 1, Forwarding classes: ef
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : 0 0 pps
 RL-dropped packets : 0 0 pps
 RL-dropped bytes : 0 0 bps
 RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 2, Forwarding classes: af
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps

```

```
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : 0 0 pps
RL-dropped packets : 0 0 pps
RL-dropped bytes : 0 0 bps
RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 3, Forwarding classes: nc
Queued:
Packets : 22231 482 pps
Bytes : 11849123 2057600 bps
Transmitted:
Packets : 22231 482 pps
Bytes : 11849123 2057600 bps
Tail-dropped packets : 0 0 pps
RL-dropped packets : 0 0 pps
RL-dropped bytes : 0 0 bps
RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
```

## show interfaces voq

**Syntax** `show interfaces voq interface-name  
<forwarding-class forwarding-class-name>  
<non-zero>  
<source-fpc source-fpc-number>`

**Release Information** Command introduced in Junos OS Release 14.1 for the PTX Series Routers  
Command introduced in Junos OS Release 15.1X53-D20 for QFX10000 switches.

**Description** Display the random early detection (RED) drop statistics from all ingress Packet Forwarding Engines associated with the specified physical egress interface. In the VOQ architecture, egress output queues (shallow buffers) buffer data in virtual queues on ingress Packet Forwarding Engines. In cases of congestion, you can use this command to identify which ingress Packet Forwarding Engine is the source of RED-dropped packets contributing to congestion.



**NOTE:** On the PTX Series routers and QFX10000 switches, these statistics include tail-dropped packets.

**Options** `interface interface-name`—Display the ingress VOQ RED drop statistics for the specified egress interface.

`forwarding-class forwarding-class-name`—Display VOQ RED drop statistics for a specified forwarding class.

`non-zero`—Display only non-zero VOQ RED drop statistics counters.

`source-fpc source-fpc-number`—Display VOQ RED drop statistics for the specified source FPC.

**Additional Information**

- On PTX Series routers, you can display VOQ statistics for only the WAN physical interface.
- VOQ statistics for aggregated physical interfaces are not supported. Statistics for an aggregated interface are the summation of the queue statistics of the child links of that aggregated interface. You can use the **show interfaces queue** command to identify the child link which is experiencing congestion and then view the VOQ statistics on the respective child link using the **show interfaces voq** command.

For information on virtual output queuing on PTX routers, see *Understanding Virtual Output Queues on PTX Series Packet Transport Routers*. For information on virtual output queueing on QFX10000 switches, see [“Understanding CoS Virtual Output Queues \(VOQs\) on QFX10000 Switches” on page 6724](#).

**Required Privilege Level** view

- Related Documentation**
- [Understanding Virtual Output Queues on PTX Series Packet Transport Routers](#)
  - [Understanding CoS Virtual Output Queues \(VOQs\) on QFX10000 Switches on page 6724](#)

**List of Sample Output**

[show interfaces voq \(For a Specific Physical Interface\) \(PTX Series Routers\) on page 7381](#)  
[show interfaces voq \(For a Specific Physical Interface\) \(QFX10000 Switches\) on page 7385](#)  
[show interfaces voq et-7/0/0 \(For a Specific Forwarding Class\) on page 7387](#)  
[show interfaces voq et-5/0/12 \(For a Specific Source FPC\) on page 7388](#)  
[show interfaces voq et-5/0/12 \(For a Specific Forwarding Class and Source FPC\) on page 7390](#)  
[show interfaces voq et-7/0/0 \(Non-Zero\) on page 7390](#)  
[show interfaces voq et-7/0/0 \(For a Specific Forwarding Class and Non-Zero\) on page 7391](#)

**Output Fields** [Table 615](#) lists the output fields for the show interfaces queue command. Output fields are listed in the approximate order in which they appear.

**Table 642: show interfaces voq Output Fields**

| Field Name          | Field Description                                                                                                                                                                            |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Physical interface  | Name of the physical interface.                                                                                                                                                              |
| Enabled             | State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .                                                         |
| Interface index     | Physical interface's index number, which reflects its initialization sequence.                                                                                                               |
| SNMP ifIndex        | SNMP index number for the interface.                                                                                                                                                         |
| Queue               | Egress queue number.                                                                                                                                                                         |
| Forwarding classes  | Forwarding class name.                                                                                                                                                                       |
| FPC number          | Number of the Flexible PIC Concentrator (FPC) located on ingress.                                                                                                                            |
| PFE                 | Number of the Packet Forwarding Engine providing virtual output queues on the ingress.                                                                                                       |
| RED-dropped packets | Number of packets per second (pps) dropped because of random early detection (RED).<br><br><b>NOTE:</b> On the PTX Series routers, these statistics include tail-dropped packets.            |
| RED-dropped bytes   | Number of bytes per second dropped because of RED. The byte counts vary by interface hardware.<br><br><b>NOTE:</b> On the PTX Series routers, these statistics include tail-dropped packets. |

## Sample Output

### show interfaces voq (For a Specific Physical Interface) (PTX Series Routers)

The following example shows ingress RED-dropped statistics for the egress Ethernet interface configured on port 0 of Physical Interface Card (PIC) 0, located on the FPC in slot 7.

The sample output below shows that the cause of the congestion is ingress Packet Forwarding Engine PFE 0, which resides on FPC number 4, as denoted by the count of RED-dropped packets and RED-dropped bytes for egress queue 0, forwarding classes best-effort and egress queue 3, forwarding class network control.

```

user@host> show interfaces voq et-7/0/0
Physical interface: et-7/0/0, Enabled, Physical link is Up
 Interface index: 155, SNMP ifIndex: 699

Queue: 0, Forwarding classes: best-effort

FPC number: 1
 PFE: 0
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
 PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
 PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
 PFE: 3
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

FPC number: 4
 PFE: 0
 RED-dropped packets : 19969426 2323178 pps
 RED-dropped bytes : 2196636860 2044397464 bps
 PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
 PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
 PFE: 3
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

FPC number: 6
 PFE: 0
 RED-dropped packets : 19969424 2321205 pps
 RED-dropped bytes : 2196636640 2042660808 bps
 PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
 PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
 PFE: 3
 RED-dropped packets : 0 0 pps

```

```

 RED-dropped bytes : 0 0 bps
PFE: 4
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 5
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 6
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 7
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

FPC number: 7
PFE: 0
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 3
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

Queue: 1, Forwarding classes: expedited-forwarding

FPC number: 1
PFE: 0
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 3
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

FPC number: 4
PFE: 0
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 3
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

FPC number: 6
PFE: 0
 RED-dropped packets : 0 0 pps

```

```

 RED-dropped bytes : 0 0 bps
PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 3
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 4
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 5
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 6
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 7
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

FPC number: 7
PFE: 0
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 3
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

Queue: 2, Forwarding classes: assured-forwarding

FPC number: 1
PFE: 0
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 3
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

FPC number: 4
PFE: 0
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 2

```

```

 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 3
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

FPC number: 6
PFE: 0
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 3
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 4
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 5
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 6
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 7
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

FPC number: 7
PFE: 0
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 3
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

Queue: 3, Forwarding classes: network-control

FPC number: 1
PFE: 0
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 3
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
```



```

FPC number: 4
PFE: 0
 RED-dropped packets : 16338670 1900314 pps
 RED-dropped bytes : 1797253700 1672276976 bps
PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 3
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

FPC number: 6
PFE: 0
 RED-dropped packets : 16338698 1899163 pps
 RED-dropped bytes : 1797256780 1671263512 bps
PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 3
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 4
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 5
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 6
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 7
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

FPC number: 7
PFE: 0
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 3
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

```

#### show interfaces voq (For a Specific Physical Interface) (QFX10000 Switches)

The sample output below shows congestion on ingress PFE 1 on FPC number 0, and on ingress PFE 2 on FPC number 1, as denoted by the count of RED-dropped packets and RED-dropped bytes for best-effort egress queue 0.

```
user@host> show interfaces voq et-1/0/0
```

```
Physical interface: et-1/0/0, Enabled, Physical link is Up
```

```
Interface index: 659, SNMP ifIndex: 539
```

```
Queue: 0, Forwarding classes: best-effort
```

```
FPC number: 0
```

```
PFE: 0
```

```
RED-dropped packets : 0 0 pps
```

```
RED-dropped bytes : 0 0 bps
```

```
PFE: 1
```

```
RED-dropped packets : 411063248 16891870 pps
```

```
RED-dropped bytes : 52616095744 17297275600 bps
```

```
PFE: 2
```

```
RED-dropped packets : 0 0 pps
```

```
RED-dropped bytes : 0 0 bps
```

```
FPC number: 1
```

```
PFE: 0
```

```
RED-dropped packets : 0 0 pps
```

```
RED-dropped bytes : 0 0 bps
```

```
PFE: 1
```

```
RED-dropped packets : 0 0 pps
```

```
RED-dropped bytes : 0 0 bps
```

```
PFE: 2
```

```
RED-dropped packets : 411063012 16891870 pps
```

```
RED-dropped bytes : 52616065536 17297275376 bps
```

```
Queue: 3, Forwarding classes: fcoe
```

```
FPC number: 0
```

```
PFE: 0
```

```
RED-dropped packets : 0 0 pps
```

```
RED-dropped bytes : 0 0 bps
```

```
PFE: 1
```

```
RED-dropped packets : 0 0 pps
```

```
RED-dropped bytes : 0 0 bps
```

```
PFE: 2
```

```
RED-dropped packets : 0 0 pps
```

```
RED-dropped bytes : 0 0 bps
```

```
FPC number: 1
```

```
PFE: 0
```

```
RED-dropped packets : 0 0 pps
```

```
RED-dropped bytes : 0 0 bps
```

```
PFE: 1
```

```
RED-dropped packets : 0 0 pps
```

```
RED-dropped bytes : 0 0 bps
```

```
PFE: 2
```

```
RED-dropped packets : 0 0 pps
```

```
RED-dropped bytes : 0 0 bps
```

```
Queue: 4, Forwarding classes: no-loss
```

```
FPC number: 0
```

```
PFE: 0
```

```
RED-dropped packets : 0 0 pps
```

```
RED-dropped bytes : 0 0 bps
```

```
PFE: 1
```

```
RED-dropped packets : 0 0 pps
```

```
RED-dropped bytes : 0 0 bps
```

```

PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

FPC number: 1
PFE: 0
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

Queue: 7, Forwarding classes: network-control

FPC number: 0
PFE: 0
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

FPC number: 1
PFE: 0
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

```

#### show interfaces voq et-7/0/0 (For a Specific Forwarding Class)

```

user@host> show interfaces voq et-7/0/0 forwarding-class best-effort
Physical interface: et-7/0/0, Enabled, Physical link is Up
Interface index: 155, SNMP ifIndex: 699

```

```

Queue: 0, Forwarding classes: best-effort

FPC number: 1
PFE: 0
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 3
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

```

```

FPC number: 4
PFE: 0
 RED-dropped packets : 66604786 2321519 pps
 RED-dropped bytes : 7326526460 2042936776 bps
PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 3
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

FPC number: 6
PFE: 0
 RED-dropped packets : 66604794 371200 pps
 RED-dropped bytes : 7326527340 326656000 bps
PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 3
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 4
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 5
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 6
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 7
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

FPC number: 7
PFE: 0
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 1
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 2
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
PFE: 3
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

```

#### show interfaces voq et-5/0/12 (For a Specific Source FPC)

```

user@host> show interfaces voq et-5/0/12 source-fpc 0
Physical interface: et-5/0/12, Enabled, Physical link is Up
 Interface index: 166, SNMP ifIndex: 1104

Queue: 0, Forwarding classes: best-effort

```

FPC number: 0

PFE: 0

|                       |   |       |
|-----------------------|---|-------|
| RED-dropped packets : | 0 | 0 pps |
| RED-dropped bytes :   | 0 | 0 bps |

PFE: 1

|                       |   |       |
|-----------------------|---|-------|
| RED-dropped packets : | 0 | 0 pps |
| RED-dropped bytes :   | 0 | 0 bps |

PFE: 2

|                       |   |       |
|-----------------------|---|-------|
| RED-dropped packets : | 0 | 0 pps |
| RED-dropped bytes :   | 0 | 0 bps |

PFE: 3

|                       |   |       |
|-----------------------|---|-------|
| RED-dropped packets : | 0 | 0 pps |
| RED-dropped bytes :   | 0 | 0 bps |

Queue: 1, Forwarding classes: expedited-forwarding

FPC number: 0

PFE: 0

|                       |   |       |
|-----------------------|---|-------|
| RED-dropped packets : | 0 | 0 pps |
| RED-dropped bytes :   | 0 | 0 bps |

PFE: 1

|                       |   |       |
|-----------------------|---|-------|
| RED-dropped packets : | 0 | 0 pps |
| RED-dropped bytes :   | 0 | 0 bps |

PFE: 2

|                       |   |       |
|-----------------------|---|-------|
| RED-dropped packets : | 0 | 0 pps |
| RED-dropped bytes :   | 0 | 0 bps |

PFE: 3

|                       |   |       |
|-----------------------|---|-------|
| RED-dropped packets : | 0 | 0 pps |
| RED-dropped bytes :   | 0 | 0 bps |

Queue: 2, Forwarding classes: assured-forwarding

FPC number: 0

PFE: 0

|                       |   |       |
|-----------------------|---|-------|
| RED-dropped packets : | 0 | 0 pps |
| RED-dropped bytes :   | 0 | 0 bps |

PFE: 1

|                       |   |       |
|-----------------------|---|-------|
| RED-dropped packets : | 0 | 0 pps |
| RED-dropped bytes :   | 0 | 0 bps |

PFE: 2

|                       |   |       |
|-----------------------|---|-------|
| RED-dropped packets : | 0 | 0 pps |
| RED-dropped bytes :   | 0 | 0 bps |

PFE: 3

|                       |   |       |
|-----------------------|---|-------|
| RED-dropped packets : | 0 | 0 pps |
| RED-dropped bytes :   | 0 | 0 bps |

Queue: 3, Forwarding classes: network-control

FPC number: 0

PFE: 0

|                       |   |       |
|-----------------------|---|-------|
| RED-dropped packets : | 0 | 0 pps |
| RED-dropped bytes :   | 0 | 0 bps |

PFE: 1

|                       |   |       |
|-----------------------|---|-------|
| RED-dropped packets : | 0 | 0 pps |
| RED-dropped bytes :   | 0 | 0 bps |

PFE: 2

|                       |   |       |
|-----------------------|---|-------|
| RED-dropped packets : | 0 | 0 pps |
| RED-dropped bytes :   | 0 | 0 bps |

PFE: 3

```

RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

### show interfaces voq et-5/0/12 (For a Specific Forwarding Class and Source FPC)

```

user@host> show interfaces voq et-5/0/12 forwarding-class best-effort source-fpc 5
Physical interface: et-5/0/12, Enabled, Physical link is Up
Interface index: 166, SNMP ifIndex: 1104

```

Queue: 0, Forwarding classes: best-effort

FPC number: 5

PFE: 0

```

RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

PFE: 1

```

RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

PFE: 2

```

RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

PFE: 3

```

RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

PFE: 4

```

RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

PFE: 5

```

RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

PFE: 6

```

RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

PFE: 7

```

RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

### show interfaces voq et-7/0/0 (Non-Zero)

```

user@host> show interfaces voq et-7/0/0 non-zero

```

```

Physical interface: et-7/0/0, Enabled, Physical link is Up
Interface index: 155, SNMP ifIndex: 699

```

Queue: 0, Forwarding classes: best-effort

FPC number: 4

PFE: 0

```

RED-dropped packets : 95862238 2301586 pps
RED-dropped bytes : 10544846180 2025396264 bps

```

FPC number: 6

PFE: 0

```

RED-dropped packets : 95866639 2322569 pps
RED-dropped bytes : 10545330290 2043860728 bps

```

Queue: 3, Forwarding classes: network-control

FPC number: 4

PFE: 0

```

RED-dropped packets : 78433066 1899727 pps
RED-dropped bytes : 8627637260 1671760384 bps

FPC number: 6
PFE: 0
RED-dropped packets : 78436704 1900628 pps
RED-dropped bytes : 8628037440 1672553432 bps

```

#### show interfaces voq et-7/0/0 (For a Specific Forwarding Class and Non-Zero)

```
user@host show interfaces voq et-7/0/0 forwarding-class best-effort non-zero
```

```
Physical interface: et-7/0/0, Enabled, Physical link is Up
```

```
Interface index: 155, SNMP ifIndex: 699
```

```
Queue: 0, Forwarding classes: best-effort
```

```

FPC number: 4
PFE: 0
RED-dropped packets : 119540012 2322319 pps
RED-dropped bytes : 13149401320 2043640784 bps

FPC number: 6
PFE: 0
RED-dropped packets : 119540049 2322988 pps
RED-dropped bytes : 13149405390 2044229744 bps

```





## CHAPTER 264

# Operational Commands (Data Center Bridging and Flow Control)

- `show class-of-service congestion-notification`
- `show dcbx`
- `show dcbx neighbors`

## show class-of-service congestion-notification

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show class-of-service congestion-notification                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Display whether priority-based flow control (PFC) is enabled for each IEEE 802.1p code point.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <b>none</b> —Display the PFC state for all IEEE 802.1p code points.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 6905</a></li> <li>• <a href="#">Example: Configuring CoS PFC for FCoE Traffic on page 6504</a></li> <li>• <a href="#">Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch) on page 6969</a></li> <li>• <a href="#">Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 6977</a></li> <li>• <a href="#">Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 6986</a></li> <li>• <a href="#">Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI) on page 7000</a></li> <li>• <a href="#">Example: Configuring PFC Across Layer 3 Interfaces on page 6923</a></li> <li>• <a href="#">Understanding CoS Flow Control (Ethernet PAUSE and PFC) on page 6493</a></li> </ul> |
| <b>Output Fields</b>            | <a href="#">Table 643</a> describes the output fields for the <b>show class-of-service congestion-notification</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

**Table 643: show class-of-service congestion-notification Output Fields**

| Field Name          | Field Description                                                                                                       |
|---------------------|-------------------------------------------------------------------------------------------------------------------------|
| <b>Type</b>         | Type of interfaces on which congestion notification is applied. Congestion notification is applied on input interfaces. |
| <b>Index</b>        | Index of this congestion notification profile.                                                                          |
| <b>Name</b>         | Name of the congestion notification profile.                                                                            |
| <b>Cable Length</b> | Length of the attached physical cable in meters. The default value is 100 meters.                                       |
| <b>Priority</b>     | IEEE 802.1p code point.                                                                                                 |
| <b>PFC</b>          | State of PFC for the corresponding code point, either <b>enabled</b> or <b>disabled</b> .                               |

Table 643: show class-of-service congestion-notification Output Fields (*continued*)

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>MRU</b>                 | <p>Maximum receive unit of the interface in bytes. (Incoming traffic that exceeds the MRU size of an interface is dropped.) The default values are:</p> <ul style="list-style-type: none"> <li>• 2500 bytes for priority 3 traffic</li> <li>• 9216 bytes for priority 4 traffic</li> </ul> <p><b>NOTE:</b> If you configure flow control on a priority that is not one of the default flow control priorities, the default MRU value is 2500 bytes. For example, if you configure flow control on priority 5 and you do not configure an MRU value, the default MRU value is 2500 bytes.</p> |
| <b>Flow-Control-Queues</b> | <p>Output queue mapping to IEEE 802.1p code points (priorities). Explicit output queue to priority mapping overwrites the default configuration, and only explicitly mapped queues are displayed in the output. Flow control is only enabled on a queue when you enable PFC on the corresponding priority in the input stanza of the congestion notification profile.</p>                                                                                                                                                                                                                    |

## Sample Output

### show class-of-service congestion-notification

```

user@switch> show class-of-service congestion-notification
Name: fcoe_p3_cnp, Index: 12037
Type: Input
Cable Length: 100 m
 Priority PFC MRU
 000 Disabled
 001 Disabled
 010 Disabled
 011 Enabled 2500
 100 Enabled 9216
 101 Disabled
 110 Disabled
 111 Disabled
Type: Output
 Priority Flow-Control-Queues
 000
 001 0
 010 1
 011 2
 100 3
 101 4
 110 5
 111 6
 7

Name: fcoe_p3_p5_cnp, Index: 46484
Type: Input
Cable Length: 100 m
 Priority PFC MRU

```

|              |                     |      |
|--------------|---------------------|------|
| 000          | Disabled            |      |
| 001          | Disabled            |      |
| 010          | Disabled            |      |
| 011          | Enabled             | 2240 |
| 100          | Disabled            |      |
| 101          | Enabled             | 2240 |
| 110          | Disabled            |      |
| 111          | Disabled            |      |
| Type: Output |                     |      |
| Priority     | Flow-Control-Queues |      |
| 011          |                     |      |
|              | 3                   |      |
| 101          |                     |      |
|              | 5                   |      |

show dcbx

|                          |                                                                                                                                                                            |
|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | show dcbx                                                                                                                                                                  |
| Release Information      | Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                            |
| Description              | List DCBX status (enabled or disabled) and the interfaces on which DCBX is enabled.                                                                                        |
| Required Privilege Level | view                                                                                                                                                                       |
| Related Documentation    | <ul style="list-style-type: none"><li>• <a href="#">show dcbx neighbors on page 6581</a></li><li>• <a href="#">Configuring DCBX Autonegotiation on page 6524</a></li></ul> |
| Output Fields            | <a href="#">Table 523</a> lists the output fields for the <b>show dcbx</b> command. Output fields are listed in the approximate order in which they appear.                |

Table 644: show dcbx output fields

| Field Name | Field Description                                                                                                                                                                                                                                                  |
|------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DCBX       | Status of DCBX on the switch or for the specified interface: <ul style="list-style-type: none"><li>• Enabled—DCBX is enabled on the switch or on the specified interface</li><li>• Disabled—DCBX is disabled on the switch or on the specified interface</li></ul> |
| Interface  | Name of the interface                                                                                                                                                                                                                                              |

Sample Output

show dcbx

```
user@switch> show dcbx
DCBX : Enabled
Interface DCBX
xe-0/0/9.0 enabled
xe-0/0/32.0 enabled
xe-0/0/36.0 enabled
```

## show dcbx neighbors

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show dcbx neighbors</b><br><b>&lt;interface interface-name&gt;</b><br><b>&lt;terse&gt;</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.<br>Command introduced in Junos OS Release 11.3 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Display information about Data Center Bridging Capability Exchange protocol (DCBX) neighbor interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <b>none</b> —Display information about all DCBX neighbor interfaces.<br><br><b>interface-name</b> —(Optional) Display information for the specified interface.<br><br><b>terse</b> —Display the specified level of output.                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring DCBX Autonegotiation on page 6524</a></li> <li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 6535</a></li> <li>• <a href="#">Example: Configuring an FCoE Transit Switch</a></li> <li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li> <li>• <a href="#">Understanding DCB Features and Requirements on page 6490</a></li> <li>• <a href="#">Understanding Data Center Bridging Capability Exchange Protocol for EX Series Switches</a></li> <li>• <a href="#">dcbx on page 6563</a></li> </ul>                   |
| <b>List of Sample Output</b>    | <a href="#">show dcbx neighbors interface (QFX Series, DCBX Version 1.01 Mode) on page 7411</a><br><a href="#">show dcbx neighbors interface (QFX Series, IEEE DCBX Mode) on page 7413</a><br><a href="#">show dcbx neighbors terse (QFX Series) on page 7415</a><br><a href="#">show dcbx neighbors (EX4500 Switch: FCoE Interfaces on Both Local and Peer with PFC Configured Compatibly) on page 7415</a><br><a href="#">show dcbx neighbors (EX4500 Switch: DCBX Interfaces on Local and Peer Are Configured Compatibly with iSCSI Application) on page 7416</a><br><a href="#">show dcbx neighbors (EX4500 Switch: Includes ETS) on page 7417</a> |
| <b>Output Fields</b>            | <a href="#">Table 524</a> lists the output fields for the <b>show dcbx neighbors</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

Table 645: show dcbx neighbors Output Fields

| Field Name | Field Description      |
|------------|------------------------|
| Interface  | Name of the interface. |

Table 645: show dcbx neighbors Output Fields (*continued*)

| Field Name             | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Parent Interface       | Name of the link aggregation group (LAG) interface to which the DCBX interface belongs.                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Active-application-map | Name of the application map applied to the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Protocol-Mode          | <p>(QFX Series) DCBX protocol mode the interface uses:</p> <ul style="list-style-type: none"> <li>• IEEE DCBX Version—The interface uses IEEE DCBX mode.</li> <li>• DCBX Version 1.01—The interface uses DCBX version 1.01.</li> </ul> <p><b>NOTE:</b> On interfaces that use the IEEE DCBX mode, the <b>show dcbx neighbors interface <i>interface-name</i></b> operational command does not include application, PFC, or ETS operational state in the output.</p>                                        |
| Protocol-State         | <p>(DCBX Version 1.01 only) DCBX protocol state synchronization status:</p> <ul style="list-style-type: none"> <li>• <b>in-sync</b>—The local interface received an acknowledge message from the peer to indicate that the peer received a state change message sent by the local interface.</li> <li>• <b>ack-pending</b>—The local interface has not yet received an acknowledge message from the peer to indicate that the peer received a state change message sent by the local interface.</li> </ul> |
| Local-Advertisement    | <p>(DCBX Version 1.01 only)</p> <p>Status of advertisements that the local interface sends to the peer.</p>                                                                                                                                                                                                                                                                                                                                                                                                |
| Operational version    | Version of the DCBX standard used.                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| sequence-number        | <p>Number of state change messages sent to the peer.</p> <p>If the interface <b>Protocol-State</b> value is <b>in-sync</b>, this number should match the <b>acknowledge-id</b> number in the <b>Peer-Advertisement</b> section.</p> <p>If the interface <b>Protocol-State</b> value is <b>ack-pending</b>, this number does not match the <b>acknowledge-id</b> number in the <b>Peer-Advertisement</b> section.</p>                                                                                       |
| acknowledge-id         | <p>Number of acknowledge messages received from the peer.</p> <p>If the <b>Protocol-State</b> value is <b>in-sync</b>, this number should match the <b>sequence-number</b> value in the <b>Peer-Advertisement</b> section.</p> <p>If the <b>Protocol-State</b> value is <b>ack-pending</b>, this number does not match the <b>sequence-number</b> value in the <b>Peer-Advertisement</b> section.</p>                                                                                                      |

Table 645: show dcbx neighbors Output Fields (*continued*)

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Peer-Advertisement</b>  | (DCBX Version 1.01 only)<br><br>Status of advertisements that the peer sends to the local interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Operational version</b> | Version of the DCBX standard used.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>sequence-number</b>     | Number of state change messages the peer sent to the local interface.<br><br>If this number matches the <b>acknowledge-id</b> number in the <b>Local-Advertisement</b> field, this indicates that the local interface has acknowledged all of the peer's state change messages and is synchronized.<br><br>If this number does not match the <b>acknowledge-id</b> number in the <b>Local-Advertisement</b> field, this indicates that the peer has not yet received an acknowledgment for a state change message from the local interface.            |
| <b>acknowledge-id</b>      | Number of acknowledge messages the peer has received from the local interface.<br><br>If this number matches the <b>sequence-number</b> value in the <b>Local-Advertisement</b> field, this indicates that the peer has acknowledged all of the local interface's state change messages and is in synchronization.<br><br>If this number does not match the <b>sequence-number</b> value in the <b>Local-Advertisement</b> field, this indicates that the peer has not yet sent an acknowledgment for a state change message from the local interface. |



Table 645: show dcbx neighbors Output Fields (*continued*)

| Field Name                        | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Feature: PFC</b>               | Priority-based flow control (PFC) feature DCBX state information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Protocol-State</b>             | (DCBX Version 1.01 only)<br><br>DCBX protocol state synchronization status: <ul style="list-style-type: none"> <li>• <b>ack-pending</b>—The local interface has not yet received an acknowledge message from the peer to indicate that the peer received a PFC state change message sent by the local interface.</li> <li>• <b>in-sync</b>—The local interface received an acknowledge message from the peer to indicate that the peer received a PFC state change message sent by the local interface.</li> <li>• <b>not-applicable</b>—PFC autonegotiation is disabled.</li> </ul> |
| <b>Operational State</b>          | (DCBX Version 1.01 only)<br><br>Operational state of the feature: <b>enabled</b> or <b>disabled</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Local-Advertisement</b>        | Status of advertisements that the local interface sends to the peer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Enable</b>                     | (DCBX Version 1.01 only)<br><br>State that the local interface advertises to the peer: <ul style="list-style-type: none"> <li>• <b>Yes</b>—The feature is enabled.</li> <li>• <b>No</b>—The feature is disabled.</li> </ul>                                                                                                                                                                                                                                                                                                                                                          |
| <b>Willing</b>                    | Willingness of the local interface to learn the PFC configuration from the peer using DCBX: <ul style="list-style-type: none"> <li>• <b>Yes</b>—The local interface is willing to learn the PFC configuration from the peer.</li> <li>• <b>No</b>—The local interface is not willing to learn the PFC configuration from the peer.</li> </ul>                                                                                                                                                                                                                                        |
| <b>Mac auth Bypass Capability</b> | (IEEE DCBX only)<br><br>(QFX Series) Media access controller (MAC) authentication bypass provides access to devices based on MAC address authentication. This is not supported, so the only value seen in the local advertisement field is <b>no</b> .                                                                                                                                                                                                                                                                                                                               |
| <b>Error</b>                      | (DCBX Version 1.01 only)<br><br>Configuration compatibility error status: <ul style="list-style-type: none"> <li>• <b>No</b>—No error detected. Local and peer configuration are compatible.</li> <li>• <b>Yes</b>—Error detected. Local and peer configuration are not compatible.</li> </ul>                                                                                                                                                                                                                                                                                       |

Table 645: show dcbx neighbors Output Fields (*continued*)

| Field Name                                            | Field Description                                                                                                                                                                                                                                                                                                                                    |
|-------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Operational State</b>                              | <p>PFC operational state on the interface:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—PFC is enabled on the interface</li> <li>• <b>Disabled</b>—PFC is disabled on the interface</li> </ul>                                                                                                                                        |
| <b>Maximum Traffic Classes capable to support PFC</b> | <p>Largest number of traffic classes the local interface supports for PFC:</p> <ul style="list-style-type: none"> <li>• <b>6</b> (EX Series switches)</li> <li>• <b>6</b> (QFX Series)</li> </ul>                                                                                                                                                    |
| <b>Code Point</b>                                     | <p>PFC code point, which is specified in the 3-bit class-of-service field in the VLAN header.</p>                                                                                                                                                                                                                                                    |
| <b>Admin Mode</b>                                     | <p>PFC administrative state for each code point on the local interface:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—PFC is enabled for the code point.</li> <li>• <b>Disabled</b>—PFC is disabled for the code point.</li> </ul>                                                                                                     |
| <b>Operational Mode</b>                               | <p>(QFX Series) PFC operational mode for each code point:</p> <ul style="list-style-type: none"> <li>• <b>Enable</b>—PFC is enabled on the code point.</li> <li>• <b>Disable</b>—PFC is disabled on the code point.</li> </ul>                                                                                                                       |
| <b>Peer-Advertisement</b>                             | <p>Status of advertisements that the peer sends to the local interface.</p>                                                                                                                                                                                                                                                                          |
| <b>Enable</b>                                         | <p>(DCBX Version 1.01 only)</p> <p>State that the peer advertises to the local interface:</p> <ul style="list-style-type: none"> <li>• <b>Yes</b>—The feature is enabled.</li> <li>• <b>No</b>—The feature is disabled.</li> </ul>                                                                                                                   |
| <b>Willing</b>                                        | <p>Willingness of the peer to learn the PFC configuration from the local interface using DCBX:</p> <ul style="list-style-type: none"> <li>• <b>Yes</b>—The peer is willing to learn the PFC configuration from the local interface.</li> <li>• <b>No</b>—The peer is not willing to learn the PFC configuration from the local interface.</li> </ul> |
| <b>Error</b>                                          | <p>(DCBX Version 1.01 only)</p> <p>Configuration compatibility error status:</p> <ul style="list-style-type: none"> <li>• <b>No</b>—No error detected. Local and peer configuration are compatible.</li> <li>• <b>Yes</b>—Error detected. Local and peer configuration are not compatible.</li> </ul>                                                |

Table 645: show dcbx neighbors Output Fields (*continued*)

| Field Name                                            | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Operational State</b>                              | <p>PFC operational state on the interface:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—PFC is enabled on the interface</li> <li>• <b>Disabled</b>—PFC is disabled on the interface</li> </ul>                                                                                                                                                                                                                                                                                            |
| <b>Mac auth Bypass Capability</b>                     | <p>(IEEE DCBX only)</p> <p>(QFX Series) Media access controller (MAC) authentication bypass provides access to devices based on MAC address authentication. Although the QFX Series does not support this feature, the connected peer might support it. This field reports the peer state:</p> <ul style="list-style-type: none"> <li>• <b>Yes</b>—The connected peer supports MAC authentication bypass.</li> <li>• <b>No</b>—The connected peer does not support MAC authentication bypass.</li> </ul> |
| <b>Maximum Traffic Classes capable to support PFC</b> | <p>Largest number of traffic classes the peer supports for PFC:</p> <ul style="list-style-type: none"> <li>• <b>6</b> (EX Series switches)</li> <li>• <b>8</b> (QFX Series)</li> </ul>                                                                                                                                                                                                                                                                                                                   |
| <b>Code Point</b>                                     | <p>PFC code point, which is specified in the 3-bit class-of-service field in the VLAN header.</p>                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Admin Mode</b>                                     | <p>PFC administrative state for each code point on the peer:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—PFC is enabled for the code point.</li> <li>• <b>Disabled</b>—PFC is disabled for the code point.</li> </ul>                                                                                                                                                                                                                                                                    |

Table 645: show dcbx neighbors Output Fields (*continued*)

| Field Name                  | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Feature: Application</b> | State information for the DCBX application.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Protocol-State</b>       | <p>(DCBX Version 1.01 only)</p> <p>DCBX protocol state synchronization status:</p> <ul style="list-style-type: none"> <li>• <b>in-sync</b>—The local interface received an acknowledge message from the peer to indicate that the peer received an FCoE state change message sent by the local interface.</li> <li>• <b>ack-pending</b>—The local interface has not yet received an acknowledge message from the peer to indicate that the peer received an FCoE state change message sent by the local interface.</li> <li>• <b>not-applicable</b>—The local interface is set to <b>no-auto-negotiation</b> (autonegotiation is disabled). If the interface is associated with an FCoE forwarding class, the interface advertises FCoE capability even if the connected peer does not advertise FCoE capability.</li> </ul> |
| <b>Local-Advertisement</b>  | <p>Status of advertisements that the local interface sends to the peer.</p> <p>If the local interface is set to <b>no-auto-negotiation</b> (autonegotiation is disabled), the local advertisement portion of the output is not shown.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Enable</b>               | <p>(DCBX Version 1.01 only)</p> <p>State that the local interface advertises to the peer:</p> <ul style="list-style-type: none"> <li>• <b>Yes</b>—The feature is enabled.</li> <li>• <b>No</b>—The feature is disabled.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Willing</b>              | <p>(DCBX Version 1.01 only)</p> <p>Willingness of the local interface to learn the FCoE interface state from the peer using DCBX:</p> <ul style="list-style-type: none"> <li>• <b>Yes</b>—The local interface is willing to learn the FCoE interface state from the peer.</li> <li>• <b>No</b>—The local interface is not willing to learn the FCoE interface state from the peer.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Error</b>                | <p>(DCBX Version 1.01 only)</p> <p>Configuration compatibility error status:</p> <ul style="list-style-type: none"> <li>• <b>No</b>—No error detected. The local and peer configuration are compatible.</li> <li>• <b>Yes</b>—Error detected. The local and peer configuration are not compatible.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Appl-Name</b>            | Name of the application:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

Table 645: show dcbx neighbors Output Fields (*continued*)

| Field Name                            | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Ethernet-Type</b>                  | <p>(DCBX Version 1.01 only)</p> <p>Ethernet type (EtherType) of the application. For example, <b>0x8906</b> indicates the EtherType for the FCoE application. Either the EtherType (for Layer 2 applications) or the Socket Number (for Layer 4 applications) of the application is displayed in the output.</p>                                                                                                                                                                                                                                                            |
| <b>Socket-Number</b>                  | <p>Destination port socket number of the application, if applicable. Either the EtherType (for Layer 2 applications) or the Socket Number (for Layer 4 applications) of the application is displayed in the output.</p>                                                                                                                                                                                                                                                                                                                                                     |
| <b>Priority-Field or Priority-Map</b> | <p>Priority assigned to the application.</p> <p>For EX Series switches, the priority of the FCoE application is determined by the PFC congestion notification profile that has been configured and associated with the FCoE interface. For other applications, the priority is based on the application map.</p>                                                                                                                                                                                                                                                            |
| <b>Status</b>                         | <p>(DCBX Version 1.01 only)</p> <p>Local status when autonegotiation is enabled:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—The application feature is enabled on both the local interface and the peer interface. (The local configuration and the peer configuration match.)</li> <li>• <b>Disabled</b>—The local configuration and the peer configuration do not match.</li> </ul> <p><b>NOTE:</b> If there is a configuration mismatch in one application between the switch and the peer, all the other applications including FCoE are disabled.</p> |
| <b>Peer-Advertisement</b>             | <p>Status of advertisements that the peer sends to the local interface.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Enable</b>                         | <p>(DCBX Version 1.01 only)</p> <p>State that the peer advertises to the local interface:</p> <ul style="list-style-type: none"> <li>• <b>Yes</b>—The feature is enabled.</li> <li>• <b>No</b>—The feature is disabled.</li> </ul>                                                                                                                                                                                                                                                                                                                                          |
| <b>Willing</b>                        | <p>(DCBX Version 1.01 only)</p> <p>Willingness of the peer to learn the FCoE interface state from the local interface using DCBX:</p> <ul style="list-style-type: none"> <li>• <b>Yes</b>—The peer is willing to learn the FCoE interface state from the local interface.</li> <li>• <b>No</b>—The peer is not willing to learn the FCoE interface state from the local interface.</li> </ul>                                                                                                                                                                               |

Table 645: show dcbx neighbors Output Fields (*continued*)

| Field Name                            | Field Description                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Error</b>                          | (DCBX Version 1.01 only)<br><br>Configuration compatibility error status: <ul style="list-style-type: none"> <li>• <b>No</b>—No error detected. Local and peer configuration are compatible.</li> <li>• <b>Yes</b>—Error detected. Local and peer configuration are not compatible.</li> </ul>                                                                           |
| <b>Appl-Name</b>                      | Name of the application: <ul style="list-style-type: none"> <li>• <b>FCoE</b>—Fibre Channel over Ethernet</li> </ul>                                                                                                                                                                                                                                                     |
| <b>Ethernet-Type</b>                  | Ethernet type (EtherType) of the application. For example, <b>0x8906</b> indicates the EtherType for the FCoE application. Either the EtherType (for Layer 2 applications) or the Socket-Number (for Layer 4 applications) of the application is displayed in the output.                                                                                                |
| <b>Socket-Number</b>                  | Destination port socket number of the application, if applicable. Either the EtherType (for Layer 2 applications) or the Socket Number (for Layer 4 applications) of the application is displayed in the output.                                                                                                                                                         |
| <b>Priority-Field or Priority-Map</b> | Priority assigned to the application.                                                                                                                                                                                                                                                                                                                                    |
| <b>Status</b>                         | (DCBX Version 1.01 only)<br><br>Peer interface status: <ul style="list-style-type: none"> <li>• <b>Enabled</b>—The application feature is enabled on both the local interface and the peer interface. (The local configuration and the peer configuration match.)</li> <li>• <b>Disabled</b>—The local configuration and the peer configuration do not match.</li> </ul> |

Table 645: show dcbx neighbors Output Fields (*continued*)

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Feature: ETS</b>        | Enhanced Transmission Selection (ETS) DCBX state information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Protocol-State</b>      | (DCBX Version 1.01 only)<br><br>ETS protocol state synchronization status: <ul style="list-style-type: none"> <li>• <b>in-sync</b>—The local interface received an acknowledge message from the peer to indicate that the peer received an ETS state change message sent by the local interface.</li> <li>• <b>ack-pending</b>—The local interface has not yet received an acknowledge message from the peer to indicate that the peer received an ETS state change message sent by the local interface.</li> </ul>                                                          |
| <b>Operational State</b>   | (DCBX Version 1.01 only)<br><br>Operational state of the feature, <b>enabled</b> or <b>disabled</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Local-Advertisement</b> | Status of advertisements that the local interface sends to the peer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Enable</b>              | (DCBX Version 1.01 only)<br><br>State that the local interface advertises to the peer: <ul style="list-style-type: none"> <li>• <b>Yes</b>—The feature is enabled.</li> <li>• <b>No</b>—The feature is disabled.</li> </ul>                                                                                                                                                                                                                                                                                                                                                  |
| <b>TLV Type</b>            | (IEEE DCBX only)<br><br>Type of ETS TLV: <ul style="list-style-type: none"> <li>• <b>Configuration</b>—Advertises the Configuration TLV, which communicates the local ETS configuration to the peer but does not ask the peer to use the configuration.</li> <li>• <b>Recommendation</b>—Advertises the Recommendation TLV, which communicates the local ETS configuration to the peer, and if the peer is “willing,” configures the peer interface to match the local ETS configuration.</li> <li>• <b>Recommendation-or-Configuration</b>—Advertises both TLVs.</li> </ul> |
| <b>Willing</b>             | Willingness of the local interface to learn the ETS state from the peer using DCBX (EX Series switches always advertise <b>No</b> for this field): <ul style="list-style-type: none"> <li>• <b>Yes</b>—Local interface is willing to learn the ETS state from the peer.</li> <li>• <b>No</b>—Local interface is not willing to learn the ETS state from the peer.</li> </ul>                                                                                                                                                                                                 |
| <b>Credit Based Shaper</b> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

Table 645: show dcbx neighbors Output Fields (*continued*)

| Field Name                                            | Field Description                                                                                                                                                                                                                                                                                               |
|-------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                                       | (IEEE DCBX only)                                                                                                                                                                                                                                                                                                |
|                                                       | Alternative method of flow control to buffer-to-buffer credit. The QFX Series does not support a credit-based shaper, so the value of this field is always <b>No</b> .                                                                                                                                          |
| <b>Error</b>                                          | (DCBX Version 1.01 only)<br><br>Configuration error status: <ul style="list-style-type: none"> <li>• <b>No</b>—No error. This should always be the switch ETS error state.</li> <li>• <b>Yes</b>—Error detected.</li> </ul>                                                                                     |
| <b>Maximum Traffic Classes capable to support PFC</b> | (DCBX Version 1.01 only)<br><br>Largest number of traffic classes the local interface supports for PFC.                                                                                                                                                                                                         |
| <b>Maximum Traffic Classes supported</b>              | (IEEE DCBX only)<br><br>Largest number of traffic classes the local interface supports for ETS. (EX Series switches support only one traffic class for ETS. However, a different value might be shown for this field.)                                                                                          |
| <b>Code Point</b>                                     | PFC code point, which is specified in the 3-bit class-of-service field in the VLAN header.                                                                                                                                                                                                                      |
| <b>Priority-Group</b>                                 | Class-of-service (CoS) priority group (forwarding class set) identification number.                                                                                                                                                                                                                             |
| <b>Percentage B/W</b>                                 | Configured minimum percentage of link bandwidth allocated to the priority group. Only explicitly configured values appear in this output column. If the link bandwidth is the default percentage, it is not shown. (EX Series switches allocate 100% of link bandwidth to the default priority group, group 7.) |
| <b>Transmission Selection Algorithm</b>               | (IEEE DCBX only)<br><br>The transmission selection algorithm used by the interface. The QFX Series supports ETS but does not support using the credit-based shaper algorithm, so the only value shown in this field is <b>ETS</b> .                                                                             |
| <b>Peer-Advertisement</b>                             | Status of advertisements that the peer sends to the local interface.                                                                                                                                                                                                                                            |
| <b>Enable</b>                                         |                                                                                                                                                                                                                                                                                                                 |



Table 645: show dcbx neighbors Output Fields (*continued*)

| Field Name                                            | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                                       | (DCBX Version 1.01 only)<br><br>State that the peer advertises to the local interface:<br><br><ul style="list-style-type: none"> <li>• <b>Yes</b>—The feature is enabled.</li> <li>• <b>No</b>—The feature is disabled.</li> </ul>                                                                                                                                                                                                                                                                                                                                               |
| <b>TLV Type</b>                                       | (IEEE DCBX only)<br><br>Type of ETS TLV:<br><br><ul style="list-style-type: none"> <li>• <b>Configuration</b>—Advertises the Configuration TLV, which communicates the local ETS configuration to the peer but does not ask the peer to use the configuration.</li> <li>• <b>Recommendation</b>—Advertises the Recommendation TLV, which communicates the local ETS configuration to the peer, and if the peer is “willing,” configures the peer interface to match the local ETS configuration.</li> <li>• <b>Configuration/Recommendation</b>—Advertises both TLVs.</li> </ul> |
| <b>Willing</b>                                        | Willingness of the peer to learn the ETS state from the local interface using DCBX:<br><br><ul style="list-style-type: none"> <li>• <b>Yes</b>—Peer is willing to learn the ETS state from the local interface.</li> <li>• <b>No</b>—Peer is not willing to learn the ETS state from the local interface.</li> </ul>                                                                                                                                                                                                                                                             |
| <b>Credit Based Shaper</b>                            | (IEEE DCBX only)<br><br>Alternative method of flow control to buffer-to-buffer credit. The QFX Series does not support a credit-based shaper, so the value of this field is always <b>No</b> .                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Error</b>                                          | (DCBX Version 1.01 only)<br><br>Configuration error status of the peer:<br><br><ul style="list-style-type: none"> <li>• <b>No</b>—No error in peer ETS TLV.</li> <li>• <b>Yes</b>—Error in peer ETS TLV.</li> </ul>                                                                                                                                                                                                                                                                                                                                                              |
| <b>Maximum Traffic Classes capable to support PFC</b> | (DCBX Version 1.01 only)<br><br>Largest number of traffic classes the local interface supports for PFC.                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Maximum Traffic Classes supported</b>              | (IEEE DCBX only)<br><br>Largest number of traffic classes the local interface supports for ETS. (EX Series switches support only one traffic class for ETS. However, a different value might be shown for this field.)                                                                                                                                                                                                                                                                                                                                                           |
| <b>Code Point</b>                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

Table 645: show dcbx neighbors Output Fields (*continued*)

| Field Name                              | Field Description                                                                                                                                                                                                                                                                                                                                                                                            |
|-----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                         | PFC code point, which is specified in the 3-bit class-of-service field in the VLAN header.                                                                                                                                                                                                                                                                                                                   |
| <b>Priority-Group</b>                   | CoS priority group (forwarding class set) identification number.                                                                                                                                                                                                                                                                                                                                             |
| <b>Percentage B/W</b>                   | Configured minimum percentage of link bandwidth allocated to the priority group. (EX Series switches allocate 100% of link bandwidth to the default priority group, group 7.)                                                                                                                                                                                                                                |
| <b>Transmission Selection Algorithm</b> | (IEEE DCBX only)<br><br>Transmission selection algorithm used by the interface. The QFX Series supports ETS but does not support using the credit-based shaper algorithm, so the only value shown in this field is <b>ETS</b> .                                                                                                                                                                              |
| <b>PFC</b>                              | (QFX Series, <b>terse</b> option only) DCBX TLV advertisement state for PFC: <ul style="list-style-type: none"> <li>Disabled—PFC configuration matches the configuration on the connected peer and PFC is disabled</li> <li>Enabled—PFC configuration matches the configuration on the connected peer and PFC is enabled</li> <li>Not Advt—Interface does not advertise PFC to the connected peer</li> </ul> |
| <b>ETS</b>                              | ( <b>terse</b> option only) Local DCBX TLV advertisement state for ETS: <ul style="list-style-type: none"> <li>Advt—Interface advertises ETS TLVs</li> <li>Disabled—ETS is disabled on the interface (interface does not advertise ETS)</li> </ul>                                                                                                                                                           |
| <b>ETS Rec</b>                          | ( <b>terse</b> option only) DCBX TLV peer advertisement state for ETS (state received from the connected DCBX peer): <ul style="list-style-type: none"> <li>Advt—Peer interface advertises ETS TLVs</li> <li>Not Advt—Peer interface does not advertise ETS</li> </ul> <p><b>NOTE:</b> When the DCBX mode is DCBX version 1.01, no peer information is displayed.</p>                                        |

Table 645: show dcbx neighbors Output Fields (*continued*)

| Field Name | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Version    | <p>(<b>terse</b> option only) The DCBX version used on the interface and whether the DCBX version was autonegotiated or explicitly configured:</p> <ul style="list-style-type: none"> <li>• <b>IEEE</b>—The interface uses IEEE DCBX.</li> <li>• <b>1.01</b>—The interface uses DCBX version 1.01.</li> </ul> <p>When the DCBX version used is the result of autonegotiation, the term (<b>Auto</b>) appears next to the version. For example, <b>IEEE (Auto)</b> indicates that the interface autonegotiated with the connected peer to use IEEE DCBX. Autonegotiation is enabled by default.</p> |

## Sample Output

### show dcbx neighbors interface (QFX Series, DCBX Version 1.01 Mode)

```

user@switch> show dcbx neighbors interface xe-0/0/0
Interface : xe-0/0/0.0 - Parent Interface: ae0.0
Active-application-map: app-map-1
Protocol-State: in-sync
Protocol-Mode: DCBX Version 1.01

Local-Advertisement:
 Operational version: 1
 sequence-number: 130, acknowledge-id: 102

Peer-Advertisement:
 Operational version: 1
 sequence-number: 102, acknowledge-id: 130

Feature: PFC, Protocol-State: in-sync

Operational State: Enabled

Local-Advertisement:
 Enable: Yes, Willing: No, Error: No
 Maximum Traffic Classes capable to support PFC: 8

Code Point Admin Mode Operational Mode
000 Disabled Disable
001 Disabled Disable
010 Disabled Disable
011 Enabled Enable
100 Enabled Enable
101 Disabled Disable
110 Disabled Disable
111 Disabled Disable

Peer-Advertisement:
 Enable: Yes, Willing: No, Error: No
 Maximum Traffic Classes capable to support PFC: 8

Code Point Admin Mode
000 Disabled

```

|     |          |
|-----|----------|
| 001 | Disabled |
| 010 | Disabled |
| 011 | Enabled  |
| 100 | Enabled  |
| 101 | Disabled |
| 110 | Disabled |
| 111 | Disabled |

Feature: Application, Protocol-State: in-sync

Local-Advertisement:

Enable: Yes, Willing: No, Error: No

| App1-Name | Ethernet-Type | Socket-Number | Priority-Map | Status  |
|-----------|---------------|---------------|--------------|---------|
| FCoE      | 0x8906        |               | 00001110     | Enabled |
| iSCSI     |               | 3260          | 10000000     | Enabled |

Peer-Advertisement:

Enable: Yes, Willing: Yes, Error: No

| App1-Name | Ethernet-Type | Socket-Number | Priority-Map | Status  |
|-----------|---------------|---------------|--------------|---------|
| FCoE      | 0x8906        | N/A           | 00001110     | Enabled |

Feature: ETS, Protocol-State: in-sync

Operational State: Enabled

Local-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 8

| Code Point | Priority-Group |
|------------|----------------|
| 000        | 0              |
| 001        | 7              |
| 010        | 7              |
| 011        | 7              |
| 100        | 0              |
| 101        | 1              |
| 110        | 1              |
| 111        | 7              |

| Priority-Group | Percentage B/W |
|----------------|----------------|
| 0              | 40%            |
| 1              | 5%             |

Peer-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 8

| Code Point | Priority-Group |
|------------|----------------|
| 000        | 0              |
| 001        | 7              |
| 010        | 7              |
| 011        | 7              |
| 100        | 0              |
| 101        | 1              |
| 110        | 1              |

|                |                |
|----------------|----------------|
| 111            | 7              |
| Priority-Group | Percentage B/W |
| 0              | 40%            |
| 1              | 5%             |

### show dcbx neighbors interface (QFX Series, IEEE DCBX Mode)

user@switch> **show dcbx neighbors interface xe-0/0/0**

Interface : xe-0/0/0.0 - Parent Interface: ae0.0

Active-application-map: app-map-1

Protocol-Mode: IEEE-DCBX Version

Feature: PFC

Local-Advertisement:

Willing: No

Mac auth Bypass Capability: No

Operational State: Enabled

Maximum Traffic Classes capable to support PFC: 8

| Code Point | Admin Mode |
|------------|------------|
| 000        | Disabled   |
| 001        | Disabled   |
| 010        | Disabled   |
| 011        | Enabled    |
| 100        | Enabled    |
| 101        | Disabled   |
| 110        | Disabled   |
| 111        | Disabled   |

Peer-Advertisement:

Willing: No

Mac auth Bypass Capability: No

Operational State: Enabled

Maximum Traffic Classes capable to support PFC: 8

| Code Point | Admin Mode |
|------------|------------|
| 000        | Disabled   |
| 001        | Disabled   |
| 010        | Disabled   |
| 011        | Enabled    |
| 100        | Enabled    |
| 101        | Disabled   |
| 110        | Disabled   |
| 111        | Disabled   |

Feature: Application

Local-Advertisement:

| App1-Name | Ethernet-Type | Socket-Number | Priority-field |
|-----------|---------------|---------------|----------------|
| FCoE      | 0x8906        |               | 00001110       |
| iSCSI     |               | 3260          | 10000000       |

Peer-Advertisement:

| App1-Name | Ethernet-Type | Socket-Number | Priority-field |
|-----------|---------------|---------------|----------------|
|-----------|---------------|---------------|----------------|

|      |        |     |          |
|------|--------|-----|----------|
| FCoE | 0x8906 | N/A | 00001110 |
|------|--------|-----|----------|

Feature: ETS

Local-Advertisement:

TLV Type: Configuration/Recommendation

Willing: No

Credit Based Shaper: No

Maximum Traffic Classes supported: 3

| Code Point | Priority-Group |
|------------|----------------|
| 000        | 0              |
| 001        | 7              |
| 010        | 7              |
| 011        | 7              |
| 100        | 0              |
| 101        | 1              |
| 110        | 1              |
| 111        | 7              |

| Priority-Group | Percentage B/W |
|----------------|----------------|
| 0              | 40%            |
| 1              | 5%             |

| Priority-Group | Transmission Selection Algorithm |
|----------------|----------------------------------|
| 0              | Enhanced Transmission Selection  |
| 1              | Enhanced Transmission Selection  |

Peer-Advertisement:

TLV Type: Configuration

Willing: No

Credit Based Shaper: No

| Code Point | Priority-Group |
|------------|----------------|
| 000        | 0              |
| 001        | 7              |
| 010        | 7              |
| 011        | 7              |
| 100        | 0              |
| 101        | 1              |
| 110        | 1              |
| 111        | 7              |

| Priority-Group | Percentage B/W |
|----------------|----------------|
| 0              | 40%            |
| 1              | 5%             |

| Priority-Group | Transmission Selection Algorithm |
|----------------|----------------------------------|
| 0              | Enhanced Transmission Selection  |
| 1              | Enhanced Transmission Selection  |

Peer-Advertisement:

TLV Type: Recommendation

| Code Point | Priority-Group |
|------------|----------------|
| 000        | 0              |
| 001        | 7              |
| 010        | 7              |
| 011        | 7              |
| 100        | 0              |

|                |                                  |
|----------------|----------------------------------|
| 101            | 1                                |
| 110            | 1                                |
| 111            | 7                                |
| Priority-Group | Percentage B/W                   |
| 0              | 40%                              |
| 1              | 5%                               |
| Priority-Group | Transmission Selection Algorithm |
| 0              | Enhanced Transmission Selection  |
| 1              | Enhanced Transmission Selection  |

**show dcbx neighbors terse (QFX Series)**

```

user@switch> show dcbx neighbors terse
Interface Parent PFC ETS ETS Version
Interface Rec
xe-0/0/8.0 - Enabled Advt Advt IEEE (Auto)
xe-0/0/9.0 - Disabled Disabled 1.01
xe-0/0/11.0 ae0.0 Enabled Advt Advt IEEE (Auto)
xe-0/0/12.0 ae0.0 Enabled Advt Advt IEEE (Auto)
xe-0/0/32.0 - Enabled Advt Not Advt IEEE
xe-0/0/36.0 - Not Advt Advt Advt IEEE

```

**show dcbx neighbors (EX4500 Switch: FCoE Interfaces on Both Local and Peer with PFC Configured Compatibly)**

```

user@switch> show dcbx neighbors interface xe-0/0/14

Interface : xe-0/0/14.0 - Parent Interface: ae0.0
Protocol-State: in-sync

Local-Advertisement:
 Operational version: 0
 sequence-number: 6, acknowledge-id: 6

Peer-Advertisement:
 Operational version: 0
 sequence-number: 6, acknowledge-id: 6

Feature: PFC, Protocol-State: in-sync

Operational State: Enabled

Local-Advertisement:
 Enable: Yes, Willing: No, Error: No
 Maximum Traffic Classes capable to support PFC: 6

Code Point Admin Mode
000 Disabled
001 Disabled
010 Disabled
011 Enabled
100 Disabled
101 Disabled
110 Disabled
111 Disabled

```

## Peer-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 6

| Code Point | Admin Mode |
|------------|------------|
| 000        | Disabled   |
| 001        | Disabled   |
| 010        | Disabled   |
| 011        | Enabled    |
| 100        | Disabled   |
| 101        | Disabled   |
| 110        | Disabled   |
| 111        | Disabled   |

Feature: Application, Protocol-State: in-sync

## Local-Advertisement:

Enable: Yes, Willing: No, Error: No <<< Error bit will not be set as there is no miss configuration between local and peer.

| Appl-Name | Ethernet-Type | Socket-Number | Priority-Map | Status  |
|-----------|---------------|---------------|--------------|---------|
| FCoE      | 0x8906        |               | 00001000     | Enabled |

## Peer-Advertisement:

Enable: Yes, Willing: No, Error: No

| Status  | Appl-Name | Ethernet-Type | Socket-Number | Priority-Map |
|---------|-----------|---------------|---------------|--------------|
| Enabled | FCoE      | 0x8906        |               | 00001000     |

**show dcbx neighbors (EX4500 Switch: DCBX Interfaces on Local and Peer Are Configured Compatibly with iSCSI Application)**

user@switch&gt; show dcbx neighbors interface xe-0/0/14

Interface : xe-0/0/14.0 - Parent Interface: ae0.0

Protocol-State: in-sync

Active-application-map: iscsi-map

## Local-Advertisement:

Operational version: 0

sequence-number: 9, acknowledge-id: 12

## Peer-Advertisement:

Operational version: 0

sequence-number: 12, acknowledge-id: 9

Feature: PFC, Protocol-State: in-sync

Operational State: Enabled



## Local-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 6

| Code Point | Admin Mode |
|------------|------------|
| 000        | Disabled   |
| 001        | Disabled   |
| 010        | Disabled   |
| 011        | Enabled    |
| 100        | Disabled   |
| 101        | Disabled   |
| 110        | Disabled   |
| 111        | Disabled   |

## Peer-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 6

| Code Point | Admin Mode |
|------------|------------|
| 000        | Disabled   |
| 001        | Disabled   |
| 010        | Disabled   |
| 011        | Enabled    |
| 100        | Disabled   |
| 101        | Disabled   |
| 110        | Disabled   |
| 111        | Disabled   |

Feature: Application, Protocol-State: in-sync

## Local-Advertisement:

Enable: Yes, Willing: No, Error: No

| Appl-Name | Ethernet-Type | Socket-Number | Priority-Map | Status  |
|-----------|---------------|---------------|--------------|---------|
| FCoE      | 0x8906        |               | 00001000     | Enabled |
| iscsi     |               | 3260          | 00100000     | Enabled |

## Peer-Advertisement:

Enable: Yes, Willing: No, Error: No

| Appl-Name | Ethernet-Type | Socket-Number | Priority-Map | Status  |
|-----------|---------------|---------------|--------------|---------|
| FCoE      | 0x8906        |               | 00001000     | Enabled |
| iscsi     |               | 3260          | 00100000     | Enabled |

**show dcbx neighbors (EX4500 Switch: Includes ETS)**

user@switch&gt; show dcbx neighbors interface xe-0/0/3

Interface : xe-0/0/3.0  
 Protocol-State: in-sync  
 Active-application-map: map\_iscsi

## Local-Advertisement:

Operational version: 0

sequence-number: 1, acknowledge-id: 5

Peer-Advertisement:

Operational version: 0

sequence-number: 5, acknowledge-id: 1

Feature: PFC, Protocol-State: in-sync

Operational State: Enabled

Local-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 6

| Code Point | Admin Mode |
|------------|------------|
| 000        | Enabled    |
| 001        | Enabled    |
| 010        | Disabled   |
| 011        | Disabled   |
| 100        | Disabled   |
| 101        | Disabled   |
| 110        | Disabled   |
| 111        | Disabled   |

Peer-Advertisement:

Enable: Yes, Willing: Yes, Error: No

Maximum Traffic Classes capable to support PFC: 8

| Code Point | Admin Mode |
|------------|------------|
| 000        | Enabled    |
| 001        | Disabled   |
| 010        | Disabled   |
| 011        | Disabled   |
| 100        | Enabled    |
| 101        | Disabled   |
| 110        | Disabled   |
| 111        | Disabled   |

Feature: Application, Protocol-State: in-sync

Local-Advertisement:

Enable: Yes, Willing: No, Error: No

| App1-Name | Ethernet-Type | Socket-Number | Priority-Map | Status  |
|-----------|---------------|---------------|--------------|---------|
| FCoE      | 0x8906        |               | 00000001     | Enabled |
| iscsi     |               | 3260          | 00000010     | Enabled |

Peer-Advertisement:

Enable: Yes, Willing: Yes, Error: No

| App1-Name | Ethernet-Type | Socket-Number | Priority-Map | Status  |
|-----------|---------------|---------------|--------------|---------|
| FCoE      | 0x8906        |               | 00010000     | Enabled |
| iscsi     |               | 3260          | 00010000     | Enabled |

Feature: ETS, Protocol-State: in-sync

Operational State: Enabled

## Local-Advertisement:

Enable: Yes, Willing: No, Error: No  
Maximum Traffic Classes supported : 3

| Code Point | Priority-Group |
|------------|----------------|
| 000        | 7              |
| 001        | 7              |
| 010        | 7              |
| 011        | 7              |
| 100        | 7              |
| 101        | 7              |
| 110        | 7              |
| 111        | 7              |

| Priority-Group | Percentage B/W |
|----------------|----------------|
| 7              | 100%           |

## Peer-Advertisement:

Enable: Yes, Willing: Yes, Error: No  
Maximum Traffic Classes supported : 8

| Code Point | Priority-Group |
|------------|----------------|
| 000        | 0              |
| 001        | 1              |
| 010        | 0              |
| 011        | 0              |
| 100        | 2              |
| 101        | 0              |
| 110        | 0              |
| 111        | 0              |

| Priority-Group | Percentage B/W |
|----------------|----------------|
| 0              | 30%            |
| 1              | 40%            |
| 2              | 30%            |



## PART 95

# Index

- [Index on page 7423](#)



# Index

## Symbols

|                                       |                                                                                      |
|---------------------------------------|--------------------------------------------------------------------------------------|
| !                                     | regular expression<br>operator.....1827, 1828, 1849, 1850<br>system logging.....1552 |
| #                                     | comments in configuration statements.....cxlii                                       |
| \$                                    | regular expression<br>operator.....1827, 1828, 1849, 1850<br>system logging.....1552 |
| ( )                                   | regular expression<br>operator.....1827, 1828, 1849, 1850<br>system logging.....1552 |
| ( ), in syntax descriptions.....      | cxlii                                                                                |
| *                                     | regular expression operator.....1827, 1849<br>system logging.....1552                |
| +                                     | regular expression operator.....1827, 1849<br>system logging.....1552                |
| .                                     | regular expression operator.....1827, 1849<br>system logging.....1552                |
| /var/log/mib2d file.....              | 1518                                                                                 |
| /var/log/snmpd file.....              | 1518                                                                                 |
| 802.1Q VLANs                          | mixed VLAN tagging.....2368, 2370                                                    |
| 802.3ad statement.....                | 3046                                                                                 |
| < >, in syntax descriptions.....      | cxlii                                                                                |
| ?                                     | regular expression operator<br>system logging.....1552                               |
| [ ]                                   | regular expression operator<br>system logging.....1552                               |
| [ ], in configuration statements..... | cxlii                                                                                |
| \                                     | regular expression<br>operator.....1827, 1828, 1849, 1850                            |

|                                       |                                                                                      |
|---------------------------------------|--------------------------------------------------------------------------------------|
| ^                                     | regular expression<br>operator.....1827, 1828, 1849, 1850<br>system logging.....1552 |
| { }, in configuration statements..... | cxlii                                                                                |
|                                       | regular expression operator<br>system logging.....1552                               |
| (pipe), in syntax descriptions.....   | cxlii                                                                                |

## A

|                                               |                                                                      |
|-----------------------------------------------|----------------------------------------------------------------------|
| ABRs See area border routers                  |                                                                      |
| accept-data statement.....                    | 2684                                                                 |
| usage guidelines.....                         | 2648                                                                 |
| accept-remote-nexthop statement.....          | 3966                                                                 |
| usage guidelines.....                         | 3833                                                                 |
| accept-remote-source statement                | usage guidelines.....5602                                            |
| access option                                 | show route protocol command.....3519, 6404                           |
| access privilege levels                       | configuration example.....1831<br>operational mode commands.....1832 |
| configuring.....                              | 1809                                                                 |
| operational mode commands.....                | 30                                                                   |
| login classes.....                            | 1799                                                                 |
| user accounts.....                            | 1822, 1846                                                           |
| access statement.....                         | 1911                                                                 |
| access-end statement                          | login class.....270                                                  |
| access-internal option                        | show route protocol command.....3519, 6404                           |
| access-start statement                        | login class.....270                                                  |
| accounting statement.....                     | 1912                                                                 |
| authentication                                | usage guidelines.....1891                                            |
| IGMP.....                                     | 5618                                                                 |
| IGMP interface.....                           | 5618                                                                 |
| accounting-options statement.....             | 1913                                                                 |
| accounting-port statement                     | RADIUS servers.....314                                               |
| accounting-server statement.....              | 1915                                                                 |
| accounting-stop-on-access-deny statement..... | 1916                                                                 |
| accounting-stop-on-failure statement.....     | 1917                                                                 |
| action statement.....                         | 6100                                                                 |
| actions                                       | routing policy.....4529                                              |

|                                                |            |                                             |            |
|------------------------------------------------|------------|---------------------------------------------|------------|
| activate command                               |            | advertise-interval statement.....           | 2685       |
| usage guidelines.....                          | 65         | usage guidelines.....                       | 2643       |
| active routes.....                             | 3793       | advertise-peer-as statement.....            | 3970       |
| active statement                               |            | usage guidelines.....                       | 3773       |
| aggregate routes.....                          | 3270       | advertisement intervals, VRRP.....          | 2643       |
| generated routes.....                          | 3270       | advertisement messages, LDP.....            | 4808       |
| static routes.....                             | 3270       | advertisement-hold-time statement.....      | 5027       |
| active-source-limit statement.....             | 5768       | advertisement-interval statement.....       | 1918       |
| usage guidelines.....                          | 5604       | agent-address statement.....                | 1605, 1919 |
| adaptive.....                                  | 2656       | agent-id statement.....                     | 1591       |
| adaptive rerouting.....                        | 5024       | aggregate routes.....                       | 3272       |
| adaptive statement.....                        | 5024       | graceful restart.....                       | 2579       |
| add-path statement                             |            | preferences.....                            | 3786       |
| BGP                                            |            | aggregate statement.....                    | 3272       |
| usage guidelines.....                          | 3844       | RSVP.....                                   | 5306       |
| address (tracing flag).....                    | 4877       | aggregated Ethernet interfaces without link |            |
| address statement.....                         | 2973       | protection                                  |            |
| anycast RPs.....                               | 5689       | periodically rebalancing subscribers .....  | 2875       |
| usage guidelines.....                          | 5552, 5610 | rebalancing subscribers.....                | 2875       |
| local RPs.....                                 | 5690       | aggregated-devices statement.....           | 3042       |
| SNMPv3.....                                    | 1604       | aggregated-ether-options statement.....     | 3043       |
| static RPs.....                                | 5691       | aggregation, RSVP.....                      | 5306       |
| usage guidelines.....                          | 5547       | aggregator statement.....                   | 3274       |
| usage guidelines.....                          | 2824       | AIGP                                        |            |
| address-mask statement.....                    | 1605       | BGP.....                                    | 3725       |
| addresses                                      |            | aigp statement                              |            |
| associating with LSPs.....                     | 5048       | BGP                                         |            |
| egress router address.....                     | 5101       | usage guidelines.....                       | 3725       |
| ingress router.....                            | 5043       | aigp-originate statement                    |            |
| IP addresses.....                              | 34         | BGP                                         |            |
| multicast.....                                 | 5394       | usage guidelines.....                       | 3725       |
| router source addresses.....                   | 20, 222    | alarm conditions.....                       | 18         |
| adjacency.....                                 | 2617, 2627 | Alarm MIB.....                              | 1474       |
| adjust-interval statement.....                 | 5025       | alarm statement.....                        | 2977       |
| adjust-threshold statement.....                | 5025       | RMON.....                                   | 1606       |
| adjust-threshold-overflow-limit statement..... | 5026       | usage guidelines.....                       | 18         |
| adjust-threshold-underflow-limit               |            | alarms, displaying                          |            |
| statement.....                                 | 5026       | chassis.....                                | 459        |
| admin-down statement.....                      | 5027       | health monitor.....                         | 1749       |
| admin-group statement                          |            | RMON.....                                   | 1759       |
| bypass LSPs.....                               | 5304       | system.....                                 | 1027       |
| administrative distance.....                   | 3786       | alert (system logging severity level).....  | 3333       |
| BGP See preference statement                   |            | algorithm statement                         |            |
| administrative scoping.....                    | 5394       | BFD authentication.....                     | 5692       |
| advertise-external statement.....              | 3967       | alias option for static-host-mapping        |            |
| usage guidelines.....                          | 3770       | statement.....                              | 238        |
| advertise-inactive statement.....              | 3969       | alias statement.....                        | 238        |
| usage guidelines.....                          | 3770, 3788 |                                             |            |



- 
- all (tracing flag).....2679, 3384, 4084, 5102
    - class of service.....7058
    - RSVP.....5338
  - all statement
    - IGMP snooping.....5644, 5666
  - all-failures (tracing flag)
    - spanning-tree protocols.....2357
  - allow-commands statement.....271
    - usage guidelines.....1802
  - allow-configuration statement.....272
  - allow-configuration-regexps statement
    - usage guidelines.....1802
  - allow-fragmentation statement.....2962
  - allow-subnet-mismatch statement.....4842
    - usage guidelines.....4838
  - allow-transients statement.....1574
  - allowed-days statement
    - login class.....272
  - allowed-mac statement.....6138
  - allowing commands to login classes.....1802
  - always compare, BGP MED option.....3667
  - always-compare-med option.....3793
  - Analyzer MIB.....1474, 1482
  - analyzer statement.....6262
  - annotate command.....65
  - announcement statement.....273
    - usage guidelines.....126
  - announcements
    - system login.....126
  - any (system logging facility).....1545, 1546
  - any (system logging severity level).....1546, 1547
  - any-sender statement
    - RIP.....4764
  - anycast RP.....5558
    - overview.....5551
  - anycast-pim statement.....5693
    - usage guidelines.....5552, 5610
  - application statement.....6561, 7127
  - apply-groups statement.....3973
  - apply-groups-except statement.....3973
  - apply-macro statement .....1575
  - apply-path statement
    - firewall filter match condition.....3909
  - archival configuration
    - displaying.....1057, 1391
  - archival statement.....1360, 1920
    - usage guidelines.....1334
  - archive router configuration.....1333
  - archive statement
    - all system log files.....1692
    - individual system log file.....1694
    - usage guidelines.....1543
  - archive-sites statement
    - configuration files.....1361, 1921
    - system log files.....1694
    - system logging
      - usage guidelines.....1543
    - usage guidelines.....1334
  - archiving files.....342, 1370
  - area border routers
    - backbone area See backbone area
    - description.....4404
    - overview.....4398
  - area statement .....4580
    - usage guidelines, backbone.....4406
    - usage guidelines, multiarea.....4408
  - area-range statement.....4582
    - usage guidelines.....4452
  - areas See area border routers; backbone area; NSSAs; stub areas
    - overview.....4397
  - arithmetic and relational operators
    - for monitor traffic command.....1717
  - arp.....217
    - statistics, displaying.....2394
    - status information, displaying.....2394
  - arp statement.....217
  - AS boundary routers
    - overview.....4398
  - AS external link advertisements.....4393
  - AS path
    - ignoring in route selection.....3796
  - AS paths
    - distribution of, displaying.....3392
    - domain information, displaying.....3396
    - matching regular expressions,
      - displaying.....3419
    - summary of, displaying.....3399
  - as-path (tracing flag).....4081
  - as-path statement.....3274
  - as-path-ignore
    - usage guidelines.....3793, 3796
  - as-path-ignore option.....4058
  - asm-override-ssm statement.....5619, 5761
  - ASN
    - BGP community routes, displaying.....3426

|                                           |                              |  |
|-------------------------------------------|------------------------------|--|
| ASs                                       |                              |  |
| configuring.....                          | 3276                         |  |
| paths                                     |                              |  |
| aggregate routes.....                     | 3274                         |  |
| generated routes.....                     | 3274, 3291                   |  |
| static routes.....                        | 3274                         |  |
| private, removing.....                    | 4065                         |  |
| ASs (autonomous systems)                  |                              |  |
| area border routers.....                  | 4404                         |  |
| breaking into confederations.....         | 3890                         |  |
| paths.....                                | 3596                         |  |
| operations, tracing.....                  | 4081                         |  |
| private, removing.....                    | 3804, 3805                   |  |
| stub areas See stub areas                 |                              |  |
| assert (tracing flag).....                | 5755                         |  |
| assert timeout                            |                              |  |
| configuring.....                          | 5588                         |  |
| assert-timeout statement.....             | 5694                         |  |
| usage guidelines.....                     | 5588                         |  |
| associate-backup-pe-groups statement..... | 5028                         |  |
| asymmetric-hold-time statement.....       | 2686                         |  |
| asynch (tracing flag)                     |                              |  |
| class of service.....                     | 7058                         |  |
| auth (tracing flag).....                  | 4784                         |  |
| authentication configuration              |                              |  |
| BFD.....                                  | 3255, 3822, 4490, 4711, 5502 |  |
| algorithm                                 |                              |  |
| BGP.....                                  | 3898                         |  |
| IS-IS.....                                | 4255                         |  |
| BGP.....                                  | 3597, 3971, 3974             |  |
| IS-IS.....                                | 4255                         |  |
| keychains                                 |                              |  |
| BGP.....                                  | 3898                         |  |
| IS-IS.....                                | 4255                         |  |
| MD5                                       |                              |  |
| BGP.....                                  | 3898                         |  |
| NTP authentication keys.....              | 132                          |  |
| order.....                                | 1859, 1869, 1884             |  |
| RADIUS.....                               | 1804, 1847, 1873, 1893, 1899 |  |
| RIPv2, MD5.....                           | 4699                         |  |
| RIPv2, plain-text passwords.....          | 4698                         |  |
| root password.....                        | 145, 1824, 1850              |  |
| RSVP.....                                 | 5305                         |  |
| shared user accounts.....                 | 1899                         |  |
| simple                                    |                              |  |
| RIP.....                                  | 4693                         |  |
| TACACS+ .....                             | 1804, 1847, 1899             |  |
| users.....                                | 1804, 1847                   |  |
| authentication statement                  |                              |  |
| BFD.....                                  | 5697                         |  |
| BFD protocol.....                         | 5695                         |  |
| login.....                                | 274, 294                     |  |
| usage guidelines.....                     | 125, 1821, 1845, 1853        |  |
| authentication, VRRP.....                 | 2642                         |  |
| authentication-algorithm statement        |                              |  |
| BGP.....                                  | 3976, 4843                   |  |
| usage guidelines.....                     | 3898                         |  |
| IS-IS                                     |                              |  |
| usage guidelines.....                     | 4256                         |  |
| usage guidelines.....                     | 3647                         |  |
| authentication-key statement.....         | 284, 2687                    |  |
| BGP                                       |                              |  |
| usage guidelines.....                     | 3898                         |  |
| BGP and BMP.....                          | 3979                         |  |
| IS-IS.....                                | 4263                         |  |
| usage guidelines.....                     | 4183, 4256                   |  |
| LDP.....                                  | 4845                         |  |
| usage guidelines.....                     | 4835                         |  |
| MSDP.....                                 | 5769                         |  |
| RIP.....                                  | 4765                         |  |
| usage guidelines.....                     | 4693                         |  |
| RSVP.....                                 | 5305                         |  |
| usage guidelines.....                     | 132, 2642, 3647              |  |
| authentication-key-chain                  |                              |  |
| usage guidelines.....                     | 3647                         |  |
| authentication-key-chain                  |                              |  |
| statement.....                            | 3980, 4264, 4846             |  |
| BGP                                       |                              |  |
| usage guidelines.....                     | 3898                         |  |
| IS-IS                                     |                              |  |
| usage guidelines.....                     | 4256                         |  |
| authentication-key-chains statement.....  | 6298                         |  |
| authentication-md5 statement.....         | 1607                         |  |
| authentication-none statement.....        | 1608                         |  |
| authentication-order statement.....       | 1922                         |  |
| usage guidelines.....                     | 1859, 1869, 1884             |  |
| authentication-password statement.....    | 1609                         |  |
| authentication-server statement.....      | 1923                         |  |
| authentication-sha statement.....         | 1610                         |  |
| authentication-type statement.....        | 2688                         |  |
| IS-IS.....                                | 4265                         |  |
| usage guidelines.....                     | 4183                         |  |
| RIP.....                                  | 4766                         |  |
| usage guidelines.....                     | 4693                         |  |
| usage guidelines.....                     | 2642                         |  |
| authorization See permissions             |                              |  |

authorization (system logging facility).....1545, 1546  
     option to facility-override statement.....1548  
 authorization statement.....1611, 1924  
     usage guidelines.....1503  
 auto-bandwidth statement.....5029  
 auto-negotiation statement.....2978  
 auto-RP  
     overview.....5561  
 auto-rp statement.....5696  
     usage guidelines.....5561  
 autobw-state (tracing flag).....5102  
 autoinstallation, displaying the status of.....324  
 automatic bandwidth allocation.....5029  
     LSPs.....5122  
     MPLS, statistics.....4986  
 autonomous system number *See* ASN  
 autonomous system paths *See* AS paths  
 autonomous-system statement.....3276  
 auxiliary statement.....218

**B**

backbone area  
     configuring.....4406  
     description.....4404  
     overview.....4398  
 backbone router  
     overview.....4398  
 backup.....2617, 2627  
 backup router configuration.....1333  
 backup-pe-group statement.....3278, 5030  
 backups statement.....3279  
 bandwidth  
     RSVP reservations.....5334  
 bandwidth statement.....3280  
     DDoS protection.....6146  
     fast reroute.....5031  
     link protection.....5307  
     multiclass LSPs.....5031  
     RSVP.....5307  
     signaled LSPs.....5031  
 bandwidth, allocating for LSPs.....5122  
 bandwidth-based metrics  
     OSPF.....4462  
 bandwidth-based-metrics statement  
     usage guidelines.....4466  
 bandwidth-level statement.....6121  
 bandwidth-limit statement.....6100  
 bandwidth-model statement.....5032  
 bandwidth-percentage statement.....6122

bandwidth-scale statement  
     DDoS protection.....6147  
 bandwidth-threshold statement.....2689  
     usage guidelines.....2646  
 Beacon  
     operation of, controlling.....364  
 best routes, displaying.....3421  
 BFD  
     authentication  
         configuration.....3255, 3822, 4490, 4711, 5502  
         configuring.....4484  
         protocol.....3237, 3811, 3812, 4482, 4701, 4702, 5500  
         revert timer.....5086  
     traceoptions statement  
         usage guidelines.....3241  
     with IBGP.....3811, 3812  
 BFD authentication  
     algorithm statement.....5692  
     authentication statement.....5695  
     for IS-IS.....4250  
     key-chain statement.....5721  
     loose-check statement.....5724  
     OSPFv2.....4488  
 bfd-liveness-detection statement  
     BFD  
         threshold.....4079  
         transmit-interval.....4086  
     BGP.....3981  
         usage guidelines.....3812  
     IS-IS.....4266  
     OSPF.....4586  
         usage guidelines.....4484  
     PIM.....5697, 5703  
         minimum-interval.....5728  
         threshold.....5753  
         transmit-interval.....5754  
         usage guidelines.....5500  
     RIP.....4767  
         usage guidelines.....4702  
     static routes.....3281  
         usage guidelines.....3237  
     usage guidelines.....3246  
 bfd-per-address-family  
     IS-IS.....4266  
 BGP *See* multipath  
     administrative distance.....3786  
     advertising multiple paths to a  
         destination.....3651, 3844  
     aggregator path attribute.....4052

|                                                        |                                                |                                                         |                        |
|--------------------------------------------------------|------------------------------------------------|---------------------------------------------------------|------------------------|
| AIGP attribute.....                                    | 3725                                           | NLRI.....                                               | 3597                   |
| AS numbers, peers.....                                 | 4060                                           | nonstop active routing.....                             | 2601                   |
| ASs See ASs                                            |                                                | open messages.....                                      | 3597, 3919, 4057       |
| authentication.....                                    | 3597, 3898, 3979                               | outbound route filter.....                              | 3773                   |
| authentication algorithm.....                          | 3976, 4843                                     | outbound route filters                                  |                        |
| authentication keychain.....                           | 3980                                           | interoperability.....                                   | 3285, 3986             |
| best external route                                    |                                                | overview.....                                           | 3593, 4955             |
| advertising.....                                       | 3770                                           | packets, tracing.....                                   | 4081                   |
| BFD.....                                               | 3811, 3812, 3971, 3974, 3981, 4018, 4030, 4070 | passive mode.....                                       | 3919                   |
| community ASN, displaying routes.....                  | 3426                                           | path attributes.....                                    | 3596, 3597             |
| community name, displaying routes.....                 | 3428                                           | peers.....                                              | 3594, 3665, 4046, 4956 |
| confederations.....                                    | 3293                                           | point-to-point peer session (configuration editor)..... | 3602                   |
| damping parameters                                     |                                                | policy, routing.....                                    | 3996, 4012             |
| clearing.....                                          | 4092                                           | precedence.....                                         | 3769                   |
| displaying.....                                        | 4130                                           | preferences.....                                        | 3786, 3787, 4063       |
| damping routes, displaying.....                        | 3430, 4132                                     | private AS.....                                         | 3804, 3805             |
| description.....                                       | 3625                                           | route reflection.....                                   | 3988, 4053             |
| EBGP IPv6 peering.....                                 | 3609                                           | router identifier.....                                  | 3369                   |
| enabling on router.....                                | 3985                                           | routes.....                                             | 3596                   |
| external (EBGP).....                                   | 3594, 4956                                     | routing tables                                          |                        |
| graceful restart.....                                  | 2579, 2662, 4001                               | delays in exchanging routes.....                        | 3772, 4054             |
| groups.....                                            | 3601, 3665, 4003                               | nonactive routes.....                                   | 3770, 3788, 3969       |
| general information, displaying.....                   | 4100                                           | retaining routes.....                                   | 4016                   |
| hold time.....                                         | 3597, 4010                                     | session drops.....                                      | 3919                   |
| identifier.....                                        | 3597                                           | set local AS number.....                                | 4023                   |
| ignoring the AS path attribute in route selection..... | 3796                                           | summary information, displaying.....                    | 4124                   |
| indirect next hops.....                                | 3833                                           | system log messages.....                                | 3919                   |
| injecting OSPF routes into BGP.....                    | 3766, 4529                                     | table                                                   |                        |
| internal.....                                          | 3624                                           | clearing.....                                           | 4096                   |
| internal (IBGP).....                                   | 3594, 4956                                     | TCP.....                                                | 3593, 4955             |
| IP address.....                                        | 3597                                           | TCP segment size.....                                   | 3912                   |
| IPv6                                                   |                                                | tracing operations.....                                 | 4081                   |
| logical systems.....                                   | 3609                                           | 4-byte AS events.....                                   | 3949                   |
| keepalive messages.....                                | 3598, 4081                                     | BFD protocol events.....                                | 3949                   |
| local address.....                                     | 3625, 4020                                     | damping operations.....                                 | 3949                   |
| local AS.....                                          | 3704, 3709, 3719                               | description.....                                        | 3949                   |
| local preference.....                                  | 3652, 3653                                     | graceful restart.....                                   | 3949                   |
| MED.....                                               | 3668, 3680, 3694                               | keepalive messages.....                                 | 3949                   |
| messages.....                                          | 3597                                           | NSR synchronization.....                                | 3949                   |
| monitoring.....                                        | 3959                                           | open PDUs.....                                          | 3949                   |
| MP-BGP.....                                            | 3997                                           | policy processing.....                                  | 3949                   |
| MTU discovery.....                                     | 4040                                           | protocol task processing.....                           | 3949                   |
| multihop sessions.....                                 | 4042                                           | protocol timer processing.....                          | 3949                   |
| multipath configuration.....                           | 3827, 3828                                     | refresh PDUs.....                                       | 3949                   |
| neighbors                                              |                                                | route information.....                                  | 3949                   |
| clearing connections.....                              | 4093                                           | state transitions.....                                  | 3949                   |
| displaying.....                                        | 2710, 4107                                     | update PDUs.....                                        | 3949                   |
| neighbors BGP, peers See BGP, peers                    |                                                | update messages.....                                    | 3597                   |

- version supported.....3593, 4955
- VPNs
  - preventing session flaps.....3919
  - with BFD.....3811, 3812
- BGP (Border Gateway Protocol)
  - confederations *See* BGP confederations
  - export policy for CLNS.....4237
  - internal peer session (configuration editor).....3625
  - peering sessions *See* BGP peers; BGP sessions
  - point-to-point internal peer session
    - logical systems.....3636
  - route reflectors *See* BGP route reflectors
  - route-flap damping.....3926
- BGP confederations
  - creating (configuration editor).....3891
  - description.....3890
  - route-flap damping.....3926
- BGP groups
  - confederations (configuration editor).....3891
- BGP LSPs
  - ping interval.....5112
- BGP Monitoring Protocol.....3287
  - displaying
    - statistics.....4098
- BGP monitoring protocol version 3 *See* BMP version 3
- BGP peers
  - external (configuration editor).....3602
  - internal.....3636
  - internal (configuration editor).....3625
  - point-to-point connections.....3601
- BGP route reflectors
  - cluster of clusters.....3875
  - creating (configuration editor).....3876
  - description.....3873
  - multiple clusters.....3874
- BGP routing information.....3959
- BGP sessions
  - internal.....3636
  - internal (configuration editor).....3625
  - point-to-point external (configuration editor).....3602
  - sample peering session.....3601
- bgp statement.....3985
- bgp-orf-cisco-mode
  - usage guidelines.....3773
- bgp-orf-cisco-mode statement.....3285, 3986
- Bidirectional Forwarding Detection *See* BFD
- Bidirectional Forwarding Detection (BFD), nonstop
  - active routing.....2604
- binding (tracing flag).....4877
- binding state of DHCP client
  - clearing.....6324, 6360
  - displaying.....6341, 6370
- binding state of DHCPv6 client
  - clearing.....6329, 6365
  - displaying.....6350, 6378
- BMP
  - authentication algorithm.....3976, 4843
  - authentication keychain.....3980
  - tracing operations.....4084
- bmp statement.....3287
  - usage guidelines.....3647
- BMP version 3
  - configuration.....3647
- boilerplate
  - commit scripts.....1434
  - op scripts.....1436
- boot messages, displaying.....1038
- boot server
  - NTP.....137
- boot-server statement
  - NTP.....285
  - usage guidelines.....137
- bootstrap (tracing flag).....5755
- bootstrap messages.....5567
- bootstrap routers
  - overview.....5567
- bootstrap routers, displaying.....5881
- bootstrap statement.....5698
- bootstrap-export statement.....5699
- bootstrap-import statement.....5700
- bootstrap-priority statement.....5701
- Border Gateway Protocol *See* BGP *See* BGP
- bpdu (tracing flag).....2357
- BPDU protection for spanning-tree protocols
  - on individual ports.....2205
- bpdu-block statement.....2324, 2331
- bpdu-block-on-edge statement.....2325, 2332
- bpdu-timeout-action statement.....2326, 2333
- braces, in configuration statements.....cxlii
- brackets
  - angle, in syntax descriptions.....cxlii
  - square, in configuration statements.....cxlii
- bridge domain
  - dual VLAN tags.....2310

|                                                      |                       |                                                    |                                 |
|------------------------------------------------------|-----------------------|----------------------------------------------------|---------------------------------|
| bridge-detection-state-machine (tracing flag).....   | 2357                  | call admission control See CAC                     |                                 |
| bridge-options statement.....                        | 2304                  | categories statement.....                          | 1613, 1925                      |
| bridge-priority statement.....                       | 2327, 2334            | usage guidelines.....                              | 1504                            |
| brief statement.....                                 | 3289                  | CBT                                                |                                 |
| system logging.....                                  | 1703                  | defined.....                                       | 5399                            |
| usage guidelines.....                                | 1538                  | issues.....                                        | 5491                            |
| broadcast                                            |                       | CCC                                                |                                 |
| NTP.....                                             | 131, 133, 134         | route forwarding-table, displaying.....            | 5146, 5166                      |
| synchronizing NTP.....                               | 135                   | centralized ppm statement.....                     | 3290                            |
| broadcast messages, synchronizing NTP.....           | 287                   | certificates                                       |                                 |
| broadcast statement.....                             | 286                   | installed, displaying.....                         | 1052                            |
| usage guidelines.....                                | 134, 2824             | certificates statement.....                        | 6302                            |
| broadcast-client statement.....                      | 287                   | certification-authority statement.....             | 6303                            |
| usage guidelines.....                                | 135                   | change-log (system logging facility).....          | 1545, 1546                      |
| BSR                                                  |                       | change-type statement.....                         | 295                             |
| policy, import.....                                  | 5716                  | usage guidelines.....                              | 1824, 1850                      |
| bucket-size statement.....                           | 1612                  | chassis                                            |                                 |
| ICMPv4.....                                          | 225                   | alarm conditions, displaying.....                  | 459                             |
| usage guidelines.....                                | 86                    | configuration                                      |                                 |
| buffer-size statement.....                           | 7092                  | alarm conditions.....                              | 18                              |
| buffers, displaying system.....                      | 1045                  | craft interface display messages                   |                                 |
| built-in Ethernet interfaces.....                    | 2796                  | clearing the display of.....                       | 333                             |
| See also Fast Ethernet ports; Gigabit Ethernet ports |                       | environmental information, displaying.....         | 476                             |
| burst statement                                      |                       | firmware version, displaying.....                  | 692                             |
| DDoS protection.....                                 | 6148                  | installed hardware, displaying.....                | 743                             |
| burst-scale statement                                |                       | location, displaying.....                          | 945                             |
| DDoS protection.....                                 | 6149                  | MAC addresses, displaying.....                     | 949                             |
| burst-size-limit statement.....                      | 6101                  | serial numbers, displaying.....                    | 743                             |
| bypass LSPs, testing.....                            | 5347                  | switch fabric errors                               |                                 |
| bypass statement                                     |                       | displaying.....                                    | 580                             |
| RSVP.....                                            | 5308                  | switch fabric status                               |                                 |
| static LSP.....                                      | 5033                  | FPCs, displaying.....                              | 585                             |
| bypass-aggregate statement                           |                       | SIBs, displaying.....                              | 629                             |
| DDoS protection.....                                 | 6150                  | switch fabric topology, displaying.....            | 647                             |
|                                                      |                       | Chassis Definitions for Router Model MIB.....      | 1474                            |
| <b>C</b>                                             |                       | Chassis MIB.....                                   | 1474, 1482                      |
| ca-name statement.....                               | 6301                  | chassis software process.....                      | 1794                            |
| cable-length statement.....                          | 7128                  | chassis statement.....                             | 3045                            |
| cables                                               |                       | chassis-scheduler (tracing flag)                   |                                 |
| console port, connecting.....                        | 147, 1811, 1819, 1841 | class of service.....                              | 7058                            |
| Ethernet rollover,                                   |                       | chassisd process.....                              | 1794                            |
| connecting.....                                      | 147, 1811, 1819, 1841 | check-zero statement.....                          | 4769                            |
| CAC                                                  |                       | checksum                                           |                                 |
| displaying for LSPs.....                             | 5140, 5160            | calculating for a                                  |                                 |
| cache (tracing flag).....                            | 5755                  | file.....                                          | 344, 345, 346, 1372, 1373, 1374 |
| cache-size statement.....                            | 6299                  | displaying values for all files.....               | 1030                            |
| cache-timeout-negative statement .....               | 6300                  | for IS-IS.....                                     | 4239                            |
|                                                      |                       | checksum statement.....                            | 219, 1576, 4268                 |
|                                                      |                       | circuit cross-connect (CCC), graceful restart..... | 2579                            |

- circuit-id statement.....6132
- Cisco non-deterministic, BGP MED option.....3667
- cisco-non-deterministic option.....3793, 4058
- class statement.....7062, 7063
  - assigning to user.....276
  - login.....275
  - usage
    - guidelines.....125, 1821, 1823, 1845, 1848, 1853
- Class-of-Service MIB.....1474, 1482
- class-of-service statement.....7054
  - bypass LSPs.....5309
  - ingress routers.....5034
  - signaled LSPs.....5034
  - static LSPs.....5034
- classifiers statement.....7064
- classifiers, CoS.....7231
- cleanup, storage space.....399
- clear
  - snmp history.....1738
- clear (ospf | ospf3) database command.....4618
- clear (ospf | ospf3) database-protection
  - command.....4621
- clear (ospf | ospf3) io-statistics command.....4622
- clear (ospf | ospf3) neighbor command.....4623
- clear (ospf | ospf3) overload command.....4627
- clear (ospf | ospf3) statistics command.....4625
- clear bgp damping command.....4092
- clear bgp neighbor command.....4093
- clear bgp table command.....4096
- clear chassis display message command.....333
- clear ddos-protection protocols command.....6180
- clear dhcp relay binding command.....6360
- clear dhcp relay statistics command.....6362
- clear dhcp server binding command.....6324
- clear dhcp server statistics command.....6327
- clear dhcpv6 relay binding command.....6365
- clear dhcpv6 relay statistics command.....6368
- clear dhcpv6 server binding command.....6329
- clear dhcpv6 server statistics command.....6331
- clear error bpdu command.....2418
- clear ethernet-switching port-error
  - command.....6178
- clear fip vlan-discovery statistics command.....6576
- clear firewall command.....6166
- clear igmp membership command.....5790
- clear igmp statistics command.....5793
- clear igmp-snooping membership command.....5814
- clear igmp-snooping statistics command.....5815
- clear ipv6 neighbors command.....3391
- clear isis adjacency command.....4330
- clear isis database command.....4332
- clear isis overload command.....4334
- clear isis statistics command.....4336
- clear ldp neighbor command.....4884
- clear ldp session command.....4885
- clear ldp statistics command.....4886
- clear lldp neighbors command.....1994
- clear lldp statistics command.....1995
- clear log command.....332, 1368
- clear mpls container-lsp command.....5342, 5353
- clear msdp cache command.....5932
- clear msdp statistics command.....5933
- clear multicast bandwidth-admission
  - command.....5827
- clear multicast scope command.....5829
- clear multicast sessions command.....5830
- clear multicast statistics command.....5831
- clear pim join command.....5832
- clear pim register command.....5834
- clear pim statistics command.....5836
- clear rip general-statistics command.....4790
- clear rip statistics command.....4791
- clear rsvp session command.....5344
- clear rsvp statistics command.....5346
- clear sflow collector statistics command.....1730
- clear snmp history command.....1738
- clear snmp statistics command.....1739
- clear spanning-tree statistics command.....2419
- clear system commit command.....336, 1369
- clear system reboot command.....337
- CLI
  - command history
    - displaying.....435
  - comparing configuration versions.....1340, 1349
  - configuration mode
    - description.....63
  - current working directory
    - displaying.....434
  - exiting to create UNIX-level shell.....1266
  - permissions, displaying.....430
  - settings, displaying.....428
- client list
  - adding to SNMP community.....1505
- client mode, NTP.....131, 133
- client-list statement.....1613, 1925
  - usage guidelines.....1505
- client-list-name statement.....1614, 1926
  - usage guidelines.....1505



|                                             |                  |                                             |                           |
|---------------------------------------------|------------------|---------------------------------------------|---------------------------|
| clients statement.....                      | 1614, 1926       | committed-information-rate statement.....   | 6105                      |
| usage guidelines.....                       | 1503             | communities                                 |                           |
| CLNP (Connectionless Network Protocol)..... | 4157             | aggregate routes.....                       | 3291                      |
| CLNS (Connectionless Network Service) VPNs  |                  | static routes.....                          | 3291                      |
| BGP export policy.....                      | 4237             | community ASN, displaying routes.....       | 3426                      |
| IS-IS.....                                  | 4237             | community name, displaying routes.....      | 3428                      |
| cluster statement.....                      | 3988             | community statement                         |                           |
| usage guidelines.....                       | 3876             | aggregate routes.....                       | 3291                      |
| clusters <i>See</i> BGP route reflectors    |                  | generated routes.....                       | 3291                      |
| code-point statement.....                   | 7066, 7129, 7130 | RMON.....                                   | 1617                      |
| code-point-aliases statement.....           | 7067             | SNMP.....                                   | 1616, 1928                |
| code-points statement.....                  | 7068             | usage guidelines.....                       | 1503                      |
| collector statement.....                    | 1592             | static routes.....                          | 3291                      |
| color statement                             |                  | community string, SNMP.....                 | 1503                      |
| aggregate routes.....                       | 3345             | community-name statement.....               | 1618                      |
| generated routes.....                       | 3345             | compact flash                               |                           |
| static routes.....                          | 3345             | displaying usage.....                       | 167                       |
| color-aware statement.....                  | 6102             | compare command                             |                           |
| color-blind statement.....                  | 6103             | usage guidelines.....                       | 1340, 1349                |
| command output                              |                  | comparing files.....                        | 347, 1375                 |
| filtering                                   |                  | complete sequence number PDUs, IS-IS.....   | 4160                      |
| comparing configuration                     |                  | compress-configuration-files statement..... | 220                       |
| versions.....                               | 1340, 1349       | usage guidelines.....                       | 14, 1351                  |
| command statement.....                      | 1577             | compressing configuration files.....        | 14, 220, 1351             |
| commands                                    |                  | compressing files.....                      | 342, 1370                 |
| allowing or denying to login classes.....   | 1802             | confederation statement.....                | 3293                      |
| filenames, specifying.....                  | 40               | usage guidelines.....                       | 3891                      |
| options.....                                | 62               | confederations.....                         | 3293, 3890 <i>See</i> BGP |
| URLs, specifying.....                       | 40               | c o n f e d e r a t i o n s                 |                           |
| comments, in configuration statements.....  | cxlii            | config-internal (tracing flag).....         | 2679, 3384                |
| commit command.....                         | 329              | configuration                               |                           |
| usage guidelines.....                       | 65               | activating.....                             | 1338                      |
| commit operations, pending                  |                  | comparing with previous.....                | 1340, 1349                |
| clearing.....                               | 336, 1369        | deleting rescue configuration.....          | 382, 1386                 |
| displaying.....                             | 1054, 1388       | displaying                                  |                           |
| commit scripts                              |                  | archival configuration.....                 | 1057, 1391                |
| boilerplate.....                            | 1434             | previous configuration.....                 | 1302, 1394                |
| commands for monitoring.....                | 1445             | rescue configuration.....                   | 1058, 1392                |
| flow of operation illustrated.....          | 1424             | files <i>See</i> configuration files        |                           |
| multiple.....                               | 1429             | merging current and new.....                | 1326                      |
| output, displaying.....                     | 1445             | previous, displaying.....                   | 1339                      |
| persistent configuration changes.....       | 1430             | replacing.....                              | 1326                      |
| transient configuration changes.....        | 1430             | saving rescue configuration.....            | 383, 1387                 |
| using multiple.....                         | 1429             | saving to file.....                         | 1331, 1343                |
| commit statement.....                       | 1578             | storage of previous.....                    | 1314                      |
| commit synchronize command.....             | 329              | syntax, verifying.....                      | 1268, 1396                |
| commit synchronize statement.....           | 2677             | configuration files                         |                           |
| commit-delay statement.....                 | 1615, 1927       | compressing.....                            | 14, 220, 1351             |
| committed-burst-size statement.....         | 6104             | filename, specifying.....                   | 40                        |



- saving to files.....1331, 1343
- URL, specifying.....40
- Configuration Management MIB.....1475, 1483
- configuration mode commands
  - commit script.....1445
- configuration mode, CLI
  - commands
    - activate.....65
    - annotate.....65
    - commit.....65
    - copy.....65
    - deactivate.....65
    - delete.....65
    - edit.....65
    - exit.....65
    - extension.....65
    - help.....65
    - insert.....65
    - load.....65
    - paste.....66
    - quit.....66
    - rollback.....66, 1337
    - run.....66
    - save.....66
    - set.....66
    - show.....66
    - status.....66
    - top.....66
    - up.....66
    - update.....66
  - configuration hierarchy, description.....68
  - description.....63
  - identifier, description.....67
  - statement
    - container.....68
    - description.....67
    - leaf.....68
  - top-level statements, interpreting.....67
- configuration statement.....1363, 1929
  - usage guidelines.....1333
- configuration-name statement.....2328, 2335
- configure command
  - usage guidelines.....59
- configuring
  - VLANs.....2120
- conflict-log (system logging facility).....1545, 1546
- congestion-notification-profile statement.....7132
- connection (tracing flag).....5102
- connection-detail (tracing flag).....5102
- connection-limit statement.....1566, 1930
- connection-mode statement.....3989
  - usage guidelines.....3648
- connections
  - IP sockets, displaying active.....1060
  - SSH, opening.....2028
  - testing
    - general connections.....360, 1724
    - MPLS BGP connections.....5112
    - MPLS Layer 2 circuit connections.....5114
    - MPLS Layer 3 VPN connections.....5117
    - MPLS LDP connections.....4887
    - MPLS LSP-endpoint connections.....5120
    - MPLS RSVP connections.....5347
- connectivity
  - unidirectional (RIP).....4681
- console port
  - adapter.....147, 1810, 1819, 1841
- console statement
  - physical port.....221
  - system logging.....1695
  - usage guidelines.....1534
- Constrained Shortest Path First See CSPF
- constrained-path LSPs
  - computation
    - disabling.....5060
    - overview.....4971
- Constraint-based routing.....5290
- contact statement.....1619, 1931
- Container LSP
  - configuration statements.....5294
  - implementation.....5279
  - implementation overview.....5277
  - link protection.....5284
  - LSP merging.....5282
  - LSP splitting.....5280
  - network performance impact.....5298
  - node protection.....5284
  - normalization.....5285
  - overview.....5272
  - protection.....5284
  - sampling.....5291
  - supported features.....5298
  - terminology.....5279
  - unsupported features.....5298
- container-label-switched-path.....5310
- context-identifier statement
  - OSPF.....4585
- control-channel statement.....2260

|                                                |                  |
|------------------------------------------------|------------------|
| conventions                                    |                  |
| text and syntax.....                           | cxli             |
| copy command                                   |                  |
| usage guidelines.....                          | 59, 65           |
| copy-tos-to-outer-ip-header statement.....     | 2962             |
| Core Based Trees See CBT                       |                  |
| core dump files                                |                  |
| usage guidelines.....                          | 29               |
| viewing.....                                   | 34               |
| core-dumps                                     |                  |
| usage information, displaying.....             | 1079             |
| corouted-bidirectional statement               |                  |
| usage guidelines.....                          | 5035             |
| corouted-bidirectional-passive statement       |                  |
| usage guidelines.....                          | 5035             |
| CoS                                            |                  |
| byte count by interface hardware               |                  |
| displaying.....                                | 3152, 7346       |
| classifiers, monitoring.....                   | 7231             |
| code point aliases, monitoring.....            | 7236             |
| congestion-notification, displaying.....       | 7394             |
| forwarding class sets, displaying.....         | 7245             |
| forwarding classes, monitoring.....            | 7232             |
| forwarding table, displaying                   |                  |
| classifier information.....                    | 7253             |
| code point value to loss priority.....         | 7251             |
| code point value to queue number.....          | 7251             |
| configuration.....                             | 7247, 7296       |
| RED drop profiles.....                         | 7300             |
| rewrite rules.....                             | 7255             |
| rewrite rules, table identifiers.....          | 7257             |
| scheduler map information.....                 | 7302             |
| inputs and outputs overview.....               | 6639             |
| interfaces, displaying.....                    | 7167, 7258, 7304 |
| interfaces, monitoring.....                    | 7161             |
| loss priority.....                             | 7292             |
| mapping, displaying                            |                  |
| code point aliases to bit patterns.....        | 7241             |
| code point value to forwarding                 |                  |
| class.....                                     | 7238             |
| code point value to loss priority.....         | 7238             |
| forwarding classes to queue                    |                  |
| numbers.....                                   | 7243             |
| schedulers to forwarding classes.....          | 7333             |
| MIB.....                                       | 1474, 1482       |
| packet loss priority.....                      | 7292             |
| RED profile information                        |                  |
| displaying.....                                | 7293             |
| rewrite rules, displaying.....                 | 7288             |
| rewrite rules, monitoring.....                 | 7235             |
| scheduler map information, displaying.....     | 7333             |
| scheduler maps, monitoring.....                | 7291             |
| shared buffer, displaying.....                 | 7196             |
| traffic-control profile information,           |                  |
| displaying.....                                | 7336             |
| cos-adjustment (tracing flag)                  |                  |
| class of service.....                          | 7058             |
| cost statement.....                            | 2329, 2336       |
| CPU utilization, displaying.....               | 165              |
| craft interface                                |                  |
| alarm conditions                               |                  |
| overview.....                                  | 18               |
| craft interface display messages               |                  |
| clearing.....                                  | 333              |
| craft-lockout statement.....                   | 2981             |
| creating emergency boot device.....            | 157, 175         |
| credibility-protocol-preference                |                  |
| traffic engineering                            |                  |
| IS-IS.....                                     | 4324             |
| critical (system logging severity level).....  | 3333             |
| crl statement                                  |                  |
| ES PIC.....                                    | 6304             |
| csn (tracing flag).....                        | 4321             |
| csnp-interval statement.....                   | 4269             |
| CSPF                                           |                  |
| statistics, displaying.....                    | 5142, 5162       |
| cspf (tracing flag).....                       | 5103             |
| CSPF algorithm                                 |                  |
| online path computation                        |                  |
| disabling.....                                 | 5060             |
| cspf-link (tracing flag).....                  | 5103             |
| cspf-node (tracing flag).....                  | 5103             |
| curly braces, in configuration statements..... | cxlii            |
| current time, displaying.....                  | 1185             |
| current working directory                      |                  |
| displaying.....                                | 434              |
| customer support.....                          | cxliii           |
| contacting JTAC.....                           | cxliii           |
| <b>D</b>                                       |                  |
| daemon (system logging facility).....          | 1545, 1546       |
| option to facility-override statement.....     | 1548             |
| daemons See processes, software                |                  |
| damping.....                                   | 3990, 4081       |
| damping (tracing flag).....                    | 4081             |
| damping parameters, BGP                        |                  |
| clearing.....                                  | 4092             |
| displaying.....                                | 4130             |

- 
- damping routes, BGP
    - displaying.....3430, 4132
  - damping statement.....3990
    - usage guidelines.....3927, 3936
  - data flow.....5263
  - data-channel statement.....2262
  - data-encapsulation statement.....5770
    - usage guidelines.....5604
  - database description packets.....4392
  - database-protection statement
    - OSPF.....4590
  - dcbx statement.....6563, 7134
  - DDoS protection
    - clearing statistics and violations.....6180
    - configuration example.....6070
    - configuration overview.....6069
    - disabling policers and logging globally.....6074
    - flags for tracing operations.....6080
    - log file access for tracing operations.....6080
    - log file size and number.....6079
    - log filenames.....6079
    - logging
      - disabling globally.....6074
    - message severity levels for tracing
      - operations.....6081
    - overview.....6065
    - policers
      - aggregate.....6066
      - disabling globally.....6074
      - hierarchy.....6067
      - protocol.....6066
      - scaling.....6067
    - protocol information, displaying.....6182
    - protocol policer information, displaying.....6201
    - protocol policer violation information,
      - displaying.....6221
    - protocol statistics information,
      - displaying.....6208
    - regular expressions for tracing
      - operations.....6080
    - statistics, displaying.....6218
    - traffic priority
      - aggregate.....6066
    - verifying configuration.....6083
    - version, displaying.....6220
  - DDoS protection statements
    - bandwidth.....6146
    - bandwidth-scale.....6147
    - burst.....6148
    - burst-scale.....6149
    - bypass-aggregate.....6150
    - ddos-protection.....6151
    - disable-fpc.....6152
    - disable-logging.....6153
    - fpc.....6154
    - global.....6155
    - priority.....6156
    - protocols.....6157
    - traceoptions.....6162
  - ddos-protection statement
    - DDoS protection.....6151
  - deactivate command
    - usage guidelines.....65
  - dead-interval statement
    - usage guidelines.....4476
  - deaggregate statement.....4847
  - debug (system logging severity level).....3333
  - default route
    - configuring on logical systems.....4232
  - default setting
    - CoS.....6628, 6630
  - default-address-selection statement.....222
    - usage guidelines.....20
  - default-lsa statement
    - usage guidelines.....4418
  - default-metric statement
    - usage guidelines.....4414, 4418
  - default-peer statement.....5771
    - usage guidelines.....5604
  - defaults statement
    - aggregate statement.....3272
    - generate statement.....3309
    - static statement.....3377
  - delay statement
    - IS-IS.....4319
  - delay-med-update statement
    - usage guidelines.....3694
  - delete command
    - usage guidelines.....65
  - deleting
    - current rescue configuration (CLI configuration
      - editor).....1345
    - files.....351, 1379
    - licenses (CLI).....118
    - software packages.....1286
  - demand-circuit statement
    - OSPF
      - usage guidelines.....4439

|                                        |                        |
|----------------------------------------|------------------------|
| denial-of-service attacks              |                        |
| protecting against See DDoS Protection |                        |
| dense-groups statement.....            | 5702                   |
| usage guidelines.....                  | 5527                   |
| deny-commands statement.....           | 276                    |
| usage guidelines.....                  | 1802                   |
| deny-configuration statement.....      | 277                    |
| deny-configuration-regexps statement   |                        |
| usage guidelines.....                  | 1802                   |
| denying commands to login classes..... | 1802                   |
| description statement.....             | 2275, 2982, 3294, 3992 |
| MPLS.....                              | 5036                   |
| op script arguments.....               | 1579                   |
| op scripts.....                        | 1579                   |
| RMON.....                              | 1620                   |
| SNMP.....                              | 1619                   |
| usage guidelines.....                  | 3625                   |
| designated router.....                 | 5512                   |
| configuring.....                       | 4402                   |
| controlling election.....              | 4402                   |
| OSPF.....                              | 4400                   |
| destination statement.....             | 315                    |
| tunnels.....                           | 2963                   |
| usage guidelines.....                  | 1891, 2824             |
| destination-port statement             |                        |
| SNMP.....                              | 1620                   |
| usage guidelines.....                  | 1504                   |
| destination-profile statement          |                        |
| usage guidelines.....                  | 2824                   |
| detail (tracing flag modifier)         |                        |
| LDP.....                               | 4878                   |
| RSVP.....                              | 5339                   |
| detection-time statement               |                        |
| BFD.....                               | 3281                   |
| BGP.....                               | 3981, 3993             |
| IS-IS.....                             | 4266                   |
| PIM.....                               | 5703                   |
| device-count statement.....            | 3047                   |
| dfc (system logging facility).....     | 1545, 1546             |
| DHCP                                   |                        |
| extended server binding                |                        |
| clearing.....                          | 6324                   |
| displaying.....                        | 6341                   |
| extended server statistics             |                        |
| clearing.....                          | 6327                   |
| displaying.....                        | 6347                   |
| relay binding                          |                        |
| clearing.....                          | 6360                   |
| relay binding state                    |                        |
| displaying.....                        | 6370                   |
| relay statistics                       |                        |
| clearing.....                          | 6362                   |
| displaying.....                        | 6375                   |
| DHCP client                            |                        |
| binding state                          |                        |
| clearing.....                          | 6324, 6360             |
| displaying.....                        | 6341, 6370             |
| statistics                             |                        |
| clearing.....                          | 6327, 6362             |
| displaying.....                        | 6347                   |
| DHCP local server                      |                        |
| reconfigure client bindings.....       | 6332                   |
| verifying configuration.....           | 6256                   |
| DHCP local server statements           |                        |
| dhcp-local-server.....                 | 6278                   |
| DHCP option 82.....                    | 6132                   |
| DHCP relay                             |                        |
| verifying configuration.....           | 6257                   |
| DHCP relay agent statements            |                        |
| dhcp-relay.....                        | 2276, 6284             |
| forward-only.....                      | 6295                   |
| DHCP server                            |                        |
| overview.....                          | 6037                   |
| DHCP service                           |                        |
| statement.....                         | 6289                   |
| dhcp-local-server statement.....       | 6278                   |
| dhcp-relay statement.....              | 2276, 6284             |
| DHCPv6                                 |                        |
| extended server binding                |                        |
| clearing.....                          | 6329                   |
| displaying.....                        | 6350                   |
| extended server statistics             |                        |
| clearing.....                          | 6331                   |
| displaying.....                        | 6356                   |
| relay binding                          |                        |
| clearing.....                          | 6365                   |
| relay binding state                    |                        |
| displaying.....                        | 6378                   |
| relay statistics                       |                        |
| clearing.....                          | 6368                   |
| displaying.....                        | 6384                   |

- 
- DHCPv6 client
    - binding state
      - clearing.....6329
      - displaying.....6350, 6378
    - statistics
      - clearing.....6331, 6368
      - displaying.....6356
  - DHCPv6 local server
    - reconfigure client bindings.....6334
    - verifying configuration.....6257
  - DHCPv6 relay
    - verifying configuration.....6257
  - DHCPv6client
    - binding state
      - clearing.....6365
  - diagnosis
    - verifying RIP host reachability .....4691
  - DiffServ
    - classes, displaying for MPLS.....5144, 5164
  - diffserv-te statement.....5037
  - direct-access statement.....1579
  - directories
    - usage information, displaying.....1093
    - working, displaying.....434
  - disable (DCBX) statement.....6566, 7137
  - disable (tracing flag modifier).....4878
  - disable option to traceoptions statement
    - LDP.....4877
    - RSVP.....5338
  - disable statement.....2658, 4081
    - BGP.....3994
    - IGMP.....5619
    - usage guidelines.....5429
    - IGMP snooping.....5645, 5667
    - IS-IS.....4270
      - graceful restart.....4274
    - LDP.....4848
      - usage guidelines.....4810
    - link protection.....5311
    - LLDP.....1931
    - MLD
      - usage guidelines.....5466
    - MPLS.....5038
    - MSDP.....5772
    - OSPF.....4592
    - PIM family.....5704
    - PIM interfaces.....5704
    - PIM protocol.....5704
    - RSVP.....5311
    - RSVP graceful restart.....5311
    - STP.....2330, 2337
    - traceoption.....2679, 3384
    - usage guidelines
      - BGP.....2581
      - ES-IS.....2582
      - global.....2580
      - IS-IS.....2582
      - OSPF .....2583
      - OSPFv3.....2583
      - PIM.....2585
      - RIP.....2585
      - RIPng.....2585
  - disable-fpc statement
    - DDoS protection.....6152
  - disable-logging statement
    - DDoS protection.....6153
  - disable-timeout statement
    - STP.....2338
  - discard statement
    - aggregate routes.....3295
    - generated routes.....3295
  - discovery messages, LDP.....4807
  - disk space available, displaying.....1177
  - disk volume
    - FreeBSD upgrade.....94
  - distance-vector routing protocols.....4677
    - See also RIP
  - distributed denial-of-service See DDoS protection
  - distribution trees
    - RPT.....5579
    - shared.....5579
  - DNS
    - hostnames, displaying.....1002
  - DNS name servers.....15, 26
  - do-not-fragment statement
    - tunnel.....2964
  - documentation
    - comments on.....cxliii
  - dod-request-policy statement.....4849
  - Domain Name System See DNS
  - domain names on routers.....26
  - domain-name statement
    - router.....223
    - usage guidelines.....26
  - domain-search statement.....223
    - usage guidelines.....26
  - domains to be searched.....26, 223
  - downstream on demand, LDP.....4820

|                                                |          |                                                     |                                  |
|------------------------------------------------|----------|-----------------------------------------------------|----------------------------------|
| downstream-on-demand statement.....            | 4849     | enable statement                                    |                                  |
| usage guidelines.....                          | 4820     | routing options.....                                | 3319                             |
| dr-election-on-p2p statement.....              | 5705     | encapsulate-inner-vlan statement.....               | 2526                             |
| PIM                                            |          | encoding statement.....                             | 6304                             |
| usage guidelines.....                          | 5499     | encrypted passwords.....                            | 145, 1824, 1850                  |
| dr-register-policy statement.....              | 5705     | encrypted-password option.....                      | 145, 1824, 1850                  |
| usage guidelines.....                          | 5577     | engine-id statement                                 |                                  |
| drop profiles                                  |          | SNMPv3.....                                         | 1621                             |
| displaying data points.....                    | 7293     | enhanced-transmission-service                       |                                  |
| drop profiles, displaying data points.....     | 7293     | statement.....                                      | 6567, 7138                       |
| drop-profile statement.....                    | 7098     | enrollment-retry statement.....                     | 6305                             |
| drop-profile-map statement.....                | 7098     | enrollment-url statement.....                       | 6305                             |
| drop-profiles statement.....                   | 7099     | environmental information                           |                                  |
| dscp ipv6 statement.....                       | 7071     | chassis, displaying.....                            | 478                              |
| dscp statement.....                            | 7069     | FPC, displaying.....                                | 540                              |
| DVMRP                                          |          | PEMs, displaying.....                               | 566                              |
| defined.....                                   | 5398     | Routing Engines, displaying.....                    | 575                              |
| groups, displaying.....                        | 5878     | erasing user data.....                              | 409                              |
| dynamic (tracing flag)                         |          | error (system logging severity level).....          | 3333                             |
| class of service.....                          | 7058     | error (tracing flag).....                           | 4084                             |
| dynamic IGMP statements                        |          | IS-IS.....                                          | 4321                             |
| promiscuous-mode                               |          | LDP.....                                            | 4877                             |
| interface.....                                 | 5630     | MPLS.....                                           | 5103                             |
| dynamic overload bit, resetting for IS-IS..... | 4334     | OSPF.....                                           | 4611                             |
|                                                |          | RIP.....                                            | 4784                             |
|                                                |          | RSVP.....                                           | 5338                             |
| <b>E</b>                                       |          | ES-IS                                               |                                  |
| east-interface statement.....                  | 2263     | graceful restart.....                               | 2579                             |
| EBGP See BGP                                   |          | ether-options statement.....                        | 2987                             |
| EBGP (external BGP)                            |          | Ethernet interfaces                                 |                                  |
| route-flap damping.....                        | 3926     | mixed VLAN tagging.....                             | 2368, 2370                       |
| EBGP IPv6 peering, BGP.....                    | 3609     | status information, displaying                      |                                  |
| ecmp statement.....                            | 4850     | Gigabit Ethernet.....                               | 3116, 3185                       |
| edge statement.....                            | 2339     | Ethernet MAC MIB.....                               | 1475                             |
| edit command                                   |          | Ethernet rollover cable, connecting the device to a |                                  |
| usage guidelines.....                          | 65       | management device.....                              | 147, 1811, 1819, 1841            |
| egress policy, loopback address.....           | 4831     | ethernet                                            |                                  |
| egress routers                                 |          | statement.....                                      | 224, 230, 2983, 2993, 3021, 3047 |
| static LSPs.....                               | 5049     | oam.....                                            | 2943                             |
| egress statement                               |          | ethernet-ring statement.....                        | 2264                             |
| port mirroring.....                            | 6263     | ethernet-switch-profile statement.....              | 2984                             |
| egress-policy statement.....                   | 4850     | eui-64 statement.....                               | 2988                             |
| usage guidelines.....                          | 4831     | usage guidelines.....                               | 2824                             |
| embedded-rp statement.....                     | 5706     | event (tracing flag)                                |                                  |
| emergency (system logging severity level)..... | 3333     | LDP.....                                            | 4877                             |
| emergency boot device                          |          | RSVP.....                                           | 5338                             |
| creating.....                                  | 157, 175 | Event MIB.....                                      | 1475                             |
| empty paths.....                               | 5073     |                                                     |                                  |
| enable IGMP static group membership.....       | 5418     |                                                     |                                  |
| enable MLD static group membership.....        | 5454     |                                                     |                                  |

|                                                  |                        |  |
|--------------------------------------------------|------------------------|--|
| event recording                                  |                        |  |
| IGMP.....                                        | 5425                   |  |
| MLD.....                                         | 5462                   |  |
| event statement.....                             | 1622                   |  |
| event threshold                                  |                        |  |
| frame error.....                                 | 2946                   |  |
| frame period.....                                | 2947                   |  |
| frame period summary.....                        | 2947                   |  |
| symbol period                                    |                        |  |
| ...227, 2940, 2946, 2951, 2952, 2956, 2957, 3017 |                        |  |
| events                                           |                        |  |
| link event rate.....                             | 2952                   |  |
| events (tracing flag)                            |                        |  |
| spanning-tree protocols.....                     | 2358                   |  |
| events statement.....                            | 316                    |  |
| usage guidelines.....                            | 1891                   |  |
| example                                          |                        |  |
| tracing ospf traffic.....                        | 4570                   |  |
| excess-burst-size statement.....                 | 6106                   |  |
| excess-rate statement.....                       | 7100                   |  |
| exclude statement                                |                        |  |
| IGMP.....                                        | 5620                   |  |
| usage guidelines.....                            | 5418                   |  |
| MLD                                              |                        |  |
| usage guidelines.....                            | 5454                   |  |
| exclude-cmd-attribute statement.....             | 1982                   |  |
| exclude-srlg statement.....                      | 5040                   |  |
| exit command                                     |                        |  |
| usage guidelines.....                            | 65                     |  |
| exp statement.....                               | 7073                   |  |
| explicit-congestion-notification statement.....  | 7101                   |  |
| explicit-null statement                          |                        |  |
| LDP.....                                         | 4851                   |  |
| usage guidelines.....                            | 4834                   |  |
| MPLS.....                                        | 5041                   |  |
| RSVP.....                                        | 5041                   |  |
| explicit-path LSPs                               |                        |  |
| computation, disabling.....                      | 5060                   |  |
| overview.....                                    | 4971                   |  |
| explicit-priority statement.....                 | 1696                   |  |
| usage guidelines                                 |                        |  |
| single-chassis system.....                       | 1536                   |  |
| export route information, displaying.....        | 3455                   |  |
| export statement.....                            | 4851                   |  |
| BGP.....                                         | 3996                   |  |
| usage guidelines.....                            | 3770                   |  |
| forwarding table.....                            | 3296                   |  |
| usage guidelines.....                            | 3828                   |  |
| IS-IS.....                                       | 4271                   |  |
| MSDP.....                                        | 5773                   |  |
| OSPF.....                                        | 4594                   |  |
| PIM.....                                         | 5707, 5708             |  |
| configuring.....                                 | 5574                   |  |
| PIM RP                                           |                        |  |
| usage guidelines.....                            | 5567                   |  |
| RIP.....                                         | 4770                   |  |
| usage guidelines.....                            | 4685                   |  |
| usage guidelines.....                            | 4813                   |  |
| export statement, for routing policies.....      | 3765                   |  |
| export-rib statement.....                        | 3297                   |  |
| extended MAM.....                                | 5032                   |  |
| extension command                                |                        |  |
| usage guidelines.....                            | 65                     |  |
| external-preference statement                    |                        |  |
| IS-IS.....                                       | 4272                   |  |
| OSPF.....                                        | 4595                   |  |
| usage guidelines.....                            | 4468                   |  |
| external-router-id option.....                   | 3793                   |  |
| <b>F</b>                                         |                        |  |
| facilities (system logging)                      |                        |  |
| alternate for remote machine.....                | 1548                   |  |
| default for remote machine.....                  | 1548                   |  |
| for local machine.....                           | 1545, 1546             |  |
| facility-override statement.....                 | 1696                   |  |
| system logging                                   |                        |  |
| usage guidelines.....                            | 1549                   |  |
| failed LSPs                                      |                        |  |
| fast reroute.....                                | 5031, 5042             |  |
| failover-delay statement.....                    | 2690                   |  |
| falling-event-index statement.....               | 1623                   |  |
| falling-threshold statement                      |                        |  |
| health monitor.....                              | 1624, 1932             |  |
| RMON.....                                        | 1625                   |  |
| falling-threshold-interval statement             |                        |  |
| RMON.....                                        | 1626                   |  |
| family multiservice statement                    |                        |  |
| usage guidelines.....                            | 2891                   |  |
| family statement.....                            | 2986, 2989, 3004, 3005 |  |
| BGP.....                                         | 3997                   |  |
| bootstrap.....                                   | 5709                   |  |
| firewall filters.....                            | 6088                   |  |
| IS-IS.....                                       | 4273                   |  |
| link services.....                               | 2965                   |  |
| local RP.....                                    | 5711                   |  |
| PIM interfaces.....                              | 5710                   |  |
| PIM protocol.....                                | 5710                   |  |

|                                        |                 |                                              |                                 |
|----------------------------------------|-----------------|----------------------------------------------|---------------------------------|
| Fast Ethernet interfaces               |                 | files                                        |                                 |
| loopback mode.....                     | 2831            | archiving.....                               | 342, 1370                       |
| fast reroute                           |                 | calculating                                  |                                 |
| configuring.....                       | 5042            | checksum.....                                | 344, 345, 346, 1372, 1373, 1374 |
| detours.....                           | 4967            | comparing.....                               | 347, 1375                       |
| fast-interval statement.....           | 2691            | compressing.....                             | 342, 1370                       |
| usage guidelines.....                  | 2644            | configuration files, compressing.....        | 220                             |
| fast-reroute statement.....            | 5042            | configuration, compressing.....              | 14, 1351                        |
| RSVP.....                              | 5312            | contents, displaying.....                    | 356, 1384                       |
| fate-sharing                           |                 | deleting.....                                | 351, 1379                       |
| signaled LSPs.....                     | 3300            | list of, displaying.....                     | 352, 1380                       |
| fate-sharing statement.....            | 3300            | log file, clearing.....                      | 332, 1368                       |
| fault tolerance                        |                 | renaming.....                                | 354, 1382                       |
| advertising multiple paths to a        |                 | saving configurations to files.....          | 1331, 1343                      |
| destination.....                       | 3651, 3844      | system log messages, archiving.....          | 1543                            |
| fc-map statement.....                  | 6134            | files statement.....                         | 1698                            |
| fcoe-trusted statement.....            | 6136            | archiving of all system log files.....       | 1692                            |
| feature licenses See licenses          |                 | archiving of individual system log file..... | 1694                            |
| FEB                                    |                 | system logging                               |                                 |
| firmware version, displaying.....      | 692             | usage guidelines.....                        | 1543                            |
| FEC filters                            |                 | filter statement.....                        | 2994                            |
| displaying for LDP.....                | 4899            | firewall filters.....                        | 6089                            |
| fec statement                          |                 | interfaces.....                              | 6090                            |
| usage guidelines.....                  | 4852            | vlan.....                                    | 2281, 6091                      |
| FECs.....                              | 4804            | filter-duplicates statement.....             | 1626, 1932                      |
| file archive command.....              | 342, 1370       | filter-interfaces statement.....             | 1627                            |
| file checksum md5 command.....         | 344, 1372       | filter-specific statement.....               | 6107                            |
| file checksum sha-256 command.....     | 346, 1374       | filtering received labels.....               | 4811, 4858                      |
| file checksum sha1 command.....        | 345, 1373       | firewall (system logging facility).....      | 1545, 1547                      |
| file command.....                      | 341             | Firewall MIB.....                            | 1475                            |
| logical systems.....                   | 3951            | firewall statement.....                      | 6092, 6108                      |
| usage guidelines.....                  | 59              | firmware                                     |                                 |
| file compare command.....              | 347, 1375       | chassis, displaying.....                     | 693                             |
| file delete command.....               | 351, 1379       | flags                                        |                                 |
| file list command.....                 | 352, 1380       | login class.....                             | 1799                            |
| file rename command.....               | 354, 1382       | user permissions.....                        | 1799                            |
| file show command.....                 | 356, 1384, 1535 | flap damping.....                            | 3926                            |
| file statement                         |                 | parameters.....                              | 3927                            |
| commit scripts.....                    | 1580            | flash (tracing flag).....                    | 2679, 3384                      |
| op scripts.....                        | 1581            | flexible-vlan-tagging statement.....         | 2368                            |
| security certificate.....              | 6306            | flood reduction.....                         | 4459                            |
| system logging.....                    | 1697            | flooding (tracing flag).....                 | 4611                            |
| usage guidelines.....                  | 1532            | flow routes.....                             | 3301                            |
| file systems                           |                 | flow statement.....                          | 3301, 3997                      |
| checksum values, displaying.....       | 1030            | flow-control statement.....                  | 2996, 7140                      |
| free disk space, displaying.....       | 1177            | flow-control-queue statement.....            | 7142                            |
| filenames, specifying in commands..... | 40              | flow-map statement.....                      | 3302                            |
|                                        |                 | font conventions.....                        | cxli                            |
|                                        |                 | force-version statement.....                 | 2340                            |



- format statement.....296
- forward-delay statement.....2341
- forward-only statement.....6295
- forwarding equivalence class See FEC
- forwarding equivalence classes See FECs
- forwarding options
  - DHCP relay agent.....6360, 6362, 6370, 6375
  - DHCPv6 relay agent.....6365, 6368, 6378, 6384
- forwarding software process.....1794
- forwarding table
  - aggregate routes.....3289
  - CoS information, displaying
    - classifier information.....7253
    - code point value to loss priority.....7251
    - code point value to queue number.....7251
    - configuration.....7247, 7296
    - loss priority per classifier.....7251
    - RED drop profiles.....7300
    - rewrite rules.....7255
    - scheduler map.....7302
  - generated routes.....3289
  - multicast information, displaying.....5863
  - policy, routing.....3296
  - route entries, displaying.....3477, 5227
  - static routes.....3270, 3316, 3358
- forwarding-cache statement
  - flow maps.....3303
  - multicast.....3304
- forwarding-class statement
  - class of service.....7075, 7076, 7104
- forwarding-class-set statement
  - class of service.....7105
- forwarding-classes statement.....7077
- forwarding-options statement.....2282, 6290
- forwarding-table statement.....3308, 5043
  - usage guidelines.....3828
- FPC
  - environmental information, displaying.....540
  - firmware version, displaying.....692
  - installed, displaying list.....744
  - operation of, controlling.....366
  - status, displaying.....704
  - switch fabric status, displaying.....580, 585
- fpc statement.....2997
  - DDoS protection.....6154
- fragmentation
  - avoiding.....3912
- frames
  - multicast snooping.....5395
- free disk space, displaying.....1177
- FreeBSD upgrade
  - downgrading Junos OS.....102
  - Junos OS.....89
  - Junos OS disk volumes.....94
  - Junos OS package names.....92
  - Junos OS snapshots.....95
  - rebooting Junos OS.....262
  - upgrading Junos OS.....96
- freeing up storage space.....399
- from statement.....3339, 5076, 6093
  - MPLS.....5043
- ftp (system logging facility).....1545, 1547
  - option to facility-override statement.....1548
- full mesh requirement
  - fulfilling with confederations.....3890
  - fulfilling with route reflectors.....3873
- full names, in user accounts.....1822, 1845
- full statement.....3289
- full-name statement.....1933
  - usage guidelines.....125, 1821, 1845, 1853
- full-neighbors statement
  - OSPF
    - usage guidelines.....4484
- fwdd process.....1794
- G**
  - ge-0/0/0, management interface.....2796
    - See also Gigabit Ethernet ports
  - general (tracing flag).....2679, 3384
    - BMP.....4084
  - generate statement.....3309
  - generated routes.....3309
    - preferences.....3786
  - Gigabit Ethernet interfaces
    - diagnostics information, displaying.....3102
    - loopback mode.....2831
    - status information, displaying.....3116, 3185
  - global statement
    - DDoS protection.....6155
  - global tracing operations.....1409
  - global-mac-table-aging-time statement.....2315
  - GMPLS
    - link-management information, displaying
      - all.....5127
      - peers.....5131
      - routing process.....5133
      - statistics.....5136, 5156
      - traffic-engineered links.....5138, 5158

|                                          |                              |                                           |            |
|------------------------------------------|------------------------------|-------------------------------------------|------------|
| gpip statement.....                      | 5044                         | OSPF.....                                 | 2664, 4596 |
| graceful restart.....                    | 4282                         | usage guidelines.....                     | 4495       |
| commands, operational mode.....          | 2707                         | BGP.....                                  | 2581       |
| concepts.....                            | 2579                         | global.....                               | 2580       |
| configuration.....                       | 2627                         | grace period.....                         | 4497       |
| configuration procedure                  |                              | helper mode.....                          | 4501, 4504 |
| routing protocols.....                   | 2580                         | IS-IS.....                                | 2582       |
| disabling.....                           | 4497                         | OSPF.....                                 | 2583       |
| disabling strict LSA checking.....       | 4507                         | OSPFv3.....                               | 2583       |
| enabling.....                            | 4497                         | PIM.....                                  | 2585       |
| grace period interval.....               | 4497                         | RIP.....                                  | 2585       |
| LDP.....                                 | 4853                         | RIPng.....                                | 2585       |
| OSPFv2 helper mode                       |                              | strict LSA checking.....                  | 4507       |
| disabling.....                           | 4501                         | graceful-switchover statement.....        | 2673       |
| re-enabling.....                         | 4501                         | usage guidelines.....                     | 2627       |
| OSPFv3 helper mode                       |                              | graft (tracing flag)                      |            |
| disabling.....                           | 4504                         | PIM.....                                  | 5755       |
| re-enabling.....                         | 4504                         | gratuitous-arp-reply.....                 | 2998       |
| overview.....                            | 2617, 4495                   | GRE interfaces                            |            |
| MPLS protocols.....                      | 4969                         | status information, displaying.....       | 3131       |
| PIM                                      |                              | GRES (graceful Routing Engine switchover) |            |
| sparse-dense mode.....                   | 5527                         | overview.....                             | 2617       |
| protocol support.....                    | 2579                         | GRES (graceful routing engine switchover) |            |
| trace options.....                       | 2586                         | configuration.....                        | 2627       |
| verifying operation of.....              | 2707                         | group joins                               |            |
| graceful Routing Engine switchover       |                              | limiting.....                             | 5426, 5464 |
| concepts.....                            | 2617                         | group membership                          |            |
| DPC support.....                         | 2626                         | SSM maps.....                             | 5541       |
| enabling.....                            | 2627                         | group statement                           |            |
| feature support.....                     | 2624                         | BGP.....                                  | 4003       |
| overview.....                            | 2617                         | IGMP.....                                 | 5621       |
| PIC support.....                         | 2626                         | usage guidelines.....                     | 5418       |
| platform support.....                    | 2623                         | IGMP snooping.....                        | 5645, 5667 |
| subscriber access support.....           | 2626                         | MLD                                       |            |
| system architecture.....                 | 2619                         | usage guidelines.....                     | 5454       |
| system requirements.....                 | 2623                         | MSDP.....                                 | 5774       |
| verifying status of.....                 | 2629                         | PIM RPF selection.....                    | 5712       |
| graceful switchover, displaying.....     | 2738                         | RIP.....                                  | 4771       |
| graceful-deletion-timeout statement..... | 5312                         | SNMPv3 (for access privileges).....       | 1627       |
| graceful-restart (tracing flag).....     | 5103                         | SNMPv3 (for configuring).....             | 1628       |
| IS-IS.....                               | 4321                         | group-count statement                     |            |
| OSPF.....                                | 4611                         | IGMP.....                                 | 5622       |
| graceful-restart                         |                              | usage guidelines.....                     | 5418       |
| statement.....                           | 2660, 2662, 3311, 4001, 5313 | MLD                                       |            |
| BGP.....                                 | 2662, 4001                   | usage guidelines.....                     | 5454       |
| IS-IS.....                               | 4274                         |                                           |            |
| LDP.....                                 | 4853                         |                                           |            |
| usage guidelines.....                    | 4828                         |                                           |            |

|                                               |            |  |
|-----------------------------------------------|------------|--|
| group-increment statement                     |            |  |
| IGMP.....                                     | 5622       |  |
| usage guidelines.....                         | 5418       |  |
| MLD                                           |            |  |
| usage guidelines.....                         | 5454       |  |
| group-limit statement                         |            |  |
| configuring.....                              | 5426       |  |
| IGMP interface.....                           | 5623       |  |
| IGMP snooping.....                            | 5646, 5668 |  |
| MLD                                           |            |  |
| usage guidelines.....                         | 5464       |  |
| group-policy statement                        |            |  |
| IGMP.....                                     | 5623       |  |
| usage guidelines.....                         | 5414       |  |
| MLD                                           |            |  |
| usage guidelines.....                         | 5452       |  |
| group-ranges statement.....                   | 5713       |  |
| usage guidelines.....                         | 5547       |  |
| groups                                        |            |  |
| BGP                                           |            |  |
| general information, displaying.....          | 4100       |  |
| DVMRP, displaying.....                        | 5878       |  |
| IGMP membership, displaying.....              | 5795       |  |
| OSPF areas.....                               | 4408       |  |
| PIM                                           |            |  |
| general information, displaying.....          | 5886       |  |
| usage information, displaying.....            | 5878       |  |
| RIP routers.....                              | 4685       |  |
| SSM.....                                      | 3375, 5763 |  |
| groups statement.....                         | 6313       |  |
| guaranteed-rate statement.....                | 7106       |  |
| guard-interval statement.....                 | 2265       |  |
| <b>H</b>                                      |            |  |
| halts                                         |            |  |
| pending                                       |            |  |
| clearing.....                                 | 337        |  |
| displaying.....                               | 1126       |  |
| requesting.....                               | 384        |  |
| hardware, installed, displaying.....          | 743        |  |
| hardware-database (tracing flag)              |            |  |
| class of service.....                         | 7058       |  |
| hash-key statement.....                       | 3054, 3071 |  |
| health monitor alarms, displaying.....        | 1749       |  |
| health-monitor statement.....                 | 1629, 1933 |  |
| hello (tracing flag)                          |            |  |
| IS-IS.....                                    | 4321       |  |
| PIM.....                                      | 5755       |  |
| hello interval                                |            |  |
| LDP.....                                      | 4825, 4854 |  |
| RSVP.....                                     | 5315       |  |
| hello messages.....                           | 4807       |  |
| hello packets                                 |            |  |
| IS-IS.....                                    | 4160       |  |
| OSPF.....                                     | 4391       |  |
| hello-acknowledgements statement.....         | 5314       |  |
| hello-authentication-key statement.....       | 4275       |  |
| hello-authentication-key-chain statement..... | 4276       |  |
| hello-authentication-type statement.....      | 4277       |  |
| hello-interval statement                      |            |  |
| IS-IS.....                                    | 4278       |  |
| LDP.....                                      | 4854       |  |
| usage guidelines.....                         | 4825       |  |
| OSPF                                          |            |  |
| usage guidelines.....                         | 4476       |  |
| PIM.....                                      | 5714       |  |
| usage guidelines.....                         | 5494       |  |
| RSVP.....                                     | 5315       |  |
| hello-padding statement.....                  | 4280       |  |
| hello-time statement.....                     | 2342       |  |
| help command                                  |            |  |
| usage guidelines.....                         | 65         |  |
| helper-disable statement.....                 | 2666, 5316 |  |
| IS-IS.....                                    | 4274       |  |
| LDP.....                                      | 4855       |  |
| usage guidelines.....                         | 4828       |  |
| usage guidelines                              |            |  |
| IS-IS.....                                    | 2582       |  |
| OSPF .....                                    | 2583       |  |
| OSPFv3.....                                   | 2583       |  |
| high availability.....                        | 2617, 2627 |  |
| history statement                             |            |  |
| RMON.....                                     | 1630       |  |
| history, CLI commands                         |            |  |
| displaying.....                               | 435        |  |
| hold time                                     |            |  |
| LDP.....                                      | 4826, 4856 |  |
| signaled LSPs.....                            | 5027       |  |
| hold-down statement.....                      | 4006       |  |
| usage guidelines.....                         | 3647       |  |
| hold-down-interval                            |            |  |
| BGP.....                                      | 4008       |  |
| hold-interval statement                       |            |  |
| Ethernet ring protection switching.....       | 2266       |  |
| hold-multiplier statement.....                | 1934       |  |

|                                      |            |                                              |            |
|--------------------------------------|------------|----------------------------------------------|------------|
| hold-time statement.....             | 2692       | idle time, displaying.....                   | 167        |
| BGP.....                             | 4010       | idle timeout values                          |            |
| IS-IS.....                           | 4282       | login classes.....                           | 22         |
| LDP.....                             | 4856       | idle-timeout statement.....                  | 278, 1935  |
| usage guidelines.....                | 4826       | usage guidelines.....                        | 22         |
| Physical Interface.....              | 2999       | ieee-802.1 statement.....                    | 7143, 7144 |
| PIM.....                             | 5715       | ieee-802.1 statement.....                    | 7079       |
| holddown (tracing flag).....         | 4784       | if-exceeding statement.....                  | 6109       |
| holddown statement                   |            | ifd process.....                             | 1794       |
| IS-IS.....                           | 4319       | IGMP.....                                    | 5799       |
| RIP.....                             | 4773       | configuration statements.....                | 5408       |
| usage guidelines.....                | 4749       | configuring.....                             | 5408       |
| holddown-interval statement          |            | disabling.....                               | 5429       |
| BFD                                  |            | enabling.....                                | 5409, 5624 |
| static routes.....                   | 3281       | event recording.....                         | 5425       |
| hop count, maximizing.....           | 4680       | group membership                             |            |
| See also RIP                         |            | SSM maps for different groups to different   |            |
| hop-limit statement.....             | 5045, 5317 | sources.....                                 | 5541       |
| host reachability                    |            | group membership, displaying.....            | 5795       |
| verifying for RIP network hosts..... | 4691       | host-query message interval.....             | 5411, 5631 |
| Host Resources MIB.....              | 1476       | interface group limit.....                   | 5623       |
| host statement.....                  | 1699       | interfaces, displaying.....                  | 5801       |
| system logging                       |            | last-member query interval.....              | 5412, 5632 |
| usage guidelines for single-chassis  |            | overview.....                                | 5406       |
| system.....                          | 1533       | PIM-to-IGMP message translation information, |            |
| host-name statement.....             | 224        | displaying.....                              | 5859       |
| usage guidelines.....                | 16         | query response interval.....                 | 5416, 5633 |
| host-only-interface statement.....   | 5647, 5669 | robustness variable.....                     | 5417, 5634 |
| hostnames                            |            | static group membership.....                 | 5418       |
| IS-IS, displaying.....               | 4359       | statistics, displaying.....                  | 5805       |
| hostnames, DNS, displaying.....      | 1002       | tracing operations.....                      | 5427       |
| hosts, reachability                  |            | version.....                                 | 5410, 5641 |
| general connections.....             | 360, 1724  | IGMP snooping                                |            |
| MPLS BGP LSPs.....                   | 5112       | configure the switch to be an IGMP           |            |
| MPLS Layer 2 circuits.....           | 5114       | querier.....                                 | 5651, 5673 |
| MPLS Layer 3 VPN connections.....    | 5117       | enabling.....                                | 5648, 5670 |
| MPLS LDP LSPs.....                   | 4887       | group limit.....                             | 5646, 5668 |
| MPLS LSP endpoints.....              | 5120       | group statement.....                         | 5645, 5667 |
| MPLS RSVP LSPs.....                  | 5347       | host-only interface.....                     | 5647, 5669 |
|                                      |            | host-query message interval.....             | 5653, 5676 |
|                                      |            | last-member query interval.....              | 5654, 5677 |
|                                      |            | query response interval.....                 | 5655, 5678 |
|                                      |            | source address.....                          | 5657, 5680 |
|                                      |            | static statement.....                        | 5645, 5667 |
|                                      |            | igmp statement.....                          | 5624       |
|                                      |            | usage guidelines.....                        | 5409       |
|                                      |            | IGMP statements                              |            |
|                                      |            | promiscuous-mode                             |            |
|                                      |            | interface.....                               | 5630       |

## I

IBGP See BGP

IBGP (internal BGP)

    full mesh (configuration editor).....

    overview.....

icmpv4-rate-limit statement.....

    usage guidelines.....

identifiers

    BGP See BGP, identifier

- 
- igmp-querier statement.....5647, 5658, 5669, 5681
  - igmp-snooping statement.....5648, 5670
  - IGMPv3.....5408
    - interoperability with older versions.....5408
  - ignore-attached-bit statement.....4283
  - ignore-lsp-metrics statement.....4857
    - usage guidelines.....4835
  - IGP plus MED, BGP option.....3667
  - igp shortcuts
    - overview.....4511
  - IGP synchronization, LDP.....4838
  - igp-synchronization statement.....4857
    - usage guidelines.....4838
  - immediate-leave statement
    - IGMP.....5626
      - usage guidelines.....5413
    - IGMP snooping.....5649, 5671
    - MLD
      - usage guidelines.....5451
  - import statement.....7081
    - BGP.....4012
      - usage guidelines.....3770
    - bootstrap.....5716
      - usage guidelines.....5567
    - LDP.....4858
      - usage guidelines.....4811
    - MSDP.....5775
    - OSPF.....4598
    - PIM.....5717
      - usage guidelines.....5575
    - RIP.....4774
      - usage guidelines.....4719
    - route resolution.....3312
  - import statement, for routing policies.....3765
  - import-policy statement.....3313
  - import-rib statement.....3314
  - include-mp-next-hop statement.....4014
  - incoming metric (RIP)
    - description.....4725
  - indirect next hop.....3315
  - indirect-next-hop statement.....3315
  - inet statement.....238
    - usage guidelines.....23
  - infinity statement.....5718
    - usage guidelines.....5590
  - info (system logging severity level).....3333
  - informs SNMP See SNMP informs
  - ingress routers
    - address configuration.....5043
    - path connection retry information.....5085
  - ingress statement.....6267
    - static LSP.....5047
  - init (tracing flag)
    - class of service.....7059
  - initialization (tracing flag).....4878
  - initiation-message statement.....4015
    - usage guidelines.....3648
  - input statement.....6268
    - forwarding table.....6094
  - insecure statement.....15, 221, 1411
  - insert command
    - usage guidelines.....65
  - install statement.....3316
    - MPLS.....5048
  - installation
    - licenses (CLI).....117
  - installing software.....1276, 1277
  - instance-export statement.....3317
  - instance-import statement.....3317
  - integrated routing and bridging interfaces
    - integrated routing and bridging interfaces, displaying.....3140
  - inter-area-prefix-export statement
    - usage guidelines.....4544
  - inter-area-prefix-import statement
    - usage guidelines.....4554
  - interactive-commands (system logging facility).....1545, 1547
  - interface
    - broadcast
      - configuring.....4432
      - monitoring.....167, 2844
    - point-to-multipoint
      - configuring.....4438
    - point-to-point
      - configuring.....4432
  - interface (DCBX) statement.....6569, 7146
  - interface (from operator, LDP).....4811
  - interface (Storm Control) statement.....6120
  - interface addresses
    - logical interfaces.....2824
  - interface lists.....5397
  - Interface MIB.....1476, 1483
  - interface range
    - interfaces.....2795
  - interface software process.....1794

|                                                              |                                                |
|--------------------------------------------------------------|------------------------------------------------|
| interface statement                                          |                                                |
| Ethernet switching options.....                              | 3061                                           |
| IGMP.....                                                    | 5627                                           |
| usage guidelines.....                                        | 5409                                           |
| IGMP snooping.....                                           | 5650, 5651, 5672, 5673                         |
| IS-IS.....                                                   | 4284                                           |
| LDP.....                                                     | 4859                                           |
| usage guidelines.....                                        | 4810                                           |
| lfm.....                                                     | 2950, 2953, 2955                               |
| LLDP.....                                                    | 1936                                           |
| MLD                                                          |                                                |
| usage guidelines.....                                        | 5448                                           |
| multicast.....                                               | 3321                                           |
| multicast via static routes.....                             | 3319                                           |
| OSPF.....                                                    | 4599                                           |
| usage guidelines.....                                        | 4432, 4435                                     |
| PIM.....                                                     | 5719                                           |
| usage guidelines.....                                        | 5526                                           |
| port                                                         |                                                |
| mirroring.....                                               | 6264, 6265, 6266, 6269, 6270, 6271, 6273, 6274 |
| rate limiting.....                                           | 6123                                           |
| RMON history.....                                            | 1632                                           |
| RSVP.....                                                    | 5318                                           |
| SNMP.....                                                    | 1631                                           |
| usage guidelines.....                                        | 1506                                           |
| spanning tree.....                                           | 2344                                           |
| static LSPs.....                                             | 5049                                           |
| STP.....                                                     | 2343                                           |
| VRRP.....                                                    | 2693                                           |
| usage guidelines.....                                        | 2646                                           |
| interface statistics, real-time, displaying.....             | 3090                                           |
| interface-mac-limit statement.....                           | 2288                                           |
| interface-mode statement.....                                | 2290, 3006                                     |
| interface-range statement.....                               | 3008                                           |
| interface-routes statement.....                              | 3322                                           |
| interface-specific statement.....                            | 6094                                           |
| interface-type statement                                     |                                                |
| usage guidelines.....                                        | 4435                                           |
| interfaces                                                   |                                                |
| broadcast.....                                               | 4431                                           |
| demand circuit.....                                          | 4431                                           |
| descriptive text.....                                        | 2982                                           |
| mixed VLAN tagging.....                                      | 2368, 2370                                     |
| monitor interface traffic command.....                       | 1416                                           |
| NBMA.....                                                    | 4431                                           |
| overview.....                                                | 4431                                           |
| passive.....                                                 | 4431                                           |
| passive traffic engineering mode.....                        | 4431                                           |
| peer.....                                                    | 4431                                           |
| point-to-multipoint.....                                     | 4431                                           |
| point-to-point.....                                          | 4431                                           |
| system logging.....                                          | 1555                                           |
| tracing operations.....                                      | 1410                                           |
| interfaces descriptive text.....                             | 3294                                           |
| interfaces limiting SNMP access.....                         | 1506                                           |
| interfaces statement.....                                    | 1593, 3010                                     |
| class of service.....                                        | 7082, 7147                                     |
| intermediate routers                                         |                                                |
| configuring for static LSPs.....                             | 5049                                           |
| Intermediate System-to-Intermediate System See IS-IS         |                                                |
| internal routers                                             |                                                |
| description.....                                             | 4397                                           |
| overview.....                                                | 4398                                           |
| Internet Group Management Protocol See IGMP                  |                                                |
| internet-options statement.....                              | 226                                            |
| usage guidelines.....                                        | 86                                             |
| interpolate statement.....                                   | 7107                                           |
| interval statement                                           |                                                |
| health monitor.....                                          | 1633, 1937                                     |
| RMON alarm.....                                              | 1634                                           |
| RMON history.....                                            | 1634                                           |
| invalid routes, displaying.....                              | 3508                                           |
| IP addresses.....                                            | 23                                             |
| router mapping.....                                          | 34                                             |
| router names, mapping.....                                   | 23                                             |
| IP multicast                                                 |                                                |
| announced sessions, displaying.....                          | 5875                                           |
| bandwidth admission                                          |                                                |
| clearing.....                                                | 5827                                           |
| flow map information, displaying.....                        | 5850                                           |
| forwarding table, displaying.....                            | 5863                                           |
| interface information, displaying.....                       | 5852                                           |
| network information, displaying.....                         | 5854                                           |
| next-hop table, displaying.....                              | 5856                                           |
| PIM-to-IGMP message translation information, displaying..... | 5859                                           |
| PIM-to-MLD message translation information, displaying.....  | 5861                                           |
| RPF calculations, displaying.....                            | 5869                                           |
| scope, clearing.....                                         | 5829                                           |
| scoped information, displaying.....                          | 5873                                           |
| sessions, clearing.....                                      | 5830                                           |

- statistics
  - clearing.....5831
- tracing routes
  - from the receiver to the source.....5839
  - from the source to the gateway
    - router.....5847
  - from the source to the receiver.....5842
  - listen for responses.....5845
- IP packets
  - router source addresses.....20, 222
- IP sockets, displaying active.....1060
- IPv4 Protocol family
  - interface addresses.....2824
- ipv4-multicast statement.....4286
  - usage guidelines.....4203
- ipv4-multicast-metric statement.....4287
  - usage guidelines.....4203
- IPv6
  - EBGP link-local peering.....3609
  - logical systems.....3609
  - neighbor cache information
    - clearing.....3391
    - displaying.....3401
  - router advertisements
    - displaying.....3403
  - tunneling over MPLS.....5002
  - ipv6-multicast statement.....4288
  - ipv6-multicast-metric statement.....4288
  - ipv6-tunneling statement.....5050
  - ipv6-unicast statement.....4289
    - usage guidelines.....4193
  - ipv6-unicast-metric statement.....4290
- IS-IS
  - addresses.....4158
  - adjacency database entries, clearing
    - all.....4330
  - authentication.....4183, 4255, 4263, 4300
    - CSNP.....4300
    - hello.....4301
    - hitless keychain.....4256
    - PSNP.....4306
  - authentication keychain.....4264, 4276
  - authentication, displaying.....4342, 5124
  - backup coverage
    - displaying.....4344
  - backup MPLS LSPs.....4346
  - backup paths
    - SPF calculations.....4348
  - BFD.....4266
  - complete sequence number
    - PDU.....4160, 4269, 4321
  - configuring on logical systems.....4232
  - designated router.....4316
  - disabling.....4270
    - IPv4 multicast topology.....4203
    - IPv4 unicast topology.....4203
    - IPv6 multicast topology.....4203
    - IPv6 unicast topology.....4193
  - dynamic overload bit, resetting.....4334
  - enabling.....4291
  - errored packets.....4321
  - graceful restart.....2579, 4274
  - hello
    - interval.....4278
    - packet authentication.....4277
    - packet authentication key.....4275
    - PDU.....4160, 4321
  - hold time.....4282
  - hold-down timer
    - disabling.....4299
  - hostname database, displaying.....4359
  - interfaces.....4284
  - interfaces, displaying.....4361
  - IP fast reroute.....4165, 4218, 4221
  - IPv4 unicast topology.....4307
  - IPv6 unicast topology.....4289, 4306
  - level properties
    - global.....4292
  - link protection.....4167, 4219
  - link-protection statement.....4293
  - link-state database entries
    - clearing.....4332
    - displaying.....4352
  - link-state PDUs.....4160, 4321
    - interval.....4294
    - lifetime.....4295
    - tracing.....4321
  - loop-free alternate routes.....4165, 4218, 4221
  - loose authentication.....4184, 4293
  - mesh groups.....4297
  - metrics.....4317
    - IPv6.....4290
    - multicast.....4287, 4288
    - normal.....4298
    - wide.....4327
  - multicast topologies.....4202, 4203, 4286, 4288
    - IPv4.....4302
    - IPv6.....4304



|                                                    |                  |                                          |               |
|----------------------------------------------------|------------------|------------------------------------------|---------------|
| multilevel.....                                    | 4175             | Juniper-Allow-Configuration attribute    |               |
| neighbors, displaying.....                         | 4338             | (RADIUS).....                            | 1889          |
| network PDUs.....                                  | 4158             | Juniper-Authentication-Type.....         | 1890          |
| no-eligible-backup statement.....                  | 4301             | Juniper-Configuration-Change attribute   |               |
| node link protection.....                          | 4168, 4220, 4308 | (RADIUS).....                            | 1890          |
| nonstop active routing.....                        | 2601             | Juniper-Deny-Commands attribute          |               |
| NSAP.....                                          | 4158             | (RADIUS).....                            | 1889          |
| overloaded, marking router as.....                 | 4309             | Juniper-Deny-Configuration attribute     |               |
| padding.....                                       | 4280             | (RADIUS).....                            | 1889          |
| partial sequence number PDUs.....                  | 4160, 4321       | Juniper-Interactive-Command attribute    |               |
| PDUs.....                                          | 4160             | (RADIUS).....                            | 1889          |
| point-to-point interface.....                      | 4313             | Juniper-Local-User-Name attribute        |               |
| policy, routing.....                               | 4271             | (RADIUS).....                            | 1889          |
| preferences.....                                   | 3786, 4272, 4314 | Juniper-Session-Port .....               | 1890          |
| prefix limit.....                                  | 4315             | Juniper-User-Permissions attribute       |               |
| redistributing OSPF routes into.....               | 4185             | (RADIUS).....                            | 1890          |
| routes, displaying.....                            | 4369             | juniper.conf file, compressing.....      | 14, 220, 1351 |
| RSVP LSP backup paths.....                         | 4168, 4221       | JUNOS Internet software                  |               |
| SPF calculations, displaying.....                  | 4373             | version, displaying.....                 | 166           |
| SPF delay calculations.....                        | 4321             | Junos OS                                 |               |
| topology.....                                      | 4320             | alarms, displaying.....                  | 1027          |
| tracing operations.....                            | 4321             | autoinstallation status, displaying..... | 324           |
| traffic engineering support.....                   | 4270, 4324       | boot messages, displaying.....           | 1038          |
| traffic statistics                                 |                  | buffers, displaying.....                 | 1045          |
| clearing.....                                      | 4336             | checksum values, displaying.....         | 1030          |
| displaying.....                                    | 4378             | core-dumps, displaying.....              | 1079          |
| IS-IS (intermediate system to intermediate system) |                  | directory usage, displaying.....         | 1093          |
| protocol                                           |                  | disk space, displaying.....              | 1177          |
| configuration.....                                 | 2627             | downgrading from upgraded FreeBSD.....   | 102           |
| IS-IS (Intermediate System-to-Intermediate System) |                  | FreeBSD upgraded.....                    | 89            |
| BFD authentication.....                            | 4250             | halt, requesting a.....                  | 384           |
| checksum.....                                      | 4239             | loaded extensions, displaying.....       | 1132          |
| configuring.....                                   | 4169             | pending reboots                          |               |
| for CLNS route exchange.....                       | 4237             | clearing.....                            | 337           |
| isis statement.....                                | 4291             | displaying.....                          | 1126          |
| ISO                                                |                  | processes, displaying.....               | 1097          |
| addresses.....                                     | 4158             | rebooting.....                           | 6336          |
| system identifier.....                             | 4158             | rebooting with FreeBSD upgrade.....      | 262           |
| ISO Protocol family.....                           | 2824             | reinstalling                             |               |
|                                                    |                  | using removable media.....               | 159           |
| <b>J</b>                                           |                  | SRC client, displaying.....              | 1131          |
| join (tracing flag).....                           | 5755             | system statistics, displaying.....       | 1140          |
| join states, clearing PIM.....                     | 5832             | upgrading FreeBSD.....                   | 96            |
| join-load-balance statement.....                   | 5720             | uptime, displaying.....                  | 1185          |
| usage guidelines.....                              | 5513             | version, displaying                      |               |
| join-prune-timeout statement.....                  | 5721             | general.....                             | 1253          |
| Juniper-Allow-Commands attribute                   |                  | virtual memory, displaying.....          | 1195          |
| (RADIUS).....                                      | 1889             | Junos OS disk volumes                    |               |
|                                                    |                  | FreeBSD upgraded.....                    | 94            |



- Junos OS package names
  - FreeBSD upgraded.....92
- Junos OS snapshots
  - FreeBSD upgraded.....95
- JUNOS Software
  - bundles, deleting.....1286
  - packages, deleting.....1286
  - powering off.....391
  - rolling back.....1292
  - upgrade
    - performing.....1276, 1290
    - validating candidate.....1297
- JUNOS software
  - overview.....1793
  - Packet Forwarding Engine.....1793
  - processes.....1794
  - Routing Engine.....1793
- Junos-FIPS
  - password requirements.....146, 1823, 1847
- K**
  - keep multiplier, RSVP.....5319
  - keep statement.....4016
    - usage guidelines.....3772
  - keep-multiplier statement.....5319
  - keepalive (tracing flag)
    - BGP.....4081
    - MSDP.....5785
  - keepalive messages.....3598
  - keepalive-interval statement.....4860
    - usage guidelines.....4827
  - keepalive-timeout statement.....4861
    - usage guidelines.....4828
  - keepalives
    - interval.....4827, 4860
    - timeout.....4828, 4861
  - kernel
    - downgrading from upgraded FreeBSD.....102
    - FreeBSD upgrade.....89
    - upgrading FreeBSD.....96
  - kernel (system logging facility).....1545, 1547
    - option to facility-override statement.....1548
  - kernel (tracing flag).....2679, 3384
  - kernel memory usage, displaying.....1195
  - kernel replication state, displaying.....2738
  - key statement.....6307
  - key-chain statement.....6308
    - BFD authentication.....5721
  - key-chain-name
    - BGP.....4018
  - keyboard sequences
    - used with monitor interface command.....3090
    - used with monitor interface traffic
      - command.....3091
      - used with monitor mpls command.....5109
  - keychain
    - BGP.....3898
    - IS-IS.....4255
    - overview.....3897
- L**
  - l2-learning statement.....2525
  - l2-querier statement
    - IGMP snooping.....5651, 5673
  - l2-smart-policy statement.....4861
  - L2CKT, nonstop active routing.....2606
  - l3-interface statement.....2295
  - label (tracing flag).....4878
  - Label Distribution Protocol *See* LDP
  - label filtering.....4811, 4858
  - label-switched-path statement
    - MPLS.....5051
    - MPLS with RSVP.....5051
  - label-withdrawal-delay statement.....4862
    - usage guidelines.....4838
  - labels
    - operations.....4806
  - laptop *See* management device
  - Layer 2 BPDU loop prevention
    - clearing errors.....2418
  - Layer 2 bridging
    - MAC address table, displaying.....2558
  - Layer 2 circuits
    - nonstop active routing.....2601
    - reachability, testing.....5114
  - Layer 2 VPNs
    - nonstop active routing.....2601
  - Layer 3 VPNs
    - nonstop active routing.....2601
    - reachability, testing.....5117
  - ldap-url statement.....6309
  - LDP
    - authentication algorithm.....3976, 4843
    - authentication keychain.....4846
    - carrier-of-carriers VPNs.....4834
    - configuring.....4859, 4863
    - database entries, displaying.....4890

|                                   |                        |                                             |                              |
|-----------------------------------|------------------------|---------------------------------------------|------------------------------|
| disabling.....                    | 4810, 4848             | tracing LSPs.....                           | 4924                         |
| downstream on demand.....         | 4820                   | tracing operation of.....                   | 4817, 4877                   |
| egress policy.....                | 4831                   | traffic statistics, displaying.....         | 4920                         |
| enabling.....                     | 4810                   | tunneling through RSVP                      |                              |
| example configuration             |                        | LSPs.....                                   | 4805, 4834, 4866, 5053, 5264 |
| received label filtering.....     | 4812                   | ultimate-hop popping.....                   | 4834, 4851                   |
| tracing.....                      | 4819                   | LDP LSPs                                    |                              |
| Explicit Null label.....          | 4834                   | ping interval.....                          | 4887                         |
| FEC filters, displaying.....      | 4899                   | ldp statement.....                          | 4863                         |
| graceful restart.....             | 2579, 4809, 4828, 4853 | usage guidelines.....                       | 4810                         |
| hello interval.....               | 4825, 4854             | LDP-based VPLS, nonstop active routing..... | 2606                         |
| hello messages.....               | 4807                   | ldp-synchronization statement.....          | 4866                         |
| hold time.....                    | 4826, 4856             | usage guidelines.....                       | 4837                         |
| IGP synchronization.....          | 4838                   | ldp-tunneling statement.....                | 4866, 5053                   |
| Implicit Null label.....          | 4834                   | usage guidelines.....                       | 4834                         |
| interfaces, displaying.....       | 4900                   | lease-time statement.....                   | 252                          |
| Junos implementation.....         | 4804                   | leave (tracing flag)                        |                              |
| keepalive                         |                        | IGMP.....                                   | 5639                         |
| interval.....                     | 4827, 4860             | level statement                             |                              |
| timeout.....                      | 4828, 4861             | IS-IS                                       |                              |
| label operations.....             | 4806                   | protocol.....                               | 4292                         |
| LSPs, displaying.....             | 4904                   | usage guidelines.....                       | 4175                         |
| message types.....                | 4807                   | lfp                                         |                              |
| metrics.....                      | 4833                   | action profile.....                         | 2941                         |
| minimum configuration.....        | 4810                   | event threshold.....                        | 2945                         |
| multiple instances.....           | 4834                   | link discovery.....                         | 2951                         |
| neighbors                         |                        | negotiation options.....                    | 2942, 2954                   |
| clearing connections.....         | 4884                   | no allow link events.....                   | 2954                         |
| displaying.....                   | 4902                   | pdu threshold.....                          | 2955                         |
| nonstop active routing.....       | 2601                   | remote loopback.....                        | 2956                         |
| OAM periodic traceroute.....      | 4832                   | license infringement                        |                              |
| operations.....                   | 4804                   | verifying license usage.....                | 121                          |
| overview.....                     | 4804                   | verifying licenses installed.....           | 120                          |
| policy filters.....               | 4858                   | license keys                                |                              |
| received label filtering.....     | 4811, 4858             | components.....                             | 112                          |
| route preferences.....            | 4828, 4871             | licenses.....                               | 112                          |
| routes, displaying.....           | 4906                   | adding.....                                 | 438                          |
| session protection                |                        | adding (CLI).....                           | 117                          |
| configuration.....                | 4836                   | deleting.....                               | 439                          |
| overview.....                     | 4808                   | deleting (CLI).....                         | 118                          |
| sessions                          |                        | displaying.....                             | 441                          |
| clearing.....                     | 4885                   | displaying (CLI).....                       | 120                          |
| displaying.....                   | 4910                   | displaying usage.....                       | 121                          |
| statistics                        |                        | overview.....                               | 107                          |
| clearing.....                     | 4886                   | saving.....                                 | 440                          |
| displaying.....                   | 4916                   | saving (CLI).....                           | 119                          |
| synchronization with the IGP..... | 4837                   | verifying.....                              | 120                          |
| targeted hello messages.....      | 4807                   | link hello messages, LDP.....               | 4854                         |
| timer.....                        | 4825, 4854             | Link protection.....                        | 5284                         |

- 
- link-mode statement.....3018
  - link-protection statement.....3051
    - IS-IS.....4167, 4219, 4293
    - RSVP.....5320
    - static LSPs.....5054
  - link-speed statement.....3019
  - link-state acknowledgment packets *See* OSPF, link-state acknowledgment packets
  - link-state advertisements *See* OSPF, link-state advertisements
  - link-state paths *See* LSPs
  - link-state PDUs.....4160
  - lldp statement.....1938
  - lldp-configuration-notification-interval statement.....1940
  - lmp (tracing flag).....5338
  - lo0 interface.....20, 222
  - load balancing
    - advertising multiple paths to a destination.....3651, 3844
    - for PIM join.....5513
  - load command.....358
    - usage guidelines.....65
  - load merge command
    - usage guidelines.....1326
  - load override command
    - usage guidelines.....1326
  - load set command
    - usage guidelines.....1327
  - load-balance statement.....5321
    - usage guidelines.....3828
  - load-key-file command
    - usage guidelines.....125, 146, 1821, 1845, 1853
  - load-key-file statement.....279
    - usage guidelines.....125, 145, 1821, 1824, 1845, 1850, 1853
  - local AS
    - BGP.....3704, 3709, 3719
  - local statement.....6310
    - PIM.....5722
      - usage guidelines.....5545
  - local-address statement.....3323
    - BFD.....3281
      - usage guidelines.....3237
    - BGP.....4020
    - BMP.....4022
    - MSDP group.....5776
    - MSDP peer.....5776
    - PIM.....5723
      - usage guidelines.....3625, 3648
  - local-as statement.....4023
    - usage guidelines.....3709, 3719
  - local-engine statement.....1635
  - local-interface statement
    - usage guidelines.....3609
  - local-port statement
    - BMP.....4025
      - usage guidelines.....3648
  - local-preference statement.....4026
    - usage guidelines.....3653
  - local0 - local7 (options to facility-override statement).....1548
  - location statement.....228
    - SNMP.....1636, 1940
      - usage guidelines.....29
  - location, chassis.....945
  - log files
    - clearing contents of.....332, 1368
    - contents, displaying.....1003, 1784, 2727
    - filenames for DDoS protection.....6079
    - number of DDoS protection.....6079
    - size of DDoS protection.....6079
    - specifying properties.....1543
  - log-out-on-disconnect statement.....15, 221, 1411
  - log-prefix statement
    - system logging.....1701
      - usage guidelines.....1531
  - log-updown statement.....4027
    - BGP
      - usage guidelines.....3919
    - LDP.....4867
      - usage guidelines.....4837
    - MPLS.....5055
  - logging
    - routing protocol process.....3333
  - logging in as root.....1413
  - logging operations
    - tracing operations.....1409
  - logging out users.....390
  - logical interfaces
    - interface addresses.....2824
  - logical operators
    - for monitor traffic command.....1716
  - logical systems
    - configuring default route.....4232
    - configuring IS-IS.....4232
    - configuring routing policy.....4232

|                                                 |                                   |                                            |            |
|-------------------------------------------------|-----------------------------------|--------------------------------------------|------------|
| EBGP                                            |                                   | LSP splitting                              |            |
| with IPv6 interfaces.....                       | 3609                              | constraints.....                           | 5281       |
| internal BGP.....                               | 3636                              | overview.....                              | 5280       |
| viewing system files on.....                    | 3951                              | supported criteria.....                    | 5281       |
| login announcements, system.....                | 126                               | triggers.....                              | 5282       |
| login classes                                   |                                   | lsp-attributes statement.....              | 5056       |
| access privilege levels.....                    | 1799                              | lsp-generation (tracing flag).....         | 4321       |
| commands, allowing or denying.....              | 1802                              | lsp-history (tracing flag).....            | 5103       |
| defining.....                                   | 1823, 1848                        | lsp-interval statement.....                | 4294       |
| idle timeout values.....                        | 22                                | lsp-lifetime statement.....                | 4295       |
| login messages, system.....                     | 126                               | lsp-ping-interval statement                |            |
| login statement.....                            | 280                               | RSVP LSPs.....                             | 5067       |
| usage                                           |                                   | lsping (tracing flag).....                 | 5103       |
| guidelines.....                                 | 125, 1821, 1823, 1845, 1848, 1853 | LSPs                                       |            |
| login time, displaying.....                     | 167                               | adaptive rerouting.....                    | 5024       |
| login-alarms statement.....                     | 229                               | associating addresses.....                 | 5048       |
| usage guidelines.....                           | 20                                | automatic bandwidth.....                   | 5194       |
| login-tip statement.....                        | 281                               | automatic bandwidth allocation.....        | 5029       |
| logout, users.....                              | 390                               | bandwidth allocation, adjusting.....       | 5122       |
| loopback address, egress policy.....            | 4831                              | BGP, ping interval.....                    | 5112       |
| loopback mode.....                              | 2831                              | CAC information, displaying.....           | 5140, 5160 |
| loopback statement                              |                                   | clearing.....                              | 5107       |
| Aggregated Ethernet, Fast Ethernet, and Gigabit |                                   | configuration statements.....              | 5051       |
| Ethernet.....                                   | 3020                              | constrained-path See constrained-path LSPs |            |
| Fast Ethernet interfaces                        |                                   | egress routers.....                        | 5049       |
| usage guidelines.....                           | 2831                              | explicit-path See explicit-path LSPs       |            |
| Gigabit Ethernet interfaces                     |                                   | fast reroute.....                          | 5031, 5042 |
| usage guidelines.....                           | 2831                              | fate-sharing.....                          | 3300       |
| loops statement                                 |                                   | hold time.....                             | 5027       |
| BGP address family.....                         | 4028                              | ingress routers.....                       | 5043       |
| loose-authentication-check statement            |                                   | intermediate routers.....                  | 5049       |
| IS-IS.....                                      | 4293                              | LDP, displaying.....                       | 4904       |
| usage guidelines.....                           | 4184                              | LDP, ping interval.....                    | 4887       |
| loose-check statement                           |                                   | metrics.....                               | 5057       |
| BFD authentication.....                         | 5724                              | MPLS, displaying.....                      | 5176       |
| BGP.....                                        | 4030                              | named paths.....                           | 5073       |
| loss priority, CoS.....                         | 7292                              | path                                       |            |
| loss-priority statement                         |                                   | connection retry information.....          | 5085       |
| class of service.....                           | 7084                              | length.....                                | 5045, 5317 |
| system drop profiles.....                       | 7108                              | preference levels.....                     | 5081       |
| loss-priority statement (rewrite rules)         |                                   | primary.....                               | 5082       |
| class of service.....                           | 7085                              | reoptimization.....                        | 5068, 5071 |
| LSAs See OSPF, link-state advertisements        |                                   | RSVP, ping interval.....                   | 5347       |
| lsp (tracing flag).....                         | 4321                              | RSVP, real-time status.....                | 5109       |
| LSP merging                                     |                                   | RSVP-signaled.....                         | 4971       |
| constraints.....                                | 5283                              | secondary.....                             | 5088       |
| overview.....                                   | 5282                              | standby state.....                         | 5091       |
| triggers.....                                   | 5283                              | static See static LSPs                     |            |
|                                                 |                                   | text description.....                      | 5036       |

- TTL decrementing, disabling.....5061, 5063
- tunneling through RSVP
  - LSPs.....4805, 4834, 4866, 5053, 5264
- M**
- MAC address table
  - Layer 2, displaying.....2558
- MAC addresses
  - displaying.....949
- mac statement.....2296, 3020
- mac-limit statement.....2316
- mac-move-limit statement.....6137
- mac-notification statement.....2317, 2321
- mac-statistics statement.....2318
- mac-table-size statement.....2320
- MAM.....5032
- management access, sample task.....1809
- management device
  - recovering root password
    - from.....147, 1810, 1818, 1840
- management interfaces
  - overview.....2796
- management software process.....1794
- manuals
  - comments on.....cxlili
- mapping-agent-election statement.....5725
  - usage guidelines.....5561
- mappings
  - SSM.....3376, 5765
- martian addresses.....3324
- martians statement.....3324
- martians, displaying.....3508
- match conditions
  - for monitor traffic command.....1714
  - routing policy.....4528
- match statement.....1701
  - usage guidelines.....1551
- max-age statement.....2345
- max-areas statement.....4296
- max-bypasses statement.....5322
- max-configurations-on-flash statement.....231
- max-hops statement.....2346
- maximum hop count, RIP.....4680
- maximum statement
  - MSDP.....5777
    - usage guidelines.....5604
- maximum-bandwidth statement.....3325, 5057
- maximum-certificates statement.....6311
- maximum-ecmp statement
  - usage guidelines.....3870
- maximum-helper-recovery-time statement.....5323
- maximum-helper-restart-time statement.....5324
- maximum-length statement.....297
  - usage guidelines.....1824, 1850
- maximum-neighbor-recovery-time
  - statement.....4868
    - usage guidelines.....4830
- maximum-paths statement.....3326
- maximum-prefixes statement.....3328
- maximum-recovery-time statement.....4868
- maximum-rps statement.....5726
- maximum-transmit-rate statement
  - IGMP.....5628
    - usage guidelines.....5418
  - MLD
    - usage guidelines.....5454
- MBGP MVPNs.....3936
- MD5 authentication.....3898
  - BGP.....3898
- MD5 checksum, calculating.....344, 1372
- MED See BGP
- MED (multiple exit discriminator)
  - always compare option.....3667
  - Cisco non-deterministic option.....3667
  - plus IGP option.....3667
- med-igp-update-interval statement.....3329
  - usage guidelines.....3694
- med-plus-igp statement.....4058
  - usage guidelines.....3793
- member statement.....3022
- member-range statement.....3023
- members statement.....2297
  - usage guidelines.....3891
- memory utilization, displaying.....165
- mesh groups.....4297
  - MSDP.....5604
- mesh-group statement.....4297
- message statement.....281
  - usage guidelines.....126
- message-processing-model statement.....1636
- message-size statement.....4775
- messages
  - boot, displaying.....1038
  - broadcast messages, NTP.....135, 287
  - LDP message types.....4807
  - MPLS system log.....5055
  - multicast, NTP.....136

|                                           |            |                                     |            |
|-------------------------------------------|------------|-------------------------------------|------------|
| redirect.....                             | 18         | views                               |            |
| system login.....                         | 126        | SNMP.....                           | 1506       |
| user screens, displaying on.....          | 381        | VLAN.....                           | 1477       |
| metric                                    |            | minimum-bandwidth statement.....    | 5058       |
| traffic engineering.....                  | 4511       | minimum-changes statement.....      | 298        |
| metric statement                          |            | usage guidelines.....               | 1824, 1850 |
| aggregate routes.....                     | 3330       | minimum-interval.....               | 4034       |
| BGP                                       |            | BFD, transmit-interval.....         | 4036       |
| usage guidelines.....                     | 3680       | PIM.....                            | 5728       |
| generated routes.....                     | 3330       | usage guidelines.....               | 3812       |
| IS-IS.....                                | 4298       | minimum-interval statement          |            |
| MPLS.....                                 | 5057       | BFD.....                            | 3281       |
| OSPF                                      |            | usage guidelines.....               | 3237       |
| usage guidelines.....                     | 4462       | BGP.....                            | 3981       |
| static routes.....                        | 3330       | IS-IS.....                          | 4266       |
| metric-in statement                       |            | OSPF.....                           | 4586       |
| RIP.....                                  | 4776       | usage guidelines.....               | 4484       |
| usage guidelines.....                     | 4729       | PIM.....                            | 5727       |
| metric-out statement                      |            | usage guidelines.....               | 5500       |
| BGP.....                                  | 4032       | RIP.....                            | 4767       |
| usage guidelines.....                     | 3668       | usage guidelines.....               | 4702       |
| RIP.....                                  | 4777       | minimum-length statement.....       | 299        |
| metrics                                   |            | usage guidelines.....               | 1824, 1850 |
| IS-IS.....                                | 4317       | minimum-lower-cases statement.....  | 300        |
| LDP tracking IGP.....                     | 4833       | minimum-numeric statement.....      | 301        |
| OSPF.....                                 | 4461, 4608 | minimum-punctuations statement..... | 302        |
| static LSP metric.....                    | 5057       | minimum-receive-interval            |            |
| mgd process.....                          | 1794       | BFD.....                            | 4038       |
| MIBs                                      |            | minimum-receive-interval statement  |            |
| Alarm.....                                | 1474       | BFD.....                            | 3281       |
| Analyzer.....                             | 1474, 1482 | usage guidelines.....               | 3237       |
| Chassis.....                              | 1474, 1482 | BFD (BGP)                           |            |
| Chassis Definitions for Router Model..... | 1474       | usage guidelines.....               | 3812       |
| Class-of-Service.....                     | 1474, 1482 | BGP.....                            | 3981       |
| Configuration Management.....             | 1475, 1483 | IS-IS.....                          | 4266       |
| Ethernet MAC.....                         | 1475       | OSPF.....                           | 4586       |
| Event.....                                | 1475       | PIM.....                            | 5697, 5729 |
| Firewall.....                             | 1475       | usage guidelines.....               | 5500       |
| Host Resources.....                       | 1476       | RIP.....                            | 4767       |
| Interface.....                            | 1476, 1483 | usage guidelines.....               | 4702       |
| Ping.....                                 | 1476       | minimum-receive-ttl statement       |            |
| Power Supply.....                         | 1483       | BFD.....                            | 3281       |
| RMON Events and Alarms .....              | 1477       | BFD (static routes)                 |            |
| SNMP object values, displaying.....       | 1756       | usage guidelines.....               | 3237       |
| standards documents.....                  | 1469       | minimum-upper-cases statement.....  | 303        |
| Structure of Management Information.....  | 1477       | mixed VLAN tagging.....             | 2368, 2370 |
| System Log.....                           | 1477       | MLD                                 |            |
| Utility.....                              | 1460       | disabling.....                      | 5466       |
|                                           |            | event recording.....                | 5462       |

- group membership
  - SSM maps for different groups to different sources.....5541
- host-query message interval.....5449
- immediate-leave host removal
  - configuring.....5451
- last-member query interval.....5450
- overview.....5441, 5444
- PIM-to-MLD message translation information, displaying.....5861
- query response interval.....5450
- robustness variable.....5453
- static group membership.....5454
- mld
  - enabling.....5448
- MLD snooping
  - host-query message interval.....5653, 5676
- mld statement
  - usage guidelines.....5447, 5448
- mld-snooping statement.....5674
- mode statement.....2347, 6564, 7135
  - MSDP.....5778
    - usage guidelines.....5604
  - PIM.....5729
    - usage guidelines.....5512, 5527
- monitor
  - interface command output fields, table.....1417
- monitor interface command.....3090
- monitor interface traffic command.....1416
- monitor label-switched-path command.....5109
- monitor statement
  - BMP.....4039
    - usage guidelines.....3649
- monitor traffic command.....1712
- monitor-bandwidth statement.....5058
- monitoring
  - interface.....167, 2844
  - routing tables.....
  - system process information.....165
  - system properties.....166
- monitoring tools
  - tracing operations.....1409
- MOSPF, defined.....5398
- MP-BGP.....3997
- MPLS
  - automatic bandwidth allocation.....5029
  - automatic bandwidth allocation, statistics.....4986
- BGP-signaled LSP connections
  - operability, checking.....5112
- CSPF statistics, displaying.....5142, 5162
- DiffServ classes, displaying.....5144, 5164
- fast reroute.....5031, 5042
- firewall filter.....4977
- interfaces, displaying.....5154, 5174
- IPv6.....5002
- labels, displaying routes.....3503
- Layer 2 circuit connections
  - operability, checking.....5114
- Layer 3 VPN connections
  - operability, checking.....5117
- LDP-signaled LSP connections
  - operability, checking.....4887
- link-management information, displaying
  - all.....5127
  - peers.....5131
  - routing process.....5133
  - statistics.....5136, 5156
  - traffic-engineered links.....5138, 5158
- LSP endpoint connections
  - operability, checking.....5120
- LSPs, displaying.....5197
- protocol family.....2824
- route forwarding-table, displaying.....5146, 5166
- SNMP traps.....5055
- static.....5049
- static LSPs, displaying.....5241
- statistics output.....4980
- system log messages.....5055
- tracing LSPs.....4924, 5382
- tracing protocol operations.....5102
- traffic statistics.....4980, 5094
- tunneling
  - IPv6.....5002
  - ultimate-hop popping.....5041
- MPLS, graceful restart.....2579
- mru statement.....7149
- ms-chapv2
  - changing password ms-chapv2.....1875, 1895
- MSDP
  - active source limit.....5768
    - maximum.....5777
    - per-source.....5783
    - threshold.....5784
  - authentication.....5769
  - cache entries, clearing.....5932
  - configuration statements.....5598



|                                             |            |                                              |            |
|---------------------------------------------|------------|----------------------------------------------|------------|
| configuring.....                            | 5598       | leaf and branch.....                         | 5394       |
| data-encapsulation.....                     | 5770       | NTP messages.....                            | 136        |
| default peer.....                           | 5604, 5771 | packet replication.....                      | 5401       |
| enabling.....                               | 5779       | protocols                                    |            |
| general information, displaying.....        | 5934       | group membership.....                        | 5405       |
| groups.....                                 | 5774       | reverse-path forwarding (RPF).....           | 5393       |
| local address.....                          | 5776       | routing protocols.....                       | 5398       |
| message source information, displaying..... | 5936       | compared, table.....                         | 5400       |
| mode.....                                   | 5778       | scoping.....                                 | 3371       |
| nonstop active routing.....                 | 2601       | shortest-path tree (SPT).....                | 5393       |
| peer statistics                             |            | snooping.....                                | 5395       |
| clearing.....                               | 5933       | SSM groups.....                              | 3375, 5763 |
| displaying.....                             | 5941       | SSM mapping.....                             | 3376, 5765 |
| policy, routing.....                        | 5773, 5775 | terminology.....                             | 5392       |
| routing tables.....                         | 5782       | uses.....                                    | 5391       |
| source-active cache, displaying.....        | 5938       | mcast filters.....                           | 5571       |
| tracing operations.....                     | 5600       | MAC filters.....                             | 5572       |
| msdp statement.....                         | 5779       | MSDP SA messages.....                        | 5600       |
| msti statement.....                         | 2348       | RP/DR register messages.....                 | 5572       |
| MSTP                                        |            | configuring.....                             | 5577       |
| configuration                               |            | mcast group joins                            |            |
| displaying.....                             | 2431       | limiting.....                                | 5426, 5464 |
| configuring.....                            | 2170       | Multicast Listener Discovery See MLD         |            |
| VLAN.....                                   | 2360       | Multicast Open Shortest Path First See MOSPF |            |
| mstp statement.....                         | 2349       | Multicast Source Discovery Protocol See MSDP |            |
| usage guidelines.....                       | 2170       | mcast statement.....                         | 3331       |
| MSTP, nonstop bridging.....                 | 2592       | mcast-client statement.....                  | 287        |
| mt (tracing flag).....                      | 5755       | usage guidelines.....                        | 136        |
| mtrace (tracing flag)                       |            | mcast-group statement.....                   | 2526       |
| IGMP.....                                   | 5427       | mcast-router-interface statement             |            |
| mtrace command.....                         | 5839       | IGMP snooping.....                           | 5652, 5675 |
| mtrace from-source command.....             | 5842       | multihop                                     |            |
| mtrace monitor command.....                 | 5845       | BGP.....                                     | 3777       |
| mtrace to-gateway command.....              | 5847       | multihop statement.....                      | 4042       |
| MTU signaling, in RSVP.....                 | 5264       | usage guidelines.....                        | 3777       |
| mtu statement.....                          | 3024       | multilevel IS-IS.....                        | 4175       |
| mtu-discovery statement.....                | 4040       | multipath statement                          |            |
| mtu-signaling statement.....                | 5059       | usage guidelines.....                        | 3828       |
| multiarea adjacency                         |            | multiple commit scripts.....                 | 1429       |
| OSPF.....                                   | 4423       | Multiple Spanning Tree Protocol See MSTP     |            |
| multiarea network.....                      | 4408       | multiplier                                   |            |
| multicast                                   |            | BFD.....                                     | 4044       |
| addresses.....                              | 5394       | multiplier statement                         |            |
| administrative scoping.....                 | 5394       | BFD.....                                     | 3281       |
| anycast RP.....                             | 5551       | usage guidelines.....                        | 3237       |
| auto-RP.....                                | 5561       | BFD (BGP)                                    |            |
| bootstrap router.....                       | 5567       | usage guidelines.....                        | 3812       |
| defined.....                                | 5389       | BGP.....                                     | 3981       |
| Layer 2 frames.....                         | 5395       | IS-IS.....                                   | 4266       |



|                                                      |                  |                                                        |      |
|------------------------------------------------------|------------------|--------------------------------------------------------|------|
| OSPF.....                                            | 4586             | sample BGP route reflectors (cluster of clusters)..... | 3875 |
| usage guidelines.....                                | 4484             | sample BGP route reflectors (multiple clusters).....   | 3875 |
| PIM.....                                             | 5697, 5730       | sample distance-vector routing.....                    | 4678 |
| usage guidelines.....                                | 5500             | sample multiarea OSPF routing.....                     | 4404 |
| RIP.....                                             | 4767             | sample OSPF network with stubs and NSSAs.....          | 4413 |
| usage guidelines.....                                | 4702             | sample poison reverse routing.....                     | 4681 |
| multiprotocol BGP (MP-BGP).....                      | 3997             | sample split horizon routing.....                      | 4680 |
| multiservice statement.....                          | 3056             | sample unidirectional routing.....                     | 4681 |
| <b>N</b>                                             |                  |                                                        |      |
| name servers, DNS.....                               | 15, 26           | next hop (from operator, LDP).....                     | 4811 |
| name statement.....                                  | 1637, 1942, 3060 | next hop, displaying.....                              |      |
| name-server statement.....                           | 231              | next hops                                              |      |
| usage guidelines.....                                | 15, 26           | multicast entries, displaying.....                     | 5856 |
| named paths                                          |                  | PFE, displaying.....                                   | 7201 |
| empty paths.....                                     | 5073             | resolution database, displaying.....                   | 3542 |
| names                                                |                  | routes sent to, displaying.....                        | 3510 |
| domain names on routers.....                         | 26               | next-generation RIP See RIPng                          |      |
| names.....                                           | 34               | NLRI, BGP.....                                         | 3597 |
| router .....                                         | 23               | nlri-route-type statement                              |      |
| nas-ip-address statement .....                       | 1942             | usage guidelines.....                                  | 3936 |
| native-vlan-id statement.....                        | 2298, 2370       | no-accept-data statement.....                          | 2684 |
| neighbor (from operator, LDP).....                   | 4811             | usage guidelines.....                                  | 2648 |
| neighbor statement                                   |                  | no-accounting statement                                |      |
| BGP.....                                             | 4046             | IGMP.....                                              | 5618 |
| OSPF                                                 |                  | no-adaptation                                          |      |
| usage guidelines.....                                | 4438             | BFD.....                                               | 4050 |
| RIP.....                                             | 4778             | BFD (BGP)                                              |      |
| neighbor-policy statement.....                       | 5730             | usage guidelines.....                                  | 3812 |
| usage guidelines.....                                | 5573             | PIM.....                                               | 5731 |
| neighbors                                            |                  | no-adaptation statement                                |      |
| BGP.....                                             | 3594, 4956       | BFD.....                                               | 3281 |
| OSPF.....                                            | 4435             | BFD (static routes)                                    |      |
| network layer reachability information See BGP, NLRI |                  | usage guidelines.....                                  | 3237 |
| network link advertisements.....                     | 4393             | BGP.....                                               | 3981 |
| network PDUs.....                                    | 4158             | IS-IS.....                                             | 4266 |
| network service access point.....                    | 4158             | OSPF.....                                              | 4586 |
| Network Time Protocol See NTP                        |                  | RIP.....                                               | 4767 |
| network-summary-export statement                     |                  | usage guidelines.....                                  | 4702 |
| usage guidelines.....                                | 4544             | no-adjacency-holddown statement.....                   | 4299 |
| network-summary-import statement                     |                  | no-advertise-peer-as statement.....                    | 4051 |
| usage guidelines.....                                | 4554             | usage guidelines.....                                  | 3773 |
| networks                                             |                  | no-aggregate statement.....                            | 5306 |
| sample BGP confederations.....                       | 3891             | no-aggregator-id statement.....                        | 4052 |
| sample BGP MED use.....                              | 3666             | no-allow-url statement                                 |      |
| sample BGP peer session.....                         | 3601             | op scripts.....                                        | 1582 |
| sample BGP route reflector (one cluster).....        | 3874             | no-authentication-check statement.....                 | 4300 |
|                                                      |                  | usage guidelines.....                                  | 4183 |

|                                                |            |                                                   |            |
|------------------------------------------------|------------|---------------------------------------------------|------------|
| no-broadcast statement.....                    | 6124, 6125 | no-record statement.....                          | 5064, 5084 |
| no-check-zero statement.....                   | 4769       | no-redirects statement.....                       | 233        |
| no-client-reflect statement.....               | 4053       | usage guidelines.....                             | 18         |
| no-cmd-attribute-value statement.....          | 1982       | no-reliable statement.....                        | 5332       |
| no-compress-configuration-files statement..... | 220        | no-remote-tracing.....                            | 1567       |
| usage guidelines.....                          | 14, 1351   | no-retain statement.....                          | 3358       |
| no-csnp-authentication statement.....          | 4300       | no-rfc-1583 statement.....                        | 4602       |
| usage guidelines.....                          | 4183       | usage guidelines.....                             | 4428       |
| no-cspf statement.....                         | 5060, 5325 | no-root-port statement.....                       | 2351       |
| no-decrement-ttl statement.....                | 5061       | no-saved-core-context statement.....              | 236        |
| no-delegate-processing statement               |            | usage guidelines.....                             | 29         |
| periodic packet management.....                | 3343       | no-strict-lsa-checking statement.....             | 2667       |
| no-eligible-backup statement                   |            | usage guidelines                                  |            |
| IS-IS.....                                     | 4301       | OSPF.....                                         | 2583       |
| no-forwarding statement.....                   | 4869       | OSPFv3.....                                       | 2583       |
| usage guidelines.....                          | 4833       | no-trap statement.....                            | 5065       |
| no-gratuitous-arp-request statement.....       | 3025, 6139 | no-unicast-topology statement.....                | 4307       |
| no-hello-authentication statement.....         | 4301       | no-unknown-unicast statement.....                 | 6126       |
| usage guidelines.....                          | 4183       | no-world-readable option to traceoptions          |            |
| no-install statement.....                      | 3316       | statement                                         |            |
| no-install-to-address statement.....           | 5062       | LDP.....                                          | 4878       |
| no-interface-hello statement.....              | 5325       | MPLS.....                                         | 5103       |
| no-ipv4-multicast statement.....               | 4302       | RSVP.....                                         | 5339       |
| no-ipv4-routing statement.....                 | 4303       | no-world-readable statement                       |            |
| no-ipv6-multicast statement.....               | 4304       | archiving of all system log files.....            | 1692       |
| no-ipv6-routing statement.....                 | 4305       | archiving of individual system log file.....      | 1694       |
| no-ipv6-unicast statement.....                 | 4306       | system logging                                    |            |
| no-local-reversion statement.....              | 5326       | usage guidelines.....                             | 1543       |
| no-loopback statement.....                     | 3020       | Node protection.....                              | 5284       |
| usage guidelines.....                          | 2831       | node-hello statement.....                         | 5328       |
| no-multicast-echo statement.....               | 259        | node-link-protection statement.....               | 4308, 5066 |
| PIM                                            |            | usage guidelines                                  |            |
| usage guidelines.....                          | 5495       | IS-IS.....                                        | 4168, 4220 |
| usage guidelines.....                          | 86         | non-link-protected aggregated Ethernet interfaces |            |
| no-node-id-subobject statement.....            | 5327       | periodically rebalancing subscribers.....         | 2875       |
| no-nssa-abr statement.....                     | 4601       | rebalancing subscribers.....                      | 2875       |
| usage guidelines.....                          | 4418       | nonstop active routing                            |            |
| no-p2mp-sublsp statement.....                  | 5327       | concepts.....                                     | 2597       |
| no-ping-record-route statement.....            | 232        | overview.....                                     | 2617       |
| no-ping-time-stamp statement.....              | 232        | platform support.....                             | 2600       |
| no-preempt statement.....                      | 2694       | protocol support.....                             | 2601       |
| no-prepend-global-as statement                 |            | system architecture.....                          | 2597       |
| usage guidelines.....                          | 3709       | system requirements.....                          | 2600       |
| no-propagate-ttl statement.....                | 5063       | nonstop bridging                                  |            |
| no-psnp-authentication statement.....          | 4306       | concepts.....                                     | 2589       |
| usage guidelines.....                          | 4183       | platform support.....                             | 2591       |
| no-qos-adjust statement.....                   | 3332       | protocol support.....                             | 2592       |
| no-readvertise statement.....                  | 3352       | system architecture.....                          | 2589       |
| no-recommendation-tlv statement.....           | 6570       | system requirements.....                          | 2591       |

- 
- nonstop-bridging statement.....2675
  - nonstop-routing statement.....2676
  - nonvolatile statement.....1637, 1943
  - normal (tracing flag).....2679, 3384
    - BMP.....4084
  - Normalization
    - auto-bandwidth inter-operation.....5286
    - computed range.....5288
    - constraints.....5285
    - feasible range.....5288
    - overview.....5285
  - not-so-stubby areas *See* NSSAs
    - configuring.....4418
  - notice (system logging severity level).....3333
  - notification (tracing flag).....4878
  - notification messages
    - LDP.....4808
  - notify statement.....1638
  - notify-duration statement.....2668
    - usage guidelines.....2583
  - notify-filter statement
    - for applying to target.....1639
    - for configuring.....1639
  - notify-view statement.....1640
  - NPDU.....4158
  - NSAP.....4158
  - NSR *See* nonstop active routing
  - nsr-synchronization (tracing flag).....5103, 5756
  - nsr-synchronization-detail (tracing flag).....5103
  - nssa
    - configuring.....4418
  - nssa statement
    - usage guidelines.....4418
  - NSSAs (not-so-stubby areas)
    - description.....4413
    - overview.....4399
  - NTP
    - authentication keys.....132
    - boot server.....137
    - broadcast mode.....131, 133, 134
    - client mode.....131, 133
    - configuration example.....139
    - configuring.....137
    - listening
      - for broadcast messages.....135, 287
      - for multicast messages.....136
    - peer status, displaying.....450
    - peer values, displaying.....452
    - security configuration example.....142
    - server mode.....135
    - symmetric active mode.....131, 133, 134
  - ntp statement.....288
    - usage guidelines.....137
  - O**
    - OAM periodic traceroute, LDP.....4832
    - oam statement.....2948
      - RSVP LSPs.....5067
    - oid statement
      - SNMP.....1640, 1943
        - usage guidelines.....1506
      - SNMPv3.....1641
    - oif-map statement
      - IGMP.....5628
    - op scripts
      - boilerplate.....1436
      - flow of operation illustrated.....1436
      - using.....1436
    - op statement.....1583
    - open messages, BGP.....3597
    - Open Shortest Path First *See* OSPF *See* OSPF
    - operating system *See* JUNOS software
    - operational mode commands
      - displaying output from commit scripts.....1445
    - operators, regular
      - expression.....1827, 1828, 1849, 1850
      - system logging.....1552
    - optimize-aggressive statement.....5068
    - optimize-timer statement
      - bypass LSPs.....5328
      - MPLS.....5071
    - optimizing LSPs.....5068, 5071
    - optional statement.....234, 1584
    - options statement.....3333
    - order statement.....1944
    - ORF
      - BGP.....3773
    - OSPF
      - adjacencies.....4580
      - area border routers.....4404, 4451 *See* area border routers
      - areas.....4404
        - configuring.....4580
        - See also* area border routers; backbone area; NSSAs; stub areas
      - AS external link advertisements.....4393
      - authentication.....4584
      - backbone.....4406, 4580

|                                             |                        |
|---------------------------------------------|------------------------|
| backbone area See backbone area             |                        |
| backup coverage                             |                        |
| displaying.....                             | 4628                   |
| backup neighbor paths.....                  | 4631                   |
| bandwidth-based metrics                     |                        |
| configuring.....                            | 4466                   |
| BFD.....                                    | 4482, 4586             |
| configuring, router identifier.....         | 4400                   |
| context identifier, displaying.....         | 4633                   |
| controlling designated router election..... | 4402                   |
| cost See OSPF, metrics                      |                        |
| database description packets.....           | 4392                   |
| database protection                         |                        |
| configuring.....                            | 4524                   |
| overview.....                               | 4523                   |
| statement.....                              | 4590                   |
| dead interval.....                          | 4476                   |
| default route.....                          | 4398                   |
| demand circuits.....                        | 4439                   |
| designated router.....                      | 4400, 4435             |
| enabling.....                               | 4431, 4599, 4603       |
| enabling, description.....                  | 4400, 4406, 4408       |
| error packets.....                          | 4611                   |
| graceful restart.....                       | 2579, 2664, 4495, 4596 |
| hello interval.....                         | 4476                   |
| hello packets.....                          | 4391                   |
| interfaces, displaying.....                 | 4643                   |
| link-state                                  |                        |
| acknowledgment packets.....                 | 4392                   |
| advertisements.....                         | 4392, 4393, 4476       |
| flooding packets.....                       | 4611                   |
| request packets.....                        | 4392                   |
| update packets.....                         | 4392                   |
| link-state database entries, displaying     |                        |
| version 2.....                              | 4635                   |
| LSAs See OSPF, link-state advertisements    |                        |
| metrics.....                                | 4461, 4608             |
| multiarea adjacency                         |                        |
| configuring.....                            | 4424                   |
| overview.....                               | 4423                   |
| multiarea network.....                      | 4408                   |
| NBMA networks.....                          | 4435                   |
| neighbors.....                              | 4435                   |
| clearing connections.....                   | 4623                   |
| displaying.....                             | 4654                   |
| network link advertisements.....            | 4393                   |
| nonbroadcast, multiaccess networks.....     | 4435                   |
| nonstop active routing.....                 | 2601                   |
| overload bit.....                           | 4604                   |
| overview                                    |                        |
| displaying.....                             | 2731, 4660             |
| packets.....                                | 4391, 4393, 4611       |
| passive mode.....                           | 4442                   |
| policy, routing.....                        | 4594, 4598             |
| network summaries.....                      | 4544, 4554             |
| network summaries, overview.....            | 4543                   |
| route install priority.....                 | 4539                   |
| preferences.....                            | 3786, 4595, 4606       |
| prefix limit.....                           | 4457, 4607             |
| redistributing routes into IS-IS.....       | 4185                   |
| refresh.....                                | 4459                   |
| route cost See OSPF, metrics                |                        |
| route preference.....                       | 4388                   |
| route selection.....                        | 4461                   |
| route summarization.....                    | 4451, 4582             |
| router identifier.....                      | 3369, 4400             |
| router link advertisements.....             | 4393                   |
| routing algorithm.....                      | 4388                   |
| routing table entries, displaying.....      | 4665                   |
| single-area network.....                    | 4406                   |
| SPF.....                                    | 4388, 4611             |
| SPF calculations, displaying.....           | 4651                   |
| statistics, general                         |                        |
| clearing.....                               | 4625                   |
| displaying.....                             | 4671                   |
| statistics, I/O                             |                        |
| clearing.....                               | 4622                   |
| displaying.....                             | 4649                   |
| summary link advertisements.....            | 4393                   |
| timers.....                                 | 4475                   |
| topological database.....                   | 4386                   |
| tracing operations.....                     | 4611                   |
| database description PDUs.....              | 4569                   |
| demand circuit extensions.....              | 4569                   |
| error PDUs.....                             | 4569                   |
| event.....                                  | 4569                   |
| graceful restart.....                       | 4569                   |
| hello PDUs.....                             | 4569                   |
| LDP synchronization.....                    | 4569                   |
| link-state acknowledgement                  |                        |
| PDUs.....                                   | 4569                   |
| link-state analysis PDUs.....               | 4569                   |
| link-state PDUs.....                        | 4569                   |
| link-state request PDUs.....                | 4569                   |
| link-state updates PDUs.....                | 4569                   |
| NSR synchronization.....                    | 4569                   |
| packet dump.....                            | 4569                   |
| policy processing.....                      | 4569                   |

|                                               |            |
|-----------------------------------------------|------------|
| protocol task processing.....                 | 4569       |
| protocol timer processing.....                | 4569       |
| restart-signaling.....                        | 4569       |
| route information.....                        | 4569       |
| SPF calculations.....                         | 4569       |
| state transitions.....                        | 4569       |
| traffic control.....                          | 4461       |
| traffic engineering                           |            |
| features.....                                 | 4614       |
| transmission delay.....                       | 4476       |
| OSPF (Open Shortest Path First)               |            |
| injecting OSPF routes into BGP.....           | 3766, 4529 |
| monitoring.....                               | 4567       |
| OSPF database protection.....                 | 4523       |
| OSPF metric                                   |            |
| configuring.....                              | 4462       |
| OSPF reference bandwidth                      |            |
| configuring.....                              | 4462       |
| OSPF routing information.....                 | 4567       |
| ospf statement.....                           | 4603       |
| ospf3 statement.....                          | 4603       |
| OSPFv2                                        |            |
| restart-signaling.....                        | 4611       |
| OSPFv3                                        |            |
| multiple address families                     |            |
| configuring.....                              | 4446       |
| understanding.....                            | 4446       |
| nonstop active routing.....                   | 2601       |
| overview.....                                 | 4390       |
| other-routing-engine option to host           |            |
| statement.....                                | 1699       |
| out-delay statement.....                      | 4054       |
| usage guidelines.....                         | 3772       |
| out-of-band management interfaces.....        | 2796       |
| outbound-route-filter                         |            |
| usage guidelines.....                         | 3773       |
| outbound-route-filter statement               |            |
| BGP.....                                      | 4056       |
| outgoing metric (RIP)                         |            |
| description.....                              | 4725       |
| output control keys                           |            |
| for monitor interface command.....            | 3090       |
| for monitor interface traffic command.....    | 3091       |
| for monitor mpls command.....                 | 5109       |
| output statement                              |            |
| forwarding table.....                         | 6095       |
| port mirroring.....                           | 6272       |
| output-traffic-control-profile statement..... | 7108       |
| output-vlan-map statement                     |            |
| Gigabit Ethernet IQ.....                      | 2371       |
| overload bit, resetting for IS-IS.....        | 4334       |
| overload statement                            |            |
| IS-IS.....                                    | 4309       |
| OSPF.....                                     | 4604       |
| usage guidelines.....                         | 4472       |
| overloaded                                    |            |
| configuring routing devices.....              | 4472       |
| routing devices.....                          | 4471       |
| override-interval                             |            |
| PIM.....                                      | 5732       |
| override-interval statement                   |            |
| usage guidelines.....                         | 5517       |
| overview                                      |            |
| CoS inputs and outputs.....                   | 6639       |
| owner statement.....                          | 1642       |
| <b>P</b>                                      |            |
| P2MP LSPs, testing.....                       | 5347       |
| P2MP LSPs, tracing.....                       | 5382       |
| p2mp statement.....                           | 5072       |
| Packet Forwarding Engine.....                 | 1793       |
| packet headers, transmitted, displaying.....  | 1712       |
| packet loss priority, CoS.....                | 7292       |
| packet-action statement.....                  | 2299       |
| packet-dump (tracing flag).....               | 4611       |
| packet-rate statement                         |            |
| ICMPv4.....                                   | 225        |
| usage guidelines.....                         | 86         |
| packets                                       |            |
| router source addresses.....                  | 20, 222    |
| packets (tracing flag)                        |            |
| BGP.....                                      | 4081       |
| BMP.....                                      | 4084       |
| IGMP.....                                     | 5640       |
| IS-IS.....                                    | 4321       |
| LDP.....                                      | 4878       |
| OSPF.....                                     | 4611       |
| PIM.....                                      | 5756       |
| RIP.....                                      | 4784       |
| RSVP.....                                     | 5338       |
| parameters statement.....                     | 1642       |
| parentheses, in syntax descriptions.....      | cxlii      |
| parse (tracing flag).....                     | 2679, 3384 |
| partial sequence number PDUs.....             | 4160       |
| passive ARP learning, VRRP.....               | 2649       |

|                                                 |                       |  |
|-------------------------------------------------|-----------------------|--|
| passive statement                               |                       |  |
| aggregate routes.....                           | 3270                  |  |
| BGP.....                                        | 4057                  |  |
| usage guidelines.....                           | 3919                  |  |
| generated routes.....                           | 3270                  |  |
| IGMP.....                                       | 5629                  |  |
| IS-IS.....                                      | 4312                  |  |
| OSPF                                            |                       |  |
| usage guidelines.....                           | 4442                  |  |
| static routes.....                              | 3270                  |  |
| passive traffic-engineering mode                |                       |  |
| support.....                                    | 4519                  |  |
| password                                        |                       |  |
| for RIPv2 authentication.....                   | 4698                  |  |
| password statement                              |                       |  |
| login.....                                      | 304                   |  |
| passwords                                       |                       |  |
| RADIUS.....                                     | 1873, 1893            |  |
| root.....                                       | 145, 1824, 1850       |  |
| root password,                                  |                       |  |
| recovering.....                                 | 147, 1810, 1818, 1840 |  |
| passwords statement                             |                       |  |
| usage guidelines.....                           | 1824, 1850            |  |
| paste command                                   |                       |  |
| usage guidelines.....                           | 66                    |  |
| path                                            |                       |  |
| calculation                                     |                       |  |
| constrained-path computation.....               | 5060                  |  |
| connection retry information.....               | 5085                  |  |
| length, LSP.....                                | 5045, 5317            |  |
| path (tracing flag)                             |                       |  |
| LDP.....                                        | 4878                  |  |
| RSVP.....                                       | 5338                  |  |
| path attributes, BGP.....                       | 3596, 3597            |  |
| path cost metrics                               |                       |  |
| for RIP routes, description.....                | 4725                  |  |
| for RIP routes, modifying.....                  | 4726, 4728            |  |
| path statement                                  |                       |  |
| MPLS.....                                       | 5073                  |  |
| RSVP.....                                       | 5329                  |  |
| path-count statement                            |                       |  |
| BGP                                             |                       |  |
| usage guidelines.....                           | 3844                  |  |
| path-length statement.....                      | 6311                  |  |
| path-mtu statement.....                         | 5074                  |  |
| path-selection statement.....                   | 4058                  |  |
| usage guidelines.....                           | 3793                  |  |
| pathtear (tracing flag).....                    | 5338                  |  |
| PC See management device                        |                       |  |
| PDU.....                                        | 4160                  |  |
| peak-burst-size statement.....                  | 6111                  |  |
| peak-information-rate statement.....            | 6112                  |  |
| peer interfaces                                 |                       |  |
| configuring.....                                | 4444                  |  |
| peer statement.....                             | 289                   |  |
| MSDP.....                                       | 5781                  |  |
| peer-as statement.....                          | 4060                  |  |
| peer-interface statement                        |                       |  |
| usage guidelines.....                           | 4444                  |  |
| peering sessions See BGP peers; BGP sessions    |                       |  |
| PEM, environmental information, displaying..... | 566                   |  |
| Per-LSP autobandwidth adjustments.....          | 5286                  |  |
| per-packet load balancing.....                  | 3228, 3828            |  |
| per-packet statement                            |                       |  |
| usage guidelines.....                           | 3828                  |  |
| performance-monitor (tracing flag)              |                       |  |
| class of service.....                           | 7059                  |  |
| periodic (tracing flag).....                    | 4878                  |  |
| periodic statement.....                         | 3052                  |  |
| periodic-traceroute statement                   |                       |  |
| usage guidelines.....                           | 4832                  |  |
| permission flags                                |                       |  |
| login class.....                                | 1799                  |  |
| user.....                                       | 1799                  |  |
| permissions statement.....                      | 304                   |  |
| usage guidelines.....                           | 1799                  |  |
| permissions, CLI, displaying.....               | 430                   |  |
| persistent configuration changes                |                       |  |
| compared to transient changes.....              | 1430                  |  |
| overview.....                                   | 1430                  |  |
| persistent-learning statement.....              | 6139                  |  |
| pfc input statement.....                        | 7151                  |  |
| PFE                                             |                       |  |
| firewall filter TCAM, displaying.....           | 6175, 7199            |  |
| next hops, displaying.....                      | 7201                  |  |
| routing table, displaying.....                  | 7206                  |  |
| terse information, displaying.....              | 7215                  |  |
| version, displaying.....                        | 7217                  |  |
| pfe (system logging facility).....              | 1545, 1547            |  |
| physical interfaces                             |                       |  |
| descriptive text.....                           | 2982, 3294            |  |
| loopback mode.....                              | 2831                  |  |
| mixed VLAN tagging.....                         | 2368, 2370            |  |
| physical interfaces, descriptive text.....      | 2982, 3294            |  |
| pic statement.....                              | 3026                  |  |

- 
- PICs
- installed, displaying list.....744
  - operation of, controlling.....371
  - status
    - displaying for a specific PIC.....954
    - displaying FPCs and PICs.....703
- PIM
- anycast RP.....5693, 5747
  - assert timeout.....5694, 5750
    - configuring.....5588
  - background.....5491
  - BFD.....5500, 5697, 5727, 5729, 5730, 5758
  - bidirectional mode
    - defined.....5398
  - bootstrap messages import and export.....5567
  - bootstrap routers.....5567
  - bootstrap routers, displaying.....5881
  - configuring.....5494
  - dense mode.....5523, 5526
    - defined.....5399
  - designated router.....5512
  - embedded RP.....5706
  - enabling.....5733
  - filters See multicast filters
  - graceful restart.....2579
    - sparse-dense mode.....5527
  - groups
    - general information, displaying.....5886
    - usage information, displaying.....5878
  - hello interval.....5494
  - hold-time period.....5715
  - incoming join filter policy, applying.....5575
  - interfaces
    - displaying.....5883
  - join load balancing
    - configuring.....5513
  - join states, clearing.....5832
  - join suppression
    - configuring.....5517
  - join-prune-timeout.....5721
  - maximum RPs.....5726
  - mixing modes.....5525
  - neighbors, displaying.....5907
  - network components.....5493
  - nonstop active routing.....2606
  - outgoing join filter policy, applying.....5574
  - overview.....5491
  - PIM-to-IGMP message translation information,
    - displaying.....5859
  - PIM-to-MLD message translation information,
    - displaying.....5861
  - policy, routing.....5717
  - prune states, clearing.....5832
  - register
    - clearing.....5834
  - remote source.....5602
  - rendezvous point tree.....5580
  - routing tables.....5743
  - RPF, displaying source state.....5918
  - RPs.....5511, 5545, 5561, 5579, 5741, 5744
    - anycast.....5693
    - anycast RP.....5551
    - displaying.....5911
    - embedded.....5706
    - mapping options.....5511
    - maximum.....5726
    - source registration.....5581
    - SPT cutover control.....5588
  - sparse mode.....5509, 5512
    - defined.....5399
  - sparse-dense mode.....5525, 5702
    - defined.....5399
  - SSM.....5530, 5531, 5537
  - statistics
    - clearing.....5836
    - displaying.....5921
    - version.....5494, 5512, 5759
  - pim statement.....5733
    - usage guidelines.....5494
  - PIM-RP
    - SPT
      - configuring threshold cutover
        - policy.....5590
  - pim-to-igmp-proxy statement.....3334
  - pim-to-mld-proxy statement.....3335
  - ping command.....360, 1414, 1724
  - Ping MIB.....1476
  - ping mpls bgp command.....5112
  - ping mpls l2circuit command.....5114
  - ping mpls l3vpn command.....5117
  - ping mpls ldp command.....4887
  - ping mpls lsp-end-point command.....5120
  - ping mpls rsvp command.....5347
  - plain-text passwords.....145
    - configuring for root logins.....1812
    - for user accounts.....1822, 1846
    - root password.....145, 1824, 1850
  - plain-text-password option.....145, 1824, 1850



|                                                       |                  |                                                                                 |            |
|-------------------------------------------------------|------------------|---------------------------------------------------------------------------------|------------|
| point-to-point statement.....                         | 4313             | port-mode statement.....                                                        | 2302       |
| poison reverse technique.....                         | 4680             | port-receive-state-machine (tracing flag)<br>spanning-tree protocols.....       | 2358       |
| policer statement.....                                | 6113             | port-role-select-state-machine (tracing flag)<br>spanning-tree protocols.....   | 2358       |
| policer, single-rate two-color<br>example.....        | 5541             | port-role-transit-state-machine (tracing flag)<br>spanning-tree protocols.....  | 2358       |
| policing statement.....                               | 4870, 5075       | port-state-transit-state-machine (tracing flag)<br>spanning-tree protocols..... | 2358       |
| policy See routing policies                           |                  | port-transmit-state-machine (tracing flag)<br>spanning-tree protocols.....      | 2358       |
| policy (tracing flag).....                            | 2679, 3384       | ports<br>RADIUS servers.....                                                    | 1873, 1893 |
| BMP.....                                              | 4084             | ports statement.....                                                            | 234        |
| policy filters, LDP.....                              | 4858             | post-policy statement.....                                                      | 4061       |
| policy statement<br>aggregate routes.....             | 3336             | usage guidelines.....                                                           | 3649       |
| flow map.....                                         | 3337             | postinheritance, defined.....                                                   | 1425       |
| generated routes.....                                 | 3336             | Power Entry Module.....                                                         | 566        |
| SSM map.....                                          | 5762             | Power Supply Unit MIB.....                                                      | 1483       |
| policy, import<br>BSR.....                            | 5716             | powering off routing software<br>requesting a system power off.....             | 391        |
| policy, routing<br>BGP.....                           | 3996, 4012       | ppm statement.....                                                              | 3343, 3344 |
| forwarding table.....                                 | 3296             | ppmd (tracing flag)<br>spanning-tree protocols.....                             | 2358       |
| IS-IS.....                                            | 4271             | pps.....                                                                        | 2656       |
| MSDP.....                                             | 5773, 5775       | pre-policy statement.....                                                       | 4062       |
| OSPF.....                                             | 4594, 4598       | usage guidelines.....                                                           | 3649       |
| network summaries.....                                | 4544, 4554       | Precision Time Protocol Overview.....                                           | 152        |
| network summaries, overview.....                      | 4543             | preempt statement.....                                                          | 2694       |
| PIM.....                                              | 5717             | preemption statement.....                                                       | 5330       |
| PIM join filter.....                                  | 5574, 5575       | preference levels<br>LDP routes.....                                            | 4828, 4871 |
| precedence.....                                       | 3769             | signaled LSPs.....                                                              | 5081       |
| RIP.....                                              | 4770, 4774       | preference statement<br>aggregate routes.....                                   | 3345       |
| routing instance.....                                 | 3317             | BGP.....                                                                        | 4063       |
| policy-options statement.....                         | 3338, 6571, 7152 | usage guidelines.....                                                           | 3787       |
| policy-statement statement.....                       | 3339, 5076       | generated routes.....                                                           | 3345       |
| poll-interval statement<br>usage guidelines.....      | 4435             | IS-IS.....                                                                      | 4314       |
| polling-interval statement.....                       | 1594             | LDP.....                                                                        | 4871       |
| pop statement<br>Gigabit Ethernet IQ interfaces.....  | 2372             | usage guidelines.....                                                           | 4828       |
| MPLS.....                                             | 5080             | OSPF.....                                                                       | 4606       |
| Port Activation License Support<br>MX104.....         | 115              | usage guidelines.....                                                           | 4468       |
| port statement<br>RADIUS.....                         | 1945             | RIP.....                                                                        | 4779       |
| SNMPv3.....                                           | 1643             | signaled LSPs.....                                                              | 5081       |
| TACACS+.....                                          | 316              | static LSPs.....                                                                | 5081       |
| usage guidelines.....                                 | 1873, 1893       | static routes.....                                                              | 3345       |
| port-error-disable statement.....                     | 6140             |                                                                                 |            |
| port-information-state-machine (tracing<br>flag)..... | 2358             |                                                                                 |            |
| port-migration-state-machine (tracing flag).....      | 2358             |                                                                                 |            |



- preferences
  - active routes.....3793
  - aggregate routes.....3345
    - generated routes.....3786
  - default.....3786
  - IS-IS.....3786, 4272, 4314
  - modifying
    - with configuration statements.....3786
  - OSPF.....4595, 4606
  - RIP.....3786
  - static routes.....3345, 3786
- prefix limit
  - IS-IS.....4315
  - OSPF.....4457, 4607
- prefix list
  - adding to SNMP community.....1505
- prefix list statement
  - firewall filter match condition.....3909
- prefix statement.....3346
- prefix-based
  - usage guidelines.....3773
- prefix-export-limit statement
  - IS-IS.....4315
  - OSPF.....4607
    - usage guidelines.....4457
- prefix-list statement
  - PIM RPF selection.....5736
- prefix-policy statement
  - BGP
    - usage guidelines.....3844
- primary LSPs.....5082
- primary paths
  - revert timer, BFD.....5086
- primary statement
  - MPLS.....5082
- priorities
  - system logging, including in log message
    - for single-chassis system.....1536
- priority
  - PIM RPs.....5739
- priority statement.....2695
  - BMP.....4064
  - bootstrap.....5737
  - class of service.....7109
  - DDoS protection.....6156
  - IS-IS.....4316
  - OSPF
    - usage guidelines.....4402
  - PIM.....5738
    - usage guidelines.....5498
  - RSVP.....5331
  - spanning tree.....2352
    - usage guidelines.....2641, 3649, 5567
- priority-cost statement.....2696
  - usage guidelines.....2646
- priority-flow-control statement.....6572, 7153
- priority-hold-time statement.....2697
  - usage guidelines.....2646
- privacy-3des statement.....1644
- privacy-aes128 statement.....1645
- privacy-des statement.....1646
- privacy-none statement.....1646
- privacy-password statement.....1647
- private statement
  - usage guidelines.....3719
- process (tracing flag)
  - class of service.....7059
- process ID, displaying.....165
- process information, system, monitoring.....165
- process owner, displaying.....165
- process start time, displaying.....165
- process state, displaying.....165
- processes
  - displaying information.....1097
  - restarting.....414
- processes, software
  - chassis process.....1794
  - forwarding process.....1794
  - interface process.....1794
  - management process.....1794
  - routing protocol process.....1794
- profile statement.....1946
- promiscuous-mode statement
  - IGMP
    - interface.....5630
    - usage guidelines.....5415
- propagation, suppressing.....3926
- propagation-delay statement.....5740
  - usage guidelines.....5517
- properties, system, monitoring.....166
- Protocol Independent Multicast *See* PIM *See* PIM
- protocol statement.....2353, 7110
- protocol-specific tracing operations.....1409
- protocol-version statement.....1568, 1960
  - usage guidelines.....1413

|                                                |                        |
|------------------------------------------------|------------------------|
| protocols                                      |                        |
| distance vector See RIP                        |                        |
| group membership.....                          | 5405                   |
| multicast routing.....                         | 5398                   |
| compared, table.....                           | 5400                   |
| originating, displaying.....                   |                        |
| OSPF, monitoring.....                          | 4567                   |
| redirect messages.....                         | 18                     |
| RIP See RIP                                    |                        |
| RIP, monitoring.....                           |                        |
| routing protocols, monitoring.....             | 3959                   |
| protocols statement.....                       | 1947, 2354, 3347       |
| DDoS protection.....                           | 6157                   |
| prune (tracing flag)                           |                        |
| PIM.....                                       | 5756                   |
| prune states, clearing PIM.....                | 5832                   |
| pseudowires                                    |                        |
| mixed VLAN tagging.....                        | 2368                   |
| psn (tracing flag).....                        | 4321                   |
| PSNP IS-IS.....                                | 4160                   |
| push statement                                 |                        |
| Gigabit Ethernet IQ interfaces.....            | 2373                   |
| MPLS.....                                      | 5083                   |
| <b>Q</b>                                       |                        |
| QoS .....                                      | 5263                   |
| qualified-next-hop statement.....              | 3350                   |
| usage guidelines.....                          | 3246                   |
| quality-of-service See QoS                     |                        |
| query-interval statement                       |                        |
| IGMP.....                                      | 5631                   |
| usage guidelines.....                          | 5411                   |
| IGMP snooping.....                             | 5653, 5676             |
| MLD                                            |                        |
| usage guidelines.....                          | 5449                   |
| MLD snooping.....                              | 5653, 5676             |
| query-last-member-interval statement           |                        |
| IGMP.....                                      | 5632                   |
| usage guidelines.....                          | 5412                   |
| IGMP snooping.....                             | 5654, 5677             |
| MLD                                            |                        |
| usage guidelines.....                          | 5450                   |
| query-response-interval statement              |                        |
| IGMP.....                                      | 5633                   |
| usage guidelines.....                          | 5416                   |
| IGMP snooping.....                             | 5655, 5678             |
| MLD                                            |                        |
| usage guidelines.....                          | 5450                   |
| queue-num statement.....                       | 7086                   |
| queues                                         |                        |
| forwarding class mapping, displaying.....      | 7243                   |
| quit command.....                              | 59                     |
| usage guidelines.....                          | 66                     |
| <b>R</b>                                       |                        |
| RADIUS accounting.....                         | 1891                   |
| RADIUS authentication.....                     | 1804, 1847, 1873, 1893 |
| security configuration example.....            | 1878                   |
| RADIUS authorization See RADIUS authentication |                        |
| radius statement.....                          | 1962                   |
| accounting.....                                | 317                    |
| RADIUS templates                               |                        |
| security configuration example.....            | 1879, 1899             |
| radius-options statement .....                 | 1963                   |
| radius-server statement.....                   | 1964                   |
| usage guidelines.....                          | 1873, 1893             |
| Rapid Spanning Tree Protocol See RSTP          |                        |
| rapid-runs statement                           |                        |
| IS-IS.....                                     | 4319                   |
| rate-limit statement.....                      | 1569, 1965             |
| RDM.....                                       | 5032                   |
| reachability                                   |                        |
| verifying for a RIP network.....               | 4691                   |
| read-view statement.....                       | 1648                   |
| readvertise statement.....                     | 3352                   |
| real-time monitoring                           |                        |
| interfaces.....                                | 3090                   |
| IP multicast paths.....                        | 5839                   |
| RSVP LSPs.....                                 | 5109                   |
| traffic.....                                   | 1712                   |
| realm statement                                |                        |
| usage guidelines.....                          | 4446                   |
| rebooting router software                      |                        |
| pending reboots                                |                        |
| clearing.....                                  | 337                    |
| displaying.....                                | 1126                   |
| requesting a system reboot.....                | 6336                   |
| requesting a system reboot with FreeBSD        |                        |
| upgrade.....                                   | 262                    |
| receive (tracing flag modifier)                |                        |
| LDP.....                                       | 4878                   |
| RSVP.....                                      | 5339                   |
| receive statement                              |                        |
| BGP                                            |                        |
| usage guidelines.....                          | 3844                   |
| RIP.....                                       | 4780                   |
| usage guidelines.....                          | 4735                   |
| received label filtering.....                  | 4858                   |

- 
- recommendation-tlv statement.....6574, 7155
  - reconnect-time statement.....4872
    - usage guidelines.....4830
  - record statement.....5064, 5084
  - recovery software installation
    - procedures.....159
  - recovery-time statement.....4873
    - usage guidelines.....4830
  - red alarm conditions.....18
  - RED drop profiles
    - displaying.....7293
  - redirect messages
    - disabling.....18
  - redirected routes.....3786
  - redistributing routes
    - OSPF into IS-IS.....4185
  - redundant-sources statement.....3353
  - redundant-trunk-group statement.....3063
  - reference-bandwidth statement.....4608
    - IS-IS.....4317
    - usage guidelines.....4462
  - reflective-relay statement.....2377
  - refresh statement
    - commit scripts.....1584
    - op scripts.....1585
  - refresh-from statement
    - commit scripts.....1585
    - op scripts.....1586
  - refresh-time statement.....5332
  - regex-parse (tracing flag).....2679, 3384
  - register (tracing flag).....5756
  - register-probe-time statement.....5741
  - regular expression
    - operators.....1827, 1828, 1849, 1850
    - system logging.....1552
  - regular expressions
    - AS paths, displaying matching routes.....3419
    - IP multicast scope
      - clearing.....5829
    - IP multicast sessions
      - clearing.....5830
      - displaying.....5875
    - LSPs, clearing.....5107
  - reinstalling Junos OS.....159
  - reject option to static statement.....3377
  - relative option.....1327
  - reliable statement.....5332
  - Remote Monitoring.....1759
  - remote-debug-permission statement.....1966
  - remote-engine statement.....1649
  - removable media
    - reinstalling Junos OS, using.....159
  - remove-private statement.....4065
    - usage guidelines.....3805
  - removing
    - files.....351, 1379
    - software packages.....1286
  - renaming files.....354, 1382
  - rendezvous points See RPs See PIM and RP
  - reoptimizing LSPs.....5068, 5071
  - replace option.....1326
  - replication
    - multicast packet.....5401
  - replication errors, displaying.....2738
  - report (tracing flag)
    - IGMP.....5640
  - request app-engine cleanup command.....1778
  - request app-engine file-copy (crash | log)
    - from-jhost to-vjunos command.....1776
  - request chassis beacon command.....364
  - request chassis fpc command.....366
  - request chassis pic command.....371
  - request chassis routing-engine master
    - command.....376
  - request command
    - usage guidelines.....59
  - request component login command.....1996
  - request dhcp server reconfigure command.....6332
  - request dhcpv6 server reconfigure
    - command.....6334
  - request message command.....381
  - request mpls container-lsp command.....5352
  - request mpls lsp adjust-autobandwidth
    - command.....5122
  - request snmp spoof-trap command.....1741
  - request system configuration rescue delete
    - command.....382, 1342, 1345, 1346, 1386
  - request system configuration rescue save
    - command.....383, 1342, 1346, 1387
  - request system halt command.....384
  - request system license add command.....117, 438
  - request system license add terminal
    - command.....117
  - request system license delete command.....118, 439
  - request system license save command.....119, 440
  - request system logout command.....390
  - request system power-off command.....391

|                                                |            |                                           |            |
|------------------------------------------------|------------|-------------------------------------------|------------|
| request system reboot (FreeBSD upgrade)        |            | resv (tracing flag).....                  | 5339       |
| command.....                                   | 262        | resvtcar (tracing flag).....              | 5339       |
| request system reboot command.....             | 6336       | retain statement.....                     | 3358       |
| request system snapshot command (upgraded      |            | retransmission-attempt statement.....     | 253        |
| FreeBSD).....                                  | 265        | retransmit-interval statement             |            |
| request system software add                    |            | usage guidelines.....                     | 4476       |
| command.....                                   | 1276, 1290 | retry information.....                    | 5085       |
| request system software delete command.....    | 1286       | retry statement.....                      | 235, 1967  |
| request system software rollback command.....  | 1292       | usage guidelines.....                     | 1873, 1893 |
| request system software validate command.....  | 1297       | retry-count statement.....                | 1651       |
| request system storage cleanup command.....    | 399        | retry-limit statement.....                | 5085       |
| request system zeroize command.....            | 409        | retry-options statement.....              | 305        |
| request-type statement.....                    | 1650       | reverse path forwarding See RPF           |            |
| rerouting LSPs                                 |            | reverse-oif-mapping statement.....        | 3359       |
| adaptive rerouting.....                        | 5024       | reverse-path forwarding See RPF           |            |
| fast reroute.....                              | 5031, 5042 | revert-timer statement.....               | 5086       |
| rescue configuration                           |            | reverting to earlier software.....        | 1293       |
| deleting.....                                  | 382, 1386  | revision-level statement.....             | 2355       |
| displaying.....                                | 1058, 1392 | rewrite rules, displaying                 |            |
| saving.....                                    | 383, 1387  | information.....                          | 7288       |
| reset-tracking-bit statement.....              | 5742       | mapping of code point value to loss       |            |
| usage guidelines.....                          | 5517       | priority.....                             | 7255       |
| resolution statement.....                      | 3354       | mapping of code point value to queue      |            |
| resolution-ribs statement.....                 | 3355       | number.....                               | 7255       |
| resolve statement.....                         | 3356       | table identifiers.....                    | 7257       |
| Resource Reservation Protocol (RSVP), graceful |            | rewrite-rules statement.....              | 7088       |
| restart.....                                   | 2579       | RFC 1583                                  |            |
| restart (tracing flag)                         |            | disabling.....                            | 4428       |
| class of service.....                          | 7059       | OSPFv2.....                               | 4428       |
| restart command.....                           | 414        | rib statement                             |            |
| usage guidelines.....                          | 59         | route resolution.....                     | 3363       |
| restart-duration statement.....                | 2670, 3357 | routing tables.....                       | 3361       |
| IS-IS.....                                     | 4274       | rib-group statement.....                  | 3364       |
| usage guidelines                               |            | IS-IS.....                                | 4318       |
| ES-IS.....                                     | 2582       | MSDP.....                                 | 5782       |
| global.....                                    | 2580       | OSPF.....                                 | 4609       |
| IS-IS.....                                     | 2582       | PIM.....                                  | 5743       |
| OSPF.....                                      | 2583       | usage guidelines.....                     | 5526       |
| OSPFv3.....                                    | 2583       | RIP.....                                  | 4781       |
| PIM.....                                       | 2585       | usage guidelines.....                     | 4742       |
| restart-signaling (tracing flag)               |            | rib-groups statement.....                 | 3365       |
| OSPFv2.....                                    | 4611       | ring-protection-link-end statement.....   | 2269       |
| restart-time statement.....                    | 2671, 4067 | ring-protection-link-owner statement..... | 2269       |
| usage guidelines                               |            | RIP                                       |            |
| BGP.....                                       | 2581       | authentication.....                       | 4693, 4765 |
| RIP.....                                       | 2585       | basic network (configuration editor)..... | 4685       |
| RIPng.....                                     | 2585       | BFD.....                                  | 4701, 4702 |
| restarting                                     |            | disabling address checks.....             | 4764       |
| software processes.....                        | 414        | distance vector protocol.....             | 4677       |

- efficiency techniques.....4680
- enabling.....4781
- general statistics
  - clearing.....4790
  - displaying.....4792
- graceful restart.....2579
- hold-down timer.....4773
- maximum hop count.....4680
- metrics.....4776, 4777
- neighbors.....4778
  - displaying.....4794
- nonstop active routing.....2601
- overview.....4677
- packets.....4679
- poison reverse technique.....4680
- policy, routing.....4770, 4774
- preferences.....3786, 4779
- reserved fields.....4769
- rib-group messages.....4742
- rib-group statement
  - usage guidelines.....4742
- route timeout.....4782
- routing policy (configuration editor).....4685
- split horizon technique.....4680
- statistics
  - clearing.....4791
  - displaying.....4796
- tracing operations
  - authentication.....4757
  - description.....4757, 4758
  - error PDUs.....4757
  - hold-down processing.....4757
  - NSR synchronization.....4757
  - policy processing.....4757
  - protocol task processing.....4757
  - protocol timer processing.....4757
  - request PDUs.....4757
  - route expiration processing.....4757
  - route information.....4757, 4758
  - state transitions.....4757
  - trigger updates.....4757
  - update PDUs.....4757
- traffic control with metrics,
  - configuring.....4726, 4728
- UDP, use of.....4678
- unidirectional limitations.....4681
- update interval.....4787
- update messages.....4735, 4775
- verifying host reachability .....4691
- RIP (Routing Information Protocol)
  - authentication (RIPv2 only).....4693
  - authentication (RIPv2 only),
    - configuring.....4698, 4699
  - overview.....4685
  - path cost metrics See path cost metrics
  - traffic control with metrics See path cost metrics
  - metrics
- RIP routing information.....
- rip statement.....4781
- RIPng
  - graceful restart.....2579
  - nonstop active routing.....2601
- rising-event-index statement.....1652
- rising-threshold statement
  - health monitor.....1653, 1968
  - RMON.....1654
- RJ-45-to-DB-9 serial port
  - adapter.....147, 1810, 1819, 1841
- rlogin service, configuring.....1970
- rmon
  - history.....1763
- RMON alarms.....1465, 1466
- RMON alarms and events, displaying.....1759
- RMON events.....1464
- RMON Events and Alarms MIB.....1477
- RMON history control.....1466
- rmon statement.....1655
- robust-count statement.....5656, 5679
  - IGMP.....5634
  - usage guidelines.....5417
- MLD
  - usage guidelines.....5453
- rollback
  - displaying.....1302, 1394
  - requesting.....1292
- rollback command.....1301, 1337
  - usage guidelines.....66
- rollover cable, connecting the console
  - port.....147, 1811, 1819, 1841
- root logins
  - configuring plain-text passwords for.....1812
- root password.....145, 1824, 1850
- root password recovery.....147, 1810, 1818, 1840
- root-authentication statement.....307, 308
  - usage guidelines.....145, 1824, 1850
- root-login statement.....1969
  - usage guidelines.....1413

|                                           |                              |                                    |                  |
|-------------------------------------------|------------------------------|------------------------------------|------------------|
| route                                     |                              | route-socket (tracing flag)        |                  |
| static statement.....                     | 3377                         | class of service.....              | 7059             |
| route (tracing flag).....                 | 5339                         | route-timeout statement            |                  |
| BMP.....                                  | 4084                         | RIP.....                           | 4782             |
| LDP.....                                  | 4878                         | usage guidelines.....              | 4749             |
| MSDP.....                                 | 5785                         | router advertisements              |                  |
| routing.....                              | 2679, 3384                   | IPv6                               |                  |
| route advertisements                      |                              | displaying.....                    | 3403             |
| stub areas and NSSAs, to control.....     | 4413                         | router functionality.....          | 4397             |
| route authentication                      |                              | router identifier.....             | 3369             |
| peering sessions.....                     | 3897                         | configuring.....                   | 4400             |
| route distinguisher.....                  | 3367                         | router link advertisements.....    | 4393             |
| route forwarding-table See MPLS           |                              | router-id statement.....           | 3369             |
| route injection.....                      | 3766, 4529, 4532, 4535, 4539 | routers                            |                  |
| route limit, configuring                  |                              | DNS name servers, configuring..... | 15, 26           |
| paths.....                                | 3326                         | domain names.....                  | 26               |
| prefix.....                               | 3328                         | domains to be searched.....        | 26, 223          |
| route preference                          |                              | login classes.....                 | 1823, 1848       |
| external routes.....                      | 4462                         | names                              |                  |
| internal routes.....                      | 4462                         | configuring.....                   | 34               |
| route preferences                         |                              | mapping to IP addresses.....       | 23, 34           |
| LDP.....                                  | 4828, 4871                   | NTP.....                           | 137              |
| signaled LSPs.....                        | 5081                         | physical system location.....      | 29               |
| route recording.....                      | 3368                         | ports                              |                  |
| route redistribution.....                 | 3766, 4529, 4532, 4535, 4539 | RADIUS servers.....                | 1873, 1893       |
| route reflectors See BGP route reflectors |                              | redirect .....                     | 18               |
| BGP.....                                  | 3876                         | root login, controlling.....       | 1413             |
| route selection                           |                              | source addresses.....              | 20, 222          |
| OSPF.....                                 | 4461                         | time zone setting.....             | 24               |
| preference.....                           | 4462                         | user accounts.....                 | 125, 1853        |
| static routes, defining.....              | 3246                         | routes                             |                  |
| route statement                           |                              | route preferences.....             | 4828, 4871, 5081 |
| aggregate statement.....                  | 3272                         | routes, displaying                 |                  |
| generate statement.....                   | 3309                         | active.....                        | 3406             |
| VRRP.....                                 | 2698                         | active path.....                   | 3412             |
| usage guidelines.....                     | 2646                         | all.....                           | 3417             |
| route summarization                       |                              | AS paths                           |                  |
| configuring.....                          | 4452                         | distribution of.....               | 3392             |
| route, displaying                         |                              | domain information.....            | 3396             |
| next-hop.....                             | 3510                         | regular expressions, matching..... | 3419             |
| route-distinguisher-id statement.....     | 3367                         | summary of.....                    | 3399             |
| route-flap damping.....                   | 3926                         | best.....                          | 3421             |
| parameters.....                           | 3927                         | brief information.....             | 3424             |
| route-monitoring (tracing flag)           |                              | community ASN.....                 | 3426             |
| BMP.....                                  | 4084                         | community name.....                | 3428             |
| route-monitoring statement.....           | 4068                         | damping, BGP.....                  | 3430, 4132       |
| usage guidelines.....                     | 3649                         | detailed information.....          | 3435, 4137       |
| route-record statement.....               | 3368                         | extensive information.....         | 3458, 6387       |
|                                           |                              | flow validation.....               | 3475             |

- in a prefix range.....3529
  - in a specific routing table.....3558, 5199
  - in the forwarding table.....3477, 5227
  - in the LDP internal topology table.....4906
  - inactive path.....3491
  - inactive prefix.....3494
  - instances.....2016, 3496
  - learned from a specific address.....3548
  - learned from a specific protocol.....3519, 6404
  - learned from snooping.....3545
  - LSP.....3506
  - martian.....3508
  - matching the specified address.....3453
  - MPLS labels.....3503
  - next-hop resolution.....3542
  - not associated with a community.....3516
  - policy-based route export.....3455
  - received through a neighbor.....3533
  - summary statistics.....3554
  - terse.....3586
  - to specified network host.....1269
- routing
  - in one AS with RIP.....4685
  - RIP See RIP
- Routing Engine
  - software component.....1793
- Routing Engine switchover effects
  - comparison of high availability features.....2622
- Routing Engines
  - connections, displaying.....1060
  - environmental information, displaying.....575
  - operation of, controlling.....376
  - status, displaying.....971
  - switchover information, displaying.....2738
- Routing Information Protocol See RIP
- Routing Information Protocol next generation See RIPng
- routing instances
  - router identifier.....3367
- routing policies
  - applying.....3765
  - BGP export, for CLNS.....4237
  - configuration
    - tasks.....3766, 3927, 4529, 4532, 4535, 4539
  - export statement.....3765
  - import statement.....3765
  - injecting routes from one protocol into
    - another.....3766, 4529
  - OSPF import policy.....4535, 4539
  - redistributing static routes into OSPF.....4532
  - reducing update messages with flap
    - damping.....3926
  - RIP routing policy (configuration
    - editor).....4685
  - route
    - redistribution.....3766, 4529, 4532, 4535, 4539
  - route-flap damping.....3926
- routing policy
  - actions.....4529
  - configuring on logical systems.....4232
  - match conditions.....4528
  - terms.....4528
- routing protocol software process.....1794
- routing solutions
  - BGP confederations, for scaling
    - problems.....3891
  - BGP route reflectors, for scaling
    - problems.....3876
  - controlling RIP traffic with the incoming
    - metric.....4726
  - controlling RIP traffic with the outgoing
    - metric.....4728
  - multicast administrative scoping.....5394
  - multicast reverse-path forwarding
    - (RPF).....5393
  - multicast shortest-path tree (SPT).....5393
  - NSSAs, to control route advertisement.....4413
  - poison reverse, for traffic reduction.....4680
  - reducing update messages with flap
    - damping.....3926
  - split horizon, for traffic reduction.....4680
  - stub areas, to control route
    - advertisement.....4413
- routing table
  - sample distance-vector routing.....4678
  - updates, limitations in RIP.....4681
- routing tables
  - BGP, delays in exchanging routes.....3772
  - creating.....3361
  - group.....3297, 3314, 3364, 3365, 4609
  - host routes, installing.....5048
  - import policy.....3313
  - MSDP.....5782
  - nonactive routes, exchanging with
    - BGP.....3770, 3788, 3969
    - PIM.....5743
  - policy, routing.....3317



|                                                   |            |
|---------------------------------------------------|------------|
| routing-instance statement                        |            |
| tunnel.....                                       | 2966       |
| usage guidelines.....                             | 1873, 1893 |
| routing-options statement.....                    | 3371       |
| RP                                                |            |
| anycast.....                                      | 5693       |
| embedded.....                                     | 5706       |
| rp (tracing flag).....                            | 5756       |
| rp statement.....                                 | 5744       |
| rp-register-policy statement.....                 | 5746       |
| usage guidelines.....                             | 5577       |
| rp-set statement.....                             | 5747       |
| usage guidelines.....                             | 5552, 5610 |
| rp-d process.....                                 | 1794       |
| RPF                                               |            |
| calculations, displaying.....                     | 5869       |
| PIM source state, displaying.....                 | 5918       |
| RPF (reverse-path forwarding)                     |            |
| description.....                                  | 5393       |
| RPF check, multicast                              |            |
| RPF policy.....                                   | 3360       |
| rpf-check-policy statement.....                   | 3360       |
| rpf-selection statement                           |            |
| PIM.....                                          | 5748       |
| RPs                                               |            |
| displaying.....                                   | 5911       |
| maximum.....                                      | 5726       |
| RPT.....                                          | 5579       |
| rstp statement.....                               | 2356       |
| RSTP, nonstop bridging.....                       | 2592       |
| RSVP                                              |            |
| aggregation.....                                  | 5306       |
| authentication.....                               | 5305       |
| bandwidth                                         |            |
| reserving.....                                    | 5334       |
| hello interval.....                               | 5315       |
| interfaces, displaying.....                       | 5355       |
| keep multiplier.....                              | 5319       |
| LSP connections                                   |            |
| operability, checking.....                        | 5347       |
| LSPs, real-time status.....                       | 5109       |
| MTU signaling in.....                             | 5264       |
| multipath extensions.....                         | 5272       |
| multipath implementation.....                     | 5273       |
| neighbors, displaying.....                        | 5360       |
| overview.....                                     | 5263       |
| preferences.....                                  | 3786       |
| reservation request and data flow,<br>figure..... | 5263       |

|                                     |                              |
|-------------------------------------|------------------------------|
| sessions                            |                              |
| clearing.....                       | 5344                         |
| displaying.....                     | 5365                         |
| signaled LSPs.....                  | 4971                         |
| soft state.....                     | 5263                         |
| statistics                          |                              |
| clearing.....                       | 5346                         |
| displaying.....                     | 5375                         |
| timers.....                         | 5332                         |
| tracing LSPs.....                   | 5382                         |
| tracing protocol traffic.....       | 5338                         |
| tunneling LDP LSPs through RSVP     |                              |
| LSPs.....                           | 4805, 4834, 4866, 5053, 5264 |
| version, displaying.....            | 5379                         |
| RSVP LSP metrics                    |                              |
| ignoring.....                       | 4511                         |
| RSVP LSPs                           |                              |
| ping interval.....                  | 5347                         |
| rsvp-error-hold-time statement..... | 5087                         |
| run command                         |                              |
| usage guidelines.....               | 66                           |

## S

|                                          |                |
|------------------------------------------|----------------|
| sample-rate statement.....               | 1595           |
| sample-type statement.....               | 1656           |
| Sampling.....                            | 5291           |
| save command.....                        | 425            |
| usage guidelines.....                    | 66, 1331, 1343 |
| saved-core-context statement.....        | 236            |
| usage guidelines.....                    | 29             |
| saved-core-files statement.....          | 236            |
| usage guidelines.....                    | 29             |
| saving licenses (CLI).....               | 119            |
| scan-interval.....                       | 2656           |
| SCB                                      |                |
| firmware version, displaying.....        | 692            |
| scc-master option to host statement..... | 1699           |
| scheduler-map statement.....             | 7111           |
| scheduler-maps statement.....            | 7112           |
| schedulers statement.....                | 7113           |
| scope statement.....                     | 3371           |
| scope-policy statement.....              | 3372           |
| scoping, administrative.....             | 5394           |
| scoping, multicast.....                  | 3371           |
| with scope policy.....                   | 3372           |
| scripts statement.....                   | 1587           |
| secondary                                |                |
| LSPs.....                                | 5088           |



- 
- secondary statement.....5088
    - usage guidelines.....4423, 4424
  - secret statement.....6312
    - authentication.....318
    - usage guidelines, RADIUS.....1873, 1893
  - security
    - MD5 authentication for RIPv2.....4699
    - password authentication for RIPv2.....4698
  - security-level statement
    - for access privileges.....1657
    - for SNMP notifications.....1658
  - security-model statement
    - for access privileges.....1659
    - for groups.....1660
    - for SNMP notifications.....1660
  - security-name statement
    - for community string.....1661
    - for security group.....1662
    - for SNMP notifications.....1663
  - security-to-group statement.....1664
  - select statement.....5089
  - send (tracing flag modifier)
    - LDP.....4878
    - RSVP.....5339
  - send statement
    - BGP
      - usage guidelines.....3844
    - RIP.....4783
      - usage guidelines.....4735
  - serial number
    - routing platform.....166
  - serial numbers, displaying.....743
  - server mode, usage guidelines.....135
  - server statement
    - NTP.....290
    - RADIUS accounting.....319
    - TACPLUS+.....319
    - usage guidelines.....133
  - server-address statement.....255
  - service-id statement.....2303
  - service-name statement.....1982
  - services statement.....1970
  - Session and Resource Control.....1131
  - session messages, LDP.....4808
  - session protection, LDP
    - configuration.....4836
    - overview.....4808
  - session statement.....4874
    - usage guidelines.....4835
  - session-mode statement
    - BGP.....4070
  - session-protection statement.....4875
    - usage guidelines.....4836
  - set command
    - usage guidelines.....66
  - set date ntp.....136
  - set option.....1327
  - setting date and time (CLI).....136
  - setup-protection statement.....5333
  - sFlow.....1453, 1457
  - sflow statement.....1596
  - sFlow technology
    - tracing operations.....1598
  - SFM
    - firmware version, displaying.....692
  - sha-256 checksum, calculating.....346, 1374
  - shaping-rate statement.....7114
  - shared trees.....5579
  - SHA-1 checksum, calculating.....345, 1373
  - shortest path first.....4373
  - Shortest Path First See SPF algorithm
  - shortest-path tree.....5393
  - shortest-path trees.....5584
    - See also SPT
  - show (ospf | ospf3) backup coverage
    - command.....4628
  - show (ospf | ospf3) backup neighbor.....4631
  - show (ospf | ospf3) interface command.....4643
  - show (ospf | ospf3) io-statistics command.....4649
  - show (ospf | ospf3) log command.....4651
  - show (ospf | ospf3) neighbor command.....4654
  - show (ospf | ospf3) overview
    - command.....2731, 4660
  - show (ospf | ospf3) route command.....4665
  - show (ospf | ospf3) statistics command.....4671
  - show (tracing flag)
    - class of service.....7059
  - show analyzer command.....6320
  - show app-engine crash command.....1780
  - show app-engine logs command.....1782
  - show as-path command.....3392
  - show as-path domain command.....3396
  - show as-path summary command.....3399
  - show bgp bmp command.....4098
  - show bgp group command.....4100
  - show bgp neighbor
    - command.....2710, 3894, 3959, 4107
    - explanation.....3895

|                                                                             |                  |
|-----------------------------------------------------------------------------|------------------|
| show bgp summary command.....                                               | 3896, 3959, 4124 |
| explanation.....                                                            | 3896             |
| show bridge mac-table command.....                                          | 2558             |
| show chassis alarms command.....                                            | 459              |
| show chassis environment command.....                                       | 476              |
| show chassis environment fpc command.....                                   | 540              |
| show chassis environment pem command.....                                   | 566              |
| show chassis environment routing-engine<br>command.....                     | 575              |
| show chassis fabric errors command.....                                     | 580              |
| show chassis fabric fpcs command.....                                       | 585              |
| show chassis fabric sibs command.....                                       | 629              |
| show chassis fabric topology command.....                                   | 647              |
| show chassis fan command.....                                               | 678              |
| show chassis firmware command.....                                          | 692              |
| show chassis fpc command.....                                               | 703              |
| show chassis hardware command.....                                          | 743              |
| show chassis location command.....                                          | 945              |
| show chassis mac-addresses command.....                                     | 949              |
| show chassis pic command.....                                               | 954              |
| show chassis routing-engine command.....                                    | 970              |
| show class-of-service classifier<br>command.....                            | 7231, 7238       |
| show class-of-service code-point-aliases<br>command.....                    | 7236, 7241       |
| show class-of-service command.....                                          | 7163             |
| show class-of-service congestion-notification<br>command.....               | 7394             |
| show class-of-service drop-profile<br>command.....                          | 7293             |
| show class-of-service forwarding-class<br>command.....                      | 7232, 7243       |
| show class-of-service forwarding-class-set<br>command.....                  | 7245             |
| show class-of-service forwarding-table classifier<br>command.....           | 7251             |
| show class-of-service forwarding-table classifier<br>mapping command.....   | 7253             |
| show class-of-service forwarding-table<br>command.....                      | 7247, 7296       |
| show class-of-service forwarding-table drop-profile<br>command.....         | 7300             |
| show class-of-service forwarding-table rewrite-rule<br>command.....         | 7255             |
| show class-of-service forwarding-table rewrite-rule<br>mapping command..... | 7257             |
| show class-of-service forwarding-table<br>scheduler-map command.....        | 7302             |
| show class-of-service interface<br>command.....                             | 7167, 7258, 7304 |
| show class-of-service multi-destination<br>command.....                     | 7287             |
| show class-of-service rewrite-rule<br>command.....                          | 7288             |
| show class-of-service scheduler-map<br>command.....                         | 7333             |
| show class-of-service shared-buffer<br>command.....                         | 7196             |
| show class-of-service traffic-control-profile<br>command.....               | 7336             |
| show cli authorization command.....                                         | 430              |
| show cli command.....                                                       | 428              |
| show cli directory command.....                                             | 434              |
| show cli history command.....                                               | 435              |
| show command<br>usage guidelines.....                                       | 66               |
| show dcbx command.....                                                      | 6580, 7397       |
| show dcbx neighbors command.....                                            | 6581, 7398       |
| show ddos-protection protocols command.....                                 | 6182             |
| show ddos-protection protocols parameters<br>command.....                   | 6201             |
| show ddos-protection protocols statistics<br>command.....                   | 6208             |
| show ddos-protection protocols violations<br>command.....                   | 6221             |
| show ddos-protection statistics command.....                                | 6218             |
| show ddos-protection version command.....                                   | 6220             |
| show dhcp relay binding command.....                                        | 6370             |
| show dhcp relay statistics command.....                                     | 6375             |
| show dhcp server binding command.....                                       | 6341             |
| show dhcp server statistics command.....                                    | 6347             |
| show dhcpv6 relay binding command.....                                      | 6378             |
| show dhcpv6 relay statistics command.....                                   | 6384             |
| show dhcpv6 server binding command.....                                     | 6350             |
| show dhcpv6 server statistics command.....                                  | 6356             |
| show ethernet-switching interfaces<br>command.....                          | 1998, 2382       |
| show ethernet-switching mac-learning-log<br>command.....                    | 2406             |
| show ethernet-switching statistics aging<br>command.....                    | 2410             |
| show ethernet-switching statistics mac-learning<br>command.....             | 2412             |
| show ethernet-switching table command.....                                  | 2386             |
| show fip vlan-discovery command.....                                        | 6577             |
| show firewall command.....                                                  | 6167             |
| show firewall policer command.....                                          | 6171             |

|                                              |            |                                           |                              |
|----------------------------------------------|------------|-------------------------------------------|------------------------------|
| show host command.....                       | 1002       | show link-management routing command..... | 5133                         |
| show igmp group command.....                 | 5795       | show link-management statistics           |                              |
| show igmp interface command.....             | 5801       | command.....                              | 5136, 5156                   |
| show igmp statistics command.....            | 5805       | show link-management te-link              |                              |
| show igmp-snooping membership                |            | command.....                              | 5138, 5158                   |
| command.....                                 | 5816       | show lldp command.....                    | 2002                         |
| show igmp-snooping route command.....        | 5819       | show lldp local-info command.....         | 2007                         |
| show igmp-snooping statistics command.....   | 5821       | show lldp neighbors command.....          | 2009                         |
| show igmp-snooping vlans command.....        | 5823       | show lldp statistics command.....         | 2014                         |
| show interfaces (10-Gigabit Ethernet)        |            | show log command.....                     | 1003, 1535, 1784, 2727, 4241 |
| command.....                                 | 3185       | show log messages command .....           | 1555                         |
| show interfaces (GRE) command.....           | 3131       | show mpls call-admission-control          |                              |
| show interfaces diagnostics optics           |            | command.....                              | 5140, 5160                   |
| command.....                                 | 3102       | show mpls cspf command.....               | 5142, 5162                   |
| show interfaces filters command.....         | 6173       | show mpls diffserv-te command.....        | 5144, 5164                   |
| show interfaces ge- (Gigabit Ethernet)       |            | show mpls interface command.....          | 5154, 5174                   |
| command.....                                 | 3116       | show mpls lsp autobandwidth command.....  | 5194                         |
| show interfaces irb command.....             | 3140       | show mpls lsp command.....                | 5176                         |
| show interfaces queue command.....           | 3146, 7340 | show mpls path command.....               | 5197                         |
| show ipv6 neighbors command.....             | 3401       | show mpls static-lsp command.....         | 5241                         |
| show ipv6 router-advertisement command.....  | 3403       | show msdp command.....                    | 5934                         |
| show isis adjacency brief command.....       | 4173       | show msdp source command.....             | 5936                         |
| show isis adjacency command.....             | 4338       | show msdp source-active command.....      | 5938                         |
| show isis adjacency extensive command.....   | 4174       | show msdp statistics command.....         | 5941                         |
| explanation.....                             | 4174       | show multicast flow-map command.....      | 5850                         |
| show isis authentication command.....        | 4342       | show multicast interface command.....     | 5852                         |
| show isis backup coverage.....               | 4344       | show multicast minfo command.....         | 5854                         |
| show isis backup spf results.....            | 4348       | show multicast next-hops command.....     | 5856                         |
| show isis database command.....              | 4352       | show multicast pim-to-igmp-proxy          |                              |
| show isis hostname command.....              | 4359       | command.....                              | 5859                         |
| show isis interface command.....             | 4361       | show multicast pim-to-mld-proxy           |                              |
| show isis overview command.....              | 4365       | command.....                              | 5861                         |
| show isis route command.....                 | 4369       | show multicast route command.....         | 5863                         |
| show isis spf command.....                   | 4373       | show multicast rpf command.....           | 5869                         |
| show isis statistics command.....            | 4378       | show multicast scope command.....         | 5873                         |
| show isis-backup label-switched-path         |            | show multicast sessions command.....      | 5875                         |
| command.....                                 | 4346       | show multicast usage command.....         | 5878                         |
| show lacp statistics interfaces command..... | 3209       | show ntp associations command.....        | 450                          |
| show ldp database command.....               | 4890       | show ntp status command.....              | 452                          |
| show ldp fec-filters command.....            | 4899       | show oam ethernet link-fault-management   |                              |
| show ldp interface command.....              | 4900       | command.....                              | 3084                         |
| show ldp neighbor command.....               | 4902       | show ospf context-identifier command..... | 4633                         |
| show ldp path command.....                   | 4904       | show ospf database command.....           | 4635                         |
| show ldp route command.....                  | 4906       | show ospf interfaces command.....         | 4567                         |
| show ldp session command.....                | 4910       | show ospf neighbors command.....          | 4567                         |
| show ldp statistics command.....             | 4916       | show ospf statistics command.....         | 4567                         |
| show ldp traffic-statistics command.....     | 4920       | show pfe filter command.....              | 6175, 7199                   |
| show link-management command.....            | 5127       | show pfe next-hop command.....            | 7201                         |
| show link-management peer command.....       | 5131       | show pfe route command.....               | 7206                         |

|                                             |                        |                                            |                        |
|---------------------------------------------|------------------------|--------------------------------------------|------------------------|
| show pfe terse command.....                 | 7215                   | show rsvp interface command.....           | 5355                   |
| show pfe version command.....               | 7217                   | show rsvp neighbor command.....            | 5360                   |
| show pim bootstrap command.....             | 5881                   | show rsvp session command.....             | 5365                   |
| show pim interfaces command.....            | 5883                   | show rsvp statistics command.....          | 5375                   |
| show pim join command.....                  | 5886                   | show rsvp version command.....             | 5379                   |
| show pim neighbors command.....             | 5907                   | show security keychain command.....        | 5124                   |
| show pim rps command.....                   | 5911                   | show sflow collector command.....          | 1733                   |
| show pim source command.....                | 5918                   | show sflow command.....                    | 1731                   |
| show pim statistics command.....            | 5921                   | show sflow interface command.....          | 1734                   |
| show policy damping command.....            | 4130                   | show snmp health-monitor command.....      | 1749                   |
| show protocols igmp command.....            | 5799                   | show snmp inform-statistics command.....   | 1754                   |
| show ptp global-information command.....    | 456                    | show snmp mib command.....                 | 1516, 1756, 1796       |
| show redundant-trunk-group command.....     | 3212                   | show snmp rmon command.....                | 1759                   |
| show rip general-statistics command.....    | 4792                   | show snmp rmon history command.....        | 1763                   |
| show rip neighbor command.....              | 4794                   | show snmp statistics                       |                        |
| show rip statistics command.....            | 4796                   | command.....                               | 1516, 1764, 1796, 2020 |
| show route active-path command.....         | 3412                   | show snmp v3 command.....                  | 1772                   |
| show route all command.....                 | 3417                   | show spanning-tree bridge command.....     | 2420                   |
| show route aspath-regex command.....        | 3419                   | show spanning-tree interface command.....  | 2425                   |
| show route best command.....                | 3421                   | show spanning-tree mstp configuration      |                        |
| show route brief command.....               | 3424                   | command.....                               | 2431                   |
| show route command.....                     | 3406                   | show spanning-tree statistics command..... | 2433                   |
| show route community command.....           | 3426                   | show subscribers command.....              | 1007                   |
| show route community-name command.....      | 3428                   | show system alarms command.....            | 1027                   |
| show route damping command.....             | 3430, 4132             | show system audit command.....             | 1030                   |
| show route detail command.....              | 3435, 4137             | show system autoinstallation status        |                        |
| show route exact command.....               | 3453                   | command.....                               | 324                    |
| show route export command.....              | 3455                   | show system boot-messages command.....     | 1038                   |
| show route extensive command.....           | 3458, 6387             | show system buffers command.....           | 1045                   |
| show route flow validation command.....     | 3475                   | show system certificate command.....       | 1052                   |
| show route forwarding-table                 |                        | show system commit command.....            | 1054, 1388             |
| command.....                                | 3477, 5146, 5166, 5227 | show system configuration archival         |                        |
| show route inactive-path command.....       | 3491                   | command.....                               | 1057, 1391             |
| show route inactive-prefix command.....     | 3494                   | show system configuration rescue           |                        |
| show route instance command.....            | 2016, 3496             | command.....                               | 1058, 1392             |
| show route label command.....               | 3503                   | show system connections command.....       | 1060                   |
| show route label-switched-path command..... | 3506                   | show system core-dumps command.....        | 1079                   |
| show route martians command.....            | 3508                   | show system directory-usage command.....   | 1093                   |
| show route next-hop command.....            | 3510                   | show system license command.....           | 120, 441               |
| show route no-community command.....        | 3516                   | explanation.....                           | 120                    |
| show route protocol command.....            | 3519, 6404             | show system license usage command.....     | 121                    |
| show route range command.....               | 3529                   | explanation.....                           | 121                    |
| show route receive-protocol command.....    | 3533                   | show system processes command.....         | 165, 1097              |
| show route resolution command.....          | 3542                   | show system reboot command.....            | 1126                   |
| show route snooping command.....            | 3545                   | show system resource-cleanup processes     |                        |
| show route source-gateway command.....      | 3548                   | command.....                               | 1129                   |
| show route summary command.....             | 3554                   | show system rollback command.....          | 1302, 1394             |
| show route table command.....               | 3558, 5199             | show system services service-deployment    |                        |
| show route terse command.....               | 3586                   | command.....                               | 1131                   |

- show system snapshot command (FreeBSD upgrade).....267
- show system software command.....1132
- show system statistics arp command.....2394
- show system statistics
  - command 336, 381, 382, 1054, 1129, 1140, 1185, 1266, 1268, 1369, 1386, 1388, 1396
- show system statistics igmp command.....5808
- show system storage command.....1177
- show system switchover command.....2738
- show system uptime command.....1185
- show system users command.....1190
- show system virtual-memory command.....1195
- show task replication command.....2744, 2748
- show ted database command.....5244
- show ted link command.....5252
- show ted protocol command.....5255
- show version command.....1253
- show vlans command.....2395
- show vpls mac-table command.....2551, 2563
- show bfd session extensive command.....4253
- show isis interface brief command.....4172
- show isis interface detail command.....4172
  - explanation.....4173
- SIB
  - switch fabric status, displaying.....629
  - switch fabric topology, displaying.....647
- signal-bandwidth statement.....5089
- signaled LSPs.....5073
  - adaptive rerouting.....5024
  - associating addresses.....5048
  - configuration statements.....5051
  - constrained-path computation
    - disabling.....5060
  - egress router address.....5101
  - fast reroute.....5031, 5042
  - fate-sharing.....3300
  - hold time.....5027
  - ingress router address.....5043
  - metrics.....5057
  - path
    - connection retry information.....5085
    - length.....5045, 5317
    - preference levels.....5081
    - primary .....5082
    - reoptimization.....5068, 5071
    - secondary.....5088
    - standby state.....5091
    - TTL decrementing.....5061, 5063
- single-area network, OSPF.....4406
- single-connection statement.....320
- single-rate statement.....6114
- size statement.....1702
  - archiving of all system log files.....1692
  - archiving of individual system log file.....1694
  - system logging
    - usage guidelines.....1543
- smart-optimize-timer statement.....5090
- snapshots
  - FreeBSD upgrade.....95
- SNMP
  - adding client lists and prefix lists.....1505
  - community string.....1503
  - configuration
    - versions 1 and 2.....1499
  - health monitor alarms, displaying.....1749
  - inform statistics, displaying.....1754
  - limiting interface access.....1506
  - MIB object values, displaying.....1756
  - MIB views.....1506
  - RMON alarms and events, displaying.....1759
  - show commands.....1516, 1796
  - statistics
    - clearing.....1739
    - displaying.....1764, 2020
  - system location.....1636, 1940
  - tracing operations.....1518
  - trap groups.....1504
  - v3.....1460
  - version 3 configuration, displaying.....1772
- snmp
  - rmon history.....1763
- snmp (tracing flag)
  - class of service.....7059
- snmp history
  - clear.....1738
- SNMP inform statistics, displaying.....1754
- SNMP informs.....1515
- snmp statement.....1665, 1971
  - usage guidelines
    - SNMPv1 and SNMPv2.....1499
- SNMP traps
  - MPLS.....5055
  - spoofing.....1741
- snmp-community statement.....1669
- SNMPv2
  - Passive Monitoring Traps MIB.....1504

|                                            |            |                                           |                 |
|--------------------------------------------|------------|-------------------------------------------|-----------------|
| SNMPv3.....                                | 1460       | source-increment statement                |                 |
| informs, configuring.....                  | 1515       | IGMP.....                                 | 5637            |
| minimum configuration.....                 | 1461       | usage guidelines.....                     | 5418            |
| snooping                                   |            | MLD                                       |                 |
| IGMP and VLANs.....                        | 5435       | usage guidelines.....                     | 5454            |
| multicast.....                             | 5395       | source-ip statement.....                  | 1597            |
| snooping routes, displaying.....           | 3545       | source-port statement.....                | 237             |
| sockets, displaying active IP.....         | 1060       | usage guidelines.....                     | 19              |
| soft state, RSVP .....                     | 5263       | source-routing statement.....             | 3374            |
| soft-preemption statement                  |            | source-specific multicast <i>See</i> SSM  |                 |
| RSVP.....                                  | 5335       | Spanning Tree Protocol <i>See</i> STP     |                 |
| software.....                              | 1793       | spanning-tree protocols                   |                 |
| version, displaying.....                   | 166        | tracing protocol operations               |                 |
| <i>See also</i> JUNOS software             |            | traceoptions statement.....               | 2357            |
| software package names                     |            | speed statement.....                      | 3029            |
| FreeBSD upgrade.....                       | 92         | chassis.....                              | 2979            |
| source filtering.....                      | 5408       | spf (tracing flag)                        |                 |
| source gateway addresses, displaying.....  | 3548       | IS-IS.....                                | 4321            |
| source statement                           |            | OSPF.....                                 | 4611            |
| commit scripts.....                        | 1589       | SPF algorithm                             |                 |
| IGMP.....                                  | 5635       | overview.....                             | 4388            |
| usage guidelines.....                      | 5418       | SPF calculations, displaying.....         | 4651            |
| MLD                                        |            | spf-options statement                     |                 |
| usage guidelines.....                      | 5454       | IS-IS.....                                | 4319            |
| MSDP.....                                  | 5783       | split horizon technique.....              | 4680            |
| op scripts.....                            | 1589       | splitting-merging.....                    | 5336            |
| PIM RPF selection.....                     | 5731, 5749 | SPT.....                                  | 5584            |
| SSM.....                                   | 3373       | configuring threshold cutover policy..... | 5590            |
| usage guidelines.....                      | 5535       | cutover control.....                      | 5588            |
| tunnel.....                                | 2966, 3028 | SPT (shortest-path tree).....             | 5393            |
| source-active (tracing flag).....          | 5785       | spt-threshold statement.....              | 5750            |
| source-active-request (tracing flag).....  | 5785       | usage guidelines.....                     | 5590            |
| source-active-response (tracing flag)..... | 5785       | SRC client information, displaying.....   | 1131            |
| source-address statement.....              | 1669       | SSB                                       |                 |
| IGMP snooping.....                         | 5657, 5680 | firmware version, displaying.....         | 692             |
| NTP                                        |            | ssh command.....                          | 2028            |
| usage guidelines.....                      | 138        | usage guidelines.....                     | 59              |
| NTP, RADIUS, system logging, TACACS+.....  | 291        | SSH key files.....                        | 145, 1824, 1850 |
| RADIUS                                     |            | SSH service                               |                 |
| usage guidelines.....                      | 1876, 1896 | configuring.....                          | 1412            |
| RADIUS and TACACS+.....                    | 291        | root login.....                           | 1413            |
| system logging.....                        | 291        | SSH protocol version.....                 | 1413            |
| usage guidelines                           |            | ssh statement.....                        | 1570, 1975      |
| usage guidelines, RADIUS.....              | 1873, 1893 | usage guidelines.....                     | 1412            |
| source-count statement                     |            | SSH, opening a connection.....            | 2028            |
| IGMP.....                                  | 5636       | ssh-known-hosts statement.....            | 6314            |
| usage guidelines.....                      | 5418       | ssh-rsa statement.....                    | 308             |
| MLD                                        |            |                                           |                 |
| usage guidelines.....                      | 5454       |                                           |                 |

- 
- SSM.....5530, 5537
    - configuring.....5533
    - domains.....5534
    - mapping.....5535
  - SSM maps.....5541
    - example.....5541
  - SSM maps for different groups to different
    - sources.....5541
  - ssm-groups statement.....3375, 5763
    - usage guidelines.....5537
  - ssm-map statement
    - IGMP.....5764
      - usage guidelines.....5535
    - MLD
      - usage guidelines.....5535
    - SSM.....3376, 5765
      - usage guidelines.....5535
  - ssm-map-policy statement
    - IGMP interface.....5766
  - stale-routes-time statement.....2672, 4071
    - usage guidelines.....2581
  - standards documents
    - SNMP and MIBs.....1469
  - standby state, signaled LSPs.....5091
  - standby statement.....5091
  - start shell command.....1266
  - start-time statement.....6315
    - system log file archiving.....1694
    - system logging
      - usage guidelines.....1543
  - startup messages, displaying.....1038
  - startup period, VRRP.....2643
  - startup-alarm statement.....1670
  - startup-silent-period statement.....2699
    - usage guidelines.....2643
  - state
    - BMP.....4084
  - state (tracing flag)
    - LDP.....4878
    - MPLS.....5103
    - routing protocols.....2679, 3384
    - RSVP.....5339
  - state-machine-variables (tracing flag)
    - spanning-tree protocols.....2358
  - stateless firewall filters
    - accepting Routing Engine traffic from trusted
      - sources
        - example: blocking TCP access.....3904
        - example: blocking Telnet and SSH
          - access.....3909
      - examples
        - blocking TCP access.....3904
        - blocking Telnet and SSH access.....3909
  - statement-name statement.....2287
  - static (tracing flag).....5103
  - static LSPs
    - egress routers.....5049
    - intermediate routers.....5049
    - MPLS, displaying.....5241
    - overview.....4971
    - static LSP metric.....5057
  - static routes.....3377
    - BFD.....3237, 3281
    - defining route selection.....3246
    - graceful restart.....2579
    - preferences.....3786
  - static statement.....2305, 3377
    - IGMP.....5638
      - usage guidelines.....5418
    - IGMP snooping.....5645, 5659, 5667, 5682
    - MLD
      - usage guidelines.....5454
    - PIM.....5751
      - usage guidelines.....5547
  - static-host-mapping statement.....238
  - static-label-switched-path statement
    - static LSP.....5092
  - static-mac statement.....2305
  - station statement.....4072
    - usage guidelines.....3647
  - station-address statement.....4073
    - usage guidelines.....3647
  - station-port statement.....4074
    - usage guidelines.....3648
  - statistics
    - interfaces, real-time.....3090
    - MPLS traffic.....4980, 5094
    - output file.....4980
    - protocol-related, displaying.....1140
  - statistics (tracing flag)
    - BMP.....4084
  - statistics statement.....5094
    - usage guidelines.....4980



|                                              |      |                                               |                  |
|----------------------------------------------|------|-----------------------------------------------|------------------|
| statistics-timeout statement.....            | 4075 | switch fabric                                 |                  |
| usage guidelines.....                        | 3649 | errors, displaying.....                       | 580              |
| status command                               |      | status, displaying.....                       | 585, 629         |
| usage guidelines.....                        | 66   | switch fabric topology, displaying.....       | 647              |
| storage space, freeing.....                  | 399  | switches                                      |                  |
| storing previous configurations.....         | 1314 | user accounts.....                            | 1821, 1845       |
| storm-control statement.....                 | 6128 | switching                                     |                  |
| storm-control-profiles statement.....        | 6129 | configuring.....                              | 2120             |
| STP                                          |      | switching-type statement.....                 | 5096             |
| bridge                                       |      | switchover, displaying information about..... | 2738             |
| displaying.....                              | 2420 | symmetric active mode, NTP                    |                  |
| interface                                    |      | configuring.....                              | 134              |
| displaying.....                              | 2425 | defined.....                                  | 131, 133         |
| statistics                                   |      | synchronizing Routing Engines                 |                  |
| clearing.....                                | 2419 | graceful Routing Engine switchover.....       | 2627             |
| displaying.....                              | 2433 | syntax conventions.....                       | cxli             |
| STP, nonstop bridging.....                   | 2592 | syntax of configuration files, verifying..... | 1268, 1396       |
| strict-targeted-hellos statement.....        | 4875 | sysid statement.....                          | 238              |
| usage guidelines.....                        | 4810 | syslog.....                                   | 1559             |
| Structure of Management Information MIB..... | 1477 | syslog statement                              |                  |
| structured-data statement.....               | 1703 | routing options.....                          | 3333             |
| usage guidelines.....                        | 1538 | system processes.....                         | 1704             |
| stub areas                                   |      | usage guidelines.....                         | 1530             |
| configuring.....                             | 4414 | syslog-subtag statement.....                  | 1671             |
| description.....                             | 4413 | system authentication                         |                  |
| overview.....                                | 4398 | authentication order.....                     | 1859, 1869, 1884 |
| stub statement                               |      | RADIUS                                        |                  |
| usage guidelines.....                        | 4414 | configuring.....                              | 1873, 1893       |
| sub-ASs, BGP.....                            | 3890 | System Control Board See SCB                  |                  |
| subautonomous systems, BGP.....              | 3890 | system identification, displaying.....        | 166              |
| subscriber access                            |      | system identifier.....                        | 4158             |
| subscriber information, displaying.....      | 1007 | system location, SNMP.....                    | 1636, 1940       |
| subscriber-leave-timer statement.....        | 3379 | system log                                    |                  |
| subscribers                                  |      | logging operations.....                       | 1556             |
| displaying.....                              | 1007 | system log file                               |                  |
| subscribing to bandwidth.....                | 5334 | displaying.....                               | 1535             |
| subscription statement.....                  | 5334 | system log messages                           |                  |
| summaries statement                          |      | MPLS.....                                     | 5055             |
| usage guidelines.....                        | 4414 | routing protocol process.....                 | 3333             |
| summary LSA.....                             | 4393 | System Log MIB.....                           | 1477             |
| summary LSAs                                 |      | system logging                                |                  |
| advertising LSP metric.....                  | 4511 | disabling.....                                | 1534             |
| support, technical See technical support     |      | facilities                                    |                  |
| Supported features.....                      | 5298 | alternate for remote machine.....             | 1548             |
| swap statement                               |      | default for remote machine.....               | 1548             |
| Gigabit Ethernet IQ interfaces.....          | 2374 | for local machine.....                        | 1545, 1546       |
| MPLS.....                                    | 5095 | files, archiving.....                         | 1543             |



- messages, displaying
      - generated by Junos process.....1542
      - structured-data format.....1539
    - regular expression filtering.....1551
    - regular expression operators.....1552
    - single-chassis system.....1528
    - timestamp, modifying.....22, 1537
  - system login.....126
  - system overview
    - software.....1793
  - system process information, displaying.....165
  - system requirements
    - graceful Routing Engine switchover.....2623
    - nonstop active routing.....2600
    - nonstop bridging.....2591
  - system services
    - SSH.....1412
  - system statement.....239, 1976
  - system storage, displaying.....167
  - system time, displaying.....166
- T**
- TACACS+ authentication
    - overview.....1804, 1847
  - tacplus-options statement.....1982
  - tag statement.....1671, 1672, 3380
  - tag-list statement.....1672
  - tag-protocol-id statement
    - TPIDs expected to be sent or received.....3030
  - tags
    - aggregate routes.....3380
    - generated routes.....3380
    - static routes.....3380
  - target-address statement.....1673
  - target-parameters statement.....1674
  - targeted hello messages.....4807
  - targeted hello messages, LDP.....4854
  - targeted-hello statement.....4876
    - usage guidelines.....4826, 4827
  - targets statement.....1675, 1983
    - usage guidelines.....1504
  - task (tracing flag).....2679, 3384
    - BMP.....4084
  - TCC, graceful restart.....2579
  - tcp-mss statement.....4076
    - BGP
      - usage guidelines.....3912
  - te-class-matrix statement.....5097, 5100
  - te-metric statement
    - usage guidelines.....4518
  - technical support
    - contacting JTAC.....cxliii
  - TED See traffic engineering database
  - ted-export (tracing flag).....5103
  - ted-import (tracing flag).....5103
  - telnet command
    - usage guidelines.....59
  - telnet statement.....1571
  - term statement.....6096
  - terminal type.....15, 221, 1411
  - terms
    - routing policy.....4528
  - test configuration command.....1268, 1396
  - test msdp command.....5945
  - text message on craft interface
    - clearing.....333
  - then statement.....3339, 5076, 6097, 6115
  - three-color-policer statement.....6116
  - threshold
    - BFD.....4079
    - BGP.....4077
    - PIM.....5752, 5753
  - threshold statement
    - BFD (BGP)
      - usage guidelines.....3812
    - BGP.....3981
    - forwarding cache.....3381
    - IS-IS.....4266
    - MSDP.....5784
      - usage guidelines.....5604
    - RIP
      - usage guidelines.....4702
  - time
    - configuration example.....139
    - security configuration example.....142
  - time zone setting, routers.....24
  - time, displaying.....1185
  - time-format statement.....244, 1706
    - usage guidelines.....22, 1537
  - time-zone statement.....245
    - usage guidelines.....24
  - timeout statement.....321, 1675
    - authentication
      - usage guidelines, RADIUS.....1873, 1893
    - flow map.....3382
    - forwarding cache.....3383

|                                                             |            |                                   |                  |
|-------------------------------------------------------------|------------|-----------------------------------|------------------|
| timer (tracing flag).....                                   | 2679, 3384 | traceoptions statement.....       | 1676, 2958, 7058 |
| BMP.....                                                    | 4084       | BFD                               |                  |
| MPLS.....                                                   | 5103       | usage guidelines.....             | 3241             |
| timer, LDP.....                                             | 4825, 4854 | BGP.....                          | 4081             |
| timers                                                      |            | description.....                  | 3949             |
| OSPF.....                                                   | 4475       | BMP.....                          | 4084             |
| RSVP.....                                                   | 5332       | DDoS protection.....              | 6162             |
| timers (tracing flag)                                       |            | graceful restart                  |                  |
| spanning-tree protocols.....                                | 2358       | usage guidelines.....             | 2586             |
| timestamp-and-timezone statement.....                       | 1982       | IGMP.....                         | 5639             |
| to statement.....                                           | 3339, 5076 | usage guidelines.....             | 5427             |
| MPLS.....                                                   | 5101       | IGMP snooping.....                | 247, 5660, 5683  |
| tolerance.....                                              | 2656       | interfaces.....                   | 3032             |
| top command                                                 |            | IS-IS.....                        | 4321             |
| usage guidelines.....                                       | 66         | LDP.....                          | 4877             |
| topic1                                                      |            | usage guidelines.....             | 4817             |
| sub-topic 1994,1998,2002,2014,2382,2386,2395,2406,6320,7163 |            | LLDP.....                         | 1984             |
| topic2                                                      |            | MPLS.....                         | 5102             |
| sub-topic 1994,1998,2002,2014,2382,2386,2395,2406,6320,7163 |            | MSDP.....                         | 5785             |
| topologies statement.....                                   | 4320       | usage guidelines.....             | 5600             |
| topology                                                    |            | nonstop active routing.....       | 2679, 3384       |
| sample BGP confederations.....                              | 3891       | OSPF.....                         | 4569, 4611       |
| sample BGP MED use.....                                     | 3666       | PIM.....                          | 5755             |
| sample BGP peer session.....                                | 3601       | usage guidelines.....             | 5496             |
| sample BGP route reflector (one                             |            | RIP.....                          | 4757, 4758, 4784 |
| cluster).....                                               | 3874       | routing protocols.....            | 2679, 3384       |
| sample BGP route reflectors (cluster of                     |            | RSVP.....                         | 5338             |
| clusters).....                                              | 3875       | security.....                     | 6317             |
| sample BGP route reflectors (multiple                       |            | sFlow technology.....             | 1598             |
| clusters).....                                              | 3875       | SNMP                              |                  |
| sample distance-vector routing.....                         | 4678       | usage guidelines.....             | 1518             |
| sample multiarea OSPF routing.....                          | 4404       | spanning-tree protocols.....      | 2357             |
| sample OSPF network with stubs and                          |            | usage guidelines.....             | 3956             |
| NSSAs.....                                                  | 4413       | VRRP.....                         | 2700             |
| sample poison reverse routing.....                          | 4681       | traceroute command.....           | 1269             |
| sample split horizon routing.....                           | 4680       | traceroute monitor command.....   | 1273             |
| sample unidirectional routing.....                          | 4681       | traceroute mpls ldp.....          | 4924             |
| topology statement                                          |            | traceroute mpls ldp command.....  | 4924             |
| multitopology routing                                       |            | traceroute mpls rsvp.....         | 5382             |
| OSPF.....                                                   | 4610       | traceroute mpls rsvp command..... | 5382             |
| topology-change-state-machine (tracing flag)                |            | Traceroute page                   |                  |
| spanning-tree protocols.....                                | 2358       | results for RIP.....              | 4691             |
| totally stubby areas                                        |            | tracing.....                      | 1572             |
| configuring.....                                            | 4414       | destination-override.....         | 1572             |
| description.....                                            | 4413       | tracing flag modifiers            |                  |
| trace files                                                 |            | detail                            |                  |
| logical systems                                             |            | LDP.....                          | 4878             |
| .....                                                       | 3951       | RSVP.....                         | 5339             |
|                                                             |            | disable.....                      | 4878             |

|                                     |                              |  |
|-------------------------------------|------------------------------|--|
| receive                             |                              |  |
| LDP.....                            | 4878                         |  |
| RSVP.....                           | 5339                         |  |
| send                                |                              |  |
| LDP.....                            | 4878                         |  |
| RSVP.....                           | 5339                         |  |
| tracing flags                       |                              |  |
| address.....                        | 4877                         |  |
| all.....                            | 2357, 2679, 3384, 4084, 5102 |  |
| class of service.....               | 7058                         |  |
| RSVP.....                           | 5338                         |  |
| all-failures                        |                              |  |
| spanning-tree protocols.....        | 2357                         |  |
| as-path.....                        | 4081                         |  |
| assert.....                         | 5755                         |  |
| asynch                              |                              |  |
| class of service.....               | 7058                         |  |
| auth.....                           | 4784                         |  |
| automatic bandwidth.....            | 5102                         |  |
| binding.....                        | 4877                         |  |
| bootstrap.....                      | 5755                         |  |
| bpdu.....                           | 2357                         |  |
| bridge-detection-state-machine..... | 2357                         |  |
| cache, PIM.....                     | 5755                         |  |
| chassis-scheduler                   |                              |  |
| class of service.....               | 7058                         |  |
| config-internal.....                | 2679, 3384                   |  |
| connection.....                     | 5102                         |  |
| connection-detail.....              | 5102                         |  |
| cos-adjustment                      |                              |  |
| class of service.....               | 7058                         |  |
| csn.....                            | 4321                         |  |
| cspf.....                           | 5103                         |  |
| cspf-link.....                      | 5103                         |  |
| cspf-node.....                      | 5103                         |  |
| damping.....                        | 4081                         |  |
| down.....                           | 4084                         |  |
| dynamic                             |                              |  |
| class of service.....               | 7058                         |  |
| error.....                          | 4084                         |  |
| IS-IS.....                          | 4321                         |  |
| LDP.....                            | 4877                         |  |
| MPLS.....                           | 5103                         |  |
| OSPF.....                           | 4611                         |  |
| RIP.....                            | 4784                         |  |
| RSVP.....                           | 5338                         |  |
| event.....                          | 4084                         |  |
| LDP.....                            | 4877                         |  |
| RSVP.....                           | 5338                         |  |
| events                              |                              |  |
| spanning-tree protocols.....        | 2358                         |  |
| flash.....                          | 2679, 3384                   |  |
| flooding.....                       | 4611                         |  |
| general.....                        | 2679, 3384, 4084             |  |
| graceful restart                    |                              |  |
| IS-IS.....                          | 4321                         |  |
| OSPF.....                           | 4611                         |  |
| graceful-restart.....               | 5103                         |  |
| graft                               |                              |  |
| PIM.....                            | 5755                         |  |
| hardware-database                   |                              |  |
| class of service.....               | 7058                         |  |
| hello                               |                              |  |
| IS-IS.....                          | 4321                         |  |
| PIM.....                            | 5755                         |  |
| holddown.....                       | 4784                         |  |
| init                                |                              |  |
| class of service.....               | 7059                         |  |
| initialization.....                 | 4878                         |  |
| join.....                           | 5755                         |  |
| keepalive                           |                              |  |
| BGP.....                            | 4081                         |  |
| MSDP.....                           | 5785                         |  |
| kernel.....                         | 2679, 3384                   |  |
| label.....                          | 4878                         |  |
| leave                               |                              |  |
| IGMP.....                           | 5639                         |  |
| lmp.....                            | 5338                         |  |
| lsp.....                            | 4321                         |  |
| lsp-generation.....                 | 4321                         |  |
| lsp-history.....                    | 5103                         |  |
| lsping.....                         | 5103                         |  |
| mt.....                             | 5755                         |  |
| mtrace                              |                              |  |
| IGMP.....                           | 5427                         |  |
| normal.....                         | 2679, 3384, 4084             |  |
| notification.....                   | 4878                         |  |
| nsr-synchronization.....            | 5103, 5756                   |  |
| nsr-synchronization-detail.....     | 5103                         |  |
| packet-dump.....                    | 4611                         |  |
| packets.....                        | 4084                         |  |
| BGP.....                            | 4081                         |  |
| IGMP.....                           | 5640                         |  |
| IS-IS.....                          | 4321                         |  |
| LDP.....                            | 4878                         |  |
| OSPF.....                           | 4611                         |  |
| PIM.....                            | 5756                         |  |

|                                     |                  |
|-------------------------------------|------------------|
| RIP.....                            | 4784             |
| RSVP.....                           | 5338             |
| parse.....                          | 2679, 3384       |
| path                                |                  |
| LDP.....                            | 4878             |
| RSVP.....                           | 5338             |
| pathtear.....                       | 5338             |
| performance-monitor                 |                  |
| class of service.....               | 7059             |
| periodic .....                      | 4878             |
| policy.....                         | 2679, 3384, 4084 |
| port-information-state-machine..... | 2358             |
| port-migration-state-machine.....   | 2358             |
| port-receive-state-machine          |                  |
| spanning-tree protocols.....        | 2358             |
| port-role-select-state-machine      |                  |
| spanning-tree protocols.....        | 2358             |
| port-role-transit-state-machine     |                  |
| spanning-tree protocols.....        | 2358             |
| port-state-transit-state-machine    |                  |
| spanning-tree protocols.....        | 2358             |
| port-transmit-state-machine         |                  |
| spanning-tree protocols.....        | 2358             |
| ppmd                                |                  |
| spanning-tree protocols.....        | 2358             |
| process                             |                  |
| class of service.....               | 7059             |
| prune                               |                  |
| PIM.....                            | 5756             |
| psn.....                            | 4321             |
| regex-parse.....                    | 2679, 3384       |
| register.....                       | 5756             |
| report                              |                  |
| IGMP.....                           | 5640             |
| restart                             |                  |
| class of service.....               | 7059             |
| resv.....                           | 5339             |
| resvtear.....                       | 5339             |
| route.....                          | 4084, 5339       |
| LDP.....                            | 4878             |
| MSDP.....                           | 5785             |
| routing.....                        | 2679, 3384       |
| route-monitoring.....               | 4084             |
| route-socket                        |                  |
| class of service.....               | 7059             |
| rp.....                             | 5756             |
| show                                |                  |
| class of service.....               | 7059             |

|                               |                  |
|-------------------------------|------------------|
| snmp                          |                  |
| class of service.....         | 7059             |
| source-active.....            | 5785             |
| source-active-request.....    | 5785             |
| source-active-response.....   | 5785             |
| spf                           |                  |
| IS-IS.....                    | 4321             |
| OSPF.....                     | 4611             |
| state.....                    | 4084             |
| LDP.....                      | 4878             |
| MPLS.....                     | 5103             |
| routing protocols.....        | 2679, 3384       |
| RSVP.....                     | 5339             |
| state-machine-variables       |                  |
| spanning-tree protocols.....  | 2358             |
| static.....                   | 5103             |
| statistics.....               | 4084             |
| task.....                     | 2679, 3384, 4084 |
| timer.....                    | 2679, 3384, 4084 |
| MPLS.....                     | 5103             |
| timers                        |                  |
| spanning-tree protocols.....  | 2358             |
| topology-change-state-machine |                  |
| spanning-tree protocols.....  | 2358             |
| trigger.....                  | 4784             |
| up.....                       | 4084             |
| update                        |                  |
| RIP.....                      | 4784             |
| util                          |                  |
| class of service.....         | 7059             |
| write.....                    | 4084             |
| tracing IP multicast path     |                  |
| from receiver to source.....  | 5839             |
| from router to gateway.....   | 5847             |
| from server to router.....    | 5842             |
| tracing operations.....       | 1409             |
| BGP.....                      | 3949, 4081       |
| BMP.....                      | 4084             |
| IGMP.....                     | 5427, 5639       |
| IS-IS.....                    | 4321             |
| LDP.....                      | 4817, 4877       |
| MPLS.....                     | 5102             |
| MSDP.....                     | 5600, 5785       |
| OSPF.....                     | 4569, 4611       |
| PIM.....                      | 5755             |
| RIP.....                      | 4757, 4758, 4784 |
| routing protocols.....        | 2679, 3384       |
| RSVP.....                     | 5338             |

|                                             |            |                                              |                       |
|---------------------------------------------|------------|----------------------------------------------|-----------------------|
| SNMP.....                                   | 1518       | traffic-statistics statement.....            | 4880                  |
| STP.....                                    | 2357       | usage guidelines.....                        | 4815                  |
| tracing ospf                                |            | transfer interval                            |                       |
| configuring.....                            | 4570       | usage guidelines.....                        | 1334                  |
| tracing routes                              |            | transfer-interval statement                  |                       |
| from the receiver to the source.....        | 5839       | archiving of configuration.....              | 1364, 1986            |
| from the source to the gateway router.....  | 5847       | system log file archiving.....               | 1694                  |
| from the source to the receiver.....        | 5842       | system logging                               |                       |
| monitoring.....                             | 5845       | usage guidelines.....                        | 1543                  |
| track statement.....                        | 2702       | transfer-on-commit statement.....            | 1365, 1987            |
| usage guidelines.....                       | 2646       | usage guidelines.....                        | 1334                  |
| track-igp-metric statement.....             | 4879       | transferring router configuration to archive |                       |
| usage guidelines.....                       | 4833       | site.....                                    | 1333                  |
| traffic                                     |            | transient configuration changes              |                       |
| controlling with incoming RIP metric.....   | 4726       | compared to persistent changes.....          | 1430                  |
| controlling with outgoing RIP metric.....   | 4728       | overview.....                                | 1430                  |
| statistics.....                             | 4980, 5094 | transit areas                                |                       |
| traffic control                             |            | overview.....                                | 4399                  |
| OSPF.....                                   | 4461       | transit-delay statement                      |                       |
| traffic engineering                         |            | usage guidelines.....                        | 4476                  |
| advertising LSP metric in summary LSAs..... | 4511       | transit-lsp-association statement            |                       |
| database credibility value.....             | 4511       | MPLS-TP.....                                 | 5099, 5105            |
| enabling.....                               | 4513       | transitions                                  |                       |
| ignoring RSVP LSP metrics.....              | 4511       | advertising.....                             | 5027                  |
| igp shortcuts.....                          | 4511       | translational cross-connect See TCC          |                       |
| metric.....                                 | 4511, 4518 | transmit-interval                            |                       |
| ospf protocol preference.....               | 4511       | BFD.....                                     | 4086                  |
| support.....                                | 4511       | PIM.....                                     | 5754                  |
| Traffic engineering                         |            | transmit-interval statement                  |                       |
| challenges.....                             | 5274       | BFD.....                                     | 3281                  |
| traffic engineering database                |            | BGP.....                                     | 3981                  |
| database entries, displaying.....           | 5244       | IS-IS.....                                   | 4266                  |
| link information, displaying.....           | 5252       | transmit-rate statement.....                 | 7117                  |
| OSPF support.....                           | 4614       | transparent clock.....                       | 311                   |
| protocols learned from, displaying.....     | 5255       | transport-address statement.....             | 4882                  |
| traffic, real-time monitoring.....          | 1712       | usage guidelines.....                        | 4815                  |
| traffic-control-profiles statement.....     | 7116       | trap groups, SNMP.....                       | 1504                  |
| traffic-engineering statement               |            | trap-group statement.....                    | 1678, 1988            |
| IS-IS.....                                  | 4324       | usage guidelines.....                        | 1504                  |
| lsp-metric-info-summary                     |            | trap-options statement.....                  | 1679, 1989            |
| usage guidelines.....                       | 4513       | traps                                        |                       |
| MPLS.....                                   | 5098, 5104 | spoofing.....                                | 1741                  |
| OSPF.....                                   | 4614       | trigger (tracing flag).....                  | 4784                  |
| OSPF passive TE mode                        |            | Trivial Network Protocol family              |                       |
| usage guidelines.....                       | 4520       | interface addresses.....                     | 2824                  |
| shortcuts                                   |            | troubleshooting                              |                       |
| usage guidelines.....                       | 4513       | general.....                                 | 1816, 1838            |
| usage guidelines.....                       | 4513       | resources.....                               | 1814, 1836            |
|                                             |            | root password recovery.....                  | 147, 1810, 1818, 1840 |

|                                        |                              |                                            |                       |
|----------------------------------------|------------------------------|--------------------------------------------|-----------------------|
| trunk interface.....                   | 2290, 3006                   | upgrading software                         |                       |
| trunk port.....                        | 2290, 3006                   | performing.....                            | 1276, 1290            |
| trusted DHCP server                    |                              | upstream-interface statement.....          | 3387                  |
| overview.....                          | 6037                         | uptime, displaying.....                    | 1185                  |
| trusted port.....                      | 6037                         | URLs, specifying in commands.....          | 40                    |
| trusted-key statement.....             | 292                          | user (system logging facility).....        | 1545, 1547            |
| usage guidelines.....                  | 132                          | option to facility-override statement..... | 1549                  |
| TTL decrementing                       |                              | user access                                |                       |
| disabling.....                         | 5061, 5063                   | login classes.....                         | 1823, 1848            |
| TTY, displaying.....                   | 167                          | user accounts.....                         | 125, 1821, 1845, 1853 |
| tunnel statement.....                  | 2967, 2968                   | user accounts                              |                       |
| tunnel-services statement.....         | 5340                         | configuring.....                           | 125, 1821, 1845, 1853 |
| tunneling, MPLS                        |                              | user authentication                        |                       |
| RSVP LSPs.....                         | 4805, 4834, 4866, 5053, 5264 | methods for.....                           | 1804, 1847            |
| RSVP LSPs, heterogeneous networks..... | 4835                         | user data, erasing.....                    | 409                   |
| two-rate statement.....                | 6117                         | user identifiers See UIDs                  |                       |
| type statement.....                    | 1680                         | user permission flags.....                 | 1799                  |
| auxiliary port.....                    | 15, 1411                     | user statement                             |                       |
| console port.....                      | 15, 1411                     | access.....                                | 282, 309, 1990        |
|                                        |                              | usage guidelines.....                      | 125, 1821, 1845, 1853 |
|                                        |                              | SNMPv3.....                                | 1681                  |
|                                        |                              | system logging.....                        | 1707                  |
|                                        |                              | usage guidelines.....                      | 1533                  |
|                                        |                              | username                                   |                       |
|                                        |                              | displaying.....                            | 167                   |
|                                        |                              | users                                      |                       |
|                                        |                              | CLI permissions, displaying.....           | 430                   |
|                                        |                              | displaying.....                            | 167                   |
|                                        |                              | logged in, displaying.....                 | 1190                  |
|                                        |                              | logging users out.....                     | 390                   |
|                                        |                              | logs, displaying.....                      | 1003, 1784, 2727      |
|                                        |                              | messages, displaying for.....              | 381                   |
|                                        |                              | usm statement.....                         | 1682                  |
|                                        |                              | util (tracing flag)                        |                       |
|                                        |                              | class of service.....                      | 7059                  |
|                                        |                              | Utility MIB.....                           | 1460                  |
|                                        |                              |                                            |                       |
|                                        |                              | <b>V</b>                                   |                       |
|                                        |                              | v3 statement.....                          | 1684                  |
|                                        |                              | vacm statement.....                        | 1686                  |
|                                        |                              | validating software.....                   | 1297                  |
|                                        |                              | var/log/mib2d file.....                    | 1518                  |
|                                        |                              | var/log/snmpd file.....                    | 1518                  |
|                                        |                              | variable statement.....                    | 1687                  |
|                                        |                              | vendor-id statement.....                   | 6142                  |
|                                        |                              | verification                               |                       |
|                                        |                              | active licenses.....                       | 120                   |
|                                        |                              | BFD for IS-IS.....                         | 4248                  |
|                                        |                              | BGP session flap prevention.....           | 3925                  |
| U                                      |                              |                                            |                       |
| udp-port statement.....                | 1599                         |                                            |                       |
| uid statement.....                     | 249                          |                                            |                       |
| usage guidelines.....                  | 125, 1821, 1845, 1853        |                                            |                       |
| UIDs.....                              | 1822, 1846                   |                                            |                       |
| unit statement                         |                              |                                            |                       |
| class of service.....                  | 7089                         |                                            |                       |
| interfaces.....                        | 3035                         |                                            |                       |
| tunnel.....                            | 2969                         |                                            |                       |
| UNIX-level shell, creating.....        | 1266                         |                                            |                       |
| unknown-unicast-forwarding statement   |                              |                                            |                       |
| rate limiting.....                     | 6130                         |                                            |                       |
| unreachable-vtep-aging-timer.....      | 2528                         |                                            |                       |
| Unsupported features.....              | 5298                         |                                            |                       |
| up (tracing flag)                      |                              |                                            |                       |
| BMP.....                               | 4084                         |                                            |                       |
| up command                             |                              |                                            |                       |
| usage guidelines.....                  | 66                           |                                            |                       |
| update (tracing flag)                  |                              |                                            |                       |
| RIP.....                               | 4784                         |                                            |                       |
| update command                         |                              |                                            |                       |
| usage guidelines.....                  | 66                           |                                            |                       |
| update messages                        |                              |                                            |                       |
| BGP.....                               | 3597                         |                                            |                       |
| update-interval statement              |                              |                                            |                       |
| RIP.....                               | 4787                         |                                            |                       |
| usage guidelines.....                  | 4749, 4750                   |                                            |                       |
| update-server statement.....           | 256                          |                                            |                       |
| update-threshold statement.....        | 5340                         |                                            |                       |

- 
- IS-IS policy.....4235
  - license usage.....121
  - licenses .....120
  - multicast topology for IS-IS.....4208
  - network interfaces.....3757, 3864
  - RIP host reachability.....4691
  - static route.....3245, 3250, 3259
  - tracing.....3955
  - verifying syntax of configuration file.....1268, 1396
  - version
    - BFD.....4088
    - firmware, displaying.....692
    - RSVP, displaying.....5379
    - software, displaying.....166
    - general.....1253
  - version statement
    - BFD.....3281, 5758
    - usage guidelines.....3237
    - BFD (BGP)
      - usage guidelines.....3812
    - BGP.....3981
    - IGMP.....5641
    - usage guidelines.....5410
    - IS-IS.....4266
    - MLD
      - usage guidelines.....5449
    - OSPF.....4586
    - PIM.....5759
    - usage guidelines.....5494, 5500, 5512, 5547
    - RIP.....4767
    - usage guidelines.....4702
    - SNMP.....1688, 1991
    - usage guidelines.....1504
  - View Events page
    - field summary (filtering log messages).....
  - view statement
    - SNMP (associating with community).....1690
    - usage guidelines.....1503
    - SNMP (configuring MIB view).....1689
    - usage guidelines.....1506
  - views, MIB
    - SNMP.....1506
  - virtual link, through the backbone area.....4404
  - virtual links
    - overview.....4398
  - virtual memory, displaying.....1195
  - virtual switch
    - dual VLAN tags.....2310
  - virtual-address statement.....2703
    - usage guidelines.....2641
  - VLAN MIB.....1477
  - vlan statement
    - IGMP snooping.....5663, 5686
    - usage guidelines.....5435
    - port mirroring.....6275
  - vlan-id statement.....2306
  - vlan-id-list statement.....2307, 2375
  - vlan-rewrite statement.....2308
  - vlan-tags statement.....2310
  - VLANs
    - configuring.....2120, 2311
    - IGMP snooping.....5435
  - vpls statement.....2311
  - vni.....2528
  - VPLS
    - nonstop active routing.....2601
  - vpls protocol family
    - interface addresses.....2824
  - VPNs
    - BGP
      - preventing session flaps.....3919
  - VPNs, graceful restart.....2579
  - VRRP
    - advertisement interval.....2643
    - authentication.....2642
    - basic configuration.....2641
    - group
      - inheritance.....2650
      - passive ARP learning.....2649
      - tracking logical interface status.....2646
      - tracking route status.....2646
  - vrp
    - failover-delay.....2690
  - vrp-group statement.....2704
    - usage guidelines.....2641
  - VSTP
    - VLAN.....2361, 2362
  - vstp statement.....2364
  - vtep-source-interface statement.....2529
  - vxlان statement.....2529
- W**
- warning (system logging severity level).....3333
  - west-interface statement.....2270
  - wide-metrics-only statement.....4327
  - wildcard-source statement
    - PIM RPF selection.....5760

|                                                 |      |
|-------------------------------------------------|------|
| working directory                               |      |
| current, displaying.....                        | 434  |
| world-readable option to statistics statement   |      |
| MPLS.....                                       | 5094 |
| world-readable option to traceoptions statement |      |
| LDP.....                                        | 4879 |
| MPLS.....                                       | 5103 |
| RSVP.....                                       | 5339 |
| world-readable statement                        |      |
| archiving of all system log files.....          | 1692 |
| archiving of individual system log file.....    | 1694 |
| system logging                                  |      |
| usage guidelines.....                           | 1543 |
| write (tracing flag)                            |      |
| BMP.....                                        | 4084 |
| write-interval statement.....                   | 6143 |
| write-view statement.....                       | 1690 |
| <br><b>Y</b>                                    |      |
| yellow alarm condition.....                     | 18   |